OCAI GO: developed content and upcoming sessions

Expanded content inventory, checked 2026-10-02. This describes available material and production status. It is not a record of attendance, completed delivery, or approved future bookings.

Complete Session 1 content supplied with this collection

The collection includes *s2-ocai-go-session-1-full-content.pdf*, a conversion of the complete learner page at *content/Session 1/session-1-post-session.html*. It retains the lesson text, activities, copyable prompts, source checks, reflection, takeaways, and resource links. It is the full lesson, not an overview extract. The PDF provides enough substantive content for a Project to inspect teaching approaches, reuse actual prompts, and compare the foundation with later session needs.

Session 1 is *From a Prompt to a Reusable AI Teammate*. It develops the idea that useful context is more than a single request. Context can include the current task, standing instructions, conversation history, examples, files, and information reached through tools. Different products retain and retrieve these differently. Saving a role or creating a Project supplies context; it does not change the underlying model through retraining.

Thought partner

Participants build a role that helps them think through a real decision, surface assumptions, ask sharper questions, and consider other perspectives. The person brings the situation and retains the judgment. The test is whether the dialogue improves the person's thinking and respects the intended role.

The learner page includes the instructions and practice needed to create and test the thought partner, then save it in a reusable home. A fresh conversation tests whether the useful setup carries forward. The full PDF retains the actual activity and prompts, so a course-management assistant can inspect what participants were asked to do rather than rely on this short description.

Editor

Participants build an editor for their own writing. Its role is critique, with authorship and revision decisions left to the writer. A useful editor should follow the requested boundaries, notice meaningful weaknesses, and give actionable feedback. Fluent replacement prose is not the sole success criterion.

The lesson supports saving the editor, adding useful knowledge when appropriate, and testing whether the role remains useful beyond the first exchange. It gives participants a reusable artifact as well as a way to evaluate it. The thought partner and editor are foundations for the maintained context introduced in Session 2.

Optional critic, reflection, and later use

A hypercritical critic provides an optional extension. Reflection asks participants to articulate what they delegated, how they judged the response, and what remains human work. The takeaway is a configuration they can use again and a next-use test. Building it once can make it Level 3-ready; observed reuse is the stronger evidence.

The full page also supplies preparation and resources. Product-specific instructions are teaching surfaces, not guarantees about another person's account. Actual persistence and available features should be checked when rehearsing the live path.

Historical June pilot and supporting course material

The June 23, 2026 trial retains historical Session 0 filenames and includes a recording, deck, learner page, and delivery assets. It is the historical foundation for the live program's Session 1. It does not create a current learner-facing Session 0 or override the revised program sequence.

A course shell contains a start page, schedule, resource library, and help. The presence of local source and rendered files establishes that material exists locally. It does not establish that the latest version has been published or that every scheduled offering occurred.

Session 2 storyboard overview

The current storyboard is *04-storyboard/session-2-projects-and-notebooks.md*, a draft for review dated 2026-10-01, revision 1. Its compact summary records the same design. It has not become an approved production outline.

Intended promise and outcomes

The draft promises a reusable workspace for one body of work: define its boundary, select trustworthy sources, test two related tasks, and record how the workspace will stay current. Its durable distinction is that assistants carry a role and Projects carry a body of work.

The four outcomes are to design a scoped workspace, justify source choices, evaluate both task results against actual passages, and explain maintenance and later use. Source judgment includes authority, relevance, currency, permission, and unresolved conflict. A participant should record whether an answer is accepted, needs revision, or remains unresolved, with reasons.

Draft activity design

The draft uses the fictional Customer Onboarding Refresh as its shared example. The saved artifact consists of a workspace brief, two short sources, and a test/maintenance record. Task A is a status brief; Task B is preparation for the next working meeting.

Its eight beats move from recalling recurring context, through the workspace model and source selection, to configuration, the two task tests, maintenance, and sharing. The 55-minute plan allocates 11 minutes to facilitator content and 44 minutes to learner activity, including interface operations and transitions. These figures describe the existing draft; they are not a newly rehearsed timing plan.

The floor is the saved brief, sources, two task attempts, source checks, and maintenance note. A portable packet and prepared results cover access failure while recording that persistent configuration remains pending. The stretch adds an authoritative update, supersedes the affected source, and reruns the affected question. Later reuse is still needed to establish Level 3.

Newer instructor direction that still needs incorporation

Subsequent planning changes the delivery to two demonstration/build pairs: Project demonstration and participant Project build, then Microsoft 365 Copilot Notebook demonstration and participant Notebook build. The Project product remains to be selected. Four authentic examples provide alternatives for instructor and non-instructor work:

• Course management: this OCAI GO collection, for a Project.

• Community planning: fuller GenAI CoP purpose, planning, and historical records, for a Project.

• Course content: the complete Learning and Memory section of Discover Psychology 2.0, for a Notebook.

