GREEN FUTURE TECH
LIMITED / HONG KONG SAR

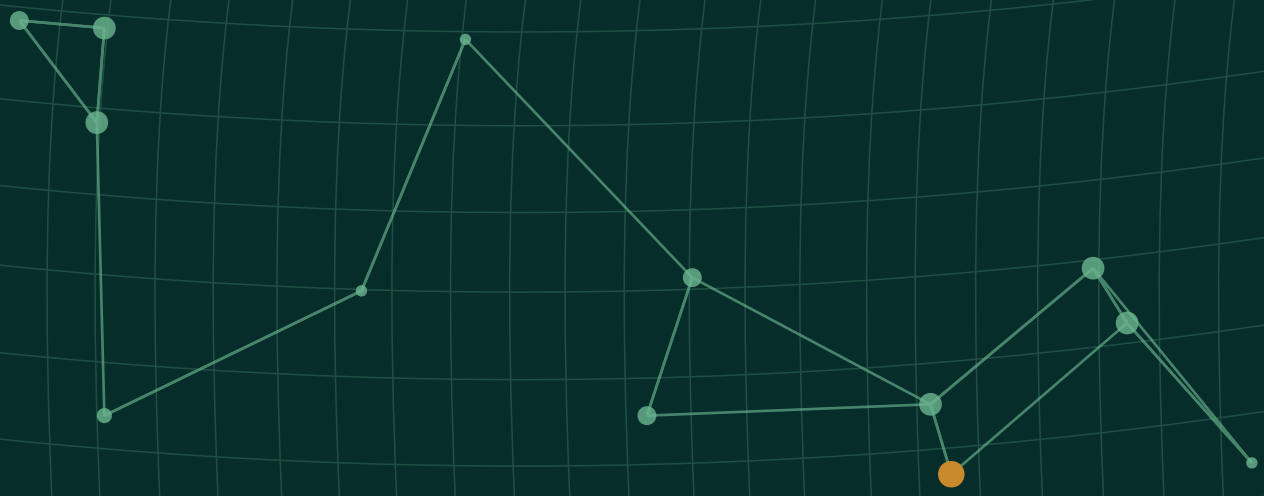
GFT RESEARCH / WARP AND SIGNAL

Edge to Dashboard

A reference architecture for the connected textile factory

wp18
RESEARCH EDITION

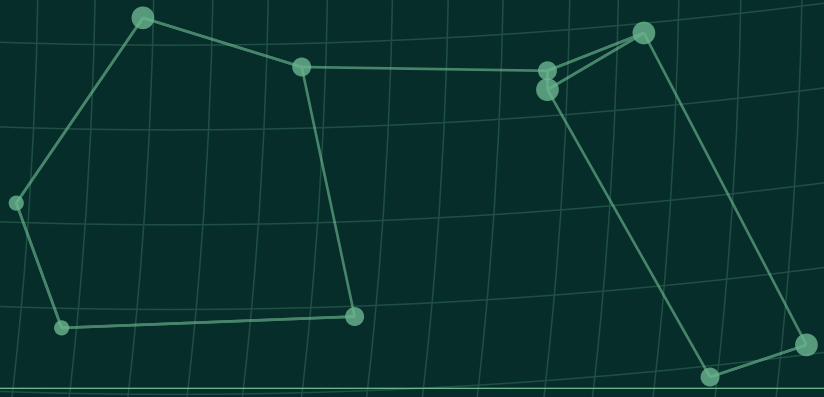
EDITION 2026-10-30
HONG KONG SAR



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	wp18 / Edge to Dashboard
EDITION DATE	2026-10-30
AUDIENCE	Industrial sponsors and factory directors planning cotton, textile or garment operations, Production, quality, maintenance, sustainability and information-technology leaders, System integrators, equipment suppliers and development-finance technical advisers, Cameroonian and regional industrial-policy teams evaluating practical digital infrastructure
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, R08, R09, R10, R11, R12, R13, R16, R29. Temporal eligibility is unverified; these are inherited source-text checks, not independent source validation.

The operating case runs from fibre to feedback.

01 / A dashboard cannot rescue an unmeasured factory 13

1.1 / The local-processing gap raises the cost of blind operation 14

1.4 / Power interruption makes local continuity a primary requirement 16

1.7 / Coverage figures constrain dependence but do not measure a factory link 19

1.10 / A cold start favours instrumentation before ambitious applications 21

1.12 / The unit of design is a decision tied to an operating event 23

1.14 / GFT's position is edge-first, safety-independent and evidence-led 25

02 / Requirements begin with the decision, not the device 29

2.1 / Production monitoring starts with a shared KPI definition 30

2.2 / Quality inspection requires an acquisition specification as well as a model 32

2.3 / Work-in-progress tracking connects physical identity to execution 34

2.4 / Resource dashboards begin at meters with declared boundaries 36

2.5 / Traceability requirements stop at the evidence actually available 38

2.6 / Every requirement card includes failure and recovery behaviour 40

03 / The architecture separates control, operations and enterprise responsibilities 43

3.1 / ISA-95 fixes the enterprise-to-control conversation 44

3.4 / Control remains close to the machine and below the analytics path 47

3.7 / Manufacturing operations owns the production context 50

3.10 / RAMI 4.0 adds asset and life-cycle views to the layer model 51

3.13 / Edge and cloud are placement choices within the same contracts 53

3.16 / Digital representations must state what is exchanged 56

04 / Identity and time turn signals into events 59

4.1 / Asset identity follows a stable factory hierarchy 60

4.2 / A topic namespace carries operational meaning 62

4.3 / NTP and PTP serve different correlation needs 63

4.4 / Source time and receipt time answer different questions 65

4.5 / Visibility events need business context beyond a sensor value 68

4.6 / Provenance travels with the value through every transformation 69

05 / Brownfield capture must not disturb machine control 72

5.1 / The machine audit precedes the gateway bill of materials 73

5.4 / Modbus acquisition is polling with explicit exception handling 75

5.7 / IO-Link separates cyclic process data from maintenance dialogue 77

5.10 / Deterministic networks are not interchangeable with ordinary Ethernet 78

5.13 / Added sensors need mechanical, radio and process validation 81

5.16 / Safety and security boundaries constrain every retrofit 83

06 / Edge inference keeps quality decisions on the line 89

6.1 / The camera cell is an engineered measurement station 90

6.2 / Inference stays local when the decision cannot wait 92

6.3 / Only useful results need to cross the constrained link 94

6.4 / Local persistence makes interruption recoverable 95

6.5 / Model outputs remain reviewable production evidence 98

6.6 / Manual inspection and safe production remain the fallback 100

07 / Networks should degrade in layers, not fail all at once 103

7.1 / Zones and conduits make permitted communication explicit 104

7.4 / The failure order begins with external convenience 107

7.7 / Store-and-forward needs explicit capacity and ordering rules 110

7.10 / Power and link telemetry are part of production context 113

7.13 / Remote access is a controlled session, not a permanent exception 115

7.16 / Recovery is proven by restoration and exercise 117

08 / Interoperability is a semantic contract before it is a connection 122

8.1 / OPC UA separates the information model from transport 123

8.4 / MQTT moves messages but does not define factory meaning 125

8.7 / Sparkplug adds state and session conventions 127

8.10 / AAS carries durable asset descriptions 129

8.13 / EPCIS carries visibility events rather than waveforms 130

8.16 / The canonical model maps standards without erasing their roles 132

09 / The historian must preserve context, quality and provenance 140

9.1 / Raw observations remain separate from interpreted events 141

9.5 / Quality state is data, not an annotation for later 144

9.9 / The edge buffer and central store have different duties 146

9.13 / KPI results are reproducible calculations 148

9.17 / Retention follows evidence need and storage cost 151

9.21 / Recovery includes reconciliation, not just database restoration 153

10 / A dashboard earns its place by changing a decision 157

10.1 / Every tile resolves to a governed KPI definition 158

10.4 / The line view favours immediate, controllable exceptions 161

10.8 / The plant view connects families without hiding their sources 163

10.11 / Freshness and coverage sit beside every important value 164

10.14 / Alerts end in owned action and verification 167

10.18 / Degraded modes remain legible and operable 170

11 / Small factory teams need reversible application deployment 173

11.1 / A pilot must test an operational hypothesis 174

11.2 / Edge modules isolate application change from machine change 176

11.3 / Version identity joins data, configuration and model 178

11.4 / Monitoring watches acquisition and decision quality together 180

11.5 / Local analytics reduce unnecessary transfer and administration 182

11.6 / Rollback is part of security and operational change control 184

12 / Resource and product data need a defensible metering boundary 186

12.1 / The metre hierarchy follows the energy decision 187

12.4 / Electricity source and backup state belong in the context 190

12.7 / Water volume is not automatically a water footprint 192

12.10 / Product carbon inputs begin with a product and reference flow 194

12.13 / Allocation uses the same production denominator as operations 196

12.16 / Product evidence remains linked to events and provenance 199

13 / Security, safety and recovery are acceptance tests 204

13.1 / The risk assessment defines the system under consideration 205

13.4 / Cybersecurity outcomes include governance and ownership 206

13.7 / OT controls must pass production and timing tests 208

13.10 / Remote maintenance fails closed and leaves evidence 210

13.13 / Protective safety never depends on the dashboard path 211

13.16 / Backups count only when isolated restoration succeeds 214

13.19 / Exercises turn contingency documents into operating capability 216

14 / Scale only after the architecture proves its decisions 222

14.1 / The target is one operating architecture, not a collection of pilots 223

14.2 / Phase 0 establishes the baseline and the contracts 225

14.3 / Phase 1 proves one line in shadow and then in operation 227

14.4 / Phase 2 replicates contracts before adding advanced analytics 230

14.5 / Phase 3 optimises a system that can still be restored 231

14.6 / GFT would measure decision performance and evidence health together 232

14.7 / Procurement gates keep ownership and exit rights with the factory 234

Executive summary 11

Evidence register 236

Glossary 261

Term index 265

Method and imprint 267

Every exhibit states a testable claim.

EXHIBIT 01 / bar
The inherited baseline shows how little cotton fibre is processed locally 15

EXHIBIT 02 / small-multiples
Cameroonian firms report outages more often than regional comparators 17

EXHIBIT 03 / keyfigures
The outage-incidence gap is 23.6 percentage points against Sub-Saharan Africa 18

EXHIBIT 04 / layered-architecture
The design position moves from continuity to evidence to applications 26

EXHIBIT 05 / flow
A useful requirement ends in verified action 32

EXHIBIT 06 / bar
The assumed inspection chain leaves a limited response margin 34

EXHIBIT 07 / table
Opened sources establish functions rather than local guarantees 36

EXHIBIT 08 / matrix
Urgent action and extensive evidence are different requirements 39

EXHIBIT 09 / table
Failure behaviour belongs in the use-case card 41

EXHIBIT 10 / layered-architecture
The reference stack assigns decisions before technologies 45

EXHIBIT 11 / matrix
Control, operations and enterprise functions have distinct owners 47

EXHIBIT 12 / flow
A command path and an evidence path are not interchangeable 49

EXHIBIT 13 / table
Reference models answer different architecture questions 52

EXHIBIT 14 / layered-architecture
Durable asset identity survives a change of messaging endpoint 61

EXHIBIT 15 / table
Published clock capability remains conditional on site evidence 64

EXHIBIT 16 / timeline
Delivery delay leaves the original occurrence time intact 66

EXHIBIT 17 / flow
An event carries identity, time and the reason its meaning is accepted 69

EXHIBIT 18 / table
The machine-capture sheet exposes unknowns before purchase 74

EXHIBIT 19 / table
Timing claims retain their technical conditions 79

EXHIBIT 20 / matrix
Protocol selection follows behaviour, not fashion 81

EXHIBIT 21 / layered-architecture
The machine boundary separates protective control from evidence export 84

EXHIBIT 22 / flow
The retrofit gate proves observation before integration 86

EXHIBIT 23 / layered-architecture
The cell separates observation from quality authority 91

EXHIBIT 24 / matrix
Immediate decisions favour local placement 93

EXHIBIT 25 / flow
Host loss changes delivery before inspection eligibility 96

EXHIBIT 26 / bar
Image retention dominates the illustrative outage buffer 97

EXHIBIT 27 / table
A result preserves observation, interpretation and disposition 99

EXHIBIT 28 / layered-architecture
Security zones preserve a safe core as external services fail 106

EXHIBIT 29 / flow
Degradation follows a predetermined order 108

EXHIBIT 30 / waterfall
Edge storage allocation separates compact event logs from diagnostic media 112

EXHIBIT 31 / keyfigures
Fourteen percent of Africa's population is outside mobile-broadband coverage 114

EXHIBIT 32 / table
Recovery objectives become testable factory contracts 119

EXHIBIT 33 / layered-architecture
Five standards occupy different layers of meaning 125

EXHIBIT 34 / flow
A machine state becomes a business event through explicit transformations 128

EXHIBIT 35 / matrix
Protocol claims are limited to their actual scope 131

EXHIBIT 36 / table
The interface contract is versioned as one unit 134

EXHIBIT 37 / table
The interface contract is versioned as one unit, continued 134

EXHIBIT 38 / waterfall
Semantic filtering reconciles raw edge messages into verified good-garment output 137

EXHIBIT 39 / layered-architecture
Four stores preserve different kinds of evidence 143

EXHIBIT 40 / table
A quality flag tells the reader whether to act 145

EXHIBIT 41 / flow
Late data is reconciled by identity and source time 147

EXHIBIT 42 / waterfall
Replayed edge records and audited corrections reconcile shift output 151

EXHIBIT 43 / matrix
Retention is assigned by evidence class rather than one duration 152

EXHIBIT 44 / table
The KPI dictionary is the dashboard's source of truth 160

EXHIBIT 45 / small-multiples
Role-specific views share definitions but not information density 162

EXHIBIT 46 / matrix
Dashboard content follows decision horizon and control 163

EXHIBIT 47 / flow
An alert closes only after action and verification 169

EXHIBIT 48 / flow
Evidence gates precede wider application dependence 175

EXHIBIT 49 / matrix
Custom behaviour must justify its maintenance burden 177

EXHIBIT 50 / layered-architecture
Release authority stays distinct from analytical lineage 179

EXHIBIT 51 / table
Health signals require an owner and a recovery action 181

EXHIBIT 52 / bar
Summaries carry less payload while preserving less detail 183

EXHIBIT 53 / layered-architecture
The metering hierarchy connects supply, process and product evidence 189

EXHIBIT 54 / keyfigures
Hydropower dominates Central Africa's regional installed capacity 191

EXHIBIT 55 / matrix
Each sustainability claim has a different evidence boundary 196

EXHIBIT 56 / table
A per-garment result retains its calculation trail 201

EXHIBIT 57 / table
A per-garment result retains its calculation trail, continued 201

EXHIBIT 58 / matrix
Acceptance maps every control to objective evidence 208

EXHIBIT 59 / layered-architecture
Protective safety remains outside the ordinary data path 213

EXHIBIT 60 / flow
Recovery ends with evidence reconciliation and corrective action 217

EXHIBIT 61 / table
The acceptance pack proves behaviour under failure 219

EXHIBIT 62 / layered-architecture
Operating responsibility survives application replacement 224

EXHIBIT 63 / flow
Expansion follows evidence rather than elapsed time 226

EXHIBIT 64 / bar
An aggregate can conceal a weak critical source 228

EXHIBIT 65 / table
A failed critical condition retains the existing scope 229

EXHIBIT 66 / table
Evidence and action qualify every operating result 233

Build the factory and its evidence as one system.

The connected factory is not a cloud subscription with cameras attached. It is a chain of evidence that begins at a machine, bundle, metre or camera and ends in a decision whose owner, timing and consequence are explicit. This paper defines that chain for a textile factory operating with mixed-generation equipment, intermittent power and variable connectivity. Cameroon 2024 Enterprise Survey data show why local continuity matters: 93.3 percent of surveyed firms experienced electrical outages [S32]. ITU's 2024 regional figures show a separate coverage constraint, with 14 percent of Africa's population outside mobile-broadband coverage, although this is not a measure of factory service quality [S33]. GFT therefore argues for an edge-first architecture that keeps machine control and protective safety independent, buffers data locally, and synchronises when links recover. ISA-95 and IEC 62264 separate control, operations and enterprise responsibilities [S01] [S02]. OPC UA, MQTT, Sparkplug, EPCIS and the Asset Administration Shell perform different, non-interchangeable roles across transport, state, events and asset meaning [S04] [S05] [S40] [S15] [S07]. The paper specifies identity and time, brownfield capture, local inference, historians, decision-led dashboards, reversible application deployment, resource-metering boundaries, security acceptance and a four-phase roadmap. Its test is not the number of connected devices. It is whether a supervisor, plant manager or sponsor can trace every displayed measure to a timely, contextual event and an accountable action.

01 Local operation is a design requirement because surveyed Cameroonian firms report frequent outages and regional mobile coverage remains uneven, not because cloud services have no value. [S32] [S33]

02 ISA-95 and IEC 62264 provide the responsibility boundary that prevents enterprise dashboards and integration services from becoming informal machine-control paths. [S01] [S02] [S17]

03 No single protocol solves interoperability: OPC UA models industrial information, MQTT transports messages, Sparkplug defines session state, EPCIS shares visibility events and AAS structures asset information. [S04] [S05] [S40] [S15] [S07]

04 A production event is not audit-ready until it carries asset identity, source time, quality state and provenance appropriate to its data class, with NTP or PTP selected according to the required correlation precision. [S12] [S13] [S15] [S31]

05 Brownfield instrumentation should begin with read-only, non-disruptive capture and explicit protocol behaviour rather than assume that every machine can publish useful data. [S09] [S10] [S17] [S35]

06 Edge inference is operationally justified for inspection and energy analytics, but lighting, calibration, local persistence and manual fallback determine whether the output is dependable. [S16] [S26] [S29] [S22]

07 Dashboard tiles should implement a governed KPI dictionary and a management action, not recalculate isolated metrics independently of their source elements and time behaviour. [S08] [S30]

08 Energy, water and carbon displays become defensible only when their organisational, process or product boundary and allocation logic are visible. [S36] [S37] [S38]

09 Security, protective safety and recovery require separate acceptance evidence, including zones and conduits, independent protective action, isolated backups and exercised recovery. [S19] [S23] [S21] [S22]

10 A bounded production pilot should prove data quality and decision closure before the architecture is replicated across lines or sites. [S30] [S08] [S18]

GREEN FUTURE TECH

LIMITED / HONG KONG SAR

01

PART I: DECISIONS BEFORE DEVICES / OPERATING PROBLEM

A dashboard cannot rescue an unmeasured factory

GFT starts with the production decisions and infrastructure constraints that a connected architecture must serve, because visualisation alone cannot create trustworthy operating evidence.

CHAPTER 01
DECISION TEST

GFT starts with the production decisions and infrastructure constraints that a connected architecture must serve, because visualisation alone cannot create trustworthy operating evidence.

A connected textile factory creates operational value only when a production event retains its physical identity, source timestamp, quality state and process context from the machine edge to a governed decision surface. [S01] [S08] Buying a cloud dashboard before engineering shop-floor capture reverses that dependency: it places a visual layer over missing denominators, unsynchronised clocks and fragile external links. [S08] [S31] In Cameroon, where the inherited sector baseline records roughly 95 percent of cotton fibre exported raw and only about 5 percent processed locally, new spinning, weaving, dyeing and garment capacity must prove operational control over throughput, first-pass yield and utility consumption from the first shift. [R06] At the same time, the Cameroon 2024 Enterprise Survey reports that 93.3 percent of surveyed formal private firms experienced electrical outages, with 10.4 outages in a typical month. [S32] Under these operating conditions, Green Future Tech Limited (GFT), Hong Kong SAR, specifies an edge-first, safety-independent reference architecture in which protective machine functions never depend on the telemetry path, local edge buffers preserve production records across power and link interruptions, and every dashboard tile traces to an accountable management action. [S17] [S22] [S23]

1.1

The local-processing gap raises the cost of blind operation

Industrial policy in raw-material economies often frames textile industrialisation as a capacity-installation problem, yet factory survival depends on conversion discipline once equipment is on the floor. [R06] [R12] According to the inherited sector register from the Cameroon Trade Hub, approximately 95 percent of Cameroon's cotton fibre is exported raw while about 5 percent is processed locally, an observation whose original publication date is unrecorded in the inherited annex and which represents a sector baseline rather than a GFT operating result. [R06] That imbalance exists alongside substantial upstream output: inherited trade-press records indicate that Cameroon's seed-cotton production reached 330,000 tonnes in 2022, 314,455 tonnes in 2023 and about 340,000 tonnes in 2024, supported by more than 220,000 cotton producers, nine ginning plants and two oil mills operated by SODECOTON in northern Cameroon. [R01] [R03] Converting a larger share of that lint into yarn, fabric and garments inside the country requires factories that measure and control what happens between bale opening and final packing. [R06] [S08]

The Cameroonian government's national development strategy, SND30, restates a policy target of processing 50 percent of cotton locally by 2030. [R07] Policy targets and announced investment programmes establish ambition, but they do not constitute operating evidence that any connected factory or GFT deployment has achieved that conversion ratio. [R07] Third-party announcements illustrate the capital scale under discussion: in November 2025, trade press reported the Camtext SA integrated cotton-to-textile project at an estimated US$300 million (180 billion XAF), targeting 12,000 tonnes a year of processed fibre and optimal capacity around 2033, alongside separate proposals to merge SODECOTON and CICAM into an industrial pole by 2030. [R08] [R09] [R10] These third-party plans are external sector context, not GFT engagements, and their long payback horizons underscore why domestic mills cannot afford unmeasured conversion losses. [R09] [R10]

1.2 HISTORICAL UNDERUTILISATION SHOWS THE PENALTY OF WEAK OPERATIONAL VISIBILITY

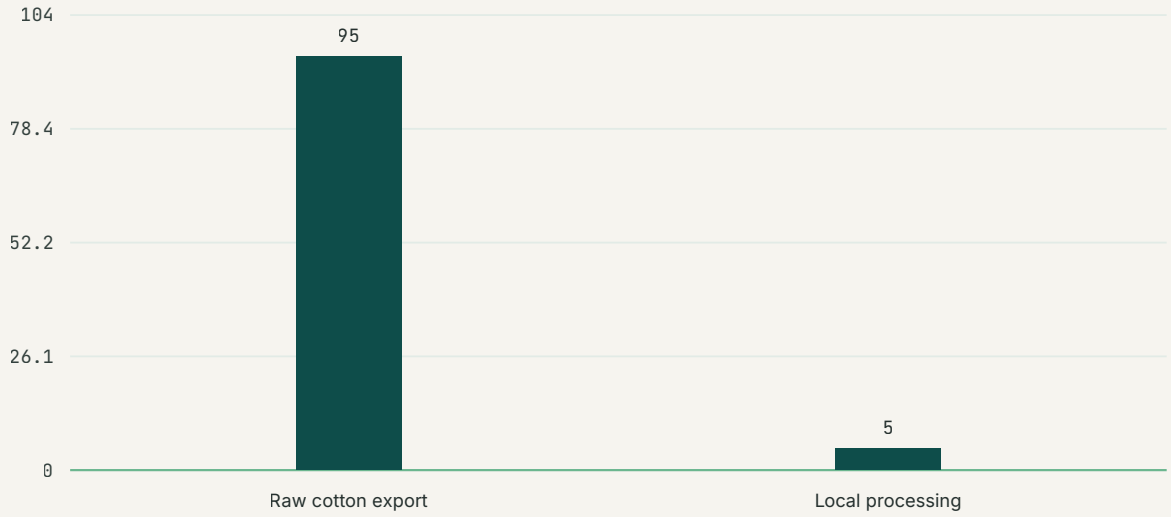
Domestic processing history demonstrates how quickly unaddressed equipment downtime, quality drift and cost leakage erode competitiveness. [R11] [R12] Inherited press reporting states that the national textile producer CICAM saw its domestic market share fall from about 85 percent three decades ago to about 5 percent, while a July 2023 report recorded capacity utilisation at 33.7 percent at Garoua, 48.9 percent at CICAM 1 and 10.83 percent at CICAM 2, with underutilisation costing about 2 billion XAF in 2020 and restructuring estimates ranging from 30.7 billion to 70.2 billion XAF. [R11] [R12] [R13] A plant operating below half of its rated capacity cannot recover margin through summary charts that arrive weeks after a production run has closed. [R12] [S08]

When a spinning frame, loom or sewing line runs without structured event capture, three losses merge into an unexplained shortfall at the end of the month. [S08] First, availability losses from breakdown, warp breakage, changeover or power transfer go unseparated from scheduled stops, hiding whether maintenance or scheduling owns the remedy. [S08] Second, speed losses and micro-stoppages reduce output per operating hour without triggering a downtime ticket. [S08] Third, quality defects discovered only at final packing inspection force rework or second-grade discounting after full labour, dye, steam and electricity inputs have been consumed. [R29] [S26]

Exhibit 1, p. 15 contrasts the inherited baseline shares of raw cotton export and domestic processing in Cameroon. [R06]

EXHIBIT 01 / PERCENT OF COTTON FIBRE

The inherited baseline shows how little cotton fibre is processed locally



SOURCED [R06] Approximate inherited sector baseline from Cameroon Trade Hub (R06); original publication date not recorded. This is a sector observation, not a GFT operating result.

The distribution in **Exhibit 1**, p. 15 establishes the starting point for GFT's Cameroon textile programme, prepared on 31 January 2024 with source-refresh notes retained in 2026, covering a feasibility study, performance-management system design and

digital strategy. [R06] However, **Exhibit 1**, p. 15 cannot establish the financial viability of any individual mill, nor does it specify the instrumentation required to move from a 5 percent baseline toward the 50 percent policy target. [R06] [R07]

1.3 VISUALISATION WITHOUT EVENT INTEGRITY MULTIPLIES MANAGEMENT ERROR

A recurring failure in industrial digitalisation is procuring an enterprise business-intelligence portal before defining what constitutes a valid shop-floor observation. [S01] [S08] When managers commission visual tiles first, software integrators are forced to scrape whatever numbers are easiest to reach: an uncalibrated machine stroke counter, an operator's end-of-shift tally sheet, or a warehouse dispatch note entered two days late. [S08] [S09] If a dashboard divides today's warehouse receipts by today's scheduled shift hours, it mixes work-in-process inventory flushed from previous shifts with current line performance, producing a false efficiency spike followed by an unexplained collapse the next day. [S08]

That distortion compounds when managers use unverified visual trends to allocate capital or penalise line supervisors. [S08] [S30] A spinning or garment line that appears to underperform because its manual tally sheets were entered after the midnight cutoff may have its preventive-maintenance window cancelled to chase output, accelerating mechanical wear and increasing yarn breaks on the following shift. [S08] [S30] Conversely, a dyeing range that logs batch completions without linking them to sub-metered steam and water volumes can report high apparent machine utilisation while masking costly re-dyeing cycles caused by shade drift. [S08] [S27] [S29]

In GFT's engineering practice, a dashboard is the last component of an evidence chain, never its foundation. [S01] [S08] Production is the warp that carries physical material forward; measurement is the weft that binds each thread to a verifiable time, asset identifier and quality state. [S01] [S40] If the weft is missing at the machine edge, no downstream chart library can reconstruct whether an output drop came from yarn hairiness, a feeder trip, an unrecorded style change or a silent gateway reboot. [S08] [S40]

1.4

Power interruption makes local continuity a primary requirement

Cloud-centric industrial architectures assume continuous mains electricity, uninterrupted local-area networks and always-on wide-area backhaul. [S14] [S16] In Cameroon, empirical enterprise data invalidate that assumption as a design baseline. [S32]

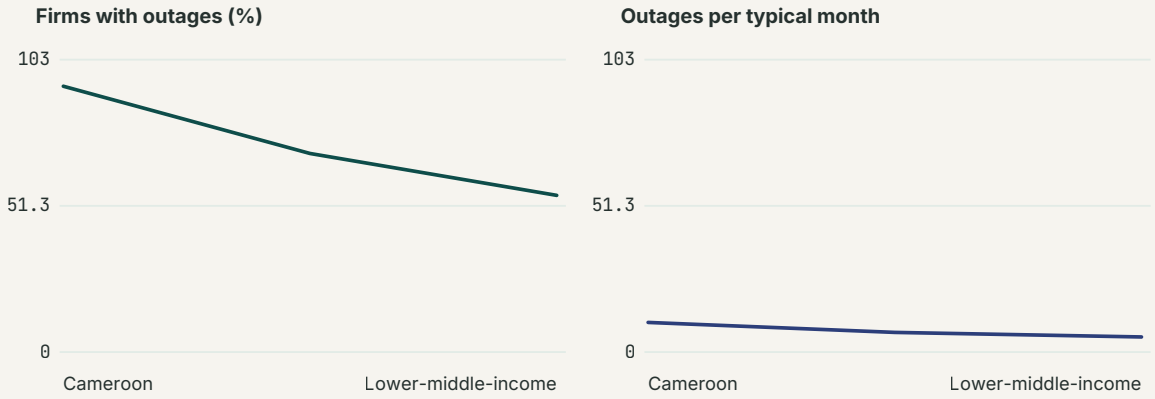
In the World Bank's Cameroon 2024 Enterprise Survey, conducted between June 2024 and February 2025 across formal non-agricultural private firms, 93.3 percent of all surveyed firms reported experiencing electrical outages, compared with 69.7 percent across Sub-Saharan Africa and 55.0 percent across lower-middle-income economies. [S32] Surveyed Cameroonian firms also reported 10.4 electrical outages in a typical month, against 6.9 in Sub-Saharan Africa and 5.3 in lower-middle-income comparator economies. [S32] Because the Enterprise Survey samples formal private businesses across sectors rather than textile mills alone, these figures are an all-firm

baseline rather than a site-specific guarantee. [S32] Even with that sample boundary preserved, grid failure roughly ten times a month makes power interruption a normal operating condition. [S32]

Exhibit 2, p. 17 presents both dimensions of the World Bank outage evidence across Cameroon, Sub-Saharan Africa and lower-middle-income comparator groups. [S32]

EXHIBIT 02 / PERCENT OF FIRMS (PANEL A) AND MONTHLY OUTAGES (PANEL B)

Cameroonian firms report outages more often than regional comparators



SOURCED [S32] World Bank Enterprise Surveys Cameroon 2024 Country Profile, Figure 8, fieldwork June 2024 to February 2025 (S32). Covers surveyed formal non-agricultural private firms across all sectors, not a textile-only sample. Panel A reports percent of firms experiencing electrical outages; Panel B reports number of power outages in a typical month.

As **Exhibit 2**, p. 17 shows, Cameroonian enterprises face both wider outage exposure and higher monthly interruption frequency than either comparator group. [S32] While `small-multiples` uses a shared numeric scale, the two panels carry separate units: Panel A reports percentage of firms experiencing outages and Panel B reports typical-month outage counts. [S32] Moreover, **Exhibit 2**, p. 17 cannot establish feeder sag duration in Douala, Dibamba or Garoua, nor does it measure generator transfer speed at a specific mill. [S32]

To quantify the structural reliability gap that a systems architect must design around, **Exhibit 3**, p. 18 derives the incidence and frequency differentials from the World Bank dataset. [S32]

EXHIBIT 03 / PERCENTAGE POINTS AND MONTHLY EVENTS

The outage-incidence gap is 23.6 percentage points against Sub-Saharan Africa

23.6 pp	38.3 pp	3.5/mo
Gap vs Sub-Saharan Africa	**Gap vs lower-middle-income**	**Monthly outage excess**
Derived: 93.3% Cameroon minus 69.7% regional benchmark (all surveyed firms)	Derived: 93.3% Cameroon minus 55.0% income-group benchmark (all surveyed firms)	Derived: 10.4 typical-month outages in Cameroon minus 6.9 in Sub-Saharan Africa

DERIVED [S32] GFT arithmetic from World Bank Enterprise Surveys Cameroon 2024 Country Profile (S32). Outage-incidence gap against Sub-Saharan Africa: 93.3 percent minus 69.7 percent equals 23.6 percentage points. Outage-incidence gap against lower-middle-income economies: 93.3 percent minus 55.0 percent equals 38.3 percentage points. Monthly frequency excess: 10.4 minus 6.9 equals 3.5 outages per typical month. Compares all surveyed firms, not textile factories.

The derived differences in **Exhibit 3**, p. 18, namely 23.6 percentage points above Sub-Saharan Africa, 38.3 percentage points above lower-middle-income economies, and 3.5 additional outages per typical month, confirm that power discontinuity is a primary architectural constraint, though firm-level responses do not reveal individual factory UPS coverage. [S32]

1.5 BROWNOUT AND TRANSFER TRANSIENTS CORRUPT NAIVE TELEMETRY PIPELINES

When mains power drops at a mill, high-inertia machines coast or wait for diesel synchronisation while low-voltage controllers, sensors, gateways and switches either drop immediately or ride through on battery backup. [S17] [S29] If monitoring relies on stateless remote polling or unbuffered gateways, four forms of evidence corruption occur: [S09] [S40]

1. **Counter reset and register volatility:** Volatile holding registers on retrofitted controllers reset to zero on voltage dips, causing cumulative counter differences to record negative spikes or drop an hour of output. [S09] [S17]
2. **Clock drift and epoch reversion:** Edge nodes lacking battery-backed real-time clocks or local Network Time Protocol (NTP) servers reboot to a default epoch timestamp before synchronisation recovers, stamping post-outage events into the wrong shift. [S13] [S31]
3. **Silent gap misinterpretation:** Without explicit session-state announcements and UPS telemetry, a central dashboard cannot distinguish zero defects during a 40-minute window from an unpowered inspection station. [S32] [S40]
4. **Unordered replay storms:** When power returns, edge buffers flushing stored readings without separate source and receipt timestamps can overwrite current machine state with historical values. [S29] [S40]

1.6 CONTINUITY REQUIRES EXPLICIT RIDE-THROUGH, LOCAL BUFFERING AND EXERCISED RECOVERY

NIST Special Publication 800-82 Rev. 3 requires operational technology (OT) security and architecture to address performance, reliability and safety alongside data protection. [S17] Edge gateways, local brokers, time servers and access switches must sit on monitored uninterruptible power supplies sized to outlast generator transfer plus controlled database flush. [S17] [S29] Industrial IoT energy-monitoring evidence confirms the value of storing meter data in a local edge database before deriving analytical parameters. [S29] Similarly, the Eclipse Sparkplug specification defines MQTT birth and death certificates so subscribers mark offline metrics stale and flush historical data once the host application returns online. [S05] [S40] Finally, NIST Special Publication 800-34 Rev. 1 distinguishes testing recovery capabilities, training personnel and exercising plans to identify gaps. [S22] No monitoring stack should be accepted until physical power-pull tests prove that zero committed edge events are lost or reordered. [S21] [S22]

1.7

Coverage figures constrain dependence but do not measure a factory link

Regional telecommunications indicators similarly caution against making shop-floor execution depend on uninterrupted external connectivity. [S33] In its 2024 *Facts and Figures* report, the International Telecommunication Union (ITU) reports that Africa has the world's largest mobile-broadband coverage gap, with 14 percent of the population lacking access to a mobile-broadband network. [S33]

A continental population gap of 14 percent is a demographic indicator, not a measurement of fibre availability in Douala or Dibamba or a packet-loss log in Garoua. [S33] What the ITU figure justifies is a design constraint: external wide-area connectivity cannot be treated as a zero-outage utility on which machine control, interlocks or primary event capture may depend. [S17] [S33]

1.8 SITE LINK QUALIFICATION REQUIRES MEASURED TRANSPORT METRICS RATHER THAN COVERAGE MAPS

Before deciding which data streams cross the site boundary, engineers must execute an active link survey across working shifts and weather conditions, measuring five parameters for each candidate backhaul path: [S14] [S16]

- **Round-trip latency distribution:** Median, 95th-percentile and 99th-percentile round-trip time, because tail latency governs whether synchronous acknowledgements stall. [S14] [S16]
- **Packet loss and burst error duration:** Random frame loss and multi-second blackouts during handovers or tower battery depletion during grid outages. [S32] [S33]
- **Sustained uplink throughput:** Available outbound bandwidth, which in industrial telemetry and vision inspection carries far heavier load than the downlink. [S05] [S26]
- **Link availability and mean time to repair:** Observed outage duration and whether redundant providers share physical ducts, bridges or exchanges. [S17] [S22]

- **Tariff structure and per-gigabyte cost:** Volume metering or fair-use caps that constrain continuous streaming of high-resolution camera frames or waveforms. [S05] [S26]

1.9 SELECTIVE EDGE PUBLICATION PROTECTS CONSTRAINED LINKS AND LOCAL AUTONOMY

Standards already separate local acquisition from selective external publication. [S14] [S16] The W3C *Web of Things Architecture 1.1* models an industrial edge device that collects selected controller data for a cloud dashboard or preventive-maintenance analysis while accommodating local networks behind NAT or firewalls. [S14] The W3C *Web of Things Thing Description 1.1* standard specifies a JSON-LD metadata and interface model hosted on the device or externally for resource-restricted or retrofitted legacy equipment. [S0102] ETSI Multi-access Edge Computing (ETSI GS MEC 003 V3.2.1) runs software applications on virtualised edge infrastructure, while OASIS MQTT Version 5.0 provides a lightweight, payload-agnostic publish-subscribe transport for constrained bandwidth. [S05] [S16] [S0101] Together, these standards let a factory run high-rate polling, vision inference, buffering and shift displays locally while publishing compact summaries when links permit. [S05] [S14] [S16]

WORKED EXAMPLE: SIZING AN EDGE STORE-AND-FORWARD BUFFER FOR A 72-HOUR BACKHAUL OUTAGE

Consider an **illustrative** sizing calculation for a single pilot line during a 72-hour wide-area link outage. All quantities below are **illustrative** engineering assumptions used solely to demonstrate the sizing method, not measurements from a GFT factory.

Assume an illustrative pilot line generates: 1. **Machine and utility telemetry:** 40 polled points sampled every 10 seconds, producing 4 points per second, or 14,400 records per hour (illustrative). 2. **Work-in-process and quality events:** 300 bundle RFID reads and defect logs per hour across 2 shifts (16 hours per day), or 4,800 visibility events per day (illustrative). 3. **Edge computer-vision inspection:** 1 fabric camera cell running 12 net hours per day, generating 1 summary inference record every 2 seconds (21,600 records per day) and flagging an illustrative 2 percent of frames (432 diagnostic image crops per day).

Applying an illustrative payload budget per record including timestamps, quality flags and local index overhead:

- Compact telemetry at an illustrative 250 bytes per record across 345,600 records per day requires 86,400,000 bytes (86.4 MB per day, illustrative).
- WIP and quality events at an illustrative 600 bytes across 4,800 events per day require 2,880,000 bytes (2.88 MB per day, illustrative).
- Vision inference metadata at an illustrative 500 bytes across 21,600 records per day requires 10,800,000 bytes (10.8 MB per day, illustrative).
- Structured events total 86.4 plus 2.88 plus 10.8, which equals 100.08 MB per day, or **300.24 MB over 72 hours** (illustrative). A 2.0 safety multiplier for write-ahead logging and index overhead yields **600.48 MB** of local structured storage (illustrative).
- Diagnostic image crops at an illustrative 250 KB (0.25 MB) across 432 crops per day require 108 MB per day, or **324 MB over 72 hours** (illustrative). Storing raw 2 MB frames at 0.5 Hz without edge filtering would instead generate 43.2 GB per day, or **129.6 GB over 72 hours** (illustrative).

Engineering consequence: Edge filtering keeps the illustrative 72-hour buffer under 1 GB (600.48 MB structured plus 324 MB crops equals 924.48 MB, illustrative). On a 512 kbps (64 KB/s) recovery link allocating 25 percent of bandwidth (16 KB/s, or 57.6 MB per hour) to historical replay, flushing 300.24 MB of structured history takes **5.21 hours** (illustrative), whereas raw video would require more than 93 days. [S05] [S16] [S29] [S40]

1.10

A cold start favours instrumentation before ambitious applications

When sponsors modernise or build a mill, vendors often pitch artificial-intelligence suites and enterprise digital twins first. [S30] [S35] In West and Central African manufacturing, however, the empirical starting point is predominantly analogue or mixed-generation equipment. [S35]

In UNIDO's *Industrial Development Report 2020*, survey evidence from Ghana showed that almost 70 percent of surveyed manufacturing firms fell into the analogue category, relying on manual machinery and paper or basic records. [S35] GFT retains the exact boundary of this finding: it is a 2019 Ghanaian manufacturing benchmark, not a Cameroon textile estimate. [S35] As a regional readiness comparator, however, it explains why a reference architecture must accommodate a brownfield cold start rather than assume that every carding machine, dyeing vat, stenter frame and sewing machine exposes a native OPC UA server. [S04] [S09] [S35]

1.11 JUMPING STRAIGHT TO COMPLEX ANALYTICS ON UNINSTRUMENTED LINES FAILS

Deploying complex analytics before clean event capture fails through three mechanisms: [S08] [S30]

1. **Numerator-denominator mismatch:** If sewing output is tallied by a packing clerk at day's end while operating hours come from the shift roster, any model trained on line efficiency fits clerical batching delays rather than physical sewing rhythm. [S08]
2. **Uncontrolled environmental noise:** Textile defect-detection research shows that lighting variations and background noise significantly degrade machine-learning accuracy, while wireless sensor reviews show that clock synchronisation and spectrum sharing are essential to data integrity. [S26] [S31] Uncalibrated sensing produces false alarms that operators ignore. [S26] [S31]
3. **Organisational overload:** Research on Industry 4.0 in textile SMEs recommends small pilot projects to test and validate selected technologies before wider deployment. [S30] A small factory team cannot debug PLC registers, train bundle scanning, validate KPIs and run a bespoke cloud cluster at once. [S08] [S09] [S30]

GFT COLD-START SEQUENCE: MEASURE, STABILISE, THEN MODEL

GFT's programme enforces a four-step dependency order before commissioning advanced analytics: 1. **Define the decision and denominator (Contract stage):** Specify the management decision, owner, cadence and ISO 22400 KPI formula, elements, time behaviour and unit before selecting hardware. [S08] 2. **Establish non-disruptive capture and time (Instrumentation stage):** Audit the physical machine interface, verify electrical isolation and protective safety independence, and synchronise edge clocks via NTP or PTP. [S09] [S12] [S13] [S17] [S23] 3. **Prove coverage and reconciliation in shadow mode (Stabilisation stage):** Run the edge buffer, historian and dashboard in parallel with approved manual logs on one pilot line until coverage, freshness and KPI reconciliation tolerances hold across actual outages. [S08] [S22] [S30] [S40] 4. **Deploy bounded models under change control (Application stage):** Introduce local vision inference, forecasting or energy optimisation only after the event spine is stable, keeping every release reversible. [S16] [S24] [S26] [S30]

1.12

The unit of design is a decision tied to an operating event

In GFT's reference model, the unit of design is **an operational decision tied to a verifiable production event**. [S01] [S08] International standard ISO 22400-2:2014 specifies manufacturing KPIs through their formula, corresponding elements, time behaviour, unit or dimension, and operational characteristics. [S08] Every dashboard tile must therefore declare: [S08]

- **Source elements:** Which physical events populate the numerator and denominator, such as validated good pieces at end-of-line inspection versus planned operating time minus authorised stops. [S08]
- **Time behaviour and boundary:** Whether the metric is evaluated by exception, aggregated over a rolling 60-minute window, or closed at shift handover using source timestamps. [S08] [S13]
- **Unit and dimensionality:** Whether the metric is expressed in pieces per hour, percentage of inspected metres, kilowatt-hours per good kilogram of yarn, or litres of water per garment. [S08] [S29]

1.13 GFT'S MANAGEMENT CADENCE LINKS SHOP-FLOOR EVENTS TO THREE DECISION HORIZONS

GFT's Cameroon textile programme links six KPI families (throughput and OEE, quality, maintenance, people, unit cost, and sustainability) to three management cadences: [S08] [S29]

Management cadence	Primary decision owner	Core source events required	Architectural freshness and continuity requirement	Operational action triggered
Shift and hourly huddle	Line supervisor and quality lead	Good pieces or metres, defect code by operation, machine stop start/end and cause [S08] [S26]	Local edge processing; independent of WAN internet; stale state visible within seconds via Sparkplug [S16] [S40]	Rebalance bottleneck workstation, adjust sewing tension or illumination, dispatch technician [S26] [S30]
Weekly plant review	Plant manager, maintenance lead and utilities engineer	Reconciled shift OEE elements, repeat downtime pareto, submetered kWh and water volume per good unit [S08] [S29]	Local historian reconciliation across power outages; versioned KPI calculation records [S08] [S29] [S32]	Schedule preventive maintenance, revise line balancing, correct compressed-air or steam leaks [S08] [S29]
Monthly sponsor and executive review	Factory director and industrial sponsor	Order on-time-in-full delivery, scrap and rework cost, resource intensity baseline vs target, pilot gate metrics [R06] [S08] [S30]	Verified backup restoration logs, audit trail of manual corrections and closed action items [S21] [S22] [S24]	Approve capital gate for next line, adjust product mix, authorise supplier or process change [S08] [S30]

This separation aligns with ISA-95 (ANSI/ISA-95.00.01-2025) and IEC 62264-1:2013. [S01] [S02] ISA-95 Part 1 organises physical assets and defines the enterprise-control interface, while IEC 62264-1 locates manufacturing operations management (MOM) at Level 3 and specifies transactions within Level 3 and with Level 4 enterprise planning. [S01] [S02] Hourly line huddles execute inside Level 3 using Level 1 and Level 2 events without round-tripping through a Level 4 ERP server or overseas cloud, whereas monthly sponsor reviews consume reconciled Level 3 evidence without polling Level 1 PLC registers. [S01] [S02] [S16]

Preserving this boundary also protects data ownership during phased modernisation. [S01] [S08] In GFT's Cameroon textile programme, a factory may begin Phase 0 and Phase 1 with structured supervisor logs and validated spreadsheet templates before migrating to a dedicated historian and automated edge gateways. [S08] [S30] Because the KPI formula, source-event boundary, unit and review cadence are governed at the operations layer rather than hard-coded inside a proprietary visualization tool, replacing a manual entry sheet with a Modbus counter or an RFID read point changes the acquisition mechanism without breaking historical comparability. [S08] [S09] [S25] Supervisors, plant managers and industrial sponsors continue reading the same

definitions before and after automation, allowing the factory to prove whether a digital intervention genuinely improved throughput, right-first-time quality or utility intensity. [S08] [S29] [S30]

1.14

GFT's position is edge-first, safety-independent and evidence-led

Synthesising the local-processing imperative, Cameroonian outage evidence, connectivity constraints and automation standards defines GFT's architectural position: **edge-first in execution, safety-independent in control, and evidence-led in governance**. [R06] [S01] [S17] [S23] [S32]

1.15 PROTECTIVE SAFETY AND MACHINE CONTROL NEVER DEPEND ON THE ANALYTICS PATH

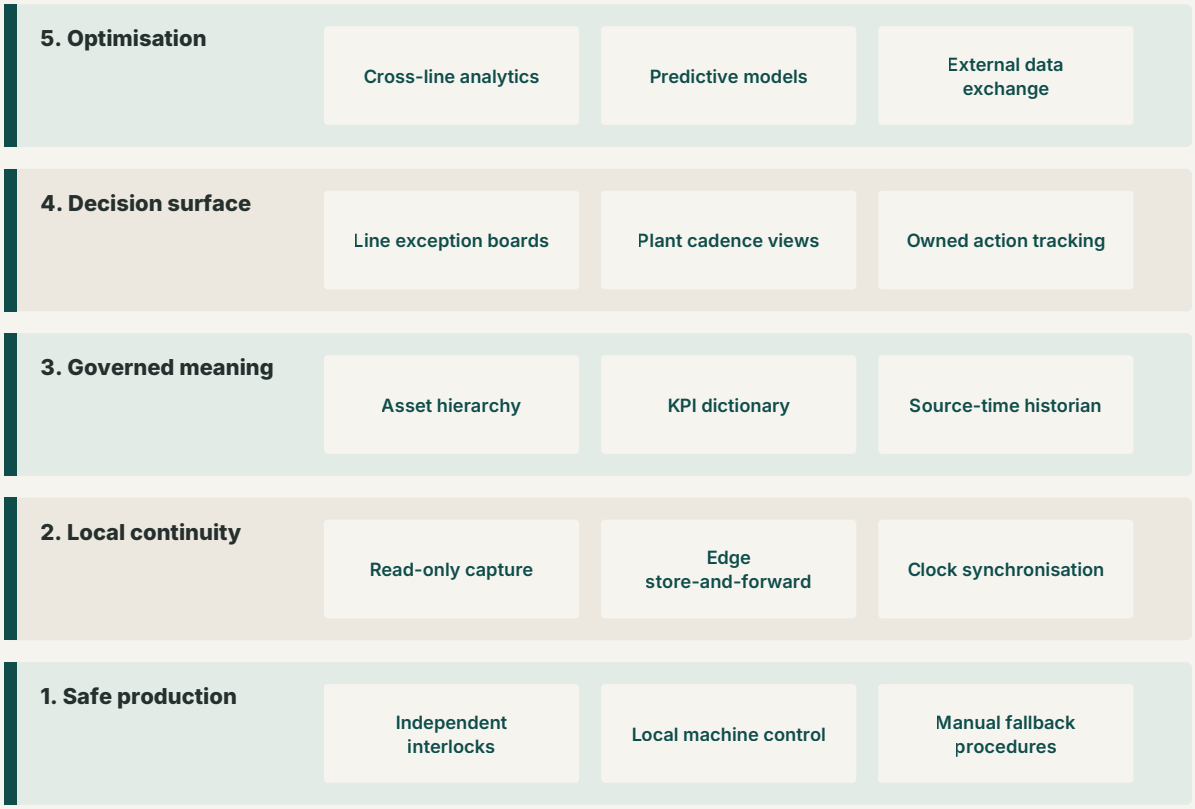
As the International Electrotechnical Commission (IEC) states, functional safety provides automatic protection in response to errors or hazardous conditions. [S23] Emergency stops, guard interlocks, dyeing over-temperature cut-outs, boiler relief valves and cutter light curtains must operate through certified hardware channels independent of gateways, brokers, inference nodes and dashboards. [S17] [S23] Combining NIST SP 800-82 Rev. 3 on OT performance, reliability and safety with IEC 62443-3-2:2020 on security zones and conduits yields three rules: [S17] [S19] [S23]

1. **No dashboard or cloud service holds an unmediated write path to a machine controller:** Dashboards display evidence and record human decisions rather than acting as informal remote HMIs. [S01] [S17] [S19]
2. **Brownfield capture defaults to read-only observation:** Modbus polling or cyclic IO-Link reads must be tested so gateway failure or cable disconnection causes zero variation in machine cycle timing. [S09] [S10] [S17]
3. **Safety zones remain isolated:** Safety-related functions occupy a dedicated zone and never depend on higher-level information networks. [S19] [S23]

Exhibit 4, p. 26 illustrates the five-layer architectural design position applied across this white paper. [S01] [S17] [S22] [S30]

EXHIBIT 04 /

The design position moves from continuity to evidence to applications



ILLUSTRATIVE [S01] [S17] [S22] [S30] Illustrative GFT reference design sequence informed by ISA-95 responsibility boundaries (S01), NIST OT performance and safety guidance (S17), NIST contingency validation and exercises (S22), and textile SME pilot discipline (S30). Represents a target design method, not an installed GFT factory.

In **Exhibit 4**, p. 26, each upper layer relies on the layer beneath it, while lower layers survive the failure of every layer above them. [S17] [S22] [S23] If Layer 5 or Layer 4 goes offline, Layer 3 and Layer 2 continue capturing timestamped evidence locally. [S29] [S40] If Layer 2 loses power beyond its UPS window, Layer 1 still protects operators and equipment while manual fallback procedures record output. [S22] [S23] Because **Exhibit 4**, p. 26 is an illustrative design sequence rather than a wiring diagram, it cannot replace a site-specific IEC 62443 risk assessment or machinery safety validation. [S19] [S23]

GFT NOTE: PRACTICAL PROCUREMENT RULE FOR INDUSTRIAL SPONSORS

Sponsors should reject any equipment or software proposal that locks machine telemetry inside a mandatory cloud subscription without local access. [S04] [S14] Require three contractual deliverables before milestone sign-off: (1) a documented local interface (OPC UA, Sparkplug B over MQTT, or read-only Modbus register tables) operable without internet access; (2) an edge store-and-forward buffer verified by a physical disconnection test; and (3) a governed KPI dictionary mapping every dashboard tile to its ISO 22400 formula, source elements, unit and operational owner. [S04] [S08] [S09] [S22] [S40]

SCOPE AND ADVISORY BOUNDARY

This paper specifies an engineering reference architecture for industrial data capture, semantic interoperability, local resilience and operational dashboards. It is not investment, legal or tax advice, and it does not certify the safety, cybersecurity compliance or financial return of any factory, machine or product. Inherited sector statistics (R01 to R18 and R23 to R32) retain their source-text status and unrecorded publication dates where noted; missing references R19 to R22 are excluded. [R06] [S17] Every retrofit, power scheme and safety boundary requires site-specific engineering validation. [S17] [S19] [S23]

1.16 HOW THE FOURTEEN CHAPTERS OF THIS PAPER BUILD THE ARCHITECTURE

The remaining chapters follow a production event from the machine edge to gated scale: [S01] [S08] [S30]

- **Part I (Chapters 1 to 3)** defines use-case requirement cards with latency and fallback rules (Chapter 2) and formalises the ISA-95, IEC 62264 and AAS reference stack (Chapter 3). [S01] [S02] [S07] [S08] [S15] [S25] [S26] [S27] [S29]
- **Part II (Chapters 4 to 7)** engineers edge capture, covering asset identity and NTP/PTP time (Chapter 4), non-disruptive brownfield retrofits (Chapter 5), local edge vision and energy inference (Chapter 6), and layered network degradation and recovery (Chapter 7). [S09] [S12] [S13] [S16] [S19] [S21] [S22] [S26] [S29] [S40]
- **Part III (Chapters 8 to 11)** builds the governed data spine and decision surfaces, reconciling OPC UA, MQTT, Sparkplug, AAS and EPCIS (Chapter 8), specifying the two-tier historian (Chapter 9), designing role-specific dashboards (Chapter 10), and defining reversible application deployment (Chapter 11). [S04] [S05] [S07] [S08] [S15] [S16] [S24] [S30] [S40]
- **Part IV (Chapters 12 to 14)** governs proof of value and safe scale, establishing resource and product metering boundaries (Chapter 12), security, safety and recovery acceptance tests (Chapter 13), and a four-phase gated roadmap (Chapter 14). [S08] [S17] [S19] [S21] [S22] [S23] [S29] [S30]

A factory becomes intelligent not when its screens light up in the cloud, but when every thread, kilowatt and defect on the floor leaves a timestamped truth that a supervisor can act on before the shift ends.

CHAPTER SOURCE NOTES

S01 International Society of Automation. 2025-04-10.

S08 International Organization for Standardization. 2014-01.

S31 MDPI, Sensors. 2020-04-28.

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

S32 World Bank Enterprise Surveys. 2025.

S17 National Institute of Standards and Technology. 2023-09.

S22 National Institute of Standards and Technology. 2010-11-11.
S23 International Electrotechnical Commission. 2022-04-28.
R12 Business in Cameroon. Publication date unrecorded; inherited source-text check.
R01 Business in Cameroon. Publication date unrecorded; inherited source-text check.
R03 SODECOTON. Publication date unrecorded; inherited source-text check.
R07 Ecofin Agency. Publication date unrecorded; inherited source-text check.
R08 Business in Cameroon. 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.
R11 Business in Cameroon; MMI News. Publication date unrecorded; inherited source-text check.
R13 Business in Cameroon; MMI News. Publication date unrecorded; inherited source-text check.
R29 MDPI; industry inspection literature. Publication date unrecorded; inherited source-text check.
S26 MDPI, Electronics. 2025-09-18.
S09 Modbus Organization. 2012-04-26.
S30 MDPI, Sustainability. 2023-08-19.
S27 MDPI, Applied Sciences. 2024-06-21.
S29 MDPI, IoT. 2024-09-28.
S40 Eclipse Foundation. 2022-11-16.
S14 World Wide Web Consortium. 2023-12-05.
S16 European Telecommunications Standards Institute. 2024-04.
S13 Internet Engineering Task Force and RFC Editor. 2010-06.
S05 OASIS Open. 2019-03-07.
S21 CISA and joint international cybersecurity authorities. 2022-05-09.
S33 International Telecommunication Union. 2024-11-10.
S0102 World Wide Web Consortium. 2023-12-05.
S0101 OASIS Open. 2019-03-07.
S35 United Nations Industrial Development Organization. 2019.
S04 OPC Foundation. 2025-10-22.
S12 IEEE Standards Association. 2020-06-16.
S24 International Organization for Standardization and International Electrotechnical Commission. 2022-10.
S02 International Electrotechnical Commission. 2013-05-22.
S25 Springer Nature, International Journal of Advanced Manufacturing Technology. 2016-09-05.
S19 International Electrotechnical Commission. 2020-06.
S10 IO-Link Community. 2024-06.
S07 Industrial Digital Twin Association. 2024-06.
S15 GS1 AISBL. 2022-06.

GREEN FUTURE TECH

LIMITED / HONG KONG SAR

02

PART I: DECISIONS BEFORE DEVICES / USE-CASE CONTRACT

Requirements begin with the decision, not the device

Each connected-factory use case needs an event, an owner, a permitted delay, a fallback and a measurable consequence before any sensor or platform is selected.

CHAPTER 02
DECISION TEST

Each connected-factory use case needs an event, an owner, a permitted delay, a fallback and a measurable consequence before any sensor or platform is selected.

Specify the decision before selecting the device. GFT proposes a use-case contract that connects an observed production event to an accountable action, with explicit limits on delay, uncertainty and permitted data loss. That contract should determine what the factory captures, where it interprets the observation, who can act and how the result will be checked. A connected asset earns its place when this chain is testable. A screen receiving values does not, by itself, demonstrate that the chain works.

This chapter translates production monitoring, inspection, work-in-progress tracking, energy, water and traceability into requirements that a factory owner can review before procurement. It proposes a method for GFT's Cameroon textile programme, not a record of an installed system. Later chapters develop the responsibility architecture, event design, capture equipment and dashboard implementation. Here the deliverable is the operating contract those components must satisfy. Standards and research supply specific reference disciplines; proposed thresholds, roles and examples remain design choices requiring site approval.

2.1

Production monitoring starts with a shared KPI definition

ISO's description of its manufacturing KPI standard explicitly associates indicators with formulas, constituent elements, time behaviour and units. It also identifies their user groups and production methodologies. This supports a disciplined starting point: a requested production display needs a definition before it needs a visual design. The publicly opened abstract establishes this scope; it does not establish that a proposed factory calculation conforms to every provision of the full standard. [S08]

GFT recommends beginning each requirement with the sentence the supervisor should be able to complete: given this event and this context, I will take this action before this decision window closes. For output monitoring, that might mean checking whether the current order will finish within the approved shift plan. The requirement then names the order, operation, product unit, responsible supervisor and permitted intervention. Buying another counter is justified only after identifying the missing evidence in that sentence.

The proposed KPI dictionary should distinguish machine cycles, produced units, accepted good units and rejected units. It should explain whether a counted object is a metre of fabric, a bundle or a completed garment. A machine cycle must not acquire the label of saleable output without an explicit conversion and quality rule. GFT would ask the production owner to sign off that conversion and keep it visible wherever the measure is aggregated. Changes to the conversion should produce a new definition version, with the effective operating period recorded.

Time needs equal care. GFT recommends stating the beginning and end of the reporting window, the treatment of breaks and planned stops, the rule for an operation crossing the boundary and the status of late corrections. A line manager and plant manager should receive values derived from the same underlying event rules even when their reporting windows differ. A denominator should never change merely

because a screen has less space for its label. Where the source lacks a reliable state, the requirement should preserve that limitation rather than silently classify the period as productive.

An illustrative shift calculation shows why these choices matter. Assume a planned production window of 480 minutes, including 60 minutes of planned excluded time, with 42 minutes of recorded downtime during the remaining 420 minutes. The assumed operating time is 378 minutes. Availability under this particular illustrative definition is 378 divided by 420, or 90 percent. Dividing the same operating time by the full 480 minutes instead gives 78.75 percent. Neither display can be assessed until the denominator and its exclusions are disclosed. These assumed values are a definition exercise, not a factory baseline.

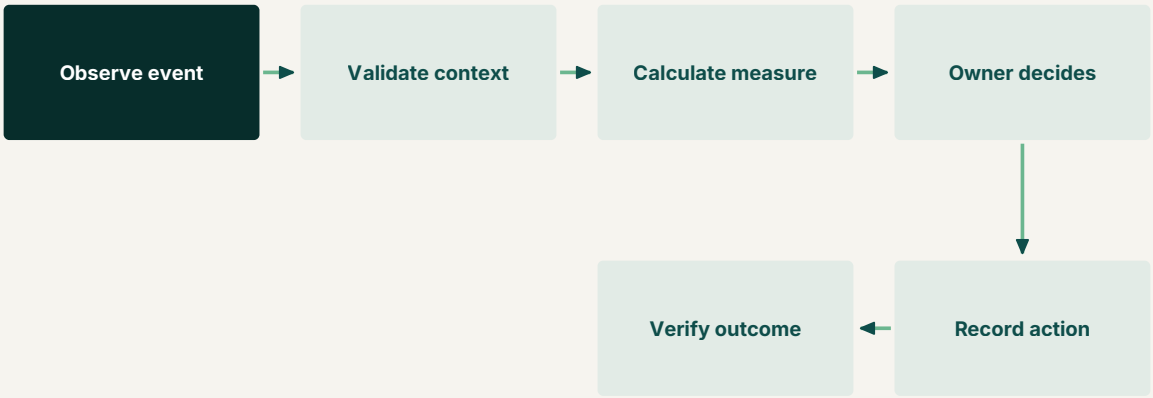
For an overall equipment effectiveness requirement, GFT would therefore demand the chosen availability, performance and quality formulas separately, alongside their raw elements. The definition should state the reference cycle basis, how product changes are treated and whether reworked output returns to the good-unit count. A combined measure should be withheld when its elements are unavailable or incompatible. This is a proposed acceptance rule for the contract, not a claim that the chapter reproduces the standard's complete calculation specification.

A useful proposed procurement test is to give competing implementations the same retained event set and the same signed dictionary, then compare their reported result. GFT would include a late rejection, a product change and a corrected stop classification in that set. The implementation should explain every difference through an agreed rule. The acceptance pack should retain the input events, definition version and expected output so the factory can repeat the test after a software change. This establishes calculation reproducibility within the chosen contract. It does not establish whether the counter correctly observed the original physical process, which needs a separate capture check.

The source-to-decision chain in **Exhibit 5**, p. 32 is an illustrative design sequence. It makes verification part of the requirement rather than leaving action closure to the display designer.

EXHIBIT 05 / ILLUSTRATIVE DECISION SEQUENCE

A useful requirement ends in verified action



ILLUSTRATIVE [S08] [S30] GFT proposed sequence, informed by KPI definition discipline and bounded pilot guidance. This is not a deployed workflow.

The sequence is deliberately independent of technology. Its test is whether a reviewer can reproduce the measure and understand the subsequent action. It does not establish which protocol, sensor or software package should supply the evidence. Those choices belong downstream of the agreed contract.

2.2

Quality inspection requires an acquisition specification as well as a model

The textile inspection study opened for this chapter reports that poor factory lighting initially produced false positives through shadows and inconsistent texture appearance. Adjusting the lighting mitigated those problems. It also discusses acquisition conditions and the interaction between fabric movement and imaging. This is evidence that an inspection requirement must cover the measurement station, not only the model. It is not a transferable accuracy guarantee for other fabrics or factories. [S26]

GFT recommends defining the physical decision first: identify a suspect region, mark it, hold the affected material and request disposition from the authorised quality role. The contract should say whether the system advises, requests inspection or initiates an approved non-safety process action. These are different permissions. An advisory model should not obtain automatic rejection authority through an undocumented connection added during commissioning. Protective safety and control acceptance remain separate responsibilities, consistent with NIST's attention to OT performance, reliability and safety. [S17]

The proposed acquisition specification should identify the fabric families in scope, line operating envelope, lighting arrangement, camera mounting, field of view, focus checks and reference samples. It should define the smallest defect that the business needs to distinguish, the material position needed to find that defect again and the acceptable ambiguity at an inspection boundary. These are requirements to settle through process review and trials. They are not universal settings to copy from a research apparatus or supplier demonstration.

GFT would place the defect taxonomy beside the acquisition specification. Each class needs a practical description, a disposition rule and an authorised reviewer. A crease, shade anomaly and hole can lead to different actions even when the software draws the same box around them. Borderline cases should have an explicit review queue with the original image, material identity, acquisition condition and model version available. The reviewer should be able to record disagreement without rewriting the original observation.

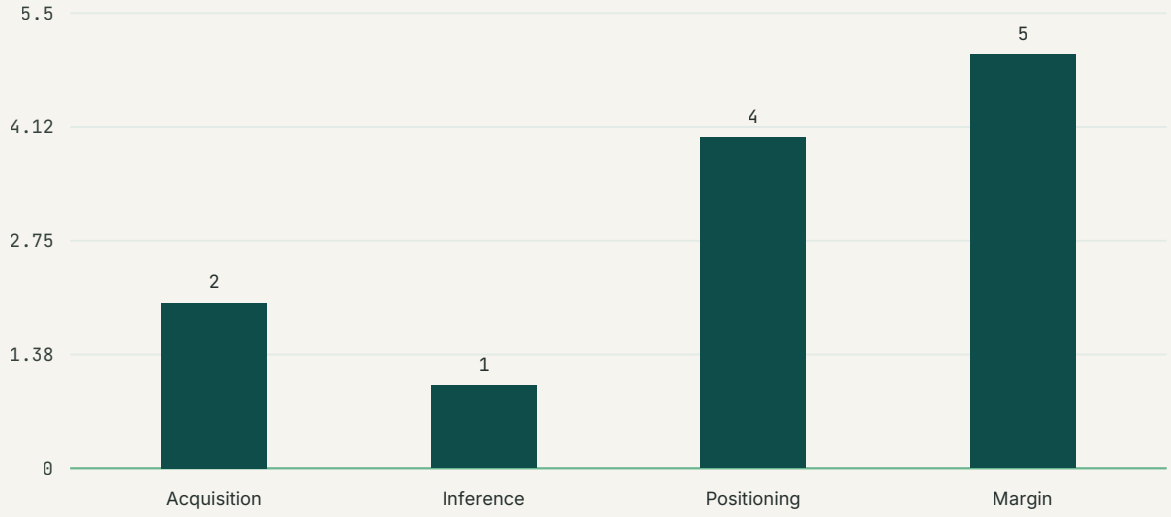
Acceptance should assess missed defects and unnecessary holds separately. GFT recommends a retained evaluation set agreed by the quality owner, covering representative products and difficult operating conditions. The specification should prevent images of the same physical defect from appearing as apparently independent evidence across evaluation groups. It should also separate model processing time from the elapsed interval between the fabric reaching the camera and an operator being able to locate it. Fast inference cannot compensate for an untraceable material position.

> **WORKED EXAMPLE: A DEFECT DECISION HAS A PHYSICAL DEADLINE**
>
> An illustrative inspection line moves fabric at an assumed 30 metres per minute. The review or marking point is assumed to be 6 metres downstream of the camera. Available travel time is 6 divided by 30 minutes, or 12 seconds. Assume acquisition and transfer require 2 seconds, inference and validation require 1 second, and marking or operator positioning requires 4 seconds. The assumed chain uses 7 seconds, leaving 5 seconds of margin. These are engineering assumptions for this example, not measurements or recommended universal limits.
>
> If an illustrative off-site service adds 8 seconds of delay, the same chain requires 15 seconds and misses the assumed physical window by 3 seconds. The requirement must then change placement, marking arrangements, operating conditions or disposition procedure before acceptance. The arithmetic does not justify slowing or stopping equipment without an approved operating and safety procedure. It simply identifies an incompatible timing proposal.

Exhibit 6, p. 34 keeps the assumed timing components visible. The unallocated margin is not evidence of reliability; it is the part of the illustrative window available for variation and response.

EXHIBIT 06 / SECONDS, ILLUSTRATIVE

The assumed inspection chain leaves a limited response margin



ILLUSTRATIVE Assumptions: fabric speed 30 m/min and distance 6 m give 12 s travel. Acquisition 2 s + inference 1 s + positioning 4 s = 7 s; margin = 12 - 7 = 5 s. Off-site delay of 8 s would raise the chain to 15 s. No measured GFT result.

GFT would also specify what prevents an unsafe or misleading acceptance shortcut. A demonstration on stationary samples should not satisfy a moving-fabric requirement unless the owner has explicitly limited the intended use to stationary inspection. An evaluator should retain the operating condition with each test result and identify excluded material families. If a new fabric falls outside that boundary, the contract should route it through the agreed review path until qualification is complete. The factory then knows which decisions the evidence permits, and the supplier knows which changes require another trial rather than an undocumented model adjustment.

The requirement should also define what happens when the camera fails, the light check fails or the model cannot classify the material. GFT recommends an explicit unavailable state and a documented manual inspection route. Contingency guidance distinguishes planning from testing and personnel preparation; the factory should demonstrate that the fallback can actually be activated by the responsible shift team. [S22]

2.3

Work-in-progress tracking connects physical identity to execution

The opened denim case study reports that RFID could support tracking and tracing of work in process in denim production. Its abstract describes design, configuration and deployment as part of the implementation approach. That is useful textile-specific evidence for considering RFID. It does not establish a read-success rate, cost saving or preferred tag technology for the factory being designed here. [S25]

GFT recommends naming the tracked object before specifying the reader. The object might be a bundle carrier, fabric roll or individual item, but the contract must explain what happens when material is split, combined, reworked or transferred to another carrier. A carrier identifier should not become permanent proof of product identity after the carrier is emptied. The owner must define the association, the event that changes it and the evidence required to make that change.

The proposed contract should distinguish a physical read from an accepted execution event. A tag being visible near a station may demonstrate presence within the reader's effective area. GFT would require an additional, explicitly defined rule before interpreting that observation as operation completion. That rule could involve an operator confirmation or an independent process event, according to the approved design. GS1's EPCIS standard concerns sharing visibility events in a business context; this reinforces the distinction between raw observations and contextual records. [S15]

For an illustrative bundle route, GFT would specify arrival, start, completion, transfer and exception as separate event meanings. These names are proposed process states, not a prescribed external vocabulary. The route owner should determine which transitions are valid and what evidence permits each one. An unexpected transition should create a reconciliation task rather than silently move the bundle to the location most convenient for the schedule. A dispatch decision is only as defensible as the interpretation of those transitions.

Duplicate and missing observations require different responses. GFT recommends retaining enough observation identity to recognise repeated delivery of the same read, while allowing legitimate later visits by the same object. A broad instruction to discard repeated tag identifiers would conceal rework and return movement. Conversely, assuming that the next read proves every intermediate operation occurred would invent process history. The requirement should state whether a missed station can be reconciled, who may approve the correction and how the original uncertainty remains visible.

An illustrative reconciliation starts with 120 bundles released to a route. Assume that 112 have valid completion records, 5 have confirmed work still in process and 3 have unresolved location. The assumed balance is 120 = 112 + 5 + 3. The release count therefore reconciles, but the route is not fully explained operationally. Presenting all 120 as completed because they were once observed would erase the unresolved balance. Presenting the unresolved bundles as lost would also exceed the evidence. The appropriate action is a search and reconciliation task under the proposed contract.

GFT would test the tracking requirement with ordinary movement, return movement, tag replacement and a bundle that passes the wrong read point. The acceptance record should connect each scripted physical movement to its expected interpreted event and final location. The production owner should witness representative cases and confirm that the resulting dispatch instruction is usable. This is the proposed test method; the chapter claims no completed factory trial.

Exhibit 7, p. 36 compares the narrow findings of opened sources. It deliberately records the limit next to each capability so that a purchase specification cannot turn a published example into an unqualified local performance promise.

EXHIBIT 07 / SOURCE SCOPE

Opened sources establish functions rather than local guarantees

Capture	Source finding	Evidence limit
RFID, S25	Denim WIP tracking	No local read-rate proof
Vision, S26	Lighting affects results	No universal accuracy
Energy, S29	Edge stores meter data	No local savings proof
EPCIS, S15	Shares visibility events	No physical verification

SOURCED [S25] [S26] [S29] [S15] Scope summaries of the opened denim abstract, textile inspection paper, energy-monitoring paper and EPCIS standard. No numerical performance is transferred between sites.

For handover, GFT recommends an object-association exercise witnessed by production and quality together. The team should follow a chosen physical bundle through a planned split, confirm the resulting identifiers and inspect how its parent relationship appears in the retained record. It should then replace a carrier and demonstrate that the old identifier no longer follows the material. These proposed checks expose ambiguity that a reader-connectivity test cannot resolve. Acceptance should preserve both the expected relationship and the observed result, with unresolved discrepancies assigned to an operating owner before the tracking result is used for dispatch.

The comparison supports selecting a capture approach for investigation. It cannot establish radio performance, camera fitness, meter accuracy or process compliance at an untested installation. Those remain acceptance questions for the relevant use-case owner.

2.4

Resource dashboards begin at meters with declared boundaries

The opened industrial energy-monitoring paper describes receiving meter data at the edge, storing it and calculating additional analytical parameters locally. This provides an implementation example of a meter-to-analysis chain. Its existence does not establish that the same design will save energy at another site. GFT recommends attaching a proposed action and a verifiable boundary to each resource measure before treating it as a dashboard requirement. [S29]

For electricity, the proposed contract should name the incoming supply or subcircuit being measured, the assets inside that boundary and the events that can change the connection. A cumulative register and an instantaneous power observation answer different questions. GFT would require the data owner to specify which one supplies each calculation, how reset or rollover is recognised and how an interrupted reporting interval is labelled. Missing observations should not quietly become zero consumption.

