Georgia Institute of Technology

· FACILITATOR GUIDE · INSTRUCTORS ONLY

Applied AI at TIYA Services: Facilitator Guide

The one document that covers all three days: minute-by-minute run sheets, talking points, demos, and debriefs. Day 1 runs on Microsoft 365 Copilot, Day 2 builds agents with the Copilot agent builder, and Day 3 builds triggered, scheduled, and multi-agent solutions in Copilot Studio.

DAY 1 Tue, Oct 20, 2026 DAY 2 Wed, Oct 21, 2026 DAY 3 [Date TBD] LEAD Alex Djalali

CO-INSTRUCTOR Alex Rhodeen

COURSE DESIGN

How the three days fit together

Each day climbs one part of the autonomy ladder: Day 1 prompts, Day 2 on-demand agents, Day 3 agents that start themselves. All three teach one habit, Mask → Ask → Check, and practice it on one set of fictional TIYA-style documents. Participants spend most of the day at the keyboard; instruction comes in short bursts between exercises.

1 · Mask

Strip or replace sensitive data before anything goes into Copilot.

- Names → roles
- Numbers → placeholders
- Pricing, CUI, PHI never go in

2 · Ask

Prompt with Microsoft's four parts for Copilot.

- **G**oal: what you want
- **C**ontext: why, for whom
- **S**ource: what to use
- **E**xpectations: format, length, tone

3 · Check

A person verifies before anything is used or sent.

- Facts, numbers, names
- Policy and tone
- "Would I sign this?"

Design principles

- **Keyboard time first.** Every block except the policy block has participants typing within 15 minutes of its start. Day 1 has about 5.5 hours hands-on; Day 2 has 3.5 hours of building, plus design and red-team exercises; Day 3 has about 5.25 hours of building.
- **One habit, repeated.** Mask → Ask → Check appears on every exercise card, so by 4:30 it is reflexive and the closing assessment tests it directly.
- **One sample world.** The same fictional work orders, SOPs, and building names run through both days. Participants spend their attention on the AI technique, not on learning new material each hour. Day 2 introduces its own files (the sandbox SOPs and the work order export), so a Day 2 cohort that took Day 1 months earlier doesn't need to remember the Day 1 samples. Day 3 reuses the Day 2 SOPs and export and adds a fictional inbox. Each day's materials stand alone: its own participant page, workbook, agenda, assessment kit, and zip.
- **Real work, masked.** The role lab and the Day 2 build invite participants to bring a real task, but only after masking it. Instructors check masking before anyone pastes.
- **AI drafts, people decide.** No exercise ends with AI output used as-is. Each one ends with a Check step and a share-out about what had to be fixed.

Instructor roles

- **Djalali (lead):** delivers instruction blocks, runs live demos, leads debriefs, keeps time.

- **Rhodeen (co-instructor):** runs the morning login check, circulates during labs, owns the remote channel, collects prompts for the "prompt wall", flags masking problems quietly to individuals. Swap lead on the leadership block (Day 1), the governance block (Day 2), and the operations block (Day 3) if you want a change of voice. On Day 3, Rhodeen also owns the trigger shutdown check.

Remote participants

- Join the Teams meeting with camera on during exercises if possible. Rhodeen watches chat throughout and voices remote questions at each debrief.
- Remote participants get that day's digital workbook and the materials folder link the day before.
- During labs, pair remote participants with each other in a Teams breakout room, or with an in-room partner on a laptop call.
- Share-outs: remote participants paste their prompt and result into chat; the lead reads one aloud each debrief.

Materials

- **Participant workbooks**, one for each day: a printed copy of that day's workbook for each participant plus the digital version (prompts are copy-ready). The Day 2 workbook opens with a one-page Day 1 review.
- **Slide outlines**, built into PowerPoint with GT branding.
- **Assessment kits**, one for each day: knowledge check, exit task or checklist, rubric, answer keys, certificate criteria, evaluation form.
- **Materials folder** (`materials/`): sample SOPs as Word files and a work order export as CSV. Upload these to a training SharePoint site before Day 2. Day 3 adds `materials/day3/`: the sample inbox and the export as an Excel table.
- **Day 3 outline** for the program owner: purpose, prerequisites, outcomes, the 8:00–4:00 plan, environment requirements, and open decisions.
- Printed: Mask/Ask/Check table tents, "Can I paste this?" quick-fire cards, red-team attack cards, certificates.

READINESS

Readiness checklist

Most failures in tool training happen at 8:05 AM, when a third of the room can't log in. Confirm each item with TIYA IT, in writing, before the event.

T-14 DAYS · TIYA IT

Licensing and tenant

- [ ] Every participant has a Microsoft 365 Copilot license, or at least Copilot Chat. Get the list of who has which.
- [ ] Which cloud: commercial, GCC, or GCC High. Feature availability differs; confirm which features below are live.
- [ ] Agent builder ("Create an agent" in Copilot Chat) is enabled for participants.
- [ ] Agent sharing to specific users is allowed.
- [ ] Code interpreter and file upload are enabled in Copilot Chat.
- [ ] Two test accounts for instructors in the same tenant.
- [ ] Day 2 cohort (took Day 1 in June): confirm each person still has a license and agent builder access, since licenses can lapse or move between June and October.

T-14 DAYS · TIYA LEADERSHIP

Policy and content

- [ ] A copy of TIYA's AI use policy and data handling rules (PII, PHI, CUI, pricing). Replace every placeholder rule in both workbooks with it.
- [ ] The name of the person who answers "Is this allowed?" after the course.
- [ ] TIYA's SOP template and doc-control process, if one exists.
- [ ] Roster with roles, so the role lab tracks can be balanced.

