Proof in the Field: Mission Evidence and Export Controls in Autonomous-Defence M&A

A transaction framework for programme backlog, operational evidence, responsible autonomy and transferable market access

Authored by Chennakeshav Adya

Independent Researcher

September 2026

Abstract

Autonomous-defence businesses combine sensors, communications, software, artificial intelligence, platforms and human command to perform missions in contested and safety-critical environments. Buyers can be attracted by demonstrations, prototype speed, programme announcements and apparent geopolitical demand. Those indicators do not answer the transaction question: which capabilities have been tested against an authorised mission, accepted by a customer, supportable across a defined configuration, and transferable within applicable export, security and end-use controls?

This paper develops an Autonomous-Defence Acquisition Framework. It tests mission definition, programme authority, backlog quality, system configuration, field evidence, testing independence, data provenance, model validation, human judgement, safety, cybersecurity, interoperability, production readiness, intellectual-property rights, export classification, licensing, sanctions, foreign investment, delivery economics and transaction structure. The framework separates demonstration performance from operational evidence, contract value from funded backlog, and technical capability from transferable market access.

A wholly hypothetical case covers 24 government and prime-contractor customers; USD 31.5 million of annual revenue; USD 146 million of management-reported backlog; 38 deployed or trial systems; and 11 countries described as addressable. Management identifies USD 20.0 million of annual customer value and USD 13.2 million of buyer revenue and synergy. Evidence gates retain USD 62 million of funded, executable backlog, 16 systems with controlled mission evidence and four currently supportable export markets. After performance, licence, concentration, production, delivery-cost and cash-conversion adjustments, the framework retains USD 7.6 million of customer value and USD 5.0 million of buyer revenue and synergy. Every figure is an illustrative management assumption. It is not observed company data, a market benchmark, a forecast, a military assessment or a valuation opinion.

The evidence base includes U.S. Department of Defense Directive 3000.09; NATO's revised Artificial Intelligence Strategy; U.S. Government Accountability Office work on weapons testing; the United Nations Convention on Certain Conventional Weapons process; current U.S. ITAR and EAR resources; UK strategic export-control guidance; NIST artificial-intelligence risk-management guidance; and current accounting requirements for revenue, intangible assets, business combinations and fair value. The conclusion is that mission evidence becomes transaction value only when it is configuration-specific,

independently reviewable, connected to customer acceptance and supportable within current legal permissions.

JEL Classification: G32, G34, L64, O32, O38, P48

Keywords: defence M&A, autonomous systems, mission evidence, export controls, programme backlog, responsible AI, testing and evaluation, technology valuation

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

1. DEFINE THE ACQUISITION DECISION

The buyer must decide how much value to assign to an autonomous-defence business whose products can influence surveillance, protection, logistics, navigation, decision support or other authorised missions. The valuable asset is a controlled capability and delivery system. A successful demonstration has limited transaction value when the buyer cannot establish the mission, configuration, test conditions, customer acceptance, legal permissions and support burden.

Autonomy spans different functions and levels of consequence. A navigation aid, maintenance optimiser, uncrewed logistics platform and weapon system present different technical, legal and governance questions. The framework does not assume that every autonomous product is a weapon. It requires the target to define the function and decision boundary.

The acquisition framework tests six propositions. The mission and accountable authority are defined. Evidence is specific to the production configuration. Human judgement and system limits are governed. Contracts and backlog are funded and executable. intellectual property and technical data are transferable. Export and security permissions support the commercial plan. Failure at one layer changes the amount, timing or certainty of value.

Table 1. Evidence required before autonomous-defence capability receives transaction value

Evidence layer	Core question	Minimum record	Transaction consequence
Mission	What authorised task must the system perform?	User, environment, constraints and success criteria	Defines relevance
Configuration	Which hardware, software and data were tested?	Baseline, version, supplier and change history	Defines comparability
Field evidence	Does performance persist under relevant conditions?	Test plan, results, limitations and independent review	Supports reliance
Programme	Is customer demand funded and executable?	Authority, contract, acceptance and schedule	Supports backlog
Export and security	Can technology, service and support transfer?	Classification, licence, end user and handling	Supports market access
Economics	Does delivery produce scalable cash?	Margin, working capital, service burden and collection	Supports transaction value

Mission, configuration, programme and market-access evidence should be tested together.

2. SEPARATE MISSION EFFECT, SYSTEM FUNCTION AND AUTONOMY

Mission effect describes the outcome sought by the authorised customer. System function describes what the product does. Autonomy describes how functions are performed with varying degrees of human input. These concepts should not be collapsed into one label.

The U.S. Department of Defense Directive 3000.09 defines policy for autonomous and semi-autonomous weapon systems and requires appropriate levels of human judgement over the use of force. NATO's responsible-use principles include lawfulness, responsibility and accountability, explainability and traceability, reliability, governability and bias mitigation. These public frameworks demonstrate that technical performance and responsible use must be considered together.

Diligence should map each feature to mission, user, authority and consequence. Marketing terms such as autonomous, AI-enabled or combat-proven should be decomposed into observable functions and evidence. Transaction value should follow the capability that the buyer can lawfully support and reproduce.