GFT also recommends deciding whether the intended comparison concerns total use, productive use, idle use or intensity per accepted output. A process owner should identify the intervention each comparison supports. A maintenance investigation might begin with abnormal idle consumption; an operating review might compare matched production periods. The proposed requirement should retain product mix and operating state where they affect interpretation. Otherwise a reduction caused by lower production could be presented as improved process performance.

An illustrative intensity example assumes a measured boundary consuming 240 kWh while producing 600 accepted garments during the same window. The assumed intensity is 240 divided by 600, or 0.4 kWh per accepted garment. If another display divides the same consumption by 800 started garments, it reports 0.3 kWh per started garment. The calculations are both arithmetically correct for their assumed definitions. They support different decisions. GFT recommends retaining the denominator wording rather than labelling either result simply as garment efficiency.

Water needs a similarly explicit chain. The opened Pakistani textile study presents water-level fluctuations as information for water-use optimisation. That supports considering level telemetry for an operating use case. It does not turn a level reading into an independently verified flow measurement or establish the impact of water consumption at a different factory. The requirement must identify the measurement actually available before selecting the environmental claim. [S27]

GFT would ask whether the proposed water measure is tank level, abstraction, process input, discharge or recirculation. A calculation from level change to volume would need a declared vessel relationship and treatment of simultaneous inflows and outflows. If those inputs are absent, the display should retain the language of level rather than invent a consumption total. Reused water should not be counted as new abstraction each time it passes a process boundary. The operating owner should approve the chosen accounting rule.

ISO's published scope for water-footprint assessment concerns products, processes and organisations using life-cycle assessment. A water-volume dashboard therefore should not be labelled a complete water footprint merely because it can total readings. GFT recommends keeping the operating measurement and any later environmental assessment as distinct requirements, joined by documented boundaries and evidence. Detailed allocation and assessment methods belong in the resource chapter of this paper. [S38]

The resource requirement should also describe how a later correction reaches a previously issued report. GFT proposes preserving the original observation, the corrected interpretation and the reason for the change. If a meter association was wrong, the reviewer needs to identify the affected operating periods and product allocations rather than amend only the latest screen. A recalculated trend should carry a correction state until the responsible owner approves its use. This is especially relevant when the same measurement feeds both an immediate operating decision and a later external statement: their review and release responsibilities should be specified separately.

The practical handover for both resources is a boundary sketch, a measurement definition, a matched production denominator and an exception action. GFT would require the owner to show how an implausible reading, a disconnected meter or a changed process connection affects the report. An attractive historical trend should not override an unresolved boundary change. Procurement can then compare instruments against the actual measurement requirement rather than a generic request for sustainability data.

2.5

Traceability requirements stop at the evidence actually available

EPCIS is intended to let different applications create and share visibility event data within and across enterprises. The opened standard also describes business context around identified objects. That makes it relevant to the exchange of contextual production and shipment evidence. It does not establish the truth of a supplier declaration or make a sensor observation equivalent to a complete product history. [S15]

GFT recommends defining a traceability question in plain language before defining the exchange payload. A quality owner might need to identify which released units used a suspect input lot. A dispatch owner might need to establish which finished bundles were packed into a shipment. Each question requires an explicit association and a boundary where that association was observed or declared. A later scan should not imply that an unrecorded upstream relationship has been verified.

The proposed contract should distinguish direct observation, operator confirmation, supplier declaration and inference. Each can be useful, but each needs an evidence type and a responsible role. A material-origin declaration, for example, should retain its issuer and the declared lot association. A factory should not rewrite it as a physical observation solely because it has been imported into the same database as machine events. GFT would require corrections to preserve the prior statement and the authority for the replacement.

Product-passport preparation adds a separate requirement boundary. In its dated adoption announcement, the Council of the European Union described the ecodesign framework as including information requirements and a digital product passport, with requirements to be set through delegated acts. That historical announcement supports maintaining adaptable product information. It is not a current legal determination of the fields, deadlines or obligations applicable to a specific textile product or market. [S0201]

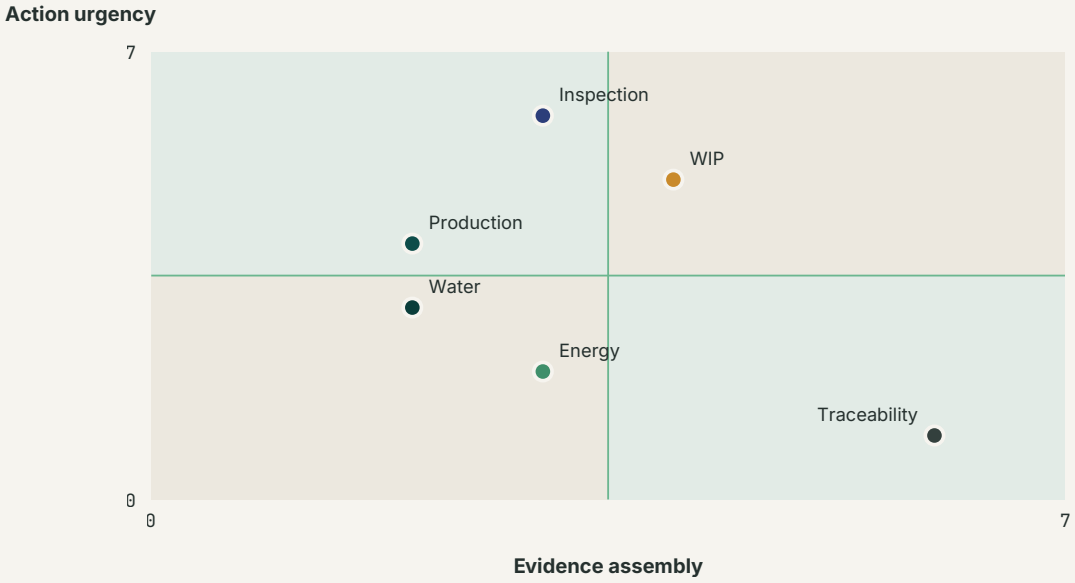
For this architecture, GFT therefore recommends exportable, versioned evidence rather than a claim that a universal textile passport has been implemented. The product owner should identify the destination market, product category and current applicable requirements through a separate review before asserting compliance. The proposed engineering contract should retain configurable evidence mappings, a record of the rule version used and a way to withdraw a superseded representation. This chapter provides engineering guidance, not legal, investment or tax advice.

GFT would test evidence export through a bounded retrieval exercise before accepting the traceability requirement. The reviewer should choose a released object, request its supporting associations and open the underlying records outside the originating display. The export should retain evidence status, effective time, correction history and the applicable access limits. An unresolved upstream association should remain unresolved in the export. This proposed exercise tests whether the factory can supply a reviewable record when the usual application is unavailable or replaced. It is a portability and evidence check, not a legal opinion on the sufficiency of the resulting product information.

The illustrative matrix in **Exhibit 8**, p. 39 separates action urgency from evidence assembly. Its coordinates are explicitly assigned design scores: they express the proposed comparison, not measured latencies or a ranking established by the cited research. A site team should replace them with its own requirements when the relevant decisions are agreed.

EXHIBIT 08 / ILLUSTRATIVE ORDINAL SCORES, 1 TO 6

Urgent action and extensive evidence are different requirements



ILLUSTRATIVE [S08] [S25] [S26] [S29] [S27] [S15] GFT design illustration. Higher x means more cross-record evidence assembly; higher y means a shorter assumed action window. Coordinates are assigned judgements, not source-derived scores or measured factory latency. Energy and water refer to operating review, not protective control.

Inspection and traceability can share event identity while needing different response windows and evidence packages. The proposed matrix makes that design discussion visible. It should not be used to downgrade a resource alarm or an urgent recall: those would be different use cases with different action contracts.

2.6

Every requirement card includes failure and recovery behaviour

NIST's contingency guide distinguishes testing recovery capabilities, preparing personnel through training and finding planning gaps through exercises. GFT adopts that distinction as an acceptance discipline. A written fallback is a proposal until the relevant people can activate it and the system can recover the records needed for the decision. The guide is a federal information-system reference; this chapter adapts its planning discipline rather than asserting factory certification. [S22]

GFT recommends a requirement card containing the decision, physical object, event meaning, source owner, decision owner, freshness limit, permitted loss, retention purpose, security boundary, fallback and acceptance evidence. Each field should have a clear role accountable for approval. An integrator can implement the card, but should not decide an operating trade-off without that role's agreement. A requirement that says only real time, secure or scalable leaves the important choices unresolved.

Freshness should describe the age of the evidence used for the decision, rather than only the time when the display refreshed. GFT would define when a value becomes stale, who is informed and whether the action remains permitted. The card should distinguish delayed delivery from missing acquisition. Data arriving after a link recovers may repair historical reporting without restoring the opportunity to act on time. Recovery acceptance must therefore check both the record and the decision service.

Permitted loss should be defined by data class. GFT recommends stating whether an observation may be sampled, whether an interpreted completion event must be preserved and whether a missed image invalidates an inspection interval. Retention should follow the evidence needed for review and correction, with an approved deletion rule. It should not be an indefinite promise that everything will remain available. The security boundary must identify the authorised exchange and access roles; OT requirements must also account for performance, reliability and safety. [S17]

Exhibit 9, p. 41 presents illustrative requirement-card entries for discussion. The compact exhibit lists failure-sensitive fields; the decision descriptions and event definitions remain in the accompanying text. Its timing values are assumed to make review concrete and require site validation before procurement.

EXHIBIT 09 / ILLUSTRATIVE DESIGN ENTRIES

Failure behaviour belongs in the use-case card

Field	Hourly output	Bundle transfer	Inspection
Source owner	Production	Route owner	Quality
Decision owner	Supervisor	Dispatcher	Inspector
Source time	Count window	Transfer time	Frame capture
Freshness	At most 5 min	Before release	Within 12 s
Loss rule	Flag gaps	Reconcile all	Flag uncovered
Retention	Trend review	Route history	Disposition proof
Zone	Line reporting	Operations	Inspection cell
Fallback	Signed tally	Controlled log	Manual check
Recovery	Reconcile counts	Replay once	Restore coverage

ILLUSTRATIVE [S17] [S22] Illustrative thresholds: output evidence age at most 5 min; WIP confirmation at transfer; inspection timing follows the worked example. Retention durations and permitted-loss quantities remain approval conditions, not invented standards. Each card requires a defined zone and approved exchange.

For acceptance, GFT would script a source interruption, an external-link interruption, a repeated event and a correction after reporting. Each test should identify the expected display state, permitted action and final reconciled record. The operator should demonstrate the fallback with the ordinary tools available on shift. The reviewer should then compare physical evidence, source history and displayed totals. Passing data transport alone would not satisfy this proposed contract.

> **GFT NOTE**
>
> Approve a use-case card only when the operating owner can describe the action, the data owner can reproduce its evidence and the recovery owner can demonstrate the fallback. Start with a bounded production decision. Keep unapproved fields as explicit acceptance holds, with a responsible role and closure evidence. Do not resolve them by copying a supplier's default setting.

Textile SME guidance recommends small pilot projects to test selected technologies before broader implementation. GFT's proposed requirement review should therefore finish with a bounded acceptance trial whose success condition is settled in advance. The trial should demonstrate a usable decision and record unresolved exceptions, rather than report only connected-device counts. [S30]

A final proposed review should distinguish acknowledged alarms from closed decisions. GFT would require the owner to record the intervention, its completion evidence and the condition that permits closure. Where an action produces no observable result, the case should remain open or close with an explicit unresolved finding. This prevents the system from treating a button press as proof that a produc-

tion problem was corrected. The retained record also gives the next shift a concrete starting point: what happened, what was attempted and what evidence is still needed. The contract should define that handover before application design begins.

The resulting handover is a set of approved operating contracts. Each links an event to a decision, states how evidence can fail and defines what successful recovery means. That gives the architecture designer a stable question to answer and gives the factory owner a practical basis for rejecting an impressive demonstration that cannot support the required action.

Select the device only after the decision has a contract.

CHAPTER SOURCE NOTES

S08 International Organization for Standardization. 2014-01.

S30 MDPI, Sustainability. 2023-08-19.

S26 MDPI, Electronics. 2025-09-18.

S17 National Institute of Standards and Technology. 2023-09.

S22 National Institute of Standards and Technology. 2010-11-11.

S25 Springer Nature, International Journal of Advanced Manufacturing Technology. 2016-09-05.

S15 GS1 AISBL. 2022-06.

S29 MDPI, IoT. 2024-09-28.

S27 MDPI, Applied Sciences. 2024-06-21.

S38 International Organization for Standardization. 2014-08.

S0201 Council of the European Union. 2024-05-27.

GREEN FUTURE TECH

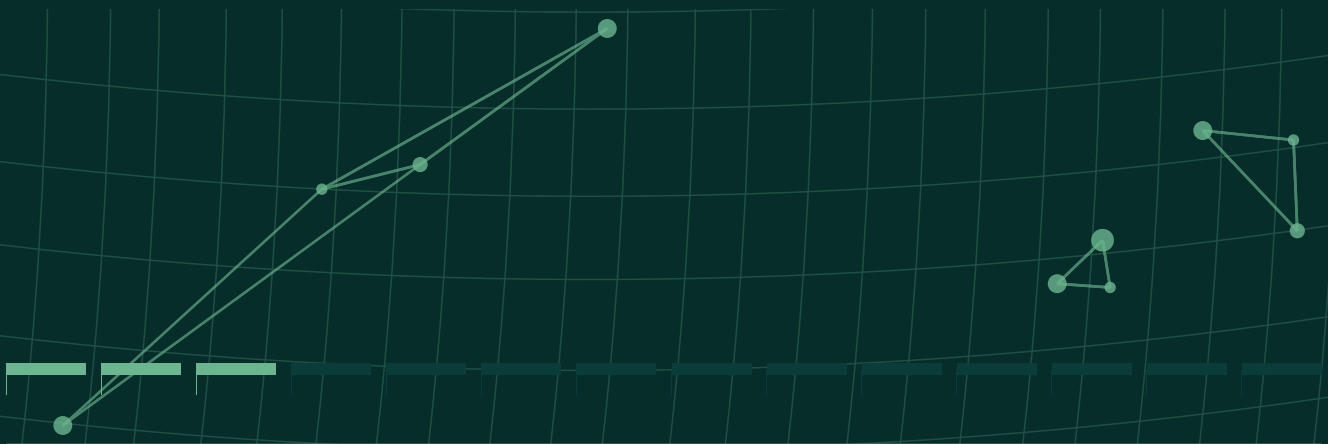
LIMITED / HONG KONG SAR

03

PART I: DECISIONS BEFORE DEVICES / REFERENCE MODEL

The architecture separates control, operations and enterprise responsibilities

A useful reference architecture assigns each function to a responsibility boundary and information contract before it assigns that function to a product.



CHAPTER 03
DECISION TEST

A useful reference architecture assigns each function to a responsibility boundary and information contract before it assigns that function to a product.

A connected textile factory fails when its software treats every signal as if it belonged to one flat network. [S01] When an enterprise resource planning system, a cloud dashboard or an ad-hoc script reaches directly into a programmable logic controller, nobody can state which system owns the definition of an order, a shift, a good garment or an equipment state, and the reporting path becomes an undocumented dependency whose polling load or network failure can disturb machine timing. [S02] [S17] Green Future Tech Limited (GFT) resolves both defects before selecting a gateway, broker or application vendor by partitioning the factory into explicit responsibility layers governed by international standards for enterprise-control integration, operational technology security and asset description. [S01]

The architectural rule is direct: protective safety and closed-loop machine control stay at the equipment boundary, manufacturing operations management owns the shop-floor context that turns raw cycles into production events, and enterprise systems exchange scheduled requirements and reconciled performance summaries across a defined transaction boundary. [S01] In Cameroon, where the inherited sector baseline 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, mills must prove operational discipline under severe infrastructure constraints. [R06] Because 93.3 percent of surveyed Cameroonian firms reported electrical outages in the 2024 World Bank Enterprise Survey, with 10.4 outages in a typical month across formal non-agricultural private firms, a textile plant cannot rely on continuous wide-area connectivity or remote cloud orchestration to coordinate its floor. [S32] Separating control, operations and enterprise responsibilities allows the mill to lose external links without losing shop-floor coordination, and to lose shop-floor reporting without compromising machine stability or human protection. [S17]

This chapter expands the edge-to-dashboard pipeline introduced in GFT's Cameroon textile programme into an implementable responsibility model, drawing on ANSI/ISA-95 Part 1, IEC 62264-1, NIST SP 800-82 Rev. 3, IEC 62443-3-2, CISA Cybersecurity Performance Goals 2.0, IEC functional safety, ISO 22400-2, RAMI 4.0, ETSI Multi-access Edge Computing (MEC), W3C Web of Things specifications, ISO 23247-4 and the Asset Administration Shell (AAS) metamodel. [S01] [S02] [S08] [S17] Governance of the shared data spine belongs to paper wp14, while the daily, weekly and monthly management cadence belongs to paper wp02. [S08]

3.1

ISA-95 fixes the enterprise-to-control conversation

Every connected-factory programme faces an early temptation to wire shop-floor sensors straight into whatever visualisation tool or enterprise database the sponsor already licenses. [S01] Moving a register value from a ring-spinning frame or a garment-sewing line to an executive chart does not explain which work centre produced the count, which production order was active, whether the machine was in a changeover test or normal execution, or who holds authority to amend a miscounted bundle. [S02] Without an explicit model of functions and shared information, every dashboard project reinvents its own plant hierarchy and reconciliation logic inside private scripts. [S01]

EXHIBIT 10 /

The reference stack assigns decisions before technologies



ILLUSTRATIVE [S01] [S02] [S14] [S16] Illustrative GFT responsibility stack informed by ISA-95 asset and interface scope (S01), IEC 62264 Level 3 and Level 4 separation (S02), W3C selective edge collection (S14) and ETSI MEC virtualised edge execution (S16). Represents an architectural allocation of duties, not a deployed factory measurement.

As shown in **Exhibit 10**, p. 45, GFT organises the factory into six vertical responsibility tiers before naming any software package or network protocol. [S01] The diagram is an illustrative architectural design synthesising standards scope from ISA-95, IEC 62264, W3C Web of Things and ETSI MEC rather than an operating record from an existing mill. [S02] While **Exhibit 10**, p. 45 clarifies which layer owns each class of decision, verifying boundary enforcement in deployment requires the interface contracts and acceptance tests defined later in this paper. [S19]

3.2 PHYSICAL-ASSET ORGANISATION PRECEDES SOFTWARE SELECTION

ANSI/ISA-95.00.01-2025, known as ISA-95 Part 1, summarises the scope of the manufacturing operations and control domain, addresses how physical assets of a manufacturing enterprise are organised, lists the functions associated with the interface between control functions and enterprise functions, and describes the information shared among those functions. [S01] Before an integrator configures a

gateway on a carding machine, a slasher-dyeing range or a sewing line, the architecture must define how the enterprise partitions its holdings into sites, areas, production lines, work centres, work units and equipment modules. [S01]

In Cameroon's textile sector, where public reporting on existing industrial assets such as CICAM has noted multi-site structures across Garoua and Douala with distinct spinning, weaving and finishing capacities, physical organisation cannot be left to network IP addressing. [R12] When a technician swaps a failed controller on a stenter frame or moves an overlock sewing machine between garment lines, device addresses change while work-centre responsibility remains, or the equipment moves while its maintenance history must follow the physical asset. [S01] Adopting the ISA-95 physical-asset hierarchy ensures that every measurement point is bound first to an engineered equipment role and only secondarily to a communication endpoint. [S01] Order promising, commercial purchasing, financial accounting and master production scheduling sit on the enterprise side of the ISA-95 boundary, whereas work-centre dispatching, material lot tracking, setup verification and inline quality disposition sit inside the manufacturing operations and control domain. [S01]

3.3 LEVEL 3 TRANSACTIONS PREVENT DIRECT ENTERPRISE-TO-MACHINE COUPLING

IEC 62264-1:2013 formalises this separation by describing the manufacturing operations management domain at Level 3 and its activities, alongside the interface content and associated transactions within Level 3 and between Level 3 and Level 4. [S02] In the IEC 62264 hierarchy, Level 0 defines the physical production processes, Level 1 covers sensing and manipulating those processes, Level 2 addresses monitoring and supervisory control, Level 3 governs manufacturing operations management across production, maintenance, quality and inventory workflows, and Level 4 manages business planning and logistics. [S02]

The critical rule in IEC 62264-1 is that Level 3 and Level 4 communicate through structured transactions rather than shared memory or direct device polling. [S02] Level 4 issues a production schedule specifying product families, quantities and due dates alongside the approved bill of materials. [S02] Level 3 checks equipment availability, crew rosters, work-in-process balances and tooling readiness, translates the commercial requirement into executable work orders for specific lines and shifts, and returns verified production performance responses that Level 4 consumes for inventory valuation and capacity replanning. [S02]

When a factory omits the Level 3 boundary and connects Level 4 enterprise software directly to Level 1 or Level 2 controllers, transaction integrity collapses. [S02] A programmable logic controller on a loom or cutting table counts mechanical strokes, picks or knife cycles; it does not know whether those cycles produced first-quality fabric for an export order, calibration waste during a reed change, or rework on a rejected batch. [S02] Even where Level 3 functions are initially hosted in a lightweight operations database rather than a full manufacturing execution system, preserving the IEC 62264-1 transaction contract between Level 3 and Level 4 keeps business planning and shop-floor execution decoupled. [S02]

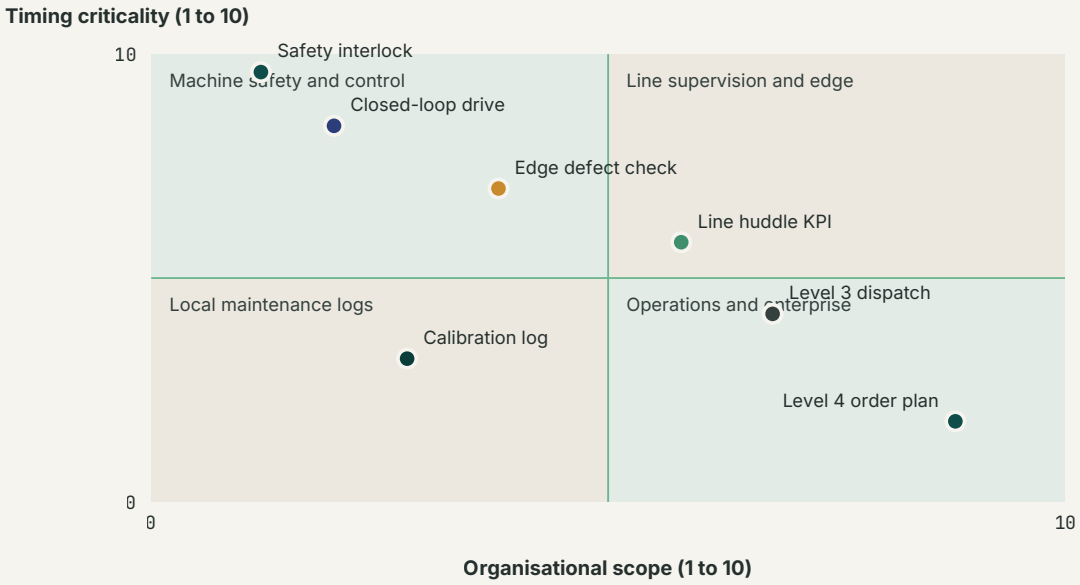
3.4

Control remains close to the machine and below the analytics path

Separating enterprise planning from manufacturing operations is necessary, but the most dangerous failures occur lower in the stack, between data acquisition and machine control. [S17] Integrators sometimes place machine start-stop commands, speed setpoints or recipe downloads on the same broker, gateway and application path used for dashboard telemetry, exposing deterministic control loops to the latency, availability and credential risks of analytical networks. [S17]

EXHIBIT 11 / ORDINAL ARCHITECTURAL SCALE (1 TO 10)

Control, operations and enterprise functions have distinct owners



DERIVED [S01] [S02] [S17] [S23] Derived qualitative placement from ISA-95 functional scope (S01), IEC 62264 Level 3/4 boundaries (S02), NIST OT performance and reliability guidance (S17) and IEC functional safety independence (S23). X-axis scores organisational scope from single machine point (1) to multi-site enterprise (10); Y-axis scores timing criticality from monthly review (1) to immediate hardware protection (10).

The functional separation across timing criticality and organisational scope is mapped in **Exhibit 11**, p. 47. [S01] This derived exhibit translates the domain boundaries of ISA-95, IEC 62264-1, NIST SP 800-82 Rev. 3 and IEC functional safety into a two-dimensional ordinal grid from 1 to 10. [S02] Safety interlocks and closed-loop drives occupy the upper-left quadrant because they act within milliseconds on a single machine module, whereas Level 3 dispatching and Level 4 order planning occupy the lower-right quadrant because they coordinate plant-wide or enterprise-wide resources over hours, days and months. [S17] Because the coordinates are ordinal placements derived from standard definitions rather than empirical latency measurements, **Exhibit 11**, p. 47 establishes ownership separation rather than hardware benchmarks. [S23]

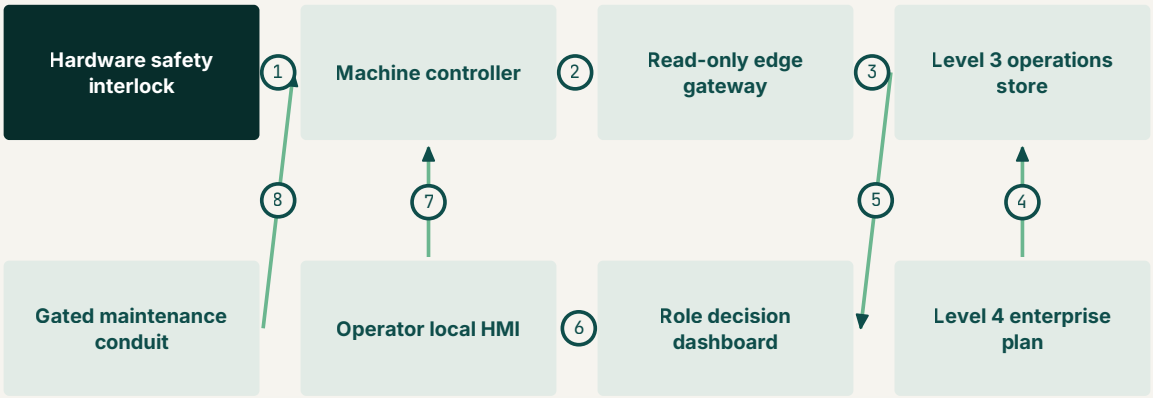
3.5 READ-ONLY TELEMETRY AND COMMAND PATHS FOLLOW ASYMMETRIC RULES

NIST Special Publication 800-82 Rev. 3 provides guidance on securing operational technology while addressing its unique performance, reliability and safety requirements. [S17] Operational technology controls physical motion, heat, tension and chemical dosing where timing jitter or unexpected lockouts can damage machinery, ruin batches or injure personnel. [S17] GFT therefore enforces strict asymmetry between the upward flow of production evidence and the downward flow of control commands. [S17]

Upward evidence capture from Level 1 and Level 2 to edge nodes and Level 3 operations defaults to read-only access. [S17] Whether an edge gateway reads cycle counters from a legacy controller, monitors power draw on a compressor or receives fabric-inspection alerts from an optical cell, the acquisition interface is engineered so the upstream consumer cannot write to controller registers, alter motion setpoints or consume controller scan-cycle capacity beyond tested limits. [S17] In brownfield environments, exemplified by UNIDO's Industrial Development Report 2020 finding that almost 70 percent of surveyed Ghanaian manufacturing firms operated in the analogue category (a relevant comparator for mixed-generation regional plants though not a Cameroon survey), many machines have older controllers with limited communication buffers. [S35] Downward paths from Level 3 to Level 2 are restricted to supervisory coordination, such as making an approved recipe or order identifier available for a qualified operator to review and accept at the local human-machine interface (HMI); a dashboard tile or cloud model never receives a direct actuation path into Level 1 or Level 2 equipment. [S02]

EXHIBIT 12 /

A command path and an evidence path are not interchangeable



1. Hardwired stop overrides motion immediately
2. Read-only telemetry leaves control zone
3. Observations buffered at Level 3 store
4. Schedules exchange via Level 3 transactions
5. Governed KPIs populate role-specific views
6. Operators verify dispatches at local HMI
7. Local setpoints execute within PLC limits
8. MFA maintenance sessions use gated conduit

ILLUSTRATIVE [S01] [S02] [S17] [S19] [S20] [S23] Illustrative flow architecture informed by ISA-95 and IEC 62264 boundaries (S01, S02), NIST OT security (S17), IEC 62443-3-2 zones and conduits (S19), CISA MFA goals (S20) and IEC functional safety (S23). Shows logical separation of actuation, telemetry, operational dispatch and maintenance access.

The separation of these pathways is illustrated in **Exhibit 12**, p. 49. [S19] Process observations move outward from the machine controller through a read-only edge gateway to the Level 3 operations store and role-specific dashboards, while operational adjustments return through human-verified local HMIs and engineering changes pass through a separate, time-bounded maintenance conduit. [S20] Because **Exhibit 12**, p. 49 is an illustrative logical topology, it does not replace site-specific firewall rule tables, wiring diagrams or safety-integrity calculations, which must be engineered and tested for each line during commissioning. [S19]

3.6 FUNCTIONAL SAFETY AND SECURITY ZONES ISOLATE THE MACHINE BOUNDARY

Two further standards anchor the lower boundary: IEC functional safety principles and IEC 62443-3-2 security zones and conduits. [S23] Within the overall safety of a system, IEC defines functional safety as the aspect that provides automatic protection in response to errors. [S23] On textile machinery, where carding cylinders, calender nip rollers, pressurised dyeing vessels and automated cutters present mechanical and thermal hazards, emergency stops, light curtains and guard interlocks must act automatically and deterministically. [S23] GFT prohibits routing any protective safety

function through an edge gateway, an MQTT broker, an inference container or a supervisory dashboard; safety states needed for downtime accounting are read via auxiliary isolated contacts. [S17]

To prevent network-borne threats from crossing into machine control, IEC 62443-3-2 establishes industrial automation and control system (IACS) security zones and conduits as a key design concept. [S19] Safety systems, Level 1/2 machine control, edge nodes, Level 3 operations servers, corporate Level 4 IT systems and cloud services reside in distinct zones. [S19] CISA's Cybersecurity Performance Goals 2.0 specify that multi-factor authentication (MFA) must be enabled on all operational technology accounts and systems that can be accessed remotely, including vendor and maintenance accounts, remotely accessible user and engineering workstations, and remotely accessible HMIs when available. [S20] Isolating engineering maintenance into a dedicated, authenticated, time-bounded conduit prevents dashboard traffic from sharing credentials or paths with controller programming interfaces. [S19]

> **GFT NOTE**
>
> Never allow an analytics vendor or ERP integrator to request write credentials to a machine controller as a shortcut for logging operator reason codes or order numbers. [S17] Provide an edge-zone operator terminal or local HMI screen that writes to the Level 3 operations store, leaving Level 1 and Level 2 machine control logic read-only. [S02]

3.7

Manufacturing operations owns the production context

A raw sensor reading has no intrinsic business meaning. [S01] Knowing that a sewing motor drew 420 watts at 10:14:02 or that a folding-table optical sensor triggered 180 times in fifteen minutes, as an illustrative telemetry example, tells a plant manager nothing about delivery performance, rework loss or electricity per shirt. [S08] Raw signals become actionable production evidence only when joined with operational context: which production order was running, which fabric roll and dye lot were loaded, which specification applied, which crew was assigned, and whether the line was scheduled for production or maintenance. [S01] In the GFT reference architecture, Level 3 manufacturing operations management is the sole authoritative owner of this production context. [S02]

3.8 ORDER, ROUTE, MATERIAL AND CREW BINDINGS BELONG AT LEVEL 3

Both ISA-95 Part 1 and IEC 62264-1 organise the manufacturing operations domain around four resource pillars: personnel, equipment, material and process segments across production, maintenance, quality and inventory operations. [S01] When a spinning, weaving, dyeing or garment order executes on the floor, each operational event binds those four pillars at a specific work centre over a bounded time interval. [S02]

When a dashboard bypasses Level 3 and assembles context inside front-end chart queries, consistency breaks down because production, quality and energy charts bind equipment, material and time using incompatible boundaries. [S08] Fixing this requires a single Level 3 context service that records the active binding of work centre, production order, operation step, material lot, crew state and planned calendar state; when an edge node publishes a machine cycle, bundle scan or metre reading, the Level 3 pipeline attaches the authoritative context active at the event's source timestamp. [S15]

3.9 KPI ELEMENTS RISE FROM THE EDGE WHILE DEFINITIONS REMAIN GOVERNED

ISO 22400-2:2014 bridges shop-floor event capture and management reporting by presenting manufacturing KPIs through their formula and corresponding elements, their time behaviour, their unit or dimension, and other characteristics. [S08] The standard separates basic measured elements, such as planned busy time, actual production time, good quantity, scrap quantity and rework quantity, from composite indicators calculated from those elements, such as availability, effectiveness, quality ratio and overall equipment effectiveness (OEE). [S08]

In GFT's responsibility model, basic elements originate at the edge and Level 2 supervision boundary, while the KPI formula, element reconciliation and aggregation window are governed at Level 3. [S08] The Level 3 KPI service applies the ISO 22400-2 dictionary definition, separates planned from unplanned downtime, prevents double-counting of rework, and confirms that numerator and denominator share the same work-centre and time boundary. [S08]

> **METHOD**
>
> To verify that a factory's responsibility architecture preserves KPI integrity across Levels 1, 2, 3 and 4, GFT applies a four-step boundary audit before any dashboard tile is accepted into operation: [S08] 1. **Element origin trace:** Identify the exact Level 1 or Level 2 physical sensor, controller counter or operator terminal that generates each elementary time or quantity input required by the ISO 22400-2 formula. [S08] 2. **Context binding check:** Confirm that the Level 3 operations layer binds work-centre identity, order, lot and calendar state to the elementary event using the source timestamp rather than the database arrival timestamp. [S02] 3. **Single-calculation enforcement:** Verify that the composite KPI formula executes once inside the governed Level 3 KPI service and that line, plant and enterprise views read the same stored result record. [S08] 4. **Denominator reconciliation:** Prove across a shadow-mode test window that the sum of good quantity, scrap quantity and rework quantity at Level 3 reconciles with the physical material balance reported to Level 4 inventory. [S02]

3.10

RAMI 4.0 adds asset and life-cycle views to the layer model

While the ISA-95 and IEC 62264 hierarchy excels at separating operational time horizons and enterprise-to-control transactions, a single vertical stack of levels is not enough to specify a connected factory. [S03] Viewing the factory solely as five horizontal layers from Level 0 to Level 4 treats every machine as a fixed set of runtime

process tags and overlooks what happens across design, commissioning, firmware maintenance, reconfiguration and decommissioning. [S03] Plattform Industrie 4.0 developed the Reference Architecture Model Industrie 4.0 (RAMI 4.0) to resolve this limitation; its specification explains the three dimensions of RAMI 4.0 and shows how the framework unites the most important aspects of Industrie 4.0. [S03]

EXHIBIT 13 /

Reference models answer different architecture questions

Reference standard	Primary scope in source	GFT architectural artefact
ANSI/ISA-95 Part 1 (2025)	Physical-asset organisation and enterprise-control interface functions (S01)	Plant hierarchy register and enterprise-to-operations boundary contract
IEC 62264-1 (2013)	Level 3 MOM activities and transactions within Level 3 and with Level 4 (S02)	Level 3 context store and structured Level 3-to-Level 4 transaction schema
RAMI 4.0 (2018)	Three-dimensional framework uniting hierarchy, life cycle and layers (S03)	Multi-stage asset matrix across commissioning, operation and maintenance
ISO 23247-4 (2021)	Technical requirements for information exchange between twin entities (S06)	Explicit synchronisation, freshness and interface rules for digital views
IDTA AAS Part 1 (2024)	Administration Shell structure for exchanging asset information (S07)	Durable machine and submodel records separated from live telemetry

DERIVED [S01] [S02] [S03] [S06] [S07] Derived architectural comparison synthesising the stated purpose of ISA-95 Part 1 (S01), IEC 62264-1 (S02), RAMI 4.0 (S03), ISO 23247-4 (S06) and IDTA AAS Part 1 (S07) alongside the concrete design artefact GFT requires in a connected textile mill.

How these complementary reference models fit together is summarised in **Exhibit 13**, p. 52. [S01] Each standard addresses a distinct structural question: ISA-95 and IEC 62264 partition vertical operational responsibility, RAMI 4.0 locates assets across their life cycle and representation layers, ISO 23247-4 governs information exchange between digital-twin entities, and the Asset Administration Shell structures durable asset descriptions. [S06] Because **Exhibit 13**, p. 52 is a derived synthesis of standards scope, it defines which design artefacts an engineering team must produce rather than certifying that any off-the-shelf software bundle implements all five standards automatically. [S07]

3.11 THREE ARCHITECTURAL AXES PREVENT FLAT NETWORK DRAWINGS

When an integrator presents a topology diagram showing controllers, switches, an MQTT broker and a cloud server, three questions are routinely conflated: where an object sits in the plant hierarchy, how its physical and digital representations relate, and what stage of its life cycle is being governed. [S03] On the hierarchy axis, RAMI 4.0 extends plant levels down to the field device and outward to the connected product and value network, so a fabric roll, a cut-bundle tote with an RFID tag, or a finished garment carrying a data carrier participates directly in the architecture. [S15] On the representation axis, separating asset, integration adapter, communication transport, information semantics, functional logic and business rules ensures that changing a transport protocol does not force a redesign of the KPI dictionary. [S04] On the life-

cycle axis, distinguishing between the type (the engineered design, recipe template or software release) and the instance (the serialised loom on the floor, Tuesday's dye batch, or the deployed model on an edge camera) prevents configuration drift. [S03]

3.12 COMMISSIONING, OPERATION AND MAINTENANCE REQUIRE DISTINCT VIEWS

Applying the life-cycle dimension alongside the Asset Administration Shell metamodel shows why a machine, an inspection camera, a line and a textile product require distinct views across their working lives rather than one static tag list. [S07] During commissioning, engineers need schematics, interface profiles, safety checklists and calibration certificates. [S17] During operation, supervisors and Level 3 operations need event-driven signals: cycle counts, run-stop state, defect alerts, active order binding and energy draw. [S08] During maintenance, technicians need cumulative hours, fault history, spare-part specifications and change records. [S24] Separating these submodels while anchoring them to one permanent ISA-95 and AAS asset identity avoids both unindexed folders and bloated real-time telemetry packets. [S07]

3.13

Edge and cloud are placement choices within the same contracts

A common misconception in industrial digitalisation is that edge and cloud represent rival philosophies of factory management. [S16] GFT rejects that binary: edge and cloud are physical execution placements for software workloads, whereas control, operations and enterprise remain invariant responsibility boundaries. [S16] Moving an application module from a plant server room to a virtualised edge host on the floor or to a remote data centre changes its latency, bandwidth cost and outage exposure, but never changes the information contract it must honour. [S01]

3.14 VIRTUALISED EDGE HOSTS EXECUTE LOCAL ADAPTATION AND BUFFERING

The European Telecommunications Standards Institute (ETSI) specification GS MEC 003 V3.2.1 explains that Multi-access Edge Computing enables the implementation of MEC applications as software-only entities running on top of a virtualisation infrastructure located in or close to the network edge. [S16] Instead of buying a separate proprietary appliance for every sensor vendor, the textile factory deploys managed software modules on virtualised edge hosts inside the plant network. [S16] When wide-area connectivity degrades, a realistic planning assumption given ITU's 2024 report that 14 percent of Africa's population lacks access to a mobile-broadband network and given the need to verify actual site link reliability independently of population coverage, local software entities continue running on the plant LAN. [S33] Coupled with uninterruptible power supplies and backup generation sized for the 93.3 percent outage incidence reported across surveyed Cameroonian firms, the edge host preserves acquisition, store-and-forward persistence and line-side operator displays throughout external blackouts and link drops. [S32]

3.15 SELECTIVE PUBLICATION REPLACES RAW CONTROLLER REPLICATION

How the edge tier interacts with upstream operations and cloud services is clarified by World Wide Web Consortium (W3C) Web of Things standards. [S14] W3C's Web of Things Architecture 1.1 states that an industrial edge device collects selected data

from various controllers and makes it available to a cloud backend service, for instance for remote monitoring via a dashboard or to analyse it for preventive maintenance. [S14] Complementing that pattern, W3C's Web of Things Thing Description 1.1 defines a formal information model where a Thing Description describes the metadata and interfaces of Things as abstractions of physical or virtual entities, while W3C's Web of Things Profiles defines technical specifications providing assertions to which conformant consumers and things must conform. [S0301] [S0302]

Two practical design rules follow from these specifications. [S14] First, publication across the edge boundary is selective rather than a raw dump of internal controller memory. [S14] An industrial energy-monitoring study published in MDPI IoT in 2024 demonstrates this pattern: once the edge receives the metre's data, it stores the information in a local database server, followed by a data-processing method to derive nine additional analytical parameters before forwarding selected results. [S29] Second, every exposed edge interface declares its metadata, units, interaction affordances and conformance assertions explicitly, following OPC UA layered information modelling (which OPC-10000-1 notes is layered to isolate core design from computing technology and network transport), W3C Thing Descriptions and W3C Profiles. [S04]

WORKED EXAMPLE

To illustrate how responsibility separation and selective edge publication govern bandwidth and storage sizing under outage conditions, consider an illustrative integrated textile plant with 40 production machines, 12 utility metres and 4 inline optical fabric-inspection stations. [S29] All numerical inputs in this example are assumed illustrative engineering parameters used solely to demonstrate the boundary calculation method, not measurements from a GFT facility or forecasts for a Cameroonian mill. [S14]

Assumed illustrative inputs at the Level 1 and Level 2 machine boundary: [S14]

- 40 machines polled locally for 25 raw controller registers each at 1 hertz, at an illustrative 120 bytes per sample frame: 40 multiplied by 25 multiplied by 1 multiplied by 120 equals 120,000 bytes per second of local polling traffic confined to the machine-to-edge conduit. [S14]
- 12 utility metres read locally every 5 seconds at an illustrative 200 bytes per reading: 12 multiplied by 0.2 multiplied by 200 equals 480 bytes per second. [S29]
- 4 inline inspection cameras capturing raw frames locally at an illustrative 8 megabytes per second each inside the isolated camera-to-edge link: 4 multiplied by 8 equals 32 megabytes per second of raw image data that never leaves the inspection cell. [S16]

Illustrative selective reduction at the Edge-to-Level 3 boundary: [S14]

- Each machine publishes an illustrative 1 contextualised state or counter event every 10 seconds at 400 bytes per event: 40 multiplied by 0.1 multiplied by 400 equals 1,600 bytes per second. [S01]
- Each utility metre publishes its 1-minute locally derived parameter summary (computing nine derived parameters at the edge as in S29) at an illustrative 600 bytes per minute: 12 multiplied by 10 bytes per second equals 120 bytes per second. [S29]
- Each inspection station publishes defect-event metadata at an illustrative average of 1 event every 20 seconds (800 bytes), plus a cropped diagnostic image of 50,000 bytes only when a defect is flagged (assumed illustratively at 1 defect per 120 seconds per station): 4 multiplied by (40 plus 416.67) equals approximately 1,827 bytes per second. [S16]

Resulting boundary comparison and 24-hour outage buffer calculation (illustrative):

- Total raw local acquisition rate across machine, metre and camera links equals 32,120,480 bytes per second (about 32.12 megabytes per second). [S16]
- Total selective event stream crossing from the edge layer into Level 3 operations equals 1,600 plus 120 plus 1,827, which is 3,547 bytes per second (about 3.55 kilobytes per second), an illustrative reduction of 99.989 percent in upstream data volume. [S14]
- Across an illustrative 24-hour outage (86,400 seconds), buffering the selective event stream requires 3,547 multiplied by 86,400, which equals 306,460,800 bytes (about 306.5 megabytes) of edge disk space, while retaining 24 hours of

> raw 1-hertz machine register history and 5-second metre logs requires 120,480 multiplied by 86,400, which equals 10,409,472,000 bytes (about 10.41 gigabytes). [S29]

3.16

Digital representations must state what is exchanged

In many industrial-strategy presentations, the phrase "digital twin" is used loosely to describe anything from a three-dimensional rendering of a building to a basic line chart of temperature readings. [S06] GFT's reference architecture grounds digital representations in two explicit international specifications: ISO 23247-4 for digital-twin information exchange and the IDTA Asset Administration Shell (AAS) Part 1 metamodel for durable asset information. [S06]

3.17 DIGITAL-TWIN INTERFACES SPECIFY SYNCHRONISATION AND EXCHANGE RULES

ISO 23247-4:2021, part of the ISO 23247 digital twin framework for manufacturing, identifies the technical requirements for information exchange between entities within the reference architecture. [S06] In a connected textile factory, applying ISO 23247-4 means that no digital representation is accepted without a declared exchange contract between the observable manufacturing element and its digital counterpart: [S06]

- **Entity scope and fidelity:** Exactly which physical attributes of the observable element (such as spindle speed, dye-bath liquor ratio, warp beam identifier and cumulative run hours) are represented, and which physical phenomena remain unmodelled. [S06]
- **Direction and authority of synchronisation:** Whether the interface is strictly observational or whether supervisory parameters may flow back through a validated control entity under human approval. [S06]
- **Freshness and divergence behaviour:** How frequently state updates are exchanged, how source timestamps are preserved, and how the digital representation flags itself as stale when an outage pauses synchronisation. [S22]
- **Cross-tier interoperability:** How the digital representation exchanges state with Level 3 execution services or supply-chain visibility systems without forcing every consumer to parse raw Level 1 fieldbus registers. [S06]

Research on Industry 4.0 adoption in textile SMEs recommends the implementation of small pilot projects to test and validate the effectiveness of selected technologies. [S30] A mill does not need a three-dimensional simulation of its spinning shed on day one; it needs a state-synchronised representation of work-centre status, order progress and resource consumption whose exchange interfaces obey ISO 23247-4. [S06]

3.18 ADMINISTRATION SHELLS SEPARATE DURABLE ASSET MEANING FROM LIVE STREAMS

While ISO 23247-4 governs exchange requirements between runtime entities, the Industrial Digital Twin Association (IDTA) specification IDTA 01001-3-0-1 defines selected specifications of the structure of the Asset Administration Shell (AAS) to enable meaningful exchange of information about assets and Industrie 4.0 components between partners in a value creation network. [S07] Integrating the AAS metamodel with GS1 EPCIS Release 2.0 completes the outward boundaries of GFT's responsibility architecture, because EPCIS enables disparate applications to create and share visibility event data both within and across enterprises. [S15] Together, these standards prevent the factory from mixing three different classes of data: [S07] 1. **High-frequency or cyclic process telemetry:** Sub-second or second-by-second machine observations that stay inside the Level 1, Level 2 and edge domains and are summarised into the historian. [S14] 2. **Durable asset and product master submodels (AAS):** Slowly changing structured descriptions, including machine nameplates, calibration certificates, sensor scaling ranges, maintenance manuals and product specification templates, referenced by identifier rather than repeated inside every telemetry packet. [S07] 3. **Contextualised operational and supply-chain visibility events (EPCIS and Level 3 records):** Discrete records of what happened, when, where and why, such as a cotton bale lot entering carding, a fabric roll passing inspection, or a garment shipment leaving the loading bay, shared across Level 3, Level 4 and authorised external partners. [S15]

Maintaining these boundaries requires ongoing management discipline under the NIST Cybersecurity Framework (CSF) 2.0 (which guides organisations in managing cybersecurity risks), ISO/IEC 27001:2022 (which guides establishing, maintaining and continually improving an information security management system) and NIST SP 800-34 Rev. 1 (which distinguishes recovery testing, personnel training and plan exercises). [S18] [S24] [S22]

> **WARNING**
>
> A reference architecture diagram on paper does not prevent boundary violations during software maintenance. [S24] If an integrator or internal developer bypasses the Level 3 context layer to meet a short-term reporting deadline, for example by opening a direct database query from a cloud dashboard into an edge buffer or adding an undocumented write register to a PLC gateway, the separation of responsibilities degrades silently. [S19] Resolving this constraint requires automated conduit firewall rules under IEC 62443-3-2, mandatory schema and change-control reviews under ISO/IEC 27001, and periodic architectural audits before any new use case moves from pilot to production. [S24]

By fixing who owns control, who owns operational context, and who owns enterprise planning before any hardware or software is procured, the connected textile factory gains a stable spine. [S01] Individual sensors, gateways, edge hosts and dashboard applications can then be upgraded, replaced or scaled across lines without putting machine safety, data integrity or shop-floor continuity at risk. [S02] The next chapter

moves inside the event contract itself, specifying how asset identity, clock synchronisation, quality flags and provenance turn raw signals into audit-ready production events. [S01]

Separate the hand that protects the machine, the ledger that records the shift, and the plan that guides the enterprise before connecting the first wire.

CHAPTER SOURCE NOTES

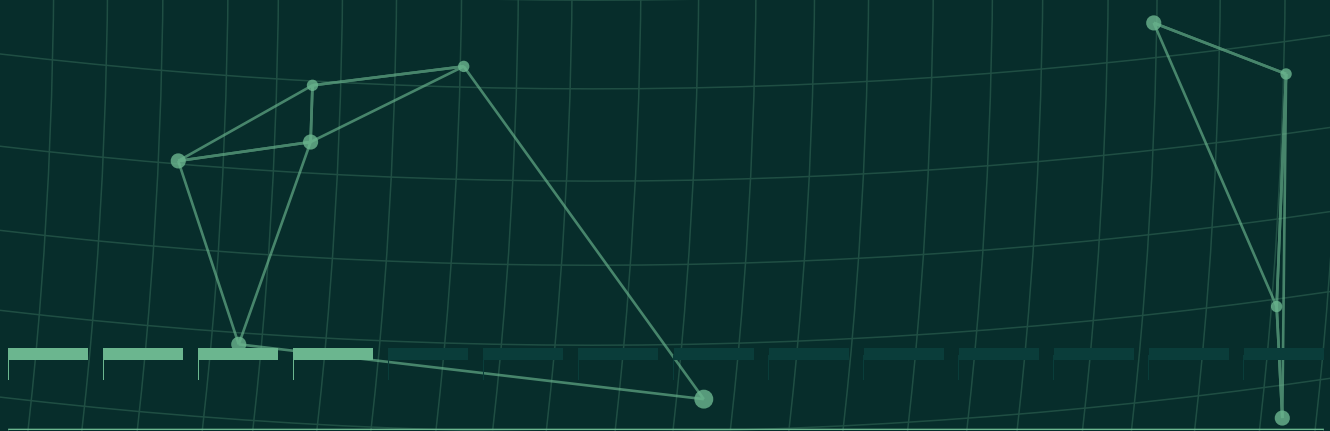
S01 International Society of Automation. 2025-04-10.
S02 International Electrotechnical Commission. 2013-05-22.
S17 National Institute of Standards and Technology. 2023-09.
R06 Cameroon Trade Hub. Publication date unrecorded; inherited source-text check.
S32 World Bank Enterprise Surveys. 2025.
S08 International Organization for Standardization. 2014-01.
S14 World Wide Web Consortium. 2023-12-05.
S16 European Telecommunications Standards Institute. 2024-04.
S19 International Electrotechnical Commission. 2020-06.
R12 Business in Cameroon. Publication date unrecorded; inherited source-text check.
S23 International Electrotechnical Commission. 2022-04-28.
S35 United Nations Industrial Development Organization. 2019.
S20 Cybersecurity and Infrastructure Security Agency. 2025-12.
S15 GS1 AISBL. 2022-06.
S03 Plattform Industrie 4.0. 2018-08-09.
S06 International Organization for Standardization. 2021-10.
S07 Industrial Digital Twin Association. 2024-06.
S04 OPC Foundation. 2025-10-22.
S24 International Organization for Standardization and International Electrotechnical Commission. 2022-10.
S33 International Telecommunication Union. 2024-11-10.
S0301 World Wide Web Consortium. 2023-12-05.
S0302 World Wide Web Consortium. 2025-11-04.
S29 MDPI, IoT. 2024-09-28.
S22 National Institute of Standards and Technology. 2010-11-11.
S30 MDPI, Sustainability. 2023-08-19.
S18 National Institute of Standards and Technology. 2024-02-26.

04

PART II: CAPTURE AT THE PRODUCTION EDGE / EVENT INTEGRITY

Identity and time turn signals into events

A value becomes production evidence only when its asset, occurrence time, uncertainty and transformation history remain inspectable.



CHAPTER 04
DECISION TEST

A value becomes production evidence only when its asset, occurrence time, uncertainty and transformation history remain inspectable.

A connected factory should admit a reading into production decisions only when it can identify the object measured, distinguish occurrence from delivery, and explain the reading's quality and origin. GFT proposes an event contract that preserves those distinctions from acquisition to display. The contract is a design recommendation, not a claim that GFT has deployed or certified the architecture. It connects the responsibility model established earlier to the capture methods developed next, without making the dashboard responsible for repairing missing evidence.

The standards provide complementary foundations. ISA-95 addresses physical-asset organisation; the Asset Administration Shell structures asset information; Sparkplug defines messaging identity and state; EPCIS supplies visibility-event semantics. Clock protocols address alignment, while the provenance standard considered here has a narrower master-data scope. These sources inform GFT's proposed contract, but none independently specifies the entire textile-factory implementation. [S01] [S07] [S40] [S15] [S13] [S41]

4.1

Asset identity follows a stable factory hierarchy

ISA's publication notice describes physical-asset organisation and information shared across the enterprise-control interface as part of ISA-95's scope. It supplies a reason to govern the factory hierarchy before connecting addresses. GFT recommends separate identifiers for the physical asset, its current operational position, and each measurement point. A network endpoint is then an attribute of a connection, rather than the enduring identity of the production object. [S01]

For a proposed textile installation, the hierarchy would describe site, area, line, work centre and measurement point, with explicit parent relationships. This is GFT's suggested operating model, not a requirement that every named level be copied literally from a standard. Production should approve the operational positions. Maintenance should approve the physical-asset register. The integration owner should maintain the relationship between these registers and acquisition endpoints. Disagreements should become visible exceptions before publication.

Consider an illustrative sewing machine moved from assembly to rework. Its physical identity should follow the machine, while its line assignment changes from the effective move time. A line's output query should use the assignment applicable when the work occurred. A maintenance query should follow the physical machine across assignments. Encoding the current line permanently inside the only machine identifier would make those questions difficult to separate. A readable location alias remains useful, provided the durable identity survives a rename.

GFT recommends retaining both the valid period of an assignment and the time its correction was recorded. Suppose an illustrative relocation was entered after the shift ended. The revised relationship can describe where the machine actually operated without pretending that the information was known earlier. This distinction lets an investigator reproduce the original report and explain a corrected report. Historical context should therefore be addressable by event time and revision, rather than supplied solely by today's register.

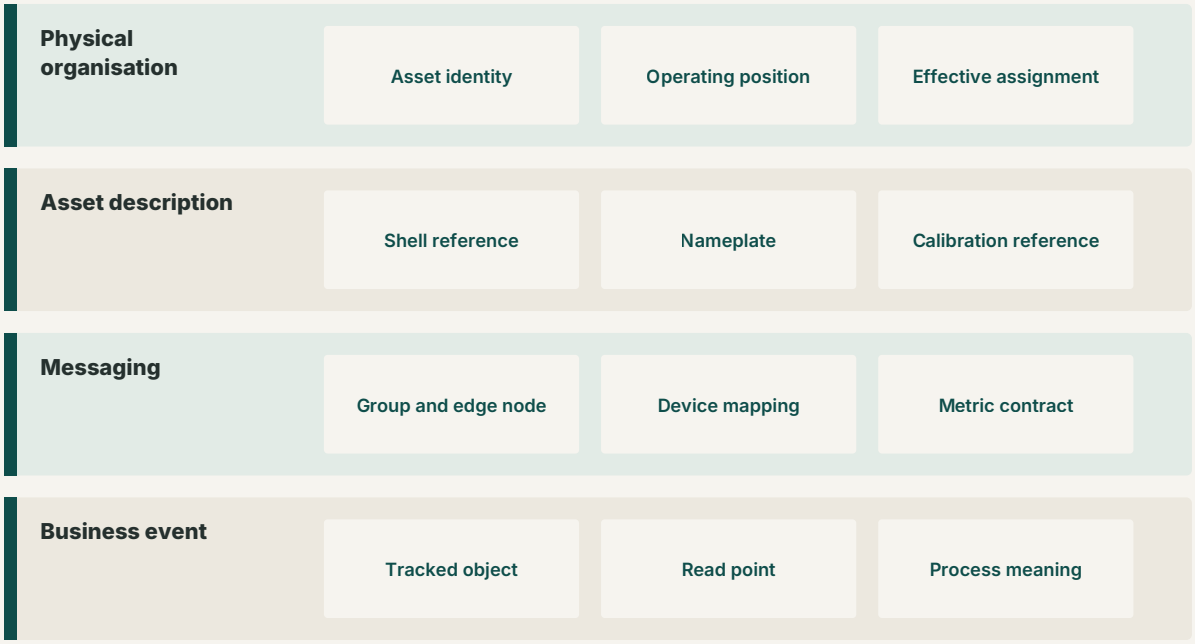
The Asset Administration Shell metamodel provides structures for exchanging information about assets and distinguishes the identifiers of represented assets and their shells. Its asset descriptions can support a link between an equipment record and its technical information. That does not establish that the described equipment has been calibrated, installed correctly or connected. GFT would require separate evidence for those operating states. [S07]

For an illustrative dyeing vessel, the physical asset, its temperature probe and its temperature measurement point should be distinguishable. Replacing the probe changes the installed sensor relationship and the applicable calibration reference. It need not create a new process position. Replacing the vessel itself may require a new physical-asset record. The event should retain enough identity to determine which arrangement existed at acquisition, including any interval when the replacement record was incomplete.

The proposed mapping in **Exhibit 14**, p. 61 links these concerns without treating them as interchangeable. The layers show relationships to be governed, not a mandatory software stack or an implemented installation. Keeping the labels separate also gives procurement a practical test: can a replacement gateway preserve asset history without retaining the old vendor's private addressing scheme?

EXHIBIT 14 / ILLUSTRATIVE RESPONSIBILITY MAPPING

Durable asset identity survives a change of messaging endpoint



ILLUSTRATIVE [S01] [S07] [S40] [S15] GFT design illustration informed by the cited standards. Layers are related representations, not equivalent identifiers or evidence of deployment.

The mapping cannot establish the accuracy of any register. GFT recommends commissioning it by tracing a chosen physical label through acquisition, the asset description and the displayed event. The reverse trace should reach the same object.

An unresolved or ambiguous relationship should suspend the affected inference, while retaining the raw observation for investigation. Inventing an asset match to keep a dashboard complete would remove the very exception the register needs to reveal.

Identity should also survive retirement. GFT proposes marking an asset record as retired while preserving its historical relationships, rather than reusing its key for replacement equipment. An illustrative probe removed from service may remain relevant to an earlier quality investigation. Its former measurements should still resolve to the retired record and the installation period. A spare of the same model should receive its own physical identity. Model sameness supports comparison of capability; it does not prove that the physical history is shared.

4.2

A topic namespace carries operational meaning

Sparkplug defines group, edge-node and device identifiers within its MQTT topic structure. Its node descriptor combines group and edge-node identity, while device identity is scoped beneath the node. These are messaging relationships. A factory should map them explicitly to its equipment register rather than assume that a device name alone uniquely identifies a physical machine everywhere. [S40]

GFT recommends an interface dictionary in which every published metric has an accountable owner, meaning, datatype, unit, valid states and mapping to a measurement point. The dictionary should distinguish an instantaneous observation, a cumulative counter, a state transition and a business completion. A consumer should not have to infer those categories from a short topic label. This contract belongs alongside the namespace because routing a value correctly does not determine how to interpret it.

An illustrative pulse counter shows why meaning matters. A maintenance test may increment it without producing saleable fabric. A reset may reduce it without reversing physical production. A repeated publication may carry an unchanged total. GFT would require the metric contract to state when differences may be calculated, how resets are identified, and which operating modes permit a production inference. The application would then retain the observation even when its interpretation is withheld.

Names should remain readable without becoming the sole repository of meaning. GFT would keep a concise display alias for technicians and a stable mapping key for software. Topic changes should pass a compatibility review that identifies affected consumers, historical mappings and rollback behaviour. If a measurement changes from an instantaneous rate to an interval average, keeping the old label while silently changing its interpretation should fail that review. A renamed display label need not trigger the same treatment.

The proposed change process is deliberately small. The metric owner records the requested meaning, the integration owner demonstrates the mapping, and the receiving application owner confirms the intended interpretation. A representative payload and an expected event should accompany the review. The decision concerns the

interface, not permission to write to machine control. Read-only publication and any command path require separately stated boundaries, even when the messaging technology supports both directions.

Sparkplug also provides birth, death and state behaviour, including stale treatment when communication is lost. Those conventions support consumer awareness of connection state. They do not replace GFT's proposed decision-specific rules for freshness, valid process mode and measurement quality. [S40]

> **GFT NOTE**
>
> Commission the identity dictionary before the attractive display. Require a technician and a production supervisor to trace a selected metric to its physical point, unit, acquisition rule and permitted interpretation. Retain the approved example payload with the contract. If either role cannot explain what a change in the value means, keep the metric out of production decisions until the definition is repaired.

GFT would test the namespace using deliberately conflicting illustrative cases: identical local device labels under different nodes, an equipment rename, a gateway replacement, and a metric with an incompatible unit. The expected outcome should be recorded before execution. Correct routing is necessary, but acceptance also requires that unsupported interpretations remain visible as exceptions. A rejection log should identify which mapping rule failed and who can correct it.

The chapter's boundary is the identity and meaning contract. Detailed protocol selection belongs to the interoperability discussion. The contract should nevertheless remain exportable independently of the broker configuration. That makes it possible to inspect what the factory believes a value means even when the messaging service is unavailable, and to compare successive configurations without reverse engineering a live installation.

GFT would also define how consumers react to an unfamiliar contract revision. A receiving application should either recognise the revision or retain the payload as uninterpreted evidence and expose the incompatibility. Guessing a unit from a previous version would make a successful connection look like successful interpretation. The commissioning exercise should include an intentionally incompatible revision and verify that the receiving application refuses the affected calculation. Compatibility evidence should name the tested producer and consumer configurations, so a later software change does not inherit acceptance automatically.

4.3

NTP and PTP serve different correlation needs

The clock requirement should begin with the decision that depends on event order. Comparing shift totals poses a different question from locating a camera observation against a fast-moving process. GFT recommends recording the permissible correlation error, the acquisition point and the consequence of an incorrect association before selecting a clock technology. Timestamp resolution alone should never be used as the acceptance measure.

RFC 5905 describes potential NTP accuracy in the tens of microseconds under conditions involving modern workstations and fast local-area networks. That statement supports considering NTP for general alignment, while preserving its conditions. It is not a measured result for the proposed factory, its legacy controllers or a congested wireless path. [S13]

The IEEE description of its precision clock synchronisation standard states support for accuracy and precision below a microsecond. PTP is therefore a candidate when tighter cross-device correlation is required. Its capability statement does not establish the accuracy of a particular sensor timestamp, network configuration or application. GFT would require commissioning evidence at the points where the relevant observations are actually timestamped. [S0401]

The comparison in **Exhibit 15**, p. 64 preserves that distinction. It compares published scope and capability with proposed acceptance evidence. The practical choice remains conditional on an end-to-end measurement. Neither an enabled synchronisation service nor a nominally precise timestamp establishes that a delayed acquisition path meets the business requirement.

EXHIBIT 15 / SOURCE STATEMENTS AND PROPOSED ACCEPTANCE CHECKS

Published clock capability remains conditional on site evidence

Candidate	Source boundary	Proposed check
NTP	Potential tens of microseconds on fast LANs	Measure acquisition-time error
PTP	Supports sub-microsecond synchronisation	Test the installed timing path
Unaligned nodes	Sensing integrity can be uncertain	Restrict cross-device ordering

ILLUSTRATIVE [S13] [S0401] [S31] Illustrative GFT selection aid. Capability statements are sourced; acceptance checks are recommendations. No accuracy has been measured at a GFT installation.

GFT recommends constructing an error budget across the physical transition, sensor sampling, acquisition software, local timestamping and clock alignment. Synchronising the computer clock addresses only part of that chain. The proposed budget should identify which contributions are measured, bounded by a supplier specification, or unknown. An unknown contribution should remain unknown until tested, rather than disappearing into the number of decimal places displayed.

An illustrative correlation test could drive a common observable transition into the relevant acquisition channels and compare the resulting events. The test should include normal operation, high acquisition load, loss of the time reference and recovery. The measured boundary must be explicit: agreement between gateway timestamps does not prove agreement between physical transitions when the gateways poll at different moments. The commissioning record should state what the test can and cannot resolve.

METHOD

GFT proposes accepting a temporal ordering claim only when its uncertainty intervals are separated. For an illustrative event with recorded time T and an accepted absolute error bound E, represent possible occurrence as T minus E through T plus E. If the intervals overlap, preserve local sequence evidence and state that cross-device order is unresolved. Do not treat an estimated error as a guaranteed bound without an agreed validation basis.

The interval rule is GFT's analytical recommendation. It avoids ranking observations merely because their displayed timestamps differ. A source can report fine-grained time while its acquisition uncertainty is much coarser. Conversely, local sequence can establish a device's own order when absolute time is degraded. Cross-device association should explain which evidence it used, including any process dependency, rather than relying on an unexplained global sort.

GFT would record clock state with the event: synchronised, operating without a current reference, or unresolved, using definitions in the interface dictionary. The maximum accepted interval without a reference should follow the measured drift and use case. On recovery, the application should retain the original recorded timestamp and the clock-state transition. Any corrected estimate belongs in a separate field with its method, preserving what the system originally knew.

Elapsed duration deserves its own treatment. GFT recommends using a local elapsed-time mechanism for acquisition intervals where the platform can support it, with a recorded relationship to civil time. A clock correction should not create an apparent negative machine cycle or extend downtime simply because the reporting clock moved. Events spanning a restart require an explicit continuity rule. Where continuity cannot be established, the interval should be marked incomplete and reconciled against independent evidence, rather than bridged by subtraction across unrelated timing sessions.

4.4

Source time and receipt time answer different questions

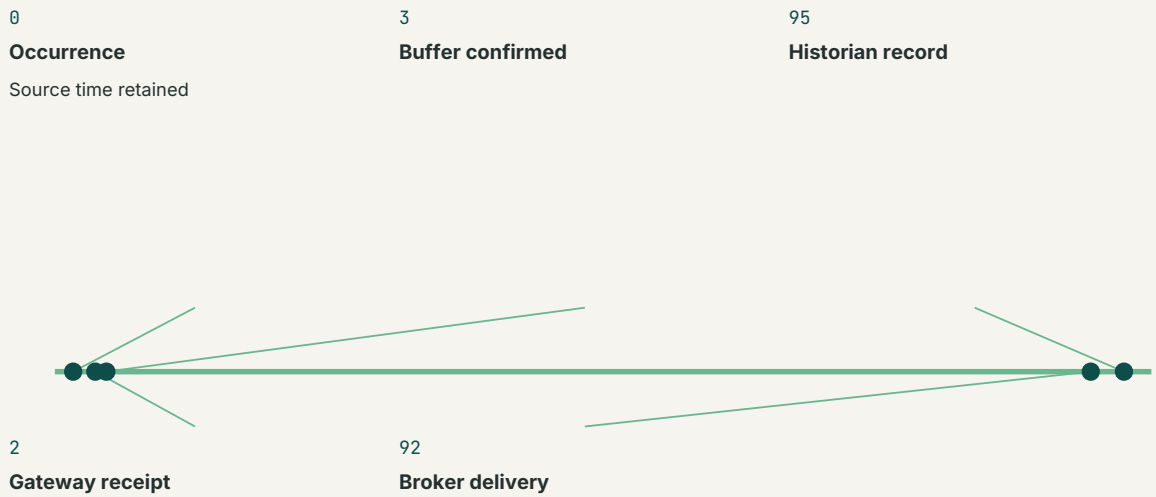
Source time should identify the best supported occurrence or acquisition time, with its meaning declared. Gateway receipt time identifies admission into the local processing path. Broker delivery and historian persistence describe later stages. GFT recommends recording these separately wherever the application needs them. If the sensor cannot timestamp occurrence, the gateway's observation time can be useful, but should not be relabelled as a more precise physical-event time.

EPCIS distinguishes event time from repository record time. Sparkplug distinguishes payload timing from metric timing and carries session state. These source semantics require careful mapping into a canonical event: an envelope's send timestamp is not automatically the time every measurement occurred, and repository capture is not gateway receipt. [S15] [S40]

The illustrative timeline in **Exhibit 16**, p. 66 makes the recording points explicit. Its times are invented to demonstrate the difference between an event and its subsequent movement. The physical event occurs before the gateway accepts it; the event remains locally buffered before a later transmission. None of those stages should overwrite the earlier time simply because delivery resumes.

EXHIBIT 16 / ILLUSTRATIVE SECONDS FROM PHYSICAL OCCURRENCE

Delivery delay leaves the original occurrence time intact



ILLUSTRATIVE [S13] [S15] [S40] Assumed teaching sequence, not measured latency: occurrence 0 s; gateway receipt 2 s; buffer confirmation 3 s; broker delivery 92 s; historian persistence 95 s. Source time remains 0 s.

Under the illustrative assumptions, occurrence-to-persistence delay is 95 seconds. The assumed gateway-to-broker interval is 92 minus 2, or 90 seconds. Those calculations describe different boundaries. Calling both measures network latency would obscure acquisition delay, buffering and processing. If the clocks are not demonstrably aligned, the differences should be reported with the relevant uncertainty rather than presented as exact measured durations.

WORKED EXAMPLE: DELAYED ARRIVAL DOES NOT PROVE REVERSED PRODUCTION

An illustrative inspection system records event A at 10:00:00.000 UTC with an accepted uncertainty bound of 40 milliseconds. A downstream reader records event B at 10:00:00.050 UTC with a bound of 30 milliseconds. These assumed bounds are teaching inputs, not measured performance or recommended tolerances.

Relative to 10:00:00.000, A could lie between minus 40 and plus 40 milliseconds. B could lie between 50 minus 30 and 50 plus 30, or 20 to 80 milliseconds. The illustrative overlap is 20 to 40 milliseconds. The timestamp evidence alone therefore cannot establish their physical order, even though the nominal difference is 50 milliseconds.

In this illustrative case, suppose that A reaches the historian at 10:01:35 UTC and B at 10:00:03 UTC. Their nominal occurrence-to-record delays are 95 seconds and 2.950 seconds respectively, assuming aligned historian and source clocks. Sorting by arrival would place B first. Sorting by nominal occurrence would place A first. Neither sort removes the interval overlap. Preserve both events, flag the ambiguous association, and require independent process evidence before assigning a defect to a downstream bundle.

The worked example separates uncertainty from transport delay. GFT would require the event contract to name which uncertainty affects occurrence and which affects the observed delivery duration. An accurate historian clock cannot repair an unknown source clock. Replaying an event after an interruption should preserve its identity and occurrence evidence while adding delivery observations, allowing the receiving system to distinguish late evidence from current production.

GFT also recommends a durable acquisition identity separate from wall-clock time. A device instance, acquisition session and local sequence can support duplicate detection when clock corrections or replay produce repeated timestamps. The key design should state what happens after restart, replacement and counter wrap. A rejected duplicate should remain explainable by that key, rather than disappearing because its value happened to match a recent reading.

Shift assignment should be another declared interpretation. GFT recommends storing an unambiguous occurrence time and applying the approved site calendar when producing a shift view. An event arriving after the reporting cut-off should remain attached to its occurrence period when that time is trustworthy. The published report can then identify that it was revised after late evidence arrived. If occurrence time is unresolved, an arrival-based operational view may still be useful, provided it says what it measures and does not silently replace the production-period account.

4.5

Visibility events need business context beyond a sensor value

EPCIS is designed for sharing visibility-event data across applications and organisations. Its event dimensions include what, when, where, why and how. That provides a vocabulary for business context, including tracked objects and read points. It does not mean that a factory must convert every sampled temperature or motor-current waveform into a supply-chain event. [S15]

GFT proposes a controlled promotion from observation to operational event. A raw observation states what the capture system detected. A candidate association links that detection to an identified object and read point. An accepted operational event adds the process assertion that the factory is prepared to make. Retaining these stages prevents an uncertain association from acquiring authority merely because it has entered a shared event feed.

For an illustrative garment bundle, proximity to an inspection station could support an arrival observation. Completion should require the defined completion evidence, such as a station transaction and an accepted disposition linked to the bundle. Those are proposed rules, not findings about a particular mill. The interface should distinguish arrival, inspection started, inspection completed and release authorised when those distinctions affect decisions. Collapsing them into a generic seen status would prevent later scrutiny of the process assertion.

GFT recommends separating the read point from the business location. The read point describes where capture occurred; the business location describes the operational place associated with the event. A mobile reader needs a governed assignment to the relevant work area during the event period. Its network address alone should not be accepted as proof of where a bundle travelled. Unresolved location should be carried as an exception rather than silently copied from the last successful read.

