

GCC Agentic AI Acquisitions Arabic Workflows and Regulated Sector Access

A Transaction Framework for Localisation Distribution Governance and Sustainable Value

Authored by Chennakeshav Adya

Independent Researcher

September 2026

Abstract

Agentic artificial intelligence businesses entering the Gulf Cooperation Council can appear strategically valuable because they combine a fast-growing technology category with Arabic capability and access to regulated customers. The transaction case can still fail when the target provides an Arabic interface without local workflow competence, depends on data it cannot transfer, relies on founders or partners for market access, or cannot preserve human accountability in consequential decisions. A buyer therefore needs evidence that connects language performance, domain controls and distribution rights to accepted customer outcomes and recurring cash.

This paper develops a transaction framework for acquisitions of agentic AI businesses serving GCC markets. It distinguishes Modern Standard Arabic, dialect, terminology, cultural context and sector-specific process knowledge. It examines data provenance, evaluation, deployment location, customer permissions, financial-services controls, healthcare safeguards, government procurement, energy and infrastructure operations, identity, delegation, human review, model dependence, channel economics and key-person continuity. The framework translates these findings into sustainable EBITDA, evidence-weighted synergies, valuation, deal protections and a first-hundred-day plan.

The analysis draws on UAE and Saudi AI strategies, data-protection materials, financial-sector guidance, Arabic language-model research, NIST risk-management resources, competition rules and accounting standards [1-50]. These sources describe public frameworks and research results. They do not establish the quality, compliance, customer access or value of a particular acquisition target. Every target claim requires current technical, legal, commercial, accounting and regulatory verification.

A hypothetical case illustrates the method. Reported annual revenue is USD 48 million and reported EBITDA is USD 12 million. Normalising Arabic data-rights remediation, local governance and control, regulated-sector sales and support, evaluation operations and retention reduces sustainable EBITDA to USD 6.5 million. Gross annual synergy of USD 10 million becomes USD 4.1 million after continuing control, revalidation, migration, retention and client-remediation costs. An illustrative valuation bridge starts with eleven times sustainable EBITDA, adds USD 15 million of evidence-weighted synergy present value and deducts USD 27.5 million of integration, control and market-access risk, producing USD 59 million. All amounts are management assumptions used solely to demonstrate the framework.

The analysis finds that transferable value resides in accepted Arabic workflows, lawful and reproducible data use, documented decision authority, durable customer permissions and a local operating system that survives a change of ownership. Language scores, customer logos and regulatory references provide incomplete evidence in isolation. The buyer should release value only after representative workflows pass language, domain, control, customer-acceptance and cash gates.

JEL Classification: G24, G34, L86, M15, O33, Z13

Keywords: agentic AI, Arabic workflows, GCC, mergers and acquisitions, regulated sectors, localisation, data rights, distribution, valuation

Introduction

The GCC presents an attractive setting for agentic AI acquisitions. Governments and regulated industries are investing in artificial intelligence, Arabic language technology is improving, and local institutions increasingly expect digital services that reflect regional language, law and operating practice. These conditions can support growth, but they also create a diligence problem. A target may use the word Arabic for translated screens, the word agent for automated sequences and the word access for a relationship that does not survive the seller.

Arabic workflow localisation reaches beyond translation. A customer interaction can move among Modern Standard Arabic, Gulf dialects, English, mixed-script text, local terminology and institution-specific forms. The accepted result can require an authorised employee, a recorded rationale and evidence that remains available for complaint handling, audit or supervision. A target that performs well on general language benchmarks may still fail a customer-specific task with legal or financial consequences.

Public frameworks reinforce the need for an operating view. The UAE National Strategy for Artificial Intelligence 2031 identifies governance, regulation, data, talent and priority sectors as enablers [1-2]. UAE financial regulators require governance, accountability, validation, transparency and appropriate data management for material AI use [3-6]. Saudi materials address AI ethics, adoption and personal-data governance [7-12]. NIST resources organise risk work around governance, mapping, measurement and management [13-15]. These materials inform diligence; applicability remains a matter for current advice and target facts.

This paper is designed for boards, strategic buyers, financial sponsors, lenders and management teams. It treats Arabic workflow capability and regulated-sector access as testable acquisition assets. The decision standard is whether the combined business can continue to deliver accepted outcomes, preserve rights and controls, retain customers and earn cash after ownership, people, models and distribution arrangements change.

1 Define the acquisition decision

The investment committee should state the proposed value mechanism before opening the data room. Possible mechanisms include acquiring Arabic workflow capability, accelerating access to regulated customers, adding sovereign or private deployment, obtaining domain data rights, strengthening local delivery or combining the target with the buyer's distribution. Each mechanism needs separate evidence and should have a named owner, timing, cost and failure condition.

The perimeter should identify legal entities, products, model and agent layers, datasets, evaluation assets, integrations, licences, deployments, customer contracts, approvals, partners and people. The buyer should distinguish assets owned by the target from assets supplied by founders, cloud providers, model vendors, systems integrators or customers. The transaction thesis is vulnerable when a critical component cannot be transferred, replaced or operated after closing.

No-go findings should be defined in advance. Examples include Arabic training or evaluation data without demonstrable rights, regulated workflows that lack accountable human authority, customer agreements that require consent to a change of control, material revenue tied to one relationship, or deployment architecture that conflicts with customer data-location commitments. Pre-agreed failure conditions improve discipline when competitive auction pressure rises.

