Who Authorises the Agent? Delegation and Revocation in AI Identity

An M&A framework for valuing control infrastructure for autonomous systems

Authored by Chennakeshav Adya

Independent Researcher

September 2026

Abstract

AI agents can authenticate to tools, call external services, create downstream agents and continue long-running work after the initiating user has left the session. The commercial opportunity is substantial. The control problem is equally concrete: a business needs to know which principal authorised an agent, which agent received the authority, which actions and resources were included, whether the authority could be delegated again, when it expired, how it was revoked and whether the resulting business effect can be reconstructed. An identity record answers only part of that question. Authority is a changing relationship among a principal, an agent, a task, a policy, a resource and an enforcement point.

This paper develops an Agent Authority Control Framework for boards, investors and acquirers. It separates agent identity, user or system delegation, policy evaluation, runtime enforcement, evidence, revocation and business finality. The framework converts those elements into a delegation graph, a control-latency model, a diligence evidence room, contribution economics, scenario valuation, transaction protections and a first-one-hundred-days integration plan.

Current primary sources establish a credible foundation and an active standards frontier. NIST launched its AI Agent Standards Initiative in February 2026 and published a concept paper focused on agent identification, authorization, auditing and non-repudiation. NIST later described agent registries, credentials, rights delegation, policy management and governance as developing enterprise capabilities. RFC 8693 defines OAuth token exchange and actor semantics for delegation. RFC 7009 and RFC 7662 provide token revocation and introspection mechanisms. OpenID AuthZEN 1.0 standardises the exchange of authorization requests and decisions. The OpenID Shared Signals Framework and Continuous Access Evaluation Profile reached final-specification status in 2025. Emerging IETF work addresses delegation chains, task-level authorization, asynchronous workflows, cascading revocation and agent identity management. These drafts are evidence of active design work; their draft status means they cannot be treated as settled interoperability.

A wholly hypothetical acquisition case applies the framework to an enterprise control platform used by financial, industrial and professional-services customers. The target reports 9,600 registered agents, 4,100 active agents, 280 million annual policy evaluations and 38 million consequential operations. The model prices current value from contracted recurring revenue, verified enforcement coverage, gross contribution and customer retention. It assigns separate option value to emerging standards alignment and discounts revocation gaps, shared credentials, incomplete delegation evidence and concentration. Every company, amount, rate and operating assumption in the case is illustrative management input. It is not observed market data, a valuation opinion, a financing proposal or a forecast.

The central conclusion is that agent-control value resides in enforceable authority, not agent count. A valuable platform binds each material action to a distinct agent, an accountable principal, a constrained purpose, current policy and a non-bypassable enforcement point. Revocation has value only when it reaches every relevant credential, token, queue, sub-agent and effect boundary within the required time. An acquirer should value verified control coverage and recurring contribution as current economics; price standards convergence as an option; reserve for remediation and migration; and connect deferred consideration to evidence that survives change of control.

JEL Classification: G34, L22, L86, M15, O31, O33

Keywords: AI agents, delegated authority, revocation, non-human identity, OAuth, workload identity, agent governance, M&A, autonomous systems, continuous authorization

This paper is an educational and structural analysis prepared for research purposes. It is not investment advice, an offer, or a solicitation. It contains no client information.

1. DEFINE THE ACQUISITION DECISION

The acquisition decision is whether the target controls autonomous authority across real business workflows or provides an administrative layer that can be bypassed. A platform can maintain an impressive registry of agents while individual tools continue to accept shared user credentials, static API keys or locally cached permissions. The buyer needs transaction-level evidence that policy decisions reach the point where action becomes possible.

The commercial thesis should identify the decision being improved. A control platform can accelerate enterprise deployment, lower integration effort, reduce manual approval cost, improve auditability and support regulated use. Each benefit requires an observable baseline, an implemented control and a measurable outcome.

The diligence sequence begins with consequential operations. The buyer selects payments, data exports, contract changes, production deployments or customer communications and reconstructs who initiated them, which agent acted, which authority was presented, which policy applied, which enforcement point permitted the action and which evidence remained.

Table 1. Agent-authority value layers and acquisition evidence

Value layer	Minimum evidence	Principal exposure	Transaction response
Agent identity	Unique cryptographic identity and lifecycle record	Shared credentials or identity drift	Remediation condition
Delegation	Named principal, purpose, scope, limits and expiry	Impersonation or overbroad consent	Value constrained grants
Policy decision	Versioned decision with relevant context	Hidden local rules and inconsistent outcomes	Reconcile policy estate
Enforcement	Non-bypassable checks at tools and resources	Advisory decisions ignored downstream	Price verified coverage
Revocation	Tested propagation to credentials, tokens and queues	Stale authority after withdrawal	Reserve and service-level covenant
Evidence	Tamper-evident action and decision records	Incomplete attribution or repudiation	Closing evidence gate
Economics	Revenue and contribution by controlled workflow	Volume without monetised control	Underwrite recurring contribution

Every layer needs a defined owner, production evidence and a transferability test.

2. SEPARATE IDENTITY FROM AUTHORITY

Identity establishes which entity is interacting with a system. Authority defines which actions that entity may perform under current conditions. Authentication can prove that a request came from Agent A. It does not establish that Agent A may transfer funds, read a customer file or create Agent B for the current task.

An agent may act for a person, a legal entity, a software service or itself within pre-approved operating limits. The authority source changes the evidence required. User delegation needs informed consent and scope. Enterprise delegation needs accountable ownership, policy and employment or role context. System delegation needs workload identity, service ownership and change control.

