

Defensible FM Procurement in the GCC

Why spreadsheet-based tendering has reached its limits — and how a governed procurement platform, powered by an ontological and epistemological AI layer, can raise the quality, speed and auditability of RFP tender packages and the clarifications process.

FM Procurement

Governance

Applied AI

GCC Property

AUTHOR

Douglas Masuku

BSc Facilities Management · MSc Occupational Health & Safety · MSc Intelligent Buildings
CIWFM · MIFMA

ifm-procure.solutions

Executive summary

Facilities management in the Gulf Cooperation Council is now a market approaching **USD 60 billion** and growing at more than 5% a year, yet the way FM services are bought has not kept pace with the scale, complexity or governance expectations of the region's giga-projects and public-sector reform agendas. A significant share of tendering is still run on spreadsheets and email — a method that is error-prone, difficult to audit, and increasingly hard to defend when a decision is challenged.

This paper argues three things. First, that the cost of informal, spreadsheet-based FM procurement in the GCC is no longer administrative inconvenience but a governance and value-for-money risk. Second, that a purpose-built procurement platform can directly serve an organisation's strategic objectives — value, compliance, speed and defensibility — rather than merely digitising paperwork. Third, that the most consequential advance is not automation for its own sake but an **ontological and epistemological AI layer**: a shared, machine-readable model of FM procurement knowledge (the ontology) coupled with disciplined rules for how the system knows and justifies what it asserts (the epistemology). Together these let AI draft higher-quality RFP tender packages and manage the clarifications process with traceability regulators and auditors can trust.

01 The GCC procurement inflection point

The GCC facilities management market was valued at approximately **USD 60.1 billion in 2025** and is forecast to reach **USD 77.5 billion by 2030**, a compound annual growth rate of about 5.2%. Saudi Arabia and the United Arab Emirates account for the majority of that spend, with the Saudi market alone projected to move from roughly USD 26 billion in 2024 to nearly USD 39 billion by 2030. This growth is not organic drift; it is being manufactured by deliberate state policy — Saudi Vision 2030, the UAE's diversification agenda and Oman Vision 2040 — through giga-projects, tourism assets, healthcare expansion and smart-city programmes that all generate long-duration, service-intensive FM demand.

\$60.1bn

GCC FM market size, 2025 (MarkN-tel Advisors)

\$77.5bn

Forecast market size by 2030, at ~5.2% CAGR

~60%

Share concentrated in KSA & UAE, the two reform frontrunners

Two structural shifts make procurement — not service delivery — the pinch point. The first is **professionalisation of the buy-side**. Government entities are consolidating spend on national e-procurement rails such as Saudi Arabia's Etimad platform, and are attaching local-content, transparency and value-for-money conditions to award decisions. The second is the **rise of outcome-based and bundled FM contracts**, where an owner buys a performance standard across hard, soft, risk and administrative services rather than a schedule of tasks. Both shifts raise the intellectual difficulty of writing a good tender and evaluating the responses fairly. The scope must be unambiguous, the evaluation criteria must be defined in advance and applied consistently, and the whole chain of reasoning must survive scrutiny. That is precisely what informal tooling cannot guarantee.

THE CORE TENSION

Procurement volume, contract complexity and governance expectations are all rising together — while the dominant tools for running FM tenders in the region remain the spreadsheet, the shared inbox and the individual buyer's memory.

02 Why spreadsheet-based tendering has reached its limit

Spreadsheets are the default because they are universal, cheap and flexible. That same flexibility is the problem. A tender run across a workbook and an email thread has no enforced structure, no reliable version history, no separation of duties, and no native record of *why* a decision was made. In a low-stakes purchase this is tolerable. In a multi-year, multi-million-dirham FM contract subject to audit and potential challenge, it is a latent liability.

