



AF 9100:2026

Aerospace Quality Management System

PROCEDURE

Software, AI, Models and Trusted Updates

AF-PR-9100-027

FIRST EDITION

30 September 2026 | English master text

Publication identity and application

Control	Publication value
Title	Software, AI, Models and Trusted Updates
Document code	AF-PR-9100-027
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 technical and security authorities
Primary source anchors	12, 13, N, O
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 Trigger, inputs and handover	4
3 Operating sequence	4
4 Decision, interruption and escalation	5
5 Required records and receiving checks	5
6 Worked case and adverse challenge	5
Appendix A — Controlling source and evidence	7
Source tables — complete field and decision payloads	13
Appendix B — Local implementation and release	15
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

Applies to the responsible function and exact evaluated object selected from the source. Use only the duties triggered by actual activity and retained responsibility.

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 Trigger, inputs and handover

Trigger: Digital development, model-based acceptance, deployment, update, drift or end of support.

Accountable process owner: Digital technical and security authorities. Actual named appointments and deputies are maintained locally.

Input to verify	Required output / receiving function
Exact request, object and source profile	Accepted scope and controlled work package, or a reasoned hold/refusal.
Current people, authority, data and equipment	Authorized assignment and evidence-access boundary.
Original execution, adverse evidence and source history	Attributable result, identified limitations and required action.
Relevant source clocks and actual event times	Genuine due/completion history and timely recipient communication.

3 Operating sequence

3.1 Define reliance and authority

Identify requirements, criticality, architecture, domain, users, permitted and prohibited actions, interfaces and recovery. Assess whether the proposed tool is necessary and sufficiently verifiable.

3.2 Control source and data

Preserve source, logic, build definitions, toolchain, dependency inventory and loaded configuration. For AI preserve dataset rights, labeling, partitions, weights and preprocessing without evaluation leakage.

3.3 Verify the actual target and domain

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.

3.4 Check defensive integrity

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

3.5 Release and update predictably

Approve the exact artifact, target, installation, compatibility, interruption response and rollback limits. Verify the deployed state and protect against inappropriate reintroduction of a known unacceptable defect.

3.6 Monitor, constrain and preserve

Define drift, changed exposure, unavailable inputs and response authority. Reassess adaptation outside the approved envelope; preserve source/build access, records and support beyond supplier or platform exit.

4 Decision, interruption and escalation

Unknown target configuration, invalid evaluation or failed essential protection blocks consequential deployment.

State or event	Action and limit
Evidence valid and criteria met	Record the bounded result and pass it to the authorized receiving function; do not enlarge scope.
Evidence missing, invalid or inconclusive	Keep the dependent conclusion not yet demonstrated / HOLD, preserve the reason and define competent next work.
Valid evidence demonstrates failure	Record nonconformity and protect the affected population; follow actual technical disposition and source status controls.
Trigger genuinely absent	Record the exclusion and its approved responsibility basis. Resource shortages do not justify N/A.
Interrupted or changed work	Preserve prior state and actual configuration; complete impact and revalidation before claiming completion.

5 Required records and receiving checks

Document	Record purpose
AF-FR-9100-019	Software, AI, Model and Technical-Tool Release
AF-FR-9100-029	Cyber Incident, Update and Assured Recovery
AF-FR-9100-005	Technical Result and Decision Rule
AF-FR-9100-007	Independent Review and Attestation Decision

At handover the receiver checks object, scope, current version, original evidence, open conditions, true event chronology and the action they may authorize. An email sent, field marked complete or invoice paid is not acceptance of that handover.

6 Worked case and adverse challenge

Synthetic scenario: A model's overall accuracy is high but a separately tested difficult defect set has low recall.

Required reasoning: Retain both populations and limitations. Resolve the critical domain or enforce a technically validated restriction; do not pool it into a favorable overall score.

Challenge before local use	Expected behavior
Missing essential input	Dependent conclusion remains held; no favorable default.
Wrong configuration or performer	Detect the mismatch and reassess relevant evidence.

