AI NEGOTIATIONS

What Procurement Gets Wrong and How to Fix It

Chad Johnson

Former Chief Procurement Officer | Pfizer | McKesson | Humana | Centene

All opinions expressed in this paper are the author's own.

Why I Wrote This

I have negotiated hundreds of enterprise contracts across some of the largest healthcare and pharmaceutical organizations in the world. I have renegotiated most software agreements, rebuilt supplier frameworks from scratch, and sat in more rooms than I can count where Finance looked at procurement's numbers and said prove it.

In 2024, I found myself sitting on AI governance committees at the enterprise level. And I will be honest about what that experience felt like: I had no idea what I was looking at. I was a former Fortune 20 Chief Procurement Officer with over two decades of sourcing experience, and I felt like I had nothing of value to add. The conversations were full of terms I did not know, concepts I could not evaluate, and decisions I was not equipped to push back on.

At the same time, I was being asked to review and approve contracts with AI embedded in them. No return on investment framework. No governance structure. No baseline. Just a vendor deck and a request to sign. The same failures I had spent my career eliminating from traditional procurement were showing up all over again, just wrapped in new language.

So, I did what I have always done when I did not understand something well enough to negotiate it. I took it apart. I read everything I could find. I talked to people who understood technology. And I started building a framework that made sense to me as a procurement professional, not as a technologist.

This paper is the result of that work. It is written for procurement leaders who feel the same way I did sitting in those governance meetings. You do not need to be a technologist to negotiate AI well. You need to apply the same discipline you have always applied, with a clear understanding of what is different about these deals and what is not.

> ***I was a former Fortune 20 Chief Procurement Officer with over two decades of sourcing experience, sitting in an AI governance meeting, and I felt completely lost. If that is where you are right now, this paper is for you.***

Executive Summary

Enterprises are signing AI contracts they do not fully understand, cannot benchmark, and cannot exit cleanly. The pricing model is unfamiliar, the data risks are underestimated, and the results are reported in terms Finance will never validate.

This paper gives procurement leaders a practical framework to negotiate AI deals the same way Finance expects everything else to be bought: with clear baselines, measurable outcomes, and accountability built into the contract before anyone signs.

It also addresses a scenario most AI negotiation guides ignore entirely: what happens when the AI comes to you. Your existing software vendors are quietly embedding AI into products you already own and calling it an upgrade. No new contract trigger. No procurement review. Just a terms update at renewal and suddenly your data is feeding someone else's model. That gap is covered here.

Section 1: Know What You Are Buying

Before you can negotiate an AI deal, you need to understand what kind of AI you are actually procuring. Most enterprises treat AI as a single category. It is not. There are now four distinct types, and each requires a completely different negotiation approach.

Type 1: Direct AI Infrastructure

Examples include OpenAI, Anthropic Claude, and Google Gemini. When you buy access directly to one of these, you are buying infrastructure, similar to a cloud compute contract. Pricing is based on tokens. These deals look simple on the surface and are not.

Type 2: AI-Powered Applications

Contract review tools, spend analytics platforms, and sourcing automation tools all fall into this category. You are buying a finished solution, not raw computing power. The underlying AI cost is buried inside the vendor's margin, and most buyers never see it.

Type 3: Agentic AI Systems

This is the fastest growing and least understood category. According to a 2026 Futurum Group survey of 830 enterprise decision makers, autonomous agents surged 31.5 percent year over year as the top technology priority. These systems do not wait for human instructions at each step. They plan, execute, and complete multi-step workflows on their own. The commercial model is still evolving, with the industry moving toward per-action pricing where you pay for a completed task rather than a token count.

Type 4: Embedded AI in Existing Software

This is the category most procurement teams are not watching. Your existing software vendors, including platforms like Salesforce, Workday, ServiceNow, Oracle, and Microsoft, are adding AI capabilities to products you already own. It shows up as a new module in a renewal conversation, a terms of service update you clicked through, or a feature announced at the vendor's annual conference that is now live in your environment.

No new purchase order is generated. No new vendor relationship is created. No procurement trigger fires. The risk lands in your renewal terms and your existing data processing agreement, and most organizations discover it after the fact.

> ***Getting the category right matters because the negotiation is completely different for each one. Confusing them is how procurement ends up with the wrong contract structure and a return on investment conversation it cannot win.***

Section 2: The Embedded AI Problem Nobody Is Watching

The most overlooked AI negotiation happening right now is not a new vendor pitch. It is your existing software renewal.

