

STRATEGY.AI

Pricing & Engagement Guidebook

A comprehensive reference on how we price, why we price this way, and answers to the 36 most important questions we receive – from procurement, CFOs, CIOs, functional leaders, and legal teams.

CONTENTS

- 01 – Introduction
- 02 – Pricing philosophy (5 principles)
- 03 – The three tiers (Pilot / Functional / Enterprise)
- 04 – The AI connection model
- 05 – Implementation & onboarding
- 06 – 36 Q&As across 8 categories
- 07 – Quick reference table
- 08 – Closing note

INTERNAL & CLIENT USE · CONFIDENTIAL

SECTION 01 – INTRODUCTION

Why a guidebook on pricing?

Pricing conversations are often where enterprise deals become complicated – not because the numbers are wrong, but because the conversation hasn't established the right framing first. This guidebook exists to fix that.

It explains the logic behind every number we publish, the reasoning that shaped each tier, the implementation model your team will work through, and thorough answers to the questions we hear most often – from procurement, from CFOs, from CIOs, from functional leaders who want to move independently, and from legal teams who need to understand liability.

This is a living document. Use it before a client conversation, during a commercial negotiation, or as a reference when a question comes in that you haven't heard before. If something is missing, add it.

WHO THIS GUIDE IS FOR

Sales team members preparing for pricing conversations. Business development and pre-sales contacts.

Client-facing team members who need to explain or defend our pricing model. Leadership who need a complete reference for board or investor conversations.

SECTION 02 – PRICING PHILOSOPHY

Five principles behind our pricing

Principle 1: Price reflects commitment, not features

We do not use feature-gated pricing. Every tier includes everything required to achieve the outcome described. We do not withhold governance tools, audit trails, or support tiers to create artificial upsell pressure.

What changes across tiers is scope – the number of AI connections, the depth of the strategy artifact intake, the breadth of the governance deployment, and the level of ongoing support. Price reflects how much we are deploying, not what we are willing to unlock.

Principle 2: The Pilot is paid for a deliberate reason

We do not offer free trials. A free trial produces low engagement from both sides – the client commits minimal internal resource, we have no commercial alignment to invest heavily, and the result is an inconclusive evaluation that neither party can act on.

The \$20,000 Pilot fee is calibrated to sit below most single-signature procurement thresholds while being meaningful enough to signal serious intent from both parties. At that level, you get a real implementation team, real configuration, and a real working system – not a demo, not a sandbox, not a self-serve setup. You leave with something you can show your leadership.

THE \$20K FRAMING

We often describe the Pilot as 'small enough to approve below the PO threshold, large enough to take seriously.' In practice this means a functional leader with a discretionary budget – CMO, CFO, CHRO, COO – can typically approve it without a formal procurement process. That is intentional.

Principle 3: The unit of value is the AI connection, not the user

We deliberately avoided per-user or per-seat pricing. The reason is simple: the value of strategy-aware AI is not determined by how many people use a tool – it is determined by which tools are governed and what decisions they influence.

Microsoft Copilot deployed to 400 people is still one AI connection. Salesforce Einstein is one connection. A custom procurement agent is one connection. Pricing at the tool level aligns our fee to the scope of governance we are performing – not to your headcount, not to your company size, and not to your licence count with another vendor.

Principle 4: Full transparency on the day rate

For Functional and Enterprise tiers, onboarding is included within a defined scope. When clients ask for work outside that scope – additional document intakes, custom integrations, extended training, post-year-one re-ingestion cycles – we charge \$2,000 per day. This is disclosed upfront, on the pricing page, in every proposal, and in every contract.

We do this because the most common source of commercial friction in enterprise software is scope creep that was never discussed. We would rather have that conversation at the beginning than after the invoice arrives.

Principle 5: Every engagement ends with a clear decision

The Pilot is structured to give you a go/no-go decision at the end of 90 days. Not a vague feeling about whether it was useful. A formal readout with a working system, a governed claim set, usage analytics, and a specific recommendation on how to scale.

This discipline benefits both parties. If the approach does not fit your organisation, we would rather know at 90 days than at month 18 of an unhappy annual contract. If it does fit, you have everything you need to make a confident expansion commitment.

