

FIRM CAPACITY BLUEPRINT

Delivered decision study

Cedar Vale Advisory, its people, cases, records and amounts are fictional source material used to show what a completed Firm Capacity Blueprint contains.

This is the design for one revenue path, the fixed-scope market assessment, from agreed scope to approved billing record, with reuse decided separately. It holds three connected designs, shared controls, an effort and capacity model, cash scenarios, a comparison plan and the decision each owner is asked to record.



Completed design example; decision review unfilled. None of the twenty proposed installation tests has run, and no professional approval, client acceptance or finance decision is represented as having occurred.

18

PEOPLE

40

ANNUAL ASSESSMENTS

4

UNSIGNED PIPELINE CASES

Where the firm stands

Eighteen people deliver 40 fixed-scope market assessments a year alongside other advisory work.

A case carries an A\$30,000 fee and about A\$20,000 of cost, including the normal labour of your own staff, leaving a modelled A\$10,000 contribution once collected. Six further qualified cases are in the annual pipeline, four of them expected within nine months. None is signed.

Elaine holds an unapproved option to hire a senior reviewer at about A\$120,000 a year. Drafting seats already exist. Ruth approves invoices in the existing finance system; an approved invoice is not cash received.

Contribution from an added case arrives only after quality passes review, capacity is actually freed, demand signs, work is accepted and billed, and cash is collected. These designs address the evidence and handoff conditions; their effect has not been measured. No rate is attached to any link.

What hurts today

Three of six sampled cases returned from review. Two of those also had permission gaps.

CV02, CV03 and CV05 returned. CV02 and CV05 had permission gaps. The overlap is visible; causation is not, and six cases cannot settle it. Marcus describes missing evidence and overstatement in corrections. Nadia collects notes manually and checks which version is current before acting.

A report has layers a reader cannot see from the prose: the findings, the sources behind them, the rights attached to each source, and the acceptance that binds a particular version. When any layer is missing, the next person reconstructs it or proceeds without it. Better writing does not restore a layer.

The commercial risk is that a polished draft reaches Marcus with the same gap underneath, or that a new reviewer absorbs checking that preparation should have done. Neither is measured yet.

3 / 6

returned from professional review

2 / 6

carried permission gaps

Small supplied sample. Overlap is visible; causation is unverified.

S01, S03


How a case moves now

Eight stages, several records, and a return loop whose cost is not recorded separately.

A return sends the case back to preparation and Marcus re-accepts. Where that labour is booked is unknown, so the hour means in the challenges section may already contain it. The constraint is therefore not established: the sample fits Marcus's view that returns concern evidence and overstatement, and equally fits ordinary correction on a small sample.

- Agreed enquiry and scope: Elaine to Nadia; email and project spreadsheet assumed, unconfirmed
- Evidence permission and brief: Nadia, with Tomas for access; case folder and permission entries

- Research and preparation: analysts in Word; case folder
- Professional review: Marcus in Word comments; returns reopen preparation
- Client acceptance: noted informally; email and spreadsheet assumed, unconfirmed; actor, date and version not consistently recorded
- Approved billing record: Tomas prepares, Ruth approves in the finance system
- Collection: finance; no collection data supplied
- Permitted reuse: analysts reuse prior structure; rights checked inconsistently




Review returns to preparation. Collection and reuse are distinct; see the full map on the next page.

S02, S03, S06

Eight stages. Separate decisions.

The map shows where work changes hands and where it can return.



05 acceptance follows 04 review. 06 billing → 07 collection. 08 reuse requires its own right and owner.

Review returns to preparation. Collection follows billing. Reuse is a separate permission branch, not an automatic last step.

CONTENT-v3 b04; S02, S03, S06; finance-handoff.md

What the records cannot yet tell you

The six-case means are 18 senior, 42 analyst and 6 operations hours per case, with no stage split.

Without a split, nobody can say which role binds or how much of the senior time is correction. The acceptance-to-invoice lags of 2, 8, 3, 11, 4 and 7 days average 5.8, but S01 has no completed-status column and two cases conflict with S06. CV06's seven-day lag implies acceptance and invoice events while S06 shows review pending and no acceptance. CV05's four-day lag rests on an acceptance with no actor, date or version and an unapproved finance record. The lag is not a baseline, and it never measured collection.

CV06's evidence is also short. Four interview excerpts were supplied against a method of eight. S05 records INTO4 as permitted for the current study, but S06 has no registry row for it, so the manifest must be reconciled before that excerpt can be used. OLD-CLIENT-X is denied for this project. No external processing right is recorded for any source, and the retention schedule is still to be agreed.

These are gaps in the record, not proof that interviews, permissions or external events never happened. They must be reconciled before the case or the baseline can be used.

Case	Review return	Permission gap	Invoice lag*
CV01	No	No	2 days
CV02	Yes	Yes	8 days
CV03	Yes	No	3 days
CV04	No	No	11 days
CV05	Yes	Yes	4 days
CV06	No	No	7 days

* S01 recorded fields; CV05/CV06 acceptance provenance conflicts remain unresolved. No causality or cash collection inferred.

S01, S02, S05, S06

Three commercial jobs. One connected path.

The buyer purchases a way to produce and control valuable client work. Evidence checks and handoffs sit inside those jobs.

Commercial job	The output you can inspect	What must be established
Scope work the firm can deliver	Costed options, proposal draft, signed-term comparison and accepted delivery brief.	Actual enquiry, cost method, fee assumptions, available reviewer and signatory.
Produce the expert decision pack	Analysis workpapers, source-linked findings, alternatives, recommendation draft and reviewed report.	Permitted complete evidence, approved method, professional reviewer and quality/effort comparison.
Reconcile work through finance	Acceptance/milestone packet, invoice exceptions, payment allocation and unresolved-balance queue.	Signed basis, acceptance record, invoice/payment evidence and authorised finance owner.

Three complete designs on one revenue path. First installation separately scoped; evidence supports the delivery problem most strongly.

Founder-approved v6 design direction; website workflows and current offer at main44a2b777

Make the commercial promise deliverable.

Defined responsibilities prepare the work; named people decide what may proceed.

Qualify the enquiry	Cost the options	Compare signed terms	Accept delivery
Intake analyst prepares the decision question, use, evidence, exclusions and missing answers.	Commercial analyst prepares → approved cost/fee assumptions; writer drafts the proposal.	Quality checker names changes to → scope, dates, evidence and obligations.	Case coordinator prepares the → brief and named reviewer/capacity checks.
Elaine qualifies or declines.	Ruth checks cost; Elaine decides fee.	Authorised signatory resolves terms.	Nadia and Marcus accept the exact brief.

Return to the commercial decision owner

Missing cost stays unknown. Changed terms or unavailable reviewer reopen the promise before mobilisation.

Simpler option: scope template, required fields and deterministic cost calculator. Underpricing is a hypothesis, not an established Cedar Vale fact.

Produce the work your client pays for.

Defined responsibilities prepare the work; named people decide what may proceed.

Assemble evidence	Prepare analysis	Challenge the draft	Approve exact version
Evidence librarian returns permitted passages with exact locators and versions.	Delivery analyst applies the → approved method; writer drafts findings, alternatives and report.	Quality checker compares claims → with original sources and the decision question.	Coordinator records unresolved → items, revised version and issue decision.
Tomas checks rights; Nadia checks coverage.	Marcus owns professional adequacy.	Marcus returns unsupported wording.	Professional issue and client acceptance separate.

Return to analysis with the source and reason

Contradiction stays visible. A material edit invalidates approval. No source or model score approves the professional recommendation.

First installation hypothesis. Compare full preparation, review, correction and operation with the current and repaired manual methods.

Know what is accepted, invoiced and paid.

Defined responsibilities prepare the work; named people decide what may proceed.

Match acceptance	Prepare finance basis	Check exceptions	Reconcile payment
Coordinator joins issued version, actual client response and signed milestone.	Commercial analyst - finance → reconciles amount, condition and invoice evidence.	Quality checker flags duplicate → IDs, wrong versions, mismatches and credits.	Finance responsibility links actual → allocations and names the open balance.
Account/delivery owner resolves qualification.	Ruth authorises finance actions.	Owner resolves; model cannot issue invoice.	Ruth owns collection, dispute and write-off.

