European Industrial AI Platforms Cross Plant Data without Customer Leakage

An Acquisition Framework for Interoperability Data Rights and Integration Value

Authored by Chennakeshav Adya

Independent Researcher

September 2026

Abstract

Industrial artificial intelligence platforms promise to transfer learning across factories while allowing each manufacturer to retain control of sensitive production data. Acquisition value depends on whether that promise survives technical, contractual and operational scrutiny. A model that performs in one plant may fail when equipment, product mix, maintenance policy, sensor calibration, operator practice or enterprise systems change. A data-sharing contract may permit service delivery while prohibiting model training, benchmarking or use after a change of control. A common interface may transport values while losing the semantics required to interpret them.

This paper develops an acquisition framework for European industrial-AI platforms that operate across plants, customers and equipment estates. It treats interoperability, data rights, leakage controls, model transfer, workflow adoption and continuing integration cost as linked sources of value. The framework separates access from rights, connectivity from semantic interoperability, pooled learning from customer-confidential information and reported software margin from the recurring engineering work needed to keep models useful in production.

The analysis draws on the European Union's Data Act, Data Governance Act, Artificial Intelligence Act, General Data Protection Regulation, Trade Secrets Directive, Cyber Resilience Act and NIS2 Directive; European data-space policy and guidance; manufacturing data-space initiatives; and industrial interoperability and AI-risk standards [1-50]. These sources define relevant duties and design principles. They do not establish the compliance, performance, economics or value of any target. A transaction requires current legal, competition, cyber, technical, commercial, accounting and tax diligence in each relevant jurisdiction.

An illustrative case demonstrates the method. The target group reports EUR 180 million of annual revenue and EUR 38 million of EBITDA. Normalising customer-specific engineering, connector maintenance, model monitoring, cyber operations, data-governance work and retention cost reduces sustainable EBITDA to EUR 24 million. Gross annual synergy claims of EUR 34 million become EUR 12 million of recurring net cash after adoption, rights, integration and assurance costs. An illustrative bridge applies sixteen times sustainable EBITDA, adds EUR 52 million of evidence-weighted synergy present value and deducts EUR 86 million for remediation, migration and customer-consent risk, producing EUR 350 million. Every amount is a management assumption used only to demonstrate the method.

The central conclusion is that buyers should value transferable industrial workflows and recurring cash supported by enforceable data rights, reproducible semantics and accepted plant operations. Deal terms should defer value where cross-customer learning lacks a clear legal basis, semantic mapping is customer-specific, integrations depend on scarce individuals, cybersecurity obligations are underfunded or model performance cannot be reproduced across the intended plant cohort.

JEL Classification: G24, G34, L23, L60, M15, O32, O33

Keywords: industrial AI, manufacturing data, data spaces, model interoperability, data rights, mergers and acquisitions, platform integration, Europe

Introduction

European manufacturers are deploying machine learning for quality inspection, predictive maintenance, process optimisation, energy management, scheduling and supply-chain resilience. Platform vendors seek to convert these deployments into repeatable products. Strategic buyers and private-capital investors may therefore see an opportunity to combine point solutions, industrial connectors, domain models and distribution into a broader platform.

The acquisition problem is more difficult than a conventional software roll-up. Industrial data reflects physical processes, local controls, equipment histories and customer operating methods. Production data can reveal throughput, yield, recipes, downtime, energy intensity, supplier quality and product design. Sharing can improve models and benchmarks. Uncontrolled sharing can expose confidential information, trade secrets or operational vulnerabilities.

The European legal and policy environment creates both an opportunity and a diligence perimeter. The Data Act establishes rules for access to and use of data from connected products and related services [1-3]. The Data Governance Act provides mechanisms intended to support trustworthy data sharing and data-intermediation services [4-5]. The Artificial Intelligence Act creates risk-based obligations for AI systems [6-8]. GDPR applies where personal data is present [9-10]. The Trade Secrets Directive protects qualifying confidential business information [11]. NIS2 and the Cyber Resilience Act address cybersecurity duties relevant to many industrial operators and digital products [12-14].

European data-space programmes aim to make data available under common governance and technical arrangements [15-20]. Manufacturing initiatives such as Catena-X, Manufacturing-X and sector data spaces provide useful architectures and operating lessons [21-27]. Industrial standards such as OPC UA, ISA-95, Asset Administration Shell specifications and ISO or IEC management standards provide technical reference points [28-39]. None removes the need to test the target's actual rights, mappings, customer acceptance and economics.

This paper is for boards, strategic buyers, private-equity investors, lenders and management teams evaluating industrial-AI combinations. It provides a transaction framework. It does not provide legal, regulatory, competition, accounting, tax, technical-safety or valuation advice.

1 Define the acquisition thesis by controlled industrial work

The thesis should begin with a production decision that improves after closing. Examples include scheduling a maintenance intervention, identifying a likely quality defect, adjusting an energy-intensive process, prioritising a production constraint or recommending a process parameter for authorised review. Each decision has a distinct economic owner, tolerance for error, time horizon and evidence requirement.

The target asset should be described as an operating system of rights, interfaces, models and accepted work. It may include edge connectors, semantic mappings, feature pipelines, model libraries, deployment tooling, monitoring, customer contracts, plant integrations, domain engineers and distribution relationships. The buyer's contribution may include a wider installed base, adjacent workflow products, stronger cyber operations, a data-space connector, capital for productisation or access to new industrial sectors.

Every value mechanism needs a baseline, owner, evidence, continuing cost, timing and failure condition. A claim that the combination will improve predictive maintenance should identify the equipment cohort, failure mode, observation window, existing maintenance policy, false-alert cost, avoided-downtime method and accountable plant manager. A claim that cross-plant learning will improve accuracy should

identify which data and parameters can lawfully move, how customer identity is protected, which features remain comparable and what validation is required before deployment.

The investment committee should distinguish four assets. The first is contractual permission to access and use data. The second is technical ability to connect and interpret it. The third is a model and workflow that performs under controlled conditions. The fourth is customer acceptance and recurring cash. Weakness in any asset can break the value chain.

