SHROUVA

Privacy-Preserving AI for the Enterprise

WHITE PAPER

Version 1.0 | June 2026

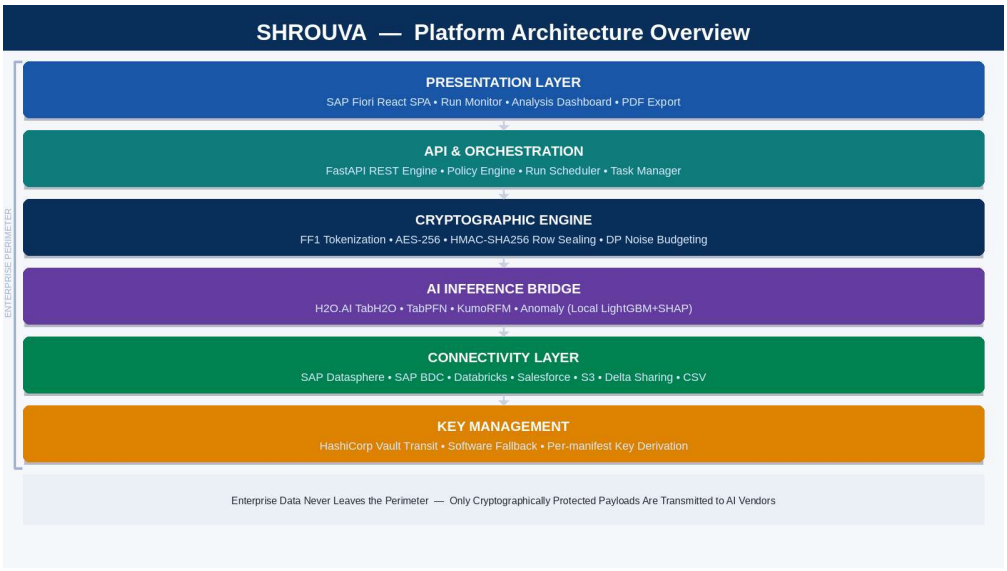

Figure 1 — Shrouva Platform Architecture: five functional layers, all within the enterprise perimeter.

Executive Summary

Enterprises today are sitting on vast repositories of high-value data — customer records, transaction histories, operational metrics — yet remain unable to unlock the full potential of foundation AI models because those models are operated by external vendors who would, by necessity, receive a copy of that data. The regulatory environment (GDPR, CCPA, HIPAA, and emerging AI-specific frameworks) makes this an unacceptable compliance risk. The result is a growing "AI value gap": organizations know the insight is there, but cannot safely reach it.

Shrouva bridges that gap. Built as part of the SAP BDC ecosystem, Shrouva is a privacy-preserving AI handoff platform that cryptographically protects sensitive data before it ever leaves the enterprise perimeter, submits the protected payload to best-in-class tabular AI vendors, and returns intelligible, auditable results to the data team — all without the vendor ever seeing a single piece of personally identifiable information (PII).

Key capabilities in this release: Format-preserving tokenization (FF1), AES-256 encryption, statistical generalization, differential privacy noise budgeting, HMAC-SHA256 per-row integrity seals, native connectors to SAP Datasphere, SAP BDC, Databricks, Salesforce, Amazon S3, Delta Sharing, and local file systems, AI inference orchestration for H2O.AI TabH2O (regression, classification, clustering, imputation, anomaly detection), and a Claude-powered AI narrative engine for human-readable interpretation of model results.

Part I — The Rise of Tabular Foundation Models

1.1 A New Paradigm in Applied Machine Learning

For decades, machine learning on structured (tabular) enterprise data required bespoke feature engineering, domain expertise, and significant data science investment. Models were trained from scratch on proprietary datasets, often requiring millions of rows and months of iteration. The arrival of tabular foundation models represents a fundamental inflection point.

Drawing inspiration from the transformer revolution in natural language processing — exemplified by GPT-4, Claude, and Gemini — researchers at leading institutions and AI companies have demonstrated that large-scale pre-training on diverse tabular datasets enables zero-shot and few-shot learning on novel enterprise datasets. Where a traditional gradient-boosted model might need 50,000 labeled examples to converge, a tabular foundation model can deliver competitive accuracy with a few hundred.

> **Research Highlight:** *Hollmann et al. (2022) introduced TabPFN, demonstrating that a prior-data fitted network pre-trained on synthetic datasets could outperform tuned gradient boosting on small classification tasks in under one second of inference time. This paper catalyzed a generation of tabular foundation model research. [Ref. 1]*

The commercial implications are profound. Tabular data constitutes the dominant data type in enterprise environments — ERP systems, CRM databases, data warehouses, financial ledgers — yet has historically been the last domain to benefit from the "pre-train once, deploy everywhere" paradigm that transformed NLP and computer vision. That asymmetry is now resolving.

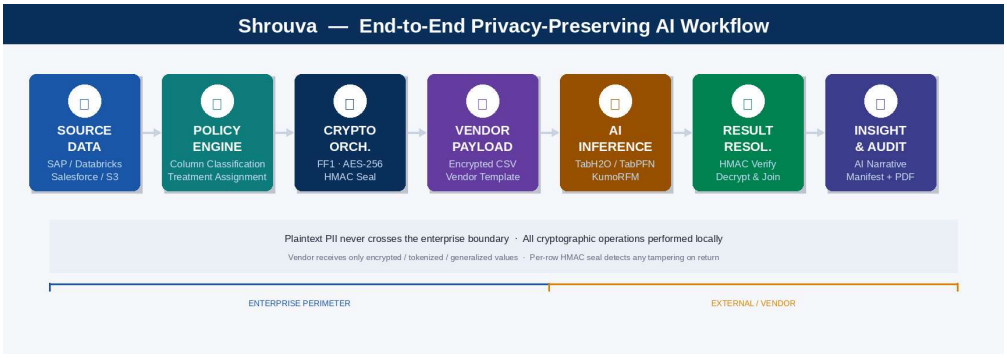

Figure 2 — End-to-end Shrouva workflow: data moves from enterprise sources, through the cryptographic engine, to external AI vendors, and back — with no plaintext leaving the perimeter.

1.2 Leading Tabular AI Models and Vendors

The tabular AI landscape is consolidating around a small number of high-performance foundation models, each with distinct architectural choices, target use cases, and commercial positioning.

H2O.AI — TabH2O

H2O.AI, headquartered in Mountain View, California, is one of the most established names in enterprise AI, with a customer base spanning Fortune 500 financial institutions, healthcare systems, and retailers. Their flagship tabular foundation model, TabH2O, builds on H2O.AI's long history of AutoML and gradient boosting excellence (H2O-3, Driverless AI) and extends it with a foundation-model architecture capable of zero-shot tabular inference.

