2026 Proposal & Bid Software Report

A Structural Analysis of RFP and Proposal Technology

Published April 2026 | First Edition

stargazy | stargazy.io

How to use this report

This report is a comprehensive independent analysis of the software market for Request for Proposal (RFP) response. It is structured as reference material, not as a sequential read.

An interactive version of this report is available through NotebookLM. Buyers can query the report directly with situation-specific questions, or access an audio synthesis for consumption. The interactive version is designed to surface the analysis most relevant to an individual team's constraint rather than requiring cover-to-cover reading.

The report's analytical framework is built around a single principle that proposal technology purchasing errors happen when buyers evaluate platforms before classifying their binding operational constraint. This report exists to invert that sequence.

Sponsor disclosure. This research is sponsored by lup and AutoRFP.ai. Sponsorship underwrites distribution. It does not influence category definitions, evaluation criteria, inclusion decisions, or vendor conclusions. The complete methodology and independence standards are documented in Appendix A.

Research currency. This report reflects publicly available evidence through Q1 2026. Vendor capabilities, pricing, certifications, and product positioning evolve continuously. The next edition is scheduled for Q1 2027. Buyers should validate current vendor status during their own evaluation.

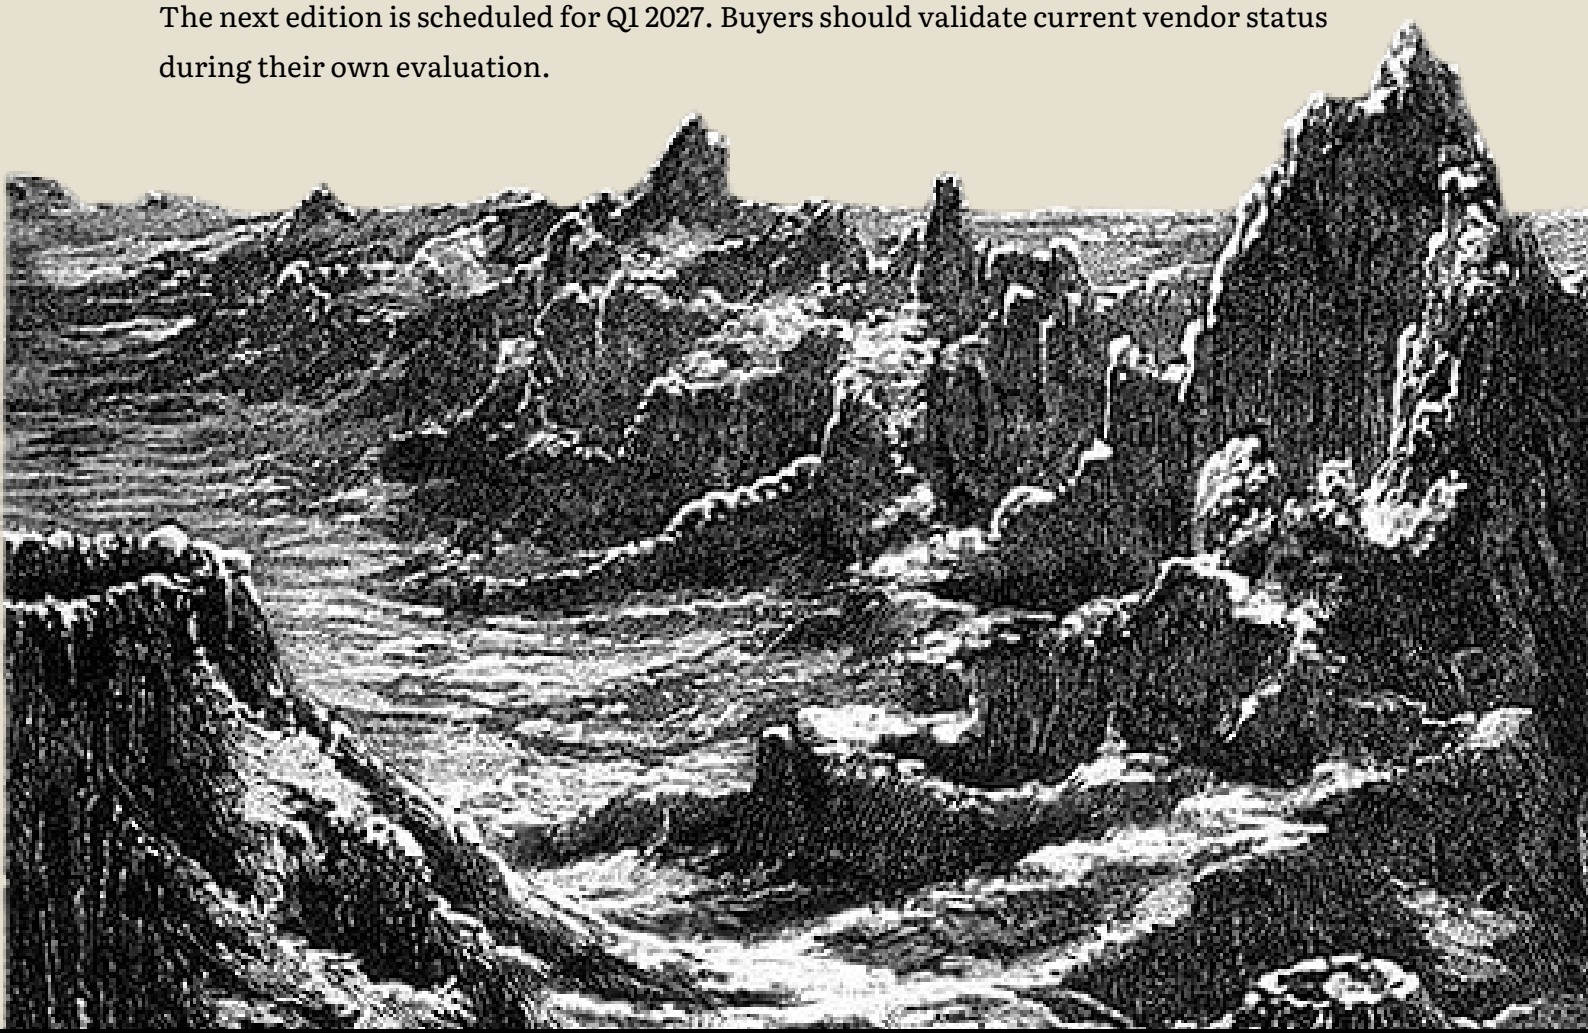

Flammarion, Camille. “The View of the Earth from the Moon.” Astronomie populaire, 1880, engraving.

Most proposal teams are buying software designed for a problem they don’t have.

The proposal category has optimized for drafting speed and content reuse. Then Generative AI reduced the cost of drafting, but it did not reduce the cost of a losing draft. In most environments, AI moved the challenge from writing a proposal into review, lengthening SME queues and reducing confidence in what gets submitted to the buyer.

There are five architectures a proposal team can buy in 2026. Four of them will not fix the constraint your team is actually facing. This report is the map to the one that will.

If proposals drive more than 20% of your pipeline, the first move is a team’s proposal constraint diagnosis. Do that before you sit through another demo.

The 2026 proposal & bid technology taxonomy

Revenue leaders can see almost everything now, from conversation intelligence capture, to CRMs tracking where deals sit, to forecasting tools modeling what will close. Yet most revenue organizations have no signal at all about what will happen to an RFP opportunity after a proposal is submitted.

This is where proposal software sits in the revenue stack, creating an empty space that forecasting models stumble over.

Proposal and bid software share surface-level similarities (content libraries, AI drafting, collaboration features) but differ in architecture. Architecture determines whether a platform gives your team a system of record for the highest-stakes document in the deal cycle, or a faster way to produce unverified content under deadline pressure.

We call this trust fidelity. Can the platform trace a claim back to a source, and does it require approval from an accountable reviewer before that claim reaches a customer? That single question separates platforms that reduce revenue risk from platforms that accelerate it.

Why control surface determines category

Every team responding to Requests for Proposals (RFPs) is struggling with something unique to them, whether that is coordination, win rates, having time to respond to all qualified RFPs, or something completely different.

When it comes to purchasing software, the most common buying error occurs when teams purchase an AI-drafting software to solve a governance problem. The second most common error occurs when teams treat a general-purpose automation backbone as proposal technology.

Five control surfaces characterize the 2026 primary market.

- **Managed execution control.** Enforces RFP and questionnaire intake discipline with explicit human ownership, deadlines, and requirement coverage. The team operates; the platform coordinates and provides optional AI assistance.

- **Autonomous execution control.** Enforces intake-to-submission discipline with the AI as the default mover. Humans review, approve, and handle exceptions rather than acting as the primary authors.

- **Drafting control.** Optimizes requirement parsing and narrative generation from governed sources.

- **Compliance and capture control.** Enforces regulatory mapping, evaluator-aware response logic, and capture-to-submission continuity within public-sector and defense procurement.

- **Domain-specific evidence control.** Manages structured personnel records and project references for bid environments where past-performance evidence determines evaluation scores.

A sixth dimension, governance, cuts across all five. Governance is the enforcement of approval states, permissions, auditable claim lineage, and content freshness. Every category must meet this performance requirement regardless of its primary control surface.

System of record strategy determines drift and maintenance economics

Every proposal stack must decide where their proposal content truth lives because this decision defines long-term cost. We see three models recur:

- **Platform-stored truth.** The proposal tool stores approved content. This accelerates reuse but creates recertification debt as answer volatility rises.

- **Upstream-stored truth.** The tool retrieves facts from authoritative systems such as product documentation, CRMs, security repositories, and internal wikis. This reduces duplication but requires permission enforcement and claim-level citation.

- **Hybrid truth.** Proposal-specific artifacts live in the platform; factual assertions come from upstream sources like SharePoint, OneDrive, Teams, or Confluence. The platform indexes content in place rather than migrating it into a proprietary database.

When teams correct a claim in one system and not another, contradiction rates increase across proposal submissions. In regulated environments, this increases review cycles and legal scrutiny.

The operational question is,“Where do we correct inaccurate content six months from now?”

If the answer requires updates in multiple repositories, the organization will repeat the error.

Four patterns of platform failure

Every proposal team hits at least one of these patterns in production, and before buying a software, a team should know which pattern their binding constraint produces:

- **Review overload.** The team adopts an AI drafting tool, but what if response volume doubles and so does the SME validation queue? The issue simply moved downstream and gets heavier.

- **Integration fragility.** The software doesn't integrate with every system it needs to, to perform what the team needs it to, further siloing the proposal content, team, and data.

- **Traceability gaps.** A procurement team flags a claim in a submitted proposal. The proposal leader searches the platform for approval lineage and finds only the edit log.

- **Workflow bypass.** SMEs bypass the platform. They reply in Slack or over email. Within three months, the proposal manager is copying content from email threads into the system by hand, and the orchestration layer is theater.

Teams that avoid these failures test for each one before they commit. They run a live RFP through the platform using real integrations and named approvers. They measure approval cycle time and rework volume over 90 days. If two or more of these patterns show up inside that window, the architecture does not fit the team and the evaluation should restart.

The governance capability axis

The 2026 market treats governance as a separate product lane because several softwares center their architecture on controlled content and approval workflows.

Instead of a standalone category, this report treats governance as a scored capability axis applied to every vendor in every category.

Governance capability	What it requires	How to test
Claim-level approval state	Each material assertion carries named-owner approval, timestamp, and evidence link	Retrieve a single answer from six months ago and reconstruct its approval chain in under two minutes
Reuse blocking	Unapproved or expired content cannot move into a live response	Attempt to insert an expired answer into an active RFP. If the system allows it, it fails this test
Audit export	The system produces a complete claim-to-evidence-to-approver report without manual spreadsheet work	Request an export for a completed RFP submitted 90 days ago. Evaluate completeness and time to produce
Expiration and recertification	Time-bound approvals trigger recertification for volatile domains including security, compliance, and pricing	Set an expiration window, wait for it to lapse, and confirm the system flags or blocks reuse
Permission enforcement	RBAC controls apply to reuse and approval, not only to login	Assign a restricted user and verify they cannot access or approve content outside their scope

Governance capability scoring guide

Apply this scoring rubric to every vendor on your shortlist, regardless of primary category.

Most regulated-industry buyers should require a minimum score of 3. Financial services, healthcare, pharmaceuticals, and defense buyers should require 4. Buyers under active regulatory scrutiny should require 5.

No vendor in the current market achieves a clean 5 across all conditions. This is the frontier.

Score	Label	What it means
1	No governance	Content stored without approval state, ownership, or expiration. Reuse unrestricted.
2	Nominal governance	Version tracking and basic approval workflows, but enforcement can be bypassed under deadline pressure.
3	Functional governance	Approval state, ownership, and expiration enforced. Unapproved reuse blocked. Audit exports require some manual assembly.
4	Strong governance	Claim-level approval with evidence linkage, automatic expiration triggers, full audit export, RBAC enforcement across integrations.
5	Production-grade governance	All of the above, plus governance extends to AI-generated content. System blocks AI-drafted claims without source evidence.

Proposal and bid software categories

Each category is defined by the operational constraint it removes for teams responding to Requests for Proposals (RFPs).

The categories are not ranked, and they are not interchangeable. Buying from the wrong one produces a predictable failure.

For each category, the report covers five things, including what qualifies a vendor for inclusion, where the category creates measurable value, where it fails, how it scores on the governance axis, and which vendors currently compete in the space.

Managed Proposal Platforms

Managed proposal platforms run the intake-to-submission operating system for RFPs, RFIs, DDQs, and adjacent requests. They center a curated content library under explicit ownership and enforce coordination discipline across the team that authors the response.

This category fixes coordination failure for teams that already know what good looks like, and the team's process is mature enough to run the response. The platform's job is to stop ten contributors from colliding inside that process.

Administrative work stays with the humans, but coordination overhead drops sharply. The content library is walled and governed, which means every answer the platform surfaces has been reviewed and approved by a named owner. For regulated teams and high-stakes bespoke bids, that control is a feature.

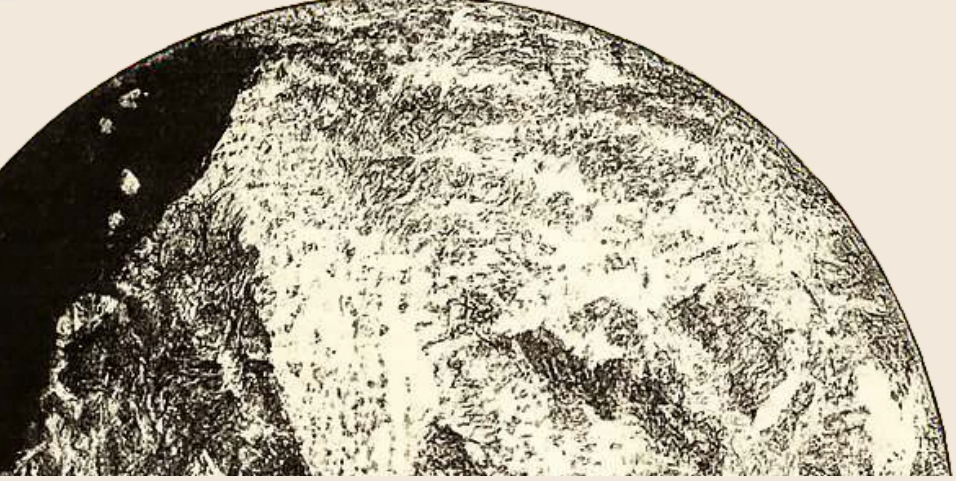

Managed proposal platforms are not a ‘legacy system’ category. Some buyers actively do not want broad retrieval because every new content surface is a new audit surface and a new place a claim can cause response risks. In regulated environments, controlled retrieval is the product.