The valuation model should therefore avoid counting authenticated agents as controlled agents. Current value belongs to agents whose material operations pass through enforced, contextual and revocable authority.

3. MODEL THE AUTHORITY TUPLE

The authority tuple contains principal, agent, action, resource, purpose, constraints, validity period and evidence. Some deployments also need jurisdiction, counterparty, data class, risk tier, monetary limit, approval threshold and delegation depth.

The tuple should be machine-readable and stable enough for policy evaluation. Natural-language prompts remain useful context; they are ambiguous as an enforcement boundary. A structured proposal can capture the operation that a principal reviews and the resource server later evaluates.

The buyer should sample tuples from production and reconcile them to contracts, policies, user interfaces and logs. Missing attributes indicate either a control gap or a business decision being made elsewhere. Both affect integration scope and valuation.

4. DRAW THE DELEGATION GRAPH

The delegation graph represents principals, agents, sub-agents, authorization servers, policy decision points, enforcement points, tools and resources. An edge records who delegated which authority to whom. A valid chain preserves provenance and narrows or maintains authority at each hop.

Graph depth creates operating leverage and risk. Multi-agent orchestration can allocate specialist work efficiently. It can also obscure the initiating principal and allow authority to expand through inconsistent policies. The buyer should test strict-subset rules, audience restriction, proof of possession, expiry and maximum depth.

Graph concentration matters. A single identity provider, policy engine or gateway may protect many workflows. That central position can create durable distribution and recurring revenue. It can also create common-mode failure, customer concentration and change-of-control dependence.

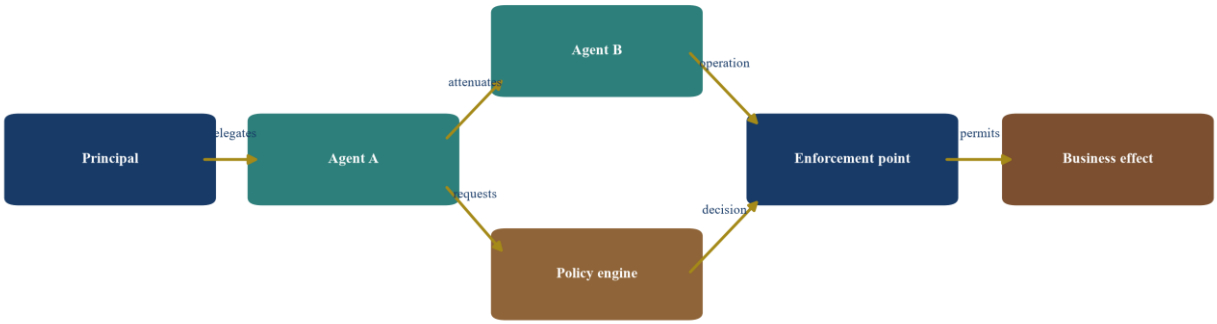

Figure 1. Delegation graph from principal to consequential effect

Every authority edge should preserve provenance, narrow scope and terminate at an enforcement point.

5. DISTINGUISH DELEGATION FROM IMPERSONATION

RFC 8693 distinguishes delegation from impersonation. In delegation, the subject retains its identity and an actor acts on its behalf. The actor claim can preserve that relationship. Impersonation collapses the distinction and can make an agent appear to be the user.

This distinction changes audit, consent and liability. A bank, employer or marketplace may accept a person acting directly and apply a different policy when software acts for that person. Shared credentials remove the signal needed to make that decision.

The buyer should examine token claims, logs and user interfaces for explicit actor identity. Revenue generated through impersonation patterns can require redesign before regulated customers expand deployment. The migration cost belongs in the valuation bridge.

6. REGISTER EACH MATERIAL AGENT

An agent registry should identify owner, developer, model, version, deployment, credentials, permitted environments, risk tier and lifecycle state. It should also link to policy, monitoring, incident and retirement records.

Registration creates inventory and accountability. It does not create enforcement by itself. The diligence test compares the registry with runtime observations, credential issuance, service accounts, API gateways and tool logs. Unregistered agents and orphaned credentials indicate control coverage below reported levels.

The buyer should define materiality. A deterministic read-only helper may receive a lighter control profile than an agent that can move money or deploy software. Risk-tiering supports pricing, service levels and operating efficiency.

7. BIND AGENT IDENTITY TO WORKLOAD IDENTITY

Enterprise agents run as workloads in containers, functions, applications and orchestration services. Workload identity can provide short-lived credentials and cryptographic proof that the expected software instance is acting. NIST identifies SPIFFE, OAuth and emerging WIMSE work as relevant foundations.

The identity should bind to deployment and key material. A copied token should have limited value without the bound key. Proof-of-possession methods can reduce replay exposure. Attestation can add evidence about the execution environment where customers require it.

The buyer should test issuance, rotation, recovery, deployment transfer and revocation. Long-lived secrets in configuration files create remediation cost and weaken the platform's claim to control autonomous systems.

8. IDENTIFY THE ACCOUNTABLE PRINCIPAL

Every material authority chain needs an accountable principal. The principal can be an individual, an enterprise role, a service or a legal entity. Ownership should remain visible when people change roles, systems are reorganised or workflows cross companies.

Accountability cannot depend solely on the employee who first configured the agent. Enterprise authority needs a business owner, technical owner and control owner. Customer-facing agents may also need a clear relationship to the contracting entity.