3. DEFINE THE PROGRAMME AND CUSTOMER AUTHORITY

Defence demand can originate from a funded programme, research contract, operational trial, framework agreement, memorandum or public announcement. These instruments do not carry equal commercial certainty. The target should identify the contracting authority, appropriation or budget, scope, configuration, quantities, acceptance, options and termination rights.

The buyer should trace each programme claim to executed documents and customer records. Public statements can support context while contract and funding evidence determine backlog. Security restrictions may limit disclosure, so diligence can use controlled data rooms, counsel review and customer confirmations without moving protected material into unrestricted systems.

Programme value also depends on schedule and dependencies. A prime-contractor relationship, government-furnished equipment, range access or certification can control delivery. The transaction model should distinguish awarded work, options, bids and addressable programmes. Pipeline should not become backlog through terminology.

4. BUILD THE PROGRAMME-BACKLOG REGISTER

The backlog register should record customer, programme, contract, funding, deliverable, quantity, price, acceptance, schedule, export condition, supplier dependency, cancellation and cash milestone. It should reconcile with accounting and project systems. Management forecasts should preserve the difference between contracted backlog and opportunity.

Multi-year authorisations and licences can cover broad activities without proving actual transfer or revenue. The U.S. Department of State has noted in public reporting that authorisation values do not correlate precisely with articles transferred. The same discipline applies in M&A: permission, contract, delivery and revenue are separate evidence states.

Diligence should sample high-value, overdue and low-margin programmes. It should identify unfunded options, customer convenience termination, liquidated damages and acceptance risk. The

buyer should value executable backlog after delivery cost and working capital, not headline contract value.

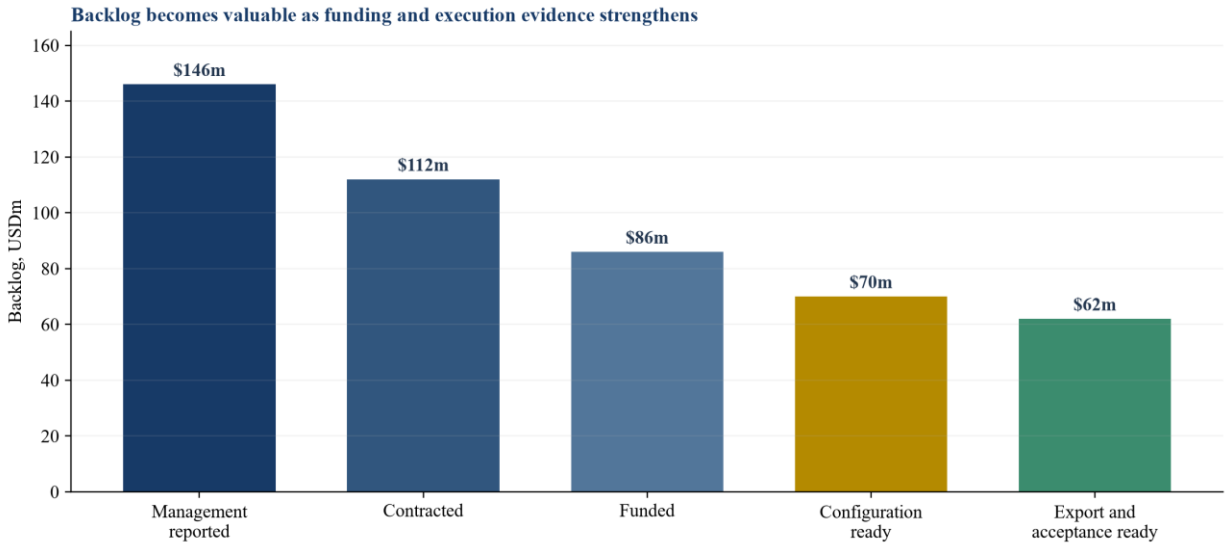

Figure 1. Hypothetical programme backlog by evidence state

Values in USD millions are illustrative management assumptions, not observed company data.

5. FREEZE THE SYSTEM CONFIGURATION

Autonomous-system performance depends on hardware, software, sensors, models, training data, communications, payload, environment and operator procedure. A test result belongs to the configuration tested. Later code, supplier or component changes can affect comparability.

The target should maintain configuration baselines, software bills of materials, model versions, data lineage, firmware, calibration, interfaces and release approvals. Customer-specific forks and field modifications should be visible. Serial-number and deployment records should connect delivered systems to controlled baselines.

The buyer should reproduce a sample from source code and components to test result and delivered unit. A platform assembled for a demonstration may differ from the producible system. Technical debt and obsolete components belong in cost and schedule forecasts. Value should follow the production configuration, with prototype learning treated as supporting evidence.

6. DEFINE THE OPERATIONAL ENVELOPE

The operational envelope describes where and how the system is intended to function. It can include terrain, weather, lighting, electromagnetic conditions, communications, speed, mission duration, operator skill and adversarial interference. Performance outside the envelope should not be implied by a broad success rate.

The target should state boundaries, fallback modes and prohibited uses. Environmental and threat assumptions should be versioned. A system tested in controlled conditions may require additional evidence before use in a contested or degraded environment.

