The Cyber-Physical Integration Plan: Combining Plants, Networks and Operational Technology Safely

A transaction framework for identity, connectivity, control systems, vendors, safety and recovery across industrial integration

Authored by Chennakeshav Adya
Independent Researcher
15 September 2026

Abstract

An acquisition that combines industrial plants also combines physical processes, control systems, engineering workstations, identities, communications links, vendors and recovery dependencies. A conventional information-technology integration schedule can create unacceptable production and safety exposure when it treats operational technology as another set of applications to migrate.

This paper presents a transaction framework for integrating cyber-physical operations while preserving safe production. It begins with the operating process and minimum viable service, builds a definitive architecture and critical-asset map, separates identity and connectivity decisions, and sequences change through evidence-based integration waves. It also connects vendor access, configuration integrity, backup validation, incident command and recovery testing to transaction governance and value creation.

A wholly hypothetical illustration covers six plants, 1,460 inventoried cyber-physical assets and 184 critical dependencies. It prioritises 72 high-criticality assets, places 38 proposed connections behind additional engineering gates, and models four integration waves. Every company, amount, rate and outcome in the illustration is hypothetical. A live integration requires verified plant, safety, engineering, legal, regulatory, cyber security, insurance and financial evidence, together with qualified professional advice.

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

Keywords: operational technology, cyber-physical integration, industrial control systems, M&A, post-merger integration, cyber security, plant safety, network segmentation, recovery

1. Define the transaction decision

The board decision is whether the combined group can capture the intended operating and financial value without creating an unacceptable path to physical harm, environmental loss, prolonged outage or loss of product quality. The integration plan should identify which systems must be combined, which should remain separated, and which decisions need more evidence. It should also state the production, safety and recovery conditions that govern timing.

Operational technology includes programmable systems and devices that monitor or change physical processes. NIST identifies industrial control systems, building automation, transportation systems and physical access systems among its examples. It also emphasises performance, reliability and safety requirements that differ from ordinary corporate computing. [1] These characteristics mean that a plant integration must start from process consequences and engineered safeguards.

The acquisition case should distinguish value sources. Some benefits depend on shared procurement or consolidated reporting and may require limited OT change. Others depend on common production planning,

remote operations, data platforms or plant optimisation and may require new connectivity. Each benefit should therefore have a technical dependency, risk owner, implementation cost and earliest safe date.

The investment committee should approve a change perimeter. It can authorise low-risk corporate integration while reserving production-network changes for site-level engineering approval. A gated perimeter prevents transaction momentum from becoming an implicit instruction to connect systems before the combined group understands them.

The final decision paper should answer five questions. Which physical services must continue? Which cyber assets and dependencies support them? Which proposed changes alter safety or recovery? What evidence permits each wave? Which value benefits remain available when a connection or migration is delayed?

2. Begin with the physical process

A cyber-physical integration map begins with material, energy, water, movement or another physical process. It traces the process from inputs through control points to safe output and shutdown. The team should identify hazards, quality constraints, environmental limits, operator interventions and the minimum set of equipment needed to sustain a critical service.

Process engineers should describe normal, degraded and emergency states. A control signal that appears minor in a network diagram may command a valve, breaker, drive, robot or dosing system. The consequence of an unavailable historian differs from the consequence of an unavailable safety instrumented function. Asset criticality should therefore reflect physical effect and recovery dependency.

The joint international Principles of Operational Technology Cyber Security state that safety is paramount and that knowledge of the business is crucial. The guidance connects cyber decisions to human life, plant, environment, reliability and essential services. It also asks whether security tools, backups and recovery processes behave predictably in the physical environment. [2]

Integration workshops should include operations, engineering, process safety, maintenance, cyber security, information technology, quality, legal and business continuity. Each function sees a different part of the system. Operators understand workarounds and local dependencies; engineers understand logic and fail states; cyber teams understand trust and exposure; finance understands the value case and outage cost.

The process map becomes the governing reference for later technical decisions. It defines which connections can be interrupted, which tests require a planned shutdown, and which assets need independent operation when corporate services are unavailable.

3. Establish a definitive architecture

The combined group needs a current, trusted view of both estates before changing connectivity. The architecture should include sites, process areas, control zones, safety systems, servers, engineering workstations, controllers, network devices, wireless links, remote access, vendor services, cloud connections and data flows. It should preserve enough detail to support decisions while protecting sensitive engineering information.

The UK National Cyber Security Centre recommends defined processes for maintaining a definitive OT architecture, an information-security management programme, categorised assets, documented connectivity and documented third-party risk. [3] A transaction adds urgency because ownership, contracts, personnel and support routes can change while the technical environment remains in service.

The architecture should state confidence. Verified physical inspection, controller exports and passive network observation provide stronger evidence than an inherited spreadsheet. Conflicts should remain visible until resolved. A diagram dated before a plant expansion may still be useful, but it cannot support a migration gate without validation.

Engineering configuration data deserves restricted handling. The joint OT principles describe its enduring operational and adversarial value. [2] The integration data room should separate board-level dependency views from detailed logic, addressing and configuration material. Access should follow role, purpose and time limits.

A definitive architecture is a maintained decision asset. Every approved integration change should update it. A temporary connection, vendor tunnel or dual-run interface should have an owner, expiry condition and removal evidence.