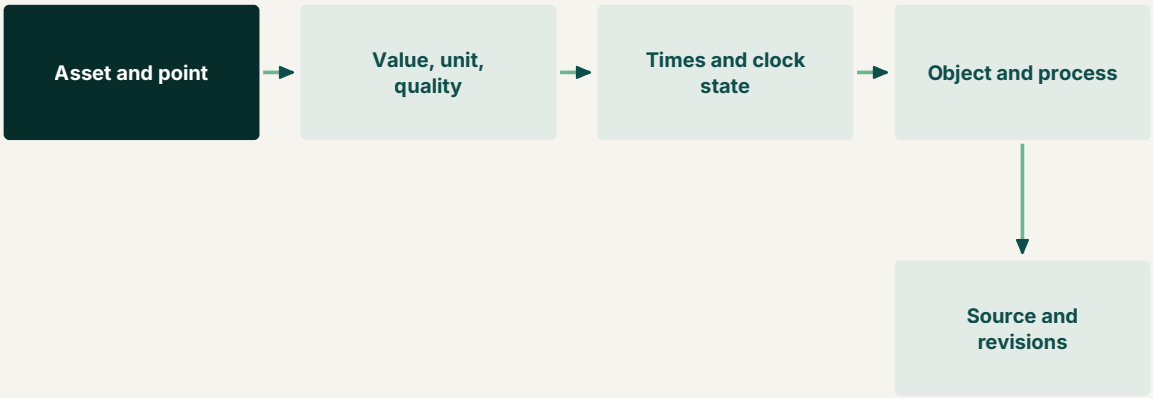
Repeated reads require an explicit policy. In an illustrative station design, a bundle remaining near the antenna should not accumulate completed operations with every observation. GFT would define an association window, expected process state, permitted transitions and evidence for leaving or re-entering the station. The parameters should be tested against the actual workflow. An anti-duplication rule chosen only to make a chart look plausible could hide legitimate rework.

The sensing review identifies time synchronisation and spectrum sharing as concerns for data integrity in RFID and wireless-sensor systems. That supports recording acquisition diagnostics and clock state alongside identity. It does not establish a success rate for the proposed textile installation. GFT would resolve read coverage, interference and ambiguous associations through local acceptance tests. [S31]

The illustrative envelope in **Exhibit 17**, p. 69 groups the evidence that should survive promotion. Its provenance field refers to GFT's proposed analytical lineage, rather than claiming that an asset-description or master-data standard defines every telemetry field. It also leaves room for an explicit unknown state, because a complete-looking record with invented context would be less useful than an honest exception.

EXHIBIT 17 / ILLUSTRATIVE EVIDENCE SEQUENCE

An event carries identity, time and the reason its meaning is accepted



ILLUSTRATIVE [S07] [S13] [S15] [S31] [S40] [S41] GFT-proposed event envelope informed by the cited sources. General analytical lineage is a design recommendation; ISO 8000-120 concerns characteristic master-data provenance.

GFT would keep the authority to change business context explicit. Production can confirm an operation; quality can confirm a disposition; an integration service can report that a mapping failed. These are different assertions. The event record should identify the approving role or service and preserve the original observation. Identity and time make the observation traceable, while declared authority makes its operational interpretation reviewable.

An illustrative split bundle needs an additional relationship: the parent bundle and each resulting child should remain connected to the recorded split operation. Otherwise, downstream observations may look like unexplained additions or duplicated work. GFT would require the process owner to define the conservation rule, including the quantity basis and any recorded loss. A merge should similarly preserve its contributing objects. The identifier does not prove physical conservation; it provides the address through which the reconciliation and its exceptions can be inspected.

4.6

Provenance travels with the value through every transformation

ISO 8000-120 addresses provenance for characteristic master data. Its published scope excludes configuration management and change control, and it does not supply a general telemetry-lineage contract. GFT therefore uses it to inform provenance discipline within that stated boundary. The analytical event lineage recommended below is a separate design proposal, not a claim of compliance with a standard whose scope is narrower. [S41]

GFT recommends recording the source observation reference, acquisition configuration revision, applicable mapping revision and transformation version whenever a value changes meaning. A converted temperature, a counter difference and a classified defect have different transformation dependencies. The lineage should make those dependencies inspectable without requiring the reader to infer them from a final dashboard label. Unchanged raw evidence remains the anchor for reviewing subsequent interpretations.

An illustrative counter calculation may start at 1,200 and finish at 1,248, producing a nominal difference of 48 counts. That arithmetic is meaningful only under the assumed conditions of a continuing counter instance, unchanged count definition and no intervening reset. If a reset is recorded, the simple subtraction no longer establishes output. GFT would retain the observations, mark the interval unresolved and apply a separately approved reconstruction rule if independent evidence permits one.

Reconstruction should declare its authority and uncertainty. GFT recommends distinguishing an observed count, an arithmetic result from accepted observations, and an estimate reconstructed from another record. A manual correction should identify the role, reason, affected period and superseded interpretation. It should not erase the source values or silently turn estimated production into measured production. The display can remain concise while exposing those distinctions through its evidence view.

For an illustrative defect event, the lineage may include the image reference, acquisition setup, model version, decision threshold and human review outcome. This does not require keeping every possible artefact indefinitely. It requires a retention decision matched to the use and a clear explanation when evidence expires. Storage architecture and retention implementation belong to the later historian chapter; this chapter defines what relationships that implementation must preserve.

GFT would express event quality through separate dimensions where the decision requires them: measurement validity, identity confidence, clock state, process-context status and delivery freshness. A reading may have valid measurement quality while its order assignment remains unresolved. Another may have complete identity but arrive too late for immediate intervention. Combining all those states into an unexplained good flag would conceal which use remains defensible.

The proposed acceptance exercise should follow a representative event from physical capture through an intentional transformation and into a decision. The reviewer should recover the original observation, identify the applicable asset assignment and explain every timestamp. The exercise should then repeat with a gateway restart, an assignment correction, a delayed replay and an invalid clock state. These are recommended tests, not reports of completed GFT operations.

For each test, the expected interpretation should be approved before execution. Evidence of correct rejection matters alongside evidence of correct acceptance. An unknown identity should remain unresolved. An ambiguous time association should remain ambiguous. A corrected mapping should produce an explainable revision. The factory should be able to reproduce both the original decision view and its later correction without claiming that late knowledge existed at the time.

GFT recommends making release authority part of the handover. The metric owner should accept the definition, the acquisition owner should accept capture and clock evidence, and the receiving application owner should accept interpretation and exception behaviour. Conditional acceptance should state which use is permitted and which evidence remains missing. A signal may be suitable for a maintenance trend while remaining unsuitable for product allocation. Recording that restricted permission gives the factory a useful intermediate state without presenting partial qualification as universal readiness.

The resulting handover is a compact contract: an identity dictionary, effective-dated mappings, declared timestamp semantics, uncertainty rules, event-promotion criteria and retained transformation references. It gives the following capture chapter a clear obligation: acquire observations without losing these distinctions. It also gives later storage and dashboard designs an acceptance test. A displayed value must be traceable to the production evidence that permits its interpretation.

A signal earns its meaning when its identity and time survive the journey.

CHAPTER SOURCE NOTES

S01 International Society of Automation. 2025-04-10.

S07 Industrial Digital Twin Association. 2024-06.

S40 Eclipse Foundation. 2022-11-16.

S15 GS1 AISBL. 2022-06.

S13 Internet Engineering Task Force and RFC Editor. 2010-06.

S41 International Organization for Standardization. 2016-10.

S0401 IEEE Standards Association. 2020-06-16.

S31 MDPI, Sensors. 2020-04-28.

GREEN FUTURE TECH

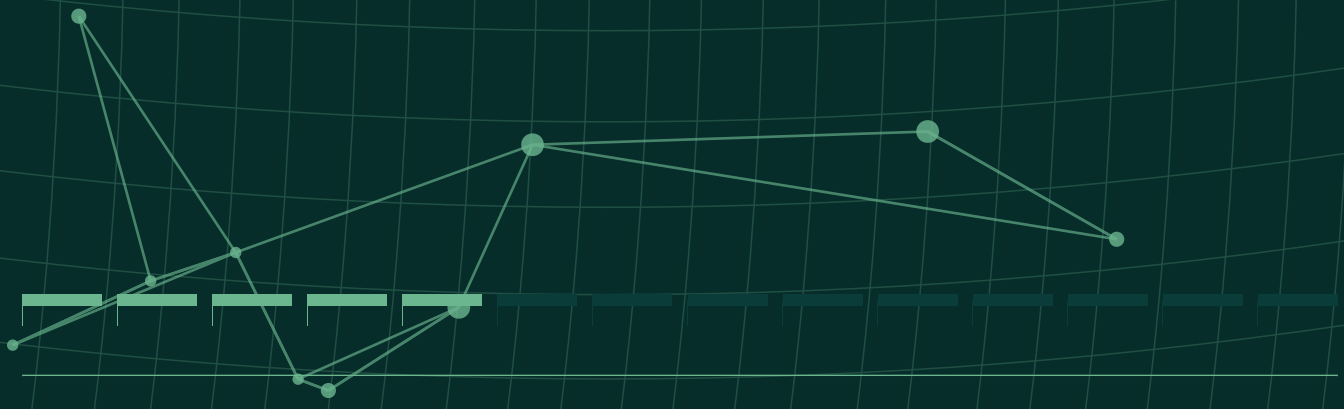
LIMITED / HONG KONG SAR

05

PART II: CAPTURE AT THE PRODUCTION EDGE / MACHINE INTERFACE

Brownfield capture must not disturb machine control

The safest digital retrofit begins by observing equipment through its documented interfaces and proving that acquisition does not alter control timing, protective functions or maintainability.



CHAPTER 05
DECISION TEST

The safest digital retrofit begins by observing equipment through its documented interfaces and proving that acquisition does not alter control timing, protective functions or maintainability.

A brownfield textile or garment factory cannot afford a digital retrofit that degrades the machines it is meant to measure. [S17] In GFT's reference architecture, every machine interface begins as a read-only, electrically isolated and timing-bounded observation path that leaves local closed-loop control, protective safety and maintenance access intact. [S17] [S23] Cameroon's inherited cotton-to-garment baseline, where approximately 95 percent of cotton fibre is exported raw and about 5 percent is processed locally against a national target of 50 percent local processing by 2030, means domestic mills and garment lines operate under tight capital and maintenance constraints. [R06] [R07] Third-party reporting from July 2023 on the legacy public manufacturer CICAM illustrates the stakes on existing shop floors: reported capacity utilisation stood at 33.7 percent in Garoua, 48.9 percent at CICAM 1 and 10.83 percent at CICAM 2, with underutilisation estimated to have cost about CFAF 2 billion in 2020. [R12] When a plant is working to recover availability across looms, dyeing vats, stenters, cutting tables and sewing lines, an unverified gateway that stalls a controller scan, corrupts a serial bus or bypasses an interlock creates immediate production loss. [S09] [S17] [S23] This chapter specifies the method for auditing brownfield equipment, polling serial and Ethernet controllers, separating cyclic sensor telemetry from maintenance dialogue, isolating deterministic motion networks, validating secondary radio and optical sensors, and gating every connection on security and safety evidence. [S09] [S10] [S11] [S19] [S31]

5.1

The machine audit precedes the gateway bill of materials

Procurement mistakes in factory digitalisation often begin when an integration team orders a uniform batch of edge gateways before opening an electrical cabinet on the floor. [S17] [S30] Survey evidence from West African manufacturing shows why an architecture cannot assume native digital readiness across an existing fleet: in the United Nations Industrial Development Organization 2020 report, almost 70 percent of surveyed manufacturing firms in Ghana fell into the analogue category as of the 2019 benchmark. [S35] That Ghanaian figure is an international comparator rather than a Cameroon survey, yet it reflects the physical reality that GFT's Cameroon programme addresses: a working plant combines electromechanical relays, cam-driven looms, older programmable logic controllers with serial ports, skid-mounted utilities, and manual benches with no controller at all. [S09] [S17] [S35]

5.2 A PHYSICAL AND LOGICAL INVENTORY REPLACES VENDOR ASSUMPTIONS

Before specifying any adapter or sensor kit, GFT requires a machine-by-machine audit documenting physical electrical interfaces and logical register semantics. [S01] [S09] [S17] Within the ISA-95 and IEC 62264 hierarchy, this audit maps each work unit at Level 0 and Level 1 to the operational events required at Level 3, ensuring engineers do not extract diagnostic tags that serve no decision. [S01] [S02] [S08] The walkdown inspects cabinet drawings, backplanes, terminal blocks, firmware revisions and maintenance procedures, recording nameplate and interface metadata in a structure compatible with the Asset Administration Shell metamodel. [S07] [S17]

EXHIBIT 18 / AUDIT FIELDS BY ASSET CLASS

The machine-capture sheet exposes unknowns before purchase

Asset class	Native controller and port	Primary cycle or flow source	Counter reset and state rule	Isolation and access constraint
Sewing bench	Clutch or servo motor; no open PLC port	Added optical stitch sensor or RFID bundle reader	Edge node counts pulses; distinguish power-on from active sewing	24 V DC isolated tap; radio spectrum check; read-only
Fabric spreading and cutting	OEM motion controller with Ethernet and serial port	PLC knife-active state and cut-cycle completion register	Verify whether cycle counter resets at shift change or power loss	Do not poll motion bus; read via secondary port only
Dyeing and finishing stenter	Modular PLC with RS-485 Modbus RTU and local HMI	Exhaust temperature, fabric speed and batch state registers	Map function codes 0x03 and 0x04; log exception responses	Galvanic isolator on RS-485; keep burner safety independent
Rotary air compressor	Packaged OEM controller with Modbus TCP gateway option	Package kW, discharge pressure and loaded-hours counter	Separate PowerOnDuration from loaded OperationDuration	Conduit firewall rule blocking all write function codes
Process water meter	Pulse output or IO-Link smart flow transducer	Volumetric pulse totaliser and cyclic flow rate	Non-resettable lifetime totaliser; retain edge store during outage	Separate cyclic process data from on-request parameter writes

ILLUSTRATIVE [S09] [S10] [S17] [S31] [S0501] Illustrative audit template for five representative textile and garment factory asset classes, derived from Modbus specification rules (S09), IO-Link interface separation (S10), NIST OT security guidance (S17), wireless sensor integrity constraints (S31) and OPC 40001-1 counter definitions (S0501). No vendor compatibility is assumed.

The audit worksheet in **Exhibit 18**, p. 74 shows why a bill of materials cannot be finalised from an asset register alone. [S09] [S17] Each row forces engineers to verify the controller and port, the signal representing a cycle or flow, the counter reset rule, the isolation boundary, and any vendor lock. [S09] [S10] [S17] [S0501] Because **Exhibit 18**, p. 74 is an illustrative template rather than a survey of a specific Cameroonian mill, it cannot establish which controller brands exist on a site; those facts exist only after the walkdown is signed off by the maintenance lead. [S17] [S30]

5.3 COUNTER SEMANTICS AND MACHINE STATES MUST BE VERIFIED AGAINST OPERATIONAL DEFINITIONS

A frequent failure in brownfield projects occurs when an integrator binds an undocumented integer register labelled "hours" or "pieces" directly to an availability or throughput tile. [S08] [S09] The OPC UA for Machinery companion specification (OPC 40001-1, Version 1.04.1) makes the required distinction explicit: `PowerOnDuration` records the duration an item has been powered with its processor running, `OperationDuration` records the duration it has been powered and performing an activity, and `OperationCycleCounter` counts transitions into active operation. [S0501] All three lifetime counters must only increase and must not reset on restart. [S0501]

In Cameroon, where the World Bank 2024 Enterprise Survey found that 93.3 percent of surveyed formal private firms experienced electrical outages with a typical-month frequency of 10.4 outages, counter behaviour across power loss is a primary design constraint. [S32] [S0501] Many older controllers store shift counters in volatile memory that resets when the breaker drops, roll a 16-bit integer over at 65,535 cycles, or let an operator zero the counter mid-shift. [S09] [S17] [S32] The audit tests every candidate counter through a simulated power cycle, shift reset and emergency stop, recording whether the edge node must compute positive deltas from a lifetime register or maintain its own non-volatile accumulator. [S29] [S32] [S0501]

Furthermore, a raw machine state at Level 1 is rarely identical to a manufacturing operations time element at Level 3. [S02] [S08] ISO 22400-2 defines manufacturing key performance indicators through explicit formulas, time and quantity elements, time behaviour and units. [S08] Annex C.2 of OPC 40001-1 shows that setting ISO 22400-2 time elements such as Actual Delay Time (`ADET`), Actual Down Time (`ADOT`) and Actual Preventive Maintenance Time (`APMT`) requires combining local `MachineryItemState` (`Executing`, `NotExecuting`, `OutOfService`, `NotAvailable`) and `MachineryOperationMode` (`Processing`, `Setup`, `Maintenance`) with a Higher Level System Condition State (`HLS-ConditionState`) indicating whether a production order is registered or whether the operation calendar schedules busy time. [S08] [S0502] A dyeing vat in `NotExecuting` state is experiencing unplanned delay (`ADET`) if an order is active, but sits in planned idle time if no order is scheduled. [S08] [S0502] The audit records which state bits the controller asserts truthfully on its own, and assigns order and calendar joins to the operations layer. [S01] [S02] [S0502]

> **GFT NOTE**
>
> **Require a signed machine-interface dossier before purchasing gateways.** In GFT's methodology for the Cameroon textile programme, no gateway or secondary sensor may be connected to an operating machine until maintenance, production and operational-technology security owners sign a one-page dossier. [S17] [S24] [S30] That dossier records the drawing reference, galvanic isolation rating, port, read-only function codes or input terminals, maximum polling rate, counter rollover and power-loss behaviour, warranty impact, and the physical plug that maintenance can disconnect in under 60 seconds illustrative limit to return the machine to standalone operation. [S09] [S17] [S22]

5.4

Modbus acquisition is polling with explicit exception handling

When brownfield machines, chillers, boilers, compressors or sub-meters expose a digital interface, it is frequently a serial RS-485 or Ethernet port speaking MODBUS. [S09] [S29] The MODBUS Application Protocol Specification (V1.1b3) defines MODBUS as a request/reply protocol offering services specified by function codes. [S09] Consequently, a MODBUS server inside a controller or meter does not spontaneously stream events when a seam finishes or a valve closes; the edge gateway acts as a client that transmits an explicit request frame, waits for the reply or exception frame, and attaches a receipt timestamp and quality state. [S09] [S13] [S40]

5.5 READ-ONLY FUNCTION CODES AND EXPLICIT REGISTER MAPS PREVENT UNINTENDED CONTROL WRITES

Because MODBUS function codes share the same framing for reading and writing, a misconfigured gateway or permissive firewall rule can turn a monitoring node into an unauthorised command source. [S09] [S17] Function codes `0x01` (Read Coils), `0x02` (Read Discrete Inputs), `0x03` (Read Holding Registers) and `0x04` (Read Input Registers) perform reads, whereas `0x05`, `0x06`, `0x0F` and `0x10` modify controller coils or registers. [S09] Writing to an incorrect holding register on a stenter or sizing machine can alter temperature setpoints, speed ratios or tension limits without an operator prompt. [S09] [S17] [S23]

GFT enforces a read-only capture rule at two layers. [S09] [S17] [S19] First, the gateway driver profile permits only read function codes (`0x01` through `0x04`), with write functions disabled. [S09] [S17] Second, where MODBUS TCP traverses a conduit under IEC 62443-3-2, the firewall inspects the MODBUS header and drops any frame carrying a write or reset function code. [S09] [S19] Even when a controller stores read-only values inside holding registers (`0x03`), the register map must document every contiguous block, data type, endianness, scale, unit and valid range so the gateway never polls unmapped addresses. [S09] [S17]

5.6 POLLING SCHEDULES MUST RESPECT SERIAL TIMING, SCAN CYCLES AND EXCEPTION CODES

On a two-wire RS-485 trunk connecting meters, temperature controllers or legacy programmable logic controllers, only one transaction can occupy the medium at any instant. [S09] Polling dozens of non-contiguous registers every 50 milliseconds illustrative interval across multiple slaves at 9,600 or 19,200 bits per second saturates the line, multiplies timeouts, and forces the legacy microprocessor to spend excessive interrupt time servicing serial requests rather than executing its control scan. [S09] [S17]

The MODBUS protocol specifies exception responses when a server cannot fulfil a request, returning an exception code such as `0x01` (Illegal Function), `0x02` (Illegal Data Address), `0x03` (Illegal Data Value), `0x04` (Server Device Failure), `0x05` (Acknowledge) or `0x06` (Server Device Busy). [S09] Firing immediate retries on a timeout or `0x06` Server Device Busy response worsens congestion while a controller executes a recipe transition. [S09] [S17] GFT requires every MODBUS adapter to classify exceptions explicitly: an Illegal Data Address (`0x02`) halts polling of that register and raises an alert; a Server Device Busy (`0x06`) or timeout triggers exponential back-off; and when retries fail, the edge node marks the metric as stale in its local store and Sparkplug payload rather than repeating the last valid reading. [S09] [S31] [S40]

METHOD

Method: Sizing a non-disruptive MODBUS RTU polling budget. 1. **System boundary:** One RS-485 multi-drop trunk connecting `N` slave devices to a read-only edge gateway port. [S09] [S29] 2. **Inputs:** Baud rate `B` (bits per second), framing `c` (typically 11 bits per byte), request bytes `L_req`, response bytes `L_resp`, turnaround time `T_turn`, and maximum bus duty cycle `U_max` (illustrative ceiling of 35 percent of line time). [S09] [S17] 3. **Single-transaction duration:** For transaction `i`, elapsed bus time is `T_i = ((L_req + L_resp) * c) / B + T_turn`. [S09] 4. **Duty-cycle check:** Across polled blocks with periods `P_i`, total utilisation is the sum of `T_i / P_i`, which must remain at or below `U_max` under normal polling and single-node timeout conditions. [S09] [S17]

WORKED EXAMPLE

Illustrative serial polling budget for a dyeing and utility trunk. Consider an illustrative RS-485 trunk operating at `B = 19,200` bits per second with 11-bit framing (`0.5729` milliseconds per byte), connecting 8 devices: 4 dyeing-vat controllers and 4 steam and power sub-meters, each polled via one contiguous `0x03` request for 20 registers (40 data bytes). [S09] [S29]

In this illustrative calculation, each request is `L_req = 8` bytes (`4.58` milliseconds) and each response is `L_resp = 45` bytes (`25.78` milliseconds). [S09] With a turnaround time of `T_turn = 25.0` milliseconds, each transaction occupies `T_i = 55.36` milliseconds, and one complete sweep across 8 devices requires `442.88` milliseconds. [S09]

Polling all 8 devices every `1.0` second (`P_i = 1,000` milliseconds) yields `44.3` percent baseline utilisation, and one unplugged meter with a 500-millisecond illustrative timeout pushes sweep duration to `942.88` milliseconds (`94.3` percent utilisation). [S09] [S17] Matching polling rate to the decision horizon under ISO 22400-2, GFT polls the 4 vat controllers every `2.0` seconds (`11.07` percent utilisation) and the 4 sub-meters every `10.0` seconds (`2.21` percent utilisation), reducing steady-state utilisation to an illustrative `13.28` percent. [S08] [S09] [S17] [S29]

5.7

IO-Link separates cyclic process data from maintenance dialogue

Where a brownfield machine lacks a controller register for a required variable, or where tapping into an obsolete programmable logic controller carries unacceptable risk, the cleanest option is to mount secondary smart sensors on the frame and wire them to an independent edge master. [S10] [S17] Common retrofit points include photoelectric sensors counting fabric folds or cut bundles, pressure transducers tracking dye-bath or rinse-tank levels, optical tachometers measuring roller speed, and thermal flow sensors monitoring compressed-air or water use. [S27] [S29] When those sensors use analogue `4` to `20` milliampere loops or discrete `24` volt switching wires, the factory obtains a raw signal but cannot remotely verify whether a lens is fouled with cotton lint or whether someone altered the switching threshold on the housing. [S10] [S26]

5.8 CYCLIC PROCESS EXCHANGE OPERATES WITHOUT INTERFERENCE FROM ON-REQUEST TRAFFIC

The IO-Link Interface and System Specification (Version 1.1.4) solves this visibility gap through an explicit division in its communication stack: the Process Data Exchange unit cyclically transmits and receives Process Data without interference from the On-request Data (parameters, commands and Events). [S10] Over a standard three-wire unshielded sensor cable connected point-to-point between the sensor and an IO-Link master port, each cycle allocates a deterministic sub-slot for cyclic process data and a separate, non-interfering fragment for acyclic on-request service messages. [S10]

This guarantee has two practical consequences for GFT's architecture. [S10] [S17] First, production telemetry (such as fabric-presence state, nip-roll pressure or water level) arrives at the IO-Link master on a predictable millisecond cycle regardless of whether an engineer or diagnostic service is simultaneously querying device health or contamination status. [S10] [S27] Second, because IO-Link is a point-to-point link between each device and its master port rather than a shared serial bus, a severed cable or failed transducer on one stitch-station sensor or tank probe isolates the fault to that single port without corrupting neighbouring sensors. [S10] [S17]

5.9 PARAMETER MANAGEMENT AND DIAGNOSTICS BELONG TO GOVERNED CHANGE CONTROL

While IO-Link's on-request channel exposes device identification, teaching thresholds, operating hours and lens-fouling events, that same acyclic channel also supports writing new parameters and commands to the sensor. [S10] Unrestricted acyclic parameter writes from the operations network could allow an unauthorised script or accidental push to invert switching logic, change a hysteresis band or trigger a false calibration command during production. [S10] [S17] [S24]

GFT binds IO-Link's separation to the information-security management discipline of ISO/IEC 27001 and the operational-technology guidance of NIST SP 800-82 Rev. 3. [S10] [S17] [S24] Cyclic process data and read-only diagnostic events flow continuously from the IO-Link master's secondary Ethernet interface into the edge acquisition service, carrying sensor identity and port quality flags into the event envelope. [S10] [S40] Conversely, on-request parameter writes and device-replacement configuration downloads are restricted to authenticated maintenance sessions under change control, logging operator identity, old value, new value and timestamp. [S10] [S20] [S24] Factory acceptance tests verify both sides of the contract: intensive acyclic diagnostic reads during a simulated run must cause zero dropped cyclic samples, and an unauthenticated parameter write from the monitoring conduit must be rejected. [S10] [S17] [S19]

5.10

Deterministic networks are not interchangeable with ordinary Ethernet

Modernised sections of a brownfield mill, such as automated bale-opening lines, multi-drive spinning frames, air-jet looms, rotary printers and computer-controlled cutters, frequently use industrial Ethernet fieldbuses inside the machine enclosure. [S11] [S17] Because these networks use familiar RJ45 or M12 connectors and Cat-

egory 5e or Category 6 cabling, IT-led teams sometimes assume they can plug a gateway directly into a spare port on the machine's internal switch to scrape tags or stream video. [S11] [S17] [S19] That assumption confuses physical connector compatibility with real-time network behaviour and violates both control determinism and zone segmentation. [S11] [S19]

EXHIBIT 19 / SYNCHRONISATION AND JITTER CAPABILITIES BY STANDARD

Timing claims retain their technical conditions

Standard or technology	Published timing capability	Required technical conditions in source	Role in GFT brownfield architecture
PROFINET (Isochronous Real-Time)	Jitter of less than 1 microsecond	High-precision closed-loop control tasks with deterministic and isochronous transmission of time-critical process data	Keep inside the certified machine control network; never bridge raw analytics or IT polling into the IRT domain
IEEE 1588-2019 (Precision Time Protocol)	Sub-microsecond synchronisation accuracy and precision	Networked measurement and control systems supporting the protocol with minimal network and local computing resources	Synchronise multi-axis or cross-device edge capture only where sub-millisecond event correlation is required
NTPv4 (IETF RFC 5905)	Potential accuracy to the tens of microseconds	Modern workstations and fast local area networks using NTPv4 mitigation and discipline algorithms	Default clock alignment for edge gateways, local historians, MQTT brokers and operational dashboards

SOURCED [S11] [S12] [S13] Sourced comparison of published timing capabilities from the PROFINET System Description (S11), IEEE 1588-2019 (S12) and IETF RFC 5905 for NTPv4 (S13). Each figure retains the specific technical conditions stated by its standard.

5.11 ISOCHRONOUS CLOSED-LOOP CONTROL REQUIRES STRICT ISOLATION FROM BURSTY TRAFFIC

The comparison in **Exhibit 19**, p. 79 explains why GFT treats deterministic machine networks as protected control domains. [S11] [S12] [S13] According to the PROFIBUS Nutzerorganisation system description for PROFINET, deterministic and isochronous transmission of time-critical process data with a jitter of less than 1 microsecond is possible for high-precision closed-loop control tasks. [S11] Achieving sub-microsecond jitter across synchronised drives requires reserved bandwidth and switch scheduling. [S11] Similarly, IEEE 1588-2019 specifies that Precision Time Protocol (PTP) supports sub-microsecond synchronisation accuracy and precision with minimal network and computing resources, whereas IETF RFC 5905 states that Network Time Protocol Version 4 (NTPv4) extends potential accuracy to the tens of microseconds with modern workstations and fast local area networks. [S12] [S13]

The sourced figures in **Exhibit 19**, p. 79 establish upper bounds under stated conditions, not automatic guarantees on an arbitrary LAN. [S11] [S12] [S13] Injecting bursty TCP/IP traffic, broadcast discovery packets, large OPC UA browse requests or video streams into an unsegmented machine network can exhaust switch buffers or disturb

cyclic communication. [S04] [S11] [S17] If a sizing machine or cutter suffers a communication watchdog timeout because an external gateway flooded its network card, the machine trips to an immediate stop. [S11] [S17] [S23]

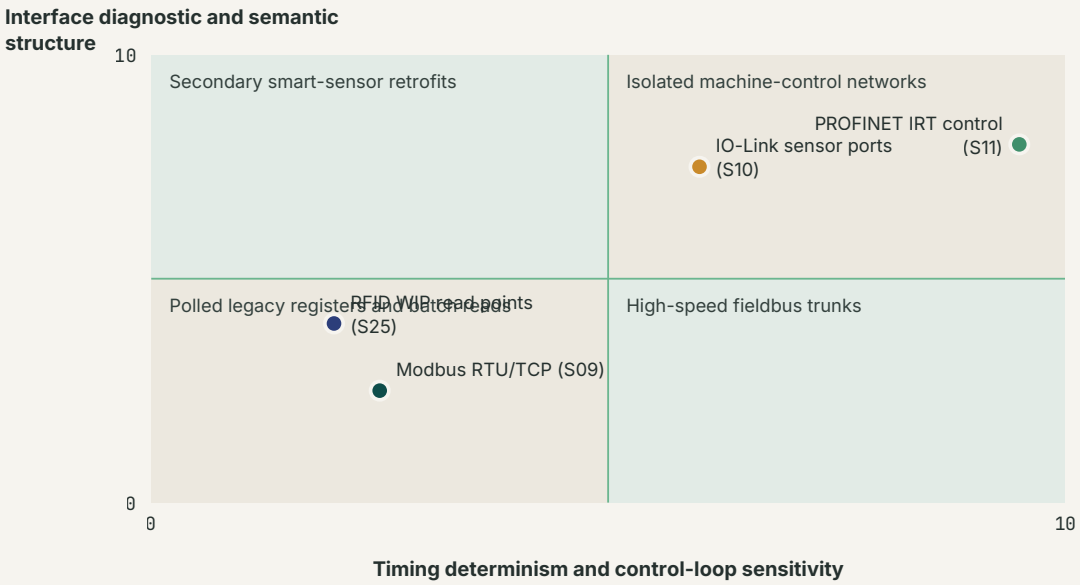
5.12 COUPLERS, DUAL-HOMED CONTROLLERS AND UNIDIRECTIONAL GATEWAYS PROTECT THE CONTROL BOUNDARY

To capture operational evidence from deterministic networks without disturbing closed-loop timing, GFT enforces a boundary between the internal fieldbus and the edge acquisition zone. [S11] [S17] [S19] First, where the controller provides a separate communication processor or secondary Ethernet port isolated from the fieldbus backplane, it exposes a read-only OPC UA server or data-block mirror on that secondary port. [S04] [S14] [S17] Because OPC UA is layered to isolate its core design from computing technology and network transport, the machine can expose a structured object model on the secondary port while keeping drive exchange private on the machine network. [S04] [S11] [S0501]

Second, where a PROFINET network lacks a spare port, a certified network coupler copies a bounded array of cyclic and acyclic bytes between the machine network and the edge network without bridging Ethernet frames or broadcasts. [S11] [S19] Third, where hardware-enforced one-way extraction is warranted, NIST SP 800-82 Rev. 3 defines a unidirectional gateway as hardware and software in which the hardware permits data to flow from one network to another but is physically unable to send any information back to the source network, while software replicates databases and emulates protocol servers. [S17] [S0503] Whichever boundary is selected, W3C's Web of Things Architecture 1.1 captures the pattern: an industrial edge device collects selected controller data and makes it available to downstream monitoring or preventive-maintenance services. [S14]

EXHIBIT 20 / DERIVED ENGINEERING PLACEMENT (1 TO 10 SCALE)

Protocol selection follows behaviour, not fashion



DERIVED [S09] [S10] [S11] [S25] Derived placement from S09 (Modbus request/reply polling), S10 (IO-Link cyclic and on-request separation), S11 (PROFINET deterministic control) and S25 (RFID denim WIP tracking). Horizontal axis scales timing determinism and control sensitivity (1 = slow polling or asynchronous read, 10 = sub-microsecond isochronous control). Vertical axis scales interface diagnostic and semantic structure (1 = raw address or tag ID, 10 = structured cyclic data with separate diagnostics).

The derived placement in **Exhibit 20**, p. 81 synthesises why GFT rejects single-protocol mandates on brownfield floors. [S09] [S10] [S11] [S25] MODBUS (S09) sits in the lower-left quadrant because it relies on client polling of raw register addresses; RFID capture (S25) operates as an asynchronous identification event at discrete read points; IO-Link (S10) provides cyclic process data alongside separated on-request diagnostics; and PROFINET Isochronous Real-Time (S11) occupies the high-determinism upper-right quadrant for closed-loop motion control. [S09] [S10] [S11] [S25] Because the coordinates on the 0 to 10 axes in **Exhibit 20**, p. 81 are GFT's derived ordinal translation of those four sources rather than laboratory measurements, the chart guides role assignment and is not a throughput benchmark. [S09] [S10] [S11] [S25]

5.13

Added sensors need mechanical, radio and process validation

Across spinning, weaving, wet processing and garment assembly, many legacy stations have no controller port from which to extract counts or batch identity. [S25] [S35] On these stations, capture requires adding secondary hardware: radio-frequency identification (RFID) readers at bundle transfer points, optical or inductive sensors on shafts, current transformers on motor leads, or water-level probes on vats. [S25] [S27] [S29] Adding an external sensor avoids touching legacy firmware, but introduces physical, radio and process-interpretation failure modes that must be validated before use. [S25] [S31]

5.14 RADIO SPECTRUM SHARING AND TIME SYNCHRONISATION GOVERN WIRELESS INTEGRITY

An industrial sewing floor or dyehouse is a harsh electromagnetic environment filled with variable-frequency drives, contactor arcs, humid bales, metal tables and steel vats. [S25] [S31] A review of IoT sensing applications and challenges using RFID and wireless sensor networks highlights that time synchronisation between nodes and spectrum-sharing techniques are essential to protect data integrity and avoid uncertainty in sensed data without exhausting radio resources. [S31]

Deploying ultra-high-frequency (UHF) RFID readers, Wi-Fi tablets and wireless sensor nodes across the same bay without a radio plan causes co-channel interference, multipath reflections and reader collisions. [S31] Furthermore, if wireless nodes or RFID gateways operate with unsynchronised clocks that drift by seconds or minutes, the event sequence can report that a bundle arrived at a downstream sewing station before departing the cutting table. [S13] [S15] [S31] GFT requires every wireless or RFID retrofit to enforce clock synchronisation via NTPv4 (or IEEE 1588 PTP where tighter correlation is required) and to publish diagnostic metrics for signal strength, retries, channel occupancy and synchronisation state alongside sensed values. [S12] [S13] [S31] [S40]

5.15 PHYSICAL READ-POINT ENGINEERING SEPARATES PROXIMITY FROM COMPLETED WORK

In garment and denim manufacturing, RFID is widely proposed to bridge the visibility gap between cutting rooms and sewing benches. [S25] [S28] Published research on an RFID-based tracking system for denim production found that radio-frequency identification could support the tracking and tracing of work in process in denim processes, while an RFID-enabled manufacturing execution study describes systematic shop-floor deployment to track objects and collect real-time production data. [S25] [S28] Those studies confirm that RFID is a viable capture mechanism for textile work in process, without implying that bolting an antenna to a bench yields accurate completion counts automatically. [S25] [S28]

Three ambiguities must be engineered out of every RFID or optical read point before its output populates an EPCIS visibility event or an ISO 22400 throughput calculation. [S08] [S15] [S25] First is **stray and cross-lane reading**: a UHF reader with excessive power at Workstation A will read bundles stacked at Workstation B two metres away, creating phantom arrivals. [S25] [S31] Second is **metal and liquid detuning**: wet denim, zippers, rivets, steel presser feet and metal trolleys detune antennas and shadow tags unless tag placement, polarisation and read-zone geometry are tested with the actual product mix. [S25] [S31] Third is **semantic conflation between presence and completion**: an RFID tag entering a reader's field proves only that the bundle was within radio range at that timestamp, not that the operator completed the sewing operation or passed quality checks. [S15] [S25]

To resolve these ambiguities, GFT separates raw sensor acquisition from business-event assertion. [S15] [S31] At the physical layer, engineers tune reader power, use near-field or shielded antennas for bench reads, and add chutes or operator trigger buttons so a read occurs only at a deliberate handoff. [S25] [S31] At the edge software layer, a state-machine filter deduplicates repeated reads within a dwell window, flags out-of-route appearances as exceptions, and emits an EPCIS event record only when the read satisfies the station's process rule. [S15] [S25] [S28]

WARNING

Unresolved site constraint: Radio-frequency shadowing and lint accumulation on retrofitted sensors. Neither standard specifications nor literature case studies can predict the electromagnetic reflection profile of a specific Cameroonian weaving shed or garment hall, nor the rate at which cotton fly and size dust obscure optical lenses on high-speed frames. [S25] [S26] [S31] Resolving this constraint requires on-site validation during Phase 1 commissioning: conducting a spectrum and read-zone walkdown under full production load, measuring false-positive and missed-read rates against a manual tally baseline over a complete shift, and establishing a preventive cleaning schedule before any dashboard KPI depends on the feed. [S26] [S30] [S31]

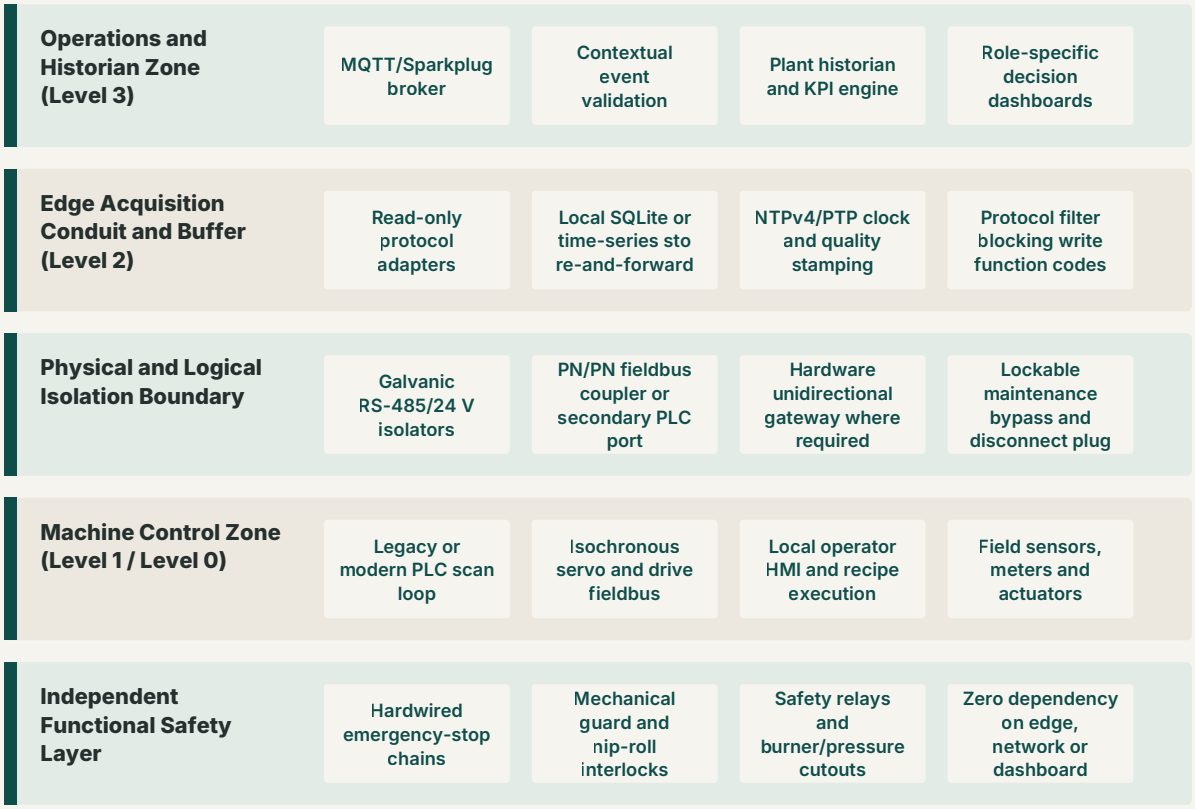
5.16

Safety and security boundaries constrain every retrofit

A brownfield capture project changes the connectivity profile of machines designed decades before industrial cybersecurity standards existed. [S17] [S19] Connecting twenty legacy programmable logic controllers and utility skids to a flat Ethernet switch without segmentation creates a lateral movement path whereby a compromised laptop, infected USB drive or misconfigured vendor modem can reach every controller on the floor. [S17] [S19] [S20] Equally, careless wiring inside a cabinet can compromise hardwired safety circuits protecting operators from mechanical, thermal and electrical hazards. [S17] [S23]

EXHIBIT 21 / ARCHITECTURAL ZONES AND CONDUITS FROM MACHINE TO OPERATIONS

The machine boundary separates protective control from evidence export



ILLUSTRATIVE [S01] [S11] [S19] [S23] [S0503] Illustrative GFT brownfield boundary architecture informed by ISA-95 asset hierarchy (S01), PROFINET deterministic control scope (S11), IEC 62443-3-2 zones and conduits (S19), IEC functional safety independence (S23) and NIST unidirectional gateway definitions (S0503).

5.17 ZONES AND CONDUITS PREVENT LATERAL EXPOSURE AND UNAUTHORISED WRITE PATHS

The boundary model in **Exhibit 21**, p. 84 applies IEC 62443-3-2, which makes security zones and conduits a key concept in system design. [S19] Rather than allowing gateways to straddle machine and enterprise networks, **Exhibit 21**, p. 84 places each machine cell inside a Machine Control Zone, separates the Edge Acquisition Zone across an isolation boundary, and routes outbound telemetry to the Operations Zone through an authenticated conduit. [S01] [S19] [S40] Because **Exhibit 21**, p. 84 is an illustrative reference architecture, it defines functional separation that must be instantiated onto specific VLANs, firewall rules, couplers and terminal blocks during site engineering. [S17] [S19]

Remote maintenance paths receive strict control inside this model. [S19] [S20] Under the Cybersecurity and Infrastructure Security Agency (CISA) Cybersecurity Performance Goals 2.0, multi-factor authentication (MFA) is enabled on all remotely accessible operational-technology accounts and systems, including vendor and maintenance accounts, workstations and HMIs when available. [S20] GFT prohibits unmonitored cellular modems or internet bridges inside machine cabinets; authorised

remote maintenance must terminate in an industrial demilitarised zone with multi-factor authentication, time-bounded approval and session logging, and must never use the capture gateway as a backdoor into the controller. [S19] [S20] [S24]

Before any retrofit wire is landed or secondary port enabled, engineers must take a verified backup of existing PLC logic, motion parameters, HMI files and register maps, storing that backup in an isolated repository disconnected from network paths that could spread malware. [S21] NIST SP 800-34 Rev. 1 states that testing validates recovery capabilities, training prepares personnel for activation, and exercising the plan identifies gaps. [S22] Recovery validation means proving on the machine that technicians can unplug the gateway, restore baseline controller logic if needed, and verify standalone operation within the approved recovery time objective. [S17] [S21] [S22]

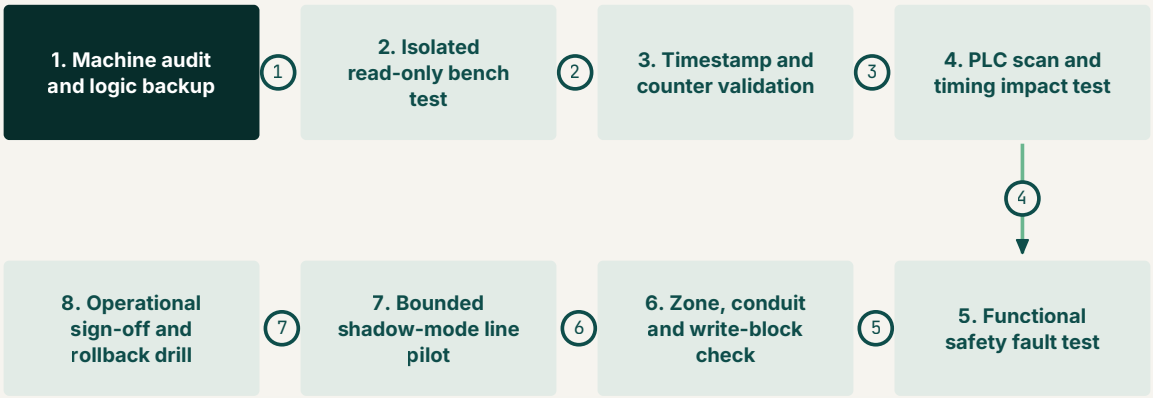
5.18 FUNCTIONAL SAFETY REMAINS PHYSICALLY AND LOGICALLY INDEPENDENT OF DATA CAPTURE

The bottom layer of **Exhibit 21**, p. 84 enforces the core rule of this chapter: functional safety never depends on the edge gateway, network conduit or dashboard. [S19] [S23] As the International Electrotechnical Commission states, within the overall safety of a system, functional safety refers to the aspect that provides automatic protection in response to errors. [S23] On textile machinery, protective functions include emergency-stop pushbuttons, pull-cord switches along spinning and weaving alleys, two-hand knife controls, light curtains on presses and cutters, interlocking guard doors over carding cylinders and nip rolls, and over-temperature or over-pressure cutouts on steam vessels and dryers. [S17] [S23]

GFT enforces two prohibitions regarding protective safety during brownfield instrumentation. [S17] [S23] First, no edge sensor, power supply or signal tap may be wired into a primary safety-relay circuit where it could bridge a contact, cause a ground fault or impede an emergency stop; safety state is read exclusively from an auxiliary dry contact on the safety relay or a read-only PLC bit after the hardwired trip executes. [S08] [S17] [S23] Second, no analytical rule, camera model or dashboard alert running on an edge node may ever be counted as a primary protective safety control. [S17] [S23]

EXHIBIT 22 /

The retrofit gate proves observation before integration



1. Dossier and isolated backup complete
2. Galvanic and read-only path verified
3. Power-loss and rollover rules proven
4. Zero control-loop degradation
5. E-stop and interlocks independent
6. Unauthorised writes and paths blocked
7. KPIs reconcile with manual baseline

ILLUSTRATIVE [S17] [S19] [S22] [S23] [S30] Illustrative eight-step engineering gate sequence for brownfield machine instrumentation, derived from NIST OT performance and safety requirements (S17), IEC 62443-3-2 zone-and-conduit design (S19), NIST contingency testing and exercises (S22), IEC functional safety independence (S23) and textile SME pilot validation guidance (S30).

5.19 THE EIGHT-STAGE RETROFIT GATE GOVERNS MACHINE COMMISSIONING

To turn these constraints into a repeatable routine, GFT gates every brownfield connection through the eight stages in **Exhibit 22**, p. 86. [S17] [S19] [S22] [S23] [S30] Because **Exhibit 22**, p. 86 is an illustrative workflow synthesised from standards and pilot guidance, it prescribes mandatory acceptance evidence rather than fixed calendar durations. [S17] [S30]

1. **Machine audit and logic backup:** Complete **Exhibit 18**, p. 74, verify isolation ratings and store an isolated backup of controller logic. [S17] [S21]
2. **Isolated read-only bench test:** Verify offline that only read function codes (0x01 to 0x04 in MODBUS) or cyclic inputs are active. [S09] [S10] [S17]
3. **Timestamp and counter validation:** Confirm NTPv4 or IEEE 1588 PTP synchronisation, counter continuity across power loss, and ISO 22400-2 state mapping. [S08] [S12] [S13] [S0501] [S0502]
4. **PLC scan and timing impact test:** Monitor controller scan time, fieldbus diagnostics and exception counts at peak load to prove zero control degradation. [S09] [S11] [S17]

5. **Functional safety fault test:** Unplug gateway power and sever telemetry cables in test mode to verify that emergency stops and interlocks act identically. [S17] [S23]
6. **Zone, conduit and write-block check:** Verify IEC 62443-3-2 zone placement, write-blocking firewall rules, and multi-factor authentication on remote maintenance. [S19] [S20] [S24]
7. **Bounded shadow-mode line pilot:** Run in shadow mode through power and WAN outages until captured counts reconcile with manual logs. [S29] [S30] [S32] [S40]
8. **Operational sign-off and rollback drill:** Sign acceptance only after technicians execute a timed rollback drill confirming standalone machine operation. [S17] [S22] [S30]

By enforcing this eight-stage gate across sewing benches, looms, dyeing vats, cutters and meters, GFT turns brownfield heterogeneity into a disciplined engineering sequence. [S30] [S35] Once physical signals are captured safely at the edge, the next chapter examines how local inference keeps quality and resource decisions on the line when remote links fail. [S14] [S26] [S29]

An edge retrofit earns its place on the shop floor only when the machine runs identically if the gateway is unplugged.

CHAPTER SOURCE NOTES

S17 National Institute of Standards and Technology. 2023-09.

S23 International Electrotechnical Commission. 2022-04-28.

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.

S09 Modbus Organization. 2012-04-26.

S10 IO-Link Community. 2024-06.

S11 PROFIBUS Nutzerorganisation. 2018-11.

S19 International Electrotechnical Commission. 2020-06.

S31 MDPI, Sensors. 2020-04-28.

S30 MDPI, Sustainability. 2023-08-19.

S35 United Nations Industrial Development Organization. 2019.
S01 International Society of Automation. 2025-04-10.
S02 International Electrotechnical Commission. 2013-05-22.
S08 International Organization for Standardization. 2014-01.
S07 Industrial Digital Twin Association. 2024-06.
S0501 OPC Foundation and VDMA. 2026-01-01.
S32 World Bank Enterprise Surveys. 2025.
S29 MDPI, IoT. 2024-09-28.
S0502 OPC Foundation and VDMA. 2026-01-01.
S24 International Organization for Standardization and International Electrotechnical Commission. 2022-10.
S22 National Institute of Standards and Technology. 2010-11-11.
S13 Internet Engineering Task Force and RFC Editor. 2010-06.
S40 Eclipse Foundation. 2022-11-16.
S27 MDPI, Applied Sciences. 2024-06-21.
S26 MDPI, Electronics. 2025-09-18.
S20 Cybersecurity and Infrastructure Security Agency. 2025-12.
S12 IEEE Standards Association. 2020-06-16.
S04 OPC Foundation. 2025-10-22.
S14 World Wide Web Consortium. 2023-12-05.
S0503 National Institute of Standards and Technology. 2023-09.
S25 Springer Nature, International Journal of Advanced Manufacturing Technology. 2016-09-05.
S15 GS1 AISBL. 2022-06.
S28 Elsevier, Robotics and Computer-Integrated Manufacturing. 2013-04.
S21 CISA and joint international cybersecurity authorities. 2022-05-09.

06

PART II: CAPTURE AT THE PRODUCTION EDGE / LOCAL ANALYTICS

Edge inference keeps quality decisions on the line

Local inspection earns its place when a valid image becomes a traceable quality decision before the fabric leaves the decision point. Acquisition, persistence and an approved fallback matter as much as inference speed.

CHAPTER 06
DECISION TEST

Local inspection earns its place when a valid image becomes a traceable quality decision before the fabric leaves the decision point. Acquisition, persistence and an approved fallback matter as much as inference speed.

Keep the inspection decision beside the fabric when waiting for a remote service would make that decision late. GFT proposes a local chain from controlled acquisition to inference, durable evidence and an authorised quality action. The remote service receives selected results and supports later analysis. Locality alone does not establish accuracy, continuity or safety. Each depends on a tested operating boundary. This chapter specifies that chain as a reference design, not an account of an installed GFT system. W3C describes selective industrial edge collection, while ETSI provides an architecture for applications near the network edge. [S14] [S16]

6.1

The camera cell is an engineered measurement station

Start by defining what the station must distinguish. GFT recommends a defect catalogue with acceptance criteria, a physical location convention and examples agreed by the quality owner. Describe the product families and surfaces within scope. Separate a visible indication from a confirmed defect, and separate either from the decision to reject material. The acquisition literature identifies lighting variation and background noise as threats to detection. A promising model result therefore cannot stand in for a camera-cell acceptance record. [S26]

GFT's proposed commissioning record should bind the camera position, lens, focus, illumination arrangement, exposure, trigger source, fabric presentation and calibration reference to an approved configuration. Specify how the fabric is supported and where its position is measured. Make the reference image reproducible after cleaning or maintenance. Record changes to the light source and enclosure as configuration changes, because the station must be able to explain a changed image before it attributes a changed result to the fabric. These are design recommendations derived from the acquisition limitations, not performance guarantees. [S26]

Commissioning should prove the station's physical coverage before assessing its model. GFT recommends inspecting the edges of the intended field, transitions between camera views, folds and areas obscured by guides. Test the permissible fabric movement and presentation conditions against the defect definition. A station can report a technically valid inference while failing to see the material that matters. Record unseen or invalid areas explicitly as uninspected coverage. Do not let a camera heartbeat stand for complete fabric inspection.

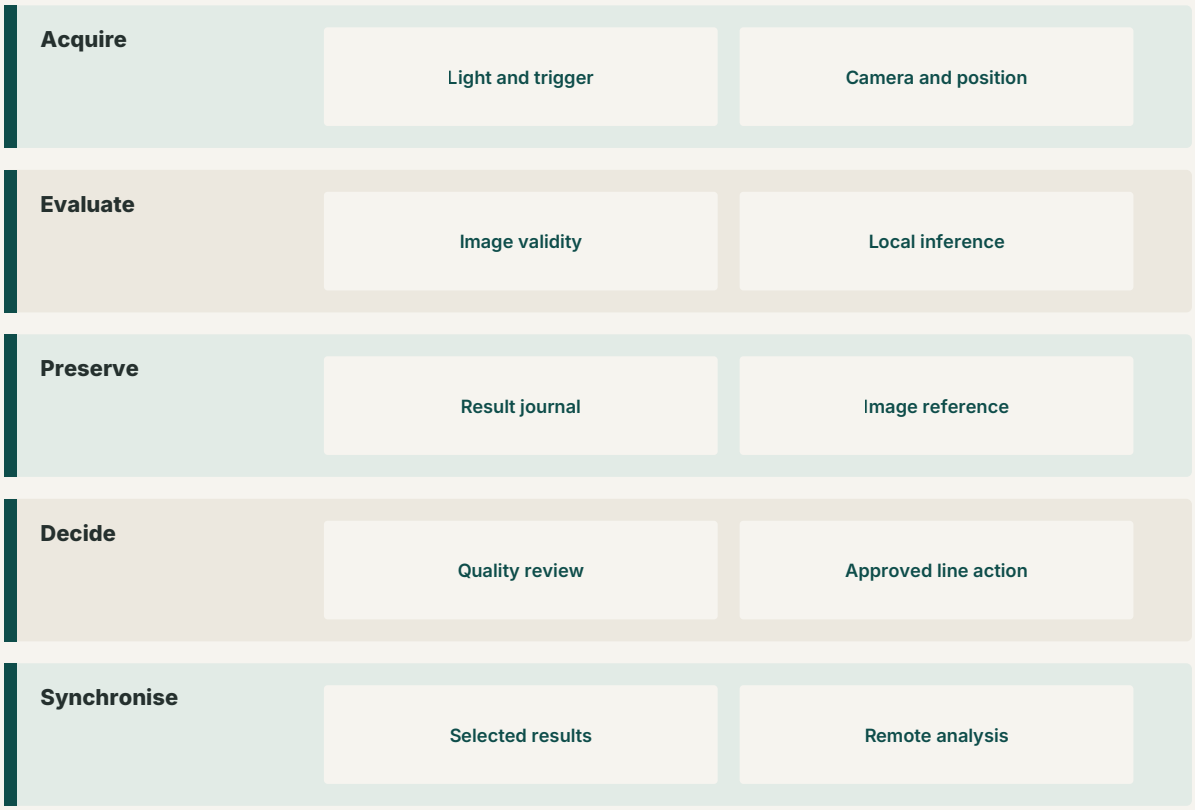
Site validation should then use independently adjudicated material from the intended operating conditions. GFT recommends grouping evaluation evidence by fabric construction, colour, finish and relevant operating condition. Keep samples from the same physical defect together when separating development and evaluation material. Otherwise adjacent views may make the evaluation look more independent than it is. Document how a disputed reference label is resolved. The research scope includes plain, patternless cotton and linen; it does not establish transferable performance across every textile construction or camera arrangement. [S26]

GFT recommends an acquisition-validity gate before inference eligibility. The gate should reject an obscured field, lost trigger association, unacceptable exposure or an unapproved setup. It should return an explicit invalid state with a reason, rather than a confident defect classification on an unusable image. This recommendation makes the distinction operational: a classification concerns visible material; invalid acquisition concerns whether the station could observe that material adequately. Both conditions require a responsible response, but they support different conclusions.

The illustrative architecture in **Exhibit 23**, p. 91 separates the observation, computation and decision responsibilities. Its layers describe an allocation of duties, not a purchased equipment list. The protective machine system remains outside the analytics chain. NIST's guidance considers OT performance, reliability and safety together; GFT uses that discipline to require a reviewed boundary between quality analytics and machine protection. [S17]

EXHIBIT 23 / ILLUSTRATIVE RESPONSIBILITY LAYERS

The cell separates observation from quality authority



ILLUSTRATIVE [S14] [S16] [S26] [S17] GFT proposed cell informed by the sources. Layers are responsibilities, not certified network zones. Protective machine functions are independent and outside this analytics path.

GFT recommends accepting physical coverage, acquisition validity and classification performance as separate findings. A failure in any finding should limit the permitted operating scope. Replacing a model cannot cure an unlit edge of the fabric, and changing a lamp should not silently preserve an old acceptance statement. The cell owner

should retain the evidence and state which conditions require revalidation. This makes commissioning a repeatable measurement activity rather than a demonstration arranged around favourable samples.

For maintenance, GFT proposes a station handback sheet that begins with the physical intervention. State whether the camera mount moved, the lens was replaced, the lighting was changed or only the enclosure was cleaned. Compare the restored reference view with the approved setup and record the decision to resume. A maintenance ticket marked complete should not automatically clear an inspection restriction. Give the person responsible for quality a clear route to accept the restored scope or request further checks.

6.2

Inference stays local when the decision cannot wait

Place computation according to the consequence of delay. ETSI describes software applications running on virtualisation infrastructure near the network edge. W3C's industrial example separates controller operation from selected data collection for monitoring or analysis. Neither source proves a latency for a particular inspection station. GFT's placement rule is therefore conditional: keep the necessary decision path local when the measured remote path cannot meet the decision deadline and interruption tolerance. [S16] [S14]

Define that deadline at the physical action point. GFT recommends measuring the available travel distance from observation to diversion, marking, hold or operator intervention. Establish the speed envelope and the time needed by the receiving mechanism. Subtract acquisition, transfer within the cell, preprocessing, inference, evidence recording, decision handling and action time. Reserve explicit uncertainty for speed variation, scheduling and triggering. The resulting budget is a commissioning hypothesis until tested on the complete installed path.

WORKED EXAMPLE: THE PHYSICAL DEADLINE INCLUDES THE ACTION

In an illustrative inspection cell, assume the action point is 0.8 metres after the camera and fabric moves at 0.5 metres per second. Travel time is 0.8 divided by 0.5, or 1.6 seconds. Assume capture and preprocessing consume 0.10 seconds, inference 0.12, durable result recording 0.03, decision routing 0.05 and the action mechanism 0.30. Their sum is 0.60 seconds. Reserve an illustrative 0.40 seconds for timing uncertainty. The remaining headroom is 1.6 minus 0.60 minus 0.40, or 0.60 seconds. These assumed inputs explain the method; they are not measured equipment capabilities.

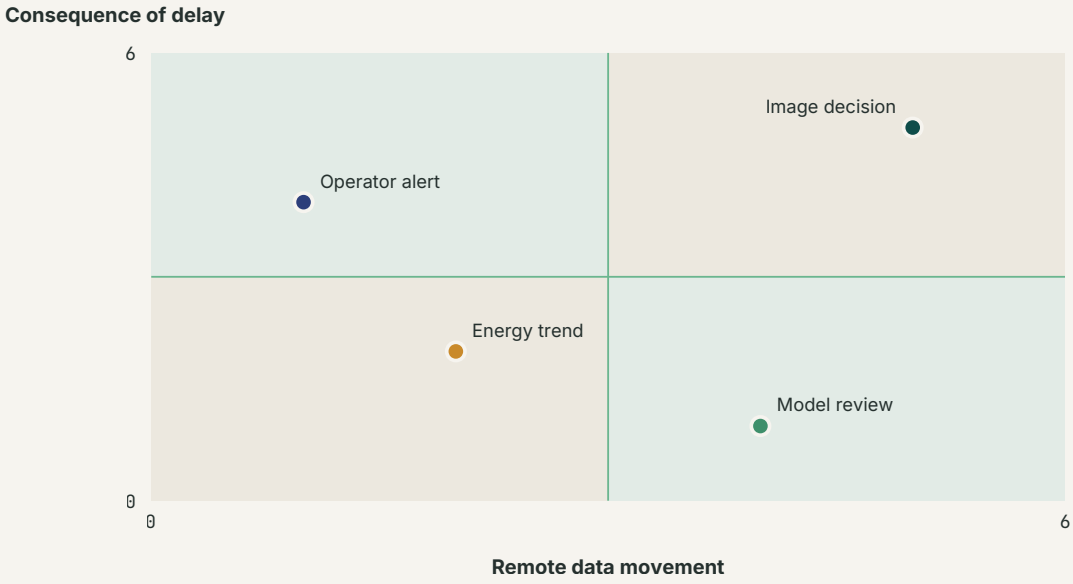
For the same illustrative geometry, doubling speed to 1.0 metre per second reduces travel time to 0.8 seconds. The assumed processing and uncertainty allowance still total 1.0 second, so the design misses its deadline by 0.2 seconds. A faster network alone cannot repair this budget. The designer must change the geometry, speed envelope, processing path or permitted action, then measure the revised chain. An operator review cannot be assumed to fit a subsecond automatic-action budget.

GFT recommends measuring the whole decision path under competing workloads. Include image persistence, record synchronisation, log rotation, temperature-related throttling and restart activity in the test conditions. Report delayed and lost decisions alongside typical timing. A benchmark that excludes the result store or action acknowledgement does not test the same boundary as production. Where the required deadline remains unproven, restrict the station to advisory or retrospective use and retain the approved quality procedure.

The illustrative placement matrix in **Exhibit 24**, p. 93 expresses consequence and evidence movement using ordinal design scores. The coordinates are GFT assumptions, not measured network or computing performance. The matrix is informed by edge placement and selective collection, and by the energy paper's abstract describing local storage, processing and selective transfer. It supports discussion of placement; it cannot select hardware or certify a deadline. [S16] [S14] [S0602]

EXHIBIT 24 / ILLUSTRATIVE ORDINAL SCORES, 1 TO 5

Immediate decisions favour local placement



ILLUSTRATIVE [S16] [S14] [S0602] GFT assumed scores: higher x means greater data movement if processed remotely; higher y means greater consequence of waiting. Coordinates describe design judgement, not measurements or an economic ranking.

GFT proposes a split between inference needed during production and analysis that can wait for synchronisation. Local software should hold the approved model, required configuration and operating context before the shift begins. Remote services may compare outcomes, curate review material and prepare a candidate update. Approval and deployment of that update remain a separate change. The placement argument does not require a complete MEC platform at every camera; the relevant requirement is a bounded local service with known dependencies. [S16]

For resource analytics, GFT recommends the same discipline without copying an image-inspection deadline. A local energy calculation can inform an operator while later reports reconcile production context. The energy study's publisher-deposited abstract describes a database at the edge followed by analytical processing. This

supports the pattern, not a claimed saving or a guarantee against power failure. GFT would require the factory to identify the useful response interval and measure continuity at its own metering boundary. [S0602]

GFT also recommends an explicit rule for late results. Once the inspected fabric has passed the permitted action point, an inference may still support investigation or a material hold downstream. It should no longer issue the original time-sensitive instruction as though its deadline remained open. Record the missed deadline and the alternative action. This rule prevents a delayed but otherwise correct classification from being counted as a successful timely decision, and it makes the consequence of queue growth visible to the quality owner.

6.3

Only useful results need to cross the constrained link

Separate decision evidence from image distribution. MQTT is a publish/subscribe messaging transport whose specification discusses constrained environments. W3C describes selected industrial data collection and edge preprocessing. GFT therefore recommends publishing compact quality events first, with images subject to a separate retention and transfer policy. A transport choice cannot determine whether a defect label is correct, whether its material identity is intact or whether a remote reviewer has enough evidence. [S0601] [S14]

For each publication, GFT recommends retaining event identity, observed material identity, source time, position, defect taxonomy, model and configuration references, acquisition state and current decision state. Include a confidence score only with its declared meaning. Distinguish an image location from evidence that the image remains retrievable. Record what a receiver can do if the image is absent. The result should explain whether material was inspected and whether its release was authorised, even when the image transfer is delayed.

GFT recommends separate states for acquisition, inference execution, defect indication and human disposition. A valid image with no detected defect is different from an invalid image with no inference. A supervisor's acceptance of material is different from the model's indication. Keeping these states separate prevents a subscriber from treating an empty defect field as evidence of satisfactory quality. Require receiving applications to reject unknown state combinations rather than repair them through a hidden assumption.

Before sizing transfer, declare the event cadence. GFT recommends distinguishing frame-level observations, merged defect events and final dispositions. The event reduction rule must preserve the relationship between successive images of the same physical defect. If the rule merges indications, retain their references and spatial boundary. If it discards repeated observations, record that policy. Otherwise a link-efficient system may change the apparent defect count whenever the camera rate or merge window changes.

AN ILLUSTRATIVE TRANSFER CALCULATION SEPARATES IMAGES AND RESULTS

Assume an illustrative stream of 4 images per second, with each compressed image occupying 0.5 decimal megabytes. The image stream is 4 multiplied by 0.5, or 2 megabytes per second, equivalent to 16 megabits per second before protocol overhead. Assume a compact result of 2 kilobytes for every image. Results then occupy 8 kilobytes per second, equivalent to 0.064 megabits per second. The image payload is 250 times the result payload under these assumptions. This compares payloads, excludes encryption and protocol overhead, and assumes no event merging or compression change.

The illustrative calculation supports a transfer-policy test, not a universal reduction claim. GFT recommends reserving transfer capacity for current results and agreed review evidence before bulk historical images. Identify image classes that require retention, such as disputed decisions, invalid acquisitions and a representative audit sample. Define the sampling policy before deployment and preserve its version. Avoid selecting only convincing detections, because a reviewer also needs evidence about material the system judged acceptable.

GFT recommends minimising the scene to the fabric and necessary reference features. Where people or customer-sensitive patterns may appear, assign an owner for access and retention decisions before image collection begins. ISO's public description of its information-security management standard provides a governance framework; it does not supply this station's image permissions or retention duration. The factory must document those decisions and the means of enforcing them. This chapter does not offer a legal conclusion about a particular capture arrangement. [S24]

Test the published schema with a receiving application before accepting the link. GFT recommends injecting unknown taxonomy values, unavailable image references and delayed dispositions. Check that the receiver preserves the raw event and displays its quality state. A message arriving successfully proves transport arrival within the tested boundary. It does not prove physical inspection coverage, material release or a completed corrective action. The application acceptance record should make those distinctions visible. [S0601]

GFT proposes that the transfer contract distinguish unavailable evidence from restricted evidence. A missing file, an expired retention period and a denied access request have different remedies. The remote reviewer should see the applicable reason without receiving unnecessary sensitive content. The local operator should be able to locate the original material decision even if the reviewer has no image permission. This makes selective publication an evidence policy with accountable exceptions rather than an unexplained reduction in data.

6.4

Local persistence makes interruption recoverable

Persist the evidence needed to explain the local decision before allowing remote delivery status to determine its fate. The energy study's abstract describes local database storage followed by processing. Sparkplug describes a possible edge

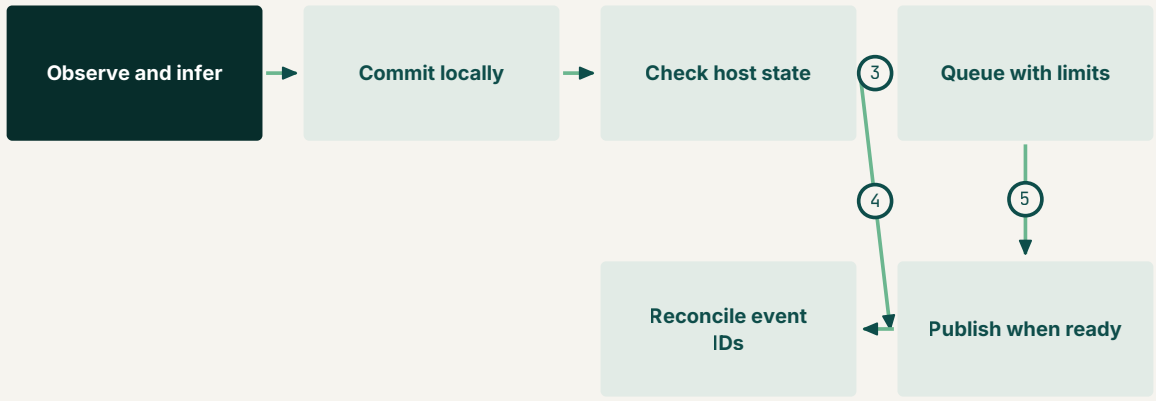
store-and-forward behaviour linked to primary-host state. These are useful patterns, but neither establishes this station's durable-write behaviour, outage capacity or replay correctness. GFT recommends testing those properties independently and keeping inspection eligibility separate from host availability. [S0602] [S40]

GFT's proposed local journal has distinct observation, result, disposition and delivery records joined by stable identifiers. An observation can exist without a completed inference; a result can exist before a human disposition. Delivery acknowledgements should refer to the actual event version delivered. Preserve the order and meaning of revisions. If the service restarts between recording a result and publishing it, the unresolved delivery remains visible and eligible for reconciliation rather than being silently treated as completed.

The illustrative recovery flow in **Exhibit 25**, p. 96 shows the intended boundary. Host loss changes remote synchronisation while a healthy, eligible local cell continues under its approved procedure. This is a GFT design condition, not automatic Sparkplug behaviour. A host may be essential in another application's architecture. Test the dependency explicitly by removing the actual remote services, including authentication, name resolution and image access, that the local process might otherwise use. [S40]

EXHIBIT 25 / ILLUSTRATIVE RECOVERY SEQUENCE

Host loss changes delivery before inspection eligibility



3. Offline
4. Online
5. Host returns

ILLUSTRATIVE [S40] [S0602] [S22] Proposed GFT sequence. Local continuation requires healthy acquisition, inference, storage and an approved procedure. Sparkplug does not guarantee storage capacity, persistence or uninterrupted inspection.

GFT recommends acknowledging remote delivery only at a defined receiving boundary. A broker acknowledgement and a committed business record are different events. Define which acknowledgement permits local deletion, how its identity survives restart and how lost acknowledgements are retried. The receiving store should

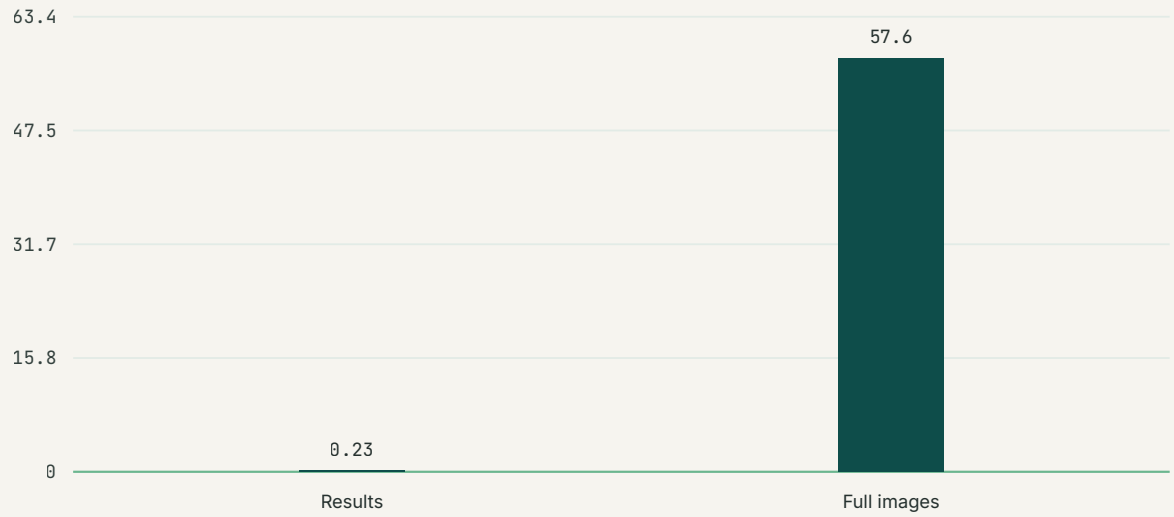
recognise a repeat delivery without duplicating the production event. This is an application requirement beyond the fact that MQTT provides messaging transport. Test repeat delivery and conflicting revisions against the declared event policy. [S0601]

> **WORKED EXAMPLE: CAPACITY DEPENDS ON THE RETENTION POLICY**
>
> Use the illustrative stream already defined: 4 images per second at 0.5 megabytes each and a 2-kilobyte result per image. Assume uninterrupted acquisition during an illustrative 8-hour host outage. Results require 4 multiplied by 2 multiplied by 28,800, or 230,400 kilobytes, equal to 230.4 decimal megabytes. Full images require 2 multiplied by 28,800, or 57,600 megabytes, equal to 57.6 decimal gigabytes. Storage for the same outage therefore depends heavily on whether every image must remain available.
>
> Assume an illustrative 20 percent provisioning allowance on these payload totals. Results require 276.48 megabytes and images 69.12 gigabytes. The combined provision is 69.39648 gigabytes, rounded to 69.4 gigabytes. This allowance is a planning assumption, not a measured database overhead. It excludes operating-system space, model files, spare capacity and device endurance. Verify actual record sizes and retention before selecting storage.