Salesforce, Workday, ServiceNow, Oracle, Microsoft, and virtually every major enterprise software vendor are aggressively embedding AI into their existing platforms. This is not accidental. These companies are under significant investor pressure to demonstrate AI roadmaps, and embedding AI into products customers already own is the fastest way to show momentum. The features arrive packaged as upgrades, announced at user conferences, and bundled into renewals before most organizations have formed a coherent view on whether they want them.

The honest assessment of where most of this embedded AI sits today: the technology is real, the vendor ambition is genuine, and the immediate return on investment is often unclear. These are not mature capabilities with validated enterprise outcomes. They are early-stage features being deployed at scale into production environments, and the organizations buying them are largely figuring out the value after the fact.

That does not mean the answer is always no. It means procurement needs to get ahead of the conversation with a clear internal position before the vendor arrives at the renewal table.

The Question You Must Answer Before the Renewal Conversation

The single most important thing procurement can do before any embedded AI renewal is align with Finance and the business on one strategic question: are we buying a business technology or are we making an innovation investment?

These are not the same thing, and they require completely different governance frameworks.

A business technology purchase follows proven logic. There is a defined problem, a measurable baseline, a Finance-validated return on investment, and contractual accountability for outcomes. The standard is clear: show me the value before we sign.

An innovation investment follows different logic. You are buying optionality. You are making a calculated bet that early adoption will create competitive advantage, organizational learning, or future capability that is difficult to quantify today. The standard is different: what is the cost of not participating, and what do we need to learn?

Both are legitimate. But they require different conversations with Finance, different contract structures, and different tolerance for ambiguity. The failure mode procurement keeps repeating is applying innovation logic during the sales conversation and business technology logic during the ROI review six months later. The vendor sold you a vision. Finance wants a spreadsheet. Nobody aligned on which game was being played.

Get that alignment in writing before the renewal. A one-page internal position paper co-signed by Finance and the business unit sponsor is enough. It does not need to be complicated. It needs to exist.

> ***The vendor sold you a vision. Finance wants a spreadsheet. The failure mode is applying innovation logic during the sales conversation and business technology logic during the ROI review. Align on which game you are playing before anyone signs anything.***

How Embedded AI Enters Your Environment

Once your organization has aligned on the strategic question, procurement needs to understand the three paths through which embedded AI typically arrives. The first is a terms of service update that quietly grants the vendor expanded rights to use your data, including the right to use transactional data to train or improve their AI models. The second is a renewal conversation where AI capabilities are bundled into a higher tier or a price increase. The third is a product update announcement where AI features are enabled by default without requiring explicit customer approval.

In all three cases, the question procurement should be asking is the same: what rights did we just grant this vendor over our data, and did we agree to that intentionally?

The Return on Investment Question Is Different

When you buy new AI, the return on investment question is straightforward: what will this replace, reduce, or improve, and can Finance validate the number. When a vendor embeds AI into software you already own, the return on investment question has a different structure. You are not evaluating a new investment. You are evaluating whether the incremental value of the AI capability justifies the incremental cost being added to your renewal, and whether the data rights you are granting in exchange are acceptable.

Three questions define the return on investment analysis for embedded AI:

- What am I being asked to pay more for at renewal, and what specific AI capability is bundled into that increase? If the vendor cannot quantify the value of the AI feature separately from the base product, the price increase is not justified.
- Is the AI feature actually replacing a workflow or just augmenting one that still requires the same headcount and effort? A feature that assists a human is not the same as a feature that eliminates a process. Only one of those generates a Finance-validated return.
- What data rights am I granting by accepting the new terms, and what is the cost of those rights? Proprietary supplier pricing, negotiation strategies, contract terms, and spend data have real commercial value. Granting a vendor the right to use that data to train their models is a transfer of value that does not appear on any invoice.

What the Contract Review Must Cover

Before any renewal that includes embedded AI capabilities, procurement must review the updated terms against three specific provisions.

The first is model training rights. The updated terms must explicitly prohibit the vendor from using your data to train, fine-tune, or improve their AI models without your prior written approval. This provision should already exist in your data processing agreement. If it does not, add it at renewal.

The second is data residency. If your organization operates in regulated industries or serves customers in jurisdictions with data protection laws, confirm that the AI processing does not route your data outside your required geography. Embedded AI features sometimes use infrastructure that is separate from the base product's data residency commitments.

