Turning Process Data into Margin: Chemicals and Energy AI Acquisitions

An M&A framework for mapping control authority, process safety, savings persistence and integration economics

Authored by Chennakeshav Adya

Independent Researcher

September 2026

Abstract

Industrial AI businesses serving chemicals and energy companies often claim value from better yield, lower energy use, higher throughput, reduced downtime and safer operation. An acquirer cannot price those claims from model accuracy, connected-tag counts or vendor case studies alone. Process plants operate within engineering envelopes, alarm systems, operating procedures, management-of-change controls, safety instrumented systems and local authority structures. A model may recommend a profitable operating move while the plant declines to implement it because the evidence, permissions, safety case or control-system integration is incomplete. The relevant acquisition question is therefore how process data becomes an authorised operating decision and whether the resulting economic effect persists after ownership, personnel, feedstock, equipment and market conditions change.

This paper develops a Process AI Acquisition Framework for chemicals and energy transactions. It separates data access, model validity, advisory use, supervisory control and closed-loop execution. It maps the independence of the basic process control system and safety instrumented system, tests whether model outputs remain inside approved operating limits, and links each claimed benefit to control actions, production records and finance-approved outcomes. The framework covers historian completeness, sensor quality, model drift, operator adoption, cyber access, intellectual-property rights, customer concentration, revenue recognition, service intensity and post-close integration.

A wholly hypothetical case covers an acquired industrial-AI platform deployed across 72 sites, 214 process units and 1,240 modelled tags. Management identifies USD 12.4 million of annual customer benefit and USD 5.6 million of buyer synergy. Evidence review retains USD 6.9 million of customer benefit after removing unimplemented recommendations, overlapping effects, unapproved baselines and benefits that did not persist through a full operating cycle. The transaction model retains USD 2.1 million of annual buyer synergy after integration engineering, site validation, cyber remediation, customer-change approvals and specialist retention. Every figure is a hypothetical management assumption created to demonstrate the method. It is not observed company data, a market benchmark, a forecast or a valuation opinion.

The evidence base includes the International Energy Agency's 2025 analysis of AI in energy and industrial efficiency; current US Department of Energy work on AI-enabled industrial manufacturing; ISA and IEC functional-safety standards; OSHA process-safety and management-of-change requirements; ISO asset-management standards; NIST and CISA operational-technology security *guidance; API process-safety indicators; current public disclosures concerning industrial-automation transactions; IFRS requirements for business combinations, intangible assets, fair value and revenue recognition; and current AI governance standards. The conclusion is that process-data value becomes transaction value only when the buyer can trace a governed chain from measurement to authorised action, realised outcome and durable cash flow.*

JEL Classification: G32, G34, L65, L71, L95, M15, O32, Q40

Keywords: industrial AI, process control, chemicals, energy, M&A, process safety, optimisation, digital twins, operational technology, valuation

1. DEFINE THE ACQUISITION DECISION

The acquisition decision is how much value a buyer can assign to an industrial-AI business whose products influence chemical or energy operations. The buyer needs to know which applications improve a physical process, which remain advisory demonstrations and which depend on customer-specific engineering that cannot be scaled economically. This distinction affects stand-alone value, synergy, integration cost, deal protection and the sequence of post-close investment.

Process plants convert feedstocks, energy and operating capacity into products under constraints. Pressure, temperature, flow, composition, equipment condition, emissions limits and safety barriers interact. An optimisation recommendation has value only when the receiving plant can verify the measurement, understand the model, authorise the operating change and observe a durable outcome. A dashboard that identifies an opportunity does not create margin until the plant changes its operation and the result survives normal variability.

The framework tests six propositions. The process and economic boundary are defined. The data represents the operating state. The model is valid for the target regime. Control authority is explicit. Process safety and cybersecurity remain intact. Finance can reconcile the effect to cash flow without double counting. Failure at any layer reduces the amount, timing or certainty of value.

Table 1. Evidence required before process AI receives transaction value

Evidence layer	Core question	Minimum record	Transaction consequence
Process identity	Which unit, constraint and decision does the model address?	Process boundary and equipment register	Defines the economically relevant population
Measurement integrity	Do tags represent the physical state at sufficient quality?	Historian, calibration and bad-data records	Determines whether the model input is reliable
Model validity	Does the model perform across approved operating regimes?	Versioned validation and regime map	Supports technical efficacy
Control authority	Who can implement, reject or override the recommendation?	Authority matrix and change record	Establishes operational adoption
Safety and security	Can the application operate without weakening protection layers?	Hazard review and OT security assessment	Constrains deployment and liability
Economic persistence	Does the result persist and reconcile to cash flow?	Baseline, action, outcome and finance approval	Supports valuation and synergy

Each layer requires a named operating owner and a traceable source record.

2. IDENTIFY THE ACQUIRED PRODUCT STACK

An industrial-AI target can include connectors, historians, data lakes, feature pipelines, first-principles models, machine-learning models, optimisation solvers, digital twins, alarm analytics, maintenance applications, dashboards, edge servers and managed services. The same product name may cover very different architectures at different customers. Diligence should therefore trace the deployed components, versions, licences and service obligations at each material site.

Revenue can combine subscriptions, perpetual licences, implementation, engineering, support, cloud consumption, hardware and outcome-linked fees. A recurring invoice may conceal recurring manual work by process engineers or data scientists. A software licence may depend on proprietary connectors, third-party simulation engines or customer-owned models. The buyer should allocate delivery effort and gross margin to the component that actually creates and sustains the result.