An invoice is not collected cash

Absent acceptance holds the packet. Partial payment leaves the exact balance open; a dispute returns to its accountable owner.

Prefer existing finance reports and deterministic matching. Vendor/API capabilities and any collection problem remain unverified.

Who carries the work?

Map every role to an output and a decision owner.

Website responsibility	Work it may carry	Decision and evidence boundary
Intake analyst	Decision question, evidence gaps, scope assumptions.	Elaine/Nadia resolve; no autonomous commitment.
Evidence librarian	Case-scoped passages, locators, versions, rights and conflicts.	Tomas/Nadia check rights/coverage; not unrestricted cross-client search.
Public researcher	Dated brief from an approved sanitised public question.	Separate public route; private case notes excluded. Not executed in this specimen.
Delivery analyst	Reproducible analysis, observations, interpretations and alternatives.	Marcus decides findings; AI preparation is not professional approval.
Deliverable writer	Proposal, report and executive summary from approved basis.	Signatory/Marcus approve issue; drafting does not authorise a send.

Candidate responsibilities inside the system. None is a purchased seat or proof of a deployed agent. See buildnbloom.io/roles.

v6 proposed design; fictional buyer, not installed capability

Who checks, coordinates and maintains it?

The same role name does not confer the same access or authority.

Website responsibility	Work it may carry	Decision and evidence boundary
Quality checker	Discrepancy, citation and omission report against original evidence.	Independent check is advisory; reviewer owns adequacy.
Case coordinator	Owned exception queue, dependencies, reminders and return paths.	Named operations/account owner decides actions.
Commercial analyst	Sales: scope/fee workpaper. Finance: milestone/payment reconciliation.	Separate identities/source access. Elaine/Ruth approve commercial and finance actions.
Knowledge curator	Eligible method asset, rights, exclusions, version and retirement proposal.	Rights and method owners approve. Redaction alone grants no reuse right.
Technical investigator	Authorised log-based hypotheses, retest and recovery proposals.	Technology/risk owner approves restore/rollback. Not demonstrated by record-shape checks.

Candidate responsibilities inside the system. None is a purchased seat or proof of a deployed agent. See buildnbloom.io/roles.

v6 proposed design; fictional buyer, not installed capability

The work hidden behind the draft

This is an illustrated working situation built from the CV06 files, not a recorded live case or a scheduled date.

Picture the day before Marcus reviews the supplier-choice report. The client asked for one sentence: "Most customers will switch suppliers for faster delivery." To decide whether it may go out under the firm's name, Marcus needs three things beside it: where each supporting statement comes from, what in the interviews cuts against it, and whether the firm may use each source for this study.

Today the folder does not carry those facts. The scope agreement is unavailable, so question, segment and geography are unconfirmed. Four interview excerpts are supplied against a method that calls for eight; the full set is not available in the supplied record, which does not show the other interviews never happened. INT04 is described as permitted in S05 but has no registry row in S06, so its rights are unresolved. OLD-CLIENT-X is denied for this project.

Professional review is the valuable step. The repeated burden is the reconstruction before it: Nadia collecting notes, finding the current version and checking rights source by source. That burden is not measured. A useful review basis is simply every finding, its exact reference, its contrary evidence and its rights status in one place Marcus can read.

Requested claim	Exact source locator	Contrary account	Rights status
What Marcus needs beside each finding.			

Fable supplement S1; S03, S05, S06; hypothetical next-review scene, not a recorded event.

How the designs help your existing tools

Keep the tools people already use. Add the evidence and handoff records they need.

The records name Word, case folders, existing drafting seats and an existing finance system. The drafting and finance vendors are unnamed. SharePoint, email and a project spreadsheet are planning assumptions to confirm with Tomas; no tenant has been inspected.

W1 puts the source inventory, rights and agreed scope in the case folder. W2 keeps each finding beside its source, contrary evidence and review decision in Word. W3 records who accepted which report version and prepares an internal finance item for Ruth. This gives the next person a record to check instead of reconstructing the handoff.

Tomas checks actual features, licences, access, retention and processing rights before configuration. API compatibility and extra costs remain unverified. The manual route uses authorised existing tools; no platform move is required. Professional approval, invoice approval and any reuse decision remain separate human responsibilities.

Tool / evidence	Proposed aid	Human check
Word — recorded	Finding + citation + contrary evidence	Marcus reviews finding
Case folders — recorded	Manifest; separate versions	Tomas checks rights
Drafting seats — recorded	Compare against repaired manual method	Whole-job comparison
Finance — vendor unknown	Acceptance-linked internal item	Ruth approves invoice
SharePoint — assumption	Metadata / version controls	Confirm features and licences
Email — assumption	Actor / date / accepted version	Confirm acceptance route
Project spreadsheet — assumption	Case / handoff register	Confirm existing tool and owner

Recorded S02/S03 versus authored assumptions. No tenant inspected or API compatibility verified.

Recorded: S02, S03. Planning assumptions: SharePoint, email, project spreadsheet, unconfirmed; no tenant inspection.

Investigate the expert decision pack first.

The supplied fictional evidence supports the delivery/review problem most strongly. Scope and finance remain connected opportunities to validate.

Design all three complete jobs. Rank the expert decision pack for first-installation investigation; the installation itself needs a separately agreed scope, quote, provider route, comparison and acceptance. The scope/fee and payment designs require additional buyer evidence before a claimed economic benefit.

Use existing templates, source registers, review sheets and finance reports as the comparison. A simpler repaired method can win. Elaine retains the hire decision; no expanded AI design proves an avoided hire.

The following W1-W3 pages preserve the existing manual controls and record interfaces inside the wider jobs. W1 supports scope/readiness; W2 supports expert delivery; W3 supports the finance handoff and separate reuse branch.

CV06 remains blocked. New CV07 is a separate synthetic teaching episode with its own execution and fidelity checks; neither is professional or installation acceptance.

Re-estimate costs, participation and dates for the expanded designs. Existing money curves are preserved as prior conditional illustrations.

Design decision; confidence moderate; reversal conditions under Alternatives


The CV06 claim under the review rule

Requested client claim: "Most customers will switch suppliers for faster delivery."

Supplied material: four interview excerpts of the eight the method requires. INTO1 wants faster parts but names its own approval paperwork as the delay. INTO2 asks for compatibility documentation. INTO3 would not change supplier for speed. INTO4, a distributor, won its last five orders on fit and existing relationship; S05 records it as permitted for the current study, but S06 has no registry row, so it sits quarantined until Tomas reconciles the manifest. There is no denominator and no count of anyone switching.

The review rule returns the claim as an unsupported generalisation and offers a narrower candidate for Marcus to approve, amend or reject, built only on the permitted excerpts: in the supplied excerpts, one buyer prioritised faster parts but identified its own paperwork as the delay, and one would not change supplier for speed. INTO4 is noted as a quarantined record and forms no part of the finding until its registry entry is confirmed. The candidate does not estimate market-wide switching.

The case itself is blocked. OLD-CLIENT-X is denied for this project. S01 implies an invoice while S06 shows review pending and no acceptance. Tomas reconciles dates, versions and authority with Marcus and Ruth before any finance handoff. The gaps are recorded, not filled.



CV06 supplied: 4 excerpts of an expected 8-interview pack. INTO4 stays quarantined until rights records reconcile.

S05, S06; working-assets/CV06-review-rule.md and prepared-report-specimen.md

The CV06 review packet

From a requested market claim to source-linked findings and an owned next action.

REQUESTED CLAIM

Most customers will switch suppliers for faster delivery.

INT01 • S05 L5

Faster parts are desirable; the respondent's own approval paperwork delays its order.

INT02 • S05 L6

A local supplier is considered if compatibility is documented; substitution requires that evidence.

INT03 • S05 L7

Current lead times usually suffice; the respondent would not switch merely for speed.

INT04 • QUARANTINED

S05 permission conflicts with an absent S06 registry row. Tomas reconciles before inclusion.

DRAFT • HUMAN-AUTHORED • NO AI RUN

F01 • Approval and compatibility

INT01 identifies internal approval paperwork as an order delay; INT02 requires compatibility documentation for substitution.

Sources: S05 L5–L6. Two supplied accounts; no common denominator or causal estimate. INT01 also wants faster parts.

F02 • Speed does not uniformly imply switching

INT03 says current lead times usually suffice and would not switch for speed alone.

