GREEN FUTURE TECH
LIMITED / HONG KONG SAR

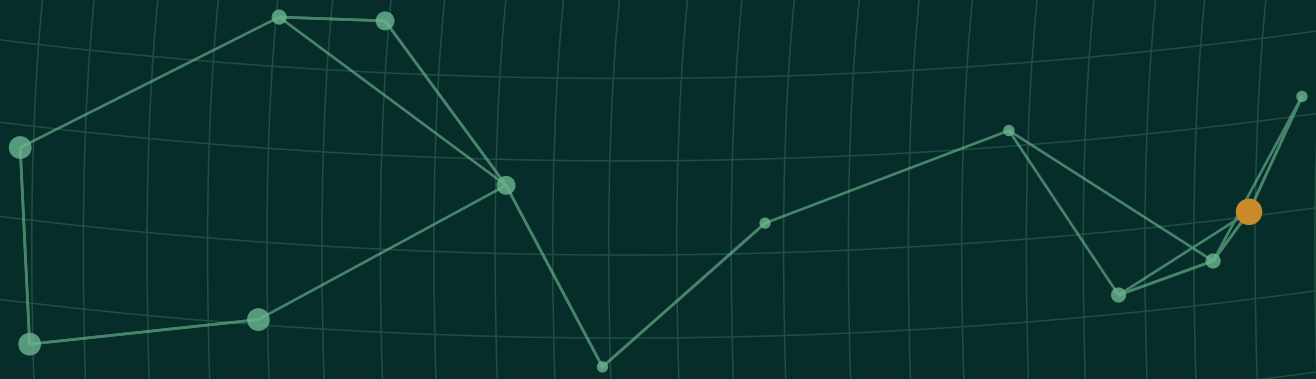
GFT RESEARCH / WARP AND SIGNAL

# Factories on a Weak Grid

*Designing industrial digital systems for intermittent power and variable connectivity*

wp07
RESEARCH EDITION

EDITION 2026-10-30
HONG KONG SAR

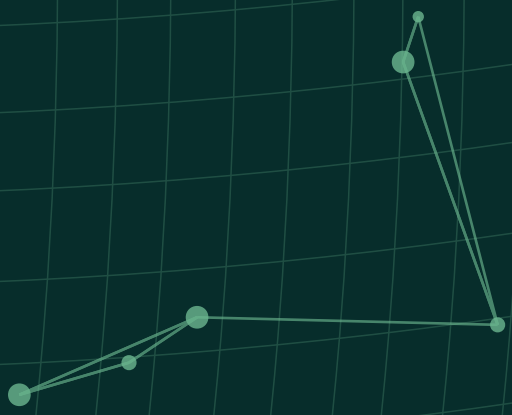


THE GFT THESIS

# 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** | wp07 / Factories on a Weak Grid |
| **EDITION DATE** | 2026-10-30 |
| **AUDIENCE** | Factory sponsors and plant managers assessing industrial digital resilience, Operations, maintenance and electrical engineering leads defining continuity requirements, Factory IT and OT integrators implementing local capture and recovery, Quality leaders and supervisors responsible for offline working, Procurement and finance teams evaluating evidence and ownership cost |
| **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: R03, R06, R07, R09, R10, R11, R12, R13, R14, R15, R16, R32. Temporal eligibility is unverified; these are inherited source-text checks, not independent source validation.

# The operating case runs from fibre to feedback.


## 01 / The factory must keep its record when the grid fails 13

1.1 / A weak supply makes continuity part of the production design 14

1.2 / A production stop and a lost record are different failures 16

1.3 / Cloud reachability must not decide whether the line can record work 18

1.4 / Continuity is a contract with a measurable boundary 19

1.5 / The programme supplies a starting specification, not an operating proof 22

1.6 / The case for resilience must survive its counterargument 23

1.7 / This paper owns failure behaviour across the power and data boundary 24

## 02 / Site measurements must replace national averages 27

2.1 / Country statistics establish context and cannot size a site UPS 28

2.3 / Power quality needs an event log as well as an energy meter 30

2.5 / A reachable modem does not establish usable connectivity 32

2.7 / The joint failure history matters more than separate averages 33

2.8 / Tail durations decide retention and autonomy requirements 35

2.10 / Restart and backlog recovery belong in the baseline 37

2.12 / The site dossier must expose what remains unknown 38

## 03 / Local service classes turn resilience into a testable contract 42

3.1 / Classify tasks by the consequence of interruption 43

3.2 / Availability and durability need separate budgets 45

3.3 / Every dependency must inherit an offline requirement 47

3.4 / Safe degradation must be an explicit state 49

3.5 / Bound the retained record by time, volume and priority 51

3.6 / Manual continuity requires the same identifiers 53

3.7 / Acceptance is evidence of behaviour under a named failure 54

## 04 / A protected digital circuit should have a measured duty 56

4.1 / Separate production energy from digital continuity loads 57

4.2 / Measure consumption in the relevant operating modes 59

4.3 / Size energy from duration and explicit derating inputs 60

4.4 / Transfer behaviour decides whether the computer actually rides through 62

4.5 / Controlled shutdown is a fall-back with dependencies 64

4.6 / Environmental and maintenance assumptions change usable autonomy 66

4.7 / Procurement must buy the promised protected function 67

## 05 / A working local network should outlast a failed backhaul 69

5.1 / Local operation begins with local dependencies 70

5.4 / Measure useful application capacity rather than advertised bandwidth 73

5.7 / The transport protocol must fit device and failure behaviour 75

5.10 / A second carrier is useful only if its failure path differs 77

5.12 / Retry policy must preserve capacity for urgent work 79

5.15 / Local administration must remain controlled while offline 81

5.18 / Network qualification should reproduce the actual failure modes 83

## 06 / The local event must survive the acknowledgement 86

6.1 / An event envelope should carry enough context to recover 87

6.3 / Write the event and outbox in one transaction 89

6.6 / Consistency does not guarantee the latest commit is durable 91

6.10 / Keep a local database on a supported storage path 94

6.13 / Broker persistence is a separate configuration decision 98

6.17 / Full storage must trigger a visible and bounded policy 100

6.21 / A durable record needs a demonstrated recovery path 105

## 07 / A bounded queue makes disconnection manageable 110

7.1 / Calculate retention from event production and outage duration 111

7.3 / Separate the durable source queue from the uplink spool 112

7.5 / Use idempotent ingestion to make duplicates harmless 114

7.7 / Queue-drain capacity must exceed live arrivals 115

7.9 / Priority and expiry rules must be agreed before saturation 118

7.10 / Backpressure must reach a responsible person 120

7.11 / Replay success is a reconciled ledger, not an empty buffer 121

## 08 / A restarted clock must not rewrite production history 124

8.1 / Timestamp formatting and timestamp accuracy are different 125

8.3 / Identity must survive a lack of online allocation 127

8.5 / Clock holdover must be measured on the chosen device 128

8.7 / Time recovery can invalidate an unsafe ordering assumption 131

8.9 / Authenticated time introduces a bootstrap dependency 133

8.11 / Manual and automated records need a conflict policy 135

8.13 / The dashboard should show unresolved uncertainty 137

## 09 / Local inference still needs a complete operating envelope 140

9.1 / An edge model is one dependency within an inspection function 141

9.4 / Offline startup must include model and licence availability 143

9.7 / Inference delay must be bounded by the operating task 145

9.10 / Human fallbacks must preserve quality authority 147

9.13 / A supervisor needs the local view before the remote dashboard 150

9.16 / Energy scarcity should trigger a declared compute policy 151

9.19 / Model updates must have an offline recovery route 152

## 10 / Disconnection changes security dependencies rather than removing them 155

10.1 / Map trust boundaries to the local operating duties 156

10.3 / Local access needs expiry and recovery rules 158
10.4 / Certificate and time failures can block restoration 159
10.7 / Integrity checks must cover firmware and local software 161
10.9 / Offline updates require controlled media and staged activation 163
10.12 / Local logs need an explicit persistence and capacity policy 164
10.15 / An incident response must preserve safe plant control 167

**11 / Recovery finishes when the record and the process agree** 171

11.1 / Power return should start a recovery sequence 172
11.2 / Service supervision needs bounded retries 174
11.3 / A backup must be consistent before it can be useful 175
11.5 / Replacement equipment needs a tested identity handover 177
11.6 / Replay after restore can reintroduce old events 179
11.7 / Operator release requires an understandable recovery record 181
11.8 / Recovery capability should improve after each exercise 183

**12 / The economic case should price avoidable loss and restoration** 185

12.1 / Start with the loss the protected service can actually avoid 186
12.4 / Use contribution margin when estimating lost output value 189
12.7 / Expose the ownership cost across the selected period 191
12.9 / Test sensitivity instead of publishing an unsupported payback 193
12.12 / A vendor claim should become an acceptance obligation 196
12.14 / Choose a maintainable option for the actual local team 198
12.16 / Contractual exit must preserve the production history 199

**13 / Commissioning must challenge the system at its failure boundaries** 203

13.1 / Map each requirement to an observable failure experiment 204
13.3 / Simulation and hardware exercises answer different questions 206
13.4 / A deterministic event ledger reveals missing and duplicate work 207
13.6 / Combine faults that share a plausible dependency 209
13.8 / Recovery timing should be measured from agreed boundaries 211
13.9 / A pilot needs operating observations as well as lab results 213
13.10 / Acceptance should name versions, evidence and residual uncertainty 214

**14 / Scale the design only after continuity is demonstrated** 217

14.1 / Begin by agreeing the production and record boundaries 218
14.2 / Measure the site before fixing autonomy and retention 219
14.3 / Prove a minimal durable local record before broad integration 221
14.4 / Add inspection and remote reporting without weakening the boundary 223
14.5 / Expand only when the local team can restore service 224
14.6 / Measure survival, freshness and recovery in the PMS cadence 225
14.7 / GFT should report verified outcomes and the gaps that remain 227

Executive summary 11
Evidence register 230
Glossary 251
Term index 254
Method and imprint 256

# Every exhibit states a testable claim.

EXHIBIT 01 / bar
Outages affected most firms in the Cameroon survey 14

EXHIBIT 02 / table
Each function needs a separate interruption decision 17

EXHIBIT 03 / flow
Local acceptance can precede restored remote reporting 18

EXHIBIT 04 / waterfall
Different acceptance boundaries produce different counts 21

EXHIBIT 05 / bar
Outage frequency is a separate survey measure 29

EXHIBIT 06 / heatmap
Coincident failures reveal the real design envelope 34

EXHIBIT 07 / timeline
Connectivity recovery can lag restoration of mains 37

EXHIBIT 08 / table
Every site unknown has a collection method 39

EXHIBIT 09 / matrix
Consequence and freshness require separate judgements 44

EXHIBIT 10 / bar
The illustrative outage uses only part of its time allowance 47

EXHIBIT 11 / layered-architecture
An offline inspection duty includes its supporting services 48

EXHIBIT 12 / flow
Returning dependencies must pass a recovery hold 50

EXHIBIT 13 / table
Capacity pressure requires an approved record policy 52

EXHIBIT 14 / stacked
The record-only load includes six supporting components 58

EXHIBIT 15 / table
The nominal energy result exposes every assumption 61

EXHIBIT 16 / range
Duty and derating assumptions move the energy requirement 62

EXHIBIT 17 / flow
Orderly shutdown still ends at crash-recovery checks 65

EXHIBIT 18 / layered-architecture
Local work can survive the loss of an external route 71

EXHIBIT 19 / matrix
Protocol selection depends on the required delivery contract 75

EXHIBIT 20 / table
Common-cause failures defeat naive dual-carrier availability math 78

EXHIBIT 21 / timeline
Backhaul restoration is a sequence of dependencies 81

EXHIBIT 22 / flow
The acknowledgement follows the durable transaction 91

EXHIBIT 23 / matrix
Persistence claims have different failure boundaries 96

EXHIBIT 24 / waterfall
Retention must reserve space beyond event payload 102

EXHIBIT 25 / hbar
Outbox and index structures multiply the raw event footprint 103

EXHIBIT 26 / table
A known event ledger makes loss and duplication visible 107

EXHIBIT 27 / area
An outage creates a backlog that a fast link must clear 117

EXHIBIT 28 / heatmap
A link that cannot exceed arrivals never catches up 118

EXHIBIT 29 / table
Retention priorities make overload behaviour reviewable 119

EXHIBIT 30 / flow
Reconciliation closes the delivery loop 121

EXHIBIT 31 / line
Clock holdover consumes the timestamp error budget linearly with offline duration 130

EXHIBIT 32 / timeline
Source sequence and boot identity preserve local event order through a clock step 132

EXHIBIT 33 / table
Authenticated time and certificate checks depend on persisted clock state 135

EXHIBIT 34 / matrix
Conflicts between manual and automated records require an explicit reconciliation decision 137

EXHIBIT 35 / layered-architecture
Inspection continuity includes its full dependency chain 142

EXHIBIT 36 / table
Physical fabric speed sets the hard ceiling on edge inference latency 146

EXHIBIT 37 / matrix
The operator must see why a decision is unavailable 148

EXHIBIT 38 / flow
An interrupted update returns to a known-good service 153

EXHIBIT 39 / matrix
Each trust boundary needs an offline access rule 157

EXHIBIT 40 / flow
Trust and time must recover together 161

EXHIBIT 41 / table
Recoverable evidence requires configured local retention 165

EXHIBIT 42 / waterfall
A dedicated security partition protects ledger capacity 167

EXHIBIT 43 / gantt
Service restoration follows dependencies rather than a common restart command 173

EXHIBIT 44 / flow
A restored backup must pass ledger checks before release 176

EXHIBIT 45 / table
A spare is useful only if its handover is rehearsed 178

EXHIBIT 46 / waterfall
Reconciliation adds recovered work rather than replay volume 181

EXHIBIT 47 / waterfall
Only attributable losses belong in the benefit case 188

EXHIBIT 48 / table
Cost ownership should remain explicit after procurement 192

EXHIBIT 49 / heatmap
The preferred option changes when outage and recovery assumptions change 195

EXHIBIT 50 / matrix
The purchase must specify how continuity will be proved 197

EXHIBIT 51 / matrix
Every failure requirement should have an evidence-producing test 205

EXHIBIT 52 / table
Acceptance counts reconciled records rather than green process indicators 208

EXHIBIT 53 / timeline
A recovery trace separates service return from ledger closure 212

EXHIBIT 54 / flow
Unexplained loss blocks expansion of the pilot 214

EXHIBIT 55 / gantt
The roadmap advances through evidence gates 219

EXHIBIT 56 / flow
The minimal record earns each additional dependency 222

EXHIBIT 57 / table
Every continuity measure needs a numerator and denominator 226

EXHIBIT 58 / matrix
Completion status must state the evidence actually obtained 228

# Build the factory and its evidence as one system.

Industrial digital systems in Cameroon must be designed around the conditions in which production actually runs. Intermittent electricity and variable connectivity can interrupt inspection, erase an inadequately protected record and leave management reading a dashboard whose apparent precision hides missing events. National enterprise statistics establish the importance of the problem, but they cannot determine the battery, retention capacity or recovery policy of a particular factory. GFT argues for a safe local operating boundary and a durable production record before dependence on remote analytics. This paper turns that position into a site survey, service classes, protected digital loads, local network requirements, transaction boundaries, bounded queues and explicit recovery states. It distinguishes consistent storage from durable commits, protocol delivery from application acceptance, timestamp formatting from clock accuracy, and a restarted process from an authorised return to production. The research combines primary institutional publications, protocol specifications and implementation documentation with clearly labelled illustrative calculations. It develops commissioning experiments and procurement obligations that can expose hidden online dependencies and unexplained record loss. The recommendations connect GFT's Cameroon textile programme to an evidence-gated sequence: measure the site, prove the minimal local record, qualify inspection and synchronisation, rehearse restoration and expand only after the results support it. No site availability, equipment performance or financial return is claimed without operating evidence.

01 Historical enterprise evidence makes outages a material design condition, but site measurement is still required. [S01] [S02]

02 The protected digital load must match the service promised during a production outage. [S03] [S05]

03 A local durable transaction is the essential acknowledgement boundary for the production record. [S13] [S16] [S24]

04 Consistent storage can still lose recent committed work after abrupt power loss. [S16] [S15] [S29]

05 Offline buffers require both a retention limit and spare recovery capacity. [S08] [S18]

06 Protocol delivery guarantees do not prove a downstream business action happened exactly once. [S12] [S19] [S20] [S24]

07 Clock quality and source identity must survive a disconnected restart. [S10] [S11] [S27] [S30] [S31] [S32]

08 Inspection continuity depends on powered inputs, available software and an authorised human fallback. [S03] [S04] [S05]

09 Recovery must reconcile the ledger and restore an approved operating state. [S06] [S17] [S21] [S22]

10 An evidence-producing fault programme should gate scale-up. [S05] [S06] [S26] [S28]

GREEN FUTURE TECH

LIMITED / HONG KONG SAR

01

PART I: SPECIFY THE INTERRUPTION / THE POSITION

# The factory must keep its record when the grid fails

*GFT argues that continuity means retaining a safe operating state and an accountable record through interruption and recovery.*

CHAPTER 01
DECISION TEST

GFT argues that continuity means retaining a safe operating state and an accountable record through interruption and recovery.

A factory should be able to preserve the record of accepted work when electricity or connectivity fails, even when production itself must stop. GFT proposes a local boundary within which the operator can identify work, record an authorised decision and retain its acceptance status without waiting for a remote service. That boundary must have a defined failure response, a responsible person and a demonstrated recovery path. Keeping a computer powered is only a component of the answer. The test is whether people can explain what happened to the material and its record after interruption, including what remains unknown.

1.1

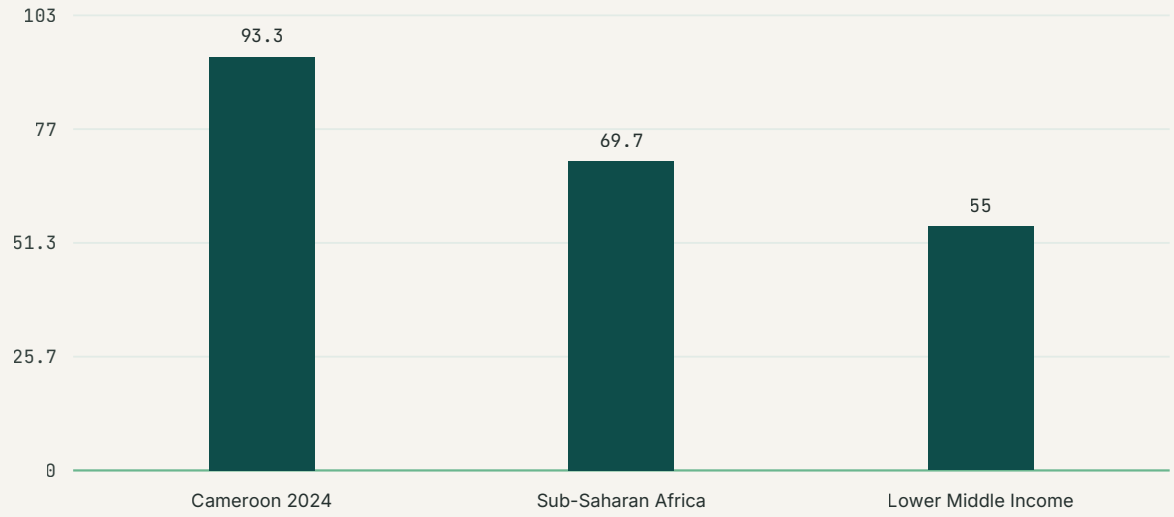
## A weak supply makes continuity part of the production design

The World Bank's Cameroon 2024 Enterprise Survey country profile reports that 93.3 percent of surveyed firms experienced electrical outages and records 10.4 outages in a typical month. These are historical enterprise indicators, not forecasts for a proposed textile plant. They establish a reason to investigate interruption early; neither figure supplies an outage duration, a battery duty or a site's availability. [S01]

The comparison in **Exhibit 1**, p. 14 preserves the published geographical and income groups. The Cameroon observation is labelled by its survey vintage. The comparator groups carry the labels printed in the profile; they are not represented here as contemporaneous textile samples. The chart compares exposure to outages, without attempting to explain its causes or predict the operating conditions of a particular industrial location. [S01]

EXHIBIT 01 / PERCENT OF FIRMS EXPERIENCING ELECTRICAL OUTAGES

### Outages affected most firms in the Cameroon survey




SOURCED [S01] World Bank Cameroon 2024 Country Profile, published 2025, printed page 6, Figure 8. Comparator groups and vintages are as published, not a textile-only or site-level forecast.

The profile covers formal private enterprises across manufacturing and selected services. Its observations therefore cannot be assigned directly to a garment line, an industrial feeder or an individual carrier. GFT's inference is limited: treat interruption as a design question before selecting the digital operating model. The site investigation must then establish the relevant local evidence. [S01]

The supplied sector register describes about 95 percent of Cameroon's cotton fibre as exported raw and about 5 percent as processed locally. It also records a national target of 50 percent local processing by 2030. These inherited observations and targets have unrecorded original publication dates and have not been independently refreshed for this chapter. They establish programme context, with that limitation, rather than a measured opportunity captured by GFT. [R06] [R07]

GFT's proposed connection is operational. Processing material locally requires the operator to account for its conversion into usable output. A work record should connect the identified material to the relevant operation, quality decision and movement. If an interruption severs that connection, the operator needs an explicit hold or reconstruction procedure. Continuity cannot by itself deliver a national processing target, but the production design should explain how accountability survives adverse operating conditions.

The practical design question is therefore more precise than whether the site has backup power. Ask which decisions remain possible during a supply interruption and which observations still exist afterwards. A stopped machine with a clear work state may be easier to restart responsibly than an apparently operating application whose last accepted event is uncertain. This is a proposed decision principle, to be qualified against the process, rather than a comparative claim about measured plant performance.

> **GFT NOTE**
>
> Before approving a digital package, write down the smallest record the supervisor must retain through an interruption. Name the source, acceptance signal, protected equipment, manual alternative and release authority. Require a demonstration against that record. Extend the protected scope only when the proposed production benefit and local operating capability justify it.

GFT recommends asking the plant team to describe an interruption in ordinary working language before drawing the system diagram. Where does the unfinished material wait? Which document tells the next shift whether it can move? Who can resolve a disagreement between a physical count and an accepted event? Those answers make the recording duty inspectable. They also reveal whether the proposed continuity service protects a real decision or merely keeps a screen available.

This approach also changes the project conversation. Electrical engineers define the supply and protective boundary; production leaders define the acceptable material state; quality leaders define release conditions; the information-system supplier defines where acceptance becomes durable. GFT recommends bringing those descriptions together before procurement. An attractive interface cannot resolve a disagreement about which party owns an event when the site loses power.

1.2

## A production stop and a lost record are different failures

NIST's operational technology guidance treats performance, reliability and safety as requirements that need attention alongside security. GFT applies that principle by separating the safe physical response from the survival of information. A powered gateway has no authority to defeat a protective stop, and a production restart needs its own approval conditions. [S03]

Consider an illustrative cutting-to-sewing bundle. Its identifier is assigned before transfer, and a supervisor records its release from cutting on a locally powered terminal. Mains supply then fails before sewing begins. Assume that the approved machine response brings the process to a safe stop while the terminal, its local network path and the record store remain powered. These are explicit scenario assumptions, not descriptions of an installed GFT factory.

In that illustrative situation, the accepted cutting-release event can remain available even though the next production operation has not occurred. The supervisor can locate the bundle, distinguish waiting material from completed sewing and identify the event still awaiting remote acceptance. A management screen may continue to show an older count until communication returns. The screen's age and the bundle's state answer different questions; neither should silently substitute for the other.

GFT proposes that inspection availability be stated separately. If a camera, illumination unit or associated computer is outside the protected circuit, assume its inspection result is unavailable during that failure unless another qualified method exists. A surviving record store cannot create a measurement that the inspection system never made. An operator should see whether a missing result means not inspected, inspection failed, communication pending or inspection held for review.

The proposed state table in **Exhibit 2**, p. 17 makes these distinctions explicit. It assumes that local safety functions are independently approved and that a protected record terminal remains available during mains loss. A LAN failure means the ordinary terminal-to-store path is unavailable; the suggested local-entry alternative applies only where a separately qualified terminal store exists. The dispositions are illustrative requirements to discuss with the operator, not permission to run a process.

EXHIBIT 02 / ILLUSTRATIVE OPERATING DISPOSITION

## Each function needs a separate interruption decision

| Function | Mains loss | LAN loss | WAN loss | Restart |
|---|---|---|---|---|
| Process | Safe stop | Hold | Local control | Hold |
| Inspection | Stop | Hold | Local if qualified | Hold |
| Record | Protected | Local entry | Local capture | Check store |
| Reporting | Mark age | Mark age | Delay upload | Reconcile |

ILLUSTRATIVE [S03] [S05] GFT proposed example, not site approval. Local control assumes independently qualified safety. Local entry assumes a qualified terminal store. Inspection stops on mains loss in this example; no protected camera chain is assumed.

A useful illustrative contrast is a gateway that stays alive while its input sensors go dark. Its health message may still arrive, but that proves only that the gateway could send that message. GFT recommends recording the affected source set and the beginning of the observation gap. Do not convert absent machine pulses into a claim that the machine produced nothing. The operator may be unable to distinguish stopped production from lost observation until local checks resolve the uncertainty.

The reverse distinction also matters in the proposed design. Some work might remain physically possible while the authorised recording route is unavailable. The correct response could be a controlled manual procedure or a production hold, depending on the consequence of losing traceability. Define that choice before an outage. Otherwise the shop floor must invent a policy while trying to restore service, and different shifts may attach different meanings to apparently comparable counts.

GFT recommends defining the meaning of a hold at the work location. Identify the affected material, the reason for withholding release and the role that can clear it. A hold should remain visible when reporting is delayed. Its cancellation should refer to the evidence that resolved the condition. This makes the proposed record useful to the person moving the bundle, rather than only to the person reviewing totals after the shift.

For the illustrative bundle journey, recovery begins with physical identification and the retained event, not with the colour of a network icon. The supervisor checks that cutting release was accepted, sewing remains pending and any manually observed movement has a documented status. Only then should the reporting system incorporate the reconciled events. The lesson is a boundary rule: an application restart, an inspection restart and permission to resume production must each have an explicit meaning.

1.3

## Cloud reachability must not decide whether the line can record work

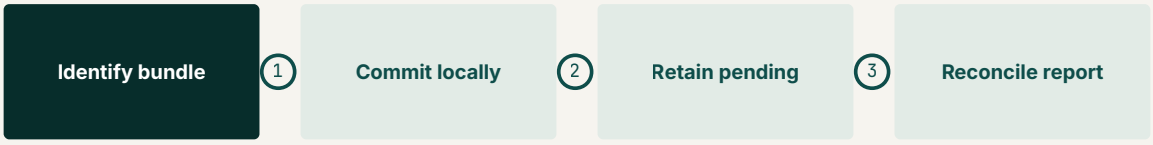
The delay-tolerant networking architecture describes persistent storage as a way to address network interruption. It offers a relevant store-and-forward principle, not a ready-made factory acceptance contract. GFT proposes applying that principle to locally accepted events while leaving the protocol choice and implementation to the engineering design. [S08]

The proposed local record should preserve the identified work item, operation, responsible role, acceptance state and origin of any later correction. Work instructions needed during disconnection should be available through an approved local copy with its revision visible. A supervisor should know whether a record has been accepted locally, remains only on an entry screen, or has also been accepted upstream. These are proposed functional requirements; the chapter does not claim a particular software product provides them.

**Exhibit 3**, p. 18 sketches an illustrative path from the cutting-to-sewing handover to reconciled reporting. The material handover initiates the recording obligation. The local commit and pending-delivery entry form the proposed local responsibility boundary. A later upload is a separate transition. The arrows express a designed workflow, not elapsed time, production output or a completed GFT implementation.

EXHIBIT 03 / ILLUSTRATIVE EVENT JOURNEY

### Local acceptance can precede restored remote reporting




1. Record authorised handover
2. Keep delivery obligation
3. Verify remote acceptance

ILLUSTRATIVE [S08] [S13] GFT proposed cutting-to-sewing handover. Accept the event and pending-delivery intent together locally, retain during disconnection, and reconcile after remote durable acceptance. This is not a certified implementation or a material-flow quantity.

SQLite's atomic-commit documentation explains the all-or-nothing transaction property and sets out hardware assumptions beneath it. Its publication date was not recorded in the inherited outline; the page was reopened for this chapter and is retained as undated. GFT's design inference is to keep the event and its delivery obligation inside an appropriate local transaction. Atomicity alone must not be presented as evidence that every recent acknowledgement survives arbitrary equipment failure. [S13]

The local boundary should therefore be tested while the remote route is deliberately unavailable. GFT recommends checking work identification, authorised entry, revision access, acceptance feedback and subsequent recovery. An application that can display an old screen offline may still require a remote login or lookup before saving. The proposed test should exercise the complete task, including a fresh start under the agreed disconnected conditions, rather than assuming that a previously open session establishes independence.

Device constraints and network constraints are distinct in the IETF terminology for constrained-node networks. GFT draws a related practical distinction: preserving already captured data does not extend the powered life or sensing capability of its source. Design the retained subset explicitly, and treat unavailable sources as gaps rather than inventing observations from neighbouring equipment. [S18]

A proposed local acceptance indicator should also state when the system cannot promise persistence. If a save fails, the operator needs an understandable instruction: retain the identified material, use the approved alternative or seek the named supervisor. Avoid presenting a pending entry as completed work simply because the interface remains responsive. The continuity promise belongs at the moment the person decides whether to move on.

A bounded local store is part of this proposal. Define what happens as capacity becomes scarce, who sees the warning and which work must pause before acceptance promises become false. The detailed queue calculation belongs later in the paper. At this stage the required answer is operational: there must be an agreed disposition when the local system can no longer retain what it is being asked to accept.

GFT also recommends separating incoming operational authority from outgoing historical evidence. A delayed event can remain useful for accounting after its capture time, while a delayed instruction may be unsafe or obsolete for the current work item. The local operator needs an approved rule for instruction validity during disconnection. Treating every queued message as equally fit for later use would confuse preservation of history with permission to direct current production.

1.4

## Continuity is a contract with a measurable boundary

NIST's contingency-planning guidance provides a basis for evaluating system requirements and recovery priorities. GFT proposes turning that planning discipline into distinct commitments for the factory record. Recovery time, allowable committed-record loss, maximum disconnected operation and permissible data age should each be specified by the operator. They answer different questions and should not be compressed into a general promise of uptime. [S05]

For recovery time, state the starting event and the service whose restoration ends the interval. For record loss, state whether the promise covers only durably accepted events or also unacknowledged inputs. For disconnected operation, state the work-

load and the services that remain permitted. For data age, state which decision is restricted when its information becomes too old. These are GFT's proposed contract fields, with numerical limits left for site assessment and risk acceptance.

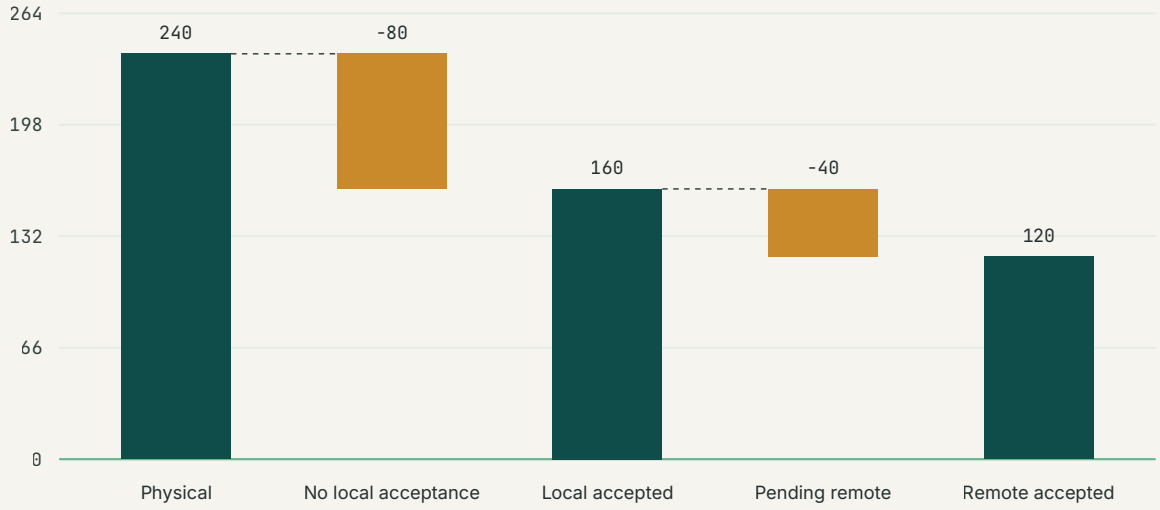
The acknowledgement boundary needs equally plain wording. GFT proposes distinguishing capture at the source, durable local acceptance and durable remote acceptance. A transmission attempt establishes none of those states by itself. TCP provides an ordered reliable byte stream to applications; GFT's application-level inference is that transport delivery does not prove that a receiving business system validated and durably accepted a production event. [S24]

> **WORKED EXAMPLE: THE COUNT DEPENDS ON ITS ACCEPTANCE BOUNDARY**
>
> Assume an illustrative handover of 12 bundles with 20 garments in each. Physical inspection confirms 240 garments. Assume that 8 bundle events have been durably accepted locally, representing 8 multiplied by 20, or 160 garments. The remaining 4 bundles, representing 80 garments, have no local acceptance and must remain an explicitly unresolved recording group. These are constructed inputs for a reconciliation exercise, not observed production.
>
> Assume in the same illustrative case that the remote application has durably accepted only 6 of the locally accepted bundle events, representing 120 garments. The remote backlog is therefore 2 bundles, or 40 garments: the locally accepted 160 less the remotely accepted 120. Successful replay should bring the remote accepted count to 160, provided nothing else changes. It does not establish acceptance of the unresolved 80 garments.
>
> The illustrative control equation is physical count 240 equals remotely accepted 120 plus locally accepted but pending remote 40 plus physical work pending local acceptance 80. These are mutually exclusive groups for this handover. A supervisor must investigate the last group against material identity and authorised source records. A later correction needs its own provenance; repeating the same event must not create additional production.

**Exhibit 4**, p. 21 displays the illustrative reconciliation as changes in reporting scope. The deductions identify excluded records at narrower acceptance boundaries, not scrap or missing physical garments. This distinction is essential: an information difference must not be presented as a material loss without supporting investigation.

EXHIBIT 04 / GARMENTS REPRESENTED AT EACH ILLUSTRATIVE BOUNDARY

## Different acceptance boundaries produce different counts




ILLUSTRATIVE Illustrative assumed inputs: 12 bundles x 20 garments = 240 physically checked; 8 × 20 = 160 locally accepted; 6 × 20 = 120 remotely accepted. Exclusions are 80 awaiting local acceptance and 40 awaiting remote acceptance. Deductions are record-scope differences, not scrap.

GFT recommends preserving the denominator alongside a reported rate. In the illustrative case, a productivity calculation using remotely accepted work would describe a different coverage boundary from a calculation using locally accepted work. Neither should silently become the physical-output rate. The report should state which events are included, what remains pending and whether a later reconciliation changes a previously issued result. This is a reporting rule, not a forecast of the size of any correction.

An acceptance contract should also define the evidence that settles a dispute. GFT recommends retaining a stable event identity, its acceptance status and the relevant failure record so a test can reconcile the same work across boundaries. The supplier should explain the fate of events interrupted before, during and after acceptance. A disappearing progress indicator or an empty upload queue is insufficient evidence of the final business state.

NIST's recovery guidance covers planning, playbooks, testing and improvement. GFT adopts the need to rehearse recovery as a design principle, while recognising that the source addresses cybersecurity recovery. For an ordinary outage, the proposed factory exercise should still finish with a documented release decision and unresolved exceptions. A cyber incident may require containment and investigation before the same restoration path is safe to use. [S06]

1.5

# The programme supplies a starting specification, not an operating proof

The programme design basis supplied to this paper joins feasibility, performance-management system design and digital strategy. Its feasibility work identifies power reliability as an operating risk. Its performance-management design starts with disciplined point-of-work capture. Its digital strategy proposes edge processing and local buffering under constrained connectivity. These are GFT programme proposals read for this chapter, not evidence that a factory has been commissioned or that its record has survived an outage.

GFT recommends preserving that relationship when translating the programme into procurement. The feasibility risk becomes a question about the physical and financial consequences of interruption. The performance-management design defines which work events must remain accountable. The digital design identifies equipment and software dependencies that could support the required local service. Each proposal needs an owner, a demonstrable behaviour and a reason for its priority.

The programme's proposed progression starts with manual capture and bought inspection capability, then introduces forecasting when operating data can support it. GFT's continuity recommendation is to qualify the recording boundary before making a more complex application dependent on it. This does not mean that manual entry is inherently complete or that a bought inspection package is automatically suitable. Both need to be tested against the operator's materials, working conditions and recovery procedure.

For inspection, this paper asks whether a qualified result remains available under the specified failure. The related inspection paper owns evaluation of detection performance and its production consequences. Keeping those questions separate avoids an unjustified inference from a model's nominal accuracy to the availability of a complete inspection service. GFT recommends proving the camera, lighting, computing, instruction and human-review chain together within the intended operating boundary.

For the performance-management system, GFT proposes that the supervisor see accepted output, provisional observations and missing coverage as distinct states. A count should carry enough context to establish its operation, period and acceptance boundary. Where the record was reconstructed after interruption, preserve the reconstruction status. The reporting design should permit a useful correction without rewriting the earlier uncertainty as if it had never existed.

The illustrative bundle reconciliation supplies a modest commissioning objective for that progression. The owner can inspect whether the retained record agrees with the identified material, whether remote reporting catches up without duplicated output and whether the unresolved group remains visible. Passing such an exercise would establish only the behaviour tested under its named conditions. It would not qualify all products, shifts, equipment failures or future software versions.

GFT recommends assigning these records to operating roles rather than a supplier account alone. The shift supervisor should understand acceptance feedback and holds. The quality lead should understand which evidence authorises release. The

maintenance and information-system leads should understand the protected dependency chain and restoration route. Handover should require those roles to explain the agreed procedure in their own terms, with any unresolved disagreement recorded before expansion.

The programme should consequently be described in terms of proposed capability until operating evidence exists. A specification says what is wanted; a demonstration shows a bounded behaviour; a commissioning record identifies what was accepted; ongoing observations show what happened in service. GFT recommends retaining those distinctions in internal approvals and external research. They prevent a design recommendation from acquiring the appearance of a verified operating result through repetition.

1.6

## The case for resilience must survive its counterargument

The strongest counterargument is that a simpler manual procedure may be sufficient. For a low-criticality activity, an approved paper form, visible material identifier and scheduled reconciliation could meet the operator's requirement. GFT recommends testing that alternative before adding protected digital dependencies. The comparison should include the consequence of delay, the staff effort required to reconcile and the credibility of the resulting record, rather than assuming that digital continuity always has greater value.

This is consistent with using contingency planning to establish priorities rather than applying identical protection to every system. NIST's guidance supports evaluating information-system and operational requirements; the decision to use a manual fallback here is GFT's proposed application. The activity owner must agree what the fallback preserves and what it deliberately postpones. [S05]

An adequate proposed manual procedure needs available forms, identifiers that match the material, a responsible custodian and a rule for later entry. GFT recommends recording who observed the event, who entered the reconstructed record and who resolved any disagreement. A handwritten time should retain its actual confidence rather than being promoted into an exact automated timestamp. Where a safety or quality decision requires unavailable evidence, the procedure must specify a hold.

The digital alternative must face the same scrutiny. Its proposed benefit could be faster local reconciliation, clearer acceptance feedback or reduced dependence on later transcription. Those are hypotheses until measured. GFT recommends comparing the alternatives during an exercise with the actual shift roles and workload. Include recovery work, training, replacement and exception handling. Exclude any production or financial benefit that the comparison cannot attribute to the continuity design.

Remote services still have a suitable place in the proposed operating model. Non-urgent reporting and remote optimisation can wait where the operator has agreed a freshness limit and a valid local procedure. Network constraints can vary independ-

ently of device resources, as the IETF terminology explains; the contract should therefore distinguish unavailable communication from unavailable local computation. [S18]

GFT recommends marking delayed information where the user makes the decision, including exported summaries where relevant. A remote manager should know whether a displayed total is current, reconciled through an earlier boundary or missing a named source. The appropriate response might be to defer a planning change or contact the local supervisor. It should not depend on a manager inferring an outage from an unusually flat chart.

The proposed comparison should also consider a shift change during disconnection. A fallback that depends entirely on the outgoing supervisor remembering which bundles moved has not preserved an accountable record. Require the incoming supervisor to locate the material and interpret the retained acceptance state using the agreed procedure. This exercise challenges both manual and digital alternatives without presuming which will be preferable at the site.

A staged decision is therefore preferable to a general commitment to maximise digital autonomy. Protect the activities whose interruption consequences justify the duty, retain an authorised simpler route where it meets the requirement, and permit remote work to wait where delay is acceptable. The unresolved question is local consequence, not whether a cloud service or a manual form is intrinsically superior. The later economic analysis should price those consequences using site evidence.

1.7

## This paper owns failure behaviour across the power and data boundary

This paper's proposed scope is the behaviour of production information through loss, degraded operation and recovery. It asks which inputs remain available, which events are accepted, how uncertainty is shown and how responsibility returns to normal operation. It also defines what observations would support those claims. The scope reaches across the electrical and information boundary because a promised local service depends on both.

The paper on resource efficiency owns wider energy systems and generation economics. The governance paper owns policy, data stewardship and model lifecycle responsibilities. The connected-factory paper owns the fuller integration reference architecture. Here those subjects appear only where they create an interruption dependency or a release condition. That boundary should help the operator combine the papers without receiving competing specifications for the same equipment or approval.

The evidence hierarchy must remain visible. The opened World Bank profile provides historical enterprise context. Opened NIST publications support planning and the need to respect operational requirements. Opened protocol and implementation documentation supplies defined mechanisms and limitations. GFT's recommendations

translate those materials into proposed factory decisions. The illustrative exercises explain reasoning with declared inputs; they are not site measurements, supplier guarantees or performance results. [S01] [S03] [S05]

The inherited cotton-processing entries retain their unknown original dates. The undated SQLite page is identified as such, with an access date separate from publication. No source has been supplied for this chapter that establishes the proposed site's outage distribution, protected-load consumption, carrier behaviour or application recovery time. Accordingly, this chapter makes no site availability claim and selects no equipment rating. [R06] [R07] [S13]

GFT recommends leaving this opening chapter with a short decision record: identify the material and work event to preserve; name the local acceptance boundary; describe permitted degraded work; name the physical and information holds; assign restoration and release responsibility. Then ask the site-survey and engineering chapters to supply the evidence and limits needed to make that record testable. A missing answer should remain an explicit requirement rather than a reassuring adjective.

The proposed decision record should identify its exclusions in the same place as its promise. For example, preserving accepted bundle events may leave detailed inspection images outside the protected scope. The operator must decide whether those images are necessary to support release or only later analysis. Naming the exclusion prevents a narrow success in event retention from being mistaken for complete preservation of every production artefact.

That decision record should be revisited when the workload, source equipment or operating procedure changes. A new inspection stream or remote approval step may alter what must survive disconnection. GFT proposes treating such changes as reasons to review the continuity boundary and repeat the relevant exercise. Confidence should travel with a defined configuration and operating duty, rather than attach permanently to the label of a system.

The recommendations are research and proposed engineering methods, not investment, legal or tax advice. Equipment choice, protective design and permission to operate require site-specific professional responsibility. The intended outcome is an accountable basis for those decisions: a factory can stop safely, retain what it actually accepted, expose what it could not observe and demonstrate how the record becomes trustworthy again.

*The factory may lose its supply; it must not lose the meaning of its work.*

**CHAPTER SOURCE NOTES**

S01 World Bank. 2025.

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

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

S03 NIST. 2023.

S05 NIST. 2010.

S08 IETF / RFC Editor. 2007.

S13 SQLite. n.d..

S18 IETF / RFC Editor. 2014.

S24 IETF / RFC Editor. 2022.

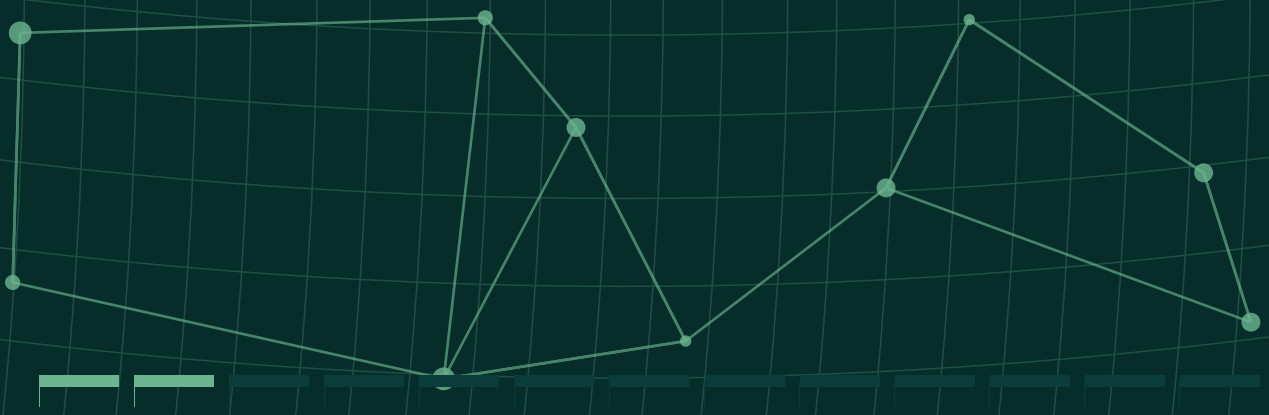
S06 NIST. 2016.

02

PART I: SPECIFY THE INTERRUPTION / SURVEY THE FAILURE

# Site measurements must replace national averages

*A useful design envelope records the duration, coincidence and recovery of power and network failures at the intended site.*



CHAPTER 02
DECISION TEST

A useful design envelope records the duration, coincidence and recovery of power and network failures at the intended site.

National enterprise surveys establish that electrical instability is a material operating constraint across Cameroon, where 93.3 percent of surveyed formal firms experienced power outages and reported 10.4 outages in a typical month. [S01] Yet an average monthly outage count cannot size an uninterruptible power supply (UPS), dimension an offline event buffer or verify whether a candidate backhaul link survives when utility mains drop. [S01] [S05] GFT argues that before an industrial sponsor selects edge hardware, battery autonomy or synchronisation intervals for a textile plant, the project team must replace macroeconomic averages with a time-synchronised site survey that logs the duration, coincidence and recovery lag of electrical and network failures at the actual plant boundary. [S03] [S05] Until those feeder-level and carrier-level observations are collected under a documented protocol, every autonomy duration, queue budget and availability figure remains a conditional design scenario rather than an observed site fact. [S02] [S05]

2.1

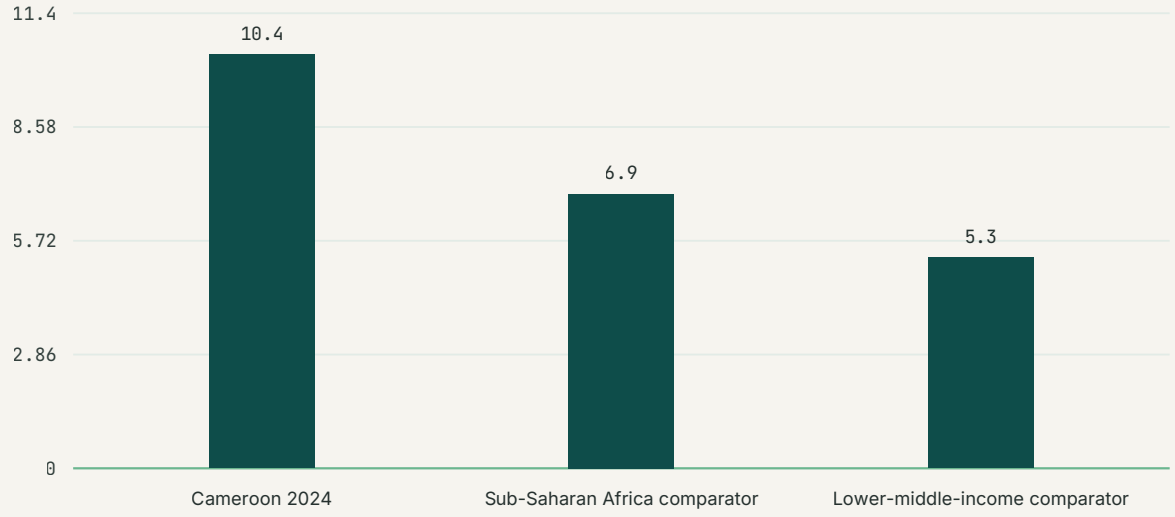
## Country statistics establish context and cannot size a site UPS

Macroeconomic business-environment indicators justify why digital resilience belongs in the capital-expenditure scope of a Cameroonian factory, while simultaneously demonstrating why aggregate survey numbers cannot serve as engineering inputs. [S01] [S02] In the World Bank Enterprise Survey for Cameroon 2024, field teams interviewed business owners and top managers across 615 formal private non-agricultural firms employing at least five workers between June 2024 and February 2025. [S01] Across that national sample, 93.3 percent of firms reported experiencing electrical outages, compared with 69.7 percent for the Sub-Saharan Africa comparator group and 55.0 percent for the lower-middle-income comparator group. [S01] When respondents selected the single biggest obstacle to their business from 15 constraints, 12.1 percent of all firms and 25.3 percent of large firms (100 or more employees) identified electricity as their primary obstacle, ahead of political instability at 19.7 percent and access to finance at 11.3 percent. [S01]

The survey also reports a second, distinct reliability indicator: the number of electrical outages experienced in a typical month, presented in **Exhibit 5**, p. 29. [S01] Surveyed firms in Cameroon reported 10.4 electrical outages in a typical month, against 6.9 in the Sub-Saharan Africa comparator and 5.3 in the lower-middle-income comparator. [S01] By establishment size, outage incidence reached 95.7 percent among small firms (5 to 19 employees), 90.6 percent among medium firms (20 to 99 employees) and 77.9 percent among large firms (100 or more employees), whereas typical monthly frequency stood at 10.6 outages for small firms, 9.4 for medium firms and 11.8 for large firms. [S01] Firms applying for an electrical connection waited an average of 35.5 days across all surveyed establishments. [S01]

EXHIBIT 05 / OUTAGES IN A TYPICAL MONTH

## Outage frequency is a separate survey measure




SOURCED [S01] World Bank Cameroon 2024 Country Profile, printed page 6 Figure 8 and page 11 Appendix. Survey population covers 615 formal non-agricultural private firms (5+ employees). Frequency does not measure outage duration or feeder tail risk.

As **Exhibit 5**, p. 29 shows, the published frequency of 10.4 outages per month is nearly double the lower-middle-income benchmark of 5.3 outages per month. [S01] However, an engineering team must separate what **Exhibit 5**, p. 29 proves from what it cannot establish. [S01] First, a monthly count of 10.4 events carries no information about the distribution of event durations. [S01] Ten interruptions lasting three minutes each impose 30 minutes of cumulative monthly downtime in an illustrative calculation and require only short ride-through storage, whereas ten interruptions containing two eight-hour feeder faults impose a completely different battery and data-retention duty. [S05] Second, the Country Profile headline chart reports incidence and monthly count on separate axes and does not publish a site-level duration percentile, phase-imbalance count or voltage-sag profile. [S01]

### 2.2 SURVEY STRATIFICATION PREVENTS DIRECT EXTRAPOLATION TO A TEXTILE FEEDER

The World Bank Enterprise Surveys methodology explains why national point estimates cannot be assigned to a specific industrial plot or textile mill. [S02] Under the global methodology applied since 2005 to 2006 and updated in August 2022 to use International Standard Industrial Classification (ISIC) Revision 4 codes, the formal-sector survey samples registered companies with five or more workers and at least 1 percent private ownership across manufacturing (ISIC codes 10 to 33), construction, retail, wholesale, hospitality, transport, storage, communications, information technology, and professional, scientific and technical activities. [S02] Agriculture, fishing, mining, public utilities, financial intermediation, public administration, education, health and 100 percent state-owned enterprises are excluded. [S01] [S02]

In the Cameroon 2024 sample of 615 establishments, only 161 firms belonged to manufacturing (49 food-manufacturing and 112 other-manufacturing firms), alongside 126 retail and 328 other-services firms, stratified across the Centre (207 firms), Littoral (199), North (146) and West (63) regions. [S01] Furthermore, because the

World Bank tracks panel enterprises across survey waves, sampling weights are adjusted to account for altered probabilities of inclusion in the sample frame. [S02] Even within manufacturing, outage incidence reached 100.0 percent in food manufacturing and 87.3 percent in other manufacturing, with typical monthly counts of 11.9 and 10.4 outages respectively. [S01]

Consequently, the 10.4 monthly outage figure is a sampling-weighted, multi-sector national mean, not a textile-sector measurement or a feeder log for Douala, Dibamba or Garoua. [S01] [S02] While Cameroon's cotton-to-garment opportunity rests on moving local fibre processing from roughly 5 percent (with about 95 percent exported raw) toward the national 50 percent target by 2030, and while historic reports on CICAM cited July 2023 capacity utilisation of 33.7 percent at Garoua, 48.9 percent at CICAM 1 and 10.83 percent at CICAM 2, neither source provides feeder-specific electrical or telecoms telemetry. [R06] [R07] [R12] Treating the national 10.4 monthly average as a factory design envelope would conflate a macroeconomic statistic with a site engineering specification. [S01] [S02]

2.3

## Power quality needs an event log as well as an energy meter

A conventional utility billing meter or monthly kilowatt-hour totaliser measures cumulative energy transfer over a billing cycle, which is useful for tariff reconciliation and resource accounting but blind to the transient disturbances that disrupt digital systems. [S03] [S05] Operational technology (OT) and edge-computing equipment, including programmable logic controllers, machine-vision cameras, industrial switches, time servers and local database hosts, fail on timescales of milliseconds to seconds when exposed to voltage sags, single-phasing, frequency excursions or slow contactor transfer intervals. [S03] A plant energy meter can record normal daily consumption while a 400-millisecond illustrative voltage dip or an uncoordinated generator changeover reboots an unprotected switch and severs a transaction in flight. [S03] [S13]

To specify the continuity boundary required by NIST Special Publication 800-82 Revision 3 for operational technology reliability and safety and by NIST Special Publication 800-34 Revision 1 for information-system contingency planning, the site baseline must capture a discrete power-event log rather than a simple energy total. [S03] [S05]

METHOD

**Proposed site power-event logging boundary and schema.** GFT proposes that the pre-design site survey instrument the incoming low-voltage distribution board and the candidate digital sub-distribution circuit to record eight mandatory fields for every electrical disturbance: [S03] [S05] 1. **Event identifier (`event_id`):** A locally generated 128-bit Universally Unique Identifier (UUID Version 7 or Version 4) assigned by the power-quality recorder so that power events join deterministically with network and application logs without a central allocator. [S31] 2. **Boundary timestamps (`t_start_utc`, `t_end_utc`):** Event onset and restoration times formatted under RFC 3339 with an explicit offset to Coordinated Universal Time (UTC) and sub-second resolution. [S32] 3. **Recorder clock state (`clock_quality`):** Metadata stating whether the recorder's clock was synchronised via Network Time Protocol (NTP), operating in local holdover, or unsynchronised after a cold boot. [S10] [S32] 4. **Phase and waveform status (`phase_state`):** Individual status for phases L1, L2 and L3, distinguishing three-phase blackout, single-phase loss, phase imbalance, brownout (sustained under-voltage), swell (over-voltage) and frequency deviation. [S03] 5. **Protective device action (`trip_status`):** Indication of whether an upstream utility breaker, site incoming breaker, residual-current device or surge-protection path tripped. [S03] 6. **Transfer interval (`transfer_ms`):** The measured interruption gap in milliseconds between loss of utility mains and assumption of load by the standby generator, automatic transfer switch (ATS) or UPS. [S03] [S05] 7. **Affected circuit boundary (`circuit_id`):** Tag naming whether the event affected the main motive-power bus, auxiliary bus or protected digital continuity circuit. [S03] 8. **Observer survival flag (`recorder_ride_through`):** A boolean and voltage log confirming whether the recording instrument remained continuously powered or suffered a reset that right-censored the event record. [S05] [S22]

Distinguishing brownouts, single-phasing and transfer intervals from complete blackouts is essential for textile operations. [S03] During changeover from utility mains to a site diesel generator through an open-transition automatic transfer switch, the motive-power bus experiences a dead interval of several seconds to several minutes while the generator starts and stabilises voltage and frequency. [S03] [S05] If an edge gateway or network switch sits upstream of an online UPS, that open-transition transfer interval reboots the digital node just like an eight-hour illustrative grid blackout. [S05] Moreover, although the Linux `ext4` filesystem enables write barriers by default (`barrier=1`) to enforce on-disk journal ordering and defaults to a five-second commit interval (`commit=5`), a sudden loss of DC power on an unbacked host can still lose up to five seconds of un-fsynced metadata, and SQLite databases in write-ahead logging (WAL) mode with `PRAGMA synchronous=NORMAL` can roll back recently committed transactions after a power loss. [S16] [S29]

### 2.4 ELECTRICAL ENGINEERING AUTHORITY REMAINS DISTINCT FROM DIGITAL LOGGING

A site power-event log defines the disturbance environment that the digital system must survive, but it does not replace qualified electrical engineering design. [S03] GFT's recording method specifies what telemetry IT and OT architects need to classify services and size digital backup autonomy; it does not prescribe earthing topologies, short-circuit ratings, arc-flash boundaries or protective relay settings. [S03] The project sponsor must commission a licensed electrical engineer to inspect the site transformer, verify neutral-to-earth bonding, specify overcurrent and surge protection, and certify that any temporary power-quality logger is installed safely. [S03]

2.5

## A reachable modem does not establish usable connectivity

Just as a monthly energy meter hides power transients, a router link light or a simple Internet Control Message Protocol (ICMP) echo test hides the network failure modes that stall industrial data synchronisation. [S18] [S24] In practice, an edge outbox attempting to forward committed production events to a central repository depends on a multi-layer protocol chain, and failure at any intermediate layer renders the backhaul unusable for application recovery. [S12] [S23] [S24]

GFT decomposes backhaul usability into six sequential verification layers that the site survey probe must measure and log independently during scheduled production shifts: [S05] [S18]

1. **Local physical and data-link readiness:** The edge computer's local area network (LAN) interface, access switch port and site router interface must maintain carrier link without excessive frame errors or port flapping. [S03] [S18]
2. **Domain Name System (DNS) resolution:** DNS is a query-response protocol associating domain names with IP addresses. [S0204] Under RFC 9520, which updates RFC 2308 to require negative caching for DNS resolution failures, resolvers cache resolution failures to reduce latency and authoritative server load. [S0201] Consequently, even after a physical cellular or fibre link recovers, cached negative responses or stale forwarder states on the site router can prevent the edge client from resolving the ingestion hostname until the negative cache timer expires. [S0201] [S0204]
3. **Transport-layer connection establishment:** Under RFC 9293, the Transmission Control Protocol (TCP) provides a reliable, in-order byte-stream service using a three-way handshake (`SYN`, `SYN-ACK`, `ACK`). [S24] A bare TCP connection proves only that three small control segments traversed the path; it does not prove bulk payload capacity or remote application health. [S24]
4. **Cryptographic session establishment:** Industrial ingestion endpoints require Transport Layer Security (TLS) Version 1.3 as specified in RFC 9846 (published July 2026, obsoleting RFC 8446). [S23] If high packet loss fragments the certificate chain or if an unbacked local clock has reset to an epoch date after a power cut, certificate validation fails and the TLS 1.3 handshake aborts even though raw TCP packets pass. [S23] [S32]
5. **Sustained payload goodput under load:** RFC 7228 notes that constrained and challenged networks exhibit lossy channels, limited and unpredictable bandwidth, and severe penalties for larger packets. [S18] Under the RFC 6349 framework for TCP throughput testing, end-to-end TCP throughput depends on round-trip time, packet loss rate, maximum segment size and window limits. [S0202] A link that passes a 64-byte illustrative keepalive probe can suffer retransmission timeouts when transmitting a 500-byte illustrative event batch or an inspection image archive. [S18] [S0202]
6. **Application-level durable commit acknowledgement:** Under HTTP semantics (RFC 9110) and atomic commit rules (SQLite atomic commit), usable connectivity requires that the remote ingestion service validate the payload, commit the record

to non-volatile storage, and return an explicit application-level receipt (such as HTTP `200 OK` or `201 Created` carrying the committed event ID) across the open TLS session before disconnection. [S12] [S13] [S23]

### 2.6 ACTIVE PROBING MUST RESPECT SHIFT CALENDARS AND RETAIN MONITORING GAPS

To measure usable connectivity without inventing carrier performance figures, the site survey should deploy a lightweight software probe on the candidate site network that executes this six-step chain against a test endpoint at a regular cadence during agreed production hours, testing representative payload sizes under RFC 6349 principles. [S05] [S0202] [S0203] Because no site carrier trace has been collected for a GFT textile plant in Cameroon as of the edition date, this paper publishes zero Cameroon bandwidth, latency or packet-loss claims. [S18]

Three rules govern the connectivity probe log. [S10] [S22] [S0203] First, the probe must record its own state using persistent local logging such as `systemd-journald` with `Storage=persistent` so records survive reboots. [S22] Second, the probe host must log clock synchronisation status and offset under RFC 5905 and RFC 8633 (`chronyd` with `rtcsync` or `rtcfile`). [S10] [S11] [S30] Third, whenever the probe host is unpowered or isolated from the LAN, the interval must be flagged explicitly as a monitoring gap (`observer_offline`), never interpolated as available or failed. [S05] [S0203]

2.7

## The joint failure history matters more than separate averages

Even when a project team measures site electricity and site connectivity with proper instruments, analysing the two logs in isolation produces a dangerous sizing error. [S03] [S05] Reliability engineers who multiply a standalone mains-availability fraction by a standalone WAN-availability fraction implicitly assume that electrical outages and telecommunications outages are statistically independent. [S05] In industrial zones, that assumption fails because power feeders and telecoms networks share physical and electrical dependencies: [S03] [S18]

- **Shared utility feeder at the local access tower or exchange:** When a medium-voltage feeder trips across an industrial district, cellular base stations and fibre aggregation cabinets serving that district lose mains power at the same instant as the factory. [S03] If a tower battery exhausts after 30 to 60 minutes in an illustrative scenario, WAN connectivity fails midway through the factory's mains outage. [S05] [S18]
- **Unprotected customer-premises equipment inside the plant:** Leaving the ISP optical network terminal (ONT), microwave power-over-Ethernet injector or wiring-closet switch on an unprotected circuit turns every mains outage into an immediate 100 percent WAN outage. [S03] [S05]
- **Shared civil ducts and overhead poles:** Vehicle impacts, excavations, storms or fires along plant approach corridors frequently sever overhead power lines and aerial fibre cables on the same poles simultaneously. [S03] [S05]

- **Post-outage congestion and reconnection storms:** Simultaneous sector restoration causes thousands of nearby devices to re-authenticate on local cellular cells at once, degrading goodput just as the factory's outbox attempts to drain its backlog. [S08] [S18]

To make these interactions visible, the site baseline must cross-tabulate power state and network state over identical time slices to form a joint outage envelope, illustrated in **Exhibit 6**, p. 34. [S05] [S18]

EXHIBIT 06 / ILLUSTRATIVE SHIFT-HOURS OUT OF 160

## Coincident failures reveal the real design envelope

| | Mains OK / WAN OK | Mains OK / WAN Down | Mains Down / WAN OK | Mains Down / WAN Down |
|---|---|---|---|---|
| Unbacked CPE router | 136 | 4 | 0 | 20 |
| Shared-feeder tower battery | 136 | 4 | 8 | 12 |
| Independent dual-path target | 138 | 2 | 18 | 2 |

Cell scale: 0 to 138

ILLUSTRATIVE [S05] [S18] Illustrative synthetic 160-hour (four-week, single-shift) site log demonstrating the joint-state matrix method. Values are assumed scenario counts that sum to 160 hours per row; they are not observed Cameroon statistics.

In **Exhibit 6**, p. 34, three illustrative site topologies are compared across a synthetic 160-hour production window in which utility mains are unavailable for 20 hours (12.5 percent of shift hours, an illustrative scenario input). [S05] [S18] In the first row (`Unbacked CPE router`), because the site modem loses power whenever utility mains fail, the system spends zero hours in `Mains Down / WAN OK` and 20 hours in the joint failure state (`Mains Down / WAN Down`), bringing total WAN unavailability to 24 hours out of 160. [S03] [S05] In the second row (`Shared-feeder tower battery`), protecting the plant router on a local UPS yields 8 hours of connected operation during short cuts, but the external tower shares the utility feeder and exhausts its battery during longer cuts, leaving 12 hours in `Mains Down / WAN Down`. [S05] [S18] Only in the third row (`Independent dual-path target`), where the router is protected and secondary backhaul avoids the local feeder failure mode, does the plant remain connected for 18 of the 20 mains-outage hours. [S05] **Exhibit 6**, p. 34 is methodological: it shows why separate marginal averages hide the exact cell (`Mains Down / WAN Down`) that governs simultaneous battery and storage duty. [S02] [S05]

WORKED EXAMPLE

**Illustrative calculation: independent marginals versus shared-feeder coincidence.** Consider an illustrative 30-day site baseline covering 480 scheduled production hours (two eight-hour shifts per day). Suppose the power logger records 48.0 hours of mains outage during scheduled shifts (a 10.0 percent illustrative unavailability fraction), and the connectivity probe records 36.0 hours of unusable WAN during the same 480 shift hours (a 7.5 percent illustrative unavailability fraction). [S05]

If an analyst treats mains and WAN availability as statistically independent, expected joint failure duration $T_{\text{joint, indep}}$ over the 480 illustrative shift hours is: [S05] $$T_{\text{joint, indep}} = 480\text{ h} \times 0.100 \times 0.075 = 3.6\text{ hours}$$ Under that naive assumption, the engineer expects only 3.6 hours per month of simultaneous power and network loss, assuming that during the remaining $48.0 - 3.6 = 44.4\text{ hours}$ of mains cuts the edge node can stream events live over a working WAN. [S05]

Now inspect the actual joint log under an illustrative shared-feeder condition where the local telecoms access node loses mains power with the factory and its battery lasts 45 minutes (0.75 hours). [S03] [S18] Suppose the 48.0 mains-outage hours comprise eight short interruptions averaging 0.5 hours each (4.0 hours total, during which the tower battery holds) and eleven extended feeder outages averaging 4.0 hours each (44.0 hours total, during each of which the tower fails after 0.75 hours). The actual coincident `Mains Down / WAN Down` duration is: $$T_{\text{joint, actual}} = 11 \times (4.0\text{ h} - 0.75\text{ h}) = 35.75\text{ hours}$$ Here, 35.75 of the 36.0 WAN-outage hours occur directly inside mains outages, nearly ten times the 3.6 hours predicted by independent marginals. [S05] The local edge store cannot rely on off-site streaming during extended power cuts and must hold the full production backlog in local non-volatile storage. [S08] [S13] This worked example uses assumed illustrative inputs to demonstrate the accounting method, not measured GFT factory data. [S05]

Each of the four cells in **Exhibit 6**, p. 34 places a distinct demand on the continuity architecture. [S03] [S05] `Mains OK / WAN OK` represents normal connected operation. [S08] `Mains OK / WAN Down` represents isolated operation while machinery runs at full speed on utility power, generating production events at peak shift rate and governing maximum local queue intake. [S03] [S08] `Mains Down / WAN OK` allows a battery-backed digital node to flush final transactions and send a power-loss alert before shutting down or switching to generator power. [S05] [S14] `Mains Down / WAN Down` represents compound isolation, where all point-of-work capture must commit atomically to local storage without external DNS, NTP, identity or cloud dependencies. [S08] [S10] [S13]

2.8

## Tail durations decide retention and autonomy requirements

Averages conceal the extremes that break bounded queues and exhaust batteries. [S05] [S08] Whether a local SQLite WAL partition fills up or a protected digital circuit brown-outs depends on the upper tail of the outage-duration distribution: the longest single offline interval and the longest cluster of back-to-back outages whose recovery gap is too short to recharge batteries or drain queues. [S05] [S08] [S16]

To characterise outage durations without inventing parametric curves, the site survey should compile observed durations into an empirical distribution function (EDF) as defined in RFC 2330. [S0203] For $N$ scalar duration measurements, $F(x)$ gives the fractional proportion of measurements less than or equal to $x$, and the $p$-th percentile is the smallest observed value $x$ for which $F(x) \ge p$ percent. [S0203]

### 2.9 SAMPLE SIZE LIMITS THE PERCENTILES THAT A BASELINE CAN CLAIM

RFC 2330 states a strict mathematical rule for empirical percentiles: any percentile summary must report the sample size $N$, because any percentile difference finer than $1/N$ is below the resolution of the sample. [S0203] If a four-week site survey observes $N = 10$ electrical outages (near the World Bank Cameroon 2024 typical monthly count of 10.4 outages), the sample resolution is $1/10 = 10$ percentage points. [S01] [S0203] On 10 events, claiming to distinguish a 95th percentile from a 99th percentile is an interpolation artefact, and even a 90-day baseline with an illustrative $N = 30$ outages ($1/30 \approx 3.33$ percent resolution) cannot resolve a true 99th percentile. [S0203]

To prevent false statistical precision when sizing autonomy and retention in Chapters 4 and 7, GFT requires the site duration table to report four explicit items: [S02] [S05] [S0203]

1. **Exact event count ($N$) and observation window:** Discrete disturbances observed, calendar dates and total monitored hours. [S0203]
2. **Shift-calendar stratification:** Separation of disturbances during active production shifts from unstaffed nights or weekends, because an eight-hour illustrative Sunday outage when lines are stopped adds zero production events to the outbox. [S05] [S08]
3. **Right-censored observations:** Any outage still underway at survey close or truncated when a recorder battery dies after an illustrative four hours must be logged as right-censored (`> 4.0 h, censored`), never dropped or recorded as a completed 4.0-hour event. [S05]
4. **Inter-arrival and recharge gaps:** Elapsed operating time between consecutive outages. [S05] Two three-hour illustrative mains cuts separated by 20 minutes of restored power act as a single compound drain on a UPS battery requiring several hours to recharge, and can starve SQLite WAL checkpointing if an uplink reader holds an open read cursor across back-to-back WAN failures (where SQLite defaults to a 1,000-page automatic passive checkpoint threshold). [S05] [S08] [S14]

Because pre-commissioning samples are small, the design offline interval (such as the illustrative two-hour protected digital autonomy in Chapter 4 and eight-hour event retention in Chapter 7) must be set as an explicit management risk-acceptance boundary informed by observed maxima, censored events and operator interviews, never extrapolated from the World Bank's 10.4 monthly frequency figure. [S01] [S05]

2.10

# Restart and backlog recovery belong in the baseline

An external utility or telecommunications outage is shorter than the operational disruption it causes. [S06] Between physical supply restoration and a reconciled remote ledger lies a sequence of hardware boot, name resolution, cryptographic handshake, transaction verification and queue-drain delays. [S13] [S23] [S24] **Exhibit 7**, p. 37 illustrates why usable remote synchronisation lags the physical return of mains voltage. [S06] [S23] [S24]

EXHIBIT 07 / ILLUSTRATIVE RELATIVE MINUTES FROM MAINS RETURN

## Connectivity recovery can lag restoration of mains

0

**Mains voltage returns**

Utility feeder re-energises at t = 0 min; ATS waits for voltage stability.

5

**Link up and DNS recovery**

Carrier link establishes; negative DNS cache entries expire and hostname resolves.

8

**First remote durable commit**

Cloud endpoint accepts first replayed batch and returns committed-ID receipt.



2

**Stable power and CPE boot**

Transfer switch closes; unbacked ISP modem and access switch cold-boot.

7

**TCP and TLS 1.3 handshake**

Edge client completes TCP 3WHS, TLS 1.3 certificate validation and auth.

