Trust the Model's Supply Chain: Provenance in AI Security M&A

An Acquisition and Valuation Framework for Model Lineage, Attestation, Remediation Economics and Integration

Authored by Chennakeshav Adya

Independent Researcher

September 2026

Abstract

Artificial-intelligence systems combine data, code, model weights, third-party components, training infrastructure, evaluation assets, deployment configuration and operational policy. Each component can change the behaviour, rights, security and commercial usability of the resulting system. An acquirer that cannot reconstruct this supply chain may inherit models whose origins, training conditions, licences, vulnerabilities or approvals cannot be demonstrated. A seller can present model cards, bills of materials and signatures while leaving decisive gaps between declared evidence and the artifact used by customers.

This paper develops an acquisition and valuation framework for provenance in AI-security M&A. The proposed unit of value is a verified model release that meets explicit customer expectations and produces an accountable operating decision at complete cost. The framework tests lineage completeness, artifact identity, attestations, signatures, data and model rights, dependency exposure, evaluation reproducibility, release approval, runtime continuity, customer adoption and remediation economics.

NIST's Secure Software Development Framework requires organisations to collect and share provenance data for software-release components. NIST SP 800-218A extends secure-development practices to generative AI and dual-use foundation models, including model and component provenance. The NIST AI Risk Management Framework addresses third-party software, data and supply-chain risk. SLSA defines provenance as verifiable information describing where, when and how an artifact was produced. CISA's SBOM guidance emphasises consumption processes that convert component transparency into risk decisions. In-toto, Sigstore, SPDX and CycloneDX provide complementary mechanisms for attestations, signing and machine-readable component records.[1][2][3][4][5][6][7][8]

A hypothetical acquisition illustrates a provider that inventories AI assets, generates and verifies attestations, governs releases and supports regulated customers. Every revenue, customer, cost, probability, performance and valuation figure in the 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 evidence continuity and remediation capacity before assigning value to provenance coverage. Six figures and seven tables convert the framework into diligence tests, a valuation bridge, consideration protection and a 180-day integration programme.

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

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

Keywords: AI provenance, model supply chain, cybersecurity M&A, model lineage, attestations, SBOM, valuation, integration

1. Define the acquisition decision

The board should identify the customer decision that the target improves. AI-provenance firms may inventory models, trace training runs, sign artifacts, generate bills of materials, verify build attestations, enforce release policy, monitor model changes or investigate incidents. These activities serve a common trust objective while producing different evidence and economics.

The transaction thesis should state whether the buyer seeks proprietary lineage technology, regulated-customer access, integration with development platforms, scarce security expertise, a compliance control plane or a consolidation platform. Each source of value needs an observable test. Lineage claims require reconstructed releases. Distribution claims require deployed customer workflows, renewal and collections.

The board should compare acquisition with licensing, partnership, minority investment and internal development. Ownership may matter when value depends on controlling the evidence graph, verification policy, integrations and security team. A narrower arrangement can be proportionate when the principal benefit is access to a standard or channel.

Evidence timing should shape terms. Controlled release reconstruction can occur before signing. Production coverage, customer acceptance and remediation cost may require later access. Base consideration should follow evidence available at closing; contingent value should follow completed customer and integration milestones.

Figure 1. Model-release-to-value evidence chain

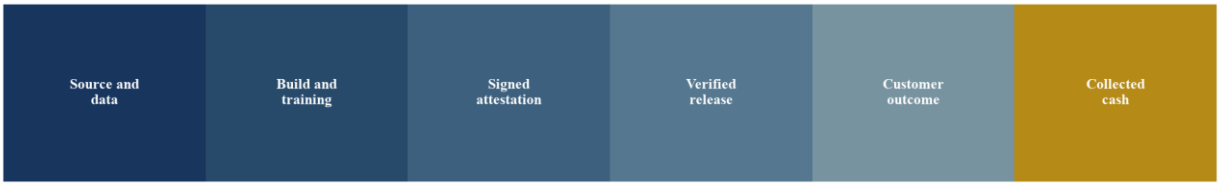

The proposed chain connects model lineage to a verified release, customer decision and collected cash.

2. Define the unit of value

The proposed unit of value is a verified model release that meets explicit customer expectations and produces an accountable operating decision at complete cost. A release record should bind the model digest to source revisions, datasets, dependencies, training instructions, builder identity, evaluation results, approval, package and deployment configuration.

Complete cost includes metadata capture, artifact storage, signing, key custody, verification, policy operation, integrations, remediation, customer support, security review, compliance and working capital. A platform can appear scalable while customer engineers manually reconstruct missing evidence. The acquisition model should include every activity required to support the promised decision.

Metadata volume is an incomplete denominator. A million lineage records create limited value if they cannot prove which artifact reached production or whether its licence and evaluation satisfy policy. Buyers should measure verified releases, failed verifications, time to resolution, customer action, avoided rework and retained contribution.

3. Map the AI supply chain

The supply chain begins before training. Data collection, cleaning, labelling and transformation can affect rights, bias, safety and reproducibility. Code, libraries, frameworks, base models, adapters, prompts, evaluation sets, hardware, training services and deployment components can each introduce dependency and control risk.

The diligence team should map first-party, third-party and open-source elements. It should identify who selected each component, which rights were obtained, where it was processed, how changes were approved and what evidence survives. A model registry should not be treated as a complete supply-chain graph without this reconstruction.

Derivatives require special attention. Fine-tuning, quantisation, distillation, merging and retrieval augmentation can change behaviour and rights while preserving a familiar model name. The buyer should trace each production artifact to its exact parents and transformation instructions.

Table 1. AI supply-chain components and acquisition tests

Component	Required evidence	Principal exposure	Acquisition test
Training data	source, rights, transformations	infringement, privacy, quality	sample lineage reconstruction
Code and libraries	revision, dependency and licence	vulnerable or restricted component	reproducible build
Base model	digest, supplier terms and evaluation	change, access or licence restriction	artifact and contract match
Fine-tuning	dataset, method and run record	behaviour and rights drift	reproduce approved checkpoint
Evaluation	versioned set, method and result	non-comparable performance	rerun sealed tests
Deployment	package, policy and configuration	wrong artifact in production	runtime-to-release reconciliation

Each component creates a distinct evidence and remediation obligation.

