



AF 9100:2026

Aerospace Quality Management System

TECHNICAL PROFILE BOOK

Software, AI and Digital Engineering — Technical Application Book

AF-TP-9100-002

FIRST EDITION

30 September 2026 | English master text

Publication identity and application

Control	Publication value
Title	Software, AI and Digital Engineering — Technical Application Book
Document code	AF-TP-9100-002
Edition / language	FIRST EDITION / English master text
Companion edition identity	30 September 2026
Controlling source	AF 9100:2026 — FIRST EDITION — 29 September 2026
Responsible owner function	Digital design authority
Primary source anchors	12, N, D
Delivery status	Prepared implementation companion. Actual approval, adoption, effectivity and operational decisions are recorded separately.

Use and document-control boundary

This companion implements the supplied AF 9100 requirement basis without changing the original standard. Its process allocation, examples, checklists and record design are authored implementation material. Actual mission or product limits, specialist procedures, permissions and genuine observations must be completed before consequential use. Applicable source minima and independent decision conditions cannot be reduced by this companion.

The printed edition date identifies the companion series; it is not evidence that later preparation, review or operational decisions occurred on that date. New support codes in this set are proposed allocations for controlled registration. Established AF-IG-9100-001, AF-FAQ-9100-001 and AF-TR-9100-001 identities are retained; earlier files, the corporate master list and ASTRA codes are not overwritten.

Equivalent combined or digital records are permitted when all required fields, relationships, authority and history remain intact. A completed PDF field or typed name is not evidence of an authenticated approval. Blank records in this delivery contain no invented applicant, signature, event or result.

Contents and reading route

Publication identity and application	2
Contents and reading route	3
1 Scope, responsibility and source priority	4
2 Technical profile and prospective parameter completion	4
3 Subject-by-subject implementation	4
4 Representative evaluation and adverse evidence	6
5 Interface records, decisions and continuing use	6
Appendix A — Controlling source and evidence	7
Source tables — complete field and decision payloads	11
Appendix B — Local implementation and release	14

Select actual scope → complete application criteria → use the applicable procedure and technical book → capture authentic evidence → evaluate every applicable condition → obtain the required independent decision. Document availability does not establish implementation or conformity.

1 Scope, responsibility and source priority

Software and programmable logic, generated artifacts, datasets, models, digital twins, intended-use validation and long-life digital support.

Boundary	Required implementation decision
Object and route	Identify ORG, PRD, SVC, PRC, PER, SUB or CAB separately. Combined records do not merge outcomes.
Activity	Apply source technical modules by performed or retained responsibility, including relevant outsourced work.
Authority	Technical release, employer task permission, independent attestation and statutory authorization remain distinct.
Source priority	AF 9100:2026 and all compatible binding obligations control. This companion does not import ASTRA programme rules or external normative schemes.

2 Technical profile and prospective parameter completion

Target/configuration; source/build identity; data partitions; domain; essential scenarios; timing/resource limits; error metrics; drift and adaptation bounds

Parameter record	Content required before evaluating results
Identity and trigger	Stable criterion ID, source anchor, owner, actual object, phase, mode and reason it applies.
Limit or attribute	Value/range with compatible units and boundary convention, or an observable state with every essential step.
Technical basis	Actual environment, duty, consequence, derivation, qualified evidence and relevant approval.
Method and representation	Inspection, analysis, demonstration or test; method version, article/population, configuration and validity limits.
Uncertainty and decision	Corrections, model, covariance, coverage, pre-agreed rule, sample/scoring logic and hold conditions.
Authority and change	Technical and independent approvals, actual effectivity, preserved earlier basis and reopening triggers.

The fields above are completed for the actual application. A blank limit, undefined domain or unvalidated method prevents the dependent claim; the book deliberately does not invent one severity, tolerance or task threshold for unrelated applications.

3 Subject-by-subject implementation

12.1 — Digital lifecycle and architectural assurance

Source p. 29; complete conditions reproduced in Appendix A.

Implementation record: Digital lifecycle plan; allocated requirements; architecture; fault-containment and timing analyses.