Challenge before local use	Expected behavior
Expired or conflicted authority	Block the prohibited assignment or decision; later renewal cannot validate earlier unauthorized work.
Failed mandatory condition with high total score	No compensation; preserve the failure and required action.
Correction without original history	Reject an apparent closure that erases relevant earlier evidence.

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.

13.1 Security scope and integrated risk assessment

SOURCE p. 31 | Normative when applicable

13.1(a) The organization shall identify digital assets, trust boundaries, dependencies, users and interfaces affecting assured outcomes. Threat scenarios shall be evaluated together with safety and mission consequences, including loss of integrity, availability, authenticity and confidentiality.

13.1(b) Security assumptions shall be reviewed across suppliers, remote services, maintenance access and physical access. The organization shall identify where a security measure could impair emergency operation or safety and resolve the conflict through a joint technical decision.

Required evidence: Asset and trust-boundary inventory; threat assessment; safety-security interface decisions.

13.2 Identity, access and privileged actions

SOURCE p. 31 | Normative when applicable

13.2(a) Access shall be based on verified identity, least necessary privilege and approved roles. Privileged access, remote maintenance and release-authority actions shall have strengthened authentication and traceable authorization appropriate to risk.

13.2(b) Access shall be removed or changed promptly when responsibilities end or change. Shared credentials shall not obscure accountability. Emergency access shall be controlled, logged, time-bounded and reviewed after use. Service accounts and machine identities shall be governed as well as human accounts.

Required evidence: Access matrix; privileged approvals; identity lifecycle; emergency-access reviews.

13.3 Secure development and supply-chain integrity

SOURCE p. 31 | Normative when applicable

13.3(a) Security requirements shall be integrated into development, procurement and maintenance. The organization shall assess dependency provenance, vulnerability exposure, secure configuration, supported versions and mechanisms for receiving security information.

13.3(b) A software bill of materials or equivalent controlled inventory shall identify relevant components and dependencies. Vulnerability disposition shall consider exploitability, exposure, safety consequences and available mitigations. Presence of a component in an inventory shall not itself demonstrate that it is secure.

Required evidence: Secure-development plan; SBOM; dependency reviews; vulnerability decisions and mitigations.

13.4 OT and connected-production safeguards

SOURCE p. 31 | Normative when applicable

13.4(a) Industrial networks and connected equipment shall be designed to limit unauthorized changes and propagation of failures. Appropriate segmentation, access control, backup, monitoring and safe degraded operation shall be established based on the process risk.

13.4(b) Patches, remote connections and security tools shall be evaluated for compatibility with timing, availability and safety. Where immediate patching is inappropriate, a time-bounded compensating-control decision shall be documented. Production data and machine-program integrity shall remain verifiable.

Required evidence: OT architecture; authorized connections; patch assessments; compensating controls; restoration tests.

13.5 Secure updates, rollback and cryptographic agility

SOURCE p. 31 | Normative when applicable

13.5(a) Updates shall verify origin, integrity, compatibility, authorization and intended target configuration before installation. The organization shall define interruption handling, recovery, anti-rollback needs and conditions for safe rollback without reintroducing an unacceptable vulnerability.

13.5(b) Cryptographic dependencies shall be inventoried and reviewed for supported algorithms, key management, certificate validity and migration capability. Future algorithm replacement shall be planned according to product life and threat exposure; this standard does not mandate a particular algorithm or claim universal post-quantum readiness.

Required evidence: Update architecture; signed release manifest; interruption and rollback tests; cryptographic inventory and migration plan.

13.6 Security testing and vulnerability handling

SOURCE p. 31 | Normative when applicable

13.6(a) Security testing shall have documented authorization, boundaries, safety precautions, data handling and restoration requirements. Methods shall be selected for the architecture and threat model. Test-induced operational or physical hazards shall be assessed before execution.

13.6(b) The organization shall provide a channel for vulnerability reports, assess credibility and impact, track remediation and communicate relevant actions to affected parties. Security findings shall remain connected to product configuration and safety assessment; closure shall require evidence, not only a changed status field.

