The Analyst Hour: Pricing Detection Quality and Productivity in AI SOC Acquisitions

An Acquisition Framework for Detection Quality, Analyst Productivity, Customer Outcomes and Integration Risk

Authored by Chennakeshav Adya

Independent Researcher

September 2026

Abstract

Security operations centres increasingly use machine learning, large language models, automation and agentic workflows to prioritise alerts, enrich investigations, recommend actions and draft reports. Acquirers can therefore encounter persuasive claims about faster triage, lower analyst effort, broader detection and autonomous response. Those claims do not by themselves establish transaction value. A system may create many alerts while finding few material threats, improve measured speed by closing difficult cases early, or reduce visible analyst time while shifting work into tuning, engineering, customer support and incident remediation.

This paper develops an acquisition framework for pricing detection quality and analyst productivity in AI-enabled security operations. The investable unit is a verified security outcome delivered at complete cost: a material behaviour is observed, converted into an actionable case, investigated to an evidenced decision, escalated or contained within the agreed service level, and reconciled to customer and financial records. The framework links ATT&CK-aligned coverage, detection precision, false-positive testing, time measures, model governance, telemetry rights, content operations, analyst cohorts, customer retention and integration readiness.[1][2][3]

NIST's preliminary Cyber AI Profile identifies AI-enabled cyber defence as a distinct focus area and discusses advanced threat detection, reduction of false positives and false negatives, automated response and security-help-desk efficiency. MITRE ATT&CK Evaluations now distinguishes detection coverage, precision and speed, including dedicated benign-activity validation. NIST's Cybersecurity Framework 2.0 and AI Risk Management Framework provide complementary governance structures for cyber outcomes and trustworthy AI.[1][2][4][5]

A hypothetical acquisition illustrates an AI SOC platform serving 180 customers with annual recurring revenue of USD 54 million. Every customer count, price, cost, detection measure, productivity input, probability and valuation figure in the example is a management assumption created solely to demonstrate the method. It is not a forecast, market average or representation of a named business.

The paper concludes that an acquirer should build a detection-to-cash ledger before accepting an AI premium. The ledger connects representative attack behaviours, telemetry, alerts, cases, analyst decisions, response actions, customer outcomes, service credits and cash. Six figures and seven tables convert this evidence into diligence, valuation, consideration design and a 180-day integration programme.

Cybersecurity, privacy, AI governance, employment, competition, foreign investment, accounting, tax, insurance and securities decisions require current advice from qualified specialists in each 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: cybersecurity M&A, security operations centre, artificial intelligence, detection quality, analyst productivity, managed detection and response, valuation, integration

1. Define the acquisition decision

The board should begin with the specific capability it intends to buy. An AI SOC target may sell software, managed detection and response, threat intelligence, incident response, security analytics, case management or a mixture of these. Each business has a different cost base, evidence set and integration problem. A product company can depend on customer analysts to investigate. A managed service can own the investigation outcome while depending on third-party telemetry and tools. A transaction thesis that treats both as the same recurring software asset can misprice delivery risk.

The acquisition memorandum should state the customer problem, the operating perimeter and the expected source of value. The buyer may seek better detection, lower service cost, access to a customer segment, proprietary content, scarce analysts, a data advantage or a platform for consolidation. Each thesis needs an observable test. A claim that AI improves productivity should identify whose time changes, which work is included, what quality remains constant and whether the saving reaches gross margin and cash.

The board should also compare acquisition with partnership, licensing, internal build, minority investment and commercial integration. Acquisition can be appropriate when value depends on coordinated ownership of product, data rights, content teams, customer contracts and delivery operations. A narrower arrangement can be preferable when interoperability creates most of the benefit and the buyer does not need to assume model, service or customer liabilities.

The decision gate should separate proven capability from planned capability. Production telemetry, reproduced benchmark results, accepted customer cases, contracted revenue and trained operating teams can support the verified case. Product road maps, model demonstrations, non-binding pipeline and assumed labour removal belong in an option case until evidence is complete.

The buyer should assign an accountable owner to each claimed source of value and identify the evidence available before signing, before closing and after integration. This timing matters because some tests require customer permission, production data or a controlled environment. Missing evidence can be addressed through conditions, retained consideration, covenants or an explicit exclusion from price. The board should receive the same definition of success that product, security, operations and finance teams will use after closing.

Figure 1. Detection-to-cash evidence chain

The proposed chain connects observable security behaviour to an evidenced decision, customer outcome and collected cash.

2. Define the unit of value

The proposed unit of value is a verified security outcome delivered at complete cost. A verified outcome begins with behaviour inside the contracted monitoring boundary. The platform observes sufficient telemetry, creates or enriches a case, supports an analyst or automated control in reaching an evidenced decision, and initiates an authorised response. The outcome includes documentation, customer communication and post-incident learning where these are within scope.

Complete cost includes data ingestion, storage, third-party licences, model inference, detection engineering, threat research, analyst labour, quality review, incident escalation, customer success, service credits, infrastructure, compliance and working capital. A vendor can appear efficient when it excludes costly telemetry or relies on customers to complete investigation and response. The acquisition model should preserve the actual service boundary.

The analyst hour is an important denominator because it can be observed and priced. It should be allocated across triage, investigation, escalation, response, tuning, threat hunting, customer communication, quality assurance, training and administration. Time saved in one category can reappear elsewhere. Generative summaries may reduce writing time while increasing verification effort. Automated closure may reduce queues while raising missed-detection risk.

Value therefore requires a paired measure: outcome quality and complete effort. The buyer should compare cases of similar severity, telemetry, customer environment and service level. Aggregate alerts per analyst can reward noisy products. Incidents closed per analyst can reward premature closure. Verified outcomes per complete analyst hour, supported by quality and customer measures, gives a more defensible basis for pricing.

3. Normalise the product and service perimeter

The diligence team should map every product module and service promise to an operating responsibility. Modules can include endpoint, identity, email, cloud, network, application, data, operational technology, threat intelligence, user behaviour, automation and case management. Services can include monitoring, triage, investigation, threat hunting, containment, recovery advice, compliance reporting and incident response retainers.

The map should identify which capabilities are proprietary, licensed, open source or supplied by the customer. It should record data sources, deployment model, supported environments, geographic restrictions, service hours, escalation paths and authorised response. A demonstration using the target's preferred telemetry can overstate performance for customers with incomplete instrumentation or different technology stacks.

