

UK Manufacturing Roll-Ups: Valuing Predictive-Maintenance Data as an Integration Asset

An evidence-led framework for testing whether combined failure histories and sensor coverage can support measurable uptime and margin improvement

Authored by Chennakeshav Adya

Independent Researcher

September 2026

Abstract

Manufacturing roll-ups are commonly underwritten on procurement, footprint, overhead and commercial synergies. Predictive-maintenance data can support another source of value: combining asset identities, operating histories, work orders, failures, sensor signals and maintenance outcomes across acquired businesses may improve reliability decisions and reduce avoidable downtime. The economic benefit is conditional. Data may be incomplete, legally constrained, technically incompatible or too weakly labelled to support a repeatable model. A buyer that treats raw records as a valuable asset without testing provenance, coverage and decision usefulness can pay for a capability that does not transfer or scale.

This paper develops a Predictive-Maintenance Integration Asset Framework for UK manufacturing acquisitions. The framework tests whether the acquired data can be controlled, combined, interpreted and converted into actions that improve cash flow. It separates accounting recognition from transaction value, maps the evidence chain from equipment to financial outcome, and values the capability through buyer-specific incremental cash flows after integration cost, operating cost, ramp time and execution risk. It also incorporates asset-management, operational-technology security, safety, privacy, competition and post-merger governance requirements.

The method is demonstrated through a hypothetical platform-and-target transaction. Management assumptions include platform revenue of GBP 210 million and EBITDA of GBP 26 million; target revenue of GBP 85 million and EBITDA of GBP 9 million; 340 target critical assets; and uneven identity, sensor and work-history coverage. The illustrative steady-state gross annual benefit is GBP 4.9 million, recurring programme cost is GBP 0.9 million and one-time integration investment is GBP 5.8 million. A five-year cash-flow model with a management-assumed fourteen-percent discount rate produces an illustrative buyer-specific net present value of approximately GBP 3.2 million before tax, financing, working capital, purchase-price allocation or negotiated value sharing. These figures demonstrate the method only. They are not observed company data, a valuation opinion, accounting advice, engineering advice, legal advice or investment advice.

JEL Classification: C55, G34, L60, M21, M41, O33

Keywords: UK manufacturing roll-ups, predictive maintenance, industrial data, M&A integration, synergy valuation, asset management, sensor coverage, failure history, manufacturing analytics, post-merger integration

1. Define the investment decision

The investment decision is whether predictive-maintenance data acquired through a UK manufacturing roll-up can create incremental, controllable and durable cash flow for the combined group. The question is narrower than whether each company owns sensors, a maintenance system or an analytics pilot. The buyer

needs evidence that the combined records can improve a specified decision, such as when to inspect a bottleneck, how to prioritise a shutdown, which spare to hold, or whether a recurring failure warrants redesign.

The answer affects valuation, diligence scope, integration cost, synergy timing, transaction protections and the first hundred days. It can influence the maximum price the buyer is prepared to pay, the conditions attached to a synergy case, the treatment of missing data, the transition plan for operational systems and the post-close investment budget. It can also show that a claimed data advantage has no separable transaction value because the same capability could be built independently at lower cost or because the evidence does not support reliable deployment.

The framework therefore treats data as an integration asset only when five conditions are met. The buyer must have lawful and contractual control; the records must be attributable to identifiable assets and operating states; the combined population must improve a defined maintenance decision; the organisation must possess the engineering and operating capability to act; and the incremental cash benefit must exceed implementation, governance and continuing operating costs. Failure at any condition reduces or removes the buyer-specific value case.

Table 1. Predictive-maintenance integration evidence hierarchy

Evidence layer	Typical record	Question answered	Transaction use
Asset identity	Functional location, serial number, model, site and line	Which physical asset generated the record?	Population and perimeter control
Operating context	Duty, load, product, shift, environment and process state	Under what conditions was the asset operating?	Comparability and model segmentation
Condition signal	Vibration, temperature, acoustic, oil, current or process data	What observable change preceded intervention or failure?	Detection and diagnostic testing
Maintenance action	Work order, inspection, part, labour and completion record	What action was taken, when and by whom?	Intervention design and cost evidence
Outcome	Failure avoided, downtime, quality, safety and recurrence	Did the action improve the operating result?	Benefit attribution and validation
Financial mechanism	Contribution margin, labour, contractor, spares and inventory	How did the operating result affect cash flow?	Synergy valuation and tracking

Records become transaction evidence only when identity, operating context, control and economic consequence can be traced.

2. Distinguish records from an integration asset

Raw records are inputs. An integration asset is a governed capability that joins records across time, systems and acquired entities to improve repeatable decisions. The difference matters because data volume can grow while economic usefulness remains unchanged. Millions of unlabelled sensor readings may have less value than a smaller history in which asset identity, operating regime, work performed and verified outcome are linked.

The capability also depends on transferability. A target may use vendor-hosted software under a non-transferable licence. Machine data may be stored in a proprietary format, retained for only a short period or accessible solely through an original equipment manufacturer. Employee or contractor records can contain personal information. Customer agreements may restrict use of production or quality data. The

buyer must test rights, interfaces, retention and consent before including the records in an integration thesis.