Diligence should compare claimed missions with test conditions and customer concepts of operation. It should inspect edge cases, aborts and unavailable data. Market opportunity across new domains should remain separate from demonstrated capability. A clear envelope can strengthen value because it supports responsible deployment and customer trust.

7. BUILD THE MISSION-EVIDENCE REGISTER

The evidence register should link requirement, configuration, test method, environment, result, anomaly, reviewer, limitation and customer decision. Evidence can include laboratory, simulation, hardware-in-the-loop, range, exercise and operational records. Each type has a defined purpose.

Simulation can explore many scenarios while inheriting model assumptions. Controlled trials can isolate variables while underrepresenting field complexity. Operational evidence can be realistic while affected by selection, reporting and security constraints. The target should preserve these distinctions.

The buyer should inspect raw and summarised results under appropriate access controls. It should test whether adverse outcomes and incomplete missions were retained. A curated video or selected success does not establish repeatability. Evidence quality depends on sampling, comparability and independent review.

8. REQUIRE INDEPENDENT TESTING AND EVALUATION

Testing should involve reviewers independent of development in authority and evidence, with rigor proportionate to consequence. The U.S. GAO has identified modernisation needs in Department of Defense testing policy and highlighted early tester involvement, iterative testing, digital threads and user feedback as leading practices. These principles provide a useful diligence lens.

Validation should assess effectiveness, suitability, reliability, human factors, safety, cybersecurity and logistics within relevant conditions. It should identify open deficiencies and accepted limitations. A vendor-funded test can be informative while independence, scope and customer acceptance should be explicit.

The buyer should reproduce headline metrics and inspect negative results. It should map unresolved findings to programme milestones and valuation. Price should reflect the evidence state that exists at signing rather than the broadest expected future performance.

9. TEST MISSION REPRESENTATIVENESS

A test can be technically valid and commercially weak when its conditions do not represent the customer mission. Representativeness includes environment, operator, adversary, network, payload, duration and maintenance. Short scripted demonstrations can omit failure accumulation and recovery.

The target should document why each test supports a requirement. Whole-mission trials can complement component tests. Repeated runs and blind scenarios can reduce selection effects. Where security limits detail, authorised reviewers can assess sufficiency and record a non-sensitive conclusion.

Diligence should compare test populations with proposed exports and variants. A land system, maritime system and aerial system face different dynamics and controls. Performance should not be transferred across domains without evidence. The valuation should separate proven configurations from options that require new testing.

10. MEASURE RELIABILITY, AVAILABILITY AND MAINTAINABILITY

Mission success depends on more than task accuracy. Reliability, availability, maintainability, battery or fuel, spares, software support and operator training affect field value. A highly capable prototype can fail commercially when support requirements exceed customer capacity.

The target should report mission hours, failures, aborts, mean time between relevant failures, repair time, availability and support actions. Definitions should be consistent across customers. Scheduled maintenance and cannibalised units should remain visible.

The buyer should reconcile field logs with warranty, service and inventory records. It should evaluate the cost of maintaining deployed configurations and secure update channels. Support burden belongs in gross margin and backlog execution. A recurring sustainment business can add value when obligations, rights and cash are clear.

11. VALIDATE PERCEPTION, NAVIGATION AND DECISION MODELS

Autonomous systems can use perception, localisation, planning, control and decision-support models. Each component has different data, error and consequence. Aggregate mission success can conceal systematic weaknesses in rare conditions.

The target should document development data, labels, synthetic data, simulation, exclusions, leakage controls, performance and uncertainty. Whole-platform validation should complement component metrics. Geographic, sensor and environmental shifts should be tested.

Diligence should examine false positives, false negatives, calibration and threshold effects. It should identify which outputs inform humans and which initiate machine action. Model performance should be connected to mission and safety consequences. Transaction value should follow validated decision use rather than model novelty.

12. QUANTIFY UNCERTAINTY AND FAILURE MODES

Autonomous performance contains uncertainty from sensors, models, environment, communications and human interaction. Point estimates can hide tail conditions. The system should detect relevant uncertainty and transition to an approved safe or degraded state where designed.

The target should maintain hazard analyses, failure modes, fault trees, safety cases and incident records appropriate to its system. Common-cause and cascading failures should be considered. A redundant component does not create resilience when both channels share software, data or power.

The buyer should inspect whether failure evidence changed design, training or operating limits. It should test the completeness of anomaly closure. Accepted residual risk should have an accountable owner. Valuation should include the cost and schedule of unresolved safety and reliability work.

13. GOVERN HUMAN JUDGEMENT AND INTERVENTION

Human roles should be defined across planning, activation, supervision, intervention and review. DoD Directive 3000.09 requires appropriate levels of human judgement over the use of force for covered systems. NATO's responsible-use principles also emphasise responsibility, accountability and governability.

The target should document authority, interface, alerting, workload, training, override, abort and logging. A nominal human role can be ineffective when time, information or interface design prevents meaningful action. Human factors testing should reflect operational conditions.

