



Agentic AI in Procurement

Post-Event Report

Chief Clubhouse, San Francisco, CA

May 14, 2026

Conversational intelligence tool

contented

Foreword

It was a genuine pleasure to welcome such an engaged group to our latest event, and the energy in the room was clear from the start of the day.

A few themes stood out: AI adoption works best when it starts from a real problem and an existing workflow rather than the technology itself, and the most valuable use cases are the ones embedded directly into how people already work. We also heard a recurring message that the bigger challenge isn't the tools, it's readiness, trust, and giving people the confidence to experiment. The podcast conversation on AI infrastructure and data centre economics struck a particular chord, opening up a topic that many in the room are only just beginning to think about, and one we expect to return to.

None of this would have been possible without the generosity of our speakers, who shared genuinely candid insights and frameworks; our sponsors, whose support continues to make these events possible; and our attendees, who brought sharp questions, honest debate, and a real appetite to keep pushing the agentic AI conversation forward.

Thank you all for making the day so valuable and we look forward to continuing the conversation with you.



Claire Barnes

Co-Founder, KonnectHouse

NB: All names of individuals, companies, sponsors and vendors referenced in the original recordings have been removed or replaced with generic role descriptions.

“

“The difference between someone who knows how to bring their art to the technology and someone who is relying on the technology to make art is the difference between music and noise.”

Thank You to the Event Sponsors



Event Summary Report

Executive Summary

Agentic AI in Procurement, San Francisco brought together procurement, technology, AI, supplier risk, transformation and infrastructure perspectives around one central question: how should procurement evolve in an AI-enabled world?

The strongest message from the day was that AI is not simply a new toolset for procurement. It is reshaping the role, operating model, skills, governance expectations and value proposition of the function. Participants repeatedly returned to the same themes: procurement must move from process ownership to outcome ownership; from tactical execution to orchestration; from category policing to enterprise value creation; and from experimenting with AI to embedding it responsibly into workflows.

The event also surfaced an important tension. There is strong excitement about agentic AI, automation and intelligent orchestration, but there is also anxiety around governance, data quality, compliance, workforce impact, model accuracy, auditability and infrastructure constraints. The most credible path forward is not “AI everywhere”, but targeted, problem-led adoption with clear human ownership, strong governance, measurable value and continuous learning.

The day closed with a broader discussion on the physical infrastructure behind AI. Data centres, power availability, water usage, edge computing and device-level AI are no longer abstract technology issues. They are becoming procurement, risk, sustainability and negotiation issues.



Main Body of Learnings

1.1 Procurement's Evolution and the Case for AI

- Procurement has evolved from a manual, ERP-driven back-office function (illustrated through a nostalgic walk-through of legacy ERP processes that took 20–30 minutes per purchase order) to an increasingly automated, strategic function.
- Despite this evolution, many practitioners still feel procurement has not fully earned its seat at the executive table; the function continues to be siloed under finance, IT or operations in many organisations.
- The CPO role is gradually moving into the C-suite alongside CTO/CFO peers, driven by the scale of spend under management, the savings opportunity, and rising data-privacy and litigation risk.
- A consistent post-pandemic shift was noted: a move away from “growth at all costs” toward cost discipline and profitability, which has elevated procurement's strategic relevance.

1.2 How AI Is Reshaping the Procurement Team

- Teams are becoming leaner and more cross-functional, with technical/transactional work increasingly automated and human effort redirected toward strategy, sustainability, and supplier relationship management.
- A recurring theme: most organisations manage only a small fraction of their total supplier base actively — AI is seen as a way to extend meaningful relationship management across a much larger supplier population.
- Skill-gap concerns were raised but tempered by a “behaviour over skills” hiring philosophy: adaptability and willingness to learn matter more than current technical proficiency.
- A practical example given: AP invoice-matching staff who previously spent ~30 hours/week on manual matching could see that workload reduced to a handful of exceptions, freeing time for higher-value analysis — provided the freed-up time is deliberately reinvested rather than simply absorbed.

1.3 Where AI Is Already Delivering Value

- AI use cases were grouped into two buckets: (1) accelerating tasks already being done (e.g. invoice matching, contract review), and (2) enabling tasks that previously weren't done at all due to time constraints (e.g. systematic adverse-media/supplier-risk screening).
- Intake validation was highlighted as a high-value use case: AI can flag inconsistencies in intake forms (e.g. a vendor contract implying data-sharing when the requester indicated none), improving the quality of data flowing into downstream workflows and building trust in automated routing.

- Exception-based working was repeatedly cited as the target operating model — AI surfaces the handful of genuinely risky transactions out of large volumes (e.g. T&E, AP), allowing scarce expert resource (especially security/compliance) to focus where it matters.
- A concrete contract-risk example: AI comparing a new vendor MSA against the prior version can detect newly introduced AI-related clauses and automatically route the contract to security review.
- The single biggest factor determining whether AI tools get used was described as embedding AI directly into existing workflows, rather than requiring staff to manually copy data between a chat interface and other systems.

1.4 Compliance, Governance and the “Accuracy Bar” Debate

- Compliance/governance was repeatedly named as the top concern around AI adoption, especially for regulated and publicly listed organisations (SOX, audit, contract-compliance obligations).
- A live debate emerged over what accuracy standard AI should be held to: should it be held to a near-zero-error bar (the autonomous-vehicle analogy — a single high-profile accident leading to a programme being shut down despite a strong overall safety record), or judged against the (often much lower) accuracy of the human process it replaces?
- Practical governance question raised: which controls should allow AI to take autonomous action versus only provide recommendations, and how will auditors view AI-driven decisions in regulated environments?
- Consensus: organisations need a clear, risk-tiered framework for where full automation is acceptable, where human-in-the-loop is required, and where AI is excluded entirely (“friction gates”).