Contract language should be reconciled with delivery. Marketing may describe continuous monitoring while contracts exclude selected assets, log sources, threat classes or response actions. The buyer should sample statements of work, onboarding records, asset inventories, data-ingestion health and case histories. Unsupported systems and unmonitored periods should be visible in the cohort analysis.

The perimeter also determines regulatory and liability exposure. Automated isolation, identity suspension and cloud changes can affect production systems and employees. Advice-only services have a different risk profile. The target should maintain authority matrices, playbooks and customer approvals that match the actions its platform can initiate.

4. Build the detection-to-outcome ledger

The ledger should give the acquirer a single evidence model across product, service and finance. Each sampled case records the relevant behaviour, technique, asset, data source, detection content, model

version, alert, enrichment, analyst actions, decision, response, customer communication, service-level result and commercial consequence. The purpose is traceability, not surveillance of individual employees.

The ledger should preserve negative evidence. Missed behaviours, duplicate alerts, suppressed signals, incomplete telemetry, reverted automations and customer-disputed cases are necessary for a balanced view. A data room containing only showcase incidents cannot support population-level conclusions. Samples should cover customers, products, regions, severities, telemetry maturity and periods of operational stress.

NIST CSF 2.0 organises cyber risk outcomes across Govern, Identify, Protect, Detect, Respond and Recover. That structure can help connect detection operations with enterprise risk and customer responsibility.[1] The ledger can also map observed behaviours to ATT&CK techniques while retaining the actual evidence behind the mapping.[6]

Finance should connect the same cohorts to contracted annual recurring revenue, consumption charges, professional services, credits, renewals, expansion and collections. This allows the buyer to ask whether superior outcomes produce retention, pricing power or lower delivery cost. It also exposes cases where apparently attractive revenue depends on extensive unpriced analyst work.

Table 1. Detection-to-outcome diligence ledger

Evidence field	Required record	Primary owner	Valuation use
Behaviour and asset	timestamped activity and affected environment	customer and telemetry owner	defines test population
Detection content	rule, model, version and change history	detection engineering	tests repeatability
Case evidence	alert context, enrichment and linked activity	product and SOC	measures actionability
Analyst decision	disposition, rationale and reviewer	SOC operations	measures quality and effort
Response	authorised action and completion	customer and responder	measures operational outcome
Service result	SLA, communication, credit and dispute	service management	measures contract performance
Commercial record	revenue, cost, renewal and cash	finance	connects evidence to value

The ledger connects technical evidence, human work, customer acceptance and transaction value.

5. Test detection coverage

Coverage asks whether the product observes and identifies the behaviours relevant to the customer's risk and contracted boundary. It should be tested at the technique, data-source and environment levels. A broad ATT&CK heat map can conceal that coverage is inferred from documentation, depends on unavailable telemetry or produces only low-context visibility.

MITRE ATT&CK provides a common knowledge base of adversary tactics and techniques, while ATT&CK Evaluations exercise vendor capabilities against documented scenarios.[2][6] The 2026 evaluation structure distinguishes detection coverage, precision and speed. An acquirer can use the concepts without converting an external evaluation into a universal ranking. Test configuration, scope and customer environment matter.

The target should maintain a coverage register linking each supported behaviour to required data, content, validation date, expected alert or case, known limitations and owner. The buyer should reproduce a

representative subset in a controlled environment. The test should include multi-stage activity, identity and cloud behaviours, living-off-the-land techniques and expected benign lookalikes.

Coverage quality includes context. A case that says a technique may have occurred has less operational value than one that identifies who acted, what happened, when and where it occurred, how the activity developed, why it matters and what action is recommended. Missing context transfers effort to the analyst and lengthens customer response.

6. Test precision and false positives

Precision measures how much of the surfaced work is useful. It should be calculated at alert, case and customer levels because correlation can convert many low-level alerts into one actionable case. A product can report low alert false-positive rates while analysts spend substantial time joining fragments. Another product can produce fewer cases but omit weak signals that are useful in combination.

The buyer should construct benign-activity tests that reflect normal administration, software deployment, remote access, cloud storage, scripting and data processing. MITRE's current evaluation methodology explicitly includes legitimate business activity in precision testing.[2] Customer-specific allow lists and suppression rules should be recorded because highly tuned demonstrations may not transfer after acquisition.

False negatives require controlled testing and retrospective incident review. No diligence programme can prove the absence of misses. It can test selected behaviours, examine known incidents, compare threat-hunting findings with alert history and review customer escalations that originated outside the platform. The residual limitation should be stated rather than hidden behind a single detection rate.

Precision also has a severity dimension. A false positive that triggers automated isolation can cause more harm than an informational alert. A missed low-impact reconnaissance event differs from a missed destructive action. The target should measure errors by decision consequence and service obligation.

Figure 2. Detection-quality measurement stack

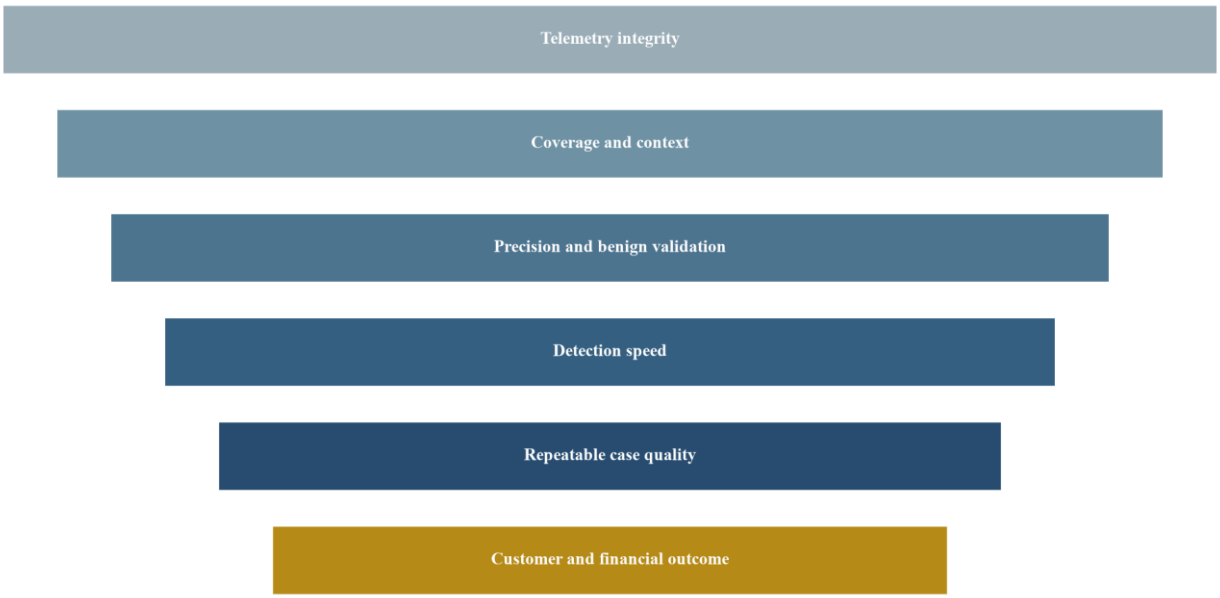

The proposed stack requires coverage, context, precision, speed and repeatability before customer and financial outcomes are attributed.

