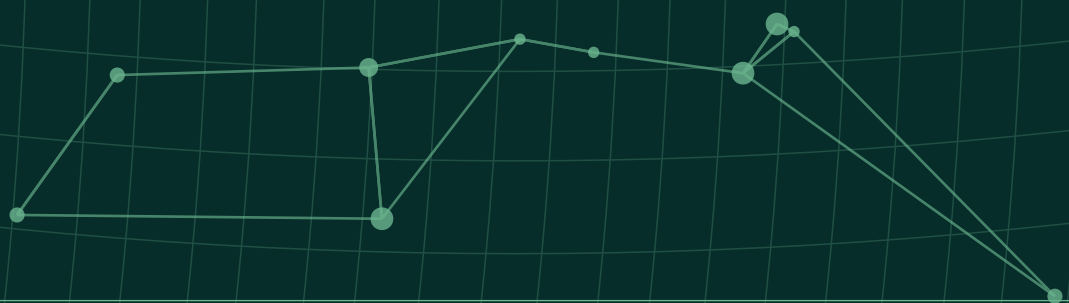
GREEN FUTURE TECH

LIMITED / HONG KONG SAR

GFT RESEARCH / WARP AND SIGNAL

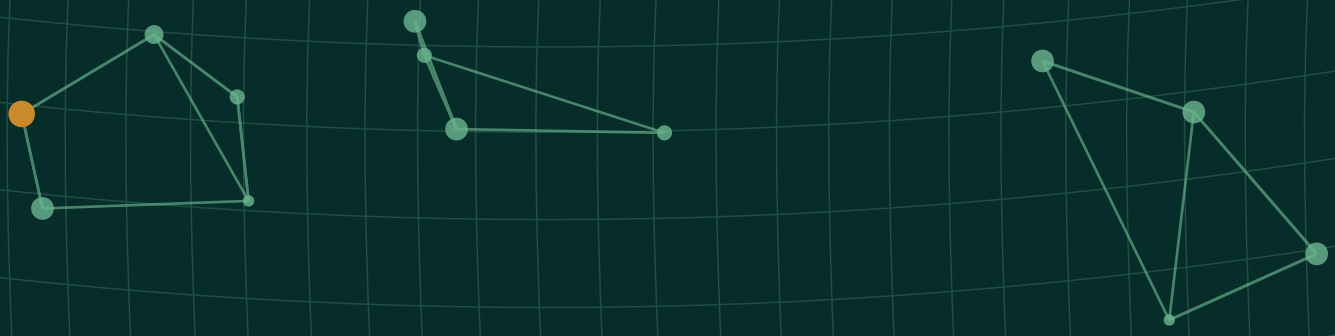
The Industrial Data Spine

Data governance and lightweight MLOps for small factory teams



wp14
RESEARCH EDITION

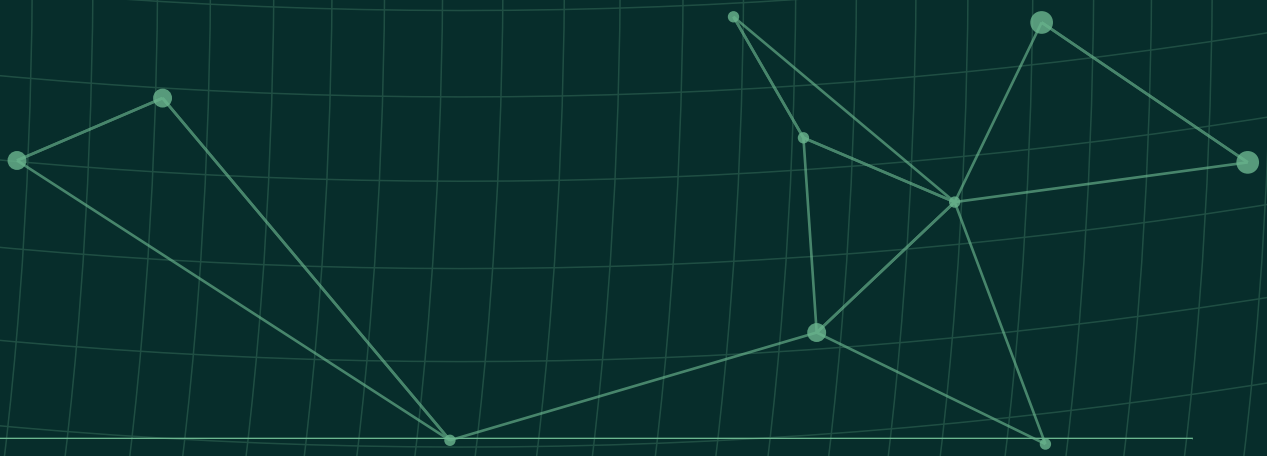
EDITION 2026-10-30
HONG KONG SAR



THE GFT THESIS

Value is made where it is measured.

GFT engineers sustainable factories that capture their own data from the first shift, so raw-material economies can keep the value they grow.



Read the evidence before acting on the recommendation.

ISSUER	Green Future Tech Limited, Hong Kong SAR
DOCUMENT	wp14 / The Industrial Data Spine
EDITION DATE	2026-10-30
AUDIENCE	Factory directors and operations managers in small and mid-sized manufacturers, Quality, maintenance, industrial-engineering and sustainability leads, IT and OT practitioners responsible for factory systems, Data analysts and machine-learning practitioners working with constrained teams, Technology vendors and integrators serving emerging-market factories, Development-finance and technical-assistance teams assessing digital readiness
STATUS	Research edition. Generated locally. Publication and operating-performance acceptance are separate decisions.
PROGRAMME	GFT's Cameroon textile programme: feasibility, performance-management system design and digital strategy. Prepared 31 January 2024; source-refresh annotations retained in 2026.

IMPORTANT NOTICE

This document is not investment, legal or tax advice. Proposed controls and illustrative calculations require project-specific validation. Targets and announced third-party projects are not evidence of delivered operating performance.

The confidential programme PDFs are not reproduced or published. No client identity or engagement code is included.

EVIDENCE REMAINS GRADED

A: official. B: institutional or major research. C: trade or local press and market-research vendors. D: unverified; excluded from use. A grade describes the supplied register and does not certify accuracy.

Inherited publication dates unrecorded: R01, R03, R06, R07, R09, R10, R11, R12, R14, R15, R16, R26, R27, R28, R29, R32. Temporal eligibility is unverified; these are inherited source-text checks, not independent source validation.

The operating case runs from fibre to feedback.

01 / A factory becomes reliably data-driven only when its records can be trusted 14

1.1 / Data debt starts on the first shift 16

1.4 / Disconnected logs fracture one production event 18

1.7 / Small teams need decision rights more than platforms 20

1.10 / Reuse turns preparation from a tax into infrastructure 21

1.12 / Trust includes safety, legality and worker dignity 23

1.16 / GFT's position is to govern the record before scaling the model 24

1.19 / The paper judges success by better decisions 26

02 / Named decision rights keep a small data team accountable 29

2.1 / Every critical domain receives one accountable owner 30

2.5 / Ownership, stewardship and custody remain distinct 32

2.9 / Line stewards need time and escalation authority 36

2.12 / Definition changes follow a controlled decision 38

2.15 / Model risk receives independent challenge 40

2.18 / An internal champion sustains adoption 42

2.21 / Data decisions use, but do not redesign, factory meetings 43

03 / One operational vocabulary joins the line to the ledger 47

3.1 / Stable identifiers anchor every record 48

3.3 / ISA-95 clarifies the boundary between control and business 50

3.4 / A KPI dictionary makes formulas contractual 52

3.5 / Events carry context that sensor values lack 54

3.6 / Transport and semantics must not be confused 56

3.8 / Canonical taxonomies preserve local detail 57

3.9 / Data contracts make compatibility an explicit promise 59

04 / Quality gates make fitness visible before errors compound 64

4.1 / Quality dimensions become operational tests 65

4.4 / Validation belongs at capture and ingestion 67

4.7 / Missing is not zero and estimated is not measured 69

4.10 / Rules and exceptions are versioned 72
4.13 / Lineage connects every derived value to its origins 73
4.16 / Quality improves in measured iterations 75
4.19 / Upstream checks limit downstream cascades 77

05 / The smallest durable spine separates capture, context and consumption 81

5.1 / Capture remains close to the work 82
5.4 / Raw evidence and curated records serve different purposes 83
5.7 / A canonical core replaces repeated point translation 85
5.10 / Transport follows the operating requirement 87
5.13 / Consumption receives controlled products, not raw access 89
5.16 / WP14 governs records while WP18 designs the topology 91
5.19 / Buy commodity capability but retain semantic control 92

06 / Time, identity and state turn sensor readings into industrial evidence 97

6.1 / Event time and ingestion time answer different questions 98
6.5 / Master data stabilise the factory's objects 100
6.8 / Genealogy links inputs, transformations and outputs 102
6.10 / Explicit states prevent incompatible downtime and defect counts 104
6.13 / Replay must not create a second reality 106
6.16 / Corrections append rather than erase 109
6.19 / Retention follows purpose and obligation 110

07 / Security and privacy are design constraints, not aftercare 113

7.1 / OT consequences set the control objective 114
7.3 / Least privilege covers people and machine identities 117
7.5 / A small-team baseline prioritises essential cyber hygiene 119
7.6 / Recovery evidence matters more than backup completion 120
7.7 / Incident response is practised as factory work 121
7.8 / Worker data are limited to a stated legitimate purpose 123
7.10 / Cameroonian legal duties require current local advice 124
7.11 / Equipment contracts must preserve lawful data access 126

08 / Continuity depends on recovery and portability, not perpetual connectivity 129

8.1 / Store-and-forward buys time during disconnection 130
8.4 / Local decisions continue when external services stop 132

8.7 / Recovery objectives vary by data class 134
8.10 / Open representations make exit practical 136
8.13 / Recovery includes the instructions that make data intelligible 137
8.16 / Migration preserves history from the manual system 139
8.19 / Deliberate failure tests expose hidden dependence 140

09 / A use case earns admission only when a decision and baseline exist 145

9.1 / Start with the decision and the person who acts 146
9.4 / A baseline makes improvement falsifiable 147
9.7 / Readiness tests data, rights and operating capacity 148
9.11 / The simplest useful comparator sets the bar 150
9.14 / Risk tier determines scrutiny and authority 151
9.17 / Cost includes the whole operating lifecycle 152
9.19 / Stop and postpone are successful governance outcomes 155
9.22 / GFT's sequence favours vision before history-hungry forecasting 157

10 / Reproducible training turns experiments into auditable candidates 162

10.1 / Version the complete experiment, not only the model file 163
10.2 / Dataset snapshots prevent silent training-set change 164
10.3 / Features inherit governed operational definitions 167
10.4 / Execution metadata make comparison possible 168
10.5 / A registry records status and authority 170
10.6 / A lightweight stack avoids platform sprawl 171
10.7 / Bought models face the same evidence discipline 173

11 / Validation and staged release keep models subordinate to operations 175

11.1 / People set the acceptance metrics and thresholds 176
11.4 / Local slices reveal aggregate success that hides failure 178
11.7 / Operational comparison includes the cost of being wrong 180
11.10 / Shadow and assisted modes earn autonomy gradually 183
11.13 / Rollback and manual fallback are release artefacts 185
11.16 / Approval separates building from acceptance 188
11.19 / Worker-affecting outputs retain meaningful human review 189
11.22 / The release package tells the floor what changed 190

12 / Monitoring must trigger diagnosis, not automatic retraining 194

12.1 / Service health and data quality come before accuracy 195

12.4 / Input drift and concept drift require different interpretation 197

12.7 / Delayed labels demand honest interim evidence 199

12.10 / Alerts need thresholds, owners and runbooks 201

12.13 / Retraining follows a diagnosed need and a new approval 204

12.15 / Feedback loops belong in the monitoring design 206

12.18 / One incident taxonomy joins data, model, security and OT response 207

12.20 / Review and retirement prevent model accumulation 210

13 / Adoption and value are operating metrics, not launch claims 214

13.1 / A champion and local stewards sustain use 215

13.3 / Training happens around tangible work 216

13.6 / Adoption measures record behaviour, not screen activity 218

13.9 / Operational outcomes use the factory's existing measures 219

13.13 / Financial value is net of lifecycle effort 221

13.15 / Trust and harm sit beside benefit 223

13.16 / A benefit register drives fix, scale or stop decisions 224

13.17 / Quality disclosure improves with the evidence 225

14 / A gated 18-month roadmap builds capability without a data department 227

14.1 / The first 30 days establish authority and scope 228

14.4 / Days 31 to 90 make the manual spine trustworthy 230

14.7 / Months 4 to 6 create reusable data products 232

14.10 / Months 7 to 9 prove one model in shadow mode 234

14.13 / Months 10 to 12 prove controlled operation 235

14.16 / Months 13 to 18 reuse the spine before adding complexity 237

14.19 / GFT should measure seven capabilities every month 238

14.20 / The closing recommendation is disciplined portability 240

Executive summary 12

Evidence register 243

Glossary 270

Term index 273

Method and imprint 275

Every exhibit states a testable claim.

EXHIBIT 01 / keyfigures
Local-processing ambition raises the cost of untrusted evidence 16

EXHIBIT 02 / flow
A small capture defect can compound into a wrong management action 17

EXHIBIT 03 / table
Disconnected shop-floor logs split one stoppage into four irreconcilable records 19

EXHIBIT 04 / matrix
Four admission tests separate a data spine from a data pile 25

EXHIBIT 05 / matrix
A six-role charter covers governance without creating a department 34

EXHIBIT 06 / table
Decision rights become testable when each has an artefact and escalation route 35

EXHIBIT 07 / hbar
A bounded weekly stewardship budget fits inside a standard 40-hour shift week 37

EXHIBIT 08 / flow
Controlled definition change prevents silent metric redefinition across shifts 39

EXHIBIT 09 / matrix
Data escalation adds decisions without creating a second management cadence 44

EXHIBIT 10 / layered-architecture
Identity and meaning sit between physical signals and decisions 51

EXHIBIT 11 / table
A minimum production event needs more than a timestamp and value 55

EXHIBIT 12 / table
A minimum production event needs more than a timestamp and value, continued 55

EXHIBIT 13 / flow
One governed event serves three views without three redefinitions 59

EXHIBIT 14 / matrix
Compatibility rules distinguish safe evolution from breaking change 62

EXHIBIT 15 / heatmap
A rule-coverage map reveals where fitness is assumed rather than tested 67

EXHIBIT 16 / flow
Quarantine keeps bad records visible without contaminating curated data 68

EXHIBIT 17 / flow
Lineage reconstructs a KPI and a model prediction on separate paths 74

EXHIBIT 18 / table
A quality register records the issue, consequence and next correction 76

EXHIBIT 19 / table
A quality register records the issue, consequence and next correction, continued 77

EXHIBIT 20 / layered-architecture
Three separations keep the spine understandable 84

EXHIBIT 21 / dumbbell
A canonical core cuts interface mappings as connected endpoints grow 87

EXHIBIT 22 / flow
Local buffering preserves capture during a disconnected interval 88

EXHIBIT 23 / table
Controlled data products isolate source evidence from consumer views 90

EXHIBIT 24 / matrix
Build-against-buy changes by layer, not by ideology 94

EXHIBIT 25 / timeline
Late arrival changes processing order without changing event time 100

EXHIBIT 26 / table
An identity registry records origin, authority and permitted reuse 102

EXHIBIT 27 / flow
Event genealogy preserves lot history through split and rework 103

EXHIBIT 28 / flow
A governed state machine stops one interval becoming two losses 106

EXHIBIT 29 / waterfall
Deduplication and event-time allocation reconcile the 09:00 bucket 109

EXHIBIT 30 / layered-architecture
Protection follows the path from edge identity to restored record 116

EXHIBIT 31 / matrix
Access rights differ sharply by role and data product 118

EXHIBIT 32 / timeline
An exercised response contains the incident and preserves safe production 122

EXHIBIT 33 / table
A processing register separates operational data from personal-data obligations 125

EXHIBIT 34 / table
Structured operational evidence requires modest local buffer capacity compared with raw media 132

EXHIBIT 35 / range
Recovery objectives should reflect consequence rather than data volume 135

EXHIBIT 36 / matrix
Portability requires meaning, evidence and operating instructions 137

EXHIBIT 37 / gantt
A recovery exercise proves the whole chain, not only the backup 142

EXHIBIT 38 / matrix
An admission scorecard can reject a technically attractive use case 150

EXHIBIT 39 / waterfall
Lifecycle cost can reverse an apparent pilot return 153

EXHIBIT 40 / flow
Every candidate exits through deploy, redesign, postpone or stop 156

EXHIBIT 41 / table
Vision and forecasting need different evidence before admission 158

EXHIBIT 42 / layered-architecture
Small capabilities keep the experiment reconstructable 164

EXHIBIT 43 / flow
Candidate identity connects the record to review 166

EXHIBIT 44 / bar
Mean error alone favours the peak-period failure 169

EXHIBIT 45 / table
The candidate record separates knowledge from authority 171

EXHIBIT 46 / table
Four threshold families and slice rules prevent aggregate masking 178

EXHIBIT 47 / dumbbell
Candidate improvement must survive comparison with current practice 182

EXHIBIT 48 / gantt
Staged release makes each expansion of authority conditional 185

EXHIBIT 49 / flow
Rollback restores a known operating state and preserves evidence 187

EXHIBIT 50 / matrix
Higher consequence demands more independent release authority 189

EXHIBIT 51 / line
Drift is a diagnostic signal, not a retraining instruction 198

EXHIBIT 52 / table
Six production anomalies produce statistical shift but require opposite remedies 199

EXHIBIT 53 / flow
An alert becomes controlled action through a diagnosed runbook 203

EXHIBIT 54 / matrix
Monitoring ownership follows the failing layer 209

EXHIBIT 55 / timeline
Model status remains visible from candidate to retirement 211

EXHIBIT 56 / range
Time to value includes learning and evidence accumulation 217

EXHIBIT 57 / small-multiples
Adoption, operations and trust can move in different directions 221

EXHIBIT 58 / waterfall
Net value remains after recurring lifecycle work is counted 222

EXHIBIT 59 / table
A benefit register turns observed evidence into a portfolio decision 225

EXHIBIT 60 / gantt
Eighteen months of capability building remain conditional on evidence 230

EXHIBIT 61 / table
The first 90 days produce decisions and evidence, not a transformation slogan 232

EXHIBIT 62 / matrix
No stage opens until its predecessor leaves usable evidence 235

EXHIBIT 63 / keyfigures
Seven capability measures replace a vanity count of dashboards 239

Build the factory and its evidence as one system.

Cameroon's local-processing ambition requires factories whose daily records can support operating, customer and sustainability decisions. GFT's programme evidence base, prepared 31 January 2024, records about 95 percent of cotton fibre exported raw and about 5 percent processed locally, against a 50 percent by 2030 target. These entries were not refreshed in 2026 and their original source dates are unresolved (R06, R07). This paper argues for an industrial data spine, not a large data department. The spine is one governed chain from event capture and operational meaning to quality control, model release, monitoring and value review. It assigns decision rights, defines identifiers and KPIs, preserves lineage, and keeps defined data services available during network interruption without providing power or protective safety. Drawing on 44 opened sources, the paper distinguishes raw events from curated records, separates transport from semantics, and places quality gates before training. [R01] [S28] Light-weight MLOps adds only controls a small team can sustain: versioned data and code, an experiment and model register, human-set thresholds, shadow release, rollback, drift monitoring and practised incident response. Security, privacy and Cameroonian legal screening shape the design. Value is measured against an operating baseline and net of preparation, integration and maintenance. The closing illustrative 18-month roadmap sets decision gates and seven capability measures. [R09] [R14] GFT's Cameroon textile programme supplies the starting point: capture cleanly from the first shift, then earn the right to model.

01 Governance begins with decision rights and a named internal champion, not with a platform purchase. [S02] [S23] [S37]

02 A shared operational vocabulary must precede integration because transport protocols do not supply business meaning. [S05] [S06] [S07] [S38]

03 Data quality is a lifecycle discipline, and upstream defects can compound into delayed downstream failures. [S01] [S03] [S04] [S19]

04 Lineage from source record to dataset, model and decision is the minimum evidence needed for reproducibility and audit. [S09] [S15]

05 Small teams should adopt MLOps in stages, with automation earned by operational need rather than copied from large technology firms. [S16] [S20] [S42]

06 Monitoring should open a diagnosed response, not trigger indiscriminate automatic retraining. [S17] [S21] [S27]

07 OT safety, least privilege, tested recovery and worker-data limits belong in the initial design. [S22] [S25] [S26] [S29] [S30]

08 Offline buffering and portable records protect continuity better than an assumption of permanent connectivity or vendor access. [S32] [S33] [S05] [S08]

09 The business case must count data preparation, integration and maintenance, then compare net value with a pre-model baseline. [S18] [S36] [S39] [S43]

10 GFT's manual data spine is the right starting system, but scale should follow passed ownership, quality, recovery and value gates. [S02] [S04] [S16] [S44]

GREEN FUTURE TECH

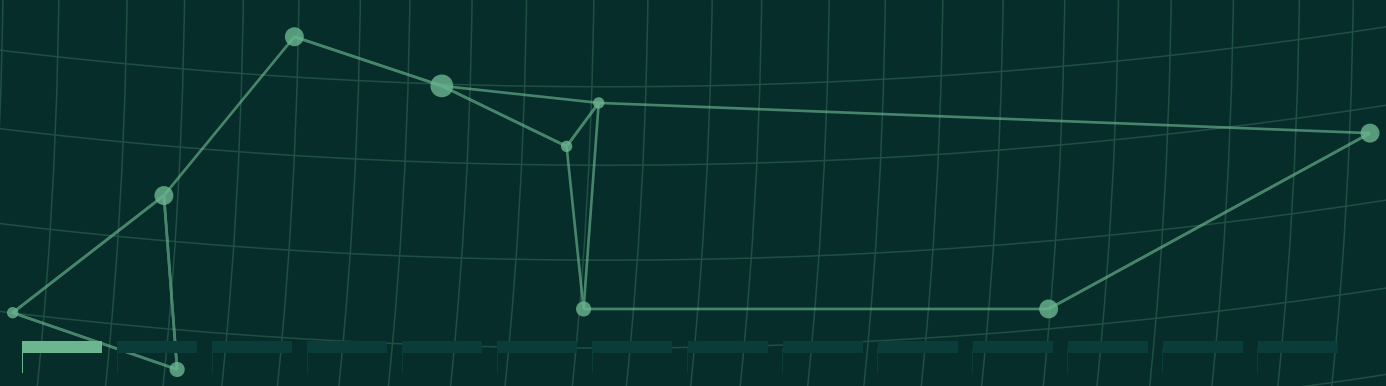
LIMITED / HONG KONG SAR

01

PART I: GOVERN MEANING BEFORE SCALE / THE PROBLEM

A factory becomes reliably data-driven only when its records can be trusted

The binding constraint on industrial AI is often the broken chain between work, record and decision, not the absence of an algorithm.



CHAPTER 01
DECISION TEST

The binding constraint on industrial AI is often the broken chain between work, record and decision, not the absence of an algorithm.

A factory becomes reliably data-driven only when a supervisor, an engineer and a machine-learning pipeline can read the same operational record and know what physical event occurred, who owns its definition, whether the value passed quality checks and how to recover the history if a system fails. [S01] [S02] [S09] [S26] In small and mid-sized manufacturing plants, the binding constraint on industrial analytics is rarely the absence of a statistical model. [S19] [S0101] The constraint is the broken chain between shop-floor work, recorded evidence and management action. [S05] [S08] [S43] Green Future Tech Limited (GFT), Hong Kong SAR, anchors its work in the principle of Warp and Signal: physical production is the warp, and operational measurement is the weft. In GFT's three-part Cameroon textile programme, prepared on 31 January 2024 with source-refresh annotations retained in 2026, the sector baseline records that about 95 percent of Cameroon's cotton fibre is exported raw while about 5 percent is processed locally, against a national industrialisation target of 50 percent local processing by 2030. [R06] [R07] Those baseline and target entries come from the inherited programme register, were not independently refreshed in 2026 and carry unrecorded original publication dates, so they frame the scale of the industrial ambition rather than certifying current mill output. [R06] [R07] Moving from raw-fibre export to competitive yarn, fabric and garment production requires factories that can defend their daily numbers on yield, quality, downtime, order trace and resource use without hiring an enterprise data department. [R06] [S02] [S38] This paper specifies that alternative: an industrial data spine built on named decision rights, explicit operational meaning, enforced quality gates, reversible model releases and measured operational value. [S01] [S04] [S15] [S36]

The urgency of governed records rises sharply when a sector seeks to rebuild domestic processing capacity under tight capital and margin constraints. [R06] [R12] Inherited sector evidence illustrates the cost of operating without reliable utilisation and market alignment: CICAM's reported domestic market share fell from about 85 percent three decades ago to about 5 percent in the inherited register, while a July 2023 report recorded capacity utilisation at 33.7 percent in Garoua, 48.9 percent at CICAM 1 and 10.83 percent at CICAM 2, with underutilisation costing about CFAF 2 billion (XAF 2 billion) in 2020. [R11] [R12] Meanwhile, announced third-party industrial expansion, such as the Camtext SA integrated cotton-to-textile project reported on 12 November 2025 at US$300 million (180 billion XAF) with a planned processing volume of 12,000 tonnes a year, remains an announced third-party investment target reaching optimal capacity around 2033 rather than realised factory performance. [R09] [R10] Whether an operating team is rehabilitating an existing mill or commissioning a new garment line, capital equipment alone does not produce control. [S05] [S36] As **Exhibit 1**, p. 16 summarises, the gap between the inherited 5 percent local-processing baseline and the 50 percent national target by 2030 places a high penalty on untrusted operational evidence. [R06] [R07]

EXHIBIT 01 / PERCENT OF COTTON FIBRE

Local-processing ambition raises the cost of untrusted evidence

95%

Raw cotton fibre exported

Approximate inherited sector baseline (R06; original date unrecorded)

5%

Cotton fibre processed locally

Approximate inherited sector baseline (R06; original date unrecorded)

50%

National processing target

SND30 local-processing target by 2030 (R07; policy target, not realised output)

SOURCED [R06] [R07] Inherited GFT programme evidence base prepared 31 January 2024. R06 and R07 were not refreshed in 2026 and their original publication dates are unrecorded. Figures describe sector-level baseline and policy ambition, not a current GFT factory measurement or causal productivity claim.

Exhibit 1, p. 16 establishes the industrial context and its evidentiary limits. [R06] [R07] It shows why local value retention is a strategic priority in Cameroon, but a macroeconomic target cannot tell a spinning supervisor why roving breaks rose on the afternoon shift or tell a garment director whether an automated fabric-inspection camera is reducing customer claims. [R06] [R07] [S38] Answering those operational questions requires a governed chain from the shop floor to the management review table. [S05] [S38]

1.1

Data debt starts on the first shift

Data debt is the future correction, reconciliation and redesign cost that a factory accumulates when it captures operational numbers with undefined fields, inconsistent identifiers and no named owner. [S01] [S02] [S18] Sculley and colleagues showed that real-world machine-learning systems incur heavy ongoing maintenance costs because models depend on unstable upstream data feeds, undeclared consumers and improvised correction scripts. [S18] In a factory, that obligation starts on the first shift when an operator writes a bundle count on a clipboard, a technician logs a fault in a chat message or a supervisor edits a spreadsheet formula to force weekly output to match the packing list. [S01] [S18] [S38]

1.2 UPSTREAM CAPTURE DEFECTS COMPOUND SILENTLY ACROSS THE LIFECYCLE

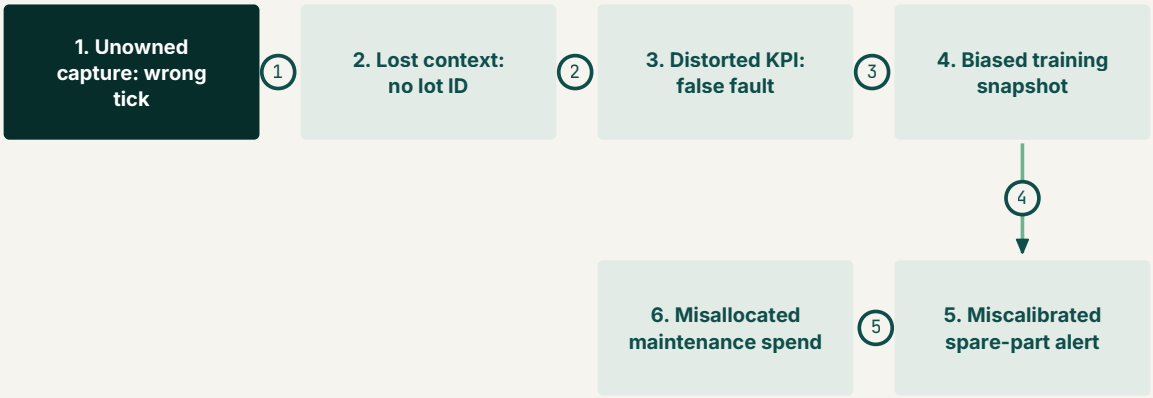
When initial capture lacks explicit rules for syntax, meaning and provenance, defects multiply across reporting and analytical layers. [S01] [S09] [S19] Sambasivan and colleagues studied high-stakes artificial-intelligence deployments and defined data cascades as compounding events caused by upstream data problems that produce negative downstream consequences. [S19] Among the high-stakes AI practitioners they interviewed, 92 percent reported experiencing at least one data cascade, noting

that such defects were typically invisible at capture, delayed in impact and avoidable through deliberate data practices. [S19] That 92 percent figure describes the interviewed sample of high-stakes AI practitioners rather than a prevalence estimate for manufacturing firms or small and medium-sized enterprises (SMEs), yet the mechanism it identifies operates directly on factory floors. [S19] Treating data entry as clerical overhead rather than engineered measurement hides upstream omissions until a management decision or model fails in production. [S04] [S19]

Consider how an unowned capture defect propagates through a garment plant, as traced in the illustrative sequence of **Exhibit 2**, p. 17. [S18] [S19] Suppose a sewing line loses forty minutes of output because a batch of dyed fabric arrives from finishing with uneven shade across the roll width, forcing operators to re-sort cut panels at the line feeder (an illustrative shop-floor scenario). If the line tally sheet offers only two downtime checkboxes, mechanical breakdown and operator absence, the clerk ticks mechanical breakdown to avoid leaving a mandatory field blank (illustrative). That forced misclassification corrupts every downstream stage. [S18] [S19]

EXHIBIT 02 /

A small capture defect can compound into a wrong management action



1. No schema check
2. Unverified aggregation
3. Silent ingestion
4. Corrupted ground truth
5. Unchallenged threshold

ILLUSTRATIVE [S18] [S19] Illustrative sewing-line event traced from shop-floor capture through context, KPI calculation, training snapshot, model output and management decision. Failure mechanics are derived from technical-debt entanglement in S18 and upstream data cascades in S19.

As **Exhibit 2**, p. 17 illustrates, the error does not correct itself. [S18] [S19] At the second node, the link between the sewing stoppage and the upstream fabric lot disappears because the tally sheet records only the sewing line number. [S08] [S09] At the third node, the weekly key performance indicator (KPI) calculation charges the lost forty minutes against equipment availability under ISO 22400-2 time elements rather than material starvation or quality hold. [S38] At the fourth and fifth nodes, an analyst

training a predictive-maintenance classifier on twelve months of logs ingests those mislabelled stoppages as mechanical failures. [S04] [S18] [S19] At the sixth node, management approves spare-parts expenditure for sewing heads while the dye-lot shade problem remains untouched. [S18] [S19] What **Exhibit 2**, p. 17 cannot establish is the exact frequency of such errors in a given mill before an audit; rather, it shows why ISO/IEC 5259-3 requires data-quality management across the full analytics and machine-learning lifecycle rather than as a late cleaning step. [S04] [S19]

1.3 INFORMAL SPREADSHEET WORKAROUNDS MASK STRUCTURAL GAPS

In the early phases of GFT's Cameroon textile programme, manual spreadsheets and shift logs serve as the starting system of record. Structured manual logs are an appropriate choice when capital and connectivity are constrained, provided those sheets enforce field definitions, identifier stability and version control. [S01] [S04] [S16] Data debt arises when workbooks mutate privately across shifts: a morning supervisor inserts a column, an afternoon supervisor merges cells for a split order and a clerk pastes values over formulas. [S01] [S09] [S18] Reconciling those files at month-end consumes engineering hours needed for process improvement, and any dashboard built on top inherits their contradictions. [S18] [S0101]

1.4

Disconnected logs fracture one production event

A manufacturing plant is a physical system in which material, machines, people and orders interact simultaneously, yet most small factories record those interactions in separate silos. [S05] [S08] When a disruption occurs on the floor, it leaves partial footprints in three or four departmental logs that cannot be joined without manual investigation. [S05] [S08] [S0101]

1.5 ONE SEWING STOPPAGE PRODUCES FOUR IRRECONCILABLE ACCOUNTS

Consider a single operational event reconstructed across four typical factory logs: the sewing-line hourly board, the maintenance technician's note, the quality inspector's defect sheet and the commercial order tracker (an illustrative reconstruction grounded in IEC 62264-1 integration boundaries and GS1 EPCIS 2.0 visibility events). [S05] [S08] Suppose an overlock machine on Line B begins skipping stitches at 10:12 on a Tuesday while running a cotton polo-shirt order. [S05] [S08]

1. **The line supervisor's hourly sheet** records output only on the clock hour, noting that Line B produced 45 garments between 10:00 and 11:00 against an illustrative target of 80 garments, with a note saying "machine trouble" and no asset identifier or exact stop timestamp. [S05] [S38]
2. **The maintenance technician's log** records a repair on asset `OVL-14` from 10:25 to 10:50 for a worn looper, without recording which line `OVL-14` was mounted on or which order was interrupted. [S05] [S08]
3. **The quality inspector's sheet** logs 18 rejected polo shirts under style text "Navy Polo" during the morning shift for seam skip, without recording the bundle identifier, machine identifier or whether the 18 pieces went to rework or scrap. [S08] [S38]

4. **The commercial order spreadsheet** shows Customer Order ORD-2026-114 short by 18 pieces at packing two days later, logging "sewing short-ship" with no link to OVL-14 or the morning quality hold. [S05] [S08]

Exhibit 3, p. 19 contrasts these four departmental fragments with the unified record structure required by enterprise-control integration and visibility-event standards. [S05] [S08] [S38]

EXHIBIT 03 / RECORD ATTRIBUTES CAPTURED

Disconnected shop-floor logs split one stoppage into four irreconcilable records

Record source	Time grain	Object identifier	Recorded outcome	Missing link (S05, S08, S38)
Line hourly sheet	Clock hour (10:00 to 11:00)	Line B only	35-unit shortfall	No asset ID, lot ID or exact stop duration
Maintenance note	Repair window (10:25 to 10:50)	Machine OVL-14	Looper replaced	No line assignment, wait time or order ID
Quality tally sheet	Morning shift aggregate	Style description text	18 seam-skip defects	No bundle ID, machine ID or rework vs scrap state
Order progress tracker	Packing date (two days later)	Order ORD-2026-114	18-unit short-ship	No link to quality hold or machine fault event
Governed spine event	Event + ingestion timestamps	Line B + OVL-14 + Lot + Order	State transitions + 18 rework/scrap	Reconciles time, object, location and business step

DERIVED [S05] [S08] [S38] Derived comparison applying IEC 62264-1 enterprise-control interface domains (S05), GS1 EPCIS 2.0 visibility event dimensions (S08) and ISO 22400-2 KPI time and quantity elements (S38) to a single shop-floor stoppage.

As **Exhibit 3**, p. 19 shows, none of the four logs is falsified, yet the factory cannot reconstruct the event from them. [S05] [S08] The maintenance note shows 25 minutes of active repair (10:25 to 10:50), whereas the line lost 38 minutes between the stop at 10:12 and restart clearance at 10:50 (illustrative timestamps). [S38] Because the 13-minute wait from 10:12 to 10:25 is unrecorded in maintenance and merged into a clock hour on the line sheet, management cannot separate technician response delay from active wrench time, nor can quality prove whether the 18 defective pieces came from OVL-14 just before 10:12. [S08] [S38] What **Exhibit 3**, p. 19 does not prescribe is the transport protocol or database engine; Chapter 3 specifies the vocabulary contract and Chapter 5 the minimum architecture. [S05] [S07] [S08]

1.6 INTERFACE STANDARDS AND VISIBILITY EVENTS SUPPLY THE MISSING GRAMMAR

Two international standards explain how a small team can repair these broken links without buying a monolithic software suite. [S05] [S08] First, IEC 62264-1 (the ISA-95 standard) defines models and terminology for the interfaces between enterprise business systems and manufacturing operations and control systems. [S05] It organ-

ises work schedules, production performance, equipment, material lots and personnel into explicit interface domains so that an order identifier and a machine identifier meet at a defined boundary. [S05]

Second, the GS1 EPCIS 2.0 standard enables disparate applications to create and share visibility event data within and across enterprises. [S08] Rather than overwriting a status cell, an EPCIS-aligned event records what object or lot was involved, when the event occurred and was captured, where it happened down to the line or station, and why it occurred under a defined business step and disposition. [S08] Adopting that four-part event grammar even in initial structured sheets links the stoppage, repair, quality hold and order variance into one inspectable chain. [S05] [S08] [S09]

1.7

Small teams need decision rights more than platforms

When production, quality and finance numbers conflict in a weekly meeting, management is often tempted to buy a dashboard tool or an enterprise resource planning (ERP) module. [S02] [S34] [S0101] Software procurement cannot fix an accountability vacuum. [S02] [S23] If nobody holds the right to decide whether reworked garments count inside hourly output, a new platform merely displays the dispute faster and at higher cost. [S02] [S38]

1.8 ROLE ASSIGNMENT PRECEDES TOOL SELECTION

ISO 8000-150:2022 specifies the key considerations for organisations establishing roles and responsibilities for data-quality management. [S02] Its core principle is organisational: data quality improves only when specific roles are accountable for defining data requirements, monitoring conformance and executing corrections. [S02] Similarly, the NIST Cybersecurity Framework 2.0 states that governance activities are critical for incorporating risk management into broader enterprise strategy. [S23] A factory must therefore assign who owns domain meaning, who stewards daily capture quality and who holds technical custody before comparing software vendors. [S02] [S23]

In a small factory, assigning decision rights does not create a separate bureaucracy. [S02] [S0101] NIST's smart-manufacturing programme observes that a typical SME cannot afford a dedicated data-analytics specialist on staff and that commercial tools remain difficult to integrate with shop-floor routines. [S0101] Governance must therefore sit inside existing operating roles: the production manager owns output and time definitions, the quality lead owns defect taxonomies and dispositions, the maintenance lead owns asset hierarchies and fault codes, and the commercial lead owns order and lot identifiers. [S02] [S05] [S38] Chapter 2 formalises this six-role charter with conflict-of-interest rules where one person holds multiple roles. [S02] [S11] [S23]

1.9 A NAMED INTERNAL CHAMPION DETERMINES WHETHER ADOPTION SURVIVES

Paper role descriptions fail if daily adoption is left to external consultants. [S37] In a May 2025 World Bank case study of SME digital technology adoption in Viet Nam, firms that successfully adopted the platform assigned an internal point person who

championed adoption within the firm. [S37] That champion coordinates training, follows up on late logs, escalates cross-domain definition disputes and protects line stewards' time. [S02] [S37]

Designers must also avoid assuming that integrated enterprise software already exists on the ground. [S34] In the OECD's June 2026 Western Balkans Enterprise Survey, tools supporting integrated operations, such as cloud computing, ERP and customer relationship management (CRM) systems, were used by only 10 to 20 percent of surveyed SMEs, with supplier-customer systems even rarer. [S34] That 10 to 20 percent finding describes the surveyed Western Balkans SME sample rather than a Cameroon estimate, but it warns against assuming that an SME possesses a clean ERP backbone ready for machine learning. [S34] [S42] As SCITEPRESS Industry 4.0 research notes, the starting technical infrastructure and target degree of automation strongly govern MLOps complexity. [S42]

> **GFT NOTE**
>
> **Appoint domain owners and an internal champion before buying software.** In GFT's Cameroon textile programme, no dashboard, camera or forecasting tool should be procured until the factory director names one internal adoption champion and assigns explicit ownership for production, quality, maintenance, order and sustainability records. [S02] [S37] When two managers present conflicting numbers in a review meeting, the resolution is a versioned update to the KPI dictionary signed by the domain owner, not a second spreadsheet. [S02] [S38]

1.10

Reuse turns preparation from a tax into infrastructure

When a factory treats data as a by-product of isolated projects, every new question triggers another custom extraction, cleaning script and identifier mapping. [S10] [S39] [S0101] One effort builds a sheet for overall equipment effectiveness (OEE), a second pulls defect tallies for a camera trial and a third extracts utility bills for carbon reporting, each redefining shift boundaries, rework and good output. [S10] [S38] [S39]

1.11 APPLICATION-SPECIFIC PIPELINES MULTIPLY PREPARATION COSTS

In a 2025 NIST workshop report on AI and scaled data sharing in semiconductor manufacturing, manufacturers reported that data processing for a single application can reach up to 70 percent of total cost. [S39] That 70 percent figure is a manufacturer-reported upper-bound warning from semiconductor fabrication rather than a universal SME average, yet it highlights the penalty of building separate pipelines for every application. [S39] NIST's smart-manufacturing research confirms that substantial time goes into data acquisition and formatting because analytics tools use proprietary interfaces and specialised data formats. [S0101]

The alternative is to capture each operational event once with governed identity, timestamps, units and quality states, catalogued under a standard metadata model such as the W3C Data Catalog Vocabulary (DCAT) Version 3, and publish controlled

views to multiple consumers. [S08] [S09] [S10] [S0102] Across the GFT white-paper library, one governed record spine feeds four downstream uses without re-cleaning the floor logs:

- **Line performance and OEE management (WP02, WP15):** uses validated availability, cycle-time and good-count events under ISO 22400-2 definitions for shift reviews and energy intensity. [S38]
- **Automated visual inspection (WP03):** uses canonical defect codes, roll identifiers and inspector dispositions to test computer-vision cameras against manual inspection baselines (where inherited literature records manual accuracy of 60 to 70 percent against computer vision above 90 percent across more than 40 defect types). [R29] [S04] [S08]
- **Demand and order forecasting (WP04):** uses clean order, dispatch and shortship records once history accumulates, preventing stockouts from distorting demand signals. [R27] [S04] [S43]
- **Traceability and carbon accounting (WP05, WP16):** uses lot genealogy and utility intervals across iterative data-quality improvement cycles. [S08] [S09] [S41]

WORKED EXAMPLE: DUPLICATED PROJECT EXTRACTS VERSUS ONE GOVERNED SPINE

Consider an illustrative arithmetic comparison over a 52-week year in a 400-worker garment plant running four analytical workflows: weekly OEE, quality defect analysis, buyer traceability packs and monthly energy-water intensity. All hours and wage figures below are illustrative method assumptions, not measured GFT factory results.

Assumed inputs (illustrative): 1. **Fragmented pipelines (four separate extracts):** Each workflow independently cleans shift sheets, reconciles lot and asset codes, and resolves meeting disputes. At an illustrative 6 hours per week of preparation plus 2 hours of dispute reconciliation per workflow, 4 workflows consume `4 × (6 + 2) = 32 hours per week`, or `32 × 52 = 1,664 hours per year` (illustrative). 2. **Governed spine (capture once, reuse four times):** Defining canonical identifiers, the KPI dictionary and ingestion checks takes an illustrative 160 hours upfront. Ongoing stewardship takes an illustrative 8 hours per week for the shared spine plus 1 hour per week for each of the 4 views (`8 + 4 = 12 hours per week`, or `624 hours per year`, illustrative). Year 1 totals `624 + 160 = 784 hours`, and Year 2 onward requires `624 hours per year` (illustrative).

Resulting arithmetic (illustrative): The governed spine saves `1,664 - 784 = 880 hours` in Year 1 (47.1 percent reduction, illustrative) and `1,664 - 624 = 1,040 hours per year` from Year 2 onward (62.5 percent reduction, illustrative). Valued against the inherited Cameroon cross-sector average salary benchmark of about US$359 a month across a 173-hour standard schedule (about US$2.08 an hour, alongside the statutory SMIG minimum wage of XAF 43,969 a month from 23 February 2024 and reference rates near XAF 600 per US$1), 1,040 hours equals roughly US$2,160 (about XAF 1.30 million) in direct staff time each year, while freeing 20 scarce engineering hours every week to fix physical bottlenecks (illustrative). [R14] [R15] [R16] [R32] [S18]

1.12

Trust includes safety, legality and worker dignity

A factory record is not trustworthy simply because its schema validates. [S22] [S29] [S30] An architecture that disrupts machine reliability, exposes control networks, collects excessive personal telemetry on workers or breaches national data-protection law creates severe operational exposure. [S22] [S29] [S31]

1.13 OPERATIONAL TECHNOLOGY SAFETY OVERRIDES DATA CONVENIENCE

Factory systems interact with rotating machinery, dyeing vessels, boilers and switchgear. [S22] NIST Special Publication 800-82 Revision 3 guides operational technology (OT) security around performance, reliability and safety requirements. [S22] Data collection and model inference must remain subordinate to physical process safety: no gateway, broker or model may write unvalidated commands into safety-critical control loops, and any data-layer failure must leave machines operating safely under local control. [S22] [S33] During network outages, edge devices use local store-and-forward buffering to retain event records without disturbing shop-floor control. [S22] [S33]

1.14 DATA MINIMISATION AND HUMAN REVIEW PROTECT THE WORKFORCE

Shop-floor logs often touch operator sign-ins, bundle scans, skills matrices and camera feeds. [S29] [S30] Two rules govern these records. [S29] [S30] First, Regulation (EU) 2016/679 (GDPR, Recital 39) states the data-minimisation principle: personal data should be adequate, relevant and limited to what is necessary for the stated purpose. [S29] Line balance, availability, yield and energy metrics require station, line, shift or role identifiers rather than personal identity; where personal records are required (such as high-voltage maintenance certification), access follows least privilege under NIST SP 800-53 Revision 5. [S25] [S29] [S38]

Second, Principle 5.5 of the International Labour Office (ILO) 1997 code of practice states that decisions concerning a worker should not be based solely on automated processing of that worker's personal data. [S30] While the ILO code is non-binding guidance, GFT enforces it as an operating rule: no productivity score, defect tally or camera output may automatically trigger discipline, pay changes or dismissal without informed human review of material, machine and line context. [S11] [S30]

1.15 CAMEROONIAN STATUTORY COMPLIANCE REQUIRES QUALIFIED LOCAL COUNSEL

Factories in Cameroon must also address Law No. 2024/017 of 23 December 2024 relating to personal data protection in Cameroon. [S31] Section 19(1) of Law No. 2024/017 provides that the processing of personal data shall be subject to prior authorisation by the Personal Data Protection Authority. [S31]

LEGAL SCREENING AND REGULATORY QUALIFICATION

This paper provides engineering governance guidance, not legal advice. While Section 19(1) of Cameroon's Law No. 2024/017 of 23 December 2024 sets out the statutory prior-authorisation rule, the current operational status of the Personal Data Protection Authority, implementing decrees and filing procedures must be verified with qualified Cameroonian legal counsel before collecting, linking or exporting identifiable personal data. [S31] No dataset containing personal data may enter cloud analytics or model training until its legal basis and authorisation status are recorded in the processing register (Chapter 7). [S29] [S31]

1.16

GFT's position is to govern the record before scaling the model

Across GFT's Cameroon textile programme, the sequencing principle is strict: capture cleanly from the first shift, establish one governed vocabulary and admit machine learning only after records pass explicit readiness gates. [S02] [S04] [S16]

1.17 FOUR ADMISSION TESTS SEPARATE A DATA SPINE FROM A DATA PILE

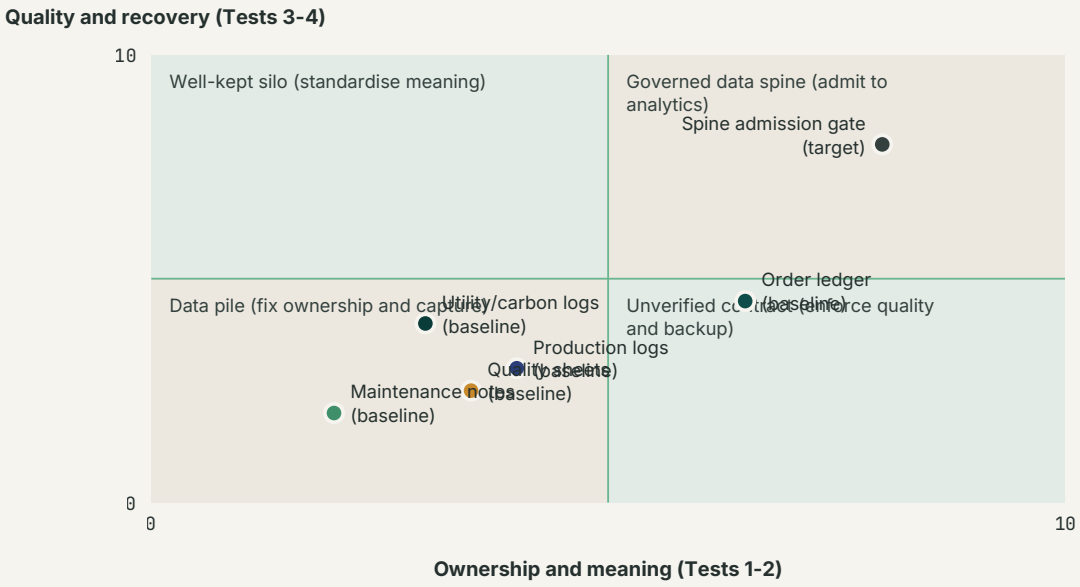
GFT translates international standards into four admission tests that separate a governed data spine from an uncurated data pile: [S01] [S02] [S05] [S09] [S26]

1. **Test 1: Named ownership and decision rights (ISO 8000-150, NIST CSF 2.0).** Every domain has one accountable business owner who approves purpose, definitions, access and retention, supported by a line steward. [S02] [S23]
2. **Test 2: Explicit operational meaning and contract (IEC 62264-1, GS1 EPCIS 2.0, ISO 22400-2).** Every record uses stable site, line, asset, lot and order identifiers, separates event time from ingestion time, and binds KPIs to versioned formulas and units. [S05] [S08] [S38]
3. **Test 3: Enforced quality gates and lineage (ISO 8000-1, ISO/IEC 25012, ISO/IEC 5259-3, W3C PROV-O).** Capture and ingestion check validity, completeness, consistency, timeliness and uniqueness; quarantine failures; flag missing or estimated states; and preserve provenance to raw events. [S01] [S03] [S04] [S09]
4. **Test 4: Tested recoverability and portability (NIST SP 800-82 Rev 3, CISA #StopRansomware, edge buffering).** Raw records, schemas, master data and KPI logic have offline encrypted backups with tested restoration and local store-and-forward buffering during disconnection. [S22] [S26] [S33]

Exhibit 4, p. 25 maps five factory record domains across two composite axes derived from these four tests: ownership and meaning (Tests 1 and 2) on the horizontal axis, and quality, lineage and recovery (Tests 3 and 4) on the vertical axis. [S01] [S02] [S05] [S09] [S26]

EXHIBIT 04 / DERIVED GOVERNANCE READINESS SCORE (0 TO 10 SCALE)

Four admission tests separate a data spine from a data pile



DERIVED [S01] [S02] [S05] [S09] [S26] Derived readiness framework scoring five factory record domains across two composite axes: Ownership and Meaning (Tests 1 and 2, from S02 and S05; 0 = unowned free text, 10 = contracted owner, IDs and KPI formula) and Quality, Lineage and Recovery (Tests 3 and 4, from S01, S09 and S26; 0 = unchecked/unbacked, 10 = gated, provenance-linked and restoration-tested). Point coordinates represent an illustrative pre-governance factory baseline versus the minimum admission threshold (score >= 7 on both axes) required to enter the governed spine.

In **Exhibit 4**, p. 25, a typical pre-governance baseline leaves production, quality, maintenance and utility logs in the lower-left "Data pile" quadrant while order ledgers sit near the middle (illustrative baseline coordinates on the derived 0 to 10 rubric). [S01] [S02] [S05] Training models on lower-left records imports unowned definitions and unflagged missingness directly into predictions. [S04] [S18] [S19] Reaching the upper-right "Governed data spine" gate (8.0, 8.0) requires passing all four admission tests during an on-site audit using the checklists in Chapters 2 through 8. [S01] [S02] [S09] [S26]

1.18 VISION AND FORECASTING REMAIN SUBORDINATE TO RECORD DISCIPLINE

These four tests also govern the sequence between bought computer vision (WP03) and demand forecasting (WP04): [S04] [S16] [S44]

- **Purchased fabric and garment vision (WP03):** Pre-trained optical inspection can be trialled early because detection relies on image physics rather than years of order history, but it still requires a canonical defect taxonomy (Test 2), local validation against human inspectors on Cameroonian weaves and shades (Test 3), human-set thresholds (NIST AI RMF 1.0) and a manual inspection fallback (Test 4). [R29] [S04] [S11] [S22]
- **Demand and order forecasting (WP04):** Although McKinsey's *The State of Fashion 2025* reports that AI at scale can cut forecast errors by 20 to 50 percent and inventory costs by 35 to 40 percent, those results depend on multi-year stockout-corrected order histories. [R27] A new or restructuring mill must first accumulate

governed order, dispatch and short-ship records and compare any forecaster against a simple moving-average baseline before release (Chapter 9). [R12] [S04] [S43] [S44]

1.19

The paper judges success by better decisions

An industrial data spine is judged by whether operational decisions improve, not by gigabytes stored, MQTT message rates or dashboard counts. [S07] [S36] [S43]

1.20 CYCLE TIME FROM ABNORMAL EVENT TO VERIFIED ACTION IS THE PRIMARY TEST

World Bank guidance on MSME digitalisation stresses embedding monitoring and evaluation throughout design and implementation to track impact and improve learning. [S36] On the shop floor, the primary operational measure is the cycle time from an abnormal physical event to a verified corrective action across five steps: [S10] [S36] [S38]

1. **Detection latency:** elapsed time from physical deviation to timestamped capture. [S08] [S38]
2. **Validation and contextualisation:** automated checking and linking to order, lot and asset identifiers. [S04] [S05] [S08]
3. **Diagnostic clarity:** provenance showing whether the fault lies in machine, material, method or data feed. [S09] [S17] [S21]
4. **Action execution:** bounded response by the named owner or supervisor. [S02] [S11] [S15]
5. **Verified effect:** confirmation in subsequent records that the KPI returned to its baseline band. [S36] [S38]

1.21 TECHNICAL ACCURACY WITHOUT OPERATIONAL FIT FAILS THE LIFECYCLE GATE

In manufacturing data mining, the CRISP-DM lifecycle runs through six phases: business understanding, data understanding, data preparation, modelling, evaluation and deployment. [S43] A model with high offline accuracy in phase four still fails in evaluation or deployment if alerts arrive after fabric is cut, if operators cannot act on the output or if data-preparation and maintenance costs exceed gross savings. [S18] [S39] [S43] [S0101] When resources, skills, costs or expected benefits are not met, postponing or stopping an AI project is a valid governance decision. [S44]

1.22 HOW THE FOURTEEN CHAPTERS OF WP14 BUILD THE SPINE

The fourteen chapters of this paper build that governed chain across four parts:

- **Part I (Chapters 1 to 4) governs meaning before scale:** assigning decision rights (Chapter 2), defining the operational vocabulary and contracts (Chapter 3) and enforcing quality gates and PROV-O lineage (Chapter 4). [S02] [S03] [S04] [S05] [S08] [S09] [S38]
- **Part II (Chapters 5 to 8) builds the minimum durable spine:** separating capture, raw evidence and curated views alongside WP18 (Chapter 5), governing time, identity, genealogy and state (Chapter 6), embedding OT safety, privacy and

Cameroonian legal screening (Chapter 7), and proving offline continuity, restoration and portability (Chapter 8). [S06] [S10] [S15] [S22] [S25] [S26] [S27] [S29] [S30] [S31] [S32] [S33]

- **Part III (Chapters 9 to 12) controls models as production change:** enforcing model admission and baselines (Chapter 9), reproducible training and registries (Chapter 10), human-threshold validation and shadow release with rollback (Chapter 11), and diagnostic drift monitoring (Chapter 12). [S11] [S13] [S15] [S16] [S17] [S21] [S43] [S44]
- **Part IV (Chapters 13 and 14) proves value and scales deliberately:** measuring behavioural adoption and net lifecycle benefit (Chapter 13) and sequencing an 18-month gated roadmap with seven monthly capability metrics (Chapter 14). [S18] [S35] [S36] [S39] [S42] [S44]

A factory does not become data-driven by collecting more numbers; it becomes data-driven when every number it acts upon can defend its origin, its meaning and its fitness for the decision at hand.

CHAPTER SOURCE NOTES

S01 International Organization for Standardization. 2022.
S02 International Organization for Standardization. 2022.
S09 World Wide Web Consortium. 2013-04-30.
S26 CISA, MS-ISAC, NSA and FBI. 2023.
S19 Association for Computing Machinery, CHI 2021. 2021.
S0101 National Institute of Standards and Technology. 2025-03-26.
S05 International Electrotechnical Commission. 2013.
S08 GS1 AISBL. 2022.
S43 Frontiers in Artificial Intelligence. 2021-06-14.
R06 Cameroon Trade Hub. Publication date unrecorded; inherited source-text check.
R07 Ecofin Agency. Publication date unrecorded; inherited source-text check.
S38 International Organization for Standardization. 2014.
S04 ISO/IEC. 2024.
S15 Google Cloud. 2024-08-28.
S36 World Bank. 2023.
R12 Business in Cameroon. Publication date unrecorded; inherited source-text check.
R11 Business in Cameroon; MMI News. Publication date unrecorded; inherited source-text check.
R09 Ecofin Agency. Publication date unrecorded; inherited source-text check.
R10 Ecofin Agency. Publication date unrecorded; inherited source-text check.
S18 Advances in Neural Information Processing Systems. 2015.
S16 Microsoft Learn, Azure Architecture Center. 2026-06-03.
S07 OASIS Open. 2019-03-07.
S34 OECD. 2026-06-26.
S23 National Institute of Standards and Technology. 2024-02-26.
S11 National Institute of Standards and Technology. 2023-01-26.

S37 World Bank. 2025-05-20.

S42 SCITEPRESS, IN4PL 2022. 2022.

S10 National Institute of Standards and Technology. 2016-02-23.

S39 National Institute of Standards and Technology. 2025.

S0102 World Wide Web Consortium. 2024-08-22.

R29 MDPI; industry inspection literature. Publication date unrecorded; inherited source-text check.

R27 McKinsey, The State of Fashion 2025. Publication date unrecorded; inherited source-text check.

S41 GHG Protocol, World Resources Institute and World Business Council for Sustainable Development. 2013.

R14 WageIndicator. Publication date unrecorded; inherited source-text check.

R15 wage.is. Publication date unrecorded; inherited source-text check.

R16 WageIndicator. Publication date unrecorded; inherited source-text check.

R32 Wise; XE. Publication date unrecorded; inherited source-text check.

S22 National Institute of Standards and Technology. 2023.

S29 European Parliament and Council of the European Union. 2016-04-27.

S30 International Labour Office. 1997.

S31 Presidency of the Republic of Cameroon. 2024-12-23.

S33 Microsoft. 2024-06-10.

S25 National Institute of Standards and Technology. 2020-12-10.

S03 ISO/IEC. 2008.

S44 The Journal of Technology Transfer, Springer Nature. 2024-07-27.

S17 Microsoft Learn, Azure Machine Learning. 2026-01-27.

S21 arXiv, author manuscript of an IEEE Transactions on Knowledge and Data Engineering article. 2020-04-13.

S06 OPC Foundation. 2025-10-22.

S27 National Institute of Standards and Technology. 2025.

S32 European Parliament and Council of the European Union. 2023-12-13.

S13 ISO/IEC. 2023.

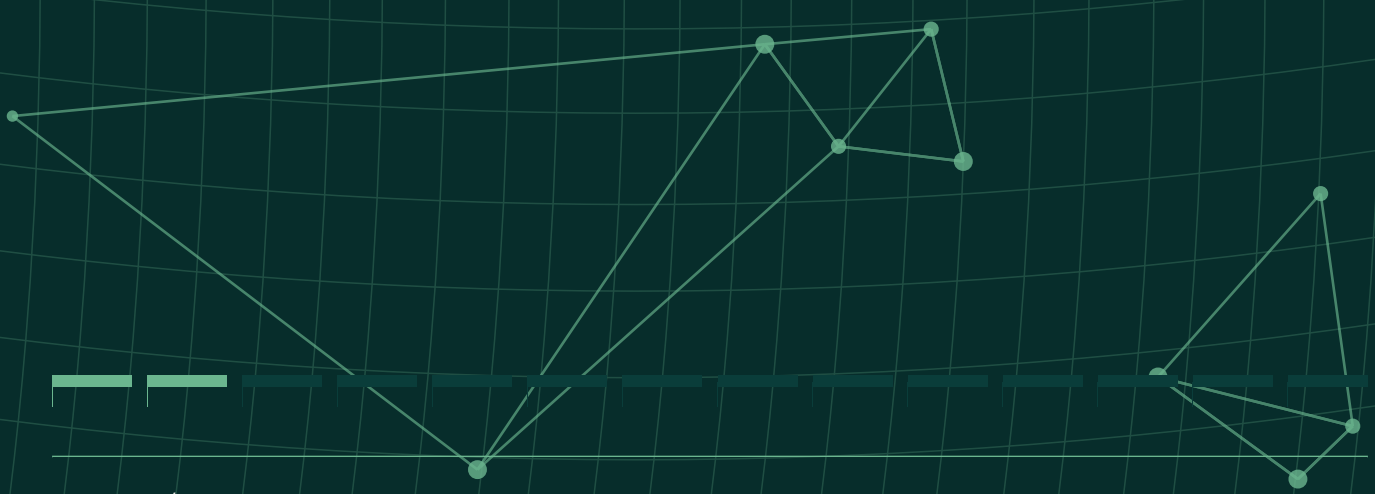
S35 OECD. 2022-11-21.

02

PART I: GOVERN MEANING BEFORE SCALE / OPERATING OWNERSHIP

Named decision rights keep a small data team accountable

A workable governance model fits into existing factory roles and makes every definition, exception and release somebody's explicit decision.



CHAPTER 02
DECISION TEST

A workable governance model fits into existing factory roles and makes every definition, exception and release somebody's explicit decision.

A small factory does not need a standalone data department to govern its records and models; it needs named decision rights attached to existing operational roles. [S02] When a mill or garment plant records production, quality, maintenance and order events, ambiguity over who may define a metric, alter a schema, grant access, quarantine a broken feed or approve a model release corrupts the evidence base. [S01] In GFT's Cameroon textile programme, whose evidence register prepared 31 January 2024 records about 95 percent of cotton fibre exported raw and about 5 percent processed locally against a 50 percent local-processing target by 2030 (inherited entries not refreshed in 2026 with unrecorded original publication dates, stating strategic ambition rather than current factory output), digitalisation exists to make domestic processing reliable and investable. [R06] [R07] Meeting that challenge with a lean team requires one business owner, one domain steward and one technical custodian for every critical dataset, indicator and model, with explicit escalation into existing factory meetings. [S02] [S23]

Without explicit decision rights, data work defaults to convenience. [S19] An engineer modifies a spreadsheet column for a new style, an analyst excludes rework from first-pass yield to tidy a slide, or a vendor updates an inspection threshold during maintenance without recording the change. [S18] [S38] Such informal edits break provenance between physical work and management action. [S09] International data-quality, cybersecurity and AI risk standards require organisations to establish clear roles and responsibilities before relying on technical controls, keeping technical choices subordinate to operational risk appetite. [S02] [S11] [S23] [S0201]

This chapter translates those standards into an operating charter for small and mid-sized factories. [S02] [S14] It assigns one owner to each of six factory domains, separates ownership from stewardship and custody, protects line stewards with a weekly time budget and quarantine authority, subjects definition changes and model releases to versioned approval and independent challenge, and routes cross-domain decisions through WP02's management cadence. [S01] [S11] [S23] [S36] [S38]

2.1

Every critical domain receives one accountable owner

Accountability fails whenever a dataset is owned by "everyone" or left to whichever technician administers the database. [S02] Under ISO 8000-150:2022, establishing appropriate roles and responsibilities is foundational for data quality management. [S02] In a small or mid-sized textile and garment plant, creating new executive posts for data governance is neither affordable nor necessary; instead, the factory assigns data ownership to the existing heads of six bounded operational domains: production, quality, maintenance, orders, people and sustainability. [S02] [S05]

2.2 SIX OPERATIONAL DOMAINS MATCH THE PHYSICAL BOUNDARIES OF FACTORY WORK

Dividing the data spine into six domains aligns governance with the interfaces between enterprise planning and shop-floor execution in IEC 62264-1 (ISA-95). [S05] Each domain covers a distinct family of operational events, master records and management decisions: [S02] [S05]

- **Production:** Assigned to the production manager, covering shift schedules, machine states, hourly line output, work-in-progress balances, lot splits and downtime codes. [S05] [S38]
- **Quality:** Assigned to the quality assurance lead, covering incoming fibre and fabric inspection, defect logs, laboratory results, rework routing, quarantine holds, customer claims and inspection-model validation records. [S04] [S08]
- **Maintenance:** Assigned to the chief maintenance engineer, covering the asset register, preventive schedules, fault codes, work orders, spare-parts use and operational-technology network boundaries. [S05] [S22]
- **Orders and commercial:** Assigned to the planning or commercial manager, covering customer orders, style specifications, bills of materials, delivery dates, dispatch events and sales history for forecasting. [S05] [S08]
- **People and attendance:** Assigned to the human-resources manager, covering shift rosters, skill matrices, training logs and attendance records, under strict purpose limitation and legal screening. [S29] [S30] [S31]
- **Sustainability and utilities:** Assigned to the utilities and environmental lead, covering electricity, steam, water and air sub-metering, chemical inventories, effluent checks and fibre-waste balances. [S02] [S05]

2.3 DOMAIN OWNERS GOVERN MEANING AND RIGHTS WHILE CAPTURE STAYS AT THE LINE

Assigning a domain owner does not mean a department head keys in hourly logs or writes queries. [S02] Transaction capture belongs at the point of work, where operators, inspectors, storekeepers and technicians observe the event. [S05] [S08] The domain owner holds decision authority over four structural questions: [S02]

1. **Purpose and scope:** Determining why a record is collected, which decision or obligation it serves, and whether a proposed field is needed. [S02] [S29]
2. **Definition and calculation:** Approving the contractual meaning of every entity, state and KPI in the domain, including units, time boundaries and exclusions under ISO 22400-2. [S01] [S38]
3. **Access and sharing:** Specifying which internal roles, pipelines and external partners may view, consume or export curated data products under least privilege. [S25]
4. **Retention and archival:** Setting how long raw evidence, curated records and summaries are kept for operational reuse, audit and compliance. [S09] [S29]

Separating governance decisions from routine entry prevents a recurring failure in turnarounds. [R12] [S19] Public reporting on Cameroon's legacy textile manufacturer CICAM in July 2023 noted capacity utilisation of 33.7 percent at Garoua, 48.9 percent

at CICAM 1 and 10.83 percent at CICAM 2, with underutilisation costing about CFAF 2 billion in 2020 (an inherited grade-C report not refreshed in 2026). [R12] Under utilisation pressure, managers need clarity on whether lost output stems from material starvation, mechanical fault, power loss or rework; if supervisors define stoppage categories differently on local sheets, leadership cannot direct engineering time to the binding constraint. [R12] [S02] [S38]

2.4 PERSONAL DATA IN THE PEOPLE DOMAIN REQUIRE HEIGHTENED LEGAL AND ETHICAL CONTROLS

The people domain demands stricter boundaries than machine telemetry because its records describe identifiable workers. [S29] [S30] Under the data-minimisation principle in Recital 39 of Regulation (EU) 2016/679 (GDPR), personal data must be adequate, relevant and limited to what is necessary for the stated purposes. [S29] Tracking line-level or bundle-level completion supports scheduling and traceability, whereas continuous individual biometric or video tracking is rarely justified for line balancing. [S08] [S29]

Furthermore, Principle 5.5 of the International Labour Office (ILO) code of practice on workers' personal data states that decisions concerning a worker should not be based solely on automated processing of that worker's personal data, which GFT treats as a mandatory engineering constraint. [S30] In Cameroon, Section 19(1) of Law No. 2024/017 of 23 December 2024 relating to personal data protection provides that processing personal data is subject to prior authorisation by the Personal Data Protection Authority. [S31] Because current enforcement requires local legal verification, the people-domain owner must obtain qualified Cameroonian legal advice before authorising collection, transfer or algorithmic analysis of worker records; this paper provides engineering design, not legal advice. [S31]

2.5

Ownership, stewardship and custody remain distinct

A small factory team often collapses all data work into one job description, hiring a single "IT and data officer" expected to define metrics, clean shift sheets, maintain servers, build dashboards and validate models. [S18] [S20] That concentration fails because a technical administrator lacks operational authority to fix capture habits on the floor ([S02], [S19]), and allowing the person who administers storage or builds a model to alter KPI formulas and approve releases removes independent quality and risk control. [S11] [S14]

2.6 A SIX-ROLE CHARTER SEPARATES BUSINESS AUTHORITY FROM TECHNICAL EXECUTION

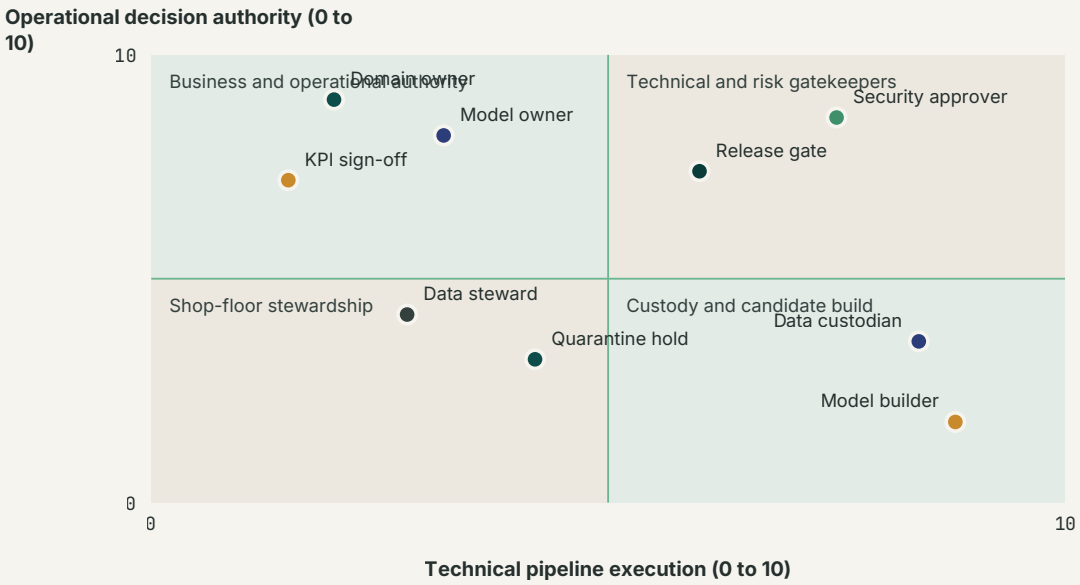
To preserve checks and balances without inflating headcount, GFT defines six functional roles across the data spine, combining role principles from ISO 8000-150 with risk-governance structures from NIST CSF 2.0, NIST AI RMF 1.0 and ISO/IEC 23894. [S02] [S11] [S14] [S23] These roles are bundles of decision rights, not separate full-time salaries: [S02] [S0201]

1. **Accountable domain owner:** A department head owning the purpose, definition, access policy and retention schedule for a domain's datasets and indicators. [S02] [S38]
2. **Domain data steward:** A practitioner close to the work, such as an industrial engineer, quality supervisor or maintenance planner, who monitors daily completeness and validity, manages quarantine queues, coordinates corrections and drafts change proposals. [S02] [S04]
3. **Technical data custodian:** An IT, OT or systems engineer operating storage, brokers, schema enforcement, access controls, encrypted backups and restoration, without authority to alter business meaning. [S02] [S25]
4. **Model builder (or technical integrator):** The analyst, data scientist or vendor team preparing training snapshots, configuring features, training candidate models and compiling evaluation evidence. [S13] [S15]
5. **Operating model owner:** The line manager whose decision uses a model's output, holding authority to set thresholds, accept or reject candidates, trigger rollback and retire models. [S11] [S14]
6. **Security and OT approver:** The practitioner verifying that any new feed, edge agent, interface or model complies with OT safety constraints, least privilege and personal-data rules before activation. [S22] [S23] [S25]

The illustrative governance map in **Exhibit 5**, p. 34 positions these six functional roles alongside three decision anchors across operational decision authority (vertical axis) and technical pipeline execution (horizontal axis). [S02] [S11] [S14] [S23]

EXHIBIT 05 / ILLUSTRATIVE GOVERNANCE COORDINATES (0 TO 10 SCALE)

A six-role charter covers governance without creating a department



ILLUSTRATIVE [S02] [S11] [S14] [S23] Illustrative mapping of the six functional governance roles and three decision anchors across technical execution closeness (x-axis) and operational decision authority (y-axis), derived from role assignment in S02, human threshold setting in S11, AI risk integration in S14 and enterprise risk governance in S23. Coordinates are illustrative design positions for a small factory, not empirical measurements.

As **Exhibit 5**, p. 34 illustrates, high operational authority sits in the upper quadrants with the domain owner, model owner and security approver, while technical execution sits in the lower-right quadrant with the custodian and builder. [S02] [S11] [S23] The chart clarifies which hat a person wears for each decision and which combinations create a conflict of interest. [S02] [S14]

2.7 ROLE COMBINATION RULES PREVENT SELF-APPROVAL IN CONSTRAINED TEAMS

In a small factory with an illustrative workforce of 150 to 500 staff and two or three technical specialists, one person will hold more than one role in the charter. [S02] [S16] For instance, the quality assurance manager may serve as both quality domain owner and operating model owner for an inspection system, while an industrial engineer acts as both production and quality steward. [S02] [S14] Combining compatible roles keeps overhead low, provided the factory enforces three separation rules from NIST AI RMF Playbook guidance and ISO/IEC 23894: [S11] [S14] [S0201]

- **No silent definition changes by technical custodians:** The custodian may never alter a KPI formula, unit conversion, state mapping or allowed-value list without recorded approval from the domain owner. [S02] [S38]
- **No self-approval of model releases:** The model builder who trains, tunes or integrates a model may never act as the operating model owner or security approver who signs off its production release. [S11] [S14] [S0201]
- **Declared conflict and peer review when roles overlap:** Whenever an analyst acts as both steward and builder for a low-risk internal tool, the conflict is recorded in the register and validation sign-off is elevated to the domain owner. [S02] [S04] [S0201]

2.8 SEVEN RECURRING DECISIONS LINK AUTHORITY TO VERIFIABLE ARTEFACTS

Decision rights remain abstract unless each right is tied to a trigger, an approver, a required artefact and an escalation forum. [S02] [S0201] Synthesising role assignment from ISO 8000-150 ([S02]), provenance modelling from W3C PROV-O ([S09]), pipeline lineage from MLOps architecture ([S15]) and incident-response verification from NIST SP 800-61 ([S27]), **Exhibit 6**, p. 35 specifies the seven lifecycle decisions governing the data spine.

EXHIBIT 06 / GOVERNED LIFECYCLE DECISION CLASSES

Decision rights become testable when each has an artefact and escalation route

Decision	Proposer or executor	Accountable approver	Required evidence artefact	Escalation forum
1. Define KPI or schema	Domain steward	Domain owner	Versioned KPI or contract entry with impact note	Weekly operations review
2. Grant data access	Data custodian	Domain owner and security approver	Scoped access ticket with purpose and expiry	Monthly governance review
3. Quarantine feed	Domain steward or automated gate	Domain steward (immediate hold)	Quarantine log with failed rule ID and batch scope	Daily shift tier-two meeting
4. Correct historical record	Source role and steward	Domain owner	Appended superseding event linked to original ID	Weekly operations review
5. Release model version	Model builder	Model owner and security approver	Candidate record, local slice metrics and fallback test	Monthly risk and S&OP gate
6. Roll back active model	Model owner, steward or OT lead	Model owner (immediate revert)	Incident ticket, alert trigger and restored baseline pointer	Daily shift meeting and post-mortem
7. Retire dataset or model	Domain or model owner	Domain owner and custodian	Archival manifest, lineage freeze and consumer notice	Monthly governance review

DERIVED [S02] [S09] [S15] [S27] Derived governance specification linking seven lifecycle decisions across the industrial data spine to their proposing or executing role, accountable approver, mandatory evidence artefact and escalation forum, synthesised from S02, S09, S15 and S27.