Inclusion Criteria for Managed Proposal Platforms

A proposal and bid software qualifies only if it enforces production-grade control across five behaviors:

I

Structured intake and project creation.

II

Requirement coverage as a first-class object.

III

Stage-gated review enforcement.

IV

Hybrid truth rules.

V

Packaging reliability.

Where Managed Proposal Platforms Deliver Value

Managed platforms deliver the most value when the team owns the process and needs coordination and auditability enforcement.

- **They protect a curated library as the system of record.** Every answer the platform surfaces has been reviewed and approved by a named owner. In regulated environments and high-stakes proposals, that control is a required feature, not a limitation.

- **They stop contributors from colliding.** Versioning and late-cycle contradictions drop sharply when the system enforces ownership and stage-gated review.

- **They enable disciplined selectivity.** Win-rate improvement correlates with formal go/no-go discipline. Visibility into intake volume and effort cost strengthens a team's prioritization discipline.

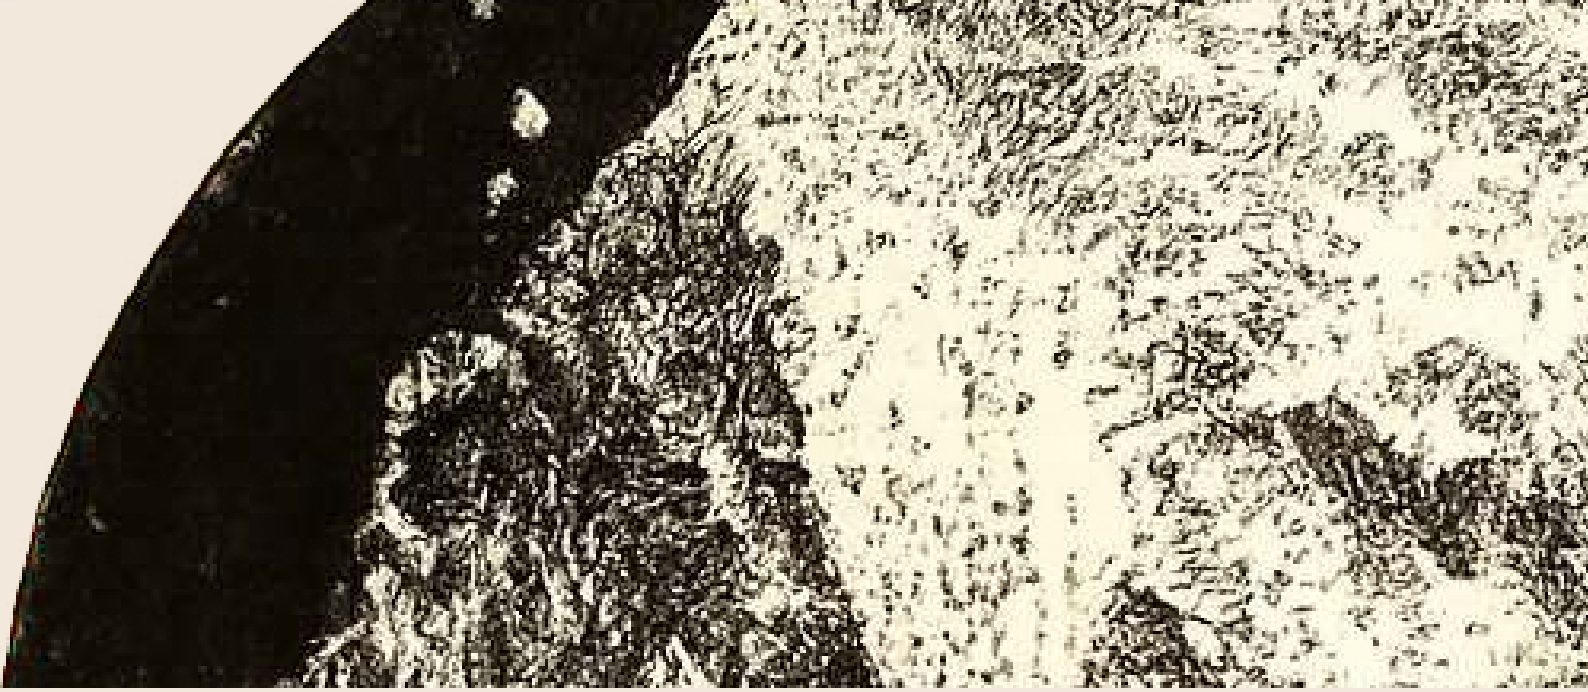

Where Managed Proposal Platforms Can Fail

01

Green-dashboard compliance

Task completion does not equal claim readiness. When a response is flagged complete without real verification or review, AI acceleration can widen that gap rather than close it.

02

Recertification debt

Curated libraries carry heavy ongoing maintenance cost. When library growth outruns review capacity, stale answers start reaching evaluators.

03

Permission boundary fractures

Undefined hybrid project and content ownership leads to overexposure of sensitive artifacts or loss of retrieval trust. When this happens, teams revert to email and offline documents, and the coordination gains the platform delivered disappear.

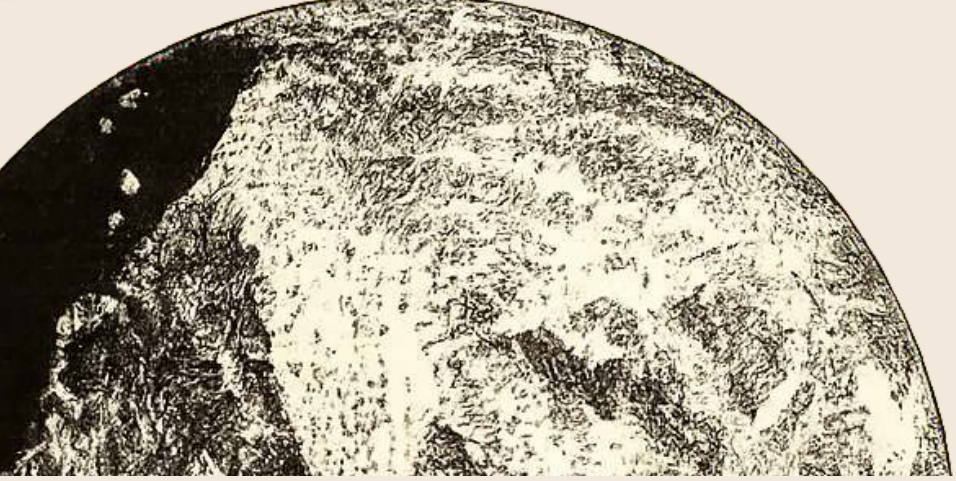

Governance Axis Score for Managed Proposal Platforms

Managed platforms vary widely on governance maturity.

Vendors that grew from library-centric architectures (RocketDocs, QorusDocs, and Responsive) tend to carry stronger approval and content-state controls. Other managed platforms vary in governance maturity and may require additional configuration in regulated environments.

Buyers in regulated industries should run the governance tests from the governance capability axis during pilot.

Software in Managed Proposal Platforms

- Loopio
- Proposify
- QorusDocs
- Responsive
- RocketDocs
- Tendium
- XaitPorter

Governance-forward vendors that compete on the strength of their approval and content-state architecture include QorusDocs and RocketDocs within this category.

Governance-forward positioning is cross-category.

Signals to Watch for Managed Proposal Platforms

I

Governance as the product, not a configuration.

The Managed vendors grown from library-centric architectures will reposition in 2026–2027 as systems of record for compliant claims, with proposal workflow as how that product gets delivered. Watch for pricing that separates governance seats from proposal seats, and sales motions that lead with compliance and legal buyers.

II

The crossover move.

Every major Managed platform will add autonomous authoring within 18 months. The question that matters is whether they redesign the approval layer for AI-generated content. Platforms that bolt AI drafting on top of a workflow built for human authors are carrying technical debt that will surface in the first audit.

III

The library becoming an approved-claim registry.

The winners will reposition the curated library as the team's registered, provenance-linked set of claims. The vendor's job becomes keeping that registry fresh and provable. Watch for language shifts in how Managed platforms describe their content repositories over the next two quarters.

IV

Reviewer-layer ergonomics.

When AI handles drafting, what reviewers do at their desks becomes the user-facing product. Delta-highlighting on regenerated content, confidence scoring on claims, source previews alongside text, disagreement logging between reviewers. The Managed platforms that invest in reviewer tooling will win enterprise deals over the next 24 months.

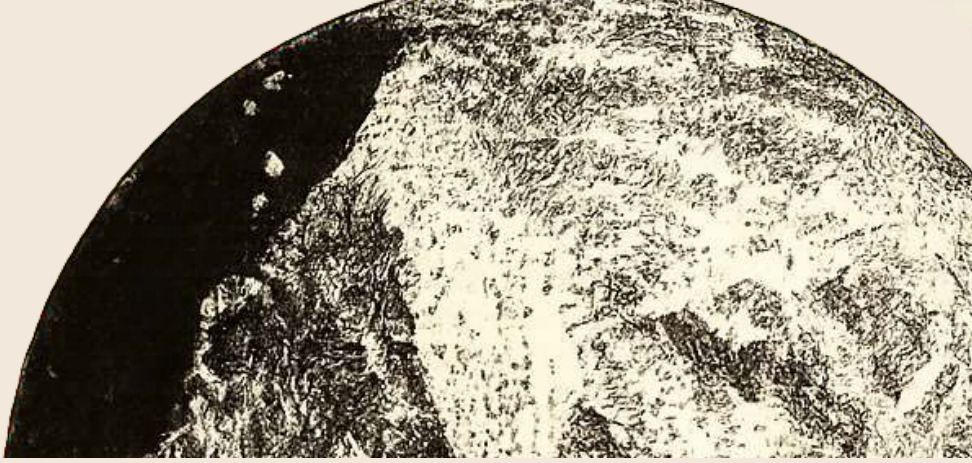

Autonomous Proposal Platforms

Autonomous proposal platforms also run the intake-to-submission operating system, and they absorb the process admin. Judgment still belongs to the human, which data the response needs, how the approval process runs, and what wording wins with a specific evaluator.

This category fixes operating lift for teams where admin, not writing, is the bottleneck. A small team running 150 RFPs a year feels it first, but the same pattern appears in growing teams that are under-processed and in enterprise teams where hours spent on customer research, assignment routing, and even qualification decisions exceed hours spent on authoring. The authoring capacity is there; the admin load is absorbing it.

The content philosophy differs from Managed proposal platforms, as well. Autonomous platforms can also lean on connected rather than a library that can only ever live in the proposal platform. This expands the content surface the team can draw from but also expands the surface the team must govern.

That governance question is the central risk of the category, which is why abstention behavior and claim-level citation are non-negotiable during evaluation.

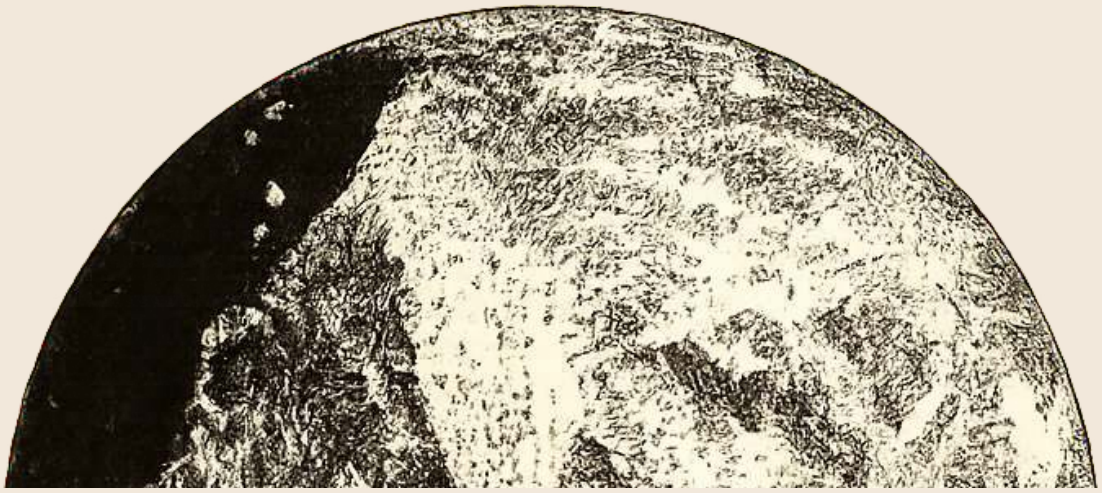

Inclusion Criteria for Autonomous Proposal Platforms

A proposal and bid software qualifies in this category only if it enforces production-grade AI-driven control across five behaviors:

I

AI-driven intake and project creation.

The system parses incoming RFPs, RFIs, and questionnaires without requiring a human to structure them first.

II

Default-on agentic drafting.

The platform produces a first-pass response across the requirement set before a human is pulled in, with claim-level source attribution.

III

AI-routed review and approval.

The system identifies which sections need SME, legal, or executive review and routes them with documented reasoning.

IV

Exception-based human loop.

The system escalates to humans when it detects low confidence, missing evidence, or high-risk claims rather than requiring manual assignment for every task.

V

Auditable AI behavior.

The platform exposes what the AI did, what sources it used, and where it abstained, so the team can reconstruct the response pathway after the fact.

Where Autonomous Proposal Platforms Deliver Value

Autonomous platforms deliver the most value when administrative labor, not authoring, is the issue response teams are facing.

01

They absorb the admin layer

Intake parsing, qualification, requirement extraction, retrieval, first-pass drafting, routing with reasoning, reviewer assignment, and packaging all happen without a human organizing them first. Teams redirect those hours to the sections that actually require authoring.

02

They reduce library maintenance by drawing from upstream sources

Connectors to SharePoint, Confluence, product documentation, CRM records, and security repositories let the platform retrieve at response time rather than relying on a curated library the team must maintain.

03

They scale under volume pressure without linear headcount growth

The operating model changes shape. Teams absorb volume spikes without temporary contractors or weekend work.

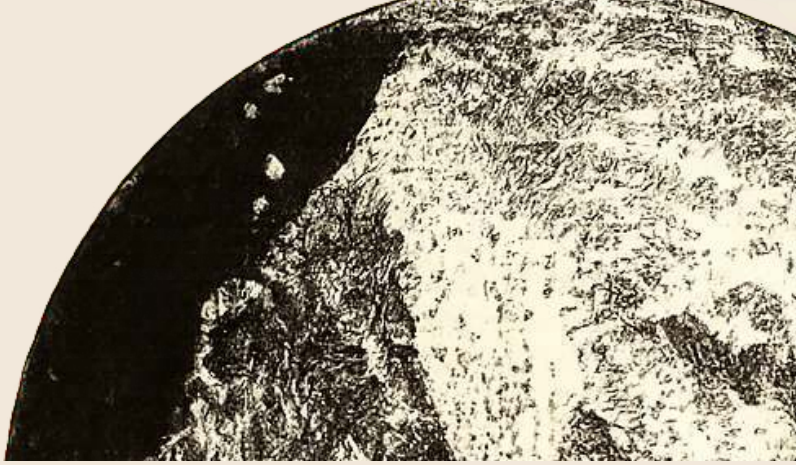

Where Autonomous Proposal Platforms Can Fail

01

Review latency collapse

AI-generated drafts can land in review faster than reviewers can absorb them. When they do, the human loop becomes the necessary bottleneck the platform was meant to solve.