45

**Backlog clearance complete**

Outbox drains accumulated offline events while live shift arrivals continue.

ILLUSTRATIVE [S06] [S23] [S24] Illustrative event sequence in relative minutes from physical voltage return (t = 0). Timings are hypothetical scenario values showing protocol and operational ordering under RFC 9293 (TCP), RFC 9846 (TLS 1.3) and NIST SP 800-184; they are not measured site timings.

As shown in **Exhibit 7**, p. 37, six milestones separate grid fault clearance from full recovery in an illustrative sequence: [S06] [S23] [S24]

1. **External supply return ($t_0$, 0 minutes in the illustrative timeline):** Feeder voltage returns while transfer switches verify stability. [S03] [S05]
2. **Stable power confirmation and cold boot ($t_1$, illustrative 2 minutes):** Unbacked routers or exhausted edge nodes cold-boot, replay `ext4` journals and start `systemd.service` units. [S21] [S22] [S29]
3. **First accepted local event ($t_{\text{local}}$):** Local capture resumes once storage mounts and SQLite recovers any hot journal or WAL. [S13] [S21]

4. **Link, DNS and cryptographic session recovery ($t_2$ to $t_3$, illustrative 5 to 7 minutes):** Carrier bearer link establishes, RFC 9520 negative DNS cache entries expire, and TCP (RFC 9293) and TLS 1.3 (RFC 9846) handshakes complete. [S23] [S24] [S0201]
5. **First verified remote durable commit ($t_4$, illustrative 8 minutes):** The remote endpoint commits the first replayed batch and returns an HTTP application receipt. [S12] [S13]
6. **Oldest pending record clearance ($t_5$, illustrative 45 minutes):** Because the outbox must drain accumulated backlog $B = r_{\text{in}} \times D_{\text{outage}}$ while live shift arrivals continue at rate $r_{\text{in}}$, clearance takes $B / (r_{\text{up}} - r_{\text{in}})$ and finishes only when net upload goodput $r_{\text{up}} - r_{\text{in}}$ clears the queue. [S06] [S08] [S0202]

### 2.11 PARTITIONING OUTAGE COMPONENTS PREVENTS ECONOMIC DOUBLE-COUNTING

Recording all six timestamps allows the project team to partition every incident into non-overlapping intervals for Chapters 4, 7 and 12: [S05] [S06]

- **External fault duration ($t_0 - t_{\text{onset}}$):** Carrier or utility failure at the site demarcation point. [S05]
- **Local capture outage ($\max(0, t_{\text{local}} - t_{\text{local_stop}})$):** Interval when local capture was unpowered or unready; zero when battery ride-through succeeds. [S03] [S05] [S13]
- **Equipment restart and handshake delay ($t_4 - t_0$):** Post-restoration lag to re-establish an authenticated commit channel. [S06] [S23] [S24]
- **Queue-drain and staleness recovery ($t_5 - t_4$):** Interval when the remote dashboard is connected but lagging real time while backlog drains. [S06] [S08]

Separating these intervals ensures that Chapter 12 never attributes motive-power stoppage to digital communication lag and never double-counts the minutes between $t_0$ and $t_5$. [S03] [S05]

2.12

## The site dossier must expose what remains unknown

Before freezing service classes in Chapter 3 or sizing UPS batteries in Chapter 4, the engineering team should compile a formal Site Continuity Dossier covering six domains: [S03] [S05]

1. **Single-line electrical diagram and feeder provenance:** Utility voltage, transformer rating, generator capacity, ATS transition type, earthing, and the breaker panel feeding the protected digital load. [S03]
2. **Telecommunications demarcation and physical routing:** Entry duct or pole route from plant boundary to cabinet, base station or fibre PoP identity, and CPE power protection status. [S03] [S05] [S18]

3. **Production shift calendar and point-of-work density:** Planned shifts per day, working days per week, seasonal peaks, and active capture stations per hall. [S03] [S05]
4. **Candidate protected digital load inventory:** Every device required during a mains cut (edge host, broker, switches, time source, operator display and, only where required offline, vision cameras and lighting) with measured power draw across operating modes. [S03] [S05]
5. **Environmental envelope at the equipment location:** Ambient temperature, humidity, cotton-lint exposure and enclosure ventilation across peak shift temperatures. [S03] [S05]
6. **Local maintenance and escalation ownership:** Named electrical technician, IT/OT administrator or shift supervisor responsible for logs, fallback and recovery sign-off. [S03] [S05] [S06]

**Exhibit 8**, p. 39 provides an illustrative template for this dossier, pairing each baseline metric with its recorder, operator cadence, missing-data flag and approval owner. [S02] [S03] [S05]

EXHIBIT 08 / SITE BASELINE SPECIFICATION

## Every site unknown has a collection method

| Metric and boundary | Unit | Recorder or source | Operator cadence | Missing-data flag | Approval owner |
|---|---|---|---|---|---|
| Mains voltage, phase and frequency events | V, Hz, ms | Three-phase power quality logger at LV board | Continuous event trigger plus 1 min summary | recorder_unpowered | Site electrical engineer |
| Generator and UPS transfer interval | ms, s | Digital sub-panel logger and UPS telemetry | Every changeover event plus weekly test | transfer_uninstrumented | Maintenance lead |
| End-to-end WAN commit probe (DNS, TLS, HTTP) | ms, bytes/s, status | Persistent edge probe to test ingestion endpoint | Operator-chosen (e.g. every 60 s in shift) | observer_offline | IT and OT integrator |
| Joint power-and-WAN state occupancy | shift-hours | Joined power and probe ledger (UTC / UUIDv7) | Compiled daily and at baseline close | clock_unsynchronised | Plant operations manager |
| Enclosure ambient temperature and dust | deg C, RH percent | Calibrated data logger inside edge cabinet | Operator-chosen (e.g. every 10 min) | sensor_absent | Reliability engineer |
| Carrier path and tower power dependency | route audit record | Physical duct walkdown and ISP questionnaire | Pre-contract and after any route change | path_unverified | Procurement lead |

ILLUSTRATIVE [S02] [S03] [S05] Illustrative site-survey dossier template derived from NIST SP 800-82 Rev. 3 (S03), NIST SP 800-34 Rev. 1 (S05) and World Bank survey boundary discipline (S02). Cadences are operator-selected planning inputs, not observed GFT results.

As **Exhibit 8**, p. 39 emphasises, the fifth column (`Missing-data flag`) is as important as the measurement itself. [S02] [S05] Whenever a logger is absent, loses power during a long fault, or lacks a synchronised clock, the dossier records the explicit flag (`recorder_unpowered`, `observer_offline`, `clock_unsynchronised` or `path_unverified`) rather than substituting a national average or vendor brochure figure. [S02] [S10] [S22]

> **WARNING**
>
> **Unverified path independence and unrecorded tail risk.** The most common survey failure is accepting a vendor's verbal assurance that a backup link is "independent" without walking the cable duct and checking whether the backup base station shares the factory's utility substation. [S03] [S05] Until `path_unverified` is cleared by a physical route audit and joint outage log, the design must assume that external backhaul can fail simultaneously with utility mains. [S03] [S05]

> **GFT NOTE**
>
> **Practical recommendation for programme tools and cross-paper handoff.** In GFT's Cameroon textile programme, this site continuity dossier supplies the empirical input parameters for the `site-selection`, `edge-architecture`, `solar-storage` and `risk-matrix` decision tools. [S05] Observed joint outage durations and enclosure temperatures feed the protected-load and battery sizing calculations in `edge-architecture` and `solar-storage`, while any parameter still carrying `path_unverified` or `recorder_unpowered` is registered as an active exposure in `risk-matrix` with a named owner and measurement action. [S03] [S05] Even though third-party industrial projects and existing mills operate or are announced in Douala, Dibamba and Garoua, including the announced US$300 million (180 billion XAF) Camtext SA project targeting 12,000 tonnes a year processed with hubs at Dibamba and Garoua, and CICAM's historic plants in Douala and Garoua, GFT refuses to rank candidate locations using fictional feeder scores. [R09] [R10] [R12] [S02] Whole-plant motive power, solar generation yield and thermal resource efficiency are handed to `wp06`, while `wp07` retains ownership of the failure envelope, the protected digital load and proof of local record survival. [S03] [S05]

With the site survey protocol, joint-state matrix and recovery timestamps defined, the next step is to translate the measured interruption envelope into a binding operational contract. [S03] [S05] Chapter 3 establishes how local service classes, recovery-time objectives, recovery-point budgets and explicit degraded operating states turn site observations into testable engineering requirements before any hardware or cloud platform is purchased. [S03] [S05] [S06]

*National averages justify the question, but only a joint site log of power, network and recovery lag can size the answer.*

**CHAPTER SOURCE NOTES**

S01 World Bank. 2025.

S05 NIST. 2010.
S03 NIST. 2023.
S02 World Bank. 2026-10-30.
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.
S13 SQLite. 2026-10-30.
S31 IETF / RFC Editor. 2024.
S32 IETF / RFC Editor. 2002.
S10 IETF / RFC Editor. 2010.
S22 systemd project. 2026-10-30.
S16 SQLite. 2026-10-30.
S29 Linux kernel project. 2026-10-30.
S18 IETF / RFC Editor. 2014.
S24 IETF / RFC Editor. 2022.
S12 IETF / RFC Editor. 2022.
S23 IETF / RFC Editor. 2026.
S0204 IETF / RFC Editor. 2024-03.
S0201 IETF / RFC Editor. 2023-12.
S0202 IETF / RFC Editor. 2011-08.
S0203 IETF / RFC Editor. 1998-05.
S11 IETF / RFC Editor. 2019.
S30 chrony project. 2026-10-30.
S08 IETF / RFC Editor. 2007.
S14 SQLite. 2026-10-30.
S06 NIST. 2016.
S21 systemd project. 2026-10-30.
R09 Ecofin Agency. Publication date unrecorded; inherited source-text check.
R10 Ecofin Agency. Publication date unrecorded; inherited source-text check.

03

PART I: SPECIFY THE INTERRUPTION / SPECIFY THE BOUNDARY

# Local service classes turn resilience into a testable contract

*Criticality, allowable staleness and restart dependencies should be assigned before choosing equipment or a cloud platform.*

CHAPTER 03
DECISION TEST

Criticality, allowable staleness and restart dependencies should be assigned before choosing equipment or a cloud platform.

A factory should buy continuity against a defined operating duty, rather than against a general promise of resilience. GFT proposes a service contract that names the task, its permissible interruption, its record-loss boundary, its offline dependencies and the person authorised to resume it. Local production records, inspection decisions and remote planning need different contracts. The classification below is a proposed engineering method, not a report of an installed GFT system. It translates contingency priorities and recovery planning into requirements that a factory team can approve and test. [S05] [S06]

3.1

## Classify tasks by the consequence of interruption

Begin with the decision the task enables. Ask what happens if its result is absent, stale, incorrect or irrecoverable. A missing inspection result may require material to be held. A missing dispatch forecast may permit planning to continue with a visible qualification. An incorrect machine-control signal raises a different engineering responsibility. NIST's operational technology guidance explicitly recognises performance, reliability and safety requirements. GFT therefore proposes classifying functions by their consequences before grouping them by application or supplier. [S03]

The proposed safety-related class belongs to the machinery and process engineering authority. Its boundary should identify which controls establish the permissible plant state, which alarms remain available locally and which information may pass to management systems. This chapter does not prescribe protective settings, safety logic or a safety integrity classification. Analytics should receive an explicit interface contract; their availability must not be treated as evidence that the machinery protection function works. The responsible engineer must determine the permitted behaviour on lost communication. [S03]

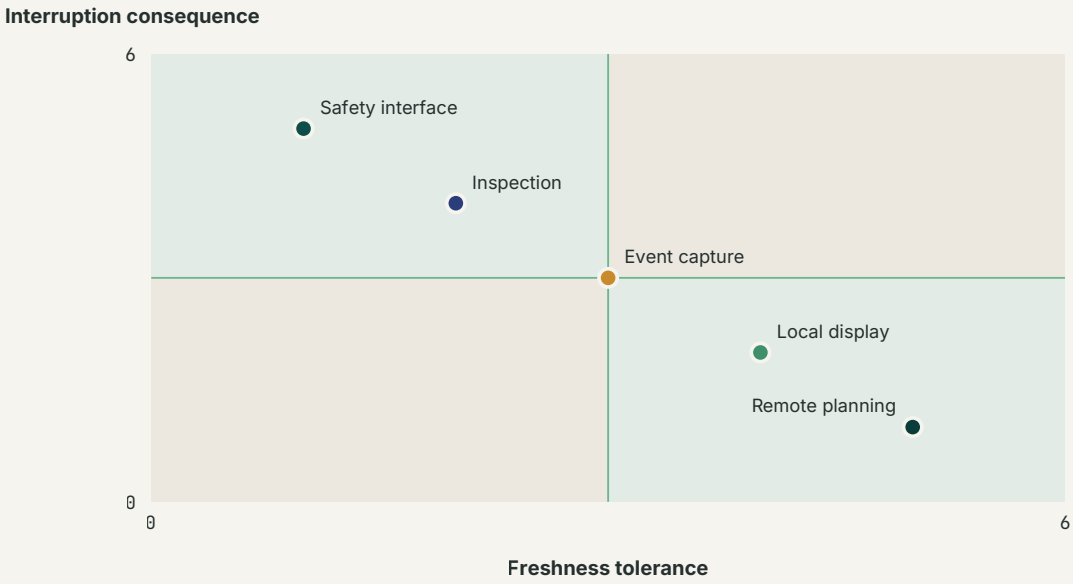
The inspection class concerns permission to accept, reject or hold material. GFT recommends naming the required sensing, illumination, model or rule set, result association and human quality authority. The event-capture class concerns retaining an accountable account of production actions. Its minimum useful output is an identifiable, recoverable record, even when the remote destination is unavailable. These functions can share equipment without sharing an acceptance result: a camera failure might disable inspection while manual production recording remains permissible. That separation is a proposed operating rule requiring site approval.

Local visual management concerns what the supervisor can use at the workplace. Define the latest trustworthy production state, the visible age of that state and how uncertainty changes the decision. Remote planning concerns aggregation and future allocation. It may tolerate delayed receipt while still requiring complete history later. GFT recommends documenting these as distinct duties even where a supplier presents a common dashboard. A visual interface is not the boundary of every service it displays.

The illustrative matrix in **Exhibit 9**, p. 44 separates the consequence of interruption from freshness tolerance. Its coordinates express an assumed workshop judgement, not measured risk or a numerical safety standard. The site team should replace these relative positions with approved consequences and explicit requirements. Equal proximity on the drawing does not mean equal equipment, budgets or authority.

EXHIBIT 09 / ILLUSTRATIVE ORDINAL POSITIONS

## Consequence and freshness require separate judgements




ILLUSTRATIVE [S03] [S05] Illustrative GFT workshop model, not measured risk. Freshness tolerance: 1 immediate, 5 planning horizon. Consequence: 1 deferrable, 5 requires protective authority. Positions are assumed discussion prompts; neither axis is a safety rating or a service-level target.

For each class, GFT proposes a requirement card with the process owner, observation boundary, permitted modes, recovery time, recovery point, maximum trustworthy age and offline operating interval. Separate compulsory fields from undecided fields. An unresolved tolerance should block acceptance of that duty rather than disappear into a supplier assumption. Contingency planning provides a basis for setting priorities; it does not supply a universal factory target. [S05]

Do not let the classification hide different boundaries within an operation. GFT proposes distinguishing permission to keep making material, permission to record a completed action and permission to release that material. The same interruption can suspend release while leaving capture available. The requirement card should identify the affected duty directly, together with the physical location of held material and the means of preventing an unauthorised handover. This makes the classification usable by operators rather than leaving them to interpret a general criticality label during a fault.

A useful challenge is to ask whether a function is genuinely essential during the interruption. Protecting an inspection workstation is unnecessary if the process must stop and no material can move. Conversely, declaring inspection non-essential is

indefensible if dispatch can release unverified output. The answer must follow the material flow and the authorised operating method. This classification is the input to the protected-load design, not a substitute for measuring that load.

3.2

## Availability and durability need separate budgets

Recovery time defines how long a specified service may remain unavailable. Recovery point defines the historical boundary to which its data must be restored. NIST's contingency guide treats these as different planning concepts. GFT adds a record-count check because elapsed time alone can conceal gaps in an irregular event stream. State whether loss is measured in time, missing acknowledged events, unresolved material movements or a combination, and record the applicable failure envelope. [S05]

For the proposed contract, availability means the duty can produce its required result, not simply that its process is running. A screen that answers a health request while its capture path rejects transactions is unavailable for capture. Durability means the accepted record survives the named failure through the configured software and storage path. SQLite documents that transactions committed with WAL and synchronous=NORMAL can roll back after a power loss or system crash. A consistent database is therefore insufficient evidence of the promised record survival. [S16]

Define the acknowledgement vocabulary before testing. Received means the application has taken responsibility for examining an event. Validated means the event passed the agreed checks. Locally committed means the configured local transaction completed. The contract must additionally state what durability that configuration claims. Remote accepted means the receiving application validated and durably committed the identified event. These are proposed application states; TCP's reliable byte stream does not define the downstream business commit. [S24]

GFT recommends acknowledging a production record only at the promised acceptance boundary and retaining its identity through onward delivery. SQLite's atomic-commit description supports grouping related changes within a transaction; the project must still qualify persistence under its own failure conditions. Detailed storage implementation belongs later in this paper. Here, the purchasing obligation is to name the boundary and expose evidence that it was crossed. [S13]

### WORKED EXAMPLE: A RESTORED INTERFACE CAN CONCEAL A FAILED RECORD CONTRACT

This illustrative case assumes a scheduled duty of 480 minutes, an availability target of 99.5 percent and an event rate of 120 accepted events per minute before an injected failure. These values explain the method; they are not Cameroon observations, a GFT result or recommended factory targets. The assumed outage allowance is 480 × (1 − 0.995) = 2.4 minutes. No scheduled period is excluded from the denominator.

Suppose the illustrative test returns the capture service to verified readiness after 1.5 minutes and its recovery-time objective is 2 minutes. Availability is (480 − 1.5) / 480 × 100 = 99.6875 percent, leaving 2.4 − 1.5 = 0.9 minutes of the assumed outage allowance. The example assumes production pauses during that interruption, so no additional production events are created while capture is unavailable.

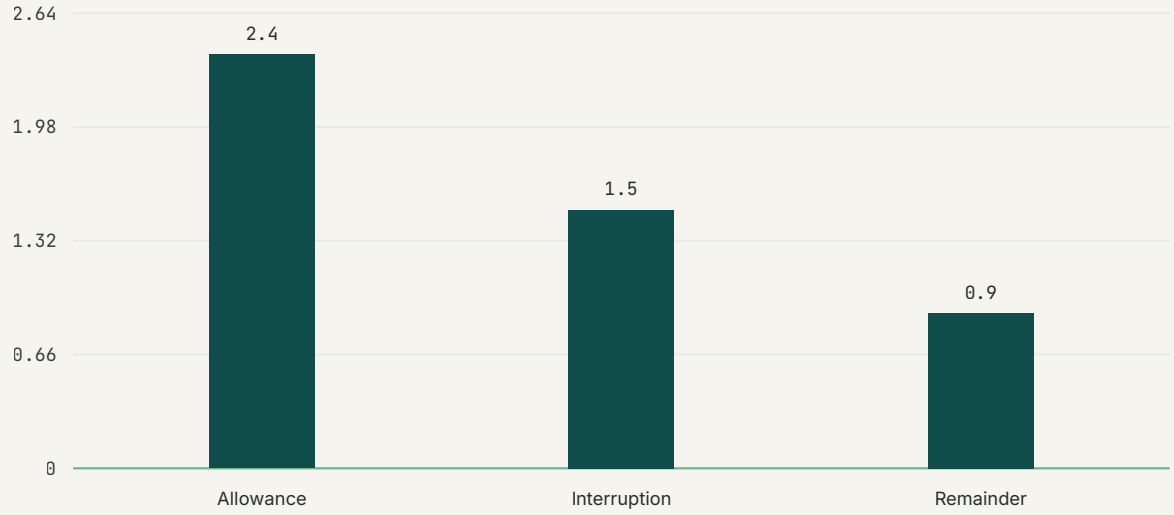
Now assume the same illustrative implementation had acknowledged the preceding 90 seconds of events from a volatile queue. A comparison of the synthetic input ledger with the recovered ledger finds 120 × 90 / 60 = 180 acknowledged events absent. Recovery time passes, but the assumed requirement to retain every acknowledged event fails. The lost interval is 90 seconds in this uniformly spaced example; the missing count remains essential because real arrivals need not be uniform.

This illustrative result calls for correcting the acknowledgement and persistence boundary, then repeating the failure test. A faster screen restart would not repair the missing history. Availability, recovery time and acknowledged-record survival must receive separate acceptance decisions.

The illustrative timing exhibit in **Exhibit 10**, p. 47 reconciles the permitted interruption and the amount consumed. It says nothing about the lost-record result. Keeping that result outside the timing axis prevents incompatible quantities from being combined into an apparently favourable resilience score.

EXHIBIT 10 / MINUTES WITHIN AN ILLUSTRATIVE SCHEDULED DUTY

## The illustrative outage uses only part of its time allowance




ILLUSTRATIVE Illustrative inputs: scheduled duty 480 minutes, availability target 99.5 percent, measured-in-scenario interruption 1.5 minutes. Allowance = 480*(1-0.995) = 2.4 minutes; remainder = 2.4-1.5 = 0.9 minutes. The scenario separately loses 180 acknowledged events and fails its record contract. No operating result is asserted.

A no-loss commitment must also name the failures it covers. GFT proposes separating process termination, abrupt power loss and loss of the storage device itself. Demonstrating survival after a process restart does not establish survival after physical destruction of the only retained copy. The requirement should say whether its promise concerns acknowledged local events, remotely accepted events or reconstructible manual evidence. Keep any excluded failure visible in procurement acceptance and in the operational instructions. A bounded promise can be tested; an unqualified promise that all records always survive cannot be established by a successful restart.

Failure budgets also require ownership. GFT proposes recording who can approve temporary deviation, which duty is suspended and what evidence closes the exception. An availability allowance must never become permission to bypass a required quality check or protective action. State the production calendar, the fault-start marker and the readiness-end marker alongside the result. Otherwise, changing the accounting boundary can improve the reported percentage without improving the service.

3.3

## Every dependency must inherit an offline requirement

A local application earns an offline claim only for the complete duty it supports. GFT recommends tracing required power, storage, local networking, name resolution, identity, licensing, time, software assets and human permission. For each depend-

ency, ask whether it is available during continued operation, after a process restart and after a cold boot. These are different scenarios. An existing session continuing to run is not evidence that the station can start again without external access.

The illustrative stack in **Exhibit 11**, p. 48 makes the inspection boundary inspectable. Its content is a proposed dependency model, informed by the distinct reliability and safety concerns in OT guidance and by the existence of authenticated time mechanisms. It does not claim that a specific supplier has these dependencies or that any station has passed qualification. [S03] [S27]

EXHIBIT 11 / ILLUSTRATIVE DEPENDENCY GROUPS

## An offline inspection duty includes its supporting services




ILLUSTRATIVE [S03] [S27] Illustrative GFT inspection dependency model. Each required item must be available or have an approved substitute within the declared failure envelope. Local execution alone proves neither offline restart nor continued authorisation. This is not an implemented topology.

GFT's proposed procurement review should examine licence renewal, activation, model retrieval and account verification explicitly. Require the supplier to demonstrate the installed version starting within the agreed offline envelope, using the intended credentials and retained assets. If a renewal interval limits the claim, write that limit into the requirement. Do not invent a vendor's grace period or assume that a cached token remains valid. These are qualification questions, not allegations about a named product.

Time and identity should appear as separate dependencies. Network Time Security specifies a security mechanism for network time synchronisation. Its presence is a reason to test the selected implementation's bootstrap and disconnected behaviour,

not a guarantee of it. GFT recommends recording the clock's trust state and the permitted actions when external time is absent. Later clock analysis can establish the device-specific uncertainty; the service card must already say which decisions depend on trustworthy time. [S27]

The dependency review should include circular prerequisites. Ask whether identity verification needs a network service that itself waits for a clock or account service. Ask whether the local display requires a remote name even though the underlying data resides on site. A proposed mitigation must be tested through a cold boot rather than judged from a configuration diagram. Avoid assuming that bypassing authentication is an acceptable way to preserve continuity. Define an approved restricted local role instead, with its own authorisation boundary.

Human capacity is also part of the proposed offline contract. A manual release requiring an absent approver is unavailable even if the screen works. Assign an available role, an authorised substitute and an escalation route. Verify that local instructions and required credentials can be accessed under the failure being tested. A document stored behind the failed external service cannot serve as the sole recovery instruction. This follows GFT's requirement to specify an executable duty rather than a list of functioning components.

Test the end of the approved offline interval as well as its beginning. GFT proposes checking whether an authorisation expires, whether a required asset is refreshed and whether the selected application changes behaviour after extended isolation. The supplier should explain how these limits are displayed before they prevent the promised duty. If continued operation requires a controlled renewal visit or signed offline package, include that activity in the service contract with its responsible role. A dependency is not removed by renaming its planned maintenance as an exceptional support task.

The practical handover is a dependency register that distinguishes required, substitutable and optional items. For each required item, retain its supported offline interval, restart behaviour, failure indication and evidence reference. Leave untested dependencies visibly unqualified. The wider reference architecture can use this register later; it should not silently broaden an accepted record-keeping duty into an accepted inspection duty.

3.4

## Safe degradation must be an explicit state

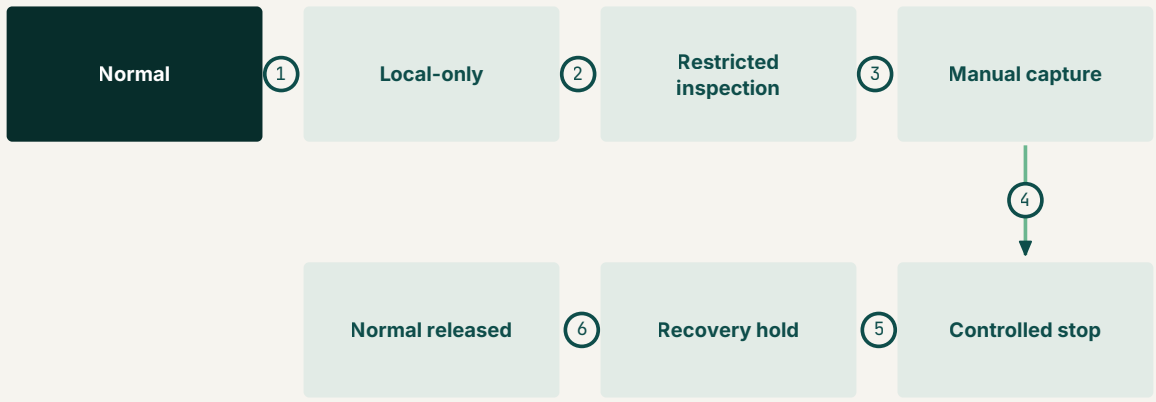
GFT proposes an operating-state model that describes what work remains permitted as dependencies disappear. Normal means the declared duty and its required communication paths are available. Local-only means the approved local duty remains available while upstream delivery is deferred. Restricted inspection means the quality authority has approved a narrower inspection method or material route. These are proposed modes, not automatically safe states created by a network status flag. Their permission must follow the process and equipment assessment. [S03]

Manual capture means operators use an approved substitute record while digital intake is unavailable. It does not mean production may continue regardless of inspection or machinery conditions. Controlled stop means the responsible process authority has suspended the relevant operation under its approved procedure. Recovery hold means dependencies are returning but the service has not yet satisfied its release checks. GFT recommends displaying permitted actions, the reason for restriction and the responsible role in plain local instructions.

The illustrative flow in **Exhibit 12**, p. 50 includes restricted inspection explicitly. The arrows show a representative route through the states, not every permissible transition. Direct recovery from local-only, restricted inspection or manual capture must also enter recovery hold. A failed release check leaves the duty on hold or invokes the approved stop. These are permission conditions, not executable safety logic. The diagram is a proposed continuity playbook structure, adapting the planning and testing approach of NIST's recovery guidance to ordinary interruption. It does not imply that every power failure is a cybersecurity incident. [S06]

EXHIBIT 12 / ILLUSTRATIVE OPERATING STATES

## Returning dependencies must pass a recovery hold




1. WAN lost; local ready
2. Quality approves limit
3. Capture fallback approved
4. Fallback permission fails
5. Recovery permitted
6. Checks and release pass

ILLUSTRATIVE [S03] [S06] Illustrative GFT route through the operating states, not executable safety logic or an exhaustive transition graph. Quality and capture fallbacks need approval. Every mode must stop when its operating permission fails. Returning dependencies require recovery hold and authorised release. Direct recovery routes and failed-release handling are specified in the text. The final node repeats the normal state after authorised release.

Each state needs an entry trigger that can be observed locally. A missed upstream heartbeat may justify marking remote information stale; it does not by itself diagnose the cause. A failed persistence check should prevent an event from being labelled

durably accepted. An unavailable inspection result should create the approved hold or substitute route. GFT recommends keeping uncertainty visible until the responsible role resolves it, rather than replacing a failed reading with the last apparently normal value.

Local maintenance alarms should remain within the declared local duty where that duty depends on timely intervention. The proposed console should distinguish an active alarm, an acknowledged alarm, a missing sensor and an unavailable alarm path. A silent upstream dashboard must not count as evidence of normal plant conditions. The site engineering authority should decide which alerts require independent local indication, consistent with the differing reliability and safety needs of operational technology. [S03]

State transitions also need records. GFT proposes retaining the trigger, observing component, prior state, new state, approving role and unresolved conditions. Record a rejected transition as well as an accepted transition. This makes it possible to investigate why production remained restricted after apparent technical restoration. The recovery owner should close the discrepancy through an approved decision, without editing the historical transition into a more convenient account.

> **GFT NOTE**
>
> Approve the degraded operating method before procuring its supporting hardware. Put the permitted work, local warning, record obligation and release authority on the same requirement card. If the team cannot explain what operators should do when a dependency fails, the service class is not ready for acceptance.

3.5

## Bound the retained record by time, volume and priority

Offline continuity needs an explicit storage policy as well as an operating mode. Delay-tolerant networking describes persistent storage as a response to network interruption. Constrained-network terminology also recognises limited and unpredictable capacity. These sources support planning for retained work; they do not establish the necessary storage capacity or acceptable loss at a particular factory. GFT recommends specifying the required record interval and measuring the workload before choosing the device. [S08] [S18]

Define retention dimensions separately. The design offline interval describes how long the service is expected to operate without its normal upstream acceptance. The byte budget describes how much local capacity that workload may consume. Maximum age describes whether a retained item remains useful for a decision or must still be kept for traceability. A planning recommendation may expire while the production event that informed it remains necessary. Expiry is therefore a business rule, not merely a storage setting.

The illustrative policy table in **Exhibit 13**, p. 52 separates record categories and possible pressure responses. It leaves retention durations unset because no site workload, contractual requirement or approved policy has been supplied. An unset duration is a required design decision, not permission to delete. The entries are GFT's proposed discussion structure, not statutory retention guidance.

EXHIBIT 13 / ILLUSTRATIVE POLICY TEMPLATE

## Capacity pressure requires an approved record policy

| Category | Retained boundary | Pressure response |
|---|---|---|
| Production events | Accepted history | Restrict intake |
| Audit events | State and authority | Preserve reserve |
| Inspection images | Approved evidence set | Hold or approved reduction |
| Diagnostics | Agreed investigation window | Count approved expiry |

ILLUSTRATIVE [S08] [S18] Illustrative GFT retention template. No durations or legal obligations are inferred. Allocate time, bytes and maximum age separately for each category. Priority is an operator-approved decision; no mandatory production or audit record is silently discarded.

GFT recommends agreeing the pressure sequence before storage fills. Reduce optional diagnostic detail first only where that reduction is approved and counted. Preserve capacity for the records that explain the restriction itself. If a required production record can no longer be durably accepted, expose that failure and apply the authorised manual route or stop. Silently overwriting acknowledged history would contradict a contract that promised its survival. An alert is insufficient unless a role can act within the remaining operating interval.

Priority must not be confused with delivery order. A retained event may be low priority for transmission but essential to reconcile a later material movement. Conversely, an urgent alarm may need immediate local action even when its remote copy must wait. The proposed class card should state which actions need freshness, which records need completeness and what ordering relationships the receiver requires. Detailed queue-drain calculations belong to the later queue chapter; this chapter establishes the obligations those calculations must satisfy.

Deletion should leave evidence of the deletion rule even where the original payload is no longer required. GFT proposes retaining the category, covered interval, reason, approving policy revision and discarded count, with a clear distinction between intentional expiry and unexplained loss. A completeness display should then show which categories remain comparable across the interruption. If image detail was reduced under an approved policy, the later reviewer must see that qualification alongside the associated inspection record. Otherwise, a full event count could conceal a reduced evidential basis.

An offline envelope also needs an exit condition when interruption exceeds the design assumption. GFT proposes a visible reserve threshold, an accountable escalation and a permitted restricted duty. State whether the remaining capacity is estimated from

current arrival rate and how changing image volume affects that estimate. Treat an unmeasured rate as unresolved. A nominal storage size cannot establish how much approved operation remains under the actual workload.

3.6

## Manual continuity requires the same identifiers

A manual fallback should preserve the meaning of the event, not imitate the appearance of the screen. GFT proposes a controlled paper form and a matching offline spreadsheet schema with bundle identifier, operation, shift, operator role, local sequence and reconciliation state. Add the original event time, its time basis and the physical material location. Retain the distinction between the production bundle and the event: a bundle can undergo different operations and corrections without becoming a different bundle.

Identifiers must be available without an online allocator. The UUID specification describes identifiers that require no central registration process, but correct generation remains an implementation responsibility. GFT recommends either a qualified local generator or pre-issued controlled form identifiers for manual use. A timestamp alone should not carry identity. An offline spreadsheet copied to another device must not silently create a new identifier namespace that can collide with earlier manual entries. [S31]

Record time in a format with an explicit relationship to UTC when that relationship is known. The timestamp specification addresses that interchange relationship; it does not prove that the operator's clock was correct. GFT proposes keeping an uncertain original time, a confidence flag and the later import time as separate fields. Do not manufacture precise production times from spreadsheet row order or overwrite the original observation with the time at which connectivity returned. [S32]

For an illustrative form entry, use bundle BUNDLE-A, operation INSPECT, shift SHIFT-A, role QUALITY, sequence SEQ-A and state PENDING-IMPORT. These symbolic values demonstrate the field relationships without suggesting a real factory transaction. A matching import candidate should carry the controlled form identity and its original line reference. Mark a correction as a linked correction with its reason and authority, rather than erasing the original entry.

GFT recommends importing through a staging area that validates identifiers, required fields, permitted operations and material references before posting to the performance-management system of record. Compare the candidate against digital events created before the failure and against earlier imports. If the same event identity has a different payload, hold the conflict for a responsible reviewer. A duplicate candidate is not a second production action merely because it arrived through a different channel.

The import contract should make repeating an approved import harmless to the intended business result. HTTP defines idempotent method semantics, but a method name alone cannot establish factory-level duplicate handling. GFT proposes retaining a durable import receipt with the accepted identity and outcome, using a transaction where related local records must change together. The application must demonstrate that retrying a timed-out import does not duplicate the business action. [S12] [S13]

Close manual continuity by reconciling issued forms, used forms, cancelled forms and posted entries. Record unresolved gaps explicitly. Keep the original evidence available for the approved retention period and record who authorised closure. This chapter defines the continuity interface; detailed access policy and governance belong in the governance paper. Manual capture is a qualified substitute only when the team can both use it during the interruption and reconcile it afterwards.

3.7

## Acceptance is evidence of behaviour under a named failure

GFT proposes translating every service card into a failure experiment. Name the stimulus, starting mode, workload, accepted-event ledger, permitted loss, readiness condition, restart owner and retained proof. Link each observed result to the exact requirement revision. Recovery guidance supports realistic scenarios and improvement through testing; a planned test is nevertheless different from a completed test, and neither should be reported as a deployed factory result. [S06]

The proposed acceptance set should distinguish loss of external connectivity, application termination, local storage exhaustion and a controlled loss of device power. Select combined cases where dependencies plausibly coincide, including offline cold boot and recovery while a queue remains occupied. Electrical interruption work requires the responsible site engineer and the equipment's permitted test method. A software restart cannot establish the behaviour of the power and storage path during abrupt supply loss.

SQLite documents simulated crash testing and compound failure testing. That is useful implementation evidence, but GFT recommends project tests on the selected configuration as well. Compare the identified input ledger, acknowledgements and recovered ledger; inspect missing, duplicated and conflicting outcomes. Retain the actual trace rather than a summary that merely says the application restarted. This is a proposed acceptance method, not a claim that the chapter author has performed hardware fault tests. [S28]

For each experiment, report availability, recovery time, record survival and permitted production mode separately. An unresolved record gap should fail a completeness requirement even if the interface meets its recovery time. If power returns but inspection remains restricted, distinguish technical readiness from authorised process release. Include the time spent waiting for local authority where the contract promises an operable duty, rather than measuring only the application's start time.

Procurement acceptance should require a signed site-specific requirement record and a matching evidence index. GFT proposes recording the software versions, hardware configuration, tested failure envelope, exclusions and residual issues. The Cybersecurity Framework provides outcomes without prescribing their implementation. A reference to that framework is therefore not a certificate, a substitute for this evidence or proof that a supplier's continuity claim applies to the selected installation. [S26]

Before signing, the owner should be able to follow a requirement through the stimulus, raw observation, calculation and acceptance decision without asking the supplier to reconstruct the story. GFT proposes preserving failed attempts and the change that preceded a successful repeat. Keep unresolved exclusions attached to the accepted duty, including any missing measurement of the site envelope. A laboratory success can qualify the tested configuration within its stated conditions while still leaving deployment approval pending. The final decision should say which boundary has been proved and which further observation would justify extending it.

Use the resulting cards as inputs to the edge-architecture and maturity tools. A planning duty can be accepted while inspection remains unqualified, provided the distinction stays visible. Protected-load engineering should inherit the duties and dependencies approved here. The performance-management system should report unresolved gaps alongside successful tests, so management can see which production decisions the evidence actually supports. GFT's practical recommendation is to expand the promised boundary only after its new dependencies and recovery obligations have been demonstrated.

*Resilience becomes a contract when the permitted work and the surviving record can both be tested.*

**CHAPTER SOURCE NOTES**

S05 NIST. 2010.
S06 NIST. 2016.
S03 NIST. 2023.
S16 SQLite. 2026-10-30.
S24 IETF / RFC Editor. 2022.
S13 SQLite. 2026-10-30.
S27 IETF / RFC Editor. 2020.
S08 IETF / RFC Editor. 2007.
S18 IETF / RFC Editor. 2014.
S31 IETF / RFC Editor. 2024.
S32 IETF / RFC Editor. 2002.
S12 IETF / RFC Editor. 2022.
S28 SQLite. 2026-10-30.
S26 NIST. 2024-02-26.

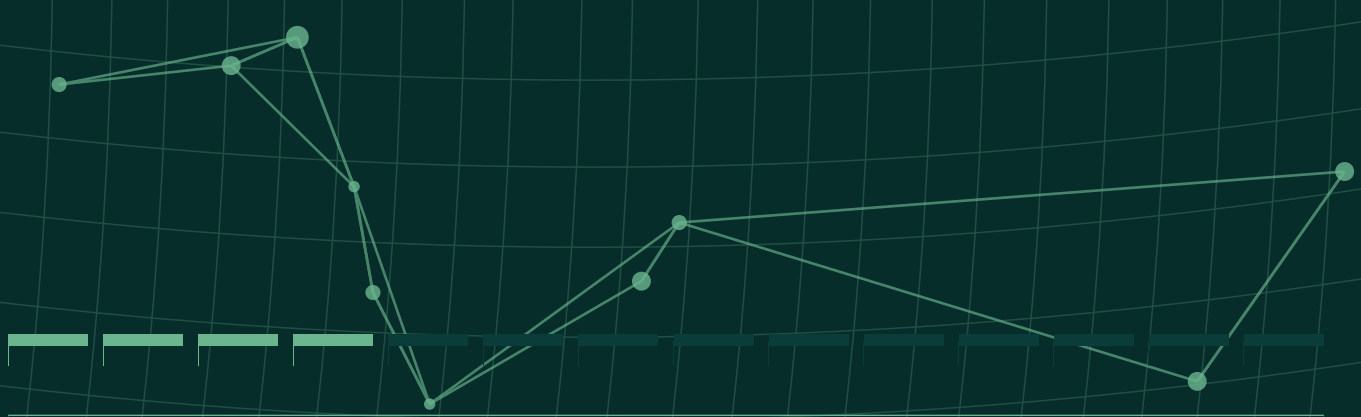
GREEN FUTURE TECH
LIMITED / HONG KONG SAR

04

PART II: KEEP THE LOCAL RECORD ALIVE / POWER THE RECORD

# A protected digital circuit should have a measured duty

*Protecting a selected record-keeping load can be practical even when the production machinery must stop.*



CHAPTER 04
DECISION TEST

Protecting a selected record-keeping load can be practical even when the production machinery must stop.

A factory on an unreliable grid should protect a declared digital duty, not an indiscriminate collection of electronic equipment. The useful boundary is the smallest complete chain that can capture, retain and display the records the operator has chosen to preserve during a power interruption. That chain may remain useful after production machinery stops, provided its sensors, network, compute, storage and operator interface share a suitable protected supply. GFT recommends measuring that chain in its actual operating modes, sizing energy from an agreed interruption duty and qualifying both transfer and exhaustion. The result is a testable continuity specification, not a claim that a UPS keeps the factory producing.

This distinction prevents two common design errors. A circuit sized only for an edge computer may lose the record because a switch, time source or local console has stopped. A circuit expanded to include every digital device may consume its reserve on work that could safely wait. Contingency planning begins with operational priorities, while OT design must retain performance, reliability and safety requirements. The protected-load decision therefore belongs to the production owner, electrical engineer and digital system owner together. [S03] [S05]

4.1

## Separate production energy from digital continuity loads

Start with the work that must remain possible. Record capture can sometimes continue when machines are stationary because operators still need to identify completed bundles, record their state and preserve evidence of the interruption. Inspection is different. An automated inspection duty needs its required camera, illumination, trigger, compute and disposition interface. If any required input is unpowered, a live model server cannot produce a valid inspection result. Machinery control is different again and remains under the responsible process and electrical engineering authorities. This chapter does not place production drives or protective functions on the digital continuity circuit. [S03]

GFT proposes a load inventory that follows the selected duty from observation to acknowledgement. For manual or supervisor-assisted capture, inspect the edge computer, local broker, access switch, router where local routing requires it, operator display, time source and any barcode or identification device. For automated inspection, add the complete sensing and illumination chain. For local reporting, include every display and network component needed to make the record understandable at the workplace. An item belongs in the inventory because the duty fails without it, not because it happens to use electricity.

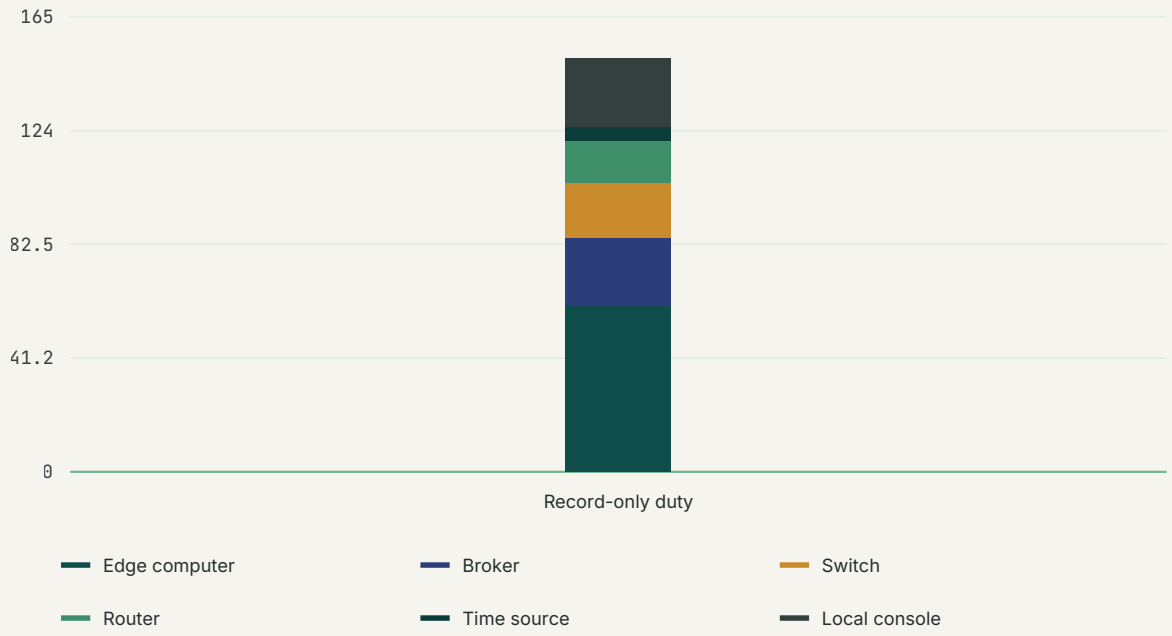
The survey should also expose hidden auxiliaries. Check power injectors, media converters, small floor switches, external storage, licence or security appliances, and the supply used by the monitoring function itself. Record the physical socket, protective device, upstream board and normal supply mode for each item. Identify which UPS outlets are actually backed, since a device plugged into an unprotected outlet on a protected cabinet does not inherit continuity. Treat these as site verification tasks rather than assumptions about a particular product.

For every load, record whether it is required continuously, only during capture, during recovery or not at all during the declared outage. A remote-upload radio may be unnecessary while the wide-area path is known to be unavailable, but the local switch that carries events to the store remains essential. Non-urgent displays and diagnostic equipment can be candidates for an approved load-shedding state. The operating instruction must say who authorises that state and how its reduced coverage is displayed.

The illustrative stack in **Exhibit 14**, p. 58 defines a record-only duty. It includes a local console and time source, but excludes inspection cameras and lighting. Its component values are assumptions chosen to demonstrate load accounting. They are neither measurements from a factory nor recommended equipment ratings.

EXHIBIT 14 / WATTS, ILLUSTRATIVE STEADY OPERATING SCENARIO

## The record-only load includes six supporting components




ILLUSTRATIVE [S03] [S05] Illustrative GFT load boundary. Assumed values: edge computer 60 W, broker 25 W, switch 20 W, router 15 W, time source 5 W and local console 25 W. Total = 60+25+20+15+5+25 = 150 W. Cameras, illumination, machinery and production control are excluded. This protects proposed record capture and local reporting only; it is not an equipment selection or an operating result.

The stack makes the service boundary inspectable. Adding an inspection duty would change both the component list and the operating modes to measure. Removing the console might be acceptable only if another powered and authorised means of supervising capture remains. The exhibit cannot establish that the listed devices are compatible, that their combined peak demand is acceptable or that the proposed duty is safe for a particular process.

A useful boundary also names what happens beyond the protected circuit. If production stops, identify the approved method for recording material already in progress and the point at which manual capture replaces automated input. If the record

depends on an unprotected sensor, mark the resulting interval as unknown rather than filling it with inferred production. Protected storage can preserve events it receives; it cannot reconstruct observations that were never made.

> **GFT NOTE**
>
> Approve a one-page protected-duty card before selecting a UPS. Name the permitted work, included devices, excluded loads, required local display, manual fallback and authorised stop. Procure the power path against that card. A battery attached to an undefined digital load protects expenditure, not a production record.

4.2

## Measure consumption in the relevant operating modes

A nameplate is a procurement clue, not an autonomy measurement. It may describe a maximum input, a typical output or a supply rating rather than the power drawn by the configured device in the promised duty. GFT recommends measuring at the boundary that the continuity system will actually feed, with the intended software, peripherals and network state. Record the instrument, its operating range, the sampling method and the measurement uncertainty alongside the trace. The responsible electrical engineer should confirm that the method is suitable for the load and supply waveform.

Measure a set of modes that follows the production record. Idle capture establishes the standing duty. Active ingestion adds storage and communication work. Inference, if included, may create a different compute profile. Boot and service restart can produce a temporary peak while several components initialise together. Backup, database checkpointing and log rotation can change storage demand. Communication bursts can add radio or network load. These modes should be observed with their simultaneous auxiliaries rather than reconstructed from unrelated device readings.

The measurement plan should retain mode identity as carefully as watts. Each interval needs the software state, active peripherals, capture rate, network condition and any operator action that explains a change. A high reading during boot must not disappear into an average, and a quiet period with inspection disabled must not represent full inspection duty. Where the recorder itself loses power, mark the interval missing. Do not turn a gap in the baseline into an assumed low-load period.

Energy and power answer different questions. Power, expressed in watts, is the instantaneous rate of use. Energy, expressed in watt-hours, accumulates power over time. A suitable energy estimate integrates the measured load profile across the declared duty, while a transfer and supply-capacity assessment must consider the relevant peaks and equipment limits. Multiplying one convenient watt reading by the outage duration is defensible only when that reading represents the whole duty or is explicitly adopted as a conservative scenario.

Use a mode-weighted calculation when operation is repetitive and the mode durations are evidenced. For each mode, multiply measured power by its duration, sum the resulting energy and divide by the total interval to obtain average power. Retain the maximum observed simultaneous demand separately. If the factory has no represent-

ative trace, present a bounded illustrative scenario and commission measurement before procurement. A larger arbitrary margin can hide ignorance without proving that the power path will start or carry the load.

The electrical boundary also determines which losses belong in the arithmetic. Measure or document whether the load figure describes device input, UPS output or UPS input. Conversion losses between those points cannot be counted twice, and a device-level total cannot be applied directly to a battery without an identified conversion path. Supplier efficiency and surge information should apply to the intended model, configuration and operating condition. Where no suitable documentation exists, retain an explicit assumption and make its verification an acceptance action.

Finally, preserve the raw trace and its context as design evidence. A summary table should show the mode, observation window, mean demand, relevant peak, missing intervals and measurement qualification. It should distinguish observations from values assumed for sizing. This record lets the electrical engineer assess supply behaviour while the digital owner checks that the measured mode really performed the promised capture duty. A watt number without that context is not yet a protected-load specification.

4.3

## Size energy from duration and explicit derating inputs

The energy calculation should begin only after the load and duty are explicit. For a simplified constant-load case, let $P$ be the protected load in watts and $t$ the required autonomy in hours. Their product is the energy required at the protected output boundary. Divide by the applicable conversion efficiency $\eta$, the permitted usable fraction $u$ and the remaining capacity fraction $c$ at the declared end-of-life condition. The illustrative relationship is $E = P \times t / (\eta \times u \times c)$, where $E$ is nominal battery energy in watt-hours.

Each denominator term answers a different question. Conversion efficiency represents losses across the modelled path. Usable fraction represents the portion of nominal capacity the design permits itself to use under its operating policy. Remaining capacity represents the condition at which the owner still expects the system to perform. These are not universal battery constants. They must come from measured evidence, applicable supplier documentation or a declared scenario, and their boundaries must not overlap. A supplier figure that already incorporates one allowance should not receive the same allowance again.

Autonomy itself is an operational requirement, not a national outage average. Choose it from the site outage envelope, the approved manual fallback and the time needed for an orderly stop. State whether the duty includes continued capture, local display, controlled shutdown or all three. A longer requirement can be evaluated transparently, but it cannot be justified by an invented duration distribution. Contingency priorities supply the reason for the requirement; they do not supply its numerical value. [S05]

**WORKED EXAMPLE: NOMINAL ENERGY IS A CALCULATED SCENARIO, NOT AN EQUIPMENT CHOICE**

This illustrative case assumes a constant critical load of 150 W, an autonomy duty of 2 h, conversion efficiency of 0.90, usable fraction of 0.80 and end-of-life capacity fraction of 0.80. All five inputs are scenario assumptions, not measurements or recommended ratings. The contingency source supports setting recovery requirements before design, not any of these values. [S05]

The protected-load energy is 150 W × 2 h = 300 Wh. The combined denominator is 0.90 × 0.80 × 0.80 = 0.576. Nominal energy is therefore 300 Wh / 0.576 = 520.833 Wh, shown as 521 Wh after rounding to the nearest whole watt-hour. This illustrative arithmetic is dimensionally consistent because watts multiplied by hours give watt-hours and the denominator is dimensionless.

The result does not qualify a battery, UPS or protective setting. It does not establish transfer continuity, peak-current capability, recharge behaviour, thermal suitability or the survival of committed records. Those questions require separate evidence. If the illustrative load were 180 W with every other input unchanged, the same formula would require 625 Wh, showing why the measured duty must precede selection.

The calculation record in **Exhibit 15**, p. 61 preserves the inputs and their status. It deliberately places the assumed basis beside every value so a later reviewer cannot mistake the result for a supplier specification.

EXHIBIT 15 / ILLUSTRATIVE INPUTS AND CALCULATED NOMINAL ENERGY

## The nominal energy result exposes every assumption

| Input or output | Value | Unit | Basis |
|---|---|---|---|
| Critical load | 150 | W | Assumed constant duty |
| Autonomy | 2 | h | Assumed interruption duty |
| Conversion efficiency | 0.9 | ratio | Assumed path loss factor |
| Usable fraction | 0.8 | ratio | Assumed operating window |
| Remaining capacity | 0.8 | ratio | Assumed end-of-life condition |
| Nominal energy | 521 | Wh | Calculated scenario output |

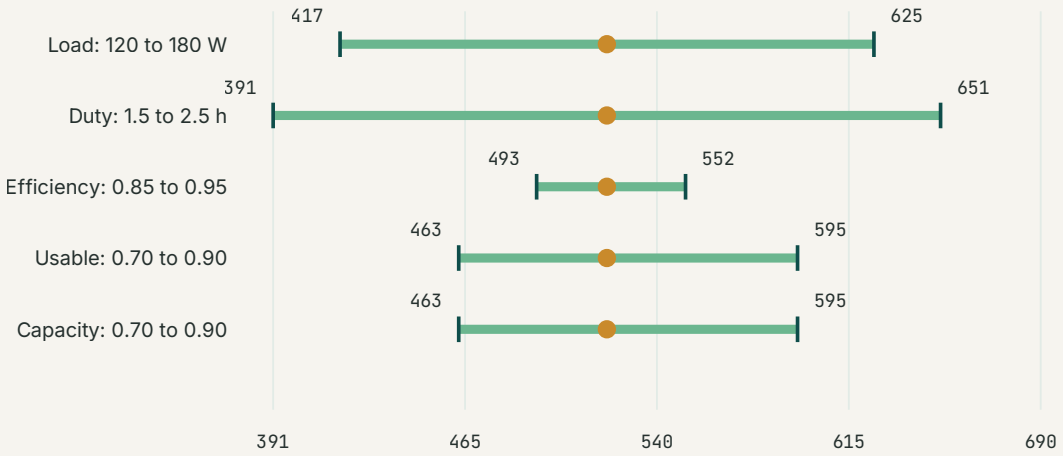
DERIVED [S05] Derived from illustrative inputs only. E_Wh = P_W*t_h/(efficiency*usable_fraction*remaining_capacity_fraction) = 150*2/(0.90*0.80*0.80) = 520.833 Wh. S05 supports contingency requirement setting, not the numerical inputs. The result is not a measured battery capacity or equipment selection.

Sensitivity analysis is more useful than an unexplained design margin. Change one declared input at a time and retain the others from the worked case. This shows which measurement or procurement uncertainty has the greatest effect on nominal energy. It also exposes asymmetric reasoning: a higher assumed efficiency reduces the calculated requirement, while a lower remaining capacity increases it. The direction must be checked before the scenario is used to justify a smaller device.

**Exhibit 16**, p. 62 applies that one-factor method to the load, duty and three denominator terms. The low and high ends are scenario outputs, not statistical confidence limits. They do not describe the probability of any condition, and they should not be combined into a new apparently precise rating.

EXHIBIT 16 / NOMINAL WATT-HOURS, ILLUSTRATIVE ONE-FACTOR SCENARIOS

## Duty and derating assumptions move the energy requirement




ILLUSTRATIVE [S05] Illustrative sensitivity around E = 150 2 / (0.90 0.80 * 0.80) = 520.833 Wh. Each row varies only the named input, with all other worked-case inputs fixed. Load range 120 to 180 W; duty 1.5 to 2.5 h; efficiency 0.85 to 0.95; usable fraction 0.70 to 0.90; remaining capacity 0.70 to 0.90. Low is the lowest resulting energy requirement and high the greatest, not an uncertainty interval. Values rounded to 0.1 Wh. S05 does not supply these inputs.

The exhibit identifies where better evidence would change the decision. If the duty is uncertain, resolve the operating requirement rather than hiding the choice in battery capacity. If conversion performance is uncertain, obtain applicable documentation or measure the path. If remaining capacity is uncertain, define the maintenance condition at which replacement is required. Sensitivity does not remove the need for a decision; it makes the consequence of that decision visible.

A fixed nominal battery can be examined with the inverse relationship, `t = E × η × u × c / P`, under the same simplified assumptions. That calculation estimates the autonomy associated with a scenario, not the time until every connected function stops gracefully. The operator still needs a reserve for the shutdown sequence and a rule for loads whose demand changes as services are shed. Capacity arithmetic is one gate in a wider qualification.

4.4

# Transfer behaviour decides whether the computer actually rides through

A battery-energy balance says nothing about the interruption experienced while supply changes source. The connected computer, network devices and storage must tolerate the actual transfer behaviour of the selected power path. GFT therefore recommends a controlled qualification that observes the complete protected duty

through loss of normal supply, operation on reserve, restoration and exhaustion. Electrical tests belong under the site electrical engineer, within equipment limits and an approved isolation and restoration procedure. This chapter specifies evidence, not switching instructions or protective settings. [S03]

Begin with a known operating state. Generate an identified sequence of synthetic production events, open the local console and record the software and firmware versions, active peripherals, database settings and starting ledger. Note which devices are expected to remain live and which are expected to restart. A pass criterion should describe the promised function, such as continued durable capture or an authorised manual fallback, rather than the vague survival of a glowing status lamp.

During each test, retain three aligned records. The electrical record should capture the supply event and relevant transfer observation using a method approved by the electrical engineer. The device record should show restarts, interface state, alarms and service readiness. The application record should show accepted event identities, local commits, rejected intake and any interval in which capture was unavailable. Aligning these records reveals whether an apparently successful power transfer still caused a network reset, a missed sensor observation or an uncertain application acknowledgement.

Test restoration as a separate event. Returning supply can restart devices, recharge batteries and initiate deferred work at the same time. The acceptance record should show whether the power path remains stable, whether services start in their required order and whether remote replay is prevented from displacing local capture. Cold-start behaviour also needs its own test because a system that remained live through a short transfer has not demonstrated that it can boot from an exhausted or disconnected state without external services.

Battery exhaustion is part of the same envelope. Establish what warning the system emits, which component receives it and what happens if the warning is late or absent. Observe whether the controlled stop completes within the available reserve and whether the last acknowledged event remains recoverable. Do not infer a shutdown window solely from nominal autonomy. The actual reserve depends on the measured duty and the state of the protected system at the time of the event.

SQLite's crash-testing documentation explains why software projects simulate power failures, but a supplier's simulation does not qualify the factory's physical power path. Conversely, a successful hardware transfer does not qualify the application's storage settings. The project needs both forms of evidence on the selected configuration. [S28] In particular, SQLite documents that a transaction committed in WAL mode with `synchronous=NORMAL` might roll back following power loss or a system crash. A continuous screen and a durable database are therefore separate acceptance results. [S16]

NIST recovery guidance links recovery planning with playbooks, testing and improvement. [S06] Record every unexpected restart and every unexplained ledger discrepancy as a defect to investigate, not as an acceptable cost of testing. Preserve failed trials alongside successful repeats and identify the configuration change between them. The resulting evidence should state the tested boundary precisely: which load,

which transfer events, which software versions and which record-loss criterion. No test conducted for this chapter is claimed here; these are proposed commissioning obligations.

4.5

## Controlled shutdown is a fall-back with dependencies

An orderly shutdown is a second continuity path, not a substitute for ride-through or crash recovery. It begins when a trustworthy signal indicates that the approved operating reserve is being consumed and ends only after the required record and shutdown state are recoverable. Network UPS Tools describes `upsmon` as a UPS monitoring and shutdown controller, which illustrates the role of software between the power device and the operating system. The presence of such a component does not prove that its signals, permissions, timing or application actions are correct on the selected installation. [S0401]

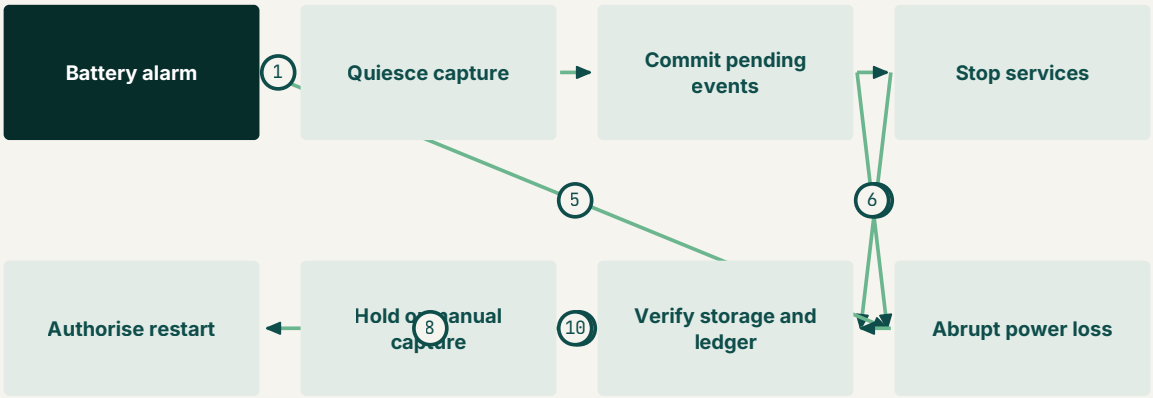
GFT proposes an explicit sequence. Validate the battery or shutdown signal and record its source. Stop admitting work that cannot be completed within the remaining reserve. Quiesce capture while preserving the identity of events already accepted. Commit the event and its pending-delivery record at the promised transaction boundary, then close dependent services in an order that prevents new writes from arriving after their destination has stopped. Record the reason, outcome and unresolved work before powering down the remaining compute duty.

The sequence needs deadlines and an owner. A service supervisor can provide configured stop commands and stop timeouts, but those settings must be read back and exercised rather than inferred from a default. [S21] Diagnostic evidence also needs an appropriate storage policy. `systemd-journald` distinguishes volatile and persistent journal storage, so a reboot may remove evidence that the team expected to inspect. [S22] Confirm the effective settings on the actual image and retain enough local evidence to explain why shutdown succeeded, timed out or never began.

**Exhibit 17**, p. 65 shows the controlled route and its abrupt-loss alternative. Both routes converge on storage and ledger verification before an authorised restart. The diagram is a proposed operating contract. It does not imply that a database transaction guarantees the behaviour of every storage device or that an operator may release production solely because services have started.

EXHIBIT 17 / ILLUSTRATIVE OPERATING SEQUENCE

## Orderly shutdown still ends at crash-recovery checks




1. Signal valid
4. Power returns
5. Late or absent signal
6. Reserve exhausted
8. Checks pass
9. Checks fail
10. Evidence restored

ILLUSTRATIVE [S0401] [S13] [S15] [S21] [S22] Illustrative GFT shutdown and recovery contract. The battery signal, capture quiesce, durable commit, service stop and retained shutdown state all require configuration evidence. The abrupt branch covers late or absent warning and reserve exhaustion. Neither route permits production release until storage and ledger checks pass. This is not a completed test or executable safety logic.

The abrupt branch must receive the same design attention as the graceful route. The warning cable can be disconnected, the monitoring process can fail, a service can exceed its stop deadline or the remaining reserve can be smaller than expected. SQLite states that atomic commit makes all changes within a transaction occur or none occur, but it also warns that the database is not immune to corruption. [S13] [S15] The recovery procedure must therefore inspect the complete storage path and reconcile identified events rather than assume that either a clean shutdown flag or a database opening successfully proves record survival.

Filesystem and device behaviour sit beneath that check. The ext4 documentation describes write barriers as enforcing on-disk ordering of journal commits when volatile disk write caches are used. [S29] A UPS can keep voltage present while an application completes a write, but it cannot by itself establish that the storage device honoured the required flush. Qualification should combine the actual power path, filesystem, storage hardware and database configuration, with unresolved assumptions retained in the acceptance record.

On restoration, hold new production acknowledgement until the required checks are complete. Verify storage integrity, effective durability settings, the latest committed identities and any incomplete shutdown state. If the evidence conflicts, retain the

affected work in recovery hold and invoke the approved manual record. An orderly shutdown reduces the number of uncertain cases; the crash-recovery path determines whether the remaining cases are detected and explained.

4.6

## Environmental and maintenance assumptions change usable autonomy

The calculated energy case becomes credible only inside an observed operating envelope. Collect the supplier's permitted conditions for the selected UPS, battery, computer, network devices and enclosure. Record where each item will sit, how air moves through it, what dust or other contamination it may encounter and how a technician can inspect and replace it. These are site-design inputs, not reasons to import a generic degradation rate or battery life into the model.

Measure the conditions that can alter the promised duty. Retain enclosure temperature, protected-load power, battery or shutdown status and relevant service alarms on a common maintenance record. Where the power system reports an estimated remaining time, preserve it as a device estimate and compare it with observed discharge behaviour during approved tests. Do not present that estimate as an independently measured autonomy. Missing or implausible telemetry should create an unqualified state rather than a reassuring zero.

Maintenance access is part of the continuity design. A battery that cannot be isolated and replaced under the site's approved procedure may turn a planned replacement into a longer interruption. Check physical clearances, connector identity, replacement compatibility, safe handling instructions and the availability of the required spare through the intended support channel. Assign the inspection, test, replacement and disposal responsibilities to named organisational roles without inventing local regulations or supplier intervals.

Ageing should be handled as a changing design input. Record installation identity, manufacture or batch information where supplied, commissioning date, replacements, test evidence and observed faults. Review whether the remaining-capacity assumption still reflects the condition at which the owner requires the duty to work. Network UPS Tools includes a battery-replacement notification among its defined UPS events, but an alarm definition is not proof that a particular battery needs replacement or that the site received the alarm. [S0401]

Environmental review also covers the record around the power device. Dust can contaminate more than the UPS, and an enclosure that protects one component may create a different thermal condition for another. Keep these interactions visible in the survey rather than assigning an unsupported numerical penalty. Where the combined supplier envelopes are incompatible with the intended location, change the location, protection or equipment through the responsible engineering process. A larger nominal battery does not resolve an unsuitable installation condition.

The maintenance record should close with a current qualification status. Distinguish as-designed capacity, factory documentation, site acceptance, latest observed test and overdue evidence. Repeat the relevant transfer and recovery checks after a

material change to the battery, power path, connected load or control software. This keeps the autonomy claim attached to the system that exists now rather than to the configuration that was originally purchased.

4.7

## Procurement must buy the promised protected function

The purchase schedule should describe the function to be proved, not only the equipment to be delivered. Attach the protected-duty card, measured load trace, outage duty, energy assumptions and acceptance tests to the request for quotation. Require the supplier to identify the exact hardware, firmware, monitoring software, licences, protected outlets and accessories included in its response. A family-level brochure cannot establish the behaviour of a different model, firmware version or outlet configuration.

GFT proposes an acceptance schedule with seven evidence rows. First, map every required device to a protected port and record excluded outlets. Second, document the power path and the applicable operating envelope. Third, demonstrate the agreed transfer cases under the approved test method. Fourth, demonstrate low-battery notification, controlled shutdown and abrupt-loss recovery. Fifth, demonstrate cold start with the declared external services unavailable. Sixth, prove replacement-battery compatibility and the maintenance procedure. Seventh, deliver the raw test traces, effective configuration and maintenance record needed for later review.

Each row needs a requirement, supplier response, evidence reference, witnessed result, exception and responsible approver. Mark unsupported claims as unresolved rather than converting them into assumed compliance. Where the supplier offers an estimated runtime, retain its load and battery conditions and compare them with the site's measured duty. Where it offers shutdown software, require the versions and permissions used in the demonstration. Acceptance should concern the complete chain that performs the record-keeping function.

Firmware and recovery assets belong in that chain. NIST's platform-firmware guidance addresses protection, detection and recovery against potentially destructive attacks, providing a reason to request authentic recovery images and controlled update behaviour rather than treating firmware as invisible infrastructure. [S04] This does not make an ordinary power outage a cybersecurity incident. It ensures that replacement and recovery do not depend on an unverified image or an undocumented external download.

Commercial evaluation should retain the cost of the whole duty. Request itemised equipment, integration, support, replacement, testing and maintenance offers with their dates and exclusions. No prices are asserted in this chapter, and this procurement framework is not financial, legal or tax advice. A lower initial equipment price can omit the switch, monitoring interface, replacement path or acceptance work required by the protected-duty card. Compare only offers that address the same evidence schedule.

The boundary with wider energy design must remain explicit. This chapter determines the protected digital load, its interruption behaviour and its record-survival evidence. Generation sizing, solar yield, recharge strategy and resource-efficiency optimisation belong to the wider energy work. The load register and sensitivity model can supply inputs to that work without pretending that the digital UPS case is a factory energy system.

Before approval, reconcile the delivered system with the original duty. Confirm that included loads remain within the measured and qualified envelope, excluded loads cannot silently migrate onto protected outlets, and the operator can see the current operating mode. Retain the approved manual fallback and recovery hold for failures outside that envelope. Expansion to powered inspection or another digital duty should reopen the load, transfer, energy and recovery assessments rather than inherit acceptance by association.

GFT's recommendation is deliberately narrow: protect the smallest complete record-keeping chain whose continuity has operational value, measure it under the promised duty, and buy evidence of its behaviour through transfer, exhaustion and restoration. That discipline gives the edge-architecture and energy-sizing work a transparent input. More importantly, it gives the factory a power boundary it can explain, test and maintain without claiming that production itself has become uninterruptible.

*Protect the duty, measure the load and prove the record survives the power path.*

**CHAPTER SOURCE NOTES**

S03 NIST. 2023.
S05 NIST. 2010.
S28 SQLite. 2026-10-30.
S16 SQLite. 2026-10-30.
S06 NIST. 2016.
S0401 Network UPS Tools (NUT) Project. 2026-10-28.
S21 systemd project. 2026-10-30.
S22 systemd project. 2026-10-30.
S13 SQLite. 2026-10-30.
S15 SQLite. 2026-10-30.
S29 Linux kernel project. 2026-10-30.
S04 NIST. 2018.

GREEN FUTURE TECH
LIMITED / HONG KONG SAR

05

PART II: KEEP THE LOCAL RECORD ALIVE / CONNECT LOCALLY

# A working local network should outlast a failed backhaul

*Local name resolution, switching, identity and observability need their own continuity plan when external connectivity disappears.*

CHAPTER 05
DECISION TEST

Local name resolution, switching, identity and observability need their own continuity plan when external connectivity disappears.

A factory network on a weak grid remains useful during a wide-area outage only if every service required to address, authenticate, route and observe local traffic lives inside the plant boundary and shares a protected power path [S03] [S18]. Many industrial sites protect an edge server or wide-area router with an uninterruptible power supply, yet lose internal operations when the backhaul drops or a feeder trips because floor switches, address leases, name resolvers, broker credentials or operator consoles depend on unbacked distribution boards or external cloud endpoints [S03] [S05]. Designing for intermittent connectivity starts on the local area network (LAN) before it touches a wide-area carrier. When local switching, naming, time distribution, transport sessions and maintenance access function without an internet route, the plant can record production events, display local alarms and coordinate floor workstations for hours or days until upstream links return [S06] [S08].