Two asymmetric rules in **Exhibit 6**, p. 35 deserve emphasis. [S02] [S27] Opening a new exposure, such as changing a definition, granting access or releasing a model, requires prior approval backed by documented evidence. [S02] [S15] Conversely, containing a failure, whether quarantining a corrupted feed or rolling back a misbehaving model to manual fallback, is a pre-authorised protective action that the steward, model owner or OT lead executes immediately on the shift, with formal review at the next scheduled meeting. [S15] [S22] [S27]

2.9

Line stewards need time and escalation authority

Data stewardship fails when management treats it as an unbudgeted chore added to a full workload. [S19] [S35] Sambasivan et al. found in high-stakes AI systems that upstream data cascades were pervasive (92 percent of studied cases), invisible and delayed because organisations undervalued data work. [S19] Although that study examined high-stakes AI broadly rather than Cameroonian manufacturing, the shop-floor mechanism is identical: without protected weekly hours to inspect rejected records, verify master-data codes and coach operators, defects accumulate until reports and models fail. [S04] [S19]

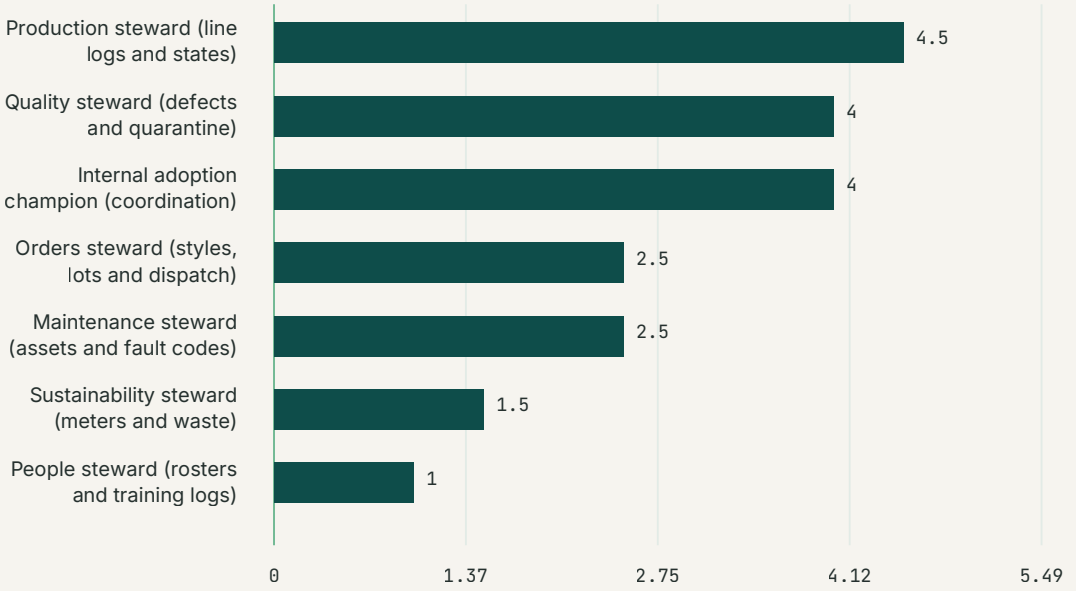
2.10 STEWARDSHIP REQUIRES A PROTECTED WEEKLY TIME BUDGET INSIDE STANDARD SHIFT HOURS

Rather than hiring full-time data bureaucrats, a small factory protects weekly working time for existing supervisors, planners and engineers to perform stewardship. [S02] [S35] In Cameroon, labour reference data record a standard 40-hour week, a 48-hour maximum and overtime at 1.2 times base pay, while the statutory minimum wage (SMIG) rose 5 percent from XAF 41,875 in 2023 to XAF 43,969 a month from 23 February 2024 (about US$76), and multi-sector average wages sit near US$359 a month in inherited grade-C sources. [R14] [R15] [R16] Pushing stewardship beyond the 40-hour week without workload relief adds overtime cost or leaves checks undone when orders run late. [R16] [S35]

To make the commitment explicit before a site baselines its own defect rates, **Exhibit 7**, p. 37 shows an illustrative weekly allocation across the six stewards and internal champion inside a 40-hour week. [R16] [S02] [S37]

EXHIBIT 07 / ILLUSTRATIVE PROTECTED HOURS PER WEEK

A bounded weekly stewardship budget fits inside a standard 40-hour shift week



ILLUSTRATIVE [R16] [S02] [S37] Illustrative weekly time budget allocated to domain stewardship and adoption coordination within Cameroon's 40-hour standard working week (R16), applying role assignment from S02 and internal champion coordination from S37. Values are illustrative planning allowances to be resized after a 30-day factory workload baseline, not universal staffing benchmarks.

As shown in **Exhibit 7**, p. 37, the illustrative allowance totals 20.0 hours per week across seven staff members, equivalent to 0.50 of a full-time equivalent (FTE) spread across the plant rather than a new hire. [R16] [S02] Production and quality stewards receive the largest allocations (4.5 and 4.0 illustrative hours per week) because hourly line logs and defect sheets carry the highest transaction volume. [S04] [S38] A plant migrating legacy spreadsheets needs more stewardship time during its first 90 days than one with stabilised barcode capture. [S35] [S41]

WORKED EXAMPLE

Sizing and valuing a small factory's monthly stewardship capacity. Consider an illustrative Cameroonian garment facility on the standard 40-hour week ([R16]), assigning the six steward roles and one champion role ([S02], [S37]) to existing supervisors and engineers with the weekly budget in **Exhibit 7**, p. 37.

1. **Monthly protected hours:** Summing the seven allowances gives `4.5 + 4.0 + 4.0 + 2.5 + 2.5 + 1.5 + 1.0 = 20.0` illustrative hours per week. Across 4.33 weeks per month, the plant budgets `20.0 * 4.333 = 86.7` illustrative hours per month, or `0.50` of one 40-hour FTE ([R16]).
2. **Illustrative monthly value in XAF:** Rather than the statutory minimum wage of XAF 43,969 a month ([R14]), assume mid-level supervisors earn an illustrative gross salary of XAF 220,000 a month (near the inherited multi-sector average of US$359 a month in [R15] at roughly XAF 600 per US dollar). Across 173.3 standard monthly hours, the illustrative rate is `220,000 / 173.3 = XAF 1,269` per hour, valuing 86.7 protected hours at an illustrative `XAF 110,000` per month.
3. **Avoided overtime and cascades:** Paying those 86.7 hours as overtime at the 1.2x statutory rate ([R16]) costs an illustrative `XAF 132,000` per month while leaving bad records live all week ([S19]). Replacing manual slide formatting and duplicate entry funds stewardship inside normal hours ([S02], [S19]). All hours and XAF figures are illustrative method calculations, not observed GFT results.

2.11 QUARANTINE AND ESCALATION RIGHTS GIVE STEWARDS OPERATIONAL TEETH

Protected time is wasted if a steward lacks authority to stop a broken feed from polluting downstream reports and training sets. [S04] [S19] Under ISO/IEC 5259-3:2024, managing data quality for analytics and machine learning requires systematic lifecycle controls. [S04] GFT grants every domain steward three operational rights: [S02] [S04]

1. **Right to quarantine a defective batch or feed:** When a check detects impossible units, missing asset IDs, unmapped defect codes or clock drift, the steward places the batch into quarantine immediately, blocking it from curated KPIs and training snapshots while preserving raw storage. [S04] [S09] [S15]
2. **Right to issue a source-correction task:** A steward never overwrites a shop-floor entry silently, instead returning a ticket to the originating role for an appended superseding record with actor identity and reason under W3C PROV-O. [S09]
3. **Right to escalate recurring capture defects:** If a quarantined feed remains uncorrected after one shift (or 24 hours for daily logs), or fails three times in a week, the steward escalates it to the daily tier-two board or weekly operations review. [S02] [S36]

2.12

Definition changes follow a controlled decision

Nothing erodes trust faster than two managers arriving at a meeting with different numbers for the same indicator because a formula changed without notice. [S01] [S38] ISO 22400-2:2014 specifies that manufacturing KPIs must be defined by formula,

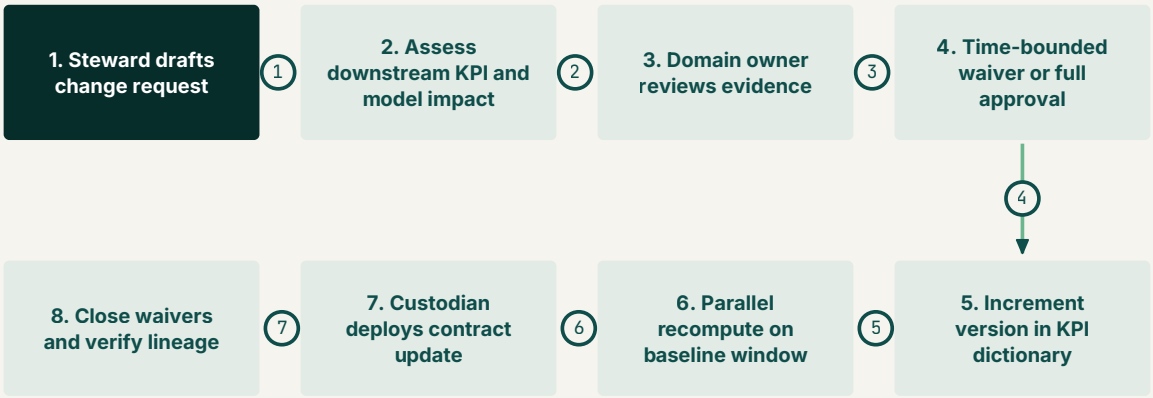
inputs, time behaviour, unit and characteristics. [S38] ISO 8000-1:2022 frames quality around explicit requirements, while W3C's Data on the Web Best Practices requires a version history explaining each change. [S01] [S0202]

2.13 EVERY SCHEMA OR KPI MODIFICATION REQUIRES A FIVE-PART CHANGE PROPOSAL

In GFT's governance model, a KPI definition, taxonomy or event schema is treated as an engineering drawing: neither a manager nor an analyst may alter a calculation in place, and every change follows the sequence in **Exhibit 8**, p. 39. [S01] [S02] [S38] [S0202]

EXHIBIT 08 / GOVERNED CHANGE-CONTROL SEQUENCE

Controlled definition change prevents silent metric redefinition across shifts



1. Record reason, formula diff and fields
2. Check downstream reports and ML features
3. Approve change or grant 14-day waiver
4. Assign semantic version and shift date
5. Compare old and new formula on 30 days
6. Deploy schema update and notify owners
7. Verify lineage and close expired waivers

DERIVED [S01] [S02] [S38] [S0202] Derived eight-step workflow for modifying a KPI formula, event schema or taxonomy mapping, combining quality requirements from S01, role accountability from S02, manufacturing KPI structure from S38 and version-history documentation from S0202.

As **Exhibit 8**, p. 39 shows, every change proposal must record five elements before the owner approves it: [S01] [S38] [S0202]

1. **Operational reason:** Why the current definition or schema is unfit, such as adding a rework station or aligning with a buyer protocol. [S01] [S38]
2. **Exact specification difference:** The change in formula, numerator, denominator, window, grain, unit or allowed values under ISO 22400-2. [S38]
3. **Downstream impact assessment:** Every report, dataset, carbon calculation or model feature consuming the field, compiled with the custodian. [S09] [S15] [S18]

4. **Semantic version and effective date:** A new version ID (such as `first_pass_yield` moving from `v1.2` to `v2.0`) and effective shift timestamp. [S09] [S0202]
5. **Named approver and dual-run reconciliation:** Owner sign-off with an illustrative 30-day parallel calculation showing how the new formula shifts the baseline. [S36] [S38]

2.14 TEMPORARY WAIVERS HANDLE SHOP-FLOOR URGENCY WITHOUT CREATING PERMANENT DEBT

Factories face urgent disruptions that cannot wait for a weekly review: an emergency fabric lot, a failed sub-meter or an unscheduled split shift. [S05] [S18] Without a lawful fast path, operators revert to unrecorded sheets. [S19] [S37]

Step 4 of **Exhibit 8**, p. 39 provides a controlled temporary waiver. [S01] [S02] A domain owner may authorise a provisional code, estimated meter factor or temporary override on the same day under three conditions: [S02] [S04]

- **Explicit quality-state tagging:** Every record captured under the waiver is tagged `provisional-waiver` or `estimated` rather than `observed`. [S04] [S09]
- **Named owner and hard expiry:** Each waiver carries one owner and an automatic expiry set at an illustrative maximum of 14 calendar days. [S02] [S0201]
- **Mandatory conversion or reversion:** Within 14 days, the steward and owner must either promote a permanent versioned contract or revert the rule and reconcile provisional records. [S01] [S0202] Overdue waivers appear at the top of the weekly action log. [S18]

2.15

Model risk receives independent challenge

When a factory moves from descriptive reporting to predictive or vision models, the consequences of a flawed decision right multiply. [S11] [S18] A model compresses historical records into recommendations or actions. [S13] [S15] Left unchecked, the builder who trained a model has a bias to emphasise aggregate accuracy and downplay local failures on specific styles, night shifts or lighting conditions. [S04] [S19]

2.16 HUMAN JUDGEMENT SETS ACCEPTANCE THRESHOLDS AND HOLDS VETO AUTHORITY

Standards place human accountability at the centre of AI risk control. [S11] [S12] [S14] NIST AI RMF 1.0 states that human judgement should decide specific metrics and threshold values for AI trustworthiness. [S11] ISO/IEC 23894:2023 integrates risk management into AI activities ([S14]), while ISO/IEC 42001:2023 and ISO/IEC 5338:2023 govern management systems and lifecycle processes ([S12], [S13]).

Independent challenge separates the **model builder** from the **operating model owner** (the line manager accountable for the affected process). [S11] [S14] [S0201] The operating model owner holds four veto rights: [S11] [S44]

1. **Setting acceptance thresholds:** Specifying maximum false rejections, false-alarm burden, review time per shift and minimum gain over a simple comparator before training. [S11] [S43]

2. **Rejecting operationally unfit models:** Even if a vision model exceeds 90 percent aggregate accuracy across defect types in literature ([R29]), the quality owner rejects release if false alarms slow inspection or fallback fails. [S11] [S43]
3. **Ordering suspension or rollback:** Suspending a model and reverting to manual or statistical rules immediately if drift, overrides or yield losses appear. [S15] [S27]
4. **Stopping or postponing a project:** Postponing or stopping an AI project whenever resources, skills, costs or expected benefits are not met. [S44]

2.17 CHALLENGE INTENSITY SCALES WITH OPERATIONAL, WORKER AND SAFETY CONSEQUENCE

Applying identical bureaucracy to a monthly yarn forecast and a real-time machine-control loop wastes management attention. [S11] [S14] GFT scales challenge across four consequence tiers, combining NIST AI RMF Playbook lines-of-defence concepts ([S0201]), operational-technology safety guidance in NIST SP 800-82 Revision 3 ([S22]) and the ILO worker-protection principle ([S30]).

METHOD

Method: Consequence-tiered model challenge and separation of duties. Before development or procurement, the domain owner and champion classify the use case into one of four tiers ([S11], [S14], [S43]):

1. **Tier 1: Advisory planning analytics (low consequence).** Weekly or monthly planning tools, such as demand forecasting or spare-parts suggestions ([S43]). *Challenge:* Steward review of lineage and comparator results; domain owner approval ([S11], [S15]).
2. **Tier 2: Shop-floor operator assistance (moderate consequence).** Shift-level aids, such as vision defect highlighting confirmed by an inspector ([R29], [S16]). *Challenge:* Shadow evaluation across style and shift slices ([S04], [S16]), rehearsed fallback ([S15]), and sign-off by the quality owner and custodian ([S11]).
3. **Tier 3: Worker-affecting analytics (high human and legal consequence).** Any metric or model evaluating operator productivity, skill, scheduling, pay or discipline ([S29], [S30]). *Challenge:* Compliance with ILO Principle 5.5 against sole automated decisions ([S30]), data minimisation under [S29], Cameroonian prior-authorisation screening under Section 19(1) of Law No. 2024/017 with legal counsel ([S31]), and sign-off by the HR owner, plant director and worker forum ([S11], [S30]).
4. **Tier 4: Control-connected automation (high physical and OT consequence).** Any model writing setpoints or altering OT behaviour without human confirmation ([S22]). *Challenge:* OT safety verification under NIST SP 800-82 ([S22]), hard-wired interlocks ([S22]), disconnect drills ([S27]), and sign-off by the OT engineer, production owner and plant director ([S22], [S23]).

Binding challenge rules to these four tiers concentrates scrutiny where physical safety ([S22]), worker rights ([S30], [S31]) or throughput ([R29]) are at stake, while letting low-risk advisory analytics move once baseline and quality gates pass. [S04] [S16] [S43]

2.18

An internal champion sustains adoption

Role charters fail if nobody drives cross-functional follow-through under daily production pressure. [S35] [S37] World Bank case-study evidence on SME digital adoption in Viet Nam found that firms successfully taking up the platform assigned an internal point person who championed adoption. [S37] OECD SME research similarly notes that converting traditional processes into digital ones requires sizeable effort and a longer learning process rather than an immediate payoff. [S35]

2.19 THE CHAMPION COORDINATES FOLLOW-THROUGH WITHOUT STRIPPING DOMAIN OWNERS OF ACCOUNTABILITY

In GFT's governance architecture, the **internal adoption champion** coordinates execution across the six domains, reporting to the factory director without replacing domain owners or becoming a solitary "chief data officer". [S02] [S23] [S37] Instead, the champion holds three cross-cutting responsibilities: [S36] [S37]

- **Cross-domain coordination:** When an identifier or formula spans departments, such as linking a spinning lot ID ([S08]) to a sewing order ID ([S05]) and a quality claim ([S38]), the champion convenes owners and stewards to agree on one shared contract. [S02] [S37]
- **Workplace training around real tasks:** Rather than classroom lectures, the champion coordinates practical training so staff practise event capture, quarantine clearance, control-chart reading and model fallback on live work. [S27] [S35] [S0201]
- **Cadence hygiene and benefit tracking:** The champion ensures overdue waivers, recurring quarantine failures and monitoring alerts reach the weekly and monthly agendas, and maintains the embedded monitoring and evaluation log recommended by World Bank MSME guidance. [S36] [S37]

2.20 PAIRING THE CHAMPION WITH A DEPUTY AND PORTABLE REGISTERS PREVENTS HERO DEPENDENCE

If the digitalisation programme lives inside one champion's laptop, the system collapses when that person falls ill or leaves. [S18] [S37] NIST AI RMF Playbook guidance underscores that roles, responsibilities and communication lines must be documented across the organisation, with ownership metadata attached to work products. [S0201]

To prevent single-person dependency, the factory enforces two safeguards from day one: [S37] [S0201] [S0202]

1. **Named deputy rotation:** The champion is paired with a deputy (typically the quality or production steward) who chairs every fourth weekly review and co-signs registry updates during the first six months. [S27] [S37]
2. **Zero private governance files:** Every role assignment, KPI formula, schema contract, waiver, quarantine ticket and model record resides in the shared, versioned governance register with complete change history ([S09], [S0202]), never on a personal drive. [S15] [S0201]

GFT NOTE

Practical recommendation for GFT's Cameroon textile programme. In a greenfield or turnaround Cameroonian mill, appoint the industrial-engineering lead or quality-systems manager as internal adoption champion at kick-off ([S37]), paired with the production planner as deputy ([S02]). Do not form a separate data committee before the first shift runs. Embed the six domain owners and stewards into the initial manual KPI sheets so every column has one named owner, one ISO 22400-2 formula ([S38]) and one quarantine rule ([S04]) before any database, broker or vision camera is commissioned.

2.21

Data decisions use, but do not redesign, factory meetings

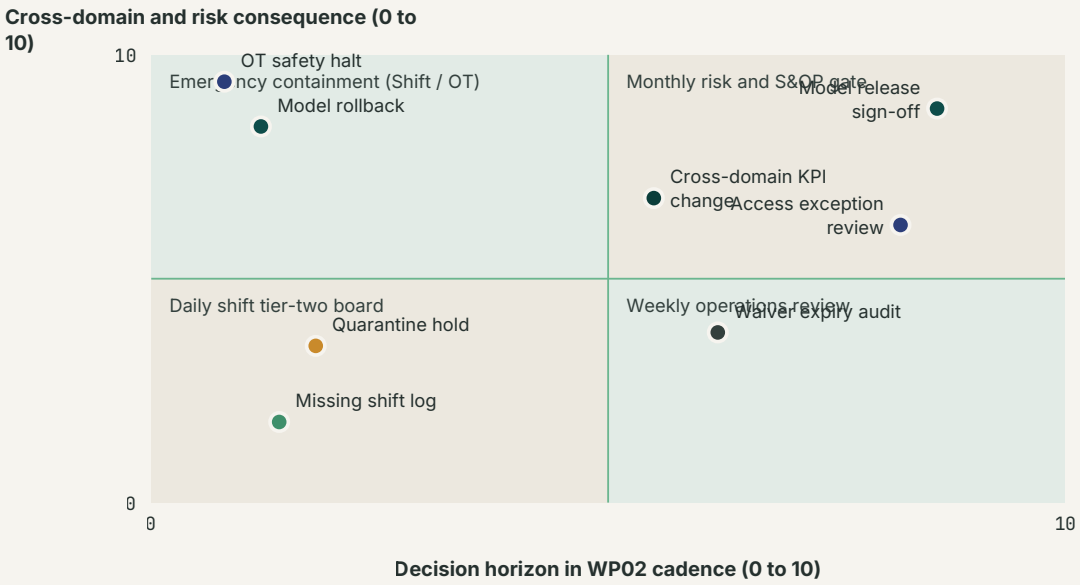
Scheduling a second tier of "data committee" meetings that duplicate the daily shift handover, weekly operations review and monthly sales and operations planning (S&OP) cycle breeds resentment. [S18] [S35] In the GFT library, **WP02** owns the operating cadence, visual boards and meeting method; **WP14** refuses to redesign those forums or invent a parallel hierarchy. [S23] [S36]

2.22 FOUR CLASSES OF DATA AND MODEL ESCALATION ROUTE INTO EXISTING WP02 FORUMS

WP14 injects decision-ready escalations into the WP02 cadence only when an issue crosses domains or exceeds a steward's local authority. [S02] [S23] [S0201] Drawing on ISO 8000-150 role escalation ([S02]), NIST AI RMF threshold authority ([S11]) and NIST CSF 2.0 risk governance ([S23]), **Exhibit 9**, p. 44 maps eight recurring escalations across decision horizon (x-axis) and cross-domain consequence (y-axis).

EXHIBIT 09 / HORIZON (X) AND CONSEQUENCE (Y) INDEX (0 TO 10)

Data escalation adds decisions without creating a second management cadence



DERIVED [S02] [S11] [S23] Derived routing matrix mapping eight recurring data and model escalations across decision horizon (x-axis: 0 = immediate shift action, 10 = monthly portfolio review) and cross-domain risk consequence (y-axis: 0 = single-line capture fix, 10 = plant-wide risk or policy gate) into WP02's existing management forums, based on S02, S11 and S23.

As **Exhibit 9**, p. 44 shows, each WP02 forum handles a defined subset of decisions with a standard WP14 artefact: [S02] [S11] [S23]

1. **Daily shift and tier-two board (5 minutes in handover):** Reviews active quarantine holds blocking daily KPIs, missing shift logs after 24 hours, and any model rollback or OT disconnect from the prior shift. [S04] [S22] [S27] *Approver:* Line supervisor and steward (or model owner for rollback). [S02] [S11]
2. **Weekly operations review (10-minute slot):** Resolves capture defects breaching the three-failure threshold, approves or rejects cross-domain KPI and schema changes, and audits waivers nearing 14-day expiry. [S01] [S02] [S38] *Approver:* Domain owners chaired by the operations manager, logged by the champion. [S02] [S37]
3. **Monthly S&OP and risk review (15-minute gate):** Decides model admission, shadow-to-assisted promotions, retirements, least-privilege access exceptions ([S25]), personal-data register updates ([S29], [S31]), and fix, scale, postpone or stop decisions on pilots ([S36], [S44]). *Approver:* Plant director, model owner, domain owner and security/OT approver. [S11] [S14] [S23]

2.23 ONE SHARED ACTION LOG CLOSES THE LOOP AND FEEDS THE GFT TOOL SUITE

Every decision in the daily, weekly or monthly forum is recorded in the single WP02 action log, never in a separate data-team sheet. [S36] [S37] Each entry carries four fields: artefact or defect ID, named owner, due date and verification condition (such as seven consecutive shifts with zero unmapped downtime codes on weaving line two, verified by the production steward). [S02] [S36] [S38]

This governance structure connects directly to three diagnostic tools in the GFT platform: [S02] [S14] [S38]

- `kpi-cascade` **tool:** Uses the six-domain ownership assignments and ISO 22400-2 contracts ([S38]) to trace how capture defects or unapproved formula changes propagate into plant OEE, first-pass yield and margin. [S02] [S38]
- `maturity` **tool:** Evaluates named domain owners, protected weekly stewardship hours, versioned change control and practised rollback authority before scoring readiness under staged MLOps maturity ([S16]). [S02] [S16]
- `risk-matrix` **tool:** Operationalises the four-tier consequence classification and challenge rules ([S11], [S14], [S22], [S30]) to weigh operational, worker-privacy and OT-safety exposure before model admission or release. [S14] [S23]

> **WARNING**
>
> **Constraint: Governance collapses if dual logs or unowned exceptions are tolerated.** Two habits defeat this charter even when job descriptions are signed ([S18], [S19]). First, if department heads bring private spreadsheets into the weekly review alongside the governed KPI product, staff learn that the governed spine is optional ([S19], [S37]): any figure outside a governed contract or active waiver cannot justify a decision in the meeting ([S02], [S38]). Second, renewing temporary waivers past their 14-day illustrative expiry without root-cause correction turns the exception log into a shadow schema ([S18]). Auditing overdue waivers and recurring quarantine holds in the monthly review ([S36]) prevents both failures.

Anchoring every dataset, formula, access grant and model release to a named role inside the factory's management rhythm gives a small team institutional control without bureaucracy. [S02] [S23] Once decision rights are clear, those owners need one shared language across machines, lines and ledgers, which the next chapter establishes. [S05] [S38]

A metric without a named owner is an opinion; a model without an independent challenger is an unhedged bet.

CHAPTER SOURCE NOTES

S02 International Organization for Standardization. 2022.

S01 International Organization for Standardization. 2022.

R06 Cameroon Trade Hub. Publication date unrecorded; inherited source-text check.

R07 Ecofin Agency. Publication date unrecorded; inherited source-text check.

S23 National Institute of Standards and Technology. 2024-02-26.

S19 Association for Computing Machinery, CHI 2021. 2021.

S18 Advances in Neural Information Processing Systems. 2015.

S38 International Organization for Standardization. 2014.

S09 World Wide Web Consortium. 2013-04-30.

S11 National Institute of Standards and Technology. 2023-01-26.

S0201 National Institute of Standards and Technology. 2023-01-26.

S14 ISO/IEC. 2023.

S36 World Bank. 2023.

S05 International Electrotechnical Commission. 2013.

S04 ISO/IEC. 2024.

S08 GS1 AISBL. 2022.

S22 National Institute of Standards and Technology. 2023.

S29 European Parliament and Council of the European Union. 2016-04-27.

S30 International Labour Office. 1997.

S31 Presidency of the Republic of Cameroon. 2024-12-23.

S25 National Institute of Standards and Technology. 2020-12-10.

R12 Business in Cameroon. Publication date unrecorded; inherited source-text check.

S20 arXiv, author version of an IEEE Access article. 2022-05-14.

S13 ISO/IEC. 2023.

S15 Google Cloud. 2024-08-28.

S16 Microsoft Learn, Azure Architecture Center. 2026-06-03.

S27 National Institute of Standards and Technology. 2025.

S35 OECD. 2022-11-21.

R14 WageIndicator. Publication date unrecorded; inherited source-text check.

R15 wage.is. Publication date unrecorded; inherited source-text check.

R16 WageIndicator. Publication date unrecorded; inherited source-text check.

S37 World Bank. 2025-05-20.

S41 GHG Protocol, World Resources Institute and World Business Council for Sustainable Development. 2013.

S0202 World Wide Web Consortium. 2017-01-31.

S12 ISO/IEC. 2023.

S44 The Journal of Technology Transfer, Springer Nature. 2024-07-27.

S43 Frontiers in Artificial Intelligence. 2021-06-14.

R29 MDPI; industry inspection literature. Publication date unrecorded; inherited source-text check.

GREEN FUTURE TECH

LIMITED / HONG KONG SAR

03

PART I: GOVERN MEANING BEFORE SCALE / MEANING AND CONTRACTS

One operational vocabulary joins the line to the ledger

Shared identifiers, event semantics and KPI formulas let unlike systems describe the same work without hiding their differences.

CHAPTER 03
DECISION TEST

Shared identifiers, event semantics and KPI formulas let unlike systems describe the same work without hiding their differences.

A factory joins the production line to the financial ledger only when every record shares one governed operational vocabulary of immutable identifiers, boundary-aware event semantics, contractual key performance indicator (KPI) formulas and versioned data contracts. [S05] [S08] [S38] Assigning decision rights to domain owners and stewards, as established in Chapter 02, determines who approves a definition. [S02] Authority alone cannot stop spreadsheets, a programmable logic controller (PLC) and an accounting ledger from using the same word to measure different physical realities. When a supervisor logs output as garments lifted from the table, the quality lead logs garments passing first inspection, and the accountant books packed cartons in inventory, the factory holds three incompatible definitions behind one label. [S05] [S38]

The stakes of semantic discipline are immediate in Green Future Tech Limited (GFT) and its Cameroon textile programme. According to the programme's inherited evidence base, prepared on 31 January 2024 with entries not refreshed in 2026 and unrecorded publication dates, roughly 95 percent of Cameroon's cotton fibre is exported raw while about 5 percent is processed locally, against a target of 50 percent by 2030. [R06] [R07] Moving from raw-fibre export toward spinning, weaving, dyeing and garment assembly requires daily records supporting cost control, compliance, maintenance and machine-learning (ML) applications. [R06] [S04] In small teams, those records start as manual shift sheets before brokers or cameras arrive. If each workbook invents its own identifiers and rules, the plant accumulates semantic debt: manufacturers in workshop settings have reported that data processing for one application can consume up to 70 percent of total cost, a warning about repeated translation rather than a universal average. [S39] Establishing one vocabulary before scaling integration turns preparation into shared infrastructure. [S10] [S39]

3.1

Stable identifiers anchor every record

Every operational record makes a claim about physical objects interacting at a place and time, and that claim cannot be audited unless the objects carry stable identifiers. [S05] [S08] In an ungoverned spreadsheet, operators type free-text strings: a sewing line appears as `Line 2`, `L-02` or `Garoua Sewing B`, while a fabric batch appears as `Indigo Lot 4` or `Lot-04-Nov`. [S0302] String mismatches cause silent drops or false joins across production, quality and maintenance logs. Worse, renaming `Line 2` to `Uniform Assembly Line 1` makes historical records keyed by display label lose continuity or attribute past output to the wrong configuration. [S01] [S05]

To eliminate identity drift, the data spine separates an entity's immutable surrogate identifier from its mutable attributes. [S05] [S09] The identifier is a permanent token assigned once by a single authority when registered, carrying no embedded mutable attributes. [S02] [S08] Descriptive properties such as display names, bay coordinates, supervisor assignments or rated speeds live in versioned master-data tables linked to the immutable identifier by effective dates, preserving history across shifts and migrations. [S05] [S06] [S09]

A garment or textile factory requires ten identifier classes across manual entry and automated capture: [S05] [S08]

1. **Site identifier (`site_id`)**: designates a facility, distinguishing a northern ginning or spinning site from a coastal garment plant. [S05] [S08]
2. **Area identifier (`area_id`)**: designates a zone within a site, such as opening, spinning, dyeing, cutting, sewing or packing. [S05]
3. **Line identifier (`line_id`)**: designates a production unit scheduled as a capacity resource, such as a sewing line or finishing range. [S05]
4. **Asset identifier (`asset_id`)**: designates a serialised machine, utility meter, edge gateway or inspection table, remaining constant when relocated. [S05] [S06]
5. **Production order identifier (`order_id`)**: designates an authorisation to manufacture a stated quantity of a style by a scheduled date. [S05]
6. **Material lot identifier (`lot_id`)**: designates a traceable quantity of lint, yarn, fabric roll, dye bath or cut bundle sharing uniform history. [S08]
7. **Product style identifier (`style_id`)**: designates the versioned specification, bill of materials, grading and tolerance sheet for an item. [S05]
8. **Operation identifier (`operation_id`)**: designates a transformation, inspection or handling step within a routing sequence. [S05] [S08]
9. **Defect code identifier (`defect_id`)**: designates a controlled entry in the canonical taxonomy of product faults or downtime causes. [S01] [S04]
10. **Shift instance identifier (`shift_id`)**: designates a bounded operating window on a date and site, linking utility readings and targets. [S05] [S38]

3.2 WORKER IDENTIFIERS ARE EXCLUDED BY DEFAULT TO PROTECT PRIVACY AND LEGAL COMPLIANCE

Absent from the ten default identifiers is a mandatory person or operator identifier (`person_id`). Vendor templates often key every piece count, defect tally and stoppage to an individual worker badge. GFT rejects indiscriminate worker-level keying in the core spine for legal and analytical reasons. [S29] [S31] Under Recital 39 of Regulation (EU) 2016/679 (GDPR), personal data must be adequate, relevant and limited to what is necessary for the purposes for which they are processed. [S29] In Cameroon, Section 19(1) of Law No. 2024/017 of 23 December 2024 states that processing personal data is subject to prior authorisation by the Personal Data Protection Authority. [S31] Because implementing procedures under Law No. 2024/017 require project-specific verification by local counsel (and this paper is not legal advice), designing shop-floor schemas without unnecessary personal data reduces exposure from the first shift. [S29] [S31]

Excluding `person_id` from standard production, quality and machine events does not weaken engineering analysis. Line balancing, bottleneck diagnosis, maintenance, yield reconciliation and forecasting depend on `line_id`, `asset_id`, `operation_id`, `style_id` and `shift_id`, not on the personal identity of the operator at a sewing head. [S05] [S38] Where a documented purpose requires human attribution, such as verifying

that a safety lock-out or quarantine release was performed by an authorised role, recording `role_id` in an access-controlled audit table satisfies accountability under ISO 8000-150 without turning the store into a surveillance log. [S02] [S29]

3.3

ISA-95 clarifies the boundary between control and business

Once identifiers are stable, the factory must decide which system mints each identifier and originates each operational fact. Without an explicit boundary, engineers stuff commercial order metadata (customer names, currencies, shipping dates) into PLC registers or MQTT topics, while accountants try to reconstruct stoppages from weekly payroll and inventory journals. [S05] [S07] Both fail because control systems and business systems run on different time scales, state models and reliability constraints. [S05]

Part 1 of IEC 62264-1 (ANSI/ISA-95) provides standard models and terminology for describing the interfaces between business systems and manufacturing operations and control systems: [S05]

- **Level 4 (Business planning and logistics)**: establishes the plant schedule, customer orders, procurement, cost accounting and inventory valuation over days to months. [S05]
- **Level 3 (Manufacturing operations management, or MOM)**: manages work-order dispatching, resource allocation, inline quality, maintenance execution and lot genealogy across shifts to minutes. [S05]
- **Level 2 (Supervisory control)** and **Level 1/0 (Sensing and production)**: sense variables, execute closed-loop control and count cycles over milliseconds to seconds. [S05]

Applying IEC 62264-1 to a small factory does not require buying a monolithic enterprise resource planning (ERP) suite and a separate commercial manufacturing execution system (MES). [S05] Even when Level 4 is an accounting ledger and Level 3 is a set of governed tables in GFT's manual data spine, every exchanged element has one originating boundary: [S05] [S0302]

- **Capability and resources**: asset hierarchies (`site_id`, `area_id`, `line_id`, `asset_id`) and routing steps (`operation_id`) originate at Level 3 and publish available capacity upward to Level 4. [S05]
- **Schedules and orders**: order quantities, delivery windows and style allocations (`order_id`, `style_id`) originate at Level 4 and cross downward as work requests. [S05]
- **Performance and actuals**: piece counts, run and stop durations, lot consumption (`lot_id`), defect observations (`defect_id`) and dispositions originate at Level 3 and cross upward to Level 4 as reconciled actuals. [S05] [S08]

As shown in **Exhibit 10**, p. 51, the operational vocabulary acts as the translating membrane across these levels, placing governed event contracts, entity master data and contractual KPI formulas between physical signals and management or model decisions. [S05] [S06] [S07] [S08] [S38]

EXHIBIT 10 / SEMANTIC ARCHITECTURE LAYERS

Identity and meaning sit between physical signals and decisions



DERIVED [S05] [S06] [S07] [S08] [S38] Derived semantic stack mapping physical signals, governed events, operational entities, contractual KPIs and decisions to IEC 62264-1 (S05), OPC UA (S06), MQTT (S07), GS1 EPCIS 2.0 (S08) and ISO 22400-2 (S38). Depicts logical information responsibilities, not physical network topology or an installed GFT plant.

The five layers in **Exhibit 10**, p. 51 enforce upward contextualisation. [S05] [S08] Layer 1 captures physical measurements or manual tallies via OPC UA, MQTT or tabular uploads. [S06] [S07] [S0302] Layer 2 wraps them into immutable visibility events recording what, when, where and why. [S08] Layer 3 binds those events to stable identifiers on the IEC 62264-1 boundary. [S05] Layer 4 applies contractual ISO 22400-2 formulas over those entities and time elements, and Layer 5 delivers the output to managers and approved ML models. [S04] [S38] Physical network topology and edge hardware sizing belong to wp18, which must obey the semantic contracts defined here.

3.4

A KPI dictionary makes formulas contractual

Even when two departments share the same event table, they compute conflicting numbers if aggregation formulas exist only inside private spreadsheet macros or chart queries. [S38] [S0302] In manufacturing, terms such as *availability*, *efficiency*, *first-pass yield* and *capacity utilisation* are elastic. Without one governed specification, a production supervisor and a financial controller can bring conflicting charts to the weekly review from the same underlying table. [S02] [S38]

Consider the public evidence on industrial capacity in Cameroon's textile sector. A July 2023 report cited in the inherited evidence register recorded capacity utilisation at CICAM's plants at 33.7 percent for Garoua, 48.9 percent for CICAM 1 and 10.83 percent for CICAM 2, noting that underutilisation cost about CFAF 2 billion in 2020 (an inherited source-text finding not refreshed in 2026 and representing an external report on CICAM, not a GFT operating measurement). [R12] A bare percentage cannot be interpreted without its denominator: is utilisation measured against 24-hour nameplate spindle speed, scheduled shift hours, or mechanically available equipment? [R12] [S38]

To make metrics comparable across lines and shifts, GFT registers every indicator in a contractual KPI dictionary structured under ISO 22400-2:2014 (*Key performance indicators for manufacturing operations management, Part 2: Definitions and descriptions*). [S38] Under ISO 22400-2, a KPI is presented through its formula, base elements, time behaviour, unit and contextual characteristics. [S38] Combining that standard with ISO 8000-150 accountability [S02] and W3C PROV-O lineage [S09], every KPI entry specifies nine attributes:

- **KPI identifier and name**: a unique code (`KPI-OEE-AVAIL-V1`) and approved title. [S38]
- **Decision purpose**: the decision informed and the role authorised to act. [S02] [S38]
- **Algebraic formula**: the equation in named base elements (`actual_production_time / planned_busy_time`). [S38]
- **Input base elements**: numerator and denominator definitions linked to event types (`event_type`) and states. [S08] [S38]
- **Unit and scale**: the unit (`percent`, `pieces / hour`, `kWh / kg`, `XAF / piece`) and valid domain (`0.0 to 100.0`). [S01] [S38]
- **Time behaviour**: shift close, rolling hour, or lot completion, plus rules for splitting boundary-straddling events. [S08] [S38]
- **Aggregation grain**: valid entity dimensions (`site_id`, `line_id`, `style_id`, `shift_id`) and sum-before-divide rollup rules. [S38]
- **Explicit exclusions**: rulings on reworked units, scheduled breaks and utility outages. [S38]
- **Owner and version**: the domain owner (`S02`), effective date and version (`v1.0.0`) logged with each value. [S02] [S09]

WORKED EXAMPLE: HOW UNCONTRACTED DEFINITIONS CREATE A 10.8 PERCENTAGE-POINT OEE MIRAGE ON ONE SEWING SHIFT

Consider an illustrative 8-hour shift (`480 minutes` calendar duration) on a garment line sewing woven cotton shirts. All figures are explicitly illustrative inputs constructed to demonstrate the ISO 22400-2 method, not measurements from a GFT facility. [S38]

Suppose the physical events on the line during this illustrative 480-minute shift are identical for every observer:

- Calendar shift length: `480 minutes`, with a scheduled meal break of `30 minutes` and scheduled shift-handover maintenance of `18 minutes` (illustrative).
- Unplanned utility power changeover: `24 minutes`, plus unplanned thread-break and feed-dog jam stoppages: `36 minutes` (illustrative).
- Actual running time (`actual_production_time`): `480 - 30 - 18 - 24 - 36 = 372 minutes` at a standard cycle time (`ideal_cycle_time`) of `0.75 minutes per shirt` (`45 seconds`, illustrative).
- Total shirts reaching the end of the line (`produced_quantity`): `460 shirts`, comprising `414 shirts` passing first inspection (`first_time_good_quantity`), `36 shirts` repaired in rework before shift close (`reworked_in_shift_quantity`), and `10 shirts` rejected as scrap (`414 + 36 + 10 = 460 shirts`, illustrative).

Compare how two analysts calculate Overall Equipment Effectiveness (`OEE = Availability x Performance x Quality`) without a shared ISO 22400-2 dictionary: [S38]

First, the contractual ISO 22400-2 calculation (`KPI-OEE-LINE-V1`, illustrative): [S38]

- *Availability*: Excluding scheduled breaks (`30 minutes`) and maintenance (`18 minutes`) gives `planned_busy_time = 480 - 48 = 432 minutes` (illustrative). Power changeover (`24 minutes`) and jams (`36 minutes`) remain unplanned losses, so `Availability = 372 / 432 = 86.11 percent` (illustrative). [S38]
- *Performance*: `Performance = (460 x 0.75) / 372 = 345.0 / 372 = 92.74 percent` (illustrative). [S38]
- *Quality (First-Pass Yield)*: Only first-pass units (`414 shirts`) enter the numerator while reworked shirts (`36 shirts`) are tracked separately, giving `Quality = 414 / 460 = 90.00 percent` and `Contractual OEE = 0.86111 x 0.92742 x 0.90000 = 71.88 percent` (illustrative). [S38]

Second, the uncontracted local-workbook calculation (illustrative):

- The supervisor's workbook excludes power changeover (`24 minutes`) from the denominator (`adjusted_busy_time = 432 - 24 = 408 minutes`, giving an illustrative `Availability = 372 / 408 = 91.18 percent`).
- Counting all packed shirts (`414 + 36 = 450 shirts`) as good output gives an illustrative `Quality = 450 / 460 = 97.83 percent` and `OEE = 0.91176 x 0.92742 x 0.97826 = 82.72 percent`, a gap of `10.84 percentage points` above the governed figure from the same shift. [S38]

Locking the formula, base elements and exclusions in the KPI dictionary keeps power stops visible as an infrastructure loss (`24 minutes` illustrative) and rework visible as a quality loss (`36 shirts` illustrative). [S38]

3.5

Events carry context that sensor values lack

A contractual KPI dictionary is only as reliable as the atomic records feeding its base elements. [S08] [S38] When factories digitise logs or connect sensors, they often model data in one of two fragile ways: as bare time-series tuples (`timestamp`, `tag_name`, `numeric_value`), or as mutable status rows where a cell such as `order_status` is overwritten whenever state changes. [S06] [S08] Both strip away shop-floor context.

A bare sensor tuple such as `("2026-10-30T09:14:22Z", "SEW_HEAD_14_SPEED_RPM", 3200)` records that a shaft rotated at 3,200 revolutions per minute at an instant, an illustrative reading. [S06] Alone, that tuple cannot tell an analyst whether the machine was sewing a customer order (`order_id`), running test scrap, re-stitching a rejected collar (`defect_id`), or consuming a specific cotton lot (`lot_id`). [S05] [S08] Conversely, a mutable row for bundle `LOT-2026-0842` whose `status` cell is overwritten from `CUT` to `SEWN`, `DEFECT_HOLD`, `REPAIRED` and `DISPATCHED` erases the intermediate defect and repair, hiding first-pass failures and introducing target leakage into ML training sets. [S04] [S08]

To preserve context without mutating history, the data spine models shop-floor occurrences as immutable visibility events under the GS1 Electronic Product Code Information Services (EPCIS) Release 2.0 standard across four dimensions: [S08]

- **What**: physical or digital objects identified by immutable IDs (`asset_id`, `lot_id`, `order_id`, `style_id`) and quantities with units. [S08]
- **When**: temporal coordinates separating operational `event_time` from system `ingestion_time`, so store-and-forward buffering never distorts hourly rates. [S08]
- **Where**: spatial location (`site_id`, `area_id`, `line_id`) and destination read point. [S05] [S08]
- **Why**: business context, comprising process step (`operation_id`), disposition (`conformant`, `quarantined_for_rework`, `scrapped`), governing order (`order_id`) and fault code (`defect_id`). [S08]

Under this discipline, `defect_observed`, `repair_completed`, `lot_split` and `dispatch_confirmed` are distinct, append-only event records linked by shared identifiers and provenance under W3C PROV-O. [S08] [S09] If a bundle of 50 panels is split into 45 conformant panels and 5 stained panels, a transformation event records the parent `lot_id` and two child lot IDs (`45 pieces` and `5 pieces`, illustrative quantities). [S08] Current work-in-progress state is derived as a view over the event stream, while the immutable log remains intact for KPIs and model training. [S08] [S09]

To make this structure concrete, **Exhibit 11**, p. 55 specifies the minimum ten-field envelope required for every visibility event before it enters the curated store. [S07] [S08] [S09]

EXHIBIT 11 / REQUIRED EVENT SCHEMA FIELDS

A minimum production event needs more than a timestamp and value

Field name	EPCIS / PROV role	Constraint and format	Operational failure prevented
event_id	Identity (PROV entity)	Immutable UUID or deterministic hash	Prevents duplicate counting during store-and-forward replay
event_type	Why (business step)	Controlled taxonomy verb-noun code	Stops overwritten status cells from erasing rework history
event_time	When (occurrence)	ISO 8601 UTC with local offset	Anchors KPIs to actual shop-floor shift rather than upload time
ingestion_time	When (record time)	ISO 8601 UTC assigned by broker/store	Exposes network buffering lag and late manual sheet entry
asset_id	Where / What (equipment)	Registered site-area-line-asset ID	Survives machine relocation and line display-name changes
order_or_lot_id	What / Why (object)	Registered order_id and/or lot_id	Links physical cycle counts to ledger orders and cotton lots
unit_and_value	What (quantity)	Numeric value paired with explicit UCUM/ISO unit	Prevents mixing metres, kilograms, pieces and cartons
source_method	Provenance (PROV agent)	Enum: manual_sheet, plc_counter, vision_model	Separates observed facts from model predictions or estimates

DERIVED [S07] [S08] [S09] Derived ten-field minimum production event contract synthesising transport framing from MQTT 5.0 (S07), visibility-event dimensions from GS1 EPCIS 2.0 (S08) and provenance attribution from W3C PROV-O (S09). Field names and constraints are GFT architectural specifications for the industrial data spine.

EXHIBIT 12 / REQUIRED EVENT SCHEMA FIELDS

A minimum production event needs more than a timestamp and value, continued

Field name	EPCIS / PROV role	Constraint and format	Operational failure prevented
schema_version	Contract lineage	Semantic version (e.g. v1.2.0)	Allows consumers to parse historical payloads under then-active rules
quality_state	Fitness gate	Enum: observed, estimated, corrected, quarantined	Prevents missing or imputed values from masquerading as measurements

DERIVED [S07] [S08] [S09] Derived ten-field minimum production event contract synthesising transport framing from MQTT 5.0 (S07), visibility-event dimensions from GS1 EPCIS 2.0 (S08) and provenance attribution from W3C PROV-O (S09). Field names and constraints are GFT architectural specifications for the industrial data spine.

Every field in **Exhibit 11**, p. 55 blocks a documented integration failure. [S07] [S08] [S09] Carrying both `event_time` and `ingestion_time` ensures that when an edge gateway buffers counts during a two-hour network outage (illustrative duration), the delayed upload is credited to the shift in which the garments were sewn (`event_time`) while the lag remains visible (`ingestion_time - event_time`). [S08] Requiring `source_method` and `quality_state` prevents manual entries, optical counters and imputed estimates from being silently merged. [S09] Wire serialisation and cryptographic signing belong to the transport and security layers examined next and in Chapter 07.

3.6

Transport and semantics must not be confused

A frequent procurement error in industrial digitalisation is assuming that adopting a communication protocol automatically creates a shared vocabulary. [S06] [S07] [S10] Vendors often claim that connecting machines to an MQTT broker or an OPC UA server unifies the factory data model. Moving packets across a wire and agreeing on the business meaning of the payload inside those packets are separate engineering tasks. [S06] [S07] Confusing transport with semantics produces an automated pipeline of incompatible numbers. [S10]

The contrast between OPC Unified Architecture (OPC UA, IEC 62541 / OPC 10000-1) and MQTT Version 5.0 defines the boundary between transport and meaning: [S06] [S07]

- **OPC UA (`OPC 10000-1`, version 1.05.06)**: OPC UA provides a framework for modelling information, secure communication and interoperability across devices. [S06] Its meta-model exposes typed object nodes, variables, engineering units, valid ranges and hierarchical references (including companion specifications for machinery, energy metering and textiles) in an addressable graph, exposing contextual nodes rather than raw memory addresses. [S06]
- **MQTT (`Version 5.0`, OASIS Standard)**: MQTT is a client-server publish-subscribe transport protocol suited to constrained links. [S07] MQTT 5.0 is payload-agnostic: a payload is an opaque byte array, and a topic name is a UTF-8 routing string separated by slashes. [S07]

When an engineer publishes `{"val": 42, "ts": 1793336400}` to `cameroon/plant1/weaving/loom07/picks`, MQTT 5.0 governs delivery according to Quality of Service (QoS 0, 1 or 2). [S07] It does **not** verify whether `42` is picks per minute, hundreds of picks since reset, or a fault code; whether `ts` is UTC seconds or local milliseconds; or which `lot_id` and `order_id` are running on `loom07`. [S07] [S08] A topic path is a routing address, not a governed data contract. [S05] [S07]

3.7 DECOUPLING PAYLOAD SEMANTICS FROM THE WIRE PROTOCOL PROTECTS SMALL TEAMS

Separating transport from payload semantics lets the same governed event contract (**Exhibit 11**, p. 55) travel over three transport mechanisms as the plant matures, without altering a downstream KPI formula or ML feature definition: [S06] [S07] [S0302]

1. **Stage 1 (Manual tabular transport)**: Operators record shift events on forms uploaded as CSV tables annotated with W3C Tabular Data Metadata (CSVW) schemas, which provide metadata describing how tabular data should be interpreted and bind every column to **Exhibit 11**, p. 55 with `source_method = "manual_sheet"`. [S08] [S0302]
2. **Stage 2 (Publish-subscribe transport)**: Optical counters or sub-meters on an edge gateway publish messages over MQTT 5.0 (S07) validated against a JSON Schema embodying the identical ten-field contract, setting `source_method = "plc_counter"`. [S07] [S08]
3. **Stage 3 (Contextual machine interoperability)**: Machinery with native OPC UA servers (S06) maps OPC UA nodes (`asset_id`, unit, status and timestamp) into the same canonical event envelope. [S06] [S08]

Because consumers read the governed event contract rather than binding to CSV columns, MQTT topics or PLC registers, upgrading a line from manual sheets to sensors is a non-breaking change in `source_method`. [S05] [S10] [S38]

3.8

Canonical taxonomies preserve local detail

Even when events arrive in a ten-field envelope, categorical payloads (`event_type`, `defect_id`, downtime cause, waste category and unit) fragment unless governed by canonical taxonomies. [S01] [S04] Different consumers need different granularity from the same observation:

- **Supervisors and technicians** need mechanism-specific codes: distinguishing `skip_stitch_needle_deflection` (DEF-SEW-014) from `skip_stitch_looper_timing` (DEF-SEW-015) or `oil_spot_presser_foot` (DEF-SEW-022), and `unplanned_stop_yarn_break_warp` (DT-WEV-004) from `unplanned_stop_rapier_gripper_wear` (DT-WEV-009). [S05] [S38]
- **Computer-vision engineers (wp03)** need morphological labels. According to the inherited evidence register, industry literature reports manual fabric inspection at 60 to 70 percent accuracy (missing 20 to 30 percent of defects), whereas computer vision reaches above 90 percent accuracy across more than 40 defect types at speeds up to 60 metres a minute (an inherited benchmark not refreshed in 2026 and not a GFT plant result). [R29] Training or validating vision models under ISO/IEC 5259-3 requires recording specific codes (`broken_pick`, `slub`, `reed_mark`, `oil_stain`, `shade_bar`), not `fabric_fault`. [S04] [R29]
- **Buyers and quality leads** require Acceptable Quality Limit (AQL) severity bands (`critical`, `major`, `minor`) and order summaries (wp02, wp05). [S08] [S38]

- **Sustainability analysts (`wp16`)** require offcuts, dye discharges and utility use rolled up into mass-balance and greenhouse-gas categories (`cotton_cutting_waste_kg`, `reworked_batch_steam_kwh`, `rejected_garment_scrap_kg`). [S05] [S38]

Logging only coarse categories (`Major Defect`) at entry destroys diagnostic signal: once 40 defect types are collapsed into `Major`, no query or ML pipeline can recover oil stains versus skipped stitches. [S04] [S09] Letting each team maintain private mapping tables is equally destructive. [S01] [S39]

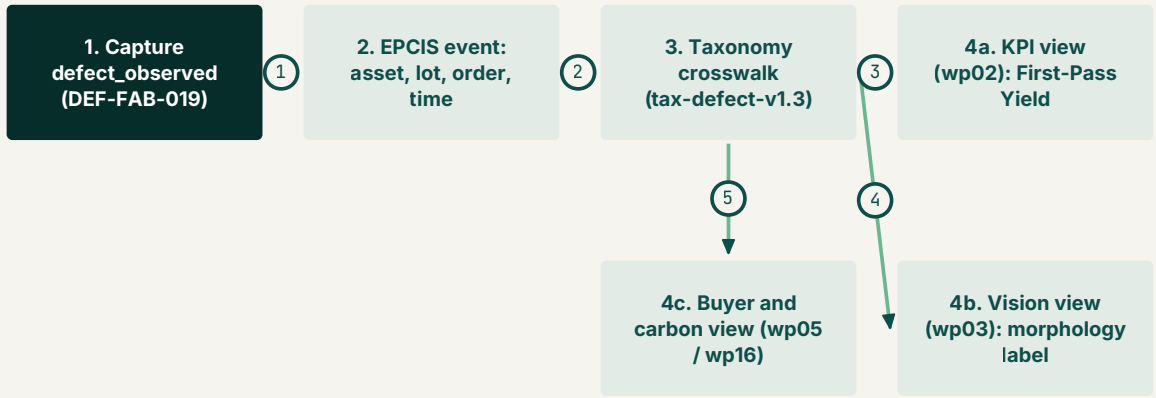
GFT enforces a two-part rule: **capture at the atomic leaf code, and project external views through versioned crosswalk tables**: [S08] [S09] [S0301]

1. **Every event stores the immutable leaf code (`defect_id`) and active taxonomy version (`taxonomy_version`) at capture time**, while each workstation displays only the 8 to 12 leaf codes relevant to its `operation_id` (illustrative subset size). [S05] [S08] [S09]
2. **Crosswalk mappings link each leaf code to each view** (`kpi_disposition_group`, `vision_class_id`, `buyer_aql_severity`, `carbon_waste_category`), with effective dates and provenance under W3C PROV-O and W3C DCAT-3. [S09] [S0301]
3. **Aggregations never overwrite the leaf code**, outputting both the rollup category and lineage to the leaf records. [S09]

As shown in **Exhibit 13**, p. 59, one governed defect observation event feeds three views without three collection routines. [S08] [S38]

EXHIBIT 13 / SEMANTIC ROUTING PATH

One governed event serves three views without three redefinitions



1. validate contract
2. retain leaf code
3. disposition rule
4. visual label map
5. AQL & mass map

ILLUSTRATIVE [S08] [S38] Illustrative process and semantic flow showing a single inline defect observation event (GS1 EPCIS 2.0 structure, S08) routed through a versioned taxonomy crosswalk to serve the ISO 22400-2 quality KPI (S38), the wp03 vision-model evaluation dataset and the wp05 customer lot-traceability view.

In **Exhibit 13**, p. 59, Node 1 records leaf code `DEF-FAB-019` (such as an illustrative weft slub above 4 millimetres). [S08] Node 2 validates the ten-field envelope (**Exhibit 11**, p. 55), binding `asset_id`, `lot_id`, `order_id` and `event_time`. [S08] Node 3 applies the versioned crosswalk (`tax-defect-v1.3`) without altering the stored event. [S09] [S0301] Node 4a feeds the shift's ISO 22400-2 First-Pass Yield calculation (`wp02`); [S38] Node 4b supplies a morphological label with lot context for vision evaluation (`wp03`); [S04] and Node 4c projects buyer AQL severity and rejected mass for traceability and carbon accounting (`wp05` and `wp16`). [S08] Verifying consistent human coding across shifts requires the quality gates in Chapter 04 and `wp03`. [S04]

3.9

Data contracts make compatibility an explicit promise

Identifiers, IEC 62264-1 boundaries, KPI formulas, EPCIS envelopes and taxonomies join in one instrument: the **data contract**. [S04] [S10] Pipelines fail when an upstream producer silently alters a feed that a consumer relies upon: changing a unit, reusing order IDs, renaming `reject_count` to `scrap_total`, or adding an unmapped downtime code. [S01] [S04] Unannounced changes make dashboards display falsehoods and feed corrupted features into ML models. [S04] [S15]

NISTIR 8107 notes that standards facilitate delivering the right information at the right time, enable actions based on that information, and reduce adoption risk. [S10] ISO/IEC 5259-3:2024 defines requirements for managing data quality in analytics and machine learning. [S04] GFT translates these standards into a **data contract**: a versioned, machine-testable specification stored in version control governing a published dataset between a named producer role and registered consumer roles. [S02] [S04] [S0302]

Every published dataset or event feed registers a contract with seven clauses: [S01] [S02] [S04] [S0301] [S0302]

1. **Identity, ownership and boundary**: dataset ID, IEC 62264-1 boundary (S05), owner and steward (S02), and registered consumers (KPI IDs, views or ML pipelines). [S02] [S05]
2. **Structural schema**: field names, types, encodings and key references (`site_id`, `line_id`, `asset_id`, `order_id`, `lot_id`) in a JSON Schema or W3C Tabular Data Metadata (CSVW) descriptor. [S0302]
3. **Required fields and quality states**: mandatory fields (`NOT NULL`), explicit `quality_state` values replacing magic numbers (`-999` or `0`), and active taxonomy versions. [S01] [S04]
4. **Physical and monetary units**: mandatory units for numeric fields (`mm`, `m/min`, `kg`, `kWh`, `pieces`, `XAF`), barring unitless numbers or unlabelled currencies. [S01] [S38]
5. **Temporal semantics**: timezone (`UTC` with offset), capture cadence (such as upload within 30 minutes of shift close, an illustrative window), and maximum lag before an alert fires. [S08] [S38]
6. **Quality gates and quarantine**: ingestion assertions routing failures to quarantine for steward correction rather than silent deletion. [S01] [S04]
7. **Semantic versioning**: `MAJOR.MINOR.PATCH` version and criteria separating backward-compatible (`MINOR`) from breaking (`MAJOR`) changes. [S0301]

3.10 SEMANTIC VERSIONING DISTINGUISHES SAFE EVOLUTION FROM BREAKING CHANGE

The W3C Data Catalog Vocabulary (DCAT Version 3) shows how a specification evolves safely: DCAT 3 relaxes constraints and adds classes and properties while preserving backward compatibility so changes do not break previous term definitions. [S0301] GFT applies three version levels to factory contracts: [S09] [S15] [S0301]

- **Patch version (`v1.0.0 to v1.0.1`)**: clarifies documentation without altering any field, unit, taxonomy code or KPI formula. [S0301]
- **Minor version (`v1.0.0 to v1.1.0`, backward-compatible)**: adds an optional field; registers a new leaf code with complete crosswalk mappings; or adds a new asset or style. [S08] [S0301] Existing KPI queries and ML pipelines run unchanged with identical historical results. [S15] [S0301]

- **Major version (`v1.1.0 to v2.0.0`, breaking change)**: renames or deletes a required field; alters a unit; redefines a KPI formula or exclusion (`S38`); or splits a leaf code without a deterministic crosswalk. [S38] [S0301] Major versions require owner approval (`S02`), a migration transform in W3C PROV-O lineage (`S09`), and ML pipeline re-validation (`S15`). [S02] [S09] [S15]

CONTRACT VERIFICATION BEFORE RELEASE AND AT INGESTION

GFT enforces contracts across workbooks and automated feeds in four steps: [S04] [S09] [S15] [S0302]

1. **Schema binding**: Every shift workbook (`CSV`) and edge topic pairs with `contract-domain-name-v1.0.json` using W3C Tabular Data Metadata (`S0302`) or JSON Schema. [S0302]
2. **Pre-release test**: Before promoting any template, tag, crosswalk or KPI change, an automated check runs synthetic records and frozen historical shifts through producer schemas and consumer queries, blocking unversioned `MAJOR` changes. [S04] [S15] [S38]
3. **Ingestion gate**: Incoming batches or streams are validated before curated storage; valid records receive a W3C PROV-O stamp (`schema_version`, `validated_at`), while failures route to quarantine for the steward (`S02`). [S01] [S02] [S04] [S09]
4. **Lineage binding**: Every report (`S38`) and ML snapshot (`S04`, `S15`) logs the `schema_version`, `taxonomy_version` and `kpi_dictionary_version` of its inputs. [S09] [S15] [S38]

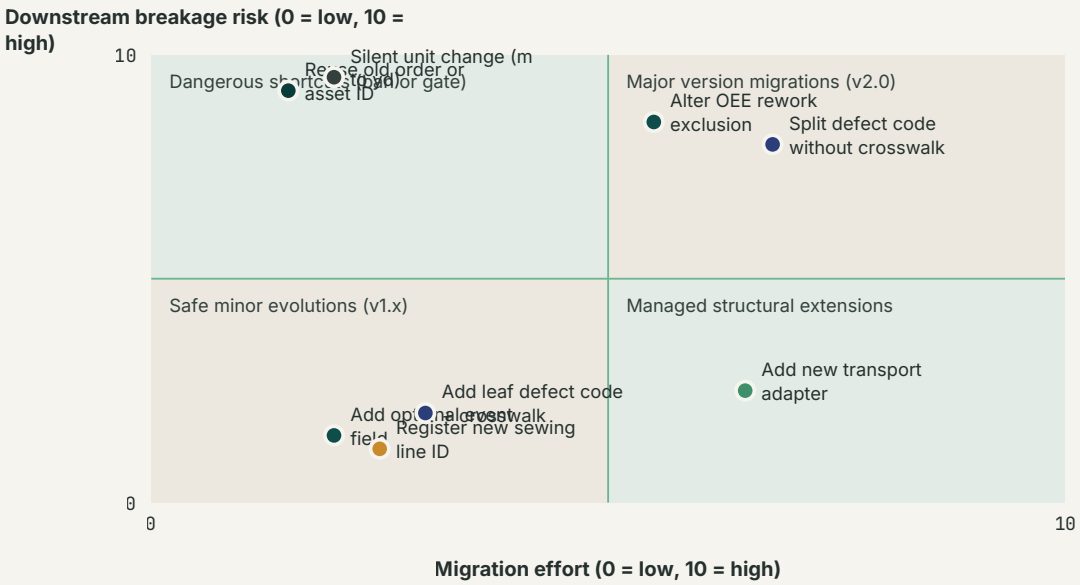
GFT NOTE: TURN SPREADSHEET KPI SHEETS INTO VERSIONED CONTRACTS FROM DAY ONE

In GFT's Cameroon textile programme, tracking starts with manual shift, quality, maintenance and order sheets before cloud warehouses or vision lines exist. [S10] [S0302] GFT recommends treating each workbook as a contracted producer from shift one: lock column headers (`site_id`, `line_id`, `shift_id`, `order_id`, `lot_id`, `event_time`, `defect_id`), restrict categorical cells via drop-downs bound to a locked taxonomy tab, ban merged cells and text in numeric columns, and keep a W3C Tabular Data Metadata (`CSVW`) JSON file beside the workbook. [S38] [S0302] This lets the team validate uploads with a 50-line script (illustrative size), feed the `kpi-cascade` and `oee-line` tools directly, and migrate later to the edge topology in `wp18` without changing an identifier or KPI formula. [S05] [S38]

To guide owners and stewards when evaluating proposed changes, **Exhibit 14**, p. 62 places eight modifications on migration effort (`x`) and downstream breakage risk (`y`). [S04] [S10]

EXHIBIT 14 / ILLUSTRATIVE GOVERNANCE INDEX (0 TO 10)

Compatibility rules distinguish safe evolution from breaking change



ILLUSTRATIVE [S04] [S10] Illustrative governance assessment of eight typical schema, taxonomy and KPI modifications on implementation/migration effort (x, 0 to 10) and silent downstream breakage risk (y, 0 to 10), derived from data-quality management principles in ISO/IEC 5259-3 (S04) and standards-integration risk in NISTIR 8107 (S10). Coordinates are illustrative ordinal scores to show contract classification rules, not empirical frequency measurements.

The four quadrants of **Exhibit 14**, p. 62 translate contract rules into shop-floor practice. [S04] [S10] In the lower-left quadrant (*Safe minor evolutions*), adding an optional event field (`x = 2.0, y = 1.5`, illustrative), registering a line ID (`x = 2.5, y = 1.2`, illustrative), or adding a leaf defect code with crosswalk mappings (`x = 3.0, y = 2.0`, illustrative) requires only a `MINOR` bump (`v1.x`) and automated contract validation. [S04] [S0301] In the lower-right quadrant (*Managed structural extensions*), adding an MQTT or OPC UA adapter that maps into the canonical event envelope (`x = 6.5, y = 2.5`, illustrative) requires engineering effort but preserves downstream compatibility. [S06] [S07]

By contrast, the upper-left quadrant (*Dangerous shortcuts*) exposes changes that cause severe damage because they look effortless at entry: reusing an old `order_id` or `asset_id` (`x = 1.5, y = 9.2`, illustrative) or changing a unit from metres to yards without renaming the field (`x = 2.0, y = 9.5`, illustrative). [S01] [S04] The contract rejects upper-left shortcuts at ingestion. [S04] Finally, the upper-right quadrant (*Major version migrations*) covers legitimate evolutions that alter meaning, such as revising how rework enters an ISO 22400-2 OEE formula (`x = 5.5, y = 8.5`, illustrative) or splitting a defect code without a one-to-one mapping (`x = 6.8, y = 8.0`, illustrat-

ive), which require a MAJOR release (v2.0), dual-version reporting and model re-evaluation. [S15] [S38] Executing automated quality gates, quarantining bad records and tracing provenance across transformations is the task of Chapter 04.

A factory does not share a fact until the line and the ledger mean the same thing by its name.

CHAPTER SOURCE NOTES

S05 International Electrotechnical Commission. 2013.

S08 GS1 AISBL. 2022.

S38 International Organization for Standardization. 2014.

S02 International Organization for Standardization. 2022.

R06 Cameroon Trade Hub. Publication date unrecorded; inherited source-text check.

R07 Ecofin Agency. Publication date unrecorded; inherited source-text check.

S04 ISO/IEC. 2024.

S39 National Institute of Standards and Technology. 2025.

S10 National Institute of Standards and Technology. 2016-02-23.

S0302 World Wide Web Consortium. 2015-12-17.

S01 International Organization for Standardization. 2022.

S09 World Wide Web Consortium. 2013-04-30.

S06 OPC Foundation. 2025-10-22.

S29 European Parliament and Council of the European Union. 2016-04-27.

S31 Presidency of the Republic of Cameroon. 2024-12-23.

S07 OASIS Open. 2019-03-07.

R12 Business in Cameroon. Publication date unrecorded; inherited source-text check.

R29 MDPI; industry inspection literature. Publication date unrecorded; inherited source-text check.

S0301 World Wide Web Consortium. 2024-08-22.

S15 Google Cloud. 2024-08-28.

GREEN FUTURE TECH

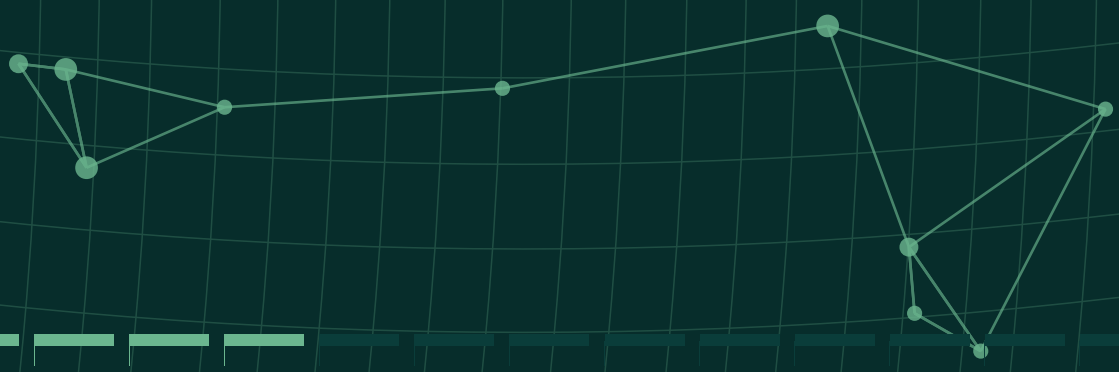
LIMITED / HONG KONG SAR

04

PART I: GOVERN MEANING BEFORE SCALE / FITNESS AND LINEAGE

Quality gates make fitness visible before errors compound

A small set of enforced checks, visible exceptions and complete lineage is more useful than an unmeasured claim that the data are clean.



CHAPTER 04
DECISION TEST

A small set of enforced checks, visible exceptions and complete lineage is more useful than an unmeasured claim that the data are clean.

A factory record is fit for use only when its defects, omissions and transformations are explicitly measured and preserved before the record enters a management report or a training pipeline. [S01] In a small manufacturing team, declaring a dataset clean by inspection is an operational hazard. Unchecked capture errors do not vanish when rows are aggregated into shift summaries or fed into machine-learning features; they compound silently into distorted key performance indicators (KPIs), misallocated maintenance hours and brittle model predictions. [S18] The practical remedy is not an enterprise data-cleansing department. It is a compact set of automated and steward-enforced quality gates placed at capture and ingestion, paired with explicit quality states and end-to-end provenance links connecting every derived metric or prediction back to the shop-floor events that produced it. [S04] When a record fails a rule, the industrial data spine isolates the offending event in a visible quarantine queue and routes a repair task to the originating domain steward, keeping curated tables uncontaminated while preserving raw evidence. [S09]

This discipline is foundational for Green Future Tech Limited (GFT), Hong Kong SAR, in its three-part Cameroon textile programme covering feasibility, performance-management system design and digital strategy. According to the programme's inherited evidence base, prepared on 31 January 2024 and not refreshed in 2026, with original source publication dates unrecorded, about 95 percent of Cameroon's cotton fibre is exported raw and about 5 percent is processed locally, against a national target of 50 percent local processing by 2030. [R06] [R07] Across existing domestic mills, a July 2023 report cited in that inherited register recorded capacity utilisation at 33.7 percent in Garoua, 48.9 percent at CICAM 1 and 10.83 percent at CICAM 2, with underutilisation costing about CFAF 2 billion in 2020. [R12] Those sector figures describe external industrial conditions rather than GFT operating results, yet they establish why data quality cannot wait. A factory closing such utilisation gaps on tight margins cannot base capital allocation, shift scheduling, quality disposition or demand forecasting on unverified spreadsheet entries or silently patched logs. [S36]

Where earlier chapters established why trusted records matter, who holds decision rights and how operational vocabularies define shared meaning, this chapter specifies how the data spine tests whether incoming records honour those contracts. [S02] It translates international data-quality and provenance standards into seven shift controls: turning quality dimensions into grain-specific tests, enforcing validation and quarantine at capture and ingestion, separating missing and estimated values from observations, versioning rules alongside schemas, recording provenance across KPIs and pipelines, improving capture iteratively through a quality register, and halting upstream data cascades before they corrupt models. [S03] [S15]

4.1

Quality dimensions become operational tests

Data quality is the degree to which stated characteristics of data fulfil explicit requirements for an intended task. [S01] Standardised frameworks such as ISO/IEC 25012 categorise data-quality requirements and measures so that data can be evaluated independently from the applications that process them. [S03] Similarly, the World Wide Web Consortium (W3C) Data Quality Vocabulary emphasises that pub-

lishing consistent quality measures allows a dataset consumer to judge fitness for purpose rather than relying on an unverified claim of completeness. [S0401] For analytics and machine learning, ISO/IEC 5259-3 extends this requirement across training, validation and operational production feeds. [S04]

4.2 FIVE PRACTICAL DIMENSIONS EVALUATE FITNESS WITHOUT CLAIMING AN EXHAUSTIVE TAXONOMY

Standards catalogues list over a dozen quality characteristics, but a small team scoring every theoretical dimension across every table will produce administrative theatre rather than safer decisions. [S03] Within the evaluation approach of ISO/IEC 25012, this paper selects five operational dimensions as its practical test set: validity, completeness, consistency, timeliness and uniqueness. [S03] This five-part selection is an engineering design choice for small industrial teams, not a claim that international standards mandate only these five categories. [S04]

Each dimension answers a distinct question between shop-floor capture and curated storage:

- **Validity** tests whether a value conforms to its contracted data type, physical range, unit and allowed code list under the active schema. [S01] A sewing-line record reporting negative 15 garments, a roll width without a unit, or an unrecognised fault code fails validity immediately. [S04]
- **Completeness** tests whether required fields are populated and whether all expected records for a scheduled interval have arrived. [S03] An hourly log omitting the style code fails field completeness, whereas a shift with only 6 of 8 scheduled hourly logs submitted fails grain completeness. [S38]
- **Consistency** tests whether representations of the same physical reality agree across records or systems without contradiction. [S03] If a loom state log reports 60 minutes of mechanical stoppage during an hour in which the output log reports 45 metres of woven fabric on that same asset, the records fail consistency. [S05]
- **Timeliness** tests whether the delay between physical occurrence (event time) and validated availability (ingestion time) falls within the decision window. [S08] A defect log entered 36 hours after shift handover fails timeliness for morning quality disposition. [S36]
- **Uniqueness** tests whether each physical event or entity is represented once at the stated grain, preventing message replay or double-entered sheets from inflating counts. [S08]

4.3 EVERY QUALITY TEST BINDS TO A SPECIFIC GRAIN, DENOMINATOR AND ACCOUNTABLE ROLE

An operational quality test is executable only when it specifies the grain of evaluation, the pass condition, the summary denominator and the owner who acts when the test fails. [S02] In ISO 22400-2, manufacturing KPIs require explicit formulas, time behaviours and boundaries; quality tests on underlying inputs need the same precision. [S38] Specifying that every closed corrective work order (`grain: work_order_id`) must contain a valid asset identifier, fault code, repair duration in minutes and restart timestamp before closure (`owner: Maintenance Steward`) turns an aspiration into a deterministic gate. [S02] [S05]

Before building dashboards or training models, a factory should map active rule coverage across its core domains, as illustrated in **Exhibit 15**, p. 67. [S03] [S04]

EXHIBIT 15 / ILLUSTRATIVE COVERAGE SCORE (0 TO 4)

A rule-coverage map reveals where fitness is assumed rather than tested



ILLUSTRATIVE [S03] [S04] [S0401] Illustrative scoring model (0 = untested, 1 = manual spot check, 2 = entry syntax check, 3 = automated gate with quarantine, 4 = versioned gate with full lineage) across four factory record sets and five quality dimensions, derived from ISO/IEC 25012 (S03), ISO/IEC 5259-3 (S04) and W3C DQV (S0401). Values are illustrative and not a measured GFT factory result.