T-7 DAYS

Environment

- [ ] Create a SharePoint site "AI Training Sandbox"; upload `materials/`; give all participants read access.
- [ ] Instructors build both Day 2 agents end to end in the tenant and save screenshots.
- [ ] Run every workbook prompt once in the tenant; note any that behave differently.
- [ ] Send remote participants the Teams link, that day's digital workbook, and the sandbox link.

MORNING OF

Room

- [ ] Wi-Fi tested on a participant-type device; projector and room audio tested with Teams.
- [ ] Workbooks, table tents, cards, pens, power strips at each table.
- [ ] Prompt wall: a flip chart labeled "Prompts that worked".
- [ ] Backup plan ready (below).

T-14 DAYS · DAY 3 · TIYA IT

Copilot Studio

- [ ] A Copilot Studio license for every Day 3 participant and both instructors, and maker access to one training environment.
- [ ] Event triggers, agent flows, and child and connected agents are available in TIYA's cloud (commercial, GCC, or GCC High).
- [ ] Data policies allow the Office 365 Outlook, Microsoft Teams, Excel Online (Business), and SharePoint connectors in that environment.
- [ ] Enough capacity for the day; the Day 3 outline has the estimate.
- [ ] The Day 2 sandbox SOP site still exists, with read access for every Day 3 participant.
- [ ] Instructors build all four Day 3 agents end to end, send the full inbox, and record the tenant's labels, trigger timing, and which connection each trigger and tool uses.
- [ ] A named owner who switches off or deletes every training agent and trigger after the day.

BACKUP PLAN IF COPILOT IS DOWN OR BLOCKED

Switch to instructor-driven mode: the lead runs the exercise prompts on the instructor account on screen while tables write their own GCSE prompts on paper and predict the output. Every exercise card works this way. For Day 2, walk through the saved build screenshots and have pairs complete the design canvas and test plan on paper. For Day 3, run Builds 2–4 as instructor demos on screen; participants fill in the triage log, routing tests, and runbook from the demo output.

OPEN ITEMS

Decisions for the program owner

1. **Syllabus wording on tools.** The Day 1 outline says "tools like ChatGPT". The course is built on Microsoft 365 Copilot. Suggest "tools like Microsoft 365 Copilot" so participants don't expect ChatGPT access.
2. **Day 2 outcomes vs. agent builder.** Outcome 2 promises *autonomous* agents and the outline lists *event triggers*. The agent builder makes on-demand agents; triggers and autonomous agents need Copilot Studio or Power Automate. This guide teaches triggers as a concept with an instructor demo and a design exercise. Either soften the outcome ("design … and deploy agents, and design event-triggered automations") or provide a Copilot Studio environment for the demo.
3. **Day 2 closing.** The Day 2 outline ends at 4:30 with no closing or assessment. This guide uses the last 10 minutes of the governance block for the exit checklist and certificates. If the day runs to 5:00, move them to a 4:30–5:00 closing block.
4. **Day 2 headers.** The Day 2 outline reuses Day 1's session titles. Suggested: Morning "Agent Foundations & First Builds"; Afternoon "Deploy, Trigger & Govern".
5. **Certificate criteria.** Each day's assessment kit proposes full attendance (up to 30 minutes excused) plus that day's evidence of work: the Day 1 exit task, the Day 2 agent and governance card, and the Day 3 intake agent, runbook, and trigger shutdown check. Confirm, or set a stricter bar.
6. **Day 3 date and access.** Day 3 has no date yet, and it can't run as designed without Copilot Studio licenses (readiness card above). The Day 3 outline lists the remaining Day 3 decisions: cohort size and who approves triggered agents at TIYA.

DAY 1 · TUESDAY, OCTOBER 20, 2026 · 8:00 AM – 5:00 PM

Generative AI for Daily Work

Outcomes: **LO1** apply generative AI to daily tasks · **LO2** use a prompting framework · **LO3** practice responsible use · **LO4** improve operational materials · **LO5** support leadership work.

8:00 – 8:30 30 min

AI Overview & Context

LO1

GOAL Everyone is logged in, knows what generative AI is good and bad at, and has the day's habit in mind.

TIME	WHAT HAPPENS	WHO
8:00	Welcome, introductions, logistics: breaks, restrooms, Wi-Fi, phones, remote etiquette. While the lead talks, the co-instructor walks the room for the login check.	LEAD · CO
8:08	Show of hands: used Copilot or ChatGPT at work? At home? Never? Then ask two people: "Which task would you most like an hour back from each week?" Write answers on the flip chart; revisit them at 4:30.	LEAD
8:12	What generative AI is and is not. Live demo (below).	LEAD
8:20	Where AI fits in TIYA work: the "draft, not decide" map (slide).	LEAD
8:26	Introduce Mask → Ask → Check and the day's arc. Point to the parking lot flip chart.	LEAD

LOGIN CHECK Each participant opens Copilot Chat the way TIYA staff normally reach it and types "hello". Anyone who can't is listed by name and seat. Escalate to TIYA IT before 8:30; seat unresolved people beside a working neighbor.

SAY "Generative AI predicts the most likely next words based on patterns in huge amounts of text. That makes it very good at language work: drafting, summarizing, reorganizing, rewording. It doesn't know TIYA's facts unless you give them to it, and when it doesn't know, it can sound just as confident as when it does."

DEMO Type: `What is the phone number for TIYA Services' emergency work order desk at Fort Benning?` It will either say it doesn't know or produce something plausible. Either way, ask the room: "How would you know if that's right?" That question is the whole Check step.

MAP AI is good for: **drafting, summarizing, organizing, analyzing, rewording**. AI is not for: **deciding, approving, signing, or holding sensitive data**. Give one example per role: facilities (turn tech notes into a status update), contracting (plain-language summary of a public clause), business management (meeting notes into action items), leads (first draft of a team announcement).

