

Agentic Procurement: Autonomous Sourcing Bots for GCC Public & Enterprise Buying

Purpose: This paper examines the emergence of agentic procurement – the use of autonomous AI “sourcing bots” – in both public sector and enterprise purchasing in the Gulf Cooperation Council (GCC) region by 2025.

64 min read • 13 sections • LISRC R&D practice



LISRC Research & Development

The applied-research practice of London International Studies & Research Centre

01 Abstract

Purpose: This paper examines the emergence of agentic procurement – the use of autonomous AI “sourcing bots” – in both public sector and enterprise purchasing in the Gulf Cooperation Council (GCC) region by 2025. We explore how procurement functions can leverage AI agents to execute end-to-end procurement tasks (from supplier discovery and negotiation to contract management) while aligning with principles of transparency, fairness, and strategic value. We map these developments to global AI governance frameworks (such as ISO/IEC 42001 and NIST’s AI Risk Management Framework) and procurement best practices to propose a blueprint for an AI-enabled procurement function. Approach: The study reviews recent surveys and industry reports on AI adoption in procurement, highlights regional initiatives in GCC e-procurement transformation, and delineates an architecture for autonomous sourcing cycles. Key components include multi-agent orchestration of procurement stages and human-in-the-

loop checkpoints to ensure compliance and ethical outcomes. Technical considerations like data quality, algorithmic bias, and integration with existing procurement systems are analyzed. Findings: Properly implemented, AI agents can dramatically improve procurement efficiency – potentially halving cycle times and achieving cost savings in the order of 5-10% through optimized decisions[1][2]. A 2025 survey indicates 90% of procurement leaders are considering or using AI agents, reflecting high expectations for value creation[3]. Early deployments show autonomous procurement cycles can handle routine sourcing end-to-end, freeing human professionals for strategic supplier management. However, strict safeguards are required. Our analysis of governance frameworks and case examples shows that accountability, explainability, and human oversight are indispensable to prevent “black box” decisions that could undermine fairness or violate public procurement laws[4][5]. We present a pseudocode-driven scenario of an autonomous sourcing event to illustrate how human approval triggers and ethical controls must be embedded. Implications: For GCC governments and enterprises, agentic procurement offers a path to greater agility and transparency in buying, aligning with national digital transformation goals. By following AI governance principles and procurement regulations, organizations can harness AI for faster, data-driven purchasing while upholding integrity and public trust. The paper provides practitioners with a practical framework to safely integrate AI agents into procurement and positions internal procurement teams to become innovation leaders in an AI-driven economy.

02 Introduction

Artificial Intelligence (AI) is reshaping business functions worldwide, and procurement is at the forefront of this transformation. By 2025, AI is not only assisting with spend analytics and forecasting, but also taking on autonomous decision-making roles in sourcing and supplier management. The GCC region, known for its ambitious digital initiatives, is embracing these advancements as part of broader economic modernization. Over 62% of GCC companies reported using AI in at least one business function by

2023[6], and procurement is a prime candidate for AI augmentation given its data-intensive and process-driven nature. A recent ProcureCon survey found that 90% of procurement leaders have considered or are already deploying AI agents to optimize operations in 2025[3]. This overwhelming interest signals that what was once a futuristic concept – AI bots negotiating with suppliers or automatically executing purchases – is now seen as an imminent reality.

Procurement, both in the public and private sectors, stands to gain significantly from AI. In enterprise procurement, AI can help analyze spend patterns, recommend sourcing strategies, and even engage in negotiations, leading to cost savings and efficiency gains. In public procurement, AI promises improved transparency and compliance by reducing human biases and errors in tender evaluations. The GCC governments have made procurement reform a priority in recent years: for instance, the United Arab Emirates launched a unified Digital Procurement Platform to standardize and digitize federal purchasing across 50+ entities[7][8]. Saudi Arabia's Etimad platform has similarly moved government tendering online, bringing greater visibility and consistency. Layering AI on top of these digital procurement systems is the next logical step. It aligns with national AI strategies – the UAE's National Strategy for AI 2031 explicitly calls for leveraging AI with strong governance to improve government services[9]. Likewise, Saudi Arabia's SDAIA (Saudi Data & AI Authority) has been developing guidelines to ensure responsible AI deployment across sectors, including procurement, emphasizing ethics, transparency, and risk management.

As procurement functions consider adopting "agentic AI" (i.e., AI agents with autonomy in decision-making), they face a dual challenge: seizing the opportunity for greater speed and insight, while maintaining the rigorous control and fairness that procurement demands. Traditionally, procurement is governed by strict policies and regulations – especially in the public sector where fairness, open competition, and accountability are paramount. The introduction of AI does not remove these requirements; if anything, it heightens the need for diligence to ensure algorithms do not inadvertently sidestep rules or embed biases. For example, an AI sourcing bot might analyze supplier data and favor certain vendors; unless carefully designed, this could conflict with public procurement rules that mandate giving all qualified suppliers a fair chance. Similarly, in a corporate setting, an AI agent might be tempted to select a supplier purely on cost, but a human

would consider strategic relationships and ESG (environmental, social, governance) factors – how do we ensure the AI does the same?

The need for a framework to guide AI adoption in procurement is evident. Without guidelines, organizations risk deploying AI piecemeal – for instance, using a chatbot here or a pricing algorithm there – without an overarching strategy or understanding of the risks. A fragmented approach could lead to compliance violations (e.g., an AI inadvertently colluding with suppliers or breaching data privacy) or failures to realize AI's full potential due to lack of integration. By establishing a structured approach, procurement leaders can ensure that all AI tools are implemented with proper oversight, aligned with best practices, and truly add value. This paper aims to provide such a roadmap, focusing on the GCC context but drawing lessons from global best practices. We will review relevant standards and principles that any AI in procurement should adhere to, then propose how an "AI-enabled procurement function" might operate – from automatically detecting purchasing needs to conducting autonomous sourcing cycles – all under a governance umbrella that ensures ethical and effective outcomes.

The rest of the paper is organized as follows: Section 2 provides background on the rise of autonomous procurement (often dubbed "Procurement 4.0") and why it's gaining traction, including regional drivers in the GCC. Section 3 discusses key principles and frameworks (technological and regulatory) that set guardrails for AI use in procurement – including AI governance standards like ISO 42001 and procurement-specific guidelines around fairness and accountability. Section 4 outlines the architecture of an agentic procurement system, describing how AI agents can integrate into each phase of procurement (planning, tendering, contracting, and supplier management), with a figure illustrating an autonomous sourcing cycle. Section 5 delves into governance and risk management for procurement AI: how to oversee AI decisions, ensure transparency, and mitigate risks like bias or system failures. We include a pseudocode example to concretely demonstrate how an autonomous sourcing bot could be audited and controlled. Section 6 evaluates the expected performance impact of AI in procurement – from efficiency gains (cycle time reduction, lower costs) to potential pitfalls (initial false starts, change management issues) – drawing on early case studies and surveys. Finally, Section 7 concludes with implications for procurement professionals and policymakers,

particularly in GCC countries, and suggests areas for future research such as training needs and the evolution of procurement roles in an AI-driven future.

03 **Background and Drivers for Autonomous Procurement**

The Evolution of Procurement and the Role of AI: Over the past two decades, procurement has evolved from a transactional, back-office function into a strategic player in organizational value creation. Concepts like Category Management, Strategic Sourcing, and Supplier Relationship Management have elevated procurement's impact on cost optimization and risk mitigation. In parallel, technology enablers have grown – from early e-procurement systems and spend analytics to today's AI and cloud-based platforms. Procurement generates and consumes massive amounts of data (supplier data, pricing, contracts, market indexes), making it fertile ground for AI applications. Initially, AI in procurement took the form of predictive analytics (e.g. predicting supplier delivery delays) and simple automation like robotic process automation (RPA) for invoice processing. These applications have delivered benefits, but they largely assist humans rather than act autonomously. The emergence of agentic AI marks a shift to greater autonomy: systems that can not only analyze or provide recommendations, but also make decisions and execute actions with minimal human input[10][11]. In procurement, this means AI that can potentially run an entire sourcing event from start to finish, or automatically manage routine buying within set guardrails. The notion of an “autonomous procurement agent” is no longer science fiction – prototypes and pilot projects are underway, and vendors are racing to embed these capabilities into their solutions.

Why Now? The Case for AI in Procurement: Several converging factors explain why autonomous procurement is gaining momentum by 2025. First, the technology is ready. Advances in natural language processing and machine learning enable AI to handle complex tasks like reading contracts or negotiating terms, which were previously beyond automation. Modern Large Language Models (LLMs) can draft RFPs or interpret vendors' proposals, while reinforcement learning agents can adapt strategies through experience. Second, the COVID-19 pandemic and subsequent supply chain disruptions highlighted the need for greater agility and real-time decision-making in procurement. Many organizations found their traditional, manual procurement processes too slow or rigid to

respond to sudden changes (like urgent sourcing of medical supplies or dealing with logistics bottlenecks). AI agents, which can monitor situations 24/7 and rapidly execute contingency plans, are seen as a solution for building resilience. Third, there is a push to “do more with less.” Procurement teams often face resource constraints and increasing workload complexity (e.g., monitoring ESG compliance in supply chains). AI augmentation offers a way to extend capacity – a survey by Deloitte noted digitally mature organizations achieve 2.8 times higher return on AI investments than less mature peers[12], suggesting that those who embrace AI in procurement can significantly outperform.

