A STRATEGIC PROPOSALS AND VISIBLETHREAD GUIDE

From AI experimentation to bid orchestration

Building a coherent proposal AI strategy

Strategic Proposals | VisibleThread

Why this guide, and why these two organisations

This is a joint publication, and the two halves answer different questions.

STRATEGIC PROPOSALS · SECTIONS 1 TO 3

The maturity question

Strategic Proposals contributes the AI Maturity in Bidding benchmark and the Proposal AI Operating Model. Sections 1 to 3 use them to ask what a mature proposal AI capability actually looks like, and why so few organisations have one.

VISIBLETHREAD · SECTIONS 4 TO 6

The quality question

VisibleThread contributes the quality and compliance argument, and the deterministic tooling behind it. Sections 4 to 6 ask a narrower question: once AI is producing content at speed, how do you know the output is any good, and what do you do about it on Monday morning?

The guide is written to be read in order, because the second half assumes the model built in the first.

CONTENTS

SECTION	WHAT IT COVERS	PAGE
1. The AI inflection point	Why the adoption question is settled, and what replaced it	3
2. Why AI maturity stalls	Five patterns that can stall maturity, and how organisations progress	6
3. The proposal AI operating model	Five cumulative layers, and why the order matters	10
4. Where AI still falls short	Failure modes, task-shape risk, grounding, and the rules-versus-AI split	13
5. Two teams, one bid	The same pursuit run with and without structure	19
6. Where to start	Four sequenced steps and a 90-day shape	24
7. From tools to transformation	What the argument comes down to	28

01

The AI inflection point

Why the adoption question is settled, and what replaced it



The conversation has moved on

For much of the last two years, discussions about AI in bidding and proposal management have centred on a single question: **should we be using AI?** Today that question is largely irrelevant.

Whether organisations have formally adopted it or not, artificial intelligence has already entered the bid process. Proposal managers are using it to create outlines and first drafts. Subject matter experts are using it to summarise technical concepts and accelerate response development. Sales teams are using it to analyse opportunities, generate briefing notes and prepare customer-facing content. It is no longer an emerging capability being tested on the sidelines. It is embedded, in one form or another, in most proposal organisations.

The real question is no longer whether bid teams should use AI. It is whether they are using it in a way that creates sustainable competitive advantage. For many organisations, the answer remains uncertain.

Strategic Proposals' recent AI Maturity in Bidding benchmark revealed a profession experiencing rapid adoption but comparatively low maturity. AI activity is increasing quickly across the bid lifecycle, yet relatively few organisations have established the governance, capability and measurement disciplines required to scale that activity into consistent business value.

> **AI adoption is accelerating faster than AI maturity.**

The advantage is in the knowledge, not the model

The distinction matters because the next competitive advantage in bidding is unlikely to come from technology alone. As access to increasingly capable models becomes widespread, differentiation will depend less on which model an organisation uses and more on the knowledge, context and expertise it can bring to that model.

For decades, proposal organisations have focused on helping people find information. AI changes the equation. Increasingly, success depends on how effectively the right information can be surfaced, trusted and applied at the right point in the bid process.

Large language models are becoming available to everyone. Access is commoditising. Organisational knowledge is not.

> **Today, organisations compare tools. Tomorrow, they will compare knowledge.**

A familiar pattern in a new technology cycle

The bid and proposal profession has seen technology revolutions before, from content management systems to collaborative authoring tools. AI is different in kind. Previous waves improved how work was managed. This one changes how work is created, because unlike traditional software, AI effectiveness depends on factors that sit outside the technology:

01	The quality of organisational content
02	Governance and controls
03	The maturity of knowledge management practices
04	User capability
05	Workflow design
06	Measurement and accountability

Without those foundations, AI amplifies existing weaknesses rather than solving them. The organisations seeing the greatest value are the ones that invested in the systems, processes and capabilities needed to support AI at scale.

The three waves of proposal AI

Across the industry, three waves of maturity are emerging.

WAVE ONE

Experimentation

AI enters informally, as individuals explore tools on their own. Few standards, little governance, minimal measurement. Success judged on anecdote.