Table 1 GCC agentic AI acquisition diligence perimeter

Value claim	Required evidence	Decision question	Principal risk
Arabic workflow capability	cohort tests accepted outputs and native-speaker review	does the system perform the customer's real task	translated interface without workflow competence
Regulated-sector access	contracts procurement records permissions and active use	does access survive ownership and integration	relationship or approval is nontransferable
Local data asset	provenance licences consent and deletion controls	can data be lawfully used after closing	remediation restriction or loss of model value
Distribution advantage	pipeline conversion renewals and channel terms	can the buyer reproduce customer acquisition	founder or partner dependence
Agent autonomy	identity authority approval logs and recovery	are consequential actions bounded and attributable	uncontrolled action or weak evidence
Economics	revenue margin support control and migration cost	what earnings and cash remain after full cost	reported margin excludes local operating burden

Proposed structure; target-specific legal regulatory technical commercial accounting and security review is required.

2 Distinguish an Arabic interface from a localised workflow

An Arabic interface changes visible language. A localised workflow preserves meaning, authority and outcome across the complete process. It captures the customer's terminology, document conventions, exception routes, approval hierarchy, communication style and evidence requirements. The buyer should test whether the target can move from instruction to accepted result without an undisclosed human translation or repair layer.

Four layers deserve separate scores. Linguistic localisation covers script, grammar, dialect and code-switching. Domain localisation covers terminology, forms and decision rules. Institutional localisation covers roles, delegations, records and escalation. Commercial localisation covers procurement, support, service levels, pricing and customer acceptance. Strength in one layer cannot compensate automatically for failure in another.

The diligence team should select representative production cohorts and compare Arabic, bilingual and English paths. It should record first-pass acceptance, corrections, escalation, latency, cost and downstream outcomes. A polished demonstration has limited evidential weight when it uses prepared prompts, a narrow dialect or expert intervention that is absent from normal operations.


3 Map the Arabic workflow evidence chain

The evidence chain begins with the customer's instruction and ends with an accepted action, artifact or decision. Between those points sit language detection, dialect interpretation, retrieval, policy selection, model inference, tool calls, human review, system updates and communication. Each step can change meaning or authority. The transaction team should map the chain for every material revenue cohort.

The map should identify where Arabic content is created, translated, normalised, embedded, stored and reviewed. It should show the model and version, retrieved sources, data region, tool permissions, human approver and record retained. The buyer should reconcile diagrams with logs, configuration and sampled cases. Unmapped manual work often becomes recurring cost after acquisition.

Evidence should connect technical performance to commercial value. A high response score matters when it reduces handling time, improves acceptance, lowers complaint cost or creates revenue. The buyer should require cohort-level links among task volume, accepted outcome, invoice, renewal and cash. This protects the valuation from activity metrics that do not represent customer value.

Figure 1 Arabic workflow evidence chain



Proposed transaction map; actual gates should reflect customer consequence sector and jurisdiction.

4 Segment Arabic by dialect domain and consequence

Arabic is not a single operating cohort. Modern Standard Arabic is common in formal documents, while customers and employees can use Gulf, Saudi, Emirati, Egyptian, Levantine or other dialects in speech and informal text. Code-switching with English is frequent in commercial settings. The target should define the language populations it serves and avoid extrapolating from one benchmark or market.

Domain terminology can be more important than general fluency. Banking, insurance, healthcare, government, energy and legal workflows use specialised words, abbreviations and forms. A plausible answer can still be unacceptable when it uses the wrong obligation, product term or institutional role. Evaluation therefore needs native domain reviewers and consequences proportionate to the use case.

The buyer should build a matrix by language variety, input mode, domain, customer, outcome and consequence. It should sample high-volume and high-risk cases, including ambiguous instructions, mixed language, spelling variation and speech recognition. Segment-level evidence should drive product perimeter, integration priority and any price adjustment.

Table 2 Arabic localisation evidence matrix

Cohort	Evidence	Outcome measure	Transaction implication
Formal Modern Standard Arabic	source-grounded document and form tests	accuracy acceptance and correction	supports institutional document workflows
Gulf dialect conversation	native-speaker audio and text cohorts	intent capture escalation and satisfaction	supports service use only where proven
Arabic-English code-switching	real mixed-language interactions	task completion and exception rate	tests normal commercial behaviour
Sector terminology	domain corpus and qualified review	material error and approved output	determines regulated workflow perimeter
Customer-specific language	templates glossary and operating rules	first-pass acceptance and rework	supports retention and switching claims
High-consequence decision	trace approval rationale and appeal	authorised outcome and complaint rate	may require a human-controlled boundary

Proposed matrix; thresholds should be set for the specific workflow and customer consequence.

5 Verify data rights and provenance

Arabic capability can depend on scarce corpora, transcripts, documents, annotations and customer records. The buyer should identify the source, licence, consent, purpose, geography, retention and deletion rule for each material dataset. It should distinguish public availability from permission to train, fine-tune, evaluate or operate a commercial system.

Provenance should extend into derived assets. Embeddings, synthetic examples, preference data, evaluation sets and fine-tuned weights can preserve characteristics of source data. A target may be unable to reconstruct which sources influenced an acquired model or knowledge base. That uncertainty can restrict use, create remediation cost or weaken the reliability of representations.

The transaction team should sample the data inventory against contracts, consent records and deployed stores. It should test deletion, correction and tenant separation. UAE and Saudi data-protection frameworks include duties and transfer considerations that require current legal analysis [9-12,16-18]. The valuation should recognise only rights and operating uses that the buyer can substantiate.

6 Evaluate benchmarks and production outcomes

ArabicMMLU, ALLaM, Jais and related research show meaningful progress and continuing evaluation challenges [19-25]. General benchmarks are useful for comparison but cannot establish fitness for a regulated customer workflow. They may focus on Modern Standard Arabic, multiple-choice knowledge or a particular dataset and may not represent dialect, speech, current rules or tool use.