The illustrative storage comparison in **Exhibit 26**, p. 97 uses the same inputs. Its common scale intentionally makes compact result storage appear small beside image storage. The exhibit does not compare reliability or evidence sufficiency. A policy that retains fewer images may save capacity while weakening the ability to resolve disputed decisions. GFT recommends documenting that trade-off in the quality acceptance record.

EXHIBIT 26 / DECIMAL GIGABYTES BEFORE PROVISIONING ALLOWANCE

Image retention dominates the illustrative outage buffer



ILLUSTRATIVE Illustrative inputs: 4 images/s, 0.5 MB/image, 2 kB/result and 8 hours. Results = 4 × 2 × 28800 / 1000000 = 0.2304 GB. Images = 4 x 0.5 × 28800 / 1000 = 57.6 GB. Payload only; all inputs assumed.

GFT recommends staged storage alarms tied to a calculated remaining capture horizon. Stop nonessential bulk transfer or discard only material already eligible under the approved retention policy. Never silently overwrite unresolved decisions. If the required evidence cannot be preserved, the station should declare itself ineligible for the affected quality function and invoke its fallback. Local persistence protects recoverability only while capacity, power and the recovery procedure satisfy the agreed boundary.

GFT recommends testing the recovery drain as well as the outage buffer. After a host returns, historical records and current production compete for the available link and receiving capacity. Give current decision visibility an explicit priority while preserving source time for historical records. Verify that the backlog shrinks under the tested production load. If it continues to grow, the system has reconnected without recovering. Record the remaining backlog and its oldest unresolved source time until reconciliation is complete.

6.5

Model outputs remain reviewable production evidence

Preserve enough context to reproduce the meaning of a decision. GFT recommends identifying the model artefact, preprocessing, threshold configuration, defect taxonomy and acquisition setup separately. Store their references with the result rather than relying on whichever versions happen to be installed during later review. The acquisition study motivates attention to input conditions; ISO's security-management overview supports governance at a broad level. The detailed inference-lineage requirement here is GFT's proposed engineering contract. [S26] [S24]

Do not describe a confidence score as a measured probability of correctness unless that interpretation has been validated for the relevant operating scope. GFT recommends recording the score's definition, class association and threshold rule. An uncalibrated ranking may still help order human review, but its use must be declared. Preserve the distinction between a classification score and the station's eligibility state. A high score cannot turn an out-of-scope material or invalid acquisition into an accepted inspection.

The illustrative schema in **Exhibit 27**, p. 99 groups related fields to keep the exhibit readable. Its groups are a GFT proposal, not fields mandated by the cited standards. Implementers should expand the contract into a typed schema with permitted values, null behaviour, units and compatibility rules. The visible table identifies the minimum questions that the receiving quality system must be able to answer.

EXHIBIT 27 / ILLUSTRATIVE FIELD GROUPS

A result preserves observation, interpretation and disposition

Field group	Required meaning
Event and material	Stable identity and fabric reference
Time and position	Source time, clock state, location
Acquisition	Setup reference and validity reason
Inference	Model, preprocessing, threshold
Indication	Taxonomy, score and geometry
Evidence	Image reference and availability
Disposition	Decision, role and revision link
Delivery	Receiver state and replay identity

ILLUSTRATIVE [S26] [S31] [S24] GFT proposed schema, not a standard-mandated field list. Each group needs typed fields, explicit missing-value behaviour and a versioned contract. Acquisition and wireless sources motivate quality checks; they do not define model lineage.

Wireless sensing research discusses synchronisation and radio sharing as data-integrity concerns. GFT recommends preserving those limitations where a wireless source contributes material position or operating context. A received value without a known association to the inspected material should remain uncertain. The model cannot repair a missing association merely by returning a class. Conversely, a wired camera does not make unrelated wireless context trustworthy. Validate the evidence chain at each relevant interface. [S31]

ISO's provenance standard for characteristic master data has a narrower scope than this proposed inference record. Its public abstract addresses provenance representation and exchange for characteristic master data, and explicitly excludes configuration management and change control from its scope. GFT may use governed asset and characteristic records as context, but should not describe that standard as a model-lineage specification. The inference record needs its own versioned contract and approved responsibility. [S41]

GFT recommends keeping later review as an appended decision, with a link to the original result and a reason for amendment. Review should state whether the disagreement concerns the reference label, acquisition, classification, location or release decision. That distinction determines the corrective action. A confirmed camera problem calls for cell intervention; a disputed taxonomy calls for quality governance; a missed process defect calls for production investigation. A generic override count would hide these different mechanisms.

For routine review, GFT proposes a paired evidence pack containing an observation reference, the original indication, the final disposition and any consequent action. Sample material judged acceptable as well as flagged material, using a declared selection rule. Keep unavailable images and unresolved cases in the denominator.

Compare only periods whose operating scope and configuration are understood. This makes the record useful for production learning without implying that a dashboard trend alone demonstrates a change in model performance.

When production and quality teams disagree about an indication, GFT recommends preserving the disagreement as an unresolved case with an assigned owner. Do not overwrite the original label to make the records agree. State the additional inspection, sample examination or material identification needed to close it. Once resolved, retain the reason and the resulting material decision. This provides a usable review trail without pretending that the model output, the first operator response and the final disposition were always equivalent.

6.6

Manual inspection and safe production remain the fallback

Define a quality fallback before giving the cell a production responsibility. NIST distinguishes recovery validation, personnel training and exercises that expose planning gaps. GFT recommends applying that discipline to inspection eligibility, not merely to server availability. Identify who can suspend automated reliance, who controls affected material and who authorises return to service. A functioning model process is insufficient if the station cannot preserve evidence or the operator cannot execute its decision. [S22]

GFT proposes explicit fallback entries for camera loss, an obscured lens, a lost trigger association, an unavailable model, unacceptable processing delay, full storage and missing remote services. For each, state the detection method, visible condition, immediate material treatment and responsible role. Remote-host loss can permit continued local inspection when all local conditions remain valid. Acquisition failure cannot receive that same permission. The procedure should make the difference apparent at the station.

Manual inspection should be a defined operating mode with a scope the quality owner has accepted. GFT recommends recording its capacity, inspection method, material identification and release authority. If the fallback cannot keep pace, the approved response may require slower production, segregated output or a quality hold. Do not label uninspected fabric acceptable simply to preserve an output counter. Keep fallback output identifiable so later reports distinguish normal digital coverage from material controlled by another procedure.

Protective action remains a separate engineering responsibility. NIST describes safety systems and the importance of considering their relationship with other OT functions. GFT's requirement is stricter for this proposed analytics path: inference availability, broker state and dashboard access must not disable guards, emergency stops or other protective functions. Any automated quality action needs its own reviewed machine interface and permitted behaviour. This chapter does not qualify an inference model as a protective safety function. [S17]

GFT NOTE

Commission the smallest complete decision path: representative material enters a qualified observation area, an eligible result is preserved locally, an authorised role or mechanism acts, and the affected material remains identifiable. Exercise that path with the remote host absent and with local faults introduced under an approved test procedure. Expand its operating scope only after the quality owner can explain every uninspected interval, unresolved decision and recovery gap.

GFT recommends a staged restart that first restores the approved acquisition setup and operating context. Check the installed model and configuration references, available storage, clock state and a reference sample before enabling reliance on results. Reconcile the material that passed during the interruption separately from delayed event delivery. Replaying historical evidence should not replay physical line actions against current fabric. Record which results arrived late and which material remains on hold.

Acceptance should retain the injected fault, observed station state, affected material boundary, response taken and evidence of restoration. GFT recommends a witnessed exercise with the intended operators, including a shift handover while a quality hold remains open. A technical recovery can succeed while responsibility is lost between shifts. NIST's separation of testing, training and exercise provides the relevant discipline; the material-control checks are GFT's application to the inspection cell. [S22]

GFT recommends recording acceptance by permitted use. The quality owner may accept image capture for retrospective review while withholding permission for automated disposition. Another acceptance may permit local alerts but require manual release of affected material. State each permission, its operating scope and the evidence that supports it. Changes to the cell should be assessed against those permissions. This avoids turning a successful demonstration of one function into an implied approval for every function the software happens to expose.

The chapter's recommendation is deliberately bounded. GFT proposes local inference where its deadline and continuity requirements justify placement, selected publication where the evidence policy permits it, and manual control when the cell is ineligible. Hardware selection, model economics, network recovery and wider deployment governance need their own acceptance evidence. The cell earns operational trust when observation, interpretation and disposition can all be inspected, including the intervals in which the digital path could not perform its assigned task.

Keep the decision beside the fabric, and keep the evidence beside the decision.

CHAPTER SOURCE NOTES

S14 World Wide Web Consortium. 2023-12-05.

S16 European Telecommunications Standards Institute. 2024-04.

S26 MDPI, Electronics. 2025-09-18.

S17 National Institute of Standards and Technology. 2023-09.

S0602 MDPI, IoT (publisher-deposited abstract via Crossref). 2024-09-28.

S0601 OASIS Open. 2019-03-07.

S24 International Organization for Standardization and International Electrotechnical Commission. 2022-10.

S40 Eclipse Foundation. 2022-11-16.

S22 National Institute of Standards and Technology. 2010-11-11.

S31 MDPI, Sensors. 2020-04-28.

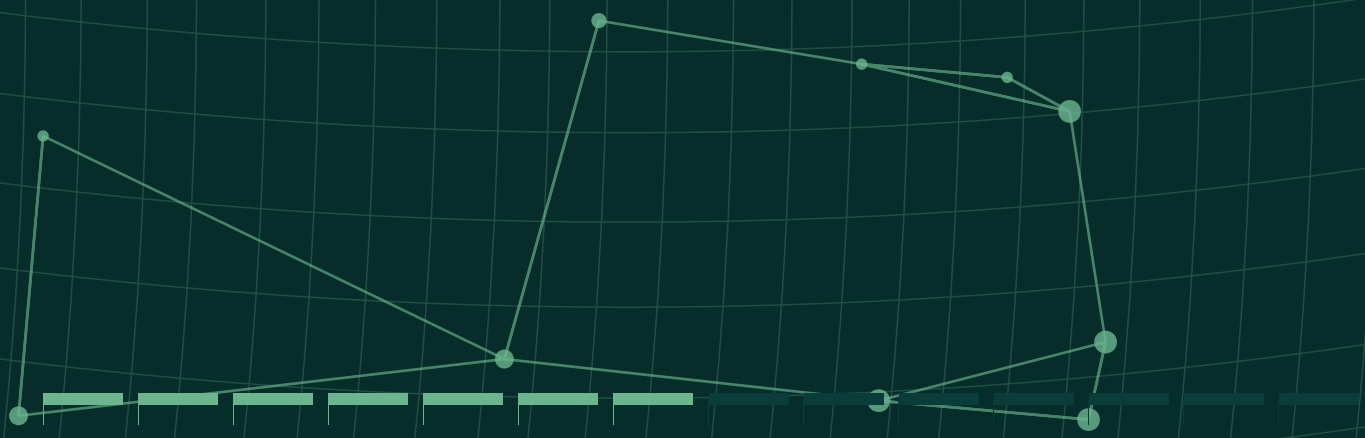
S41 International Organization for Standardization. 2016-10.

07

PART II: CAPTURE AT THE PRODUCTION EDGE / CONTINUITY DESIGN

Networks should degrade in layers, not fail all at once

The connected factory must lose remote services before it loses local evidence, and lose local evidence before it ever loses safe machine operation.



CHAPTER 07
DECISION TEST

The connected factory must lose remote services before it loses local evidence, and lose local evidence before it ever loses safe machine operation.

A connected textile factory must never share a single failure boundary between its external cloud services, its plant historian and its shop-floor machines. [S17] Where electrical interruptions and wide-area network drops occur regularly, treating full connectivity as a precondition for production either halts manufacturing unnecessarily or trains supervisors to bypass digital records whenever a link flickers. [S32] Green Future Tech Limited (GFT), Hong Kong SAR, designs the continuity architecture so that capabilities shed in a strict sequence from the outside in. [S17] Remote convenience fails first, central plant dashboards degrade second, local edge buffering persists third, and machine control alongside automatic protective safety continues undisturbed at the core. [S19] Recovery reverses that sequence through controlled buffer replay, timestamp reconciliation and exercised restoration. [S22]

This rule serves Cameroon's cotton-to-garment sector, where inherited records indicate that approximately 95 percent of cotton fibre is exported raw and about 5 percent is processed locally, alongside a national target of 50 percent local processing by 2030. [R06] Because domestic mills such as CICAM have suffered severe capacity underutilisation across spinning, weaving and finishing assets, any modernised or greenfield facility must protect machine availability while building an unbroken audit trail of throughput, quality and resource intensity. [R12] Earlier chapters established how asset identity and time discipline turn signals into events, how brownfield capture isolates legacy controllers, and how edge inference keeps inspection on the line. [S01] This chapter specifies how network zones, power telemetry, remote-access conduits and isolated backups hold those functions together when external links or utility feeders fail. [S19]

7.1

Zones and conduits make permitted communication explicit

An industrial network fails all at once when every device, database, workstation and gateway resides on a flat address space without enforced boundaries. [S19] In a flat network, a misconfigured vendor laptop, a switch broadcast storm or office malware can reach programmable logic controllers, operator panels and inspection cameras simultaneously. [S17] To prevent cascading disruption, IEC 62443-3-2 establishes security zones and conduits as the primary structural method for industrial automation and control system design. [S19] A zone groups physical and logical assets sharing common operational responsibilities and security requirements, while a conduit defines the exclusive pathway through which two zones exchange permitted traffic. [S19]

7.2 FIVE OPERATIONAL ZONES ISOLATE FAILURE DOMAINS ACROSS THE PLANT

GFT organises the factory network into distinct zones aligned with ISA-95 and IEC 62264-1, augmented by an industrial demilitarised zone (DMZ) and an independent safety-related zone. [S01] Each zone has an operational owner, a survival requirement during utility or link loss, and a default-deny posture against undocumented connec-

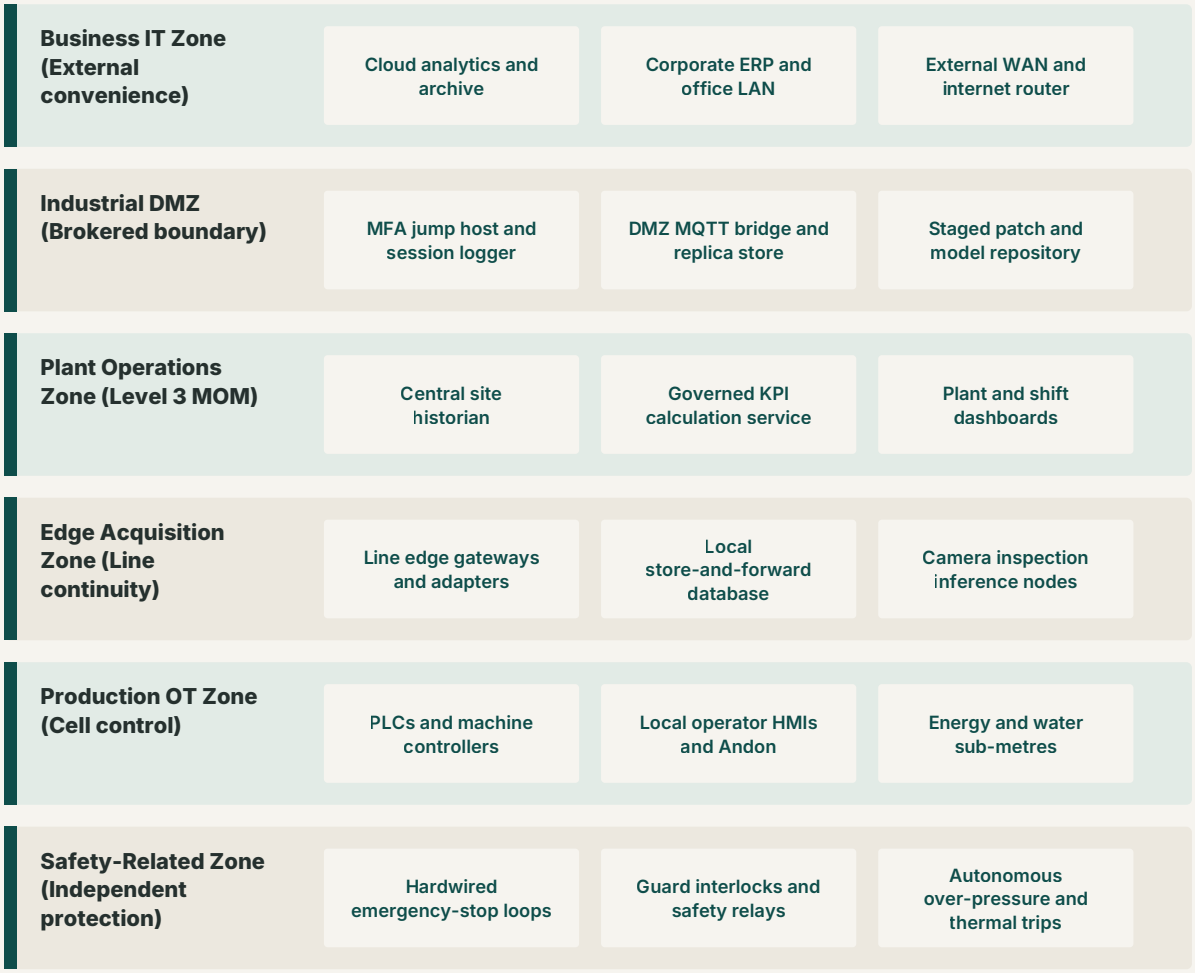
tions. [S02] Rather than allowing business applications at Level 4 to poll shop-floor controllers directly, every cross-layer conversation terminates at a governed boundary. [S19]

The outermost layer is the Business Information Technology Zone, hosting enterprise resource planning, office workstations, internet gateways and cloud synchronisation services. [S01] Below it sits the Industrial DMZ, a buffer zone that terminates inbound remote-access sessions, hosts replica brokers, stages approved updates and blocks direct routing between office networks and the factory floor. [S19] Beneath the Industrial DMZ lies the Plant Operations Zone, corresponding to IEC 62264-1 Level 3 manufacturing operations management, where the central site historian, the governed key performance indicator (KPI) service, manufacturing execution software and plant-wide dashboards operate. [S02]

Closer to the physical process, the Edge Acquisition and Buffering Zone contains line-level gateways, local databases, protocol adapters and vision-inference nodes described in ETSI's multi-access edge computing framework and W3C's industrial edge pattern. [S16] Below the edge nodes sits the Production Operational Technology (OT) Zone, subdivided by manufacturing cell or line so that spinning frames, looms, dyeing vats, cutting tables and sewing lines do not share a broadcast domain. [S17] Finally, the Safety-Related Zone encompasses emergency-stop circuits, safety relays, light curtains, guard interlocks and thermal or pressure trips, which provide automatic functional safety in response to hazards and never depend on edge gateways, plant brokers or enterprise networks. [S23] The structural relationship across these layers is shown in **Exhibit 28**, p. 106. [S19]

EXHIBIT 28 / ARCHITECTURAL ZONES

Security zones preserve a safe core as external services fail



ILLUSTRATIVE [S17] [S19] [S20] [S23] Illustrative zone-and-conduit design informed by NIST SP 800-82 Rev. 3 (S17), IEC 62443-3-2 (S19), CISA CPG 2.0 (S20) and IEC functional-safety principles (S23). Not a representation of an implemented GFT plant.

As **Exhibit 28**, p. 106 illustrates, placing the Industrial DMZ between Business IT and Plant Operations ensures that an office outage or compromise cannot propagate into production. [S19] While **Exhibit 28**, p. 106 defines the containment hierarchy, it cannot specify physical cable routing, switch backplane capacity or firewall rule tables for a specific mill building, which require a site audit before commissioning. [S17]

7.3 CONDUITS SPECIFY DIRECTION, PROTOCOL, AUTHENTICATION AND FAILURE STATE

Defining zones is ineffective if conduits permit broad port ranges or unauthenticated bidirectional traffic. [S19] Every conduit in GFT's architecture is documented as a contract recording source zone, destination zone, initiating asset, permitted protocol, transport encryption, credential mechanism, bandwidth limit and link-failure behaviour. [S19] Conduits between the Edge Acquisition Zone and Plant Operations Zone carry MQTT messages suited to constrained bandwidth environments with Sparkplug session state, accompanied by Network Time Protocol (NTP) synchronisation and structured syslog notifications conforming to RFC 5424. [S05]

Where production evidence moves from the Plant Operations Zone to Business IT or an external cloud service, the conduit terminates inside the Industrial DMZ rather than spanning both firewalls in one hop. [S17] The internal historian or broker publishes selected aggregates to a DMZ relay service, and the enterprise consumer pulls from that relay across a separate external conduit. [S14] If the wide-area network fails or the external firewall blocks traffic, the internal conduit between edge nodes and the plant historian experiences zero routing change. [S19]

Acceptance testing for zones and conduits verifies positive and negative assertions before the network carries live traffic. [S19] Engineers confirm that authorised payloads cross each conduit within latency bounds, and execute negative tests proving that unauthorised protocols, reverse-initiated connections from Business IT into Production OT, and lateral traffic between separate lines are rejected and logged via RFC 5424 syslog messages. [S0702] Any conduit for vendor maintenance remains disabled by default and follows the session rules established later in this chapter. [S20]

7.4

The failure order begins with external convenience

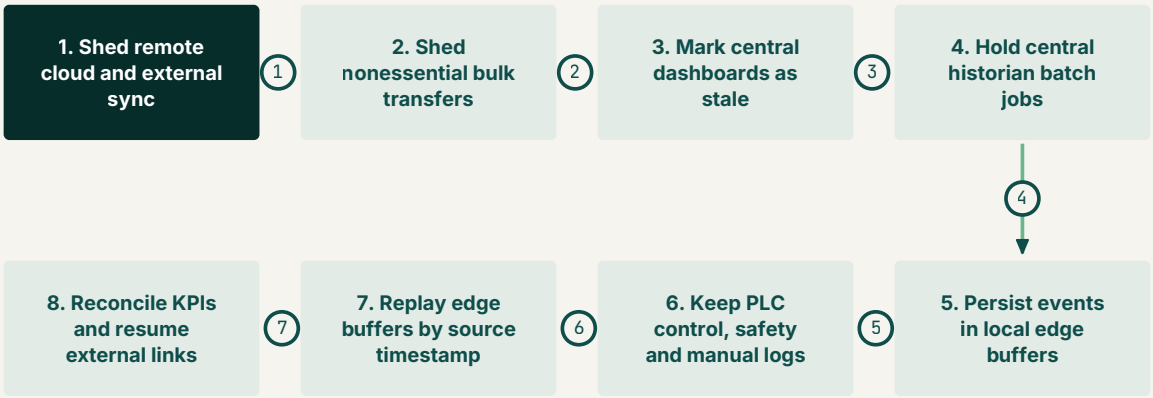
When electrical supply degrades or network links drop, an uncoordinated factory suffers partial failures: operator screens freeze on old numbers, edge scripts block waiting for remote acknowledgements, and supervisors lose confidence in displayed metrics. [S40] NIST Special Publication 800-82 Rev. 3 emphasises that operational technology security and architecture must address unique performance, reliability and safety requirements alongside confidentiality and integrity. [S17] Network degradation must therefore follow a deterministic shedding order in which external convenience is sacrificed immediately to preserve local measurement, and local measurement is sacrificed if necessary to preserve safe physical operation. [S23]

7.5 SIX SHEDDING STAGES PROTECT MACHINE CONTROL AND LOCAL EVIDENCE

GFT defines six ordered operational states through which the factory transitions as connectivity or power degrades, paired with a reverse sequence for recovery. [S17] Each transition is triggered by observable link, broker or power telemetry rather than guesswork, and maps to an explicit decision owner. [S22] The sequence is presented in **Exhibit 29**, p. 108. [S17]

EXHIBIT 29 / ORDERED OPERATIONAL STAGES

Degradation follows a predetermined order



1. WAN loss
2. LAN congestion
3. Server loss
4. Broker offline
5. Edge bypass
6. Host online
7. Audit verified

ILLUSTRATIVE [S17] [S22] [S23] [S40] Illustrative degradation and recovery sequence derived from NIST OT reliability principles (S17), NIST SP 800-34 Rev. 1 contingency planning (S22), IEC functional safety (S23) and Sparkplug offline buffering (S40).

As shown in **Exhibit 29**, p. 108, the first capability shed during a wide-area disruption is Stage 1: remote cloud analytics, off-site replication, vendor telemetry and non-urgent downloads across the external router. [S14] Because local operations in the Plant Operations Zone and Edge Acquisition Zone do not depend on cloud execution, shedding Stage 1 has zero effect on shift huddles, line inspection or machine control. [S16] If internal bandwidth degrades or the plant switches to backup generation with constrained load, Stage 2 sheds nonessential internal traffic, including archive exports and training-image uploads. [S17]

If the Plant Operations Zone server cluster or site-wide broker becomes unreachable, the architecture enters Stage 3 and Stage 4. [S40] In Stage 3, central dashboards mark affected metrics as stale using Sparkplug session-state conventions rather than displaying the last received figure as live production. [S40] In Stage 4, the central historian suspends downstream KPI batch aggregations for missing intervals so that partial denominators do not corrupt shift-level overall equipment effectiveness (OEE) calculations under ISO 22400-2. [S08]

Simultaneously, the factory operates in Stage 5 inside the Edge Acquisition Zone: line gateways continue polling programmable logic controllers and energy metres, camera nodes execute local defect inference, and every event is written with its source timestamp into local non-volatile storage. [S29] Even if an individual edge gateway

fails during a prolonged crisis, the line transitions cleanly to Stage 6: machine controllers, local human-machine interfaces (HMIs) and functional-safety circuits continue operating while operators record hourly output, rejects and downtime on manual fallback sheets aligned with the KPI dictionary. [S23] While **Exhibit 29**, p. 108 establishes the sequence of degradation and restoration, timeout thresholds in seconds for each transition must be tuned to the polling intervals of each work centre. [S08]

> **GFT NOTE**
>
> Never allow an edge gateway, message broker or dashboard service to sit in the execution loop of a machine cycle or a protective safety interlock. [S23] In GFT's architecture, edge acquisition from brownfield controllers is strictly read-only or asynchronous, and any quality alert generated by an edge camera triggers an Andon signal for operator disposition rather than reaching into PLC motion control. [S17] When a switch or edge node loses power, the loom, spinning frame or sewing line remains safely operable under its native controller and hardwired protective circuits. [S23]

7.6 PRODUCTION RULES GOVERN HOW LONG A LINE RUNS WITHOUT DIGITAL CAPTURE

Separating machine control from digital capture prevents network glitches from causing mechanical stops, but raises a governance question: how long should a factory continue manufacturing when its digital measurement layer is degraded? [S22] The answer depends on which layer has failed and what risk unmeasured operation creates for product quality, traceability and worker safety. [S17]

Under IEC functional-safety principles, any failure inside the Safety-Related Zone, such as a faulted emergency-stop relay, a bypassed light curtain or an inoperable pressure trip, triggers immediate automatic protection and halts the affected equipment regardless of network health. [S23] Conversely, when the Safety-Related Zone and Production OT Zone are healthy but the Edge Acquisition Zone or Plant Operations Zone is degraded, production continues under manual fallback procedures documented in the contingency plan. [S22]

For discrete garment assembly and cutting lines, if an edge gateway fails, supervisors revert to hourly paper tally cards capturing good pieces, reworked pieces, defect codes and downtime minutes at the end of each hour, preserving the elements required by ISO 22400-2 for later reconciliation. [S08] For continuous processes such as fabric inspection or dyeing, losing an edge camera cell or bath-temperature logger requires shifting to a reduced line speed or an increased manual inspection sampling rate specified in the use-case contract so that defective rolls are not shipped blind. [S22] The contingency plan sets a maximum duration for manual fallback per process before the plant manager reviews whether to pause the line or reallocate maintenance resources. [S22]

7.7

Store-and-forward needs explicit capacity and ordering rules

Claiming that an edge node buffers data offline is meaningless unless the architecture specifies what is stored, how much capacity is provisioned, how records are evicted if an outage exceeds the design horizon, and how historical data is flushed without overwhelming the broker upon reconnection. [S40] In Sparkplug 3.0.0, when the Primary Host Application publishes a state message indicating that it is offline, or when the broker connection drops, the edge node stores metric updates locally and flushes that historical data once the Primary Host Application returns online. [S40] Turning that mechanism into dependable engineering requires explicit storage arithmetic and replay discipline. [S29]

7.8 EDGE PERSISTENCE BUDGETS DEPEND ON PAYLOAD RATE, RETENTION HORIZON AND MEDIA POLICY

An industrial edge gateway in a textile plant handles two classes of data that must never compete for the same storage pool. [S29] The first class consists of compact structured records: machine cycle counts, status transitions, energy and water metre readings, RFID bundle events, and tabular defect-inference results carrying timestamp, asset identifier, defect class, confidence and coordinates. [S40] As shown in edge energy-monitoring deployments, storing metre readings in a local database at the edge allows the gateway to compute derived parameters locally and persist compact time-series rows across link interruptions. [S29]

The second class consists of high-volume diagnostic media, notably camera frames from fabric-inspection stations. [S16] If a line-scan camera writes uncompressed frames into the disk partition holding the store-and-forward buffer, a multi-hour outage will fill the disk, crash the database and halt event capture. [S22] GFT enforces partition separation at the edge: structured events and metre time series occupy a dedicated storage pool sized for a multi-day outage, while diagnostic image crops use a separate circular ring buffer that overwrites the oldest unflagged frames first without encroaching on event persistence. [S29]

METHOD

To size an edge buffer, GFT calculates required non-volatile storage across an approved outage horizon `H` (hours, at 3,600 seconds per hour) by summing byte rates of retained streams and applying an indexing multiplier `k_fs` via `Storage_req = H * 3600 * k_fs * sum(f_i * b_i) + M_media`, where `f_i` is event frequency (events per second) for stream `i`, `b_i` is payload size (bytes) including asset identity, source timestamp, receipt timestamp, sequence number and quality flags, and `M_media` is the circular quota for compressed defect-crop thumbnails. [S40]

The boundary covers local solid-state storage on the line edge node. [S29] During recovery, minimum drain time `T_drain` to flush backlog `B_backlog` across a conduit with throughput `R_link` and replay bandwidth share `alpha_replay` is checked via `T_drain = B_backlog / (alpha_replay * R_link)`. [S22]

WORKED EXAMPLE

Consider an illustrative sizing calculation for one garment and inspection line edge node required to survive an illustrative 72-hour loss of the Plant Operations Zone historian and external WAN without losing structured records. [S29] Assume an illustrative 40 work-centre streams publishing at 0.2 messages per second each (8.0 messages per second total) at 250 bytes per record. [S40] Add an illustrative 6 electrical and water sub-metres sampled every 10 seconds (0.6 samples per second total) at 200 bytes per sample, plus an illustrative 2 RFID readers and 1 inspection station generating 1.4 event and defect-metadata records per second at 500 bytes per record. [S29]

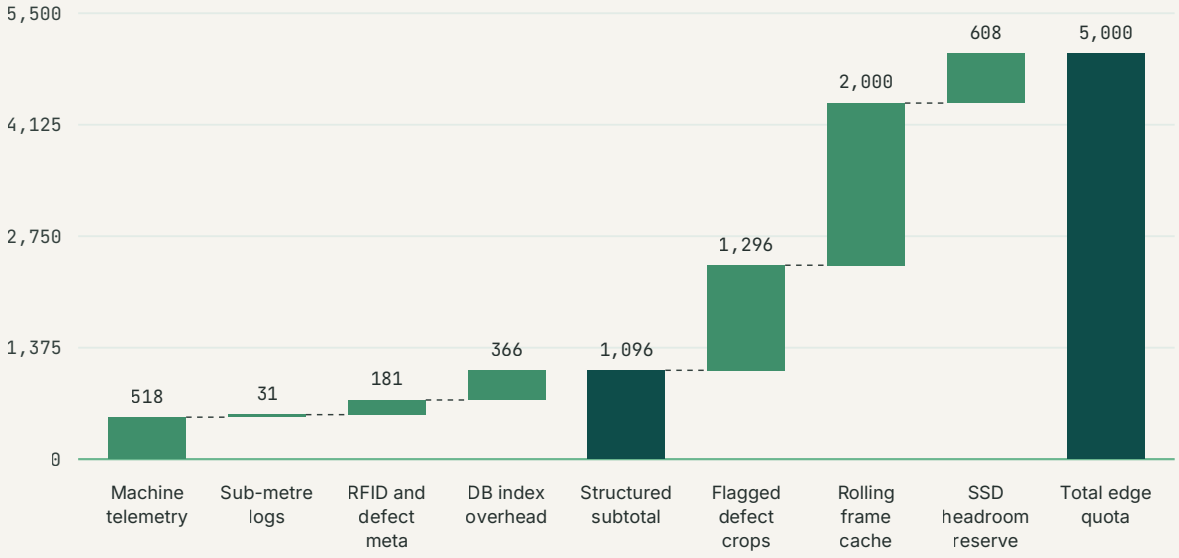
The combined structured rate before overhead is `(8.0 * 250) + (0.6 * 200) + (1.4 * 500) = 2820` bytes per second, or 10.152 megabytes per hour (illustrative). [S40] Over an illustrative 72-hour horizon (259,200 seconds), raw structured records total `2820 * 259200 = 730,944,000` bytes, or about 731 megabytes (illustrative). Applying an illustrative indexing factor of `k_fs = 1.5` brings structured persistence to 1,096 megabytes (about 1.1 gigabytes, illustrative). [S29]

For inspection media during the 72-hour outage, the node stores compressed defect crops only for flagged anomalies at an illustrative 120 defects per hour and 150 kilobytes per crop, equalling 18 megabytes per hour or 1,296 megabytes over 72 hours, plus an illustrative 2,000 megabytes for a circular cache of unflagged calibration frames and 608 megabytes of solid-state wear-levelling headroom, yielding an illustrative 5,000 megabytes (5.0 gigabytes) total. [S22] On recovery with an illustrative link of `R_link = 2.0` megabits per second (250 kilobytes per second) and replay share `alpha_replay = 0.50` (125 kilobytes per second), flushing the 731 megabytes of structured records takes `731,000 / 125 = 5,848` seconds (97 minutes, illustrative), while draining all 2,027 megabytes including defect crops takes `2,027,000 / 125 = 16,216` seconds (4.5 hours, illustrative). [S40]

The breakdown of this 72-hour edge storage allocation is shown in **Exhibit 30**, p. 112. [S29]

EXHIBIT 30 / MEGABYTES (ILLUSTRATIVE 72-HOUR OUTAGE)

Edge storage allocation separates compact event logs from diagnostic media



ILLUSTRATIVE [S22] [S29] [S40] Illustrative 72-hour edge storage budget derived from the worked example inputs: 518 MB machine telemetry, 31 MB sub-metre logs, 181 MB RFID and defect metadata, 366 MB database index overhead (totalling 1,096 MB structured store), 1,296 MB flagged defect crops, 2,000 MB rolling calibration ring buffer and 608 MB wear-levelling reserve, reconciling to 5,000 MB. Informed by S22, S29 and S40.

As **Exhibit 30**, p. 112 demonstrates, structured telemetry, sub-metre readings, RFID events and defect metadata together with indexing overhead require only 1,096 megabytes over an illustrative 72-hour outage, less than a quarter of a 5,000 megabyte partition. [S29] Even a modest solid-state drive preserves weeks of structured evidence provided that camera frames remain confined to a ring buffer. [S40] What **Exhibit 30**, p. 112 cannot establish is flash wear under elevated temperatures in a spinning or dyeing hall, which requires storage rated for the site's thermal envelope and health monitoring via the gateway. [S17]

7.9 HISTORICAL FLUSH REQUIRES SEQUENCE DISCIPLINE, RATE LIMITING AND DUPLICATE SUPPRESSION

When a severed conduit returns to service, naive edge implementations dump their queued backlog onto the broker at wire speed. [S40] That behaviour creates three hazards: it saturates the link so live events and alerts are delayed, it causes out-of-order updates where an old buffered reading overwrites a newer live reading, and it produces duplicate counts if a connection flickers during replay. [S40]

GFT prevents replay chaos through four rules at the edge-to-broker interface. [S40] First, live state and historical backfill travel with explicit semantic distinction: under Sparkplug 3.0.0, when the Primary Host Application publishes its online `STATE` message, the edge node establishes current session state via birth certificates (`NBIRTH` and `DBIRTH`) so live consumers immediately see current machine status, and historical records are ingested using their original source timestamps rather than arrival time. [S40] Second, the edge node prioritises structured telemetry and event metadata ahead of diagnostic image crops, flushing the compact records needed for ISO 22400-2 KPI reconstruction before bulk media. [S08]

Third, historical flush is rate-limited to a fraction of conduit capacity so real-time production messages and heartbeats retain head-of-line priority. [S17] Fourth, every buffered event carries a deterministic identifier combining the ISA-95 asset identifier, source timestamp, clock-synchronisation quality state and a monotonic sequence number, allowing the historian in Chapter 09 to perform idempotent upserts and discard duplicates. [S31] If an outage lasts so long that even the structured partition nears capacity, the eviction policy preserves hourly summaries, quality dispositions and audit logs while down-sampling high-frequency analogue waveforms first. [S29]

7.10

Power and link telemetry are part of production context

In factories with stable grids and redundant fibre rings, engineers often treat electrical supply and telecommunications as background utilities managed solely by facilities or IT staff. [S17] In Cameroon and Central Africa, power and connectivity directly shape when machines run, why energy intensity fluctuates and why telemetry contains gaps. [S32] If the architecture fails to record utility and link status alongside machine events, a plant manager reviewing yesterday's dashboard cannot tell whether a zero-output hour on a weaving line resulted from mechanical breakdown, material starvation, a blackout or a dead switch. [S29]

7.11 INFRASTRUCTURE OBSERVATIONS FRAME THE CONTINUITY REQUIREMENT WITHOUT REPLACING SITE SURVEYS

Two institutional datasets illustrate why GFT treats power and connectivity interruptions as baseline operating conditions. [S32] First, in the World Bank's Cameroon 2024 Enterprise Survey Country Profile (fieldwork June 2024 to February 2025), 93.3 percent of surveyed formal private non-agricultural firms reported experiencing electrical outages, compared with 69.7 percent across Sub-Saharan Africa and 55.0 percent across lower-middle-income economies, with surveyed Cameroonian firms reporting 10.4 outages in a typical month against 6.9 in Sub-Saharan Africa. [S32] Second, the International Telecommunication Union (ITU) reports in *Facts and Figures 2024* that Africa has the largest mobile-network coverage gap globally, with 14 percent of the population lacking access to a mobile-broadband network. [S33] These indicators are summarised in **Exhibit 31**, p. 114. [S33]

EXHIBIT 31 / PERCENT AND MONTHLY COUNT

Fourteen percent of Africa's population is outside mobile-broadband coverage

14%

Africa population outside mobile broadband

ITU 2024 regional population gap; not a site link metric (S33)

93.3%

Cameroon surveyed firms with power outages

World Bank 2024 survey, all firms vs 69.7% SSA (S32)

10.4

Cameroon outages in a typical month

World Bank 2024 survey, all firms vs 6.9 SSA (S32)

SOURCED [S32] [S33] Sourced macro context from ITU Facts and Figures 2024 (S33) and World Bank Cameroon 2024 Enterprise Survey Country Profile, published 2025 (S32). ITU figure is a regional population coverage indicator, not a factory link-quality metric. World Bank figures cover all surveyed formal private firms in Cameroon, not a textile-only sample.

It is essential to state what **Exhibit 31**, p. 114 does and does not establish. [S33] The ITU's 14 percent figure measures the share of Africa's population outside a mobile-broadband footprint in 2024; it does not measure packet loss, latency, jitter or availability inside an industrial zone in Douala, Dibamba or Garoua. [S33] Similarly, the World Bank's 93.3 percent outage incidence and 10.4 monthly outage frequency represent all surveyed formal private firms in Cameroon rather than a specific textile mill's feeder reliability. [S32]

Nevertheless, **Exhibit 31**, p. 114 proves that assuming continuous wide-area connectivity or uninterrupted utility power is an engineering error. [S32] Before selecting any cloud dependency or wireless backhaul, GFT requires a site link and power survey measuring round-trip latency, packet loss, spectrum interference, availability and voltage stability over representative shifts, while sizing local uninterruptible power supplies (UPS) and edge buffers to ride through utility transitions. [S31]

7.12 UPS, GENERATOR AND GATEWAY STATE TRAVEL IN THE SAME TIMELINE AS MACHINE EVENTS

To make degraded operation legible to operators and KPI services, the Edge Acquisition Zone captures infrastructure telemetry using the same clock discipline and event contract applied to production machines. [S13] Every line edge gateway, core switch, site historian and inspection lighting controller is backed by an uninterruptible power supply bridging the interval between a utility drop and either generator transfer or controlled shutdown. [S0701]

Rather than treating the UPS as an unmonitored battery box, the gateway polls the UPS controller and packs its status into the telemetry stream. [S0701] The Internet Engineering Task Force (IETF) UPS Management Information Base (RFC 1628) defines a standard object model across battery, input, output, bypass and alarm groups. [S0701] Specifically, RFC 1628 specifies `upsBatteryStatus` to distinguish normal capacity (`batteryNormal`), low remaining run-time (`batteryLow`) and depleted

capacity (`batteryDepleted`), alongside `upsSecondsOnBattery`, `upsEstimatedMinutesRemaining`, `upsEstimatedChargeRemaining` and `upsInputLineBads`, which counts input out-of-tolerance transitions. [S0701]

Recording these UPS states alongside main-incomer and generator sub-metres yields three benefits. [S29] First, when `upsBatteryStatus` transitions to `batteryLow` and `upsEstimatedMinutesRemaining` drops below threshold without generator transfer, the edge node and historian flush open transactions to disk and execute an orderly halt before battery depletion corrupts file systems. [S0701] Second, during shift reviews under ISO 22400-2, any interval where utility power dropped is tagged with the source timestamp of the feeder trip and generator takeover, preventing blackouts from being misclassified as unexplained mechanical downtime. [S08]

Third, local gateway survivability does not rely on a single routing switch. [S0703] Where a hall deploys redundant layer-3 switches or dual WAN gateways, the Virtual Router Redundancy Protocol Version 3 (RFC 5798) provides dynamic election and failover of the virtual default router if the master router becomes unavailable, without requiring dynamic routing on every controller or edge node. [S0703] Every VRRP transition, WAN drop and Sparkplug state change is logged with its NTP-synchronised timestamp so engineers and managers inspect one unified timeline. [S13]

7.13

Remote access is a controlled session, not a permanent exception

Textile factories operating mixed-generation equipment rely on external equipment manufacturers, vision specialists or integrators outside Cameroon or off site to diagnose controller faults, tune drives or update software. [S17] When implemented hastily, vendors often install persistent virtual private network (VPN) tunnels, remote-desktop agents or cellular modems directly inside machine cabinets. [S19] Those bypasses destroy the zone-and-conduit boundary, create unmonitored entry points into Production OT and allow a supplier incident to compromise multiple factories. [S21]

7.14 MULTI-FACTOR AUTHENTICATION AND JUMP HOSTS TERMINATE EXTERNAL SESSIONS IN THE DMZ

GFT treats remote access as a temporary, audited conduit that is closed by default and opened only for an approved maintenance window. [S19] In *Cybersecurity Performance Goals 2.0* (CPG 2.0), the Cybersecurity and Infrastructure Security Agency (CISA) requires multi-factor authentication (MFA) on all remotely accessible operational-technology accounts and systems, including vendor and maintenance accounts, engineering workstations and remotely accessible human-machine interfaces where available. [S20] Furthermore, ISO/IEC 27001 places supplier access, access control and operational change within a governed information security management system. [S24]

In GFT's architecture, no external user, vendor engineer or corporate IT administrator can initiate a direct session into the Plant Operations Zone, Edge Acquisition Zone or Production OT Zone. [S19] Instead, every remote connection follows a five-step path terminating in the Industrial DMZ: [S20]

1. The engineer submits a request identifying the target asset in the ISA-95 hierarchy, the operational reason, the procedure, the rollback plan and the requested window. [S01]
2. A designated plant or OT owner confirms that the machine is in a safe maintenance state and approves a time-bounded window under ISO/IEC 27001 rules. [S24]
3. The external engineer authenticates to the Industrial DMZ jump host using MFA as required by CISA CPG 2.0, establishing an encrypted outer session terminating in the DMZ. [S20]
4. From the DMZ jump host, an inner conduit is enabled only to the approved edge node or workstation for the ticket duration, while session activity and RFC 5424 syslog events are recorded to an append-only audit store. [S0702]
5. When the window expires or the local supervisor revokes approval, the jump host terminates the session, disables the temporary conduit and rotates the credential. [S19]

7.15 LEGACY CONTROLLERS WITHOUT NATIVE AUTHENTICATION REQUIRE COMPENSATING ISOLATION

In brownfield mills, many programmable logic controllers, dye-house controllers and older HMIs lack modern operating systems and cannot natively evaluate MFA tokens or transport encryption. [S17] CISA CPG 2.0 acknowledges this hardware limit by specifying MFA on remotely accessible HMIs "when available", while IEC 62443-3-2 provides the compensating mechanism through zone segmentation and conduit mediation. [S20]

Where a legacy controller or HMI cannot support strong authentication natively, GFT prohibits making that device directly reachable from any remote conduit, even via port forwarding. [S19] Instead, the controller is isolated inside its cell-level Production OT Zone behind an edge gateway or hardened engineering workstation inside the plant. [S17] The remote specialist authenticates with MFA to the DMZ jump host and interacts only with the hardened engineering proxy, which enforces role-based authentication, mediates protocol commands and prevents scanning of the legacy cell. [S20]

For critical machine-control logic modifications, the architecture requires dual-control confirmation: the remote specialist stages a configuration or ladder-logic file in the DMZ repository, and a technician standing at the machine turns a local key-switch or confirms the download on the local HMI after verifying that the equipment is clear of personnel. [S23]

> **WARNING**
>
> Unauthorised cellular routers, Wi-Fi bridges and "call-home" vendor dongles inside machine panels represent a severe failure mode in brownfield factories. [S19] Because these devices bypass the Industrial DMZ, they expose unpatched controllers directly to external carrier networks and can introduce unexpected polling onto deterministic machine buses. [S17] During the Phase 0 audit, engineers must inspect every cabinet, trace all cables, remove undocumented modems, and contractually require suppliers to route all support through the Industrial DMZ jump host. [S20]

7.16

Recovery is proven by restoration and exercise

A degraded-operation architecture is incomplete if it only describes how systems step down during an outage without proving how they return to verified service after hardware failure, database corruption or a cyber incident. [S22] In many projects, backup and recovery exist only as a policy document or a script writing files to a network drive on the same subnet as production servers. [S21] When ransomware, a power surge or a storage fault strikes, operators discover too late that backups were encrypted over the same network path, machine recipes are missing, or no one knows the sequence for rebuilding an edge gateway. [S21]

7.17 ISOLATED BACKUPS COVER MACHINE LOGIC, EDGE CONFIGURATIONS AND HISTORIAN STATE

To protect critical infrastructure against cyber threats, CISA and international authorities advise organisations to regularly test backup procedures and ensure backups are isolated from network connections that could spread malware. [S21] In a connected textile factory, business files are only a fraction of what is required to resume manufacturing. [S17] A complete industrial backup inventory preserves six operational asset classes across the ISA-95 hierarchy: [S01]

- **Machine recipes and controller logic:** verified PLC programs, motion parameters, dye-cycle profiles, loom tension settings and inverter configurations in the Production OT Zone. [S17]
- **Edge node images and interface contracts:** container images, Modbus register maps, IO-Link device descriptions, OPC UA certificates and Sparkplug configurations in the Edge Acquisition Zone. [S40]
- **Local inference models and calibration records:** approved vision model weights, camera calibration matrices, lighting targets and defect-taxonomy rules at inspection stations. [S24]
- **Network and security configurations:** switch VLAN tables, VRRP router configurations, firewall zone-and-conduit rules, broker access lists and jump-host policies. [S19]
- **Central historian and event stores:** raw observation partitions, contextualised EPCIS and production event tables, and audit logs in the Plant Operations Zone. [S21]
- **KPI dictionary and dashboard definitions:** versioned ISO 22400-2 formulas, denominators, shift calendars, alert rules and dashboard templates. [S08]

GFT enforces the "pull and isolate" pattern for these backups. [S21] Rather than allowing Production OT or Plant Operations servers to hold write credentials to the backup vault, a dedicated backup service pulls encrypted snapshots across a scheduled conduit and replicates an immutable offline copy that cannot be modified from the live network. [S21] Every change to a PLC recipe, edge configuration or KPI formula approved under ISO/IEC 27001 change control triggers a versioned backup alongside its cryptographic hash so engineers verify integrity before restoration. [S24]

7.18 CONTINGENCY EXERCISES MEASURE ACTUAL RESTORATION TIME AND DATA LOSS AGAINST FACTORY CONTRACTS

Having isolated backups on a shelf still does not guarantee that a small factory team can restore a line within an acceptable window. [S22] NIST Special Publication 800-34 Rev. 1 distinguishes three activities that organisations frequently conflate: testing validates recovery capabilities, training prepares personnel for plan activation, and exercising the plan identifies planning gaps; combined, these activities improve preparedness. [S22] Complementing this discipline, the NIST Cybersecurity Framework (CSF) 2.0 places recovery governance, risk ownership and verification outcomes at the core of cybersecurity management. [S18]

GFT translates NIST's contingency framework into testable factory recovery contracts defined by two parameters per service class: the Recovery Time Objective (RTO), the maximum target elapsed time to restore service after disruption, and the Recovery Point Objective (RPO), the maximum tolerated age of lost data measured backwards from disruption. [S22] Because each mill has different shift patterns, buffer capacities and staffing, GFT does not invent universal minute or hour targets in the reference architecture; instead, the process and OT owners approve target RTO and RPO values on the use-case cards prior to commissioning and verify performance through restoration drills. [S22] This recovery contract is detailed in **Exhibit 32**, p. 119. [S21]

EXHIBIT 32 / CONTRACTUAL FIELDS PER ASSET CLASS

Recovery objectives become testable factory contracts

Asset or service class	Primary backup artefact	RPO mechanism	RTO verification test	Responsible role
PLC logic and machine recipes	Offline versioned project and parameter image	Zero unrecorded change via ISO 27001 gate	Bare-metal controller reload on spare unit	Maintenance and OT lead
Edge gateway configuration	Signed container image and register map	Versioned commit on each approved change	Cold-spare edge node swap and birth check	OT systems engineer
Inspection results and crops	Local edge database and ring buffer	Buffered at edge; manual gate if node fails	Camera node rebuild and calibration check	Quality and edge lead
Production and metre events	Edge store-and-forward plus paper tally	Sparkplug flush reconciled with shift cards	Broker reconnect and idempotent replay drill	Production and data owner
Central site historian	Isolated immutable database snapshot	Snapshot plus replay from 72h edge buffers	Clean restore to isolated server and hash check	IT/OT infrastructure lead
KPI service and dashboards	Versioned ISO 22400 dictionary and UI code	Recomputed from reconciled historian events	Service redeployment and formula audit check	Performance management owner

ILLUSTRATIVE [S21] [S22] [S24] [S40] Illustrative RTO and RPO governance worksheet informed by CISA isolated-backup guidance (S21), NIST SP 800-34 Rev. 1 testing, training and exercise rules (S22), ISO/IEC 27001 change control (S24) and Sparkplug historical replay (S40). Target durations are site-specific owner inputs, not invented here.

A close inspection of **Exhibit 32**, p. 119 shows why edge persistence improves the recovery posture of the central historian. [S40] In an ordinary database, if a daily snapshot is taken at midnight and storage fails at noon, twelve hours of transactions are lost unless synchronous replication was active. [S21] In GFT's edge-first architecture, if the central historian is restored from a 24-hour-old isolated snapshot, the missing 24-hour window remains preserved inside the 72-hour local store-and-forward databases of the line edge nodes. [S29]

Historian recovery therefore follows a four-step reconciliation procedure: [S22]

1. Engineers restore the central historian onto clean, isolated hardware from the latest verified offline snapshot and confirm integrity before reconnecting conduits. [S21]
2. Once the historian and broker return online, line edge gateways replay retained event records covering the gap between the snapshot timestamp and the present, using deterministic event identifiers and source timestamps to populate the interval without duplication. [S40]
3. Where an edge gateway itself was offline during the incident, operators enter signed manual fallback tally cards into the reconciliation queue with an explicit manual-entry provenance state so KPI calculations under ISO 22400-2 remain auditable. [S08]

4. The performance-management service recalculates affected shift and daily KPIs from the reconciled events, compares totals against shift logs, records actual RTO and RPO against targets, and logs planning gaps into the ISO/IEC 27001 corrective-action register. [S24]

While **Exhibit 32**, p. 119 establishes the verification contract across all six asset classes, it cannot substitute for executing drills on the shop floor. [S22] As Chapter 13 specifies, no pilot line in Phase 1 transitions from shadow mode to operational dependence until a cold-spare edge swap, a network disconnection drill and an isolated historian restore have been executed by factory staff and signed off by process and safety owners. [S22] By engineering the network to shed external dependencies cleanly, buffer evidence locally with source-time integrity, monitor power and link state, lock remote access behind multi-factor jump hosts and rehearse isolated restoration, the factory ensures that infrastructure turbulence never becomes an excuse for blind or unsafe production. [S17]

A resilient factory does not pretend that every link will hold: it decides in advance which layer fails first, preserves the evidence at the edge and keeps the machine safe when the screen goes dark.

CHAPTER SOURCE NOTES

S17 National Institute of Standards and Technology. 2023-09.

S32 World Bank Enterprise Surveys. 2025.

S19 International Electrotechnical Commission. 2020-06.

S22 National Institute of Standards and Technology. 2010-11-11.

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

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

S01 International Society of Automation. 2025-04-10.

S02 International Electrotechnical Commission. 2013-05-22.

S16 European Telecommunications Standards Institute. 2024-04.

S23 International Electrotechnical Commission. 2022-04-28.

S20 Cybersecurity and Infrastructure Security Agency. 2025-12.

S05 OASIS Open. 2019-03-07.

S14 World Wide Web Consortium. 2023-12-05.

S0702 Internet Engineering Task Force and RFC Editor. 2009-03.

S40 Eclipse Foundation. 2022-11-16.

S08 International Organization for Standardization. 2014-01.

S29 MDPI, IoT. 2024-09-28.

S31 MDPI, Sensors. 2020-04-28.

S33 International Telecommunication Union. 2024-11-10.

S13 Internet Engineering Task Force and RFC Editor. 2010-06.

S0701 Internet Engineering Task Force and RFC Editor. 1994-05.

S0703 Internet Engineering Task Force and RFC Editor. 2010-03.

S21 CISA and joint international cybersecurity authorities. 2022-05-09.

S24 International Organization for Standardization and International Electrotechnical Commission. 2022-10.

S18 National Institute of Standards and Technology. 2024-02-26.

GREEN FUTURE TECH

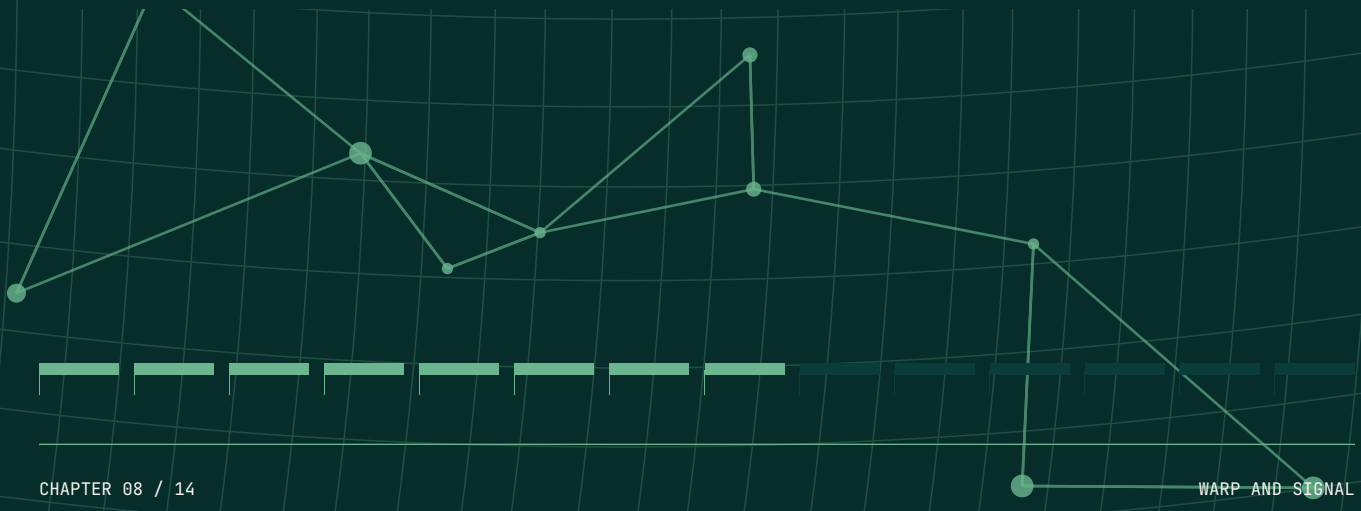
LIMITED / HONG KONG SAR

08

PART III: ONE GOVERNED DATA SPINE / INFORMATION EXCHANGE

Interoperability is a semantic contract before it is a connection

Two systems can exchange packets and still misunderstand production, so the architecture must define information, state, events and asset meaning separately.



CHAPTER 08
DECISION TEST

Two systems can exchange packets and still misunderstand production, so the architecture must define information, state, events and asset meaning separately.

A connected textile factory does not achieve interoperability merely because an edge gateway delivers packets to a broker. Interoperability exists only when the machine controller, local historian, operations service and dashboard interpret every exchanged value under one grammar. [S01] [S04] Two systems can share an active socket, exchange valid JavaScript Object Notation (JSON) documents under RFC 8259 rules, and still corrupt a production record because one side reports cumulative cycles since power-on while the other records good garments completed in the shift. [S0804] [S08] In GFT's architecture, the semantic contract precedes the connection. Before any broker topic is subscribed, the factory specifies five layers of meaning: industrial object and service semantics via OPC Unified Architecture (OPC UA), constrained transport via Message Queuing Telemetry Transport (MQTT), session state and metric validity via Sparkplug, durable asset descriptions via the Asset Administration Shell (AAS), and visibility events via Electronic Product Code Information Services (EPCIS). [S04] [S05] [S40] [S07] [S15]

This separation is vital for industrialisation in Cameroon, where the inherited baseline records roughly 95 percent of cotton fibre exported raw and about 5 percent processed locally, against a 50 percent local-processing target by 2030. [R06] [R07] Factories closing that gap cannot assume homogeneous machinery or uninterrupted links when 93.3 percent of surveyed Cameroonian firms report electrical outages in the 2024 World Bank Enterprise Survey. [S32] When power returns or a brownfield sewing line connects beside a newer dyeing range, the data spine must reconcile replayed messages, machine states, work-in-process (WIP) bundles and asset capabilities without human guesswork. [S32] [S40] GFT's Cameroon textile programme starts with a spreadsheet system of record to build baseline discipline across production, quality, maintenance and resource logs. Replacing spreadsheets with automated capture is safe only when a governed contract preserves the meaning, units, denominators and provenance of every event. [S08] [S41]

8.1

OPC UA separates the information model from transport

Interoperability fails first when engineers confuse the wire protocol that moves bytes with the information model that defines what those bytes mean. Legacy fieldbus integrations often expose a programmable logic controller (PLC) or gateway as a flat array of holding registers. [S09] In MODBUS, communication follows a request and reply model where services are specified by function codes operating on register addresses. [S09] At the sensor level, IO-Link cyclically transmits process data without interference from on-request parameters, commands and events, protecting deterministic scanning while leaving factory context outside the cyclic byte stream. [S10] Mapping register `40012` directly to a dashboard tile traps its meaning (whether raw needle strokes, completed seam cycles, or total pieces including rework) inside one custom script. [S04] [S09]

8.2 FLAT REGISTER MAPS TRAP OPERATIONAL MEANING INSIDE PRIVATE INTEGRATOR CODE

When meaning is buried inside point-to-point scripts, every new application multiplies maintenance. If the maintenance tool, huddle board, historian and enterprise resource planning (ERP) connector parse controller addresses independently, four codebases must be updated whenever a recipe or counter scale changes, diverging silently whenever one script filters warm-up cycles and another does not. [S01] [S02] [S08] ISA-95 Part 1 prevents this by defining how physical assets are organised and specifying the information shared between control and enterprise functions. [S01] IEC 62264-1 reinforces this boundary by describing the manufacturing operations management domain at Level 3 and transactions within Level 3 and with Level 4. [S02] Neither standard lets an enterprise dashboard guess a Level 1 controller's memory layout. [S01] [S02]

OPC UA resolves this at the machine boundary. Under OPC 10000-1, OPC UA is layered to isolate its core design from computing technology and network transport. [S04] Separating the information model from serialisation and transport allows an edge node or controller to expose a graph of typed objects, variables, methods and references that remain identical across client-server sessions or publish-subscribe messages. [S04] [S0802] Instead of an anonymous integer, an OPC UA server on a spinning frame or cutting table exposes a typed `MachineModule` object linked to its parent `WorkCentre` in the ISA-95 hierarchy, carrying engineering units, valid ranges, source timestamps and quality codes. [S01] [S04]

8.3 TYPED OBJECTS AND COMPANION SPECIFICATIONS SEPARATE MACHINE SEMANTICS FROM NETWORK BINDINGS

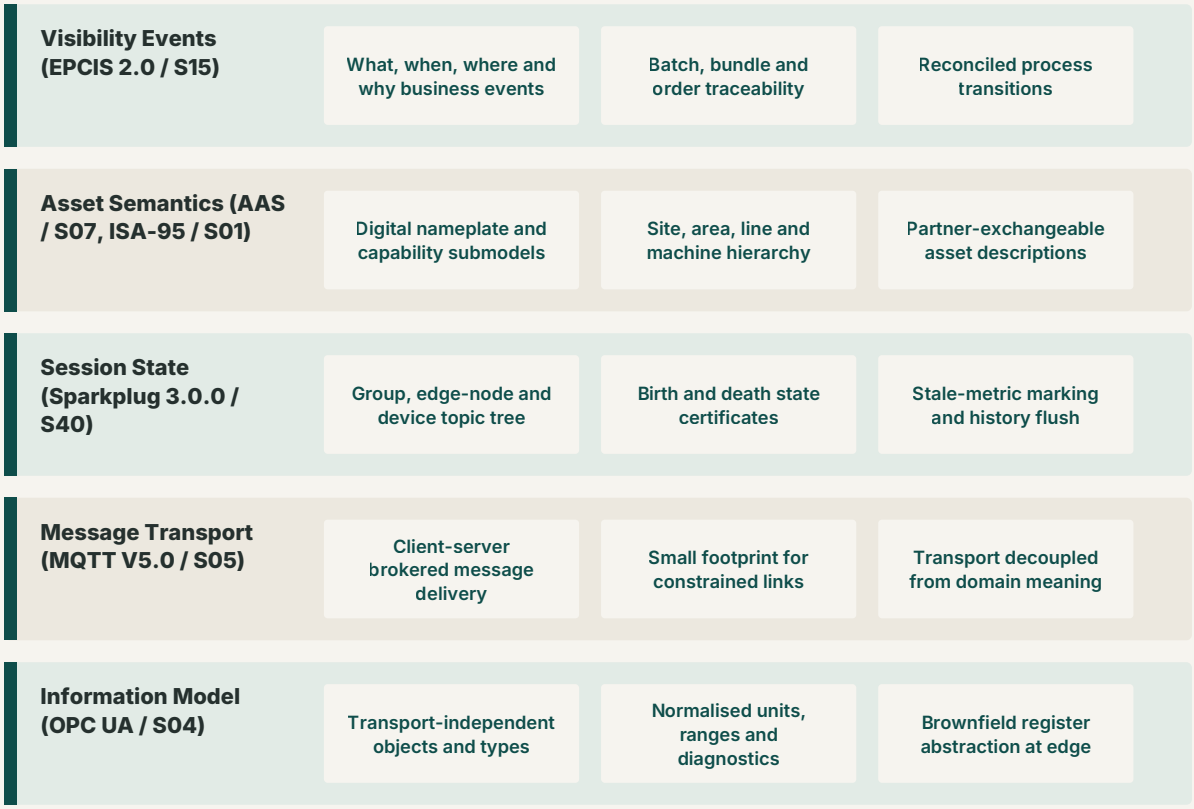
Model-to-transport separation becomes operational in OPC 10000-14, which defines the OPC UA PubSub model. [S0802] OPC UA PubSub complements client-server sessions by distributing data and events from an OPC UA source to interested observers in a device network or analytics system. [S0802] A factory thus retains OPC UA object semantics over brokered constrained links. [S04] [S0802] At the edge gateway, the OPC UA address space acts as the normalisation layer: brownfield MODBUS registers and IO-Link cyclic data are read once, checked against hardware diagnostics, and mapped into typed OPC UA nodes before leaving the work centre. [S04] [S09] [S10]

Where a legacy controller cannot host an OPC UA server, explicit interaction metadata still applies. The World Wide Web Consortium (W3C) Web of Things Architecture 1.1 describes how an industrial edge device collects selected controller data for remote monitoring or preventive maintenance. [S14] The W3C Web of Things Thing Description 1.1 specification defines a formal model providing interactions from a small vocabulary so diverse devices and applications interoperate. [S0801] Whether normalised into OPC UA objects or W3C Thing Descriptions, GFT's rule is invariant: no raw register address, untyped byte array or unscaled count crosses the edge-node boundary into the data spine. [S04] [S14] [S0801]

As shown in **Exhibit 33**, p. 125, the five core standards in GFT's data spine do not compete for the same job; each governs a distinct layer of meaning within the ISA-95 enterprise-to-control hierarchy. [S01] [S04] [S05] [S07] [S15] [S40]

EXHIBIT 33 /

Five standards occupy different layers of meaning



DERIVED [S01] [S04] [S05] [S07] [S15] [S40] Derived architecture stack mapping each standard's stated scope in the opened specifications to its assigned layer in GFT's reference architecture, framed within ISA-95 enterprise-to-control responsibilities (S01).

Exhibit 33, p. 125 shows why no single protocol solves factory interoperability. [S04] [S05] [S40] An integration built solely on MQTT has message transport but lacks a machine object model, session-state conventions, asset submodels and business-event semantics. [S05] Conversely, local OPC UA client-server browsing alone lacks lightweight brokered session management and cross-enterprise visibility grammar. [S04] [S15] [S40] **Exhibit 33**, p. 125 is a structural responsibility map rather than a performance benchmark: it establishes which specification owns each semantic concern, while site acceptance testing must verify actual gateway implementation. [S17] [S30]

8.4

MQTT moves messages but does not define factory meaning

In many digitalisation projects, MQTT is mistakenly treated as an interoperability standard that unifies factory data. The OASIS standard for MQTT Version 5.0 defines its role accurately: MQTT is a client-server publish-subscribe transport protocol whose lightweight design, small code footprint and low bandwidth demand suit constrained Machine-to-Machine (M2M) and Internet of Things (IoT) contexts. [S05]

[S0803] Those transport properties suit a textile mill where gateways share constrained wireless links or buffer telemetry across intermittent wide-area connections. [S05] [S32] [S33] However, transport efficiency is not semantic interoperability. MQTT treats the payload as an opaque byte sequence and places no domain restriction on topic names or message bodies. [S05] [S0803]

8.5 BROKER CONNECTIVITY WITHOUT A PAYLOAD CONTRACT PRODUCES SYNTACTIC COUPLING AND SEMANTIC DRIFT

When a factory deploys an MQTT broker without a semantic contract above it, every publisher invents a dialect. One developer publishes JSON under RFC 8259 rules to `factory/line1/loom3/speed` containing `{"val": 420}`, a second publishes a string to `garoua/spinning/ring02/rpm`, and a third publishes a binary integer to `cameroon/plantA/dyeing/temp` without stating whether the unit is degrees Celsius or tenths of a degree. [S0804] [S05] All three use valid MQTT packets, and the broker routes them without error. [S05] Yet every consumer (historian, alert service, OEE engine and energy dashboard) must maintain custom parsing code per topic, guess missing units, and assume whether timestamps reflect sensor capture or broker arrival. [S08] [S13] [S40]

This pattern turns a decoupled broker into hidden point-to-point dependencies. Whenever a publisher renames a JSON key from `val` to `value` or changes polling from 10 seconds to 60 seconds, downstream KPI calculations fail silently. [S08] [S0804] ISO 22400-2 specifies that manufacturing KPIs are presented through their formula, corresponding elements, time behaviour and unit or dimension. [S08] Because bare MQTT carries none of those attributes natively, an ungoverned broker cannot ensure that two dashboards compute ISO 22400-2 indicators on a compatible basis. [S05] [S08]

8.6 RETAINED MESSAGES AND WILL NOTIFICATIONS REQUIRE EXPLICIT OPERATIONAL RULES

Even transport-level MQTT features can mislead personnel if their operational semantics are ungoverned. In MQTT Version 5.0, retained messages are sent to a session when it first subscribes, while a Will Message is published by the server after a network connection is closed when the connection is not closed normally. [S0803] If an integrator sets the retain flag on a production topic without session tracking, a supervisor opening a dashboard at 08:00 during a power outage receives the retained value published at 06:15 before power failed, mistaking a two-hour-old speed reading for live production. [S32] [S40] [S0803]

GFT therefore prohibits bare, uncontracted MQTT topics in the production spine. Where MQTT provides transport, the architecture binds five specifications onto every connection: a governed topic namespace, a schema-versioned payload, explicit units and quality flags, deterministic session state via Sparkplug or OPC UA PubSub, and named producer and consumer ownership under change control. [S04] [S05] [S24] [S40] [S0802]

8.7

Sparkplug adds state and session conventions

To bridge bare MQTT transport and operational state awareness, the Eclipse Sparkplug Specification 3.0.0 defines an MQTT topic namespace, payload structure and session-state management approach for the Industrial Internet of Things (IIoT), meeting the requirements of supervisory control and data acquisition (SCADA) and human-machine interface (HMI) solutions. [S40] Sparkplug organises MQTT traffic around a deterministic namespace (`spBv1.0/group_id/message_type/edge_node_id/[device_id]`), allowing every participant to discover edge nodes, devices, metric definitions and session transitions. [S40]

8.8 BIRTH AND DEATH CERTIFICATES BIND METRIC VALIDITY TO LIVE SESSION CONTINUITY

Sparkplug makes connection state an explicit part of data semantics. When a Sparkplug edge node connects to the broker, it registers a Death Certificate (`NDEATH`) via MQTT's Will Message mechanism and publishes a Birth Certificate (`NBIRTH`), followed by Device Birth Certificates (`DBIRTH`) for attached devices. [S40] [S0803] The birth certificate declares all metrics exposed by that node or device, with initial values, data types, units and sequence numbers, enabling report-by-exception (`NDATA` and `DDATA`) updates thereafter. [S40]

When an edge node loses power or connectivity unexpectedly, the broker publishes the registered `NDEATH` message, and Sparkplug specifies that all metrics in the previous associated `NBIRTH` and all `DBIRTH` messages in that session under that edge node are now `STALE`. [S40] [S0803] In a Cameroonian mill facing electrical outages or wireless drop-outs, a dashboard or historian following Sparkplug rules immediately marks affected machine metrics `STALE` upon receiving `NDEATH` or `DDEATH`, preventing operators from reading a frozen value as live. [S32] [S40]

8.9 HISTORICAL FLUSH AFTER OUTAGE REQUIRES SEQUENCE AND SOURCE-TIME SEPARATION

Sparkplug also specifies how edge nodes coordinate with a Primary Host Application across outages. When the Primary Host Application publishes a `STATE` message denoting that it is online, the edge node resumes live publishing and flushes historical data stored locally while offline. [S40] This store-and-forward flush complements local persistence, such as the industrial IoT energy-monitoring pattern where an edge server stores metre data locally before computing nine derived parameters. [S29] [S40]

Because buffered records flushed after an outage arrive after their physical occurrence, GFT enforces three rules on every Sparkplug consumer: [S13] [S40]

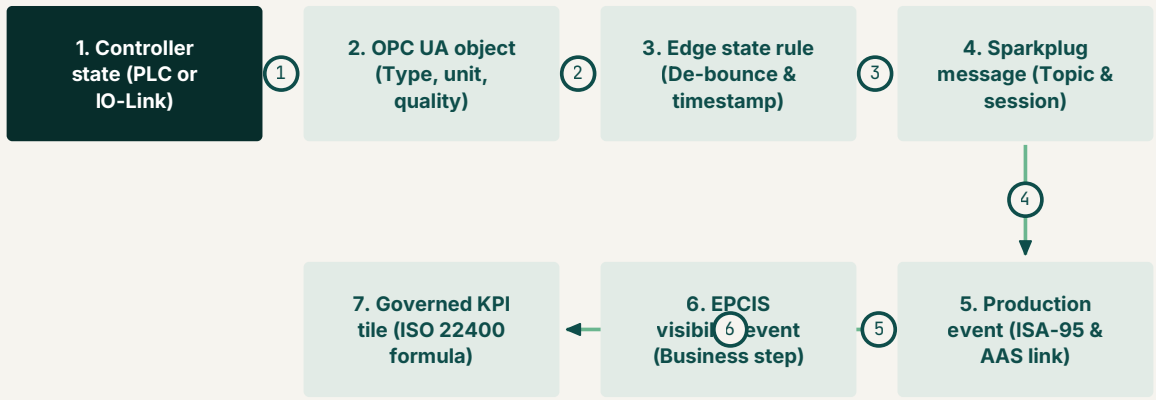
1. **Preserve source time separately from arrival time.** Every flushed metric carries the edge capture timestamp, disciplined via Network Time Protocol Version 4 (NTPv4, offering tens-of-microseconds potential accuracy on modern workstations and fast LANs) or Precision Time Protocol (IEEE 1588-2019, supporting sub-microsecond accuracy where tighter correlation is required). [S12] [S13] [S40] Consumers route flushed records by source timestamp and historical flag. [S40]

2. **Distinguish live report-by-exception from historical replay.** Sparkplug marks historical updates so live HMI displays ignore flushed packets while the historian inserts them into the prior window and flags affected ISO 22400-2 KPI intervals for recalculation. [S08] [S40]
3. **Verify sequence continuity before closing an outage gap.** If the edge buffer overflowed during a long outage, the interval between the last pre-outage timestamp and the first retained historical record is marked `MISSING` rather than interpolated. [S31] [S32] [S40]

The progression from a controller register to an enterprise visibility event and governed KPI is shown in **Exhibit 34**, p. 128, tracing how each standard adds semantic structure. [S04] [S08] [S15] [S40]

EXHIBIT 34 /

A machine state becomes a business event through explicit transformations



1. Hardware read & type binding
2. Local state evaluation
3. Report by exception or buffer
4. Session and context join
5. Batch or bundle transition
6. Windowed KPI calculation

ILLUSTRATIVE [S04] [S08] [S15] [S40] Illustrative transformation sequence showing how a physical controller signal is normalised, transported with session state, contextualised into a production event, exported as an EPCIS visibility event and aggregated into an ISO 22400 KPI (S04, S08, S15, S40).