The buyer should review orphaning controls. Termination, role change, business-unit transfer and corporate separation should trigger reassessment or withdrawal of authority. Slow ownership updates create hidden authority that survives the reason it was granted.

9. CAPTURE USER INTENT AS STRUCTURED AUTHORIZATION

User intent is relevant when an agent performs a specific operation. A broad consent screen can be inadequate for a material payment, disclosure or commitment. Emerging IETF work explores structured operation proposals and authorization tokens that record the approved operation without treating raw prompts as policy.

The design should show the principal a clear action, resource, counterparty, amount, timing and relevant constraints. The authorization record should bind that reviewed proposal to the executing agent and protected resource.

Draft proposals are evidence of design direction. They do not establish settled standards. An acquirer should value production controls and interoperability already demonstrated; future alignment belongs in the option case.

10. USE LEAST AUTHORITY AT EACH HOP

Each delegation hop should preserve or narrow authority. A travel-planning agent allowed to arrange a journey should not give a payment sub-agent unlimited merchant and amount access. Scope, audience, time, purpose and delegation depth should contract as work becomes specific.

Strict-subset validation can be enforced by an authorization or delegation service. Offline attenuation can reduce latency and central dependency; online issuance can support synchronous

policy and revocation. The architecture choice should reflect the consequence and operating environment.

The buyer should test for scope inflation across real chains. Textual scope names are insufficient where the downstream service interprets them differently. The evidence room should include policy definitions and resource-server behaviour.

11. SEPARATE POLICY DECISION FROM ENFORCEMENT

OpenID AuthZEN separates policy decision points from policy enforcement points through a standard authorization API. This separation can support consistent policy across gateways, applications and tools. It also creates a requirement: the enforcement point must actually call or apply the decision.

The buyer should inventory decision points and enforcement points. It should identify cached decisions, fallback modes, local overrides and tools that can be reached through another route. Advisory policy that can be ignored has limited control value.

Commercial value increases when one policy plane governs diverse resources without adding excessive latency. Customer evidence should show deployment breadth, reliability, decision volume and incident outcomes.

12. PUT POLICY AT THE CONSEQUENCE BOUNDARY

The best enforcement point sits where an operation can be stopped before the consequential effect. For a payment, this can be the payment initiation or approval boundary. For a data export, it can be the service that releases the data. For production change, it can be the deployment controller.

An upstream orchestration check can be bypassed when an agent holds direct credentials. A downstream audit log records the event after control has been lost. The architecture should connect current policy to the non-bypassable boundary.

The buyer should test alternate routes, retries, queues and emergency modes. Coverage should be measured by consequential operations protected, not APIs nominally integrated.

13. CLASSIFY AUTHORITY BY CONSEQUENCE

Authority tiers can align control effort with potential loss. Low-consequence actions include search, summarisation and draft creation. Medium-consequence actions can change workflow state or contact customers. High-consequence actions include money movement, legal commitments, privileged access, destructive operations and regulated decisions.

Each tier should define required identity, consent, policy, approval, token lifetime, monitoring, evidence and revocation target. Higher tiers can require step-up approval, dual control or a human checkpoint.

This tiering creates a product architecture and a commercial packaging structure. The buyer should verify that premium controls correspond to customer willingness to pay and observable risk reduction.

14. DESIGN TIME-BOUNDED AUTHORITY

Agent authority should expire on a schedule matched to the task. Short-lived execution tokens reduce the window for misuse. Long-running workflows need renewal mechanisms that preserve the principal, actor, task and policy context.

Expiry alone cannot handle immediate withdrawal. It remains a useful first boundary and reduces the burden on revocation systems. Token lifetime should reflect sensitivity, connectivity, queue duration and the time required to obtain fresh authority.

The buyer should compare configured lifetimes with actual task duration. Excessively short lifetimes can drive insecure workarounds. Excessively long lifetimes expand stale-authority exposure.

15. PRESERVE CONTEXT IN ASYNCHRONOUS WORKFLOWS

Agents frequently place work on queues, schedule jobs or resume after the original access token expires. The system needs a controlled way to obtain downstream authority while preserving the original subject and actor relationship.

Emerging OAuth work addresses asynchronous delegation and the continuity of actor semantics across refreshed or exchanged tokens. Current implementations may use proprietary grants, durable workflow credentials or service accounts. Each pattern should be assessed for scope, renewal and revocation.

The buyer should select a delayed operation and trace authority across every pause and resume. A fresh service token that loses the initiating principal can break accountability even when the job remains technically authenticated.

16. TREAT TOOL ACCESS AS A SEPARATE CONTROL SURFACE

Agents obtain capability through tools. A model with no direct privilege can still perform material actions through a browser, database connector, payment API or code runner. Tool catalogues, permissions and runtime routing therefore belong in the authority model.

The control platform should identify each tool, owner, resource, consequence tier and permitted agent population. Tool descriptions supplied to a model should align with enforceable policy at the service boundary.

The buyer should test hidden tools, duplicate connectors and direct endpoints. A polished orchestration layer can coexist with unmanaged access paths that bypass the central policy engine.

17. PREVENT CONFUSED-DEPUTY BEHAVIOUR

A confused deputy uses its own authority to serve a request that the initiating principal could not perform directly. Agents increase this risk because they interpret instructions, select tools and delegate tasks dynamically.

