INDUSTRIAL INTELLIGENCE PLATFORM

Success Criteria, ROI Framework, Technology Adoption, and Job Redefinition

A companion reference to The Industrial Intelligence Platform article | March 2026

This document defines the measurable success criteria for an industrial intelligence programme across four dimensions: financial implementation and ROI, new technology adoption, and job redefinition. Each section closes with a structured checklist that programme teams, sponsors, and reviewers can use at stage gates to assess readiness, track progress, and verify outcomes.

Section A: Financial Implementation and ROI Success Criteria

Financial success in industrial AI is not measured by technology deployment. It is measured by whether the platform delivers quantifiable improvements in margin, cost, throughput, and capital efficiency within agreed timeframes. The following criteria define what financial success looks like at each investment horizon.

A.1 Horizon One Returns (Months 0 to 18)

The first horizon is a proof-of-value period. Financial targets must be achievable on one or two pilot lines with a limited technology footprint. If the economics do not close at this scale, the programme should be redesigned before advancing.

Financial Metric	Minimum Target	Strong Target	Measurement Approach
Unplanned downtime reduction on targeted assets	10% reduction	25% reduction	Compare 90-day pre vs post deployment on same assets
First-pass yield improvement on pilot line	3% improvement	8% improvement	Daily quality records, pre vs post baseline
Scrap and rework cost reduction	5% reduction	15% reduction	Material waste log and rework labour hours
Maintenance cost per asset	Flat or declining	8% reduction	CMMS actual vs prior year same period
Payback period on Horizon One investment	Within 24 months	Within 15 months	Total deployment cost vs annualised operational savings
Platform ROI at end of Horizon One	Positive	Greater than 1.5x	Net present value of proven savings vs total cost

A.2 Horizon Two Returns (Months 18 to 42)

The second horizon shifts financial targets from cost reduction to a combined cost and revenue picture. Digital twin scenario planning, engineering compression, and AI-assisted scheduling all contribute to the top line as well as the cost base.

Financial Metric	Minimum Target	Strong Target	Measurement Approach
Throughput improvement on AI-optimised lines	8% improvement	20% improvement	Units per shift pre vs post, sustained over 90 days

Financial Metric	Minimum Target	Strong Target	Measurement Approach
Engineering iteration cycle time reduction	20% reduction	40% reduction	Design change to validated production release, tracked per project
Capex avoidance from simulation-first validation	5% of programme capex	15% of programme capex	Physical prototypes avoided, tooling rework prevented
Schedule reliability improvement	5% fewer missed delivery windows	15% improvement	On-time completion rate, rolling 13-week window
Platform cost per plant	Declining with each new site	40% lower than first site	Full deployment cost including integration labour
Net platform ROI across Horizon Two	Greater than 2x	Greater than 3.5x	Portfolio NPV across all deployed sites

A.3 Horizon Three Returns (Month 36 and Beyond)

The third horizon measures platform economics: the compounding returns that emerge when a reusable intelligence stack deploys across a growing portfolio with declining marginal cost per site.

Financial Metric	Minimum Target	Strong Target	Measurement Approach
Annual platform recurring revenue	Growing year on year	Greater than 20% annual growth	Licensing and managed service contract value
AI-native contract manufacturing margin	Above sector average	2x sector average	Gross margin per production unit vs comparable conventional plants
Multi-plant ROI vs single-site deployment	20% lower cost per site from site 3 onward	50% lower cost per site from site 5 onward	Deployment cost log across all sites
Engineering simulation service revenue	New revenue line established	Greater than 10% of total platform revenue	Contract value and utilisation of simulation capacity
Return on total capital deployed	Greater than 15% IRR	Greater than 25% IRR	Board-level portfolio financial model

> *Financial discipline is not a constraint on ambition. It is the mechanism by which ambition earns the right to scale. Every horizon gate requires evidence, not intention.*

Section B: New Technology Adoption Success Criteria

Technology adoption in a manufacturing environment is not just a technical deployment question. It is an operational readiness question. A model that runs without being used by the people closest to the process has not been adopted. A robot cell that requires constant manual intervention has not been adopted. The criteria below define what real adoption looks like across the platform layers.

B.1 Data Infrastructure Adoption

Adoption Criterion	Target	How to Assess
Critical asset sensors connected and streaming	Greater than 95% uptime on priority assets	Connectivity dashboard, data gap reports
Asset graph coverage across pilot scope	100% of in-scope assets registered and versioned	Asset graph completeness audit
Data latency for operational queries	Sub-3 seconds for common queries	Response time log across application layer
Data quality score on priority data streams	Greater than 90% completeness and accuracy	Automated data contract validation reports
IT and OT teams actively using platform data	Both teams using shared dashboards weekly	Login and query activity audit

B.2 AI Model and Analytics Adoption

Adoption Criterion	Target	How to Assess
Predictive maintenance alerts reviewed and acted on	Greater than 80% alert review rate within 24 hours	Alert management log, action completion tracking
Vision inspection replacing or augmenting manual check	Active on all in-scope quality gates	Inspection throughput and false-positive rate log
AI scheduling recommendations accepted by planners	Greater than 70% acceptance rate	Recommendation log with accept or override tracking
Model confidence scores reviewed before action	100% of A2 and above actions reviewed	Action verification audit trail

Adoption Criterion	Target	How to Assess
Model retraining triggered on performance degradation	Automated monitoring active with defined thresholds	Model registry performance dashboard
Root-cause investigations using platform analytics	Greater than 60% of investigations use platform data	Maintenance and quality investigation log

B.3 Digital Twin and Simulation Adoption

Adoption Criterion	Target	How to Assess
Production changes validated in twin before deployment	100% of in-scope process changes twin-validated	Change management log with twin validation records
Engineers using twin for scenario planning	Active users among engineering team on a monthly basis	User activity log, scenario run count
Twin accuracy vs actual process behaviour	Within 5% variance on primary KPIs	Calibration reports comparing twin predictions to actuals
Simulation used to evaluate capital decisions	All capex above defined threshold assessed in twin first	Investment decision records with simulation evidence