The third is opt-out rights. Embedded AI features are sometimes enabled by default. Your contract should give you the right to disable AI processing for specific data types or workflows without losing access to the base product or triggering a pricing penalty.

Four Questions Before Accepting Any Embedded AI Upgrade

- What data does this AI feature access and process?
- Does the vendor have the right to use our data to train or improve their models under the new terms?
- What is the incremental cost and what specific, measurable value does the AI feature deliver in return?
- Can we disable AI processing for sensitive workflows without penalty?

Section 3: The Token Problem

For any deal involving direct access to an AI language model, tokens are the unit of cost. Most procurement teams have never negotiated a token-based contract, and most AI vendors are counting on that.

AI language models do not scale like traditional software. They scale like computing infrastructure, and that changes everything about how enterprises should budget for them. Token consumption compounds at scale in ways that are not obvious at the pilot stage. A workflow processing thousands of interactions daily can generate dramatically different costs depending on which model you use and how the contract is structured.

The Variables Procurement Must Understand Before Signing

- Input versus output token cost. Output tokens cost significantly more than input tokens.
- Reasoning tokens. Used by advanced models for complex tasks, billed separately and more expensively.
- Prompt caching discounts. These exist but must be contractually specified, not assumed.
- Model tiering. Routing simpler tasks to less expensive models can reduce cost significantly.

Organizations that actively manage these levers typically achieve 30 to 50 percent reductions in AI-related expenses, according to Binadox's 2025 analysis.

The Negotiation Basics

- Push for volume commitment tiers with defined discount thresholds.
- Negotiate usage caps with renegotiation rights if you exceed them.
- Require itemized pricing rather than bundled rates.
- Benchmark across providers before committing.

> ***Token spend must be converted into a cost-per-outcome metric Finance can track over time. Paying a rate per million tokens is operationally meaningless. Paying a defined amount per contract reviewed is auditable.***

Section 4: The Confidentiality Risk Nobody Is Negotiating

When a company runs sensitive data through a cloud-hosted AI tool, that data leaves the corporate network and enters infrastructure controlled by a third party. Supplier pricing, negotiation strategies, contract terms, financial

projections, patient records, and merger analysis are all examples of information that enterprises routinely input into AI tools without fully understanding where it goes.

The Samsung Warning: A Real-World Case Study

In March 2023, Samsung semiconductor engineers inadvertently exposed proprietary company data in three separate incidents within a single month, all by using a commercial AI chat tool to assist with their work. One engineer entered source code to debug a database program. A second entered code related to identifying defective equipment. A third converted a recorded internal meeting into a document and entered it to produce meeting minutes.

Because the commercial AI interface used input data to train and improve its models, the proprietary information was retained on the vendor's servers. Samsung had no means to retrieve or delete it. Samsung responded by banning external AI tools company-wide. JPMorgan Chase, Verizon, and Amazon issued similar warnings or restrictions around the same period.

The lesson is simple: an AI governance policy that arrives after the breach is not a policy. It is a response.

The Shared Infrastructure Problem

Standard cloud AI deployments are multi-tenant. Many organizations share the same underlying infrastructure. Research published in 2025 found that implementing AI language models in shared environments poses serious security risks, specifically that sensitive information can be leaked through memory caches in ways that expose data to other users of the same system.

The On-Premise Requirement for Sensitive Workloads

For any use case that touches confidential data, procurement must evaluate on-premise or private cloud deployment as the baseline requirement, not an optional upgrade. Unlike cloud-hosted AI, which requires data to be sent offsite, on-premise models allow organizations to process AI workloads entirely within their own infrastructure.

Full on-premise deployment at enterprise scale carries real capital cost and operational complexity. The most effective approach for most enterprises is a hybrid structure: local models deployed on internal servers handle the sensitive daily workload, while complex tasks that do not involve confidential data are routed to a secure cloud service.

Four Questions Before Any AI Tool Processes Enterprise Data

- Where does the data go during processing?
- Is the infrastructure shared with other organizations?
- Is the vendor contractually prohibited from using your data to improve its own models?
- Is on-premise or private deployment available for your most sensitive use cases?

Section 5: Require Return on Investment Before You Sign, Not After

The ROI problem exists across all four AI categories, but it is most acute in Agentic AI. This is where the commercial model is least mature, vendor claims are most aspirational, and Finance has the least precedent to

validate against. You are being asked to fund autonomous systems that will execute workflows end to end, often at significant scale, based on pilot results that were designed by the technology team rather than structured to answer the questions Finance will ask six months later.