The target should provide datasets, ownership, sampling, instructions, graders, model versions, thresholds and failed cases. Results should be segmented by customer, language variety, task and consequence. The buyer should test for contamination, translated questions, synthetic-data dependence and repeated tuning against the test set. Independent native and domain review should be included for material claims.

Production measures should include first-pass acceptance, correction effort, escalation, false action, latency, unit cost, complaint and recurrence after model change. A release should be gated when a model, prompt, retrieval source or tool changes. The buyer should price the continuing evaluation function as operating cost rather than a one-time diligence exercise.

7 Establish regulated-sector access as a transferable asset

Regulated-sector access can arise from a licence, customer procurement, an approved vendor list, a sandbox, a partner, a deployed reference or staff credibility. These sources have different durability. The buyer should describe the exact right or relationship and determine whether it transfers on a change of control, product change, hosting change or subcontractor change.

Customer contracts should be reconciled with the operating reality. A master agreement can coexist with pilot-only use, narrow data access or a prohibition on autonomous decisions. Procurement registration does not prove active adoption. A reference logo does not prove permission to market the relationship. Revenue and renewal evidence should be tied to the precise deployed workflow.

Regulated access also creates obligations. Security reviews, model validation, audit rights, incident reporting, data location, business continuity and named personnel can increase cost. The acquirer should model the resources needed to maintain access after integration. Unsupported access claims should receive no synergy credit.

Table 3 Regulated-sector access and dependency tests

Access route	Evidence	Change event to test	Value risk
Direct customer contract	executed terms invoices and active use	change of control product and hosting	consent termination or narrowed use
Approved vendor status	current registration scope and reviews	new owner security and subcontractors	requalification delay
Regulatory sandbox	admission conditions testing record and exit	commercial scale and new use case	pilot status mistaken for approval
Channel partner	contract pipeline conversion and economics	exclusivity control and termination	nontransferable distribution
Public procurement	award scope acceptance and payment	assignment localisation and staffing	delayed renewal or compliance cost
Named relationship	meeting pipeline and succession evidence	founder departure and account transfer	access disappears after closing

Proposed diligence structure; legal and regulatory applicability requires current advice.

8 Read customer contracts at workflow level

The buyer should create a contract-to-workflow schedule. For each customer it should record permitted use, data categories, decision boundary, human review, model and subcontractor disclosure, hosting region, service level, audit, incident notice, intellectual property, assignment and termination. Broad product descriptions should be reconciled with statements of work and actual configuration.

Agentic systems can create actions that were absent when older contracts were signed. The target may have introduced tools, models or autonomous steps through product updates without refreshed customer consent. Diligence should compare contract dates, release notes, model changes and audit records. Material differences can require disclosure or remediation before integration.

Revenue quality depends on permission and acceptance. The team should test invoices, service credits, complaints, renewal discussions and cash. A buyer should segregate contracted recurring revenue from pilots, pass-through model charges, professional services and discretionary extensions. Synergy assumptions should use only customer cohorts whose rights and willingness support the proposed combined product.

9 Test deployment data location and operational control

GCC customers can require cloud, private-cloud, sovereign-cloud or on-premises deployment. Data location can vary across prompts, logs, embeddings, backups, support tools and model endpoints. A local application server does not establish local processing when inference, telemetry or support crosses borders. The buyer should map every material data flow and processor.

Operational control includes encryption, keys, identity, patching, model updates, incident response, backup and exit. The target should show who can access production, where privileges are recorded and how a customer can stop or isolate an agent. The buyer should test revocation, outage, model unavailability and restoration from backup.

Architecture affects margin. Multiple deployment patterns can create duplicated engineering, security and support. Customer-specific forks can delay releases and increase control risk. These recurring costs belong in sustainable EBITDA. The integration plan should preserve customer commitments while converging infrastructure only after permission and acceptance.

10 Verify identity delegation and human accountability

An agent can act for a customer, employee, organisation or another agent. The system should record the initiating principal, delegated authority, policy, tool, approval and resulting action. Shared credentials or generic service accounts weaken attribution. The buyer should reconstruct material decisions from raw records and compare them with policy.

Delegation should be bounded by purpose, resource, time, amount and consequence. A customer-service agent authorised to draft a response should not gain authority to alter a financial record when a tool fails. Human review should present the evidence, proposed action and uncertainty needed for a meaningful decision. Approval volume and override rates reveal whether the control is workable.

UAE financial-sector materials emphasise governing-body and senior-management accountability for AI outcomes [3-6]. NIST guidance addresses governance, measurement and management throughout the lifecycle [13-15]. The buyer should translate these principles into target-specific controls, people and cost. Responsibility cannot be transferred to the model vendor.

11 Apply a financial-services evidence standard

Financial-services use cases include customer communication, onboarding, fraud review, credit support, claims, compliance and operations. Material decisions can affect access to products or customer rights. The buyer should identify which tasks advise, recommend, decide or execute and what qualified person remains accountable.

Testing should examine explainability, fairness, data quality, model validation, complaint handling, monitoring and change control. The target should show performance by language and customer cohort. Arabic quality matters for disclosures and explanations as well as model prediction. A translated explanation that omits a decisive factor can create conduct and trust risk.

The commercial model should include validation staff, compliance, audit support, customer-specific testing and incident response. A financial institution's procurement and deployment can take longer than a technology sales forecast assumes. Pipeline probability should reflect completed control gates and budget authority, not meetings or demonstrations.

12 Apply a healthcare evidence standard

Healthcare workflows combine language, clinical context, sensitive data and professional responsibility. An Arabic-speaking patient can use dialect, incomplete history or culturally specific descriptions. A general conversational score cannot establish safe triage, coding, documentation or clinical support. The buyer should define the intended use and prohibited use precisely.