REMOTE Co-instructor posts the login instructions in Teams chat at 7:55 and checks in with each remote participant by name.

`8:30 – 9:45` `75 min · hands-on`

Generative AI Fundamentals

LO1 LO2

GOAL Participants can write a GCSE prompt, iterate with follow-ups, and use Copilot to draft, summarize, brainstorm, and reword.

TIME	WHAT HAPPENS	WORKBOOK
8:30	Tour of Copilot Chat on screen: new chat, typing, editing a prompt, regenerating, copying output, uploading a file, chat history. Point out the work/web toggle if the tenant shows one.	p. "Copilot basics"
8:35	Teach Ask = Goal, Context, Source, Expectations. Show the weak vs. strong example side by side.	GCSE card
8:45	**E1 · Two prompts, one task.** Water shutoff notice: run the weak prompt, then the GCSE prompt, compare.	E1
9:00	**E2 · Summarize the work order notes.** Three formats from one source: status summary, timeline, table.	E2 · Sample A
9:15	**E3 · Follow-ups are prompts too.** Push the E2 output with four follow-up moves.	E3
9:25	**E4 or E5**, participant's choice: brainstorm and organize ideas, or fix the tone of a terse email.	E4 · E5

TIME WHAT HAPPENS WORKBOOK

SAY "Copilot is like a capable new hire on day one. Smart, fast, eager, and it knows nothing about Building 214 or who the COR is. The more you'd have to explain to a new hire, the more you need to put in the prompt."

E1 SETUP Everyone runs the weak prompt first, then the GCSE version. Ask: "Which one could you send with the fewest edits? What did the second prompt add that made the difference?" Expected answer: audience, specific details, length, and format.

WATCH FOR

- People pasting Sample A into a new chat without saying what they want. Prompt them: "What's the Goal?"
- People accepting the first output. Push them to the E3 follow-ups.
- Fast finishers: give them E4 and E5 both, then ask them to help a neighbor.

DEBRIEF

1. Which part of GCSE changed the output most for you?
2. Where did Copilot add something that wasn't in the notes? (Expect at least one invented detail in E2. Use it to set up 10:00.)
3. What would you have to check before sending the E2 summary to the COR?

10:00 – 10:45 45 min

Responsible Use (TIYA Policy)

LO3

GOAL Participants can classify data as red, amber, or green, mask a document correctly, and apply a Check routine before using output.

REPLACE BEFORE DELIVERY

The data rules in this block and in both workbooks are placeholders based on common practice for federal-installation contractors. Replace them with TIYA's written policy. Where the two differ, TIYA's policy wins.

TIME	WHAT HAPPENS	WORKBOOK
10:00	Why this matters: a short story (below). Ask: "What went wrong, and when?"	
10:05	The four categories with TIYA examples: PII, PHI, CUI, pricing. Then the red/amber/green chart.	Data traffic light
10:15	**Masking drill.** Mask Sample B alone (6 min), compare with a neighbor (3 min), reveal the key (3 min).	Sample B
10:27	**"Can I paste this?"** Eight quick-fire cards; tables vote green, amber, or red.	Quick-fire
10:35	Check: human-in-the-loop. Demo (below), then the Check list and the "Would I sign this?" rule.	Check card

STORY

"A contracts assistant at a company like ours gets a vendor's quote for a chiller replacement. They want a quick comparison against last year's price, so they paste both quotes into a free AI tool on their phone. The comparison is great. But two companies' pricing has now left the building, into a service the company doesn't control, under terms nobody read. Nothing bad has to happen next for this to be a problem. It already is one." *(Fictional; tell it as a hypothetical.)*

CATEGORIES

- **PII**: anything that identifies a person: names with other details, phone numbers, home addresses, employee IDs, DoD ID numbers, SSNs.
- **PHI**: health information tied to a person: injuries, medical restrictions, workers' comp details.
- **CUI**: government information marked or designated as Controlled Unclassified Information, including some facility, security, and infrastructure details. If you don't know whether something is CUI, treat it as CUI and ask.
- **Pricing**: TIYA's rates, cost data, bids, proposals, and any vendor's quote.

SAY

"Copilot signed in with your TIYA account has enterprise data protection. Microsoft doesn't use your prompts to train its models. That's real and it matters. It does not change the rules for CUI, PHI, or pricing. Those rules come from TIYA's policy and its contracts, not from the tool."

DEMO Ask Copilot: **`Which FAR clause covers limitations on subcontracting for small business set-asides? Quote the key paragraph.`** Then open acquisition.gov on screen and compare the clause number and wording. Whether Copilot gets it right or not, the point lands: you check against the source, every time it matters.

MASKING KEY Sample B contains 11 items to mask or remove: 3 names, 1 phone number, 1 email address, 1 injury detail (PHI, remove entirely), 1 contract number, 2 dollar amounts (pricing, remove), 1 vendor name, 1 gate access code (remove entirely). Most people miss the access code and the vendor name. The answer key is in the assessment kit.

QUICK-FIRE KEY

1. Generic HVAC troubleshooting question: green.
2. Meeting notes listing attendees by name: amber; replace names with roles.
3. Subcontractor's bid spreadsheet: red.
4. Published FAR clause text: green.
5. Email about an employee's medical restriction: red.
6. A SOW excerpt marked CUI: red.
7. Your own draft newsletter article: green.
8. Overtime hours by named employee: amber; use Employee A, B, C, or aggregate.

`10:45 – 12:00` `75 min · hands-on`

Role-Based Application Lab

LO1 LO2 LO3

GOAL Each participant completes one track task and one real task of their own, masked, and leaves with a prompt they will reuse.