B.4 Robotics and Bounded Autonomy Adoption

Adoption Criterion	Target	How to Assess
Robot cells operating at defined autonomy level	No cell operating above approved autonomy level	Safety case and autonomy level certification record
Exception handling rate in autonomous cells	Less than 5% of cycles requiring human intervention	Cell exception log, operator intervention record
Safety incident rate in robot-cell zones	Zero incidents attributable to AI-driven action	Incident report log, safety case review record
Rollback events and recovery time	Rollback available and tested on	Rollback test log, mean time to recovery

Adoption Criterion	Target	How to Assess
	every A3 or above cell	
Operator confidence in robot cell supervision	Positive score in structured operator survey	Bi-monthly operator confidence survey results

> *Technology adoption is measured by behaviour, not by deployment. The test is not whether the system is running. It is whether the people closest to the process trust it, use it, and prefer it to what they did before.*

Section C: Job Redefinition and New Role Descriptions

Industrial AI does not simply replace roles. It restructures them. The tasks most amenable to automation are the most repetitive, the most physically demanding, and the most cognitively narrow. What remains, and what grows, are the roles that require judgement, contextual knowledge, system understanding, and the ability to supervise, interpret, and correct intelligent systems.

The following describes the primary role transformations that accompany a mature industrial intelligence deployment, organised by function.

C.1 Role Transformation Overview

Previous Role	Redefined Role Title	Primary Shift	New Capability Required
Maintenance Technician	Asset Intelligence Technician	Reactive repair to predictive intervention	AI alert interpretation, root-cause analysis, digital work orders
Quality Inspector	Quality Intelligence Analyst	Manual sampling to AI-augmented continuous monitoring	Vision system supervision, defect pattern analysis, model feedback
Production Planner	AI-Assisted Production Strategist	Spreadsheet planning to scenario-based optimisation	Digital twin operation, AI recommendation review, exception handling
Process Engineer	Simulation-First Process Engineer	Physical trial and error to virtual validation	Physics simulation, surrogate modelling, digital twin calibration
Plant Operator	System Supervisor and Exception Handler	Equipment operation to intelligent system supervision	Copilot interaction, anomaly review, safe-state intervention
Data Analyst	Industrial AI Product Owner	Reporting to platform management	AI model governance, data contract management, use case prioritisation
IT Systems Manager	Industrial Platform Architect	Infrastructure management to intelligent edge-cloud operations	Edge compute, OT security, data pipeline governance
Robot Cell Operator	Physical AI Cell Coordinator	Manual task execution to robot fleet supervision	Autonomy level management, exception routing, safety boundary monitoring

C.2 Detailed Role Descriptions

Asset Intelligence Technician

Previously focused on reactive maintenance, this role shifts to operating as the human partner of the predictive maintenance AI. The technician reviews AI-generated risk scores, interprets anomaly patterns, prioritises interventions based on failure consequence and operational context, and closes the feedback loop by recording outcomes that improve model accuracy.

Core responsibilities: review and act on AI maintenance alerts, validate or override model risk classifications, perform targeted inspections based on AI recommendations, document intervention outcomes, and identify patterns the model has not yet learned.

New capability requirements: working knowledge of the predictive maintenance cockpit, ability to interpret vibration, temperature, and current anomaly signatures, understanding of confidence scores and their limitations, and willingness to challenge model recommendations when physical observation conflicts with AI output.

Quality Intelligence Analyst

This role evolves from manual inspection and sampling to continuous monitoring, AI model supervision, and defect pattern investigation. The analyst is responsible for ensuring the vision inspection system is performing accurately, for reviewing cases where the model confidence is below threshold, and for translating quality data into process improvement actions.

Core responsibilities: supervise AI defect classification on active lines, review and adjudicate borderline cases flagged by the model, conduct root-cause analysis using correlated process and quality data, feed confirmed defect cases back into model training pipelines, and report quality trends to engineering and operations.

New capability requirements: understanding of vision system performance metrics including precision, recall, and confidence calibration; ability to use the quality cockpit for pattern analysis; structured problem-solving methods applied to AI-surfaced defect clusters.

AI-Assisted Production Strategist

The production planner role expands from managing schedules in spreadsheets to operating as the human decision-maker in an AI-augmented planning system. The strategist uses digital twin scenario analysis to evaluate production options, reviews AI scheduling recommendations before release, and manages the exception cases that fall outside the model's operating envelope.

Core responsibilities: review and approve AI-generated production schedules, run scenario simulations for demand changes, supply disruptions, and line configurations, override recommendations where operational knowledge conflicts with model output, maintain planning logic documentation, and report planning performance against AI baselines.

New capability requirements: ability to operate the scenario planning cockpit and interpret digital twin outputs, structured approach to accepting or overriding AI recommendations, understanding of scheduling model parameters and their operational meaning.

Simulation-First Process Engineer

Process engineering shifts from physical prototype-led development to a simulation-first discipline. The engineer's primary design validation work occurs in digital space before any physical trials are run, reserving physical tests for final confirmation of simulation-validated configurations.

Core responsibilities: build and maintain process digital twins for assigned product families, validate design changes in simulation before physical implementation, calibrate surrogate models against production data, define and document confidence thresholds for simulation-validated changes, and manage the version history of engineering knowledge assets.

New capability requirements: working proficiency with physics simulation tools, understanding of surrogate model construction and accuracy assessment, ability to interpret simulation outputs and translate them into production recommendations, and disciplined version control of engineering models.

