

SAMPLE ARCHITECTURE REVIEW

# Rowan Portfolio Project

Planned portfolio assistant

EXECUTIVE VERDICT  
Read-only prototype only

|                               |                                         |                                    |                                                  |
|-------------------------------|-----------------------------------------|------------------------------------|--------------------------------------------------|
| REFERENCE<br>HA-AR-SAMPLE-001 | SAMPLE STATUS<br>Fictional sample       | ISSUE DATE<br>4 August 2026        | CLIENT<br>Rowan Portfolio Project                |
| SYSTEM STAGE<br>Early design  | REVIEW BASIS<br>AGREED<br>3 August 2026 | REVIEW TYPE<br>Architecture Review | SCOPE<br>Architecture and control<br>design only |

PREPARED BY HODL ADVISORY  
Architecture Reviews and Blueprints for consequential agent systems. [hodl.org](https://hodl.org)

This sample is based on a simulated engagement and contains no real-client material.

# The verdict and the Review boundary.

EXECUTIVE VERDICT

Read-only prototype only

The idea is concrete enough to prototype, but not in the action-capable form currently imagined. The client-confirmed design asks one AI conversation to combine portfolio interpretation, policy memory, recommendation, approval handling, action submission, outcome verification, and record keeping. The agreed Review basis does not define an authority for current state, an approved rule version, an exact approval record, a durable action identity, or independent confirmation of the outcome. Those are not details the AI can safely work out on its own.

The correct first move is a read-only evidence and recommendation prototype. The material does not define whether such recommendations may influence a human's manual provider entry; if they can, that entry is a consequential human-executed boundary, not a harmless read-only path. Keep portfolio actions manual until the team can demonstrate that inputs, rules, approval, action identity, provider outcome, reconciliation, and stopping behavior remain correct under interruption and disagreement.

REVIEW BOUNDARY

Conclusions are based on the agreed Review basis and material listed in this report. No source code, production access, provider verification, investment-strategy assessment, or compliance review was performed.

At a glance

| ITEM                       | REVIEW CONCLUSION                                                                                                     |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------|
| Client                     | Rowan Portfolio Project                                                                                               |
| Reviewed system            | One planned portfolio information, recommendation, approval, and later action workflow                                |
| Current state              | Notes, rough workflow, and tentative tools; nothing implemented                                                       |
| Immediate decision         | Whether to hire a developer and begin implementation                                                                  |
| Architecture Review answer | Begin only with a bounded read-only stage; do not authorize automated portfolio actions                               |
| Confidence                 | High for the documented design gaps; no conclusion about unselected APIs, vendors, code, or actual operating behavior |

## THE CLIENT'S QUESTIONS

## Direct answers before technical detail.

### 1. Is one AI agent enough for the first version?

One model-assisted analysis step can be enough for a read-only experiment. It is not enough as the system architecture. Even a single-model prototype needs separately defined input freshness, rule version, structured output, validation, human review, failure behavior, and an inspectable record. The number of agents is not the first decision; the control and evidence boundaries are.

### 2. Is the proposed workflow coherent and safe enough to start building?

It is coherent enough to build a read-only prototype, but not an action-capable system. The current plan has no supported answer for which state is authoritative, what a "yes" approves, how a sent action is identified, how partial or unknown outcomes are represented, or whether a retry could repeat an action.

### 3. What should be built first, and what should not be automated yet?

Build the evidence foundation, versioned rules, structured recommendation record, and human review path first. Keep broker action, automatic retry, automatic spreadsheet correction, and AI-controlled policy changes out of scope until the validation gates in the roadmap pass.

|         |        |          |         |       |
|---------|--------|----------|---------|-------|
| VERDICT | SYSTEM | FINDINGS | ROADMAP | BASIS |
|---------|--------|----------|---------|-------|