Table 1 Industrial AI acquisition diligence perimeter

Value claim	Required evidence	Decision question	Principal risk
reusable plant connectivity	connector inventory test records version support and maintenance effort	can the platform connect without repeated bespoke engineering	one-off integration is presented as a product
semantic interoperability	information models mappings units lineage and conformance tests	do values retain meaning across equipment plants and customers	transport works while context is lost
transferable data rights	contracts lawful basis purpose limits trade-secret controls and change-of-control terms	can the buyer continue each material use	access exists without training benchmarking or transfer rights
cross-plant model value	cohort definitions external validation drift tests and acceptance records	does performance transfer to the intended plant cohort	local correlation is presented as general intelligence
sustainable earnings	full connector support data engineering cyber monitoring and retention cost	what recurring cash remains under proper control	essential engineering is capitalised or omitted
integration value	product overlap migration plan customer approvals and net cash model	which synergies survive rights and adoption gates	forced pooling causes leakage or customer loss

Proposed structure; target-specific legal competition technical cyber commercial accounting and tax review is required.

2 Map the data rights chain before valuing the model

An industrial platform can receive data through equipment interfaces, customer systems, sensors, service agreements, partner feeds or data spaces. Physical access does not determine the legal right to train a model, create a benchmark, combine records across customers or retain derived information after termination. The diligence team should build a rights ledger that follows data from origin to every material use.

The ledger should identify the data holder, user, controller or processor where relevant, equipment manufacturer, service provider, data-space participant and platform operator. It should record purpose, territory, duration, permitted recipients, confidentiality, trade-secret treatment, security, audit rights, deletion, derived-data treatment, model-output rights and change-of-control consequences. It should also identify statutory rights that may qualify or override contract terms.

The Data Act applies to defined data generated by connected products and related services and addresses access, use and sharing [1-3]. Its effect depends on the product, data, user and contractual relationship. The Data Governance Act regulates specified data-intermediation activities and establishes governance mechanisms [4-5]. GDPR remains relevant when production data contains information about employees, contractors or identifiable individuals [9-10]. Trade-secret protection depends on the information, its secrecy, commercial value and reasonable protective steps [11].

The buyer should test whether a customer contract authorises service delivery only, improvement of that customer's instance, pooled learning, aggregated benchmarking or broader product development. These are different rights. A clause allowing use of anonymised data needs technical evidence that the process prevents identification or disclosure of sensitive plant characteristics. Industrial data can remain commercially revealing even when personal identifiers are absent.

Figure 1 Industrial data rights and value chain

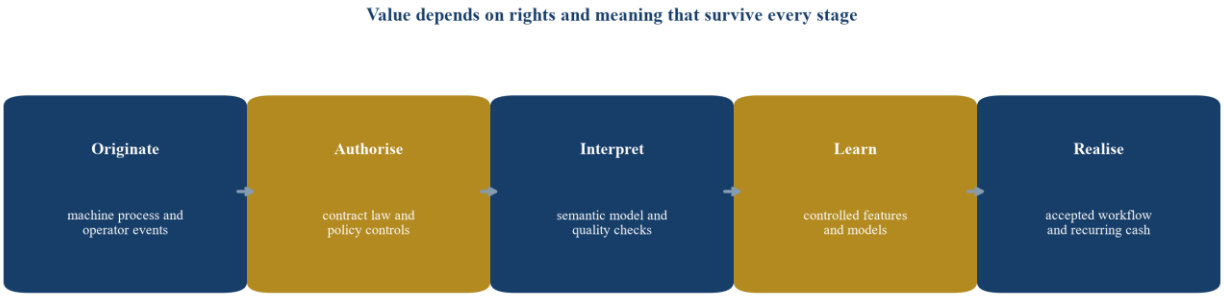

Proposed acquisition map; actual rights depend on applicable law contracts data provenance and technical controls.

3 Test interoperability at five distinct layers

Interoperability is frequently described as one capability. Acquisition diligence should divide it into transport, syntax, semantics, identity and workflow. A connector can move data while leaving units, timestamps, asset hierarchy or operating state ambiguous. A shared schema can still fail when plants use different naming conventions, maintenance codes or product definitions. A technically correct prediction can remain unusable when it does not fit authorised work.

Transport interoperability concerns protocols, network paths and reliable exchange. Syntactic interoperability concerns format and structure. Semantic interoperability concerns the meaning of assets, variables, events and relationships. Identity interoperability concerns consistent representation of organisations, sites, equipment, users and credentials. Workflow interoperability concerns how an output enters an accountable production decision.

OPC UA provides a platform-independent architecture for information exchange and includes information, message, communication and conformance models [28-31]. Companion specifications can add domain information models [31]. ISA-95 provides concepts for integration between enterprise and control activities [32]. Asset Administration Shell specifications provide another approach to representing industrial assets and submodels [33-35]. Data-space blueprints add identity, catalogue, policy, trust and exchange capabilities [18-20].

The buyer should run conformance tests on a representative plant cohort. The test should include legacy controls, modern equipment, multiple vendors, common historians, manufacturing execution systems and enterprise systems. It should measure time to connect, mapping exceptions, missing context, manual intervention, update resilience and continuing support cost. Sales demonstrations and connector counts do not substitute for this evidence.

The test should also examine time. Industrial signals can be sampled, aggregated, delayed or corrected at different stages. A vibration value recorded by a controller may reach a historian under a different

timestamp, time zone or clock quality. A quality result may arrive after the production event that created it. Models trained on aligned historical data can fail when production pipelines introduce delay or reordering. Diligence should therefore reproduce event-time handling, late-arriving data, daylight-saving changes, clock drift and the recovery of missed transmissions.

Units and physical context deserve equal attention. Pressure, temperature, energy and flow can be represented through different units, scales and reference conditions. Equipment tags can be reused across lines or changed during maintenance. A semantic mapping should preserve the source, transformation, unit, calibration, asset, operating mode and quality flag. Automated conversion should be verified against engineering documentation and sampled plant records. A platform that infers these relationships manually may remain valuable, although the labour and error rate should be incorporated in sustainable economics.

Version resilience is a separate product test. Equipment firmware, control logic, historian schemas, enterprise applications and customer policies change. The target should show how it detects change, tests compatibility, communicates impact and restores service. The acquisition model should include the supported version matrix and the cost of keeping it current. Unsupported legacy estates can create attractive initial revenue while accumulating a growing service obligation.

Table 2 Five-layer interoperability record

Layer	Evidence	Failure that affects value
transport	protocol support latency reliability security and version tests	data cannot move consistently or safely
syntax	schemas types units timestamps and validation results	values arrive in incompatible structures
semantics	asset models taxonomies relationships lineage and mappings	the same field means different things across plants
identity	participant site asset user and credential resolution	data or permissions attach to the wrong entity
workflow	role decision interface override escalation and record	prediction does not become accepted accountable work

Proposed record; each material workflow should be tested on representative plant configurations.