Audience restriction, subject and actor claims, resource-specific tokens and contextual policy can reduce the exposure. The resource should evaluate the principal and agent relationship, not only the agent's service identity.

Diligence should include adversarial tests involving prompt injection, cross-user context, tool redirection and untrusted content. Security findings should connect to control coverage and remediation cost rather than remain a separate technical appendix.

18. BUILD REVOCATION AS A CONTROL SYSTEM

Revocation is a system of triggers, decisions, distribution and enforcement. A principal can withdraw consent. An administrator can disable an agent. A fraud engine can raise risk. A model version can be quarantined. A task can finish. A counterparty can become prohibited.

RFC 7009 defines token revocation. RFC 7662 allows a resource server to query whether a token remains active. Shared Signals and CAEP support event-driven changes in security state. These mechanisms can complement short-lived credentials and policy re-evaluation.

The buyer should map each trigger to affected credentials, grants, sessions, queues, sub-agents and resources. Unmapped descendants create partial revocation.

19. MEASURE REVOCATION LATENCY

Revocation latency runs from the decision to withdraw authority until the last relevant enforcement point rejects it. Components include detection, policy update, event delivery, cache expiry, token invalidation, queue cancellation and resource response.

One headline service level is inadequate. A payment agent may require seconds. A low-risk reporting task may tolerate minutes. Offline environments need bounded token lifetimes and a defined reconnect policy.

The platform should produce percentile evidence by control tier. The buyer should test actual revocation, including an in-flight operation and a downstream sub-agent. Design documentation cannot substitute for observed propagation.

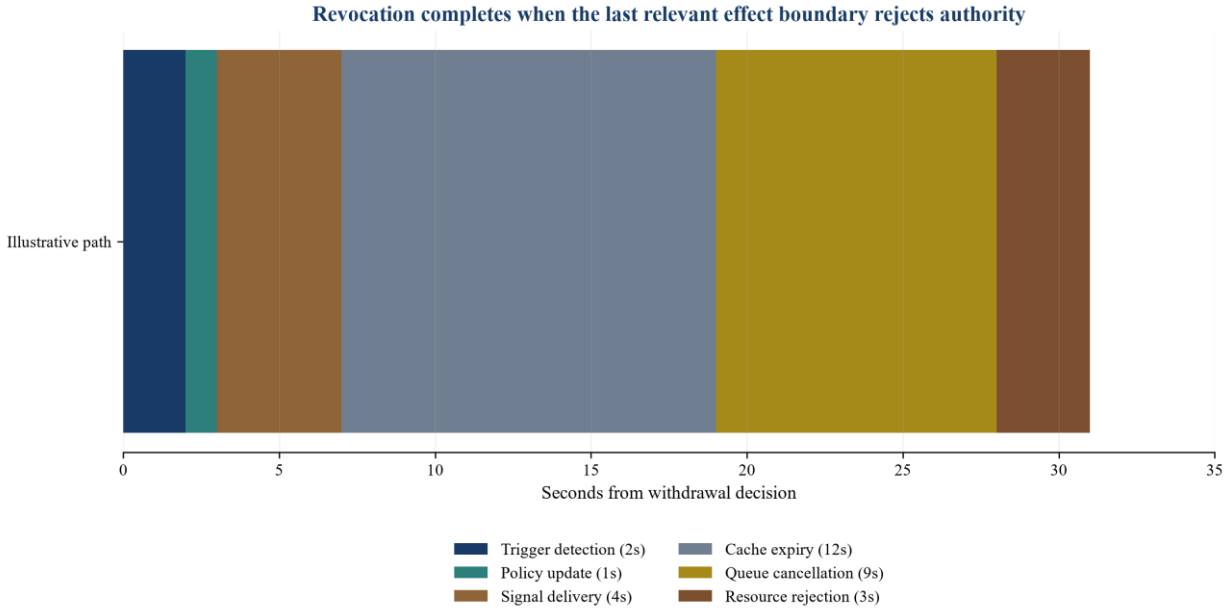

Figure 2. Revocation path and latency budget

Illustrative timings show where stale authority can survive after a withdrawal decision.

20. DECIDE BETWEEN ONLINE AND OFFLINE VERIFICATION

Online introspection can provide current token state and central policy. It adds latency, availability dependency and data about resource activity. Offline signed tokens provide scale and resilience. They remain valid until expiry unless the verifier receives status or revocation information.

The design can combine short-lived offline tokens with event-driven signals, status lists or online checks for higher-consequence actions. The decision should be explicit by control tier.

The buyer should examine cache settings and failure behaviour. A resource that permits action when introspection fails can create the largest exposure during an outage. A resource that blocks every action can create operational loss. Contracts and service levels should reflect the chosen balance.

21. REVOKE THE CHAIN, NOT ONE TOKEN

A multi-agent task can produce several access tokens, refresh tokens, delegation grants, queued jobs and derived credentials. Revoking one visible token can leave descendants active. The authority inventory should support root, branch, task, agent, user and resource-level withdrawal.

Emerging IETF drafts explore cascading revocation and attenuated delegation chains. Their design choices reveal unresolved trade-offs around offline verification, synchronous state and interoperability. Production implementations need their own tested semantics today.

The buyer should document which objects are revoked together and which remain independent. Customer-facing controls should communicate the actual effect clearly.

22. CLOSE THE AUTHORIZATION-TO-EFFECT GAP

An operation can be authorised, queued and executed after authority changes. A transfer may pass policy and then wait for settlement. A deployment may be approved and later reach production. A message may be generated and released by another system.