SECTION 03 – THE THREE TIERS

Pilot, Functional, Enterprise

The three tiers are not a good/better/best ladder. They represent three genuinely different deployment scenarios – different buyer types, different organisational contexts, different governance scope. A Functional deployment for a CFO is not a downgraded Enterprise – it is a complete, self-contained deployment for a specific function. A Pilot is not a partial Functional – it is a structured proof-of-concept with a defined outcome and a defined decision point.

Pilot

\$20K

flat fee · 90 days · all-inclusive

- ✕ 3 AI system connections (e.g. Copilot, Salesforce Einstein, custom agent)
- ✕ Strategy context build – delivered by our implementation team (up to 30 governing sources)
- ✕ Governance console setup and configuration
- ✕ Leadership governance review and claim approval workflow
- ✕ Hands-on implementation team included throughout
- ✕ Formal readout: working system, approved claim set, expansion roadmap

ONBOARDING & SUPPORT

All implementation, configuration, and onboarding included within the 90-day period. No additional fees during the Pilot.

OUTCOME

A working Strategy Context Layer with 3 governed AI connections – and a clear decision: expand to Functional or Enterprise, or don't.

Functional

\$13K

per AI connection · per year · minimum 4 connections (\$52K/yr minimum)

- ✕ All contracted AI connections configured and governed
- ✕ Strategy context build – delivered by our implementation team (up to 50 governing sources)
- ✕ Governance console with human authorization workflow
- ✕ Full audit trail and provenance chain
- ✕ Governance training session for your leadership team
- ✕ Annual strategy refresh cycle (one full ingestion cycle per year)
- ✕ Email and Slack support during business hours

ONBOARDING & SUPPORT

Onboarding included within defined scope. Additional implementation available at \$2,000/day. No CSM – shared support.

OUTCOME

Your function's AI tools operating from your approved, governed functional strategy — a complete deployment, not a stepping stone.

Enterprise

\$75K + \$9K/connection

platform fee plus per-connection fee · per year · no connection cap

- ☑ Everything in Functional, expanded to all functions
- ☑ Corporate strategy layer connecting all functional deployments
- ☑ Dedicated Customer Success Manager
- ☑ Strategy context build — delivered by our implementation team (up to 150 governing sources across all functions)
- ☑ Quarterly strategic alignment reviews with your leadership team
- ☑ Board-level reporting and claim provenance dashboards
- ☑ Priority support with SLA on strategy layer availability
- ☑ Custom claim taxonomy aligned to your strategy architecture
- ☑ Private cloud deployment option (own AWS/Azure/GCP environment)

ONBOARDING & SUPPORT

Expanded onboarding included. Additional implementation at \$2,000/day. CSM manages ongoing change as your strategy evolves.

OUTCOME

Every AI tool across every function drawing from one authoritative, governed source of strategic truth.

SECTION 04 – THE AI CONNECTION MODEL

What is an AI connection?

An AI connection is a single AI system configured to query the Strategy Context Layer via our MCP/API integration layer. When that AI system generates a recommendation, prioritises a course of action, or responds to a user query, it first calls the Strategy Context Layer to retrieve the relevant governing context — approved strategic direction, authority levels, constraints, time horizons — before generating its response.

The connection is invisible to end users. Your team continues using the tools they already have — Copilot, Salesforce, Workday, a custom agent. The Strategy Context Layer sits beneath those tools as an infrastructure service, called silently at inference time.

KEY DISTINCTION

Connections are priced per tool type, not per user or seat. Microsoft Copilot deployed to 400 people is one

connection. Salesforce Einstein is one connection. A custom procurement agent is one connection. Adding 200 more users to Copilot does not change your connection count or your annual fee.

Examples of AI connections by function

- Microsoft 365 Copilot – connected for a marketing or finance function
- Salesforce Einstein – connected for sales deal qualification and forecasting
- Workday AI – connected for CHRO workforce planning and headcount decisions
- ServiceNow AI – connected for IT operations and incident prioritisation
- Power BI AI narratives – connected for finance reporting and board commentary
- Custom LLM agent – any proprietary agent built on GPT-4, Claude, Gemini, or similar
- Kinaxis RapidResponse – connected for supply chain scenario analysis
- Oracle SCM AI – connected for procurement and supplier risk decisions
- HubSpot AI – connected for marketing campaign and segment prioritisation
- Workiva AI – connected for board reporting and governance document generation

Connection count by tier

In the Pilot, 3 connections are included in the flat fee – enough to demonstrate real governance value across meaningful use cases. In Functional, the minimum is 4 connections ($\$13K \times 4 = \$52K/\text{year}$). In Enterprise, the \$75K platform fee covers the infrastructure and corporate strategy layer; each connection is then \$9K/year with no cap on the total number.