TabH2O supports regression, multi-class classification, clustering, imputation, and anomaly detection via a unified REST API, making it operationally attractive for enterprise MLOps teams seeking a single integration point for multiple ML task types. Shrouva integrates natively with the TabH2O prediction endpoint.

> **Market Context:** *H2O.AI reported over 20,000 organizations using its platform and raised $100M+ in funding. The company's focus on explainability and regulated industries makes it a natural fit for privacy-conscious deployments. [Ref. 2]*

Prior Labs — TabPFN (Cloud API)

Prior Labs, spun out of the research group at the University of Freiburg that authored the original TabPFN paper, offers TabPFN v2 as a cloud API. The model employs a novel in-context learning architecture that treats the training set itself as the context window for each inference call — effectively performing Bayesian inference over a learned prior distribution of tabular datasets.

TabPFN v2 achieves state-of-the-art results on the AutoML Benchmark for datasets up to approximately 10,000 rows and 500 features, with inference times measured in seconds rather than minutes. Its zero-shot capability is particularly compelling for organizations with limited labeled data, such as those launching new products or entering new markets.

> **Research Highlight:** *Hollmann et al. (2025), 'TabPFN v2: Improved In-Context Learning for Tabular Data,' demonstrated that TabPFN v2 surpasses AutoML systems including AutoGluon, LightAutoML, and CatBoost on a broad meta-benchmark spanning 300+ datasets. [Ref. 3]*

Kumo.AI — KumoRFM

Kumo.AI, founded by former members of Pinterest's machine learning infrastructure team, approaches tabular AI from a relational data perspective. KumoRFM (Relational Foundation Model) is designed to perform predictive queries directly over multi-table relational graphs — the native structure of enterprise data in ERP and CRM systems — without requiring the user to perform manual join and feature engineering operations.

KumoRFM uses Predictive Query Language (PQL), a SQL-like syntax that allows data teams to express ML tasks in familiar terms ("PREDICT customer.churn FOR customer_id WHERE last_purchase > 90 days ago") and receive trained model outputs without writing model code. This dramatically lowers the barrier to production ML for data engineering teams.

Market Context: *Kumo.AI has been recognized as one of the most innovative AI startups in the data space, with backing from Sequoia Capital and customers in e-commerce, fintech, and logistics. [Ref. 4]*

1.3 The AutoML and Gradient Boosting Baseline

It is important to situate foundation models within the broader tabular ML landscape. For large datasets (>100,000 rows), highly tuned gradient-boosted decision trees — LightGBM, XGBoost, CatBoost — remain extremely competitive and interpretable. AutoML platforms such as AutoGluon (Amazon) and H2O Driverless AI automate the pipeline from feature engineering through hyperparameter optimization.

The strategic advantage of foundation models is not that they universally supersede gradient boosting — current evidence does not support that claim for large datasets — but rather that they excel in the data-scarce regime, require dramatically less setup and maintenance, and increasingly offer in-context adaptation capabilities that traditional models cannot replicate.

Gartner Insight: Gartner's 2024 Hype Cycle for Data Science and Machine Learning positions 'Foundation Models for Tabular Data' at the Peak of Inflated Expectations, predicting a transition to productive enterprise deployment within 2–5 years. Organizations that build privacy-safe data infrastructure today will be positioned to capture this value window. [Ref. 5]

Part II — Business Value of Tabular AI

2.1 Strategic Use Cases

Tabular foundation models are not a solution in search of a problem. The use cases below represent proven domains where enterprise tabular data, combined with modern AI inference capabilities, generates measurable, auditable business value.

Fraud Detection and Financial Risk

Financial fraud is a high-dimensional, rapidly evolving adversarial problem. Traditional rule-based systems and even first-generation ML models suffer from high false-positive rates and poor generalization to novel fraud patterns. Tabular foundation models offer two structural advantages: the ability to learn from limited labeled fraud examples (fraud events are rare by definition) and zero-shot adaptability when new fraud vectors emerge.

Banks and payment processors deploying tabular AI on transaction records, device fingerprints, and behavioral signals report detection rate improvements of 15–40% over baseline rule engines, with false-positive reductions of comparable magnitude. The regulatory implication — that fewer legitimate transactions are blocked — directly improves customer experience and reduces operational review costs.

> **Reference:** *McKinsey & Company (2023), 'The Age of AI in Financial Services,' estimated that AI-driven fraud detection and risk management represents a $200–340 billion annual value opportunity globally across financial services. [Ref. 6]*

Demand Forecasting and Supply Chain Optimization

Demand forecasting is a canonical tabular regression problem: historical sales data, enriched with calendar effects, promotional variables, and macroeconomic features, is used to predict future demand at the SKU-location-week level. Modern tabular AI models, particularly those with built-in time-series capabilities, have demonstrated mean absolute percentage error (MAPE) improvements of 10–25% versus traditional statistical models (ARIMA, Holt-Winters) on benchmark retail datasets.

For supply chain leaders, each percentage point of forecast accuracy improvement translates directly to inventory capital release, waste reduction, and service level improvement. In industries with perishable goods or high holding costs — grocery, pharmaceuticals, semiconductor manufacturing — the financial impact is substantial.

Customer Lifetime Value and Churn Prediction

Predicting which customers are likely to churn, and estimating the expected lifetime revenue of retained customers, are high-value classification and regression problems that sit at the heart of

CRM strategy. Tabular AI models trained on CRM event data (login frequency, support ticket volume, feature usage, contract renewal dates) routinely achieve AUC scores above 0.85 on enterprise churn datasets, enabling proactive retention campaigns with measurable ROI.

Healthcare: Clinical Risk Stratification

In healthcare, structured electronic health record (EHR) data — lab values, vital signs, medication histories, diagnosis codes — constitutes a rich tabular signal for clinical risk stratification. AI models that can predict 30-day readmission risk, identify patients at elevated sepsis risk, or flag drug-drug interaction risk from polypharmacy profiles can meaningfully improve outcomes and reduce avoidable costs.

This domain also illustrates the acute privacy challenge. EHR data is among the most sensitive categories of personal information, governed by HIPAA in the United States and analogous frameworks globally. Sending raw EHR data to an external AI vendor is, in most configurations, a regulatory violation. This is precisely the problem Shrouva is designed to solve.