Revocation at the token layer does not automatically cancel an effect already admitted into a downstream process. High-consequence workflows need a finality boundary or a second check close to irreversible execution.

The buyer should trace the interval from authorization to effect. It should identify cancellation capability, idempotency, rollback and compensation. This control gap can become a material source of liability.

23. USE CONTINUOUS AUTHORIZATION SIGNALS

Security state changes during a session. Account compromise, device posture, employment status, counterparty risk and policy can invalidate an earlier decision. OpenID Shared Signals and CAEP provide a framework for communicating such events among cooperating systems.

Continuous authorization should define which signals change agent authority, how quickly they propagate and how recipients acknowledge or apply them. Signal quality and false positives affect both security and workflow continuity.

The buyer should review event coverage, delivery reliability, deduplication, replay protection and customer adoption. A standards-compatible publisher has limited value when protected resources ignore the events.

24. RECORD DECISION AND ACTION EVIDENCE

Audit requires more than an action log. The evidence should link principal, agent, delegated scope, policy version, input attributes, decision, enforcement point, tool call, resource response and resulting business state.

Evidence should support investigation and customer assurance while respecting data minimisation. Raw prompts may contain sensitive or irrelevant material. Structured operation records and references can preserve accountability without copying every input into every token.

The buyer should test whether a material event can be reconstructed after system changes and staff turnover. Evidence retention, integrity, access and export affect enterprise value.

25. ALLOCATE LIABILITY ACROSS THE CHAIN

Loss can arise from the principal, agent developer, model provider, orchestration platform, identity provider, policy engine, tool owner, resource server or customer configuration. Contracts should allocate responsibility according to control and failure mode.

The platform should distinguish identity assurance, policy service, enforcement integration and business outcome. Broad assurances that an agent is secure can exceed the actual service boundary.

The buyer should compare contractual caps, indemnities, insurance, incident history and technical telemetry. Liability that cannot be matched to evidence can warrant a reserve or specific transaction protection.

26. MAP REGULATORY ACCOUNTABILITY

The EU AI Act requires record keeping and human oversight for specified high-risk systems. The exact obligations depend on system classification, role and use case. Agent authorization evidence can support accountability; it does not determine regulatory compliance by itself.

Financial services, healthcare, employment, critical infrastructure and public administration can add sector-specific duties. A global control platform needs a regulatory map by workflow and jurisdiction.

The buyer should verify customer representations, product documentation and legal analyses. Revenue attributed to regulated readiness should be supported by contracted requirements and implementation evidence.

Table 2. Authority-control architecture and diligence tests

Control domain	Established foundation	Developing frontier	Diligence test
Delegation	OAuth and RFC 8693 token exchange	Agent-specific grants and chains	Preserve subject and actor across hops
Decision	ABAC concepts and AuthZEN API	Rich task and purpose semantics	Reproduce policy decision from evidence
Revocation	RFC 7009 and RFC 7662	Cascading and task-level revocation	Withdraw root and test descendants
Continuous change	Shared Signals and CAEP	Agent-specific security events	Measure delivery and enforcement latency
Workload identity	SPIFFE and enterprise workload practices	WIMSE agent identity models	Bind identity to deployment and key
Audit	Enterprise logging and signed tokens	Standard agent authorization evidence	Reconstruct consequential operation

Standards maturity and production implementation should be assessed separately.

27. DEFINE CONTROL COVERAGE

Control coverage is the share of consequential operations that pass through the required identity, decision, enforcement, evidence and revocation controls. The denominator should be operations, not agents or integrations.

Coverage can be weighted by consequence. Protecting every low-risk search does not compensate for an unmanaged payment or production-deployment path. A weighted measure helps the board prioritise remediation.

The buyer should reconcile gateway logs, tool logs, resource logs and agent inventory. Reported coverage should exclude sandbox, read-only and duplicate traffic when the acquisition thesis concerns high-consequence control.

28. BUILD CUSTOMER-LEVEL COHORTS

Customer cohorts should begin at production activation. Measures include time to register agents, connect identity, express policy, integrate enforcement, test revocation and expand to additional workflows.

Maturing products should show falling implementation effort, higher policy reuse and broader control coverage. Bespoke professional services can create revenue and limit software scalability.

Retention should be analysed by controlled workflow and contract. Customers can retain a platform while leaving material actions outside its enforcement boundary. Expansion quality matters more than licence-seat growth.

29. PRICE BY CONTROLLED OUTCOME

Pricing can follow agents, policy evaluations, protected resources, consequential operations, customer tiers or risk modules. Agent counts are easy to explain and weakly connected to value when agents vary widely in activity and consequence.

Policy-evaluation volume can align with infrastructure use. It can encourage caching or bypass when every call creates cost. Pricing by protected workflow or consequence tier can align value and create clearer packaging.

The buyer should calculate realised price and gross contribution by customer and workflow. Free development traffic and duplicated evaluations should remain outside current monetised demand.

30. MEASURE DECISION ECONOMICS

Decision economics include revenue per evaluation or protected workflow, compute, storage, event delivery, support, professional services and third-party identity cost. Low-latency policy can require regional infrastructure and high availability.

Gross margin can improve through reusable connectors and policy templates. It can deteriorate through customer-specific logic, noisy logs, long retention and around-the-clock incident commitments.

The buyer should separate software contribution from implementation contribution. Future margin expansion should depend on observed cohort improvement rather than a generic software multiple.