At 9+ connections in an Enterprise deployment, the economics compare favourably: $\$75K + (\$9K \times 9) = \$156K$ versus the equivalent Functional rate of $\$13K \times 9 = \$117K$. The premium reflects the corporate strategy layer, the dedicated CSM, the priority SLA, and the private cloud option.

SECTION 05 – IMPLEMENTATION & ONBOARDING

How an engagement actually runs

Every engagement follows the same six-step delivery framework regardless of tier. The difference is scope – the number of connections, the volume of strategy documents ingested, and the breadth of the governance deployment. Steps 1–6 are always completed in order; the timeline compresses or expands based on the complexity of the AI tool environment and the availability of your governance stakeholders.

Step	Week	What happens
01	Week 1	Kickoff & discovery – we map your AI tool inventory, strategy document sources, and governance stakeholders. We agree the claim taxonomy and document the authority hierarchy for each strategic domain.
02	Weeks 2–3	Strategy artifact intake – approved documents (strategy decks, OKRs, board papers, approved policies, budget frameworks) are ingested, structured, and tagged against the claim taxonomy.
03	Weeks 4–5	Claim structuring & leadership review – extracted strategic claims are reviewed by your governance team. Each claim is approved, rejected, or flagged for revision. Authority levels and time horizons are set.
04	Weeks 5–6	API/MCP configuration – the Strategy Context Layer is deployed. Connections to your AI systems are configured and tested in a sandboxed staging environment before going live.
05	Weeks 7–9	Live governance testing – the system runs in a monitored live state. We track inference calls, review sample outputs, refine claims, and validate that governed responses are accurate, useful, and proportionate.
06	Weeks 10+	Formal readout & decision – a structured session with your project sponsor presenting the working system, governed claim analytics, and a specific expansion recommendation. This is the go/no-go decision point.

What we need from you

A successful engagement requires genuine investment from your side. Here is an honest account of the internal resource typically needed for a Pilot:

- A named project lead who can coordinate internal stakeholders and schedule document access – 3 to 5 hours per week during implementation
- A curated shortlist of 15–30 governing sources your governance lead considers currently authoritative – we help with selection at kickoff
- Two governance stakeholders with authority to approve strategic claims – typically the CSO or equivalent, plus one functional leader – for a combined 8–10 hours across the engagement
- One IT contact for a single day of API/MCP configuration support in Step 4 – no special access to AI systems required
- A 2-hour leadership claim review session in Weeks 4–5

WHAT WE DO NOT NEED

We do not need access to your AI systems themselves – only the ability to configure a context layer that sits beneath them. We do not modify your Copilot, Salesforce, or Workday deployments. We work via standard MCP/API integration protocols. Your internal user experience does not change.

SECTION 06 – QUESTIONS & CHALLENGES

36 questions we get asked

These are real questions from real conversations – from procurement teams, CFOs, CIOs, functional leaders, and legal teams. They are organised into eight categories. Read them in order or jump to the category most relevant to your current situation.

CATEGORY 1 OF 8 – BUDGET & PROCUREMENT**Q01 Where does the Pilot fee typically sit in a budget?**

The \$20,000 Pilot fee is deliberately sized to sit below most single-manager or single-signature approval thresholds in mid-to-large enterprise organisations. In practice it can usually be approved as professional services spend under existing budget authority – without a formal procurement process or committee review. We have seen Pilots approved from IT budgets, strategy consulting budgets, marketing technology budgets, innovation funds, and transformation programmes. The key is that it sits with a functional leader who has the appetite and the authority – not with a centralised procurement gate.

Q02 Can we avoid a full procurement process for the Pilot?

In most cases, yes. The Pilot is structured as a professional services engagement with a defined start and end date, a fixed fee, and no auto-renewal. Most enterprise procurement frameworks have a simplified process for fixed-scope professional services below a threshold – typically \$25K to \$50K. We provide a concise Statement of Work designed to clear a lightweight approval process. If your organisation requires a preferred vendor list, we can discuss the fastest path to qualification. We have done this before and know what information procurement teams typically need.

Q03 The Pilot is small enough to approve, but the Functional contract isn't. How do we plan for that?

