

CMG Live Stream · Monday 28 September 2026

AI Agents for the Workplace

Enda Rochford

AI Training Manager

RandTrad Consulting

Today's schedule

Time	Session
09:30 – 11:00	Session 1 — Understanding & Identifying Agents
11:00 – 11:15	Coffee break
11:15 – 13:00	Session 2 — Designing & Building Agents
13:00 – 14:00	Lunch
14:00 – 15:30	Session 3 — Governance, Risk & The Future
15:20 – 15:30	Q&A and close

Glossary

Term	Meaning
Agent	AI that takes actions to complete a task
Chatbot	AI that answers questions
Single-use agent	One context, one task, one output
Autonomous agent	Persistent, self-improving, runs across tools
Vibe working	Creating documents by chatting with AI
Prompt injection	Malicious instructions hidden in input
Hallucination	Confident but incorrect AI output
Human-in-the-loop	A person reviews and owns the output
AI literacy	Staff understanding of safe AI use (EU AI Act Art 4)

How familiar are you with AI agents?

Quick temperature check: which sounds most like you?

- | | | |
|----------|-------------------------------|--------------------------------------|
| 1 | Agent curious | I haven't used an agent yet. |
| 2 | Tried an agent | I've experimented once or twice. |
| 3 | Using agents regularly | Agents already help with my work. |
| 4 | Building my own agents | I've created or customised an agent. |

Pop 1, 2, 3 or 4 in the chat. All starting points welcome.

What you'll be able to do by end of today

Recognise

Tell an AI agent apart from a chatbot

Build

A single-use agent with tools you already have

See

Microsoft Office Agent create documents from a prompt

Govern

Agents responsibly under the EU AI Act

No hype. No jargon. Practical, evidence-led.

SESSION 1

09:30 – 11:00

Understanding & Identifying Agents



BY THE END OF THIS SESSION YOU WILL BE ABLE TO:

1

Define an AI agent and tell it apart from a chatbot

2

Describe the agent spectrum, from single-use to autonomous

3

Recognise that single-use agents are one concept across platforms

4

Identify where agents could fit in your own workplace

A chatbot answers. An agent acts.

Works towards a goal

**Uses tools — search,
files, apps**

**Chains multiple steps
together**

**Iterates — checks its
own work and refines**

Agent vs chatbot vs assistant

	Chatbot	Assistant	Agent
Does what?	Answers questions	Helps with tasks	Completes goals
Tool use?	No	Limited	Yes
Multi-step?	No	Some	Yes
Iterative?	No	No	Yes
Example	Website bot	Siri / Alexa	Office Agent, Custom GPT

The agent spectrum

Level 1 — Single-use

One context, one task, one output

Custom GPT, Gemini Gem, M365 notebook
· Low autonomy

Level 2 — Multi-step

Orchestrates several steps

Office Agent, n8n workflow · Medium autonomy

Level 3 — Autonomous

Persistent, self-improving, runs across tools

Claude CoWork, Hermes, OpenClaw · High autonomy

Custom GPTs, Gemini Gems and M365 notebooks are the same thing

Specific instructions

Optional context or knowledge

One defined task

One context, one task, one output. The platform changes, the concept doesn't.

Single-use agents — same concept, different platforms

Custom GPT

ChatGPT

Built for builders. Needs ChatGPT Plus to create.

Gemini Gem

Google

Built for users, simpler and more personal. Integrates with Gmail, Docs, Drive. Free tier available.

M365 Notebook

Microsoft

Same concept in the Microsoft ecosystem. Part of your existing M365 licence.

In Session 2 I'll build the same agent in all three, live, so you can see this for yourself.

Where agents fit in workplace workflows

Agents are best at draft, summarise, analyse, create — not at decide.

Drafting

Reports, proposals,
emails, briefings

Summarising

Long documents,
meetings,
consultations

Analysing

Spreadsheets, tender
responses, feedback

Creating

Presentations,
structured documents

Reviewing

Checking against
criteria, compliance

Irish public sector examples

- Drafting responses to consultations
- Summarising policy documents
- Preparing board papers
- Analysing consultation responses
- Creating briefing decks for management

The pattern: agents draft, humans decide.

Where in your own work could a single-use agent help?

Drafting

Summarising

Analysing

Creating

Hold that thought — in Session 2 I'm going to build one live.

