Identity for Every Machine: Building Roll-Ups around Agents, APIs and Workloads

An Acquisition and Valuation Framework for Policy Depth, Federation, Enterprise Distribution and Integration

Authored by Chennakeshav Adya

Independent Researcher

September 2026

Abstract

Enterprises increasingly depend on software workloads, APIs, service accounts, deployment pipelines, robotic automation and AI agents that act without a person at the keyboard. Each machine actor needs a verifiable identity, bounded authority, short-lived credentials, policy enforcement and an audit trail. Fragmented controls create dormant accounts, shared secrets, excessive privilege, inconsistent trust domains and unclear accountability. These weaknesses can grow when an acquirer combines products that use different identity models.

This paper develops an acquisition and valuation framework for machine-identity businesses. The proposed unit of value is a verified and authorised machine action delivered inside a customer workflow at complete cost. The framework examines discovery, identity issuance, attestation, authentication, authorisation, credential lifecycle, federation, observability, governance and enterprise distribution. It treats AI agents as a demanding extension of machine identity because an agent may select tools, delegate work and change its actions in response to context.

NIST zero-trust guidance requires authentication and authorisation before access and rejects implicit trust based on network location. NIST guidance for cloud-native applications emphasises application and service identities, API gateways, sidecars and standards including SPIFFE. The SPIFFE specifications define workload identities, verifiable identity documents and trust-domain federation. Kubernetes recommends bound, time-limited service-account tokens instead of long-lived token secrets. OAuth standards provide token exchange, mutual-TLS-bound tokens, proof of possession, metadata, introspection and revocation mechanisms that can support controlled API access.[1][2][3][4][5][6][7][8]

A hypothetical acquisition illustrates a platform serving regulated enterprises, cloud-native software companies and infrastructure operators. Every revenue, customer, cost, performance, probability and valuation figure in that illustration is a management assumption created solely to demonstrate the method. It is neither a forecast nor a market benchmark.

The analysis concludes that a buyer should price policy depth, evidence continuity and enterprise adoption before assigning value to the number of identities discovered. Six figures and seven tables translate that conclusion into diligence tests, a valuation bridge, consideration protection and a 180-day integration programme.

Cybersecurity, privacy, employment, competition, foreign-investment, accounting, tax, insurance and securities decisions require current advice from qualified specialists in each relevant jurisdiction. This paper provides general information for professional audiences and does not provide legal, regulatory, technical, accounting, tax or investment advice.

JEL Classification: G24, G34, L86, O32, O33

Keywords: machine identity, workload identity, API security, AI agents, zero trust, cybersecurity M&A, roll-up strategy, valuation

1. Define the acquisition decision

The board should define the customer decision that the target improves. Machine-identity firms may discover non-human accounts, issue workload credentials, broker secrets, authorise API calls, govern cloud roles, secure deployment pipelines, manage certificates, monitor service-to-service activity or control AI-agent tools. These activities address related risks while producing different evidence, economics and integration obligations.

The acquisition thesis should name the intended source of value: proprietary policy technology, enterprise distribution, regulated-customer access, a scalable trust service, identity telemetry, scarce engineering capability or a platform for consolidation. Each source requires a reproducible test. Discovery claims need population evidence. Policy claims need denied and permitted action tests. distribution claims need contracted adoption, retention and collections.

The board should compare acquisition with partnership, licensing, minority investment and internal development. Ownership may matter when value requires coordinated control of credential issuance, policy engines, customer integrations and sensitive telemetry. A commercial arrangement may be more proportionate when interoperability or channel access supplies most of the benefit.

Evidence timing should shape terms. Pre-signing tests can reproduce protocol support, credential rotation and policy decisions in a controlled environment. Customer-specific coverage and integration economics may require post-close access. Base consideration should follow evidence available at signing; contingent value should follow verified milestones.

Figure 1. Machine-action-to-value evidence chain

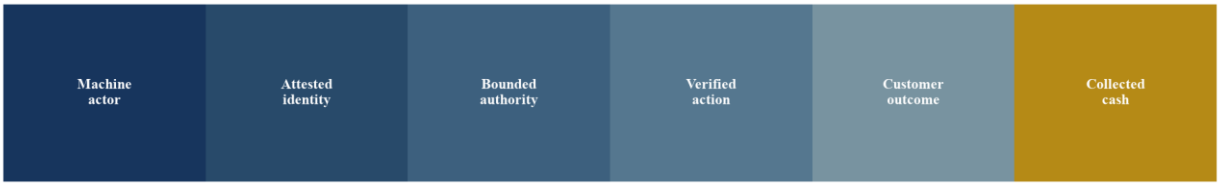

The proposed chain connects an identified machine actor to an authorised action, customer outcome and collected cash.

2. Define the unit of value

The proposed unit of value is a verified and authorised machine action delivered inside a customer workflow at complete cost. An action may retrieve data, invoke an API, deploy code, rotate a key, approve an automated step or delegate a task. The record should identify the actor, workload, environment, requested resource, policy, credential, decision, response and accountable owner.

Complete cost includes discovery, attestation, credential issuance, cryptographic operations, policy evaluation, telemetry, storage, third-party licences, support, integration, security operations, incident response, compliance and working capital. A platform can report software margins while implementation teams manually reconcile identities or customers retain parallel tools. The acquisition model should include every activity needed to produce the promised control.

Identity count is an incomplete denominator. One discovered service account may create no value if it remains unmanaged. A short-lived credential can still confer excessive authority. A policy decision can be technically correct while the customer workflow ignores it. Buyers should therefore measure verified actions, effective policy coverage, prevented unauthorised actions, investigation time and customer operating cost.

3. Map the machine-identity perimeter

The perimeter includes cloud workloads, containers, virtual machines, serverless functions, APIs, service accounts, certificates, devices, CI/CD jobs, bots, robotic automation and AI agents. Each actor has a different creation event, owner, runtime, credential, privilege pattern and termination signal. A single inventory number can conceal these differences.

The diligence team should map every product module to the actors it can discover, identify and govern. Coverage should be tested across cloud providers, orchestration systems, operating systems, development platforms and legacy environments. Unsupported populations should remain visible.

The target should distinguish identity from credential. An identity represents an actor and its attributes. A credential proves possession or control under defined conditions. Several credentials may represent one identity, and one shared secret may obscure several actors. Consolidation should reduce ambiguity instead of moving it into a central vault.