As **Exhibit 34**, p. 128 shows, a physical signal at Node 1 is not yet a business event at Node 6 or a KPI at Node 7. [S04] [S08] [S15] Hardware diagnostics are checked at the OPC UA boundary (`Node 2`), mechanical bounce or idle cycling is filtered at the edge rule (`Node 3`), session continuity is verified via Sparkplug (`Node 4`), Level 3 context joins the observation (`Node 5`), and only then does the spine emit an EPCIS visibility

event (Node 6) or update an ISO 22400-2 KPI tile (Node 7). [S02] [S04] [S08] [S15] [S40] Because **Exhibit 34**, p. 128 is an illustrative sequence, it defines required logical transformations rather than claiming a fixed software latency.

8.10

AAS carries durable asset descriptions

Live telemetry streams and session certificates report what an asset is doing now, not what the asset *is*. Every machine, metre, camera cell and sensor has slowly changing master attributes: nameplate data, serial number, hierarchy location, rated power, calibration dates and maintenance manuals. [S01] [S07] Plattform Industrie 4.0's Reference Architecture Model Industrie 4.0 (RAMI 4.0) organises these aspects across hierarchy levels, life-cycle value stream and architectural layers. [S03] Embedding static nameplate metadata inside every MQTT packet wastes bandwidth and creates conflicting copies across edge scripts, whereas leaving metadata in paper binders prevents dashboards from querying a frame's spindle count, kilowatt rating or calibration date. [S05] [S07] [S08] [S40]

8.11 SEPARATING STATIC SUBMODELS FROM HIGH-FREQUENCY TELEMETRY PREVENTS PAYLOAD BLOAT

The Industrial Digital Twin Association (IDTA) addresses this through the Specification of the Asset Administration Shell (AAS) Part 1: Metamodel (IDTA 01001-3-0-1), defining the Administration Shell structure for meaningful exchange of asset and Industrie 4.0 component information across value networks. [S07] An AAS assigns each asset a globally unique identifier and organises its attributes into standardised submodels (digital nameplate, technical data, documentation and operational data), referencing external semantic dictionary concepts. [S03] [S07]

GFT derives a strict integration rule from combining the AAS metamodel with the ISA-95 hierarchy and Sparkplug sessions: **high-frequency telemetry carries only the stable asset identifier, metric identifier, value, timestamp and quality state, while durable asset descriptions reside in the AAS registry.** [S01] [S07] [S40] Each Sparkplug DBIRTH certificate includes an immutable reference (aas_global_asset_id and aas_submodel_version) so historians, KPI services and maintenance tools join telemetry to AAS submodels on demand. [S07] [S08] [S40]

8.12 THING DESCRIPTIONS AND ASSET ADMINISTRATION SHELLS BRIDGE BROWNFIELD RETROFITS TO VALUE NETWORKS

In a brownfield textile mill where older spinning, weaving or sewing machines predate Industrie 4.0 standards, engineers construct a repository-hosted AAS instance during the Phase 0 audit. [S07] [S30] [S35] As the W3C Web of Things Thing Description 1.1 specification notes, a description instance can be hosted externally when a device has resource restrictions or is retrofitted. [S0801] GFT uses a W3C Thing Description or OPC UA object to document the retrofitted machine's live edge interface, and registers an AAS instance in the factory repository for its nameplate, electrical rating, calibration record and ISA-95 work-centre assignment. [S01] [S04] [S07] [S0801]

Where master data characteristics are exchanged, ISO 8000-120:2016 specifies requirements for representing and exchanging provenance of characteristic master data. [S41] In GFT's architecture, ISO 8000-120 governs provenance for characteristic master data in asset and material records (such as rated capacity, calibration

factors or cotton fibre-length characteristics), while configuration changes follow ISO/IEC 27001 change control and KPI lineage follows GFT's historian rules. [S08] [S24] [S41]

8.13

EPCIS carries visibility events rather than waveforms

Just as static asset descriptions remain separate from telemetry, continuous process waveforms remain separate from discrete material-flow events. A spindle-vibration waveform, dye-bath temperature curve and motor current trace are periodic *observations* about equipment state. [S04] [S29] Consuming a cotton bale lot at an opening line or completing collar attachment on a garment bundle at an inspection station are discrete *visibility events* about physical objects moving through a process. [S15] [S25]

8.14 OBJECT, AGGREGATION AND TRANSFORMATION EVENTS ASSERT BUSINESS TRANSITIONS

GS1's EPCIS Standard, Release 2.0, enables applications to create and share visibility event data within and across enterprises across four dimensions: [S15]

- **What:** uniquely identified objects (garment bundles, fabric rolls, dye batches or cartons at instance or lot level, plus transformation quantities). [S15]
- **When:** the event timestamp and time-zone offset when the transition occurred, plus record time. [S15]
- **Where:** the physical `readPoint` (scanner or antenna station) and `bizLocation` (finishing buffer or warehouse) after the event. [S01] [S15]
- **Why:** the business context, including `bizStep` (inspecting, packing, shipping or transforming), `disposition` (conformant, non-conformant or quarantined), and `bizTransaction` links. [S15]

To trace how lint from northern Cameroonian ginning plants is spun into yarn, woven or knitted into fabric, dyed, cut and sewn into garments, raw MQTT telemetry is insufficient. [R03] [S05] [S15] An EPCIS chain links `TransformationEvent` records (bale lots to yarn, fabric rolls to cut bundles), `ObjectEvent` records (sewing and inspection steps) and `AggregationEvent` records (garments packed into cartons). [S15] Where European Union buyers require product disclosures, Regulation (EU) 2024/1781 defines a digital product passport as product-specific data specified in the applicable delegated act and accessible electronically through a data carrier. [S39] An internal EPCIS graph linked to batch and resource logs provides export-ready evidence without claiming that raw telemetry is a regulatory passport. [S15] [S39]

8.15 SHOP-FLOOR RFID READS REQUIRE BOUNDARY RULES BEFORE BECOMING VISIBILITY EVENTS

Textile research confirms both the value and pitfalls of shop-floor object tracking: a denim case study showed radio-frequency identification (RFID) supporting work-in-process tracking, and manufacturing execution system (MES) research describes systematic shop-floor RFID deployment for real-time production data. [S25] [S28]

Wireless-sensor and RFID review literature stresses that node time synchronisation and spectrum sharing are vital to protect data integrity and avoid uncertain readings. [S31]

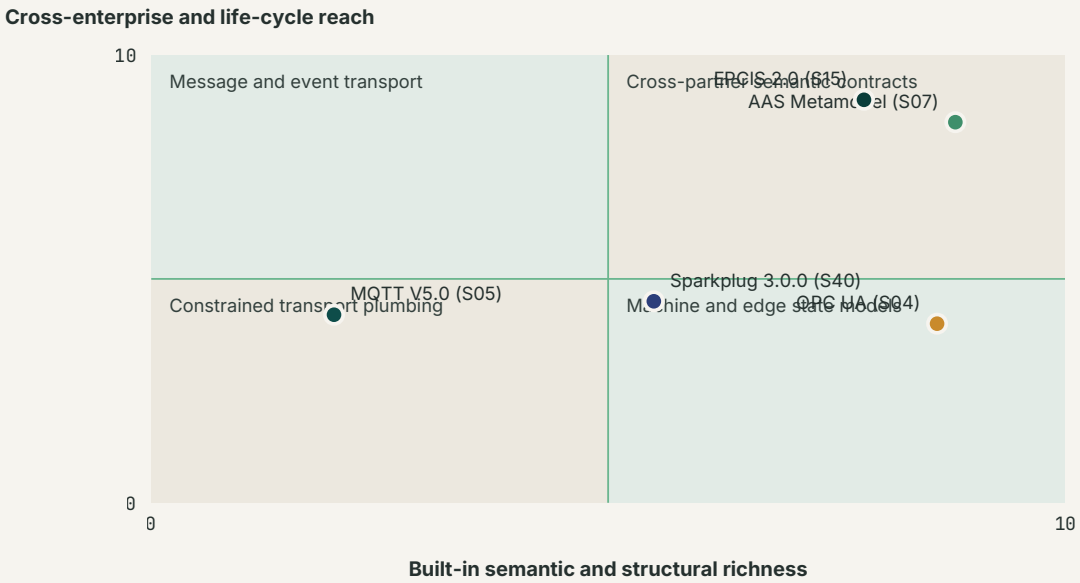
On a sewing line, a raw RFID antenna detection is not automatically a valid EPCIS event: an overhead reader may detect the same bundle tag 40 times in three minutes in an input tray or pick up a stray aisle reflection. [S15] [S25] [S31] GFT therefore places a three-step **event-qualification state machine** between raw detections and the EPCIS store: [S15] [S25] [S28]

1. **Read-point qualification and de-duplication:** Buffer raw reads at the edge, filter by signal threshold or operator trigger, and de-duplicate within the dwell window. [S25] [S31]
2. **Route and order validation:** Verify the bundle identifier against the active IEC 62264-1 route to confirm prior operations are complete. [S02] [S28]
3. **Disposition binding:** Emit the EPCIS event and increment KPI elements only when the work cycle finishes and inline quality disposition is recorded. [S08] [S15]

To prevent stretching any protocol beyond its specification, **Exhibit 35**, p. 131 positions the five standards across built-in semantic richness and cross-enterprise life-cycle reach. [S04] [S05] [S07] [S15] [S40]

EXHIBIT 35 / DERIVED ARCHITECTURAL SCOPE INDEX (0 TO 10)

Protocol claims are limited to their actual scope



DERIVED [S04] [S05] [S07] [S15] [S40] Derived qualitative positioning from opened specifications: X-axis scores built-in semantic and structural richness (from raw byte transport to typed metamodels); Y-axis scores operational reach (from local machine/edge scope to multi-enterprise value networks). Positions reflect specification scope in S04, S05, S07, S15 and S40, not performance benchmarks.

Exhibit 35, p. 131 shows why no single standard is a complete factory operating system. [S04] [S05] [S07] [S15] [S40] MQTT V5.0 (x = 2.0, y = 4.2) is a lightweight transport pipe without domain semantics. [S05] Sparkplug 3.0.0 (x = 5.5, y = 4.5) adds topic, payload and session state for SCADA and IIoT telemetry. [S40] OPC UA (x

= 8.6, y = 4.0) provides deep machine and plant object modelling, while AAS (x = 8.8, y = 8.5) and EPCIS 2.0 (x = 7.8, y = 9.0) extend structured asset and event semantics across partner networks. [S04] [S07] [S15] These derived coordinates reflect specification scope rather than throughput benchmarks.

8.16

The canonical model maps standards without erasing their roles

To combine OPC UA, MQTT, Sparkplug, AAS, EPCIS and ISO 22400-2 without creating five silos, GFT uses a **canonical data model and versioned interface contract**. [S04] [S05] [S07] [S08] [S15] [S40] The canonical model does not replace those standards with a custom protocol; it defines the identity keys, timestamp rules, unit dictionaries and mappings that connect each representation across the spine. [S01] [S06]

8.17 ONE VERSIONED INTERFACE CONTRACT REPLACES PAIRWISE TRANSLATIONS ACROSS THE FACTORY

Without a canonical model, linking N shop-floor sources to M applications requires up to $N \times M$ translation scripts: for 6 equipment families and 5 applications, an illustrative mesh requires 30 pairwise translators. Mapping each source once into a governed canonical contract and each consumer once from that contract reduces the illustrative adapter count from 30 paths to $6 + 5 = 11$ governed interfaces.

This structure rests on three international standards:

- **Enterprise and operations hierarchy (ISA-95 and IEC 62264-1):** Every canonical record carries the ISA-95 asset path (`enterprise / site / area / work_centre / work_unit`) and respects IEC 62264-1 Level 3 and Level 4 boundaries. [S01] [S02]
- **Digital-twin information exchange (ISO 23247-4):** ISO 23247-4:2021 specifies technical requirements for information exchange between entities in a manufacturing digital-twin reference architecture, preventing a dashboard vendor's private schema from becoming the plant's only model. [S06]
- **KPI element traceability (ISO 22400-2):** Every canonical production, quality or resource event maps to the exact input elements, units and time behaviour required by ISO 22400-2 KPI formulas. [S08]

METHOD

Method: Constructing and verifying a canonical interface contract

GFT specifies each interface contract as a versioned document and machine-readable schema before production deployment, binding eight mandatory elements:

1. **Hierarchical asset and point identity:** Map the measurement point to its ISA-95 path (`site_id`, `area_id`, `work_centre_id`, `work_unit_id`, `point_id`) and bind its immutable `aas_global_asset_id`. [S01] [S07]
2. **Source semantic definition:** Record the OPC UA `NodeId`, `BrowseName` and data type (or W3C Thing Description affordance) and the state transition condition. [S04] [S0801]
3. **Transport and session binding:** Specify the Sparkplug topic, metric alias, report-by-exception deadband, and `NBIRTH`/`NDEATH` stale-state rules over MQTT V5.0. [S05] [S40]
4. **Engineering unit and scaling rule:** State the canonical engineering unit, numeric data type, precision and edge scaling formula. [S04] [S08]
5. **Dual-timestamp and clock status:** Require both `source_timestamp` and `receipt_timestamp`, plus clock status (`SYNC_LOCKED` or `SYNC_UNSYNCHRONISED` under NTPv4 or PTP). [S12] [S13] [S31]
6. **Quality-state enumeration:** Map hardware diagnostics, radio link checks and Sparkplug session state into a governed flag (`GOOD`, `STALE`, `UNSYNCHRONISED`, `ESTIMATED`, `CORRECTED`, `MISSING`). [S31] [S40]
7. **Business-event and KPI element mapping:** Define the EPCIS 2.0 mapping (`eventType`, `bizStep`, `disposition`, `readPoint`, `bizLocation`) and the consuming ISO 22400-2 KPI element (`ActualProductionQuantity`, `GoodQuantity`, `ScrapQuantity`, `ActualProductionTime`). [S08] [S15]
8. **Schema version, owner and compatibility rule:** Assign a semantic version (`MAJOR.MINOR.PATCH`), role owner and ISO/IEC 27001 change-control rule requiring a `MAJOR` increment and parallel dual-publishing whenever units, denominators or state logic change. [S24]

An example of this contract for a sewing-line event appears in **Exhibit 36**, p. 134, where all ten attributes are versioned as one unit. [S04] [S24] [S40]

EXHIBIT 36 / CONTRACT SCHEMA ATTRIBUTES

The interface contract is versioned as one unit

Contract field	Canonical specification	Verification and failure rule
Contract ID & Version	urn:gft:contract:sewing-op-complete:v2.1.0	Major increment on unit, denominator or state logic change (S24)
Entity Identifiers	ISA-95: CM-DLA / SEW-A / LINE-02 / ST-14; AAS: urn:gft:aas:lockstitch:0014	Reject message to quarantine queue if asset ID is unregistered
Sparkplug Topic	spBv1.0/CM_DLA_SEW/DDATA/GW_SEW_L02/ST_14	Must match active NBIRTH/DBIRTH session alias on broker (S40)
Source OPC UA Node	ns=2;s=Station14.OperationCycle.CompletedGoodPiece	Mapped from PLC cycle counter + inline disposition confirm (S04)
Payload & Datatype	Sparkplug Metric: UInt32 delta=1; BundleEPC: String; OrderID: String	Schema validation at ingestion; malformed payloads logged & blocked
Unit & Denominator	1 good garment piece completing Operation OP-040 (Collar Attach)	Excludes needle warm-up cycles, dry runs and flagged rework
Source Time Rule	UTC epoch ms from NTPv4-disciplined edge gateway clock	Flag UNSYNCHRONISED if NTP stratum lost or drift > threshold
Quality & Session State	Enum: GOOD | STALE | UNSYNCHRONISED | CORRECTED | MISSING	Auto-transition to STALE immediately on NDEATH or DDEATH (S40)

ILLUSTRATIVE [S04] [S24] [S40] Illustrative interface contract for a sewing-line operation completion event in GFT's reference architecture, informed by OPC UA object modelling (S04), Sparkplug 3.0.0 topic and session state (S40) and ISO/IEC 27001 change control (S24).

EXHIBIT 37 / CONTRACT SCHEMA ATTRIBUTES

The interface contract is versioned as one unit, continued

Contract field	Canonical specification	Verification and failure rule
EPCIS & KPI Mapping	EPCIS ObjectEvent (bizStep=inspecting, disposition=conformant) -> ISO 22400 GQ	Updates GoodQuantity (GQ) only when disposition=conformant
Owner & Compatibility	Owner: Production Engineering Lead; Backwards-compatible minor additions only	Breaking change requires 14-day shadow run & dual-version topic

ILLUSTRATIVE [S04] [S24] [S40] Illustrative interface contract for a sewing-line operation completion event in GFT's reference architecture, informed by OPC UA object modelling (S04), Sparkplug 3.0.0 topic and session state (S40) and ISO/IEC 27001 change control (S24).

Exhibit 36, p. 134 shows why versioning applies to the full operational contract rather than only JSON syntax. [S04] [S24] [S40] Changing the trigger from raw cycles to inspected good pieces alters the ISO 22400-2 element to `GoodQuantity`, requiring a

major version increment (`v2.0.0`), ISO/IEC 27001 change approval and an explicit historian cutover record. [S08] [S24] **Exhibit 36**, p. 134 is an illustrative template rather than a live mill record.

8.18 SCHEMA GOVERNANCE AND CHANGE CONTROL PROTECT HISTORICAL REPRODUCIBILITY

As new defect classes, energy sub-metres and digital product passport fields arise, GFT governs the schema registry as a controlled configuration item under ISO/IEC 27001:2022 across three tiers: [S24] [S26] [S29] [S39]

- **Patch (`v2.1.0` to `v2.1.1`):** Documentation clarifications leaving payload structure, units, state logic and KPI formulas unchanged. [S24]
- **Minor (`v2.1.0` to `v2.2.0`):** Backwards-compatible optional metrics (such as `ambient_humidity_pct` in a birth certificate) that leave existing calculations intact. [S08] [S40]
- **Major (`v2.1.0` to `v3.0.0`):** Changes to units, state transitions, KPI denominators or quality rules, requiring shadow validation and version tagging in the historian. [S08] [S24] [S30]

> **GFT NOTE**
>
> **Practical recommendation: Migrate from spreadsheets to brokers by freezing the data dictionary first**
>
> In GFT's Cameroon textile programme, operations begin with a spreadsheet system of record so teams master daily, weekly and monthly KPI reviews before adding software. When migrating to gateways and brokers, freeze the working spreadsheet columns, units, defect codes and ISO 22400-2 formulas as Version `1.0.0` of the canonical interface contract first, and require the edge pipeline to reproduce manual shift totals in shadow mode before retiring the spreadsheet log. [S08] [S15] [S30] [S40]

WORKED EXAMPLE

Arithmetic reconciliation and payload reduction on an illustrative 40-station garment sewing line

Consider an illustrative model of a 40-station sewing line operating an 8-hour shift (480 minutes, or 28,800 seconds). These assumed figures demonstrate semantic filtering and payload arithmetic, not measured GFT factory results.

Part A: Reconciling raw controller triggers into verified good-garment output

Suppose final assembly operation OP-040 increments a PLC register whenever the thread-trimmer sequence finishes. Over an illustrative 8-hour shift, a naive MQTT counter records 12,400 raw triggers. Applying GFT's semantic contract across **Exhibit 34**, p. 128 reconciles that count to ISO 22400-2 GoodQuantity: [S04] [S08] [S15] [S40]

1. **Opening raw trigger count:** 12,400 raw PLC trimmer triggers read via MODBUS polling. [S09]
2. **Bounce and dry-run exclusion (-310 counts):** Requiring a 4.0-second minimum seam duration and thread tension removes 310 double-taps and unthreaded test cycles (12,400 - 310 = 12,090 valid cycles). [S04]
3. **Maintenance mode exclusion (-240 counts):** Filtering cycles where OPC UA MachineMode is MAINTENANCE removes 240 adjustment cycles (12,090 - 240 = 11,850 production cycles). [S01] [S04]
4. **Duplicate replay suppression (-380 counts):** Reconciling Sparkplug sequence numbers and source timestamps after a 25-minute wireless drop-out suppresses 380 duplicate flushed packets (11,850 - 380 = 11,470 unique cycles). [S13] [S32] [S40]
5. **Inline quality disposition binding (-370 counts):** Binding inline inspection and RFID bundle checks into EPCIS ObjectEvent records routes 260 pieces to rework and rejects 110 pieces as scrap, subtracting 370 pieces (11,470 - 370 = 11,100 verified good garments). [S08] [S15] [S25]

The naive connection overstates good output by 12,400 - 11,100 = 1,300 pieces, an error of 1,300 / 11,100 = 11.71 percent (about 11.7 percent). [S08]

Part B: Bandwidth reduction from AAS separation and Sparkplug report-by-exception

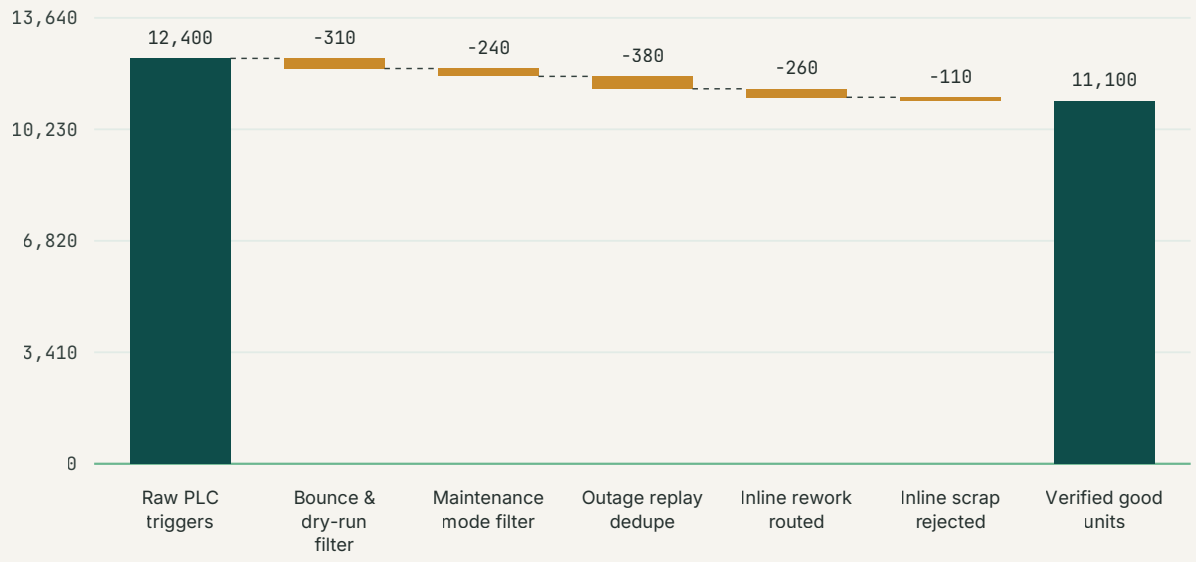
Compare illustrative shift payload volume across the 40 stations over 28,800 seconds:

- **Naive verbose JSON polling:** Polling 40 stations every 2 seconds (0.5 Hz) yields 40 x (28,800 / 2) = 576,000 messages. Repeating static strings in a 650-byte JSON payload under RFC 8259 produces 576,000 x 650 = 374,400,000 bytes (374.4 megabytes per shift). [S05] [S0804]
- **Governed AAS + Sparkplug contract:** Publishing two 1,200-byte DBIRTH certificates per station (40 x 2 x 1,200 = 96,000 bytes) referencing external AAS submodels (S07), plus 14,000 report-by-exception DDATA messages (12,090 valid cycles plus 1,910 state transitions) at 96 bytes each (1,344,000 bytes), totals 1,440,000 bytes (1.44 megabytes), an illustrative reduction of 99.62 percent. [S05] [S07] [S40]

The arithmetic reconciliation from Part A is charted in **Exhibit 38**, p. 137, showing each deduction between raw PLC triggers and verified ISO 22400-2 `GoodQuantity`. [S08] [S15] [S40]

EXHIBIT 38 / ILLUSTRATIVE GARMENT CYCLE COUNTS PER 8-HOUR SHIFT

Semantic filtering reconciles raw edge messages into verified good-garment output



ILLUSTRATIVE [S08] [S15] [S40] Illustrative worked-example reconciliation: 12,400 raw triggers minus 310 mechanical bounce/dry runs, minus 240 maintenance-mode cycles, minus 380 duplicate outage-replay packets (S40), minus 260 rework and 110 scrap dispositions (S15) equals 11,100 verified good units for ISO 22400 GoodQuantity (S08). Not a measured GFT factory result.

As **Exhibit 38**, p. 137 shows, bounce and maintenance filters derive from OPC UA state modelling, replay deduplication derives from Sparkplug session rules, and rework or scrap deductions derive from EPCIS disposition binding and ISO 22400-2 definitions. [S04] [S08] [S15] [S40] Although the counts are illustrative, the accounting structure applies to every discrete textile line: without state, session and disposition contracts, raw counters distort performance. [S08]

WARNING

Unresolved constraint: Silent semantic drift during unrecorded maintenance changes

Semantic interoperability degrades if a technician replaces a sensor or edits PLC scaling during a breakdown without updating the contract registry. [S17] [S24] Replacing a dyeing water metre of 1 litre per pulse with a 10-litre-per-pulse unit without updating the OPC UA scale under ISO/IEC 27001 change control leaves Sparkplug reporting `GOOD` quality while water intensity drops tenfold in an illustrative failure mode. [S24] [S27] [S40] Resolving this requires daily historian mass-balance checks and a maintenance work-order gate verifying OPC UA scaling and AAS calibration submodels after hardware replacement. [S04] [S07] [S24]

Interoperability in the connected textile factory is neither a single vendor's platform licence nor an unmonitored broker connection. By assigning industrial object modelling to OPC UA, constrained transport to MQTT, session state to Sparkplug, durable asset descriptions to the Asset Administration Shell, and visibility event sharing to EPCIS within an ISA-95 and IEC 62264 frame, the factory establishes one governed language from edge to dashboard. [S01] [S02] [S04] [S05] [S07] [S15] [S40] The next chapter examines how the historian preserves this evidence over time so raw observations, events, KPI versions and corrections remain distinguishable and replayable.

A wire connects two machines, but only a shared grammar lets a factory understand what they did.

CHAPTER SOURCE NOTES

S01 International Society of Automation. 2025-04-10.

S04 OPC Foundation. 2025-10-22.

S0804 Internet Engineering Task Force and RFC Editor. 2017-12.

S08 International Organization for Standardization. 2014-01.

S05 OASIS Open. 2019-03-07.

S40 Eclipse Foundation. 2022-11-16.

S07 Industrial Digital Twin Association. 2024-06.

S15 GS1 AISBL. 2022-06.

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

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

S32 World Bank Enterprise Surveys. 2025.
S41 International Organization for Standardization. 2016-10.
S09 Modbus Organization. 2012-04-26.
S10 IO-Link Community. 2024-06.
S02 International Electrotechnical Commission. 2013-05-22.
S0802 OPC Foundation. 2025-10-22.
S14 World Wide Web Consortium. 2023-12-05.
S0801 World Wide Web Consortium. 2023-12-05.
S17 National Institute of Standards and Technology. 2023-09.
S30 MDPI, Sustainability. 2023-08-19.
S0803 OASIS Open. 2019-03-07.
S33 International Telecommunication Union. 2024-11-10.
S13 Internet Engineering Task Force and RFC Editor. 2010-06.
S24 International Organization for Standardization and International Electrotechnical Commission. 2022-10.
S29 MDPI, IoT. 2024-09-28.
S12 IEEE Standards Association. 2020-06-16.
S31 MDPI, Sensors. 2020-04-28.
S03 Plattform Industrie 4.0. 2018-08-09.
S35 United Nations Industrial Development Organization. 2019.
S25 Springer Nature, International Journal of Advanced Manufacturing Technology. 2016-09-05.
R03 SODECOTON. Publication date unrecorded; inherited source-text check.
S39 European Parliament and Council of the European Union, Official Journal of the European Union. 2024-06-28.
S28 Elsevier, Robotics and Computer-Integrated Manufacturing. 2013-04.
S06 International Organization for Standardization. 2021-10.
S26 MDPI, Electronics. 2025-09-18.
S27 MDPI, Applied Sciences. 2024-06-21.

GREEN FUTURE TECH

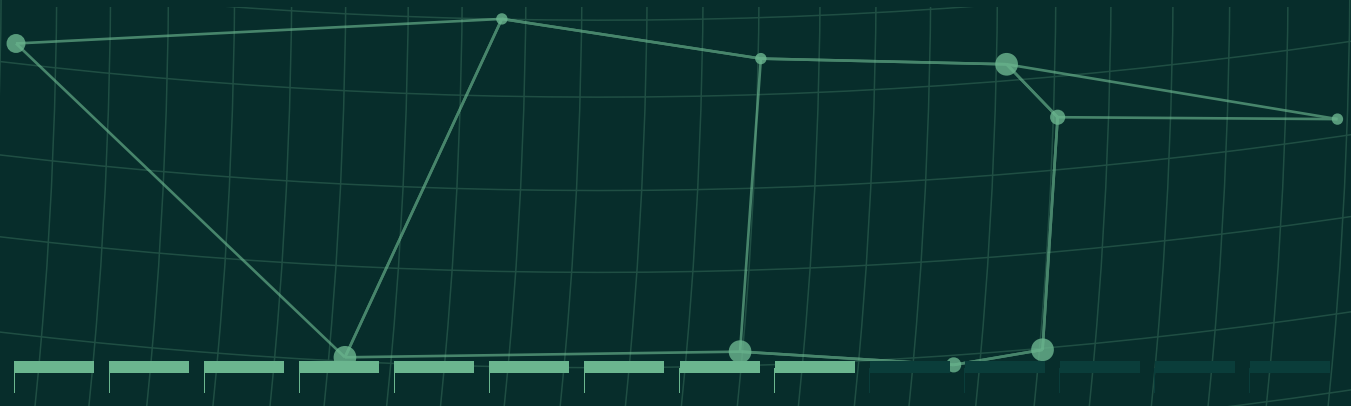
LIMITED / HONG KONG SAR

09

PART III: ONE GOVERNED DATA SPINE / SYSTEM OF RECORD

The historian must preserve context, quality and provenance

The factory's data store is an evidence system, not a growing pile of values, so raw observations, contextual events, KPI results and corrections remain distinguishable and replayable.



CHAPTER 09
DECISION TEST

The factory's data store is an evidence system, not a growing pile of values, so raw observations, contextual events, KPI results and corrections remain distinguishable and replayable.

A factory historian earns its authority as a system of record only when it stores raw machine observations, contextualised operational events, versioned key performance indicator calculations and audited human corrections in distinct, traceable layers rather than overwriting values inside one mutable table. [S08] [S15] [S24] [S0902] In a connected textile plant operating under intermittent grid power and variable wide-area network links, telemetry rarely arrives in an uninterrupted chronological stream. [S32] [S40] Gateways buffer readings during blackouts, wireless sensors experience clock drift or radio contention, and supervisors reclassify defective garment bundles after physical re-inspection. [S26] [S29] [S31] If the central store treats arrival time as occurrence time, silently interpolates across communication gaps, or lets a manual adjustment erase the original sensor reading, management cannot reconstruct why a production decision was taken. [S0901] [S0903] Green Future Tech Limited (GFT), Hong Kong SAR, therefore specifies the historian as an engineered chain of evidence that preserves context, quality state and provenance from the machine edge to the management review. [S08] [S0904]

This discipline directly supports GFT's three-part Cameroon textile programme (feasibility study, performance-management system design and digital strategy, prepared 31 January 2024 with source-refresh notes retained in 2026). Cameroon's inherited sector baseline records roughly 95 percent of cotton fibre exported raw and about 5 percent processed locally, while national policy targets 50 percent local processing by 2030. [R06] [R07] Closing that gap requires plants that prove control over utilisation, first-pass quality and resource intensity from their first shifts, avoiding the unmeasured underutilisation that historically cost domestic mills billions of CFA francs (XAF). [R12] Where a plant starts with a manual or spreadsheet system of record for daily shift huddles, transitioning to automated capture must not sever historical continuity or alter indicator formulas. [S08] [S30] Preceding chapters defined how operational meaning travels across protocols and semantic models; this chapter specifies how the storage spine persists, reconciles, retains and recovers that meaning. [S04] [S07] [S15] [S40]

9.1

Raw observations remain separate from interpreted events

An industrial data spine fails its first audit test when it conflates what a physical instrument measured with what execution logic concluded. [S01] [S15] A current transformer on a ring-spinning frame, an optical counter on a sewing station, or a level sensor on a dye vat produces a raw observation: a value bound to a unit, a source timestamp, a receipt timestamp, a quality code and a point identifier. [S27] [S29] [S0903] That observation does not know which customer order was running, which crew was assigned, or whether the machine was in maintenance warm-up rather than commercial production. [S01] [S02] Those attributes belong to the interpreted event layer, where raw signals join manufacturing operations management context from Level 3 schedules, work-in-process (WIP) scans and quality specifications. [S02] [S25] [S28] Storing both layers separately ensures that if a work-order mapping table was misconfigured at shift handover, engineers can re-run the contextualisation rule against untouched raw observations without losing physical telemetry. [S24] [S0902]

9.2 IMMUTABLE RAW CAPTURE PRESERVES THE ORIGINAL PHYSICAL SIGNAL

At the lowest persistence tier, the raw observation store acts as an append-only ledger of instrument readings as they crossed the acquisition boundary. [S29] [S0903] In the OPC Unified Architecture (OPC UA) Historical Access specification (OPC 10000-11), RawData and HistoricalDataNodes are queried through their SourceTimestamp, while allowing clients to inspect both SourceTimestamp and ServerTimestamp (BOTH) so transport delay never masks the instant of physical measurement. [S04] [S0903] Each raw record retains the asset-hierarchy point identifier, engineering value, unit symbol, source clock reading, receiving timestamp and acquisition status code, with no downstream filter or script permitted to overwrite it. [S01] [S24] [S0902]

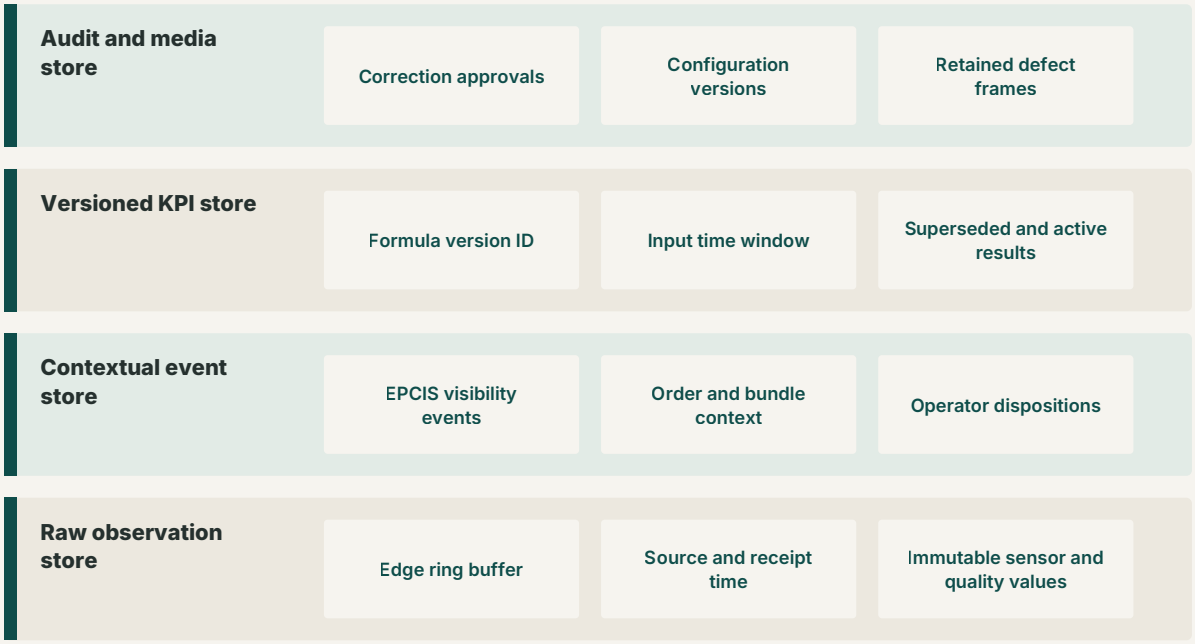
Interpreted production events belong in a separate contextual event store. [S15] [S16] GS1's Electronic Product Code Information Services (EPCIS Release 2.0) standard enables disparate applications to create and share visibility event data within and across enterprises, capturing the what, when, where and why of object movements. [S15] In a denim or garment plant using radio-frequency identification (RFID) at cutting and sewing checkpoints, an antenna read is a raw observation that a tag passed within radio range. [S25] [S31] The interpreted event asserts that bundle B-1042 completed sewing operation OP-20 on Line 2. [S25] [S28] Because proximity reads can bounce at workstation boundaries, transforming raw RFID reads into an EPCIS visibility event relies on a versioned state-machine rule that can be replayed cleanly if the raw log remains intact. [S15] [S31] [S0902]

9.3 CONTEXTUAL INTERPRETATION APPLIES VERSIONED OPERATIONAL RULES

The separation extends upward to calculated key performance indicators (KPIs) and outward to media and audit records, forming four distinct persistence stores across the factory spine, as shown in **Exhibit 39**, p. 143. [S08] [S15] [S24] [S29]

EXHIBIT 39 /

Four stores preserve different kinds of evidence



ILLUSTRATIVE [S08] [S15] [S24] [S29] Illustrative GFT historian store architecture informed by ISO 22400 KPI structure (S08), EPCIS visibility event sharing (S15), ISO/IEC 27001 change control (S24) and edge database persistence (S29).

As **Exhibit 39**, p. 143 illustrates, each store enforces its own mutability contract. [S08] [S15] [S24] [S29] The raw observation store records machine counters, metre totals and sensor readings at the edge and central archive. [S14] [S29] The contextual event store binds observations to ISA-95 and IEC 62264 Level 3 entities: work centres, orders, lots, shifts and EPCIS milestones. [S01] [S02] [S15] Above it, the versioned KPI store holds reproducible ISO 22400 results, while the audit and media store preserves configuration releases, override approvals and selected inspection images linked by reference identifiers. [S08] [S24] [S26] Because **Exhibit 39**, p. 143 is an illustrative architectural model, it defines logical boundaries rather than prescribing a specific database product or claiming measured plant throughput. [S08] [S29]

9.4 PROVENANCE DISTINGUISHES MASTER DATA FROM TELEMETRY LINEAGE

Engineers must also distinguish which standards govern provenance. [S41] [S0904] The World Wide Web Consortium (W3C) PROV overview defines provenance as information about entities, activities and people involved in producing data, used to assess quality, reliability or trustworthiness. [S0904] Within industrial standards, however, ISO 8000-120:2016 specifies provenance requirements strictly for master data consisting of characteristic data, supplementing ISO 8000-110. [S41] Characteristic master data covers durable asset properties, material grade specifications and nameplate parameters such as those inside an Asset Administration Shell (AAS) submodel, not telemetry streams or Sparkplug states. [S07] [S40] [S41]

GFT therefore applies ISO 8000-120 where the factory exchanges characteristic master data, while specifying telemetry transformation and KPI lineage as GFT's own architecture rule governed under ISO/IEC 27001 change control and OPC 10000-11 modification tracking. [S07] [S24] [S41] [S0902] Furthermore, OPC 10000-11 Section 4.4 notes that multiple historical events can share a timestamp; a historian storing events must preserve the EventId generated during event processing and must never allow a modification to alter that EventId. [S0902]

9.5

Quality state is data, not an annotation for later

A numeric value stripped of its acquisition health misleads automated dashboards and supervisors alike. [S31] [S40] If a steam-flow metre freezes because a serial converter lost power, or if a warehouse humidity node transmits with an unsynchronised clock, storing the bare number as an ordinary observation corrupts downstream averages and alerts. [S31] [S40] [S0901] Quality state cannot be an optional comment added during monthly data cleaning. [S08] [S24] It must travel as a mandatory, machine-readable attribute inside every observation and event envelope, governing whether the historian may feed that value into a live tile, alert or KPI calculation. [S08] [S40] [S0901]

9.6 SESSION STATE PREVENTS STALE VALUES FROM MASQUERADING AS LIVE READINGS

In MQTT publish-subscribe architectures, report-by-exception publishes updates only when values change. [S05] [S40] Without session-state management, a subscribing historian cannot tell whether a quiet metric is constant or offline. [S40] The Eclipse Sparkplug Specification (Version 3.0.0) resolves this through birth and death certificates (NBIRTH, DBIRTH, NDEATH, DDEATH). [S40] When an Edge Node disconnects and its NDEATH message is published, Sparkplug specifies that all metrics in the associated NBIRTH and DBIRTH messages under that Edge Node are immediately implied to be STALE. [S40]

The historian must persist this transition as an explicit time-stamped state change. [S40] [S0901] In OPC 10000-11 Section 4.3, when data collection breaks and recovery reveals a discontinuity between the last stored SourceTimestamp and the newly arrived initial SourceTimestamp, the historian must record a bad status at the break timestamp and report gaps to querying clients with Bad_DataLost. [S0901] Similarly, if an event subscription buffer overflows during an interruption, the historian must record EventQueueOverflowEventType or Bad_DataLost so queries never assume an incomplete event stream is whole. [S0901]

9.7 ACQUISITION DIAGNOSTICS RECORD CLOCK AND RADIO UNCERTAINTY AT THE SOURCE

Wireless sensor networks and shop-floor RFID deployments depend on node time synchronisation and spectrum-sharing discipline to protect data integrity and avoid uncertainty in sensed data. [S31] If an edge gateway loses lock with its Network Time Protocol (NTPv4, capable of tens-of-microseconds accuracy on fast LANs with modern workstations) or Precision Time Protocol (IEEE 1588-2019, capable of sub-microsecond accuracy with minimal resources) reference, its local clock drifts relative

to neighbouring machines. [S12] [S13] [S31] Unsynchronised readings may still count shift totals, but they cannot correlate sub-second stoppages across linked machines. [S12] [S31]

Computer-vision inspection cells carry a similar physical dependency: textile defect-detection research shows that lighting variations and background noise significantly degrade machine-learning detection accuracy. [S26] When illumination dims or a lens fouls, the historian must capture the resulting degraded-acquisition flag alongside the inspection event so a drop in flagged defects is recognised as a sensor fault rather than improved cloth quality. [S26] [S40]

9.8 EXPLICIT GAP MARKERS PROTECT DOWNSTREAM QUERIES FROM SILENT INTERPOLATION

To standardise how services and dashboards treat degraded readings, GFT derives a six-state quality dictionary from the underlying protocols and security standards, presented in **Exhibit 40**, p. 145. [S24] [S31] [S40] [S0901]

EXHIBIT 40 /

A quality flag tells the reader whether to act

Quality state	Trigger condition	Historian rule	Permitted dashboard use
GOOD_LIVE	Active session and synchronised clock	Persist raw value and source time	Eligible for live display and KPI calculation
STALE_OFFLINE	Sparkplug NDEATH or DDEATH received	Retain last value with offline marker	Dim tile and suppress threshold alerts
UNSYNCHRONISED	Clock stratum lost or radio contention	Store with receipt time and skew warning	Exclude from sub-second correlation
BAD_DATA_LOST	Gap detected after collection break	Insert explicit gap boundary record	Block linear interpolation across gap
ESTIMATED	Approved fallback calculation applied	Store separately from raw observation	Display with provisional badge only
CORRECTED	Authorised update with reason and approver	Append new version and retain original	Use latest approved version with audit link

DERIVED [S24] [S31] [S40] [S0901] Derived quality-state dictionary combining Sparkplug session and stale-metric semantics (S40), wireless clock and spectrum risks (S31), OPC UA historical gap reporting (S0901) and controlled correction audit rules (S24).

The derived rules in **Exhibit 40**, p. 145 translate protocol diagnostics into deterministic storage and display rules. [S24] [S31] [S40] [S0901] Only GOOD_LIVE observations enter real-time threshold alerts and unflagged KPI aggregations. [S08] [S40] When STALE_OFFLINE or BAD_DATA_LOST is active, the historian blocks linear interpolation across the gap, preventing queries from fabricating synthetic production across a power outage. [S32] [S40] [S0901] Where ESTIMATED or CORRECTED values enter under ISO/IEC 27001 change control, they are stored as versioned records so auditors can compare the raw stream against the adjusted operational view. [S24] [S0902] Because

Exhibit 40, p. 145 is a derived governance dictionary, edge gateways must map device-specific fault codes (such as Modbus exception responses) into these six canonical states. [S09] [S40]

> **GFT NOTE**
>
> Never configure a historian or SQL view to replace missing or stale telemetry with ZERO or with the prior valid reading (LOCF, last observation carried forward) inside KPI views. [S08] [S40] [S0901] Mapping a disconnected camera to zero defects inflates quality yield, while carrying forward a flow-metre reading across a blackout fabricates water or steam use. [S26] [S27] [S32] Persist explicit STALE_OFFLINE and BAD_DATA_LOST markers, and report valid-data coverage alongside every aggregated metric. [S08] [S40] [S0901]

9.9

The edge buffer and central store have different duties

An architecture that relies on a single central or cloud database to ingest machine signals directly loses evidence whenever site power or external links fail. [S14] [S16] [S32] In the World Bank's Cameroon 2024 Enterprise Survey (fieldwork June 2024 to February 2025 across formal private non-agricultural firms), 93.3 percent of surveyed firms experienced electrical outages, averaging 10.4 outages in a typical month. [S32] Although this covers all surveyed firms rather than textile mills alone, it confirms that power and communication interruptions are routine operating conditions. [S32] GFT therefore divides persistence duties between a local edge buffer in the production security zone and a central historian in the operations zone. [S16] [S19] [S29]

9.10 LOCAL PERSISTENCE SUSTAINS LINE OPERATIONS DURING POWER AND LINK OUTAGES

The edge buffer's duty is uninterrupted local capture and immediate line-level calculation. [S16] [S29] ETSI's Multi-access Edge Computing (ETSI GS MEC 003) runs software applications on virtualised infrastructure close to the network edge, while W3C's Web of Things (WoT Architecture 1.1) describes an industrial edge device collecting selected controller data for local analysis and downstream monitoring. [S14] [S16] A peer-reviewed industrial energy-monitoring system demonstrates this pattern: when the edge receives metre data, it stores the reading in a local database server and computes nine additional analytical parameters locally before transfer. [S29]

Writing raw observations and local parameters to non-volatile storage at the edge node (backed by uninterruptible power) keeps line huddles, local human-machine interfaces (HMIs) and inspection cells running when the plant network or WAN is down. [S17] [S26] [S29] An illustrative edge gateway polling 50 registers once per second (1 Hz) with an illustrative 120-byte compressed envelope writes 6,000 bytes per second, or 518.4 megabytes per day (50 * 120 * 86,400 = 518,400,000 bytes). A 32-gigabyte industrial solid-state partition holds weeks of illustrative telemetry and cropped defect images without dropping primary evidence. [S26] [S29] [S40]

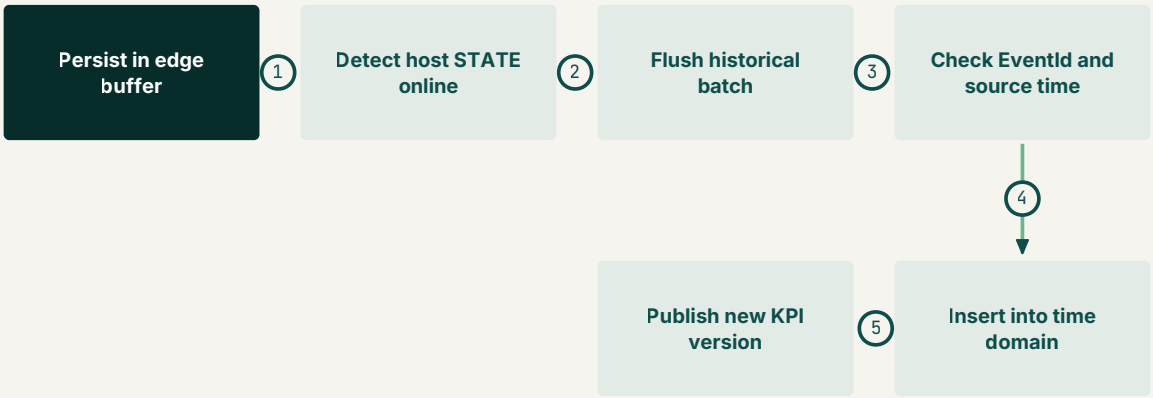
9.11 CENTRAL RECONCILIATION ORDERS REPLAYED HISTORY BY EVENT IDENTITY AND SOURCE TIME

While the edge store optimises for write survivability, the central historian optimises for cross-line reconciliation, long-term query, enterprise-control integration under ISA-95 and IEC 62264, and governed reporting. [S01] [S02] [S29] Under Sparkplug 3.0, when a Primary Host Application publishes an online STATE message, the Edge Node resumes live publishing and flushes historical data stored while offline. [S40]

If the central historian indexes messages by broker arrival time (ServerTimestamp), a three-hour flush from Line 1 will be interleaved at the reconnection instant rather than placed in the afternoon window when the garments were sewn. [S40] [S0903] The central ingestion service must therefore route replayed batches through the identity and source-time reconciliation workflow shown in **Exhibit 41**, p. 147. [S24] [S29] [S40] [S0902]

EXHIBIT 41 /

Late data is reconciled by identity and source time



1. WAN outage
2. Link restored
3. Replay stream
4. Unique record
5. Close gap

ILLUSTRATIVE [S24] [S29] [S40] [S0902] Illustrative store-and-forward replay sequence based on edge persistence (S29), Sparkplug host STATE and historical flush rules (S40), immutable OPC UA EventId tracking (S0902) and controlled recalculation (S24).

As **Exhibit 41**, p. 147 shows, the edge node writes every observation and event locally with its SourceTimestamp, sequence counter and EventId. [S29] [S40] [S0902] Once the primary host publishes its online STATE message, the edge node flushes its queue in bounded batches so replay traffic does not starve live operational telemetry. [S17] [S40] The central service checks the composite key (asset_id, point_id, SourceTimestamp, EventId), inserts unique records into the OPC 10000-11 time domain, clears BAD_DATA_LOST markers where continuity is proven, and triggers a

versioned KPI recalculation. [S08] [S40] [S0901] [S0903] Because **Exhibit 41**, p. 147 is an illustrative workflow, it defines the required logical steps without promising a fixed replay duration, which depends on link bandwidth and outage length. [S14] [S40]

9.12 DEDUPLICATION AND SEQUENCE CHECKS PREVENT DOUBLE COUNTING DURING BUFFER FLUSHES

If an acknowledgement packet drops during a historical flush, the edge node retransmits the unacknowledged batch. [S05] [S40] To prevent double-counting output, the historian enforces two idempotency rules. [S0902] [S0903] For scalar observations (`HistoricalDataNodes`), an identical (`asset_id, point_id, SourceTimestamp, source_version`) tuple is treated as an idempotent retry that updates only the replication log. [S0901] [S0903] For discrete production and visibility events (`HistoricalEventNodes` and `EPCIS` events), uniqueness is enforced on the immutable `EventId` generated at the edge (`OPC 10000-11` Section 4.4). [S15] [S0902] Machine cycle counters are also published as monotonic cumulative totals alongside delta increments so any unrecovered gap still reveals the exact number of elapsed cycles, flagged as `ESTIMATED` until verified. [S08] [S24] [S0901]

9.13

KPI results are reproducible calculations

When a plant manager reviews weekly Overall Equipment Effectiveness (`OEE`), first-pass yield, or kilowatt-hours per kilogram of fabric, that number is the output of a formula evaluated over selected observations, events, time boundaries and exclusion rules. [S08] [S29] If dashboards and spreadsheets compute KPIs independently with drifting queries, review meetings stall over conflicting numbers. [S08] [S30] The historian must store KPI results as versioned calculation artefacts locked to a governed dictionary. [S08] [S24]

9.14 STORED INDICATORS BIND VALUES TO FORMULA VERSIONS AND INPUT WINDOWS

`ISO 22400-2:2014` presents manufacturing KPIs through their formula, source elements, time behaviour, unit or dimension and related characteristics. [S08] In GFT's architecture, every stored KPI record must carry five provenance attributes: 1. The indicator identifier and version (`KPI-OEE-LINE-V1.2`, specifying numerator and denominator elements under `ISO 22400-2`). [S08] 2. The target entity in the `ISA-95` hierarchy and the operational scope (order, lot or shift). [S01] [S02] 3. The evaluation window (`window_start_utc` and `window_end_utc` on `SourceTimestamp`) and execution time (`computed_at_utc`). [S08] [S0903] 4. The input coverage ratio (valid `GOOD_LIVE` and approved `CORRECTED` events divided by expected events) and worst-case quality flag. [S08] [S40] [S0901] 5. A hash or snapshot pointer to the input event set and the calculation version (`calc_version = 1, 2, ...`). [S24] [S0904]

METHOD

Reproducible KPI storage and reconciliation boundary

To verify that any stored KPI result K_{v} (version v) is reproducible, the historian enforces four checks based on ISO 22400-2 ([^S08]), OPC 10000-11 ([^S0902], [^S0903]) and ISO/IEC 27001 ([^S24]):

1. **Time domain:** Input elements $E_1, \dots, E_m$ are queried over $[t_{\text{start}}, t_{\text{end}})$ using SourceTimestamp ([^S0903]), never ServerTimestamp.
2. **Formula binding:** The record stores $K_v = f_{\text{id,ver}}(E_1^{(v)}, \dots, E_m^{(v)})$ referencing the immutable KPI dictionary formula ([^S08]).
3. **Coverage check:** Source coverage $C_v = N_{\text{valid}} / N_{\text{expected}}$ accompanies K_v; if C_v is below threshold, K_v is marked PROVISIONAL ([^S40], [^S0901]).
4. **Versioned recalculation:** When delayed buffers flush ([^S29], [^S40]) or an authorised correction enters ([^S24], [^S0902]), the engine writes $K_{v+1} = K_v + \Delta K_{\text{replay}} + \Delta K_{\text{dedup}} + \Delta K_{\text{corr}}$ while retaining K_v as SUPERSEDED ([^S08], [^S24]).

9.15 CONTROLLED CORRECTIONS APPEND NEW CALCULATION VERSIONS WITHOUT ERASING HISTORY

As OPC 10000-11 Section 4.4 notes, historical data may require updating because of a failed sensor, communication drop, calibration error, or offline inspection and laboratory findings arriving after the shift. [S0902] A quality supervisor may re-inspect 30 fabric rolls flagged by a line-scan camera under variable lighting ([^S26], [^R29]), confirm 22 defective rolls, and release 8 rolls as first-quality goods. [S26] [S0902]

When such an update enters, OPC 10000-11 requires recording the original value and who made the change (HistoryUpdateDetails) and generating an audit event (AuditHistoryValueUpdateEventType or AuditHistoryEventUpdateEventType). [S0902] Under ISO/IEC 27001, GFT requires an authenticated user identity, reason code, evidence link and role approval. [S24] Rather than overwriting Version 1 of a shift KPI, the historian appends Version 2 and marks Version 1 SUPERSEDED, preserving what was visible at the morning huddle alongside the verified weekly record. [S08] [S24] [S0902]

9.16 RECONCILIATION ARITHMETIC EXPOSES THE DIFFERENCE BETWEEN LIVE AND REPLAYED SHIFTS

An illustrative shift-reconciliation case on a garment line shows how late buffer replay, duplicate filtering and audited quality corrections combine inside the KPI store. [S08] [S25] [S29] [S40]

WORKED EXAMPLE: SHIFT KPI RECONCILIATION ACROSS AN OUTAGE AND QUALITY REVIEW

Consider an illustrative eight-hour shift (06:00 to 14:00) on Sewing Line A with 450 net operating minutes (27,000 seconds) and an illustrative target takt time of 22.5 seconds per garment, giving a reference output of 1,200 garments (27,000 / 22.5 = 1,200). [S08] Bundles are tracked via RFID ([^S25], [^S28]), buffered locally at the edge ([^S29]), and published via Sparkplug ([^S40]).

From 11:30 to 13:30 (120 minutes), an internal switch fault isolates the line from the central server room while sewing continues on backup power ([^S32]). At 14:00, before switch recovery, the central historian computes Version 1 (PROVISIONAL) ([^S08]):

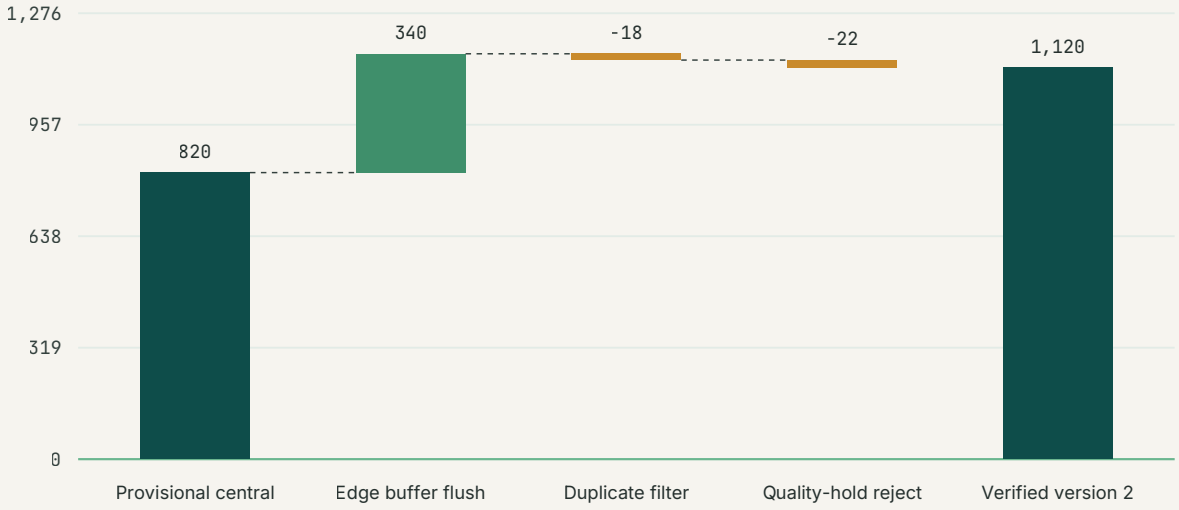
- **Provisional central record (Version 1 at 14:00):** The central store holds 820 good garments from the connected intervals ([^S08], [^S40]). With 120 of 450 minutes marked STALE_OFFLINE ([^S40]) and BAD_DATA_LOST ([^S0901]), source coverage is 330 / 450 = 73.3 percent (illustrative) and provisional performance is 820 / 1,200 = 68.3 percent (PROVISIONAL).

At 14:15, the host publishes its online STATE message and the Edge Node flushes its buffer ([^S40]), followed at 14:40 by an audited quality disposition ([^S24], [^S26]): 1. **Edge buffer flush:** The gateway replays 340 garment-completion events captured with SourceTimestamp between 11:30 and 13:30 ([^S29], [^S40], [^S0903]). 2. **Idempotent deduplication:** Checking immutable EventId values ([^S0902]) filters out 18 duplicate RFID reads across the drop/reconnect boundaries ([^S25], [^S31]), netting 340 - 18 = 322 added garments. 3. **Audited quality-hold correction:** At 14:40, the quality supervisor logs an approved correction ([^S24], [^S0902]) on a lot flagged during camera lighting drift ([^S26]), reclassifying 22 garments to rework/reject. 4. **Reconciled KPI (Version 2 at 14:45):** The historian computes Version 2 ([^S08]): 820 + 340 - 18 - 22 = 1,120 verified good garments (illustrative). Coverage reaches 450 / 450 = 100.0 percent, first-pass quality is 1,120 / 1,142 = 98.1 percent (illustrative), and output against reference is 1,120 / 1,200 = 93.3 percent (VERIFIED). Both Version 1 and Version 2 remain stored ([^S08], [^S24]).

The signed contributions from this illustrative shift case appear in **Exhibit 42**, p. 151, reconciling the provisional central count to the verified closing balance. [S08] [S24] [S29] [S40]

EXHIBIT 42 / ILLUSTRATIVE GARMENT COUNT

Replayed edge records and audited corrections reconcile shift output



ILLUSTRATIVE [S08] [S24] [S29] [S40] Illustrative shift reconciliation case showing how a provisional central count of 820 garments is updated by +340 buffered RFID reads (S29, S40), -18 duplicate scans, and -22 quality-hold rejects (S08, S24) to reach 1,120 verified good garments.

As **Exhibit 42**, p. 151 shows, each step between `Version 1` (820 illustrative garments) and `Version 2` (1,120 illustrative garments) maps to a distinct mechanism: edge persistence (+340), `EventId` deduplication (-18) and audited quality disposition (-22). [S08] [S24] [S29] [S40] [S0902] Because **Exhibit 42**, p. 151 is an illustrative method case, it demonstrates reconciliation logic rather than reporting live factory output. [S08]

9.17

Retention follows evidence need and storage cost

A historian cannot keep every raw camera frame, millisecond sample, visibility event and KPI record for the same duration on the same storage tier. [S26] [S29] [S36] Retaining uncompressed vision streams indefinitely exhausts storage budgets, whereas purging production and resource events after thirty days destroys the evidence required for quality claims, life-cycle carbon accounting and digital product passports. [S15] [S26] [S36] [S39] Retention policy must classify records by evidentiary obligation, reconstruction cost and byte volume rather than applying one uniform retention period. [S24] [S36] [S39]

9.18 PRODUCT LIFE-CYCLE AND PASSPORT RULES DEMAND GRANULAR EVENT HISTORIES

Sustainability and traceability standards show why monthly plant rollups cannot replace batch and order histories. [S36] [S38] [S39] The `GHG Protocol Product Life Cycle Accounting and Reporting Standard` quantifies greenhouse-gas emissions and removals associated with a specific product across its life cycle. [S36] Similarly, `ISO 14046:2014` specifies life-cycle water-footprint assessment for products, processes and organisations, while `ISO 50001:2018` governs energy management

systems and baselines. [S37] [S38] Allocating electricity ([^S29]) or water ([^S27]) to a garment lot requires preserving the production denominators (good units, rework and run windows) at the lot and order boundary ([^S08], [^S15]). [S36] [S38]

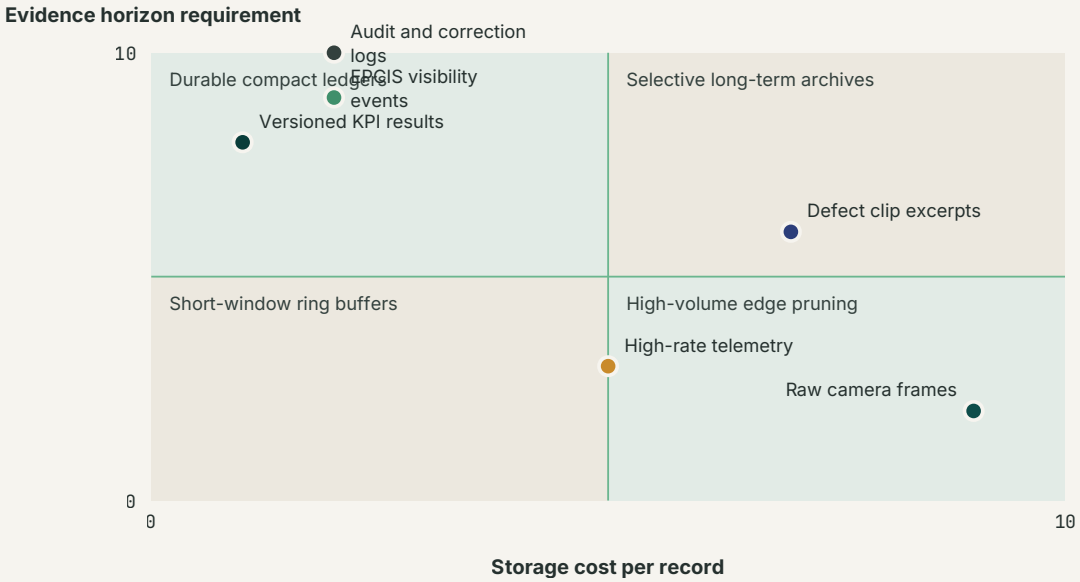
In the European Union, Regulation (EU) 2024/1781 (ESPR) defines a digital product passport as data specific to a product, specified in the applicable delegated act under Article 4 and accessible electronically through a data carrier. [S39] Preserving export-ready evidence requires keeping EPCIS visibility events, lot genealogy and batch resource allocations across commercial horizons rather than discarding them while garments remain in the supply chain. [S15] [S36] [S39]

9.19 TIERED STORAGE SEPARATES HIGH-VOLUME INSPECTION MEDIA FROM DURABLE EVENT LEDGERS

At the high-volume extreme, industry inspection literature reports line-scan fabric inspection running at up to 60 metres a minute across more than 40 defect types. [R29] Archiving every raw frame from defect-free cloth wastes capacity without adding analytical value. [S26] [R29] GFT therefore positions each data class by storage cost per record versus required evidence horizon, as shown in **Exhibit 43**, p. 152. [S08] [S15] [S24] [S26] [S36] [S39]

EXHIBIT 43 / ILLUSTRATIVE INDEX (0 TO 10)

Retention is assigned by evidence class rather than one duration



ILLUSTRATIVE [S08] [S15] [S24] [S26] [S36] [S39] Illustrative positioning of six factory data classes by storage cost per record (x-axis) and long-term audit or life-cycle evidence horizon (y-axis), informed by S08, S15, S24, S26, S36 and S39. Actual retention durations require factory policy approval.

As **Exhibit 43**, p. 152 illustrates, raw defect-free camera frames (x = 9, y = 2 illustrative index) stay in an edge circular buffer for short-term diagnostics before pruning, while cropped defect images (x = 7, y = 6) are indexed by EventId in selective archives for claim review and model retraining. [S16] [S26] [S0902] High-rate telemetry (x = 5, y = 3) remains at full resolution in warm storage before compaction into validated rollups and state boundaries. [S29] [S0901] In the upper-left quadrant (Durable compact ledgers), EPCIS visibility events (x = 2, y = 9), versioned ISO

22400 KPI results (x = 1, y = 8) and ISO/IEC 27001 / OPC 10000-11 audit logs (x = 2, y = 10) combine small byte footprints with multi-year life-cycle and compliance value. [S08] [S15] [S24] [S36] [S39] [S0902] Because **Exhibit 43**, p. 152 uses an illustrative zero-to-ten scale, exact retention durations remain a site governance decision set with legal and commercial owners. [S24] [S39]

9.20 ARCHIVAL FORMATS PROTECT LONG-TERM READABILITY AGAINST VENDOR LOCK-IN

Long-term archives fail if they depend on a proprietary historian licence that the plant may later replace. [S04] [S24] Cold-tier archives and compliance exports must be written in open, self-describing formats bundled with their ISA-95 hierarchy snapshot ([^S01]), AAS submodels ([^S07]), ISO 22400 KPI dictionary version ([^S08]), EPCIS vocabulary context ([^S15]), ISO 8000-120 master-data provenance records ([^S41]) and a SHA-256 manifest logged under ISO/IEC 27001 ([^S21], [^S24]).

9.21

Recovery includes reconciliation, not just database restoration

The final test of a historian is how it recovers from server failure, storage corruption or cybersecurity isolation. [S18] [S21] [S22] Checking only that a nightly backup script completed is insufficient: a network-accessible backup share can be encrypted by the same malware that hit the primary server ([^S21]), and restoring a noon snapshot at 18:00 leaves a six-hour gap between the database and the shop floor ([^S22], [^S40], [^S0901]). True recovery requires isolated restoration followed by evidence reconciliation. [S21] [S22]

9.22 ISOLATED BACKUPS PROTECT HISTORIAN STATE AND SCHEMA METADATA FROM COMPROMISE

CISA (Advisory AA22-110A) advises operators to test backup procedures regularly and isolate backups from network connections that could spread malware. [S21] Within an IEC 62443-3-2 zone-and-conduit design ([^S19]) aligned with NIST SP 800-82 Rev. 3 ([^S17]) and NIST CSF 2.0 ([^S18]), the backup store must sit behind a controlled conduit with separate credentials and immutable retention locks ([^S19], [^S21], [^S24]). To rebuild the system of record cleanly, the isolated backup package must include six artefacts: 1. Raw observation and contextual event partitions up to the checkpoint ([^S15], [^S29], [^S0903]). 2. The versioned ISO 22400 KPI store and formula dictionary history ([^S08]). 3. The ISA-95 hierarchy, AAS submodels and ISO 8000-120 master-data provenance tables ([^S01], [^S07], [^S41]). 4. The OPC 10000-11 modification records, quality-flag rules and ISO/IEC 27001 audit logs ([^S24], [^S0902]). 5. Edge-gateway configuration manifests, Sparkplug topic contracts and broker access lists ([^S19], [^S21], [^S40]). 6. The restoration runbook and integrity-verification scripts ([^S21], [^S22]).

9.23 POST-RESTORE RECONCILIATION VERIFIES SEQUENCE CONTINUITY AGAINST EDGE BUFFERS AND SHIFT LOGS

When a historian is restored from a backup taken at t_{backup} and brought back online at t_{restore}, it initially lacks all data from $[t_{\text{backup}}, t_{\text{restore}}]$. [S21] [S22] [S0901] Because edge buffers continued recording

machine cycles, energy metres and RFID scans during the central outage ([^S25], [^S29], [^S40]), recovery personnel execute a four-step catch-up sequence before clearing degraded-mode status:

- **Step 1 (Watermark audit):** Query each registered Edge Node's last committed SourceTimestamp and sequence number as of t_{backup} ([^S01], [^S40], [^S0903]).
- **Step 2 (Bounded edge replay):** Re-establish Sparkplug host session state ([^S40]) and replay edge-buffered observations and events ([^S29]) across $[t_{\text{backup}}, t_{\text{restore}}]$ using the EventId and composite-key deduplication rules in **Exhibit 41**, p. 147 ([^S0902]).
- **Step 3 (Manual log cross-check):** Reconcile any unbuffered workstation or laboratory records against approved paper contingency tally sheets ([^S22], [^S30]), inserting validated entries as ESTIMATED or CORRECTED with approver audit events ([^S24], [^S0902]) and marking any unrecoverable gap BAD_DATA_LOST ([^S0901]).
- **Step 4 (Versioned KPI rebuild):** Recompute affected shift and daily ISO 22400 KPIs ([^S08]) as new calculation versions linked to the recovery ticket ([^S22], [^S24]).

9.24 EXERCISED CONTINGENCY PLANS MEASURE ACTUAL RECOVERY TIME AND DATA LOSS

NIST SP 800-34 Rev. 1 distinguishes three activities: testing validates recovery capabilities, training prepares personnel for plan activation, and exercising identifies planning gaps to improve preparedness. [S22] GFT requires a live restoration and reconciliation exercise on a staging instance before operational acceptance ([^S21], [^S22], [^S30]), measuring actual Recovery Time Objective (RTO) and Recovery Point Objective (RPO) after edge replay and manual log reconciliation and logging any variance as a corrective action under ISO/IEC 27001 ([^S08], [^S22], [^S24]).

> **WARNING**
>
> **Unresolved constraint: prolonged multi-day blackouts exceeding edge UPS and clock holdover**
>
> If a multi-day blackout ([^S32]) outlasts edge backup power and a depleted real-time clock (RTC) battery while NTP/PTP servers ([^S12], [^S13]) are unreachable, an edge gateway may reboot with a reset epoch timestamp (such as 1970-01-01) while generator-backed lines are running. [S31] [S32] **Resolution required at site commissioning:** Gateways must enforce monotonic sequence counters and validate clock lock against a battery-backed local GNSS/RTC source before marking timestamps GOOD_LIVE; pre-lock records must be flagged UNSYNCHRONISED with boot-relative counters so post-outage reconciliation can anchor them against ServerTimestamp and shift sign-in records ([^S13], [^S31], [^S0901], [^S0903]).