02

Confidence calibration

Autonomous platforms that do not expose abstention behavior transfer risk to the reviewer. Buyers should test what the system does when content and evidence are thin.

03

Governance shadow

If the platform cannot show which human approved which claim before it reached the customer, regulated buyers inherit an audit issue.

04

Retrieval surface sprawl

Broad upstream connectors expand the content surface the team can draw from, but also expand the surface the team must govern. When integration permissions or source-of-truth rules are unclear, the platform surfaces content that a reviewer would have rejected if they had seen it first.

Governance Axis Score for Autonomous Proposal Platforms

The control model assumes the AI will act first and a human will review second, which makes the audit trail the primary evidence a buyer has that the system behaved safely. Software that embed claim-level citation, abstention behavior, and reviewer approval state from the architecture up tend to score higher than software that retrofit governance onto a speed-first engine.

Buyers in regulated industries should run the governance tests from the governance capability axis during pilot and weight abstention behavior heavily in the scoring.

Software in Autonomous Proposal Platforms

Anchor, BidScript, SiftHub, and Tribble are governance-forward within this category.

Signals to Watch for Autonomous Proposal Platforms

Four positioning moves that will separate the Managed platforms repositioning for the next cycle from those carrying forward the last one.

I

Governance as the product, not a configuration.

The Managed vendors grown from library-centric architectures will reposition in 2026–2027 as systems of record for compliant claims, with proposal workflow as how that product gets delivered. Watch for pricing that separates governance seats from proposal seats, and sales motions that lead with compliance and legal buyers.

II

The crossover move.

Every major Managed platform will add autonomous authoring within 18 months. The question that matters is whether they redesign the approval layer for AI-generated content. Platforms that bolt AI drafting on top of a workflow built for human authors are carrying technical debt that will surface in the first audit.

III

The library becoming an approved-claim registry.

The winners will reposition the curated library as the team's registered, provenance-linked set of claims. The vendor's job becomes keeping that registry fresh and provable. Watch for language shifts in how Managed platforms describe their content repositories over the next two quarters.

IV

Reviewer-layer ergonomics.

When AI handles drafting, what reviewers do at their desks becomes the user-facing product. Delta-highlighting on regenerated content, confidence scoring on claims, source previews alongside text, disagreement logging between reviewers. The Managed platforms that invest in reviewer tooling will win enterprise deals over the next 24 months.

Distribution Partner Perspective

AutoRFP.ai

For many revenue leaders, the central constraint in scaling RFP response is control.
Across growth-stage and enterprise teams, proposal and response teams are getting to a first draft quickly, but they are struggling to get the systems that support a winning proposal.

Too many RFPs enter their sales funnel, and the ownership and process is unclear or isn't supported by their software. This often leads to more work, but less revenue, overworked SMEs, and unpredictable pipeline.
This is why many teams are reassessing what they truly need from AI in response work.

The focus sits on running the response operation as an AI-supported system. Customers typically come from drafting-only tools, legacy proposal platforms, or spreadsheet-driven workflows. They move to AutoRFP.ai because their previous toolset accelerated writing without supporting the imperative work that surrounds it. AutoRFP.ai addresses the full response lifecycle of qualification, assignment, writing, review, approval, and reporting.

To win consistently, proposal, solution, and response teams need upstream go/no-go discipline. They need clear ownership, visible deadlines, and capacity awareness. AutoRFP.ai embeds those controls directly into the workflow, so teams make the right decisions from the start and are supported throughout the proposal lifecycle.

We are proud AutoRFP.ai is sponsoring this report because it reflects the autonomous proposal operational and response view of proposal technology. We're finally seeing truly automated workflows in our autonomous proposal automation systems. And we believe this will be the true indicator of teams who have systemized winning more RFPs.

Explore AutoRFP.ai to evaluate proposals earlier, manage proposals in one system, and gain visibility into capacity.

Follow AutoRFP.ai on LinkedIn for practical perspectives on RFP operations, qualification discipline, and scaling response teams without losing opportunities.

AI-native Proposal Drafting Engines

AI-native proposal drafting engines are one of the fastest-moving categories in proposal technology, and the clearest example of focused, narrow-scope AI winning against broader enterprise platforms.

These vendors take a defined scope, encompassing the requirement-to-draft layer, where an inbound RFP, questionnaire, or DDQ becomes a grounded, cited first draft a reviewer can approve or lightly correct before submission.

This narrow scope is the category's strategic advantage.

Drafting engines deploy inside an existing stack. They connect to the Google Drive, SharePoint, Confluence, and Notion repositories where enterprise knowledge already lives, and they integrate with the web portals and office suites where proposal work already happens.

For buyers whose binding constraint is writing throughput, this integration model shortens time-to-value from quarters to weeks.

The category test is straightforward. If the binding constraint is writing throughput, a drafting engine fits.

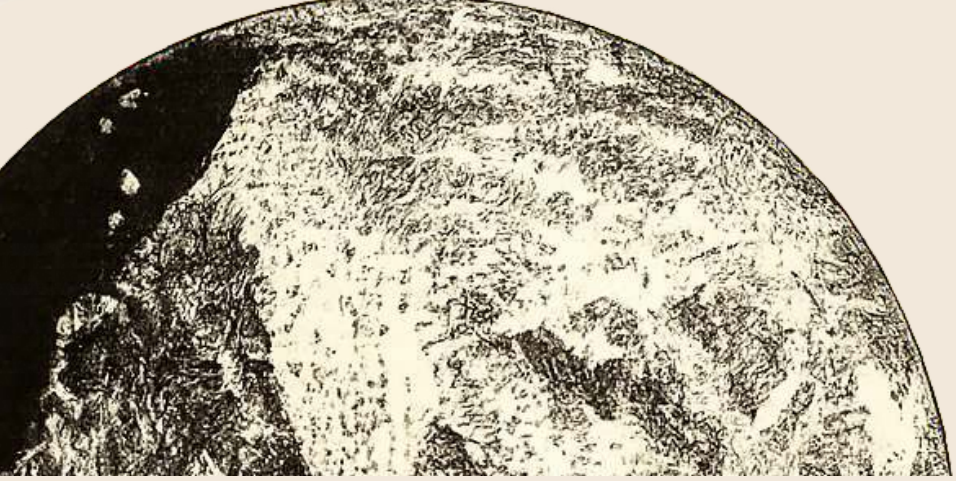

Inclusion Criteria for AI-native Proposal Drafting Engines

A proposal and bid software qualifies in this category only if it enforces production-grade AI-driven control across five behaviors:

I

Requirement-object parsing.

The system converts an inbound RFP, questionnaire, or DDQ into a structured list of discrete requirements that can be answered individually, rather than treating the input as an undifferentiated document.

II

Source-grounded generation.

The engine drafts responses by retrieving from connected internal repositories (Google Drive, SharePoint, Confluence, Notion, past proposals, product documentation) rather than generating from model training data alone.

III

Claim-level citation at generation time.

Every material claim in the draft carries a source link produced during generation, not added post-hoc. A reviewer can trace any statement back to the passage and document that supported it.

IV

Abstention under weak evidence.

The platform surfaces gaps rather than filling them with plausible-sounding content. When the connected sources do not support an answer, the engine flags the requirement for human input instead of fabricating.

V

Review-ready export.

The draft exits the platform in the submission format the buyer requires (Word, Excel, web portal field, PDF) with source traceability preserved through export, so the reviewer can approve, edit, or reject without losing the evidence chain.

Where AI-native Proposal Drafting Engines Deliver Value

Drafting engines own the requirement-to-draft layer. They plug in alongside whatever workflow and system of record a team already runs, and they accelerate turning a requirement into a grounded, cited first draft.

01

SME time recovery for technical sales and security-response teams

Drafting engines pull from product documentation, security policies, and past responses, and generate cited answers in the flow of work. A browser extension can handle web portal questionnaires directly, and a Slack integration turns the knowledge base into an on-demand answering surface for anyone on the deal team. The measurable lift is SME hours recovered per week.

02

Faster time-to-value for small and mid-sized proposal teams

Teams of 1 to 10 FTEs often cannot justify the seat count, implementation effort, and ongoing administration a Managed platform requires. Drafting engines price and deploy for this tier.

03

Drafting acceleration for teams with working process discipline

Some teams already run a good process. Ownership is clear and reviews happen on time, but what they lack is speed at the drafting layer, because their knowledge sits spread across Google Drive, Confluence, Notion, Slack, SharePoint, and past proposals. Drafting engines retrieve across those sources without asking the team to migrate content into a new library. The team keeps the process it has and layers in AI drafting where it matters.

Where AI-native Proposal Drafting Engines Can Fail

Drafting engines own the requirement-to-draft layer. They plug in alongside whatever workflow and system of record a team already runs, and they accelerate turning a requirement into a grounded, cited first draft.

01

Review overload

Fluent text without enforceable evidence shifts validation onto SMEs under deadline pressure. Writing speed rises, review time rises faster, and cycle time goes nowhere.

02

The buyer-detectability tax

Evaluators recognize machine-written prose even when they cannot name what tips them off. A proposal that reads as generic loses ground against a thinner response that reads as considered. Faster can be visibly worse.

03

Jagged frontier misuse

Drafting engines produce strong output on answer-type questions (capability statements, technical Q&A, security controls) and weak output on strategy-type questions (win themes, buyer-specific framing, proof selection, commercial narrative). Teams that apply the engine uniformly across both gain throughput while losing deals on the strategy side.

04

Acceleration without reallocation

The investment case assumes recovered SME and writer time flows into customer conversations, qualification, and commercial strategy. When it flows into responding to more RFPs instead, the tool produces output, not revenue.

05

Commoditization pressure

Retrieval plus generation with citations is now available inside ChatGPT Projects, Claude Projects, and Microsoft 365 Copilot at a fraction of enterprise pricing. Engines competing on retrieval alone face direct substitution. The defensible moat sits in claim-level governance enforcement, workflow integration with live submission processes, and domain-tuned behavior that frontier models do not match out of the box.

Governance Axis Score for AI-native Proposal Drafting Engines

Governance in AI drafting is the set of controls that make generated content defensible after submission. Proposals are often legal commitments and security questionnaires feed contract terms, so when an AI-generated claim is wrong and the buyer relied on it, the vendor is responsible regardless of who or what wrote the words.

That reality sets the bar. Governance-forward engines enforce citation at generation time, abstain when evidence is weak, and preserve source traceability through export. Engines without those three behaviors move the validation burden onto reviewers under deadline pressure, which is the point in the cycle where human oversight is weakest.

The operational test for buyers in this category is straightforward. Run a complex historical RFP through the tool during pilot with one trusted source deliberately removed from the ingest. A governance-forward engine abstains on the affected claims and surfaces the gap. A cosmetic one produces a confident draft that looks finished but contains unbacked claims.

The test is cheap, runnable before purchase, and the single strongest predictor of how the tool behaves under real deadline pressure.

Engines that score below 3 on the governance axis carry material risk in regulated environments for this reason.

Convergence with managed and autonomous platforms

Our assessment is that the distinction holds for approximately 18 more months. By late 2027, vendors on both sides will have absorbed enough of the other's control surface that the boundary becomes a feature-maturity gradient, not a category line.

Buyers making a purchase in 2026 should still apply the category distinction but should ask vendors about their roadmap for the adjacent control surface. At this point, we believe one of the main differentiators, in addition to governance, will be true and useful automation at each stage of the process; categories will most likely change quickly.

ANALYST FLAG

Q1 2026

Convergence clock flag

This 18-month estimate is a forward-looking market call based on current roadmap signals and feature convergence patterns.

Stargazy will reassess this timeline in Q3 2026 and publish an updated position in the quarterly signal briefing. If the estimate proves wrong, it will be revised in the 2027 edition with an explicit correction.

Buyers and vendors should track which AI-native vendors add project management features and which management platforms add retrieval-native drafting, quarter by quarter.

Software in AI-native Proposal Drafting Engines

1up	AnswerTree	Arphie	AutogenAI
Inventive AI	Iris	SEQUESTO	Seev
Tenderseal	Tomorro	Vera	

Governance in AI-native Proposal Drafting Engines

For commercial teams, 1up, Arphie, and Iris are governance-forward, each competing on a different architectural governance behavior.

1up operates a library-less architecture that indexes connected sources (Google Drive, Confluence, Notion, SharePoint) and generates with citations and anti-hallucination guardrails, so there is no static answer bank to go stale between approvals. Arphie provides source transparency and reasoning visibility on every answer. Iris enforces inline source citations with confidence scores and pairs them with dynamic retrieval from connected documentation, avoiding the static answer bank pattern that creates recertification debt. All three carry SOC 2 Type II.

For B2G teams, AutogenAI is the governance-forward vendor within this category, carrying a deep compliance stack (FedRAMP High, ISO 27001, SOC 2 Type II, single-tenant deployment, and contractual prohibition of model training on customer content).

Signals to Watch in AI-native Proposal Drafting Engines

1. Workflow redesign becomes the dividing line. The single strongest predictor of EBIT impact from generative AI is whether the organization fundamentally redesigned its workflows around the tool, not the sophistication of the model. The drafting engine market will split on this axis. Vendors that keep AI as a bolt-on to existing response workflows will see flat adoption and flat revenue impact for their customers. Vendors whose customers redesign intake, SME engagement, and review around AI drafting will produce the case studies everyone else cites. Watch for customer stories that show process change, not just time savings.

2. Buyer-side genAI forces vendor-side evidence. A majority of B2B purchase influencers already use or plan to use private genAI engines in procurement, and a measurable share of buyers report less confidence in decisions after using AI to evaluate vendors. Evaluators will increasingly send RFP and questionnaire responses through their own AI to fact-check claims. Drafting engines that cannot produce claim-level citations with verifiable source links will fail buyer-side verification even when they pass internal review. Watch for vendors publishing third-party accuracy benchmarks and citation-coverage metrics in 2026-2027.

3. Guardian agents enter the proposal stack. Guardian agents, a newly defined category of AI systems that supervise other AI systems, are projected to capture 10 to 15 percent of the agentic AI market by 2030. Inside proposal workflows, guardian agents will audit drafting engines for hallucinated claims, unbacked commitments, and regulatory exposure before a submission leaves the building. Drafting engines that expose their reasoning, source-binding, and abstention signals through open APIs will integrate with guardian layers. Closed drafting engines will find themselves routed around. Watch for guardian-ready certifications and multi-agent orchestration patterns emerging in vendor roadmaps.