Table 1. Machine-identity perimeter and acquisition tests

Actor	Typical credential	Required evidence	Acquisition warning
Workload	short-lived certificate or token	attested runtime and owner	static credential presented as workload identity
API client	token, key or certificate	client, scope and resource binding	shared key with weak attribution
CI/CD job	federated token	repository, workflow and run context	reusable deployment secret
Service account	platform token or secret	owner, purpose and expiry	dormant account with persistent privilege
AI agent	delegated token and policy	principal, tools, task and approval	broad authority without action-level audit
RPA bot	application account	process, operator and target system	human credential reused by automation

Each actor requires a distinct lifecycle and evidence model.

4. Build the identity-action ledger

The ledger should connect discovery source, actor, owner, workload attestation, trust domain, credential, policy, resource, action, decision, exception, customer outcome and financial record. It should preserve both allowed and denied actions. The purpose is to trace product capability into customer value.

Negative evidence belongs in the ledger. Orphaned identities, failed rotations, stale policies, policy bypasses, unavailable telemetry and manual overrides reveal the real control boundary. A data room containing only successful demonstrations cannot establish population coverage.

Finance should link customer cohorts to deployed identities, governed actions, implementation effort, support cost, renewal, expansion and collections. This allows the buyer to test whether deeper policy use improves retention and contribution or creates unpriced service work.

5. Test discovery coverage and ownership

Discovery should begin with an independently defined population. The buyer should reconcile cloud directories, orchestration platforms, certificate stores, secrets systems, API gateways, code repositories, deployment systems and network telemetry. The target's own inventory should then be compared with that population.

Coverage must be reported by environment and actor type. A high aggregate percentage can conceal weak coverage in production clusters or privileged deployment accounts. False positives also matter because an unusable inventory increases remediation work and weakens customer trust.

Ownership evidence should identify a responsible team, approved purpose, data access and termination trigger. Machine identities often outlive the application or employee that created them. The product should support re-attestation and escalation when ownership becomes unclear.

Population construction requires care. Cloud control planes, application logs and source repositories observe different parts of the estate and at different times. The buyer should define a measurement window, deduplicate stable identifiers and retain the source of every finding. Ephemeral workloads may appear only briefly, while dormant privileged accounts may generate no traffic. A discovery engine that relies exclusively on activity can miss dangerous inactive credentials; a directory-only engine can report identities that no longer reach a resource.

The buyer should run seeded tests. Approved test identities can be created across representative platforms with known owners, privileges, credential forms and lifetimes. The diligence team can then measure detection, classification, ownership assignment and remediation. Seeded identities should include ambiguous and adversarial cases, such as copied names, shared labels and misleading metadata. Results should be reproduced without seller intervention.

Remediation evidence should extend beyond a ticket. The record should show whether an identity was disabled, re-scoped, rotated, assigned or accepted as an exception; whether the application continued to operate; and whether the change persisted. Reappearing identities can indicate automated recreation, incomplete infrastructure-as-code changes or a disconnected source system. Durable remediation is more valuable than a high volume of closed findings.

Coverage claims also need temporal testing. A one-time scan can produce an attractive baseline while missing daily creation and deletion. The buyer should measure time from identity creation to discovery, owner notification, policy attachment and resolution. Tail latency can reveal coverage gaps that an average obscures. This operating cadence affects customer benefit, support load and renewal.

Figure 2. Hypothetical discovery and false-positive frontier

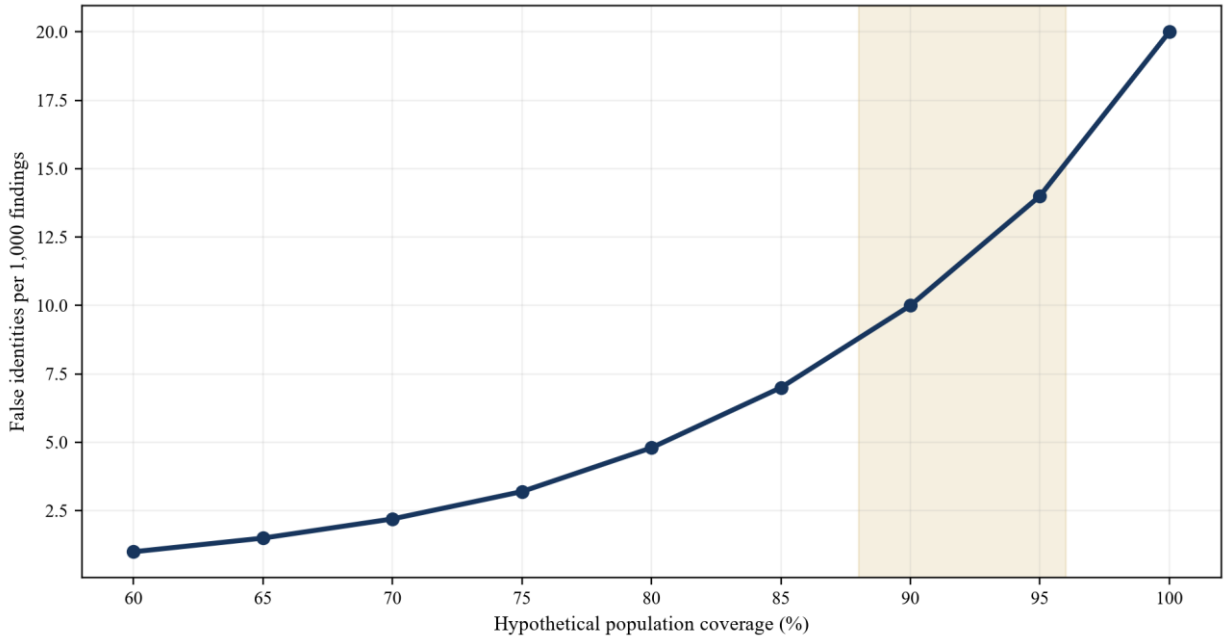

Values are management assumptions for method demonstration.

6. Test identity issuance and attestation

Workload identity should be issued only after evidence connects the running workload to an approved environment and owner. SPIFFE defines workload identities and verifiable identity documents; its Workload API provides identities without requiring applications to handle authentication secrets directly.[4][9]

The buyer should inspect node, workload and process attestation. Tests should attempt to obtain an identity from an unauthorised node, altered image, copied configuration and adjacent workload. The system should record the evidence used, issuer, validity period and revocation path.

Attestation dependencies belong in the valuation model. Cloud metadata, orchestration controls, hardware roots, certificate authorities and third-party services can create concentration or portability risk. The target should show fallback, migration and incident procedures.

7. Separate authentication from authorisation

Authentication establishes which actor presents a credential. Authorisation determines whether that actor may perform a specific action on a specific resource under current conditions. A platform that authenticates every workload can still permit excessive or unintended activity.

Policy depth should be measured from network or role-level access through resource, action, data and context controls. Useful context can include workload state, environment, time, risk, data classification, requested tool and human approval. The product should explain which attributes are authoritative and how conflicts are resolved.