That is intentional. The Pilot is sized to be approvable at the function level. If it works, the Functional or Enterprise contract – which does require procurement involvement – is supported by a working system, governance analytics, and a formal readout that gives procurement something concrete to evaluate. We consistently hear that having a live pilot outcome makes the downstream procurement process significantly faster: the evaluation question shifts from 'does this work?' to 'how much should we spend scaling it?' We are

happy to brief your procurement team at the readout stage and will provide all standard vendor documentation.

Q04 Can we split the Pilot fee across two budget lines or cost centres?

Yes. We can issue a split invoice against two purchase orders or cost centres. We will need the allocation confirmed in the Statement of Work, but this is a routine request and adds no complexity on our side. For clients where the Pilot spans two departments – for example, when the CIO and the CMO are co-sponsoring – split billing is straightforward.

Q05 Is the annual Functional contract paid upfront or in instalments?

The default is annual upfront. For Functional contracts above \$78,000 per year (six or more connections), we offer quarterly instalments on request. Enterprise contracts are structured case by case – annual upfront is the default; quarterly is available with a standard 3% instalment adjustment to cover cash flow cost. We do not offer monthly billing for any tier.

Q06 What happens at end of year one – is there an auto-renewal?

No auto-renewal. Functional and Enterprise contracts require active renewal at the end of each term. We contact you 90 days before renewal to confirm scope and pricing for the next year. If your strategy has evolved significantly, we discuss what a refresh cycle looks like. If you decide not to renew, we provide a clean data export – your governed claim set, your strategy data schema, and your usage logs – and a formal offboarding process. Nothing is locked in beyond the contracted term.

CATEGORY 2 OF 8 – COMMERCIAL & NEGOTIATION**Q07 Is the pricing negotiable?**

The published prices are our standard rates. We do not negotiate the per-connection price or the platform fee as a matter of policy – not because we are inflexible, but because our pricing is calibrated to what it actually costs to deliver the outcome we describe. Where we do have flexibility: payment terms, instalment schedules, multi-year commitment discounts (available at year 2+), and the scope of included onboarding. If there is a genuine budget constraint that makes the standard pricing difficult, we would rather have that conversation directly than spend three months pretending it is not a constraint.

Q08 What if we want more than 3 connections in the Pilot?

The Pilot is intentionally scoped to 3 connections. This is a delivery discipline, not a pricing constraint. Three connections is enough to demonstrate real governance value, produce meaningful analytics, and give your team a clear view of what expansion looks like. Adding more connections to the Pilot does not improve the decision you are trying to make – it makes the engagement harder to manage without improving the outcome. If you know from day one that you want 8 connections, we recommend going directly to a Functional contract. We can discuss what a modified timeline looks like if you want to move faster.

Q09 Can we lock in the per-connection rate at \$13K for multiple years?

Yes. We offer a rate-lock commitment for clients who sign a two- or three-year agreement at the time of the initial Functional contract. In exchange for the term commitment, we guarantee no price increase during the contracted period beyond a CPI adjustment after year 2. A three-year rate-lock includes a 5% discount on the per-connection fee – bringing it to \$12,350 per connection per year. Multi-year Enterprise commitments are available with bespoke commercial terms.

Q10 What if we need to add connections mid-year?

Mid-year additions are charged on a prorated basis from the month of activation. If you add a fifth connection in month 7 of a 12-month contract, you pay 5 twelfths of the annual \$13K for that connection in year one, then the full rate from year 2. We issue a co-termination order so all connections share the same renewal date and you receive one annual invoice rather than multiple staggered ones.

Q11 What if we want to pause or suspend the contract for a quarter?

We do not offer pause provisions in annual contracts. The Strategy Context Layer is a persistent infrastructure service – governance, audit trails, and strategy claims need to remain live and current while your AI tools are operating. Suspending the contract while your AI tools continue to run would create a governance gap that would undermine the purpose of the deployment. If there is a genuine business disruption – a major restructuring, a merger, an executive transition – we can discuss options case by case. We have handled these situations before and they do not require adversarial negotiation.