Evidence should include qualified clinical review, data rights, performance by population, escalation, adverse-event handling and human sign-off. The team should distinguish administrative productivity from clinical decision support. It should inspect product claims, customer training and incident records for divergence from the approved or contracted perimeter.

Healthcare value can come from reducing documentation burden, improving access or supporting bilingual service. It should be measured through accepted records, time saved, error, patient outcome where appropriate and cash. The acquisition model should fund continuing clinical governance and local language evaluation.

13 Apply a government and public-services standard

Government deployments can offer scale, reference value and recurring demand. They can also depend on procurement cycles, sovereign requirements, local content, security clearance and policy priorities. The

buyer should inspect award documents, acceptance certificates, payment records, renewal options and obligations attached to public data or infrastructure.

Arabic public-service workflows require accessibility, consistency and appeal. The system should identify when a citizen is interacting with AI, preserve the source of a response and route consequential matters to authorised staff. Dialect and literacy variation should be included in testing. The target should demonstrate how policy updates reach every deployed agent.

Public-sector access should be valued after mobilisation, support, localisation and receivable cost. A flagship pilot can have strategic relevance while contributing little sustainable earnings. The investment committee should separate reference value from contracted cash and assign evidence weights to each.

14 Apply an energy and infrastructure standard

Energy and infrastructure agents can support maintenance, field service, procurement, permitting, engineering documents and customer operations. The consequences of a wrong tool call or misunderstood instruction can be physical, financial or operational. The buyer should classify tasks by criticality and keep safety-critical authority within approved control systems and accountable roles.

Arabic capability can improve field communication and access to procedures, but technical terminology and mixed-language documentation require rigorous source grounding. Tests should include degraded connectivity, speech variation, outdated manuals, conflicting instructions and emergency escalation. Evidence should show that the agent refuses or escalates beyond its approved boundary.

Value should be tied to measurable downtime, work-order quality, maintenance productivity, inventory, safety or revenue. Integration with operational technology can create cyber and continuity risk. The buyer should assess separation, identity, logging and rollback before recognising synergy.

15 Test distribution and channel economics

Distribution can come from direct enterprise sales, systems integrators, cloud marketplaces, government partners or sector specialists. The target should show lead source, conversion, sales cycle, implementation, discount, partner share, renewal and cash by channel. Pipeline labels should be reconciled with customer evidence.

A channel can accelerate trust and implementation while controlling the customer relationship. The buyer should inspect exclusivity, territory, referral ownership, data access, pricing, support, change of control and termination. A partner-dependent business may need a valuation discount or closing condition when the agreement cannot be confirmed.

Arabic localisation can strengthen distribution only when it reduces a customer barrier. The team should compare win rates and deployment outcomes with and without the capability. Marketing statements, demonstration attendance and memoranda of understanding provide weak evidence of repeatable acquisition economics.

16 Measure customer acceptance and retention

Customer acceptance should be defined at workflow level. Measures can include first-pass completion, correction, escalation, user adoption, complaint, service credit, renewal and expansion. The buyer should compare the target's reported metrics with raw tickets, logs, invoices and customer interviews authorised for diligence.

Retention can depend on embedded integrations or high switching cost rather than satisfaction. The buyer should distinguish productive attachment from unresolved migration burden. It should identify customers

using a discontinued model, custom fork or founder-managed process. These cohorts can require disproportionate post-close investment.

The strongest evidence links accepted outcomes to recurring cash and provides a reproducible operating path. Customer concentration, public-sector receivables and discretionary pilots should be shown separately. Synergy should be deferred until the combined product is accepted by the relevant customer cohort.

17 Protect people knowledge and local authority

Arabic workflow quality often depends on product managers, linguists, domain specialists, customer-success staff and engineers whose knowledge is weakly documented. Regulated access can depend on named officers, security contacts or senior relationships. The buyer should map critical roles, succession, incentives and restrictive obligations.

Key-person analysis should trace every material workflow, dataset, customer and control to at least one accountable owner and a documented operating record. Bus-factor testing can ask another qualified team member to operate, investigate and update the system. Failure reveals integration cost and continuity risk.

Retention packages should align with knowledge transfer, customer continuity and control outcomes. Headcount should be included in sustainable economics. Replacing experienced local staff with central functions can weaken customer trust or control performance even when salary cost falls.

18 Identify model vendor and infrastructure dependence

The target may depend on a proprietary model, Arabic-centric model, cloud service, vector database, speech provider or orchestration framework. The buyer should map contractual rights, pricing, data use, region, service levels, change notice, exit and technical substitution. A model-independent marketing claim should be tested through comparable workflows.

Arabic model research shows different training approaches and benchmark performance [19-25]. No paper establishes performance for the target's customer tasks. The team should test alternative models using the same accepted-outcome criteria. It should measure quality, latency, cost, safety and support after substitution.

Dependence is not automatically negative when the supplier is reliable and the economics remain attractive. The transaction issue is whether dependence is understood, contractually supported and reflected in valuation. A contingency plan should cover price changes, regional unavailability, policy changes and model retirement.

19 Normalise reported EBITDA to sustainable earnings


Reported EBITDA can exclude the full cost of data rights, Arabic evaluation, local compliance, customer-specific deployment, partner support and key-person retention. The buyer should reconstruct cost by workflow and customer cohort. Capitalised development and pass-through model costs should be reviewed consistently with accounting policy.

The hypothetical target reports USD 12 million of EBITDA. Five continuing adjustments reduce sustainable EBITDA to USD 6.5 million: USD 1.4 million for data-rights and provenance operations, USD 1.1 million for local control and compliance, USD 1.3 million for regulated-sector sales and support, USD 0.9 million for evaluation and release assurance, and USD 0.8 million for critical-person and partner continuity. These are management assumptions, not observed benchmarks.