This chapter defines the controls that keep a plant network operational during external disconnection and disciplined during link recovery. Following the power boundary of Chapter 4 and preceding the storage rules of Chapter 6, it maps external and electrical dependencies on the LAN [S03] [S18], replaces advertised carrier bandwidth with useful application goodput measured at durable commit [S12] [S24], compares transport protocols under channel loss [S09] [S12] [S19] [S24], evaluates multi-carrier diversity through common-cause failure paths [S05] [S18], bounds retry and backlog admission during recovery [S06] [S19] [S20], establishes controlled offline administration under cached identity and bounded clock uncertainty [S23] [S27], and specifies a fault-injection qualification protocol before production cutover [S06] [S28].

5.1

## Local operation begins with local dependencies

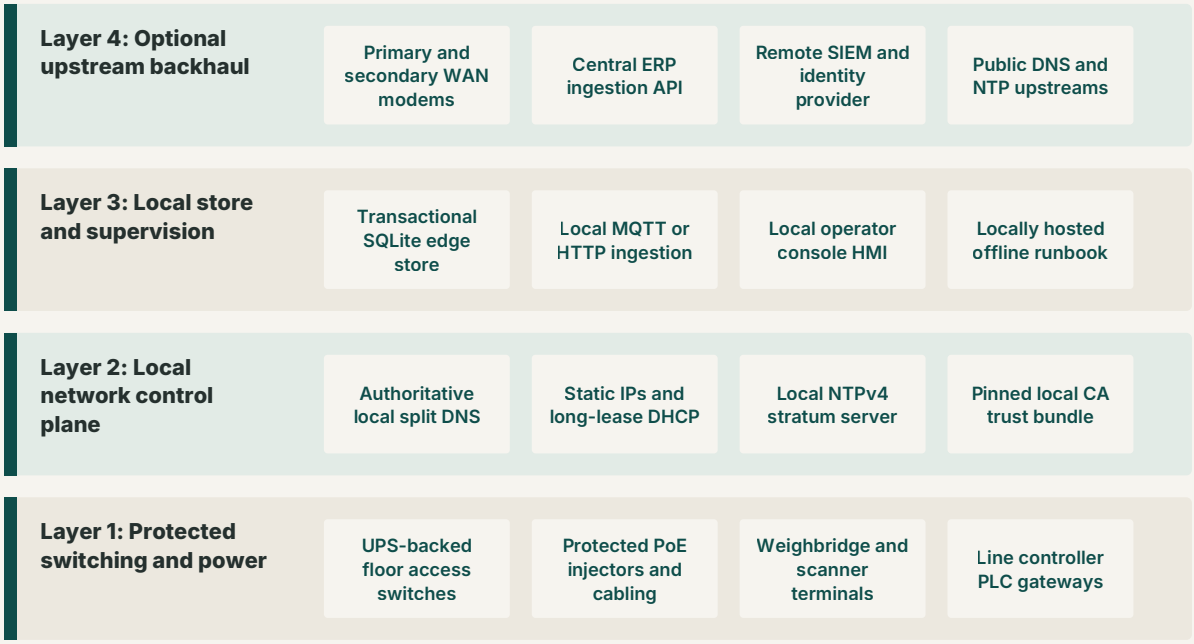
Industrial digital systems frequently fail during a backhaul outage not because the application logic requires remote compute, but because basic network plumbing relies on external reachability or unprotected floor power [S03] [S18]. In a typical commercial deployment, a barcode terminal or programmable logic controller gateway resolves an edge broker hostname through a domain name system (DNS) forwarder pointing to a public resolver, obtains its dynamic host configuration protocol (DHCP) lease from an office router, validates transport certificates against an online revocation responder, and authenticates operator sessions through a remote identity provider [S03] [S23]. While the wide-area link is healthy and every distribution board is energised, this chain appears local to the shop floor. When upstream fibre is cut or a brownout drops an intermediate switch cupboard, the chain breaks inside the factory even though the edge server remains powered [S03] [S05].

Keeping the local network useful requires treating the plant LAN as an autonomous survivorship zone with its own power, addressing, naming, time and session authorities [S03] [S11]. NIST SP 800-82 Rev. 3 emphasises that operational technology networks must be segmented and engineered so that loss of external connectivity or corporate IT services does not impair essential local monitoring and control [S03]. Similarly, RFC 7228 notes that constrained networks combine node limitations with channel loss, unpredictable bandwidth and dynamic topology changes [S18]. As

shown in **Exhibit 18**, p. 71, every layer between the capture terminal and the local edge store must remain complete inside the site boundary when wide-area carriers are dark.

EXHIBIT 18 / ILLUSTRATIVE ARCHITECTURE LAYERS

## Local work can survive the loss of an external route




ILLUSTRATIVE [S03] [S18] Illustrative plant LAN dependency stack derived from NIST SP 800-82 Rev. 3 zone segmentation principles (S03) and RFC 7228 constrained-network topology considerations (S18). Every component in Layers 1 to 3 must operate with Layer 4 physically disconnected.

### 5.2 DISTINGUISH A BACKED-UP ROUTER FROM A PROTECTED LOCAL AREA NETWORK

A common procurement mistake on industrial sites is placing an uninterruptible power supply under the wide-area router and edge server while leaving intermediate access switches on raw lighting circuits [S03] [S05]. Under Dependency Inference, a path is only as available as the minimum of its series elements. If a line scanner connects through an unpowered 8-port floor switch on a pillar, or through a wireless access point whose power-over-Ethernet (PoE) injector sits on an unprotected sub-board, the edge server is unreachable during a mains sag even if its own battery holds for 60 minutes (illustrative site configuration) [S03] [S05].

Mapping the physical LAN path requires tracing every active device between the data source and the local store: media converters, PoE midspans, floor switches, wireless bridges and industrial firewalls [S03]. Each intermediate device must share the protected DC or UPS circuit of the service class it supports (Chapter 3) or be bypassed by a direct home-run cable to a protected switch [S03] [S05]. Where a floor switch serves both essential terminals and non-essential office loads, shedding non-essential PoE ports during mains loss prevents peripheral devices from draining the switch backup supply (illustrative operational rule) [S03].

### 5.3 KEEP ADDRESSING, NAMING AND SESSION ESTABLISHMENT STRICTLY INTERNAL

Even when every switch remains powered, software dependencies can sever local communication after a link failure or cold reboot. Four control-plane traps must be eliminated by design [S03] [S11] [S23]:

1. **Forwarding-only DNS resolution:** If edge clients connect to `broker.plant.example` and the local DNS resolver acts only as a forwarding cache, cached records expire once their time-to-live (TTL) elapses [S03]. A 300-second TTL (illustrative setting) means that five minutes after the backhaul drops, any restarting client fails with a name-resolution error even though the broker sits on the same VLAN [S03]. The plant network must run a local authoritative DNS zone for internal service names, backed by static host entries on critical gateways [S03].
2. **Short DHCP leases without local authority:** If dynamic addressing is served across a wide-area virtual private network (VPN) or by an unbacked office router, clients whose leases pass the 50 percent renewal (T1) and 87.5 percent rebinding (T2) marks drop their IP addresses upon lease expiry [S03]. Critical servers, brokers, controllers and consoles must use static IP assignments in the site address register, and mobile pools must be served by a protected local DHCP instance with 24 to 72 hour leases (illustrative target) [S03] [S05].
3. **Synchronous online certificate revocation checks:** Under TLS 1.3 (RFC 9846) [S23], clients or brokers configured to require a live Online Certificate Status Protocol (OCSP) responder or external certificate revocation list (CRL) abort the handshake when the WAN is down [S23]. Local TLS must validate against a locally pinned certificate authority bundle without mandatory online OCSP fetches [S03] [S23].
4. **Cloud-redirected operator login:** A local web console that redirects the browser to an external single sign-on domain or fetches scripts from public content delivery networks fails during the very outage the operator needs to diagnose [S03] [S06]. The local console must authenticate against a local credential store and serve all user-interface assets locally [S03].

> **GFT NOTE: THE PULL-THE-PLUG COLD-START RULE**
>
> A local network is independent of its backhaul only if it passes a cold-start test while the wide-area demarcation cable is physically unplugged [S06]. Networks often survive a warm disconnection because ARP tables, DNS caches, DHCP leases and TLS session tickets sit in volatile memory [S23] [S24]. If a power transfer restarts floor switches and edge clients while the WAN remains down, those caches vanish. GFT architectural reviews require that every local capture, broker and supervision service boot from cold power and establish normal local operation with the WAN interface physically disconnected [S03] [S06].

5.4

# Measure useful application capacity rather than advertised bandwidth

Carrier quotations state nominal peak bit rates, such as a 10 megabits per second (Mbit/s) wireless link or a 50 Mbit/s shared fibre circuit (illustrative offerings) [S18]. On a weak grid or across contended last-mile infrastructure, advertised link bandwidth is a poor predictor of whether an edge store can clear its backlog before the next outage window [S08] [S18]. RFC 7228 warns that constrained networks exhibit limited, unpredictable bandwidth alongside lossy channels [S18]. Engineers must separate raw channel capacity from **effective application goodput**: the volume of unique, schema-valid application payload durably committed and acknowledged by the receiving endpoint per unit of elapsed wall-clock time [S12] [S13] [S24].

Between reading queued records from a local SQLite store and receiving remote confirmation of durable acceptance, a TCP stream (RFC 9293) must complete a three-way handshake, negotiate TLS 1.3 parameters (RFC 9846) and open its congestion window [S13] [S23] [S24]. Packet loss or wireless fading forces TCP to retransmit segments, reduce its congestion window and stall delivery until in-order sequence continuity is restored [S24]. If the bearer drops before the receiving application commits the batch and returns an application-layer confirmation (an HTTP `2xx` status under RFC 9110 or an MQTT `PUBACK` under OASIS MQTT v5.0), the partially transmitted payload must be sent again when the link recovers [S12] [S19] [S24].

### 5.5 DEFINE EFFECTIVE UPLOAD GOODPUT ACROSS INTERRUPTION CYCLES

Useful upload capacity over an observation window $T_{\text{window}}$ (such as a 1-hour recovery interval or 24-hour operating day, illustrative intervals) is defined at the application commit boundary [S12] [S13]:

$$G_{\text{eff}} = \frac{B_{\text{committed}}}{T_{\text{window}}}$$

Here $B_{\text{committed}}$ is the total number of unique application payload bytes (excluding duplicate replays, aborted partial transfers, TLS record headers, TCP/IP headers and link framing) for which the local sender receives an explicit, durable commit acknowledgement during $T_{\text{window}}$ (illustrative formula convention) [S12] [S19] [S23] [S24]. This metric penalises four sources of waste hidden by interface counters [S18] [S24]:

- **Handshake tax:** Link flaps force new DNS lookups, TCP handshakes and TLS 1.3 negotiations (1 round-trip time under RFC 9846) before application payload moves [S23] [S24].
- **Transport retransmission and head-of-line blocking:** TCP guarantees a reliable, in-order byte stream [S24]; packet loss generates duplicate segments that register as interface traffic without advancing application state [S24].
- **Uncommitted partial transfers:** If a 5-megabyte inspection image (illustrative payload) fails at 90 percent completion without byte-range resumption, those 4.5 megabytes contribute zero to $B_{\text{committed}}$ [S12] [S24].

- **Application-layer duplicate delivery:** Under at-least-once delivery contracts (MQTT QoS 1 in OASIS MQTT v5.0 or retried HTTP `PUT` requests in RFC 9110), lost acknowledgements trigger retransmission of already-stored records, consuming uplink capacity even when idempotent deduplication preserves data integrity [S12] [S19].

### 5.6 KEEP SIX MEASUREMENT SERIES DISTINCT DURING SITE SURVEYS

No carrier performance figures are asserted in this paper for any region or operator, because links vary by tower load, weather, power state and time of day [S01] [S18]. Before freezing batch sizes or backlog clearance budgets on a target site, the engineering team must instrument a local probe across working shifts and outage windows to record six distinct time series [S18] [S24]:

1. **Physical and link-layer modulation rate (bit/s):** Negotiated speed between the modem and first carrier hop, exposing signal fade or base-station power curtailment [S18].
2. **Network-layer throughput (bit/s):** Total IP bytes per second across the WAN interface, including framing overhead and retransmissions [S24].
3. **Round-trip latency and jitter (milliseconds):** Active probing to the target ingestion endpoint, bounding throughput when in-flight windows are small [S09] [S24].
4. **Packet loss and connection reset rate (percentage and resets/hour):** Lost probe fraction and frequency of TCP `RST` or bearer drops, governing batch size between interruptions [S18] [S24].
5. **Useful application goodput ($G_{\text{eff}}$, bytes/s):** Net rate of durably acknowledged unique payload bytes over 1-minute, 15-minute and 60-minute rolling windows (illustrative buckets) [S12] [S19].
6. **End-to-end application commit delay (seconds):** Elapsed time from local event creation to confirmed remote commit at the 50th, 95th and 99th percentiles (illustrative percentiles) [S12] [S32].

> **METHOD**
>
> **Site link characterisation protocol before production cutover**
>
> Run an automated synthetic replay harness on the site gateway for 7 to 14 consecutive days (illustrative qualification window) [S06] [S18], transmitting representative batches (1-kilobyte production records and 250-kilobyte diagnostic attachments, illustrative sizes) to a staging endpoint performing real TLS 1.3 termination and durable disk commits [S13] [S23]. Correlate $G_{\text{eff}}$ with local mains power logs. On weak grids, wide-area goodput frequently drops 15 to 45 minutes after a neighbourhood power cut (illustrative failure window) as surrounding users shift onto mobile data while cell towers throttle transmitters on battery reserve [S01] [S18]. Backlog clearance calculations (`C02` in the calculation register) must use the lower-decile post-outage goodput measured during these degraded windows, never the peak link rate [S18].

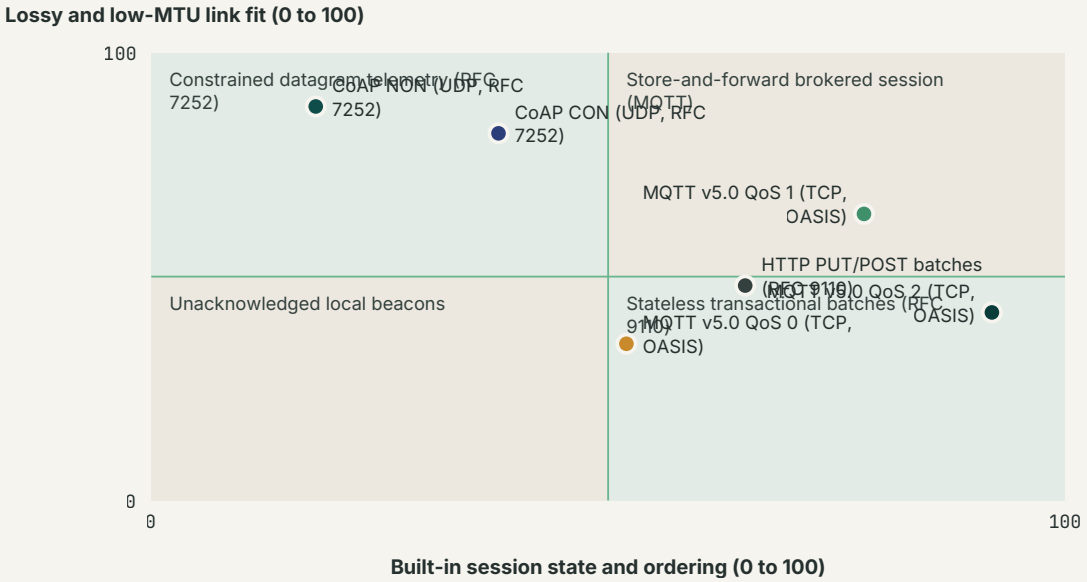
5.7

# The transport protocol must fit device and failure behaviour

No single transport protocol is universally superior across both the shop floor and the wide-area backhaul [S09] [S12] [S19]. Each protocol embodies specific assumptions about transport reliability, framing overhead, state retention, message ordering and congestion control [S09] [S12] [S19] [S24]. A protocol suited to 32-byte sensor readings on a local mesh can perform poorly when draining 50,000 queued production records over a flapping cellular link (illustrative workload contrast) [S09] [S18] [S19]. As mapped in **Exhibit 19**, p. 75, selection depends on the constraint and loss profile of the network path and the delivery contract required between endpoints [S09] [S12] [S19] [S24].

EXHIBIT 19 / RELATIVE DESIGN SCORE (0 TO 100)

## Protocol selection depends on the required delivery contract




SOURCED [S09] [S12] [S19] [S24] Source-based comparison of protocol mechanics specified in RFC 7252 (CoAP, S09), RFC 9110 (HTTP Semantics, S12), OASIS MQTT v5.0 (S19) and RFC 9293 (TCP, S24). Coordinates represent relative design placement (0 to 100); exact wire specifications and GFT outbox recommendations are tabulated in the text.

### 5.8 COMPARE WIRE MECHANICS, RETRY SEMANTICS AND OVERHEAD FROM PRIMARY SPECIFICATIONS

Engineers must separate what each standard guarantees on the wire from what the surrounding storage architecture must provide [S09] [S12] [S13] [S19]:

- **Constrained Application Protocol (CoAP, RFC 7252):** Designed for constrained nodes and lossy channels [S09] [S18], baseline CoAP binds to UDP with a 4-byte fixed header, a 0 to 8 byte Token, options and payload [S09]. CoAP defines Confirmable (`CON`), Non-confirmable (`NON`), Acknowledgement (`ACK`) and Reset (`RST`) messages matched by a 16-bit Message ID [S09]. For `CON` messages, RFC 7252 specifies stop-and-wait retransmission with exponential backoff under default parameters (`ACK_TIMEOUT` = 2 s, `ACK_RANDOM_FACTOR` = 1.5, `MAX_RETRANSMIT` = 4, `NSTART` = 1, `DEFAULT_LEISURE` = 5 s, `PROBING_RATE` = 1 byte/s, `MAX_TRANSMIT_SPAN` = 45 s and `MAX_TRANSMIT_WAIT` = 93 s) and recommends fitting messages inside a

single IP packet (1,280-byte IP MTU, 1,152-byte message bound and 1,024-byte payload upper bound, or 127-byte IEEE 802.15.4 frames) [S09]. CoAP fits constrained local radio links, but NSTART = 1 and the 1,024-byte payload bound limit bulk WAN replay without block-wise tuning [S09].

- **Message Queuing Telemetry Transport (MQTT v5.0, OASIS Standard):** Operating over TCP/IP (RFC 9293), MQTT uses a 2-byte minimum fixed header (1 control byte plus a 1 to 4 byte Variable Byte Integer encoding Remaining Length up to 268,435,455 bytes) [S19] [S24]. OASIS MQTT v5.0 defines three Quality of Service levels: QoS 0 ("At most once"), QoS 1 ("At least once", via PUBLISH and PUBACK where duplicates can occur) and QoS 2 ("Exactly once", via PUBLISH, PUBREC, PUBREL, PUBCOMP) [S19]. It also defines Clean Start and Session Expiry Interval for session state, Receive Maximum to cap concurrent unacknowledged QoS 1/2 packets, and Topic Alias to compress repeated topic strings into a 2-byte integer [S19].
- **HTTP Semantics (RFC 9110):** HTTP provides stateless request-response semantics over TCP and TLS [S12] [S23] [S24]. RFC 9110 specifies that GET, HEAD, OPTIONS and TRACE are safe, while PUT, DELETE and all safe methods are idempotent: a client can safely repeat an unacknowledged PUT batch after a connection reset [S12]. POST is not idempotent unless keyed by an application UUID [S12] [S31]. HTTP also supports compression (Content-Encoding) and byte-range resumption (Range) [S12].

| Protocol | Framing | Ordering and delivery | Flow control | Weak-grid strengths | Misapplication traps |
|---|---|---|---|---|---|
| **CoAP** [S09] | UDP; 4-byte header, 0 to 8 byte Token, 1,024-byte payload bound [S09] | CON/ACK stop-and-wait; no cross-datagram ordering [S09] | ACK_TIMEOUT = 2 s, MAX_RETRANSMIT = 4, NSTART = 1 [S09] | No TCP head-of-line blocking on lossy radio links [S09] [S18] | NSTART = 1 and 1,024-byte bound stall bulk replay [S09] |
| **MQTT v5.0** [S19] | TCP (RFC 9293); 2-byte min header plus 1 to 4 byte length [S19] [S24] | Ordered per TCP stream; QoS 0, QoS 1 (PUBACK) or QoS 2 [S19] [S24] | Receive Maximum, Session Expiry Interval, Topic Alias [S19] | Persistent pub/sub session; low overhead; local queuing [S19] [S20] | PUBACK confirms broker receipt, not database commit [S13] [S19] |
| **HTTP** [S12] | TCP (RFC 9293) and TLS 1.3 (RFC 9846); headers and body [S12] [S23] [S24] | Ordered per stream; explicit status codes (2xx, 4xx, 5xx) [S12] [S24] | Idempotent PUT/DELETE/GET; Range resumption [S12] | End-to-end confirmation from target API; range tooling [S12] | Verbose per-record headers; POST duplicates without UUID [S12] [S31] |

### 5.9 PROTOCOL DELIVERY IS HOP-BY-HOP, NOT END-TO-END DURABLE INGESTION

Assuming that MQTT QoS 1/2 or CoAP `CON` guarantees end-to-end database durability is an architectural error [S09] [S13] [S19]. Both protocols define delivery across a single hop [S09] [S19]. In MQTT, the broker returns `PUBACK` as soon as it accepts the packet under its local configuration [S19]. If Eclipse Mosquitto runs with `persistence false` in `mosquitto.conf`, or flushes state to `mosquitto.db` only on periodic `autosave_interval` cycles, a power cut on the broker host immediately after `PUBACK` loses the message after the publisher has discarded its copy [S19] [S20]. Even when `mosquitto.db` is persisted, a bridge worker that crashes after acknowledging the broker packet but before committing to its SQLite store drops the record [S13] [S19] [S20].

GFT's reference architecture separates local collection from wide-area synchronisation (GFT architectural recommendation) [S08] [S13]. Inside the plant LAN, use MQTT v5.0 QoS 1 over protected Ethernet/Wi-Fi (with Mosquitto `persistence true` and an explicit `persistence_location`) or CoAP `CON` for constrained sensor subnetworks [S09] [S19] [S20]. Across the intermittent WAN, commit records atomically to a local SQLite outbox (Chapter 6) and drain that outbox using batched, compressed, idempotent HTTP `PUT` requests (RFC 9110) or MQTT v5.0 QoS 1 batches paired with an explicit application-level commit receipt naming the committed event UUIDs [S12] [S13] [S19] [S31].

5.10

## A second carrier is useful only if its failure path differs

When a single wide-area connection proves unreliable, the standard response is to buy a second subscription from another operator or insert a second cellular SIM into a dual-WAN router [S05]. Decision-makers are often shown a naive calculation: if Carrier A and Carrier B each offer 95 percent standalone availability, combining them in active-passive failover supposedly yields $1 - (0.05 \times 0.05) = 99.75\%$ availability, cutting expected downtime from 36 hours per month to 1.8 hours per month (illustrative naive calculation) [S05].

On a weak grid, that multiplication is invalid because the two carriers are not statistically independent [S05] [S18]. Reliability engineering separates independent random faults from **common-cause failures** where a single physical or electrical event disables both paths simultaneously [S05]. NIST SP 800-34 Rev. 1 emphasises that telecommunications contingency planning must evaluate physical routing diversity, shared central offices, local loop dependencies and backup power at intermediate facilities [S05]. Two fibres in the same street trench fail under the same excavator bucket; two mobile operators on the same roadside mast fail together once a prolonged feeder outage exhausts the mast's shared rectifier batteries [S01] [S05] [S18].

### 5.11 AUDIT FIVE COMMON-CAUSE FAILURE DOMAINS BEFORE BUYING A BACKUP LINK

Before contracting a secondary backhaul, engineers must require written path disclosures and conduct a physical walkthrough across five failure domains [S03] [S05] [S18]:

1. **On-site termination and power boundary:** Terminating both links in one dual-SIM router fed by a single power brick creates a single point of failure; separate primary and secondary modems onto distinct interfaces powered from fused branches of the essential UPS bus [S03] [S05].
2. **Last-mile physical exit and civil route:** Trace wired conduits from the patch panel to the street chamber; cables sharing a trench, culvert or pole line are vulnerable to a single civil cut [S05].
3. **Radio base station and tower power reserve:** Record the physical cell ID (PCI) or tower coordinates to verify whether operators co-locate on the same mast and how long tower batteries last during blackouts [S01] [S05] [S18].
4. **Regional backhaul and metropolitan aggregation:** Retail service providers frequently lease wholesale metro transport from the same tier-1 fibre carrier or exchange point (IXP) [S05].
5. **Logical and contractual failure modes:** Shared data caps, simultaneous quota exhaustion or reliance on the same upstream DNS resolver or VPN concentrator can silence both links together [S03] [S05].

As shown in **Exhibit 20**, p. 78, common-cause downtime dominates combined wide-area unavailability whenever physical or power paths overlap.

EXHIBIT 20 / HOURS PER 720-HOUR MONTH (ILLUSTRATIVE)

## Common-cause failures defeat naive dual-carrier availability math

| Diversity scenario | Carrier A (h/mo) | Carrier B (h/mo) | Shared fault (h/mo) | Naive math (h/mo) | Actual dual-WAN (h/mo) |
|---|---|---|---|---|---|
| 1. Dual SIMs, same mast and grid feeder | 36.0 (5.0%) | 36.0 (5.0%) | 24.0 (3.33%) | 1.8 (0.25%) | 24.2 (3.36%) |
| 2. Two retail ISPs in shared street trench | 21.6 (3.0%) | 28.8 (4.0%) | 14.4 (2.00%) | 0.9 (0.12%) | 14.6 (2.02%) |
| 3. Fibre + cellular on shared metro backbone | 21.6 (3.0%) | 36.0 (5.0%) | 7.2 (1.00%) | 1.1 (0.15%) | 7.9 (1.09%) |
| 4. Terrestrial fibre + UPS-backed satellite | 21.6 (3.0%) | 14.4 (2.0%) | 0.7 (0.10%) | 0.4 (0.06%) | 1.1 (0.15%) |

DERIVED [S05] [S18] Illustrative reliability derivation applying NIST SP 800-34 Rev. 1 telecommunications diversity principles (S05) and RFC 7228 channel constraints (S18). Actual combined downtime equals shared downtime plus independent simultaneous downtime. All values are illustrative scenarios.

**WORKED EXAMPLE: WHY TWO SIM CARDS IN ONE ROUTER STILL LEAVE A PLANT DARK DURING A FEEDER OUTAGE**

Consider a hypothetical bottling plant where the utility feeder trips for 4 hours twice a week (illustrative scenario) [S01]. The plant installs a dual-SIM router with Operator A (95% standalone monthly availability, or 36 hours of downtime in a 720-hour month) and Operator B (95% standalone availability, 36 hours of downtime) [S05]. Management expects combined WAN downtime of $0.05 \times 0.05 = 0.25\%$, or 1.8 hours per month (illustrative naive expectation).

A site walkthrough reveals that both operators serve the estate from the same roadside mast 1.2 kilometres away, fed by the same distribution feeder as the factory [S05] [S18]. The mast's rectifier batteries hold for 90 minutes before shutting down both base stations. Out of each operator's 36 hours of monthly unavailability, 24 hours (3.33% of the month) occur simultaneously when feeder outages exceed 90 minutes, while only 12 hours (1.67%) represent independent core or radio faults (illustrative decomposition). Decomposing total unavailability $U$ in this illustrative scenario into $U_{\text{common}} = 0.0333$ and independent components $U_{A,\text{ind}} = 0.0167$ and $U_{B,\text{ind}} = 0.0167$, true combined unavailability is $U_{\text{combined}} = U_{\text{common}} + (1 - U_{\text{common}}) \times U_{A,\text{ind}} \times U_{B,\text{ind}} \approx 0.0333 + (0.9667 \times 0.0167 \times 0.0167) \approx 0.0336$, or 24.2 hours per month (illustrative result) [S05].

Instead of cutting downtime twentyfold (to 1.8 hours), the second SIM trims downtime by less than 33% (to 24.2 hours) and fails during every extended blackout [S05]. Genuine diversity requires pairing terrestrial links with an orthogonal medium independent of local utility transformers, such as a satellite terminal powered from the plant's own essential UPS bus, while sizing the local store (Chapter 6) to survive 100% WAN loss [S05] [S08].

5.12

## Retry policy must preserve capacity for urgent work

The most fragile moment for a weak-grid network is not the middle of a disconnection, when local services run inside the LAN, but the first ten minutes after a degraded wide-area link returns [S06] [S18]. During a four-hour outage, edge gateways, floor terminals and diagnostic agents accumulate thousands of queued production events, telemetry points, journal entries and images (illustrative backlog) [S08] [S22]. If every process detects link restoration simultaneously and flushes its backlog at maximum concurrency, the resulting **reconnection and replay storm** saturates uplink buffers, inflates round-trip latency into multiple seconds, triggers TCP retransmissions and TLS timeouts, and starves live high-priority alarms [S06] [S19] [S24]. NIST SP 800-184 emphasises that recovery planning must establish prioritised restoration sequences so critical operational functions recover before secondary workloads [S06]. On a factory network, that requires traffic prioritisation, bounded in-flight concurrency and jittered exponential backoff [S09] [S19] [S20].

### 5.13 SEPARATE TRAFFIC LANES BY OPERATIONAL URGENCY AND PAYLOAD SIZE

When the wide-area link transitions from `DOWN` to `DEGRADED` or `UP`, outbound traffic from the plant edge must pass through a scheduler or broker bridge enforcing four priority lanes (illustrative GFT scheduling policy aligned with Chapter 3 service classes) [S03] [S06]:

- **Lane P0 (Live safety, security and plant health alarms):** Compact state transitions, threshold breaches and heartbeat summaries (under 1 to 2 kilobytes per message, illustrative size), prioritised ahead of all queued history so remote supervisors see current plant state within seconds [S03] [S06].
- **Lane P1 (Current and queued transactional production records):** Idempotent records of weighbridge tickets, batch completions and operator sign-offs, drained in watermarked batches once Lane P0 is clear [S12] [S13].
- **Lane P2 (Historical environmental telemetry and diagnostic summaries):** High-frequency sensor series downsampled or batched into compressed chunks [S18].
- **Lane P3 (Bulk binary blobs, camera images, full logs and firmware pulls):** Multi-megabyte payloads scheduled only when P0 and P1 queues sit below low-water marks, using resumable HTTP `Range` requests (RFC 9110) so link flaps do not waste transferred bytes [S12] [S22].

### 5.14 BOUND IN-FLIGHT CONCURRENCY AND APPLY JITTERED BACKOFF

Pipelining hundreds of unacknowledged requests across a thin wide-area link merely moves the queue from the local edge disk (where it can be inspected and reprioritised) into carrier router buffers (where it causes head-of-line blocking) [S18] [S24]. Edge configurations must enforce explicit concurrency bounds [S09] [S19] [S20]:

- **Bound in-flight messages on MQTT bridges:** OASIS MQTT v5.0 provides the `Receive Maximum` property to cap concurrent unacknowledged QoS 1/2 `PUBLISH` packets [S19]. In Eclipse Mosquitto (`mosquitto.conf`), `max_inflight_messages` defaults to `20` (`max_inflight_bytes` defaults to `0`, and `max_queued_messages` defaults to `1000`) [S20]. On a lossy WAN bridge, tightening `max_inflight_messages` to `1` to `5` with an explicit byte cap (illustrative tuning) prevents a stalled batch from filling kernel TCP buffers [S19] [S20] [S24].
- **Apply exponential backoff with random jitter:** RFC 7252 sets CoAP `NSTART = 1` by default and chooses initial Confirmable (`CON`) retransmission timeouts randomly between `ACK_TIMEOUT` (2 seconds) and `ACK_TIMEOUT * ACK_RANDOM_FACTOR` (3 seconds), doubling up to `MAX_RETRANSMIT` (4 retries) [S09]. Custom HTTP or TCP reconnect loops must apply equivalent jittered backoff so rebooting devices do not hammer gateways in lockstep [S09] [S12].
- **Gate bulk replay behind a link stability soak timer:** Because modems often flap repeatedly during carrier fault clearing [S18], the edge synchroniser should verify DNS and time, send a P0 status beacon, and observe a configurable soak window (such as 30 to 120 seconds of stable reachability, illustrative policy) before opening P1/P2 backlog drain, as shown in **Exhibit 21**, p. 81 [S06] [S23] [S24].

EXHIBIT 21 / RELATIVE RECOVERY STEP INDEX (0 TO 6)

## Backhaul restoration is a sequence of dependencies

0

**Bearer and IP route restored**

WAN modem re-establishes bearer; router verifies gateway without disturbing LAN.

1

**Control-plane check (DNS and NTS)**

Gateway validates upstream DNS and checks NTS/NTP offset without stepping clocks.

2

**Jittered TLS 1.3 session setup**

Coordinator applies random jitter and completes TCP and TLS 1.3 negotiation.

3

**Lane P0 live alarms admitted**

Current plant state, active safety alarms and queue depth metrics transmit first.

4

**Link stability soak probe**

Coordinator observes packet loss and RTT over a 30 to 120 s soak window.

5

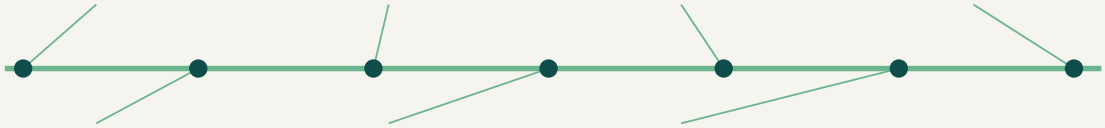
**Lane P1 throttled outbox drain**

Committed SQLite outbox batches drain using bounded in-flight windows (1 to 5).

6

**Lanes P2 and P3 background sync**

Historical sensor series and resumable HTTP byte-range blobs transfer last.



ILLUSTRATIVE [S06] [S23] [S24] Illustrative recovery choreography synthesised from NIST SP 800-184 recovery prioritisation (S06), RFC 9293 TCP connection semantics (S24) and RFC 9846 TLS 1.3 handshake requirements (S23). Numeric positions represent dependency order (0 to 6) rather than fixed wall-clock minutes.

5.15

# Local administration must remain controlled while offline

When a factory is cut off from corporate headquarters and cloud portals for thirty-six hours, local technicians still need to inspect queue depths, replace a failed scanner, rotate logs or restart a wedged PLC connector [S03] [S06]. If the security architecture requires real-time authorisation from a cloud identity provider or remote bastion host, the site team faces an unacceptable choice during an outage: let production stall, or bypass security controls by enabling shared root passwords and opening unauthenticated management ports [S03]. NIST SP 800-82 Rev. 3 requires that operational technology maintenance and access controls account for isolated or degraded network modes without sacrificing least privilege, authentication or auditability [S03]. Offline administration must therefore be engineered as a permanent, auditable operating mode [S03] [S06].

### 5.16 DESIGN AN AUDITABLE LOCAL MAINTENANCE PATH AND OFFLINE RUNBOOK

A disciplined offline administration architecture rests on four local controls [S03] [S06] [S22] [S23]:

1. **Pre-provisioned, role-scoped local maintenance accounts:** Provision individual named accounts or locally verifiable SSH/TLS client certificates signed by the plant's internal authority so every offline action is attributable to an individual technician [S03] [S23].
2. **Locally hosted documentation and runbooks:** Serve a read-only local HTML or PDF copy of the network diagram, IP address register, LED fault codes, safe-shutdown procedure and recovery playbook directly from the local appliance alongside a physical cabinet binder [S05] [S06].
3. **Persistent, bounded local audit logging:** Record every offline login, configuration change or queue action locally and queue audit records for upstream sync [S03] [S22]. On Linux gateways using `systemd-journald`, `journald.conf` must set `Storage=persistent` (writing to `/var/log/journal` rather than volatile `/run/log/journal`), enforce `SystemMaxUse=` and `SystemKeepFree=` (defaulting to 10 percent and 15 percent of the file system, capped at 4 gigabytes) so logs cannot exhaust disk space, and enable `Seal=yes` for tamper-evident Forward Secure Sealing [S22]. Because `journald` flushes `CRIT`, `ALERT` and `EMERG` messages immediately while syncing lower priorities over `SyncIntervalSec=` (default 5 minutes), avoid `ForwardToConsole=yes` on slow serial consoles where a hung console can block `journald` [S22].
4. **No perimeter rule weakening during outages:** Never expose management listeners (`SSH`, `HTTPS`, database or broker admin ports) to untrusted Wi-Fi or cellular interfaces when cloud tunnels fail; restrict maintenance to designated local management ports [S03].

### 5.17 HANDLE CERTIFICATE VALIDITY AND TIME SECURITY UNDER PROLONGED DISCONNECTION

Under TLS 1.3 (RFC 9846), certificate validation checks that wall-clock time falls between `notBefore` and `notAfter` [S23]. Similarly, Network Time Security (NTS, RFC 8915) uses TLS-based NTS Key Establishment (NTS-KE) to verify certificates and issue AEAD cookies for authenticated NTPv4 time exchange [S10] [S27]. On a weak grid, two time-related failure modes must be prevented [S11] [S23] [S27]:

- **Clock reset on cold boot without battery-backed real-time clock (RTC):** Gateways lacking a monitored RTC battery may cold-boot after a blackout with system time reset to the Unix epoch (`1970-01-01T00:00:00Z` in RFC 3339 format), causing every valid TLS 1.3 certificate to fail as "not yet valid" (`current_time < notBefore`) and breaking local HTTPS consoles, MQTT-over-TLS and NTS-KE simultaneously [S23] [S27] [S32]. Every local time server and edge gateway must include a monitored hardware RTC (or local GNSS receiver) and persist the latest timestamp to disk as a monotonic floor across reboots [S11] [S27].
- **Short-lived certificate expiration during multi-day cuts:** If a site uses 7-day leaf certificates (illustrative policy) and a fibre cut isolates the plant for 10 days, local TLS handshakes fail mid-outage [S23]. Certificate lifetimes and renewal triggers on

intermittent sites must exceed the maximum credible disconnection window by a wide margin (such as 90-day to 365-day validity renewed at one-third elapsed lifetime, illustrative policy) [S03] [S23].

> **WARNING**
>
> **Never disable TLS verification (`verify=false`) as an offline workaround**
>
> When a clock jump or unreachable OCSP responder breaks a TLS connection during an outage, field staff are often tempted to set `tls_insecure true` or `verify=false` "temporarily" [S20] [S23]. Temporary security bypasses in factories routinely become permanent and expose credentials to interception across wireless and carrier networks [S03] [S23]. Instead, remove brittle external dependencies by deploying a local NTPv4 stratum server (RFC 5905, RFC 8633), disabling mandatory online OCSP fetching for internal endpoints, pinning the local CA bundle, and flagging clock uncertainty in application payloads rather than disabling transport cryptography [S10] [S11] [S23].

5.18

## Network qualification should reproduce the actual failure modes

A factory network designed for weak-grid resilience cannot be qualified by checking that `ping` succeeds on a sunny afternoon when mains power and fibre backhaul are healthy [S06] [S28]. Just as SQLite's testing guide notes that power-loss durability requires simulated crash testing across the exact failure boundary [S28], qualifying a plant network requires deliberate fault injection that reproduces weak-grid failure modes and records evidence in both local store logs and remote ingestion audit trails [S06] [S22] [S28].

### 5.19 SEVEN MANDATORY FAULT-INJECTION TESTS FOR SITE ACCEPTANCE

The site acceptance test (SAT) checklist must pass seven fault scenarios without losing a single locally committed record or requiring an unscripted process restart [S06] [S13] [S28]:

1. **Hard WAN demarcation pull:** Unplug all WAN cables and cellular antennas during active floor capture [S06]. *Pass criteria:* Local capture latency, broker delivery and console loads show zero errors while the local outbox increments cleanly [S03] [S19].
2. **Cold-start LAN boot while WAN is dark:** With all WAN links unplugged, cold-boot floor switches, capture terminals, local DNS/DHCP, local NTP server and edge gateway [S06] [S28]. *Pass criteria:* Every device obtains its IP, resolves internal names, syncs to local time, completes local TLS handshakes and resumes capture within the recovery time objective (under 5 minutes, illustrative target) with zero external packets [S03] [S11] [S23].
3. **Blackhole routing and zombie TCP session test:** Keep the local WAN interface electrically `UP` while dropping 100 percent of packets two hops upstream [S24]. *Pass criteria:* The edge synchroniser detects application-layer timeout, marks the path degraded and continues local capture without hanging threads [S12] [S24].

4. **Upstream DNS and NTP silence:** Block upstream DNS (port 53) and NTP/NTS-KE (ports 123 and 4460) while leaving IP routing open [S10] [S27]. *Pass criteria:* Internal names resolve authoritatively, the local clock enters holdover without stepping, and the console raises a `DNS_FAILURE` diagnostic alert [S03] [S11].
5. **Degraded, high-loss link simulation:** Inject 15 percent packet loss, 600 milliseconds delay, 200 milliseconds jitter and a 64 kilobits per second cap on the WAN test port (illustrative degraded parameters) [S18] [S24]. *Pass criteria:* Bounded in-flight windows prevent bufferbloat, Lane P0 alarms arrive within budget, Lane P3 bulk transfers pause, and zero duplicate records commit remotely [S12] [S19] [S24].
6. **Reconnection storm and backlog drain test:** Accumulate a synthetic backlog of 10,000 production records and 50 diagnostic attachments (illustrative test load) offline, then restore the WAN under a 256 kilobits per second ceiling while generating live P0 alarms [S06] [S28]. *Pass criteria:* Live P0 alarms arrive ahead of historical P1/P2 batches and the backlog drains monotonically following **Exhibit 21**, p. 81 [S06] [S19].
7. **Dual-carrier failover and failback under load:** Trip the primary WAN mid-batch to force failover to the secondary carrier, then restore the primary path 10 minutes later (illustrative timing) [S05] [S06]. *Pass criteria:* Interrupted batches retry idempotently without duplicating records, and failback respects a hold-down timer [S05] [S12].

### 5.20 REQUIRE AN EXPLICIT ROLLBACK PLAN FOR EVERY NETWORK CHANGE

Pushing a mistyped firewall rule, VLAN tag or static route over a remote link can sever both the backhaul and the shop-floor LAN [S03] [S06]. Every change to routing, DNS, DHCP, firewall rules or broker bridges must therefore use an automated commit-confirm rollback timer that reverts to the prior known-good configuration if local and upstream health checks fail to confirm within 5 minutes (illustrative safeguard) [S03] [S06]. While full cross-system enterprise commissioning is addressed in `wp18`, no plant LAN should enter service until it passes the local autonomy and degraded-backhaul qualification tests above [S06].

In summary, a factory network on a weak grid survives external disconnection when every dependency required for local capture, addressing, naming, time, security and operator visibility resides inside the plant's protected power boundary, when transport protocols and retry policies match actual channel loss and preserve headroom for live alarms, and when dual-carrier paths and recovery sequences are qualified by physical fault injection rather than brochure math [S03] [S05] [S06] [S18]. With the local

network kept alive and disciplined, the next requirement is ensuring that data arriving over that network is committed to non-volatile storage before any sender is told the record is safe, which is the subject of Chapter 6 [S13].

*A plant network that needs the internet in order to talk to the machine across the aisle is not a local network; it is a remote outage waiting for a cable cut.*

## CHAPTER SOURCE NOTES

S03 NIST. 2023.

S18 IETF / RFC Editor. 2014.

S05 NIST. 2010.

S06 NIST. 2016.

S08 IETF / RFC Editor. 2007.

S12 IETF / RFC Editor. 2022.

S24 IETF / RFC Editor. 2022.

S09 IETF / RFC Editor. 2014.

S19 OASIS. 2019-03-07.

S20 Eclipse Mosquitto. 2026-10-30.

S23 IETF / RFC Editor. 2026.

S27 IETF / RFC Editor. 2020.

S28 SQLite. 2026-10-30.

S11 IETF / RFC Editor. 2019.

S13 SQLite. 2026-10-30.

S01 World Bank. 2025.

S32 IETF / RFC Editor. 2002.

S31 IETF / RFC Editor. 2024.

S22 systemd project. 2026-10-30.

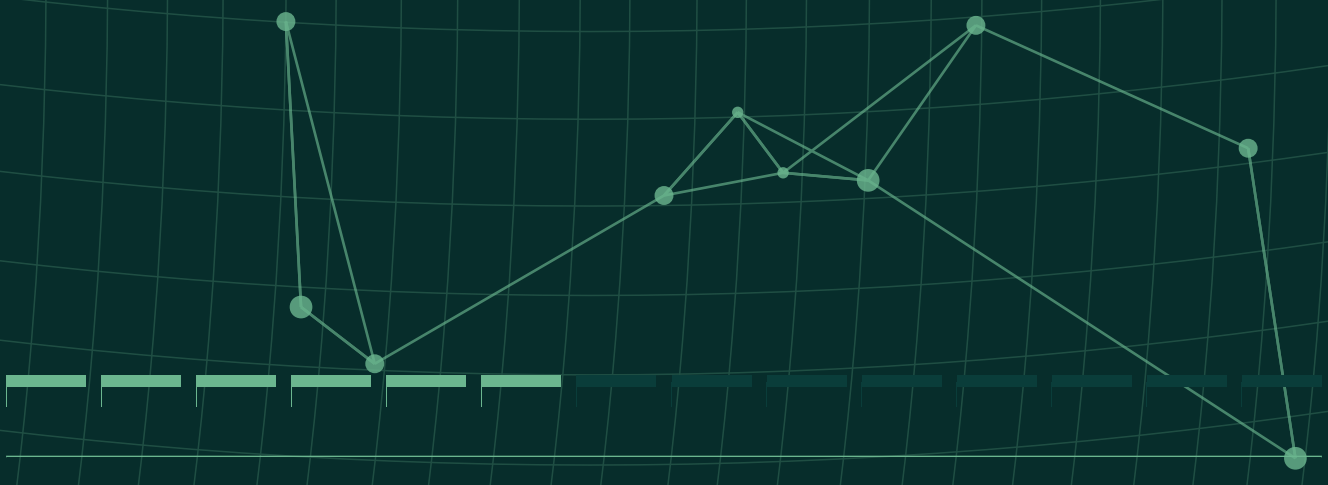
S10 IETF / RFC Editor. 2010.

06

PART II: KEEP THE LOCAL RECORD ALIVE / MAKE CAPTURE DURABLE

# The local event must survive the acknowledgement

*The transaction boundary, storage path and acknowledgement semantics decide whether a production record survives abrupt loss.*



CHAPTER 06
DECISION TEST

The transaction boundary, storage path and acknowledgement semantics decide whether a production record survives abrupt loss.

A factory event is captured only when the local node has validated its envelope, written both the business event and its pending-delivery outbox record within a single atomic transaction, and confirmed that the storage stack has flushed that commit to non-volatile media before returning an acknowledgement to the operator or machine controller. [S13] [S16] In an industrial environment where 93.3 percent of surveyed Cameroonian firms reported electrical outages and experienced 10.4 outages in a typical month according to the 2024 World Bank Enterprise Survey, treating an in-memory buffer write or a transport socket handoff as an accepted record guarantees silent data loss whenever power fails between the acknowledgement and the physical disk flush. [S01] [S24] Green Future Tech Limited (GFT), Hong Kong SAR, argues in its Cameroon textile programme (prepared 31 January 2024, sources refreshed 2026) that closing the gap between roughly 95 percent raw cotton export and the 50 percent local-processing target by 2030 requires point-of-work records that survive the exact moment of electrical interruption. [R06] [R07] This chapter specifies the local transaction contract, the difference between database consistency and power-loss durability, the single-host storage rules for write-ahead logging, the boundary between message-broker queues and application commits, the partition reserve budget for bounded retention, and the fault-injection ledger required to prove recovery before production relies on the record. [S14] [S20] [S28]

6.1

## An event envelope should carry enough context to recover

A bare production value such as a sewing bundle count, a loom pick total or a dye-bath temperature reading cannot be reconciled after an unscheduled reboot unless the event carries its own recovery metadata from the instant of creation. [S31] [S32] When an edge gateway or line terminal loses mains supply and restarts without wide-area network reachability, its local operating system may come up with an unsynchronised real-time clock, a reset process counter and an unconfirmed transmission queue. [S03] [S22] If the capture application relies on a central cloud database to assign auto-increment primary keys or ingestion timestamps, local operation either blocks completely during a backhaul outage or generates colliding local integers that overwrite history when connectivity returns. [S31] Data governance ownership, schema lifecycle approval and long-term archival retention rules belong to the wider governance framework in paper wp14; this chapter defines the nine mandatory envelope fields required at the local capture boundary so that any interrupted node can detect gaps, reject corrupted payloads and deduplicate replayed records without external assistance. [S13] [S31]

### 6.2 NINE RECOVERY FIELDS SEPARATE REPLAY FROM NEW PRODUCTION

Every production event written to the local store should be wrapped in a fixed contextual envelope before the database transaction begins. [S13] [S31] Rather than trusting transport headers that disappear across protocol bridges, the capture service embeds nine fields directly into the committed record: [S19] [S31]

1. **Event identifier (`event_id`):** A locally generated 128-bit Universally Unique Identifier (UUID) formatted according to RFC 9562, requiring no central registration authority or online coordinator. [S31] Using time-ordered UUID version 7 (`UUIDv7`), which places a 48-bit Unix Epoch millisecond timestamp in the most significant bits followed by version, variant, optional monotonic counter and pseudorandom bits, preserves B-tree index locality during high-rate inserts while preventing identifier collisions across disconnected edge nodes. [S31]
2. **Source identifier (`source_id`):** A stable, provisioned identifier naming the specific physical capture boundary, such as a garment bundle scanner station, a spinning-frame programmable logic controller bridge or a supervisor tablet assigned to a named production line. [S03] [S31]
3. **Boot identifier (`boot_id`):** A 128-bit identifier generated fresh each time the operating system kernel or capture daemon initialises, allowing downstream reconciliation to see immediately whether two events with low sequence numbers came from the same continuous run or sit on opposite sides of a power reset. [S22] [S31]
4. **Per-source sequence number (`source_seq`):** A strictly increasing 64-bit integer maintained per `source_id` and `boot_id` pair (or persisted across restarts in the local transaction store), which exposes missing records as arithmetic gaps even when the local clock jumps backward or forward during recovery. [S13] [S32]
5. **Source event timestamp (`event_time_utc`):** The instant at which the physical observation or operator action occurred, formatted under the RFC 3339 profile of ISO 8601 with an explicit Coordinated Universal Time (UTC) offset (`Z` or numeric offset), never an unqualified local wall-clock string. [S32]
6. **Local ingest timestamp (`ingest_time_utc`):** The RFC 3339 UTC timestamp recorded when the local database transaction accepts the event, preserving the delay between peripheral sensing and local persistence whenever a scanner or controller buffers a reading briefly before handoff. [S13] [S32]
7. **Clock quality indicator (`clock_quality`):** An enumerated state flag recording whether the local system clock at `event_time_utc` was locked to an authenticated or verified local reference, running in bounded holdover after reference loss, or unsynchronised following a cold boot without battery-backed time. [S03] [S32]
8. **Schema version (`schema_version`):** An explicit integer or semantic version tag identifying the structure and engineering units of the payload so that queued offline events remain parseable even if an application update occurs before the backlog drains. [S05] [S13]
9. **Payload integrity digest (`payload_sha256`):** A cryptographic checksum or cyclic redundancy check computed over the canonical serialised payload bytes and key envelope headers at capture time, enabling the local store and remote consumer to distinguish application-level payload corruption from valid zero-value readings. [S13] [S15]

Embedding these nine attributes inside the local transaction record means that an operator scanning a cut-work bundle on a sewing line in Douala or Garoua produces a self-describing evidence unit. [S03] [S31] If a voltage dip restarts the station three seconds later, the next event carries a new `boot_id` and a verifiable `source_seq` continuity trail, allowing the performance-management system (PMS) proposed in GFT's Cameroon textile programme to separate genuine production stops from instrumentation restarts. [S03] [S22]

6.3

## Write the event and outbox in one transaction

Capturing an event locally and queuing it for upstream transmission are two views of the same operational commitment, so they must succeed or fail as a single atomic database transaction. [S13] [S0602] In SQLite, atomic commit guarantees that either all database modifications within a single transaction occur together or none of them occur, making writes across multiple tables appear instantaneous and simultaneous even when interrupted by an operating-system crash or power failure. [S13] When the local capture service writes the validated event into the `local_events` ledger table and simultaneously inserts a delivery pointer into the `outbox_queue` table inside one `BEGIN IMMEDIATE ... COMMIT` block, the node eliminates any state in which an event is recorded locally but forgotten by the uplink forwarder, or forwarded upstream while absent from the local shift ledger. [S13] [S0602]

### 6.4 DUAL WRITES OPEN A SPLIT-BRAIN WINDOW BETWEEN STORAGE AND TRANSPORT

Many industrial integration scripts violate atomicity by executing two separate actions: they insert a record into a local file or database and then publish a message over a Transmission Control Protocol (TCP) socket or message-broker client before returning success to the station. [S19] [S24] That dual-write pattern creates two distinct failure windows under intermittent power. [S13] [S15] If the process writes to the local database first and loses power one millisecond before publishing to the broker, the record exists on disk, but nothing in the communication path knows that it still needs to be transmitted once the wide-area network recovers. [S08] [S13] Conversely, if the process publishes the event to a network socket first and crashes before the local database commit completes, the remote dashboard or downstream enterprise service may record a bundle completion that is missing from the local supervisor console, violating the single-record principle required for shop-floor accountability. [S03] [S13]

Nor does a successful socket write prove that a remote or local peer has stored the event. [S24] As RFC 9293 specifies, TCP provides a reliable, in-order byte-stream service between network endpoints, not an application-layer transaction guarantee. [S24] A `send()` or `write()` call on a TCP socket returns as soon as bytes are copied into the local kernel socket buffer, long before the receiving application has parsed the payload, validated the schema or flushed a database transaction to non-volatile storage. [S24] [S0603] By replacing dual writes with a single local database transaction that writes both `local_events` and `outbox_queue`, the edge node decouples point-of-work acknowledgement from network state. [S08] [S13] An independent local forwarder process then reads committed rows from `outbox_queue` where `delivery_state =`

'PENDING', transmits them when connectivity permits, and updates delivery_state = 'ACKNOWLEDGED' only after receiving an explicit application-level commit receipt containing the matching event_id. [S12] [S13]

### 6.5 UNACKNOWLEDGED COMMITS REQUIRE DETERMINISTIC CLIENT RETRY HANDLING

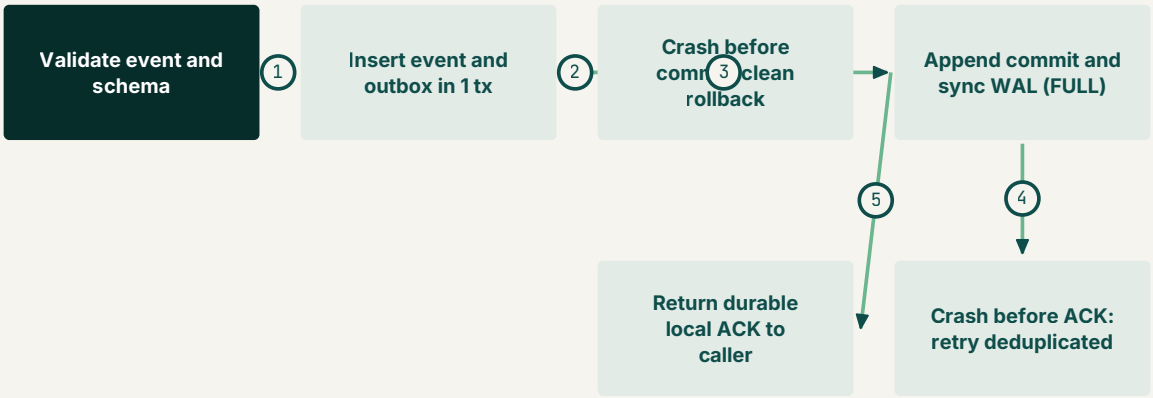
Even with an atomic local transaction, a narrow failure window remains between the instant the storage device completes the commit and the instant the calling peripheral or operator terminal receives the local acknowledgement packet. [S12] [S13] Suppose a barcode terminal at an inspection table submits event_id = 0192dce0-0002 to the local edge service, the SQLite transaction commits durably to flash storage, and the local access switch or terminal loses power before the HTTP or local inter-process response arrives. [S12] [S13] When the terminal recovers, it cannot tell from the missing response alone whether the transaction aborted before commit or failed after commit. [S12] [S13]

If the terminal generates a brand-new random identifier on retry, the local database will insert a second row for the same physical garment bundle, inflating shift output. [S19] [S31] To close this ambiguity, the peripheral or station client must retain the deterministic event_id (or the deterministic tuple of source_id, boot_id and source_seq) across its retry loop. [S12] [S31] When the edge capture service receives the retried submission, its transaction checks the unique index on event_id (and on (source_id, boot_id, source_seq)). [S13] [S31] If a row with identical payload_sha256 is already committed in local_events, the service does not insert a duplicate row or a duplicate outbox_queue entry; it simply returns a durable local acceptance receipt citing the original ingest_time_utc, making the local capture endpoint idempotent in the sense defined by RFC 9110. [S12] [S13] If a retried event_id arrives with a different payload_sha256, the service logs an explicit integrity conflict rather than overwriting the prior record. [S13] [S15]

The transaction sequence in **Exhibit 22**, p. 91 illustrates where crash points fall across validation, atomic insertion, write-ahead log synchronisation and caller acknowledgement. [S13] [S16]

EXHIBIT 22 / TRANSACTION SEQUENCE

## The acknowledgement follows the durable transaction




1. valid
2. power cut
3. commit
4. link drop
5. fsync ok

ILLUSTRATIVE [S13] [S16] [S0603] Illustrative transaction and outbox sequence with crash boundaries before insert, before commit, after commit and before client acknowledgement. Guarantees follow SQLite atomic commit (S13), WAL synchronous modes (S16) and fsync behaviour (S0603).

As **Exhibit 22**, p. 91 shows, any interruption prior to the completion of the write-ahead log synchronisation step (n4) rolls back both the event and the outbox row cleanly upon reboot, leaving zero orphaned state. [S13] [S16] Any interruption after n4 but before n6 preserves the committed event and outbox row on disk, while the idempotent retry rule on `event_id` prevents duplicate insertion when the caller resends the unacknowledged frame. [S12] [S13] What this flow diagram cannot establish on its own is whether the underlying hardware honours the synchronisation call at step n4; that requires explicit database pragma configuration and storage verification. [S15] [S16]

6.6

# Consistency does not guarantee the latest commit is durable

A database engine that recovers to a structurally consistent state after a power cut can still silently discard the last several seconds of committed production transactions unless its synchronisation mode forces an input/output barrier on every commit. [S15] [S16] Software teams frequently select SQLite Write-Ahead Logging (`PRAGMA journal_mode=WAL`) because it allows concurrent readers and writers and replaces random page rewrites with sequential log appends. [S14] [S16] However, many deploy-

ments pair WAL mode with PRAGMA synchronous=NORMAL to avoid blocking the writer thread on disk flush operations, without realising that NORMAL in WAL mode sacrifices power-loss durability. [S14] [S16]

### 6.7 WRITE-AHEAD LOG SYNCHRONISATION MODES TRADE SPEED AGAINST POWER-LOSS SURVIVAL

In SQLite's traditional rollback-journal mode (DELETE, TRUNCATE or PERSIST), committing a transaction copies original unchanged pages into the rollback journal, flushes the journal to mass storage, writes modified pages into the main database file, flushes the database file, and then deletes, truncates or zeroes the header of the rollback journal to mark the commit instant. [S13] [S0604] Write-Ahead Logging inverts that sequence: original pages remain untouched in the main database file while new page versions and a commit record are appended sequentially to a separate -wal file, and a later checkpoint operation transfers committed frames from the -wal file back into the main database file. [S14] [S0604]

Because the -wal file is sequential, SQLite offers four distinct settings for PRAGMA synchronous: OFF (0), NORMAL (1), FULL (2) and EXTRA (3). [S16] Their exact failure boundaries under application crash versus operating-system crash or power loss are documented in the SQLite pragma specification: [S16]

- synchronous=OFF (0): SQLite hands data to the operating system and continues immediately without issuing any fsync() or FlushFileBuffers() system call. [S15] [S16] While transactions remain safe across an application-process crash if the operating system stays up, an operating-system crash or power loss before buffers reach non-volatile media will likely corrupt the database file in both rollback and WAL modes. [S15] [S16]
- synchronous=NORMAL (1): In WAL mode, SQLite synchronises the -wal file before each checkpoint, synchronises the main database file after each completed checkpoint, and synchronises the -wal header when reusing a reset log file, but it performs **no sync operations during ordinary transaction commits**. [S14] [S16] Consequently, transactions committed in WAL mode with synchronous=NORMAL remain atomic, consistent and isolated, and survive application crashes, but **a transaction committed with synchronous=NORMAL might roll back following a power loss or system crash**. [S16] Between checkpoints (which default to running when the -wal file reaches 1,000 pages, or roughly 4 megabytes at the standard 4,096-byte page size), hundreds of acknowledged production events sitting only in volatile page cache can vanish cleanly on a power cut without raising a corruption error. [S14] [S16]
- synchronous=FULL (2): In WAL mode, SQLite invokes the Virtual File System xSync method (fsync() on Unix) on the -wal file after every single transaction commit in addition to the checkpoint syncs, making WAL transactions fully atomic, consistent, isolated and durable (ACID) across power loss provided the storage hardware honours the flush. [S14] [S16] In rollback-journal mode, FULL prevents corruption across power loss, though on some filesystems an unlinked rollback journal directory entry can still reappear after an immediate power cut unless the containing directory is also synced. [S16] [S0603]

- `synchronous=EXTRA (3)`: In `WAL` mode, `EXTRA` is identical to `FULL`. [S16] In rollback-journal `DELETE` mode, `EXTRA` adds an explicit sync of the directory containing the rollback journal after unlinking it, closing the Linux `fsync(2)` caveat where syncing a file descriptor does not automatically guarantee that its directory entry update has reached the disk surface. [S16] [S0603]

For a factory service class that promises durable local acceptance at the moment an operator scans a bundle or logs a quality disposition, `PRAGMA journal_mode=WAL` with `PRAGMA synchronous=NORMAL` violates the contract. [S05] [S16] A station screen that flashes green while its database runs `WAL` with `synchronous=NORMAL` is issuing an optimistic memory receipt, not a power-safe commitment. [S13] [S16] Where the service class requires zero committed-record loss across abrupt mains failure (and where the UPS or protected circuit can still drop unexpectedly due to transfer failure or battery exhaustion), the local event store must run `PRAGMA journal_mode=WAL` with `PRAGMA synchronous=FULL` (or rollback mode with `synchronous=EXTRA`). [S05] [S16]

### 6.8 STARTUP READBACK VERIFIES EFFECTIVE DATABASE PRAGMAS

Configuring durability requires more than placing a SQL string in an initialisation script, because SQLite deliberately generates no error message when an unknown or misspelled `PRAGMA` statement is issued; unknown pragmas are silently ignored. [S16] Furthermore, while `PRAGMA journal_mode=WAL` is persistent across database connections once succeeded, `PRAGMA synchronous`, `PRAGMA foreign_keys`, `PRAGMA busy_timeout` and `PRAGMA wal_autocheckpoint` are per-connection settings that revert to compile-time defaults whenever a database connection closes and reopens. [S14] [S16] If a VFS lacks shared-memory primitives and the connection has not been placed into `PRAGMA locking_mode=EXCLUSIVE` prior to entering `WAL`, an attempt to set `PRAGMA journal_mode=WAL` silently leaves the database in its prior rollback mode and returns that prior mode string. [S14] [S16]

METHOD

### 6.9 STARTUP DURABILITY CONFIGURATION AND READBACK CHECK

Every local capture service connection must execute an explicit setup and readback sequence on every startup and reconnection before accepting any production event: [S14] [S16]

1. **Apply connection pragmas:** Execute `PRAGMA locking_mode=EXCLUSIVE;` (when designed as a single-process writer), `PRAGMA journal_mode=WAL;`, `PRAGMA synchronous=FULL;`, `PRAGMA foreign_keys=ON;`, `PRAGMA busy_timeout=5000;`, `PRAGMA wal_autocheckpoint=1000;` and `PRAGMA journal_size_limit=16777216;` (an illustrative 16-megabyte post-checkpoint truncation cap). [S14] [S16]
2. **Read back and assert effective values:** Immediately query `PRAGMA journal_mode;`, `PRAGMA synchronous;`, `PRAGMA foreign_keys;` and `PRAGMA compile_options;`. [S16] Verify programmatically that `journal_mode` returns `'wal'`, `synchronous` returns integer `2` (`FULL`) or `3` (`EXTRA`), and `foreign_keys` returns `1`. [S16]
3. **Run startup structural check:** After an unclean shutdown (detected via a persistent state marker or a changed kernel `boot_id`), run `PRAGMA quick_check;` (which executes in `O(N)` time) or `PRAGMA integrity_check;` (which executes in `O(N log N)` time and verifies unique and index consistency) and confirm that the single returned row equals `'ok'` before opening the intake socket. [S16] [S22]
4. **Fail closed on mismatch:** If any readback value differs from the required durability specification or if the integrity check reports an anomaly, hold the capture service in an explicit degraded maintenance state and alert the supervisor console rather than accepting events into an unverified store. [S03] [S06]

6.10

## Keep a local database on a supported storage path

Database transaction guarantees hold only when the database files reside on a local storage path whose locking, shared-memory, write-ordering and sector-overwrite semantics match the assumptions of the database engine. [S13] [S15] Placing an embedded transactional store on a convenience network share or on unverified consumer flash media can defeat even `PRAGMA synchronous=FULL`. [S14] [S15]

### 6.11 SHARED-MEMORY INDEXING RULES OUT NETWORK FILESYSTEMS FOR LOCAL WAL CAPTURE

Plant integrators sometimes attempt to centralise resilience by mounting a Network File System (NFS) or Server Message Block (SMB) share from a back-office storage appliance onto multiple line-side terminals and pointing each terminal's SQLite database connection at the network share. [S14] [S15] SQLite's documentation rules out this architecture explicitly: all processes using a `WAL` database must run on the same host computer, and `WAL` does not work over a network filesystem. [S14] To allow readers to locate page versions in the `-wal` file without scanning the entire log on every query, SQLite maintains a `wal-index` in a memory-mapped `-shm` file in the same directory as the database. [S14] [S0604] Processes running on separate physical

machines cannot share that memory mapping, and network filesystems frequently contain locking bugs or cache-coherency delays that break POSIX advisory locks and corrupt the database. [S14] [S15]

Even on a single host, concurrency boundaries require disciplined process design. [S14] [S0602] In SQLite, writes are serialised so that only a single writer can modify a database at a time. [S14] [S0602] While multiple local processes can open a WAL database concurrently, multi-connection write and checkpoint coordination introduces avoidable complexity and historical corner cases. [S14] [S15] For example, on 3 March 2026 the SQLite team documented and fixed the WAL-reset race condition (present in versions 3.7.0 from July 2010 through 3.51.2 from January 2026, and fixed in 3.51.3 on 13 March 2026 with backports to 3.44.6 and 3.50.7), in which two separate connections concurrently writing and checkpointing at a tight timing boundary could leave a wal-index header field stale and cause a subsequent checkpoint to skip part of a committed transaction. [S14] [S15] By confining all writes and checkpoints for the local event store to a single dedicated capture daemon process on the local host (or running a single-process connection with PRAGMA locking_mode=EXCLUSIVE, which also allows WAL to operate without shared memory if only one process accesses the file) and upgrading SQLite to at least version 3.51.3, the edge architecture eliminates cross-process write races entirely. [S14] [S0602]

Equally important, the main database file, its -wal write-ahead log, its -shm shared-memory index and any hot rollback -journal file form an indivisible state set. [S14] [S15] If an administrator or faulty cleanup script deletes, renames or separates a hot -journal or -wal file from its database file after a power failure, SQLite cannot play back or recover the interrupted state on the next sqlite3_open(), resulting in silent transaction loss or immediate database corruption. [S14] [S15]

### 6.12 FILESYSTEM BARRIERS AND FLASH CONTROLLERS SIT BENEATH DATABASE GUARANTEES

Below the SQLite Virtual File System layer sit the operating-system filesystem and the physical mass-storage controller. [S13] [S29] SQLite assumes by default (since version 3.7.9 in November 2011) that the underlying storage provides the **powersafe overwrite** property: when an application writes a range of bytes in a file, no bytes outside that range will change even if a crash or power failure occurs just before or during the write. [S13] [S0601] On flash media where the minimum erasable block or sector (typically 4,096 bytes or larger) exceeds the byte range being modified, a controller lacking powersafe overwrite can suffer collateral sector damage during an interrupted write-back, corrupting adjacent database pages that were not even being modified. [S15] [S0601] Even worse, SQLite's corruption guide notes that some low-cost flash memory controllers use wear-levelling logic that can inflict random filesystem damage across unrelated files if power is cut during any write, and that consumer USB flash sticks frequently report fsync() completion while data still sits in volatile controller RAM (visibly indicated by an LED that continues flashing for several seconds after COMMIT returns). [S15]

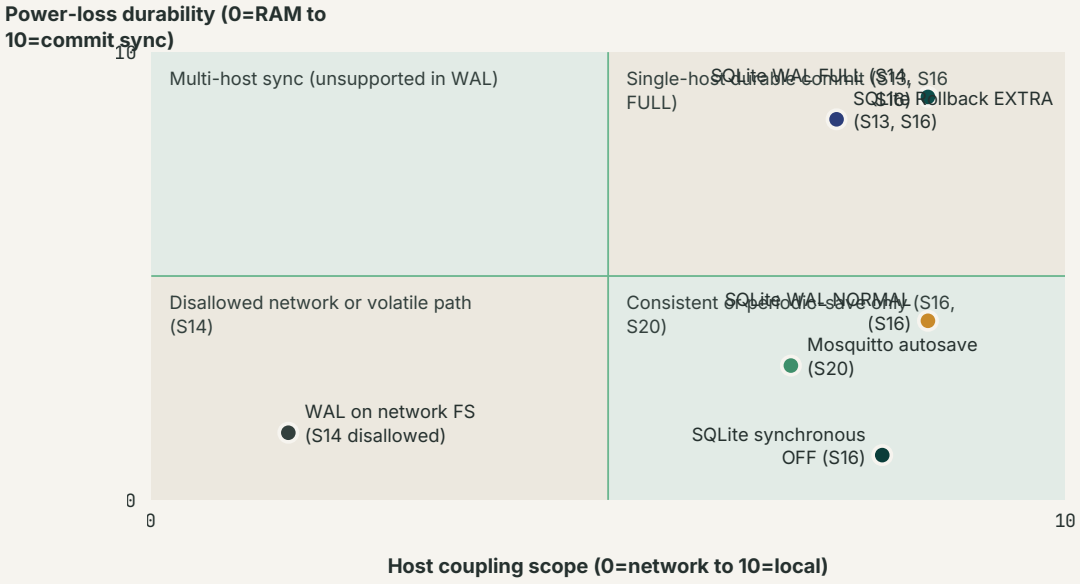
On Linux edge computers, the filesystem mount configuration must preserve write ordering down to the block device. [S29] The Linux kernel ext4 documentation explains that ext4 enables write barriers (barrier=1) and data=ordered mode by default. [S29] In data=ordered mode, ext4 forces all modified file data blocks out to the main

filesystem prior to committing the associated metadata transaction to the journal, and write barriers enforce proper on-disk ordering of journal commits so that volatile disk write caches remain safe to use. [S29] By default, `ext4` commits running metadata transactions every `commit=5` seconds, whereas an explicit `fsync()` call from SQLite (`synchronous=FULL`) forces an immediate flush and barrier wait on the file descriptor. [S29] [S0603] If an integrator mounts the partition with `nobarrier` (`barrier=0`) without a battery-backed storage controller, or uses `data=writeback` where metadata can commit before data blocks, an abrupt power loss can reorder writes across journal commits and defeat database recovery. [S15] [S29]

The source-based matrix in **Exhibit 23**, p. 96 maps these documented persistence configurations across two operational dimensions: host coupling scope and power-loss durability. [S13] [S14] [S16] [S20]

EXHIBIT 23 / ORDINAL BOUNDARY INDEX (0 TO 10)

## Persistence claims have different failure boundaries




SOURCED [S13] [S14] [S16] [S20] Source-based mapping of documented persistence boundaries from SQLite atomic commit (S13), WAL same-host shared-memory constraints (S14), PRAGMA synchronous modes (S16) and Mosquitto mosquitto.db persistence (S20). Coordinates represent ordinal scope indices (0 to 10) for visual placement; exact source rules and GFT acceptance checks are tabulated in the text.