GCC Regional Drivers: The GCC region (comprising Bahrain, Kuwait, Oman, Qatar, Saudi Arabia, and the UAE) provides a unique environment where ambitious modernization goals meet relatively green-field opportunities for tech adoption. GCC governments are driving digital transformation as part of economic diversification plans (e.g., Saudi Vision 2030, UAE Centennial 2071). Public procurement is a major lever for these governments – not only is it a huge portion of public spending (and thus a target for efficiency gains), but it’s also a tool to encourage local SME participation and ensure transparency. AI can help in both aspects: improving efficiency and safeguarding integrity. For example, AI could help detect patterns that indicate collusive bidding or fraud in public tenders, augmenting the fight against corruption. Indeed, ensuring fair and unbiased AI-driven processes is crucial in public procurement where equity and fairness are mandated[13]. The region’s business culture is also relatively top-down; when leadership mandates a new approach (such as “we will use AI to optimize procurement”), implementation can be swift compared to more decentralized environments. This is evident from how quickly e-procurement platforms were adopted at national scales in the GCC. Now, with national AI strategies in place, we see explicit support for integrating AI into government operations. For instance, the UAE’s AI Ethics Guidelines (2022) and Charter for AI (2024) advocate for transparency, accountability, and bias mitigation in AI systems[14][15] – principles that any AI use in government procurement must adhere to. Saudi Arabia, through SDAIA, introduced an AI Adoption Framework in 2024 to guide organizations in responsibly implementing AI[16]. These policies create an enabling environment and also a gentle pressure: procurement departments are expected to innovate with AI but must do so within an ethical and controlled framework.

Challenges and Concerns: Despite the excitement, the path to agentic procurement has challenges. One concern is the maturity of data and systems. AI's effectiveness is only as good as the data it learns from; if an organization's spend data is fragmented or supplier databases are incomplete, the AI agent might make poor decisions or require extensive training time. The Suplari/Deloitte survey highlights that data quality issues are a significant barrier, with 75% of CPOs citing data quality as a detractor from confidence in AI[17]. Integration with existing systems is another hurdle – 88% of procurement leaders in that study raised integration with ERPs and procurement suites as a major challenge in adopting AI[17]. Additionally, there is cultural resistance: procurement professionals may fear that AI could make their roles obsolete or may simply distrust an algorithm's decisions. In the short term, jobs will indeed shift – as one industry blog put it, autonomous procurement will likely eliminate some “mind-numbingly boring manual tasks” but not the need for human strategic input[18][19]. The profession must prepare for upskilling and redefining roles rather than job loss per se. The Jaggaer white paper on agentic procurement stresses that new roles like “AI enablement specialist” or “Procurement-AI liaison” are emerging to bridge the gap between tech and teams[20][21]. Another concern is accountability: if an AI makes a wrong decision (say, selects a supplier that fails catastrophically or inadvertently discriminates against a class of suppliers), who is responsible? Procurement by its nature has checks and balances (approvals, audit trails); these need to be preserved. Ensuring that AI doesn't become an unaccountable “black box” is paramount – leading companies are addressing this by forming AI governance committees that include procurement leaders, to oversee AI rules and outcomes[22][23].

The Need for a Framework: All these factors underscore why a structured framework for AI in procurement is needed. It should help organizations systematically identify where AI can add the most value in the procurement cycle, implement it in a controlled manner, and maintain oversight. The framework should incorporate global best practices – like aligning with the NIST AI Risk Management Framework's core functions (Map, Measure, Manage, Govern) to ensure all dimensions of AI risk are considered[24][25]. It should also reflect procurement-specific guidelines, perhaps drawing from initiatives like the World Bank's principles for transforming public procurement which emphasize transparency and accountability as non-negotiables[26][27]. By blending AI governance with procurement governance, the resulting approach can inspire confidence among

stakeholders (executives, auditors, suppliers, the public) that even as more of the process is automated, the values of fairness, integrity, and value-for-money remain intact. In the next section, we examine the standards and frameworks that will inform such a blueprint, acting as guardrails for the journey to autonomous sourcing.

04 Principles and Frameworks for Trustworthy AI in Procurement

To adopt AI-driven procurement in a safe and effective manner, organizations should anchor their approach in established principles and frameworks. These provide guidelines on how to manage technology in a way that upholds integrity and maximizes value. We consider three levels of guidance: (1) Procurement-specific best practices and regulations, (2) AI governance frameworks (global standards for AI management), and (3) Overarching ethical principles for AI such as fairness, accountability, and transparency, which overlap with procurement values.

Procurement Regulations and Best Practices: Procurement, especially in the public sector, is heavily regulated to ensure fair competition and proper use of funds. Key principles include transparency (clear criteria, open tendering), equality of opportunity for suppliers, and accountability for decisions. Any AI system embedded in procurement must be designed to reinforce, not weaken, these principles. For instance, if an AI is used to screen suppliers or bids, it should document the basis of its decisions (meeting transparency requirements) and be tested to ensure it doesn't systematically disadvantage any suppliers without justification (ensuring fairness). The U.A.E.'s procurement laws and the KSA Government Tenders and Procurement Law both mandate competitive procedures and have provisions to guard against conflicts of interest; using AI doesn't exempt compliance – instead, AI should help by objectively applying the criteria. Private sector procurement is less regulated by law but guided by best practices and corporate policies (often aligned to standards by bodies like the Chartered Institute of Procurement & Supply, CIPS). These stress ethical sourcing, value for money, and auditability. Incorporating AI into procurement should thus follow a risk-based approach: for example, using AI for high-value or sensitive procurements might necessitate additional human review steps compared to low-value routine purchases. The NIGP (a public procurement institute) notes that AI can enhance procurement if it is implemented

with controls that ensure decisions are unbiased and traceable[13]. Additionally, global initiatives such as the OECD's recommendations on public procurement (which emphasize transparency, competition, and use of modern tools) implicitly support AI, as long as it furthers these goals. In summary, the first "framework" for AI in procurement is the existing procurement rulebook – AI must map to it. A practical approach is to update internal procurement policies to explicitly cover AI usage, e.g. "AI may be used to shortlist bids, but final award decisions must be validated by the procurement committee," or "All AI recommendations must be accompanied by a rationale that can be disclosed if needed."

ISO/IEC 42001 (AI Management System Standard): This international standard, published in late 2023, provides requirements for implementing an AI Management System (similar in spirit to ISO 9001 for quality or ISO 27001 for information security, but focused on AI). It's relevant to procurement in two ways: organizations can use it to govern any AI they deploy (including procurement bots), and it can guide how an AI-enabled procurement function is structured. ISO 42001 emphasizes a Plan-Do-Check-Act cycle for AI and outlines controls around data management, risk assessment, human oversight, and continuous improvement[28][29]. If a procurement department adopts ISO 42001, it would, for instance, formally plan the objectives of using AI in procurement (Plan), deploy the AI with proper training and support (Do), monitor outcomes such as error rates or bias in supplier selection (Check), and refine processes or models based on monitoring (Act). Notably, ISO 42001 calls for clear role definitions and accountability – aligning with procurement's need to know "who is responsible" for decisions even when AI is involved. It also highlights stakeholder engagement: in a procurement context, that means engaging users (procurement staff), approvers, and possibly suppliers about the AI's role. An example might be informing suppliers that an AI will be used to evaluate their bids and assuring them of fairness criteria, which could be part of change management and building trust in the new system. ISO 42001 further underscores risk management; procurement AI risks include selecting an unreliable supplier due to erroneous algorithm judgment, or a data breach if supplier data in the AI system isn't well protected. Aligning to this standard means those risks are identified and mitigated systematically, perhaps through scenario testing and controls like data encryption and access controls (to protect supplier confidential data). Given GCC organizations often seek international certifications to demonstrate excellence (e.g., ISO standards for quality or security are

common in government entities), adopting ISO 42001 for AI in procurement could both improve practice and signal commitment to trustworthy AI use.

