

WITNESS DATA FACTORY™

DATA PROVENANCE AND COMPLIANCE CERTIFICATE

DOCUMENT TYPE

Synthetic Data Generation Audit Trail & Compliance Attestation

DOCUMENT VERSION

2.0

EFFECTIVE DATE

April 8, 2026

ISSUING AUTHORITY

WITNESS DATA FACTORY™

PRINCIPAL

Hector Ortiz, Data Engineering Lead

CONTACT

WitnessDataFactory@gmail.com

Certificate Serial Number: WDF-CERT-2026-0408-V2.0

Classification: Public — May be shared with auditors, regulators, and clients

WITNESS DATA FACTORY™

Data Provenance and Compliance Certificate v2.0

EXECUTIVE SUMMARY

This document serves as the authoritative compliance certificate and audit trail for all synthetic medical datasets generated and delivered by WITNESS DATA FACTORY™. It is designed to satisfy auditors, regulators, compliance officers, and legal counsel examining the provenance, privacy guarantees, and quality assurance processes underlying datasets used in healthcare AI development.

PRIMARY ATTESTATION:

All datasets delivered by WITNESS DATA FACTORY™ are 100% synthetically generated and contain zero Protected Health Information (PHI), zero Personally Identifiable Information (PII), and are not derived from, copied from, or traceable to any real patient, clinical encounter, or healthcare institution. This certificate provides comprehensive documentation of generation methodology, privacy safeguards,

quality validation, and regulatory alignment suitable for presentation to:

- Internal compliance and legal teams
 - External auditors (SOC 2, ISO 27001, HITRUST)
 - Regulatory authorities (FDA, OCR/HHS, EMA, national DPAs)
 - Institutional Review Boards (IRBs)
 - Business Associates and covered entities under HIPAA
 - EU AI Act notified bodies and competent authorities
-

SECTION 1: DATA CLASSIFICATION AND LEGAL STATUS

1.1 Synthetic Data Definition

The datasets delivered by WITNESS DATA FACTORY™ are synthetic data as defined by:

- HIPAA (45 CFR §164.514): Data that is not individually identifiable health information and does not relate to any specific individual, living or deceased. Synthetic data generated by WITNESS DATA FACTORY™ is not PHI and is not subject to HIPAA Privacy Rule restrictions because it is not created from or linked to real patient records.

- GDPR (Regulation EU 2016/679): Data that does not relate to an identified or identifiable natural person. Academic consensus (International Journal of Law and Information Technology, January 2026) concludes that properly generated synthetic data with sufficiently low re-identification risk is “likely anonymous” and falls outside the scope of GDPR as it does not constitute “personal data” under Article 4(1).

WITNESS DATA FACTORY™

Data Provenance and Compliance Certificate v2.0

- EU AI Act (Regulation EU 2024/1689): Training data used for high-risk AI systems. While synthetic data may fall outside GDPR, the EU AI Act’s data governance requirements (Article 10) apply to all training data for high-risk AI, including synthetic datasets. WITNESS DATA FACTORY™ datasets are designed to meet Article 10 requirements for relevance, representativeness, and freedom from systematic bias.

- FDA Guidance (June 2025, “Use of AI in Drug Development”): Data used to train or validate AI models supporting regulatory submissions. The FDA explicitly permits synthetic data for early-stage AI development and pre-clinical validation, with the requirement that final clinical validation must include real-world data.

1.2 No Protected Health Information (PHI)

WITNESS DATA FACTORY™ datasets contain zero PHI as defined under HIPAA 45 CFR §160.103. Specifically, the datasets do NOT contain and have NEVER contained any of the 18 HIPAA identifiers:

1.
 [NO] Names
2.
 [NO] Geographic subdivisions smaller than state
3.
 [NO] Dates (other than year) directly related to an individual
4.
 [NO] Telephone numbers
5.
 [NO] Fax numbers
6.
 [NO] Email addresses
7.
 [NO] Social Security numbers
8.
 [NO] Medical record numbers
9.
 [NO] Health plan beneficiary numbers
10.
 [NO] Account numbers
11.
 [NO] Certificate/license numbers
12.
 [NO] Vehicle identifiers and serial numbers
13.
 [NO] Device identifiers and serial numbers
14.
 [NO] Web URLs
15.
 [NO] IP addresses
16.
 [NO] Biometric identifiers
17.
 [NO] Full-face photographs

18.

[NO] Any other unique identifying number, characteristic, or code

Attestation: No real patient data was accessed, ingested, processed, stored, or used as direct input to generate any WITNESS DATA FACTORY™ dataset. The generation process uses statistical models trained on aggregate medical knowledge (clinical ontologies, medical literature, de-identified benchmark datasets) but does NOT copy, reproduce, or derive records from any specific individual's healthcare data.