Figure 1. Proposed cyber-physical integration architecture

Integrate through controlled conduits while preserving physical safeguards

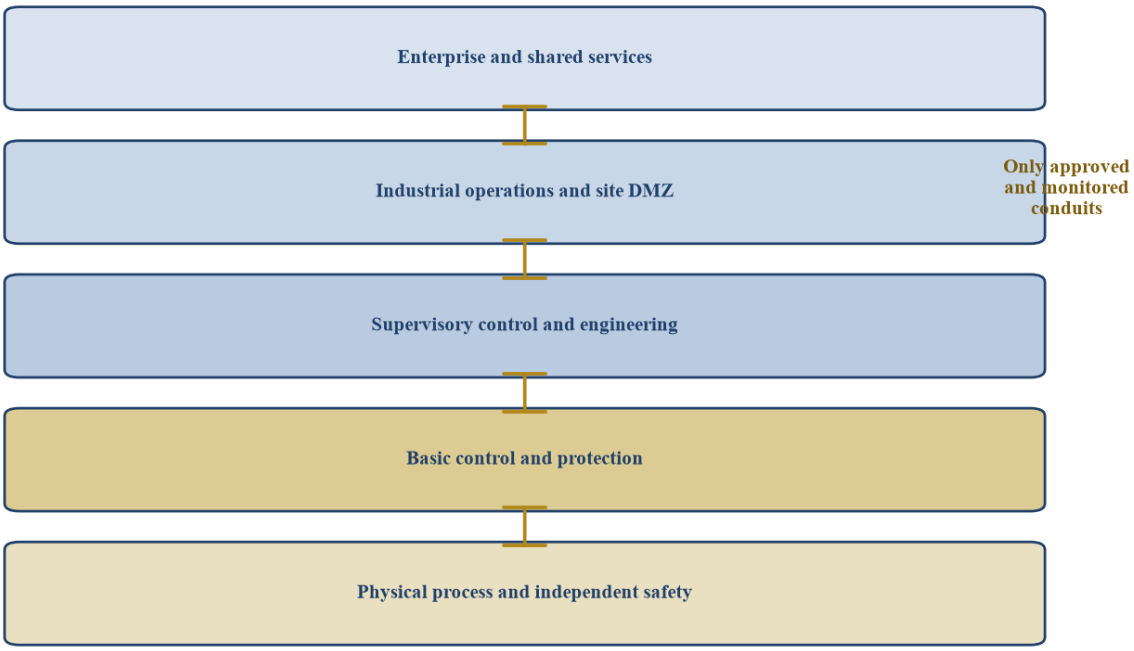

Original framework. Zones, controls and physical dependencies require site-specific engineering verification.

4. Build the asset and dependency register

The register should connect each physical function to the cyber assets, people, utilities and suppliers required to operate and recover it. Hardware inventory alone is insufficient. Firmware, logic, licences, certificates, accounts, time sources, naming services, engineering files, spare parts, communications circuits and support agreements can all become critical dependencies.

CISA's Cross-Sector Cybersecurity Performance Goals recommend a regularly updated inventory of assets with an Internet Protocol address, including OT, and place the practice within a broader set of prioritised outcomes for critical infrastructure. [4] The 2025 joint Foundations for OT Cybersecurity guidance develops an asset-inventory approach for owners and operators. [5]

The transaction team should use passive and engineering-approved discovery methods. Active scanning can disrupt fragile devices or communications. Site personnel should approve tools, rates, timing and rollback. The register should record how each asset was discovered and when it was last verified.

Criticality needs several dimensions. Safety criticality captures potential harm. Production criticality captures output and outage consequence. Cyber criticality captures privilege and reach. Recovery criticality captures whether the asset or record is required to restore others. A modest engineering workstation may score highly because it contains the only current configuration and trusted programming route.

Dependency mapping should identify common-mode failure. Two plants may appear independent while relying on one remote-access platform, telecom provider, directory, cloud licence server or specialist engineer.

Acquisition integration can increase this concentration when local arrangements are replaced by one shared service.

Table 1. Proposed critical-asset and dependency register

Field	Required evidence	Decision use	Typical exception
Physical function	Process diagram, operating procedure and hazard analysis	Establish consequence	Asset has no documented process owner
Cyber asset	Passive observation, configuration and site validation	Establish operating role	Duplicate, dormant or undiscovered device
Connectivity	Network flow, firewall rule and circuit record	Establish trust path	Temporary route became permanent
Identity	Named user, service account, credential store and privilege	Establish control authority	Shared or vendor-owned credential
Recovery dependency	Backup, configuration, spare, licence and test record	Establish restoration path	Backup exists without a proven restore
Supplier dependency	Contract, support route, personnel and remote access	Establish continuity	Service cannot transfer at completion

Original framework. Ratings and evidence thresholds require plant-specific approval.

5. Reconcile identity and authority

Identity integration can change who has authority over physical processes. The team should inventory human users, service accounts, local controller credentials, shared operator accounts, certificates, keys, vendor identities and emergency access. It should map each identity to a person or system, privilege, approval path, authentication method and revocation route.

Corporate directory consolidation may be desirable, but some plants depend on local operation during wide-area or enterprise outages. The target state should preserve necessary local resilience. Authentication changes should be tested against shift work, emergency response, disconnected operation and vendor support.

Privileged access requires special treatment. Engineering tools can download logic, change protection settings or alter recipes. Remote administration infrastructure can create a path across many sites. The integration should establish named access, strong authentication where technically feasible, approved jump paths, session logging and time-bounded elevation while respecting safety and vendor constraints.