4. Build the provenance ledger

The ledger should connect requirement, source, data, code, dependency, training run, model digest, evaluation, approval, package, signature, deployment, customer use, incident and financial record. It should preserve changes and superseded artifacts instead of overwriting history.

Negative evidence belongs in the ledger. Missing parents, unsigned artifacts, failed builds, expired keys, unresolved licences, unapproved evaluations and emergency overrides reveal the actual control boundary. A data room containing only successful release records cannot support a population conclusion.

Finance should link customer cohorts to verified releases, integrations, support effort, renewal, expansion and collections. This shows whether provenance depth reduces customer friction or creates unpriced services work.

The ledger needs a controlled vocabulary for relationships. Terms such as derived from, trained on, evaluated by, packaged with, approved by and deployed as should have precise meanings. Free-text links can make a graph look complete while preventing automated verification. Schema changes should be versioned and migration should preserve earlier interpretations.

Evidence custody should also be recorded. Some customers require metadata to remain within their environment; others allow a vendor control plane to store it. The buyer should identify where evidence is

generated, transmitted, retained, backed up and deleted. Customer encryption, residency and access commitments affect both architecture and delivery cost.

Reconciliation should operate continuously. A production artifact that appears without an approved release, or an attestation that names an unknown builder, should create an exception with an owner and due date. Closure should include the technical action and customer decision. A silent exception backlog can conceal extensive manual acceptance of unverified artifacts.

The buyer should sample the ledger in both directions. Starting from a production artifact, it should reach all required sources, approvals and evaluations. Starting from a vulnerable dependency or restricted dataset, it should identify every affected derivative and customer. These traversals test the graph's practical value for prevention and response.

5. Measure lineage completeness

Completeness begins with an independently defined population of production and customer-delivered artifacts. The buyer should reconcile model registries, object stores, container registries, repositories, deployment platforms, cloud accounts and customer manifests. The target's lineage graph should then be compared with that population.

Coverage should be measured at required fields and edges. An artifact may appear in inventory while its training data, parent model or approval remains unknown. The buyer should distinguish discovered, identified, linked, attested, verified and policy-compliant states.

Seeded tests can expose blind spots. The diligence team can create approved artifacts with known parents, ambiguous names, copied metadata and altered packages. It should measure discovery, graph construction, conflict handling and remediation without seller intervention.

Figure 2. Hypothetical coverage and unresolved-gap curve

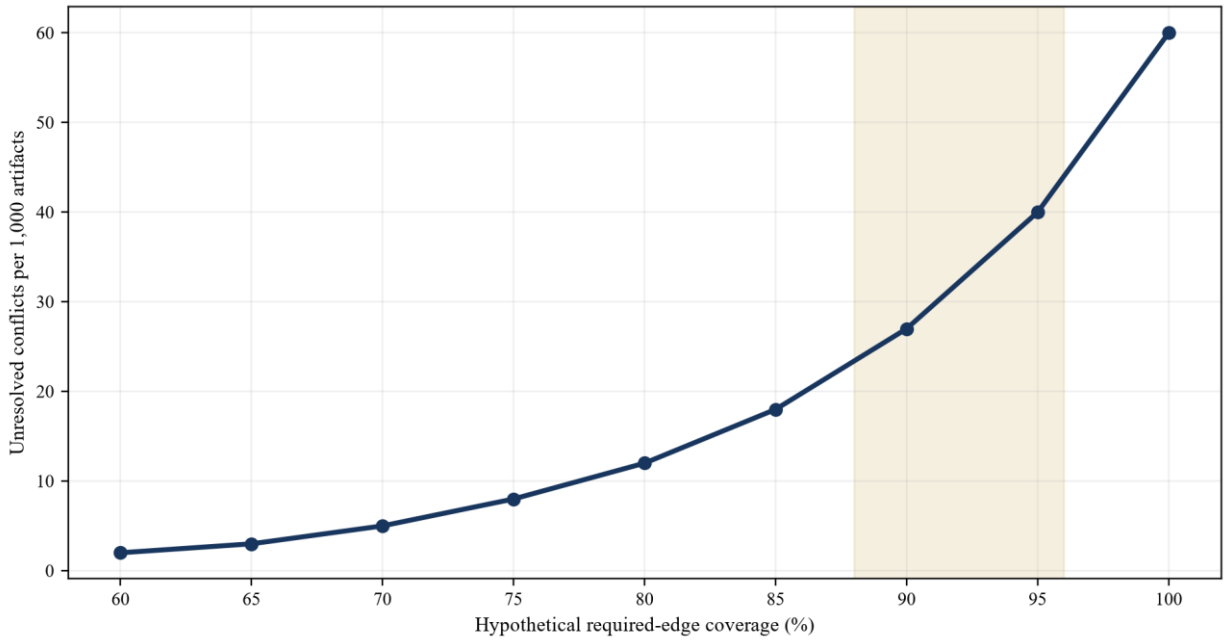

Values are management assumptions for method demonstration.

6. Bind identity to artifacts

Every material artifact should have a stable cryptographic digest. Names, paths and tags can change or be reused. The buyer should verify that source revisions, datasets, model weights, packages and deployment images are bound to the identifiers recorded in attestations.

The system should distinguish identity from location. Copying a model to another registry should preserve the artifact digest while changing custody and policy context. Rebuilding from the same source may produce a different digest where training is stochastic or the environment is not reproducible.

Tests should include substitution, tag reuse, partial download, altered metadata and repackaging. Verification should fail safely and produce evidence that an operator can investigate.

7. Evaluate attestations and signing

An attestation is a signed statement about an artifact or process. SLSA provenance can describe source, builder and external parameters through an in-toto predicate. Sigstore supports signing and verification workflows using transparency services and identity-linked certificates.[4][6][9]

The buyer should inspect issuer identity, signing policy, key or certificate lifecycle, transparency evidence, revocation, timestamping and verification expectations. A valid signature proves that a key signed a statement; it does not prove that the statement is complete or truthful.

Attestations should be generated by controlled systems rather than reconstructed manually after release. The diligence team should attempt to forge provenance, use an unapproved builder, alter external parameters and replay an old attestation against a new artifact.