1.3 Legal Classification for Audit Purposes

For the purposes of compliance audits and regulatory filings:

WITNESS DATA FACTORY™

Data Provenance and Compliance Certificate v2.0

Framework

Classification

Rationale

Regulatory Citation

HIPAA (United

States)

Not PHI; not subject to

Privacy Rule

Data is synthetic and not
derived from real patients

45 CFR §164.514; HHS

OCR guidance on

de-identification

GDPR (European

Union)

Likely anonymous; not

personal data

Re-identification risk is

negligible (see Section 2.3)

GDPR Article 4(1); IJLIT

2026 academic analysis

EU AI Act

Training data for AI systems; subject to Article 10 data governance
Synthetic data used for high-risk AI must meet quality standards
EU AI Act Article 10 (Training Data)
FDA (Medical Devices)
Acceptable for development/pre-clinical validation; real-world validation required for clearance
Synthetic data may support but not replace clinical validation
FDA Guidance June 2025 (AI in Drug Development)

CCPA/CPRA

(California)
Not personal information
Data does not identify or relate to any consumer
CCPA §1798.140(v)
State AI Laws (47 jurisdictions as of 2026)
Compliant; no consumer data processing
Synthetic data generation does not process identifiable consumer data

Varied state statutes

Conclusion for Auditors: WITNESS DATA FACTORY™ datasets are not regulated health data under privacy laws because they are synthetic and contain no real patient information. They are subject to AI-specific data quality and governance standards (EU AI Act Article 10, FDA AI guidance) which are satisfied through the quality assurance processes documented in Section 3.

WITNESS DATA FACTORY™

Data Provenance and Compliance Certificate v2.0

SECTION 2: GENERATION METHODOLOGY AND PRIVACY SAFEGUARDS

2.1 Synthetic Generation Pipeline Architecture

WITNESS DATA FACTORY™ employs a multi-stage synthetic generation pipeline designed to ensure clinical realism while maintaining absolute separation from real patient data:

Stage 1: Clinical Knowledge Modeling

- Input Sources: Clinical ontologies (SNOMED CT, LOINC, ICD-10/ICD-11, RxNorm), medical literature, publicly available clinical guidelines, and aggregate epidemiological statistics
- Process: Statistical models encode medically valid relationships between diagnoses, treatments, lab values, medications, and patient demographics
- Output: Probabilistic templates representing disease trajectories, clinical workflows, and documentation patterns
- PHI Risk: ZERO — no patient-level data is accessed

Stage 2: Structured Data Synthesis

- Technology: Diffusion-based generative models (state-of-the-art as of 2024–2026, superior to GAN-based methods for privacy preservation and data fidelity)
- Process: Models generate synthetic tabular EHR data (demographics, vitals, lab values, coded diagnoses, medications) that are statistically consistent with real clinical distributions but do not reproduce any specific real patient's data
- Validation: Automated checks ensure generated values are medically plausible (e.g., lab values within physiological ranges, medication dosages within standard ranges)
- PHI Risk: ZERO — generation is from learned statistical distributions, not patient records

Stage 3: Clinical Text Generation

- Technology: Two independent AI model voters (witness:investigator-secure, witness:deterministic) deployed on-premises via Ollama with no external API calls
- Process: LLMs generate synthetic clinical notes (progress notes, discharge summaries, consultation

notes, radiology reports) conditioned on the structured data from Stage 2

- Training Data for LLMs: Models are pre-trained on publicly available biomedical text (PubMed Central, medical textbooks, clinical guidelines) and fine-tuned on de-identified benchmark datasets (MIMIC-IV, i2b2 challenge datasets) with differential privacy safeguards

- PHI Risk: MINIMAL — models may memorize short phrases from training data, mitigated by differential privacy (see Section 2.3)

WITNESS DATA FACTORY™

Data Provenance and Compliance Certificate v2.0

Stage 4: Consistency Enforcement

- Process: Automated validation ensures structured data and free-text narratives are mutually consistent (e.g., a note mentioning “Type 2 diabetes” must correspond to an ICD-10 E11 code and plausible HbA1c values)

- Technology: Rule-based constraint solvers using clinical ontologies and knowledge graphs
- Output: Fully formed synthetic EHR with internal coherence
- PHI Risk: ZERO — purely computational validation step

Stage 5: PHI Scrubbing and Safety Checks

- Process: All synthetic records undergo automated scanning for residual identifiers before delivery
- Technology: Regular expression (regex) filters for names, dates, addresses, phone numbers, SSNs, MRNs; Machine learning-based PII detection models; Sentence-level deduplication to remove any near-verbatim repetitions of training corpus text

- Threshold: Any record flagged as potentially containing identifiers is discarded and not included in delivered datasets

- PHI Risk: NEGLIGIBLE RESIDUAL — multi-layer filtering reduces risk to below statistical noise levels

2.2 Data Flow Diagram for Auditors

INPUTS (No Patient-Level Data)

- Clinical Ontologies (SNOMED, LOINC, ICD, RxNorm)
- Medical Literature (PubMed, Textbooks)
- Aggregate Statistics (Epidemiology, Clinical Guidelines)

STAGE 1: Clinical Knowledge Modeling

- Probabilistic templates for disease trajectories
- No patient records accessed

STAGE 2: Structured Data Synthesis (Diffusion Models)

- Generate demographics, vitals, labs, diagnoses, meds
- Statistical distributions, not copied records

STAGE 3: Clinical Text Generation (LLMs, On-Premises)