Table 2. Representative detection benchmark

Test dimension	Evidence	Measure	Diligence warning
Technique coverage	replayed representative behaviours	evidenced detections divided by tested behaviours	claimed coverage without required telemetry
Case context	seven decision-relevant context elements	complete-context cases divided by cases	technique label without investigation value
Precision	mixed malicious and benign population	correct decisions divided by decisions	tuning based on known test scripts
Case consolidation	linked multi-stage activity	alerts and fragments per accepted case	low alert rate caused by missing visibility
Speed	event, alert, case and escalation timestamps	distribution by severity and cohort	averages that exclude unresolved cases
Repeatability	repeated tests after updates	stable results across versions	one-off demonstration environment

The benchmark design separates the attack population, benign population, context and required decision.

7. Measure speed without gaming

Speed should be measured as a sequence rather than one headline number. Relevant timestamps include behaviour occurrence, telemetry arrival, automated detection, case creation, analyst opening, decision, escalation, containment and customer acknowledgement. Each interval has a different owner and economic meaning.

Mean time can be distorted by a long tail. The buyer should inspect medians, upper percentiles, unresolved cases and severity cohorts. Queue pauses, customer waiting time and data latency should be reported separately. A target may control detection speed while the customer controls containment. The contract and value model should reflect this boundary.

Automation can improve speed by enriching entities, retrieving context, grouping related events and preparing playbook steps. It can also create fast but weak dispositions. The diligence sample should pair time with correctness, evidence completeness and review outcomes. A rapid closure later reopened as an incident is not a productivity gain.

Service levels should be tested against system records rather than management reports alone. Case-platform exports, message timestamps and customer records should reconcile. Exclusions, planned maintenance and severity changes need consistent rules. The buyer should quantify credits, disputes and unpriced remediation associated with missed levels.

8. Measure analyst productivity

Analyst productivity is the amount of accepted security work completed at defined quality and complete labour cost. The workforce ledger should classify roles, shifts, seniority, location, employment model, productive hours, leave, training, management, quality assurance and engineering support. Contractor and offshore capacity should be included where it contributes to delivery.

The work ledger should allocate time to triage, investigation, response, hunting, detection tuning, customer communication, reporting, onboarding and administration. If direct time capture is incomplete, the buyer can triangulate workflow timestamps, staffing schedules, sampled observation and manager records. The limitation should remain explicit.

AI assistance should be tested through comparable cohorts. One useful design compares cases with and without a feature while controlling severity, telemetry and analyst experience. Measures include time,

correctness, evidence completeness, escalation and rework. The product team should disclose model changes and feature adoption because nominal access does not establish use.

Productivity value reaches the acquisition model only when it affects capacity, service quality, customer growth or cost. Saved minutes that remain fragmented across a shift may create resilience without reducing headcount. That benefit can still be valuable, but it should be modelled as capacity, quality or avoided overtime rather than immediate labour removal.

Table 3. Analyst productivity ledger

Measure	Numerator	Denominator	Required control
Accepted cases per hour	reviewed cases meeting evidence standard	complete productive hours	severity and telemetry cohort
Investigation effort	direct and support minutes	accepted case	include rework and escalation
Decision accuracy	correct dispositions after review	reviewed dispositions	independent quality sample
Escalation quality	accepted material escalations	escalations	customer confirmation and severity rules
Rework	reopened or materially corrected cases	closed cases	consistent reopen window
Capacity value	additional controlled workload	available team hours	maintain quality and service level
Cash conversion	collected recurring cash less delivery cost	contracted service cohort	reconcile credits and support costs

The ledger pairs complete effort with accepted quality and commercial consequences.

9. Validate AI model claims

The target should maintain an inventory of models and AI-enabled features. Each entry identifies purpose, owner, provider, version, training or retrieval data, deployment boundary, human decision role, evaluation, monitoring, fallback and retirement. Features marketed under one AI label may rely on rules, supervised models, foundation-model services and analyst workflows with different risks.

NIST AI RMF organises governance through Govern, Map, Measure and Manage. Its Generative AI Profile adds actions for risks associated with generative systems.[4][5] NIST's Cyber AI Profile addresses securing AI, using AI for defence and thwarting AI-enabled attacks.[3] These sources support a diligence structure; they do not certify a target's claims.

Evaluation should reflect the actual task. Summarisation can be tested for factual consistency and omitted material. Classification can be tested for precision, recall and calibration. Recommendations can be reviewed for authority, evidence and harmful actions. Agentic workflows require tool permissions, approval points, state management, audit logs and safe failure.

The buyer should assess prompt injection, data leakage, poisoned retrieval, model drift, adversarial inputs, hallucinated evidence and vendor dependency. NIST's adversarial machine-learning taxonomy provides terminology for attacks and mitigations.[7] Production monitoring should detect material change in input, output, error and human-override patterns.

10. Protect telemetry, data and model rights

AI SOC value can depend on access to customer telemetry, case history, threat intelligence, annotations and detection content. The rights register should distinguish ownership, permitted processing, model

improvement, aggregation, retention, residency, transfer, deletion and post-termination use. Customer data cannot be treated as a transferable training asset without enforceable rights.

The acquirer should map data lineage from source to storage, features, prompts, model outputs, cases and reporting. It should identify personal data, secrets, regulated records, cross-border transfers and customer-specific restrictions. Encryption, access, segregation, logging, retention and deletion controls should be tested against architecture and practice.