The error problem

The reliability of manual spreadsheets has been studied for decades, and the findings are consistent: audits of real-world operational spreadsheets repeatedly find that a large majority contain at least one material error. When a weighted evaluation model, a price-normalisation formula or a linked summary sheet sits at the centre of a tender award, a single mis-keyed cell reference can quietly change the ranking of bidders — and no one may ever know. Unlike a controlled system, a spreadsheet does not warn you when its own logic has broken.

The governance problem

Defensible procurement rests on a simple principle: decisions must be made against criteria set in advance, applied evenly to every bidder, and evidenced by a record that a third party can reconstruct. RICS's professional standard on the procurement of facility management and the broader body of RICS and IWFM guidance all point to the same disciplines — clear scope, transparent and proportionate evaluation, and a robust audit trail. Spreadsheets and inboxes struggle against every one of these. Criteria drift as the workbook is edited; evaluator scores are entered without a locked rationale; the clarifications sent to one bidder are not systematically shared with all; and the final audit trail has to be reassembled after the fact from files and emails rather than captured as the process runs.

Requirement of a defensible tender	Spreadsheet & email reality	Governed platform
Criteria fixed before bids are opened	Weightings editable at any time, often after responses arrive	Criteria and weightings locked and time-stamped pre-issue
Consistent, evidenced scoring	Free-text or blank rationale; scores overwritten silently	Structured scores with mandatory, retained justification
Equal treatment on clarifications	Q&A scattered across inboxes; uneven distribution	Single clarifications register, published to all bidders
Reconstructable audit trail	Assembled retrospectively from files and emails	Captured automatically as an immutable event log
Separation of duties	One workbook, shared credentials, no role control	Role-based access with recorded actions per user

The strategic point is that these are not IT inconveniences. When a losing bidder challenges an award, when an internal auditor tests value-for-money, or when a regulator reviews a public entity's spend, the organisation must *prove* that the process was fair. A spreadsheet cannot carry that burden of proof; a governed system is designed to.

03 What a procurement platform must do for the organisation

A platform earns its place only if it advances the organisation's own objectives rather than adding a layer of software administration. For an FM owner, developer or public entity in the GCC, those objectives are consistent: extract genuine value for money, stay compliant with tightening regulation, move faster without cutting corners, and be able to defend every decision. A well-designed platform maps to each.

- **Value for money, evidenced.** Structured scope and standardised evaluation make like-for-like comparison possible, so award goes to the most economically advantageous tender rather than the best-formatted bid.
- **Compliance by construction.** The process enforces the sequence — locked criteria, sealed responses, equal clarifications, retained scoring — so that following the workflow *is* following the policy.
- **Speed without sacrifice.** Reusable templates, a managed clarifications register and assisted document generation compress the calendar of a tender from weeks to days.
- **Defensibility as a by-product.** Because every action is captured as it happens, the audit trail is a live artefact, not a post-hoc reconstruction.

These are necessary but, on their own, no longer sufficient. Digitising a broken process only makes it faster. The differentiating capability — and the reason FM procurement is now an AI problem, not merely a workflow problem — lies in the intelligence layer that helps buyers actually write and adjudicate the tender.

04 The ontological and epistemological AI layer

THE CORE OF THE PLATFORM

The temptation with generative AI is to bolt a chatbot onto a document store and let it draft tenders from raw text. In a governed procurement context this is dangerous: a general-purpose language model will produce fluent output that is plausible, inconsistent and unaccountable. It may invent a clause, misread a scope boundary, or answer a bidder's clarification in a way that contradicts the RFP — and it cannot tell you why it said what it said. For procurement, *how the system knows* matters as much as *what it produces*. That is why the platform is built on two disciplined layers rather than a single model.

The ontological layer — *what exists*

A formal, machine-readable model of the FM procurement domain: the entities (assets, service lines, SLAs, KPIs, scope items, evaluation criteria, bidders, clarifications), their attributes, and the rules that connect them. It encodes, for example, that a "reactive maintenance response time" is a KPI attached to a hard-FM service line, measured in hours, and referenced by a specific evaluation criterion. The ontology gives every AI action a shared, unambiguous vocabulary that means the same thing to the buyer, the model and the auditor.