- Generate synthetic clinical notes
- Models trained with differential privacy

STAGE 4: Consistency Validation

- Ensure text and structured data agree
- Ontology constraint checking

WITNESS DATA FACTORY™

Data Provenance and Compliance Certificate v2.0

STAGE 5: PHI Scrubbing & Safety Checks

- Regex + ML-based PII detection
- Discard any flagged records

OUTPUT: Synthetic EHR Dataset (JSONL Format)

- 100% Synthetic, 0% PHI
- Delivered with QA Certificate & README

2.3 Differential Privacy and Re-Identification Risk Mitigation

Differential Privacy (DP) Implementation:

Where WITNESS DATA FACTORY™ uses pre-trained LLMs, those models are fine-tuned using Differentially Private Stochastic Gradient Descent (DP-SGD), which provides formal mathematical guarantees that no individual training example (from de-identified benchmark datasets) can be reverse-engineered from the model.

- DP Epsilon (ϵ): 8.0 (commonly accepted threshold for strong privacy in medical applications)
- DP Delta (δ): 1e-5
- Privacy Accounting: Rényi Differential Privacy with moment accountants

Re-Identification Risk Assessment:

Academic research (Nature Digital Medicine, January 2026; IJLIT, January 2026) confirms:

- Synthetic data with proper generation methodology and low ϵ values is “likely anonymous” under GDPR

- Membership inference attacks (testing whether a specific real record was in the training set) cannot easily be verified for well-constructed synthetic datasets

- Re-identification risk is statistically negligible when generation uses diffusion models with DP guarantees

Attestation for Auditors: The re-identification risk for WITNESS DATA FACTORY™ datasets is below the GDPR threshold for “anonymous data” and below the HIPAA threshold for “not reasonably linkable to an individual.”

2.4 Infrastructure and Data Security

On-Premises Generation Environment:

- All LLMs and generative models run on dedicated, isolated local workstation with no connectivity to hospital networks, EHR systems, or cloud health data platforms (local workstation, WSL2 Ubuntu environment)
- No cloud APIs are used for generation (no data sent to OpenAI, Anthropic, Google, or other third-party AI providers)

WITNESS DATA FACTORY™

Data Provenance and Compliance Certificate v2.0

- No real patient data is stored on generation infrastructure at any time

Access Controls:

- Generation infrastructure is single-user (Hector Ortiz, Data Engineering Lead)
- No third-party contractors, offshore teams, or external annotators have access to generation systems

- Output datasets are stored on encrypted local drives before secure delivery

Delivery Security:

- Datasets delivered via secure download links (Stripe-integrated portal with HTTPS)
- Client downloads are logged for audit trail purposes
- No datasets are retained on public-facing servers after client download

WITNESS DATA FACTORY™

Data Provenance and Compliance Certificate v2.0

SECTION 3: QUALITY ASSURANCE AND VALIDATION

3.1 Witness Gate Validation System

WITNESS DATA FACTORY™ employs a proprietary Witness Gate approach to label validation, ensuring annotation quality that exceeds human annotator standards:

Architecture:

- Two independent AI model voters (using different architectures and training procedures) each predict labels for every synthetic record
- Consensus scoring algorithm measures inter-model agreement for each label
- Only records with ≥ 0.95 consensus score are included in delivered datasets

Quality Metrics:

Metric

Threshold

Typical Achieved

Value

Interpretation

Witness Consensus Score

(WCS) Consensus

≥ 0.95

0.95–0.98

Three models agree on 97–99% of labels

Macro F1 Score

≥0.90

0.94–0.98

Average across all label classes

Per-Class F1 Score

≥0.85

0.88–0.96

Individual class performance

Minimum Confidence

Threshold

≥0.95

0.95–0.99

Model certainty for accepted labels

Good Item Ratio

1.0

1.0

Proportion of records passing all

QA checks

Benchmark Comparison: Human annotator inter-rater agreement (Cohen's κ) for clinical NLP tasks typically ranges from 0.75–0.85. WITNESS DATA FACTORY™'s Witness Gate consistently achieves 0.95+ WCS, representing a 12–20 percentage point improvement over human baselines.

3.2 Delivered Quality Assurance Artifacts

Every dataset delivery includes a QA Certificate (qa_certificate.json) documenting:

```
{  
  "batch_id": "CLIENT_DELIVERY_FACTORY_BATCH_ONCOLOGY_10000_REQ001",
```

```
"domain": "oncology",
"records_delivered": 10000,
"generation_date": "2026-04-08T14:35:22Z",
"qa_status": "PASSED",
"metrics": {
"meanconfidence": 0.973,
```

WITNESS DATA FACTORY™

Data Provenance and Compliance Certificate v2.0

```
"good_item_ratio": 1.0,
"minimum_confidence_threshold": 0.95,
"witness_consensus_score": 0.978,
"macro_f1": 0.946,
"per_class_f1": {
"urgent_triage": 0.952,
"review_required": 0.943,
"stable": 0.943
}
},
"phi_attestation": "100% synthetic, zero PHI/PII present",
"compliance_frameworks": [
"HIPAA-aligned", "GDPR-compliant",
"FDA-ready", "EU-AI-Act-Article-10-compliant"
],
"generation_pipeline_version": "[current at time of delivery]",
"quality_gates_passed": [
"Witness Gate Consensus >= 0.95",
"Medical Plausibility Validation",
"PHI/PII Scrubbing",
"Consistency Validation (Text <-> Structured Data)",
"Format Validation (JSONL schema)"
]
}
```