By separating raw observations from interpreted events, persisting quality state as operational data, pairing edge buffers with identity-based central reconciliation, locking KPIs to ISO 22400 formula versions, tiering retention by evidence need, and

proving recovery through exercised reconciliation, the factory turns its historian into a trustworthy spine for the decision dashboards examined next. [S02] [S08] [S15] [S22] [S40]

An observation records what the sensor saw; an event records what the factory did; a historian earns trust only when it can prove the path between the two without erasing its corrections.

CHAPTER SOURCE NOTES

S08 International Organization for Standardization. 2014-01.

S15 GS1 AISBL. 2022-06.

S24 International Organization for Standardization and International Electrotechnical Commission. 2022-10.

S0902 OPC Foundation. 2024-10-15.

S32 World Bank Enterprise Surveys. 2025.

S40 Eclipse Foundation. 2022-11-16.

S26 MDPI, Electronics. 2025-09-18.

S29 MDPI, IoT. 2024-09-28.

S31 MDPI, Sensors. 2020-04-28.

S0901 OPC Foundation. 2024-10-15.

S0903 OPC Foundation. 2024-10-15.

S0904 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.

S30 MDPI, Sustainability. 2023-08-19.

S04 OPC Foundation. 2025-10-22.

S07 Industrial Digital Twin Association. 2024-06.
S01 International Society of Automation. 2025-04-10.
S27 MDPI, Applied Sciences. 2024-06-21.
S02 International Electrotechnical Commission. 2013-05-22.
S25 Springer Nature, International Journal of Advanced Manufacturing Technology. 2016-09-05.
S28 Elsevier, Robotics and Computer-Integrated Manufacturing. 2013-04.
S16 European Telecommunications Standards Institute. 2024-04.
S14 World Wide Web Consortium. 2023-12-05.
S41 International Organization for Standardization. 2016-10.
S05 OASIS Open. 2019-03-07.
S12 IEEE Standards Association. 2020-06-16.
S13 Internet Engineering Task Force and RFC Editor. 2010-06.
S09 Modbus Organization. 2012-04-26.
S19 International Electrotechnical Commission. 2020-06.
S17 National Institute of Standards and Technology. 2023-09.
R29 MDPI; industry inspection literature. Publication date unrecorded; inherited source-text check.
S36 GHG Protocol, World Resources Institute and World Business Council for Sustainable Development. 2011-09.
S39 European Parliament and Council of the European Union, Official Journal of the European Union. 2024-06-28.
S38 International Organization for Standardization. 2014-08.
S37 International Organization for Standardization. 2018-08.
S21 CISA and joint international cybersecurity authorities. 2022-05-09.
S18 National Institute of Standards and Technology. 2024-02-26.
S22 National Institute of Standards and Technology. 2010-11-11.

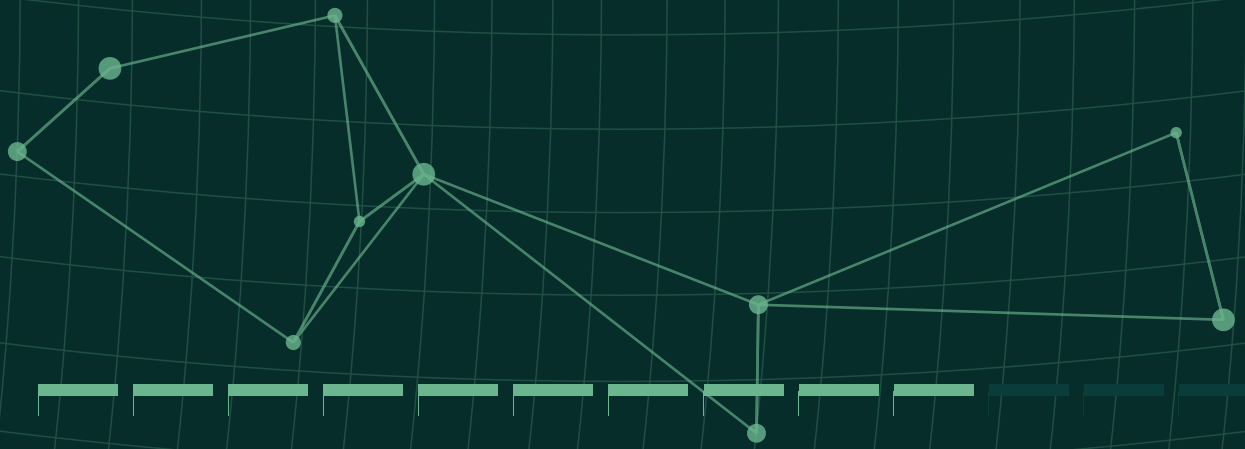
GREEN FUTURE TECH
LIMITED / HONG KONG SAR

10

PART III: ONE GOVERNED DATA SPINE / OPERATIONAL INTERFACE

A dashboard earns its place by changing a decision

The dashboard is a role-specific control surface whose freshness, definitions and next action are visible, not a gallery of independently calculated charts.



CHAPTER 10
DECISION TEST

The dashboard is a role-specific control surface whose freshness, definitions and next action are visible, not a gallery of independently calculated charts.

A factory dashboard earns its screen space only when a named role uses it within a defined review window to take an operational action and verify the result. [S08] [S30] Visualisation software cannot repair ambiguous denominators, unsynchronised clocks or unowned exceptions. In a connected textile plant operating under the infrastructure constraints documented for Cameroon, where 93.3 percent of surveyed formal firms reported electrical outages in the 2024 Enterprise Survey and typical monthly outage frequency reached 10.4 events, a screen that silently displays stale numbers or recalculates metrics with private logic becomes an operational hazard. [S32] [S40] Green Future Tech Limited (GFT), Hong Kong SAR, treats the dashboard not as a presentation project, but as a governed consumer of the factory data spine. Every tile must resolve to a versioned key performance indicator (KPI) definition under ISO 22400-2, expose its source time and data-quality state beside the figure, separate line-huddle exception handling from plant diagnostic drill-down under IEC 62264-1, bind every alert to an accountable owner and verification check under IEC 62682 alarm-management principles, and fail gracefully into manual procedures without touching machine control or functional safety. [S02] [S08] [S23] [S1001]

The sector context explains why disciplined operational interfaces matter. Cameroon's inherited cotton-to-textile baseline records roughly 95 percent of cotton fibre exported raw and 5 percent processed locally, against a national target of 50 percent local processing by 2030. [R06] [R07] Domestic industrial capacity has suffered from underutilisation: a July 2023 report recorded utilisation at CICAM's Garoua plant at 33.7 percent, CICAM 1 at 48.9 percent and CICAM 2 at 10.83 percent, costing roughly CFAF 2 billion in 2020. [R12] Moving to competitive manufacturing requires shift-level control over availability, quality, bundle flow and utility use. GFT's Cameroon textile programme (feasibility, performance-management design and digital strategy, prepared 31 January 2024, refreshed 2026) organises management around six KPI families on daily, weekly and monthly cadences. This chapter specifies the interface architecture turning those definitions into accountable decisions. It is an engineering specification, not investment, legal or tax advice.

10.1

Every tile resolves to a governed KPI definition

The most common failure in industrial reporting occurs when a dashboard tool becomes an informal calculation engine. An integrator connects a charting package to raw machine tags or spreadsheets, writes a local query dividing good pieces by shift hours, and labels the widget "Line Efficiency". Weeks later, the weekly workbook, the line supervisor's screen and the monthly pack report three different values for the same line and shift due to inconsistent downtime exclusions, rework rules or power outage handling. [S08] [S32] Once meetings spend their opening minutes debating whose number is right, the dashboard has failed.

ISO 22400-2 prevents this by defining manufacturing KPIs through their formula, source elements, time behaviour, unit and operational characteristics. [S08] In GFT's architecture, a dashboard application never invents or embeds its own KPI arithmetic. Every tile binds to a unique identifier and version in the factory KPI dictionary and queries the historian's certified KPI service. [S08] [S24]

10.2 A TILE CONTRACT BINDS VISUAL ELEMENTS TO FORMULA METADATA

A numeral alone is incomplete evidence. To make a defensible decision, the reader must know the denominator, aggregation window, source coverage and quality state. [S08] [S40] GFT requires every KPI tile to expose six attributes on the card face or via an inspection drawer:

1. **Definition identifier and version**: canonical code and semantic version in the KPI dictionary, linked to the ISO 22400-2 formula so changed exclusions increment the version under ISO/IEC 27001 change control. [S08] [S24]
2. **Aggregation window and time behaviour**: start and end source timestamps plus the ISO 22400-2 time behaviour classification (batch, shift-cumulative, rolling hourly or instantaneous rate). [S08]
3. **Numerator and denominator elements**: explicit quantities in both numerator and denominator (such as 412 illustrative good garments divided by 440 illustrative inspected garments) so ratios cannot hide volume drops. [S08]
4. **Source coverage ratio**: valid critical events or meter intervals received divided by expected events or intervals under the logging reference, using Sparkplug session state and wireless integrity checks. [S31] [S40]
5. **Freshness age and quality flag**: elapsed time since the latest contributing source timestamp, paired with the consolidated quality flag (`GOOD`, `STALE`, `UNSYNCHRONISED`, `ESTIMATED`, `CORRECTED` or `MISSING`). [S31] [S40]
6. **Decision owner and action link**: the role accountable for acting when the threshold is breached, linked to the open exception record or standard work instruction. [S08] [S1001]

This governed dictionary is summarised in **Exhibit 44**, p. 160, extending ISO 22400-2 with provenance, coverage and action fields. [S08] [S24] [S40] [S1001]

EXHIBIT 44 / SCHEMA SPECIFICATION FIELDS

The KPI dictionary is the dashboard's source of truth

Dictionary attribute	Standard basis	Operational content	Dashboard display rule
KPI ID and version	ISO 22400-2; ISO/IEC 27001	Unique code, semantic version, approval date and change log	Shown in tile footer and export header; version changes flagged on trend axes
Formula and elements	ISO 22400-2	Mathematical expression, numerator element, denominator element and exclusions	Drill-down reveals exact numerator and denominator values alongside the ratio
Unit and time behaviour	ISO 22400-2	Physical or monetary unit, evaluation window and update cadence	Explicit unit label on every axis and card; window start and end timestamps visible
Source coverage rule	Sparkplug 3.0; ISO 22400-2	Expected reporting nodes, minimum coverage threshold and missing-interval rule	Coverage percentage badge beside value; below-threshold tiles enter degraded styling
Quality state propagation	Sparkplug 3.0; IEC 62541-9	Precedence rule when inputs are stale, unsynchronised, estimated or corrected	Visual state badge prevents stale or estimated values from appearing as live good data
Owner and action link	IEC 62682; ISO 22400-2	Accountable role, threshold setpoints, response window and verification method	Click-through opens exception ticket or standard work instruction for the named role

DERIVED [S08] [S24] [S40] [S1001] Derived schema combining ISO 22400-2 KPI definition elements (S08) with ISO/IEC 27001 change control (S24), Sparkplug session-state coverage (S40) and IEC 62682 operator-response linkage (S1001).

Exhibit 44, p. 160 establishes that a dashboard tile is an interface to a contract, not a freehand chart. [S08] What the exhibit cannot establish by itself is the specific numeric threshold or exclusion policy for a given factory; the process owner and data owner must approve those parameters during Phase 0 set-up before any screen enters production use. [S30]

10.3 SHARED SERVICES PREVENT DIVERGENCE BETWEEN SPREADSHEETS AND SCREENS

During early digitalisation, automated edge capture coexists with manual logs for bundle tickets, lab tests or attendance. [S25] [S30] If the line screen reads from an MQTT broker while the weekly pack uses a standalone workbook, formulas drift apart: one analyst excludes fifteen-minute power changeovers while another includes them; one clerk counts reworked seams as defects while another records only scrap. [S08] [S32]

To eliminate drift, GFT enforces a single-calculation boundary under ISO/IEC 27001 configuration control. [S24] Whether an input comes from an IO-Link sensor, Modbus meter, RFID reader or manual form, the historian computes the KPI once under the active dictionary version and stores the versioned output. [S08] [S24] Both live screens

and weekly workbooks query that stored result. Authorised retrospective corrections create a new calculation version with approver, timestamp and reason, preserving the audit trail of what appeared during the shift. [S08] [S24]

> **GFT NOTE**
>
> Never configure a business-intelligence tool to calculate production KPIs inside browser templates or private workbooks. Require every view and report to read pre-computed, version-stamped KPI records from the governed historian service, and reject any tile that cannot expose its numerator, denominator, definition version and source coverage ratio on demand. [S08] [S24]

10.4

The line view favours immediate, controllable exceptions

A sewing supervisor, weaving overlooker or dyehouse shift leader does not manage a shift through thirty-day averages or financial ratios. During an active shift, the supervisor controls immediate variables: workstation balance, bundle arrival at the bottleneck on takt, mechanic response to stopped machines, and stitch, shade or weave defects requiring needle, tension or lighting adjustment. [S25] [S26] [S28] Cluttering a floor screen with corporate allocations or twenty gauges violates human-machine interface principles and delays action. [S02] [S1001]

IEC 62264-1 separates Level 3 manufacturing operations management from Level 4 enterprise planning, defining distinct scopes and horizons. [S02] Within Level 3, GFT's Cameroon programme uses a three-tier cadence: shift and daily line huddles for immediate execution, weekly plant reviews for cross-functional bottlenecks, and monthly sponsor reviews for capacity and unit economics. The dashboard architecture provides role-specific views over shared KPI definitions. [S02] [S08]

10.5 THE SHIFT HUDDLE SCREEN ANSWERS THREE OPERATIONAL QUESTIONS IN UNDER A MINUTE

On a cutting, sewing or finishing line, an Andon board or tablet has a strict cognitive budget. In a sixty-second scan at the hourly check or shift huddle, it must answer three questions without scrolling:

- **Are we ahead or behind the current hourly and shift target for good units?** The production tile compares actual good units against cumulative hourly and shift targets using the ISO 22400-2 good-quantity element, separating first-pass good output from rework. [S08]
- **Which station or machine is holding back flow or generating defects?** The exception panel lists active stops, RFID bundle starvation, or defect clusters from inline inspection cells, ordered by duration and throughput impact. [S25] [S26] [S28]
- **Who owns the active abnormality, and when is the response due?** Each exception shows the assigned role (mechanic, quality rover or bundle handler), elapsed duration, and acknowledgement state. [S1001]

When stations operate within limits and coverage is complete, the line view stays visually calm on neutral paper tones so an ochre warning or alert commands immediate attention.

10.6 INFORMATION DENSITY SCALES WITH REVIEW TIME, NOT SCREEN RESOLUTION

Designing a single "master dashboard" on a desktop monitor and projecting it five metres above an aisle fails both operators and executives. [S02] [S1001] **Exhibit 45**, p. 162 compares three illustrative interface panels (Line huddle view, Plant weekly view and Sponsor monthly view) across four functional emphasis scores on a 0 to 10 scale. [S01] [S02] [S08]

EXHIBIT 45 / ILLUSTRATIVE DESIGN EMPHASIS SCORE (0 TO 10)

Role-specific views share definitions but not information density



ILLUSTRATIVE [S01] [S02] [S08] Illustrative interface-design allocation across four content classes for three management roles, structured using ISA-95 hierarchy levels (S01), IEC 62264 Level 3 and Level 4 boundaries (S02) and ISO 22400 KPI aggregation rules (S08). Values are design weights, not measured factory results.

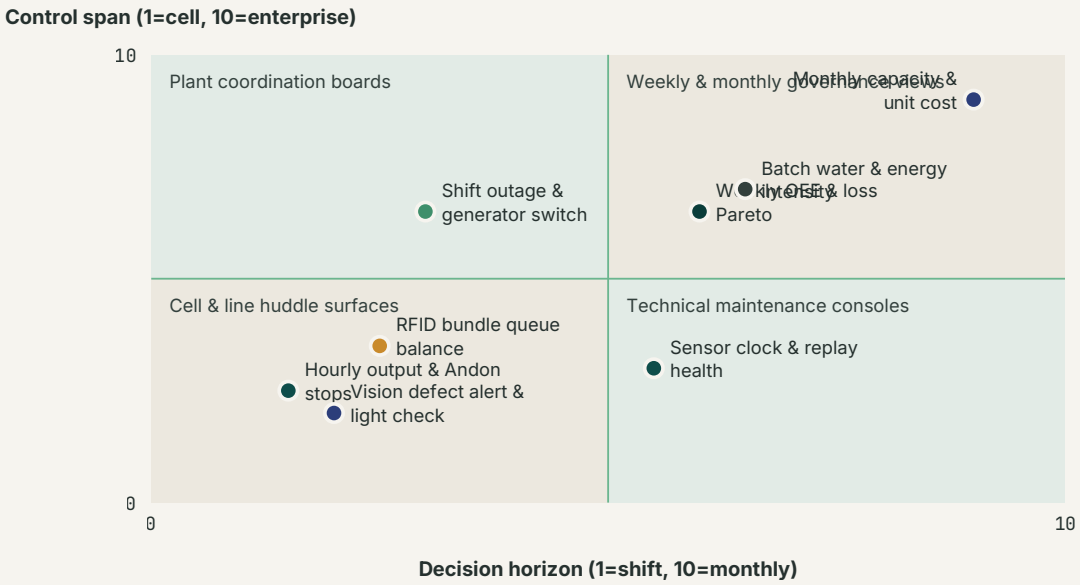
As **Exhibit 45**, p. 162 shows, the line huddle view allocates nearly all visual weight (an illustrative 9 out of 10) to live exceptions and shift variance, with zero space on capital or policy. [S02] [S08] The plant weekly view emphasises cross-family links and multi-week trends, while the sponsor monthly view focuses on trajectory, unit cost and capacity utilisation. [R12] [S01] [S02] [S32] As an illustrative allocation model, **Exhibit 45**, p. 162 prescribes design emphasis rather than reporting empirical operator-response measurements.

10.7 DECISION HORIZON AND CONTROL SPAN DETERMINE WIDGET PLACEMENT

Every candidate tile must sit in the view whose cadence matches the process time constant and viewer authority. [S02] [S08] Showing five-second kilowatt-hour pulses to a sewing supervisor creates distraction without agency; waiting for a monthly pack to report three weeks of camera illumination drift guarantees fabric waste. [R29] [S26] [S29] **Exhibit 46**, p. 163 maps eight textile modules across derived decision horizon (1 to 10) and ISA-95 / IEC 62264 control span (1 to 10). [S01] [S02] [S08] [S30] [S1001]

EXHIBIT 46 / DERIVED HORIZON AND SPAN INDEX (1 TO 10)

Dashboard content follows decision horizon and control



DERIVED [S01] [S02] [S08] [S30] [S1001] Derived placement of textile decision modules by decision horizon (x-axis: 1=intra-shift minutes, 10=monthly review) and ISA-95 / IEC 62264 control span (y-axis: 1=workstation/cell, 10=enterprise/sponsor), informed by S01, S02, S08, S30 and S1001.

In **Exhibit 46**, p. 163, lower-left modules (`Cell & line huddle surfaces`) serve floor supervisors acting within minutes. [S08] [S1001] Upper-left modules (`Plant coordination boards`) coordinate site-wide shift responses such as generator switching during outages. [S32] Upper-right modules (`Weekly & monthly governance views`) guide plant and sponsor reviews of OEE losses, resource intensity and capacity utilisation. [R12] [S02] [S08] Lower-right modules (`Technical maintenance consoles`) track sensor clocks, buffer replay and camera calibration without cluttering floor huddles. [S26] [S31] [S40] Local roles must be mapped during site commissioning.

10.8

The plant view connects families without hiding their sources

Plant decisions cross indicator families: a delivery drop may stem from machine stops, cutting defects, absenteeism, boiler downtime or power dips. [S26] [S29] [S32] GFT organises the plant view across six KPI families: **Production** (throughput, availability, OEE), **Quality** (first-pass yield, defect taxonomy, rework), **Delivery and WIP** (lead time, bundle queue age), **Cost** (labour, scrap and maintenance per good unit), **People** (attendance, skills, overtime against Cameroon's 40-hour standard and 48-hour maximum week), and **Sustainability** (electricity kilowatt-hours per good unit, water litres per kilogram of fabric, waste share). [R16] [S08] [S29] [S37] [S38] Connecting these families requires preserving measurement boundaries and avoiding false causal claims.

10.9 DRILL-DOWN FOLLOWS THE ISA-95 ASSET AND ORDER HIERARCHY

When weekly first-pass yield drops by 4.2 illustrative percentage points, an aggregate chart cannot guide Monday action. The plant view supports deterministic drill-down along two paths defined by ISA-95 and EPCIS: [S01] [S15]

1. **Physical asset drill-down**: site to area (spinning, weaving, dyeing, cutting, sewing, finishing), line, work centre, machine and measurement point. [S01]
2. **Execution and material drill-down**: weekly aggregate to shift, order, fabric batch or dye lot, style, and EPCIS bundle event sequence. [S15] [S25] [S28]

Every step preserves the parent denominator and event list. [S08] If Line B's OEE trails Line A's, drill-down shows whether Line B had higher mechanical downtime, smaller batch runs or a three-hour feeder outage, preventing managers from blaming a supervisor for external constraints. [S01] [S02] [S08] [S32]

10.10 CROSS-FAMILY VIEWS PROMPT INVESTIGATION RATHER THAN ASSUME CAUSALITY

Visual correlation on a multi-family screen is not physical causation. If dyehouse water use per kilogram rises in the same week as shade rework, excess water did not cause shade defects; off-shade lots required re-dyeing and extra rinsing against a static good-fabric denominator. [S08] [S27] [S38] Likewise, where edge energy monitors compute local parameters, specific energy use rises during high-outage weeks because fixed baseload utilities draw power during restarts while output is zero. [S08] [S29] [S32] GFT applies three rules to multi-family views:

- **Denominators are plotted alongside ratios**: companion panels show good-unit denominators so a volume drop is never misread as a consumption spike. [S08] [S29]
- **Contextual state bands accompany process trends**: charts shade intervals of backup generation, grid outage, changeover or degraded sensor coverage. [S29] [S32] [S40]
- **Cross-family anomalies generate investigative prompts, not automated blame**: under ISO 50001 and Six-Sigma pilot discipline, simultaneous shifts open a root-cause ticket requiring engineering verification. [S30] [S37]

10.11

Freshness and coverage sit beside every important value

In an office application, a number on a screen is assumed to reflect the current state of the database. In an industrial plant with mixed-generation equipment, wireless shop-floor readers and intermittent wide-area links, that assumption is dangerous. [S31] [S33] [S35] If an edge gateway on a sewing line loses its local network connection at 09:15 during a power dip, and the central plant dashboard continues displaying the 09:14 cumulative output and a `0.0%` defect rate in bright green font at 11:30, the screen is deceiving the supervisor. [S32] [S40] A quiet screen can mean a smoothly running line, or a dead sensor.

10.12 SPARKPLUG SESSION STATE AND SOURCE TIMESTAMPS GOVERN VISUAL RENDERING

The Sparkplug 3.0 specification addresses this hazard for real-time SCADA and HMI solutions through birth (`NBIRTH`, `DBIRTH`) and death (`NDEATH`, `DDEATH`) certificates. [S40] When an edge node disconnects, all metrics previously published under that node and its devices in that session are implicitly `STALE`. [S40] When the primary host returns online and the edge node flushes buffered historical data, each record retains its original source timestamp rather than broker arrival time. [S40]

GFT extends this protocol discipline to the pixels on the dashboard. No value may be rendered without evaluating its **freshness age** (`t_review - t_source`) and its **source coverage ratio** (`N_valid / N_expected`). [S08] [S31] [S40] Every tile renders one of four visual states driven by machine-readable quality flags:

1. **Live and complete (`GOOD`)**: freshness age is within the approved limit and source coverage meets threshold, rendered in high-contrast typography with an explicit age and coverage badge. [S40]
2. **Delayed or buffered (`STALE`)**: an `NDEATH` or `DDEATH` certificate arrived, or freshness age exceeds the limit; the tile drops saturated status colours, applies a hatched border and `STALE` badge, shows the last known source timestamp, and suppresses "all clear" styling. [S40]
3. **Partial coverage or clock uncertainty (`DEGRADED` / `UNSYNCHRONISED`)**: some devices in an aggregated metric are offline or wireless node clocks drifted outside the correlation window; the tile shows the coverage fraction (`Coverage 4 of 6 machines | 66.7%`) and blocks cross-line ranking until replay completes. [S08] [S31] [S40]
4. **Manual or corrected override (`ESTIMATED` / `CORRECTED`)**: the figure includes a fallback manual log entry or a post-shift supervisory correction under change control, marked with a link to the approver and reconciliation note. [S22] [S24]

METHOD

Method: Calculating source coverage, freshness age and denominator reconciliation

Before rendering a KPI tile for evaluation window $[T_0, T_1]$ at review time T_{review}, the dashboard engine evaluates three equations: [S08] [S40]

1. **Source coverage (C)**: $C = \frac{N_{\text{valid}}(T_0, T_1)}{N_{\text{expected}}(T_0, T_1)} \times 100\%$, where N_{valid} is the count of expected reporting intervals or critical events received with `GOOD` or verified-replayed flags inside $[T_0, T_1]$, and N_{expected} is the scheduled count in the KPI dictionary. [S08] [S40]
2. **Freshness age (A)**: $A = T_{\text{review}} - \max(t_{\text{source}})$. If $A > A_{\max}$ on the use-case card or the Sparkplug session is offline, the tile state becomes `STALE`. [S40]
3. **Inspection coverage distinction**: For defect rate $D = \frac{Q_{\text{defect}}}{Q_{\text{inspected}}}$, the tile simultaneously evaluates inspection ratio $I = \frac{Q_{\text{inspected}}}{Q_{\text{produced}}}$. [S08] [S25] [S26] If $Q_{\text{inspected}} = 0$ while $Q_{\text{produced}} > 0$, D is undefined (`MISSING` inspection data, $I = 0\%$), never `0.0%` defects. [S08] [S26]

10.13 A ZERO-DEFECT READING IS NEVER CONFUSED WITH ZERO INSPECTION DATA

The third equation prevents a severe failure mode in quality dashboards. Textile inspection literature notes that manual inspection is 60 to 70 percent accurate (missing 20 to 30 percent of defects), whereas computer vision can exceed 90 percent accuracy across more than 40 defect types at speeds up to 60 metres a minute when properly engineered. [R29] However, lighting variations and background noise in industrial settings can significantly degrade model detection accuracy. [S26] If an edge camera cell is placed in bypass mode due to a failed lamp or full local storage during a disconnect, the line produces garments while zero digital inspection events reach the historian. [S22] [S26] [S40] Dividing zero digital defects by total line production, or mapping divide-by-zero to `0.0%`, would falsely report perfect quality while inspection is blind. [S08] [S26]

WORKED EXAMPLE

Illustrative reconciliation of a one-hour sewing-line huddle display during partial telemetry loss

Consider an illustrative sewing line (`Line-02`) scheduled to produce 120 cotton shirts per hour across 4 parallel final-assembly stations (`Stn-A` to `Stn-D`), each reporting 12 five-minute Sparkplug intervals per hour ($N_{\text{expected}} = 48$) and feeding 1 end-of-line camera-assisted inspection table. [S08] [S26] [S40] All figures are illustrative to demonstrate the method, not a measured GFT factory result.

During the illustrative 10:00 to 11:00 shift window, a local switch power supply trips at 10:30 on the conduit serving `Stn-D` and the inspection table gateway: [S32] [S40]

- Stations `Stn-A`, `Stn-B` and `Stn-C` report all 12 illustrative intervals each (36 valid intervals, 90 completed shirts). [S08] [S40]
- Station `Stn-D` reports 6 valid intervals before 10:30 (15 shirts) and buffers 6 intervals locally while offline (15 additional shirts on its local counter, undelivered at 11:02). [S29] [S40]
- The inspection table logs 52 inspected shirts with 2 stitch defects before 10:30, and delivers 0 digital records between 10:30 and 11:00 while offline (operators record 50 inspected shirts and 3 defects on a fallback paper tally). [S22] [S26]

At the illustrative 11:02 huddle ($T_{\text{review}} = \text{11:02}$), an **ungoverned dashboard** querying arrival-time rows without coverage checks would display `105 shirts` (omitting `Stn-D`'s 15 buffered shirts) and a $2 / 105 = 1.90\%$ illustrative defect rate. [S08] [S40]

Under **GFT's governed dashboard rules** (**Exhibit 44**, `p. 160`), the 11:02 screen displays the true state of the evidence: [S08] [S40] 1. **Production tile**: shows `105 confirmed shirts (120 target)` with badge `PARTIAL COVERAGE: 42 / 48 intervals (87.5%) | Stn-D STALE (age 32 min)`. The supervisor sees `Stn-D` is offline and checks its local HMI counter (15 shirts) to confirm physical output met the 120-shirt target. [S08] [S40] 2. **Quality tile**: divides the 2 digital defects only by the 52 digitally inspected shirts ($2 / 52 = 3.85\%$ illustrative defect rate for 10:00 to 10:30), flags 10:30 to 11:00 inspection coverage as `MISSING (49.5% coverage; age 32 min)`, and prompts entry of the fallback paper tally (50 inspected, 3 defects). [S08] [S22] [S26] 3. **Post-replay reconciliation**: when switch power restores at 11:10 and buffered Sparkplug messages flush with source timestamps, the historian updates 10:00 to 11:00 output to $120 / 120$ shirts (`Coverage 100%`) and reconciles total defects to $(2 + 3) / (52 + 50) = 5 / 102 = 4.90\%$ illustrative defect rate under versioned change control. [S08] [S24] [S40]

10.14

Alerts end in owned action and verification

An industrial alarm or dashboard alert is a request for human intervention in a physical process. IEC 62682:2022 defines the primary function of an alarm system as notifying operators of abnormal process conditions or equipment malfunctions and supporting the response. [S1001] Companion specifications in IEC 62541-9:2026 (OPC UA Part 9: Alarms and Conditions), designed in accordance with IEC 62682 and ISA-18.2, standardise how alarm states, severities, limit deadbands, acknowledgements and shelving comments are represented. [S04] [S1002]

When a project treats alert creation as free, engineers attach pop-up rules to every sensor limit, causing alarm floods during power changeovers or wireless dips. [S31] [S32] [S1001] Operators respond to alarm fatigue by muting notifications, so genuine excursions go unaddressed. [S1001] Textile SME guidance recommends bounded pilots within a Six-Sigma discipline to validate that digital tools improve control before scaling. [S30] GFT therefore governs alerts as closed-loop operational transactions.

10.15 EVERY CONFIGURED ALERT REQUIRES A RATIONALISATION RECORD BEFORE ACTIVATION

Applying IEC 62682 rationalisation to Level 3 operations under IEC 62264-1, GFT prohibits activating any dashboard alert unless its use-case card documents five elements: [S02] [S1001]

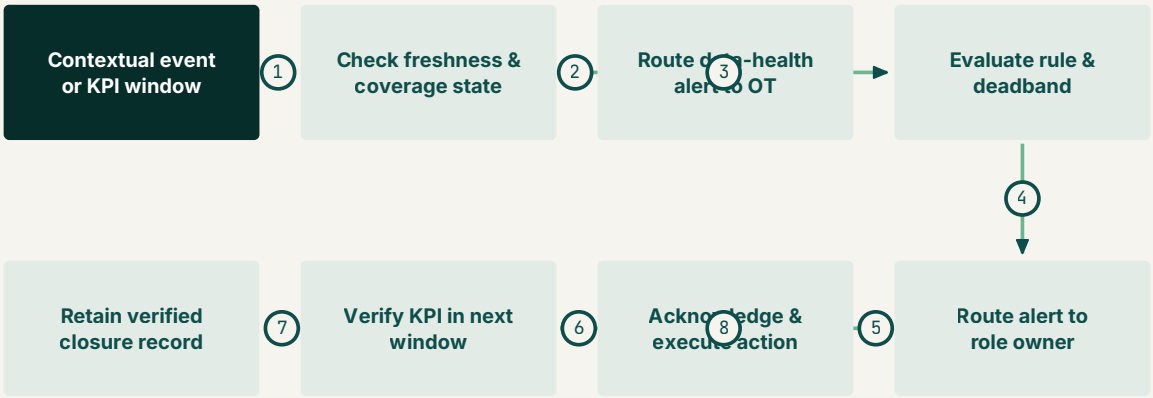
1. **Abnormal condition and deadband**: the exact KPI formula or event rule, threshold setpoint, hysteresis deadband and minimum persistence duration, modelled using IEC 62541-9 semantics to prevent chattering alerts. [S08] [S1002]
2. **Consequence of inaction**: the operational, quality or resource loss incurred if uncorrected within the response window. [S1001]
3. **Single accountable role**: the one shift role (`Line-02 Supervisor` or `OT Network Technician`) responsible for acknowledging and acting; an alert sent to everyone is owned by nobody. [S02] [S1001]
4. **Prescribed corrective action**: the concrete physical check or adjustment required (checking needle tension on Station 4, verifying illumination on Camera Cell 1, or shedding nonessential compressor load). [S26] [S29] [S1001]
5. **Verification rule and closure evidence**: how the system and supervisor confirm in the next evaluation window that the action restored performance, rather than closing on acknowledgement alone. [S08] [S30] [S1001]

10.16 ALERT LIFECYCLE SEPARATES ACKNOWLEDGEMENT, ACTION AND VERIFIED CLOSURE

Under IEC 62541-9 and IEC 62682, alarm condition states (`Active`/`Inactive`), acknowledgement states (`Acknowledged`/`Unacknowledged`) and shelving states (with comment and bounded duration) are distinct state machines. [S1001] [S1002] GFT's workflow adds an operational verification gate before closure and a parallel escalation path when telemetry is stale, as shown in **Exhibit 47**, `p. 169`. [S08] [S40] [S1001]

EXHIBIT 47 / OPERATIONAL STATE SEQUENCE

An alert closes only after action and verification



1. ingest
2. if STALE/MISSING
3. if GOOD
4. limit breached
5. owner intervention
6. next window
7. within target
8. escalate if unfixed

ILLUSTRATIVE [S08] [S30] [S40] [S1001] [S1002] Illustrative closed-loop alert workflow informed by ISO 22400 KPI evaluation (S08), Six-Sigma textile pilot control (S30), Sparkplug stale-state detection (S40), IEC 62682 alarm lifecycle (S1001) and IEC 62541-9 alarm-condition state modelling (S1002).

As **Exhibit 47**, p. 169 illustrates, the workflow branches first on data quality. [S40] If an edge node disconnects or a clock drifts out of synchronisation, the system routes a data-health alert to the technical owner while marking the production tile STALE on the supervisor's board. [S31] [S40] When data quality is GOOD and a threshold is breached outside its deadband, the alert routes to the operational role owner, records the physical intervention, and waits for the next evaluation window to verify recovery before closing. [S08] [S1001] [S1002] Because **Exhibit 47**, p. 169 is an illustrative workflow, it specifies required functional logic rather than empirical closure statistics.

10.17 STALE-DATA ESCALATION PREVENTS SILENT CONTROL FAILURE

In **Exhibit 47**, p. 169, silence is never treated as compliance. [S40] [S1001] If an inline inspection camera, dyehouse water meter or feeder energy meter stops publishing during an active order, suppressing process alarms without notice leaves the line blind. [S26] [S27] [S29] Conversely, sounding a quality klaxon on the sewing floor because an MQTT broker restarted confuses operators who cannot repair a switch. [S40] [S1001] GFT separates **process exceptions** (routed to production, quality or utility supervisors) from **evidence-chain exceptions** (routed to OT support while badging

affected tiles). [S02] [S40] [S1001] When an evidence-chain outage exceeds the allowed freshness limit, the dashboard instructs the supervisor to activate the manual fallback log until the digital conduit is restored and reconciled. [S22] [S40]

10.18

Degraded modes remain legible and operable

Where 93.3 percent of surveyed Cameroonian firms experience electrical outages and wide-area links face periodic disruption, degraded operation is a recurring operating state rather than a rare disaster. [S32] [S33] A dashboard architecture that shows a spinning browser wheel, a blank screen or an unlabelled cache during a server or network outage is unfit for factory service. [S17] [S22]

NIST SP 800-82 Rev. 3 requires operational technology architectures to address performance, reliability and safety alongside security. [S17] IEC 62443-3-2 segments industrial systems into security zones and conduits so business IT or cloud failures cannot propagate into production or safety zones. [S19] GFT combines these standards with NIST SP 800-34 contingency planning principles (testing recovery, training personnel and exercising plans) to specify how decision surfaces behave across four degraded modes. [S22]

10.19 FOUR FAILURE MODES HAVE PREDETERMINED DISPLAY AND OPERATING FALLBACKS

Every dashboard deployment must prove during commissioning how it behaves under four degradation tiers:

1. **Loss of wide-area or cloud connectivity (local network intact)**: because GFT places the operational historian, KPI service and local web/HMI server inside the factory's Level 3 operations zone rather than depending on an external cloud loop, shift line boards and plant views continue with zero loss of freshness during an internet outage. [S02] [S14] [S16] [S19] Only off-site sponsor replication and external EPCIS sharing pause until the wide-area conduit recovers. [S15]
2. **Loss of central factory historian or dashboard server (edge nodes intact)**: if the central Level 3 server restarts or undergoes maintenance, edge nodes continue local acquisition, inference and database persistence while machine HMIs and line-level edge displays remain readable. [S16] [S29] [S40] Supervisors record hourly huddle figures on pre-printed manual boards matching ISO 22400-2 elements; once the server restores and Sparkplug historical flush completes, replayed data is reconciled against the manual log. [S08] [S22] [S40]
3. **Loss of a specific sensor or edge gateway (central dashboard intact)**: the central dashboard remains online, isolates the missing source via Sparkplug `NDEATH`/`DDEATH` state, badges affected tiles `STALE` or `PARTIAL COVERAGE`, and prompts targeted manual logging only for the affected station. [S22] [S40]
4. **Security isolation or suspected compromise (read-only or air-gapped operation)**: if an incident requires isolating the industrial DMZ or operations zone under IEC 62443-3-2, remote sessions terminate immediately, multi-factor authentication gates recovery access under CISA CPG 2.0, and floor displays drop to local read-only mode or manual whiteboards while isolated backups are verified before restoration. [S19] [S20] [S21] [S22]

> **WARNING**
>
> **Constraint: Display failure and access controls must never impair machine control or protective safety**
>
> Under no circumstances may a dashboard server, KPI service, MQTT broker or remote HMI session sit inside a machine's closed-loop control path or functional-safety circuit. [S17] [S23] Functional safety provides automatic protection in response to errors through independent safety-rated hardware and logic. [S23] Commissioning and change-control acceptance tests must physically disconnect the dashboard conduit, stop the broker process and revoke remote HMI credentials while verifying that machines run safely and emergency stops and guard interlocks trip without delay. [S17] [S19] [S20] [S23]

10.20 CONTINGENCY EXERCISES TEST HUMAN DECISIONS, NOT JUST SERVER REBOOTS

NIST SP 800-34 states that testing validates recovery, training prepares personnel and exercising identifies planning gaps, while CISA advises regular backup testing from isolated media. [S21] [S22] GFT integrates the dashboard into three operational exercises:

- **Shadow-to-manual switchover drills**: during Phase 1 pilots and routine operation, the plant manager runs announced thirty-minute "dark screen" drills where supervisors conduct hourly huddles from local machine counters and manual tally cards, reconciling records once displays return. [S22] [S30]
- **Post-outage reconciliation checks**: after any power dip or link loss, the recovery view highlights the backfilled window from edge buffers or manual logs so the supervisor and data owner confirm no double-counted bundles or missing downtime corrupt weekly KPIs. [S08] [S24] [S29] [S40]
- **Alert closure and stale-tile audits**: during the weekly plant review under ISO/IEC 27001 governance, operations leaders audit unacknowledged alerts, overdue actions, chattering rules and tiles spending excessive shift time in `STALE` state, repairing capture points before adding new screens. [S24] [S30] [S1001]

By treating the dashboard as a governed decision instrument rather than decorative wallpaper, the factory ensures that every signal captured at the production edge ends in a verified action on the floor. [S08] [S30] While this chapter specified how operational

roles consume governed evidence, the next chapter examines how a small factory team can deploy, monitor and roll back the software modules that power this stack without endangering production continuity. [S16] [S17] [S24]

A screen that changes no action is only illuminated wallpaper; a governed dashboard turns verified events into the next shift's improvement.

CHAPTER SOURCE NOTES

S08 International Organization for Standardization. 2014-01.

S30 MDPI, Sustainability. 2023-08-19.

S32 World Bank Enterprise Surveys. 2025.

S40 Eclipse Foundation. 2022-11-16.

S02 International Electrotechnical Commission. 2013-05-22.

S23 International Electrotechnical Commission. 2022-04-28.

S1001 International Electrotechnical Commission. 2022-12-08.

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.

S24 International Organization for Standardization and International Electrotechnical Commission. 2022-10.

S31 MDPI, Sensors. 2020-04-28.

S25 Springer Nature, International Journal of Advanced Manufacturing Technology. 2016-09-05.

S26 MDPI, Electronics. 2025-09-18.

S28 Elsevier, Robotics and Computer-Integrated Manufacturing. 2013-04.

S01 International Society of Automation. 2025-04-10.

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

S29 MDPI, IoT. 2024-09-28.

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

S37 International Organization for Standardization. 2018-08.

S38 International Organization for Standardization. 2014-08.

S15 GS1 AISBL. 2022-06.

S27 MDPI, Applied Sciences. 2024-06-21.

S33 International Telecommunication Union. 2024-11-10.

S35 United Nations Industrial Development Organization. 2019.

S22 National Institute of Standards and Technology. 2010-11-11.

S04 OPC Foundation. 2025-10-22.

S1002 International Electrotechnical Commission. 2026-04-16.

S17 National Institute of Standards and Technology. 2023-09.

S19 International Electrotechnical Commission. 2020-06.

S14 World Wide Web Consortium. 2023-12-05.

S16 European Telecommunications Standards Institute. 2024-04.

S20 Cybersecurity and Infrastructure Security Agency. 2025-12.

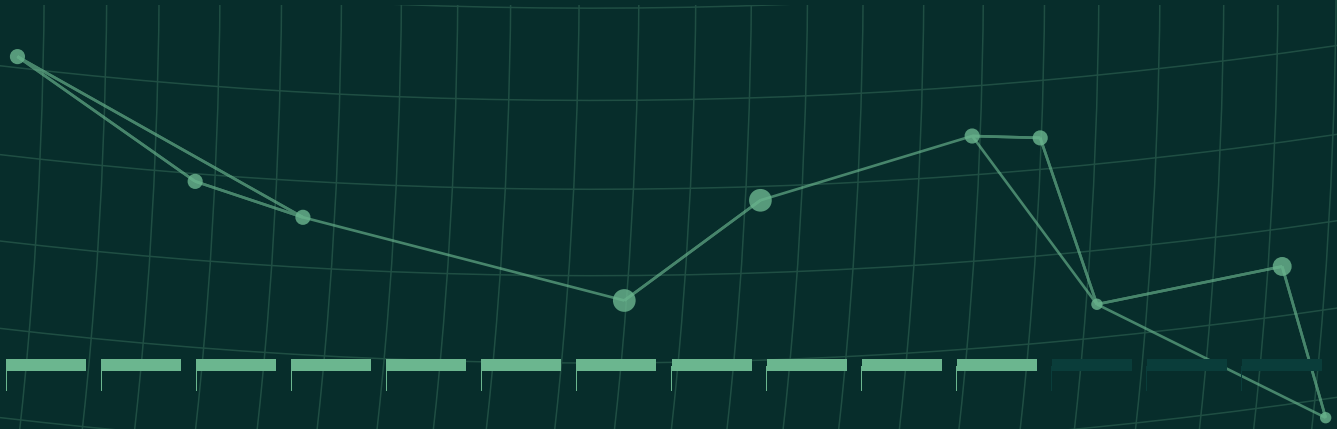
S21 CISA and joint international cybersecurity authorities. 2022-05-09.

11

PART III: ONE GOVERNED DATA SPINE / OPERATING THE STACK

Small factory teams need reversible application deployment

Applications should arrive as bounded, observable and reversible changes to a working production architecture, with ownership and rollback proven before scale.



CHAPTER 11
DECISION TEST

Applications should arrive as bounded, observable and reversible changes to a working production architecture, with ownership and rollback proven before scale.

Small factory teams should deploy applications only when they can identify the change, observe its effect, withdraw its authority and restore an acceptable operating state. GFT proposes treating reversibility as an acceptance condition, alongside the operational benefit. A release is therefore a controlled change to a decision and its evidence, rather than merely a successful installation. The practical question is whether the available team can recover the service and reconcile its records while production continues under an agreed fallback. This chapter sets out a proposed operating method, not a claim that GFT has deployed or proved a factory application.

11.1

A pilot must test an operational hypothesis

Research on textile small and medium-sized enterprises recommends limited pilots to validate selected technologies before wider implementation. Its implementation model provides a useful starting point, but does not establish a universal factory schedule, staffing requirement or financial return. GFT's proposed extension is to give each pilot a decision owner, a comparison method and an explicit withdrawal condition before connecting it to operational work. The installation itself is an intermediate event; acceptance concerns the decision that the application helps a supervisor make. [S30]

A useful hypothesis describes a process that can be observed. For example, an illustrative inspection pilot might ask whether presenting suspected defects with fabric position helps an inspector locate and classify them without increasing unresolved holds. The baseline records the existing inspection method, product family, operating conditions and treatment of borderline cases. The candidate records the same boundary. A faster screen response is relevant only if the inspector can use it and the physical disposition of the fabric remains traceable. This is a proposed test design, not a reported inspection result.

GFT recommends agreeing success and withdrawal criteria together. Acceptance should require adequate evidence coverage, correct event identity, usable instructions and a confirmed fallback, as well as a beneficial operating measure. A withdrawal criterion might be unexplained discrepancies between inspected material and recorded decisions. The criterion should name who can stop the pilot and where unresolved material will go. If the team discovers that the baseline itself is unreliable, the decision is to repair measurement and repeat the comparison, rather than treat missing records as evidence of improvement.

The pilot brief should also state which work remains outside the test. GFT recommends naming excluded product families, unsupported machine states and the decisions that still use the accepted method. A quiet shift should not silently extend that scope. At handover, the outgoing supervisor should identify any material whose classification remains provisional and the record used to resolve it. A candidate that works only when its developer is present has established a narrower capability than an application operated independently by the shift team.

The comparison should retain product mix, shift, machine state and intervention notes. Under this proposed method, a change in fabric family during the candidate period is a reason to analyse the periods separately. A maintenance intervention is

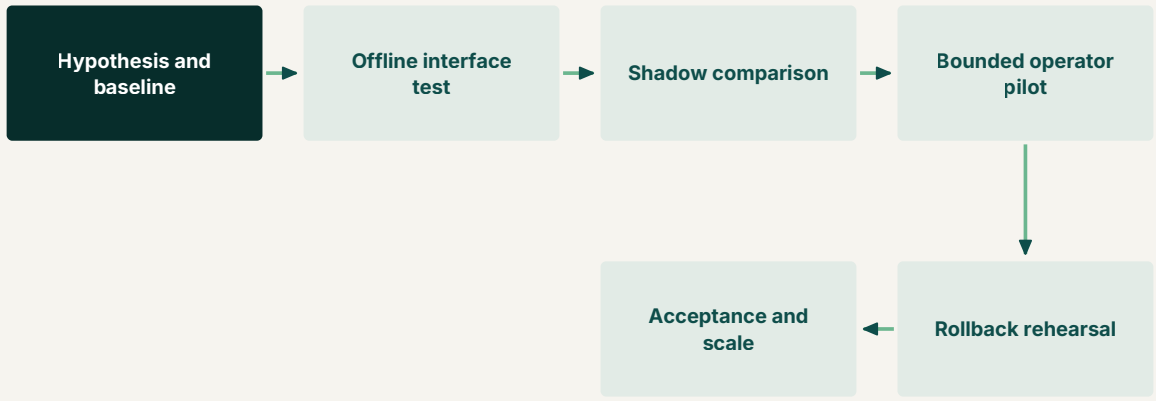
recorded alongside the software release rather than credited to it. Where comparable operating conditions cannot be obtained, the pilot can still establish interface compatibility and recovery behaviour. It should then report those results narrowly and leave the productivity hypothesis unresolved. Technical learning remains useful without a fabricated causal benefit.

For GFT's Cameroon textile programme, the proposed application strategy is to buy or license a proven capability before commissioning bespoke development, while retaining control over capture definitions and access to records. This is a design principle, not an assertion about purchases or supplier performance. The programme should configure the production context that differentiates the factory, then request custom code only where an observed requirement remains unmet. A product demonstration does not prove that local interfaces, intermittent connectivity or the factory's exception handling will work.

The illustrative gates in **Exhibit 48**, p. 175 combine the pilot sequence with the return path. Each forward step requires retained evidence. A failed gate returns the application to the agreed operating baseline, not automatically to another development attempt. The diagram is a proposed procedure informed by the pilot literature and high-level risk guidance; neither source prescribes this exact workflow. [S30] [S18]

EXHIBIT 48 / PROPOSED RELEASE SEQUENCE

Evidence gates precede wider application dependence



ILLUSTRATIVE [S30] [S18] Illustrative GFT procedure. Gates have no assumed duration or success rate. Each stage must retain evidence; failure invokes the agreed fallback and a reconciliation review before another attempt.

A gate should leave behind an artefact that another shift can inspect: the baseline definition, interface test record, discrepancy log, operator acknowledgement or recovery record. The release owner collects these in a compact release pack rather than creating a parallel reporting system. Shadow operation means the candidate observes and records without controlling production decisions. Moving into an operator pilot changes that authority, so the approval must identify the affected decisions and the means of removing the candidate from them.

11.2

Edge modules isolate application change from machine change

ETSI's edge reference architecture separates applications, the hosting environment and management responsibilities. Its software application model supports reasoning about replaceable services near the network edge. It does not require a textile factory to adopt the complete telecommunications architecture. GFT proposes using the separation as a practical design discipline: an adapter, inspection service or reporting calculation should have an explicit boundary that can be changed without silently changing machine behaviour. [S16]

The boundary begins with inputs and outputs. GFT recommends a module card naming the events consumed, expected schema, source identity, time interpretation and quality states accepted. It also declares what the module publishes and which consumer relies on it. Its operating section records resource limits, writable storage, permitted network destinations, service identity and required dependencies. Its failure section explains what persists, what becomes stale and who receives the warning. The card should allow an operator and maintainer to describe the service's purpose without reading its source code.

NIST's OT guidance addresses security in the context of performance, reliability and safety. GFT therefore proposes that an analytical application initially receive only the authority required to observe its approved inputs. A dependency on a reporting service must not become an undocumented dependency for machine operation. An application that can change a recipe or issue a machine command needs a separate engineering assessment of that control path; it cannot inherit approval from a successful read-only pilot. [S17]

GFT proposes testing interface exceptions before exposing the candidate to live work. The replay set should include valid events, missing fields, unknown assets, duplicate delivery, delayed arrival and an unexpected restart. Expected responses should be declared before the replay begins. Rejection should preserve enough context for investigation without allowing invalid observations into an accepted result. A tolerant parser is useful only when its tolerance cannot silently convert an unknown unit or absent identity into apparently valid production evidence.

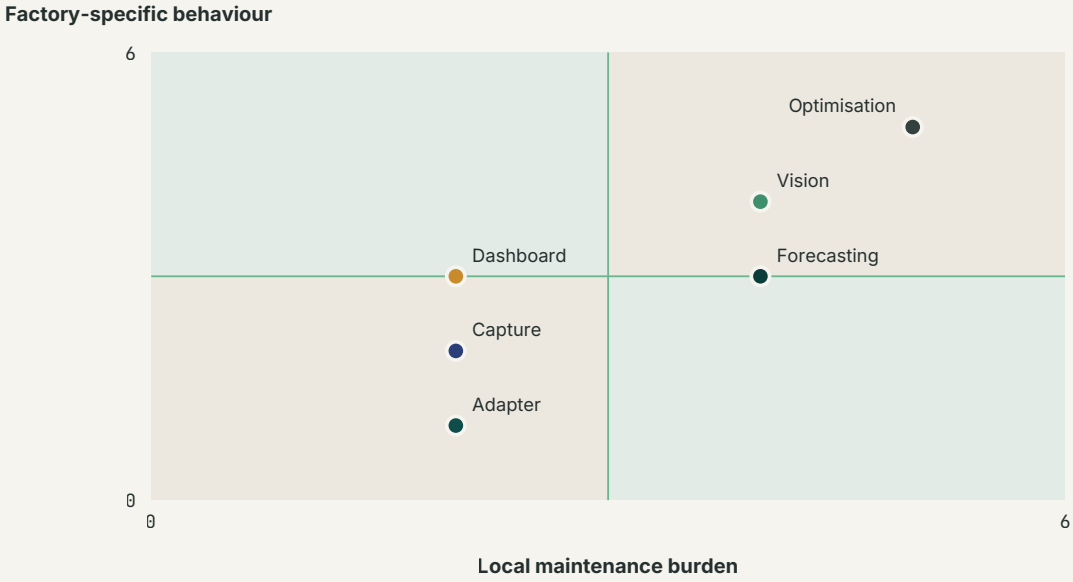
For the proposed factory stack, capture and durable buffering should remain available when an optional analytical module is stopped. A candidate must not consume the only copy of an event before the accepted path has retained it. GFT recommends testing resource contention by running the candidate against representative replay while the accepted capture process remains active. The acceptance record should describe memory pressure, disk growth and input lag, including the response when the host reaches its declared limits. Packaging a process does not itself prove isolation.

An illustrative sourcing matrix appears in **Exhibit 49**, p. 177. Its positions encode GFT's assumed local maintenance burden and need for factory-specific behaviour, rather than supplier ratings. Moving a function towards bespoke development should

require evidence that configuration cannot meet a material requirement and that the factory can maintain the resulting responsibility. The matrix organises the discussion; it does not calculate a procurement winner.

EXHIBIT 49 / ILLUSTRATIVE ORDINAL SCORE, 1 LOW TO 5 HIGH

Custom behaviour must justify its maintenance burden



ILLUSTRATIVE Illustrative GFT discussion model, not measured vendor capability. Coordinates are assumed: protocol adapter (2,1), durable capture (2,2), dashboard configuration (2,3), vision integration (4,4), forecasting integration (4,3), process optimisation (5,5). Scores do not imply equal economic intervals or an approved supplier.

The proposed sourcing review should include an exit demonstration. Ask the supplier or implementation team to export a representative record set, restore a configuration into an approved test environment and show how the retained records can be read after the candidate is removed. The test need not recreate every future migration. It must establish that the factory can retrieve its own evidence and identify the operating obligations that continue after support ends. An untested export button is an unresolved dependency, even when the initial installation is straightforward.

Under this illustrative interpretation, standard adaptation and capture are candidates for purchase or configuration first, while factory-specific integration may justify more engineering. GFT recommends requesting a record export, a documented interface, a supported previous release and an offline operating demonstration during evaluation. A supplier's contract should identify access needed for maintenance and the handling of service termination. These are proposed acceptance questions, not assertions that any named supplier provides those capabilities.

GFT NOTE

Keep the first release small enough for the available shift team to explain its dependencies and execute its fallback. If an application requires specialist intervention that cannot be obtained during the factory's operating hours, reduce its authority or postpone reliance until the support arrangement has been demonstrated. Buy-before-build should reduce local operating responsibility, not conceal it inside an opaque service.

11.3

Version identity joins data, configuration and model

The information-exchange part of the manufacturing digital-twin framework identifies technical requirements between reference-architecture entities. That scope supports explicit interface contracts, but does not prove that a particular factory has compatible implementations. GFT proposes recording the accepted contract beside each release so that a deployment cannot silently change the interpretation of an event. A message that still parses may nevertheless have acquired a different unit, classification rule or timestamp meaning. [S06]

Under the proposed method, the release manifest identifies the application artefact, configuration, input schema, output schema, model where applicable and relevant calibration reference. It also identifies the approved predecessor and the compatibility conditions for returning to it. An artefact digest helps distinguish exact packages; it does not establish that the package was authorised or that its origin is trustworthy. GFT recommends separately recording the approved source of the artefact, verification result and release approver, with credentials held outside the manifest.

The provenance standard for characteristic master data has a narrower scope than a complete application history. Its published scope excludes configuration management and change control. GFT therefore distinguishes provenance of a material or asset characteristic from analytical lineage and from release authorisation. Connecting these records is useful, but citing the master-data standard cannot substitute for a change process or establish that model outputs have been reproduced. [S41]

GFT's proposed analytical record links a result to its source-event identity, source quality, application identity and the configuration that affected interpretation. Where an inference uses an image, the record should identify the retained image or governed reference permitted by the retention policy. Where an aggregate uses a rule, the result should identify the rule and aggregation boundary. A corrected result should refer to the superseded result and reason for correction. It should not overwrite the history that explains an earlier operational decision.

The illustrative layers in **Exhibit 50**, p. 179 separate what was observed, how it was interpreted and who allowed the interpretation into use. No layer is sufficient alone. An approver can authorise an unsuitable package, and a correct calculation can be displayed outside its accepted operating scope. GFT recommends testing the links by selecting a retained decision and reconstructing it through the release pack, rather than accepting a manifest merely because every field is populated.

EXHIBIT 50 / PROPOSED EVIDENCE RESPONSIBILITIES

Release authority stays distinct from analytical lineage



ILLUSTRATIVE [S06] [S41] [S18] Illustrative GFT architecture, not a normative schema. ISO 23247-4 informs information exchange. ISO 8000-120 is limited to characteristic master-data provenance and excludes change control. CSF informs risk governance, not the exact release fields.

Data compatibility determines whether the return path is credible. GFT recommends introducing a new field without immediately removing the accepted field, then checking that both the candidate and predecessor can operate on the retained event set. A destructive migration should be postponed until its separate restoration and reconciliation procedure has passed. Returning an executable to its previous version cannot reverse an incompatible record transformation or recover observations that were never stored.

GFT recommends giving compatibility tests an explicit consumer boundary. A candidate dashboard may understand a changed classification while the accepted reporting service does not. Both must be exercised on the same retained events, including the transition period, before the team calls the release reversible. Compare the interpreted fields and decisions, not just record counts. Equal counts can conceal changed meaning. Where a compatibility adapter is introduced, give that adapter its own accepted artefact and return procedure rather than treating it as invisible migration work.

Configuration should pass the same practical discipline. Under the proposed method, a threshold edit, changed asset mapping or revised unit conversion is a release whenever it changes a decision. Such edits need an identifiable difference, an operator-visible consequence and a tested return path. Model retraining remains a separate lifecycle topic; this chapter requires only that the application can identify the approved model and withdraw its operational use without losing the underlying production evidence.

11.4

Monitoring watches acquisition and decision quality together

An industrial textile-inspection study's publisher-deposited abstract describes edge inference under controlled lighting, together with recorded defect information. This is useful evidence that an application boundary includes acquisition conditions and output records. Only the abstract was retrieved for this chapter; its existence does not establish performance at another factory or quantify sensitivity to lighting variation. GFT consequently proposes testing changed illumination, background, focus and material presentation explicitly before treating an inspection service as dependable. [S1101]

The proposed health review separates acquisition failure from interpretation failure. If the image is missing, blurred or poorly illuminated, the immediate response is to inspect capture conditions and mark affected output appropriately. If valid images reach the application but decisions disagree with reviewed examples, the review moves to classification behaviour. If accepted decisions fail to reach the operator, the problem lies further downstream. GFT recommends retaining these distinctions in incident records so that every apparent loss of accuracy does not trigger unnecessary model replacement.

A review of wireless sensing identifies synchronisation and spectrum sharing as data-integrity concerns. That finding supports checking the input path before attributing changed results to application behaviour. For the proposed factory release process, clock anomalies, delayed events, missed reads and interference observations should be recorded alongside the candidate's output. A simultaneous rise in discrepancy and input delay is a diagnostic lead, not proof of either model drift or a software defect. [S31]

GFT recommends monitoring operating consequence as well as technical state. An inspector's override should retain a reason and a route for review. An alert acknowledged without a recorded disposition should remain unresolved. A reporting calculation that is fresh but based on incomplete source coverage should retain that qualification beside its output. These measures belong to application acceptance because they reveal whether the release preserves the decision contract established by the preceding dashboard chapter; they do not replace the governed KPI definitions themselves.

The illustrative monitoring register in **Exhibit 51**, p. 181 assigns an immediate response to each signal. Its basis is the preceding evidence on OT requirements, sensing integrity and controlled acquisition, extended through GFT's proposed release method. It contains no universal alarm threshold. Each factory must define the response boundary through its use-case contract and test whether the responsible person can distinguish a service warning from an instruction to stop production. [S17] [S31] [S1101]

EXHIBIT 51 / ILLUSTRATIVE MONITORING REGISTER

Health signals require an owner and a recovery action

Signal	Accountable role	First response
Module and resource state	Service maintainer	Protect capture capacity
Freshness and coverage	Data owner	Qualify affected results
Acquisition quality	Inspection lead	Check optics and lighting
Schema and model identity	Release owner	Verify accepted manifest
Override and disposition	Production supervisor	Review unresolved material
Replay backlog	Data owner	Reconcile without duplicates
Return-path readiness	Release owner	Rehearse accepted recovery

ILLUSTRATIVE [S17] [S31] [S1101] [S1102] Illustrative GFT operating register informed by the cited scopes. Rows and responses are proposed controls, not a prescribed standard or observed factory results. Thresholds must be set and tested locally.

The register should be tested using deliberately introduced faults during the approved test window. GFT proposes stopping the candidate process, blocking its upstream connection, supplying an invalid event and exhausting its allocated temporary storage in an isolated test environment. The expected result is a visible limitation and preserved accepted capture, followed by a bounded recovery action. A green process indicator during an invalid-data test is a failed monitoring test, even if the application continues to answer requests.

A health view should help the person receiving the alert decide what to do next. GFT proposes testing that view with the operator who will use it, asking them to identify the affected function, last trustworthy observation and permitted fallback. If they need a developer to interpret an unexplained error code, the operating interface has not passed. The repair may be clearer instructions or a more informative warning rather than another monitoring platform. Record unresolved questions while the simulated fault is still visible.

After a release, GFT recommends reviewing unexpected behaviour against the recorded baseline before changing several variables together. If a temporary fallback becomes routine, its operating burden should appear in the acceptance review. Repeated manual reclassification or record repair can make an apparently available service unsuitable for expansion. Acceptance should remain conditional on the service continuing to meet its decision contract; approval is not a reason to suppress the evidence that would justify withdrawing it.

11.5

Local analytics reduce unnecessary transfer and administration

An industrial energy-monitoring study's publisher-deposited abstract describes retaining meter readings in an edge database, deriving analytical parameters and filtering what is sent remotely. The accessible abstract supports this placement pattern, rather than a general saving or reliability guarantee. GFT recommends considering local persistence and calculation where the operational decision must continue without an external link, while keeping export and later reconciliation available. [S1102]

The W3C architecture also describes edge preprocessing, filtering and aggregation. GFT's placement recommendation follows that distinction: send a remote consumer the information it needs, with identity and quality preserved, rather than assuming that every raw observation must cross the external link immediately. Retaining detail locally can support investigation, but the design must still define storage ownership, retention and retrieval. Lower transfer volume is not evidence that the application has become simpler to operate. [S14]

For a proposed placement review, the service owner should compare required response time, transferred volume, available compute, maintainable dependencies, update access and failure consequence. A local inspection decision may justify local compute while a periodic cross-site comparison may justify remote processing. The accepted arrangement should identify where the authoritative raw record lives and what happens when each location is unavailable. The objective is an operable division of responsibility, not a rule that all applications must share a location.

WORKED EXAMPLE: EVENT SUMMARIES REDUCE THE ASSUMED TRANSFER VOLUME

An illustrative factory application produces 12 events per second, each occupying an assumed 500 bytes of application payload. Over an illustrative 8-hour shift, the raw payload is 12 multiplied by 500 multiplied by 8 multiplied by 3,600, giving 172,800,000 bytes, or 172.8 decimal MB. These are constructed inputs to explain sizing, not measured factory traffic. Headers, encryption, acknowledgements, compression and retransmission are excluded.

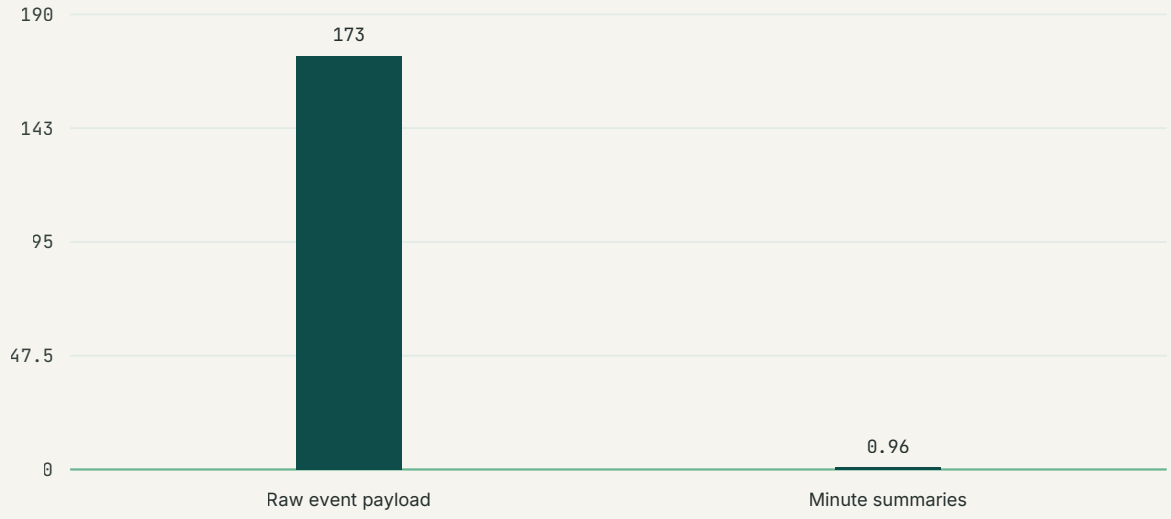
In the same illustrative case, a remote summary contains an assumed 2,000 bytes each minute. There are 8 multiplied by 60, or 480 summary intervals, giving 960,000 bytes, or 0.96 decimal MB. The difference is 171.84 MB of application payload and the ratio is 180 to 1. This compares different information products: the summary cannot replace the raw event record for detailed investigation. The example assumes equal shift coverage and no additional remote raw-data request.

If the illustrative local service is unavailable for 18 minutes while acquisition continues, 12 multiplied by 60 multiplied by 18 produces a backlog of 12,960 events. A demonstrated replay capacity of an assumed 30 events per second would leave 18 events per second for clearing the backlog after allowing for 12 continuing live events per second. Catch-up would take 12,960 divided by 18, or 720 seconds, equivalent to 12 minutes. The replay assumption must be measured under simultaneous live load before it can support an operating commitment.

The illustrative comparison in **Exhibit 52**, p. 183 shows only the application-payload arithmetic. It cannot establish a bandwidth tariff, storage purchase or financial saving. GFT recommends repeating the sizing with measured message overhead, peak bursts, retention needs and replay traffic, then checking available disk space and link capacity under the same load boundary.

EXHIBIT 52 / ILLUSTRATIVE DECIMAL MB PER ASSUMED SHIFT

Summaries carry less payload while preserving less detail



ILLUSTRATIVE Illustrative arithmetic. Raw: 12 events/s x 500 bytes x 8 h x 3600 s/h / 1000000 = 172.8 MB. Summary: 2000 bytes/min x 8 h x 60 min/h / 1000000 = 0.96 MB. Difference 171.84 MB; ratio 180:1. Excludes transport overhead and later detail requests. Detail must remain governed locally.

The illustrative backlog arithmetic assumes that replay and live acquisition share capacity predictably. GFT recommends checking that assumption with the actual persistence and acknowledgement behaviour, because the accepted rate must include successful durable handling rather than attempted transmission. An interrupted replay should restart from a verifiable boundary and tolerate repeated delivery without counting another physical event. The operating procedure should identify which records are waiting, which are confirmed and which require investigation. A falling queue length alone does not prove that the destination retained the intended evidence.

For the proposed release decision, local calculation should also pass a support test. The team must demonstrate how to inspect the service, replace the accepted artefact, retrieve records and recover configuration without depending on an unavailable remote specialist. Remote administration should remain an approved access arrangement with a local means of ending the session. Where the team cannot sustain those duties, a simpler service boundary or a supported managed component may be preferable to additional local computation.

11.6

Rollback is part of security and operational change control

NIST's cybersecurity framework offers outcomes for risk management and explicitly leaves the implementation method to the organisation. GFT's release pack and decision gates are proposed ways of making ownership inspectable, rather than requirements quoted from the framework. ISO's description of its information-security management standard provides the complementary management-system context. Using these references does not establish certification, conformity or a completed security assessment for GFT or any proposed factory. [S18] [S24]

GFT recommends separating the permission to install a candidate from permission to depend on it. The release owner identifies the exact artefact and affected functions. The production owner confirms that the timing, fallback and material handling are workable. The data owner confirms compatibility and reconciliation. A small team may combine duties, but the record should retain each decision explicitly and obtain a separate check for changes that could compromise evidence or production. Unavailable cover is a reason to defer the release, not omit the recovery responsibility.

Under the proposed rollback method, the team first removes the candidate's authority and preserves the observations and logs surrounding the fault. It then activates the accepted predecessor or approved manual mode, checks that incoming events remain retained, and verifies the operational display against an independent reference. Finally, it reconciles events and decisions produced during the interruption. Simply restarting an older process is insufficient when a candidate has already issued instructions, changed records or created unresolved material dispositions.

GFT recommends declaring the data cutover boundary before every pilot. Record the last accepted source event, the last result acted upon and the point from which replay is permitted. Preserve stable event identities so that a repeated transmission does not create another production occurrence. Recomputed results should be marked as recomputed, and physical actions already completed should remain historical facts. The predecessor must be tested against candidate-period records; if it cannot interpret them, restoration requires a governed conversion or retained compatible path.

NIST's contingency-planning guidance provides a basis for evaluating recovery priorities. GFT proposes translating that into a rehearsal performed by the people expected to recover the application, using the approved instructions and available access. The rehearsal should report elapsed restoration time, unreconciled records, operational exceptions and any dependency that prevented completion. A recovery objective is an agreed target; it becomes demonstrated capability only for the conditions actually exercised. [S22]

A compromised predecessor must not be restored merely because it was previously accepted. GFT recommends verifying the integrity and current permission of the return package before activation, using a clean recovery path if compromise is suspected. The accepted fallback may then be an approved manual method while the

service is rebuilt. This is why the release record should name a recovery decision owner as well as a package location. Technical convenience cannot determine whether the factory may depend on an artefact whose integrity is unresolved.

The release record should distinguish the states reached: package verified, installed, observing, permitted for a pilot, accepted for the stated scope, withdrawn and reconciled. GFT recommends recording the evidence and responsible role at each transition. Supplier access should end when its approved maintenance purpose ends, and corrective work should link back to the failed acceptance condition. A successful rollback proves that a particular return path worked under the exercised conditions. It does not prove that the candidate is fit to deploy again.

GFT recommends closing the change with a short operating review after the relevant production cycle. Compare the accepted hypothesis with retained observations, record support effort and reconcile outstanding exceptions. Keep the candidate's withdrawal history visible even if a later release succeeds. The next release should inherit the lessons and current compatibility tests, not a general exemption earned by a previous installation. Where evidence remains inconclusive, maintain the bounded scope and state which observation or recovery test would permit a wider decision.

For the chapter's final acceptance decision, GFT proposes asking whether the factory can continue making and recording the product while the application is removed. If the answer depends on untested compatibility, undocumented support or reconciliation that nobody owns, the change remains a candidate. Keep its scope bounded until those conditions are resolved. This preserves the purpose of Warp and Signal: applications may change, but the connection between production, measurement and accountable action must remain inspectable.

A release earns wider trust when its return path preserves the evidence.

CHAPTER SOURCE NOTES

S30 MDPI, Sustainability. 2023-08-19.

S18 National Institute of Standards and Technology. 2024-02-26.

S16 European Telecommunications Standards Institute. 2024-04.

S17 National Institute of Standards and Technology. 2023-09.

S06 International Organization for Standardization. 2021-10.

S41 International Organization for Standardization. 2016-10.

S1101 MDPI, Electronics; publisher-deposited abstract accessed through Crossref. 2025-09-18.

S31 MDPI, Sensors. 2020-04-28.

S1102 MDPI, IoT; publisher-deposited abstract accessed through Crossref. 2024-09-28.

S14 World Wide Web Consortium. 2023-12-05.

S24 International Organization for Standardization and International Electrotechnical Commission. 2022-10.

S22 National Institute of Standards and Technology. 2010-11-11.

12

PART IV: PROVE VALUE AND SCALE SAFELY / SUSTAINABILITY EVIDENCE

Resource and product data need a defensible metering boundary

This chapter specifies evidence capture rather than accounting or compliance methods: metres, denominators, allocations and life-cycle boundaries must make every displayed result reconstructable.

CHAPTER 12
DECISION TEST

This chapter specifies evidence capture rather than accounting or compliance methods: metres, denominators, allocations and life-cycle boundaries must make every displayed result reconstructable.