As **Exhibit 23**, p. 96 shows, only single-host local configurations with explicit per-commit synchronisation (`SQLite WAL FULL` and `SQLite Rollback EXTRA`) occupy the upper-right quadrant where committed records survive abrupt power loss. [S13] [S14] [S16] Configurations in the lower-right quadrant (`SQLite WAL NORMAL` and `Mosquitto autosave`) remain on a supported local host but can lose recent committed events or queued messages when power drops between checkpoints or autosave intervals. [S16] [S20] Placing `WAL` on a network filesystem sits in the lower-left disallowed zone because shared-memory indexing and network locking fail across hosts. [S14] [S15] Because a two-dimensional plot uses ordinal placement indices (`0` to `10`) to visualise relative boundaries, the exact documented source statements and the corresponding GFT procurement and commissioning questions are set out in the comparison table below. [S13] [S14] [S16] [S20]

| Mechanism and source | Documented source behaviour | Failure boundary under abrupt power loss | GFT acceptance question (proposed control) |
|---|---|---|---|
| SQLite rollback journal with `synchronous=EXTRA` [^S13] [^S16] [^S0603] | Original pages are flushed to a rollback journal before main database writes; `EXTRA` also syncs the containing directory after unlinking the journal [^S13] [^S16] | Atomic, consistent, isolated and durable (`ACID`) across power loss if the storage device honours flush and powersafe overwrite [^S16] [^S0601] | Does startup read-back confirm `synchronous=3` when rollback mode is selected, and is directory sync verified on the target filesystem? [^S16] [^S0603] |
| SQLite `WAL` mode on local disk [^S14] | Appends changes to a `-wal` file and uses a memory-mapped `-shm` `wal-index`; requires all processes to run on the same host computer [^S14] | Never supported over NFS or SMB network filesystems; requires SQLite `3.51.3` or later (or single-connection access) to avoid the multi-connection `WAL-reset` race [^S14] [^S15] | Are the database, `-wal` and `-shm` files on a local block device with a single writer process running SQLite `3.51.3` or newer? [^S14] |
| SQLite `WAL` with `synchronous=NORMAL` [^S16] | Syncs the `-wal` file before checkpoints and the main database after checkpoints, but omits `fsync` on ordinary transaction commits [^S14] [^S16] | Preserves structural consistency across power loss, but transactions committed since the last checkpoint can roll back after power loss or hard reboot [^S16] | Has the operator explicitly restricted `NORMAL` to non-critical diagnostic stores where post-crash rollback of recent events is acceptable? [^S05] [^S16] |
| SQLite `WAL` with `synchronous=FULL` [^S16] | Issues an extra `xSync` (`fsync`) of the `-wal` file on every transaction commit in addition to checkpoint syncs [^S14] [^S16] | Full `ACID` durability across power loss, bounded below by filesystem write barriers and honest hardware flush completion [^S15] [^S16] [^S29] | Does the capture service verify `journal_mode=wal` and `synchronous=2` on every connection open before issuing local acknowledgements? [^S16] |
| Linux `ext4` mount and block device [^S15] [^S29] [^S0601] | `ext4` defaults to `data=ordered`, `barrier=1` and `commit=5` seconds; consumer flash sticks and non-powersafe controllers can lie about `fsync` or corrupt adjacent sectors [^S15] [^S29] | If `nobarrier` is set without battery-backed cache or if flash media lies about flush completion, out-of-order writes can corrupt the database on power cut [^S15] [^S29] | Does `/proc/mounts` confirm `barrier=1` and `data=ordered` on industrial solid-state storage qualified via physical power-cut tests? [^S28] [^S29] |
| | `persistence` defaults to `false`; when `true`, state is written to | In-memory queued messages and session updates accu- | Does the architecture rely on a transactional SQLite |

| Mechanism and source | Documented source behaviour | Failure boundary under abrupt power loss | GFT acceptance question (proposed control) |
|---|---|---|---|
| Eclipse Mosquitto broker persistence [^S19] [^S20] | `mosquitto.db` on exit, `SIGUSR1` or every `autosave_interval` (default `1,800` seconds) [^S20] | mulated since the last `mosquitto.db` save are lost if power fails before the next periodic save [^S20] | outbox for durability rather than assuming an MQTT `PUBACK` flushed `mosquitto.db`? [^S13][^S20] |

6.13

# Broker persistence is a separate configuration decision

In an edge architecture that uses Message Queuing Telemetry Transport (MQTT) to decouple local station publishers from local consumers or upstream bridges, an MQTT protocol acknowledgement is not proof that a message has been synchronously committed to non-volatile disk. [S19] [S20] Treating a local message broker as a drop-in substitute for a transactional database confuses transport-layer session state with durable ledger storage. [S13] [S19]

### 6.14 PERIODIC BROKER SAVES LEAVE AN UNPERSISTED MEMORY WINDOW

The OASIS MQTT Version 5.0 standard defines three Quality of Service (`QoS`) levels for message delivery between a client and a server (broker): `QoS 0` (`At most once`, where messages are delivered according to the best efforts of the operating environment and message loss can occur), `QoS 1` (`At least once`, where messages are assured to arrive via a `PUBLISH` and `PUBACK` exchange but duplicates can occur), and `QoS 2` (`Exactly once`, where a four-part `PUBLISH`, `PUBREC`, `PUBREL` and `PUBCOMP` handshake assures that the message arrives once at the protocol layer). [S19] When a client connects with a persistent session (`Clean Start` set to `0` and a non-zero `Session Expiry Interval` in MQTT v5.0, or `clean_session=False` in MQTT v3.1.1), the specification requires the server and client to store session state so that undelivered `QoS 1` and `QoS 2` messages can be retried across network reconnections. [S19] [S20]

However, how and when a broker writes that session state to non-volatile disk is an implementation and configuration choice, not an automatic per-message `fsync()`. [S20] The official `mosquitto.conf` manual page for Eclipse Mosquitto reveals three critical defaults that surprise teams who have not read the documentation: [S20]

1. **Persistence is disabled by default:** The `persistence` boolean option defaults to `false`, meaning that unless explicitly configured (`persistence true` with a valid writable `persistence_location`), Mosquitto stores all connection, subscription and queued message data in volatile memory only. [S20]
2. **Disk saves are periodic, not per-message:** Even when `persistence true` is enabled so that state is written to `mosquitto.db`, Mosquitto writes the in-memory database to disk only when the broker closes cleanly, when it receives a `SIGUSR1` signal, or at periodic intervals governed by `autosave_interval`. [S20] By default,

`autosave_interval` is `1,800` seconds (`30` minutes) when `autosave_on_changes` is `false` (or `1,800` accumulated subscription, retained-message and queued-message changes when `autosave_on_changes` is `true`). [S20]

3. **Per-client broker queues are bounded and drop silently on overflow:** For disconnected persistent clients, Mosquitto holds at most `max_queued_messages` (defaulting to `1,000` `QoS 1` or `QoS 2` messages above those in flight) and `max_queued_bytes` (defaulting to `0`), and `queue_qos0_messages` defaults to `false`. [S20] As the manual states, once the `max_queued_bytes` or `max_queued_messages` limit is reached, subsequent outgoing messages for that client are **silently dropped**. [S20] Similarly, if `memory_limit` is exceeded while an incoming message is processed, the message is dropped and the publishing client is disconnected. [S20]

Consequently, if a sewing-line terminal publishes a `QoS 1` event to a local Mosquitto broker, receives a `PUBACK` packet in two milliseconds, and deletes its own copy of the event, an unscheduled power failure ten minutes later (before the `1,800`-second `autosave_interval` elapses or when `SIGKILL` / abrupt mains loss prevents a clean broker exit) erases the message from the broker's volatile RAM. [S19] [S20] Even if `autosave_interval` is lowered, rewriting the entire `mosquitto.db` snapshot file is an batch checkpoint mechanism, not an atomic per-event transaction log. [S13] [S20]

### 6.15 TRANSPORT QUALITY OF SERVICE DOES NOT EQUAL APPLICATION LEDGER DURABILITY

A second boundary error occurs between the broker and the downstream consumer application. [S19] [S24] Suppose a local broker delivers a `QoS 1` or `QoS 2` message to a local PMS ingestion worker or an upstream cloud bridge. [S19] If the consumer client library sends `PUBACK` (or completes `PUBCOMP`) automatically upon invoking the application callback in memory, before the application has committed the event into its target database, a crash inside the consumer process loses the event permanently because the broker has already deleted its queued copy upon receiving the protocol acknowledgement. [S13] [S19]

To maintain an unbroken durability chain across producer, broker and consumer, the architecture must enforce three separate boundaries: [S08] [S13] [S19]

- **Producer durability boundary:** The station or edge capture service commits the event and its `outbox_queue` row into local SQLite (`WAL`, `synchronous=FULL`) before acknowledging the station action. [S13] [S16] The producer's MQTT publisher reads from `outbox_queue` and retains the row in `'PENDING'` state across broker restarts. [S08] [S13]
- **Broker transport role:** The local MQTT broker acts as a pub/sub router for live local subscribers (such as supervisor displays and maintenance alarm panels) with `persistence true` and bounded queue settings configured to aid fast session resumption, but the broker's `mosquitto.db` is never treated as the sole system of record for unforwarded production events. [S03] [S20]
- **Consumer acceptance boundary:** Downstream consumers disable automatic client-library `PUBACK` generation, execute an idempotent database transaction keyed by `event_id` (`INSERT ... ON CONFLICT (event_id) DO NOTHING`), commit that transaction durably, and emit both the protocol acknowledgement and an

application-level receipt message (allowing the producer's `outbox_queue` to mark the `event_id` as `'ACKNOWLEDGED'`) only after the consumer commit succeeds. [S12] [S13] [S19]

> **GFT NOTE**
>
> **6.16 TREAT THE TRANSACTIONAL OUTBOX, NOT THE BROKER QUEUE, AS THE LOCAL SYSTEM OF RECORD**
>
> In GFT's Cameroon textile programme, the performance-management system (PMS) and digital strategy require point-of-work logs from cutting, sewing and quality inspection to survive both backhaul loss and local power interruptions. [S03] [S08] GFT recommends that integrators never treat an MQTT `PUBACK` or `PUBCOMP` from a local Mosquitto broker as a durable storage receipt. [S19] [S20] Keep the authoritative offline backlog inside an atomic SQLite `local_events` and `outbox_queue` store (`PRAGMA journal_mode=WAL`, `PRAGMA synchronous=FULL`) on a single-host local partition, and release an `outbox_queue` entry only when the downstream ledger returns a signed or authenticated application receipt naming the committed `event_id`. [S13] [S16] [S31] This design lets the broker restart cleanly from an empty or stale `mosquitto.db` without losing a single committed production event. [S13] [S20]

6.17

# Full storage must trigger a visible and bounded policy

Persistent local storage protects the record during an outage only until the storage partition fills up; an unbounded queue or an unthrottled diagnostic log on a shared disk will eventually cause the database engine to fail with `SQLITE_FULL` and halt local capture. [S14] [S22] Sizing the local disk solely from the raw byte length of the event payload ignores five additional consumers of non-volatile space that grow alongside the event table during an extended disconnection. [S14] [S17] [S22]

### 6.18 CHECKPOINTS, JOURNALS AND BACKUPS COMPETE WITH RAW EVENT PAYLOADS

In the paper's calculation register (`C02`, examined further for uplink drain dynamics in Chapter 7), an illustrative site with `50` event sources generating `1` event per second per source at `500` bytes per event over an `8`-hour (`28,800`-second) offline interval produces `1,440,000` events, or `720,000,000` bytes (`720` decimal megabytes, `0.72` decimal gigabytes) of raw event payload. [S08] On an actual local filesystem, committing and maintaining those `1,440,000` events requires substantial additional space across six distinct structures: [S14] [S17] [S22] [S0604]

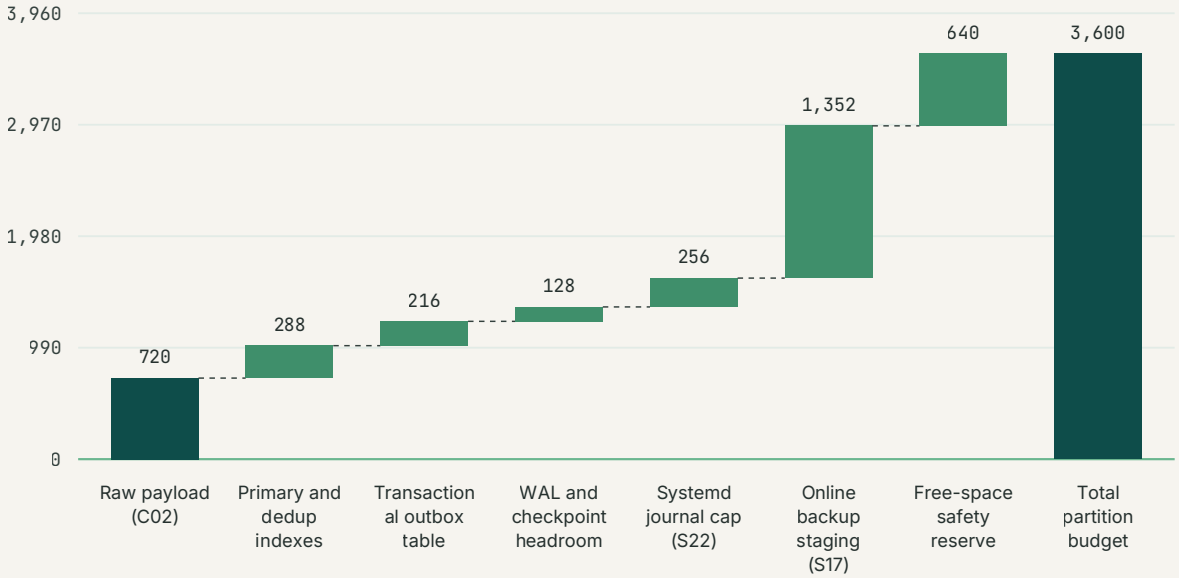
1. **Primary and secondary B-tree indexes:** Enforcing idempotency and sequence gap detection requires a unique index on `event_id` (128-bit UUID) and a composite index on `(source_id, boot_id, source_seq)` plus status indexing on `outbox_queue`. [S13] [S31] Together with B-tree page fill factors and cell headers, index pages add an illustrative 35 to 45 percent above the raw payload size. [S14] [S31]

2. **The transactional outbox table:** Until the wide-area network recovers and downstream receipts clear the backlog, `outbox_queue` holds a row for every pending event (storing `event_id`, `source_id`, retry count, next-attempt timestamp and delivery state), adding an illustrative 150 bytes per event. [S08] [S13]
3. **Write-ahead log (`-wal`) and temporary files:** By default, SQLite checkpoints the `-wal` file automatically at `1,000` pages (about `4` megabytes), but it does not truncate the `-wal` file after a checkpoint unless `PRAGMA journal_size_limit` is configured, because overwriting an existing file is faster than appending. [S14] [S16] Furthermore, if a poorly designed local dashboard or backup script holds a long-running read transaction open, **checkpoint starvation** prevents the checkpointer from resetting the `-wal` file, causing `-wal` to grow without bound until all readers close their transactions or a `RESTART` / `TRUNCATE` checkpoint forces completion. [S14] [S16] SQLite can also create up to nine kinds of temporary files (rollback journals, super-journals, `-wal` files, `-shm` shared-memory files, statement journals, `TEMP` databases, view materialisations, transient indices and `VACUUM` databases), which must have reserved headroom on their target directory. [S0604]
4. **Operating-system and service journals (`systemd-journald`):** When a backhaul fails, network retry daemons and connection watchdogs often emit verbose error logs. [S22] Under `journald.conf`, when `Storage=persistent` is active (`the default`), `systemd-journald` stores logs under `/var/log/journal` with `SystemMaxUse=` defaulting to `10` percent of the filesystem (capped at `4` gigabytes) and `SystemKeepFree=` defaulting to `15` percent (also capped at `4` gigabytes), and syncs non-critical logs every `SyncIntervalSec=5m` (`5` minutes) while syncing `CRIT`, `ALERT` and `EMERG` messages immediately. [S22] Although `RateLimitBurst=` defaults to `10,000` messages per `RateLimitIntervalSec=30s` (multiplied by a base-2 logarithm factor of free disk space), `journald` never deletes active journal files during vacuuming, and if another process fills the disk after journal files are created, `journald` stops growing but does not remove existing archived files to restore `SystemKeepFree=`. [S22]
5. **Consistent online backup staging:** As the SQLite Backup API documentation (`S17`) and corruption guide (`S15`) explain, copying a live database file with an external utility while a transaction is active or without its paired `-wal` / `-journal` file risks a corrupted backup. [S15] [S17] Creating a safe local snapshot using `sqlite3_backup_init()` / `sqlite3_backup_step()` or `VACUUM INTO 'filename'` writes a complete second copy of the active database into staging storage before the file is moved to removable or secondary media, temporarily doubling the database footprint on disk. [S15] [S17]
6. **Filesystem free-space reserve:** Flash translation layers and `ext4` multiblock allocators degrade sharply and risk `ENOSPC` during metadata allocation when a partition approaches 100 percent occupancy (`ext4` reserves `2` percent or `4,096` clusters by default in `reserved_clusters` to prevent unexpected `ENOSPC` during delayed allocation). [S29]

The waterfall chart in **Exhibit 24**, `p. 102` reconciles these seven components for an illustrative 8-hour offline design budget starting from the `720`-megabyte raw payload in calculation `C02`. [S08] [S14] [S17] [S22]

EXHIBIT 24 / MEGABYTES (MB, DECIMAL)

## Retention must reserve space beyond event payload



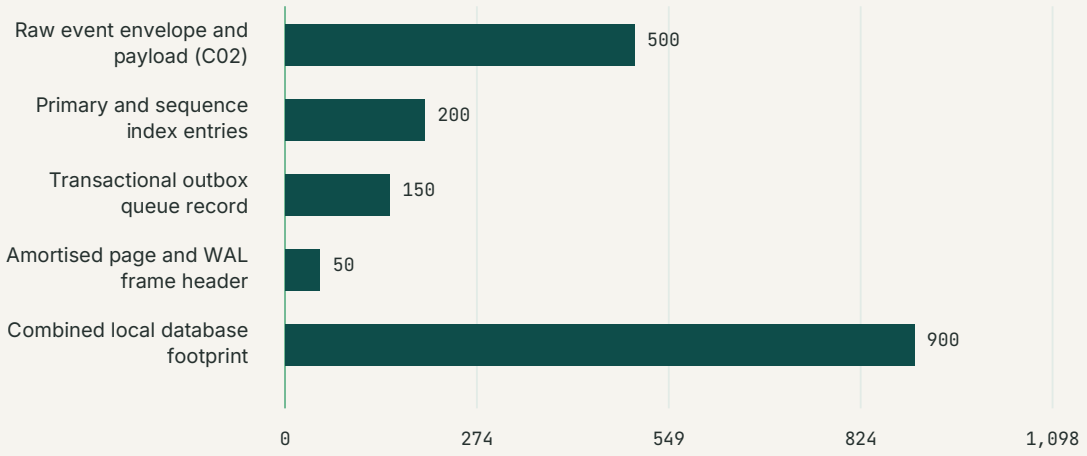

ILLUSTRATIVE [S08] [S14] [S17] [S22] Illustrative storage allocation starting from the 720 MB raw event payload baseline (C02: 50 sources, 1 event/s, 500 bytes, 8 h) and adding illustrative allowances for B-tree indexes, outbox rows, WAL and rollback journals (S14), system journals (S22), online backup staging (S17) and free-space reserve. Site sizing requires measured overhead.

As **Exhibit 24**, p. 102 demonstrates, a 720-megabyte raw event payload requires an illustrative 3,600-megabyte (3.60 decimal gigabytes) dedicated partition budget once indexes (288 MB), the transactional outbox (216 MB), WAL and temporary file headroom (128 MB), a bounded systemd-journald allocation (256 MB), a complete online backup staging slot (1,352 MB, matching the combined 1,224 MB database plus 128 MB WAL allowance) and a 640-megabyte free-space reserve are included. [S14] [S17] [S22] Note that this 3,600-megabyte figure still excludes high-resolution camera frames from automated optical inspection; if defect images are buffered locally, they must be placed on a separate partition governed by its own ring-buffer policy so that image storage can never starve the transactional event ledger. [S03] [S05]

To make the per-event database multiplier inspectable at the record level, **Exhibit 25**, p. 103 breaks down the derived local database footprint per committed event before backup staging and system logs. [S13] [S14] [S31] [S32]

EXHIBIT 25 / BYTES PER COMMITTED EVENT

## Outbox and index structures multiply the raw event footprint




DERIVED [S13] [S14] [S31] [S32] Derived from the illustrative 500-byte raw event payload in calculation C02 plus explicitly assumed local storage structures: 200 bytes for UUID and sequence B-tree index entries (S31, S32), 150 bytes for the atomic outbox state record (S13) and 50 bytes of amortised page-header and WAL frame metadata (S14), totaling 900 bytes per event before backup staging.

**Exhibit 25**, p. 103 shows why provisioning disk capacity from raw JSON or binary message length alone underestimates local database growth by nearly half: adding 200 illustrative bytes of index entries (`event_id` and composite sequence keys), 150 illustrative bytes for the `outbox_queue` coordination row, and 50 illustrative bytes of B-tree page slack and WAL frame headers raises each 500-byte event to 900 bytes inside the live SQLite store (1.80 times the raw payload) even before accounting for online backup copies or operating-system logs. [S13] [S14] [S31]

WORKED EXAMPLE

### 6.19 SIZING A PARTITION AND ITS WATERMARK THRESHOLDS FROM THE C02 BASELINE

Consider an illustrative edge capture node using the C02 register inputs: 50 active station sources, 1 event per second per source, 500 bytes of raw envelope and payload per event, and an 8-hour (28,800-second) design offline interval. [S08] None of these values represent a measured GFT factory result; they illustrate how an engineering team translates an arrival rate into partition boundaries and database pragmas: [S05] [S16]

1. **Event count over the design interval:** Multiply 50 sources by 1 event per second by 28,800 seconds to obtain 1,440,000 events per 8-hour shift disconnection. [S08]
2. **Live SQLite database footprint:** Using the derived 900 bytes per event from **Exhibit 25**, p. 103 (500 bytes payload + 200 bytes indexes + 150 bytes outbox row + 50 bytes page/WAL overhead), 1,440,000 events occupy 1,440,000 * 900 = 1,296,000,000 bytes (1,296 decimal MB, or 1,224 decimal MB of main table/index pages plus 72 to 128 MB of bounded WAL and temporary headroom). [S13] [S14]
3. **Database page cap (`PRAGMA max_page_count`):** At the standard 4,096-byte SQLite page size (S16), allocating an illustrative 2,048 decimal MB (2,048,000,000 bytes) hard ceiling for the main database file corresponds to 4,096-byte page limit `PRAGMA max_page_count = 500000` (500,000 * 4,096 = 2,048,000,000 bytes), while `PRAGMA journal_size_limit = 67108864` caps the post-checkpoint `-wal` file at 64 binary MiB (67,108,864 bytes). [S14] [S16]
4. **Partition layout and watermarks:** On the illustrative 3,600 MB dedicated capture partition from **Exhibit 24**, p. 102, setting a soft warning watermark at 70 percent occupancy (2,520 MB used) alerts the local supervisor while 1,080 MB remains, setting a high watermark at 85 percent (3,060 MB used) suspends non-essential diagnostic and verbose telemetry intake while forcing a `PRAGMA wal_checkpoint(TRUNCATE);`, and setting a critical stop/fallback threshold at 95 percent (3,420 MB used) transitions the station cleanly to authorised manual paper-and-spreadsheet capture before an unhandled disk-full crash can occur. [S03] [S05] [S16]

### 6.20 WATERMARK THRESHOLDS MUST PROTECT PRODUCTION TRACE RECORDS BEFORE DIAGNOSTICS

When an outage extends beyond the design envelope or a downstream receiver remains unreachable, storage exhaustion must follow a pre-approved priority policy rather than a first-come, first-failed crash. [S03] [S05] First, the operating system should place `/var/log/journal` (with explicit `SystemMaxUse=` and `SystemKeepFree=` caps in `journald.conf`) and any camera image spool on separate partitions or logical volumes from the `/var/lib/gft-ledger` SQLite directory, so a runaway daemon log or image burst cannot consume the database's free blocks. [S05] [S22] Second, the capture service must monitor free partition space, database page count (`PRAGMA page_count;` versus `PRAGMA max_page_count;`) and `-wal` file size continuously, running periodic `PRAGMA wal_checkpoint(RESTART);` or `PRAGMA wal_checkpoint(TRUNCATE);` during brief reader gaps so that dashboard queries never cause checkpoint starvation. [S14] [S16]

Third, when the high watermark is crossed, the intake controller sheds load strictly by service class: discretionary debug traces and high-frequency ambient sensor readings are dropped first, while mandatory bundle completions, quality dispositions and safety-relevant maintenance alarms retain exclusive access to the remaining database reserve. [S03] [S05] Crucially, the capture service must never silently overwrite or delete unacknowledged production trace events to free space. [S03] [S05] Every shed diagnostic batch and any rejected production event must increment a persistent, monotonic `rejected_event_count` register stored in a fixed-size summary row and display an immediate visual banner on the local supervisor console instructing operators to switch to the controlled manual continuity ledger defined in Chapter 3. [S03] [S05]

6.21

## A durable record needs a demonstrated recovery path

Claiming that a local capture node is crash-safe because it uses SQLite and an industrial solid-state drive is an unverified assumption until the exact software image, pragma configuration, filesystem mount and hardware storage path are subjected to repeatable fault experiments. [S06] [S28] Commissioning must verify that interrupted transactions roll back cleanly, that acknowledged transactions survive on disk, and that post-recovery replay reconciles every single expected identifier against a deterministic test ledger. [S13] [S28]

### 6.22 SIMULATED FAULT INJECTION AND PHYSICAL POWER CUTS TEST DIFFERENT ASSUMPTIONS

SQLite's own testing documentation explains why library-level verification and site hardware qualification play distinct roles: SQLite's core library (`155.8` thousand source lines of C code as of version `3.42.0`, backed by `92,053.1` thousand lines of test code across the TCL, TH3, SLT and `dbsqlfuzz` harnesses) achieves 100 percent branch and Modified Condition/Decision Coverage (`MC/DC`), but **because it is impractical to run millions of crash tests using physical power failures, SQLite's own crash testing is performed in simulation**. [S28] In the TCL crash harness, a custom Virtual File System intercepts writes, spawns a child process that crashes at random I/O points, and randomly reorders and corrupts unsynchronised writes to mimic a buffered filesystem; in the embedded TH3 harness, an in-memory VFS snapshots the filesystem after a set number of I/O operations, reverts to the snapshot, injects characteristic power-loss file damage, and runs `PRAGMA integrity_check` to confirm that every transaction either completed or rolled back completely. [S16] [S28]

That simulated VFS testing proves the correctness of SQLite's internal journalling logic under its stated hardware assumptions, but it cannot test whether a specific factory edge computer's flash controller lies about `fsync()`, whether its Linux kernel `ext4` mount honors write barriers, or whether an application developer accidentally left `PRAGMA synchronous=NORMAL` or an unawaited asynchronous promise in the capture service code. [S15] [S16] [S28] [S29] Therefore, GFT's proposed commissioning

method combines two complementary test tiers before any pilot line is accepted (noting that this white paper defines the verification method and does not claim that GFT has executed these tests at a client plant): [S06] [S28]

1. **Software fault-injection exercises:** While running a deterministic event generator at the design rate (for example, `50` simulated sources emitting sequential `UUIDv7` identifiers and known `payload_sha256` digests), the test harness terminates the capture process (`SIGKILL`), the local broker (`SIGKILL`) and the operating system (`echo c > /proc/sysrq-trigger` or virtual machine hard reset) at four specific points in the transaction lifecycle: before `INSERT`, after `INSERT` but before `COMMIT`, immediately after `COMMIT` returns inside the server process before the HTTP/IPC response is written to the socket, and during an active `WAL` checkpoint or online backup (`sqlite3_backup_step()`). [S13] [S17] [S28]
2. **Qualified physical power-cut exercises:** Under the supervision of the site electrical engineer and within the equipment's rated switching envelope, the target edge computer and its local storage media undergo controlled mains and DC supply interruptions while actively committing events and running checkpoints. [S03] [S28] After power returns and the system boots, the test harness executes `PRAGMA integrity_check;`, inspects `journalctl` for filesystem or block-layer I/O errors (`ext4` journal replay warnings or barrier failures), and reconciles the local database and downstream consumer ledger against the generator's ground-truth log. [S16] [S22] [S29]

To make the pass/fail criterion unambiguous, the test harness compares four sets of identifiers for every run: the generator's **sent ledger**, the generator's **locally acknowledged ledger** (every `event_id` for which the capture endpoint returned a durable local ACK prior to power loss), the edge node's recovered `local_events` **and** `outbox_queue` **tables**, and the **remote committed ledger** after replay completes. [S13] [S31] The illustrative audit table in **Exhibit 26**, `p. 107` demonstrates how five representative event outcomes are classified during commissioning reconciliation. [S13] [S16] [S31] [S32]

EXHIBIT 26 / ILLUSTRATIVE AUDIT LEDGER

## A known event ledger makes loss and duplication visible

| Event ID and seq | Local commit | Remote commit | Replay count | Audit classification |
|---|---|---|---|---|
| 0192dce0-0001 (seq 401) | Committed (FULL) | Committed | 0 | Verified normal delivery |
| 0192dce0-0002 (seq 402) | Committed (FULL) | Committed | 2 | ACK lost; retry deduplicated |
| 0192dce0-0003 (seq 403) | Rolled back (NORMAL) | Missing | 0 | Power-loss rollback after ACK |
| 0192dce0-0004 (seq 404) | Aborted pre-commit | Missing | 1 | Unacknowledged; re-sent as 405 |
| 0192dce0-0005 (seq 405) | Committed (FULL) | Pending outbox | 1 | Retained locally for replay |

ILLUSTRATIVE [S13] [S16] [S31] [S32] Illustrative commissioning ledger comparing deterministic source sequence numbers and UUIDs (S31), UTC timestamps (S32) and atomic commit states (S13, S16). No factory test has been executed; values illustrate reconciliation logic.

In **Exhibit 26**, p. 107, five distinct behaviours become immediately visible to the commissioning engineer: [S13] [S16] [S31]

- **Sequence** `401` **(**`0192dce0-0001`**):** Committed locally under `synchronous=FULL`, acknowledged to the station, and forwarded to the remote ledger before the fault (`0` replays); this is normal baseline operation. [S13] [S16]
- **Sequence** `402` **(**`0192dce0-0002`**):** Committed durably on disk before a fault severed the response path; the station retried twice using the same `UUIDv7` (`2` replays), and idempotent deduplication returned the original commit receipt without creating duplicate rows in either `local_events` or the remote ledger. [S12] [S13] [S31]
- **Sequence** `403` **(**`0192dce0-0003`**):** Illustrates a **critical acceptance failure** that occurs when a database is misconfigured with `WAL` and `synchronous=NORMAL`: the capture service returned a local acknowledgement to the station, power failed before the next checkpoint sync, and the committed transaction rolled back on reboot, leaving a permanent gap (`seq 403` missing locally and remotely) despite a positive acknowledgement. [S16] Any occurrence of an acknowledged `event_id` missing after recovery fails the durability gate. [S06] [S16]
- **Sequence** `404` **and** `405` **(**`0192dce0-0004` **and** `0192dce0-0005`**):** Sequence `404` was interrupted mid-transaction before commit and never acknowledged; upon restart, the station or operator re-submitted the work (or retained the original `event_id` depending on peripheral buffer capability), which committed durably as `0192dce0-0005` and sits in `outbox_queue` awaiting backhaul restoration. [S08] [S13]

**WARNING**

**6.23 CLEAN SHUTDOWNS AND SIMULATED CRASHES CANNOT CERTIFY UNVERIFIED FLASH HARDWARE**

A commissioning drill that tests offline capture only by unplugging an Ethernet cable or issuing a clean `systemctl stop` command proves nothing about power-loss durability, because clean service termination allows the operating system and Mosquitto (`persistence true`) to flush volatile memory caches to disk before exit. [S20] [S21] Likewise, software-only `SIGKILL` tests leave the Linux kernel page cache intact and never test whether the physical SSD or eMMC controller honours `fsync()` barriers or powersafe sector overwrite during a sudden voltage collapse. [S15] [S29] [S0601] To resolve this constraint before scaling a factory deployment, procurement must require industrial storage media with documented power-loss protection (or verified power-safe overwrite) and witness at least ten physical power-cut cycles under active write load with zero lost acknowledged `event_id` records and a clean `PRAGMA integrity_check` result on every reboot. [S15] [S16] [S28]

Once the local capture boundary is proven to commit every acknowledged event and its outbox entry atomically to a durable, bounded local store, the next engineering problem is managing the accumulated backlog across extended backhaul outages and draining it without overwhelming the restored link, which is the subject of Chapter 7. [S08] [S13]

*An acknowledgement is a promise that the record will still be there when the lights come back on.*

**CHAPTER SOURCE NOTES**

S13 SQLite. 2026-10-30.

S16 SQLite. 2026-10-30.

S01 World Bank. 2025.

S24 IETF / RFC Editor. 2022.

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

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

S14 SQLite. 2026-10-30.

S20 Eclipse Mosquitto. 2026-10-30.

S28 SQLite. 2026-10-30.

S31 IETF / RFC Editor. 2024.

S32 IETF / RFC Editor. 2002.

S03 NIST. 2023.

S22 systemd project. 2026-10-30.

S19 OASIS. 2019-03-07.

S05 NIST. 2010.

S15 SQLite. 2026-10-30.

S0602 SQLite. 2022-04-18.

S08 IETF / RFC Editor. 2007.

S0603 Linux man-pages project. 2026-09-09.

S12 IETF / RFC Editor. 2022.

S0604 SQLite. 2025-06-12.

S06 NIST. 2016.

S29 Linux kernel project. 2026-10-30.

S0601 SQLite. 2025-05-31.

S17 SQLite. 2026-10-30.

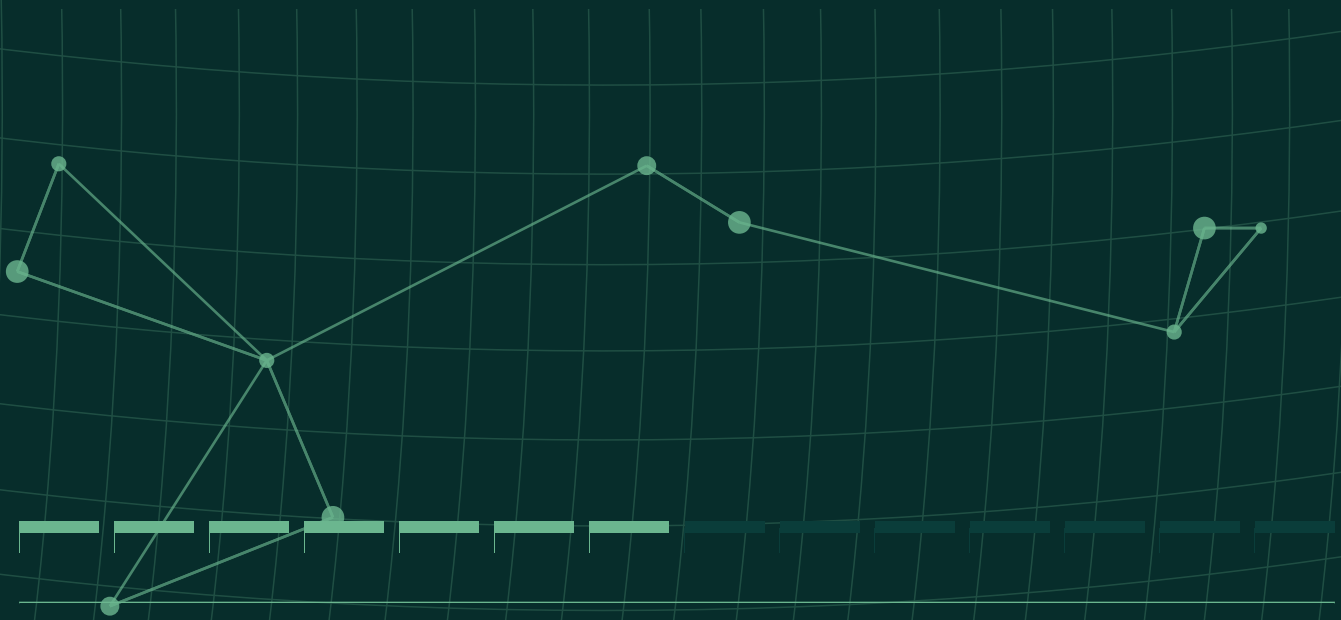
S21 systemd. 2026-10-30.

07

PART II: KEEP THE LOCAL RECORD ALIVE / CARRY THE BACKLOG

# A bounded queue makes disconnection manageable

*Offline retention and recovery capacity must be designed together so a restored connection can clear its accumulated work.*



CHAPTER 07
DECISION TEST

Offline retention and recovery capacity must be designed together so a restored connection can clear its accumulated work.

A bounded queue makes disconnection manageable only when retention, recovery and custody are designed together. The store must hold the declared offline load, the restored path must clear that load while fresh events arrive, and each record must remain identifiable until the remote ledger confirms its durable acceptance. Persistent storage is the starting point, not the whole answer. [S08] A large disk attached to a slow or intermittent uplink merely postpones saturation. Conversely, a fast uplink cannot recover events that an unbounded or volatile local queue has already lost. GFT therefore treats the queue as an operating contract with explicit limits, visible overload behaviour and a reconciled end state. The calculations below are illustrative design methods, not measurements from a GFT deployment.

7.1

## Calculate retention from event production and outage duration

Retention begins with a byte balance at the boundary that accepts production records. For each event class, measure the number of sources, their event rates in relevant production states, the serialised payload distribution and the design offline interval. The basic payload balance is `B = N x r x s x T`, where `N` is the source count, `r` is events per source per second, `s` is bytes per event and `T` is the offline interval in seconds. This is deterministic accounting. It does not predict how long a site outage will last, and it must not turn a national outage frequency into a site duration assumption.

Use separate rates for normal production, start-up, alarm bursts and recovery. A single shift average can conceal the very burst that fills the queue. Where records arrive in batches, retain both the batch count and the distribution of records per batch. Where payloads vary, calculate with an agreed upper design quantile and retain the observations behind it. The operator also needs an event-count budget, because a database can reach a row, inode or transaction limit before it reaches its byte limit. The measurement boundary remains the one established in the site survey: an event that never reaches a powered source or capture service cannot be rescued by downstream storage. [S03]

The paper's illustrative payload case uses `N = 50` sources, `r = 1` event per source per second, `s = 500` bytes per event and `T = 28,800` seconds, equivalent to an illustrative 8-hour outage. The live arrival rate is `50 x 1 x 500 = 25,000` bytes per second. The accumulated payload is `50 x 1 x 500 x 28,800 = 720,000,000` bytes, or `0.72` decimal GB, representing `1,440,000` illustrative events. Every input is assumed solely to demonstrate calculation C02. The persistent-store principle supports the architecture, but it does not supply these site inputs. [S08]

### 7.2 PROVISIONED STORAGE MUST EXCEED SERIALISED PAYLOAD

The result is a payload subtotal, not a disk specification. A durable implementation also stores event metadata, indexes, replay state and integrity fields. Its database may maintain a write-ahead log, while its broker and operating system maintain their own persistence files. SQLite's WAL documentation explains that changes are appended to a separate log before checkpointing transfers them into the main database, so the

main database file alone is not the full storage footprint. [S14] The system journal also needs configured limits and free-space treatment rather than permission to compete unnoticed with the production outbox. [S22]

Do not respond by applying a generic overhead multiplier. Benchmark the intended schema and software on the intended storage. Measure a representative event mix, index definitions, checkpoint behaviour, broker persistence, backup staging and log verbosity. Then calculate `provisioned capacity = measured stored bytes for the design load + recovery workspace + operating reserve`. Recovery workspace matters because a full partition is a poor place to perform integrity checks, compaction or restoration. Record the test version and configuration with the result. A later schema index or logging change can invalidate an apparently conservative estimate.

Images and other large attachments require a separate budget. Their byte rates can dominate structured events, yet their retention rules may differ. Keep a structured quality decision and its attachment reference in the event model, but account for the media store, its metadata and its recovery traffic independently. If the quality authority permits sampling, the retained decision must expose that policy rather than imply that every attachment remains available. This is a proposed evidence-completeness control, not a claim that sampling is acceptable for every inspection duty.

Retention also has a time dimension. A queue can have ample free bytes while its oldest mandatory event is already too old for the decision it must support. Capacity review should therefore include bytes used, events retained and oldest unacknowledged age. The design interval ends only when remote acceptance has been restored sufficiently to release custody, not when a modem first reports a signal. Unpredictable bandwidth and lossy channels make that distinction material to recovery planning. [S18]

> **GFT NOTE**
>
> **Size from a replayable event sample, not from a disk catalogue.** GFT recommends capturing representative production, alarm and restart periods; calculating payload and stored-byte budgets separately for structured events and media; and signing the design offline interval with operations and quality owners. The illustrative `0.72` decimal GB result in C02 is only the serialised payload. Procurement should require the measured overhead evidence, the recovery reserve and the test configuration behind any final capacity figure.

7.3

## Separate the durable source queue from the uplink spool

A useful queue design names the component that owns each record at every moment. GFT distinguishes the durable source queue from the uplink spool. The source queue accepts custody from the production application and retains the authoritative local copy until terminal acceptance is proved. The uplink spool manages transmission attempts, connection state and in-flight messages. A site aggregation store may sit between them when several cells must share a history or a remote route. The remote consumer becomes the terminal owner only after its application transaction commits.

This separation follows an architectural lesson from delay-tolerant networking: persistent storage and explicit responsibility help a system cross interruptions. [S08] Bundle Protocol Version 7 provides another reference for store-and-forward operation, but the factory is not thereby required to adopt that protocol. [S07] The useful decision is where custody sits and which acknowledgement moves it. A transactional outbox, a broker and a remote ingestion service can implement the boundary with ordinary industrial software if their combined behaviour satisfies the contract.

### 7.4 EACH STAGE NEEDS A DIFFERENT ACKNOWLEDGEMENT

At the producer, local acknowledgement follows the atomic commit of the business event and its pending-delivery record. SQLite defines atomic commit as all changes within a transaction occurring or none occurring, which is the required starting boundary for this proposed outbox pattern. [S13] The previous chapter owns the storage durability settings. Here the important point is that the production application does not need to wait for the remote route before it can record accepted work.

The forwarder then reads pending records and writes them into the uplink path. Its send cursor may advance when a broker or remote endpoint accepts transport responsibility, but that cursor is not permission to erase the source record. Broker persistence is configured behaviour: Mosquitto documents that connection, subscription and message data are written to its persistence database when enabled. [S20] Even with that setting, the broker remains one custody stage. Its acceptance does not prove that the downstream production ledger has committed the event.

The remote consumer issues a distinct application receipt after durable ingestion. The source queue records that receipt against the relevant event identities and advances a separate committed cursor. Three positions therefore matter: the next record to read, the records currently in flight and the records whose remote commit has been proved. After a crash or spool reset, the forwarder can resend from the last reconciled position. Repetition is expected and is handled by the ingestion contract below.

A small plant may not need a separate site aggregation store. Adding one can centralise retention and prioritisation, but it also creates another queue, another capacity budget and another recovery dependency. GFT recommends the fewest custody stages that can support the operating boundary. Every retained stage must have an owner, a durable acknowledgement meaning, a replay cursor and a saturation response. Removing an unnecessary hop reduces the number of places where a record can appear delivered while remaining uncommitted downstream.

The source queue must also remain independently inspectable. A supervisor should be able to distinguish locally committed, in-flight, remotely committed, quarantined and policy-degraded records without relying on the uplink. That visibility prevents a transport indicator from masquerading as production certainty. It also supplies the evidence needed when recovery requires manual reconciliation rather than another automatic retry.

7.5

# Use idempotent ingestion to make duplicates harmless

A sender that cannot tell whether the previous attempt committed must either risk loss by moving on or risk duplication by trying again. For mandatory production records, GFT chooses controlled retransmission and makes the receiver idempotent. Idempotency means that repeating the same identified operation has the same intended business effect as performing it once. HTTP semantics use the same basic distinction for request methods, while also warning that idempotent intent does not erase every possible side effect or log entry. [S12] The factory contract must therefore define the business state that may change, not merely label a retry safe.

MQTT does not remove this obligation. Its at-least-once quality permits duplicates, and its other delivery handshakes operate within the messaging protocol rather than proving a downstream production transaction. [S19] A broker can accept a publication, the consumer can update a ledger, and the return acknowledgement can still be lost. If the next delivery appends a second production event, the queue has transported one physical event into two business actions. Protocol delivery and business idempotency are separate acceptance boundaries.

### 7.6 IDENTITY MUST SURVIVE EVERY RETRY

The producer assigns the event identity once and preserves it through every replay. A practical envelope includes the site, source, event identifier, boot identity and source sequence. RFC 9562 states that a UUID requires no central registration process, which makes it a candidate for local generation without an online allocator. [S31] That property does not excuse weak random generation, reused counters or accidental regeneration during retry. The implementation must test identity persistence across process restart, device replacement and restoration from backup. Clock interpretation remains a separate chapter; the queue needs stable identity even when event order is disputed.

At the remote consumer, deduplication and the business mutation belong in one atomic transaction. The receiver attempts to register the event key in a processed-event table with a uniqueness constraint. If registration succeeds, the same transaction applies the production or quality mutation and prepares the receipt. If the key already exists with the same payload digest, the receiver returns the earlier committed result without applying the mutation again. Atomicity is what keeps the deduplication mark and business state from diverging. [S13]

A repeated key with a different payload is not a harmless duplicate. It is an identity conflict. The receiver should preserve both the original committed state and the conflicting evidence, place the item in a quarantine state and notify the responsible reconciliation role. Automatically overwriting the first event would destroy history. Automatically accepting the second would double-count work. Silently dropping it would hide an implementation or integrity fault.

The receipt should identify what was committed, not merely what was received. For a batch, it can return committed event identifiers or a source-sequence watermark accompanied by a count and digest. A watermark is valid only across a contiguous,

fully resolved range. Priority replay and quarantined records can create holes, so the sender must not treat the highest observed sequence as proof that every earlier sequence arrived. The response to a repeated batch should reproduce the committed result, allowing a sender whose earlier receipt was lost to finish reconciliation.

Idempotency does not establish causal order. A quality decision may arrive before the production event it references when queues drain by priority. The consumer must either resolve dependencies by source sequence, defer the dependent projection or record an explicit unresolved state. It must not invent a parent record or treat arrival order as event order. These are proposed application rules; the transport standard does not supply them.

7.7

## Queue-drain capacity must exceed live arrivals

An offline interval creates a recovery debt. When the route returns, the uplink must carry both new production and that debt. Let `lambda` be the live payload arrival rate in bytes per second, `B` the queued payload at reconnection and `C` the useful capacity for remote-committed payload after protocol overhead, retransmission and competing traffic. While the path remains up and capacity exceeds arrivals, net drain is `C - lambda`, and clearance time is `T = B / (C - lambda)`. Constrained-network documentation explicitly recognises unreliable channels and limited, unpredictable bandwidth, so nominal carrier speed is not an acceptable substitute for measured useful capacity. [S18]

If `C` equals `lambda`, the restored link only maintains the current position. If it is lower, the backlog continues to grow even though the link is online. A recovery specification should therefore reserve capacity above the live rate and test that reserve with the actual batch sizes, acknowledgements and remote commit path. The denominator is an application mass balance, not a network marketing measure.

**WORKED EXAMPLE: HOLD THE BACKLOG, THEN CLEAR IT**

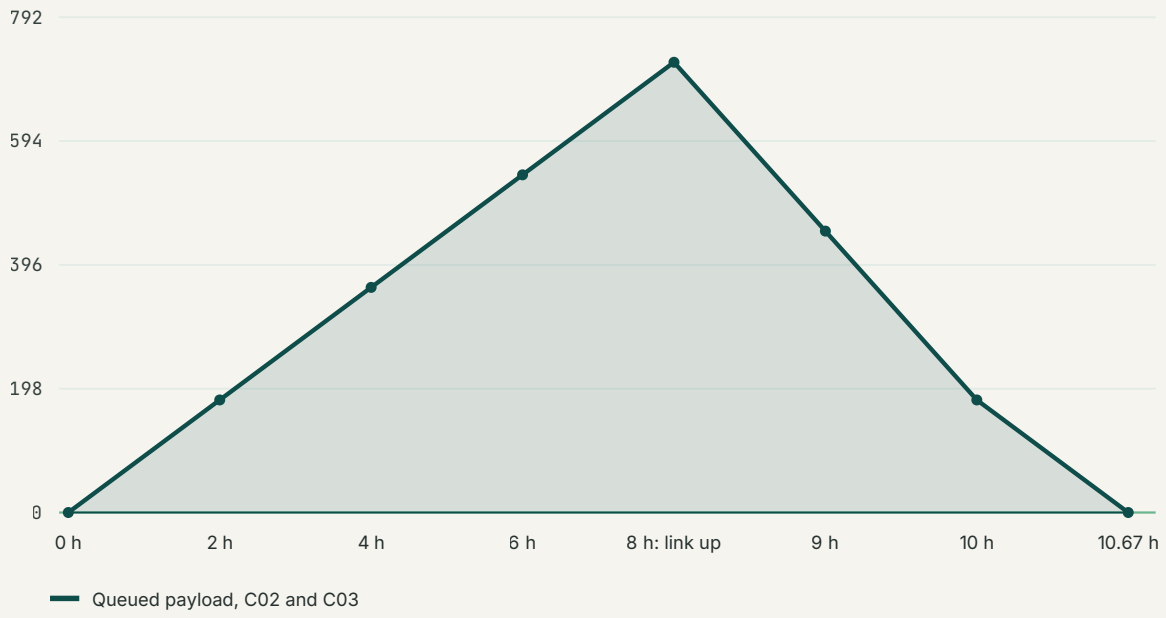
Consider an illustrative sewing and inspection cell whose structured event load follows C02. All values are assumed to demonstrate the method.

1. **Accumulate the payload.** Fifty illustrative sources each emit one event per second with a serialised size of 500 bytes. Arrival is `50 x 1 x 500 = 25,000` bytes per second. Across an illustrative 8-hour outage, or `28,800` seconds, the queue accepts `720,000,000` payload bytes. This excludes storage overhead and media.
2. **Measure useful recovery capacity.** After reconnection, an illustrative path sustains `100,000` bytes per second of remotely committed event payload after all transport and application overhead. Live arrivals consume `25,000` bytes per second, leaving `75,000` bytes per second for the backlog.
3. **Calculate steady-link clearance.** The illustrative drain time is `720,000,000 / 75,000 = 9,600` seconds, or 2 hours 40 minutes. During that period, the path also carries `240,000,000` bytes of new events, so total transmitted payload is `960,000,000` bytes. This is calculation C03, not a carrier performance claim.
4. **Close custody.** The source queue does not declare success at elapsed clearance time. It compares committed receipts with its expected event identities and retains any missing, conflicting or quarantined records for resolution. The arithmetic predicts byte movement; reconciliation establishes record completion.

**Exhibit 27**, `p. 117` shows the illustrative accumulation and clearance trajectory. Its horizontal labels mark event times, not equal plotting intervals.

EXHIBIT 27 / DECIMAL MB OF QUEUED PAYLOAD

## An outage creates a backlog that a fast link must clear




DERIVED [S08] [S18] Derived from illustrative C02 and C03 assumptions: arrival 25,000 bytes/s for 28,800 s creates 720 MB. Useful restored capacity 100,000 bytes/s less live arrivals 25,000 bytes/s gives net drain 75,000 bytes/s and clearance 9,600 s. Elapsed labels are categorical and unequally spaced; 10.67 h is rounded from 10 h 40 min. Storage overhead and later interruptions are excluded.

The area falls only after useful capacity exceeds live arrivals. It returns to zero under the stated steady-link assumption; it does not show that remote records are reconciled at that instant. Nor does it include indexes, journals or image attachments. Those omissions are deliberate so that the payload balance remains inspectable.

### 7.8 REPEATED INTERRUPTION REQUIRES A BALANCE, NOT A SINGLE QUOTIENT

A weak route may drop again before clearance. For each interval with approximately constant conditions, update the illustrative balance as `B_next = max(0, B_previous + (lambda - C) x interval_seconds)`, using `C = 0` while the path is unavailable. Any residual queue then enters the next interval. This piecewise method is less elegant than one drain quotient, but it represents the failure behaviour that matters.

For example, start with the illustrative `720` MB backlog and the `75,000` bytes-per-second net drain above. One hour of stable recovery removes `270` MB, leaving `450` MB. A further illustrative 2-hour interruption at `25,000` bytes per second adds `180` MB, raising the queue to `630` MB. At the same restored net drain, that residue needs `8,400` seconds, or 2 hours 20 minutes, to clear. All values are illustrative. Repeated short recovery windows can leave a residual backlog even when each individual link test appears successful.

**Exhibit 28**, `p. 118` generalises the steady-link case across declared illustrative arrival rates and useful capacities. The outage duration remains fixed within this sensitivity, so differences arise from arrival and capacity rather than from a new outage forecast.

EXHIBIT 28 / RECOVERY HOURS; -1 MEANS NON-DRAINING

## A link that cannot exceed arrivals never catches up

| | 25 kB/s link | 50 kB/s link | 100 kB/s link | 150 kB/s link |
|---|---|---|---|---|
| 15 kB/s arrival | 12 | 3.43 | 1.41 | 0.89 |
| 25 kB/s arrival | -1 | 8 | 2.67 | 1.6 |
| 40 kB/s arrival | -1 | 32 | 5.33 | 2.91 |
| 50 kB/s arrival | -1 | -1 | 8 | 4 |

Cell scale: -1 to 32

ILLUSTRATIVE [S18] Illustrative sensitivity after an 8 h outage. For each row, B = arrival x 28,800 s and T = B / (capacity - arrival), shown in hours only where capacity exceeds arrival. Cells marked -1 have capacity less than or equal to arrival and never drain under steady conditions. Arrival rows are 15, 25, 40 and 50 kB/s, producing 432, 720, 1,152 and 1,440 MB. The 25 kB/s arrival and 100 kB/s capacity cell is C03: 2.67 h.

Recovery time rises sharply as useful capacity approaches the arrival rate from above. This is why a small nominal capacity margin is not a practical recovery plan. Site tests should vary production rate, batching and competing traffic, then retain measured useful-capacity distributions rather than one successful speed test. If the operator requires a clearance limit, procurement must reserve sufficient headroom or reduce replay load by agreed priority and media policies. The arithmetic can expose the requirement; only the implemented path can prove it.

7.9

## Priority and expiry rules must be agreed before saturation

Every queue has a saturation point, even if the design never reaches it during the expected outage envelope. The time to decide what happens there is before production, not during a disk-full alarm. NIST contingency guidance connects requirements and priorities to operational impact, which supports making this an operations and quality decision rather than a database default. [S05] GFT proposes a signed admission and expiry policy for each record class, with a named authority, a local action and an auditable result.

Mandatory traceability events, quality dispositions and alarm transitions should not share an undifferentiated first-in, first-out partition with discretionary diagnostics. Reserve capacity for the records whose loss would make production history incomplete. Apply admission control before the reserved space is consumed. A mandatory record may justify refusing a new diagnostic record, moving to an approved manual fallback or holding the next safe production cycle. It never justifies silently deleting an earlier event needed to explain the same bundle.

Expiry also needs a semantic rule. Delay-tolerant protocols use lifetimes to limit storage and forwarding obligations, but a factory cannot infer that an expired production record has become irrelevant merely because a timer elapsed. [S07] Expiry is acceptable only where the record owner has defined the decision life and the permitted degradation. A stale diagnostic sample may be summarised. A quality hold whose parent bundle remains unresolved may become more important with age.

Where media sampling is allowed, retain the structured disposition, the event reference, the active sampling rule and an evidence-completeness flag. The reviewer must be able to distinguish an attachment that was never required from one omitted by policy and one lost through an unexplained fault. Any reduced inspection coverage needs quality approval and visible operating status; storage pressure must not quietly redefine a passed inspection. OT reliability and safety requirements remain controlling constraints on the proposed degradation. [S03]

**Exhibit 29**, p. 119 states a compact illustrative policy. Its classes are design prompts, not universal regulatory categories or retention periods.

EXHIBIT 29 /

## Retention priorities make overload behaviour reviewable

| Event class | Custody rule | Pressure response | Escalation |
|---|---|---|---|
| Trace and quality | Retain until remote commit and reconciliation | Refuse discretionary intake; never silently discard | Quality and operations owners choose fallback or controlled hold |
| Alarms and interventions | Retain state transitions and acknowledgements | Coalesce only approved repetition; preserve transitions | Shift supervisor and engineering review |
| Inspection media | Keep decision references and completeness state | Apply approved sampling or suspend discretionary media | Quality owner authorises reduced evidence |
| Diagnostics | Bound by class quota | Summarise or shed by approved rule; write a gap marker | OT platform owner records the degradation |
| Recovery evidence | Protect audit, quarantine and reconciliation records | Reserve space before ordinary replay traffic | Recovery owner resolves exceptions before release |

ILLUSTRATIVE [S05] [S08] Illustrative GFT policy template derived from contingency-priority and persistent-store principles. Actual classes, authorities and actions require site approval. No regulatory retention period or permission to discard mandatory evidence is asserted.

Priority creates its own risk. Permanent preference for fresh alarms can starve old diagnostics that explain the original fault, while aggressive media shedding can remove the evidence needed to investigate a quality escape. The policy therefore needs bounded fairness, minimum diagnostic context and periodic review of starvation by class. A queue dashboard should show both what was protected and what was degraded. Otherwise, a successful mandatory-record total can conceal an unexamined loss of explanatory evidence.

7.10

# Backpressure must reach a responsible person

Backpressure becomes useful only when someone can act before the mandatory record boundary fails. A remote dashboard cannot provide that warning during the disconnection that caused the queue to grow. The local console must therefore expose occupancy by class and total, oldest pending age, rejected or degraded event counts, and predicted time to exhaustion. These are different signals. Low occupancy with an old quality event suggests a blocked cursor or poison record. High occupancy with a young oldest event suggests rapid intake. Both require investigation, but not necessarily the same action.

Estimate time to exhaustion from remaining usable quota divided by net storage growth, with the measurement window and zero-growth treatment displayed. A short window can reveal an alarm storm; a longer window can suppress a transient burst. Neither is a forecast of outage duration. During recovery, also display the net drain condition. If useful uplink capacity does not exceed arrivals, the interface should say that the queue is not catching up rather than present an implausibly large completion estimate.

GFT proposes staged responses linked to the site's service classes. An advisory state confirms local recording and asks the supervisor to investigate the route. A pressure state applies the approved diagnostic and media policy, checks free recovery workspace and prepares manual continuity. A critical state protects mandatory intake and escalates to the operations and quality owners. If mandatory traceability can no longer be maintained, the approved response is either validated manual recording or a controlled hold at a safe process boundary. The threshold is site-specific; this chapter prescribes no universal stop percentage.

> **WARNING**
>
> **Do not connect queue saturation directly to an uncoordinated emergency stop.** The digital queue must not issue a hazard-blind machine command. Process safety, operations and quality responsibilities must define the safe cycle boundary and manual alternative. OT continuity decisions must preserve process reliability and safety, not merely keep software available. [S03]

Every alert needs an owner, acknowledgement and closure record. Supervisor acknowledgement means that the condition has been seen, not that the backlog is healthy. Closure requires the relevant recovery evidence: route restored, queue draining, degraded counts explained and mandatory events reconciled. Recovery playbooks and their tests should capture these actions so that an actual interruption does not become the first rehearsal. [S06]

7.11

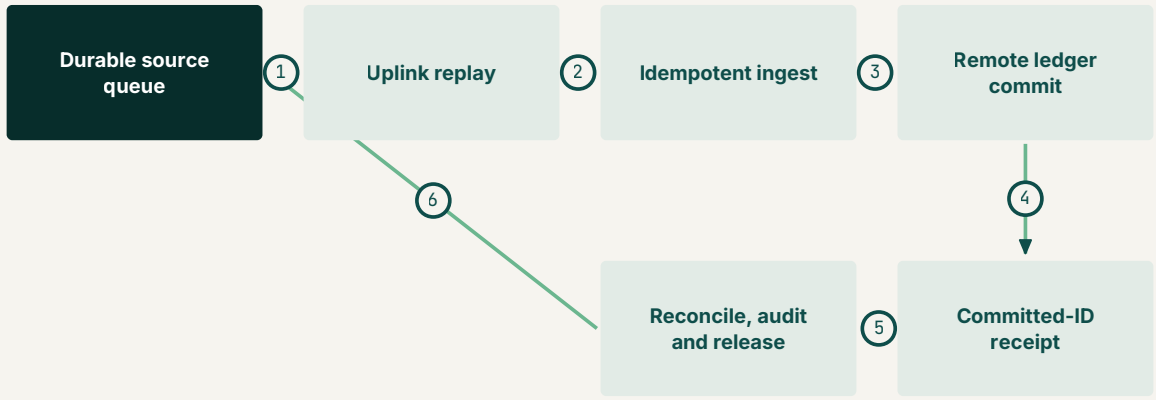
## Replay success is a reconciled ledger, not an empty buffer

Replay finishes when expected identities, remote commits and local custody agree. An empty transport spool shows only that the spool no longer holds work. It does not show that a remote consumer committed every production event, nor that dependent business actions occurred once. The source queue must retain unresolved records through duplicate responses, reordered batches and repeated interruption until an application receipt and local reconciliation close the gap. This is the final custody boundary, not an optional reporting enhancement.

**Exhibit 30**, p. 121 shows the proposed closed loop. Transport acknowledgement advances transmission state; the remote committed receipt and audit result govern retention release.

EXHIBIT 30 /

### Reconciliation closes the delivery loop




1. Pending identities
2. At-least-once batches
3. One atomic business effect
4. Durable commit result
5. IDs, counts and exceptions
6. Release only reconciled custody

ILLUSTRATIVE [S12] [S13] [S19] Illustrative GFT custody design. Atomic commit supports the remote transaction boundary; idempotent semantics support safe retry; MQTT delivery alone is not terminal business acceptance. The diagram is not a claim of an implemented GFT system.

The receipt should resolve identities rather than offer a generic success status. The local reconciler compares the committed set with the expected set, checks counts and payload digests where used, and records exceptions. Matching totals alone are insufficient: two different event sets can have the same count. Where a contiguous

watermark is used, every gap must be explained by a committed record, an authorised policy state or an open exception. Unresolved items remain in custody and visible to the operator.

A quarantined item must not block unrelated records forever. Continue other eligible work while keeping the poison record, its reason and its raw evidence bounded and visible. Conversely, do not advance a contiguous deletion watermark across that hole. The system needs both progress and an honest incomplete state. After restoration from backup or replacement of a forwarder, rebuild the send position from retained source records and receipts rather than trusting a volatile in-flight list.

> **ACCEPTANCE EVIDENCE FOR THE QUEUE CONTRACT**
>
> Test the contract against a known illustrative event ledger before site acceptance. Inject duplicates and lost receipts, then verify one business effect and repeatable committed responses. Reorder priority batches and confirm that dependency rules prevent false causal order. Introduce a corrupted or conflicting identity and verify bounded quarantine without watermark advancement. Restrict capacity below live arrivals and confirm that local backpressure reaches the responsible role. Interrupt an active drain, restart the components and reconcile expected, local and remote identities. Retain configurations, event ledgers, counters and operator actions as recovery evidence; NIST recovery guidance explicitly includes testing and improvement of recovery planning. [S06]

Release storage in a controlled maintenance operation after reconciliation, not through an unexamined delete on every receipt. Keep a compact audit record linking the local range, remote receipt, exceptions and release decision. The retention period for that audit belongs to the site's governance policy. This chapter requires its existence and meaning, not an invented universal duration.

The resulting design gives the edge-architecture tool four linked inputs: measured arrival distributions, stored-byte overhead, useful recovery capacity and class-specific saturation rules. It gives traceability design a stable event identity and gives inventory analysis an explicit point at which replayed events may affect stock. None of these interfaces establishes operating performance by itself. The queue is accepted only when the implemented system preserves mandatory records and closes its ledger under the declared faults.

*A queue earns deletion only through proof of arrival, durable acceptance and reconciliation.*

**CHAPTER SOURCE NOTES**

S08 IETF / RFC Editor. 2007.

S03 NIST. 2023.
S14 SQLite. 2026-10-30.
S22 systemd project. 2026-10-30.
S18 IETF / RFC Editor. 2014.
S07 IETF / RFC Editor. 2022.
S13 SQLite. 2026-10-30.
S20 Eclipse Mosquitto. 2026-10-30.
S12 IETF / RFC Editor. 2022.
S19 OASIS. 2019-03-07.
S31 IETF / RFC Editor. 2024.
S05 NIST. 2010.
S06 NIST. 2016.

GREEN FUTURE TECH

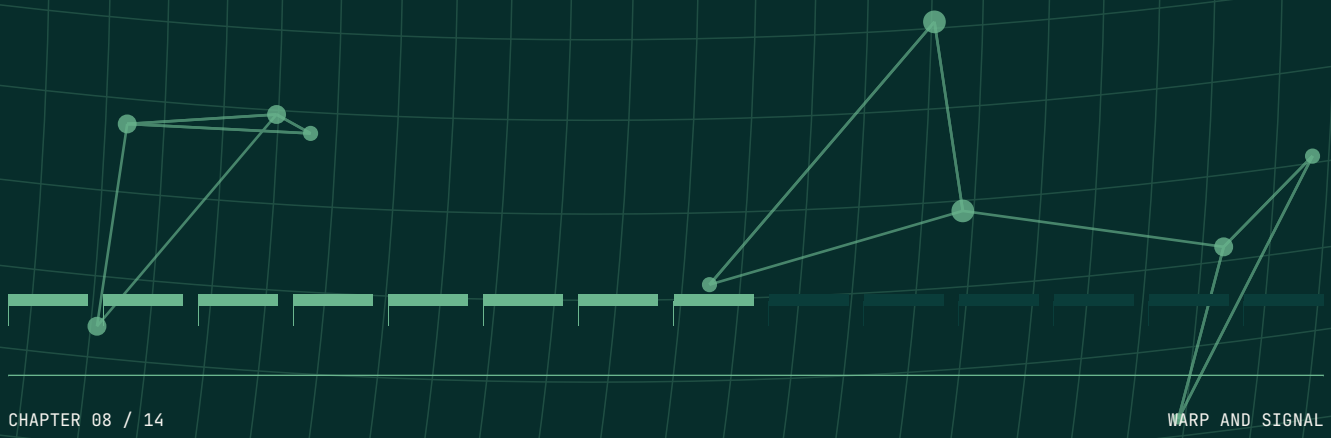
LIMITED / HONG KONG SAR

08

PART III: RECOVER A TRUSTWORTHY OPERATION / REBUILD EVENT ORDER

# A restarted clock must not rewrite production history

*A trustworthy record separates identifiers, source sequence and clock quality so reconciliation can explain what happened during disconnection.*



CHAPTER 08
DECISION TEST

A trustworthy record separates identifiers, source sequence and clock quality so reconciliation can explain what happened during disconnection.