The illustrative matrix in **Exhibit 15**, p. 67 scores each intersection from 0 (no test) to 4 (versioned gate with quarantine and lineage), adapting ISO/IEC 25012, ISO/IEC 5259-3 and the W3C Data Quality Vocabulary. [S03] [S04] [S0401] In this illustrative baseline, customer orders and production logs reach scores of 3 to 4 on syntax validity and uniqueness because entry templates enforce primary keys and numeric cells. Quality inspections and maintenance tickets score 1 on completeness and consistency because operators record free-text notes on clipboards and transcribe them after the shift without cross-checking asset states or lot identifiers. [S05] What **Exhibit 15**, p. 67 cannot establish is the actual defect rate of the records; a high coverage score proves only that a test runs, not that the shop-floor process is error-free. Its operational value is to stop a team from training a model on domains where completeness and consistency remain at level 1. [S04]

4.4

Validation belongs at capture and ingestion

Discovering a broken identifier or impossible unit weeks later during feature engineering is the costliest way to manage factory data. [S18] Upper-bound estimates from semiconductor manufacturing workshops indicate that data processing for a single application can reach 70 percent of total cost when teams repeatedly clean raw logs for each downstream tool. [S39] Under ISO 8000-1 and ISO/IEC 5259-3, validation belongs at two early boundaries: syntax enforcement at capture, followed by contract validation at ingestion before curated storage. [S01] [S04]

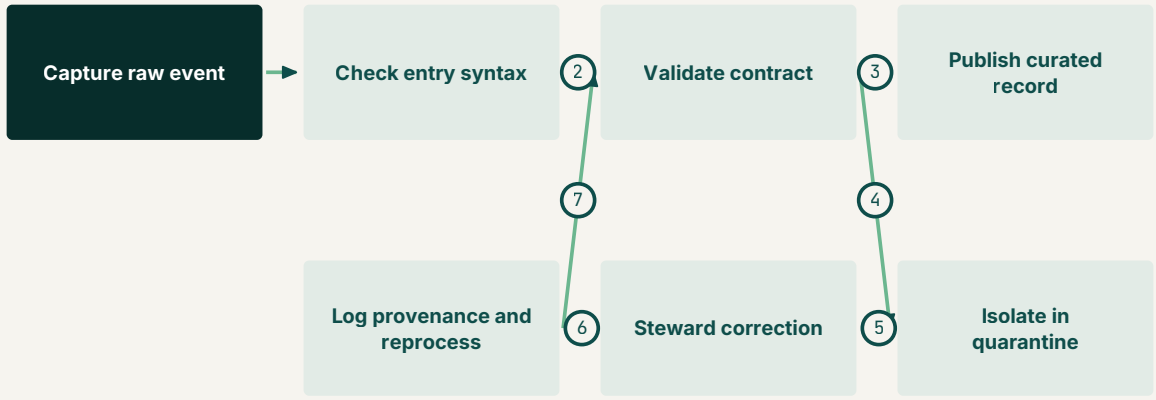
4.5 ENTRY CONSTRAINTS BLOCK IMPOSSIBLE SYNTAX WHILE INGESTION GATES ISOLATE CONTRACT FAILURES

At capture, whether an operator types into a tablet at a weaving beam, scans a bar-code on a dye lot or configures an edge gateway polling a controller, the interface should make syntactic invalidity impossible to submit. [S06] Picklists bound to master data prevent invented style codes, unit-locked inputs prevent kilogram weights from entering as bale counts, and barcode scans eliminate transposition errors. [S05]

Entry constraints cannot catch every defect. An edge buffer flushing messages after a network outage, a laboratory batch import or a cross-record check between production and maintenance can only be evaluated at ingestion. [S08] Dropping non-conforming records destroys physical evidence, whereas writing them directly into curated tables with unread warnings contaminates downstream KPIs and training sets. [S19] The solution is a formal quarantine gate between raw capture and curated publication, shown in **Exhibit 16**, p. 68. [S01] [S04] [S09]

EXHIBIT 16 /

Quarantine keeps bad records visible without contaminating curated data



2. Valid syntax
3. Pass gate
4. Fail rule
5. Assign owner
6. Superseding event
7. Re-evaluate

DERIVED [S01] [S04] [S09] Derived operational control flow linking raw capture, entry syntax check, ingestion contract gate, quarantine isolation, steward remediation and provenance-logged reprocessing, based on ISO 8000-1 (S01), ISO/IEC 5259-3 (S04) and W3C PROV-O (S09).

In **Exhibit 16**, p. 68, every incoming record first lands in an immutable raw log so shop-floor evidence is preserved. [S09] Records passing syntax checks are tested against the active contract at ingestion, verifying foreign keys, units, state transitions and uniqueness. [S04] Passing records enter the curated store; failing records enter

quarantine with the failed rule ID, contract version and rejection timestamp. [S01] While **Exhibit 16**, p. 68 defines the control path, it cannot guarantee timely resolution without an enforced steward workload allocation and escalation routine. [S02]

4.6 ROUTING REPAIR TASKS TO THE ORIGINATING STEWARD STOPS SILENT DOWNSTREAM PATCHING

When a record enters quarantine, a central analyst or IT administrator must never edit the number quietly in the database to clear an error. [S02] The analyst was not on the floor and cannot know whether an anomalous figure reflects a typo, a split lot or an unlogged fault. [S19] Moreover, silent downstream repair shields the originating department from its own capture defects, guaranteeing recurrence. [S18]

Under ISO 8000-150 role accountability, each quarantine rule routes a repair task to the originating domain steward. [S02] The task cites the record ID, failed rule (for example, illustrative check `CONSISTENCY_04: output of 120 units exceeds maximum speed of 95 units per hour for style ST-104`), shift and asset. [S38] The steward verifies the physical tally and submits a superseding correction event citing the raw event ID, reason code and role identity, which re-enters the contract gate in **Exhibit 16**, p. 68 before curated release. [S09]

4.7

Missing is not zero and estimated is not measured

In factory spreadsheets and poorly designed databases, a blank cell, a sensor drop and a genuine zero-output hour often collapse into `0` or an untyped `NULL` coerced to zero. [S04] In manufacturing, that conflation is destructive. An hour in which a ring-spinning frame produced zero kilograms of yarn because it stopped for a scheduled bearing replacement differs fundamentally from an hour in which the frame ran normally but the clerk forgot the tally, or an hour in which a power dip dropped the gateway while the machine ran on backup generation. [S38]

4.8 FIVE EXPLICIT QUALITY STATES SEPARATE PHYSICAL STOPPAGE FROM UNRECORDED OBSERVATION

To prevent absence or approximation from masquerading as measured production, every value-bearing record and time bucket must carry an explicit quality-state attribute alongside its numerical value. [S0402] Following W3C Data on the Web Best Practices on documenting quality issues and data unavailability in metadata, every record is tagged with one of five mutually exclusive states: [S0402]

1. `OBSERVED`: Captured directly from the primary instrument, scan or verified tally within the timeliness window and passed all validation gates. [S01]
2. `MISSING`: An expected observation for a scheduled interval or mandatory field did not arrive by cutoff; the value field remains empty rather than defaulting to zero. [S0402]
3. `ESTIMATED`: Imputed using a documented, approved rule when primary measurement was unavailable, such as deriving fabric length from beam revolutions or allocating daily steam across shifts by run-time hours. [S41]

4. CORRECTED: Superseded by an authorised steward correction with a recorded reason code and provenance link after failing a gate or reconciliation. [S09]
5. NOT_APPLICABLE: The asset or line was formally unscheduled or outside the measure boundary during the interval, so no observation was expected. [S38]

4.9 PROPAGATING STATE METADATA INTO AGGREGATES PREVENTS FALSE PRECISION IN MANAGEMENT REVIEWS

Storing quality states at the event level helps only if those states propagate through KPI formulas, management tables and feature pipelines. [S04] When an hourly KPI engine aggregates eight shift hours under ISO 22400-2, it must compute both the metric value and the quality-state composition across contributing intervals. [S38]

METHOD

Quality-state propagation rule for operational aggregates and model features

1. **System boundary**: Applies to any roll-up from atomic or hourly shop-floor records (event_id or hour_bucket) to shift, daily, weekly or lot-level KPIs and feature vectors. [S38]
2. **State precedence in roll-ups**:
3. Marked OBSERVED only when 100 percent of scheduled constituent intervals are OBSERVED or CORRECTED with a verified audit trail. [S04]
4. Marked ESTIMATED if any constituent interval is ESTIMATED and none are MISSING, carrying the estimated input share. [S41]
5. Marked INCOMPLETE if any scheduled constituent interval is MISSING, excluded from automated model ingestion, and displaying both the partial sum and missing-interval count without substituting zero. [S0402]
6. **Verification check**: Before freezing any dataset snapshot, a gate verifies that OBSERVED, CORRECTED, ESTIMATED, MISSING and NOT_APPLICABLE intervals sum to the exact calendar capacity of the window. [S03]

A numerical example shows how separating MISSING from 0 protects both management action and model training. [S38]

WORKED EXAMPLE

Illustrative sewing-line availability and output distortion from conflating missing logs with zero production

Consider an illustrative 8-hour shift on sewing line `Line S-02` scheduled for all 8 hours (480 minutes) at a target rate of 50 shirts per hour (400 shirts per shift), evaluated under ISO 22400-2 definitions. [S38] Assume the following illustrative sequence across the 8 one-hour buckets:

- **Hours 1 to 5 (5 hours)**: In this illustrative normal run, the line logs 48, 52, 50, 49 and 51 shirts (`status: OBSERVED`), with 60 run minutes per hour (300 minutes total; 250 shirts). [S38]
- **Hour 6 (1 hour)**: In this illustrative stoppage, a needle-bar mechanism jams for the full hour, logging 0 run minutes and 0 shirts (`status: OBSERVED`, code `MECH_JAM`). [S38]
- **Hours 7 and 8 (2 hours)**: The mechanic clears the jam at the end of Hour 6, and the line resumes at an illustrative 50 shirts per hour (100 shirts across 2 hours, 120 run minutes), but tablet Wi-Fi drops and the final 2 hourly logs are not synced (`status: MISSING`). [S0402]

Case A: Naive pipeline (missing buckets coerced to zero) In this illustrative naive calculation, a script treats the 2 missing hourly logs as 0 run time and 0 output across 8 scheduled hours, yielding 300 run minutes out of 480 (`62.5%` apparent availability) and 250 shirts (`31.25` shirts per scheduled hour). [S38] Management sees an illustrative availability drop to 62.5 percent and 3 hours of stoppage, triggering a false maintenance escalation, while a forecasting model learns a downwardly biased output rate. [S19]

Case B: Governed quality-state spine (missing buckets isolated) Under the governed spine, the system retains `MISSING` on Hours 7 and 8 and refuses to coerce them to zero: [S0402]

- **Observed coverage**: Reports 6 of 8 scheduled hours (`75.0%` illustrative coverage, with 2 hours `MISSING` isolated in quarantine). [S04]
- **Observed-window availability**: Over the 6 observed hours, illustrative availability is 300 run minutes divided by 360 observed scheduled minutes (`83.3%`, not `62.5%`). [S38]
- **Observed-window output**: Records 250 shirts across 6 observed hours (`41.67` shirts per observed scheduled hour, or `50.0` shirts per running hour). [S38]
- **Decision consequence**: Instead of reprimanding maintenance for an illustrative phantom 3-hour fault, the morning review sees a 75.0 percent completeness alert, prompts the steward to sync Hours 7 and 8, and restores the true illustrative 8-hour record (420 run minutes out of 480, or `87.5%` availability and 350 shirts) before publication. [S04] [S38]

4.10

Rules and exceptions are versioned

Factory operations change continuously. A mill introduces blended yarn counts, recalibrates moisture sensors before the rainy season, updates defect thresholds for a new buyer or splits a downtime code. [S38] Overwriting validation rules and code lists in place without version control makes historical records uninterpretable, causing later audits to reject valid historical rows or misclassify passes. [S04]

4.11 VALIDATION THRESHOLDS TRAVEL WITH SCHEMA VERSIONS SO HISTORICAL AUDITS REMAIN REPRODUCIBLE

Under ISO/IEC 5259-3 and W3C Data on the Web Best Practices, every validation rule set, range threshold and code list must be versioned alongside the data contract it enforces. [S04] [S0402] Each gate evaluation stores not only `pass` or `fail`, but the tuple `(schema_version, rule_set_version, evaluated_at)`. [S09]

Suppose an illustrative contract `fabric_roll_inspection` runs under `schema_v1.1` and `rules_v1.3` from Month 1 to Month 4, with valid cotton twill grammage set at 180 to 220 grams per square metre (`g/m2`). In Month 5, the mill runs a poplin order at an illustrative 135 `g/m2`. Rather than mutating `rules_v1.3` globally, the owner and steward publish `rules_v1.4`, binding grammage bounds to the style master record (`style_id → [min_gsm, max_gsm]`) from an explicit timestamp. [S02] [S05] Historical rolls from Month 2 remain valid under `rules_v1.3`, new poplin rolls validate under `rules_v1.4`, and training queries spanning Months 1 to 6 inspect rule-version metadata to verify how each slice was screened. [S15]

4.12 TIME-BOUNDED WAIVERS PREVENT TEMPORARY PRODUCTION WORKAROUNDS FROM BECOMING PERMANENT BLIND SPOTS

Urgent situations occasionally require relaxing a strict gate temporarily, such as when a laboratory scale interface fails mid-week or a trial lot runs before its buyer style code enters the order table. [S05] If the gate offers only a binary choice between halting capture and permanently disabling the check, operational pressure will disable the check, leaving much of the suite inactive within months. [S18]

To preserve discipline during exceptions, the spine uses time-bounded waivers governed by named decision rights. [S02] [S11] Every waiver requires four fields in the governance log:

- **Scope and rule identifier**: The exact rule, asset or lot boundary affected (for example, waiving `FK_BUYER_STYLE` only for illustrative trial order `TR-2026-09` on `Line W-01`). [S04]
- **Accountable approver**: The domain owner authorising the exception after assessing downstream risk under NIST AI RMF principles. [S11]
- **Compensating control and quality state**: Admitted records are tagged `ESTIMATED` or `WAIVED_PENDING_RECONCILIATION`, excluding them from automated model training until reconciled. [S0402]
- **Hard expiry timestamp**: An automatic expiration date (such as an illustrative 7 calendar days from approval) at which the gate reactivates unless reconciliation is formally closed. [S02]

4.13

Lineage connects every derived value to its origins

Even when every event passes a versioned gate, raw and curated events are only the start of the chain. Before a figure reaches a review or a model outputs a defect probability, data undergo filtering, unit conversion, time-bucket joins, KPI calculation, snapshotting and feature transformation. [S15] [S38] When an unexpected number appears or a model makes a false rejection, the team must walk backward from that output to the source records, code versions and approvals that created it. [S09]

4.14 ENTITY, ACTIVITY AND AGENT RELATIONS RECORD HOW RAW EVENTS BECOME CURATED FACTS

To represent lineage without vendor lock-in, the spine adopts the core model of the W3C PROV Ontology (PROV-O) and the W3C PROV Model Primer: [S09] [S0403]

- **Entities (`prov:Entity`)**: Physical or digital things with fixed aspects, such as an illustrative raw visibility event (`raw_event_88412`), a curated hourly table (`curated_hourly_prod_2026W14`), a frozen training snapshot (`ds_vision_train_v03`), or a model binary (`model_defect_cnn_v1.1`). [S09] [S0403]
- **Activities (`prov:Activity`)**: Processes occurring over a time interval that act upon or generate entities, such as an ingestion validation run, a steward correction, an ISO 22400-2 KPI job, or a training pipeline run. [S09] [S38]
- **Agents (`prov:Agent`)**: Persons, roles or services bearing responsibility for an activity or entity, such as `Quality Steward`, `Production Owner`, or `ingestion-worker-v2.1`. [S02] [S0403]

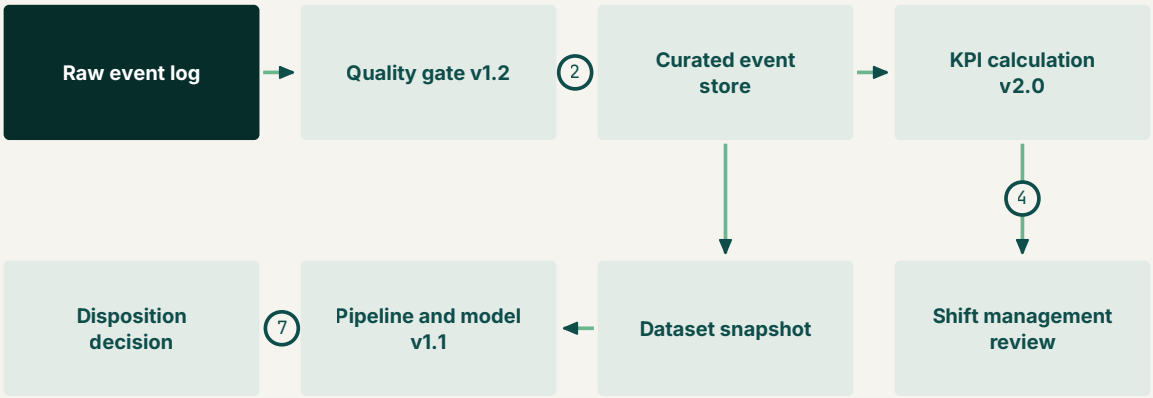
Recording core provenance relations (`wasGeneratedBy`, `used`, `wasDerivedFrom`, `wasAttributedTo` and `wasAssociatedWith`) as relational rows alongside pipeline runs captures complete lineage without a separate graph database, letting consumers judge trust and context accurately. [S09] [S0402]

4.15 SHARED PROVENANCE LINKS ALLOW A BOARDROOM KPI AND A MODEL SCORE TO BE AUDITED ON THE SAME SPINE

A unified data spine lets management reporting and machine-learning operations (MLOps) share the same upstream provenance graph rather than maintaining disconnected silos. [S15] Recording execution metadata for every ML pipeline run supports data and artefact lineage, reproducibility and comparison. [S15] Linking pipeline metadata to curated operational identifiers lets the factory reconstruct both a management KPI and a model decision down to the same raw events, as shown in **Exhibit 17**, p. 74. [S09] [S15] [S0403]

EXHIBIT 17 /

Lineage reconstructs a KPI and a model prediction on separate paths



2. Validated
4. KPI path
7. Model path

DERIVED [S09] [S15] [S0403] Derived provenance graph linking raw visibility event, versioned quality gate, curated record, KPI formula execution, shift review, frozen dataset snapshot, pipeline run and model decision, based on W3C PROV-O (S09), Google Cloud MLOps lineage requirements (S15) and W3C PROV Primer (S0403).

In **Exhibit 17**, p. 74, a single observation enters the raw log, passes `Quality gate v1.2` and lands in the curated store with its provenance hash. [S09] Two auditable paths branch outward: `KPI calculation v2.0` produces the shift management figure under ISO 22400-2, while a frozen `Dataset snapshot` feeds `Pipeline and model v1.1` under S15 to support a live `Disposition decision`. [S15] [S38] While **Exhibit 17**, p. 74 cannot prove by itself that the KPI formula or feature code is conceptually sound, it allows the custodian to query forward (`wasDerivedFrom`) whenever a sensor fault is found, identifying every KPI report, snapshot and model candidate that absorbed the tainted interval. [S09] [S15]

To keep capture lightweight, the spine logs provenance at four checkpoints: (1) ingestion from raw batch to curated partition (`schema_version`, `rule_set_version`), (2) steward correction from quarantined event to superseding event (`steward_role`, `reason_code`), (3) KPI calculation (`kpi_formula_version`), and (4) MLOps execution (`snapshot_id`, `pipeline_run_id`, `model_version_id`). [S09] [S15] Standard SQL queries across those four append-only tables satisfy ISO/IEC 5259-3 and NIST AI RMF reproducibility obligations without enterprise catalogue costs. [S04] [S11]

4.16

Quality improves in measured iterations

Small factories often stall digitalisation by treating data quality as a binary precondition: either waiting until every legacy sheet is flawless, or ignoring quality and rushing into pilots. [S36] Waiting paralyses action, while ignoring quality produces outputs that operators quickly abandon. [S19]

4.17 AN ITERATIVE QUALITY LOOP PRIORITISES RECURRING DECISION RISK OVER COSMETIC COMPLETENESS

The practical alternative follows the Greenhouse Gas (GHG) Protocol's Scope 3 guidance: collecting data, assessing data quality and improving data quality is an iterative process. [S41] Applied to factory records, the team captures essential events, assesses where defects threaten high-value decisions, and invests limited stewardship hours in fixing the top recurring failure modes. [S36] [S41] A typo in an optional comment has near-zero consequence, whereas a dye-lot mismatch between the slasher-dryer sheet and weaving-beam ticket breaks lot genealogy (`wp05` and `wp16`) and corrupts shade analysis (`wp03`). [S08] Each iteration follows four steps inside existing reviews (`wp02`): [S36] [S41]

- **Measure baseline gate failures**: Count weekly quarantined, `MISSING`, `ESTIMATED` and `CORRECTED` records by domain and rule, plus steward repair hours. [S03] [S04]
- **Rank by decision consequence**: Score failures by whether they distort an active KPI, block a model gate (`wp03` or `wp04`), or threaten compliance (`wp05` or `wp16`). [S11] [S38]
- **Execute structural capture fixes**: Replace manual typing with barcode scans or lock picklists to master data for the top issues. [S01] [S10]
- **Verify recurrence reduction**: Confirm after two to four weeks that rule failures dropped below target before closing the item. [S36] [S41]

4.18 A LIVE QUALITY REGISTER TURNS RECURRING QUARANTINE EVENTS INTO FUNDED ENGINEERING FIXES

To make iteration visible, the steward maintains a live quality register, shown in **Exhibit 18**, p. 76, linking recurring quarantine patterns to affected decisions, owners and verification windows. [S02] [S19] [S41]

EXHIBIT 18 / ILLUSTRATIVE REGISTER ENTRIES

A quality register records the issue, consequence and next correction

Issue and domain	Quality state	Affected decision	Weekly count	Owner	Corrective action and verify window
QR-01 Production	MISSING	Shift OEE review	14	Prod Steward	Enable tablet offline queue alert (14d)
QR-02 Quality	INVALID	Roll disposition	19	Qual Steward	Lock defect picklist to v2.1 taxonomy (10d)
QR-03 Maintenance	INCONSISTENT	Downtime Pareto	11	Maint Steward	Reconcile stop time with loom state (21d)
QR-04 Production	LATE	Morning line balance	22	Line Lead	Move tally cutoff to shift handover (7d)
QR-05 Orders	DUPLICATE	Dye-lot scheduling	6	Order Steward	Enforce unique order-line key (14d)
QR-06 Quality	MISSING	Vision training set	17	Qual Owner	Require roll barcode scan at camera (21d)
QR-07 Sustainability	ESTIMATED	Batch energy intensity	8	Sust Steward	Install sub-meter on stenter frame (45d)
QR-08 Maintenance	MISSING	Spare-parts reorder	9	Maint Steward	Require part code on ticket close (14d)

ILLUSTRATIVE [S19] [S41] [S0402] Illustrative twelve-issue factory quality register showing issue ID and domain, quality state, affected decision, illustrative weekly recurrence count, accountable owner and structural corrective action with verification window. Issue tracking and iterative improvement structure derived from Sambasivan et al. (S19), GHG Protocol iterative quality guidance (S41) and W3C DWBP metadata disclosure (S0402). All counts and windows are illustrative.

EXHIBIT 19 / ILLUSTRATIVE REGISTER ENTRIES

A quality register records the issue, consequence and next correction, continued

Issue and domain	Quality state	Affected decision	Weekly count	Owner	Corrective action and verify window
QR-09 Production	INCONSISTENT	Piece-weight yield	12	Prod Owner	Standardise tare weight on scale (10d)
QR-10 Quality	LATE	Lab shade approval	7	Lab Lead	Import spectrop hotometer CSV live (28d)
QR-11 Orders	INVALID	Delivery promise date	5	Order Owner	Block past-date entry on amendment (7d)
QR-12 Sustainability	MISSING	Effluent pH compliance	4	Sust Owner	Add shift-end check before drain (14d)

ILLUSTRATIVE [S19] [S41] [S0402] Illustrative twelve-issue factory quality register showing issue ID and domain, quality state, affected decision, illustrative weekly recurrence count, accountable owner and structural corrective action with verification window. Issue tracking and iterative improvement structure derived from Sambasivan et al. (S19), GHG Protocol iterative quality guidance (S41) and W3C DWBP metadata disclosure (S0402). All counts and windows are illustrative.

The twelve illustrative rows in **Exhibit 18**, p. 76 turn abstract dimensions into engineering tasks. [S41] Illustrative issue QR-02 logs 19 weekly quarantine events for obsolete defect codes (status: INVALID), threatening roll disposition and wp03 evaluation sets; the fix is a 10-day task locking the picklist to taxonomy v2.1. [S04] Illustrative issue QR-07 discloses that 8 weekly batch energy allocations use ESTIMATED steam splits, guiding a 45-day stenter sub-meter installation. [S41] [S0402] While **Exhibit 18**, p. 76 cannot resolve cross-departmental trade-offs automatically, tracking remediation hours exposes hidden preparation costs: eliminating an illustrative 4 hours a week of manual reconciliation saves about 200 hours a year of recurring toil. [S02] [S18] [S39]

4.19

Upstream checks limit downstream cascades

Empirical evidence on AI failures reinforces why quality gates precede modelling. Studying high-stakes AI practitioners at ACM CHI 2021, Sambasivan et al. defined *data cascades* as compounding events causing negative downstream effects from data issues, finding them pervasive (92 percent prevalence in their interviewed sample), invisible, delayed, but often avoidable. [S19] Though that 92 percent figure describes their qualitative sample rather than a statistical survey of Cameroonian manufacturers, the mechanism is a direct warning for industrial MLOps. [S19]

4.20 DELAYED GROUND TRUTH AND SILENT FIELD CHANGES TRIGGER COMPOUNDING DATA CASCADES

Data cascades arise when upstream carelessness in calibration, labelling or schema evolution interacts with production over time. [S19] In textile inspection, literature notes that manual visual inspection is roughly 60 to 70 percent accurate and misses 20 to

30 percent of defects, whereas computer vision exceeds 90 percent accuracy across more than 40 defect types at speeds up to 60 metres a minute (inherited entry R29, prepared 31 January 2024, not refreshed in 2026, original date unrecorded). [R29] Yet validating a vision model against manual logs lacking a governed taxonomy or roll-encoder consistency gate makes the validation score an illusion. [S04] [S19]

Three cascade paths are especially destructive in small factories:

1. **Silent unit or sensor drift**: A weaving-shed humidity sensor drifts out of calibration, or a firmware update alters scaling from tenths of a millimetre to whole millimetres. [S06] Without an ingestion range gate, feature pipelines ingest shifted numbers as process changes, triggering false alarms until operators mute the system. [S18]
2. **Label contamination from unlogged workarounds**: During an end-of-month shipment push, supervisors pass minor slub defects on secondary panels without logging a waiver or changing the grade code. [S02] Retraining on those passed rolls as defect-free ground truth teaches the model to ignore slub defects on premium orders later. [S19]
3. **Delayed outcome joins across broken identifiers**: In demand forecasting (wp04) or return analysis, outcomes arrive weeks after production. [S43] If order amendments or lot splits overwrote the original identifier instead of appending a provenance-linked event under EPCIS and PROV-O rules, delayed outcomes join to the wrong production record. [S08] [S09]

> **WARNING**
>
> **Unresolved upstream quality failures must halt model promotion**
>
> When a candidate model shows poor validation accuracy or shadow-mode drift, teams are tempted to tune hyperparameters or trigger automated retraining. [S18] Under this paper's governance rules and the stop/postpone principle in SME AI cases, no model may be retrained or promoted to compensate for an unresolved quarantine backlog or unverified schema change. [S44] If lineage links a performance drop to a surge in MISSING or ESTIMATED inputs, the model owner must freeze model changes and escalate the capture defect to the domain owner. [S11]

4.21 ROOT-CAUSE LINEAGE STOPS TEAMS FROM RETRAINING MODELS TO MASK BROKEN CAPTURE

When quality gates and provenance graphs work together, diagnosing an anomaly becomes a fast lookup rather than a week-long forensic project. [S15] Following the CRISP-DM manufacturing lifecycle, evaluation and deployment remain coupled to data understanding and preparation. [S43] Before opening a modelling notebook to investigate a degraded forecast or erratic KPI, the spine runs three checks along the lineage graph:

- **Step 1: Inspect upstream quality-state composition**: Did MISSING, ESTIMATED or quarantined shares in feeding tables exceed the contracted threshold during the window? [S04]

- **Step 2: Check contract and rule-version transitions**: Did a schema version, KPI formula version (S38), master-data mapping (S05) or threshold change along wasDerivedFrom? [S09] [S38]
- **Step 3: Audit manual correction and waiver logs**: Did a temporary waiver or post-shift CORRECTED surge alter input or label distributions? [S02]

Only when all three checks confirm stable capture, semantics and gates does the team treat the anomaly as process or concept drift under Part III. [S15] [S43]

> **GFT NOTE**
>
> **Convert manual spreadsheet checks into governed gates before buying software**
>
> In GFT's Cameroon textile programme, the initial performance-management system starts with shift sheets and spreadsheet logs before edge brokers or vision cameras are deployed. A factory need not wait for a cloud platform to apply this chapter's controls. [S36] On day one, the steward can add the five quality states (OBSERVED, MISSING, ESTIMATED, CORRECTED, NOT_APPLICABLE) as a mandatory workbook column, lock picklists to a master tab, move failed rows to a Quarantine_Log sheet rather than deleting them, and review the 12-row quality register weekly. [S02] [S41] When the plant migrates to a database and MLOps pipeline (wp18), those exact rules, states and provenance obligations transfer directly into SQL constraints and pipeline assertions. [S04] [S15]

By making data fitness an explicit, versioned and provenance-backed property of every record, the factory completes Part I of the industrial data spine: decision rights are named, vocabulary is contracted, and quality gates guard curated storage. [S01] [S02] Part II turns to the minimum durable architecture required to store, contextualise, protect and recover those records under factory constraints.

Clean data are not found by analysts at the end of the pipeline; they are kept clean by visible gates and traceable lineage at the point of work.

CHAPTER SOURCE NOTES

S01 International Organization for Standardization. 2022.

S18 Advances in Neural Information Processing Systems. 2015.

S04 ISO/IEC. 2024.

S09 World Wide Web Consortium. 2013-04-30.

R06 Cameroon Trade Hub. Publication date unrecorded; inherited source-text check.

R07 Ecofin Agency. Publication date unrecorded; inherited source-text check.

R12 Business in Cameroon. Publication date unrecorded; inherited source-text check.

S36 World Bank. 2023.

S02 International Organization for Standardization. 2022.

S03 ISO/IEC. 2008.

S15 Google Cloud. 2024-08-28.

S0401 World Wide Web Consortium. 2016-12-15.

S38 International Organization for Standardization. 2014.

S05 International Electrotechnical Commission. 2013.

S08 GS1 AISBL. 2022.

S39 National Institute of Standards and Technology. 2025.

S06 OPC Foundation. 2025-10-22.

S19 Association for Computing Machinery, CHI 2021. 2021.

S0402 World Wide Web Consortium. 2017-01-31.

S41 GHG Protocol, World Resources Institute and World Business Council for Sustainable Development. 2013.

S11 National Institute of Standards and Technology. 2023-01-26.

S0403 World Wide Web Consortium. 2013-04-30.

S10 National Institute of Standards and Technology. 2016-02-23.

R29 MDPI; industry inspection literature. Publication date unrecorded; inherited source-text check.

S43 Frontiers in Artificial Intelligence. 2021-06-14.

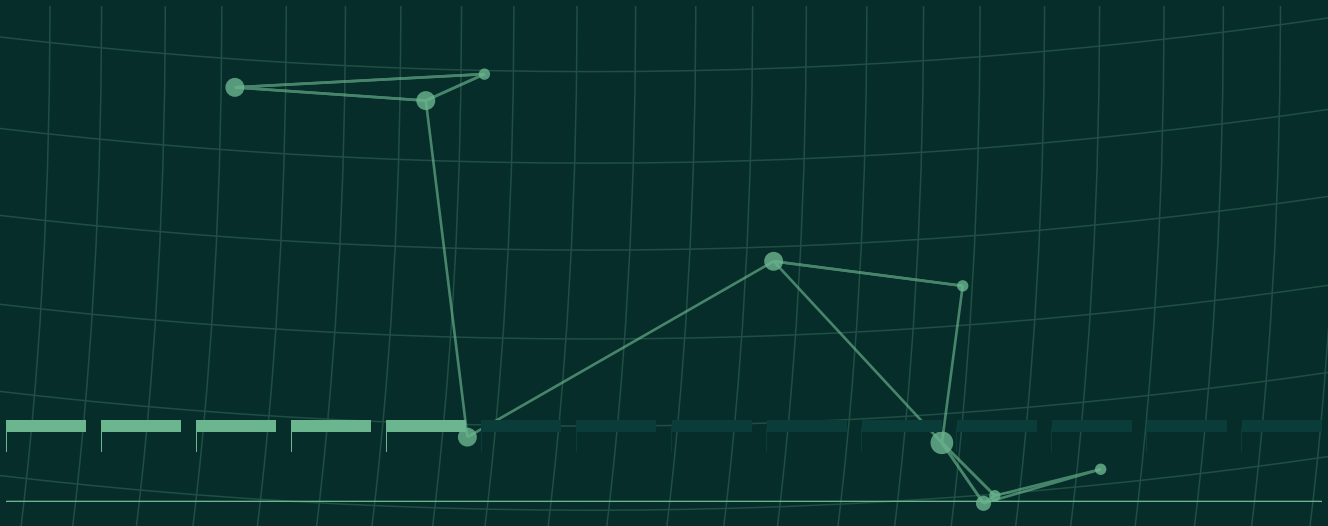
S44 The Journal of Technology Transfer, Springer Nature. 2024-07-27.

05

PART II: BUILD THE MINIMUM DURABLE SPINE / MINIMUM ARCHITECTURE

The smallest durable spine separates capture, context and consumption

A durable architecture preserves source evidence, adds governed context once and lets each decision consume a controlled product.



CHAPTER 05
DECISION TEST

A durable architecture preserves source evidence, adds governed context once and lets each decision consume a controlled product.

The smallest durable industrial data spine is a three-stage separation of responsibilities, not a large software suite. It separates capture close to physical work, context assembly in a canonical operational core, and consumption through read-only, purpose-specific data products. [S05] [S10] In Cameroon, where the inherited sector baseline prepared on 31 January 2024 records about 95 percent of cotton fibre exported raw and about 5 percent processed locally against a 50 percent target by 2030 (original dates unrecorded and not refreshed in 2026), domestic mills cannot afford brittle data estates. [R06] [R07] Reported capacity utilisation at domestic textile plants in July 2023 ranged from 10.83 percent at CICAM 2 to 33.7 percent at Garoua and 48.9 percent at CICAM 1, with underutilisation costing about CFAF 2 billion in 2020. [R12] Connecting shop-floor logs, machine counters, quality sheets and order schedules directly to spreadsheets or dashboards without separating raw evidence from business transformation means every formula change alters history and every new consumer multiplies integration effort. [S18] [S39] Separating capture, context and consumption lets a lean team preserve immutable shop-floor evidence, attach governed meaning once, and serve management reporting, traceability and machine-learning pipelines from one auditable foundation. [S08] [S09] [S15]

5.1

Capture remains close to the work

Capture must occur at the physical boundary where materials move, machines cycle and operators inspect garments, and it must continue working when external networks fail. [S22] [S33] An event loses diagnostic fidelity when an operator must remember a stoppage until shift end or transcribe a chalkboard tally into a terminal hours later. [S19] Keeping capture adjacent to work means placing manual entry forms at the sewing line or inspection table, reading programmable logic controllers and utility meters at the cell boundary, and running edge vision inference next to the fabric frame rather than across an unstable wide-area link. [R29] [S05] [S22] Because operational technology environments have strict performance, reliability and safety constraints, capture components must observe production passively or through defined interfaces without delaying control loops or introducing unsafe dependencies. [S05] [S22]

5.2 ADJACENT RECORDING PREVENTS TRANSLATION DRIFT AT THE BOUNDARY

Recording an event where it happens requires local resilience against network interruptions. [S33] Where telecommunications or switch faults occur intermittently, a capture node that depends on continuous cloud or central-server availability drops production evidence when operational stress is highest. [S22] [S33] Product-specific edge architectures illustrate the required mechanism: when upstream connectivity is lost, a local parent device buffers device-to-cloud messages on durable storage until the connection returns. [S33] In GFT's minimum spine, that store-and-forward principle applies across all three capture channels: a line-side tablet caches signed shift entries locally, an edge gateway writes polled machine registers to a local ring buffer, and an optical inspection unit stores defect events and image hashes on local solid-state storage before upstream synchronisation. [R29] [S08] [S33]

5.3 PROVENANCE TAGS DISTINGUISH HUMAN, SENSOR AND MODEL ORIGINS

Adjacency alone does not make captured records trustworthy unless the spine records how each value entered the system. [S09] [S0501] Under the W3C PROV Data Model, provenance is information about entities, activities and agents involved in producing a piece of data, used to assess its quality, reliability and trustworthiness. [S0501] If a table stores a fabric defect count or stoppage duration without recording the generating agent and activity, downstream analysts and training pipelines cannot tell whether the figure came from a calibrated optical sensor, a manual tally sheet, a default fallback value or a model's inference. [S04] [S09] [S19]

Every captured record therefore carries an explicit capture-method and source-agent header alongside its payload. [S08] [S09] Manual entries record the workstation role, entry mode and submission timestamp while respecting worker data minimisation by avoiding unnecessary personal identifiers. [S02] [S29] Sensor and meter readings record the device identifier, interface protocol and polling interval. [S05] [S06] Edge model outputs, such as automated fabric defect classifications where literature reports computer vision exceeding 90 percent accuracy across more than 40 defect types at line-scan speeds up to 60 metres a minute compared with 60 to 70 percent for manual inspection, must record the model identifier, version, confidence score and threshold applied. [R29] [S15] Keeping manual, sensor and model-generated values structurally distinct at capture prevents a later pipeline from training a new model on the unverified predictions of an older model. [S04] [S15] [S18]

> **GFT NOTE**
>
> In GFT's Cameroon textile programme, early logging begins with structured manual sheets and line-side tablets before full machine telemetry is installed. Manual origin is never a reason to skip provenance metadata. Every manual record should carry its capture channel, workstation identifier, event timestamp and ingestion timestamp from day one so that when automated meters or cameras are added later, the spine distinguishes instrumentation changes from real shop-floor change. [S08] [S09] [S38]

5.4

Raw evidence and curated records serve different purposes

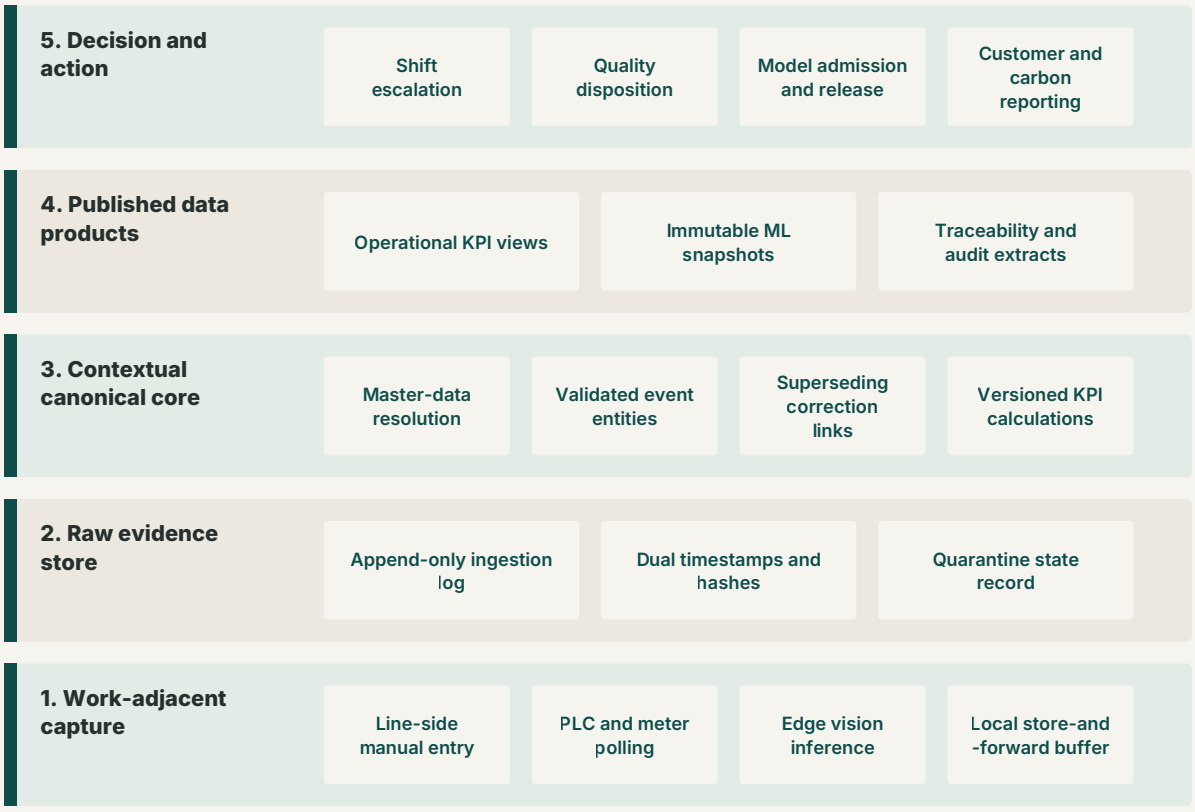
Once an event leaves the capture boundary, the spine must separate the duty to remember what was observed from the duty to present clean, contextualised data for daily decisions. [S09] [S15] In many small factories, one database or shared workbook serves both roles: scripts or operators write incoming numbers into a table, and supervisors edit that same table in place when they spot a typo, a miscalibrated meter or a split order. [S18] [S19] In-place mutation destroys the audit trail. When a downstream key performance indicator (KPI) shifts unexpectedly or a forecasting model degrades, the team cannot reconstruct what the capture layer reported on the day of production or test whether a cleaning rule introduced bias. [S04] [S15] [S19]

5.5 IMMUTABLE RAW RETENTION PROTECTS DIAGNOSTIC RECONSTRUCTION

The minimum durable spine enforces a strict boundary between the immutable raw evidence log and the curated contextual store, as shown in **Exhibit 20**, p. 84. [S05] [S09] [S15] The raw evidence layer accepts every receivable payload from capture, assigns an ingestion timestamp and cryptographic content hash, and writes the record to append-only storage. [S08] [S09] [S26] No user, steward or transformation script may update or delete rows in the raw layer during normal operations. [S25] Even when a record fails schema validation or violates a physical range check, the raw payload is preserved so engineers can inspect the exact malformed message, diagnose whether a controller register shifted or a manual form broke, and replay the raw stream once the parser or rule is fixed. [S01] [S04] [S15]

EXHIBIT 20 /

Three separations keep the spine understandable



DERIVED [S05] [S09] [S10] [S15] [S0501] Derived minimum architecture separating capture, immutable raw evidence, contextual canonical records, published data products and operational decisions, based on interface boundaries in S05, provenance modelling in S09 and S0501, smart manufacturing standards in S10 and pipeline lineage in S15.

As **Exhibit 20**, p. 84 illustrates, the five functional layers embody three architectural separations: capture is isolated from storage by local buffering and provenance headers; raw evidence is isolated from curated context by append-only immutability and quality gates; and curated context is isolated from consumption by read-only published products. [S05] [S09] [S10] [S15] What **Exhibit 20**, p. 84 cannot establish on its own is the physical deployment topology, such as whether layers two, three and four

run inside a single relational database engine with schema-level access controls or across separate server nodes; that physical sizing belongs to the infrastructure design in WP18 rather than the governance architecture. [S05] [S25]

5.6 SUPERSEDING RECORDS REPLACE DESTRUCTIVE IN-PLACE EDITS

Curated records support rapid, reliable consumption without forcing every dashboard query or feature pipeline to re-parse raw payloads and re-apply historical correction logs from scratch. [S04] [S15] [S0502] In the contextual layer, raw events are validated against their data contracts, linked to the master-data version active at the event time, and tagged with an explicit quality state such as observed, estimated, missing or manually corrected. [S03] [S04] [S08]

When a domain steward discovers that an operator entered the wrong lot code or that a power meter reported a spike during maintenance testing, the steward never overwrites the raw record and never silently edits the curated row in place. [S02] [S09] Instead, the spine records a superseding correction event that references the original event identifier, states the corrected value and reason code, identifies the approving steward role, and links the new curated state to both the original raw entity and the correction activity using PROV-O derivation relationships. [S02] [S09] [S0501] Curated views expose the current effective state for routine reporting while retaining the full derivation chain so any historical report or model training snapshot can be reproduced as of the timestamp it was originally generated. [S09] [S15] [S0503]

5.7

A canonical core replaces repeated point translation

Connecting each source directly to each consuming application creates an unmanageable web of bilateral translations. [S10] [S18] When a sewing plant has line-side logs, an order table, a quality sheet, a maintenance log and an energy meter, and feeds an overall equipment effectiveness (OEE) tracker, a buyer portal, a carbon calculator and a demand forecaster through point-to-point scripts, each script embeds private assumptions about shift boundaries, units, scrap codes and order identifiers. [S05] [S38] [S39] NIST's smart manufacturing analysis notes that standards are fundamental to delivering the right information at the right time, enabling action, and reducing adoption risk. [S10] Manufacturers in a 2025 NIST semiconductor workshop reported that data processing for a single application can reach 70 percent of total application cost when representations are not shared. [S39]

5.8 BILATERAL ADAPTERS MULTIPLY INTEGRATION MAINTENANCE

Point-to-point translation fails small teams because maintenance effort scales with the product of sources and consumers rather than their sum. [S10] [S18] When an upstream machine tag or spreadsheet column changes, the custodian must update multiple extraction scripts, each written at a different time with incompatible filter rules. [S18] [S19] A canonical core replaces bilateral adapters with a hub-and-spoke contract grounded in IEC 62264-1 (ISA-95) enterprise-control models and GS1 EPCIS visibility event structures. [S05] [S08] [S10] Each capture source maps once into the canonical model; each downstream product reads from that core. [S05] [S10]

METHOD

To quantify the maintenance difference between bilateral translation and a canonical core, define:

- S as the number of distinct capture sources (manual shift logs, machine counters, quality stations, maintenance logs and utility meters). [S05]
- C as the number of distinct consumption views or applications (shift OEE dashboards, quality logs, customer provenance extracts, carbon calculators and forecasting datasets). [S08] [S38]
- Under bilateral integration where each consumer maintains its own extraction and translation from each source, the number of interface mappings M_{direct} is:

$$M_{\text{direct}} = S \times C$$

- Under a canonical core where each source maps once to the governed model and each consumer reads through one product contract, the number of interface mappings $M_{\text{canonical}}$ is:

$$M_{\text{canonical}} = S + C$$

- Check boundary: this count measures distinct schema and semantic contracts tested when definitions change; it does not claim identical engineering hours per mapping. [S10] [S39]

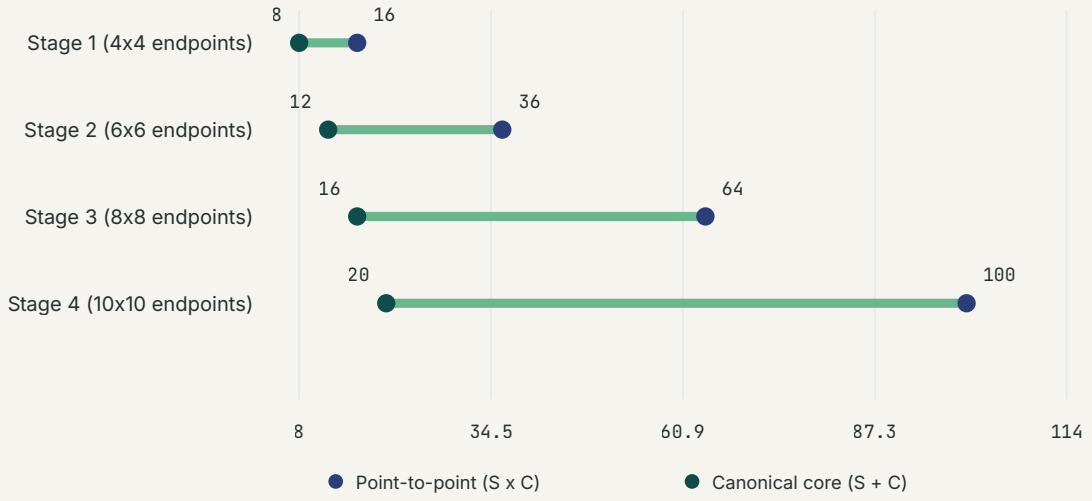
WORKED EXAMPLE

Consider an illustrative progression across four stages as a small textile plant expands from manual tracking to multi-domain analytics. [S05] [S10] In Stage 1 (an illustrative baseline of 4 sources and 4 views), bilateral integration requires $4 \times 4 = 16$ illustrative mappings, whereas a canonical core requires $4 + 4 = 8$ illustrative mappings, a reduction of 8 mappings (50 percent). In Stage 2 (an illustrative 6 sources and 6 views), bilateral integration requires $6 \times 6 = 36$ illustrative mappings against $6 + 6 = 12$ illustrative canonical mappings, saving 24 mappings (66.7 percent). In Stage 3 (an illustrative 8 sources and 8 views), bilateral integration reaches $8 \times 8 = 64$ illustrative mappings against $8 + 8 = 16$ illustrative canonical mappings, saving 48 mappings (75 percent). By Stage 4 (an illustrative 10 sources and 10 views), bilateral scripts require $10 \times 10 = 100$ illustrative mappings compared with $10 + 10 = 20$ illustrative canonical mappings, an illustrative reduction of 80 mappings (80 percent). If each mapping contract requires an illustrative 4 hours of steward and custodian verification per year during schema reviews, the Stage 3 plant spends $64 \times 4 = 256$ illustrative hours a year maintaining point-to-point scripts versus $16 \times 4 = 64$ illustrative hours on the canonical spine, saving 192 illustrative engineering hours a year. [S02] [S10] [S39]

The divergence between bilateral translation and a canonical core across those four illustrative stages is shown in **Exhibit 21**, p. 87. [S05] [S10] [S39]

EXHIBIT 21 / NUMBER OF INTERFACE MAPPING CONTRACTS (ILLUSTRATIVE)

A canonical core cuts interface mappings as connected endpoints grow



ILLUSTRATIVE [S05] [S10] [S39] Illustrative arithmetic comparing S + C canonical mappings with S * C bilateral point-to-point mappings across four symmetrical factory integration stages (4, 6, 8 and 10 sources and consumers). Structural rationale derived from S05, S10 and S39; counts are mathematical identities on assumed endpoint counts, not empirical measurements from a specific plant.

As **Exhibit 21**, p. 87 demonstrates, the payoff of a canonical core is modest with one source and one report, but compounds rapidly once shop-floor records serve operations, quality, maintenance, sustainability and machine learning. [S05] [S10] [S39] However, **Exhibit 21**, p. 87 cannot prove that an expansive enterprise model is desirable; designing a schema for hundreds of theoretical entities before proving shop-floor value stalls implementation. [S16] [S35] [S40]

5.9 SELECTIVE SCOPE GOVERNS HIGH-REUSE OPERATIONAL ENTITIES FIRST

To avoid premature complexity, a small team must restrict the canonical core to cross-domain records carrying repeated decision value. [S05] [S08] [S38] Rather than mapping every internal controller register into the central model, the spine standardises only records crossing domain boundaries: site and line hierarchy, order and lot identity, shift calendar, output event, state transition (downtime) event, quality inspection and defect event, maintenance event, and utility consumption interval. [S05] [S08] [S38] High-frequency internal diagnostic signals that matter only to one machine's local maintenance remain at the edge or in the raw log until a validated use case requires promotion. [S06] [S16] [S33]

5.10

Transport follows the operating requirement

Moving records from capture nodes to the raw store and canonical core requires transport suited to factory networks, but transport selection must never be confused with governance. [S06] [S07] Teams often equate installing a message broker with building a data architecture, assuming that once payloads flow across a publish-subscribe bus, the factory has a usable spine. [S07] [S10] That confusion leaves consumers parsing undocumented strings, guessing units, and misinterpreting retried messages as duplicate output. [S04] [S07] [S08]

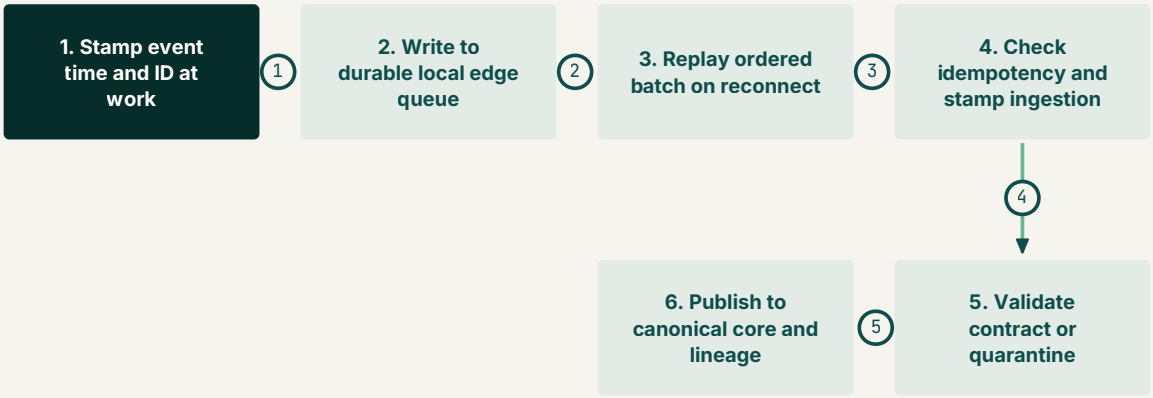
5.11 PUBLISH-SUBSCRIBE BROKERS DECOUPLE PRODUCERS FROM CONSUMERS

Where multiple services need timely access to shop-floor events without polling controllers repeatedly, a publish-subscribe transport protocol such as MQTT Version 5.0 is an effective choice. [S07] A line-side gateway or vision unit publishes an event once to a broker topic, and the raw writer, operational monitor and alerting service subscribe without placing extra connection load on the shop-floor controller. [S07] [S22] Where richer machine-to-machine contextual interoperability and browseable object hierarchies are needed at the device boundary, OPC Unified Architecture (OPC UA) provides an information-modelling and secure communication framework. [S06] For low-frequency shift summaries or batch test results, transactional hypertext transfer or scheduled file ingestion with checksums is often simpler for a small team to maintain than a streaming broker. [S16] [S20]

When network links between a production hall and central storage drop, the transport contract must coordinate with local edge buffering, as traced in **Exhibit 22**, p. 88. [S07] [S08] [S09] [S33]

EXHIBIT 22 /

Local buffering preserves capture during a disconnected interval



1. Link down or normal write
2. Connection restored
3. At-least-once delivery
4. Unique raw event logged
5. Contract passed

DERIVED [S07] [S08] [S09] [S33] Derived store-and-forward governance flow combining local edge buffering during disconnection from S33 with transport delivery from S07, event identity and dual timestamps from S08, and provenance state tracking from S09.

As **Exhibit 22**, p. 88 shows, surviving a network outage requires six steps across capture, transport and ingestion. [S07] [S08] [S09] [S33] First, the capture node assigns a globally unique event identifier and records observation time before transmission. [S08] Second, the event is written to a persistent local queue on the edge parent device so a link failure does not drop the record. [S33] Third, when connectivity

returns, the transport client replays buffered messages in order. [S07] [S33] Fourth, because at-least-once delivery can send a message more than once during reconnection, the receiver checks the event identifier for idempotency and stamps an ingestion timestamp before appending the unique payload to the raw store. [S07] [S08] Fifth, the schema and quality gate validates the payload and routes failures to quarantine. [S04] Sixth, validated records advance to the canonical core with provenance showing delayed replay rather than live streaming. [S09] Note that **Exhibit 22**, p. 88 covers data continuity during communication loss; it does not provide backup power or machine safety interlocking. [S22] [S33]

5.12 DELIVERY GUARANTEES DO NOT VALIDATE PAYLOAD SEMANTICS

OASIS defines MQTT Version 5.0 strictly as a client-server publish-subscribe messaging transport protocol. [S07] The protocol transports bytes between publishers and subscribers according to configured delivery levels, session expiry intervals and topic filters, without knowing whether a payload labelled `line1/output` represents finished first-quality garments, total pieces including rework, or cumulative counter ticks since a controller reboot. [S07] [S38]

Treating broker throughput or topic hierarchy as proof of semantic governance is a recurring error. [S06] [S07] [S19] Every transport topic or ingestion endpoint must be bound to a versioned data contract specifying the schema, unit, time grain, event semantics, quality state rules and accountable owner. [S02] [S04] [S08] If a publisher alters its payload structure or unit without an approved contract version update, the validator quarantines the message even though the broker delivered the packet without error. [S04] [S07]

> **WARNING**
>
> Unbounded edge queues and unthrottled reconnect storms can disrupt both edge capture and central ingestion when a prolonged outage ends. [S22] [S33] Every store-and-forward node must enforce a documented local storage ceiling, an explicit overflow policy (preserving summary totals and fault events before high-frequency diagnostic traces), and rate-limited replay so reconnection does not saturate bandwidth or starve live monitoring. [S22] [S33]

5.13

Consumption receives controlled products, not raw access

The third separation stands between the canonical core and the people or systems that consume data. [S05] [S0502] Granting analysts, dashboard authors, external auditors or scripts direct credentials to query raw ingestion tables or mutate canonical tables causes two failures: ad hoc queries diverge from the governed KPI dictionary, resurrecting conflicting figures for output, downtime and defects; and unthrottled queries or accidental writes contend with live ingestion and violate least privilege. [S18] [S22] [S25] [S38]

5.14 ROLE-SCOPED VIEWS ISOLATE SOURCE EVIDENCE FROM CONSUMER MUTATION

In GFT's minimum spine, consumers never receive write access to raw evidence or canonical operational tables, and ordinary consumers do not query raw payloads directly. [S09] [S25] Instead, the spine publishes controlled, read-only data products tailored to distinct consumption patterns, following W3C Data on the Web Best Practices for consistent dataset publication and the Data Quality Vocabulary principle of exposing fitness-for-purpose metadata alongside published datasets. [S0502] [S0503] Access to each published product is governed by least privilege under NIST SP 800-53 Revision 5, granting each human role or automated service account only the permissions required to accomplish its assigned task. [S25]

The five core consumption product classes required by a small factory and their governed contracts are detailed in **Exhibit 23**, p. 90. [S05] [S08] [S15] [S25] [S38] [S0502]

EXHIBIT 23 /

Controlled data products isolate source evidence from consumer views

Product class	Primary consumer	Grain and cadence	Access control	Lineage and state rule
1. Shift operational view	Line supervisors and plant manager	Hourly and shift aggregates	Read-only by operational role	Uses active ISO 22400-2 KPI version; flags missing or buffered intervals
2. Quality and steward queue	Domain stewards and quality leads	Event-level exceptions and quarantine	Read queue; append correction event	Links every superseding correction to raw event ID and steward role
3. ML training snapshot	Model builder and validation pipeline	Frozen point-in-time dataset extract	Immutable read by pipeline identity	Records manifest hash, schema version, quality states and split window
4. Traceability and carbon view	Sustainability and customer compliance	Lot, order and resource balances	Read-only scoped to external reporting	Joins EPCIS lot events with utility intervals; excludes worker data
5. Raw diagnostic archive	Data custodian and incident responder	Unmodified ingestion payloads	Restricted privileged audit access	Append-only storage with cryptographic hashes and offline backup copy

DERIVED [S05] [S08] [S15] [S25] [S38] [S0502] Derived product contract table mapping five factory consumption classes to their canonical inputs, refresh grain, access permissions and lineage guarantees, synthesised from S05, S08, S15, S25, S38 and S0502.

As **Exhibit 23**, p. 90 specifies, each consumer interacts with a bounded contract rather than an undifferentiated data store. [S15] [S25] [S0502] Line supervisors reviewing shift performance consume pre-aggregated views that implement the exact formula, time elements and units registered in the factory's ISO 22400-2 KPI dictionary, with visual badges whenever an hourly bucket contains buffered, missing or estimated

records. [S04] [S38] [S0503] Machine-learning pipelines never train against live mutable tables; they consume frozen, content-hashed snapshots whose manifest records the query, canonical schema version and quality-rule version used to assemble the extract. [S04] [S09] [S15]

5.15 SHARED IDENTIFIERS ALIGN OPERATIONAL, ANALYTICAL AND COMPLIANCE VIEWS

Publishing separate data products for operations, quality, sustainability and machine learning does not duplicate data pipelines. [S10] [S39] Because every published product in **Exhibit 23**, p. 90 is derived from the same canonical core and uses the same governed identifiers for site, line, asset, production order, material lot and shift, figures reconcile across departments by construction. [S05] [S08] [S38]

When the plant manager examines first-pass yield in the daily cadence (WP02), the quality engineer evaluates an optical model's false-rejection rate (WP03), the planner prepares an order forecast (WP04), the compliance lead traces a cotton bale to a garment lot (WP05), and the sustainability lead allocates kilowatt-hours and water per kilogram of fabric (WP16), all five decisions reference identical lot identifiers, order identifiers and validated output quantities. [S05] [S08] [S38] Fit-for-purpose transformations happen in the published product layer, such as converting piece counts into fabric mass for carbon accounting or lagging order histories into features for forecasting, while the underlying canonical facts remain single-sourced and versioned. [S09] [S15] [S38]

5.16

WP14 governs records while WP18 designs the topology

A disciplined library separates information governance from physical network and hardware engineering so neither concern is diluted. Within GFT's research architecture, WP14 owns the governance of the data spine: decision rights, operational vocabulary, data contracts, quality gates, provenance lineage, capture-context-consumption separation, recovery objectives and staged MLOps controls. [S01] [S02] [S09] [S15] Companion paper WP18 owns the physical and logical reference topology: network segmentation, switch and firewall placement, edge compute sizing, container orchestration, storage engine selection and link engineering. [S06] [S22] [S28]

5.17 GOVERNANCE OBLIGATIONS CONSTRAIN PHYSICAL AND NETWORK SELECTION

Separating WP14 from WP18 does not make them independent; governance sets the non-negotiable invariants that every topology choice in WP18 must satisfy. [S05] [S10] [S22] In small and medium-sized manufacturers, where the 2026 OECD Western Balkans survey illustrates that integrated tools such as enterprise resource planning (ERP), cloud computing and customer relationship management (CRM) are used by only 10 to 20 percent of surveyed firms (an external SME benchmark rather than a Cameroon prevalence measure), infrastructure teams face pressure to improvise quick hardware or cloud fixes. [S34] Without governance constraints from WP14, an infrastructure deployment can break record integrity, for example by configuring an

edge gateway to overwrite local ring buffers silently without logging dropped intervals, or by deploying a database that lacks role-based separation between raw and curated tables. [S09] [S25] [S33]

Every topology and component design under WP18 must therefore pass six record-governance obligations defined here: 1. **Boundary safety and non-interference:** Network taps, polling gateways and edge inference nodes must respect NIST SP 800-82 Revision 3 operational technology performance, reliability and safety constraints, ensuring a data-spine fault cannot stall or alter machine control. [S22] 2. **Disconnection durability:** Any physical link between a capture cell and central storage must include durable local store-and-forward buffering sized for expected outages, preserving event identifiers and observation timestamps across link drops. [S08] [S33] 3. **Immutable raw separation:** The storage engine and access-control configuration must enforce append-only retention for raw payloads, preventing in-place updates or unlogged deletions by normal users or transformation jobs. [S09] [S25] 4. **End-to-end provenance capture:** Pipeline runners, brokers and schemas must record source agent, capture method, dual timestamps, schema version and transformation metadata from ingestion through serving. [S08] [S09] [S15] 5. **Least-privilege product isolation:** Authentication and authorisation infrastructure must support separate identities for operators, stewards, custodians, analysts, ingestion services and pipelines, restricting consumers to read-only published products. [S25] 6. **Tested offline recovery and portable export:** Backup and storage components must maintain offline, encrypted backup copies and support complete export of raw records, curated tables, schemas, KPI definitions and model artefacts in open formats that can be restored and verified independently. [S10] [S26] [S32]

5.18 INTERFACE VERIFICATION TESTS TOPOLOGY AGAINST RECORD INVARIANTS

To keep the boundary between WP14 and WP18 operational, the custodian and domain stewards verify infrastructure changes against record invariants before any new network segment, broker or storage node enters production. [S02] [S27] When WP18 upgrades an edge gateway or migrates storage from a local server to a hybrid cloud instance, the acceptance test is not merely that packets cross the network or that a server reports healthy processor utilisation. [S17] [S27] The acceptance test is whether a test batch of events injected at capture arrives in the raw store with intact hashes, passes or quarantines deterministically under the active contract, resolves against the canonical master-data version, reproduces historical KPI values in the operational view, and restores cleanly from an offline backup into an isolated test environment. [S04] [S09] [S26] [S38]

5.19

Buy commodity capability but retain semantic control

A small factory team cannot build and maintain custom database engines, message brokers or experiment-tracking platforms, nor should it surrender the meaning of its operations to an external vendor's black-box schema. [S16] [S18] [S32] Resolving the

build-versus-buy decision requires disaggregating the spine by architectural layer rather than treating build and buy as an ideological choice across the factory. [S16] [S20]

5.20 COMMODITY INFRASTRUCTURE REDUCES OPERATIONAL TOIL FOR SMALL TEAMS

Transport brokers, relational or analytical storage engines, object storage, version control, backup encryption utilities and basic dashboard rendering are mature commodity capabilities. [S07] [S16] [S20] Writing custom software for message queuing, time-series compression or role-based authentication diverts scarce factory engineering hours into low-differentiation maintenance, compounding the hidden technical debt described in machine-learning systems literature. [S18] [S20] In line with the Azure MLOps Maturity Model's guidance to develop capability gradually rather than facing full implementation complexity up front, small teams should buy or adopt proven, standard-compliant commodity software for transport, storage, backup and basic registry mechanics. [S16] [S42] Equally, World Bank case study evidence on SME digital adoption in Viet Nam underscores that successful platform uptake depends on assigning a named internal point person who champions adoption within the firm; buying commodity software never removes the need for internal operational leadership. [S37]

Conversely, a factory must never outsource its operational semantics. [S05] [S38] The identifier taxonomy (how sites, lines, machines, cotton lots, garment styles and shifts are coded), canonical event schemas, ISO 22400-2 KPI formulas, quality-gate thresholds and provenance lineage graph are the core of the factory's management system. [S04] [S08] [S09] [S38] If a vendor's turnkey dashboard or vision appliance computes OEE or defect rates using proprietary internal definitions and refuses to expose underlying event records, the factory cannot reconcile those figures with its order ledger, cannot audit changes after firmware updates, and loses its historical baseline if it switches suppliers. [S10] [S32] [S38]

This layer-by-layer allocation of build, configure and buy choices is mapped in **Exhibit 24**, p. 94. [S20] [S32] [S37]

EXHIBIT 24 / ILLUSTRATIVE SCORE (0 TO 10 SCALE)

Build-against-buy changes by layer, not by ideology



ILLUSTRATIVE [S20] [S32] [S37] Illustrative architectural positioning of seven spine capabilities across commodity availability (x-axis, 0 = factory-specific to 10 = standardised commodity) and need for internal semantic control (y-axis, 0 = generic utility to 10 = core business meaning). Rationale derived from MLOps architecture in S20, connected-product data access considerations in S32 and internal champion ownership in S37.

As **Exhibit 24**, p. 94 shows, capabilities in the lower-right quadrant (message brokers, storage, and encrypted backup utilities) combine high commodity availability with low need for bespoke business semantics, making them prime candidates to buy or adopt off the shelf. [S07] [S20] [S26] Capabilities in the upper-left quadrant (identifier schemes, canonical event contracts and KPI formulas) carry the factory's core operational meaning and must be authored and owned internally by the accountable domain owners and stewards. [S02] [S08] [S38] Capabilities in the upper-right quadrant (quality rule execution, lineage graphs, model registries and BI presentation) sit in the middle: the factory should adopt lightweight tools that support open standards, then configure them with its own governed definitions and metadata. [S09] [S15] [S20] Because **Exhibit 24**, p. 94 is an illustrative structural framework rather than a cost survey, each site must evaluate vendor pricing, local support and staff skills through GFT's build-or-buy tool before signing a contract. [S37] [S44]

5.21 CONTRACTUAL DATA RIGHTS AND OPEN SCHEMAS PREVENT SEMANTIC LOCK-IN

Retaining semantic control when purchasing machinery, edge sensors or software platforms requires explicit procurement clauses before commissioning. [S10] [S32] Where the European Union Data Act (Regulation (EU) 2023/2854) applies, for instance when a connected product or service falls within the territorial and material scope of Union rules, Recital 5 states that users in the Union can access, in a timely manner, the data generated by the use of that product or service and share those data with third parties of their choice. [S32] For a factory in Cameroon, the EU Data Act is not domestic law and must not be cited as an automatic statutory right; instead, it serves as a practical procurement benchmark for the access clauses the factory should write into its commercial equipment and software contracts. [S32]

Specifically, every procurement specification for machinery, optical inspection units, meters or cloud analytics in the spine should require four contractual deliverables:

- Unfettered, no-fee local access to raw machine and sensor events at the documented observation grain via standard protocols such as OPC UA or MQTT, with complete payload documentation and units. [S06] [S07] [S32]
- Explicit right to ingest, store, transform and share those operational records with any third-party maintenance provider, auditor or analytics tool chosen by the factory. [S32]
- Bulk export of all historical raw events, curated tables, schemas, user logs and model artefacts in open formats (such as documented CSV, Parquet, JSON-LD or SQL dumps) on demand and at contract end. [S09] [S10] [S0502]
- Advance notification of any firmware or software update that alters payload schemas, units or embedded model versions, so the data steward can run contract compatibility checks before production release. [S04] [S15]

Combining purchased commodity infrastructure with internally owned identifiers, contracts and export rights gives a small factory team a durable data spine that fits its staffing limits today while preserving architectural freedom as production scales. [S10] [S16] [S37]

Preserve the raw event, govern the context once, and let every decision read a product it cannot corrupt.

CHAPTER SOURCE NOTES

S05 International Electrotechnical Commission. 2013.

S10 National Institute of Standards and Technology. 2016-02-23.

R06 Cameroon Trade Hub. Publication date unrecorded; inherited source-text check.

R07 Ecofin Agency. Publication date unrecorded; inherited source-text check.

R12 Business in Cameroon. Publication date unrecorded; inherited source-text check.

S18 Advances in Neural Information Processing Systems. 2015.

S39 National Institute of Standards and Technology. 2025.

S08 GS1 AISBL. 2022.

S09 World Wide Web Consortium. 2013-04-30.

S15 Google Cloud. 2024-08-28.

S22 National Institute of Standards and Technology. 2023.

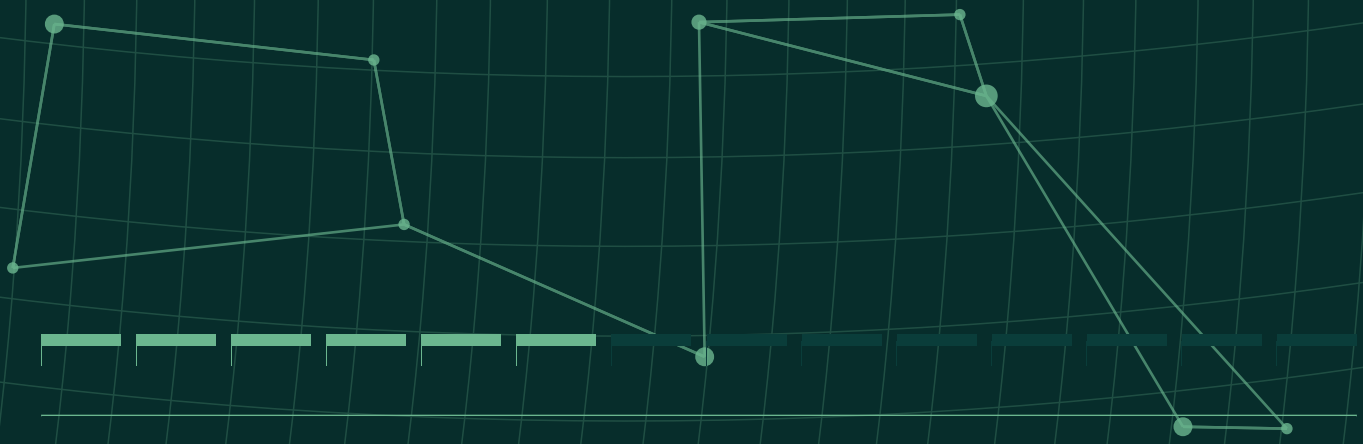
S33 Microsoft. 2024-06-10.
S19 Association for Computing Machinery, CHI 2021. 2021.
R29 MDPI; industry inspection literature. Publication date unrecorded; inherited source-text check.
S0501 World Wide Web Consortium. 2013-04-30.
S04 ISO/IEC. 2024.
S02 International Organization for Standardization. 2022.
S29 European Parliament and Council of the European Union. 2016-04-27.
S06 OPC Foundation. 2025-10-22.
S38 International Organization for Standardization. 2014.
S26 CISA, MS-ISAC, NSA and FBI. 2023.
S25 National Institute of Standards and Technology. 2020-12-10.
S01 International Organization for Standardization. 2022.
S0502 World Wide Web Consortium. 2017-01-31.
S03 ISO/IEC. 2008.
S0503 World Wide Web Consortium. 2016-12-15.
S16 Microsoft Learn, Azure Architecture Center. 2026-06-03.
S35 OECD. 2022-11-21.
S40 International Labour Organization and UNESCO. 2020.
S07 OASIS Open. 2019-03-07.
S20 arXiv, author version of an IEEE Access article. 2022-05-14.
S28 International Society of Automation. 2026.
S34 OECD. 2026-06-26.
S32 European Parliament and Council of the European Union. 2023-12-13.
S27 National Institute of Standards and Technology. 2025.
S17 Microsoft Learn, Azure Machine Learning. 2026-01-27.
S42 SCITEPRESS, IN4PL 2022. 2022.
S37 World Bank. 2025-05-20.
S44 The Journal of Technology Transfer, Springer Nature. 2024-07-27.

06

PART II: BUILD THE MINIMUM DURABLE SPINE / INDUSTRIAL EVIDENCE

Time, identity and state turn sensor readings into industrial evidence

A value becomes useful evidence only when the factory can establish what it describes, when it happened and which state or transformation it represents.



CHAPTER 06
DECISION TEST

A value becomes useful evidence only when the factory can establish what it describes, when it happened and which state or transformation it represents.

A sensor reading becomes industrial evidence only when it is bound to an event time, a stable operational identity and an explicit state or transformation. A counter value alone cannot establish whether output occurred, which asset or lot produced it, whether the line was running, or whether a later replay has already counted it. The chapter's answer is therefore a governed evidence chain: keep event time separate from ingestion time, resolve events against versioned master data, represent genealogy and state transitions explicitly, make replay idempotent, append corrections, and retain each record class for a stated purpose. [S05] [S08] [S09] The sensor value remains an observation. Evidence is the reconstructable account of what that observation meant when a person, KPI or model used it.

This distinction lets a factory migrate from manual hourly logs to machine ingestion without rewriting its operating history. In Cameroon, where GFT's programme evidence base prepared on 31 January 2024 records about 95 percent of cotton fibre exported raw and about 5 percent processed locally against a 50 percent target by 2030 (with original source dates unrecorded and not refreshed in 2026), expanding domestic processing depends on operational records that hold up across shifts, power interruptions and buyer audits. [R06] [R07] Preserving time, identity and state prevents a delayed connection from masquerading as a production stop, a renamed asset from losing its past, a split lot from acquiring a fictitious single batch identity, and a retried message from becoming additional output. [S08] [S33] [S38] It also gives later forecasting and inspection models the same temporal and operational definitions used on the floor, rather than allowing each training pipeline to invent them again. [S04] [S15]

6.1

Event time and ingestion time answer different questions

Every governed event needs two timestamps. Event time records when the represented operational occurrence happened. Ingestion time records when the receiving system accepted its record. EPCIS makes the equivalent distinction explicit: its `recordTime` is a bookkeeping mechanism and does not describe the real-world event. [S0601] A factory should also distinguish processing time, when a particular transformation ran, from both. A stoppage that occurred during one shift but arrived during the next belongs to the first shift's operating account and to the second shift's data-latency account. Those are different questions, and neither timestamp can answer both. [S08] [S33]

6.2 TWO CLOCKS PROTECT THE PRODUCTION ACCOUNT FROM THE NETWORK ACCOUNT

The raw record should preserve the source's event timestamp, its stated time-zone offset and the receiving system's ingestion timestamp without overwriting either. Normalising an additional representation to Coordinated Universal Time makes cross-system ordering consistent, but the original local representation remains useful evidence about what the source asserted. [S08] Ingestion time supports freshness monitoring, replay diagnosis and reconstruction of what a dashboard could have known at a given moment. Event time supports output windows, state intervals, genealogy and features that must not use future information. [S15] [S38]

This separation changes operational behaviour. When a gateway reconnects, the custodian investigates the interval between event and ingestion timestamps as a continuity issue. The production steward does not automatically treat that interval as missing output or unplanned downtime. Conversely, a feed can arrive promptly with a wrong event timestamp because a device clock has drifted. Low ingestion latency therefore proves delivery speed, not temporal accuracy. [S0601] [S33]

6.3 CLOCK SOURCE AND UNCERTAINTY MAKE A TIMESTAMP INSPECTABLE

GFT's event contract adds clock source, observed clock offset or uncertainty, and a source sequence number where the device provides one. These fields are engineering metadata, not decorative compliance. A line-side terminal, programmable controller and manual back-entry form can assert different degrees of temporal confidence even when their displayed timestamps have the same precision. Recording the source and an estimated uncertainty lets a steward decide whether two events can be ordered safely or need review. [S08] [S09]

The estimate needs a method. At commissioning and after a gateway restart, the custodian can compare the device with the site's chosen time reference, record the observed offset and retain the check date. Any uncertainty attached to a record should be labelled an estimate based on that check, not presented as a measured guarantee for every later event. Where the uncertainty crosses an hourly or shift boundary, the record should carry a boundary-review flag rather than a falsely precise allocation. A source sequence can help diagnose order within one device, but it cannot establish order across independent devices and cannot replace event time. [S05] [S0601]

Manual entries require the same honesty. If a supervisor records a stoppage after the event, the form should ask when it happened and separately capture when the entry was submitted. It should not silently replace the first with the second. Back-entry status and the submitting role become provenance, while any unnecessary worker identity remains outside the production feed. [S09] [S29]

6.4 AGGREGATION WINDOWS NEED AN EXPLICIT LATENESS RULE

An hourly bucket is a governed calculation, not a natural property of incoming messages. The KPI contract must state its event-time window, inclusion rule, provisional period, late-event treatment and formula version, consistent with ISO 22400's emphasis on formula, elements, time behaviour and unit. [S38] A live operational view may publish a provisional total so the shift can act, then issue a new calculation version when an accepted late event arrives. A monthly report may use a longer closure period. Neither waiting period is a universal constant; each is a site decision based on decision urgency and observed delay.