Application check: Use traced tests, boundary and failure conditions, resource and timing cases, meaningful coverage and independent validation data. A high aggregate score does not close a failed essential scenario.

Review boundary: Unknown target configuration, invalid evaluation or failed essential protection blocks consequential deployment. Record each source condition separately in the linked checklist and preserve approved exclusions.

12.2 — Implementation, review and build control

Source p. 29; complete conditions reproduced in Appendix A.

Implementation record: Source and review records; build manifest; dependency inventory; reproducibility and artifact identity checks.

Application check: Use traced tests, boundary and failure conditions, resource and timing cases, meaningful coverage and independent validation data. A high aggregate score does not close a failed essential scenario.

Review boundary: Unknown target configuration, invalid evaluation or failed essential protection blocks consequential deployment. Record each source condition separately in the linked checklist and preserve approved exclusions.

12.3 — Verification, test coverage and independence

Source p. 29; complete conditions reproduced in Appendix A.

Implementation record: Verification plan; traceable tests; coverage analysis; anomaly disposition; target-representativeness assessment.

Application check: Use traced tests, boundary and failure conditions, resource and timing cases, meaningful coverage and independent validation data. A high aggregate score does not close a failed essential scenario.

Review boundary: Unknown target configuration, invalid evaluation or failed essential protection blocks consequential deployment. Record each source condition separately in the linked checklist and preserve approved exclusions.

12.4 — AI and machine-learning use-case governance

Source p. 30; complete conditions reproduced in Appendix A.

Implementation record: AI use-case inventory; intended-use statement; alternatives assessment; authority limits; oversight validation.

Application check: Use traced tests, boundary and failure conditions, resource and timing cases, meaningful coverage and independent validation data. A high aggregate score does not close a failed essential scenario.

Review boundary: Unknown target configuration, invalid evaluation or failed essential protection blocks consequential deployment. Record each source condition separately in the linked checklist and preserve approved exclusions.

12.5 — Dataset quality, lineage and evaluation design

Source p. 30; complete conditions reproduced in Appendix A.

Implementation record: Dataset manifest; labeling assurance; partition controls; scenario coverage; stratified evaluation and uncertainty.

Application check: Use traced tests, boundary and failure conditions, resource and timing cases, meaningful coverage and independent validation data. A high aggregate score does not close a failed essential scenario.

Review boundary: Unknown target configuration, invalid evaluation or failed essential protection blocks consequential deployment. Record each source condition separately in the linked checklist and preserve approved exclusions.

12.6 — Model release, drift and bounded adaptation

Source p. 30; complete conditions reproduced in Appendix A.

Implementation record: Model release card; drift thresholds and rationale; fallback tests; update approvals; monitoring records.

Application check: Use traced tests, boundary and failure conditions, resource and timing cases, meaningful coverage and independent validation data. A high aggregate score does not close a failed essential scenario.

Review boundary: Unknown target configuration, invalid evaluation or failed essential protection blocks consequential deployment. Record each source condition separately in the linked checklist and preserve approved exclusions.

12.7 — Generative AI and engineering evidence

Source p. 30; complete conditions reproduced in Appendix A.

Implementation record: Approved use cases; data-access rules; review records; verified calculations and references.

Application check: Use traced tests, boundary and failure conditions, resource and timing cases, meaningful coverage and independent validation data. A high aggregate score does not close a failed essential scenario.

Review boundary: Unknown target configuration, invalid evaluation or failed essential protection blocks consequential deployment. Record each source condition separately in the linked checklist and preserve approved exclusions.

12.8 — Digital twins, MBSE and model-based acceptance

Source p. 30; complete conditions reproduced in Appendix A.

Implementation record: Model authority map; twin identity and update controls; validation; model-based acceptance rationale.

Application check: Use traced tests, boundary and failure conditions, resource and timing cases, meaningful coverage and independent validation data. A high aggregate score does not close a failed essential scenario.