One-time remediation should be separated from recurring cost. Data cleanup, contract refresh and control implementation may require an integration reserve, while the ongoing staff and testing remain in

sustainable earnings. The investment committee should receive both views and avoid classifying necessary operating capability as temporary synergy spend.

Figure 2 Hypothetical reported to sustainable EBITDA bridge



Management assumptions in USD millions; figures are illustrative and are not a forecast or benchmark.

Table 4 Hypothetical sustainable EBITDA normalisation

Item	USD million	Treatment	Evidence required
Reported EBITDA	12.0	starting point	audited and management accounts
Data rights and provenance	(1.4)	recurring	inventory licences deletion and review team
Local control and compliance	(1.1)	recurring	accountable roles validation audit and reporting
Regulated-sector sales and support	(1.3)	recurring	procurement implementation and service cohorts
Evaluation and release assurance	(0.9)	recurring	datasets graders regression and release gates
Critical people and partners	(0.8)	recurring	retention succession and channel continuity
Sustainable EBITDA	6.5	valuation earnings	reconciled target operating model

Management assumptions in USD millions; target-specific financial accounting tax and operational diligence is required.

20 Build evidence-weighted synergies


Synergy should begin with a specific combined action. Buyer distribution can increase sales only when the target product passes buyer and customer controls. Arabic workflow capability can support cross-sell only when the relevant cohort accepts it. Shared technology can reduce cost only after migration preserves service and rights. Faster market access needs contractual and regulatory evidence.

The hypothetical gross annual synergy is USD 10 million: USD 3.4 million from buyer distribution, USD 2.6 million from Arabic workflow cross-sell, USD 2 million from shared technology and USD 2 million from faster regulated-sector access. Continuing controls of USD 2.1 million, revalidation and migration of

USD 1.6 million, retention of USD 1 million and client remediation of USD 1.2 million reduce annual net synergy to USD 4.1 million. Every amount is a management assumption.

Each synergy should have an owner, baseline, action, customer cohort, cost, timing and cash measure. Evidence weights should reflect completed tests and permissions. The buyer should report gross, cost and net amounts separately and prevent the same benefit from appearing in both base earnings and synergy.

Figure 3 Hypothetical gross to net annual synergy bridge



Management assumptions in USD millions; figures are illustrative and are not a forecast or benchmark.


21 Value the asset through evidence and risk

IFRS 13 describes fair value as an exit-price measurement using market-participant assumptions [26]. IFRS 3 addresses business combinations and acquired assets and liabilities [27]. IAS 36 and IAS 38 address impairment and intangible assets [28-29]. These accounting frameworks do not determine transaction price, but they reinforce the need to identify rights, useful lives, risk and supportable assumptions.

The illustrative bridge applies eleven times to USD 6.5 million of sustainable EBITDA, producing USD 71.5 million. It adds USD 15 million for the present value of evidence-weighted synergy, deducts USD 18 million for integration and control investment and deducts USD 9.5 million for market-access and concentration risk. The resulting illustrative enterprise value is USD 59 million. The multiple, cash flows and deductions are management assumptions.

Sensitivity should vary sustainable earnings, customer retention, partner continuity, evaluation cost, model price and synergy timing. Scenario probability should depend on evidence gates. A valuation that depends on nontransferable access or untested localisation should be restructured through deferred consideration, earn-outs or closing conditions.

Figure 4 Hypothetical evidence-weighted enterprise value bridge



Management assumptions in USD millions; the illustration is not a valuation conclusion or recommendation.

Table 5 Hypothetical valuation bridge

Component	USD million	Basis	Principal sensitivity
Sustainable EBITDA value	71.5	6.5 at an assumed 11 times	recurring cost and retained revenue
Evidence-weighted synergy present value	15.0	discounted implemented benefits	timing acceptance and cash conversion
Integration and control investment	(18.0)	systems data governance and migration	remediation scope and deployment variety
Access and concentration risk	(9.5)	customer partner and permission exposure	transferability and renewal
Illustrative enterprise value	59.0	resulting bridge	all underlying assumptions

Management assumptions in USD millions; actual valuation requires target-specific evidence and professional advice.

22 Review competition and market-access risk

Acquiring a material route to customers, data or distribution can create competition questions. UAE competition materials address economic concentration and notification thresholds, while Saudi guidance addresses control and review of economic concentrations [30-33]. Application depends on current facts, market definition, sector rules and jurisdictional advice.

The buyer should assess whether the combination controls an important dataset, channel, platform or procurement route. It should examine exclusivity, bundling, self-preference, interoperability, access terms and information gained from partners or customers. Localisation can be pro-competitive when it expands access, yet the supporting distribution or data arrangement can still merit review.

Timing and remedy risk belong in the transaction plan. Regulatory conditions, customer consents and sector approvals can interact. The valuation should include delay, separation or behavioural costs where supported. A filing threshold is not a substitute for a substantive competition assessment.

23 Select transaction protections that follow the evidence

Diligence findings should change price, structure, conditions and covenants. Confirmed rights, transferable customer relationships and reproducible accepted workflows can support base value.

Unproven access or localisation can support deferred consideration. Material data, security or authority gaps can require remediation before closing or exclusion from the perimeter.

Representations can address ownership, licences, data provenance, privacy, model and open-source dependencies, customer permissions, incidents, compliance and partner terms. Schedules should describe the deployed reality. Technical claims need objective definitions and version references. Legal drafting requires current qualified advice.

Earn-outs should use retained recurring revenue, accepted workflow cohorts, gross margin after full control cost and collected cash. Raw prompts, users or demonstration volume are vulnerable to manipulation and weakly connected to value. Audit rights, customer attribution and treatment of buyer decisions need clear design.