Completion creates a personnel event. Departing employees, seller administrators, integrators and vendors may retain knowledge or credentials. The closing plan should identify which access transfers, which is revoked, which must continue under a transition service, and who can authorise emergency restoration.

Service accounts and certificates often outlive documentation. Changing them without dependency testing can stop production. Leaving them unchanged can preserve unknown access. The decision should use staged discovery, monitored rotation, rollback and evidence that dependent devices and applications continue functioning.

6. Design zones and conduits

Segmentation creates boundaries that limit failure and make permitted communication explicit. The design should group assets by physical function, consequence and trust. It should then define conduits between zones, including source, destination, protocol, purpose, owner, monitoring and removal condition.

The joint OT principles recommend segmenting and segregating OT from other networks. [2] Recent Australian secure-connectivity guidance extends the concept to least privilege, inspection and controlled administration. [6] A transaction should use these principles to evaluate proposed connections rather than assume that every shared service belongs in every plant.

The site demilitarised zone can mediate selected data and administrative flows between enterprise and plant environments. Historians, patch staging, remote access, file transfer and monitoring need deliberate placement. The design should prevent a compromised corporate identity or management platform from automatically controlling critical OT.

Legacy protocols may lack authentication or encryption. Segmentation, unidirectional transfer, application proxies, strict allowlists and physical procedures can reduce exposure when replacement is impractical. Compensating controls should have owners and review dates.

Every new conduit should have a value dependency. If a proposed data link supports a synergy, the business case should identify the benefit. If the benefit does not justify the residual risk and control cost, the integration plan should preserve separation.

7. Separate data integration from control integration

The combined group can often obtain reporting and planning benefits by moving selected data outward without creating an inward control path. The architecture should distinguish telemetry export, historian replication, analytical access, remote observation, remote support and remote control. Each category carries a different consequence.

Data integration should define content, frequency, destination, retention and authority. Engineering configuration, alarm data, production data and quality records have different confidentiality and operational significance. The data pipeline should avoid exposing detailed control information beyond the people and systems that need it.

Control integration changes command authority. A central operating centre, optimisation engine or maintenance platform may improve performance, but it can also create common-mode dependence. The design should identify safe manual fallback, local override, communications-loss behaviour and the conditions under which a central command is rejected.

Analytics and artificial intelligence require additional controls when they influence OT. The 2025 international principles for secure AI integration in OT recommend clear business cases, safety and security risk management, controlled data, secure deployment and continuing monitoring. [7] An integration should keep advisory models separate from direct control until the plant has verified behaviour, boundaries and failure modes.

The transaction model should assign benefits to the least invasive technical route that supports them. A read-only, delayed data feed may deliver fleet benchmarking without the risk of immediate control convergence.

8. Govern vendors and the supply chain

Industrial estates depend on original equipment manufacturers, systems integrators, maintenance contractors, telecom providers, cloud services and specialist engineers. The buyer should map which suppliers can access which systems, who owns credentials and tools, and whether contracts transfer at completion.

The joint OT principles place supply-chain security among six core principles and emphasise that a small supplier can still create material exposure. [2] Vendor criticality should therefore reflect access and dependency rather than contract value alone.

The diligence team should review remote-support methods, background checks where applicable, incident notification, subcontracting, secure development, vulnerability handling, product support life, software provenance, recovery assistance and exit rights. Evidence should include actual connection records and support practices.

Transition services can preserve continuity when the seller owns a support platform or licence. Each service needs a technical boundary, permitted users, data rights, security obligations, incident coordination, change control, duration and tested exit. An agreement that describes a service commercially without its cyber-physical dependencies leaves operational uncertainty.

Procurement integration should avoid forced standardisation before engineering review. A group-wide endpoint, network or identity product may be unsuitable for a legacy controller, safety system or isolated site. Product selection should account for plant certification, deterministic behaviour, support and rollback.

9. Protect engineering configuration and change control

Control logic, set points, protection settings, recipes, firmware, drawings and system configurations define plant behaviour. The integration should establish authoritative versions, approved storage, change history, cryptographic integrity where feasible and a controlled path from engineering development to production.

The buyer should identify undocumented changes and reconcile running configurations with approved baselines. A file in a document repository may not match the controller. Comparison tools and vendor methods can support verification, subject to engineering approval.

Change control should record purpose, affected physical function, test evidence, safety review, cyber review, implementation window, backout condition and responsible engineer. Emergency changes require retrospective evidence and independent review. Group integration deadlines should never replace site authority over hazardous work.

Configuration repositories should remain usable during enterprise outages. They also need protection because they can enable precise manipulation of the process. The design should combine restricted access, offline or isolated copies, integrity checks and recovery instructions.

Acquisition teams should include configuration quality in value assessment. Missing source files, unsupported controllers or sole-person knowledge can require remediation capital and extend synergy timing. These issues belong in the transaction model and completion plan.

10. Establish the safety case for each wave

Each integration wave should have a concise safety and operational case. It should describe the change, affected hazards, safeguards, test method, production window, personnel, communications, rollback and acceptance evidence. The site owner and relevant engineering authority should approve it.

The case should account for interactions. A new identity platform may affect operator login, engineering access, alarm acknowledgement and vendor support. A network change may alter latency, time synchronisation or failover. A backup agent may consume bandwidth or processing capacity. The review should follow the physical consequences through the architecture.