Current public disclosures illustrate the strategic importance of process-control software. Emerson completed its acquisition of the remaining AspenTech shares in March 2025 and reports the business within Control Systems and Software [10][11]. The commercial logic of such combinations does not establish that every model, site or customer can be integrated. Transaction value still requires application-level evidence.

3. DEFINE THE PROCESS POPULATION

The denominator begins with sites, process units, equipment trains, operating modes, products, feedstocks, grades, campaigns, constraints and model versions. Management may report thousands of connected tags while only a smaller set affects the controlled decision. The buyer should reconcile contracted sites to commissioned applications, actively used models, authorised recommendations and outcomes with sufficient observation periods.

Process plants spend substantial time outside steady-state conditions. Start-up, shutdown, grade change, catalyst ageing, fouling, maintenance, feedstock change and abnormal operation can move the process beyond the model's validated range. A system that performs well during one stable campaign may provide little value during the periods that determine lost production or safety exposure.

The population register should identify every regime in which the model is permitted, restricted or disabled. It should also identify the relevant production and economic denominator. A throughput model may apply to one bottlenecked unit and have no plant-wide effect when downstream capacity is constrained. Connected-site counts and aggregate tag volumes can materially overstate the addressable benefit.

4. MAP THE INDUSTRIAL DATA LINEAGE

Process data moves through instruments, transmitters, controllers, distributed control systems, historians, integration layers, model services and operator interfaces. Each transformation can change timestamps, units, quality flags, sampling frequency or context. The diligence team should trace every material feature back to the physical measurement and forward to the model output and operating decision.

Historian completeness is not equivalent to measurement quality. A tag can be present while a sensor is biased, frozen, poorly calibrated or operating outside its accurate range. Substitution values and manual entries can conceal outages. The buyer should inspect calibration records, bad-data handling, missingness, time alignment, unit conversion and the treatment of maintenance periods.

Data lineage also defines portability. Customer contracts may allow analysis at one site while prohibiting cross-site training, offshore processing or retention after termination. A model derived from customer data may contain confidential process knowledge. The buyer should map ownership and permitted use for raw data, features, trained parameters, first-principles models and benchmark libraries.

5. SEPARATE PREDICTION FROM PROCESS CONTROL

Industrial AI applications occupy different positions in the control hierarchy. A monitoring model may detect an anomaly. An advisory model may recommend a set-point change. A supervisory optimiser may write targets to the basic process control system. A closed-loop model may directly manipulate an actuator. Each step increases integration value and the evidence required for safety, reliability and governance.

The buyer should identify the exact action path for each material application. It should observe whether operators receive the recommendation, how often they accept it, what changes are executed and why recommendations are rejected. A high model-availability figure can coexist with low economic adoption when the system is bypassed during difficult operating periods.

Authority must remain legible after integration. Model confidence cannot override the plant's operating procedures, alarm philosophy or engineering limits. The acquired business should be able to state which actions are advisory, which require operator confirmation, which pass through supervisory controls and which are prohibited from affecting safety functions.

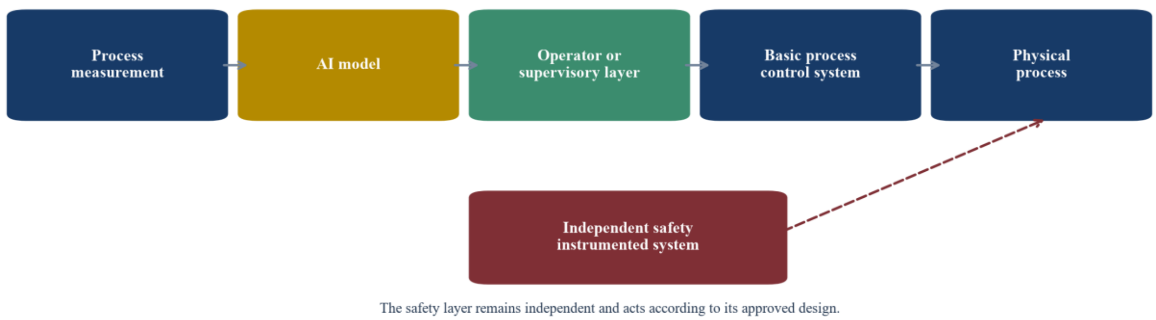

Figure 1. Process AI authority and protection-layer chain

Economic value is released only after the recommendation passes plant authority without weakening independent protection layers.

6. TEST CONTROL AUTHORITY

Control authority should be documented by application, site and operating mode. The record should identify who can approve a model, change a constraint, accept a recommendation, write a set point, suspend the service or return the process to manual operation. It should also record the evidence and approvals required before authority expands.

Acquired products often combine software capability with customer-specific permissions. One refinery may permit a supervisory optimiser to write targets. Another may use the same product only for offline analysis. A portfolio metric that groups both deployments as active optimisation hides a material difference in realised value and liability.

Integration can unintentionally change authority through common interfaces, identity systems or remote-support tools. The buyer should test role definitions, privileged accounts, approval workflows and audit logs before connecting environments. Any expansion of model authority should pass the site's management-of-change process and relevant hazard review.

7. PRESERVE FUNCTIONAL SAFETY

Safety instrumented systems provide independent protection for specified hazardous events. ISA-84 and IEC 61511 define the process-industry framework for safety instrumented systems [4]. A process-AI acquisition should not assume that better prediction permits a reduction in independent protection layers, test intervals or operating discipline.

The buyer should map every model output that can affect a safety-relevant variable, alarm, interlock or operating limit. It should determine whether the model is inside or outside the safety lifecycle and whether any dependency can impair a safety function. Shared infrastructure, clocks, networks or data sources can introduce common-cause failure even when the software components appear separate.