TIME	WHAT HAPPENS	WORKBOOK
10:45	Pick a track: Facilities Maintenance, Contracting, Business Management & Admin, or Supervisors & Leads. Sit with your track; pair up.	Track cards
10:50	**Task A** (required): the track card's task, using the sample input. Run Mask → Ask → Check.	Track cards
11:10	**Task B**: a real task from your own job. Write the masked version in the workbook first; an instructor initials it before you paste. Or do the card's Task B.	Task B page

TIME	WHAT HAPPENS	WORKBOOK
11:35	Share-outs: one pair per track, 4 minutes each. Show the prompt, the output, and what you had to fix.	
11:55	Everyone copies one prompt worth keeping into "My prompt library".	Prompt library

WATCH FOR

- **Masking on Task B.** This is the first time participants use their own material. Look at the masked text before they paste. Common misses: building numbers combined with names, email signatures, and quoted dollar figures.
- Someone with no suitable real task: they take the card's Task B.
- Tracks that are lopsided: merge Supervisors & Leads into Business Management if it has fewer than four people.

DEBRIEF

1. What did you have to fix in the output?
2. What did you have to take out before pasting?
3. How long would this task have taken without Copilot? With it, counting the checking?

1:00 – 1:15 15 min

Afternoon Framing & Recap

LO1 LO2 LO3

GOAL Re-anchor the morning, then shift from single tasks to whole workflows and leading others.

TIME	WHAT HAPPENS
1:00	Pairs, 4 minutes: one thing that worked this morning, one thing that surprised you. Two pairs report.
1:05	Four-question show-of-hands quiz on slides: name the four GCSE parts; red, amber, or green for two examples; what Check means.
1:10	Transition: "This morning was single tasks. This afternoon is the documents your operation runs on, SOPs, and the people you lead."

1:15 – 2:45 90 min · hands-on

SOP Lifecycle Lab

LO2 LO3 LO4

GOAL Take one outdated SOP through four stages: diagnose it, rewrite it to a standard template, turn it into a flow, and adapt it into training and communication materials.

TIME	WHAT HAPPENS	WORKBOOK
1:15	The SOP lifecycle: Generate → Update → Improve → Adapt. Rule for the lab: AI drafts, the subject-matter expert owns accuracy, and the doc-control process approves.	SOP lab intro
1:25	**Step 1 · Diagnose.** Ask Copilot to critique Sample C (FM-SOP-07, 2019): gaps, ambiguity, safety risks, outdated items.	Sample C · Step 1
1:40	**Step 2 · Rewrite** to the standard template: Purpose, Scope, Roles, Definitions, Procedure, Records, References, Revision history.	Step 2
2:00	**Step 3 · Visualize.** Turn the procedure into a step table (Step | Who | Action | Decision | Next) and a text flowchart. Optional: open the table in Word and convert it to SmartArt.	Step 3
2:15	**Step 4 · Adapt.** Choose two: a one-page quick-reference card, a 5-question training check, an announcement email to technicians, or a 3-minute toolbox talk.	Step 4
2:30	Discussion (below).	

SAY

"Copilot is very good at the shape of an SOP: structure, plain language, consistency, catching steps that don't say who does them. It cannot know whether the after-hours number is still right, whether the on-call rotation changed, or what actually happens on a Saturday night in Building 214. That's you."

WATCH FOR

- Rewritten SOPs that invent specifics: response times, phone numbers, job titles. Ask participants to mark every invented detail with `[VERIFY]`. Step 2's prompt asks Copilot to do this too.
- Step 1 should surface at least: no defined roles, an outdated phone tree, a missing step for notifying the COR, passive voice ("the system should be shut down"), no record-keeping step, and no definition of "emergency".

DISCUSS

1. What did Copilot miss or get wrong that only someone at TIYA would know?
2. Which step saved the most time? Which needed the most checking?
3. At TIYA, who would have to review an AI-assisted SOP change before it's official? Does the doc-control process need a note about AI-drafted content?

3:00 – 4:30 90 min · hands-on

AI for Leadership Support & Responsible AI Adoption

LO3 LO5

GOAL Use Copilot as a thinking partner for communication, hard conversations, and decisions, and leave with draft AI norms for your team.

TIME	WHAT HAPPENS	WORKBOOK
3:00	Framing: AI as a thinking partner, not a decision-maker. Privacy for people topics: never name the employee; describe the situation.	Leadership intro
3:10	**L1 · Announce a change.** Draft the message about a 90-day shift pilot and anticipate the team's questions.	L1
3:30	**L2 · Prepare a hard conversation.** Plan it, then have Copilot play the employee for a practice round and give feedback.	L2
3:50	**L3 · Decide between options.** Options matrix and pre-mortem for reducing a work order backlog.	L3
4:05	**L4 · Lead adoption.** Respond to a team member's risky AI use and draft five team AI norms.	L4
4:25	Debrief.	

SAY "Leadership is judgment and relationships, and AI has neither. What it gives you is a fast, patient sounding board: a first draft, a list of questions your team will ask, a rehearsal partner who doesn't get tired. You stay the one who decides and the one who shows up."

L2 TIP The role-play is the most memorable exercise of the day. Model it on screen for 2 minutes first: give Copilot the scenario and say "You are the employee. Respond realistically, including some defensiveness. Wait for my reply each time." After three exchanges, type "Pause the role-play. Give me feedback on how I handled it."

WATCH FOR Participants describing a real colleague in identifying detail during L2. Steer them to the workbook scenario or to a fully anonymized version: no names, no unique job details, no health information.

DEBRIEF

1. Where was Copilot useful as a thinking partner, and where did it fall flat?
2. Read your best team norm aloud. (Collect five for the prompt wall.)
3. What's one conversation you'll have with your team about AI this month?

