



CATHA PRACTICAL AI FOR CONSTRUCTION LIBRARY

What Is Under the Hood of Construction AI?

AI Architecture for Construction Leaders

Guide 12

Practical AI for Construction. Responsible. Outcome-driven.

Catha Construction Services | 2026 Edition

How to Use This Guide

Executives need a simple mental model for how business AI systems are assembled. This guide explains the common components without requiring engineering knowledge.

A construction AI system typically sits between people and systems such as project management, shared drives, email, CRM, estimating, accounting/job cost and field reporting. This guide explains that chain without requiring engineering knowledge.

Best use

Use this map when discussing a proposal so each component has a clear purpose, owner and control boundary.

The Basic Flow

User → Interface → Orchestrator → Model → Company Data/Tools → Controls → Action → Logs → Human Oversight

Every production AI system is some variation of this chain. Complexity should be added only where the business requirement needs it.

Core Components

Component	Plain-English meaning	Why it matters
LLM/model	The reasoning/generation engine	Capability, cost, speed and failure characteristics
Prompt/instructions	The job definition and rules	Shapes behavior but is not a security boundary
RAG/retrieval	Fetches relevant company knowledge	Keeps changing information grounded in sources
Structured data/API	Authoritative records and transactions	Needed for exact values and actions
Orchestrator/workflow engine	Coordinates steps, tools and state	Turns a model into a repeatable process
Agent	A system that can choose steps/tools within a goal	Adds flexibility and governance burden
Memory/state	Stores context across steps or time	Useful but can create privacy/staleness risk
Identity/permissions	Controls what the system may access/do	Primary authority boundary
Evaluation/monitoring	Tests and observes behavior	Shows whether the system is reliable
Logs/audit	Records inputs, decisions and actions	Enables investigation and accountability

Reference Architecture Questions

Check	Item	Notes
☐	Where does the user enter the workflow?	_____
☐	Which model(s) are used?	_____
☐	What data is retrieved?	_____
☐	Which system is authoritative?	_____
☐	Which tools can be called?	_____
☐	What identity is used?	_____
☐	Where is approval required?	_____
☐	What is logged?	_____
☐	What happens on failure?	_____
☐	How is the system stopped?	_____

Architecture Anti-Patterns

- Using prompts as the only control.
- Giving one agent broad access to every system.
- Allowing the model to “remember” authoritative facts instead of retrieving them.
- No separation between drafting and execution.
- No deterministic checks around irreversible actions.
- No observability beyond chat transcripts.

Construction Applications

The safest pattern separates retrieval, reasoning, approval and execution. Prompts are instructions, not security boundaries. Identity, permissions, deterministic checks, audit logs and authoritative systems remain essential.