Management should define forbidden combinations. Two redundant assets should not be changed in the same window without explicit engineering justification. A safety system and its monitoring path may require separate change control. Central and local recovery routes should not be disabled together.

The plan should use observable entry and exit criteria. Entry can require verified backups, spares, test environment, staffing and an approved rollback. Exit can require stable process conditions, alarm review, configuration comparison, security monitoring and operator sign-off.

The board should receive aggregated risk and value information. Site-level engineering evidence should remain available for audit and incident learning. This arrangement preserves accountable oversight without asking directors to approve technical details beyond their competence.

11. Sequence integration waves

Wave design should follow dependency and risk. A common sequence begins with governance and observation, then introduces low-risk shared services, controlled data exchange, selected identity changes and finally justified control or platform convergence. Sites with weak evidence may remain in discovery while others proceed.

The first wave should establish owners, architecture, asset registers, incident contacts and monitoring. It can also address exposed access and unsupported transition dependencies where safe. The objective is to improve visibility and control without broad production change.

The second wave can connect selected data through designed conduits and test central reporting. The third can address identity, vendor access and supported infrastructure where local resilience is proven. The fourth can implement plant-system convergence tied to a specific value case.

Waves should be reversible. Dual operation, preserved configurations and tested rollback can reduce the consequence of error. Reversibility has a cost, including temporary licences, people and infrastructure, and should be included in the integration budget.

The schedule should respect outages, seasonal demand, regulatory inspections, customer commitments and maintenance windows. Financial synergy timing should follow the verified engineering schedule. Pressure to meet an accounting period does not change physical risk.

Figure 2. Proposed integration waves and evidence gates

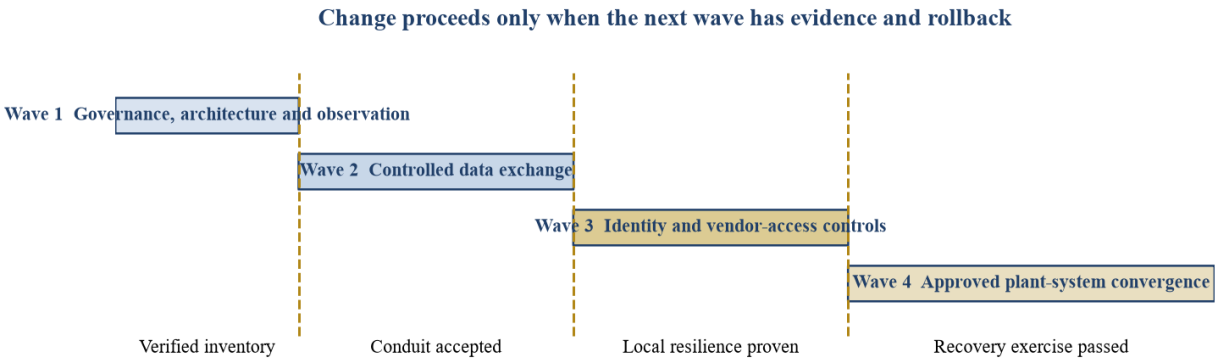

Original framework. Timing is illustrative and requires transaction-specific planning.

12. Build recovery before convergence

Recovery should be designed before the integration creates new dependencies. The team should identify the minimum equipment, configurations, software, licences, people, communications and utilities needed to restore each critical physical service. It should also define safe shutdown and degraded operation.

NIST SP 1339 provides an OT backup quick-start guide and focuses on planning, implementing, testing and maintaining backup capabilities for operational environments. [8] A transaction should apply this discipline to both inherited estates and the proposed combined architecture.

Backup existence does not prove recoverability. The team should test restoration to representative hardware or an approved environment, verify configuration integrity, confirm firmware and licences, and demonstrate operator and engineer procedures. Safety validation should follow restoration before production resumes.

Recovery dependencies can cross ownership boundaries. The seller may retain a domain, cloud tenant, telecom contract, code repository or specialist. These dependencies need transfer, replacement or a controlled transition service. The plan should test recovery under the expected post-completion ownership model.

The recovery strategy should assume that enterprise systems may be unavailable. Plants need approved local procedures, contact information and access to essential records. The degree of independence depends on process and risk, and should be established through engineering analysis.

13. Exercise incident command

The combined group needs one incident-command model that respects site safety authority. The plan should define who can stop production, isolate networks, engage vendors, notify regulators and customers, preserve evidence and authorise recovery. It should identify alternates for nights, weekends and cross-border operations.

Cyber containment can create physical consequences. Disconnecting a link may remove monitoring or support. Rebooting a device may change state. Incident responders should coordinate with operators and engineers before actions that affect OT, except where emergency authority permits immediate protective action.

NIST CSF 2.0 organises outcomes across Govern, Identify, Protect, Detect, Respond and Recover and adds explicit emphasis on governance. [9] The integration programme can use these outcomes to align corporate and site responsibilities while retaining process-specific controls.

Exercises should cover loss of corporate identity, compromised vendor access, malicious configuration change, unavailable historian, ransomware in the enterprise, telecom failure and suspected manipulation of a safety-relevant system. Each exercise should test decisions and communications as well as technical restoration.

Evidence from exercises should change the integration plan. Missing contacts, unclear authority, inaccessible backups and untested manual procedures should become funded actions with owners and deadlines.