Where agents fit — by sector

Construction & Built Environment

- Drafting method statements and risk assessments
- Summarising site reports and inspection notes
- Reviewing compliance documentation against standards
- Generating progress reports from raw data

Marketing & Communications

- Drafting campaign briefs and creative concepts
- Summarising campaign performance reports
- Analysing engagement data across channels
- Generating first-draft content for review

Public Sector & Government

- Drafting responses to consultations
- Summarising policy documents
- Preparing board papers
- Analysing consultation responses

General Office Workflows

- Meeting minutes → structured output
- Long documents → executive summaries
- Tender responses → compliance checking
- Spreadsheets → narrative analysis

The task changes. The pattern doesn't: agents draft, humans decide.

Coffee Break

11:00 – 11:15

SESSION 2

11:15 – 13:00

Designing & Building Agents



BY THE END OF THIS SESSION YOU WILL BE ABLE TO:

1

Understand how Microsoft Office Agent creates documents from a prompt

2

Describe the “vibe working” concept: start with a prompt, steer iteratively

3

Recognise that building a single-use agent is the same process across platforms

4

Identify n8n as a no-code automation platform, for awareness

Session 2 — what we'll do

- 1 Kick off Office Agent — start a document generation (it takes time)
- 2 Build the same agent live across Custom GPT, Gemini Gem and M365 notebook
- 3 Return to Office Agent — show the result and steer it
- 4 A quick look at n8n — an automation platform, for awareness
- 5 Take-home exercises — templates you can use after the course

Microsoft Office Agent

What it is

- Microsoft's chat-first document creation tool
- Creates PowerPoint decks and Word documents from a prompt
- Powered by Anthropic models

Availability

- M365 Personal, Family or Premium
- Rolling out via the Frontier program
- Check your own account access

“Vibe working”

The office equivalent of vibe coding

1. Start

Start with a prompt

2. Steer

Steer iteratively — “make it more visual”,
“add a section on X”

3. Hand off

Hand off to PowerPoint or Word for final touches

You don't write the document. You direct it.

Office Agent — the three-step process

Step 1 — Clarify intent

It asks about length, theme, focus, audience

Step 2 — Deep research

Web research with a live preview. Shows its chain of thought — this takes time

Step 3 — Produce content

Generates a ready-to-share document. Hand off to PowerPoint or Word

Live demo: kick off Office Agent

- 1** Open officeagent.microsoft.com
- 2** Enter a prompt: “Create a deck summarising the top 5 trends in the Irish public sector for 2027”
- 3** Answer the clarifying questions
- 4** Let it run — the deep research step takes several minutes

While Office Agent is working, we'll build the same single-use agent across three platforms, so you can see the concept in action.

Building the same agent across three platforms

The process is the same every time: name it, describe it, write the instructions, add context (optional), test it.

The agent

- Meeting Minutes Summariser

The task

- Raw meeting notes → structured minutes with attendees, decisions, action items and open questions

Watch how the fields are identical — only the logo changes.

Build #1: Custom GPT

chatgpt.com → Explore GPTs → Create

Field	Content
Name	Meeting Minutes Summariser
Description	Turns raw meeting notes into clean, structured minutes
Instructions	Produce structured minutes with attendees, decisions, action items (owner + deadline), and open questions. Neutral tone. Flag anything unclear. Never invent facts.
Knowledge	Optional — your organisation's minutes template
Capabilities	Web browsing off, code interpreter off

Live: enter the fields, save, test with a sample input.

Build #2: Gemini Gem

gemini.google.com → Gem manager → Create

Field	Content
Name	Meeting Minutes Summariser
Description	Turns raw meeting notes into clean, structured minutes
Instructions	Same as the Custom GPT — word for word
Knowledge	Optional — your organisation's minutes template

Same fields. Same instructions. Same result. Different logo.

Build #3: M365 Notebook

M365 → Notebooks → Create

Field	Content
Name	Meeting Minutes Summariser
Description	Turns raw meeting notes into clean, structured minutes
Instructions	Same again — word for word
Context	Optional — your organisation's minutes template

Three platforms. One process. One concept.

I just built the same agent in three different platforms using the exact same process

Same name

Same description

Same instructions

Same test input

Same quality of
output

The hardest part is writing good instructions, and that's platform-independent.

Return to Office Agent

Let's check on our deck...

- Show the completed result, or near-complete
- Show the chain of thought — the research it did
- Steer it live: “make it more visual”, “add a slide on AI adoption”, “shorten it to 5 slides”
- Hand off to PowerPoint for final touches

This is vibe working — you direct, the agent builds.

n8n — just so you know it exists

What it is

- Open-source, no-code automation platform
- Connect nodes on a visual canvas to build workflows
- Native AI integration
- Can be self-hosted, for data control

What it looks like

- Trigger node → Process node → Output node
- Example: a new form submission triggers AI to summarise it, then an email is sent

The “connect things together” tool. Not today's focus, but worth knowing about.