1.5 The “Future-Proof Practitioner” Framework

- A detailed framework was presented describing procurement work as a pyramid: transactional work (largely automatable), cognitive/decision-support work (significantly impacted by AI), and strategic work (judgment, relationships, negotiation — remaining distinctly human, though augmented).
- A stark figure was cited: an estimated 60–80% of practitioners are not currently ready for this shift, and the function itself is expected to shrink and change shape.
- A three-layer capability model for individuals was proposed: (1) a foundational layer of “procuring knowledge” and guarding against cognitive debt (think before prompting, do your own analysis first, set spend/category thresholds where humans must lead, periodically practise tasks manually); (2) an enabling layer balancing AI literacy with disciplined use; (3) a differentiating layer comprising orchestration (systems thinking, sequencing work across humans/agents/systems), business acumen (market, strategy, customer, operations, internal stakeholder and supplier-ecosystem literacy), and human leverage (self-awareness, relationship management, persuasion/storytelling, negotiation, creative problem-solving, and continuous curiosity).
- Closing advice to attendees: don’t wait for the organisation to drive this — individuals should start experimenting now, tolerate failure, assess their own role against the risk framework, and focus on orchestration and outcomes rather than tools.

1.6 Laying the Foundations — The “Three Mirrors” and ABCDFG Models

- A “three mirrors” framework for starting an AI programme: (1) look at the problem first — never start with “let’s use AI” before defining what is broken or costly; (2) look at your people — build confidence before competence, since staff are often “emotionally retired” the moment automation is announced; (3) look at your process — redesign broken processes before automating them, since AI will simply execute a bad process faster.
- A live survey statistic was shared (and challenged) suggesting roughly 60% of companies expect to reduce headcount for staff not using AI — with a counter-prediction that future workforce reductions may instead target over-reliance on AI without sufficient human oversight.
- On the “clean data” debate: one school of thought says fix data quality before deploying AI; the other says start small, let AI use deployment expose where the data problems actually are, and clean iteratively (the analogy used: a baby doesn’t wait for a clean floor before learning to walk).
- A memorable change-management framework, “ABCDFG”, was shared: Alignment with a sponsor; Belongingness (involve people from day one); Communication; Data and execution; Fail fast (treat failure as a companion of innovation, not an end-state); and Govern/celebrate small wins.
- Definition of AI success was reframed around three measures: adoption as habit (not just “it works”), whether time saved is genuinely reinvested into higher-value work (not just absorbed), and whether decision quality and practitioner confidence improve.
- A “digital mindset” was defined using the Nokia-versus-smartphone analogy: the critical skill is not learning new tools but unlearning old assumptions quickly, moving from a perfectionist mindset to an experimental one, and shifting focus from owning tasks to owning outcomes.

1.7 Scaling Agentic AI Without Scaling Headcount (Panel)

- The dominant theme across panellists was “more with less”: stabilising headcount while increasing output, using a combination of agentic AI and improved process/workflow design.
- On build-vs-buy for agents: panellists were divided. Some favoured a hybrid model — using a platform/vendor backbone for stability and embedded best-practice knowledge from across many customers, supplemented by custom-built agents for specific needs. The risk of a pure “build it yourself” approach was likened to a chick taking short flights — workable for a while, but raising sustainability questions over time.
- A vendor perspective argued that off-the-shelf platforms bring an aggregated base of cross-industry experience that a single organisation cannot replicate by building agents internally, even with capable general-purpose AI coding tools.
- On who should own enterprise AI/agent orchestration: views diverged between IT (security, data, technical ownership) and the business function itself (procurement, finance, operations), with some panellists proposing a dedicated “centre of intelligence” role that bridges business problem-owners and technical capability-providers. No consensus emerged — described explicitly as “an open question”.

- On governance: organisations are at very different maturity levels — from no formal AI access policy, to curated lists of approved tools, to (at the more aggressive end) companies marketing “unlimited AI token budgets” as a recruitment perk. AI cost forecasting was described as extremely difficult, even for large enterprises.
- On workforce sentiment: panellists noted growing pressure from boards (often driven by general business press coverage rather than deep technical understanding) for every department to “do something” about AI, with the practical response varying widely by industry and company maturity.

1.8 Creativity and AI

- Giving teams free rein to build AI tools without direction, which initially produced a lot of low-value output (“noise”) before more focused, valuable use cases emerged — framed as a necessary part of the learning curve.
- Closing audience poll: when asked what attendees would do differently in the next 90 days, the most common responses were aligning stakeholders/leadership around AI, and launching or expanding a pilot use case.

“

“In procurement we have convinced ourselves that we cannot take any decision until we have clean data and I call it the most sophisticated procrastination strategy.”