Table 6 Transaction response by finding

Finding	Value effect	Potential deal response	Post-close measure
Verified Arabic workflow acceptance	supports revenue and cross-sell	base value or evidence-weighted synergy	accepted cohorts retention and cash
Unclear data rights	restricts product use and model value	remediation condition escrow or exclusion	verified provenance and permitted operation
Nontransferable access	weakens growth and retention	consent condition deferred value or price change	confirmed customer and partner continuity
Weak human accountability	raises conduct and control cost	control remediation and perimeter limit	authorised actions exceptions and complaints
Key-person dependence	threatens operation and relationships	retention succession and deferred consideration	knowledge transfer and customer continuity
Model or cloud dependence	affects margin and resilience	price protection exit plan or reserve	substitution outcome cost and service

Proposed framework; actual instruments require current legal tax accounting regulatory and financial advice.

24 Build a decision-led evidence room

The evidence room should connect every material investment claim to source records, tests, owners and findings. Technical materials should include architecture, data lineage, model and prompt versions, tools, identity, evaluations, releases, incidents, deployment and recovery. Commercial materials should include contracts, procurement, usage, acceptance, invoices, credits, renewals, pipeline and cash.

Arabic evidence should include cohort definitions, native and domain reviewers, failed cases, correction effort and change history. Regulated-sector evidence should include the precise access route, scope, conditions and transfer implications. Finance should reconcile model, cloud, evaluation, compliance, implementation and support cost to customer cohorts.

Access to sensitive information should be controlled. Customer data, credentials and protected records should remain in approved review environments. The room should preserve dates, hashes or version identifiers so conclusions can be reproduced. A board summary should distinguish verified facts, management assumptions and matters requiring advice.

Table 7 Evidence gates for releasing acquisition value

Gate	Minimum evidence	Decision	Measure after release
Language and domain	representative cohorts and qualified acceptance	approve workflow perimeter	first-pass acceptance correction and escalation

Gate	Minimum evidence	Decision	Measure after release
Rights	traceable data licences consent and deletion	permit continued use	exceptions deletion and audit findings
Authority	identity delegation human review and recovery	approve consequential action	authorised actions and incidents
Access	transferable contracts permissions and partners	include revenue or synergy	retention renewal and sales conversion
Economics	full local control delivery and model cost	set sustainable earnings	contribution margin and cash
Integration	migrated cohort with customer acceptance	release synergy	recurring net cash and residual risk

Proposed governance; thresholds should be approved for the transaction and customer consequences.


25 Execute the first hundred days through controlled cohorts

Days one to fifteen should preserve code, data inventories, model versions, prompts, evaluations, logs, contracts, approvals, incidents and customer commitments. The buyer should confirm access controls and stop undocumented changes to consequential workflows while allowing necessary security responses.

Days sixteen to thirty-five should reconcile architecture, rights and customer schedules. The team should run representative Arabic and bilingual cohorts, test authority and identify customer or regulator engagement requirements. Finance should establish the cost baseline. Days thirty-six to sixty-five should remediate priority controls, revalidate affected workflows and seek required customer permissions.

Days sixty-six to one hundred should migrate approved cohorts, activate distribution and release synergy only where acceptance and cash evidence exist. Governance should report language outcome, customer status, control, economics and incidents together. Unproven cohorts should remain ring-fenced or deferred.

Figure 5 First hundred days for GCC agentic AI integration



Proposed sequence; actual timing should reflect customers regulation people security and systems.

26 Compare target archetypes

An Arabic model developer can own valuable research, weights and talent while lacking repeatable enterprise distribution. Its value depends on rights, model performance, compute economics and a route to

accepted workflows. A horizontal agent platform can have strong tools and orchestration while depending on third parties for Arabic and domain accuracy.

A vertical regulated-sector business can show accepted outcomes and customer access. It can also carry concentration, bespoke implementation and key-person risk. A systems integrator with agentic capability can possess relationships and delivery capacity while owning limited reusable technology. A channel-led reseller can grow quickly while depending on supplier and partner terms.

The buyer should value each archetype through its own evidence. Model benchmarks, software recurring revenue, professional-services backlog and customer access cannot share one multiple without adjustment. The preferred target is the one whose assets complement the buyer's weaknesses and remain transferable after closing.

27 Limitations and conclusion

AI technology, Arabic model performance and GCC regulation continue to evolve. Public sources reviewed for this paper describe the position available on the publication date. Customer contracts, data rights, legal duties and regulatory expectations require current target-specific verification. Hypothetical numbers illustrate method and provide no forecast, benchmark, investment recommendation or valuation conclusion.

Arabic workflow capability creates acquisition value when it produces accepted outcomes across relevant dialects and domains with lawful data, bounded authority and reproducible controls. Regulated-sector access creates value when it rests on transferable permissions, durable customer relationships and an operating model that can maintain its obligations after closing.

The transaction model should use sustainable EBITDA after language, governance, delivery and people cost. Synergies should be weighted by customer and implementation evidence. Deal protections should defer uncertain value, and the first hundred days should preserve evidence, revalidate cohorts and release integration value only after outcomes and cash are demonstrated.

The strongest acquisition case is measurable. The combined business can explain what the agent understood, which source and policy it used, who authorised the action, why the customer accepted it and how the result converted into recurring cash. That evidence makes localisation and market access investable rather than rhetorical.

Frequently asked questions

What makes Arabic capability valuable in an agentic AI acquisition

Value comes from accepted customer workflows across relevant language varieties, domains and consequences. The buyer needs evidence of data rights, reproducible outcomes, bounded authority, continuing cost and recurring cash.

