Multi Agent Platform M&A Interoperability as the Defensible Asset

A Transaction Framework for Orchestration Portability Identity Observability and Sustainable Platform Value

Authored by Chennakeshav Adya

Independent Researcher

September 2026

Abstract

Multi agent platforms promise to coordinate specialised artificial intelligence agents across models, tools, data sources and organisations. Their acquisition cases often treat protocol support, connector breadth or developer adoption as evidence of a durable platform advantage. The economic asset is narrower and more demanding: a controlled ability to move tasks, context, authority, evidence and recovery state across components while preserving customer outcomes. A buyer that acquires nominal interoperability can inherit brittle adapters, undocumented semantics, privileged credentials, opaque failures and dependence on one model or cloud.

This paper develops a transaction framework for deciding whether interoperability is a defensible asset in multi agent platform M&A. It separates protocol compatibility from operating interoperability and examines discovery, capability declaration, task lifecycle, artifacts, context, memory, tools, identity, delegation, human approval, observability, conformance, recovery and switching. It links the technical evidence to customer retention, gross margin, sustainable EBITDA, synergies, valuation, competition risk, transaction protections and the first hundred days.

The framework draws on the NIST AI Agent Standards Initiative, Agent2Agent protocol materials, the Model Context Protocol, OAuth and IETF authorization standards, OpenTelemetry semantic conventions, CloudEvents, the NIST AI Risk Management Framework, the EU Data Act, the EU AI Act, competition authority materials and accounting standards [1-33]. These sources describe evolving standards and legal requirements. They do not establish the quality, market position or value of a particular target.

A hypothetical target illustrates the method. Reported annual revenue is USD 54 million and reported EBITDA is USD 13.5 million. Normalising connector maintenance, identity and security, observability and evaluation, protocol migration and retention reduces sustainable EBITDA to USD 7.5 million. Gross annual synergy of USD 9.4 million becomes USD 3.6 million after continuing control, migration, retention and compatibility costs. A separate illustrative valuation bridge starts with ten times sustainable EBITDA, adds evidence weighted synergy present value and deducts integration and platform risk, producing USD 62 million. Every amount is a management assumption for method illustration, not an observation, forecast, benchmark or valuation conclusion.

The analysis finds that durable value comes from reproducible outcomes, bounded authority, portable evidence and switching that works under production conditions. Open specifications can reduce dependency while increasing the importance of implementation quality, governance and network participation. The buyer should release value only when representative workflows pass conformance, security, recovery, customer acceptance and cash tests.

JEL Classification: G24, G34, L13, L86, M15, O33

Keywords: multi agent platforms, agentic AI, interoperability, M&A, orchestration, identity, observability, platform valuation, due diligence

Introduction

Multi agent platforms coordinate software agents that can plan, communicate, call tools, exchange artifacts and act on enterprise systems. Their strategic appeal comes from reuse. A platform can connect many models and applications, allocate work to specialised agents and make complex workflows easier to assemble. The same architecture creates transaction risk because value can depend on undocumented adapters, shared memory, privileged credentials, hidden state and the operational knowledge of a small engineering team.

Interoperability has become a central design objective. NIST launched an AI Agent Standards Initiative around industry standards, open protocols and research in identity, authorization, security and evaluation [1-2]. The Agent2Agent protocol defines concepts for agent discovery, messages, tasks, artifacts and capability declaration [5-8]. The Model Context Protocol defines ways for applications to expose tools, resources and prompts, with continuing work on authorization, tasks and agent identity [9-15]. OpenTelemetry and CloudEvents provide complementary conventions for observable operations and portable event envelopes [20-23].

Protocol adoption does not answer an acquisition question. A target can advertise support while retaining proprietary context formats, tool semantics, policy engines, recovery logic and commercial restrictions. A technically open interface can still be expensive to operate, difficult to switch and unsafe to combine. A defensible asset therefore requires verified operating outcomes and an economic system around them.

This paper provides a diligence and valuation method for strategic buyers, financial sponsors, boards, lenders and management teams. It treats interoperability as an evidence chain from declared capability to accepted customer outcome and cash. It also recognises that standards are evolving. The transaction structure should preserve value when protocols, models, regulation and customer requirements change.

1 Define the acquisition decision

The investment committee should begin with a precise decision statement. It should identify the target entities, products, protocols, customer workflows, deployment modes, data rights, people and intellectual property being acquired. It should then state the proposed value mechanism: accelerated distribution, lower integration cost, higher customer retention, a broader developer network, stronger security, model neutrality or access to an enterprise control plane.

Each mechanism requires different evidence. Connector count can support distribution only when connectors are maintained, used and commercially relevant. Model neutrality requires comparable outcomes across supported models. A developer network needs verified active participation and contribution. A control plane needs identity, policy, audit and recovery that operate across the acquired perimeter.

The committee should define failure conditions before diligence. Examples include context that cannot move between agents, tools that behave differently behind the same schema, shared credentials that prevent attribution, recovery that depends on inaccessible vendor state, or customers whose contracts restrict migration. These conditions should become tests, price adjustments or closing requirements.

The acquisition perimeter should include protocols, software development kits, connectors, schemas, registries, orchestration engines, memory stores, policy systems, evaluation sets, telemetry, credentials, infrastructure, contracts and community governance. The buyer should distinguish owned components from open source, licensed and customer specific components. Ownership alone does not establish defensibility; transferability and continued operability are the relevant transaction facts.

2 Separate protocol support from operating interoperability

Protocol support shows that two components can exchange a conforming message. Operating interoperability shows that a complete workflow can discover a capability, authenticate, delegate authority, exchange context, execute a task, produce an artifact, record evidence, handle failure and resume without material loss of meaning or control. The second standard is the relevant acquisition test.

Four layers should be distinguished. Syntactic interoperability concerns formats and transport. Semantic interoperability concerns the shared meaning of tasks, fields, tools and artifacts. Operational interoperability concerns identity, policy, observability, error handling and service levels. Commercial interoperability concerns rights, pricing, support, switching and customer acceptance. Weakness in any layer can prevent the expected platform outcome.

The diligence team should avoid a single yes or no classification. It should test representative workflows across heterogeneous models, clouds, tools and counterparties. Each test should record configuration, versions, inputs, outputs, authority, exceptions, latency, cost and recovery. A successful demonstration on a prepared path provides limited evidence about production variance.

Open specifications can reduce implementation ambiguity while leaving substantial room for differentiation. Agent discovery, planning quality, policy, memory, evaluation and operational support may remain proprietary. The buyer should identify which proprietary elements create customer value and which create avoidable dependence.

3 Map the multi agent control plane

The control plane determines how agents are registered, discovered, assigned, authorised, observed and stopped. A transaction map should show every boundary among the orchestration service, agent runtimes, model providers, tools, data stores, identity providers, approval systems and customer environments. It should identify where state and authority persist.