14. Connect regulation, disclosure and insurance

Legal and regulatory obligations depend on geography, sector, entity and incident. The transaction should map licences, critical-infrastructure obligations, data rules, safety reporting, environmental reporting, contractual notification and securities disclosure. Qualified counsel should determine applicability.

The EU NIS2 Directive establishes cyber risk-management and incident-reporting obligations for covered entities and addresses supply-chain security, business continuity and vulnerability handling. [10] The EU Cyber Resilience Act establishes requirements for products with digital elements and staged application dates. [11] These frameworks can affect manufacturers, operators and suppliers in a combined group.

The US Securities and Exchange Commission requires covered registrants to disclose material cybersecurity incidents and information about risk management, strategy and governance under its 2023 rules. [12] An acquisition can change the systems and processes that support materiality assessment and disclosure control.

Insurance diligence should identify notification requirements, exclusions, security representations, business-interruption measurement and forensic-provider conditions. The buyer should avoid changing a represented control without understanding policy consequences. Coverage and claim outcomes require policy-specific advice.

The regulatory map should influence wave design. A planned connection can create new data transfers, reporting dependencies or product obligations. The integration committee should review these effects before approval.

15. Quantify value, cost and risk

The value model should connect each synergy or capability to a defined technical change. Shared monitoring may reduce duplicated spend and improve detection. Central planning may improve utilisation. Consolidated procurement may reduce unit cost. Remote operations may improve specialist coverage. Each benefit needs a baseline, owner, timing and evidence.

The cost model should include discovery, engineering, licences, network changes, test environments, outages, dual running, vendor support, spares, training, compliance, recovery exercises and contingency. Legacy constraints can make apparently simple integration expensive.

Risk should be expressed through scenarios rather than a false single probability. Relevant consequences include lost output, damaged equipment, quality loss, environmental release, safety event, regulatory action, customer interruption and delayed synergy. The model should distinguish insured, recoverable and enduring effects.

Transaction structure can allocate known uncertainty. A purchase-price adjustment, escrow, warranty, indemnity, covenant or transition service may address a defined exposure, subject to legal advice and negotiation. The integration budget should remain separate from purchase accounting judgments.

The board should compare three paths: preserve separation, create controlled interoperability, or converge platforms. Each path should show value, cash cost, execution time, residual dependency and recovery evidence. The appropriate answer can differ by site and system.

16. Apply a hypothetical six-plant model

Consider a wholly hypothetical acquisition involving six industrial plants in three jurisdictions. The initial records list 1,280 cyber assets. Passive discovery, configuration review and site validation identify 1,460 active assets. The 180-asset difference consists of undocumented devices, duplicate records removed and assets added through recent projects.

The dependency analysis identifies 184 relationships that can interrupt a critical physical service. Seventy-two assets are classified as high criticality because of safety, production, cyber privilege or recovery function. Thirty-eight proposed connections require additional engineering evidence before approval. These values describe no actual organisation.

The central plan uses four waves across 180 days. Wave one costs a hypothetical USD 1.8 million for architecture, asset reconciliation, monitoring and urgent access controls. Wave two costs USD 2.4 million for designed conduits and data exchange. Wave three costs USD 3.1 million for identity, vendor access and supported infrastructure. Wave four costs USD 4.7 million for selected plant-system convergence and recovery exercises.

The gross annual benefit is hypothetically USD 8.6 million. Engineering review defers USD 2.1 million of benefit associated with premature control convergence. The risk-adjusted first-year benefit is therefore USD 3.4 million after timing and implementation effects, while the full run-rate depends on accepted gates. These figures are calculation examples rather than forecasts.

The model shows that delayed convergence can preserve value. Procurement, reporting and selected analytics benefits proceed while higher-consequence control changes wait for evidence. The board retains the option to abandon a connection whose cost or residual risk exceeds its benefit.

Figure 3. Hypothetical critical-asset and proposed-connection map

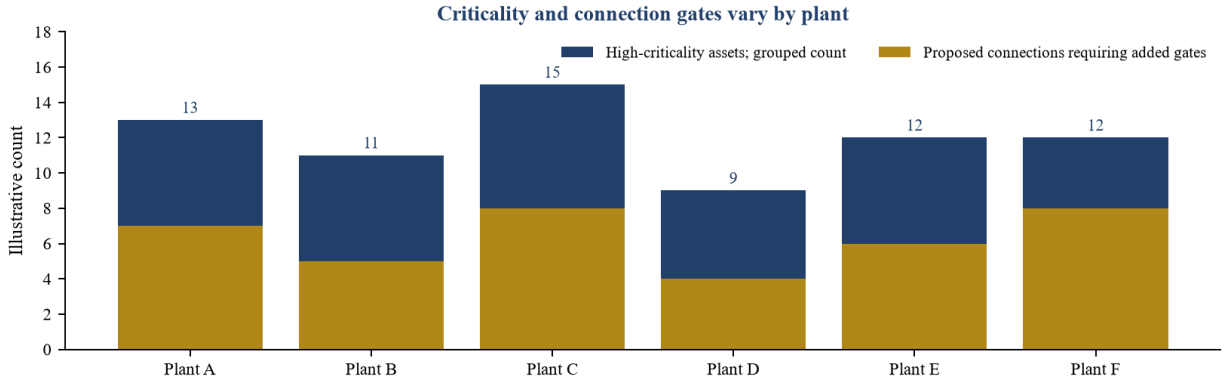

Original illustrative model. Counts describe no organisation.