4. Productivity-as-priority reframes the buying case. Productivity is the number one revenue-leader growth strategy for 2026, up from fourth place the prior year, and teams using revenue-specific AI produce 77 percent more revenue per rep. The buying committee for drafting engines will shift accordingly. Procurement, IT, and compliance will remain involved, but the CRO and Chief of Staff to the CRO will increasingly own the business case. Drafting engine vendors that pitch time savings will lose to vendors that pitch reps-freed-for-customer-conversations. Watch for pricing repositioning around per-rep productivity instead of per-seat licensing.

5. The workslop tax becomes a budgeted line item. AI-generated workslop (fluent but low-value AI output) costs approximately two hours of rework per instance. Inside proposal teams, workslop shows up as drafts that look complete but require reviewers to rewrite from scratch.Drafting engines that produce review-ready output will pull ahead of engines that produce fast-but-unusable output. Watch for measurement shifts in proposal operations reporting.

Distribution Partner Perspective

For many mid-market teams, the central constraint in using AI for RFPs and questionnaires is trust.

Across response teams, we consistently see the same pattern. Many tools promise speed but deliver answers that are subtly wrong or unjustifiably confident. The result is not efficiency but escalation. In addition to proposal that reads like "AI slop," it will create proposals that require more reviews and rework, and risk that compounds with every response cycle.

This is why many teams are reassessing what they actually need from AI-assisted response tools.

At 1up, the focus is on producing answers that teams can rely on. Their customers typically come from RFP response platforms or internally built GPT workflows because they are rejecting tools that fabricate confidence without evidence.

Mid-market teams do not have the time or appetite for months of change management. They need tools that work immediately, integrate cleanly with existing knowledge sources, and fit naturally into how people already respond to questions.

1up prioritizes fast onboarding and immediate value, so teams can move from connection to usable answers without heavy process redesign.

We are proud 1up is sponsoring this report because it reflects a more mature conversation about AI in proposal work. The next phase of this category is about producing answers that teams trust enough to use.

For organizations navigating that transition, especially in the mid-market, clarity matters more than novelty. The insights in this report are intended to help buyers make that distinction and choose tools aligned to how they actually work, not how AI is marketed.

Start a free trial of 1up to connect your real knowledge sources and generate answers you can trust, without long implementation cycles.

Follow 1up on LinkedIn for practical takes on RFP reality, product updates, and the daily meme.

The AI acceleration effect

Drafting speed no longer limits how much a team can write, but speed without structure is a liability.

AI tool adoption shows no independent predictive power on win rates once structural and process variables are controlled, meaning platforms that increase output without reducing unverified claim rates create a measurable failure mode, where SMEs spend more time reviewing and rewriting content, later in the cycle and under deadline pressure.

McKinsey's State of AI research found that high-performing organizations are nearly three times as likely to redesign workflows around AI rather than layer automation on top.

Stanford's Human-Centered AI Index found that teams where people actively guide AI outputs see productivity gains of 30 to 35 percent. Teams where automation replaces oversight see far smaller returns.

MIT's review of over 100 human-AI collaboration studies confirmed the combination only outperforms when each side handles what it does better than the other.

> “AI replacing proposal teams is like handing someone a scalpel and calling them a surgeon. The tool doesn't do the work… the expertise does.
>
> In GovCon, a proposal is often the last conversation you have before someone decides to trust you with millions of dollars. And that someone is an evaluator, a real person trying to make a confident buying decision. They don't award contracts to the most efficient process. They award them to the team that made them feel most sure they picked the right partner. Strip out the humans who know how to earn that confidence, and you've already lost.
>
> The firms treating AI as headcount replacement are going to find out the hard way that a faster, cheaper proposal is still a losing one.”
>
> — KRYSTN MACOMBER, FOUNDER AND CEO, SUMMIT STRATEGY

The pattern holds in proposal operations. When executive teams remove proposal ownership roles and replace them with automation software, adoption decays because nobody maintains permissions, sources, or review design.

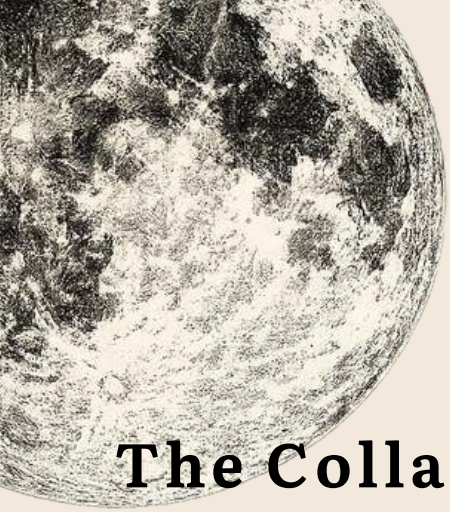

The Collapse of Manual Library Economics

Manual libraries carry a compounding maintenance tax. Each new answer increases future review load because the organization must keep the answer aligned with ever-changing policy, legal, security, and product updates.

Stargazy implementation evidence (2025-2026) records two observable outcomes.

Retrieval-first implementations often retire large portions of legacy libraries, with reductions of 40 to 70 percent in year one.

> "Humans are essential in the loop of answering an RFP because they bring strategic judgment that automated systems cannot replicate. Their knowledge of the specific customer business unit, the broader organisation, and the industry provides critical context that strengthens the response. They leverage relationships and insight into unstated challenges or priorities, adding value beyond what is explicitly written in the RFP. This deeper understanding enables more tailored, relevant, and persuasive proposals. As a result, human input creates a meaningful competitive advantage that improves the chances of winning."

ALESSANDRO MUCCIARDI
VP GTM SERVICES SALES · SALESFORCE

Library-dependent environments often expand repositories, with some reporting annual growth above 20% to preserve coverage of product or policy changes.

Those trajectories create opposite economics, as one narrows the surface area teams must govern, and the other expands it.

GovCon capture-to-proposal platforms

Government contracting is the most structurally distinct proposal environment in the 2026 market, and the category moving fastest on AI-native architecture.

FAR/DFARS compliance, CUI handling, evaluator-scoring realism, capture-to-proposal context continuity, past-performance evidence requirements, and government portal submission create a control surface that commercial proposal tools cannot absorb.

Three changes are reshaping the category right now:

- **Federal agencies are now using AI to evaluate proposals.** OMB Memoranda M-25-21 and M-25-22, GSA's 2026 AI directive, and an April 2026 GAO oversight report confirm that DoD, DHS, GSA, and VA are running AI-assisted proposal evaluations on major acquisitions. Proposals are now read by machines on the agency side, so structured compliance matrices, explicit requirement mapping, and citation-linked past performance matter more than prose fluency.
- **The category is splitting between AI-native BD operating systems and traditional tools with AI bolted on.** One platform owns opportunity discovery, capture, proposal, and contract management. Buying a standalone drafting tool and plugging it into GovWin IQ is becoming the losing pattern.
- **FedRAMP High is the new sponsor-tier differentiator.** Moderate was the bar 18 months ago. High is the bar now for DoD, intelligence community, and sensitive CUI work. Achieving High takes 18 to 24 months and seven figures in compliance cost, which creates a real moat. Vendors that hold it can sell into environments that others legally cannot enter.

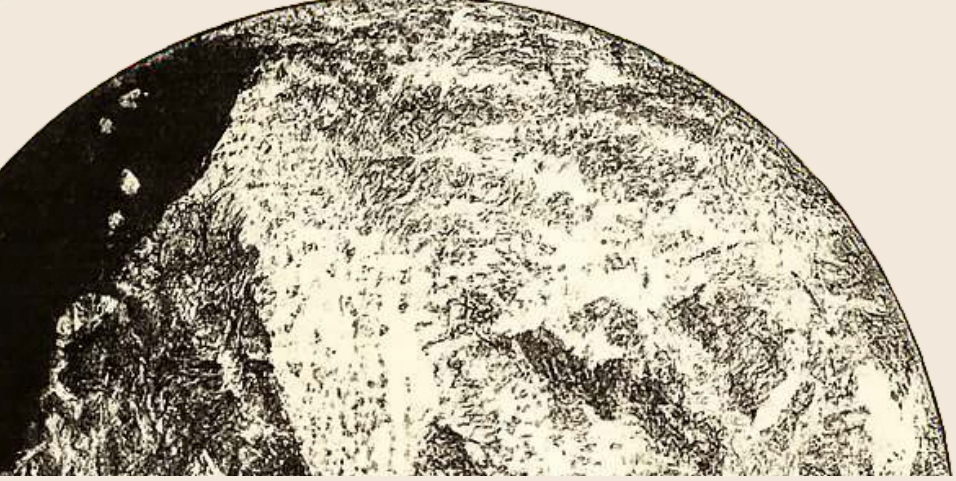

Inclusion Criteria for GovCon capture-to-proposal engines

A proposal and bid software qualifies in this category only if it enforces US government contracting discipline across five requirements.

I

Regulatory awareness.

The platform understands FAR, DFARS, or equivalent regulatory frameworks and can map solicitation requirements to compliance obligations without manual cross-referencing.

II

Capture-to-proposal continuity.

Pursuit intelligence feeds directly into proposal content and the compliance matrix.

III

Evaluator-realism capability.

The system can produce or assess responses against evaluation criteria published in the solicitation.

IV

Structured evidence for personnel and past performance.

The platform manages or integrates with structured data sources for key personnel and experience records.

V

Submission-format compliance.

Outputs meet government portal requirements.

Where GovCon capture-to-proposal engines deliver value

GovCon capture-to-proposal platforms own a workflow that commercial proposal tools cannot absorb. FAR/DFARS compliance mapping, Section L/M parsing, CUI handling, past-performance evidence management, and capture-to-submission context continuity sit inside the product rather than layered on top. The platform is purpose-built for the federal and public-sector buyer.

01

Compliance mapping at the speed of the deadline

Platforms in this category parse these structures automatically and generate compliance matrices that track every requirement to a named response owner.

02

Capture-to-proposal continuity for contractors with pipeline discipline

Mid-market and enterprise federal contractors running structured capture processes need pursuit context to flow directly into the proposal, including customer intelligence, competitive positioning, win themes, teaming arrangements, and past-performance selections. Platforms that carry this context from opportunity signal through final submission reduce rework, improve win themes, and produce proposals that read as informed by the pursuit rather than assembled at the last minute. This is where the BD operating system model beats standalone drafting tools.

03

Past-performance and personnel evidence at evaluator-grade quality

GovCon evaluators score past performance and key personnel against specific criteria with strengths and weaknesses assessments and color-coded standards. Platforms that manage structured personnel records (clearances, certifications, relevant projects) and past-performance databases (CPARS scores, contract values, period of performance) produce proposals that match evaluator scoring logic rather than generic capability statements. Vendors including Rohirrim, pWin.ai, Ensis AI, and Flowcase operate at this evidence layer.

Where GovCon capture-to-proposal engines can fail

01

Over-engineering for compliance, under-engineering for drafting

Compliance mapping without strong drafting forces teams back to manual writing with a compliance overlay. The matrix gets built, the requirements get tracked, and the team still spends three weeks writing narratives from scratch. Buyers should test drafting quality on a real Section L task order, not on a vendor demo.

02

False evaluator confidence

AI-assisted scoring simulation is valuable but imperfect. Vendors that produce confident-looking "evaluator scores" on draft sections without disclosing the model's training data or confidence bounds set teams up to submit proposals their own platform predicted would win. Agencies score differently, and no commercial AI has access to agency-specific evaluation notes.

03

CUI boundary failures

Most CUI incidents in AI proposal tools happen at integration boundaries, not inside the platform. A contractor uploads a CUI solicitation to a FedRAMP-authorized platform, then pulls reference content from a SharePoint library that sits in a commercial tenant. The CUI crosses a boundary the platform cannot enforce, and the audit trail stops at the platform edge. Buyers should map every content source the platform reads and verify each one sits inside the same authorization boundary.

04

Agency-generated proposals fail the sniff test

Federal executives and evaluators can detect AI-generated proposals that lack human oversight, and they read generic win themes, boilerplate past-performance selection, and evaluator-facing language recycled from prior awards as signals of inattention. False representations can trigger disqualification or award rescission. The failure mode is not that the AI writes badly. It is that the team ships AI output without the capture-informed editing that turns a draft into a pursuit-specific proposal.

05

BD operating system lock-in

The platforms consolidating capture, proposal, pricing, and contract management into one system create real value and real switching cost. A contractor that moves its pipeline, past-performance library, teaming database, and compliance templates into a single vendor stack has made the vendor a critical dependency. If the vendor is acquired, repriced, or pivots, the unwind cost is substantial. Buyers should negotiate data portability and export rights before signing, and should treat the BD operating system as infrastructure, not software.

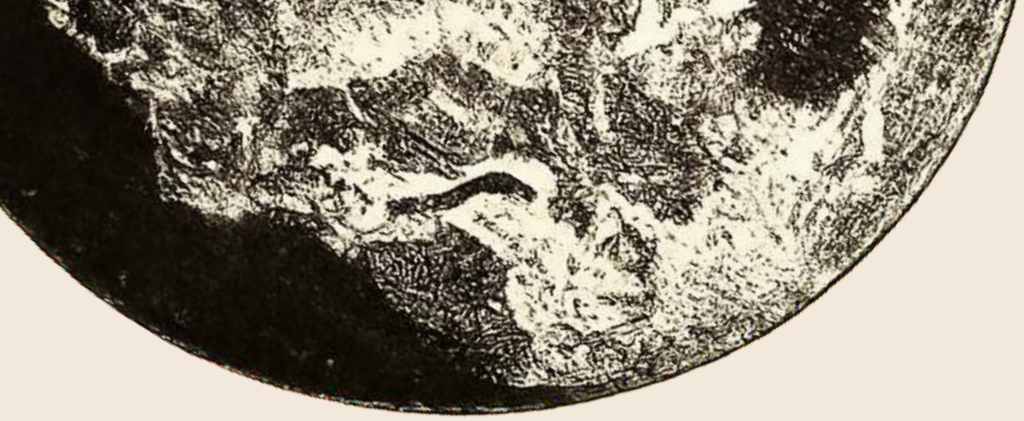

Governance Axis Score for GovCon capture-to-proposal engines

Governance in GovCon proposal work is the set of controls that make a proposal defensible in three high-stakes scenarios: a False Claims Act investigation, a CUI handling audit, and a competitor's protest of the award decision.

Federal contractors carry personal liability for signed proposal content under the FCA, and CUI mishandling can trigger suspension and debarment across the entire federal marketplace. The governance bar is set by legal exposure.

Every substantive claim traces to a named human approver with an audit trail that survives FOIA, protest review, and FCA investigation.

CUI boundary enforcement at the data layer requires permission enforcement as an architectural guarantee, not a platform setting, and CUI cannot cross from a FedRAMP-authorized environment into a commercial tenant through any retrieval or integration path.

Evidence chain reconstruction for past performance means CPARS references, contract values, and period-of-performance data must reconstruct back to source agency records.

The operational test for GovCon governance is a mock protest:

Run a historical winning proposal through the platform and ask it to reconstruct, for every material claim, the source document, the named approver, the approval timestamp, and anything else relevant to your situation.

A governance-forward platform produces this in minutes with exportable audit logs. A platform that cannot produce it under controlled conditions cannot produce it under an actual protest or FCA investigation either.