System Supervisor and Exception Handler

Plant operators in AI-enabled environments shift from direct equipment operation to intelligent system supervision. Their primary value is contextual judgement: recognising when the AI is misreading a situation, intervening correctly, and providing the feedback that improves system performance over time.

Core responsibilities: monitor AI system status and alert conditions across assigned assets, review operator copilot recommendations before acting on them, intervene in exception conditions that fall outside normal operating parameters, escalate safety-boundary events according to defined protocols, and provide structured feedback on AI guidance quality.

New capability requirements: confidence in interacting with the operator copilot interface, clear understanding of which decisions require human approval and why, ability to distinguish between AI recommendation and definitive instruction, and familiarity with safe-state procedures for each assigned system.

Industrial AI Product Owner

A new role that did not exist in traditional manufacturing organisations. The industrial AI product owner sits at the intersection of data, operations, and AI capability. They prioritise the use case backlog, manage model governance, define data contracts, and ensure that the platform continues to deliver value as the operational environment evolves.

Core responsibilities: maintain the AI use case roadmap and prioritise against operational value, manage model performance monitoring and retraining triggers, own data quality standards and data contract documentation, coordinate between plant teams, IT, OT, and AI development resources, and report platform value delivery to programme sponsors.

New capability requirements: operational understanding of manufacturing processes, ability to frame operational problems as AI use cases, working knowledge of model lifecycle management, structured stakeholder communication skills, and comfort with quantitative performance measurement.

C.3 Job Redefinition Success Criteria

Success Criterion	Target	Measurement Approach
Affected roles have defined new job descriptions before platform launch	100% of in-scope roles documented	HR role documentation audit pre-launch
Retraining completion for all transition roles	Greater than 90% completion before autonomy level advances	Training completion records by role and site
Operator confidence in working with AI systems	Positive score in structured bi-monthly survey	Survey results tracked over time by role group
Voluntary attrition in affected roles	No increase above pre-programme baseline	HR attrition tracking by function and site
Grievances or escalations related to AI deployment	Zero unresolved after 30 days	HR grievance log and resolution tracking
New role capability assessments completed	100% of transition roles assessed at 6 and 12 months	Line manager assessment records
Internal promotions into AI-adjacent roles	Growing proportion over programme lifetime	HR talent movement tracking

Section D: Programme Checklists

The following checklists are designed for use at stage gates, programme reviews, and site-level deployment assessments. Each checklist maps to the success criteria defined in Sections A through C. Items are grouped by domain and labelled with their applicable horizon.

Checklist 1: Financial Implementation and ROI

Foundation Economics (Horizon One Gate)

- ☐ **Pilot line throughput baseline established and documented before deployment**
  Required for valid before vs after comparison
- ☐ **Asset-level downtime and maintenance cost baseline recorded for all targeted assets**
  Minimum 90 days of pre-deployment data
- ☐ **First-pass yield and scrap rate baseline recorded by line and product family**
- ☐ **Total Horizon One investment cost captured including integration labour, sensors, and licences**
- ☐ **Payback model completed with conservative, base, and optimistic scenarios**
- ☐ **Stage-gate financial threshold defined and approved by CFO or finance sponsor**
- ☐ **Pilot line improvements sustained for 90 days before Horizon Two capital approved**
  Critical gate: improvement must be sustained, not a single-period result
- ☐ **Positive platform ROI demonstrated and independently verified**

Scaled Programme Economics (Horizon Two Gate)

- ☐ **Throughput improvements recorded across all Horizon Two lines** *Minimum 8% target for gate passage*
- ☐ **Engineering cycle time reduction documented with project-level evidence**
- ☐ **Capex avoidance from simulation-first validation quantified and recorded**
- ☐ **Platform deployment cost for second and third sites tracked against first site** *Declining cost confirms platform playbook is working*
- ☐ **Schedule reliability metric baseline established and improvement tracked**
- ☐ **Multi-site ROI model updated with actual data from each deployed site**
- ☐ **External validation of financial claims completed before portfolio expansion**

Platform Economics (Horizon Three Gate)

- ☐ **Recurring platform revenue established and growing**
- ☐ **Deployment cost per new site is at least 20% lower than site one by site three**
- ☐ **AI-native contract manufacturing margin measured against sector comparables**
- ☐ **Engineering simulation service operating as a separate revenue line**

- ☐ **Board-level portfolio IRR model completed and reviewed against targets**
- ☐ **Capital recycling from proven platform ROI funding next portfolio acquisitions**

Checklist 2: Technology Adoption

Data and Infrastructure Readiness

- ☐ **All priority assets connected and streaming with greater than 95% uptime**
- ☐ **Asset graph complete and versioned for all in-scope assets** *Gaps must be resolved before AI layer is activated*
- ☐ **Data quality validation active on all critical data streams**
- ☐ **IT and OT teams jointly reviewing dashboards on a weekly cadence**
- ☐ **Data latency within specification for all operational applications** *Sub-3 seconds for operator-facing queries*
- ☐ **Security baseline implemented: segmented IT/OT network, role-based access, signed updates**
- ☐ **Edge compute operational and failover tested for local resilience**

AI Models and Analytics

- ☐ **Predictive maintenance models live on all priority assets**
- ☐ **Alert review rate above 80% with documented action outcomes**
- ☐ **Vision inspection system active on all in-scope quality gates**
- ☐ **False-positive rate below defined threshold on all inspection streams**
- ☐ **AI scheduling recommendations integrated into planning workflow**
- ☐ **Planner acceptance rate above 70% for AI scheduling proposals**
- ☐ **Model performance monitoring active with automated degradation alerts**
- ☐ **Model retraining process defined, tested, and documented**
- ☐ **Action verification gate active on all A2 and above recommendations**
- ☐ **Audit trail complete and immutable for all model-driven actions**