The economic asset is therefore the controlled evidence chain, combined with people, operating routines and systems that convert evidence into action. It may include taxonomies, mapping rules, validated features, failure labels, maintenance playbooks, model documentation and feedback loops. Its value can disappear if a key engineer leaves, an interface is retired, the target loses access to historical data or the merged group cannot standardise decisions across sites.

3. Establish the roll-up thesis

A roll-up should state why combining businesses creates a maintenance advantage that each business could not achieve as effectively alone. Scale alone is insufficient. A credible thesis may arise from repeated equipment families, complementary failure histories, broader operating regimes, specialist engineering coverage, shared spare-parts economics or the ability to fund a common data and reliability platform.

The thesis should identify the competitive and operational boundary. If acquired factories make unrelated products with different machinery, loads and maintenance practices, pooled records may not support generalisation. If several sites operate the same compressors, pumps, furnaces, robots or machine tools, the combined population may improve failure-mode coverage and allow specialist knowledge to travel. The buyer should specify where pooling is permitted and where local models or local engineering judgement remain necessary.

The acquisition rationale should also compare non-transaction alternatives. The platform could license a vendor solution, form a data-sharing arrangement, hire reliability engineers or build the capability organically. The UK Competition and Markets Authority assesses merger efficiencies through evidence that addresses timeliness, likelihood, sufficiency and merger specificity. Internal board papers, integration plans, operating data and prior delivery records can carry more weight than a post hoc assertion that data combination will create value [1-3].

4. Map the asset and data perimeter

The diligence population should begin with the physical assets that matter to production, safety, quality, energy and customer commitments. Each asset needs a persistent identifier, site, line, function, manufacturer, model, age where known, rated duty, actual duty, criticality and dependency map. The data perimeter should then show which systems hold condition signals, alarms, inspections, work orders, parts, labour, production losses and quality events.

Roll-up targets often use different naming structures. One business may organise records by cost centre, another by production line and another by equipment hierarchy. Similar assets can have different local names, while the same tag can be reused at several sites. A crosswalk must preserve original identifiers, record the mapping rationale and flag unresolved matches. Forced matching can contaminate the pooled history and create false patterns.

The perimeter should include interfaces and hidden dependencies. Sensor gateways, historians, manufacturing execution systems, computerised maintenance-management systems, spreadsheets, original-equipment-manufacturer portals and contractor databases can all hold material evidence. Ownership, access credentials, export capability, licence terms, retention settings and data quality should be recorded for each source. The transaction plan must say which source transfers, which remains under a transitional service, which requires a new interface and which cannot be relied upon.

5. Build the evidence-to-value architecture

The architecture should allow a reviewer to move from a claimed synergy to the underlying asset population, source records, transformation logic, model output, maintenance action and financial result. The chain should preserve source timestamps, units, operating state, quality flags and version history. It should also retain human decisions, including overrides, because the economic outcome usually depends on how engineers and planners respond to an alert.

Data lineage is essential during integration. A temperature value without a calibrated sensor reference, sampling interval and operating load may be unusable. A failure label without the replaced component, root-cause assessment and return-to-service record can teach the wrong relationship. A downtime record without the production constraint and contribution margin cannot support a financial claim. The architecture should expose these gaps rather than silently impute them.

Figure 1. Predictive-maintenance data-to-value architecture



The integration asset links controlled operating evidence to human-approved maintenance actions and verified financial outcomes.

6. Test data control and transfer rights

The buyer should classify every material dataset by legal controller, contractual owner, permitted purpose, security classification, retention period and transfer mechanism. Industrial records can contain several rights at once. The factory may own raw operating data, a vendor may own software or derived features, an employee may be identifiable in shift and maintenance records, and a customer may control product or quality information.

The Information Commissioner's Office states that data sharing must be considered during merger due diligence when a transaction changes or adds a controller. The parties should establish what personal data is transferring, why it was obtained, the lawful basis, the information provided to individuals, retention and security arrangements [4]. The Data (Use and Access) Act 2025 may change parts of the UK data-protection framework and associated guidance; transaction counsel should confirm the current position for the specific records and processing [5].

Contracts should address extraction, continued access, deletion, model training, derived data, subcontractors, audit, incident handling and exit. A warranty that data exists does not establish its usability. The buyer may require a pre-close export, a test restoration, interface documentation and evidence that the seller can grant the promised rights. Material uncertainty should appear in the valuation case and transaction protections.

7. Assess identity and taxonomy quality

Identity quality determines whether histories can be combined. The buyer should calculate the proportion of critical assets with stable identifiers, confirmed make and model, functional location, parent-child

hierarchy and mapping to work orders and sensors. It should also test whether identifiers survived prior system migrations or were recreated, causing one physical machine to appear as several records.

Taxonomies require similar control. Failure mode, cause, remedy, maintenance type, priority and downtime reason may use local codes or free text. A shared dictionary should map the original codes without discarding source meaning. Ambiguous terms should remain unresolved until a domain specialist confirms them. A one-to-many mapping may be necessary where a broad legacy code covers several distinct failure mechanisms.