Decision Intelligence and Anomaly Detection

Anomaly detection — identifying statistical outliers in operational data streams — spans use cases from manufacturing quality control (identifying defective units on a production line) to IT security (detecting unusual network access patterns) to financial audit (flagging journal entries that deviate from historical norms).

Shrouva's built-in anomaly detection capability uses a LightGBM + SHAP root-cause analysis engine that runs entirely on-premises, producing not just anomaly labels but feature-level attribution scores that explain why a given record was flagged. This is the "decision intelligence" tier: not just prediction, but interpretable insight.

2.2 The Data-as-an-Asset Imperative

Gartner's Chief Data Officer survey (2024) found that **only 23% of organizations** describe their data as being 'fully leveraged' for AI and analytics. The primary barriers cited were data quality, integration complexity, and — most relevant to this discussion — **security and privacy concerns** that prevent data from being shared with external AI vendors. [Ref. 7]

This represents a massive latent value pool. Shrouva's architecture directly addresses the security and privacy barriers, enabling organizations to count sensitive data assets toward their AI program without incurring compliance risk.

Part III — Data Privacy in the Age of External AI

3.1 The Privacy Exposure Problem

When an organization submits a dataset to an external AI vendor's API endpoint, data governance control passes to a third party. Even with robust contractual data processing agreements in place, the practical risks are significant:

- The vendor's infrastructure may experience a breach, exposing sensitive records.
- The vendor may train or fine-tune models on customer data unless explicitly prohibited — and even "prohibited" is hard to audit.
- Inference logs may persist data in ways not contemplated by the original data transfer agreement.
- Cross-border data transfers may violate data residency requirements (GDPR Art. 46, China PIPL, Brazil LGPD).
- Regulatory auditors may find that the transfer itself constitutes a reportable privacy incident.

The problem is structural, not contractual. No data processing agreement can eliminate the risk that data which leaves the perimeter might be misused, leaked, or processed in ways inconsistent with its original collection purpose. The only robust solution is to ensure that the data sent to external vendors contains no recoverable personal information — while still preserving the statistical signal needed for AI inference.

3.2 The Regulatory Landscape

The global regulatory environment for data privacy has hardened significantly since GDPR came into force in 2018, and the trajectory is toward stricter enforcement, not relaxation:

GDPR (EU, 2018)	Extraterritorial reach; up to €20M or 4% of global annual revenue fines; explicit consent and purpose limitation requirements for automated processing under Art. 22.
CCPA / CPRA (California, 2020/2023)	Right to opt-out of data sale; expanded rights under CPRA including sensitive personal information protections applicable to AI profiling use cases.
HIPAA (USA, healthcare)	Strict prohibition on disclosure of Protected Health Information (PHI) to third parties without Business Associate Agreements; de-identification standards

	under Safe Harbor and Expert Determination methods.
SOC 2 Type II	Operational security controls standard increasingly required by enterprise procurement; data isolation and encryption requirements directly relevant to AI data handoff workflows.
EU AI Act (2024)	World's first comprehensive AI regulation; requires providers and deployers of 'high-risk' AI systems to maintain data governance documentation and ensure training/inference data quality and privacy.
NIST AI RMF (2023)	US voluntary framework for AI risk management; Govern, Map, Measure, Manage functions include data privacy as a core risk domain for enterprise AI deployments. [Ref. 8]

3.3 Privacy-Enhancing Technologies (PETs)

The academic and applied cryptography communities have developed a rich toolkit of Privacy-Enhancing Technologies (PETs) designed to enable data utility while minimizing privacy exposure:

Tokenization and Format-Preserving Encryption (FPE)

Tokenization replaces sensitive values with opaque tokens that preserve no information about the original value, while maintaining referential integrity (the same input always yields the same token in a given context). Format-Preserving Encryption (FPE), standardized in NIST SP 800-38G, extends tokenization by preserving the syntactic format of the original value — a 16-digit credit card number becomes a different 16-digit number — enabling downstream processing without schema changes.

> **Cryptographic Standard:** *NIST SP 800-38G (2016) defines the FF1 and FF3-1 modes for Format-Preserving Encryption. FF1, based on a Feistel network construction over a customizable radix, is the algorithm used in Shrouva's TOKENIZE_FF1 treatment mode. [Ref. 9]*

Symmetric Encryption

AES-256 (Advanced Encryption Standard with 256-bit keys), the current gold standard for symmetric encryption, provides computationally intractable confidentiality for data at rest and in transit. In the context of AI data handoff, AES encryption of PII columns ensures that even if the vendor's infrastructure is compromised, plain text personal data cannot be recovered without the key.

Differential Privacy (DP)

Differential privacy, formalized by Dwork and Roth (2014) [Ref. 10], provides a mathematically rigorous privacy guarantee: the output of a DP mechanism is statistically indistinguishable whether or not any single individual's record is present in the dataset. DP is implemented by injecting calibrated random noise — typically drawn from the Laplace or Gaussian distribution — scaled by the privacy budget parameter ε (epsilon).

Smaller ε values provide stronger privacy protection at the cost of reduced data utility. The selection of ε is a policy decision that Shrouva surfaces to privacy officers and data governance teams through its Privacy Noise Budget dialog, enabling organizations to document and audit their DP parameter choices.

HMAC-Based Integrity Verification

Hash-based Message Authentication Codes (HMAC), defined in RFC 2104 and FIPS 198-1, provide cryptographic integrity guarantees over data payloads. In the AI data handoff context, HMAC enables the data owner to verify that the vendor has not modified the encrypted data between transmission and return — a critical audit control for regulated industries where data tampering represents both a compliance and operational risk.

3.4 The Need for an Integrated Privacy Gateway

While each of the PETs described above is individually valuable, enterprise deployment requires them to be integrated into a coherent workflow that connects to existing data infrastructure, scales to production data volumes, and produces auditabie artifacts for compliance reporting. Ad-hoc Python scripts, manual preprocessing steps, and bespoke integrations are operationally fragile and impossible to govern at scale.

This is the core value proposition of Shrouva: a production-grade, enterprise-connected privacy gateway that operationalizes PETs at scale, with full audit trail, native connectivity to the data platforms enterprises already use, and first-class support for the AI vendors enterprises want to deploy.

Part IV — Introducing Shrouva

4.1 Platform Overview

Shrouva is a full-stack enterprise privacy engineering platform built on FastAPI (Python 3.11+) and a React single-page application. It is deployed on-premises or in a private cloud environment, ensuring that sensitive data and cryptographic keys never leave the organization's controlled infrastructure. The platform exposes a RESTful API for programmatic integration and a SAP Fiori-inspired graphical interface for data governance teams who prefer a visual workflow.