Digital Twin and Simulation

- ☐ **Process twins calibrated and validated for all in-scope asset families**
- ☐ **Twin accuracy verified within 5% variance on primary KPIs**
- ☐ **Change management policy requires twin validation before physical deployment**
- ☐ **Engineering team trained and actively using twin for scenario planning**
- ☐ **Simulation evidence required and present for all in-scope capex decisions**
- ☐ **Surrogate models operational where full physics simulation is latency-impractical**

Robotics and Bounded Autonomy

- ☐ Safety case completed and approved for all robot cells before A3 activation
- ☐ Each cell operating at or below its approved autonomy level
- ☐ Exception handling rate below 5% of cycles on all A3 cells
- ☐ Rollback procedure tested and documented for each autonomous cell
- ☐ Zero safety incidents attributable to AI-driven actions
- ☐ Operator confidence survey showing positive scores on supervised cells
- ☐ A4 autonomy advancement subject to formal safety case review

Checklist 3: Job Redefinition and Workforce

Before Deployment

- ☐ All roles affected by the deployment identified and mapped
- ☐ Redefined job descriptions written and approved for all transition roles
- ☐ Retraining curriculum defined for each transition role
- ☐ Retraining delivery schedule agreed with HR and line management
- ☐ Union or works council consultation completed where applicable
- ☐ Baseline workforce confidence and sentiment survey completed
- ☐ No deployment autonomy advances past A1 until role redesign is complete

During Deployment

- ☐ Retraining completion rate above 90% for all transition roles before each autonomy advance
- ☐ Operator copilot interfaces reviewed with users before go-live
- ☐ Feedback mechanism active for reporting AI guidance quality
- ☐ Attrition monitored monthly in all affected role groups
- ☐ Grievance escalation path defined and communicated to all affected staff
- ☐ Line manager capability assessments completed at 6 months for all transition roles

After Deployment

- ☐ Role capability assessments repeated at 12 months
- ☐ Operator confidence survey showing improvement from pre-deployment baseline
- ☐ Voluntary attrition in affected roles at or below pre-programme baseline
- ☐ Internal promotions into AI-adjacent or supervisory roles tracked and reported

- ☐ Job design reviewed and updated as platform autonomy level advances
- ☐ Workforce ROI documented: productivity per head, capability improvement, retention

Checklist 4: Governance and Audit Readiness

Governance Infrastructure

- ☐ Policy engine operational and enforcing defined action rules
- ☐ Human approval workflow active for all A2 and above actions
- ☐ Rollback mechanism tested for all governed action types
- ☐ Immutable audit log active for all model-driven actions
- ☐ Model version registry operational with approval records
- ☐ Safety boundary definitions documented and enforced in platform
- ☐ Cyber incident response plan tested and up to date

Audit Readiness

- ☐ Audit trail covers data source, model version, confidence score, and approval for every action
- ☐ Regulatory compliance review completed for all safety-critical applications
- ☐ Third-party safety case review completed for A3 and above autonomy deployments
- ☐ Board-level AI governance policy in place and reviewed annually
- ☐ Incident review process defined and tested with at least one dry run

> *A programme that passes all four checklists has not just deployed technology. It has built a governed, trusted, and financially justified platform that earns the right to expand. That is the only path to portfolio scale.*

Summary: Programme Health Scorecard

The scorecard below provides a single-page view of programme status across all four dimensions. It is designed for monthly programme reviews and quarterly board reporting. Status ratings are: On Track, At Risk, and Requires Action.

Dimension	Key Indicator	Horizon	Status
Financial ROI	Pilot line payback within target period	H1	
Financial ROI	Throughput improvement sustained 90 days	H1	
Financial ROI	Capex avoidance from simulation validated	H2	
Financial ROI	Platform deployment cost declining with each site	H2	
Financial ROI	Portfolio IRR above 15 percent	H3	
Technology Adoption	Asset connectivity above 95 percent	H1	
Technology Adoption	Alert review rate above 80 percent	H1	
Technology Adoption	Twin validation active on all change types	H2	
Technology Adoption	Robot cell exception rate below 5 percent	H2-H3	
Job Redefinition	Role redesign complete before deployment	Pre-launch	
Job Redefinition	Retraining completion above 90 percent	H1	
Job Redefinition	Operator confidence score improving	H1-H2	
Job Redefinition	Attrition at or below baseline	H1-H3	
Governance	Action verification gate active	H1	
Governance	Audit trail complete and immutable	H1	
Governance	Safety case approved for A3 cells	H2	
Governance	Regulatory compliance review complete	H2	

Status definitions: On Track means the criterion is met or the programme is within acceptable variance of the target. At Risk means the criterion is behind schedule or below target but has a credible recovery plan. Requires Action means the criterion is materially off target, there is no current recovery plan, and the programme sponsor must be engaged within five working days.

This document should be reviewed and updated at each programme stage gate. Checklist items that are not applicable to the current deployment scope should be marked as deferred with a rationale, not removed. The deferred list forms part of the expansion scope for subsequent horizons.

Section E: Localisation Success Criteria

An industrial intelligence platform deployed across multiple regions, acquired plants, or sovereign manufacturing programmes cannot use a single uniform deployment template. Localisation is not a translation exercise. It is a structured process of adapting the platform architecture, governance model, workforce approach, regulatory posture, and supply chain relationships to the specific legal, cultural, industrial, and political conditions of each operating territory.

Failing to localise does not simply create operational friction. It creates compliance exposure, workforce rejection, regulatory risk, and in strategic sectors such as defence and semiconductors, it can invalidate the entire deployment for sovereign procurement purposes.

E.1 Regulatory and Compliance Localisation

Every jurisdiction applies a different legal framework to AI systems operating in physical environments. The criteria below apply regardless of sector and must be assessed before any deployment enters the A2 autonomy level or above.