Safety evidence should include hazard studies, safety requirements, proof-test records, bypass logs, alarm rationalisation and change approvals. The integration plan should preserve independence, failsafe behaviour and manual fallback. Savings that depend on operating closer to a safety limit require explicit engineering evidence and should not be valued as ordinary software uplift.

8. APPLY MANAGEMENT OF CHANGE

OSHA's process-safety standard requires covered facilities to manage changes to process chemicals, technology, equipment and procedures [5]. Its guidance identifies computer-program revisions, alarms, interlocks and operating-condition changes as matters that may require management-of-change control [22]. Similar obligations arise through local regulations, company standards and customer procedures in other jurisdictions.

Model deployment and retraining can change how a plant operates even when no physical equipment is replaced. A new objective function, constraint, feature, threshold or user interface can alter operator behaviour. The diligence team should review how the target decides whether a model change triggers technical review, training, documentation and pre-startup checks.

Backlogs matter. A target may have many models awaiting approval while recognising subscription revenue. The buyer should reconcile proposed, validated, approved and operating versions. Integration synergy should include the cost and time required to clear change controls across sites, especially where customers require their own review.

9. VALIDATE THE MODEL BY OPERATING REGIME

Validation should reflect the regimes in which the model will be used. Random train-test splits can place adjacent time periods and similar campaigns in both sets, overstating generalisation. A stronger design separates by time, feedstock, product grade, catalyst condition, equipment state or site. It should also include abnormal and transition conditions when the model claims to support them.

Error metrics need operating meaning. A small average temperature error may be immaterial in one regime and critical near a constraint. A yield prediction can appear accurate because production varies little while failing to identify the incremental effect of a controllable change. The buyer should test performance against the decision the customer actually makes.

Validation should include first-principles consistency, uncertainty bounds, sensitivity to missing inputs and response outside the training domain. A model should abstain or revert safely when evidence is inadequate. The transaction record should preserve failed tests and rejected deployments as well as successful cases.

10. MEASURE OPERATOR ADOPTION

Operators mediate most process-AI value. They understand equipment condition, maintenance work, unusual feedstock and local constraints that may not appear in the model. Diligence should measure recommendation display, acknowledgement, acceptance, implementation, reversal and reason for rejection. Login counts and dashboard views are weak adoption evidence.

Low acceptance can indicate poor recommendations, missing context, slow response, difficult interfaces or weak trust. High acceptance can also be misleading if recommendations are trivial or recorded after the operator independently made the change. The buyer should sample events and interview operators, control engineers and process engineers at representative sites.

The post-close plan should preserve local authority and feedback. Integration targets should reward validated operating outcomes rather than recommendation volume. A model that operators routinely override may still have diagnostic value, while its control and savings claims require adjustment.

Table 2. Process AI authority categories and diligence consequences

Authority category	Model role	Required evidence	Valuation treatment
Monitoring	Detects conditions without recommending action	Detection performance and response workflow	Diagnostic value only
Advisory	Recommends an action for human decision	Recommendation, acceptance and outcome trail	Value based on implemented actions
Supervisory	Writes targets or constraints to process control	Approved authority, validation and rollback	Value after site approval and operating evidence
Closed loop	Directly manipulates the process	Full control design, safety review and failover	Highest evidence burden and liability review
Safety related	Influences a protection function or safety limit	Functional-safety lifecycle evidence	Separate engineering and legal assessment

Classification follows the deployed decision path at each customer site.

11. RECONCILE THE CLAIMED BASELINE

Every savings claim begins with a counterfactual. The target should state what would have happened without the model, over which period and under which operating conditions. Baselines can use historical cohorts, matched operating periods, engineering models, controlled trials or approved budgets. Each method has limitations that should remain visible in the transaction model.

Production economics move with feedstock price, product price, grade mix, ambient conditions, maintenance state and plant utilisation. A benefit recorded during a favourable campaign may disappear when these variables change. The buyer should reproduce the baseline calculation from source data and test sensitivity to the material drivers.

Finance approval improves evidence quality. The plant should distinguish physical effects from price effects and separate gross opportunity from realised cash. A yield gain that increases intermediate output may have limited value when the downstream unit is constrained. The transaction model should retain the bottleneck and market conditions that allowed the benefit to be monetised.

12. PREVENT SAVINGS OVERLAP

Industrial-AI vendors often report benefits across energy, yield, throughput, maintenance, emissions and avoided downtime. These categories can describe the same underlying event. Lower steam consumption can reduce energy cost and emissions. Higher availability can increase throughput only when demand, feedstock and downstream capacity are available.

The diligence ledger should assign each outcome to a primary physical effect and record secondary consequences separately. It should identify whether the customer, another vendor or an internal improvement programme contributed to the same result. A benefit cannot be counted twice because it appears in two dashboards or supports two commercial narratives.

Portfolio aggregation creates another risk. Site case studies may cover selected successful units and use different baselines. The buyer should aggregate only comparable, non-overlapping effects and show excluded cases. Missing or unsuccessful deployments are part of the evidence needed to estimate probability-weighted value.

13. TEST SAVINGS PERSISTENCE

Short trials can establish technical possibility while providing limited evidence of durable economics. Chemicals and energy assets experience seasonal demand, feedstock changes, catalyst ageing, fouling, maintenance cycles and operator rotation. A claimed benefit should pass through a representative operating cycle before it supports recurring valuation.

