

BOOTCAMP

Project Controls in the Age of AI: Building AI Agents for You and Your Team

Bootcamp Speakers



Abhi Datta

Managing Partner
ClarityAxis Group



Hiren Soni

Managing Partner
ClarityAxis Group



Edmund Chew

Deputy Director - Planning
Project Controls Institute

Date

23 November 2026

Time

09:00 – 17:30

Venue

**Melbourne Cricket
Ground**

Course Overview

AI is transforming the way organisations plan, deliver and govern projects. As AI rapidly evolves from simple chatbots to intelligent agents capable of analysing information, automating repetitive tasks and supporting complex decision-making, project professionals have an opportunity to fundamentally change how project controls are performed and how value is delivered.

This is a practical, hands-on workshop that will equip participants with the knowledge and skills to understand, build and apply AI agents. Participants will leave with working AI agents and practical means to roadmap and adopt AI within their teams and organisations. By the end of the workshop, participants will be able to:



Distinguish clearly between chatbots, copilots and agents, and explain what has changed in the latest generation of AI.

Build, test and deliberately break a working AI agent using no-code tooling.

Decompose any agent into its six working parts: goal, instructions, knowledge, tools, guardrails and output.

Apply an autonomy-tier governance model and a one-page usage policy to their own agents.

Select and prioritise AI use cases on value and feasibility rather than enthusiasm.

Produce an adoption roadmap covering vision, use cases, people, data, technology, governance, measures and timeline.

The workshop is structured to provide a working understanding of key concepts, extended hands-on practice with industry tools, and a real implementation case study to reinforce learning. Every concept is tested against real project controls work.

Approach:

Build in the Morning, Plan in the Afternoon.

The day is deliberately structured in two halves. The morning is about theory and building. Short sessions leading into an extended hands-on block where every participant constructs a working agent. The afternoon is a plan: governance and assurance, then a facilitated workshop where participants produce a twelve-month adoption roadmap.

This structure was chosen because the gap between reading about agents and having shipped one is where most organisations stall. Closing that gap in a single facilitated day and pairing it with the governance and adoption work that makes it stick. No coding experience is required, and all build activity uses platforms that most organisations already have.

What you'll get:

A working AI agent, built by you during the day and configured against your own workflow, not a demonstration of someone else's work.

A completed one-page roadmap canvas for a real team, with a stop rule and a named first action.

A six-part prompting framework and a reusable instruction library for common project controls tasks, including variance commentary, schedule integrity checks and change triage.

A one-page AI usage policy template, ready to issue as version 0.1, covering approved tools, data rules, autonomy tiers, ownership, disclosure and review.

The full slide pack, canvas template, policy template and sample data files.

Agenda

Opening

08:30 – 09:00

Registration, Tech Check and Wi-Fi Setup

09:00 – 09:15

Welcome and How Today Works

Setting the scope of the day and getting every participant build-ready.

A short orientation covering the structure of the day, the two hands-on blocks, and the takeaways participants will leave with.

- Introductions and the thread for the day: automate the mechanical eighty per cent, return the judgement twenty per cent.
- Quick poll of the room on current AI maturity, tools already in use, and organisational constraints.

Session 1

09:15 – 10:00

Evolution of AI – Context for Project Controls and Tools Overview

Cutting through the hype to a working definition that leaders can actually use.

The term AI agent is used loosely and often incorrectly. This session runs as a structured conversation rather than a presentation, with Edmund putting the questions the room is already thinking to two practitioners who build and deploy these tools for a living.

Four generations of AI – rules, prediction, generation and agency

- Where agents already land in a project controls day, across schedule, cost, risk, reporting, contracts and assurance.
- The tools landscape in three shelves: general assistants, agent builders, and project-controls-native platforms.
- What most teams already own and licence today, and what genuinely requires a purchase decision.

Session 2

10:00 – 10:30

AI Strategy for Project-Oriented Organisations

Why tools are the wrong place to start.

A deliberately short session that gives its time to the agent build. The purpose is to give the room a way of thinking about where AI belongs in a project business before anyone builds anything.

- Three failure modes: tool-led pilots, data-starved use cases, and owner-less initiatives.
- A value-versus-feasibility screen for sorting quick wins, strategic bets, hygiene tasks and the demo graveyard.
- Build, buy or partner – the trade-offs of each route and when a mix is the right answer.
- The four foundations of an operating model: people, data, governance and platform.

Agenda

10:30 – 10:45

Morning Break

Session 3

Hands-On Agent Building Workshop

10:45 – 12:30

From a prompt to a working agent, live and in the room. Laptops open.

The core build block of the day, run in two parts: a live end-to-end build that participants watch, followed by their own build in small groups. Nothing is pre-prepared, and the stumbles are part of the content.