Localisation Criterion	Minimum Standard	Assessment Approach
Local AI and automation regulatory mapping complete	All applicable regulations identified and documented	Legal review against EU AI Act, local health and safety law, sector-specific codes
Data residency requirements met	All plant data stored within required jurisdiction	Data architecture review against local data sovereignty laws
Product liability framework defined for AI-driven actions	Clear accountability chain documented for each action type	Legal counsel review of action verification logs and approval records
Export control compliance for AI and robotics components	All hardware and software screened against applicable export rules	Procurement and legal sign-off before deployment of controlled components
Local safety standards met for robot cells and autonomous systems	Certified to applicable standard in each jurisdiction	Third-party safety certification by approved local body
Privacy rules met for operator monitoring and biometric systems	Compliance confirmed where operator interaction data is collected	DPO or privacy counsel review of data collection scope

E.2 Infrastructure and Technology Localisation

Platform infrastructure must reflect local connectivity conditions, power reliability, edge compute availability, and the maturity of the local vendor ecosystem. A deployment architecture designed for a well-connected Western European plant will not perform identically in a reshoring facility in a region with limited grid stability or restricted cloud access.

Localisation Criterion	Minimum Standard	Assessment Approach
Edge-first architecture deployed where cloud connectivity is unreliable	Local inference and buffering operational before go-live	Connectivity reliability test over 30-day period pre-deployment
Local vendor ecosystem assessed for support and maintenance	At least two qualified local support partners identified per site	Vendor capability assessment and SLA review
Power and environmental specifications met for edge hardware	Hardware rated for local temperature, humidity, and power variance	Site survey and hardware specification sign-off
Language localisation of operator interfaces	All operator-facing screens and alerts in local primary language	UI review with native speakers from the plant workforce
Local network segmentation architecture approved by local security authority	IT/OT segmentation reviewed and approved before go-live	Security assessment by locally accredited assessor
Spare parts and replacement hardware available locally	Critical components available within 48-hour replacement window	Inventory and supplier lead time audit
Local simulation data calibrated to regional process norms	Twin models reflect local material specs, energy costs, and process conditions	Model calibration review against local production data

E.3 Workforce and Cultural Localisation

Workforce adoption is always local. The barriers to adoption in a German automotive plant are different from those in a Malaysian electronics facility, a French aerospace site, or a US defence contractor. Cultural attitudes toward AI authority, local labour law on automation, union structures, and the local skills base all shape what an effective adoption programme looks like.

Localisation Criterion	Minimum Standard	Assessment Approach
Local labour law compliance for automation deployment	Legal review completed before any role change takes effect	Employment counsel sign-off per jurisdiction

Localisation Criterion	Minimum Standard	Assessment Approach
Works council or union consultation completed	Formal consultation process completed where legally required	Consultation records and agreed implementation terms documented
Retraining curriculum adapted to local education baseline	Training materials validated with local workforce sample before rollout	Pilot training session with feedback scoring above 70 percent comprehension
Local training delivery partners identified	At least one accredited local training provider engaged	Provider capability assessment and curriculum review
Cultural attitudes to AI authority assessed and addressed	Workforce sentiment survey completed and findings incorporated into UX design	Pre-deployment survey and debrief with local plant leadership
Local management capability to oversee AI systems	Site leadership trained on platform governance and escalation	Management training completion records

E.4 Supply Chain and Sourcing Localisation

In strategic sectors and sovereign programmes, the sourcing provenance of AI components, sensors, compute hardware, and robotics is subject to political scrutiny and procurement rules. A platform that cannot demonstrate local or allied-nation sourcing for critical components will be excluded from an increasing number of defence, semiconductor, and government-adjacent manufacturing programmes.

Localisation Criterion	Minimum Standard	Assessment Approach
Critical component sourcing mapped by country of origin	Full bill of materials with country of origin for all AI and robotics components	Procurement audit against applicable sourcing rules
Restricted-origin components identified and alternatives assessed	No restricted-origin components in safety-critical or data-handling roles	Legal and security review of component sourcing map
Local content requirements met for government-adjacent programmes	Local content percentage meets programme threshold	Procurement records and local content calculation submitted
Supplier resilience assessment completed for critical components	At least two qualified suppliers for all critical components	Supplier risk register reviewed quarterly

Localisation Criterion	Minimum Standard	Assessment Approach
Software supply chain integrity verified	All third-party software components reviewed for provenance and vulnerability	Software bill of materials audit and security scan

E.5 Sovereign and Geopolitical Positioning

For programmes operating in or adjacent to strategic national industries, the platform must demonstrate alignment with national industrial policy objectives. This is not a soft criterion. It directly affects access to government contracts, subsidies, export licence eligibility, and regulatory fast-track pathways.

Localisation Criterion	Minimum Standard	Assessment Approach
National industrial policy alignment documented	Platform deployment mapped to relevant national AI or industrial strategy	Policy alignment briefing reviewed and approved by government affairs lead
Data sovereignty architecture approved for strategic programmes	Classified or controlled data never leaves approved national infrastructure	Security architecture review by competent national authority
Reshoring or domestic capacity contribution quantified	Jobs created, local spend, and domestic production capacity calculated	Economic impact report prepared for each site
Government or defence customer security clearance requirements met	All personnel with system access cleared to required level	Security clearance verification records maintained
Platform positioned for applicable national innovation incentives	Relevant R&D tax credits, grants, or subsidies identified and applied for	Finance team review of local incentive landscape

> *Localisation is a competitive advantage, not a compliance cost. The platform that can demonstrate sovereign capability, local workforce benefit, and regulatory alignment will access markets that a purely technology-led approach cannot reach.*

Checklist 5: Localisation Readiness

Regulatory and Legal

- ☐ **Local AI and automation regulatory mapping completed before deployment begins**
  Covers EU AI Act, local health and safety, sector codes
- ☐ **Data residency architecture confirmed compliant with local law**
- ☐ **Product liability chain documented for all governed action types**
- ☐ **Export control screening completed for all hardware and software components**
- ☐ **Robot cell safety certification by locally approved body**
- ☐ **Privacy compliance confirmed where operator interaction data is collected**