Is an Arabic user interface sufficient evidence of localisation

No. The buyer should test dialect, terminology, institutional process, human approval, exception handling and customer acceptance across the complete workflow. Interface translation proves only a narrow part of the operating system.

How should a buyer test Arabic language-model claims

Use representative customer cohorts with native and domain-qualified review. Record first-pass acceptance, correction, escalation, latency, cost and downstream outcome. Include mixed language, dialect, ambiguous instructions and model changes.

What proves regulated-sector access

The evidence can include executed contracts, approved vendor status, procurement awards, active deployment, customer acceptance and payment. The buyer should confirm scope, conditions, change-of-control treatment and continuing obligations.

Which finding can block an acquisition

A material workflow can become unsuitable when the target lacks data rights, accountable decision authority, transferable customer permission or a secure and recoverable operating path. The board should define no-go conditions before diligence.

How should recurring localisation cost be treated

Native and domain review, evaluation, governance, customer-specific deployment, compliance and support belong in sustainable operating economics. One-time remediation should be separately identified and funded.

How should Arabic workflow synergies be valued

Tie each benefit to a customer cohort, owner, action, cost, timing and accepted outcome. Apply an evidence weight and present value. Release the benefit after implementation produces recurring net cash.

What should the buyer do during the first hundred days

Preserve rights and evidence, reconcile architecture and customer permissions, revalidate representative cohorts, remediate priority controls, migrate approved workflows and release value through outcome and cash gates.

References

- [1] UAE Artificial Intelligence Office. UAE National Strategy for Artificial Intelligence 2031. <https://ai.gov.ae/wp-content/uploads/2023/05/AI-Report-EN-v4.pdf>
- [2] UAE Artificial Intelligence Office. UAE Charter for the Development and Use of Artificial Intelligence. <https://ai.gov.ae/>
- [3] Central Bank of the UAE. Big Data Analytics and Artificial Intelligence. <https://rulebook.centralbank.ae/en/rulebook/big-data-analytics-and-artificial-intelligence-ai>
- [4] Central Bank of the UAE. Guidelines for Financial Institutions Adopting Enabling Technologies. https://rulebook.centralbank.ae/sites/default/files/en_net_file_store/CBUAE_EN_2413_VER1.pdf
- [5] Central Bank of the UAE. Guidance Note on Consumer Protection and Responsible Adoption and Use of AI and ML. <https://rulebook.centralbank.ae/en/entiresection/6960>
- [6] Dubai Financial Services Authority. UAE Regulatory Authorities Jointly Issue Guidelines for Financial Institutions Adopting Enabling Technologies. <https://www.dfsa.ae/news/uae-regulatory-authorities-jointly-issue-guidelines-financial-institutions-adopting-enabling-technologies>
- [7] Saudi Data and Artificial Intelligence Authority. AI Ethics Principles. <https://dgp.sdaia.gov.sa/wps/wcm/connect/4c56ed1c-1b82-447d-ac29-638f5f99c12e/ai-principles-EN.pdf?MOD=AJPERES>
- [8] Saudi Data and Artificial Intelligence Authority. AI Adoption Framework. <https://sdaia.gov.sa/en/SDAIA/about/Files/AIAAdoptionFramework.pdf>
- [9] Saudi Data and Artificial Intelligence Authority. Personal Data Protection Law Knowledge Center. <https://dgp.sdaia.gov.sa/wps/portal/pdp/knowledgecenter/>
- [10] Saudi Data and Artificial Intelligence Authority. Regulation on Personal Data Transfer Outside the Kingdom. <https://dgp.sdaia.gov.sa/wps/portal/pdp/knowledgecenter/>
- [11] Saudi Data and Artificial Intelligence Authority. Guidelines for Binding Common Rules for Personal Data Transfer. <https://dgp.sdaia.gov.sa/wps/portal/pdp/knowledgecenter/details/GuidelinesforBindingCommonRules/>
- [12] Saudi Data and Artificial Intelligence Authority. Laws and Regulations. <https://sdaia.gov.sa/en/SDAIA/about/Pages/RegulationsAndPolicies.aspx>
- [13] National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework. <https://www.nist.gov/itl/ai-risk-management-framework>
- [14] National Institute of Standards and Technology. AI Risk Management Framework Playbook. <https://www.nist.gov/itl/ai-risk-management-framework/nist-ai-rmf-playbook>