Review boundary: Unknown target configuration, invalid evaluation or failed essential protection blocks consequential deployment. Record each source condition separately in the linked checklist and preserve approved exclusions.

12.9 — Long-term digital maintainability and obsolescence

Source p. 30; complete conditions reproduced in Appendix A.

Implementation record: Digital support plan; escrow or access arrangements where needed; migration tests; obsolescence decisions.

Application check: Assign custodians, lawful rights, readable formats, technical dictionaries and replacement expertise. Test whether receiving users can interpret the evidence without the original author.

Review boundary: Loss of essential support knowledge or source access holds the claim that depends on it until a valid alternative exists. Record each source condition separately in the linked checklist and preserve approved exclusions.

4 Representative evaluation and adverse evidence

Software, AI, Models and Trusted Updates

Witness or trace: Assess trust, privileges, vulnerabilities, OT compatibility and safety-security conflicts. Intrusive tests require written safe authorization; signatures alone do not establish technical suitability.

Synthetic challenge: A model's overall accuracy is high but a separately tested difficult defect set has low recall. Retain both populations and limitations. Resolve the critical domain or enforce a technically validated restriction; do not pool it into a favorable overall score.

Technical Tools, Data Integrity and Evidence Portability

Witness or trace: Reconcile ingestion to the actual source, object, version and event time. Identify generated, simulated and reconstructed material; a hash proves neither physical truth nor engineering validity.

Synthetic challenge: A migration copies every file, but a test report links to the wrong serial. Reject semantic reconciliation despite matching file counts. Repair relationships, repeat the failing case and review conclusions already relying on the incorrect link.

5 Interface records, decisions and continuing use

Record	Required implementation purpose
AF-FR-9100-019	Software, AI, Model and Technical-Tool Release
AF-FR-9100-029	Cyber Incident, Update and Assured Recovery

Qualification, instance acceptance, current operating capability and independent attestation retain separate evidence and decisions. A new site, material, supplier, digital load, use environment or duty requires a relevant difference and requalification review. No previous result transfers automatically because the name is unchanged.

Appendix A — Controlling source and evidence

Source: AF 9100:2026, Aerospace Quality Management System, FIRST EDITION, 29 September 2026; supplied 131-page PDF. Source page references below locate the original publication. These excerpts are read with the complete applicable standard; local implementation sequences do not replace it.

12.1 Digital lifecycle and architectural assurance

SOURCE p. 29 | Normative when applicable

12.1(a) The organization shall establish a lifecycle for software and programmable hardware covering requirements, architecture, implementation, integration, verification, release, maintenance and retirement. Safety and mission allocations shall be reflected in architecture and verification independence.

12.1(b) The architecture shall address partitioning, timing, resource exhaustion, fault containment, recovery and interfaces. An assumption that software does not wear out shall not be used to dismiss systematic defects, data corruption, configuration errors or changes in the environment.

Required evidence: Digital lifecycle plan; allocated requirements; architecture; fault-containment and timing analyses.

12.2 Implementation, review and build control

SOURCE p. 29 | Normative when applicable

12.2(a) Source code, programmable logic, generated code and build definitions shall be version controlled with attributable changes. Review and analysis shall be suitable for the language, platform and criticality. Coding or modeling constraints shall be justified and enforced by competent review and appropriate tools.

12.2(b) Builds shall identify source revisions, dependencies, toolchain versions, compiler settings, configuration and resulting artifact hashes. Reproducibility shall be demonstrated where necessary to establish that the released binary corresponds to reviewed sources. Secrets and unapproved dependencies shall not be embedded in released artifacts.

Required evidence: Source and review records; build manifest; dependency inventory; reproducibility and artifact identity checks.

12.3 Verification, test coverage and independence

SOURCE p. 29 | Normative when applicable

12.3(a) Digital verification shall address requirements, interfaces, boundary values, abnormal inputs, timing, failure handling and regression. Structural or functional coverage criteria shall be chosen based on the assurance rationale and applicable domain obligations, not on a universal percentage in this standard.