SWOT Analysis

Strengths	Weaknesses
• Strong practitioner-led knowledge base across the event, with speakers drawing on decades of real-world procurement, supply chain, and technology experience • Clear articulation of AI's practical value in procurement — from invoice matching and contract review to supplier due diligence — with concrete examples shared by multiple presenters • Recognition that craft, empathy, and point of view (not just technology) drive meaningful outcomes, grounding AI adoption in human-centered thinking • Demonstrated 80% reduction in time for supplier due diligence when AI is applied with domain expertise and customer empathy (Freedom) • Behavioral agent automation approach (Liminal) that observes user patterns and builds automations without requiring users to become engineers • Strong awareness among attendees of AI governance, data privacy risks, and the need for human oversight in regulated environments • Progressive, self-selected audience already actively piloting AI tools and thinking critically about procurement's strategic future	• 60–80% of procurement practitioners are estimated to be unprepared for the skill and role changes AI will bring • Many AI initiatives within organizations produce “noise” rather than value — teams building tools they like rather than tools customers actually need • Data quality remains a persistent barrier; poor or unclear data fed into AI systems compounds problems rather than solving them • Lack of clarity on who owns AI orchestration within organizations — IT, procurement, finance, or a transformation team — creates governance gaps • Token-based AI pricing models are unpredictable and difficult to forecast, making budget planning and CFO conversations challenging • Shadow AI usage is widespread, with 75% of enterprise data breaches now reportedly occurring through model providers, often driven by employees using personal tools when corporate options are inadequate • Junior procurement professionals face significant risk of being left behind as transactional and cognitive decision-support work shifts to machines

“

“The vast majority of (AI) users are still at the crawl stage.

SWOT Analysis

Opportunities	Threats
• AI creates a real path for procurement to shift from spend policing and firefighting to genuine strategic value creation and enterprise-level influence • Agentic AI platforms that observe behavioral patterns and automate repetitive workflows can free up significant practitioner time for higher-value work • Mid-market organizations (500–5,000 employees) represent an underserved segment with strong demand for secure, governed, cost-effective AI tooling • The Human Edge Matrix framework offers a structured way for procurement leaders to assess which roles and tasks remain human, enabling smarter workforce planning • Outcome-focused operating models — leaner, more modular teams orchestrating humans, agents, and systems — represent a significant redesign opportunity for procurement functions • Edge computing and on-device AI models may reduce dependence on expensive cloud infrastructure, lowering long-term token costs for enterprises • Organizations that invest now in AI literacy, cognitive discipline, and human skills (influence, negotiation, stakeholder management) will build a durable competitive advantage • Intake validation and workflow orchestration tools can dramatically improve front-end data quality, driving better downstream outcomes across invoicing, compliance, and supplier management	• AI model providers' interests are not aligned with enterprise data privacy — overt, covert, and in some cases illegal data retention practices create real regulatory and reputational exposure • Subsidized AI pricing models are unsustainable; the shift to token-based billing (as signaled by GitHub's June 2026 announcement) will significantly increase costs for many organizations • Goldman Sachs projections suggest AI token costs could reach 100% of headcount costs for organizations that lean heavily into AI — meaning the tool meant to reduce costs may match or exceed the cost of the people it replaces • Open-license AI policies are as risky as no policy at all — without governance, employees will take data off-network and use personal tools, increasing breach exposure • CFOs who don't see procurement's strategic value may opt to automate 60–70% of the function at a fraction of the cost, cutting headcount without capturing the innovation and judgment that human practitioners provide • The procurement function risks losing the growth trajectory it has built over the past 30 years if practitioners don't proactively develop the skills and mindset needed for a post-AI environment • Community and political resistance to data center expansion, combined with power grid constraints and water usage concerns, may slow AI infrastructure growth and drive up costs unpredictably

2. Podcast on AI Economics — Data Centres, Power and Pricing

The live podcast discussion explored the economic pressures behind the AI boom, focusing on pricing, infrastructure, power demand, and the future location of AI compute.

A key takeaway was that current AI pricing models are unlikely to remain sustainable. Speakers suggested that many AI subscriptions are effectively subsidised, with the market beginning to shift from flat-rate pricing towards consumption, token-based, outcome-based, or blended models. This creates new challenges for buyers, as AI usage and infrastructure costs are difficult to forecast due to the variable compute required for different tasks.

The discussion also highlighted power availability as a major constraint on AI growth. Access to electricity is increasingly becoming more important than cost for data-centre operators, while water use for cooling and local environmental concerns are driving greater community and political pushback against new data-centre developments.

Alternative infrastructure models were discussed, including space-based data centres and distributed residential compute, but both were viewed as highly experimental and subject to major practical, security, and trust barriers. Audience polling suggested limited appetite for hosting compute infrastructure in private homes.

A counter-view was that more AI workloads may move away from centralised data centres and towards edge computing, including smartphones, wearables, and on-premise enterprise systems. The likely future is therefore hybrid: routine AI tasks may increasingly run locally, while intensive generative workloads remain cloud-based.

Overall, the session framed AI economics as entering a period of recalibration, where pricing, power access, infrastructure strategy, and public acceptance will all shape the next phase of AI adoption.

“

“AI model providers are not aligned with your interests. The number one thing standing between them and the next generation model is better data.”