The same rule must reach model datasets. If a training feature describes an hour, its manifest should show whether late events were included up to a defined ingestion cut-off and which recalculation version produced the value. Otherwise a model can learn from a more complete history than will exist when it makes a live prediction. [S04] [S15] **Exhibit 25**, p. 100 isolates the temporal mechanism before the worked example adds quantities.

EXHIBIT 25 / MINUTES AFTER 09:00 (ILLUSTRATIVE)

Late arrival changes processing order without changing event time

3

A occurs and arrives

Event minute 3; ingestion minute 3.

42

C occurs into buffer

Event minute 42; no central ingestion yet.

65

D occurs and arrives

D belongs to the next event-time hour.



17

B occurs and arrives

Event minute 17; ingestion minute 17.

60

Retry follows hour cut

Duplicate delivery of A is detected; provisional bucket remains provisional.

67

C arrives and recomputes

Event minute 42; ingestion minute 67; earlier bucket receives a new calculation version.

ILLUSTRATIVE [S08] [S33] [S0601] Six illustrative timeline entries for events A, B, C and D, one transport retry and one KPI recomputation. All times are invented to demonstrate the method. Dual-time semantics derive from S08 and S0601, and buffering from S33. The exhibit does not set a recommended lateness threshold.

The important result in **Exhibit 25**, p. 100 is not the arrival order. It is the preservation of two valid histories: what happened in production, and what the information system knew at each point. C does not move to D's hour merely because it arrived after reconnection, and the retry of A does not create a new production event. The exhibit cannot tell a factory how long to keep a bucket provisional, nor can it prove that the source clocks were correct. Those require site rules and clock evidence. [S0601]

6.5

Master data stabilise the factory's objects

An event identifier says which record is being described. A master identifier says which persistent object the record concerns. The second remains stable while names, locations, assignments and other attributes change. IEC 62264 supplies role-based equipment and physical-asset hierarchies, while OPC UA provides an information-modelling framework for objects and their relationships; neither removes the factory's duty to decide which record is authoritative. [S05] [S06] GFT therefore separates the slow-changing master record from the high-frequency event stream and gives every master class an identified origin, steward and effective history.

6.6 IDENTITY PERSISTS WHILE ATTRIBUTES RECEIVE EFFECTIVE DATES

A sewing asset can keep one immutable asset identity when its display name changes, when it moves to another line, or when its controller is replaced. Those are three different changes. Renaming alters an attribute. Moving alters a role-based assignment with a start and end time. Replacing the controller may alter a technical relationship without replacing the production asset, provided the asset owner confirms that operational continuity remains valid. [S05] [S06] The event never carries a mutable display name as its only link to the object.

Each derivation then resolves the master version effective at event time, not whichever version happens to be current when a query runs. A historical output total remains assigned to the line on which the asset operated at that time. A current capacity view uses the current assignment. Both views can coexist because they state their temporal rule. The curated record should retain the master-data version or effective-dating evidence used in the join, so a later hierarchy edit cannot silently redistribute history. [S05] [S09]

6.7 REGISTRY AUTHORITY PREVENTS ONE OBJECT BECOMING SEVERAL IDENTITIES

Master data also need a merge and retirement rule. If two registers contain the same physical asset, the steward does not delete one and leave dangling events. The registry records a mapping from the superseded identifier to the surviving identifier, the approval, the effective time and the reason. Existing raw events retain their original identifier; curated products apply the governed mapping and preserve that transformation in lineage. [S02] [S09] An identifier that has been retired must not be issued to a different object, because reuse would make old events indistinguishable from new ones.

Exhibit 26, p. 102 turns these rules into a compact registry design. Its entries are illustrative classes and controls, not a claim that a particular factory already operates this registry.

EXHIBIT 26 / ILLUSTRATIVE REGISTRY DESIGN

An identity registry records origin, authority and permitted reuse

Entity class	Authoritative origin	Identity and effective history	Permitted reuse
Asset	Engineering asset register	Immutable asset ID; dated line and controller links	Events, maintenance and resource allocation
Order	Order-management record	Order ID with versioned style, quantity and due attributes	Scheduling, output and delivery reconciliation
Lot	Material receipt or transformation event	Lot ID plus parent and child event links	Genealogy, quality and provenance
Style	Product-definition register	Style ID with dated operation and unit versions	Routing, costing and feature derivation
Operation	Versioned routing definition	Operation ID effective within a style version	Work-centre events and cycle-time analysis
Defect code	Quality taxonomy owner	Stable code with dated label and mapping history	Inspection, rework and model labels
Worker role	Authorised role register	Role ID only; no person ID in production events	Authority checks and shift-level aggregation

ILLUSTRATIVE [S05] [S06] [S29] Illustrative master-identity registry based on equipment and enterprise boundaries in S05, object and relationship modelling in S06, and personal-data minimisation in S29. Rows are recommended design examples, not observed implementations or GFT operating results.

The registry in **Exhibit 26**, p. 102 deliberately gives worker role a different treatment from machine and material identities. A production event may need to show which role was authorised to declare a transition, without attaching a persistent person identifier to every output record. [S29] The table also shows why master data cannot be reduced to a list of codes: each class has an origin, an effective history and a permitted use. What it does not supply is the site-specific code format, which should follow existing operating needs rather than an imposed naming convention.

6.8

Genealogy links inputs, transformations and outputs

Genealogy is a chain of event relationships, not a batch field copied from one database row to the next. EPCIS is designed to share visibility event data about objects in a business context, while PROV-O represents entities, the activities that transform them and the agents responsible. [S08] [S09] Together they support a practical factory model: receipt creates an identified input lot, splitting relates children to their parent, joining records all inputs to a new output, and rework creates a traceable transformation rather than silently changing the original lot's status.

6.9 PARENT AND CHILD LINKS SURVIVE SPLIT, JOIN AND REWORK

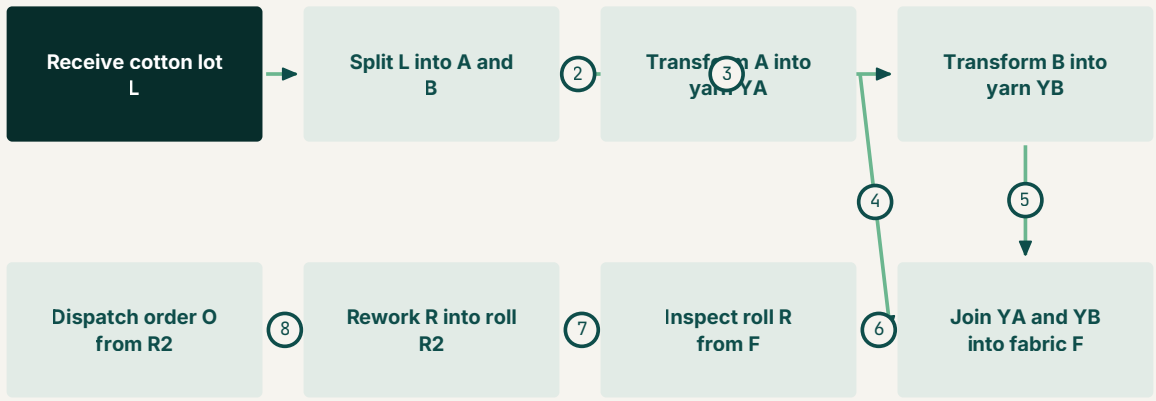
A single `batch_id` on a garment order fails as soon as production has more than one input or output. If a cotton lot is split, each child needs its own identity and a link to the parent. If yarn from two lots enters one fabric batch, the output needs links to both. The graph must be many-to-many, with the transformation event carrying the event time, location, operation and responsible role needed to interpret the relationship. [S08] [S09] Copying only the most recent lot identifier would discard the earlier inputs and make provenance dependent on memory.

Rework needs equal care. The original inspection event and rejected state remain evidence. A rework event identifies the affected object, operation and outcome; a subsequent inspection creates a new disposition. If the physical item retains its identity, the events carry its changing state. If the process creates a new roll, lot or bundle, the new identifier receives a derivation link. The choice follows physical practice, not database convenience. [S08] [S09]

Exhibit 27, p. 103 shows the minimum relationship pattern. It is a derived event design, not a mass balance and not an observed production history.

EXHIBIT 27 / EVENT AND ENTITY RELATIONSHIPS

Event genealogy preserves lot history through split and rework



2. Parent to child
3. Parent to child
4. Input link
5. Input link
6. Fabric roll created
7. Rejected state retained
8. Derived output

DERIVED [S08] [S09] Derived genealogy pattern using visibility events from S08 and entity, activity and derivation relationships from S09. Identifiers are illustrative and links carry no quantities. The diagram does not demonstrate material conservation, legal chain of custody or carbon allocation.

The flow in **Exhibit 27**, p. 103 preserves why illustrative order O has a more complicated ancestry than one lot code can express. It also separates an inspection observation from the transformation that followed it. The links must carry quantities in the actual event records where allocation requires them, but the diagram deliberately makes no conservation claim. WP05 owns provenance and chain-of-custody method; WP16 owns carbon allocation. This chapter supplies the temporal and identity evidence those methods need, not their accounting conclusions. [S08] [S09]

> **GFT NOTE**
>
> GFT should migrate hourly production logs into a small event ledger before connecting more sensors. For each receipt, split, join, inspection, rework and dispatch, retain the event identifier, event type, event and ingestion times, subject identifiers, parent or input links, operation, source and master-data version. Start with the few transformations needed to explain current customer and production decisions. A complete historical graph is less useful than a small chain whose physical meaning has been checked with the people who perform the work. [S08] [S09]

6.10

Explicit states prevent incompatible downtime and defect counts

A state is a governed interpretation of an entity during an interval or after an event, not a free-text label on a dashboard. Machine-time states and unit-disposition states answer different questions. A sewing line can be running while a garment is in rework, and a line can be stopped while an operator inspects a unit. Treating `rework` as a machine state or `blocked` as a defect disposition mixes those dimensions and makes downtime and defect totals irreconcilable. [S38] GFT therefore keeps separate state machines for asset time and material disposition, linked by events but never added together as if they shared a denominator.

6.11 STATE CONTRACTS DEFINE ENTRY, EXIT AND TRANSITION AUTHORITY

Each state definition needs an entity class, meaning, entry condition, exit condition, permitted reason codes, overlaps and authoritative role. The event that changes state records the previous and next values, event time, ingestion time, source and transitioning role. A supervisor may authorise a planned-stop classification, a maintenance role may verify restart after a fault, and a quality role may authorise rework or scrap. The small-team role assignments remain explicit rather than being buried in software permissions. [S02] [S09]

Allowed overlaps must be equally precise. Under an exclusive asset-time model, one asset cannot be running and in an unplanned fault during the same interval. It can, however, be in an unplanned fault while an order is awaiting material and a unit is held for inspection, because those states belong to different entities and dimensions. Where management wants simultaneous technical and business states, the contract should publish separate fields rather than force both into one status column. [S38]

Late and back-entered transitions follow the same temporal rule. If a fault began before it was entered, its interval begins at the evidenced event time and its ingestion record shows the delay. A missing end transition leaves an open interval flagged for

steward action, not an automatic zero-duration stop. If a later correction changes the start or end, the original transition remains available through the append-only correction mechanism described below. [S0601] [S0602]

6.12 KPI MAPPINGS CONSUME EACH INTERVAL AND UNIT ONCE

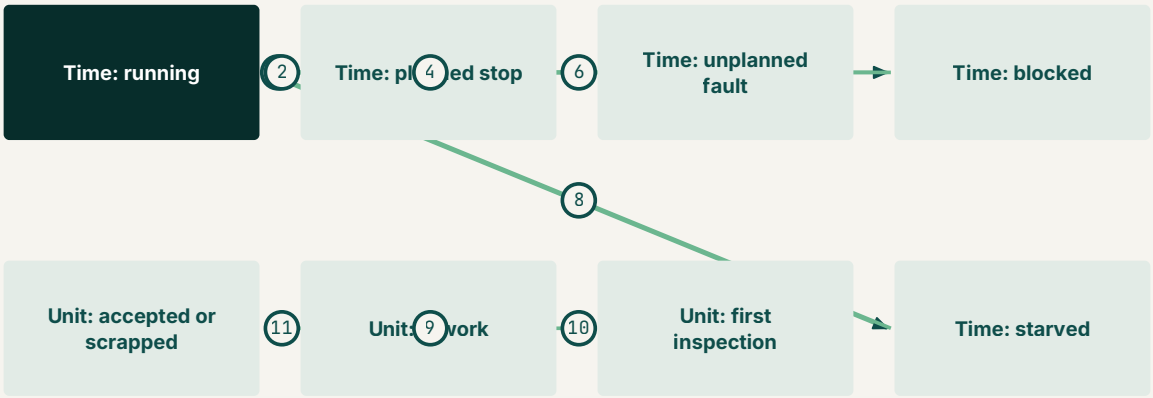
For each asset-time model, the base states should form a complete and non-overlapping partition of the time being measured. Reason codes add diagnosis without creating a second duration. A fault interval carries one fault reason; it does not also create a generic downtime interval that is added to the fault duration. Planned stop, blocked and starved intervals enter availability or other KPI calculations only through the exclusions and time elements in the active KPI version. [S38] The mapping is part of the definition, not an analyst's later preference.

Unit disposition follows a separate event path. First-pass acceptance, rework, reinspection and scrap describe the fate of a unit or lot. Defect counts need a stated unit and inspection boundary, just as downtime needs a stated asset and time boundary. A second inspection after rework must not silently become another first-pass unit. These controls make the KPI numerator and denominator reconstructable rather than merely reproducible as code. [S04] [S38]

Exhibit 28, p. 106 places the two state machines side by side. It illustrates transition logic and does not prescribe one universal KPI formula.

EXHIBIT 28 / ILLUSTRATIVE STATE TRANSITIONS

A governed state machine stops one interval becoming two losses



1. Authorised stop
2. Authorised restart
4. Repair verified
6. Downstream released
8. Input available
9. Accept or scrap directly
10. Repairable defect
11. Reinspection decides

ILLUSTRATIVE [S38] Illustrative asset-time and unit-disposition state machines. Transition design and the requirement to state KPI time elements derive from S38. Reason codes, transition authority and acceptance thresholds remain site decisions. The two tracks must not be added as if they had a common denominator.

The upper logic in **Exhibit 28**, p. 106 allocates asset time once. The lower logic tracks material disposition without creating another downtime category. A blocked interval and a starved interval can carry different operational causes while receiving the same treatment in a particular availability formula, but that formula must be cited by version. [S38] The exhibit cannot determine whether a planned stop should be excluded at a given factory, nor can it establish that a source system's transitions are complete. Acceptance requires interval reconciliation against a known shift calendar and sampled physical observations.

6.13

Replay must not create a second reality

Buffered delivery and transport retries protect availability, but they also create a governance risk: the same application event can arrive more than once or after later events. MQTT supplies publish-subscribe transport, not a factory-specific rule for recognising production identity. [S07] The capture boundary must therefore issue a stable event identifier that survives every retry and buffer cycle, such as a time-ordered RFC 9562 UUIDv7 value generated when the occurrence is recorded. [S0603]

A repeated transmission keeps that identifier; a genuinely new occurrence receives a new one even if its value, asset and timestamp happen to match another record. [S08] [S33]

6.14 STABLE EVENT IDENTIFIERS MAKE CONSUMERS IDEMPOTENT

Idempotency means that processing the same event again produces no additional unintended effect. The raw store accepts one identity once, and every derived consumer either tracks the event identifiers already applied or recomputes from a defined input set rather than blindly incrementing a total. A dataset manifest records the included identifiers and rule versions so a training snapshot can be rebuilt without double counting. [S04] [S15] Deduplication is therefore a chain-wide obligation, not only a broker setting.

A conflict needs a different response from a duplicate. If two messages carry the same event identifier and the same canonical payload hash, the consumer can recognise a retry. If they carry the same identifier but different hashes, it must quarantine an identity conflict and preserve both payloads for investigation. Choosing the latest would allow a transmission fault or unauthorised change to rewrite evidence without explanation. [S04] [S09]

6.15 DETERMINISTIC REPLAY REBUILDS PRODUCTS RATHER THAN REWRITING THEM

A controlled replay reads the preserved raw set, applies a named schema, master-data, state and KPI version, then publishes a new derived version with lineage to the execution. It does not mutate an old product in place and pretend that consumers saw the same result all along. [S09] [S15] Where the same inputs and rules produce a different output, the team has found nondeterminism or an unrecorded dependency. That defect should block acceptance until explained.

METHOD

Replay acceptance tests for a governed feed

- Deliver the same event identifier and payload twice. The raw identity and operational effect must occur once, while the delivery log retains both attempts.
- Deliver the same event identifier with a different payload hash. The feed must create an identity-conflict exception, not select a winner silently.
- Deliver an earlier event after the hourly cut. The consumer must allocate it by event time and issue a traceable recalculation under the active lateness rule. [S0601]
- Reverse two state-transition messages during a test. The state service must apply its documented ordering rule or raise an exception, not manufacture a plausible interval. [S38]
- Rebuild a dataset from an unchanged raw manifest and rule set. The rebuilt content hash must match, or the pipeline must identify the changed dependency before release. [S15]

These tests should run before a feed enters operational or training products and whenever its contract changes. They use synthetic events with known outcomes, so no production order or worker record is needed for the test. [S04] [S29]

WORKED EXAMPLE

Consider an illustrative line with four genuine output events, A, B, C and D, each representing 20 pieces. Their illustrative event times are 09:03, 09:17, 09:42 and 10:05. C is buffered locally. By the illustrative 10:00 hourly cut, A and B have arrived, along with one retried copy of A. A naive message sum gives `3 x 20 = 60` illustrative pieces, but identifier deduplication gives the correct provisional event sum `2 x 20 = 40` illustrative pieces. [S07] [S33]

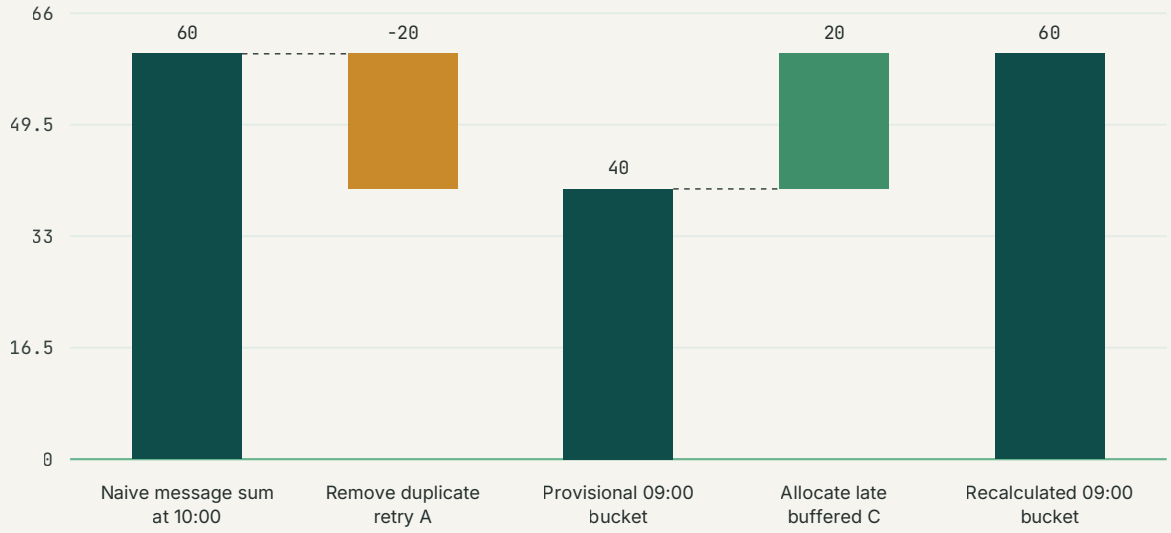
At an illustrative 10:07, C arrives after D. The completed event-time allocation is `A + B + C = 3 x 20 = 60` illustrative pieces in the 09:00 hour and `D = 20` illustrative pieces in the 10:00 hour. The provisional 09:00 total therefore moves from 40 to 60 illustrative pieces through a new calculation version, while the event itself remains at 09:42. Across both hours, the factory has `4 x 20 = 80` illustrative pieces, not the `5 x 20 = 100` obtained by counting the retry as production. [S38]

The two failure modes can now be separated. Counting duplicate A would overstate the illustrative 09:00 hour by 20 pieces. Allocating C by ingestion time instead would understate that hour by 20 pieces and overstate the illustrative 10:00 hour by the same amount. These are constructed values used only to show the arithmetic; they are not output measurements from a GFT client or a benchmark for replay delay. [S0601]

The worked example shows why replay acceptance must reconcile identities, event-time buckets and calculation versions together. Deduplication alone would still misallocate C. Event-time allocation alone would still count the retry. A versionless dashboard could calculate the correct final total while hiding the fact that its earlier figure had changed. **Exhibit 29**, p. 109 summarises that reconciliation for the 09:00 hourly bucket across both adjustments. Operational monitoring should therefore retain provisional revision history, duplicate-conflict counts and lateness distributions alongside the published KPI. [S04] [S38]

EXHIBIT 29 / ILLUSTRATIVE PIECES IN 09:00 BUCKET

Deduplication and event-time allocation reconcile the 09:00 bucket



ILLUSTRATIVE [S07] [S33] [S38] [S0601] [S0603] Illustrative reconciliation from the worked example: 3 arrived messages x 20 pieces = 60 naive pieces at the 10:00 cut, less 20 pieces for the duplicate retry of A (giving 40 provisional pieces), plus 20 pieces when buffered event C arrives at 10:07 and is allocated to its 09:42 event time (giving 60 recalculated pieces). Values are invented for method demonstration, not GFT operating results.

As **Exhibit 29**, p. 109 makes visible, the naive message sum at 10:00 and the recalculated event-time bucket after 10:07 share the same illustrative value of 60 pieces, yet the first figure is wrong twice: it includes a duplicate retry of A and omits buffered event C. [S07] [S33] [S0601] Without identifier deduplication and dual-time accounting, those two errors mask each other in the 09:00 bucket while inflating the two-hour total to 100 illustrative pieces. [S38] [S0603] What **Exhibit 29**, p. 109 cannot establish is the site-specific outage distribution or the cut-off window for closing provisional shift reports, which must be set from measured plant connectivity. [S33]

6.16

Corrections append rather than erase

A correction changes the best supported account of an event; it must not destroy the earlier assertion or the fact that people acted on it. EPCIS error declarations follow this principle directly: the prior event is not modified, and subsequent queries return both the prior event and the declaration. [S0602] In W3C PROV-O, a revision is modelled explicitly as a derivation in which the resulting entity is a revised version of an original entity. [S0604] GFT applies the same pattern across manual and machine records. A superseding event identifies what it corrects, states the reason and changed fields, and records the submitting and approving roles with their declaration times. The original raw payload remains untouched. [S02] [S09]

6.17 CORRECTION TIME AND EFFECTIVE TIME PRESERVE TWO TRUTHS

A supervisor can discover today that yesterday's lot code was wrong. The corrected lot relationship takes effect for the current operational view, but the factory still needs to reconstruct yesterday's view when the original code was believed correct.

Correction time records when the new assertion entered the system. Event time records the production occurrence being corrected. Keeping both prevents a present-day repair from pretending that past users possessed information they did not have. [S0601] [S0602]

The curated model therefore supports two questions without confusing them. “What is our best current account?” follows the supersession chain to the effective record. “What did this report or model know at its execution?” uses the records and correction cut-off captured in its lineage. PROV-O derivation and revision relationships provide a standards-based way to represent those links, while the factory's own schema supplies reason, authority and approval detail. [S09] [S15]

A correction can itself be wrong. The chain must allow a later superseding assertion with the same evidence requirements rather than editing the first correction. Current-state queries need a deterministic rule for selecting the effective tip of that chain, and any branching corrections require steward resolution. This avoids both destructive editing and an ambiguous collection of competing truths.

6.18 RECALCULATION IS A NEW DERIVED VERSION, NOT SILENT HISTORICAL DRIFT

Every material correction triggers an impact assessment. The steward and custodian identify affected state intervals, genealogy paths, KPI windows, published products and dataset snapshots before releasing the new effective state. The pipeline then produces new calculation versions with links to the correction event and the exact rule versions used. A previously issued management report retains its original value and receives an addendum where the operating owner decides that restatement is required. [S02] [S09] [S38]

Frozen training data follow the same rule. The pipeline does not mutate an old snapshot and retain its identifier. It creates a new manifest that records corrected event identifiers, supersession links and the ingestion cut-off, then evaluates whether the change affects a candidate or deployed model. [S04] [S15] This deliberate path distinguishes legitimate historical revision from silent drift and makes the operational consequence of a correction measurable.

6.19

Retention follows purpose and obligation

Retention is a decision about each record class, not a promise to keep everything forever or delete everything on one schedule. Raw evidence may be needed to reconstruct a customer inquiry or diagnose a recurring fault. Curated events need their schemas and master versions. Labels, snapshots and model artefacts need the lineage required to explain historical outputs. Access logs serve security investigation and have a different exposure and purpose. [S09] [S15] [S29] Each class therefore receives a purpose, owner, review trigger, deletion method and, where applicable, contractual or legal basis.

6.20 A DIFFERENTIATED SCHEDULE PRESERVES EVIDENCE WITHOUT INDEFINITE ACCUMULATION

The retention register should address the classes separately:

- **Raw machine and material events:** retain for the period justified by operational reconstruction, warranty, customer and applicable legal needs, with non-personal asset and lot identifiers preferred at capture. [S29]
- **Curated events and KPI results:** retain the effective record together with the schema, master-data, state and calculation versions required to interpret it, rather than preserving a number whose definition has been lost. [S38]
- **Labels, dataset snapshots and model artefacts:** retain for as long as historical predictions or releases require explanation, with the retention decision linked to consequence, contractual need and reusable validation value. [S15]
- **Access and correction logs:** restrict them to their accountability and security purposes, because administrator and approver identifiers can themselves be personal data. [S29]
- **Worker-identifiable records:** set the shortest justified period after the documented purpose and current Cameroonian legal review, using erasure or appropriately assessed anonymisation when that period ends. [S29] [S31]

These are decision fields, not proposed retention durations. A small factory should set actual periods only after reviewing operational reuse, customer terms, litigation or incident holds, and applicable law. Indefinite retention is not a conservative default when it increases privacy exposure and storage burden without a stated decision need. [S29]

6.21 DELETION MUST REACH COPIES WITHOUT FABRICATING A NEW HISTORY

An executable retention procedure covers primary stores, replicas, local buffers, analytical exports and backup cycles. It records a minimal disposal receipt containing record class, scope, authority, execution time and verification result, not the deleted payload. Where immediate deletion from every backup is not feasible, the procedure isolates expired backups from ordinary use and ensures that a restoration cannot silently return deleted personal records to production. A legal or incident hold is scoped, approved and reviewed rather than becoming an invisible permanent exception. [S29] [S31]

Aggregates require the same discipline. Removing direct identifiers does not automatically make a small-group or event-level dataset non-personal. The owner must obtain an appropriate assessment before treating an aggregate as exempt from the personal-data schedule. Where retention ends for raw operational evidence, derived records should not claim a lineage that can no longer be inspected; their metadata should state the disposal boundary honestly. [S09] [S29]

WARNING

Cameroon's Law No. 2024/017 of 23 December 2024 includes a prior-authorisation requirement for personal-data processing, but this chapter does not determine current implementing practice, exemptions or a statutory retention period. [S31] Qualified Cameroonian counsel must review identifiable-worker records, cross-border copies, holds and deletion routines before the factory fixes those controls. This is engineering guidance, not legal advice.

Time, identity and state turn readings into evidence because they make every later use answerable to its original context. Replay, correction and retention then preserve that answerability as systems change. The factory can reconstruct not only a value, but also the object, occurrence, transformation, rule and knowledge state behind it.

Keep the observation, its time and its identity together, or the number is only an assertion.

CHAPTER SOURCE NOTES

S05 International Electrotechnical Commission. 2013.

S08 GS1 AISBL. 2022.

S09 World Wide Web Consortium. 2013-04-30.

R06 Cameroon Trade Hub. Publication date unrecorded; inherited source-text check.

R07 Ecofin Agency. Publication date unrecorded; inherited source-text check.

S33 Microsoft. 2024-06-10.

S38 International Organization for Standardization. 2014.

S04 ISO/IEC. 2024.

S15 Google Cloud. 2024-08-28.

S0601 GS1 AISBL. 2022.

S29 European Parliament and Council of the European Union. 2016-04-27.

S06 OPC Foundation. 2025-10-22.

S02 International Organization for Standardization. 2022.

S0602 GS1 AISBL. 2022.

S07 OASIS Open. 2019-03-07.

S0603 Internet Engineering Task Force. 2024-05.

S0604 World Wide Web Consortium. 2013-04-30.

S31 Presidency of the Republic of Cameroon. 2024-12-23.

GREEN FUTURE TECH

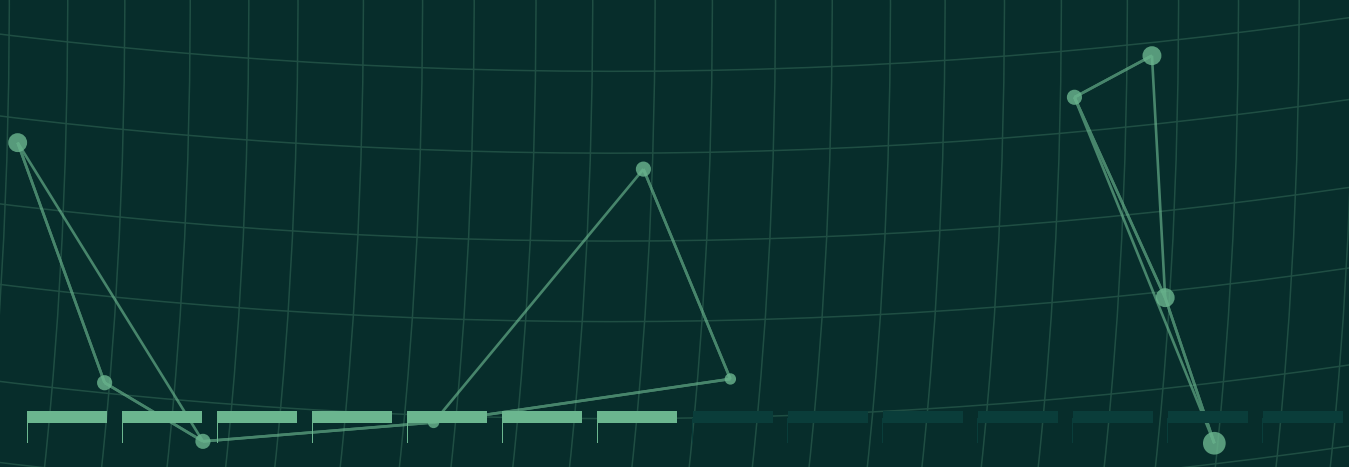
LIMITED / HONG KONG SAR

07

PART II: BUILD THE MINIMUM DURABLE SPINE / RIGHTS AND PROTECTION

Security and privacy are design constraints, not aftercare

The data spine must protect production, people and legitimate access at the same time, with controls proportionate to a small team.



CHAPTER 07
DECISION TEST

The data spine must protect production, people and legitimate access at the same time, with controls proportionate to a small team.

A factory data spine is trustworthy only when its security, privacy and data-access controls are built into the capture boundary, identity model and recovery routine before the first operational feed is connected to analytics. [S22] In GFT's Cameroon textile programme, prepared in January 2024 with source-refresh notes retained in 2026, the sector baseline records about 95 percent of cotton fibre exported raw and about 5 percent processed locally, against a 50 percent local-processing target by 2030 (these inherited entries were not refreshed in 2026 and their original publication dates are unrecorded). [R06] Closing that gap requires factories that capture production, quality, maintenance and energy evidence from the first shift without exposing machinery to unsafe interference, workers to unlawful surveillance, or the enterprise to vendor lock-in. [R07] For a small team without a security department, protection must operate as a compact set of engineering constraints enforced from edge device identity to offline backup restoration. [S24]

When digital projects treat protection as aftercare bolted on after dashboards or machine-learning pilots succeed, three failure modes appear on the shop floor. [S22] First, direct polling or unauthenticated brokers bridged into operational technology (OT) networks can degrade controller scan times, interrupt human-machine interfaces, or expose physical equipment to unauthorised commands. [S22] Second, shared convenience accounts across operators, contractors and ingestion scripts destroy provenance and let one compromised credential alter both historical evidence and live model artefacts. [S25] Third, indiscriminate collection of worker-identifiable productivity or camera records triggers statutory data-protection duties and erodes workforce trust before a lawful basis or human-review safeguard exists. [S29] Treating security, privacy and contractual data rights as primary design constraints prevents these failures within the capacity of a small plant team. [S23]

7.1

OT consequences set the control objective

Industrial data architecture begins with the physical consequences of failure on the factory floor, not with the convenience of cloud ingestion. [S22] Operational technology covers the programmable logic controllers, drive systems, looms, dyeing controllers, cutting tables and environmental monitors that move matter and energy through a plant. [S22] As the National Institute of Standards and Technology (NIST) sets out in its OT security guidance, control environments differ from corporate IT systems because performance, reliability and physical safety are primary operating requirements. [S22] A delayed packet on an office server causes minor annoyance; a delayed or malformed packet on a dye-house temperature loop or spinning frame can spoil a batch, damage actuators, or endanger operators. [S22] Every component of the data spine touching the shop floor must therefore be evaluated first by its worst-case effect on physical safety and production continuity. [S28]

7.2 THE INTERFACE BOUNDARY SEPARATES CONTROL EXECUTION FROM DATA COLLECTION

Standardising the boundary between control execution and enterprise data management stops analytical workloads from destabilising equipment. [S05] Under the IEC 62264 enterprise-control model, real-time sensing, safety interlocks and closed-loop control belong in the lower control domains, while manufacturing operations management, event aggregation and business planning sit across defined interfaces.

[S05] The data spine must never reach directly into safety controllers or bypass local interlocks to extract telemetry. [S22] Instead, machine data cross an explicit collection boundary where an edge gateway reads exposed variables at a bounded polling rate or receives outbound publish-subscribe messages without granting external systems write access to control logic. [S06]

Where machinery supports OPC Unified Architecture (OPC UA), the spine should use both its information modelling and its security model to authenticate applications, sign messages and restrict read versus write permissions at the session and node level. [S06] The OPC UA security specification addresses threats across the physical, hardware and software environments in which industrial servers run, including message spoofing, alteration, replay and resource exhaustion from message flooding. [S0703] Where publish-subscribe transport uses MQTT, engineers must recognise that MQTT is a transport protocol rather than a complete security architecture. [S07] Topic namespaces, transport encryption, client certificates and payload validation must be configured on the broker so an analytical subscriber cannot publish commands back to a machine topic or exhaust broker memory. [S07]

EXHIBIT 30 / CONTROL LAYERS

Protection follows the path from edge identity to restored record



DERIVED [S22] [S24] [S25] [S26] [S0703] Derived six-layer protection stack mapping operational technology safety, zero-trust identity, transport security, storage immutability, model governance and tested offline recovery to the industrial data spine.

As shown in **Exhibit 30**, p. 116, protection runs as a continuous chain from the physical machine boundary to the offline recovery vault. [S22] Each layer enforces an engineering constraint so a failure in an upper analytical component cannot propagate downward into production control. [S28] Across this stack, the factory documents two mandatory rules before any data agent is commissioned: no data-collection service, broker or model may become a single point of failure for machine safety or local operation, and uncontrolled experimentation or active network scanning on control networks is prohibited. [S22] What **Exhibit 30**, p. 116 cannot establish on its own is the site-specific network topology or firewall rule set, which must be engineered against each plant's physical wiring and controller inventory. [S28]

7.3

Least privilege covers people and machine identities

In a small factory, access control often fails not because encryption is weak, but because convenience encourages shared accounts and over-privileged service tokens. [S25] When supervisors share a terminal password, or when an ingestion script uses the database administrator's master credentials, the spine loses both security and provenance. [S09] If a KPI threshold is altered or a batch record is corrected, the audit log cannot show which person or automated job made the change. [S02] Applying least privilege under NIST security and privacy controls requires granting each human user and automated process only the minimum authorisations needed for its assigned task. [S25]

Modern industrial architecture reinforces least privilege with zero-trust principles: no device, user account or service receives implicit trust merely because it sits inside the local area network or on a company workstation. [S0701] Authentication and authorisation of both subject and device occur as discrete checks before a session is established to any data resource. [S0701] This requires separating human identities from machine identities across every stage of the spine. [S25]

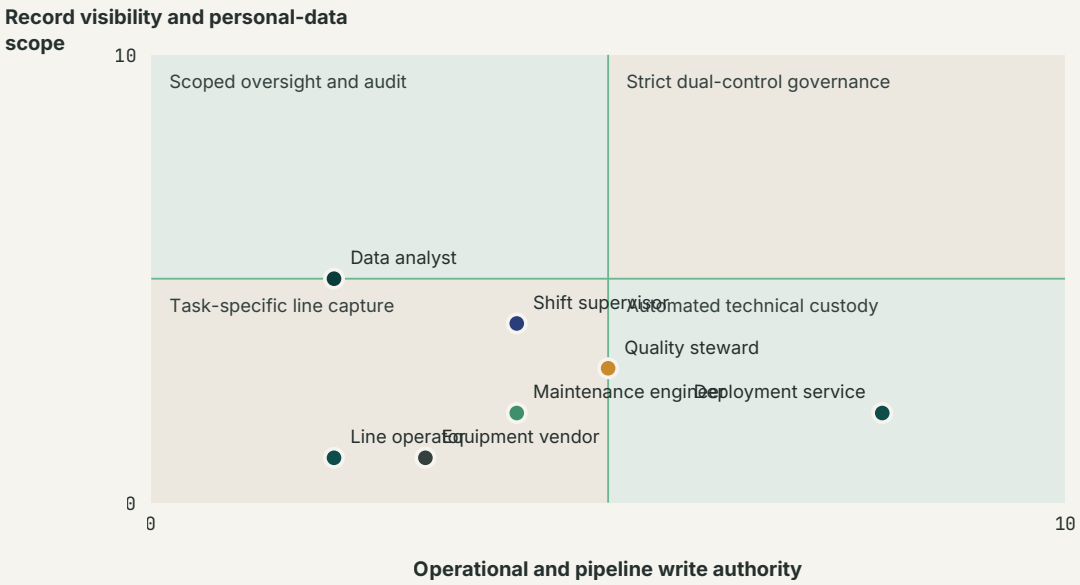
7.4 HUMAN ROLES AND AUTOMATED SERVICES REQUIRE DISTINCT PERMISSION BOUNDARIES

For human personnel, permissions follow the governance roles established in this paper: line operators record and view shift events for their work centre; domain stewards review quarantine queues and submit correction events; technical custodians manage storage, backups and network conduits without authority to alter business definitions; and domain owners approve schema versions, KPI formulas and access exceptions. [S02] External equipment vendors never receive standing administrative credentials; their access is provisioned only for approved maintenance windows through an audited jump host and revoked immediately afterward. [S24]

For non-human processes, every edge gateway, MQTT publisher, ingestion worker, transformation job, training pipeline and registry writer receives its own cryptographic identity and narrowly scoped permission set. [S15] An edge collector on spinning line A needs permission to publish signed telemetry to the raw topic for line A; it has no reason to read order margins, query attendance records, or write to the curated KPI store. [S25] Similarly, a training pipeline needs read-only access to frozen dataset snapshots and write access to the candidate experiment log, never credentials that overwrite the production model pointer or mutate raw events. [S15] Narrowly scoping machine credentials also limits the blast radius of adversarial machine-learning threats such as training-data poisoning or model tampering across the lifecycle stages defined in NIST's adversarial machine-learning taxonomy. [S0702]

EXHIBIT 31 / ILLUSTRATIVE SCORE (0 TO 10)

Access rights differ sharply by role and data product



ILLUSTRATIVE [S25] [S29] [S0701] Illustrative permission mapping across seven factory roles on two governance dimensions: x measures operational and model write authority (0 = read/append local only, 10 = full pipeline or infrastructure write), y measures scope of sensitive and identifiable record visibility (0 = anonymised or single-asset only, 10 = plant-wide or personal-data review). Positions are illustrative design targets based on S25, S29 and S0701, not empirical measurements.

The permission separation in **Exhibit 31**, p. 118 shows why no single role should combine broad pipeline write authority with unrestricted visibility into identifiable worker or commercial records. [S25] In this illustrative matrix, a deployment service holds high write authority to promote approved model containers (x = 8 on the illustrative zero-to-ten scale), yet its visibility into personal records stays near the floor (y = 2) because it processes only governed operational features. [S29] Conversely, a data analyst inspects cross-line curated views (y = 5 on the illustrative scale) while holding minimal write authority (x = 2), preventing silent manual edits to production tables. [S25] What **Exhibit 31**, p. 118 cannot replace is a concrete access-control list in the identity provider and database engine, reviewed whenever staff change role or leave the factory. [S25]

To handle genuine emergencies without weakening daily controls, the spine includes a documented break-glass procedure. [S25] If an urgent night-shift fault requires elevated access when the primary custodian is unavailable, a time-locked emergency credential can be activated by the duty manager. [S27] Activation triggers an audit alert, expires after an illustrative four-hour maintenance window, and requires a formal review at the next morning meeting to verify every command executed. [S25]

7.5

A small-team baseline prioritises essential cyber hygiene

Small manufacturers rarely possess a security operations centre, and copying the control catalogue of a bank guarantees shelfware that nobody follows. [S24] The Cybersecurity and Infrastructure Security Agency (CISA) designed its Cross-Sector Cybersecurity Performance Goals to help small and medium organisations establish a prioritised starting posture. [S24] Combined with the governance function in the NIST Cybersecurity Framework 2.0, which embeds cybersecurity accountability into leadership and enterprise risk management, a small plant can maintain a durable baseline without a specialist department. [S23] Rather than treating these references as one universal checklist, GFT adapts them into six hygiene disciplines within the ISA/IEC 62443 industrial control lifecycle. [S28]

1. **Maintain a live hardware, software and data-link inventory.** You cannot protect or recover an edge box, PLC interface or broker that is absent from the plant register. [S24] The custodian records each device's location, network address, software version, connected control assets, feeds produced and owner. [S22]
2. **Enforce multifactor authentication and eliminate default credentials.** Every remote connection, cloud console, model registry and privileged account requires multifactor authentication, and all default passwords on gateways, cameras and switches are replaced before commissioning. [S24]
3. **Segment operational, corporate and guest networks.** Production controllers and edge collectors sit in segmented zones separated from office workstations, payroll computers and visitor wireless networks by stateful firewalls or conduits. [S22] Data cross from the OT zone to the spine through whitelisted ports using outbound-initiated or unidirectional flows. [S28]
4. **Coordinate patching with maintenance windows.** Unpatched edge systems and broker libraries create exposure, yet rebooting a gateway mid-shift can disrupt batch tracking. [S22] The custodian aligns patching with scheduled mechanical stops, testing updates in staging before plant-wide release. [S28]
5. **Centralise tamper-evident security and access logs.** Authentication failures, privilege changes, firewall blocks, schema approvals and break-glass activations are forwarded to an append-only log store separated from the monitored systems. [S25]
6. **Gate and monitor all third-party vendor access.** Equipment and vision vendors often request permanent remote access for diagnostics. [S22] The baseline prohibits always-on vendor tunnels (`0` permanent unmonitored remote connections as an illustrative target), using time-bounded jump-host sessions approved by the maintenance owner and logged end to end. [S24]

METHOD

System boundary and verification checks for the small-team security baseline

- **Boundary:** Covers all edge capture nodes, local store-and-forward buffers, message brokers, raw and curated stores, model registries and jump hosts in the data spine. [S28]
- **Quarterly verification checks:**
- *Inventory reconciliation:* Sample an illustrative 10 physical ports or wireless clients in the production zone and confirm that 100 percent match registered assets. [S24]
- *Credential audit:* Verify that zero default passwords remain active and 100 percent of departed or reassigned staff have had obsolete permissions revoked. [S25]
- *Conduit test:* Attempt an unauthorised connection from the office network to an OT controller port during a scheduled test window and verify that the firewall blocks and logs the packet. [S22]
- *Vendor session review:* Reconcile every external vendor login in the jump-host log against a signed maintenance work order. [S24]

Placing these six disciplines inside the ISA/IEC 62443 lifecycle keeps security active across specification, commissioning, operation, modification and retirement. [S28]

7.6

Recovery evidence matters more than backup completion

A green checkmark on a nightly backup scheduler does not prove that a factory can recover its data spine after ransomware, hardware failure or operator error. [S26] Joint guidance from CISA, MS-ISAC, NSA and the FBI emphasises two requirements: maintain offline, encrypted backups of critical data, and regularly test their availability and integrity in a disaster-recovery scenario. [S26] Backing up only the raw sensor table is insufficient; if schemas, master-data hierarchies, KPI dictionary versions, transformation scripts, provenance graphs and approved model artefacts are lost, the restored numbers are uninterpretable. [S15]

The recovery scope therefore encompasses the full semantic and operational chain: the append-only raw event store, versioned master data, the KPI dictionary and contracts, PROV-O lineage records, model registry metadata and deployment configurations. [S09] At least one encrypted copy is held offline or in an immutable vault whose retention lock cannot be shortened or deleted using the credentials of the active database administrator or pipeline service. [S26]

Proof of recoverability comes from restoration drills into an isolated environment, never from backup logs alone. [S26] During a drill, the custodian and a domain steward take the offline backup media, provision a clean isolated host, restore the raw records and metadata, recompute a historical shift KPI sheet and re-run a historical model

batch. [S15] The test passes only when actual recovery time meets the recovery time objective (RTO), the missing-data interval meets the recovery point objective (RPO), and restored KPI outputs match the signed management report. [S26]

WORKED EXAMPLE

Illustrative reconciliation of backup intervals, local edge buffering and restoration time

Consider an illustrative garment and spinning plant operating 3 shifts of 8 hours each (24 hours a day) across 6 production lines, generating an average of 500 governed events per line per hour (`6 x 500 = 3,000` events per hour across the plant, an illustrative rate used solely to demonstrate the sizing method). Suppose the central plant database takes encrypted snapshots to an offline vault every 6 hours, while each line's edge gateway maintains a local encrypted store-and-forward disk buffer capable of holding 72 hours of local events (`72 x 3,000 = 216,000` events, taking approximately 216 megabytes at an illustrative 1 kilobyte per signed JSON event). [S33]

At 14:30 on a Tuesday (2.5 hours after the 12:00 offline snapshot), a storage-array failure corrupts the central database. Because the edge gateways isolate cleanly and continue buffering locally during the outage, the lines continue operating safely under local visual controls. [S22] The custodian initiates a restoration into a replacement server from the 12:00 offline encrypted backup:

1. **Offline vault transfer and decryption (illustrative):** 45 minutes to mount and verify cryptographic hashes of the 12:00 snapshot. [S26]
2. **Database, schema and KPI dictionary restoration (illustrative):** 35 minutes to restore tables, master data and lineage indices. [S09]
3. **Edge store-and-forward replay and deduplication (illustrative):** Once the central store is online at 15:50 (80 minutes after the 14:30 fault), the 6 edge gateways replay buffered events generated between 12:00 and 15:50 (`3.83 hours x 3,000 events/hour = 11,500` buffered events). At an illustrative ingestion rate of 500 events per second with idempotent event-ID checking, replay completes in 23 seconds (`11,500 / 500 = 23` seconds), rounded to 5 minutes including steward verification. [S33]
4. **Steward reconciliation sign-off (illustrative):** 15 minutes to verify that the recomputed morning shift KPIs match the pre-fault supervisor log. [S02]

Result: Total verified recovery time is `45 + 35 + 5 + 15 = 100` minutes (1 hour and 40 minutes), against an illustrative target RTO of 120 minutes. Without local edge buffering, the missing-data window would have been 2.5 hours (`7,500` lost events between 12:00 and 14:30); with tested offline restoration plus idempotent edge replay, verified data loss is reduced to `0` committed events. This calculation is illustrative and shows why backup frequency, edge buffer capacity and restoration rehearsal must be engineered as a single equation. [S26]

7.7

Incident response is practised as factory work

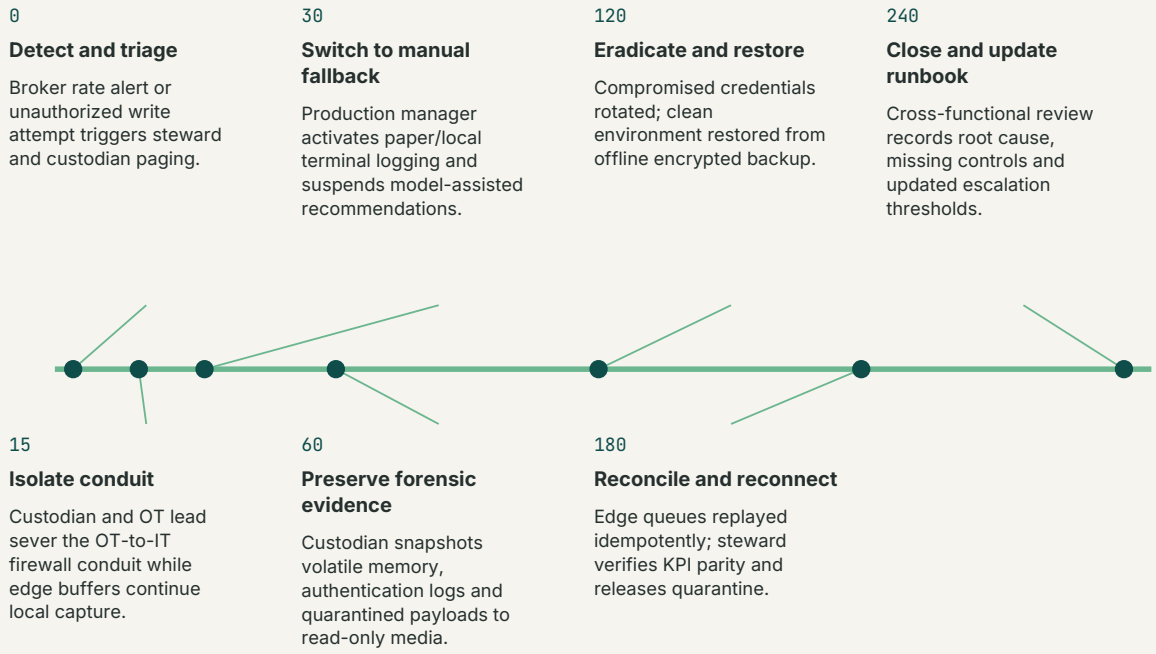
When an anomaly strikes a connected factory, ambiguity over who has authority to disconnect a network conduit or suspend an automated model costs more time than the technical fault itself. [S27] If an IT administrator hesitates to isolate a compromised

broker for fear of halting a weaving room, or if a line supervisor continues following corrupted model recommendations because no manual fallback rule was issued, a containable incident escalates into an operational shutdown or quality recall. [S22] NIST's incident-response guidance aligned with the Cybersecurity Framework 2.0 stresses that response procedures must be tested and exercised periodically both to verify their accuracy and to train personnel before a crisis occurs. [S27]

In a manufacturing environment, incident response is not an abstract tabletop discussion for software engineers; it is a physical shop-floor drill comparable to a fire drill or a lockout-tagout safety exercise. [S22] The response plan must predefine clear decision authority across five roles: the production manager (who owns physical line operation and manual fallback execution), the OT/maintenance lead (who owns controller safety and physical network disconnection), the data custodian (who owns broker isolation, evidence preservation and backup restoration), the domain steward (who owns data quarantine and post-recovery reconciliation), and the plant director (who owns external legal, regulatory, customer and vendor notifications). [S23]

EXHIBIT 32 / ELAPSED MINUTES FROM DETECTION (ILLUSTRATIVE)

An exercised response contains the incident and preserves safe production



ILLUSTRATIVE [S22] [S26] [S27] Illustrative seven-step incident-response and recovery sequence measured in minutes from initial anomaly detection (minute 0 to minute 240). Timings are illustrative drill targets derived from the procedural stages in S22, S26 and S27, not empirical benchmarks.

The sequence in **Exhibit 32**, p. 122 shows how containment and production continuity run in parallel. [S22] Within an illustrative 15 minutes of confirming an unauthorised write or broker anomaly, the OT lead and data custodian isolate the firewall conduit between the shop floor and the server room. [S22] Because the edge architecture is engineered for disconnection, machines remain running safely under local loop control and operators switch by minute 30 (illustrative) to rehearsed fallback sheets or

local terminal logging. [S33] Forensic evidence is captured at minute 60 before any system is wiped, clean restoration from offline backups occurs at minute 120, and idempotent replay reconciles the buffered interval by minute 180, followed by review at minute 240. [S26] While the minute marks in **Exhibit 32**, p. 122 are illustrative drill targets, the ordering is mandatory: isolate the control boundary first, stabilise safe manual production second, preserve evidence third, and reconnect only after clean restoration and steward reconciliation. [S27]

Rehearsing this sequence twice a year exposes hidden dependencies that diagrams miss, such as a fallback sheet lacking the current lot-identifier format, a backup decryption passphrase locked in an absent manager's desk, or an inspection station whose lighting controller depended on the office network switch. [S27] Every defect discovered during a drill enters the governance action register with a named owner and due date. [S02]

7.8

Worker data are limited to a stated legitimate purpose

A factory data spine captures events generated by machines and human operators working together, so operational telemetry can easily cross into personal-data processing. [S29] Attendance logs, operator-tagged sewing outputs, maintenance sign-offs, skill matrices and overhead cameras aimed at inspection tables can all identify individual workers directly or indirectly. [S29] In Cameroon, where the inherited register notes that informal employment accounts for about 90 percent of the labour market alongside a standard 40-hour workweek and 48-hour maximum with overtime at 1.2 times normal pay (inherited entry not refreshed in 2026; original publication date unrecorded), formal industrial employers carry a clear duty to show that digitalisation respects labour standards and worker dignity. [R16]

Two principles govern worker-related information in the spine: strict data minimisation at collection, and mandatory human authority over employment decisions. [S29] Under Recital 39 of the European Union General Data Protection Regulation (GDPR), personal data must be adequate, relevant and limited to what is necessary for the stated purpose. [S29] On the shop floor, minimisation means the spine defaults to asset-level, lot-level and shift-level identifiers rather than person-level identifiers wherever the operational decision does not require knowing the individual worker. [S29]

7.9 OPERATIONAL MEASUREMENT DOES NOT REQUIRE UBIQUITOUS WORKER PROFILING

To calculate line overall equipment effectiveness (OEE), track cotton-lot genealogy, balance work-in-progress buffers, or allocate electricity and steam per kilogram of dyed fabric, the spine needs the machine ID, line ID, style code, order number, shift code and event timestamp. [S05] None of those calculations requires a worker ID on every piece-count message or continuous video of operators in a central archive. [S29] Where cameras support automated fabric-defect detection, the optical field of view should be framed mechanically onto the moving textile web and inspection platen, excluding workers' faces and bodies at the lens boundary rather than relying on downstream blurring. [S29]

Where a legitimate purpose requires identifiable worker records, such as statutory payroll and overtime compliance, safety incident reporting, or certified operator qualification for hazardous dyeing and boiler plant, those records belong in a restricted people-domain store with separate access controls and shorter retention schedules. [S25] Analytical views published to engineers or machine-learning pipelines strip direct identifiers and aggregate metrics to the team, line or shift grain so operational analytics cannot become covert surveillance. [S29]

Equally important is the boundary on automated workplace decisions. [S30] Principle 5.5 of the International Labour Office (ILO) code of practice on workers' personal data states that decisions concerning a worker should not be based solely on automated processing of that worker's personal data. [S30] Although the ILO code is non-binding guidance rather than statute, GFT adopts Principle 5.5 as a hard constraint across all factory analytics and MLOps workflows. [S30] No KPI score, vision defect attribution or model output may automatically trigger discipline, pay deduction, bonus forfeiture, demotion or dismissal. [S30] Any workforce-related indicator serves at most as an inquiry prompt for a trained supervisor, who must inspect machine conditions, material quality and tooling before reaching a judgement under the human-oversight principles of NIST's AI Risk Management Framework and ISO/IEC 23894. [S11]

7.10

Cameroonian legal duties require current local advice

Beyond engineering minimisation and internal rules, a factory in Cameroon must comply with national data-protection law. [S31] On 23 December 2024, the Presidency of the Republic of Cameroon enacted Law No. 2024/017 relating to personal data protection in Cameroon. [S31] That statute regulates the collection, recording, storage, modification, transfer and destruction of personal data within the national territory. [S31] Crucially for industrial digitalisation, Section 19(1) of Law No. 2024/017 states verbatim: "The processing of personal data shall be subject to prior authorization by the Personal Data Protection Authority." [S31]

Section 19(1) has immediate consequences for how the data spine is phased and documented. [S31] An engineering team cannot switch on biometric time-clocks, operator-linked productivity tracking, cloud human-resources analytics or cross-border replication of identifiable employee logs and assume internal policy notes satisfy the law. [S31] Before any identifiable personal data are collected, linked across domains or transmitted outside Cameroon, the enterprise must obtain current local legal advice from qualified Cameroonian counsel on prior-authorisation filings, applicable exemptions, the operational status and procedures of the Personal Data Protection Authority, and cross-border transfer rules. [S31] This white paper provides engineering design guidance only; it is not legal, regulatory, investment or tax advice. [S31]

EXHIBIT 33 / GOVERNANCE REGISTER SCHEMA

A processing register separates operational data from personal-data obligations

Record class	Stated purpose	Identifiability	Automated use rule	Legal and transfer gate
Machine cycle and energy logs	OEE, maintenance and carbon intensity	Non-personal (asset and shift ID only)	Permitted for anomaly and drift alerts	Verify vendor contract rights (S32 where applicable)
Fabric roll vision images	Automated defect detection and grading	Non-personal (camera framed on fabric web)	Operator-assisted roll disposition	Physical lens check; local edge storage default
Order and lot genealogy	Customer traceability and rework control	Non-personal (lot, order and station ID)	Rule-based hold and quarantine routing	Commercial confidentiality; export format contract
Operator attendance and payroll	Statutory wage, overtime and shift safety	Personal data (named worker identity)	No automated discipline or pay deduction (S30)	Law No. 2024/017 Sec. 19(1) review before processing
Safety and machine qualification	Verify certified operator for hazardous plant	Personal data (worker ID and licence date)	Hard interlock check only; human sign-off	Local counsel authorisation review; local hosting
Custodian and break-glass audit logs	Security forensics and privilege accountability	Personal data (administrator account ID)	Alert routing to security owner only	Defined retention window; restricted vault access

DERIVED [S29] [S30] [S31] [S32] Derived processing-register structure separating non-personal machine and lot records from identifiable worker records across purpose, identifiability, access, retention, automated-decision limits, Cameroonian legal screening and cross-border routing.

Maintaining the processing register in **Exhibit 33**, p. 125 gives management, legal counsel and workforce representatives one inspectable inventory of every dataset in the factory. [S29] Separating non-personal machine, energy, vision and lot records from the narrow set of personal records required for payroll, safety certification and security auditing reduces the statutory compliance surface while protecting worker privacy by design. [S29] For every row classified as personal data, **Exhibit 33**, p. 125 binds the dataset to the ILO non-automation safeguard and holds collection or reuse behind an explicit legal-review gate under Cameroon's Law No. 2024/017. [S31] What **Exhibit 33**, p. 125 cannot do is replace a formal legal opinion or regulatory filing with the Cameroonian authority, which must be executed by local counsel before deployment. [S31]

WARNING

Unresolved statutory implementation details require pre-collection legal verification

While Law No. 2024/017 of 23 December 2024 establishes the prior-authorisation requirement in Section 19(1), administrative procedures, filing schedules and cross-border adequacy mechanisms evolve through regulatory decrees and authority practice. [S31] A factory team must never assume that anonymisation claims, standard contract clauses or foreign cloud certifications satisfy Cameroonian law. [S31] Until local counsel confirms authorisation status for a personal-data workflow, the spine must operate strictly on non-personal asset, lot and shift identifiers. [S29]

7.11

Equipment contracts must preserve lawful data access

The final protection constraint covers commercial and legal control over operational data generated by purchased machinery and software. [S32] Emerging-market textile plants acquire spinning frames, looms, dyeing vessels, compressors and inspection rigs from international original equipment manufacturers (OEMs). [S06] Modern machinery frequently ships with embedded sensors and proprietary gateways that transmit telemetry to the vendor's cloud portal while locking the factory out of the raw stream, charging recurring fees to export the plant's own operating history, or barring the factory from sharing diagnostics with an independent maintenance or analytics team. [S32]

When an OEM contract restricts access to machine telemetry, the data spine is crippled at capture. [S05] The plant cannot join machine states to its order genealogy, audit how the vendor's dashboard calculates availability, or migrate to a new analytics tool without losing history. [S09] Within the European Union, Regulation (EU) 2023/2854 (the EU Data Act) addresses this directly: Recital 5 states that users of a connected product or related service in the Union can access, in a timely manner, the data generated by the use of that product or service and share them with third parties of their choice. [S32]

7.12 PROCUREMENT MUST CONVERT STATUTORY CONCEPTS INTO EXPLICIT CONTRACT CLAUSES

For a factory in Cameroon or structured through Hong Kong SAR, the EU Data Act does not apply automatically by territory; its obligations depend on whether the user, product or service falls within the Regulation's territorial and material scope. [S32] An engineering team therefore never assumes foreign statutory rights will protect a Cameroonian mill automatically. [S32] Instead, GFT translates those access principles and industrial interoperability standards into five procurement clauses negotiated before any machinery or software order is signed: [S06]

- **Unencumbered local access to telemetry:** The factory retains the irrevocable right to read all operational, diagnostic, energy and quality data generated by the equipment at the local network boundary, free of volume surcharges. [S32]

- **Documented open interface standards:** Equipment must expose telemetry and event states through documented open standards, such as OPC UA information models or MQTT schemas with unit, timestamp and quality definitions, rather than undocumented binary registers. [S06]
- **Right to share operational data with chosen third parties:** The factory may share its machine-generated operational records and fault logs with independent maintenance engineers, auditors or analytics advisors without vendor consent, while respecting valid trade-secret protections for proprietary control algorithms. [S32]
- **Separation of warranty from read-only collection:** Connecting a read-only edge collector to the documented interface or isolating the vendor's remote modem outside maintenance windows must not void the equipment warranty. [S22]
- **Complete exit and historical export obligations:** Where equipment includes a hosted portal or bought vision software, the vendor must provide a tested bulk export of historical records, calibration parameters, defect labels and schemas in open machine-readable formats on schedule and at contract exit. [S15]

> **GFT NOTE**
>
> **Practical recommendation for GFT's Cameroon textile programme**
>
> In GFT's Cameroon textile programme, the data spine starts with governed manual and local digital logs before scaling to bought computer vision and later demand forecasting. [R06] GFT recommends four protection gates before any network collector or external tool is approved: 1. Complete an OT consequence and conduit check under NIST SP 800-82 so no data agent can interrupt machine safety or local production. [S22] 2. Enforce separate human and service identities with multifactor authentication on privileged accounts and offline encrypted backups verified by a quarterly restoration drill. [S25] 3. Strip worker-identifiable fields from production, quality and energy feeds by default, prohibit automated employment decisions under ILO Principle 5.5, and hold any personal-data processing behind a written legal review under Cameroon's Law No. 2024/017. [S31] 4. Insert standard local data-export and open-interface clauses into every request for quotation for spinning, weaving, dyeing and vision equipment. [S32]
>
> These controls connect to the power and utility resilience designs in `wp07`, labour governance in `wp09`, regulatory screening in `wp17`, and edge topology in `wp18`, with site exposure scored in GFT's `risk-matrix` tool. [S14]

Designing security, privacy and contractual data rights into the spine from the outset prevents the delays caused by safety trips, ransomware lockouts, labour disputes, regulatory stop-orders and vendor lock-in. [S23] With the protection stack, access

matrix, incident-response drill, processing register and procurement clauses in place, the factory can now address the next requirement of Part II: keeping the spine continuous during network outages and portable across vendor changes. [S33]

A factory owns its evidence only when it can collect it safely, restrict it fairly, restore it cleanly and export it freely.

CHAPTER SOURCE NOTES

S22 National Institute of Standards and Technology. 2023.

R06 Cameroon Trade Hub. Publication date unrecorded; inherited source-text check.

R07 Ecofin Agency. Publication date unrecorded; inherited source-text check.

S24 Cybersecurity and Infrastructure Security Agency. 2025.

S25 National Institute of Standards and Technology. 2020-12-10.

S29 European Parliament and Council of the European Union. 2016-04-27.

S23 National Institute of Standards and Technology. 2024-02-26.

S28 International Society of Automation. 2026.

S05 International Electrotechnical Commission. 2013.

S06 OPC Foundation. 2025-10-22.

S0703 OPC Foundation. 2025-10-22.

S07 OASIS Open. 2019-03-07.

S26 CISA, MS-ISAC, NSA and FBI. 2023.

S09 World Wide Web Consortium. 2013-04-30.

S02 International Organization for Standardization. 2022.

S0701 National Institute of Standards and Technology. 2020-08.

S15 Google Cloud. 2024-08-28.

S0702 National Institute of Standards and Technology. 2024-01.

S27 National Institute of Standards and Technology. 2025.

S33 Microsoft. 2024-06-10.

R16 WageIndicator. Publication date unrecorded; inherited source-text check.

S30 International Labour Office. 1997.

S11 National Institute of Standards and Technology. 2023-01-26.

S31 Presidency of the Republic of Cameroon. 2024-12-23.

S32 European Parliament and Council of the European Union. 2023-12-13.

S14 ISO/IEC. 2023.

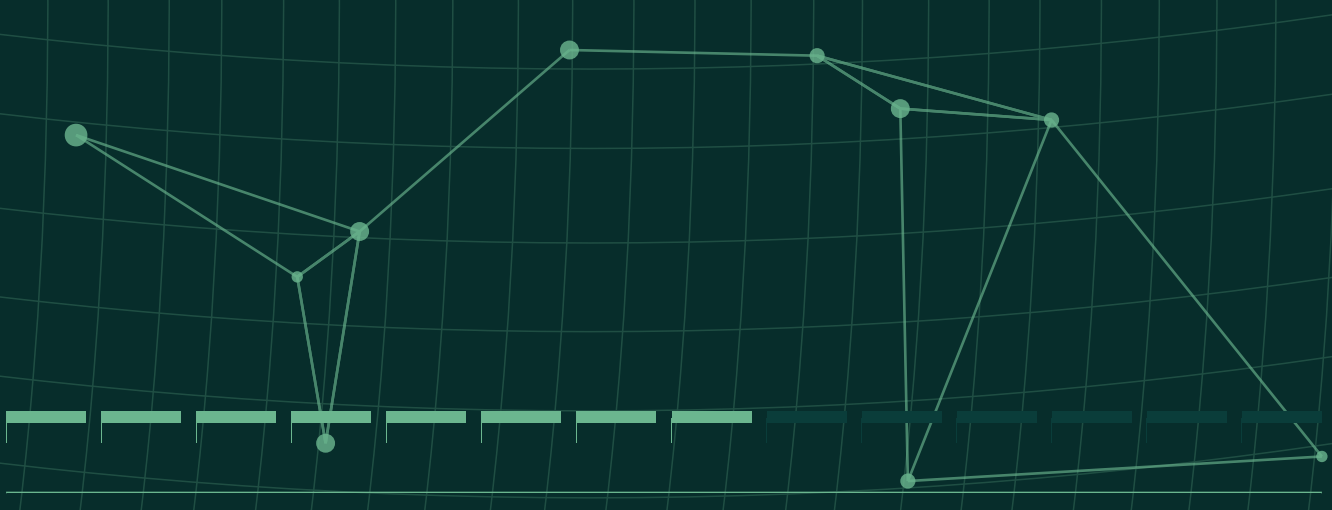
GREEN FUTURE TECH
LIMITED / HONG KONG SAR

08

PART II: BUILD THE MINIMUM DURABLE SPINE / CONTINUITY BY DESIGN

Continuity depends on recovery and portability, not perpetual connectivity

A factory proves resilience by operating through disconnection, restoring its records and leaving a vendor without losing its history.



CHAPTER 08
DECISION TEST

A factory proves resilience by operating through disconnection, restoring its records and leaving a vendor without losing its history.

A factory in an emerging industrial corridor cannot design its data architecture around uninterrupted wide-area connectivity or permanent reliance on one vendor. [S22] [S33] In Cameroon, where the inherited programme evidence base prepared on 31 January 2024 records about 95 percent of cotton fibre exported raw and 5 percent processed locally against a 50 percent target by 2030 (entries not refreshed in 2026 and with original publication dates unrecorded), new spinning, weaving, dyeing and garment plants must maintain high utilisation to stay solvent. [R06] [R07] Earlier programme evidence also notes that capacity underutilisation at legacy domestic textile facilities carried heavy financial penalties in the July 2023 reporting record. [R12] While the industrial data spine neither supplies power nor replaces physical safety interlocks on operational technology, it must guarantee that whenever the line runs, the plant can capture work events, make local shop-floor decisions, recover corrupted stores to a verified state, and migrate historical records across platforms without losing operational meaning. [S05] [S22] [S26] Continuity is an engineering property verified through store-and-forward buffering, class-specific restoration testing, semantic portability and fault injection, rather than a cloud contract promise. [S27] [S0801]

8.1

Store-and-forward buys time during disconnection

Telecommunications interruptions, expired certificates and routing faults routinely sever the link between a factory floor and remote cloud environments. [S22] [S33] Relying on synchronous remote writes to record lots, inspections or stoppages forces a bad choice during an outage: halt capture or stall operator terminals. [S08] [S22] Store-and-forward architecture eliminates this coupling by placing a durable local queue at the site boundary so capture proceeds at line speed regardless of upstream link state. [S07] [S33]

8.2 DURABLE LOCAL QUEUES DECOUPLE LINE CAPTURE FROM UPSTREAM LINKS

In a hierarchical edge topology, child sensors, terminals and controllers publish events to a local parent gateway or broker inside the factory network rather than addressing external endpoints directly. [S22] [S33] Commercial edge runtimes illustrate this: when connectivity drops, the parent device stores device-to-cloud messages locally until the connection returns while managing local module communication and child-device authentication. [S33] The data spine specifies store-and-forward behaviour as a vendor-neutral requirement across all brokers and ingestion nodes. [S07] [S10]

Four technical rules govern a durable local buffer. [S07] [S08] [S09] First, the broker must persist incoming payloads to non-volatile disk upon acknowledgement rather than holding them in volatile memory. [S07] [S33] Second, every buffered message must carry both the capture timestamp and the gateway receipt timestamp so delayed draining is never mistaken for late production. [S08] [S09] Third, each event must carry a deterministic edge identifier so retries are processed idempotently without double-counting bundles or downtime. [S07] [S08] Fourth, reconnection draining must follow a rate-limited priority schedule that sends active alerts and lot dispositions ahead of bulk diagnostics. [S22] [S33]

8.3 BUFFER SIZING IS SITE-SPECIFIC ARITHMETIC FROM MEASURED EVENT RATES

Local storage capacity cannot be copied from a vendor datasheet because event frequency and payload size vary by orders of magnitude across capture methods. [S08] [S33] A factory logging barcode scans, bundle completions and machine state changes generates compact text records, whereas a line retaining raw camera frames from automated fabric inspection can exhaust an undersized drive within one shift. [R29] [S08] Buffer sizing is an explicit arithmetic exercise derived from the plant's event rates, payload schemas, media retention policy and an assumed outage duration. [S33] [S0801]

WORKED EXAMPLE: SIZING A 72-HOUR STORE-AND-FORWARD BUFFER FOR A GARMENT AND INSPECTION FLOOR

Consider an illustrative garment and fabric-inspection facility sizing local edge storage for an assumed 72-hour telecommunications outage across three 8-hour shifts per day. Inputs are illustrative engineering parameters, not measured GFT results. [S07] [S08] [S33]

1. **Structured production and quality visibility events:** Suppose 12 sewing lines and 4 cutting and finishing stations generate an illustrative 2,500 bundle-tracking, defect-coding and downtime events per hour. Each event is serialised as JSON at an illustrative average size of 1.2 kilobytes (kB), plus a 50 percent index and log overhead (1.8 kB effective storage per event). Over 72 hours, those 180,000 events require 324,000 kB, or 0.32 gigabytes (GB). [S08] [S09]
2. **Machine, power and environmental telemetry:** Suppose 80 shop-floor meters and machine interfaces publish status and energy readings every 10 seconds (28,800 messages per hour). At an illustrative footprint of 0.5 kB per reading, 1 hour requires 14.4 megabytes, and 72 hours require 1.04 GB. [S06] [S07]
3. **Shift KPI snapshots and operator correction logs:** Suppose supervisors and stewards record 120 hourly line summaries, reconciliation notes and quarantine decisions per shift (360 records per day) at an illustrative 8 kB per record including provenance metadata. Over 3 days, those 1,080 records consume 8.64 megabytes, or 0.01 GB. [S09] [S38]
4. **Computer-vision inspection artefacts:** Suppose 2 fabric-inspection tables run local optical defect detection at up to 60 metres per minute. [R29] Buffering the structured defect event alongside a compressed 250 kB image crop for each flagged anomaly (at an illustrative rate of 400 anomalies per hour across both tables) requires 28,800 images over 72 hours, totalling 7.20 GB. By contrast, buffering continuous raw video (at an illustrative 15 GB per hour per camera) across 2 cameras would demand 2,160 GB (2.16 terabytes), overflowing a 256 GB drive inside 9 hours. [R29] [S33]

Combining structured records and selective anomaly crops yields an illustrative 72-hour raw requirement of 8.57 GB; applying a 3.0 safety factor for system overhead and compaction yields an illustrative 25.71 GB allocation on a 128 GB or 256 GB industrial drive, provided raw video is excluded. [S33] [S0801]

As shown in **Exhibit 34**, p. 132, separating lightweight operational evidence from bulky diagnostic media is the decisive control in local buffer design. [S07] [S08] [S33] Structured events and meter telemetry for a three-day disconnection occupy less than two gigabytes under illustrative plant assumptions, whereas indiscriminate media buffering creates an artificial storage crisis. [S08] [S33]

EXHIBIT 34 / GIGABYTES (72-HOUR BUFFER)

Structured operational evidence requires modest local buffer capacity compared with raw media

Record stream	Illustrative hourly volume	Unit footprint (kB)	72-hour buffer (GB)	Queue exhaustion policy
Visibility and quality events	2,500	1.8	0.32	Never drop; halt only on hardware fault
Machine and utility telemetry	28,800	0.5	1.04	Downsample to 60s averages above 85% disk
KPI snapshots and audit logs	15	8	0.01	Never drop; retain in primary transaction store
Vision anomaly crops (selective)	400	250	7.2	Evict oldest unlabelled crops first above 75% disk
Unfiltered raw camera video	2	15m	2,160	Prohibited in continuity queue; local ring only

DERIVED [S07] [S08] [S33] Derived store-and-forward sizing model using event semantics from S08, transport persistence concepts from S07, and local parent-device buffering from S33. Event rates and payload sizes are explicitly illustrative plant engineering inputs (2,500 visibility events/h at 1.8 kB; 28,800 telemetry readings/h at 0.5 kB; 15 KPI/audit logs/h at 8 kB; 400 vision defect crops/h at 250 kB versus 30 GB/h raw video).