12.3(b) Test results shall identify the target environment and configuration. Differences between simulation, host execution and deployed hardware shall be assessed. Uncovered behavior, unresolved anomalies and tool limitations shall be evaluated explicitly before release.

Required evidence: Verification plan; traceable tests; coverage analysis; anomaly disposition; target-representativeness assessment.

12.4 AI and machine-learning use-case governance

SOURCE p. 30 | Normative when applicable

12.4(a) Every AI use case affecting assured outcomes shall have an accountable owner, intended use, prohibited use, operational domain, performance criteria, hazard assessment and monitoring plan. The organization shall assess whether AI is necessary and whether simpler, more verifiable alternatives meet the need.

12.4(b) Authority granted to the model shall be bounded. Human oversight shall be designed so that the responsible person has adequate information, time and ability to intervene. A nominal human approval step shall not be treated as effective control when the person cannot meaningfully review the output.

Required evidence: AI use-case inventory; intended-use statement; alternatives assessment; authority limits; oversight validation.

12.5 Dataset quality, lineage and evaluation design

SOURCE p. 30 | Normative when applicable

12.5(a) Datasets shall have documented provenance, rights, collection methods, labeling procedures, relevant populations, environmental coverage and known limitations. Training, tuning and evaluation partitions shall be controlled to prevent leakage or unjustified overlap.

12.5(b) Evaluation shall include relevant rare conditions, degraded sensors, out-of-domain inputs and important subgroups. Performance shall be reported with denominators, uncertainty and failure analysis. Aggregate accuracy shall not conceal unacceptable failure in a critical operating condition. Annex N specifies a minimum model and dataset evidence card.

Required evidence: Dataset manifest; labeling assurance; partition controls; scenario coverage; stratified evaluation and uncertainty.

12.6 Model release, drift and bounded adaptation

SOURCE p. 30 | Normative when applicable

12.6(a) Model weights, preprocessing, inference configuration, dependencies and decision thresholds shall be controlled as a single released configuration. The organization shall demonstrate safe behavior for invalid, missing, adversarial or out-of-distribution inputs within the defined threat and hazard model.

12.6(b) Post-release monitoring shall detect relevant drift, performance degradation and unexpected behavior. Online learning or autonomous adaptation shall not change safety-relevant behavior outside an approved envelope. Material changes shall undergo a new impact assessment, verification and authorized release.

12.6(c) Fallback and rollback shall be technically feasible and tested. A model shall not be retained merely because replacing it would disrupt operations when its approved performance basis is no longer valid.

Required evidence: Model release card; drift thresholds and rationale; fallback tests; update approvals; monitoring records.

12.7 Generative AI and engineering evidence

SOURCE p. 30 | Normative when applicable

12.7(a) Generative AI used in engineering, procurement, inspection, coding or technical documentation shall have controlled inputs, access and review requirements. Generated claims, citations, calculations and design details shall be verified by competent personnel before becoming controlled evidence.

12.7(b) The organization shall prevent unauthorized disclosure of proprietary or controlled technical data to external services. Prompt history or a generated explanation shall not substitute for a validated calculation, traceable test or approved design rationale. The human or organizational approval responsibility shall remain identifiable.

Required evidence: Approved use cases; data-access rules; review records; verified calculations and references.

12.8 Digital twins, MBSE and model-based acceptance

SOURCE p. 30 | Normative when applicable

12.8(a) Model-based engineering artifacts shall have defined authority, scope, interfaces, configuration and exchange rules. A digital twin shall state its decision purpose, physical association, update mechanism and validated domain. Synchronization errors and stale data shall be detectable.

12.8(b) When a model replaces or reduces physical testing, the organization shall provide a documented equivalence argument, uncertainty assessment, validation evidence and required independent approval. Model convenience shall not be sufficient justification for removing a test.

Required evidence: Model authority map; twin identity and update controls; validation; model-based acceptance rationale.

12.9 Long-term digital maintainability and obsolescence

SOURCE p. 30 | Normative when applicable

12.9(a) The organization shall plan preservation of source, build environments, licenses, tools and technical knowledge needed to maintain or reconstruct long-life digital products. Dependence on a discontinued cloud service or supplier shall be assessed before release.