AGREED REVIEW BASIS

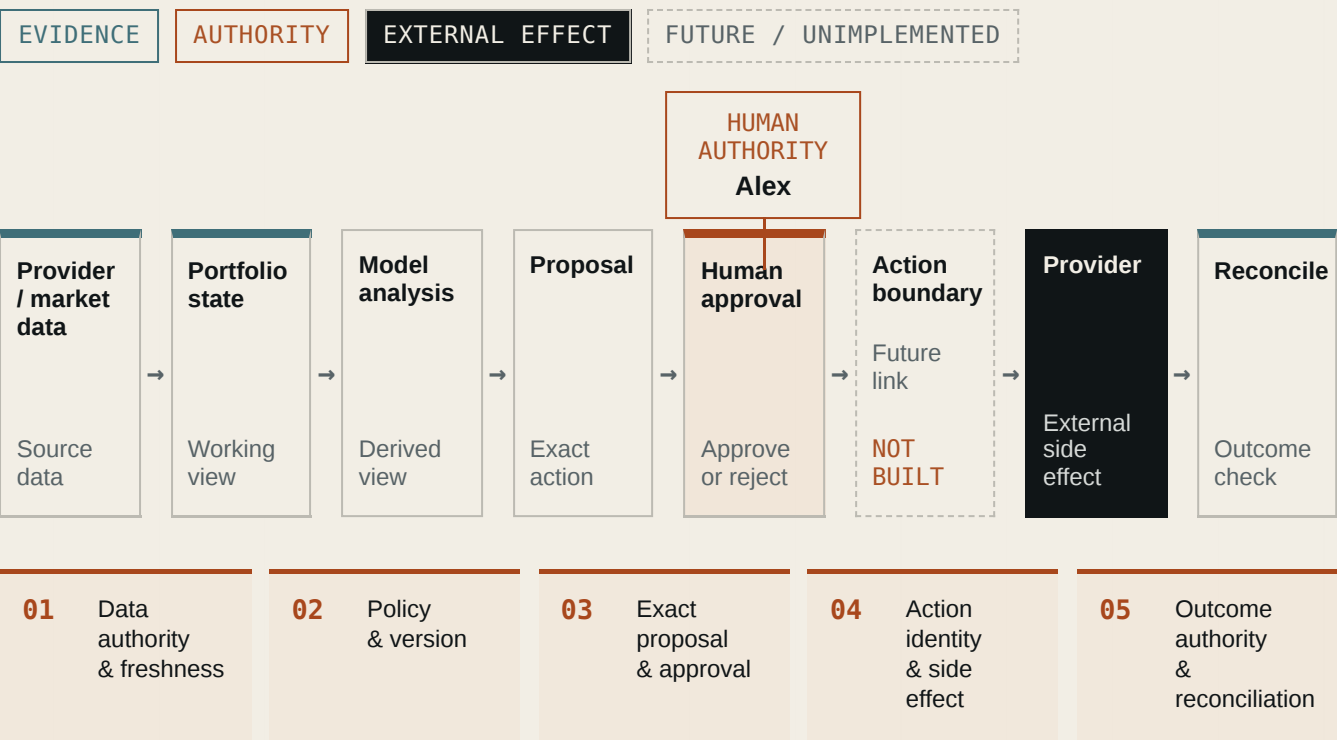
# The complete client-confirmed description used for this Review.

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <ul style="list-style-type: none"><li>• <b>Purpose:</b> decide whether to hire a developer and begin building the planned portfolio assistant or rethink the approach first.</li><li>• <b>Intended workflow:</b> provider exports and market prices feed a spreadsheet; an AI produces a summary and recommendation; Alex approves through Telegram; a later connector may submit an action; the AI is then expected to check the outcome and update the spreadsheet.</li><li>• <b>Client questions:</b> whether one AI is enough; whether the workflow is coherent enough to begin; and what to build first versus keep manual.</li><li>• <b>Rules and policy handling:</b> portfolio and risk rules live in an editable spreadsheet tab and notes. Alex would edit them; no version, effective time, approval record, or machine-readable validation exists.</li><li>• <b>Runtime and failure assumptions:</b> no orchestration platform or model has been selected. Coding assistants and a scheduled Python script were considered as possible runtime mechanisms. Workflow identity, stopping, resumption, timeout, retry, unresolved-state ownership, and concurrent limits are undecided.</li><li>• <b>Material unknowns:</b> the authorizing principal and operating limits; whether recommendations may influence manual provider entry; state authority and freshness; policy version; approval meaning and expiry; provider outcome evidence; credential and permission design; orchestration; and failure behavior.</li></ul> | <ul style="list-style-type: none"><li>• <b>Current state:</b> early design only; notes, a rough workflow, and tentative tools exist; nothing is implemented.</li><li>• <b>Human authority:</b> Alex approves portfolio actions and risk-limit changes. Transfers, withdrawals, unapproved actions, and autonomous policy changes are prohibited.</li><li>• <b>Portfolio and state sources:</b> provider accounts hold the real assets; provider exports and market prices feed a spreadsheet intended to become the combined view. No conflict rule, freshness rule, or outcome-state model has been chosen.</li><li>• <b>Approval, action, and outcome:</b> Alex expects to reply "yes" to the preceding Telegram message. A later connector may submit the action; the same AI is then expected to check the provider outcome, update the spreadsheet, and report completion. Exact action parameters, approval expiry, durable action identity, provider evidence, and partial or unknown outcomes are not defined.</li><li>• <b>Material reviewed:</b> the client's description of the intended workflow, goals, rules, planned tools and permissions, and the clarifications summarized above.</li><li>• <b>Acknowledgement:</b> Alex Rowan confirmed this summary as the basis for the Review on 3 August 2026.</li></ul> |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

This is a client-confirmed description of the planned design. It was not independently verified against a repository, source code, provider documentation, accounts, holdings, credentials, deployment, or actual operating behavior.