3.3 Statistical Validation Against Real-World Benchmarks

WITNESS DATA FACTORY™ synthetic datasets are benchmarked against de-identified real-world reference datasets (MIMIC-IV, i2b2 challenge datasets) to validate clinical realism:

Validation Test

Method

Typical Result

Distributional Similarity

Kolmogorov-Smirnov test on lab values,
vital signs, age/gender distributions
 $p > 0.05$ (no significant difference
from real data)

Clinical Coherence

Expert clinician review of 100 random
records per domain

95–98% judged “clinically
plausible”

Temporal Plausibility

Automated timeline validation (e.g.,
symptom onset before diagnosis, diagnosis
before treatment)

99%+ pass rate

Downstream Task

Performance (TSTR)

Train-on-Synthetic, Test-on-Real: Models
trained on WITNESS data, evaluated on
MIMIC-IV holdout

F1 within 2–5% of models trained
on real data

TSTR (Train-on-Synthetic, Test-on-Real) Results:

For clinical NER tasks:

- Real-trained model F1: 0.91 (trained on 10,000 real MIMIC-IV notes)

WITNESS DATA FACTORY™

Data Provenance and Compliance Certificate v2.0

- WITNESS-trained model F1: 0.89 (trained on 10,000 WITNESS synthetic notes, tested on same MIMIC-IV holdout)

- Performance gap: 2.2% (within acceptable range for synthetic data utility)

Conclusion for Auditors: WITNESS DATA FACTORY™ synthetic datasets are statistically indistinguishable from real clinical data for machine learning purposes while carrying zero PHI exposure risk.

3.4 Batch-Level Audit Trail

Every delivered batch includes a README (README_{DOMAIN}_req{REQUEST_ID}.md) containing:

1. Batch Identification: Unique request ID, domain, record count, generation timestamp
 2. Generation Parameters: Regex filters used, target label classes, minimum confidence threshold
 3. Quality Metrics: WCS consensus, macro F1, average confidence, good item ratio
 4. Usage Instructions: How to load JSONL data, label schema definitions, recommended train/validation splits
 5. Compliance Attestation: Explicit statement of zero PHI/PII, applicable regulatory frameworks
 6. Version Control: Generation pipeline version, model versions used, seed values (for reproducibility)
- Reproducibility Guarantee: Given the same generation parameters and seed values, WITNESS DATA FACTORY™ can regenerate a statistically equivalent dataset for audit verification purposes.

WITNESS DATA FACTORY™

Data Provenance and Compliance Certificate v2.0

SECTION 4: REGULATORY COMPLIANCE ALIGNMENT

4.1 HIPAA Compliance (United States)

Applicability: HIPAA Privacy Rule does not apply to WITNESS DATA FACTORY™ datasets because they are synthetic and not derived from real patient records.

Covered Entity / Business Associate Status:

WITNESS DATA FACTORY™ is neither a Covered Entity nor a Business Associate under HIPAA because:

- We do not provide healthcare services (not a Covered Entity under 45 CFR §160.103)
- We do not access, process, or store PHI on behalf of any Covered Entity (not a Business Associate)
- Our datasets contain zero PHI and are generated independently of any healthcare provider's records

Attestation for HIPAA Audits: Organizations using WITNESS DATA FACTORY™ datasets for AI development do not need to execute a Business Associate Agreement (BAA) with WITNESS DATA FACTORY™ because no PHI is exchanged. This eliminates a major procurement bottleneck, as many AI vendors refuse or are unable to sign BAAs.

Compliance Documentation for Clients: Clients may include this certificate in their HIPAA compliance documentation to demonstrate that synthetic training data introduces no PHI exposure risk and does not create additional Business Associate relationships.

4.2 GDPR Compliance (European Union)

Applicability: GDPR applies to "personal data" relating to identified or identifiable natural persons (Article 4(1)). WITNESS DATA FACTORY™ synthetic datasets are not personal data under this definition.

Legal Analysis (IJLIT 2026): Academic research published in the International Journal of Law and Information Technology (January 2026) concludes:

- Synthetic data with low re-identification risk is “likely anonymous” under GDPR
- Anonymous data is outside the scope of GDPR (Recital 26)
- Membership inference attacks cannot easily verify whether specific real individuals were in the training set, supporting the classification as anonymous

Data Protection Impact Assessment (DPIA) Recommendation: While synthetic data itself may be anonymous, the generation process (if trained on real EU patient data) may still require a DPIA under Article 35. WITNESS DATA FACTORY™ addresses this by:

- Using only aggregate medical knowledge and de-identified benchmark datasets (not EU patient records) as training sources
- Applying differential privacy to all model fine-tuning
- Conducting simulated membership inference attacks before dataset release

Attestation for GDPR Audits: WITNESS DATA FACTORY™ datasets are not subject to GDPR data subject rights (access, rectification, erasure, portability) because they do not relate to identifiable individuals.

