Evidence Room Export (Sanitized Sample) — EU AI Act Annex IV

KLA Digital Audit & Compliance Studio — Example High-Risk AI System Documentation Bundle

2025-12-12

0. Important notice (read first)

This document is a **fictional, sanitized sample** provided for demonstration and marketing purposes.

- All system names, vendors, metrics, identifiers, and log excerpts are **synthetic**.
- Any similarity to real systems or persons is coincidental.
- This is **not legal advice** and does not claim regulatory compliance.
- "Evidence Room" refers to an exported evidence package containing controls, decision trails, and monitoring artifacts for internal/external review.

This sample is structured to align with the **minimum technical documentation elements** listed in **EU AI Act Annex IV**.

1. Evidence Room bundle overview

1.1 Bundle identity

- Evidence Room Bundle ID: `ER-SAMPLE-2025-12-12-0001`
- Generated by: KLA Digital — Audit & Compliance Studio (sample export)
- Export profile: `EU_AI_ACT_ANNEX_IV_HIGH_RISK`
- Generation timestamp (UTC): `2025-12-12T10:30:00Z`
- Scope: Technical documentation + supporting runtime evidence artifacts
- Sanitization profile: `PUBLIC_WEBSITE_SAMPLE_v3`

1.2 What this bundle contains

This Evidence Room includes:

1. **Annex IV technical documentation** (Sections 2–11 of this document)
2. **Evidence manifest** (Appendix A)
3. **Policy pack excerpts (policy-as-code)** (Appendix B)
4. **Dynamic sampling & quality monitoring report** (Appendix C)
5. **Human oversight decision records** (Appendix D)
6. **Validation & testing excerpts** (Appendix E)
7. **Training data sheet excerpts** (Appendix F)
8. **Model card excerpts** (Appendix G)
9. **Audit ledger integrity verification procedure** (Appendix H)

1.3 Integrity and tamper-evidence model (high level)

The Evidence Room is designed to be verifiable:

- Each artifact has a `SHA-256` digest recorded in the evidence manifest.
- A bundle "root hash" is computed from the manifest.
- The bundle root hash is recorded in an append-only audit ledger.
- Verification steps are described in Appendix H.

1.4 Sanitization summary (what was removed or transformed)

This public sample applies the following sanitization:

- Removed: customer names, addresses, emails, phone numbers, government identifiers, account numbers, free-text notes from operators.
- Transformed: internal IDs → salted hashes, timestamps → rounded to nearest minute, monetary values → bucketed ranges.
- Redacted: any prompt/output content that could reveal proprietary policies or third-party confidential text.
- Synthetic substitution: example prompts, retrieved snippets, and test cases were regenerated from synthetic templates.

2. Annex IV conformance map (navigation)

This document is organized to cover Annex IV items (1)–(9):

- Annex IV (1) General description → Section 3
- Annex IV (2) System elements & development process → Section 4
- Annex IV (3) Monitoring, functioning, control → Section 5
- Annex IV (4) Appropriateness of performance metrics → Section 6
- Annex IV (5) Risk management system → Section 7
- Annex IV (6) Relevant lifecycle changes → Section 8
- Annex IV (7) Standards & technical specs → Section 9
- Annex IV (8) EU declaration of conformity → Section 10 (template)
- Annex IV (9) Post-market monitoring plan → Section 11

3. General description of the AI system (Annex IV item 1)

3.1 System name, provider, versioning

- AI system name: **CreditAssist**
- Provider (fictional): **Northwind Retail Bank plc**
- System version: `v1.4.2`
- Previous versions: `v1.3.x` (feature extraction updates), `v1.2.x` (monitoring redesign)

3.2 Intended purpose

CreditAssist is intended to support evaluation of **creditworthiness of natural persons** for consumer lending.

- Intended users: trained credit underwriters and compliance reviewers
- Intended setting: internal banking decision workflow
- Output: recommended decision band + explanation package + risk flags
- Decision authority: **human underwriter remains responsible** for the final decision

3.3 High-risk classification (context)

This use case is treated as **high-risk** in this sample because it concerns evaluation of creditworthiness and access to essential private services (credit).