The integration plan should prioritise decision-critical mappings. A complete enterprise taxonomy can take years. The first hundred days can focus on bottleneck equipment, high-cost failures, statutory inspections and repeated asset families. A controlled minimum vocabulary can support early value while the broader hierarchy is improved.


8. Measure sensor and event coverage

Coverage has several dimensions. Asset coverage asks whether the relevant equipment is instrumented. Time coverage asks whether histories span enough seasons, loads and maintenance cycles. Signal coverage asks whether the measured variables relate to the failure modes. Event coverage asks whether enough verified failures and interventions exist to test performance. Operating coverage asks whether the data represent the conditions in which the model will be used.

Sampling frequency and synchronisation also matter. Slowly sampled temperature may support trend monitoring but miss a transient event. Vibration channels may use different units, filter settings or mounting points. Clock drift can misalign an alarm with the operating event that caused it. Missingness can be informative, such as a sensor going offline during a shutdown, but it can also create false patterns if the model treats absence as a normal value.

The buyer should report coverage by asset family and use case. A portfolio-wide average can hide a weak bottleneck population. The report should state which assets have usable identity, signals, work history and outcome labels, and which depend on new instrumentation. New sensor investment belongs in the integration cost and timing model.

Figure 2. Illustrative evidence-readiness heat map



Management-assumed scores demonstrate how identity, signals, work history and outcomes can vary by asset family.

9. Reconstruct failure and maintenance histories

Predictive-maintenance value depends on verified outcomes. The buyer should link each material event to symptoms, alarms, operating state, inspection, work order, component, labour, parts, downtime, root cause and return-to-service evidence. Repeated text copied between work orders can overstate the apparent sample size. Planned replacements can be misclassified as failures. A repair may restore operation without resolving the underlying cause.

The reconstruction should distinguish failure, functional degradation, quality drift, safety intervention, planned overhaul and opportunistic work during a shutdown. These categories have different economics and different lead times. An early warning that arrives two hours before a failure may have little value if the necessary part has a twelve-week lead time. A warning several weeks ahead can support procurement and production planning.

Maintenance records should be reconciled to financial and production evidence. Contractor invoices, parts issues, overtime, line logs, rejected product and customer delivery records can corroborate the event. The objective is a defensible history for selected high-value use cases, not a cosmetic cleanup of every legacy record.

10. Define the maintenance decision

Each use case needs a precise decision, decision owner, lead time, action set and cost of error. Examples include whether to inspect a bearing at the next planned stop, whether to replace a spindle before a customer campaign, whether to rebalance a compressor duty cycle, or whether to hold a critical spare. A model output that does not change an operating decision cannot support an incremental value claim.

The decision window should match operational reality. A remaining-useful-life estimate may appear accurate on average while being too unstable for planning. A simpler anomaly score may be more useful if it reliably identifies which assets require engineering review. The threshold should reflect the cost of missed failure, unnecessary intervention, lost production and safety exposure.

The buyer should document the existing baseline. If experienced engineers already detect most failures through routine inspection, a new model may shift timing without generating material cash. If the target relies on reactive work and fragmented records, a governed decision process may create greater benefit. Value is measured against the credible counterfactual, not against doing nothing forever.

11. Evaluate technical performance

Performance should be measured on the intended asset population using frozen test data. Suitable metrics depend on the decision and can include precision, recall, false-negative rate, alert lead time, calibration, stability across loads and economic cost per decision. Aggregate accuracy is rarely sufficient. The test should show results by asset family, site, operating regime and failure mode.

Public datasets can help test methods and engineering pipelines. NASA's C-MAPSS and N-CMAPSS datasets support remaining-useful-life research for simulated turbofan degradation, while MIMII provides industrial-machine sound data for malfunction investigation [6-8]. These datasets do not validate performance in an acquired UK factory. Local equipment, sensors, environment, maintenance policy and failure prevalence determine deployment performance.

Independent validation should challenge leakage, labels, feature construction, thresholds, implementation and monitoring. Records from the same asset or event should not appear in both training and test populations. Repeated tuning against a nominal test set destroys its independence. Unknown equipment, missing channels and out-of-distribution conditions should trigger abstention and human review.

12. Protect operational technology

Predictive-maintenance integration increases connectivity between operational technology, enterprise systems and analytical environments. That connectivity can create additional attack paths and operational dependencies. The National Cyber Security Centre advises organisations to maintain a definitive view of operational-technology architecture, including asset criticality, exposure, availability constraints and the information held or shared [9-11].

The buyer should map gateways, remote access, vendor connections, historians, edge devices, cloud services, credentials and data-transfer paths. Active discovery can disrupt older industrial systems; collection methods should be approved, tested and appropriate to the equipment. Data extraction should use controlled media or interfaces, malware scanning, encryption, least privilege and auditable access.