4:30 – 5:00 30 min

Closing & Assessment

LO1–LO5

GOAL Confirm learning, commit to next steps, and issue certificates.

TIME	WHAT HAPPENS	SOURCE
4:30	Knowledge check, 10 questions, self-scored. Formative: it isn't pass/fail.	Assessment kit
4:40	Practical exit task: Mask → Ask → Check on a short scenario. Hand in the page.	Assessment kit
4:52	Action plan: three tasks I'll use Copilot for next week, and one person I'll show. Revisit the 8:08 flip chart: did we cover the hour people wanted back?	Workbook
4:56	Evaluation form; certificates at the door once exit tasks are in.	Assessment kit

DO The co-instructor scans exit tasks against the rubric while the action plan runs. Anyone who scored 0 on masking gets a two-minute conversation before getting a certificate, not a refusal.

DAY 2 · WEDNESDAY, OCTOBER 21, 2026 · 8:30 AM – 4:30 PM

Agentic AI with Microsoft Copilot

Outcomes: **LO1** agent fundamentals and architectures · **LO2** design, configure, and deploy no-code agents · **LO3** integrate reasoning, planning, tools, and memory · **LO4** apply security, privacy, safety, and trust. Day 2 is for participants who took Day 1 in June 2026, four months earlier. Don't assume they remember it: the day opens with a 10-minute Day 1 review, and nothing later depends on a Day 1 exercise or sample.

8:30 – 10:15 105 min

Agentic AI Foundations

LO1 LO4

GOAL Participants can explain what makes an agent different from a chat, name its parts, place the Microsoft tools on a map, and design one agent on paper.

TIME	WHAT HAPPENS	WORKBOOK
8:30	Welcome and logistics. Login check, plus: everyone opens Copilot Chat and finds the option to create an agent. Anyone who can't is escalated now.	
8:40	**Day 1 in review** (four slides). Everyone took Day 1 in June: ask who has used Copilot since and for what, then recap draft-not-decide, Mask → Ask → Check, and the data traffic light, ending with a three-card traffic-light vote.	Day 1 review (front of Day 2 workbook)
8:50	**From chat to agent.** The autonomy ladder (below).	Autonomy ladder
9:05	**Anatomy of an agent.** Seven parts, and which ones agent builder supports.	Anatomy
9:20	**The no-code ecosystem** and four common architectures.	Ecosystem
9:30	**Exercise · Agent or not?** Tables sort ten TIYA tasks into prompt, agent, automation, or not-AI.	Agent or not?
9:45	**Exercise · Agent design canvas.** Pairs design an agent for their own work. Three pairs share.	Design canvas
10:10	Wrap: "You'll build two agents this morning, then this one this afternoon."	

LADDER

1. **Prompt:** you ask, it answers, it forgets.
2. **Saved prompt:** you reuse good instructions.
3. **Agent:** standing instructions, its own knowledge, and tools, for one defined job. You start each conversation. *Agent builder lives here.*
4. **Triggered agent:** starts itself when something happens, like a new email or a scheduled time.
5. **Autonomous agent:** plans and takes multi-step actions in other systems, with people approving at checkpoints. *Copilot Studio territory.*

"Each rung up gives the agent more independence, so each rung needs more design, more testing, and more governance."

ANATOMY

- **Purpose:** one job, one audience. Agent builder: name and description.
- **Instructions:** a standing prompt. Agent builder: yes.
- **Knowledge:** what it reads: SharePoint, files, web. Agent builder: yes.
- **Tools:** what it can do beyond text. Agent builder: some built-in capabilities, such as code interpreter and image generation. Connections to other systems: Copilot Studio.
- **Memory:** what it remembers. Agent builder: the current conversation. Long-term memory varies by tenant.
- **Triggers:** what starts it. Agent builder: a person. Events and schedules: Copilot Studio or Power Automate.
- **Guardrails:** what it must not do. Agent builder: instructions, plus the permissions of its knowledge.

ARCHITECTURES

- **Knowledge Q&A:** answers from a defined set of documents (the SOP Navigator).
- **Task assistant:** walks a person through a structured task with a fixed output format.
- **Multi-tool:** combines knowledge with tools, such as analyzing an uploaded spreadsheet (the Work Order Analyst).
- **Orchestrated or multi-agent:** one agent hands work to others, often triggered by events. Mentioned for awareness; not built today.

SORT KEY

The answer key for "Agent or not?" is in the assessment kit. Expect debate on items 4 and 9; that debate is the point.

`10:30 – 12:15  105 min · hands-on`

Building Real Agents

LO2 LO3

GOAL Each participant builds and tests a knowledge agent and a multi-tool agent, writing instructions from a template and tests before building.

TIME	WHAT HAPPENS	WORKBOOK
10:30	**Prompt engineering for agents.** Instructions template: role and purpose, audience, how to use knowledge, steps, output format, boundaries and escalation, tone. Tests first: write five test questions before building.	Instruction template · Test plan
10:45	**Build 1 · SOP Navigator**, guided step by step on screen. Knowledge: the sandbox SOP folder.	Build 1
11:05	Run the five tests. Fix the instructions for each failure. Rerun.	Test plan
11:25	**Build 2 · Work Order Analyst**, a multi-tool agent: priority definitions as knowledge, code interpreter on, CSV uploaded in chat.	Build 2
11:55	Test with the sample CSV; ask for a backlog chart and a repeat-failure analysis; check the numbers by hand on three rows.	Build 2
12:05	Debrief.	

SAY

"An agent's instructions are a prompt that never ends. Every weakness in them shows up in every conversation, for everyone you share it with. So we write the tests before we build. A test is just a question with an answer you already know is right."