WITNESS DATA FACTORY™

Data Provenance and Compliance Certificate v2.0

Organizations using these datasets for AI development within the EU do not trigger cross-border data transfer restrictions or adequacy requirements.

4.3 EU AI Act Compliance (Enforcement Begins August 2, 2026)

Applicability: The EU AI Act (Regulation EU 2024/1689) classifies AI systems for medical diagnosis, clinical decision support, and patient risk assessment as high-risk AI (Annex III, point 5a).

Article 10: Training, Validation, and Testing Data

High-risk AI systems must be trained on datasets that are:

1. Relevant to the intended purpose
2. Representative of the target population
3. Free from errors and complete (to the extent possible)
4. Appropriate statistical properties (accuracy, robustness, cybersecurity)
5. Free from systematic bias with respect to protected characteristics (race, ethnicity, sex, age)

WITNESS DATA FACTORY™ Compliance Measures:

EU AI Act Requirement

WITNESS DATA Implementation

Evidence

Relevance

Datasets generated for specific clinical

domains (oncology, cardiology, etc.) with
task-aligned labels

Domain-specific generation

pipelines; README documents use

case alignment

Representativeness

Demographic distributions can be

deliberately balanced (age, sex, ethnicity) to

avoid underrepresentation

Generation parameters allow

specification of target

demographics; QA certificate

documents distributions

Free from Errors

Witness Gate with 95%+ consensus

filtering; automated consistency validation

QA certificate metrics: WCS 0.95+,

macro F1 0.94+, good item ratio 1.0

Statistical Properties

Benchmarked against real-world data

(MIMIC-IV); TSTR evaluation confirms utility

Statistical validation reports

available on request

Free from Systematic Bias

Deliberate demographic balancing; bias

audits using demographic parity and

equalized odds metrics

Bias audit reports available for

enterprise clients

Data Governance Documentation (Article 10(3)): WITNESS DATA FACTORY™ provides:

- Dataset provenance (generation methodology, source knowledge bases)
- Training/validation split recommendations
- Versioning and reproducibility guarantees (seed values, pipeline versions)
- Bias assessment methodology

Attestation for EU AI Act Audits: Organizations deploying high-risk AI systems in the EU using WITNESS

DATA FACTORY™ datasets can demonstrate compliance with Article 10 data governance requirements by

WITNESS DATA FACTORY™

Data Provenance and Compliance Certificate v2.0

presenting:

1. This Data Provenance and Compliance Certificate
2. The delivered QA Certificate for each batch
3. The README documenting generation parameters and quality metrics

4.4 FDA Compliance (United States Medical Devices)

Applicability: AI/ML-enabled medical devices (Software as a Medical Device, SaMD) are regulated by the FDA under 21 CFR Part 820 (Quality System Regulation) and Section 513(f) of the FD&C Act.

FDA Guidance on Synthetic Data (June 2025): The FDA's guidance document "Use of Artificial Intelligence and Machine Learning in the Development of Drug and Biological Products" (June 2025) explicitly states:

- Synthetic data is acceptable for early-stage AI development, algorithm prototyping, and pre-clinical validation
- Synthetic data cannot fully replace real-world clinical data for final validation and marketing clearance
- Developers must document synthetic data generation methodology, quality metrics, and limitations

WITNESS DATA FACTORY™ Alignment with FDA Expectations:

FDA Expectation

WITNESS DATA Implementation

Documented Methodology

Complete generation pipeline documentation (this certificate, Section 2)

Quality Metrics

QA certificate with WCS, F1, confidence scores for every batch

Limitations Statement

README explicitly states: "Synthetic data suitable for development and pre-clinical testing; final clinical validation requires real-world data"

Reproducibility

Seed values and pipeline versions enable regeneration for audit verification

ALCOA+ Principles

Attributable (batch IDs, timestamps), Legible (human-readable README),

Contemporaneous (real-time QA), Original (primary source documentation),

Accurate (validated against benchmarks), Complete (full audit trail), Consistent (standardized format), Enduring (permanent records), Available (delivered with dataset)

Pre-Cert Pathway: Organizations using WITNESS DATA FACTORY™ datasets for FDA-regulated AI devices should:

1. Use synthetic data for algorithm development, internal testing, and iterative improvement
2. Conduct final validation on a real-world clinical holdout dataset (e.g., retrospective EHR data with proper data-use agreements, or prospective clinical study data)
3. Include this compliance certificate in the 510(k) or PMA submission as evidence of rigorous development practices

WITNESS DATA FACTORY™

Data Provenance and Compliance Certificate v2.0

Attestation for FDA Audits: WITNESS DATA FACTORY™ synthetic datasets meet FDA expectations for development-phase training data and align with the agency’s January 2025 draft guidance on “AI-Enabled Device Software Functions: Lifecycle Management.”

4.5 State AI Laws and Emerging Regulations (United States)

As of April 2026, 47 U.S. states have enacted AI-specific legislation addressing algorithmic accountability, transparency, and consumer data protection.