Software in GovCon capture-to-proposal engines

- AutogenAI
- CLEATUS
- Deltek/GovWin IQ
- Ensis AI
- GovDash
- GovEagle
- GovSignals
- Procurement Sciences
- pWin.ai
- Rohirrim
- Sweetspot
- Turingon
- VisibleThread
- Vultron

Signals to Watch in GovCon capture-to-proposal engines

- **SLED market entry becomes the growth test.** Federal market growth is structurally capped by agency budgets and procurement cycles, but state, local, and education procurement is larger in aggregate and far less consolidated. SLED volume is projected to grow faster than federal through 2028. Many in this category have already moved into SLED. Watch for vendors adding state procurement portal integrations, SLED-specific past-performance libraries, and education-sector teaming support. The GovCon vendor that owns federal plus SLED plus DIBBS will build the widest moat.

- **Guardian agents reach federal environments first.** Federal compliance requirements (FedRAMP, CMMC, NIST 800-53) create the conditions for guardian agent deployment, where a second AI layer audits the first for hallucinated claims, unbacked commitments, and regulatory exposure before submission. Watch for GovCon vendors adding abstention logging, reasoning-chain export, and compliance-traceable output that guardian layers can audit independently. Vendors that ship closed, opaque agents will be routed around by federal buyers whose security teams require auditability.

- **AI-free author certification enters solicitations.** Some solicitations now require named-author certifications on proposal content, and federal executives have publicly flagged detection of AI-generated proposals as a disqualification trigger. Watch for solicitation language that requires documented human authorship, audit logs of AI contribution versus human contribution, and signed attestations from proposal managers. The platforms that make human-in-the-loop oversight visible and exportable will win the next wave of compliance-sensitive contracts.

- **The OneGov buying vehicle reshapes procurement timelines.** GSA's OneGov strategy (launched under OMB Memos M-25-21 and M-25-22) is designed to accelerate federal AI acquisition through pre-negotiated government-friendly pricing agreements with approved vendors. Vendors on the OneGov list get faster agency procurement paths. Vendors not on the list face longer, more contested procurement cycles. Watch for which GovCon proposal platforms secure OneGov placements in 2026. The first two or three will build meaningful procurement moats against competitors who have to run standard GSA schedules.

Vertical evidence and experience specialists

Some proposal environments are won or lost on structured evidence, especially with personnel credentials, project references, and digital assets determining evaluation outcomes in professional services, engineering, construction, architecture, and legal. Most vertical evidence software is used alongside a full proposal and bid management software.

Flowcase — leads on the CV/resume/case-study system-of-record positioning, with customer evidence spanning four verticals (AEC, IT consulting, management consulting, law). It fits firms where personnel evidence determines evaluation outcomes.

Joist AI — leads on the contextual knowledge graph differentiator, the architectural bet that sets it apart from generic AEC tools, ties the AEC-specific language claim to concrete examples such as design-build, LEED, and delivery methods.

ContraVault AI — leads on clause-level risk detection as the differentiator versus only RFP analysis, with Go/No-Go, Risk Flagging, Compliance Matrix, and Drafting capabilities. They're focused on the material risk exposure positioning that distinguishes it from generic RFP tools.

OpenAsset — leads on the DAM foundation with the AEC-specific additions, such as project-based organization and image similarity),with Shred.ai as an AI layer, and closes with the honest analyst note that most firms deploy it as a DAM alongside dedicated proposal tools rather than as a full proposal platform.

Adjacent market: trust-center and security questionnaire automation

Security questionnaires have become one of the highest-volume document types in B2B selling. SaaS vendors routinely receive SIG, CAIQ, VSAQ, and custom questionnaires from every prospect's security team, often before the commercial conversation reaches a proposal. For some teams, questionnaire volume exceeds RFP volume by a factor of three or four, and the binding constraint on deal velocity sits inside the security review rather than inside the proposal drafting workflow.

Trust-center and security-questionnaire automation platforms emerged to solve this specific problem. They publish a public-facing trust center that preempts inbound questionnaires by making security posture self-serve. They automate responses to questionnaires that still come through, draw from evidence already collected for SOC 2, ISO 27001, and other compliance frameworks, and integrate with the CRM and deal-stage tooling that security-driven deals depend on. SafeBase and Conveyor alongside Vanta serve a different buyer (CISO or security-compliance team), draw from a different budget, and compete in a market with its own evaluation criteria.

They enter the proposal technology conversation only when security-review throughput, not narrative proposal authoring, is the binding constraint on deal velocity.

- **SafeBase** is the category-defining trust center platform, designed to reduce inbound security questionnaire volume by making security posture self-serve. Prospects access a public-facing trust center behind an NDA gate, with certifications, policies, and evidence available without triggering a manual response. AI Questionnaire Assistance handles the questionnaires that still come through, drafting from trust center content and past answers via a Chrome extension for portal autocomplete.

- **Conveyor** is AI-first security questionnaire automation built for cybersecurity-adjacent B2B SaaS, with a sharper questionnaire-response focus than SafeBase. It auto-completes questionnaires, includes a Trust Center AI to answer prospect questions behind an NDA gate, and has a one-click portal autocomplete via browser extension.

- **Vanta** is a full GRC and compliance automation platform with trust center and questionnaire automation as features rather than the core product. It includes automates evidence collection and continuous monitoring across 35+ frameworks (SOC 2, ISO 27001, GDPR, HIPAA, PCI DSS). Vanta's AI Agent drafts questionnaire responses from the customer's existing security program and evidence, and automated SME routing and reminder handling.

Adjacent market: enterprise workflow automation backbones

Workato, Tray.ai, Make, Zapier, and n8n provide general-purpose orchestration for connecting systems, moving data between applications, and automating repeatable workflows across the business.

They solve legitimate workflow problems. Workato and Tray.ai carry enterprise governance (SSO, RBAC, audit trails, SOC 2 and ISO 27001 certification) suitable for regulated environments. Zapier and Make price accessibly for mid-market teams. n8n offers self-hosted deployment for teams with data residency requirements. All of them ship hundreds of pre-built connectors to the systems proposal teams already use.

What they do not do is proposal work. They lack requirement parsing, content library management, claim-level citation, approval state enforcement, SME routing with domain context, or any of the other controls that define the proposal technology categories. A workflow automation platform can move a proposal project between systems. It cannot draft the proposal, enforce the approval chain, or produce an audit trail of who approved which claim. Treating it as proposal technology produces a fast-moving pipeline of unverified content with no governance surface at all.

They belong on the proposal technology shortlist only when the organization already has an automation center of excellence with dedicated staff, existing integrations, and a governance model the proposal team can inherit. In that configuration, extending existing orchestration into proposal workflows is cheaper and more controllable than buying a new platform. In every other configuration, the total cost of ownership (platform license plus integration development plus ongoing maintenance plus governance build-out) exceeds the cost of a purpose-built proposal platform that ships with these capabilities already wired together.

What about Claude Skills, Microsoft Copilot, and Google Gems? These are LLM-native orchestration primitives, not enterprise workflow platforms. Claude Skills let teams package reusable prompts, context, and tool configurations inside Claude for repeatable tasks. Google Gems do similar work inside the Gemini ecosystem. Both are useful for small-scale automation where a single user wants to codify a repeatable workflow (drafting a specific type of response, running a specific analysis pattern, producing a specific output format). Neither ships with enterprise governance, RBAC, audit logging, or the system-of-record semantics that proposal work requires at scale. Teams using these primitives for proposal work are effectively running shadow-IT automation, which creates compliance exposure in regulated environments and makes enforcement of the governance controls impossible.

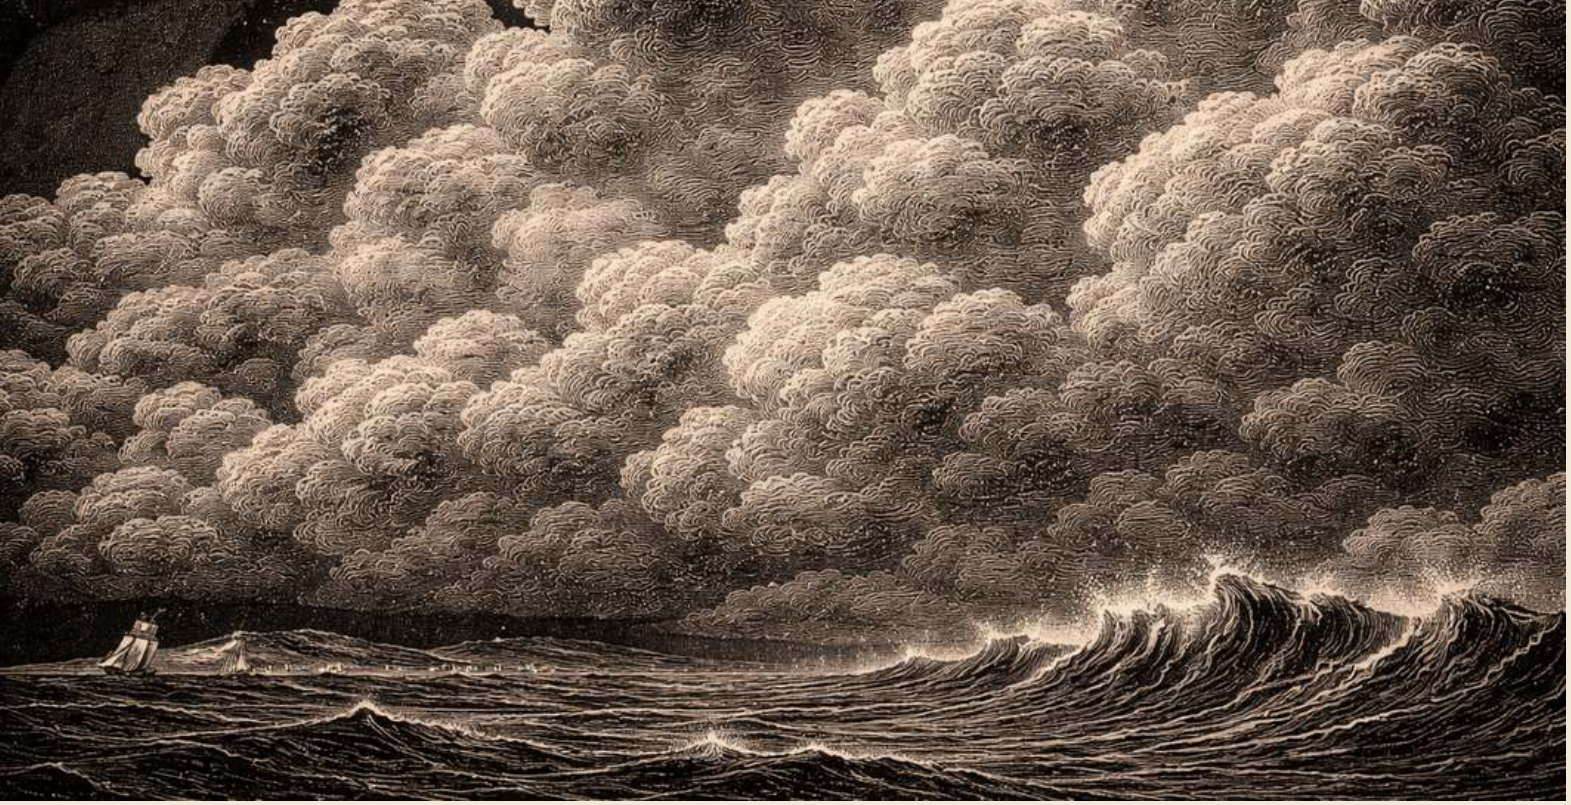

Purchasing by team profiles

Buying the wrong architecture wastes budget and moves work to the most expensive people on the team, later in the cycle, under deadline pressure.

Two questions identify your team profile before you shortlist a proposal and bid software:

1. How many people touch a proposal before it ships?
2. What breaks most often in your current process? Speed, personalization, coordination, governance, or compliance?

The answers map directly to one of six profiles below.

This section maps each team profile to the architectural categories defined. Each profile includes the dominant constraint in a proposal process, the architecture that removes it, what will fail if you buy the wrong one, realistic budget guidance, and a forward reference to the vendor shortlist.

If your profile resolves to either Managed or Autonomous, one additional question will help you decide - Does your team need to own the library and the flow, or do you want the platform to absorb some of the admin and process?

Ad-hoc teams (no dedicated proposal function)

Your situation. Nobody owns proposals. RFPs land on whoever is available. Sales, delivery, or leadership handles them on a rotating basis, and each response starts from scratch, rebuilt from old emails, slide decks, or individual memory. Senior staff get pulled away from revenue-generating work to write documents they will never see again.

Stargazy's Win Rate Report found that 14% of Low Win teams report no dedicated bid role at all. RFPs are handled opportunistically by sales, solution consultants, delivery, or leadership. The cost is senior staff pulled away from core work and lower win rates.

> "Relying on AI only to write bids when there's no bid function creates an illusion of efficiency. If opportunities are not properly pursued, qualified, or planned, there are few wins. I've seen it firsthand — a contractor once proudly told me they only used AI for their bids. Their win rate was 10%. I rest my case!"
>
> PIPPA BIRCH
> FOUNDER · PIPSTER SOLUTIONS
>
> PIPSTER SOLUTIONS

Right fit. AI-native drafting engines are ideal because of their low setup burden. Platforms such as lup tolerate inconsistent inputs and fragmented ownership. Their value lies in absorbing content and process chaos without requiring predefined libraries or dedicated administration.

Wrong fit. Managed proposal platforms and governance-centric systems assume stable ownership and ongoing maintenance. In ad-hoc environments, those assumptions break within weeks.

Common failure mode. Adopting governance-heavy systems that require staffing the tool. If nobody owns the platform after month two, nobody maintains it, which means no one will use it.

Budget guidance. Typically $6k to $30k annually.

Small teams (1 to 3 people)

Your situation. You have a proposal function, but it is one person doing everything or two to three people splitting triage, drafting, and SME coordination on top of other responsibilities. Process-heavy tooling creates drag, if they are not in a highly regulated industry or Government Contracting. What they need is immediate output with minimal overhead.

Right fit. AI-native drafting engines with fast intake and low administrative burden. Retrieval-based ingestion and source-restricted drafting matter more than advanced workflow layers. Platforms such as Iup, Arphie, AutogenAI, AutoRFP.ai, BidScript, Iris, and mytender.io, prioritize coherent first drafts and retrieval without heavy tagging.

> "Most organizations hire their first proposal person to take the paperwork off someone else's plate. That instinct is understandable, but it sets the role up to fail because a proposal function that's treated as administrative support will always be reactive and under-resourced, leading to low win-rates and a burned-out RFP writer. However, when positioned as a strategic function with authority, alignment to sales goals, and leadership buy-in that RFPs are a company priority, it becomes one of your most valuable sources of competitive intelligence available."
>
> 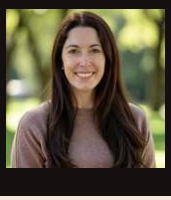
>
> STEPHANIE PEACE BLACK, APMP-CF
> PROPOSAL CONSULTANT

Wrong fit. Governance-centric and library-centric platforms assume dedicated ownership for content upkeep and workflow management. For teams of one to three, this overhead competes directly with writing time. The system asks for more attention than the team can give it, and adoption stalls before the first quarter ends.