WAVE TWO

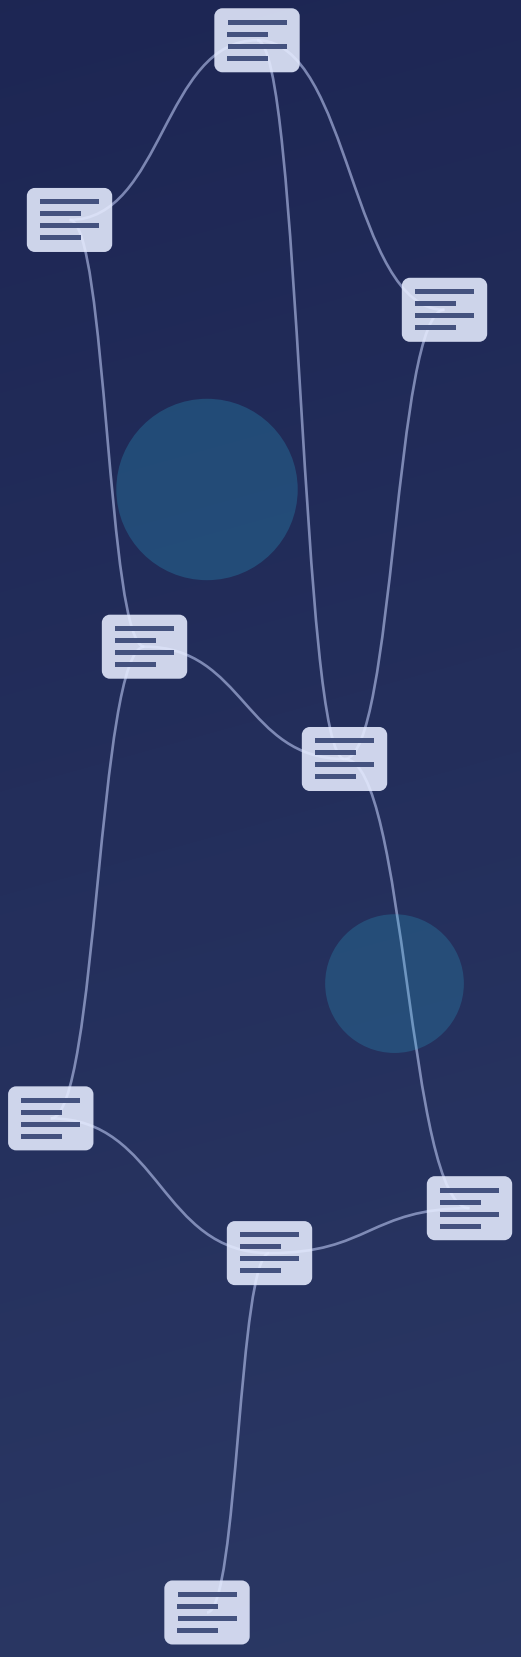
Adoption

Experimentation spreads. Licences, prompt libraries, specialist copilots and agents, training and early workflow automation become more common. Activity is widespread, but different teams may still use different tools, approaches and standards.

WAVE THREE

Orchestration

Content is structured and governed, roles evolve, processes repeatable. AI becomes part of an operating model, not an isolated initiative.



02

Why AI maturity stalls and how organisations move beyond it

Five patterns that can stall maturity, and what organisations do differently as they progress

Why organisations see different levels of value

If AI is now embedded in the profession, an obvious question follows: why are organisations seeing such different levels of value from it?

Some organisations are already approaching proposal AI strategically: establishing governance, curating knowledge, redesigning workflows and developing capability across their teams. Others remain earlier in the journey, with AI activity growing faster than the operating model around it.

The distinction matters. AI creates the greatest value when it is treated not simply as a technology deployment, but as an operational transformation involving people, process, knowledge, governance and measurement.

Activity is not the same as maturity

7% **of organisations report linking AI activity to measurable ROI and bid performance outcomes**

Licences purchased, active users, prompts generated, agents deployed: all are useful indicators of activity. But on their own, they reveal relatively little about whether AI is creating meaningful business value.

The more useful question is not simply how much AI is being used, but whether that use is improving proposal quality, reducing effort, strengthening consistency and contributing to better outcomes.

The distance between adoption and value has therefore become one of the defining challenges, and opportunities, facing proposal leaders today.