Table 2. Attestation maturity model

Level	Capability	Evidence	Value limitation
1	metadata inventory	artifact record	no integrity assurance
2	signed statement	signature and issuer	statement may be incomplete
3	controlled generation	builder and process identity	limited consumer expectations
4	policy verification	approved source, builder and parameters	integration effort
5	continuous enforcement	admission, monitoring and response	governance and availability burden

Value increases when signed evidence is verified against explicit expectations and drives action.

8. Test reproducibility and verifiability

Reproducibility asks whether the same inputs and process produce the same output. Many AI-training workflows contain stochastic operations, hardware differences and external services that limit bit-for-bit reproduction. The buyer should define what can be reproduced and what evidence supports equivalent behaviour.

Verifiability can still be strong when exact reproduction is impractical. Controlled builders, immutable inputs, signed run records, retained checkpoints and independent evaluation can establish a reliable chain. The target should explain uncertainty rather than claiming universal reproducibility.

The diligence team should rebuild representative software components, rerun selected training or fine-tuning steps and reproduce evaluations. Variance should be recorded and connected to approved tolerances.

Figure 3. Hypothetical evidence-decay curves

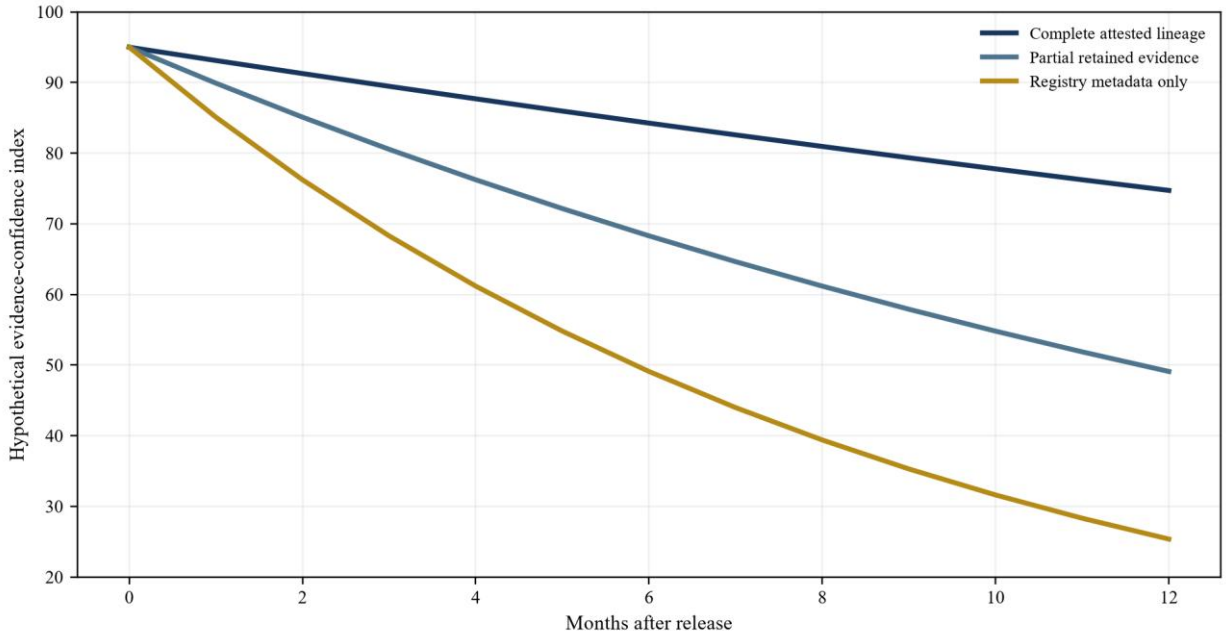

The curves illustrate how retained evidence affects confidence after platform and dependency change; values are management assumptions.

9. Evaluate bills of materials

SPDX and CycloneDX provide machine-readable formats for software and broader component information. AI extensions can record models, datasets and relationships. CISA emphasises that SBOM value depends on consumption processes that turn component data into risk action.[5][7][8]

The buyer should test completeness, version accuracy, dependency depth, identifiers, licences and vulnerability mapping. A generated bill can miss dynamically loaded, hosted or customer-supplied elements. The product should state the observation boundary.

An AI bill of materials should complement, not replace, provenance. A list describes components; provenance explains how a specific artifact was produced. Verification policy needs both relationships and approved expectations.

10. Diligence data provenance and rights

Data lineage should connect source, collection basis, permission, licence, transformation, labelling, filtering, retention and use. The buyer should sample records from production models back to source evidence. Aggregated descriptions are insufficient for high-risk populations.

Rights may differ across training, evaluation, fine-tuning, retrieval and output use. Contract language, open licences, privacy obligations and customer restrictions require qualified legal review. Technical controls should reflect approved use rather than assuming that possession permits processing.

The target should show removal and retraining procedures where rights expire or a source must be excluded. Remediation cost depends on data isolation, model dependency and availability of substitutes.

Dataset identity requires more than a filename. Versioned manifests should record included objects, hashes or stable references, transformation code, filtering rules and label provenance. Where privacy or contractual limits prevent retaining raw data, the system should retain sufficient controlled evidence to support approved use and later review. The buyer should test whether a training run can be connected to the exact dataset state that existed at the time.

Derived and synthetic data need their own lineage. A generated dataset may depend on a source model, prompt process, sampling rules, human review and original reference material. Synthetic origin does not remove rights, quality or security questions. The provenance record should preserve the derivation and approved purpose.

Data suppliers and annotation vendors create third-party risk. Contracts, security controls, worker access, quality review and change notification should match the technical record. A vendor name in a model card does not establish which data was delivered or how it was used. The diligence sample should reconcile invoices, delivery manifests, storage records and training configurations.

Privacy and deletion requests can propagate through caches, derived datasets, checkpoints and deployed models. Current technical methods may not allow removal of an individual record's influence from a trained model with certainty. The buyer should examine the target's legal position, retraining capability, documentation and customer communication rather than assuming a complete technical remedy.

11. Diligence model and dependency rights

Base-model terms can restrict commercial use, redistribution, fine-tuning, regulated applications or deployment geography. Open-source labels do not replace licence analysis. The buyer should reconcile model digests with the terms applicable when each artifact was obtained.

Dependencies include training frameworks, tokenisers, evaluation libraries, safety filters, container images and hosted APIs. A change in one dependency can alter security, performance, cost or rights. The product should preserve version and source evidence.