4 Build plant cohorts before claiming model transfer

Cross-plant learning is valuable when performance transfers across a defined cohort. The cohort should reflect the variables that affect the physical process and the model. These may include equipment family, age, controller, sensor package, calibration, product mix, materials, climate, maintenance policy, shift pattern, operator practice and failure prevalence.

The platform should maintain a cohort register for every material model. The register should identify training plants, validation plants, excluded plants, operating envelope, target variable, reference standard, performance thresholds and known limitations. It should also record data quality, missingness, sampling frequency, label construction and changes in process or equipment.

Transfer can be tested through external validation, leave-one-plant-out analysis, prospective shadow deployment and controlled production release. Aggregate performance can conceal weak results at a particular customer, site or operating state. The diligence team should examine distributions, calibration, false alarms, missed events, economic consequences and subgroup stability.

Federated learning, parameter sharing, synthetic data and secure computation can reduce movement of raw records. Each approach introduces assumptions and residual risks. Model updates can encode information about a participant. Synthetic data can reproduce sensitive patterns or fail to represent rare

events. Secure computation can reduce exposure while increasing cost, latency and operational complexity. The target should provide threat models, privacy or confidentiality tests and evidence that controls match the promised use.

Figure 2 Plant cohort transferability matrix

Similarity must be demonstrated by dimension

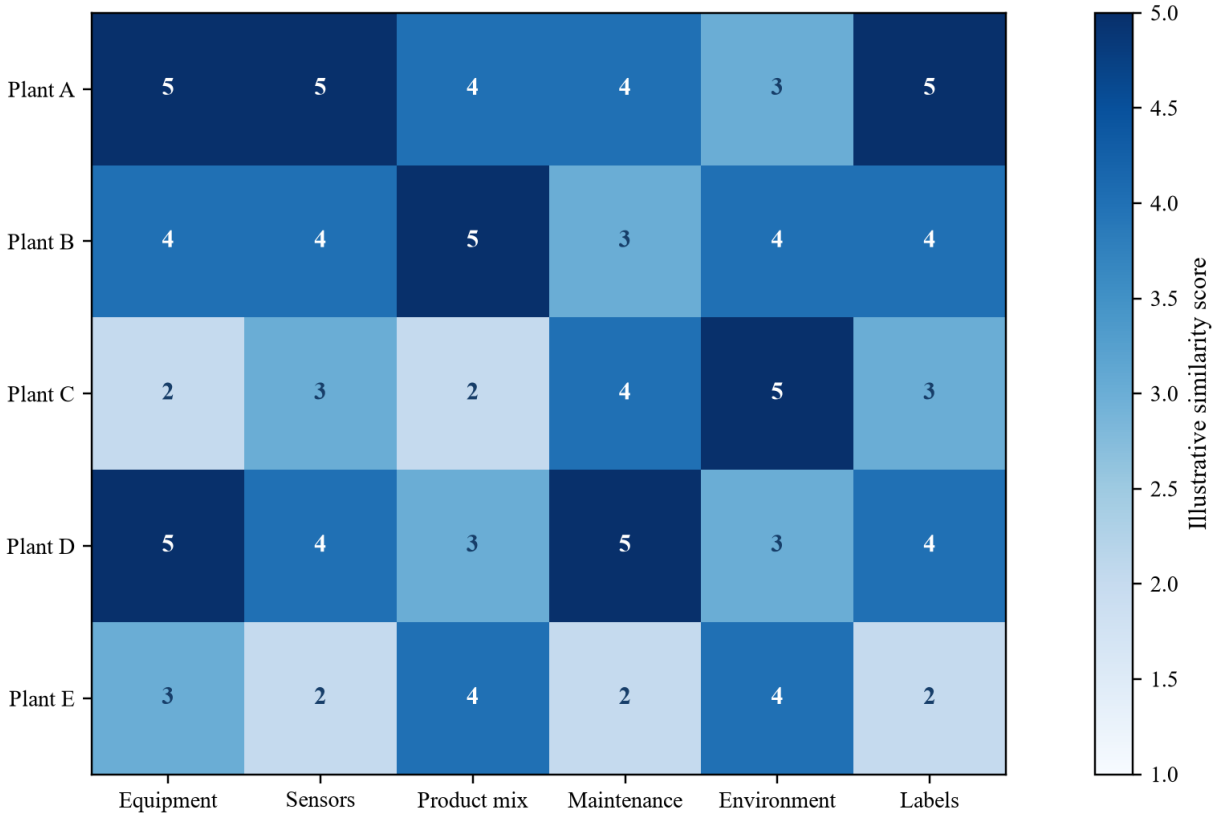

Illustrative scoring framework; scores are management assumptions used only to demonstrate the method.

5 Separate customer isolation from platform learning

Customer isolation should be designed across storage, computation, features, models, operations and people. Logical tenancy alone may be insufficient when engineers can query pooled logs, support staff can export records or model pipelines combine customer features without enforceable policy.

The target should document data-flow diagrams, tenancy boundaries, encryption, key ownership, access-control models, privileged access, logging, retention, deletion, backup, incident response and subcontractor access. It should show how policies are enforced in training, evaluation, deployment and support. The buyer should test the controls through configuration review, sampled logs, penetration testing and scenario exercises.

Model artefacts require separate treatment. Base code may be common. Customer-specific parameters, feature definitions, fine-tuned weights or thresholds may embody confidential information. A global model may be permissible only when contribution rights and leakage controls are established. Benchmarking may reveal relative productivity or quality even when names are removed.

The strongest architecture aligns technical boundaries with the rights ledger. A policy engine should determine which data can be discovered, accessed, transformed, trained upon, combined, retained and exported. Audit records should connect the decision to the controlling contract and policy version. Data-space blueprints describe machine-readable policy negotiation and enforcement as relevant capabilities [18-20].

The diligence team should run a controlled leakage exercise. The exercise can use synthetic plant records and should attempt cross-tenant retrieval, unauthorised feature reuse, model extraction, membership inference, benchmark re-identification, privileged support access and recovery of deleted data from backups. The purpose is to test the target's stated control design without exposing real customer information. Findings should be classified by exploitability, commercial sensitivity, operational consequence and time to remediate.