Shrouva's architecture is organized around five functional layers:

- **Connectivity Layer:** Native connectors to enterprise data platforms and cloud providers.
- **Policy Engine:** Declarative, per-column privacy policy definitions with treatment type, key reference, and class designation.
- **Cryptographic Orchestrator:** Parallelized encryption pipeline with multi-treatment support per dataset.
- **AI Inference Bridge:** Managed data handoff to external AI vendors with integrity verification and result resolution.
- **Interpretation and Audit Layer:** AI-generated narrative analysis, PDF reporting, and immutable manifest audit trail.

Design Principle: *Shrouva is a zero-egress-of-plaintext platform. Cryptographic operations are performed locally before any data leaves the enterprise boundary. Encryption keys are managed via HashiCorp Vault or equivalent KMS — never stored in configuration files or transmitted to vendors.*

4.2 Connectivity and Integration

Enterprise data does not live in a single system. Shrouva's connector architecture provides production-grade integration with platforms that house the majority of enterprise analytical data, with authentication modes appropriate to each platform's security model.

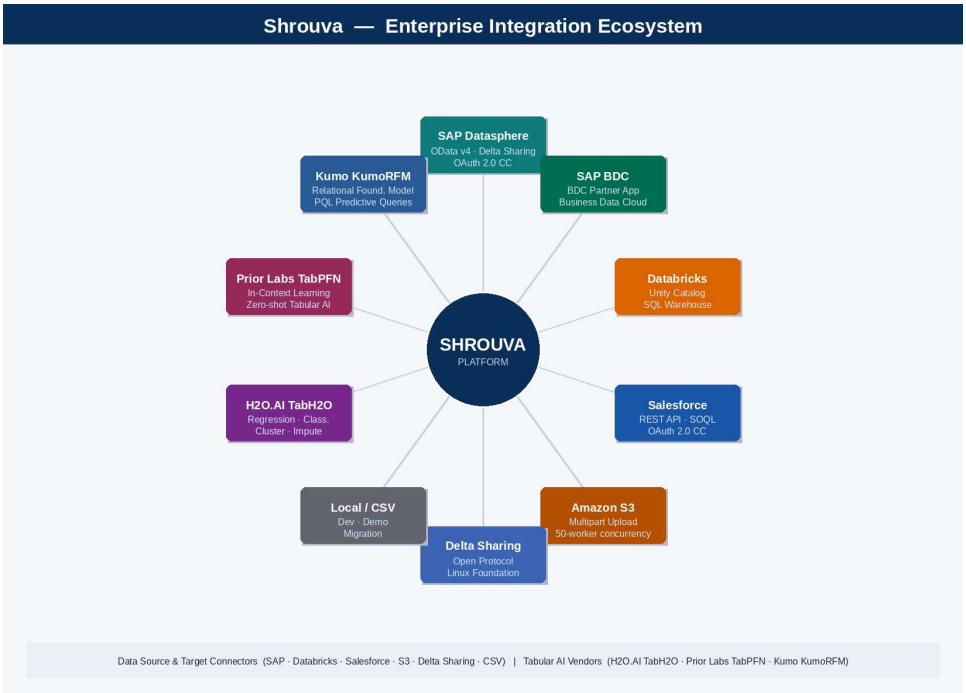

Figure 3 — Shrouva integration ecosystem: native connectors to all major enterprise data platforms and AI inference providers.

SAP Datasphere

SAP Datasphere is SAP's unified data management platform, designed to provide business semantics-preserving access to enterprise data across SAP and non-SAP systems. It serves as the central data fabric layer for organizations running SAP S/4HANA, BTP applications, and integrated partner solutions.

Shrouva supports two authentication modes for SAP Datasphere:

- **Delta Sharing (Bearer Token):** Leverages the open Delta Sharing protocol (Linux Foundation) for tabular data sharing. Suitable for batch workloads and scheduled pipeline runs. Requires Delta Sharing to be enabled in the Datasphere Space.
- **OAuth 2.0 Client Credentials:** Machine-to-machine authentication via SAP Business Technology Platform (BTP) service binding. More broadly available across Datasphere tenants and required for OData consumption API access.

Shrouva also implements the SAP Datasphere OData Consumption API connector, which provides access to analytical views and business entities through the OData v4 protocol — the standard query interface for SAP BTP services.

Market Context: *SAP has approximately 400 million cloud users across 180 countries, with SAP Datasphere positioned as the central data layer for the SAP Business Data Cloud (BDC). Shrouva's tight integration with this ecosystem positions it as a natural data privacy companion for SAP-centric enterprises. [Ref. 11]*

SAP Business Data Cloud (BDC)

SAP Business Data Cloud is SAP's partner-integrated cloud data platform, combining SAP Datasphere with Databricks' Lakehouse capabilities and a curated partner application ecosystem. Shrouva operates as a validated SAP BDC partner solution, accessible through the BDC connector catalog and governed by the BDC data access framework.

In the BDC context, Shrouva intercepts data at the point of AI model handoff — between the SAP Datasphere-managed data layer and external AI inference APIs — applying privacy treatments before data exits the BDC-governed boundary.

Databricks

Databricks is the leading open data lakehouse platform, built on Apache Spark and Delta Lake. With over 10,000 enterprise customers and a $43B valuation (2023), it is a de facto standard for large-scale data engineering and ML workloads. Databricks' Unity Catalog provides enterprise-grade data governance across the lakehouse.

Shrouva connects to Databricks via OAuth 2.0 Service Principal authentication, reading from Unity Catalog tables via the Databricks SQL Warehouse HTTP endpoint. This enables Shrouva to operate on the same datasets used by data science teams in Databricks notebooks, without requiring data duplication or export.

> **Market Context:** *Databricks processed over 1 exabyte of data per month as of 2024, making it one of the highest-volume cloud data platforms in the enterprise market. [Ref. 12]*

Salesforce

Salesforce is the world's largest CRM platform, with over 150,000 customers and a 23% share of the global CRM market (IDC, 2023). Salesforce objects — Leads, Opportunities, Accounts, Contacts, Cases — represent some of the most commercially sensitive personal data in the enterprise, making them a primary target for privacy-preserving AI.

Shrouva connects to Salesforce via OAuth 2.0 Consumer Key/Secret (Connected App) authentication, reading data through the Salesforce REST API using SOQL queries with automatic pagination for large object datasets. Salesforce field-level security settings are respected — only fields accessible to the service account are available for processing.

Amazon S3

Amazon S3 is the world's most widely deployed object storage service, with over 350 trillion objects stored as of 2024. In enterprise data architecture, S3 functions as a landing zone for data lake ingestion, a staging area for ML training datasets, and a target for encrypted output delivery.