Persistence should be tested by cohort. The buyer should compare the initial result, six-month retention, twelve-month retention and performance after a material regime change. It should identify how much engineering effort was required to preserve the result. A benefit sustained through frequent manual retuning has different margin and scalability from a stable model.

Customer renewal provides commercial evidence, although it does not prove the size of the physical benefit. Contracts may renew because the application supports compliance, engineering workflow or broader platform access. The transaction case should connect physical persistence, customer use, service cost and retained revenue without treating any one measure as a complete proxy.

14. MEASURE SENSOR AND HISTORIAN DEBT

Industrial-AI performance depends on existing plant instrumentation and data infrastructure. Sensor bias, calibration gaps, unreliable analysers, inconsistent tag names, low sampling rates and historian compression can reduce model value. A target may spend substantial implementation effort repairing these conditions while reporting the work as software deployment.

The buyer should create a data-debt register by site. It should include missing instruments, bad actors, manual measurements, clock drift, gaps, quality flags, unit inconsistencies, obsolete interfaces and unsupported historians. Remediation cost should be separated from reusable product development.

Data debt can also become a commercial opportunity if the target has a repeatable method for assessing and improving customer readiness. The evidence should show standard tools, cycle time, staffing and gross margin. Bespoke remediation that depends on a few senior engineers should be valued as a service capability with capacity constraints.

15. TEST THE DIGITAL TWIN

The term digital twin can refer to a first-principles simulator, a calibrated hybrid model, a statistical surrogate, an equipment model or a visual asset representation. Diligence should define the model's purpose, state variables, boundary conditions, calibration process and decision role. A label alone provides no evidence of fidelity or portability.

ISO 23247 provides a digital-twin framework for manufacturing [15]. The buyer should map the target's architecture to observable entities, data collection, modelling and services. For process industries, the review should also test mass and energy balance, thermodynamic consistency and response across operating regimes.

Twin maintenance is part of the economics. Equipment modifications, control tuning, catalyst changes and instrument replacement can make a calibrated model stale. The target should record recalibration events, owners and acceptance tests. A portfolio of twins without funded maintenance can become a growing liability after acquisition.

16. PROTECT OPERATIONAL TECHNOLOGY

Process-AI platforms can connect historians, engineering workstations, controllers, cloud services and remote-support environments. NIST SP 800-82 Rev. 3 addresses OT security while recognising performance, reliability and safety requirements [7]. CISA guidance emphasises that business decisions can affect OT security and continuity [8].

The buyer should inventory assets, accounts, trust relationships, protocols, certificates, remote-access routes and third-party dependencies. It should identify where the acquired platform crosses security zones or introduces outbound connectivity. Unsupported operating systems, shared credentials and unmanaged edge devices can convert a local weakness into a portfolio risk.

Security controls must respect process operation. Patching, identity changes and network segmentation can affect latency, availability and validation. The integration plan should coordinate cybersecurity, control engineering, safety and production. Recovery tests should demonstrate that sites can return to a known safe operating mode when the AI service or its network connection fails.

17. GOVERN MODEL CHANGES

A combined business may inherit multiple code bases, feature pipelines, solvers and local patches. The buyer should create a model register linking each deployed version to source code, training data, process units, validated regimes, constraints, approvals, dependencies, monitoring and rollback. An undocumented model can carry operating risk even when customers report good results.

Version control should extend beyond model weights. Changes to data cleaning, timestamps, unit conversion, objective functions, economic coefficients, optimiser bounds and user interfaces can alter decisions. The model register should capture the complete decision chain and the evidence that supports each production release.

NIST's AI Risk Management Framework organises work around govern, map, measure and manage [13]. ISO/IEC 42001 provides requirements for an AI management system [14]. These frameworks support transaction governance, although plant-specific safety and control obligations remain decisive for process applications.

18. ASSESS PROCESS-SAFETY INDICATORS

API RP 754 provides a tiered framework for process-safety indicators in refining and petrochemicals. The fourth edition, published in August 2026, continues the use of leading and lagging indicators to identify weaknesses and learn from events [9]. An AI vendor's safety claim should connect to the indicators and protection activities used by the customer.

The buyer should separate early warning, barrier monitoring, incident prevention and post-event analysis. A model that predicts a precursor may support a leading indicator without proving that it prevented a release. The evidence should show alert quality, response time, completed action and the condition of the relevant barrier.

Safety benefits require conservative valuation. Severe events are infrequent, site-specific and influenced by multiple controls. The transaction model should avoid monetising avoided catastrophes as recurring ordinary income. Safety capability can support customer retention, risk reduction and strategic value while remaining subject to separate engineering, insurance and legal assessment.

19. RECONCILE REVENUE QUALITY

Revenue should be separated into licences, subscriptions, implementation, engineering, support, cloud usage, hardware and outcome-linked consideration. IFRS 15 requires revenue recognition based on performance obligations and transfer of promised goods or services [17]. The diligence model should follow contracts, invoices, collections and delivery evidence.

Reported recurring revenue may depend on recurring engineering, customer-specific model maintenance or third-party licences. The buyer should reconcile active sites, authorised applications, service tickets, engineering time, renewal and contribution margin. Revenue from a model that is installed but disabled or outside its validated regime requires specific review.

Outcome-linked contracts need the underlying benefit calculation and approval trail. A fee can be capped, disputed or delayed even when the technical application performs. The buyer should test whether the target's revenue, the customer's recognised benefit and the transaction synergy use the same economic event, preventing double counting.