31. EVALUATE DATA AND MODEL DEPENDENCE

Authorization decisions can depend on identity, role, device, transaction, risk, model, tool and environmental attributes. Data freshness and provenance affect control quality. A sophisticated policy engine cannot correct stale or unauthoritative inputs.

Machine-learned risk can support decisions. Deterministic rules and accountable owners remain important for high-consequence boundaries. The buyer should identify where a model influences denial, approval or escalation.

Data rights, portability and customer-specific attributes affect transferability. The diligence model should cost rebuilding features that cannot move with the transaction.

32. TEST CHANGE-OF-CONTROL TRANSFERABILITY

The target may depend on cloud marketplaces, identity providers, trust domains, code-signing keys, model agreements, security certifications and customer approvals. A change of control can trigger consent, reassessment or termination.

The buyer should map critical relationships to revenue and control coverage. Non-transferable keys or designations need a transition plan that preserves continuous authorization and evidence.

Closing conditions can include customer consent, credential reissuance, key custody, policy export and successful revocation tests. Contingent consideration can align value where transferability remains uncertain.

33. CONSTRUCT THE HYPOTHETICAL CASE

The hypothetical target provides agent registry, delegated authorization, policy evaluation, enforcement integrations, event signals and evidence retention. It serves financial, industrial and professional-services customers.

Management reports 9,600 registered agents, 4,100 active agents, 280 million annual policy evaluations and 38 million consequential operations. The platform records USD 18.4 million of annual recurring revenue and USD 2.6 million of implementation and support revenue.

Every input is illustrative. The case demonstrates the framework. It does not describe a named company, market benchmark or expected result.

34. RECONCILE THE HYPOTHETICAL CONTROL FUNNEL

The case begins with 280 million policy evaluations. Deduplication and system checks identify 96 million agent-initiated business operations. Thirty-eight million are classified as consequential. Thirty-two million reach a verified non-bypassable enforcement point with complete decision evidence.

The gap of six million consequential operations includes advisory-only decisions, direct credentials, legacy integrations and incomplete evidence. Management estimates that four million can migrate within eighteen months. This estimate requires customer plans and delivery evidence.

Current control coverage for valuation is 84 percent of consequential operations. The remaining coverage belongs in remediation and option scenarios.

35. CALCULATE THE HYPOTHETICAL BASE ECONOMICS

The case assumes USD 21.0 million of total revenue. Direct cloud, event, storage, third-party identity and support cost is USD 5.9 million. Direct implementation cost is USD 1.4 million. Contribution before product development and central overhead is USD 13.7 million.

The largest five customers produce 44 percent of recurring revenue. One cloud identity integration supports 37 percent of protected operations. Contracted gross retention is 94 percent; verified workflow retention is assumed at 90 percent after excluding dormant deployments.

These amounts and rates are hypothetical management inputs. A real valuation requires ledger, contract, invoice, telemetry and customer reconciliation.

36. BUILD AN EVIDENCE-WEIGHTED CONTROL SCORE

The illustrative score grades agent identity, principal binding, least authority, enforcement coverage, revocation latency, chain revocation, evidence and transferability. Current evidence is scored separately from the post-remediation target.

The score is a diligence organiser rather than a valuation multiple. Each point should link to a source, owner and required action. A low score in non-bypassable enforcement can dominate higher scores elsewhere.

The investment committee should view both average and floor. Minimum control for the highest-consequence workflow can determine the risk appetite and closing conditions.