Vendor contracts deserve particular attention. Foundation-model, cloud, enrichment and threat-intelligence providers can restrict training, redistribution, benchmarking or change of control. Consumption pricing can make apparently high-margin features expensive at scale. Service changes or termination can impair a core workflow.

Data advantage should be valued only where rights, quality, representativeness and operational use are established. A large event count can be dominated by duplicate or low-information records. Annotations may reflect inconsistent analyst decisions. The buyer should test whether data improves a defined task and whether the benefit survives customer and regulatory obligations.

11. Assess detection engineering and threat research

Detection content is a living product. The target should show how threat intelligence becomes hypotheses, data requirements, rules, models, tests, releases, monitoring and retirement. The process should include peer review, version control, reproducible tests, customer applicability and rollback.

Content economics include research, engineering, test infrastructure, telemetry validation, tuning, documentation and support. A large rule count is not evidence of quality. Rules can be duplicate, disabled, obsolete or dependent on unavailable fields. The buyer should analyse enabled content, firing distribution, case contribution, maintenance frequency and customer coverage.

Open standards can improve portability. STIX and TAXII support structured threat-information exchange, Sigma provides a generic signature format, YARA supports pattern matching and OCSF provides a security-event schema.[8][9][10][11] The target's value may lie in validated logic, operating data and deployment workflows rather than proprietary syntax.

Key-person risk is material when a small research team maintains critical content or customer knowledge. Documentation, review, release authority and succession should be assessed. Earn-outs tied only to new rule volume can reward activity without improving outcomes.

12. Test architecture and integration

The architecture review should trace data ingestion, normalisation, storage, enrichment, detection, correlation, case management, automation and reporting. It should identify tenant isolation, identity boundaries, secrets, administrative access, resilience, backup, recovery and observability. Diagram claims should be reconciled with deployed configurations.

Integration value depends on connectors and schemas. The buyer should inventory supported sources, versions, authentication methods, field mappings, data loss, latency, rate limits and maintenance ownership. Revenue attached to bespoke connectors can carry hidden engineering cost.

The integration plan should preserve detection quality during platform change. Migrating telemetry, identity, cases or content can alter fields and suppress evidence. Dual running, shadow testing and cohort cutovers allow comparison. A rapid forced migration may save duplicated cost while increasing customer and incident risk.

NIST SP 800-207 provides zero-trust architecture principles and SP 800-53 provides a broad control catalogue.[12][13] These can inform architecture review alongside product-specific threat modelling. The buyer should state which controls apply and how evidence is produced.

Figure 3. Acquisition measurement architecture

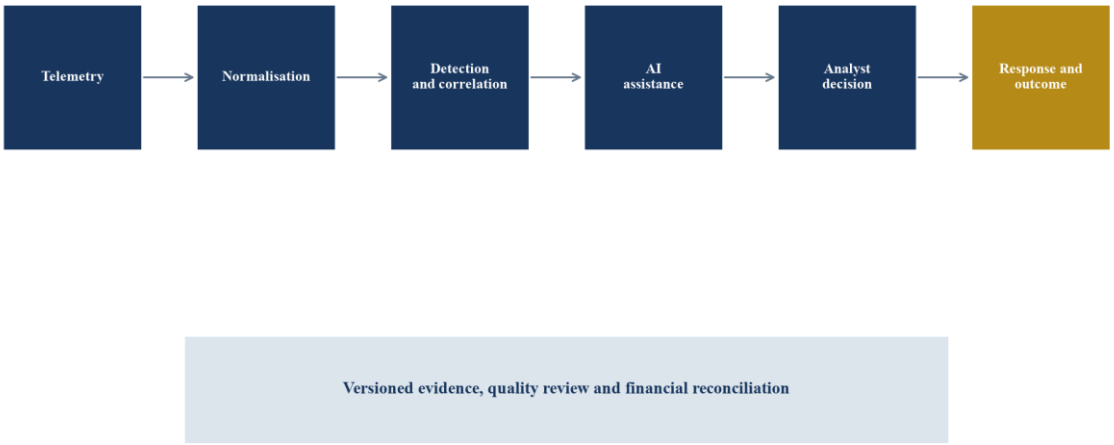

The proposed architecture keeps source telemetry, detection logic, AI assistance, human decisions and customer outcomes independently auditable.

13. Analyse customer cohorts

Customer quality should be analysed by product, service, size, industry, geography, telemetry maturity, contract term and acquisition channel. Annual recurring revenue should reconcile to executed contracts, invoices, deferred revenue and cash. Usage-based and professional-services elements should be separated.

Retention measures need transparent rules. Gross retention shows the recurring base preserved before expansion. Net retention adds expansion and contraction. Logo retention can conceal loss of large customers. Cohorts should include service credits, concessions, free periods and support cost. A renewal achieved through extensive unpriced work may preserve revenue while reducing value.

Security outcomes can influence retention, but causation requires evidence. The buyer can compare renewal, expansion, satisfaction, escalations and incident performance across controlled cohorts. Product adoption, account management and contract changes should be considered. Testimonials and selected wins do not establish population results.

Concentration and change-of-control rights should be modelled. A target may rely on a channel partner, cloud marketplace or a small number of regulated customers. Consent, termination, price review, data transfer and security-assessment obligations can affect closing and integration.

Table 4. Customer-cohort evidence matrix

Cohort dimension	Commercial measure	Operating evidence	Transaction question
Product and service	ARR, margin and expansion	enabled modules and delivered scope	which offer creates retained value
Telemetry maturity	price and support cost	source health and coverage	does performance depend on ideal data
Industry and geography	retention and contract term	regulatory and threat profile	can the model transfer

Cohort dimension	Commercial measure	Operating evidence	Transaction question
Customer size	ACV and sales efficiency	case volume and service effort	is delivery cost scalable
Incident experience	renewal and credits	detection, response and dispute	do outcomes support trust
Acquisition channel	retention and commission	partner dependency	who controls the relationship

The matrix links customer economics to delivery evidence rather than aggregate revenue alone.

14. Price managed service delivery

Managed detection and response combines software economics with people, process and contractual accountability. The buyer should construct a service-cost model by customer cohort. It includes analysts, threat hunters, incident responders, service managers, detection engineers, platform cost, data retention, third-party licences, quality assurance and on-call capacity.

Shared operations can create scale when workload is pooled and standardised. They can also create tail risk because severe incidents cluster and require senior attention. Staffing models should include surge, absence, training, attrition and regional coverage. Follow-the-sun arrangements need clear handover and consistent decision authority.