Change-of-control, assignment and sublicensing provisions affect integration. The buyer should identify consents and replacement options before assuming that the acquired platform can be combined or redistributed.

Table 3. Rights and replacement tests

Asset	Rights evidence	Replacement test	Value consequence
Dataset	source and permitted use	isolate and substitute	retraining cost and delay
Base model	exact terms and digest	alternative model evaluation	margin and performance change
Library	licence and dependency tree	rebuild with approved version	engineering and security effort
Hosted API	contract and service terms	portable interface and fallback	concentration and pricing risk
Evaluation set	ownership and permitted reuse	recreate comparable benchmark	evidence continuity

The matrix connects rights evidence to commercial continuity.

12. Test evaluation provenance

Performance claims should bind the tested artifact, dataset, method, environment, metric, threshold and result. A model card that reports an unlinked score cannot prove the performance of the deployed package.

The buyer should rerun sealed evaluations and compare results. It should test data contamination, repeated tuning, changed prompts, post-processing and customer-specific configuration. Differences should be investigated rather than averaged away.

Evaluation approval should include intended use, limitations, risk tolerance and accountable sign-off. NIST's AI RMF treats testing, evaluation, verification and validation as continuing lifecycle work.[3][10]

13. Verify release and deployment continuity

The release gate should compare the artifact and attestations with approved expectations. SLSA verification includes artifact identity, signature, builder, source and external parameters. Verification without an action path creates limited protection.[4][11]

The buyer should trace sampled customer deployments back to approved releases. It should inspect admission controls, exceptions, emergency deployment, rollback and runtime monitoring. Customer-managed deployments need evidence that survives outside the vendor environment.

The system should identify drift after release, including changed configuration, adapters, retrieval sources or safety policy. A verified model can become an unverified system when surrounding components change.

Release expectations should be explicit and versioned. They may include approved repositories, builders, model families, licences, evaluation thresholds, regions, risk classifications and signatories. Unknown fields or parameters should fail or require authorised exception rather than being ignored. The buyer should test whether expectations are controlled through reviewed code or an equivalent auditable mechanism.

Exception governance affects commercial value. Emergency releases may be necessary, but they should identify the approver, reason, scope, expiry and compensating control. The product should prevent a temporary waiver from becoming a permanent bypass. Cohort analysis should show the volume, age and recurrence of exceptions by customer and product.

Customer deployment models change the evidence boundary. A software-as-a-service vendor can control release admission centrally. An on-premises or air-gapped customer may verify evidence locally and report only a result. The target should demonstrate how policy, trust roots, revocation and audit updates reach each model without relying on unsupported remote access.

Runtime reconciliation should compare observed digests and configurations with the approved release. It should detect shadow deployments, copied models and unauthorised adapters. Alerts need an operational response; unresolved differences should appear in service reporting and customer governance.

14. Test security and abuse resistance

The provenance platform is privileged infrastructure. Compromise can sign malicious artifacts, alter lineage, suppress failures or disclose sensitive architecture. The buyer should review threat models, code, build systems, key custody, privileged access and tenant isolation.

Scenarios should include stolen signing identity, compromised builder, poisoned dependency, malicious insider, transparency-log outage, policy bypass and denial of service. Each needs prevention, detection, containment and recovery evidence.

NIST SP 800-218 and SP 800-218A provide a secure-development baseline. The acquisition team should connect claimed practices to repositories, build logs, approvals and incident records.[1][2]

Figure 4. Proposed provenance-control architecture

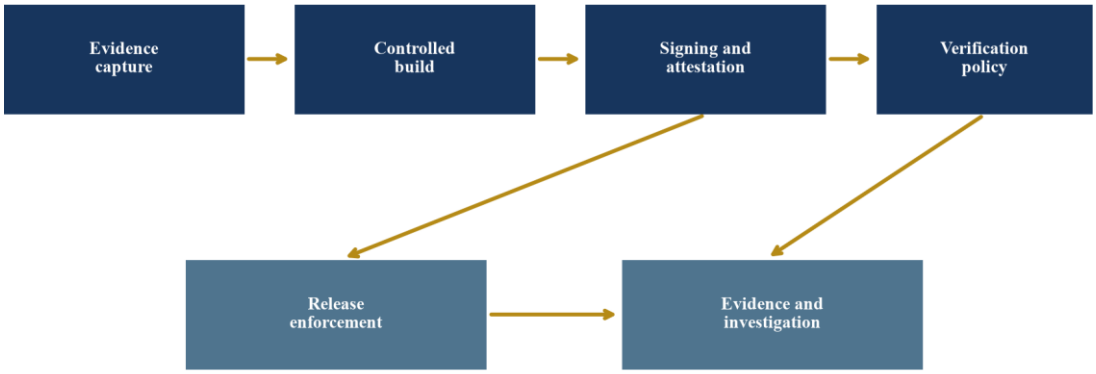

The architecture separates evidence capture, signing, verification, enforcement and investigation.

15. Quantify remediation economics

Remediation begins with exposure discovery. The buyer should identify affected artifacts, customers, rights, dependencies and environments. It should then estimate replacement, retraining, retesting, migration, communication, legal review, credits and incident cost.

Cost varies by graph position. Replacing a leaf library may require a rebuild and regression test. Replacing a base model or dataset can affect every derivative, evaluation and contract. The provenance graph should support impact analysis.

The model should include time and cash. Engineering capacity diverted to remediation can delay roadmap and sales. Customer disruption can reduce renewal before direct costs appear.

Impact analysis should distinguish a disclosed weakness from an exploitable production exposure. Component presence, execution path, configuration, compensating controls and customer use affect priority. Provenance helps narrow the affected population, but the buyer should test the accuracy of that narrowing before recognising cost savings.

Replacement paths may change performance and economics. Substituting a base model can alter inference cost, latency, accuracy, safety and data-location obligations. Replacing a library can require code changes and new evaluation. The remediation model should include requalification and customer acceptance, not only engineering hours.

Rights remediation can require licence purchase, data removal, retraining, settlement or withdrawal from a use case. Each route has different timing and cash. Where facts are uncertain, the acquisition case should use scenarios and retain a reserve instead of presenting one point estimate.

Incident remediation should include investigation, evidence preservation, regulator and customer communication, legal advice, service credits, insurance deductibles and increased support. Insurance recovery should be recognised only when policy terms and claim facts support it. A security event can reduce renewal and pipeline while increasing delivery cost.