The synergy case should include security cost. Network segmentation, hardened gateways, identity management, monitoring, backup, recovery tests and supplier assurance may be required before records can be pooled. A model that depends on continuous data from a fragile interface can reduce resilience. Integration design should preserve safe local operation when central analytics are unavailable.

13. Maintain safety and engineering accountability

The Health and Safety Executive states that work equipment should be maintained in an efficient state, efficient order and good repair; maintenance logs should be kept up to date where they exist; and maintenance operations should be carried out safely [12]. Predictive analytics can support prioritisation, but it does not replace statutory inspection, competent persons, manufacturer instructions, safe isolation or engineering judgement.

The diligence team should identify safety-critical assets and failure modes separately. A model optimised for production loss may underweight low-frequency events with severe safety or environmental consequences. Alerts must not encourage maintenance while equipment is energised or bypass established permit-to-work controls. The system should route safety-related observations to accountable specialists under approved procedures.

HSE research on emerging machinery technologies recognises potential applications for sensors and artificial intelligence in predictive maintenance while identifying safety implications that require careful control [13]. The transaction plan should therefore record intended use, prohibited use, evidence limitations, reviewer competence and escalation. The investment committee remains accountable for commercial decisions; engineers and responsible persons retain technical and safety duties.

14. Align with asset-management governance

ISO 55001:2024 frames asset management as a management system that aligns asset decisions with organisational objectives [14]. Predictive maintenance should operate inside that system. The combined group needs an asset-management policy, decision criteria, roles, lifecycle plans, risk controls, information requirements and performance review. A collection of isolated models does not create a durable integration capability.

Condition-monitoring standards can support a disciplined programme. ISO 17359 addresses general procedures for condition monitoring and diagnostics of machines; ISO 13374 addresses data processing, communication and presentation; ISO 14224 provides reliability and maintenance data principles for equipment in petroleum, petrochemical and natural-gas industries and can inform data structure beyond its core sector when applied carefully [15-17]. The buyer should select standards that match its equipment, sector and obligations.

Governance should define who approves an asset for monitoring, who changes thresholds, who owns an alert, who authorises intervention and who closes the feedback loop. Changes to sensors, firmware, maintenance policy or production duty can affect performance. The record should preserve configuration and decision history.

15. Separate accounting recognition from transaction value

Financial reporting and deal valuation answer different questions. IFRS 3 requires an acquirer to recognise identifiable acquired assets and liabilities at the acquisition date, subject to the standard's requirements. IAS 38 defines an identifiable intangible asset through separability or contractual or other legal rights and sets recognition and measurement principles [18-19]. Whether a particular dataset, software right or technology relationship qualifies for separate recognition requires transaction-specific accounting analysis.

A buyer can still attribute economic value to an integration capability even when no separately recognised asset appears in the acquired balance sheet. Conversely, accounting recognition does not prove that the data will generate incremental cash for this buyer. Purchase-price allocation, goodwill, useful life, impairment and disclosure should be handled by qualified accounting and valuation professionals using the applicable reporting framework.

For UK entities applying FRS 102, the Financial Reporting Council's business-combination materials describe the identification and fair-value measurement of acquired assets and liabilities under the relevant requirements [20]. The transaction model should maintain a clear boundary between the buyer's synergy analysis, any identifiable intangible valuation and the final accounting treatment.

16. Select the valuation approach

The primary approach is an incremental cash-flow model. It estimates the cash that the combined group can generate because the acquired and platform records are integrated, after deducting the costs required to deliver and sustain the capability. The model should include baseline failure cost, achievable operating change, ramp time, intervention cost, system and sensor investment, data engineering, security, validation, training and continuing governance.

A relief-from-royalty approach is usually difficult to support for internal maintenance data because comparable licences may not exist and the economic benefit depends on the buyer's operating context. A cost approach can estimate the expense of recreating records and systems, but historical collection cost does not establish economic value. A market approach may be impractical when transaction evidence does not isolate data from the wider business. International Valuation Standards emphasise the quality of data and inputs, model selection, documentation and the specific requirements for businesses, intangible assets and plant and equipment [21-23].

The cash-flow model should avoid capitalising the same benefit twice. Reduced downtime can increase output, lower overtime and improve customer service; these may describe one economic event. The model should identify the limiting mechanism and count only the incremental cash effect. Benefits already included in a broader integration case, capital expenditure plan or margin forecast should be removed or explicitly reconciled.

17. Build the synergy baseline

The baseline should represent credible performance without the proposed integration asset. It should use recent failure frequency, duration, production constraint, maintenance labour, contractor cost, parts usage, scrap and service outcomes. Abnormal shutdowns, temporary labour shortages or one-off supply disruptions should be treated separately. Management should explain which period is representative and why.

The baseline must distinguish recorded downtime from economic downtime. A machine can stop without reducing output if spare capacity absorbs the loss. A short failure at a bottleneck can affect an entire line. Contribution margin, inventory buffers, recovery production, customer penalties and overtime determine the cash consequence. Site-level accounting averages may obscure the relevant mechanism.