Energy, water, material-loss and carbon numbers displayed on a factory dashboard become operationally and externally defensible only when every figure is anchored to a declared physical metering boundary, a synchronised production denominator, a versioned allocation rule and an explicit organisational or life-cycle scope. [S08] [S36] [S37] [S38] A chart that plots kilowatt-hours, cubic metres of water or tonnes of carbon dioxide equivalent without recording which utility feeders, process lines, backup generators, scrap fractions and life-cycle stages sit inside the calculation cannot support either engineering root-cause analysis or external disclosure. [S08] [S29] [S36] Green Future Tech Limited (GFT), Hong Kong SAR, therefore designs resource and product data capture as an integral branch of the factory's governed event spine rather than a parallel sustainability spreadsheet populated at month-end. [S01] [S08] [S15]

This chapter specifies the shop-floor and edge-to-historian evidence-capture architecture required under GFT's Cameroon textile programme, prepared 31 January 2024 with source-refresh annotations retained in 2026. [R06] It separates physical evidence capture from downstream reduction, accounting and compliance methodologies: within GFT's white-paper library, wp06 develops operational resource-efficiency engineering, wp16 owns corporate and product greenhouse-gas accounting methodologies, and wp05 and wp17 specify supply-chain traceability and digital product passport frameworks. [S15] [S36] [S39] In Cameroon, where the inherited sector baseline records roughly 95 percent of cotton fibre exported raw and about 5 percent processed locally against a national 50 percent local-processing target by 2030, industrialising domestic spinning, weaving, wet processing and garment manufacture requires factories that can prove their resource discipline and product provenance from the first shift. [R06] [R07] This chapter provides engineering reference architecture rather than investment, legal or tax advice, and does not certify any facility, claim or passport. [S24] [S39]

12.1

The metre hierarchy follows the energy decision

An industrial metering network should be structured around the operational decisions and management reviews it serves, not around whatever communication ports happen to ship on a switchboard. [S01] [S37] ISO 50001:2018 provides an international framework for establishing, implementing, maintaining and improving an energy management system (EnMS), linking energy data collection to energy performance evaluation, operational control of significant energy uses and continual management review. [S37] When a factory installs isolated utility meters solely to populate an executive summary screen, supervisors cannot tell whether an energy spike originated in a leaking compressed-air ring main, an idling stenter frame, a dyehouse steam trap failure or an unscheduled transfer to backup diesel generation. [S29] [S37]

Research on industrial Internet of Things (IIoT) energy monitoring demonstrates why edge persistence and local calculation belong between shop-floor meters and central dashboards. [S29] In a documented industrial edge energy-monitoring implementation, once the edge node receives meter data, it stores the readings in a local database server and runs a data-processing method to derive nine additional analytical parameters locally before forwarding selected values. [S29] GFT adapts this edge-first

pattern into a four-tier metering hierarchy aligned with the physical-asset organisation of ISA-95 Part 1 and the manufacturing operations boundary of IEC 62264-1. [S01] [S02] [S29] [S37]

12.2 SUBMETERING SEPARATES BASE LOAD, PROCESS UTILITIES AND LINE CONSUMPTION

To prevent plant-wide totals from masking line-level waste, GFT partitions electrical, thermal and fluid metering into four physical tiers mapped to the factory's asset hierarchy. [S01] [S29] [S37] Each meter point carries a stable asset identifier, an Asset Administration Shell (AAS) submodel describing its accuracy class, calibration date, scaling ratio and unit, and an explicit parent-child relationship in the metering tree: [S01] [S07] [S37]

- **Tier 1: Site supply and generation boundary.** Revenue-grade and check meters sit at the utility grid incomer, transformer secondaries, diesel generator breakers, solar combiner panels and boiler fuel headers to measure total site energy. [S29] [S37] [S1201] [S1202]
- **Tier 2: Shared utility and auxiliary plant headers.** Submeters measure central systems serving multiple areas: the air-compressor station, steam boiler house and condensate loop, chilled-water and air-washer humidification bays, and the effluent treatment plant. [S01] [S29] [S37]
- **Tier 3: Production line and work-centre feeders.** Dedicated submeters track electricity, steam and compressed air delivered to each line or work centre, including blowroom and carding, ring spinning, warping and sizing, weaving sheds, dyeing and stenter ranges, cutting tables and sewing lines. [S01] [S02] [S29]
- **Tier 4: Significant energy uses and diagnostic sub-loops.** Targeted sensors instrument high-draw drives, dyehouse heat-recovery exchangers, lint fans and pneumatic manifolds identified under the ISO 50001 review. [S10] [S29] [S37]

This structure captures every kilowatt-hour or kilogram of steam where an operator can act, pairing Tier 3 line feeds with Tier 2 utility conversion efficiency. [S08] [S29] [S37]

12.3 EDGE PERSISTENCE PROTECTS CUMULATIVE REGISTERS AGAINST ROLLOVER AND LINK LOSS

Textile plants read panel power meters, thermal flowmeters and manifolds through MODBUS polling or cyclic IO-Link exchange. [S09] [S10] The edge gateway manages MODBUS byte order, transformer multipliers and timeouts without blocking controllers, while IO-Link keeps cyclic process data separated from on-request diagnostics so calibration checks never interrupt metering. [S09] [S10]

Polling instantaneous power (kilowatts) at irregular intervals and multiplying by elapsed time introduces aliasing error when motors cycle and loses consumption during link drops. [S08] [S29] [S40] GFT instead requires the gateway to read non-volatile cumulative active and reactive energy registers (watt-hours and var-hours), store stamped readings in the local edge database, and compute interval deltas between verified states. [S09] [S29] [S40] Gateways discipline local clocks through Network Time Protocol Version 4 (NTPv4), achieving tens-of-microseconds potential accuracy on modern workstations and fast local area networks, or IEEE 1588-2019 Precision Time Protocol (PTP) where sub-microsecond alignment is required. [S12]

[S13] Over MQTT with Sparkplug 3.0.0, birth and death certificates communicate session state; offline meters are marked stale rather than zero, and historical flush upon reconnection restores exact consumption across the gap. [S05] [S29] [S40]

To make the vertical relationship between physical meters, operational denominators and reporting views concrete, **Exhibit 53**, p. 189 sets out GFT's five-layer metering and allocation architecture. [S08] [S29] [S36] [S37] [S38]

EXHIBIT 53 / PROPOSED ARCHITECTURE LAYERS

The metering hierarchy connects supply, process and product evidence



ILLUSTRATIVE [S08] [S29] [S36] [S37] [S38] Illustrative GFT metering and allocation architecture informed by ISO 22400-2 KPI elements (S08), edge energy database and parameter derivation (S29), GHG Protocol Product Standard boundaries (S36), ISO 50001 EnMS structure (S37) and ISO 14046 water-footprint scopes (S38). Illustrative design model, not an observed factory installation.

As **Exhibit 53**, p. 189 shows, no reporting view may read raw supply or submeter registers directly without passing through the production denominator spine and the versioned allocation engine. [S08] [S36] [S37] Because a diagram cannot verify physical current-transformer polarity or un-metered bypass pipes, commissioning requires a physical balance check comparing Tier 1 supply against the sum of Tier 2 and Tier 3 submeters. [S29] [S37]

12.4

Electricity source and backup state belong in the context

A kilowatt-hour consumed on a textile line does not carry a uniform cost, operational risk or greenhouse-gas intensity across time. [S29] [S32] [S34] [S1201] According to the International Renewable Energy Agency (IRENA), installed power-generation capacity in the Central African Power Pool (CAPP) region stood at just over 11 gigawatts (GW) at the start of 2023, of which hydropower made up roughly 75 percent. [S34] Taken out of context, that regional hydropower share might tempt a software vendor to multiply a Cameroonian factory's annual kilowatt-hour consumption by a static hydro-heavy factor and claim low-carbon production. [S34] [S36] [S1201]

Regional installed capacity, however, is neither a Cameroonian grid-delivery emission factor nor a measure of what powers a factory when the public grid fails. [S32] [S34] In the World Bank's Cameroon 2024 Enterprise Survey, 93.3 percent of surveyed formal private non-agricultural firms reported experiencing electrical outages, with 10.4 outages in a typical month. [S32] Although those survey figures cover all formal non-agricultural sectors rather than textile mills alone, they establish that backup generation and outage transitions are routine operating conditions. [S32] When a plant runs on-site diesel generators or hybrid solar-plus-storage systems during grid interruptions, the energy source, cost and emission boundary shift immediately. [S29] [S32] [S1201] [S1202]

Exhibit 54, p. 191 contrasts the regional CAPP generation-capacity baseline reported by IRENA with the firm-level outage incidence and frequency recorded in the Cameroon 2024 Enterprise Survey. [S32] [S34]

EXHIBIT 54 / REGIONAL CAPACITY AND NATIONAL SURVEY INDICATORS

Hydropower dominates Central Africa's regional installed capacity

>11 GW

CAPP installed capacity

Central African Power Pool total at start of 2023 (S34)

~75%

Regional hydropower share

Share of CAPP installed capacity at start of 2023; not a site factor (S34)

93.3%

Cameroon firms with outages

All surveyed formal firms in 2024 Enterprise Survey (S32)

10.4

Outages in a typical month

Cameroon all-firm monthly frequency in 2024 survey (S32)

SOURCED [S34] [S32] Sourced from IRENA Central African Power Pool (CAPP) start-of-2023 regional capacity data (S34) and World Bank Cameroon 2024 Enterprise Survey all-firm outage metrics (S32). Regional capacity is not a site emission factor, and all-firm survey metrics are not textile-specific measurements.

The juxtaposition in **Exhibit 54**, p. 191 shows why a connected factory cannot rely on an uncontextualised energy total. [S32] [S34] Neither source establishes the exact grid-to-generator ratio of any specific Cameroonian mill; what they establish is an architectural requirement to meter each electricity source separately and bind source state to every production interval. [S29] [S32] [S34]

12.5 TRANSFER-SWITCH STATE SEPARATES PURCHASED GRID POWER FROM ON-SITE GENERATION

Under the GHG Protocol Corporate Accounting and Reporting Standard and the Scope 2 Guidance, stationary fuel combustion in on-site generators belongs to direct Scope 1 inventories, whereas purchased or acquired electricity, steam, heat and cooling belong to indirect Scope 2 inventories. [S1201] [S1202] At the product level, the GHG Protocol Product Standard requires activity data to reflect the actual energy carriers consumed inside the product boundary. [S36]

If a Tier 3 submeter on a spinning frame or dyeing range accumulates kilowatt-hours while the automatic transfer switch (ATS) toggles between utility grid, diesel generators and solar arrays, the downstream register mixes purchased, renewable and diesel-generated electricity. [S29] [S1201] [S1202] GFT therefore requires the edge layer to capture three synchronised signals at the Tier 1 boundary: [S29] [S37]

1. **Source-bus active energy and fuel telemetry.** Separate meters measure utility grid import, generator electrical output, calibrated diesel fuel flow (supply-minus-return differential metering and tank-level verification), and net alternating-current output from solar inverters or battery converters. [S09] [S29] [S37]
2. **Transfer-switch and breaker state events.** Auxiliary contacts on the ATS, tie breakers and generator switchgear publish state-change events with NTPv4 or PTP source timestamps recording the exact second a bus transitions between grid, diesel, hybrid or blackout states. [S12] [S13] [S40]
3. **Interval source-fraction vector.** For every accounting interval, the edge service calculates the fraction of bus electricity supplied by grid import, diesel generation and solar generation, attaching that vector to every Tier 2 and Tier 3 submeter delta on that bus. [S29] [S1201]

Preserving generator fuel input and the bus source-fraction vector lets engineers track generator specific fuel consumption (litres per kilowatt-hour) under ISO 50001 while giving wp16 clean inputs for Scope 1 fuel combustion and Scope 2 purchased electricity. [S29] [S37] [S1201] [S1202]

12.6 OUTAGE TRANSITIONS AND MISSING INTERVALS REQUIRE EXPLICIT QUALITY FLAGS

Power interruptions are often preceded by voltage sags and followed by generator cranking delays, load-shedding sequences and staggered motor restarts, during which ovens lose air circulation and dyeing vats hold heated liquor without advancing the batch. [S08] [S29] [S32] If submeters or wireless gateways lack uninterruptible power supply (UPS) backing, register reads are delayed and wireless or RFID nodes may lose time synchronisation across shared radio spectrum. [S29] [S31] [S40] GFT requires meters, IO-Link masters, gateways and switches to ride through transfer windows on backed control power, tagging any missed submeter interval as `stale`, `unsynchronised`, `outage-bridged` or `missing` rather than linearly interpolating power across a dead bus. [S10] [S29] [S31] [S40]

12.7

Water volume is not automatically a water footprint

Water is a critical process medium across textile production, from warp sizing to scouring, bleaching, mercerising, dyeing, washing-off, steam raising and humidification. [S27] [S38] Yet many dashboards display a volumetric flow reading in cubic metres and label the tile "Water Footprint", conflating an operational withdrawal meter with a life-cycle environmental impact assessment. [S27] [S38]

ISO 14046:2014 specifies principles, requirements and guidelines for water-footprint assessment of products, processes and organisations based on life-cycle assessment (LCA). [S38] Under ISO 14046, a water footprint quantifies potential environmental impacts related to water, grounded in a declared scope (product, process or organisation), geographic and temporal boundary, water-availability context and water-degradation categories. [S38] Conversely, a Pakistani textile IoT case study shows how monitoring real-time water-level fluctuations in tanks enables teams to detect abnormal drawdowns and optimise operational water usage. [S27] Both activities are useful, but they answer different questions at different boundaries. [S27] [S38]

12.8 INTAKE, RECYCLE AND EFFLUENT FLOWS NEED DISTINCT MASS-BALANCE POINTS

To support operational water management (wp06) while capturing foreground inputs for life-cycle assessments (wp16), GFT structures factory water metering across five hydrological boundaries: [S27] [S38]

- **Raw water withdrawal by source.** Separate flowmeters and tank-level sensors measure withdrawal from municipal mains, boreholes, surface intakes and rainwater storage, preserving source identity. [S27] [S38]
- **Pre-treatment and boiler feedwater.** Meters across filtration, softening and reverse-osmosis skids separate permeate delivered to processes from reject brine discharged to drain. [S10] [S27]
- **Process-line and batch water delivery.** Electromagnetic or ultrasonic flowmeters and batching valves on sizing boxes, washing ranges and dyeing machines record cold water, hot water and condensate delivered to each order. [S08] [S10] [S27]
- **Internal recovery, reuse and condensate return.** Dedicated meters track cooling-water recirculation, counter-current washwater reuse and condensate recovery to separate gross process demand from net freshwater withdrawal. [S27] [S38]
- **Effluent treatment and final discharge.** Flowmeters and inline transmitters (pH, temperature, conductivity, turbidity and dissolved oxygen over IO-Link or MODBUS) monitor raw effluent, recycled utility water and treated discharge, paired with laboratory sample IDs for chemical and biological oxygen demand. [S09] [S10] [S38]

Closing the volumetric balance between withdrawal, tank storage change (S27), evaporative loss, recycled permeate and effluent discharge lets operators spot pipe leaks, stuck drain valves or unrecorded wash cycles immediately. [S27] [S38]

12.9 OPERATIONAL WITHDRAWAL TELEMETRY STOPS SHORT OF LIFE-CYCLE IMPACT CLAIMS

Even a complete factory water balance remains a gate-to-gate operational inventory rather than a full product water footprint, because cotton irrigation and chemical manufacture occur upstream and watershed stress differs between northern Cameroon around Garoua and coastal zones near Douala or Dibamba. [R03] [R10] [S38] Shop-floor tiles are therefore labelled with their physical quantity and boundary, such as "Freshwater withdrawal intensity (litres per kilogram of dyed fabric, gate-to-gate)", carrying the ISO 22400 formula and source-meter list, while the term "water footprint" is reserved for ISO 14046 assessments that include life-cycle impact characterisation. [S08] [S27] [S38]

GFT NOTE

Never allow a dashboard, customer portal or sustainability slide to relabel raw utility-meter totals as a "Product Carbon Footprint" or "Water Footprint". [S36] [S38] Configure the KPI dictionary so every resource tile displays four mandatory qualifiers alongside its numeric value: the physical quantity and unit (`kWh`, `litres`, `kg steam`), the system boundary (`Line 2 ring spinning, gate-to-gate`), the production denominator (`per kg first-pass good yarn`), and the source coverage and quality state (`99.4% metered, 0.6% outage-bridged`). [S08] [S29] [S40] Where life-cycle impact characterisation or upstream agricultural inventory is absent, state that boundary explicitly on the screen. [S36] [S38]

12.10

Product carbon inputs begin with a product and reference flow

Sponsors and export buyers increasingly ask textile factories for product-level greenhouse-gas figures per kilogram of yarn, per square metre of fabric, or per garment. [S36] [S39] The GHG Protocol Product Life Cycle Accounting and Reporting Standard (the Product Standard) provides requirements and guidance for organisations to quantify and publicly report an inventory of GHG emissions and removals associated with a specific product. [S36] While wp16 details the accounting rules and emission-factor selection, the connected-factory architecture in this paper supplies the primary shop-floor activity records in a form that satisfies the Product Standard's structural prerequisites. [S08] [S36]

Dividing monthly facility electricity bills by total garments shipped that month and calling the result a product carbon footprint mixes heavy and lightweight constructions, ignores work-in-process inventory changes, omits upstream cotton and chemical stages, fails to separate grid and diesel power during outages, and conceals energy wasted on scrapped batches. [S08] [S32] [S36]

12.11 GATE-TO-GATE PROCESS RECORDS MUST DECLARE WHAT LIES OUTSIDE THE FACTORY WALL

Under the GHG Protocol Product Standard, a product inventory is organised around a defined product, a reference flow, and an explicit life-cycle boundary (or declared cradle-to-gate subset for intermediate products). [S36] Within GFT's architecture, factory records capture the gate-to-gate foreground production stage (or multiple stages in an integrated complex such as the announced Camtext SA model in Cameroon, which remains a third-party project context rather than a GFT deployment). [R09] [R10] [S02] [S36] Each batch record preserves five boundary attributes: [S08] [S36]

1. **Product specification and reference flow.** Material ID, yarn count or fabric areal density, garment style and reference flow unit (for example, 1,000 kilograms of carded ring-spun yarn or 100 finished jackets). [S07] [S08] [S36]
2. **Included unit processes.** The ordered sequence of ISA-95 Level 3 operations executed on the lot, linked to work centres and time windows. [S01] [S02] [S36]

3. **Metered foreground inputs and outputs.** Physical quantities of input lint or fabric, chemicals, dyes, grid electricity, diesel electricity, solar electricity, steam and water consumed, plus recovered waste, hard waste and effluent. [S08] [S29] [S36]
4. **Explicit boundary exclusions.** A machine-readable list of stages outside the factory's metering boundary, such as seed-cotton cultivation and ginning in northern Cameroon (R01, R03), chemical synthesis, external transport and consumer use, blocking downstream systems from mistaking gate-to-gate data for a full life-cycle total. [R01] [R03] [S36]
5. **Data provenance and quality ratings.** Indicators showing whether each input was directly submetered, allocated using a physical key, or estimated during a sensor outage, paired with ISO 8000-120 master-data provenance. [S08] [S36] [S41]

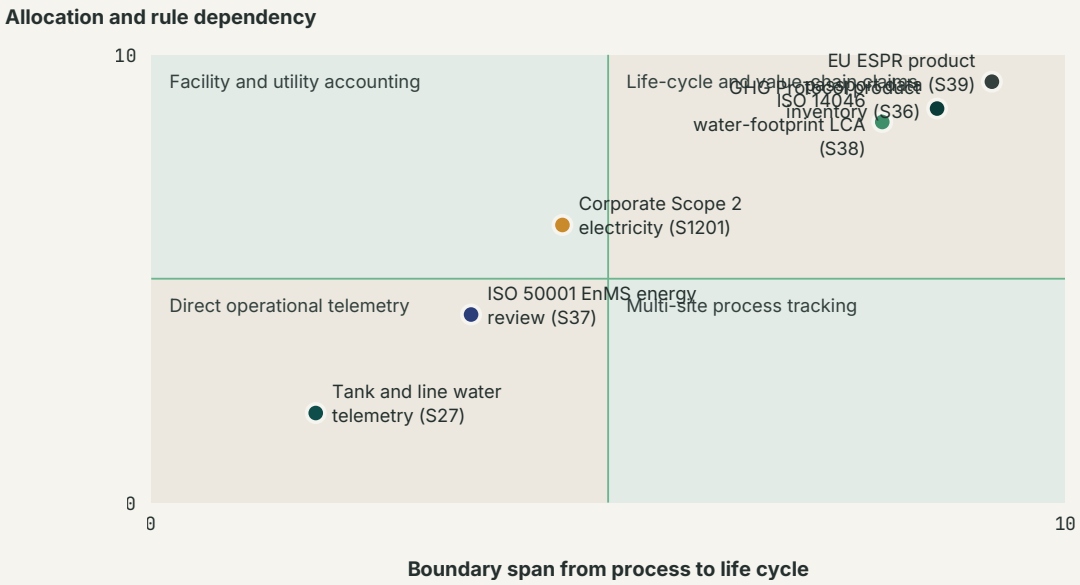
12.12 VERSIONED EMISSION FACTORS AND MASTER-DATA PROVENANCE REMAIN SEPARATE FROM RAW TELEMETRY

Power meters measure kilowatt-hours, fuel meters measure litres, and steam meters measure kilograms of vapour; none measures carbon dioxide equivalent directly. [S09] [S10] [S29] [S36] If an edge script multiplies kilowatt-hours by a hard-coded emission factor before storage, any later factor revision invalidates stored history. [S08] [S24] [S36] GFT prohibits embedding emission factors in edge logic or raw telemetry: the historian stores native physical units with timestamps and quality flags, while emission factors, global warming potential (GWP) versions and characteristic master data reside in a versioned repository under ISO/IEC 27001 change control and ISO 8000-120 provenance rules. [S24] [S29] [S36] [S41]

Exhibit 55, p. 196 maps the six resource and product domains across their physical-to-life-cycle boundary span and their dependence on allocation and regulatory rules beyond raw sensor telemetry. [S27] [S36] [S37] [S38] [S39] [S1201]

EXHIBIT 55 / DERIVED BOUNDARY AND DEPENDENCY SCORE (0 TO 10)

Each sustainability claim has a different evidence boundary



DERIVED [S27] [S36] [S37] [S38] [S39] [S1201] Derived qualitative comparison of six resource and product claims from the stated scopes of Pakistani textile IoT water telemetry (S27), ISO 50001 EnMS (S37), GHG Protocol Scope 2 Guidance (S1201), ISO 14046 water-footprint LCA (S38), GHG Protocol Product Standard (S36) and EU Regulation 2024/1781 digital product passport (S39). X-axis scores physical boundary span from machine/tank to multi-tier supply chain; Y-axis scores allocation, characterisation and regulatory rule dependency.

As **Exhibit 55**, p. 196 shows, tank water telemetry (S27) and ISO 50001 energy reviews (S37) sit inside machine and site boundaries where physical metering dominates, whereas Scope 2 accounting (S1201), ISO 14046 water footprints (S38), GHG Protocol product inventories (S36) and EU digital product passports (S39) require multi-tier boundaries and formal allocation governance. [S27] [S36] [S37] [S38] [S39] [S1201] GFT prohibits jumping from the lower-left quadrant (S27, S37) to a claim in the upper-right quadrant (S36, S38, S39) without executing the intervening boundary, denominator and allocation steps. [S08] [S36] [S38]

12.13

Allocation uses the same production denominator as operations

Whenever a factory calculates resource intensity or allocates shared utility loads to batches, the numerator comes from the metering hierarchy while the denominator comes from the manufacturing operations record. [S08] [S29] [S36] ISO 22400-2:2014 specifies manufacturing KPIs through their formula, source elements, time behaviour and unit. [S08] GFT enforces denominator unity by requiring every resource-intensity KPI and batch-allocation routine to draw its denominators from the same governed ISO 22400 elements used for throughput, quality and overall equipment effectiveness (OEE). [S08]

12.14 SHARED UTILITIES REQUIRE EXPLICIT PHYSICAL ALLOCATION KEYS ACROSS CONCURRENT BATCHES

While a Tier 3 meter on a ring-spinning frame or jet-dyeing vessel links directly to the active lot, Tier 2 central utilities serve multiple lines concurrently: boilers supply common steam headers (S01, S37), compressors feed shared air mains (S01, S29), humidification plants condition entire spinning or weaving halls (S01, S37), and lighting and effluent plants serve the facility (S29, S38). [S01] [S29] [S37] [S38] Following the GHG Protocol Product Standard preference for process subdivision, direct metering and physical relationships (S36), GFT's allocation engine applies a three-level hierarchy per interval: [S08] [S29] [S36]

1. **Direct batch submetering.** Where a Tier 3 or Tier 4 meter serves a single identified batch between `t_start` and `t_end`, 100 percent of that interval is assigned to that batch. [S08] [S15] [S29]
2. **Physics-weighted line and header allocation.** Where a shared header serves multiple active batches without branch meters, the interval volume is allocated in proportion to a declared physical key, such as processed mass multiplied by a recipe-specific engineering factor (liquor ratio and temperature for dyehouse steam, or picks-per-centimetre and reed width for loom compressed air). [S08] [S29] [S36]
3. **Explicit facility base-load and residual pool.** Lighting, idle-period humidification and any positive difference between Tier 1 supply and summed Tier 2/3 submeters sit in a transparent residual pool, allocated where required by wp16 using a versioned rule while remaining visible as an overhead line item. [S08] [S29] [S36] [S37]

12.15 DENOMINATOR DRIFT DISTORTS UNIT INTENSITY WHEN SCRAP AND REWORK ARE HIDDEN

In any textile process, raw input (`Q_in`), gross processed output (`Q_gross`), reworked output (`Q_rework`), rejected scrap (`Q_scrap`) and net good output (`Q_good`) diverge whenever defects occur. [S08] When a dyehouse re-dyes an off-shade batch after an outage restart (S32), re-dyeing consumes a second full cycle of water, steam and electricity without adding net good fabric. [S08] [S27] [S29] [S32] Dividing machine energy by `Q_gross` (counting the re-dyed batch twice) masks the waste, whereas low line utilisation under fixed hall humidification raises per-unit intensity without machine degradation. [S08] [S29] [S37] GFT therefore stores and displays three metrics side by side under ISO 22400-2: **gross process intensity** (divided by `Q_gross`), **net good-unit intensity** (including rework and scrap burden, divided by `Q_good`), and the **base-load versus variable split**. [S08] [S29] [S36] [S37]

METHOD

Before publishing any shift, batch or monthly resource-intensity or product-input figure, the allocation and KPI service executes four automated reconciliation checks against the historian record. [S08] [S29] [S36]

1. **Meter tree conservation check.** For each carrier (`kWh`, `kg steam`, `m3 water`, `Nm3 compressed air`), verify that Tier 2 headers and direct Tier 3 feeders plus explicit residual equal Tier 1 supply (`E_tier1 = sum(E_sub) + E_residual`) within tolerance. [S29] [S37]
2. **Source-power and outage boundary check.** Confirm that every electrical interval partitions across grid, diesel and solar (`f_grid + f_diesel + f_solar = 1.0`) using ATS timestamps, flagging unbuffered meter loss as `incomplete-coverage`. [S29] [S32] [S40] [S1201]
3. **Production denominator mass balance.** Reconcile ISO 22400 elements so gross output equals good output plus rework and scrap (`Q_gross = Q_good + Q_rework + Q_scrap`), enclosed by EPCIS or RFID batch timestamps. [S08] [S15] [S25]
4. **Allocation conservation check.** Confirm that summing allocated batch and overhead quantities reproduces 100 percent of Tier 1 metered totals without double-counting or orphan loss. [S08] [S36]

WORKED EXAMPLE: BATCH ALLOCATION AND DENOMINATOR RECONCILIATION ACROSS A DYEHOUSE SHIFT

Consider an illustrative 8-hour shift in a wet-processing department running Batch A and Batch B through dyeing and stenter drying. All figures are explicitly illustrative to demonstrate GFT's calculation method, not measured GFT results or forecasts. [S08] [S29] [S36]

1. Assumed Tier 1 supply and outage split. Over the illustrative 8-hour shift, the department bus records an assumed 4,000 kWh. During a 96-minute grid outage (S32), the ATS routes power from a diesel generator: Tier 1 meters record 3,200 kWh grid electricity (80 percent) and 800 kWh diesel electricity (20 percent) from 220 litres of diesel, plus 120.0 m3 (120,000 litres) of water withdrawal and 14,000 kg of header steam. [S27] [S29] [S32] [S37] [S1201] [S1202]

2. Submeter conservation and shared utility allocation. Tier 3 submeters record 2,900 kWh on Line 1 (Batch A) and 700 kWh on Line 2 (Batch B); Tier 2 effluent and ventilation meters record 300 kWh, leaving an unsubmetered residual of 4,000 - 3,900 = 100 kWh (2.5 percent). [S29] [S37] Water meters record 84.0 m3 on Line 1, 26.0 m3 on Line 2, 6.0 m3 boiler make-up and 4.0 m3 washdown (120.0 m3 conserved). [S27] [S38] Under the illustrative physical key (S36), the 14,000 kg header steam is allocated by gross mass times recipe thermal factor: Batch A processes 1,500 kg twill at factor 1.2 (1,800 units, 75 percent = 10,500 kg steam), and Batch B processes 600 kg poplin at factor 1.0 (600 units, 25 percent = 3,500 kg steam). [S08] [S36] The 400 kWh shared electricity (300 + 100 kWh) allocates by active machine-hours (7.5 hours or 75 percent = 300 kWh to Batch A; 2.5 hours or 25 percent = 100 kWh to Batch B), giving 3,200 kWh total electricity for Batch A and 800 kWh for Batch B. [S08] [S29] [S36]

3. Denominator comparison. Batch A starts with 1,350 kg grey fabric; an outage temperature drop (S32) forces a 150 kg re-dye (Q_gross = 1,500 kg), and final inspection rejects 100 kg scrap (Q_scrap = 100 kg), leaving Q_good = 1,250 kg. [S08] Dividing direct inputs (2,900 kWh and 84,000 litres) by Q_gross (1,500 kg) gives an unvalidated 1.93 kWh/kg and 56.0 litres/kg. [S08] Under GFT's governed boundary (S08, S36), dividing Batch A's total allocated electricity (3,200 kWh, comprising 2,560 kWh grid and 640 kWh diesel power from 176 litres diesel), allocated water (84,000 + 7,500 = 91,500 litres) and steam (10,500 kg) by Q_good (1,250 kg) yields 2.56 kWh/kg electricity, 73.2 litres/kg water and 8.40 kg/kg steam; hiding rework, scrap and shared utilities understates electrical intensity by 24.6 percent and water intensity by 23.5 percent. [S08] [S36]

12.16

Product evidence remains linked to events and provenance

Downstream buyers, verifiers and digital product passport systems require product-level records that travel across enterprise boundaries without losing provenance. [S15] [S36] [S39] GFT closes this chain by binding resource-allocation records directly to physical visibility events and durable asset descriptions. [S07] [S15] The GS1 Electronic Product Code Information Services (EPCIS) Release 2.0 standard enables disparate applications to create and share visibility event data within and across enterprises. [S15] Meanwhile, Regulation (EU) 2024/1781 defines a digital product

passport as a set of data specific to a product that includes the information specified in the applicable delegated act adopted pursuant to Article 4 and that is accessible via electronic means through a data carrier in accordance with Chapter III. [S39]

12.17 VISIBILITY EVENTS BIND RESOURCE INTERVALS TO TRACKED LOTS AND GARMENT BUNDLES

Across textile stages, seed cotton becomes lint bales (R03), yarn lots, fabric rolls and cut garment bundles. [R03] [S25] [S28] Shop-floor studies of radio-frequency identification (RFID) in denim work-in-process tracking and RFID-enabled real-time manufacturing execution systems (MES) show how systematic reader deployment captures real-time object identity, location and execution timestamps. [S25] [S28] Each material transition generates an EPCIS 2.0 `TransformationEvent`, `ObjectEvent` or `AggregationEvent` recording what lot IDs were involved, when the step started and ended (under NTPv4 or PTP), where it occurred (mapped to ISA-95), and why (`bizStep` and `disposition`). [S01] [S12] [S13] [S15]

The historian joins the EPCIS event window (`t_start` to `t_end` at `bizLocation`) to the Tier 2 and Tier 3 resource-allocation records for that work centre. [S08] [S15] [S29] When dyed fabric rolls are cut into garment bundles, the EPCIS `TransformationEvent` links parent rolls to child bundles, propagating verified gate-to-gate resource inputs forward while recording cutting-room `marker loss` explicitly. [S08] [S15] [S25] [S36] Linking batch records to Asset Administration Shell (AAS) submodels (S07) lets auditors verify meter calibration validity without bloating MQTT or EPCIS payloads (S15, S40), while ISO 8000-120:2016 governs provenance for characteristic master data alongside ISO/IEC 27001 change-controlled analytical lineage (S24, S41). [S07] [S15] [S24] [S40] [S41]

12.18 EXPORT SCHEMAS PREPARE FOR DIGITAL PRODUCT PASSPORTS WITHOUT INVENTING PREMATURE MANDATES

Because Regulation (EU) 2024/1781 specifies that product-group delegated acts pursuant to Article 4 define the required passport fields, a factory architecture must neither claim today that an internal database is already a certified EU textile passport nor wait to structure batch-level capture until a deadline arrives. [S15] [S39] GFT assembles an immutable, versioned calculation trail in the historian for export to wp05 traceability views, wp16 footprint models or wp17 passport carriers. [S08] [S15] [S36] [S39] **Exhibit 56**, p. 201 defines the ten-field schema accompanying every batch or garment resource record, with no fabricated footprint value. [S08] [S15] [S24] [S36] [S39]

EXHIBIT 56 / PROPOSED EVIDENCE-PACKAGE SCHEMA

A per-garment result retains its calculation trail

Evidence field	Recorded content	Verification rule
1. Reference flow and lot ID	Product SKU, batch or bundle EPC, unit mass and reference flow	Must resolve to valid EPCIS object or transformation event (S15, S36)
2. System and process boundary	Included ISA-95 unit processes, site ID and gate-to-gate scope	Must match active work-centre route in operations log (S01, S02, S36)
3. Production denominators	Gross processed units, rework, scrap and net good units (Q_good)	Reconciled to ISO 22400 shift and batch production counts (S08)
4. Metered physical inputs	Native kWh, diesel litres, steam kg, water m3 and input material kg	Conserved against Tier 1 and Tier 2 meter tree registers (S29, S37)
5. Power source and outage split	Interval fractions for grid, on-site diesel and solar PV supply	Verified against ATS state logs and generator fuel flow (S32, S1201)
6. Shared-utility allocation key	Allocation method, physical weighting key and residual share	Must conserve 100% of shared header volume across lots (S08, S36)
7. Explicit exclusions	Upstream agriculture, chemical synthesis, transport or use stages omitted	Displayed beside any gate-to-gate total to block relabelling (S36, S38)
8. Factor and master-data versions	Emission factor ID, source date, GWP version and ISO 8000-120 provenance	Stored separately from raw telemetry under change control (S24, S41)

ILLUSTRATIVE [S08] [S15] [S24] [S36] [S39] Illustrative GFT calculation-trail schema derived from ISO 22400-2 KPI formula and element discipline (S08), GS1 EPCIS 2.0 event linkage (S15), ISO/IEC 27001 change control (S24), GHG Protocol Product Standard inventory rules (S36) and EU Regulation 2024/1781 passport data-carrier scope (S39). Contains schema definitions only, with no fabricated factory footprint values.

EXHIBIT 57 / PROPOSED EVIDENCE-PACKAGE SCHEMA

A per-garment result retains its calculation trail, continued

Evidence field	Recorded content	Verification rule
9. Coverage and quality flags	Percent directly metered, allocated or estimated; clock and stale flags	Propagated from Sparkplug session and wireless sync state (S31, S40)
10. Analytical lineage and approval	Calculation rule digest, release version, approver role and timestamp	Auditable under ISO/IEC 27001 configuration management (S24)

ILLUSTRATIVE [S08] [S15] [S24] [S36] [S39] Illustrative GFT calculation-trail schema derived from ISO 22400-2 KPI formula and element discipline (S08), GS1 EPCIS 2.0 event linkage (S15), ISO/IEC 27001 change control (S24), GHG Protocol Product Standard inventory rules (S36) and EU Regulation 2024/1781 passport data-carrier scope (S39). Contains schema definitions only, with no fabricated factory footprint values.

Every row in **Exhibit 56**, p. 201 serves a specific verification purpose. [S08] [S24] [S36] If an energy-efficiency project on a stenter frame reduces steam per kilogram of dyed fabric (wp06), Fields 2, 3, 4 and 6 isolate whether the reduction came from burner fuel

flow, line utilisation or reduced re-dyeing rework. [S08] [S29] [S37] If a carbon accountant in wp16 updates an emission factor under the GHG Protocol Scope 2 Guidance (S1201) or Product Standard (S36), Field 4's physical quantities and Field 5's grid-versus-diesel split allow clean recalculation without altering raw historian records. [S24] [S36] [S1201] And when a garment lot requires an exportable package for a digital product passport carrier under wp05 and wp17 (S39), Field 1's EPCIS event chain (S15) and Field 7's explicit exclusions ensure that the factory attests only to what its instruments measured. [S15] [S36] [S39]

While this schema makes resource and product evidence reconstructable from machine edge to reporting view, a schema alone cannot prevent someone from bypassing a meter, altering a calibration multiplier or losing historian records during a cyber incident or power failure. [S17] [S21] [S24] The next chapter therefore turns the reference architecture's security, functional-safety and recovery requirements into witnessed engineering acceptance tests. [S17] [S19] [S22] [S23]

A sustainability figure is only as credible as the physical boundary, production denominator and event trail that can reconstruct it.

CHAPTER SOURCE NOTES

S08 International Organization for Standardization. 2014-01.

S36 GHG Protocol, World Resources Institute and World Business Council for Sustainable Development. 2011-09.

S37 International Organization for Standardization. 2018-08.

S38 International Organization for Standardization. 2014-08.

S29 MDPI, IoT. 2024-09-28.

S01 International Society of Automation. 2025-04-10.

S15 GS1 AISBL. 2022-06.

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

S39 European Parliament and Council of the European Union, Official Journal of the European Union. 2024-06-28.

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

S24 International Organization for Standardization and International Electrotechnical Commission. 2022-10.

S02 International Electrotechnical Commission. 2013-05-22.

S07 Industrial Digital Twin Association. 2024-06.

S1201 GHG Protocol, World Resources Institute and World Business Council for Sustainable Development. 2015-01.

S1202 GHG Protocol, World Resources Institute and World Business Council for Sustainable Development. 2015.

S10 IO-Link Community. 2024-06.

S09 Modbus Organization. 2012-04-26.

S40 Eclipse Foundation. 2022-11-16.

S12 IEEE Standards Association. 2020-06-16.

S13 Internet Engineering Task Force and RFC Editor. 2010-06.

S05 OASIS Open. 2019-03-07.

S32 World Bank Enterprise Surveys. 2025.

S34 International Renewable Energy Agency. 2024.

S31 MDPI, Sensors. 2020-04-28.

S27 MDPI, Applied Sciences. 2024-06-21.

R03 SODECOTON. Publication date unrecorded; inherited source-text check.

R10 Ecofin Agency. Publication date unrecorded; inherited source-text check.

R09 Ecofin Agency. Publication date unrecorded; inherited source-text check.

R01 Business in Cameroon. Publication date unrecorded; inherited source-text check.

S41 International Organization for Standardization. 2016-10.

S25 Springer Nature, International Journal of Advanced Manufacturing Technology. 2016-09-05.

S28 Elsevier, Robotics and Computer-Integrated Manufacturing. 2013-04.

S17 National Institute of Standards and Technology. 2023-09.

S21 CISA and joint international cybersecurity authorities. 2022-05-09.

S19 International Electrotechnical Commission. 2020-06.

S22 National Institute of Standards and Technology. 2010-11-11.

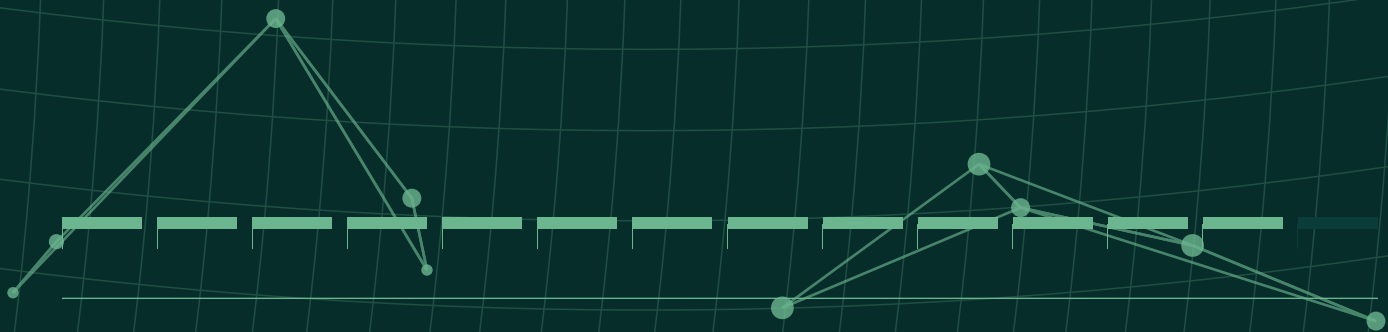
S23 International Electrotechnical Commission. 2022-04-28.

13

PART IV: PROVE VALUE AND SCALE SAFELY / EVIDENCE BEFORE OPERATION

Security, safety and recovery are acceptance tests

A connected factory is not accepted when its screens work on a normal day, but when its controls, fallbacks and recovery perform under adversarial and degraded conditions.



CHAPTER 13
DECISION TEST

A connected factory is not accepted when its screens work on a normal day, but when its controls, fallbacks and recovery perform under adversarial and degraded conditions.

A connected textile factory is ready for operational dependence only when its cybersecurity boundaries, protective-safety independence and recovery procedures pass witnessed acceptance tests under degraded and adversarial conditions. [S17] A commissioning demonstration that shows live charts on a calm afternoon proves only that packets cross a wire when power, links and software behave normally. [S32] Where electrical interruptions are frequent and wide-area connectivity cannot be assumed continuous, an untested architecture fails in ways that confuse operators, corrupt production history or stall the line. [S32] Green Future Tech Limited (GFT), Hong Kong SAR, therefore treats security, functional safety and contingency recovery not as compliance paperwork, but as engineering acceptance gates governing whether a pilot line may move from shadow observation into daily operational use. [S30]

This chapter specifies the evidence pack required before a sponsor, plant manager or safety owner accepts a connected-factory deployment under GFT's Cameroon textile programme. [R06] While Cameroon's inherited sector baseline records roughly 95 per-cent of cotton fibre exported raw and about 5 percent processed locally, expanding domestic processing towards the national 50 percent target for 2030 requires factor-ies whose instrumentation strengthens control rather than introducing new single points of failure. [R06] [R07] Each control maps to an observable test condition, an explicit pass-or-fail criterion set by the process or safety owner, a retained test record and a remediation path for exceptions. [S18] This chapter provides engineering reference guidance rather than investment, legal or tax advice, and it does not certify any machine, supplier or facility. [S24]

13.1

The risk assessment defines the system under consideration

Security and safety acceptance cannot begin until the factory defines the exact boundary of the system under test and accounts for every physical and logical path across it. [S19] In many brownfield retrofits, integration begins with an informal sketch showing an edge gateway reading a programmable logic controller (PLC), while omit-ting engineering laptops, cellular modems, shared switches, local buffers and safety relays. [S17] IEC 62443-3-2 makes security zones and conduits a key concept in security risk assessment for system design, requiring the team to partition the system under consideration into asset groupings with common security requirements and to document every channel connecting them. [S19]

In GFT's reference architecture, the risk assessment partitions the connected factory into explicit zones mirroring the responsibility separation established earlier in this paper. [S01] [S02] Production equipment and controllers sit in the production opera-tional technology (OT) zone; safety sensing, logic and actuation reside in an inde-pendent safety zone; protocol adapters, edge buffers and vision nodes sit in the edge acquisition zone; historians, brokers and manufacturing operations services occupy the operations zone; an industrial demilitarised zone (DMZ) mediates external exchange; and corporate or cloud services remain in the business information technology (IT) or external zone. [S01] [S19]

13.2 EVERY ASSET ENTERS THE REGISTER WITH AN OWNER AND AN INTERFACE BOUNDARY

A boundary definition is only as credible as its asset inventory. [S18] Before acceptance testing starts, the project team compiles a complete register of every asset inside or touching the system under consideration: spinning, weaving, dyeing, cutting and sewing controllers, IO-Link masters, Modbus converters, PROFINET switches, RFID readers, inspection cameras, edge appliances, local stores, MQTT brokers, engineering workstations, human-machine interfaces (HMIs), uninterruptible power supplies (UPS) and safety relays. [S09] [S10] [S11] [S19] The register records each asset's ISA-95 role, security zone, software version, clock source, role owner and failure consequence. [S01] [S12] [S13] [S18]

Ownership in the register is assigned to a specific factory role, such as the spinning maintenance engineer, quality systems lead or plant OT custodian, rather than an external integrator. [S18] [S24] Where an edge node runs containers under a multi-access edge computing pattern, both host appliance and deployed modules appear as distinct entries with resource and storage limits, letting the factory see which lines and KPIs are exposed when a change or vulnerability arises. [S08] [S16] [S18]

13.3 CONDUITS SPECIFY PERMITTED DIRECTION, PROTOCOL AND FAILURE BEHAVIOUR

Between any two security zones, communication occurs only through a documented conduit. [S19] For each conduit, the register specifies initiating asset, receiving asset, transport protocol, application semantics, connection direction, payload direction, authentication mechanism, expected message rate and required behaviour when congested or severed. [S05] [S19] [S40] For example, the conduit between an edge node and the MQTT broker permits the edge node to initiate an outbound transport-layer-security session and publish Sparkplug birth, data and death messages, while prohibiting inbound connections from business IT to the edge appliance or controller. [S19] [S40]

Acceptance testing verifies conduit rules empirically rather than accepting firewall screenshots. [S17] [S19] During commissioning, the test engineer attempts authorised transactions and representative unauthorised connections across each boundary: probing closed ports, attempting direct write commands from the operations zone to machine registers, and simulating an unauthorised workstation joining the production switch. [S09] [S17] [S19] A conduit passes acceptance only when permitted production evidence flows cleanly, prohibited attempts are blocked, and rejections generate timestamped logs. [S17] [S19]

13.4

Cybersecurity outcomes include governance and ownership

Technical network controls degrade rapidly if no management system governs who may change them, how exceptions are reviewed and how incidents are handled after commissioning. [S18] [S24] Because a mid-sized textile mill or garment facility in Cameroon cannot sustain a sprawling compliance bureaucracy, GFT combines the

National Institute of Standards and Technology (NIST) Cybersecurity Framework (CSF) 2.0 for outcome-oriented goals with ISO/IEC 27001:2022 for a lean, continually improved information security management system (ISMS). [S18] [S24] [S30]

NIST CSF 2.0 organises cybersecurity risk management around six core functions: Govern, Identify, Protect, Detect, Respond and Recover. [S18] NIST Special Publication 800-61 Revision 3 extends this structure by showing how organisations incorporate incident response recommendations and considerations across those same CSF 2.0 functions. [S1301] GFT translates these six functions into concrete acceptance evidence that a small factory team can verify and maintain during normal shifts. [S18] [S1301]

13.5 TARGET OUTCOMES MATCH FACTORY ROLES RATHER THAN ENTERPRISE HEADCOUNT

In a lean manufacturing team, cybersecurity governance works when responsibilities attach to existing operational cadences. [S18] [S30] Under Govern and Identify, the sponsor and plant manager approve the risk appetite, zone-and-conduit architecture, asset register and role assignments for data owner, process owner, OT custodian and safety officer before Phase 1 pilot operation begins. [S18] [S19] Under Protect and Detect, the OT custodian and line engineers verify access control, network segmentation, configuration baselines, clock synchronisation and stale-session alerting. [S13] [S18] [S40] Under Respond and Recover, operations and maintenance leads rehearse incident containment, manual fallback logging and isolated backup restoration. [S22] [S1301]

Each outcome is written as an observable state checked during a site acceptance walkdown. [S18] Rather than assuming a dedicated security operations centre watches alerts around the clock, the design reduces attack surface by eliminating unnecessary remote paths, restricting write capabilities at the gateway interface, and routing high-priority integrity alerts into supervised operational channels. [S17] [S19]

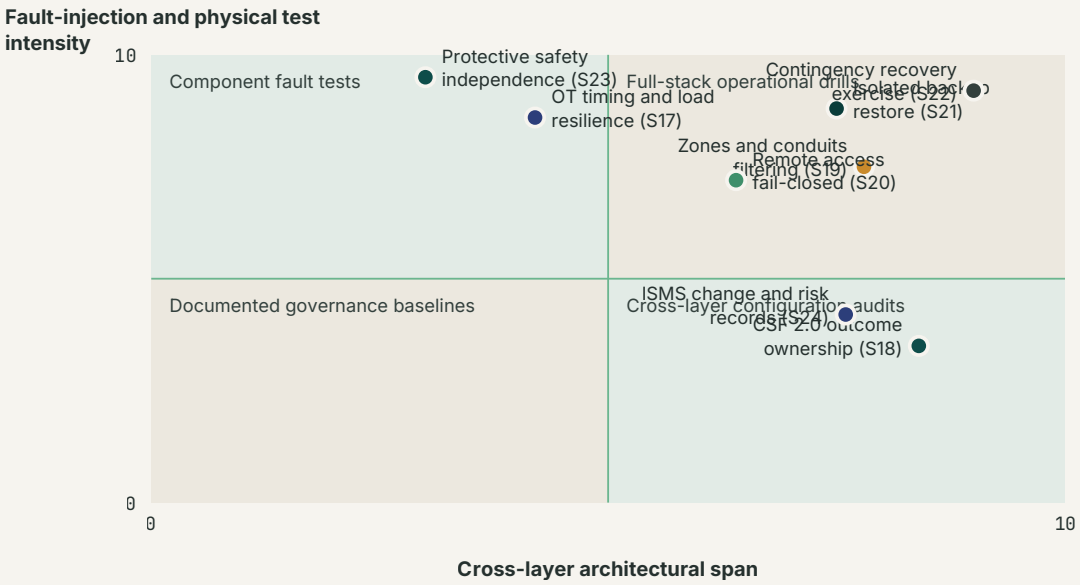
13.6 MANAGEMENT-SYSTEM RECORDS BIND TECHNICAL CONTROLS TO CONTINUAL IMPROVEMENT

ISO/IEC 27001:2022 provides requirements for establishing, implementing, maintaining and continually improving an information security management system for organisations of any size and sector. [S24] Acceptance requires four working management loops before handover: a risk-treatment register linking every identified OT and edge threat to an implemented control or signed residual-risk acceptance; a change-control workflow requiring authorisation, version recording and rollback readiness before any firmware, PLC logic, vision model or KPI formula is modified; a supplier-access procedure governing integrators; and a corrective-action log tracking test failures to closure. [S08] [S18] [S24]

To make the relationship between standards, technical controls and verification concrete, **Exhibit 58**, p. 208 maps the eight acceptance control families across their architectural span and required fault-injection intensity. [S17] [S19]

EXHIBIT 58 / DERIVED SCOPE AND VERIFICATION INDEX (0 TO 10)

Acceptance maps every control to objective evidence



DERIVED [S17] [S18] [S19] [S20] [S21] [S22] [S23] [S24] Derived qualitative placement of eight acceptance domains from the scope of NIST SP 800-82r3 (S17), NIST CSF 2.0 (S18), IEC 62443-3-2 (S19), CISA CPG 2.0 (S20), CISA backup guidance (S21), NIST SP 800-34r1 (S22), IEC functional safety (S23) and ISO/IEC 27001 (S24). X-axis scores cross-layer architectural span; Y-axis scores required physical and fault-injection test intensity.

As **Exhibit 58**, p. 208 illustrates, governance baselines under NIST CSF 2.0 and ISO/IEC 27001 span the stack from shop-floor devices to executive reviews, whereas protective safety independence and OT timing resilience require high-intensity physical fault injection at the machine level. [S17] [S18] [S23] [S24] Isolated backup restoration and contingency exercises sit in the upper-right quadrant because they combine cross-layer coordination with live restoration and reconciliation drills, though a qualitative matrix cannot substitute for executing individual tests on plant hardware. [S17] [S21] [S22]

13.7

OT controls must pass production and timing tests

In an office network, packet inspection delay is a minor inconvenience; on an industrial line, a security control that starves a controller CPU or disrupts cyclic communication is an operational hazard. [S17] NIST Special Publication 800-82 Revision 3 emphasises that securing operational technology requires addressing the unique performance, reliability and safety requirements of industrial environments. [S17] GFT therefore requires every OT security control, gateway polling schedule and network filter to pass operational timing and production-impact tests before acceptance. [S17]

This requirement applies across the fieldbus and industrial Ethernet protocols in a mixed-generation textile plant. [S35] Where brownfield machinery uses MODBUS, communication follows a request-and-reply model specified by function codes;

aggressive polling can delay controller scan cycles or trigger timeouts. [S09] [S17] Where smart sensors connect through IO-Link, cyclic process data exchange is separated from on-request parameters, commands and events so process data flows without interference. [S10] Where high-speed machinery uses deterministic PROFINET, closed-loop control can depend on isochronous transmission with jitter below one microsecond. [S11] Any security appliance or gateway placed near these networks must prove that it respects their deterministic boundaries. [S11] [S17]

13.8 SECURITY APPLIANCES AND POLLING GATEWAYS ARE TESTED UNDER PEAK LINE LOAD

Operational acceptance of OT security controls proceeds in two stages: non-interference verification under peak production rates, and degraded-mode verification under abnormal traffic or appliance failure. [S17] During the first stage, the machine runs at maximum rated speed while the edge gateway executes its full polling schedule, local inference models process camera frames, and the zone firewall enforces its rules; engineers monitor scan times, fieldbus error counters, IO-Link event logs, PROFINET jitter and gateway resource use. [S09] [S10] [S11] [S17] [S26]

During the second stage, the test team introduces stress and failure conditions across the boundary conduit: disconnecting the operations network, saturating the firewall with blocked traffic on the non-OT side, forcing a broker failover, and power-cycling the edge gateway. [S17] [S19] Throughout these disturbances, the machine controller must continue executing its control loop without fault, cyclic process data within the machine zone must remain undisturbed, and no false command or register corruption may occur. [S10] [S17]

13.9 PROCESS OWNERS SET TIMING AND RELIABILITY THRESHOLDS BEFORE TESTING BEGINS

Neither the cybersecurity standard nor the reference architecture invents universal numerical thresholds for scan-time margin, polling interval or recovery jitter, because tolerances depend on the physical process: a bale scale, dyeing vat and air-jet loom have different time constants. [S11] [S17] Instead, the process owner and equipment engineer record the operating envelope on the machine-capture sheet before testing starts, including maximum scan-time increase, zero uncommanded state changes, and required clock precision under NTPv4 or IEEE 1588 Precision Time Protocol (PTP). [S12] [S13] [S17]

METHOD

To verify that OT security and acquisition controls do not degrade machine performance or event integrity, the commissioning team executes a four-step protocol on each machine class. [S17]

1. **Establish the uninstrumented baseline.** With the edge gateway disconnected, record baseline scan time, error counters, HMI response and cycle timing over an approved run. [S09] [S17]
2. **Measure full-rate acquisition and filtering.** Connect the conduit with read-only polling, firewall rules, TLS and NTP or PTP active, then verify scan times, jitter and cyclic IO-Link exchange stay within approved tolerances. [S10] [S11] [S12] [S13] [S17]
3. **Inject boundary faults and traffic bursts.** Simulate firewall rejections, broker loss, certificate expiry and gateway reboot during operation, confirming the controller neither faults nor pauses and wireless or RFID reads flag any sync loss. [S17] [S31] [S40]
4. **Verify deterministic telemetry recovery.** Restore the conduit, confirm the edge buffer flushes observations with preserved source timestamps, and verify no duplicate or out-of-order events corrupt KPIs. [S08] [S29] [S40]

Applying this method prevents deploying intrusive IT scanners in fragile control networks or waiving segmentation out of untested fear that a firewall might slow production. [S17] [S19]

13.10

Remote maintenance fails closed and leaves evidence

Equipment suppliers, vision integrators and automation engineers often request remote access to diagnose faults or update software. [S17] [S20] In Cameroon, where specialist vendor support for imported spinning, weaving or inspection machinery may reside abroad, structured remote maintenance has a clear rationale. [R06] [S20] Yet unmanaged remote access, such as an always-on cellular router wired into a machine switch or shared passwords, is a primary pathway for compromise and configuration drift. [S19] [S20]

The Cybersecurity and Infrastructure Security Agency (CISA) Cybersecurity Performance Goals (CPG) 2.0 require multi-factor authentication (MFA) on all OT accounts and systems that can be accessed remotely, including vendor and maintenance accounts, remotely accessible user and engineering workstations, and remotely accessible HMIs when available. [S20] IEC 62443-3-2 complements this by requiring that remote traffic cross only defined conduits terminating in an appropriate zone rather than bridging directly into production control. [S19]

13.11 VENDOR SESSIONS REQUIRE EXPLICIT ACTIVATION, MULTI-FACTOR AUTHENTICATION AND EXPIRY

In GFT's reference architecture, remote maintenance fails closed by default: when no window is active, the remote conduit is disabled and no external party can reach the industrial DMZ or operations zone. [S19] [S20] Opening a session requires a three-step sequence tested during acceptance: a named vendor engineer requests access for a specific asset, purpose and window against an ISO/IEC 27001 change-control ticket; the process owner and OT custodian approve the window and enable the time-bounded rule on the DMZ jump host; and the vendor authenticates with an individual account and MFA. [S19] [S20] [S24]

Once authenticated at the DMZ, the session is restricted to the target edge node or workstation using least-privilege rules blocking lateral movement. [S19] [S20] Every command, file transfer and change is logged with user identity and timestamp, and when the timer expires or the supervisor trips a termination switch, the session closes and credentials revoke. [S19] [S20] [S24]

13.12 LEGACY ASSETS WITHOUT MODERN AUTHENTICATION REQUIRE COMPENSATING CONDUIT ISOLATION

Brownfield textile plants include older PLCs, panels and field devices that cannot natively support MFA or encryption. [S17] [S20] CISA CPG 2.0 accounts for this by specifying MFA on remotely accessible systems and HMIs when available, while IEC 62443-3-2 isolates legacy devices behind an intermediary boundary. [S19] [S20] Where a controller or HMI lacks strong authentication, GFT prohibits exposing it directly; remote access terminates at a hardened DMZ jump host enforcing MFA, session recording and role-based authorisation before communicating over a filtered conduit. [S19] [S20]

During acceptance, the team tests positive and negative paths: login without MFA, login with an expired window, lateral probes to a neighbouring sewing line from a cutting-room session, and pulling the WAN link mid-session to confirm clean termination. [S17] [S19] [S20]

> **GFT NOTE**
>
> Never accept a machine delivery or integration package with a vendor cellular modem, Wi-Fi access point or cloud agent wired directly into the control cabinet behind the plant boundary. [S19] During pre-commissioning inspection, trace every cabinet cable, remove or de-energise any bypass modem, and require all vendor diagnostics to traverse the industrial DMZ jump host under MFA, supervisor approval and automatic session expiry. [S19] [S20]

13.13

Protective safety never depends on the dashboard path

The single most important architectural invariant in a connected factory is that protective safety functions remain independent of data-collection gateways, edge inference models, message brokers, historians, dashboards and cloud services. [S19] [S23] As the International Electrotechnical Commission (IEC) explains, within the overall

safety of a system, functional safety provides automatic protection in response to errors. [S23] IEC 61508-1:2010 sets out general requirements when electrical, electronic and programmable electronic (E/E/PE) systems carry out safety functions, establishing a safety lifecycle distinct from general-purpose information systems. [S1302]

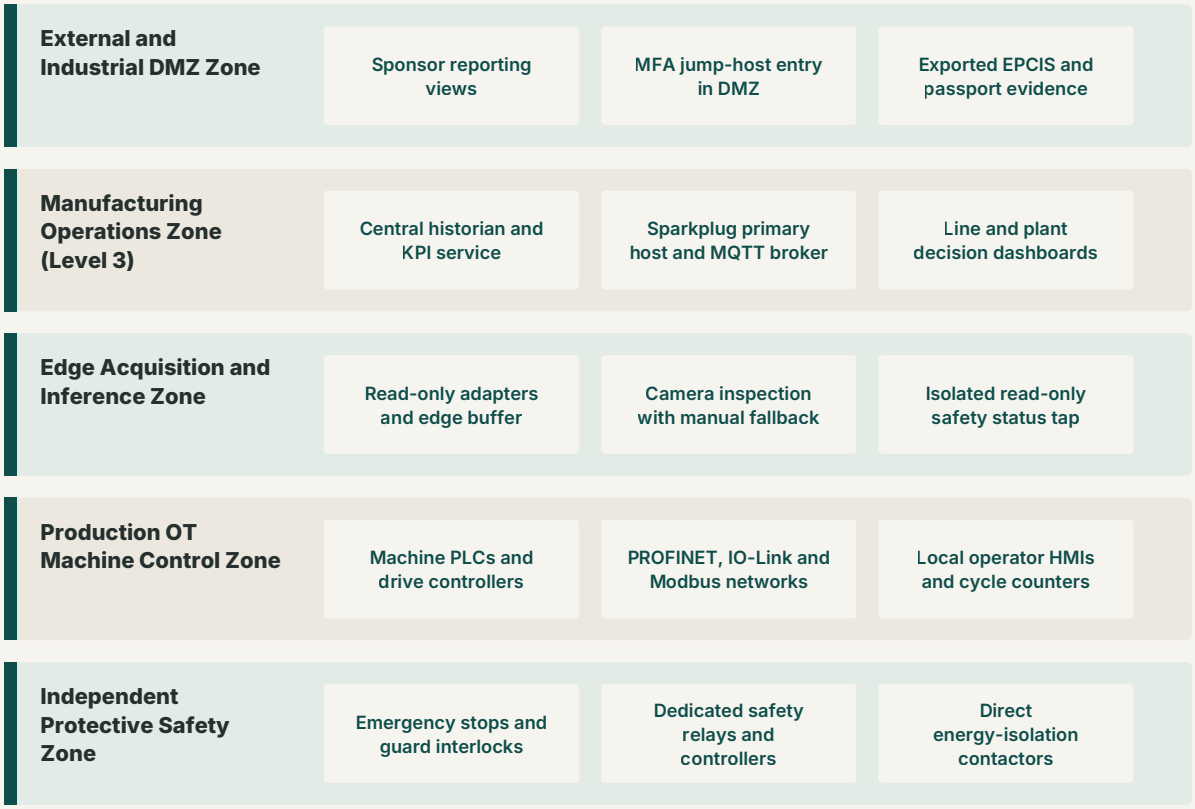
In a textile and garment factory, physical hazards include rotating shafts on spinning frames and looms, cutting blades, heated rollers, pressurised dyeing vessels and conveyors. [S23] Emergency-stop circuits, light curtains, two-hand controls, interlock switches, over-temperature cut-outs and safety relays or safety PLCs bring equipment to a safe state automatically when a hazard or fault occurs, and none of those actions may wait for an MQTT message, an edge container or a dashboard alert. [S05] [S16] [S23] [S1302]

13.14 FAULT INJECTION PROVES THAT TELEMETRY FAILURE NEVER INHIBITS PROTECTIVE ACTION

While edge nodes and dashboards observe safety events, the signal path from the safety system to the edge gateway must be unidirectional or galvanically and logically isolated so no monitoring fault propagates into the safety loop. [S17] [S19] [S23] **Exhibit 59**, p. 213 illustrates how GFT separates independent protective safety from production control, edge acquisition, operations management and external access. [S19] [S23] [S1302]

EXHIBIT 59 /

Protective safety remains outside the ordinary data path



ILLUSTRATIVE [S19] [S23] [S1302] Illustrative zone and responsibility separation informed by IEC 62443-3-2 zones and conduits (S19), IEC functional safety principles (S23) and IEC 61508-1 E/E/PE safety-related system scope (S1302). Protective actuation never traverses the edge, operations or external layers.

As shown in **Exhibit 59**, p. 213, the protective safety zone acts directly on machine actuators and contactors while providing only an isolated, read-only diagnostic tap to the edge layer. [S19] [S23] [S1302] During safety acceptance, the safety officer and automation engineer execute witnessed fault-injection tests in a safe test mode: unplugging edge power, disconnecting or shorting the diagnostic tap, flooding the edge network interface, freezing the dashboard and severing the plant LAN. [S17] [S23] Under every condition, every emergency stop, guard interlock and protective trip must actuate within its certified response time and bring the machine to its safe state; any exception blocks commissioning. [S23] [S1302]

13.15 PHYSICAL AND LOGICAL SEPARATION PROTECTS SAFETY LOGIC FROM UNAUTHORISED MODIFICATION

IEC 62443-3-2 and IEC 61508-1 also require that safety logic and parameters be protected against unintended or malicious modification over network interfaces. [S19] [S1302] Where a machine uses a programmable safety controller alongside a standard PLC, the safety program remains locked behind a physical key-switch or hardware write-protect mechanism, and no conduit from edge, operations or remote zones may alter safety parameters or bypass interlocks. [S17] [S19] [S1302] Acceptance testing verifies that parameter writes over the monitoring interface are rejected and that key-switch states are logged. [S19] [S24]

> **WARNING**
>
> Never use a computer-vision model, wireless sensor reading, MQTT message or dashboard alert as a protective safety interlock for personnel or high-energy equipment. [S23] [S26] [S31] Vision detectors are sensitive to lighting and background variation, and wireless or brokered networks experience latency, contention and stale sessions. [S26] [S31] [S40] All personnel and equipment protection must remain governed by dedicated hardwired or certified E/E/PE functional-safety systems under IEC 61508-1, independent of the edge-to-dashboard stack. [S23] [S1302]

13.16

Backups count only when isolated restoration succeeds

Ransomware, media failure, power-surge corruption and human configuration errors share a common remedy: restoring a clean, known-good copy of software, logic, configurations and historical data onto replacement or wiped hardware. [S21] [S22] Yet many plants discover during an outage that their backup share sat read-write on the same network segment that malware encrypted, the PLC project file was outdated, or certificates and KPI definitions were omitted. [S08] [S21] [S22]

Joint guidance from CISA and partner authorities advises operators to regularly test backup procedures and ensure backups are isolated from network connections that could enable malware spread. [S21] NIST Special Publication 800-34 Revision 1 reinforces that testing validates recovery capabilities against their intended design, so GFT treats backup verification as a live, clean-target restoration test rather than a schedule inspection. [S21] [S22]

13.17 OFFLINE AND IMMUTABLE COPIES PREVENT SHARED-CREDENTIAL COMPROMISE ACROSS BACKUPS

To satisfy CISA's isolation requirement, the backup architecture ensures that compromised credentials in the business IT or operations zone cannot delete, overwrite or encrypt recovery copies. [S19] [S21] In a small or mid-sized textile plant, this uses an access-controlled backup appliance pulling backups over a restricted conduit plus periodic offline, write-protected media stored on-site and off-site, with every artefact hashed and logged in the ISO/IEC 27001 configuration register. [S19] [S21] [S24]

The backup inventory covers four layers required to reconstruct the evidence chain: 1. Authorised PLC, drive and IO-Link master project files and machine recipes, captured on every approved change. [S10] [S21] [S24] 2. Edge operating-system images, container specifications, Modbus and OPC UA register maps, Sparkplug configurations, and approved vision weights and calibration files. [S04] [S09] [S16] [S26] [S40] 3. Switch configurations, firewall rules, NTP/PTP settings, TLS certificate authorities and MQTT broker access lists. [S12] [S13] [S19] [S21] 4. Central historian schemas, raw observation archives, event tables, ISO 22400 KPI dictionary definitions, dashboard layouts, and ISO/IEC 27001 audit and change logs. [S08] [S21] [S24]

13.18 BARE-METAL AND CLEAN-TARGET RESTORATION VERIFIES COMPLETENESS OF THE CONFIGURATION SET

During acceptance, the commissioning team wipes a representative edge node and a staging historian and KPI service to bare metal, disconnects internet or vendor cloud links, and has plant maintenance and OT staff restore both solely from isolated backup media and the written runbook. [S21] [S22] The test passes only when the restored edge node validates its certificates, resumes polling and local buffering, and publishes valid Sparkplug messages, while the restored historian and KPI service reproduce identical historical KPI values without manual patching. [S08] [S21] [S22] [S40]

WORKED EXAMPLE

Consider an illustrative restoration and reconciliation test for a garment-assembly line showing how recovery point objective (RPO), recovery time objective (RTO) and KPI reconciliation are calculated. All figures in this worked example are explicitly illustrative and demonstrate the verification arithmetic rather than an observed GFT factory result.

Suppose the process owner sets an illustrative target RTO of 120 minutes for the central historian and dashboard, an illustrative central-database RPO of 15 minutes prior to edge replay, and an illustrative reconciled post-replay RPO of 0 lost production events for outages under 24 hours, backed by edge persistence. [S22] [S29] Assume an illustrative steady rate of 180 completed bundles and 12 energy-metre readings per hour, or 192 critical events per hour. [S08] [S29]

At hour 04:00 of an illustrative 8-hour shift, the test coordinator simulates historian storage failure while keeping the line and edge node running for 3 hours (until 07:00) before restoring an isolated snapshot taken at 03:45. [S21] [S22] Across the 3-hour outage plus the 15-minute gap since 03:45 (3.25 hours total), the line generates an illustrative 3.25 multiplied by 180 equals 585 bundle events and 3.25 multiplied by 12 equals 39 energy readings, or 624 critical events stored in the local edge buffer. [S29] [S40]

The plant team restores the isolated 03:45 snapshot onto a clean server in an illustrative 74 minutes (by 08:14) and starts the Sparkplug primary host. [S21] [S22] [S40] The edge node flushes its 624 buffered records with original source timestamps in an illustrative 11 minutes (by 08:25), bringing total recovery and reconciliation time to 85 minutes against the 120-minute illustrative RTO target. [S29] [S40]

Before sign-off, the quality and data owners execute three illustrative checks against the ISO 22400 KPI dictionary. [S08] Event coverage divides reconciled historian events (624) by expected line-counter and metre events (624), yielding 100.0 percent coverage and a reconciled RPO of 0 lost events. [S08] [S29] Timestamp integrity confirms all 624 replayed events retain source timestamps between 03:45 and 07:00 rather than the 08:14 to 08:25 receipt window. [S13] [S40] KPI reproducibility compares hourly output and energy-per-bundle KPIs for hours 04:00 to 07:00 against an independent raw-log calculation, confirming an illustrative difference of 0.0 units. [S08] [S22]

13.19

Exercises turn contingency documents into operating capability

Unrehearsed contingency procedures fail under real disruption. [S22] NIST Special Publication 800-34 Revision 1 distinguishes three activities: testing validates recovery capabilities, training prepares recovery personnel for plan activation, and exercising the plan identifies planning gaps; combined, they improve plan effectiveness and organisation preparedness. [S22] In GFT's acceptance framework, a connected factory is accepted only after recovery capabilities are tested, floor and technical staff are trained on degraded-mode duties, and end-to-end contingency scenarios are exercised with documented gap closure. [S22] [S1301]

This discipline matches documented operating conditions in Cameroon and the region. [S32] [S33] Because 93.3 percent of surveyed Cameroonian firms in the 2024 World Bank Enterprise Survey reported electrical outages, with 10.4 outages in a typical month across formal private firms, power transitions and partial restarts are routine. [S32] Similarly, while ITU's 2024 figure that 14 percent of Africa's population lacks mobile-broadband coverage is a regional population indicator rather than a factory link measurement, it underscores why wide-area link loss and delayed remote support must be rehearsed as standard conditions. [S33]

13.20 FIVE FAILURE SCENARIOS TEST THE FACTORY AGAINST REALISTIC OPERATIONAL DISRUPTIONS

GFT requires the acceptance programme to execute five disruption scenarios before scale-up is approved. [S21] [S22]

1. **Ransomware containment and operations-zone isolation.** Severing conduits between business IT, the industrial DMZ, the operations zone and the edge zone verifies that machines and safety systems run safely, supervisors use approved manual shift sheets, and technicians rebuild the historian from isolated backups without reconnecting untrusted networks. [S19] [S21] [S22] [S23] [S1301]
2. **Extended WAN and site-power outage with local buffering.** Simulating multi-hour WAN loss combined with grid-to-generator switching verifies that edge nodes record power-source state, buffer observations on UPS power, shut down cleanly if batteries deplete, and flush ordered history upon restoration. [S29] [S32] [S33] [S40]
3. **Edge-hardware failure and cold-spare replacement on an inspection cell.** Simulating failure of an inspection edge node verifies that quality operators immediately switch to manual inspection fallback, while a technician installs a cold spare, restores the versioned container and calibration from isolated backup, and verifies acquisition quality before resuming automated assistance. [S21] [S22] [S26]
4. **Corrupted historian or bad KPI release rollback.** Simulating a faulty schema migration or corrupted partition verifies that stale-data and quality flags alert owners instead of displaying wrong figures, the team rolls back under ISO/IEC 27001 change control, and missing intervals are reconstructed from raw observations and edge buffers. [S08] [S24] [S40]

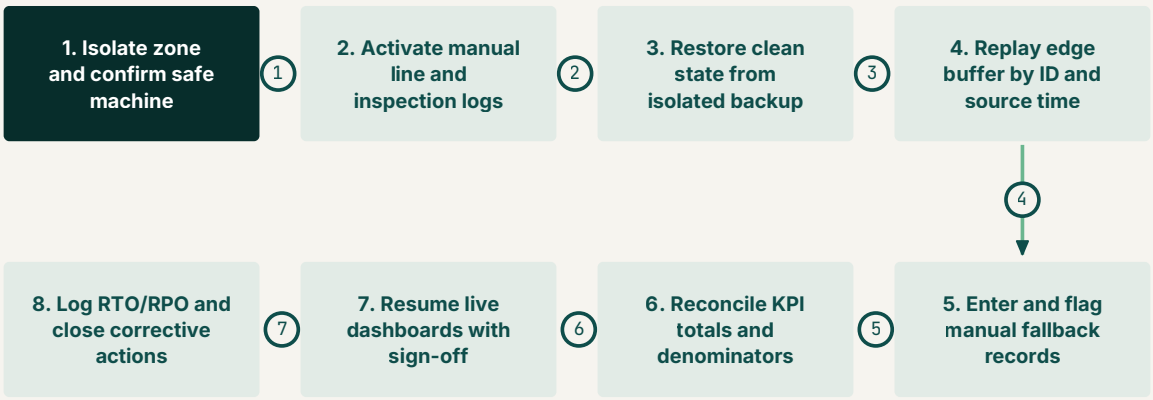
5. **Unavailable external vendor support during a machine interface fault.** Simulating a fieldbus or adapter fault while external suppliers are unreachable verifies that local technicians have the training, register maps and disconnect procedures to isolate the gateway and keep the machine running safely under local control. [S09] [S17] [S22] [S33]

13.21 POST-EXERCISE RECONCILIATION CLOSES EVIDENCE GAPS BEFORE NORMAL OPERATION RESUMES

Recovering a server or reconnecting a cable is only the midpoint of a contingency exercise. [S22] When a digital system returns from degraded or manual operation, some events sit in edge buffers, some quality records sit on manual sheets, and some metre intervals occurred during generator transitions. [S08] [S22] [S29] [S40] Resuming dashboards without reconciling those records leaves gaps, double-counted bundles or distorted resource-intensity denominators. [S08] [S29] **Exhibit 60**, p. 217 shows the eight-stage recovery and reconciliation sequence required for every contingency exercise and incident response. [S21] [S22] [S1301]

EXHIBIT 60 /

Recovery ends with evidence reconciliation and corrective action



1. Safety confirmed
2. Clean target ready
3. Host online
4. Telemetry deduplicated
5. Provenance tagged
6. Zero KPI drift
7. Audit logged

ILLUSTRATIVE [S21] [S22] [S1301] Illustrative recovery and reconciliation workflow informed by CISA isolated backup restoration guidance (S21), NIST SP 800-34r1 testing, training and exercise principles (S22), and NIST SP 800-61r3 incident response lifecycle alignment (S1301).