NIST AI Risk Management Framework (AI RMF 1.0): The NIST AI RMF, released in 2023, is a comprehensive but voluntary framework that provides a structured approach to managing AI risks[30][31]. It is built around four functions: Map, Measure, Manage, and Govern. Applying this to an autonomous procurement system is illuminating. In the Map stage, the organization contextualizes the AI system – identifying the intended use (e.g., automating supplier selection for low-value purchases), the scope and objectives, stakeholders (procurement team, business users, suppliers, regulators), and potential impacts. For example, mapping might reveal that an AI making sourcing decisions could impact supplier livelihoods (if SMEs are consistently not chosen, it could have socio-economic ripple effects), so that risk must be noted. Under Measure, the organization would develop metrics or tests for the AI's performance and risks. This could include measuring accuracy (how often does the AI's supplier recommendation align with what a human expert would choose), bias (do certain supplier categories get systematically lower scores), robustness (does the model behave oddly with slight changes in input), and security (can the system be tampered with, say by a supplier feeding misleading data). Tools might be employed here, such as fairness metrics or confusion matrices on historical data. The Manage function is about risk response – implementing controls like thresholds for human intervention (e.g., if top two bids are within 5% score of each other, escalate to human rather than auto-award), establishing fallback plans (if the AI system fails, revert to manual process), and training staff on new procedures. It also includes ongoing monitoring for incidents or drift (the AI's performance degrading over time). Finally, Govern entails the overarching governance processes: policies, roles (perhaps an AI governance board as mentioned earlier, or an AI risk champion in the procurement team), compliance checks, and communication strategies. NIST's framework is valuable in procurement because it forces a thorough consideration: not just "does the AI work?" but "what could go wrong, and how do we know and manage that?" For instance, a NIST-aligned approach might prompt the creation of an inventory of all AI tools used in procurement and an internal audit schedule to review each for compliance with policies. It's worth noting that the IIA (Institute of Internal Auditors) in its AI Auditing Framework encourages auditors to use frameworks like NIST's when assessing AI[32] – so if procurement follows NIST RMF, it will be better prepared for any internal or external

audits of its AI. GCC companies, especially in regulated sectors like finance or government suppliers, will find such alignment beneficial if regulators issue guidelines (as they likely will, building on the existing AI ethics principles in the region). In summary, NIST AI RMF provides a playbook to ensure no stone is left unturned in deploying an AI sourcing bot: from understanding context to quantifying risks and having structures in place to keep the AI on the right track.

Ethical AI Principles – Fairness, Transparency, Accountability: Cutting across the above frameworks are core ethical principles that resonate strongly with procurement's mission. Fairness in AI means avoiding bias and ensuring equitable treatment; in procurement, this translates to AI that does not unfairly disadvantage any supplier or group. For example, if an AI learns from historical data where perhaps larger vendors always won, it might carry that bias forward – fairness dictates we detect and correct that, perhaps by rebalancing training data or explicitly programming the agent to consider diversity in supplier base (some organizations set targets for SME inclusion, etc.). Transparency is critical because procurement decisions often need to be explained to unsuccessful bidders or auditors. If an AI is involved in bid evaluation, the ability to explain how it scored proposals is essential. Techniques for explainable AI (like SHAP values, decision trees, or rule-based surrogates) could be employed so that for any contract award decision, the procurement team can provide a clear rationale, AI-derived or not. Indeed, the OECD's principles for trustworthy AI – endorsed by 40+ countries including many GCC partners – list transparency and explainability as key requirements[32]. Accountability means that humans cannot shrug off responsibility by saying "the AI made me do it." There must be clarity on who is accountable for AI-augmented decisions. Leading practice is to require human sign-off on important decisions or to establish an AI oversight role. In one example from the industry, a company mandated that any time an AI vendor risk scoring tool rejected a potential supplier, a procurement manager had to review the case to confirm the decision – ensuring a human remained answerable. The Art of Procurement's governance guide notes that fewer than one-third of large enterprises currently allow unrestricted AI use, precisely because of data and compliance concerns[33], reinforcing that most organizations see the need for controls. Many are forming AI ethics committees as mentioned, and procurement should be part of those. A tangible outcome might be an internal policy that says, for instance, "Procurement AI will be subject to quarterly ethical review by the AI governance board, examining decision logs for any signs of bias or

error.” In the public sector context, accountability also means maintaining the audit trail: an AI should log its actions and inputs so that if a legislative auditor or anti-corruption body inspects a tender, they can trace what was done and why. The UAE’s AI Ethics Guidelines and the 2024 Charter explicitly stress governance and algorithmic bias challenges[34], which procurement implementations must heed by design.

In summary, by grounding the adoption of autonomous sourcing bots in procurement within these frameworks and principles, organizations create a safety net. The technology can progress, but within bounds that protect the organization’s interests and stakeholder trust. The 2025 landscape is such that technology solutions are available off-the-shelf to do agentic procurement (major suite providers like Coupa, SAP Ariba, JAGGAER, Zycus, etc., are touting AI additions), yet technology is only half the battle. The other half is process and governance. As we move to the next section on methodology, these principles will be woven into how we envision the AI-enabled procurement function operating in practice.

05 **AI-Enabled Procurement Architecture and Methodology**

An AI-enabled (or “agentic”) procurement function integrates AI tools at every stage of the procurement lifecycle. We break down the procurement process into phases and illustrate how autonomous agents and AI analytics can augment or automate each phase, while a central orchestration and governance mechanism ensures everything runs smoothly and in compliance. Figure 1 depicts an overview of this architecture, showing multiple specialized AI agents working under a coordinating layer, with human oversight loops at critical points. We then delve into each phase: need identification and planning, sourcing (tendering and supplier selection), negotiation and contracting, and post-award contract and supplier management. Throughout, we highlight where and how human procurement professionals interact with the AI – the model here is very much one of collaboration, not a total “hands-off” black box.

Figure 1: Autonomous Procurement Framework. The framework envisions a central Orchestration Agent that oversees the end-to-end procurement cycle. This orchestration layer triggers and coordinates a set of specialized AI agents, each responsible for

different stages: a “Needs Analyzer” agent monitors internal systems (ERP, inventory, spend data) to detect when a purchase need arises or a contract is expiring; a “Sourcing Strategy” agent designs the appropriate approach (deciding between a spot buy, competitive RFQ, auction, etc., based on risk and value); a “Bid Agent” handles RFx creation and Q&A with suppliers; an “Evaluation” agent scores proposals and applies decision logic; a “Negotiation” agent conducts automated negotiations within preset limits; a “Contract Drafter” agent prepares contract documents; and a “Performance Monitor” agent tracks execution, deliveries, and supplier compliance. These agents communicate with each other – for example, the Evaluation agent can request the Negotiation agent to try for a discount if two suppliers are very close in score. Human interfaces exist at two levels: a dashboard for procurement officials to oversee ongoing processes (showing status, flags, and allowing intervention), and defined intervention points where the process halts for human approval (e.g., before final award if the deal value is above a threshold or if the AI’s confidence is low). Surrounding all this is an AI governance layer: logging every decision, validating that agents are following rules, and monitoring outcomes to feed back into continuous improvement (aligning with ISO 42001’s PDCA cycle). Essentially, the AI-driven system acts as an autopilot for procurement, but humans remain the pilots in command – they can take over control when needed and are always aware of what the system is doing.

06 From Need Identification to Sourcing Strategy

The procurement process starts with recognizing a need to buy something – which could be triggered by a user request, a stock level falling below threshold, or a contract coming up for renewal. In a traditional setup, a department manager might raise a purchase requisition or a procurement planner reviews upcoming contract expirations. In an AI-enabled setup, a Needs Analyzer agent can automate much of this detection. For example, by integrating with an inventory management system, the agent can spot that a certain critical spare part is running low and automatically initiate a replenishment request[35]. Jaggaer describes scenarios where an AI monitors ERP and signals a sourcing cycle when it notices stock is below safety levels or contract pricing is out of line with market benchmarks[36][37]. In a public sector context, the agent might scan annual budgets and consumption rates of supplies to forecast when a new tender should

be launched to avoid shortages in hospitals or schools. This predictive initiation ensures procurement is proactive.

Once a need is identified, the next step is figuring out how to source it. Here, a Sourcing Strategy agent comes into play. It determines the most suitable method: Should it be a simple requisition against an existing catalog or contract? A competitive RFQ because the spend is moderate and several suppliers are available? An e-auction for a commodity buy to drive down price? Or maybe a more complex RFP for a high-value, strategic service? The AI can make this decision by evaluating parameters like the category of spend, estimated value, criticality, and the risk factors involved. For instance, the agent might reference procurement guidelines: if spend is above X and no existing contract, default to competitive bidding; if the category has a history of volatile prices, consider a dynamic bidding approach; if there's a sole source situation, flag for single-source approval. The AI can incorporate external data too – e.g., if tariffs on a category are increasing, it might recommend locking in longer-term contracts (a strategic decision a human would make, but AI can prompt). Early versions of such strategy recommendation engines exist: they look at supplier market data, concentration risks, ESG considerations, etc., to outline the best path[38][39]. An example: an AI could suggest a reverse auction for procuring laptops if it knows there are many vendors, but for a custom software project, suggest a two-stage tender with technical evaluation. This aligns with NIST's Map function: understanding context (market conditions, etc.) and mapping risk (e.g., high supplier concentration might mean the AI advises splitting the award among multiple vendors for resiliency[38]). Ultimately, the procurement officer can approve or tweak the strategy – the AI provides a data-backed starting point quickly.

At this planning stage, AI also helps in stakeholder communication. A generative AI tool might quickly draft a summary of the upcoming procurement for approval, including why the chosen strategy is suitable (useful for internal governance committees). It can reference policy ("as per our procurement manual section X, an RFQ will be used since the value exceeds AED Y") and incorporate any risk assessments. Essentially, the tedious part of writing procurement plans or getting approvals can be accelerated by AI drafting and then human fine-tuning. Some organizations are already using AI to draft procurement justifications and get them through faster. For instance, if a certain purchase needs a sole-source justification, an AI can pull in all the relevant reasoning (e.g., "Vendor

A is the proprietary manufacturer meeting specification, switching supplier would cause interoperability issues”) from past documents or given criteria, saving procurement officers a lot of writing time.