3. Key Learnings & Highlights

- AI adoption succeeds when it starts from a clearly defined problem and existing workflow, not from the technology itself — multiple speakers independently arrived at variations of “start with the problem, not the tool”.
- The “exception-based working” model (AI surfaces the few genuinely risky items; humans focus there) was the most concretely actionable and broadly endorsed operating pattern discussed.
- Embedding AI directly into existing workflows — rather than as a separate chat tool requiring manual data transfer — was repeatedly identified as the single biggest determinant of real adoption.
- There is a meaningful gap between enthusiasm for AI and organisational readiness: the 60–80% “not ready” figure, combined with reports of emotionally disengaged staff at the moment automation projects are announced, points to change management as the binding constraint, not technology.
- The day surfaced a genuinely novel and timely topic for the procurement community: AI/data-centre economics, token-based pricing shifts, and the link between electricity/power access and enterprise AI strategy — an area most procurement teams are only beginning to grapple with.
- High engagement signals: live polling throughout the day, a high proportion of attendees already paying for multiple AI subscriptions personally, and an enthusiastic response to the live podcast-recording format.



4. Risks & Open Questions Raised

- Governance and compliance uncertainty: no clear consensus exists on what accuracy/error threshold AI should be held to in regulated environments, or how auditors will treat AI-assisted or AI-driven decisions.
- Cost unpredictability: the shift toward token/consumption-based AI pricing introduces budgeting and forecasting risk that procurement teams are not yet well equipped to manage — several real-world examples of budgets being exhausted early were shared.
- Vendor/ecosystem concentration and dependency risk: the financial interdependencies between AI labs and infrastructure providers were flagged as a systemic risk that could affect pricing and availability if growth assumptions don't materialise.
- Ownership ambiguity: the question of who owns AI agent orchestration within an enterprise (IT vs. business function vs. a new “centre of intelligence” role) remains unresolved, creating risk of duplicated effort, shadow IT (“25 different skills doing the same thing”), and unclear accountability.
- Workforce polarisation risk: a divide between staff who have embraced AI and those who haven't was described as a growing internal risk, with implications for both productivity gaps and employee relations/retention.
- Infrastructure and community risk: data centre projects face real-world delivery risk from power constraints, water/environmental concerns, and local political opposition — all factors procurement and sourcing teams may increasingly need to factor into vendor and location risk assessments.
- Speculative/early-stage alternative infrastructure models (space-based and residential/distributed data centres) carry significant unresolved security, liability and trust barriers, and should be treated as directional signals rather than near-term sourcing options.

“

“Goldman Sachs has reported that some businesses are spending 10% of headcount costs on AI tokens. That's projected to rise to 100% of the cost of headcount. So the tool that was supposed to replace people is now costing as much or more than the people.”

KEY RISK FINDINGS

Risk 1: AI Data Privacy and Unauthorized Data Retention by Model Providers

Consequence: Major

Multiple speakers flagged that AI model providers — including OpenAI, Anthropic, and Google — are retaining user data in ways that organizations may not fully understand or have consented to. Specific examples included Anthropic updating its terms of service to retain data for up to five years and use it for model training, the federal government requiring OpenAI to indefinitely retain API logs, and Google losing a \$400 million class action lawsuit over 20 years of undisclosed Android data collection. For organizations handling sensitive commercial data, trade secrets, patient records, or regulated financial information, this creates real exposure. A hospital system using OpenAI to support patient diagnosis, for example, could face HIPAA violations if that data ends up retained and accessible by third parties.

Likelihood: Almost Certain

This isn't a theoretical risk. The examples cited in the transcript have already occurred. The incentive structure for model providers is misaligned with enterprise data protection — they need user-generated data to train next-generation models, and that pressure isn't going away. Unless organizations actively implement governance controls, data will continue flowing to these providers.

Overall Risk Level: Extreme

The combination of a major consequence and an almost certain likelihood makes this extreme. Regulatory exposure, reputational damage, and potential legal liability are all live risks for any organization using frontier AI tools without proper data governance in place.

Risk 2: Shadow AI and Uncontrolled Employee Use of External AI Tools

Consequence: Major

When organizations either restrict AI access too tightly or fail to provide adequate tools, employees find their own solutions. The transcript notes that 75% of enterprise data breaches are now occurring through model providers, largely driven by employees taking data off-network and using personal or unsanctioned AI tools. This creates a situation where sensitive business data — contracts, financial models, supplier information — is being processed through platforms the organization has no visibility into and no governance over.

Likelihood: Likely

The transcript describes this as an active and widespread pattern, not a future concern. With AI tools freely available and employees under pressure to perform, the behavior is already happening across industries. Organizations that haven't addressed this are almost certainly already experiencing it to some degree.

Overall Risk Level: High

A major consequence paired with a likely likelihood lands this firmly in the high category. The risk is compounded

by the fact that it's largely invisible — organizations often don't know it's happening until a breach surfaces.

Risk 3: Unsustainable AI Pricing Models and Unpredictable Token Costs

Consequence: Moderate

AI model providers have been operating on subsidized pricing that doesn't reflect actual infrastructure costs. GitHub's announcement of a move to token-based billing — the first major public pivot of this kind — signals that the current pricing model is changing. The transcript references Goldman Sachs projections that AI token costs could reach 100% of headcount costs for businesses that lean heavily into AI. Uber reportedly burned through its entire annual AI token budget in just a few months. For procurement teams trying to build business cases, forecast spend, and negotiate contracts, this unpredictability creates real planning and budgeting challenges.

Likelihood: Almost Certain

The shift is already underway. GitHub has announced it publicly, effective within weeks of the event date. Other providers are expected to follow. The subsidized model was never sustainable, and the market is correcting.

Overall Risk Level: High

Even though the consequence is moderate rather than major, the near-certainty of this shift happening — and the difficulty of forecasting costs in a token-based world — pushes this to high. Organizations that haven't started thinking about consumption governance and AI spend management are going to be caught off guard.