20. CONSTRUCT THE HYPOTHETICAL CASE

The illustrative target serves 72 sites and 214 process units. Its platform receives 1,240 modelled tags and operates 168 customer applications. Management classifies 126 applications as active optimisation and reports USD 12.4 million of annual customer benefit. The initial buyer case assumes that a common deployment and support model can create USD 5.6 million of annual synergy.

Evidence review first reconciles the population. Of 126 active-optimisation applications, 101 have complete action logs, 84 have finance-approved baselines, 73 show implemented recommendations and 59 retain an observed benefit through a representative operating cycle. The remaining applications may have diagnostic or strategic value, while their savings claims are excluded from the evidence-adjusted total.

The hypothetical review retains USD 6.9 million of customer benefit after removing unimplemented recommendations, overlapping effects, price-only movements, unapproved baselines and benefits that failed to persist. Integration costs and probability adjustments reduce buyer synergy to USD 2.1

million. These amounts are hypothetical management assumptions for demonstrating the framework and do not represent any observed company or transaction.

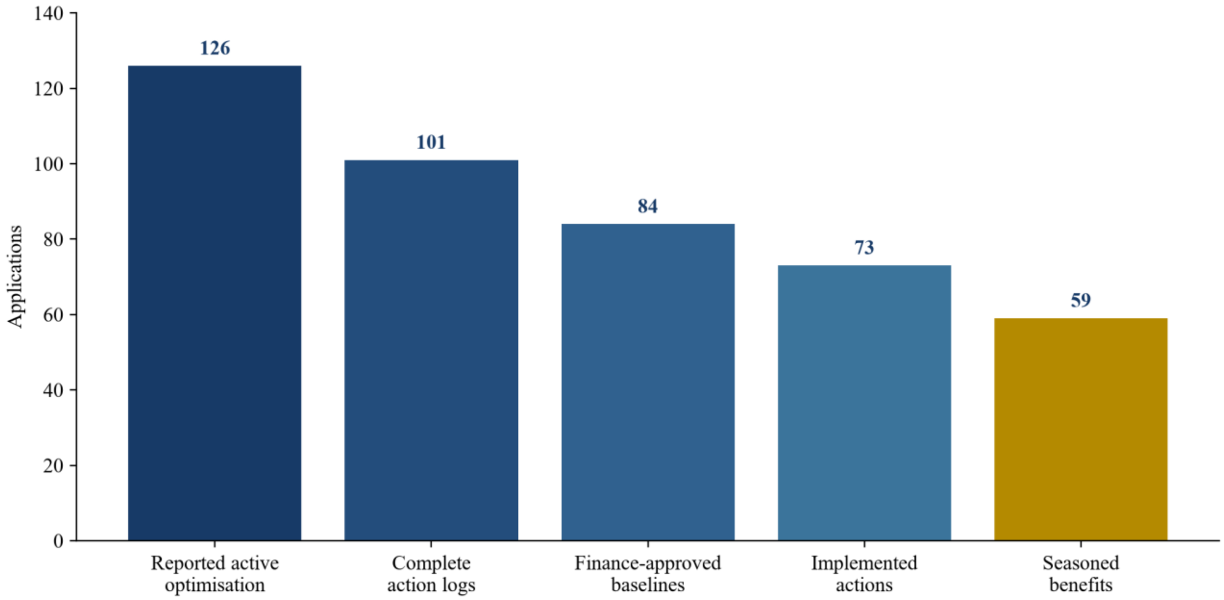

Figure 2. Hypothetical process AI evidence funnel

Counts are management assumptions created solely to demonstrate the diligence method.

Table 3. Hypothetical evidence-adjusted process AI economics

Item	Management case	Evidence adjustment	Retained annual value
Customer energy and yield benefit	7.2	Baseline, overlap and persistence	4.1
Customer throughput and availability benefit	5.2	Bottleneck, implementation and seasoning	2.8
Buyer deployment and support synergy	2.0	Site engineering and customer approvals	0.8
Buyer platform and infrastructure synergy	1.6	Architecture and cyber remediation	0.7
Buyer product and commercial synergy	2.0	Retention, rights and execution probability	0.6
Total buyer synergy	5.6	Combined evidence and execution deductions	2.1

All amounts are management assumptions in USD millions and do not represent an observed transaction.

21. TRANSLATE EVIDENCE INTO VALUATION

IFRS 13 defines fair value as a market-based measurement using assumptions that market participants would use [18]. The transaction case should separate stand-alone cash flow, required reinvestment, integration cost and buyer-specific synergy. An opportunity that depends on the buyer's customer relationships or installed control systems does not become stand-alone value merely because the buyer can pursue it.

The stand-alone model should begin with revenue and contribution margin by customer, site and application. It should incorporate renewal, engineering effort, third-party costs, cloud consumption, product maintenance and specialist dependence. Evidence adjustments reflect data rights, validated regimes, control authority, cyber posture, safety obligations, concentration and benefit persistence.

IFRS 3 and IAS 38 govern business combinations and identifiable intangible assets [19][20]. Accounting allocation and commercial valuation answer related questions through different methods. The diligence record should support both while preserving the operating evidence needed to set price, structure protections and approve integration capital.

22. PRICE IMPLEMENTATION CAPACITY

Process-AI growth can be constrained by scarce engineers who understand control systems, thermodynamics, process safety and customer operations. The buyer should measure implementation hours, site travel, commissioning cycles, validation effort, support load and the share of work performed by founders or named specialists.

Backlog quality depends on customer readiness. A signed order may wait for historian access, instrumentation repair, cyber approval, plant shutdown or management-of-change review. Revenue forecasts should distinguish contracted work, technically ready sites, scheduled commissioning and benefits contingent on customer action.