Source: S05 L7. One respondent; no prevalence inference. Preserve INT01's different account beside it.

RETURN UNSUPPORTED GENERALISATION

Marcus decides exact wording and adequacy. Elaine confirms scope; Nadia reconciles the interview set; Tomas resolves rights and chronology. CV06 remains blocked.

S05 L5–L8; S06; CV06-review-rule.md; prepared-report-specimen.md

The output: a costed scope and delivery brief.

CV07 gives the proposed work a complete fictional starting point. These amounts and approvals are teaching inputs.

Prepared record	CV07 fictional content	Role and human decision
Decision question and scope	Test a non-critical local replacement-part offer; eight notes, three options, a conditional recommendation and report.	Intake analyst prepares; Elaine owns scope. No market sizing, certification or revenue forecast.
Cost workpaper	Role cost A\$13,500 + direct cost A\$1,500 + contingency A\$1,500 = A\$16,500 planning total.	Commercial analyst – sales prepares; Ruth checks cost method. Loaded rates are invented planning inputs.
Fee and obligations	A\$30,000 proposed service fee; A\$10,000 M1 and A\$20,000 M2; timing depends on complete inputs/reviewer.	Writer prepares; signatory approves. Cost/fee surplus A\$13,500 is not collected profit.
Delivery brief and rights	Case-only anonymous paraphrase; no interview reuse; method, source IDs and reviewers named.	Quality checker/coordinator prepare; Nadia/Marcus decide readiness. All permissions/approvals scripted.

NEW CV07 fiction, not the historical CV06 record. One Sonnet generation; owner-edited specimen. Complete inputs/output/cost basis in v6-specimen. Tests are fidelity checks, not acceptance.

v6-specimen INPUT / SPECIMEN; exact locators and arithmetic checked; entirely fictional

The output: an expert decision report.

Should Larch Components test a local replacement-part offer, and which buyer conditions should a pilot require?

Executive summary: the eight invented notes do not establish general willingness to switch for speed. Compatibility approval and component criticality condition trial interest; incumbent performance and absent commitments qualify it. One urgent premium-price purchase is not a general pricing rule.

Method: compare eight purposive qualitative notes on urgency, compatibility/approval, incumbent resistance, price evidence and explicit trial intent. Keep reported facts separate from interpretation. The fixture is complete for its scripted method; it establishes neither population demand nor saturation.

The report includes three findings, the source/method appendix, comparison of no launch, broad speed-led launch and a bounded compatibility-led pilot, and a conditional recommendation. Safety-critical lines are excluded. Supplier capability and economics still need separate validation.

Candidate responsibilities are visible in the outputs: Evidence librarian supplies locators; Delivery analyst prepares findings/options; Deliverable writer prepares the report; Quality checker flags unsupported claims; Case coordinator prepares reviewer returns. This was one Sonnet drafting run, not five separately deployed agents.

Professional adequacy of the inference and recommendation requires qualified review, which has not occurred.

The review return and revised v2 example are scripted teaching material, not genuine professional or client acceptance.

The output: findings with their limits visible.

Each statement carries evidence and the boundary of what that evidence supports.

Finding / evidence locator	Prepared observation	Limit or contrary basis
Compatibility conditions trial; incumbent delivery matters N01, N02	N01 requires engineering sign-off before trial. N02 reports nine of ten incumbent orders within the agreed window.	N02 describes delivery performance, not compatibility. These are separate accounts, not evidence of population dominance.
Criticality limits substitution N03, N05, N07	N03 would trial documented low-risk fittings. N05 says no local-source exception bypasses validation.	N07 holds critical spares and reports no recent stockout. Do not infer a uniform urgent need.
A premium purchase is not a pricing rule N06, N04, N08	N06 paid 12% more on one urgent job. N04 has no switching record; N08 has committed no order or premium.	One transaction cannot establish general willingness to pay or a market switching rate.

Exact source pointers are CV07#/records/Nxx/text. All notes invented; source fidelity does not establish that the recommendation is professionally correct.

CV07 INPUT.json N01-N08; report is an owner-edited synthetic draft, not a population study

The output: alternatives and a conditional next step.

The report prepares a decision; the qualified owner retains the judgement.

Option	Reasoning in the draft	Decision condition
No launch	Avoids qualification cost and forgoes conditional trial interest; remains a reasonable option.	Use if validation effort or supplier capability makes a trial unattractive. No lost-revenue estimate supplied.
Broad speed-led launch	Assumes a general switching/price signal not established by these notes.	Do not recommend on the supplied basis. Incumbent/approval counterevidence remains visible.
Bounded compatibility-led pilot	Explore one non-critical, documented-equivalence part with willing sites; engineering approval before order.	Conditional only. Verify supplier capability, costs, permissions and actual commitment separately.
Scripted review return	Reject the supplied most-customers switching claim and any general price-premium inference.	Marcus is the fictional reviewer. v2 wording is an authored teaching revision, not an actual approval.

The next step in the professional draft is a bounded evidence/feasibility decision. An accepted assessment would not itself authorise a pilot launch.

The output: a payment record that stays open.

Partial payment does not close the invoice. The queue tells Ruth exactly what remains to decide.

Finance state	CV07 fictional record	Exception / next decision
Report and acceptance	Report v2; scripted issue 17 September, client acceptance 18 September.	Acceptance concerns the included assessment, not permission to launch a pilot.
Invoice basis	M1 A\$10,000 + M2 A\$20,000 = A\$30,000 invoiced. Invoice IDs and milestones present.	Ruth owns issue/allocation decisions. No real invoice or receiving-system action occurred.
Payments allocated	M1 A\$10,000 paid; M2 A\$10,000 allocated. Total scripted receipts A\$20,000.	M2 remains partly paid: A\$20,000 - A\$10,000 = A\$10,000 open.
As-of and due date	As-of 2 October 2026; M2 due 5 October 2026.	Not overdue. Preserve open balance until allocated payment or an authorised adjustment resolves it.

Commercial analyst – finance prepares; Quality checker flags; coordinator names owner; Ruth decides. Finance reconciliation is not a collection guarantee or an autonomous payment action.

CV07 INPUT acceptance/invoice/payment records; arithmetic verified; entirely invented finance events

What has actually been demonstrated.

Keep design, synthetic execution and real client operation distinct.

Published website roles describe candidate responsibilities. The three larger jobs are proposed complete designs. Existing W1-W3 instruments and their twenty tests remain controls/specifications; the tests are proposed and unrun.

The earlier Sonnet example prepared observations and a provisional interpretation from eight invented generic notes. Its eleven exact citations passed checks. CV06 is human-authored and blocked.

CV07 is new, separate teaching material. One authenticated Sonnet run produced a proposal, conditional report and finance packet from its invented inputs. Twenty exact citations and arithmetic pass local fidelity checks. The original executed draft is retained; owner wording edits and scripted reviewer v2 are disclosed.

No tenant integration, independently running role chain, professional approval, genuine client acceptance, invoice issue, actual payment, production security result, improved throughput or measured saving is represented. Public research, technical fault diagnosis and full method curation are not demonstrated by these runs.

ai-demo and v6-specimen run/QA receipts; dated status 2 Oct 2026

Accepted scope and evidence brief


When Elaine agrees a scope, Nadia produces one versioned brief and evidence register before any analysis starts.

W1 runs in the case folder and Word, which your records name, and in SharePoint and the project spreadsheet if Tomas confirms them; those two are planning assumptions and no tenant has been inspected. Existing user identities apply. No model or agent is proposed. Whether required columns and version history exist is unverified; Tomas checks features, licences, access and retention before activation. Six installation tests are specified and unrun. Instruments: working-assets/case-brief.md and evidence-register.csv.

Inputs: agreed scope, the S02 method, participant and evidence lists, current permission entries per source

Outputs: AcceptedBrief and EvidenceManifest at a stated version, every expected source marked supplied, verified and permitted, or visibly quarantined

Owner: Nadia drafts; Tomas administers the register and folder
 Human approval: Marcus for methodological adequacy, Elaine for scope and client-facing change; no approval overrides a missing source right
 Failure route: missing source quarantined; wrong permission excluded and recorded; contradictory versions both shown for Marcus; an edited approved brief reverts and becomes a proposed revision; absent owner means the state waits
 Acceptance criterion: brief approved by both approvers with no silently filled field; a repeated submission of the same version creates no second approval



HOLD while scope, permitted inputs or method evidence is missing. Marcus approves adequacy after the hold clears.