Risk 4: Data Breach Risk from AI Model Provider Infrastructure

Consequence: Major

The transcript cites 1,400 data breaches across the top five model providers in a single year. When organizations send data to these platforms — whether through APIs or end-user interfaces — that data is retained and becomes subject to the provider's security posture. A breach at the provider level can expose trade secrets, customer data, financial information, and regulated data that the organization had no intention of making accessible to anyone outside its walls.

Likelihood: Likely

Given the volume of breaches already documented and the scale at which data is being retained across these platforms, this is not a remote possibility. The more data flowing through these systems, the larger the attack surface becomes.

Overall Risk Level: High

Major consequence and likely likelihood combine to make this a high-rated risk. The organization has limited direct control over the security practices of third-party model providers, which makes mitigation dependent on governance decisions made before data ever leaves the organization.

Risk 5: Procurement Workforce Readiness Gap in an AI-Driven Environment

Consequence: Major

Omar Abdullah's presentation made a pointed case that 60–80% of procurement practitioners are not ready for the changes AI is bringing to the function. Transactional work is already being automated. A significant portion of cognitive decision-support work — RFP management, spend analysis, data synthesis — is following. What remains is strategic activity, and many practitioners don't yet have the skills to operate effectively at that level. If the function can't demonstrate strategic value, CFOs will make the calculation that machines can deliver 60–70% of procurement's output at a fraction of the cost, and headcount will follow.

Likelihood: Likely

The informal surveys referenced in the transcript, combined with the observable trend of procurement functions shrinking and restructuring, suggest this gap is real and present. It's not a future problem — it's already affecting hiring decisions, role definitions, and how procurement is perceived at the leadership level.

Overall Risk Level: High

The consequence is major — this is an existential question for the function — and the likelihood is high enough that organizations can't afford to treat it as a long-term concern. The window to build capability and reposition the function is narrowing.



Risk 6: AI Governance and Compliance Failures in Regulated Industries

Consequence: Major

Several speakers highlighted the tension between deploying AI quickly and maintaining compliance in regulated environments. Industries including banking, healthcare, insurance, and life sciences face regulatory bodies that can take action when data is misused or when AI-driven decisions aren't properly governed. The transcript specifically calls out HIPAA violations as a live risk for healthcare organizations using frontier AI tools, and notes that SOX compliance in procurement processes requires careful thinking about which controls AI can touch versus which require human accountability.

Likelihood: Possible

Not every organization will face a regulatory action, but the conditions for it are present. Many organizations are still in early-stage AI adoption without mature governance frameworks. The gap between what's technically possible and what's compliant is real, and regulators are paying attention.

Overall Risk Level: High

A major consequence — regulatory penalties, reputational damage, legal liability — combined with a possible likelihood is enough to warrant a high rating. The risk is particularly acute for organizations in regulated sectors that are moving fast on AI without a corresponding investment in governance infrastructure.

“

“If the data is not available, cleansed, it's the old garbage in, garbage out scenario. If we're extracting data from your source systems and that data is not clean, it's not been normalized, it's not a package that can be ingested by an AI and utilized correctly, you really just compounded your problems as you go down the line.”

Risk 7: AI Tool Adoption Without Clear Problem Definition

Consequence: Moderate

Ashish Gupta made the observation that “the most expensive sentence in the English language is ‘let's use AI.’” The pattern he described — organizations deploying AI tools before identifying the problem they're trying to solve — leads to low adoption, wasted investment, and what he called “technology zombies.” Justin's presentation echoed this from a different angle: his team built things they were excited about that didn't actually serve customer needs. The result is spend on tools that don't deliver ROI and a workforce that loses confidence in AI initiatives.

Likelihood: Likely

This is described as a common and recurring pattern across organizations of all sizes. The excitement around AI creates pressure to deploy something, and that pressure often outpaces the strategic thinking needed to deploy something useful.

Overall Risk Level: High

Moderate consequence and likely likelihood combine to make this high. The financial waste is real, but the deeper risk is organizational — failed pilots erode trust in AI initiatives and make it harder to get buy-in for future investments that might actually deliver value.

Risk 8: Dependency on a Small Number of AI Model Providers**Consequence: Major**

The transcript raises the financial fragility of the current AI ecosystem. OpenAI reportedly needs \$852 billion in revenue to service the Stargate project. If OpenAI can't meet its obligations to Oracle, Oracle can't service its debt. The dependency chain runs deep, and a significant portion of enterprise AI infrastructure is built on top of a small number of providers whose long-term financial viability is not guaranteed. Organizations that have embedded these tools deeply into their workflows face significant disruption if a provider fails, pivots, or dramatically changes its terms.

Likelihood: Possible

The financial projections cited in the transcript are concerning, but these are large organizations with significant backing. A complete failure is not imminent, but a major restructuring, acquisition, or service disruption is plausible within a multi-year horizon.

Overall Risk Level: High

The consequence of a major provider failure or disruption would be significant for organizations that have built workflows around those tools. Paired with a possible likelihood, this sits at high. Vendor diversification and contingency planning are reasonable mitigations that most organizations haven't yet prioritized.

“

“You don’t clean data before you use it, you clean data when you use it.”