Exhibit 34, p. 132 cannot prove that a drive will survive ambient heat, power cycling or write wear on the floor. [S22] [S33] The custodian must pair capacity sizing with disk-health telemetry and eviction rules so diagnostic images are overwritten in a ring buffer before governed production or traceability events are threatened. [S08] [S22] [S33]

8.4

Local decisions continue when external services stop

Supervisors, inspectors and technicians must also make informed decisions on the floor while external connectivity is absent. [S22] [S0801] Industrial systems have strict performance, reliability and safety requirements that cannot depend on fragile external links. [S22] If an operator terminal requires a cloud call to validate a bundle ticket, read a bill of materials, or quarantine a defective fabric roll, a severed internet cable stalls production. [S05] [S22]

8.5 ESSENTIAL SHOP-FLOOR FUNCTIONS REMAIN RESIDENT AT THE PLANT BOUNDARY

Following IEC 62264-1 (ISA-95), work schedules, active production orders, product specifications, defect taxonomies and lot-genealogy capture belong inside the site boundary during execution. [S05] [S08] [S0803] Four functional categories must run locally when the wide-area network is disconnected: [S05] [S22] [S38]

- **Order and lot execution capture:** Scanning cotton or fabric roll receipts, splitting or merging bundles, logging piece counts and recording shift handovers must validate against a locally cached copy of active master data and write to the local store. [S05] [S08]
- **Quality disposition and quarantine:** Inspectors and vision stations must flag defects, place suspect lots into quarantine, and enforce hold rules locally so defective material does not advance to downstream stations during an outage. [S08] [S22]
- **Line visual management and hourly KPIs:** Shift supervisors need live calculation of line availability, first-pass yield and hourly output against target using the local KPI engine under ISO 22400-2, so shift reviews proceed without external servers. [S38]
- **Safety-adjacent and equipment-protection alerts:** Local threshold monitoring on motor temperatures, dye-bath pressures, boiler states and electrical loads must execute on local controllers and supervisory nodes, with protective interlocks wired in hardware or deterministic control logic rather than routed through data-spine software. [S22]

8.6 STALE REMOTE ANALYTICS DECLARE THEIR AGE RATHER THAN MASQUERADING AS LIVE

Functions that aggregate multi-site demand history, execute retraining runs, or query external supplier portals may pause when wide-area connectivity fails because they do not control minute-by-minute sewing or spinning execution. [S05] [S33] [S0803] When external services become unreachable, the danger is that a local screen displays yesterday's cached forecast or replenishment recommendation as if it were current. [S22] [S0802]

Every screen and operator panel must display the freshness state and watermark timestamp of its feed. [S08] [S22] When the wide-area link drops, any view relying on remote computation must switch to a *disconnected/cached* banner showing the age of the last synchronised state, while local views update in real time. [S33] [S38]

> **GFT NOTE**
>
> In GFT's Cameroon textile programme, factory sites face periodic telecommunications instability and utility fluctuations alongside constrained on-site IT staffing. [R06] [S22] GFT's engineering rule is that no operator action on the spinning, weaving, dyeing or garment floor may block on an external network request, and no cloud-hosted model may sit inside a real-time control loop. [S22] [S33] When a model is admitted for shop-floor assistance (such as optical fabric inspection), the inference runtime and approved model artefact execute locally on the plant edge node, while external links serve only for asynchronous telemetry export, backup replication and controlled downloads of approved model versions. [R29] [S15] [S33]

8.7

Recovery objectives vary by data class

Hardware failure, accidental deletion, faulty schema migrations and ransomware can corrupt or destroy local data stores. [S26] [S0802] NIST contingency planning guidance (SP 800-34 Rev. 1) requires evaluating systems and operations through a business-impact assessment to set contingency priorities, recovery point objectives (RPO, maximum tolerable data-loss interval) and recovery time objectives (RTO, maximum tolerable restoration time). [S0801] Applying one uniform RPO and RTO to every store either wastes budget on low-consequence archives or leaves critical production records underprotected. [S26] [S0801]

8.8 BUSINESS IMPACT SEPARATES CONTROL-ADJACENT RECORDS FROM ANALYTICAL ARCHIVES

A disciplined industrial data spine classifies its stores by operational consequence and reproducibility cost rather than by storage volume. [S15] [S0801] Four record classes require separate recovery targets: [S08] [S15] [S38] [S0801]

1. **Control-adjacent operational events and active lot genealogy:** These records track which fabric roll entered which dyeing vat or cutting table, which quality holds are active, and which customer order is on the line. [S05] [S08] Losing active genealogy breaks traceability and risks shipping quarantined goods, while a prolonged outage forces the floor onto paper tally sheets. [S08] [S0801] These records warrant tight local replication (such as write-ahead log shipping to a secondary on-site node) and rapid restoration. [S26] [S0801]
2. **Daily production, maintenance and quality shift logs:** Hourly piece counts, downtime reason codes and inspection tallies feed the daily and weekly management cadence. [S38] If an outage occurs mid-shift, operators and supervisors can reconstruct a brief interval from machine counters and tally sheets, provided restoration completes before the next daily review meeting. [S38] [S0801]
3. **Curated KPI views and management reports:** Because curated datasets and KPI tables are deterministic transformations of raw events and master data, they do not require near-zero RPO hardware replication if raw events and transformation code are intact; they can be recomputed from the restored event store within a few hours. [S09] [S15] [S38]
4. **Model registries, training snapshots and historical experiment logs:** Approved model artefacts, configuration files and dataset manifests needed for rollback or audit change only when a candidate is trained or promoted. [S15] Their RPO aligns with release events rather than seconds of clock time, though an immutable offline copy of the deployed model package must remain accessible so an edge node can be re-imaged quickly after hardware failure. [S15] [S26]

The bands in **Exhibit 35**, p. 135 are illustrative planning ranges showing how recovery targets scale with operational consequence, not universal benchmarks or measured GFT outcomes; each plant must set its own targets after a business-impact assessment and restoration tests. [S26] [S0801]

EXHIBIT 35 / ILLUSTRATIVE TARGET HOURS

Recovery objectives should reflect consequence rather than data volume



ILLUSTRATIVE [S26] [S0801] [S0802] Illustrative planning ranges for Recovery Point Objective (RPO, maximum data-loss interval) and Recovery Time Objective (RTO, maximum time to restore verified service) across four factory record classes. Method derived from contingency planning priorities in S0801, offline backup testing in S26, and data-integrity recovery in S0802. Values are explicitly illustrative design ranges, not benchmarks; final site targets must follow a factory business-impact assessment.

Exhibit 35, p. 135 shows why active lot genealogy requires tight bounds (an illustrative 0.08 to 1.0 hour RPO and 0.5 to 4.0 hour RTO) whereas historical archives tolerate an illustrative 8.0 to 72.0 hour restoration window. [S08] [S0801] However, **Exhibit 35**, p. 135 cannot prove that backups meet those targets until a technician restores an offline copy into an isolated sandbox and measures actual elapsed time and record completeness. [S26] [S0802]

8.9 RESTORATION TESTS VERIFY ACCURACY AND PRECISION RATHER THAN BACKUP JOB LOGS

Industrial teams often confuse backup completion with recovery capability. [S26] A completed backup script does not prove that the archive is uncorrupted, that encryption keys are accessible outside the primary domain controller, or that restored tables pass integrity checks. [S26] [S0802] Multi-agency guidance (CPG 2.R) requires offline, encrypted backups of critical data and regular testing of their availability and integrity in a disaster recovery scenario, while NIST SP 1800-11 stresses trusting the accuracy and precision of recovered data. [S26] [S0802]

To satisfy both requirements, the custodian must store at least one encrypted backup copy offline or on an immutable storage target isolated from operational credentials, and execute a scheduled restoration drill into a sandbox environment. [S25] [S26] The

acceptance metric is whether the restored tables contain the expected event counts up to the snapshot boundary, preserve cryptographic hashes on model artefacts, and reconcile against shift production totals. [S15] [S26] [S0802]

8.10

Open representations make exit practical

At the commercial layer, vendor lock-in, licence price increases, product discontinuation or service withdrawal threaten continuity just as severely as hardware failure. [S10] [S32] [S0803] NIST's cloud computing synopsis (SP 800-146) highlights data portability and lock-in risk when adopting external services. [S0803] If historical production records, quality labels and KPI definitions are trapped inside undocumented vendor tables, leaving that vendor resets institutional memory to zero. [S05] [S10] [S0803]

8.11 SEMANTIC STANDARDS PREVENT PROPRIETARY LOCK-IN DISGUISED AS OPEN SYNTAX

Exporting CSV or JSON files without semantics leaves a trap of opaque column names and undocumented status codes. [S05] [S06] [S10] True portability requires that exported records conform to open semantic models across three layers: [S05] [S06] [S08]

- **Enterprise-control boundary semantics (IEC 62264-1 / ISA-95):** Production schedules, work performance actuals, material lot definitions and equipment hierarchies map to the standard terminology and object models of IEC 62264-1 so moving between manufacturing execution or enterprise systems does not force a redesign of core domain entities. [S05]
- **Machine and device information modelling (OPC UA):** Where equipment controllers and gateways expose structured process variables, alarms and asset metadata, using OPC Unified Architecture (OPC 10000-1) information models ensures that tag relationships, units and security contexts are self-describing rather than buried in a vendor PLC tool. [S06]
- **Visibility event semantics (GS1 EPCIS 2.0):** Representing lot receipts, bundle splits, inspections and shipments through EPCIS 2.0 event dimensions (what object, when, where, why, and under which business step and disposition) allows any standards-compliant traceability or analytics store to ingest the historical event stream without custom guesswork. [S08]

8.12 COMMERCIAL CONTRACTS MUST SPECIFY COMPLETE EXPORT AND EXIT VERIFICATION

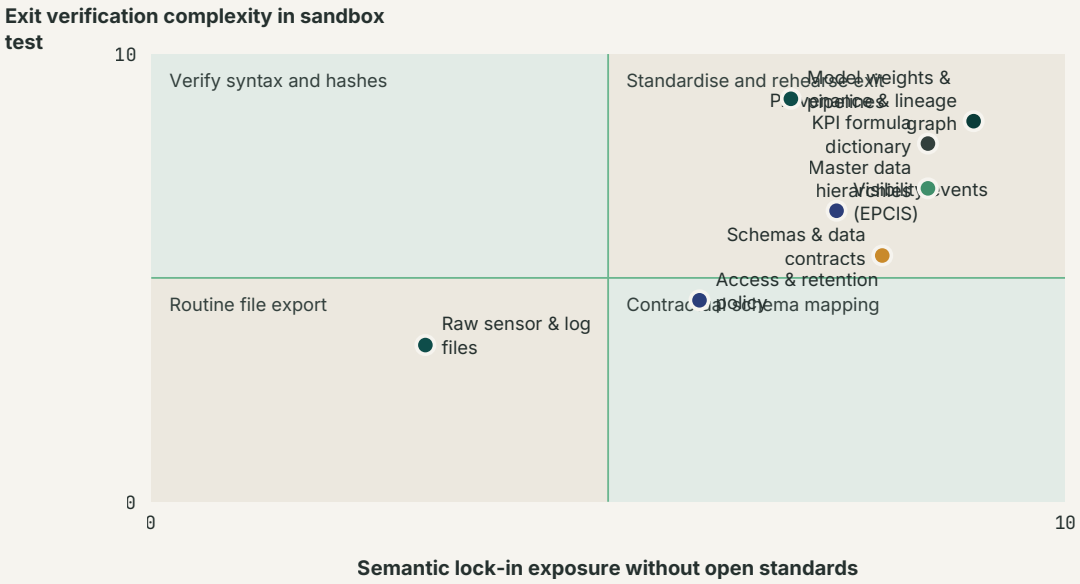
Commercial terms govern whether a factory can extract its data when a contract ends. [S32] Where the European Union Data Act (Regulation (EU) 2023/2854) applies to connected products or related services in the Union, users can access generated data in a timely manner and share them with chosen third parties. [S32] Because a manufacturer in Cameroon cannot assume European statutory rights apply extraterritorially to local procurement (and this paper is not legal advice), territorial scope requires review by qualified counsel. [S32]

GFT's procurement rule writes explicit data-access, export format, metadata and exit clauses into every equipment and software contract before signature, requiring full self-service export of raw records, curated products, schemas, lineage and models during normal operation. [S09] [S10] [S15] [S32] [S0803]

Exhibit 36, p. 137 maps the eight artefact classes required for a platform exit across two dimensions: semantic dependency risk if left in a proprietary format, and verification rigor required to prove that an exported copy works outside the originating vendor's environment. [S05] [S06] [S08] [S09] [S15] [S32]

EXHIBIT 36 / DERIVED INDEX (0 TO 10)

Portability requires meaning, evidence and operating instructions



DERIVED [S05] [S06] [S08] [S09] [S15] [S32] Derived portability assessment across eight required export artefact classes, combining enterprise-control models (S05), OPC UA information models (S06), EPCIS event semantics (S08), PROV-O lineage (S09), MLOps pipeline artefacts (S15) and connected-product data access principles (S32). X-axis scores semantic lock-in exposure if left unstandardised (0=low, 10=high); Y-axis scores independent restoration test complexity required before accepting an export routine (0=low, 10=high).

Exhibit 36, p. 137 places raw log files in the lower-left quadrant because bulk files are simple to copy, while master data, visibility events, KPI dictionaries, provenance graphs and model pipelines sit in the upper-right quadrant (`Standardise and rehearse exit`). [S05] [S08] [S09] [S15] Because **Exhibit 36**, p. 137 cannot replace a live drill, an export routine is proven only when imported into an independent environment to reproduce operational reports without vendor APIs. [S27] [S38] [S0803]

8.13

Recovery includes the instructions that make data intelligible

Database recoveries often succeed technically while failing operationally because the backup captured only transactional rows and omitted the interpretive instructions that turn those rows into governed evidence. [S09] [S15] [S38] A raw counter reading of 480 from a sewing line means nothing unless the recovery environment also restores the

master-data version identifying the line and style, the schema version defining the unit as stitched collars rather than bundles, the quality-gate rule filtering out test runs, and the ISO 22400-2 KPI formula version specifying how rework and planned breaks were treated on that date. [S05] [S09] [S38]

8.14 RAW EVENTS WITHOUT SCHEMAS, MASTER DATA AND KPI FORMULAS ARE UNINTELLIGIBLE BYTES

Provenance standards (W3C PROV-O) and MLOps pipeline-lineage guidance emphasise that reconstructing an outcome requires recording the entities, activities, agents, parameters and execution artefacts that produced it. [S09] [S15] Under CISA's offline-backup principle and NIST's data-integrity framework, the custodian must treat six instruction sets as first-class recovery payloads alongside the raw event store: [S09] [S15] [S25] [S26] [S38] [S0802]

1. **Schema and data-contract registry:** Historical versions of JSON schemas, table definitions, unit specifications and validation rules used to gate ingestion. [S08] [S09]
2. **Temporal master-data snapshots:** Versioned equipment trees, line configurations, bills of material, shift calendars and defect-code taxonomies with effective date intervals. [S05] [S06]
3. **KPI dictionary and calculation code:** Scripts implementing each version of the factory's ISO 22400-2 indicators, including numerator, denominator, grain and exclusion rules. [S15] [S38]
4. **Provenance and lineage logs:** Execution manifests linking raw event batches to cleaned tables, feature snapshots, model runs and published management views. [S09] [S15]
5. **Model registry and serving packages:** Immutable container images or model weights, feature-transformation code, human-set decision thresholds, and rollback pointers for deployed and suspended models. [S15]
6. **Role-based access and retention policies:** Least-privilege permission matrices and processing-register constraints under NIST SP 800-53 so a restored environment does not expose restricted records during an emergency rebuild. [S25]

8.15 HISTORICAL FIGURE REPRODUCTION IS THE ACCEPTANCE TEST FOR COMPLETE RECOVERY

To verify that all six instruction sets are captured and reassembled coherently, GFT requires a deterministic *historical figure reproduction test* in every quarterly recovery drill. [S15] [S26] [S38] The recovery team selects a closed historical reporting period and instructs the restored sandbox environment to recompute the shift's Overall Equipment Effectiveness (OEE), first-pass quality yield and active model predictions from the restored raw events and historical instruction versions. [S09] [S15] [S38]

HISTORICAL KPI AND MODEL-STATE RECONSTRUCTION CHECK

- **System boundary:** Isolated recovery sandbox with zero network links to live production databases or vendor APIs, populated solely from offline encrypted backup archives and immutable repositories. [S15] [S26]
- **Required inputs:** (1) Offline backup of raw visibility and telemetry events for period T; (2) archived master-data snapshot effective at T; (3) archived schema and quality-rule version effective at T; (4) archived KPI formula version v_k effective at T; (5) archived model package and threshold configuration m_j active at T; (6) signed-off management report and model audit log from period T. [S05] [S09] [S15] [S38]
- **Execution steps:** Mount the restored raw event store; apply schema validation and quarantine filters using version T rules; join with master-data snapshot T; execute KPI transformation v_k and model inference batch m_j; compare regenerated outputs cell-by-cell against the historical signed-off report. [S09] [S15] [S38]
- **Verification gate:** Recomputed KPI figures must match historical reported values within rounding tolerance (0.01 percent), quarantine counts must match, and restored model outputs on the test slice must match logged production predictions. Any discrepancy reveals an unbacked-up manual override, an unversioned master-data mutation, or a missing configuration dependency that must be fixed before certification. [S15] [S38] [S0802]

8.16

Migration preserves history from the manual system

Continuity across technology generations starts when the plant migrates from its manual spreadsheet logs to a structured database. [S05] [S09] In GFT's Cameroon textile programme, operations begin with a manual data spine of shift logs, defect tally sheets, lot cards and a governed KPI workbook. [R06] [S38] When transaction volume justifies a database or operations platform, treating migration as a clean-slate replacement destroys the baseline history needed to evaluate process improvements and seasonal demand. [S09] [S15] [S38]

8.17 THE INITIAL SPREADSHEET REGISTER SERVES AS A FORMAL MIGRATION SPECIFICATION

A governed spreadsheet system of record is the living specification of the factory's domain identifiers, state definitions and calculation logic. [S05] [S38] If the manual system enforces stable identifiers for lines, assets, styles, lots and defect codes, migrating to a database is a change of storage engine rather than business meaning. [S05] [S08] If an integrator introduces a new platform with incompatible asset codes, renumbered defect categories and altered downtime buckets without a mapping table, historical records collected prior to cutover become unusable for trend analysis or feature engineering. [S08] [S09] [S18]

To preserve continuity during migration, the data owner and steward require four artefacts before cutover approval: [S05] [S08] [S09] [S38]

- **Bijective identifier crosswalk:** A versioned mapping table linking every legacy spreadsheet identifier (line, machine, style, supplier lot, defect code) to its new system identifier, retaining the legacy code as an immutable attribute. [S05] [S08]
- **Historical event backfill with provenance:** An ingestion script converting historical spreadsheet rows into immutable records in the new store, tagged with a provenance activity (`migrated_from_manual_spine_v1`), original capture date, migration date and responsible steward under PROV-O semantics. [S09]
- **Explicit quality-state preservation:** Carrying missing-shift, estimated or corrected quality states into the new database rather than imputing them as clean sensor observations or replacing them with zeros. [S08] [S09]
- **Read-only archival of the source workbooks:** Cryptographically hashed, read-only copies of the cutover spreadsheets stored in the offline backup vault so an auditor or engineer can verify the pre-migration source text. [S26] [S0802]

8.18 PARALLEL RECONCILIATION PROVES SEMANTIC EQUIVALENCE BEFORE CUTOVER

Subtle differences in how new software handles time-zone boundaries, shift changeovers, partial bundles or planned breaks can shift calculated KPIs. [S38] If line availability jumps from an illustrative 72 percent to 79 percent on cutover morning solely because the new software excludes meal breaks from scheduled time while the spreadsheet included them, management cannot tell whether performance improved or the ruler changed. [S38]

Migration cutover therefore requires a mandatory parallel run across an agreed operational window (for example, an illustrative two to four weeks of shifts) during which both the legacy workbook and the new platform compute the factory's ISO 22400-2 indicators from the same shop-floor events. [S38] Any divergence must be diagnosed as either an implementation bug in the new software (fixed before cutover) or an approved upgrade to the KPI definition (requiring a new KPI version number and backward recalculation of the historical series). [S09] [S38] Only when parallel reconciliation closes within tolerance does the new platform become the system of record. [S38] [S0802]

8.19

Deliberate failure tests expose hidden dependence

Architectural diagrams and written contingency plans assume fallback paths work as drawn. [S27] [S0801] Hidden dependencies surface when a component fails: an edge gateway configured for offline operation relies on an external Domain Name System (DNS) server to resolve local hostnames; a local MQTT broker rejects messages because a TLS client certificate expired during an outage; or a fallback spreadsheet sits on a shared network drive unreachable when the main switch loses power. [S07] [S22] [S33] NIST's incident response guidance (SP 800-61 Rev. 3) emphasises testing and exercising procedures periodically to verify accuracy and train personnel. [S27]

8.20 SIX FAULT-INJECTION SCENARIOS TEST THE SPINE BEFORE A REAL OUTAGE OCCURS

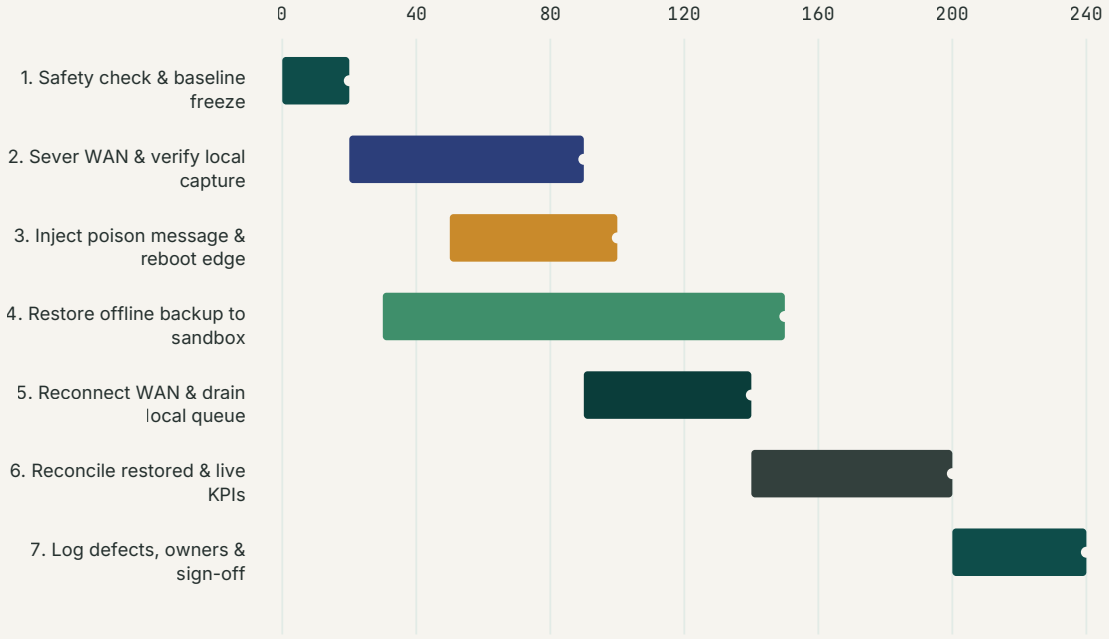
Rather than waiting for an unplanned blackout, fibre cut or vendor dispute, the custodian and stewards should run scheduled failure exercises during maintenance windows or controlled shifts. [S22] [S27] [S0801] Six fault scenarios test the continuity and portability boundary of a small factory's data spine: [S07] [S26] [S27] [S32] [S33]

1. **Wide-area network severance:** Disconnect the external internet router for an illustrative 4-hour shift window. Verify that local capture, visual management boards, quality holds and edge model inference continue without latency spikes, remote-dependent screens display stale-data watermarks, and reconnection drains the buffer in priority order without duplicates. [S07] [S08] [S33]
2. **Unplanned edge-node power cycle:** Force an abrupt restart on the local ingestion gateway mid-stream. Verify that the write-ahead log and persistent message store recover cleanly on boot, committed events are preserved, and child devices reconnect automatically. [S07] [S22] [S33]
3. **Buffer saturation and high-water eviction:** Restrict the buffer partition or inject synthetic diagnostic payloads until disk usage crosses the illustrative 75 percent and 85 percent thresholds. Confirm that the ring-buffer policy evicts old unlabelled image crops or downsamples high-frequency telemetry while preserving 100 percent of governed production, quality and traceability events. [S08] [S33]
4. **Expired credential or certificate fault:** Deploy a short-lived test certificate expiring during an active capture session. Confirm that expiry raises a pre-expiry alert, fails safely without crashing control-adjacent software, and has a documented local rotation runbook independent of cloud authentication. [S22] [S25] [S27]
5. **Poison-pill corrupted message injection:** Publish a malformed JSON payload or invalid schema version into the ingestion queue. Verify that the ingestion gate quarantines the bad record into the dead-letter store with provenance metadata rather than stalling the queue. [S07] [S08] [S09]
6. **Clean-room backup restoration and vendor export import:** Take the latest offline encrypted backup and vendor self-service export archive into an isolated sandbox container. Execute the historical figure reproduction method to confirm that a technician following the runbook can restore the data, schemas, master data, KPI dictionary and model artefacts within the target RTO and reproduce shift figures accurately. [S15] [S26] [S32] [S38] [S0802]

Exhibit 37, p. 142 lays out an illustrative 4-hour (240-minute) integrated continuity and restoration exercise that a small factory team can run within a single half-day maintenance and operational window, combining live disconnection testing on the floor with parallel clean-room backup restoration. [S26] [S27] [S0801]

EXHIBIT 37 / ELAPSED EXERCISE MINUTES

A recovery exercise proves the whole chain, not only the backup



ILLUSTRATIVE [S26] [S27] [S0801] Illustrative 240-minute factory continuity and recovery drill sequence derived from contingency testing guidance in S0801, periodic incident-response exercises in S27, and offline backup restoration testing in S26. Timings are explicitly illustrative planning durations for a half-day drill, not measured GFT operational benchmarks.

8.21 MEASURED RECOVERY GAPS DRIVE ENGINEERING REMEDIATION BEFORE MODEL ADMISSION

As **Exhibit 37**, p. 142 shows, a half-day exercise tests both tracks simultaneously: while the shop-floor steward verifies that disconnection, poison-message quarantine, edge restart and reconnection draining work cleanly on the live buffer (Tasks 2, 3 and 5), the technical custodian restores the offline encrypted backup into an isolated sandbox (Task 4), and both roles converge to reconcile historical and live KPIs before signing off corrective actions (Tasks 6 and 7). [S26] [S27] [S38] What **Exhibit 37**, p. 142 cannot guarantee is that the first drill will pass without friction; the value of a scheduled exercise is to surface missing passwords, undocumented scripts, slow restoration steps and silent duplicate events under calm daytime conditions rather than during a real emergency. [S27] [S0801]

UNEXERCISED EXIT AND RECOVERY PATHS ARE LATENT SINGLE POINTS OF FAILURE

A factory that records backups without testing sandbox restoration, or that signs data-export clauses without importing a test archive into an independent environment, carries an unmeasured continuity risk. [S26] [S32] [S0802] Common failure modes discovered only during live crises include backup encryption keys stored exclusively inside the primary server being recovered, cloud export APIs that throttle bulk historical downloads, edge brokers whose offline queues overwrite the oldest production events when diagnostic logs fill the disk, and restored datasets that fail to reproduce historical OEE because master-data changes were overwritten in place rather than versioned. [S05] [S15] [S33] [S38] This constraint is resolved only by making a passed restoration and export drill a mandatory gate before any production system is declared operational or any machine-learning use case is admitted to the floor. [S15] [S26] [S27]

Every continuity exercise concludes with a governance record logging the scenario tested, measured RPO and RTO, duplicate or quarantined counts during reconnection, historical KPI reproduction accuracy, and defect owners and due dates. [S26] [S27] [S38] [S0801] When a factory can disconnect its router for a shift without losing a scan, restore its store and KPI dictionary from an offline vault in an afternoon, and import its history into an open schema without vendor help, its spine is ready to support production modelling. [S15] [S22] [S26] [S33]

Continuity is proven not while the network is up, but when the link fails, the backup restores and the history survives a change of vendor.

CHAPTER SOURCE NOTES

S22 National Institute of Standards and Technology. 2023.

S33 Microsoft. 2024-06-10.

R06 Cameroon Trade Hub. Publication date unrecorded; inherited source-text check.

R07 Ecofin Agency. Publication date unrecorded; inherited source-text check.

R12 Business in Cameroon. Publication date unrecorded; inherited source-text check.

S05 International Electrotechnical Commission. 2013.

S26 CISA, MS-ISAC, NSA and FBI. 2023.

S27 National Institute of Standards and Technology. 2025.

S0801 National Institute of Standards and Technology. 2010-11-11.

S08 GS1 AISBL. 2022.

S07 OASIS Open. 2019-03-07.

S10 National Institute of Standards and Technology. 2016-02-23.

S09 World Wide Web Consortium. 2013-04-30.

R29 MDPI; industry inspection literature. Publication date unrecorded; inherited source-text check.

S06 OPC Foundation. 2025-10-22.

S38 International Organization for Standardization. 2014.

S0803 National Institute of Standards and Technology. 2012-05-29.

S0802 National Institute of Standards and Technology. 2020-09.

S15 Google Cloud. 2024-08-28.

S25 National Institute of Standards and Technology. 2020-12-10.

S32 European Parliament and Council of the European Union. 2023-12-13.

S18 Advances in Neural Information Processing Systems. 2015.

GREEN FUTURE TECH

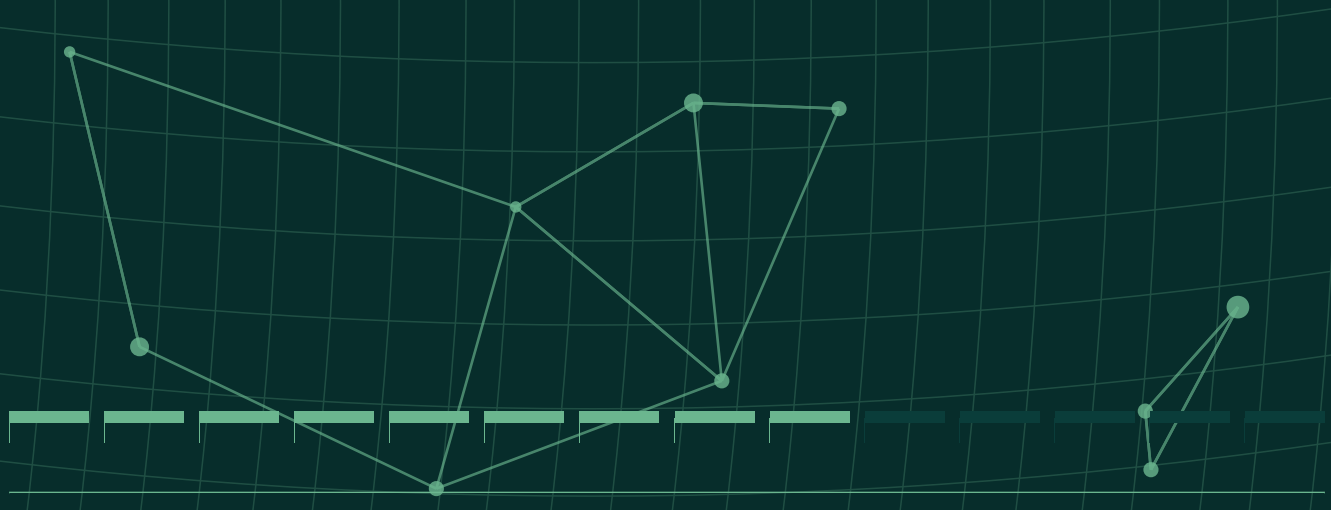
LIMITED / HONG KONG SAR

09

PART III: CONTROL MODELS AS PRODUCTION CHANGE / MODEL ADMISSION

A use case earns admission only when a decision and baseline exist

The first model gate is operational and economic: nobody should train before the action, baseline, data rights and stop conditions are explicit.



CHAPTER 09
DECISION TEST

The first model gate is operational and economic: nobody should train before the action, baseline, data rights and stop conditions are explicit.

A machine-learning candidate earns admission to development only when the factory can name the operational decision it will change, the role authorised to act on its output, the frozen baseline it must beat, the lawful and governed records it will consume, and the condition under which the work will stop. [S43] [S44] In a small or mid-sized manufacturer, the scarcest resource is not algorithmic variety; it is the engineering, stewardship and supervisory attention required to turn raw production events into dependable routines. [S02] [S18] Under Green Future Tech Limited's (GFT's) Cameroon textile programme, prepared on 31 January 2024 with source-refresh notes retained in 2026, inherited sector evidence records roughly 95 percent of cotton fibre exported raw and about 5 percent processed locally, against a national target of 50 percent local processing by 2030 (these inherited entries were not refreshed in 2026 and their original publication dates are unrecorded). [R06] [R07] Moving a spinning, weaving, dyeing or garment plant along that path demands strict capital discipline. Admitting an ungrounded model diverts stewards from capture quality, creates maintenance debt, and risks automating an unverified assumption. [S18] [S19]

Part I established how a factory governs operational meaning, decision rights and quality gates. Part II defined the minimum durable data spine that preserves time, identity, state, security and offline continuity. Part III treats machine learning as a controlled production change. Before reproducible training (wp14 Chapter 10), staged release (wp14 Chapter 11) or drift monitoring (wp14 Chapter 12) apply, a candidate must pass an operational and economic admission gate. [S13] [S43] This chapter specifies that gate, shows how lifecycle cost and simple comparators filter out speculative proposals, legitimises stopping or postponing candidates whose prerequisites fail, and explains why GFT admits bought fabric-inspection vision (wp03) ahead of history-dependent demand forecasting (wp04). [S43] [S44]

9.1

Start with the decision and the person who acts

Model admission begins with the physical or commercial decision that the factory intends to improve, never with an available dataset or a vendor demonstration. [S43] In the Cross-Industry Standard Process for Data Mining (CRISP-DM) lifecycle applied to manufacturing, business understanding is the opening phase and deployment is the closing phase that makes analytical work operationally real. [S43] ISO/IEC 5338 similarly frames artificial-intelligence lifecycle processes around defined system purpose, stakeholder requirements and operating context before technical build begins. [S13] Reversing that order creates an orphan model with no governed place in daily or weekly factory management. [S18] [S43]

9.2 AN OUTPUT WITHOUT AN AGREED ACTION IS A VANITY SCORE

Every candidate proposal must state five operational facts before any training extract is approved: the specific decision to be taken, the named role accountable for taking it, the physical or commercial action that follows, the decision horizon, and the escalation route when the signal is contested. [S02] [S43] Consider a model that predicts a sewing-line bottleneck four hours ahead. If the line supervisor has neither the authority to rebalance operators across styles nor a buffer of pre-cut bundles for an alternative operation, the score cannot alter production. [S43] Under GFT's admission

rule, a candidate whose sole output is an interesting risk score, heat map or anomaly index with no pre-agreed operational response fails the first gate immediately. [S43] [S44]

9.3 DECISION TIMING BOUNDS THE TECHNICAL SPECIFICATION

Defining the action and actor first fixes the technical constraints that the data spine must satisfy. [S08] [S43] A cloth-inspection alert on a rolling frame requires a decision within seconds so the operator can mark the fault, log the defect code and cut or downgrade the roll before dyeing (wp03). [R29] [S08] By contrast, a yarn-purchasing or style-allocation forecast operates on a weekly or monthly planning cycle aligned with sales orders and ginning deliveries (wp04). [S43] The first decision requires low-latency edge inference that survives a wide-area network outage; the second tolerates batch execution but demands reconciled order, inventory and lead-time history across enterprise and operations boundaries. [S08] [S22] Naming the decision owner and horizon at admission prevents building streaming pipelines for a monthly planning meeting, or overnight batch reports for a line fault already sewn into three hundred garments by morning, an illustrative timing mismatch. [S02] [S43]

9.4

A baseline makes improvement falsifiable

No model can prove its worth unless the factory first records how the decision performs without it. [S36] [S0902] Practical machine-learning engineering guidance states the rule directly: before formalising what a machine-learning system will do, track as much as possible in the current system. [S0902] World Bank evidence on micro, small and medium enterprise digitalisation reinforces the same requirement, showing that embedding monitoring and evaluation throughout design and implementation is essential for tracking impact and learning what works. [S36] Collecting defect rates, downtime minutes or forecast errors only from the day a model launches blends measurement visibility, supervisory attention and the model itself. [S36] [S38]

9.5 FREEZE THE KPI CONTRACT BEFORE DEVELOPMENT STARTS

To make claimed improvement falsifiable, the domain owner and data steward must freeze the governing key performance indicator (KPI) definition in the KPI dictionary before candidate development begins. [S02] [S38] Following ISO 22400-2, that frozen record specifies the indicator's exact formula, input elements, unit, time behaviour, grain and exclusions. [S38] For a fabric-inspection candidate, the baseline states whether defect catch accuracy is measured per hundred linear metres or per roll, whether minor slubs within buyer tolerances are excluded, and how manual re-inspection confirms missed faults. [R29] [S38] For a demand-forecasting candidate, the baseline locks the forecast horizon, the style-colour grain, and whether cancelled orders or stockouts are adjusted when computing actual demand. [S38] [S0901] Changing the KPI formula mid-pilot invalidates the comparison and forces the candidate back to baseline collection. [S36] [S38]

9.6 ISOLATE OPERATIONAL CONFOUNDERS IN THE BASELINE WINDOW

A factory baseline is a time-stamped operating record that captures the conditions under which the pre-model process ran. [S08] [S38] In a Cameroonian cotton-to-garment operation, fibre characteristics shift across ginning lots, ambient humidity

changes between seasons, grid power fluctuates, and operator proficiency varies across shifts. [R03] [S38] Evaluating a candidate during a four-week run of plain-weave sheeting against a baseline window dominated by dyed twills confuses product mix with model skill, an illustrative confounding case. [S36] [S38] The admission dossier must specify the baseline window, the stratification variables (shift, line, fabric construction, cotton lot and order class) from the operational vocabulary (wp14 Chapter 03), and the counterfactual check separating model contribution from external shifts. [S08] [S36] [S38]

9.7

Readiness tests data, rights and operating capacity

Once a candidate has a named decision owner and a frozen baseline, the admission review tests whether the factory possesses the data quality, legal rights and human capacity required to build and sustain it. [S04] [S29] [S44] Practitioners in high-stakes artificial intelligence report that upstream data cascades, triggered by neglected capture quality and improvised labelling, are pervasive (92 percent prevalence in the cited study), invisible at first, and costly downstream. [S19] Admitting a model onto unvalidated records accelerates the propagation of capture errors into operational decisions. [S04] [S19]

9.8 DISTINGUISH POOR HISTORY FROM INSUFFICIENT HISTORY

ISO/IEC 5259-3 defines requirements and guidance for managing the quality of data used in analytics and machine learning. [S04] At the admission gate, the data steward and model builder apply those requirements to separate *corrupted history* from *insufficient history*. [S04] [S09] Corrupted history occurs when events were logged for months, but identifiers were reused without versioning, missing meter readings were stored as zero output, or defect labels lacked a shared taxonomy (wp14 Chapters 03 and 04). [S04] [S08] Corrupted history requires upstream capture remediation before records enter a training snapshot. [S04] [S19]

Insufficient history occurs when the capture spine is clean, but the factory has not yet operated through enough order cycles, seasonal transitions or machine faults to populate a representative training distribution. [S04] [S43] A pre-trained computer-vision model for weave defects (wp03) does not require three years of local sales tables; it requires a curated, well-lit local validation set of fabric rolls with verified defect coordinates to test transfer fitness and calibrate thresholds. [R29] [S04] Conversely, a demand forecaster (wp04) cannot substitute vendor pre-training for missing local order, lead-time and stockout history. [S04] [S0901]

9.9 VERIFY LAWFUL ACCESS, WORKER-DATA LIMITS AND EQUIPMENT RIGHTS

Technical data availability does not equal lawful permission to model. [S29] [S31] Every admission review checks the processing register (wp14 Chapter 07) against three boundaries:

- **Personal-data minimisation:** Under Recital 39 of the European Union General Data Protection Regulation (GDPR), personal data must be adequate, relevant and limited to what is necessary for the stated purpose. [S29] Ingesting operator badge scans or workstation video when line or bundle aggregates suffice is rejected. [S29] [S30]
- **Cameroonian statutory authorisation:** Under Section 19(1) of Law No. 2024/017 of 23 December 2024 in Cameroon, processing personal data is subject to prior authorisation by the Personal Data Protection Authority. [S31] Because implementing practice requires current local legal verification (this paper is not legal advice), no candidate touching personal data passes admission until local counsel confirms authorisation status. [S31]
- **Connected-equipment data rights:** When a candidate relies on sensor streams from spinning frames, looms or stenters, the custodian must confirm contractual rights to extract and reuse those records, applying conditional access principles where the European Union Data Act applies and explicit contract clauses elsewhere. [S32]

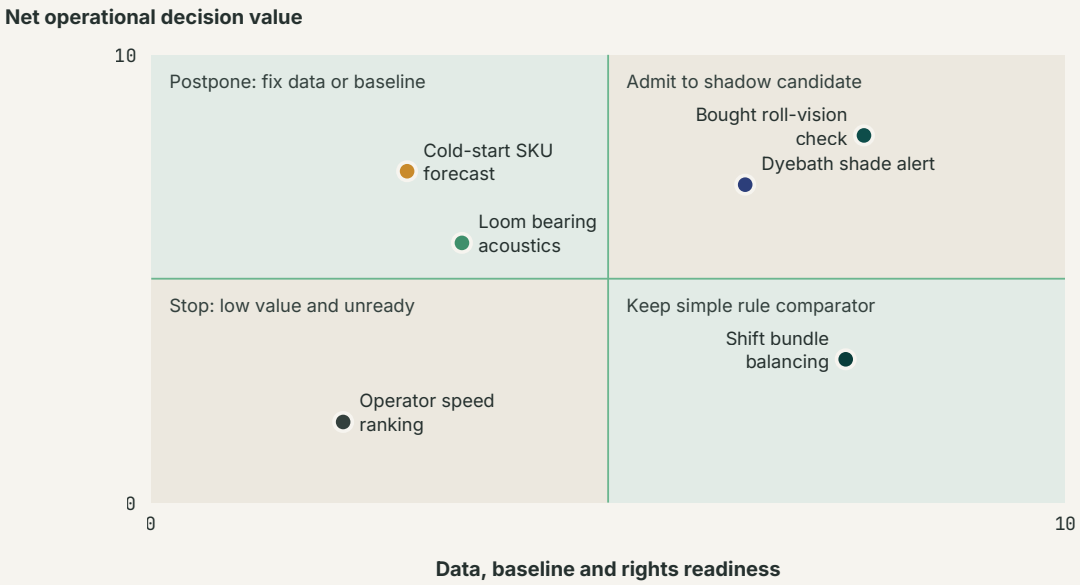
9.10 CONFIRM STEWARDSHIP AND CUSTODY CAPACITY

A model is an ongoing operating commitment, not a one-off script. [S16] [S18] Industry 4.0 MLOps research shows that starting technical infrastructure and target automation levels strongly influence complexity. [S42] Before admitting a candidate, the internal champion and domain owner must confirm who will maintain features, review quarantined inputs each shift, and hold custody of the registry and rollback routine. [S02] [S37] [S42] Admitting a second model when the sole analyst is fully committed to core production and energy logs guarantees decay. [S02] [S18] [S44]

Exhibit 38, p. 150 places six illustrative candidates on a two-dimensional admission matrix comparing data, baseline and rights readiness against net operational decision value.

EXHIBIT 38 / ILLUSTRATIVE SCORE (0 TO 10)

An admission scorecard can reject a technically attractive use case



ILLUSTRATIVE [S04] [S11] [S14] [S36] [S44] Illustrative admission scoring of six candidate factory use cases across data, baseline and rights readiness (x-axis, 0 to 10) and net operational decision value after lifecycle cost (y-axis, 0 to 10). Evaluation dimensions derive from S04, S11, S14, S36 and S44; coordinates are illustrative method examples, not measured factory results.

As **Exhibit 38**, p. 150 illustrates, cold-start stock-keeping-unit (SKU) demand forecasting falls into the upper-left *Postpone* quadrant when order history and stockout records are immature (x = 2.8, y = 7.4 illustrative). [S04] [S44] Shift bundle balancing has clean data (x = 7.6 illustrative) but low incremental model value (y = 3.2 illustrative) because a work-in-progress rule already solves the supervisor's problem, placing it in *Keep simple rule comparator*. [S43] [S0902] Individual operator speed ranking scores poorly on both axes (x = 2.1, y = 1.8 illustrative) and fails legal and worker-dignity screening. [S29] [S30] Only upper-right candidates clear the human-set thresholds to advance, and a visual score never replaces a signed KPI baseline and legal check. [S11] [S31] [S38]

9.11

The simplest useful comparator sets the bar

A recurring mistake in industrial analytics is comparing a complex model against nothing, or against a weak straw man. [S43] [S0901] Complexity is an operating liability that must earn its keep. [S16] [S18] Every admitted candidate must be tested against two mandatory baselines during evaluation: *current operating practice* in the frozen baseline window, and the *simplest useful rule or statistical comparator* built on the same governed dataset. [S43] [S0901] [S0902]

9.12 SIMPLE HEURISTICS AND BENCHMARK METHODS SOLVE A LARGE SHARE OF THE PROBLEM

Engineering guidance for machine-learning systems notes that basic heuristics should launch first when data are scarce; a simple heuristic captures a large fraction of the attainable benefit while establishing the measurement pipeline, whereas a

premature complex model is difficult to maintain. [S0902] Forecasting theory enforces the same rule through four standard benchmarks: the historical mean, the naive method (setting future forecasts to the last observation), the seasonal naive method (setting each forecast to the last observation from the same season of the prior cycle), and the drift method. [S0901] As Hyndman and Athanasopoulos state, any new forecasting method must be compared against these simple benchmarks to ensure that it beats them; if not, the new method is not worth considering. [S0901]

In a factory context, simple comparators are transparent and low-cost: a seasonal naive or moving-average model with confirmed contract backlog for demand planning (wp04); fixed run-hour or metre-count preventive intervals from manuals and stoppage logs for maintenance (wp02 and wp15); and statistical process control (SPC) limits or four-point manual perch sampling for quality control (wp03 and wp06). [R29] [S05] [S38] [S43] [S0901]

9.13 COMPLEXITY MUST CLEAR A NET-VALUE HURDLE, NOT A DECIMAL-POINT METRIC GAIN

Suppose a gradient-boosted tree improves mean absolute scaled error by 2 percent over a seasonal naive comparator in an offline holdout test, an illustrative result. [S0901] In a small factory, that 2 percent statistical gain may produce zero change in yarn purchase orders because supplier minimum order quantities and container increments are coarse. [S43] Meanwhile, the complex model introduces feature-pipeline fragility and maintenance costs that the seasonal naive formula avoids. [S18] [S42] Under GFT's admission standard, the domain owner and model owner set a minimum operational superiority margin before training begins. [S11] If the candidate cannot beat the simple comparator by enough to change physical decisions and recover its extra lifecycle cost, the factory deploys the simple comparator and closes the model candidate. [S44] [S0901]

9.14

Risk tier determines scrutiny and authority

Not every factory model carries the same consequence when wrong. [S11] [S14] NIST's Artificial Intelligence Risk Management Framework (AI RMF 1.0) stresses that human judgement must be employed when choosing trustworthiness metrics and setting their precise threshold values in light of consequence. [S11] ISO/IEC 23894 similarly guides organisations to integrate AI risk management into ordinary operating and risk functions. [S14] At admission, GFT classifies every candidate into one of four consequence tiers (wp14 Chapters 02 and 11). [S11] [S14]

9.15 TIER 1 AND TIER 2 COVER PLANNING AND QUALITY DISPOSITION

Tier 1 (Advisory planning): Models generating batch demand projections or maintenance lists that a planner reviews during normal hours with full access to source records. [S11] [S14] If a Tier 1 model stalls, the planner reverts to the spreadsheet or seasonal naive comparator with no threat to safety, product integrity or worker status. [S16] [S0901] Admission requires domain-owner and steward sign-off, a frozen KPI baseline, and a documented manual fallback. [S02] [S38]

Tier 2 (Quality and customer-affecting disposition): Models influencing roll grading, cutting, shade-lot grouping, or buyer carbon and compliance disclosures (wp03, wp05 and wp16). [S04] [S14] A false negative allows defective cloth into dyeing or cutting, multiplying scrap cost; excessive false positives halt inspection frames and breed alert fatigue. [R29] [S19] Admission requires independent challenge by the quality lead and model owner (wp14 Chapter 02), explicit error-cost weights, a stratified validation plan across styles and lighting, and a mandatory shadow-mode trial. [S04] [S11] [S16]

9.16 TIER 3 AND TIER 4 GOVERN WORKER-AFFECTING AND CONTROL-CONNECTED USES

Tier 3 (Worker-affecting analytics): Any model touching identifiable workers, skill assessments, ergonomics, shift allocation, pay incentives or discipline. [S29] [S30] Principle 5.5 of the International Labour Office (ILO) code of practice on workers' personal data states that decisions concerning a worker should not be based solely on automated processing of that worker's personal data. [S30] While the ILO code is non-binding guidance, GFT adopts Principle 5.5 as a binding internal rule alongside GDPR data minimisation and Cameroon's statutory prior-authorisation requirement in Law No. 2024/017. [S29] [S30] [S31] Ranking individual workers automatically for pay penalties or termination is prohibited at admission. [S30] Legitimate ergonomic or training analytics require role-aggregated inputs where feasible, documented legal review, and substantive human review where a supervisor has time, context and authority to reject any output. [S29] [S30] [S31]

Tier 4 (Control-connected OT automation): Any candidate proposed to write set-points, adjust line speeds, or interface directly with programmable logic controllers (PLCs), drives, boilers or dyehouse valves. [S22] NIST Special Publication 800-82 Revision 3 underscores that operational technology (OT) systems have strict performance, reliability and safety requirements where latency or unintended state changes can damage equipment or injure staff. [S22] In a small factory building its initial data spine, GFT rejects closed-loop Tier 4 model actuation at admission and redesigns the candidate as an open-loop Tier 2 advisory display decoupled from the control network. [S16] [S22] Any future control-connected candidate requires a formal process-hazard and OT safety review, independent hardwired interlocks, and joint sign-off by the factory director, OT custodian and safety officer. [S22] [S23]

9.17

Cost includes the whole operating lifecycle

Vendor brochures and macro-studies routinely present artificial intelligence through headline benefit percentages while omitting the local engineering effort required to prepare data and keep models alive. [S18] [S39] In the inherited evidence register, McKinsey's *The State of Fashion 2025* reports sector-level estimates where AI at scale cuts inventory cost by 35 to 40 percent, lifts forecast accuracy by 20 to 30 percent, and reduces forecast errors by 20 to 50 percent. [R27] Those figures describe potential outcomes in digitally mature fashion enterprises with established enterprise resource planning (ERP) pipelines; they are not a baseline return that an industrialising Cameroonian mill can copy into an investment case. [R27] [S42]

9.18 COUNT PREPARATION, INTEGRATION AND HIDDEN MAINTENANCE DEBT

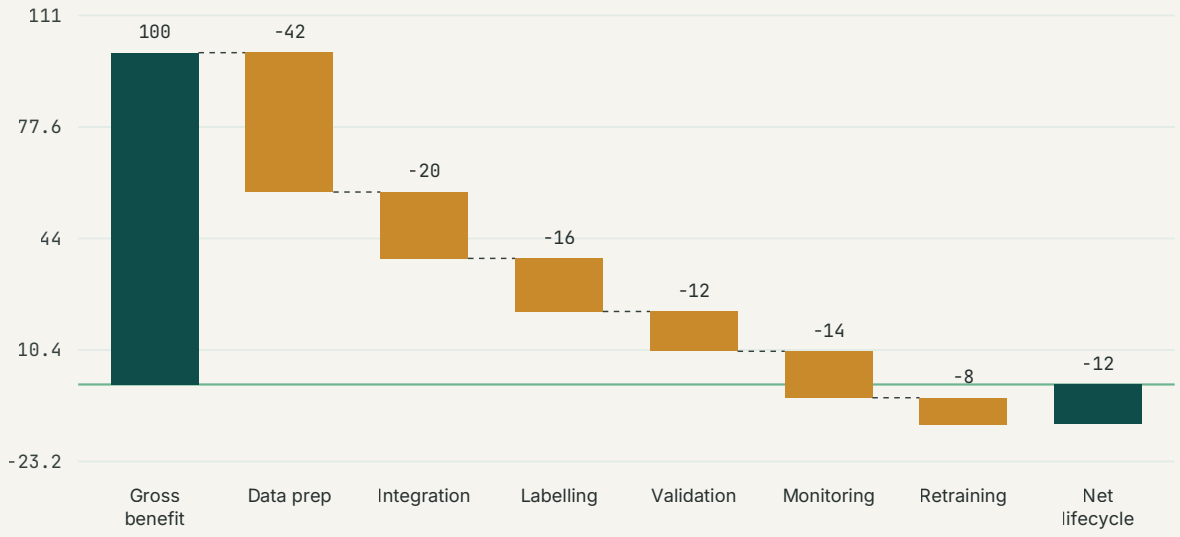
Research on machine-learning systems warns that model code is a tiny fraction of a production system, surrounded by data collection, feature extraction, verification and monitoring infrastructure that accumulate heavy ongoing maintenance costs. [S18] In manufacturing specifically, a 2025 NIST workshop report on semiconductor artificial intelligence notes that manufacturers have reported data-processing costs for a single application as high as 70 percent of total cost. [S39] That upper-bound warning highlights the trap small factories must avoid: building bespoke data-cleaning pipelines for each model instead of reusing a governed data spine (wp14 Chapters 01 and 05). [S18] [S39] Empirical OECD case evidence from small and medium enterprise digitalisation likewise shows that converting traditional processes into digital routines demands sizeable effort and a learning curve that does not pay off immediately. [S35]

To prevent hidden deficits, GFT's admission gate requires a whole-lifecycle cost sheet covering seven explicit buckets: capture and preparation increments beyond the baseline spine; ground-truth labelling by inspectors or engineers; integration and manual-fallback engineering; shadow-mode validation (wp14 Chapter 11); operator and supervisor floor training; ongoing drift diagnosis and retraining (wp14 Chapter 12); and contractual exit archival (wp14 Chapter 08). [R29] [S04] [S08] [S09] [S16] [S18] [S22] [S32] [S35] [S39] [S40] [S43]

Exhibit 39, p. 153 shows how counting these recurring lifecycle obligations can turn a positive gross pilot saving into a net negative return when a model requires heavy bespoke preparation and maintenance.

EXHIBIT 39 / ILLUSTRATIVE INDEX (GROSS BENEFIT = 100)

Lifecycle cost can reverse an apparent pilot return



ILLUSTRATIVE [S18] [S39] Illustrative reconciliation from gross annual operational benefit (100 index units) to net lifecycle contribution (-12 index units) for a custom history-dependent model with bespoke pipeline preparation. Cost categories derive from S18 and the upper-bound data-processing warning in S39; values are illustrative arithmetic, not a universal factory ratio.

In **Exhibit 39**, p. 153, a proposed pilot promises an illustrative gross operational benefit indexed at 100 units. Deducting application-specific data preparation (-42 units), shop-floor integration (-20 units), domain labelling (-16 units), shadow validation (-12 units), drift monitoring (-14 units) and periodic retraining (-8 units) brings total lifecycle costs to 112 units, leaving a net loss of -12 units (`100 - 42 - 20 - 16 - 12 - 14 - 8 = -12`). [S18] [S39] This waterfall shows why admission must evaluate *net* value after lifecycle effort, and why reusing a governed data spine across decisions (wp14 Chapter 01) is essential to shrink preparation and integration costs. [S18] [S39]

WORKED EXAMPLE: ADMISSION ARITHMETIC FOR A FABRIC-INSPECTION CANDIDATE AGAINST A COLD-START DEMAND FORECASTER

Consider an illustrative Cameroonian weaving and garment mill evaluating two candidates in Month 6 of its data-spine rollout. Operational volumes and defect rates are illustrative assumptions; wage and exchange-rate anchors come from the inherited register. [R14] [R15] [R29] [R32] Cameroon statutory minimum wage (SMIG) rose to XAF 43,969 a month from 23 February 2024 (about US$76), while the cross-sector average salary is about US$359 a month (roughly XAF 215,400 at an illustrative XAF 600 per US dollar conversion, alongside the euro peg of XAF 655.957). [R14] [R15] [R32] Literature reports manual defect inspection at 60 to 70 percent accuracy (missing 20 to 30 percent of defects), versus computer vision above 90 percent accuracy across more than 40 defect types at up to 60 metres a minute. [R29]

Candidate A (Illustrative bought edge vision on two frames, wp03 link): Inspecting 120,000 linear metres of greige cloth a month at a 65 percent manual catch midpoint leaves missed weave defects that cause 1,800 metres a month of downgraded dyed cloth, costing XAF 1,200 a metre (XAF 25,920,000 a year). [R29] Upgrading frame lighting and tally sheets recovers an estimated 25 percent of missed defects without software, saving XAF 6,480,000 a year less XAF 1,080,000 in lighting and supervision (net comparator benefit: XAF 5,400,000 a year). Bought camera assistance (Tier 2) lifts detection to an estimated 88 percent, catching 750 additional metres a month worth `750 * 12 * XAF 1,200 = XAF 10,800,000` a year gross. [R29] Setup and local labelling require 160 skilled hours at XAF 2,500 an hour (XAF 400,000 upfront), annualised hardware and licence costs of XAF 4,800,000, and 6 hours a week of steward checks (`312 * XAF 2,500 = XAF 780,000` a year), totalling XAF 5,980,000 in Year 1. [R15] Incremental illustrative net benefit over the comparator is `XAF 10,800,000 - XAF 5,980,000 = XAF 4,820,000` in Year 1. **Verdict: Admit to shadow candidate.**

Candidate B (Illustrative custom SKU demand forecaster, wp04 link): The mill holds 5 months of order history across 40 styles, dominated by 3 uniform contracts, with stockouts separated only in Month 4. A seasonal naive and backlog sheet costs zero new software outlay. [S0901] Reconstructing short history and maintaining bespoke pipelines costs an estimated XAF 9,600,000 in Year 1, while incremental inventory savings over the backlog comparator are near zero (illustrative range: XAF 0 to XAF 1,500,000). [S18] [S39] [S0901] Net value is negative (`-XAF 8,100,000` optimistic bound) and coverage fails ISO/IEC 5259-3. [S04] [S44] **Verdict: Postpone for 12 to 18 months of clean history.**

9.19

Stop and postpone are successful governance outcomes

In many digital programmes, admitting that a use case lacks clean data or positive economics is treated as a failure by the project team. [S19] [S44] That culture breeds zombie pilots: models that never reach production reliability but consume months of analyst effort. [S18] Empirical research on supporting small and medium enterprises in mapping artificial-intelligence potential states the corrective principle clearly: in information-technology investments, the decision to postpone an AI project is valid when resources, skills, costs or expected benefits are not met. [S44]

9.20 FOUR EXPLICIT EXIT ROUTES GOVERN EVERY ADMISSION REVIEW

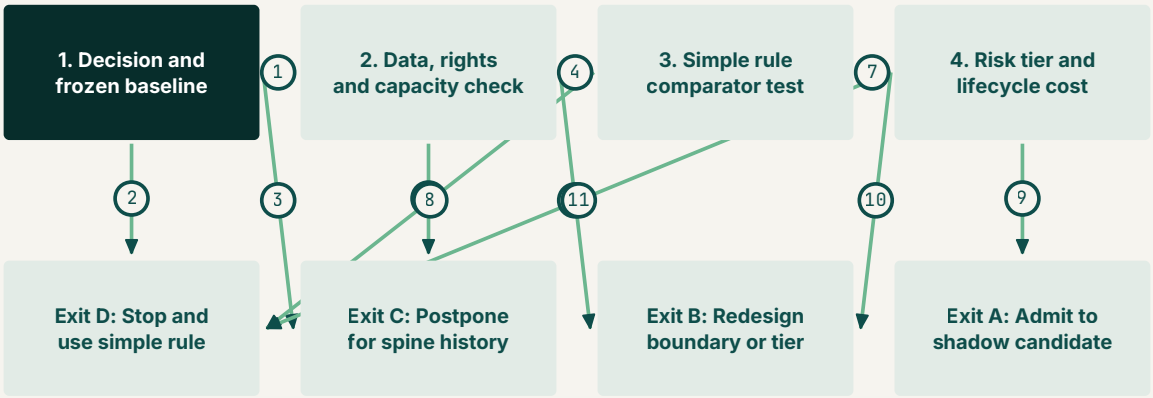
GFT designs the admission gate so that every candidate concludes with one of four documented dispositions under the AI management system (wp14 Chapter 02): [S12] [S43] [S44]

- **Admit to shadow candidate:** All prerequisites pass (named decision and actor, frozen KPI baseline, verified data fitness and rights, defined comparator, assigned risk tier, and positive net value). The candidate receives a registry ID (wp14 Chapter 10) and enters shadow evaluation. [S13] [S43]
- **Redesign scope or boundary:** The operational problem is valuable, but the proposed form is unsafe, unlawful or over-engineered. [S22] [S29] [S42] Closed-loop loom control becomes an open-loop supervisor alert (Tier 4 to Tier 2), or individual worker tracking becomes an anonymised line-balance report (Tier 3 to Tier 1). [S22] [S29] [S30]
- **Postpone with a prerequisite trigger:** The decision has strong potential, such as multi-season demand forecasting (wp04), but lacks clean history or steward capacity today. [S04] [S44] The register records the exact data trigger that re-opens review, such as fifty-two weeks of validated EPCIS order events (wp14 Chapter 06). [S08] [S44]
- **Stop permanently (and deploy the simple comparator):** Either no action follows the prediction, harm cannot be mitigated, or a simple rule captures the benefit at lower cost. [S30] [S44] [S0901] The model closes, and any useful rule or KPI clarification goes to operations. [S38] [S43]

Exhibit 40, p. 156 traces the derived decision path from problem statement to one of these four dispositions.

EXHIBIT 40 /

Every candidate exits through deploy, redesign, postpone or stop



1. Action and KPI frozen
2. No action
3. Baseline missing
4. Fit and lawful
5. History short
6. Rights or OT issue
7. Model needed
8. Rule suffices
9. Net value positive
10. Reduce risk tier
11. Cost exceeds value

DERIVED [S43] [S44] Derived stage-gate path from operational decision and frozen baseline through data and rights readiness, simple comparator test and risk-adjusted lifecycle value check to four explicit exit outcomes, based on the lifecycle phases in S43 and the SME postponement and stop criteria in S44.

As **Exhibit 40**, p. 156 shows, a candidate must pass all sequential checks to reach *Exit A*, whereas failing any single check routes the proposal immediately to *Redesign*, *Postpone* or *Stop* before training compute or vendor fees are spent. [S43] [S44] What **Exhibit 40**, p. 156 cannot enforce by itself is governance discipline; if a sponsor bypasses the gate without a recorded risk exception (wp14 Chapter 02), the chart becomes decorative. [S02] [S14]

9.21 PRESERVE REUSABLE LEARNING WHEN A CANDIDATE STOPS OR WAITS

When a candidate exits through *Postpone* or *Stop*, the admission review remains useful if recorded in the governance register. [S12] [S13] The dossier archives the frozen KPI definition (wp14 Chapter 03), the data defects uncovered during screening (wp14 Chapter 04), the simple comparator formula, and the failed assumption. [S04] [S38] [S44] Preparing a candidate for admission often exposes a broken identifier or ambiguous downtime state that the steward can fix in the core spine, improving daily management reporting (wp02) even when no model is built. [S19] [S36]

ADMISSION DOSSIER VERIFICATION CHECKLIST

Before the domain owner and model owner sign an admission certificate, the dossier must record six inspectable artefacts in version control (wp14 Chapter 10): (1) named operational decision, acting role, action, latency bound and escalation forum; (2) frozen KPI dictionary ID, version, ISO 22400-2 formula, baseline window and stratification slices; (3) dataset contract IDs, ISO/IEC 5259-3 quality pass rates, provenance link, GDPR data-minimisation check and Cameroon Law No. 2024/017 legal review status where personal data apply; (4) formula or rule for current practice and the simplest useful benchmark, plus required superiority margin; (5) risk classification into Tier 1 to Tier 4, human thresholds under NIST AI RMF 1.0, ILO Principle 5.5 compliance and NIST SP 800-82 safety isolation; and (6) itemised lifecycle cost ranges and explicit stop triggers. [S02] [S04] [S09] [S11] [S13] [S18] [S22] [S29] [S30] [S31] [S38] [S39] [S43] [S44] [S0901] [S0902]

9.22

GFT's sequence favours vision before history-hungry forecasting

GFT's Cameroon textile programme pairs an early governed data spine with a staged model sequence: admit bought computer-vision fabric inspection early where local validation and economics pass (wp03), and defer history-dependent machine-learning demand forecasting (wp04) until clean transactional history has been captured. [R28] [R29] [S16] [S44] The admission gate explains the engineering reason behind that sequence without repeating the optical design of wp03 or the forecasting equations of wp04. [S04] [S43]

9.23 WHY BOUGHT VISION CLEARS ADMISSION EARLIER THAN CUSTOM FORECASTING

Executives often assume that tabular demand forecasting is simpler than computer vision and should launch first. Applying the admission gates shows why the opposite holds in a newly modernising textile plant: [S04] [S43]

- **Data dependency and cold-start horizon:** A pre-trained fabric-inspection vision model (wp03) arrives with feature extractors trained on generic weave textures. [R28] [R29] Its local admission requirement is a stratified validation and threshold-tuning set of several hundred local fabric frames under mill lighting, paired with a governed defect taxonomy (wp14 Chapter 03). [R29] [S04] By contrast, a demand forecaster (wp04) cannot borrow another firm's seasonal cycle, Cameroonian school-uniform calendar, or export lead times. [R26] [S04] [S0901] Until the order, shipment and stockout spine records one to two seasonal cycles cleanly, complex forecasters overfit noise and lose to a backlog-plus-seasonal-naive comparator. [S04] [S0901]
- **Observability and label delay:** In fabric inspection, ground truth is immediate: when the camera flags a warp break or stain on Roll 1042, the perch operator inspects the cloth within seconds and confirms or overrides the mark, making shadow evaluation (wp14 Chapter 11) fast. [R29] [S08] In demand forecasting, a twelve-week ahead prediction waits twelve weeks for ground truth, and stockouts or promotions confound the label (wp14 Chapter 12). [S19] [S0901]

- **Actionability and containment:** A vision alert triggers a bounded, reversible action on one roll before dyeing adds chemical and energy cost. [R29] An erroneous demand forecast that triggers an oversized yarn purchase locks up working capital for a quarter. [S14] [S43]

Exhibit 41, p. 158 contrasts the admission evidence required for bought fabric-inspection vision (wp03) and demand forecasting (wp04). [S04] [S43]

EXHIBIT 41 /

Vision and forecasting need different evidence before admission

Admission dimension	Bought fabric-inspection vision (wp03)	History-dependent demand forecasting (wp04)
Operational decision and actor	Perch or frame operator marks fault, grades roll or cuts before dyeing	Commercial planner sets yarn purchase lots and line capacity allocation
Frozen baseline KPI (ISO 22400-2)	Confirmed defect catch rate and second-quality dyed metres per 1,000 m	Scaled forecast error (MASE) and stockout or excess-inventory days by style
Minimum data readiness (ISO/IEC 5259-3)	Controlled lighting, governed defect codes, stratified local roll validation set	12 to 24 months of clean order, shipment, cancellation and stockout events
Mandatory simple comparator	Standardised manual 4-point inspection with upgraded frame lighting	Confirmed contract backlog plus seasonal naive and moving-average rules
Risk tier and primary hazard	Tier 2 (quality disposition): false negatives enter dyeing; false alarms slow line	Tier 1 or 2 (planning): working-capital lockup in wrong yarn count or shade
Typical initial gate disposition	Admit to shadow candidate once spine defect codes and roll IDs pass quality gates	Postpone during cold start; use simple comparator until order spine matures

DERIVED [S04] [S43] [S44] [R29] [S0901] Derived admission comparison between early bought computer-vision fabric inspection (wp03) and later demand forecasting (wp04) under GFT's Cameroon textile programme, derived from S04, S43, S44, R29 and S0901.

As **Exhibit 41**, p. 158 summarises, neither use case receives an unconditional pass. Bought vision avoids the multi-season time-series cold start, but cannot bypass admission: if the mill has not stabilised roll identifiers, lighting and defect codes on the data spine, admitting a camera system produces unlinked image alerts that cannot be reconciled to looms or dye lots. [R29] [S04] [S08] Conversely, postponing machine-learning demand forecasting does not leave commercial planning unmeasured; it deploys the governed order-backlog and seasonal-naive comparator immediately while accumulating the clean event history needed for a statistical or machine-learning forecaster to earn admission later. [S44] [S0901]

GFT NOTE: ELAPSED CALENDAR TIME NEVER SUBSTITUTES FOR PASSED ADMISSION GATES

In GFT's 18-month implementation roadmap (wp14 Chapter 14), bought vision is scheduled for shadow evaluation earlier than history-dependent forecasting. Sponsors must treat those roadmap months as conditional planning windows, never as automatic permission to deploy. [S16] [S44] If a factory reaches Month 13 on the calendar but its order records still fail completeness checks or merge unfulfilled demand with shipped invoices, the forecasting candidate remains in *Postpone* status until the spine defect is fixed. [S04] [S44] Conversely, if an existing mill already holds three years of audited, item-level order and stockout records passing ISO/IEC 5259-3 quality gates, it can test a forecasting comparator without waiting for an artificial milestone. [S04] [S43] Evidence, not the clock, opens the gate.

GUARD AGAINST VENDOR PILOT TRAPS DURING ADMISSION

Technology vendors frequently offer "free" 30-day proof-of-concept pilots that bypass the factory's data spine, ingesting ad-hoc CSV extracts into a proprietary notebook and reporting accuracy on a vendor-selected slice. [S18] [S32] Accepting such an offer violates three admission rules: it leaves the factory without a frozen internal baseline, obscures true preparation and integration cost (wp14 Chapter 05), and creates lock-in where the factory cannot reproduce the evaluation. [S18] [S32] [S39] Every vendor pilot must consume governed products from the factory's spine, evaluate against the frozen KPI and simple comparator, and disclose full multi-year licence, hardware and exit terms before admission. [S32] [S38] [S43]

9.24 CONNECT THE ADMISSION GATE TO THE FACTORY'S ANALYTICAL TOOLS

To make this admission discipline repeatable without bureaucracy, the rules in this chapter map directly to three operational tools in the GFT suite: the `use-case-portfolio` tool implements the admission scorecard (**Exhibit 38**, p. 150) and stage-gate routing (**Exhibit 40**, p. 156); the `inspection-economics` tool operationalises the Tier 2 fabric-vision admission calculation (wp03) using local roll throughput and

manual catch rates; and the `cold-start-forecast` tool operationalises the comparator gate for `wp04`, benchmarking candidate algorithms against mean, naive, seasonal naive and drift baselines while penalising short history. [R29] [S0901]

When a candidate clears all six checks in the admission dossier, it becomes an authorised engineering work item. The next chapter (`wp14` Chapter 10) specifies how that admitted candidate is trained, versioned and recorded so that another competent practitioner can reconstruct every dataset snapshot, feature transformation, parameter and evaluation result that produced it. [S09] [S13]

A factory earns the right to train a model only after it has named the decision, frozen the baseline, tested a simple rule and priced the full cost of keeping the model alive.

CHAPTER SOURCE NOTES

S43 Frontiers in Artificial Intelligence. 2021-06-14.

S44 The Journal of Technology Transfer, Springer Nature. 2024-07-27.

S02 International Organization for Standardization. 2022.

S18 Advances in Neural Information Processing Systems. 2015.

R06 Cameroon Trade Hub. Publication date unrecorded; inherited source-text check.

R07 Ecofin Agency. Publication date unrecorded; inherited source-text check.

S19 Association for Computing Machinery, CHI 2021. 2021.

S13 ISO/IEC. 2023.

S08 GS1 AISBL. 2022.

R29 MDPI; industry inspection literature. Publication date unrecorded; inherited source-text check.

S22 National Institute of Standards and Technology. 2023.

S36 World Bank. 2023.

S0902 Google for Developers. 2025-08-25.

S38 International Organization for Standardization. 2014.

S0901 OTexts. 2021-05-31.

R03 SODECOTON. Publication date unrecorded; inherited source-text check.
S04 ISO/IEC. 2024.
S29 European Parliament and Council of the European Union. 2016-04-27.
S09 World Wide Web Consortium. 2013-04-30.
S31 Presidency of the Republic of Cameroon. 2024-12-23.
S30 International Labour Office. 1997.
S32 European Parliament and Council of the European Union. 2023-12-13.
S16 Microsoft Learn, Azure Architecture Center. 2026-06-03.
S42 SCITEPRESS, IN4PL 2022. 2022.
S37 World Bank. 2025-05-20.
S11 National Institute of Standards and Technology. 2023-01-26.
S14 ISO/IEC. 2023.
S05 International Electrotechnical Commission. 2013.
S23 National Institute of Standards and Technology. 2024-02-26.
S39 National Institute of Standards and Technology. 2025.
R27 McKinsey, The State of Fashion 2025. Publication date unrecorded; inherited source-text check.
S35 OECD. 2022-11-21.
S40 International Labour Organization and UNESCO. 2020.
R14 WageIndicator. Publication date unrecorded; inherited source-text check.
R15 wage.is. Publication date unrecorded; inherited source-text check.
R32 Wise; XE. Publication date unrecorded; inherited source-text check.
S12 ISO/IEC. 2023.
R28 DataIntelo. Publication date unrecorded; inherited source-text check.
R26 Statista. Publication date unrecorded; inherited source-text check.

10

PART III: CONTROL MODELS AS PRODUCTION CHANGE / REPRODUCIBILITY

Reproducible training turns experiments into auditable candidates

A model becomes a release candidate only when another competent person can reconstruct what data, code, configuration and evaluation produced it.

CHAPTER 10
DECISION TEST

A model becomes a release candidate only when another competent person can reconstruct what data, code, configuration and evaluation produced it.

Reproducible training means preserving an inspectable chain from the factory record to the candidate artefact and its evaluation. GFT proposes that a candidate cannot enter release review until a competent colleague can retrieve that chain, rerun the agreed checks and explain any unresolved difference. A successful training job is evidence of execution. It does not establish operational fitness or permission to deploy. The preceding chapter determines whether modelling should begin; the following chapter determines whether a candidate may affect operations. This chapter defines the evidence package between those decisions.

10.1

Version the complete experiment, not only the model file

Google's MLOps architecture connects pipeline execution metadata with lineage, reproducibility and comparison. The research overview by Kreuzberger and colleagues treats MLOps as practices, concepts and development culture as well as architecture. Together they support preserving the experiment as a connected record rather than handing over an unexplained model file. [S15] [S20]

GFT's proposed experiment record starts with the admitted purpose and intended output. It identifies the code revision, configuration file, dependency lock, data manifest, feature definitions and evaluation procedure. The execution environment records the operating system, runtime, relevant numerical libraries, processing device and driver where material. Parameters include stopping rules and initialisation settings, including random seeds where used. References point to retrievable artefacts, rather than to somebody's recollection of what was on a laptop. Credentials remain outside this record.

Separate an experiment plan from a run and from a candidate. The plan describes the comparison being attempted. A run records an actual execution, including interruption or failure. A candidate is the immutable artefact bundle selected for review. A rerun receives its own run identifier even if every declared input matches. An exported, quantised or otherwise transformed model receives a linked candidate identity because its actual operating artefact has changed. Retaining the parent prevents a conversion step from disappearing between training and evaluation.

Allocate the candidate identifier before results are reviewed, then bind it to a manifest of content digests. The manifest should cover weights, preprocessing, label mapping, configuration and scoring wrapper. GFT recommends keeping both a readable identifier and digests: the identifier supports conversation, while the digests make later integrity checks explicit. A digest is an integrity comparison, not proof that the training data were correct or that the person claiming approval had authority. Permissions and approval records remain separate controls.

Define what a reconstruction must match. GFT recommends exact agreement for inputs, partition membership, parameter files and stored artefact digests. For newly executed numerical outputs, the experiment plan should specify whether exact agreement is required or a declared tolerance is acceptable. Record the reason for any tolerance and compare the operating decision as well as the summary metric.

Never introduce a wider tolerance after an inconvenient rerun. An unresolved discrepancy leaves the package incomplete even when the original score remains attractive.

The minimum stack in **Exhibit 42**, p. 164 is an illustrative design, informed by the sources, rather than a claim that GFT operates a deployed training platform. It separates durable records from the authority to change their status. [S15] [S16]

EXHIBIT 42 / PROPOSED CAPABILITY LAYERS

Small capabilities keep the experiment reconstructable



ILLUSTRATIVE [S15] [S16] [S20] Illustrative GFT design informed by execution lineage and incremental MLOps. Layers describe responsibilities, not an implemented system or mandatory software products.

The layers may share a modest storage service. Their separation is logical: editing a descriptive note must not replace a model, and completing a run must not grant release authority. This distinction gives the factory a practical way to review work without purchasing a separate product for every function.

10.2

Dataset snapshots prevent silent training-set change

The public description of the data-quality management standard covers managing data used in analytics and machine learning. The provenance ontology provides a vocabulary for tracing entities, activities and responsibility. GFT applies these ideas to snapshots: the extract, its creation activity and the accountable roles should remain identifiable together. [S04] [S09]

A snapshot should preserve the actual records used, subject to approved retention and access rules. A saved query alone is insufficient for the proposed control because its underlying tables may subsequently change. Retain the query, extraction time, source versions, schema, row identifiers, exclusions and quality-rule version alongside the immutable extract or an equivalent reconstructable historical source. Record reconciliation totals at the same grain as the extract. If later correction changes a label, create a successor snapshot and explain its relationship to the earlier version.

Define prediction time before defining partitions. For an order-planning model, specify what the planner could have known when issuing the forecast. GFT recommends recording both the event time and the time a record became available to the decision process. An event that occurred earlier but arrived after the forecast cannot silently enter the historical input. This rule applies equally to corrections, final quality disposition and later cancellations. Keep the target's eventual outcome separate from the inputs available when the decision was made.