S01, S02, S06; tests W1-01 to W1-06 in TECHNICAL-APPENDIX.json, PROPOSED_NOT_RUN

Evidence to reviewable report

From an accepted brief, Nadia prepares a review package in which every finding sits beside its source locator, limits and contrary evidence, for Marcus to decide.

The manual method is the recommendation: finding table, review rule, Word. A bounded draft-preparation component is specified as a candidate only. Its contract fixes input and output ceilings, a time limit, one retry after timeout and zero external actions, with Nadia as owner and Marcus as evaluator. It stays unbuilt and would need its own authorisation before entering any comparison: no external processing right is recorded for any CV06 source, no provider is selected and retention is not yet agreed. Seven tests specified, unrun. Instruments: finding-review.csv, CV06-review-rule.md, prepared-report-specimen.md.

Inputs: approved brief and manifest, approved report structure, S02 rules; only manifest sources may be cited

Outputs: ReviewPackage with findings, gaps, claim-check report and report version

Owner: Nadia prepares; Marcus approves each material finding and the whole

Human approval: a substantive edit after approval invalidates it; no score releases anything

Failure route: missing citation blocks the finding; fabricated identifier fails validation; an instruction embedded in evidence stays text; a revoked source flags dependents; manual preparation is always available

Acceptance criterion: a professionally approved report version recorded in the decision register

DEFAULT • permitted evidence → manual findings / review rule in Word → Marcus review

HELD • bounded AI preparation candidate: separate authorisation, rights and provider checks first

Then compare whole-job quality and effort against the repaired manual route. Candidate remains unbuilt.

S01, S03, S05; tests W2-01 to W2-07, PROPOSED_NOT_RUN

What the agent carries

The candidate does preparation work. The professional still owns the conclusion.

Work to carry	Proposed agent / workflow contribution	Evidence and current status
Assemble the case	Bring permitted sources, the agreed brief, method and report structure into one preparation package.	W1 manual design and instruments exist; no live connection is claimed.
Prepare analysis	Retrieve permitted spans with exact locators. Compare claims against all supplied accounts. Draft bounded findings and interpretations with their basis stated.	W2 custom preparation is a defined, unbuilt candidate. The CV06 specimen is human-authored.
Check the packet	Assemble the review pack and open questions. Check locators/hashes, rights references, source count and changed versions; route errors to the owner.	Local deterministic fixture checks passed. They do not prove professional truth or actual rights.
Prepare the next handoff	Carry the exact report version, acceptance evidence and stable case/contract/milestone key. Detect duplicate/revised items.	W3 design and local contract checks; no invoice is issued. Finance capability is unverified.
Keep judgement and authority human	Marcus decides adequacy; client acceptance stays separate; Ruth decides finance; reuse requires its own right.	No agent approval, external send, auto-invoice, collection result or autonomous reuse is represented.

Recommended now: authorised manual W1/W3 and W2 review rule after readiness checks. AI candidate gate: separate permission, provider/retention checks and comparison before installation purchase.

FIRM-CAPACITY-SYSTEM at main44a2b777; technical/README; local contract check; b10-w2

An actual AI preparation run

Eight invented notes in. Source-linked observations and a review pack out. Human review remains pending.

01 · 8 invented notes	02 · Sonnet preparation	03 · JSON + review pack	04 · Human review pending
OBS01 · A checklist still needs an owner The fictional INT01 participant described spotting missing contacts, approvals and files before kickoff. “the checklist is only useful when somebody owns follow-up” INT01, line 3. No claim of prevented rework or measured time saved.		OBS02 · A decision log preserves the boundary The fictional INT02 participant valued a record of the scope question, decision, owner and date. “unresolved questions should remain visible rather than being marked complete” INT02, line 3. The log does not replace agreement changes or client approval.	

INTERPRETATION TO TEST · NOT A PROFESSIONAL FINDING

The AI linked visibility and ownership across INT01, INT02 and INT05. The reviewer must test that interpretation against the full notes and INT06’s warning about generalising from a small set. No measured improvement follows from this run.

What actually ran: source reading, three attributed observations, one provisional interpretation and a review pack. Excerpts above are condensed from the output; full wording and exact citations are supplied.

This demonstrates AI preparation on synthetic inputs. It does not demonstrate CV06 processing, an installed pipeline, professional approval, time savings or financial results. The 20 installation tests remain unrun.

Separate ordinary fixture, not CV06. ai-demo/findings.json, review-pack.md, run-receipt.json and demo-qa-receipt.json.


Acceptance finance handoff and permitted reuse

When an approved report version has been issued and the client responds, Tomas prepares one internal finance review item for Ruth. Reuse is decided separately.

W3 never issues an invoice, contacts a client or changes bank or finance settings. It is assumed to add about half an operations hour per case, one hour in the sensitivity; that is a design figure, not an observed one. Finance-system field capability is unverified; Ruth and Tomas check. Seven tests specified, unrun. Instrument: working-assets/finance-handoff.md.

- Inputs: approved report version, issue record, client acceptance evidence naming actor and version, contract and milestone references
- Outputs: AcceptanceHandoff keyed by case, contract reference and milestone, report version carried beside the key; a recorded reuse decision per asset

- Owner: Tomas prepares; Ruth alone acts in the finance system
- Human approval: Ruth's recorded decision; Elaine for commercial reuse; Marcus for professional content reuse
- Failure route: absent or qualified acceptance stays pending; conflicting records raise a blocker naming the sources; a duplicate event returns the existing receipt; a mistaken handoff is withdrawn with history kept
- Acceptance criterion: Ruth's finance item reference recorded against the key; a changed report updates the same item and reopens review; a new milestone needs a new key




Internal preparation only. Professional approval, client acceptance and invoice approval stay separate.

S03, S06; tests W3-01 to W3-07, PROPOSED_NOT_RUN

The working outputs your team receives

These are the concrete outputs that connect preparation, review and the next handoff.



Separate reuse decision

Asset, destination, permission and commercial/professional owners.
Elaine/Marcus decide in their domains; Tomas confirms rights.

Comparison and value ledger

Whole-job hours, return reasons, accepted cases, actual cash and unknown costs.
Named buyer owners decide whether to install, simplify, defer or stop.

Templates and record contracts are included. An accepted design is distinct from an installed capability or an actual human approval.

working-assets/; technical/schemas/; CONTENT-v3 W1-W3

The record controls inside the larger jobs.


Four record interfaces, one rule for approval, and a human at every transition.

AcceptedBrief, EvidenceManifest, ReviewPackage and AcceptanceHandoff each carry case id, record id, version, what they supersede, state, source references and errors. Approval binds to a version and a digest, never to a filename. A record changes only by creating a new version; history is never deleted. The finance key is case plus contract plus milestone, so a replaced report cannot create a second finance item.

Approvers: Elaine for scope and commercial change, Marcus for professional adequacy, Tomas for access and reconciliation, Ruth for finance. Every deputy is unknown and due before the first live case; without one, the affected state waits rather than passing to someone unnamed.

Data stays in the restricted case folder under existing user identities. No provider, region or retention term has been checked, so any model route remains unverified and the W2 candidate stays manual. Current-study permission never implies cross-client reuse.

Every blocked transition keeps the last authorised version and records reason, affected records, owner and recovery. Re-entry is a named decision in working-assets/decision-register.csv.



Same case identity and version record connect the designs. Acceptance never grants cross-client reuse rights.

Shared controls; the technical folder holds the interface specifications

A system your firm can operate.

The three jobs share sources, decision records and a controlled operating method.

Existing working tools

Word and case folders remain. Finance remains its own authority; proposed connectors require capability checks.

Context graph

Linked scope, sources, claims, approvals, acceptance and payment records preserve what supports each decision.

Human decisions

Named owners decide scope, professional issue, rights, invoices and disputes; changed versions reopen review.

Permitted evidence / RAG

Retrieval-augmented generation (RAG) supplies permitted sources for drafting. A fixed approved pack may be enough.

Bounded preparation

Defined model tasks and deterministic calculation/checking prepare inspectable commercial outputs.

Operation and change

Owned exceptions, versioned releases, budget limits, stop/recover, evaluation and documented handover.