From AI islands to organisational capability

Across many organisations, AI adoption initially develops in pockets. Proposal managers create their own prompts, agents or workflow automations; SMEs use different platforms; sales and capture teams develop separate approaches.

That experimentation is valuable. It surfaces use cases quickly and helps organisations understand where AI can make a difference. The challenge comes when organisations try to scale what works. Different approaches can produce different outputs, quality levels and standards.

Expertise can remain concentrated in individuals rather than becoming organisational capability, while teams may duplicate work by solving similar problems independently. The opportunity is to capture what works, standardise it where appropriate and make successful approaches available across the organisation.

From individual capability to organisational capability

Consider two proposal managers in the same organisation. Same platform, same content, same training session. One consistently produces first drafts that survive review. The other does not.

> **The technology is identical. The capability is not.**

Individual experimentation and learning are essential, particularly while organisations are establishing what works. But enterprise capability cannot ultimately depend on individual expertise alone.

As organisations mature, successful practices need to become repeatable: embedded in shared workflows, guidance, knowledge, agents and ways of working so that performance is less dependent on who happens to be staffed on a particular opportunity.

The knowledge challenge

Historically, proposal content libraries were often viewed primarily as repositories. AI is changing the importance of proposal knowledge management.

For years, organisations could tolerate content that was difficult to find, inconsistently structured or spread across multiple sources because experienced people knew where to look and how to interpret it. AI makes the quality and organisation of that knowledge much more consequential.

Consider two organisations developing their AI capability. One concentrates principally on increasingly sophisticated prompting and automation. The other also invests in curating, validating and organising the knowledge that sits behind those tools. The second is building something more difficult for competitors to replicate: a trusted organisational knowledge base that can support consistently stronger outputs.

> **AI does not remove the need to manage knowledge well. It increases it.**

Governance enables scale

In the early stages of adoption, governance can understandably lag experimentation. Organisations want teams to explore the technology and identify where it can create value. As adoption expands, however, consistent operating principles become increasingly important.

Governance is not about restricting innovation. It is what allows successful experimentation to become trusted, repeatable organisational capability.

56% **of organisations operate with inconsistent guardrails**

Measure value, not just activity

Prompt volumes, user numbers, agent usage and adoption rates help organisations understand activity. Mature measurement goes further and asks what that activity is achieving.

- Has SME effort been reduced?
- Has proposal quality improved?
- Has rework decreased?
- Has time to first draft accelerated?
- Are teams spending more time on strategy and differentiation?
- Are bid outcomes improving?

Faster does not always mean better

Much of the current conversation about proposal AI remains focused on generation. Yet some of the greatest opportunities may sit earlier in the process.

Mature organisations increasingly use AI not simply to accelerate writing, but to improve proposal preparation: finding relevant content, locating evidence, connecting SMEs with knowledge, analysing customer requirements, establishing response foundations and preparing contributors.

03

The proposal AI operating model

Five cumulative layers, and why the order matters

Most proposal AI initiatives begin in the same place: identify a use case, deploy a tool, encourage adoption, measure activity. Many begin at the wrong layer of maturity. Every struggling programme has typically overlooked at least one of the foundational layers required for maturity, because AI maturity is cumulative rather than isolated.

AI scales inconsistency faster than it scales performance.

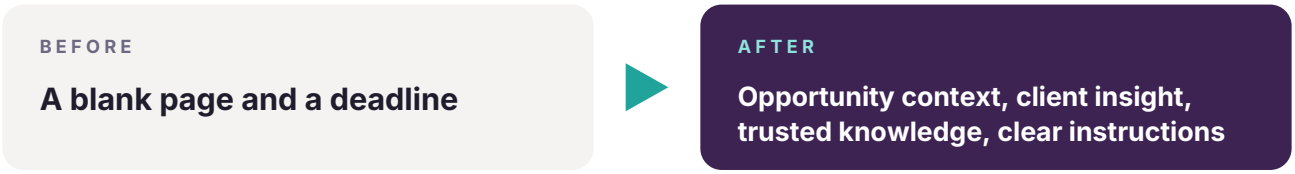


The model is cumulative rather than optional. Trusted knowledge creates the foundation. Governance creates consistency. Skills create capability. Workflow creates scale. Measurement creates continuous improvement. Together they create orchestration. **Miss a layer and AI may still function. It is far less likely to transform performance.**

How the work changes

As these layers mature, the proposal process itself changes.