The map should trace at least three workflow classes: routine internal work, customer facing work and a consequential workflow with approval or regulatory significance. Each trace should begin with a human or system instruction and end with an accepted artifact, recorded action or collected cash. The map should show alternative and failure paths.

Central control can improve consistency and provide a valuable system of record. It can also create a concentration point for outage, compromise or commercial leverage. Distributed control can improve resilience while increasing semantic drift and evidence reconciliation. The target should explain the chosen architecture and demonstrate its operation.

The buyer should reconcile architecture diagrams with source code, deployment configuration, telemetry and customer implementation records. Differences often reveal manual procedures, legacy components or customer specific forks. Those differences can become continuing cost after acquisition.

Table 1 Multi agent platform diligence perimeter

Component	Required evidence	Decision question	Principal risk
Discovery	agent cards registries versions and uptime	can capabilities be found and trusted	stale or self asserted capability
Tasks and artifacts	schemas lifecycle logs and accepted outputs	can work move without losing meaning	syntactic exchange without useful outcome
Context and memory	provenance retention export and deletion	can state move lawfully and accurately	hidden lock in or contamination

Component	Required evidence	Decision question	Principal risk
Tools	contracts permissions tests and rollback	can actions be reproduced and bounded	excessive authority or semantic drift
Identity	principals tokens delegation and revocation	who acted for whom and with what authority	shared credentials and weak attribution
Operations	traces evaluations incidents and recovery	can the platform prove and restore service	opaque failure and recurring support cost

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

4 Test agent discovery and capability declaration

Agent2Agent materials use an Agent Card to describe identity, endpoint, capabilities, skills and authentication requirements [5-8]. This creates a useful diligence starting point. The buyer should test whether declarations are complete, current, machine readable and bound to a trustworthy operator. A registry of stale cards can create more integration risk than value.

Capability names should map to measurable inputs, outputs, constraints and service levels. Broad statements such as research, analyse or transact are insufficient. The target should show test cases, schema versions, supported media, latency ranges, error states, geographic limits and approval requirements. Capability changes should trigger versioning and compatibility procedures.

Discovery also has a commercial dimension. The platform may rank or route agents according to sponsorship, internal ownership, cost or performance. Diligence should identify these rules and determine whether customers or partners can understand them. Undisclosed preference can weaken trust and raise competition concerns when a platform controls access to complementary services [34-38].

The buyer should examine spoofing, substitution and downgrade risks. A trusted capability declaration should be connected to verifiable identity, signed software or controlled deployment where appropriate. Tests should include revoked agents, changed endpoints, unsupported versions and conflicting descriptions.

5 Test task lifecycle and artifact portability

A task lifecycle should distinguish submitted, working, input required, completed, failed and cancelled states. The platform should preserve a stable task identifier, context, messages and artifacts through these transitions. Diligence should test whether status is consistent across transports and whether the customer can export evidence of the completed work.

Artifacts are the economically relevant output. They can include documents, code, structured data, decisions or executed system changes. Portability requires content, metadata, provenance, schema and acceptance information. A file alone may be unusable if the buyer cannot reconstruct the sources, permissions and transformations that produced it.

The team should test long running work, interruptions, partial results, duplicate delivery and cancellation. Idempotency matters when an agent can make payments, create records or alter infrastructure. Replayed messages should not cause repeated consequential actions. Compensation or rollback should be defined when an action cannot be reversed.

Commercial evidence should connect artifacts to customer acceptance, billing and renewal. High task volume can have limited value when outputs are experimental, rejected or included without separate revenue. The target should show accepted workflows by customer, version and deployment mode.

6 Test context and memory portability

Context includes instructions, conversation history, retrieved data, policies, tool results and intermediate reasoning available to a workflow. Memory includes persisted facts, preferences, embeddings, summaries and state reused across sessions. Both can create high switching cost because their semantics may depend on proprietary storage and retrieval behaviour.

The buyer should identify every context and memory store, its owner, retention rule, region, encryption, access policy and export format. It should trace provenance to source data and determine whether customers have rights to move, delete or reuse it. The EU Data Act places emphasis on cloud switching, open interfaces and machine readable export, subject to its detailed scope and application [24-26].

Portability tests should move representative state to an alternative runtime and compare outcomes. Exact model outputs are rarely expected. The test should examine completeness, grounded facts, policy application, customer acceptance and error rates. It should also test selective deletion and tenant separation.

Memory can preserve incorrect or sensitive information after the originating record changes. The target should demonstrate correction, expiry and conflict resolution. A buyer should price the remediation of undocumented memory stores and embeddings that cannot be traced to permitted sources.

7 Examine tool access and action semantics

MCP defines tools as one of its core primitives, alongside resources and prompts [9-15]. A schema can describe a tool call while leaving business semantics, side effects and authority outside the interface. Two tools with similar names may create different records, validation, pricing or rollback behaviour.

The diligence team should catalogue tools by consequence. Read only retrieval, draft preparation, record creation, financial execution and infrastructure control require different safeguards. Each tool should have an owner, version, permitted principals, input validation, output validation, rate limits, monitoring and a failure response.

Tests should include malformed inputs, stale schemas, unavailable dependencies, duplicated requests and partial completion. The team should inspect whether an agent can choose a more privileged tool when a lower privilege option fails. Tool descriptions and examples can become an attack surface if untrusted content influences selection or parameters [16-19].

Tool portability requires more than connector replacement. The new tool must preserve the accepted business outcome and evidence. The buyer should measure the effort to substitute a tool and the resulting change in latency, unit cost, failure and customer acceptance.

8 Verify identity authorization and delegation

An agent can act for a user, service, organisation or another agent. The platform should represent these principals and their delegation chain. NIST's work on software and AI agent identity and authorization highlights OAuth extensions, policy based access control and token based mechanisms as relevant building blocks [3-4]. IETF resource indicators and protected resource metadata help bind authorization to intended resources [17-18].

Diligence should reconstruct who initiated each consequential action, what authority was granted, which policy applied and when the authority expired or was revoked. Shared API keys weaken attribution and can create a hidden integration project. Long lived credentials increase the consequence of compromise.

Delegation should be bounded by purpose, resource, action, time and amount where relevant. An agent that can read a customer record does not automatically need authority to modify or transmit it. Step up approval should be available when consequence rises or evidence is incomplete.

The buyer should test revocation, employee departure, customer termination, incident containment and cross tenant isolation. Authorization documentation should match deployed policy. A platform whose commercial proposition relies on autonomous execution requires especially strong proof of bounded authority.