The same Futurum survey [1] found that direct financial impact combining revenue growth and profitability has nearly doubled as the primary measure of AI return, while productivity gains have declined as the leading indicator of success. Finance leaders are no longer accepting time saved as a return on investment. They want income statement impact. In the agentic context, that means procurement must build the measurement framework into the contract before anyone signs, not after the system goes live and the vendor declares success.

Require a Pre-Agreed Measurement Framework Before Signing

Three elements are non-negotiable:

- A documented baseline of current-state metrics including processing time, error rates, and labor hours.
- A defined go-live milestone with specific performance criteria.
- A 90-day post-live review with Finance co-signature on the results.

Research indicates companies with clearly established baselines are three times more likely to achieve positive AI investment returns.

Structure Payment to Performance Milestones

Milestone-based payment schedules tied to go-live and 90-day performance validation protect the organization and create the right incentives for the vendor. This is especially critical in agentic deals where the vendor is selling outcomes but pricing on activity. A vendor that cannot define what a successfully completed workflow looks like before signing is telling you something important.

The Pilot Trap

More than 70 percent of enterprises have run AI pilots, but fewer than 20 percent push them into full production. The reason is almost always the same: the pilot was not structured to generate the validated data needed to justify the production investment. Agentic AI pilots are particularly vulnerable to this because the technology is genuinely impressive in demonstration and genuinely difficult to scale without clean data, defined scope, and governance infrastructure that most organizations have not yet built. Procurement should design the pilot to answer the questions Finance will ask, not just the questions the technology team finds interesting.

> ***A vendor that cannot define what a successfully completed workflow looks like before signing is telling you something important.***

Section 6: Contract and Governance Basics

Every AI contract, regardless of type, should include these provisions before signature.

No Training on Your Data

The contract must explicitly prohibit the vendor from training or fine-tuning models on your data without prior written approval. This restriction applies to every vendor regardless of how reasonable their data policies appear at first read.

Data Residency in Writing

Data residency is no longer optional. It is a baseline requirement for any AI tool processing commercial agreements or sensitive business information. If your organization operates in regulated industries or serves customers subject to data protection laws, this provision is not negotiable.

Breach Notification With a Timeline

Security incidents affecting your data should require vendor notification within 24 hours, with documented mitigation steps. Standard contract language referencing a reasonable timeframe is not sufficient.

Model Change Notification

Vendors may change the underlying AI model without informing customers, which can alter how the system behaves and affect the quality of outputs your organization depends on. Require advance notice of major model changes and a testing window before they go into production.

Exit Provisions

Require that your data be returnable in a usable format upon contract termination and that proprietary dependencies do not make migration prohibitively expensive. Include a clause allowing mid-term rate renegotiation if usage exceeds agreed thresholds.

Governance Cadence

Set a quarterly business review with vendor and Finance representation. Track performance against the baseline established at signing. Documented, repeatable, and co-signed by Finance. This is the same governance model that makes procurement savings credible to the chief financial officer, and it works the same way in AI.

What To Do With Each Type of AI: A Practical Checklist

The checklist below is designed to be used at the negotiation table. Apply the relevant section to every AI deal your organization reviews.

Type 1: Direct AI Infrastructure Contracts

✓ Require itemized token pricing for every model you plan to use. Input, output, and reasoning tokens listed separately.

✓ Negotiate volume commitment tiers and lock in discount thresholds before you scale.

✓ Include a usage cap with a renegotiation right if consumption exceeds agreed thresholds.

✓ Require prompt caching discounts to be specified in the contract, not assumed.

✓ Convert token cost to a cost-per-transaction metric Finance can benchmark.

✓ Prohibit use of your data for model training in the data processing addendum.

✓ Require advance notice of any model version changes before they go into production.

Type 2: AI-Powered Application Contracts

✓ Require the vendor to disclose which underlying AI model powers the product and who owns that relationship.

✓ Demand visibility into the token cost layer embedded in your subscription or usage fee.

✓ Establish a documented baseline of current-state process metrics before go-live.

✓ Tie payment milestones to go-live performance criteria, not just software delivery.

✓ Require a 90-day post-live return on investment review co-signed by Finance.

✓ Confirm data residency, breach notification timelines, and training prohibitions in the contract.

✓ Build in exit provisions that allow data portability without penalty.

Type 3: Agentic AI Contracts

✓ Negotiate on outcomes, not tokens. If the vendor cannot define what a successfully completed workflow looks like, the deal is not ready to sign.