12.9(b) Obsolescence plans shall address migration, emulation, replacement, security support and requalification. A digitally signed record shall remain interpretable after algorithm or certificate changes through an approved preservation strategy; cryptographic signing alone shall not guarantee permanent readability.

Required evidence: Digital support plan; escrow or access arrangements where needed; migration tests; obsolescence decisions.

D.1 Requirement-record integrity

SOURCE p. 55 | Normative when applicable

D.1(a) Every controlled product or service requirement shall have an identifier, statement, source, owner, revision, rationale, applicability, verification method and acceptance criteria. Derived and safety-related requirements shall include their parent link and approving authority.

D.1(b) The organization shall maintain traceability through allocation, implementation, verification, anomaly resolution and release. Traceability completeness shall be evaluated for meaning and validity, not only the presence of hyperlinks. Broken or circular links shall be investigated.

D.2 Identifiers, versioning and time

SOURCE p. 55 | Normative when applicable

D.2(a) Identifiers shall remain unique within the defined repository and resolvable after migration. Revision and status shall be distinct; an approved obsolete version shall not appear current. Time records shall identify the relevant time basis and sequence sufficient to reconstruct events.

D.2(b) Units, coordinate systems, reference frames and data encodings shall be explicit where misinterpretation could affect conformity. Imported supplier data shall be checked for mapping errors, truncated precision and semantic differences. Automated exchange shall retain validation results.

D.3 Evidence provenance and tamper indication

SOURCE p. 55 | Normative when applicable

D.3(a) An evidence record shall identify the source system or person, acquisition method, processing, transformations and approvals. Cryptographic hashes or signatures may support integrity but shall not be treated as proof that the original observation was correct.

D.3(b) Corrections shall preserve the prior evidence and explain why it changed. Access and modification logs shall be protected in accordance with consequence and obligations. Migration shall verify relationships as well as files; a readable PDF without its underlying configuration links may be insufficient. The filled chain demonstrates structure only. It deliberately does not supply a universal stopping-time value, load or speed; those values require a valid system-specific safety basis.

N.1 Digital release card

SOURCE p. 75 | Normative when applicable

N.1(a) Each assured digital release shall have the fields in Table N-1 or an equivalent linked record. The card shall identify the exact evaluated configuration and the limits of the claim. Missing critical evidence shall block release rather than being masked by a summary score.

N.1(b) The release decision shall address unresolved anomalies, security exposure, compatibility, installation, rollback, support and required independent approval. Subsequent updates shall retain lineage to the prior release.

N.2 AI model and dataset cards

SOURCE p. 75 | Normative when applicable

N.2(a) An AI assurance card shall identify use case, human authority, prohibited uses, operational domain, data lineage, evaluation design, performance criteria and monitoring. Training and test-data provenance shall be sufficiently documented to evaluate representativeness and rights.

N.2(b) The organization shall retain important subgroup and scenario results, uncertainty, known failure modes, drift triggers and fallback evidence. An aggregate accuracy figure shall not replace this evidence.

N.3 Acceptance and change-control rules

SOURCE p. 75 | Normative when applicable

N.3(a) Performance criteria shall be established before final evaluation and linked to consequences. A model shall meet each mandatory criterion for the intended domain; superior performance on one metric shall not compensate for an unacceptable critical failure rate.

N.3(b) A change to model weights, data, preprocessing, threshold, sensor or deployment environment shall trigger an impact assessment. Safety-relevant adaptation outside the approved envelope shall require renewed assurance. The card shall be updated before the changed configuration is released.

N.4 Informative confusion-matrix example

SOURCE p. 76 | Definition, explanation or informative example; no new universal limit

N.4(a) A synthetic visual-inspection model is evaluated on 1,000 independent labeled cases: 190 true positives, 15 false positives, 10 false negatives and 785 true negatives. Here “positive” means the defect class. These values are examples, not AF or supplier test results.

