

Beyond the AI Buzzwords


A practical guide to the enterprise AI lifecycle

Prepare the data, choose how the work runs, test the results, and compare changes.



01 GROUND What must the system know?

Useful context starts with current sources and agreed business definitions. Connect the relevant records and documents so the system can retrieve evidence that answers the question and that the user is allowed to access.



Cortex Search retrieves relevant text, while **Cortex Analyst with semantic views** supports structured questions using defined business meaning. Enforce permissions at retrieval and query time.

02 BUILD How should the work happen?

Define the outcome and what counts as done. Use the simplest design that can complete the task.

CHOOSE THE SHAPE OF THE WORK


Explicit rules sufficient?	Yes →	Code, SQL or deterministic workflow
No ↓		
One bounded model call sufficient?	Yes →	Focused model call or AI function
No ↓		
Can code prescribe the steps and branches?	Yes →	Workflow with AI in selected steps
No ↓		
Must the model choose the next action from intermediate results?	Yes →	Bounded agent loop Tools, budget, stopping conditions, escalation
No ↓		
Clarify the task.		

The distinction is who chooses the next step. A workflow follows code-defined branches, even when those branches depend on new evidence. An agent lets the model choose an action within enforced limits, and can itself run as one part of a larger workflow.

AI Functions handle bounded model tasks; **Cortex Agents** provides managed multi-step orchestration.

Inside an agent loop


The model uses the available context to propose an answer or action. The runtime checks that proposal before execution, then feeds the result back into context so the next decision reflects what happened.



04 IMPROVE

What should change, and did it help?

Use failed cases and traces to locate the cause before choosing a fix. An incorrect policy answer may come from an outdated source or missed retrieval, so first check the evidence the system used. That diagnosis determines whether to repair GROUND or change how BUILD handles the task.



Compare the proposed change with the current version on the same test cases, checking both the original failure and regressions elsewhere. Promote it only after reviewing the effect on quality, latency and cost, then observe real use again.

A run's state records what happened in that attempt; it is not automatically a lesson for future work. If the evidence suggests reusable guidance, review it before adding it to a skill or durable memory, test the resulting behavior, and remove guidance when it becomes stale.

Choosing the next use case

Start with five to ten real business problems, each with an owner, a desired outcome and a way to recognize success. Compare their business value with immediate feasibility before choosing a product or deciding how much autonomy to give the system.