The buyer should compare the target's historical estimates with completed remediation. Variance in scope, duration, cost and customer impact reveals planning quality. A platform that produces rapid impact analysis can create value through narrower, faster action, provided that the result is reproduced during diligence.

16. Diligence customer cohorts and distribution

Customers should be segmented by industry, deployment model, regulated status, evidence depth, verification policy, contract, support burden, renewal and collections. A customer that stores metadata should not receive the same valuation as one that blocks unverified releases.

The buyer should reconstruct adoption from connector installation through first inventory, signed release, enforced policy and steady-state use. Time to value and open exceptions influence contribution and retention.

Distribution partnerships require sourced pipeline, conversion and economics. Integration with a development platform can widen reach while increasing platform dependence and pricing pressure.

The implementation funnel should run from signed order through connector installation, inventory coverage, first attestation, first verified release, enforced policy and steady-state governance. The buyer should measure elapsed time, professional-services effort and open exceptions at each stage. Contracted revenue that remains at inventory may have lower durability than the reported subscription suggests.

Expansion should be decomposed into artifact volume, additional teams, new environments and deeper enforcement. Volume growth can follow customer activity without demonstrating more value. Deeper enforcement can increase customer reliance while raising integration and support requirements. Net retention should therefore be analysed with retained contribution and control depth.

Customer outcome evidence can include faster incident scoping, reduced manual release review, fewer unauthorised deployments, improved audit preparation and shorter remediation. Each measure needs a baseline, defined population and source. Testimonials and calculated savings should remain separate from observed operating records.

Contracts may restrict assignment, telemetry transfer, hosting changes and use of customer metadata. The buyer should map change-of-control consents, data localisation, customer-controlled keys and audit obligations before assuming platform consolidation. A migration that breaks an evidence chain can create contractual and operational risk.

Table 4. Customer-cohort evidence matrix

Cohort	Deployment evidence	Economic test	Principal risk
Regulated enterprise	enforced verification and audit export	retained contribution	long implementation
AI developer	release attestations and policy	expansion and support	tool consolidation
Critical infrastructure	controlled deployment and rollback	contract duration	operational liability
Platform customer	integrated admission control	net revenue	channel dependence
Metadata-only customer	inventory coverage	migration potential	limited workflow adoption

Cohorts should be valued according to enforced use, contribution and durability.

17. Reconstruct complete delivery economics

Revenue should be reconciled from contract through invoice and bank receipt. The buyer should separate subscription, usage, implementation, managed remediation and pass-through services. Annual recurring revenue should exclude unsupported or non-recurring amounts.

Cost includes storage, graph processing, signing services, transparency infrastructure, vulnerability data, support, security review and customer engineering. Labour that repeatedly repairs missing lineage belongs in delivery economics.

Unit economics should use verified releases, active integrations and evidence volume alongside customer cohorts. Pricing by model count can discourage complete capture or misalign with verification value.

The buyer should rebuild gross margin from source records. Customer engineering, recurring schema mapping, evidence repair and audit support may be classified as product development while functioning as cost of service. Cloud credits and minimum commitments can temporarily improve reported margin. Normalisation should retain the resources required to deliver the current promise.

Infrastructure cost should be traced to graph storage, artifact retrieval, signing, transparency queries, vulnerability feeds, policy evaluation and retention. Peaks can occur during enterprise onboarding or incidents. Average cost per customer can conceal a small cohort with unusually complex evidence and support.

Pricing models should be tested for behavioural effects. Per-artifact pricing can discourage complete inventories. Per-verification pricing can align with enforcement while creating bill uncertainty. Enterprise subscription can support adoption while transferring volume and retention risk to the vendor. Contracts should be reviewed for minimums, overages, service credits and indexation.

Sales efficiency needs a full-cycle view. Security review, proof of concept, procurement, integration and policy approval can extend well beyond signature. The buyer should measure cash acquisition cost from initial pursuit through collected contribution and compare cohorts by channel and regulated status.

Working capital should connect deployment, billing and collections. Large customers may delay payment until acceptance or audit completion while the target funds integrations. The valuation model should reflect cash conversion, not only recognised revenue.

18. Build a hypothetical acquisition case

Assume a target with USD 17.0 million recurring revenue, USD 3.5 million implementation revenue and USD 1.0 million remediation revenue. Management estimates USD 11.8 million retained recurring contribution after direct delivery and support. The largest ten customers represent 46 per cent of recurring revenue. These figures are hypothetical.

Evidence review attributes USD 6.8 million of contribution to customers enforcing verified releases, USD 3.1 million to signed provenance without enforcement and USD 1.9 million to inventory-only customers. Each layer receives different confidence.

Management identifies USD 2.2 million potential cross-sell contribution and USD 1.4 million duplicated cost. The base valuation excludes both until customer acceptance and delivery evidence exist.

The target reports ninety enterprise customers. Diligence confirms that thirty-two enforce provenance policy in production, twenty-six verify signatures without blocking release, twenty use the product mainly for inventory and twelve remain in implementation. These counts are hypothetical. The buyer should avoid applying one retention or margin assumption to all four groups.

The enforced cohort has longer contracts and higher implementation cost. The inventory cohort has lower support cost but weaker evidence of customer dependence. Finance should calculate retained contribution

by cohort after cloud, signing, support, customer engineering and partner share. Customer concentration should be presented within each adoption layer.

The transaction model assumes that half of the signature-only cohort reaches enforcement over two years. This is a management scenario, not an observed probability. Consideration for that migration should follow completed adoption and collected contribution. The integration budget should include connector work, policy design, customer security review and audit migration.

An identified rights issue affects one dataset connector used by six customers. The hypothetical base case reserves USD 2.0 million for replacement and customer work. The adverse case assumes slower substitution, additional legal cost and one customer loss. This treatment keeps the known exposure visible instead of netting it against broad synergies.

Table 5. Hypothetical evidence-layered contribution

Layer	Retained contribution	Evidence status	Valuation treatment
Enforced verified releases	6.8	deployed and renewed	base case subject to retention
Signed provenance	3.1	deployed without full enforcement	adoption-adjusted
Inventory only	1.9	limited workflow value	contingent or option value
Potential cross-sell	2.2	management plan	excluded from base price
Duplicated cost opportunity	1.4	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 signing compromise, incomplete lineage, licence removal, platform change, customer loss, slower enforcement adoption and higher remediation cost. Correlated events require specific treatment.