Rather than beginning with a blank page, teams increasingly give AI the right opportunity context, client insight, trusted knowledge and clear instructions about what strong proposal content should look like for that customer and opportunity. AI can then accelerate analysis, structure and first-draft development.



The role of the proposal professional shifts increasingly towards orchestration, direction and review:

- Shaping the response strategy
- Determining what the AI needs to know
- Validating its output
- Strengthening differentiation
- Applying judgment

Subject matter experts can spend less time recreating information the organisation already knows and more time validating, challenging and adding expertise.

> **Crucially, this does not reduce the need for people to understand what excellent proposal content looks like. It increases it.**

Teams need that expertise both to instruct AI effectively and to recognise when its output is not good enough.

04

Where AI still falls short

Failure modes, task-shape risk, grounding, and the rules-versus-AI split

Fluency and accuracy do not travel together

Generative AI has solved the blank page. It has not solved the harder problem, which is knowing whether what came back is any good. A language model predicts the most plausible next words. Obvious errors get caught. Confident, well structured, entirely invented claims do not.

> *"Everybody knows AI is valuable. But they don't quite know how to leverage it best, where it can be applied, when it can be trusted, and when it cannot."*
>
> FERGAL MCGOVERN · CPO, VISIBLETHREAD

Four failure modes, none of them visible in the document

01

Fabrication

A certification, a metric, a past performance detail that is clean but does not exist, or does not say what the draft implies.

02

No audit trail

Without a traceable path back to real documentation, correct content and invented content look identical.

03

Drift

Run the same prompt twice and you get two different answers. Across a long pursuit that variability compounds.

04

Automation bias

Reviewers approve fluent output faster than they read it, and the more dependable a system proves, the less carefully it gets checked.

AI generates well. It judges badly.

The instinctive fix is to use AI to check AI. This fails for a structural reason: you are asking the same probabilistic engine to grade its own output, using the same weighting that produced the problem in the first place. The model treats every word in your content estate as equally authoritative.

> *"If you point AI at a big corpus of content, much of which is out of date, AI has no way to differentiate old stuff from new stuff. Every word is treated equally."*
>
> FERGAL MCGOVERN · CPO, VISIBLETHREAD

Risk is a function of task shape

Error risk tracks the structure of the task far more closely than the sophistication of the model. Task structure is something you design: add scaffolding and grounding to a task and you move it down the gradient.



LOWER RISK — HIGHER RISK

Extraction
Answer already exists in the source. Lowest risk, fully traceable.

Open-ended generation
No single correct answer. Needs human validation.

Chained sequences
Errors propagate through every downstream step. Deserves the most caution.

A misclassified requirement in step one does not remain a step one problem. It propagates into the compliance matrix, into the outline, into the draft, and into the review that checks the draft against a matrix that was already wrong.

Two kinds of quality, two different mechanisms

Bid quality is not one thing. The mistake most organisations make is asking AI to handle the second category. It will produce something that looks like a compliance check. It will not be complete, and it will not be the same twice.

SUBJECTIVE QUALITY

Stays with your people

Whether the win theme lands, whether the discriminator is genuinely a discriminator. This is judgment. It requires context the model does not have, and it stays with your people.

OBJECTIVE QUALITY

Rules have right answers

Whether every requirement is addressed, whether terminology is consistent across twelve authors. These are not judgment calls. They are rules, and rules have right answers.

The four properties of grounding

Grounding is not a setting. It is four properties, and a system either has them or it does not.

PROPERTY 1

Scope

Draws only from a defined set of sources you control.

PROPERTY 2

Attribution

Every passage traces back to the content it came from.

PROPERTY 3

Freshness

Knows how current its sources are, never prefers a superseded version.

PROPERTY 4

Refusal

Says so when approved content does not cover a requirement.