The acquirer should inspect operator performance, automation bias and mode confusion. It should verify that system limits are communicated and that customer training matches configuration. Human oversight should not become an undocumented repair layer for weak autonomy.

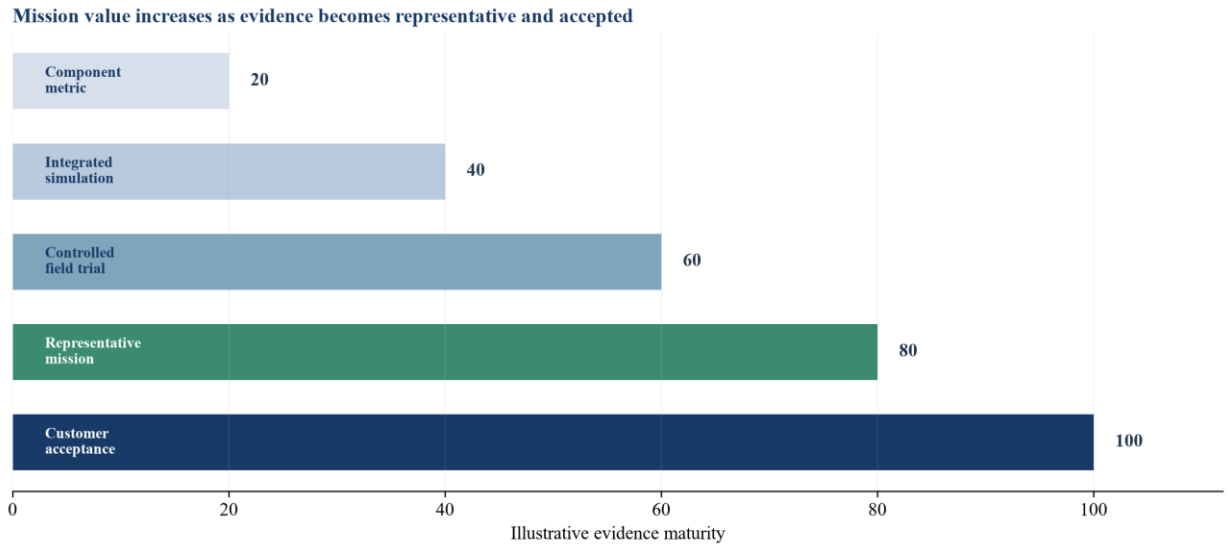

Figure 2. Autonomy evidence ladder from component result to mission acceptance

The ladder shows increasing decision relevance; higher levels require controlled configuration and customer evidence.

14. APPLY RESPONSIBLE AI GOVERNANCE

Responsible AI governance should be specific to system function and consequence. NATO's revised strategy identifies lawfulness, responsibility and accountability, explainability and traceability, reliability, governability and bias mitigation, and calls for testing, evaluation, verification and validation capability.

The vendor should document purpose, data, model, performance, limitations, monitoring, security, human roles and change. Foundation or generative models require separate controls for output reliability, data exposure and prompt-based manipulation. Marketing should identify actual use rather than using AI as a broad label.

The buyer should test whether governance evidence can satisfy customer review without exposing protected information. Reusable assurance can support procurement and renewal. Governance assembled separately for each bid creates cost and inconsistency.

15. TEST CYBERSECURITY AND ADVERSARIAL RESILIENCE

Autonomous systems can be affected by software vulnerabilities, compromised updates, spoofing, jamming, data poisoning, model extraction and supply-chain compromise. Security evidence should reflect the mission and architecture. A general corporate security certification does not prove platform resilience.

The target should maintain threat models, secure development, component provenance, access control, cryptographic management, update integrity, vulnerability handling and incident response. Tests should include authorised adversarial assessment within safe and lawful boundaries.

Diligence should inspect unresolved vulnerabilities, exceptions and customer obligations. It should identify foreign-source components and support dependencies. Cyber remediation, accreditation and secure facilities can materially affect schedule and cost. Transaction value should reflect the configuration that can be securely delivered and sustained.

16. CONTROL DATA, SOFTWARE AND TECHNICAL-DATA RIGHTS

Defence products can combine company intellectual property, government-funded development, third-party software, open-source code, supplier designs and customer data. Ownership and use rights can differ. Possession of source code does not establish unrestricted commercial or export rights.

The target should map invention, funding, contract clauses, licences, markings, restrictions and employee or contractor assignments. Training data and simulation environments require provenance. Open-source obligations should be evaluated within the distribution and security model.

The buyer should compare claimed proprietary assets with deliverables and customer rights. It should identify escrow, source-access and interface obligations. Technical-data restrictions can constrain integration, support and export. These limits belong in valuation and transaction covenants.

17. ESTABLISH PRODUCTION READINESS

Prototype success does not establish repeatable production. The buyer should evaluate bill of materials, supplier qualification, tooling, test equipment, quality, yield, lead time, obsolescence and capacity. Export-controlled components and single sources can constrain delivery.

The target should reconcile planned quantities with purchase commitments and manufacturing evidence. Configuration changes must flow through production and field support. First-article and acceptance failures should be visible.