Table 2 Identity and delegation control tests

Control	Evidence	Test	Failure implication
Principal identity	registered user service and agent identities	trace one action to initiating principal	weak attribution and dispute risk
Delegation	scope purpose resource and expiry	exceed delegated amount or resource	excessive agency and control failure
Token audience	issuer audience and resource metadata	replay token at another resource	credential misuse and lateral movement
Revocation	policy update token invalidation and logs	revoke during active task	continuing unauthorised action
Human approval	named approver evidence and threshold	trigger consequential action	uncontrolled execution or delay
Tenant isolation	keys stores logs and policies by tenant	attempt cross tenant access	confidentiality and platform risk

Proposed test set; actual controls should reflect consequence jurisdiction and customer commitments.

9 Assess human approval and accountable authority

Human approval should be designed as a control activity rather than a generic button. The approver needs the proposed action, evidence, uncertainty, affected resource and available alternatives. The system should record who approved, what was approved and whether the executed action matched the approval.

Approval fatigue can weaken a platform that routes too many low quality exceptions to senior staff. Diligence should measure approval volume, response time, override, rejection and later incident by workflow. A low rejection rate can indicate stable quality or routine confirmation; sampling is needed to distinguish them.

The EU AI Act includes obligations concerning documentation, logging, quality management and human oversight for relevant systems and roles [27-28]. Applicability depends on the use case and legal analysis. The target should show how its design supports customers that bear these duties.

The buyer should identify decisions that cannot be delegated under contract, policy or regulation. It should also identify the cost of maintaining qualified accountable people. Automation value should be measured after this continuing cost.

10 Measure observability and evidence portability

Observability converts platform activity into evidence. Logs show events, metrics show aggregate behaviour and traces show causality across services. OpenTelemetry semantic conventions include attributes for generative AI systems, agents, models, operations and token usage [20-21]. CloudEvents provides a common envelope for events across systems [22-23].

The target should demonstrate end to end traces across orchestration, agent, model, tool and artifact boundaries. Trace identifiers should persist through asynchronous and external steps where feasible. Sensitive prompts and outputs require controlled capture, retention and access.

Evidence portability means that a customer or buyer can export enough information to reproduce a material event, investigate an incident and support contractual or regulatory reporting. Proprietary dashboards without exportable underlying records can create lock in and weaken assurance.

The team should measure telemetry completeness, sampling, delay, retention, cost and tenant isolation. It should reconcile reported service levels with raw records and customer incidents. Observability cost belongs in sustainable margin because agentic workflows can generate high event and trace volumes.

11 Test evaluation and conformance

Conformance tests whether an implementation follows a specification. Evaluation tests whether it produces acceptable outcomes for defined work. A target needs both. A protocol compliant agent can still be unreliable, unsafe or commercially unusable.

The buyer should obtain the conformance suite, evaluation datasets, expected results, version history and failure thresholds. It should confirm rights to use the data and examine leakage or overfitting. Evaluation should include normal, boundary, adversarial and recovery cases.

Results should be segmented by model, customer, language, tool, workflow and version. Aggregate scores can conceal failure in a commercially important cohort. Consequential workflows should include human review and downstream outcome measures.

The platform should explain how specification changes enter release management. Compatibility windows, deprecation notices and migration tools can be valuable assets. Unfunded backward compatibility can also become a growing support burden.

Table 3 Conformance and outcome evidence matrix

Layer	Evidence	Minimum test	Transaction relevance
Syntax	protocol suite and schema validation	valid and invalid message exchange	basic connector reliability
Semantics	task tool and artifact definitions	same intent across two implementations	usable interoperability
Authority	identity policy and delegation records	permitted denied revoked and expired cases	bounded action and liability
Outcome	accepted workflow cohort	quality latency cost and customer acceptance	revenue and margin support
Recovery	checkpoint replay rollback and incident files	outage duplicate and partial completion	resilience and remediation cost
Switching	export substitution and customer migration	alternate model tool or runtime	dependency and valuation risk

Proposed matrix; pass thresholds require board and customer specific approval.

12 Examine failure recovery and state resumption

Multi agent workflows fail in more ways than linear applications. An agent can time out, return an ambiguous artifact, call a failing tool, lose context, duplicate an action or wait indefinitely for another agent. Recovery should be a designed state machine with ownership and evidence.

The target should identify checkpoints, replay boundaries, idempotency keys, compensation steps and manual intervention. It should demonstrate recovery from model outage, tool outage, credential revocation, corrupted memory and network partition. Recovery time should be measured from customer impact to accepted restoration.

State resumption is especially important for long running tasks. The platform should preserve approved inputs and avoid repeating consequential actions. A replacement agent should understand what has completed, what remains and which authority is still valid.

Incident records should distinguish platform defect, customer configuration, vendor dependency and malicious input. The buyer should reconcile incident cost, service credits, support effort and churn with reported margins.

13 Quantify vendor and model dependence

Model neutrality can be overstated when prompts, evaluations, context windows, tool calling and safety behaviour are tuned to one provider. The buyer should test alternative models using the same accepted workflow and measure quality, latency, cost, failure and support effort.

Cloud dependence can arise through identity, queues, databases, telemetry and managed AI services. A deployment described as portable may require substantial reengineering. The target should provide infrastructure definitions, dependency inventories, exit plans and observed migration experience.

Vendor concentration should be measured across spend, revenue dependency, service criticality and contractual rights. Price changes, usage limits or product withdrawal can alter unit economics. Contract terms should be reviewed for assignment, data use, audit, liability, continuity and termination.

Dependence is not automatically negative. A specialised vendor can provide superior economics and innovation. The transaction question is whether the dependence is understood, contractually supported and reflected in value.

14 Separate open specification from proprietary implementation

Open specifications can expand adoption and reduce customer concern. They can also make interface features easier to replicate. Defensibility may therefore sit in implementation quality, distribution, data rights, evaluations, governance, operating evidence and network participation.

The buyer should examine licences, contributor agreements, trademarks, patents and governance rights. It should identify code copied or modified from open source projects and confirm compliance. Community goodwill can be valuable while remaining difficult to own or control.

Fork risk matters when a buyer plans to change licensing, pricing or governance. Contributors and customers may move to an alternative implementation. The target should show why participants remain: service quality, compatibility, certification, enterprise support, marketplace liquidity or trusted governance.

The acquisition case should separate maintainable proprietary advantage from a temporary lead. Protocol leadership can create influence, but unilateral control can weaken adoption and attract scrutiny.

15 Analyse developer and participant network effects

A multi agent platform may connect developers, tool providers, model providers, customers and agents. Network effects exist when participation improves value for other participants. Registration counts provide weak evidence because inactive or duplicate participants do not create liquidity.

The buyer should measure active developers, published capabilities, successful cross party tasks, repeat use, accepted artifacts, customer concentration and participant retention. It should identify which side subsidises the network and whether pricing can change without reducing participation.

Quality governance is part of the network asset. Certification, reputation, dispute handling and removal of harmful participants affect trust. The cost of these functions should be included in sustainable economics.