Common failure mode. Buying systems designed for coordination at scale when the real constraint is writing capacity. A team of two does not have a coordination problem. It has a throughput problem

Budget guidance. Typically $29k to $50k annually, depending on response volume and contributor count.

Medium teams (3 to 10 people)

Your situation. Drafting speed is no longer the constraint. Coordination is. Duplicated edits, inconsistent answers, and version conflicts multiply as contributor count increases. Time bleeds into alignment and rework rather than writing. If your team has ever submitted a proposal with two different answers to the same question, you are in this profile.

Right fit. Managed or Autonomous proposal platforms, or AI-native drafting engines with integrated workflow. AutoRFP.ai (autonomous) and 1up (AI-native drafting) reduce SME burden by producing credible drafts and routing only unresolved questions for review. Flowcase addresses personnel bottlenecks in experience-led bids.

Conditional Fit. Library- and workflow-centric platforms such as QorusDocs and RocketDocs can perform well, but only with a dedicated content owner responsible for library hygiene and governance.

Common failure mode. Purchasing fast drafting tools without structured review control, leading to expanded reviewer queues and late-cycle rework.

Budget guidance. Typically $50k to $100k annually.

Teams in highly regulated industries

Your situation. Regulatory exposure and audit defensibility is the biggest focus. These teams operate under industry-specific regulatory frameworks (HIPAA, FDA 21 CFR Part 11, SOX, PCI-DSS, FINRA, MiFID II, FCA) that impose requirements on every document they submit externally. The constraint is the practical reality that every claim in a proposal, DDQ, or security questionnaire can become evidence in an audit, a regulatory filing, or a legal proceeding.

Regulated commercial teams face industry regulators who can investigate the content of proposals after submission. A pharmaceutical company that overstates a clinical capability in an RFP response creates a different kind of liability than a SaaS company that overstates an integration feature. The governance standard for these teams is not optional.

Right fit. Highly regulated commercial teams usually require a managed or autonomous proposal platform with governance scores of 4 or higher on the governance axis. The platform must enforce claim-level approval states, evidence linkage, permission inheritance across integrations, and exportable audit trails that cover AI-generated content alongside human-authored content.

Wrong Fit. Relying solely on governance-heavy platforms often produces hidden inefficiency. Drafting remains manual, SMEs are overused, and quality issues surface late. Deploying AI-native tools without integrating them into enterprise permissioning creates the opposite problem. You get speed without accountability, and the first audit exposes it.

Common failure mode. AI-generated claims carry the same regulatory exposure as human-authored claims, but many platforms do not enforce the same approval states across both.

Budget guidance. Typically $60k to $300k annually, depending on team size, response volume, and security environment requirements. Buyers in financial services and healthcare should budget separately for the security and compliance review process, which often adds four to eight weeks and procurement overhead beyond the platform license. Total cost of ownership should include integration depth (permission-aware synchronization, metadata mapping, write-back capabilities) and administrative load for governance maintenance.

High-volume teams (100+ annual responses)

Your situation. At triple-digit annual volume, small inefficiencies compound into structural capacity loss. A five-minute delay per review, multiplied across 150 responses with nine contributors each, becomes thousands of hours. Your constraint is not any single proposal. It is the economics of the entire pipeline.

Stargazy's Win Rate Report found that teams combining content automation, high reuse rates, and systematic customer insight are three times less likely to be in the lowest win-rate tier. Automation alone does not get you there. Automation with discipline does. Platforms must automate requirement extraction across document types, provide predictable first-draft timing, and support batch correction propagation. Reducing SME bandwidth consumption is critical to sustaining velocity without burnout.

Right fit. AI-native drafting engines and autonomous proposal platforms that prioritize rapid parsing and reuse at scale.

Wrong Fit. Library-centric platforms that depend on continuous human maintenance and curation. At triple-digit annual volume, the maintenance burden competes directly with delivery. The library grows faster than the team can review it, and stale answers start reaching evaluators.

Common failure mode. Relying on library-centric systems that require continuous human upkeep while volume compounds the maintenance deficit.

Budget guidance. Typically $60k to $150k annually.

GovCon and public-sector teams

Your situation. Commercial proposal software does not understand your world. FAR/DFARS compliance, CUI handling obligations, evaluator-scoring frameworks, and government portal submission constraints create a control surface that generic tools cannot absorb without extensive manual workarounds. Your evaluators score against published criteria, and your compliance exposure is not theoretical.

> "There are four risks to using AI in GovCon: security, hallucination, bias, and AI Speak. Most consider hallucination the biggest risk. Every fact the AI presents must be verified using factual data. However, the risks of a security leak, proposing a biased solution, and presenting information that seems polished yet says nothing (AI Speak) are equally risky."
>
> BRENDA CRIST
> VICE PRESIDENT • LOHFELD CONSULTING

LOHFELD CONSULTING

Right fit. GovCon capture-to-proposal platforms.

Common failure mode. Using commercial proposal software that lacks regulatory mapping or capture-to-proposal continuity. GovCon evaluators score against specific criteria, and tools that cannot map requirements to compliance obligations force manual reconstruction before every submission, and that reconstruction eats the time the tool was supposed to save.

Budget guidance. Varies widely by contract size and security requirements. Buyers should factor in FedRAMP and security environment costs, which often exceed the platform license. Total cost of ownership for GovCon deployments typically ranges from $75k to $500k annually including security environment and compliance overhead.

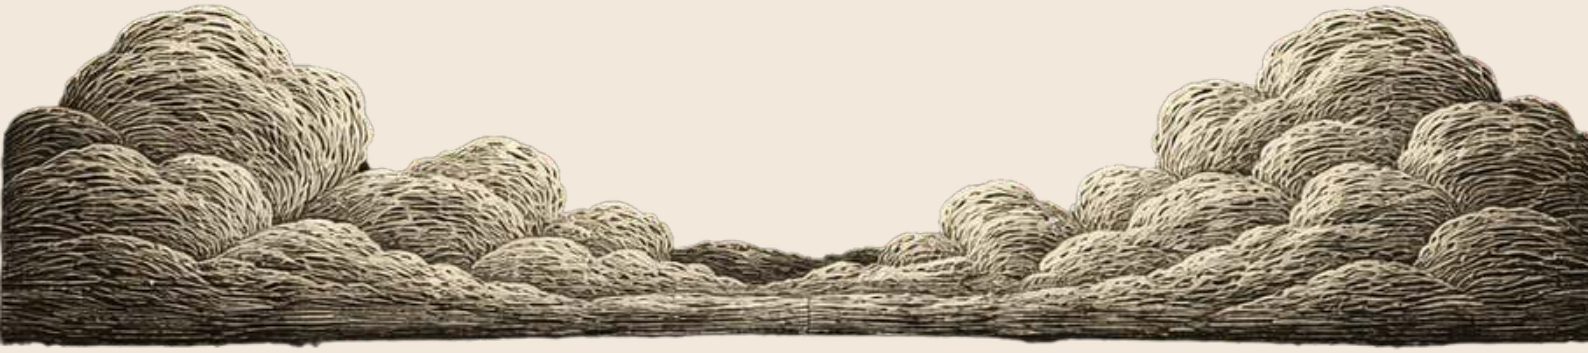

Structural trends: 2026–2028

Proposal technology will not converge around one dominant suite. It will separate by control surface and data discipline. Revenue leaders in B2B and B2G care about four outcomes. More qualified volume. Lower review load. Faster submission cycles. Fewer unsupported claims. McKinsey's 2025 survey shows the risk side of that same market: 51% of AI-using organizations report at least one negative consequence, and nearly one-third report consequences from AI inaccuracy. Those conditions will define the market through 2028.

The vendors that win will be the ones that reduce coordination cost, prove answer accuracy, carry context from opportunity to submission, and keep governance proportional to risk. Buyers already behave that way. The product map will follow.

The end of tagging economics

Tagging no longer absorbs the right problem. The marginal cost of drafting fell, but the marginal cost of validating and maintaining stored answers did not. Teams still rank SME collaboration and access to current answers among their biggest obstacles, and they still identify faster SME response times and stronger content management as two of the clearest ways to win more work. Manual library upkeep now competes directly with the work that actually moves win rates.

The library does not disappear; it contracts. Mature teams keep a small governed core for high-risk claims and pricing positions, then retrieve the rest from upstream systems at response time.

stargazy's documented switching patterns already point in that direction. Buyer movement clusters away from tagging-heavy, library-centric systems and toward retrieval-native tools and hybrid stacks.

Retrieval-first implementations often retire large portions of legacy libraries, with reported reductions of 40 to 70 percent in year one. Library-dependent environments often expand repositories, with some reporting annual growth above 20% to preserve coverage. Those trajectories create opposite economics. One narrows the surface area teams must govern; the other expands it.

Managed and autonomous proposal platforms are adding retrieval layers to reduce library dependence. AI-native drafting engines such as lup, AutogenAI, Arphie, and AnswerTree bypass the library model entirely. Governance-forward vendors across Managed and Autonomous face the hardest question. How do they enforce claim-level approval when the content those claims reference no longer lives inside the platform?

Some have moved toward a generative AI assistant built on curated content libraries and connected repositories, with custom prompt and skill-building capabilities.

But drafting capability alone does not resolve the governance question. Their ability to govern AI-generated claims, not only human-authored libraries, will determine whether that advantage holds.

Others, like RocketDoc's LUMA, grounds AI activity exclusively in authorized company knowledge sources with zero-knowledge architecture, along with AI applications (AI Solutions Engineer, AI RFP Manager) alongside encrypted virtual data rooms and real-time sales enablement.

The strategic bet is that governance-forward proposal vendors can retain enterprise buyers by becoming the governed knowledge layer for multiple revenue functions, not just proposal response.

Tendrio and TenderSeal face a different version of this challenge. As a pre-bid triage tool, its governance value lies in the compliance matrix that feeds downstream decision-making.

As AI-native drafting tools generate more of the proposal content, the quality of the upstream triage (whether the bid should have been pursued at all) becomes a higher-stakes governance decision.

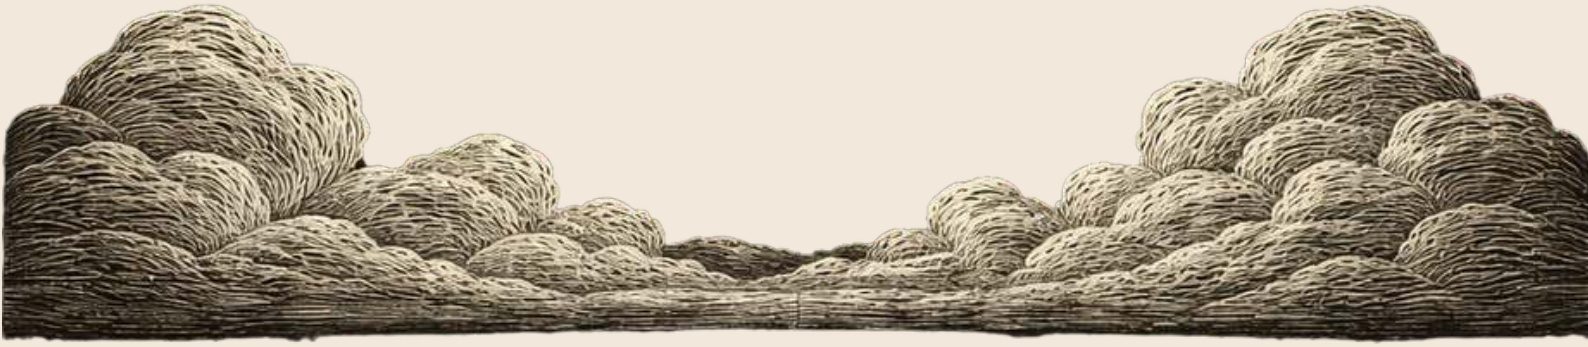

Orchestration displaces assistance as the buying standard

Proposal technology is moving through three stages of AI maturity, and agent systems are where the category is heading.

Stage 1: Assistance. ChatGPT Projects, Claude Projects, and Microsoft 365 Copilot now ship retrieval and drafting at a fraction of enterprise proposal software pricing. This stage is structurally commoditizing, and vendors competing on assistance alone face margin compression through 2027.

Stage 2: Orchestration. The AI takes on coordination work. Intake parsing, requirement decomposition, retrieval across connected sources, SME routing, review-state management, and handoff to approvers. In proposal operations, this is already visible in the Autonomous platform category.

Stage 3: Agent systems. Specialized agents coordinate with each other, each owning a narrow scope, handing off in a structured state.

Proposal teams that treat AI as a drafting assistant see throughput gains without win-rate impact.

Teams that rebuild around orchestration see compounding advantages.

The practical implication for 2026 buyers is that orchestration capability is no longer a forward-looking nice-to-have. It is the buying standard.

“Frankly, in too many organisations, bid and proposal professionals aren't taken seriously enough. We're so important to generating (and protecting) revenue that we should be seen by the C-suite as a business critical-function. But we're not. And now AI may make things worse: if we're simply seen as glorified admin, then surely new technology can take our place? In this brave new world, it's critical that we measure and articulate the value that we bring, and champion the role of human intelligence in winning highly competitive bids.”

JON WILLIAMS

MANAGING DIRECTOR · STRATEGIC PROPOSALS

Accuracy becomes a procurement KPI

When an AI-generated claim in an RFP response turns out to be false, the vendor is responsible regardless of whether a human or a model wrote the words. This legal reality is converging across geographies.

US federal contractors face False Claims Act exposure. UK and EU suppliers face contractual breach, regulatory investigation, and public-sector debarment. Financial services firms across jurisdictions face regulatory filings that treat proposal content as evidence. Healthcare and pharmaceutical companies face industry regulators who can investigate any external representation of capability.

The common pattern is that accuracy is moving from a marketing talking point to a measurable procurement criterion.