The rules half: objective evaluation

Governance defines what good looks like. A deterministic quality layer is what proves it, on every bid, without depending on who happens to be available to check. In practice it means checking:

01	**Requirement coverage**	Every will, shall and must extracted from the tender and mapped to a response, with the exact page and paragraph. Not a sample. All of them.
02	**Amendment traceability**	Every change between versions surfaced line by line, so you can prove the change you asked for actually made it in.
03	**Language discipline**	Watchword lists that flag claims you cannot evidence and terms your legal team will not sign off on.
04	**Readability against the reader**	Grade level, sentence length, passive voice, complexity, tested against the person who has to read it once.
05	**Acronym integrity**	Undefined on first use, defined inconsistently, used before definition, conflicting definitions across volumes.
06	**Terminology consistency**	One organisation’s voice across every contributor, held automatically rather than by one person reading the document four times.

Deterministic means repeatable and auditable wherever a rule can be explicitly defined. It does not mean automatically correct. What it will not do is return a different answer on Friday than it returned on Monday.

The principle

Speed and control are not in tension, provided you stop asking one technology to deliver both. Use generative AI where variability is acceptable and a head start is the point. Use deterministic, rules-based technology where the answer must be identical every time and complete on the first pass. Keep human judgment for the questions that decide whether you win.

GENERATIVE AI

Where you need a head start

- Outlines
- Research
- First drafts
- Restructuring
- Synthesis

DETERMINISTIC RULES

Where you need proof

- Extraction
- Compliance matrices
- Version comparison
- Acronym and language checks

HUMAN JUDGMENT

Where the bid is won

- Win themes
- Discriminators
- Response strategy
- The questions that decide whether you win

> *"You need rules to satisfy the need in many use cases for 100% accuracy and 100% repeatability. AI is always predictive and it's variable. If I do it today it'll be different from tomorrow."*
>
> FERGAL MCGOVERN · CPO, VISIBLETHREAD

Rules where you need proof. AI where you need a head start.