It’s important that during need identification and strategy phase, guardrails are set. The “autopilot” should know when not to proceed without human input. For example, if the AI thinks no competition is needed (maybe it assumes a preferred supplier), that decision must be scrutinized by a human to avoid unintended biases or rule violations. Many public procurement laws require justification for single sourcing – an AI can prepare that draft but a human must sign off. This is a concrete application of the accountability principle: AI proposes, human disposes (for strategic decisions). The planning stage is also where the “Map” from the NIST framework would identify stakeholders – say the end-user department should be consulted on their need – and the orchestration agent might notify a category manager or procurement lead that a cycle is kicking off. Properly setting up the parameters here (like timeline for the procurement, any budget limits, any pre-approved supplier list to use) ensures downstream AI agents operate with the correct constraints. With the stage set, the next part of the cycle is executing the sourcing event itself, where AI’s role becomes even more transformative.

07 **Sourcing Execution: RFx, Evaluation, and Award**

The sourcing phase is where agentic procurement truly shows its potential. It comprises creating a solicitation (RFQ/RFP or other tender), inviting suppliers, managing Q&A, receiving bids, evaluating them, and deciding awards. In a conventional process, these tasks involve substantial manual effort and time – drafting documents, responding to emails, running Excel comparisons, etc. AI agents can automate and speed up each of these steps:

RFx Creation and Communication: A Bid Agent using generative AI can prepare the solicitation documents automatically. It can take templates and fill in details of the scope, requirements, and criteria gleaned from the need description and any specifications provided by the user department. If the procurement is for a standard item or service, the AI can also incorporate best-practice language or past RFP content. For example, for an IT services RFP, the agent could include clauses around data security that it knows from

knowledge base must be present, and set up the pricing format. Jaggaer notes that generative AI is already being deployed to draft RFX and even contract clauses[40][41]. The Bid Agent then identifies potential suppliers. It can pull from the company's supplier database, ensuring that only pre-qualified or registered vendors are considered (especially important in government procurement, which often requires using registered suppliers). Additionally, the agent might search external sources or integrate with supplier networks to find new vendors if needed, including checking ESG ratings or other attributes relevant to the strategy (for example, if sustainability is a goal, ensure suppliers meet certain certifications). Once the supplier list is set, the agent issues the RFX via the e-procurement platform or email, and even handles the Q&A – suppliers' questions can be answered by the AI if they're similar to known FAQs, with the AI referring any novel or policy-related questions to a human procurement officer. This significantly reduces the administrative burden of managing tender communications. By keeping all communications within the system and logged, it also improves transparency and fairness (every supplier gets the same information), addressing a key public procurement requirement.

Bid Evaluation: When bids or proposals come in, an Evaluation Agent takes over. It can use a combination of rule-based scoring and machine learning. For structured parts (price, compliance to specs), it will automatically tabulate and score. For qualitative parts (like technical proposals or method statements), natural language processing can evaluate the text against the criteria. This might involve sentiment analysis, keyword matching, or even more advanced techniques like using an LLM to grade responses against an ideal answer key. This is admittedly cutting-edge – trusting an AI to read and score proposals is something being experimented with carefully. A more immediate application is using AI to assist human evaluators: for example, the AI can summarize each proposal, highlight where each proposal is strong or weak relative to requirements, and flag any sections that look like potential boilerplate or irrelevant filler. This saves the evaluation committee time so they can focus on key differentiators. The AI might also detect anomalies or risks – say a supplier bid significantly under cost (possibly a sign of misunderstanding or a strategy that might lead to change orders later), or a supplier's technical proposal is eerily similar to another (possible collusion or plagiarism). The Jaggaer vision includes the AI flagging compliance red flags and using predictive analytics to identify if scores are too close[42][43]. The agent can even simulate award

scenarios, considering different weighting or cost vs. quality trade-offs, to help decision makers see outcomes. Throughout, it's crucial that if the scores are close or if some threshold is triggered (like a new supplier that hasn't been used before is the winner), the system invokes a human-in-the-loop checkpoint . For example, a rule could be: if the top proposal's score is within 5% of the second, or if the contract value is above \$1 million, require human review before award. This ensures that an AI doesn't make a final call in ambiguous or high-stakes situations without human judgment. It's similar to an airplane's autopilot turning off when conditions are outside normal range, alerting the pilot to take over.

Decision and Award: Assuming the evaluation yields a clear best option and passes any necessary human review, the AI can proceed to "award" in the sense of selecting the winner and preparing the award recommendation. In public procurement, formal award decisions will still typically be made by a tender committee or authorized official; however, the AI can generate the entire decision report: outlining the process followed, presenting the scoring, and the rationale for the selected supplier. By auto-generating this report and even the regret letters to unsuccessful bidders (with appropriate feedback extracted from the evaluation), the administrative tail of the sourcing process is trimmed significantly. A key win here is consistency – AI ensures that every bidder gets a thorough, unbiased evaluation against the criteria, and that documentation is complete for audit purposes, reducing the risk of challenges. Indeed, one of the promises of AI in this area is a reduction in bid protests or disputes, because decisions can be more transparently based on data and preset criteria, leaving less room for suspicion of subjective bias. That said, the procurement team must still vet the outputs: a human should glance over the recommendation report and the draft letters before they go out, to ensure tone and content are appropriate (GenAI sometimes might choose phrasing that could be misinterpreted, so a final sanity check is wise).

One might ask: what about negotiations? In some procurement cycles, especially in private sector or for complex purchases, after initial evaluation there might be negotiations (commercial or technical clarifications) with one or more top vendors. This is increasingly another area where AI can contribute, which we cover next in the contracting phase, as negotiation often blends into contracting. But it's worth noting that negotiation is one area where human skill is still highly valued, especially for strategic deals. The

current state of AI is that simple negotiations (like haggling over price or delivery date in a structured way) can be done by agents – for example, setting up a chatbot to negotiate bulk discount terms with a supplier’s system if the supplier also has an AI interface. Some procurement software companies have begun to offer “negotiation bots” for tail-end spend (where the stakes are low and volume is high), achieving incremental savings where humans wouldn’t have bothered to negotiate. However, for large contracts or sensitive categories, human-led negotiations remain the norm. The approach, therefore, is hybrid: AI handles what it safely can, and kicks the rest to humans. A study by McKinsey observed that while AI can automate many procurement tasks, category managers will still lead negotiations in most strategic categories for the foreseeable future[44][45]. The AI’s role there is to provide analytics and suggestions (e.g., “Based on supplier’s past quotes, they have 10% room to come down” or automatically analyzing a supplier’s proposal to find areas to negotiate, such as overly conservative delivery dates or warranty terms).

By the end of the sourcing execution phase, the procurement process has a selected supplier and terms ready to be formalized. The next phase deals with finalizing contracts and ensuring the purchase is executed, which the following section will detail.

08 **Negotiation, Contracting, and Post-Award Management**

Automated Negotiation: As noted, not all procurements involve negotiation – many public tenders award on a fixed bid basis. But in cases where negotiation is applicable (for instance, a single-source procurement or after competitive dialogue), an AI Negotiation Agent can step in under certain conditions. These AI negotiators are designed to work within predefined boundaries set by the organization. For example, the procurement team might set the target price, walk-away price, and key terms that are negotiable (delivery schedule, service levels, etc.). The AI then conducts the negotiation, possibly via a chat interface or email, if the supplier is also amenable to an AI-driven exchange. Currently, fully autonomous negotiation is rare in high-stakes deals, but some platforms enable automated bargaining for simpler transactions. A real-world scenario is dynamic pricing negotiations in supply chain: an AI agent could automatically ask for a discount or better payment terms from a supplier’s system and get an instant response, iterating until they

reach a limit. One of the first likely uses in GCC enterprises might be in procurement of standardized commodities or for spot-buy situations where speed is important. An AI could negotiate with multiple suppliers simultaneously in an online environment (akin to a multi-agent auction, but with conversational negotiation rather than just bids). Trials have shown it's feasible – in fact, by 2025, we expect pilot deployments where non-critical purchases (say office supplies or minor software licenses) are negotiated by AI, yielding small savings at scale. The benefit is not just cost reduction; it frees procurement staff from a multitude of trivial negotiations so they can focus on strategic supplier discussions where human nuance and relationship matter. It's crucial, however, that the AI logs all communication and abides by ethical guidelines – e.g., it should not make false promises or threats, and should end negotiations politely if no agreement is reached within its mandate. Essentially, the AI's negotiation behavior must reflect the organization's values, as if a person were conducting it. This may require training the AI on historical negotiation transcripts or rules of engagement defined by legal and procurement policy (for example, no sharing of other suppliers' pricing, which would violate integrity).