- [15] National Institute of Standards and Technology. Generative Artificial Intelligence Profile. <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf>
- [16] UAE Government. Federal Decree Law 45 of 2021 Regarding the Protection of Personal Data. <https://u.ae/en/about-the-uae/digital-uae/data/data-protection-laws>
- [17] Dubai International Financial Centre. Data Protection Law DIFC Law 5 of 2020. <https://www.difc.com/business/registrars-and-commissioners/commissioner-of-data-protection>
- [18] Abu Dhabi Global Market. Data Protection Regulations 2021. <https://www.adgm.com/operating-in-adgm/office-of-data-protection>
- [19] Bari et al. ALLaM Large Language Models for Arabic and English. <https://arxiv.org/abs/2407.15390>
- [20] Sengupta et al. Jais and Jais-chat Arabic-Centric Foundation and Instruction-Tuned Open Generative Large Language Models. <https://arxiv.org/abs/2308.16149>
- [21] Anwar et al. Jais 2 A Family of Arabic-Centric Open Large Language Models. <https://arxiv.org/abs/2608.13580>
- [22] Koto et al. ArabicMMLU Assessing Massive Multitask Language Understanding in Arabic. <https://arxiv.org/abs/2402.12840>
- [23] Stanford Center for Research on Foundation Models. Holistic Evaluation of Language Models Arabic. <https://crfm.stanford.edu/helm/arabic/latest/>
- [24] Habash et al. ArabLegalEval A Multitask Benchmark for Assessing Arabic Legal Knowledge in Large Language Models. <https://aclanthology.org/2024.arabicnlp-1.20/>
- [25] Technology Innovation Institute. Arabic Large Language Model Benchmarks. <https://github.com/tiiuae/Arabic-LLM-Benchmarks>
- [26] IFRS Foundation. IFRS 13 Fair Value Measurement. <https://www.ifrs.org/issued-standards/list-of-standards/ifrs-13-fair-value-measurement/>
- [27] IFRS Foundation. IFRS 3 Business Combinations. <https://www.ifrs.org/issued-standards/list-of-standards/ifrs-3-business-combinations/>
- [28] IFRS Foundation. IAS 36 Impairment of Assets. <https://www.ifrs.org/issued-standards/list-of-standards/ias-36-impairment-of-assets/>
- [29] IFRS Foundation. IAS 38 Intangible Assets. <https://www.ifrs.org/content/dam/ifrs/publications/pdf-standards/english/2021/issued/part-a/ias-38-intangible-assets.pdf?bypass=on>
- [30] UAE Ministry of Economy. Regulation of Competition. <https://www.moec.gov.ae/regulation-of-competition>
- [31] UAE Ministry of Economy. Cabinet Resolution 3 of 2025 Thresholds for Economic Concentration. <https://www.moec.gov.ae/en/-/ministry-of-economy-reviews-cabinet-resolution-on-thresholds-and-controls-associated-with-competition-law>
- [32] UAE Ministry of Economy. Cabinet Decision 3 of 2025. <https://www.moec.gov.ae/documents/20121/0/Cabinet%2BDecision%2BNo.%2B%283%29%2Bof%2B2025.pdf>
- [33] Saudi General Authority for Competition. Economic Concentration Review Guidelines. <https://gacbp.gac.gov.sa/cms/b9376edc-79a1-4573-a36d-4f3effaba838.pdf>
- [34] Organisation for Economic Co-operation and Development. OECD AI Principles. <https://oecd.ai/en/ai-principles>
- [35] International Organization for Standardization. ISO IEC 42001 Artificial Intelligence Management System. <https://www.iso.org/standard/81230.html>
- [36] International Organization for Standardization. ISO IEC 23894 Artificial Intelligence Risk Management. <https://www.iso.org/standard/77304.html>
- [37] MITRE. Adversarial Threat Landscape for Artificial Intelligence Systems. <https://atlas.mitre.org/>
- [38] OWASP Foundation. OWASP Top 10 for Large Language Model Applications. <https://genai.owasp.org/llm-top-10/>
- [39] Internet Engineering Task Force. RFC 8707 Resource Indicators for OAuth 2.0. <https://www.ietf.org/ietf-ftp/rfc/rfc8707.pdf>
- [40] Internet Engineering Task Force. RFC 9728 OAuth 2.0 Protected Resource Metadata. <https://www.ietf.org/rfc/rfc9728.pdf>
- [41] OpenTelemetry. Generative AI Attributes. <https://opentelemetry.io/docs/specs/semconv/registry/attributes/gen-ai/>
- [42] OpenTelemetry. Semantic Conventions. <https://opentelemetry.io/docs/specs/semconv/>
- [43] Cloud Native Computing Foundation. CloudEvents Specification. <https://github.com/cloudevents/spec/blob/main/cloudevents/spec.md>
- [44] Model Context Protocol. Authorization. <https://modelcontextprotocol.io/specification/draft/basic/authorization>
- [45] Agent2Agent Project. Protocol Specification. <https://a2a-protocol.org/latest/specification/>
- [46] United States Department of Justice and Federal Trade Commission. 2023 Merger Guidelines. <https://www.justice.gov/atr/2023-merger-guidelines>
- [47] United Kingdom Competition and Markets Authority. Merger Assessment Guidelines. <https://www.gov.uk/government/publications/merger-assessment-guidelines>
- [48] European Union. Regulation 2024 1689 Artificial Intelligence Act. <https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX%3A32024R1689>
- [49] European Commission. AI Act. <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>
- [50] European Commission. Data Act Explained. <https://digital-strategy.ec.europa.eu/en/factpages/data-act-explained>

About the Author

Chennakeshav (CK) is a corporate finance and investment banking executive with 25+ years of global experience in deal origination, structuring and execution across M&A, growth capital and corporate strategy. He has led value-creation mandates for founders, corporates and funds — bridging the boardroom view to hands-on execution and close.

His career spans Morgan Stanley, HSBC, Lloyds Banking Group, EWEC, ADQ portfolio companies and Emirates Growth Fund, across TMT, real estate, fintech, deeptech, cleantech, infrastructure and energy. He has partnered with C-suite leaders, private equity and venture funds, sovereign wealth funds and family offices to finance complex fund raises and scale-up ventures, and has led M&A due diligence, post-merger integration and business-transformation initiatives to create value.

At Matchpoint Partners he is Managing Partner, leading the firm's corporate finance, M&A and capital-raising practice. He holds an MBA from London Business School, an engineering degree from VTU and a Master of Laws (LLM, in progress) from UCL London.

An active start-up mentor, CK mentors at Techstars, DIFC FinTech Hive, Startup Grind, Founder Institute and IN5, serves as Entrepreneur Mentor in Residence (EMiR) at London Business School, and judges the Entrepreneurship World Cup.

<https://www.linkedin.com/in/ckadya/>

<https://www.matchpoint-partners.com/team/ck-adya.html>

This paper is part of a continuing series on the structure of private and alternative markets. The views expressed are the author's own. The paper is for information only, describes market structure in general terms, and does not constitute investment, legal, tax or regulatory advice or a recommendation in respect of any security, vehicle or counterparty.