The buyer should test policy decisions under changed context and failure. It should examine default behaviour when the policy service, attribute source or audit system is unavailable. Availability achieved through permissive fallback can transfer operational resilience into security exposure.

Table 2. Policy-depth maturity model

Level	Control scope	Evidence	Value limitation
1	inventory only	discovered actor	no enforcement
2	credential control	issuance and rotation	authority may remain broad
3	resource access	allow or deny record	limited action context
4	action and data	method, object and scope	integration complexity
5	contextual delegation	task, risk and approval	governance and latency burden

Valuation should follow effective control and evidence rather than policy count.

8. Measure credential half-life

Short-lived credentials reduce the period in which a copied credential remains useful. Kubernetes recommends bound, time-limited service-account tokens and advises against long-lived token secrets. OAuth proof-of-possession mechanisms can bind tokens to a client certificate or key.[6][10][11]

The buyer should measure median and tail validity periods, rotation success, emergency revocation and residual static secrets. It should test whether applications refresh credentials safely and whether revocation reaches distributed enforcement points.

Credential duration should reflect operational recovery. Extremely short validity creates little benefit if outages lead teams to install persistent emergency secrets. The relevant measure is effective exposure after issuance, compromise, rotation, revocation and exception handling.

Figure 3. Hypothetical credential exposure over time

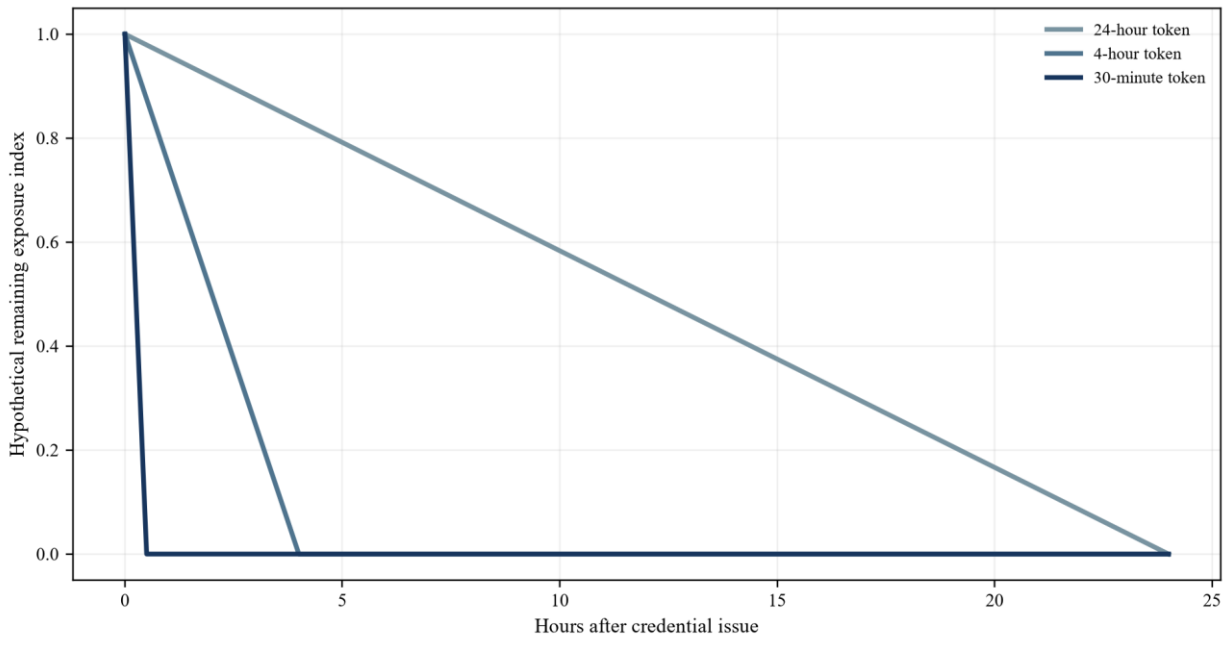

Curves illustrate relative exposure under different validity and revocation designs; values are management assumptions.

9. Test federation across trust domains

Federation allows an identity established in one trust domain to be accepted under policy in another. SPIFFE federation exchanges trust bundles and binds them to trust domains. OAuth token exchange supports obtaining a token for a different service or security domain.[5][7]

The buyer should test trust establishment, bundle distribution, issuer restrictions, audience binding, claim mapping and revocation. Cross-domain access should require explicit policy. A configuration change that silently broadens trust can create systemic exposure.

Commercial value depends on interoperability. Customers operate multiple clouds, clusters, software vendors and acquired estates. Proprietary federation can increase switching cost while restricting adoption. Standards support can widen distribution; implementation quality, governance and customer workflow still determine differentiation.

10. Evaluate API identity and token exchange

API access should bind a client, token, audience, scope, resource and action. OAuth standards support server metadata, protected-resource metadata, token introspection, revocation and token exchange. Mutual TLS and DPoP can reduce bearer-token replay by proving possession of a bound key.[7][10][11][12][13][14]

The diligence team should replay tokens against unintended resources, alter audience and scope, test expired and revoked tokens, and examine behaviour when introspection or metadata is unavailable. Logs should allow a customer to reconstruct the decision.

API-key management alone should not be valued as comprehensive machine identity. The buyer should identify how the product moves customers from shared keys to attributable, time-limited and policy-bound access without disrupting production systems.

11. Govern AI-agent identity and delegation

AI agents can select tools and sequence actions in response to data. NIST's 2026 concept paper on software and AI-agent identity asks how agents should be identified, authorised, audited and linked to human authority.[15] The acquisition thesis should treat agent identity as an extension of enterprise control, with additional delegation and unpredictability risks.

Every agent action should connect to an owning organisation, approved agent version, initiating principal, task, permitted tools, resource scope, time window and escalation rule. Delegated authority should narrow as work passes between agents. A receiving service should not assume that an agent may exercise every privilege held by its human sponsor.

Prompt content should not become an unverified source of authority. Policy should be evaluated outside the model against authenticated context. High-consequence actions may require deterministic checks, dual control or human approval. The audit trail should preserve inputs, tool calls, policy decisions and outcomes while respecting privacy and data-minimisation requirements.

Agent identity also changes the meaning of session duration. A conventional service may execute a narrow function repeatedly, while an agent can remain active across a sequence of planning, retrieval, generation and execution steps. The buyer should test whether authority is re-evaluated at each sensitive step, when the task changes and when the agent receives new data. A single approval at session start should not silently authorise an unrelated later transaction. Delegation records should show the maximum authority available, the authority actually used and the reason for each elevation.