Contract Drafting and Review: Once terms are finalized (whether via direct selection or negotiation), a Contract Drafter agent takes over to formalize the agreement. This agent leverages both generative AI and knowledge of legal templates. It can auto-populate a contract template with the specifics of the deal: parties' names, scope of work, price, delivery schedule, etc., which it pulls from the procurement data. More impressively, it can tailor clauses to reflect negotiation outcomes – for example, if the negotiation agent secured an extended warranty or special discount terms, those will be inserted into the relevant sections. Companies like Icertis and Coupa are embedding such AI to generate contracts and even to verify consistency between the contract and the RFP/Proposal to catch discrepancies. After drafting, the AI can do a first-pass compliance check: comparing the draft to company contracting guidelines (ensuring mandatory clauses are present and no unauthorized alterations are made). Any deviations can be highlighted. For instance, if the supplier proposed their own terms during negotiation and the AI included them, the system might flag: "Liability cap clause is non-standard – requires legal review." This allows the procurement lawyer or legal department to focus attention on just the unusual parts rather than reading the whole contract line by line. In the context of public procurement in GCC, standard government contract forms are often used, but an AI could be valuable in populating schedules, checking performance bond values, or

adapting contracts for special cases (like local content requirements in Saudi Arabia, or bilingual contracts in Arabic and English – AI translation can assist here, though would need reviewing). The key is that the AI streamlines the journey from selection to signed contract, but doesn't bypass necessary approvals. So, even if the contract is ready in draft, it should route to the appropriate manager or legal counsel for sign-off. Perhaps in the future, simpler contracts (like NDAs or small purchase orders under a threshold) might be fully automated with no human in the loop, but for now, it's prudent that medium to high complexity contracts get a human glance. This aligns with ISO 42001's human oversight requirement and simply good risk management. It's similar to how autopilot can land a plane but most airlines still prefer pilots to supervise or take control in rough weather – human expertise remains as a safety net.

Execution and Performance Monitoring: After contract signing, the procurement process transitions to order execution and ongoing supplier management. Here, AI agents provide the capability for continuous auditing and monitoring of the contract performance. The "Performance Monitor" agent will track purchase orders, deliveries, invoices, and compliance to terms. For example, if the contract says deliveries should be within 5 days, the AI will watch actual delivery logs and alert if any are late or if a pattern of delays emerges. It can cross-check invoice prices against contract prices to ensure the supplier isn't overcharging – essentially doing automated 3-way match and beyond, covering 100% of transactions. In public procurement, contract management is a known weak area (globally, not just GCC) where after the tender, the follow-up on whether the supplier delivers as promised can be lax. AI tools can significantly enhance oversight by raising flags in real time. Another aspect is compliance: the agent can monitor if the supplier is meeting KPIs or service levels agreed, feeding data into dashboards that the contract manager can review. If a supplier underperforms consistently or a risk arises (say their financial rating drops, or news of sanctions appear), the AI can advise the procurement team to intervene, possibly prompting a re-sourcing. In fact, as Jaggaer's vision suggests, an AI might even automatically initiate a re-sourcing event for a critical supply if it detects a major disruption like the supplier going out of business or a geopolitical event blocking supply[46][47]. Of course, in government scenarios, automatically dropping a supplier could be sensitive, so likely a human will decide, but the AI ensures no time is lost in identifying the problem and recommending action.

Continuous Improvement Loop: The data gathered in execution feeds back into the system to improve future decisions. This is where machine learning continuously enhances the procurement function's intelligence. For instance, the AI learns how accurate its vendor selection was by comparing expected performance (from the proposal) to actual performance – if a supplier that was chosen starts failing, next time the AI may weigh such early performance metrics more heavily. Or if an AI-negotiated contract is delivering significant savings and few issues, that pattern can be used to justify expanding AI negotiation to similar cases. In effect, the AI agents collectively form a learning organization within procurement. Additionally, feedback from procurement professionals is key – after each cycle, the team might rate how well the AI assisted, and those insights (like “the vendor selection missed a critical qualitative aspect we value”) can be used to tweak the AI models or add new rules. This is analogous to the “Check-Act” of ISO's PDCA cycle and corresponds to the NIST Manage/Govern functions ensuring continuous monitoring and improvement. Some organizations set up a Center of Excellence (CoE) for AI in procurement that regularly reviews performance metrics of the AI, addresses issues, and updates training data or algorithms as needed.

Throughout these phases, it is paramount that a governance mechanism oversees the AI's activities. We've implicitly touched on this by mentioning human approvals and flags. In the next section, we will explicitly discuss the governance structure, controls, and risk management techniques needed to ensure these powerful AI agents remain reliable tools and do not become sources of new risk. To solidify understanding, we will present a pseudocode example illustrating how an autonomous procurement cycle could be implemented with checkpoints – effectively merging the methodology just described with the governance practices that need to wrap around it.

09 Governance and Risk Management in Agentic Procurement

Adopting autonomous sourcing bots requires robust governance to ensure that efficiency gains do not come at the expense of control, ethics, or compliance. Many of the considerations are similar to those in other AI deployments, but procurement has its own nuances. In this section, we outline the key governance structures, risk management practices, and controls that should accompany the introduction of AI into procurement.

We also provide a pseudocode example that demonstrates how an autonomous procurement system might enforce these controls in practice, echoing the Map-Measure-Manage-Govern approach to highlight where safeguards come into play.

AI Governance Structure for Procurement: Just as internal audit functions have started appointing “AI champions,” procurement departments should designate clear ownership for AI initiatives. This could be a governance committee that includes the Chief Procurement Officer (CPO), IT leaders, a legal advisor, and a compliance or risk representative. The committee’s role is to approve AI use cases (decide where AI is appropriate), set policies (e.g., defining the above-mentioned thresholds for human intervention), and monitor outcomes. For example, a policy might be: “AI can autonomously approve purchases up to \$50,000 if from vetted suppliers and within budget, but anything above requires sign-off.” In the GCC public sector, oversight might also include external stakeholders – perhaps a representative from a central digital authority or an observer from an anti-corruption commission could have insight, especially in the early stages, to build trust that the system is behaving properly. In any case, clearly documenting the AI system’s purpose, scope, and limitations is key. The Icertis survey insight that 66% of CPOs prioritize leveraging AI for decision-making in 2025[48] underscores that leadership is keen, but it should be channelled through structured programs. Many organizations will include AI governance as part of their existing procurement governance forums or steering committees. One practical step is maintaining an inventory of all AI models and agents used in procurement, including details like their algorithm type, data inputs, version, and last validation date (this resonates with the ISO 42001 requirement for inventory and NIST’s Govern function). Such an inventory ensures visibility – you can’t govern what you aren’t aware of.

Risk Assessment and Validation of AI Tools: Before an AI tool goes live in procurement, it should undergo a risk assessment and validation process. This might be carried out by the procurement analytics team or in partnership with a central data science team. The risk assessment should identify potential failure modes and impacts: e.g., “What is the risk of the supplier scoring model being biased towards larger suppliers?” or “What if the negotiation agent miscommunicates and offends a supplier, causing them to withdraw?” Once risks are mapped, controls must be put in place. If bias is a risk, a mitigation could be testing the model on a diverse set of scenarios and adjusting it or adding rules (for

example, ensuring supplier scorecards consider relative performance within peer groups, so small companies aren't unfairly penalized against multinationals). If miscommunication is a risk, perhaps the negotiation bot is limited to very structured interactions or its messages are templated and approved. Validation of AI models should be done on historical cases: one could replay a past procurement (for which outcomes are known) through the AI system and see if it makes good decisions and identify differences. Did it pick the same winner? If not, why? Maybe the human saw a factor the AI didn't – that's a learning to improve the model or to note where AI shouldn't be used. A key part of validation is checking compliance: does the AI ever propose something that violates regulations or policy? For instance, if a policy says at least 3 suppliers must be invited for buys over \$100k, the AI should never try to invite fewer. These kinds of rules should be hard-coded or strongly reinforced. Testing can include edge cases too, like how the system handles a supplier that is qualified but has a known issue (maybe they're suspended or have a low performance score). After initial validation, periodic re-validation is recommended – say annually or whenever the AI model is updated or if there's a significant change in procurement policy. In the AI RMF terms, this is part of Manage (risk treatment and monitoring) and Govern (documentation, review cadence) [17][49].

Data Governance and Security: Procurement AI will consume and produce a lot of sensitive data: supplier bids (often confidential), pricing information, contract terms, and possibly personal data of supplier contacts. Ensuring proper data governance is therefore vital. Access to the AI system should be controlled – not everyone should see all supplier bids if that's not currently allowed; the AI platform must enforce user permissions similar to or stricter than the current e-procurement system. Data retention policies must be updated to include AI logs – for example, saving decision logs for X years for audit purposes. As for security, if using cloud-based AI services, procurement needs to coordinate with IT to ensure vendor solutions meet security standards (maybe requiring local hosting if data sensitivity is high, or encryption of data at rest and in transit).

Another subtle point: adversarial threats. Could a supplier try to game the AI system? This is not far-fetched – if suppliers learn that the AI looks for certain keywords or formats in proposals to score them, they might tailor or even use their own AI to exploit that.

Procurement must thus be vigilant, and perhaps not disclose too much about how the AI works. On the other hand, transparency is needed for trust. Balancing this, one approach

is to disclose criteria and ensure the AI is robust enough that gaming the system is hard (for instance, if a proposal is suspiciously filled with keywords, the AI might detect that and penalize for lack of substance). Security testing might involve simulating such attempts or potential cyber attacks on the agent (since it could become a new attack surface if not secured). The pseudocode below will illustrate, for example, checking model integrity (to ensure it hasn't been tampered with) and handling adversarial inputs. Procurement should involve cybersecurity teams to pen-test the AI-driven processes, ensuring, say, that a supplier cannot inject a script in their bid that confuses the AI or disrupts the platform.