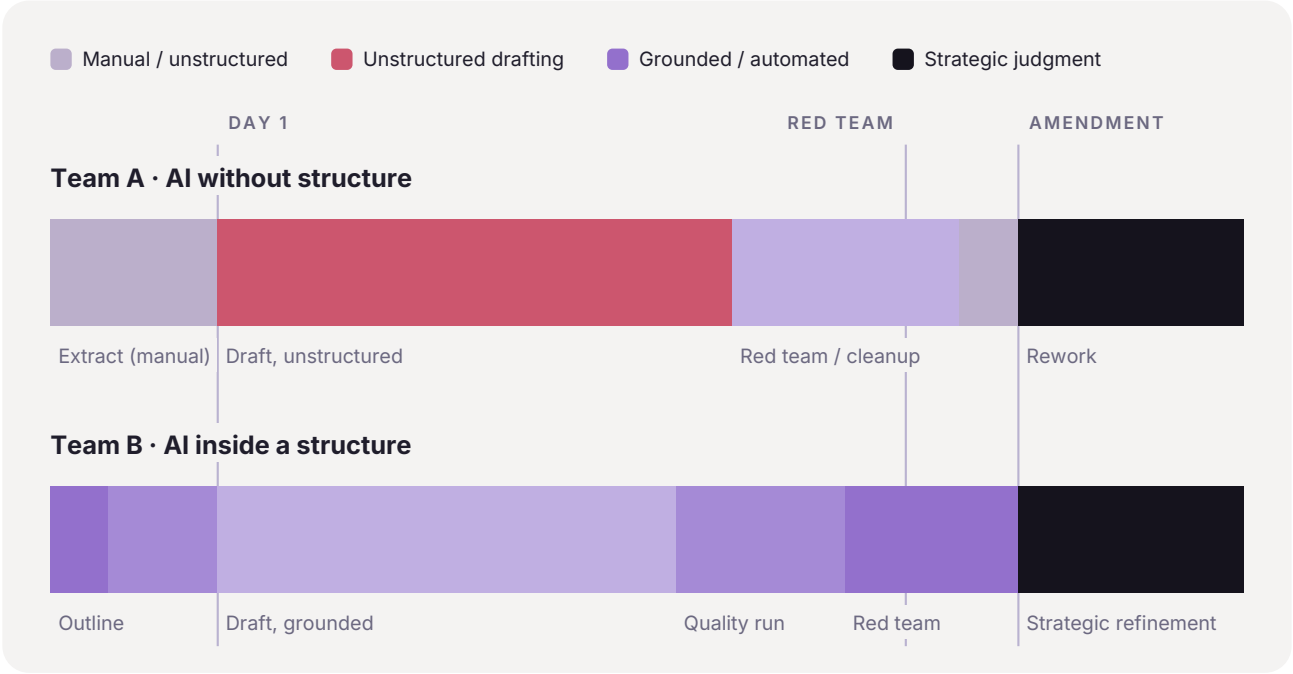
05

Two teams, one bid

The same pursuit run with and without structure

The scenarios below are illustrative rather than drawn from a single account, but every failure and fix within them is one both organisations encounter repeatedly. *The procurement terminology is US federal. The pattern is not.*

The setup. A 900-page solicitation. Twenty-one days to submit. A bid manager, four subject matter experts, a price author and a capture lead who has been tracking the opportunity for eight months. Two amendments will land during the cycle, one of them on day seventeen.



Note what did not change: headcount, deadline, or the volume of the documentation. The difference is entirely in where the work went, and in whether the quality layer was a rule or a hope.

Where the difference actually shows up

	TEAM A	TEAM B
Requirement extraction	3 days manual, completeness unproven	Same day, 100% repeatable, page-referenced
Content sourcing	Whatever each author could find	Curated, approved, cited to source
Amendment handling	Manual, one miss	Automated, line by line, audit trailed
Quality review	During red team, competing with strategy	Before red team, cleared in advance
Senior time on strategy	Roughly 2 days	Roughly 8 days
Post-submission defensibility	Reconstructed from memory	On the record

> **The gain across these examples is not that people work faster. It is that the expensive people work on the expensive problems.**

The same pattern in practice

Marieke Bland, Director of Proposals at KMS Solutions, has run this shift with proposal teams at Octo, IBM and KMS. Manual version comparison became a repeatable step anyone on the team could run. Compliance matrices, previously built by hand on every bid, were reduced to a near-instant step.

> *"The teams that have figured out AI adoption have not done it overnight. If you are not starting now, you are already behind. VisibleThread gave us the structure to do it right."*
>
> MARIEKE BLAND · DIRECTOR OF PROPOSALS, KMS SOLUTIONS

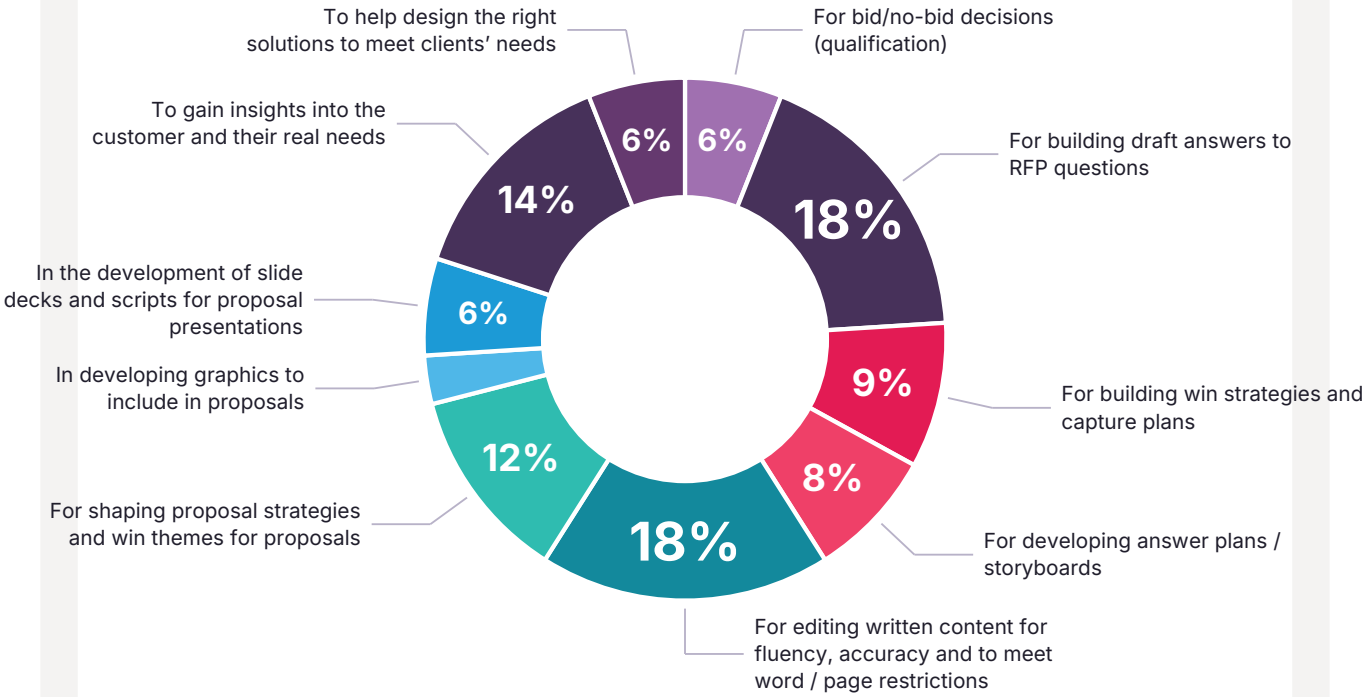
The same pattern holds after award. Parker Hannifin runs contract review across its Aerospace divisions on the deterministic side of this model: customised dictionaries flag critical terms automatically, contracts are checked against regulatory clauses, and document comparison surfaces version differences without the manual scroll.

