

AI Agents for Construction

From repeatable workflow to controlled pilot

A practical guide to where agents fit, where they do not, and how to start safely.



What an AI Agent Actually Is

Three distinct levels. Each workflow sits on one — not the whole firm.



Chat

Workflow

Agent

📌 **Bottom line:** An agent is not magic. It is a controlled workflow upgrade — with boundaries, not autonomy.

Why "Agent" Gets Overused

✗ NOT THIS

- A chatbot with a system prompt
- A well-crafted prompt template
- A report generator with fixed fields
- A single-step document summariser

✓ YES THIS

Real agents only make sense when the workflow is:

Repeatable

Structured


Checkable

Connected

- Do not jump to agents before the workflow is stable.


Where Agents Actually Help

Four categories where construction teams see repeatable, reviewable value.



Repetitive Reporting

Weekly update packs,
progress summaries,
executive dashboards




Document Extraction

Tender review summaries,
specification gap analysis,
clause extraction



Fixed-Format Drafting

QA and compliance packs,
method statements,
inspection notices



Multi-Step Preparation

Issue tracking, action routing,
risk register updates

 **Rule of thumb:** Agents work best where inputs are known, format is repeatable, and review stays human-led.

Where Agents Should Not Be Trusted Alone

Final Contractual Judgement

Agents cannot hold professional accountability. A qualified person must own the decision.

Unreviewed Commercial Recommendations

Risk and pricing decisions require commercial context an agent cannot fully assess.

Safety-Critical Decisions

No AI output should be the final word on anything affecting life safety or structural integrity.

Black-Box Outputs

If you cannot see the reasoning or cite the source, it is not enterprise-ready.

Incomplete or Unverified Source Material

Garbage in, garbage out — with professional liability attached.



Bottom line: Agents can support professional judgement. They do not replace accountability.

Agents Are Already Shipping in Your Tools

The question is no longer whether to build one. It is which ones you allow, on what data, and who signs off.

PROCORE • JUL 2026
01
Starter agents
Deep Search, RFI, Submittal Review, Daily Log, Contract Review
02
New specialist agents
Site Safety, Schedule Analyst, Change Analysis
03
Skills
Teach it your own procedures (rolling out from August)
04
Control Tower
Admin view of AI use by agent, project and person
05
Early results
Haskell: submittal review from 7 days to 10 minutes (vendor case)

MICROSOFT 365
01
Model choice
Claude models in Word, Excel, PowerPoint and Copilot Chat
02
Researcher
Research across the web and company files, emails and chats
03
Copilot Cowork
Longer, multi-file tasks across tools
04
Copilot Studio
Build your own agents on company data
05
Admin control
IT decides which models and agents are switched on

The Agents Are the Easy Part

Five boxes. The two gates carry the load — and they are the two nobody draws.

Contracts / project lead • a person who sets scope, standards and outputs

INPUT GATE Nothing runs unless the data has provenance • a classification • an owner

Coordinator agent • routes work, summarises results

Queries & RFIs
technical query triage

Commissioning
tests, hold points, certificates

Commercial
tenders, change, valuations

Reporting
weekly status packs

SIGN-OFF GATE A named human • an audit trail • a data boundary. Nothing leaves without it.

Vendors sell the grey middle and the admin switches. Your provenance rules and named signatory are yours to build.



Net: anybody can draw the lanes. The two gates decide whether it is usable on a live project.

From Workflow to Agent Upgrade

The logic stays the same. The system becomes more timely, consistent, and scalable.

A WORKFLOW ALREADY HAS

01

Role

Who does this task

02

Inputs

Defined source documents and data

03

Steps

Sequence of actions to complete

04

Checks

Quality gates and review points

05

Output

Defined format and destination

AN AGENT ADDS

01

Trigger

Time-based, event-based, or manual start

02

Memory

Persistent context across runs

03

Tool Access

Live files, systems, and project data

04

Routing

Automated handoffs to the right people

05

Approval Boundary

Defined gates before output is used

Example: Weekly Reporting Agent

BEFORE — MANUAL

A PM manually pulls together:

- Programme extract
- Progress notes from site
- Issues log
- Action list
- Drafts the weekly summary manually

AFTER — AGENT-ASSISTED

Every Thursday 4pm, the agent:

- Reads approved source materials
- Compares with previous week
- Drafts executive summary
- Flags missing information
- Prepares risks and actions

Separates facts from assumptions

Cites source sections

PM sign-off required before
circulation

Example: Tender Review Agent

WORKFLOW

01

Reads tender documents

Full package ingested against defined scope

02

Extracts key commercial and scope points

Structured into standard review format

03

Identifies clarifications and exclusions

Flags gaps that require a question or caveat

04

Drafts review note

Standard template, ready for commercial team

CONTROLS

Data terms agreed in writing first

Sets no position, approves no risk

Flags ambiguity — never guesses

Human commercial review required



Strong first pilot — once the data questions are answered in writing.

Example: Commissioning & Handover Agent

INPUTS

Commissioning tracker • IST test sheets • snag list • O&M and certificate register

WHAT THE AGENT PRODUCES

REVIEW

Contracts manager or commissioning manager approval required before distribution.

📄 The agent structures and routes information. It does not replace technical judgement.

→ Open items by system, prioritised

→ Tests at risk next week

→ O&M manual and certificate gaps

→ Missing information flags

What's Realistic Now, Later, and Never


Now — 2 to 6 Weeks	Later — 3 to 9 Months	Never
One narrow workflow with a clear owner	Better system integration and automated triggers	Fully autonomous contractual decisions
Standard templates and defined prompts	Shared project context across the team	Final approvals without human review
Human-in-the-loop review at every gate	Metrics, improvement cycles, and workflow library	Black-box outputs with no traceability
Sanitised documents only	Automated routing and handoff logic	Uncontrolled access to live project systems
Simple, measurable success metric	Reusable workflows across projects	Scaling before the pilot is trusted



Start narrow. Prove value. Keep accountability with people.

Sign-off Gate — Connectors, MCP, and Approval Gates

The moment an agent touches live systems, governance is no longer optional.



WHY THIS MATTERS

- Tool access increases usefulness — and risk
- System access changes the governance requirement immediately

RULES TO START SAFELY

Least privilege by default

Start read-only

No send / publish without approval

Credentials controlled centrally

Input Gate — Data Hygiene, Audit, and Traceability



GDPR / Data Hygiene

- Minimise personal and sensitive data in inputs
- Use sanitised examples first — always
- Do not paste confidential disputes or uncontrolled pricing into casual tools
- Public-sector and client work requires stricter discipline, not looser



Audit and Retention

- Know what was input, what was generated, what sources were used
- Know who approved it and where the record is stored
- Define retention and deletion rules before deployment



Practical Rule

If the workflow matters, it must be traceable.

A useful agent without auditability is not enterprise-ready.

Final exercise: Pick 1 Pilot, Not 10

Design your pilot against these nine fields — then hold to them.

Workflow	One specific, named task
Owner	Named individual, not a team
Trigger	Time, event, or manual start
Inputs	Defined, sanitised documents only
Excluded data	What must never enter the agent
Output	Fixed format, defined destination
Checks / controls	Quality gates before use
Approval step	Who signs off, and how
Success metric	One measurable outcome
30-day scope	Narrow, testable, reviewable

RECOMMENDED FIRST PILOTS



Weekly Reporting Agent

Thursday trigger • CM sign-off • one project



Tender Review Agent

One package • data terms in writing first



Notices & Early Warnings

Minutes + change register • CM sign-off



The smartest first move is not the most ambitious one. It is the one your team can **trust, review, and improve**.