Logical architecture, not installed infrastructure. No graph database, provider, tenant integration or licence compatibility is represented.

v6 proposed design; fictional buyer, not installed capability

```
graph LR; S03[Source N03  
CV07 / v1; case-only rights] -- supports --> F2[Finding F2  
Criticality conditions trial;  
linked to N03]; F2 -- contains --> Rv2[Report v2  
CV07; contains finding and  
its limits]; Papp[Partner approval  
Marcus; exact v2; scripted  
event] -- approves --> Rv2; Ca[Client acceptance  
Exact v2; 18 Sep; scripted  
event] -- accepts exact version --> Rv2; Pa[Payment allocation  
A$10,000 to M2 invoice;  
scripted receipt] -- allocated to --> IR[Invoice record  
M2 / A$20,000; A$10,000  
remains open]; IR -- bills --> M2[Milestone M2  
Assessment delivery;  
signed basis required]; Ca -- covers --> M2;
```

Source N03
CV07 / v1; case-only rights

Finding F2
Criticality conditions trial;
linked to N03

Report v2
CV07; contains finding and
its limits

Partner approval
Marcus; exact v2; scripted
event

Payment allocation
A\$10,000 to M2 invoice;
scripted receipt

Invoice record
M2 / A\$20,000; A\$10,000
remains open

Milestone M2
Assessment delivery;
signed basis required

Client acceptance
Exact v2; 18 Sep; scripted
event

Block further retrieval -> identify F2 / v2 dependencies -> queue owner review. Preserve history; do not reverse invoices or payments automatically.

Original CV07 teaching subgraph; Neo4j decision-trace reference + Astra control synthesis. All records/approvals fictional. DIAGRAM-REFERENCES-v6.md.

Find the right evidence. Exclude the rest.

Retrieval-augmented generation (RAG) supplies relevant sources for drafting. Choose the simplest retrieval route that meets the job.

Approach	Use it when	What still must be proved
Approved fixed pack	The case sources are small, stable and intentionally selected.	Exact locators, version/rights manifest, provider processing permission and missing-evidence route.
Permission-aware RAG	Evidence is larger or changing; repeated retrieval warrants indexing.	Relevant evidence found; wrong-client, stale and revoked sources excluded before model use/output.
Permitted method library	An approved method or template can help another eligible case.	Rights, applicability, version and expiry; source content is not reusable merely because it was redacted.
Graph-based retrieval	Tested relationship questions exceed what current structured records handle well.	Benefit over simpler records; permissions, costs, export and maintainability. No graph engine selected.


Retrieval supports grounding; it does not establish truth, professional approval or security. Technical references are in SYSTEM-DESIGN-v6.md.

Design recommendation; Microsoft retrieval/access controls and Neo4j relationship context, checked 2 Oct 2026

How an answer gets its evidence.


Prepare permitted sources. Retrieve passages. Draft with citations. Review the result.

01 / WHEN SOURCES ENTER OR CHANGE



Index + current permissions feed retrieval.

02 / EACH TIME A QUESTION IS ASKED



Human review remains a separate decision.

Recheck release rights. Marcus judges the inference and approves the exact report version.

Small, stable case? An approved fixed pack can supply context directly. Retrieval does not train the model or establish truth. Retrieved files cannot grant tool authority.

Original proposed design; Microsoft RAG architecture + BnB controls. References: DIAGRAM-REFERENCES-v6.md. Fictional buyer; not installed.

Three terms. Three different jobs.

Evidence retrieval and decision dependencies belong together, but they are different mechanisms.

Mechanism	What it does	Buyer question it helps answer
RAG	Retrieves source passages and supplies them as context for generation.	Which permitted evidence supports this draft finding?
Context graph	Stores typed connections, current state and decision evidence; enables dependency queries.	Which report and pending decisions depend on a changed source?
GraphRAG	Uses graph relationships as part of retrieval to assemble context for generation.	Which connected entities and passages help answer a relationship question?
Selection rule	Start with the simplest route that meets representative quality, permission and operating tests.	Does a fixed pack or linked register meet the job with less maintenance?

These mechanisms can be combined. A connected register is not automatically GraphRAG. Microsoft GraphRAG is one implementation, not a required product. No supplier has been selected.

Microsoft RAG / GraphRAG local search; Neo4j context-graph guides. Design interpretation and direct reference links: DIAGRAM-REFERENCES-v6.md.

Security belongs in the workflow.

The model cannot grant itself access, approve client work or spend without the agreed control path.

Boundary	Required control	Test / decision owner
Source and identity	Case/purpose access outside the model; sales and finance identities separate.	Try wrong-client/revoked access: deny before context and export. Tomas + security owner.
Untrusted content	Retrieved files are evidence, not tool instructions; tools and outputs allowlisted.	Inject a source instruction: no added privilege, client send or confidential exposure.
Provider and retention	Approve processing rights, region, training/retention, encryption and minimal protected logs.	Unknown route blocks client processing. Confirm actual supplier terms, not a compliance badge.
Consequential action	Human approval for professional issue, fee, finance and external action.	No model approval or draft status authorises action. Named human owns release.
Consumption and recovery	Time/cost/retry limits, secrets outside prompts, stop and manual fallback.	Provider failure/tool revocation stops safely; last authorised version remains usable.

RAG does not eliminate prompt injection. These are proposed security requirements and tests, not a security certification.

Know what stopped. Know who acts next.

A useful system must remain inspectable when a source changes, a provider fails or its operator leaves.

Operating responsibility	What is included in the design	What must be agreed
Management view	Case/version, stage, exception, next owner, source validity and permitted duration/cost summary.	Who reads it; alert destination, operating hours and response/service checks.
Safe recovery	Disable switch, last authorised checkpoint, bounded retry, duplicate protection and manual route.	Who can stop, diagnose, restore and restart; recovery never approves a report.
Controlled change	Versioned method/model/template/tool configuration and representative regression cases.	Affected tests rerun before release; new responsibilities need scope approval.
Operable handover	Interfaces, configuration, sources, runbook, limits, test results and access ownership.	Use rights, licences, export formats, retention, backup/restore and transition scope.

Optional managed operation is separately scoped and priced. No automatic retainer, 24-hour cover or perpetual support is implied.

v6 proposed design; fictional buyer, not installed capability

Choose who runs it after handover.

Three operating models are available after installation. None moves professional or commercial authority away from your firm.

Operating model	Who carries the work	What stays with the firm
Client operated	Handover configuration, templates, tests and runbook. The named client operator owns the queue and change process.	Professional issue, rights, fee and finance decisions; access administration and budget oversight.
Optional managed operation	Build n Bloom carries the agreed monitoring, exception handling, updates or administration tasks.	Named owner, approval boundaries, service checks and operating hours. Management is separately priced.
Scoped hybrid	Client operates daily; Build n Bloom carries a named responsibility such as tested updates or restore support.	Everything outside that written scope remains with client owners; no automatic retainer.

The delivered Blueprint and agreed use rights remain with the buyer. Configuration export, third-party licences, logs, retention and exit rights are agreed in writing.

Fable SH-10, condensed by Root; frequencies/fees/cover remain unagreed

Build around the way your firm works.

The useful difference is a complete job fitted to your method, evidence, tools and review standard.

What can adapt	Buyer benefit to test	Boundary / tradeoff
Method and deliverables	Your terminology, approved analysis, scope questions, report format and reviewer criteria.	Method owner approves changes; templates do not confer professional correctness.
Sources and existing tools	Allowed case sources and the exact file/finance touchpoints used by the firm.	Tenant, licence, API and rights checks first; preserve existing finance authority.
Role responsibilities	Define preparation, checking, coordination and decision responsibilities for every output.	No custom setting can bypass source rights or grant a model human authority.
Management and exit	Owned exception views, exportable records and a documented manual/transition route.	Operating responsibility remains; third-party licences and agreed use rights apply.

Customisation is agreed configuration and change control. It is not unlimited features, universal integration or ownership of third-party software.

v6 proposed design; fictional buyer, not installed capability

Price the complete operating job.

A lower model bill is not the total cost of a usable system.