Required evidence: Authorized test plans; results; vulnerability register; remediation verification; customer communication.

13.7 Incident response and recovery assurance

SOURCE p. 32 | Normative when applicable

13.7(a) The organization shall establish incident triage, containment, evidence preservation, technical escalation, communication and recovery responsibilities. Plans shall consider compromised engineering records, manufacturing programs, product updates and supplier credentials, not only office systems.

13.7(b) Recovery shall verify authenticity and consistency of restored data and configurations. A system shall not be returned to assured service solely because it is operational. Exercises shall demonstrate recovery time, recoverable data age and product-impact assessment under representative conditions.

Required evidence: Incident plan; exercise records; RTO/RPO basis; integrity verification; return-to-service decision.

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: $(N-1) \text{ Precision} = TP/(TP+FP)$; $\text{Recall} = TP/(TP+FN)$; $F1 = 2PR/(P+R)$ 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.

O.1 Security evidence baseline

SOURCE p. 77 | Normative when applicable

O.1(a) The organization shall maintain the records in Table O-1 for the relevant scope. Evidence shall be versioned and linked to the product, service or infrastructure configuration. Confidential security information shall have controlled access without losing authorized auditability.

O.1(b) The security review shall identify unresolved dependencies and limitations. A security certification or penetration-test report for a different scope shall not establish assurance of the actual deployed system.

O.2 Update release checklist

SOURCE p. 77 | Normative when applicable

O.2(a) Before update deployment, the organization shall verify the approved build, target compatibility, cryptographic origin/integrity, regression evidence, hazard/security impact and required authorization. The deployment plan shall include interrupted installation and loss-of-power scenarios where relevant.

O.2(b) After installation, the actual version, configuration and protective functions shall be verified. Failed or partial deployment shall remain visible in fleet or asset status. Rollback shall be evaluated for compatibility and vulnerability consequences rather than assumed always safe.

O.3 Incident-to-product linkage

SOURCE p. 77 | Normative when applicable

O.3(a) When an incident could affect engineering data, process programs, test records or released products, the organization shall assess the affected technical population. Cyber response and product-quality response shall share a controlled evidence trail.

O.3(b) Return to operation shall require both operational restoration and justified integrity. The authority making that decision shall understand remaining uncertainty. Unverified records shall not be silently reclassified as authentic because the server has been restored.

O.4 Filled illustrative recovery evaluation

SOURCE p. 77 | Definition, explanation or informative example; no new universal limit

O.4(a) Example recovery objectives are RTO=4 hours and RPO=1 hour for a technical-record service. An exercise restores service in 3.2 hours from a backup 45 minutes old. These two objectives are met in the example, but one configuration-link integrity check fails.

O.4(b) The exercise outcome is therefore “service restored; assured use held pending integrity correction”. Passing time objectives does not establish trustworthy evidence. The record retains the failed check, affected records, corrective action and successful repeat verification.

Source tables — complete field and decision payloads

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.
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

Table O-1 — Required security and integrity records

Source p. 77 | Required structure when applicable

Record	Minimum content
Asset/trust model	Assets, identities, interfaces, zones, data flows, dependencies and trust assumptions.
Threat and impact assessment	Credible threats, exposure, safety/mission consequences, controls and residual decisions.
Access record	Role, identity, privileges, approver, start/end, emergency use and periodic review.
SBOM and dependencies	Components, versions, provenance, known exposure, support status and disposition.
Vulnerability record	Report source, affected versions, severity basis, exploitability context, mitigation, deadline and verification.
Update manifest	Target identity, old/new configuration, origin/integrity evidence, compatibility, authorization and rollback conditions.
Incident record	Timeline, evidence preservation, affected population, containment, communication, root cause and recovery approval.
Recovery test	RTO/RPO basis, backup identity, restoration duration, integrity checks, configuration reconciliation and result.
Crypto-agility record	Algorithms, keys/certificates, support horizon, dependencies, migration needs and verification.
Supplier interface	Access conditions, notification duties, update responsibilities, evidence rights and termination controls.

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.