✓ Require a pre-agreed measurement framework including baseline, go-live criteria, and Finance validation at 90 days.

✓ Insist on on-premise or private cloud deployment for any workflow that touches confidential data.

✓ Define the scope of what the AI is authorized to do and not do in writing. Scope boundaries belong in the contract, not just in the configuration settings.

✓ Require human review checkpoints for any workflow that involves financial commitments, supplier decisions, or regulatory obligations.

✓ Build in a pilot-to-production gate with Finance co-signature before full deployment funding is released.

Type 4: Embedded AI in Existing Software

✓ Review updated vendor terms at every renewal for expanded data rights, including any language granting the vendor rights to use your data for model training or improvement.

✓ Require the vendor to quantify the value of any AI capability bundled into a price increase separately from the base product renewal.

✓ Confirm that the AI feature's data processing complies with your existing data residency requirements. Embedded AI sometimes uses separate infrastructure.

✓ Negotiate opt-out rights for AI processing on sensitive data types and workflows without penalty.

✓ Document whether the AI feature is replacing a process or augmenting one. Only replacement generates a Finance-validated return on investment.

✓ Confirm the existing data processing agreement explicitly prohibits model training on your data and update it if it does not.

Conclusion

I started this paper feeling genuinely lost. Sitting in governance meetings I could not contribute to, approving contracts I could not properly evaluate, and watching the same structural failures I had spent my career fixing show up in a brand new category.

What I found, once I took the time to understand the landscape, is that AI is not a different discipline for procurement. It is the same discipline applied to an unfamiliar commercial model. The fundamentals that made me effective as a Chief Procurement Officer at Humana and Centene, validated baselines, Finance partnership, governance before go-live, contractual accountability, apply here exactly as they apply everywhere else.

The embedded AI scenario deserves special attention because it is the one most likely to bypass procurement entirely. When your existing vendors embed AI into products you already own, no alarm goes off. No purchase order is created. No contract review is triggered. The risk lands in a renewal conversation that looks routine, and the data rights you give up in exchange do not appear on any invoice. That is the gap this paper is designed to close.

Procurement that negotiates AI contracts the way Finance expects them to be negotiated, that insists on private deployment for sensitive workloads, that ties payment to validated outcomes, and that builds governance before go-live will not just protect the organization. It will earn the credibility the function has always deserved.

You do not need to be a technologist to do this well. You need to be a good procurement professional. That part you already know how to do.

> ***Procurement credibility is not built in the boardroom. It is built in the contract, the Finance review, and the 90-day performance gate. Get that process right and the seat at the table follows.***

References and Sources

All opinions expressed in this paper are the author's own. Sources are cited for statistical claims and factual assertions only.

[1] Futurum Group. '1H 2026 Enterprise Software Decision Maker Survey Report.' February 2026. futurumgroup.com

[2] AIMultiple Research. 'LLM Pricing: Top 15+ Providers Compared.' 2025. research.aimultiple.com

[3] XAiGi Tech. 'The Price of Intelligence: Mastering LLM Pricing and Enterprise AI Cost Optimization in 2025.' October 2025. xaigi.tech

[4] Binadox. 'LLM API Pricing Comparison 2025: Complete Cost Analysis Guide.' August 2025. binadox.com

[5] Redress Compliance. 'The Enterprise Guide to Negotiating OpenAI Contracts.' August 2025. redresscompliance.com

[6] Lasso Security. 'LLM Data Privacy: Protecting Enterprise Data in the World of AI.' February 2026. lasso.security

[7] CIO Dive / Cybersecurity Dive. ‘Samsung employees leaked corporate data in ChatGPT: report.’ April 2023. cybersecuritydive.com

[8] TechRadar. ‘Samsung workers made a major error by using ChatGPT.’ April 2023. techradar.com

[9] SamMobile. ‘Samsung lets employees use ChatGPT again after secret data leak in 2023.’ May 2025. sammobile.com

[10] IJISRT. ‘Entropy-Based Obfuscation Framework for Multi-Tenant LLM Security.’ Volume 10, Issue 9, September 2025. ijisrt.com

[11] Zhang, Q. et al. ‘Burn-After-Use for Preventing Data Leakage through a Secure Multi-Tenant Architecture in Enterprise LLM.’ arXiv:2601.06627, January 2026. arxiv.org

[12] QAnswer. ‘Rethinking Enterprise LLM: Secure, Cost-Effective AI.’ 2025. qanswer.ai

