GOVERNANCE / EVIDENCE PACK

PAYMENT_ANTIFRAUD Synthetic Release

Audit-ready evidence bundle for the dataset processed under project Payment Fraud Synthetic. This document summarizes lineage, utility, privacy testing, policy checks, and release readiness for internal.

DATASET

PAYMENT_ANTIFRAUD v1.2.0

1,250,000 records from 847,392-row source.

EVIDENCE PACK ID

EP-PRJ-001

Sealed — by Alex Carter

TECHNIQUE

DP-CTGAN

Differential privacy budget ε=1.24, δ=1e-05.

RELEASE DECISION

PASS

internal

PCI-DSS Compliance Framework

Evidence level EL_2

EXECUTIVE SUMMARY

Dataset profile, provenance, and release posture

FINAL VERDICT

PASS

The dataset was assembled from a governed source, transformed through a differential privacy synthesis workflow, and evaluated against customer-defined privacy and governance thresholds before publication.

DATASET FACTS		GOVERNANCE CONTEXT	
Source dataset	**PAYMENT_ANTIFRAUD_RAW**	Workspace	**Compliance Hub**
Output dataset	**PAYMENT_ANTIFRAUD**	Owner	**Alex Carter**
Source rows	**847,392**	Applied policies	**3**
Published rows	**1,250,000**	Evidence retention	**2555 days**
Fields	**20**	Lineage snapshot	**2026-03-01T14:45:00Z**
Format	**Parquet**	Packaging hash	**sha256:a3f1c8b2e4d6f8a0c...**

CORE METRICS

POST-TRANSFORM RISK

14.0%

k-anonymity k=7. No singling-out or linkage risks detected.

PRIVACY BUDGET USED

24.8%

1.24 of total epsilon 5.0 consumed.

LINEAGE OVERVIEW

VALIDATION RESULTS

Fidelity, privacy testing, and policy checks

The final release candidate passed all configured privacy thresholds while preserving downstream utility. Scores below are the values captured in the evidence pack payload.

METRIC	OBSERVED	INTERPRETATION
D1 Validity	**94.2**	All validity checks passed. Null rates within tolerance.
D2 Fidelity	**87.5**	KS/JS divergence within bounds. Correlation structure preserved.
D4 Privacy Risk	**14.0**	k-anonymity k=7. No singling-out or linkage risks detected.
D6 Composite Score	**85.0**	Weighted composite (D1×0.2 + D2×0.3 + D4×0.5) = 85.0

METRIC	OBSERVED	THRESHOLD	STATUS
Privacy Risk	**14.0%**	**30.0%**	PASS

Risk Reduction

Privacy risk at 14.0%

k-anonymity k=7. No singling-out or linkage risks detected.

Policy Alignment

3 policies applied

PAMOLA Standard Gate Policy v2.1, DP Budget Allocation Policy, PCI-DSS Compliance Framework

Usage Constraints

Release approved

Release context: internal. Threshold profile: Default Synthetic.

PRIVACY EVIDENCE

Frozen metrics behind each gate check

The privacy metrics this run sealed, each feeding one gate check — the same evidence the gate decision consumed · run #7.

PRIVACY CHECK	METRIC	MEASURED	DIRECTION	THRESHOLD	WITHIN BOUND	CONFIDENCE
Singling-out resistance	`singling_frac`	0.003	lower is safer	—	—	94%
Linkage resistance	`cvpl_score`	0.008	lower is safer	—	—	92%
Membership-inference resistance · weakest	`mia_advantage`	0.02	lower is safer	—	—	90%
Privacy budget	`dp_budget_state`	0.11 / 5.0	spent / ceiling	5.0	**yes**	98%
Attack-suite coverage	`attack_suite_coverage`	0.97	higher is safer	0.7	**yes**	97%
Special-category protection	`race_leak_gap`	0.005	lower is safer	—	—	95%

Evidence digest sha256:c69d5b52030ccb03... — matches the live gate decision for this run.

ATTACK SIMULATION RESULTS

Per-execution findings from the attack battery

Normalized results from the G4 attack battery run against this subject. Each row is one execution with its verdict, severity, observed metric, and pass threshold. Legacy "WARNING" values are normalized to "WARN".

WARN **4 executions**

ATTACK TEMPLATE	METRIC	VALUE	THRESHOLD	VERDICT	SEVERITY	ADVERSARY	CONFIDENCE · N
atk-linkage-naive	linkage_success_rate	0.03	0.1	PASS	LOW	ADV-EXT-001	92% · n=1400
atk-singling-out	singling_out_rate	0.007	0.05	PASS	LOW	ADV-EXT-001	95% · n=1400
atk-mia-shadow-model	mia_advantage	0.058	0.05	WARN	MEDIUM	ADV-INT-002	81% · n=800
atk-attribute-inference	attribute_inference_accuracy	0.62	0.75	PASS	LOW	ADV-EXT-001	88% · n=1400

MERKLE ROOT: sha256:c1d2e3f4a5b6c7d8e9f0a1b2c3d4e5f6c1d2e3f4a5b6c7d8e9f0a1b2c3d4e5f6

APPENDIX

Artifact manifest and integrity details

The evidence bundle ties together the synthetic dataset, lineage snapshot, attack reports, policy decisions, and packaging metadata needed for audit and handoff conversations.

ASSEMBLED FROM	
Dataset version	**v1.2.0**
Lineage snapshot	**2026-03-01T14:45:00Z**
Attack reports	**4 executions · WARN**
Policy decisions	**3 rules applied**
Budget log	**Included**
Package size	**139 KB**

INTEGRITY & RETENTION	
sha256:a3f1c8b2e4d6f8a0c...	`sha256:a3f1c8b2e4d6f8a0c2e4f6a8b0d2e4f6a8c0b2d4e...`
Evidence retention	**2555 days**
Primary use	**Compliance documentation**
Secondary use	**External audit, regulator submission**
Assessor	**Alex Carter**
Workspace	**Compliance Hub**

ARTIFACT MANIFEST

Sealed artifacts (3)

Full list of artifacts sealed inside the Evidence Pack. Each entry maps to a SHA-256 hash recorded inside the cryptographic seal — tampering with any artifact invalidates the Ed25519 signature.

ARTIFACT	DESCRIPTION	SHA-256	SIZE	STATUS
evidence_pack_prj001_run1.json	evidence_pack (1a2b3c4d5e6f1a2b...)			INCLUDED
gate_pack_run_001.json	gate_pack (e5f6a1b2c3d4e5f6...)			INCLUDED
attack_sim_mia_report.json	attack_report (f6a1b2c3d4e5f6a1...)			INCLUDED

8 Recommendations

Safe Use Cases

This evidence pack supports: internal analytics, ML model development, privacy-safe data sharing within workspace scope, and audit submission.

Residual Risk

Residual risk is within acceptable thresholds. Ongoing monitoring recommended: re-evaluate if source data distribution changes, monitor attack-success metrics, review privacy budget.

Refresh Schedule

Evidence Retention	2555 days
Next Review	On next pipeline run or before retention expiry