Take-home exercises

You have the templates and links in your attendee pack

Exercise 1 — Build your own single-use agent

- Pick one task you do repeatedly
- Use the template: name, description, instructions, context
- Build it in Custom GPT, Gemini Gem or M365 notebook
- Test, refine, repeat

Exercise 2 — Try Office Agent

- Go to officeagent.microsoft.com
- Create a deck or document from a prompt
- Steer it with follow-up commands
- Verify the output

You don't need to do these live — do them when you're back at your desk.

Session 2 — recap

- Office Agent — kicked off a prompt, let it research, steered the result
- Single-use agents — built the same agent across three platforms, same process every time
- n8n — an automation platform, for awareness
- Take-home exercises — build your own when you're ready

You've seen agents built live. Now, how do you govern them responsibly? That's Session 3.

Lunch

13:00 – 14:00

SESSION 3

14:00 – 15:30

Governance, Risk & The Future



BY THE END OF THIS SESSION YOU WILL BE ABLE TO:

- 1 Classify your agents under the EU AI Act risk framework
- 2 Apply transparency obligations under Article 50
- 3 Explain human oversight and accountability for agent outputs
- 4 Identify the four key risks and their mitigations
- 5 Describe what's coming, and what it means for the future of work
- 6 Leave with a concrete action plan

The EU AI Act — a risk-based framework

Regulation (EU) 2024/1689

Risk level	Examples	Obligations
Prohibited	Social scoring, manipulation	Banned
High risk	Recruitment, credit scoring, education	Full obligations from Dec 2027
Limited risk	Chatbots, content generation	Transparency, Article 50
Minimal risk	Most workplace agents	General principles only

Most workplace agents are minimal or limited risk.

EU AI Act timeline

Obligation	Date	Status
AI Act entered into force	1 Aug 2024	Done
Article 4 — AI literacy	2 Feb 2025	In force
Article 50 — Transparency	2 Aug 2026	In force
Annex III — high risk	2 Dec 2027	Pending
Annex I — product-embedded	2 Aug 2028	Pending

Article 4, AI literacy, is why this course exists — training your staff is a direct legal response.

Article 50 — transparency obligations

In force since 2 August 2026

The rules

- 50(1): tell people when they're interacting with AI, unless obvious
- 50(2): mark AI-generated content in a machine-readable way
- 50(4): disclose deepfakes
- AI-generated text on public interest matters must be labelled

Practical actions

- Map where your organisation generates AI content
- Add disclosures where people interact with AI
- Mark AI-generated content

Human oversight — the core principle

Human-in-the-loop

A person reviews and owns the output

Sign-off process

Important outputs require human approval

Accountability

The person who reviews or publishes owns the output, not the tool

Audit trail

Keep records of what was asked, produced, and changed

"AI advises, humans decide."

The four key risks

Risk	What it is	Mitigation
Hallucination	Confident but wrong	Verify against source, cite, cross-check
Prompt injection	Malicious instructions hidden in input	Don't paste untrusted text
Data leakage	Sensitive data exposed	Know your data, no personal data in public tools
Over-reliance	Trusting output without checking	Verification culture, human sign-off

Governance walkthrough — the Meeting Minutes Summariser

Let's run the agent built in Session 2 through the checklist

Step 1 — Classify

Limited risk: content generation, interacts with users

Step 2 — Transparency

Disclose that the minutes were AI-assisted. Mark the output as AI-generated.

Step 3 — Human oversight

The meeting chair or secretary reviews the minutes before circulation, and owns the output once approved.

Step 4 — Risks

Hallucination: verify against the recording. Data leakage: keep confidential points out of public tools. Over-reliance: human sign-off before circulation.

This is how you do it, for every agent, every time.

Data protection (GDPR)

- Lawful basis for processing any personal data
- Data minimisation — only input what's needed
- No personal data in public tools (ChatGPT, Gemini free tier)
- Know where your data goes — organisational controls vs public tools
- DPIA may be needed for higher-risk processing
- Retention — know how long data is held by the provider

Penalties (for awareness)

**€15M
or 3%**

Maximum for Article 50 transparency breaches, whichever is higher

**€35M
or 7%**

Maximum for prohibited practices under Article 5, whichever is higher

The regulator's starting point is compliance support, not fines — but the stakes are real.

Working with sensitive data — a deeper look

The question: *"How do I use AI agents when I'm dealing with sensitive data?"*

The principle: Know what data you're putting in, and know where it goes.