[13] Accrets Cloud. ‘The Executive Playbook to On-Premise LLM Deployment in 2026.’ February 2026. accrets.com

[14] Harvard Business Review. As cited in Monetizely. ‘How to Measure ROI for AI Agent Implementations.’ July 2025. getmonetizely.com

[15] Ampcome. ‘Enterprise AI Agents ROI Framework: The 2025 Guide.’ October 2025. ampcome.com

[16] Tascon Legal. ‘AI Clauses in Contracts: The Practical Guide for 2025.’ 2025. tasconlegal.com

[17] European Business Magazine. ‘Only 35% of EU Firms Can Use AI Contract Tools Without Breaking GDPR.’ February 2026. europeanbusinessmagazine.com

[18] Redress Compliance. ‘AI Procurement in 2025: How OpenAI Is Changing Software Negotiations Forever.’ August 2025. redresscompliance.com

Appendix: Key Terms

These definitions are provided for reference. You do not need a technology background to apply the frameworks in this paper, but understanding these terms commercially will sharpen every negotiation conversation.

Artificial Intelligence	Software that can perform tasks that normally require human thinking, such as reading documents, answering questions, writing content, or making recommendations based on data.
Large Language Model	The engine that powers most modern AI tools. It is trained on massive amounts of text and learns to understand and generate language. When you use ChatGPT, Claude, or Gemini, you are using a large language model. Think of it as the brain behind the product.
Generative AI	AI that creates new content including text, summaries, analysis, and recommendations, rather than just retrieving stored information. Most enterprise AI tools today are generative AI applications built on top of a large language model.
Agentic AI	The next generation of AI that does not just answer questions but takes action autonomously. An agentic system can receive a goal, break it into steps, execute those steps across multiple systems, and complete a workflow without human input at each stage. Examples include AI that can run a sourcing event, onboard a supplier, or process invoices end to end.
Embedded AI	Artificial intelligence capabilities added by a vendor to a software product you already own, either included in your existing contract, bundled into a renewal,

	or offered as an upsell module. No new vendor relationship is created. The risk shows up in your renewal terms, not a new purchase order.
Token	The unit of measurement AI language models use to process text. A token is roughly three to four characters. When you send a question to an AI and it receives a response, both are measured in tokens and billed accordingly. This is how most direct AI infrastructure contracts are priced.
Input Token	Text that you send to the AI, such as a question, a document, or a prompt. Input tokens are cheaper to process.
Output Token	Text the AI generates in response. Output tokens cost more because generating a response requires significantly more computing power than reading an input.
Reasoning Token	An additional token type used by more advanced AI models that think through complex problems before responding. These are billed separately and can meaningfully increase cost in contracts that do not account for them.
Multi-Tenant	A cloud setup where multiple organizations share the same underlying infrastructure. Most commercial AI services are multi-tenant by default, meaning your data is processed in a shared environment that other companies also use.
Single-Tenant	A dedicated infrastructure environment where only your organization's data is processed. More secure and more expensive. Often available as a premium option from major AI vendors.
On-Premise Deployment	Running AI software on servers your organization owns and controls, inside your own network. Data never leaves your environment. The highest level of data security available.
Data Residency	The requirement that data be stored and processed within a specific geographic location or jurisdiction. Critical for organizations subject to privacy laws including those covering health information.
Data Processing Addendum	A legal agreement between your organization and a vendor that specifies how your data will be handled, stored, protected, and deleted. Every AI contract should include one.
Model Training	The process by which an AI learns from data. If a vendor uses your data to train or improve its model, your proprietary information can become part of that vendor's knowledge base. This must be explicitly prohibited in the contract.
Fine-Tuning	Training an AI model further on your organization's specific data to improve its performance for your use cases. Creates significant data ownership and privacy considerations that must be addressed in the contract.
Service Level Agreement	Contractual commitments from the vendor on uptime, response time, and performance. AI services should be held to the same performance standards as any other enterprise software.
Total Cost of Ownership	The full cost of an AI solution over its lifecycle, including licensing or token costs, implementation, integration, training, maintenance, and the cost of any data security infrastructure required.

About the Author

Chad Johnson is a former Chief Procurement Officer at Humana and Centene, where he led enterprise procurement transformations delivering $650M+ in validated savings across Fortune 20 healthcare organizations. He previously

held global category management and procurement leadership roles at Pfizer and McKesson. All opinions expressed in this paper are the author's own. This paper is intended for educational purposes.