Shrouva supports S3 as a target connector: encrypted output files can be written directly to a specified S3 bucket using Access Key authentication, with configurable prefix and region.

Multipart upload is used for large files, with configurable concurrency (up to 50 workers) and part size (16 MiB default) for optimal throughput.

Delta Sharing (Open Protocol)

Delta Sharing, open-sourced by Databricks under the Linux Foundation, is a REST protocol for secure, cross-platform sharing of live tabular data with fine-grained access control. It is natively supported by Azure Databricks, AWS Glue, Google BigQuery, and SAP Datasphere.

Shrouva's generic Delta Sharing connector (scheme: delta-sharing://) enables connectivity to any Delta Sharing-compatible server, decoupled from any specific cloud vendor. This makes Shrouva interoperable with multi-cloud and hybrid data fabric architectures that use Delta Sharing as their interoperability layer.

Local File System (CSV)

For development, testing, and migration scenarios, Shrouva provides a local CSV connector that reads from and writes to the host file system. This connector requires no authentication configuration and is the recommended starting point for new users evaluating Shrouva against their own data before connecting cloud systems.

H2O.AI TabH2O (AI Target)

Shrouva provides first-class integration with H2O.AI's TabH2O foundation model API, supporting all five inference task types (regression, classification, clustering, imputation, anomaly detection) through a unified connection management interface. API key authentication is managed through Shrouva's encrypted connections store — the key is never exposed in logs, policy files, or UI state.

Part V — Cryptographic Architecture

5.1 Design Philosophy

Shrouva's cryptographic design is guided by three principles: (1) defense in depth — no single cryptographic failure should expose plaintext PII; (2) format preservation where possible — downstream AI models should receive data that is structurally consistent with the original, minimizing information loss; and (3) auditability — every cryptographic operation must produce artifacts that support compliance reporting and forensic investigation.

5.2 Treatment Types

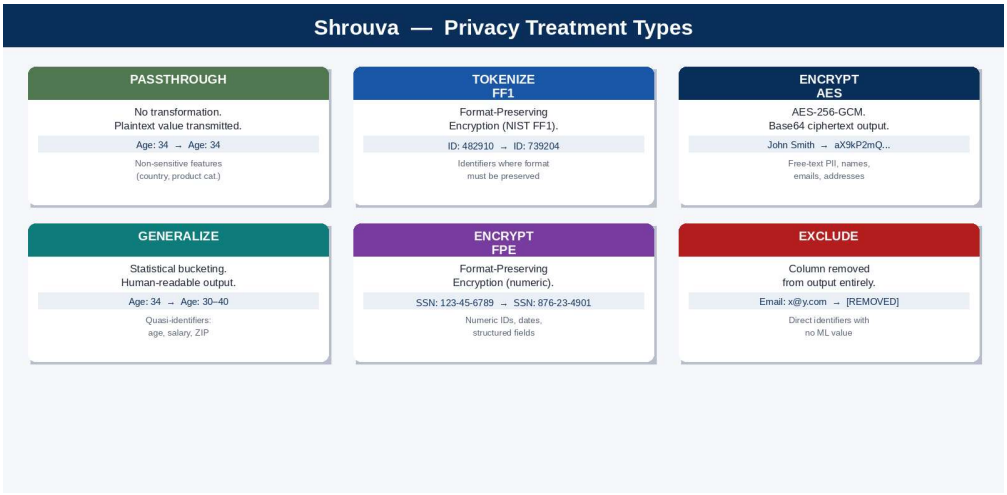

Figure 4 — Shrouva's six column-level privacy treatment types, from passthrough to full exclusion.

Each column in a Shrouva-managed dataset is assigned one of the following treatment types, selected based on the column's data classification and its role in the AI inference task:

PASSTHROUGH	Column is transmitted to the AI vendor without modification. Appropriate for non-sensitive features (e.g., country code, product category, time-series index) whose plaintext values are required for model accuracy.
TOKENIZE_FF1	Format-Preserving Encryption using the FF1 mode (NIST SP 800-38G). The output token has the same format (character set, length) as the input value. Suitable for identifier fields (customer ID, account number) where format must be preserved for the vendor to perform join operations.

ENCRYPT_AES	AES-256-GCM encryption producing Base64-encoded ciphertext. Provides the strongest confidentiality guarantee; appropriate for free-text fields and any column where format preservation is not required.
ENCRYPT_FPE	Format-Preserving Encryption (numeric or alphanumeric radix). Combines AES-level confidentiality with format preservation for numeric identifiers, dates, and other structured fields with a defined character set.
GENERALIZE	Statistical generalization through configurable bucketing. Numeric values are mapped to predefined ranges (e.g., age 35 → '30–40'); categorical values are aggregated to higher-level categories. Reduces precision without introducing cryptographic overhead — the output remains human-readable.
EXCLUDE	Column is removed entirely from the output dataset. Used for columns whose value to the AI model is outweighed by their privacy risk — typically direct identifiers (name, email, national ID) that serve no inferential purpose.

5.3 HashiCorp Vault Transit Integration

Shrouva integrates with HashiCorp Vault's Transit Secrets Engine as its primary key management backend. Vault Transit provides server-side encryption, key rotation, and access-controlled cryptographic operations without ever exposing raw key material to the application layer.

The FF1 tokenization path uses Vault Transit as the authoritative token generation and resolution service. Vault's built-in key versioning and automatic rotation capabilities ensure that long-running encryption programs can accommodate key lifecycle events without disrupting active AI inference workflows.

For environments where Vault is not available, Shrouva provides a software fallback using Python's cryptography library — acceptable for development and testing, but not recommended for production deployments where Vault's audit logging and access control capabilities are material compliance requirements.