> *"Doing this manually would take a lot longer and would be subject to user fatigue. By the end of contracts, eyes get blurry and you start missing things."*
>
> JENNIFER ALVARADO · CONTRACT ADMINISTRATOR, PARKER HANNIFIN AEROSPACE

AI is being used, but maturity is what creates value.

The organisations achieving the greatest impact are not simply adopting AI. They're embedding it into repeatable processes, governance and bid-winning ways of working.

WHERE AI IS BEING USED IN CURRENT BID PROCESSES



Source: Strategic Proposals, AI Maturity in Bidding benchmark



06

Where to start

Four sequenced steps and a 90-day shape

This is not a multi-year transformation programme. The organisations getting value from proposal AI are not the ones with the largest budgets. They are the ones that sequenced four things properly, in the order the operating model implies.

1 Assess what is already happening

You almost certainly have more AI in your bid process than you have on record. Run this as a factfind, not an enforcement exercise: the unauthorised use is the useful signal, telling you exactly where the process is failing people badly enough that they built their own workaround.

OUTPUT an honest map of current usage, with the confidential-data exposures ranked.

2 Audit content maturity before touching tooling

AI amplifies whatever content estate you already have. Assess findability, currency, modularity, ownership and evidence. The highest-leverage move is curation over volume: a collection of your fifteen best case studies produces better grounded output than a pointer at all one hundred.

OUTPUT a prioritised content backlog and two or three curated collections you would trust to ground a live bid.

3 Define governance as a set of decisions, not a policy document

This fails when written as principles. It works when written as decisions someone can act on at 4pm on a Friday: what AI generates, what deterministic systems validate, what humans own outright, where content may come from, who signs off, and what must never leave the governed environment. Write it on one page. If it needs twenty, it will not be followed.