Risk 9: Energy and Power Constraints Limiting AI Infrastructure Growth

Consequence: Moderate

The data center discussion flagged that new data centers face five to seven year grid connection waits in major US markets. Power availability — not compute or data — is now the primary constraint on AI infrastructure expansion. This has downstream implications for AI service availability, pricing, and the pace at which organizations can scale AI-dependent operations. Communities are also pushing back against data center construction, adding political and regulatory friction to an already constrained supply picture.

Likelihood: Almost Certain

This is not a projected risk — it's a current reality. PG&E reportedly wouldn't upgrade a residential service connection, and the conference discussion confirmed that power availability is already the number one consideration in data center siting decisions, with cost a secondary concern.

Overall Risk Level: High

A moderate consequence combined with an almost certain likelihood produces a high rating. The constraint is real and structural, and it will affect the cost and availability of AI services for the foreseeable future. Organizations planning AI infrastructure investments need to factor this in now.

“

“You need to build confidence before competence. Let people try it out, let people edit, let people see how it's helping them so that the confidence is good.”

Risk 10: Junior Talent Development Gap Caused by AI Automation

Consequence: Moderate

Omar Abdullah raised a concern that doesn't get enough attention: the impact of AI on junior practitioners. Historically, junior roles have been the training ground where people develop judgment, category knowledge, and business acumen through hands-on work. As AI takes over the transactional and cognitive decision-support tasks that junior staff typically handle, the pipeline for developing experienced senior practitioners narrows. The long-term consequence is a function that struggles to develop the next generation of strategic leaders.

Likelihood: Likely

The trend is already visible. Hiring at junior levels is declining

anecdotally, and the work that used to develop early-career practitioners is increasingly being automated. Without deliberate intervention, this gap will widen.

Overall Risk Level: High

Moderate consequence and likely likelihood make this high. It's a slow-moving risk that's easy to deprioritize, but the compounding effect over five to ten years could leave organizations without the experienced practitioners they need to operate at a strategic level.

Risk 11: Overreliance on AI Leading to Cognitive Skill Erosion

Consequence: Moderate

Both Ashish Gupta and Omar Abdullah raised the concept of “cognitive debt” — the risk that practitioners who outsource their thinking to AI tools gradually lose the judgment, analytical capability, and domain expertise that makes them valuable. If people stop doing their own analysis before prompting, stop challenging AI outputs, and stop practicing the hard work of thinking through complex problems, the quality of human judgment in the loop degrades over time. This undermines the very human edge that justifies keeping people in the process.

Likelihood: Possible

The behavior is already observable — Ashish described it as “AI fasting” and noted he’d written about whether people are becoming “AI obese.” It’s not universal, but the conditions that drive it are present in most organizations deploying AI tools without accompanying cognitive discipline frameworks.

Overall Risk Level: Moderate

Moderate consequence and possible likelihood land this at moderate. It’s worth monitoring and building guardrails around, but it’s not yet at the level of an immediate operational threat. The risk grows over time if left unaddressed.

Risk 12: Inaccurate AI Outputs Leading to Compliance or Operational Errors

Consequence: Moderate

The transcript references AI agent accuracy sitting around 90% for automated workflow creation, and multiple speakers acknowledged that AI systems hallucinate and make errors. In procurement contexts — contract review, invoice matching, supplier onboarding, spend analysis — a 10% error rate can translate into real compliance failures, financial discrepancies, and audit findings. The Cruise autonomous vehicle example was used to illustrate the tension: even when AI performs better than humans on average, a single high-profile failure can shut down an entire program.

Likelihood: Likely

AI systems making errors is not a hypothetical — it’s a documented and expected behavior of current-generation models. The question is not whether errors will occur but whether the governance structures are in place to catch them before they cause harm.

Overall Risk Level: High

Moderate consequence and likely likelihood combine to make this high. The risk is particularly acute in regulated environments where errors have legal and financial consequences, and where auditors may not yet have clear frameworks for evaluating AI-assisted processes.

Risk Matrix

	Insignificant	Minor	Moderate	Major	Extreme
Almost Certain			Risk 3: Unsustainable AI Pricing Models and Unpredictable Token Costs Risk 9: Energy and Power Constraints Limiting AI Infrastructure Growth	Risk 1: AI Data Privacy and Unauthorized Data Retention by Model Providers	
Likely			Risk 7: AI Tool Adoption Without Clear Problem Definition Risk 10: Junior Talent Development Gap Caused by AI Automation Risk 12: Inaccurate AI Outputs Leading to Compliance or Operational Errors	Risk 2: Shadow AI and Uncontrolled Employee Use of External AI Tools Risk 4: Data Breach Risk from AI Model Provider Infrastructure Risk 5: Procurement Workforce Readiness Gap in an AI-Driven Environment	
Possible			Risk 11: Overreliance on AI Leading to Cognitive Skill Erosion	Risk 6: AI Governance and Compliance Failures in Regulated Industries Risk 8: Dependency on a Small Number of AI Model Providers	
Unlikely					
Rare					

At a Glance: Event Summary in Q&A

Q: How is the shape of the modern procurement team changing and what new roles are we seeing emerge?

A:

- The team is leaning more toward strategic and data transformation capabilities rather than purely operational ones
- Key emerging focus areas include sustainability strategy, supplier relationship management, and AI prompt engineering
- AI is increasingly handling supplier onboarding, contract management, and relationship tracking at scale — particularly important for organizations managing thousands of suppliers
- The shift is toward humans defining what AI should do, managing outputs, and focusing on strategic decisions rather than executing routine tasks

Q: Is there a skill gap as tactical roles get automated, and can existing team members adapt?