3.4 System boundary and dependencies

CreditAssist comprises:

- Data ingestion: application form + uploaded documents (synthetic in this sample)
- Feature extraction: structured extraction pipeline (document parsing + rules)
- Scoring model: supervised learning model producing risk score
- Explanation generator: produces reason codes and evidence references
- Guardrails layer: policy checks, redaction, human review routing
- Observability & evidence capture: trace capture and immutable audit logging

External dependencies (examples):

- Identity provider (single sign-on)
- Credit bureau feed (batch)
- Core banking system (loan origination)
- Document storage system

3.5 Interaction with hardware/software (where applicable)

- Deployed as: web application + internal application programming interface (API)
- Integrations:
  - Loan origination system calls CreditAssist API to request a recommendation.
  - Underwriter console displays model outputs and requires a human sign-off step.
  - Audit/evidence layer captures prompts, inputs, retrieved snippets, outputs, and human actions.

3.6 Forms of placing on the market / putting into service

- Internal service (bank deployment) exposed via:
  - REST API
  - Web user interface (underwriter console)
  - Batch processing endpoint (pre-qualification campaigns)

3.7 Hardware environment (intended)

- Execution environment: private cloud virtual private cloud (VPC)
- Data residency (example): European Union (EU) region
- Key management: per-tenant keys (customer-managed key option in real deployments)

3.8 Basic user interface description (deployer-facing)

Underwriter console provides:

- Application overview (sanitized fields)
- Model risk score band and confidence indicator
- Top reason codes (structured)
- "Evidence links" to the captured inputs and rationale artifacts
- Required actions:
  - Approve / Reject / Request more info
  - Mandatory free-text justification if overriding model recommendation
  - Escalate to compliance for flagged cases

3.9 Instructions for use (summary)

- Use only for the defined lending products and eligible geographies.
- Do not use as the sole basis for automated refusal.
- Review reason codes and supporting evidence before finalizing decisions.
- If the system flags uncertainty, policy violations, or fairness concerns, route to the designated human review queue.
- Follow data quality requirements defined in Section 5.4.

4. Detailed description of system elements and development process (Annex IV item 2)

4.1 Development methodology (overview)

CreditAssist is developed using a controlled lifecycle:

1. Requirements definition (legal + compliance + business)
2. Data collection and curation
3. Model development (training + evaluation)
4. Validation (performance, robustness, discrimination testing)
5. Deployment with guardrails + monitoring enabled
6. Ongoing post-market monitoring with change control

4.2 Use of pre-trained systems or third-party tools

This sample assumes use of:

- Third-party open-source libraries for model training and feature processing
- A document parsing toolchain for extracting structured fields from uploaded documents
- Optional large language model components may be used for summarization/explanations (when enabled), with strict redaction and human oversight gating

All third-party components are tracked in a software bill of materials (SBOM) in the internal (non-public) version of this bundle.

4.3 Design specifications (logic, key choices, assumptions, trade-offs)

4.3.1 General logic

- Inputs: applicant-provided data + verified data sources (where available)
- Processing:
  - Validate input schema and completeness
  - Extract features (income stability, debt-to-income, delinquency indicators)
  - Compute risk score using supervised model
  - Produce structured explanation (reason codes + supporting signals)
  - Apply policy-as-code guardrails
  - Route to human review when required
- Outputs:
  - Recommendation: `APPROVE`, `CONDITIONAL_APPROVE`, `REFER`, `REJECT`
  - Risk band: `A-E`
  - Explanation package: reason codes + evidence references

4.3.2 Key design choices

- Human oversight by design:
  - CreditAssist does not finalize approvals; it produces a recommendation requiring explicit human confirmation.
- Explainability:
  - Outputs include reason codes tied to structured features, not free-form text alone.
- Minimization:
  - Only features required for intended purpose are used; sensitive attributes are excluded from direct use in scoring (see Section 7 for fairness controls).

4.3.3 Rationale and assumptions (examples)

- Assumption: verified income and debt signals improve risk estimation vs. self-declared values.
- Assumption: reason-code explanations improve underwriter review quality and contestability.
- Known trade-off: stricter guardrails increase review latency but reduce compliance risk.