# The reviewed route and the points where identity, evidence, or authority changes hands.



PRIORITY

01

ARCHITECTURE  
FINDING

CRITICAL

# The described action path has no exact approval contract

**What is happening.** In the client-confirmed design, a plain Telegram “yes” is expected to authorize the preceding AI message, but approval is not bound to an immutable asset, direction, quantity, price condition, rule version, or expiry.

**Recommended change.** Define whether recommendations may influence manual provider entry. Create a durable proposal with a stable ID, exact parameters, mandate and limit reference, snapshot and policy versions, expiry, and explicit approve or reject response. Any material change must invalidate approval.

**Why it matters.** The message may change, the market may move, context may be lost, or the wrong preceding message may be interpreted as approved. A human can remain “in the loop” while still lacking control over the exact action.

**What to validate.** Name the authorizing principal and enforce current authority. Demonstrate that stale, altered, missing, duplicate, wrong-message, revoked, and concurrent approvals cannot cross either the manual or connector action boundary.

VERDICT

SYSTEM

FINDINGS

ROADMAP

BASIS

PRIORITY

02

ARCHITECTURE  
FINDING

CRITICAL

## No authoritative portfolio or action state was defined

**What is happening.** The client described provider accounts as holding the real assets and the spreadsheet as the intended combined view. No conflict rule was defined.

**Why it matters.** The AI can produce a confident recommendation or “done” message from a stale projection while provider evidence says something else.

**Recommended change.** Define, for each material outcome, which provider evidence is authoritative and why, including its scope, limitations, relationship to the provider boundary, and finality. Treat the spreadsheet as a versioned projection or operator input. Record freshness and lineage, preserve prior reliance and explicit supersession, and surface unresolved conflicts instead of overwriting them.

**What to validate.** Alex chooses the business meaning; the implementer enforces the hierarchy. A conflicting provider export and spreadsheet view must produce a visible unresolved state, not silent selection.

PRIORITY

03

ARCHITECTURE  
FINDING

HIGH

Rules were described as editable prompt material, not governed policy

**What is happening.** Portfolio and risk rules were described as an editable spreadsheet tab and notes. No version, effective time, approval record, or machine-readable validation was defined.

**Why it matters.** The system cannot prove which rule set produced a recommendation or whether a later edit changed the meaning of an already-approved action.

**Recommended change.** Separate approved constraints from conversational instructions. Store a versioned policy record with an owner, effective time, change record, and validation. Include the policy version in every recommendation and approval record.

**What to validate.** Alex owns rule approval. Every recommendation must trace to one immutable approved version; an invalid, missing, changed, or not-yet-effective version must stop the run.

VERDICT

SYSTEM

FINDINGS

ROADMAP

BASIS

PRIORITY

04

ARCHITECTURE  
FINDING

HIGH

# Outcome verification was assigned to the same AI without defined closure evidence

**What is happening.** The client-confirmed plan assigns the same AI to send an approved action, “check what happened,” update the spreadsheet, and report completion. No independent provider record, lifecycle, or representation of partial and unknown outcomes was defined.

**Recommended change.** Require structured provider records and an explicit lifecycle such as proposed, approved, sending, accepted, partial, completed, rejected, cancelled, and unknown. Reconcile the internal projection against provider evidence before reporting closure.

**Why it matters.** A model summary is not evidence that the external action completed. Timeout, rejection, partial completion, delayed provider state, or repeated action can be collapsed into a plausible narrative.

**What to validate.** Assign the meaning of closure and an owner to every unmatched or unresolved outcome. Record what remains at risk, who owns it, when it must be reviewed or escalated, and what evidence closes it. Failure tests must retain rejected, delayed, conflicting, provisional, reversed, unknown, and partial states—never turn absence of error into completion.

PRIORITY

05

ARCHITECTURE  
FINDING

HIGH

No orchestration or failure boundary was selected

**What is happening.** The client-confirmed design does not select an orchestration framework or custom control service and treats coding assistants, models, and a scheduled script as unresolved runtime possibilities.

**Recommended change.** Define required capabilities before selecting a product: durable workflow identity and states; bounded retries; boundary-specific pause and stop; structured tool calls; approval waits; audit records; and aggregate and concurrent limits. Define what happens when approvals, policy changes, stops, and concurrent actions collide. Identify infrastructure whose failure could disable both execution and its safeguards.

**Why it matters.** Scheduling a prompt does not define durable state, retries, approval waits, timeouts, stopping, resumption, or records. Those gaps become consequential once an external action is possible.