The buyer should test data completeness before extrapolation. Missing work orders or inconsistent downtime coding can bias the baseline. Cross-checks can use production logs, energy profiles, quality records, parts issues, contractor invoices and operator reports. Uncertainty should be expressed through ranges and sensitivities rather than concealed in a single point estimate.

18. Translate maintenance change into cash flow

Value mechanisms fall into four groups. Availability benefits arise when avoidable failures fall or repairs are completed faster. Maintenance benefits arise from better planning, fewer emergency callouts, improved labour deployment and less repetitive diagnosis. Working-capital benefits arise from rationalised critical spares and more predictable demand. Quality and customer benefits arise when equipment condition reduces scrap, rework or delivery disruption.

Table 2. Predictive-maintenance value mechanisms

Mechanism	Operating measure	Financial translation	Common double-counting risk
Availability	Avoided constrained downtime hours	Incremental contribution after recovery and variable cost	Counting output, overtime and service benefit for the same event
Maintenance planning	Emergency work, contractor callouts and overtime	Avoided incremental cash cost	Treating redeployed salaried labour as immediate cash saving
Repair effectiveness	Repeat failures and mean time to repair	Parts, labour, downtime and quality effect	Counting root-cause benefit in several asset programmes
Spares	Stockout risk, lead time and inventory position	Released cash less service-risk reserve	Treating inventory release as recurring earnings
Quality	Scrap, rework and process deviation	Material, labour and contribution effect	Attributing product or operator changes to the model
Customer service	Late orders and disrupted campaigns	Avoided penalties or protected contribution	Assuming every avoided delay creates new revenue

Each benefit requires a baseline, an operating action, evidence of causality and a financial owner.

Benefits should be classified as cost avoided, cash released, revenue protected or incremental revenue. The categories have different evidence and valuation treatment. Avoided future cost requires a credible forecast. Cash released from spares is generally one-time. Revenue protected needs evidence that the failure would have constrained saleable output. Incremental revenue also requires demand and capacity evidence.

The model should retain a ledger from operational metric to cash. For example: verified compressor alert; planned bearing change during scheduled downtime; avoided eight-hour bottleneck outage; constrained output of a stated quantity; contribution per unit; less accelerated part and intervention cost. This chain allows engineers, finance teams and investment committees to challenge the same claim.

19. Apply evidence and execution adjustments

The valuation should reflect evidence maturity without disguising judgement as precision. One method is to model separate central, downside and failure cases. Another is to apply transparent recognition factors to specific benefits based on coverage, technical validation, intervention capability and delivery ownership. The factor should not replace a scenario model where timing and outcomes are correlated.

Execution dependencies include sensor installation, interface completion, identity mapping, label quality, engineer availability, maintenance windows, spare lead times and operating adoption. Benefits can be delayed even when the model performs well. The cash-flow schedule should link each benefit to the earliest credible decision cycle and the required capital.

The buyer should compare the proposed programme with its prior integration record. Actual delivery from earlier acquisitions can inform ramp time, staffing and governance assumptions. The CMA's updated guidance gives greater evidentiary weight to ordinary-course operational and financial data, integration plans and track records when assessing merger efficiency claims [1]. The same discipline improves internal underwriting.

20. Demonstrate the hypothetical case

The illustrative transaction combines a UK manufacturing platform and a target. Management assumptions set platform revenue at GBP 210 million and EBITDA at GBP 26 million; target revenue at GBP 85 million and EBITDA at GBP 9 million. The target has 340 critical assets. Initial diligence finds stable identifiers for sixty-two percent, usable sensor histories for fifty-four percent, linked work and failure histories for forty-one percent and verified outcome labels for thirty-six percent. These are constructed assumptions used solely to demonstrate the method.

The first integration wave covers repeated equipment families that constrain output across several sites. Management assumes annual gross benefit of GBP 2.4 million from availability, GBP 1.1 million from maintenance labour and contractor efficiency, GBP 0.6 million from spare-parts and inventory actions, and GBP 0.8 million from quality and scrap improvement. The gross total is GBP 4.9 million. Recurring data, engineering, security, validation and governance cost is GBP 0.9 million, producing a management-assumed steady-state net pre-tax benefit of GBP 4.0 million.


One-time investment is assumed at GBP 5.8 million before completion of the programme. A further GBP 1.2 million is assumed in the first year for instrumentation and integration work. Net pre-tax benefits are assumed at GBP 0.8 million in year one, GBP 2.8 million in year two and GBP 4.0 million in years three to five. Discounting the cash flows at a management-assumed fourteen percent produces an illustrative five-year net present value of approximately GBP 3.2 million. The model excludes terminal value, tax, financing, working capital, purchase-price allocation and any value beyond year five.

Table 3. Hypothetical integration-asset valuation case

Item	Timing	Management assumption	Valuation treatment
Platform integration investment	Completion	5.8	Initial cash outflow
Additional instrumentation and integration	Year 1	1.2	Year-one cash outflow
Net operating benefit	Year 1	0.8	Discounted inflow
Net operating benefit	Year 2	2.8	Discounted inflow
Net operating benefit	Year 3	4.0	Discounted inflow
Net operating benefit	Year 4	4.0	Discounted inflow
Net operating benefit	Year 5	4.0	Discounted inflow
Illustrative discount rate	Five-year model	14 percent	Management assumption
Illustrative net present value	Five-year model	3.2	Buyer-specific value before excluded items

All figures are management assumptions in GBP millions and demonstrate the framework only.