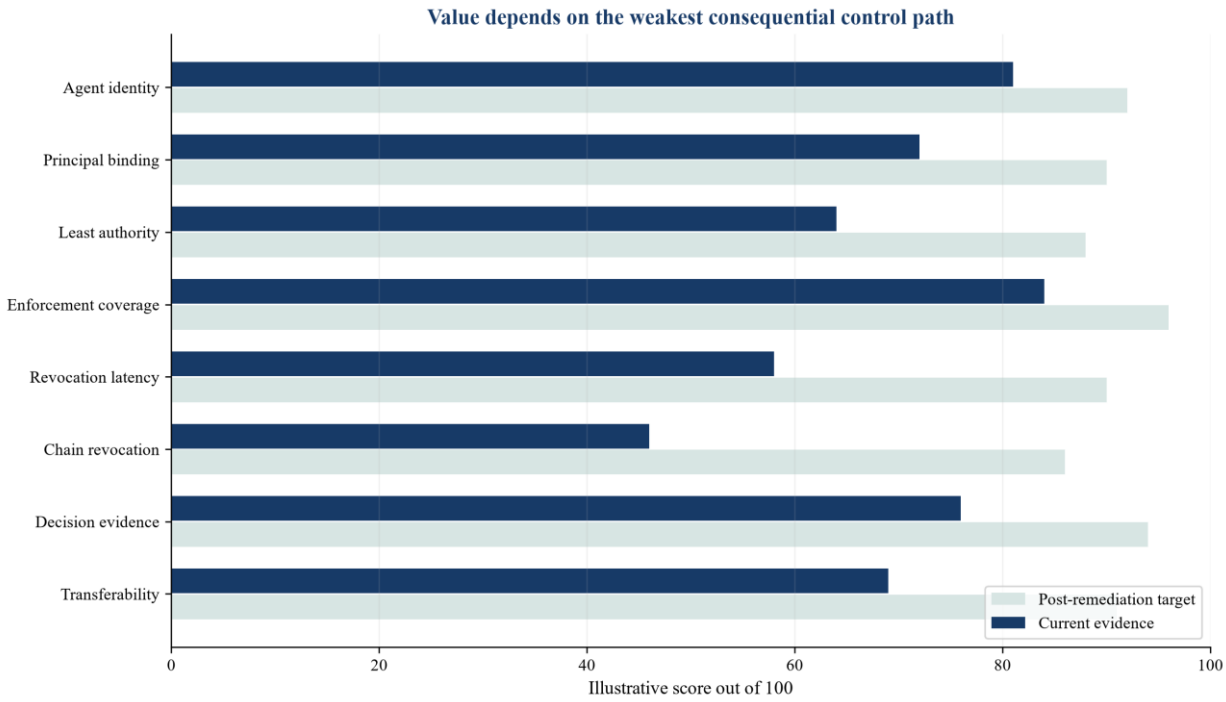

Figure 3. Evidence-weighted agent-authority score

The illustrative profile separates current production evidence from the post-remediation target.

37. BUILD THE VALUATION BRIDGE

Current recurring contribution can support upfront value when contracts, workflow retention and control coverage are verified. Remediation cost should include identity migration, connector changes, policy reconciliation, revocation engineering, evidence retention and customer assurance.

Standards alignment can lower integration cost and broaden distribution. Draft-dependent features should receive probability-weighted option value. Revenue associated with unmanaged operations should be discounted until enforcement is demonstrated.

Table 3. Hypothetical valuation bridge

Component	Illustrative basis	Value logic	Transaction treatment
Current contribution	USD 13.7m before product and central overhead	Verified recurring workflows	Cash-flow valuation
Coverage uplift	Four million operations planned for migration	Additional protected demand	Milestone consideration
Standards option	AuthZEN, Shared Signals and emerging agent protocols	Lower future integration cost	Probability-weighted earn-out
Revocation gap	Slow or partial withdrawal across descendants	Operational and liability exposure	Remediation reserve
Concentration	Five customers and one identity integration	Revenue and infrastructure dependence	Discount and consent condition
Transferability	Customer, key and trust-domain migration	Ability to preserve service after closing	Condition precedent

All amounts, weights and probabilities are illustrative management inputs.

38. SELECT DEAL PROTECTIONS

Representations should address agent inventory, material credentials, delegated authority, policy coverage, revocation tests, incidents, customer commitments, intellectual property and standards claims. Definitions should match technical evidence.

Specific indemnities can address known authority failures or regulatory investigations. Escrow can support remediation and customer consent. Earn-outs should use verified protected workflow revenue or contribution, not raw agent or evaluation counts.

The buyer should retain access to telemetry and customer evidence during any measurement period. Product changes and integrations should have agreed treatment to reduce dispute.

39. BUILD THE DILIGENCE EVIDENCE ROOM

The evidence room should contain registry exports, credential inventory, delegation records, policy versions, authorization decisions, enforcement configurations, tool catalogues, revocation tests, signal delivery, queue controls, action evidence, incidents, customer contracts, invoices and contribution schedules.

Samples should cover successful, denied, expired, revoked, retried and disputed operations. The buyer should reconstruct one high-consequence chain from principal through final effect and one emergency withdrawal from trigger through every descendant.

Evidence should be graded by source, completeness, period and reconciliation. Slides can explain architecture. Production records and independent customer evidence should substantiate performance.

40. PLAN THE FIRST ONE HUNDRED DAYS

The first phase should freeze uncontrolled authority expansion and establish a unified inventory. The buyer should protect keys, revocation services, policy availability, customer commitments and incident response.

The combined company should define one authority taxonomy, consequence model, delegation graph and control-coverage denominator. High-consequence bypasses and long-lived shared credentials should receive priority.

Product integration should preserve customer policy and evidence. Broad consolidation can wait until transferability, revocation and rollback are tested.

41. GOVERN THE STANDARDS ROADMAP

The standards roadmap should distinguish final specifications, adopted drafts, proprietary profiles and experimental work. Each dependency needs a business owner, implementation owner, migration plan and customer impact.

NIST's initiative and active IETF work indicate continued development. The board should avoid presenting draft alignment as completed interoperability. Production conformance and independent counterparty acceptance provide stronger evidence.

Roadmap investment should link to customer activation, control coverage, integration cost, retention or regulatory readiness. Technical activity without a decision outcome belongs in product operations rather than valuation.

42. USE A BOARD AND INVESTMENT-COMMITTEE SCORECARD

The scorecard should track active agents, accountable ownership, high-consequence operations, enforced coverage, least-authority exceptions, revocation percentiles, descendant withdrawal, policy availability, decision evidence, incidents, retention and contribution.

Metrics need definitions and reconciliation. A rising agent count can coexist with falling control coverage. Faster decisions can coexist with stale authorization when caching expands.

Board action should connect to risk appetite, pricing, product investment, customer migration, transaction protection or capital allocation.

Table 4. Board and investment-committee scorecard

Domain	Core measure	Warning signal	Decision response
Identity	Active agents with unique bound credentials	Shared keys or orphaned ownership	Suspend and reissue
Delegation	Chains with named principal and constrained scope	Missing actor or expanding scope	Redesign grant
Enforcement	Consequential operations at protected boundary	Advisory-only policy or bypass route	Migrate connector
Revocation	P95 withdrawal to last rejection	Descendant or queue remains active	Tighten lifetime and cascade
Evidence	Reconstructable decisions and effects	Missing policy version or resource outcome	Repair evidence pipeline
Economics	Contribution by protected workflow	Volume expands without monetised control	Reprice package
Transferability	Customer, key and trust relationships portable	Consent or credential reissue unresolved	Condition closing

Measures should be segmented by consequence tier and verified from production evidence.