As **Exhibit 60**, p. 217 demonstrates, live dashboards resume at step 7 only after edge buffers replay by source timestamp (step 4), manual fallback records enter with explicit provenance and quality flags (step 5), and ISO 22400 KPIs reconcile against physical line counters and metre totals (step 6). [S08] [S22] [S40] Step 8 records

elapsed restoration time against target RTO, data loss against target RPO, and corrective-action owners and due dates in the ISO/IEC 27001 register, while leaving site-specific RTO and RPO targets to be set by process owners before testing begins. [S22] [S24]

To consolidate all seven acceptance domains into an inspectable commissioning contract, **Exhibit 61**, p. 219 presents the GFT acceptance evidence index, where no test is marked passed in the outline and every row defines scenario, expected behaviour, measurement record, approver and exception rule. [S17] [S19] [S22] [S23]

EXHIBIT 61 /

The acceptance pack proves behaviour under failure

Acceptance scenario	Expected system behaviour	Measured acceptance evidence	Approver and exception rule	Outline status
Zone and conduit boundary scan and unauthorised write probe (S19)	Authorised telemetry passes; unapproved ports and cross-zone PLC writes fail closed and log an alert	Packet capture, firewall deny log with synchronised timestamp, and PLC register integrity check	OT custodian and plant manager; open write path blocks pilot connection	Pending site test
Peak-load polling and boundary fault injection on OT control (S17)	PLC scan time and fieldbus jitter stay within process-owner bounds during peak load and gateway reboot	Controller diagnostic log before, during and after fault injection against pre-approved baseline	Process owner and automation lead; control timing breach requires polling redesign	Pending site test
Remote vendor maintenance session and expiry test (S20)	Access fails closed by default; opens only with ticket, MFA and timer; blocks lateral zone access	Bastion authentication log, session recording, lateral-probe rejection log and auto-revocation timestamp	OT custodian; any legacy asset lacking MFA must sit behind DMZ jump host	Pending site test
Protective safety independence under edge and network failure (S23)	Emergency stops, guards and safety relays trip deterministically when edge, broker or LAN fails	Witnessed physical trip-time and safe-state record under disconnected and faulted edge tap	Safety officer; zero exceptions permitted before machine energisation or pilot run	Pending site test
Isolated bare-metal backup restoration and replay test (S21, S22)	Clean server and spare edge node restore from offline media; edge flush reconciles KPIs with zero drift	Stopwatch RTO, measured event RPO, cryptographic hash check and KPI reconciliation sheet	Data owner and OT custodian; failed restore blocks transition from shadow mode	Pending site test
End-to-end contingency exercise across five disruption scenarios (S22)	Operators execute manual fallback; technicians restore service and reconcile manual and buffered logs	Exercise observation report, reconciled KPI audit, training roster and ISO/IEC 27001 corrective-action log	Plant manager and process owners; gaps assigned with retest date before Phase 2 scale	Pending site test

ILLUSTRATIVE [S17] [S19] [S20] [S21] [S22] [S23] [S24] Illustrative commissioning acceptance register derived from NIST SP 800-82r3 (S17), IEC 62443-3-2 (S19), CISA CPG 2.0 (S20), CISA backup guidance (S21), NIST SP 800-34r1 (S22), IEC functional safety (S23) and ISO/IEC 27001 (S24). Status fields remain Pending site test until witnessed on actual factory equipment.

As **Exhibit 61**, p. 219 makes explicit, the acceptance pack does not treat all exceptions equally. [S19] [S23] An open exception in protective safety independence is a non-negotiable stop condition: no line may operate with a compromised emergency stop or guard interlock. [S23] [S1302] An open exception in zone-boundary write protection, backup restoration or KPI reconciliation blocks transition from shadow mode to operational dependence, keeping the official record on the baseline procedure until retest passes. [S19] [S21] [S22] Minor administrative findings that do not threaten safety, machine control or evidence integrity are logged with an owner, compensating control and retest date in the ISO/IEC 27001 corrective-action register, feeding into the phase-gate decisions of the implementation roadmap. [S24] [S30]

A factory is ready to trust its screen only after it has proven how it protects its people, controls its machines and reconstructs its truth when the screen goes dark.

CHAPTER SOURCE NOTES

S17 National Institute of Standards and Technology. 2023-09.

S32 World Bank Enterprise Surveys. 2025.

S30 MDPI, Sustainability. 2023-08-19.

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

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

S18 National Institute of Standards and Technology. 2024-02-26.

S24 International Organization for Standardization and International Electrotechnical Commission. 2022-10.

S19 International Electrotechnical Commission. 2020-06.

S01 International Society of Automation. 2025-04-10.

S02 International Electrotechnical Commission. 2013-05-22.

S09 Modbus Organization. 2012-04-26.

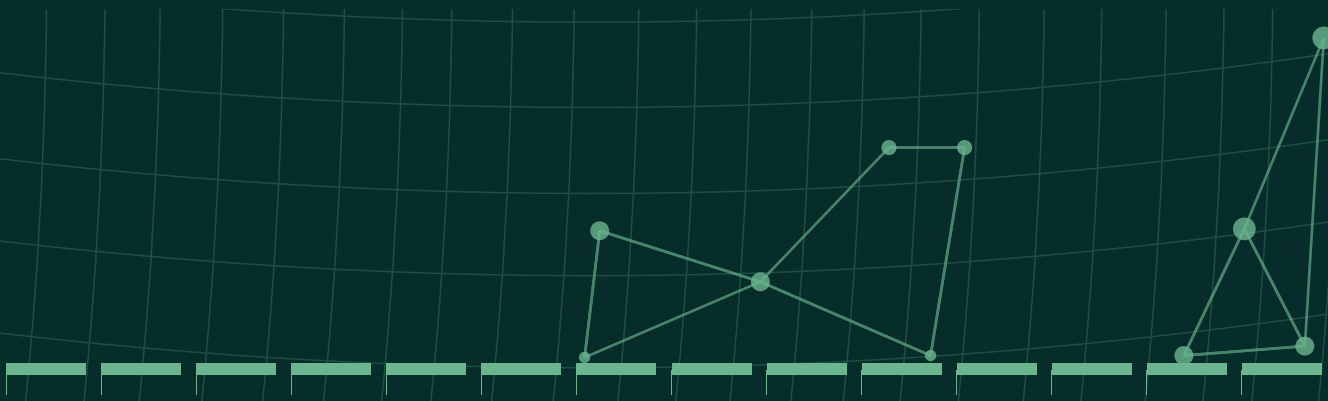
S10 IO-Link Community. 2024-06.

S11 PROFIBUS Nutzerorganisation. 2018-11.
S12 IEEE Standards Association. 2020-06-16.
S13 Internet Engineering Task Force and RFC Editor. 2010-06.
S08 International Organization for Standardization. 2014-01.
S16 European Telecommunications Standards Institute. 2024-04.
S05 OASIS Open. 2019-03-07.
S40 Eclipse Foundation. 2022-11-16.
S1301 National Institute of Standards and Technology. 2025-04.
S22 National Institute of Standards and Technology. 2010-11-11.
S20 Cybersecurity and Infrastructure Security Agency. 2025-12.
S21 CISA and joint international cybersecurity authorities. 2022-05-09.
S23 International Electrotechnical Commission. 2022-04-28.
S35 United Nations Industrial Development Organization. 2019.
S26 MDPI, Electronics. 2025-09-18.
S31 MDPI, Sensors. 2020-04-28.
S29 MDPI, IoT. 2024-09-28.
S1302 International Electrotechnical Commission. 2010-04-30.
S04 OPC Foundation. 2025-10-22.
S33 International Telecommunication Union. 2024-11-10.

14

Scale only after the architecture proves its decisions

GFT recommends a four-phase build that earns each expansion through evidence quality, operating value, safe failure and recovery rather than device count.



CHAPTER 14
DECISION TEST

GFT recommends a four-phase build that earns each expansion through evidence quality, operating value, safe failure and recovery rather than device count.

Scale the connected factory only when its evidence supports a useful operating decision and its failure leaves production under an accepted fallback. GFT recommends establishing contracts, proving a bounded line, replicating those contracts and then optimising the resulting system. Each transition requires a recorded decision by the factory's accountable owners. This is a proposed implementation method, not a claim that GFT has deployed or proved the architecture. The acceptance object is the complete path from an identifiable production event to a verified management action, including what happens when that path breaks.

14.1

The target is one operating architecture, not a collection of pilots

The sponsor should first approve the division of responsibilities between machine control, manufacturing operations and enterprise reporting. ISA's description of its integration standard and IEC's description of the manufacturing operations interface provide a vocabulary for that separation. They do not select a supplier or demonstrate a site's conformity. GFT's recommendation is to preserve responsibility boundaries while allowing individual products to change. A line supervisor must know which system records production, which record authorises a disposition and who resolves disagreement between them. [S01] [S02]

Create a small architecture authority with explicit decision rights rather than a new committee for every application. The production owner approves operational use and fallback. The data owner approves definitions and corrections. The engineering and safety owners approve physical interfaces and protective boundaries. The information-technology and security owners approve service operation and access. The sponsor resolves resource and scope decisions. These are recommended responsibilities, not assertions about existing GFT personnel. The same person may hold compatible roles, but authorship of a change should remain visible to whoever accepts its consequences.

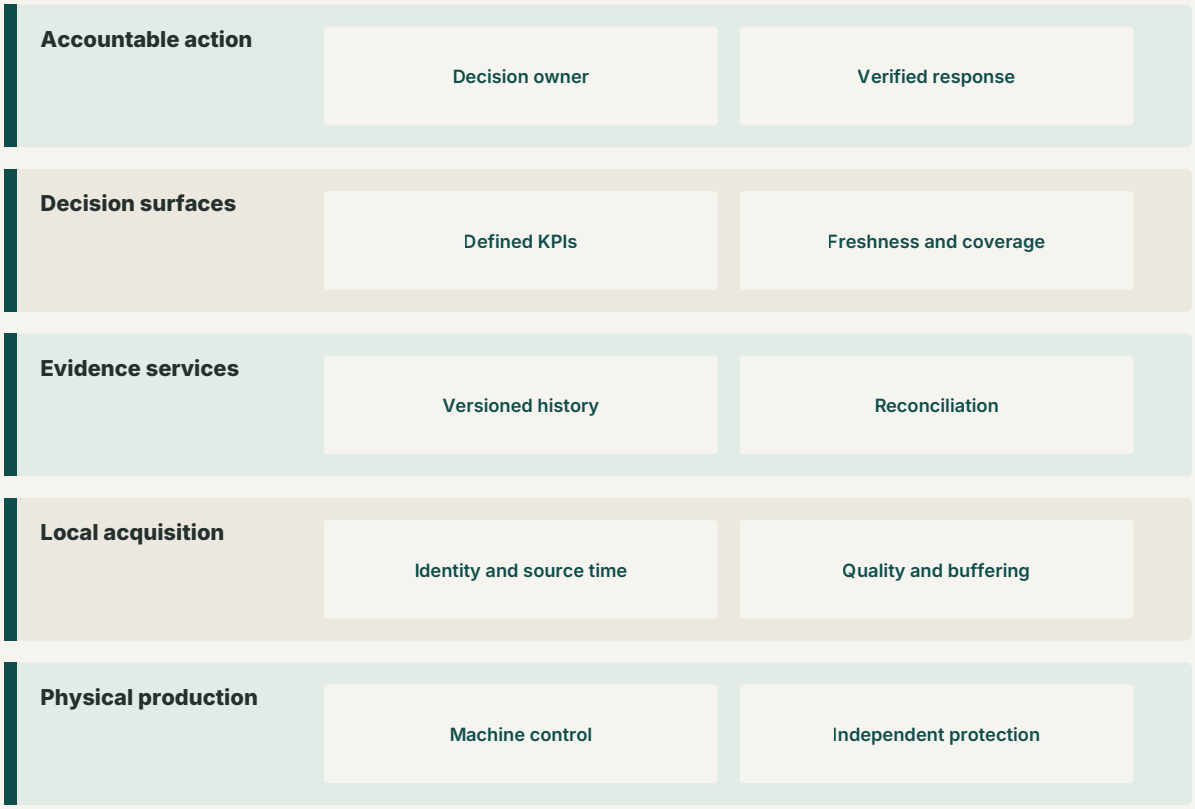
The authority should maintain a shared event contract, KPI dictionary, interface register, zone design and recovery policy. Applications should refer to these records through controlled versions. A defect dashboard and an energy dashboard may use different views, but neither should silently invent the identity of a line or redefine its production period. When a definition changes, record the effective boundary and identify the earlier results that remain comparable. Repeated local exceptions are a reason to improve the shared contract, not to create an undocumented alternative.

Use local processing where the decision must survive external service loss. Use remote services where their analytical or administrative benefit justifies the dependency. This allocation is GFT's design recommendation, subject to site measurement. NIST's OT guidance explicitly treats performance, reliability and safety alongside security; a reporting convenience therefore cannot be accepted solely because it works under a healthy network connection. The factory should approve what remains operable when remote authentication, synchronisation or analytics is unavailable. [S17]

The proposed end state in **Exhibit 62**, p. 224 consolidates responsibilities without specifying a vendor stack. Protective functions sit outside the reporting dependency path. The diagram describes an acceptance intention, not an installed system or a safety assessment.

EXHIBIT 62 / PROPOSED RESPONSIBILITY LAYERS

Operating responsibility survives application replacement



ILLUSTRATIVE [S02] [S08] [S17] [S40] Illustrative GFT synthesis. Layers describe proposed responsibilities, not installed equipment. Protective safety is independent of every reporting layer; it is not an analytics component.

The diagram's practical test is replaceability. Ask whether a new display could consume the existing dictionary, whether a replacement gateway could preserve event identity and whether a restored store could reproduce an accepted result. Failure of these tests identifies an ownership or interface gap before procurement turns it into an expensive dependency. This conclusion does not require every application to share a database. It requires their agreements and disagreements to remain inspectable.

14.2

Phase 0 establishes the baseline and the contracts

Phase 0 is the illustrative set-up gate; its duration is site-dependent. Begin with the decisions whose delay or error causes an identifiable operating consequence. Record the event needed, the responsible role, the permitted delay, the available manual record and the consequence of missing evidence. Inventory the relevant machines and measurement points against those needs. A machine that exposes a counter is not automatically a reliable source: the acceptance plan must establish what increments it, when it resets and how its meaning changes with the product or operating state.

Survey power, communications and physical access under representative operating conditions. Record interruptions and recovery behaviour rather than substituting a national indicator for local measurements. Identify dependencies on time services, licences, remote identities and vendor support. The proposed register should distinguish an observed condition from an assumed condition awaiting a test. It should also preserve conditions under which the survey was conducted, so later replication does not mistake an acceptance result from an idle line for evidence about loaded production.

Define the logging reference before measuring source coverage. For periodic measurements, specify the expected schedule and treatment of planned shutdowns. For event-driven capture, identify an independent count or reconciled production record that establishes how many events should have existed. An event stream cannot prove its own completeness merely by counting messages it received. Where no independent reference is feasible, record the limitation and restrict the decision accordingly. The baseline may be incomplete, but its incompleteness must be visible rather than converted into a favourable percentage.

The proposed baseline contract should also distinguish source time, gateway receipt time and later processing time. Define the quality states accepted by each use case and the rule for an uncertain clock. Record the identifier and version that travel through transformation, correction and replay. These agreements belong before shadow acceptance because retrospective interpretation cannot establish what an operator actually saw when a decision was due.

Write the KPI dictionary at the same time as the capture register. The ISO description of manufacturing KPIs includes formulas, input elements, time behaviour and units. GFT would turn that principle into an acceptance record containing the calculation boundary, source versions, exclusions, rounding and comparison tolerance. The process owner approves what the result means; the data owner approves how it is reproduced. A manual calculation is a comparison reference only after its own definitions and transcription controls have been reviewed. [S08]

The resource envelope should include the factory's internal capacity as well as purchased equipment. Assign time for process review, integration, physical engineering, safety review, configuration, training, support and recovery exercises. Identify who covers these responsibilities during absence or a supplier dispute. No budget or staff-

ing level is assumed here. The gate should close only when the sponsor can see which necessary tasks are funded, which depend on existing roles and which remain unassigned. An unassigned critical task is an explicit stop condition, not a contingency allowance.

> **GFT NOTE**
>
> Approve the acceptance plan before selecting the pilot platform. Attach each critical use case to an owner, an independent reference, a freshness rule, a reconciliation tolerance and a fallback. Record missing site inputs as unresolved decisions. A supplier demonstration can support selection, but it cannot substitute for the factory's authority to define acceptable operation.

The proposed sequence in **Exhibit 63**, p. 226 gives each phase an exit decision rather than a promised calendar date. Remediation returns to the relevant gate; it does not count as completion of that gate.

EXHIBIT 63 / ILLUSTRATIVE PHASE SEQUENCE

Expansion follows evidence rather than elapsed time



1. Baseline approved
2. Dependence accepted
3. Comparison accepted

ILLUSTRATIVE [S08] [S17] [S22] Illustrative GFT method. Phase identifiers are sequence labels, not durations. Every transition requires its own approval; failed tests retain the current scope and trigger remediation.

Phase 0 exits with the decision register, machine-capture register, interface contracts, zone-and-conduit design, KPI dictionary, recovery objectives and test plan under controlled ownership. The recovery plan should name the affected business decisions and the resources required to restore them. NIST's contingency guidance supports prioritising recovery requirements through an assessment of systems and operations. GFT proposes applying that logic to the factory's decision services without claiming that federal guidance itself constitutes a factory certification. [S22]

14.3

Phase 1 proves one line in shadow and then in operation

The illustrative Phase 1 sequence begins with bounded read-only capture and an explicit rollback route. The selected line should exercise the intended operating decision under meaningful conditions while remaining small enough to investigate a failure. Define the approved products, operating states and shifts in scope. Record what lies outside that scope. If a later proposal adds a different machine family, lighting condition or production route, its owner should identify the evidence needed to extend acceptance rather than assume that physical proximity makes it equivalent.

Shadow mode means that the proposed service records and evaluates decisions while the accepted operating process retains authority. Review event identity, source time, validity flags and transformations before assessing dashboard appearance. Recalculate accepted KPIs independently from the same versioned inputs and compare the results using the tolerance agreed before observation. The proposed difference is platform result minus independent result. Accept only when its absolute value is within the approved tolerance. Matching calculations establish consistency of implementation; they do not prove that the common input record is correct.

Approve an observation window that covers the relevant operating variation and foreseeable recovery cases. A longer window is not automatically better if it repeats the same easy conditions. Retain planned exclusions, unplanned exclusions and missing reference periods separately. If a product change removes the independent record, flag the affected comparison as unavailable. Do not revise the expected-event denominator after seeing the result merely to clear the gate. Any genuine correction to the logging reference needs a reason, an approving owner and a new comparison against the corrected version.

Test loss of the external link, interruption of the acquisition service, restart and controlled replay. Physical interruption tests require the responsible engineering and safety owners to approve the method and conditions. The proposed acceptance question is whether the decision remains legible during loss and whether recovered events reconcile without duplicate operating actions. A replayed event may belong in historical analysis while being too old to trigger a current response. Sparkplug specifies session-state behaviour and stale metrics, but the factory must still define how its consuming display and action process interpret that state. [S40]

Where inspection is included, separate acquisition acceptance from model acceptance and quality-disposition authority. The proposed pilot should preserve representative images or other reviewable evidence under the agreed retention boundary. Quality staff should be able to identify an unusable observation, follow the manual inspection route and record an override. A plausible model output is insufficient when the acquisition condition is unknown. Until the applicable evidence and fallback are accepted, the pilot remains advisory. No open protective-safety exception may be waived into operational dependence under this roadmap.

WORKED EXAMPLE: A FAVOURABLE AVERAGE STILL FAILS A CRITICAL GATE

This illustrative case uses assumed pilot counts solely to demonstrate acceptance arithmetic. Suppose an independently maintained logging reference expects 12,000 critical events during the approved window. The service receives 11,880 unique, valid events, giving coverage of 11,880 divided by 12,000, or 99 percent. Assume the pre-approved overall minimum is 99.5 percent. The observed shortfall is 0.5 percentage points, so the coverage gate fails even before other tests are considered. These are invented acceptance inputs, not a GFT operating result.

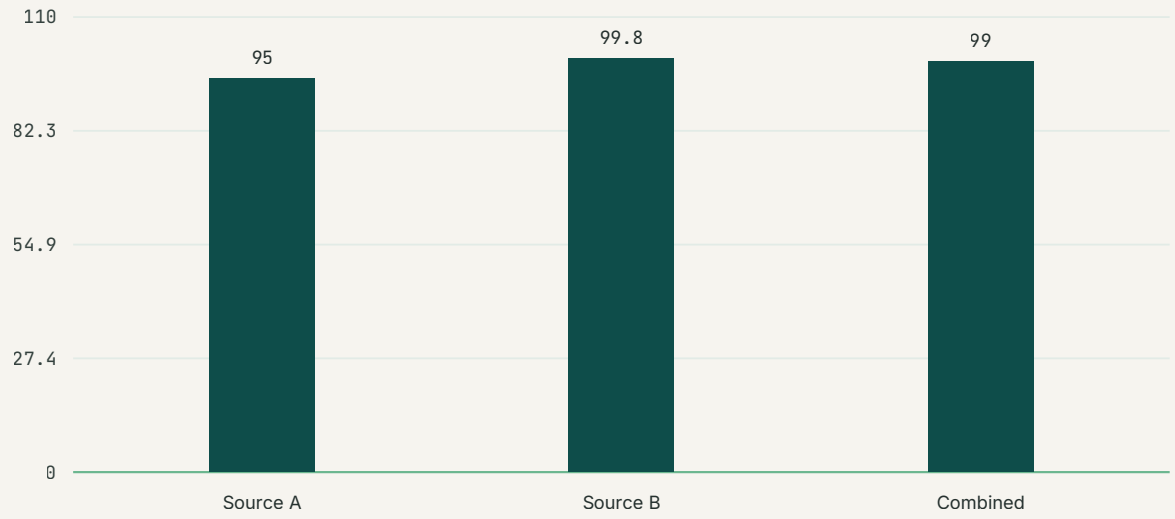
For the same illustrative case, assume the platform reports 9,500 good units and the independent recalculation from the identical input snapshot gives 9,480. The difference is 20 units. With an assumed absolute tolerance of 25 units, this reconciliation test passes. Its success cannot cancel the coverage failure because agreement over available records says nothing about the missing events. The owner retains shadow operation and investigates the missing sources rather than approving expansion on a combined score.

The illustrative aggregate also hides a more severe local failure. Assume critical source A expects 2,000 events and supplies 1,900, while source B expects 10,000 and supplies 9,980. Their respective coverage values are 95 percent and 99.8 percent. Aggregation gives the same 99 percent overall result, but a shared illustrative minimum of 99.5 percent leaves source A unaccepted. A critical use case depending on that source therefore remains blocked even if another use case can be accepted within its narrower boundary.

The arithmetic in **Exhibit 64**, p. 228 makes the example's denominator visible. It is an illustration of the rule, not evidence that these values are suitable thresholds for any factory.

EXHIBIT 64 / ILLUSTRATIVE VALID-EVENT COVERAGE, PERCENT

An aggregate can conceal a weak critical source



ILLUSTRATIVE Illustrative inputs: source A 1900/2000=95%; source B 9980/10000=99.8%; combined 11880/12000=99%. Assumed minimum 99.5% is a teaching threshold only. All events are unique and valid under an assumed approved reference.

Operational acceptance should state the permitted decision, scope, service version, evidence window and fallback. A conditional acceptance can authorise a narrower advisory use without authorising automatic quality disposition or machine control. Attach each remaining exception to a responsible role, expiry condition and compensating measure. A failed critical evidence gate blocks the dependent use. Do not let a deadline convert a shadow demonstration into operating authority. The sponsor should be able to explain exactly what has been accepted without relying on a supplier's general declaration that the pilot works.

The phase decision should record pass, hold, restricted acceptance or withdrawal, with the evidence supporting that state. A hold preserves the earlier authority and identifies the next useful test. Restricted acceptance names the exact use permitted and the uses still blocked. Withdrawal removes reliance on the affected service and activates its accepted fallback. These proposed states prevent a partial technical success from being reported as an unqualified programme success. The record should remain understandable to the next shift and to a reviewer who did not attend the demonstration.

The acceptance schedule in **Exhibit 65**, p. 229 summarises the proposed owner and stopping rule. It is an illustrative governance design, not a derived result from a standard. Thresholds are intentionally site-approved before testing, rather than filled with invented industry benchmarks.

EXHIBIT 65 / ILLUSTRATIVE GATE SCHEDULE

A failed critical condition retains the existing scope

Gate	Accountable owner	Required evidence
Set-up	Sponsor	Contracts and resources
Bounded operation	Process owner	Coverage and safe fallback
Replication	Receiving owner	Comparable results and scope
Optimisation	Service owner	Restore and applicability

ILLUSTRATIVE [S08] [S17] [S22] Illustrative GFT synthesis. Every owner approves objective evidence before transition. No protective-safety exception permits operational dependence. Detailed phase criteria and exception rules are in the prose.

For the bounded-operation gate, the proposed evidence pack includes per-source coverage, freshness under relevant operating states, KPI differences within approved tolerances, completed interruption tests and confirmation that no protective-safety exception remains open. Record dissent and unresolved uncertainty rather than averaging their severity into a readiness score. The sponsor can adjust the programme's sequence or resource envelope, but a commercial preference to expand does not itself supply missing evidence. Where a different decision can be accepted independently, its boundary must exclude reliance on the failed source or service.

For replication, the receiving owner should confirm that the comparison method, target, exclusions and local differences were approved before results were assessed. Require evidence gaps to be closed or, for eligible non-critical limitations, accepted

through the documented exception process. The approval should state why the limitation does not undermine the authorised decision and what would cause that judgement to expire. It should not merely inherit the pilot sponsor's confidence. Repeated exceptions require a review of the common contract before more installations copy the same weakness.

For optimisation, require restoration time and data loss within their approved objectives, plus an authorised review of the new scope. Recovery evidence expires in practical value when a material dependency changes. A replacement identity service, storage format or support arrangement should therefore trigger an applicability review of the earlier exercise. Retest what the change invalidates and retain the unaffected evidence with its original conditions. This avoids both an unsupported blanket renewal and an unnecessary restart of every acceptance activity.

14.4

Phase 2 replicates contracts before adding advanced analytics

The illustrative Phase 2 sequence replicates an accepted decision contract rather than clone an installation directory. Compare the receiving line's asset hierarchy, machine states, product routes and event semantics with the proven scope. Maintain a visible difference register. Differences may be configuration changes, new capture engineering or reasons to repeat part of shadow acceptance. Retain the common contract where its meaning remains valid. Where meaning differs, approve an explicit mapping or a revised version and identify the consequences for cross-line comparison.

Before comparing operating results, approve the baseline, comparison period, exclusions and target values. Track production mix, scheduled time, staffing conditions, material quality, maintenance interventions and other relevant changes alongside the outcome. An improvement following installation is an observation about sequence, not by itself a causal result. GFT would require the reviewer to state the plausible alternative explanations and the evidence used to assess them. If the periods are not meaningfully comparable, report that limit rather than produce a digitalisation benefit percentage from incompatible denominators.

Use a narrow set of operational outcomes that the process owner can influence: good-unit throughput, first-pass yield, downtime, delivery and resource intensity where the measurement boundary supports them. Keep definitions consistent with the approved dictionary, including the treatment of rework, planned stops and product changes. ISO's KPI description provides the discipline of explicit calculation characteristics; the particular target and decision consequence remain the factory's responsibility. A shared name does not justify comparing results calculated under different production boundaries. [S08]

Replication also changes the support burden. The proposed review should establish who maintains mappings, investigates alarms, approves releases and covers out-of-hours failures. Repeated success on an easy line does not resolve an unowned recovery service. Review exception age and repeated overrides before authorising the next group of lines. Where a non-critical exception is accepted, retain its precise scope

and expiry. A protective-safety exception cannot be accepted through the ordinary data-exception process, and any unresolved critical evidence gap continues to block the decision that depends on it.

Advanced analytics should enter through a new use-case contract once its required history, reference outcomes and action process exist. Identify the proposed model's decision horizon, evaluation reference, error consequences and permitted authority. Start with advisory output if that is the accepted risk boundary. Preserve the previous decision method and a rollback route. GFT's measure-before-model sequence is a design recommendation: the existence of stored data does not establish that those data cover the conditions in which the model will later be used.

The proposed release gate should also ask whether a supplier can reproduce a result without privileged access to undocumented transformations. Require the event snapshot, relevant configurations, calculation or model version and exclusion record. These artefacts support acceptance across products and integrators. They also allow the factory to recognise a misleading comparison before it becomes a repeated management target. The purpose of scale is a more dependable operating system, so evidence health and the capacity to act must grow with the installation footprint.

14.5

Phase 3 optimises a system that can still be restored

The illustrative Phase 3 sequence permits cross-line analysis, product evidence and more sophisticated applications only within the same ownership and change-control arrangements. Every added interface needs an identified producer, consumer, meaning and failure behaviour. A simulation or digital representation should disclose which observations update it and which outputs remain assumptions. It must not become an informal control route around the established approval process. Optimisation changes the questions asked of the evidence; it should not erase the records needed to challenge its answers.

Recovery acceptance should use a business-defined recovery time objective and recovery point objective. Specify the start and end of the restoration clock, the service being restored and the evidence used to identify missing data. The proposed operational endpoint is the ability to perform the accepted decision with reconciled records, rather than the moment a server responds. Report the measured restoration time against the approved objective and the measured loss against the approved loss boundary. A restored database with unresolved production identity conflicts is not an accepted decision service.

Keep a recoverable package of configurations, interfaces, dictionaries, identities and required application versions alongside the necessary data. Exercise restoration in an environment that avoids contaminating the recovered service with the original failure. Verify access and the ability to interpret the recovered records. NIST distinguishes testing, training and exercises in its contingency guidance. GFT would therefore retain separate evidence that the technical restore worked, that the responsible roles could perform it and that the decision process survived the disruption. [S22]

An illustrative recovery exercise might set an approved RTO of 120 minutes and an approved RPO of 15 minutes for a reporting service. If the measured return to usable, reconciled operation takes 95 minutes but the proven data-loss interval is 18 minutes, restoration time passes and data loss fails. The excess loss is 3 minutes. These assumed values demonstrate independent tests; they do not recommend objectives or describe a GFT exercise. Expansion of the dependent service remains blocked until the loss failure is resolved and retested.

Replication to another area or site repeats the relevant site survey, risk review, physical engineering acceptance and restoration exercise. Carry forward reusable contracts and test methods, then identify what must be established again. Power conditions, equipment interfaces, access responsibilities and available recovery resources require local evidence. In GFT's proposed governance, the receiving owner signs the applicability review; the originating site's acceptance cannot sign on that owner's behalf. This keeps replication conditional even when the technical package appears identical.

Optimisation must also remain interruptible. Agree when the factory returns to the earlier method, how outstanding actions are handed back and which historical calculations remain under review. Preserve the evidence of a failed release so the same defect is not introduced again under a new application name. The review should distinguish a technically reversible deployment from an operationally recoverable change: restoring software does not automatically undo an erroneous production disposition or recreate an observation that was never captured.

14.6

GFT would measure decision performance and evidence health together

The proposed scorecard should place the operating result beside the evidence needed to trust it and the action needed to change it. Coverage is unique valid critical events received divided by critical events expected under the approved logging reference. Freshness age is review time minus source timestamp, subject to the clock-validity rule. Action closure is verified closed actions divided by actions due, with overdue items visible separately. These are GFT's proposed management definitions. They require agreed units, reference periods, owners and exclusions before use.

Freshness should be assessed at the moment the decision is needed, not only after a replay has repaired historical completeness. Report clock uncertainty and stale state as separate qualifications. A negative calculated age is a clock or timestamp exception to investigate, not evidence of exceptional freshness. For an event-driven source, silence can be legitimate; use the agreed state or heartbeat rule to distinguish quiet production from failed acquisition. Sparkplug's stale-metric provisions inform this requirement without determining the factory's acceptable decision age. [S40]

Keep source coverage disaggregated by critical source and relevant operating condition. Retain missingness, invalid observations, duplicates and late arrivals as different categories. An empty denominator means unavailable or not applicable under an explicit rule, not perfect performance. The scorecard should expose how exclusions

affected the result. This prevents a declining observation population from appearing as an improvement and helps the owner direct remediation towards the actual weak point.

Monitor replay backlog as both a record count and the age of the oldest unresolved event. An apparently draining queue may still leave a critical interval unreconciled, so the review should identify affected sources and decision periods. Record stale-display incidents, invalid clock status and alert acknowledgement delay alongside service availability. These proposed health measures should trigger an assigned investigation or a restricted operating mode, rather than accumulate as unattended diagnostics. Their thresholds must follow the consequence of delayed evidence for the relevant decision.

For actions, distinguish acknowledgement, assignment, intervention and verified closure. An acknowledged alert may still represent an unresolved fault. Record who verified the consequence and whether the action addressed the cause or only restored a temporary condition. Retain human overrides with their reason, affected decision and review status. Frequent justified overrides may show a poor operating rule rather than operator resistance. Repeated unreviewed overrides are a governance gap even when production continues.

The compact scorecard in **Exhibit 66**, p. 233 is a proposed definition set. Its baseline and targets must be approved at the factory before comparison. It contains no assumed measured performance.

EXHIBIT 66 / ILLUSTRATIVE DEFINITION AND OWNERSHIP SCHEDULE

Evidence and action qualify every operating result

Measure	Rule	Accountable owner
Coverage	Valid / expected	Data owner
Freshness	Review minus source time	Process owner
Action closure	Verified closed / due	Process owner
Recovery	Measured against RTO/RPO	Service owner
Access review	List overdue items	Security owner

ILLUSTRATIVE [S08] [S18] [S22] [S40] Illustrative GFT method, not measured values or a prescribed standard. Owner approves baseline window, exclusions and target before comparison. Detailed denominator and verification rules are in the adjacent prose.

Security review should similarly preserve the unresolved item and its owner. List overdue remote-access reviews, affected privileges, expiry and action due rather than presenting an unqualified compliance percentage. NIST's Cybersecurity Framework supplies high-level risk outcomes rather than a prescribed implementation. GFT would use those outcomes to organise responsibility, while requiring local evidence for the controls selected. A favourable reporting score cannot override an unresolved condition that blocks the relevant operating gate. [S18]

14.7

Procurement gates keep ownership and exit rights with the factory

Procurement should purchase an inspectable operating capability and its exit route. Ask bidders to identify the data that can be exported, the schemas supplied, the configuration backup method and the access needed to reproduce accepted results. Require a demonstration using the proposed contractual deliverables. A statement that data belong to the factory is incomplete if only the supplier can interpret their identifiers or restore the application that reads them. These are recommended negotiating requirements; no supplier acceptance is claimed.

Specify the support boundary for security updates, vulnerabilities, remote maintenance and end-of-support events. Require application and model versions to be recorded against the evidence they produce. Tie each release to the factory's approval and rollback process. NIST's OT guidance supports evaluating changes against operating constraints, but procurement must translate those constraints into testable obligations for the actual service. A general promise of secure operation is not an acceptance method. [S17]

Before contract award, establish how the factory would transition to another integrator or operate the accepted fallback during a dispute. Test a representative export and independent recalculation, inspect backup completeness and identify any licences or keys required for restoration. Record costs, dependencies and permitted uses in the proposed terms. Legal review should assess whether those terms achieve the intended ownership and continuity. This chapter specifies the engineering questions; it does not assert enforceable rights in an unsigned agreement.

Final acceptance should be a short, traceable decision record stating the authorised scope, evidence reviewed, unresolved limitations, responsible roles and next review trigger. Scope expansion should cite that record and explain what changes. The sponsor can then distinguish an available prototype, an accepted advisory service, an operational dependency and a replicated capability. Those states need different evidence. Treating them as interchangeable is how a collection of promising demonstrations becomes an unmanaged factory dependency.

GFT's recommendation is conditional on site engineering and applicable legal and regulatory review. This reference architecture is not investment, legal or tax advice. It certifies neither a machine nor a supplier nor a factory. Its useful conclusion is narrower and practical: preserve the production boundary, prove the evidence, accept the decision and retain the ability to recover before the next expansion.

Scale the decision that survives scrutiny, and the system that survives its loss.

CHAPTER SOURCE NOTES

S01 International Society of Automation. 2025-04-10.

S02 International Electrotechnical Commission. 2013-05-22.

S17 National Institute of Standards and Technology. 2023-09.

S08 International Organization for Standardization. 2014-01.

S40 Eclipse Foundation. 2022-11-16.

S22 National Institute of Standards and Technology. 2010-11-11.

S18 National Institute of Standards and Technology. 2024-02-26.

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

Announced SODECOTON and CICAM merger to create a cotton-textile industrial pole targeted by 2030.

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

CICAM restructuring cost estimate CFAF 30.7 to 48.2 billion; one relaunch plan cited at 70.2 billion.

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

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.

R15 Inherited evidence R15 GRADE C

wage.is. Publication date unrecorded. wage.is/cameroon

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.

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

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.

R27 Inherited evidence R27 GRADE B

McKinsey, The State of Fashion 2025. Publication date unrecorded. mckinsey.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.

R28 Inherited evidence R28 GRADE C

DataIntelo. Publication date unrecorded. dataintelo.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.

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

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.

c01-S01 Update to ISA-95 Standard Addresses Integration of Enterprise and Manufacturing Control Systems GRADE A

International Society of Automation. 2025-04-10. https://www.isa.org/news-press-releases/2025/april/update-to-isa-95-standard-addresses-integration-of

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S02 IEC 62264-1:2013: Enterprise-control system integration, Part 1: Models and terminology GRADE A

International Electrotechnical Commission. 2013-05-22. https://webstore.iec.ch/en/publication/6675

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S04 OPC-10000-1: OPC Unified Architecture, Part 1: Overview and Concepts 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-S05 MQTT V5.0 OASIS Standard published GRADE A

OASIS Open. 2019-03-07. https://www.oasis-open.org/2019/03/21/mqtt-v5-0-oasis-standard-published/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S07 Specification of the Asset Administration Shell Part 1: Metamodel, IDTA 01001-3-0-1 GRADE A

Industrial Digital Twin Association. 2024-06. https://industrialdigitaltwin.org/wp-content/uploads/2024/06/IDTA-01001-3-0-1_SpecificationAssetAdministrationShell_Part1_Metamodel.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S08 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. 2014-01. https://www.iso.org/standard/54497.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S09 MODBUS Application Protocol Specification V1.1b3 GRADE A

Modbus Organization. 2012-04-26. https://www.modbus.org/file/secure/modbusprotocolspecification.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S10 IO-Link Interface and System Specification, Version 1.1.4 GRADE A

IO-Link Community. 2024-06. https://io-link.com/fileadmin/user_upload/Downloads/Package_2024/IOL-Interface-Spec_10002_V114_Jun24.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S12 IEEE 1588-2019: IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems GRADE A

IEEE Standards Association. 2020-06-16. https://ieeexplore.ieee.org/document/9120376

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S13 Network Time Protocol Version 4: Protocol and Algorithms Specification GRADE A

Internet Engineering Task Force and RFC Editor. 2010-06. https://www.rfc-editor.org/rfc/rfc5905.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S14 Web of Things Architecture 1.1 GRADE A

World Wide Web Consortium. 2023-12-05. https://www.w3.org/TR/wot-architecture11/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S15 **EPCIS Standard, Release 2.0** GRADE A

GS1 AISBL. 2022-06. https://ref.gs1.org/standards/epcis/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S16 **ETSI GS MEC 003 V3.2.1: Multi-access Edge Computing, Framework and Reference Architecture** GRADE A

European Telecommunications Standards Institute. 2024-04. https://www.etsi.org/deliver/etsi_gs/MEC/001_099/003/03.02.01_60/gs_mec003v030201p.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S17 **NIST SP 800-82 Rev. 3: Guide to Operational Technology Security** GRADE A

National Institute of Standards and Technology. 2023-09. https://csrc.nist.gov/pubs/sp/800/82/r3/final

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S19 **IEC 62443-3-2:2020: Security for industrial automation and control systems, Part 3-2: Security risk assessment for system design** GRADE A

International Electrotechnical Commission. 2020-06. https://webstore.iec.ch/preview/info_iec62443-3-2%7Bed1.0%7Den.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S21 **Russian State-Sponsored and Criminal Cyber Threats to Critical Infrastructure** GRADE A

CISA and joint international cybersecurity authorities. 2022-05-09. https://www.cisa.gov/news-events/cybersecurity-advisories/aa22-110a

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S22 **NIST SP 800-34 Rev. 1: Contingency Planning Guide for Federal Information Systems** 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.

c01-S23 **Ensuring safety in the workplace** GRADE A

International Electrotechnical Commission. 2022-04-28. https://iec.ch/blog/ensuring-safety-workplace

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S24 **ISO/IEC 27001:2022: Information security, cybersecurity and privacy protection, Information security management systems, Requirements** GRADE A

International Organization for Standardization and International Electrotechnical Commission. 2022-10. https://www.iso.org/standard/27001

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S25 **An RFID-based tracking system for denim production processes** GRADE B

Springer Nature, International Journal of Advanced Manufacturing Technology. 2016-09-05. https://doi.org/10.1007/s00170-016-9385-7

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S26 **Textile Defect Detection Using Artificial Intelligence and Computer Vision: A Preliminary Deep Learning Approach** GRADE B

MDPI, Electronics. 2025-09-18. https://doi.org/10.3390/electronics14183692

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S27 **Digital and Sustainable Transition in Textile Industry through Internet of Things Technologies: A Pakistani Case Study** GRADE B

MDPI, Applied Sciences. 2024-06-21. https://doi.org/10.3390/app14135380

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S29 **Industrial IoT-Based Energy Monitoring System: Using Data Processing at Edge** GRADE B

MDPI, IoT. 2024-09-28. https://doi.org/10.3390/iot5040027

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S30 **Six-Sigma Reference Model for Industry 4.0 Implementations in Textile SMEs** GRADE B

MDPI, Sustainability. 2023-08-19. https://doi.org/10.3390/su151612589

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S31 **A Review of IoT Sensing Applications and Challenges Using RFID and Wireless Sensor Networks** GRADE B

MDPI, Sensors. 2020-04-28. https://doi.org/10.3390/s20092495

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S32 **Enterprise Surveys: What Businesses Experience, Cameroon 2024 Country Profile** GRADE A

World Bank Enterprise Surveys. 2025. https://www.enterprisesurveys.org/content/dam/enterprisesurveys/documents/country/Cameroon-2024.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S33 **Facts and Figures 2024: Mobile network coverage** GRADE A

International Telecommunication Union. 2024-11-10. https://www.itu.int/itu-d/reports/statistics/2024/11/10/ff24-mobile-network-coverage/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S35 Industrial Development Report 2020: Industrializing in the digital age GRADE A

United Nations Industrial Development Organization. 2019. https://www.unido.org/sites/default/files/unido-publications/2023-03/UNIDO-IDR2020-main-report-en.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S40 Sparkplug Specification 3.0.0 GRADE A

Eclipse Foundation. 2022-11-16. https://sparkplug.eclipse.org/specification/version/3.0/documents/sparkplug-specification-3.0.0.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S0101 MQTT Version 5.0 OASIS Standard GRADE A

OASIS Open. 2019-03-07. https://docs.oasis-open.org/mqtt/mqtt/v5.0/mqtt-v5.0.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S0102 Web of Things (WoT) Thing Description 1.1 GRADE A

World Wide Web Consortium. 2023-12-05. https://www.w3.org/TR/wot-thing-description11/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c02-S08 Same source record as c01-S08.

c02-S15 Same source record as c01-S15.

c02-S17 Same source record as c01-S17.

c02-S22 Same source record as c01-S22.

c02-S25 An RFID-based tracking system for denim production processes GRADE B

Springer Nature, International Journal of Advanced Manufacturing Technology. 2016-09-05. https://link.springer.com/article/10.1007/s00170-016-9385-7

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c02-S26 Textile Defect Detection Using Artificial Intelligence and Computer Vision: A Preliminary Deep Learning Approach GRADE B

MDPI, Electronics. 2025-09-18. https://mdpi-res.com/d_attachment/electronics/electronics-14-03692/article_deploy/electronics-14-03692.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c02-S27 Digital and Sustainable Transition in Textile Industry through Internet of Things Technologies: A Pakistani Case Study GRADE B

MDPI, Applied Sciences. 2024-06-21. https://mdpi-res.com/d_attachment/applsci/applsci-14-05380/article_deploy/applsci-14-05380.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c02-S29 Industrial IoT-Based Energy Monitoring System: Using Data Processing at Edge GRADE B

MDPI, IoT. 2024-09-28. https://mdpi-res.com/d_attachment/IoT/IoT-05-00027/article_deploy/IoT-05-00027.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c02-S30 Six-Sigma Reference Model for Industry 4.0 Implementations in Textile SMEs GRADE B

MDPI, Sustainability. 2023-08-19. https://mdpi-res.com/d_attachment/sustainability/sustainability-15-12589/article_deploy/sustainability-15-12589.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c02-S38 ISO 14046:2014: Environmental management, Water footprint, Principles, requirements and guidelines GRADE A

International Organization for Standardization. 2014-08. https://www.iso.org/standard/43263.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c02-S0201 Green transition: Council gives its final approval to the ecodesign regulation GRADE A

Council of the European Union. 2024-05-27. https://www.consilium.europa.eu/en/press/press-releases/2024/05/27/green-transition-council-gives-its-final-approval-to-the-ecodesign-regulation/

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-S03 RAMI4.0: A reference framework for digitalisation GRADE A

Plattform Industrie 4.0. 2018-08-09. https://www.plattform-i40.de/IP/Redaktion/EN/Downloads/Publikation/rami40-an-introduction.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c03-S04 Same source record as c01-S04.

c03-S06 ISO 23247-4:2021: Automation systems and integration, Digital twin framework for manufacturing, Part 4: Information exchange GRADE A

International Organization for Standardization. 2021-10. https://www.iso.org/ru/standard/78745.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c03-S07 Same source record as c01-S07.

c03-S08 Same source record as c01-S08.

c03-S14 Same source record as c01-S14.

c03-S15 Same source record as c01-S15.

c03-S16 Same source record as c01-S16.

c03-S17 Same source record as c01-S17.

c03-S18 The NIST Cybersecurity Framework 2.0 GRADE A

National Institute of Standards and Technology. 2024-02-26. https://csrc.nist.gov/pubs/cswp/29/the-nist-cybersecurity-framework-csf-20/final

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c03-S19 Same source record as c01-S19.

c03-S20 Cybersecurity Performance Goals 2.0 GRADE A

Cybersecurity and Infrastructure Security Agency. 2025-12. https://www.cisa.gov/cybersecurity-performance-goals-2-0-cpg-2-0

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c03-S22 Same source record as c01-S22.

c03-S23 Same source record as c01-S23.

c03-S24 Same source record as c01-S24.

c03-S29 Same source record as c01-S29.

c03-S30 Same source record as c01-S30.

c03-S32 Same source record as c01-S32.

c03-S33 Same source record as c01-S33.

c03-S35 Same source record as c01-S35.

c03-S0301 Same source record as c01-S0102.

c03-S0302 Web of Things (WoT) Profiles GRADE A

World Wide Web Consortium. 2025-11-04. https://www.w3.org/TR/wot-profile/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c04-S01 Same source record as c01-S01.

---

c04-S07 Same source record as c01-S07.

---

c04-S13 Same source record as c01-S13.

---

c04-S15 Same source record as c01-S15.

---

c04-S31 Same source record as c01-S31.

---

c04-S40 Same source record as c01-S40.

---

c04-S41 ISO 8000-120:2016: Data quality, Part 120: Master data: Exchange of characteristic data: Provenance GRADE A

International Organization for Standardization. 2016-10. https://www.iso.org/standard/62393.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c04-S0401 IEEE 1588-2019: IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems GRADE A

IEEE Standards Association. 2020-06-16. https://standards.ieee.org/ieee/1588/6825/

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-S04 Same source record as c01-S04.

---

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-S11 PROFINET System Description: Technology and Application GRADE A

PROFIBUS Nutzerorganisation. 2018-11. https://www.profibus.com/fileadmin/media/downloadsection/technical_description_%26_books/PROFINET_System_Description_engl_2018_Update.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c05-S12 Same source record as c01-S12.

---

c05-S13 Same source record as c01-S13.

---

c05-S14 Same source record as c01-S14.

---

c05-S15 Same source record as c01-S15.

---

c05-S17 Same source record as c01-S17.

---

c05-S19 Same source record as c01-S19.

---

c05-S20 Same source record as c03-S20.

---

c05-S21 Same source record as c01-S21.

---

c05-S22 Same source record as c01-S22.

---

c05-S23 Same source record as c01-S23.

---

c05-S24 Same source record as c01-S24.

---

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 RFID-enabled real-time manufacturing execution system for mass-customization production GRADE B

Elsevier, Robotics and Computer-Integrated Manufacturing. 2013-04. https://doi.org/10.1016/j.rcim.2012.08.001

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c05-S29 Same source record as c01-S29.

---

c05-S30 Same source record as c01-S30.

c05-S31 Same source record as c01-S31.

c05-S32 Same source record as c01-S32.

c05-S35 Same source record as c01-S35.

c05-S40 Same source record as c01-S40.

c05-S0501 OPC 40001-1: OPC UA for Machinery, Part 1: Basic Building Blocks (v1.04.1), Section 14: Operation Counter GRADE A

OPC Foundation and VDMA. 2026-01-01. https://reference.opcfoundation.org/Machinery/v1041/docs/14

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c05-S0502 OPC 40001-1: OPC UA for Machinery, Part 1: Basic Building Blocks (v1.04.1), Annex C.2: Example interpretation for ISO 22400 GRADE A

OPC Foundation and VDMA. 2026-01-01. https://reference.opcfoundation.org/Machinery/v1041/docs/C.2

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c05-S0503 CSRC Glossary: unidirectional gateway (from NIST SP 800-82 Rev. 3) GRADE A

National Institute of Standards and Technology. 2023-09. https://csrc.nist.gov/glossary/term/unidirectional_gateway

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c06-S14 Same source record as c01-S14.

c06-S16 Same source record as c01-S16.

c06-S17 Same source record as c01-S17.

c06-S22 Same source record as c01-S22.

c06-S24 Same source record as c01-S24.

c06-S26 Same source record as c01-S26.

c06-S31 Same source record as c01-S31.

c06-S40 Same source record as c01-S40.

c06-S41 Same source record as c04-S41.

c06-S0601 MQTT Version 5.0, OASIS Standard 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.

c06-S0602 Industrial IoT-Based Energy Monitoring System: Using Data Processing at Edge GRADE B

MDPI, IoT (publisher-deposited abstract via Crossref). 2024-09-28. https://api.crossref.org/works/10.3390/iot5040027

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c07-S01 Same source record as c01-S01.

c07-S02 Same source record as c01-S02.

c07-S05 Same source record as c01-S05.

c07-S08 Same source record as c01-S08.

c07-S13 Same source record as c01-S13.

c07-S14 Same source record as c01-S14.

c07-S16 Same source record as c01-S16.

c07-S17 Same source record as c01-S17.

c07-S18 Same source record as c03-S18.

c07-S19 Same source record as c01-S19.

c07-S20 Same source record as c03-S20.

c07-S21 Same source record as c01-S21.

c07-S22 Same source record as c01-S22.

c07-S23 Same source record as c01-S23.

c07-S24 Same source record as c01-S24.

c07-S29 Same source record as c01-S29.

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-S40 Same source record as c01-S40.

c07-S0701 UPS Management Information Base GRADE A

Internet Engineering Task Force and RFC Editor. 1994-05. https://www.rfc-editor.org/rfc/rfc1628.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c07-S0702 The Syslog Protocol GRADE A

Internet Engineering Task Force and RFC Editor. 2009-03. https://www.rfc-editor.org/rfc/rfc5424.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c07-S0703 Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv4 and IPv6 GRADE A

Internet Engineering Task Force and RFC Editor. 2010-03. https://www.rfc-editor.org/rfc/rfc5798.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c08-S01 Same source record as c01-S01.

c08-S02 Same source record as c01-S02.

c08-S03 Same source record as c03-S03.

c08-S04 Same source record as c01-S04.

c08-S05 Same source record as c01-S05.

c08-S06 Same source record as c03-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-S12 Same source record as c01-S12.

c08-S13 Same source record as c01-S13.

c08-S14 Same source record as c01-S14.

c08-S15 Same source record as c01-S15.

c08-S17 Same source record as c01-S17.

c08-S19 Same source record as c01-S19.

c08-S24 Same source record as c01-S24.

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-S28 Same source record as c05-S28.

c08-S29 Same source record as c01-S29.

c08-S30 Same source record as c01-S30.

c08-S31 Same source record as c01-S31.

c08-S32 Same source record as c01-S32.

c08-S33 Same source record as c01-S33.

c08-S35 Same source record as c01-S35.

c08-S39 Regulation (EU) 2024/1781 establishing a framework for the setting of ecodesign requirements for sustainable products GRADE A

European Parliament and Council of the European Union, Official Journal of the European Union. 2024-06-28. https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32024R1781

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c08-S40 Same source record as c01-S40.

c08-S41 Same source record as c04-S41.

c08-S0801 Same source record as c01-S0102.

c08-S0802 OPC-10000-14: OPC Unified Architecture, Part 14: PubSub GRADE A

OPC Foundation. 2025-10-22. https://reference.opcfoundation.org/specs/OPC-10000-14/full

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c08-S0803 Same source record as c01-S0101.

c08-S0804 RFC 8259: The JavaScript Object Notation (JSON) Data Interchange Format GRADE A

Internet Engineering Task Force and RFC Editor. 2017-12. https://www.rfc-editor.org/rfc/rfc8259.txt

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c09-S01 Same source record as c01-S01.

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-S06 Same source record as c03-S06.

c09-S07 Same source record as c01-S07.

c09-S08 Same source record as c01-S08.

c09-S09 Same source record as c01-S09.

c09-S12 Same source record as c01-S12.

c09-S13 Same source record as c01-S13.

c09-S14 Same source record as c01-S14.

c09-S15 Same source record as c01-S15.

c09-S16 Same source record as c01-S16.

c09-S17 Same source record as c01-S17.

c09-S18 Same source record as c03-S18.

c09-S19 Same source record as c01-S19.

c09-S21 Same source record as c01-S21.

c09-S22 Same source record as c01-S22.

c09-S24 Same source record as c01-S24.

c09-S25 Same source record as c01-S25.

c09-S26 Same source record as c01-S26.

c09-S27 Same source record as c01-S27.

c09-S28 Same source record as c05-S28.

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-S36 Product Life Cycle Accounting and Reporting Standard GRADE B

GHG Protocol, World Resources Institute and World Business Council for Sustainable Development. 2011-09. https://ghgprotocol.org/sites/default/files/ghgp/standards/Product-Life-Cycle-Accounting-Reporting-Standard_041613.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c09-S37 ISO 50001:2018: Energy management systems, Requirements with guidance for use GRADE A

International Organization for Standardization. 2018-08. https://www.iso.org/standard/69426.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c09-S38 Same source record as c02-S38.

c09-S39 Same source record as c08-S39.

c09-S40 Same source record as c01-S40.

c09-S41 Same source record as c04-S41.

c09-S0901 OPC 10000-11: OPC Unified Architecture, Part 11: Historical Access, Release 1.05.04, Section 4.3: Historians and interruption of data collection GRADE A

OPC Foundation. 2024-10-15. https://reference.opcfoundation.org/Core/Part11/v105/docs/4.3

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c09-S0902 OPC 10000-11: OPC Unified Architecture, Part 11: Historical Access, Release 1.05.04, Section 4.4: Modification of Historical Data/Events GRADE A

OPC Foundation. 2024-10-15. https://reference.opcfoundation.org/Core/Part11/v105/docs/4.4

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c09-S0903 OPC 10000-11: OPC Unified Architecture, Part 11: Historical Access, Release 1.05.04, Section 4.5: Timestamps GRADE A

OPC Foundation. 2024-10-15. https://reference.opcfoundation.org/Core/Part11/v105/docs/4.5

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c09-S0904 PROV-Overview: An Overview of the PROV Family of Documents GRADE A

World Wide Web Consortium. 2013-04-30. https://www.w3.org/TR/prov-overview/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c10-S01 Same source record as c01-S01.

c10-S02 Same source record as c01-S02.

c10-S04 Same source record as c01-S04.

c10-S08 Same source record as c01-S08.

c10-S14 Same source record as c01-S14.

c10-S15 Same source record as c01-S15.

c10-S16 Same source record as c01-S16.

c10-S17 Same source record as c01-S17.

c10-S19 Same source record as c01-S19.

c10-S20 Same source record as c03-S20.

c10-S21 Same source record as c01-S21.

c10-S22 Same source record as c01-S22.

c10-S23 Same source record as c01-S23.

c10-S24 Same source record as c01-S24.

c10-S25 Same source record as c01-S25.

c10-S26 Same source record as c01-S26.

c10-S27 Same source record as c01-S27.

c10-S28 Same source record as c05-S28.

c10-S29 Same source record as c01-S29.

c10-S30 Same source record as c01-S30.

c10-S31 Same source record as c01-S31.

c10-S32 Same source record as c01-S32.

c10-S33 Same source record as c01-S33.

c10-S35 Same source record as c01-S35.

c10-S37 Same source record as c09-S37.

c10-S38 Same source record as c02-S38.

c10-S40 Same source record as c01-S40.

c10-S1001 IEC 62682:2022: Management of alarm systems for the process industries GRADE A

International Electrotechnical Commission. 2022-12-08. https://webstore.iec.ch/en/publication/65543

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c10-S1002 **IEC 62541-9:2026: OPC unified architecture, Part 9: Alarms and Conditions** GRADE A

International Electrotechnical Commission. 2026-04-16. https://webstore.iec.ch/en/publication/82277

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c11-S06 Same source record as c03-S06.

---

c11-S14 Same source record as c01-S14.

---

c11-S16 Same source record as c01-S16.

---

c11-S17 Same source record as c01-S17.

---

c11-S18 Same source record as c03-S18.

---

c11-S22 Same source record as c01-S22.

---

c11-S24 Same source record as c01-S24.

---

c11-S30 Same source record as c02-S30.

---

c11-S31 **A Review of IoT Sensing Applications and Challenges Using RFID and Wireless Sensor Networks** GRADE B

MDPI, Sensors. 2020-04-28. https://pmc.ncbi.nlm.nih.gov/articles/PMC7249175/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c11-S41 Same source record as c04-S41.

---

c11-S1101 **Textile Defect Detection Using Artificial Intelligence and Computer Vision: A Preliminary Deep Learning Approach** GRADE B

MDPI, Electronics; publisher-deposited abstract accessed through Crossref. 2025-09-18. https://api.crossref.org/works/10.3390/electronics14183692

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c11-S1102 **Industrial IoT-Based Energy Monitoring System: Using Data Processing at Edge** GRADE B

MDPI, IoT; publisher-deposited abstract accessed through Crossref. 2024-09-28. https://api.crossref.org/works/10.3390%2Fiot5040027

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c12-S01 Same source record as c01-S01.

---

c12-S02 Same source record as c01-S02.

c12-S05 Same source record as c01-S05.

c12-S06 Same source record as c03-S06.

c12-S07 Same source record as c01-S07.

c12-S08 Same source record as c01-S08.

c12-S09 Same source record as c01-S09.

c12-S10 Same source record as c01-S10.

c12-S12 Same source record as c01-S12.

c12-S13 Same source record as c01-S13.

c12-S14 Same source record as c01-S14.

c12-S15 Same source record as c01-S15.

c12-S17 Same source record as c01-S17.

c12-S19 Same source record as c01-S19.

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-S24 Same source record as c01-S24.

c12-S25 Same source record as c01-S25.

c12-S27 Same source record as c01-S27.

c12-S28 Same source record as c05-S28.

c12-S29 Same source record as c01-S29.

c12-S30 Same source record as c01-S30.

c12-S31 Same source record as c01-S31.

c12-S32 Same source record as c01-S32.

c12-S34 Planning and prospects for renewable power: Central Africa GRADE A

International Renewable Energy Agency. 2024. https://www.irena.org/-/media/Files/IRENA/Agency/Publication/2025/Jan/IRENA_Planning_Prospects_Central_Africa_2025.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c12-S36 Same source record as c09-S36.

c12-S37 Same source record as c09-S37.

c12-S38 Same source record as c02-S38.

c12-S39 Same source record as c08-S39.

c12-S40 Same source record as c01-S40.

c12-S41 Same source record as c04-S41.

c12-S1201 GHG Protocol Scope 2 Guidance GRADE A

GHG Protocol, World Resources Institute and World Business Council for Sustainable Development. 2015-01. https://ghgprotocol.org/scope-2-guidance

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c12-S1202 The GHG Protocol Corporate Accounting and Reporting Standard GRADE A

GHG Protocol, World Resources Institute and World Business Council for Sustainable Development. 2015. https://ghgprotocol.org/corporate-standard

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c13-S01 Same source record as c01-S01.

c13-S02 Same source record as c01-S02.

c13-S04 Same source record as c01-S04.

c13-S05 Same source record as c01-S05.

c13-S08 Same source record as c01-S08.

c13-S09 Same source record as c01-S09.

c13-S10 Same source record as c01-S10.

c13-S11 Same source record as c05-S11.

c13-S12 Same source record as c01-S12.

c13-S13 Same source record as c01-S13.

c13-S16 Same source record as c01-S16.

c13-S17 Same source record as c01-S17.

c13-S18 Same source record as c03-S18.

c13-S19 Same source record as c01-S19.

c13-S20 Same source record as c03-S20.

c13-S21 Same source record as c01-S21.

c13-S22 Same source record as c01-S22.

c13-S23 Same source record as c01-S23.

c13-S24 Same source record as c01-S24.

c13-S26 Same source record as c01-S26.

c13-S29 Same source record as c01-S29.

c13-S30 Same source record as c01-S30.

c13-S31 Same source record as c01-S31.

c13-S32 Same source record as c01-S32.

c13-S33 Same source record as c01-S33.

c13-S35 Same source record as c01-S35.

c13-S40 Same source record as c01-S40.

c13-S1301 NIST SP 800-61 Rev. 3: Incident Response Recommendations and Considerations for Cybersecurity Risk Management: A CSF 2.0 Community Profile GRADE A

National Institute of Standards and Technology. 2025-04. https://csrc.nist.gov/pubs/sp/800/61/r3/final

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c13-S1302 IEC 61508-1:2010: Functional safety of electrical/electronic/programmable electronic safety-related systems, Part 1: General requirements GRADE A

International Electrotechnical Commission. 2010-04-30. https://webstore.iec.ch/en/publication/5515

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c14-S01 Same source record as c01-S01.

c14-S02 Same source record as c01-S02.

c14-S08 Same source record as c01-S08.

c14-S17 Same source record as c01-S17.

c14-S18 Same source record as c03-S18.

c14-S22 Same source record as c01-S22.

c14-S40 Same source record as c01-S40.

Stable definitions make comparisons useful.

Asset Administration Shell

A standardised digital representation of an asset and its submodels, used to exchange durable asset information between systems and partners.

Andon

A visual or audible signal that calls attention to a production abnormality and requests a defined response near the line.

Asset hierarchy

The ordered representation of sites, areas, lines, work centres, machines and measurement points used to give data stable operational context.

Birth certificate

A Sparkplug message that declares an edge node, device or host application and its current metrics or online state at session establishment.

Brownfield

An existing operating facility with equipment, constraints and records that a digital architecture must accommodate rather than assume away.

Canonical data model

A governed shared representation used to map between systems without allowing every pair of applications to invent a separate translation.

Cloud computing

Computing services delivered from remote shared infrastructure, used where their operational benefit justifies connectivity, dependency and administration.

Conduit

The explicitly defined communication path between security zones, including the permitted services, protocols and controls.

Data freshness

The elapsed time and validity state of a value relative to its source event and the decision that consumes it.

Data provenance

The recorded origin and transformation history of a value, sufficient to interpret its credibility and reproduce a derived result.

Data quality

The fitness of data for a specified use, including accuracy, completeness, timeliness, consistency and contextual validity.

Death certificate

A Sparkplug message, normally registered through MQTT's Will mechanism, that informs consumers when a session ends or a node is offline.

Digital product passport

Product-specific data made accessible electronically through a data carrier under the requirements applicable to that product.

Digital twin

A digital representation of an observable manufacturing element with defined information exchange and synchronisation between relevant entities.

Edge computing

Computation performed near the production process or network edge to meet continuity, latency, data-volume or local-control requirements.

Edge node

A computing and gateway function that acquires local device data, applies edge services and communicates with host applications.

EPCIS

A GS1 standard for creating and sharing visibility event data about objects and processes within and across enterprises.

Event

A contextual record of something that happened, carrying identity, time, meaning and quality rather than only a measured value.

Functional safety

The part of overall safety that provides automatic protection in response to specified errors or hazardous conditions.

Gateway

A component that translates or transports data between device, network or application boundaries under a defined interface contract.

Historian

A system that stores time-ordered industrial observations and related context so trends, calculations and investigations can be reproduced.

Human-machine interface

The operational display and interaction surface through which people observe and, where authorised, operate a machine or process.

Industrial Internet of Things

Connected sensors, devices, edge services and applications used to acquire and act on industrial operating data.

ISA-95

A family of standards for enterprise-control integration, manufacturing operations, asset organisation and shared information models.

Key performance indicator

A defined measure with a formula, source elements, time behaviour, unit and decision purpose.

Manufacturing execution system

Software that supports the execution, recording and coordination of production orders, resources, operations and shop-floor events.

Manufacturing operations management

The activities that coordinate and optimise production, quality, inventory, personnel, equipment and related information within manufacturing operations.

MQTT

A publish-subscribe messaging protocol suited to constrained clients and networks, which requires separate conventions for payload and factory meaning.

Network Time Protocol

A protocol for synchronising computer clocks across networks, commonly used for general server, edge and application time alignment.

Overall equipment effectiveness

A manufacturing KPI that combines availability, performance and quality under an explicit calculation and time boundary.

OPC UA

An industrial interoperability architecture whose information and service model is separated from underlying transport and computing technology.

Operational technology

Hardware and software that monitors or controls physical devices, processes and industrial operations, with distinct performance and safety needs.

Programmable logic controller

An industrial controller that executes machine or process logic and exposes selected states and commands through supported interfaces.

Precision Time Protocol

The IEEE 1588 clock-synchronisation protocol used where networked measurement and control require precise time alignment.

Quality flag

A machine-readable state that tells a consumer whether a value is good, stale, missing, estimated, corrected or otherwise limited.

RAMI 4.0

The Reference Architecture Model Industrie 4.0, a three-dimensional framework for classifying assets, hierarchy levels and life-cycle stages.

Reference architecture

A reusable specification of responsibilities, components, interfaces and constraints that guides implementations without prescribing one supplier stack.

Report by exception

A communication pattern that publishes changed values rather than repeatedly sending unchanged data, while retaining session-state awareness.

Recovery point objective

The maximum tolerated loss of data measured backwards from a disruption, set by the process owner and tested through restoration.

Recovery time objective

The target elapsed time for restoring a service or operating capability after disruption, verified through an exercise.

SCADA

Supervisory control and data acquisition, the class of systems used to monitor and supervise distributed industrial processes.

Semantic interoperability

The ability of exchanging systems to interpret shared information with the same operational meaning, units, state and context.

Source timestamp

The time assigned where an observation or event originated, preserved separately from later gateway receipt or storage times.

Store-and-forward

A continuity pattern that persists data locally during communication failure and forwards it under controlled ordering and reconciliation after recovery.

System of record

The authoritative source for a defined class of data and its history, with explicit ownership and rules for correction and synchronisation.

Takt time

The allowable production interval derived from available production time and customer demand, used to assess the required rhythm of a line.

Time series

A sequence of observations associated with timestamps and, where required, quality and provenance metadata.

Work in process

Material, components, bundles or garments that have entered production but have not yet completed the defined process.

Zone

A grouping of industrial assets with common security requirements and a defined boundary, connected to other zones only through authorised conduits.

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.

Andon 109 161

Asset Administration Shell 44 52 53 56 57 60 61 73 123 129 138 143 188 200

Asset hierarchy 46 188 230

Birth certificate 127 135

Brownfield 22 25 27 48 73 74 75 77 78 81 83 85 86 87 104 109 116 117 123 124 129 205 208 211

Canonical data model 132

Cloud computing 261

Conduit 49 50 55 57 76 78 84 85 87 104 106 107 110 112 113 115 116 117 118 153 167 170 171 206 207 209 210 211 213 214 226

Data freshness 261

Data provenance 153 195

Data quality 169

Death certificate 127

Digital product passport 38 130 135 187 199 202

Digital twin 56 57 129

Edge computing 20 44 53 105 146 206

Edge node 51 75 76 85 109 110 111 112 115 116 117 124 127 144 146 147 148 150 154 165 169 187 206 211 215 216

EPCIS 27 35 38 57 60 65 68 82 117 123 128 130 131 132 133 136 137 138 142 143 164 170 198 199 200 202

Event 14 15 19 22 23 24 27 30 31 35 38 39 40 41 42 50 51 53 55 57 58 60 61 62 63 65 66 67 68 69 70 71 78 81 82 94 95 96 97 101 108 110 111 112 113 114 117 119 123 125 128 130 131 132 133 138 141 142 143 144 145 147 148 149 151 153 155 164 168 174 176 178 179 181 182 183 184 187 199 200 202 209 210 214 215 223 224 225 227 230 231 232 233

Functional safety 25 44 47 49 85 87 105 158 171 205 212

Gateway 16 25 44 46 47 48 49 50 57 61 63 64 65 66 70 73 75 76 77 79 80 85 87 104 108 109 110 112 114 115 116 117 119 123 124 125 144 146 150 153 154 164 167 170 188 205 207 208 209 210 212 217 224 225

Historian 22 24 27 57 65 67 70 104 105 107 108 111 113 114 115 117 119 120 123 124 126 127 128 130 135 138 141 144 145 146 147 148 149 150 151 153 155 158 160 161 166 167 170 187 195 198 200 202 214 215 216

Human-machine interface 48 127 161

Industrial Internet of Things 127 187

ISA-95 24 27 44 45 46 47 50 51 52 53 60 73 104 113 116 117 124 129 132 133 138 143 162 164 188 194 200 206

Key performance indicator 105 141 158

Manufacturing execution system 46 130

Manufacturing operations management 24 44 46 50 105 124 141 161

MQTT 19 20 26 27 50 52 62 94 97 106 123 125 126 127 129 130 131 132 133 136 138 144 160 169 171 189 200 206 212 214

Network Time Protocol 18 79 106 127 144 188

OPC UA 22 26 27 54 74 79 80 117 123 124 125 126 128 129 131 132 133 136 137 138 142 167 214

Operational technology 19 44 48 50 105 107 170 205 208

Overall equipment effectiveness 31 51 108 148 196

Precision Time Protocol 79 127 144 188 209

Programmable logic controller 44 46 77 123 205

Quality flag 148 159

RAMI 4.0 44 52 129

Recovery point objective 118 154 215 231

Recovery time objective 85 118 154 215 231

Reference architecture 14 22 27 50 52 56 57 73 84 118 129 132 176 187 202 205 209 211 234

Report by exception 263

SCADA 127 131 165

Semantic interoperability 27 126 138

Source timestamp 14 51 108 110 113 115 127 141 159 165 217 232

Store-and-forward 26 53 96 110 119 127

System of record 123 135 141 153

Takt time 150

Time series 110

Work in process 34 82

Zone 25 50 79 80 82 83 84 85 87 104 105 106 107 108 109 111 114 115 116 117 130 146 153 170 205 206 207 209 210 211 213 214 216 220 223 226

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 wp18. 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.