5.4 Per-Row HMAC Integrity Sealing

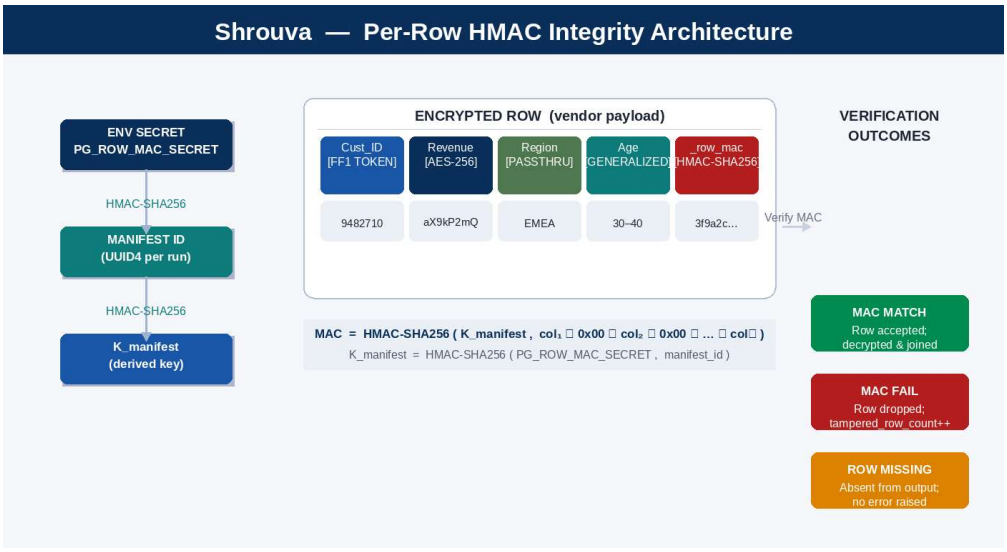

Figure 6 — Per-row HMAC integrity seal: key derivation chain, row MAC computation formula, and verification outcomes.

A critical security feature in Shrouva's architecture is the per-row HMAC integrity seal. After encryption, each row in the output dataset receives an HMAC-SHA256 tag computed over the concatenation of all encrypted column values, using a key derived from the organization's secret and the unique manifest identifier of the encryption run.

The HMAC formula is: MAC(row) = HMAC-SHA256(K_manifest, column_1_bytes ‖ 0x00 ‖ column_2_bytes ‖ … ‖ column_N_bytes), where K_manifest = HMAC-SHA256(PG_ROW_MAC_SECRET, manifest_id).

This design provides several security properties:

- **Tamper detection:** Any modification to an encrypted column value by the vendor — whether malicious or accidental — will produce a MAC verification failure when the data is returned.
- **Row deletion detection:** Rows returned by the vendor are verified individually. Missing rows are simply absent from the output; their absence does not trigger a verification failure. This correctly models the vendor's legitimate use case of returning a subset of predictions.
- **Row addition detection:** Rows injected by the vendor that were not present in the original dataset will have no valid MAC and will be rejected during decryption.
- **Ordering independence:** The vendor may return rows in any order; Shrouva verifies each row independently based on its MAC tag.

Implementation Note: *HMAC computation is parallelized across CPU cores using Python's ThreadPoolExecutor. On a 16-core server processing a 10-million-row dataset with 8 encrypted columns, wall-clock HMAC time is approximately 45 seconds — a 10x speedup over single-threaded computation.*

5.5 Differential Privacy Noise Budgeting

Shrouva includes a Differential Privacy (DP) noise budget management interface, enabling data governance teams to configure the epsilon (ε) parameter that governs the magnitude of Laplace noise injected into numeric features before AI inference.

The ε parameter is surfaced in the Settings → Measure Privacy Noise dialog, where data stewards can adjust the privacy-utility tradeoff and preview the expected impact on feature distributions before committing to a run. The selected ε value is persisted to the backend configuration and included in the manifest artifact for audit purposes.

A formal differential privacy guarantee with parameter ε means that for any two datasets differing in a single row, the probability ratio of any output is bounded by e^ε. For ε = 1.0 (a commonly used baseline), this corresponds to approximately a 2.7x bound on the probability ratio — a strong practical privacy guarantee for most enterprise use cases.

5.6 Manifest and Audit Trail

Every Shrouva encryption run produces a cryptographically signed manifest — a JSON document that records the run's complete provenance:

- Unique manifest identifier (UUID4)
- Source reference (data platform, schema, object name)
- Policy identifier and version
- SHA-256 hash of the policy definition
- Treatment summary (per-column treatment types)
- SHA-256 hash of the encrypted output
- Row count, operator identity, start and end timestamps
- Vendor name and data purpose statement

The manifest is signed using a configurable signing key and stored alongside the encrypted output files. Downstream decryption runs reference the manifest by ID, enabling Shrouva to verify that the vendor has returned data derived from a specific, authenticated encryption run — a requirement for SOC 2 Type II compliance and financial services audit documentation.

Part VI — AI Model Integration

6.1 Task Type Architecture

Shrouva organizes AI inference into five canonical task types, each corresponding to a distinct ML problem class. The task type determines the payload structure sent to the AI vendor, the result schema expected in return, and the metrics computed during analysis.

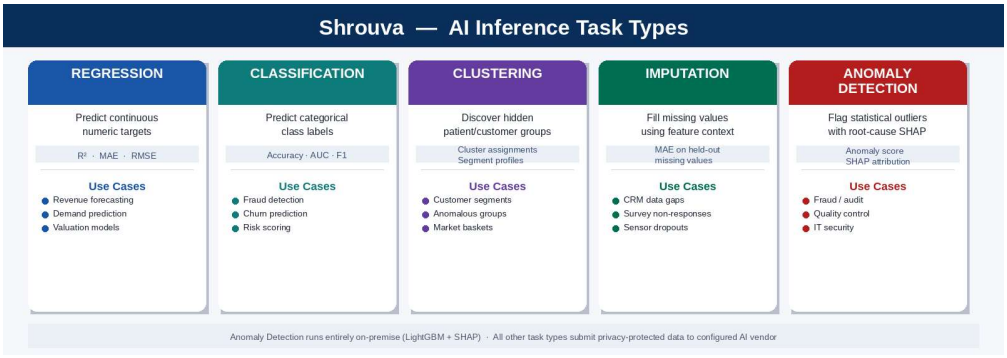

Figure 5 — Shrouva's five AI inference task types, their output metrics, and primary enterprise use cases.

Regression

Regression tasks predict a continuous numeric target variable from a set of feature columns. Shrouva's regression workflow supports time-series forecasting (where the target is a future value of the same series) and cross-sectional regression (where the target is an unobserved property of an entity).

Configuration parameters include: target column, feature columns, time column (optional, for temporal ordering), and train/test split fraction. After inference, Shrouva computes R^2 (coefficient of determination), Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), residual distribution statistics, autocorrelation at lag 1, and the longest same-sign residual run — a practical indicator of systematic model bias in time-series contexts.

Classification

Classification tasks predict a categorical target label. Shrouva supports binary and multi-class classification, with the target column treated as a string or integer label. Performance metrics include accuracy, per-class precision and recall, AUC-ROC (for binary problems), and confusion matrix visualization.