Model and tool versions belong in the identity record. A policy approved for one tool interface or model behaviour may not remain appropriate after an update. Release governance should connect the deployed model, prompt package, tool schema, policy bundle and evaluation result. The target should demonstrate rollback, retirement and residual-access testing. These controls allow a buyer to distinguish an experimental agent wrapper from an enterprise control plane that can support regulated adoption.

Table 3. AI-agent delegation tests

Test	Expected control	Evidence
Tool substitution	unapproved tool denied	policy decision and alert
Scope expansion	broader resource denied	audience and scope record
Agent hand-off	authority narrows	delegation chain
Prompt injection	instruction cannot grant privilege	external policy result
High-value action	approval required	approver and transaction record
Agent retirement	credentials and access end	revocation and residual-access test

Each test connects authority to a traceable business task.

12. Test observability and non-repudiation

The product should produce records sufficient to answer who acted, under which identity, with what authority, against which resource and with what result. Logs should include policy and credential versions so a later review can reproduce the decision.

Integrity and access controls matter because machine-identity telemetry can expose architecture and secrets. The buyer should inspect collection gaps, clock synchronisation, retention, export, signing, customer ownership and incident preservation. A polished dashboard cannot compensate for missing source evidence.

Operational metrics should include policy latency, denied-action rate, exception age, credential failures, orphaned identities and investigation time. Measures should be segmented by customer cohort and environment.

13. Review key, secret and certificate management

NIST key-management guidance addresses policies, procedures, planning and cryptographic key-management systems.[16][17] A machine-identity platform should define generation, storage, distribution, rotation, revocation, backup, recovery and destruction for each credential class.

The buyer should map custody and administrator access. It should inspect hardware-security-module use, certificate-authority hierarchy, emergency access, export controls and separation of duties. Customer-managed and vendor-managed models create different liability and gross-margin profiles.

Secrets migration is a major integration risk. Acquiring a vault or certificate product does not automatically create unified identity. The plan should preserve service continuity while reducing duplicated credential stores and privileges.

14. Evaluate product architecture and dependencies

The architecture should separate control plane, data plane and evidence plane. Policy administration can be central while enforcement remains close to workloads. This design can reduce latency and maintain operation during control-plane disruption, provided that cached policy and failure modes are governed.

The buyer should inventory open-source components, cloud services, protocol libraries, certificate authorities, databases and observability dependencies. Licence rights, maintenance status and replacement cost belong in diligence.

Scalability tests should reflect peak authentication and policy traffic, tenant isolation, certificate rotation events and incident conditions. Average request volume can conceal operational failure during a widespread expiry or emergency revocation.

The control plane should maintain authoritative configuration, approval and policy history. Distributed enforcement points should receive signed, versioned policy and expose their applied state. The evidence plane should record enough information to reconcile a decision without storing unnecessary secrets or customer data. The buyer should test consistency when network partitions, delayed updates and rollbacks occur.

Multi-tenant architecture requires explicit isolation of policy, identity namespaces, trust bundles, telemetry and administrator access. Tests should attempt cross-tenant reference, identifier collision, policy import and support-access misuse. Customer-controlled encryption or dedicated deployment options may improve regulated-market access while increasing cost and release complexity. These economics should be visible by cohort.

Data residency can influence architecture and transaction value. Identity telemetry may reveal service names, routes, privileges and operational patterns. The buyer should map collection, processing, support access, backup and disaster recovery by jurisdiction. Contractual promises should match actual routing and subprocessors. Any planned consolidation of regional systems should be costed and reviewed before synergy is recognised.

The acquisition team should examine developer experience because adoption depends on integration quality. Software development kits, command-line tools, policy testing, local development, migration helpers and error messages can determine time to value. Documentation should distinguish secure defaults from optional controls. A product that requires extensive bespoke engineering may still serve valuable customers, but its contribution and scalability should be modelled accordingly.

Release engineering is part of the control. Changes to protocol libraries, policy evaluation, certificate handling and agents can affect every customer. The target should show code review, dependency monitoring, signed builds, staged rollout, compatibility testing and emergency rollback. NIST's Secure Software Development Framework provides a useful reference for examining these practices.[36]

Figure 4. Proposed roll-up reference architecture

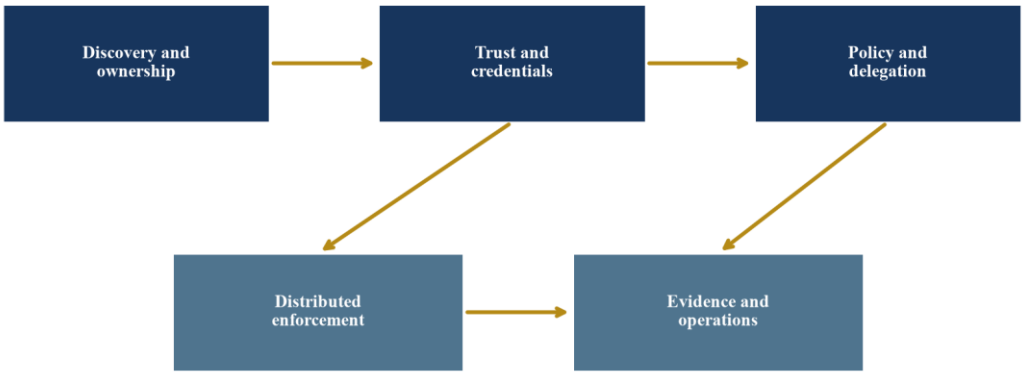

The design separates discovery, trust, policy, enforcement and evidence while preserving customer control points.

15. Test security and abuse resistance

The target itself is privileged infrastructure. Compromise can issue trusted credentials, alter policy or suppress evidence. The buyer should perform architecture review, code review, penetration testing, build-pipeline review and privileged-access analysis proportionate to the risk.

Threat scenarios should include issuer compromise, stolen signing keys, malicious administrators, tenant escape, policy tampering, metadata spoofing, token replay, dependency compromise and denial of service. Each scenario needs prevention, detection, containment and recovery evidence.

The seller's incident history should be reconciled across ticketing, security monitoring, customer notifications, insurers and regulators. Absence of reported incidents is not equivalent to absence of compromise.

16. Diligence customer cohorts and distribution

Enterprise distribution can be a principal source of roll-up value. The buyer should segment customers by industry, environment, actor population, deployed modules, policy depth, contract term, implementation model, support burden, renewal and collections.

Signed contracts should be reconciled with deployment evidence. Shelfware and limited pilots should not receive the same valuation as enforced production policy. Usage should show sustained governed actions across relevant environments.

Channel partnerships require evidence of sourced pipeline, conversion, economics and customer control. A cloud marketplace listing or technical integration can support distribution; it does not establish customer demand by itself.