Leakage is a documented reproducibility concern in machine-learning research. Kapoor and Narayanan examine methodological failures and propose reporting that exposes how information enters model development. Their research is a warning about evaluation design, not a measured failure rate for textile factories in Cameroon. [S1001]

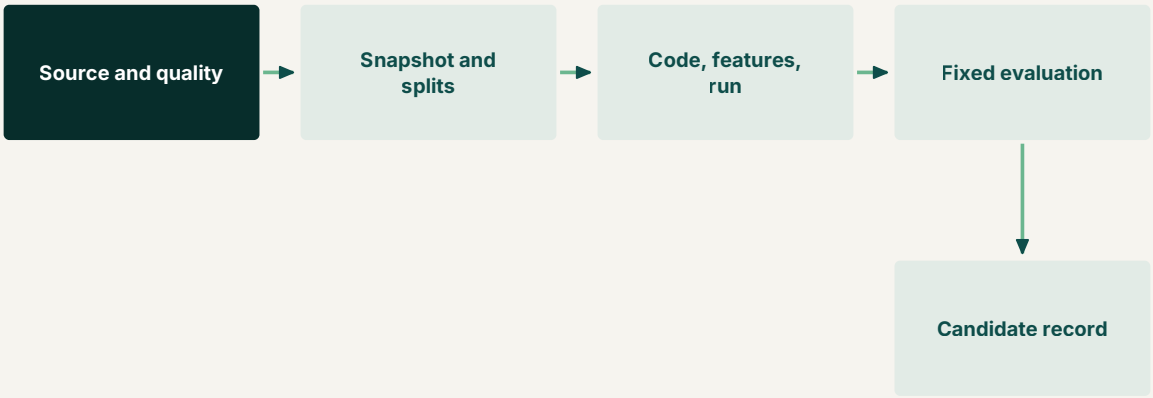
For the factory protocol proposed here, use chronological partitions for forecasting and define any gap around the prediction horizon. Keep overlapping target windows from transferring future outcomes into earlier training examples. For visual inspection, allocate related images at the physical entity boundary, such as a roll or production lot, before generating crops or augmentations. Record the grouping key and check its intersections across partitions. If the intended claim concerns a new style or machine, reserve that entity accordingly; otherwise describe the narrower scope actually tested.

Fit imputers, scalers, feature selection and other learned preparation only on the training partition. Apply their frozen outputs to validation and test records. Tune using validation evidence and keep the final test sealed until the selection procedure is complete. GFT recommends logging each opening of the test set. If its results lead to another tuning decision, retain those results but withdraw the claim that the set remains untouched. Obtain a new independent evaluation boundary before making that stronger claim again.

The lineage in **Exhibit 43**, p. 166 is a derived mapping of provenance and pipeline-metadata concepts into proposed factory records. Its inputs are identifiable source entities, recorded transformation activities and execution artefacts. It is a traceability checklist, not evidence that a particular dataset has passed it. [S09] [S15]

EXHIBIT 43 / EVIDENCE DEPENDENCIES

Candidate identity connects the record to review



DERIVED [S09] [S15] Derived mapping: provenance entities and activities from S09 plus pipeline execution records from S15 map to source, snapshot, run, evaluation and candidate records. Arrows show dependencies, not approval or elapsed time.

Reconstruction should stop at a broken dependency. A missing snapshot cannot be cured by a screenshot of its row count, and a reproducible training script cannot establish which labels were supplied. Preserve uncertainty in the candidate record and repair the missing link before describing the package as complete.

Label history deserves its own boundary in this proposed design. Keep the annotation instructions, vocabulary version, adjudication process and correction history with the snapshot. Record whether a label represents an initial operator judgement, a later laboratory result or a final customer disposition. Those outcomes answer different questions. If disputed examples are excluded, preserve their identifiers and reasons so a reviewer can assess the resulting scope. A tidy extract should not conceal an unresolved disagreement about what the target means.

For a supplier-generated dataset, GFT recommends obtaining a manifest that makes permitted use and provenance inspectable without demanding unnecessary personal information. Establish whether duplicate material can appear across deliveries and whether the evaluation records were previously available to the supplier for tuning. Where that cannot be established, qualify the independence claim. Preserve the restriction in the candidate record rather than relegating it to a procurement conversation that the eventual release reviewer may never see.

10.3

Features inherit governed operational definitions

The manufacturing KPI standard describes indicators through their formula, elements, time behaviour and units. GFT extends that definition discipline to model features, while distinguishing a model input from a management indicator. A feature may be useful without being a KPI, but it still needs a stable operational meaning. [S38]

Define each feature's source event, entity key, unit, aggregation window, time basis and missing-value treatment. Include its admissible range, correction policy and owner. A feature called recent downtime should state whether it measures stopped minutes, stop counts or lost scheduled time. It should also state whether changeover and planned maintenance are included. Avoid allowing a model notebook to create an alternative meaning for a term that production managers already use. A changed definition becomes a new feature version with an impact review.

GFT recommends distinguishing an observed zero from an unavailable value. A machine that reported no stoppage and a machine that sent no report provide different evidence. Preserve a missingness indicator where appropriate and define what the scorer does when information is late or absent. If the operating procedure requires the model to abstain, training and evaluation should retain that path. Do not fill every blank merely to make the input matrix acceptable to an algorithm. The reason for absence may matter more than the imputed value.

Retain the preparation artefacts that convert governed fields into numerical inputs. These include category dictionaries, normalisation parameters, vocabulary, image transformations and the order of output classes. GFT's proposed rule is that inference consumes the same versioned transformation as evaluation. A separate implementation requires an equivalence check on agreed fixtures before it can replace the original. Otherwise the candidate record would identify the weights precisely while leaving their input interpretation uncertain. The wrapper belongs in the artefact bundle for this reason.

For example, an illustrative feature calculates unplanned-stop share from an assumed 24 stopped minutes and 480 scheduled minutes: 24 divided by 480 equals 0.05, or 5 percent. Its unit is a fraction of scheduled time, its grain is the completed shift, and its assumed denominator excludes planned shutdown. This is a constructed definition example, not a GFT factory observation. If the decision occurs midway through the shift, its final totals are unavailable and this completed-shift feature is ineligible for that decision.

Use that example to challenge timing as well as arithmetic. A feature with the right formula can still have the wrong availability boundary. GFT recommends storing an availability rule alongside every lookback window: the newest record must be usable by the scoring cut-off, rather than merely carry an earlier event date. Where historical availability cannot be reconstructed, flag the evaluation limitation explicitly and collect prospective evidence. Do not substitute a convenient finalised history and call it an exact replay of a past decision.

METHOD

For each feature fixture, preserve the source rows, prediction cut-off, expected unit, missing-data state and expected transformed value. Run the fixture through the training preparation path and the scoring path, then compare the outputs under the declared tolerance. Include an unseen category, a late correction and an absent observation when those conditions belong to the intended use. Record the outcome against the feature and wrapper versions. This is GFT's proposed equivalence method, not a certification test or a substitute for local model validation.

Feature ownership should remain with the operational definition owner, supported by the model builder. The builder may propose a different window or encoding, but should not silently change what a stoppage or accepted unit means. This keeps the data spine connected to factory work when a later experiment replaces the algorithm.

10.4

Execution metadata make comparison possible

Execution records provide the link between an intended pipeline and what actually ran. Google's architecture explicitly uses pipeline metadata for comparison and diagnosis as well as lineage. GFT adopts that principle and requires a run to record its start, completion state, inputs, outputs and diagnostic references, including unsuccessful executions. [S15]

In the proposed log, distinguish completed, failed, cancelled and invalid executions. A completed job may still be invalid for comparison if it used the wrong snapshot or an undeclared feature. Retain the reason and superseding run rather than deleting the inconvenient entry. Record material manual intervention, including a resumed checkpoint, changed resource allocation or a corrected configuration. Never overwrite an earlier run's metrics after repairing a bug. The repair creates new evidence and should preserve the link to the defect it resolves.

Lock the comparison protocol before selecting the winner. Candidate, simple comparator and existing planning rule should face identical evaluation records, target definitions, horizons and weighting. Record the metric implementation as code, including exclusions and handling of missing predictions. GFT recommends a prediction ledger containing the evaluation key, actual outcome, prediction, abstention state and candidate identifier. This lets a reviewer recompute a summary and investigate a local failure without trusting a chart exported by the builder.

WORKED EXAMPLE: A LOWER MEAN CAN CONCEAL A WORSE PEAK

This illustrative planning comparison assumes actual order quantities of 100, 120, 80 and 140 units over equally spaced planning periods. A simple comparator predicts 110, 100, 90 and 120 units, giving absolute errors of 10, 20, 10 and 20. Its mean absolute error is (10 + 20 + 10 + 20) divided by 4, or 15 units. Candidate A predicts 105, 110, 85 and 130 units, giving (5 + 10 + 5 + 10) divided by 4, or 7.5 units. Every value is an assumption for demonstrating an auditable calculation.

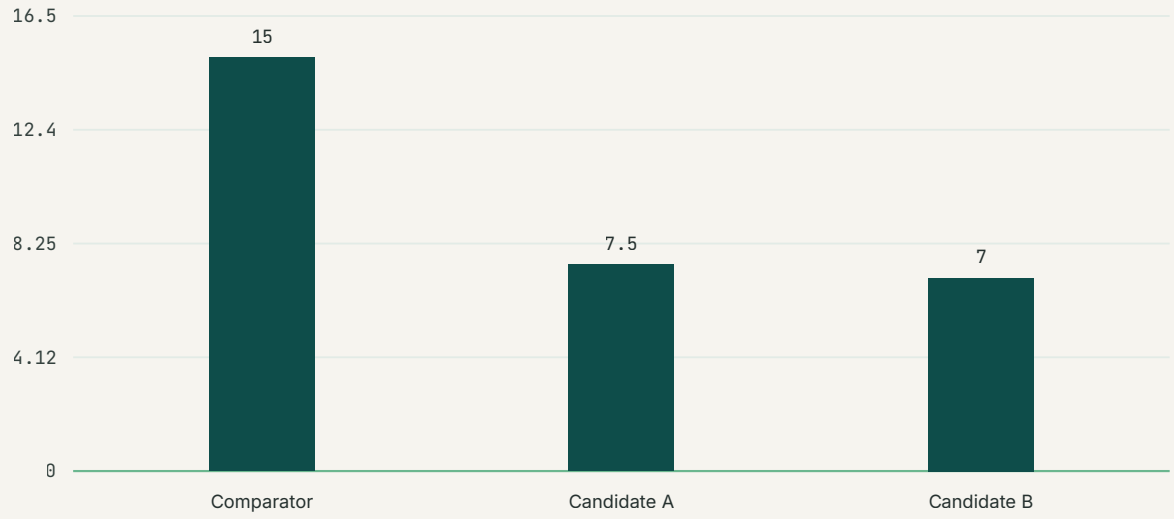
Candidate B in the same illustrative example predicts 100, 120, 80 and 112 units. Its absolute errors are 0, 0, 0 and 28, so its mean absolute error is 28 divided by 4, or 7 units. Relative to the comparator, A reduces mean error by (15 minus 7.5) divided by 15, or 50 percent. B reduces it by (15 minus 7) divided by 15, or approximately 53.3 percent. B has the better mean but the larger peak-period shortfall.

The illustrative sample supports arithmetic, not a reliability or business-value claim. If a peak-period shortfall were an agreed acceptance constraint, it would need to be reviewed alongside the mean. A reviewer needs the predictions and period labels to see the trade-off. The owner sets that constraint before candidate selection; it cannot be inferred from the average after the results arrive. NIST assigns judgement over metrics and threshold values to people. [S11]

The illustrative mean errors in **Exhibit 44**, p. 169 use only the assumptions and calculations above. Lower bars show smaller mean error; they do not show the distribution, uncertainty or permission to act.

EXHIBIT 44 / ILLUSTRATIVE MEAN ABSOLUTE ERROR, UNITS

Mean error alone favours the peak-period failure



ILLUSTRATIVE Illustrative assumed actuals [100,120,80,140]. Comparator predictions [110,100,90,120]; A [105,110,85,130]; B [100,120,80,112]. Mean absolute errors are 60/4=15, 30/4=7.5 and 28/4=7. B's final-period shortfall is 28 units. No observed factory result or forecast benefit is claimed.

GFT recommends keeping computational cost and elapsed duration in the same run record, with units and measurement boundaries. Treat estimated costs as estimates and distinguish active processing from queueing or human preparation. These records

support later lifecycle costing without turning this chapter into a business-case approval. A cheaper run is not necessarily a cheaper operating system, and a faster experiment does not justify weaker evidence retention.

Preserve the selection history as well as the selected run. GFT proposes an experiment summary that lists the candidate families explored, the tuning boundary and the rule used to choose the submitted artefact. If several random initialisations are evaluated, retain their results and explain how the reported candidate was chosen. Do not present the most favourable observed run as though it were the predetermined procedure. The summary need not reproduce every diagnostic message, but it should make discretionary choices visible enough to challenge.

A reconstruction discrepancy should open a bounded investigation. First compare the declared inputs and artefact digests, then compare feature outputs and record-level predictions before interpreting the summary metric. Identify the earliest point of disagreement and retain the original package. GFT recommends distinguishing a documentation defect, an environment dependency and a genuine difference in the model's decisions. Each needs a different repair. A new successful execution resolves the issue only when the reason for the earlier disagreement is recorded or its remaining uncertainty is explicitly accepted for the stated scope.

10.5

A registry records status and authority

The AI lifecycle standard describes processes for managing AI systems and explicitly covers acquisition as well as development. GFT's proposed registry translates lifecycle control into named candidate states. The state names below are a factory design choice; they are not a certification claim or a prescribed state machine taken from the standard. [S13]

Candidate means the bundle has been registered, including known gaps. Validated means the agreed checks passed for a stated scope, with the evidence linked. Approved means the authorised owner accepted a particular scope and set of conditions. Deployed means that an identified operating endpoint has confirmed that version is active. Suspended means use is halted while the record is retained. Retired means the candidate is withdrawn from further authorised use, subject to the retention policy. GFT recommends recording rejection as a decision outcome against a candidate, rather than erasing the candidate itself.

Keep release readiness distinct from active use. An approval record should specify the candidate digest, permitted line or planning process, conditions, approver and effective window. A deployment record should identify the actual service, device or batch job and confirm its loaded wrapper and configuration. An approved model can have no deployment, and a deployment may fail to activate the approved bundle. Until endpoint confirmation exists, report the state that is evidenced. Do not infer production status from a registry label alone.

For a small factory, GFT proposes a current-state table backed by an append-only transition history. Every transition records the prior state, new state, reason, actor, time and supporting evidence. Corrections append a resolving entry instead of rewrit-

ing the original decision. The technical custodian protects storage and access; the operating owner controls scope and acceptance. Where a person holds overlapping roles, record the arrangement and obtain the review required by the consequence of the intended use. The next chapter defines release scrutiny in detail.

The model-record fields in **Exhibit 45**, p. 171 are a derived design checklist. Their inputs are lifecycle control, provenance, data-quality management, indicator definition and human judgement over thresholds. Mapping each principle to an inspectable record does not establish that a candidate already satisfies it. [S04] [S09] [S11] [S13] [S38]

EXHIBIT 45 / REQUIRED RECORD GROUPS IN A PROPOSED DESIGN

The candidate record separates knowledge from authority

Record group	Inspectable evidence
Purpose and scope	Decision, owner, exclusions
Data and rights	Snapshot, quality, access basis
Feature meaning	Definition, window, unit
Run and artefact	Code, environment, digests
Evaluation	Protocol, predictions, limitations
Authority and exit	State history, approval, fallback

DERIVED [S04] [S09] [S11] [S13] [S38] Derived mapping of source principles to GFT record groups: data quality to scope and snapshots; provenance to lineage; human judgement to approval; lifecycle control to status and exit; KPI definitions to feature meaning. This is a proposed template, not an implemented registry or standard-conformity claim.

Retention is part of the model record. GFT recommends recording whether worker-level information is present, where restricted evidence is held and who may access it. Do not copy personal records into an unrestricted registry to make lineage convenient. Preserve the minimum permitted evidence needed to explain decisions, and link to controlled material where appropriate. If deletion or contractual limits prevent later reconstruction, record what can no longer be reproduced and reassess continued use. Legal applicability and retention periods require qualified advice; this chapter supplies an engineering record design, not legal advice.

10.6

A lightweight stack avoids platform sprawl

Microsoft's maturity model recommends gradual capability development and describes both organisational and technical dimensions. GFT uses that advice to separate necessary evidence from optional automation. A factory should choose software only after defining which recurring record or decision it must sustain. [S16]

The proposed starting stack comprises version control for code and definitions, controlled storage for snapshots and artefacts, a structured run log and a registry with transition history. These are capabilities, not a mandatory procurement list. A disciplined directory structure and validated structured files may initially provide several

capabilities together. Give the custodian a documented export and restoration procedure. If a platform is adopted later, preserve the identifiers and relationships so historical candidates remain intelligible outside its interface.

Plain portable records need an explicit schema. GFT recommends mandatory fields for identity, parent identity, record type, creation time, responsible role, content location, digest and access classification. Domain-specific records then add partitions, metrics or approval conditions. State the schema version in every export. Validate references before accepting an import and flag a missing dependency rather than dropping it silently. A portable model file without its label mapping, feature contract or evaluation keys is an incomplete transfer under this proposed protocol.

Make recovery part of the reconstruction exercise. Ask a colleague to start from an approved identifier on a clean permitted environment, retrieve the package, verify integrity and rerun its evaluation. Record missing instructions, inaccessible files and unexpected manual choices. The colleague need not repeat a costly full training job every time to reveal a broken evidence chain. Define when training reruns are warranted, such as an unexplained result difference or material environment change, and distinguish that exercise from rescoring the stored artefact.

> **GFT NOTE**
>
> Start with the smallest record system that a colleague can actually reconstruct. Require an identified candidate, frozen data boundary, governed feature definitions, retrievable execution evidence and a separate authority record. Automate the most frequent evidenced failure in that chain first. Keep a named custodian and a portable export even when a supplier operates the tooling. This is GFT's practical recommendation, not a claim that any particular software tier guarantees reproducibility.

Automation should remove repeated transcription and enforce declared rules. GFT recommends beginning with manifest generation, reference validation and automatic capture of actual run inputs. Leave disputed operational meaning and acceptance with the responsible people. A scheduler can repeat a mistaken query more efficiently; it cannot resolve who owns the definition. Similarly, an automatic registry update should not transform completion of training into approval. Record the automation's own version because it participates in the evidence chain it creates.

The factory should also define an exception route. If an urgent investigation uses an incomplete experimental package, label it exploratory and prohibit it from acquiring release status through convenience. Assign an owner and a clear condition for completing or abandoning the record. Preserve whatever learning is reusable without presenting the investigation as an accepted candidate. This keeps the system workable during operational pressure while retaining a meaningful distinction between exploratory work and evidence suitable for a release decision.

Close the reconstruction exercise with an evidence receipt. GFT proposes a short record stating who attempted it, which package and environment they used, what they could retrieve, which checks they executed and what remains unresolved. Record whether they rescored a stored artefact, rebuilt the wrapper or repeated training, because those exercises support different claims. The receipt gives the release

reviewer a bounded statement of work actually performed. It should never be titled independent validation unless the assigned reviewer also performed the separate validation work.

10.7

Bought models face the same evidence discipline

Acquiring a model changes which evidence the factory can obtain, but does not remove lifecycle responsibility. The AI lifecycle standard's public abstract explicitly includes acquired systems. GFT therefore recommends applying the same candidate identity and local evaluation discipline to a bought model while recording the limits of access to its internal development process. [S13]

For a bought inspection system, record supplier, product, model identifier, wrapper version, input requirements, class mapping, licence terms and update channel. Preserve the documentation version supplied with the product and distinguish supplier claims from locally observed results. Record whether training data, training code and weights are accessible. Unknown should remain unknown. Do not infer local reproducibility from a vendor's training claim or describe a local test as independent reconstruction of an inaccessible training pipeline.

The procurement requirement should specify notification of material changes and the evidence needed for revalidation. GFT recommends defining materiality to include preprocessing, defect taxonomy, thresholds, model weights and the service contract that determines outputs. Where an update cannot be pinned, log the provider's available version markers and dates, then describe the remaining uncertainty. An endpoint name that stays constant while behaviour changes is not sufficient candidate identity for the proposed protocol. The owner must decide whether that uncertainty is tolerable for the intended decision.

Keep the locally used configuration inside the package. A purchased model may still be tuned through thresholds, calibration, prompts or operating recipes. Preserve those inputs and any local adaptation data with their rights and lineage. A tuned derivative gets a new identity linked to its supplier base. GFT recommends running the same fixed local evaluation on the actual delivered interface, including its preprocessing and abstention behaviour. This tests the system the factory will receive rather than a selectively exported component.

Agree an exit route before dependency becomes difficult to reverse. Where rights permit, retain the last approved local artefact and the environment required to read it. Where a remote licence or service prevents that, document the feasible fallback, including a return to the approved manual procedure. Do not promise recovery of a supplier service the factory cannot restore. Link the fallback evidence to the candidate record; the following chapter addresses the operational rehearsal and authority needed before release.

The handover should expose contractual dependencies alongside technical ones. GFT recommends recording whether access to historical predictions, model versions or evaluation exports ends when a subscription lapses. Obtain an export before that

boundary and verify that its identifiers still resolve to the retained local evidence. If a service cannot provide the evidence required for the intended use, the owner can narrow the scope, negotiate a different arrangement or decline the candidate. Recording that constraint early protects the factory from mistaking a convenient interface for a controllable operating asset.

GFT's programme sequence connects early bought vision and later history-dependent forecasting through the same record discipline. This chapter proposes that common discipline; it does not claim either system has been trained, validated or deployed by GFT. The practical distinction is evidence availability. A bought system may offer limited training visibility but still needs a traceable local test. A locally trained forecaster must preserve its historical information boundary as well as its code. Both hand the release reviewer an identifiable artefact, an explicit scope and an honest account of what remains unknown.

A candidate earns trust when its evidence can travel without its author.

CHAPTER SOURCE NOTES

S15 Google Cloud. 2024-08-28.

S20 arXiv, author version of an IEEE Access article. 2022-05-14.

S16 Microsoft Learn, Azure Architecture Center. 2026-06-03.

S04 ISO/IEC. 2024.

S09 World Wide Web Consortium. 2013-04-30.

S1001 arXiv, Sayash Kapoor and Arvind Narayanan. 2022-07-14.

S38 International Organization for Standardization. 2014.

S11 National Institute of Standards and Technology. 2023-01-26.

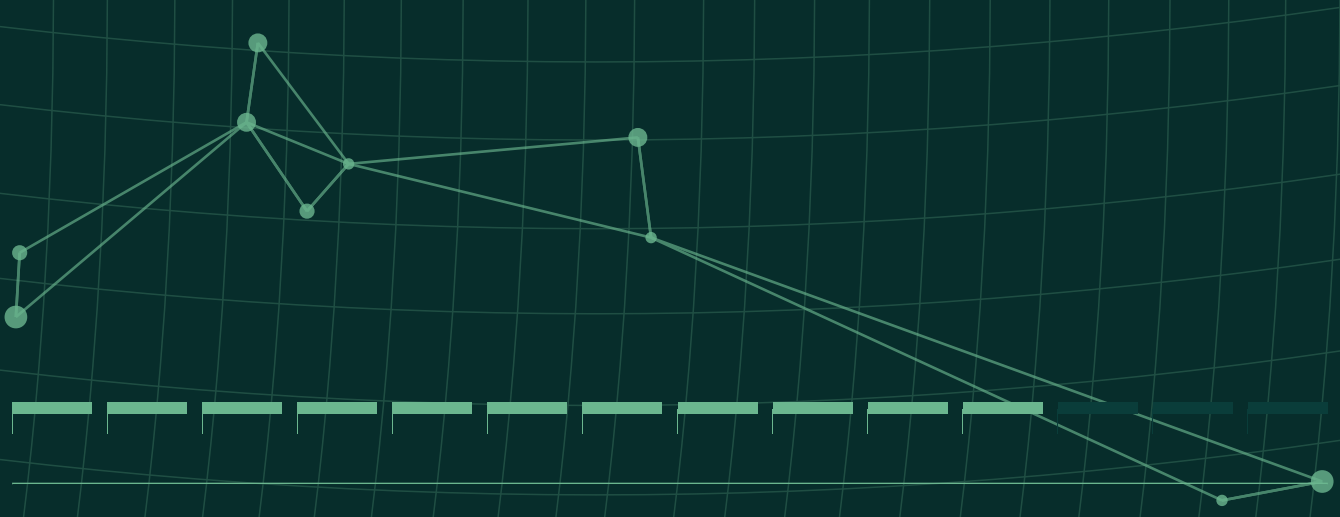
S13 ISO/IEC. 2023.

11

PART III: CONTROL MODELS AS PRODUCTION CHANGE / CONTROLLED RELEASE

Validation and staged release keep models subordinate to operations

Release authority follows demonstrated local fitness, safe fallback and accountable acceptance, not a strong aggregate validation score.



CHAPTER 11
DECISION TEST

Release authority follows demonstrated local fitness, safe fallback and accountable acceptance, not a strong aggregate validation score.

A machine-learning candidate that scores well on an offline test partition has not yet earned the right to alter factory work. Release is an operational change decision, not the automatic final step of a training script. [S11] [S15] Promoting a model into production changes how inspectors examine cloth, how planners size buffers, or how supervisors allocate maintenance across shifts. [S43] Governing that change by a single aggregate accuracy figure exposes daily execution to unmeasured subgroup failures, false-alarm fatigue, silent pipeline defects and workflow dependency. [S19] [S1101] Green Future Tech Limited (GFT), Hong Kong SAR, designed its three-part Cameroon textile programme (prepared 31 January 2024, with source-refresh notes retained in 2026) around the sector baseline in which about 95 percent of Cameroonian cotton fibre is exported raw and about 5 percent is processed locally, against a 50 percent local-processing target by 2030 (these inherited baseline entries were not refreshed in 2026 and their original publication dates are unrecorded). [R06] [R07] Moving from raw-fibre export to competitive local spinning, weaving, dyeing and garment assembly requires decision tools that line managers trust under shop-floor constraints. [R06] [S04] Validation and staged release keep every bought or locally trained model subordinate to operations through four controls: human-set thresholds tied to floor consequences, disaggregated validation across local slices, staged authority expansion from shadow mode to assisted operation, and rehearsed rollback to a known baseline. [S11] [S16] [S22]

Earlier chapters in this part established that a candidate enters development only after an operational decision, baseline and comparator are frozen, and that training becomes auditable only when code, dataset snapshots, feature definitions and execution lineage are recorded in a versioned registry. [S09] [S15] [S38] Those controls prove how a candidate was built, not that it is safe or useful on the floor. [S13] [S1102] This chapter specifies the validation and release discipline that separates technical calculation from operational acceptance, prices the cost of wrong interventions, sequences deployment through conditional stages, and treats rollback and manual fallback as tested release artefacts. [S14] [S15] [S43]

11.1

People set the acceptance metrics and thresholds

An automated pipeline can compute precision, recall, mean absolute error or calibration curves, but it cannot decide what an error costs on a weaving shed floor or in a cutting room. [S11] [S38] The National Institute of Standards and Technology (NIST) Artificial Intelligence Risk Management Framework (AI RMF 1.0) states the rule directly: human judgement must be employed when choosing the specific metrics related to AI trustworthiness characteristics and the precise threshold values for those metrics. [S11] Delegating threshold selection to software defaults or to the developer who trained the model confuses mathematical convenience with operational risk appetite. [S11] [S14] In a factory spine, the domain operating owner and designated risk owner set acceptance metrics and cut-off values jointly before validation begins, locking those thresholds into the evaluation protocol so they cannot be adjusted retroactively. [S02] [S11] [S15]

11.2 CONSEQUENCE-WEIGHTED THRESHOLDS SEPARATE TECHNICAL METRICS FROM OPERATIONAL ACCEPTANCE

Every industrial prediction error carries an asymmetric physical or commercial cost. [S11] [S43] In fabric-defect detection, a false negative lets a stained or broken-pick roll pass into dyeing or cutting, multiplying wasted labour, dye liquor and cloth. [R29] [S43] A false positive stops the frame or flags a clean roll for manual re-inspection, consuming operator minutes without destroying material. [R29] [S38] In short-horizon demand forecasting, over-prediction ties up working capital in slow-moving greige cloth, whereas under-prediction triggers short shipments or overtime changeovers. [S38] [S43] Because these consequences hit production, quality and commercial budgets, the accountable line manager defines the trade-off between missed events and false alarms. [S02] [S11]

Setting human thresholds translates statistical scores into unit quantities managed under ISO 22400-2 indicator definitions. [S11] [S38] An Area Under the Receiver Operating Characteristic curve (AUC) of 0.94 is not an operational threshold because a supervisor cannot staff around an abstract area score. [S38] [S1102] An actionable threshold states required recall at a maximum false-alarm rate per thousand linear metres of cloth, or maximum mean absolute percentage error alongside a cap on single-order under-predictions during peak shipping weeks, recorded in the model registry with its operating assumptions. [R29] [S11] [S15] [S38]

11.3 FOUR THRESHOLD FAMILIES PROTECT THE LINE FROM ONE-DIMENSIONAL OPTIMISATION

A model that clears a single accuracy hurdle can still degrade operations if its input feed is incomplete, its latency stalls a work cell, or its gains on high-volume styles are paid for by severe errors on specialty lots. [S04] [S19] [S1101] GFT structures pre-release acceptance around four mandatory threshold families: technical discrimination and calibration, operational net benefit, serving data fitness, and bounded subgroup behaviour. [S04] [S11] [S1101] Each candidate must pass all four families simultaneously. [S11] [S14]

Exhibit 46, p. 178 sets out the derived structure of these four threshold families, mapping each family to its primary metric class, required disaggregation dimensions, mandatory failure action and accountable owner. [S04] [S11] [S19] [S1101]

EXHIBIT 46 / GOVERNANCE SPECIFICATION

Four threshold families and slice rules prevent aggregate masking

Threshold family	Primary metric class	Required slice dimensions	Mandatory failure action	Accountable role
Technical performance	Recall, precision, calibration error	Defect class, fabric weight, weave structure	Block export; refine features or threshold	Model builder and model owner
Operational benefit	Net review minutes, first-pass yield, scrap delta	Line, shift, operator experience tier	Hold in shadow or reject candidate	Domain operating owner
Serving data fitness	Completeness, timeliness, valid range share	Capture device, lighting state, offline buffer	Quarantine feed; repair capture before re-test	Data steward and data custodian
Subgroup floor	Worst-slice error cap, minimum sample support	Low-volume style, night shift, degraded mode	Restrict operational scope or stop release	Risk owner and domain owner

DERIVED [S04] [S11] [S19] [S1101] Derived threshold framework combining human metric and threshold authority from S11, ML data-quality management from S04, upstream cascade awareness from S19, and disaggregated condition reporting from S1101.

As **Exhibit 46**, p. 178 illustrates, each threshold family triggers a specific remedy when a candidate falls short. [S04] [S11] A failure on serving data fitness sends the team back to capture hardware, lighting calibration or schema enforcement rather than hyperparameter tuning, preventing upstream data cascades. [S04] [S19] However, **Exhibit 46**, p. 178 is a governance contract rather than a universal numeric table: it cannot establish the exact numerical cut-offs for a specific loom or inspection table, which must come from local baseline records. [S11] [S38]

11.4

Local slices reveal aggregate success that hides failure

Aggregate evaluation scores weight every test record equally, so they are dominated by the factory's most frequent styles, best-lit day shifts and cleanest sensor feeds. [S04] [S1101] Mitchell and colleagues show through their model-card framework that benchmarked evaluation must be disaggregated across conditions, groups and intersectional combinations relevant to the application domain so unsuitable contexts are exposed before release. [S1101] On a production line, disaggregation prevents releasing a model that performs well on standard plain-weave greige cotton but fails on dyed twill, heavy canvas, humid afternoon shifts or frames running after a yarn-lot changeover. [S04] [S19] [S1101]

11.5 DISAGGREGATED EVALUATION TESTS THE OPERATING FACTORS THAT SHIFT ON THE FLOOR

Under ISO/IEC 5259-3, managing data quality for analytics and machine learning requires verifying that validation datasets represent actual operating conditions. [S04] When evaluating a candidate, a factory team partitions held-out validation and shadow datasets along six operational axes wherever applicable: [S04] [S38] [S1101]

1. Production line and machine unit, accounting for camera mounting angles, vibration profiles and mechanical wear across individual looms or sewing lines. [S04] [S1101]
2. Product style, construction and colourway, capturing how yarn count, weave density and dye shade alter visual contrast and lead-time elasticity. [R29] [S1101]
3. Shift and ambient environment, reflecting daylight leakage, voltage swings, dust and crew rotation between day and night shifts. [S04] [S19]
4. Defect or event category, separating rare high-cost failures (oil stains, warp floats, safety trips) from minor blemishes so accuracy on common classes cannot mask blindness to severe defects. [R29] [S11] [S1101]
5. Operating mode, distinguishing steady-state runs from start-up after changeover, rework and post-maintenance restarts. [S38] [S43]
6. Missing-data and degraded-feed states, testing behaviour when context fields arrive late from an offline buffer or carry estimated or missing quality markers. [S04] [S19]

Intersectional slices are equally important. [S1101] A vision candidate may pass recall thresholds across all shifts combined and across dark-dyed fabrics combined, yet fail when a dark-dyed style runs on Line B during the night shift under artificial lighting. [R29] [S1101] Pre-release slice tables that cross-tabulate style family by line and shift catch that failure before a customer rejects a lot. [S19] [S1101]

11.6 SAMPLE SUPPORT AND UNCERTAINTY BOUNDS PREVENT OVER-INTERPRETING SMALL SLICES

Disaggregating a validation set reduces the sample size in each cell. [S04] [S1101] If a specialty style appears only eight times in a snapshot, a single misclassified roll shifts the apparent error rate by 12.5 percentage points (an illustrative ratio of 1 divided by 8), risking either false confidence or premature rejection. [S04] [S11]

METHOD

Disaggregated slice validation and minimum support check

To evaluate local operating slices without mistaking sampling noise for proven fitness, the validation protocol executes four sequential checks on the frozen held-out snapshot and shadow log: [S04] [S11] [S15] [S1101]

1. **Define the mandatory slice grid.** Before scoring, the domain owner, data steward and builder list required unitary and two-way intersectional slices (such as `style_group x line_id` and `defect_severity x shift_code`), linked to governed identifiers in the KPI dictionary. [S02] [S38] [S1101]
2. **Enforce a minimum sample-support floor.** Record the exact denominator n_{slice} in each cell and apply an explicit support floor (such as an illustrative floor of $n_{\text{slice}} \ge 30$ independent events) below which a point score is marked *insufficient support for autonomous release*. [S04] [S11]
3. **Report uncertainty intervals alongside point scores.** Display a two-sided confidence or credible interval around each slice metric so the operating owner evaluates the worst-case plausible error bound. [S11] [S1101]
4. **Apply scoped release or targeted collection when support is low.** When a slice breaches its threshold at the conservative bound or lacks sample support, either extend shadow collection or issue a scoped release that excludes the under-supported slice and routes those cases to manual baseline handling. [S04] [S11] [S1102]

Checking sample support also exposes upstream data cascades: when a shift slice shows near-zero logged defects, the steward verifies whether cloth quality improved or operators stopped recording tags during busy changeovers. [S02] [S19]

11.7

Operational comparison includes the cost of being wrong

A candidate model is never evaluated against an implicit alternative of zero capability. [S43] [S1102] Even a manual factory has a working decision process with an observable yield, speed and error cost. [S36] [S38] Google's machine-learning engineering rules advise keeping the first model simple to establish baseline metrics and behaviour, and running sanity checks on held-out data right before exporting any model to serving. [S1102] In the CRISP-DM manufacturing lifecycle, the evaluation phase verifies that the model meets business objectives before deployment. [S43]

11.8 THE CANDIDATE MUST BEAT BOTH CURRENT PRACTICE AND A TRANSPARENT COMPARATOR

Every validation report places three options side by side on the same historical and shadow windows: [S38] [S43] [S1102]

- **Current operating practice (the baseline):** Actual human decisions, inspection catches or planning allocations recorded during the frozen window, scored under the governed ISO 22400-2 KPI formula. [S38]

- **A simple, transparent comparator:** A deterministic rule, moving average, seasonal naive forecast or single-threshold optical filter maintainable in a worksheet or programmable logic controller (PLC) block. [S43] [S1102]
- **The candidate model:** The versioned artefact registered under Chapter 10 reproducibility controls, evaluated with its proposed threshold and serving pipeline. [S15] [S16]

Industry literature in GFT's register notes that manual fabric defect inspection is 60 to 70 percent accurate and misses 20 to 30 percent of defects, whereas computer-vision benchmarks exceed 90 percent accuracy across more than 40 defect types at speeds up to 60 metres a minute (inherited Grade-C entry, not refreshed in 2026 and original dates unrecorded). [R29] Those figures justify a pilot, but local lint, cotton-seed trash, vibration and weave variations can push a vendor model's precision below a trained two-person inspection perch. [R29] [S04] [S19] Only local comparison settles the question. [S43] [S1102]

11.9 REVIEW BURDEN, FALSE INTERVENTIONS AND LATENCY ENTER THE NET ACCEPTANCE BALANCE

A statistical gain that slows the line or exhausts operators is an operational regression. [S22] [S38] [S43] If a vision system catches four additional true defects per shift but triggers forty false stops requiring three minutes each, it consumes 120 minutes of inspection capacity per shift (an illustrative product of 40 stops and 3 minutes) and drives operators to ignore alerts. [S18] [S19] [S38]

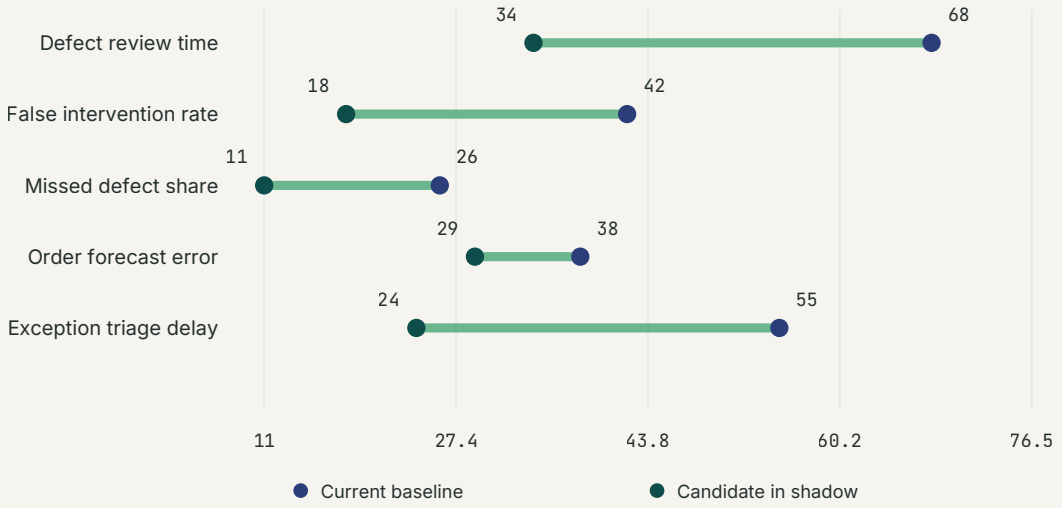
Operational validation therefore measures five quantities alongside statistical error: [S22] [S38] [S43]

1. **Inference and alert latency (mean and 95th percentile)** within physical equipment cycle times and operational technology (OT) reliability limits. [S22] [S38]
2. **False-intervention rate and review minutes per shift** consumed when human verification overturns an alert. [S38] [S43]
3. **Missed-event downstream cost** in rework, scrap or stockout penalties compared with baseline practice. [S38] [S43]
4. **Exception-triage delay**, testing whether bounding boxes or attributions speed up operator verification. [S11] [S38]
5. **Infrastructure and maintenance overhead** required from a small team to sustain calibration and pipelines. [S18] [S42]

Exhibit 47, p. 182 illustrates how a candidate model is compared against current baseline practice across five normalised operational burden and error indicators during shadow validation. [S38] [S43] [S1102]

EXHIBIT 47 / ILLUSTRATIVE INDEX (0 TO 100, LOWER IS BETTER)

Candidate improvement must survive comparison with current practice



ILLUSTRATIVE [S38] [S43] [S1102] Illustrative shadow-validation comparison across five operational indicators. Values are constructed to demonstrate paired pre-release evaluation under S38 KPI definitions, S43 lifecycle evaluation and S1102 baseline sanity checks, not measured results from a GFT facility.

In **Exhibit 47**, p. 182, lower index values represent less wasted time, fewer false stops or lower error over the same shift population. [S38] [S43] If a candidate cut missed defect share from 26 to 11 on the illustrative index while raising false intervention rate from 42 to 75, the chart would expose an unacceptable trade-off hidden by a composite F1-score. [S11] [S43] Because **Exhibit 47**, p. 182 is an illustrative demonstration, it cannot establish actual factory gains; live decisions require measured units from a frozen shadow run. [S38] [S43]

WORKED EXAMPLE

Illustrative net-benefit validation of a greige-fabric inspection candidate

Consider an illustrative validation test at a woven-cotton inspection table in Cameroon (all figures are illustrative assumptions to demonstrate the calculation method, not observed GFT results; wage and currency context references the inherited register where Cameroon's statutory minimum wage rose to XAF 43,969 a month from 23 February 2024 and the CFA franc is pegged to the euro at 655.957). [R14] [R32]

- **Assumed shift boundary and inputs (illustrative):** One perch inspects 100 greige rolls of 100 metres each (10,000 linear metres) per 480-minute shift containing 50 true major defects. Each undetected defect reaching dyeing and cutting causes an illustrative rework and downgrade loss of XAF 18,000. Each flagged stop requires an illustrative 2.5 minutes of perch time, valued at XAF 400 per minute (XAF 1,000 per stop). [S38]
- **Option 1: Current manual baseline (illustrative, informed by the 60 to 70 percent accuracy range in R29):** Catches 64 percent of the 50 defects (32 caught, 18 missed) with 6 false stops (38 total stops). Stop time is $38 \times 2.5 = 95.0\text{ minutes}$ ($\text{XAF } 38,000$), and missed-defect loss is $18 \times \text{XAF } 18,000 = \text{XAF } 324,000$, giving an illustrative baseline cost of $\text{XAF } 362,000$ per shift. [R29] [S38]
- **Option 2: Candidate Model A with uncalibrated threshold (illustrative):** Catches 92 percent of defects (46 caught, 4 missed), but lint triggers 140 false stops (186 total stops). Stop time reaches $186 \times 2.5 = 465.0\text{ minutes}$ out of 480 minutes ($\text{XAF } 186,000$) plus $\text{XAF } 72,000$ in missed-defect loss (illustrative total $\text{XAF } 258,000$). Because 465 stop minutes breach the 120-minute review ceiling and collapse throughput, Candidate A is **rejected**. [S11] [S38]
- **Option 3: Candidate Model B after calibration and lint filtering (illustrative):** Catches 88 percent of defects (44 caught, 6 missed) with 14 false stops (58 total stops). Bounding-box coordinates cut stop duration from 2.5 to an illustrative 1.8 minutes ($58 \times 1.8 = 104.4\text{ minutes}$, costing $\text{XAF } 41,760$), plus $6 \times \text{XAF } 18,000 = \text{XAF } 108,000$ in missed-defect loss. Total illustrative cost is $\text{XAF } 149,760$ per shift, saving $\text{XAF } 212,240$ per shift against baseline while keeping stop time below 120 minutes. Candidate B **passes comparison**. [S11] [S38] [S43]

11.10

Shadow and assisted modes earn autonomy gradually

Switching a live process directly from manual execution to closed-loop model control introduces unnecessary operational and safety risk. [S16] [S22] Offline snapshots cannot replicate live camera fouling, network micro-outages, clock drift or end-of-month dispatch pressure. [S19] [S22] [S42] Microsoft's MLOps maturity model advises building capability gradually rather than attempting full automation upfront, while NIST SP 800-82 Revision 3 prioritises OT performance, reliability and safety. [S16] [S22] Combining these rules creates a staged release ladder. [S16] [S22]

11.11 SILENT SHADOW OPERATION PROVES PIPELINE INTEGRITY AND LIVE DISTRIBUTION FITNESS

Every new deployment or major version upgrade begins in **silent shadow mode**: the model ingests live events or frames, computes features and predictions, and logs timestamped outputs and latency, while its outputs remain hidden from operators and cannot trigger machine or workflow actions. [S15] [S16] [S22]

Shadow operation tests four live questions: [S04] [S15] [S22] [S1102]

1. **Training-serving parity:** Verifying that live messages and edge tables compute the exact same feature values and unit conversions as historical training extracts. [S04] [S15] [S1102]
2. **Runtime resource safety:** Confirming that peak inference load on edge hardware respects processor, memory and thermal margins under NIST SP 800-82 reliability rules. [S22]
3. **Live slice accuracy against delayed ground truth:** Joining shadow predictions to completed inspection sheets or order records via shared lot identifiers across all active slices. [S04] [S09] [S15]
4. **Virtual alert volume:** Counting hourly alerts that would have fired so the operating owner sees true operator review burden before anyone is interrupted. [S11] [S38]

11.12 OPERATOR-ASSIST MODE TESTS HUMAN-MACHINE INTERACTION BEFORE ANY BOUNDED AUTOMATION

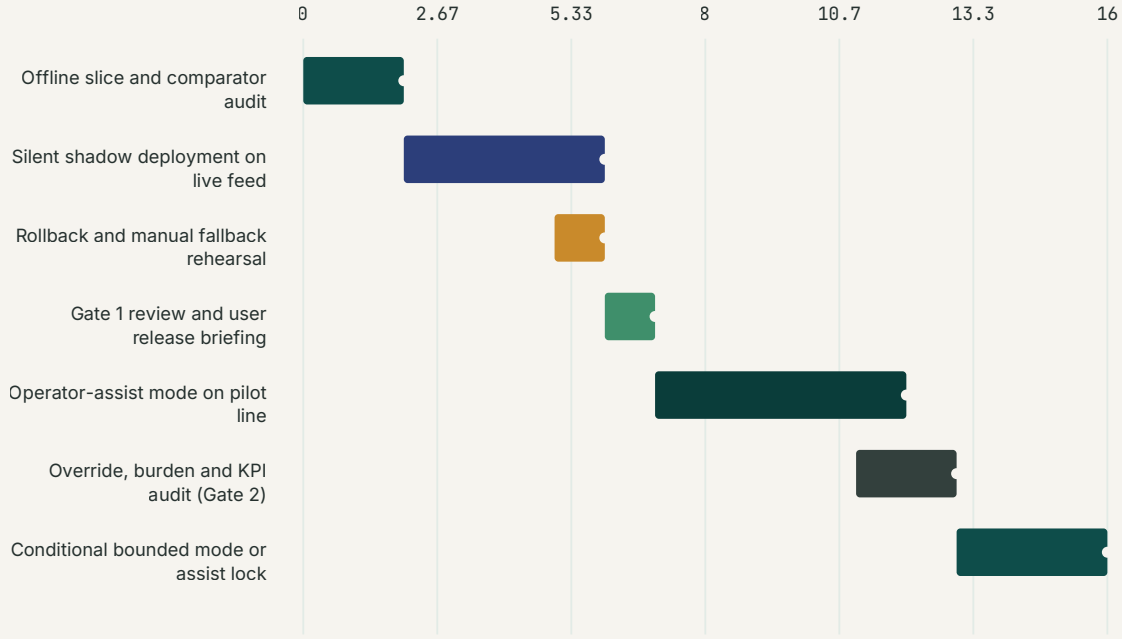
Once the shadow log passes all four threshold families over a predefined window (such as an illustrative window of two to four weekly cycles across all shifts and main style families), the operating owner may promote the candidate to **operator-assist mode**, where the model displays recommendations or bounding boxes while the operator retains authority to confirm, modify or reject each action. [S11] [S16] [S30]

Advancing from shadow to operator-assist mode is an explicit gate decision, and advancing further to **bounded automation** (acting automatically only within narrow pre-approved limits) is optional. [S16] [S22] [S42] Decisions involving complex rework trade-offs, customer concessions or safety-adjacent settings remain permanently in operator-assist mode. [S11] [S22] [S30]

Exhibit 48, p. 185 shows an illustrative 16-week staged release sequence from candidate freeze through shadow validation, rollback rehearsal, operator-assist operation and conditional bounded automation. [S15] [S16] [S22]

EXHIBIT 48 / ILLUSTRATIVE WEEKS FROM CANDIDATE FREEZE

Staged release makes each expansion of authority conditional



ILLUSTRATIVE [S15] [S16] [S22] Illustrative 16-week staged release schedule derived from S15 pipeline lineage and rollback controls, S16 incremental MLOps maturity, and S22 OT reliability constraints. Durations are planning illustrations, not fixed commitments.

As **Exhibit 48**, p. 185 shows, the rollback rehearsal occurs in Week 5 to Week 6 before Gate 1 approves operator-assist activation in Week 7, and Gate 2 in Week 11 to Week 13 audits human overrides and KPI movement before any expansion. [S11] [S15] [S27] [S38] Because **Exhibit 48**, p. 185 is an illustrative schedule, it cannot fix calendar durations across use cases; monthly planning models require longer shadow windows than high-speed inspection cameras. [S16] [S38] [S43]

11.13

Rollback and manual fallback are release artefacts

A model release is safe only if the factory can undo it faster than a bad prediction series can damage production. [S15] [S22] When deployment overwrites scripts in place and removes paper check-sheets, a bumped camera lens or corrupted feature leaves the floor unable to restore operations without hours of downtime. [S18] [S19] [S22] GFT treats technical **rollback** (reverting to prior approved version v_{k-1}) and operational **manual fallback** (reverting to the non-model baseline workflow) as mandatory release artefacts built and exercised before approval. [S15] [S22] [S27]

11.14 IMMUTABLE LINEAGE POINTERS ALLOW IMMEDIATE REVERSION TO A KNOWN STATE

Under Google Cloud's MLOps pipeline architecture, recording execution metadata and artefacts enables lineage, reproducibility and rollback. [S15] Rather than loading an unversioned file path, the serving layer resolves a registry pointer binding four immutable artefacts: [S09] [S15]

- The **model artefact digest and version identifier**, keeping prior approved version v_{k-1} staged locally. [S15]
- The **feature-transformation code version**, reverting normalisation constants alongside the model. [S15]
- The **input data contract and schema version**, matching the structure and units validated for v_{k-1}. [S04] [S15]
- The **threshold configuration**, restoring the human-approved cut-offs of version v_{k-1}. [S11] [S15]

The previous approved bundle remains staged on the local node and in offline encrypted backup under CISA recovery guidance for at least one review cycle, making rollback a rapid pointer switch. [S15] [S26]

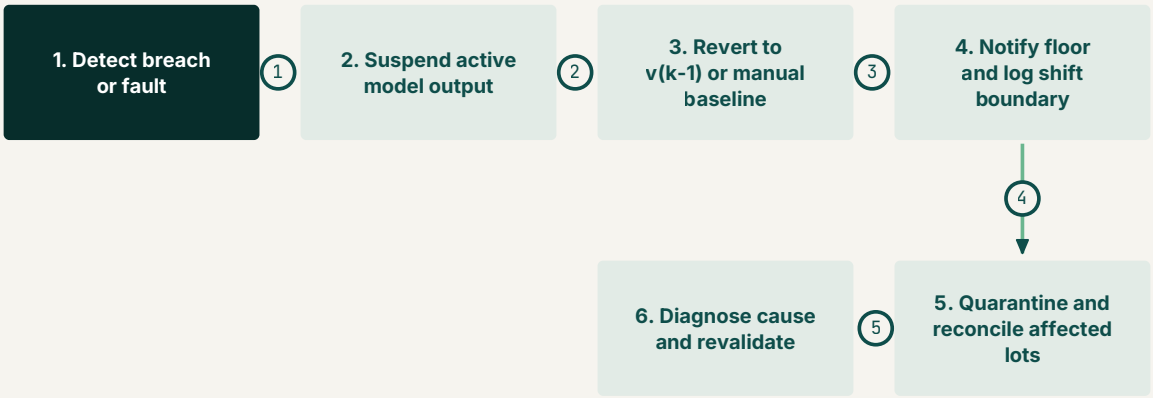
11.15 REHEARSED FALLBACK DRILLS PROVE THAT OPERATORS CAN RESUME BASELINE WORK ON COMMAND

When an upstream sensor fails or a first-time model is suspended, the line must drop back cleanly to **manual or deterministic baseline operation**. [S22] [S27] Following NIST SP 800-61 Revision 3 guidance on periodic incident exercises, a pre-release fallback drill verifies that workflows survive suspension without lost records. [S22] [S27]

Exhibit 49, p. 187 traces the derived six-stage control path from anomaly detection through suspension, reversion, floor communication, lot reconciliation and investigation. [S15] [S22] [S27]

EXHIBIT 49 / OPERATIONAL CONTROL SEQUENCE

Rollback restores a known operating state and preserves evidence



1. Threshold or steward trigger
2. Registry pointer or SOP switch
3. Confirm stable line state
4. Tag interval in event log
5. Preserve failed snapshot

DERIVED [S15] [S22] [S27] Derived rollback and fallback sequence combining S15 artefact and pipeline lineage, S22 OT safety and reliability preservation, and S27 practised incident-response procedures.

In **Exhibit 49**, p. 187, suspending the model (Steps 2 and 3) must never leave an OT controller or terminal in an undefined state, and Step 4 logs the exact timestamp and lot boundary so KPI calculations separate model-assisted intervals from fallback intervals. [S09] [S22] [S38] Step 5 reconciles lots processed just before suspension. [S04] [S27] Because **Exhibit 49**, p. 187 is a process specification, it cannot guarantee recovery times unless timed physically on the line before approval. [S22] [S27]

> **WARNING**
>
> **Untested fallback paths fail when the baseline workflow has been dismantled**
>
> If a model runs for several months and the factory stops maintaining manual inspection sheets or training new hires on baseline rules, a later hardware fault or drift suspension leaves operators unable to keep the line moving. [S18] [S27] [S35] Every release package must name the baseline Standard Operating Procedure (SOP), keep copies at the workstation, and include manual fallback in periodic shift drills. [S27] [S35]

11.16

Approval separates building from acceptance

When the analyst or vendor who trained a model also decides whether it is ready for production, optimism bias enters the gate. [S11] [S14] Both ISO/IEC 23894 and the NIST Cybersecurity Framework 2.0 place risk acceptance with accountable operational leadership rather than technical builders alone. [S14] [S23]

11.17 THREE DISTINCT SIGN-OFFS ALIGN TECHNICAL EVIDENCE, LINE ACCOUNTABILITY AND OT SAFETY

The GFT data spine divides release authority into three functional sign-offs: [S02] [S14] [S22] [S23]

1. **Builder and data-steward verification:** Certifies reproducible artefacts, complete registry lineage, passed shadow skew checks, and unmanipulated slice and uncertainty tables. [S02] [S04] [S15]
2. **Operating-owner acceptance:** The domain owner accountable for affected KPIs reviews slice results, comparator charts and shift review burden, certifying operational fitness. [S02] [S11] [S38]
3. **Risk, OT and security approval:** Verifies least-privilege identities, NIST SP 800-82 OT safety boundaries, worker-data protections, and timed fallback drill completion. [S22] [S25] [S27] [S30]

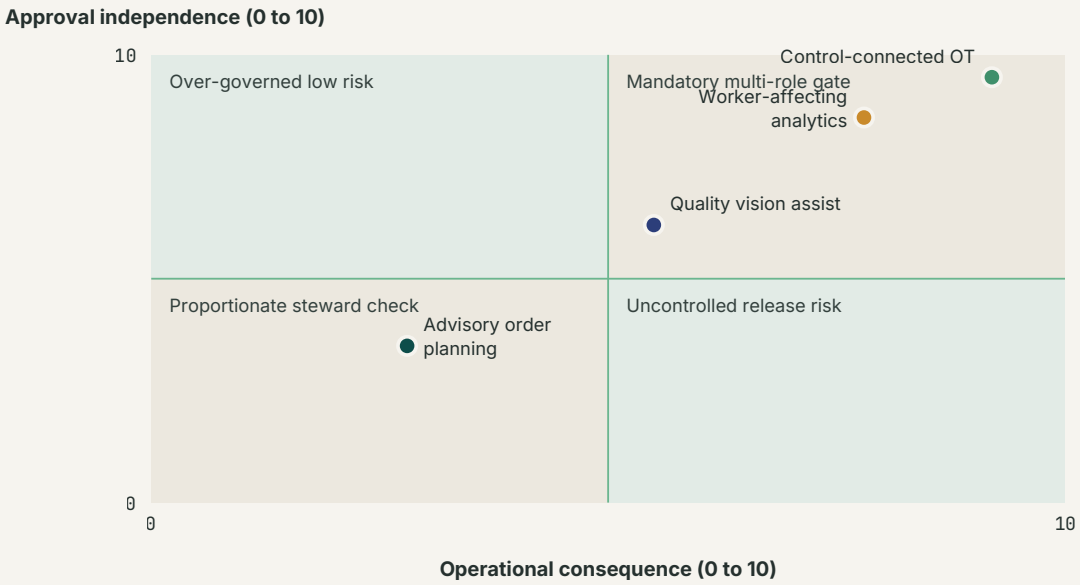
11.18 SMALL TEAMS PREVENT SELF-APPROVAL THROUGH ROLE SEPARATION AND PROPORTIONATE ESCALATION

A small Cameroonian spinning, weaving or garment plant needs no separate model-risk department to prevent self-approval: the analyst or vendor prepares the evidence package, while the quality or production manager (operating owner) and plant or maintenance manager (OT and risk approver) hold veto and sign-off rights. [S02] [S11] [S14] [S16] [S42]

Exhibit 50, p. 189 maps four factory use-case classes across operational consequence and required approval independence. [S11] [S14] [S22] [S23] [S30]

EXHIBIT 50 / DERIVED GOVERNANCE SCORE (0 TO 10 SCALE)

Higher consequence demands more independent release authority



DERIVED [S11] [S14] [S22] [S23] [S30] Derived consequence-to-authority matrix based on S11 human threshold authority, S14 integrated AI risk management, S22 OT safety constraints, S23 ERM governance, and S30 worker-data protection. Coordinates represent ordinal governance tiers rather than empirical probabilities.

As **Exhibit 50**, p. 189 shows, no deployment is permitted in the lower-right quadrant (*Uncontrolled release risk*). [S11] [S14] Advisory planning (x = 2.8, y = 3.5) requires steward verification and operating-owner sign-off, whereas worker-affecting analytics (x = 7.8, y = 8.6) and control-connected OT (x = 9.2, y = 9.5) sit in the upper-right quadrant (*Mandatory multi-role gate*) with strict builder-approver separation and safety interlocks. [S11] [S22] [S30] These ordinal coordinates illustrate governance tiers and cannot replace a machine-specific hazard assessment. [S14] [S22]

11.19

Worker-affecting outputs retain meaningful human review

Factory data spines touch human work whenever logs describe identifiable operators. [S29] [S30] Alongside GDPR Recital 39 data minimisation and the statutory prior-authorisation requirement in Section 19(1) of Cameroon's Law No. 2024/017 of 23 December 2024 (requiring local legal counsel for interpretation), Principle 5.5 of the International Labour Office (ILO) code of practice states that decisions concerning a worker should not be based solely on automated processing of that worker's personal data. [S29] [S30] [S31] GFT enforces this non-binding ILO principle as a mandatory release gate. [S11] [S30]

11.20 NOMINAL OVERRIDE BUTTONS FAIL WITHOUT TIME, CONTEXT AND AUTHORITY TO DISAGREE

An "Approve / Override" button is not meaningful review if a supervisor must clear fifty scores in three minutes without seeing machine-fault context or faces pressure not to override. [S11] [S19] [S30] Validation of any worker-affecting candidate must verify four conditions before release: [S11] [S29] [S30]

1. **Aggregation by default:** Model at machine, line, style or shift grain rather than individual worker grain wherever possible. [S29]
2. **Visible context and uncertainty:** Show measured inputs, missing-data states and confounders such as yarn-lot breaks or machine faults. [S04] [S11] [S38]
3. **Time and authority to reject:** Allocate shift review time and authorise supervisors to reject recommendations without penalty. [S02] [S11] [S30]
4. **Right of contestation and correction:** Provide a clear route through the steward or human-resources lead to correct erroneous logs via superseding provenance events. [S09] [S30]

11.21 DECISION LOGS AUDIT WHETHER AUTOMATED RECOMMENDATIONS BECOME DE FACTO COMMANDS

The operator-assist interface logs each recommendation, reviewer action, reason code and review duration under least-privilege access controls. [S15] [S25] [S29] Gate 2 and monthly reviews audit overrides: near-100 percent overrides signal poor utility, while 0 percent overrides with two-second review times warn of mechanical rubber-stamping. [S11] [S30] [S36]

11.22

The release package tells the floor what changed

A model release recorded only in a code repository leaves operators uninformed. [S12] [S13] Translating ISO/IEC 42001:2023 management-system requirements and ISO/IEC 5338:2023 lifecycle processes into shop-floor practice requires publishing a concise **release package** alongside every promoted version. [S12] [S13] [S1101]

11.23 A COMPACT RELEASE MANIFEST DOCUMENTS SCOPE, LIMITS, ALERTS AND FALLBACK ROUTES

Combining ISO/IEC 42001, ISO/IEC 5338 and Mitchell and colleagues' model cards, the release package pairs a registry manifest with a one-page workstation sheet containing eight elements: [S11] [S12] [S13] [S15] [S1101]

- **Release identity and lineage links:** Model version ID, dataset snapshot IDs, feature-pipeline version, KPI dictionary version and activation timestamp. [S15] [S38]
- **Intended operational decision and scope:** Decision supported, approved lines and styles, and authority mode (shadow, operator-assist or bounded automation). [S12] [S16] [S1101]
- **Out-of-scope conditions and limitations:** Styles, fabric weights, lighting states or modes that failed slice support or threshold gates. [S04] [S11] [S1101]

- **Validated performance summary:** Baseline, comparator and candidate results across the four threshold families and local slices, with sample support and uncertainty bounds. [S11] [S38] [S1101]
- **User instructions and review expectations:** How operators inspect context, confirm actions and log overrides. [S12] [S30] [S35]
- **Alert route and escalation owner:** Recipients of automated quality or drift alerts and the steward contact for abnormal behaviour. [S02] [S17] [S27]
- **Rollback and manual fallback instruction:** Switch to suspend the model, pointer to prior version v_{k-1}, and workstation manual SOP reference. [S15] [S22] [S27]
- **Named approvers and review date:** Signatures of the builder/steward, operating owner and risk/OT approver, plus the scheduled review date. [S12] [S13] [S14]

11.24 SHIFT-ALIGNED ACTIVATION AND USER BRIEFING PRESERVE KPI INTERPRETABILITY

In OECD case-study evidence on SME digitalisation, converting traditional processes into digital routines requires structured learning rather than yielding an instant payoff, and ILO-UNESCO guidance emphasises tangible pilots with built-in feedback and user preparation. [S35] [S40] Model activation is aligned to a clean shift, batch or planning boundary and recorded as a provenance event, preceded by a workstation briefing from the steward and adoption champion covering shadow examples, override logging and the manual fallback sheet. [S09] [S27] [S35] [S37] [S38]

GFT NOTE

Applying staged release to GFT's Cameroon textile programme sequence

In GFT's Cameroon textile programme (prepared 31 January 2024, sources refreshed 2026), shop-floor capture and KPI governance come first, purchased vision inspection is evaluated early where lighting and image capture are controlled, and demand forecasting waits until the spine accumulates clean order and sales history. [R06] [R29] [S16] [S44]

Two practical rules govern release across that sequence: [S04] [S11] [S16]

1. **Bought computer-vision inspection (`wp03` link):** Never treat a vendor's laboratory benchmark (>90 percent accuracy in industry literature) as local release approval. [R29] Run the camera in silent shadow mode above the inspection perch across greige, dyed and printed cotton lots, disaggregate by weave, dye shade and lighting, price false-stop minutes against downstream scrap in XAF, rehearse reverting to the manual SOP, and release in operator-assist mode where the inspector confirms every cut or tag. [R29] [S04] [S11] [S16]
2. **Order and demand forecasting (`wp04` link):** Never promote a forecaster until it beats both the planner's worksheet and a seasonal or moving-average comparator on held-out local cycles under frozen ISO 22400-2 definitions. [S38] [S43] [S1102] Where order records lack sample support across styles or omit cancelled orders, postpone release and fix capture first. [S19] [S44]

With the release package active, the candidate operates under controlled authority on the floor. Yet validation at release is only a snapshot: cotton fibre properties vary across harvests, dye lots shift, cameras accumulate lint, and ordering patterns change. [R01] [S18] [S21] Keeping the deployed model subordinate to operations over time requires continuous production monitoring, drift diagnosis and rehearsed incident response, examined next. [S17] [S21] [S27]

A model earns its place on the line not when it predicts well in the lab, but when operators can verify its gains, bound its errors and switch it off without losing a shift.

CHAPTER SOURCE NOTES

S11 National Institute of Standards and Technology. 2023-01-26.

S15 Google Cloud. 2024-08-28.

S43 Frontiers in Artificial Intelligence. 2021-06-14.

S19 Association for Computing Machinery, CHI 2021. 2021.

S1101 arXiv, Conference on Fairness, Accountability, and Transparency. 2019-01-14.

R06 Cameroon Trade Hub. Publication date unrecorded; inherited source-text check.

R07 Ecofin Agency. Publication date unrecorded; inherited source-text check.

S04 ISO/IEC. 2024.

S16 Microsoft Learn, Azure Architecture Center. 2026-06-03.

S22 National Institute of Standards and Technology. 2023.

S09 World Wide Web Consortium. 2013-04-30.

S38 International Organization for Standardization. 2014.

S13 ISO/IEC. 2023.

S1102 Google Developers. 2025-08-25.

S14 ISO/IEC. 2023.

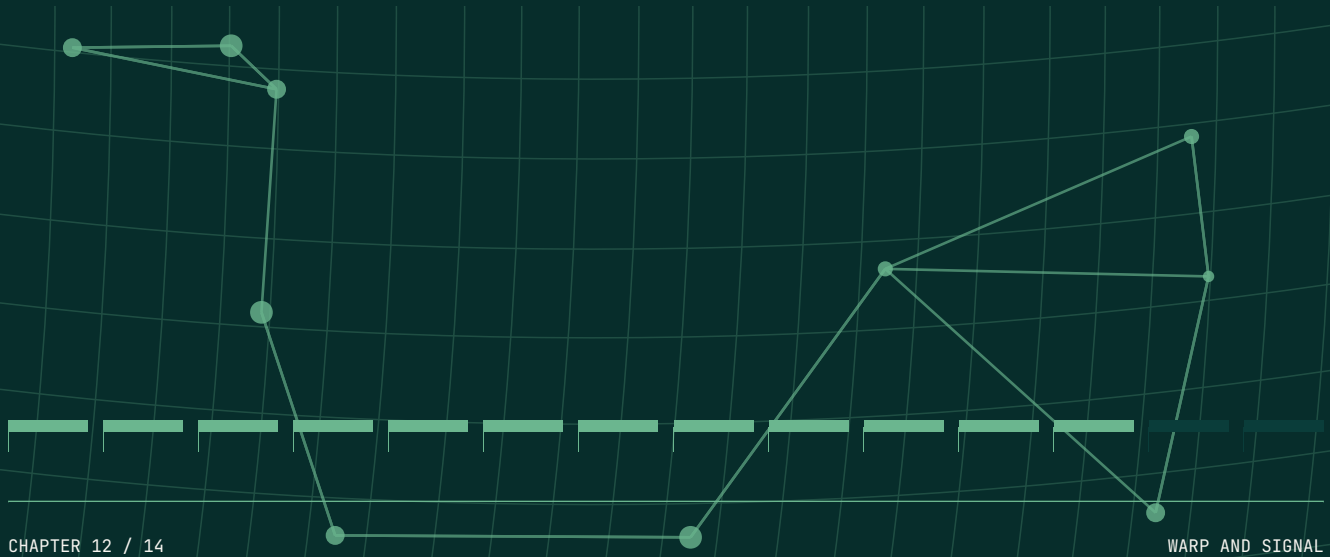
S02 International Organization for Standardization. 2022.
R29 MDPI; industry inspection literature. Publication date unrecorded; inherited source-text check.
S36 World Bank. 2023.
S18 Advances in Neural Information Processing Systems. 2015.
S42 SCITEPRESS, IN4PL 2022. 2022.
R14 WageIndicator. Publication date unrecorded; inherited source-text check.
R32 Wise; XE. Publication date unrecorded; inherited source-text check.
S30 International Labour Office. 1997.
S27 National Institute of Standards and Technology. 2025.
S26 CISA, MS-ISAC, NSA and FBI. 2023.
S35 OECD. 2022-11-21.
S23 National Institute of Standards and Technology. 2024-02-26.
S25 National Institute of Standards and Technology. 2020-12-10.
S29 European Parliament and Council of the European Union. 2016-04-27.
S31 Presidency of the Republic of Cameroon. 2024-12-23.
S12 ISO/IEC. 2023.
S17 Microsoft Learn, Azure Machine Learning. 2026-01-27.
S40 International Labour Organization and UNESCO. 2020.
S37 World Bank. 2025-05-20.
S44 The Journal of Technology Transfer, Springer Nature. 2024-07-27.
R01 Business in Cameroon. Publication date unrecorded; inherited source-text check.
S21 arXiv, author manuscript of an IEEE Transactions on Knowledge and Data Engineering article. 2020-04-13.

12

PART III: CONTROL MODELS AS PRODUCTION CHANGE / PRODUCTION MONITORING

Monitoring must trigger diagnosis, not automatic retraining

A useful monitor detects change, establishes whether it matters and routes a named person through a rehearsed response.



CHAPTER 12
DECISION TEST

A useful monitor detects change, establishes whether it matters and routes a named person through a rehearsed response.

Production monitoring exists to protect factory decisions after a model leaves validation, not to keep an automated retraining loop busy. [S17] In a small or mid-sized manufacturer, the most common causes of an apparent model anomaly are broken data feeds, unannounced schema or unit mutations, delayed outcome labels, and planned shifts in product mix or raw material lots. [S19] When a team wires drift alerts directly into an unattended retraining scheduler, it risks fitting a new model to corrupted inputs, masking repairable capture faults, and promoting an unverified artefact without human challenge. [S13] Green Future Tech Limited (GFT), Hong Kong SAR, designs its Cameroon textile programme (prepared 31 January 2024, with source-refresh notes retained in 2026) around the opposite rule: a production monitor is a diagnostic instrument that opens an owned investigation, preserves safe fallback, and requires fresh lifecycle approval before any replacement model reaches the shop floor. [S21]

That discipline matters because Cameroon's industrialisation ambition leaves little margin for opaque automation failures. In the inherited programme register, about 95 percent of Cameroon's cotton fibre is exported raw and about 5 percent is processed locally, against a national target of 50 percent local processing by 2030 (sector observations whose original publication dates were not recorded in the inherited annex and were not refreshed in 2026, rather than GFT operating results). [R06] [R07] Moving a mill from raw-fibre export to competitive yarn, fabric and garment production requires stable daily control over yield, quality and order fulfilment. [S38] Whether a factory deploys an early bought computer-vision inspector covered in companion paper wp03 or a later demand-forecasting model covered in companion paper wp04, post-release monitoring must separate transport and data-quality failures from genuine changes in the production process. [S04]

12.1

Service health and data quality come before accuracy

A model cannot be judged for predictive accuracy until the factory confirms that the serving pipeline is alive and the incoming records obey their governed contracts. [S04] Machine-learning systems fail silently when input pipelines degrade: instead of throwing a visible runtime exception, a model presented with stale, zero-filled, or unit-shifted features continues to emit plausible numeric scores. [S1201] If the monitoring layer watches only aggregate prediction drift or waits weeks for confirmed accuracy metrics, operators act on bad advice while the underlying ingestion fault goes unnoticed. [S17] Production monitoring therefore begins at the bottom of the industrial data spine, checking service health and input fitness before evaluating statistical model behaviour. [S15]

12.2 FIVE MONITORING LAYERS SEPARATE PIPELINE FAILURE FROM MODEL DECAY

To prevent engineers from debugging model weights when a sensor cable or message broker has failed, a small factory team should organise production checks in an explicit diagnostic hierarchy. [S17] Each layer tests a specific boundary of the spine and halts downstream interpretation when its precondition fails: [S04]

1. **Transport and serving availability:** Verify that edge devices, message brokers, feature jobs, and model endpoints are reachable, completing executions within their scheduled windows, and returning responses within the latency ceiling required by the work cell. [S15] During a network interruption, confirm that local store-and-forward buffers are accumulating events rather than dropping messages silently. [S33]
2. **Capture completeness and timeliness:** Measure record arrival rates by line, shift, and asset against expected production schedules, separating event time from ingestion time so a delayed batch upload is not misread as an idle machine or a sudden drop in defect frequency. [S08]
3. **Contract and schema validity:** Enforce the governed data contract at ingestion, checking field presence, data types, allowed taxonomy codes, identifier integrity, and engineering units before features are computed. [S04]
4. **Feature range and distribution sanity:** Track null shares, constant-value streaks, out-of-range physical readings, and sudden shifts in individual model inputs using threshold rules defined in the monitoring configuration. [S17]
5. **Prediction behaviour and outcome alignment:** Inspect the distribution of model scores, operator override rates, interim proxy indicators, and, once ground truth matures, realised technical and operational performance against the frozen baseline. [S1202]

12.3 BROKEN FEEDS MUST NEVER BE CONCEALED BY RETRAINING

Distinguishing the first three layers from the fifth is an essential operational safeguard. [S04] Suppose a humidity transmitter in a spinning or weaving shed loses calibration, or an enterprise software update alters a unit of measure from kilograms to tonnes without updating the downstream feature view. [S19] A drift detector watching feature distributions will rightly flag a sharp statistical departure from the training reference. [S17] If the factory's response is to trigger an automated retraining script over the recent window, the new candidate model will adjust its internal parameters to accommodate the broken sensor or rescaled unit, hiding the capture defect inside a newly fitted set of weights. [S18] When a maintenance technician later repairs the transmitter or a data custodian restores the unit mapping, the retrained model will immediately fail again. [S19] By placing service health and data-quality gates ahead of drift interpretation, the factory routes ingestion faults to the data steward and custodian for repair while holding the model constant or reverting to manual baseline rules. [S02]

12.4

Input drift and concept drift require different interpretation

Once the factory verifies that incoming records are complete, timely, and contract-compliant, the monitoring system must help the team interpret changes in the statistical behaviour of inputs and outputs. [S17] In streaming and operational settings, concept drift describes unforeseeable changes in the underlying distribution of data over time. [S21] In manufacturing practice, however, not every distribution change breaks a model, and not every model failure is accompanied by a dramatic shift in marginal input distributions. [S1201] Treating all statistical distance metrics as equivalent retraining triggers wastes engineering capacity and destabilises production control. [S1202]

12.5 COVARIATE SHIFT, PRIOR SHIFT AND CONCEPT DRIFT HAVE DISTINCT SHOP-FLOOR CAUSES

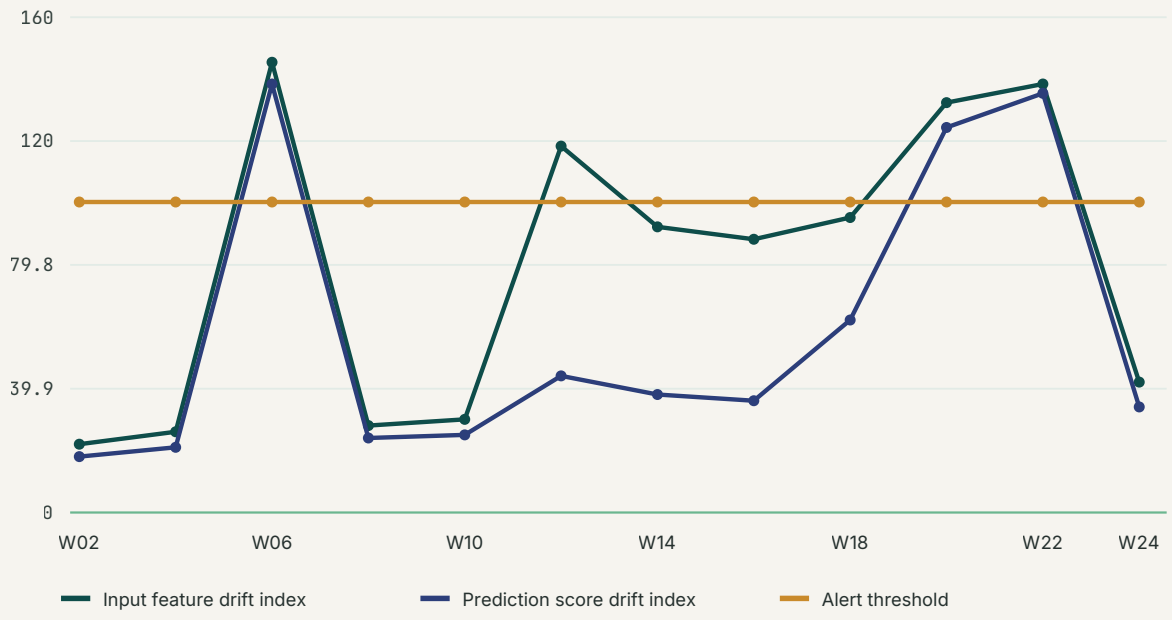
A diagnostic monitor separates three statistical phenomena because each points to a different physical or commercial event in the mill: [S21]

- **Input or covariate shift:** The distribution of input features changes, while the true physical relationship between those inputs and the target outcome remains stable. [S21] In a garment plant, switching a cutting line from lightweight woven shirting to heavy cotton twill changes fabric-weight and cutter-load features immediately, yet a well-specified model whose training scope covered both fabric classes may remain accurate. [S1203] Conversely, if the new style falls outside the model's validated domain, predictions become unreliable even though no sensor is broken. [S04]
- **Prior or label-distribution shift:** The base rate of the outcome changes across shifts, seasons, or suppliers, even when the appearance of a given defect or demand signal stays the same. [S1201] For example, processing a lower-grade seed-cotton lot during a wet harvest month may raise the true incidence of yarn slubs or fabric nep defects without changing what a slub looks like under a line-scan camera. [R29]
- **True concept drift:** The conditional relationship between inputs and the target outcome changes over time. [S21] If a dye house introduces a new reactive dye chemistry, or if customer purchasing behaviour changes after a wholesale pricing revision, the same input features now map to different physical or commercial outcomes, invalidating the model's learned function. [S21]

As shown in **Exhibit 51**, p. 198, tracking a feature drift index alongside a prediction drift index over an illustrative 24-week operating window reveals why statistical movement requires contextual diagnosis rather than a reflex model update. [S17] [S21]

EXHIBIT 51 / ILLUSTRATIVE DRIFT INDEX (ALERT THRESHOLD = 100)

Drift is a diagnostic signal, not a retraining instruction



ILLUSTRATIVE [S17] [S21] Illustrative 24-week bi-weekly monitoring trace for a factory model. Week 6 shows a schema unit error repaired at the source; Week 12 shows a planned product-mix change requiring slice validation rather than retraining; Week 20 shows sustained concept drift after a process change that justifies a governed candidate update. Values are invented to demonstrate the diagnostic method, not measured factory performance.