The buyer should model liquidity. Emergency key rotation, customer notification, rebuilds, retraining, legal review and credits can require cash before insurance or revenue recovery.

Concentration should be mapped by customer, cloud, model supplier, data source, signing system and channel. Logo diversification can conceal common dependency exposure.

Stress design should follow causal chains. A signing compromise may require trust-root rotation, release reverification, customer communication, service credits and forensic review. Sales can slow while support cost rises. Treating each effect independently can understate the combined event.

Supplier-change stress should examine model deprecation, licence revision, API pricing, regional availability and changed safety policy. The target should identify affected derivatives and customer contracts quickly. Replacement testing should include performance, cost, rights and customer approval.

Incomplete-lineage stress should assume that a material dependency cannot be proven for part of the installed base. The model should estimate discovery, evidence reconstruction, customer assurance, rebuild and possible withdrawal. The response should state which actions can occur before legal or technical certainty.

Management responses should be feasible and sequenced. Cost reduction can protect liquidity while slowing remediation. Forced migration can simplify the platform while increasing churn. The board should define triggers for additional security capacity, customer escalation, liquidity preservation and covenant engagement.

The stress pack should distinguish contractual facts, observed metrics, management estimates and scenario assumptions. Post-close results should be compared with the original cases each month so that variance changes the integration plan and contingent-value assessment.

Figure 5. Hypothetical retained-contribution stress bridge

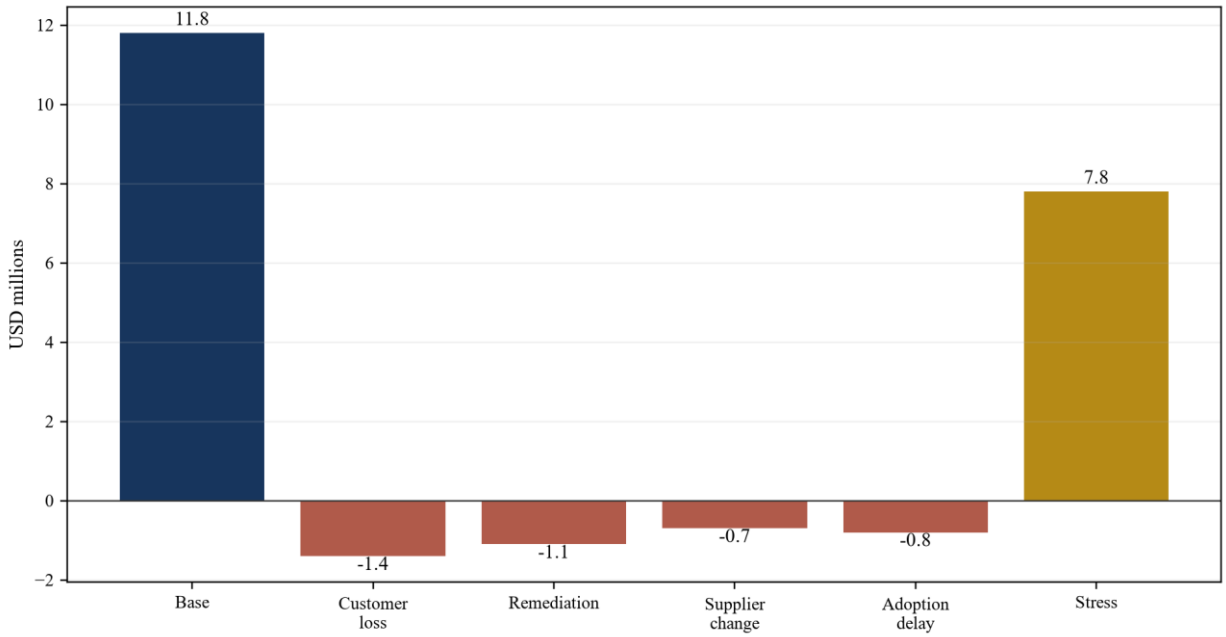

All values are management assumptions in USD millions.

20. Value the evidence layers

Valuation should begin with retained recurring contribution supported by contracts, enforced use and cash. The required return or multiple should reflect growth, retention, concentration, security exposure, remediation capacity and capital needs.

The bridge should separate enforced production value, adoption-dependent value, inventory options, delivered synergies and risk reserves. Each layer needs an owner, milestone, cost and downside case.

IFRS 3, IAS 38 and IFRS 13 may require separate recognition and measurement of technology, customer relationships and other assets. IAS 36 governs impairment assessment according to the applicable facts and advice.[12][13][14][15]

Evidence durability should influence the forecast period and required return. Customer contribution can weaken at renewal, technical evidence can decay after dependency changes, and policy integrations can break during migration. Each material layer should have a review date, leading indicator and downside response.

Strategic option value should remain separate from current cash flow. A lineage platform may support future regulatory reporting or agent governance, but additional product, sales, legal and capital requirements should be identified. An option can justify transaction structure without supporting the same amount of cash consideration at closing.

Comparable-company and transaction evidence requires normalisation. Revenue definitions, services content, growth, retention, stock compensation, cash burn and security liability vary. The valuation committee should retain a traceable bridge from observed market evidence to the company-specific conclusion.

Contingent value should use measures that the seller and buyer can verify. Suitable measures may include retained contribution from enforced customers, completed migrations and collected cross-sell. Artifact

count or metadata volume can be manipulated or disconnected from value. Definitions should address acquisitions, pricing changes, customer credits and accounting-policy changes.

The board should review value and certainty together. A higher headline price with broad unresolved rights, customer consents and security exposure can produce lower risk-adjusted value than a staged structure. The model should present consideration, remediation funding, integration investment, working capital and downside liquidity in one view.

Table 6. Hypothetical enterprise-value bridge

Component	Evidence basis	Hypothetical value USDm
Enforced customer contribution	deployed, renewed and collected	68.0
Adoption-dependent contribution	signed provenance customers	17.0
Inventory option	metadata-only customers	5.0
Delivered synergy	verified milestones	7.0
Remediation and concentration reserve	downside adjustment	-15.0
Illustrative enterprise value	sum of evidence layers	82.0

Amounts and valuation factors are management assumptions.