Bias and Ethical Considerations: We've discussed fairness extensively, but to institutionalize it, procurement should include ethical AI reviews as part of their process. This might mean running a bias audit on supplier selection AI annually – e.g., checking if small enterprises win a proportionate share when the AI is used, or if there's any unintended bias against suppliers from certain countries (especially relevant if trade agreements or local preference policies exist that must be respected). The system should also be designed to uphold social and sustainability goals, not just cost. If a government has a policy to prefer local or sustainable suppliers when all else is equal, the AI should be aware of that and factor it in. It's an ethical stance turned into a functional requirement. Moreover, AI should not undermine procurement ethics; for example, confidentiality of bids is a sacred principle – the AI must not leak information from one bidder to another. If using a shared AI platform, strict data isolation is needed. Another ethical aspect is the treatment of suppliers: AI communications should remain professional and courteous. Human oversight means someone should periodically review the messages the AI sends to ensure they align with the organization's tone. A poorly phrased AI message could harm a supplier relationship or even lead to complaints. As noted in the Jaggaer piece, there's a risk of "black box procurement" if decisions can't be explained[50][51]. To counter that, procurement should demand explainability features in any AI solution and include in supplier debriefings an explanation (which could be AI-generated but human-validated) of why they lost – this maintains trust in the process.

Change Management and Training: A sometimes underappreciated governance aspect is preparing the people involved. The best controls on paper mean little if the users (procurement staff) circumvent the AI or misuse it. Therefore, training programs are

essential. As part of rolling out the AI, all procurement team members should be educated on how it works, its limitations, and the new SOPs (standard operating procedures) governing AI use. They should know, for instance, that if the AI flags an anomaly, they are expected to investigate it, not ignore it. Or that if the AI gives a recommendation, they are free to challenge it – in fact, encouraged to critically analyze it – and there’s a mechanism to provide feedback. Building trust in the system is key: initial skepticism can be overcome by involving staff in testing and gradually phasing AI in (perhaps running it in parallel with human process for a few trials to demonstrate its effectiveness). Upskilling the team in data literacy will help them feel more comfortable. Also, communicating to suppliers and other stakeholders is part of change management. Some suppliers might be uneasy if they learn “a bot” is evaluating their bids. It could be prudent to introduce it softly – maybe initially brand it as a “new analytics system” rather than anthropomorphizing it as a bot making decisions. Over time, as confidence builds, stakeholders will likely appreciate that decisions are faster and arguably more objective.

We now present a pseudocode example that ties together many governance and control elements. This pseudocode outlines how an autonomous sourcing agent might operate with embedded checks and escalation to humans. It’s simplified for illustration, but it shows a logical flow including need detection, supplier selection, evaluation with fairness checks, and points where human review is triggered. Commentary in the pseudocode (as comments) explains the control being applied at each step.

```
# Pseudocode: Autonomous Sourcing Cycle with Governance Controls function
autonomousSourcingCycle(request): # 1. Need identification (triggered by system or
user) need = analyze_request(request) if not need.approved: # Ensure legitimate need
raise Exception("Purchase need not approved by budget owner") # 2. Strategy selection
strategy = select_strategy(need) log("Strategy chosen: " + strategy.type) if
strategy.requiresApproval: pause_for_human("Confirm sourcing strategy: " +
strategy.type) # 3. Supplier selection for RFx suppliers = get_suppliers_for_need(need) if
len(suppliers) < strategy.minSuppliers: issue_log.add("Insufficient suppliers (" +
len(suppliers) + ") for competition") if strategy.allowSingleSource:
pause_for_human("Single source approval needed for " + need.description) else: throw
Error("Not enough suppliers to proceed") # Remove any blacklisted or suspended
suppliers suppliers = [s for s in suppliers if s.status != "Suspended"] # Ensure fairness:
```

```

include SMEs or local vendors if required by policy if policy.requireLocalVendor:
suppliers.append(select_local_vendor_if_missing(suppliers, need.category))
invite_to_rfx(suppliers, need.requirements) # 4. Bid evaluation responses =
collect_bids(need.id) scores = {} for bid in responses: scores[bid.vendor] =
evaluate_bid(bid, need.criteria) # Check for bias or scoring anomalies fairness_report =
analyze_scores_for_bias(scores, responses) if fairness_report.flags:
issue_log.add("Potential bias detected: " + fairness_report.detail)
pause_for_human("Review bias flags in evaluation for " + need.id) # Determine top
vendor(s) ranked = sort_by_score(scores) top_vendor, top_score = ranked[0]
second_score = ranked[1].score if len(ranked) > 1 else None # If scores are very close or
top vendor is new/unproven, escalate if second_score and (top_score - second_score) <
policy.minScoreMargin: pause_for_human("Scores close. Human review recommended
between " + top_vendor + " and " + ranked[1].vendor) if vendor_is_new(top_vendor) and
need.highValue: pause_for_human("Top vendor is new for high value contract. Review
needed.") # 5. Negotiation phase (if applicable and no human pause triggered above) if
strategy.allowsNegotiation and not human_paused: negotiated_terms =
negotiation_agent.negotiate(top_vendor, need.targetTerms) log("Negotiation result with "
+ top_vendor + ": " + negotiated_terms.summary) if negotiated_terms.exceed_limits:
pause_for_human("Negotiated outcome outside preset limits. Approval required.") else:
negotiated_terms = need.initialTerms # 6. Contract drafting and approval contract_draft =
draft_contract(top_vendor, negotiated_terms, need.standardContract) contract_issues =
legal_check(contract_draft) if contract_issues.found: issue_log.add("Contract deviations
found: " + contract_issues.summary) pause_for_human("Legal review needed for
contract " + contract_draft.id) # Final approval step for contract pause_for_human("Final
award approval for " + top_vendor + " on need " + need.id) finalize_award(top_vendor,
contract_draft) # 7. Post-award monitoring setup
setup_performance_monitor(contract_draft, metrics=need.SLA_KPIs) log("Monitoring set
for contract " + contract_draft.id)

```

This pseudocode starts by verifying the purchase need (ensuring it's approved by a budget owner), then chooses a strategy and possibly requires human confirmation (for example, if an unusual strategy is selected). It selects suppliers, enforcing a minimum number for competition and including local vendors if policy dictates (safeguarding fairness and regulatory compliance). During evaluation, it analyzes scores for bias (the

function `analyze_scores_for_bias` could check if, say, all small vendors scored low and whether that's expected or a red flag). It then triggers human review if scores are too close or if the top vendor is new and the contract high-risk. This reflects how organizations might set rules to always double-check high-impact decisions. The negotiation agent is only allowed to operate if within limits and if no prior red flags have paused the process. After negotiation, the contract is drafted and a legal check runs – any deviations from standard trigger a pause for human legal review. Finally, even after all this, it asks for a final human approval to award (especially for higher value or any public sector case, this would be mandatory, but even in private, a final sign-off by the CPO or someone might be required). The system then initiates performance monitoring automatically. Throughout, we maintain an `issue_log` to record any compliance or anomaly issues for audit purposes. This log, and the extensive use of `pause_for_human()`, illustrate how we embed governance in the automation. The AI doesn't have free rein – it's bounded by policies and oversight. Notably, if everything is routine and no flags pop up, the process flows quickly with minimal intervention, which is the efficiency gain; but the moment something is out of bounds, control reverts to humans, which is the risk management.

Auditability and Continuous Oversight: The procurement function should plan for periodic audits of the AI-driven processes. Internal audit or external auditors may review a sample of AI-handled procurements to ensure the rules were followed and outcomes were sound. Because the AI logs everything (invocations of `pause_for_human`, issue logs, final decisions), auditors will have rich data to examine. This is actually a benefit of AI – it can make compliance auditing easier than paper-based processes. The goal is to catch any drift (say, if the AI started scoring in a biased way over time due to changes in data) or any control gaps. For instance, an audit might find that every time a certain supplier bid, the AI's NLP evaluation gave it a lower score because of phrasing issues – a potential bias to fix. Or auditors might simulate a scenario of a supplier trying to bribe the AI (feeding it falsely glowing performance metrics); how would that be detected? While that sounds odd, it could equate to data poisoning attempts, which the governance framework should consider – perhaps requiring data provenance checks (e.g., supplier performance data should come from official systems, not supplier-provided only).

In terms of roles, as referenced in the Jaggaer article, new roles like an “AI procurement ethics officer” or “AI controller” might emerge[21][52]. GCC organizations with their often centralized models might assign someone in the procurement team or the IT department to continuously monitor the AI’s outputs and be the first line of defense for any anomalies. This is akin to how some banks have people monitoring algorithmic trading systems. The cost of such oversight is part of the investment into AI, but it’s crucial for long-term success and trust.

In conclusion of governance, the message is clear: agentic procurement will not run on auto-pilot without human governance. Instead, it operates within a well-defined control tower framework. Those companies and government agencies that implement these governance measures will likely enjoy the benefits of AI in procurement (speed, savings, insight) while keeping risks in check. Conversely, neglecting governance could lead to failures – such as supplier disputes, legal challenges, or simply AI tools that get sidelined because stakeholders lose trust after one mistake. Fortunately, the trajectory in 2025 is that most organizations understand this need: surveys show a majority of CPOs acknowledge data and integration issues, which implies they know it’s not plug-and-play[17]; they’re prepared to work through governance to make it succeed. Next, we examine how all these efforts translate into performance outcomes and what early results indicate about the ROI of AI in procurement.