A factory operating through power interruptions and backhaul dropouts cannot rely on wall-clock timestamps alone to establish what happened on the shop floor, in what order it happened, or whether two records describe the same bundle. [S10] [S31] [S32] When an edge gateway, barcode station or Programmable Logic Controller (PLC) bridge loses mains power and exhausts its local reserve, its clock can drift during holdover, reset to a default epoch on cold boot, or step forward or backward when network time returns. [S10] [S11] [S30] Green Future Tech Limited (GFT) argues that a trustworthy production record must decouple event identity, local causal sequence and wall-clock time across three timestamps: source event time, site ingest time and central ingest time, each stored with clock-quality metadata and a kernel or process boot identity. [S31] [S32] When connectivity returns, reconciliation rules should preserve the source observation, append attributed correction events rather than overwriting history, and mark unresolved conflicts visibly on the Performance Management System (PMS) dashboard so missing telemetry is never mistaken for zero output or fault-free operation. [S03] [S05] [S13]

In Cameroon's cotton-to-garment value chain, where roughly 95 percent of cotton fibre is exported raw and about 5 percent is processed locally according to the inherited sector register, expanding domestic transformation toward the 50 percent national industrialisation target by 2030 under the National Development Strategy (SND30) requires operating records that managers, quality leads and buyers can audit after every outage. [R06] [R07] Those sector figures describe the national baseline rather than a GFT factory result, and GFT's three-part Cameroon textile programme (prepared 31 January 2024, with sources refreshed in 2026) is a design specification rather than a commissioned plant. [R06] [S03] Within that boundary, durable local capture and bounded queueing keep records alive during an interruption, but those records remain usable for batch traceability, piece-rate verification and Overall Equipment Effectiveness (OEE) accounting only if time and identity survive a disconnected restart. [S05] [S13] [S31]

8.1

## Timestamp formatting and timestamp accuracy are different

A syntactically valid timestamp string proves only that a serialiser formatted a register according to a text grammar; it does not prove that the clock agreed with Coordinated Universal Time (UTC) when the physical event occurred. [S10] [S32] RFC 3339 specifies an Internet profile of ISO 8601 in which every timestamp expresses a four-digit year, month, day, hour, minute, second, optional fractional second and an explicit offset or `Z` indicator relative to UTC, eliminating local time-zone ambiguity and permitting lexicographical sorting when timestamps share a common UTC offset and precision. [S32] Yet an edge controller whose uncompensated real-time clock wakes from a power cut at `1970-01-01T00:00:04.120Z` or `2024-01-31T00:00:00.000Z` emits valid RFC 3339 strings that place a current sewing completion years in the past, corrupting shift attribution, lot genealogy and dwell-time calculations once delayed events reach an upstream database. [S13] [S30] [S32]

To preserve what the capture device knew at the moment of work, GFT specifies a three-timestamp envelope for every event: `source_event_time` (when the sensor trigger, scan or operator confirmation was accepted locally), `site_ingest_time` (when the site broker or SQLite store committed the event on the Local Area Network, LAN) and `central_ingest_time` (when the upstream service committed the replayed record after backhaul recovery). [S03] [S12] [S13] [S20] [S32] None of these fields may overwrite another during replay; instead, each carries a `clock_quality` attribute recording the synchronisation state of the clock that struck it. [S10] [S11] [S32]

In the Network Time Protocol version 4 (NTPv4) specification, RFC 5905 defines roles and metrics that clarify what a clock knows. [S10] A primary server at stratum 1 synchronises directly to a reference clock traceable to UTC, such as a Global Navigation Satellite System (GNSS) receiver, secondary servers at strata 2 to 15 distribute time down a subnet, and stratum 16 designates an unsynchronised clock. [S10] At every update, an NTPv4 association tracks offset (`theta`), round-trip delay (`delta`), dispersion (`epsilon`, growing at the worst-case frequency tolerance `PHI` of 15 parts per million in RFC 5905) and jitter (`psi`), combining them into root delay (`DELTA`), root dispersion (`EPSILON`) and synchronisation distance (`LAMBDA = EPSILON + DELTA / 2`), while the Leap Indicator (`LI`) signals value 3 when unsynchronised and a zero timestamp marks unknown time. [S10]

### 8.2 A THREE-TIMESTAMP CONTRACT PRESERVES OBSERVATION PROVENANCE

When an edge station records a bundle transfer on a weak grid, its event schema should capture enough timing metadata to reconstruct both the nominal instant and its uncertainty band. [S10] [S32] GFT's proposed event envelope attaches four timing fields to the source observation:

1. `source_event_time_utc`, formatted to RFC 3339 with explicit UTC indicator `Z` (or `-00:00` where UTC is used without asserting a local zone offset) and the shift calendar identifier. [S32]
2. `monotonic_ns`, read from the kernel's monotonic counter since boot, advancing independently of wall-clock steps within that boot and interpreted only within its `boot_id`. [S30] [S31]
3. `clock_state`, distinguishing `SYNCHRONISED_LOCKED` (disciplined within declared root dispersion), `HOLDOVER_RTC` (free-running on local crystal or real-time clock), `RESTORED_MONOTONIC_FLOOR` (advanced on boot from the last disk timestamp without live NTP lock) and `UNSYNCHRONISED_EPOCH` (uncalibrated after battery loss or cold reset). [S10] [S30]
4. `estimated_max_error_ms`, recording current synchronisation distance or holdover bound at capture, or `null` in `UNSYNCHRONISED_EPOCH`. [S10]

Retaining `source_event_time_utc`, `site_ingest_time_utc` and `central_ingest_time_utc` side by side separates process duration between stations, LAN queueing delay and WAN backlog duration during backhaul recovery. [S08] [S13] [S32] If a station's real-time clock battery failed over a weekend and recorded 40 illustrative scans at an epoch date before the site NTP server booted, `site_ingest_time_utc` brackets the capture window while `UNSYNCHRONISED_EPOCH` prevents scripts from computing a multi-year negative lead time. [S10] [S30] [S32]

**GFT NOTE**

Never configure an ingestion pipeline to replace a missing or implausible `source_event_time` with `central_ingest_time` upon cloud arrival. After an eight-hour illustrative backhaul outage, thousands of buffered events may arrive within a three-minute replay burst. Overwriting source time with central arrival time collapses a full shift into a three-minute spike, falsifying line balance, piece-rate pacing and electrical demand correlation. [S05] [S13] [S32]

8.3

## Identity must survive a lack of online allocation

When external connectivity fails or a switch loses power between bays, capture stations cannot request sequential keys from a central coordinator. [S03] [S31] Relying on an online auto-increment sequence or ticket server forces the line either to stop recording work during a partition or to assign temporary numbers that collide upon reconnection. [S13] [S31] RFC 9562 defines 128-bit Universally Unique Identifiers (UUIDs) so distributed nodes can mint persistent identifiers autonomously, while noting that uniqueness and sortability depend on the UUID version, entropy quality and clock-rollback handling. [S31]

Comparing the UUID versions in RFC 9562 shows why a factory on an intermittent grid cannot treat a time-ordered UUID as a standalone substitute for sequence and identity fields. [S31] UUID version 4 (`UUIDv4`) populates 122 bits (outside the 4-bit version `0b0100` and 2-bit variant `0b10`) from a Cryptographically Secure Pseudorandom Number Generator (CSPRNG), avoiding wall-clock dependency at the cost of poor B-tree index locality. [S31] Conversely, UUID version 7 (`UUIDv7`) places a 48-bit big-endian Unix Epoch millisecond timestamp (`unix_ts_ms`) in bits 0 through 47, followed by the 4-bit version (`0b0111`), 12 bits in `rand_a`, the 2-bit variant (`0b10`) and 62 bits in `rand_b`, sorting monotonically and clustering cleanly in SQLite indexes under synchronised operation. [S13] [S31]

### 8.4 CLOCK ROLLBACKS AND EARLY-BOOT ENTROPY REQUIRE COMPOSITE KEYS

On an unreliable grid, embedding a wall-clock timestamp inside `UUIDv7` exposes two hazards. [S31] First, if an edge node reboots without a valid real-time clock or experiences a backward clock step when NTP disciplines a fast oscillator, the 48-bit `unix_ts_ms` prefix moves backward. [S10] [S31] RFC 9562 Sections 6.1 and 6.2 warn that when the clock moves backward, time-based generators break monotonicity or risk collisions unless the implementation maintains generator state (such as the last timestamp and a monotonic counter in `rand_a` or `rand_b`) and applies monotonic error checking. [S31] Second, RFC 9562 Sections 6.9 and 8 warn that uniqueness depends on a seeded CSPRNG; if an embedded node boots rapidly after power loss before gathering sufficient entropy, or if a virtualised image is cloned across stations without reseeding state, random bits can repeat. [S31]

GFT therefore separates uniqueness from causal ordering by requiring a composite identity alongside `event_id`. [S13] [S31] Even when `event_id` is minted as a 128-bit `UUIDv4` or `UUIDv7`, every locally committed record must carry `(source_id, boot_id, source_seq)`: [S31]

- `source_id` is an immutable station identifier provisioned at commissioning (such as `STN-GAR-CUT-04` for an illustrative cutting terminal) and bound to local credentials, never derived from a replaceable Media Access Control (MAC) address, which RFC 9562 warns is neither guaranteed unique across virtualised environments nor desirable to expose. [S04] [S31]
- `boot_id` is a 128-bit identifier generated for the running kernel or process invocation and scoped by `source_id`; a persistent `boot_epoch` counter in SQLite distinguishes process restarts without a kernel reboot. This is GFT's proposed identifier contract, using the decentralised identifier properties in RFC 9562 rather than assuming a particular operating-system function. [S13] [S21] [S31]
- `source_seq` is a strictly increasing 64-bit unsigned integer assigned inside the same SQLite atomic commit transaction that writes the event and its outbox row. [S13]

Combining `source_id`, `boot_id` and `source_seq` prevents post-crash counter collisions. [S13] [S31] If a station with an unrecovered journal resets its counter to 1 after an abrupt power cut while its clock resets to a default date, omitting `boot_id` causes new events `(STN-01, seq=1..50)` to collide with pre-outage events `(STN-01, seq=1..50)`, leading an idempotent consumer to drop new work or overwrite prior records. [S12] [S15] [S16] [S19] Including `boot_id` marks the restart boundary, accepts the new sequence without collision and flags any unclosed transaction from the prior boot for supervisor review. [S13] [S31]

8.5

## Clock holdover must be measured on the chosen device

When a factory's external link drops or a power cut de-energises the site router while protected stations continue capturing bundle completions on battery reserve, every surviving digital clock enters holdover. [S03] [S10] During holdover, a system clock advances by counting oscillations of its local quartz crystal or real-time clock (RTC) chip, accumulating offset across manufacturing tolerances, voltage changes, ageing and enclosure temperature swings. [S10] [S30] In RFC 5905 Section 4, time offset `T(t)` at running time `t` relative to calibration at `t0` follows `T(t) = T(t0) + R(t0) * (t - t0) + 0.5 * D(t0) * (t - t0)^2 + e`, where `T(t0)` is initial offset, `R(t0)` is fractional frequency offset (`1 ppm = 10^-6 s/s`), `D(t0)` is crystal ageing rate and `e` is stochastic noise. [S10] While the chrony documentation notes that a typical computer clock drifts by less than 100 ppm uncompensated, and RFC 5905 sets default frequency tolerance `PHI` to `15e-6 s/s` (15 ppm), actual shop-floor holdover depends on prior crystal calibration and thermal rise when ventilation stops. [S10] [S30]

### 8.6 RTC COMPENSATION AND DRIFT FILES REDUCE COLD-START AND HOLDOVER ERROR

A resilient Linux edge configuration manages both the volatile kernel system clock in RAM and the hardware RTC (`/dev/rtc`) powered by a coin cell or supercapacitor during power loss. [S30] As documented in the chrony FAQ, three directives govern clock behaviour across disconnections and power cycles: `driftfile` (recording the measured system clock frequency offset to disk so `chronyd` applies the learned correction upon restart), `rtcsync` (copying synchronised time to the hardware RTC periodically while NTP lock is healthy) and `rtcfile` combined with `-s` (tracking the hardware RTC's own rate error and compensating the RTC reading when cold-booting offline). [S30]

If an embedded controller lacks a battery-backed RTC or its coin cell fails in heat, `chronyd -s` can use the drift file's modification time to advance the boot clock towards the last recorded pre-shutdown timestamp, though the clock still lags behind UTC by the unpowered outage duration until a server responds. [S30] GFT specifies full NTPv4 supervision via `chronyd` on site gateways and edge nodes, anchored by a local Stratum 1 GNSS reference or a site server with the `local` directive so LAN nodes share a common timescale when WAN backhaul fails. [S10] [S11] [S30]

**METHOD**

To qualify whether an edge computer or PLC gateway can maintain timestamps within a service class's accuracy target during an outage, GFT specifies a three-step method based on RFC 5905 and chrony parameters: [S10] [S30]

1. **Declare the target timestamp error budget (`delta_T_max`)**: Select the maximum tolerable time offset between stations correlating events without shared sequence numbers. [S03] [S10]
2. **Compute linear free-running drift (`C04`)**: For an illustrative constant residual frequency offset `R_ppm` (in ppm) over offline duration `t_off` (in seconds), calculate accumulated drift as `error_s = R_ppm * t_off / 1,000,000`. [S10]
3. **Qualify thermal and reboot step error on target hardware**: Measure actual offset against a local GNSS pulse-per-second (PPS) reference across a disconnected thermal cycle and cold power cycle under `rtcfile -s` compensation. [S30]

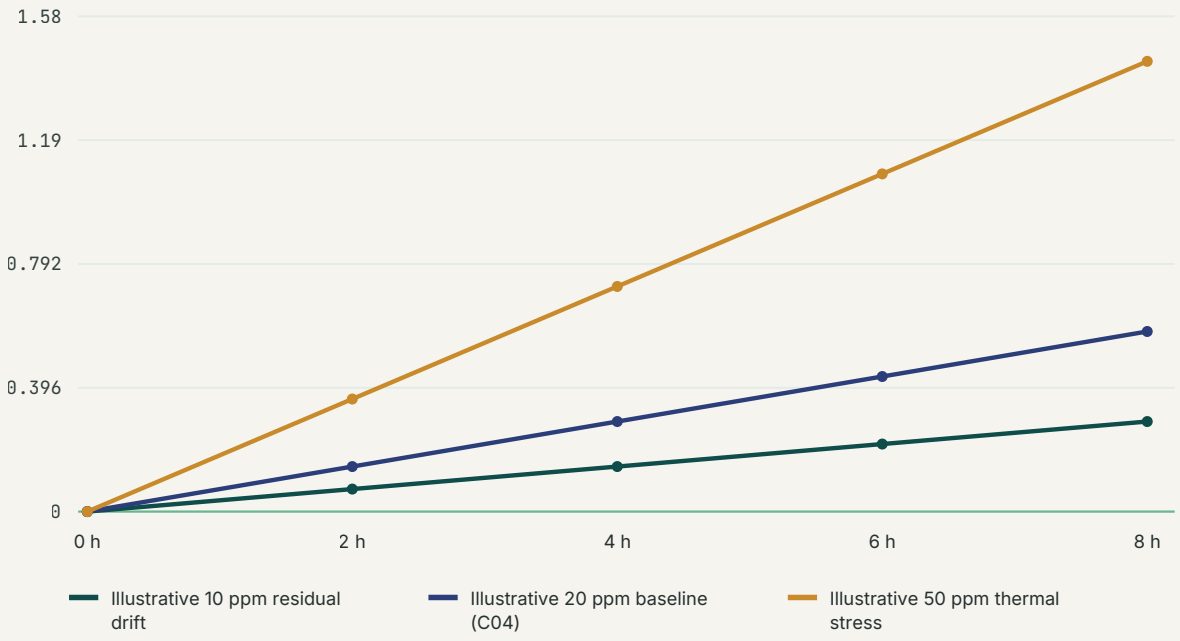
WORKED EXAMPLE

**Calculation C04: Illustrative free-running oscillator drift over an 8-hour outage.** Suppose an engineering team evaluates an edge gateway during an illustrative 8-hour (`t_off = 8 h * 3,600 s/h = 28,800 s`) interruption with an illustrative residual oscillator frequency error of `R_ppm = 20 ppm` under changing cabinet temperature. [S10] [S30] Applying Calculation C04 (`error_s = 20 ppm * 8 h * 3,600 s/h / 1,000,000 = 20 * 28,800 / 1,000,000`), the clock drifts by `0.576 seconds` (`576 milliseconds`) over the 8-hour illustrative holdover window. [S10] Two stations drifting in opposite directions by `20 ppm` diverge by `1.152 seconds` after 8 hours; at the `100 ppm` uncompensated bound in the chrony FAQ, single-node illustrative drift reaches `2.880 seconds`. [S30] A half-second to three-second error is acceptable for 8-hour shift attribution, marginal for manual barcode ordering, and unusable for sub-second camera trigger correlation without local sequence numbers. [S03] [S10] All inputs in this worked case are illustrative scenario figures rather than measured performance of a GFT factory device.

The relationship between holdover duration and accumulated timestamp error across three illustrative oscillator residual drift scenarios (10 ppm, the 20 ppm baseline from Calculation C04, and 50 ppm under thermal stress) is shown in **Exhibit 31**, p. 130. [S10] [S30]

EXHIBIT 31 / SECONDS OF ACCUMULATED CLOCK ERROR

## Clock holdover consumes the timestamp error budget linearly with offline duration




DERIVED [S10] [S30] Derived from illustrative oscillator input only (Calculation C04): error_s = ppm * t_seconds / 1,000,000. At 20 ppm over 8 hours (28,800 s), drift equals 0.576 s. The 10 ppm and 50 ppm curves illustrate sensitivity to oscillator stability and thermal drift; no measured factory device performance is implied.

As **Exhibit 31**, p. 130 demonstrates, even a disciplined crystal gradually consumes the error budget during prolonged disconnection, and the linear equation cannot capture abrupt step errors caused by a power cycle to a dead RTC battery or thermal

shock when cabinet fans stop. [S10] [S30] Site commissioning must therefore measure holdover drift on the selected hardware across realistic cabinet temperatures rather than copying a datasheet rating. [S05] [S30]

8.7

## Time recovery can invalidate an unsafe ordering assumption

When external or local NTP service returns after an outage, the clock discipline daemon reconciles the drifted or reset system clock with the reference timescale via a gradual frequency **slew** (`adjtime()`) or an abrupt **step** (`settimeofday()`). [S10] [S30] In `chronyd`, `makestep` (such as `makestep 1 3`) permits stepping the clock if the offset exceeds 1 second during the first 3 updates, whereas smaller offsets are slewed gradually. [S30]

Both modes affect event logs directly. [S10] [S30] If a station's clock ran fast during holdover and `chronyd` steps the clock **backward** by 4 illustrative seconds when the site server becomes reachable, events captured right after the step carry `source_event_time_utc` values numerically *earlier* than events captured 2 seconds before the step. [S10] [S32] Conversely, if a station booted from a stale disk floor (`RESTORED_MONOTONIC_FLOOR`) 25 illustrative minutes behind UTC and `chronyd` steps the clock **forward** by 25 minutes upon recovery, the timestamp trace shows a 25-minute void that looks identical to an unplanned stoppage if an OEE script subtracts consecutive timestamps. [S30]

**Exhibit 32**, p. 132 illustrates how a local event trace preserves exact causal order across an abrupt clock correction when each event carries `boot_id`, `source_seq`, `monotonic_ns`, `source_event_time` and `site_ingest_time`. [S10] [S31] [S32]

EXHIBIT 32 / ILLUSTRATIVE MONOTONIC SECONDS FROM SHIFT START

## Source sequence and boot identity preserve local event order through a clock step

0

**Seq 101: Bundle scan in holdover**

Boot B1; clock_state=HOLDOVER_RTC (+12 s fast).

300

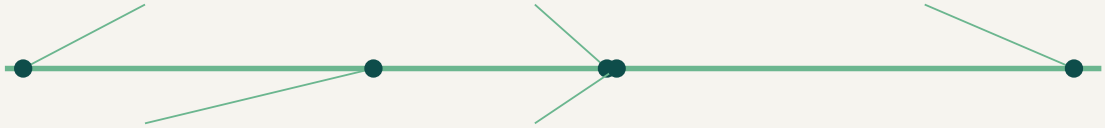
**Seq 103: NTP recovery steps clock -12 s**

Boot B1; wall clock jumps backward 12 s; audit event written.

540

**Seq 105: Reconciled outbox replay**

Boot B1; site_ingest_time_utc recorded; order preserved by Seq.



180

**Seq 102: Seam operation logged**

Boot B1; monotonic time advances +180 s.

305

**Seq 104: Quality check recorded**

Boot B1; wall time is 7 s earlier than Seq 102, but Seq 104 > 102.

ILLUSTRATIVE [S10] [S31] [S32] Illustrative event trace on a single sewing station across a 600-second window. At monotonic second 300, NTP recovery steps a fast holdover clock backward by 12 seconds; source_seq (101 to 105) and monotonic_ns preserve true local order while source_event_time_utc steps backward.

As shown in **Exhibit 32**, p. 132, if a query orders events solely by `ORDER BY source_event_time_utc`, Event `Seq 104` (recorded 5 illustrative seconds after the 12-second backward clock step) sorts *before* Event `Seq 102` (recorded 125 seconds earlier), making it appear that inspection preceded sewing. [S10] [S32] Ordering intra-station events by `(source_id, boot_id, source_seq)` and using `monotonic_ns` for elapsed cycle time within `boot_id` ensures the jump at `Seq 103` neither reorders history nor distorts the 125-second interval. [S31] [S32]

### 8.8 SOURCE SEQUENCE ORDERS ONE STATION, NOT THE ENTIRE FACTORY

While `(source_id, boot_id, source_seq)` establishes a total order for events emitted by one capture process within a single boot session, local sequence numbers cannot order events across independent devices. [S10] [S31] If Cutting Table A emits `(CUT-A, B1, seq=412)` and Sewing Line B emits `(SEW-B, B9, seq=88)` while the LAN switch between them is unpowered, neither counter constrains relative timing between the stations. [S03] [S31] Cross-station ordering during disconnected operation relies on three reconciliation rules: [S13] [S31]

1. **Physical token handoff dependencies**: When bundle traveller `BUNDLE-2026-0841` must scan out of Cutting Table A (`CUT_COMPLETE`) before scanning into Sewing Line B (`SEW_START`), the state machine defines a causal edge (`CUT_COMPLETE` → `SEW_START`) keyed by the bundle identifier, ordering the handoff by state transition and flagging any negative apparent transit time as a clock-offset anomaly. [S13] [S31] [S32]

2. **Clock-step audit events**: Whenever `chronyd` steps the clock or transitions from `HOLDOVER_RTC` to `SYNCHRONISED_LOCKED`, the agent writes a `CLOCK_STEP_RECORDED` event with the pre-step timestamp, post-step timestamp, `monotonic_ns` reading and current `source_seq`, enabling reconciliation to map pre-step timestamps onto the post-step UTC timeline without mutating `source_event_time_utc`. [S10] [S13] [S22] [S30] [S32]
3. **Explicit tie-breaking and uncertainty intervals**: When two events on different stations share no physical token and their `estimated_max_error_ms` bands overlap, reconciliation classifies their relative order as `CONCURRENT_WITHIN_CLOCK_UNCERTAINTY` rather than inventing millisecond precedence. [S10] [S32]

8.9

## Authenticated time introduces a bootstrap dependency

Protecting a factory clock against network spoofing is both a reliability requirement and a cybersecurity control. [S03] [S11] [S0802] RFC 7384 and RFC 8633 show how unauthenticated NTP packets can be forged, replayed or delayed to shift clocks, and RFC 8633 Section 5.2 warns of daemon-restart attacks in which an offset above the NTP panic threshold (`PANICT = 1000 s` in RFC 5905) crashes a daemon that `systemd` restarts with cold-start panic checks disabled. [S10] [S11] [S21] [S0802] RFC 8915 specifies Network Time Security (NTS), pairing an NTS Key Establishment (`NTS-KE`) handshake over Transport Layer Security (TLS) on TCP port 4460 with authenticated encryption (`AEAD_AES_SIV_CMAC_256`) using cookies in NTPv4 extension fields. [S23] [S27]

However, deploying NTS or TLS-dependent services on a weak grid introduces a circular dependency at cold boot. [S23] [S27] [S30] During `NTS-KE` over TLS 1.3 (RFC 9846, obsoleting RFC 8446 as of July 2026), the client verifies the server's X.509 certificate under RFC 5280 (`notBefore` and `notAfter`). [S23] [S27] [S0801] As RFC 8915 Section 8.5 and the chrony FAQ Section 3.5 explain, if an edge device boots with an exhausted RTC and its clock defaults to `1970-01-01` or a stale date outside `[notBefore, notAfter]`, strict TLS verification fails, blocking the handshake needed to set the clock. [S23] [S27] [S30] [S0801]

> **WARNING**
>
> Never resolve an NTS or TLS bootstrap failure by permanently enabling `nocerttimecheck` or stripping TLS authentication. [S03] [S27] [S30] RFC 8915 Section 8.7 warns against "NTS stripping", where an implementation falls back to plain NTP when `NTS-KE` fails, because an on-path adversary can block TCP port 4460 to force the downgrade. [S27] Disabling expiration checks globally lets an attacker holding an expired key impersonate services indefinitely. [S27] [S0801]

### 8.10 FOUR ENGINEERING CONTROLS BREAK THE TIME-AND-TLS DEADLOCK SAFELY

Both RFC 8915 Section 8.5 and the chrony documentation provide mechanisms that prevent an offline restart from triggering an NTS/TLS deadlock: [S27] [S30]

1. **Persist NTS cookies and AEAD keys across restarts**: RFC 8915 Section 5.7 specifies retaining unused cookies, `C2S` and `S2C` keys and AEAD parameters in persistent storage; in `chronyd`, `ntsdumpdir /var/lib/chrony` writes NTS cookies to disk so a rebooted client resumes authenticated UDP port 123 polling without a fresh TLS handshake on TCP port 4460. [S27] [S30]
2. **Persist a monotonic minimum time floor on disk**: RFC 8915 Section 8.5 recommends writing synchronised time to disk and rejecting certificates whose `notAfter` precedes that floor, while `chronyd -s` initialises the boot clock from the drift file's modification time even without an RTC battery, placing the client inside `[notBefore, notAfter]` when the outage is shorter than remaining certificate validity. [S27] [S30] [S0801]
3. **Maintain a protected local Stratum 1 or site boundary time server**: Following RFC 8633 Sections 3.2 and 3.3 and chrony LAN guidelines, edge stations should poll at least three or four sources anchored by a site GNSS server, using `ntstrustedcerts` and `nosystemcert` with annual rotation for internal NTS. [S10] [S11] [S30]
4. **Require an auditable operator runbook for cold-start intervention**: Following RFC 8633 Section 5.2 and NIST SP 800-82 Rev. 3 controls, `chronyd` should restrict `makestep` to boot updates and set `maxchange` to guard runtime steps, while manual bootstrap after battery and cookie expiry requires an authenticated maintenance login logged in persistent `journald` storage. [S03] [S11] [S22] [S30]

**Exhibit 33**, p. 135 summarises the protocol mechanisms alongside GFT site qualification requirements. [S10] [S11] [S23] [S27] [S30] [S0801]

EXHIBIT 33 / SPECIFICATION MECHANISM AND GFT CONTROL

## Authenticated time and certificate checks depend on persisted clock state

| Failure or mechanism | Documented protocol or daemon behaviour | GFT site requirement (proposed control) |
|---|---|---|
| X.509 validity window (S0801, S23) | TLS 1.3 checks current clock against certificate notBefore and notAfter dates | Verify cold-boot clock reaches validity window before NTS-KE or HTTPS replay |
| NTS cookie persistence (S27, S30) | RFC 8915 persistent cookies and chrony ntsdumpdir allow NTP resume without TLS NTS-KE | Configure ntsdumpdir on durable storage and test reboot with TCP 4460 blocked |
| RTC drift and disk floor (S30) | chrony rtcfile -s compensates RTC drift; disk timestamp advances clock on battery loss | Measure RTC holdover across cabinet heat and verify disk-floor advancement |
| Panic threshold and restart (S10, S11, S30) | RFC 5905 PANICT is 1000 s; RFC 8633 warns auto-restart plus ignored panic enables attack | Bound makestep to boot updates, set maxchange and alert on daemon exit |
| NTS stripping defense (S27) | RFC 8915 forbids silent fallback from NTS to unauthenticated NTP on handshake failure | Hold clock in HOLDOVER_RTC with dashboard flag rather than downgrade to plain NTP |

SOURCED [S10] [S11] [S23] [S27] [S30] [S0801] Source-based comparison of time-synchronisation and certificate-validation behaviours from RFC 5905 (S10), RFC 8633 (S11), RFC 9846 (S23), RFC 8915 (S27), chrony FAQ (S30) and RFC 5280 (S0801), paired with GFT commissioning controls.

As **Exhibit 33**, p. 135 makes explicit, the difference between a resilient factory node and one that locks itself out after a weekend power cut lies in persistent local state (`ntsdumpdir`, `driftfile`, `rtcfile` and persistent `journald` logs) verified under fault testing before sign-off. [S22] [S27] [S30]

8.11

# Manual and automated records need a conflict policy

Even with battery-backed gateways, prolonged power cuts or sensor failures periodically push stations into manual fallback. [S03] [S05] During fallback, supervisors or clerks record completions, piece counts and defect dispositions on paper travellers or offline spreadsheets carrying the bundle identifier, operation code, shift, role and paper sequence number. [S05] [S31] When power and LAN services return, the plant must reconcile manual entries against automated events committed before the outage, buffered in an outbox during holdover, or re-scanned after the terminal returned. [S12] [S13] [S31]

Without an explicit conflict policy, automated replay and manual import corrupt the ledger in two opposite ways: blindly summing manual counts with buffered scans double-counts bundles scanned just before a screen freeze, whereas deleting automated rows to match a paper sheet destroys the audit trail. [S12] [S13] [S19] [S31]

### 8.12 APPEND-ONLY RECONCILIATION STATES PRESERVE THE AUDIT TRAIL

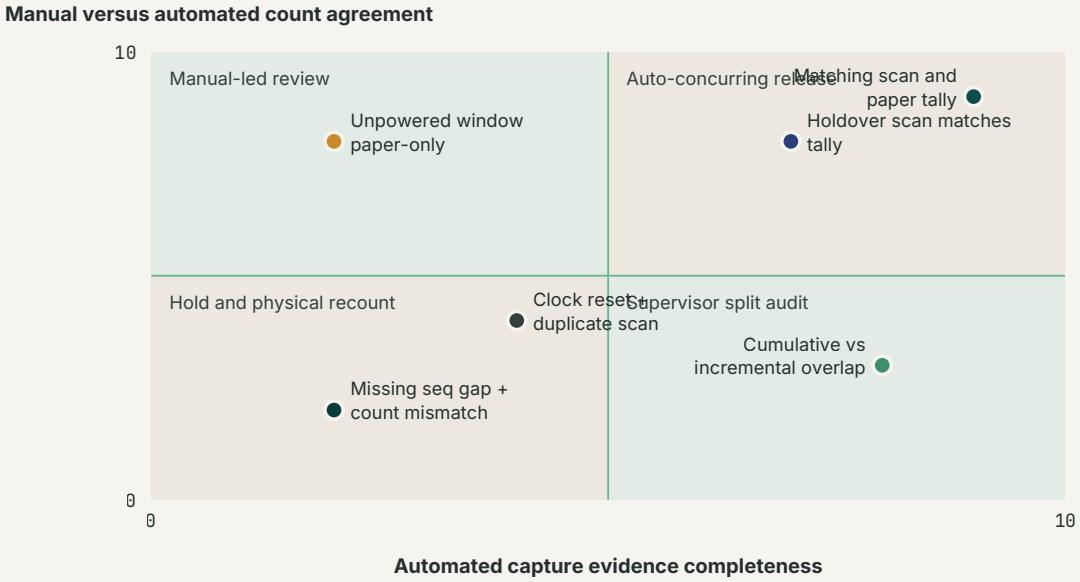
GFT resolves the boundary between manual and automated capture through an append-only ledger built on atomic transactions (S13), idempotent keys (S12, S31) and explicit authority rules. [S12] [S13] [S31] Every production event is stored as an immutable claim linked to (bundle_id, operation_id) and evaluated across four rules: [S13] [S31]

1. **Deduplication of identical business claims**: If an automated event and a manual entry reference the same (bundle_id, operation_id) with identical piece counts and compatible shift windows, the system links both raw records to one consolidated ledger entry, sets RECONCILED_CONCURRING and credits production once. [S12] [S13] [S31]
2. **Resolution of split-outage partial counts**: Suppose a station committed 35 illustrative pieces of a 50-piece bundle before UPS shutdown (SERVICE_SHUTDOWN_COMMITTED) and the operator completed 15 illustrative pieces on paper. If the paper form records the cumulative total (50 pieces) rather than the increment (15 pieces), summing both yields 85 pieces; the engine flags QUANTITY_OVERFLOW_CONFLICT whenever auto_count + manual_count > bundle_standard_qty and sets PENDING_SUPERVISOR_RECONCILIATION until the supervisor confirms whether the paper entry was cumulative or incremental. [S05] [S13]
3. **Source authority hierarchy by measurement type**: For presence and timestamp of a barcode scan in SYNCHRONISED_LOCKED state, the committed automated event holds primary timing authority over a transcribed paper time, whereas for physical piece count and defect grading completed while terminals were unpowered (MANUAL_FALLBACK_ACTIVE), the signed supervisor tally sheet holds primary count authority. [S03] [S05] [S10] [S13]
4. **Immutable correction events rather than destructive updates**: Resolving a conflict never executes an UPDATE or DELETE on raw event rows; instead, it commits an atomic RECONCILIATION_DECISION event carrying its own event_id, reviewer_role, decision_timestamp_utc, superseded_event_ids, authoritative_qty and reason_code (delegating retention and access governance to wp14 and lot genealogy to wp05). [S03] [S13] [S31] [S32]

**Exhibit 34**, p. 137 maps representative post-outage reconciliation scenarios across automated capture evidence completeness and count agreement between manual and automated records. [S13] [S31]

EXHIBIT 34 / ILLUSTRATIVE EVIDENCE AND CONSISTENCY SCORE (0 TO 10)

## Conflicts between manual and automated records require an explicit reconciliation decision




ILLUSTRATIVE [S13] [S31] Illustrative classification of six post-outage bundle reconciliation cases on a 0-to-10 scale of automated evidence completeness (x-axis) versus count agreement between manual and automated logs (y-axis). Design basis: atomic ledger integrity (S13) and independent event identification (S31).

Looking at **Exhibit 34**, p. 137, only cases in `Auto-concurring release` close automatically upon replay. [S13] [S31] Cases in `Manual-led review` post under manual authority with a `SOURCE=MANUAL_FALLBACK` tag, whereas `Supervisor split audit` and `Hold and physical recount` hold release until supervisor confirmation or a recount is committed. [S03] [S05] [S13]

8.13

## The dashboard should show unresolved uncertainty

The final link in recovering a trustworthy operation is honest visual representation on shop-floor and management dashboards. [S03] [S05] Most dashboards assume a closed world: if the database holds zero completions for Sewing Line 2 between `10:00` and `11:00`, the chart plots `0 units/hour` as a line stop; if 400 illustrative buffered events arrive at `11:04` without source-time separation, it plots a `400 units/hour` surge. [S05] [S32] On a weak grid, a silent telemetry gap during a LAN or power outage is an epistemic state (`UNKNOWN_PRODUCTION_PENDING_REPLAY`), not a measured physical zero (`MEASURED_ZERO_OUTPUT`). [S05]

Similarly, when a station captures events in `HOLDOVER_RTC` or `UNSYNCHRONISED_EPOCH`, or when an illustrative sequence gap (`source_seq` jumping from 410 to 425 on the same `boot_id`) indicates 14 unrecovered records, displaying an unannotated KPI tile misleads supervisors into acting on incomplete data. [S10] [S16] [S31] In GFT's

Cameroon textile programme, where the PMS guides shift huddles, coaching and maintenance dispatch, the interface must surface data trust alongside production volume. [S03] [S05]

### 8.14 FOUR VISUAL QUALITY STATES PREVENT FALSE OPERATIONAL CONCLUSIONS

GFT proposes that every chart, shift summary and OEE indicator in the PMS render four data-state overlays driven by the reconciliation ledger (delegating OEE formulations under partial observability to `wp15` and daily management routines to `wp02`): [S03] [S05] [S11]

- **Stale backhaul / pending replay (`STALE_TELEMETRY_HOLD`)**: When a station or gateway has not delivered a cursor update within its freshness threshold (`t_last`), the segment after its last confirmed commit is shaded with a hatched band rather than plotting zero output (`0 units`). [S05] [S08]
- **Degraded clock certainty (`CLOCK_UNSYNCHRONISED` or `CLOCK_HOLDOVER`)**: Events committed outside `SYNCHRONISED_LOCKED` (such as an 8-hour holdover at 20 ppm with a `0.576 s` C04 bound, or a dead RTC reboot) display a clock-uncertainty badge or `site_ingest_time_utc` bracket, allowing shift tallying while blocking sub-second cross-station correlation. [S10] [S30] [S32]
- **Sequence discontinuity alert (`MISSING_SEQUENCE_GAP`)**: If the ledger for `(source_id, boot_id)` contains sequence numbers `N` and `N + k` (`k > 1`) without intermediate records after replay catches up, the dashboard raises a banner naming the station, boot session and missing slot count (`k - 1`) so supervisors check bundle travellers before closing the shift. [S13] [S16] [S31]
- **Provisional count pending reconciliation (`UNRECONCILED_BUNDLE_CONFLICT`)**: Whenever a shift total includes bundles in `PENDING_SUPERVISOR_RECONCILIATION`, the KPI card displays both the confirmed baseline and the disputed range (such as `35 confirmed + 15 disputed` in an illustrative overlap), locking shift sign-off until a `RECONCILIATION_DECISION` event is committed. [S05] [S12] [S13]

Enforcing these rules at the presentation layer ensures that a power cut or clock reset cannot silently rewrite history into either a phantom crash or a phantom productivity spike. [S03] [S05] [S13] Operators and managers see what was measured on a locked clock, what was ordered by local sequence during holdover, what was recovered from paper fallback, and what awaits verification. [S10] [S31] [S32]

*A clock may lose the hour when the grid fails, but a numbered thread in a committed ledger still knows which stitch came first.*

#### CHAPTER SOURCE NOTES

S10 IETF / RFC Editor. 2010.

S31 IETF / RFC Editor. 2024.

S32 IETF / RFC Editor. 2002.

S11 IETF / RFC Editor. 2019.

S30 chrony project. 2026-10-30.

S03 NIST. 2023.

S05 NIST. 2010.

S13 SQLite. 2026-10-30.

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

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

S12 IETF / RFC Editor. 2022.

S20 Eclipse Mosquitto. 2026-10-30.

S08 IETF / RFC Editor. 2007.

S04 NIST. 2018.

S21 systemd project. 2026-10-30.

S15 SQLite. 2026-10-30.

S16 SQLite. 2026-10-30.

S19 OASIS. 2019-03-07.

S22 systemd project. 2026-10-30.

S0802 IETF / RFC Editor. 2014-10.

S23 IETF / RFC Editor. 2026.

S27 IETF / RFC Editor. 2020.

S0801 IETF / RFC Editor. 2008-05.

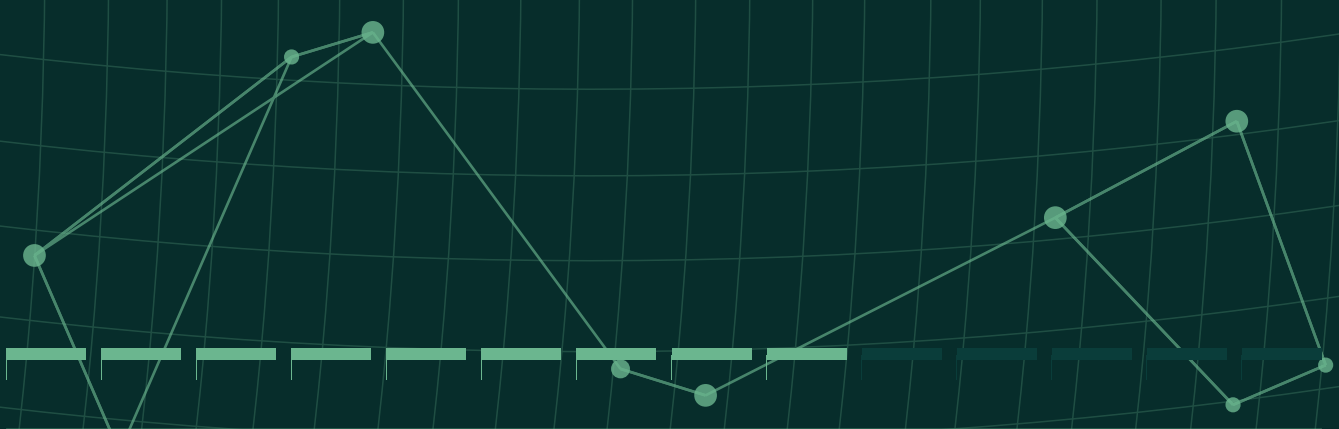
GREEN FUTURE TECH
LIMITED / HONG KONG SAR

09

PART III: RECOVER A TRUSTWORTHY OPERATION / KEEP DECISIONS USABLE

# Local inference still needs a complete operating envelope

*An edge model supports continuity only when its inputs, licence, operator interface and fallback remain usable without the cloud.*



CHAPTER 09
DECISION TEST

An edge model supports continuity only when its inputs, licence, operator interface and fallback remain usable without the cloud.

Deploying neural network weights onto a factory-floor edge processor does not guarantee that automated quality inspection will survive a power disturbance or a wide-area network blackout. An inference engine is one component inside a coupled cyber-physical loop that begins with stable LED illumination, camera capture, and rotary encoder pulses, passes through local model execution and operator display, and finishes only when a defect classification, roll-metre coordinate, and human disposition are durably committed to the local production management system (PMS) ledger and outbox [S03] [S13]. If any link in that chain relies on an unbacked power circuit, a cloud container registry, an online licence lease, or a browser interface rendered from a distant server, the inspection station fails the moment external infrastructure degrades [S03] [S05]. GFT's position is that local inference delivers operational continuity only when the entire operating envelope is engineered for disconnected autonomy: sensing peripherals must share the protected electrical boundary, the runtime must cold-boot and authenticate from immutable local flash, end-to-end latency must stay below the physical reaction window of moving cloth, operators must possess a structured local fallback when automation degrades, and energy scarcity must trigger an explicit, logged compute-shedding policy rather than a silent drop in quality authority.

In Cameroon's cotton-textile value chain, where about 95 percent of lint is exported raw and only about 5 percent is processed locally, national industrial policy targets 50 percent domestic transformation by 2030 [R06] [R07]. Moving from raw lint export into competitive woven, dyed, and finished fabric requires strict roll-level quality control so weaving and dyeing faults do not propagate into garment cutting or export shipment rejections. The internal Cameroon textile programme proposes off-the-shelf optical defect detection deployed on edge compute nodes alongside point-of-work PMS capture (`wp03` evaluates computer-vision selection and qualification; `wp14` governs model lifecycle and data governance). Those internal programme documents establish a design intent rather than an operating proof. Under grid conditions where 93.3 percent of surveyed Cameroonian firms experience electrical outages and a typical month brings 10.4 interruptions [S01], an edge vision appliance specified solely by its laboratory accuracy will fail on the shop floor unless its electrical, startup, timing, human-interface, and rollback boundaries are engineered as one system.

9.1

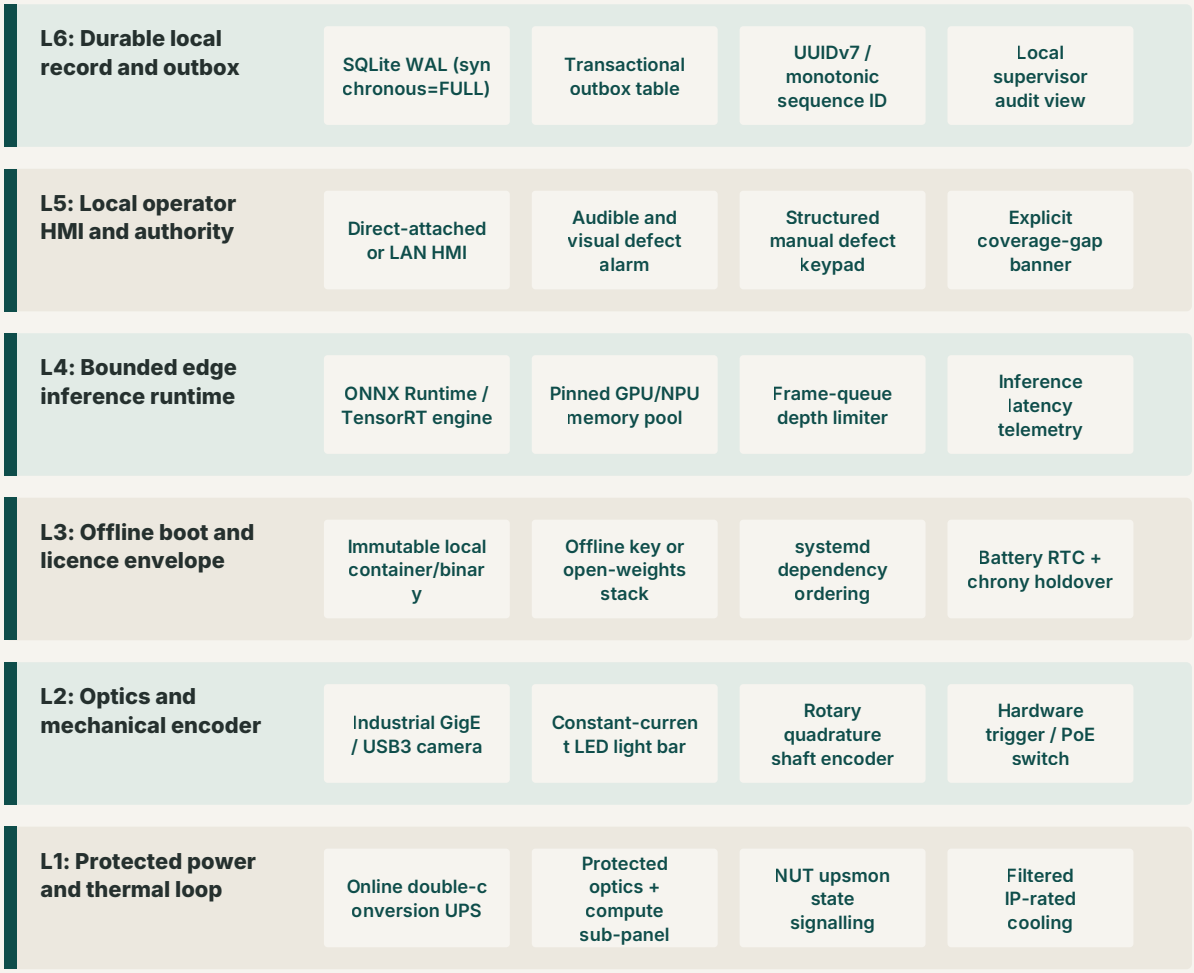
## An edge model is one dependency within an inspection function

### 9.2 THE PHYSICAL AND TRANSACTIONAL CHAIN CONNECTS OPTICS TO DISPOSITION

Vendor demonstrations of industrial computer vision routinely isolate the neural network from the factory environment. A pre-recorded video stream is fed into a graphics processing unit (GPU) or neural processing unit (NPU), bounding boxes appear on a monitor, and the software is declared edge-capable because no frames were transmitted to a remote cloud during the forward pass. On an operating weaving loom, stenter frame, or fabric inspection table, the mathematical forward pass of the model is the middle step of a six-layer dependency stack [S03]. If the layers beneath or above the inference runtime fail, the neural network either receives unusable inputs or produces classifications that vanish before they can govern product quality.

EXHIBIT 35 / QUALITATIVE ARCHITECTURE

## Inspection continuity includes its full dependency chain




ILLUSTRATIVE [S03] [S04] [S13] [S18] Illustrative six-layer physical, runtime, and transactional stack required to sustain local optical inspection during a WAN or mains interruption. Design synthesis based on NIST SP 800-82r3, SP 800-193, SQLite atomic commit, and RFC 7228.

Exhibit x09-inspection-dependency-stack maps the six dependencies that must remain intact simultaneously for an automated inspection station to perform its operational duty [S03] [S04] [S13] [S18]. At the base (Layers 1 and 2), optical inspection measures reflected or transmitted light synchronised to mechanical motion. Constant-current LED illumination arrays must maintain steady luminous flux without 50 Hz mains flicker or voltage-sag dimming (illustrative mains frequency), while a rotary quadrature encoder on the roller generates pulses proportional to cloth travel so defects are located at exact roll coordinates (such as 14.2 m or 184.6 m, illustrative coordinates). Above the physical capture layer, Layer 3 governs cold-boot readiness from local flash [S04] [S21], Layer 4 executes the forward pass within a millisecond budget dictated by roller velocity [S03], Layer 5 presents the result on the local HMI for operator action, and Layer 6 binds the automated detection and human disposition into an fsync-backed SQLite transaction [S13] [S16].

### 9.3 THE ELECTRICAL BOUNDARY MUST ALIGN WITH THE DATA BOUNDARY

A common engineering error in industrial digitalisation is protecting the compute box while leaving physical sensors and the operator terminal on raw utility power. When the edge server sits on an IT cabinet UPS while the camera PoE injector, 24 V DC encoder supply, overhead LED light bar, and touchscreen monitor plug into convenience outlets on the machinery circuit (illustrative supply voltage), a brief mains sag extinguishes the LED bar for 800 milliseconds, reboots the PoE switch for 45 seconds, and blanks the monitor (illustrative timings) [S03]. If the heavy rollers coast during that window, metres of cloth pass under a blind camera while the server logs a healthy process status.

To prevent this partial-survival trap, the protected electrical boundary defined in Chapter 4 must match the functional boundary of the inspection loop [S03] [S05]. Every active component required to capture, locate, classify, display, and commit a defect must sit on the same conditioned UPS sub-circuit whenever the roller can move. Where the inspection motors stop on mains loss, the protected digital load must remain powered long enough to process the final frames captured during coast-down, let the operator record the stopping metre mark, and commit the roll suspension state cleanly to disk [S05] [S13].

9.4

## Offline startup must include model and licence availability

### 9.5 COLD-BOOT FAILURE MODES SEVER MODEL EXECUTION DURING BACKHAUL LOSS

An extended mains failure drains local UPS batteries, halts the edge gateway, and leaves the floor unpowered until utility supply or standby generation returns. When power resumes two hours later (illustrative outage duration), telecommunications backhaul often remains dark because local cellular base stations or fibre nodes have exhausted their own batteries [S01] [S18]. The critical continuity test for an edge inspection node is therefore whether it can execute a cold boot from zero power into full inspection readiness while the wide-area network is unreachable [S04] [S05]. Commercial vision appliances fail this test through four hidden cloud dependencies:

1. **Container and weight fetching on start**: Orchestration agents with `imagePullPolicy: Always` or startup scripts checking remote object stores for `.onnx` or `.engine` weights hang or crash-loop when external DNS and HTTPS time out [S18] [S21].
2. **Volatile RAM caching**: Accelerators compiling neural network graphs into `/tmp` (`tmpfs` in RAM) that require an online package or telemetry handshake lose their cache on reboot and block startup [S04].
3. **Online cryptographic licence leases**: Proprietary vision SDKs enforcing floating licences or token refreshes against a vendor cloud server refuse to start if a lease expires during a 12-hour outage (illustrative lease window) [S03] [S23].
4. **Blocking network and time wait-states**: Service units waiting on internet-facing `network-online.target` checks or upstream Network Time Protocol (NTP) servers prevent the local application from binding its camera and HMI sockets [S21] [S30].

### 9.6 SYSTEMD ORDERING AND LOCAL SECRET CACHING ENFORCE OFFLINE READINESS

Eliminating cold-boot fragility requires strict discipline across the software bill of materials and the operating system service graph [S04] [S21]. First, GFT recommends banning runtime cloud-leased licences for shop-floor inspection functions. Procured vision software must use either perpetual node-locked cryptographic keys stored on the local filesystem (or an internal hardware security module) or an open-weights execution stack (such as ONNX Runtime or TensorRT running locally compiled models) with zero external activation logic [S03] [S04].

Second, all container images, shared libraries, camera GenICam XML descriptors, and compiled model weights must reside on a persistent, journalled local partition (`ext4` with write barriers enabled) mounted read-only during normal execution [S04] [S29]. The `systemd` unit governing the inspection daemon must declare explicit local dependencies [S21]:

- **Unit ordering (`[Unit]`)**: Declare `After=local-fs.target systemd-udevd.service chronyd.service` and `Requires=local-fs.target`, omitting `network-online.target` so startup never waits for an external gateway [S21] [S30].
- **Filesystem isolation (`[Service]`)**: Set `ReadOnlyPaths=/opt/gft/models /opt/gft/config` and `ReadWritePaths=/var/lib/gft/pms-db /var/log/gft` to prevent accidental mutation of model weights [S04] [S21].
- **Pre-start integrity gate (`ExecStartPre`)**: Run `verify-model-digest` against `/opt/gft/models/active/manifest.json` before launching `inspect-daemon --offline-ready`, and set `Type=notify` with `WatchdogSec=15s` so `systemd` marks the service ready only after the model digest check, warm-up inference, camera link check, and SQLite WAL connection succeed [S04] [S14] [S21].

#### OFFLINE COLD-BOOT QUALIFICATION PROCEDURE

Before accepting any edge inspection station for factory deployment, commissioning engineers must execute a physical disconnected cold-boot test [S04] [S05] [S21]: 1. **Isolate the WAN boundary**: Physically unplug the upstream backhaul cable and cellular modem from the site gateway so external DNS, NTP, registries, and vendor licence servers are 100 percent unreachable (illustrative isolation check) [S05]. 2. **Exhaust volatile state**: Shut down the edge node and disconnect mains and UPS power for an illustrative 5 minutes so all RAM caches and volatile registers clear completely [S04]. 3. **Execute cold start**: Restore local power on the isolated plant LAN and start a stopwatch (illustrative step 3) [S05]. 4. **Verify readiness contract**: Confirm that the OS boots, `chronyd` initialises from the battery-backed RTC with an explicit holdover flag [S30], the model digest check and warm-up inference pass, and the HMI transitions to `READY (LOCAL MODE)` within the agreed recovery time objective (for example, an illustrative 90-second target) [S05] [S21]. 5. **Advance the clock**: Advance the local RTC forward by an illustrative 30 days with the WAN disconnected, repeat the cold boot, and confirm that no certificate expiry or licence lease timeout blocks execution [S03] [S30].

9.7

# Inference delay must be bounded by the operating task

## 9.8 PHYSICAL FABRIC MOTION CONVERTS INFERENCE LATENCY INTO MISSED DEFECTS

On a moving fabric inspection table or inline loom monitor, time is directly convertible into linear distance of cloth [S03]. Every millisecond that elapses between the moment a woven defect passes beneath the camera lens and the moment the system flashes a stop alarm or actuates a selvedge marker moves the defect further downstream toward the wind-up roller. Once the defective segment winds onto the take-up beam, the operator must stop the machine, reverse the drive, unwind metres of fabric, locate the fault by eye, insert the physical tag or cut, and re-tension the roll.

Let the linear fabric speed be `v` (in metres per minute, or `v / 60` metres per second in this illustrative kinematic formulation) [S03]. With an along-roll camera field of view `L_fov` (in metres) and frame overlap fraction `alpha` (typically 0.10 to 0.25 as an illustrative range so boundary defects are captured intact), the maximum cloth advance per exposure is `Delta_x = L_fov * (1 - alpha)`, fixing the minimum frame rate `f_min = (v / 60) / (L_fov * (1 - alpha))` (illustrative relation).

To prevent an unbounded queue in memory, per-frame processing across `N_cam` cameras must stay strictly below `T_budget = 1 / f_min` [S03]. Total reaction latency `T_total` from photon exposure to operator alert and SQLite commit sums optical exposure (`T_expose`), DMA transfer (`T_transfer`), model inference (`T_infer`), SQLite WAL commit (`T_commit`) [S13] [S14], and HMI display (`T_display`). During `T_total`, the cloth advances by `d_react = (v / 60) * T_total` (illustrative reaction distance), which must remain shorter than the physical board distance `L_table` from camera axis to nip roller minus human braking reaction time `T_human` (an illustrative 0.5 to 1.0 seconds).

In keeping with GFT's evidence discipline, no generic line-speed or detection-accuracy figure is imported from external literature or the programme annex as a promised plant specification; actual speed and accuracy must be measured during the future site pilot (`wp03`) [S03]. The worked case below demonstrates the kinematic sizing arithmetic using illustrative inputs.

WORKED EXAMPLE

**Illustrative frame-rate and reaction-window arithmetic across inspection modes**

Consider an illustrative inspection station with camera field of view `L_fov = 0.20 m` (200 mm) and overlap `alpha = 0.20`, yielding an advance per frame of `Delta_x = 0.16 m` (160 mm). Distance from camera axis to wind-up nip roller is `L_table = 1.20 m`, and operator braking reaction time is `T_human = 0.80 s`. All parameters are illustrative inputs.

1. **Mode A: Slow manual-assist table (`v = 15 m/min = 0.25 m/s` illustrative)**: `f_min = 0.25 / 0.16 = 1.56 fps` per lane (`213.3 ms` budget across 3 cameras). Nip transit window is `4.80 s`, leaving `4.00 s` after `T_human`. At an illustrative `T_total = 120 ms`, cloth travels only `30 mm` [S03].
2. **Mode B: Standard power table (`v = 40 m/min = 0.67 m/s` illustrative)**: `f_min = 4.17 fps` per lane (`12.5 fps` across 3 cameras, `80.0 ms` per-frame budget). Nip transit window is `1.80 s`, leaving a ceiling of `T_total ⩽ 1.00 s`. At an illustrative p99 `T_total = 120 ms` (`15 ms` transfer, `75 ms` inference, `15 ms` SQLite commit, `15 ms` display), cloth advances `80 mm` before the alarm sounds, leaving `1.68 s` (`1.12 m` of board length) to brake and tag [S03] [S13].
3. **Mode C: High-speed stenter exit (`v = 90 m/min = 1.50 m/s` illustrative)**: `f_min = 9.38 fps` per lane (`28.1 fps` across 3 cameras, `35.6 ms` budget). Because the `0.80 s` transit across a `1.20 m` table equals `T_human = 0.80 s`, a human operator cannot stop the roll before wind-up at `90 m/min` (illustrative); control must shift to automated PLC selvedge marking and SQLite defect-map logging for cutting-table avoidance [S03] [S13].

EXHIBIT 36 / M/MIN, FPS, MS, S

## Physical fabric speed sets the hard ceiling on edge inference latency

| Operating regime (illustrative) | Speed (m/min) | Min lane fps | 3-cam budget (ms) | Nip window (s) | Feasible quality control action |
|---|---|---|---|---|---|
| Mode A: Manual-assist table | 15 | 1.56 | 213 | 4.8 | Operator foot-brake stop and manual selvedge tag |
| Mode B: Standard power table | 40 | 4.17 | 80 | 1.8 | Operator foot-brake stop if p99 T_total <= 200 ms |
| Mode C: High-speed stenter exit | 90 | 9.38 | 35.6 | 0.8 | Automated PLC selvedge marker + digital roll defect map |

DERIVED [S03] [S13] [S18] Derived frame-acquisition cadence, 3-camera compute budget, and reaction-distance limits across three illustrative fabric inspection speeds (L_fov = 0.20 m, overlap = 20%, L_table = 1.20 m). All parameters are illustrative engineering inputs.

### 9.9 TAIL LATENCY REQUIRES EXPLICIT LOAD SHEDDING AND BACKPRESSURE

Exhibit x09-latency-and-frame-budget-table shows why mean latency is an inadequate metric for operational technology [S03]. A model averaging 45 ms per frame in Mode B (80 ms illustrative budget) fails if its 99th-percentile (p99) latency spikes to 950 ms during background activity. On a shared edge gateway, three events routinely trigger tail-latency spikes if left unisolated: unthrottled reconnect replay storms encrypting thousands of buffered rows over TLS 1.3 after an eight-hour blackout (illustrative duration) [S08] [S23]; unbounded SQLite WAL checkpoints stalling writer transactions for hundreds of milliseconds (illustrative stall) [S14] [S16]; and thermal frequency throttling stepping down clock speeds by 40 to 60 percent when hall ambient exceeds 35 degrees Celsius (illustrative thermal derating) [S03].

To enforce deterministic latency bounds, the node must apply Linux cgroups v2 and systemd resource slices (CPUAffinity, CPUWeight, IOWeight, MemoryMin) so capture and inference hold priority over outbox replay workers [S03] [S21]. The frame-ingestion ring buffer must enforce a hard depth limit (for example, 3 frame triplets, an illustrative bound). If T_infer breaches T_budget and fills the buffer, the runtime must never queue frames infinitely in RAM; instead, it must drop to 1:2 frame decimation, switch the HMI banner to DEGRADED: REDUCED FRAME RATE, log the exact metre mark in SQLite, and cap roller speed at Mode A (15 m/min illustrative) or transition to MANUAL FALLBACK [S03] [S13].

9.10

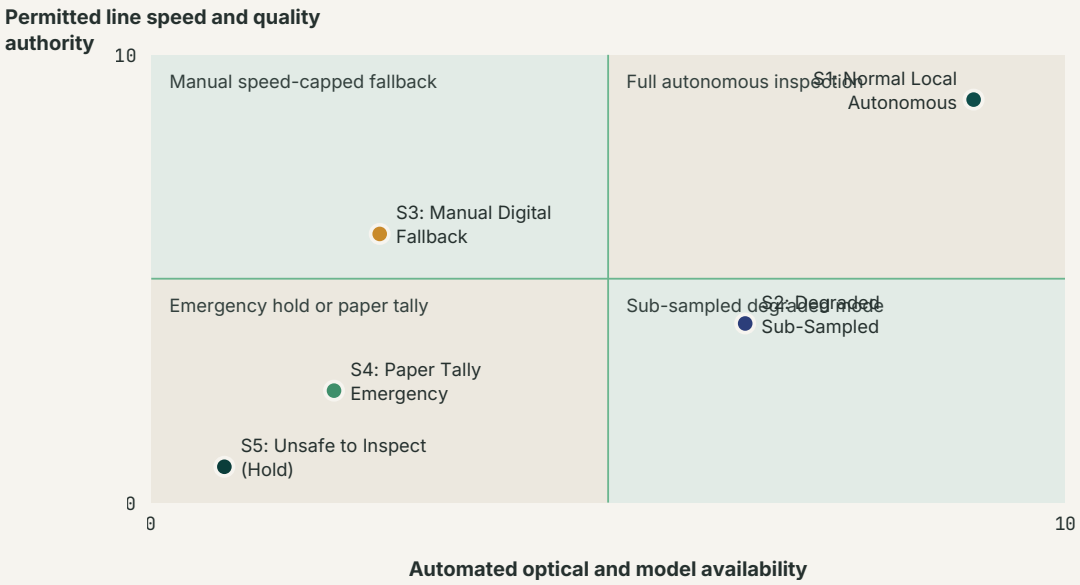
## Human fallbacks must preserve quality authority

### 9.11 DEGRADED INSPECTION STATES REQUIRE DISTINCT VISUAL SEMANTICS

When an automated inspection loop slows down, loses illumination, or drops offline, the worst design outcome is ambiguity about who holds quality authority over the moving roll. If an inspector believes the vision system is scanning the cloth when the inference worker has crashed or frozen on a stale video frame, metres of defective fabric will ship as certified first-quality goods [S03] [S05]. Continuity engineering treats human fallback as a formally engineered operating mode with explicit entry triggers, visual state signalling, and structured data capture [S03] [S05].

EXHIBIT 37 / ILLUSTRATIVE OPERATING-STATE COORDINATES (0 TO 10)

## The operator must see why a decision is unavailable




ILLUSTRATIVE [S03] [S05] [S13] Illustrative operating-state matrix placing the five station modes (S1 to S5) across automated model availability (x-axis) and permitted line-speed authority (y-axis). Coordinates are illustrative design positions.

Exhibit `x09-operator-decision-state-matrix` places the five mutually exclusive station operating states (`S1` through `S5`) across model availability and permitted speed authority [S03] [S05] [S13]:

- `S1: Normal Local Autonomous`: Optics, encoder, model, and local SQLite healthy; steady green banner (`AUTO INSPECTION ACTIVE`); records `inspection_mode='AUTO_FULL'` with `model_sha256` [S13].
- `S2: Degraded Sub-Sampled`: Thermal throttling, single-lane drop, or UPS conservation policy active; pulsing amber banner and chime (`REDUCED COVERAGE`); enforces low-speed ceiling and logs `inspection_mode='AUTO_DEGRADED'` with `coverage_ratio` [S03] [S13].
- `S3: Manual Digital Fallback`: Inference engine or camera fault while local HMI and SQLite remain healthy; high-contrast blue banner (`MANUAL INSPECTION ONLY`); operator acknowledges takeover, inspects visually at manual speed, and logs faults via touch keypad (`inspection_mode='MANUAL_DIGITAL'`) [S05] [S13].
- `S4: Paper Tally Emergency`: Edge compute or touchscreen HMI unpowered while lighted board remains live; hardware stack-light relay drops out (`DIGITAL RECORD OFFLINE`); operator logs faults on serialised paper forms for post-recovery transcription (`inspection_mode='MANUAL_PAPER_BACKFILL'`) [S05].
- `S5: Unsafe to Inspect (Hold)`: Board illumination lost, encoder dead, or roll identity unverified; flashing red beacon and continuous buzzer (`STOP ROLL`); roller halted and physical quarantine tag applied (`inspection_mode='SUSPENDED_HOLD'`) [S03].

Two hardware and interface rules enforce these states on the shop floor. First, the HMI must pair a live visual heartbeat counter (updated every 100 ms, an illustrative cadence) with a fail-safe physical stack-light relay wired to the `systemd` watchdog so that a frozen display screen immediately drops the relay and illuminates the overhead warning beacon [S03] [S21]. Second, any transition from `S1` into `S2` or `S3` must require an explicit operator acknowledgement button press before releasing the roller speed lock [S03] [S05].

### 9.12 STRUCTURED FALLBACK ENTRY PRESERVES SCHEMA PARITY AND AUDIT INTEGRITY

When the station operates in `S3 (Manual Digital Fallback)`, the local HMI switches into a fast-touch defect keypad writing to the exact same local SQLite schema and transactional outbox used during automated operation [S13]. The live encoder metre counter remains visible above high-contrast touch buttons for the plant's standard textile defect taxonomy (broken warp end, missing pick, slub, oil stain, reed mark, dye streak, selvedge tear, pinhole) and severity grades (`Minor`, `Major`, `Critical - Cut`, `Quarantine Roll`). Pressing a fault button commits an atomic SQLite transaction containing a locally generated `UUIDv7` or monotonic `event_id` [S31], `roll_id`, `metre_offset`, defect code, `inspection_mode = 'MANUAL_DIGITAL'`, `model_version = NULL`, `operator_id`, and an RFC 3339 timestamp with explicit UTC offset and `clock_quality` holdover flag [S30] [S32].

If the digital station suffers a complete hardware failure while utility power feeds the lighted table (`S4`), the operator uses pre-printed, sequentially numbered paper inspection cards mirroring the digital fields (`Card Serial #`, `Roll ID`, `Operator ID`, `Start/End Metre`, and `Metre Mark | Defect Code | Severity`). Upon station recovery, the supervisor transcribes the card into SQLite under `inspection_mode = 'MANUAL_PAPER_BACKFILL'` with the mandatory `paper_form_serial` key [S05] [S13].

> **GFT NOTE**
>
> **Record coverage intervals, not just point defects, in the local PMS ledger.** Most factory databases record only positive defect events: if an illustrative 100-metre roll has three rows in the `defects` table, downstream software assumes the remaining 99.7 metres were inspected and found fault-free. During power sags, thermal throttling, or camera reboots, that assumption is false. GFT recommends adding an explicit `inspection_coverage_spans` table to the local SQLite schema (`roll_id`, `start_metre`, `end_metre`, `inspection_mode`, `model_digest`, `mean_fps`, `operator_id`). Whenever an inspection mode changes or a camera drops out for even an illustrative 0.5 metres of cloth travel, the daemon closes the active span and opens a new `UNINSPECTED_GAP` or `MANUAL_DIGITAL` span. A roll can be cleared for automatic shipment only if the union of its verified coverage spans (`AUTO_FULL` or `MANUAL_DIGITAL`) covers 100 percent of its linear length [S03] [S13].

9.13

# A supervisor needs the local view before the remote dashboard

### 9.14 SHIFT LEADERSHIP DEPENDS ON LOCAL PLANT-LAN VISIBILITY

During a wide-area network outage or prolonged mains failure on site backup power, executive dashboards in a remote cloud freeze at the last upstream sync. For off-site planners, a four-hour lag in aggregate metrics is tolerable once labelled as stale (Chapter 7) [S05]. For the shift supervisor and quality manager walking the shop floor in Douala, Garoua, or Ngaoundere, flying blind for four hours is operationally unacceptable [S03] [S05]. Supervisors make floor control decisions every fifteen to thirty minutes (illustrative cadence): halting an upstream loom or dyeing batch when repetitive warp breaks or shade barring appear rather than weaving another 50 metres of second-grade cloth [S03]; inspecting accumulated defect penalty points per 100 square metres (illustrative quality threshold) before releasing a roll to the cutting room; and dispatching maintenance when a station drops into S2 or S3 or experiences high SQLite WAL growth [S14] [S0901]. If the supervisor UI is served from a cloud host, or hard-codes Single Sign-On (SSO) redirects to an external identity provider or public Content Delivery Networks (CDNs), a WAN cut locks the supervisor out of local machines standing twenty metres away (illustrative floor distance) [S03].

### 9.15 LOCAL TELEMETRY SCRAPING DECOUPLES FLOOR CONTROL FROM CLOUD REACHABILITY

To keep shop-floor leadership effective under all connectivity conditions, the primary operational UI must be served directly from the local plant network with zero external runtime dependencies [S03] [S21]. All HTML, JavaScript, CSS, and font assets must be packaged inside the local image and served over the plant LAN via a local reverse proxy using locally trusted TLS 1.3 certificates issued by an internal factory Certificate Authority [S23], backed by local credential verifiers [S03].

Beneath the supervisor UI, each edge inspection node exposes a local `/metrics` endpoint over the plant LAN reporting `gft_inspection_state_code` (states S1 to S5), `gft_inference_latency_ms_bucket` (p50, p95, p99), `gft_frames_dropped_total`, `gft_uninspected_metres_total`, `gft_sqlite_wal_bytes`, and `gft_ups_battery_charge_ratio` (illustrative metric names) [S0902]. A local Prometheus instance on the protected site gateway scrapes these endpoints across the wired LAN every 5 to 10 seconds (illustrative scrape interval) [S0902]. As documented in the Prometheus configuration specification, Prometheus is configured via command-line flags and a local configuration file, where command-line flags govern immutable local storage parameters and the configuration file defines scrape jobs and rule files [S0902]. Because scraping and rule evaluation execute entirely on the local host, the supervisor's local dashboard renders live latency heatmaps, roll defect tallies, and station mode alerts throughout a multi-day internet outage, and malformed configuration reloads are rejected without stopping the service [S0902].

9.16

# Energy scarcity should trigger a declared compute policy

### 9.17 BATTERY STATE OF CHARGE GOVERNS A THREE-TIER COMPUTE HIERARCHY

Optical inspection stations consume far more electrical power than simple barcode or MQTT transaction gateways. While a record-keeping gateway may draw 15 W to 25 W, an optical inspection station combines an accelerated edge computer (60 W to 150 W under GPU/NPU load), constant-current LED line lights (80 W to 200 W), industrial PoE cameras (15 W to 30 W), and a touchscreen monitor (25 W to 40 W), bringing a single station's total load to an illustrative 180 W to 420 W [S03]. As shown in Chapter 4's illustrative battery calculation (C01), sustaining a 150 W load for 2 hours requires about 521 Wh of nominal battery capacity. Running multiple optical stations at full GPU and LED load during a prolonged blackout will rapidly exhaust the UPS bank and crash the local PMS database and plant network switch [S03] [S05].

Because losing an uncommitted SQLite transaction or corrupting the event outbox is far more damaging than deferring optical inference to manual inspection [S13] [S15], the edge architecture must enforce a **three-tier compute-shedding policy**:

- **Tier 3 (Shed immediately on mains loss)**: Raw video archiving, local analytics compaction, and bulk historical backfill over high-power cellular/satellite radios, shed upon `ONBATT` [S0901].
- **Tier 2 (Sustained for bridge window, then throttled or shed)**: High-frame-rate GPU/NPU inference and high-wattage LED illumination, sustained during the initial 3 to 10 minutes of mains loss (waiting for generator transfer) or down to an illustrative 60 percent battery state of charge, then stepped down or suspended [S03] [S0901].
- **Tier 1 (Protected until low-battery shutdown)**: Local SQLite WAL database (`synchronous=FULL`), local Mosquitto MQTT broker [S20], operator keypad capture, and local supervisor read view, preserved until `LOWBATT` triggers an orderly `fsync` and halt [S13] [S16] [S0901].

### 9.18 AUTOMATED UPS TRANSITIONS RECORD EVERY COVERAGE BOUNDARY

Under Network UPS Tools (NUT), the `upsmon` daemon monitors each UPS and emits "notify events" including `ONLINE`, `ONBATT`, `LOWBATT`, and `FSD` (forced shutdown) [S0901]. GFT's reference architecture binds `upsmon` notify events and `battery.charge` polling directly to `inspect-daemon`:

1. **Mains loss (`ONBATT`, `battery.charge` ⩾ 60% illustrative)**: `upsmon` emits `ONBATT` [S0901]. The node halts Tier 3 tasks, caps GPU power limits, and displays `POWER: ON UPS BATTERY` on the HMI while keeping `S1` inspection active during the generator start window.
2. **Extended outage (`ONBATT` > 5 minutes or `battery.charge` < 60% illustrative)**: If the generator fails to take the load and the table remains operable, the controller transitions to `S2 (Degraded Sub-Sampled)` or `S3 (Manual Digital Fallback)` [S03] [S0901]. In `S3`, the daemon logs the exact encoder metre mark in

`inspection_coverage_spans`, requires operator acknowledgement, unloads the GPU worker, and cuts the high-wattage camera/LED relay while keeping the low-power keypad and SQLite ledger alive on Tier 1 power [S13].

3. **Critical battery exhaustion (`LOWBATT` or `battery.runtime < 180s` illustrative)**: `upsmon` emits `LOWBATT` [S0901]. The daemon transitions the HMI to `S5 (Unsafe to Inspect - Hold)`, instructs the operator to tag the stopping metre mark, runs `PRAGMA wal_checkpoint(TRUNCATE)` [S14] [S16], flushes filesystem write barriers via `sync()` [S29], and unmounts `/var/lib/gft/pms-db` cleanly before inverter cut-off.

> **WARNING**
>
> **Never reduce inspection coverage silently to save power.** A power-saving script that quietly lowers camera frame rates from an illustrative 12 fps to 2 fps, disables two out of three cross-width cameras, or swaps a full-resolution defect detector for a lightweight model without alerting the operator and logging the change in the database corrupts quality assurance [S03]. Every energy-driven reduction in optical or model coverage must simultaneously (1) trigger an unmistakable amber or blue mode change on the local HMI, (2) enforce the corresponding roller speed cap or manual-inspection rule, and (3) write an explicit `AUTO_DEGRADED` or `MANUAL_DIGITAL` coverage span into the permanent roll record [S03] [S13].