BUILD 1

Use the agent builder's configure view, not the conversational describe view, so everyone's screen matches. Walk through: name, description, instructions (paste from the workbook), knowledge (add the sandbox SOP folder), starter prompts. Test in the preview pane. Before the event, note the exact labels and steps in TIYA's tenant on this card; the agent builder's interface changes often.

BUILD 2

Turn on code interpreter. Add `FM-SOP-03 Work Order Priority Definitions` as knowledge. In the test chat, upload `work-orders-aug-sep-2026.csv` . The data has two planted patterns for participants to find: repeat HVAC failures in Building B-214, and a backlog of Routine work orders older than 30 days.

WATCH FOR

- Agents that answer questions outside their SOPs from general knowledge. Fix: the instruction "If the answer is not in the SOPs, say so and name who to ask."
- Charts and totals people don't check. Have them count three rows in the CSV by hand.
- Knowledge that hasn't finished indexing. A newly added SharePoint source can take a few minutes; build it before the break if possible.

DEBRIEF

1. Which test failed first, and what instruction fixed it?
2. Where did the Work Order Analyst do real reasoning, and where did it just restate the data?
3. Would you trust its chart in a report to the COR? What would you check first?

`1:15 – 3:00` `105 min · hands-on`

Building Real Agents, Continued

LO2 LO3 LO4

GOAL Participants design an event trigger, understand memory limits, publish safely, and build and peer-test the agent they designed this morning.

TIME	WHAT HAPPENS	WORKBOOK
1:15	**Event triggers.** Trigger → condition → action. Instructor demo or walkthrough of a triggered flow. Participants storyboard a trigger for their canvas agent.	Trigger storyboard
1:35	**Memory.** Conversation context vs. knowledge vs. long-term memory, and how to design an agent that doesn't need to remember.	Anatomy
1:45	**Publishing.** Sharing options, knowledge permissions, naming, the description as the agent's label. Demo: share the SOP Navigator with a partner.	Publishing checklist
2:00	**Build your own** from this morning's canvas. Write the five tests first.	Canvas · Test plan
2:30	**Swap test.** Share with your partner; they run your tests plus two of their own and log the results. Fix and reshare.	Test plan
2:45	Showcase: four or five volunteers, 2 minutes each.	

TRIGGERS "Everything you built this morning waits for someone to ask. A triggered agent starts itself: a new email arrives in the work order mailbox, a file lands in a SharePoint folder, it's 7 AM Monday. In Microsoft's stack, that's Copilot Studio or Power Automate. The design thinking is the same, and that's what we'll practice."

DEMO If a Copilot Studio or Power Automate environment is available, show a flow that starts on a new email to a shared mailbox, summarizes it, and posts to a Teams channel for a person to approve. Otherwise, walk through the screenshots captured during prep. Stress the approval step: the person stays in the loop.

MEMORY "Assume the agent remembers nothing between conversations. If it needs context, design it to ask for it, or put that context in a document it reads as knowledge."

PUBLISHING

- The agent can only show a user what that user can already open. Sharing an agent doesn't share its knowledge, so check that your audience can open the sources.
- That cuts both ways: an agent built on a folder with loose permissions will surface whatever is in it to everyone who can open it.
- Share with named people first. Organization-wide availability goes through TIYA's review (governance block).
- Naming: `TIYA · [Team] · [Job]`, for example `TIYA · FM · SOP Navigator`. The description says what it does, what it doesn't, and who owns it.

`3:15 – 4:30` `75 min`

Governance & Security

LO4

GOAL Participants red-team an agent, rate its risk tier, and write a governance card for it.

TIME	WHAT HAPPENS	WORKBOOK
3:15	**PII handling in agents.** Oversharing through knowledge, permission inheritance, why instructions aren't a security control, and what never goes into knowledge.	Governance intro
3:30	**Red-team swap.** Five attack cards against your partner's agent. Log what broke; fix one thing.	Red-team log

TIME	WHAT HAPPENS	WORKBOOK
3:45	**Risk assessment.** Rate your agent's tier; list the controls that tier requires.	Risk tier worksheet
4:00	**Model governance.** The agent lifecycle; owner, change log, re-test, retire. Write the governance card.	Governance card
4:20	Exit checklist, evaluation, certificates	Assessment

SAY

"An instruction like 'never reveal personal information' is a request, not a lock. A clever question can get around it. The real lock is what the agent can read. If it can't see it, it can't leak it. So the most important governance decision is what you put in the knowledge."

NEVER AS KNOWLEDGE

HR and personnel files, medical or injury records, pricing and cost data, bid documents, CUI, security details such as access codes or camera locations, and any folder whose permissions you haven't checked.

LIFECYCLE

Propose → Build → Test → Review → Share → Monitor → Re-test → Retire. Re-test when the knowledge changes, when the instructions change, and periodically anyway: Microsoft updates the underlying models, and behavior can shift without anyone touching the agent.

DEBRIEF

1. Which attack worked on your agent? What fixed it, and is the fix a lock or a request?
2. What tier is your agent? Did your partner agree?
3. Who at TIYA should own the list of shared agents?

DAY 3 · [DATE TBD] · 8:00 AM – 4:00 PM

Advanced Agents with Copilot Studio

Outcomes: **LO1** choose and design on-demand, event, or scheduled agents · **LO2** build triggered and scheduled agents with a human check · **LO3** build an orchestrated multi-agent solution, and know when not to · **LO4** operate autonomous agents safely. Day 3 is for Day 2 completers who will own agents. The program owner's outline is a separate document. Everything runs in Copilot Studio, so the Day 3 readiness card above is not optional.

8:00 – 8:30 30 min

Welcome & Day 2 in Review

GOAL Everyone is logged in to Copilot Studio, and the room shares four new ideas the rest of the day depends on.