Interoperability can increase multi homing because participants can use competing platforms. The buyer should model the resulting limit on take rate and exclusivity. Defensibility can come from superior operating outcomes even when participants remain free to connect elsewhere.

16 Underwrite data rights and switching

The target should provide a data register covering customer data, generated data, telemetry, evaluation sets, memory, agent cards and marketplace records. Each category needs provenance, purpose, permission, retention, export and deletion rules. The buyer should test contract and system alignment.

Switching should be evaluated through an observed exercise. The customer should be able to export relevant data and artifacts, substitute a component and continue a workflow with documented differences. The EU Data Act's switching and interoperability provisions create an important legal context for cloud services [24-26]. Detailed application requires current advice.

Data gravity can remain even when export is available. Large embeddings, proprietary indexes, policy configurations and historical traces can make migration expensive. The transaction model should include the engineering and customer success effort needed to migrate.

The buyer should also examine inbound switching. A platform that can import competitor state accurately may acquire customers faster. Import tools and mappings should be tested against real migrations rather than demonstration data.

17 Test commercial packaging and pricing

Pricing can be based on seats, agents, tasks, tokens, tools, outcomes, orchestration volume or enterprise commitments. Each basis creates a different relationship between customer value and platform cost. Diligence should reconcile contract price, usage, cloud cost, support and gross margin by workflow cohort.

Usage growth can reduce margin when telemetry, model calls, retries and support rise faster than revenue. Minimum commitments can improve predictability while creating unused capacity and renewal risk. Outcome pricing requires clear acceptance and attribution.

The buyer should identify bundled protocol support that has no separate price. It may support retention or cross sell, but its value should be evidenced. Customer interviews and renewal records should show whether interoperability influences purchasing.

Commercial packaging should allocate responsibility among the platform, agent developer, model provider and customer. Ambiguous responsibility can create expensive support and dispute. The acquisition case should fund the operating model customers actually buy.

18 Normalise sustainable EBITDA

Reported EBITDA may exclude the full cost of maintaining connectors, protocols, identity, telemetry, evaluations and compatibility. Capitalised development can also defer expense. The buyer should reconcile payroll, cloud, licences, contractors and customer support to the operating architecture.

The hypothetical target reports USD 54 million revenue and USD 13.5 million EBITDA. The illustration deducts USD 2.0 million for underrecorded connector maintenance, USD 1.2 million for identity and security, USD 0.8 million for observability and evaluation, USD 0.9 million for protocol migration and USD 1.1 million for retention. Sustainable EBITDA becomes USD 7.5 million. These values are management assumptions.

Each adjustment requires evidence. Connector maintenance should be linked to versions and incidents. Security cost should reflect the accepted architecture. Retention should identify people whose knowledge

or authority is not documented. Protocol migration should distinguish recurring compatibility work from a temporary project.

The buyer should avoid counting required remediation as synergy. Remediation protects the acquired revenue and belongs in the stand alone case or transaction funding.

Figure 1 Hypothetical reported to sustainable EBITDA bridge

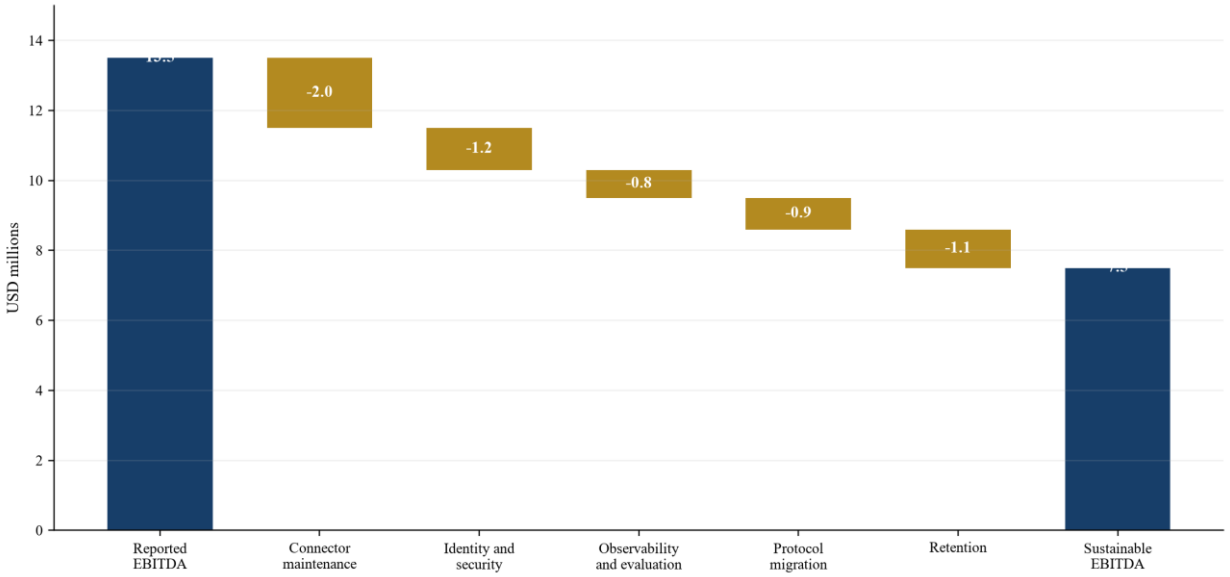

Management assumptions in USD millions; the bridge is not an observation forecast benchmark or valuation conclusion.

Table 4 Hypothetical sustainable EBITDA normalisation

Item	Amount	Evidence required	Potential treatment
Reported EBITDA	13.5	ledger management accounts and quality of earnings	starting point only
Connector maintenance	-2.0	releases incidents people and contractor cost	sustainable operating cost
Identity and security	-1.2	architecture controls incidents and roadmap	sustainable operating cost
Observability and evaluation	-0.8	telemetry tests people and infrastructure	sustainable operating cost
Protocol migration	-0.9	compatibility backlog versions and customer commitments	recurring or funded transition cost
Retention	-1.1	dependency analysis compensation and succession	operating or transaction allocation
Sustainable EBITDA	7.5	reconciled cohort economics	valuation input subject to diligence

Management assumptions in USD millions; transaction treatment requires verified facts and adviser analysis.

19 Build evidence weighted synergies

Synergies should be linked to implemented actions and accepted customer outcomes. The hypothetical gross annual synergy is USD 9.4 million: USD 3.2 million from attach and cross sell, USD 2.0 million from connector rationalisation, USD 1.8 million from shared identity and observability and USD 2.4 million from faster integration.

Continuing costs reduce the illustration by USD 2.3 million for security and control, USD 1.4 million for migrations, USD 1.0 million for partner and employee retention and USD 1.1 million for compatibility and client remediation. Net annual synergy is USD 3.6 million. These are management assumptions and exclude tax, financing and present value.