9.19

# Model updates must have an offline recovery route

### 9.20 A/B SLOT ISOLATION PREVENTS PARTIAL WRITES FROM DISABLING THE STATION

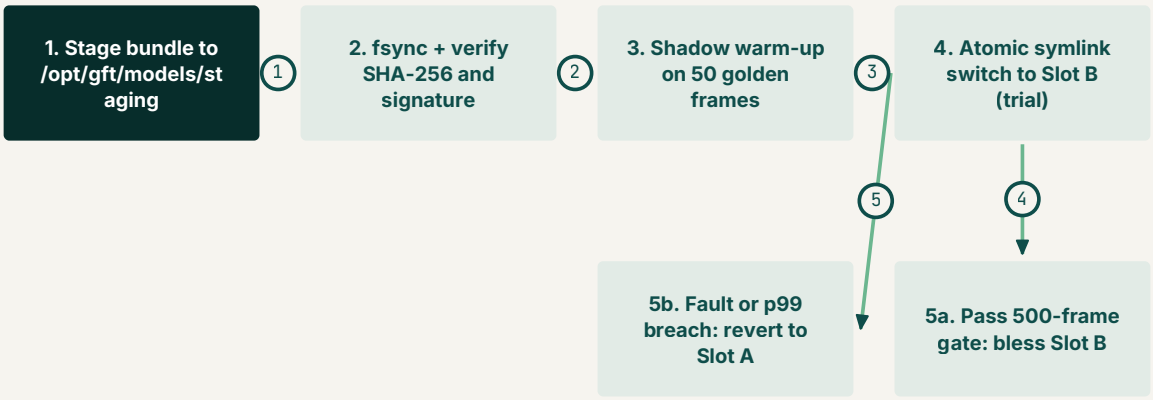
As the mill introduces new fabric weaves, changes dye recipes, or qualifies improved defect-detection weights during pilot trials (`wp03`, `wp14`), updated model bundles must be deployed to shop-floor stations [S04]. Under intermittent connectivity and unstable power, overwriting `/opt/gft/models/current.onnx` in place is a major reliability hazard: if the WAN link drops at an illustrative 78 percent transfer completion or a mains sag reboots the node mid-write, the active model is left corrupt [S04] [S15]. Even when a download completes, a tensor library mismatch can cause segfaults or double `p99` inference latency from an illustrative 75 ms to 190 ms, violating the Mode B reaction window.

Following NIST SP 800-193 (`Protect`, `Detect`, `Recover`) [S04] and NIST SP 800-184 recovery playbooks [S06], every edge inspection node must implement **immutable A/B model slot isolation**. Persistent storage maintains `/opt/gft/models/slot-a/`, `/opt/gft/models/slot-b/`, and a read-only `/opt/gft/models/slot-factory/`, with an atomic POSIX symlink (`/opt/gft/models/active`) pointing to the blessed slot. Incoming updates stream into `/opt/gft/models/staging/` (over resumable WAN transfers or signed USB media) while `inspect-daemon` continues reading from `active`. Before sta-

ging completes, the updater executes `fsync` [S29], verifies the bundle's cryptographic signature and SHA-256 digests against `manifest.json` [S04], and promotes the verified bundle into the inactive slot (`slot-b`).

EXHIBIT 38 / QUALITATIVE SEQUENCE

## An interrupted update returns to a known-good service




1. Transfer complete
2. Digest valid (between rolls)
3. p99 <= budget
4. Healthy keepalive
5. Watchdog or latency fault

ILLUSTRATIVE [S04] [S06] [S13] [S21] Illustrative A/B staging, cryptographic verification, health-gated trial, and automatic offline rollback workflow for edge inspection models. Synthesised from NIST SP 800-193, SP 800-184, SQLite atomic commit, and systemd supervision.

### 9.21 HEALTH-GATED ACTIVATION REVERTS FAULTY MODELS WITHOUT WAN ACCESS

Exhibit `x09-model-update-rollback-flow` traces the five-step transition that guarantees an edge station can always recover to a known-good inspection service without external network assistance [S04] [S06] [S13] [S21]. First, **an active fabric roll is never interrupted mid-roll to swap models**: switching weights at an illustrative metre 65 of a 120-metre roll alters detection sensitivity mid-piece, so `slot-b` waits in `STAGED_READY` until roll completion [S03]. Second, before flipping `/opt/gft/models/active` to `slot-b`, a shadow warm-up runs `slot-b` against 50 locally stored golden reference frames (illustrative test set) to verify tensor shapes, deterministic outputs, and `p99` latency below `T_budget`. Once switched, SQLite records `TRIAL_PENDING` [S13]. If the new model crashes, misses `WatchdogSec=15s` [S21], or breaches the `p99` latency ceiling during its first 500 live frames (illustrative trial gate), the startup wrapper

atomically reverts `/opt/gft/models/active` to `slot-a`, marks `slot-b` as `QUARANTINED_FAILED`, and logs `MODEL_AUTO_ROLLBACK` in SQLite and `systemd-journald` [S13] [S21] [S22].

*A model that lives on the edge but depends on the cloud for its licence, its screen, or its recovery is only a remote service waiting to fail.*

## CHAPTER SOURCE NOTES

S03 NIST. 2023.

S13 SQLite. 2026-10-30.

S05 NIST. 2010.

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

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

S01 World Bank. 2025.

S04 NIST. 2018.

S18 IETF / RFC Editor. 2014.

S21 systemd project. 2026-10-30.

S16 SQLite. 2026-10-30.

S23 IETF / RFC Editor. 2026.

S30 chrony project. 2026-10-30.

S29 Linux kernel project. 2026-10-30.

S14 SQLite. 2026-10-30.

S08 IETF / RFC Editor. 2007.

S31 IETF / RFC Editor. 2024.

S32 IETF / RFC Editor. 2002.

S0901 Network UPS Tools. 2026-10-28.

S0902 Prometheus Authors. 2026-10-30.

S15 SQLite. 2026-10-30.

S20 Eclipse Mosquitto. 2026-10-30.

S06 NIST. 2016.

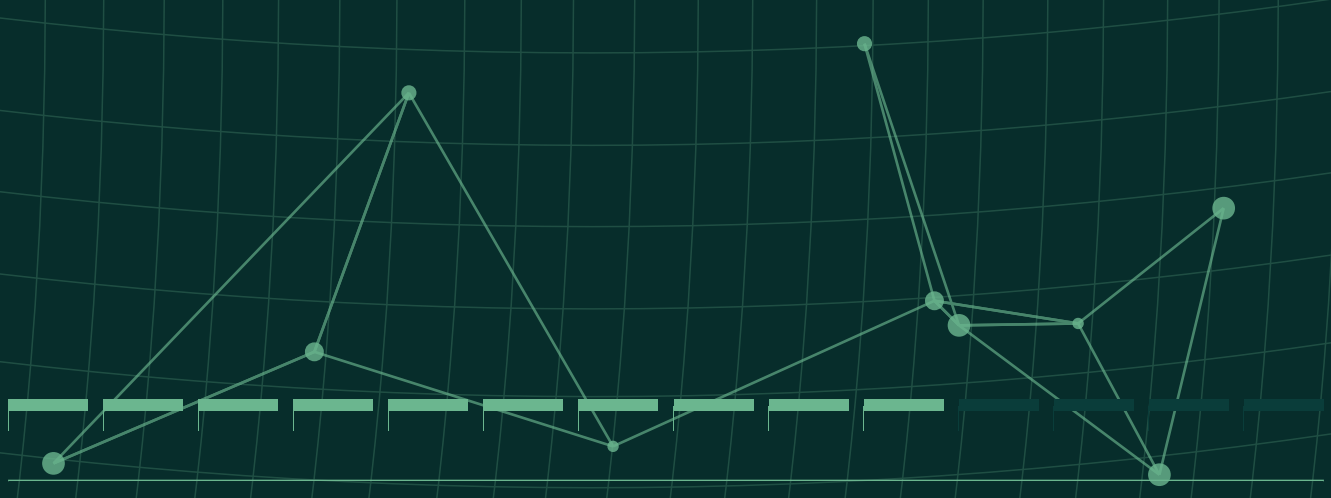
S22 systemd project. 2026-10-30.

10

PART III: RECOVER A TRUSTWORTHY OPERATION / MAINTAIN TRUST OFFLINE

# Disconnection changes security dependencies rather than removing them

*Local credentials, signed software, controlled access and recoverable evidence must keep working through communication loss.*



CHAPTER 10
DECISION TEST

Local credentials, signed software, controlled access and recoverable evidence must keep working through communication loss.

Losing the external backhaul does not remove a factory's security dependencies. It removes online verification paths while local users, maintenance interfaces, software and records still require protection. [S03] [S23] [S27] Green Future Tech Limited (GFT) therefore designs the disconnected state as a distinct security mode: local Operational Technology (OT) zones retain controlled access, cached permissions expire, certificate recovery respects clock uncertainty, signed software remains verifiable, and audit evidence persists within explicit storage limits. [S03] [S04] [S22] [S27] [S1002]

This chapter extends GFT's Cameroon textile programme, whose feasibility, Performance Management System (PMS) design and digital strategy are design work rather than evidence of a commissioned facility. Earlier chapters established event identity, clock quality, local inspection and human fallback. The present task is to maintain cryptographic trust, access control, software integrity, auditability and safe incident containment when external services are absent. Enterprise data governance remains in paper wp14, regulatory analysis in wp17, and integration architecture in wp18. [S03] [S25] [S26] This is not legal, tax or investment advice. Referencing a security framework does not assert compliance or certification. [S26]

10.1

## Map trust boundaries to the local operating duties

A flat factory network exposes machinery controllers and the production ledger to every maintenance laptop and compromised workstation, whether or not the internet router is reachable. [S03] NIST Special Publication (SP) 800-82 Revision 3 requires OT security to address performance, reliability and physical safety alongside information protection. [S03] A security control that delays a safety stop, disrupts a Programmable Logic Controller (PLC) loop or locks out a supervisor during a power transfer is therefore an operational hazard, not a harmless IT inconvenience. [S03] [S05]

GFT partitions the site into four duty-based trust zones so a security failure neither permits unauthenticated control changes nor unnecessarily stops safe local work: [S03] [S26]

1. **Zone 1: Production control and machinery safety.** PLCs, drives and hardwired interlocks accept only authorised control traffic. External services and unmediated laptops cannot issue actuation commands. [S03]
2. **Zone 2: Local monitoring and durable capture.** Scanning terminals, inspection nodes, the local Message Queuing Telemetry Transport (MQTT) broker and the SQLite event-and-outbox store continue their approved local duties without Wide Area Network (WAN) access. [S03] [S13]
3. **Zone 3: Site operations and management.** The PMS console, local time service, update repository, credential cache and persistent log collector supervise capture through locally verifiable credentials. [S03] [S04] [S22] [S30]

4. **Zone 4: External reporting and supplier gateway.** This boundary terminates backhaul, outbound replay and controlled support paths. It initiates outbound Transport Layer Security (TLS) sessions when service returns, without requiring inbound access to the other zones. [S03] [S23]

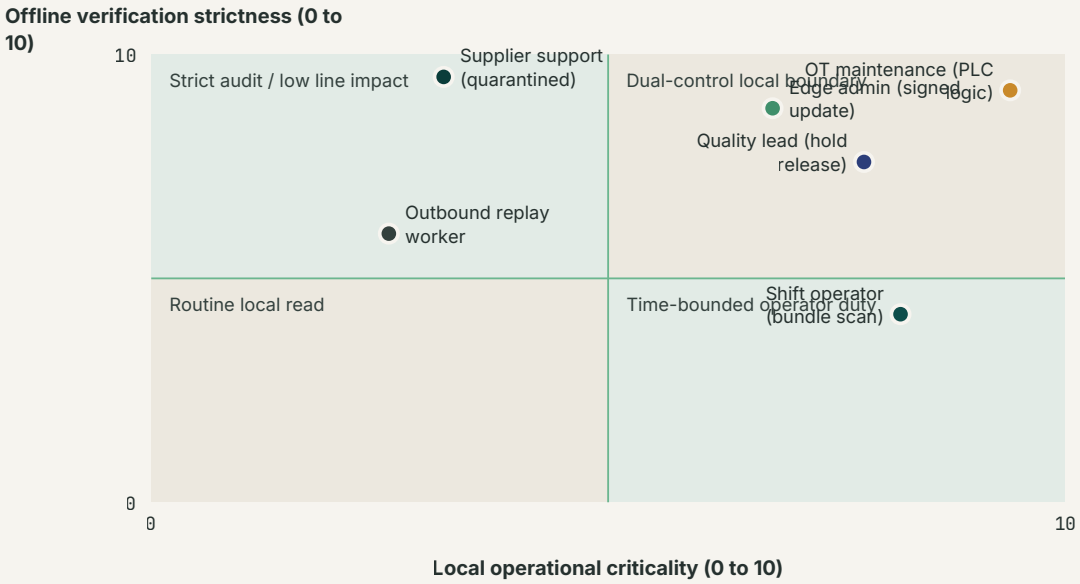
### 10.2 OUTCOME FRAMEWORKS STRUCTURE GOVERNANCE WITHOUT CERTIFYING COMPLIANCE

The NIST Cybersecurity Framework (CSF) 2.0 organises outcomes under GOVERN, IDENTIFY, PROTECT, DETECT, RESPOND and RECOVER, but does not prescribe how to achieve them. [S26] GFT uses those categories to give each proposed control an online and disconnected behaviour, an accountable role and an acceptance observation. The organisation must still compare its Current and Target Profiles against its own mission and risk. [S26] An illustrative disconnection window of 8 to 72 hours is used only to test that thinking; it is neither a site forecast nor a compliance claim. [S05] [S26]

The allocation of plant roles and zone transitions across local operational criticality and offline verification strictness is illustrated in **Exhibit 39**, p. 157. [S03] [S26]

EXHIBIT 39 / ILLUSTRATIVE SCORE (0 TO 10)

## Each trust boundary needs an offline access rule




ILLUSTRATIVE [S03] [S26] Illustrative two-dimension placement of plant roles and zone transitions on local operational criticality (x-axis, 0 to 10) and offline verification strictness (y-axis, 0 to 10), derived from OT segmentation guidance in S03 and governance outcome categories in S26. Positions are analytical design allocations for GFT's proposed architecture, not measured plant security scores.

In **Exhibit 39**, p. 157, routine bundle scanning has an illustrative criticality score of 8.2 but a verification-strictness score of 4.2: access remains simple within a named station role. PLC modification and signed updates score higher on both axes and require dual-role authorisation. Supplier support receives strict verification and low direct line authority, with diagnostics held in the site quarantine. [S03] [S04] [S26] These positions are design choices, not measured security scores, and they cannot replace a site hazard study. [S03]

10.3

# Local access needs expiry and recovery rules

When an enterprise relies on cloud single sign-on or a headquarters directory, a backhaul failure severs the path between shop-floor terminals and the identity provider. [S03] [S05] If terminals require a live online token exchange before login, an outage locks supervisors out of disposition screens and blocks technicians from acknowledging alarms; conversely, shared non-expiring accounts (such as a generic `admin` login taped to the cabinet) or anonymous broker writes destroy audit accountability. [S03] [S13] [S26]

GFT specifies a four-part offline access contract under the `PROTECT` (`PR.AA`) and `GOVERN` (`GV.RR`) categories of CSF 2.0 and NIST SP 800-82 Rev. 3: [S03] [S26]

- **Named role bindings with least-privilege scope.** Every identity is bound to one of five roles: `role:shift-operator` (records bundle events in Zone 2), `role:quality-lead` (releases or rejects held bundles and authorises restricted manual inspection fallback), `role:shift-supervisor` (annotates anomalies, initiates manual fallback and co-signs emergency windows), `role:ot-maintenance` (views diagnostics, swaps spares and applies signed updates in Zones 2 and 3), and `role:edge-admin` (manages trust stores, journal budgets and backups in Zone 3), while traffic between capture processes, the MQTT broker and SQLite uses distinct machine credentials. [S03] [S23] [S26]
- **Bounded local credential caching.** Zone 3 gateways and Zone 2 stations maintain an encrypted role-assertion cache synchronised when the upstream directory is reachable, carrying a normal online refresh interval (an illustrative 24-hour target) and a hard offline survival horizon exceeding the outage envelope (an illustrative 14-day lease for active staff against a 72-hour design outage), whereas contractor accounts carry a short offline expiry tied to the work window (an illustrative 8-hour shift lease). [S03] [S05] [S26] [S1001]
- **Local revocation propagation and sequence floors.** Under RFC 5280, a CRL carries `thisUpdate` and `nextUpdate` timestamps bounding revocation currency. [S1001] Because disconnected nodes cannot reach external Online Certificate Status Protocol (OCSP) responders or remote CRL distribution points, GFT places a local revocation list and a monotonically increasing `credential_policy_seq` counter on the Zone 3 controller; when a badge is lost offline, the administrator increments `credential_policy_seq` on Zone 3 and broadcasts the signed delta across the Local Area Network (LAN) so Zone 2 stations invalidate the cached token within the operator-approved propagation deadline. [S03] [S23] [S1001]
- **Auditable dual-control emergency access ("break-glass").** If an outage outlasts a technician's cached lease while the identity server is unreachable, a holder of `role:shift-supervisor` and a holder of `role:ot-maintenance` present their local tokens together to open a time-boxed window (an illustrative 4-hour lease), appending an `EMERGENCY_ACCESS_GRANTED` audit event to SQLite and the persistent system journal with both role IDs, the kernel `boot_id` (`sd_id128_get_boot()`), a 128-bit `UUIDv4` or `UUIDv7` identifier, the monotonic clock reading and the reason code. [S03] [S05] [S06] [S13] [S22] [S31] [S1005]

**GFT NOTE**

Never configure an edge operating system, database or local HMI to bypass authentication automatically when a remote identity provider times out. Agree the offline survival lease and dual-person break-glass procedure with plant operations before commissioning, test login flows with the WAN cable unplugged, and require every emergency elevation to generate a reconciliation ticket on the PMS dashboard until reviewed after connectivity returns. [S03] [S05] [S06]

10.4

# Certificate and time failures can block restoration

Certificate validation needs a trustworthy wall clock, while authenticated time synchronisation uses certificate-based Transport Layer Security (TLS). [S23] [S27] [S1001] After a cold boot with a failed or uninitialised real-time clock (RTC), an edge device can therefore reject a valid server certificate as expired or not yet valid, and authenticated time recovery stalls just when the plant needs to reconnect. [S27] [S30] [S1001]

### 10.5 QUALIFY CURRENT TLS 1.3 SPECIFICATIONS ON THE TARGET STACK

For restored site and remote connections, GFT specifies TLS 1.3 as defined by RFC 9846, published in July 2026, subject to qualification of the actual device library. [S23] The implementation must retain strict service identity verification and must not silently negotiate an obsolete protocol version when authentication fails. [S23] [S1003]

Two TLS 1.3 features need explicit policy. Session resumption can use pre-shared keys with an ephemeral key share for forward secrecy, while externally provisioned keys require cryptographic entropy rather than passwords. [S23] Zero round-trip time (`0-RTT`) early data lacks the same replay protection as a full handshake. [S23] Idempotent event replay still needs application deduplication, but a repeated configuration or state-reset command could cause a second business action. GFT therefore disables `0-RTT` on control, configuration and security channels. [S12] [S23] [S31]

### 10.6 BREAK THE NTS AND X.509 BOOTSTRAP DEADLOCK WITHOUT DISABLING VALIDATION

RFC 8915 separates Network Time Security (NTS) into TLS-based NTS Key Establishment (`NTS-KE`) and authenticated NTP extension fields. The server remains stateless by giving the client opaque cookies, while the client retains the keys and cookies needed for later authenticated NTP exchanges. [S10] [S27]

X.509 certificates carry `notBefore` and `notAfter` limits, so ordinary validation cannot complete when the client clock is outside the certificate window. [S1001] RFC 8915 permits a carefully bounded approach to this initialisation problem, but also warns against NTS stripping: an attacker can obstruct key establishment or induce a client to abandon authenticated time. [S27] A permanent certificate-time bypass or automatic return to plain NTP is therefore unacceptable. [S27] [S30]

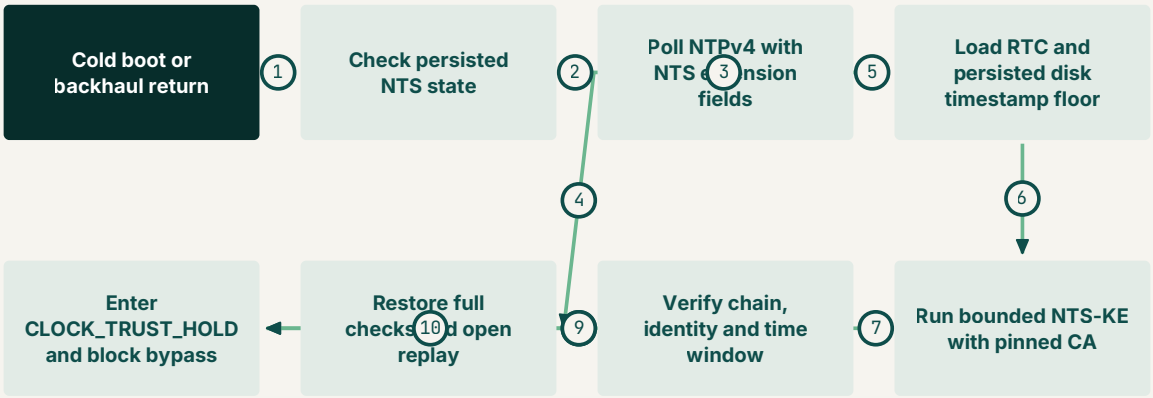
GFT uses a five-step recovery procedure that treats persisted state and clock bounds as evidence to test, not assumptions to trust blindly: [S27] [S30]

1. **Persist NTS state with tested writes.** The client stores unused cookies and keys so an ordinary restart can resume authenticated NTP without a fresh NTS-KE exchange. [S27] [S30] Because state may be written on orderly shutdown or on an explicit dump, commissioning must prove that the chosen atomic-write and dump schedule survives abrupt power loss. Absent, exhausted or rejected state returns the device to key establishment, not to plain NTP. [S27] [S30]
2. **Restore a conservative time floor.** The node records its last authenticated UTC reading and restores a clock no earlier than that saved floor. RTC synchronisation or an RTC drift file may support this process, but those chrony features do not themselves create the application floor. [S27] [S30]
3. **Use only a bounded initialisation exception.** If the restored clock prevents NTS-KE, the stack may exempt certificate-time checking for one key-establishment attempt while retaining certificate-chain and hostname verification. It confirms that `notAfter` is later than the disk floor, then accepts an authenticated NTP reply only when that reply lies inside the certificate validity window. [S27] [S1001] A stack without such a bounded control enters recovery hold.
4. **Restore full validation immediately.** After the authenticated reply sets the clock, the node removes the exception and requires ordinary `notBefore` and `notAfter` checks for later TLS connections. [S27] [S1001]
5. **Hold rather than bypass on conflict.** A failed chain, inconsistent reply or regression below the disk floor produces `CLOCK_TRUST_HOLD`. Local capture continues with sequence and clock-quality flags, while remote replay waits for authorised review. [S03] [S27] [S30]

The decision sequence coupling TLS 1.3 certificate validation and NTS time recovery after cold boot is shown in **Exhibit 40**, p. 161. [S23] [S27] [S1001]

EXHIBIT 40 /

## Trust and time must recover together




1. Start time recovery
2. Valid cookie present
3. State absent or rejected
4. AEAD tag valid
5. Refresh required
6. Apply floor
7. Chain signed by pinned CA
8. Untrusted CA or before floor
9. Consistent
10. Window mismatch

ILLUSTRATIVE [S23] [S27] [S1001] Illustrative bootstrap sequence after cold boot or WAN restoration, based on TLS 1.3 (S23), X.509 validity rules (S1001) and NTS state persistence and certificate recovery rules (S27). A one-attempt certificate-time exception, where supported, retains chain and hostname checks and ends after authenticated time recovery. Unverified time never triggers plain NTP or unverified TLS.

As **Exhibit 40**, p. 161 shows, valid persisted NTS state can avoid a fresh key-establishment exchange after a routine restart. [S27] [S30] Missing or rejected state invokes the bounded recovery path. Firmware verification protects the approved boot chain, but it does not prove that the RTC and storage were not physically replaced; that stronger claim would need separate hardware attestation and tamper evidence. [S04] [S27]

10.7

## Integrity checks must cover firmware and local software

On an intermittent grid, an abrupt power cut can force an edge gateway through an unscheduled reboot unless its protected supply rides through. [S03] [S05] If an operator assumes every reboot is benign, an adversary or corrupted update can modify the

boot loader, kernel, container image or configuration while offline and rely on the next power cycle to execute the altered code, while a brownout during an unprotected write can leave a boot partition unbootable. [S04] [S15] [S25]

NIST SP 800-193 ("Platform Firmware Resiliency Guidelines", May 2018) organises resiliency around three properties that GFT extends across platform boot firmware and the local software stack: [S04]

1. **Protection (`Protect`).** Preserves firmware and boot configuration integrity through hardware roots of trust, write-protected memory regions and signature verification before flash writes. [S04]
2. **Detection (`Detect`).** Detects corruption or unauthorised modification through a Root of Trust for Detection (RTD), or a chain anchored in one, using trusted comparison against authorised reference values. Measurements from a verified boot chain can support that decision but do not replace it. [S04] [S1002]
3. **Recovery (`Recover`).** Restores firmware and critical data to a known-good state when tampering or corruption is detected, so recovery does not depend on executing the previously running image. [S04]

### 10.8 POWER-LOSS RECOVERY AND COMPROMISE RECOVERY REQUIRE DIFFERENT PLAYBOOKS

Recovering from an electrical outage differs fundamentally from recovering from a security compromise or integrity failure. [S04] [S06] [S25] After an ordinary mains failure where verified boot measurements confirm that firmware, OS image and signed binaries match approved digests, the node executes power-loss recovery: it checks the filesystem journal, replays or rolls back the SQLite write-ahead log (WAL) as analysed in Chapter 6, verifies the last committed `source_seq` and resumes capture. [S04] [S13] [S16] [S29]

When a boot measurement, file-integrity check or journal seal verification fails, or an intrusion indicator triggers under NIST SP 800-61 Rev. 3, treating the event as a normal power-cut restart is unsafe: re-executing the mutable snapshot reloads the tampered state, and copying or replaying a SQLite database written by a compromised process via `SQLite Backup API` (S17) pollutes the central ledger. [S04] [S13] [S17] [S25] [S1004] GFT therefore enforces two architectural controls: [S04] [S06] [S25]

- **Immutable active and recovery boot slots.** Edge appliances in Zones 2 and 3 provide a hardware root of trust and dual read-only root partitions (Slot A and Slot B) plus a factory-signed recovery image in write-protected flash, keeping runtime state on separate data partitions mounted with `noexec`, `nosuid` and `nodev`. [S04] [S21] [S22] [S1002]
- **Quarantine of the mutable ledger on integrity failure.** If verified boot detects an unauthorised change and falls back to the recovery slot, startup mounts the mutable production database in read-only quarantine mode rather than starting the replay worker, raises a Zone 3 alarm and switches the bay to the manual fallback defined in Chapter 3 until the ledger is inspected against physical bundle counts. [S03] [S04] [S06] [S13] [S25]

10.9

# Offline updates require controlled media and staged activation

When an illustrative fleet of 30 edge stations pulls large OS or vision-model updates over constrained backhaul, downloads can compete with event replay and truncate on power loss. [S03] [S18] The proposed design therefore uses a local staging server in Zone 3 or controlled removable media. Because removable media and maintenance laptops can bypass OT defences under NIST SP 800-82 Rev. 3, GFT adopts the transport-agnostic signed manifest architecture in RFC 9019 ("A Firmware Update Architecture for Internet of Things", April 2021) alongside NIST SP 800-193 and SP 800-184. [S03] [S04] [S06] [S1002]

## 10.10 A SIGNED MANIFEST GOVERNS COMPATIBILITY, STAGING AND ROLLBACK

Under RFC 9019 and SP 800-193, an edge device verifies five properties before activating any update package: [S04] [S1002]

1. **Detached signature and update authority.** The manifest is signed by an offline key pinned in the read-only trust store; per RFC 9019, code-signing authority is separated from operational logins so a technician session can deliver a signed package but never install unsigned code. [S03] [S04] [S1002]
2. **Hardware and dependency compatibility.** The manifest declares device class, hardware revision, minimum bootloader version and required SQLite schema version, preventing loading a cutting-table image onto a spinning-frame gateway. [S03] [S1002]
3. **Monotonic sequence number (anti-rollback protection).** The manifest carries a `manifest_seq` that may not be lower than the device's protected security watermark. A same-version reinstall is permitted only by explicit policy; otherwise the new sequence must be strictly greater. The watermark advances only after the trial boot passes its health checks, so recovery to the previous slot remains possible. [S04] [S1002]
4. **Full payload digest verification in an inactive staging slot.** Updates imported via the Zone 3 quarantine are written exclusively to the inactive partition (Slot B while Slot A runs capture); if power fails mid-transfer, Slot A and the SQLite ledger remain untouched, and on reboot the node verifies the full SHA-256 digest before marking Slot B ready. [S04] [S13] [S1002]
5. **Trial activation with watchdog and health-check commit.** Slot B enters `TRIAL_BOOT` under readiness checks and a hardware watchdog. The software must verify its SQLite schema, open capture and pass self-test within an illustrative 120-second deadline. Only then does it commit the slot and security watermark; otherwise it reverts to Slot A and logs `UPDATE_ROLLBACK_EXECUTED`. [S04] [S06] [S13] [S21]

## 10.11 RETAIN OFFLINE MAINTENANCE EVIDENCE WITHOUT OPENING INBOUND PORTS

Before activating a staged package offline, the station writes a signed maintenance event into its SQLite audit outbox and flushes its journal (`journalctl --sync`), recording technician role ID, old and new `manifest_seq` and SHA-256 digests, pre-update

source_seq, kernel boot_id and timestamp with clock-quality state. [S13] [S22] [S1004] [S1005] When backhaul returns, the Zone 4 worker replays these records over the outbound TLS 1.3 channel alongside production events, never requiring inbound WAN ports. [S03] [S23]

10.12

## Local logs need an explicit persistence and capacity policy

Volatile logs can erase the evidence of events preceding a power cut. Unbounded logs on the production filesystem can instead crowd out the SQLite event store. [S13] [S15] [S22] The security design therefore needs both persistence and a separate capacity boundary.

### 10.13 CONFIGURE AND READ BACK JOURNALD.CONF PERSISTENCE AND LIMITS

On Linux edge gateways, GFT configures systemd-journald through drop-in files and reads back the effective settings during commissioning. [S22]

- **Persistence.** Storage=persistent selects /var/log/journal on disk, with early-boot data flushed from volatile storage when the persistent location becomes available. GFT places that directory on a non-volatile partition separate from SQLite. [S13] [S22]
- **Space controls.** SystemMaxUse=, SystemKeepFree=, SystemMaxFileSize= and SystemMaxFiles= govern rotation and disk use. These are not absolute caps or retention guarantees: vacuuming removes archived files, while active files can remain above the configured limit. The dedicated partition, not SystemMaxUse= alone, protects the production ledger from log growth. [S14] [S22]
- **Rate limits.** RateLimitIntervalSec= and RateLimitBurst= bound messages accepted from each service, and dropped-message counts must remain visible. Commissioning injects a noisy-service fault rather than assuming defaults fit the site. [S22]
- **Synchronisation.** SyncIntervalSec= controls periodic disk synchronisation, while CRIT, ALERT and EMERG events trigger immediate synchronisation. GFT also writes security events to its transactional audit outbox and disables console forwarding where a stalled console could obstruct logging. [S13] [S22]

### 10.14 DISK PERSISTENCE IS NOT TAMPER RESISTANCE: FORWARD SECURE SEALING AND ITS LIMITS

Persistence alone does not prevent a local administrator from altering or deleting records before reconnection. [S22] [S1004] Forward Secure Sealing (FSS), enabled with Seal=yes and initialised through journalctl --setup-keys, adds verifiable sealing epochs. The sealing key remains on the host and ratchets forward, while the verification key is held separately. [S22] [S1004]

journalctl --verify --verify-key=KEY can detect alterations within the sealed journal files presented to it. It cannot prove that a complete log was received, detect the absence of a deleted directory on its own, or recover deleted bytes. [S1004] GFT

therefore keeps an independent file-and-epoch inventory and replicates critical security records to the site collector when the LAN is available. Entries since the latest sealing epoch still require that continuity check. [S22] [S1004] [S1005]

The source-based controls governing firmware resiliency, persistent sealed logging and incident recovery, together with GFT's proposed commissioning evidence, are compiled in **Exhibit 41**, p. 165. [S04] [S22] [S25] [S1004]

EXHIBIT 41 /

## Recoverable evidence requires configured local retention

| Control domain | Source basis | Specified technical mechanism | Proposed GFT acceptance evidence |
|---|---|---|---|
| Firmware protection | NIST SP 800-193 (S04) | Hardware root of trust, write-protected boot flash and authenticated update verification | Readback of write-protect register and test log rejecting an unsigned update image |
| Firmware detection | NIST SP 800-193 (S04) | Cryptographic integrity measurement of boot stages before execution handover | Recorded boot measurement digest matched against signed release manifest |
| Firmware recovery | NIST SP 800-193 (S04) | Automated or controlled recovery to a known-good authenticated firmware image | Witnessed recovery drill booting Slot A after deliberate corruption of inactive Slot B |
| Journal persistence | journald.conf (S22) | Storage=persistent in /var/log/journal with systemd-journal-flush from /run/log/journal | Post-power-cut readback confirming pre-outage boot_id and CRIT entries survive on disk |
| Journal space controls | journald.conf (S22) | SystemMaxUse=, SystemKeepFree=, SystemMaxFileSize= and RateLimitBurst= bounds | Simulated log flood verifying journal rotation without encroaching on SQLite partition |
| Forward-secure sealing | journald.conf (S22); journalctl (S1004) | Seal=yes with journalctl --setup-keys ratcheting sealing key and external verification key | Clean verification of an intact log and detection of a deliberate edit |
| Incident containment | NIST SP 800-61 Rev. 3 (S25) | Risk-managed containment and evidence preservation integrated with CSF 2.0 functions | Signed drill record preserving forensic snapshot before re-imaging an isolated edge node |

SOURCED [S04] [S22] [S25] [S1004] Source-based control checklist combining platform firmware resiliency guidelines (S04), systemd journald.conf and journalctl specifications (S22, S1004) and incident response guidance (S25). The final column states GFT's proposed commissioning acceptance evidence and does not represent a completed audit of a deployed factory.

As summarised in **Exhibit 41**, p. 165, every security control that matters during an outage has both a concrete configuration mechanism in primary technical documentation (S04, S22, S25, S1004) and an observable commissioning test, whereas relying on OS defaults leaves a plant discovering after an incident that logs were held only in volatile /run/log/journal. [S04] [S22] [S25] [S1004]

METHOD

GFT uses a four-step storage method so diagnostics, updates and forensic snapshots do not compete with the production ledger for the same filesystem: [S22] [S1002] [S1004]

1. **Isolate security storage.** Put journals, update staging and forensic snapshots on `P_sec`, separate from the SQLite volume `P_prod`. [S13] [S22]
2. **Budget the offline journal.** Estimate `B_journal` from measured event size and rate across the design disconnection. Configure `SystemMaxUse` and file rotation, then test actual retention because active files and record overhead can invalidate a simple byte calculation. [S22] [S1004]
3. **Reserve recovery material.** Assign `B_update` to a signed update package and `B_forensic` to an online SQLite backup and journal export before re-imaging. [S17] [S25] [S1002]
4. **Keep a positive margin.** Set the keep-free policy and verify `B_margin = P_sec - (B_journal + B_update + B_forensic + B_keepfree) > 0`. The partition is the containment boundary; the margin is an operating reserve, not promised retention. [S22]

WORKED EXAMPLE

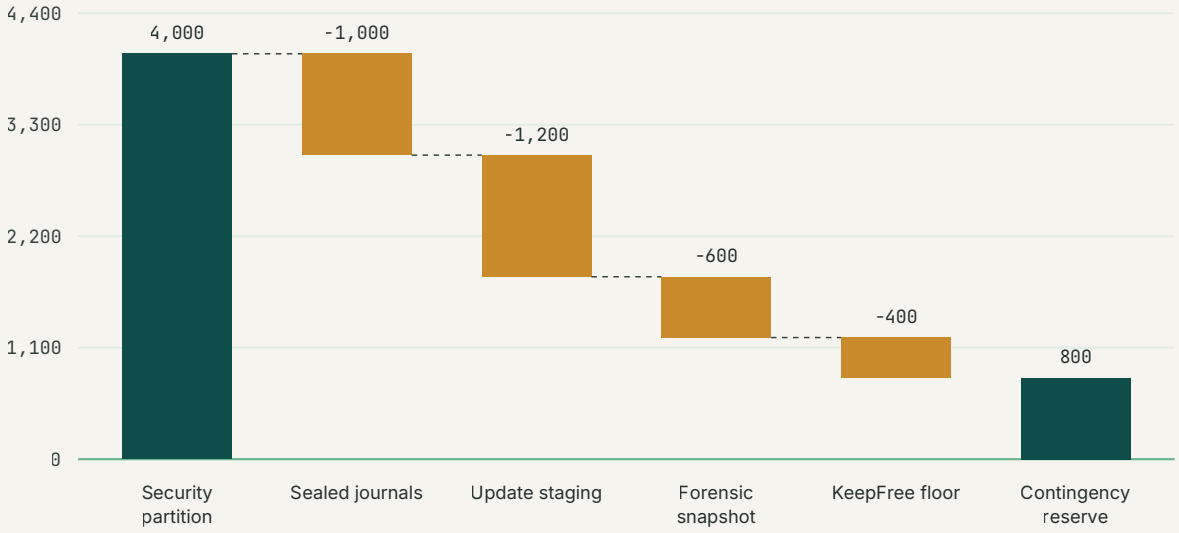
**An illustrative partition budget tests capacity without promising log retention.** Assume an illustrative `P_sec = 4,000 MB`, separate from the SQLite ledger. During an illustrative 72-hour (`259,200 s`) disconnection, local services generate `4` log entries per second at `350 bytes` each, or `1,400 bytes/s`. Raw arrival is `1,400 * 259,200 = 362,880,000 bytes`, or `362.88 decimal MB`. An illustrative 30 percent allowance for journal structure and sealing gives `362.88 * 1.30 = 471.744 MB`. [S22] [S1004] This is a sizing estimate, not a guarantee of usable retention.

Allocate an illustrative `B_journal = 1,000 MB`, `B_update = 1,200 MB`, `B_forensic = 600 MB` and `B_keepfree = 400 MB`. The remaining reserve is `4,000 - 1,000 - 1,200 - 600 - 400 = 800 MB`. [S17] [S22] [S25] [S1002] The journal allocation exceeds the estimated requirement, but `SystemMaxUse=` is not a hard cap and active files can remain outside vacuuming. The acceptance test must therefore flood the journal on the target image and confirm both the required history and the separate ledger's continued operation. [S13] [S22]

**Exhibit 42**, `p. 167` reconciles the illustrative `4,000 MB` allocation to its closing `800 MB` reserve. [S22] [S1002] [S1004]

EXHIBIT 42 / DECIMAL MEGABYTES (MB)

## A dedicated security partition protects ledger capacity




DERIVED [S22] [S1002] [S1004] Derived from an illustrative 4,000 MB local security and maintenance partition budget using the journal sizing and sealing mechanisms in S22 and S1004 and the staged update manifest architecture in S1002. Allocations subtract 1,000 MB for persistent sealed journals (SystemMaxUse), 1,200 MB for signed update staging, 600 MB for incident forensic snapshots and 400 MB for the filesystem keep-free floor, leaving an 800 MB contingency reserve.

The waterfall shows capacity allocation, not observed behaviour. Separating `P_sec` from the SQLite ledger contains the direct disk-space effect of a journal flood, while the illustrative `1,000 MB` journal budget, `1,200 MB` update reserve, `600 MB` forensic reserve and `400 MB` keep-free floor make trade-offs reviewable. [S13] [S22] Commissioning must still measure arrival rates, active-file behaviour and record survival under fault injection; no retention or availability claim follows from the arithmetic alone. [S22] [S28]

10.15

# An incident response must preserve safe plant control

When a security alert, unexpected configuration change or failed integrity seal occurs on a factory node, the response playbook cannot copy an office Information Technology (IT) procedure. [S03] [S25] On a textile shop floor, where gateways interface with spinning frames, dye vats, stenters and conveyors, an uncoordinated digital isolation or forced shutdown by an automated tool can strand heated or tensioned fabric, blind operators to alarms or cause an unsafe mechanical stop. [S03] [S25]

NIST SP 800-61 Revision 3 ("Incident Response Recommendations and Considerations for Cybersecurity Risk Management: A CSF 2.0 Community Profile", April 2025, superseding Revision 2 from August 2012) integrates incident response across the six CSF 2.0 Functions (`GOVERN`, `IDENTIFY`, `PROTECT`, `DETECT`, `RESPOND` and `RECOVER`); combined with NIST SP 800-184 (December 2016) and NIST SP 800-82 Rev. 3, any

incident response on a weak-grid factory floor must preserve safe physical control and capture local evidence before eradication overwrites the device. [S03] [S06] [S25] [S26]

> **WARNING**
>
> Never permit a remote security tool, cloud management plane or automated containment script to issue an uncoordinated process stop, PLC reset or local HMI lockout across Zone 1 or Zone 2 without the awareness and authorisation of the on-duty plant supervisor and OT authority. Under NIST SP 800-82 Rev. 3, physical safety and process stability take precedence over automated IT quarantine; containment at the Zone 4 or Zone 3 boundary must sever external and inter-zone propagation while leaving local machinery safety interlocks and local operator visibility intact until a controlled stop point is reached. [S03] [S25]

### 10.16 CONTAINMENT SEVERING FOLLOWS ZONE BOUNDARIES OUTWARD-IN

To contain a suspected compromise or corrupted update while disconnected without endangering personnel or machinery, GFT specifies a four-step outward-in sequence: [S03] [S06] [S25]

1. **Sever Zone 4 and inter-bay lateral paths first.** The Zone 3 controller blocks the affected node's credentials from writing to the site-wide MQTT broker and halts outbound WAN replay from that source ID, containing lateral spread across the LAN without cutting power to Zone 1 controllers. [S03] [S20] [S25]
2. **Transition the affected bay to declared manual fallback.** The Zone 3 console notifies the shift supervisor and switches the station's service class to the manual paper-and-spreadsheet capture state defined in Chapter 3, preserving bundle identifier, operation, shift and role continuity on physical travellers. [S03] [S05] [S31]
3. **Capture a verifiable forensic and ledger snapshot before eradication.** Before re-flashing a node or swapping a cold spare under NIST SP 800-193 and SP 800-184, the technician runs an evidence-preservation script onto `/var/lib/gft-forensics` or approved quarantine media: flush and verify sealed journals (`journalctl --sync` and `journalctl --verify --verify-key=KEY`), capture an online SQLite snapshot via the `SQLite Backup API` (S17) alongside WAL state (S14, S16) and active kernel `boot_id` (S1005), and record running firmware digests (S04) and `chrony` status (S30). [S04] [S06] [S14] [S16] [S17] [S22] [S25] [S30] [S1004] [S1005]
4. **Restore from a known-good image and reconcile in isolation.** Following Chapter 11, the node is restored from a hardware-verified known-good image (S04) or replaced with a pre-commissioned spare, re-provisioned with credentials and NTS keys from Zone 3 (S23, S27), and its quarantined events are reconciled against adjacent station ledgers before release into the central PMS. [S04] [S06] [S13] [S25]

### 10.17 REHEARSE DISCONNECTED SECURITY SCENARIOS BEFORE SCALING

NIST SP 800-184 and NIST SP 800-61 Rev. 3 stress that recovery playbooks work only when tested under realistic conditions; GFT requires three disconnected security drills prior to pilot expansion (linked to Chapter 13): [S06] [S25] [S26]

- **Drill 1: Cold boot and NTS/TLS recovery under uncalibrated clock state.** Disconnect the WAN, simulate an unpowered interval with a stepped clock, restore connectivity, and verify that the gateway enforces its disk timestamp floor (S27, S30), refuses `NTS stripping` to plain NTP (S27), and either completes authenticated `NTS-KE` over TLS 1.3 (S23) or enters `CLOCK_TRUST_HOLD` without halting capture. [S23] [S27] [S30]
- **Drill 2: Dual-control offline elevation and interrupted update rollback.** With the upstream directory unreachable, exercise a witnessed two-person break-glass login (`role:shift-supervisor` plus `role:ot-maintenance`), verify that SQLite and the sealed journal record the elevation at `CRIT` priority (S13, S22), stage a signed RFC 9019 update to Slot B (S1002), cut power during trial activation, and prove automatic rollback to Slot A (S04) without losing any committed record. [S04] [S13] [S22] [S1002]
- **Drill 3: Isolated bay containment and forensic ledger preservation.** Simulate a failed `journalctl --verify` check on one Zone 2 station while disconnected (S25, S1004), execute Zone 3 broker quarantine and manual capture handover without disturbing Zone 1 PLCs (S03), capture the `SQLite Backup API` snapshot (S17) and sealed journal archive (S1004), and verify that the PMS console holds the bay unreconciled until supervisor and quality sign-off completes. [S03] [S17] [S25] [S1004]

Trust boundaries, cached credentials, coupled time-and-certificate recovery, verified boot slots, signed offline manifests, sealed partition-bounded journals and OT-safe containment give operators a controlled alternative to halting production or abandoning security when the cloud is unavailable. [S03] [S04] [S22] [S25] [S27]

*A factory does not become secure because the cable to the outside world goes dark; it stays trustworthy when every local door, clock, update and ledger can prove its authority in the dark.*

#### CHAPTER SOURCE NOTES

S03 NIST. 2023.
S23 IETF / RFC Editor. 2026.
S27 IETF / RFC Editor. 2020.
S04 NIST. 2018.
S22 systemd project. 2026-10-30.
S1002 IETF / RFC Editor. 2021-04.
S25 NIST. 2025.
S26 NIST. 2024-02-26.

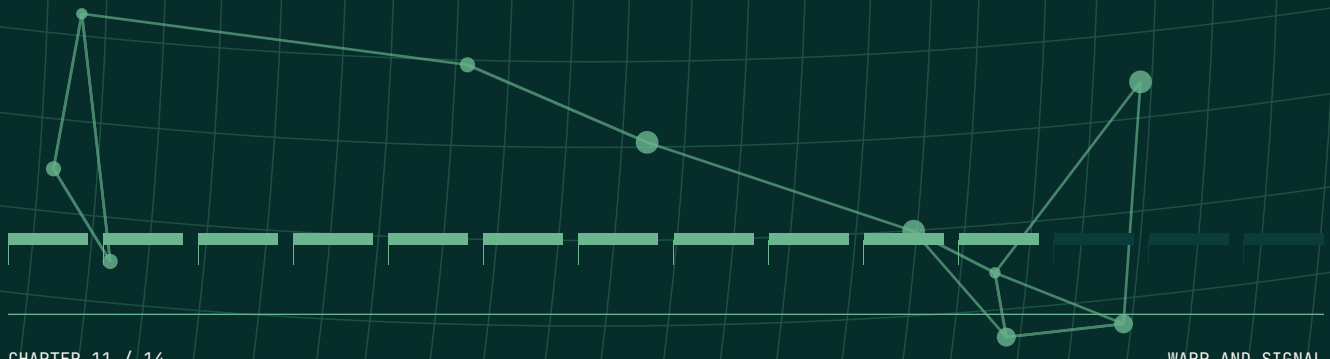
S05 NIST. 2010.
S13 SQLite. 2026-10-30.
S30 chrony project. 2026-10-30.
S1001 IETF / RFC Editor. 2008-05.
S06 NIST. 2016.
S31 IETF / RFC Editor. 2024.
S1005 systemd project. 2026-10-30.
S1003 IETF / RFC Editor. 2022-11.
S12 IETF / RFC Editor. 2022.
S10 IETF / RFC Editor. 2010.
S15 SQLite. 2026-10-30.
S16 SQLite. 2026-10-30.
S29 Linux kernel project. 2026-10-30.
S17 SQLite. 2026-10-30.
S1004 systemd project. 2026-10-30.
S21 systemd project. 2026-10-30.
S18 IETF / RFC Editor. 2014.
S14 SQLite. 2026-10-30.
S28 SQLite. 2026-10-30.
S20 Eclipse Mosquitto. 2026-10-30.

11

PART III: RECOVER A TRUSTWORTHY OPERATION / RESTORE DELIBERATELY

# Recovery finishes when the record and the process agree

*Boot order, consistent backups, authorised process release and ledger reconciliation decide whether restoration is complete.*



CHAPTER 11
DECISION TEST

Boot order, consistent backups, authorised process release and ledger reconciliation decide whether restoration is complete.

Recovery is complete only when the restored services, the physical process and the reconciled ledger describe the same authorised state. Returned mains, running processes and an empty transmission queue are useful observations, but none proves that the factory has recovered. The database may predate the last local commits, a replacement gateway may replay work already accepted elsewhere, and an inspection screen may be available before its quality disposition is authorised. Green Future Tech Limited (GFT) therefore treats restoration as a controlled production decision rather than a successful boot.

The practical answer is a recovery hold with explicit exits. Restore stable supply and storage first, establish clock and local identity state, then start local infrastructure, durable capture and inspection in dependency order. Keep process release separate from service supervision. Restore backups into isolation, reconcile automated and manual records before declaring the dashboard authoritative, and replay through idempotent consumer transactions. NIST contingency and recovery guidance supports planned priorities, exercised playbooks and improvement; the detailed sequence below is GFT's proposed application to a weak-grid factory, not a report of a completed deployment. [S05] [S06]

This boundary matters because a recovery can make an incomplete record look current. A local dashboard populated from an old backup will still respond to queries. A remote total can rise while duplicate deliveries conceal a missing sequence. A supervisor under pressure to restart may reasonably read either signal as closure. The recovery record must instead show what was restored, what remained uncertain, who resolved each exception and which authority released production.

11.1

## Power return should start a recovery sequence

Power restoration is the beginning of a dependency check, not permission for every device and process to resume at once. The site electrical engineer and equipment procedures determine when the supply is stable enough for the relevant loads. The digital recovery runbook should consume that decision rather than invent its own voltage or transfer settings. It should also distinguish a protected record-keeping circuit from production machinery: the gateway may be ready while the line remains stopped, and machinery may restart under its own control procedure while the production record remains in recovery hold. [S03] [S05]

GFT's proposed sequence begins with supply confirmation and a check that storage devices, filesystems and free space are fit for writes. The next gate establishes clock quality and local identity. Capture then starts against a known schema and a known source map. The broker and local console follow their declared dependencies, while inspection remains unavailable until its complete input, model and disposition path has passed its own readiness check. Remote replay waits until the local system can preserve new work and retain evidence of what it sends. This ordering prevents a restored uplink from becoming an unauthorised drain on a fragile restart.

The dependency graph must also cover a second interruption during recovery. If supply fails while a schema check, restore or replay is underway, the next boot must be able to identify the interrupted stage without treating its partial output as complete.

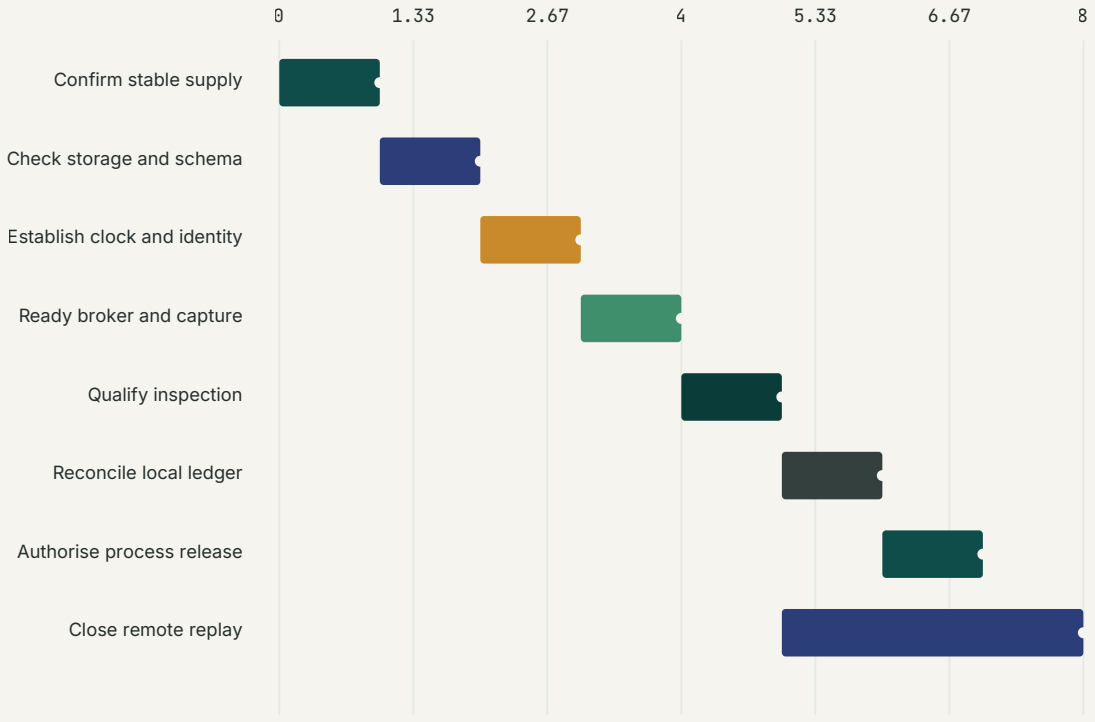
GFT recommends a small durable recovery-state record, written at each gate transition, containing the incident, restore candidate and completed checks. It should not contain a bare completed flag whose transaction boundary is unclear. A stage is complete only when its evidence and its state marker share the same durable commit or another explicitly tested persistence boundary. [S13]

This state record lets the restarted runbook resume at a known gate or deliberately repeat a non-idempotent check. It also prevents an operator from inheriting a misleading ready screen from before the second failure. Where a stage cannot safely resume, the rule should be conservative and specific: repeat storage validation, reopen the isolated candidate, or return the station to manual capture. The aim is bounded repetition of known work, not an automatic march through dependencies whose earlier results may no longer hold.

The dependency positions in **Exhibit 43**, p. 173 are deliberately not durations. Actual timing belongs to commissioning evidence and later operating records. A common restart command would erase precisely the distinctions the runbook needs: a storage check can fail while the network is healthy, and the broker can be healthy while capture has an incompatible schema. [S05] [S21]

EXHIBIT 43 / ILLUSTRATIVE DEPENDENCY POSITION, NOT ELAPSED TIME

## Service restoration follows dependencies rather than a common restart command




ILLUSTRATIVE [S03] [S05] [S21] Illustrative recovery order proposed by GFT. Equal numeric positions show precedence only and carry no measured or promised durations. Remote replay may transmit after local capture is ready, but its closure remains subject to reconciliation and release evidence.

Recovery hold must remain useful while these gates are pending. The local console should name the failed dependency in operational language, retain the last known good state and offer the authorised manual capture route. A vague starting screen

invites repeated reboots and informal workarounds. By contrast, a hold state that says storage verification pending, capture unavailable and manual batch book authorised gives the supervisor a safe action and preserves the same bundle identifiers for later reconciliation.

The runbook should also define what may happen automatically. Starting a read-only diagnostic service can be automatic. Accepting production events requires a readiness decision. Releasing held material or restarting a physical process requires the role assigned by the plant's operating and safety procedures. NIST operational technology guidance makes performance, reliability and safety part of the control problem; GFT's separation of service restoration from process authorisation follows that boundary. [S03]

11.2

## Service supervision needs bounded retries

A supervised process is not necessarily a ready service. A capture daemon can remain alive while its database is read-only, its schema migration has failed, its local credentials have expired or its event source has not completed initialisation. Restarting that process repeatedly may produce a convincing stream of log entries without restoring the promised function. Readiness must therefore test the duty, including durable recording, rather than infer it from a process identifier.

systemd provides a useful implementation distinction. With `Type=notify`, a service can send `READY=1` after finishing startup, and follow-up units wait for that notification. GFT recommends emitting readiness only after the application has checked its effective schema, source registry, storage mode, clock-quality state and required local credentials. The check should use a dedicated health record or transaction path that cannot be mistaken for factory output. A successful socket connection alone is not readiness, and a synthetic production event should never be left in the ledger as proof of boot. [S21]

Retries need a failure budget just as queues need a storage budget. The service manual exposes restart delays and start-rate limiting, so a site can version and test its chosen behaviour instead of accepting an unreviewed loop. GFT does not prescribe a universal retry count or delay. Those values depend on reserve energy, storage wear, the time needed for storage checks and the operator's manual fallback. The important rule is that repeated failure must terminate in a visible hold with diagnostics retained, not continue until the protected supply is exhausted. [S21] [S22]

A watchdog deserves the same discipline. It should represent a meaningful application obligation, such as completing a health transaction and servicing the capture loop, rather than merely confirming that a thread can send a heartbeat. If the database is blocked, the watchdog action should preserve evidence and enter the declared recovery state. Killing and restarting the process without retaining the last failure reason destroys the information needed to distinguish a transient dependency from a persistent defect. [S21] [S22]

> **GFT NOTE**
>
> Version the service definition, readiness probe, restart limit and log-retention setting together. During acceptance, force a dependency to remain unavailable and verify that the system stops retrying at the approved boundary, retains the failure evidence and offers manual capture. Do not accept a demonstration in which the service eventually starts after an unknown number of reboots.

Readiness should include the complete write path but avoid contaminating the production record. A dedicated health transaction can verify that the database accepts writes, the configured journal behaves as expected and free space remains above the operator's approved floor. Its identifier and schema must keep it outside production totals and remote replay. The probe should also confirm that the real capture path can construct a valid event from the current source registry. Testing only the health table would miss a broken source mapping; inserting an unmarked test bundle would create false work. The distinction belongs in application code and acceptance evidence, not in an operator's memory. [S13] [S21]

The local console should report four different conditions: process stopped, process starting, process running but not ready, and service ready. This vocabulary lets the supervisor distinguish a supervisor action from an application defect. It also creates clean timing boundaries for the recovery record, from supply confirmation to first ready capture and from first ready capture to authorised process release.

11.3

## A backup must be consistent before it can be useful

A backup is a recovery input, not a recovery result. SQLite's backup documentation warns that a power or operating-system failure during an external live-file copy can leave the backup corrupted after system recovery. Its Online Backup API instead copies database contents through a defined interface and can produce a destination snapshot while other database users continue. The design lesson is not that every factory must use one product. It is that the backup method must define consistency under concurrent writes and interruption. [S17]

The useful unit of restoration is larger than the database file. GFT specifies a restore set containing the consistent database snapshot, schema and migration version, application configuration, source registry, outbox cursors, retained remote receipts, manual-batch register, approved software manifest and integrity metadata. Secret material should be restored through the authorised credential process rather than copied into an ordinary backup archive. The manifest must identify which components were captured together and which represent different points in time.

That distinction prevents false completeness. A database snapshot may be internally consistent while its configuration refers to a newer schema. An outbox cursor may name a remote receipt that is not present in the archive. A source registry may describe a device replaced after the snapshot. None of these faults requires a corrupt

database, yet each can corrupt the operating interpretation. Atomic commit defines the transaction boundary within the database; it does not make an uncoordinated collection of files a consistent factory restore set. [S13]

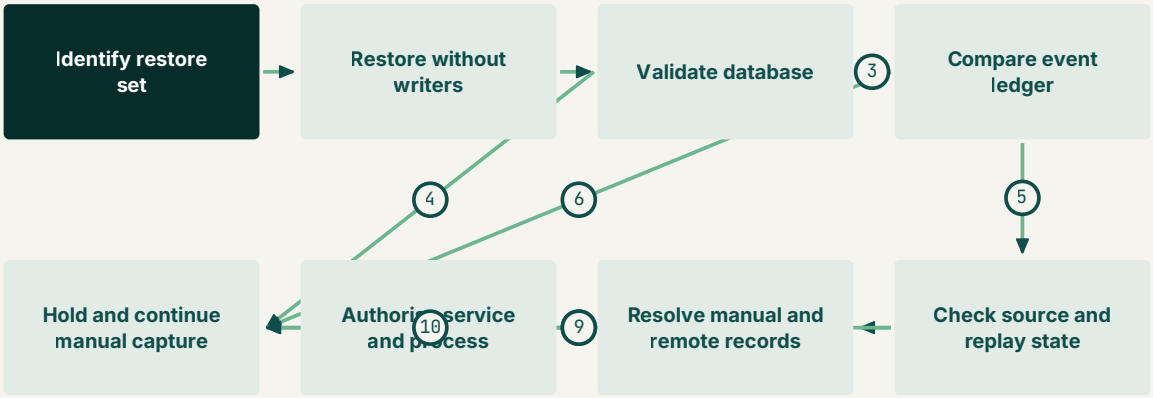
### 11.4 ISOLATION PRESERVES THE RIGHT TO SAY NO

Restore into an isolated environment with no production writer, no outbound replay and no authority to release material. Preserve the original failed media or its evidence image where the incident procedure requires it. Validate the database using the selected engine's integrity facilities, then validate the application ledger: expected source sequences, event identifiers, payload hashes, outbox references, manual-batch links and configuration compatibility. A successful database check answers one necessary question, not the whole production question.

**Exhibit 44**, p. 176 makes the decision sequence explicit. Failure returns the station to controlled manual capture and evidence preservation. It does not trigger a succession of different backups until one happens to open, because each attempted restore can change the evidence and obscure which version produced which record. NIST recovery guidance supports playbook-driven recovery and testing; GFT applies that discipline by naming the restore candidate and retaining every validation result. [S06] [S17]

EXHIBIT 44 /

## A restored backup must pass ledger checks before release




3. Pass
4. Fail
5. No unexplained gap
6. Unresolved conflict
8. Ambiguous authority
9. Evidence complete
10. Keep provisional

ILLUSTRATIVE [S06] [S13] [S17] Illustrative GFT restore procedure. Database validation and production-ledger reconciliation are separate gates. No live writer or outbound replay is permitted while the candidate remains isolated.

After the checks pass, the team should still treat the restored ledger as a historical claim to reconcile, not as the automatic truth. Later events may exist on the failed device, on manual forms or at the remote consumer. The isolated candidate supplies a known starting point against which those sources can be compared. Restoration becomes recoverable because disagreements are visible before new production is mixed into the record.

> **METHOD**
>
> Give every restore set a manifest with a unique candidate identifier, creation condition, source-device identity, software and schema versions, database consistency method, included replay state, integrity result and retention owner. During restoration, record the actual candidate used, every validation outcome and every added event source. This chain lets an auditor distinguish an event recovered from the backup, a tail recovered from local media, a manual correction and a new post-recovery observation.

The newest backup is not automatically the best candidate. It may have been created during the failure, after a configuration error or from a source whose writes were already suspect. Selection should compare the candidate's consistency evidence, recovery-point requirement, software compatibility and distance from the incident. An older validated snapshot plus a recoverable local tail can be more useful than a newer file with an uncertain write boundary. Conversely, choosing an old snapshot merely because it opens can impose needless reconciliation and lose records that a valid later candidate retains.

This selection decision belongs in the recovery certificate even when the first candidate succeeds. It explains why the team accepted a particular starting point and which later sources must still be examined. Where several candidates are compared, each receives a separate identifier and isolated validation record. The team must never merge convenient pieces from different snapshots without recording the transaction and identity rules that make the combined ledger legitimate.

11.5

## Replacement equipment needs a tested identity handover

A cold spare restores capacity only when it can assume a defined operational role without creating a second authority for the same events. Hardware fit is therefore necessary but insufficient. The replacement needs an approved boot and firmware state, compatible interfaces, a supported storage path, usable offline licences and a configuration matched to the station. NIST platform-firmware guidance provides a basis for authenticated recovery to a known-good state; it does not establish that an arbitrary imaged spare is ready for a particular factory role. [S04]

GFT separates three identities during substitution. The logical `source_id` identifies the station's event stream only when the replacement is authorised to continue that stream. A device identity identifies the physical or virtual appliance. A fresh boot identity distinguishes the new execution interval. A handover record links all three,

names the replaced device, states the fencing evidence and records the configuration and restore-set candidates. This avoids both extremes: treating every spare as a new production source, or cloning an old identity so completely that two devices can write as one.

UUIDs can be generated without a central registration process, which supports local identifier creation during disconnection. That property does not decide source authority. The plant still needs a controlled handover and sound identifier generation. If the original device may remain connected, if its last sequence is unknown, or if the replacement carries a different source map, the affected events stay in reconciliation hold until the ambiguity is resolved. [S31]

The checklist in **Exhibit 45**, p. 178 belongs in the spare kit and the recovery runbook. It is a proposed readiness record, not an assertion that GFT holds a spare inventory. A box on a shelf has recovery value only when the local team can prove these handover steps on the actual target equipment.

EXHIBIT 45 /

## A spare is useful only if its handover is rehearsed

| Handover control | Evidence before substitution | Release condition |
|---|---|---|
| Approved image | Firmware, operating system and application manifest | Integrity and compatibility checks pass |
| Physical fit | Ports, storage, power path and mounting inspection | Station dependencies operate in local mode |
| Credentials | Authorised provisioning record and trust-store version | Required local roles work without external access |
| Source authority | Original fenced or retirement evidence | Only one writer owns the logical source |
| Boot identity | New device and boot handover event | Restart cannot be mistaken for duplicate production |
| Restore rehearsal | Named backup candidate and isolated validation log | Ledger and replay state reconcile |
| Rollback route | Original-state record and manual fallback batch | Failed handover preserves capture and evidence |

ILLUSTRATIVE [S04] [S13] [S17] [S31] Illustrative GFT spare-readiness checklist. Evidence and release conditions are proposed controls for a future site, not records of existing inventory or completed commissioning.

The handover rehearsal should include an imperfect case, not only a clean swap. For example, the runbook can require the technician to respond when the original device's final local sequence is unavailable. The correct result is not a confident guess. It is a bounded hold, preservation of both evidence sets and a reconciliation decision based on event identifiers, remote receipts and physical production records.

11.6

# Replay after restore can reintroduce old events

A restored outbox describes what the backup knew, not everything the remote system later accepted. Suppose the local receipt cursor in the backup predates the last usable connection. The remote consumer may have committed later events and issued receipts that reached only the failed device. Replaying the complete restored outbox would then resend old work. Deleting every pending item because the backup looks old would create the opposite failure: events accepted locally but never committed remotely would disappear.

The safe unit of replay is an identified event with a durable consumer decision. HTTP semantics define an idempotent request as one where multiple identical requests have the same intended server effect as a single request, while also noting that a server may retain separate logs or other side effects. GFT's factory contract therefore needs more than a retryable method. The consumer should deduplicate on the agreed source and event identifiers, compare the payload hash and commit the production event and its receipt in one transaction. A repeated identifier with a different payload is a conflict for review, not a duplicate to suppress silently. [S12] [S13]

Before replay, compare three ledgers: the restored local committed set, the local pending set with its retained receipts, and the remote committed set when reachable. The comparison should produce accepted, already-remote, missing-local, conflicting and unresolved categories. If the remote ledger cannot be queried, local capture may continue under the approved service class, but remote reconciliation remains open. An old backup cannot independently prove which later business actions occurred elsewhere.

This is why an empty outbox is a weak completion signal. The queue may be empty because every event received a durable remote receipt, because a retention rule deleted it, or because restoration loaded a state in which pending work was never represented. Only the reconciled identifier sets can distinguish those outcomes. Replay telemetry should count transmissions, new remote commits, harmless duplicates and payload conflicts separately so network activity cannot be mistaken for accounting closure.

WORKED EXAMPLE

**An illustrative restored backup needs local-tail, remote and manual reconciliation.**

Consider an illustrative exercise in which the expected automated stream ends at event E1200. The consistent backup contains events through E1160, while the failed device's preserved tail supplies the following 40 events, E1161 to E1200. The restored automated ledger therefore contains 1,160 + 40 = 1,200 identified events. These are constructed test-ledger values, not observed factory production or a recovery-time claim. The backup and tail methods follow the consistency and identifier boundaries described above. [S13] [S17] [S31]

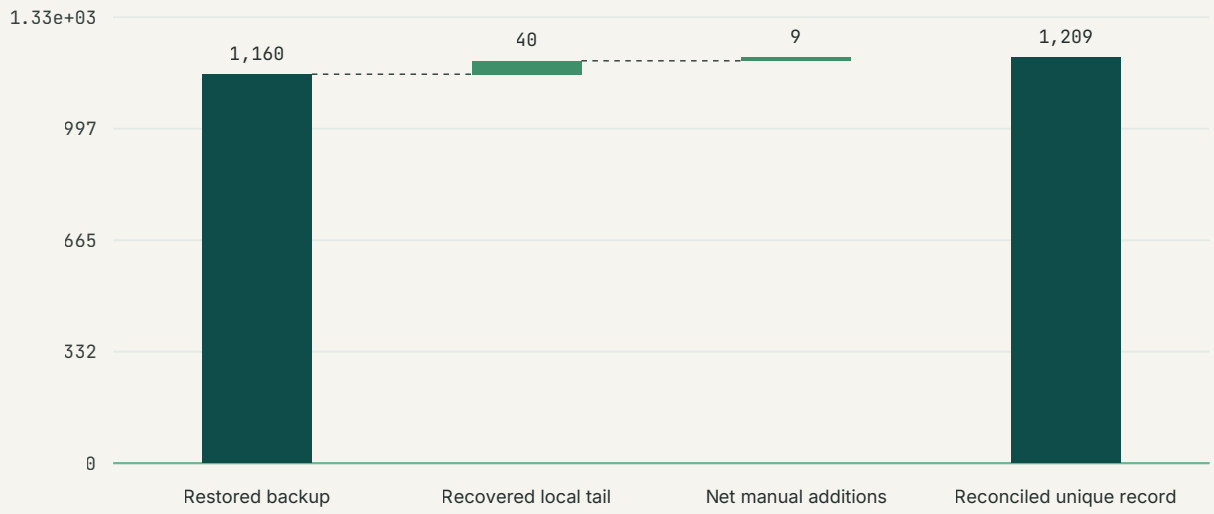
The last remote receipt retained in the backup reaches E1140, but a query of the remote consumer proves that it has actually committed through E1150. The restored pending set contains 60 candidates, E1141 to E1200. Replay therefore produces 10 harmless duplicate deliveries for E1141 to E1150 and 50 new remote commits for E1151 to E1200. The remote automated total becomes 1,200, while the transmission log records 60 attempts. Treating all attempts as new work would overstate the ledger by 10; deleting all candidates at the restored receipt cursor would have lost the 50 genuinely pending events. All values in this paragraph remain illustrative arithmetic.

The supervisor also supplies 12 manual fallback entries. Reconciliation finds that 3 carry identifiers already present in the automated tail, leaving 12 - 3 = 9 additional production records. The final unique total is consequently 1,200 + 9 = 1,209 events. The restored backup alone held 1,160, so it was 49 records short of the reconciled result even though its database was consistent. The worked case closes only because local and remote committed sets both contain the same 1,209 identifiers, all duplicate decisions are explained and no sequence remains unresolved. [S12] [S13]

**Exhibit 46**, p. 181 shows the unique-record arithmetic in the worked example. It does not count replay attempts as production, and it does not treat the backup's internal consistency as proof of completeness.

EXHIBIT 46 / ILLUSTRATIVE UNIQUE PRODUCTION EVENTS

## Reconciliation adds recovered work rather than replay volume




DERIVED [S12] [S13] [S17] [S31] Derived from the chapter's illustrative worked example: restored backup 1,160 events plus recovered local tail 40 plus 9 net manual additions equals 1,209 unique events. The 12 manual entries include 3 duplicates, so only 9 are added. Separately, 60 replay candidates produce 10 duplicate deliveries and 50 new remote commits; transmission count is not added to the production total.

The arithmetic is simple, but the evidence obligations are strict. Every added tail event needs its identifier and integrity check. Every remote claim needs a committed-identity response, not a successful connection. Every manual duplicate needs a link to the automated event it duplicates and a retained reason. The final totals agree only after those records agree at event level.

11.7

# Operator release requires an understandable recovery record

The person releasing work needs a decision record, not a collection of daemon logs. GFT proposes a recovery certificate in ordinary operational language, with links to the underlying evidence. It should identify the incident, affected stations, restore-set candidate, software and schema versions, source and boot handovers, database result, clock-quality state, inspection readiness and manual batches still open. It should then reconcile expected, locally committed, remotely committed, duplicated and unresolved identifiers by record class.