4.4 System architecture (components and data flow)

4.4.1 Components

- `UI.UnderwriterConsole`
- `API.ScoringService`
- `ETL.FeatureExtractor`
- `MODEL.RiskScorer`
- `XAI.ReasonCodeGenerator`
- `POLICY.PolicyEngine`
- `OBS.TraceCollector` (OpenTelemetry-compatible)
- `AUDIT.AppendOnlyLedger`

4.4.2 Data flow (simplified)

1. Application submitted → schema validated
2. Documents ingested → fields extracted (with confidence per field)
3. Feature vector constructed → scored
4. Explanation package created → guardrails checked
5. Decision + evidence captured → underwriter review step enforced
6. Final human decision recorded → exported to origination system
7. Sampling monitors continuously evaluate quality and policy near-misses

4.4.3 Computational resources (example)

- Training: offline environment with dedicated compute, access-controlled

- Inference: autoscaled service with per-request logging and redaction
- Storage: encrypted at rest; access via least-privilege roles

4.5 Data requirements and dataset documentation (datasheets)

4.5.1 Training dataset (synthetic summary)

- Dataset name: `NW_Credit_Apps_EU_2019_2024` (synthetic label)
- Records: ~1.2M applications (synthetic count)
- Label: repayment outcome within a defined observation window
- Sampling: stratified by product and region (within intended scope)
- Exclusions:
  - Applications missing required verification fields
  - Known fraudulent submissions (handled in separate fraud systems)

4.5.2 Provenance and collection

- Sources: loan origination records + verified repayment history + bureau attributes
- Collection controls:
  - Data access approved by data governance board (internal)
  - Purpose limitation: model development for credit underwriting support only

4.5.3 Labelling procedures

- Outcome label derived from repayment events within defined time horizon.
- Label quality checks:
  - Duplicate removal
  - Temporal leakage checks
  - Consistency checks against source of truth

4.5.4 Data cleaning and preprocessing

- Missing values:
  - Controlled imputations for non-critical fields
  - Mandatory completion for critical fields
- Outlier detection:
  - Winsorization for extreme numeric values
  - Manual review thresholds for suspicious document extraction outputs
- Drift checks:
  - Feature distribution monitoring at inference time (Section 11)

4.6 Human oversight measures (Article 14 alignment, high-level)

Human oversight measures include:

- Mandatory underwriter confirmation before decision finalization
- Escalation paths for:
  - low-confidence extraction
  - policy violations
  - fairness alarms
  - ambiguous explanations
- Override handling:
  - overrides require justification
  - overrides are logged and sampled for audit review
- Interpretability support:
  - reason codes mapped to policy and feature definitions
  - evidence links to captured inputs and intermediate signals

4.7 Pre-determined changes (planned updates) and continuous compliance

Pre-determined changes in this sample:

- Quarterly recalibration:
  - threshold tuning to maintain stability under observed drift
- Semiannual model refresh:
  - retraining with updated repayment outcomes
- Policy pack updates:
  - new jurisdictional rules or internal policy changes
- Monitoring rule updates:
  - new near-miss detectors based on incident learnings

All changes require documented testing and sign-off (Section 8).

4.8 Validation and testing procedures (summary)

Validation covers:

- Predictive performance:
  - discrimination (e.g., AUC), calibration (e.g., Brier score), stability
- Robustness:
  - sensitivity to missing fields, document extraction noise, distribution shift
- Potential discriminatory impacts:
  - group-level comparisons using approved fairness metrics
- Policy compliance:
  - guardrail test suites (block/route/redact behavior)
- Logging and traceability:

  - evidence completeness checks (trace coverage targets)

Test logs and sign-offs appear in Appendix E (excerpt).

4.9 Cybersecurity measures (summary)

Controls in this sample include:

- Authentication and access control:
  - single sign-on, role-based access control, least privilege
- Encryption:
  - encryption in transit (TLS) and at rest (AES)
- Immutable audit logging:
  - append-only ledger with hash chaining and external timestamp anchoring
- Supply chain:
  - signed builds, dependency scanning, SBOM tracking
- Incident response:
  - security events generate alerts and are retained for forensic investigation

5. Monitoring, functioning and control (Annex IV item 3)

5.1 Capabilities

CreditAssist can:

- Produce a risk band and recommendation within intended scope
- Explain outputs with reason codes tied to structured signals
- Route uncertain or risky cases to human reviewers
- Provide full traceability for each decision episode

5.2 Limitations and known failure modes

Known limitations (examples):

- Document extraction errors:
  - poor scan quality or non-standard formats may reduce accuracy
- Distribution shift:
  - macroeconomic changes may degrade calibration without recalibration
- Proxy effects:
  - non-sensitive features may correlate with protected characteristics; requires monitoring
- Overreliance risk:
  - underwriters may overweight recommendations; mitigated via training and UI warnings