Diligence should stress schedules for supplier delay, redesign and licence timing. It should inspect working capital, advance payments and inventory. Backlog value should be reduced for capacity and cash requirements. A credible industrialisation plan can support milestones while remaining distinct from proven output.

18. ASSESS INTEROPERABILITY AND INTEGRATION

Autonomous systems often depend on command, control, communications, intelligence, navigation and customer platforms. Interfaces can be proprietary, classified, regulated or customer-specific. Integration can control both mission performance and sales cycle.

The target should document standards, interface versions, test environments, accreditation and customer dependencies. NATO's AI strategy emphasises interoperability and quality data. Those aims support market relevance but do not prove a particular product's integration readiness.

The buyer should inspect completed integrations, reusable adapters and engineering effort. It should separate product capability from prime-contractor or government-furnished systems. Integration rights, schedule and margin belong in backlog valuation.

19. CLASSIFY ITEMS, SOFTWARE, TECHNOLOGY AND SERVICES

Export analysis begins with the specific item, software, technology, technical data, defence service and activity. In the United States, ITAR and EAR scope, classification, destination, end user, end use and licence exceptions or exemptions can matter. In the United Kingdom, military and dual-use lists, end-use controls and sanctions can apply.

The target should maintain written classifications with responsible counsel or specialists. Classification should be configuration-specific and updated after change. Intangible transfers, cloud access, training, source code and technical assistance require attention as well as physical shipment.

The acquirer should avoid treating a past licence as universal permission. Current official rules and transaction-specific advice are required. The valuation should identify markets that are presently supportable, markets requiring approval and markets excluded from the plan.

20. BUILD THE EXPORT-CONTROL MATRIX

The export matrix should connect product configuration, controlled content, jurisdiction, classification, destination, end user, end use, licence, proviso, technical assistance, retransfer, recordkeeping and expiry. It should identify U.S.-origin and UK-origin content within non-U.S. products where relevant.

The matrix should also cover employees, contractors, cloud regions, demonstrations, data rooms and post-close integration. M&A diligence itself can create controlled-information risk. Access should be designed before disclosure.

The buyer should test a sample of shipments and technical interactions against records. It should identify voluntary disclosures, audits, denials and unresolved questions. Export compliance is a live operational system rather than a static policy.

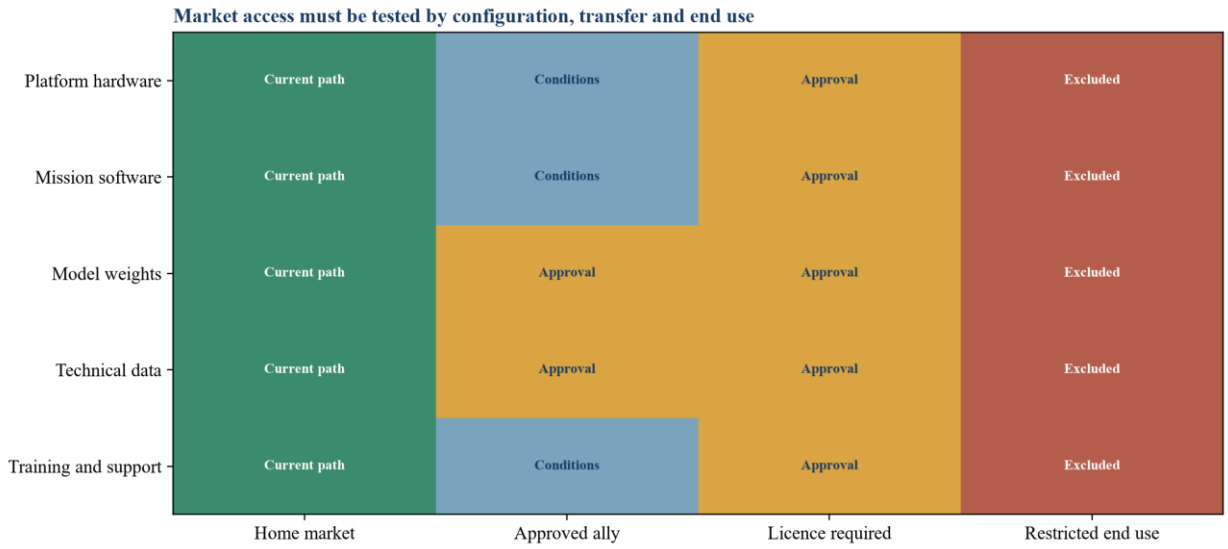

Figure 3. Illustrative export and market-access matrix

The matrix is a transaction-control example, not a legal classification or licence determination.

21. TEST LICENCES, PROVISOS AND RETRANSFER RIGHTS

A licence can be limited by quantity, value, destination, end user, end use, validity, provisos and reporting. Technical-assistance agreements and exemptions can carry separate conditions. Approval for one customer does not establish access to a broader market.

The target should reconcile licences with deliveries, services, employees and records. It should monitor expiry and changes. Retransfer and integration can require additional approval. Customer delays or configuration changes can consume licence validity.

The buyer should map each material programme to current authorisation and remaining conditions. It should include licence lead time and uncertainty in closing and backlog forecasts. Where approval is pending, contingent value or conditions precedent can align payment with outcome.

