OCAI GO: program overview

Expanded course-management demonstration source, prepared 2026-10-02 from the current program brief, objectives, and sequence. This copy provides substantial working context; the original project records remain authoritative.

Identity and purpose

OCAI GO means *Okanagan College AI Growth Opportunities*. OCAI stands for Okanagan College AI. The series is an ongoing, practical companion to *AI Fluency Foundations*. It helps members of the Okanagan College community use generative AI in their actual work, build reusable ways of working, and judge whether those ways are useful.

The central question is what a person can build and use again. A good answer to one prompt is useful, but recurring work usually needs more: standing instructions, reference material, examples, a procedure, access to appropriate information, and a clear place for human judgment. The series develops those capabilities through working artifacts rather than through a tour of product features.

Its interests include reusable AI teammates, Projects and Notebooks, editable documents and presentations, Skills, connections to workplace information, evidence-checked research, small working tools, scheduled tasks, and bounded agents. These are examples within an evolving program. A named product provides a teaching surface for a particular offering; it does not define the permanent curriculum.

OCAI GO draws on prior workshops and GenAI Community of Practice material, and on the custom-tools and agents content of AI Fluency Foundations. Past workshops supply useful evidence and resources. They are not templates that must be repeated unchanged. The series aims to keep pace with practical capabilities as tools and needs change.

Who it serves and how people participate

The intended audience includes faculty, staff, administrators, and students. Examples and core activities should make sense across that mix. Faculty may bring course materials or assessment tasks; staff may bring event guidance or a continuing administrative project; students may bring readings or a study workflow. No one role is the default participant.

People arrive with different tools, confidence, access, and reasons for attending. Motivations include professional development, improving current work, understanding emerging workplace skills, and helping peers and students use AI thoughtfully. A shared floor lets a newcomer make something useful, while a stretch gives an experienced participant further work.

Participation is voluntary and intrinsically motivated. There are no grades, badges, credit requirements, attendance gates, or fixed completion point. There is no terminal capstone. Success can be a useful artifact from one session, a personal workflow used across several sessions, or a contribution that helps someone else.

Comfort with prompting for a work task is recommended preparation. AMII's AI Workforce Readiness and Anthropic's 4Ds are suggested foundations. Prompting fundamentals, model-development theory, and the dedicated social and ethical deep dive from AI Fluency Foundations are not recurring core lessons here. Responsibility still accompanies every practical build.

Format and practical constraints

Core workshops last 55 minutes and use a browser-first approach. The brief describes online delivery, possibly hybrid, with modality determined for an offering. A cadence of roughly every two weeks is aspirational. The actual dates, times, and booking status belong in the separate pathway and schedule document; a passed date is not evidence that a workshop occurred.

The program is unfunded and supported by the instructor's time. Paid AI licences are strongly recommended in the brief, but the College does not supply those licences through this program. Actual account access, administrative restrictions, and the proposed live path must be checked before delivery. A portable artifact and prepared fallback can preserve useful work when access fails.

The brief names Brightspace as the likely host. Current production also has a Moodle website path. These are delivery arrangements, not evidence that a particular review page has been uploaded or released. Each offering should identify where participants can find preparation, recordings or equivalent walkthroughs, finished resources, and the next-use activity.

Desktop tools, command-line work, custom connectors, and heavier technical composition belong in the advanced bank. They should not quietly add requirements to the browser-first core. Technical work is appropriate when it serves a practitioner and when the person can judge the limits of what they are building or releasing.

Four Levels of AI Use

This is the only numbered learner-facing progression. It describes a particular use of AI, not a person's permanent ability.

1. **Tool basics and confidence.** The person can operate the interface and its basic features independently. Evidence is successful operation, such as finding and using a feature.

2. **Task-level competence.** The person can use AI to complete a useful task and evaluate whether its output fits the purpose. An attractive output alone does not establish competence; the judgment matters.

3. **Workflow integration.** The person has made a capability reusable or connected it to real work. Evidence requires reuse on a fresh input or later occasion without rebuilding the setup. A saved workspace in a workshop can be Level 3-ready until that reuse happens.

4. **Workflow redesign.** The process changes because AI can reliably carry a part of it. Evidence requires a documented changed, removed, or reassigned process step and the human checkpoint that remains. A proposed change can be Level 4-ready until the process change is demonstrated.

Someone may operate at different levels in different tasks. Studios, later reuse, and Demo Day provide opportunities to show stronger evidence than an initial in-session configuration.

Capability building blocks

The internal design map uses six unnumbered building blocks: effective task use, reusable context, interactive collaboration, durable configurations, composed workflows, and bounded autonomy. These organize concepts and dependencies. They are not a second maturity ladder.