Figure 3. Hypothetical steady-state annual benefit bridge



Management assumptions show how gross operating benefits become a net recurring pre-tax benefit after programme cost.

21. Set the price and negotiation boundary

Buyer-specific integration value does not automatically belong to the seller. The target's standalone value should reflect cash flows and assets that a market participant can support. Synergies that depend on the buyer's existing platform, engineers, data estate and integration investment may be specific to that buyer. The negotiation boundary should distinguish target-controlled evidence from value created by the buyer after closing.

The seller may support a higher price when it can demonstrate transferable records, validated use cases, reliable interfaces, durable rights and a delivery track record. Weak evidence should reduce the amount paid upfront. Parties can use contingent consideration or specific completion mechanics where future value depends on data delivery, interface performance or measurable operating outcomes, subject to legal, accounting and commercial advice.

The buyer should cap any premium by the lower of verified incremental value and the share it is willing to transfer through competition. The premium should also preserve a return on integration capital and a buffer for execution risk. A transaction model that pays the entire central-case net present value to the seller leaves the buyer exposed to downside with little compensation for delivery.

22. Design transaction protections

The sale agreement should translate material evidence gaps into precise obligations. These can include delivery of enumerated datasets and dictionaries, export tests, access credentials, licence consents, interface documentation, specified retention periods, cyber information, model artefacts, vendor contracts and named technical support. The schedules should use stable descriptions rather than broad promises to provide all relevant data.

Representations and warranties may address authority, rights, completeness of disclosed systems, known incidents, material changes, third-party restrictions and compliance. They should not promise model performance that has not been verified. Covenants can preserve records, prevent destructive migrations and maintain agreed interfaces between signing and completion. Conditions precedent can address consents or exports that are essential to Day One.

Where value depends on post-close cooperation, transitional services should specify service levels, security, change control, incident response, extraction, costs and exit. A holdback or contingent mechanism may be appropriate for a measurable delivery outcome, but it should avoid disputes over operating factors controlled by the buyer after completion. Legal counsel should align each protection with the transaction structure and governing law.

23. Govern clean-team and pre-close analysis

Pre-close data combination can raise confidentiality and competition concerns. The parties should define which information can be shared, for what purpose, with whom and under which controls. Competitively sensitive customer, pricing, capacity or strategy data may require clean-team arrangements, aggregation or delayed access. The scope should be set by competition counsel for the specific transaction.

The predictive-maintenance team should receive only the information required for diligence and integration planning. Asset signals and failure histories can reveal production volumes, product campaigns, customer programmes and capacity constraints. Access controls, purpose limitation, secure workspaces, output review and deletion procedures can reduce leakage. The team should avoid operational coordination before lawful completion.

The analysis should keep diligence, valuation and integration outputs separate where needed. A technical finding can inform the investment committee without disclosing unnecessary raw data. The post-close team can receive the detailed implementation package after completion under approved governance.

24. Establish decision rights and controls

The combined group needs named accountability across engineering, maintenance, operations, finance, information security, data protection, legal and valuation. Model owners cannot approve the commercial benefit of their own system. Finance should challenge the cash mechanism; engineering should challenge causality and safe action; information-security teams should challenge connectivity; operating leaders should own adoption.

Table 4. Integration governance and decision rights

Decision	Evidence owner	Independent challenge	Accountable approver	Required record
Data transfer and use	Data owner and legal team	Privacy and security review	Transaction steering committee	Rights, purpose, controls and retention
Asset and history mapping	Reliability data lead	Plant engineering review	Integration asset owner	Crosswalk, exceptions and provenance
Model release	Model owner	Independent validation	AI and engineering governance owner	Test report, limits and monitoring plan
Maintenance intervention	Site maintenance lead	Discipline engineer	Plant manager or delegated authority	Alert, evidence, action and outcome
Synergy recognition	Integration finance lead	Deal finance and technical advisers	Synergy governance committee	Baseline, bridge and supporting evidence
Valuation consequence	Deal team	Accounting and valuation specialists	Investment committee	Model, sensitivities and negotiation treatment

Roles are illustrative and should be adapted to the buyer, target, advisers and transaction documents.

Controls should include model inventory, versioning, approved uses, prohibited uses, performance thresholds, change management, override review and stop conditions. Human review must be substantive.

A reviewer should see the relevant source history, operating context, model output, uncertainty and recommended action. The record should capture the decision and subsequent outcome.

NIST's Artificial Intelligence Risk Management Framework organises controls through govern, map, measure and manage; OECD principles emphasise human-centred values, transparency, robustness, security and accountability [24-25]. These frameworks can support governance while UK legal and sector requirements remain determinative.