Figure 6. Hypothetical evidence-layered enterprise value

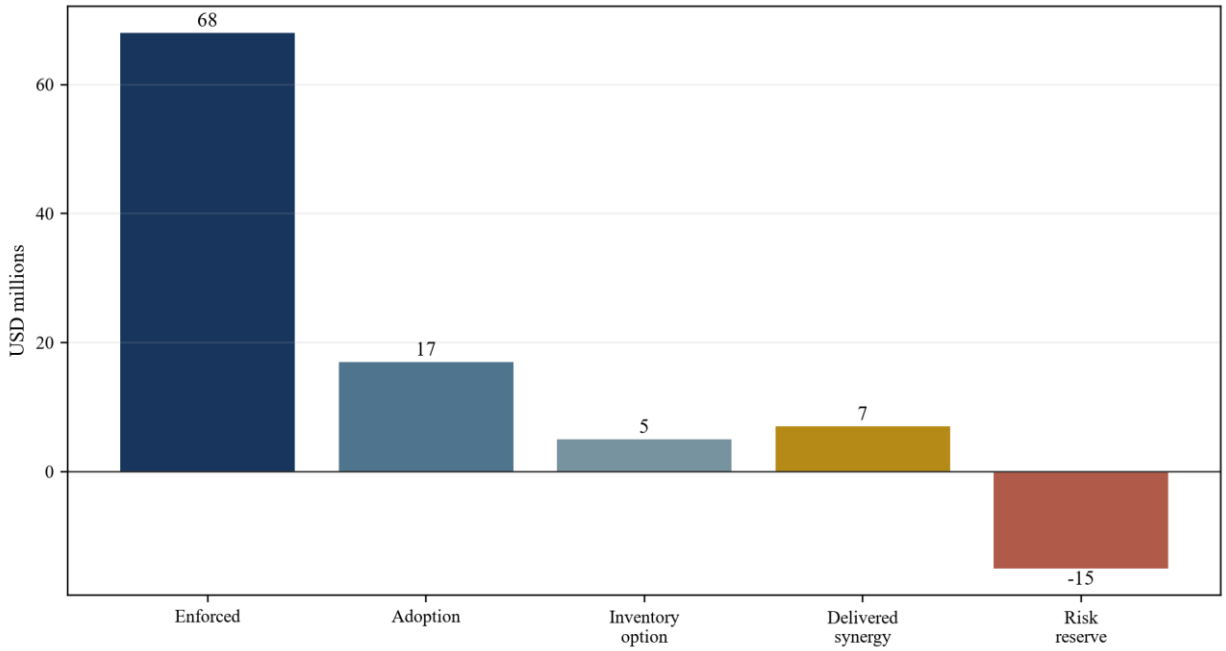

Values are management assumptions in USD millions and do not represent a market benchmark.

21. Structure consideration and integration

Base consideration should reflect reproduced technology, transferable rights, retained customer contribution and cash. Deferred value can address enforcement adoption, rights remediation, customer retention and security integration.

Representations and warranties should address intellectual property, data rights, licences, open-source use, artifact integrity, signing custody, incidents, customer commitments and compliance. Identified exposures may require conditions, escrow or specific indemnities subject to legal advice.

Integration should preserve verification continuity. The buyer should avoid replacing identifiers, trust roots or policy without mapping, equivalence tests, rollback and customer approval.

Earn-out design should avoid metrics that management can change through platform migration or accounting classification. Retained contribution from named cohorts, completed policy enforcement and collected cross-sell can be more auditable than revenue alone. The agreement should define customer credits, bundled contracts, currency, acquisitions and discontinued products.

Integration governance should assign authority for trust roots, signing policy, schema changes, release exceptions and customer communication. Security and commercial leaders should approve changes that alter customer evidence. A product roadmap should not override signed control obligations without explicit review.

Migration sequencing should begin with low-complexity cohorts while preserving support for regulated customers. Each wave should require evidence equivalence, performance testing, rollback and customer acceptance. The buyer should track duplicated cost separately from the cost needed to maintain safe parallel operation.

Table 7. Consideration and integration gates

Gate	Evidence	Transaction response
Lineage	reconstructed representative releases	supports base value
Rights	transferable data, model and software rights	condition or remediation
Customers	retained enforced contribution	deferred consideration
Security	signing custody and incident review	escrow, indemnity or condition
Migration	identity, policy and evidence equivalence	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 over signing identities, privileged access, incident response, customer escalation, artifact inventories and integration decisions. High-risk architectural changes should pause until evidence is preserved.

Days 31 to 60 should reproduce lineage, attestations, builds, evaluations and deployment reconciliation. Finance should reconcile contribution and collections by customer cohort. Legal teams should confirm critical rights and dependencies.

Days 61 to 100 should define the combined evidence architecture, verification policy and migration sequence. 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 realised benefits against the signed baseline.

The programme office should maintain one evidence register covering technical tests, rights, customers, economics, security exceptions and transaction commitments. Every material issue should have an owner, due date, decision and impact on value or integration. Closed status should require evidence of completion.

Board reporting should distinguish leading indicators from realised value. Inventory coverage, signed attestations and migration activity are leading indicators. Retained customer contribution, reduced recurring cost and collected cross-sell are realised financial results. This distinction prevents activity from being reported as synergy.

At day 180, management should decide which product components become the strategic platform, which remain supported, which are retired and which require further evidence. The decision should consider customer commitments, control quality, economics and remaining migration risk. Benefits should continue to be monitored after the initial programme.

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

23. Decision and conclusion

AI provenance creates acquisition value when a platform can reconstruct model lineage, bind evidence to exact artifacts, verify releases against explicit expectations and support remediation. Metadata volume and signatures are inputs to that outcome.

A successful acquisition requires evidence continuity. Consolidation that breaks artifact identity, trust roots, policy or customer audit history can destroy the control customers purchased.

The proposed framework links provenance to customer decisions, retained contribution and cash. It prices enforced production value, treats adoption as evidence-dependent, protects consideration and gives management a controlled integration sequence.

The board approval should state the population tested, releases reconstructed, rights confirmed, customer contribution reconciled, security exceptions accepted and milestones governing payment. Continuing monitoring should link lineage coverage, verification failures, remediation time, renewal, contribution and cash.