The certificate must make uncertainty visible. A missing sequence is not zero production. An unchecked remote total is not zero remote work. A manual form awaiting comparison is not an error to hide, and an unresolved identifier conflict is not a harmless duplicate. The dashboard should remain provisional while any required category is open. This protects later performance analysis from treating an incomplete recovery as a measured shift result.

Two signatures answer different questions. The technical recovery owner confirms that the specified services are ready, the restore candidate passed validation and the replay evidence is complete to the stated boundary. The authorised operations or quality role confirms that the physical process and disposition rules may resume. NIST operational technology guidance's attention to reliability and safety supports keeping that process authority with the plant's designated role rather than with a service manager or remote dashboard. [S03]

A release can also be partial. The site may authorise local durable capture while keeping automated inspection unavailable, or permit a restricted production mode while remote history remains provisional. The service classes established earlier determine whether that is allowed. The release record must name the restriction, its display state, the permitted work and the condition for full closure. A single recovered flag is too coarse for these decisions.

> **WARNING**
>
> Do not let a green process indicator release held material, clear a quality hold or close an incident. Those actions require the authority and evidence assigned by the operating procedure. If inspection, source identity or record reconciliation remains uncertain, show the hold and keep the authorised fallback active. [S03] [S06]

Corrections after release must preserve the original trail. If a later manual form resolves a missing bundle, append a correction or linkage event with its evidence and reviewer. Do not rewrite the earlier recovery certificate or silently change the original event timestamp. The certificate then remains an accurate account of what was known at release, while the ledger retains how the authorised understanding changed.

The record should remain reviewable after the pressure of the restart has passed. A supervisor should be able to move from each summary total to the identifiers included in it, then to the source evidence and reconciliation decision. An auditor should be able to distinguish what was known before release from what arrived afterwards. These views need not expose every technical detail on the shop-floor screen, but the export must preserve them with stable identifiers and readable reasons. A PDF signature page without the underlying event sets is not enough, and an event dump without the human release decision is equally incomplete.

Closure therefore has a retention requirement as well as a technical and operational gate. The site should retain the certificate, restore-set manifest, validation outputs, reconciliation exceptions and final committed-ID comparison under its approved evidence policy. The next incident can then start from a known recovery history rather than rediscovering which assumptions were already disproved.

11.8

## Recovery capability should improve after each exercise

A recovery exercise should change the system, not merely confirm that a team can eventually make it work. NIST SP 800-184 explicitly connects recovery planning with playbook development, testing and improvement. GFT turns each rehearsal or incident into a versioned recovery record: the initial fault, dependency order actually followed, failed assumptions, restore candidate, spare fit, manual workload, reconciliation exceptions, authorisations and corrective actions. [S06]

Measure the intervals between agreed markers rather than reporting one ambiguous restoration time. Supply confirmation, storage validation, capture readiness, inspection readiness, process release and remote ledger closure answer different operational questions. Their durations may overlap, so the runbook should show the sequence and the waiting cause. An operator-approved target can then be compared with observed evidence without converting a laboratory exercise into a claim of production availability.

The review should separate defects in the design from defects in the exercise. A missing credential is a recovery dependency defect. An expired test licence may be an exercise-preparation defect, but it still reveals a real restoration risk if production would use the same licence path. A spare with the wrong connector is not a minor administrative issue when the plan depends on substitution. Each finding needs an owner, a changed artefact and a retest condition.

The management cadence from GFT's Cameroon textile programme provides a useful place for this learning loop: review the recovery exceptions, assign the corrective action and retain the evidence alongside other production-management routines. This is a recommendation for programme implementation, not evidence that a commissioned recovery system already operates. Detailed acceptance experiments belong to the commissioning chapter; this chapter requires their results to update the restore set, runbook and release criteria.

Improvement is complete only when the changed procedure survives another bounded exercise. A revised document alone does not show that the spare now boots, the backup now includes replay state or the dashboard now exposes an unresolved sequence. The next exercise should challenge the earlier failure and verify that its correction did not weaken another dependency. Recovery capability becomes credible through that repeated chain of observed failure, explicit repair and retained proof.

*Recovery ends when the process is authorised, the ledger is reconciled and both tell the same production story.*

**CHAPTER SOURCE NOTES**

S05 NIST. 2010.

S06 NIST. 2016.

S03 NIST. 2023.

S13 SQLite. 2026-10-30.

S21 systemd project. 2026-10-30.

S22 systemd project. 2026-10-30.

S17 SQLite. 2026-10-30.

S04 NIST. 2018.

S31 IETF / RFC Editor. 2024.

S12 IETF / RFC Editor. 2022.

GREEN FUTURE TECH

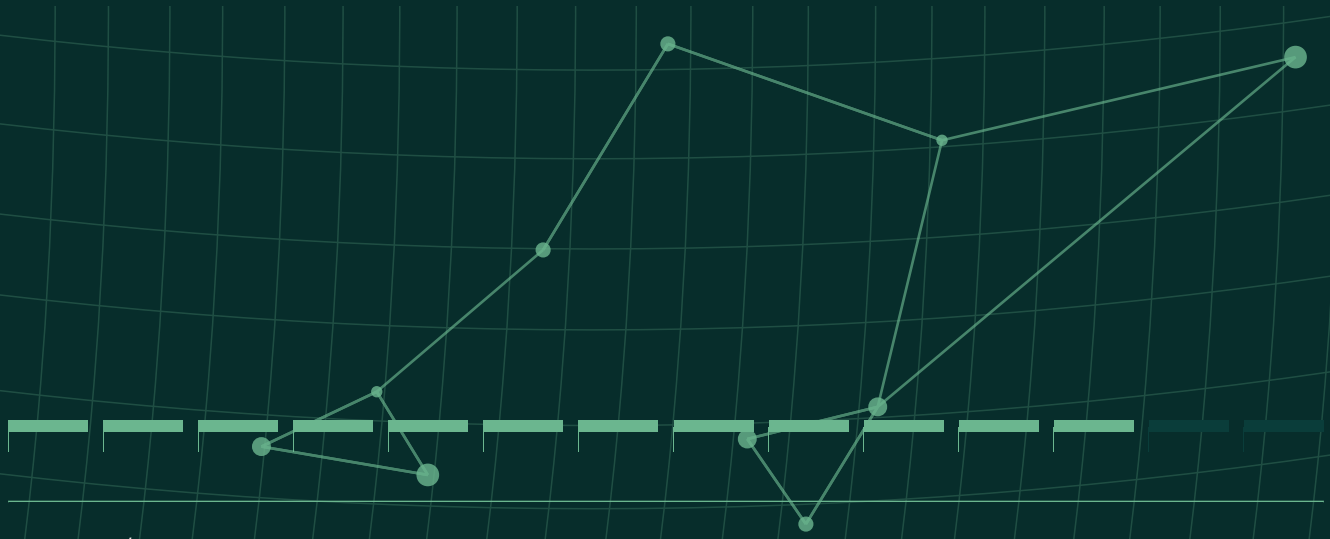
LIMITED / HONG KONG SAR

12

PART IV: PROVE AND GOVERN CONTINUITY / MAKE THE INVESTMENT TESTABLE

# The economic case should price avoidable loss and restoration

*Compare service classes using transparent operating assumptions and procurement obligations rather than promised resilience percentages.*



CHAPTER 12
DECISION TEST

Compare service classes using transparent operating assumptions and procurement obligations rather than promised resilience percentages.

An investment in digital continuity on an intermittent grid is sound only when the cost of protecting the local operating boundary and restoring a verified production record is lower than the specific operational losses that the protected service can actually prevent. [S03] [S05] A factory sponsor in Cameroon evaluating uninterruptible power for edge computing, local write-ahead storage, store-and-forward buffering, and rehearsed spare-device recovery should reject return-on-investment calculators that credit a software project with avoiding full plant downtime. [S05] [S1203] When the main distribution board trips and spinning frames, looms, dyeing vats, or sewing motors lose prime power, a battery-backed gateway and a durable local database do not keep unpowered machinery turning. [S03] What a protected digital boundary prevents is a narrower and measurable class of loss: the erasure of committed bundle histories, the quarantine of unverified fabric rolls, the clerical labour of transcribing paper dockets, the supervisory hours spent untangling duplicated or out-of-order counts, and the post-restoration delay while lines wait on an unverified ledger before resuming work. [S05] [S13] Pricing that boundary honestly, valuing prevented output loss at contribution margin rather than gross revenue, exposing multi-year ownership costs, and converting vendor promises into witnessed acceptance clauses turns resilience into an inspectable engineering contract. [S04] [S06] [S1201]

This chapter translates the technical mechanisms established across the preceding chapters into an economic and procurement decision method for GFT's Cameroon textile programme. [S05] [S1203] The inherited sector baseline shows that roughly 95 percent of Cameroon's cotton fibre is exported raw while about 5 percent is processed locally, against a national target of 50 percent local processing by 2030. [R06] [R07] Moving along that curve requires factories whose operating margins are not eroded by either unrecorded floor losses or oversized, unmaintainable information-technology spending. [R12] [S05] None of the calculations in this chapter represent observed financial returns from a commissioned GFT plant, and nothing in this paper constitutes investment, legal, or tax advice. [S26] Instead, the framework separates physical downtime from digital record loss, defines a counterfactual manual baseline, sets out a procurement cost register free of invented vendor prices, tests break-even sensitivity across correlated failure assumptions, and establishes contractual exit rights so the plant retains its production history when a supplier contract ends. [S1201] [S1203]

12.1

## Start with the loss the protected service can actually avoid

A disciplined business impact analysis begins by drawing a firm boundary between losses caused by the absence of motive power and losses caused by the failure of the information and control layer. [S03] [S05] [S1203] In a cotton-to-garment operation, a mains failure that is not bridged by standby generation halts high-kilowatt mechanical loads immediately: carding and ring-spinning frames stop, looms arrest, dye circulation pumps require thermal management, and sewing tables come to rest. [S03] As shown in Chapter 4, an illustrative 150 W protected digital circuit with 520.833 Wh of nominal battery energy is sized to carry an edge computer, message broker, local network switch, router, time source, and supervisor console for 2 hours, not to drive

production motors or high-intensity inspection lamps. [S03] [S05] Multiplying total factory hourly output by the duration of a utility outage and claiming the resulting sum as the benefit of an edge computing package is a boundary error. [S05] [S1203]

### 12.2 SEPARATE PHYSICAL STOPPAGES FROM INFORMATION-BOUNDARY LOSSES

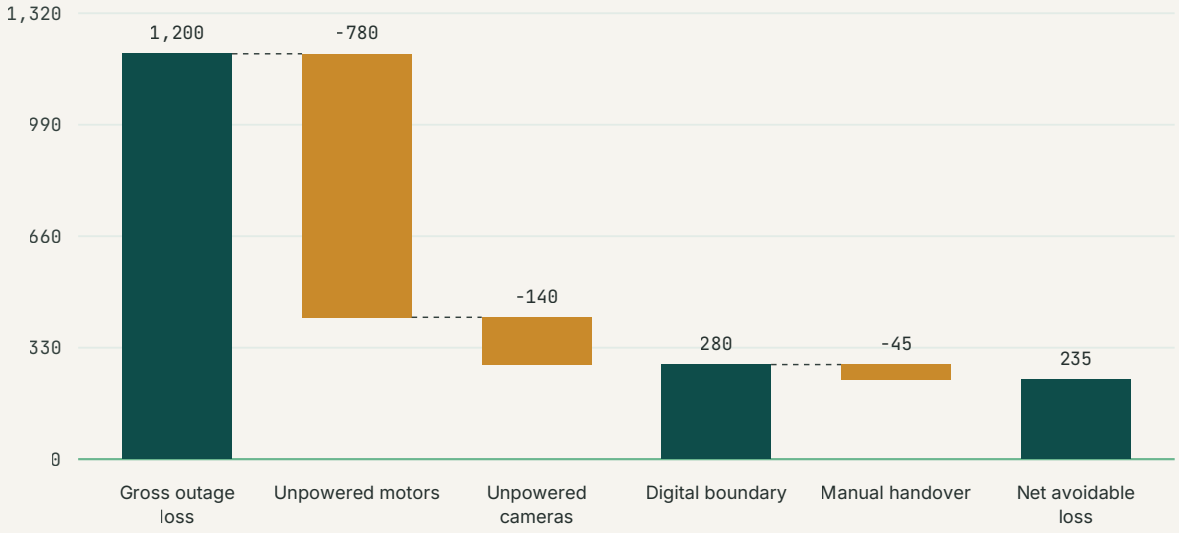
To price continuity accurately, engineering and finance teams must decompose a site disruption into five operational components and retain only those that change when the digital service class changes: [S03] [S05] [S1203]

1. **Idle production machinery during mains absence:** When prime power fails and standby generation is unavailable or withheld due to fuel economics, physical transformation ceases. [S03] The digital continuity architecture cannot avoid this stoppage, so its evaluation must assign zero benefit to unpowered machine hours (wider generation and solar-storage economics belong in wp06). [S03] [S05]
2. **Post-restoration hold awaiting digital readiness or ledger verification:** When mains power returns and mechanical lines are ready to run, production can remain held because an edge gateway failed to boot without external connectivity, a database suffered unrecovered corruption, a cached model licence expired offline, or supervisors cannot confirm which garment bundles passed inspection before the blackout. [S03] [S06] [S15] Because a local-first architecture with bounded recovery eliminates or shortens this post-mains hold, the avoided delay is attributable to the digital continuity investment. [S05] [S06]
3. **Quarantine, re-inspection, and downgrade from lost dispositions:** If an abrupt power cut rolls back committed inspection records (for example, under a write-ahead log set to `synchronous=NORMAL` rather than `FULL`) or drops in-flight quality dispositions before durable storage, quality assurance must hold the affected rolls or bundles. [S13] [S16] Re-inspecting finished bundles or downgrading unverified export lots carries a direct cost that atomic local capture prevents. [S13] [S1203]
4. **Clerical capture and manual reconciliation burden:** When electronic capture fails, operators and line clerks switch to pre-printed paper dockets, record bundle identifiers and defect codes by hand, and later key those sheets into the performance-management system (PMS) while resolving sequence gaps and duplicates. [S05] [S31] The incremental labour hours required to execute and audit that manual bridge represent a direct cost of operating without automated local survival. [S05] [S1203]
5. **Delayed management reporting and upstream scheduling lag:** When the wide-area backhaul drops while the local network stays alive, shift huddles continue on local consoles while central dashboards wait until the store-and-forward spool drains. [S08] [S18] Unless stale central visibility triggers a verified contractual penalty, emergency freight premium, or stockout within the offline window, delayed cloud reporting is a convenience degradation rather than a cash loss. [S05] [S18]

The waterfall in **Exhibit 47**, p. 188 illustrates how this boundary discipline filters an illustrative outage event before any investment comparison takes place. [S03] [S05] [S1203]

EXHIBIT 47 / ILLUSTRATIVE THOUSAND XAF PER OUTAGE EVENT

## Only attributable losses belong in the benefit case




ILLUSTRATIVE [S03] [S05] [S1203] Illustrative accounting decomposition for a single hypothetical outage event. Total gross disruption (1,200 thousand XAF) is reduced by unpowered machinery downtime (-780 thousand XAF) and unpowered optical inspection stop (-140 thousand XAF) to isolate the digital-boundary exposure (280 thousand XAF), then reduced by unavoidable changeover administration (-45 thousand XAF) to reach net avoidable record and restoration loss (235 thousand XAF). Values are illustrative method inputs, not observed GFT plant figures.

In **Exhibit 47**, p. 188, an illustrative gross outage disruption of 1,200 thousand XAF across cutting, sewing, and finishing is stripped of 780 thousand XAF in unpowered motor downtime and 140 thousand XAF in unpowered camera-bench stoppage, leaving 280 thousand XAF of exposure inside the digital boundary. [S03] [S05] Subtracting an illustrative 45 thousand XAF of unavoidable floor supervision during the power transition leaves 235 thousand XAF per qualifying event as the net avoidable loss ceiling. [S05] [S1203] What **Exhibit 47**, p. 188 cannot establish is the actual monetary loss at a specific Cameroonian mill, which requires observed site outage durations and verified accounting records. [S02] [S05]

### 12.3 DEFINE THE COUNTERFACTUAL MANUAL PROCESS HONESTLY

Every digital continuity proposal competes against a concrete operational alternative: running the factory with a structured manual paper fallback and delayed batch entry. [S05] In GFT's Cameroon textile programme, phase one of the PMS begins with structured point-of-work logs before automated edge capture and vision inspection are scaled up. [S03] [S05] During an outage, operators write the 128-bit universally unique identifier (UUID) or pre-printed barcode sequence of each bundle, the operation code, the shift identifier, and the piece count onto a traveller card, and a clerk later keys those cards into the local ledger and reconciles discrepancies against supervisor tallies. [S13] [S31]

Pricing the counterfactual manual process requires measuring four parameters during a floor drill: extra seconds per bundle to write paper dockets, the error rate of hand-written entries under shift pressure, clerical minutes per hundred bundles to key and reconcile the backlog after power returns, and batch-hold time while shipments await verified bundle genealogy. [S05] [S13] [S31] In Cameroon, statutory labour floors provide

context: the national minimum wage (SMIG) rose from 41,875 XAF per month in 2023 to 43,969 XAF per month from 23 February 2024 (about US$76, a 5 percent rise), within an economy where the standard working week is 40 hours (48 hours maximum, with overtime at 1.2 times base pay) and cross-sector average salaries sit near US$359 per month. [R14] [R15] [R16] Because ledger reconciliation and quality release consume shift supervisors, quality controllers, and IT technicians rather than entry-level operatives, the counterfactual calculation must use the loaded hourly cost of those supervisory roles during normal or overtime hours, labelled as an estimate or site-accounting input. [R14] [R16] [S05]

12.4

## Use contribution margin when estimating lost output value

When a digital failure halts output after prime power is restored (for example, when a corrupted local database, a failed certificate bootstrap, or an unverified model update blocks an inspection gate for 90 illustrative minutes after mains return), the model must value the lost production volume accurately. [S03] [S05] [S23] Valuing every delayed or lost garment bundle at its full wholesale selling price treats raw cotton lint, yarn, dyestuffs, finishing chemicals, packaging, and piece-rate consumables as if they were destroyed during the hold, whereas an idle line consumes no raw materials. [R06] [S05] [S1203]

### 12.5 PRICE DELAYED AND LOST UNITS AT CONTRIBUTION MARGIN, NOT GROSS REVENUE

Under managerial accounting, the loss from an unrecovered production hour depends on whether deferred output is made up within the scheduling window and what variable costs are saved while the line stands still. [S05] [S1203] Publicly reported sector figures show the weight of persistent underutilisation in Cameroon's textile industry: a July 2023 report on CICAM recorded capacity utilisation at 33.7 percent at Garoua, 48.9 percent at CICAM 1, and 10.83 percent at CICAM 2, noting that underutilisation cost about 2 billion XAF in 2020 alongside restructuring estimates of 30.7 billion to 48.2 billion XAF (and one relaunch plan cited at 70.2 billion XAF). [R12] [R13] Those third-party observations reflect broad industrial and equipment conditions rather than digital outage losses at a GFT facility. [R11] [R12] For a single line evaluating digital continuity, the analyst must distinguish three cases: [S05] [S1203]

- **Deferred output recovered in normal shift hours:** If the line has slack capacity later in the week and clears held bundles during regular paid shifts without missing a delivery window, lost contribution margin is zero; the only cost is temporary buffer handling. [R12] [S05]
- **Deferred output recovered through overtime shifts:** If the plant meets its delivery commitment by running authorised overtime (at or above the statutory 1.2x overtime multiplier in Cameroon), the attributable cost is the incremental overtime labour premium, auxiliary energy, and supervisory overhead required to catch up, not the unit's contribution margin. [R16] [S05]
- **Permanently lost output on a bottleneck line:** Only when the affected workstation is a bottleneck operating at full scheduled utilisation, or when a missed export shipping cutoff cancels the order, does an hour of post-restoration digital hold

forfeit sales volume. [S05] [S1203] Even then, the lost value per unit is the **contribution margin** (selling price less variable fibre, yarn, trim, dye, variable energy, and direct piece-work costs), never the gross selling price. [S1203]

METHOD

**GFT boundary for pricing avoidable outage and restoration loss**

For a candidate continuity investment `k` compared against a manual-fallback baseline `0`, the expected annual avoidable loss reduction `dL_k` across `M` scenarios `m` follows `dL_k = Sum_m [ f_m * ( (H_m0 - H_mk) * r_eff,m * c_cm + (Q_m0 - Q_mk) * c_rework + (W_m0 - W_mk) * w_sup + (D_m0 - D_mk) * c_scrap ) ]`, where each input is bounded at the site level: [S05] [S1203]

- `f_m` is the annual frequency of scenario `m` (such as mains loss followed by restored power with WAN unreachable), estimated from site event logs rather than national survey averages. [S01] [S02]
- `H_m0` and `H_mk` are post-mains-restoration hold hours during which powered machinery awaits digital readiness or ledger clearance under baseline `0` and option `k`; physical mains-off hours are excluded. [S03] [S05]
- `r_eff,m` is the unrecoverable bottleneck output rate (units per hour) that cannot be caught up in normal shifts, and `c_cm` is the unit contribution margin; if recovered via overtime, replace `r_eff,m * c_cm` with the overtime wage and utility premium per catch-up hour. [R16] [S05] [S1203]
- `Q_m0` and `Q_mk` are units placed into quality quarantine due to unverified or rolled-back inspection transactions, multiplied by unit re-inspection cost `c_rework`. [S13] [S16]
- `W_m0` and `W_mk` are clerical and supervisory hours spent recording paper dockets, keying backlogs, and resolving sequence conflicts, multiplied by loaded supervisory rate `w_sup`. [R14] [S31]
- `D_m0` and `D_mk` are units suffering physical downgrade or scrap caused directly by lost process state, valued at net scrap loss `c_scrap`. [S03] [S1203]

**Verification check:** Confirm that `dL_k = 0` whenever `H_m0 = H_mk`, `Q_m0 = Q_mk`, `W_m0 = W_mk`, and `D_m0 = D_mk`, preventing any credit for unpowered mechanical downtime. [S03] [S05]

### 12.6 KEEP SAFETY AND MANDATORY TRACEABILITY AS GATES, NOT TRADE-OFFS

Two requirements must never be traded off inside a payback spreadsheet: machinery safety interlocks and mandatory lot traceability. [S03] [S05] As NIST SP 800-82 Rev. 3 emphasises, operational technology security and continuity architectures must address unique performance, reliability, and safety requirements. [S03] Hardwired emergency stops, thermal cutouts, pressure relief valves, and guard interlocks belong to the physical safety authority and must fail safe independently of any edge computer, database, or network switch. [S03] An economic model must never evaluate whether to fund a safety interlock based on avoided downtime, nor permit an automated restart after power return to bypass physical operator clearance. [S03] [S06]

Similarly, when a buyer contract or customs regime requires complete bundle-to-bale traceability for shipment acceptance, preserving the event ledger is a market-access gate. [R06] [S13] Safety and mandatory compliance define the minimum permissible service class, while contribution-margin and labour-saving calculations determine whether to invest in higher levels of automation, battery autonomy, and hardware redundancy. [S05] [S1203] Capital budgeting, discount rates, tax treatment, and financing structures belong in wp13. [S05]

12.7

## Expose the ownership cost across the selected period

A credible economic evaluation pairs its conservative benefit boundary with a complete total cost of ownership (TCO) ledger across the evaluation horizon (such as an illustrative 36-month or 60-month window). [S04] [S05] In industrial digital projects deployed in hot, dusty, and electrically harsh environments, initial capital expenditure on edge computers and switches often represents less than half of the multi-year cash commitment. [S03] [S05] Suppliers routinely quote attractive unit hardware prices while omitting the recurring costs required to keep that hardware functional, licensed, and recoverable on a weak grid. [S04] [S1201]

### 12.8 ITEMISE LIFECYCLE, REPLACEMENT, AND VERIFICATION COSTS BEFORE SIGNING

To prevent post-award budget surprises, the procurement lead should require every bidder to populate a standardised cost schedule separating one-time acquisition from recurring obligations and condition-triggered replacements: [S04] [S05] [S1201]

- **Protected power path and battery replacement cycles:** Uninterruptible power supplies, converters, surge protection, transfer relays, and enclosure fans degrade faster under deep discharges and high ambient temperatures. [S03] [S05] Even for an illustrative 150 W digital load requiring 520.833 Wh of nominal battery capacity (at 0.90 conversion efficiency, 0.80 usable fraction, and 0.80 end-of-life capacity), the budget must include scheduled battery cell replacement and periodic load testing. [S05]
- **Local storage endurance and spare media:** Continuous write-ahead logging (`WAL`) with `synchronous=FULL` durability, MQTT broker persistence (`mosquitto.db`), systemd journal storage (`journald.conf`), and image spooling generate sustained write wear on solid-state drives. [S14] [S16] [S20] [S22] Because SQLite WAL requires same-host storage and cannot run over a network filesystem, high-endurance industrial flash media and pre-imaged spare drives belong in the bill of materials. [S14] [S15] [S29]
- **Cold-spare hardware and rehearsed handover:** As shown in Chapter 11, a recovery time objective cannot be met on an isolated site if a failed edge computer or switch requires weeks of shipping and customs clearance. [S04] [S06] Holding at least one firmware-verified spare node on site (with documented UUID handover and SQLite online backup tools) is a base-case inventory cost. [S04] [S17] [S31]

- **Backhaul retainers and secondary path fees:** Contracting a primary link alongside a secondary cellular or satellite path for store-and-forward replay (such as clearing an illustrative 0.72 decimal GB 8-hour backlog at 100,000 bytes/s of useful goodput) incurs monthly standby retainers and burst data tariffs whether or not an outage occurs. [S08] [S18] [S24]
- **Software licensing, certificates, and fault commissioning:** Commercial edge runtimes, offline licence tokens, local certificate integration under TLS 1.3 and NTS, firmware recovery signing, and witnessed fault-injection tests (Chapter 13) must be priced explicitly prior to award. [S04] [S06] [S23] [S27] [S28] [S1201]

The register in **Exhibit 48**, p. 192 provides a structured template for capturing these commitments without inventing market prices. [S04] [S05] [S1201]

EXHIBIT 48 / PROCUREMENT REGISTER FIELDS

## Cost ownership should remain explicit after procurement

| Cost category | Budget owner | Cadence | Currency and basis | Replacement or trigger condition |
|---|---|---|---|---|
| Edge compute, switches, and local time source | Factory IT / OT lead | Initial capex + spare pool | Quoted XAF or EUR landed | Hardware failure or firmware end-of-support (S04) |
| Protected digital UPS, transfer path, and cells | Electrical engineering | Initial + cell cycle | Quoted XAF landed | Measured autonomy below 0.80 end-of-life threshold (S05) |
| Industrial SSD media for WAL and broker spool | Factory IT / OT lead | Initial + wear replacement | Quoted XAF landed | SMART wear threshold or write-barrier fault (S14, S29) |
| Primary and secondary WAN backhaul retainers | Network operations | Monthly recurring | Contracted XAF / month | Shared-path audit failure or goodput breach (S18, S24) |
| Edge software, model, and offline licence tokens | Digital programme lead | Annual recurring | Quoted EUR / USD + XAF tax | Version renewal or offline cold-boot failure (S1201) |
| Witnessed fault commissioning and recovery drills | Quality and operations | Acceptance + quarterly | Internal labour + vendor fee | Failed ledger reconciliation or runbook change (S06, S28) |

ILLUSTRATIVE [S04] [S05] [S1201] Illustrative procurement cost-register template for a multi-year continuity evaluation. No market prices or vendor quotations are invented; bidders and site engineers must populate quoted values, currencies (XAF, EUR, or USD), quote dates, and replacement triggers before investment sign-off.

As **Exhibit 48**, p. 192 shows, every line item is tied to a named internal owner, a billing cadence, an explicit currency basis, and an engineering replacement condition. [S04] [S05] [S1201] Because the CFA franc (XAF) is pegged to the euro at 655.957 XAF per euro, quotations in euros carry a stable reference parity against local XAF budgets, whereas US dollar quotations introduce exchange-rate variance that finance must record with the quote date and source. [R32] What **Exhibit 48**, p. 192 cannot supply is

a substitute for live commercial tenders: customs duties, tax treatment, installation rates, and freight to Douala or Garoua must be verified by commercial, tax, and legal advisers when bids arrive. [S26] [S1201]

12.9

## Test sensitivity instead of publishing an unsupported payback

A single-point payback claim is analytically meaningless when the site failure distribution has not yet been measured. [S01] [S02] [S05] The World Bank Enterprise Surveys Cameroon 2024 Country Profile reports that 93.3 percent of surveyed firms experienced electrical outages (compared with 69.7 percent across Sub-Saharan Africa and 55.0 percent across lower-middle-income economies) and records 10.4 electrical outages in a typical month for Cameroon (against 6.9 for Sub-Saharan Africa and 5.3 for lower-middle-income comparators). [S01] However, as the Enterprise Surveys methodology documents, those indicators come from stratified, weighted responses across a broad national sample frame. [S01] [S02] The 10.4 monthly figure is not a site-level textile measurement, states no outage duration distribution, and does not record whether wide-area telecommunications failed simultaneously. [S01] [S02]

### 12.10 MODEL CORRELATED FAILURES AND BREAK-EVEN SURFACES EXPLICITLY

Inserting 10.4 outages per month into a cash-flow projection as a fixed constant hides two opposite risks. [S01] [S05] If a mill sits on a dedicated feeder that limits unbridged digital disruptions to 2 qualifying events per month, an expensive active-active cluster will never recover its incremental cost. [S05] Conversely, if a mill faces 15 power and backhaul interruptions per month during the rainy season and common-cause failures drop the local cellular tower battery alongside the factory feeder, an undersized buffer or unverified database configuration will suffer repeated rollbacks and post-restoration holds. [S05] [S16] [S18]

Rather than guessing a point payback, the engineer should evaluate a sensitivity surface and break-even thresholds across three candidate continuity options: [S05] [S1203]

- **Option A (Manual-plus-local-store baseline):** Pre-printed bundle dockets with UUID barcodes, clerical transcription into a single local SQLite database configured with `WAL` and `synchronous=FULL`, daily verified backups via the SQLite Backup API to removable media, and delayed batch upload. [S13] [S16] [S17] [S31]
- **Option B (Supervised protected edge node):** The Option A durable database augmented by automated point-of-work capture, an atomic event-and-outbox transaction boundary, an illustrative 150 W UPS providing 2 hours of digital ride-through and orderly shutdown (520.833 Wh nominal capacity), an 8-hour local spool (0.72 decimal GB payload draining in 9,600 seconds at 75,000 bytes/s net drain), a local freshness console, and one cold-spare device on site. [S05] [S08] [S13] [S18]

- **Option C (Redundant dual-node cluster with dual backhaul):** Clustered local compute nodes, automated failover supervision, dual independent WAN carriers (subject to physical path audit), and extended 8-hour battery autonomy across the switching and inspection chain. [S03] [S05] [S18]

WORKED EXAMPLE

**Illustrative break-even and sensitivity calculation across three service options**

To demonstrate the arithmetic without claiming any measured GFT factory result, consider an illustrative sewing and finishing department over a 12-month operating year, with illustrative values in thousand XAF (euro peg 655.957 XAF per euro). [R32] [S05]

- **Assumed illustrative annualised ownership costs:** Option A costs `C_A = 600` thousand XAF/year (docket printing, backup media, basic UPS upkeep, and restore drills). [S05] [S17] Option B costs `C_B = 2,400` thousand XAF/year, an incremental `dC_BA = 1,800` thousand XAF/year over Option A (covering the 520.833 Wh UPS cycle, industrial SSD wear reserve, cold-spare amortisation, and fault testing). [S04] [S05] [S28] Option C costs `C_C = 6,600` thousand XAF/year, an incremental `dC_CB = 4,200` thousand XAF/year over Option B (`dC_CA = 6,000` thousand XAF/year over Option A) for duplicate licences, secondary carrier retainers, larger battery banks, and cluster administration. [S05] [S18] [S1201]
- **Assumed illustrative per-event digital-boundary loss:** Excluding unpowered motor downtime, a qualifying disruption under Option A costs `L_eA = 95` thousand XAF per event (35 thousand XAF clerical and supervisory reconciliation labour, 40 thousand XAF unrecovered bottleneck contribution margin during a 45-minute post-restoration dispatch hold, and 20 thousand XAF re-inspection handling). [R14] [S13] [S1203] Under Option B, automated local capture rides through the 2-hour window or executes a clean commit-and-shutdown, leaving a 5-minute readiness check and residual loss `L_eB = 15` thousand XAF per event (avoiding `dL_eBA = 95 - 15 = 80` thousand XAF per event versus Option A). [S05] [S13] [S21] Under Option C, dual-node failover leaves residual loss `L_eC = 5` thousand XAF per event (avoiding an extra `dL_eCB = 10` thousand XAF per event over Option B, or `dL_eCA = 90` thousand XAF over Option A). [S05]
- **Break-even qualifying outage frequencies:** Upgrading from Option A to Option B breaks even at `f_BA = 1,800 / 80 = 22.5` qualifying events/year (`1.875` events/month), so 23 or more disruptions per year make Option B net positive. [S05] [S1203] Upgrading from Option B to Option C requires `f_CB = 4,200 / 10 = 420` qualifying events/year (`35.0` events/month), showing that the extra 10 thousand XAF saved per event cannot repay Option C's illustrative 4,200 thousand XAF incremental annual cost. [S05] [S1203]
- **Scenario net benefit comparison:** At low disruption (`f = 12` events/year), Option B yields `12 * 80 - 1,800 = -840` thousand XAF/year versus Option A, so Option A is preferred. [S05] At moderate disruption (`f = 48` events/year), Option B yields `48 * 80 - 1,800 = +2,040` thousand XAF/year over Option A, while Option C yields `48 * 90 - 6,000 = -1,680` thousand XAF/year. [S05] At high disruption (`f = 120` events/year), Option B yields `+7,800` thousand XAF/year over Option A, beating Option C (`+4,800` thousand XAF/year) by 3,000 thousand XAF/year. [S05]

The heatmap in **Exhibit 49**, p. 195 maps the net annual benefit of Option B over Option A across annual qualifying outage frequencies and per-event avoidable losses, holding the illustrative incremental annual cost at 1,800 thousand XAF. [S01] [S02] [S05]

EXHIBIT 49 / THOUSAND XAF PER YEAR

## The preferred option changes when outage and recovery assumptions change

| | 30 kXAF/event | 50 kXAF/event | 80 kXAF/event | 120 kXAF/event |
|---|---|---|---|---|
| 12 events/yr (1/mo) | -1,440 | -1,200 | -840 | -360 |
| 24 events/yr (2/mo) | -1,080 | -600 | 120 | 1,080 |
| 48 events/yr (4/mo) | -360 | 600 | 2,040 | 3,960 |
| 84 events/yr (7/mo) | 720 | 2,400 | 4,920 | 8,280 |
| 120 events/yr (10/mo) | 1,800 | 4,200 | 7,800 | 12,600 |

Cell scale: -1,440 to 12,600

ILLUSTRATIVE [S01] [S02] [S05] Illustrative sensitivity matrix computing Net Benefit = (Annual qualifying outages * Avoidable loss per event) - 1,800 thousand XAF incremental annual cost for Option B over Option A. Negative cells indicate conditions where the manual-plus-local-store baseline (Option A) is more economical; positive cells favour the supervised protected edge architecture (Option B). All figures are illustrative scenario calculations, not observed GFT results.

Inspecting **Exhibit 49**, p. 195 shows why site measurement (Chapter 2) must precede procurement freeze. [S02] [S05] At 12 qualifying events per year, Option B is negative across all four loss columns (-1,440 to -360 thousand XAF), favouring Option A's manual docket and local SQLite ledger. [S05] [S13] At 24 events per year and 80 thousand XAF avoidable loss per event, net benefit turns positive (+120 thousand XAF) and reaches +7,800 thousand XAF at 120 events per year. [S05] What **Exhibit 49**, p. 195 cannot prove is where a specific Cameroonian facility sits on this grid without site logging and a timed reconciliation drill. [S01] [S02]

### 12.11 ACCOUNT FOR CORRELATED FAILURES BEFORE PAYING FOR REDUNDANCY

Sensitivity analysis must also test **common-cause failure correlation** before allocating budget to dual-carrier network links. [S05] [S18] Vendors often justify a second cellular SIM or wireless ISP by multiplying independent unavailability figures (claiming that two links each with 5 percent illustrative unavailability yield 0.25 percent combined unavailability). [S18] If both providers share a local tower whose battery exhausts 90 minutes into a regional blackout, or cross the same duct bridge outside the industrial zone, their failures during prolonged outages are correlated. [S03] [S18] When the local edge queue already holds 8 to 24 hours of production events durably (Chapter 7), paying monthly retainers for a second carrier that fails during the same blackout buys no additional record durability. [S08] [S18]

12.12

# A vendor claim should become an acceptance obligation

Even a sound economic case fails if the purchased hardware and software do not deliver the assumed recovery behaviour when the grid drops. [S04] [S06] [S1201] Commercial proposals routinely rely on vague assurances such as "offline-capable edge AI", "enterprise persistence", or "guaranteed delivery". [S19] [S20] [S1201] Primary specifications show why none of those phrases establish an end-to-end production record contract: [S13] [S24]

- **Transport reliability is not application durability:** RFC 9293 specifies that TCP provides a reliable, in-order byte-stream service to applications, not proof that a receiving database has durably committed a bundle record to non-volatile media. [S13] [S24]
- **Protocol quality of service is not idempotent ledger posting:** OASIS MQTT Version 5.0 defines QoS 1 ("At least once", where duplicates can occur), and Eclipse Mosquitto writes state to `mosquitto.db` only when persistence is configured and checkpointed. [S19] [S20] Without transactional deduplication keyed by source UUID at the consumer (comparable to idempotent method semantics in RFC 9110 for HTTP), replaying an offline queue can double-count piecework output. [S12] [S13] [S31]
- **Database consistency is not crash durability:** SQLite documents that a transaction committed in `WAL` mode with `synchronous=NORMAL` might roll back following a power loss or crash, and full durability further depends on ext4 write barriers and storage media that honour flush commands. [S15] [S16] [S29]

### 12.13 TRANSLATE TECHNICAL REQUIREMENTS INTO MILESTONE-GATED CONTRACT SCHEDULES

Aligning with NIST SP 800-161 Rev. 1 and CSF 2.0 supply chain risk management practices (`GV.SC-05` and `ID.RA-09`), resilience and integrity requirements must be integrated into supplier contracts and verified prior to final payment. [S26] [S1201] The purchase contract should hold back an acceptance milestone payment until the supplier demonstrates six behaviours under witnessed fault injection (Chapter 13): [S06] [S28]

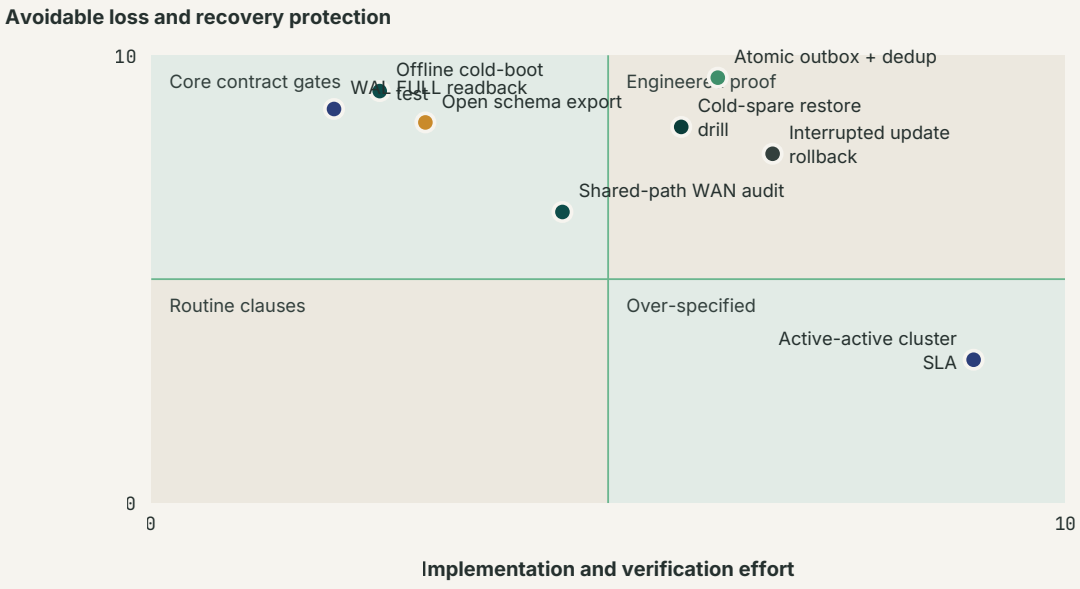
1. **Cold boot without external network reachability:** With wide-area links disconnected and upstream DNS unreachable, the edge appliance and inspection software must cold-boot, validate local licences and cached weights, authenticate operators, and record events within the recovery time objective. [S03] [S04] [S05]
2. **Readback of synchronous durability settings:** The system must provide an auditable readback proving that the local event store and outbox run on a local filesystem with `WAL` enabled, `synchronous=FULL`, and write barriers active. [S13] [S14] [S16] [S29]
3. **Zero committed-record loss under abrupt power removal:** Pulling input power during active synthetic event generation must result in zero loss of any locally acknowledged event, verified against an external test ledger. [S13] [S28]

4. **Bounded queue overflow and explicit backpressure:** Filling the offline spool to its high-watermark limit must preserve mandatory trace records, discard only pre-authorised diagnostic streams, increment a visible rejected-event counter, and alert the supervisor. [S03] [S05] [S08]
5. **Idempotent replay without duplicate business actions:** Replaying a backlog containing injected duplicate UUIDs, reordered timestamps, and interrupted transfers must produce exactly one committed business record per unique source event in the receiving PMS. [S12] [S13] [S19] [S31]
6. **Interrupted update rollback and spare handover:** Cutting power mid-update must leave the node booting cleanly into the previous signed version, and restoring an online backup onto the cold-spare unit must preserve source identity and sequence continuity. [S04] [S06] [S17] [S31]

The compliance matrix in **Exhibit 50**, p. 197 maps these procurement obligations along implementation effort versus avoidable-loss and recovery protection. [S04] [S05] [S06] [S1201]

EXHIBIT 50 / ILLUSTRATIVE SCORE (0 TO 10)

## The purchase must specify how continuity will be proved




ILLUSTRATIVE [S04] [S05] [S06] [S1201] Illustrative procurement prioritisation matrix scoring candidate contract obligations on implementation and verification effort (x-axis, 0 to 10) versus avoidable-loss and recovery protection (y-axis, 0 to 10). Upper-left items represent low-effort, high-protection mandatory contract gates; upper-right items require deeper engineering integration and witnessed fault testing. Scores are illustrative method weights, not empirical measurements.

In **Exhibit 50**, p. 197, the upper-left quadrant ("Core contract gates") highlights three low-effort, high-protection obligations that every tender should mandate: offline cold-boot testing (`x: 2.5, y: 9.2`), `WAL` and `synchronous=FULL` readback (`x: 2.0, y: 8.8`), and open schema export (`x: 3.0, y: 8.5`). [S04] [S16] [S1201] The upper-right quadrant ("Engineered proof") groups atomic outbox deduplication (`x: 6.2, y: 9.5`), cold-spare restore drills (`x: 5.8, y: 8.4`), and interrupted update rollback (`x: 6.8, y: 7.8`), which require witnessed fault testing to protect the ledger. [S06] [S13] By contrast, an active-active cluster service-level agreement (`x: 9.0, y: 3.2`) sits in the

lower-right quadrant ("Over-specified") when local skills and common-cause power failures make complex clustering fragile. [S05] What **Exhibit 50**, p. 197 cannot replace is formal legal drafting of remedies and warranty terms by qualified counsel. [S26] [S1201]

12.14

## Choose a maintainable option for the actual local team

An architecture that passes a vendor demonstration can still fail economically if its operation and post-outage restoration exceed the staffing, skills, and spare-parts reality of the plant. [S03] [S05] [S06] In Cameroon's textile sector, where industrial sites operate across regional contexts such as Garoua in the north and Douala or Dibamba near the coast, immediate access to vendor field engineers and proprietary modules cannot be assumed during an outage. [R03] [R10] [S05] When a multi-node cluster enters a split-brain state after battery exhaustion, a plant that waits 72 illustrative hours for external support suffers worse downtime than a plant running a simpler, locally maintainable design. [S05] [S06]

### 12.15 MATCH ARCHITECTURAL COMPLEXITY TO SITE RESTORATION DUTIES

To align the investment with the people who will run it (connecting to workforce training in wp09), the evaluation team should compare the three candidate options across four maintainability dimensions using the build-or-buy and use-case-portfolio decision tools: [S03] [S05]

1. **Routine shift supervision burden:** Under Option A (manual fallback), staff distribute, write, and key paper dockets for every bundle. [S05] Under Option B (supervised protected edge), operators scan bundles locally while the UPS holds the 150 W digital circuit, and the supervisor watches a single local freshness and queue indicator. [S03] [S05] Under Option C (redundant cluster), quorum or failover anomalies present diagnostic states that line supervisors cannot resolve locally. [S05] [S06]
2. **Local fault diagnosis and service restart:** A single Linux edge host managed via documented `systemd.service` units, persistent `journald.conf` logs, `chrony` RTC drift compensation (`rtcfile` with `-s`), and a local SQLite store can be inspected by a trained plant technician following a printed runbook. [S17] [S21] [S22] [S30] A proprietary appliance that hides logs behind a remote-support tunnel leaves local technicians blocked when the WAN is down. [S03] [S22]
3. **Spare-part replacement and identity handover:** Under Option B, swapping in a pre-staged identical edge box, restoring the latest verified online backup (`SQLite Backup API`), advancing the boot identity counter, and resuming capture is a deterministic 30-to-60-minute illustrative procedure. [S04] [S17] [S31] Under Option C, replacing a cluster node often requires online licence re-activation that fails while backhaul is degraded. [S04] [S23]

4. **Custom integration versus bought sub-assemblies:** Standard industrial UPS units, managed switches, and cameras should be bought with hardware warranties, whereas the transaction-and-outbox contract, UUID event schema, and reconciliation ledger should remain transparent so no black-box vendor locks up the system of record. [S13] [S31] [S1201]

> **GFT NOTE**
>
> **Practical recommendation: buy simple, inspectable local durability before paying for active-active redundancy**
>
> For a factory operating on an intermittent grid with a lean maintenance team, GFT recommends **Option B (a single supervised, battery-protected edge node per production zone, backed by an atomic SQLite WAL store, a cold-spare unit on the shelf, and a rehearsed paper-docket bridge)** over both an unbuffered cloud-dependent setup and a multi-node active-active cluster. [S05] [S13] [S17] Active-active clusters multiply licensing costs, double power draw on the critical UPS circuit, and introduce split-brain recovery modes that require specialist intervention precisely when external connectivity is down. [S05] [S06] A single well-instrumented node whose state fits in one verifiable file, paired with a cold spare that the local technician has practised restoring under the clock, delivers dependable recovery at a fraction of the lifecycle cost. [S06] [S17]

12.16

## Contractual exit must preserve the production history

The final test of a continuity investment occurs at the end of a supplier contract, software subscription, or pilot phase. [S26] [S1201] As NIST CSF 2.0 subcategory GV.SC-10 notes, cybersecurity supply chain risk management plans should include provisions for activities that occur after the conclusion of a partnership or service agreement. [S26] In industrial digital programmes, a recurring commercial failure is **traceability lock-in**: a factory pays for edge hardware and multi-year software licences, only to discover upon non-renewal that its historical production events, defect images, and quality dispositions are trapped inside a proprietary binary format or an edge agent that disables local queries when its subscription token expires. [S26] [S1201]

### 12.17 MANDATE OPEN SCHEMAS, PORTABLE DUMPS, AND SURVIVAL AFTER LICENCE EXPIRY

To ensure that accumulated evidence remains a permanent industrial asset (linking to data governance in wp14 and regulatory framing in wp17), procurement specifications should make exit portability a condition of contract award: [S26] [S1201]

- **Unencumbered ownership of production records:** The contract should state that all source events, bundle identifiers, timestamps, clock-quality metadata, inspection images, defect labels, and reconciliation audit trails generated at the site belong to the factory operator from capture. [S31] [S32] [S1201]

- **Documented, human-readable export without vendor assistance:** The local store and central PMS must support unthrottled local export in open formats at any time without requiring a vendor release key or cloud round-trip. [S17] [S1202] For SQLite-backed edge stores, the standard command-line `.dump` command converts the entire contents and schema of a database into a single portable UTF-8 text file, while online backup files and structured exports carrying RFC 9562 UUIDs and RFC 3339 UTC timestamps allow direct ingestion into a replacement system. [S17] [S31] [S32] [S1202]
- **Graceful behaviour upon licence expiry:** If a commercial subscription expires, the contract and acceptance test must forbid the software from deleting local data, revoking read access to historical logs, or halting safe local passthrough without notice. [S03] [S26] [S1201] Even if proprietary model inference stops after contract end, the local database and historical audit trail must remain readable and exportable by plant staff. [S17] [S1202]
- **Retention of recovery documentation and configuration artefacts:** Handover at commissioning must include schema definitions, `systemd.service` and `journald.conf` files, `mosquitto.conf` settings, database pragma configurations, and cold-spare restoration runbooks so plant engineers can maintain or migrate the record independently. [S16] [S20] [S21] [S22]

> **WARNING**
>
> **Unresolved commercial lock-in can erase a technically durable record**
>
> A production ledger that survives every electrical blackout through atomic `WAL` commits (`synchronous=FULL`) and battery-backed storage can still be lost if a supplier's licensing agent locks the database or withholds schema documentation at contract renewal. [S13] [S16] [S1201] Before signing any edge software, inspection, or cloud-synchronisation agreement, require a live exit rehearsal: take a populated test node, disconnect it from the vendor's licence server, expire its subscription clock, and verify that a plant technician can execute a complete `.dump` and backup export of every historical bundle record, UUID, and timestamp. [S17] [S31] [S32] [S1202] Commercial, intellectual-property, and termination clauses must be reviewed by qualified legal advisers prior to execution. [S26]

### 12.18 HAND VERIFIED ECONOMIC AND CONTRACT GATES TO COMMISSIONING

When the boundary decomposition, contribution-margin model, lifecycle cost register, break-even sensitivity surface, maintainability review, and exit portability clauses are complete, the investment decision no longer rests on optimistic vendor brochures. [S05] [S1201] [S1203] The sponsor knows which losses the protected digital circuit can avoid, how many qualifying disruptions per year are needed to justify the supervised edge architecture over a manual paper bridge, which line items belong in the multi-year budget, and which behaviours the supplier must prove before receiving

final milestone sign-off. [S05] [S06] Chapter 13 turns those contractual obligations directly into the fault-injection and commissioning programme that every candidate system must pass before a pilot is scaled across the factory. [S06] [S28]

*Buy the recovery the local team can prove, and keep the record when the contract ends.*

## CHAPTER SOURCE NOTES

S03 NIST. 2023.

S05 NIST. 2010.

S1203 NIST. 2022-11.

S13 SQLite. 2026-10-30.

S04 NIST. 2018.

S06 NIST. 2016.

S1201 NIST. 2024-11-01.

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.

S26 NIST. 2024-02-26.

S15 SQLite. 2026-10-30.

S16 SQLite. 2026-10-30.

S31 IETF / RFC Editor. 2024.

S08 IETF / RFC Editor. 2007.

S18 IETF / RFC Editor. 2014.

S02 World Bank. 2026-10-30.

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

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

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

S23 IETF / RFC Editor. 2026.

R13 Business in Cameroon; MMI News. Publication date unrecorded; inherited source-text check.

R11 Business in Cameroon; MMI News. Publication date unrecorded; inherited source-text check.

S01 World Bank. 2025.

S14 SQLite. 2026-10-30.

S20 Eclipse Mosquitto. 2026-10-30.

S22 systemd project. 2026-10-30.

S29 Linux kernel project. 2026-10-30.

S17 SQLite. 2026-10-30.

S24 IETF / RFC Editor. 2022.

S27 IETF / RFC Editor. 2020.

S28 SQLite. 2026-10-30.

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

S21 systemd project. 2026-10-30.

S19 OASIS. 2019-03-07.

S12 IETF / RFC Editor. 2022.

R03 SODECOTON. Publication date unrecorded; inherited source-text check.

R10 Ecofin Agency. Publication date unrecorded; inherited source-text check.

S30 chrony project. 2026-10-30.

S32 IETF / RFC Editor. 2002.

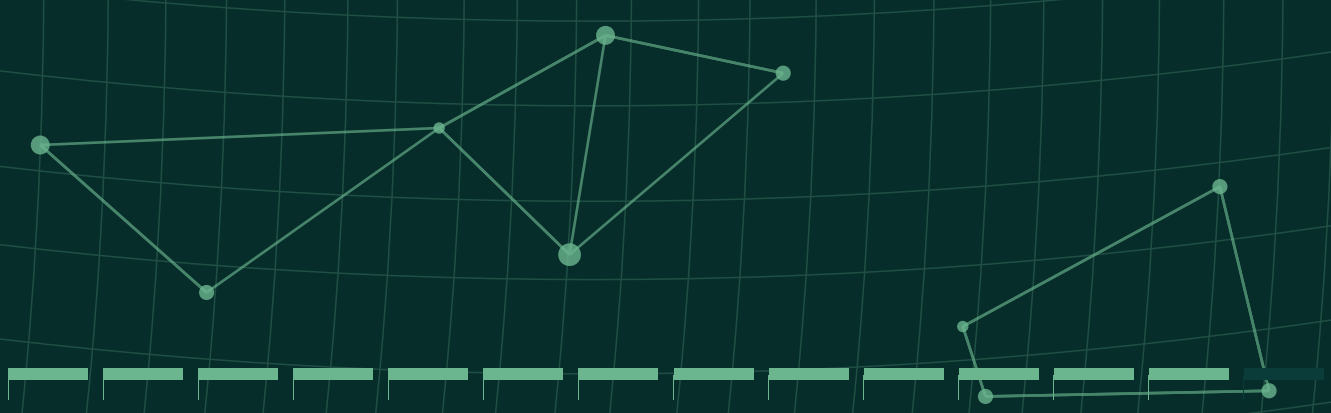
S1202 SQLite. 2026-10-30.

13

PART IV: PROVE AND GOVERN CONTINUITY / EARN THE ACCEPTANCE

# Commissioning must challenge the system at its failure boundaries

*A repeatable fault programme should verify declared losses, restart behaviour and reconciled records before a pilot is expanded.*



CHAPTER 13
DECISION TEST

A repeatable fault programme should verify declared losses, restart behaviour and reconciled records before a pilot is expanded.

Commissioning should try to break the continuity claim before production does. A weak-grid system is ready for a wider pilot only when named faults have been injected under load, every acknowledged record has been reconciled, recovery has been measured between agreed boundaries, and operators have demonstrated the fallback. A running service and a green dashboard are not acceptance evidence. This chapter proposes a test and evidence method. It does not report a completed GFT commissioning programme.

The method treats each design promise as a hypothesis. Local durability, offline capture, bounded replay and orderly restoration are accepted only for the hardware, software, configuration and operating state actually tested. Even a successful campaign does not establish the frequency of future outages or guarantee production availability. It establishes how the assembled system behaved under a recorded set of challenges, and leaves the remaining uncertainty visible.

13.1

## Map each requirement to an observable failure experiment

Every continuity requirement needs a test row that another competent team can repeat. The row should name the requirement, the fault class, the initial state, the active load, the exact stimulus, the expected behaviour, the acceptance limit and the retained artefact. It should also identify the person authorised to inject the fault and the person allowed to release the system afterwards. NIST contingency and recovery guidance supports this link between priorities, recovery playbooks, testing and subsequent improvement. [S05] [S06]

The initial state matters because an idle test can avoid the boundary that matters. A power cut after the queue has drained does not exercise recovery of a populated outbox. A broker restart without concurrent writes does not challenge the capture transaction. A backup restore without later live events does not expose replay ambiguity. The test record should therefore freeze the starting queue depth, active sources, write pattern, software state and pending update condition before the stimulus. These are test controls, not claimed factory measurements.

Expected behaviour should be written as invariants rather than impressions. If the local application acknowledged a durable commit, the acknowledged event must survive within the declared loss allowance. If acknowledgement was still in flight, the recovered state may contain either the complete committed event or no event, but never a torn business record. If a remote receiver retries delivery, repeated transport messages must not become repeated business actions. SQLite documents atomicity as all changes in a transaction occurring or none, while MQTT documentation expressly permits duplicates under at-least-once delivery. These are different boundaries and require different checks. [S13] [S19]

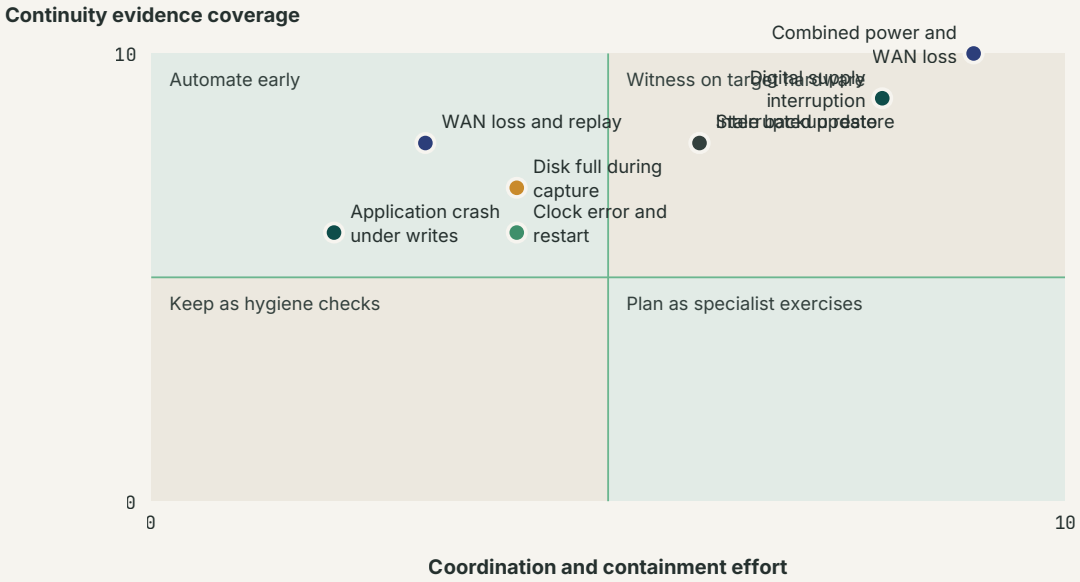
The fault catalogue should cover application termination, operating-system crash, controlled interruption of the digital supply, WAN loss, LAN partition, clock error, storage saturation, interrupted update and stale backup restoration. The sequence should begin with isolated software faults, progress to target-hardware tests and

finish with witnessed site exercises. A requirement that cannot yet be tested remains open. It must not be marked accepted because the corresponding component appears to work during normal operation.

**Exhibit 51**, p. 205 positions the proposed experiments by coordination burden and diagnostic coverage. The scores are illustrative planning judgements, not measurements of risk or reliability.

EXHIBIT 51 / ILLUSTRATIVE PLANNING SCORE FROM 0 TO 10

## Every failure requirement should have an evidence-producing test




ILLUSTRATIVE [S03] [S06] [S28] Illustrative sequencing model. Horizontal position represents coordination and containment effort; vertical position represents the breadth of continuity evidence expected from the test. Scores are GFT planning judgements, not observed failure rates or test results.

The matrix is a sequencing aid, not a reason to defer a high-consequence test until after acceptance. A mandatory case belongs in the acceptance plan regardless of where its point falls. Conversely, a difficult test with narrow diagnostic value should not displace simpler checks that expose the same requirement more directly.

### 13.2 FAULT INJECTION MUST RESPECT THE PHYSICAL OPERATING BOUNDARY

A continuity test must not create a process hazard in order to prove a data property. NIST's operational technology guidance explicitly places performance, reliability and safety requirements alongside security. [S03] The test designer should therefore separate the protected digital enclosure from production machinery, agree the safe mechanical state and obtain the responsible electrical and operations approvals before switching any physical supply. Simulation, a test rack or an isolated cabinet may answer the question without interrupting a live process.

The stimulus must also match the claim. Opening an upstream feeder would test much more than the edge record and could invalidate the exercise. A controlled interruption at the approved digital test boundary, preceded by instrumentation and a rollback

procedure, produces clearer evidence. Where the required claim concerns machinery control rather than record keeping, this chapter's method is insufficient and the responsible operational technology engineering process must govern the work.

> **WARNING**
>
> Do not generalise a digital-cabinet power test into permission to interrupt production equipment. The approved test record should identify the isolation point, connected loads, maximum test condition, emergency recovery action and responsible engineer. Stop the exercise if instrumentation, protection or the agreed safe state is lost.

13.3

## Simulation and hardware exercises answer different questions

Simulation and physical exercises are complementary because they expose different assumptions. SQLite's testing documentation states that crash testing with real power failures is impractical and describes simulation using an alternative virtual filesystem to represent the database state following a crash. [S28] That approach allows a project to repeat transaction interruption, process termination, disk errors and network failures across many controlled states. It is particularly useful before equipment reaches the site, when a defect can still be corrected without disrupting operations.

A project harness should test the application, not merely invoke the database vendor's own suite. It should interrupt capture before validation, between event insertion and outbox insertion, before commit, after commit and before the client receives acknowledgement. Each interruption should restart the actual application logic and compare the resulting records with the predeclared manifest. The harness should also repeat replay with duplicates, reordered messages, malformed payloads and another connection loss before the first replay has finished.

Simulation cannot establish what the physical device does with a flush request. SQLite's corruption guidance discusses drives that do not honour synchronisation requests and warns that the database is resistant to corruption rather than immune. [S15] The Linux `ext4` documentation states that write barriers enforce on-disk ordering of journal commits when volatile write caches are used. [S29] Consequently, the target-device test must include the selected storage media, controller, filesystem mount options, operating-system image and power path. A passing virtual filesystem test says nothing by itself about those components.

Database configuration also needs to be read back rather than inferred from a deployment file. SQLite documents that a transaction committed in WAL mode with `synchronous=NORMAL` might roll back after power loss or a system crash, even though consistency is retained. [S16] The commissioning result must therefore distinguish a structurally consistent database from conservation of every transaction that the application acknowledged as durable. The required setting is a design decision tied to the service contract, not a universal value to copy into every system.

Target-hardware exercises should repeat abrupt interruption during active writes, exhaustion of the protected supply, orderly shutdown and cold boot with external communications absent. They should record whether storage recovers without manual repair, whether services become ready in dependency order and whether the first new event receives the promised acknowledgement. Service supervision settings such as restart and stop timeouts are configuration inputs to version and exercise, not evidence that readiness has occurred. [S21]

The division of labour is clear. Simulation searches logic and timing states at low operational risk. Hardware exercises challenge the physical write path, transfer chain and boot sequence. Site exercises add the installed supply, network and operating team. None can be replaced by a supplier statement that the underlying components were tested in a different assembly.

13.4

## A deterministic event ledger reveals missing and duplicate work

A fault test needs an independent answer key. Before injection, the harness should generate a manifest containing each test event identifier, source identity, source sequence, canonical payload and payload hash. The manifest belongs on an observer outside the failure boundary, so a crash of the system under test cannot erase the record of what was offered. UUIDs provide a standard identifier format without a central registration process, but the implementation must still generate and handle them correctly. [S31]

The source observer must also record the acknowledgement boundary for every event. An event explicitly acknowledged as locally committed belongs to the expected-durable set. An event whose response was interrupted belongs to the in-flight set. An event refused with a visible rejection belongs to neither set. TCP or broker activity cannot make these classifications on its own because transport delivery does not demonstrate that the receiving application completed its business transaction.

Three ledgers then make the result inspectable. The source ledger records what was offered and acknowledged. The local committed ledger records what the recovered edge store contains, including its outbox and quarantine state. The remote committed ledger records what the receiving application accepted after validation and deduplication. Comparing identifiers and payload hashes across all three reveals defects that process status, row totals and queue depth conceal.

**Exhibit 52**, p. 208 defines the reconciliation categories and shows a synthetic result solely to demonstrate the audit format.

EXHIBIT 52 / ILLUSTRATIVE EVENT RECORDS

## Acceptance counts reconciled records rather than green process indicators

| Reconciliation check | Illustrative count | What the check establishes | Proposed response |
|---|---|---|---|
| Acknowledged and conserved everywhere | 9,420 | Identifier and payload hash match locally and remotely | Accept for the stated fault case |
| Acknowledged but missing locally | 0 | Potential breach of the local durable acknowledgement | Hold acceptance when above the declared allowance |
| In flight and completely resolved | 180 | Commit or rollback is known, with no torn record | Reconcile with the source result |
| Explicitly rejected after intake hold | 400 | No false acknowledgement; source retains retry responsibility | Check notice and bounded operator action |
| Local commit absent remotely after closure | 0 | Possible replay, validation or acknowledgement defect | Hold until every identifier is explained |
| Duplicate business action | 0 | Repeated delivery produced a repeated operational effect | Hold and repair consumer idempotency |
| Quarantined with a reason | 0 | Invalid or unresolved data remains visible | Accept only under the declared exception rule |

ILLUSTRATIVE [S12] [S13] [S19] [S31] Illustrative audit output, not a completed test. The source ledger is partitioned into acknowledged, in-flight and explicitly rejected events. A zero-loss gate applies only where the service requirement declares it; all categories still require explanation.

The table deliberately separates local committed loss, in-flight uncertainty, explicit rejection, remote stranding and duplicate business execution. These measures answer different questions. A system could have no missing local records and still strand the entire backlog remotely. It could drain the queue and still apply one event twice. Acceptance requires each category to satisfy the requirement, not a favourable average across them.

Transport duplication is not itself the same defect as duplicate business execution. MQTT's at-least-once delivery definition allows duplicates. [S19] HTTP semantics similarly distinguish methods defined as idempotent from the wider set of application requests and possible side effects. [S12] The receiver therefore needs an application contract that binds the event identifier to a transactional deduplication decision and a durable acknowledgement. A raw ingress log may show repeated delivery while the business ledger correctly records one effect.

The audit should also preserve uncertainty. If the observer cannot establish whether an acknowledgement reached the source, the event remains in the in-flight account. If a payload hash differs, the record goes to investigation rather than being counted as a successful match. If remote validation quarantines an event, the quarantine reason and retained payload become part of the result. Deleting an inconvenient record from the test set would destroy the evidence the commissioning exercise exists to produce.

**WORKED EXAMPLE: CLOSE THE EVENT BALANCE BEFORE JUDGING RECOVERY**

An illustrative test generates 10,000 events. Its independent source ledger records 9,600 events acknowledged as locally committed before the fault, 250 events in flight when acknowledgement was interrupted, and 150 events explicitly rejected after an intake hold. The illustrative source balance is therefore 9,600 + 250 + 150 = 10,000. These are constructed inputs for demonstrating the method, not observations from a factory.

After recovery, the illustrative local ledger contains all 9,600 acknowledged events and 180 complete events from the in-flight set, giving 9,780 locally committed records. The other 70 in-flight events are absent because their transactions did not commit. The remote ledger contains 9,775 valid events and five quarantined events with retained reasons, so its 9,780 records reconcile exactly with the local committed set. The full illustrative balance is 9,780 locally committed + 70 cleanly absent in-flight events + 150 explicit rejections = 10,000 generated events.

The illustrative result has zero unexplained loss of acknowledged events and zero duplicate business actions. It would satisfy those two proposed gates, subject to inspection of the raw evidence. It would not automatically pass the recovery-time, operator-release or quarantine-handling requirements. The arithmetic demonstrates conservation. It does not demonstrate the performance of any real GFT deployment.

### 13.5 THE AUDIT MUST SURVIVE LATER SCRUTINY

The evidence pack should retain the original manifest, acknowledgement log, local and remote exports, payload hashes, integrity-check output and reconciliation query results. It should record any excluded monitoring interval and any manual intervention. A summarised screenshot is not enough. Another reviewer must be able to repeat the set comparison and reach the same missing, duplicate and unresolved lists.

13.6

## Combine faults that share a plausible dependency

Single-fault tests establish a baseline, but they cannot expose every interaction between recovery mechanisms. The next stage should combine faults that the site dependency map shows could plausibly coincide. Shared supply, shared network equipment, limited storage and simultaneous restart are hypotheses to investigate, not assumptions of independence. The selected sequence must come from the site envelope where measurements exist, or be labelled as an illustrative stress case where they do not. [S05]

A combined mains and WAN exercise should begin with active local capture and an unavailable remote route. Once the outbox has accumulated work, the approved digital supply interruption should test commit and shutdown behaviour while onward delivery is impossible. Restoration should then occur in stages: local power without WAN, local durable capture with a growing retained backlog, and finally WAN return while new production events continue. This ordering tests whether replay competes with live capture and whether remote reconciliation can catch up under the declared capacity.

A storage-saturation exercise should fill the chosen partition only to its approved test condition while the system is disconnected. It should then restore the route and observe whether replay, status updates, journalling and diagnostic logging consume the remaining operating space. The expected result is not simply that intake stops. The system must preserve earlier commits, expose backpressure, maintain the reserved recovery path and leave a countable record of every refused event. The database and filesystem documentation makes clear that storage behaviour sits beneath the application promise. [S15] [S29]

An interrupted-update exercise should combine version change with process or power interruption at a declared deployment stage. It should verify the integrity of the package, the compatibility check, activation, rollback and the schema state after restart. NIST's firmware resiliency guidance supports attention to protection and recovery of platform firmware and data, but a project must still test its actual update mechanism. [S04] A previously working binary is not a sufficient result if the new configuration and old database schema cannot operate together.

Other combinations should follow the architecture rather than a fashionable checklist. Clock error during reconnection can challenge certificate validation and event ordering. A backup restore during replay can reintroduce old outbox cursors. A LAN partition during operator fallback can test whether manual and automatic records acquire conflicting identities. Each case needs a reason grounded in a shared dependency or a recovery action.

### 13.7 FAULT ORDER BELONGS IN THE EVIDENCE

The test log should record the exact order, relative offset, duration and system state of every stimulus. “Power failed during replay” is too ambiguous to reproduce. The useful record identifies which component observed each transition, what the queue contained, which version was active and when the operator intervened. It should also preserve failed attempts and aborted exercises rather than retaining only the clean final run.

Repeated synthetic cycles do not become a reliability estimate merely because there are many of them. They show behaviour under the selected sequences and implementation. Extrapolating outage frequency, annual availability or future site performance would require operating evidence with a defined denominator. Commissioning should remain precise: the campaign can retire specific design uncertainties while leaving environmental frequency and untested tails open.

13.8

# Recovery timing should be measured from agreed boundaries

A single recovery-time number conceals the services that different users need. The commissioning protocol should timestamp separate milestones using the same definitions in every run. The following boundaries turn contingency priorities into observable events rather than vendor-specific claims. [S05] [S21]

- **Fault stimulus, T0:** the approved power, network, process or storage fault is injected. The observer also records the last application-level durable acknowledgement before the fault. [S13] [S21]
- **Physical restoration, T1:** stable supply returns to the digital enclosure, or the relevant communication path returns for a network-only case. [S05]
- **Local service readiness, T2:** storage checks finish, required services are ready and local integrity checks have passed. A process merely existing does not satisfy this boundary. [S21]
- **First durable local capture, T3:** the first new post-fault event completes the promised local transaction and receives the defined acknowledgement. [S13]
- **Authorised inspection readiness, T4:** the responsible operator can see current state, staleness and any degraded-mode warning, and authorises the intended local work. [S03] [S06]
- **Remote ledger closure, T5:** eligible retained events have been replayed, retries have settled and the remote committed ledger reconciles with the local result. [S12] [S19]

These boundaries produce different measures. Local capture recovery is the interval from physical restoration to the first durable local capture. Operator recovery additionally includes inspection readiness and human release. Reporting convergence runs from restoration of usable connectivity to remote ledger closure. The last durable acknowledgement and the source ledger, rather than elapsed time alone, establish whether committed records were lost.