25. Build the first hundred day plan

The first thirty days should secure evidence and operating continuity. The team confirms asset and system ownership, preserves histories, freezes destructive migrations, validates backups, establishes access control and agrees the priority asset families. It also reconciles the synergy baseline and names owners for each financial mechanism. Safety and production controls remain unchanged unless approved through existing procedures.

Days thirty-one to sixty should build the minimum shared data model. The team maps identifiers, units, events, work orders and outcomes for selected use cases; tests extracts and interfaces; documents missingness; and completes cyber and privacy assessments. Engineers define decision windows, intervention options and stop conditions. Finance builds the evidence ledger from operating action to cash.

Days sixty-one to one hundred should run controlled pilots and approve scaling gates. Each pilot uses a frozen test population, human review and explicit baseline. The group measures alert quality, lead time, intervention rate, realised operating effect and process cost. Results are challenged before any synergy is recognised. The next tranche of instrumentation or integration capital is released only when the evidence meets the agreed gate.

Figure 4. First hundred day predictive-maintenance integration roadmap



The roadmap moves from evidence preservation to controlled use-case release and independently challenged value recognition.

26. Track value after close

Synergy tracking should start with the operating evidence, not a top-down budget. Each recognised benefit should have a baseline period, asset population, intervention, owner, financial mechanism, measurement frequency and independent review. Finance should reconcile operational measures to general-ledger cash effects where feasible and retain any bridge required for timing or allocation.

The group should separate gross improvement from integration cost and business-as-usual change. Production mix, pricing, energy cost, maintenance policy, capital investment and workforce changes can affect the same outcomes. A controlled comparison, matched asset cohort or phased release can improve

attribution. Where causality remains weak, the benefit should remain unrecognised or be reported as an operating indicator rather than cash synergy.

The investment committee should receive a quarterly view of evidence readiness, model performance, interventions, cash realised, investment spent, unresolved risks and lessons. Failed use cases should be closed or redesigned. Continuing spend requires evidence that the next tranche can improve a material decision.

27. Recognise limitations

Historical records can be incomplete, inconsistently labelled or biased toward failures that were documented. Sensor populations can change after equipment modifications, firmware updates, maintenance or acquisition. Public datasets do not represent the target factories. Model performance can deteriorate when operating regimes, products, maintenance policies or failure prevalence change.

Economic attribution also has limits. Avoided failures are counterfactual; an alert followed by maintenance does not prove the failure would have occurred. Benefits can overlap with broader reliability, capital and integration programmes. Salaried labour released from emergency work may create capacity without immediate cash saving. Inventory reduction can increase service risk if criticality and lead time are misunderstood.

The five-year case is a management-assumed illustration. It excludes terminal value, tax, financing, working capital, transaction costs outside the stated integration investment, accounting treatment and negotiated sharing. The framework does not provide a fairness opinion, purchase-price allocation, engineering certification or legal conclusion. Transaction-specific specialists should verify each material input.

28. Draw the investment conclusion

Predictive-maintenance data can support value in a UK manufacturing roll-up when the transaction creates a governed evidence base that improves material maintenance decisions across a transferable asset population. The buyer should begin with asset identity, operating context, verified failures, interventions and outcomes. Sensor volume and model sophistication follow that foundation.

The valuation should use incremental buyer-specific cash flows after integration, security, validation and continuing operating costs. It should preserve a clear baseline, prevent double counting, reflect ramp time and execution dependencies, and distinguish target-controlled evidence from value created by the buyer. The price should reflect what transfers at completion and leave the buyer a return for integration capital and risk.

The most defensible transaction case links five records: a controlled asset population, traceable operating evidence, an independently tested decision model, a human-approved maintenance action and a verified cash outcome. That chain turns data from a diligence narrative into an integration asset that can be governed, valued and tested after close.

Frequently asked questions

When does predictive-maintenance data become an integration asset?

It becomes an integration asset when the buyer controls the necessary rights and records, can link them to identifiable assets and operating states, can improve a defined maintenance decision, and can generate incremental cash after implementation and continuing costs.

Should a buyer assign a standalone value to raw sensor data?

Usually the stronger analysis values the cash-generating capability rather than raw volume. Raw data may have little value without identity, context, outcomes, rights, engineering knowledge and an operating process that acts on the evidence.

How should the buyer avoid double counting synergies?

Maintain an evidence ledger from operating event to cash. Count the limiting financial mechanism once, reconcile overlap with broader integration initiatives, and separate recurring earnings, one-time cash release and revenue protection.

Can public predictive-maintenance datasets support transaction valuation?

They can help test methods and pipelines. They do not establish performance on the acquired factories. Local equipment, sensors, operating regimes, labels and failure prevalence require independent validation.

What evidence should be secured before completion?

Secure asset hierarchies, signal dictionaries, failure and work histories, interfaces, licences, data rights, vendor contracts, model documentation, security information, backups and a tested export for the priority use cases.

How should synergy value affect the purchase price?

Separate target-controlled transferable evidence from buyer-specific value. Cap any premium by verified incremental value and preserve a return for integration capital, execution risk and uncertainty. Transaction-specific valuation advice remains necessary.