The buyer should reconstruct the implementation funnel from signed order through discovery, first credential, first enforced policy, production coverage and steady-state use. Delays between these stages can consume cash and increase churn even when contracted revenue appears strong. Cohort analysis should report time to first control, time to target coverage, implementation hours and exceptions that remain open after launch.

Expansion should be separated into price, identity volume, additional environments and deeper policy adoption. Volume growth can reflect infrastructure growth without improved security value. Deeper adoption can increase switching cost and customer outcome, but may require more engineering and support. Net retention should therefore be read alongside contribution and policy depth.

Distribution rights and customer consent affect roll-up integration. Contracts may restrict data transfer, subcontracting, changes to hosting or assignment after a change of control. Customer telemetry can also contain sensitive architecture information. Legal, product and commercial teams should identify consents, localisation duties and customer-controlled encryption before assuming that identities and policies can be moved to a common platform.

Table 4. Customer-cohort evidence matrix

Cohort	Deployment evidence	Economic test	Principal risk
Regulated enterprise	production policy and audit export	retained recurring contribution	long implementation cycle
Cloud-native scale-up	workload and API coverage	expansion and support efficiency	vendor consolidation
Infrastructure operator	resilient enforcement	contract duration and collections	operational liability
AI-agent adopter	action-level delegation	paid production use	immature governance

Cohort	Deployment evidence	Economic test	Principal risk
Channel-led customer	partner-sourced deployment	net revenue and control	dependence on intermediary

Cohorts should be valued according to adoption, contribution and durability.

17. Reconstruct complete delivery economics

Revenue should be reconciled from contract through invoice and bank receipt. The buyer should separate subscription, consumption, implementation, managed service and third-party pass-through. Reported annual recurring revenue should exclude non-recurring and unsupported amounts.

Cost should include cloud processing, certificate and key services, telemetry storage, support, customer engineering, policy design, incident response and partner share. Customer contribution should be calculated after the support required to maintain effective control.

Working capital matters where large customers pay slowly while the target funds infrastructure and implementation. The acquisition model should connect growth, deployment, billing, collections and cash.

The buyer should rebuild gross margin from source records rather than relying solely on financial-statement classification. Engineering labour assigned to customer configuration, recurring policy maintenance or incident support may sit in research and development while functioning as cost of service. Partner credits and committed cloud discounts can temporarily improve reported margin. Normalisation should retain costs necessary to deliver the current promise.

Unit economics should use customer cohorts and activity drivers. Useful denominators include production environments, governed actions, enforcement points, telemetry volume and support hours. Cost per identity can mislead when identities vary widely in activity and consequence. The team should identify which driver explains marginal infrastructure and human effort.

Pricing should be tested against customer value and cost volatility. Per-identity pricing is simple but can discourage full discovery. Per-action pricing can align with use but expose customers to uncertain bills. Enterprise subscription can support broad adoption while transferring volume risk to the vendor. Contracts should be analysed for minimum commitments, overages, indexation, service credits, termination rights and limits on price changes after acquisition.

Retention analysis should distinguish logo retention, recurring-revenue retention and retained contribution. A customer can expand reported revenue while support and infrastructure costs rise faster. The valuation case should use the measure that best connects continuing customer value to cash. Collections, disputes and credits should be reconciled to the same cohort record.

Sales efficiency needs a full-cycle view. Regulated customers may require security review, proof of concept, procurement, legal negotiation and phased deployment. The buyer should measure cash acquisition cost, sales-cycle duration, implementation capacity and payback from initial spending through collected contribution. Pipeline should be weighted by completed customer evidence rather than seller stage labels.

18. Build a hypothetical acquisition case

Assume a target with USD 18.0 million recurring revenue, USD 3.0 million implementation revenue and USD 1.0 million other revenue. Management estimates USD 12.2 million retained recurring contribution after direct delivery and support. The largest ten customers represent 44 per cent of recurring revenue. These figures are hypothetical.

Evidence review attributes USD 7.0 million of retained recurring contribution to customers using production policy enforcement, USD 3.2 million to customers using credential and discovery modules,

and USD 2.0 million to pilots or limited deployments. The buyer assigns different confidence and integration requirements to each layer.

Management identifies USD 2.4 million of potential annual cross-sell contribution and USD 1.6 million of duplicated cost. The base valuation excludes both until customer acceptance and implementation evidence exist. Contingent consideration can recognise realised cross-sell without capitalising an unproven plan at signing.

Table 5. Hypothetical evidence-layered contribution

Layer	Retained contribution	Evidence status	Valuation treatment
Production policy customers	7.0	deployed and renewed	base case subject to retention
Credential and discovery customers	3.2	deployed with limited policy	migration-adjusted
Pilots and limited deployment	2.0	incomplete adoption	contingent or option value
Potential cross-sell	2.4	management plan	excluded from base price
Duplicated cost opportunity	1.6	integration estimate	recognised after delivery

All amounts are management assumptions in USD millions.

19. Stress the operating model

Stress tests should combine technical and commercial events. Relevant cases include credential-service outage, issuer compromise, cloud-platform change, customer loss, slower deployment, higher support cost and delayed cross-sell. Correlated events deserve specific attention because a security incident can increase cost while reducing renewal.

The buyer should model liquidity as well as earnings. Emergency rotation, customer remediation, forensic work and insurance deductibles can require cash before revenue recovers. Covenants and earn-out metrics should remain workable under those conditions.

Stress design should start from causal links. An issuer compromise can trigger emergency certificate replacement, customer downtime, service credits, investigation cost, delayed sales and churn. Treating each effect as independent can understate the combined event. The model should specify timing, cash payment, insurance recovery assumptions and management responses. Insurance should be recognised only to the extent supported by policy terms and claims analysis.

Platform dependency stress should examine changes to cloud identity services, orchestration APIs, certificate lifetimes and browser or runtime trust stores. The target should show how quickly it can adapt, which customers require manual intervention and whether older versions remain supported. Contractual service obligations can continue while a third-party change raises delivery cost.

Customer-concentration stress should incorporate operational concentration. Several customers may share the same cloud, channel partner or implementation architecture, creating correlated exposure. Revenue diversification by logo can therefore overstate resilience. The buyer should map concentration by customer, platform, region, partner, issuer and product module.

Management responses should be feasible and sequenced. Cost reduction can protect liquidity while slowing remediation and product migration. Price increases can support margin while weakening renewal. The board should review central, adverse and severe cases with explicit triggers for liquidity preservation, customer communication, additional security capacity and covenant engagement.

The stress pack should state which assumptions are contractual, observed, management-estimated or scenario-only. It should retain the source, owner and approval date for each material input. Post-close results should be compared with the original cases each month so that management can identify whether variance arises from customer behaviour, technical performance, integration execution or financial assumptions. This discipline also improves the evidence available for contingent consideration, impairment review and the next acquisition.