Table 2. Hypothetical six-plant integration baseline

Measure	Initial record	Verified or proposed position	Decision implication
Inventoried cyber-physical assets	1,280	1,460	Reconcile before broad change
Critical dependencies	Not consistently recorded	184	Assign service and recovery owner
High-criticality assets	Not consistently classified	72	Apply enhanced engineering gates
Proposed new connections	64	38 require added evidence	Defer until safety and rollback are accepted
Four-wave implementation cost	Not applicable	USD 12.0 million	Fund discovery, controls, dual running and recovery
Hypothetical annual gross benefit	Not applicable	USD 8.6 million	Phase recognition with accepted technical gates

Original illustrative model. Counts and amounts describe no organisation.

17. Test the downside and break points

The central hypothetical case assumes four waves complete within 180 days and the accepted changes deliver USD 8.6 million of annual run-rate benefit. A delayed case extends the programme to 270 days, adds USD 1.6 million of dual-running and vendor cost, and defers USD 2.7 million of first-year benefit. A constrained case preserves separation for two plants and reduces annual run-rate benefit to USD 6.2 million.

The model should also test an adverse operational event. A five-day outage at one high-output plant with hypothetical contribution of USD 0.55 million per day creates USD 2.75 million of direct lost contribution before recovery cost, customer effects or insurance. This amount exceeds many individual integration savings.

Break-point analysis should identify the maximum implementation cost supported by each benefit, the outage duration that removes the first-year case, and the share of value available without control convergence. It should also test whether liquidity can fund remediation and dual running when integration is delayed.

These cases do not assign a probability to a cyber or safety event. They allow directors to understand consequence and option value. Live models need verified production economics, downtime costs, contractual exposure, insurance and site recovery estimates.

The approval should define stop conditions. Examples include an unverified recovery path, missing configuration source, unresolved safety dependency, unsupported controller, untransferable vendor service or failed rollback exercise. A stop condition protects the transaction thesis by preventing a single integration action from putting the broader asset at risk.

Table 3. Hypothetical value and downside cases

Case	Programme duration	Implementation and dual-run cost	Annual run-rate benefit	First-year interpretation
Central	180 days	USD 12.0 million	USD 8.6 million	Benefits follow accepted waves
Delayed	270 days	USD 13.6 million	USD 8.6 million	USD 2.7 million of benefit deferred
Constrained	240 days	USD 10.8 million	USD 6.2 million	Two plants remain separated

Case	Programme duration	Implementation and dual-run cost	Annual run-rate benefit	First-year interpretation
Five-day plant outage	Event case	Additional recovery cost unmodelled	USD 2.75 million direct contribution loss	Reassess value and control design

Original illustrative calculations. Amounts describe no organisation and are unsuitable as forecasts.

18. Translate findings into transaction terms

The diligence report should identify exposures that exist before signing, arise at completion or depend on integration. Pre-existing vulnerabilities, unsupported equipment and missing configurations may affect value and warranties. Completion access and transition services need closing mechanics. Future convergence belongs in the integration budget and governance.

Representations should use definitions that match available evidence. A broad claim that all systems are secure may be difficult to substantiate. Counsel can develop materiality, knowledge and disclosure structures based on the verified asset and incident record.

Covenants can preserve critical personnel, vendor support, insurance, backups and change control between signing and completion. The buyer may require notice of material incidents, plant changes and loss of key support. Regulatory and competition constraints may limit pre-completion coordination and information sharing; counsel should establish clean-team and permissible-planning rules.

Transition services should include service levels, security controls, incident cooperation, data handling, audit, change, termination and tested exit. The technical schedules should identify the systems, accounts, circuits and people behind each service.

Consideration mechanisms should avoid paying early for benefits that require uncertain convergence. A valuation range, deferred consideration or specific protection may be appropriate when a material value source depends on unresolved technical evidence. The legal and financial structure requires transaction-specific advice.

19. Organise the integration control room

The integration control room should combine a board-level decision view with site-owned technical execution. It needs a single register of physical services, assets, dependencies, changes, incidents, benefits, costs and evidence gates. Each item should have an owner and status.

The programme director should coordinate workstreams while site managers retain authority for safe operation. The cyber lead should govern architecture and access. Engineering authorities should approve changes. Finance should trace costs and benefits. Legal and regulatory advisers should manage applicable obligations. Internal audit or an independent reviewer can test evidence and process.

Weekly reviews should focus on exceptions and upcoming gates. The team should avoid using percentage-complete measures that hide unresolved critical dependencies. A wave can be 95 per cent complete while its single untested rollback makes deployment unacceptable.

Decision logs should record the evidence considered, responsible authority, residual risk, timing and review date. Temporary exceptions need expiry. Benefits should move from planned to realised only when finance can trace them and the supporting technical state is stable.

The control room should preserve learning across sites. A failed identity change, unexpected protocol dependency or recovery defect at one plant can alter the plan for the others. Lessons should be reviewed before the next wave proceeds.

20. Use 30, 90 and 180-day gates

By day 30, the combined group should have confirmed physical-service owners, incident authority, critical vendor continuity, exposed remote access and the architecture-discovery plan. It should freeze unapproved high-consequence changes and establish an evidence repository.

By day 90, it should have reconciled high-criticality assets and dependencies, designed zones and conduits, tested selected backups, exercised incident command and approved the first data-integration cases. Gaps should have funded remediation or explicit deferral.