22. SCREEN END USERS, END USES, SANCTIONS AND DIVERSION

List classification is only part of export compliance. End-user, end-use, sanctions, embargoes and diversion risk can apply. Official U.S. BIS and UK ECJU guidance requires transaction-specific review. Distributors and integrators do not remove the exporter's responsibilities.

The target should maintain screening, beneficial ownership, end-use statements, red-flag escalation, shipping controls, training and records. It should review post-delivery support and retransfer. Commercial pressure should not bypass escalation.

Diligence should inspect the operation of controls and exceptions. It should identify high-risk intermediaries and unusual routing. A written policy without transaction evidence has limited value. Market access should be valued after compliance cost and restriction.

23. ADDRESS FOREIGN INVESTMENT AND NATIONAL SECURITY REVIEW

Defence and dual-use acquisitions can trigger foreign-investment, ownership, control, classified-contract and security reviews. Requirements depend on jurisdiction, investor, technology, contracts and governance. The transaction may need notification, mitigation or structural safeguards.

The target should maintain security arrangements, facility requirements, clearances, access controls and government correspondence. The buyer should obtain current specialist advice early. Deal structure, information rights, board composition and integration can be affected.

Timing and remedy risk belong in the transaction plan. A financing commitment should account for regulatory conditions. Synergy that depends on unrestricted data or personnel access should not be included before permissions are established.

24. ASSESS INTERNATIONAL HUMANITARIAN-LAW AND POLICY CONTEXT

International humanitarian law applies to the use of weapons, including systems with autonomous functions. The UN Convention on Certain Conventional Weapons process continues to consider emerging technologies in lethal autonomous weapon systems. The policy and normative environment remains active.

The target should identify the legal and policy review responsibilities applicable to its products and customers. It should avoid presenting a technology design as a complete legal determination. Records should support authorised customer assessment of intended use, limits and human roles.

The buyer should monitor policy change and contract requirements. It should evaluate whether product architecture supports governability, traceability and updates. Normative uncertainty can affect procurement, export and reputation, and should be reflected in scenarios rather than asserted as a forecast.

25. MEASURE CUSTOMER ACCEPTANCE AND FIELD VALUE

Customer value should be connected to accepted capability and mission outcome. Measures can include availability, task completion, operator workload, response time, coverage, sustainment or avoided exposure. Each measure requires a defined mission and comparison.

Attribution is difficult because training, tactics, environment and other systems affect performance. The target should avoid claiming that autonomy caused every favourable result. Controlled comparisons, repeated missions and customer acceptance provide stronger evidence than selected anecdotes.

The acquirer should build an outcome ledger linking configuration, test, customer, acceptance, delivery, service and cash. Classified outcome detail can remain under controlled review. Valuation should use conclusions that authorised reviewers can support without exposing protected information.

26. ANALYSE DELIVERY ECONOMICS

Autonomous-defence businesses can appear software-led while carrying hardware, integration, field support, secure facilities and bid costs. These activities can create defensibility and also change margin. The target should allocate labour, components, subcontractors, warranty and support by programme.

Development, production, sustainment and non-recurring engineering should be separated. Revenue recognition and acceptance milestones should align with contracts and IFRS 15 analysis. Cash conversion can lag reported revenue because of retention, milestone and inventory.

The buyer should compare contract margin with actual programme cost and open obligations. A high-growth backlog can consume cash. Delivery capability and working capital belong in valuation rather than post-close surprise.

27. TEST CONCENTRATION, RENEWAL AND PROGRAMME PERSISTENCE

Defence revenue can be concentrated by government, prime, programme, platform and budget. Several contracts can depend on one appropriation or integration partner. The buyer should aggregate exposure at the decision-maker and funding source.

Cohort analysis should cover bids, awards, options, deliveries, renewals, sustainment and collection. Programme duration does not guarantee annual order volume. Political, budget, export and performance change can affect continuation.

The target should show how fielded systems create recurring support, software or upgrade value. Customer lock-in should be distinguished from transferable contractual rights. Revenue persistence should be supported by accepted capability, funded need and support performance.

28. BUILD THE VALUATION BRIDGE

The hypothetical target has 24 customers, USD 31.5 million of annual revenue, USD 146 million of management-reported backlog, 38 deployed or trial systems and 11 described addressable countries. Management proposes USD 20.0 million of annual customer value and USD 13.2 million of buyer revenue and synergy. These figures are illustrative management assumptions only.

Evidence gates retain USD 62 million of funded, executable backlog, 16 systems with controlled mission evidence and four currently supportable export markets. Customer value falls to USD 7.6 million after performance, acceptance and persistence adjustments. Buyer revenue and synergy fall to USD 5.0 million after export, concentration, production, delivery-cost and cash-conversion adjustments.

The bridge prevents demonstrations, authorisations and addressable programmes from entering value without evidence. The buyer can then apply its selected valuation method, discount rate and structure to retained cash flows.