What should happen during the first hundred days?

Preserve records and continuity, confirm rights and the asset perimeter, build a controlled crosswalk, reconstruct priority histories, validate models and controls, run bounded pilots, verify operating outcomes and approve further investment through evidence gates.

Who remains accountable for maintenance and investment decisions?

Named human decision-makers remain accountable. Engineers interpret condition and safety, operating leaders authorise interventions, finance validates cash effects, legal and security teams govern permitted use, and the investment committee approves valuation and capital decisions.

References

- [1] Competition and Markets Authority, Merger Assessment Guidelines updated September 2026 <https://www.gov.uk/government/publications/merger-assessment-guidelines/merger-assessment-guidelines-html-version>
- [2] Competition and Markets Authority, Merger Assessment Guidelines publication page <https://www.gov.uk/government/publications/merger-assessment-guidelines>
- [3] Competition and Markets Authority, Merging Your Business guidance <https://www.gov.uk/guidance/merging-your-business>
- [4] Information Commissioner's Office, Due Diligence When Sharing Data Following Mergers and Acquisitions <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/data-sharing/data-sharing-a-code-of-practice/due-diligence/>
- [5] UK Parliament, Data Use and Access Act 2025 <https://www.legislation.gov.uk/ukpga/2025/18/contents>
- [6] NASA Prognostics Center of Excellence, C-MAPSS and prognostics datasets <https://www.nasa.gov/content/prognostics-center-of-excellence-data-set-repository>
- [7] IEEE Industrial AI Laboratory, Predictive Maintenance Benchmark Datasets <https://ieee-ies-industrial-ai-lab.github.io/industrial-ai-hub/datasets/>
- [8] Purohit et al., MIMII Dataset Sound Dataset for Malfunctioning Industrial Machine Investigation and Inspection <https://arxiv.org/abs/1909.09347>

- [9] National Cyber Security Centre, Define Processes for Maintaining the Definitive OT Record <https://www.ncsc.gov.uk/collection/operational-technology/definitive-architecture-view/principle-1>
- [10] National Cyber Security Centre, Establish an OT Information Security Management Programme <https://www.ncsc.gov.uk/collection/operational-technology/definitive-architecture-view/principle-2>
- [11] National Cyber Security Centre, Identify and Categorise OT Assets <https://www.ncsc.gov.uk/collection/operational-technology/definitive-architecture-view/principle-3>
- [12] Health and Safety Executive, Maintenance of Work Equipment <https://www.hse.gov.uk/work-equipment-machinery/maintenance.htm>
- [13] Health and Safety Executive, Review of Machinery Safety Standards and Emerging Technologies RR1214 <https://www.hse.gov.uk/Research/assets/docs/rr1214.pdf>
- [14] International Organization for Standardization, ISO 55001 Asset Management System Requirements 2024 <https://www.iso.org/standard/83054.html>
- [15] International Organization for Standardization, ISO 17359 Condition Monitoring and Diagnostics of Machines <https://www.iso.org/standard/71194.html>
- [16] International Organization for Standardization, ISO 13374 Condition Monitoring and Diagnostics of Machines <https://www.iso.org/standard/21832.html>
- [17] International Organization for Standardization, ISO 14224 Collection and Exchange of Reliability and Maintenance Data <https://www.iso.org/standard/72764.html>
- [18] IFRS Foundation, IFRS 3 Business Combinations <https://www.ifrs.org/issued-standards/list-of-standards/ifrs-3-business-combinations/>
- [19] IFRS Foundation, IAS 38 Intangible Assets <https://www.ifrs.org/issued-standards/list-of-standards/ias-38-intangible-assets/>
- [20] Financial Reporting Council, FRS 102 Factsheet 6 Business Combinations https://media.frc.org.uk/documents/Factsheet_6_-_Business_combinations.pdf
- [21] International Valuation Standards Council, International Valuation Standards <https://ivsc.org/standards/>
- [22] International Valuation Standards Council, 2025 Standards Update and Data Inputs <https://ivsc.org/new-edition-of-the-international-valuation-standards-ivs-published/>
- [23] International Valuation Standards Council, Deciphering Technology and Technology Valuation <https://ivsc.org/perspectives-paper-deciphering-technology/>
- [24] National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework 1.0 <https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-ai-rmf-10>
- [25] OECD, OECD AI Principles updated 2024 <https://www.oecd.org/en/topics/ai-principles.html>
- [26] Office for National Statistics, Manufacturing and Production Industry <https://www.ons.gov.uk/businessindustryandtrade/manufacturingandproductionindustry>
- [27] Department for Business and Trade, Made Smarter Review <https://www.gov.uk/government/publications/made-smarter-review>
- [28] Department for Business and Trade, Factors Influencing Firms Adoption of Advanced Technologies <https://www.gov.uk/government/publications/barriers-and-enablers-to-advanced-technology-adoption-for-uk-businesses/factors-influencing-firms-adoption-of-advanced-technologies-a-rapid-evidence-review>

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 closure.

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 and leads 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.