Cross sell requires customer permission, product fit, sales capacity and accepted deployment. Connector rationalisation requires a migration path and customer support. Shared control can create value while increasing concentration risk. Faster integration needs observed cohorts.

The transaction model should apply evidence weights and timing to each synergy. Conceptual opportunities receive limited value. Contracted, implemented and collected outcomes receive greater weight.

Figure 2 Hypothetical gross to net annual synergy bridge

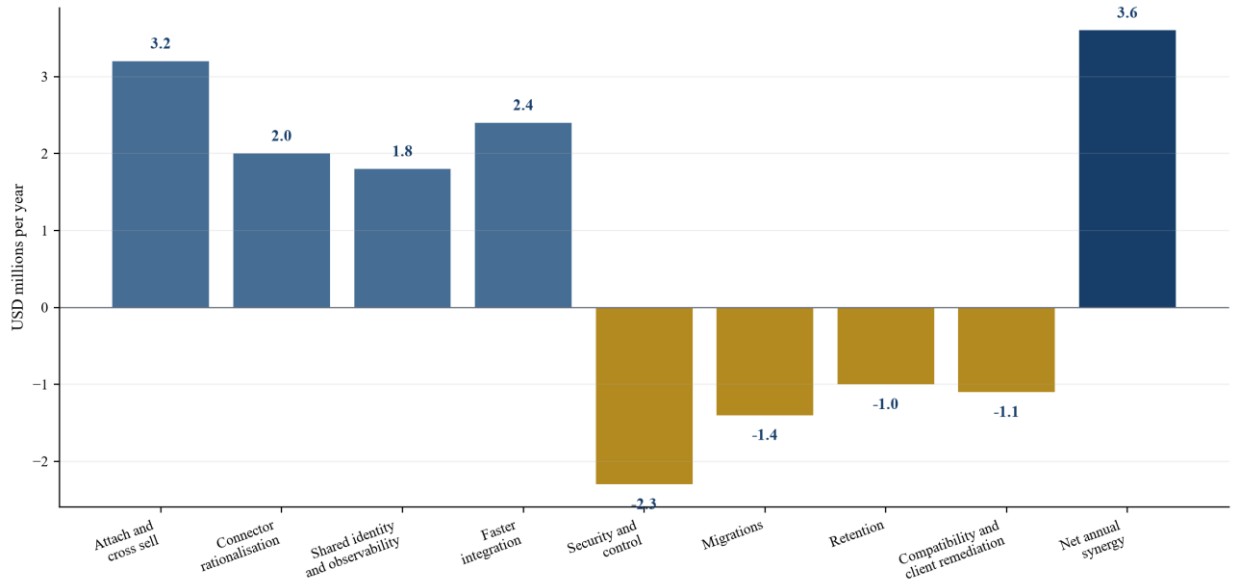

Management assumptions in USD millions; net annual synergy is USD 3.6 million before tax financing and present value.

20 Value the platform on portable economics

Valuation should begin with sustainable stand alone economics. The illustrative case applies ten times to USD 7.5 million sustainable EBITDA, adds USD 14 million of evidence weighted synergy present value, deducts USD 17 million of integration and control cost and deducts USD 10 million for platform, competition and dependency risk. The resulting illustration is USD 62 million.

The multiple is a management assumption, not a market benchmark. A buyer should select methods consistent with the asset, cash flows and available evidence. IFRS 13 describes a framework for fair value measurement, while IFRS 3, IAS 36 and IAS 38 govern relevant accounting considerations [29-33]. Transaction value and accounting measurement remain distinct exercises.

The platform asset can include software, customer relationships, data, trademarks and contractual rights. Interoperability may support these assets without becoming a separately identifiable intangible. Legal rights, separability and accounting analysis are required.

Value should be tested under protocol change, model substitution, customer migration and regulatory cost. A high strategic value can be legitimate when the buyer can execute the operating plan. The evidence should show why that buyer can convert the opportunity.

Figure 3 Hypothetical enterprise value bridge

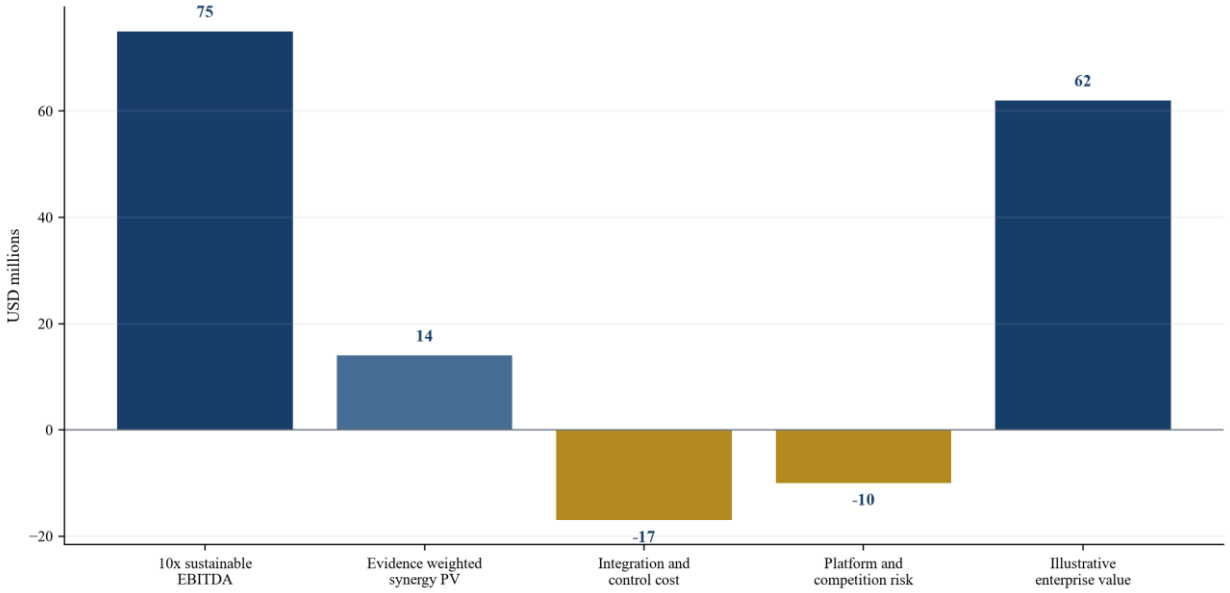

Management assumptions in USD millions; this illustration is not a valuation conclusion or recommendation.

Table 5 Hypothetical valuation bridge

Component	Amount	Basis	Evidence gate
Sustainable EBITDA	7.5	normalised platform economics	reconciled ledger and operating cohorts
Stand alone value	75.0	assumed ten times multiple	approved valuation method and sensitivity
Evidence weighted synergy present value	14.0	probability and timing adjusted	implemented action and accepted customer outcome
Integration and control cost	-17.0	migration security and operating design	costed plan owner and funding
Platform and competition risk	-10.0	dependency multi homing and conduct	legal technical and commercial diligence
Illustrative enterprise value	62.0	arithmetic bridge	investment committee approval

Management assumptions in USD millions; method and inputs require target specific evidence.