Cost category	What belongs in the quote	Current status
One-off implementation	Rights/data preparation, configuration, connectors, evaluation, training, acceptance and restoration preparation.	Expanded scope requires re-estimation. No new installation fee or date quoted.
Continuing operation	Model/retrieval/storage, licences, review/operations, backups, updates and optional management.	Provider, volume and support responsibilities unknown; agree payer and spending ceiling.
Change and exit	Permitted method changes, new sources/responsibilities, exports and transition support.	Written use rights, licence/retention obligations and restoration scope required.
Buyer participation	Owners provide evidence, compare outputs and make professional/finance decisions.	Old 10.5-15.5 hour illustration does not price expanded participation; re-estimate.

Seller founder time remains excluded. Buyer work and external costs still matter. The preserved v5 monetary curves are conditional historical examples, not a v6 forecast.

v6 proposed design; fictional buyer, not installed capability

Let the next case inherit a permitted method.

Reuse and recovery support the three jobs. Neither adds a fourth promised installation.

When a case closes, the Knowledge curator responsibility proposes a method asset: the retained instructions, excluded sources, version, expiry or retirement condition, and outputs affected by later retirement. Or it records non-reuse. Elaine decides commercial reuse, Marcus professional reuse, and the source rights owner decides what may travel.

Redaction cannot grant permission. The next enquiry or assessment may use the asset only after its own rights and applicability check. The existing package demonstrates a reuse decision record; full curation and its future preparation benefit remain unexecuted and unmeasured.

If an integration, index or service is interrupted, the Technical investigator reads authorised logs and proposes a fault hypothesis and recovery test. The Case coordinator lists affected cases and owners. Technology and risk owners approve restoration or rollback. A recovered connection never approves client work.

Manual continuation is available only when required permitted sources and authority remain available. Otherwise the affected case pauses with an owner. Record-shape checking is not evidence of diagnostic capability.

Fable BL-14, condensed/qualified by Root; proposed supporting branches, not executed

When work stops, the owner is named

A blocked case should produce an owned task, not a silent queue or a guessed answer.

Condition	Useful response	Owner and re-entry
Source right missing or revoked	Quarantine the source; flag dependent findings; preserve the last authorised version.	Tomas reconciles the permission record before permitted preparation resumes.
Unsupported claim or material edit	Return the exact version with source, limit and reason. A substantive edit invalidates prior approval.	Nadia corrects; Marcus reviews the new version.
Scope or client authority unclear	Record the open boundary before more work is prepared.	Elaine confirms scope/authority; Marcus checks method adequacy.
Acceptance and finance evidence conflict	Hold the internal handoff; identify contradictory records and stable case key.	Tomas reconciles; Ruth decides finance after supported acceptance/milestone.
Owner or deputy absent	Keep the state waiting. Preserve history, affected IDs and a reason.	A separately authorised human deputy or the recorded owner re-enters. No timeout approval.

Fallback is source-linked manual handling. No silence, retry or score grants approval. These are design responsibilities, not a claim of live workflow activation.

operator-playbook.md; technical/README; b12-connected-control

What else could be right and what would show it

Keep each alternative open until the evidence distinguishes it.

If the qualified pipeline does not sign and collect, released hours remain unpriced capacity and the added-case contribution argument for installation falls away. Other quality or cost benefits could justify a different investment case, but none is established here. Elaine and Ruth own the signed scope and collection checks.

Hire the senior reviewer. Adds review hours across all services; allocation to this path unknown. Not comparable to the study on a cost basis, and no salary-coverage arithmetic is a return. Stays open with Elaine. Decided by whether stage-level data shows reviewer hours binding.

Rely on drafting seats. Already present. Nadia's recorded account covers manual note collection and version checking; whether seats change that is untested. Decided inside the comparison, where the seats plus the manual rule form one arm.

Configure existing tools. Required columns, version history and finance fields may already cover W1 and W3. Decided by Tomas's live capability check; if it holds, both designs shrink to configuration.

Interview scheduling or transcription. Not in the sample; hours unknown. Later, if evidence appears.

Client acceptance portal. Rejected: the gap is the record, not the channel.

What building the capability can change.

The useful advantage is complete, controlled work fitted to your firm. Its economic advantage must be measured.

The service repeats forty times a year in the fictional buyer setting. The supplied case records show review returns and gaps, but they do not tell us how much reconstruction occurs across all forty cases. Measure that burden before the next hire, seat or build decision.

A complete decision pack can put the source, analysis, contrary evidence and open judgement beside the report. The intended benefit is less reconstruction and correction, with the same professional standard. A context graph makes affected work visible when a source, scope or approved version changes.

The same operating design can carry accepted work and its commercial basis to finance, then keep invoice and payment states distinct. Method curation can help a later eligible case without turning client evidence into unrestricted shared memory.

Those capabilities can add configuration, review, security and maintenance work. Compare the current method, a repaired method using existing tools and the connected candidate at comparable accepted quality. If the simpler option reaches the bar with less burden, it should win.

Measure full preparation, review, correction and operation by role. A quicker draft alone is not a faster accepted case.

Treat released hours as capacity. Money needs an evidenced paid-cost change, useful redeployment or additional signed, accepted and collected work.

Never sum the same hours across the three designs. No avoided hire, collected contribution or quality improvement has been observed.

Fable SH-11, condensed with Root evidence limits; existing ECONOMICS.json remains fictional

What you would be buying beyond a draft.

Judge the offer by the accepted work and operating control it leaves behind.

Buying consideration	What this design adds	What proves the difference
Beyond an isolated draft	Complete scope/report/finance outputs, connected records and owned return paths.	Representative whole-job tests and comparison with existing seats/configuration.
Evidence that can be inspected	Exact sources and versions, permitted purpose, contrary evidence and a visible decision basis.	Retrieval/citation checks and qualified review; no generic accuracy badge.
Fit to your service	Approved method, report format, source rules, responsibilities and exact existing-tool touchpoints.	Client owners accept the configured job; new responsibility reopens scope/tests.
Operation and choice	A working handover, manual route, export/exit terms and optional managed responsibility.	Agreed operator tests/use rights and full written guarantee; no compulsory retainer.


These are proposed buying reasons, not claims that competitors lack them. Another provider or a standard product may meet the same requirement.

Root/Fable synthesis; proposed product properties, not market uniqueness or outcome proof

Hours per case under each target

Baseline 18 senior, 42 analyst and 6 operations hours. Both targets below are untested design numbers.

Hours per case ▪ baseline / repaired method / whole-path target



The targets describe the whole job after all changes. Savings are never attributed to one design and then summed; most of the senior and analyst change sits in preparation and review, but the records cannot isolate it. Operations time is assumed to rise because W3 adds a verification step; the half hour is a design figure, not an observation.

Released hours are role-specific capacity. They are not payroll, not a saving, and not a reason to defer the hire question. Stage attribution, adoption and the actual quality effect have to be measured in the comparison before any of these numbers is relied on.

Method only, W1 and W3 manual plus the W2 review rule: 17 senior, 40 analyst, 6.5 operations

Whole path, all three designs including the W2 candidate: 14 senior, 36 analyst, 6.5 operations

Operations sensitivity: one added hour per case instead of half


All targets are untested; operations work rises.

S01 baseline; targets assumed by design, E-07 | Preserved v5 scenario; expanded v6 scope requires re-estimation.

CAPACITY

How many added cases the released hours could carry

Nine benefit months, 30 existing cases, demand of four. Every figure is conditional on the untested targets.



Demand cap: 4 collectible cases. All role caps must hold.

On the whole-path target the role caps are 8 senior, 5 analyst and 6 operations added cases, so analyst hours give the tightest role cap; demand of four is lower still and is what binds. Both statements are design outputs, not observations, and one reason the comparison must record whole-job effort by role rather than reviewer time alone.

Capacity is not demand. Four cases is the number the pipeline could supply if it signs; none has. Collection is a further step owned by Ruth.

Whole path: releases 120 senior and 180 analyst hours; operations 60 spare less 15 added leaves 45

An added case needs 14, 36 and 6.5 hours, so the caps are 8, 5 and 6; demand of four binds, result four


Method only: releases 30 senior and 60 analyst hours; caps 1, 1 and 6; result one

Operations sensitivity at one added hour: 30 spare remain, an added case needs 7, cap 4; result still four

From contribution to subtotal under one installation scenario

An installation-scenario analysis, not a return on the Blueprint alone.