CATEGORY 3 OF 8 – VALUE & ROI**Q12 How do we justify \$52,000 per year to a CFO who has never heard of a Strategy Context Layer?**

Start with the cost of the ungoverned alternative. How much has your organisation spent on AI tools in the last 12 months? For most enterprise organisations the figure is \$500K to \$5M when you include Copilot licences, Salesforce AI, Workday AI, and custom agent development. Now ask: what governance exists to ensure those tools are recommending actions aligned with your approved strategy? In most cases the answer is 'none, beyond the prompts written at setup.' At \$52K, the Functional tier governs four AI systems at less than a 10% governance premium over what you are already spending on ungoverned AI. The CFO framing is not 'why are we spending this?' – it is 'how are we not spending this already?'

Q13 What ROI metrics can we point to?

We track four categories of measurable value: (1) Decision alignment – the percentage of AI-assisted decisions that reference an approved strategic claim, measured via the audit trail; (2) Governance incident reduction – AI outputs that contradict approved policy, measured before and after deployment; (3) Leadership time savings – reduction in governance stakeholder time spent reviewing and correcting AI-generated

recommendations; (4) Procurement and capital allocation compliance — AI-influenced decisions that are within approved parameters. We provide a custom ROI model based on your AI tool usage volume and governance baseline. Most clients see quantifiable governance value within the first 60 days of live operation.

Q14 What if we go through the Pilot and decide not to proceed?

You keep everything. The approved claim set, the governance console export, the strategy data schema, and the formal readout are all yours. The \$20K fee is not refundable — it covers the cost of delivering a real implementation. But you are not locked into anything further. We have had Pilots conclude with clients who decided the timing was not right and returned 12 to 18 months later when internal readiness improved. We will not pursue a commercial relationship that is not right for your organisation. A clean 'not now' is a better outcome than a Functional contract signed before the organisation is ready.

Q15 How does this compare to hiring a strategy consultant to do something similar?

A strategy consulting engagement to map, govern, and operationalise strategic alignment across four AI systems would typically cost \$200K to \$500K for the initial engagement, produce a static report, and require a follow-on engagement to implement anything. Strategy.ai provides the platform and the implementation team together, at a fraction of the cost, and delivers a live governance system rather than a document. The ongoing maintenance cost of a bespoke consultant-built solution would be prohibitive, and the output — a static set of prompts and guidelines — would not adapt as your strategy evolves. The comparison is closer to building your own CRM versus buying Salesforce.

Q16 We are already paying for Microsoft Copilot. Why do we need to pay for governance on top?

The cost of Copilot is the cost of the AI capability. The cost of Strategy.ai is the cost of making that capability strategically safe to use at scale. Microsoft Copilot can give your team an answer — it has no mechanism to ensure that answer reflects your approved strategic priorities, governance constraints, or leadership intent. Microsoft will not add that mechanism to Copilot — it requires knowledge of your specific strategy, your authority hierarchy, and your approved decision frameworks. That is precisely what we provide. The two costs serve entirely different purposes.

CATEGORY 4 OF 8 — IMPLEMENTATION & TIMELINE

Q17 How much internal time do we need to commit during the Pilot?

Realistically: a named project lead at 3 to 5 hours per week across 90 days; two governance stakeholders for a combined 8 to 10 hours (primarily at the claim review session in weeks 4 to 5); and one IT contact for a single day of API configuration support. We do the rest. The most common reason Pilots run longer than 90 days is delayed access to strategy documents or difficulty scheduling the governance review session — both internal coordination issues, not implementation issues. We will help you avoid them by flagging them early.

Q18 Our strategy documents are scattered across SharePoint, email, and individuals' drives. Is that a problem?

It is the most common situation we encounter. We do not expect a curated library at kickoff. What we need is a shortlist of 15 to 30 governing sources your governance lead considers currently authoritative — strategic plans, OKRs, board papers, approved policies, budget frameworks. We help you identify and prioritise those during the kickoff session. The curation exercise itself often generates useful internal clarity about which sources are genuinely authoritative versus which are out of date — a byproduct of the engagement that clients consistently find valuable.

Q19 What if our strategy changes significantly during the Pilot?

The Pilot is 90 days. Major strategic shifts in most organisations happen on a quarterly or annual cadence — the Pilot timeline is designed to operate within a stable strategic window. If there is a genuine major pivot during the Pilot (a board-level strategic change, an M&A announcement, a significant leadership change), we work with the governance lead to determine whether to pause the claim review process, update the affected documents, and re-run the relevant ingestion before proceeding. This has happened; it is manageable. It is not a contractual issue — it is a practical judgment call made in partnership with your team.

Q20 What does the formal readout at the end of the Pilot include?