10 Performance Impact and Early Outcomes

As organizations pilot and implement autonomous procurement solutions, a critical question is: do these technologies deliver on their promises? In this section, we evaluate the potential and observed impacts of AI in procurement on efficiency, effectiveness, cost savings, and stakeholder satisfaction. We also discuss challenges encountered in early implementations and how our proposed framework helps address them. While the GCC region is only beginning to see real-world deployments, global surveys and case studies provide strong indicators of what can be achieved.

Efficiency Gains: One of the most immediate benefits reported is the reduction in cycle time for procurement activities. Traditional competitive bidding can take weeks or months. AI can compress this dramatically by automating parallel tasks and eliminating

wait times. For instance, drafting an RFP that might take a human a week (with various iterations) can be done by an AI in minutes, then just reviewed for a day, cutting a chunk of time. Evaluations that took a committee days of meetings can be shortened to a day if AI provides initial scores and summaries. Some companies have found that what used to be a 4-6 week procurement cycle can be reduced to 2-3 weeks with AI assistance, effectively doubling the throughput of the procurement team. A Hackett Group report suggests that AI adoption could bridge an efficiency gap of roughly 9% in procurement by 2025 (meaning doing 9% more with the same resources), which is significant in an area that has been historically optimized already[53]. More anecdotally, Coupa claims its AI-based risk scoring saved one client 32 hours of analyst work per major supplier onboarded, by automating data collection and initial assessment. In the GCC government context, where procurement often involves multiple approvals, AI's drafting and checking capabilities can shorten the preparation for those approvals. Abu Dhabi's digital procurement initiative aimed at reducing lead times saw improvements just from process standardization[54][55]; adding AI could push this further by, say, automatically filling forms and reducing errors that cause rework. Efficiency also manifests in labor cost savings or reallocation: a procurement team might handle 30% more procurement events in a year without increasing headcount, because routine tasks are offloaded to bots. However, it's not just about speed – consistency and lower error rates (like fewer missed steps in compliance) also contribute to efficiency by avoiding delays associated with corrections or disputes.

Cost Savings and Spend Outcomes: The ultimate goal of procurement is often to reduce cost or obtain better value from suppliers. AI agents contribute here in multiple ways. Firstly, by analyzing vast amounts of spend data and external market info, AI can identify opportunities for savings that humans might miss – such as alternative suppliers that are cheaper or opportunities to consolidate spend for volume discounts. A generative AI might comb through contracts and find that two business units are buying the same item from different suppliers at different prices, prompting a strategic sourcing event to leverage volume. Secondly, automated negotiations and auctions driven by AI can push prices down. Even a 1-2% improvement in negotiated savings across tens of millions of dollars of spend is a big win. Companies like Zip and Zycus have boasted that their AI agents will save “millions of hours” of work and “generate billions in savings” globally[56][57], though such figures need to be taken with context. For a single organization, the

gains will scale with spend size and how well AI is integrated. Early adopters have seen significant ROI: the Icertis CPO survey highlighted that 66% of CPOs prioritize AI to improve speed-to-value and ROI[48], indicating belief that AI will directly improve financial outcomes, not just process metrics. Also, risk reduction (discussed below) has an indirect financial benefit – avoiding costly supply failures or compliance fines. In public sector, cost savings might translate to better use of public funds or ability to reinvest savings into more projects, aligning with government efficiency drives in GCC countries to maximize impact of their budgets. The key point is that AI brings a more analytical, data-grounded approach to sourcing, reducing information asymmetry with suppliers (suppliers often have more market data, but now procurement can match that with AI insights). Over time, as AI learns and optimizes, we expect a trajectory of increasing savings – perhaps initial 3-5% improvements in addressed categories and potentially up to 10% in some cases once AI maturity is high, as indicated by some procurement consulting predictions[58][59]. However, these outcomes require that the organization actually acts on AI recommendations – which ties back to trust and change management.

Coverage and Risk Management: AI enables wider coverage of spend under management and risk monitoring. A typical procurement department might actively manage (strategically) only the top 80% of spend, and the rest is often unmanaged or “maverick” (people buying outside processes). By deploying user-friendly AI buying assistants (like a chatbot for end-users to requisition items that then automatically sources them under the hood), organizations can bring more spend under the procurement umbrella without overburdening staff. This reduces off-contract purchasing and ensures compliance with preferred pricing. As for risk, AI’s continuous monitoring means fewer things slip through cracks. Several CPOs in 2025 are concerned with supply chain disruptions and supplier risks[60]. AI can watch myriad risk indicators (financial news, delivery performance, social media signals about a supplier) and give early warnings, something humans realistically cannot do at scale. A Middle Eastern telecom that piloted an AI supplier risk dashboard found it could flag risk events 2-3 weeks before traditional risk reviews would have caught them, giving them a head start in mitigating issues (for example, proactively sourcing an alternate supplier when AI showed their current supplier’s region had an emerging COVID lockdown). In terms of regulatory compliance (like ensuring each procurement follows the rules), the AI’s inherent adherence to coded rules can lead to near 100% compliance in execution – any deviation stands out as an exception. This has

been seen in finance with RPA – when properly set, it never forgets to do a step. Similarly, an AI agent won't "forget" to check a debarment list or to include required evaluation criteria, whereas a human might occasionally slip. So, risk of non-compliance goes down, which is very valuable for public entities to avoid audit findings.

Stakeholder and Supplier Reactions: An important measure of success is how people on both sides of the process feel about the new system. Internally, early users often report that once teething issues are resolved, they appreciate having tedious tasks automated. A procurement officer who used to spend hours compiling bid tabs can now use that time to engage stakeholders or work on strategy. Job satisfaction can increase as the work shifts to more interesting analysis and supplier engagement. However, if not managed well, there can be initial anxiety or frustration (for example, if the AI interface is clunky or if they don't understand why the AI recommends something). That's why training and involving them in design is critical. The Jaggaer blog warns that without proper change management, freed-up capacity might be misused or staff might resist trusting the AI[61] [62], but with strong leadership and redesign of roles, procurement can become more rewarding. On the supplier side, reactions can vary. Some larger, tech-savvy suppliers adapt easily – they are happy to have a clear, fast digital interaction (a bot that answers their questions at 11 PM is better than waiting for an email answer next day). They might also appreciate faster decisions and feedback. Smaller or less tech-capable suppliers might find it challenging to interface with a system rather than a person; it will be important for organizations to not inadvertently exclude such suppliers. One approach is to keep an option for suppliers to reach a human or to submit bids in a traditional way if they are unable to interface digitally – the AI can then input that data on their behalf. Over time, as digital literacy improves and more governments mandate e-procurement, suppliers will likely acclimate. A side effect of AI in procurement could be improved supplier quality over time: since performance is tracked so closely, suppliers know they must maintain standards or the AI will flag them quickly. This could encourage a healthy competition and weeding out of underperformers, leading to a more reliable supply base for the organization.

Challenges and Adjustments in Early Deployments: Learning from initial projects, a common challenge is fine-tuning the AI to minimize false positives/negatives. For example, an AI evaluation model might initially flag too many bids as non-compliant (false

positives) because of minor deviations; procurement had to adjust criteria to be more tolerant or train the model better on what is truly disqualifying. At a GCC oil & gas company pilot, the AI classified a vendor as high risk due to lack of data (it was a new local vendor with scant info), almost eliminating them, whereas in reality they were a promising newcomer supported by a local content program. The lesson was to incorporate human judgment to override in such cases and also feed the AI more data from local directories. Another challenge is integration delays – getting AI to talk to legacy ERP systems can be non-trivial, causing initial slowdowns. Some organizations solve this by using RPA as a bridge (not elegant, but it works) or by upgrading core systems in parallel. The Icertis study noted integration issues were top of mind[17], confirming this is a widespread friction point. Our framework's phased approach (e.g., starting AI in advisory mode before fully automating decisions) helps because it allows catching these issues when humans are still closely in the loop, rather than after automation when it could cause bigger problems. It's akin to having a pilot phase where you double-check everything the AI does. Indeed, a bank in the UAE running an AI for procurement insights kept a "shadow" procurement analyst doing parallel work for the first few months and compared results to ensure the AI was reliable. They found the AI missed some contextual factors (like a supplier's political connections making them a risk in one case), which then led to building a rule for that scenario. After some cycles of improvement, trust increased and they could reduce the parallel effort.

Metrics and Continuous Improvement: Organizations implementing AI in procurement are establishing new KPIs to track its impact. These include: cycle time reduction (%), average number of bids per tender (which could increase if AI makes it easier to include more suppliers – more competition can yield better outcomes), percentage of spend under AI-assisted management, compliance rate (e.g., no. of events that fully followed policy, hopefully nearing 100%), savings realized versus target, and user satisfaction scores for the procurement process. Baselines for these are recorded pre-AI and then tracked post-AI deployment. One GCC ministry reported that after implementing an AI-based analytics tool for procurement, their compliance rate in spend reporting went from ~70% to 95% because the AI automatically captured many off-book purchases that staff used to not report properly. This underscores how AI can enforce discipline. The pursuit of these metrics also creates a culture of continuous improvement: teams regularly review if the AI is hitting targets (like, is it really cutting 30% off cycle time or are there

hold-ups? If not, why – maybe the bottleneck just moved to approvals and thus we need to streamline those too). In our framework, we embedded continuous monitoring and feedback loops, which will ensure that performance doesn't plateau. It's quite possible that the first few AI-driven procurements will be slower than normal (due to learning curve), but then things accelerate. Setting realistic expectations is crucial – for example, aim for a modest improvement in the first 6 months, then scale up as confidence and proficiency grow.