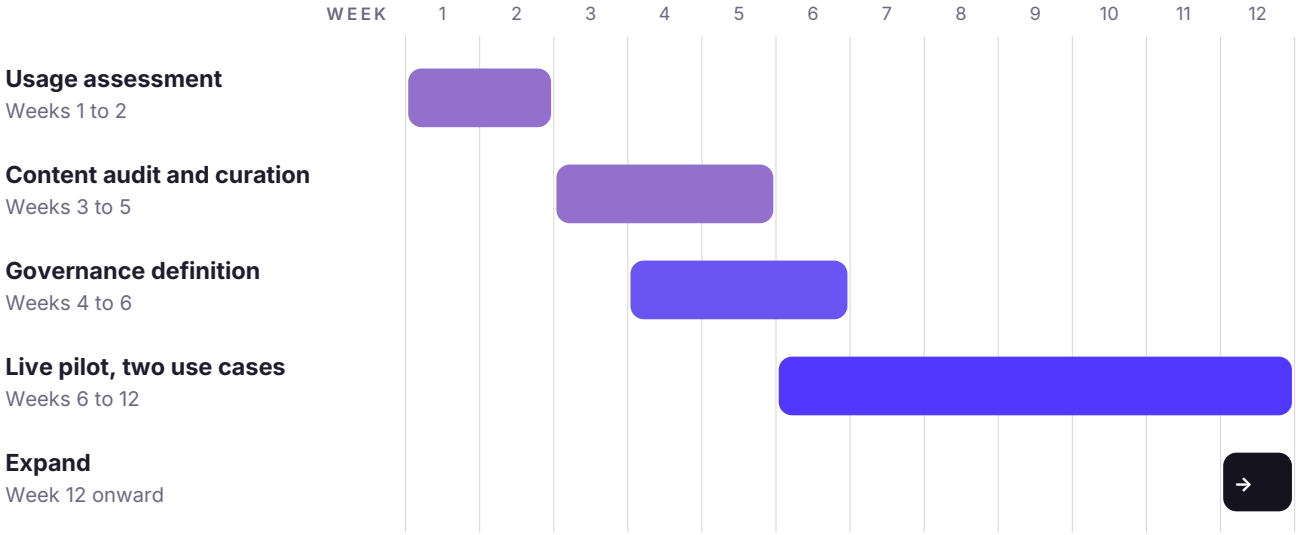
OUTPUT a one-page operating agreement your bid teams can actually apply.

4 Pilot on a live bid, not a sandbox

Pick two use cases, one deterministic and one generative, and run them on a real pursuit with a real deadline. Time the current process first: without a before number, every after number is an anecdote. Build a golden set: ten representative solicitations, fifty requirements whose answers you already know, two or three amendment scenarios. Then expand by lifecycle stage, not by department.

OUTPUT evidence from a real bid, which is the only thing that moves a sceptical proposal team.

The 90-day shape



Four ways this goes wrong

Buying tooling before auditing content

The tool will work. The output will not, and the tool gets blamed.

Writing governance nobody reads

If the rules are not usable under deadline pressure, they do not exist.

Piloting in a sandbox

Clean test data proves nothing. The problems live in real tenders with real amendments.

Treating AI adoption as a technology project

It is a process maturity project. Fix the workflow first, or you will automate the breakage.

From tools to transformation

The question this guide opened with is no longer a live one. AI is in your bid process. It arrived without a business case, and in most organisations it arrived without a policy. What is still in front of most organisations is the harder decision underneath it. Not which tools to buy. How to make the ones already in use cohere.

The constraint was never the model

Access to capable models is commoditising and will continue to. The constraint is everything the model depends on and cannot supply for itself: whether the content it draws from is current and approved, whether anybody has defined what it is permitted to do, whether its output is checked by a rule or by hope. Those are the five layers. They are cumulative, and organisations that skip to the visible ones tend to discover the missing foundation at the least convenient moment.

The asset you already own

The organisations that pull ahead over the next few years will not be the ones with the most advanced models, because everyone will have those. They will be the ones that curated the knowledge their AI works from, decided in advance what it was allowed to do, proved the things that have to be proved rather than trusting them, and spent the hours they released on the questions no model can answer.

Rules where you need proof.

AI where you need a head start.

Judgment where the bid is won.

Three things you can start this week

None of them require a vendor, a budget or a decision.

1. **Run the usage fact-find.**
   One session, no consequences attached to the answers.

2. **Curate two collections, not the whole library.**
   Your fifteen strongest past performance records and your current approved capability narratives.

3. **Write the one-page operating agreement.**
   What AI generates, what rules validate, what humans own outright, and who signs off.



When you are ready to talk to Strategic Proposals

Move beyond AI experimentation, take the AI Benchmarker to help you assess your current capabilities, benchmark your progress against your peers, and identify your next steps towards AI maturity.

Review your results with a Strategic Proposals consultant and gain practical guidance on the actions that will deliver the greatest impact.

Take the AI Benchmarker →



When you are ready to talk to VisibleThread

Bring a real solicitation, ideally a long and badly formatted one. Bring the area you scored yourself weakest on. Bring your hardest constraint. We will run the tooling against your document rather than ours. If the honest answer is that content foundations should come before a purchase, we will say so.

visiblethread.com →





strategicproposals.com · visiblethread.com