Customer communications form part of the control environment. Product documentation should explain what information is collected, where it is processed, which uses are performed, how shared models are constructed and what choices the customer retains. Contract language, architecture and sales statements should agree. A promise of complete isolation is inconsistent with undisclosed pooled learning. A promise of collective improvement requires precise boundaries so customers understand what is shared and what remains private.

The buyer should also examine internal incentives. Data science teams may be rewarded for model improvement, sales teams for rapid deployment and support teams for restoring service. Each incentive can encourage shortcuts around rights, lineage or separation. Approval workflows, training, escalation and performance measures should make compliance and customer protection part of operating performance rather than a late legal review.

Table 3 Customer leakage control matrix

Control domain	Minimum evidence	Acquisition test
tenancy and storage	architecture keys access lists backups and deletion	attempt unauthorised cross-tenant access in a controlled test
feature and training pipelines	lineage policy tags job configuration and approvals	trace each training input to a permitted purpose
model artefacts	ownership contribution records leakage tests and export policy	test whether outputs reveal customer-specific information
operations and support	privileged access workflow logs monitoring and review	sample emergency and routine access events
benchmarking	aggregation thresholds disclosure policy and customer terms	reconstruct whether a participant can be identified
exit and change of control	deletion return portability and transition evidence	determine what survives termination or acquisition

Proposed control matrix; assurance depth should match data sensitivity and operational consequence.

6 Determine the applicable European regulatory perimeter

An industrial-AI platform may sit across several legal regimes. The Artificial Intelligence Act applies according to defined roles and risk categories. Some industrial systems may be components of regulated products or safety functions. Other systems may fall outside high-risk categories while remaining subject to transparency, contractual, product-safety or sector duties. The buyer should map each system, intended purpose, deployment context and economic-operator role [6-8].

The Data Act applies from 12 September 2025 and addresses connected-product and related-service data, among other matters [1-3]. Diligence should identify whether the target is a data holder, data recipient, third party, cloud provider or other relevant actor for each service. The analysis should cover access mechanisms, user requests, trade-secret safeguards, compensation, contract fairness and switching obligations where relevant.

The Data Governance Act addresses reuse of protected public-sector data, data-intermediation services and data altruism, and created the European Data Innovation Board [4-5]. A target describing itself as a

neutral intermediary should be tested against the statutory definition, notification status, structural separation and conduct duties.

GDPR applies when industrial datasets include personal data, which can arise through user accounts, access logs, wearable devices, video, location or performance records [9-10]. NIS2 imposes cybersecurity risk-management and incident-reporting duties on covered entities and affects supply-chain expectations [12-13]. The Cyber Resilience Act creates requirements for products with digital elements, including vulnerability handling, with staged application dates [14]. Sector rules, machinery and product-safety law, competition law, export controls, employment law and national rules may also apply.

Table 4 European regulatory diligence map

Regime	Diligence focus	Transaction implication
Data Act	connected-product data access use sharing contracts and switching	rights architecture and customer terms may need remediation
Data Governance Act	intermediation role neutrality notification and governance	platform structure may affect permitted business model
Artificial Intelligence Act	system classification provider deployer obligations and controls	product roadmap and compliance cost affect value
GDPR	personal-data inventory lawful basis rights security and transfers	industrial telemetry can contain employee information
Trade Secrets Directive	secrecy commercial value and reasonable protection	pooled learning must preserve protected information
NIS2 and Cyber Resilience Act	risk management reporting product security and vulnerability handling	recurring cyber operations and liability require funding

High-level map; current counsel should determine applicability to each product entity role and jurisdiction.

7 Reconstruct sustainable earnings after interoperability cost

Industrial software revenue can contain recurring licences, usage charges, integration, engineering, support, hardware resale, cloud services and outcome-based fees. The diligence team should separate each stream by customer, plant, product, contract and delivery model. Recurring revenue should require recurring entitlement and evidence of renewal, not repeated statements of work.

Reported gross margin may omit engineering that is necessary to deploy and maintain the product. Connector development, semantic mapping, model retraining, data-quality correction, site commissioning, cybersecurity assurance and customer-success work can be classified across cost of sales, research and development, capitalised software or professional services. Sustainable earnings should include the continuing cost required to preserve contracted performance and compliance.

The buyer should build cohorts by implementation vintage, plant type, equipment family and customer segment. It should measure time to first accepted use, engineering hours, external cost, cloud or edge cost, support tickets, model interventions, renewal, expansion, contraction and cash collection. It should identify whether improvement comes from product reuse or from assigning more expert labour.

Implementation economics should be measured from signed order to stable operation. The record should capture discovery, security review, data access, connector installation, semantic mapping, model configuration, shadow operation, user acceptance, training and production release. Delays should be attributed to the vendor, customer, equipment supplier or another dependency. This allows the buyer to distinguish a product limitation from a customer-readiness problem and to estimate the capacity required for growth.

Retention analysis should connect the contractual renewal to operational use. A customer can renew because switching is difficult while active users or accepted decisions decline. Another customer can increase use while paying under a fixed enterprise licence. The buyer should examine price, volume, plant count, workflow count, gross margin, support burden and collected cash together. Expansion that requires new bespoke engineering should be evaluated as another implementation cohort.

Capitalisation policy can materially affect the view of earnings. Connector and model-development labour may qualify for accounting treatment under defined conditions. The transaction analysis should still identify which expenditure is necessary to support current revenue, remediate technical debt or create future capability. Cash requirements and engineering capacity remain relevant regardless of presentation in the income statement.

Illustrative normalisation begins with reported EBITDA of EUR 38 million. It adds EUR 3 million of genuinely non-recurring cost and deducts EUR 17 million for recurring integration labour, connector maintenance, model monitoring and controls that are necessary but incompletely reflected. Sustainable EBITDA is therefore EUR 24 million. These figures demonstrate the method and do not describe a company or market.