• Workplace onboarding: complete public OC College Relations, event protocol, and brand guidance, for a Notebook.

These examples replace the need to rely only on a fictional two-excerpt packet. The older draft's single build, proposed ChatGPT path, and Google Notebook comparison remain visible as earlier design choices. They should not be presented as the current delivery instruction. This inventory reports the discrepancy; it does not rewrite or approve the storyboard.

What exists and what remains

Session 2 has the draft storyboard, a public learner-page placeholder, a working plan, this fuller source collection, and a private source-review page. The supplied sources and practice prompts support preparation and rehearsal. A canonical production outline, finished output specifications, and an approved delivery deck are not established by those review artifacts. The actual account path, source ingestion, evidence inspection, persistence, and timing still need rehearsal.

Remaining Round One sessions

The approved sequence defines the following directions. Current learner-page placeholders are not complete lessons. Migration storyboards are raw material for adaptation.

Session 3: Documents, Data Displays & Decks

Turn useful context into an editable work product. The shared example remains a presentation, with broader document and data-display possibilities handled as alternatives or future editions. Participants should leave with an editable deliverable and a review of evidence and quality. Checks include invented claims, source fidelity, editability, accessibility, formatting, audience fit, and remaining human work. Session 1 is recommended preparation. The retained decks storyboard requires adaptation to the current sequence and verified creation path.

Session 4: Teach AI a Repeatable Move: Build a Skill

Create a reusable procedure with instructions, examples, success criteria, and a fresh-input test. Literal product Skills can be used where supported; a procedural asset is the product-neutral floor. Saving a procedure and testing it on new material makes it Level 3-ready. Session 2 is recommended preparation. A new storyboard and developed lesson are still needed.

Session 5: Connect AI to Your Real Work, and Decide What's Worth Connecting

Create a connection decision record and, where access permits, test one non-sensitive connection. The Connection Check considers task value, sensitivity, permissions, read/write scope, failure consequences, and a disconnect/review plan. Session 1 is recommended and Session 2 helpful. The earlier connectors storyboard is a migration source for this revised session; product and access claims need updating.

Workflow Studio 1

Combine two or more Round One building blocks into a personal workflow. Recommended preparation is two of Sessions 1-5. The Studio should demonstrate reuse on a fresh input or later occasion, rather than merely create another configuration. Its own storyboard and materials remain to be developed.

Round Two and recurring practice

Session 6: From Search to Evidence-Checked Research

Build a cited brief, reusable research protocol, source rules, audit checklist, and second-question test. A one-time report is task-level competence; the reusable method and another question supply the integration evidence. Session 1 is recommended. Product cards may name current research features, but the session remains tool-agnostic.

Session 7: Turn a Conversation into a Working Tool

Build one tiny internal tool and a build-versus-release boundary note. The core avoids authentication, persistent personal data, and public release. The person should explain what the tool does and what they can responsibly judge. Session 1 is recommended. A working prototype does not establish readiness for unrestricted release.

Session 8: Set It, Check It: Build a Scheduled Workflow

Create a simple scheduled AI task on a verified browser path, with monitoring and maintenance. Combining Project, Skill, connector, and schedule is a stretch because one accessible surface may not support the whole composition. Session 4 or 5 is recommended preparation. A schedule needs a review trigger and a clear response to failure or stale sources.

Session 9: Your First Agent: Give AI Bounded Autonomy

Create a bounded agent specification, stopping rule, human checkpoint, and observed test run; create a working agent where access permits. Session 6 is recommended preparation. A written specification alone is not equivalent to a tested agent. The test and recovery discussion make the delegation boundary concrete.

Workflow Studio 2 and Demo Day

Workflow Studio 2 combines two of Sessions 6-9 into a tested workflow with a before/after process map and retained human checkpoint. It can demonstrate Level 3 or Level 4 depending on the evidence. Demo Day recurs as a way to explain a useful output, observed reuse, or a documented process change, including what broke.

Sessions 6-9 and both Studios have learner-page placeholders and require developed storyboards and production. Exact winter dates are not supplied in the source schedule. Demand, capacity, and verified access should inform delivery planning before those offerings are locked.

Status words to retain in an answer

Developed means substantive material exists. *Rendered* means a local output exists. *Historical* means a retained earlier record. *Draft* means review remains. *Placeholder* means navigation or a brief description exists without the complete lesson. *Tentative* means the scheduling source has not confirmed the booking. None of these words alone establishes publication, attendance, delivery, approval, or current account access.

Original sources

Checked against *content/Session 1/session-1-post-session.html*, its post-session Markdown, *content/README.md*, *04-storyboard/README.md*, the Session 2 storyboard and summary, *03-structure/session-sequence.md*, *session-2-handoff.md*, and the current owner instructions. Use the separate pathway and schedule for dates and booking status. The complete Session 1 PDF is an intentional delivery copy of its accepted learner HTML.