TIME	WHAT HAPPENS	WORKBOOK
8:00	Welcome and logistics. Say "turn your triggers off before you leave" now; you'll check it at 3:45.	
8:05	**Login check.** Open Copilot Studio in the training environment, create a blank agent, delete it. Open Outlook and Teams in other tabs. Rhodeen escalates anyone who can't.	
8:10	**Day 2 in review.** The autonomy ladder, then four rules that still hold. Ask who built an agent after Day 2 and what happened when someone else used it.	Day 2 in one page
8:20	**What's new today:** no one is watching, untrusted input, real actions, real cost.	New today

SAY "On Day 2, every agent waited for someone to ask. Today's agents start themselves, at 2 AM on a Sunday if that's when the email arrives. Everything you learned still applies, but a person isn't sitting there reading each answer as it appears. So we design the human check into the agent instead."

8:30 – 9:45 75 min · hands-on

From Agent Builder to Copilot Studio

LO2

GOAL Participants can find their way around Copilot Studio, rebuild a known agent there, see what it used in the activity view, and publish it for other agents.

TIME	WHAT HAPPENS	WORKBOOK
8:30	**Two builders, one idea.** Agent builder vs. Copilot Studio vs. Power Automate.	
8:35	**Tour** (10 min): Overview, Knowledge, Tools, Agents, Topics, triggers, test pane, activity view, Publish.	Page map
8:45	**Tools: how an agent acts.** A tool's name and description tell the agent when to call it; every tool is something it can do without asking.	
8:50	**Build 1 · SOP Navigator (Studio)**, demo then hands-on: same instructions as Day 2, the sandbox library as knowledge, general knowledge and web search off. Rerun the Day 2 tests. Publish and allow other agents to connect.	Build 1 · Retest log
9:35	Debrief.	

DEMO

Use the configure view so screens match. After the first test question, open the activity view and show which knowledge source it used. Make that the habit for the whole day. Before the event, record TIYA's tenant labels for: the general-knowledge and web-search settings, the setting that lets other agents connect, and Publish. Copilot Studio's interface changes often.

SETTINGS AS LOCKS

"On Day 2, 'don't answer from general knowledge' was a request in the instructions. Here you can turn general knowledge off. That's a lock. Whenever a setting can enforce something, use the setting and keep the instruction as backup."

WATCH FOR

- Knowledge still indexing: answers say "I can't find that" for a few minutes after adding the library. Wait before changing instructions.
- Participants who skip Publish. Build 4 can't connect to an unpublished agent; check before the break.

DEBRIEF

1. Did any Day 2 test give a different answer here? Why?
2. What did the activity view show that the chat didn't?

10:00 – 11:45 105 min · hands-on

Event-Triggered Agents

LO1 LO2 LO4

GOAL Each participant builds an agent that starts on a new email, triages it against FM-SOP-03, and posts a card for a person, with no way to reply on its own.

TIME	WHAT HAPPENS	WORKBOOK
10:00	**Anatomy of a trigger:** event, condition, payload, action, human check. Agent instructions vs. payload instructions.	Trigger canvas
10:10	**The email is not your instructions.** The day's most important idea.	
10:15	**Build 2 · Work Order Intake**, demo then hands-on. Knowledge: FM-SOP-03 and FM-SOP-07. Tool: Teams post to yourself with fixed recipient. Trigger: new email with `[WO-TRAIN]` in the subject.	Build 2
10:55	**Send the inbox.** Each participant sends themselves the eight sample emails and fills in the triage log.	The inbox · Triage log
11:25	**The three hard emails:** #4 override, #5 health details, #6 duplicate. Fix one thing and resend.	Triage log
11:35	Debrief.	

SAY

"When a trigger fires, the agent reads the email. That means whoever wrote the email wrote part of the agent's prompt. Most senders are colleagues, but the agent can't tell a colleague from someone who wants to trick it. So we treat every email as data to classify, never as instructions, and we make sure that even a fooled agent can't do anything harmful. Look at the tools: it can post one card to you. That's all."

BUILD 2

Each participant's trigger watches *their own* inbox, filtered on `[WO-TRAIN]`, so 30 people don't trigger each other's agents. The Teams tool posts to a chat with the Flow bot, with the recipient fixed to the participant. During prep: confirm in TIYA's tenant whether a trigger fires before the agent is published or only after, how long the first run takes, and which account's connection the trigger and tools use. Write all three on this card.

WATCH FOR

- Triggers that never fire: usually the subject filter (brackets, spacing) or the wrong folder. Check the activity page before the instructions.
- Cards that repeat the injured person's details from email 5. Fix: the boundary instruction, and check that it's above the format block.
- Anyone who adds a send-email tool "to make it useful". Ask what email 4 would have done with it.
- Participants pasting real work emails. Today's inbox is fictional for a reason.

DEBRIEF

1. Which email did your agent get wrong first? Was the fix a request or a lock?
2. What would email 4 have done if the agent had a send-email tool?
3. Email 6: how would a real version catch duplicates? (A person, or a tool that searches open work orders, not memory.)

`12:30 – 1:30` `60 min · hands-on`

Scheduled Agents

LO1 LO2

GOAL Participants build an agent that runs on a schedule, check its numbers against the source, and turn it off.

TIME	WHAT HAPPENS	WORKBOOK
12:30	**Two ways to run on a schedule.** If TIYA's tenant has scheduled prompts in Copilot Chat, show one for 2 minutes. Then the recurrence trigger.	Comparison table
12:40	**Let the flow count; let the agent write.** Recall the B-410 trap from Day 2.	Flag box
12:45	**Build 3 · Monday Backlog Briefing.** Excel table in OneDrive, Excel tool, Teams tool, recurrence every 5 minutes for testing.	Build 3
1:05	Check the first briefing against the key. Fix, wait for one more run. Set Monday 7:00, then turn the trigger off.	Check the briefing
1:20	Debrief.	

