Agentic AI Design Canvas

DESIGN REVIEW

designtheagent.com
22 June 2026
Round 1

NORTH STAR

Settle simple storm-season claims in hours instead of days for customers in distress, while keeping every payout decision in human hands.

Overall

MINOR GAPS

This is a coherent, deliberately scoped design: autonomy held at prepare, the payout withheld from every agent, an audit agent over the pipeline, and a named human owning each decision. The biggest open question is that you measure speed everywhere but never measure whether the recommendation the human relies on is actually right.

COHERES The story holds together cleanly: a bounded workflow sets a prepare-level role, the money action is withheld, the audit agent and a named human close the loop. It comes apart only at the edges — what counts as a good recommendation is never measured, and the expansion plan does not yet say how the scope and oversight travel to new countries.

DO THESE NEXT

- **Target Workflow & Metrics** Add a quality metric to your Target Workflow cell (Cell 01), such as how often the human accepts the recommendation unchanged or how often it is later found correct.
- **Performance Needs** State a measurable accuracy standard in your Performance Needs cell (Cell 04) for how often the prepared recommendation must be right before a human relies on it.
- **Memory & Learning** State in your Memory & Learning cell (Cell 08) what each claim retains, for how long, and what must be deleted, especially personal and fraud-screening data.

TOP RISKS

- Every metric in your Target Workflow cell (Cell 01) is about speed or volume; nothing measures whether the prepared recommendation is correct, yet the human decision rests on it.
- Your Performance Needs cell (Cell 04) sets a five-minute speed bar but no measurable standard for how often the recommendation is right.
- Your Memory & Learning cell (Cell 08) leaves per-claim retention unspecified while the pipeline handles personal, claim and fraud data.

CROSS-CELL FLAGS

CELLS 01, 04

You measure speed and volume but set no bar for whether the recommendation is correct.

A human signs off on what the agent prepares; if no metric tracks recommendation quality, speed can rise while accuracy quietly slips and no one sees it.

CELLS 08, 07

Retention is unspecified in your Memory & Learning cell (Cell 08) and no rule in your Rules & Boundaries cell (Cell 07) covers what personal or fraud data is kept or deleted.

Holding customer and fraud-screening data with no stated expiry is a privacy and liability exposure, and it grows as you cross into new countries with different data laws.

THE CANVAS, CELL BY CELL

Cell 01 · Target Workflow & Success Metrics

NEEDS DETAIL

WHAT YOU WROTE

Workflow: food-spoilage claims under AUD 500 arising from severe-weather power outages (Allianz Australia). Metrics: end-to-end settlement time (days to hours), claims closed within a day, AI phase completed in minutes.

COACH

All three metrics measure speed or volume; none measures whether the recommendation handed to the human is correct.

How will you know the faster recommendations are also right ones, and not just quick?

Cell 02 · Users & Stakeholders

LOOKS CLEAR

WHAT YOU WROTE

Direct users: claims professionals who receive the assembled file and decide. Affected: storm-hit customers awaiting payment, claims operations, fraud teams, regulators, and reinsurers.

COACH

As you expand to new countries, which new regulators or privacy authorities join this list?

Cell 03 · AI Role & Autonomy

LOOKS CLEAR

WHAT YOU WROTE

Role: a seven-agent claims pipeline that verifies, screens, and calculates. Autonomy: prepare. It prepares the settlement recommendation; payout decisions are never automated, by design.

COACH

When you extend to motor and property claims, does prepare still hold, or will anyone push for the agent to decide?

Cell 04 · Performance Needs

NEEDS DETAIL

WHAT YOU WROTE

The AI phase completes in under five minutes. Every claim is corroborated against external weather data, and every step must be traceable enough to support an audit summary a human can act on.

COACH

You set a clear five-minute speed bar but no measurable standard for how often the recommendation is correct.

What pass-or-fail standard tells you a prepared file is good enough for a human to act on quickly?

Cell 05 · Context & Knowledge

LOOKS CLEAR

WHAT YOU WROTE

Policy records and coverage terms, the customer's claim and documentation, transaction patterns for fraud screening, and external meteorological data confirming the storm actually happened.

COACH

Who confirms the weather data source is the authoritative one, and what happens when it disagrees with the customer's account?

Cell 06 · Tools & Action Channels

LOOKS CLEAR

WHAT YOU WROTE

Reads policy and claims systems; writes a structured audit summary for the human decision-maker. The one action that touches money, paying the claim, is withheld from every agent.

COACH

Does the agent only read, or can it also update the claim record in the claims system?

Cell 07 · Rules & Boundaries

LOOKS CLEAR

WHAT YOU WROTE

Never automate a payout. Stay inside the bounded scope: food spoilage, severe-weather outages, under AUD 500. A dedicated cyber agent enforces data-security guardrails on every claim.

COACH

What must the agent do when it suspects fraud or hits a claim that sits just outside the bounded scope?

Cell 08 · Memory & Learning

NEEDS DETAIL

WHAT YOU WROTE

Learning happens by widening scope, not by silent drift: prove the bounded case, then extend to travel delay, simple motor, and property claims. Per-claim retention is not publicly specified.

COACH

Per-claim retention is unspecified while the pipeline handles personal, claim and fraud-pattern data.

What does each claim leave behind once it is settled, and what must be deleted?

Cell 09 · Oversight & Accountability

LOOKS CLEAR

WHAT YOU WROTE

An audit agent reviews every step the other six take and writes the summary; a claims professional makes every payment decision. Ultimate responsibility rests with that person, not the pipeline.

COACH

As you scale, does one named person still own outcomes across the pipeline, or only per individual claim?