Figure 3 Illustrative EBITDA normalisation bridge

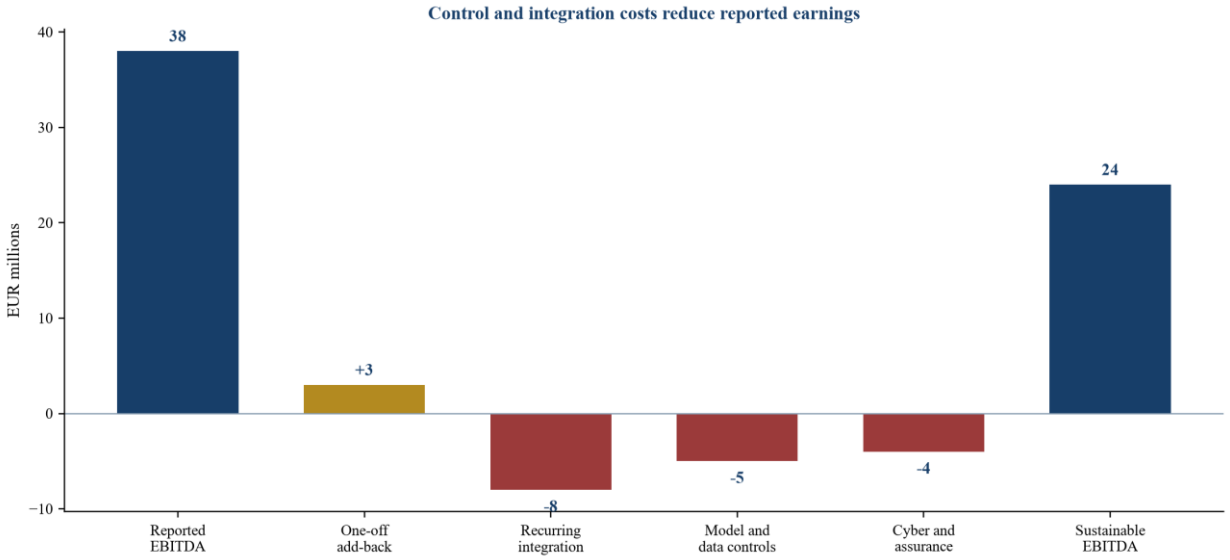

EUR millions; management assumptions used only to demonstrate the method.

8 Convert synergy claims into evidence weighted cash

Industrial-AI synergies often combine cross-selling, pooled learning, shared connectors, lower cloud cost and consolidated support. Each mechanism should be valued separately. Revenue synergy requires a named customer cohort, use case, sales owner, product fit, implementation capacity, price, probability, timing and collection assumption. Cost synergy requires an identified resource or contract that can be removed without weakening service or control.

Pooled learning can create value through better performance, faster deployment or broader coverage. The rights ledger and cohort analysis should establish whether pooling is permitted and technically valid. The model should then connect the improvement to an operational decision and cash effect. A small accuracy gain has limited value if it does not change maintenance policy, scrap, throughput, energy use or labour.

Shared connectors can reduce duplicated engineering when products use compatible architecture and customers accept migration. The buyer should compare code, protocol versions, semantic models, testing, security and support obligations. A connector count is a weak proxy because two connectors carrying the same label can differ materially in equipment coverage and reliability.

The illustrative case begins with EUR 34 million of gross annual synergy claims. Removing unsupported pipeline, rights constraints and overlapping claims reduces the amount to EUR 20 million. Implementation, continuing assurance and customer-adoption costs reduce recurring net cash to EUR 12 million. These are management assumptions for method demonstration.

Figure 4 Illustrative synergy evidence funnel

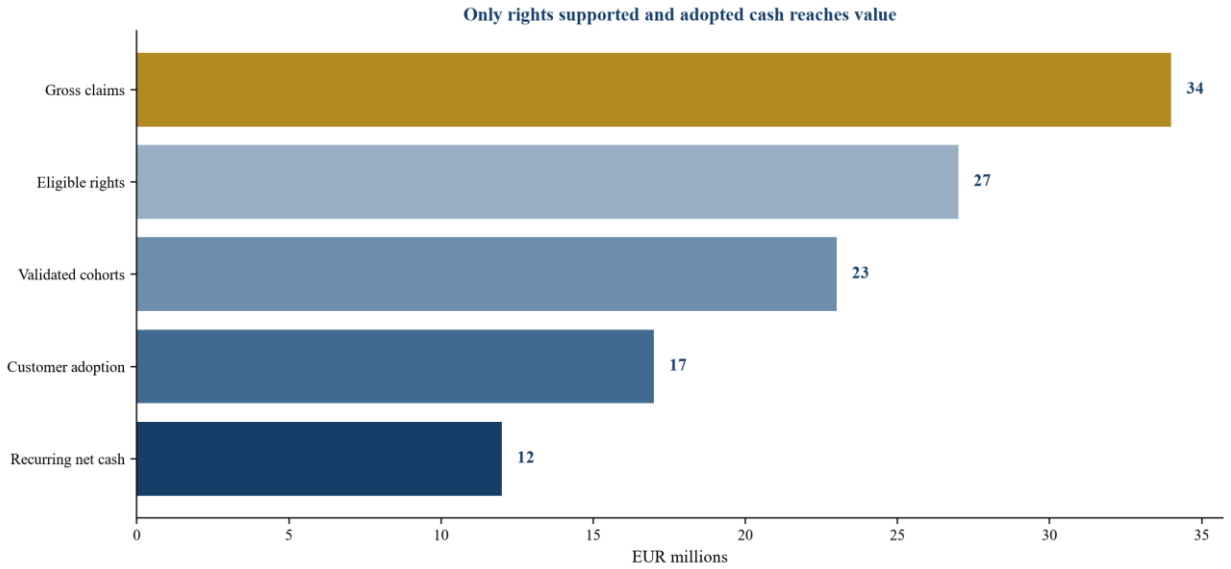

EUR millions; management assumptions used only to demonstrate the method.

9 Build a valuation bridge that prices unresolved obligations

The valuation should begin with sustainable earnings and separately recognise evidence-weighted synergy, remediation, integration and customer-consent risk. This makes assumptions visible and reduces double counting. A revenue multiple can supplement the analysis for high-growth assets, although revenue quality still depends on rights, implementation burden, retention and cash conversion.

The illustrative bridge applies sixteen times EUR 24 million of sustainable EBITDA, giving EUR 384 million. It adds EUR 52 million for the present value of evidence-weighted synergy. It deducts EUR 28 million for product and connector remediation, EUR 22 million for cyber and data-governance work, EUR 18 million for migration and EUR 18 million for customer-consent and retention risk. The result is EUR 350 million. All figures are management assumptions used only to illustrate the framework.