Table 2. Hypothetical mission-value ledger

Proposed value component	Management case	Evidence retained	Principal adjustment
Mission effectiveness	USD 8.0m	USD 3.4m	Representativeness and acceptance
Operator and response efficiency	USD 4.2m	USD 1.8m	Baseline and attribution
Availability and sustainment	USD 3.5m	USD 1.4m	Reliability and support cost
Risk and compliance	USD 4.3m	USD 1.0m	Scenario and legal dependency
Total annual customer value	USD 20.0m	USD 7.6m	Governed evidence gates

Every number is an illustrative management assumption, not observed company or programme data.

Table 3. Hypothetical buyer revenue and synergy bridge

Component	Management case	Evidence-adjusted case	Evidence gate
Revenue retained	7.5	3.5	Funding, acceptance, licence and cash
Cross-sell	2.8	0.7	Customer authority and market access
Production and delivery synergy	1.9	0.6	Capacity, quality and working capital
Data and platform synergy	1.0	0.2	Rights, security and integration
Total	13.2	5.0	No duplication with customer value

Values are illustrative management assumptions in USD millions.

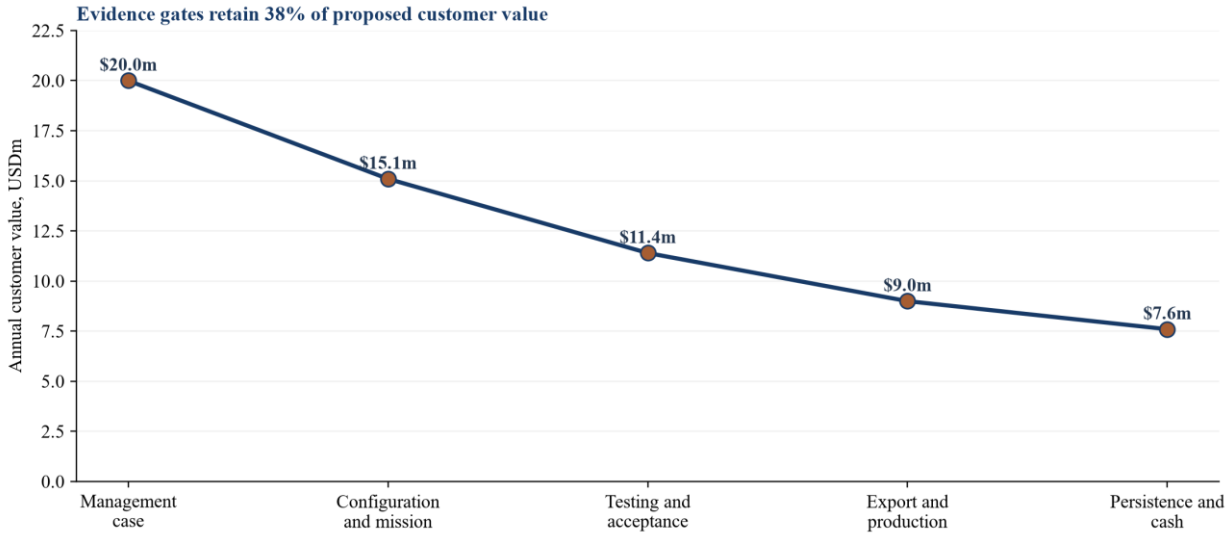

Figure 4. Hypothetical customer-value validation bridge

Values in USD millions are illustrative management assumptions.

29. TRANSLATE EVIDENCE INTO PRICE AND TERMS

Accepted, supportable capability can support value at signing. Unvalidated missions, pending licences, future variants and unfunded options can be funded through milestones. Milestones should reference objective outcomes such as independent test completion, customer acceptance, licence approval, delivery, margin or collected cash.

Representations should cover programme status, configuration, test evidence, intellectual property, technical-data rights, export compliance, sanctions, cybersecurity, security obligations and incidents. Holdbacks or indemnities can address defined exposures. The buyer should reserve integration capital for production, secure systems, accreditation and sustainment.

Earn-outs need precise definitions. Government funding and licence timing sit partly outside management control. Revenue targets should state treatment of options, pass-through hardware, delayed acceptance, claims and buyer-led cross-sell. Transaction structure should align payment with evidence becoming usable and durable.

30. EXECUTE DILIGENCE AND THE FIRST ONE HUNDRED DAYS

The diligence request should include programme and backlog registers, configuration baselines, test evidence, model inventory, data rights, production records, licences, screenings, incidents, customer acceptances, contracts, cohorts, service hours, security and financial models. Sampling should cover material programmes, recent versions, adverse tests, exports and high-risk suppliers.

The first thirty days should establish control, freeze inventories and reconcile deployments. Days thirty-one to sixty should reproduce material evidence, close priority compliance and validation gaps and confirm rights. Days sixty-one to one hundred should harmonise governance, monitoring, production, export controls and board reporting.

Management should report funded backlog, accepted systems, open deficiencies, model changes, export status, production yield, programme margin and collected cash separately. Unsupported mission claims should be removed from forecasts. Integration should preserve evidence and access controls before technical consolidation.