Pricing should reflect asset count, users, data volume, service hours, response authority, retention, compliance and incident support. A flat price can work with controlled boundaries. It can become uneconomic when customers add telemetry, assets and obligations without repricing. The diligence model should identify contract caps, overages and customary concessions.

Gross margin improvement may come from automation, standardisation, infrastructure procurement, workload routing and product simplification. Each initiative needs an implementation cost and quality guardrail. Immediate removal of senior analysts can impair escalation, customer trust and content development.

15. Evaluate security, privacy and regulatory exposure

A cybersecurity target must protect its own platform and service operations. The buyer should examine secure development, vulnerability management, penetration tests, identity, privileged access, tenant isolation, secrets, supply chain, incident response, backup and recovery. CISA's Secure by Design guidance places responsibility on technology manufacturers to design safer products and reduce customer burden.[14]

Known incidents should be traced from detection through remediation, customer notice, regulator engagement, insurance and control improvement. The SEC's cybersecurity disclosure rules focus on material impacts and governance for covered registrants.[15] Other jurisdictions and sectors impose different obligations. Counsel should determine the actual requirements.

Privacy review should cover collection, purpose, retention, employee monitoring, automated decision support, cross-border transfer and data-subject rights. Security telemetry can contain personal and sensitive information. Contract permission does not replace applicable law.

Regulatory readiness should be evidenced through control operation rather than certificate inventory alone. ISO 27001, ISO 27035 and ISO 42001 can structure information security, incident management and AI management systems.[16][17][18] Scope, exclusions, audit findings and remediation remain material.

16. Assess workforce and operating dependency

The workforce analysis should identify critical roles, experience, location, compensation, notice, retention, clearance, customer relationships and access. The target may depend on a few people who understand detection architecture, threat research, key accounts or incident response. Organisation charts rarely show this dependency.

The NICE Workforce Framework provides a common language for cybersecurity work and competencies.[19] The buyer can map roles to actual tasks and evidence. Job titles alone can overstate depth. Representative interviews and work products help distinguish supervised capacity from independent judgement.

AI changes role design. Entry-level triage may shrink while demand grows for detection engineering, model evaluation, customer advisory and complex investigation. The integration plan should preserve learning pathways. Eliminating junior work without a route to develop senior capability can create a future talent gap.

Retention arrangements should align with value and conduct. Transaction bonuses can preserve continuity, while long earn-outs tied to uncontrollable revenue can weaken motivation. Access, conflicts, intellectual property and post-employment restrictions require jurisdiction-specific advice.

17. Test revenue quality and cash conversion

Revenue quality depends on enforceable contracts, accepted delivery and repeatable economics. The buyer should reconcile bookings, contracted annual recurring revenue, reported recurring revenue, invoicing, revenue recognition, deferred balances and collections. Multi-year headline contract value should not be confused with annual recurring cash.

IFRS 15 requires revenue recognition to follow identified performance obligations and transfer of promised goods or services.[20] Cybersecurity contracts can combine licences, implementation, support, monitoring, incident response and variable usage. Accounting treatment should be reviewed by qualified advisers using the actual terms.

Cash conversion can be affected by annual prepayment, channel collection, cloud marketplace settlement, service credits, incident disputes and data-volume adjustments. The buyer should analyse days sales outstanding, renewals, refunds, credit notes and bad debt by cohort. Strong reported growth with declining cash conversion warrants investigation.

Customer acquisition cost should include sales, marketing, solutions engineering, proof of concept, security review, onboarding and partner commission. Payback and lifetime value should use cohort margin and retention, not aggregate gross margin. Integration assumptions should preserve the sales effort required for regulated customers.

18. Build a hypothetical acquisition case

Consider a hypothetical AI SOC company with 180 customers and USD 54 million of contracted annual recurring revenue. It combines a software platform with managed detection and response. The buyer expects value from stronger detection content, automation of routine investigation, cross-selling and infrastructure consolidation.

Every number in this section is a management assumption created solely to demonstrate the framework. The model assumes USD 47 million of recognised recurring revenue, USD 6 million of professional and incident-response revenue, 72 percent reported gross margin, 124 operations employees and 36 product and detection-engineering employees. It assumes no named market multiple and makes no prediction.

The diligence sample covers 36 customers representing different sectors and telemetry maturity. The model assumes that 68 percent of surfaced cases meet the complete evidence standard before remediation, 18 percent require material analyst reconstruction, 9 percent are duplicates or avoidable noise and 5 percent are unresolved within the review window. These are illustrative classifications.

The base case recognises recurring cash flow supported by contracts and cohort delivery. The evidence-backed productivity case recognises capacity where controlled tests preserve decision quality. Cross-sell and new autonomous features remain option value until customer adoption, rights, controls and economics are demonstrated.

Figure 4. Hypothetical case and analyst-hour funnel

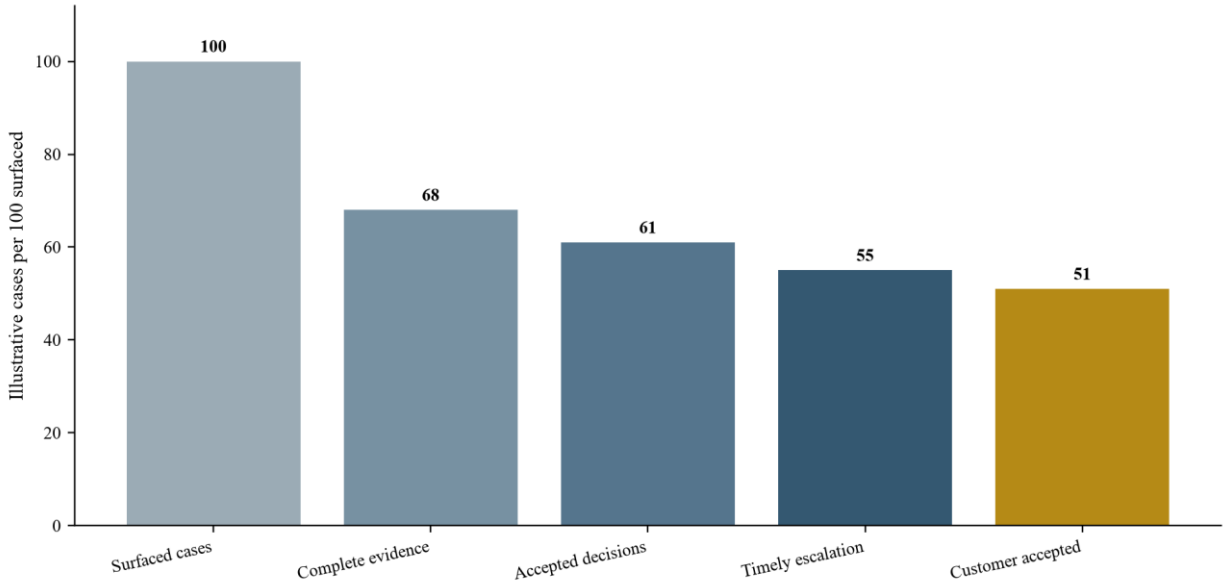

Every value in the figure is a management assumption created solely to illustrate the acquisition method.