Sensitivity should cover sustainable EBITDA, multiple, synergy conversion, remediation cost, implementation duration and customer attrition. It should also test a separation case in which customer models and datasets remain isolated for longer than expected. The board should understand how much value relies on rights that require consent, technical migration or regulatory interpretation.

Consideration can be structured around evidence gates. Deferred consideration may depend on specified customer renewals, completion of validated connector migration, lawful continuation of pooled learning, performance across defined plant cohorts or achievement of collected cash. Escrow, indemnities, price adjustment and specific covenants may address identified obligations. Current legal and tax advice is required.

Table 5 Illustrative acquisition value bridge

Component	Assumption	Value
sustainable EBITDA	24	24
base multiple	16 times	384

Component	Assumption	Value
evidence weighted synergy present value	management case	52
product and connector remediation	management case	minus 28
cyber and data governance	management case	minus 22
migration and separation	management case	minus 18
customer consent and retention risk	management case	minus 18
illustrative equity value before debt and cash	calculated	350

EUR millions except the multiple; management assumptions used only to demonstrate the method.

10 Design the transaction around rights and customer trust

Data rights and customer trust should shape transaction documents. Representations may cover authority to collect and use data, accuracy of rights schedules, compliance with purpose limits, protection of trade secrets, personal-data processing, model-training practices, cybersecurity incidents, open-source software, export controls and material customer consents.

The disclosure process should include a data-and-model schedule. It should identify each material dataset, source, rights, purpose, territory, retention, derived artefact, model and customer restriction. It should connect to the product and contract schedules. Exceptions should be resolved before closing or priced explicitly.

Change-of-control analysis is essential. A customer may permit the current supplier to process data while restricting assignment, subcontracting, new ownership or use by an affiliate. A data-space rulebook may impose participant requirements. Equipment-vendor licences may limit extraction or reuse. The buyer should determine required notices, consents, re-papering and technical separation.

Pre-closing integration planning should respect competition and confidentiality constraints. Clean teams, restricted data rooms and defined information protocols may be required. Customer data should not be pooled merely because the transaction has signed or closed. The applicable contract, law, approved architecture and customer communication should control each migration.

Table 6 Transaction protection linked to evidence

Risk	Evidence before signing	Possible protection
insufficient training rights	contract and rights-ledger review	condition remediation indemnity or excluded asset
change-of-control restriction	consent schedule and customer analysis	closing condition covenant or retention mechanism
non-transferable model performance	cohort validation and production test	earn-out tied to defined performance and cash
connector remediation	code architecture test inventory and cost plan	price adjustment escrow or completion covenant
undisclosed cyber exposure	incidents vulnerabilities testing and insurance	remediation condition indemnity and operating covenant
customer leakage	architecture logs policy and controlled testing	separation plan escrow and specific warranty

Illustrative structure; transaction counsel should tailor provisions to the identified risks.

11 Integrate in layers and preserve reversibility

Integration should proceed by customer, product and workflow cohort. Corporate systems and common administration can move early when lawful and operationally appropriate. Customer data, feature pipelines, models, connectors and plant workflows should move through evidence gates.

The first phase should establish inventories, access controls, incident coordination and operating baselines. The second should standardise identity, lineage, catalogues and policy representation. The third should migrate connectors and semantic mappings through conformance tests. The fourth should combine models only where rights, cohort transfer and leakage controls are established. Commercial cross-selling should follow product readiness and customer acceptance.

Each migration should have a rollback path. Plant operations can be safety-critical or economically consequential. A new model or connector should enter shadow mode, controlled acceptance and monitored production according to the workflow's risk. The target's former configuration, decision history and evidence should remain retrievable for audit and incident analysis.

Separation can be a deliberate integration strategy. Products may share corporate services, identity, monitoring and procurement while customer datasets and model pipelines remain distinct. This can preserve contractual boundaries and reduce customer anxiety while the buyer develops common standards. The integration case should compare the cost and value of full convergence, controlled federation and continuing separation rather than assuming that one architecture is optimal for every product.

Product rationalisation should be based on workflow evidence. Two products may appear to address predictive maintenance while serving different equipment, users or decision intervals. Removing one can destroy a channel or specialised dataset. The buyer should compare customer cohorts, functional scope, model evidence, connector coverage, unit economics and roadmap obligations before selecting a survivor. Customer migration should include feature parity, performance validation, commercial terms and support capacity.

People integration is especially important because tacit knowledge can sit with domain engineers who understand a customer's equipment, tags and operating history. The buyer should identify key-person dependencies and convert them into maintained mappings, runbooks, test suites and training. Retention arrangements can protect continuity, although long-term value requires institutional knowledge and productised controls.

Value recognition should follow accepted operation. Connector savings should be recognised after supported interfaces are retired. Infrastructure savings should follow capacity and resilience testing. Cross-sell value should follow deployment, acceptance, invoicing and collection. Pooled-learning value should follow permitted contribution, external validation and controlled production performance.

Figure 5 Gated integration sequence

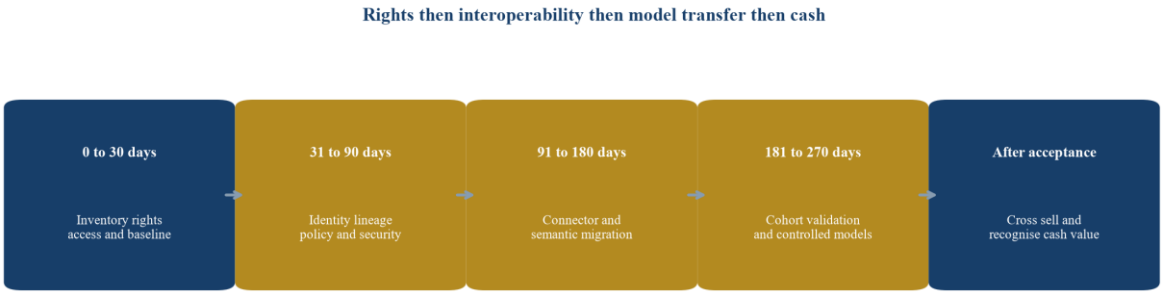

Proposed sequence; timing depends on customer contracts plant risk product architecture and regulatory duties.

12 Establish board controls and decision gates

The board should require a single acquisition register connecting each value claim to data rights, technical evidence, accountable owner, cost, timing and failure condition. The register should remain live through diligence, signing, closing and integration.