Scalability should be tested with cohorts. The buyer should compare implementation time and margin across early and recent deployments, products, geographies and plant types. Reusable connectors and templates create value when they reduce effort without weakening validation. Growth that requires proportional specialist hiring should be valued with service-business economics.

23. TEST CUSTOMER CONCENTRATION AND ACCESS

Industrial-AI targets can depend on a small number of large energy or chemical customers. One enterprise agreement may cover many sites and create an appearance of diversification. The buyer should map revenue, gross margin, data rights, references, renewal authority and termination rights at the ultimate customer level.

Customer access affects diligence. Some contracts restrict interviews, data transfer or disclosure of performance. The buyer should request permission for representative customer discussions covering operations, engineering, finance, cybersecurity and procurement. Management-selected references should be supplemented by contract and usage evidence.

Change-of-control clauses and vendor approvals can affect value. Customers may require a security review, requalification or consent before the buyer changes hosting, support locations or product integration. Synergies that require cross-selling or common data infrastructure should be probability-weighted by customer rights and approval timelines.

24. EVALUATE INTELLECTUAL PROPERTY AND DEPENDENCIES

The target's defensible capability may reside in code, domain models, process templates, data mappings, deployment tools or specialist know-how. Diligence should identify what is owned, licensed, customer-specific, open source or dependent on third-party software. It should also test assignment and change-of-control provisions.

First-principles models can incorporate public equations, vendor libraries and customer calibrations. Machine-learning models can embed customer process behaviour. The buyer should determine whether the target can redeploy a model or template without disclosing customer confidential information or violating licence terms.

Key-person dependence belongs in the valuation and retention plan. Undocumented heuristics and manual tuning can appear as proprietary software performance. The buyer should reproduce a deployment and model update using ordinary staff and controlled documentation before assigning portfolio-wide platform value.

25. STRUCTURE PURCHASE-PRICE PROTECTION

Known evidence gaps should be allocated through specific transaction terms. Holdbacks can address customer consent or unresolved data rights. Earn-outs can depend on retained recurring revenue, authorised deployments and contribution margin. Special indemnities can address disclosed cyber, intellectual-property or process-liability matters.

Milestones should be observable and resistant to manipulation. Connected tags, model accuracy and recommendation counts can improve without producing durable customer economics. Better milestones combine site acceptance, implemented actions, benefit persistence, customer retention and cost to serve over a defined population and period.

The buyer should preserve access to evidence after closing. Sellers and management may need to support customer confirmations, model provenance, benefit reconciliation and regulatory responses. Integration actions that would make a milestone impossible to measure should be addressed explicitly in the agreement.

26. PLAN THE INTEGRATION SEQUENCE

The first wave should use applications with complete lineage, stable regimes, clear authority, strong customer support and reversible deployment. Its purpose is to validate the combined operating model. Safety-related, weakly documented or highly customised applications should remain outside the first wave until their evidence gaps are resolved.

Each wave should include architecture review, data reconciliation, offline validation, shadow operation, controlled release and post-release monitoring. Rollback criteria should be approved before production. The customer retains operating authority throughout the transition.

Shared infrastructure should be introduced in layers. Identity, observability and support processes may be standardised before model logic or control interfaces. This sequence can improve visibility while limiting common-cause failure. The plan should record cost, benefit, incident response and customer acceptance for every wave.

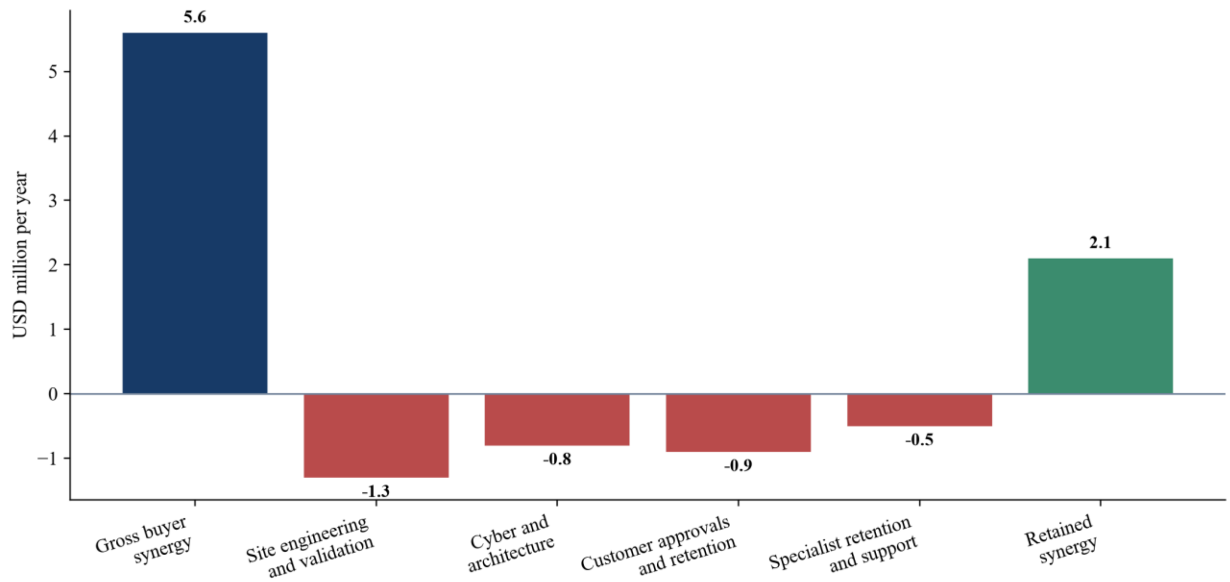

Figure 3. Hypothetical buyer synergy validation bridge

All amounts are management assumptions in USD millions per year.