By day 180, it should have completed accepted waves, tested recovery for critical services, removed or time-bounded temporary access, validated realised benefits and decided which systems remain separated. The board should receive the updated risk, value and option map.

Continuing governance should maintain the definitive architecture, asset inventory, access review, configuration integrity, supplier assurance, incident exercises and recovery tests. Acquisitions change ownership; the physical process and its hazards continue throughout.

The programme succeeds when the combined group can explain and demonstrate how each critical service operates, who can change it, how approved data crosses boundaries, how the service recovers and which value benefits depend on its integration state.

21. Design the evidence room and assurance plan

The integration evidence room should mirror the decision model. A physical-services folder should contain approved process boundaries, consequence classifications and responsible authorities. An architecture folder should contain dated diagrams, asset extracts, connectivity records and confidence assessments. Separate controlled areas should hold identities, configurations, vendor access, recovery evidence, incidents, regulatory analysis, benefits and costs.

Every extract should state its source system, extraction time, scope, owner and transformation. Where a device or configuration cannot be observed safely, the register should identify the limitation and the alternative evidence used. This discipline allows reviewers to distinguish verified plant state from inherited records and management plans.

Sampling should combine consequence, change and randomness. Review every asset that can affect an independent safety function, every proposed new conduit, every privileged remote-access route and every recovery dependency for a critical service. A random sample of lower-criticality assets can test whether the broader register is reliable. Exceptions should be traced to a cause and assessed for population-wide implications.

Technical assurance should include bidirectional tests. One test starts with the physical service and traces forward through equipment, control, network, identity and recovery. A second starts with a user, connection or configuration and traces backward to its physical consequence and authority. The two directions expose orphaned assets, unowned access and dependencies hidden by a purely system-centred review.

The evidence room should preserve decision versions. If a conduit is redesigned or a wave is deferred, the prior approval, reason and affected benefit should remain visible. This creates an audit trail for regulators, insurers, internal audit and post-incident review, subject to legal privilege and retention advice.

Independent review should focus on the highest-consequence assumptions. It can reperform asset reconciliation, inspect firewall and remote-access configuration, witness selected recovery tests, challenge benefit dependencies and examine whether site authority operated as designed. Independence and scope should be defined for the transaction.

22. Preserve workforce capability and operating knowledge

Cyber-physical resilience depends on people who understand plant behaviour, control logic, equipment history, vendor practices and safe recovery. An acquisition can destabilise this knowledge through departures, role changes, centralisation and changed contractor arrangements. The integration plan should treat critical knowledge as an operating dependency.

The team should identify who can diagnose failures, approve logic, restore configurations, operate manually, coordinate vendors and validate a safe return to service. It should distinguish formal role descriptions from demonstrated ability. Shift coverage, leave, location and contractual availability matter because an incident may occur outside normal hours.

Knowledge transfer should use evidence and practice. Current drawings, annotated procedures, configuration repositories and decision logs provide a base. Paired work, witnessed changes, simulated faults and recovery exercises show whether another qualified person can perform the task. A signed attendance sheet alone provides weak evidence of capability.

Retention arrangements should focus on genuinely critical roles and a defined transfer outcome. The buyer should consider employment, incentive, legal and cultural advice. Contractors and vendor specialists may require extended support or replacement plans when their agreements do not transfer.

Centralisation can remove local context. A group operations centre may improve coverage while still depending on site operators who recognise abnormal process behaviour. The operating model should state which decisions remain local, which can be escalated, and which can be made centrally under defined communications and plant conditions.

Training should include cyber security, process safety and the integration architecture. Corporate responders need to understand why familiar IT actions may be unsafe in OT. Operators need practical methods to recognise and report suspicious behaviour without losing focus on physical safety. Engineers need controlled routes for vendor access, configuration and evidence preservation.

Workforce evidence should feed the transaction model. Sole-person dependencies, scarce vendor skills and extensive undocumented knowledge can extend integration time and require retention or recruitment cost. These are measurable execution constraints and belong in value and liquidity planning.

Figure 4. Proposed recovery test plan

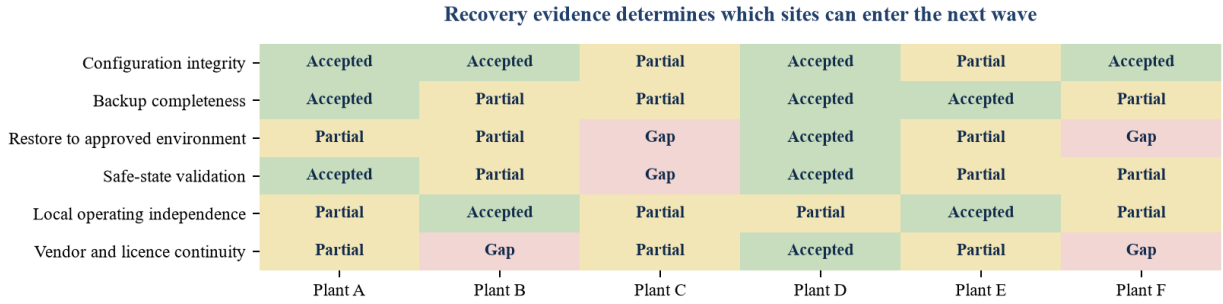

Original framework. Acceptance criteria require site-specific engineering and safety approval.