Infrastructure and Technology

- ☐ **Edge-first deployment confirmed where cloud reliability is below threshold**
- ☐ **Local support vendors identified with signed SLAs**
- ☐ **Hardware specifications validated against local environmental and power conditions**
- ☐ **All operator interfaces localised to primary plant language** *Validated by native-speaking workforce members*
- ☐ **IT/OT segmentation reviewed and approved by locally accredited security assessor**
- ☐ **Critical spare parts inventory confirmed within 48-hour replacement window**
- ☐ **Twin models calibrated to local material specifications and process norms**

Workforce and Culture

- ☐ **Local labour law compliance confirmed before any role changes take effect**
- ☐ **Works council or union consultation completed and documented where required**
- ☐ **Retraining curriculum validated with local workforce sample** *Comprehension score above 70 percent in pilot session*
- ☐ **Local accredited training delivery partner engaged**
- ☐ **Pre-deployment workforce sentiment survey completed**
- ☐ **Site leadership completed platform governance training**

Supply Chain and Sovereign Positioning

- ☐ **Full component bill of materials with country of origin completed**
- ☐ **No restricted-origin components in safety-critical or data-handling roles**
- ☐ **Local content percentage calculated and compliant with programme requirements**
- ☐ **At least two qualified suppliers confirmed for all critical components**
- ☐ **Software bill of materials audit and security scan completed**
- ☐ **National industrial policy alignment documented and approved**
- ☐ **Data sovereignty architecture approved for any strategic or defence-adjacent scope**
- ☐ **Applicable national innovation incentives and grants identified and applied for**

Section F: Industry-Specific Success Criteria

The generic success criteria in Sections A through E apply across all manufacturing deployments. The criteria in this section are additive: they supplement the generic framework with the specific quality standards, regulatory requirements, production characteristics, and risk profiles of six major industrial sectors.

Each sector sub-section covers: the defining characteristics that shape AI deployment in that sector, the additional success criteria that apply, and a supplementary checklist. Programme teams should apply the relevant sector criteria alongside, not instead of, the criteria in the preceding sections.

F.1 Defence and Aerospace

Defence and aerospace manufacturing operates under the most demanding quality, traceability, and security requirements of any industrial sector. Every component, process step, and configuration decision must be documented and defensible across the product lifetime, which may span decades. AI systems that cannot provide full provenance for every recommendation and action are not deployable in regulated defence programmes.

Sector-Specific Criterion	Standard Required	Verification Method
AS9100 or equivalent quality management system maintained alongside AI platform	Certification maintained throughout deployment	Quality management audit records
Full traceability of every AI-influenced production decision to approved design data	100% traceability for all flight-critical and safety-critical components	Configuration management system review
ITAR and EAR export control compliance for all AI and data systems	No controlled technical data accessible outside approved access list	Export control audit and access log review
Security clearance architecture for all AI system access in classified programmes	All personnel and system access points cleared to required classification level	Security authority review and certification
AI model validation against certified simulation for safety-critical processes	Model predictions validated against certified physics model, not production data alone	Airworthiness or programme authority sign-off
Design authority approval for all AI-recommended process parameter changes	No parameter change enacted without DA approval record	Change control record audit

Sector-Specific Criterion	Standard Required	Verification Method
Non-destructive testing and inspection AI validated against human baseline	Vision and AI inspection correlation above 99.5% for critical features	Validation study against certified human inspector baseline
Manufacturing execution system integration with programme master schedule	AI scheduling recommendations aligned to programme critical path	MES and programme scheduling integration test

Checklist 6a: Defence and Aerospace Supplementary

- ☐ **AS9100 or equivalent quality certification maintained and audit-ready**
- ☐ **Full traceability chain confirmed for all safety-critical and flight-critical AI-influenced decisions**
- ☐ **ITAR and EAR compliance review completed for all system components and data flows**
- ☐ **Security clearance architecture reviewed and approved by competent authority**
- ☐ **AI model validation against certified physics simulation completed for safety-critical processes**
- ☐ **Design authority approval process integrated with AI change recommendation workflow**
- ☐ **Vision inspection system validated above 99.5% correlation with certified human baseline**
- ☐ **AI scheduling integration tested against programme master schedule and critical path**
- ☐ **Domestic content and supply chain requirements met for sovereign defence programmes**

F.2 Semiconductors

Semiconductor manufacturing operates at tolerance levels, environmental sensitivities, and process complexities that exceed virtually every other industrial sector. Yield is the primary financial lever, and yield variations of fractions of a percent translate into material financial consequences at production scale. AI in semiconductor manufacturing must be exceptionally precise, extensively validated, and capable of operating in cleanroom environments with no tolerance for contamination from hardware or human intervention.

Sector-Specific Criterion	Standard Required	Verification Method
AI-driven process control validated against statistical process control baselines	Model-recommended adjustments improve Cpk, not degrade it	Statistical process control comparison before and after AI engagement

Sector-Specific Criterion	Standard Required	Verification Method
Cleanroom compatibility of all edge hardware and sensor installations	All hardware rated and certified for relevant cleanroom classification	Cleanroom qualification records for all installed hardware
Yield impact model validated before any AI-driven parameter adjustment	Predicted yield impact confirmed within tolerance in simulation before deployment	Fab simulation comparison and yield model validation report
Wafer-level traceability linked to AI decision log	Every wafer traceable to the process conditions and AI state at each step	Traceability system integration test
Geopolitical supply chain risk assessed for all AI hardware components	No single-source dependencies on restricted-jurisdiction components	Component sourcing map reviewed against applicable trade rules
Equipment maker interface protocols approved for AI integration	All SECS/GEM or equivalent interfaces validated before live use	Equipment vendor sign-off on integration protocol
AI model update and deployment process validated to prevent yield disruption	Shadow mode testing before live deployment on production tools	Model deployment protocol review and shadow mode test records