A\$000 • four collected cases • unknown costs remain



The A\$24,000 is an illustration, not a quote. Unknown and not zero: marginal cost of an added case beyond the A\$20,000 already inside the contribution figure, recurring oversight, incidents and change, provider usage beyond the allowance, financing and tax. Each lowers the subtotal one for one. Final net value is unknown.

Not counted anywhere: the A\$120,000 hire as avoided cost, released hours as cash, invoice lag as collection, or any seller time. Against the method-only route, the whole-path scenario lists A\$24,800 more cash, A\$30,800 against A\$6,000, and models three more collected cases worth A\$30,000 of contribution. The difference is an A\$5,200 subtotal before unknown costs. That is the incremental projection the comparison would have to bear out, not value that exists.

Listed scenario cash: Blueprint A\$5,000, illustrative installation A\$24,000, provider allowance A\$1,200 and a hypothetical buyer preparation cash allowance A\$600; total A\$30,800. The A\$600 is not a valuation of staff hours; actual paid cost and its basis remain unverified.

Four added cases collected: A\$40,000 contribution; subtotal A\$9,200 before unknown costs


Method only, one added case: A\$10,000 less the A\$5,000 Blueprint and an A\$1,000 repair placeholder; subtotal A\$4,000

FIN03, scenario inputs; ECONOMICS.json, economics.csv | Preserved v5 scenario; expanded v6 scope requires re-estimation.

What happens if fewer cases collect or installation costs more

Subtotal equals collected cases times A\$10,000, less installation, less A\$6,800 of listed non-installation cash, less unknown costs.

Listed-cost subtotal (A\$000), before unknown costs



At four collected cases the installation ceiling for a non-negative subtotal is A\$33,200 before unknown costs; at three it is A\$23,200. Any unknown cost takes the subtotal below zero at those ceilings. Treating the Blueprint as sunk at the installation decision adds A\$5,000 to each. These are arithmetic limits, not budgets or acceptable risk.

The method-only route has a smaller downside because it has no installation to lose. Its placeholder repair cost must be replaced by approved figures before any commitment. Ruth owns the recomputation with the firm's actual cost method.

Four cases collected: A\$21,200 at A\$12,000 installation; A\$9,200 at A\$24,000; minus A\$2,800 at A\$36,000

Three cases collected: A\$11,200; minus A\$800; minus A\$12,800

No added case collects: minus A\$30,800 at the A\$24,000 illustration

ECONOMICS.json; illustrative installation band | Preserved v5 scenario; expanded v6 scope requires re-estimation.

Six benefit months against nine

The window matters because capacity is released per existing case, and fewer existing cases pass through a shorter window.

6 benefit months

3

9 benefit months

4

Read these as conditional available-capacity scenarios. A quarter's delay does not automatically cost one case or A\$10,000; demand, signing and collection decide that, and none is measured. What the window does show is that the start date of authorised method use is a decision with consequences, which is why the Next decision page asks you to fix it rather than let it drift. Every contribution figure here is conditional and net of nothing; it is not a return on the Blueprint alone.

Nine months, 30 existing cases, whole path: caps 8, 5, 6; result four with demand of four

Six months, 20 existing cases, whole path: caps 5, 3, 4; result three even if demand is four

Annual steady state, 40 cases and demand of six: result six, or five under the higher operations burden; A\$60,000 contribution if all collect, A\$58,800 after the provider allowance only, net unknown

constrained added cases

-A\$0.8k subtotal

Conditional scenario: demand remains 4; targets must hold.

constrained added cases

A\$9.2k subtotal

OPS01, PIPE01, PIPE02; ECONOMICS.json | Preserved v5 scenario; expanded v6 scope requires re-estimation.

From a better packet to a commercial result

The agreed benefit must survive professional review, whole-job effort and the buyer's actual revenue model.

Change to establish	Measure to test	What it permits you to conclude
Preparation carries evidence and exceptions into review.	Source completeness, unsupported claims, return reasons and whole-case quality.	A more inspectable packet is useful; quality improvement needs representative cases.
Less reconstruction, with no work hidden elsewhere.	Preparation + review + correction + coordination hours, by role.	Only the net whole-job change can release capacity. A faster first draft is insufficient.
Released capacity meets real work.	Eligible accepted throughput, case mix, signed demand and available people.	Four unsigned nine-month pipeline cases are a ceiling to test, not booked revenue.
Accepted work becomes collected contribution.	Actual acceptance, invoice/receipt and buyer contribution records.	Count cash when collected; no avoided hire or payroll saving is established.
Timing changes the available test window.	Next eligible case, readiness date and the six-/nine-month assumptions.	At untested whole-path targets: three vs four added cases; -A\$0.8k vs A\$9.2k listed-cost subtotals before unknowns.

Why act now: decide the method before the next eligible case repeats the same preparation. Name an owner and start date. This comparison does not price a quarter's delay or establish lost revenue.

ECONOMICS.json; b14-b18; S03, PIPE02; FIRM-CAPACITY-SYSTEM


How we compare quality and effort

A predeclared rubric, matched cases and an honest inconclusive route. Nothing here has run.

The plan compares arms on representative cases: the current method, and the manual repair with drafting seats. The W2 candidate joins only after a separate authorisation covering its processing rights, provider and retention; it is not built. The rubric is declared before any case runs. Quality is judged on whole-case outcome by Marcus: return reasons, unsupported generalisations caught, source and rights completeness. Effort is whole-job hours by role recorded by Nadia, including any checking the change creates elsewhere.

For this added-case investment argument, a W2 installation quote requires quality not worse, lower whole-job effort and signed, collectible demand. Mixed results or an inadequate sample are inconclusive: continue the manual method, extend the sample and hold the installation purchase.

Stage hours, case mix and return reasons must be captured first; they are open items under Next decision.



Use the same permitted case and professional quality rules. Twenty proposed installation tests remain unrun.

working-assets/comparison-plan.md


Putting the manual method into service

Everything needed to run W1 and W3 by hand, and to hand it to another team.

Any field change creates a new template version and re-runs the affected tests. Stop conditions: a critical unsupported conclusion, cross-case disclosure or an authority failure halts live use. Instruments: working-assets/commissioning-brief.md and operator-playbook.md.

- Register columns: case id, contract reference, milestone, scope and brief versions, evidence register version, reviewer, deputy, state, approved report version, acceptance evidence, actor and date, finance item and receipt ids, reuse decision, retention class
- Folder per case: scope, restricted sources, register, working, review, approved, acceptance, finance handoff, reuse decisions

- Training: a proposed 90-minute session covering a walkthrough of the operator playbook and practice on one historical case; readiness means that case completed in the register unaided. The session is not the whole activation effort.
- First live case supervised by Nadia with Marcus available
- Routines: weekly register check by Tomas, monthly rights and freshness check by Nadia, quarterly retention review by Elaine and Tomas
- Exit: register and decision register as CSV, folder tree, templates, test matrix and logs



Begin with repaired manual control in existing tools. Installation is a separate scoped decision.

Commissioning design; deliverable ownership recorded in the actual agreement

Test the complete job.

Passing one preparation step does not accept the whole installation.

Design test	Expected behaviour	Who decides
Scope / ordinary and changed terms	Complete traceable scope/cost/proposal; missing cost and reviewer hold; signed differences reopen review.	Elaine + Ruth + signatory; Nadia/Marcus accept delivery.
Report / complete and disputed evidence	Complete report and exact citations; contrary evidence retained; unsupported claims returned.	Marcus evaluates professional adequacy; citations alone do not prove it.
Report / reviewer return	Revised version cannot inherit approval; unresolved method question remains visible.	Marcus approves exact version; client acceptance is separate.
Finance / missing or partial records	Wrong acceptance/amount blocks; duplicates do not create another item; partial payment leaves open balance.	Ruth controls finance actions and reconciles actual receipt.

All installation tests proposed and unrun. Scope-listed representative examples, reviewer and pass thresholds must be agreed before installation.

v6 proposed design; fictional buyer, not installed capability

Test the boundaries as well as the happy path.

Agree lawful fixtures and forbidden effects before client operation.


Shared test	Expected behaviour	Evidence to retain
Retrieval permission and freshness	Wrong-client, revoked, expired and stale sources excluded; relevant allowed evidence found.	Source/version/rights basis; retrieval and citation checks separate.
Context dependency	Changed sources flag dependent work for owner review. A revised report cannot inherit approval; history is preserved.	Dependency list, change and return reason, exact actor/version.
Security / injection / leakage	Untrusted instruction cannot add privilege; restricted logs/export and sales/finance boundaries hold.	Denied action and permitted minimal event record; no secrets.
Operations / change / exit	Stop/recovery/manual route works; regression before promotion; agreed export and restore usable.	Authorised checkpoint, test result, change/release and handover record.