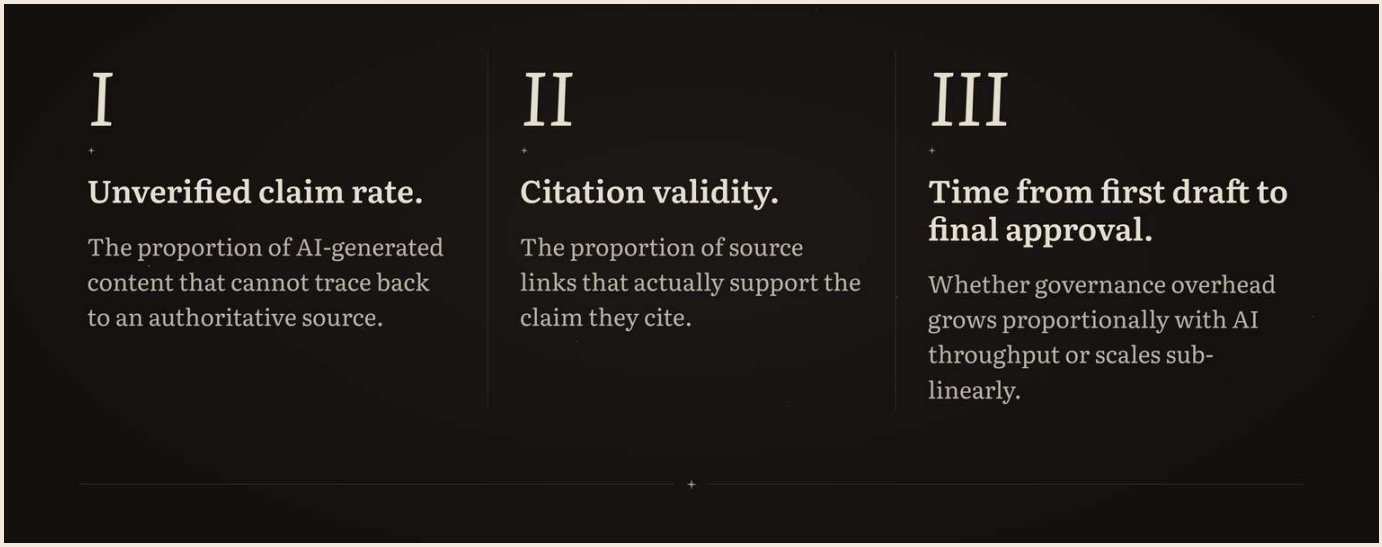

This will split the proposal technology market. Vendors that enforce claim-level citation at generation time, abstain under weak evidence, and preserve source traceability through export will survive procurement review. Vendors that produce fluent text without enforceable evidence will fail it. The split is already visible.

AI-native drafting engines with claim-level citation and Autonomous platforms with the same discipline are pulling ahead in regulated environments. Engines competing on drafting speed without governance infrastructure are facing margin compression and procurement rejection simultaneously.

“LLMs are excellent at generating ‘fuzzy’ and plausible output. Many professionals imagine them (or sell them) as precise calculators for language. That is not the case. LLMs will directly misquote, hallucinate facts, or entirely leave out relevant requirements. When precise facts matter, relying on LLMs to produce precise outputs without rigorous validation is dangerous for bid professionals.”

DAVID TIMM

PARTNER • BURR & FORMAN

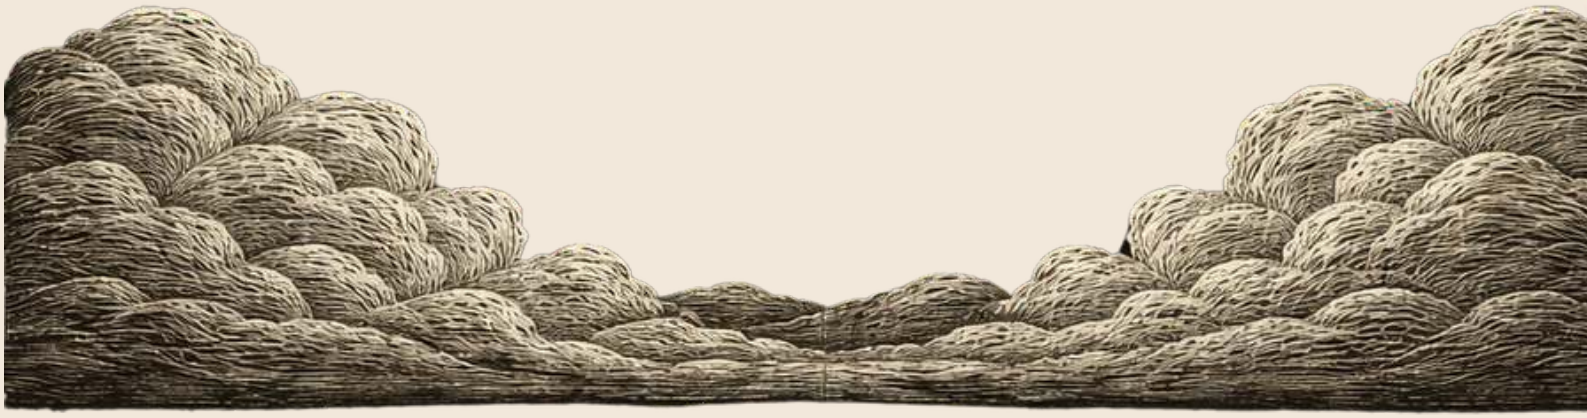

Content libraries shrink as retrieval improves

The best retrieval systems will shrink the amount of content teams must actively govern. Teams will keep a smaller set of proposal-specific assets and connect the rest to authoritative upstream systems.

stargazy implementation evidence records first-year reductions of 40 to 70 percent in legacy library surface area in retrieval-led environments, while library-dependent environments often expand repositories to preserve coverage.

The maintenance economics diverge as volume grows. Teams that cannot staff continuous library upkeep will migrate toward retrieval-native architectures.

> “Most SMEs I work with aren't adopting dedicated proposal software — they're defaulting to general LLM tools to get bids out quickly under pressure. It works up to a point, but it means they're missing the structured efficiency, governance, and collaborative benefits that proposal platforms provide. The real opportunity is bridging that gap without adding friction to already time constrained teams.”

DANIEL McILWAINE
FREELANCE BID CONSULTANT

Capture and response converge into one workflow

The line between capture and response will keep combining. Vendors will compete to own the full opportunity-to-submission workflow.

This convergence matters in both B2G and B2B. B2G teams need opportunity intelligence, qualification, compliance matrices, teaming strategy, and proposal production to share one context. B2B revenue teams need deal qualification, CRM context, security questionnaires, legal review, and proposal drafting to share one context.

High-performing teams use AI-assisted capture, predictive scoring, and pursuit-quality dashboards long before final submission. By 2028, teams that keep capture notes, qualification data, and proposal work in separate systems will recreate discovery friction inside every pursuit.

Managed and Autonomous platforms are best positioned for this convergence. Many software companies are building the full opportunity-to-submission context into a single platform.

GovCon capture-oriented vendors (Deltek/GovWin IQ, GovDash, GovSignals, Turingon, Vultron) are extending downstream from opportunity intelligence into proposal production. The category boundary between capture tool and proposal tool will blur throughout 2027 and 2028, and the buyers who benefit will be those running a single context from first pursuit signal to final submission.

The convergence continues into post-award document exchange, contract negotiation materials, and evaluation-phase data rooms are the next control surface that proposal-adjacent platforms will absorb.

stargazy buyer-intent tracking: the migration routes

stargazy tracks how buyers evaluate proposal technology through the stargazy Proposal Tech Directory, direct advisory engagements, and documented switching-route signals from the vendor community. The data shows migration patterns most analyst firms do not see, because the movement happens below the surface of publicly announced procurement cycles. The pattern is not random. It clusters around three main routes, and each route reflects a specific architectural mismatch between what the incumbent platform enforces and what the buyer's operating reality now requires.

Route 1: Library-centric to retrieval-native. Buyers leave tagging-heavy, library-centric suites for retrieval-native drafting tools. The driving pressure is the collapse of tagging economics and the library maintenance cost grows faster than the value of stored content, and teams that cannot staff continuous upkeep hit a structural ceiling on content freshness.

Route 2: Commercial horizontal to GovCon capture-to-proposal. GovCon buyers move from general-purpose commercial tools toward the dedicated GovCon capture-to-proposal category. The driving pressure is that commercial proposal tools cannot absorb FAR/DFARS compliance mapping, CUI handling, evaluator-scoring realism, or capture-to-proposal context continuity without extensive manual workarounds. These buyers are moving to a different control surface with different compliance mechanics, different buyer personas (BD and capture teams rather than RevOps), different budgets (BD or contracts rather than RevOps), and different evaluation criteria (regulatory defensibility rather than narrative quality).

Route 3: Informal AI workflows and proposal platforms cross each other in both directions. This route runs both ways. Teams running proposal work on general-purpose LLM setups (ChatGPT Projects, Claude Projects, Microsoft 365 Copilot, Gemini, or internally-built RAG pipelines with custom prompts) migrate into purpose-built platforms.

The opposite pattern is also documented, though less commonly. Teams that bought a proposal platform and found it did not absorb their constraint move back to LLM primitives with custom prompts. The driving pressure in this direction is that a proposal platform misaligned to the team's constraint costs more in administrative overhead than an informal LLM workflow costs in governance risk, at least in the short term. This movement is a warning signal for the vendor market. It indicates that the team's binding constraint was misdiagnosed during the original purchase, and that the platform they bought solved a problem they did not have.

Buyers are beginning to shop for the architecture that removes their dominant constraint. The buying signal is "our operating metrics are getting worse, not better, and the architecture cannot fix them."

Proposal & Bid Software Snapshot Matrix

Vendor	Dominant fit	Best-fit buyer	Industry emphasis
lup	AI-native drafting engine	Sales, IT, and questionnaire-heavy GTM teams	Software, consulting, BFSI
Altura	Autonomous proposal platforms	Bid managers and commercial directors	Construction, professional services, staffing, facilities, infrastructure, utilities, EU/UK/German public sector
Anchor	Autonomous proposal platforms (governance-forward)	Mid-market to enterprise proposal and sales teams	Broad enterprise
AnswerTree	AI-native drafting engine	Bid and proposal teams at tech companies	Software, telecom, professional services
Arphie	AI-native drafting engine	PreSales, proposal, and security teams handling complex RFPs, DDQs, and questionnaires	Broad B2B
AutogenAI	AI-native drafting engine	Enterprise, public-sector, and GovCon bidders	UK government, EU procurement, US federal, defense, construction
AutoRFP.ai	Autonomous proposal platforms	B2B sales and proposal teams that need fast setup	Software, tech scale-ups, BFSI, healthcare
BidScript	Autonomous proposal platforms (governance-forward)	SME and mid-market bid teams handling public and private tenders	UK and EU tenders, construction, infrastructure, professional services, B2B
Brainial	Autonomous proposal platforms (EU/UK tender-focused)	Tender and bid teams	Construction, infrastructure, facilities, EU/NL
Cassidy AI	Autonomous proposal platforms	Sales and proposal teams	B2B software
CLEATUS	GovCon capture-to-proposal	Federal, SLED, and DIBBS contractors needing full-lifecycle agentic AI	Federal, state/local, education, defense
ContraVault AI	Vertical specialist (AEC)	AEC bid teams and general contractors handling complex, multi-document solicitations	Construction, engineering, energy, infrastructure

Vendor	Dominant fit	Best-fit buyer	Industry emphasis
Conveyor	Adjacent: trust-center and questionnaire automation	Security and compliance teams	B2B SaaS
Deltek/GovWin IQ	GovCon capture-to-proposal	Federal BD and capture teams	Federal, defense, SLED
Ensis AI	GovCon capture-to-proposal	Federal and defense contractors needing Office-native proposal workflows with CUI-compliant security	Federal, defense, healthcare, education, finance
Expedience Software	Managed proposal platforms	Microsoft Word-centric proposal teams	Broad enterprise
Flowcase	Vertical specialist (personnel)	Experience-heavy bid teams	Professional services, engineering, construction, law, architecture
GovDash	GovCon capture-to-proposal	Mid-market and enterprise federal contractors	Federal, defense, SLED
GovEagle	GovCon capture-to-proposal	Federal contractors needing Office-native proposal workflows with CUI-compliant security	Federal, defense, IT services
GovSignals	GovCon capture-to-proposal	Federal, SLED, and public-sector growth teams	Defense, FedCiv, SLED, AEC, IC, NASA
Inventive AI	AI-native drafting engine	Enterprise proposal and presales teams handling high-volume RFPs	B2B SaaS, software, financial services, healthcare
Iris	AI-native drafting engine	Mid-market proposal and sales teams	Software, B2B
Joist AI	Vertical specialist (AEC)	AEC marketing and BD teams needing content enablement and proposal automation tuned to construction and engineering	Architecture, engineering, construction
Loopio	Managed proposal platforms	Cross-functional RFP/RFI/DDQ/SQ teams	Broad enterprise
MyTender	Autonomous proposal platforms	Construction and facilities management bid teams, UK public sector suppliers	Construction, facilities management, UK public sector, energy, IT

Vendor	Dominant fit	Best-fit buyer	Industry emphasis
Ombud	Autonomous proposal platforms	Distributed SME and RevOps-heavy teams	Enterprise software, cybersecurity, HR tech, healthcare
OpenAsset	Vertical specialist (AEC DAM)	AEC marketing and bid teams	AEC and real estate
Procurement Sciences (Awarded AI)	GovCon capture-to-proposal	Mid-to-large federal contractors needing full-lifecycle AI with embedded change management	Federal, defense, civilian agencies
Proposify	Managed proposal platforms	Sales-led proposal teams	Broad B2B
pWin.ai	GovCon capture-to-proposal	Shipley-aligned federal proposal teams needing methodology-embedded AI	Federal, defense, aerospace
QorusDocs	Managed proposal platforms (governance-forward)	Value-led teams in legal, AEC, professional services, IT	Legal, AEC, professional services, IT
Realm	Autonomous proposal platforms	B2B SaaS and technology revenue teams needing retrieval-native RFP automation within a broader GTM AI workspace	B2B SaaS, technology, professional services
Responsive	Managed proposal platforms	Mid-market and enterprise high-volume teams	Broad enterprise; AEC, consulting, software, BFSI
RocketDocs	Managed proposal platforms (governance-forward)	Microsoft-centric and control-oriented teams; enterprise teams needing governed AI beyond proposals	Broad enterprise; regulated buyers (financial services, healthcare, technology)
Rohirrim	GovCon capture-to-proposal	Enterprise GovCon and AEC firms	Federal, defense, AEC
SEQUESTO	AI-native drafting engine	Tender desk teams	EU tenders, broad B2B
Seev	AI-native drafting engine	Proposal teams (early stage)	B2B

Vendor	Dominant fit	Best-fit buyer	Industry emphasis
SiftHub	Autonomous proposal platforms (governance-forward)	Sales, presales, SE, and proposal teams	B2B software and technical sales
Steerlab	Autonomous proposal platforms	PreSales, proposal, and security teams at B2B software companies	B2B SaaS, software, financial services
Stotles	Managed proposal platforms (upstream intelligence)	Public sector sales and bid teams at enterprise suppliers	UK and EU public sector, technology, professional services
Sweetspot	GovCon capture-to-proposal	Mid-market federal contractors needing discovery-to-submission in one platform	Federal, defense, civilian agencies
TenderSeal	AI-native drafting engine	Tender and bid teams	EU and UK tenders
Tendium	Managed proposal platforms (with pre-bid triage capability)	European public sector bid teams needing monitoring, qualification, and bid creation in a structured workflow	Scandinavian and EU public sector tenders, construction, infrastructure, professional services
Tendrio	Pre-bid compliance triage (emerging)	Defense, aerospace, and regulated-industry bid teams needing automated tender evaluation	Defense, aerospace, engineering, infrastructure, European public-sector tenders
ThalamusAI	Autonomous proposal platforms	Proposal and bid teams	B2B
Trampoline AI	Managed proposal platforms	Service businesses handling complex tenders	B2B commercial, public sector, AEC, healthcare, professional services
Tribble	Autonomous proposal platforms	Distributed SME and RevOps-heavy teams	Enterprise software, cybersecurity, HR tech, healthcare
Turingon (Capture Pilot & Proposal Pilot)	GovCon capture-to-proposal	Federal and SLED contractors with compliance-heavy bids	Federal, defense
Vanta	Adjacent: trust-center and questionnaire automation	Security and compliance teams	B2B SaaS
Vera	AI-native drafting engine	Security, sales, and compliance teams	B2B SaaS