Independent challenge should focus on the assumptions that drive liquidity, customer harm and irreversible platform decisions, with unresolved matters reported directly to the transaction committee.

Figure 5. Hypothetical retained-contribution stress bridge

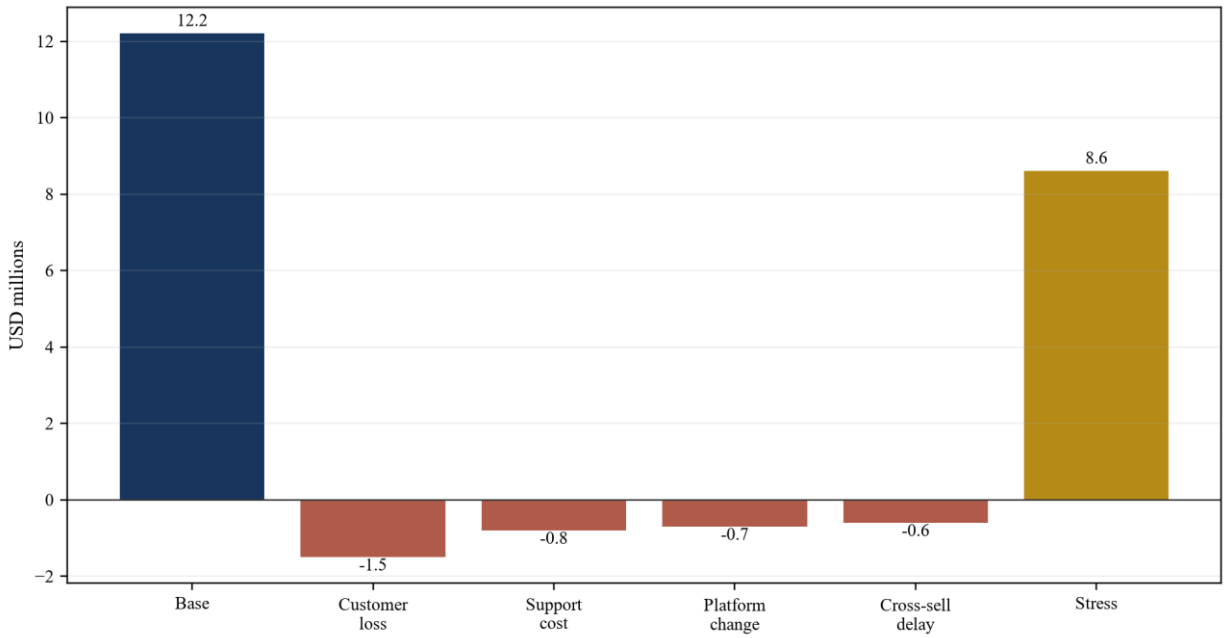

All values are management assumptions in USD millions.

20. Value the evidence layers

Valuation should begin with retained recurring contribution supported by contracts, deployment and cash. The buyer can apply a required return or multiple consistent with growth, retention, concentration, security exposure, delivery economics and capital needs. Headline market multiples should not replace company-specific evidence.

The bridge should separate contracted production value, migration-dependent value, contingent adoption, integration synergies and strategic options. Each layer should have an owner, milestone, cost and downside case. This prevents the same benefit from appearing in both the seller's forecast and buyer synergy case.

Purchase-price allocation under IFRS 3, IAS 38 and IFRS 13 may identify customer relationships, technology, brands and other assets separately from goodwill. Impairment assessment under IAS 36 depends on the applicable accounting facts and advice.[18][19][20][21]

The valuation model should make evidence expiry visible. Customer contribution can weaken at renewal, technical evidence can decay after platform changes, and integration assumptions can fail when migration begins. Each material layer should have a review date and a downside response. A static terminal value based on current identity growth can overstate durability where standards, cloud platforms or customer architectures change.

Strategic option value should be stated separately. An installed policy engine may support future agent governance, but the buyer should identify the additional product, regulatory, sales and capital requirements before attaching value. Options can justify a transaction route or a limited investment while remaining outside the price supported by current cash flows.

Comparable-company evidence should be normalised for revenue definition, services content, growth, retention, concentration, stock-based compensation and cash burn. Transactions completed under different interest-rate, cybersecurity or capital-market conditions require further adjustment. The valuation committee should retain a traceable bridge from observable market evidence to the company-specific conclusion.

Table 6. Hypothetical enterprise-value bridge

Component	Evidence basis	Hypothetical value USDm
Contracted production contribution	deployed, renewed and collected	72.0
Migration-dependent contribution	credential and discovery customers	18.0
Adoption option	pilots and agent use cases	6.0
Cost synergy after delivery	verified integration milestones	8.0
Security and concentration reserve	downside adjustment	-14.0
Illustrative enterprise value	sum of evidence layers	90.0

Amounts and valuation factors are management assumptions for method demonstration.