The epistemological layer — *how it is known*

A set of rules governing how the system justifies what it asserts: which source a statement is drawn from, how confident it is, what it may and may not infer, and when it must defer to a human. Every generated clause, answer or score carries its provenance — the standard, prior tender, policy or bidder submission it rests on. If the knowledge is not grounded, the system abstains rather than inventing. This is what turns AI output from a plausible guess into a defensible, traceable position.

The distinction is deliberate and it is the intellectual heart of the platform. An **ontology** tells the system what the world of FM procurement is made of and how its parts relate. An **epistemology** governs how the system is allowed to move from evidence to assertion — what counts as knowing, rather than guessing. Most AI failures in high-stakes domains are not ontology failures (the model has plenty of vocabulary); they are epistemology failures — the system asserts things it has no grounds to assert. By making both explicit, the platform grounds its language model in a structured knowledge graph and constrains it to reason only from sources it can cite.

An ungoverned model answers every question. A governed one knows the difference between what it can justify and what it cannot — and says so. In procurement, that difference is the whole game.

Applied to the RFP tender package

When a buyer initiates a tender, the AI layer does not free-associate a document. It assembles the RFP from the ontology: it selects the relevant service lines, pulls the associated scope items, SLAs and KPIs, and proposes evaluation criteria that are structurally linked to those requirements. Because each element is an object in the knowledge graph rather than a paragraph of prose, the package is internally consistent by construction — the evaluation criteria cannot reference a KPI the scope never defined, and the pricing schedule aligns with the services actually specified. The epistemological layer attaches to every generated clause a provenance trail: this requirement derives from that standard; this criterion reuses that clause from a prior, similar contract. The buyer reviews and approves; the system never issues on its own authority.

1 Structured intake

The buyer describes the assets, service scope and objectives; the ontology maps them to known entities and flags gaps or ambiguities before drafting begins.

2 Grounded assembly

The package — scope, SLAs, KPIs, evaluation model, pricing schedule — is generated from linked objects, so the document is coherent end to end rather than stitched from templates.

3 Provenance & review

Every clause carries its source and confidence. The buyer sees what the AI relied on, edits with full context, and formally approves before issue.

4 Locked issue

On release, scope and criteria are time-stamped and sealed, establishing the fixed baseline against which every bid and clarification is later judged.

Applied to the clarifications process

Clarifications are where tenders most often become indefensible. Questions arrive in volume, under time pressure, and an answer given to one bidder that is not shared with the rest — or that quietly changes the meaning of the scope — can taint the entire award. The AI layer treats each clarification as a first-class object linked to the specific RFP element it concerns. When a question arrives, the system retrieves the exact governing scope item or criterion from the knowledge graph and drafts a proposed answer grounded strictly in the issued documents. Crucially, the epistemological rules prevent it from introducing new requirements: if a truthful answer would change the scope, the system flags it for a formal addendum to *all* bidders rather than answering privately. The human clarifications owner approves each response, and the platform publishes it to the shared register with a full record of who asked, what was answered, on what basis, and when.

WHY THIS IS DEFENSIBLE

Every answer is traceable to an issued clause, every bidder sees the same information at the same time, and the reasoning behind each response is retained. If the award is later challenged, the clarifications record is evidence of fair and equal treatment — generated as a by-product of the work, not reconstructed under pressure.

The human remains accountable

None of this removes the professional from the decision. The ontological and epistemological layers exist to make the human faster and better-evidenced, not to automate judgement away. The AI proposes; the buyer disposes. What changes is that the buyer now works from a consistent, grounded draft with visible provenance, and spends their expertise on the decisions that matter rather than on assembling documents and chasing email threads. Accountability — legal, professional and regulatory — stays firmly with the organisation and its people.