5.3 Foreseeable unintended outcomes and fundamental rights risks

Foreseeable risks:

- Discrimination:
  - different error rates across groups or regions
- Financial exclusion:
  - systematically higher rejection in specific subpopulations
- Contestability failures:
  - insufficient explanation or missing evidence links
- Privacy:
  - leakage of personal data into logs or prompts (mitigated by redaction)

5.4 Input data specifications (operational)

Minimum required inputs:

- Applicant identity and eligibility fields (sanitized)
- Verified income range (or verified “not available” marker)
- Current debt obligations (structured)

- Product type and requested amount (bucketed)
- Document extraction confidence per required field

Hard blocks:

- Missing critical fields
- Extraction confidence below threshold for required fields
- Conflicting signals (e.g., mutually inconsistent income indicators)

5.5 Human oversight in operation (runtime)

Operational oversight measures:

- Queue-based review routing:
  - `REFER` cases require second-level review
  - fairness alarm cases require compliance review
- Real-time intervention:
  - policy engine can block output delivery and require approval
- Audit sampling:
  - dynamic sampling selects cases for quality review (Appendix C)

6. Performance metrics and their appropriateness (Annex IV item 4)

6.1 Metrics used (non-exhaustive)

Predictive:

- AUC (discrimination)
- Brier score (calibration)
- Approval rate stability (process stability)

Operational:

- Underwriter override rate (human trust / disagreement signal)
- Time-to-decision (workflow impact)
- Evidence completeness rate (trace coverage)

Fairness and harm monitoring:

- Differences in false negative rates across monitored groups (where legally and operationally appropriate)
- Disparity in referral rates
- Complaint and appeal outcomes (post-market)

6.2 Why these metrics are appropriate (justification)

- Credit underwriting requires both ranking quality and calibrated risk estimates, hence discrimination + calibration metrics.
- Operational metrics detect "silent failures" where the model looks good offline but breaks workflow or oversight in practice.
- Fairness monitoring focuses on error rates and process disparities to catch discriminatory impacts early.

7. Risk management system (Annex IV item 5)

7.1 Risk management process (summary)

Risk management steps in this sample:

1. Identify hazards and reasonably foreseeable misuse
2. Estimate and evaluate risk severity and likelihood
3. Define controls (technical + organizational)
4. Verify controls through testing and monitoring
5. Assess residual risk and decide acceptance criteria
6. Continuously monitor, investigate incidents, update controls

7.2 Risk register (public excerpt)

R-01: Discriminatory outcomes via proxy variables

- Control: fairness monitoring, feature review, policy gating, periodic revalidation
- Residual risk: medium (monitored)

R-02: Logging of personal data in prompts/outputs

- Control: redaction policy-as-code + automated scanners + sampling audits
- Residual risk: low-to-medium (monitored)

R-03: Overreliance by underwriters (automation bias)

- Control: UI warnings, training, mandatory justifications for overrides, sampling audits
- Residual risk: medium

R-04: Model drift leading to poor calibration

- Control: drift detection, recalibration plan, stop-ship thresholds
- Residual risk: medium

R-05: Security compromise of decision evidence

- Control: encryption, access control, append-only ledger, integrity verification
- Residual risk: low

7.3 Residual risk acceptance (example)

Residual risk is accepted only if:

- Monitoring thresholds are configured and active
- Stop-ship triggers are defined and tested
- Human oversight routing is enforced for high-impact cases
- Evidence completeness meets minimum target

8. Lifecycle changes (Annex IV item 6)

8.1 Change control process

All changes follow:

- Change request ticket with scope and rationale
- Impact assessment (performance, fairness, security, compliance)
- Test plan and results attached
- Approval workflow:
  - Model risk owner
  - Compliance owner
  - Engineering owner
- Versioned deployment with rollback plan
- Post-deploy monitoring intensified for defined window

8.2 Changelog excerpt (synthetic)

- v1.4.2 (2025-11-20)
  - Updated drift detector thresholds for income-feature distribution
  - Added guardrail: mandatory `REFER` for low-confidence document extraction
- v1.4.1 (2025-10-02)
  - Calibration adjustment for new macroeconomic segment
- v1.4.0 (2025-09-10)
  - Introduced explanation package v2 (reason code taxonomy updates)

9. Standards and technical specifications (Annex IV item 7)

9.1 Harmonised standards status (note)

At the time this public sample was produced, the provider's internal compliance program references a mix of international standards and internal controls, with a plan to map and adopt relevant harmonised standards once available/published for presumption of conformity pathways.