In **Exhibit 51**, p. 198, the illustrative monitoring trace crosses the alert threshold of 100 three times across 24 weeks, yet only the third excursion warrants retraining. [S17] At Week 6, both feature drift and prediction drift spike to illustrative indices of 145 and 138 because an upstream PLC tag scaling change corrupts a feed; repairing the ingestion mapping restores normal operation by Week 8 without touching the model. [S04] At Week 12, a scheduled changeover to a heavier fabric style pushes the illustrative feature drift index to 118, while prediction drift remains bounded at 44 and slice validation confirms that error rates on the heavier style stay within acceptance limits. [S1203] Only at Week 20, when a finishing-process modification alters both feature and prediction distributions and degrades verified inspection accuracy through Week 22, does the team authorise a governed retraining cycle that returns the indices to baseline by Week 24. [S13] This illustrative sequence demonstrates the diagnostic method; it cannot establish universal drift thresholds, which must be calibrated by human judgement on each site's historical variance. [S11]

12.6 OPERATIONAL MIX MUST BE LOGGED ALONGSIDE STATISTICAL DISTANCE

A statistical distance metric cannot tell the model owner whether a shift is benign or malignant unless the spine links each prediction back to its operational context. [S1201] Because every prediction event in the governed spine carries site, line, asset, style, lot, and shift identifiers under the visibility-event contract established earlier in the paper, the monitoring job can decompose aggregate drift by operational slice. [S08] Before concluding that a model has decayed, the analyst checks whether the mix of styles, yarn counts, or machine states changed during the scoring window. [S38] As

summarised in **Exhibit 52**, p. 199, six common factory anomalies all produce visible movement in monitoring charts, but five of them are harmed or left unresolved by automatic retraining. [S21]

EXHIBIT 52 / DERIVED DIAGNOSTIC CLASSIFICATION

Six production anomalies produce statistical shift but require opposite remedies

Observed anomaly	Primary signature	First diagnostic check	Correct remedy
Broken sensor or stale feed	Null spike, constant values, or sudden single-feature jump	Check transport logs, broker queues, and sensor calibration	Quarantine feed, use fallback, repair hardware; do not retrain
Schema or unit mutation	Step change in feature mean or scale at release boundary	Compare ingestion payload against data contract version	Roll back producer change or update contract; do not retrain
Planned style or lot mix shift	Multivariate feature shift aligned with order schedule	Stratify error and predictions by style, yarn count, and lot ID	Keep model if slice metrics pass; suspend if out of scope
Seasonal ambient swing	Gradual cyclical drift in temperature or moisture features	Compare against historical seasonal slice and physical limits	Adjust operating settings or validate seasonal slice coverage
Delayed or censored labels	Apparent accuracy change on partial recent ground truth	Inspect label maturity cohort and operator override sampling	Wait for matured cohort or audit unlabelled slice; do not retrain
True process or concept drift	Sustained error rise on validated inputs after process change	Confirm clean feed and verify persistent error on matured labels	Open governed retraining candidate under full lifecycle control

DERIVED [S04] [S13] [S17] [S19] [S21] [S1201] Derived diagnostic reference matching production monitoring signatures to first verification checks, containment actions, and retraining rules.

The comparison in **Exhibit 52**, p. 199 shows why diagnosis must precede adaptation. [S21] Only the sixth row, where inputs pass all contract checks and matured ground truth confirms a persistent breakdown in the input-to-output mapping, justifies training a new candidate model. [S13] Even then, the table cannot tell a factory whether the newly trained candidate will outperform a simple rule comparator; that proof requires the full evaluation and staged-release discipline defined in the preceding chapters. [S43]

12.7

Delayed labels demand honest interim evidence

Textbook monitoring tutorials often assume that ground-truth labels arrive moments after a model makes a prediction, allowing the system to plot a live accuracy or error curve. [S21] In industrial operations, confirmed outcomes almost always lag behind the prediction event, sometimes by hours, days, or months. [S1202] When a small team ignores label delay and computes performance metrics over whatever partial outcomes happen to have arrived in the current week, it introduces severe survivorship and selection bias into its monitoring reports. [S19]

12.8 LABEL LATENCY VARIES ACROSS QUALITY, MAINTENANCE AND PLANNING HORIZONS

Every factory model faces a physical or commercial clock that governs when its predictions can be verified against reality: [S38]

- **In-line fabric and garment vision inspection:** Literature benchmarks report that automated computer vision can exceed 90 percent accuracy across more than 40 defect types at line-scan speeds up to 60 metres a minute, whereas manual inspection is 60 to 70 percent accurate and misses 20 to 30 percent of defects (inherited sector literature in R29, not a GFT operating result). [R29] On a live inspection frame, however, the camera scores every metre immediately, while human verification occurs only on flagged rolls, end-of-shift quality-assurance pulls, or downstream cutting-table rejections that may be recorded hours or days later. [S08]
- **Equipment-health and maintenance advisory models:** A model that warns of bearing wear or loom stoppage risk is tested only when a technician inspects the assembly during a scheduled maintenance window, or when an unplanned failure occurs days later. [S05]
- **Order, inventory and demand forecasting:** In apparel supply chains, McKinsey's sector study reports that AI at scale has been associated with forecast errors cut by 20 to 50 percent and inventory costs reduced by 35 to 40 percent (sector-wide executive and case benchmarks in R27, not SME or Cameroon plant guarantees). [R27] Yet a four-week or twelve-week demand forecast cannot be scored against realised orders until the entire forecast horizon has elapsed and customer cancellations or returns have settled. [S43]

12.9 PROXY INDICATORS BRIDGE THE GAP ONLY WHEN THEIR LIMITS ARE DECLARED

Because operations managers cannot wait three months to discover that a model failed on day one, the monitoring design must pair delayed ground truth with immediate proxy indicators. [S17] Useful proxies include the distribution of predicted probabilities or quantities, the share of inputs falling in low-confidence regions near the decision boundary, the rate at which line operators or planners override the model's recommendation, and short-horizon physical checks such as immediate rework counts at the adjacent workstation. [S1202]

Each proxy carries a structural limitation that must be stated in the monitoring specification. [S11] A stable prediction distribution proves only that the model is emitting similar numbers, not that those numbers remain right when the process has shifted. [S1201] Conversely, a low operator override rate may reflect genuine agreement, or it may reflect automation complacency on a busy shift where workers click through recommendations without independent inspection. [S30] Proxies can therefore trigger an early diagnostic review or a protective fallback, but they must never be treated as confirmed ground-truth labels for automated retraining. [S19]

METHOD

Dual-window reconciliation for delayed factory labels

To prevent partial or unverified outcomes from corrupting performance tracking, GFT's monitoring method separates live operational health from lagged accuracy accounting across two explicit time windows: [S17] [S21]

1. **Declare the label maturity horizon (L):** For each model in the registry, record the maximum elapsed time L (in hours, shifts, or weeks) required for at least a stated threshold share (for example, an illustrative 95 percent completeness gate) of ground-truth outcomes to be recorded and validated under the quality rules of ISO/IEC 5259-3. [S04]
2. **Operate the interim proxy window ($t - L$ to t):** For predictions generated during the most recent interval where ground truth is still incomplete, publish only service health, contract validity, feature drift, prediction distribution, and override metrics. [S17] Mark accuracy cells for this window as *pending label maturity* rather than plotting partial error rates that over-represent fast-arriving failures or returns. [S19]
3. **Score the matured cohort window ($t - L - W$ to $t - L$):** Evaluate realised technical accuracy and operational KPI impact over a scoring window of length W only after the maturity horizon L has passed, joining each historical prediction to its confirmed outcome via its immutable event and candidate identifiers. [S09] [S15]
4. **Reconcile proxies against matured truth:** Each month, compare the proxy alerts raised during window ($t - L - W$ to $t - L$) against the matured error rates now observed for that same cohort. [S36] If a proxy fired repeatedly without subsequent error degradation, widen or retire the proxy threshold; if accuracy degraded without a proxy warning, add a slice check or tighten input validation. [S11]

12.10

Alerts need thresholds, owners and runbooks

A monitoring dashboard that nobody reads until a quarterly review is theatre, while an alerting engine that floods supervisors with dozens of unprioritised statistical warnings every shift breeds alert fatigue and eventual muting. [S1202] Production monitoring controls risk only when human judgement sets the thresholds, each alert maps to a single accountable role, and that person follows a documented, periodically exercised runbook. [S11] [S27]

12.11 HUMAN-SET THRESHOLDS TIE STATISTICAL ALARMS TO OPERATIONAL CONSEQUENCE

Under the NIST AI Risk Management Framework, human judgement should be employed when deciding on specific trustworthiness metrics and the precise threshold values for those metrics. [S11] Production monitoring platforms allow teams to configure thresholds on data and model metrics that fire automated notifications

through event buses or operational queues. [S17] In a small factory, those thresholds should be tiered into two actionable severities rather than a vague spectrum of colours: [S14]

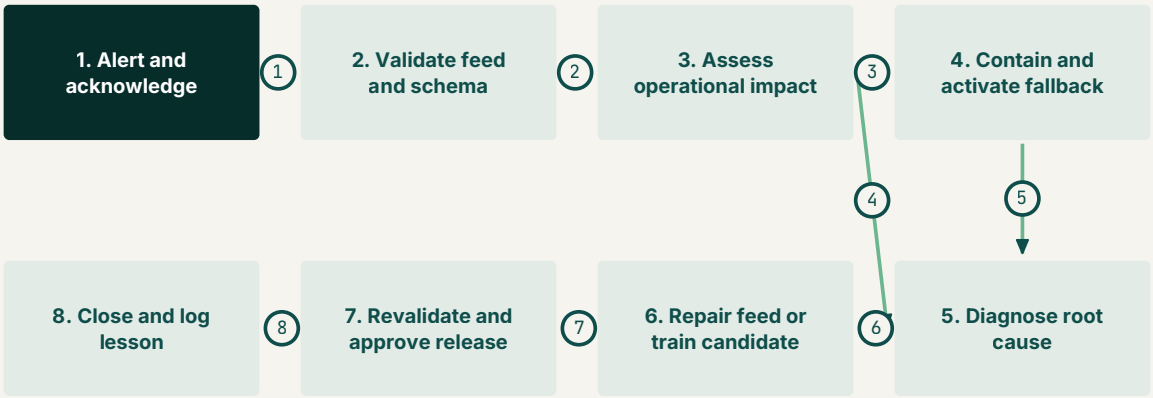
- **Warning threshold (investigate within the management cadence):** A metric has departed from its validation reference band (for example, a moderate shift in a secondary feature, a rising share of low-confidence predictions on one style, or a temporary buffer delay that recovered within its recovery-point objective), but contract rules hold and operational KPIs remain within control limits. [S17] The alert opens an assigned diagnostic ticket for the data steward or model builder to inspect before the next daily or weekly operations review. [S02]
- **Critical threshold (contain immediately on the current shift):** A precondition for safe or reliable model operation has failed, such as a broken schema contract, a null spike in a primary feature, an edge inference latency breach that stalls a line, or a verified jump in false rejections on a quality audit sample. [S22] The alert requires immediate acknowledgement and containment by the on-shift owner, including the authority to suspend the model and revert to the manual or statistical baseline. [S15]

12.12 EVERY ALERT EXECUTES A SIX-STAGE DIAGNOSED RUNBOOK

In cybersecurity and operational risk management, NIST SP 800-61 Revision 3 emphasises that response procedures can be tested or exercised periodically to verify their accuracy and train personnel. [S27] GFT applies the exact same discipline to data and model anomalies: an alert is not closed when someone clicks an acknowledgement button, but only when the runbook reaches verified operational closure. [S12] As depicted in **Exhibit 53**, p. 203, the response path forces feed validation and impact containment ahead of any modelling work. [S17] [S21]

EXHIBIT 53 / DERIVED OPERATIONAL SEQUENCE

An alert becomes controlled action through a diagnosed runbook



1. Named owner logs ticket
2. Check contract and nulls
3. Critical risk or bad feed
4. Bounded warning only
5. Baseline fallback active
6. Classify failure layer
7. Test fix on local slice
8. Owner signs off recovery

DERIVED [S15] [S17] [S21] [S27] Derived eight-node incident and drift response workflow. Containment and feed verification precede root-cause diagnosis and any candidate retraining decision.

Walking through **Exhibit 53**, p. 203 clarifies why a small team can run this control without a twenty-person operations centre. [S16] Step 1 routes the threshold breach to a single named role with an explicit acknowledgement window. [S17] Step 2 checks whether the ingestion feed, clock synchronisation, and schema contract are intact; if a feed has failed, Step 3 and Step 4 immediately quarantine the faulty records and switch the work cell to its rehearsed fallback procedure so production is protected while the custodian repairs the pipeline. [S15] Only in Step 5 does the team diagnose whether the anomaly stems from hardware, data contracts, product mix, or true concept drift. [S21] Step 6 executes the targeted remedy, Step 7 verifies that either the repaired feed or the newly trained candidate passes acceptance gates, and Step 8 records the root cause and verification date in the incident log. [S27] A flow diagram cannot guarantee that staff will follow the sequence under shift pressure; only periodic fault-injection drills make the fallback path habitual. [S27]

12.13

Retraining follows a diagnosed need and a new approval

In large consumer-internet firms, continuous retraining pipelines sometimes refit recommendation models every night on billions of clicks. [S18] Copying that pattern into a small factory is an engineering mistake. [S16] Manufacturing datasets are smaller, physical regimes are structured by equipment physics and batch schedules, and an unannounced change in model behaviour confuses operators who rely on consistent advisory thresholds across shifts. [S42] Under GFT's lightweight MLOps architecture, model retraining is a deliberate production change triggered by diagnosis, not by a calendar cron job or an raw drift score. [S13]

12.14 FOUR PRECONDITIONS GOVERN THE DECISION TO RETRAIN

Before a model builder spends compute and engineering hours fitting a new model version, the diagnostic record from Step 5 of the runbook must confirm four preconditions: [S21]

1. **Upstream data integrity is verified:** The data steward and custodian confirm that sensors, manual entry forms, unit conversions, and feature transformations obey their current contracts under ISO/IEC 5259-3, and that the observed degradation is not an artefact of missing or corrupted records. [S04]
2. **Matured ground truth confirms operational decay:** On a representative sample of verified outcomes (or an emergency physical audit pulled to bypass normal label latency), the active model's performance or operational benefit has fallen below the human-set acceptance threshold defined at release, or falls behind the simple comparator. [S11] [S43]
3. **Training coverage for the new regime exists:** The factory has accumulated enough clean, accurately labelled records from the new process condition, fabric style, or market regime to form a valid training snapshot and a separate holdout evaluation slice without data leakage. [S15] If a new style was introduced yesterday, there is not yet enough local history to retrain a history-hungry model; the right action is temporary fallback or bounded human rule adjustment until an adequate sample is collected. [S44]
4. **The new candidate passes full lifecycle validation:** Once trained, the replacement model receives a fresh candidate identifier in the registry, records its dataset snapshot and pipeline execution lineage, passes slice-level comparisons against both the previous model and the baseline comparator, and obtains sign-off from the operating owner before staged release. [S13] [S15]

WORKED EXAMPLE: DIAGNOSED SENSOR REPAIR VERSUS BLIND RETRAINING ON A TEXTILE INSPECTION LINE

Consider an illustrative weaving and inspection unit running 20,000 metres of cotton fabric a week across 50 rolls of 400 metres each, supported by an operator-assisted computer-vision model that flags defect locations alongside a moisture and tension feature feed. All figures in this case are illustrative assumptions constructed to demonstrate the arithmetic of monitoring choices, not observations from a GFT facility. [S17] [S38]

Assumed baseline operating parameters:

- True defect rate in the woven fabric: an illustrative 5.0 defects per 100 metres, or 1,000 true defects a week across 20,000 metres.
- Validated model performance under clean inputs: an illustrative 90 percent true-positive detection rate (catching 900 defects a week, consistent with the upper range in inherited literature R29) and an illustrative 2.0 false alarms per 100 metres (400 false alarms a week). [R29]
- Operator verification time: an illustrative 1.5 minutes to inspect and disposition each flagged alarm, and an illustrative penalty cost of 4,500 XAF in downstream cutting rework for each missed defect that escapes to the garment floor.
- Normal weekly verification labour: `(900 + 400) * 1.5 = 1,950` minutes (32.5 illustrative operator hours a week), with `100 * 4,500 = 450,000` XAF a week in illustrative escaped-defect rework.

The production anomaly (Week 1 of illustrative fault): A condensation fault causes the inline moisture sensor to drift upward by an illustrative 18 percentage points. Because the model uses moisture to adjust optical contrast thresholds for wet versus dry sizing, its prediction distribution shifts sharply. True-positive detection drops from 90 percent to an illustrative 70 percent (catching 700 defects and letting 300 escape), while false alarms jump from 2.0 to an illustrative 6.0 per 100 metres (1,200 false alarms a week).

Path A: Blind automatic retraining (undisciplined response):

- At the end of Week 1, an automated script detects feature drift and retrains the model on Week 1's corrupted moisture readings paired with incomplete operator overrides (where fatigued operators began ignoring alarms during the spike). [S19]
- During Week 1 and Week 2, before a technician notices the physical condensation fault during routine maintenance at the end of Week 2, the line absorbs two weeks of degraded performance: `300 - 100 = 200` extra escaped defects a week, costing an illustrative `200 * 4,500 = 900,000` XAF a week in extra rework, plus `(1,900 - 1,300) * 1.5 = 900` minutes (15 illustrative hours) a week of wasted verification time.
- Worse, when the technician cleans and recalibrates the moisture sensor at the start of Week 3, the automatically retrained model is now calibrated to the corrupted +18-point offset. It misfires again in Week 3 (causing another illustrative 900,000 XAF in extra rework) until engineers manually roll back the registry and rebuild the pipeline, incurring an illustrative 250,000 XAF in engineering remediation cost. [S18]
- **Total three-week illustrative excess loss under Path A:** `(3 weeks * 900,000 XAF) + 250,000 XAF = 2,950,000 XAF`, plus 45 illustrative hours of wasted floor inspection time.

Path B: Diagnosed runbook response (GFT method):

- Within 2 hours of the moisture-feature range breach on Day 1 of Week 1, the Layer 3 and Layer 4 monitor fires a critical alert to the data custodian and quality steward. [S17]
- Following **Exhibit 53**, p. 203, the steward quarantines the moisture feed and switches the vision unit to its pre-validated optical-only fallback profile (which achieves an illustrative 85 percent detection rate, or 850 caught and 150 escaped defects a week, with 500 false alarms). [S15]
- The maintenance technician replaces the moisture sensor desiccant on Day 2 (40 percent through the five-day week). The line runs on the fallback profile for 2 days (0.4 of a week), incurring only `0.4 * (150 - 100) = 20` extra escaped defects (`20 * 4,500 = 90,000 XAF`) and an illustrative 35,000 XAF sensor-service cost, then resumes normal validated operation for the remaining 2.6 weeks with zero model retraining required.
- **Total three-week illustrative excess loss under Path B:** `90,000 XAF + 35,000 XAF = 125,000 XAF`.
- **Net illustrative saving from diagnosis over blind retraining:** `2,950,000 - 125,000 = 2,825,000 XAF` across a single three-week incident.

12.15

Feedback loops belong in the monitoring design

A model deployed into production does not merely observe the factory; by guiding human decisions or machine settings, it changes the future data the factory records. [S18] Sculley et al. identify hidden feedback loops as one of the most expensive sources of ongoing maintenance debt in real-world machine-learning systems, because a system that shapes its own training data can gradually blind itself to the regions it stops exploring. [S18]

12.16 SELECTIVE LABELLING AND CENSORING DISTORT FUTURE TRAINING SETS

Consider what happens when an operator-assisted quality or scheduling model runs for six months. [S1202] If quality inspectors examine only the fabric rolls or garment bundles that the model flags as high risk, while passing unflagged bundles directly to packing, the factory's database records confirmed ground-truth labels only for the subset of production where the model already suspected a problem. [S19] Any false negatives that slip past the model and do not trigger an immediate customer return are silently recorded as defect-free or remain unlabelled. [S04] If a data analyst later pulls those six months of operational logs to retrain the model, the training set is censored by the previous model's own threshold: the new candidate learns that unflagged regions have zero defects, reinforcing the blind spot of its predecessor. [S18]

The same distortion appears in demand and inventory planning. [S43] When a planner follows a forecast to restrict raw-material allocation or finished-goods stock for a garment style, sales of that style cannot exceed the stocked quantity. [S38] Unless the spine records stockout events and unfulfilled order enquiries as distinct visibility

events under the EPCIS contract rather than recording only invoiced dispatches, a subsequent forecasting model mistakes supply-constrained shipments for true customer demand. [S08]

12.17 DOWNSTREAM KPIS AND WORKER SAFEGUARDS CATCH PERVERSE LOCAL OPTIMISATION

A model can also satisfy its narrow technical metric while damaging an adjacent stage of production or placing unfair pressure on staff. [S14] If a cutting-room or sewing-line advisory tool is tuned to minimise material offcuts or maximise hourly bundle output without monitoring downstream rework, sewing machine stoppages, and operator fatigue, line supervisors may push through marginal work that inflates rework hours in the finishing department. [S38] Moreover, where any monitoring metric or model output touches individual or small-team performance, the International Labour Office code of practice states that decisions concerning a worker should not be based solely on the automated processing of that worker's personal data (non-binding international guidance adopted by GFT as an operating design rule). [S30] Monitoring must therefore watch both cross-stage production KPIs and human-factors indicators such as override reasons, dispute logs, and operator workload. [S11]

> **GFT NOTE**
>
> **Keep a small unblinded audit sample and a cross-stage guard-rail KPI on every active model.** To break self-reinforcing feedback loops without building a complex experimental platform, GFT recommends two practical shop-floor rules for its Cameroon textile programme: [S04] [S38] 1. **Reserve an illustrative 2 to 5 percent stratified random audit slice:** On any inspection, sorting, or triage process where a model filters what humans review, route a small, randomly selected share of unflagged items through full manual verification each week. [S04] Tag these records in the lineage store as `audit_sample=true` so false negatives remain measurable and future retraining snapshots have an unbiased reference slice. [S09] 2. **Pair every primary model metric with one downstream guard-rail KPI:** In the model registry, link each model not only to its target indicator (for example, cutting yield or inspection throughput) but also to one governed ISO 22400-2 indicator in the next downstream process (such as sewing first-pass yield or customer claim rate), so local optimisation cannot hide a transferred defect. [S38]

12.18

One incident taxonomy joins data, model, security and OT response

When an anomaly appears on a factory screen at 06:30 on a Monday shift, the line supervisor does not know in advance whether the root cause is a loose Ethernet plug, an expired broker certificate, a corrupted vendor update, a ransomware infection, or a drifted statistical model. [S27] Maintaining separate incident procedures for IT security, OT engineering, data stewardship, and machine learning guarantees confusion during the first hour of a disruption. [S23] A small factory needs one shared incident taxonomy that classifies the failing layer, protects physical safety first, and allows each layer to be contained independently. [S22]

12.19 CONTAINMENT ACTIONS DECOUPLE THE FEED, THE MODEL AND THE MACHINE

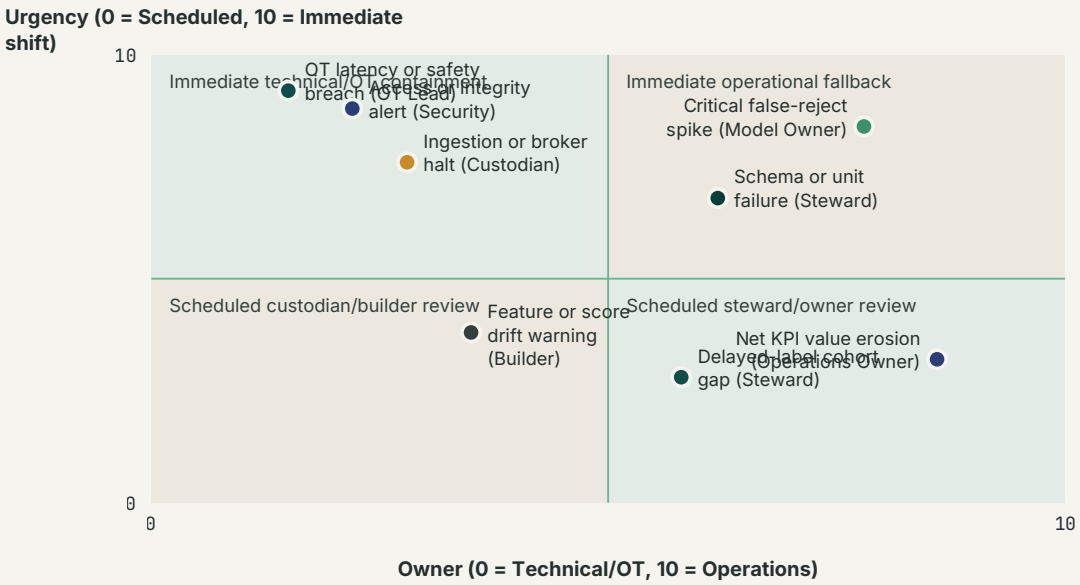
Under NIST SP 800-82 Revision 3, operational technology security and resilience must address the unique performance, reliability, and safety requirements of physical equipment. [S22] Because GFT's data spine separates capture, contextual storage, and model consumption, the incident runbook can isolate a fault at the exact boundary where it occurs without shutting down healthy production systems: [S15]

- **Layer A: Quarantine the data feed (Steward / Custodian action):** When a single sensor, manual entry terminal, or external API sends malformed, delayed, or suspiciously altered payloads, the ingestion gate routes that source to quarantine while allowing unaffected lines and local operator screens to continue operating. [S04]
- **Layer B: Suspend the model to baseline fallback (Model Owner / Operations action):** When a model's predictions become erratic, uncalibrated, or out of scope, the registry pointer is switched to `suspended` or rolled back to the prior approved version, and the work cell reverts to its documented manual SOP or statistical rule comparator without interrupting data capture. [S13] [S15]
- **Layer C: Isolate the OT or network boundary (OT / Security action):** If an alert indicates unauthorised network traffic, credential compromise, or an edge process consuming excessive CPU or memory on a control-adjacent gateway, the OT and security leads sever the external or analytical connection under the practised incident procedures of NIST SP 800-61 Revision 3, leaving machines under deterministic local control and verified offline backups intact. [S26] [S27]

As mapped in **Exhibit 54**, p. 209, assigning each monitoring signal to its primary owning layer and escalation authority prevents either technical paralysis or unauthorised shop-floor improvisation. [S02] [S23]

EXHIBIT 54 / DERIVED GOVERNANCE POSITIONING (0 TO 10 SCALE)

Monitoring ownership follows the failing layer



DERIVED [S02] [S17] [S22] [S27] Derived responsibility map placing eight recurring monitoring signals across technical-to-operational ownership (x-axis) and urgency of containment (y-axis), based on role principles in S02, monitoring alerts in S17, OT safety in S22, and incident response in S27. Coordinates are analytical placements to clarify escalation routing, not empirical measurements.

In **Exhibit 54**, p. 209, the upper-left quadrant groups signals where the OT lead, security approver, or data custodian must act immediately to protect equipment availability, network integrity, or message buffering. [S22] [S27] The upper-right quadrant covers acute operational failures, such as a critical spike in false quality rejections or an out-of-bounds scheduling recommendation, where the model owner or shift supervisor immediately triggers fallback to baseline practice. [S11] The lower quadrants hold slower-moving diagnostic signals, including moderate feature drift, delayed-label reconciliation, and monthly KPI value erosion, which are investigated during the normal weekly and monthly management cadence rather than interrupting a running shift. [S36] Because this matrix is an analytical governance map, each factory must bind the role labels to specific named individuals and deputies in its local charter before deployment. [S02]

> **WARNING**
>
> **Never allow a monitoring agent or drift response to write directly to safety-critical PLCs or override physical interlocks.** Under NIST SP 800-82 Revision 3, operational technology environments require deterministic performance and physical safety guarantees that statistical models cannot provide. [S22] A monitoring pipeline or automated fallback script must never alter machine safety interlocks, emergency stops, motor protection limits, or chemical dosing valves directly. Even in bounded automation modes, fallback means reverting to a deterministic, engineer-approved PLC setpoint or human standard operating procedure, verified during periodic incident exercises under NIST SP 800-61 Revision 3. [S27]

12.20

Review and retirement prevent model accumulation

The final failure mode of production machine learning in resource-constrained organisations is zombie model accumulation: models launched during an initial pilot remain running on servers years later after the product line has changed, the original analyst has left, and operators have quietly stopped looking at the output. [S18] Every unretired model continues to consume ingestion pipelines, storage, security patching, and cognitive attention, multiplying technical debt across the spine. [S16]

12.21 SCHEDULED AND EVENT-DRIVEN REVIEWS TEST CONTINUED FITNESS

An enduring AI management system under ISO/IEC 42001:2023 and lifecycle process control under ISO/IEC 5338:2023 require continual evaluation and explicit retirement processes alongside development and operation. [S12] [S13] To keep the model portfolio clean, each deployed artefact in the factory registry must carry two review triggers: [S1203]

1. **A mandatory calendar review date (expiry lease):** At release approval, every model is granted a fixed operating lease (for example, an illustrative 90-day or 180-day review period depending on risk tier). [S14] If the model owner and domain steward do not review its monitoring record, delayed-label accuracy, operator override logs, and net KPI value before the lease expires, the model is scheduled for suspension or formal renewal at the monthly operations review. [S36]
2. **An event-driven review trigger:** Regardless of the calendar, a mandatory review opens whenever the factory executes a material upstream change, including a new equipment installation, a revision to a governed KPI formula under ISO 22400-2, a schema or taxonomy breaking change, a vendor software upgrade on a bought model, or a regulatory update affecting worker or personal data. [S38] [S30]

As illustrated in **Exhibit 55**, p. 211, tracking explicit registry states from initial candidate admission through suspension, rollback, revalidation, and formal retirement makes a model's authority transparent at every point in its service life. [S13] [S16]

EXHIBIT 55 / ILLUSTRATIVE OPERATING MONTH FROM REGISTRATION

Model status remains visible from candidate to retirement

0
Candidate v1.0 registered
Snapshot, code hash, and baseline comparator frozen in registry.

1
Validated in shadow mode
Local slice metrics and fallback drill passed on live shift inputs.

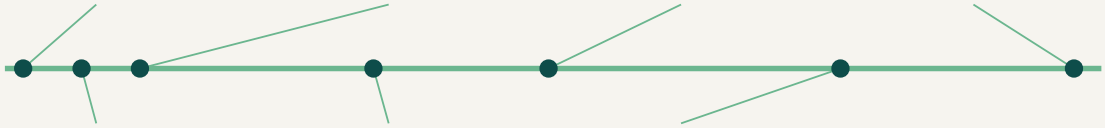
2
Approved v1.0 deployed
Operating owner signs off operator-assist release with 90-day lease.

6
Suspended and rolled back
Alert runbook detects dye-process change; baseline fallback restored.

9
Retrained v1.1 approved
New candidate trained on matured post-change cohort passes gates.

14
Lease renewed after audit
Unblinded audit slice and net KPI benefit verified at review.

18
Model v1.1 formally retired
In-line hardware sensor replaces estimation; dependencies removed.



ILLUSTRATIVE [S12] [S13] [S15] [S16] Illustrative 18-month registry history for a single factory model, showing staged release, an incident-triggered suspension and rollback at Month 6, a diagnosed retraining release at Month 9, and formal retirement at Month 18 when a process redesign removes the use case.

The sequence in **Exhibit 55**, p. 211 shows that neither suspension at Month 6 nor retirement at Month 18 is a governance failure. [S44] On the contrary, suspending `v1.0` when a physical process change invalidates its training distribution protects the shop floor from acting on uncalibrated scores, while retiring `v1.1` at Month 18 when a direct physical meter eliminates the need for statistical estimation frees the team from maintaining an obsolete pipeline. [S18] Because this timeline is an illustrative lifecycle trace, actual review intervals must be set by each site's operating cadence and risk classification. [S11]

12.22 CONTROLLED RETIREMENT REMOVES RUNTIME DEBT WHILE KEEPING AUDIT LINEAGE

Retiring a model is an engineering procedure, not an informal decision to stop checking a screen. [S13] When the operations owner decides to retire a model, whether because a simpler rule works as well, a hardware upgrade solves the problem directly, or expected benefits were not met under the stop criteria of S44, the custodian and builder execute a four-step decommissioning checklist: [S15] [S44]

- **Deactivate serving endpoints and scheduled jobs:** Shut down the model container, batch scoring cron, and feature-extraction queries that existed solely for that model so they stop consuming compute, network bandwidth, and broker topics. [S18]

- **Remove user-interface indicators and alert routes:** Update operator screens, planning spreadsheets, and alert queues so staff cannot mistake a frozen historical score for an active recommendation, and notify affected shifts of the return to the standard operating procedure. [S12]
- **Revoke service credentials under least privilege:** Disable the machine identities, API tokens, and database permissions assigned to the retired deployment pipeline. [S23]
- **Preserve the immutable registry and lineage record:** Set the registry state to `retired` with the effective timestamp, approver, and reason, while retaining the model card (S1203), training dataset manifest, pipeline execution logs (S15), and provenance graph (S09) for the applicable retention period so any past production lot, customer quality claim, or sustainability report influenced by the model remains fully reconstructable. [S09] [S1203]

By binding threshold-based monitoring to diagnosed runbooks, human release authority, and clean retirement, a small factory team keeps every model subordinate to physical operations. That discipline sets up Part IV of this paper, where Chapter 13 examines how adoption and net operational value are measured against the baseline, and Chapter 14 sequences the entire industrial data spine into a gated 18-month roadmap. [S36]

An alert is a question put to an engineer, never a blank cheque for a pipeline to rewrite itself.

CHAPTER SOURCE NOTES

S17 Microsoft Learn, Azure Machine Learning. 2026-01-27.

S19 Association for Computing Machinery, CHI 2021. 2021.

S13 ISO/IEC. 2023.

S21 arXiv, author manuscript of an IEEE Transactions on Knowledge and Data Engineering article. 2020-04-13.

R06 Cameroon Trade Hub. Publication date unrecorded; inherited source-text check.

R07 Ecofin Agency. Publication date unrecorded; inherited source-text check.

S38 International Organization for Standardization. 2014.

S04 ISO/IEC. 2024.

S1201 arXiv, Advances in Neural Information Processing Systems (NeurIPS) 2019. 2018-10-29.

S15 Google Cloud. 2024-08-28.

S33 Microsoft. 2024-06-10.
S08 GS1 AISBL. 2022.
S1202 arXiv, Proceedings of the VLDB Endowment. 2022-09-16.
S18 Advances in Neural Information Processing Systems. 2015.
S02 International Organization for Standardization. 2022.
S1203 arXiv, ACM Conference on Fairness, Accountability, and Transparency. 2018-10-05.
R29 MDPI; industry inspection literature. Publication date unrecorded; inherited source-text check.
S11 National Institute of Standards and Technology. 2023-01-26.
S43 Frontiers in Artificial Intelligence. 2021-06-14.
S05 International Electrotechnical Commission. 2013.
R27 McKinsey, The State of Fashion 2025. Publication date unrecorded; inherited source-text check.
S30 International Labour Office. 1997.
S09 World Wide Web Consortium. 2013-04-30.
S36 World Bank. 2023.
S27 National Institute of Standards and Technology. 2025.
S14 ISO/IEC. 2023.
S22 National Institute of Standards and Technology. 2023.
S12 ISO/IEC. 2023.
S16 Microsoft Learn, Azure Architecture Center. 2026-06-03.
S42 SCITEPRESS, IN4PL 2022. 2022.
S44 The Journal of Technology Transfer, Springer Nature. 2024-07-27.
S23 National Institute of Standards and Technology. 2024-02-26.
S26 CISA, MS-ISAC, NSA and FBI. 2023.

GREEN FUTURE TECH

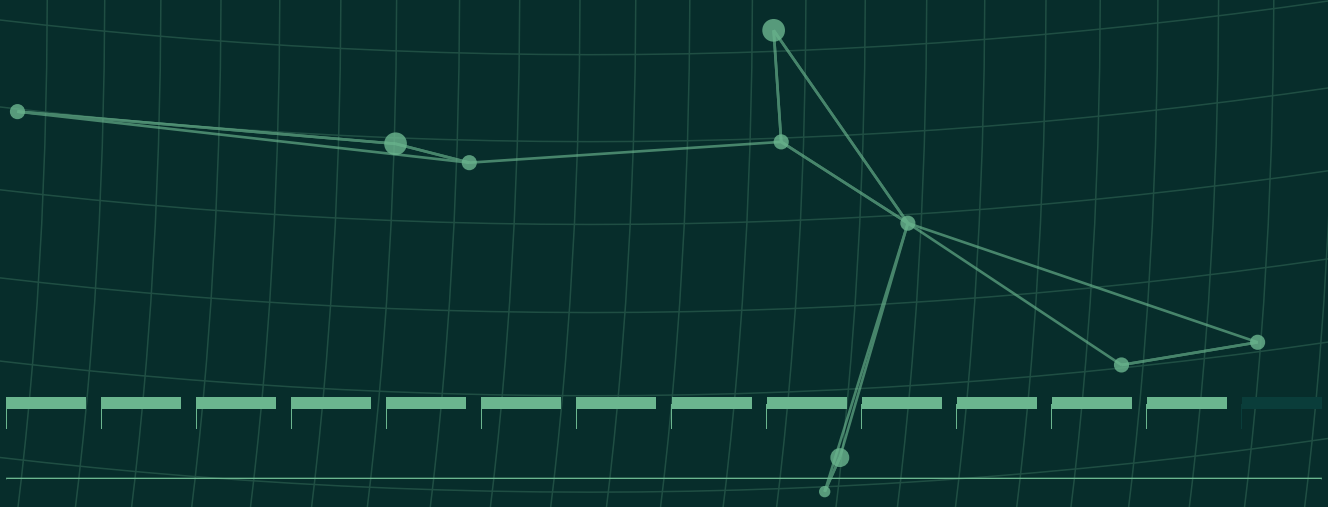
LIMITED / HONG KONG SAR

13

PART IV: PROVE VALUE AND SCALE DELIBERATELY / VALUE REALISATION

Adoption and value are operating metrics, not launch claims

The spine earns further investment when people use it, decisions improve and net operational benefit survives a comparison with the baseline.



CHAPTER 13
DECISION TEST

The spine earns further investment when people use it, decisions improve and net operational benefit survives a comparison with the baseline.

A data spine creates no operating value merely by going live. Value begins when a supervisor uses a trusted record to change a decision, when the resulting improvement survives comparison with the pre-intervention baseline, and when that benefit exceeds the recurring work required to keep the record or model useful. The launch is therefore the start of measurement, not its conclusion. GFT's recommendation is a recurring value routine with four linked tests: people perform the governed work, operational indicators improve under like-for-like conditions, net benefit remains after lifecycle effort, and trust or safety does not deteriorate. [S11] [S18] [S36] [S38]

Sustained use cannot be inferred from initial enthusiasm. In a World Bank study of an enterprise platform offered to SMEs in Viet Nam, close to 80 percent of supported firms used it after initial training, about 40 percent used it thereafter, and 35 percent were still using it about 18 months later. Firms that sustained adoption commonly assigned an internal point person to champion it. [S37] The study concerned business software in a different country and context. It is not a Cameroon factory forecast. It does show why activation and training completion are weak evidence of durable adoption.

This chapter turns that lesson into factory practice. It defines adoption as observable behaviour, measures outcomes through existing manufacturing indicators, separates cash effects from capacity and risk effects, and gives every deployment a fix, scale, maintain, postpone or stop decision. The method is an engineering governance framework. It is not investment, legal or tax advice. Site values, accounting treatment and worker-data obligations require the responsible factory functions and, where necessary, qualified local advisers.

13.1

A champion and local stewards sustain use

Adoption needs an operating owner because friction crosses functional boundaries. A capture form may belong technically to IT, its defect codes to quality, its completion time to production and its use in a shift meeting to operations. Without one person accountable for follow-through, each function can defend its own part while the complete routine fails. The World Bank's SME finding supports a named internal champion, but it does not justify dependence on a heroic individual. The champion needs management backing, a deputy, allocated time and a route into existing production decisions. [S37]

GFT would assign the champion three duties. First, make the governed workflow usable during ordinary shifts, including handover, interruption and exception handling. Second, coordinate role-based practice and confirm that users can perform their tasks without assistance. Third, bring behavioural evidence and unresolved friction to the factory's existing management cadence. Domain stewards remain close to the work. They diagnose capture defects, help operators correct records and explain why a definition or interface must change. Technical custodians keep the service available, but they do not own adoption.

13.2 WORKAROUNDS REVEAL UNMET OPERATING NEEDS

A private spreadsheet or paper tally after launch is evidence before it is a compliance issue. It may preserve a field missing from the governed record, provide a faster view at the line, allow work during disconnection or protect a local calculation that management still uses. Removing the sheet by instruction can hide the record without removing the need. The steward should ask which decision the workaround supports, what the official workflow cannot represent and which change would make the governed route more useful.

The review should also distinguish necessary local work from harmful duplication. A temporary offline contingency may be legitimate if it has an owner, a synchronisation rule and a reconciliation check. An unowned monthly extract that changes the KPI denominator is different. The former is a controlled operating mode. The latter fractures the baseline and prevents comparison. The champion records the cause, corrective action and verification result rather than reporting only the number of workarounds found.

> **GFT NOTE**
>
> Treat each workaround as a usability finding with an owner and a test. GFT's practical recommendation is to observe the task with the user, preserve any operational context the private record contains, then change the governed contract, interface or offline procedure. Retire the parallel record only after the user completes the same decision in the spine and reconciliation passes. A reduction achieved by prohibition is not evidence of adoption.

13.3

Training happens around tangible work

Training should reproduce the decisions a role will face, not tour software menus. OECD case evidence from a German bakery SME says that converting traditional processes into digital ones requires sizeable effort, does not directly pay off and involves a longer learning process. [S35] ILO and UNESCO guidance separately favours tangible pilots with practical use and built-in evidence about what works. Its two-to-three-year horizon concerns digital transformation of technical and vocational skills systems, not a universal factory payback period. [S40] Together, the sources support practice around real work and patience with learning, not a generic promise of rapid return.

A line steward should practise timely capture, an unknown defect code, a missing lot reference and correction after quarantine. A quality supervisor should practise reviewing a model flag, recording a reasoned override and checking the later physical outcome. A planner should practise comparing a forecast with the current statistical rule and documenting why a schedule changed. A custodian should practise restoring the records and instructions needed to reproduce a reported figure. Each exercise uses realistic records with personal data removed or protected according to purpose.

13.4 COMPETENCY MEANS UNAIDED TASK COMPLETION

Attendance does not demonstrate readiness. The practical test is whether the user can complete the governed task, recognise an exception and select the correct escalation route without prompting. A suitable assessment records the task, conditions, errors, assistance required and reassessment result. Failure leads first to inspection of the work instruction and interface. A confusing taxonomy or slow screen is a design defect, not a training deficiency to be repeated unchanged.

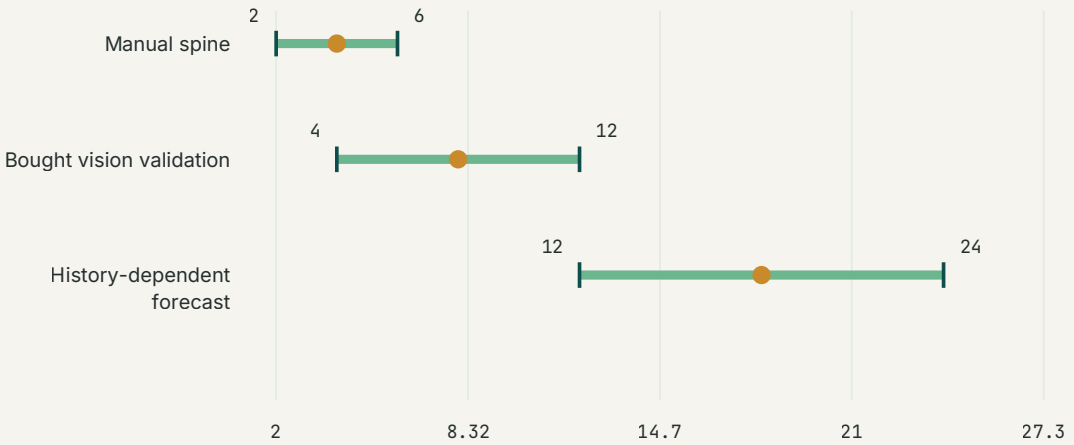
The champion should test competence again after a material process, schema or model change. This prevents an initial training certificate from becoming permanent evidence of capability. It also connects training to release control: a role that cannot execute fallback safely keeps the deployment in its current limited mode. NIST's risk framework places human judgement around metrics and thresholds, so users must also understand what a score means, when it is uncertain and when they are expected to challenge it. [S11]

13.5 EVIDENCE ACCUMULATION BELONGS IN THE PLAN

Time is required for habits to stabilise and for a fair comparison to accumulate. The ranges in **Exhibit 56**, p. 217 are illustrative planning assumptions for three different evidence problems. They are not benchmarks or commitments. A manual spine needs stable capture and repeated management use. A bought vision system needs local validation and override evidence even if no local training is planned. A history-dependent forecaster needs sufficient relevant order history and a comparator before its value can be judged. [S35] [S36] [S40] [S43]

EXHIBIT 56 / ILLUSTRATIVE MONTHS FROM APPROVAL

Time to value includes learning and evidence accumulation



ILLUSTRATIVE [S35] [S36] [S40] [S43] Illustrative planning assumptions, not observed durations or benchmarks. Low is initial routine or shadow use, mid is stable task competence with a usable comparison, and high is a net-value review with adequate evidence. S40's two-to-three-year statement concerns TVET strategy and does not set these ranges.

The ranges prevent two opposite mistakes. The first is declaring value before users and baselines have stabilised. The second is continuing a weak deployment indefinitely because learning is still supposedly under way. Each review should state which evidence is missing, when it should become available and what decision follows if it does not. Calendar time never substitutes for an evidence gate. [S36] [S44]

13.6

Adoption measures record behaviour, not screen activity

Licences issued, dashboard views and stored rows describe system activity. They do not establish that a record arrived in time for a decision, passed its quality checks or changed an operating action. The World Bank recommends embedding monitoring and evaluation throughout digitalisation programme design and implementation, which requires measures attached to the intervention's mechanism rather than a launch inventory. [S36] For the industrial data spine, GFT's proposed adoption set therefore uses five behavioural measures: timely capture, validated completeness, correction recurrence, active routine integration and documented override reasons.

Timely capture compares event time with the contract deadline for the decision. A shift total entered before handover may be useful even if a machine event arrived later through a verified offline queue. The measure must preserve that distinction rather than treating every replay as late work. *Validated completeness* counts required records that pass the applicable rules at first submission, with missing and estimated values disclosed rather than silently converted to zero. Data quality remains purpose-specific and managed across the lifecycle, consistent with ISO/IEC 5259-3. [S04]

Correction recurrence tests whether the same capture cause returns after correction. A falling number of corrected rows can signal improvement, but it can also result from a disabled rule or unreported workaround. *Active routine integration* checks whether shift, production and quality reviews cite governed records and close actions against them. *Documented override reasons* show whether a person recorded why a recommendation was rejected and whether the later outcome was reviewed. These are proposed factory measures, not universal thresholds. Each site sets its target after establishing a baseline. [S36]

> **METHOD**
>
> Write every adoption measure with a numerator, denominator, observation window, exclusions and owner. For timely capture, the denominator is the set of records required by the active production schedule, not the records that happened to arrive. For routine integration, sample the underlying action and evidence rather than accepting a meeting checklist. For correction recurrence, retain the rule version so a definition change cannot create artificial improvement. Reconcile each result with a small physical sample on the line.

13.7 ADOPTION METRICS NEED DENOMINATORS THAT RESIST GAMING

Every behavioural measure can be distorted if its denominator is convenient. Timely capture improves artificially when expected records are removed from the schedule. Completeness improves when difficult fields become optional. Routine integration improves when actions are closed without verification. The data owner should therefore approve the denominator and its exceptions before the result is used in a review. Changes remain visible through the contract version and effective date.

The review should pair each adoption indicator with a counter-check. Timely capture is read beside late corrections and physical sampling. Validated completeness is read beside quarantine volume and disabled rules. Override documentation is read beside the quality of the eventual decision. A high score that fails its counter-check opens an investigation rather than earning praise. This keeps the measure attached to factory work and discourages teams from optimising the report instead of the routine.

Adoption should also be compared across roles and shifts. An aggregate result can conceal a team that carries the capture burden for everyone else or a night shift that receives no steward support. Where the sample is too small for a conclusion, the register should say so and collect more evidence. The purpose is not to rank workers. It is to find where the operating design needs correction.

13.8 OVERRIDE FREQUENCY NEEDS A QUALITY CHECK

A low override rate can mean that a model is useful, but it can also mean that users accept its output without examining it. A high rate can mean poor model fitness, a confusing interface or a legitimate response to conditions absent from training. NIST warns that people may assume AI systems work well in every setting and calls for measurement to reflect real operational context. [S11] The adoption review should therefore sample agreements as well as overrides and compare both with delayed ground truth where available.

Three outcomes matter: a justified override catches a model error, an unjustified override rejects a sound recommendation, and an unchallenged model error passes without human intervention. The third is especially important because a simple override rate hides it. Sampling should include different lines, product families, shifts and data-quality conditions. Results lead to interface correction, threshold review, further training or suspension, not an instruction to agree with the model more often. [S11]

13.9

Operational outcomes use the factory's existing measures

A validation score answers whether a candidate meets a technical test. It does not answer whether the factory produces better work. Operational evaluation should reuse the KPI dictionary and its frozen formulas so that the baseline, intervention period and subsequent review retain the same meaning. ISO 22400-2 presents manufacturing KPIs through formulas, corresponding elements, time behaviour, units and other characteristics, making those definitions the appropriate reference point for comparison. [S38]

The selected indicators should follow the deployment's mechanism. Quality assistance may affect first-pass yield, rework and scrap. Shared stoppage evidence may affect availability and changeover duration. Maintenance advice may affect unplanned stops and repair time. Order and forecasting tools may affect schedule adherence and work in progress. Metering and allocation may affect energy, water or material intensity. A benefit claim should name the physical pathway from record to action to result; an improvement without that pathway is a hypothesis to investigate, not a validated effect. [S36] [S38]

13.10 THE COUNTERFACTUAL MUST SURVIVE CHANGES IN OPERATING CONDITIONS

The central question is what would have happened without the intervention under comparable conditions. Product mix, material grade, order urgency, equipment state and staffing can change at the same time as a release. A before-and-after average that ignores those changes can credit the tool for an easier mix or blame it for a harder one. The evaluation should hold the baseline formula constant and compare matched lines, shifts, styles or operating strata where the available evidence permits. [S36] [S43]

A simple rule helps: report the total observed change separately from the portion attributed to the intervention. The first is an operational observation. The second requires a comparison method and confidence statement. If a like-for-like comparison is impossible, the register should retain the result as unverified contribution rather than forcing a causal claim. This is particularly important when several changes share one period, such as a new maintenance routine and a model release.

13.11 A BASELINE IS A VERSIONED OPERATING ARTEFACT

A usable baseline contains more than a historical average. It records the KPI version, source records, quality state, observation window, operating conditions and current comparator. It also records known weaknesses, such as incomplete stoppage capture or an inconsistent product boundary. If those weaknesses would change the value decision, they become corrective work before a strong benefit claim is allowed. [S04] [S36] [S38]

The baseline owner should be independent enough from the deployment team to challenge a favourable reconstruction. Where new information improves an old record, the team should preserve both the original reported baseline and the corrected version, with lineage explaining the difference. This allows reviewers to distinguish a genuine operating change from an accounting repair. It also prevents a weak starting point from being selectively adjusted until the intervention appears valuable.

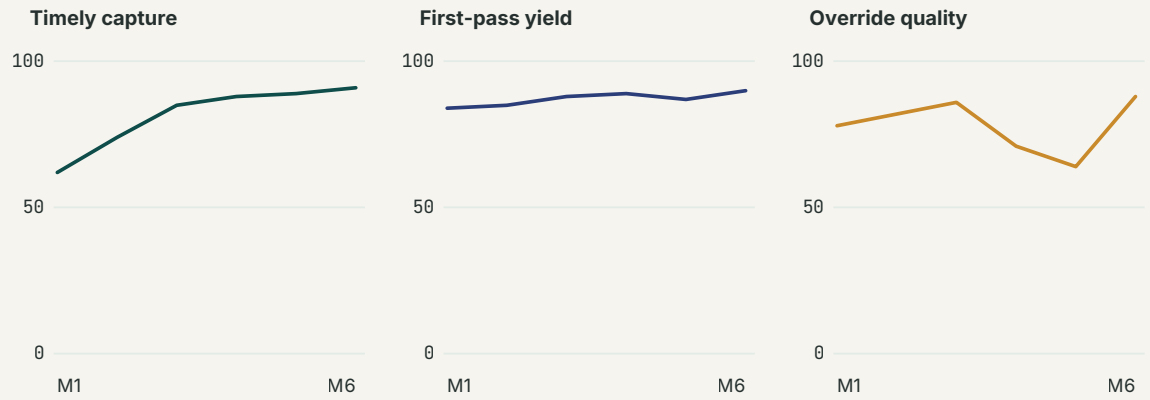
Uncertainty belongs in the outcome as well as the baseline. A narrow difference between two highly variable periods may not support a scale decision even when its arithmetic is correct. The register should state the sample support, important uncontrolled conditions and the additional observation needed. This is a decision-quality requirement, not a demand for statistical complexity that a small team cannot maintain.

13.12 ADOPTION, OPERATIONS AND TRUST BELONG ON ONE REVIEW PAGE

The illustrative panels in **Exhibit 57**, p. 221 show why these measures should not be read alone. Capture and yield can improve while override quality deteriorates. A programme that reports only the first two panels would miss the emerging oversight problem. The values are invented to demonstrate the review method and are not GFT results or factory benchmarks. [S11] [S30] [S36] [S38]

EXHIBIT 57 / ILLUSTRATIVE PERCENT ON A SHARED ZERO TO 100 SCALE

Adoption, operations and trust can move in different directions



ILLUSTRATIVE [S11] [S30] [S36] [S38] Illustrative six-month sequence. Values are invented to show diagnostic divergence. Timely capture and first-pass yield improve while audited override quality falls in M4 and M5, then recovers after intervention. These are not measured GFT or client outcomes.

In this illustrative sequence, timely capture rises from 62 to 91 percent and first-pass yield ends at 90 percent, while override quality falls from 86 to 64 percent before recovering. The correct response is not to celebrate the average direction. It is to inspect the model, alert burden and user decisions during the weaker months, then verify whether the corrective action restored informed review. The exhibit cannot establish why the measures moved; it tells the team where to look.

13.13

Financial value is net of lifecycle effort

A real operational improvement can still produce a weak investment once the work around the model is counted. Research on machine-learning systems identifies ongoing maintenance costs arising from dependencies, configuration, hidden feedback and changes in the external world. [S18] A 2025 NIST semiconductor workshop report adds a sharply bounded warning: manufacturers reported that data processing for a single application can cost as much as 70 percent of total application cost. That is a reported upper bound in a particular manufacturing context, not an average for SMEs or a prediction for a textile factory. [S39]

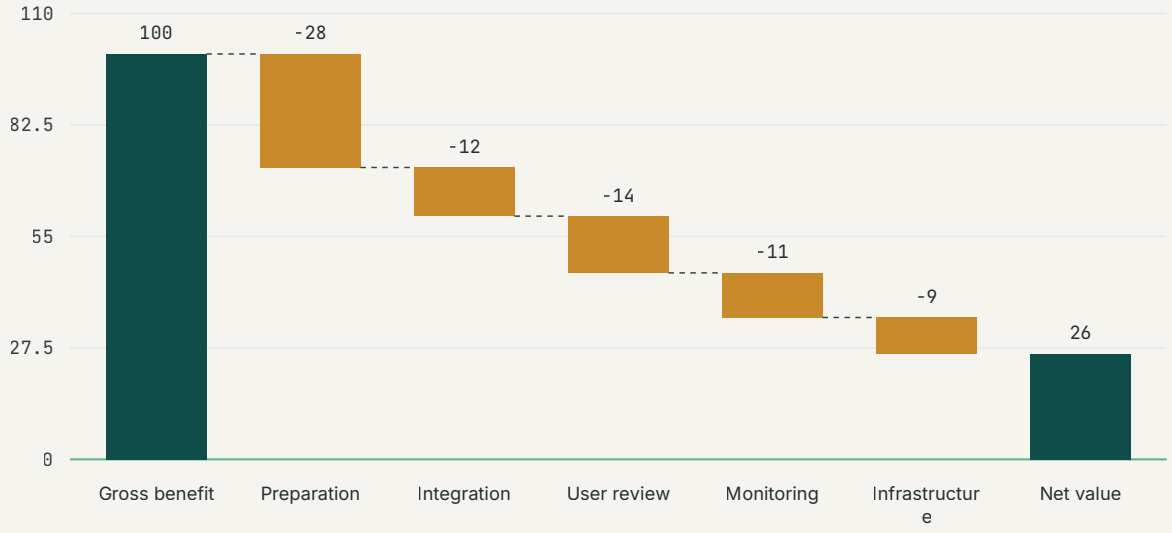
The value review should therefore begin with validated gross operating benefit and deduct the recurring effort needed to sustain it. Preparation includes entry, cleaning, labelling and reconciliation. Integration includes schema, master-data and interface

maintenance. Human review includes alert handling, false interventions and override documentation. Model operation includes monitoring, delayed-label review and revalidation. Continuity and support include recovery tests, access review, incremental infrastructure and licences. Initial implementation cost should also remain visible, but it must not be used to hide the recurring run cost. [S18] [S36] [S39]

Exhibit 58, p. 222 expresses the reconciliation as an illustrative index. It is deliberately not a return ratio. The purpose is to make every deduction visible before the remaining value is discussed. [S18] [S36] [S39]

EXHIBIT 58 / ILLUSTRATIVE ANNUAL VALUE INDEX, GROSS BENEFIT EQUALS 100

Net value remains after recurring lifecycle work is counted



ILLUSTRATIVE [S18] [S36] [S39] Illustrative reconciliation: 100 - 28 - 12 - 14 - 11 - 9 = 26. Cost categories reflect lifecycle maintenance in S18 and the application-specific processing warning in S39. Proportions are invented for method demonstration, not benchmarks or GFT results.

In the illustration, recurring lifecycle work consumes 74 of the 100 gross-benefit units and leaves 26. The appropriate management question is not whether the gross case sounded attractive. It is which deductions can be reduced through reuse and better capture without weakening control, and whether the remaining net value justifies the deployment's risk and scope. A negative result should trigger redesign, postponement or retirement rather than a search for additional unverified benefits. [S44]

13.14 CASH, CAPACITY AND AVOIDED RISK REMAIN SEPARATE

Direct cash effects include measured reductions in material waste, external rework, energy purchases or paid overtime, supported by quantities and prices for the stated period. Freed capacity includes machine or labour hours released by the change. It becomes additional revenue only when saleable demand and the remaining production constraints allow those hours to be used. Until then, it should be reported in hours or physical units, not converted into speculative cash.

Avoided risk is different again. A lower probability of a customer claim, safety event or equipment failure may matter greatly, but its monetary estimate depends on likelihood, consequence and assumptions. Keep that estimate separate from ledger savings and

disclose its basis. Do not add cash, unused capacity and avoided risk into one unqualified total. The register should let management see each category and decide which is relevant to the resource decision. [S11] [S36]

WORKED EXAMPLE

An illustrative weaving mill evaluates an advisory inspection system after a matched-style review. All quantities, prices and hours below are illustrative assumptions chosen to demonstrate the arithmetic, not observations from GFT or a client. The mill processes an illustrative 1,200,000 metres a year. Downgraded fabric falls from an illustrative 1.8 percent to 1.1 percent on comparable styles, giving an attributed improvement of 0.7 percentage points. At an illustrative XAF 1,500 a metre, validated gross annual benefit is `1,200,000 × 0.007 × XAF 1,500 = XAF 12,600,000`. [S36] [S38]

Illustrative recurring costs are XAF 1,050,000 for preparation and stewardship, XAF 480,000 for operator review, XAF 864,000 for integration and revalidation, and XAF 4,200,000 for annualised hardware and support. Their total is `XAF 1,050,000 + 480,000 + 864,000 + 4,200,000 = XAF 6,594,000`. Net annual cash value is therefore `XAF 12,600,000 - 6,594,000 = XAF 6,006,000`. Lifecycle cost equals about 52.3 percent of gross benefit by GFT's arithmetic from the illustrative inputs. [S18] [S39]

If preparation and integration require an additional illustrative XAF 7,000,000, net value becomes `XAF 6,006,000 - 7,000,000 = -XAF 994,000`. That sensitivity changes the decision from maintain to redesign or stop. Any freed finishing capacity is recorded separately in hours until confirmed demand can use it. The example excludes tax, financing and accounting treatment, which require the factory's responsible advisers. [S36] [S44]

13.15

Trust and harm sit beside benefit

A deployment can improve an operating indicator and still be unacceptable. False alerts can create unnecessary stops or inspections. A capture routine can distract an operator or encourage an unsafe shortcut. Worker-level records can become a disciplinary instrument even when their original purpose was process diagnosis. NIST treats positive and negative impacts as part of AI risk management, so harm and trust evidence belongs in the same review as benefit rather than in a separate annual exercise. [S11]

The trust review should cover false-intervention burden, safety near misses, unresolved override friction, worker concerns and contestation outcomes. Each measure needs a defined event and investigation route. A concern count alone is ambiguous: fewer complaints can mean fewer problems or less confidence that reporting will help. The more useful evidence is whether concerns were understood, whether the response changed the system and whether the same cause returned. [S11] [S30]

The ILO code of practice states that decisions concerning a worker should not be based solely on automated processing of that worker's personal data. It also says that personal data collected by electronic monitoring should not be the only factors in evaluating performance. The code is non-binding guidance, not a statement of

Cameroonian law. [S30] GFT's operating recommendation is to keep worker-affecting decisions under informed human review, provide an accessible challenge route and correct disputed records without retaliation.

Meaningful review requires time, authority and context. If a station record shows low output, the reviewer should be able to inspect material, machine, style and measurement conditions before attributing fault. If a worker challenges the record, the steward should preserve the original evidence and append the resolution. A system that cannot explain or correct its worker-affecting use has not passed the value gate, regardless of its apparent financial return. [S11] [S30]

13.16

A benefit register drives fix, scale or stop decisions

Value evidence becomes useful only when it changes a resource decision. GFT's proposed benefit register gives every active deployment one row and every review one explicit outcome. The adoption champion maintains it with the operating owner, using the factory's existing management forum rather than creating a parallel programme office. Each row links back to the admission case, the frozen baseline and the released artefact so the team evaluates the same mechanism and scope that was approved. [S36] [S43]

The register records the accountable owner, baseline definition and period, expected mechanism, behavioural adoption, observed operating effect, comparison method, gross benefit, recurring lifecycle cost, evidence confidence, adverse effects and next decision. It also records missing evidence. This prevents an uncertain estimate from being silently promoted to a measured saving between reviews. Where the KPI definition or production boundary changes, the old and new results remain distinguishable. [S36] [S38] [S41]

Four actions are available. *Scale* expands scope only after adoption, net value, recovery and trust checks pass. *Maintain* keeps a beneficial deployment within its current scope where expansion adds no proven value. *Fix* contains the deployment and assigns dated corrective work when friction, cost or override quality threatens the case. *Postpone or stop* returns to the verified manual or simpler comparator when the necessary resources, skills, costs or expected benefits do not support continuation. Research with SMEs explicitly recognises postponement as a valid investment decision under those conditions. [S44]

Exhibit 59, p. 225 shows the structure with illustrative textile deployments. The schema is derived from the cited evaluation, KPI and lifecycle sources. The row values are illustrative and are included only to show how different evidence leads to different actions. [S18] [S30] [S36] [S38] [S44]

EXHIBIT 59 / DERIVED REGISTER WITH ILLUSTRATIVE ROW ENTRIES

A benefit register turns observed evidence into a portfolio decision

Deployment and baseline	Adoption and harm	Value and confidence	Decision	Next evidence
Line log: stops 14.2% of shift	92% timely; no open concern	Stops down 2.4 pts; high	Scale	Reconcile on next line
Vision: downgrade 1.8% of metres	88% override quality	Net XAF 6.0M/yr; medium	Maintain	Repeat matched-style audit
Energy log: steam 9.4 kg/kg	71% complete; manual meters	Break-even; low	Fix	Test meter and recost
Bearing alert: 6 stops/quarter	48% override quality	Negative; medium	Stop	Compare fixed maintenance
Forecast: MAPE 28% at 8 weeks	Shadow use; 7 months history	Unproven; low	Postpone	Review after more history

DERIVED [S18] [S30] [S36] [S38] [S44] Derived register structure from embedded evaluation, governed KPI definitions, lifecycle cost and stop decisions. All deployment values are illustrative, not observed GFT results. Net value means validated operating benefit less recurring lifecycle effort. Confidence reflects comparison and input quality, not model accuracy alone.

Stopping does not erase the learning. The registry should retain why the deployment failed, which records remain reusable and which assumption would need to change before reconsideration. A stopped model may leave a valuable governed dataset and a clearer baseline. Continuing to pay for an unsupported model, by contrast, consumes the same small team's capacity for capture, maintenance and process improvement. [S18] [S44]

13.17

Quality disclosure improves with the evidence

Value measurement should start before every input is perfect, but it should show where the evidence is weak. GHG Protocol guidance describes collecting data, assessing quality and improving quality as an iterative process, with lower-quality inputs replaced over time. Although that guidance addresses Scope 3 emissions, the iterative principle is useful for factory value records when its original context remains explicit. [S41]

GFT's register should distinguish four evidence states. *Measured* contributions come from direct observations under the relevant KPI definition and comparison. *Partially measured or allocated* contributions combine observations with documented allocation rules. *Estimated or proxy* contributions depend on defaults, samples or assumptions. *Model-derived* contributions carry their validation limits and uncertainty. These labels describe the evidence behind a value claim, not a general score for the deployment. [S04] [S36] [S41]

The improvement decision follows consequence. A low-confidence input that has little influence on a decision may be adequate for the next review. A proxy that determines whether a deployment scales deserves priority for a better meter, sample

or reconciliation routine. This focuses spending on the evidence gap that can change a decision rather than on indiscriminate data collection. Each iteration should record what changed and whether earlier results need recalculation. [S36] [S41]

The same discipline can connect the paper's companion tools without pretending they share one denominator. The `kpi-cascade` and `oee-line` reviews can preserve operational definitions. The `inspection-economics` and `cold-start-forecast` reviews can expose preparation, review and comparator costs. The `footprint` review can carry measured, allocated and estimated states into resource intensity. The benefit register is the common decision layer, while each specialised method remains distinct. [S36] [S38] [S39] [S41]

The chapter's test is therefore cumulative. Launch proves availability. Adoption proves governed behaviour. Operational comparison tests the claimed mechanism. Lifecycle accounting tests net value. Trust review tests acceptability. The benefit register turns those findings into a decision, and quality disclosure makes the strength of that decision visible. Expansion is justified only when the complete chain holds.

Value is not what launch reports; it is what ordinary work can verify.

CHAPTER SOURCE NOTES

S11 National Institute of Standards and Technology. 2023-01-26.

S18 Advances in Neural Information Processing Systems. 2015.

S36 World Bank. 2023.

S38 International Organization for Standardization. 2014.

S37 World Bank. 2025-05-20.

S35 OECD. 2022-11-21.

S40 International Labour Organization and UNESCO. 2020.

S43 Frontiers in Artificial Intelligence. 2021-06-14.

S44 The Journal of Technology Transfer, Springer Nature. 2024-07-27.

S04 ISO/IEC. 2024.

S30 International Labour Office. 1997.

S39 National Institute of Standards and Technology. 2025-11-18.

S41 GHG Protocol, World Resources Institute and World Business Council for Sustainable Development. 2013.

14

PART IV: PROVE VALUE AND SCALE DELIBERATELY / RECOMMENDATIONS

A gated 18-month roadmap builds capability without a data department

GFT should scale the spine only as ownership, meaning, recovery, reproducibility and net value are demonstrated in production.

CHAPTER 14
DECISION TEST

GFT should scale the spine only as ownership, meaning, recovery, reproducibility and net value are demonstrated in production.

A small or mid-sized factory does not need a standalone data department to govern its operational records or run machine-learning models safely in production. Instead, it needs a staged, evidence-gated sequence that embeds data ownership, explicit event semantics, quality quarantine, tested offline recovery, reproducible model release and net value measurement directly into existing factory roles and management routines. [S02] [S16] [S36] In GFT's Cameroon textile programme, prepared on 31 January 2024 with source-refresh annotations retained in 2026, the sector context is defined by an acute local-processing gap: the inherited evidence register records that about 95 percent of Cameroon's cotton fibre is exported raw while about 5 percent is processed locally, against a national industrialisation target of 50 percent local processing by 2030 under the National Development Strategy. [R06] [R07] Those two inherited entries were not refreshed in 2026 and their original publication dates were not recorded in the register, so they frame the strategic imperative rather than certifying a current factory baseline. [R06] [R07] Closing that industrial gap requires factories whose daily production, quality, maintenance, order and resource records can be trusted from the first shift.

This closing chapter translates the preceding governance, architecture and MLOps controls into an illustrative 18-month implementation roadmap structured around six sequential stages and nine hard decision gates. Every duration in this chapter is an illustrative planning assumption designed to show sequencing and dependency, not an observed benchmark or fixed calendar promise. [S16] [S36] Advancement from one stage to the next is governed by inspectable operational evidence rather than elapsed calendar time. When a gate fails because capture effort is excessive, quality rules are bypassed, backups cannot be restored, or a model candidate fails to beat a simple comparator net of lifecycle cost, the required governance action is to hold, redesign, postpone or stop the expansion. [S26] [S44]

14.1

The first 30 days establish authority and scope

The opening illustrative window from day 1 to day 30 establishes human accountability, bounds the initial scope to a handful of decision-critical records, registers lawful processing purposes and freezes the baseline definitions against which every subsequent improvement claim will be judged. [S02] [S36] [S38] Small industrial teams frequently stall when they begin digitalisation by evaluating software platforms before deciding who owns a definition and which operational decision the record serves. [S02] [S37] Survey evidence from the Organisation for Economic Co-operation and Development in the Western Balkans indicates that integrated operational tools such as enterprise resource planning and cloud systems are used by only 10 to 20 percent of surveyed small and medium enterprises, showing why a roadmap cannot assume integrated systems already exist (while noting that this survey covers the Western Balkans rather than Cameroon). [S34] Under ISO 8000-150, establishing appropriate roles and responsibilities is a foundational requirement for data quality management, and World Bank case evidence from small-firm adoption confirms that firms successfully adopting digital platforms assign an internal point person who champions adoption across functions. [S02] [S37] During this illustrative first month, the factory

director appoints one internal adoption champion with a named deputy and assigns one accountable domain owner for each of six core domains: production, quality, maintenance, orders, people and sustainability. [S02] [S37]

14.2 AUTHORITY PRECEDES SOFTWARE SELECTION

Each domain owner signs a compact role charter that separates three distinct responsibilities: business accountability for meaning and access (the data owner), recurring shop-floor verification and exception handling (the data steward), and technical administration of spreadsheets, databases, brokers or backups (the data custodian). [S02] [S23] In a constrained team, one person may hold more than one role across domains, provided conflicts of interest are declared and no technical custodian can silently alter a business formula or approve a high-consequence model release alone. [S02] [S11] [S14] Line stewards receive an explicit, illustrative weekly time allowance within their normal shift duties so that reviewing missing logs or quarantining malformed entries is scheduled operational work rather than unpaid overtime. [S02] [S35]

Rather than cataloguing every machine tag or administrative spreadsheet in the enterprise, the champion and domain owners select an illustrative initial scope of five critical operational records governing daily shift performance: hourly line output, defect inspection events, unplanned machine stoppages, order and lot dispatch records, and shift attendance aggregates. [S05] [S08] [S38] For each selected record, the domain owner freezes a versioned entry in the KPI dictionary specifying purpose, formula, input elements, unit, time behaviour, grain and exclusions under ISO 22400-2. [S01] [S38] Freezing these definitions before pipeline or model work begins prevents later disputes over whether rework counts as good output or whether planned breaks reduce availability. [S38]

14.3 LEGAL SCREENING AND BASELINE MEASUREMENT START ON DAY ONE

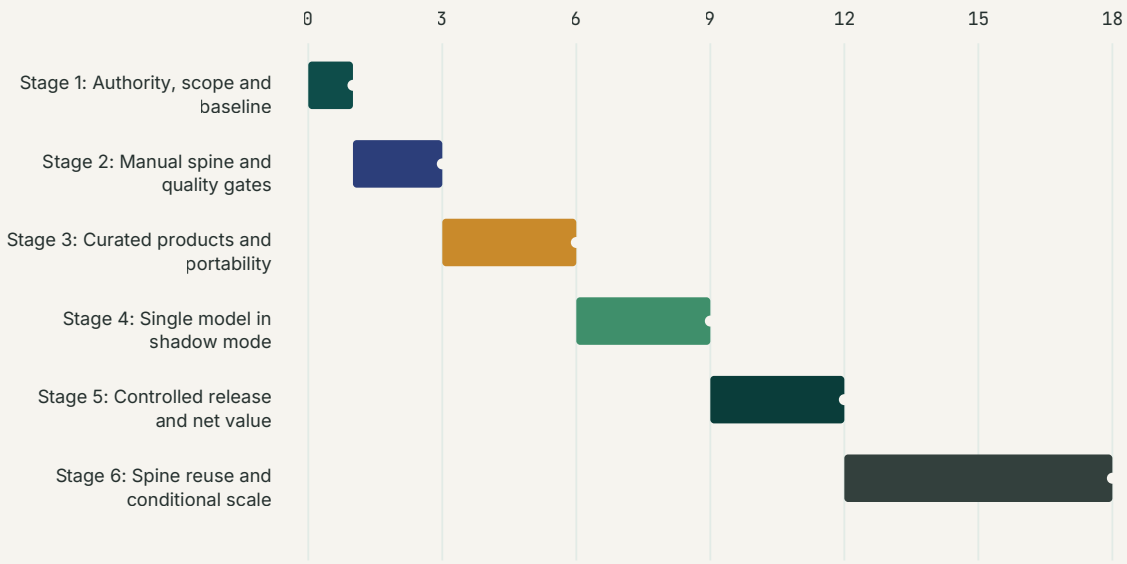
Because operational records touch workers and external partners, the first 30 days also establish the privacy and access boundary. [S29] [S31] [S1401] Under Regulation (EU) 2016/679, personal data must be adequate, relevant and limited to what is necessary for the stated purpose, while World Wide Web Consortium best practices emphasise that security, commercial sensitivity and privacy govern what data and metadata may be shared. [S29] [S1401] In Cameroon, Law No. 2024/017 of 23 December 2024 states in Section 19(1) that the processing of personal data shall be subject to prior authorisation by the Personal Data Protection Authority. [S31] Before collecting or repurposing any worker-linked telemetry, attendance record or video feed in this illustrative 30-day window, the factory records each processing purpose in a local register, strips unnecessary personal identifiers from operational aggregates, and obtains current local legal advice on statutory authorisation duties. [S29] [S30] [S31]

Finally, the team records a pre-intervention operating baseline across process performance and data handling, applying the World Bank principle that embedding monitoring and evaluation from the outset is essential for tracking impact and learning. [S36] Over an illustrative two-week to four-week observation window, stewards measure shift minutes spent entering and reconciling logs, capture-defect recurrence, correction turnaround time, backup restoration capability, and baseline levels

for first-pass yield, line availability and order fulfilment. [S26] [S36] [S38] Without this frozen baseline, neither the manual data spine nor any subsequent model can prove falsifiable operational value. [S36] [S43]

EXHIBIT 60 / MONTHS FROM APPROVAL (ILLUSTRATIVE PLANNING WINDOW)

Eighteen months of capability building remain conditional on evidence



ILLUSTRATIVE [S02] [S04] [S15] [S16] [S26] [S36] [S44] Illustrative 0 to 18 month roadmap aligned with GFT's Cameroon textile programme phases. Entry and exit gates derive from S02, S04, S15, S16, S26, S36 and S44; task durations are planning assumptions, not measured factory benchmarks.