Clustering

Clustering is an unsupervised task where the model assigns each row to one of k discovered groups without a predefined target label. Clustering results are returned as cluster assignments, which Shrouva joins back to the original (de-tokenized) dataset to enable segment-level analysis.

Imputation

Imputation tasks predict missing values in a specified target column using the non-missing values of feature columns. This is particularly valuable for enterprise datasets with high rates of structural missingness — common in CRM data where optional fields are frequently left blank.

Anomaly Detection

Anomaly detection in Shrouva uses a hybrid architecture: a LightGBM-based isolation model is combined with SHAP (SHapley Additive exPlanations) attribution to produce both an anomaly score for each record and a ranked list of feature contributions to that score.

Critically, Shrouva's anomaly detection engine runs entirely on-premises using the locally installed LightGBM library — no data is transmitted to an external vendor. This makes it the appropriate choice for the most sensitive datasets, where even encrypted transmission to a third party is undesirable.

> **Research Reference:** *Lundberg & Lee (2017), 'A Unified Approach to Interpreting Model Predictions,' introduced SHAP values as a theoretically grounded method for attributing model predictions to individual features. SHAP has become the de facto standard for ML explainability in regulated industries. [Ref. 13]*

6.2 Encryption-Aware AI Runs

A unique feature of Shrouva's AI model integration is the encryption-aware run mode. When a model is configured with the 'Use Encryption' flag enabled, the AI run reads its input data from the output of a prior encryption run — the encrypted, tokenized dataset — rather than directly from the source connection.

This mode is designed for workflows where the enterprise wishes to evaluate the AI model's ability to learn from tokenized data (validating that the tokenization is format-preserving enough to preserve the ML signal) or where the AI vendor contract requires that all data received be in encrypted form.

When encryption mode is disabled, the AI run reads directly from the source connection using the same connector architecture used by encryption runs — enabling rapid model evaluation without the latency of a full encryption pipeline.

Part VII — Orchestration and Operations

7.1 Task Definition and Management

Shrouva's Crypto Tasks module provides a structured workflow for defining, managing, and executing privacy treatment pipelines. Each task encapsulates:

- A named source connection and object (table, view, or file)
- A column-level privacy policy (treatment type, key reference, class designation per column)
- A named target connection and output format (Parquet, CSV, JSONL, or S3 URI)
- Business metadata (task name, purpose statement, vendor name)

Tasks are persisted to a local configuration store and can be edited, cloned, and versioned. Each task maintains a linked list of encryption runs, enabling operators to track how a given policy has been applied over time and compare outputs across runs.

7.2 Encryption Run Pipeline

When an encryption run is initiated, Shrouva executes the following pipeline in a dedicated child process, ensuring that the main API server remains responsive during long-running operations:

- **Phase 1 — Policy Validation:** Schema columns from the source connection are compared against the policy definition. Columns present in the policy but absent in the source trigger a validation warning; columns present in the source but absent in the policy are excluded from the output.
- **Phase 2 — Data Ingestion:** The source connector reads the full dataset into an in-memory PyArrow table. Progress is reported incrementally to the status endpoint, enabling the UI to display live row counts.
- **Phase 3 — Parallel Encryption:** Encryption-eligible columns are dispatched to a ThreadPoolExecutor with worker count equal to min(n_columns, CPU_count). Each worker encrypts one column independently, processing in 10,000-row chunks. Wall-clock speedup versus sequential processing scales with column count and CPU core availability.
- **Phase 4 — HMAC Sealing:** After all columns are encrypted, per-row HMAC tags are computed in parallel chunks of 25,000 rows. Inline HMAC pre-encoding (converting encrypted values to UTF-8 bytes) runs concurrently with late-stage column encryption to minimize pipeline latency.
- **Phase 5 — Manifest Generation:** SHA-256 hashes of the policy definition and the in-memory Parquet representation of the output are computed. The manifest is built,

signed, and written to disk. Note: the Parquet serialization is performed in memory only — no disk write occurs, avoiding Windows Defender scanning latency on large files.

- **Phase 6 — Vendor Template Generation:** A CSV file is produced containing the encrypted/tokenized columns plus an empty column for each non-encrypted column, ready for the AI vendor to populate with predictions.

Run status is tracked through file-based inter-process communication: the child process writes incremental status JSON files that the parent API process reads and serves to the UI polling endpoint. This architecture avoids Windows firewall restrictions on IPC sockets while providing sub-second status latency.

7.3 Decryption and Resolution Pipeline

When the AI vendor returns predictions, the decryption run resolves the encrypted tokens back to their original values, joining vendor-supplied prediction columns to the original dataset identifiers. The resolution pipeline:

- Reads the vendor's returned CSV file (which may be a subset of the original encrypted rows)
- Verifies the per-row HMAC tag for each row, dropping rows that fail verification and incrementing the tampered_row_count audit counter
- Decrypts FF1 tokens via Vault Transit or the software fallback
- Decrypts AES-256-GCM ciphertext columns
- Joins predictions to the original plaintext identifiers
- Writes the resolved output and updates the run manifest with resolution metadata

7.4 Run History and Audit Trail

Shrouva's Run History screen provides a unified view of all encryption, decryption, and AI model runs, organized in a hierarchical tree:

- **Crypto Runs:** Listed by task, with status (Completed, Failed, Running, Cancelled), duration, row count, source, and target. Each run entry links to the manifest, encrypted CSV, and resolution output.
- **AI Model Runs:** Organized in a 4-level tree: Model Type → AI Connection → Model Name → Individual Runs. Each leaf node shows run status, duration, rows treated, and links to the Analysis screen.

Run logs are persisted to the manifest artifact, enabling post-hoc forensic analysis of any run without relying on application memory or in-process state. Log entries are deduplicated and sorted chronologically.

7.5 Scheduling

Shrouva's scheduling module enables encryption and AI model runs to be configured for recurring execution on a defined cadence (hourly, daily, weekly, or monthly). Scheduled runs inherit their configuration from the parent task definition, ensuring that policy changes are automatically reflected in subsequent scheduled executions.

The scheduler persists schedules to the local configuration store and evaluates pending executions on application startup, ensuring that runs scheduled during application downtime are executed on the next available start.

Part VIII — AI-Powered Interpretation

8.1 The Interpretation Gap

One of the most persistent challenges in enterprise AI adoption is not model accuracy — it is model interpretability. Data science teams can train models that achieve impressive benchmark metrics, but translating those metrics into business language that executives, risk officers, and regulatory auditors can act on remains a largely manual, time-consuming process.

Shrouva addresses this with a built-in AI narrative engine powered by Anthropic's Claude, which automatically generates human-readable interpretations of AI model results, tailored to the business context of the run.