27. DESIGN THE INTEGRATION SCORECARD

The board scorecard should combine technical, operating, safety, commercial and financial evidence. Technical measures include data completeness, regime coverage, prediction error and model drift. Operating measures include recommendation acceptance, implemented actions, constraint time and fallback events. Safety measures include barrier status, management-of-change completion and incidents.

Commercial measures should include active authorised applications, retention, contribution margin, implementation backlog and customer approvals. Financial measures should include finance-approved customer benefit, integration cash cost and retained buyer synergy. Each measure needs a definition, owner, source and ageing rule.

Portfolio averages can conceal material exceptions. The scorecard should allow drill-down by customer, site, unit, application, regime and version. Exceptions should remain open until resolved, retired or accepted by the appropriate authority. A declared synergy without the required evidence stays outside realised value.

28. EXECUTE THE FIRST ONE HUNDRED DAYS

Days zero to thirty should preserve process continuity and evidence. The buyer freezes undocumented model and control changes, confirms remote access and backups, reconciles customers, sites, applications and contracts, and identifies critical engineers and customer approvals. It establishes escalation paths for safety, cybersecurity and production issues.

Days thirty-one to sixty should standardise the application register, data lineage, model inventory, authority matrix and benefit ledger. The team selects a limited integration wave and completes offline and shadow validation. It prioritises rights, cyber and safety gaps according to operating consequence.

Days sixty-one to one hundred should release approved changes with rollback and monitoring. The board receives evidence on adoption, operating effects, customer response, service effort and retained synergy. Further waves proceed when the integration method has produced accepted and reproducible results.

Table 4. One-hundred-day plan for a process AI acquisition

Period	Operating priority	Evidence deliverable	Investment decision
Days 0 to 30	Preserve operations, access, records and customer commitments	Reconciled site, application, model and contract registers	Freeze unsupported changes
Days 31 to 60	Standardise lineage, authority, governance and benefit measurement	First-wave validation and acceptance pack	Approve limited integration wave
Days 61 to 100	Deploy with shadow operation, rollback and monitoring	Site-level operating and economic evidence	Release evidence-backed synergy
Quarter 2 onward	Scale reusable components and retire unsupported variants	Seasoned portfolio scorecard	Reprice cohorts and capital plan

The sequence preserves plant authority while establishing evidence for scalable integration.

29. RECOGNISE LIMITATIONS AND RESEARCH PRIORITIES

Process performance is context dependent. A validated result applies to a defined unit, regime, model and period. It does not establish performance after feedstock, equipment, catalyst, control, market or staffing changes. Rare safety events and abnormal operations can leave substantial uncertainty even when normal-operation data is abundant.

The hypothetical case omits tax, financing, working capital, transaction fees, insurance, environmental liabilities, legal claims and many site-specific effects. Customer approvals, operating constraints and actual integration demand remain unknown until verified in a transaction. The framework supports diligence and governance; it does not produce a valuation opinion.

Further research should examine which evidence best predicts savings persistence, whether operator-rejection patterns identify model weakness earlier than aggregate accuracy, and how common infrastructure changes cyber and common-cause risk. Comparable studies require governed plant data, consistent definitions and publication of unsuccessful deployments as well as successful ones.

30. CONCLUSION

Process data becomes margin through a governed operating chain. The chain begins with a reliable measurement and a model valid for the relevant regime. It continues through explicit authority, implemented action and an observed physical result. Finance then reconciles the result to incremental cash flow and tests whether it persists.

Chemicals and energy acquisitions require particular discipline because process control, functional safety, operational technology and customer approvals constrain deployment. A common platform can create valuable reuse, support and cross-selling economics. Those benefits should be released according to site-level evidence and the cash cost of integration.

The acquisition framework gives the buyer a defensible basis for price, contractual protection, integration sequencing and capital allocation. It also protects customers and operating teams by keeping process authority, safety and rollback visible throughout the transaction.