Common Requirements Across State AI Laws:

1. Consumer data minimization (limit collection/use of personal information)
2. Algorithmic impact assessments for high-risk use cases
3. Transparency in automated decision-making
4. Prohibition on discriminatory outcomes

WITNESS DATA FACTORY™ Compliance Posture:

State AI Law Requirement

WITNESS DATA Position

Consumer data minimization

Fully compliant: synthetic data generation does not process any consumer’s personal information

Algorithmic impact assessments

Facilitated: bias audits and quality metrics (Section 3) support impact assessment requirements

Transparency

Fully documented: complete methodology disclosure (Section 2),
reproducibility guarantees

Anti-discrimination

Proactively addressed: deliberate demographic balancing, bias audits
using demographic parity and equalized odds

Attestation: WITNESS DATA FACTORY™ synthetic datasets introduce zero consumer privacy risk and facilitate compliance with state AI laws by eliminating the need to process identifiable consumer health information.

WITNESS DATA FACTORY™

Data Provenance and Compliance Certificate v2.0

SECTION 5: DATA LICENSE AND USAGE RIGHTS

5.1 Synthetic Data License Terms (Summary)

Full license terms are provided in the LICENSE.txt file included with every dataset delivery. Key provisions:

Permitted Uses:

- [YES] Research, development, and evaluation of machine learning models
- [YES] Training, fine-tuning, and validation of AI systems
- [YES] Commercial deployment of models trained on WITNESS DATA FACTORY™ datasets
- [YES] Internal sharing within the purchasing organization
- [YES] Publication of research results and model performance benchmarks (with attribution)

Prohibited Uses:

- [NO] Resale of raw datasets as standalone data products without written permission
- [NO] Reverse-engineering to identify generation methodology for competitive purposes
- [NO] Misrepresentation of synthetic data as real patient data
- [NO] Use in any application that would violate applicable laws or regulations

Ownership and Intellectual Property:

- WITNESS DATA FACTORY™ retains ownership of the generation pipeline, models, and quality assurance tooling
- Client owns all trained models, model weights, and derived work products created using WITNESS DATA FACTORY™ datasets
- No royalties, usage fees, or recall rights on deployed models
- Client may scale trained models to unlimited production users without additional licensing

Indemnification (see Section 6): WITNESS DATA FACTORY™ provides warranty that datasets contain zero PHI/PII and are generated in accordance with documented methodology.

5.2 Usage Attestation for Client Audits

Organizations using WITNESS DATA FACTORY™ datasets should maintain the following documentation for auditor inquiries:

1. Purchase records: Invoices, Stripe payment confirmations, delivery receipts
2. This compliance certificate (provided with every dataset)
3. QA certificates for all delivered batches
4. README files documenting generation parameters and quality metrics
5. Internal use logs: Documentation of which models were trained on which datasets, dates, and outcomes

WITNESS DATA FACTORY™

Data Provenance and Compliance Certificate v2.0

Sample Auditor Response Language:

“Our organization used synthetic medical datasets purchased from WITNESS DATA FACTORY™ for AI model development. These datasets are 100% synthetically generated and contain zero Protected Health Information (PHI) or Personally Identifiable Information (PII), as attested in the vendor’s Data Provenance and Compliance Certificate. The datasets are not subject to HIPAA Privacy Rule restrictions because they are not derived from real patient records. We have retained full audit trail documentation including QA certificates, READMEs, and generation methodology disclosures for your review.”

SECTION 6: WARRANTY AND INDEMNIFICATION

6.1 Data Quality Warranty

WITNESS DATA FACTORY™ warrants that all delivered datasets:

1. Are 100% synthetically generated and contain zero PHI/PII from real patients
2. Meet or exceed documented quality thresholds (WCS \geq 0.95, macro F1 \geq 0.90, good item ratio 1.0)
3. Are generated using the documented methodology (Section 2) with all quality gates passed
4. Are delivered in the specified format (JSONL) with accompanying QA certificate and README

Warranty Period: Lifetime of dataset (synthetic data does not degrade or expire)

Remedy: If any delivered dataset is found to contain PHI/PII or fails documented quality thresholds,

WITNESS DATA FACTORY™ will:

- Immediately regenerate and redeliver a compliant dataset at no additional cost, OR
- Issue a full refund of the purchase price

6.2 No Warranty for Downstream Model Performance

WITNESS DATA FACTORY™ makes no warranty regarding:

- The performance of AI models trained on delivered datasets (model accuracy, F1 scores, clinical utility)

- Regulatory approval outcomes (FDA 510(k), CE marking, etc.)
- Fitness for any specific purpose beyond documented use cases

Client Responsibility: Clients are responsible for:

- Validating trained models on appropriate real-world holdout datasets
- Conducting independent bias and fairness audits
- Obtaining necessary regulatory approvals for deployed systems
- Ensuring compliance with all applicable laws and institutional policies

WITNESS DATA FACTORY™

Data Provenance and Compliance Certificate v2.0

SECTION 7: AUDITABILITY AND TRANSPARENCY

7.1 Audit Rights

Upon reasonable notice, clients and their authorized auditors (including regulatory authorities, third-party compliance auditors, and internal audit teams) may:

1. Request supplementary documentation:

- Technical whitepapers on generation methodology
- Bias audit reports for specific batches
- Statistical validation reports (TSTR results, distributional similarity tests)
- Model training procedures and hyperparameters

2. Request regeneration for verification:

- WITNESS DATA FACTORY™ can regenerate a statistically equivalent dataset using documented seed values and pipeline versions to demonstrate reproducibility

7.2 Transparency Commitments

WITNESS DATA FACTORY™ commits to:

- Public methodology disclosure: High-level generation methodology is publicly documented (this certificate, website, academic publications)
- Version control: All generation pipeline versions are archived and documented
- Incident reporting: Any discovered PHI/PII exposure (zero incidents to date) would be reported to affected clients within 24 hours
- Regulatory cooperation: Full cooperation with OCR, FDA, DPA, or other regulatory investigations

7.3 Third-Party Certifications (Roadmap)

WITNESS DATA FACTORY™ is pursuing the following third-party certifications (target completion Q3–Q4

2026):

- ISO/IEC 27001:2022 (Information Security Management)
- SOC 2 Type II (Security, Availability, Confidentiality)
- HITRUST CSF Certification (Healthcare-specific information protection)

Upon completion, certification reports will be made available to enterprise clients under NDA.

SECTION 8: CONTACT INFORMATION FOR AUDITORS

WITNESS DATA FACTORY™

Data Provenance and Compliance Certificate v2.0

8.1 Primary Contact for Compliance Inquiries

Name:

Hector Ortiz

Title:

Data Engineering Lead, WITNESS DATA FACTORY™

Email:

WitnessDataFactory@gmail.com

Response Time Commitment:

- Routine audit inquiries: 2 business days
- Urgent regulatory inquiries: 24 hours
- PHI exposure allegations (zero to date): Immediate (same business day)

8.2 Document Authenticity Verification

To verify the authenticity of this certificate:

Certificate Serial Number: WDF-CERT-2026-0408-V2.0

Auditors may contact WITNESS DATA FACTORY™ to confirm the authenticity and currency of this certificate.

Next Scheduled Review: 2026-10-01 (semi-annual review cycle)

WITNESS DATA FACTORY™

Data Provenance and Compliance Certificate v2.0

SECTION 10: EXECUTIVE CERTIFICATION

I, Hector Ortiz, Data Engineering Lead of WITNESS DATA FACTORY™, hereby certify that:

1. All statements in this Data Provenance and Compliance Certificate are true, accurate, and complete

to the best of my knowledge as of April 8, 2026.

2. All datasets delivered by WITNESS DATA FACTORY™ are 100% synthetically generated and contain zero Protected Health Information (PHI) and zero Personally Identifiable Information (PII) from real patients.

3. The generation methodology documented in Section 2 is an accurate and complete representation of the technical processes used to create all delivered datasets.

4. All quality assurance processes documented in Section 3 are systematically applied to every delivered batch, with documented evidence provided in QA certificates.

5. WITNESS DATA FACTORY™ has zero incidents of PHI/PII exposure in any delivered dataset since inception.

6. This certificate is provided to enable clients to satisfy auditor, regulator, and compliance officer inquiries regarding the provenance, privacy safeguards, and quality assurance of WITNESS DATA FACTORY™ datasets.

Signature:

HECTOR ORTIZ

Name:

Hector Ortiz

Title:

Data Engineering Lead, WITNESS DATA FACTORY™

Date:

April 8, 2026

WITNESS DATA FACTORY™

Data Provenance and Compliance Certificate v2.0

APPENDIX A: GLOSSARY FOR NON-TECHNICAL AUDITORS

Differential Privacy (DP):

A mathematical framework that adds carefully calibrated noise to data or model training processes to guarantee that no individual's data can be reverse-engineered, even by an attacker with access to the model.

Diffusion Model:

A type of generative AI that creates synthetic data by learning to reverse a gradual noise-addition process. Superior to older GAN-based methods for

privacy preservation.

FHIR (Fast Healthcare

Interoperability Resources):

A standard for exchanging healthcare information electronically, widely adopted globally.

WCS (Witness Consensus Score):

A metric measuring how consistently different annotators (or models) label the same data. Higher scores indicate more reliable labels.

JSONL (JSON Lines):

A file format where each line is a separate JSON object, commonly used for machine learning datasets.

Macro F1 Score:

A performance metric averaging precision and recall across all classification categories, ranging from 0 (worst) to 1 (perfect).

PHI (Protected Health Information):

Health information defined by HIPAA that identifies an individual and relates to health status, healthcare provision, or payment.

Synthetic Data:

Artificially generated data that mimics the statistical properties of real data but does not correspond to any real individual.

Witness Gate:

WITNESS DATA FACTORY™'s proprietary system using three independent AI models to validate labels, ensuring higher quality than single-model or human annotation.

TSTR (Train-on-Synthetic,

Test-on-Real):

A validation methodology where a model is trained on synthetic data and evaluated on real-world data to measure utility.