21 Review competition and interoperability risk

Platforms can influence access, ranking, data and terms across multiple groups. The US merger guidelines discuss multi sided platforms, complements, visibility into rivals and conduct that can entrench a position [34-37]. The UK merger assessment guidelines also address digital market features and the competitive significance of interoperability [38]. Application depends on facts and jurisdiction.

The buyer should identify whether the target controls an important route to customers, can disadvantage competing agents or tools, gains sensitive information from participants or combines with an adjacent gatekeeper. It should review exclusivity, default ranking, self preference, tying and access terms.

Interoperability commitments can preserve competition while affecting monetisation and integration. The transaction model should include the cost of open interfaces, data separation, neutral governance or behavioural remedies if relevant. A remedy should not be assumed before authority engagement.

Competition risk can also affect the defensibility thesis. Value built on restricting substitution may be less durable than value built on reliable service and trust. The investment committee should understand which mechanism supports price and retention.

22 Select transaction protections

Diligence findings should change price, structure, conditions and covenants. Verified portable economics can support base value. Unproven migration or customer acceptance can support deferred consideration, earn outs or staged acquisition. Material security or rights gaps can require remediation before closing.

Representations can address ownership, licences, open source compliance, data rights, security incidents, customer commitments and protocol support. Indemnities, escrows and insurance require current legal advice. Technical statements should be translated into objectively testable schedules.

Earn out metrics should avoid raw task or connector counts. Suitable measures can include retained recurring revenue, accepted cross platform workflows, gross margin after full operating cost, customer migration completion and service levels. Metrics need audit rights and protection against manipulation.

The buyer should preserve evidence at signing and closing. Fast moving software can change materially during a long transaction. Version, customer and incident updates should be incorporated into closing conditions.

Table 6 Transaction response by finding

Finding	Value effect	Potential deal response	Post close measure
Verified operating interoperability	supports retention and distribution	base value or evidence weighted synergy	accepted workflows and collected revenue
Hidden connector and control cost	reduces sustainable margin	price adjustment or funded plan	full cost per accepted workflow
Weak identity or delegation	creates security and liability exposure	condition remediation escrow or perimeter change	traced actions revocation and incidents
Customer migration restriction	limits synergy and switching	consent condition deferred value or exclusion	approved migration and retention
Key person dependency	threatens continuity	retention succession and deferred consideration	knowledge transfer and service continuity
Competition concern	limits integration or conduct	covenant remedy reserve or no go	compliance and neutral access evidence

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

23 Design the first hundred days

The first phase should preserve code, configurations, registries, logs, evaluation results, credentials, contracts and customer commitments. The buyer should freeze undocumented changes to critical interfaces while allowing security fixes. Access should follow least privilege.

Days sixteen to thirty five should reconcile architecture with deployed systems and select representative workflow cohorts. The team should test identity, delegation, tool authority, context, artifacts, telemetry and recovery. Customer support and finance should link technical outcomes to renewals, credits and cash.

Days thirty six to sixty five should run controlled substitution and migration. The buyer can test alternative models, tools and runtimes, remediate identity gaps and cost the operating design. Changes should be reversible and approved by affected customers where required.

Days sixty six to one hundred should migrate accepted cohorts and release synergies only when evidence gates pass. Governance should report customer outcome, cost, risk and cash together. Unproven value remains deferred.

Figure 4 First hundred days interoperability sequence

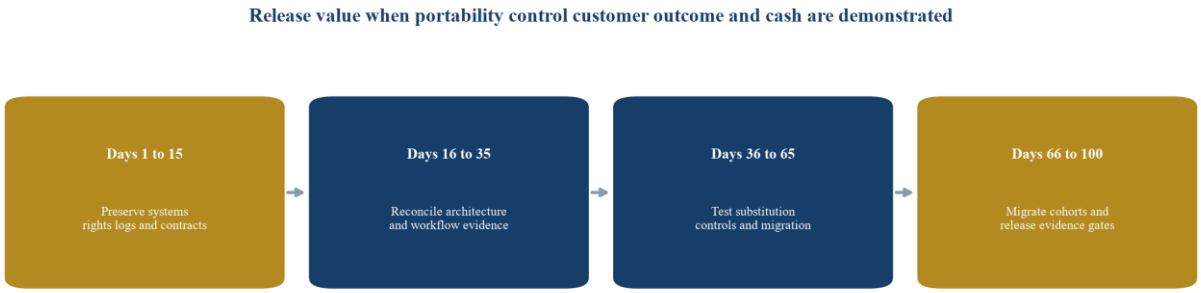

Proposed sequence; actual timing should reflect customers regulation people security and systems.

24 Build the transaction evidence room

The evidence room should be organised around decisions rather than departments. One index should connect acquisition claims to source records, tests, owners and findings. Each material claim should have a date, version and scope.

Technical materials should include architecture, dependency inventories, source repositories, releases, protocol versions, agent cards, schemas, evaluation sets, penetration tests, incidents, service levels and recovery exercises. Commercial materials should include contracts, usage, invoices, credits, renewals, churn, support and customer migration records.

Finance should reconcile cloud, model, telemetry, security, engineering and support cost to customers and workflow cohorts. People materials should identify maintainers, security authority, customer relationships and succession. Legal materials should cover ownership, licences, data, privacy, competition and regulatory commitments.

Access should be controlled and privacy preserving. Sensitive credentials and customer data should remain in secure review channels. The evidence room should retain hashes or version identifiers so conclusions can be traced to reviewed materials.

Table 7 Evidence gates for releasing transaction value

Gate	Minimum evidence	Decision	Measure after release
Capability truth	verified declarations versions and representative tests	accept or revise product perimeter	successful discovery and accepted tasks
Authority	traced principals delegation approval and revocation	approve consequential workflows	authorised actions and exceptions
Portability	export substitution migration and recovery exercise	accept switching and synergy case	migration cost quality and retention

Gate	Minimum evidence	Decision	Measure after release
Sustainable economics	full connector control telemetry and people cost	set valuation earnings	contribution and cash by workflow
Customer acceptance	contracts usage support renewal and consent	include eligible revenue	retained revenue credits and disputes
Synergy release	implemented action accepted outcome and collected cash	recognise or defer value	recurring net cash and residual risk

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

25 Compare hypothetical target archetypes

A protocol specialist may have strong standards participation and weak recurring revenue. Its value can come from talent, influence, certification or an enterprise distribution path. The buyer should avoid capitalising community participation as contracted cash flow.