Table 5. Hypothetical platform economics

Measure	Reported position	Diligence-adjusted base	Evidence-backed opportunity
Contracted ARR	USD 54m	USD 51m after concessions and scope review	USD 57m with contracted expansion
Recognised recurring revenue	USD 47m	USD 46m	USD 51m after verified deployment
Gross margin	72%	66% including complete service cost	70% after controlled automation
Operations employees	124	124	capacity equivalent to 14 roles, not immediate removal
Complete-evidence cases	not reported	68 per 100 surfaced	82 after content and workflow remediation
Upper-percentile escalation time	not reported	94 minutes	62 minutes after validated integration
Annual service credits	USD 0.4m	USD 0.7m including disputed items	USD 0.4m after service improvement

All figures are management assumptions for method illustration and are not market data or a forecast.

19. Stress the acquisition model

The model should be stressed for telemetry cost, model consumption, customer loss, false-positive growth, analyst attrition, incident surge, delayed integration and regulatory remediation. Each stress should connect to cash, service quality and capital need. A percentage reduction in revenue without an operating mechanism gives limited insight.

One downside assumes that data-ingestion and inference cost rise while customers resist repricing. Another assumes a platform migration increases duplicate cases and service credits for two quarters. A third assumes senior analyst attrition slows escalation and content release. The combined downside should test liquidity and covenant headroom.

Detection degradation can be modelled through case quality and effort. Lower precision increases triage. Missing context increases investigation. Slower escalation increases customer and liability exposure. The model should avoid inventing loss avoidance where the target lacks observed evidence.

Integration upside should also be staged. Procurement savings can begin after contracts are changed. Infrastructure savings require migration and resilience testing. Labour capacity requires validated workflow adoption. Revenue synergy requires customer consent, product readiness and a sales process. Timing and cost should be explicit.

20. Value the target in layers

The first valuation layer is verified recurring operating cash flow. It uses contract, cohort, delivery-cost and cash evidence. The second layer can recognise evidence-backed productivity where controlled tests demonstrate repeatable capacity or quality. The third layer covers contracted expansion with completed product and delivery prerequisites.

Data, models, content and integrations can support strategic value where rights, differentiation and portability are established. They should not receive a premium solely because they are labelled AI. Intangible-asset recognition, useful life and impairment require advice under IFRS 3, IAS 38 and IAS 36.[21][22][23] Fair-value measurement should reflect market-participant assumptions under IFRS 13.[24]

Liabilities and dependency costs should be deducted or protected. These can include service credits, incident exposure, customer concentration, unpriced data, model-provider dependency, privacy remediation, content debt, retention, duplicated infrastructure and integration capital. The treatment can use price, escrow, indemnity, retention, earn-out or closing conditions.

Consideration design should match the evidence gap. Revenue earn-outs can reward low-quality growth unless margin, retention and service outcomes are included. Product milestones should require deployed customer acceptance rather than feature release. Detection-quality milestones need agreed populations, data access, review authority and dispute mechanics.

Figure 5. Illustrative acquisition-value bridge

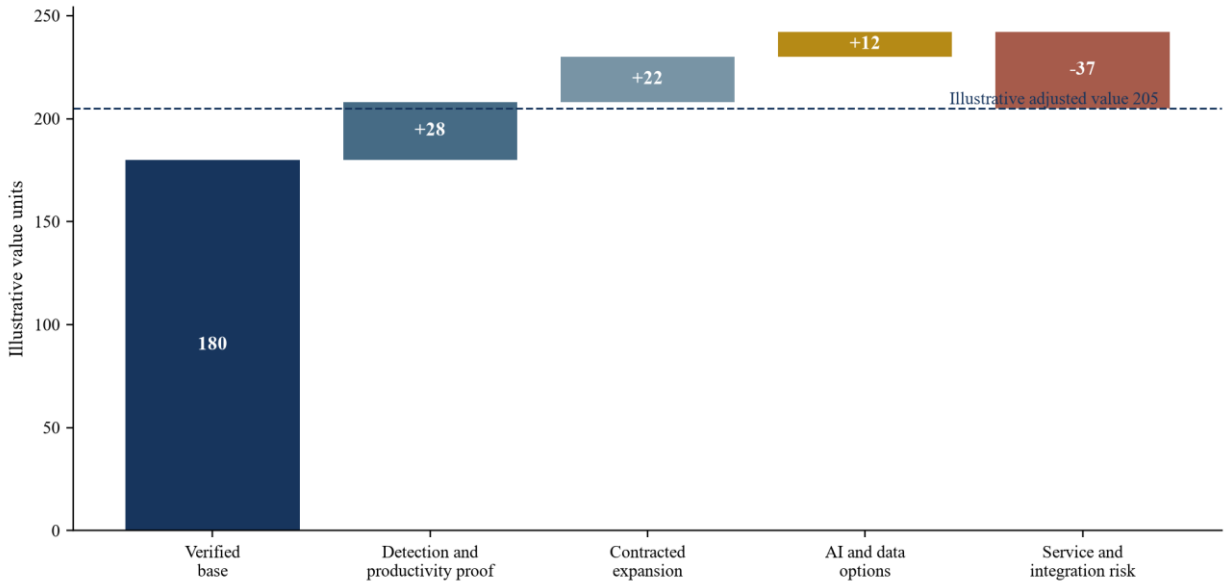

The bridge separates verified cash flow from evidenced operational value, options and risk deductions.