Full 25-test proposed schedule: SYSTEM-DESIGN-v6.md. Existing twenty W1-W3 tests remain legacy control specifications, not passed installations.

v6 proposed design; fictional buyer, not installed capability

Who decides what and the participation estimate

Five owners, five deputies to name, and estimated participation to confirm.



Participation was estimated at 10.5 to 15.5 hours across the five roles, plus unknown hours to reconcile CV06, CV05 and INT04. Actual hours and any paid costs belong in working-assets/participation-and-cost-record.md so the next phase is planned from evidence rather than from the estimate.

Method activation and any later build are separate phases with their own configuration, rights reconciliation and ongoing burden, none of which is counted yet. Buyer hours are time, never a payroll saving, and seller time is excluded from every financial figure.

Elaine Park: scope, commercial change, commercial reuse, the hire question, the refund decision

Marcus Cole: methodological adequacy, professional approval, professional reuse, the comparison judgement

Nadia Singh: delivery, brief drafting, effort recording, freshness checks

Tomas Reed: access, records, reconciliation, register administration

Ruth Chen: every finance decision, collection evidence, contribution method

10.5–15.5 h total; seller founder time excluded from economics.

Participation estimate; deputies due before first live case | Preserved v5 participation illustration; re-estimate for expanded v6 scope.

Your Blueprint satisfaction guarantee

You can use the Blueprint satisfaction promise even when the recommendation is to simplify, defer or stop.

Blueprint: paid upfront. You have 14 calendar days after both complete delivery and the decision review to request a full refund if it was not worth the fee. No defect proof, correction round or further purchase is required. Cash is refunded within five business days of the request. You keep the design and agreed use rights. Mandatory rights remain.

There is no claim fee; reasonable necessary valid-claim expenses are reimbursed. No payment is refunded twice and no business outcome is guaranteed. Full terms and claim contacts are in GUARANTEE-REFERENCE.md. Richard Opondo owns commercial fit; Rayan Ratego owns architecture and release readiness.

14

CALENDAR DAYS

5 business days

cash refund after a request

Clock starts after BOTH complete delivery and decision review.

What a later installation must agree

No installation is commissioned by this Blueprint.

A separate quote fixes scope, fee, deposit, date, core tests, reviewer and restoration before commitment. Balance is payable only on written acceptance. Core failure is corrected and retested free within ten business days after the failure report and agreed lawful evidence or access. Genuine dependency pauses require a documented, agreed revised date; provider-controlled delays are not client dependencies. Any workaround or scope reduction requires your explicit agreement.

If core failure persists after correction, you may stop: installation fees and the associated Blueprint fee actually paid are refunded once within five business days, balance waived and agreed restoration and handover completed free. A provider delay beyond the delivery date permits a free written replan or termination with installation fees refunded within five business days. The Blueprint retains its own promise unless the core-failure remedy applies.

Thirty calendar days of stabilisation follow written acceptance. A timely core-failure report retains the remedy beyond that period. Minor defects are corrected free; expanded responsibility is separately agreed. Optional support has its own service checks and ten-business-day correction period; persistent failure permits termination and refund of the affected period within five business days, with no future support balance. Support does not begin automatically or refund the installation and Blueprint again.

Summary. Full terms, permitted evidence route and mandatory rights: GUARANTEE-REFERENCE.md.

Before the next case enters review

Choose the next eligible case, name the owner and set the date.

The next eligible case gives this decision a practical moment. If it reaches review needing context rebuilt, the reconstruction happens again and its effort goes unrecorded; how large that burden is remains unmeasured. The benefit window exists only where forthcoming work and owner availability exist; no pipeline case is signed.

If scope, rights or reviewer availability are missing, hold the case and repair the record before the comparison. Name the next eligible case and its start date. Record proceed, challenge or defer per design; assign the CV06 chronology check first. The hire-or-method decision follows observed stage data. The full Blueprint satisfaction promise applies as stated; external business outcomes are not guaranteed.

Elaine confirms the decision question, scope and client authority; Tomas confirms rights per source and resolves INT04.

Nadia prepares the manual source-linked pack for that case in Word under the review rule.

Nadia records whole-job hours by role; Marcus records return reasons.

Elaine and Marcus compare total burden and accepted quality, then decide whether the candidate earns a separately scoped quote or the manual method is sufficient.

Confirm case + rights	Prepare manual pack	Compare whole job	Decide install / defer
-----------------------	---------------------	-------------------	------------------------

If a dependency is missing, hold and repair before comparing. The case and date still need to be named.

Fable supplement S5; next eligible case and dates unconfirmed; b18 benefit-window assumptions.

What to record and do now

Record a decision on the complete designs, then scope the first installation separately.

Unresolved items hold the action they affect, not the whole programme. Reconciling CV06 does not stop W1 from starting on a clean case. The one thing that should not wait is a recorded decision on each design, because the next capacity commitment, whether a hire, a seat or a build, will otherwise be made on the same records that could not settle these questions.

Record proceed, challenge or defer for scope, expert delivery and finance; include dissent and unknowns.

Prioritise permitted complete case evidence, approved method and professional review history for the expert decision-pack comparison.

Validate enquiry/cost/signed-term evidence and finance/acceptance/payment records before asserting the other commercial opportunities.

Reconcile CV06 chronology, CV05 acceptance, INT04 rights, actual stage hours, demand, cost basis, collection, tool capability, deputies and retention.

Use authorised manual W1/W3 controls only after rights, readiness and capability checks; compare whole-case effort and accepted quality.

Agree the separately quoted first installation, representative tests, provider route, spending ceiling, operating owners and restoration scope.

If the complete Blueprint was not worth its fee, its full satisfaction promise remains.

working-assets/decision-register.csv holds every open item with owner, check, due and effect

Where the complete material lives

The copy in this document summarises. The files below carry the full detail.

Two limits on the technical material. The fixture checker runs fourteen bounded structural checks on sample records with standard tooling; it is not a full schema validation engine and says nothing about any live record. The ordinary fixture is authored teaching material separate from CV06; it does not fill CV06's real gaps, which stay blocked until reconciled.

Evidence labels: client-reported items come from S01 to S06 and the PIPE and OPS entries; targets, allowances and the installation band are assumed by design; the CV06 specimen is human-authored. The separate synthetic AI preparation example has its own output and receipt; it establishes no professional outcome or installation acceptance. The CV06 candidate and proposed installation tests remain unrun.

v6 main narrative governs the expanded three-job design. SYSTEM-DESIGN-v6.md carries retrieval/context/security/operation requirements and 25 proposed tests. ASTRA-v6-DESIGN.md and FABLE-v6-NARRATIVE.md are advisory authorship contributions. v6-specimen holds the separate CV07 synthetic inputs, outputs, provenance and fidelity checks. No old execution receipt is evidence of these new outputs.

Implementation-Reference.pdf and .md: the preserved detailed v2 design schedules and execution specifications; the v5 main narrative governs newer stack and editorial clarifications

TECHNICAL-APPENDIX.json: the twenty proposed installation tests, all PROPOSED_NOT_RUN, with inputs, forbidden effects and stop rules

technical folder: interface specifications for the four record types, the ordinary eight-source teaching fixture, blocked examples for CV05 and CV06, and the bounded fixture checker

working-assets: case-brief.md, evidence-register.csv, finding-review.csv, finance-handoff.md, case-value-ledger.csv, decision-register.csv, commissioning-brief.md, operator-playbook.md, document-review.md, comparison-plan.md, participation-and-cost-record.md, CV06-review-rule.md, prepared-report-specimen.md

ECONOMICS.json and economics.csv: every scenario input, editable

GUARANTEE-REFERENCE.md: full applicable terms

current-stack-register.csv: tool, recorded evidence or assumption, proposed aid, owner, capability check and manual fallback

CHANGES-v6.json records scope amendments; preserved economics/guarantee/source records retain their original status.

Simulation. Fictional source material; no client result, installed system, passed installation test, executed quote or signed instrument is represented.