The formal readout is a structured session of approximately 90 minutes with your governance lead and project sponsor. It covers: (1) A live working system demonstration showing strategy-aware versus ungoverned AI responses side by side on your actual connections; (2) The approved claim set — a reviewed and finalised inventory of the strategic claims governing each connection, with authority levels assigned; (3) Governance analytics — volume of inference calls, percentage governed, sample output review, governance incident log; (4) An expansion roadmap — a specific recommendation on the number and priority of connections for a Functional or Enterprise deployment, estimated timeline, and indicative commercial terms. It is designed to give your leadership everything they need to make a confident expansion decision — or a confident decision not to expand.

Q21 What if we want to extend the Pilot beyond 90 days?

We offer a single 30-day extension at \$6,000 — prorated from the \$20K base. Extensions beyond 120 days are not available. If at 120 days there is still insufficient information to make a decision, that is typically a signal that the internal governance structure is not yet in place — in which case we recommend pausing and returning when internal readiness has improved. A 120-day inconclusive pilot is not a failure; it is useful information about organisational readiness.

CATEGORY 5 OF 8 — TECHNICAL INTEGRATION**Q22 We use a proprietary AI system not on your standard list. Can you connect to it?**

Yes, in most cases. Our MCP/API integration layer is designed to be protocol-agnostic. If your proprietary system accepts context at inference time — via a system prompt, a retrieval call, a context injection parameter, or a structured API call — we can configure a connection to it. We will need a technical specification of the integration surface, usually available from your internal AI team or the vendor. Custom integration work beyond our standard configuration scope is available at the \$2,000/day rate and is typically a one-time cost. Once configured, ongoing management of the connection is included in the annual fee.

Q23 Do you need access to our AI systems, or to our data?

Neither in the traditional sense. We do not access your Copilot environment, your Salesforce instance, or your internal business data. What we deploy is a context service — an API endpoint that your AI systems call at inference time to retrieve governing context. We configure the integration on our side; your IT team enables the outbound API call from your AI system to our endpoint. Your internal data stays in your environment. The only information that passes through the Strategy Context Layer is the strategic context we have ingested from the documents your governance team has approved and the inference metadata (query category, connection ID, claim matched — never the content of user queries).

Q24 What are the data security and residency requirements?

The Strategy Context Layer processes only the strategy documents and claims your governance team has approved — not user queries, user data, or proprietary business data. We are SOC 2 Type II compliant. Data residency options (EU, US, APAC) are available on request for Enterprise clients. We provide a full security questionnaire response and a Data Processing Agreement as standard for any Functional or Enterprise engagement. For the Pilot, we operate under our standard terms with a lightweight security addendum available on request. We are experienced at navigating enterprise information security review processes.

Q25 Can we deploy this in our own cloud environment rather than yours?

A private cloud deployment — your AWS, Azure, or GCP environment — is available at the Enterprise tier. This is not the default configuration; the default is a hosted managed service. For organisations with strict data sovereignty requirements, air-gap policies, or existing cloud commitments that make managed SaaS difficult, we support private deployment. This requires a custom scoping conversation, adds 4 to 6 weeks to the implementation timeline, and affects the commercial terms. It is not available for the Pilot or Functional tiers.

CATEGORY 6 OF 8 — COMPETITIVE & ALTERNATIVES**Q26 Our AI vendor says they already handle strategic alignment. What is different?**

AI vendors often claim 'customisation' or 'grounding' features. What they typically mean is: you can add content to a system prompt, or connect a document corpus to a retrieval pipeline. What they do not provide: a structured authority hierarchy for strategic claims, a governance workflow for approving and versioning those claims with named human sign-off, an audit trail that traces every AI recommendation back to a governing

source document, or a human authorization layer that ensures no claim goes live without explicit leadership approval. Those capabilities require an understanding of how enterprise strategy governance works — not just how language models work. That is the distinction.

Q27 Can't we do this with RAG (retrieval-augmented generation)?

RAG retrieves relevant text from a document corpus and includes it in the context of an AI query. It is a useful technique — but it is not strategy governance. RAG does not distinguish between an approved strategic claim and an outdated draft. It applies no authority levels: it retrieves a draft strategy deck with the same priority as a board-approved plan. It has no governance workflow, no human authorization step, and no audit trail. It cannot model the difference between a claim that is within the authority of a Finance leader versus a claim that requires Board approval. RAG is an excellent retrieval mechanism. Strategy.ai is a governance layer. In Enterprise deployments, they often coexist.