43. CONCLUDE WITH ENFORCEABLE AUTHORITY

Autonomous systems create value when they can act. They create durable enterprise value when action is attributable, constrained, current, revocable and auditable. Identity establishes the actor. Delegation establishes the relationship to a principal. Policy determines whether the proposed operation is allowed. Enforcement controls the consequence. Revocation ends authority across the full chain.

An acquirer should begin with consequential operations and trace authority backward. Current value belongs to protected workflows, recurring contribution and control evidence that survives ownership change. Emerging standards can expand interoperability and reduce future integration cost. Their value should follow demonstrated adoption and remain separate from current economics.

The disciplined sequence is to register agents, bind workload identity, name the principal, structure the grant, attenuate each hop, enforce at the effect boundary, propagate withdrawal, preserve evidence, reconcile customer economics and structure consideration around verified outcomes. That approach allows a buyer to invest in autonomous capability while maintaining accountable control.

REFERENCES

[1] National Institute of Standards and Technology. New Concept Paper on Identity and Authority of Software Agents. 5 February 2026. https://www.nist.gov/news-events/news/2026/02/new-concept-paper-identity-and-authority-software-agents

[2] National Institute of Standards and Technology. Accelerating the Adoption of Software and AI Agent Identity and Authorization Concept Paper. February 2026. https://www.nccoe.nist.gov/sites/default/files/2026-02/accelerating-the-adoption-of-software-and-ai-agent-identity-and-authorization-concept-paper.pdf

[3] National Institute of Standards and Technology. Announcing the AI Agent Standards Initiative for Interoperable and Secure Innovation. 17 February 2026. https://www.nist.gov/news-events/news/2026/02/announcing-ai-agent-standards-initiative-interoperable-and-secure

[4] National Institute of Standards and Technology. AI Agent Standards Initiative. Updated 14 August 2026. https://www.nist.gov/artificial-intelligence/ai-agent-standards-initiative

[5] National Institute of Standards and Technology. Back to the Future: Why Agentic AI Needs a Strong Identity Foundation. 27 August 2026. https://www.nist.gov/blogs/cybersecurity-insights/back-future-why-agentic-ai-needs-strong-identity-foundation

[6] National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, NIST AI 600-1. July 2024. https://doi.org/10.6028/NIST.AI.600-1

[7] Jones, M.; Nadalin, A.; Campbell, B.; Bradley, J.; Mortimore, C. OAuth 2.0 Token Exchange, RFC 8693. January 2020. https://datatracker.ietf.org/doc/rfc8693/

[8] Lodderstedt, T.; Dronia, S.; Scurtescu, M. OAuth 2.0 Token Revocation, RFC 7009. August 2013. https://datatracker.ietf.org/doc/rfc7009/

[9] Richer, J. OAuth 2.0 Token Introspection, RFC 7662. October 2015. https://datatracker.ietf.org/doc/html/rfc7662

[10] OpenID Foundation. Authorization API 1.0 Final Specification. January 2026. https://openid.net/wg/authzen/

[11] OpenID Foundation. Three Shared Signals Final Specifications Approved. 2 September 2025. https://openid.net/three-shared-signals-final-specifications-approved/

[12] OpenID Foundation. Shared Signals Framework 1.0 Final. 2025. https://openid.net/specs/openid-sharedsignals-framework-1_0-final.html

[13] OpenID Foundation. Continuous Access Evaluation Profile 1.0 Final. 2025. https://openid.net/specs/openid-caep_1_0-final.html

[14] World Wide Web Consortium. Verifiable Credentials Data Model v2.0. W3C Recommendation, 15 May 2025. https://www.w3.org/TR/vc-data-model-2.0/

[15] Liu, D.; Zhu, H.; Krishnan, S. Agent Operation Authorization, Internet-Draft draft-liu-agent-operation-authorization-02. 16 March 2026. https://datatracker.ietf.org/doc/html/draft-liu-agent-operation-authorization-02

[16] IETF WIMSE Working Group. AI Identity Management System, Internet-Draft draft-ietf-wimse-aims-00. September 2026. https://datatracker.ietf.org/doc/draft-ietf-wimse-aims/00/

[17] Asor, D. Verifiable Attenuated Delegation for AI Agent Chains, Internet-Draft draft-asor-wimse-agent-delegation-chain-01. September 2026. https://datatracker.ietf.org/doc/draft-asor-wimse-agent-delegation-chain/01/

[18] Chen, L. et al. Agent Authorization Use Cases and Gap Analysis, Internet-Draft draft-chen-oauth-agent-authz-use-cases-01. July 2026. https://datatracker.ietf.org/doc/draft-chen-oauth-agent-authz-use-cases/01/

[19] Zhu, H. et al. OAuth 2.0 Asynchronous Delegation, Internet-Draft draft-zhu-oauth-async-delegation-05. August 2026. https://datatracker.ietf.org/doc/html/draft-zhu-oauth-async-delegation-05

[20] European Parliament and Council. Regulation EU 2024/1689 laying down harmonised rules on artificial intelligence. 13 June 2024. https://eur-lex.europa.eu/eli/reg/2024/1689/oj

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.