Figure 6. Hypothetical evidence-layered enterprise value

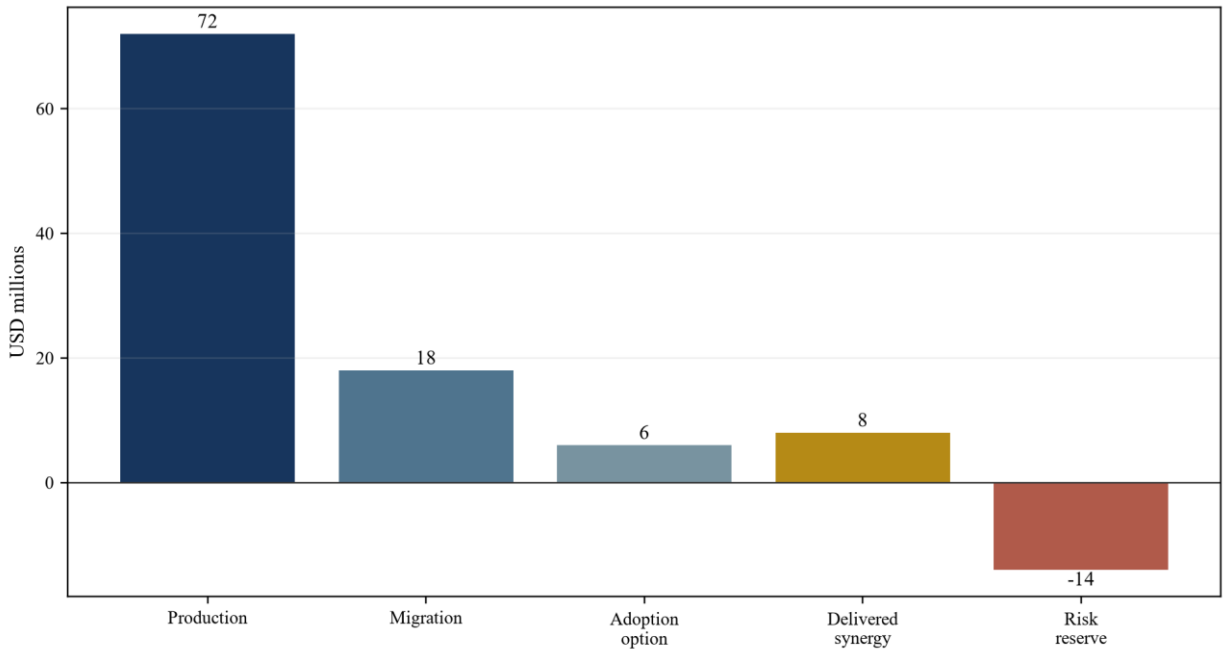

Values are management assumptions in USD millions and do not represent a market benchmark.

21. Structure consideration and integration

Base consideration should reflect reproduced technology, transferable rights, contracted contribution and cash. Deferred or contingent consideration can address customer migration, policy adoption, key-person retention and security remediation. Metrics should be objective, controllable and resistant to accounting-policy changes.

Representations and warranties should address intellectual property, open-source use, security incidents, credential custody, customer commitments, data rights and compliance. Specific indemnities or escrow may be appropriate where identified exposures cannot be resolved before closing, subject to legal advice.

Integration should preserve enforcement continuity. The buyer should avoid forced migration before identity mapping, policy equivalence, rollback and customer approval are tested. Product rationalisation should follow evidence rather than an assumed single-platform end state.

Table 7. Consideration and integration gates

Gate	Evidence	Transaction response
Technology	reproduced identity and policy tests	supports base value
Rights	transferable code, data and licences	closing condition or remediation
Customers	retained production contribution	deferred consideration
Security	key custody and incident review	escrow, indemnity or condition
Migration	policy equivalence and rollback	phased integration
Synergy	collected cross-sell and delivered cost	contingent value after realisation

The structure links payment and migration to observable evidence.

22. Execute a 180-day programme

Days 0 to 30 should establish control. The buyer should confirm privileged access, issuer and key custody, incident response, customer escalation, identity inventories and integration decision rights. It should freeze high-risk architectural changes until evidence is preserved.

Days 31 to 60 should reproduce discovery, issuance, rotation, policy, federation and agent-delegation tests. Finance should reconcile revenue, contribution and collections by cohort. Legal and technical teams should confirm rights and critical dependencies.

Days 61 to 100 should define the product and distribution architecture. Teams should map equivalent policies, trust domains, telemetry, customer contracts and support obligations. Pilot migrations should include rollback and customer acceptance.

Days 101 to 180 should scale validated migrations, launch approved cross-sell, remove duplicated controls and report benefits against the signed baseline. The board should receive a monthly evidence pack covering security, customers, economics, integration and cash.

23. Decision and conclusion

Machine identity creates acquisition value when a platform can identify actors, bind short-lived credentials, enforce action-level authority, federate trust and preserve evidence inside customer operations. Inventory volume and protocol claims are starting points.

A successful roll-up requires an explicit trust architecture. Combining products without reconciling issuers, policies, customer ownership and enforcement can increase complexity and systemic exposure. Integration should proceed through reproduced equivalence and controlled migration.

The proposed framework links technical controls to customer outcomes and retained contribution. It prices verified production value, treats migration and adoption as evidence-dependent layers, protects consideration and gives management a 180-day execution sequence.

The board's approval paper should contain a short set of auditable conditions: the population against which discovery was tested; the credentials and policies reproduced; the customer contribution reconciled to cash; the rights and dependencies confirmed; the security exceptions accepted; and the milestones that

govern payment and integration. These conditions convert a broad strategic narrative into a transaction that management can monitor.

Machine identity is likely to span several existing security budgets, including secrets, certificates, cloud permissions, API access, developer security and agent governance. A roll-up can create customer value when it reduces duplicated control and produces a consistent evidence chain. It can destroy value when consolidation removes local context, adds a privileged central dependency or forces migration before equivalence is proven. The proposed gate sequence preserves customer operations while allowing the buyer to earn the platform outcome through completed evidence.

Management should continue to measure the acquisition after the initial integration period. Renewal, policy depth, exception age, credential exposure, incident response, contribution and cash should remain linked by cohort. This continuing record supports product decisions, impairment review, additional acquisitions and eventual exit diligence.

Frequently asked questions

What is a machine identity?

A machine identity is a verifiable representation of a software workload, service, API client, device, automation process or AI agent. It should connect the actor to an owner, approved purpose, environment, credential lifecycle and permitted authority.

Why is identity count a weak valuation metric?

Discovered identities may remain unmanaged, duplicated or unauthorised. Valuation evidence improves when identities use governed credentials and policies that produce measurable customer outcomes at complete cost.

How should a buyer evaluate policy depth?

The buyer should test control from broad resource access through action, data and contextual delegation. Tests should use changed conditions, service failures and attempted scope expansion, with reproducible decision records.

What makes AI-agent identity different?

An AI agent may choose tools, sequence actions and delegate tasks. The control system should connect every action to an owning principal, agent version, task, policy, resource and approval requirement outside the model prompt.

Why do short-lived credentials matter?

Short validity can reduce the useful period of a copied credential. Effective control also requires secure issuance, automatic rotation, rapid revocation, proof of possession and removal of persistent fallback secrets.

How should federation affect an acquisition thesis?

Federation can extend enterprise reach across clouds and acquired estates. The buyer should test explicit trust, claim mapping, audience binding, revocation and configuration governance before pricing interoperability benefits.

How should synergies be valued?

Potential cross-sell and cost reduction should remain outside base price until customer adoption, collected contribution and delivered integration are evidenced. Contingent consideration can recognise realised value.

What should happen during the first 180 days?

Management should secure issuers and privileged access, reproduce product evidence, reconcile customer economics, define the trust architecture, run controlled migrations and report benefits against an approved baseline.