Checklist 6b: Semiconductors Supplementary

- ◻ **Statistical process control baseline established for all AI-managed process steps**
- ◻ **Cleanroom compatibility certification completed for all installed hardware**
- ◻ **Yield impact model validated in simulation before any live AI parameter adjustment**
- ◻ **Wafer-level traceability linked to AI decision log confirmed end to end**
- ◻ **Component sourcing map reviewed for geopolitical supply chain risk**
- ◻ **Equipment maker interface protocols approved by vendor before live integration**
- ◻ **Shadow mode testing protocol completed for all model updates before production deployment**
- ◻ **Data sovereignty review completed for all wafer-level process data**

F.3 Automotive and Mobility

Automotive manufacturing is characterised by high volume, tight cycle times, complex multi-tier supply chains, and a technology transition of exceptional scale as the industry moves from internal combustion to electric powertrains. AI deployment must support flexible assembly for rapidly changing model mixes, deliver consistent quality at volume, and integrate with a supplier ecosystem that varies widely in digital maturity.

Sector-Specific Criterion	Standard Required	Verification Method
IATF 16949 quality management alignment maintained throughout AI deployment	Certification maintained and no non-conformances attributable to AI platform	Quality system audit records
AI scheduling validated for high-mix production sequences including model changeovers	Scheduling recommendation accuracy above 90% for mixed-model sequences	Scheduling log review over 60-day production sample
Battery and electric powertrain process AI validated against OEM specifications	AI process recommendations validated against OEM-approved parameter ranges	OEM or customer sign-off on AI parameter boundary definitions
AI vision inspection validated for weld quality, sealing, and assembly completeness	Inspection correlation above 99% for safety-critical joints and seals	Correlation study vs certified human inspection standard
Supplier quality data integration active for Tier 1 inputs	AI quality platform receiving and acting on inbound quality data from key Tier 1 suppliers	Supplier data feed integration and data quality audit
Robot cell human-robot collaboration zones certified to ISO TS 15066	All collaborative robot zones certified before any human entry permitted	ISO TS 15066 certification records per cell
Gigacasting and large-format press AI process control validated in simulation first	All press parameter changes validated in digital twin before production trial	Twin validation record and sign-off by process engineering

Checklist 6c: Automotive and Mobility Supplementary

- ☐ **IATF 16949 certification status confirmed and maintained through deployment**
- ☐ **High-mix scheduling accuracy validated over 60-day production sample**
- ☐ **EV and battery process AI boundaries approved by OEM or customer engineering**

- ☐ **Vision inspection correlation study completed for all safety-critical joints and seals**
- ☐ **Tier 1 supplier quality data feed active and integrated into quality AI layer**
- ☐ **All collaborative robot zones ISO TS 15066 certified before human access**
- ☐ **Gigacasting and large-format press changes twin-validated before production trial**
- ☐ **AI platform tested against planned model changeover sequences before volume production**

F.4 Advanced Materials and Chemicals

Advanced materials and specialty chemicals manufacturing involves complex, interdependent process variables, strict environmental and safety regulations, and long development cycles between formulation and production at scale. AI can accelerate materials discovery, optimise batch processes, and manage the interaction of variables that exceed human capacity to track simultaneously, but it must do so within safety regimes designed for hazardous materials and exothermic processes.

Sector-Specific Criterion	Standard Required	Verification Method
Process safety management integration with AI action layer	AI recommendations blocked from affecting safety-instrumented systems without SIL-rated interlock	Safety instrumented system review and functional safety audit
REACH, COSHH, or local hazardous materials compliance verified for AI sensor deployments	All sensors and edge hardware validated as suitable for hazardous area	ATEX or equivalent hazardous area certification for installed equipment
AI batch optimisation validated against validated process design space	Model recommendations constrained within regulatory-approved design space	Quality by Design or equivalent validation documentation
Environmental monitoring AI linked to regulatory reporting	AI anomaly detection active on all regulated environmental outputs	Environmental reporting integration test
AI materials discovery models validated against laboratory experimental data	Predictions validated in laboratory before progressing to production scale	Lab validation records linked to model provenance
Emergency shutdown system independence from AI layer confirmed	AI has no ability to override or delay emergency shutdown systems	SIS independence review by functional safety engineer

Checklist 6d: Advanced Materials and Chemicals Supplementary

- ☐ **Safety instrumented systems confirmed independent of AI action layer** *AI cannot override or delay emergency shutdown*
- ☐ **ATEX or equivalent hazardous area certification for all installed edge hardware**
- ☐ **AI batch recommendations constrained within regulatory-approved process design space**
- ☐ **Environmental monitoring AI integrated with regulated emissions reporting**
- ☐ **Materials discovery model predictions validated in laboratory before scale-up**
- ☐ **Functional safety audit completed confirming AI-SIS independence**
- ☐ **REACH, COSHH, or local hazardous materials compliance confirmed for all sensor types**

F.5 FMCG and Food Manufacturing

Fast-moving consumer goods and food manufacturing operate at high volume and low margin with strict food safety, hygiene, and labelling regulations. AI deployment must support rapid changeovers, allergen management, shelf-life optimisation, and the traceability requirements of food safety regulations across multiple jurisdictions. Consumer trust is the ultimate constraint: any AI-related quality failure that reaches the consumer creates reputational damage that is disproportionate to the technical scale of the incident.