N.4(b) Precision= $190/(190+15)=92.6829\%$; recall= $190/(190+10)=95.0000\%$; specificity= $785/(785+15)=98.1250\%$; accuracy= $(190+785)/1,000=97.5000\%$; F1=0.938272. The 10 missed defects remain important even when aggregate accuracy is high.

Source calculation note: (N-1) Precision = TP/(TP+FP); Recall = TP/(TP+FN); F1 = 2PR/(P+R) P and R in the F1 expression denote precision and recall, not probability of product safety. Report undefined metrics explicitly when a denominator is zero.

N.5 Informative uncertainty and challenge-set example

SOURCE p. 76 | Definition, explanation or informative example; no new universal limit

N.5(a) For the 190 successes in 200 defect cases, a two-sided exact 95% binomial interval for recall is approximately [90.997%, 97.577%]. The Wilson interval is approximately [91.042%, 97.262%]. The selected interval method and its assumptions should be reported.

N.5(b) A separate challenge set—not a subset of the 1,000-case evaluation—contains 50 difficult defect cases, of which 38 are detected and 12 missed. Its recall is 76%, with a two-sided exact 95% lower endpoint of about 61.831%. The weak challenge result cannot be hidden by the 97.5% overall accuracy from the other dataset.

N.5(c) An illustrative program criterion might require a lower confidence bound of at least 90% for a particular critical scenario. That hypothetical criterion would not be met by the challenge set. The actual threshold and confidence method must be derived from the program’s hazard and decision basis; this example does not prescribe either.

N.6 Informative release interpretation

SOURCE p. 76 | Definition, explanation or informative example; no new universal limit

N.6(a) A reviewer should ask whether the model was tested on the actual defect mechanisms, materials, lighting, sensor condition and operating environment. The statistical interval describes sampling uncertainty under the model; it does not cover all uncertainty from bad labels, distribution shift or incomplete scenario coverage.

N.6(b) The release could be restricted to a demonstrated domain with a validated independent inspection path for excluded conditions. Such a restriction needs operational enforcement and monitoring; merely writing a limitation in a model card is not a physical safeguard.

Source tables — complete field and decision payloads

Table D-1 — Required lifecycle data entities

Source p. 55 | Required structure when applicable

Entity	Minimum fields	Required relationships
Requirement	ID; revision; statement; source; rationale; criticality; owner; acceptance criterion; status.	Parent need; allocated item/function; verification case; applicable obligation.
Configuration item	ID; type; revision; physical/software identity; owner; baseline; approved use.	Requirements; BOM; interfaces; change history; as-built or deployed instances.
Material or component lot	Source; manufacturer; lot/batch; quantity; condition; certificate; received date.	Incoming verification; process route; consumed serials; split/merge history.
Verification case	ID; method; requirement; configuration; procedure; environment; uncertainty; acceptance rule.	Test article; equipment; raw data; result; anomaly and disposition.
Process execution	Route revision; operation; performer; date/time; equipment; parameters; material; result.	Lot/serial; qualified procedure; inspection; nonconformity; next authorized step.
Change	ID; reason; baseline before/after; impact; approvals; effectivity; verification.	Affected requirements/items; delivered population; suppliers; support information.
Hazard or risk	ID; scenario; consequence; likelihood basis; exposure; controls; owner; residual decision.	Requirements; mitigations; verification evidence; operational limitations.
Release	ID; configuration; scope; acceptance evidence; concessions; restrictions; authorized identity; date.	Requirements baseline; product serial/lot; delivery dossier; service-life status.
Digital artifact	Artifact ID; version; source/build; dependencies; hash; signature; provenance; access class.	Requirements; verification; release; deployment instances; vulnerabilities.
Retirement	Asset identity; condition; hazardous inventory; reuse restrictions; disposition; verifier.	Life history; disposal evidence; retained records; remaining obligations.