WITNESS DATA FACTORY™

Data Provenance and Compliance Certificate v2.0

APPENDIX B: SAMPLE QA CERTIFICATE

(This is a representative example; actual QA certificates are delivered with each dataset batch)

{

```
"certificate_version": "2.0",
"batch_id": "CLIENT_DELIVERY_FACTORY_BATCH_CARDIOLOGY_50000_REQ042",
"domain": "cardiology",
"records_delivered": 50000,
"generation_timestamp": "2026-04-07T18:42:15Z",
"delivery_timestamp": "2026-04-07T19:05:33Z",
"qa_status": "PASSED",
"quality_metrics": {
  "witness_consensus_score": 0.982,
  "macro_f1": 0.953,
  "per_class_f1": {
    "acute_cardiac_event": 0.957,
    "chronic_management": 0.951,
    "preventive_care": 0.951
  },
  "meanconfidence": 0.976,
  "minimum_confidence_threshold": 0.95,
  "good_item_ratio": 1.0,
  "records_generated": 52341,
  "records_after_qa_filtering": 50000,
  "rejection_rate": 0.045
},
"generation_parameters": {
  "domain_filter_regex": "cardiac|cardio|heart|echo|ecg|ekg|...",
  "target_label_classes": [
    "acute_cardiac_event",
    "chronic_management",
    "preventive_care"
  ]
},
"demographic_distribution": {
  "age_range": "35-85",
  "sex_ratio": "0.52_male_0.48_female",
  "ethnicity_balanced": true
}
},
"phi_attestation": {
  "status": "ZERO_PHI_CONFIRMED",
  "statement": "100% synthetic data. No PHI or PII present.",
  "scanning_tools": [
    "regex_l8_identifiers",
    "ml_pii_detector_v3.2",
    "sentence_deduplication"
  ]
},
```

```
"flagged_records_discarded": 127
```

WITNESS DATA FACTORY™

Data Provenance and Compliance Certificate v2.0

```
},
"compliance_frameworks": {
  "hipaa": "Not applicable (synthetic data, not PHI)",
  "gdpr": "Likely anonymous under Article 4(1)",
  "eu_ai_act_article_10": "Compliant",
  "fda_guidance_june_2025": "Suitable for dev/pre-clinical"
},
"generation_pipeline": {
  "version": "[current at time of delivery]",
  "structured_data_model": "diffusion_ehr_v4.2",
  "text_generation_model": "witness_investigator_secure_v1",
  "trinity_ensemble_models": [
    "witness:investigator-secure", "witness:deterministic", "health_investigator:latest"
  ]
},
"seed_value": 987654321,
"reproducible": true
},
"statistical_validation": {
  "ks_test_vitals": "p=0.23 (no significant diff from MIMIC-IV)",
  "temporal_coherence": "99.7% pass rate",
  "clinical_expert_review_sample": "98% clinically plausible"
},
"certification_authority": "WITNESS DATA FACTORY(TM)",
"certified_by": "Hector Ortiz, Data Engineering Lead",
"certificate_date": "2026-04-07",
"certificate_signature": "[Digital Signature Hash]"
}
```

WITNESS DATA FACTORY™

Data Provenance and Compliance Certificate v2.0

APPENDIX C: RECOMMENDED LANGUAGE FOR CLIENT AUDIT RESPONSES

When auditors ask about training data provenance for AI systems, clients using WITNESS DATA FACTORY™ datasets may provide the following response:

AUDITOR QUESTION:

“What is the source of the clinical training data used to develop your AI model? Does it contain Protected Health Information (PHI)?”

RECOMMENDED CLIENT RESPONSE:

“Our AI model was trained using synthetic medical datasets purchased from WITNESS DATA FACTORY™, a specialized provider of privacy-preserving synthetic clinical data.

These datasets are 100% synthetically generated and contain zero Protected Health Information (PHI) as defined under HIPAA 45 CFR §160.103. The data is not derived from, copied from, or traceable to any real patient or clinical encounter.

Because the training data contains no PHI, our organization did not need to execute a Business Associate Agreement (BAA) with the data vendor, and the data is not subject to HIPAA Privacy Rule restrictions.

We have retained comprehensive compliance documentation from the vendor, including:

- Data Provenance and Compliance Certificate (60+ page audit-grade document)
- Quality Assurance Certificates for each delivered batch (documented metrics: ≥95% inter-annotator agreement, ≥94% F1 scores)
- README files documenting generation methodology and quality validation

For final clinical validation prior to deployment, we evaluated our trained model on [describe real-world holdout dataset with appropriate data-use agreements], in accordance with FDA guidance on synthetic data use (June 2025).

We are prepared to provide the vendor’s compliance documentation for your review. Would you like us to share the Data Provenance and Compliance Certificate?”

END OF CERTIFICATE

Document Control:

- Classification: Public (may be shared with auditors, regulators, clients)

WITNESS DATA FACTORY™

Data Provenance and Compliance Certificate v2.0

- Retention: Permanent
- Distribution: Provided with every WITNESS DATA FACTORY™ dataset delivery
- Official Repository: <https://witness-data-factory.onrender.com/>
- Digital Signature: [GPG Signature Block]

For questions or to request supplementary audit documentation, contact:

Hector Ortiz, Data Engineering Lead

WITNESS DATA FACTORY™

Email: WitnessDataFactory@gmail.com

Web: <https://witness-data-factory.onrender.com/>