A:

- Whether someone adapts depends largely on their personal appetite for learning and growth
- Behavior is considered more important than existing skills when hiring — someone with the right mindset will go the extra mile to learn
- Example given: an AP processor spending 30 hours a week on manual invoice matching may only need 2 hours once AI handles the bulk of it — the remaining 28 hours should be redirected to higher-value work
- Those freed-up hours can be used for data analytics, process improvement, or pushing error correction back to suppliers before invoices are submitted
- The key is helping people evolve their skills to stay valuable in a changing environment

“

“75% of enterprise data breaches are now occurring through a model provider from the enterprise.”

Q: What are some examples of AI delivering real value in procurement?

A:

- AI can be split into two buckets: tools that make existing tasks faster, and tools that enable things that weren't previously possible at all
- Adverse media searches on new vendors are a good example of the second bucket — previously too time-consuming for most companies, now executable via AI agents
- Intake validation engines can cross-reference answers on intake forms against actual contracts to flag inconsistencies (e.g., a user says they're not sharing data, but the contract says otherwise) — this cleans front-end data and drives more accurate downstream workflows
- On the invoicing side, AI can handle non-PO-backed invoices by referencing historical data, contracts, and rate tables — removing the need for AP teams to manually locate contracts
- The most effective AI deployments are embedded directly within the workflow rather than requiring users to manually move data between tools
- Exception-based processing is a major benefit — instead of reviewing 50 transactions, a team member might only need to review 3, leading to more focused and effective oversight
- AI can also flag contract scope changes with vendors, such as newly added AI capabilities in an MSA, and automatically route those to the appropriate review team

Q: What are the essential first steps to think about before embarking on an AI journey in procurement?

A:

- Start with the problem, not the tool — going straight to AI without defining the problem is like going to a pharmacy and asking for medicine without seeing a doctor
- The most expensive sentence in procurement is “let's use AI.” The most powerful is “let's first understand what really matters”
- Three mirrors framework was shared:
 - o First mirror: Look at your problem — identify where time and money are being lost, whether that's contract review, tail spend, or something else
 - o Second mirror: Look at your people — when automation is announced, employees often fear becoming redundant; build confidence before competence, and have human conversations before technical workshops
 - o Third mirror: Look at your process — if an approval process takes 20 steps, AI will just execute those 20 steps faster; redesign the process first, then apply AI
- People need to feel involved and reassured that they still matter before a pilot even launches

Q: How do you decide where to start on the AI journey, and how do you keep the focus on strategy and outcomes rather than tools?

A:

- Strategy should start with the problem, not the tool
- Don't start AI where it looks most impressive — start where it is most frustrating, then work backwards to the solution
- A time study was conducted asking the team where they were spending their time; the insights revealed the gap between where time was being spent and where it should be spent
- From there, the question becomes: do you address that gap through technology, upskilling, process redesign, or a combination?
- The best AI is the kind that saves people time and allows that time to be reinvested in higher-value work

Q: What does a digital mindset actually mean, and why does it matter?

A:

- A digital mindset has nothing to do with the tools themselves — it's about whether you're reacting or proactively adapting when technology changes
- The Nokia vs. iPhone example was used: Nokia was technically correct that the first iPhone had poor battery life and no keyboard, but they missed what actually mattered — the user experience. Nokia protected the keyboard; Apple protected the experience
- A digital mindset means being able to unlearn old things just as fast as you learn new ones — especially important in AI, where what's relevant today may not exist in three months
- It means shifting from a perfectionist mindset to an experimentation mindset — don't expect magic on day one
- It also means owning outcomes, not tasks — focusing on contract compliance or margin protection rather than “my suppliers, my category, my contract”

“

“This is making sure that we don't lose ourselves in the machines and let the machines do our thinking for us.”

“

“60 to 80% of practitioners are not ready for this change.”

Q: For organizations under pressure to grow output without growing headcount, what does an agentic operating model actually look like?

A:

- One organization shared they are currently producing 40,000 to \$50,000 USD a year with a goal of reaching \$500,000 USD in five years — without scaling spend proportionally
- Leadership is optimistic about using AI to achieve this, but the practical question of how to get there remains the central challenge
- A consistent theme across the market is “more with less” — not necessarily fewer people, but no additional headcount budget, so existing teams need to get more done
- Agentic AI is one lever; better process and workflow design is another
- The path forward is seen as a hybrid model: a backbone platform that provides structure, with the ability to build and customize specific agents on top of it

Q: How much should organizations focus on getting data clean before starting their AI journey?

A:

- Waiting for perfectly clean data is described as “the most sophisticated procrastination strategy”
- Clean data is rarely, if ever, truly achieved — those who claim to have it are often selling data cleansing solutions
- Not all data matters equally; some drives decisions, some just decorates dashboards
- The analogy used: a baby doesn’t wait for a clean floor to learn to walk — they start, fall, and try again
- The right approach is to start small and let AI surface where the data problems actually are — you clean data when you use it, not before
- That said, another perspective shared was that if data is not available, cleansed, and normalized, pointing an AI at it compounds the problem rather than solving it — so getting data actionable before deploying automation is also a valid position

Q: How is the skill set required of procurement practitioners changing as AI becomes more embedded?