Table D-2 — Filled illustrative trace chain

Source p. 55, 56 | Illustrative or explanatory source table

Record	Example value	Connection
Need	NEED-GSC-01: controlled motion during payload handling.	Allocated to SYS-GSC-R12.
Requirement	SYS-GSC-R12 rev B: stop response within a program-defined limit under stated load and speed.	Hazard H-GSC-02; verification V-GSC-12.
Configuration	Controller CTRL-07 rev C; firmware 2.4.1; cart serial GSC-004.	Build manifest BM-241; interface ICD-09 rev D.

Record	Example value	Connection
Verification	V-GSC-12 rev 3; test article GSC-004; report TR-012.	Raw records RAW-012; calibrated timer TM-08; approved decision rule DR-03.
Disposition	Anomaly AN-04 closed after independent retest.	NCR-026; verified corrective action CA-026.
Release	REL-GSC-004 rev 1; approved operating- domain restriction attached.	Release authority RA-02; receiving acknowledgement ACK-004.

Table N-1 — Software and programmable-hardware release card

Source p. 75 | Required structure when applicable

Field group	Required content
Identity	Product; version; source revision; build identity; target hardware; relevant programmable logic.
Intended function	Allocated requirements; safety/mission criticality; operating modes and domain.
Architecture	Interfaces; partitioning; timing; resources; fault containment and recovery.
Build and dependencies	Toolchain; compiler/settings; libraries; SBOM; hashes; reproducibility evidence.
Verification	Requirements coverage; test environment; structural/functional criteria; robustness; target differences.
Anomalies	Open and closed defects; severity; affected configuration; mitigations; acceptance authority.
Security	Threat-model linkage; vulnerability decisions; secure configuration; update and credential controls.
Deployment and support	Installation checks; compatibility; rollback; monitoring; maintenance and end-of-support.
Approval	Technical owner; independent reviewer where required; release authority; date and restrictions.

Table N-2 — AI-specific evidence card

Source p. 75 | Required structure when applicable

Element	Required content
Purpose and authority	Decision supported; actions allowed; human responsibilities; prohibited uses; affected people and systems.
Domain	Environment; sensors; input ranges; relevant populations; timing; degraded and excluded conditions.
Data	Sources; rights; collection; labeling; quality checks; partitions; leakage prevention; coverage and limitations.
Model	Architecture/type; version; training configuration; preprocessing; thresholds; dependencies; reproducibility.
Evaluation	Independent test design; representative scenarios; rare/challenge cases; metrics; denominators; uncertainty.
Robustness	Invalid/missing inputs; OOD detection; foreseeable misuse; relevant adversarial conditions; fallback.

Element	Required content
Deployment	Monitoring; drift; incident handling; rollback; controlled adaptation; update approval.
Decision	Requirement-by-requirement result; unresolved limitations; independent review; authorized operating boundary.

Table N-3 — Synthetic confusion matrix

Source p. 76 | Illustrative or explanatory source table

Actual class	Predicted defect	Predicted non-defect	Total
Defect	190 (TP)	10 (FN)	200
Non-defect	15 (FP)	785 (TN)	800
Total	205	795	1,000

Appendix B — Local implementation and release

Readiness condition	Evidence before consequential use
Authority and responsibility	Actual legal operator, owner, competent technical authority, independent reviewer and permitted result class are recorded.
Scope and criteria	Object, configuration, site, period, trigger, every technical value or observable attribute, units and prospective approval are complete.
Evidence and methods	Original accessible evidence, representative observations, valid methods, uncertainty, adverse cases and limitations support the intended conclusion.
People and independence	Competence is current for each task; genuine conflicts, relevant prior work and needed separation have been checked.
Records and tools	Field meanings, unique IDs, original history, access, retention, recovery and actual deployed workflows have passed local validation.
Decision and public use	All relevant blockers are resolved; required independent decision and written permissions precede the exact public claim.
Continuity and handover	Receiving custody, unresolved duties, corrections and historical decisions survive transfer, retirement or provider closure.

Record the completed local readiness and actual authorization in AF-FR-9100-032 or an equivalent controlled record. Preserve the unchanged source standard and each earlier decision basis. A blocked activity may continue only within a genuinely permitted development boundary, not as a provisional conformity certificate.