Public tools = ChatGPT free / Gemini free. Organisational tools = M365 E3/E5, Enterprise plans.

Data type	Public tools	Organisational tools
Public information	Fine	Fine
Internal business data	Check provider's data policy	Generally covered by enterprise agreements
Personal data (GDPR)	Avoid	Only with lawful basis and DPIA
Special category data	Never	Only with explicit safeguards

Practical rules

- Default to anonymisation — strip names, IDs, financial figures before pasting into any tool
- Check your organisation's AI policy — many Irish public sector bodies now have one
- Enterprise ≠ safe by default — M365 E3/E5 has data processing agreements, but Copilot features may still log content
- When in doubt, leave it out — you can always describe the task without including the real data

The agent doesn't need real data to help you. It needs the structure of the task.

The future taster — live demos

You won't have these tomorrow. But this is where it's going, and I want you to see it in action.

Claude CoWork

Anthropic's agentic workspace

Hermes

Open-source autonomous agent

OpenClaw

Open-source personal AI

Watch what each one can do, and notice how the governance principles from today apply to all of them.

Live demo: Claude CoWork

Anthropic's agentic workspace — the agent in your workday

Watch as I

- Delegate a multi-step task across files and tools
- Show it working across calendar, email and documents
- Show a remote session — it keeps going after you walk away

What to notice

- It orchestrates multiple steps without hand-holding
- It works across your actual tools, not just a chat window
- It persists — you can close your laptop and it carries on

Live demo: Hermes

Open-source, self-hosted autonomous agent

Watch as I

- Show the agent running on its own infrastructure
- Demonstrate persistent memory — it remembers across sessions
- Show self-improving skills — it learns and refines over time

What to notice

- Data sovereignty — your data never leaves your infrastructure
- It's open source — no vendor lock-in, no per-seat licensing
- This is what Level 3, Autonomous, looks like in practice

Live demo: OpenClaw

Open-source personal AI on your own infrastructure

Watch as I

- Show a persistent personal agent in action
- Demonstrate 24/7 availability — it's always running
- Show tool integration — files, messaging, calendar, web

What to notice

- Complete data sovereignty — nothing leaves your machine
- It integrates with the tools you already use
- This is the direction of travel — personal agents that work alongside you all day

What the future means

You won't have these tomorrow. But this is where it's going.

- The future of work is agents that work with you across your day
- Not just chat windows — persistent, capable, tool-using assistants
- The governance skills you're learning today apply to all of it

Classify them. Be transparent. Keep a human in the loop. The principles don't change — the tools get more powerful.

From today to tomorrow — bridging the gap

You've seen: single-use agents, Office Agent, governance

The reality for most of you: you're starting from scratch, maybe feeling behind

That's normal — and that's fine.

You need

- Not to build everything
- Not expensive tools
- One agent, one task, verified

The attendee pack has

- The single-use agent template — name, description, instructions, context
- The governance checklist — classify, transparency, oversight, risks
- Platform links for whatever you have access to

Start small. Verify everything. Build from there.

Your action plan — do this tomorrow

- 1** Pick one task you do repeatedly
- 2** Build one single-use agent (Custom GPT or Gemini Gem) for it — use the template from your attendee pack
- 3** Try Office Agent to create one document from a prompt
- 4** Verify the output before you use it
- 5** Run it through the governance checklist: classify, check transparency, check oversight, check risks

You don't need Copilot Studio or a big budget. You have the tools. Start with one agent, one task, verify, and build from there.

Q&A

Questions? Key resources from today:

- Attendee Pack — concepts, links, glossary
- Exercise Pack — templates for take-home exercises
- Compliance Checklist — EU AI Act working tool

Thank you Enda Rochford

CMG · September 2026 — “AI advises, humans decide.”

Quick self-assessment (take away)

Answer these five questions. If you answer “no” to any, you have an action item.

- 1** Do people know when they're interacting with AI? (Article 50)
- 2** Is AI-generated content marked as such?
- 3** Is there a human who reviews and owns agent outputs?
- 4** Have you assessed the four key risks for each agent?
- 5** Do your staff have sufficient AI literacy?

Platform links

Platform	What it's for	Link
Microsoft Office Agent	Create PowerPoint/Word from a prompt	officeagent.microsoft.com
Custom GPTs	Build a single-use agent	chatgpt.com
Gemini Gems	Build a single-use agent	gemini.google.com
n8n	No-code automation	n8n.io
Claude CoWork	Future — agentic workspace	claude.ai
Hermes	Future — open-source autonomous agent	hermes-agent.org
OpenClaw	Future — open-source personal AI	openclaw.ai