Key indicators should cover rights completeness, connector reuse, semantic exceptions, model transfer, customer isolation, production acceptance, incidents, retention and cash. Metrics should be defined consistently. Connector reuse might mean a supported connector deployed without code change. A model transfer pass might require pre-agreed performance and calibration across an external plant cohort. Customer isolation might require zero unauthorised cross-tenant access events plus completion of sampled control tests.

The board should maintain suspension criteria. A model should pause when performance crosses a defined threshold, input lineage breaks, a material permission expires or a leakage event occurs. Integration should pause when customer consent is missing, rollback cannot be demonstrated or plant risk exceeds the approved envelope. These criteria turn governance into operational control.

Decision gates should include thesis approval, rights completeness, technical reproducibility, sustainable earnings, transaction protection, integration readiness and value release. Each gate should identify the approving authority and evidence pack. Management optimism should remain visible as a scenario rather than entering the verified base case.

Reporting should preserve the distinction between leading evidence and realised value. Completed rights reviews, connector tests and shadow deployments can show readiness. Accepted production use, renewal, invoicing and collection show commercial realisation. A dashboard that combines these stages into one percentage can hide the point at which a value claim is failing. The board should review the conversion between stages and the time spent at each gate.

Independent assurance can be targeted to the most consequential claims. Legal review can address rights and regulatory roles. Industrial cybersecurity specialists can test architecture and incident readiness. Domain engineers can challenge semantics and operating envelopes. Data scientists can reproduce validation and leakage testing. Finance teams can reconcile contracts, invoices, cost allocations and cash. The work should converge into one decision record rather than remain as disconnected diligence reports.

Post-closing governance should continue until material value claims have either converted to recurring cash or been retired. The board should require owners to update assumptions when customer consent is delayed, a plant cohort fails validation, integration cost rises or retention changes. This discipline avoids preserving transaction assumptions after operational evidence has contradicted them.

The final decision record should state the evidence reviewed, unresolved matters, approved assumptions, accountable executives, funding requirement and next review date. It should remain available to the integration team and support later comparison between the acquisition case and realised outcomes.

Table 7 Board acquisition and integration dashboard

Dimension	Example measure	Evidence source	Board use
data rights	material uses with verified transferable rights	rights ledger contracts and counsel review	approve use and price unresolved risk
interoperability	supported deployments without bespoke code	test records engineering time and releases	validate product reuse
model transfer	plant cohorts meeting defined performance	validation reports telemetry and acceptance	control deployment and value claims
customer isolation	policy coverage exceptions and incidents	configuration logs tests and incident register	protect confidentiality and trust
sustainable economics	recurring gross margin and cash after full control cost	contracts invoices payroll systems and bank evidence	set base valuation
integration	migrations passed with rollback and customer acceptance	gate packs change records and customer sign-off	release next phase
synergy	recurring net cash realised against claim	general ledger invoices and bank records	release deferred consideration or revise plan

Proposed dashboard; thresholds should be defined for the target portfolio and plant risk.

Conclusion

European industrial-AI consolidation can create value by combining accepted workflows, interoperable industrial data and wider distribution. The value is realised when rights, semantics, model transfer and customer trust survive the combination.

Acquirers should therefore treat the platform as a controlled chain. Data must originate lawfully, remain interpretable, enter a validated model, support an accountable plant decision and convert into recurring cash. Customer isolation and cross-plant learning should be designed together. Interoperability should be tested across transport, syntax, semantics, identity and workflow. Sustainable earnings should include the engineers, controls and assurance required to keep the chain operating.

The framework also changes transaction design. Unresolved rights, customer consents, connector remediation and model transfer should be priced explicitly. Consideration can follow evidence gates. Integration should preserve separation and rollback until each migration is accepted. This approach gives boards a clearer basis for deciding what they are buying, what must remain isolated and when claimed platform value has become observable cash.

Frequently asked questions

Q: What is the most important diligence question for an industrial AI platform acquisition? A: Determine whether the target controls a repeatable production workflow supported by transferable data rights, semantic interoperability, reproducible model performance, customer acceptance and recurring cash. A model or connector alone establishes only part of the value chain.

Q: Does access to customer plant data permit pooled model training? A: Access does not automatically establish a right to train pooled models. The buyer should verify contract terms, applicable law, purpose limits, trade-secret controls, personal-data obligations, derived-data provisions and change-of-control consequences for each material use.

Q: How should a buyer test model transfer across plants? A: Define plant cohorts using equipment, sensors, product mix, maintenance, environment and label quality. Use external validation, leave-one-plant-out analysis, shadow deployment and controlled production acceptance. Review distributions and economic consequences as well as aggregate accuracy.

Q: What does interoperability mean in acquisition diligence? A: Test transport, syntax, semantics, identity and workflow separately. Data can move successfully while losing units, asset context or decision ownership. Value depends on the complete path from equipment to accepted work.

Q: How can an acquirer reduce customer-data leakage risk? A: Align contractual rights with tenancy, encryption, access control, feature and training policies, model-artifact controls, privileged-access monitoring, benchmark disclosure rules and deletion. Test those controls with traceable scenarios and sampled logs.

Q: Why can reported industrial software margins overstate sustainable earnings? A: Necessary connector work, semantic mapping, model monitoring, cyber operations, data governance and customer-specific support may sit outside cost of sales or be capitalised. Sustainable earnings should include the recurring cost required to preserve performance and compliance.

Q: How should industrial AI synergies be valued? A: Value each mechanism separately and require evidence of rights, technical feasibility, customer adoption, continuing cost and collected cash. Remove unsupported pipeline, overlapping claims and costs required to deliver or control the synergy.

Q: Which Matchpoint practice is relevant to this research? A: Matchpoint Partners' Mergers and Acquisitions practice supports strategic buyers, investors and management teams evaluating platform combinations, diligence priorities, transaction structure and integration value.