Table 4. Diligence gates and first-one-hundred-day ownership

Workstream	Pre-close evidence	Day-one control	Day-one-hundred outcome
Mission and model	Configuration and validation	Version freeze and owner	Governed capability estate
Programme	Funding, contract and acceptance	Backlog reconciliation	Executable delivery plan
Export and security	Classification, licences and access	Transfer and access controls	Controlled market-access register
Production	Suppliers, yield and capacity	Change and quality control	Validated output and sustainment
Commercial	Margin, working capital and cash	Concentration watch	Validated integration economics

Evidence owners should be named before close and accountable through integration.

31. SET BOARD AND INTEGRATION CONTROLS

Board reporting should distinguish capability evidence, programme execution, export status and commercial economics. Technical reporting can cover validation, anomalies, changes, human factors and security. Programme reporting can cover funding, acceptance, production and schedule. Commercial reporting can cover margin, concentration, working capital and cash.

One readiness score can hide the source of risk. Separate measures should reconcile through programme, configuration and licence registers. Directors should see which limitations are accepted and which remediation constrains delivery.

Integration should sequence identities, access, technical data, source code, secure systems, suppliers, licences, programmes and customer obligations. Synergy should begin after lawful access and validated migration. This sequencing preserves mission evidence and compliance.

32. RECOGNISE LIMITATIONS AND CONCLUDE

This framework does not validate a particular defence product, programme, export classification, licence or transaction. Applicable defence, export, sanctions, security, procurement, investment, accounting and humanitarian-law requirements vary by jurisdiction and fact. Current specialist advice and authorised diligence are required.

The hypothetical case does not estimate military effectiveness, demand, export approval or transaction value. Public policy and regulatory materials establish process context. They do not establish the status of a private company or programme. Protected evidence should remain within authorised systems and review channels.

The practical conclusion is that proof in the field becomes transaction value through a governed chain. The mission must be defined. The tested configuration must be reproducible. Evidence must be representative and independently reviewable. Human judgement, safety and security must be governed. Programmes must be funded and accepted. Market access must be current and lawful.

Autonomous-defence technology can create material operational value. Its consequences require disciplined evidence and accountability. The strongest acquisition case connects mission, configuration, test, customer acceptance, legal permission, production, support and cash without converting a demonstration or addressable programme into an observed fact.

REFERENCES

[1] U.S. Department of Defense. Directive 3000.09, Autonomy in Weapon Systems, January 2023. https://media.defense.gov/2023/jan/25/2003149928/-1/-1/0/dod-directive-3000.09-autonomy-in-weapon-systems.pdf

[2] U.S. Department of Defense. Responsible Artificial Intelligence Strategy and Implementation Pathway, June 2022. https://www.ai.mil/docs/RAI_Strategy_and_Implementation_Pathway_6-21-22.pdf

[3] NATO. Summary of NATO's Revised Artificial Intelligence Strategy, July 2024. https://www.nato.int/en/about-us/official-texts-and-resources/official-texts/2024/07/10/summary-of-natos-revised-artificial-intelligence-ai-strategy

[4] NATO. Alliance Digital Strategy, January 2026. https://www.nato.int/en/about-us/official-texts-and-resources/official-texts/2026/01/13/alliance-digital-strategy

[5] United Nations Office for Disarmament Affairs. CCW Group of Governmental Experts on Lethal Autonomous Weapons Systems, 2025. https://meetings.unoda.org/meeting/74853

[6] United Nations Office for Disarmament Affairs. Informal Consultations on Lethal Autonomous Weapons Systems, 2025. https://meetings.unoda.org/unoda-stu-meeting/LAWS-consultations-2024

[7] U.S. Government Accountability Office. Weapon Systems Testing: DOD Needs to Update Policies to Better Support Modernization Efforts, GAO-26-107009, December 2025. https://www.gao.gov/products/gao-26-107009

[8] U.S. Department of State, Directorate of Defense Trade Controls. International Traffic in Arms Regulations. https://www.pmddtc.state.gov/ddtc_public/ddtc_public?id=ddtc_kb_article_page&sys_id=24d528fddbfc930044f9ff621f961987

[9] Electronic Code of Federal Regulations. 22 CFR Part 121, United States Munitions List. https://www.ecfr.gov/current/title-22/chapter-I/subchapter-M/part-121

[10] U.S. Department of Commerce, Bureau of Industry and Security. Export Administration Regulations. https://www.bis.gov/ear

[11] U.S. Department of Commerce, Bureau of Industry and Security. EAR Part 734, Scope. https://www.bis.gov/regulations/ear/734

[12] U.S. Department of Commerce, Bureau of Industry and Security. EAR Part 744, End-use and End-user Controls. https://www.bis.gov/regulations/ear/744

[13] UK Government. UK Strategic Export Controls, updated April 2026. https://www.gov.uk/guidance/uk-strategic-export-controls

[14] UK Government. Export Controls: Military Goods, Software and Technology. https://www.gov.uk/guidance/export-controls-military-goods-software-and-technology

[15] UK Government. Military End-use Controls, updated May 2026. https://www.gov.uk/guidance/military-end-use-controls

[16] UK Government. National Security and Investment Act: Notifiable Acquisitions. https://www.gov.uk/government/publications/national-security-and-investment-act-guidance-on-notifiable-acquisitions

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

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

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

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

[21] 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/

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

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.