

From enquiry to quote. Where should we start?

This is a hypothetical example, not a client case study. All assumptions and scores below illustrate a decision method. They are not measured BODHIVA results or a quote for services.

1. The decision

A service business assembles quotations from three systems. Choose the smallest responsible experiment before committing to an AI build.

2. What we assume - and still need to learn

- The same information is re-entered during quote preparation. We have not measured the hours or frequency.
- Templates and approval roles can be changed internally. Client approval time is outside the team's direct control.
- Data access, contractual permissions, error rates and integration feasibility are unverified.
- Quote preparation is a controllable intervention point, not a proven system bottleneck. Delays elsewhere may dominate end-to-end turnaround.

3. Compare credible alternatives

Gates before scores: No pilot proceeds without authorised data access, an accountable owner, an agreed review process and tolerable residual risk. A high average cannot override a failed gate. A failed gate means defer or redesign that option.

Scores are provisional judgements from 1 to 5. They must be replaced after workflow observation. Higher means better in every column. We use equal weights: a simple average, with no claim of statistical precision.

Initial comparison under the assumptions above

Option	Value	Feasibility	Control	Average
Shared template + clear ownership	3	5	5	4.33
Rules-based data assembly	4	3	4	3.67
AI-assisted first draft	4	2	2	2.67
Change client approval process	3	2	2	2.33

Anchors: Value 1 = limited local convenience; 3 = plausible recurring benefit; 5 = evidence of a material operating gain. Feasibility 1 = major unknowns/dependencies; 3 = bounded integration work; 5 = readily trialled with current tools. Control 1 = weak ownership/review or high residual risk; 3 = manageable with explicit controls; 5 = clear ownership and reversible, inspectable outputs. Intermediate scores sit between the anchors.

Why these scores? A template is easy to inspect and reverse, although its value is unmeasured. Rules may remove copying but need tested access. AI may help drafting, but error behaviour, data handling and review effort remain unknown. Client approval changes require cooperation beyond this team. Ease of control does not establish the greatest possible impact.

4. Recommended first test

Trial a shared template and named owner first. Record preparation time, corrections and end-to-end turnaround against the existing process. Add rules-based assembly only if repetition and data access justify it. Compare AI assistance with that simpler baseline, including the time spent checking and correcting its output.

5. What would change the recommendation?

- Observation shows client waiting time dominates and quote preparation is already quick: test approval communication first.
- Data access is reliable and copying drives material errors: a rules-based pilot may become the first investment.
- Unstructured drafting drives effort and a representative test demonstrates net benefit after review: consider a limited AI pilot.
- Ownership, access rights or useful measurements cannot be established: defer the build and resolve those foundations.

6. Pilot ownership and stop condition

Agree one workflow owner, a representative sample and a measurement period before testing. Establish baseline preparation time and correction rate. Human review reduces risk but does not guarantee errors will be detected. Record missed errors and review effort. Stop if the experiment increases critical errors, lacks authorised data access or cannot demonstrate a useful improvement against the baseline.

7. The next decision

At the agreed review point, decide to improve, expand or stop. A full implementation is a separate decision, with scope, costs and responsibilities agreed in advance.