- » The anatomy of an agent: goal, instructions, knowledge, tools, guardrails and output.
- » The six-part instruction – role, context, task, format, examples and constraints
- » Watch a Schedule Quality Reviewer agent built start to finish, including running it on real data and deliberately breaking it.
- » Participant build: choose one of three briefs – month-end narrative drafter, change and variation triage, or schedule quality reviewer.
- » Show and tell: what worked, what broke, and what surprised you

12:30 – 13:15

Lunch

Session 4

AI Governance and Security, Ethics and Best Practices

13:15 – 14:15

The wall of things that broke this morning and what to do about each one.

Governance introduced through the failures the room produced an hour earlier. None of them are surprising and all of them are manageable.

- » Four ways it goes wrong on a real project: data leakage, confident nonsense, IP and confidentiality, and the accountability gap.
- » Five questions to ask before any AI tool touches project data: processing location, model training, access, audit trail and independent verification.
- » Three tiers of autonomy – draft, assist and automate. Why every agent starts at tier one and earns its way up with evidence.
- » Six habits of teams that use AI well, from disclosure through to verifying before trusting.
- » Building a one-page AI usage policy: approved tools, data rules, autonomy tiers, ownership, disclosure and review.

Agenda

Session 5

14:15 – 15:00

Roadmaps for Implementing AI in Your Team and Organisation

Most teams are at stage two and believe they are at stage four.

Building an agent is the easy part. Getting one adopted, governed and still running after the enthusiast rolls off the project is the work that actually changes an organisation. This session presents a maturity framework, an eight-element roadmap structure, and a real implementation journey mapped against both.

- » Five stages of maturity – aware, experiment, pilot, scale and embed.
- » Why self-assessment consistently overstates maturity, and how to use evidence rather than intent to place your team.
- » The eight elements of a roadmap that survives contact with a project: vision, use cases, people, data, technology, governance, measures and timeline.
- » Baseline measures before you start: hours returned, cycle time, error rate and adoption. The numbers that make a benefit claim defensible.
- » Ownership that outlasts individuals: a named owner and a named approver for every agent, written into the process rather than the enthusiasm.

15:00 – 15:15

Afternoon Break – Roadmap Canvases Distributed to Tables

Session 6

15:15 – 16:20

Hands-On AI Roadmap Workshop

One page, one team, twelve months

The second hands-on block, and the one participants take back to their organisation. Working in small groups, ideally with colleagues from the same organisation together, each participant builds a twelve-month AI roadmap for a real team on a single page.

- » Choose the team: a real one, named, of ten people or fewer to keep the plan honest.
- » Set the vision in one sentence, tested against the eighty-twenty thread – what is AI for in this team, and what does it give back to people?
- » Fill the canvas quarter by quarter across all eight elements, with a quick win landing in the first quarter.
- » Name an owner against every use case, and the specific data cleaning work that has to happen first.
- » Set the measures and the baseline, then write the stop rule – what result, by when, would make you pause.

Agenda

Session 7

16:20 – 16:50

Wrap-Up, Key Takeaways and Next Steps

Five things to take out of the room, and what to do with them.

Consolidation of the day, delivered jointly, followed by a staged set of next actions.

- Agents act, chatbots answer – and today you built one that can reach your systems.
- Guardrails first, instructions second. The agent that says “I am not sure” is the one you can defend.
- Strategy is choosing what not to do. Two quick wins in production beat a thousand pilots.
- Tier one is where agents live. Autonomy is earned with evidence, one tier at a time.
- The roadmap is a standardisation roadmap. The agent is only as good as the standard behind it.

16:50 – 17:00

Certificates, Feedback and Close

Who Should Attend

This program is designed for team and division managers, department heads and directors who are accountable for project delivery performance and who must make decisions about where AI belongs in their function.

It suits leaders who need more than an awareness session. It is catered to those expected to set direction, approve tooling, carry the risk, and lead their people through the change. Senior practitioners who will be asked to build and own the first agents will also get full value from the day.

Prerequisites:

► Attendees should have:

- Working familiarity with project controls or project delivery processes
- Some prior exposure to a generative AI tool at a basic level
- A laptop with internet access and permission to sign in to browser-based tools

Participants are encouraged to bring a de-identified sample of their own project data, one repetitive task they would like to automate, and one real team they would like to build a roadmap for. All three are used directly in the hands-on sessions, and this ensures attendees can meaningfully participate in the build and planning exercises rather than observing them.

Presenters:

Three practitioners, one build day. The team combines AI strategy and platform experience with long-run project controls delivery and implementation practice.

Abhi Datta, Co-founder, ClarityAxis Group – Leads the strategy session, co-presents governance and security, and co-facilitates the agent build.

Hiren Soni, Managing Partner, ClarityAxis Group – Leads the live agent build and co-presents governance.

Edmund Chew, Project Controls Institute – Presents the roadmap framework, facilitates the afternoon roadmap workshop, and drives the opening dialogue session as the questioner on behalf of the room.



BOOK NOW!

**Individual | Corporate Discounts &
Group Booking Available**