**Exhibit 53**, p. 212 shows a synthetic trace with relative times chosen only to demonstrate the measurement method.

EXHIBIT 53 / ILLUSTRATIVE SECONDS FROM FAULT INJECTION

## A recovery trace separates service return from ledger closure

0

**T0: Supply and WAN interrupted**

Capture continues locally; the observer freezes its durable-acknowledgement watermark.

300

**Protected shutdown begins**

The declared shutdown path quiesces intake and records its progress; WAN remains unavailable.

900

**T1: Supply restored**

The cabinet restarts on stable supply while the external route remains unavailable.

940

**T2: Local services ready**

Storage and integrity checks pass; the node clock remains marked unsynchronised.

948

**T3: First durable capture**

A new source event receives the specified local durable acknowledgement.

960

**T4: Operator release**

The local view shows live state, offline status and the retained backlog before work is authorised.

1,800

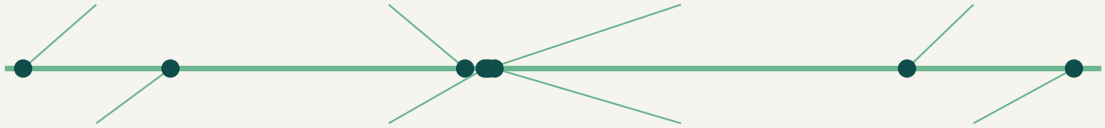
**WAN service usable**

Application-level remote acceptance becomes possible; replay starts under the live-arrival policy.

2,140

**T5: Remote ledger closed**

Local and remote identifiers reconcile, with explicit quarantine and no unresolved valid events.



ILLUSTRATIVE [S05] [S21] [S32] Illustrative timestamps for method demonstration, not measured equipment performance. An independent monotonic observer supplies the reference time; node timestamps retain clock-quality metadata. Acceptance limits must come from the site service classes.

The trace must carry clock-quality metadata. RFC 3339 requires an explicit relationship to UTC for the timestamps it defines, but correct formatting does not prove that a restarted device clock was accurate. [S32] An independent, continuously powered observer should therefore record a reference timestamp for each milestone. The node should retain its own timestamp, monotonic reading where available, boot identity and synchronisation state. Any interval that crosses a clock correction needs explanation rather than silent subtraction.

The acceptance limit comes from the service class agreed before the test. Remote reporting may have a different boundary from point-of-work capture, and inspection release may require human confirmation after the software is ready. No universal threshold is appropriate. A missed target leads to a defect, an approved waiver or a revised design, not a quiet change to the milestone definition after the result is known.

Timing evidence also needs its gaps. If the observer rebooted, lost a trace or could not distinguish physical restoration from stable supply, the affected interval remains unknown. The team should not replace it with an estimate in the acceptance record. A smaller set of trustworthy boundaries is more useful than a precise timeline built on inconsistent clocks.

13.9

# A pilot needs operating observations as well as lab results

Laboratory and site fault tests establish behaviour under controlled conditions. A bounded operational pilot adds the people, routines and disturbances that controlled tests cannot fully reproduce. It should run for an operator-approved window with explicit entry criteria, evidence duties and a scale decision at the end. NIST recovery guidance connects playbooks, testing and improvement, which makes the pilot a learning exercise rather than a delayed announcement of success. [S06]

The pilot evidence pack should include unscripted operator drills in authorised degraded modes. Observers should record whether staff recognise stale or unavailable state, select the approved fallback, preserve identifiers and complete later reconciliation without coaching from the implementation team. A technically successful fallback that supervisors cannot interpret is not an operational success. Any intervention by a remote supplier engineer should be logged as assistance, with the action and dependency made visible.

Real incidents provide a second evidence class. Power disturbances, communication losses, service restarts and storage warnings should be correlated with local journals and the event ledger. The team should ask whether each incident remained within the tested envelope and whether the system produced the expected local and remote records. A quiet pilot with no relevant disturbance demonstrates routine operation only. It does not erase an untested failure case.

Queue behaviour deserves observation under the real production cadence. The pack should retain occupancy, oldest pending age, rejected intake, storage headroom, replay duration and competition with live capture. These observations test whether the design assumptions were reasonable at the actual duty cycle. They do not turn an illustrative queue calculation into a measured forecast for another line or site.

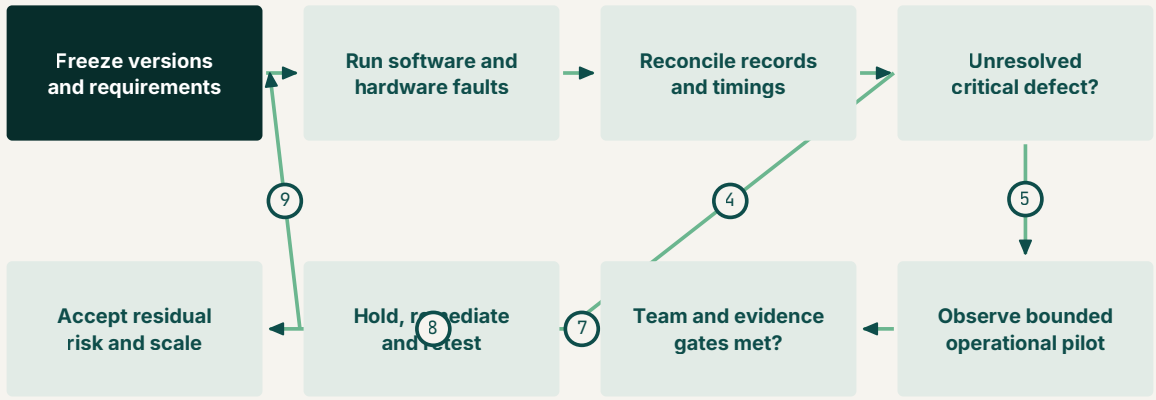
Manual reconciliation and exception handling need equal attention. Every disputed identifier, duplicate scan, quarantined payload and corrected manual entry should retain its reason, reviewer and resolution time. Maintenance actions should show whether local technicians could follow the runbook, restore a consistent backup, hand over a spare and return the system to an authorised state. Usability findings should identify alarms that were ignored, displays that were misunderstood and procedures that invited an unrecorded workaround.

> **GFT NOTE**
>
> GFT recommends a scale-up hold whenever a mandatory fault case remains untested, an acknowledged-record loss remains unexplained, or the local team cannot complete the approved fallback and reconciliation without undeclared specialist intervention. The hold should attach to a named defect and owner. Expansion resumes only after correction, a clean repeat of the affected test and review of the operating evidence. A calm pilot or an empty queue is not a substitute for that proof.

The decision path in **Exhibit 54**, p. 214 keeps technical and operational evidence inside the same gate. A defect returns the programme to a controlled baseline rather than being carried into a larger deployment.

EXHIBIT 54 / PROPOSED ACCEPTANCE WORKFLOW

## Unexplained loss blocks expansion of the pilot




4. Yes
5. No
7. No
8. Yes
9. Corrected baseline

ILLUSTRATIVE [S06] [S26] Illustrative GFT governance workflow. It is a proposed acceptance control, not a record of completed GFT tests or a certification scheme. Waivers require explicit risk ownership and cannot erase missing evidence.

The hold is not an argument against pilots. It is the reason a small pilot is valuable. A bounded deployment limits the number of people and records exposed while the organisation tests its assumptions. Expansion then uses evidence from the pilot, including unresolved exceptions, rather than treating the passage of time as acceptance.

13.10

## Acceptance should name versions, evidence and residual uncertainty

Acceptance is a decision about an identified configuration, not a permanent property of a product name. The dossier should bind every result to the equipment and software that produced it. Otherwise, a later storage substitution, firmware update or configuration change can inherit an acceptance statement that its own behaviour has not earned. Firmware resiliency and recovery guidance support attention to platform identity, recovery and improvement, while the NIST Cybersecurity Framework deliberately leaves how outcomes are achieved to the organisation. [S04] [S06] [S26]

The configuration record should include hardware and storage identities, platform firmware, the operating-system image, runtime packages, filesystem mount options and the protected supply components relevant to the test. It should retain a canonical configuration bundle and its hash, covering database durability, broker persistence, service supervision, logging, time, identity and shutdown behaviour. The purpose is not bureaucracy. It allows a reviewer to determine whether a later installation is the same tested system or a new configuration requiring further evidence.

The test provenance should sit beside that configuration. Each case needs its requirement, stimulus, initial state, load, operator approvals and source of test parameters. The dossier must say whether durations and fault combinations came from site measurements or from an illustrative stress envelope. It should retain the original manifest, acknowledgement watermarks, local and remote ledger exports, integrity checks, timing observations and reconciliation output. A summary can guide the approver, but it cannot replace the underlying artefacts.

The acceptance record should also contain a complete exception register. Each open defect, untested case, monitoring gap and deferred action needs an owner and a next decision. A waiver should state its scope, compensating control, expiry or review condition and the role accepting the residual risk. It should not convert an unexplained critical record loss into a pass. The NIST framework's outcome orientation does not turn a signed profile into certification or proof that a control works. [S26]

A useful dossier therefore contains the following connected records:

- **Tested identity:** hardware, firmware, image, packages and configuration hashes.
- **Requirement coverage:** every mandatory case, its result and any case still untested.
- **Raw evidence:** manifests, logs, ledgers, traces, hashes and reviewer queries.
- **Operating evidence:** drills, real incidents, interventions, exceptions and usability findings.
- **Residual uncertainty:** failures outside the campaign, missing observations and unresolved defects.
- **Accountable decision:** approver roles, waivers, corrective owners and conditions for retest.

### 13.11 QUALIFICATION STATUS MUST DESCRIBE THE EVIDENCE ACTUALLY OBTAINED

The report should distinguish three states. **Lab-qualified** means that the defined software and representative hardware passed the stated bench programme. **Site-accepted** means that the installed configuration passed the witnessed site exercises, ledger reconciliation and service-boundary requirements. **Operationally observed** means that the bounded pilot supplied the required shift, incident, maintenance and usability evidence. These are proposed evidence labels, not industry certifications.

A component can hold one state without holding the next. A bench result cannot establish the installed power path or local network behaviour. Site acceptance on the commissioning day cannot establish that operators will manage later exceptions

without support. Operational observation cannot cover a failure that never occurred and was never injected. Keeping the states separate prevents the strongest available evidence from being presented as something broader.

Change control should use the same boundaries. A storage-media or filesystem change directly challenges the physical write-path evidence and calls for repetition of the affected power-loss cases. A broker or consumer change challenges replay and deduplication. An operator-interface change calls for renewed fallback and release drills. The regression scope should follow the dependency analysis, with a recorded reason for cases not repeated. A new configuration hash without that analysis is only a label.

Finally, acceptance should transfer every open issue into the operating risk register and the programme's ordinary review cadence. Owners need the evidence link, compensating action and next verification condition, not merely a copy of the final report. The following roadmap can then stage expansion against the qualification state actually achieved. This chapter's acceptance method ends with a bounded decision: what was tested, what happened, what remains unknown and who accepts the consequence.

*Scale what the evidence has survived, and keep every untested boundary visible.*

**CHAPTER SOURCE NOTES**

S05 NIST. 2010.

S06 NIST. 2016.

S13 SQLite. 2026-10-30.

S19 OASIS. 2019-03-07.

S03 NIST. 2023.

S28 SQLite. 2026-04-21.

S15 SQLite. 2026-04-13.

S29 Linux kernel project. 2026-10-30.

S16 SQLite. 2026-10-30.

S21 systemd project. 2026-10-30.

S31 IETF / RFC Editor. 2024.

S12 IETF / RFC Editor. 2022.

S04 NIST. 2018.

S32 IETF / RFC Editor. 2002.

S26 NIST. 2024-02-26.

GREEN FUTURE TECH

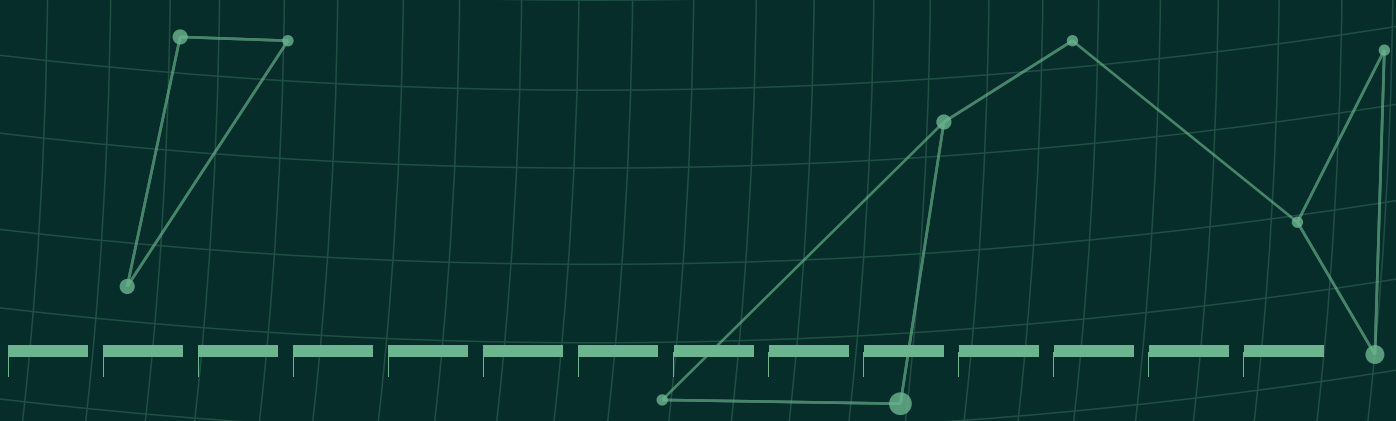
LIMITED / HONG KONG SAR

14

PART IV: PROVE AND GOVERN CONTINUITY / CLOSE WITH PROOF

# Scale the design only after continuity is demonstrated

*GFT recommends a staged programme that measures site conditions, proves a local record and expands only when recovery evidence supports it.*



CHAPTER 14
DECISION TEST

GFT recommends a staged programme that measures site conditions, proves a local record and expands only when recovery evidence supports it.

A weak-grid factory should scale its digital system only after the existing boundary has demonstrated safe local operation, conservation of acknowledged records and recovery by the local team. The correct sequence is to agree the production and record boundaries, measure the site, prove the smallest useful durable record, qualify inspection and remote reporting, then expand against operating evidence. A calendar can coordinate that work, but it cannot authorise it. Each stage earns the next through artefacts that another reviewer can inspect. This chapter sets out that proposed programme and its measures. It does not report a commissioned GFT deployment or a realised site result.

The staging also limits the cost of being wrong. A site baseline can invalidate an autonomy assumption before equipment is ordered. A minimal record pilot can expose a lost commit before inspection, remote analytics and additional lines inherit the defect. An operational pilot can reveal that a technically sound fallback is unusable by the shift team. Keeping each uncertainty inside the smallest practical boundary is therefore both an engineering control and an investment control.

14.1

## Begin by agreeing the production and record boundaries

The first deliverable is an operator-approved boundary statement, not a bill of materials. It should identify the tasks that must continue, those that may degrade and those that must stop under each declared failure. It should also distinguish physical production from the digital record. Protecting an edge computer and its network can preserve capture and supervisory visibility while machinery is stopped. It cannot make an unpowered motor turn. NIST operational technology guidance places performance, reliability and safety beside security, so these boundaries belong to the responsible operations and engineering authorities rather than to a software supplier alone. [S03]

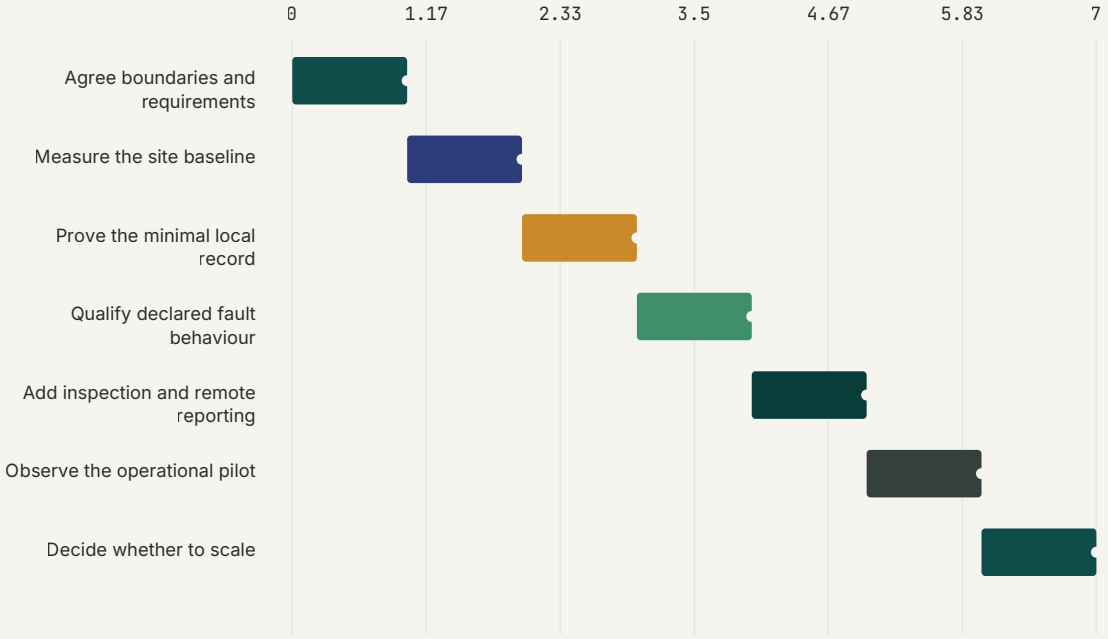
The record boundary needs equal precision. For every included task, the site should agree the event identifier, source and sequence fields, operation and bundle context, quality disposition, clock-quality fields, acknowledgement meaning and reconciliation owner. The statement should say which acknowledgements mean validated, written, locally committed or remotely accepted. It should name the approved manual process, the authority to enter it and the evidence required to leave it. Contingency planning begins by evaluating operations to determine requirements and priorities, which is why these decisions precede technology selection. [S05]

A requirement register should carry each choice into later acceptance. Its rows need a failure stimulus, permitted service state, allowable record loss, recovery boundary, evidence owner and unresolved assumption. Where the site cannot yet choose a value, the register should record the missing observation and the decision date or condition. A blank field marked unknown is more useful than a plausible default copied from another factory. The same register should define who can change a requirement after a test result is known, preventing an inconvenient failure from becoming an informal relaxation.

**Exhibit 55**, p. 219 expresses the proposed sequence as ordered positions rather than promised dates. Each transition depends on evidence from the preceding boundary.

EXHIBIT 55 / ILLUSTRATIVE ORDERED STAGE POSITION, NOT ELAPSED TIME

## The roadmap advances through evidence gates




ILLUSTRATIVE [S05] [S06] [S28] Illustrative programme logic, not a delivery schedule or a record of completed GFT work. Equal numeric spacing shows sequence only. Durations, resourcing and any permitted overlap remain operator-selected planning inputs after the site survey.

The roadmap deliberately makes the manual performance-management record the starting capability. Structured point-of-work logs can establish identifiers, operating routines and supervisory review before automated capture is expanded. Bought inspection becomes a qualified addition to that record, with its own inputs, offline startup and human fallback. Forecasting belongs later still, because a forecast cannot repair an incomplete or untrusted event history. This order follows GFT's Cameroon textile programme as a design basis, not as evidence that the capabilities have already been commissioned.

This sequence does not prohibit remote services. It prohibits making them silent prerequisites for local duties. Non-urgent planning, central comparison and model development can proceed in parallel when their absence cannot interrupt local capture or quality authority. The boundary statement should make that independence testable. If a supposedly remote service can block a cold start, a licence check or an operator login, it belongs inside the continuity scope.

14.2

## Measure the site before fixing autonomy and retention

The baseline stage should replace general concern about unreliable supply with an observed site envelope. Recorders should capture power events at the intended digital circuits, while the network survey traces switching, naming, backhaul, shared supplies and upstream dependencies. Operator interviews should reconstruct recent

failures, manual workarounds, restoration delays and record disputes. Their recollections are evidence about operating practice and investigative leads, not a substitute for instrumented durations. The resulting dossier should distinguish measurement, recollection and assumption.

National evidence establishes why this work matters but cannot set the equipment rating. In the World Bank's Cameroon 2024 country profile, 93.3 percent of surveyed firms reported experiencing electrical outages and the reported frequency was 10.4 outages in a typical month. These are firm-response indicators across the survey population, not a textile-only sample, a duration distribution or a prediction for the proposed site. [S01] The survey methodology also constructs population inference through sampling and weighting rather than measuring every factory circuit. [S02] Neither statistic can be inserted directly into a battery, queue or recovery calculation.

A useful baseline therefore needs its own measurement-quality controls. The recorder must remain alive through the events it is meant to observe, or its gaps must be visible. Its clock state must accompany event start and end. The team must know whether a communication test reached a modem, resolved a name, established an authenticated session or completed an application acceptance. Each of those observations answers a different question. A period without a valid recorder reading is unknown, not a period without an outage.

The baseline collection plan should include:

- **Power events:** start, end, affected phase or circuit, transfer behaviour, protective operation and stable restoration, interpreted by the responsible electrical engineer. [S03]
- **Connectivity events:** local link, external route, name resolution, authenticated service and application acceptance, with shared dependencies marked.
- **Operating consequences:** capture interruption, manual entry, inspection hold, delayed release, queue accumulation and later reconciliation effort.
- **Observation quality:** recorder coverage, missing intervals, clock state, configuration changes and any event still open when collection ends.

Tail events require an explicit decision rather than silent extrapolation. A short observation window may not contain the long outage that determines retention or protected energy. Extending measurement can reduce that uncertainty, but it cannot prove that the largest observed event is the largest possible event. The design freeze should therefore record the observed distribution, monitoring coverage, censored events and the additional allowance or stress case chosen for unobserved tails. The basis is a risk decision informed by contingency priorities, not a claimed statistical maximum. [S05]

Design freeze occurs only when the boundary statement and outage envelope have been reviewed together. At that point the team can size protected digital energy, storage, replay capacity and recovery targets against stated assumptions. If the evidence is insufficient, the choices are to collect more data, narrow the promised service or carry an explicitly accepted conservative assumption. Procuring first and describing the envelope afterwards would hide precisely the uncertainty the baseline is meant to resolve.

14.3

# Prove a minimal durable local record before broad integration

The first technical pilot should do one useful thing completely: capture an agreed production event, commit it locally under a declared durability contract, show its state to a supervisor and recover it after the declared faults. It need not carry every sensor, inspect every piece or feed every central dashboard. Narrowing the functional scope allows the test to concentrate on the properties that later components will inherit. It also gives operators a usable record early, rather than asking them to wait for a broad integration programme before seeing any benefit.

The event and its pending-delivery record should share one transaction. SQLite defines an atomic commit as all changes within a transaction occurring or none occurring, which supports the proposed event-and-outbox boundary. [S13] Atomicity alone, however, does not prove survival of the latest acknowledgement. SQLite also documents that a transaction committed in WAL mode with `synchronous=NORMAL` might roll back after power loss or a system crash. [S16] The pilot must therefore test the complete application, filesystem, storage and power path against the acknowledgement it actually gives. A database setting copied into a configuration file is not that proof.

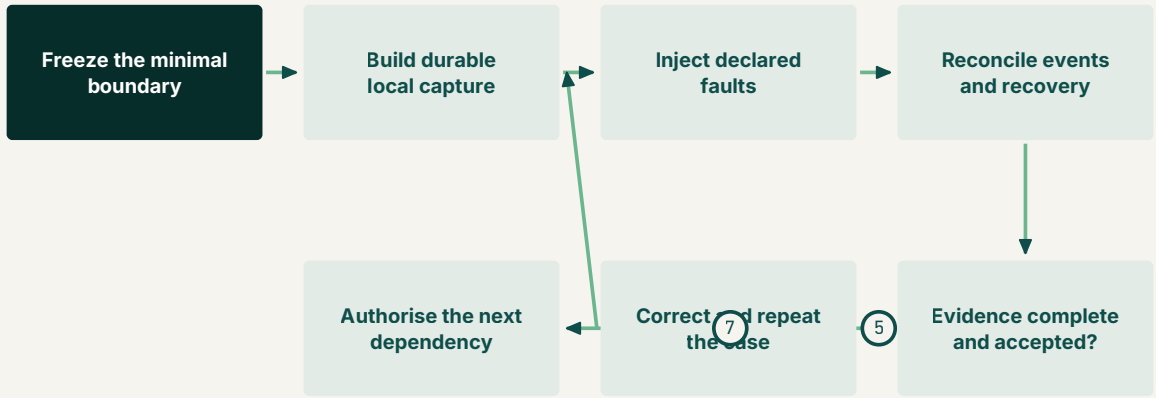
The minimal pilot should include local dashboard freshness, controlled shutdown, abrupt-loss recovery and manual import. Each path needs an observable end state. Freshness means that a supervisor can distinguish current, stale and unknown information. Shutdown means that intake is controlled and pending state is handled before power disappears. Abrupt recovery means that the system returns without hiding missing or in-flight work. Manual import means that an authorised fallback record acquires the same identity and reconciliation discipline as an automated event.

Simulation belongs early in this stage because it can repeat interruption points without disturbing production. SQLite describes simulated crash testing in its own test programme, but that does not qualify the selected factory image or hardware. [S28] The project should therefore progress from application fault injection to controlled target-device exercises and then to the installed boundary. At every level, the expected event ledger remains the answer key. A green process indicator cannot substitute for a comparison of acknowledged, stored, replayed and unresolved identifiers.

**Exhibit 56**, p. 222 shows the proposed investment gate. Additional integration begins only after the current record closes its evidence loop.

EXHIBIT 56 / PROPOSED EVIDENCE-GATE WORKFLOW

## The minimal record earns each additional dependency




5. No or unknown

7. Yes

ILLUSTRATIVE [S06] [S13] [S28] Illustrative GFT stage-control design, not a completed commissioning workflow. The pass condition requires the site's declared acceptance limits, raw event reconciliation and an authorised decision. A failed or untested mandatory case returns to remediation.

The gate should examine defects as well as passes. An unexplained missing acknowledged event, an invisible manual duplicate or an unmeasured recovery interval blocks the claim it affects. A waiver may accept a documented residual risk where the responsible authority permits it, but it cannot turn missing evidence into observed success. Recovery guidance links playbooks, testing and improvement, so remediation should update the design and the operating instructions before the case is repeated. [S06]

> **GFT NOTE**
>
> GFT recommends releasing funds for the next integration stage against an evidence pack, not against a configurator screen or a demonstration during normal operation. The pack should identify the tested configuration, expected and recovered event ledgers, fault coverage, recovery boundaries, operator actions, unresolved defects and the exact service now authorised. If a required result is unknown, the next stage remains a proposal with an owner and a test, not an assumed capability.

A minimal success also defines what has not been earned. It does not establish inspection accuracy, remote ingestion behaviour, fleet-wide manageability or forecasting value. Those claims belong to later stages with their own entry criteria. Preserving this distinction keeps a successful local pilot from becoming a general licence to scale.

14.4

# Add inspection and remote reporting without weakening the boundary

Inspection should join the accepted local record as a separately qualified function. Its camera, illumination, trigger, compute, model, licence, operator display and disposition record need to work within the declared power and connectivity envelope. The responsible quality authority must also approve the human fallback and the conditions for stopping or releasing work. NIST's operational technology guidance supports treating reliability and safety as part of the design context rather than as later security additions. [S03] A local model file by itself establishes none of these operating properties.

The entry test is an offline cold boot and recovery exercise against the already accepted local store. Inspection must start without relying on an external service that the boundary says is optional. Its events must use the established source identity, clock quality and acknowledgement semantics. If the new function saturates storage, competes for protected energy or delays the first durable capture after restart, the previous evidence no longer describes the assembled system. The affected tests must be repeated before the integrated state can be accepted.

Remote reporting should follow the same rule. HTTP defines particular methods as idempotent, but that specification does not make every factory ingestion endpoint or its downstream side effects idempotent. [S12] MQTT at-least-once delivery expressly permits duplicates. [S19] The remote receiver therefore needs an application-level identity, transactional deduplication and a durable acceptance response. Replaying a backlog must not repeat a business action or starve current local capture. These are integration requirements to demonstrate, not properties inherited automatically from a protocol name.

The remote path also needs a failure return to the local state. If name resolution, authentication, application acceptance or capacity fails, the local service should continue within its declared limits and expose the growing backlog. Restoration should prioritise current operating records and then replay under an agreed capacity policy. The central dashboard remains provisional until reconciliation closes. No remote alarm should be interpreted as proof that the corresponding local or downstream action completed.

This chapter supplies continuity acceptance inputs rather than a complete connected-factory architecture. The wider component design belongs to `wp18`, while model and data governance belong to `wp14`. Those workstreams can change an inspection model, broker or remote consumer, but they must inherit the local boundary and its regression obligations. A handover is complete only when the receiving design states which tested claims it preserves, which it invalidates and which new evidence it will produce. [S03] [S06]

If integration fails, rollback should restore the last accepted operating boundary rather than remove the whole record. That may mean suspending the new inspection function, holding remote replay or returning to the approved manual disposition while local capture continues. The rollback path itself needs an operator instruction and a reconciliation check. A stage is reversible only when that return has been exercised.

14.5

## Expand only when the local team can restore service

A system is not operationally ready while recovery depends on an unrecorded intervention by its implementer. Before expansion, the local team should demonstrate spare handover, consistent backup restoration, clock recovery, supplier escalation and supervisor release under authorised conditions. Recovery guidance connects playbook development with testing and improvement, making the people and their instructions part of the capability being assessed. [S06] The relevant question is not whether a supplier can eventually repair the system. It is whether the assigned roles can restore the promised boundary with the resources actually available.

The spare-handover drill should begin with the approved replacement, image and configuration rather than an improvised device. The team must establish firmware and software integrity, compatible storage and interfaces, trusted credentials and a controlled transfer of source identity. NIST firmware resiliency guidance supports attention to platform protection and recovery, but it does not demonstrate that a particular spare fits or boots. [S04] The drill record should show who authorised the substitution, how the original source was prevented from duplicating authority and which checks released the replacement into service.

Backup creation likewise needs a matching restoration result. SQLite's online backup mechanism supports making a consistent backup while other database users continue, but a completed backup file does not prove a usable restore. [S17] The team should restore into the approved isolated state, check database integrity and configuration, reconcile the event and outbox ledgers, and only then request operating release. The exercise should include a backup that predates later live events, because restored replay state can otherwise create duplicates or false pending work.

Clock recovery requires the same practical treatment. Chrony documents a mechanism for compensating real-time-clock drift through `rtcfile` with `-s`, but the chosen device still needs disconnected and restart measurements to establish its behaviour. [S30] The drill should record the clock state before recovery, the approved source or manual authority used to correct it, the uncertainty attached to intervening events and the point at which remote trust can be re-established. Disabling authentication or silently rewriting event history is not an acceptable shortcut.

Supervisor drills should test interpretation as well as execution. The shift team must recognise stale data, missing sequences, an unavailable inspection input and an unreconciled manual record. It should be able to select the approved degraded state, preserve the required identifiers and explain why production is running, restricted or

held. Any coaching, remote login or specialist action should appear in the evidence as assistance. Repeating a drill after coaching is useful, but the first attempt must not be erased from the learning record.

Expansion should then replicate an accepted configuration rather than improvise a new one at each line or site. The rollout pack should contain the versioned image, configuration, requirement register, test evidence, local runbook, spares list, escalation route and training record. A change in storage, firmware, network path or inspection duty triggers a dependency review and targeted repetition of the affected evidence. Technician and supervisor development belongs to the wider skills programme in wp09; this paper treats that development, spare availability and drill time as funded operating requirements, not assumed GFT capabilities. [S04] [S06]

The scale decision should also define reverse motion. If the replicated system loses its required record behaviour, if the team cannot complete restoration or if assistance exceeds the accepted operating model, the programme returns to the last supported boundary. That response is not a failure of the roadmap. It is the mechanism that prevents a local defect from becoming a fleet-wide condition.

14.6

## Measure survival, freshness and recovery in the PMS cadence

The performance-management system should show whether the continuity promise survived actual operation. A useful scorecard separates local capture availability, committed-event loss, pending-record freshness, offline coverage, clock uncertainty, restoration time and unreconciled work. These measures respond to different failure modes. A high capture-availability result can coexist with lost commits, and an empty remote queue can coexist with an unacknowledged or rejected production record. The scorecard must therefore preserve individual gates rather than hide them inside one favourable average.

**Exhibit 57**, p. 226 defines the proposed measures and their decision use. Every definition needs an event source, reporting period, exclusion rule and named owner before it enters routine review.

EXHIBIT 57 / PROPOSED METRIC DEFINITIONS

## Every continuity measure needs a numerator and denominator

| Measure | Numerator | Denominator or boundary | Decision use |
|---|---|---|---|
| Local capture availability | Seconds when capture is ready | Seconds when capture was planned | Test service survival, with monitoring gaps separate |
| Committed-event loss | Locally acknowledged IDs absent after recovery | All IDs acknowledged as locally committed | Hold the affected durability claim above its agreed allowance |
| Oldest pending age | Reference now minus oldest eligible pending capture time | Pending set at the review instant | Expose stale work and include clock uncertainty |
| Offline record coverage | Offline periods with a valid required record | All periods declared offline by the independent site log | Find invisible gaps in the continuity claim |
| Clock uncertainty | Declared bound or unbounded state for the event set | Events used for ordering or elapsed-time measures | Qualify freshness, recovery and reconciliation |
| Restore time | Elapsed time between approved restoration boundaries | Named fault and service class | Compare local readiness, release and remote closure separately |
| Unreconciled bundles | Bundles with missing, duplicate or conflicting records | Bundles requiring reconciliation in the period | Assign operational review before release or expansion |

ILLUSTRATIVE [S05] [S13] [S32] Illustrative GFT metric dictionary, not observed site results or universal targets. The operator must approve each threshold. Interchange timestamps need an explicit UTC relationship, while trustworthy elapsed time also requires clock-state and uncertainty evidence.

Denominators must remain visible beside every result. Planned capture time should follow the operating calendar and declared degraded states, not merely the time a monitoring process happened to run. Production downtime and monitoring downtime need separate flags. If the observer failed during an outage, the corresponding availability and offline-coverage results are incomplete. Treating that interval as healthy would reward the loss of the measurement system itself.

Freshness requires a defined reference and pending set. RFC 3339 timestamps state a relationship to UTC, but their format does not establish clock accuracy or prove event order. [S32] The oldest-pending calculation should therefore use an independent reference where possible and carry the source clock state and uncertainty. It should also say whether explicitly rejected, expired or quarantined records remain in the pending population. Removing them from the denominator without an auditable reason would make freshness look better by hiding work.

The cadence should move from local response to programme learning. Shift review deals with current stale records, unknown states and unreconciled bundles. Regular operational review examines repeated faults, restoration boundaries, manual effort and assistance. Stage-gate review compares the complete evidence with the require-

ments and decides whether to hold, correct or expand. Thresholds remain operator decisions derived from contingency priorities; the metric dictionary does not prescribe a universal service level. [S05]

> **WORKED EXAMPLE: CALCULATE THE GATES WITHOUT HIDING UNCERTAINTY**
>
> An illustrative shift contains 28,800 planned capture seconds, of which the local service is ready for 27,540 seconds. The illustrative local capture availability is `27,540 / 28,800 * 100 = 95.625 percent`. The independent observer covers the whole shift in this constructed case, so no interval needs to be removed as unobserved. The calculation demonstrates the definition, not a factory result.
>
> The same illustrative source ledger contains 2,400 event identifiers acknowledged as locally committed before a fault. After recovery, 2,388 of those identifiers are present with matching payloads. The illustrative committed-event loss is `12 / 2,400 * 100 = 0.5 percent`. The twelve missing identifiers remain a named exception set. They cannot be averaged against capture availability to produce a passing continuity score, and the result should hold the durability claim if the agreed allowance is zero. [S13]
>
> At an illustrative review instant, the oldest eligible pending event was captured 900 reference seconds earlier, with a declared clock uncertainty of plus or minus 20 seconds. The report should state an age of 900 seconds with that uncertainty, not a more precise value. Seven illustrative bundles also remain unreconciled and are listed individually for supervisor action. None of these constructed values supplies a pass threshold or predicts another shift. Together they show why survival, durability, freshness and reconciliation need separate decisions. [S05] [S32]

Routine reporting should retain the raw identifiers and interval records behind each indicator. A reviewer must be able to move from the displayed percentage to the missing events, from the recovery interval to its boundary timestamps and from the unreconciled count to the affected bundles. Detailed overall equipment effectiveness calculations remain in `wp15`; this chapter does not permit missing or unknown monitoring periods to be converted silently into stopped production, productive time or improved efficiency. [S05] [S32]

14.7

## GFT should report verified outcomes and the gaps that remain

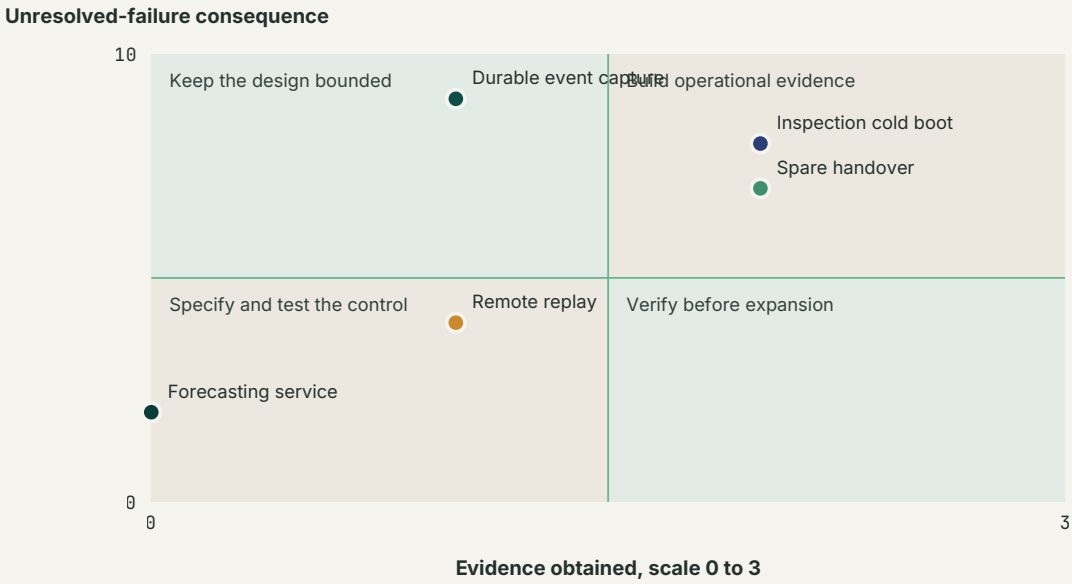
The closing report should state the strongest evidence obtained for each requirement and stop short of everything broader. **As-designed** means that the boundary, mechanism and acceptance method have been specified. **Tested** means that an identified configuration has produced retained results under declared laboratory or target-hardware faults. **Site-accepted** means that the installed boundary has passed the witnessed site requirements. **Operating-observation** means that the bounded pilot has supplied the required shift, incident, maintenance and usability evidence. These are proposed evidence labels, not certifications. The NIST Cybersecurity Framework provides outcomes without prescribing how to achieve them, and it does not supply a continuity acceptance certificate. [S26]

A scorecard row should connect each status to its proof and limits. It needs the requirement, configuration identity, evidence link, test and observation coverage, unresolved issue, owner and next action. It should also name the failures not injected, the operating conditions not observed and any assistance required during recovery. Recovery planning includes testing and improvement, so the final report is also an input to the next playbook revision rather than a static declaration of completion. [S06]

**Exhibit 58**, p. 228 shows why status and consequence should be read together. Its points are constructed planning examples, not assessments of an actual GFT system.

EXHIBIT 58 / ILLUSTRATIVE EVIDENCE AND CONSEQUENCE

## Completion status must state the evidence actually obtained




ILLUSTRATIVE [S06] [S26] Illustrative closing-scorecard example. Horizontal positions represent the strongest evidence obtained: 0 as-designed, 1 tested, 2 site-accepted, 3 operating-observation. Vertical positions are invented operator-consequence scores for layout demonstration, not risk measurements. Real rows require evidence links, owners and residual uncertainties.

The matrix does not produce a numerical readiness score. A high-consequence requirement with only design evidence cannot be offset by extensive observation of a low-consequence reporting feature. Nor does a point to the right imply that every relevant fault has occurred. The accompanying register must still list coverage and unknowns. Its practical value is to direct the next action: specify, test, accept on site, observe operation or hold the affected expansion.

Verified outcomes should be written narrowly enough to remain true outside the test room. A defensible statement identifies the configuration, fault, initial state, observed record result, recovery boundary and reviewer. It does not become a general availability percentage or a promise of financial value. Where no relevant incident occurred during the operational pilot, the report should credit routine operation and the drills actually performed, while retaining the untested incident as an open boundary. Where evidence contradicts the design expectation, the exception belongs on the main scorecard rather than in an appendix that the decision will not see.

Change control must preserve that honesty after acceptance. A new storage device can invalidate power-loss durability evidence. A new broker or consumer can invalidate replay and deduplication evidence. A new model or operator display can invalidate inspection release and fallback evidence. The status should therefore attach to a versioned configuration and regress when a material change occurs. Regression is not an accusation that the new version will fail. It is an acknowledgement that the old evidence no longer answers the same question.

The final governance pack should leave a usable trail for the operating organisation. Requirements link to tests, tests link to raw ledgers and timings, operating indicators link to exceptions, and exceptions link to owners and review conditions. Supplier escalation, training needs, spare readiness and residual risk enter the ordinary management cadence. This makes continuity a maintained operating duty rather than a commissioning document that gradually separates from the installed system.

GFT's closing recommendation is consequently restrained. Engineer the safe local record first. Budget protected power, retention and recovery capacity against measured site conditions and declared assumptions. Add inspection, remote reporting and forecasting only when their dependencies preserve that record. Scale when the local team can restore service and the evidence distinguishes what is designed, tested, accepted and observed. This research supplies the roadmap and measurement method. It claims no site performance, completed deployment or realised digital value.

*Continuity earns scale one verified boundary at a time.*

## CHAPTER SOURCE NOTES

S03 NIST. 2023.
S05 NIST. 2010.
S06 NIST. 2016.
S28 SQLite. 2026-10-30.
S01 World Bank. 2025.
S02 World Bank. 2026-10-30.
S13 SQLite. 2026-10-30.
S16 SQLite. 2026-10-30.
S12 IETF / RFC Editor. 2022.
S19 OASIS. 2019-03-07.
S04 NIST. 2018.
S17 SQLite. 2026-10-30.
S30 chrony project. 2026-10-30.
S32 IETF / RFC Editor. 2002.
S26 NIST. 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

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.

### R02 Inherited evidence R02 GRADE C

Fibre2Fashion; Business in Cameroon. Publication date unrecorded. fibre2fashion.com; businessincameroon.com

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

### R03 Inherited evidence R03 GRADE C

SODECOTON. Publication date unrecorded. sodecoton.cm

Company profile: more than 220,000 cotton producers supported in northern Cameroon; nine ginning plants; two oil mills.

Inherited source-text check; original source not independently re-opened for this edition.

### R04 Inherited evidence R04 GRADE C

Financial Fortune Media. Publication date unrecorded. financialfortunemedia.com

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

### R05 Inherited evidence R05 GRADE C

Ecofin Agency. Publication date unrecorded. ecofinagency.com

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

### R06 Inherited evidence R06 GRADE C

Cameroon Trade Hub. Publication date unrecorded. cameroontradehub.cm

About 95 percent of cotton fibre exported raw against 5 percent processed locally; industrialisation target of 50 percent local processing.

Inherited source-text check; original source not independently re-opened for this edition.

### R07 Inherited evidence R07 GRADE C

Ecofin Agency. Publication date unrecorded. ecofinagency.com

SND30 restates the 50 percent local-processing target by 2030.

Inherited source-text check; original source not independently re-opened for this edition.

### R08 Inherited evidence R08 GRADE C

Business in Cameroon. Publication date unrecorded. businessincameroon.com

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

### R09 Inherited evidence R09 GRADE C

Ecofin Agency. Publication date unrecorded. ecofinagency.com

Camtext SA integrated cotton-to-textile project, US$300 million (180 billion XAF); Panafritex and Arise with SODECOTON, CNPS and Marlo; 12 November 2025.

Inherited source-text check; original source not independently re-opened for this edition.

### R10 Inherited evidence R10 GRADE C

Ecofin Agency. Publication date unrecorded. ecofinagency.com

Camtext SA: 12,000 t a year processed; about 12,000 jobs at the Dibamba hub and 3,000 at Garoua; return on investment near 11 years; optimal capacity around 2033.

Inherited source-text check; original source not independently re-opened for this edition.

### R11 Inherited evidence R11 GRADE C

Business in Cameroon; MMI News. Publication date unrecorded. businessincameroon.com; mimimefoinfos.com

CICAM market share from about 85 percent three decades ago to about 5 percent today.

Inherited source-text check; original source not independently re-opened for this edition.

### R12 Inherited evidence R12 GRADE C

Business in Cameroon. Publication date unrecorded. businessincameroon.com

CICAM capacity utilisation, July 2023 report: Garoua 33.7 percent, CICAM 1 48.9 percent, CICAM 2 10.83 percent; underutilisation cost about CFAF 2 billion in 2020.

Inherited source-text check; original source not independently re-opened for this edition.

### R13 Inherited evidence R13 GRADE C

Business in Cameroon; MMI News. Publication date unrecorded. businessincameroon.com; mimimefoinfos.com

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

Cameroon minimum wage (SMIG) XAF 41,875 a month in 2023, raised to XAF 43,969 a month from 23 February 2024, about US$76, a rise of 5 percent.

Inherited source-text check; original source not independently re-opened for this edition.

### R15 Inherited evidence R15 GRADE C

wage.is. Publication date unrecorded. wage.is/cameroon

Average salary across sectors in Cameroon about US$359 a month.

Inherited source-text check; original source not independently re-opened for this edition.

### R16 Inherited evidence R16 GRADE C

WageIndicator. Publication date unrecorded. wageindicator.org

Informal share of employment about 90 percent; standard 40-hour week, 48-hour maximum, overtime at 1.2 times.

Inherited source-text check; original source not independently re-opened for this edition.

### R17 Inherited evidence R17 GRADE C

just-style. Publication date unrecorded. just-style.com

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

### R18 Inherited evidence R18 GRADE B

Congressional Research Service IF10149; AGOA (Wikipedia). Publication date unrecorded. congress.gov; en.wikipedia.org

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

### R23 Inherited evidence R23 GRADE C

MacMap; AfCFTA. Publication date unrecorded. macmap.org

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

### R24 Inherited evidence R24 GRADE C

Mordor Intelligence. Publication date unrecorded. mordorintelligence.com

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

### R25 Inherited evidence R25 GRADE C

Market.us. Publication date unrecorded. market.us

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

### R26 Inherited evidence R26 GRADE C

Statista. Publication date unrecorded. statista.com

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

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.

### R30 Inherited evidence R30 GRADE C

Design Journal (Taylor and Francis). Publication date unrecorded. tandfonline.com

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

### R31 Inherited evidence R31 GRADE C

arXiv 2509.04705. Publication date unrecorded. arxiv.org/abs/2509.04705

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

### R32 Inherited evidence R32 GRADE C

Wise; XE. Publication date unrecorded. wise.com; xe.com

CFA franc to US dollar reference rates; franc pegged to the euro at 655.957. The GFT study's method note places the franc near 590 to 610 per US dollar across the 2023 study window and near 563 to 577 more recently, and treats currency conversions as estimates.

Inherited source-text check; original source not independently re-opened for this edition.

### c01-S01 Enterprise Surveys: Cameroon 2024 Country Profile GRADE A

World Bank. 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-S03 SP 800-82 Rev. 3: Guide to Operational Technology (OT) Security GRADE A

NIST. 2023. 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-S05 SP 800-34 Rev. 1: Contingency Planning Guide for Federal Information Systems GRADE A

NIST. 2010. 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-S06 **SP 800-184: Guide for Cybersecurity Event Recovery** GRADE A

NIST. 2016. https://csrc.nist.gov/pubs/sp/800/184/final

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c01-S08 **RFC 4838: Delay-Tolerant Networking Architecture** GRADE A

IETF / RFC Editor. 2007. https://www.rfc-editor.org/rfc/rfc4838.txt

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c01-S13 **Atomic Commit In SQLite** GRADE A

SQLite. n.d.. https://sqlite.org/atomiccommit.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c01-S18 **RFC 7228: Terminology for Constrained-Node Networks** GRADE A

IETF / RFC Editor. 2014. https://www.rfc-editor.org/rfc/rfc7228.txt

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c01-S24 **RFC 9293: Transmission Control Protocol (TCP)** GRADE A

IETF / RFC Editor. 2022. https://www.rfc-editor.org/rfc/rfc9293.txt

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c02-S01 Same source record as c01-S01.

---

c02-S02 **Enterprise Surveys methodology** GRADE A

World Bank. 2026-10-30. https://www.enterprisesurveys.org/en/methodology

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c02-S03 Same source record as c01-S03.

---

c02-S05 Same source record as c01-S05.

---

c02-S06 Same source record as c01-S06.

---

c02-S08 Same source record as c01-S08.

c02-S10 **RFC 5905: Network Time Protocol Version 4: Protocol and Algorithms Specification** GRADE A

IETF / RFC Editor. 2010. https://www.rfc-editor.org/rfc/rfc5905.txt

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c02-S11 **RFC 8633: Network Time Protocol Best Current Practices** GRADE A

IETF / RFC Editor. 2019. https://www.rfc-editor.org/rfc/rfc8633.txt

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c02-S12 **RFC 9110: HTTP Semantics** GRADE A

IETF / RFC Editor. 2022. https://www.rfc-editor.org/rfc/rfc9110.txt

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c02-S13 **Atomic Commit In SQLite** GRADE B

SQLite. 2026-10-30. https://sqlite.org/atomiccommit.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c02-S14 **Write-Ahead Logging** GRADE B

SQLite. 2026-10-30. https://sqlite.org/wal.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c02-S16 **PRAGMA Statements** GRADE B

SQLite. 2026-10-30. https://sqlite.org/pragma.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c02-S18 Same source record as c01-S18.

c02-S21 **systemd.service** GRADE B

systemd project. 2026-10-30. https://www.freedesktop.org/software/systemd/man/latest/systemd.service.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c02-S22 **journald.conf** GRADE B

systemd project. 2026-10-30. https://www.freedesktop.org/software/systemd/man/latest/journald.conf.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

### c02-S23 RFC 9846: The Transport Layer Security (TLS) Protocol Version 1.3 GRADE A

IETF / RFC Editor. 2026. https://www.rfc-editor.org/rfc/rfc9846.txt

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c02-S24 Same source record as c01-S24.

---

### c02-S29 ext4 General Information GRADE B

Linux kernel project. 2026-10-30. https://docs.kernel.org/admin-guide/ext4.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

### c02-S30 Frequently Asked Questions GRADE B

chrony project. 2026-10-30. https://chrony-project.org/faq.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

### c02-S31 RFC 9562: Universally Unique IDentifiers (UUIDs) GRADE A

IETF / RFC Editor. 2024. https://www.rfc-editor.org/rfc/rfc9562.txt

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

### c02-S32 RFC 3339: Date and Time on the Internet: Timestamps GRADE A

IETF / RFC Editor. 2002. https://www.rfc-editor.org/rfc/rfc3339.txt

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

### c02-S0201 RFC 9520: Negative Caching of DNS Resolution Failures GRADE A

IETF / RFC Editor. 2023-12. https://www.rfc-editor.org/rfc/rfc9520.txt

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

### c02-S0202 RFC 6349: Framework for TCP Throughput Testing GRADE A

IETF / RFC Editor. 2011-08. https://www.rfc-editor.org/rfc/rfc6349.txt

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

### c02-S0203 RFC 2330: Framework for IP Performance Metrics GRADE A

IETF / RFC Editor. 1998-05. https://www.rfc-editor.org/rfc/rfc2330.txt

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c02-S0204 **RFC 9499: DNS Terminology** GRADE A

IETF / RFC Editor. 2024-03. https://www.rfc-editor.org/rfc/rfc9499.txt

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c03-S03 Same source record as c01-S03.

---

c03-S05 Same source record as c01-S05.

---

c03-S06 Same source record as c01-S06.

---

c03-S08 Same source record as c01-S08.

---

c03-S12 Same source record as c02-S12.

---

c03-S13 Same source record as c02-S13.

---

c03-S16 Same source record as c02-S16.

---

c03-S18 Same source record as c01-S18.

---

c03-S24 Same source record as c01-S24.

---

c03-S26 **The NIST Cybersecurity Framework (CSF) 2.0** GRADE A

NIST. 2024-02-26. https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.29.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c03-S27 **RFC 8915: Network Time Security for the Network Time Protocol** GRADE A

IETF / RFC Editor. 2020. https://www.rfc-editor.org/rfc/rfc8915.txt

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c03-S28 **How SQLite Is Tested** GRADE C

SQLite. 2026-10-30. https://sqlite.org/testing.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c03-S31 Same source record as c02-S31.

c03-S32 Same source record as c02-S32.

c04-S03 Same source record as c01-S03.

### c04-S04 SP 800-193: Platform Firmware Resiliency Guidelines GRADE A

NIST. 2018. https://csrc.nist.gov/pubs/sp/800/193/final

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c04-S05 Same source record as c01-S05.

c04-S06 Same source record as c01-S06.

c04-S13 Same source record as c02-S13.

### c04-S15 How To Corrupt An SQLite Database File GRADE C

SQLite. 2026-10-30. https://sqlite.org/howtocorrupt.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c04-S16 Same source record as c02-S16.

c04-S21 Same source record as c02-S21.

c04-S22 Same source record as c02-S22.

c04-S28 Same source record as c03-S28.

c04-S29 Same source record as c02-S29.

### c04-S0401 upsmon(8): UPS monitor and shutdown controller GRADE A

Network UPS Tools (NUT) Project. 2026-10-28. https://networkupstools.org/docs/man/upsmon.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c05-S01 Same source record as c01-S01.

c05-S03 Same source record as c01-S03.

c05-S05 Same source record as c01-S05.

c05-S06 Same source record as c01-S06.

c05-S08 Same source record as c01-S08.

### c05-S09 RFC 7252: The Constrained Application Protocol (CoAP) GRADE A

IETF / RFC Editor. 2014. https://www.rfc-editor.org/rfc/rfc7252.txt

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c05-S10 Same source record as c02-S10.

c05-S11 Same source record as c02-S11.

c05-S12 Same source record as c02-S12.

c05-S13 Same source record as c02-S13.

c05-S18 Same source record as c01-S18.

### c05-S19 MQTT Version 5.0 GRADE A

OASIS. 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.

### c05-S20 mosquitto.conf man page GRADE C

Eclipse Mosquitto. 2026-10-30. https://mosquitto.org/man/mosquitto-conf-5.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c05-S22 Same source record as c02-S22.

c05-S23 Same source record as c02-S23.

c05-S24 Same source record as c01-S24.

c05-S27 Same source record as c03-S27.

c05-S28 Same source record as c03-S28.

c05-S31 Same source record as c02-S31.

c05-S32 Same source record as c02-S32.

c06-S01 Same source record as c01-S01.

c06-S03 Same source record as c01-S03.

c06-S05 Same source record as c01-S05.

c06-S06 Same source record as c01-S06.

c06-S08 Same source record as c01-S08.

c06-S12 Same source record as c02-S12.

c06-S13 Same source record as c02-S13.

c06-S14 Same source record as c02-S14.

c06-S15 Same source record as c04-S15.

c06-S16 Same source record as c02-S16.

c06-S17 **SQLite Backup API** GRADE A

SQLite. 2026-10-30. https://sqlite.org/backup.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c06-S19 Same source record as c05-S19.

c06-S20 Same source record as c05-S20.

c06-S22 Same source record as c02-S22.

c06-S24 Same source record as c01-S24.

c06-S28 Same source record as c03-S28.

c06-S29 Same source record as c02-S29.

c06-S31 Same source record as c02-S31.

c06-S32 Same source record as c02-S32.

### c06-S0601 Powersafe Overwrite GRADE A

SQLite. 2025-05-31. https://sqlite.org/psow.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

### c06-S0602 Isolation In SQLite GRADE A

SQLite. 2022-04-18. https://sqlite.org/isolation.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

### c06-S0603 fsync(2) - Linux manual page GRADE A

Linux man-pages project. 2026-09-09. https://man7.org/linux/man-pages/man2/fsync.2.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

### c06-S21 systemd.service — Service unit configuration GRADE A

systemd. 2026-10-30. https://www.freedesktop.org/software/systemd/man/latest/systemd.service.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

### c06-S0604 Temporary Files Used By SQLite GRADE A

SQLite. 2025-06-12. https://sqlite.org/tempfiles.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c07-S03 Same source record as c01-S03.

c07-S05 Same source record as c01-S05.

c07-S06 Same source record as c01-S06.

c07-S07 **RFC 9171: Bundle Protocol Version 7** GRADE A

IETF / RFC Editor. 2022. https://www.rfc-editor.org/rfc/rfc9171.txt

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c07-S08 Same source record as c01-S08.

c07-S12 Same source record as c02-S12.

c07-S13 Same source record as c02-S13.

c07-S14 Same source record as c02-S14.

c07-S18 Same source record as c01-S18.

c07-S19 Same source record as c05-S19.

c07-S20 Same source record as c05-S20.

c07-S22 Same source record as c02-S22.

c07-S31 Same source record as c02-S31.

c08-S03 Same source record as c01-S03.

c08-S04 Same source record as c04-S04.

c08-S05 Same source record as c01-S05.

c08-S06 Same source record as c01-S06.

c08-S08 Same source record as c01-S08.

c08-S10 Same source record as c02-S10.

c08-S11 Same source record as c02-S11.

c08-S12 Same source record as c02-S12.

c08-S13 Same source record as c02-S13.

c08-S15 Same source record as c04-S15.

c08-S16 Same source record as c02-S16.

c08-S19 Same source record as c05-S19.

c08-S20 Same source record as c05-S20.

c08-S21 Same source record as c02-S21.

c08-S22 Same source record as c02-S22.

c08-S23 Same source record as c02-S23.

c08-S27 Same source record as c03-S27.

c08-S30 Same source record as c02-S30.

c08-S31 Same source record as c02-S31.

c08-S32 Same source record as c02-S32.

c08-S0801 **RFC 5280: Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile** GRADE A

IETF / RFC Editor. 2008-05. https://www.rfc-editor.org/rfc/rfc5280.txt

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c08-S0802 **RFC 7384: Security Requirements of Time Protocols in Packet Switched Networks** GRADE A

IETF / RFC Editor. 2014-10. https://www.rfc-editor.org/rfc/rfc7384.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-S03 Same source record as c01-S03.

c09-S04 Same source record as c04-S04.

c09-S05 Same source record as c01-S05.

c09-S06 Same source record as c01-S06.

c09-S08 Same source record as c01-S08.

c09-S12 Same source record as c02-S12.

c09-S13 Same source record as c02-S13.

c09-S14 Same source record as c02-S14.

c09-S15 Same source record as c04-S15.

c09-S16 Same source record as c02-S16.

c09-S17 Same source record as c06-S17.

c09-S18 Same source record as c01-S18.

c09-S19 Same source record as c05-S19.

c09-S20 Same source record as c05-S20.

c09-S21 Same source record as c02-S21.

c09-S22 Same source record as c02-S22.

c09-S23 Same source record as c02-S23.

c09-S29 Same source record as c02-S29.

c09-S30 Same source record as c02-S30.

c09-S31 Same source record as c02-S31.

c09-S32 Same source record as c02-S32.

### c09-S0901 upsmon(8) manual page GRADE A

Network UPS Tools. 2026-10-28. https://networkupstools.org/docs/man/upsmon.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

### c09-S0902 Prometheus Configuration Documentation GRADE A

Prometheus Authors. 2026-10-30. https://prometheus.io/docs/prometheus/latest/configuration/configuration/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c10-S03 Same source record as c01-S03.

c10-S04 Same source record as c04-S04.

c10-S05 Same source record as c01-S05.

c10-S06 Same source record as c01-S06.

c10-S10 Same source record as c02-S10.

c10-S12 Same source record as c02-S12.

c10-S13 Same source record as c02-S13.

c10-S14 Same source record as c02-S14.

c10-S15 Same source record as c04-S15.

c10-S16 Same source record as c02-S16.

c10-S17 Same source record as c06-S17.

c10-S18 Same source record as c01-S18.

c10-S20 Same source record as c05-S20.

c10-S21 Same source record as c02-S21.

c10-S22 Same source record as c02-S22.

c10-S23 Same source record as c02-S23.

### c10-S25 SP 800-61 Rev. 3: Incident Response Recommendations and Considerations for Cybersecurity Risk Management GRADE A

NIST. 2025. 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.

c10-S26 Same source record as c03-S26.

c10-S27 Same source record as c03-S27.

c10-S28 Same source record as c03-S28.

c10-S29 Same source record as c02-S29.

c10-S30 Same source record as c02-S30.

c10-S31 Same source record as c02-S31.

c10-S1001 Same source record as c08-S0801.

### c10-S1002 RFC 9019: A Firmware Update Architecture for Internet of Things GRADE A

IETF / RFC Editor. 2021-04. https://www.rfc-editor.org/rfc/rfc9019.txt

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

### c10-S1003 RFC 9325: Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS) GRADE A

IETF / RFC Editor. 2022-11. https://www.rfc-editor.org/rfc/rfc9325.txt

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

### c10-S1004 journalctl GRADE A

systemd project. 2026-10-30. https://www.freedesktop.org/software/systemd/man/latest/journalctl.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c10-S1005 **sd_id128_get_machine** GRADE A

systemd project. 2026-10-30. https://www.freedesktop.org/software/systemd/man/latest/sd_id128_get_boot.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c11-S03 Same source record as c01-S03.

c11-S04 Same source record as c04-S04.

c11-S05 Same source record as c01-S05.

c11-S06 Same source record as c01-S06.

c11-S12 Same source record as c02-S12.

c11-S13 Same source record as c02-S13.

c11-S17 Same source record as c06-S17.

c11-S21 Same source record as c02-S21.

c11-S22 Same source record as c02-S22.

c11-S31 Same source record as c02-S31.

c12-S01 Same source record as c01-S01.

c12-S02 Same source record as c02-S02.

c12-S03 Same source record as c01-S03.

c12-S04 Same source record as c04-S04.

c12-S05 Same source record as c01-S05.

c12-S06 Same source record as c01-S06.

c12-S08 Same source record as c01-S08.

c12-S12 Same source record as c02-S12.

c12-S13 Same source record as c02-S13.

c12-S14 Same source record as c02-S14.

c12-S15 Same source record as c04-S15.

c12-S16 Same source record as c02-S16.

c12-S17 Same source record as c06-S17.

c12-S18 Same source record as c01-S18.

c12-S19 Same source record as c05-S19.

c12-S20 Same source record as c05-S20.

c12-S21 Same source record as c02-S21.

c12-S22 Same source record as c02-S22.

c12-S23 Same source record as c02-S23.

c12-S24 Same source record as c01-S24.

c12-S26 Same source record as c03-S26.

c12-S27 Same source record as c03-S27.

c12-S28 Same source record as c03-S28.

c12-S29 Same source record as c02-S29.

c12-S30 Same source record as c02-S30.

c12-S31 Same source record as c02-S31.

c12-S32 Same source record as c02-S32.

c12-S1201 **SP 800-161 Rev. 1: Cybersecurity Supply Chain Risk Management Practices for Systems and Organizations** GRADE A

NIST. 2024-11-01. https://csrc.nist.gov/pubs/sp/800/161/r1/upd1/final

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c12-S1202 **Command Line Shell For SQLite** GRADE B

SQLite. 2026-10-30. https://sqlite.org/cli.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c12-S1203 **IR 8286D: Using Business Impact Analysis to Inform Risk Prioritization and Response** GRADE A

NIST. 2022-11. https://csrc.nist.gov/pubs/ir/8286/d/final

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c13-S03 Same source record as c01-S03.

c13-S04 Same source record as c04-S04.

c13-S05 Same source record as c01-S05.

c13-S06 Same source record as c01-S06.

c13-S12 Same source record as c02-S12.

c13-S13 Same source record as c02-S13.

c13-S15 **How To Corrupt An SQLite Database File** GRADE A

SQLite. 2026-04-13. https://sqlite.org/howtocorrupt.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c13-S16 Same source record as c02-S16.

c13-S19 Same source record as c05-S19.

c13-S21 Same source record as c02-S21.

c13-S26 Same source record as c03-S26.

c13-S28 **How SQLite Is Tested** GRADE A

SQLite. 2026-04-21. https://sqlite.org/testing.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c13-S29 Same source record as c02-S29.

c13-S31 Same source record as c02-S31.

c13-S32 Same source record as c02-S32.

c14-S01 Same source record as c01-S01.

c14-S02 Same source record as c02-S02.

c14-S03 Same source record as c01-S03.

c14-S04 Same source record as c04-S04.

c14-S05 Same source record as c01-S05.

c14-S06 Same source record as c01-S06.

c14-S12 Same source record as c02-S12.

c14-S13 Same source record as c02-S13.

c14-S16 Same source record as c02-S16.

c14-S17 Same source record as c06-S17.

c14-S19 Same source record as c05-S19.

c14-S26 Same source record as c03-S26.

c14-S28 Same source record as c03-S28.

c14-S30 Same source record as c02-S30.

c14-S32 Same source record as c02-S32.

# Stable definitions make comparisons useful.

### Acknowledgement boundary

The point at which a component declares acceptance, with its validation, persistence and responsibility explicitly defined.

### Application acceptance

A receiving application has validated and durably committed the event under the agreed contract.

### Atomic commit

All changes in a database transaction take effect together or none do; durability depends on the configured storage path.

### Backhaul

The external network path connecting the site network to upstream services.

### Backlog

Accepted work retained pending transmission, processing or remote durable acceptance.

### Backpressure

A controlled signal or limit that slows intake when downstream capacity is constrained.

### Boot identity

An identifier distinguishing events produced during a particular device or process start interval.

### Brownout

A reduction of supply voltage whose effect depends on equipment and protection; record it separately from complete supply loss.

### Checkpoint

An operation that transfers or consolidates persisted state, such as moving WAL contents into a main database file.

### Clock holdover

Operation of a local clock without its usual external reference, with uncertainty that must be assessed.

### Clock quality

Metadata describing the synchronisation state and uncertainty attached to a recorded time.

### CoAP

Constrained Application Protocol, designed for constrained devices and networks.

### Cold boot

Starting a device from an unpowered state rather than merely restarting an application.

### Common-cause failure

A shared dependency causes apparently separate services or paths to fail together.

### Contribution margin

Revenue less the variable costs associated with the relevant additional units, using an explicit accounting boundary.

### Data freshness

The age of the newest trustworthy information relative to the decision being made.

### Deduplication

Detection and controlled suppression of repeated processing of the same identified event.

### Degraded mode

An explicitly permitted reduced service with visible limits, recording and exit conditions.

### Durable commit

A committed change protected against the declared failure types by the complete software and storage path.

### Edge inference

Executing a trained model near the data source; local execution alone does not prove offline operability.

### Event envelope

The payload and contextual fields needed to identify, order, validate and reconcile an event.

### Goodput

Useful application payload successfully accepted per unit of elapsed time, distinguished from nominal link rate.

### Graceful shutdown

An orderly termination that attempts to commit and close state before power or execution ends.

### Idempotency

Repeating an operation has the same intended effect as carrying it out once, under the specified contract.

### Ingest time

The time a receiving component records its acceptance of an event, distinct from source event time.

### Local-first operation

A proposed operating pattern in which required site duties and records remain available within a defined local boundary.

### MQTT

A client-server publish/subscribe messaging transport with specified delivery qualities and session behaviour.

### NTP

Network Time Protocol, used to synchronise clocks across networked systems.

### NTS

Network Time Security, a cryptographic mechanism for network time synchronisation.

### Outbox

A durable set of pending-delivery records committed with the corresponding local business events.

### Outage envelope

The declared combinations, durations and consequences of failures against which a design is assessed.

### Protected digital load

The defined equipment and auxiliaries supplied by a continuity power path for specified digital duties.

### Reconciliation

Comparing retained records and resolving missing, duplicated or conflicting events with an auditable decision.

### Recovery point objective

The acceptable boundary of data loss after a disruption, defined for the relevant service and record class.

### Recovery time objective

The agreed time within which a specified service or operation should be restored after a disruption.

### Replay

Resending or reprocessing retained events after interruption, subject to deduplication and acknowledgement rules.

### Service readiness

The declared condition in which a running service can perform its promised function, rather than merely exist as a process.

### Store and forward

Retaining data durably until a suitable onward communication opportunity is available.

### Transfer interval

The interruption or changeover interval between power sources, to be qualified against the connected equipment.

### UPS

Uninterruptible power supply; continuity behaviour must be checked for the declared load and failure envelope.

### UUID

A universally unique identifier with a defined format; useful identity depends on correct generation and handling.

### WAL

Write-ahead log, a database journalling approach with implementation-specific consistency and durability properties.

### Write barrier

A storage ordering mechanism used to preserve required ordering of writes and journal commits.

# 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.

**Acknowledgement boundary** 20 207

**Application acceptance** 220 223

**Atomic commit** 32 65 89 113 128 176 221

**Backhaul** 28 32 34 40 70 72 75 78 83 84 87 100 101 107 108 125 126 127 129 138 143 144 156 157 158 163 164 187 192 193 194 198 219

**Backlog** 20 34 35 38 70 73 74 79 80 84 88 100 101 108 115 116 117 120 126 188 192 197 208 209 223

**Backpressure** 120 122 197 210

**Boot identity** 114 125 177 198 212

**Brownout** 31 70 162

**Checkpoint** 36 92 94 95 97 99 101 104 106 107 112

**Clock holdover** 251

**Clock quality** 88 156 172 223

**CoAP** 75 76 77 80

**Cold boot** 31 37 48 49 54 82 88 125 133 143 144 159 160 169 196 207 223

**Common-cause failure** 70 195

**Contribution margin** 186 189 190 194

**Data freshness** 251

**Deduplication** 74 107 114 136 159 196 197 207 208 216 223 229

**Degraded mode** 252

**Durable commit** 32 38 70 73 173 204

**Edge inference** 252

**Event envelope** 126

**Goodput** 32 34 38 70 73 74 192

**Graceful shutdown** 252

**Idempotency** 100 114 115

**Ingest time** 125

**Local-first operation** 252

**MQTT** 73 74 76 77 80 82 98 99 100 114 151 156 158 168 191 196 204 208 223

**NTP** 31 35 83 84 126 127 129 131 133 143 144 159 160 169

**NTS** 82 84 133 134 159 160 161 168 169 192

**Outage envelope** 34 60 118 158 220

**Outbox** 32 34 36 38 77 83 87 91 98 101 102 104 108 112 113 128 135 141 147 149 151 156 163 164 175 176 179 193 196 197 199 204 206 207 209 210 221 224

**Protected digital load** 38 39 40 68 143

**Reconciliation** 20 21 22 23 30 53 88 106 113 114 115 116 121 122 125 132 133 136 138 159 174 177 178 179 180 182 183 187 189 194 195 199 207 209 213 215 218 220 221 223 224 227

**Recovery point objective** 252

**Recovery time objective** 83 144 191 196

**Replay** 20 37 63 74 76 79 80 105 106 111 113 114 115 118 121 126 127 135 137 138 147 157 159 160 162 163 168 172 175 176 177 179 180 182 183 192 197 204 206 209 210 213 216 220 223 224 229

**Service readiness** 63 211

**Store and forward** 253

**Transfer interval** 31

**UPS** 28 31 34 36 38 57 59 60 61 64 65 66 68 71 78 79 93 136 143 144 148 151 193 194 198 199

**UUID** 31 53 76 88 100 114 127 188 191 193 196 199 200

**WAL** 31 35 36 37 45 63 92 104 111 144 145 147 150 151 162 168 191 199 206 221

**Write barrier** 253

# 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 wp07. 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.*