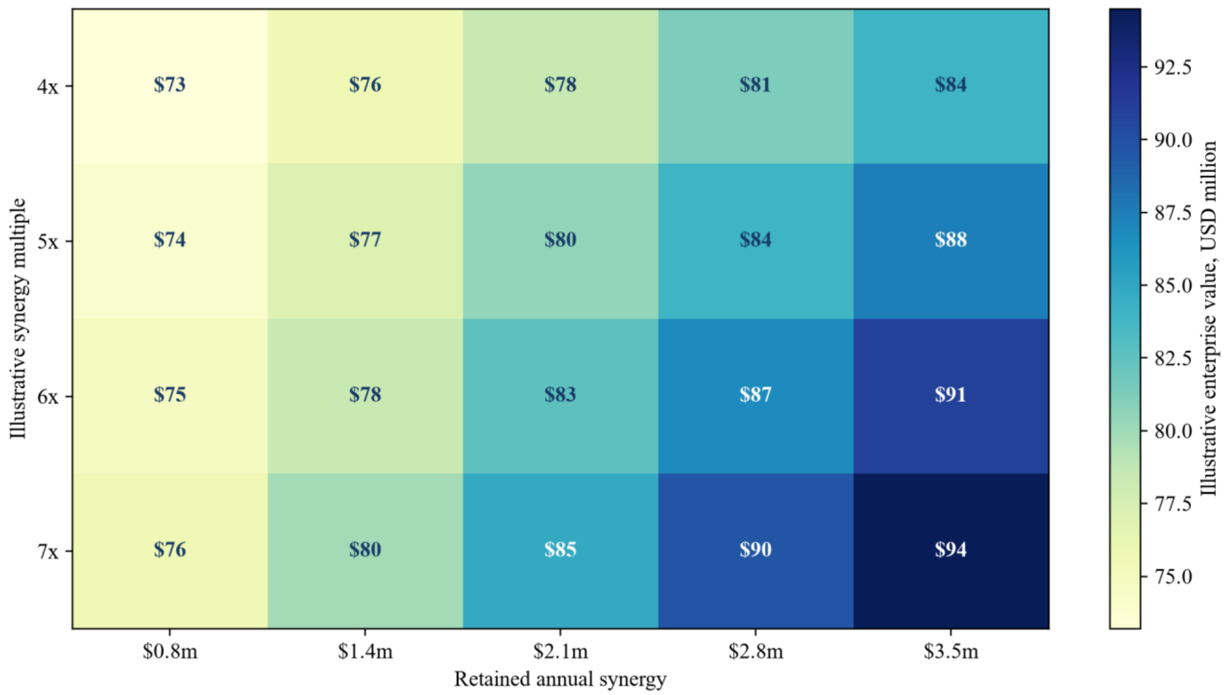

Figure 4. Hypothetical enterprise-value sensitivity to retained process AI synergy

Values are management assumptions in USD millions and exclude net debt and transaction adjustments.

REFERENCES

[1] International Energy Agency (2025). Energy and AI. https://www.iea.org/reports/energy-and-ai

[2] United States Department of Energy (2026). Reenvisioning Advanced Manufacturing and Industrial Productivity. https://www.energy.gov/undersecretaryforscience/genesis-mission/reenvisioning-advanced-manufacturing-and-industrial

[3] United States Department of Energy (2022). Agency Inventory of Artificial Intelligence Use Cases. https://www.energy.gov/sites/default/files/2022-07/DOE_Agency_Inventory_of_AI_Use_Cases.pdf

[4] International Society of Automation (2018). ISA-84 Series and IEC 61511 Functional Safety Standards. https://www.isa.org/standards-and-publications/isa-standards/isa-84-standards

[5] Occupational Safety and Health Administration. 29 CFR 1910.119 Process Safety Management of Highly Hazardous Chemicals. https://www.osha.gov/laws-regs/regulations/standardnumber/1910/1910.119

[6] International Organization for Standardization (2024). ISO 55001 Asset management systems Requirements. https://www.iso.org/standard/83054.html

[7] National Institute of Standards and Technology (2023). Guide to Operational Technology Security, NIST SP 800-82 Revision 3. https://csrc.nist.gov/pubs/sp/800/82/r3/final

[8] Cybersecurity and Infrastructure Security Agency (2024). Principles of Operational Technology Cybersecurity. https://www.cisa.gov/news-events/alerts/2024/10/01/asds-acsc-cisa-fbi-nsa-and-international-partners-release-guidance-principles-ot-cybersecurity

[9] American Petroleum Institute (2026). API RP 754 Process Safety Performance Indicators, Fourth Edition. https://www.api.org/news-policy-and-issues/news/2026/08/06/api-publishes-fourth-edition-process-safety-performance-indicators-standard

[10] Emerson Electric Co. (2025). Annual Report for the year ended 30 September 2025. https://www.sec.gov/Archives/edgar/data/32604/000003260425000105/a2025annualreport.pdf

[11] Emerson Electric Co. (2025). Completion of the Acquisition of Remaining AspenTech Shares. https://www.sec.gov/Archives/edgar/data/32604/000114036125008250/ny20044173x8_exa5x.htm

[12] Honeywell International Inc. (2026). Annual Report for the year ended 31 December 2025. https://www.sec.gov/Archives/edgar/data/773840/000077384026000031/hon2025annualreport.pdf

[13] National Institute of Standards and Technology (2023). Artificial Intelligence Risk Management Framework 1.0. https://www.nist.gov/itl/ai-risk-management-framework

[14] International Organization for Standardization (2023). ISO/IEC 42001 Artificial intelligence management system. https://www.iso.org/standard/42001

[15] International Organization for Standardization (2021). ISO 23247-1 Digital twin framework for manufacturing. https://www.iso.org/standard/75066.html

[16] OPC Foundation. OPC Unified Architecture. https://opcfoundation.org/about/opc-technologies/opc-ua/

[17] IFRS Foundation (2026). IFRS 15 Revenue from Contracts with Customers. https://www.ifrs.org/issued-standards/list-of-standards/ifrs-15-revenue-from-contracts-with-customers/

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

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

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

[21] International Energy Agency (2025). Energy Efficiency 2025 Industry. https://www.iea.org/reports/energy-efficiency-2025/industry

[22] Occupational Safety and Health Administration. Appendix C to 29 CFR 1910.119. https://www.osha.gov/laws-regs/regulations/standardnumber/1910/1910.119AppC

[23] International Society of Automation. ISA-95 Enterprise-Control System Integration. https://www.isa.org/standards-and-publications/isa-standards/isa-95-standard

[24] Cybersecurity and Infrastructure Security Agency. Industrial Control Systems Recommended Practices. https://www.cisa.gov/resources-tools/resources/ics-recommended-practices

[25] United States Department of Energy (2025). Artificial Intelligence Strategy. https://www.energy.gov/sites/default/files/2025-09/EXEC-2025-010630%20-%20250923_%20DOE%20AI%20Strategy%20VFinal.pdf

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.