As shown in **Exhibit 60**, p. 230, the six illustrative stages run across an 18-month planning window from month 0 to month 18, with each stage bounded by explicit entry and exit conditions. [S16] [S36] The chart is not an automatic calendar escalator: if Stage 2 fails to prove offline buffering and backup restoration at month 3, Stage 3 and Stage 4 remain locked until the recovery defect is fixed and re-verified. [S26] [S44] What **Exhibit 60**, p. 230 cannot establish is the exact pace of a specific site, because actual durations depend on shift complexity, staffing stability, power reliability and legal review timelines. [S35] [S42]

14.4

Days 31 to 90 make the manual spine trustworthy

Once authority, scope and baseline definitions pass Gate 1 at day 30, the illustrative second stage from day 31 to day 90 turns the factory's manual and semi-automated logs into a governed system of record. [S04] [S08] [S09] GFT's Cameroon textile programme starts with structured manual and spreadsheet-backed capture on the shop floor rather than waiting for full machine instrumentation. That approach succeeds only when manual capture enforces the same identifier discipline, event semantics, quality checks and continuity rehearsals that later automated feeds will inherit. [S04] [S08]

14.5 IDENTIFIER AND EVENT CONTRACTS STABILISE SHOP-FLOOR CAPTURE

Between illustrative days 31 and 60, the domain owners and custodian issue versioned data contracts for the critical records selected in Stage 1. [S04] [S10] Every record is anchored to immutable canonical identifiers for site, area, line, asset, order, lot, style, operation, defect code and shift, excluding personal worker identifiers unless a documented, legally reviewed purpose requires them. [S05] [S08] [S29] Following GS1 EPCIS Release 2.0, operational occurrences such as defect observations, machine faults, lot splits or carton dispatches are recorded as discrete events carrying event time, ingestion time, object identity, location and business step rather than overwriting mutable status cells. [S08] Provenance links follow the World Wide Web Consortium PROV-O and PROV-DM specifications so every record retains the entity, activity and agent involved in producing or modifying it. [S09] [S1403]

In the same illustrative window, the factory activates capture and ingestion quality gates under ISO/IEC 5259-3 and ISO/IEC 25012 across five dimensions: validity, completeness, consistency, timeliness and uniqueness. [S03] [S04] Missing readings are tagged explicitly as `missing`, engineering approximations as `estimated`, and supervisor amendments as `manually_corrected` with a superseding provenance link, never silently replaced with a numeric zero. [S04] [S09] Failed records are routed into a quarantine table and assigned to the originating line steward for correction, preventing upstream data cascades from contaminating curated summaries. [S04] [S19] Each recurring defect, its decision impact, its root cause and its verification date are tracked in a published quality register informed by the W3C Data Quality Vocabulary and iterative quality guidance. [S41] [S1402]

14.6 CONTINUITY REHEARSALS PRECEDE ANY MODELLING WORK

Between illustrative days 61 and 90, before any machine-learning model is admitted, the team proves that the manual spine survives infrastructure disruption. [S26] [S27] [S33] Drawing on store-and-forward buffering patterns documented for offline edge devices, local capture sheets and edge queues buffer timestamped records durably during network outages and replay them idempotently without duplication when connectivity returns. [S07] [S08] [S33] Simultaneously, the custodian maintains offline encrypted backups of raw logs, master tables, KPI definitions and validation rules, and executes a live restoration drill into an isolated environment alongside an incident-response exercise as recommended by CISA and NIST SP 800-61 Revision 3. [S26] [S27]

EXHIBIT 61 / ILLUSTRATIVE 90-DAY ACTION REGISTER

The first 90 days produce decisions and evidence, not a transformation slogan

Window and owner	Governed deliverable	Verification evidence	Escalation forum
Days 1 to 10: Factory director	Six-role governance charter and champion mandate	Signed charter with deputy and steward time allowance	Weekly operations review
Days 11 to 20: Domain owners	Version 1.0 KPI dictionary for priority records	Frozen formulas, units, grains and exclusions	Weekly operations review
Days 15 to 25: Champion and counsel	Local personal-data processing register	Minimised fields and Cameroon Law 2024/017 screening	Executive risk sign-off
Days 20 to 30: Line stewards	Four-week operating and data-effort baseline	Logged capture minutes, error rates and baseline KPIs	Gate 1 review at day 30
Days 31 to 50: Custodian and owners	Canonical ID registry and event data contracts	Schema check on production, defect and stop logs	Daily shift handover
Days 51 to 75: Line stewards	Active ingestion quarantine and quality register	Quarantined records resolved with PROV-O lineage	Weekly quality cadence
Days 76 to 90: Custodian and OT lead	Disconnection buffer test and backup restoration drill	Zero duplicate replay and verified isolated restore	Gate 2 review at day 90

ILLUSTRATIVE [S02] [S04] [S26] [S36] [S38] Illustrative owner-level action list for days 0 to 30 and days 31 to 90, with deliverable, verification evidence and escalation route derived from S02, S04, S26, S36 and S38. Day windows are illustrative planning ranges.

The sequence in **Exhibit 61**, p. 232 makes every deliverable in the first 90 days inspectable by naming the owner, artefact, verification proof and escalation forum. [S02] [S36] [S38] Rather than creating a second governance committee, each action routes exceptions into the factory's daily shift handover and weekly operations review (wp02). [S02] [S23] Because **Exhibit 61**, p. 232 is an illustrative template, a single-line plant can compress administrative steps, whereas a multi-site spinning, weaving and garment operation must validate identifiers across site boundaries before closing Gate 2. [S05] [S08]

14.7

Months 4 to 6 create reusable data products

Passing Gate 2 at day 90 proves that critical shop-floor events are captured under stable identifiers, checked by active quarantine rules and recoverable from offline backups. [S04] [S26] The illustrative third stage from month 4 to month 6 converts those raw logs into reusable, governed data products so downstream consumers rely on one canonical core instead of fragile, application-specific extraction scripts. [S09] [S10] [S39] Workshop evidence from the National Institute of Standards and Technology notes that manufacturers have reported data-processing costs for a single application reaching as high as 70 percent of total cost. [S39] Treated carefully as an

upper-bound warning from semiconductor manufacturing rather than a universal SME average, that 70 percent figure underscores the cost penalty of building separate pipelines for every report or model. [S18] [S39]

14.8 SEPARATING RAW EVIDENCE FROM CURATED PRODUCTS

During illustrative months 4 and 5, the custodian and domain owners separate three storage layers: immutable raw evidence, contextualised canonical records and published consumption products. [S05] [S09] [S15] Raw shift logs, inspection sheets and machine messages are stored append-only with ingestion timestamps and digests so original evidence is never overwritten. [S09] [S1403] Above the raw layer, deterministic transformations join events to versioned master data (assets, styles, calendars and defect taxonomies) and apply quality rules to publish curated operational tables. [S04] [S06] [S08] When a steward corrects a quarantined entry, the correction appends a superseding event linked via `prov:wasRevisionOf` and recomputes the curated product with full lineage. [S09] [S15]

From these curated tables, the spine publishes controlled consumption views without duplicating business formulas: production and downtime views for OEE and line balancing (`wp02`, `wp15`), defect and genealogy views for quality and traceability (`wp03`, `wp05`), and utility and material views for resource and carbon accounting (`wp06`, `wp16`). [S05] [S08] [S38] [S41] Each published view carries Data Quality Vocabulary metadata disclosing completeness, freshness and estimation share so consumers never mistake an estimate for a metered reading. [S41] [S1402]

14.9 RECONCILIATION, LEAST-PRIVILEGE ACCESS AND VENDOR EXIT READINESS

Before Stage 3 closes at month 6, the factory completes three verification checks. [S25] [S32] [S38] First, over an illustrative four-week parallel run, the team reconciles KPI figures from the curated products against legacy spreadsheet calculations under the frozen ISO 22400-2 definitions. [S09] [S38] Second, the security approver and custodian review access under NIST SP 800-53 Revision 5 least-privilege controls and CISA cybersecurity goals, confirming that operators, supervisors, analysts, vendors and automated services use separate, narrowly scoped credentials. [S24] [S25] [S1401] Third, the team executes a portability export test: applying EU Data Act principles where applicable and explicit contract clauses elsewhere, the custodian exports six months of raw events, curated tables, schemas, KPI definitions and PROV-O lineage into open formats and verifies that an independent script reproduces the monthly management report without vendor software. [S06] [S08] [S09] [S32]

METHOD

Stage-gate verification protocol for the industrial data spine. Every stage boundary in this illustrative 18-month roadmap operates as a binary gate with four mandatory checks: 1. **System boundary and artefact:** The domain owner and custodian present the live production artefact (quarantine log, isolated restore proof or shadow evaluation ledger) across the defined lines and shifts. [S02] [S15] [S26] 2. **Independent verification check:** A role that did not build the technical artefact verifies the evidence against human-set acceptance thresholds. [S11] [S14] [S23] 3. **Failure action:** If any check fails, the gate remains closed, promotion is blocked, and a remediation task is logged in the weekly operations review. [S27] [S44] 4. **Preserved audit trail:** Gate evidence, schema versions and signed approvals are archived in the governance register with provenance links. [S09] [S12] [S1403]

14.10

Months 7 to 9 prove one model in shadow mode

Only after Gate 3 confirms trusted, portable data products at month 6 does the roadmap admit a first machine-learning candidate in the illustrative window from month 7 to month 9. [S04] [S16] [S43] Industrial MLOps research shows that starting infrastructure and target automation levels drive implementation complexity, making gradual capability growth essential for small teams. [S16] [S42] Following GFT's Cameroon textile programme sequence, the recommended first use case is a bounded shop-floor problem with rapid ground-truth feedback, such as computer-vision defect inspection (wp03), whereas history-dependent demand forecasting (wp04) waits until clean multi-season order history accumulates. [R29] [S04] [S43]

14.11 USE-CASE ADMISSION, BASELINE FREEZE AND SIMPLE COMPARATOR

Month 7 opens with the model admission screen grounded in the six-phase CRISP-DM lifecycle and explicit SME project-postponement criteria. [S43] [S44] The operating owner specifies the target decision, acting role, response window, operational technology safety boundary under NIST SP 800-82 Revision 3, and worker-protection rules under the International Labour Office code of practice. [S22] [S30] [S43] The team freezes the pre-model baseline and defines a simple, transparent comparator (such as a statistical rule or standard inspection checklist) that the candidate must beat net of lifecycle cost. [S36] [S38] [S43] Builder and custodian then record the candidate in a lightweight reproducibility stack comprising version control, immutable dataset snapshots with ISO/IEC 5259-3 quality manifests, an experiment execution log, and a model registry binding data, code, vendor version and evaluation metrics to one candidate identifier. [S04] [S13] [S15] [S20]

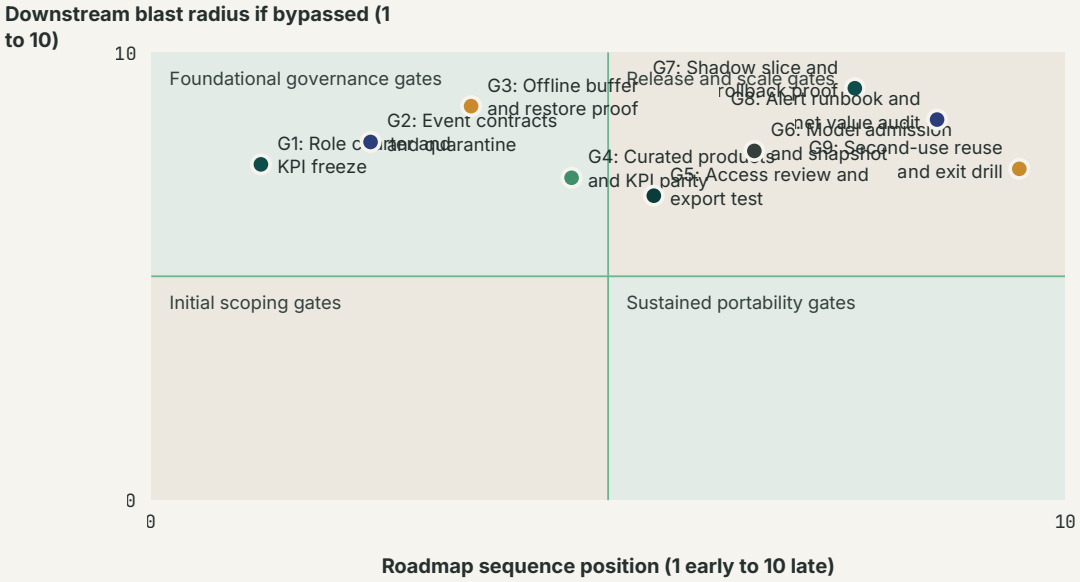
14.12 SHADOW EVALUATION AND ROLLBACK REHEARSAL

During illustrative months 8 and 9, the candidate runs exclusively in silent shadow mode on live inputs, logging timestamped predictions without displaying outputs to operators or altering shop-floor actions. [S15] [S16] [S22] Each week, the builder and steward evaluate shadow predictions against verified ground truth across local slices (fabric shade, style, shift, lighting and missing-data state) so aggregate scores cannot hide line-specific failures. [S04] [S11] [S19] Before Gate 7 closes at month 9, the team executes a live rollback and manual-fallback drill, proving that model suspension and reversion to manual operation complete cleanly within the target recovery

window. [S15] [S22] [S27] Operating and risk owners then judge shadow slice results, comparator margin, support cost and rollback evidence against human-set thresholds, postponing or stopping any candidate that fails. [S11] [S14] [S44]

EXHIBIT 62 / DERIVED ORDINAL SCALE FROM 1 (INITIAL) TO 10 (FACTORY-WIDE)

No stage opens until its predecessor leaves usable evidence



DERIVED [S02] [S04] [S09] [S11] [S15] [S17] [S26] [S32] [S36] [S44] Derived gate register placing nine programme gates across sequence position (x-axis, 1 to 10) and downstream decision and safety blast radius if bypassed (y-axis, 1 to 10), synthesised from S02, S04, S09, S11, S15, S17, S26, S32, S36 and S44.

The derived gate register in **Exhibit 62**, p. 235 maps the nine decision gates (G1 through G9) across sequence position (horizontal axis, 1 to 10) and downstream operational and safety blast radius if bypassed (vertical axis, 1 to 10). [S02] [S11] [S26] [S44] Early foundational gates (G1 to G3) score high on blast radius because unowned definitions or unrecoverable records corrupt every subsequent report and model, while G7 carries the highest score (9.2 on the derived 10-point scale) because it guards the transition from silent shadow testing to live shop-floor interaction. [S04] [S11] [S19] [S26] Coordinates in **Exhibit 62**, p. 235 are analytical ordering indices rather than empirical measurements; any control-connected use case requires even stricter scrutiny under ISA/IEC 62443 and NIST SP 800-82 Revision 3. [S22] [S28]

14.13

Months 10 to 12 prove controlled operation

Clearing Gate 7 at month 9 opens the illustrative fifth stage from month 10 to month 12, moving the candidate from shadow mode into controlled, operator-assisted production. [S12] [S13] [S16] Initial release remains advisory: the model flags a potential defect or maintenance anomaly, while a trained operator or supervisor retains explicit authority and time to inspect, accept or override the output. [S11] [S22] [S30] Where outputs relate to worker activity, the International Labour Office principle that employment decisions must not rest solely on automated personal-data processing is enforced as a hard rule with logged human review. [S30]

14.14 RELEASE PACKAGE, THRESHOLD MONITORS AND DIAGNOSED RUNBOOKS

Month 10 begins by issuing an ISO/IEC 42001 and ISO/IEC 5338 release package documenting approved model and dataset versions, intended use, slice limits, operator instructions, fallback procedures and review dates, paired with task-based workstation training anchored in tangible daily work. [S12] [S13] [S35] [S40] The custodian activates threshold monitors across service health, feature and prediction drift, and delayed label accuracy. [S04] [S17] [S21] Following Azure ML monitoring architecture and concept-drift research, crossing a statistical threshold triggers a rehearsed diagnostic runbook rather than automatic retraining: the assigned role verifies whether the shift reflects an ingestion fault, a product-mix change or genuine concept drift, and either repairs the feed, invokes manual fallback or opens a governed retraining candidate. [S13] [S17] [S21] [S27]

14.15 NET VALUE AUDIT AND PORTFOLIO DECISION AT MONTH 12

At month 12 (Gate 8), the champion, operating owner and controller audit three-month production results against the Stage 1 baseline across adoption behaviour (timely capture, completeness, correction recurrence, meeting use and override quality), operational KPIs and net financial return after subtracting ongoing preparation, review, integration and support costs. [S18] [S36] [S38] [S39]

WORKED EXAMPLE

Illustrative net-value gate calculation for a single-line defect-inspection assistant at month 12. Consider an illustrative garment plant evaluating an operator-assisted vision station at Gate 8 after three months of operation (months 10 to 12). All figures are explicitly illustrative assumptions demonstrating the accounting method, not observed GFT factory results or forecasts.

1. **Frozen baseline (illustrative inputs):** The line produces an illustrative 50,000 garments per month. Under manual inspection alone (cited in literature at 60 to 70 percent accuracy, missing 20 to 30 percent of defects), baseline rework and second-quality losses averaged an illustrative 4.0 percent (2,000 garments per month) at an illustrative XAF 1,500 per garment, or XAF 3,000,000 per month. [R29]
2. **Observed gross benefit (illustrative inputs):** With operator-assisted vision and constant product mix, rework and second-quality losses fall from the illustrative 4.0 percent baseline to an illustrative 2.6 percent (1,300 garments per month), avoiding 700 defective garments per month and generating an illustrative gross saving of 700 multiplied by XAF 1,500, or XAF 1,050,000 per month. [S36] [S38]
3. **Full lifecycle cost (illustrative inputs):** Deducting monthly recurring and amortised preparation costs under S18 and S39: steward labelling and quarantine review (illustrative 18 hours at XAF 5,000 per hour = XAF 90,000), operator false-positive review time (illustrative 25 hours at XAF 3,000 per hour = XAF 75,000), hardware, backup and licence amortisation over an illustrative 24 months (XAF 260,000 per month), and upfront contract and shadow-validation effort amortised over 24 months (illustrative XAF 3,000,000 divided by 24 = XAF 125,000 per month) yields an illustrative total lifecycle cost of XAF 550,000 per month. [S02] [S18] [S26] [S39]
4. **Gate 8 decision (illustrative arithmetic):** Subtracting the illustrative XAF 550,000 lifecycle cost from the illustrative XAF 1,050,000 gross benefit leaves an illustrative net benefit of XAF 500,000 per month (47.6 percent net retention), justifying scale to a second line in Stage 6. Had lifecycle costs exceeded XAF 1,050,000 per month, Gate 8 would require suspending or redesigning the deployment. [S18] [S36] [S44]

14.16

Months 13 to 18 reuse the spine before adding complexity

When Gate 8 confirms positive net value at month 12, the illustrative final stage from month 13 to month 18 proves the economic core of the spine: reusing governed identifiers, curated events and lineage across additional decisions without rebuilding capture pipelines. [S08] [S09] [S10] [S39] Stage 6 prevents post-pilot pipeline sprawl by requiring canonical reuse before any new analytics tool is approved. [S18] [S19]

14.17 REUSING GOVERNED RECORDS AND SELECTIVE AUTOMATION

During illustrative months 13 to 15, the team connects a second decision to existing curated tables without adding shop-floor forms: canonical order, cotton-lot and energy-meter events feed product carbon footprints (wp16) and provenance certificates (wp05), while governed downtime and quality-state events feed line balancing

and maintenance (wp02, wp07, wp15). [S05] [S08] [S38] [S41] Simultaneously, MLOps automation is added selectively where Stage 4 and Stage 5 logs show recurring manual toil, such as automating weekly slice-evaluation scripts while keeping retraining and promotion strictly human-approved under NIST AI RMF 1.0 and ISO/IEC 23894. [S11] [S14] [S16] [S42]

14.18 CONDITIONAL ADMISSION OF FORECASTING AND EXIT REHEARSAL

Months 15 to 18 also test history-dependent use cases such as demand and inventory forecasting (wp04). [S04] [S43] Although McKinsey's 2025 fashion research cites inventory-cost reductions of 35 to 40 percent and forecast-error cuts of 20 to 50 percent from AI at scale, a Cameroon factory cannot capture those gains until it accumulates clean multi-cycle order and stockout records. [R27] [S04] [S19] Under GFT's roadmap, a forecaster enters Stage 6 shadow testing only if Gate 6 verifies at least 12 to 15 months of governed order and stockout history and the candidate beats a simple moving-average or seasonal comparator. [S04] [S38] [S43] [S44] Before Gate 9 closes at month 18, the plant completes an ISO/IEC 42001 review, retires obsolete artefacts, and migrates one curated dataset and model candidate to an alternate environment using only archived contracts, PROV-O lineage and open exports. [S06] [S09] [S12] [S15] [S32]

14.19

GFT should measure seven capabilities every month

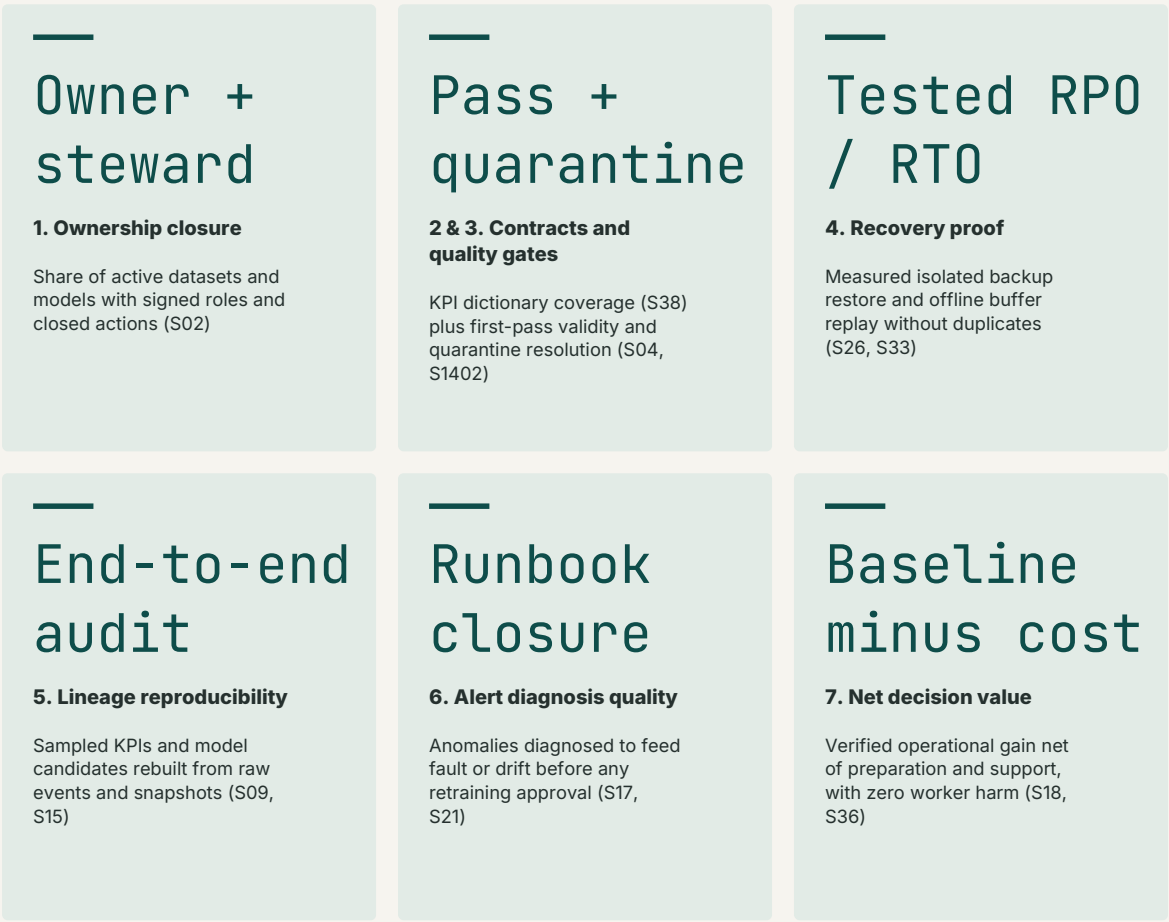
Counting sensors or dashboards rewards volume over decision quality. [S36] [S44] GFT recommends reviewing a derived seven-measure capability pack every month inside the existing operations cadence (wp02): [S02] [S04] [S15] [S17] [S26] [S36] [S1402]

1. **Ownership and stewardship closure rate:** Share of active datasets and models with a signed owner, active steward and zero overdue governance actions, derived from ISO 8000-150 and NIST CSF 2.0. [S02] [S23]
2. **Contract and KPI definition coverage:** Share of operational metrics and model features backed by a versioned ISO 22400-2 KPI dictionary entry and schema contract. [S10] [S38]
3. **Quality-gate pass and quarantine resolution rate:** First-pass validity, completeness, consistency, timeliness and uniqueness rate, median quarantine resolution time, and disclosed `missing` or `estimated` share under ISO/IEC 5259-3 and the W3C Data Quality Vocabulary. [S04] [S41] [S1402]
4. **Tested recovery and offline continuity evidence:** Pass status and measured RPO/RTO of the latest isolated backup-restore and offline store-and-forward replay drill. [S26] [S27] [S33]
5. **Candidate and KPI lineage reproducibility:** Share of sampled KPIs and model candidates reconstructed end-to-end from raw events, snapshots and code within one shift under PROV-O, PROV-DM and MLOps lineage rules. [S09] [S15] [S1403]

6. **Alert diagnosis and runbook response quality:** Share of production anomalies acknowledged on time, diagnosed between feed fault and concept drift, and closed via runbook without unscheduled automatic retraining. [S17] [S21] [S27]
7. **Verified net decision value and trust balance:** Audited operational benefit against the frozen baseline net of full preparation and support cost, paired with zero unresolved worker or customer harm reports. [S18] [S30] [S36] [S39]

EXHIBIT 63 / DERIVED MONTHLY GOVERNANCE AND MLOPS CAPABILITY PACK

Seven capability measures replace a vanity count of dashboards



DERIVED [S02] [S04] [S15] [S17] [S26] [S36] [S1402] Derived monthly measurement pack synthesised from S02, S04, S15, S17, S26, S36 and S1402. Tile 2 pairs semantic contract coverage (Measure 2) and quality-gate resolution (Measure 3) so all seven capabilities are surfaced without universal numeric targets before site baseline.

The six tiles in **Exhibit 63**, p. 239 surface all seven monthly capability measures (combining Measures 2 and 3 in Tile 2 for legibility) without prescribing arbitrary universal percentages before a site's Stage 1 baseline is recorded. [S04] [S36] [S1402] Imposing unmeasured targets on day one incentivises fabricated entries; publishing actual baselines, exception counts and estimation shares builds durable trust. [S19] [S41] [S1402]

GFT NOTE

Practical recommendation for GFT's Cameroon textile programme: Do not hire a separate data department or buy an enterprise data lake in year one. Embed the six governance roles (wp14) inside existing plant positions, start on day 1 with the manual performance logs in wp02, and enforce the nine gates in **Exhibit 62**, p. 235. Use maturity and kpi-cascade to baseline capabilities 1 to 3 by day 30, edge-architecture and risk-matrix to verify offline buffering and OT/privacy controls by day 90, and use-case-portfolio, inspection-economics and cold-start-forecast to screen models at month 7 and month 15. [S02] [S16] [S36] [S44]

14.20

The closing recommendation is disciplined portability

Across all 14 chapters of wp14, one principle governs: a small factory captures durable value from data and machine learning only by managing the path from physical event to decision as one auditable production spine. [S05] [S09] [S15] [S43] That recommendation distils into five rules of **disciplined portability**:

- **Own the operational meaning:** Keep canonical identifiers, event schemas, state machines and the ISO 22400-2 KPI dictionary under factory version control rather than inside vendor code. [S05] [S06] [S08] [S38]
- **Buy commodity capability, not black-box authority:** Favour replaceable storage, messaging and pre-trained vision components while contracting for open-format export rights so the plant can exit any vendor with its history intact. [S16] [S20] [S32] [S42]
- **Preserve immutable raw evidence and provenance:** Never overwrite shop-floor observations; append superseding records, tag missing and estimated states explicitly, and retain PROV-O and PROV-DM lineage from raw event to model and KPI. [S04] [S09] [S15] [S1403]
- **Retain human authority over thresholds and worker impacts:** Require operating and risk owners to set thresholds, inspect local slices, approve releases and prevent automated-only worker decisions or unsafe OT changes. [S11] [S14] [S22] [S30]
- **Make every release reversible and every investment falsifiable:** Rehearse offline buffering, encrypted backup recovery and model rollback before release, and evaluate every deployment against its frozen baseline net of full lifecycle cost. [S15] [S18] [S26] [S36] [S39] [S44]

WARNING

Unresolved external constraints requiring project-specific verification: Implementing this roadmap in Cameroon requires resolving three local constraints before capital commitment: obtaining current Cameroonian legal counsel on Personal Data Protection Authority authorisation procedures under Law No. 2024/017 of 23 December 2024 (S31); replacing unrefreshed or undated inherited sector baselines (R06, R07, R11, R12, R13) with audited site measurements in Stage 1; and metering local power and telecom outage durations to size offline buffers and verify recovery targets (S22, S26, S33). [R06] [R07] [R12] [S22] [S26] [S31] [S33]

14.21 HOW THE INDUSTRIAL DATA SPINE CONNECTS THE GFT LIBRARY

By governing records once, wp14 supplies the evidentiary foundation across the GFT white-paper library and tools: daily management (wp02, kpi-cascade, oee-line), automated defect inspection (wp03, inspection-economics), cold-start demand forecasting (wp04, cold-start-forecast), lot traceability (wp05), utility efficiency (wp06), Scope 1, 2 and 3 carbon accounting (wp16, footprint), maintenance reliability (wp07), workforce capability (wp09), laboratory workflows (wp12), industrial engineering (wp15), compliance (wp17) and the edge-to-cloud topology (wp18, edge-architecture, build-or-buy). [S05] [S08] [S09] [S38] [S41] A small factory earns the right to scale artificial intelligence not by building a data bureaucracy, but by capturing the truth of its work on the first shift, governing every change of meaning, proving recovery under failure, and keeping every model accountable to the people who run the line. [S02] [S11] [S26] [S36]

The factory that governs its records from the first shift does not need a data department to trust its decisions.

CHAPTER SOURCE NOTES

S02 International Organization for Standardization. 2022.

S16 Microsoft Learn, Azure Architecture Center. 2026-06-03.

S36 World Bank. 2023.

R06 Cameroon Trade Hub. Publication date unrecorded; inherited source-text check.

R07 Ecofin Agency. Publication date unrecorded; inherited source-text check.

S26 CISA, MS-ISAC, NSA and FBI. 2023.

S44 The Journal of Technology Transfer, Springer Nature. 2024-07-27.

S38 International Organization for Standardization. 2014.

S37 World Bank. 2025-05-20.

S34 OECD. 2026-06-26.

S23 National Institute of Standards and Technology. 2024-02-26.

S11 National Institute of Standards and Technology. 2023-01-26.

S14 ISO/IEC. 2023.

S35 OECD. 2022-11-21.

S05 International Electrotechnical Commission. 2013.

S08 GS1 AISBL. 2022.

S01 International Organization for Standardization. 2022.

S29 European Parliament and Council of the European Union. 2016-04-27.
S31 Presidency of the Republic of Cameroon. 2024-12-23.
S1401 World Wide Web Consortium. 2017-01-31.
S30 International Labour Office. 1997.
S43 Frontiers in Artificial Intelligence. 2021-06-14.
S04 ISO/IEC. 2024.
S15 Google Cloud. 2024-08-28.
S42 SCITEPRESS, IN4PL 2022. 2022.
S09 World Wide Web Consortium. 2013-04-30.
S10 National Institute of Standards and Technology. 2016-02-23.
S1403 World Wide Web Consortium. 2013-04-30.
S03 ISO/IEC. 2008.
S19 Association for Computing Machinery, CHI 2021. 2021.
S41 GHG Protocol, World Resources Institute and World Business Council for Sustainable Development. 2013.
S1402 World Wide Web Consortium. 2016-12-15.
S27 National Institute of Standards and Technology. 2025.
S33 Microsoft. 2024-06-10.
S07 OASIS Open. 2019-03-07.
S39 National Institute of Standards and Technology. 2025.
S18 Advances in Neural Information Processing Systems. 2015.
S06 OPC Foundation. 2025-10-22.
S25 National Institute of Standards and Technology. 2020-12-10.
S32 European Parliament and Council of the European Union. 2023-12-13.
S24 Cybersecurity and Infrastructure Security Agency. 2025.
S12 ISO/IEC. 2023.
R29 MDPI; industry inspection literature. Publication date unrecorded; inherited source-text check.
S22 National Institute of Standards and Technology. 2023.
S13 ISO/IEC. 2023.
S20 arXiv, author version of an IEEE Access article. 2022-05-14.
S17 Microsoft Learn, Azure Machine Learning. 2026-01-27.
S28 International Society of Automation. 2026.
S40 International Labour Organization and UNESCO. 2020.
S21 arXiv, author manuscript of an IEEE Transactions on Knowledge and Data Engineering article. 2020-04-13.
R27 McKinsey, The State of Fashion 2025. Publication date unrecorded; inherited source-text check.
R12 Business in Cameroon. Publication date unrecorded; inherited source-text check.

Provenance travels with the claim.

Available inherited records and all chapter-local source records are graded below. Identical source identities are consolidated without losing citation aliases. Missing annex references are excluded. Unrecorded dates are disclosed rather than invented.

R01 Inherited evidence R01 GRADE C

Business in Cameroon. Publication date unrecorded. businessincameroon.com

Cameroon seed-cotton production 2022-2024: 330,000 t (2022), 314,455 t (2023), about 340,000 t (2024).

Inherited source-text check; original source not independently re-opened for this edition.

R02 Inherited evidence R02 GRADE C

Fibre2Fashion; Business in Cameroon. Publication date unrecorded. fibre2fashion.com; businessincameroon.com

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

R03 Inherited evidence R03 GRADE C

SODECOTON. Publication date unrecorded. sodecoton.cm

Company profile: more than 220,000 cotton producers supported in northern Cameroon; nine ginning plants; two oil mills.

Inherited source-text check; original source not independently re-opened for this edition.

R04 Inherited evidence R04 GRADE C

Financial Fortune Media. Publication date unrecorded. financialfortunemedia.com

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

R05 Inherited evidence R05 GRADE C

Ecofin Agency. Publication date unrecorded. ecofinagency.com

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

R06 Inherited evidence R06 GRADE C

Cameroon Trade Hub. Publication date unrecorded. cameroontradehub.cm

About 95 percent of cotton fibre exported raw against 5 percent processed locally; industrialisation target of 50 percent local processing.

Inherited source-text check; original source not independently re-opened for this edition.

R07 Inherited evidence R07 GRADE C

Ecofin Agency. Publication date unrecorded. ecofinagency.com

SND30 restates the 50 percent local-processing target by 2030.

Inherited source-text check; original source not independently re-opened for this edition.

R08 Inherited evidence R08 GRADE C

Business in Cameroon. Publication date unrecorded. businessincameroon.com

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

R09 Inherited evidence R09 GRADE C

Ecofin Agency. Publication date unrecorded. ecofinagency.com

Camtext SA integrated cotton-to-textile project, US$300 million (180 billion XAF); Panafritex and Arise with SODECOTON, CNPS and Marlo; 12 November 2025.

Inherited source-text check; original source not independently re-opened for this edition.

R10 Inherited evidence R10 GRADE C

Ecofin Agency. Publication date unrecorded. ecofinagency.com

Camtext SA: 12,000 t a year processed; about 12,000 jobs at the Dibamba hub and 3,000 at Garoua; return on investment near 11 years; optimal capacity around 2033.

Inherited source-text check; original source not independently re-opened for this edition.

R11 Inherited evidence R11 GRADE C

Business in Cameroon; MMI News. Publication date unrecorded. businessincameroon.com; mimimefoinfos.com

CICAM market share from about 85 percent three decades ago to about 5 percent today.

Inherited source-text check; original source not independently re-opened for this edition.

R12 Inherited evidence R12 GRADE C

Business in Cameroon. Publication date unrecorded. businessincameroon.com

CICAM capacity utilisation, July 2023 report: Garoua 33.7 percent, CICAM 1 48.9 percent, CICAM 2 10.83 percent; underutilisation cost about CFAF 2 billion in 2020.

Inherited source-text check; original source not independently re-opened for this edition.

R13 Inherited evidence R13 GRADE C

Business in Cameroon; MMI News. Publication date unrecorded. businessincameroon.com; mimimefoinfos.com

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

R14 Inherited evidence R14 GRADE C

WageIndicator. Publication date unrecorded. wageindicator.org

Cameroon minimum wage (SMIG) XAF 41,875 a month in 2023, raised to XAF 43,969 a month from 23 February 2024, about US$76, a rise of 5 percent.

Inherited source-text check; original source not independently re-opened for this edition.

R15 Inherited evidence R15 GRADE C

wage.is. Publication date unrecorded. wage.is/cameroon

Average salary across sectors in Cameroon about US$359 a month.

Inherited source-text check; original source not independently re-opened for this edition.

R16 Inherited evidence R16 GRADE C

WageIndicator. Publication date unrecorded. wageindicator.org

Informal share of employment about 90 percent; standard 40-hour week, 48-hour maximum, overtime at 1.2 times.

Inherited source-text check; original source not independently re-opened for this edition.

R17 Inherited evidence R17 GRADE C

just-style. Publication date unrecorded. just-style.com

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

R18 Inherited evidence R18 GRADE B

Congressional Research Service IF10149; AGOA (Wikipedia). Publication date unrecorded. congress.gov; en.wikipedia.org

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

R23 Inherited evidence R23 GRADE C

MacMap; AfCFTA. Publication date unrecorded. macmap.org

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

R24 Inherited evidence R24 GRADE C

Mordor Intelligence. Publication date unrecorded. mordorintelligence.com

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

R25 Inherited evidence R25 GRADE C

Market.us. Publication date unrecorded. market.us

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

R26 Inherited evidence R26 GRADE C

Statista. Publication date unrecorded. statista.com

Cameroon clothing and footwear spend about US$94.70 a person in 2024 and about 8.0 pieces a person; apparel eCommerce revenue US$20.32 million (2023) to US$33.61 million (2027) at 13.41 percent a year; user penetration 18.0 percent (2023) to 22.7 percent (2027).

Inherited source-text check; original source not independently re-opened for this edition.

R27 Inherited evidence R27 GRADE B

McKinsey, The State of Fashion 2025. Publication date unrecorded. mckinsey.com

AI at scale in fashion: inventory cost down 35 to 40 percent; forecast accuracy up 20 to 30 percent; revenue per customer up 15 to 25 percent; forecast errors cut 20 to 50 percent; lost sales down up to 65 percent; 75 percent of executives prioritising; sector operating-profit potential US$150 to 275 billion over three to five years.

Inherited source-text check; original source not independently re-opened for this edition.

R28 Inherited evidence R28 GRADE C

DataIntelo. Publication date unrecorded. dataintelo.com

Fabric-defect-detection AI market US$523.4 million (2024) to US$6,522.1 million (2033) at 33.8 percent a year; Asia-Pacific about 46 percent in 2024.

Inherited source-text check; original source not independently re-opened for this edition.

R29 Inherited evidence R29 GRADE C

MDPI; industry inspection literature. Publication date unrecorded. mdpi.com

Manual defect inspection 60 to 70 percent accurate, missing 20 to 30 percent of defects; computer vision above 90 percent across more than 40 defect types; line-scan inspection up to 60 metres a minute.

Inherited source-text check; original source not independently re-opened for this edition.

R30 Inherited evidence R30 GRADE C

Design Journal (Taylor and Francis). Publication date unrecorded. tandfonline.com

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

R31 Inherited evidence R31 GRADE C

arXiv 2509.04705. Publication date unrecorded. arxiv.org/abs/2509.04705

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

R32 Inherited evidence R32 GRADE C

Wise; XE. Publication date unrecorded. wise.com; xe.com

CFA franc to US dollar reference rates; franc pegged to the euro at 655.957. The GFT study's method note places the franc near 590 to 610 per US dollar across the 2023 study window and near 563 to 577 more recently, and treats currency conversions as estimates.

Inherited source-text check; original source not independently re-opened for this edition.

c01-S01 ISO 8000-1:2022: Data quality, Part 1: Overview GRADE A

International Organization for Standardization. 2022. https://www.iso.org/obp/ui/en/#!iso:std:81745:en

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S02 ISO 8000-150:2022: Data quality, Part 150: Data quality management: Roles and responsibilities GRADE A

International Organization for Standardization. 2022. https://www.iso.org/standard/80753.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S03 ISO/IEC 25012:2008: Software engineering, Software product Quality Requirements and Evaluation, Data quality model GRADE A

ISO/IEC. 2008. https://www.iso.org/obp/ui#!iso:std:iso-iec:25012:ed-1:v1:en

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S04 **ISO/IEC 5259-3:2024: Artificial intelligence, Data quality for analytics and machine learning, Part 3: Data quality management requirements and guidelines** GRADE A

ISO/IEC. 2024. https://www.iso.org/standard/81092.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c01-S05 **IEC 62264-1:2013, Edition 2.0: Enterprise-control system integration, Part 1: Models and terminology** GRADE A

International Electrotechnical Commission. 2013. https://webstore.iec.ch/preview/info_iec62264-1%7Bed2.0%7Db.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c01-S06 **OPC Unified Architecture, Part 1: Overview and Concepts, OPC 10000-1, version 1.05.06** GRADE A

OPC Foundation. 2025-10-22. https://reference.opcfoundation.org/specs/OPC-10000-1/full

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c01-S07 **MQTT Version 5.0** GRADE A

OASIS Open. 2019-03-07. https://docs.oasis-open.org/mqtt/mqtt/v5.0/os/mqtt-v5.0-os.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c01-S08 **EPCIS Standard, Release 2.0** GRADE A

GS1 AISBL. 2022. https://ref.gs1.org/standards/epcis/2.0.0/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c01-S09 **PROV-O: The PROV Ontology** GRADE A

World Wide Web Consortium. 2013-04-30. https://www.w3.org/TR/2013/REC-prov-o-20130430/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c01-S10 **Current Standards Landscape for Smart Manufacturing Systems, NISTIR 8107** GRADE A

National Institute of Standards and Technology. 2016-02-23. https://www.nist.gov/publications/current-standards-landscape-smart-manufacturing-systems

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c01-S11 **Artificial Intelligence Risk Management Framework, AI RMF 1.0, NIST AI 100-1** GRADE A

National Institute of Standards and Technology. 2023-01-26. https://doi.org/10.6028/NIST.AI.100-1

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S13 **ISO/IEC 5338:2023: Information technology, Artificial intelligence, AI system life cycle processes** GRADE A

ISO/IEC. 2023. https://www.iso.org/standard/81118.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c01-S15 **MLOps: Continuous delivery and automation pipelines in machine learning** GRADE B

Google Cloud. 2024-08-28. https://docs.cloud.google.com/architecture/mlops-continuous-delivery-and-automation-pipelines-in-machine-learning

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c01-S16 **MLOps Maturity Model** GRADE B

Microsoft Learn, Azure Architecture Center. 2026-06-03. https://learn.microsoft.com/en-us/azure/architecture/ai-ml/guide/mlops-maturity-model

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c01-S17 **Model monitoring in production** GRADE B

Microsoft Learn, Azure Machine Learning. 2026-01-27. https://learn.microsoft.com/en-us/azure/machine-learning/concept-model-monitoring?view=azureml-api-2

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c01-S18 **Hidden Technical Debt in Machine Learning Systems** GRADE B

Advances in Neural Information Processing Systems. 2015. https://papers.nips.cc/paper_files/paper/2015/hash/86df7dcfd896fcaf2674f757a2463eba-Abstract.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c01-S19 **Everyone wants to do the model work, not the data work: Data Cascades in High-Stakes AI** GRADE B

Association for Computing Machinery, CHI 2021. 2021. https://dl.acm.org/doi/10.1145/3411764.3445518

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c01-S21 **Learning under Concept Drift: A Review** GRADE B

arXiv, author manuscript of an IEEE Transactions on Knowledge and Data Engineering article. 2020-04-13. https://arxiv.org/abs/2004.05785

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c01-S22 **Guide to Operational Technology Security, NIST SP 800-82 Revision 3** GRADE A

National Institute of Standards and Technology. 2023. https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-82r3.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S23 **The NIST Cybersecurity Framework 2.0** GRADE A

National Institute of Standards and Technology. 2024-02-26. https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.29.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S25 **Security and Privacy Controls for Information Systems and Organizations, NIST SP 800-53 Revision 5** GRADE A

National Institute of Standards and Technology. 2020-12-10. https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-53r5.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S26 **#StopRansomware Guide** GRADE A

CISA, MS-ISAC, NSA and FBI. 2023. https://www.cisa.gov/stopransomware/ransomware-guide

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S27 **Incident Response Recommendations and Considerations for Cybersecurity Risk Management: A CSF 2.0 Community Profile, NIST SP 800-61 Revision 3** GRADE A

National Institute of Standards and Technology. 2025. https://nvlpubs.nist.gov/nistpubs/specialpublications/nist.sp.800-61r3.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S29 **Regulation (EU) 2016/679, General Data Protection Regulation** GRADE A

European Parliament and Council of the European Union. 2016-04-27. https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32016R0679

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S30 **Protection of workers' personal data: An ILO code of practice** GRADE A

International Labour Office. 1997. https://www.ilo.org/media/270771/download

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S31 **Law No. 2024/017 of 23 December 2024 relating to personal data protection in Cameroon** GRADE A

Presidency of the Republic of Cameroon. 2024-12-23. https://prc.cm/files/9b/df/2c/e818fedc7d5568778f884ea2886bea7d.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S32 **Regulation (EU) 2023/2854 on harmonised rules on fair access to and use of data, Data Act** GRADE A

European Parliament and Council of the European Union. 2023-12-13. https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32023R2854

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S33 **Understand extended offline capabilities for IoT Edge devices, modules, and child devices** GRADE B

Microsoft. 2024-06-10. https://learn.microsoft.com/en-us/azure/iot-edge/offline-capabilities?view=iotedge-1.5

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S34 **Uptake of digital tools among SMEs: Western Balkans Enterprise Survey** GRADE B

OECD. 2026-06-26. https://www.oecd.org/en/publications/western-balkans-enterprise-survey_a1b14188-en/full-report/component-4.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S35 **Employee empowerment through digital tools: The case of a German bakery SME** GRADE B

OECD. 2022-11-21. https://www.oecd.org/en/publications/digital-for-smes-case-studies_d65e57f4-en/employee-empowerment-through-digital-tools-the-case-of-a-german-bakery-sme_9efbd5fa-en.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S36 **The Missing Element of Firm Digitalization: Lessons for the EU Member States: Best practices for enhancing MSMEs Digitalization and Managerial Capacity Building** GRADE B

World Bank. 2023. https://documents1.worldbank.org/curated/en/099110723083021800/pdf/P17803602b993104209912039ee0f5a13d3.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S37 **Challenges with digital technology adoption by SMEs: A case study from Viet Nam** GRADE B

World Bank. 2025-05-20. https://blogs.worldbank.org/en/allaboutfinance/challenges-with-digital-technology-adoption-by-smes--a-case-stud

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S38 **ISO 22400-2:2014: Automation systems and integration, Key performance indicators for manufacturing operations management, Part 2: Definitions and descriptions** GRADE A

International Organization for Standardization. 2014. https://www.iso.org/standard/54497.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S39 **Artificial Intelligence with Open and Scaled Data Sharing in Semiconductor Manufacturing: Workshop Report** GRADE A

National Institute of Standards and Technology. 2025. https://doi.org/10.6028/NIST.AMS.100-72

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S41 Technical Guidance for Calculating Scope 3 Emissions, version 1.0 GRADE B

GHG Protocol, World Resources Institute and World Business Council for Sustainable Development. 2013. https://ghgprotocol.org/sites/default/files/ghgp/standards/Scope3_Calculation_Guidance_0.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S42 Is MLOps different in Industry 4.0? General and Specific Challenges GRADE B

SCITEPRESS, IN4PL 2022. 2022. https://doi.org/10.5220/0011589600003329

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S43 Ensuring the Robustness and Reliability of Data-Driven Knowledge Discovery Models in Production and Manufacturing GRADE B

Frontiers in Artificial Intelligence. 2021-06-14. https://doi.org/10.3389/frai.2021.576892

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S44 Supporting SME companies in mapping out AI potential: a Finnish AI development case GRADE B

The Journal of Technology Transfer, Springer Nature. 2024-07-27. https://doi.org/10.1007/s10961-024-10122-5

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S0101 Data Analytics for Smart Manufacturing Systems GRADE A

National Institute of Standards and Technology. 2025-03-26. https://www.nist.gov/programs-projects/data-analytics-smart-manufacturing-systems

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S0102 Data Catalog Vocabulary (DCAT) - Version 3 GRADE A

World Wide Web Consortium. 2024-08-22. https://www.w3.org/TR/vocab-dcat-3/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c02-S01 Same source record as c01-S01.

c02-S02 Same source record as c01-S02.

c02-S04 Same source record as c01-S04.

c02-S05 Same source record as c01-S05.

c02-S08 Same source record as c01-S08.

c02-S09 Same source record as c01-S09.

c02-S11 Same source record as c01-S11.

c02-S12 ISO/IEC 42001:2023: Information technology, Artificial intelligence, Management system GRADE A

ISO/IEC. 2023. https://www.iso.org/standard/81230.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c02-S13 Same source record as c01-S13.

c02-S14 ISO/IEC 23894:2023: Information technology, Artificial intelligence, Guidance on risk management GRADE A

ISO/IEC. 2023. https://www.iso.org/standard/77304.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c02-S15 Same source record as c01-S15.

c02-S16 Same source record as c01-S16.

c02-S18 Same source record as c01-S18.

c02-S19 Same source record as c01-S19.

c02-S22 Same source record as c01-S22.

c02-S23 Same source record as c01-S23.

c02-S25 Same source record as c01-S25.

c02-S27 Same source record as c01-S27.

c02-S29 Same source record as c01-S29.

c02-S30 Same source record as c01-S30.

c02-S31 Same source record as c01-S31.

c02-S36 Same source record as c01-S36.

c02-S37 Same source record as c01-S37.

c02-S38 Same source record as c01-S38.

c02-S43 Same source record as c01-S43.

c02-S44 Same source record as c01-S44.

c02-S20 Machine Learning Operations: Overview, Definition, and Architecture GRADE B

arXiv, author version of an IEEE Access article. 2022-05-14. https://arxiv.org/abs/2205.02302

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c02-S35 Same source record as c01-S35.

c02-S41 Same source record as c01-S41.

c02-S0201 NIST AI RMF Playbook: Govern GRADE A

National Institute of Standards and Technology. 2023-01-26. https://airc.nist.gov/airmf-resources/playbook/govern/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c02-S0202 Data on the Web Best Practices GRADE A

World Wide Web Consortium. 2017-01-31. https://www.w3.org/TR/2017/REC-dwbp-20170131/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c03-S01 Same source record as c01-S01.

c03-S02 Same source record as c01-S02.

c03-S04 Same source record as c01-S04.

c03-S05 Same source record as c01-S05.

c03-S06 Same source record as c01-S06.

c03-S07 Same source record as c01-S07.

c03-S08 Same source record as c01-S08.

c03-S09 Same source record as c01-S09.

c03-S10 Same source record as c01-S10.

c03-S15 Same source record as c01-S15.

c03-S29 Same source record as c01-S29.

c03-S31 Same source record as c01-S31.

c03-S38 Same source record as c01-S38.

c03-S39 Same source record as c01-S39.

c03-S0301 Data Catalog Vocabulary (DCAT) - Version 3 GRADE A

World Wide Web Consortium. 2024-08-22. https://www.w3.org/TR/2024/REC-vocab-dcat-3-20240822/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c03-S0302 Metadata Vocabulary for Tabular Data GRADE A

World Wide Web Consortium. 2015-12-17. https://www.w3.org/TR/2015/REC-tabular-metadata-20151217/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c04-S01 Same source record as c01-S01.

c04-S02 Same source record as c01-S02.

c04-S03 Same source record as c01-S03.

c04-S04 Same source record as c01-S04.

c04-S05 Same source record as c01-S05.

c04-S06 Same source record as c01-S06.

c04-S08 Same source record as c01-S08.

c04-S09 Same source record as c01-S09.

c04-S10 Same source record as c01-S10.

c04-S11 Same source record as c01-S11.

c04-S15 Same source record as c01-S15.

c04-S18 Same source record as c01-S18.

c04-S19 Same source record as c01-S19.

c04-S36 Same source record as c01-S36.

c04-S38 Same source record as c01-S38.

c04-S39 Same source record as c01-S39.

c04-S41 Same source record as c01-S41.

c04-S43 Same source record as c01-S43.

c04-S44 Same source record as c01-S44.

c04-S0401 Data on the Web Best Practices: Data Quality Vocabulary GRADE A

World Wide Web Consortium. 2016-12-15. https://www.w3.org/TR/vocab-dqv/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c04-S0402 Data on the Web Best Practices GRADE A

World Wide Web Consortium. 2017-01-31. https://www.w3.org/TR/dwbp/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c04-S0403 PROV Model Primer GRADE A

World Wide Web Consortium. 2013-04-30. https://www.w3.org/TR/prov-primer/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c05-S01 Same source record as c01-S01.

c05-S02 Same source record as c01-S02.

c05-S03 Same source record as c01-S03.

c05-S04 Same source record as c01-S04.

c05-S05 Same source record as c01-S05.

c05-S06 Same source record as c01-S06.

c05-S07 Same source record as c01-S07.

c05-S08 Same source record as c01-S08.

c05-S09 Same source record as c01-S09.

c05-S10 Same source record as c01-S10.

c05-S15 Same source record as c01-S15.

c05-S16 Same source record as c01-S16.

c05-S17 Same source record as c01-S17.

c05-S18 Same source record as c01-S18.

c05-S19 Same source record as c01-S19.

c05-S20 Same source record as c02-S20.

c05-S22 Same source record as c01-S22.

c05-S25 Same source record as c01-S25.

c05-S26 Same source record as c01-S26.

c05-S27 Same source record as c01-S27.

c05-S28 ISA/IEC 62443 Series of Standards GRADE A

International Society of Automation. 2026. https://www.isa.org/standards-and-publications/isa-standards/isa-iec-62443-series-of-standards

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c05-S29 Same source record as c01-S29.

c05-S32 Same source record as c01-S32.

c05-S33 Same source record as c01-S33.

c05-S34 Same source record as c01-S34.

c05-S35 Same source record as c01-S35.

c05-S37 Same source record as c01-S37.

c05-S38 Same source record as c01-S38.

c05-S39 Same source record as c01-S39.

c05-S40 The Digitization of TVET and Skills Systems GRADE B

International Labour Organization and UNESCO. 2020. https://www.ilo.org/media/396161/download

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c05-S42 Same source record as c01-S42.

c05-S44 Same source record as c01-S44.

c05-S0501 PROV-DM: The PROV Data Model GRADE A

World Wide Web Consortium. 2013-04-30. https://www.w3.org/TR/2013/REC-prov-dm-20130430/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c05-S0502 Same source record as c02-S0202.

c05-S0503 Data Quality Vocabulary (DQV) GRADE A

World Wide Web Consortium. 2016-12-15. https://www.w3.org/TR/2016/NOTE-vocab-dqv-20161215/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c06-S02 Same source record as c01-S02.

c06-S04 Same source record as c01-S04.

c06-S05 Same source record as c01-S05.

c06-S06 Same source record as c01-S06.

c06-S07 Same source record as c01-S07.

c06-S08 Same source record as c01-S08.

c06-S09 Same source record as c01-S09.

c06-S15 Same source record as c01-S15.

c06-S29 Same source record as c01-S29.

c06-S31 Same source record as c01-S31.

c06-S33 Same source record as c01-S33.

c06-S38 Same source record as c01-S38.

c06-S0601 EPCIS Standard, Release 2.0, Section 7.4.1: EPCISEvent GRADE A

GS1 AISBL. 2022. https://ref.gs1.org/standards/epcis/2.0.0/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c06-S0602 EPCIS Standard, Release 2.0, Error declarations GRADE A

GS1 AISBL. 2022. https://ref.gs1.org/standards/epcis/2.0.0/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c06-S0603 RFC 9562: Universally Unique IDentifiers (UUIDs) GRADE A

Internet Engineering Task Force. 2024-05. https://www.rfc-editor.org/rfc/rfc9562.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c06-S0604 PROV-O: The PROV Ontology, Section 3.2: Expanded Terms GRADE A

World Wide Web Consortium. 2013-04-30. https://www.w3.org/TR/2013/REC-prov-o-20130430/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c07-S02 Same source record as c01-S02.

c07-S05 Same source record as c01-S05.

c07-S06 Same source record as c01-S06.

c07-S07 Same source record as c01-S07.

c07-S09 Same source record as c01-S09.

c07-S11 Same source record as c01-S11.

c07-S14 Same source record as c02-S14.

c07-S15 Same source record as c01-S15.

c07-S22 Same source record as c01-S22.

c07-S23 Same source record as c01-S23.

c07-S24 Cross-Sector Cybersecurity Performance Goals, Version 2.0 GRADE A

Cybersecurity and Infrastructure Security Agency. 2025. https://www.cisa.gov/sites/default/files/2025-12/CPG_Report_2.0_508c.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c07-S25 Same source record as c01-S25.

c07-S26 Same source record as c01-S26.

c07-S27 Same source record as c01-S27.

c07-S28 Same source record as c05-S28.

c07-S29 Same source record as c01-S29.

c07-S30 Same source record as c01-S30.

c07-S31 Same source record as c01-S31.

c07-S32 Same source record as c01-S32.

c07-S33 Same source record as c01-S33.

c07-S0701 Zero Trust Architecture, NIST SP 800-207 GRADE A

National Institute of Standards and Technology. 2020-08. https://csrc.nist.gov/pubs/sp/800/207/final

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c07-S0702 Adversarial Machine Learning: A Taxonomy and Terminology of Attacks and Mitigations, NIST AI 100-2 E2023 GRADE A

National Institute of Standards and Technology. 2024-01. https://csrc.nist.gov/pubs/ai/100/2/e2023/final

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c07-S0703 OPC Unified Architecture, Part 2: Security Model, OPC 10000-2, version 1.05.06 GRADE A

OPC Foundation. 2025-10-22. https://reference.opcfoundation.org/specs/OPC-10000-2/full

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c08-S05 Same source record as c01-S05.

---

c08-S06 Same source record as c01-S06.

---

c08-S07 Same source record as c01-S07.

---

c08-S08 Same source record as c01-S08.

---

c08-S09 Same source record as c01-S09.

---

c08-S10 Same source record as c01-S10.

---

c08-S15 Same source record as c01-S15.

---

c08-S18 Same source record as c01-S18.

---

c08-S22 Same source record as c01-S22.

---

c08-S25 Same source record as c01-S25.

---

c08-S26 Same source record as c01-S26.

---

c08-S27 Same source record as c01-S27.

---

c08-S32 Same source record as c01-S32.

---

c08-S33 Same source record as c01-S33.

---

c08-S38 Same source record as c01-S38.

---

c08-S0801 Contingency Planning Guide for Federal Information Systems, NIST SP 800-34 Rev. 1 GRADE A

National Institute of Standards and Technology. 2010-11-11. https://csrc.nist.gov/pubs/sp/800/34/r1/upd1/final

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c08-S0802 **Data Integrity: Recovering from Ransomware and Other Destructive Events, NIST SP 1800-11** GRADE A

National Institute of Standards and Technology. 2020-09. https://csrc.nist.gov/pubs/sp/1800/11/final

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c08-S0803 **Cloud Computing Synopsis and Recommendations, NIST SP 800-146** GRADE A

National Institute of Standards and Technology. 2012-05-29. https://csrc.nist.gov/pubs/sp/800/146/final

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c09-S02 Same source record as c01-S02.

c09-S04 Same source record as c01-S04.

c09-S05 Same source record as c01-S05.

c09-S08 Same source record as c01-S08.

c09-S09 Same source record as c01-S09.

c09-S11 Same source record as c01-S11.

c09-S12 Same source record as c02-S12.

c09-S13 Same source record as c01-S13.

c09-S14 Same source record as c02-S14.

c09-S16 Same source record as c01-S16.

c09-S18 Same source record as c01-S18.

c09-S19 Same source record as c01-S19.

c09-S22 Same source record as c01-S22.

c09-S23 Same source record as c01-S23.

c09-S29 Same source record as c01-S29.

c09-S30 Same source record as c01-S30.

c09-S31 Same source record as c01-S31.

c09-S32 Same source record as c01-S32.

c09-S35 Same source record as c01-S35.

c09-S36 Same source record as c01-S36.

c09-S37 Same source record as c01-S37.

c09-S38 Same source record as c01-S38.

c09-S39 Same source record as c01-S39.

c09-S40 Same source record as c05-S40.

c09-S42 Same source record as c01-S42.

c09-S43 Same source record as c01-S43.

c09-S44 Same source record as c01-S44.

c09-S0901 **Forecasting: Principles and Practice, 3rd edition, Section 5.2: Some simple forecasting methods** GRADE B

OTexts. 2021-05-31. https://otexts.com/fpp3/simple-methods.html

Simple benchmark forecasting methods (mean, naive, seasonal naive, drift) as mandatory comparators before adopting complex models.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c09-S0902 **Rules of Machine Learning: Best Practices for ML Engineering** GRADE B

Google for Developers. 2025-08-25. https://developers.google.com/machine-learning/guides/rules-of-ml

Instrumenting metrics and launching simple heuristics before introducing machine learning.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c10-S04 Same source record as c01-S04.

c10-S09 Same source record as c01-S09.

c10-S11 Same source record as c01-S11.

c10-S13 Same source record as c01-S13.

c10-S15 Same source record as c01-S15.

c10-S16 Same source record as c01-S16.

c10-S20 Same source record as c02-S20.

c10-S38 Same source record as c01-S38.

c10-S1001 Leakage and the Reproducibility Crisis in ML-based Science GRADE B

arXiv, Sayash Kapoor and Arvind Narayanan. 2022-07-14. https://arxiv.org/abs/2207.07048

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c11-S02 Same source record as c01-S02.

c11-S04 Same source record as c01-S04.

c11-S09 Same source record as c01-S09.

c11-S11 Same source record as c01-S11.

c11-S12 Same source record as c02-S12.

c11-S13 Same source record as c01-S13.

c11-S14 Same source record as c02-S14.

c11-S15 Same source record as c01-S15.

c11-S16 Same source record as c01-S16.

c11-S17 Same source record as c01-S17.

c11-S18 Same source record as c01-S18.

c11-S19 Same source record as c01-S19.

c11-S21 Same source record as c01-S21.

c11-S22 Same source record as c01-S22.

c11-S23 Same source record as c01-S23.

c11-S25 Same source record as c01-S25.

c11-S26 Same source record as c01-S26.

c11-S27 Same source record as c01-S27.

c11-S29 Same source record as c01-S29.

c11-S30 Same source record as c01-S30.

c11-S31 Same source record as c01-S31.

c11-S35 Same source record as c01-S35.

c11-S36 Same source record as c01-S36.

c11-S37 Same source record as c01-S37.

c11-S38 Same source record as c01-S38.

c11-S40 Same source record as c05-S40.

c11-S42 Same source record as c01-S42.

c11-S43 Same source record as c01-S43.

c11-S44 Same source record as c01-S44.

c11-S1101 Model Cards for Model Reporting GRADE B

arXiv, Conference on Fairness, Accountability, and Transparency. 2019-01-14. https://arxiv.org/abs/1810.03993

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c11-S1102 Same source record as c09-S0902.

c12-S02 Same source record as c01-S02.

c12-S04 Same source record as c01-S04.

c12-S05 Same source record as c01-S05.

c12-S08 Same source record as c01-S08.

c12-S09 Same source record as c01-S09.

c12-S11 Same source record as c01-S11.

c12-S12 Same source record as c02-S12.

c12-S13 Same source record as c01-S13.

c12-S14 Same source record as c02-S14.

c12-S15 Same source record as c01-S15.

c12-S16 Same source record as c01-S16.

c12-S17 Same source record as c01-S17.

c12-S18 Same source record as c01-S18.

c12-S19 Same source record as c01-S19.

c12-S20 Same source record as c02-S20.

c12-S21 Same source record as c01-S21.

c12-S22 Same source record as c01-S22.

c12-S23 Same source record as c01-S23.

c12-S26 Same source record as c01-S26.

c12-S27 Same source record as c01-S27.

c12-S30 Same source record as c01-S30.

c12-S33 Same source record as c01-S33.

c12-S36 Same source record as c01-S36.

c12-S38 Same source record as c01-S38.

c12-S41 Same source record as c01-S41.

c12-S42 Same source record as c01-S42.

c12-S43 Same source record as c01-S43.

c12-S44 Same source record as c01-S44.

c12-S1201 Failing Loudly: An Empirical Study of Methods for Detecting Dataset Shift GRADE B

arXiv, Advances in Neural Information Processing Systems (NeurIPS) 2019. 2018-10-29. https://arxiv.org/abs/1810.11953

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c12-S1202 Operationalizing Machine Learning: An Interview Study GRADE B

arXiv, Proceedings of the VLDB Endowment. 2022-09-16. https://arxiv.org/abs/2209.09125

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c12-S1203 Model Cards for Model Reporting GRADE B

arXiv, ACM Conference on Fairness, Accountability, and Transparency. 2018-10-05. https://arxiv.org/abs/1810.03993

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c13-S04 Same source record as c01-S04.

c13-S11 Same source record as c01-S11.

c13-S18 Same source record as c01-S18.

c13-S30 Same source record as c01-S30.

c13-S35 Same source record as c01-S35.

c13-S36 Same source record as c01-S36.

c13-S37 Same source record as c01-S37.

c13-S38 Same source record as c01-S38.

c13-S39 Artificial Intelligence with Open and Scaled Data Sharing in Semiconductor Manufacturing: Workshop Report GRADE A

National Institute of Standards and Technology. 2025-11-18. https://www.nist.gov/publications/artificial-intelligence-open-and-scaled-data-sharing-semiconductor-manufacturing

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c13-S40 Same source record as c05-S40.

c13-S41 Same source record as c01-S41.

c13-S43 Same source record as c01-S43.

c13-S44 Same source record as c01-S44.

c14-S01 Same source record as c01-S01.

c14-S02 Same source record as c01-S02.

c14-S03 Same source record as c01-S03.

c14-S04 Same source record as c01-S04.

c14-S05 Same source record as c01-S05.

c14-S06 Same source record as c01-S06.

c14-S07 Same source record as c01-S07.

c14-S08 Same source record as c01-S08.

c14-S09 Same source record as c01-S09.

c14-S10 Same source record as c01-S10.

c14-S11 Same source record as c01-S11.

c14-S12 Same source record as c02-S12.

c14-S13 Same source record as c01-S13.

c14-S14 Same source record as c02-S14.

c14-S15 Same source record as c01-S15.

c14-S16 Same source record as c01-S16.

c14-S17 Same source record as c01-S17.

c14-S18 Same source record as c01-S18.

c14-S19 Same source record as c01-S19.

c14-S20 Same source record as c02-S20.

c14-S21 Same source record as c01-S21.

c14-S22 Same source record as c01-S22.

c14-S23 Same source record as c01-S23.

c14-S24 Same source record as c07-S24.

c14-S25 Same source record as c01-S25.

c14-S26 Same source record as c01-S26.

c14-S27 Same source record as c01-S27.

c14-S28 Same source record as c05-S28.

c14-S29 Same source record as c01-S29.

c14-S30 Same source record as c01-S30.

c14-S31 Same source record as c01-S31.

c14-S32 Same source record as c01-S32.

c14-S33 Same source record as c01-S33.

c14-S34 Same source record as c01-S34.

c14-S35 Same source record as c01-S35.

c14-S36 Same source record as c01-S36.

c14-S37 Same source record as c01-S37.

c14-S38 Same source record as c01-S38.

c14-S39 Same source record as c01-S39.

c14-S40 Same source record as c05-S40.

c14-S41 Same source record as c01-S41.

c14-S42 Same source record as c01-S42.

c14-S43 Same source record as c01-S43.

c14-S44 Same source record as c01-S44.

c14-S1401 Same source record as c02-S0202.

c14-S1402 Data on the Web Best Practices: Data Quality Vocabulary GRADE A

World Wide Web Consortium. 2016-12-15. https://www.w3.org/TR/2016/NOTE-vocab-dqv-20161215/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c14-S1403 Same source record as c05-S0501.

Stable definitions make comparisons useful.

Industrial data spine

The governed chain of ownership, identifiers, events, quality controls, storage, models and decisions that connects factory work to reusable evidence.

Data governance

The decision rights, accountability, policies and controls used to manage data throughout its lifecycle.

Data owner

The role accountable for a dataset's purpose, definition, access, quality expectations and retention.

Data steward

The role that performs recurring definition, quality, correction and coordination work close to a business domain.

Data custodian

The technical role responsible for operating storage, access, backup, recovery and related controls without owning business meaning.

Data domain

A bounded subject area such as production, quality, maintenance, orders, people or sustainability used to assign accountability.

Data contract

A versioned promise describing a dataset's schema, semantics, ownership, timing, quality expectations and compatibility rules.

Canonical model

A shared representation of important entities and events that reduces repeated translation between applications while preserving source detail.

Master data

Relatively stable records describing entities such as assets, lines, styles, operations and units that give transactional events their context.

Metadata

Structured information that describes a dataset or artefact, including meaning, origin, version, quality, rights and retention.

Data lineage

The traceable path from source records through transformations to a dataset, KPI, model or decision.

Provenance

Evidence about where a record or artefact came from, which processes changed it and which agents were responsible.

Event time

The time at which the represented operational event occurred, as distinct from when its record arrived or was processed.

Ingestion time

The time at which a system received an event record, used to diagnose delay, replay and late arrival.

Idempotency

The property that processing the same event again produces no additional unintended effect, enabling safe retry and replay.

Schema

The formal structure of a record, including fields, types, required values and constraints.

Data quality

The fitness of data for a stated purpose, assessed through explicit characteristics and measures rather than a general claim of cleanliness.

Validity

The degree to which a record conforms to its defined types, domains, formats and rules.

Completeness

The degree to which required records and fields are present at the grain needed for a decision.

Timeliness

The degree to which data arrive and become usable within the period required by the decision.

Quarantine

A controlled state that isolates failed records from normal consumption while preserving them for diagnosis, correction and reprocessing.

System of record

The authoritative source for a defined class of information, with other copies explicitly subordinate or derived.

Edge computing

Processing close to machines or work activities so capture, inference or control can continue with limited dependence on remote services.

Store-and-forward

Durable local buffering of records during disconnection, followed by controlled synchronisation when communication returns.

OPC UA

OPC Unified Architecture, an industrial interoperability framework that includes information modelling and secure communication.

MQTT

A client-server publish-subscribe messaging transport that requires separate payload schemas and operational semantics.

EPCIS

A GS1 standard for creating and sharing visibility event data within and across enterprises.

ISA-95

The enterprise-control integration framework, standardised in the IEC 62264 series, used to model interfaces and terminology between business and manufacturing systems.

KPI dictionary

The governed register of each indicator's purpose, formula, inputs, unit, time basis, grain, exclusions, owner and version.

MLOps

The practices and controls used to develop, validate, release, monitor and retire machine-learning systems reliably in their operating context.

Model registry

A controlled record of model artefacts, versions, evidence, approval and lifecycle status.

Feature

A model input derived from source data under a documented definition, transformation, time window and missing-value rule.

Dataset snapshot

An immutable, identified version of a training, validation or test dataset with its manifest, quality state and lineage.

Data leakage

The use of information during training or evaluation that would not legitimately be available for the intended prediction, inflating apparent performance.

Concept drift

A change in the underlying data relationship or distribution over time that can invalidate a model's previous mapping.

Shadow deployment

Operating a model on live inputs without allowing its outputs to affect decisions, used to collect release evidence.

Rollback

A tested return to a previously approved model, configuration or operating procedure after a release is suspended.

Human oversight

Informed human authority to set thresholds, interpret limits, intervene, override and stop a system according to consequence.

Recovery point objective

The maximum targeted interval of data loss after disruption, set by business consequence and verified through recovery evidence.

Recovery time objective

The targeted time to restore a service or record set after disruption, verified through an exercised recovery procedure.

Least privilege

Granting a person or process only the access required for its assigned task, for only as long as needed.

Operational technology

Hardware and software that monitor or control physical devices and processes, where performance, reliability and safety are primary constraints.

Data minimisation

Limiting personal data to what is adequate, relevant and necessary for a stated processing purpose.

Counterfactual

The credible account of what would have happened without the model or data intervention, used to assess claimed benefit.

Model retirement

The controlled withdrawal of a model from use while preserving the records needed to explain its historical outputs and decisions.

Follow the terms back to their context.

Page references point to occurrences in chapter prose. When a term does not occur in a chapter, its glossary definition is indexed.

Canonical model 85

Completeness 24 33 66 67 71 135 159 196 201 218 219 231 233 236 238

Concept drift 79 197 203 236 239

Counterfactual 148

Data contract 56 59 60 72 89 186 196

Data custodian 33 122 196 206 209 229

Data domain 270

Data governance 30

Data leakage 204

Data lineage 270

Data minimisation 41 83 123 149 152 189

Data owner 140 219 229

Data quality 20 30 38 60 65 67 75 90 148 179 218 228 231 233 238

Data steward 33 95 147 148 180 196 202 204 229

Dataset snapshot 70 160 190 204

Edge computing 271

EPCIS 18 20 24 54 59 78 85 98 102 109 136 155 207 231

Event time 24 66 85 98 99 101 103 104 107 110 165 196 218 231

Feature 39 57 67 70 73 74 78 85 99 138 139 151 153 157 160 163 165 167 168 170 172 176 184 185 186 190 196 197 198 201 202 204 205 206 209 211 236

Human oversight 272

Idempotency 89 107

Industrial data spine 15 26 65 79 82 130 134 195 212 218 234

Ingestion time 24 66 98 104 108 196 231

ISA-95 19 31 50 85 133 136

KPI dictionary 21 22 52 53 54 89 90 120 121 138 141 143 147 157 180 190 219 229 238 240

Least privilege 23 31 33 89 90 117 212

Master data 24 51 68 75 98 101 102 120 121 133 134 137 141 233

Metadata 21 42 50 57 60 61 69 72 73 83 90 92 94 99 111 120 131 136 137 141 163 165 168 186 229 233

MLOps 21 35 45 73 74 77 79 91 93 124 138 149 163 183 186 204 228 234 238

Model registry 119 120 138 177 207 234

Model retirement 272

MQTT 26 50 51 56 57 62 88 89 95 106 115 117 127 140

OPC UA 51 56 57 62 88 95 100 115 127 136

Operational technology 23 82 92 114 130 152 181 208 209 234

Provenance 16 24 26 30 35 54 58 63 65 70 73 74 78 79 83 84 86 89 91 92 93 99 103 104 114 117 120 131 137 138 140 141 157 164 165 166 171 190 191 212 231 234 237 240

Quarantine 24 30 31 33 38 42 43 44 45 50 60 61 65 67 68 69 71 75 77 78 89 107 117 122 131 132 133 139 142 203 208 216 219 228 231 232 234 237 238

Recovery point objective 121

Recovery time objective 121

Rollback 27 33 41 44 45 134 138 149 176 184 185 186 191 210 234 240

Schema 23 30 33 39 42 44 45 57 60 61 66 77 78 79 84 85 86 87 89 91 92 107 110 111 117 119 121 134 138 139 141 143 165 172 178 186 195 196 202 203 210 217 221 224 234 238

Shadow deployment 272

Store-and-forward 23 24 54 82 89 92 120 121 130 196 231 238

System of record 18 139 140 230

Timeliness 24 66 69 196 231 238

Validity 24 33 66 67 196 201 231 238

A reproducible edition keeps its limits visible.

Production is the warp. Measurement is the weft.

Green Future Tech prepared this white paper from published sources opened during research, each recorded with its URL, publisher, date and a quoted sentence in the evidence register, and from GFT's Cameroon textile programme evidence register (R entries, transcribed from the programme's data annex and flagged where refreshed in 2026). Derived figures show their inputs; illustrative models are labelled with their basis. No figure is presented as audited, and nothing here is legal, tax or investment advice.

Source treatment preserves uncertainty.

Inherited records reproduce the supplied finding and attribution only where cited. Their source-text status is preserved. Original publication dates and exact source URLs are incomplete in that register. These records require independent validation before external reliance. New chapter-local sources carry a publication date, complete URL, publisher, exact supporting quotation and access date.

The exports share a content model.

The PDF and web edition use the same parsed chapter blocks and SVG exhibits. The build records input hashes, artifact hashes, layout measurements and elapsed time. Cross-references, source notes and the term index are generated from identifiers rather than manually typed page numbers.

IMPRINT

Issued by Green Future Tech Limited, Hong Kong SAR.
Engineering the sustainable production of tomorrow.

GFT Research. Edition 2026-10-30. Document wp14. Locally generated proof. No named author, office, certification or client endorsement is implied.

GREEN FUTURE TECH

LIMITED / HONG KONG SAR

The factory that makes the garment should make the evidence.

Engineering the sustainable production of tomorrow.