Table 4. Proposed 30, 90 and 180-day execution plan

Gate	Required evidence	Principal decision
Day 30	Service owners, incident authority, vendor continuity, exposed-access review and discovery plan	Stabilise ownership and stop unapproved change
Day 90	Critical-asset register, dependency map, conduit design, backup tests and incident exercise	Approve selected data and access changes

Gate	Required evidence	Principal decision
Day 180	Accepted integration waves, recovery evidence, temporary-access exit and traced benefits	Confirm convergence, continued separation or remediation
Continuing	Architecture maintenance, access review, supplier assurance, recovery tests and value tracking	Preserve safe operation and integration option value

Original framework. Timing and responsibility require adaptation to the transaction.

Frequently asked questions

Why does an industrial acquisition need a separate cyber-physical integration plan

Industrial systems directly monitor or change physical processes. Identity, network or software changes can affect safety, production, quality and recovery. The integration therefore needs process and engineering gates in addition to ordinary corporate technology planning.

Should every acquired plant join the buyer's corporate network

No universal answer applies. Each connection should have a defined purpose, architecture, control set, safety assessment and rollback. Some plants or systems may remain separated while selected data crosses through controlled conduits.

What belongs in a critical-asset map

The map should connect physical services to controllers, networks, engineering workstations, identities, configurations, utilities, communications, vendors, people, spares and recovery records. It should show consequence and evidence confidence.

How should vendor remote access be handled at completion

The buyer should identify every route, owner and credential; confirm contractual transfer; apply named, approved and time-bounded access where technically feasible; monitor sessions; and test revocation and emergency support.

When can identity systems be consolidated

Consolidation should follow dependency testing, local-resilience design, approved change control and rollback evidence. Plants may need local authentication during enterprise or communications outages.

What proves that an OT backup is usable

A usable backup has verified scope, integrity, required software and licences, accessible procedures and a successful restoration test in an approved environment. Safety and process validation should occur before production resumes.

How should integration synergies be valued

Each synergy should link to a technical change, baseline, owner, cost, timing and acceptance gate. The model should show benefits available through separation, controlled interoperability and convergence, together with downside and liquidity effects.

What should the board monitor after completion

The board should monitor critical-service continuity, unresolved dependencies, high-consequence changes, recovery-test results, vendor access, temporary exceptions, incidents, integration cost and benefits that finance can trace to accepted technical states.

References

[1] National Institute of Standards and Technology. Guide to Operational Technology Security, NIST SP 800-82 Revision 3. September 2023. Accessed 15 September 2026. https://doi.org/10.6028/NIST.SP.800-82r3

[2] Australian Signals Directorate's Australian Cyber Security Centre and international partners. Principles of Operational Technology Cyber Security. October 2024. Accessed 15 September 2026. https://www.cyber.gov.au/sites/default/files/2024-10/principles_of_operational_technology_cyber_security.pdf

[3] UK National Cyber Security Centre and international partners. Creating and Maintaining a Definitive View of Your OT Architecture. Version 1.0, reviewed 18 March 2024. Accessed 15 September 2026. https://www.ncsc.gov.uk/collection/operational-technology/definitive-architecture-view

[4] US Cybersecurity and Infrastructure Security Agency. Cross-Sector Cybersecurity Performance Goals. Accessed 15 September 2026. https://www.cisa.gov/cybersecurity-performance-goals

[5] US Cybersecurity and Infrastructure Security Agency and international partners. Foundations for OT Cybersecurity: Asset Inventory Guidance for Owners and Operators. August 2025. Accessed 15 September 2026. https://www.cisa.gov/sites/default/files/2025-08/joint-guide-foundations-for-OT-cybersecurity-asset-inventory-guidance_508c.pdf

[6] Australian Signals Directorate's Australian Cyber Security Centre and international partners. Secure Connectivity Principles for Operational Technology. January 2026. Accessed 15 September 2026. https://www.cyber.gov.au/business-government/secure-design/operational-technology-environments/secure-connectivity-principles-for-operational-technology

[7] Australian Signals Directorate's Australian Cyber Security Centre and international partners. Principles for the Secure Integration of Artificial Intelligence in Operational Technology. December 2025. Accessed 15 September 2026. https://www.cyber.gov.au/business-government/secure-design/operational-technology-environments/principles-for-the-secure-integration-of-artificial-intelligence-in-operational-technology

[8] National Institute of Standards and Technology. OT Backup Quick Start Guide, NIST SP 1339. June 2026. Accessed 15 September 2026. https://www.nccoe.nist.gov/publications/other/nist-sp-1339-ot-backup-quick-start-guide

[9] National Institute of Standards and Technology. The Cybersecurity Framework 2.0, NIST CSWP 29. February 2024. Accessed 15 September 2026. https://doi.org/10.6028/NIST.CSWP.29

[10] European Parliament and Council of the European Union. Directive (EU) 2022/2555 on Measures for a High Common Level of Cybersecurity Across the Union. 14 December 2022. Accessed 15 September 2026. https://eur-lex.europa.eu/eli/dir/2022/2555/oj

[11] European Parliament and Council of the European Union. Regulation (EU) 2024/2847 on Horizontal Cybersecurity Requirements for Products with Digital Elements. 23 October 2024. Accessed 15 September 2026. https://eur-lex.europa.eu/eli/reg/2024/2847/oj

[12] US Securities and Exchange Commission. Cybersecurity Risk Management, Strategy, Governance, and Incident Disclosure. Release 33-11216. 26 July 2023. Accessed 15 September 2026. https://www.sec.gov/rule-release/33-11216

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.

He is also 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.