An enterprise orchestration platform may have recurring revenue and embedded workflows. Its principal risks can be hidden implementation effort, customer specific forks and dependence on one identity or cloud stack. Representative customer migrations are critical.

An agent marketplace may show network potential. Diligence should examine active liquidity, quality governance, take rate, multi homing, dispute handling and participant concentration. Registration counts provide weak evidence.

A vertical multi agent platform may have stronger domain semantics and accepted outcomes. Its narrower market can support defensibility while limiting horizontal expansion. The buyer should test whether domain controls survive combination with a broader platform.

26 Review figures and decision indicators

An interoperability score can focus diligence while remaining an analytical tool. The hypothetical weighting assigns 10 percent to discovery, 15 percent each to task and artifact portability, context and memory, tool contracts, and identity and delegation, 10 percent to observability, 10 percent to conformance and 10 percent to switching and recovery. These weights are management assumptions.

The hypothetical target scores between 46 and 78 across components. A weighted score cannot replace individual no go findings. A weak identity result can block a consequential workflow even when the overall score appears acceptable. The committee should therefore use component thresholds and narrative findings.

Decision indicators should connect technology to economics: accepted cross platform workflows, cost per accepted workflow, migration hours, recurring support, customer retention, service credits, security incidents and collected revenue. Movement over time is more informative than one assessment.

The board should retain the evidence behind every score. A number without reproducible tests can create false precision and weaken accountability.

Figure 5 Hypothetical interoperability evidence score

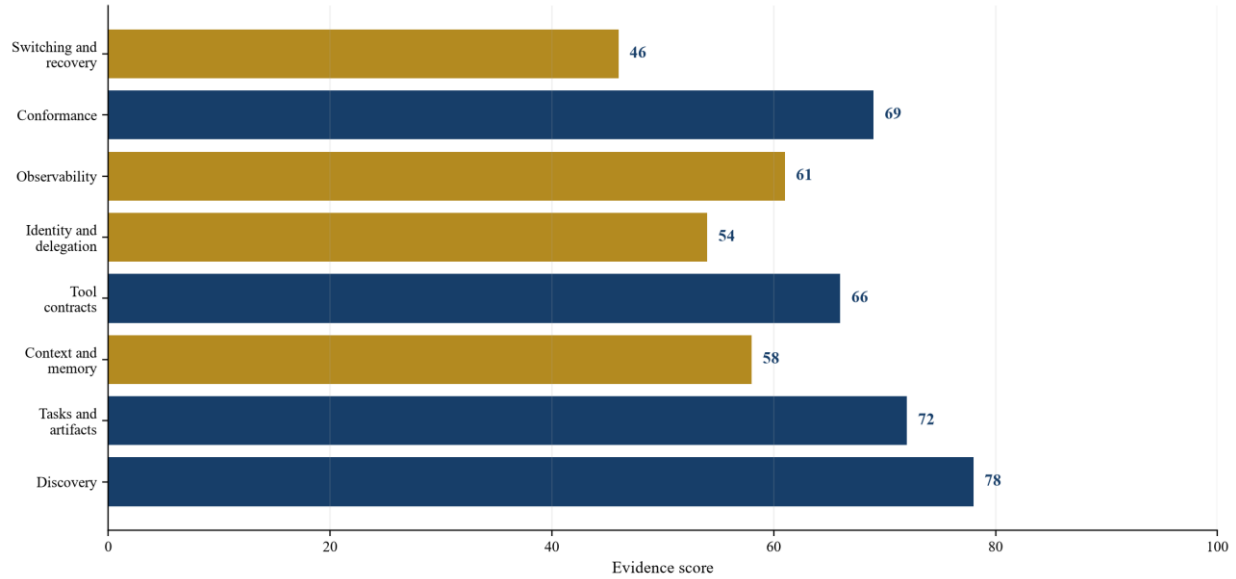

Management assumptions on a zero to one hundred scale; scores and weights are not benchmarks.

27 Limitations and conclusion

Agent standards, products and regulation are changing quickly. Sources reviewed for this paper describe the position available on the publication date. Target facts, customer terms and applicable law require current verification. Hypothetical values illustrate method and provide no forecast, benchmark or investment recommendation.

Interoperability can be a defensible acquisition asset when it produces accepted outcomes across heterogeneous systems with bounded authority, portable evidence and recoverable state. Protocol support contributes to this result while leaving major work in semantics, operations, security, governance and commercial execution.

The buyer should value the complete system. It should normalise the cost of connectors, identity, telemetry, evaluation, migration and people. It should weight synergies by implementation and customer evidence. It should structure consideration around retained economics and use the first hundred days to test substitution and migration.

The strongest acquisition case is therefore measurable: customers continue to buy, workflows continue to work, authority remains controlled, evidence survives component change and cash economics remain attractive. That evidence can support durable platform value even as individual models and protocols evolve.

Frequently asked questions

What makes interoperability defensible in an acquisition

Defensibility comes from reproducible customer outcomes, bounded authority, portable evidence, reliable recovery, trusted governance and economic switching. A protocol interface alone provides limited transaction evidence.

How should a buyer test protocol claims

Use representative production workflows across different models, tools and environments. Record discovery, authorization, task state, artifacts, context, telemetry, failure, recovery, customer acceptance and cost. Include invalid, revoked and interrupted cases.

Can an open protocol create proprietary value

Yes. Value can sit in implementation quality, certification, distribution, domain semantics, evaluations, operating evidence and trusted governance. The buyer should distinguish these assets from features that other implementations can reproduce.

Which interoperability weakness should block a deal

A weakness that prevents lawful or controlled operation of a material workflow can be a no go finding. Examples include unbounded authority, missing data rights, nontransferable customer dependence or recovery that cannot prevent repeated consequential action.

How should connector cost be treated

Recurring maintenance, testing, migration, incident and support cost belongs in sustainable operating economics. A one time remediation plan should be separately funded and should not be counted as synergy.

How should synergies be valued

Tie each synergy to an owner, action, cost, timing and accepted customer outcome. Apply evidence weights and present value. Release value after implemented actions produce recurring net cash.

What should the buyer preserve at closing

Preserve source code, configurations, protocol versions, agent cards, schemas, evaluation data, logs, incidents, credentials, contracts, rights and customer commitments. Apply secure access and privacy controls.

Which post acquisition metrics matter most

Track accepted cross platform workflows, cost per accepted workflow, migration effort, customer retention, service credits, identity and policy exceptions, recovery time, recurring incidents and collected revenue.