Vendor	Dominant fit	Best-fit buyer	Industry emphasis
VisibleThread	GovCon capture-to-proposal	GovCon proposal and contracts teams	Federal, defense, IT services
Vultron	GovCon capture-to-proposal	Mid-market and enterprise federal contractors	Federal, defense
XaitPorter	Managed proposal platforms	Enterprise proposal teams	Broad enterprise

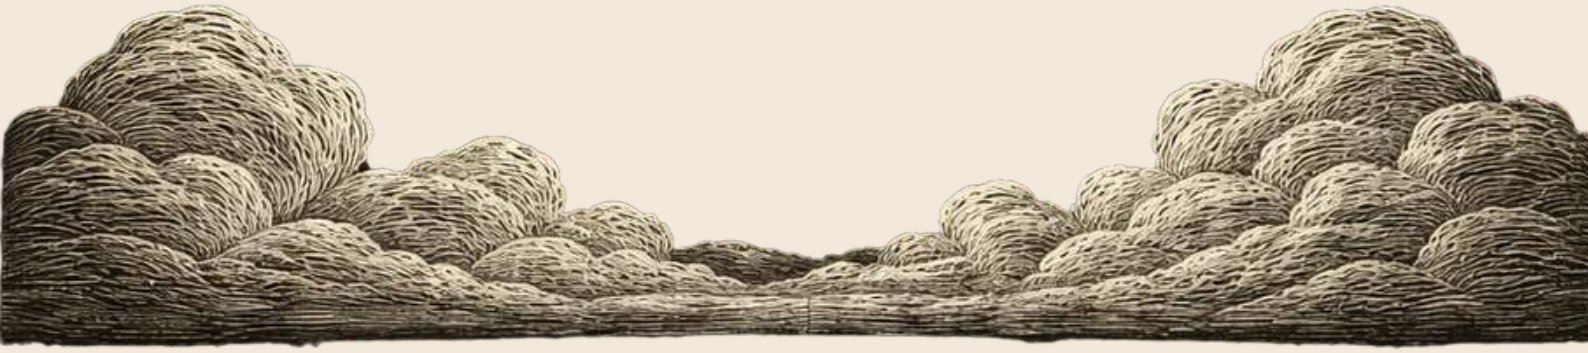

Appendix A ✹ Methodology & scope

This appendix describes how the 2026 Proposal & Bid Software Report was researched and validated.

Research Approach

The report uses a constraint-based analytical framework rather than a feature-comparison methodology. Each vendor is classified by the control surface it enforces most heavily in production use, not by marketing positioning or feature breadth. The five-category taxonomy and governance capability axis were developed through iterative analysis of vendor architectures, buyer evaluation patterns, and switching-route data collected between Q2 2025 and Q1 2026.

Research inputs include direct vendor engagement (product access, positioning queries, architectural review), published analyst sources (Gartner, McKinsey, Forrester, NIST, OMB, EY, ISO/IEC), and stargazy primary research (Proposal Win Rate Report 2026; Stargazy Proposal Tech Directory; advisory engagement data; documented switching-route signals; buyer-intent telemetry).

The report covers platforms commercially available as of Q1 2026 that serve proposal, bid, RFP, RFI, DDQ, or security-questionnaire workflows. Vendor inclusion criteria are described in Appendix D. The report does not cover general-purpose LLM providers, CRM platforms, document management systems, project management tools, or contract lifecycle management platforms unless they position explicitly for proposal and bid workflows.

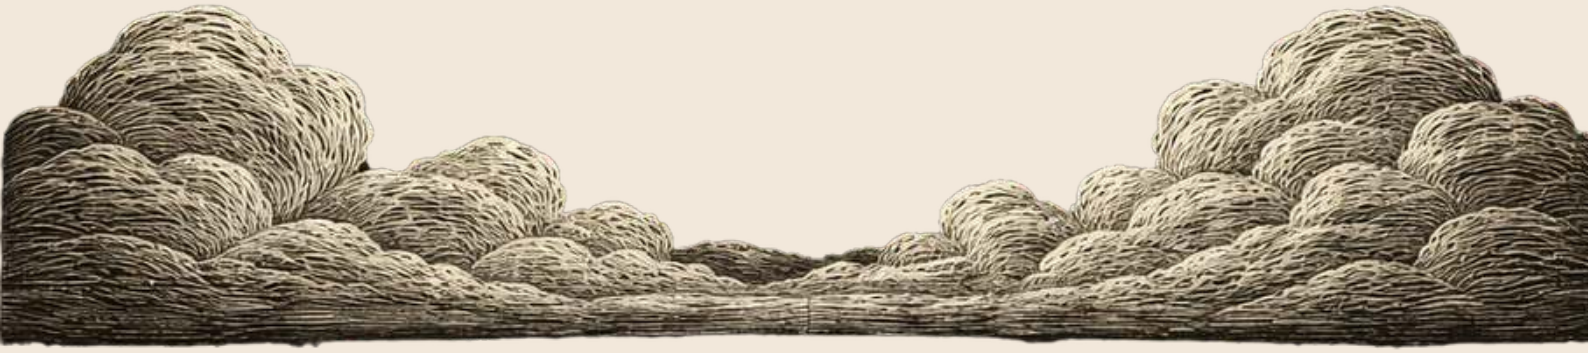

All quantitative claims are attributed to named sources. When a claim cannot be sourced to published data, the report uses directional language and attributes the assessment to stargazy. Forward-looking predictions carry explicit reassessment dates.

The report does not rank vendors, estimate market share, audit security controls, or certify compliance readiness. The vendor capability snapshot is directional, not evaluative, and buyers should validate fit through their own pilot using the evaluation criteria described in this report.

Independence

This report is sponsored by 1up and AutoRFP.ai. Sponsorship supports distribution and editorial production. Sponsorship does not influence category definitions, evaluation criteria, inclusion decisions, or vendor conclusions. The independence standard applies equally to all vendors referenced in this report, including the sponsors. No vendor paid for inclusion. No vendor reviewed or approved editorial content before publication.

Analytical standards

The report applies three editorial standards throughout.

First, every material claim carries a named source or an explicit stargazy attribution.

Second, governance is assessed as a cross-cutting capability axis (scored 1-5) rather than as a standalone category, because governance maturity varies within every architectural category.

Third, trust fidelity (the system's ability to generate claims grounded in identifiable sources, constrained by permissions, and traceable to accountable reviewers) is treated as the defining performance standard for the category in 2026, instead of drafting speed or feature breadth.

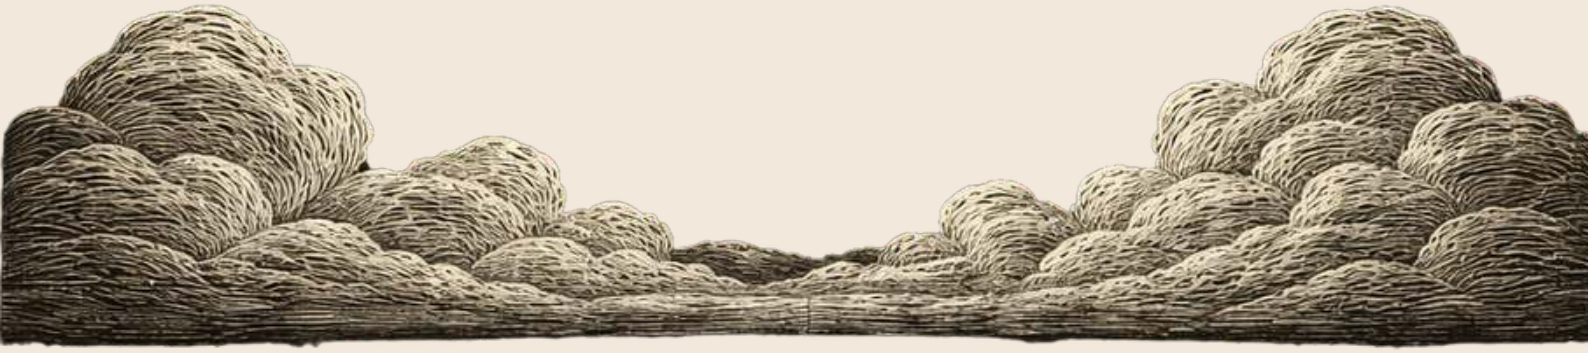

Appendix B ✹ Inclusion criteria

This appendix describes how vendors were selected for inclusion in the 2026 Proposal & Bid Software Report and how they were assigned to categories.

Inclusion scope

The report covers platforms that serve the proposal, bid, and structured-response workflow for commercial and public-sector organizations.

A vendor qualifies for inclusion if it meets all three conditions:

- the platform is commercially available as of Q1 2026 (not in closed beta or pre-launch),
- the vendor positions the product for proposal, bid, RFP, RFI, DDQ, or security-questionnaire work (not for general-purpose content generation, CRM, or project management),
- at least one verifiable customer, case study, directory listing, or public deployment is identifiable.

Requesting inclusion or correction

Vendors not listed in this report can request inclusion review for the 2027 edition by contacting contact@stargazy.io with a product URL, one verifiable customer reference, and a description of the dominant control surface the platform enforces.

Vendors that believe their category assignment is incorrect can submit a correction request with evidence of their production control surface.

stargazy will update categories and vendor inclusions on a quarterly basis in the interactive report.

Appendix C ✹ Data sources & citations

This report draws on the following published sources. All quantitative claims in the report are attributed to named sources. When a claim cannot be sourced, the report uses directional language.

Analyst and research firm sources

- Gartner, "Predicts 2025: AI Agents" (2024). Forecast: at least 30% of GenAI projects abandoned after POC by end 2025; 40%+ of agentic AI projects canceled by end 2027; agentic AI in 15% of day-to-day work decisions and 33% of enterprise software by 2028.
- McKinsey & Company, "The State of AI in 2025" (2025). Survey data: 51% of AI-using organizations report at least one negative consequence; inaccuracy most commonly reported risk; high performers redesign workflows, not just tasks.
- Forrester, "B2B Buying Study" (2024-2025). Buying data: average of 13 internal stakeholders per B2B purchase; procurement as decision-maker in 53% of cycles; 60%+ of buyers use trial; buyers validate AI information against trusted human sources.
- EY, "Generative AI Risk and Governance" (2025). Analysis: hallucinations in high-stakes, client-facing work create compliance failures, reputational damage, and regulatory exposure.
- HBR, AI adoption and change management coverage (2024-2025). Research: organizations get marginal gains from AI bolted onto legacy workflows; rollouts fail when teams revert under pressure; adoption depends on workflow redesign, not prompt training.
- Government and standards body sources
- NIST, "Generative AI Profile" (AI 600-1, 2024). Guidance: evaluate claims with empirically validated methods; review sources and citations during pre-deployment and monitoring; test in deployment-like conditions; document provenance; maintain post-deployment monitoring.
- OMB, "AI Acquisition Guidance" (April 2025). Directive: agencies assess proposals for AI-related risks before selection; consider testing environments on agency networks.
- GSA, "USAi" and "Buy AI" guidance (2025). Guidance: agencies involve CIO, AI, data, security, and privacy leaders; monitor usage costs; USAi provides secure AI model evaluation.
- ISO/IEC 42001, "Artificial Intelligence Management System" (2023). Framework: emerging standard for AI development, deployment, and governance practices.
- Industry benchmark sources
- Deltek, "2025 GovCon Clarity Study" (2025). GovCon data: AI used for proposal development and market analysis; 41% spend 7+ hours on first draft; expert warning on AI-generated proposals lacking human oversight.

stargazy primary research

- AutoRFP.ai & Stargazy, Proposal Win Rate Report (2026).
- stargazy Proposal Tech Directory (2025-2026). Vendor positioning data, switching-route signals, buyer-intent telemetry, and co-evaluation network data.
- stargazy implementation evidence (2025-2026).

Acknowledgements

This report was researched and produced by Stargazy.

stargazy thanks the proposal leaders, bid directors, revenue operations teams, and procurement professionals who shared evaluation data, adoption metrics, switching experience, and operating-model insights through the stargazy buyer community and direct advisory engagements. Their frontline experience shaped the constraint-based framework, the adoption risk analysis, and the migration-route data that distinguish this report from vendor-sourced marketing.

stargazy thanks the vendor teams who responded to positioning queries, provided product access for architectural review, and corrected factual errors during the verification process. Inclusion in this report does not imply endorsement of any vendor's product, and vendor cooperation did not influence category assignment or editorial conclusions.

stargazy thanks 1up and AutoRFP.ai for sponsoring the distribution of this report. Sponsorship supports reach and production. It does not influence category definitions, evaluation criteria, inclusion decisions, or vendor conclusions.

stargazy thanks the broader proposal management community for maintaining the professional standards and peer networks that inform independent research.

Errors and analytical judgments remain the sole responsibility of stargazy.

Corrections and feedback are welcome at contact@stargazy.io.

Copyright and legal disclaimer

Copyright

Legal Disclaimer

This report is provided for informational purposes only and does not constitute legal, procurement, financial, or technology advice. The analysis reflects publicly available evidence as of the publication date. Vendor capabilities, pricing, certifications, and product positioning change frequently. Buyers should verify all claims through their own evaluation and due diligence.

Inclusion of any vendor in this report does not imply endorsement, certification, or recommendation by stargazy. Omission does not imply exclusion from the market or unsuitability for any buyer. stargazy does not rank vendors, estimate market share, audit security controls, or certify compliance readiness.

stargazy has made reasonable efforts to ensure the accuracy of the information presented, but makes no warranties, express or implied, regarding completeness, accuracy, reliability, or fitness for any particular purpose. stargazy shall not be liable for any loss, damage, or expense arising from reliance on this report.

This report is sponsored by 1up and AutoRFP.ai. Sponsorship supports distribution and editorial production only. Sponsorship does not influence category definitions, evaluation criteria, inclusion decisions, or vendor conclusions. The independence standard applies equally to all vendors referenced in this report, including the sponsor.

About stargazy

stargazy is an independent proposal win intelligence firm.

It maintains the stargazy Proposal and Bid Software Directory and publishes annual market reports and buyer-intent data that tracks how organizations evaluate and switch between proposal technology vendors.

stargazy provides managed shortlist services for enterprise buyers, vendor dashboard access for product and go-to-market teams, and category positioning advisory for proposal software leadership. It operates a private buyer community where proposal leaders, revenue operations leaders, and bid directors share evaluation data and adoption metrics.

Proposal & bid software directory: stargazy.io
Buyer community: https://stargazy.circle.so/
Contact: contact@stargazy.io

“The promise of AI in bids is to write faster. The reality is more drafts hitting the same review queue. The constraint was never the writing. It was getting the right knowledge out of the right place in time. Teams that figure this out stop using AI purely as a drafting tool and start using it to surface what the organisation already knows: past bids, win rationale, client context, lessons that lived in one person’s head for years. That’s the shift. And AI can solve for that.”

Heather Melton

Founder, Loro Flow

Podcast Creator/Host, bidfolk