References

[1] European Union. Regulation (EU) 2023/2854 on harmonised rules on fair access to and use of data. 2023. https://eur-lex.europa.eu/eli/reg/2023/2854/oj

[2] European Commission. Data Act explained. https://digital-strategy.ec.europa.eu/en/policies/data-act

[3] European Commission. Data Act model contractual terms and standard contractual clauses. https://digital-strategy.ec.europa.eu/en/policies/data-act-model-contractual-terms-and-standard-contractual-clauses

[4] European Union. Regulation (EU) 2022/868 on European data governance. 2022. https://eur-lex.europa.eu/eli/reg/2022/868/oj

[5] European Commission. Data Governance Act explained. https://digital-strategy.ec.europa.eu/en/policies/data-governance-act

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

[7] European Commission. AI Act regulatory framework. https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai

[8] European Commission. Guidelines on the definition of an artificial intelligence system. https://digital-strategy.ec.europa.eu/en/library/commission-publishes-guidelines-ai-system-definition-facilitate-first-ai-acts-rules-application

[9] European Union. Regulation (EU) 2016/679 General Data Protection Regulation. 2016. https://eur-lex.europa.eu/eli/reg/2016/679/oj

[10] European Data Protection Board. Guidelines recommendations and best practices. https://www.edpb.europa.eu/our-work-tools/general-guidance/guidelines-recommendations-best-practices_en

[11] European Union. Directive (EU) 2016/943 on the protection of undisclosed know-how and business information. 2016. https://eur-lex.europa.eu/eli/dir/2016/943/oj

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

[13] ENISA. NIS2 Directive resources. https://www.enisa.europa.eu/topics/nis-directive

[14] European Union. Regulation (EU) 2024/2847 on horizontal cybersecurity requirements for products with digital elements. 2024. https://eur-lex.europa.eu/eli/reg/2024/2847/oj

[15] European Commission. A European strategy for data. https://digital-strategy.ec.europa.eu/en/policies/strategy-data

[16] European Commission. Common European data spaces. https://digital-strategy.ec.europa.eu/en/policies/data-spaces

[17] European Commission. Second staff working document on data spaces. 2024. https://digital-strategy.ec.europa.eu/en/library/second-staff-working-document-data-spaces

[18] Data Spaces Support Centre. Blueprint. https://blueprint.dssc.eu/

[19] Data Spaces Support Centre. Starter Kit for Data Space Designers. https://dssc.eu/space/BVE2/1071251457/Starter+Kit

[20] European Commission. Simpl cloud to edge federations empowering EU data spaces. https://digital-strategy.ec.europa.eu/en/policies/simpl

[21] Catena-X Automotive Network. Standards and governance. https://catena-x.net/en/standard-library

[22] Catena-X. Eclipse Tractus-X open-source ecosystem. https://eclipse-tractusx.github.io/

[23] Plattform Industrie 4.0. Manufacturing-X initiative. https://www.plattform-i40.de/IP/Navigation/EN/Manufacturing-X/Manufacturing-X/manufacturing-x.html

[24] European Commission. Manufacturing data spaces workshop summary. 2021. https://digital-strategy.ec.europa.eu/en/library/workshop-manufacturing-data-spaces-summary-results

[25] European Commission. Preparing for manufacturing data spaces. 2021. https://digital-strategy.ec.europa.eu/en/library/way-forward-how-prepare-manufacturing-data-spaces

[26] European Commission. Common European data spaces for smart manufacturing. 2020. https://digital-strategy.ec.europa.eu/en/events/common-european-data-spaces-smart-manufacturing

[27] Gaia-X European Association for Data and Cloud. Architecture document. https://docs.gaia-x.eu/technical-committee/architecture-document/latest/

[28] OPC Foundation. OPC Unified Architecture overview and concepts. https://reference.opcfoundation.org/Core/Part1/v105/docs/

[29] OPC Foundation. OPC UA security model. https://reference.opcfoundation.org/Core/Part2/v105/docs/

[30] OPC Foundation. OPC UA services. https://reference.opcfoundation.org/Core/Part4/v105/docs/

[31] OPC Foundation. OPC UA companion specifications. https://opcfoundation.org/about/opc-technologies/opc-ua/ua-companion-specifications/

[32] OPC Foundation. ISA-95 common object model. https://reference.opcfoundation.org/ISA-95/v100/docs/

[33] Industrial Digital Twin Association. Asset Administration Shell specifications. https://industrialdigitaltwin.org/en/content-hub/aasspecifications

[34] Industrial Digital Twin Association. Asset Administration Shell metamodel. https://industrialdigitaltwin.org/en/wp-content/uploads/sites/2/2024/06/IDTA-01001-3-0-2_SpecificationAssetAdministrationShell_Part1_Metamodel.pdf

[35] Industrial Digital Twin Association. Asset Administration Shell APIs. https://industrialdigitaltwin.org/en/wp-content/uploads/sites/2/2024/06/IDTA-01002-3-0-2_SpecificationAssetAdministrationShell_Part2_API.pdf

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

[37] ISO. ISO IEC 23894 artificial intelligence risk management. https://www.iso.org/standard/77304.html

[38] ISO. ISO IEC 27001 information security management systems. https://www.iso.org/standard/27001

[39] IEC. IEC 62443 industrial communication networks security. https://www.iec.ch/cyber-security

[40] NIST. Artificial Intelligence Risk Management Framework 1.0. 2023. https://doi.org/10.6028/NIST.AI.100-1

[41] NIST. Cybersecurity Framework 2.0. 2024. https://doi.org/10.6028/NIST.CSWP.29

[42] OECD. Recommendation of the Council on Artificial Intelligence. https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449

[43] European Commission. Standardisation request supporting safe and trustworthy artificial intelligence. 2023. https://single-market-economy.ec.europa.eu/news/commission-issues-standardisation-request-support-safe-and-trustworthy-artificial-intelligence-2023-05-22_en

[44] European Commission. Standardisation strategy. https://single-market-economy.ec.europa.eu/single-market/european-standards/standardisation-policy/standardisation-strategy_en

[45] European Commission. Digital Product Passport. https://single-market-economy.ec.europa.eu/single-market/goods/building-blocks/digital-product-passport_en

[46] European Union. Regulation (EU) 2024/1781 establishing a framework for ecodesign requirements for sustainable products. 2024. https://eur-lex.europa.eu/eli/reg/2024/1781/oj

[47] European Commission. European Alliance for Industrial Data Edge and Cloud. https://digital-strategy.ec.europa.eu/en/policies/cloud-alliance

[48] European Commission. Digital Europe Programme. https://digital-strategy.ec.europa.eu/en/activities/digital-programme

[49] European Commission. Competition policy and data. https://competition-policy.ec.europa.eu/sectors/digital-media-and-internet/data_en

[50] European Commission. EU merger control. https://competition-policy.ec.europa.eu/mergers/overview_en

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.