References

[1] National Institute of Standards and Technology. AI Agent Standards Initiative. https://www.nist.gov/artificial-intelligence/ai-agent-standards-initiative

[2] National Institute of Standards and Technology. Announcing the AI Agent Standards Initiative for Interoperable and Secure AI Agents. https://www.nist.gov/news-events/news/2026/02/announcing-ai-agent-standards-initiative-interoperable-and-secure

[3] National Cybersecurity Center of Excellence. Accelerating the Adoption of Software and AI Agent Identity and Authorization. https://www.nccoe.nist.gov/sites/default/files/2026-02/accelerating-the-adoption-of-software-and-ai-agent-identity-and-authorization-concept-paper.pdf

[4] National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework. https://www.nist.gov/itl/ai-risk-management-framework

[5] Linux Foundation. Linux Foundation Launches the Agent2Agent Protocol Project. https://www.linuxfoundation.org/press/linux-foundation-launches-the-agent2agent-protocol-project-to-enable-secure-intelligent-communication-between-ai-agents

[6] Linux Foundation. A2A Protocol Surpasses 150 Organizations. https://www.linuxfoundation.org/press/a2a-protocol-surpasses-150-organizations-lands-in-major-cloud-platforms-and-sees-enterprise-production-use-in-first-year

[7] Agent2Agent Project. A2A Protocol Specification 0.3.0. https://a2a-protocol.org/v0.3.0/specification/

[8] Agent2Agent Project. Key Concepts. https://a2a-protocol.org/latest/topics/key-concepts/

[9] Model Context Protocol. Server Concepts. https://modelcontextprotocol.io/specification/draft/server/index

[10] Model Context Protocol. TypeScript SDK Version 2. https://ts.sdk.modelcontextprotocol.io/v2/

[11] Model Context Protocol. July 2026 Specification Update. https://blog.modelcontextprotocol.io/posts/2026-07-28/

[12] Model Context Protocol. July 2026 Release Candidate. https://blog.modelcontextprotocol.io/posts/2026-07-28-release-candidate/

[13] Model Context Protocol. Roadmap. https://blog.modelcontextprotocol.io/posts/mcp-roadmap/

[14] Model Context Protocol. Authorization. https://apps.extensions.modelcontextprotocol.io/api/documents/authorization.html

[15] Model Context Protocol. First Anniversary. https://blog.modelcontextprotocol.io/posts/2025-11-25-first-mcp-anniversary/

[16] OWASP Foundation. Excessive Agency. https://genai.owasp.org/llmrisk/llm06-excessive-agency/

[17] Internet Engineering Task Force. RFC 8707 Resource Indicators for OAuth 2.0. https://www.ietf.org/ietf-ftp/rfc/rfc8707.pdf

[18] Internet Engineering Task Force. RFC 9728 OAuth 2.0 Protected Resource Metadata. https://www.ietf.org/rfc/rfc9728.pdf

[19] MITRE. Adversarial Threat Landscape for Artificial Intelligence Systems. https://atlas.mitre.org/

[20] OpenTelemetry. Generative AI Attributes. https://opentelemetry.io/docs/specs/semconv/registry/attributes/gen-ai/

[21] OpenTelemetry. Semantic Conventions. https://opentelemetry.io/docs/specs/semconv/

[22] Cloud Native Computing Foundation. CloudEvents Specification. https://github.com/cloudevents/spec/blob/main/cloudevents/spec.md

[23] Cloud Native Computing Foundation. CloudEvents Primer. https://github.com/cloudevents/spec/blob/main/cloudevents/primer.md

[24] European Commission. Data Act Explained. https://digital-strategy.ec.europa.eu/en/factpages/data-act-explained

[25] European Commission. EU Data Act Gives Users Control Over Connected Device Data. https://digital-strategy.ec.europa.eu/en/news/eu-data-act-gives-users-control-over-data-connected-devices

[26] European Commission. Cloud Computing Policy. https://digital-strategy.ec.europa.eu/en/policies/cloud-computing

[27] European Union. Regulation 2024 1689 Artificial Intelligence Act. https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX%3A32024R1689

[28] European Commission. AI Act. https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai

[29] IFRS Foundation. IFRS 3 Business Combinations. https://www.ifrs.org/issued-standards/list-of-standards/ifrs-3-business-combinations/

[30] IFRS Foundation. IFRS 13 Fair Value Measurement. https://www.ifrs.org/issued-standards/list-of-standards/ifrs-13-fair-value-measurement/

[31] IFRS Foundation. IAS 36 Impairment of Assets. https://www.ifrs.org/issued-standards/list-of-standards/ias-36-impairment-of-assets/

[32] 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

[33] IFRS Foundation. IAS 37 Provisions Contingent Liabilities and Contingent Assets. https://www.ifrs.org/issued-standards/list-of-standards/ias-37-provisions-contingent-liabilities-and-contingent-assets/

[34] United States Department of Justice and Federal Trade Commission. 2023 Merger Guidelines. https://www.justice.gov/atr/2023-merger-guidelines

[35] United States Department of Justice. Merger Guidelines Overview. https://www.justice.gov/atr/merger-guidelines/overview

[36] United States Department of Justice. Guideline 5 Mergers Can Substantially Lessen Competition by Creating a Firm That Controls Products or Services That Its Rivals May Use to Compete. https://www.justice.gov/atr/merger-guidelines/applying-merger-guidelines/guideline-5

[37] United States Department of Justice. Guideline 6 Mergers Can Substantially Lessen Competition by Entrenching or Extending a Dominant Position. https://www.justice.gov/atr/merger-guidelines/applying-merger-guidelines/guideline-6

[38] United Kingdom Competition and Markets Authority. Merger Assessment Guidelines. https://www.gov.uk/government/publications/merger-assessment-guidelines

[39] OpenAI. Agents SDK. https://developers.openai.com/api/docs/guides/agents/sdk

[40] OpenAI. Agent Orchestration. https://developers.openai.com/api/docs/guides/agents/orchestration

[41] OpenAI. Agents SDK Results. https://developers.openai.com/api/docs/guides/agents/results

[42] OpenAI. New Tools for Building Agents. https://openai.com/index/new-tools-for-building-agents/

[43] National Institute of Standards and Technology. AI Risk Management Framework Playbook. https://www.nist.gov/itl/ai-risk-management-framework/nist-ai-rmf-playbook

[44] National Institute of Standards and Technology. Generative Artificial Intelligence Profile. https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf

[45] International Organization for Standardization. ISO IEC 42001 Artificial Intelligence Management System. https://www.iso.org/standard/81230.html

[46] Model Context Protocol. Security Resources. https://blog.modelcontextprotocol.io/tags/security/

[47] Model Context Protocol. TypeScript SDK Migration Support for 2026 07 28. https://ts.sdk.modelcontextprotocol.io/v2/migration/support-2026-07-28

[48] Model Context Protocol. Go SDK Protocol Documentation. https://go.sdk.modelcontextprotocol.io/protocol/

[49] Cloud Native Computing Foundation. CloudEvents Repository. https://github.com/cloudevents/spec/blob/main/README.md

[50] OpenTelemetry. General Semantic Conventions. https://opentelemetry.io/docs/specs/semconv/general/

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.