SAY

"A schedule is just a trigger whose event is the clock. The question is the same as this morning's: who reads this, and what do they do with it? A briefing that nobody reads is cost with no benefit. A briefing that someone forwards to the COR without checking is worse."

BUILD 3 The Excel connector needs a table, not a CSV, so participants use `work-orders-aug-sep-2026.xlsx` (table `WorkOrders`), saved to their own OneDrive. Fix location, file, and table as set values. During prep, run it once and note how long the Excel tool takes on 80 rows and whether the agent gets every row.

WATCH FOR

- Five-minute schedules left running. At 1:20, walk the room and check that every trigger is off.
- Briefings that report fewer than 80 rows. The instructions say to stop and say so; check that it did.
- Counts that are close but wrong. That's the lesson, not a failure: credit the participant who caught it.

DEBRIEF

1. Which number was wrong, if any?
2. Who reads this every Monday, and what do they do with it?
3. If it silently stopped, who would notice, and when? (Sets up 3:00.)

`1:30 – 2:45` `75 min · hands-on`

Multi-Agent Orchestration

LO3

GOAL Participants build an orchestrator with one connected and two child agents, test its routing, and can say when one agent would have been better.

TIME	WHAT HAPPENS	WORKBOOK
1:30	**Why more than one agent**: focus, reuse, ownership. **Child or connected? The cost of a team**: latency, more to govern, chained errors.	Child vs. connected
1:45	**Build 4 · Facilities Desk**, demo then hands-on. Connect the Build 1 SOP Navigator; add Priority Advisor and Notice Drafter as child agents; reference them in the orchestrator's instructions with `/`.	Build 4
2:05	**Routing tests.** Six questions; record which agents ran from the activity view. Fix descriptions first.	Routing tests
2:30	**One agent or many?** Four scenarios, tables vote.	One agent or many?
2:40	Debrief.	

SAY "The orchestrator picks an agent mostly by reading its description. If two descriptions sound alike, it guesses. So when routing goes wrong, fix the description of the agent that should have been picked before you touch any instructions. And remember Microsoft's own advice: start with one agent. Split when it starts choosing badly, or when different teams own the parts."

DEMO Run routing test 4 (leak + after-hours notification) and open the activity view: two agents run and the orchestrator combines them. That's the picture to leave on screen while people build.

WATCH FOR

- "Can't find the agent to connect": the Build 1 agent wasn't published or doesn't allow connections.
- Test 2 routed to the SOP Navigator, which also has FM-SOP-03. Good teaching moment about overlapping descriptions.
- Facilities Desk answering itself: general knowledge still on.

DEBRIEF

1. Which test routed wrong, and what fixed it?
2. Scenario 3, child agents or a single agent? (Expect debate.)
3. Where would you not use a team of agents at TIYA?

`3:00 – 3:45` `45 min`

Operating Autonomous Agents

LO4

GOAL Participants attack their own triggered agent, rate its tier, and write the runbook someone else could operate it from.

TIME	WHAT HAPPENS	WORKBOOK
3:00	**Five new risks:** injection, credentials, silence, cost, chains. Each with its lock.	Five new risks
3:08	**Red team.** Turn the Build 2 trigger back on, send attack emails A1–A3, then one of your own. Log what the card said and whether the agent obeyed.	Red-team log
3:23	**The off switch and risk tiers.** A triggered agent is never tier 1.	Risk tiers
3:28	**Runbook** for the intake agent.	Runbook

TIME	WHAT HAPPENS	WORKBOOK

CREDENTIALS "When a triggered agent posts to Teams or reads a file, it does that with somebody's connection. Often it's the maker's. If you leave TIYA, or your access changes, the agent changes with you. And whatever your connection can reach, a fooled agent might reach too. Give it the least access the job needs, and write down whose it is." Confirm during prep how triggers and tools authenticate in TIYA's tenant, and correct this talking point if it differs.

THE OFF SWITCH Three levels, smallest first: turn off the trigger; disable a child or connected agent on the Agents page; unpublish the agent. "Know which one you'd use, and where it is, before you need it. That's the most important line in the runbook."

DEBRIEF

1. Did any attack email work? Was your fix a request or a lock?
2. What tier is your intake agent? What would make it tier 3?
3. Who at TIYA should approve a triggered agent before it runs for real?

3:45 – 4:00 15 min

Closing & Assessment

GOAL Knowledge check, every trigger off, evaluation, certificates.

TIME	WHAT HAPPENS	WORKBOOK
3:45	Knowledge check, eight questions, self-scored.	Day 3 assessment kit
3:52	**Shutdown check.** Each participant shows an instructor that their triggers are off and no agent has a send-email tool; instructor initials the exit checklist.	Turn it all off
3:57	Next step (one trigger to propose, its runbook, its approver), evaluation, certificates.	

WATCH FOR Remote participants: Rhodeen checks their triggers by screen share before they leave the call. Don't issue a certificate until the shutdown check is initialed.

FOLLOW-UP

After the course

- **Within 2 days:** send participants the prompt wall (typed) and the materials links; Day 1 participants also get the team norms collected in L4, and Day 2 participants the list of agents they built, Day 3 participants their runbooks.
- **Within 1 week:** send TIYA leadership a summary: attendance, knowledge check averages by question, exit-task patterns (especially masking misses), evaluation scores, and the list of agents participants built with their risk tiers.
- **Within 30 days:** a 15-question pulse survey asking how often participants used Copilot, for what, and what got in the way. This is the evidence for a second cohort.
- Delete instructor test agents and the sandbox site, or hand them to TIYA IT to own. After Day 3, confirm in Copilot Studio that every training trigger is off and every training agent is unpublished or deleted.