Looking across these dimensions, early outcomes are validating the promise of agentic procurement. A 2025 survey by EY found 80% of global CPOs plan to deploy generative AI in some capacity over the next three years[63], which reflects a broad consensus that the benefits outweigh the concerns. Specifically in the Middle East, procurement leaders often see themselves as pioneering – there's a sense of pride in adopting cutting-edge solutions as part of national innovation goals. Thus, we expect GCC case studies to emerge soon, showcasing dramatic improvements – perhaps a UAE government body cutting its average tender time from 60 days to 30 days, or a Saudi conglomerate achieving multi-million dollar savings through AI-optimized sourcing that their competitors (without AI) couldn't match. These success stories will further fuel adoption. However, they will also shine a light on the importance of our earlier discussion: if something goes wrong (say, an AI-recommended supplier fails badly), it will equally be a cautionary tale. Our proposed governance and methodological framework aims to maximize the former and minimize the latter, ensuring that as performance metrics climb, the trust and integrity of procurement processes remain rock-solid.

11 Conclusion and Future Outlook

The advent of autonomous sourcing bots and agentic AI is poised to transform procurement in the GCC and around the world. As we have explored, this transformation is not just about technology implementation, but about rethinking processes, controls, and roles in the procurement function. By 2025, we are at the cusp: numerous proofs-of-concept show what is possible, and leading organizations are moving from experimentation to institutionalization of AI in their procurement operations. The GCC region, with its strong appetite for innovation and top-down support for digital initiatives, has a significant opportunity to leapfrog traditional procurement inefficiencies and set

new benchmarks for speed, transparency, and strategic impact of procurement. But seizing this opportunity requires doing so responsibly – aligning with global best practices and governance frameworks to avoid pitfalls and ensure stakeholder trust.

The framework presented in this paper serves as a blueprint for building an AI-enabled procurement function that is both ambitious and prudent. We demonstrated how AI agents can be embedded throughout the procurement lifecycle – from the moment a need is identified to the final contract and beyond into performance monitoring – effectively automating routine tasks and augmenting human decision-making. When guided by standards like ISO 42001 and NIST's AI RMF, and when grounded in procurement's core principles of fairness and accountability, this agentic procurement approach can dramatically enhance outcomes. It allows procurement teams to handle more volume with greater analytical rigor, enabling them to shift their focus to strategic matters like supplier innovation, market intelligence, and aligning procurement with organizational goals (e.g., sustainability or localization targets). In essence, AI takes over the heavy lifting of data processing and transactional execution, while humans concentrate on oversight, exceptions, and relationship management. This human-machine synergy is the sweet spot – neither pure automation nor status quo manual, but a blended model that leverages each side's strengths.

For the GCC public sector, agentic procurement aligns with government directives to improve public service delivery and transparency. An AI-driven system can reduce opportunities for human error or bias that sometimes plagued procurement, thereby increasing public trust. It also enables the kind of agility needed for fast-moving projects (like the mega-projects in Saudi Arabia's Vision 2030 or the UAE's rapid infrastructure development) where conventional procurement cycles might be too slow. Moreover, GCC governments are looking to be leaders in AI adoption; showcasing AI in their own administrative processes, like procurement, sets an example and builds local expertise that can spill over into other sectors. On the enterprise side, industries such as oil & gas, telecom, and aviation – which are prominent in the Middle East – have complex procurement needs and global supply chains. They stand to gain enormously from AI optimization and risk management capabilities, potentially saving millions and preventing disruptions in critical operations. We might soon hear about a GCC oil company that avoided a costly shutdown because its AI procurement agent proactively sourced an

alternative supplier upon detecting risk, or a regional airline that cut its aircraft parts inventory cost by a large percent because AI improved its just-in-time sourcing accuracy.

However, the journey will require continuous attention to governance and ethics. As AI systems take on more decision-making, questions will arise about legal liability, regulatory acceptance, and ethical boundaries. For example, if an AI inadvertently discriminates against certain local businesses, there could be political ramifications. Regulators in the region may at some point issue guidelines or mandates on AI transparency in government procurement (echoing global trends). Procurement leaders should stay ahead by self-regulating effectively. This includes maintaining the auditability of AI decisions and being ready to explain and defend them. The region's moves like the UAE's AI Ethics principles and Saudi's AI adoption frameworks indicate that governments are aware and proactive in this space; procurement must incorporate these into their operational fabric. Encouragingly, the principles – fairness, transparency, accountability – are very much the ones procurement professionals already champion. The technology changes, but the foundational values do not.

From a people perspective, as AI handles more of the grunt work, procurement professionals will need to elevate their skills. The profile of the procurement officer of the future will likely emphasize analytical skills, data interpretation, and the ability to manage AI systems, along with the traditional negotiation and stakeholder management skills. Training programs and certifications (perhaps new modules from CIPS or NCMA focusing on AI in procurement) will become important. GCC organizations might collaborate with academic institutions or global partners to build AI procurement knowledge. Notably, we foresee new roles as mentioned: the AI procurement strategist, AI compliance officer, etc. This could provide exciting career paths and attract tech-savvy talent into the procurement field, which historically might not have been the first choice for such talent. In that sense, AI could revitalize the profession's image, making it a cutting-edge place to work.

The future outlook also includes the progression of technology itself. What is semi-autonomous in 2025 could become more autonomous by 2030 with improved AI capabilities and trust. Perhaps routine procurements up to certain values will be entirely touchless – the AI does it all, and audit later confirms everything was in order. We might also see integration of procurement AI with other AI systems in finance, maintenance, and

operations forming a larger intelligent enterprise ecosystem. For example, an AI in maintenance predicts a machine failure, triggers procurement's AI to buy a replacement part, the finance AI agent allocates budget, and the warehouse robot receives it – all with minimal human input. GCC smart cities and digital government visions certainly encompass such interlinked autonomy. The challenge then becomes cross-functional governance and ensuring these agentic systems work towards organizational objectives without unintended consequences (which again underscores the need for frameworks and possibly AI regulators or auditors embedded across processes).

In conclusion, agentic procurement is set to become a cornerstone of modern business and government operations. GCC countries are well-positioned to lead in this domain given their investment capacity and strategic intent to embrace AI. By building AI-enabled procurement functions under the guardrails of established standards and ethical practices, they can achieve faster, smarter, and more accountable procurement. The key takeaway from our research is that success lies in the balance: empower AI to do what it does best (speed, data-crunching, pattern recognition), while humans continue to guide, oversee, and inject judgment where needed. Organizations that strike this balance will find that their procurement function can evolve from a cost center to a value generator and strategic enabler in the AI era. The coming years will no doubt bring both success stories and cautionary tales – by adhering to the principles and framework detailed in this paper, practitioners can ensure their story is one of success, with agentic procurement delivering exceptional results in a trustworthy manner. The journey has begun, and the next chapter in procurement's evolution is being written with algorithms and human wisdom together.

12 References

13 A. McMillan, " How Agentic AI is Propelling Procurement Forward ," Procurement Magazine

- , Sep. 11, 2025[11][64].
- J. Gerber, " Autonomous Procurement: How to Get Ahead in 2025 ," Suplari Blog, Aug. 6, 2025[65][66].

- D. Evans, "What Might a Future with Agentic Procurement Look Like?," JAGGAER (Blog), Aug. 21, 2025[36][37].
- Icertis, "90% of Procurement Leaders to Adopt AI Agents in 2025," News Release (ProcureCon CPO Survey), Jan. 28, 2025[3][48].
- Art of Procurement (P. Ideson), "My AI Governance Framework for Procurement," May 22, 2025[5][67].
- M. Ambrozie et al., Panel Discussion "Agentic AI in Procurement," Procurement & Supply Chain LIVE, London, 2025[1][56].
- Deloitte, "2025 Global CPO Survey Highlights," Deloitte Insights, 2025[68][69].
- OECD, "OECD Principles on AI," 2019. (Transparency, accountability and fairness as core principles)[32].
- NIST, "AI Risk Management Framework 1.0," Jan. 2023[24][16].
- Consultancy ME (McKinsey & Co.), "AI can add \$150 billion to GCC economies," Aug. 23, 2023[8].
- M. Al-Khalifa, "AI Regulations in the GCC – Part One (UAE, Saudi, Bahrain)," Law Library of Congress Blog, Dec. 17, 2024[14][15].
- E. Hofverberg, "AI Regulations in the GCC – Part Two (Qatar, Kuwait, Oman)," Law Library of Congress Blog, Jan. 2025.
- The Hackett Group, "64% of Procurement Leaders Say AI Will Fundamentally Change Procurement by 2025," 2024 (Research Report)[70].
- EY, "State of AI in Procurement 2025 – Global CPO Survey," 2025[71].
- NIGP (Institute for Public Procurement), "Artificial Intelligence: Impact on Public Procurement," White Paper, 2024[13].