Q28 How is this different from a well-designed system prompt?

A system prompt is a static string of text prepended to an AI conversation. It has no version control, no authority hierarchy, no audit trail, no approval workflow, and no mechanism to ensure it reflects current approved strategy. Most system prompts are written by a developer or consultant who had access to a strategy document at a point in time. When strategy changes, system prompts do not update automatically. When a claim in the system prompt turns out to be wrong, there is no trace of who approved it. Strategy.ai replaces the system prompt approach with a governed, versioned, approval-tracked context layer that reflects current authorised strategy — always, for every AI system, across every function.

Q29 We have an internal AI governance team. Why can't they build this?

They could. The question is the time, cost, and quality of the alternative. Building a governed strategy context layer requires: a claim taxonomy design, a document ingestion pipeline, an authority hierarchy model, a governance approval workflow, an audit trail and provenance system, an MCP/API integration layer, ongoing version management, and a maintenance function. For most enterprise AI teams this is a 12 to 18 month build at \$500K to \$1.5M in internal resource cost, before accounting for ongoing maintenance and the distraction cost to a team that likely has a full backlog. We provide it at \$52K/year for a Functional deployment with 4 connections. The build-versus-buy conversation is a straightforward one.

CATEGORY 7 OF 8 — RISK & GOVERNANCE**Q30 What if our strategy changes significantly mid-contract?**

The annual strategy refresh cycle included in Functional and Enterprise tiers covers one full document re-ingestion per year — timed around your annual strategy cycle. If your strategy changes materially mid-year (a major pivot, a new product line, a restructuring, a leadership change that brings new priorities), additional ingestion cycles are available at \$2,000/day. In practice, most organisations find that core strategic claims are

relatively stable between annual cycles – what changes are priorities, emphasis, and resource allocation, not the fundamental direction. We work with your governance lead to determine what requires a re-ingestion and what can be handled by updating specific claims directly in the governance console.

Q31 Who owns the governance process – us or Strategy.ai?

You own the governance process entirely. Strategy.ai provides the platform, the implementation team, and the governance tooling. Every strategic claim must be reviewed and approved by a named member of your leadership team before it goes live in any AI connection. We do not have the authority to approve claims on your behalf, and we would not accept that authority if offered. The governance workflow is explicitly designed so that your organisation retains full accountability for what the Strategy Context Layer says. We are responsible for the quality and reliability of the implementation. You are responsible for the strategic content it governs.

Q32 What is the liability position if a governed AI recommendation turns out to be wrong?

This is the right question and we answer it directly. Strategy.ai is a context governance layer – we ensure that the strategic claims surfaced to AI systems are the ones your leadership has reviewed and approved. We do not make the AI recommendation, and we do not make the business decision. Our liability is limited to the accuracy and completeness of the governance layer we implement, as specified in the contract – specifically, that the claims surfaced reflect what was approved in your governance workflow. The AI vendor retains liability for the quality of their model outputs. Your organisation retains accountability for the business decisions that follow. We provide full indemnification detail in the contract and recommend your legal team review it alongside your AI vendor contracts.

Q33 We are going through a merger or acquisition. How does that affect our contract?

A merger or acquisition is handled as a contract novation – the successor entity assumes the existing contract at the same terms. If the merger results in a significantly larger AI footprint (more connections, multiple strategy layers to consolidate), we will work with your team on a revised scope and commercial arrangement. We have a specific professional services offering for post-M&A strategy consolidation: helping a combined organisation establish a unified governance layer that reflects the merged strategic direction, with a clear authority hierarchy across legacy structures. This is available as a scoped engagement and is priced at the \$2,000/day rate for work beyond the existing contract scope.

CATEGORY 8 OF 8 – SCALING & GROWTH**Q34 We started with the CMO's function. Finance now wants to deploy. How does that work commercially?**

The CMO's Functional deployment covers the marketing function's strategy layer – its own document intake, its own governance workflow, its own connections. When Finance wants to deploy, they get a separate Functional deployment: their own strategy intake, their own governance workflow, their own connections. Commercially,

this is either a separate contract or an addendum to the existing one – whichever your billing structure prefers. Finance connections are priced at \$13K each. There is no automatic cross-function discount at the Functional tier. If two or more functions are deploying simultaneously, or if a third function is considering it, that conversation is often the trigger to explore Enterprise economics – at which point the numbers become compelling.