A:

- Prompting is becoming a core skill — and prompting is not about getting answers, it's about asking the right questions
- The difference between a vague prompt and a precise one (e.g., “I want apples” vs. “I want green and red organic apples under \$5 available by tomorrow”) reflects the kind of specificity now required
- AI is becoming more human-like in some ways, which raises the question of what makes human practitioners uniquely valuable
- The answer lies in timeless skills: building relationships, influencing stakeholders, and solving complex problems
- There is an element of identity shift involved — what someone was known for may now be done by AI in minutes, which requires a rethinking of one's professional value
- This shift is happening at both the leadership level and across the broader team

Q: What are the key lessons learned from going through an AI journey, and what is one takeaway people should leave with?

A:

- A top-down approach did not work well — getting people involved from the very first step was far more successful
- The ABCDG framework was shared as a foundation to return to when in doubt:
 - o A: Alignment with the sponsor
 - o B: Belongingness — involve people from day one
 - o C: Communication
 - o D: Data
 - o E: Execution
 - o F: Fail fast
 - o G: (added as a bonus step)
- Celebrate small wins — they are bigger than they appear
- Celebrate successful failures too, because in AI, the learning itself is a success
- The one key takeaway: be an AI orchestrator — like a great music producer who doesn't play an instrument but feels the music, listens before anyone else, and brings everything together; procurement practitioners don't need to be technical experts in every tool, but they do need to orchestrate models, identify problems, and create opportunities

Q: How do you define success when it comes to AI in procurement?

A:

- Success is not just that a tool is working — that's the baseline minimum
- Real success is whether people are actually using it as a habit, not just during a project phase
- The most important measure is whether it has changed the way people work in terms of speed, quality, and confidence
- Time saved is not the metric — time reinvested is everything; if contract administration time is freed up, is that time actually going toward negotiations or supplier relationship building?
- Decision quality matters: are people making better decisions with the insights AI provides?
- Supplier resilience scoring was cited as an example — not just flagging risk, but understanding how well a supplier recovers from disruption
- Confidence is also a success metric: people move from being data aggregators to strategists, with the right information at the right time
- “Successful failure” is also a valid outcome — learning from a failed pilot is itself a success, because failure is a companion of innovation

Q: How are organizations thinking about make vs. buy decision for AI agents in procurement?

A:

- Some organizations are relying heavily on platforms and agents, including using tools for RFQ work — described as good for gaining early traction
- The challenge is sustainability over time — building your own agents may work initially but maintaining them long-term is an open question
- A hybrid model is seen as the most practical approach: a foundational platform that handles the core, with customization layered on top
- The counterpoint raised: vendors like established procurement platforms bring embedded knowledge from hundreds of other procurement organizations, which a self-built Claude or Codex deployment cannot replicate
- For tech-forward organizations, building your own is a legitimate option, particularly in early phases — but for many, the support, institutional knowledge, and single point of accountability that comes with a vendor relationship adds real value

Q: Who should own AI orchestration within a company — IT, procurement, or someone else?

A:

- This is an open and unsettled question — there is no clear consensus yet
- The case for IT: they understand data, security, and the technical infrastructure that most business users do not
- The case against IT owning it entirely: agentic systems require less deep integration than traditional enterprise software, so IT involvement is decreasing; faster implementation timelines are shifting ownership toward operational business units
- Transformation teams are often temporary and therefore not the right long-term owner
- The case for procurement: practitioners already touch legal, IT, risk, and governance — the same areas AI agents will interact with — making it a natural fit
- As AI systems increasingly provide business guidance (not just technical outputs), the risk and responsibility extends beyond IT alone
- The concept of a “Center of Intelligence” was raised as a model that could bridge the gap between procurement needs and IT/vendor capabilities
- The honest answer: it remains an open question, and organizations are still figuring out who will stand up and take accountability for what agents are doing

Q: How are companies approaching AI governance — who gets access, at what level, and how is usage managed?

A:

- Some organizations are directing employees to specific approved platforms (e.g., Gemini or similar tools) rather than allowing a free-for-all
- Some tech companies are actively recruiting with “unlimited AI token burn” as a benefit, signaling how competitive the space has become
- Others are issuing a set budget of tokens per week or month; others are leaving employees entirely on their own
- Forecasting AI consumption and cost is proving very difficult — many organizations are still wrestling with this in the early months of the year
- Governance sessions are being developed to ensure people are trained on when to use AI, how to use it, and how much
- A risk flagged: once employees are trained on prompting, you can end up with hundreds of overlapping scripts doing the same thing across departments with no coordination
- The overall picture is that governance approaches vary widely and the space is still maturing

Q: Are organizations feeling workforce tension around AI, and how are companies communicating with their teams about it?

A:

- Experiences vary significantly depending on the company and industry
- At a fast-growing company example, procurement employees are being pushed hard and are actually asking for AI help so they can focus on strategic tasks — the sentiment there is more positive
- In manufacturing and more traditional industries, the dynamic is different — there is real pressure coming from board level, driven by what boards are reading and hearing in the press, even if those boards don't fully understand what AI is
- A survey was referenced noting that 60% of companies plan to let go of employees who are not using AI — though the counterpoint was raised that eventually companies may also reduce headcount because AI has made certain roles unnecessary
- The key message for teams: build confidence before competence, involve people early, and frame AI as giving back time rather than taking away roles
- The honest acknowledgment: nobody fully knows how this will shape up





KONNECTHOUSE

Powered by

contented