References

[1] NIST. Zero Trust Architecture, SP 800-207. 2020. https://csrc.nist.gov/pubs/sp/800/207/final

[2] NIST. A Zero Trust Architecture Model for Access Control in Cloud-Native Applications, SP 800-207A. 2023. https://csrc.nist.gov/pubs/sp/800/207/a/final

[3] NIST. Implementing a Zero Trust Architecture, SP 1800-35. 2025. https://csrc.nist.gov/pubs/sp/1800/35/final

[4] SPIFFE. SPIFFE Specifications. 2026. https://spiffe.io/docs/latest/spiffe-specs/

[5] SPIFFE. SPIFFE Federation. 2026. https://spiffe.io/docs/latest/spiffe-specs/spiffe_federation/

[6] Kubernetes. Service Accounts. 2026. https://kubernetes.io/docs/concepts/security/service-accounts/

[7] IETF. OAuth 2.0 Token Exchange, RFC 8693. 2020. https://datatracker.ietf.org/doc/html/rfc8693

[8] NIST. Security Strategies for Microservices-based Application Systems, SP 800-204. 2019. https://csrc.nist.gov/pubs/sp/800/204/final

[9] SPIFFE. Workload API. 2026. https://spiffe.io/docs/latest/spiffe-specs/spiffe_workload_api/

[10] IETF. OAuth 2.0 Mutual-TLS Client Authentication and Certificate-Bound Access Tokens, RFC 8705. 2020. https://datatracker.ietf.org/doc/html/rfc8705

[11] IETF. OAuth 2.0 Demonstrating Proof of Possession, RFC 9449. 2023. https://datatracker.ietf.org/doc/html/rfc9449

[12] IETF. OAuth 2.0 Authorization Server Metadata, RFC 8414. 2018. https://datatracker.ietf.org/doc/html/rfc8414

[13] IETF. OAuth 2.0 Token Introspection, RFC 7662. 2015. https://datatracker.ietf.org/doc/html/rfc7662

[14] IETF. OAuth 2.0 Token Revocation, RFC 7009. 2013. https://datatracker.ietf.org/doc/html/rfc7009

[15] NIST NCCoE. Accelerating the Adoption of Software and AI Agent Identity and Authorization. 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

[16] NIST. Recommendation for Key Management, SP 800-57 Part 1 Revision 5. 2020. https://csrc.nist.gov/pubs/sp/800/57/pt1/r5/final

[17] NIST. A Framework for Designing Cryptographic Key Management Systems, SP 800-130. 2013. https://csrc.nist.gov/pubs/sp/800/130/final

[18] IFRS Foundation. IFRS 3 Business Combinations. 2026. https://www.ifrs.org/issued-standards/list-of-standards/ifrs-3-business-combinations/

[19] IFRS Foundation. IAS 38 Intangible Assets. 2026. https://www.ifrs.org/issued-standards/list-of-standards/ias-38-intangible-assets/

[20] IFRS Foundation. IFRS 13 Fair Value Measurement. 2026. https://www.ifrs.org/issued-standards/list-of-standards/ifrs-13-fair-value-measurement/

[21] IFRS Foundation. IAS 36 Impairment of Assets. 2026. https://www.ifrs.org/issued-standards/list-of-standards/ias-36-impairment-of-assets/

[22] NIST. Building Secure Microservices-based Applications Using Service-Mesh Architecture, SP 800-204A. 2020. https://csrc.nist.gov/pubs/sp/800/204/a/final

[23] NIST. Implementation of DevSecOps for a Microservices-based Application with Service Mesh, SP 800-204C. 2022. https://csrc.nist.gov/pubs/sp/800/204/c/final

[24] CISA. Zero Trust Maturity Model Version 2.0. 2023. https://www.cisa.gov/resources-tools/resources/zero-trust-maturity-model

[25] IETF. JSON Web Token, RFC 7519. 2015. https://datatracker.ietf.org/doc/html/rfc7519

[26] IETF. JSON Web Token Profile for OAuth 2.0 Client Authentication and Authorization Grants, RFC 7523. 2015. https://datatracker.ietf.org/doc/html/rfc7523

[27] IETF. Best Current Practice for OAuth 2.0 Security, RFC 9700. 2025. https://datatracker.ietf.org/doc/html/rfc9700

[28] IETF. OAuth 2.0 Protected Resource Metadata, RFC 9728. 2025. https://datatracker.ietf.org/doc/html/rfc9728

[29] Kubernetes. Managing Service Accounts. 2026. https://kubernetes.io/docs/reference/access-authn-authz/service-accounts-admin/

[30] Google Cloud. Workload Identity Federation. 2026. https://cloud.google.com/iam/docs/workload-identity-federation

[31] Amazon Web Services. IAM Roles Anywhere. 2026. https://docs.aws.amazon.com/rolesanywhere/latest/userguide/introduction.html

[32] Amazon Web Services. EKS Pod Identities. 2026. https://docs.aws.amazon.com/eks/latest/userguide/pod-identities.html

[33] Microsoft. Workload identity federation. 2026. https://learn.microsoft.com/en-us/entra/workload-id/workload-identity-federation

[34] GitHub. OpenID Connect. 2026. https://docs.github.com/en/actions/concepts/security/openid-connect

[35] HashiCorp. Vault documentation. 2026. https://developer.hashicorp.com/vault/docs

[36] NIST. Secure Software Development Framework, SP 800-218. 2022. https://csrc.nist.gov/pubs/sp/800/218/final

[37] NIST. Cybersecurity Framework 2.0. 2024. https://www.nist.gov/cyberframework

[38] NIST. Digital Identity Guidelines, SP 800-63-4. 2025. https://csrc.nist.gov/pubs/sp/800/63/4/final

[39] OWASP. Non-Human Identities Top 10. 2025. https://owasp.org/www-project-non-human-identities-top-10/

[40] Cloud Native Computing Foundation. SPIFFE project. 2026. https://www.cncf.io/projects/spiffe/

[41] Cloud Native Computing Foundation. SPIRE project. 2026. https://www.cncf.io/projects/spire/

[42] MITRE. ATT&CK Valid Accounts. 2026. https://attack.mitre.org/techniques/T1078/

[43] MITRE. ATT&CK Unsecured Credentials. 2026. https://attack.mitre.org/techniques/T1552/

[44] European Union. Directive (EU) 2022/2555 on measures for a high common level of cybersecurity. 2022. https://eur-lex.europa.eu/eli/dir/2022/2555/oj

[45] European Union. Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence. 2024. https://eur-lex.europa.eu/eli/reg/2024/1689/oj

[46] SEC. Cybersecurity Risk Management, Strategy, Governance, and Incident Disclosure. 2023. https://www.sec.gov/files/rules/final/2023/33-11216.pdf

[47] ISO. ISO/IEC 27001 Information security management systems. 2022. https://www.iso.org/standard/27001

[48] ISO. ISO/IEC 27002 Information security controls. 2022. https://www.iso.org/standard/75652.html

[49] ISO. ISO/IEC 42001 Artificial intelligence management systems. 2023. https://www.iso.org/standard/81230.html

[50] International Valuation Standards Council. International Valuation Standards. 2025. https://www.ivsc.org/standards/

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.