05 Governance, defensibility and the audit trail

The combined effect of a governed workflow and a grounded AI layer is that defensibility stops being an afterthought and becomes an emergent property of the system. Three characteristics matter most to auditors, regulators and boards.

Provenance. Every material artefact — a scope item, an evaluation criterion, a clarification answer, an evaluator's score — can be traced to its source and its author. Nothing in the record is unattributable. **Immutability.** Actions are captured as a time-stamped event log as they happen, so the trail cannot be quietly rewritten after a decision is questioned. **Explainability.** Because the AI reasons over a structured ontology and cites its sources, the organisation can answer the hardest question in any procurement dispute — "why did you decide this?" — with a documented chain of reasoning rather than a recollection. For public entities operating under Vision 2030-era transparency mandates, and for private owners answerable to investors and boards, this is the difference between a decision that holds and one that unravels.

Governance question	Informal process answer	Platform answer
Were criteria set before bids opened?	"We believe so" — hard to prove	Time-stamped locked baseline
Were all bidders treated equally on clarifications?	Depends on inbox discipline	Single published register, equal access
Why did this bidder win?	Reconstructed from memory and files	Scored rationale with retained evidence
Can the AI's output be trusted?	Ungoverned model, no provenance	Grounded, cited, human-approved

06 Conclusion and outlook

The GCC is building faster than almost anywhere on earth, and every asset it raises will be operated, maintained and eventually re-tendered for decades. The procurement decisions taken now set the cost and quality baseline for that entire lifecycle. Continuing to run those decisions on spreadsheets and inboxes is a false economy: it looks cheap until the first serious challenge, audit finding or value leakage exposes how little of the reasoning was ever captured.

A governed procurement platform addresses the workflow. An ontological and epistemological AI layer addresses the intelligence — helping buyers write better tender packages and run defensible clarifications, while keeping provenance and human accountability at the centre. This is not automation replacing professional judgement; it is a system that makes professional judgement faster, more consistent, and demonstrably fair. For FM leaders across the Gulf, the organisations that adopt this discipline early will not only lower their procurement risk — they will be able to prove, on demand, that every dirham and riyal was spent defensibly. In a region where scrutiny is rising as fast as spend, that proof is fast becoming the point.

THE TAKEAWAY

Digitising procurement makes it faster. Grounding it in a shared ontology and a disciplined epistemology makes it *defensible*. In the GCC's high-scrutiny, high-growth environment, defensibility is the competitive advantage.

Sources & further reading

1. MarkNtel Advisors, *GCC Facility Management Market — Size, Growth & Forecast to 2030* (2025). Market sized at USD 60.1bn in 2025, USD 77.5bn by 2030 at 5.22% CAGR.
 2. RICS, *Procurement of Facility Management*, RICS Professional Standard, 1st edition (October 2020).
 3. RICS, *Tendering Strategies*, RICS Professional Guidance (Black Book series).
 4. Arthur D. Little, *Public Procurement Transformation in the GCC Region*.
 5. JAGGAER, *Public Sector Procurement Transformation in the GCC*; Etimad national e-procurement platform, Kingdom of Saudi Arabia.
 6. Enterprise Knowledge, *Ontology and Knowledge Graph in the Age of AI and Agents*; LF AI & Data, *The Role of Ontologies in Enabling AI Transparency*.
 7. Panko, R. et al., research on spreadsheet error rates in operational use (widely cited audits finding material errors in the large majority of real-world spreadsheets).
-

Authored by Douglas Masuku (BSc Facilities Management; MSc Occupational Health & Safety; MSc Intelligent Buildings; CIWFM, MIFMA) and published by ifm-procure.solutions for information and thought-leadership purposes. Market figures are drawn from third-party research as cited and are indicative. Nothing herein constitutes legal, financial or procurement advice; organisations should validate approaches against their own regulatory obligations. © 2026 ifm-procure.solutions. All rights reserved.