Q35 When does it make sense to move from Functional to Enterprise?

Enterprise becomes the right conversation when any of the following are true: (1) Three or more functions are deploying or want to deploy; (2) Your CIO or CSO wants a corporate strategy layer connecting all functional deployments to a single authoritative source; (3) You need a dedicated Customer Success Manager rather than shared support; (4) You have board-level AI governance reporting requirements; (5) The cumulative Functional bill across multiple functions is approaching or exceeding the Enterprise platform fee. On the economics: at 9+ connections, Enterprise ($\$75K + \$9K \times 9 = \$156K$) becomes more expensive than Functional ($\$13K \times 9 = \$117K$) – the premium reflects the corporate strategy layer, the dedicated CSM, the priority SLA, and private cloud optionality. At 12+ connections, both paths converge and Enterprise becomes the natural structural choice regardless of economics.

Q36 If we eventually want 15 or 20 AI connections, is there a ceiling?

No ceiling at Enterprise. The \$75K platform fee is fixed regardless of connection count; each connection adds \$9K per year. At 20 connections, you are at $\$75K + \$180K = \$255K$ per year for an organisation with 20 governed AI systems – a meaningful governance infrastructure at a cost that most enterprise technology budgets can absorb. For context, a 20-connection Enterprise client is typically governing AI across 6 to 8 functions at a per-function cost of roughly \$30K to \$40K – significantly cheaper per function than equivalent Functional deployments at scale. The platform is architected to scale horizontally without architectural changes: adding a 20th connection is operationally identical to adding a 5th.

SECTION 07 – QUICK REFERENCE

Pricing at a glance

	PILOT	FUNCTIONAL	ENTERPRISE
Price	\$20K flat	\$13K / conn / yr	\$75K + \$9K / conn / yr
Min. commitment	None	4 connections (\$52K/yr)	None
Contract term	90-day fixed	Annual	Annual
AI connections	3 included	Min. 4, no cap	No cap

	PILOT	FUNCTIONAL	ENTERPRISE
Document intake	Up to 30	Up to 50	Up to 150
Onboarding	Fully included	Defined scope	Expanded scope
Overflow day rate	None during Pilot	\$2,000/day	\$2,000/day
Strategy refresh	N/A	1× per year	1× per year + CSM
Dedicated CSM	No	No	Yes
Priority SLA	No	No	Yes
Private cloud	No	No	Yes
Rate lock	N/A	2–3 yr available	Case by case
Instalment billing	N/A	Quarterly (6+ conns)	Quarterly available
Auto-renewal	No	No	No

Connection economics at a glance

Connections	Pilot	Functional	Enterprise
3	\$20K (flat)	N/A (min. 4)	N/A
4	–	\$52K/yr	N/A
6	–	\$78K/yr	N/A
8	–	\$104K/yr	\$147K/yr
10	–	\$130K/yr	\$165K/yr
15	–	\$195K/yr	\$210K/yr
20	–	\$260K/yr	\$255K/yr 📌

Enterprise crossover point is approximately 15–20 connections where the per-connection saving on Enterprise begins to offset the \$75K platform fee.

SECTION 08 – CLOSING NOTE

A note on pricing conversations

Pricing conversations work best when both parties are honest about constraints early. We would rather lose a deal at the first call because the budget is not there than spend three months in a negotiation that was never going to close.

If the standard pricing does not work for your situation, tell us why. We cannot always find a path, but we will always tell you honestly whether one exists.

For the right organisations – those with a meaningful AI tool footprint, a leadership team that takes strategic alignment seriously, and an internal governance structure that can support a claim review process – Strategy.ai is one of the clearest commercial decisions in the enterprise AI governance landscape.

For organisations that are not ready – because the internal strategy governance does not yet exist, because the AI tool footprint is not yet at a scale where governance is urgent, or because the budget is genuinely constrained – we would rather say that plainly and let you come back when the timing is right.

We are building for the long term.

TO START A CONVERSATION

Book a strategy briefing at strategy.ai/contact – or email your contact at Strategy.ai directly. Every engagement starts with a 30-minute conversation to understand your AI tool inventory, your governance baseline, and whether the Pilot is the right starting point for your organisation. No pricing pressure. No pitch deck. Just a conversation.