21. Design the integration plan

Integration should preserve the evidence chain while combining capabilities. The target state should identify which telemetry, content, models, case systems, analyst teams, customer contracts and brands remain, migrate or retire. Decision criteria should be agreed before political attachment to either platform dominates.

The buyer should run representative cohorts in parallel. Shadow detections, duplicate case review and controlled customer pilots can show whether coverage, precision, speed and effort improve. Differences should be investigated before broad cutover. Critical incidents require a clear system of record and command authority during transition.

People integration should protect customer and operational continuity. Detection engineers, threat researchers, senior analysts, service managers and product owners may hold tacit knowledge. Retention, documentation, paired work and succession should precede structural removal. Access and conflict controls should be updated at close.

Customer communication should explain service continuity, data processing, product changes and support. Contract consent and security assessments can govern timing. A migration plan that assumes immediate cross-sell or forced conversion can increase churn. The board should review evidence gates for each cohort.

Table 6. Integration-risk register

Integration domain	Principal risk	Required evidence	Release gate
Telemetry and schema	field loss or latency	replay and completeness comparison	representative sources pass
Detection content	coverage or precision degradation	dual-run benchmark	quality thresholds pass
AI features	changed output or cost	versioned evaluation and consumption	task and cost thresholds pass
Case platform	broken history and workflow	migration reconciliation	sampled cases match

Integration domain	Principal risk	Required evidence	Release gate
Workforce	loss of critical judgement	role and succession map	coverage and handover complete
Customers	consent, trust or churn	contract and communication plan	cohort approval complete
Finance	delayed savings or hidden cost	owner, baseline and realised cash	benefits independently reconciled

The register connects each integration choice to evidence and a release gate.

22. Execute a 180-day diligence and integration programme

Days 1 to 30 establish control. The team confirms the transaction perimeter, model inventory, telemetry rights, customer contracts, workforce map and detection-to-outcome ledger. It preserves logs, versions and case evidence. Urgent security, privacy or customer issues receive accountable owners and interim controls.

Days 31 to 75 test representative behaviour and benign activity. The team reproduces coverage, context, precision and speed measures, samples analyst work and reconciles service reports. Finance rebuilds customer cohorts, complete delivery cost, recurring revenue and cash conversion.

Days 76 to 120 stress the operating and transaction model. The team tests dependency on data sources, model providers, key people, customers and channels. It models incident surge, attrition, platform migration, data-cost change and regulatory remediation. Management assumptions remain clearly identified.

Days 121 to 180 finalise consideration protection and the integration release plan. The parties agree customer, product, people, data and platform gates. The board receives verified value, evidence-backed value, option value, liabilities, first-year investment and stop-loss conditions.

Figure 6. Evidence-gated 180-day programme

The programme moves from control and measurement to stress testing, transaction protection and cohort release.

Table 7. Evidence-gated transaction and integration plan

Gate	Evidence	Decision enabled	Failure response
Perimeter	product, service, contract and responsibility map	confirm acquisition scope	exclude or reprice unsupported scope
Detection quality	representative coverage, precision and speed	accept core capability	remediate, retain value or stop

Gate	Evidence	Decision enabled	Failure response
Productivity	controlled effort and quality cohorts	recognise capacity value	remove labour synergy
Data and AI	rights, lineage, evaluation and fallback	recognise AI and data value	restrict use or reduce value
Customer economics	contract, delivery cost, retention and cash	accept recurring base	adjust revenue and working capital
Integration readiness	dual-run plan, people coverage and consent	release cohort migration	extend parallel operation
Realised value	independently reconciled quality and cash	release retained consideration	defer or reduce payment

Each gate requires reviewable evidence before consideration, migration or cost action proceeds.

23. Decision and conclusion

AI SOC acquisition value should begin with verified security outcomes and complete delivery cost. Alert volume, model labels and analyst-seat claims provide context, but they do not establish value. The board needs an evidence chain from representative behaviour through telemetry, actionable cases, evidenced decisions, authorised response, customer acceptance and cash.

Detection quality has multiple dimensions. Coverage without context transfers work to analysts. Speed without correctness can accelerate bad decisions. Precision without representative benign testing can reflect tuning to a known script. Productivity without complete labour and quality measures can hide cost and risk.

The transaction should therefore maintain five connected records: the detection-to-outcome ledger, model inventory, data-rights register, analyst productivity ledger and customer cohort model. These records allow technical, operational and financial diligence to use the same evidence. They also create post-close control.

Consideration and integration should follow proof. Verified recurring cash flow supports the base. Reproduced detection quality and controlled productivity support additional value. Contracted expansion can be recognised when delivery prerequisites are complete. AI, data and autonomous-response options remain separately valued until rights, controls, adoption and economics are demonstrated.

This approach gives buyers and sellers a practical basis for price and execution. It preserves innovation while requiring evidence for consequential claims. The analyst hour becomes valuable when it produces a correct, timely and accepted outcome at complete cost.

Frequently asked questions

What is the most useful unit for valuing an AI SOC business?

Use a verified security outcome delivered at complete cost. The measure should connect observable behaviour, telemetry, case evidence, analyst or authorised automated decision, response, customer result and cash.

Why is alert volume a weak acquisition metric?

Alert volume can rise because visibility improves or because noise and duplication increase. It should be analysed with coverage, case consolidation, precision, investigation effort and customer outcomes.

How should a buyer test detection quality?

Use representative malicious behaviours and benign activity across relevant telemetry and customer environments. Measure coverage, context, precision, speed and repeatability, then review known misses and operational limitations.

How can AI productivity be verified?

Compare controlled case cohorts with and without the feature while holding severity, telemetry and analyst experience as consistent as practical. Measure complete effort, correctness, evidence, escalation and rework.

Should analyst time savings be valued as immediate headcount reduction?

Only where the operating plan can realise the saving while preserving quality, coverage and resilience. Fragmented minutes often create capacity or reduce overtime before they support role removal.

How should customer telemetry be treated in valuation?

Value it only where contractual and legal rights permit the relevant use, the data is sufficiently complete and representative, and it improves a defined task. Event count alone is not a data advantage.

What should an earn-out measure?

Measures can combine retained recurring revenue, margin, customer retention, deployment, detection quality and service outcomes. Definitions, populations, data access and review authority should be agreed before closing.

When should platform migration proceed?

After representative dual-run cohorts demonstrate that telemetry completeness, detection quality, case history, service levels and customer obligations remain within approved thresholds.