Sector-Specific Criterion	Standard Required	Verification Method
HACCP and food safety management system integration with AI quality layer	AI quality alerts linked to CCP monitoring and HACCP records	HACCP review and AI integration audit
Allergen changeover process AI validated for cleaning and segregation compliance	AI scheduling includes allergen status in all changeover sequencing	Changeover sequence audit and allergen risk assessment
Product traceability from raw material to finished goods linked through AI data layer	Full lot traceability confirmed within 4 hours in recall simulation	Recall simulation exercise results
AI vision inspection validated for foreign body and label compliance detection	Detection rate validated above regulatory minimum for all foreign body types in scope	Validation study records per product family

Sector-Specific Criterion	Standard Required	Verification Method
Shelf-life and cold chain AI models validated against microbiological testing data	Model predictions validated by accredited laboratory testing	Laboratory validation records linked to model provenance
Regulatory reporting integration for country-specific labelling requirements	AI label verification active for all target export markets	Label compliance audit across all product and market combinations in scope

Checklist 6e: FMCG and Food Manufacturing Supplementary

- ☐ **HACCP records linked to AI quality monitoring layer**
- ☐ **Allergen status integrated into AI scheduling and changeover sequencing**
- ☐ **Recall simulation completed: full lot trace confirmed within 4-hour target**
- ☐ **Foreign body and label compliance detection validated per regulatory minimum**
- ☐ **Shelf-life and cold chain model predictions validated by accredited laboratory testing**
- ☐ **Label compliance AI active and validated for all target export markets**
- ☐ **Consumer complaint data integrated into AI quality feedback loop**

F.6 Energy and Utilities Equipment Manufacturing

Manufacturing for the energy sector, including wind turbines, grid infrastructure, power electronics, and nuclear components, is characterised by long product lifetimes, extreme reliability requirements, and regulatory oversight that extends to the manufacturing process as well as the finished product. AI deployment must demonstrate that it improves, rather than introduces risk to, the reliability and traceability standards that are a contractual and regulatory condition of supply.

Sector-Specific Criterion	Standard Required	Verification Method
Lifetime traceability architecture for components with multi-decade service life	All manufacturing data preserved and accessible for product lifetime plus 10 years	Data retention architecture review against customer and regulatory requirements
Nuclear quality assurance compliance where applicable	AI system classified and documented under applicable nuclear QA code	Nuclear QA programme review by qualified assessor
Weld inspection and NDT AI validated to relevant energy sector standard	Inspection AI validated to ASME, EN, or applicable	Validation study against code-approved human inspection method

Sector-Specific Criterion	Standard Required	Verification Method
	welding code standard	
AI predictive maintenance validated for long-cycle equipment with sparse failure data	Model performance validated on held-out historical data over minimum 5-year window	Model validation report including false positive and miss rate analysis
Customer and regulatory witness point integration with AI quality records	AI quality evidence available in approved format for all witness point inspections	Witness point audit and customer sign-off on AI evidence format
Grid stability and export market regulatory compliance for power electronics manufacturing	All applicable IEC standards verified for AI-influenced processes	Standards compliance review per product type

Checklist 6f: Energy and Utilities Equipment Supplementary

- ☐ **Data retention architecture confirmed compliant with product lifetime requirements**
  Minimum product lifetime plus 10 years for energy sector
- ☐ **Nuclear QA programme compliance confirmed where applicable**
- ☐ **Weld and NDT inspection AI validated to applicable ASME, EN, or equivalent standard**
- ☐ **Predictive maintenance model validated on minimum 5-year historical window**
- ☐ **AI quality evidence format approved by customer for witness point use**
- ☐ **Applicable IEC or equivalent standards compliance confirmed for AI-influenced processes**
- ☐ **Long-lifecycle component traceability tested in simulated audit scenario**

> *Industry-specific criteria are non-negotiable additions to the generic framework, not optional enhancements. A programme that passes the generic checklists but fails the sector criteria is not deployable in that sector. Sector criteria represent the minimum standard for trust in that industry's regulatory and customer environment.*

Updated Summary: Programme Health Scorecard

This scorecard incorporates all six dimensions including localisation and industry-specific criteria. Sector rows should be populated with the applicable sector from Section F. Status ratings: On Track, At Risk, Requires Action.

Dimension	Key Indicator	Horizon	Status
Financial ROI	Pilot line payback within target period	H1	
Financial ROI	Throughput improvement sustained 90 days	H1	
Financial ROI	Capex avoidance from simulation validated	H2	
Financial ROI	Portfolio IRR above 15 percent	H3	
Technology Adoption	Asset connectivity above 95 percent	H1	
Technology Adoption	Alert review rate above 80 percent	H1	
Technology Adoption	Twin validation active on all change types	H2	
Technology Adoption	Robot cell exception rate below 5 percent	H2-H3	
Job Redefinition	Role redesign complete before deployment	Pre	
Job Redefinition	Retraining completion above 90 percent	H1	
Job Redefinition	Operator confidence score improving	H1-H2	
Governance	Action verification gate active	H1	
Governance	Safety case approved for A3 cells	H2	
Localisation	Regulatory mapping complete per jurisdiction	Pre	
Localisation	Language localisation and cultural validation complete	Pre	
Localisation	Works council or union consultation completed	Pre	
Localisation	Component sourcing map reviewed for restricted origins	Pre	
Localisation	Data sovereignty architecture approved	H1	

Dimension	Key Indicator	Horizon	Status
[Sector: specify]	Sector quality certification maintained	Pre	
[Sector: specify]	Sector traceability requirements met	H1	
[Sector: specify]	Sector safety validation completed	H1-H2	
[Sector: specify]	Regulatory witness or audit body satisfied	H2	

Rows marked [Sector: specify] should be populated with the applicable sector designation from Section F before the scorecard is used for programme review. Each site may carry a different sector designation where the portfolio spans multiple industries.

This document should be reviewed and updated at each programme stage gate. Checklist items not applicable to the current deployment scope should be marked as deferred with a rationale, not removed.