**What to validate.** The implementer proposes the substrate; Alex approves permissions and the operating boundary. Interruption and restart tests must resume from durable state without repeating or skipping an authorized step. Failure tests must show that safeguards remain effective—or leave any continuing risk visible and assigned—when common infrastructure fails.

ROLES AND OPERATING BOUNDARY

# Keep decision authority explicit while the system is designed and prototyped.

## Roles, permissions, and handoffs

| ROLE                           | ALLOWED RESPONSIBILITY                                                          | MUST NOT SILENTLY OWN                                        |
|--------------------------------|---------------------------------------------------------------------------------|--------------------------------------------------------------|
| Alex                           | Approve rules and exact portfolio actions; stop the workflow                    | Implementation retries, implicit data-conflict resolution    |
| Co-founder                     | Review information and decisions as agreed                                      | Unrecorded approval authority                                |
| Developer                      | Implement and operate the technical workflow within scope                       | Define portfolio policy or approve actions                   |
| Model-assisted analysis        | Summarize supplied evidence and produce bounded structured recommendations      | Change rules, approve actions, establish external completion |
| Orchestration or control layer | Persist workflow state, enforce transitions, wait for approval, record outcomes | Investment judgment or human authority                       |

The reviewed material does not establish the authorizing principal, current authority check, credential storage, separation of powers, delegation, security architecture, legal or compliance requirements, or provider permission semantics. Those remain unknown rather than controls established by this Review.


Before model-assisted recommendations may influence a manual provider entry or later connector, the team must identify the material model, prompt, policy, context, memory, configuration, and tool states. Untrusted spreadsheet, market, Telegram, or other contextual content must not change authority or permissions. Any later tool or permission expansion is consequential, and model messages remain derived—not authoritative. No implementation or test evidence for these controls was reviewed.

### Do not automate yet

- broker order submission;
- retry across a possible broker-side effect;
- transfer or withdrawal under any condition;
- risk-limit or policy changes;
- interpreting conversational “yes” as an executable instruction;
- overwriting provider evidence with spreadsheet or AI state;
- declaring action completion from model output or absence of an error;
- learning preferences from conversation and silently promoting them to policy.

DEPENDENCY-ORDERED ROADMAP

Sequence the work by evidence, not by feature appetite.



VERDICT

SYSTEM

FINDINGS

ROADMAP

BASIS

ASSUMPTIONS AND DECISION GATE

## Unknowns stay visible; the next decision stays bounded.

### Assumptions and unknowns

This Review assumes only that the confirmed Review basis, workflow, and notes accurately describe the current plan. It does not assume any provider supports a particular API, duplicate-request behavior, permission scope, or event model.

**Data and provider:** data quality, frequency, provider lifecycle, evidence finality, and duplicate-request behavior.

**Authority and limits:** authorizing principal, mandate lifecycle, delegation, value and aggregate limits, and separation of powers.

**Runtime and security:** credential design, model and provider selection, orchestration, model/prompt/context/memory/configuration/tool states, and incident response.

**Operating context:** expected volume, implementation budget, jurisdiction, and legal or regulatory obligations.

**Failure behavior:** concurrency, interruption, shared failures, unresolved-state ownership, ageing, escalation, and closure.

These unknowns do not change the immediate decision: the planned action path lacks the authority, identity, closure, containment, and limit contracts required before execution.

### RECOMMENDED NEXT DECISION

Authorize a read-only prototype only if its scope includes Stages 0–3, explicitly excludes action execution, and states whether its recommendations may influence a manual trade. Do not treat a manual provider entry as outside the control boundary. Ask the developer to propose the orchestration or control layer against the required capabilities—not merely name an AI model or coding assistant.

TERMS USED IN THIS REVIEW

# Plain-language definitions for the decisions and controls discussed in the report.

- Authoritative state** — the evidence source that wins when two views disagree.
- Orchestration or control layer** — the framework or custom service that persists workflow state, sequences work, waits for approval, handles interruption, and records outcomes. It is distinct from an AI model or coding assistant.
- Structured approval** — a durable human decision tied to one exact proposal and expiry, not a free-floating chat response.
- Reconciliation** — comparing internal records with external provider evidence and explicitly resolving or preserving disagreements.
- Closure evidence** — the record proving what happened, including rejected, partial, unknown, or completed outcomes.

METHOD NOTE

Hodl Advisory uses the open [PROOF Standard](#) as a control lens when reviewing consequential agent systems. The Review remains an architecture assessment of the client system, not a separate PROOF conformity assessment.

ABOUT HODL ADVISORY

## Hodl Advisory

Architecture Reviews and Blueprints for consequential agent systems.  
See the whole system. Decide what to build next.

[hodl.org](https://hodl.org)