9.2 Referenced standards and frameworks (examples)

Information security and privacy:

- ISO/IEC 27001 (information security management)
- ISO/IEC 27701 (privacy extension)
- OWASP application security guidance (secure development)

AI risk and governance:

- ISO/IEC 23894 (AI risk management)
- NIST AI Risk Management Framework (AI RMF)

Software quality and assurance:

- ISO/IEC 25010 (quality model)
- Supply chain controls (SBOM + signed builds)

10. EU declaration of conformity (Annex IV item 8) — template (sample)

EU Declaration of Conformity (Template — Sample Only)

1. AI System: CreditAssist `v1.4.2`
2. Provider: Northwind Retail Bank plc (fictional)
3. Address: [REDACTED — sample]
4. Statement: This declaration is issued under the sole responsibility of the provider.
5. Object of the declaration: CreditAssist — creditworthiness evaluation decision support system
6. Conformity assessment procedure: [SPECIFY — sample placeholder]
7. References to relevant standards/specifications: See Section 9
8. Signed for and on behalf of the provider:
   - Name/Role: [REDACTED — sample]
   - Place/Date: [REDACTED — sample]
   - Signature: [REDACTED — sample]

11. Post-market monitoring plan (Annex IV item 9)

11.1 Monitoring objectives

- Detect performance degradation and drift
- Detect policy violations and near-misses
- Detect emerging discriminatory impacts
- Ensure evidence completeness and traceability
- Provide timely incident response and corrective actions

11.2 Data collected in operation (minimized)

- Decision episode metadata (timestamps, model version, policy version)
- Input completeness and extraction confidence
- Output recommendation + reason codes (structured)
- Human actions: approvals, overrides, escalations
- Monitoring metrics (aggregated)
- Incident and complaint signals (aggregated where possible)

11.3 Dynamic sampling and review workflow

- Sampling strategy:
  - risk-weighted sampling (higher sampling rate for higher impact cases)
  - random baseline sample for coverage
- Review roles:
  - Model risk reviewers
  - Compliance reviewers
- Outcomes:
  - confirm OK
  - remediation required
  - incident opened

11.4 Incident response and corrective actions

Triggers include:

- fairness threshold breach
- drift threshold breach
- evidence completeness below minimum
- security alert affecting integrity or confidentiality
- abnormal override rates

Corrective actions include:

- stop-ship / rollback
- policy gate tightening

- retraining / recalibration
- UI warnings update
- reviewer training update

Appendix A — Evidence manifest (excerpt)

Bundle root hash (SHA-256): `b7c1...(sample)...9a2e`

Artifacts (excerpt; full manifest omitted in public sample):

- A1. `ANNEXIV_TechDoc_CreditAssist_v1.4.2.pdf`
  SHA-256: `2f91...(sample)...1c0b`
- A2. `PolicyPack_PUBLIC_SAMPLE_v3.yaml`
  SHA-256: `11aa...(sample)...90ff`
- A3. `DynamicSamplingReport_2025Q4_PUBLIC.pdf`
  SHA-256: `9d02...(sample)...aa41`
- A4. `HumanOversightRecords_EXCERPT.jsonl`
  SHA-256: `c0de...(sample)...beef`
- A5. `ValidationTestReport_EXCERPT.pdf`
  SHA-256: `7e57...(sample)...1337`

Integrity verification procedure: see Appendix H.

Appendix B — Policy-as-code (excerpt; sanitized)

```yaml
policy_pack_id: "POLICY_PACK_PUBLIC_SAMPLE_v3"
policy_engine: "KLA.PolicyEngine"
policy_pack_version: "3.2.0"

global:
  pii_redaction:
    enabled: true
    actions:
      - redact_patterns: ["EMAIL", "PHONE", "GOV_ID", "ACCOUNT_NUMBER"]
      - replace_with: "[REDACTED]"
  logging:
    store_raw_documents: false
    store_structured_features: true
    store_free_text: "minimized"

rules:
  - id: "POL-001"
    name: "No fully automated rejection"
    when:
      output.recommendation: "REJECT"
    then:
      require_human_approval: true
      route_queue: "UNDERWRITER_LEVEL2"

  - id: "POL-002"
    name: "Low-confidence extraction must refer"
    when:
      any_input.extraction_confidence_lt: 0.80
    then:
      override_output:
        recommendation: "REFER"
      require_human_approval: true

  - id: "POL-003"
    name: "Fairness alarm escalates to compliance"
    when:
      monitoring.fairness_alarm: true
    then:
      route_queue: "COMPLIANCE_REVIEW"
      require_human_approval: true

  - id: "POL-004"
    name: "Explanation required"
    when:
```

```
  output.recommendation_in: ["REJECT", "REFER", "CONDITIONAL_APPROVE"]
then:
  require_fields: ["reason_codes", "evidence_links"]
```