References

[1] National Institute of Standards and Technology, Cybersecurity Framework 2.0. https://www.nist.gov/cyberframework

[2] MITRE Engenuity, ATT&CK Evaluations Enterprise 2026 evaluation structure. https://attackevals.mitre.org/enterprise/er8

[3] National Institute of Standards and Technology, NIST IR 8596 Cybersecurity Framework Profile for Artificial Intelligence, preliminary draft. https://csrc.nist.gov/pubs/ir/8596/iprd

[4] National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework 1.0. https://www.nist.gov/itl/ai-risk-management-framework

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

[6] MITRE, ATT&CK Enterprise knowledge base. https://attack.mitre.org/

[7] National Institute of Standards and Technology, NIST AI 100-2e2025 Adversarial Machine Learning taxonomy and terminology. https://csrc.nist.gov/pubs/ai/100/2/e2025/final

[8] OASIS Open, STIX Version 2.1. https://docs.oasis-open.org/cti/stix/v2.1/stix-v2.1.html

[9] OASIS Open, TAXII Version 2.1. https://docs.oasis-open.org/cti/taxii/v2.1/taxii-v2.1.html

[10] SigmaHQ, Sigma generic signature format. https://sigmahq.io/

[11] Open Cybersecurity Schema Framework, OCSF schema. https://schema.ocsf.io/

[12] National Institute of Standards and Technology, SP 800-207 Zero Trust Architecture. https://csrc.nist.gov/pubs/sp/800/207/final

[13] National Institute of Standards and Technology, SP 800-53 Revision 5 Security and Privacy Controls. https://csrc.nist.gov/pubs/sp/800/53/r5/upd1/final

[14] Cybersecurity and Infrastructure Security Agency, Secure by Design. https://www.cisa.gov/securebydesign

[15] US Securities and Exchange Commission, Cybersecurity Risk Management Strategy Governance and Incident Disclosure final rule. https://www.sec.gov/files/rules/final/2023/33-11216.pdf

[16] International Organization for Standardization, ISO IEC 27001 Information security management systems. https://www.iso.org/standard/27001

[17] International Organization for Standardization, ISO IEC 27035 Information security incident management. https://www.iso.org/standard/60803.html

[18] International Organization for Standardization, ISO IEC 42001 Artificial intelligence management system. https://www.iso.org/standard/42001

[19] National Institute of Standards and Technology, NICE Workforce Framework for Cybersecurity. https://www.nist.gov/itl/applied-cybersecurity/nice/nice-framework-resource-center

[20] IFRS Foundation, IFRS 15 Revenue from Contracts with Customers. https://www.ifrs.org/issued-standards/list-of-standards/ifrs-15-revenue-from-contracts-with-customers/

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

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

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

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

[25] International Valuation Standards Council, International Valuation Standards. https://ivsc.org/standards/

[26] National Institute of Standards and Technology, SP 800-61 Revision 3 Incident Response Recommendations. https://csrc.nist.gov/pubs/sp/800/61/r3/final

[27] National Institute of Standards and Technology, SP 800-92 Guide to Computer Security Log Management. https://csrc.nist.gov/pubs/sp/800/92/final

[28] National Institute of Standards and Technology, SP 800-137 Information Security Continuous Monitoring. https://csrc.nist.gov/pubs/sp/800/137/final

[29] National Institute of Standards and Technology, SP 1353 initial public draft, Using AI for CSF Analysis and Reporting. https://csrc.nist.gov/pubs/sp/1353/ipd

[30] National Institute of Standards and Technology, AI Resource Center. https://airc.nist.gov/

[31] Cybersecurity and Infrastructure Security Agency, Roadmap for Artificial Intelligence. https://www.cisa.gov/resources-tools/resources/roadmap-artificial-intelligence

[32] Cybersecurity and Infrastructure Security Agency, Software Acquisition Guide for Government Enterprise Consumers. https://www.cisa.gov/resources-tools/resources/software-acquisition-guide-government-enterprise-consumers

[33] Cybersecurity and Infrastructure Security Agency, Known Exploited Vulnerabilities Catalog. https://www.cisa.gov/known-exploited-vulnerabilities-catalog

[34] Cybersecurity and Infrastructure Security Agency and UK National Cyber Security Centre, Guidelines for Secure AI System Development. https://www.cisa.gov/news-events/alerts/2023/11/26/cisa-and-uk-ncsc-unveil-joint-guidelines-secure-ai-system-development

[35] European Union Agency for Cybersecurity, ENISA Threat Landscape 2026. https://www.enisa.europa.eu/topics/cyber-threats/threat-landscape

[36] European Union Agency for Cybersecurity, ENISA Threat Landscape 2025. https://www.enisa.europa.eu/publications/enisa-threat-landscape-2025

[37] European Union Agency for Cybersecurity, Artificial Intelligence Cybersecurity Challenges. https://www.enisa.europa.eu/publications/artificial-intelligence-cybersecurity-challenges

[38] European Union, Directive EU 2022 2555 on measures for a high common level of cybersecurity, NIS2. https://eur-lex.europa.eu/eli/dir/2022/2555/oj

[39] European Union, Regulation EU 2022 2554 on digital operational resilience for the financial sector. https://eur-lex.europa.eu/eli/reg/2022/2554/oj

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

[41] UK National Cyber Security Centre, Guidelines for secure AI system development. https://www.ncsc.gov.uk/collection/guidelines-secure-ai-system-development

[42] UK Information Commissioner's Office, Guidance on AI and data protection. https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/

[43] US Federal Trade Commission, Keep your AI claims in check. https://www.ftc.gov/business-guidance/blog/2023/02/keep-your-ai-claims-check

[44] Saudi National Cybersecurity Authority, Essential Cybersecurity Controls. https://nca.gov.sa/en/regulatory-documents/controls-list/ecc/

[45] UAE Cybersecurity Council, National Cybersecurity Strategy. https://csc.gov.ae/en/strategies/

[46] Qatar National Cyber Security Agency, National Cyber Security Strategy. https://ncsa.gov.qa/en/strategy

[47] Bahrain National Cyber Security Centre, National Cyber Security Strategy. https://www.ncsc.gov.bh/

[48] FIRST, Common Vulnerability Scoring System. https://www.first.org/cvss/

[49] FIRST, Exploit Prediction Scoring System. https://www.first.org/epss/

[50] MITRE, D3FEND cybersecurity countermeasure knowledge graph. https://d3fend.mitre.org/

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.