Useful distinctions recur throughout the pathway. An assistant carries a role. A Project carries a body of work and what it should know. A Skill carries a repeatable procedure and how it should do the work. A connector determines what information or services it can reach. Scheduling determines when work runs. An agent specification determines what choices AI may make, when it must stop, and where a person intervenes.

Product implementations differ. Saving instructions or uploading files supplies context at runtime; it does not retrain the underlying model. What persists, what is retrieved, and what updates automatically must be tested in the actual account. A collection can support several tasks, but its usefulness depends on source selection, clear boundaries, and maintenance.

The three Ds around every build

Delegation means deciding what AI should carry and what a person should decide. For example, AI might compare source documents or draft a status brief, while a person confirms an event booking, adopts a policy, or approves release.

Discernment means evaluating the tool, process, and output. A source citation gives a location to inspect. It does not establish that the claim is supported, complete, current, or appropriate. Checks vary with the task: source fidelity, accuracy, editability, accessibility, audience fit, or a useful fresh-input test.

Diligence means attending to privacy, permissions, disclosure, oversight, maintenance, and human capability. It also includes the boundary between building something for personal or internal use and releasing it to others. People should keep the skills needed to understand and review delegated work.

These three Ds derive from Anthropic's four Ds. Description is absorbed into the wider instruction about context, Projects, Skills, connections, workflows, and agents. OCAI GO does not present the three Ds as a competing framework. Each session identifies one explicit delegation decision, one discernment test, and one diligence requirement.

Course-level outcomes

The four durable outcomes are aspirational across an ongoing series, not attendance or week-by-week gates.

• **Engineer durable context:** build reusable instructions, configurations, and sources for recurring work; explain the choices; test whether the result is more consistent than rebuilding context ad hoc.

• **Apply the three Ds:** justify delegation and tool choices, evaluate fit for purpose, protect appropriate information, disclose AI use appropriately, and retain meaningful human judgment.

• **Compose and sustain a personal workflow:** combine two or more building blocks into recurring work, demonstrate reuse, and explain how the configuration will stay current as tools and sources change.

• **Share practice and strengthen the community:** explain what was built and what broke, help a colleague or student, and contribute a reusable artifact or write-up.

The recommended pathway

Round One develops reusable teammates, maintained workspaces, editable deliverables, repeatable procedures, and appropriate connections. Workflow Studio 1 combines two or more of these into a personal workflow. Round Two develops research, small working tools, scheduling, and bounded autonomy. Workflow Studio 2 examines a changed process and its retained human checkpoint.

Demo Day is a recurring practice of showing what was built and what broke, rather than an additional capability lesson. Advanced offerings may address packaging and sharing a workflow, structured multi-step work, custom connectors, or desktop use when demand and access justify them.

Recommended prerequisites are recoverable. A catch-up path should provide a concept recap, readiness check, copyable starter artifact, and access to the earlier session. Participants can follow the sequence or recover the preparation they missed. Each session should still yield a useful standalone outcome.

How a workshop establishes useful practice

The standing design contract protects one shared core build, one floor, and one stretch. Participants follow *build → test → evaluate → share* with real, permitted material. The concept comes before the current menus; volatile product details belong in dated tool cards. Configuration is followed by a meaningful test, a human evaluation, and an explanation of what remains unresolved.

For the current Session 2 planning, the instructor has explicitly requested a short Project demonstration followed by a participant Project build, then a short Notebook demonstration followed by a participant Notebook build. That later instruction is recorded in the working plan. The existing storyboard still describes one shared build and needs reconciliation before it can govern production. The four supplied examples are choices for those two parts, not four workspaces everyone must build.

Using this collection for course management

The Project can answer questions about the program, compare existing materials with upcoming needs, draft explanations for newcomers, or identify dependencies before delivery. It should distinguish the approved program architecture, tentative schedule, developed learner material, draft storyboards, review sources, and empty placeholders. It should retain uncertainty rather than infer delivery, approval, or readiness from a filename or date.

This overview supplies substantial context. The separate schedule owns dates. The developed-content inventory owns the current production snapshot and points to the complete Session 1 lesson. The Session 2 working plan records the newer demonstration direction. The working and production guide describes how decisions become teaching artifacts.

Original sources and authority

Prepared from *00-brief.md*, *02-objectives.md*, *03-structure/session-sequence.md*, *AGENTS.md*, and the relevant production workflows, checked 2026-10-02. Program purpose and boundaries come from the brief; outcomes from the objectives; pathway and migration decisions from the sequence. Dated handoffs and later owner instructions record working direction. This demonstration copy does not silently approve or replace an upstream record.