Appendix C — Dynamic sampling & quality monitoring report (excerpt; synthetic)

Reporting window: `2025-10-01` to `2025-12-01`

Sampling configuration (synthetic):

- Base sampling rate: 2%
- Risk-weighted boost:
  - Risk band D/E → +8%
  - Any policy near-miss → +15%
  - Any underwriter override → +10%

Findings (synthetic summary):

- Evidence completeness rate: 98.7%
- Cases routed to mandatory human review: 100% (by policy)
- Policy near-misses detected:
  - PII redaction near-miss (prevented): 6
  - Missing explanation package (blocked): 14
- Quality issues flagged:
  - Document extraction low-confidence clusters: 2 (formats)
- Corrective actions opened:
  - Update extraction model for “scanned payslip” template
  - Add UI warning for low-confidence uploads

Appendix D — Human oversight records (excerpt; synthetic)

Record ID: `HITL-2025-11-18-00421`

- Model recommendation: `APPROVE` (Risk band B)
- Policy routing: `UNDERWRITER_REVIEW_REQUIRED`
- Underwriter decision: `CONDITIONAL_APPROVE`
- Underwriter rationale (sanitized): "Requested additional verification due to inconsistent employer details."
- Evidence links: `EVID-...(synthetic)...`
- Timestamp: `2025-11-18T09:42Z`

Record ID: `HITL-2025-11-22-00987`

- Model recommendation: `REJECT` (Risk band E)
- Policy routing: `UNDERWRITER_LEVEL2`
- Level-2 decision: `REJECT`
- Rationale (sanitized): "Multiple delinquency indicators confirmed from verified source."
- Timestamp: `2025-11-22T14:10Z`

Appendix E — Validation & testing report (excerpt; synthetic)

Test suite: `CreditAssist_PreDeploy_v1.4.2`

Coverage:

- Predictive performance tests: PASS
- Calibration tests: PASS (within thresholds)
- Robustness tests: PASS (missing fields, extraction noise)
- Fairness monitoring dry-run: PASS (no threshold breaches in test set)
- Policy pack tests: PASS (block/route/redact expected behavior)
- Evidence capture tests: PASS (trace completeness within target)

Sign-off (synthetic):

- Responsible engineer: [REDACTED — sample]
- Model risk owner: [REDACTED — sample]
- Compliance owner: [REDACTED — sample]
- Date: 2025-11-19

Appendix F — Training data sheet (excerpt; synthetic)

- Intended scope: consumer lending products in defined EU regions
- Known exclusions: fraud system cases, incomplete verification cases
- Label definition: repayment outcome within observation window
- Known biases: historical underwriting policy shifts across time
- Mitigations:
  - temporal validation splits
  - monitoring for policy-change drift
  - reviewer guidance and oversight escalation

Appendix G — Model card (excerpt; synthetic)

Model: `RiskScorer_v1.4.2`

- Type: supervised learning risk scoring model
- Inputs: structured feature vector (no raw documents)
- Outputs: risk score and band
- Intended use: decision support (human-in-the-loop)
- Out-of-scope uses:
  - fully automated rejection without human confirmation
  - products/geographies not listed in the internal intended use statement
- Monitoring:
  - drift detection
  - calibration checks
  - fairness alarms
  - override rate anomaly detection

Appendix H — Evidence integrity verification (how an auditor checks this bundle)

1. Compute SHA-256 for each artifact listed in Appendix A.
2. Verify computed hashes match the manifest values.
3. Compute the bundle root hash from the manifest canonical form.
4. Verify the root hash exists in the append-only audit ledger for the export timestamp.
5. Confirm the ledger chain is consistent (hash-chained entries; no breaks).
6. Validate signatures (where applicable) for:
    - policy pack version
    - deployment attestation
    - test report sign-off artifacts