Frequently asked questions

What is AI model provenance?

AI model provenance is verifiable information describing the sources, components, processes and approvals that produced a specific model artifact and its subsequent derivatives and deployments.

Is a model card sufficient provenance?

No. A model card can describe intended use and performance while remaining unbound to the exact deployed artifact, source revisions, data, builder and release approval.

What does a valid signature prove?

A valid signature proves that a key or identity signed a statement. Verification must also establish trust in the signer, binding to the artifact, completeness of the statement and compliance with expectations.

How do SBOMs and provenance differ?

A bill of materials lists components. Provenance describes how a specific artifact was produced and links it to sources, builders and parameters. Effective control often needs both.

How should a buyer test lineage completeness?

The buyer should define an independent artifact population, reconcile production and customer releases, sample required graph edges and reconstruct representative releases without seller intervention.

How should remediation be valued?

The model should include affected artifacts, replacement rights, engineering, retraining, evaluation, migration, customer disruption, legal review, credits, timing and cash.

How should synergies be treated?

Potential cross-sell and cost reduction should remain outside base value until customer adoption, collected contribution and completed integration are evidenced.

What should happen during the first 180 days?

Management should secure signing and privileged access, reproduce evidence, reconcile economics, define the combined architecture, run controlled migrations and report realised benefits.

References

[1] NIST. Secure Software Development Framework Version 1.1, SP 800-218. 2022. https://csrc.nist.gov/pubs/sp/800/218/final

[2] NIST. Secure Software Development Practices for Generative AI and Dual-Use Foundation Models, SP 800-218A. 2024. https://csrc.nist.gov/pubs/sp/800/218/a/final

[3] NIST. Artificial Intelligence Risk Management Framework 1.0. 2023. https://www.nist.gov/itl/ai-risk-management-framework

[4] SLSA. Provenance Specification. 2026. https://slsa.dev/spec/v1.2/provenance

[5] CISA. Recommended Practices for SBOM Consumption. 2024. https://www.cisa.gov/resources-tools/resources/software-bill-materials-sbom

[6] In-toto. Attestation Framework. 2026. https://in-toto.io/

[7] SPDX. SPDX 3.0 Specification. 2026. https://spdx.github.io/spdx-spec/v3.0/

[8] CycloneDX. Specification. 2026. https://cyclonedx.org/specification/overview/

[9] Sigstore. Documentation. 2026. https://docs.sigstore.dev/

[10] NIST. AI Resource Center. 2026. https://airc.nist.gov/

[11] SLSA. Verifying Artifacts. 2026. https://slsa.dev/spec/v1.2/verifying-artifacts

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

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

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

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

[16] NIST. Cybersecurity Supply Chain Risk Management Practices, SP 800-161 Rev. 1. 2022. https://csrc.nist.gov/pubs/sp/800/161/r1/final

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

[18] NIST. Adversarial Machine Learning Taxonomy, AI 100-2e2025. 2025. https://csrc.nist.gov/pubs/ai/100/2/e2025/final

[19] NIST. Generative AI Profile, AI 600-1. 2024. https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf

[20] NIST. Security and Privacy Controls, SP 800-53 Rev. 5. 2020. https://csrc.nist.gov/pubs/sp/800/53/r5/upd1/final

[21] NIST. Risk Management Framework. 2026. https://csrc.nist.gov/projects/risk-management/about-rmf

[22] CISA. Secure by Design. 2026. https://www.cisa.gov/securebydesign

[23] CISA. Software Bill of Materials. 2026. https://www.cisa.gov/sbom

[24] NTIA. Software Component Transparency. 2021. https://www.ntia.gov/page/software-bill-materials

[25] OpenSSF. Scorecard. 2026. https://scorecard.dev/

[26] OpenSSF. Security Baseline. 2026. https://baseline.openssf.org/

[27] OpenSSF. Model Signing. 2026. https://github.com/sigstore/model-transparency

[28] CNCF. Software Supply Chain Best Practices. 2021. https://project.linuxfoundation.org/hubfs/CNCF_SSCP_v1.pdf

[29] OCI. Image Specification. 2026. https://github.com/opencontainers/image-spec

[30] OCI. Distribution Specification. 2026. https://github.com/opencontainers/distribution-spec

[31] IETF. The Concise Software Identification Tags, RFC 9393. 2023. https://datatracker.ietf.org/doc/html/rfc9393

[32] IETF. Entity Attestation Token, RFC 9711. 2025. https://datatracker.ietf.org/doc/html/rfc9711

[33] IETF. Remote ATtestation procedureS Architecture, RFC 9334. 2023. https://datatracker.ietf.org/doc/html/rfc9334

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

[35] ISO. ISO/IEC 27036 Information security for supplier relationships. 2023. https://www.iso.org/standard/59648.html

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

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

[38] European Union. Regulation (EU) 2024/2847 Cyber Resilience Act. 2024. https://eur-lex.europa.eu/eli/reg/2024/2847/oj

[39] European Union. Directive (EU) 2022/2555 on cybersecurity. 2022. https://eur-lex.europa.eu/eli/dir/2022/2555/oj

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

[41] MITRE. ATLAS. 2026. https://atlas.mitre.org/

[42] MITRE. ATT&CK Software Discovery. 2026. https://attack.mitre.org/

[43] ENISA. Cybersecurity of AI and Standardisation. 2023. https://www.enisa.europa.eu/publications/cybersecurity-of-ai-and-standardisation

[44] OECD. OECD AI Principles. 2024. https://oecd.ai/en/ai-principles

[45] UK Government. AI Cyber Security Code of Practice. 2025. https://www.gov.uk/government/publications/ai-cyber-security-code-of-practice

[46] UK NCSC. Guidelines for Secure AI System Development. 2023. https://www.ncsc.gov.uk/collection/guidelines-secure-ai-system-development

[47] U.S. Department of Commerce. SBOM Minimum Elements. 2021. https://www.ntia.gov/files/ntia/publications/sbom_minimum_elements_report.pdf

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

[49] Cloud Security Alliance. AI Controls Matrix. 2026. https://cloudsecurityalliance.org/research/ai-controls-matrix

[50] OWASP. Machine Learning Security Top 10. 2026. https://owasp.org/www-project-machine-learning-security-top-10/

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.