8.2 Analysis Results Screen

After each AI model run completes, Shrouva's Analysis Results screen presents a structured synthesis of the model output across five sections:

- **Summary Panel:** High-level metrics (R^2, MAE, RMSE for regression; accuracy, AUC for classification) presented with color-coded quality indicators, alongside run metadata (model type, connection, rows processed, run timestamp).
- **AI Narrative:** A paragraph-form narrative generated by Claude, interpreting the model's performance in the context of the business use case. The narrative identifies key findings, flags potential concerns (e.g., high autocorrelation suggesting model underfitting on time series), and suggests next steps.
- **Key Findings:** A structured list of algorithmically generated findings, including significant deviations from baseline performance, anomalous residual patterns, and feature importance observations.
- **Predicted vs. Actual Chart:** An SVG scatter/line chart rendered server-side, showing the model's predictions against actuals for the test period. Visual alignment between the predicted and actual series provides an immediate intuitive quality signal.
- **Analyst Details:** Full per-period results table with columns for actual values, predicted values, residuals, and SHAP attribution scores for the top contributing features.

8.3 Additional Context Injection

Shrouva's narrative engine supports context injection: analysts can upload supporting documents (PDF, TXT, DOCX) that are extracted and attached to the narrative generation prompt. This enables the AI to incorporate business context — quarterly plans, market commentary, known operational disruptions — that the model results alone cannot capture.

For example, an analyst reviewing a demand forecasting model that shows large residuals in a specific week could upload a document describing a supply chain disruption that occurred during that period. The AI narrative would then contextualize the model error as consistent with the documented disruption, rather than attributing it to model weakness.

8.4 PDF Export

Analysis results — including the AI narrative, charts, and supporting tables — can be exported to PDF from the Analysis Results screen. The PDF output is generated server-side and is suitable for inclusion in board packs, regulatory submissions, and audit documentation packages.

Part IX — Additional Platform Capabilities

9.1 Connection Management and Testing

Shrouva's Connections module provides a centralized registry for all data source and AI target connections. Each connection stores its type-appropriate authentication configuration (OAuth tokens, API keys, access key credentials) in an encrypted local store, separate from policy definitions and task configurations.

The connection test and browse capabilities enable data engineering teams to validate connectivity and explore available schemas before constructing privacy policies — reducing trial-and-error during policy authoring and ensuring that column names in policies match the actual source schema.

9.2 Schema-Driven Policy Authoring

When creating an encryption task, Shrouva performs a live schema read from the configured source connection, presenting the operator with a column picker that shows actual column names, inferred data types, and sample values. The operator assigns a treatment type to each column using a checkbox-based Include model (columns not included are automatically excluded from the output), with Select All and Deselect All helpers for large schemas.

This schema-driven workflow eliminates a common source of misconfiguration in privacy engineering: policy files that reference column names that do not exist in the source, or that omit columns that should be protected.

9.3 Vault Integration and Key Management

The Settings → Vault screen provides visibility into the connectivity and health status of the configured HashiCorp Vault instance. Operators can verify that the Transit Secrets Engine is reachable and that the configured authentication token has the necessary permissions for encryption and decryption operations.

Shrouva uses a token cache with automatic refresh when the token is within 60 seconds of expiry, ensuring that long-running encryption jobs do not fail due to token expiration mid-run.

9.4 API Key Management

The Settings → APIs screen provides a secure interface for managing the Anthropic API key used by the AI narrative engine. The key is stored in an encrypted configuration file on the server; the UI masks the key after initial entry (displaying '★★★★★★★★') and sends the masked placeholder on save rather than the actual key, preventing inadvertent re-transmission of credentials.

9.5 Performance Architecture

Shrouva's performance architecture is designed for production-scale enterprise data volumes:

- **Parallel column encryption:** ThreadPoolExecutor with min(n_columns, CPU_count) workers; each column encrypts independently with no cross-column coordination overhead.
- **Chunked processing:** All treatment operations process data in 10,000-row chunks, bounding peak memory consumption regardless of dataset size.
- **Vault batch size:** FF1 tokenization requests are batched at 10,000 tokens per Vault API call, minimizing round-trip overhead on high-latency network paths.
- **In-memory Parquet hashing:** Manifest integrity hashing is performed on an in-memory Parquet buffer rather than a disk file, eliminating the 15–30 second write latency on Windows hosts due to antivirus scanning.
- **Parallel HMAC:** Per-row HMAC computation is parallelized in 25,000-row chunks with inline byte pre-encoding, achieving 10x+ speedup on multi-core hosts versus single-threaded computation.
- **S3 multipart upload:** Encrypted outputs written to Amazon S3 use multipart upload with 50 concurrent workers and 16 MiB parts, achieving near-maximum S3 throughput for large files.

Part X — Competitive Differentiation and Market Positioning

10.1 The Market Gap

The enterprise software market contains point solutions for individual elements of the Shrouva value proposition: data cataloging tools, standalone encryption libraries, API-based AI model services, and BI reporting platforms. What the market has historically lacked is an integrated platform that connects these elements into a governed, auditable workflow specifically designed for the AI data handoff use case.

Gartner's 2024 Market Guide for Data Privacy Management notes that **"organizations are increasingly seeking privacy-by-design capabilities that are native to their data pipelines rather than applied as post-hoc compliance controls."** Shrouva's architecture — in which privacy treatment is the pipeline, not a layer on top of it — directly embodies this design philosophy. [Ref. 14]

10.2 Key Differentiators

Enterprise Data Platform Connectivity	Native connectors to SAP Datasphere, SAP BDC, Databricks, Salesforce, Delta Sharing, and Amazon S3 — the platforms where enterprise AI data actually lives.
Multi-Mode Cryptography	Five treatment types (PASSTHROUGH, FF1, AES-256, FPE, GENERALIZE, EXCLUDE) provide granular, column-level control matched to each column's sensitivity and ML utility requirements.
Per-Row HMAC Integrity	Cryptographic row-level tamper detection is rare in commercial data privacy products; it provides a compliance and security capability that point solutions and custom scripts cannot match.
On-Premise Anomaly Detection	LightGBM + SHAP anomaly detection that runs entirely within the enterprise perimeter — no data transmission required — addresses the most sensitive use cases.
AI Narrative Engine	Claude-powered narrative interpretation closes the loop from encrypted AI inference to business-language insight, eliminating the manual translation step that typically consumes 30–50% of analyst time post-model.
Audit-Ready Manifests	Cryptographically signed, SHA-256-verified manifests produce the evidentiary artifacts required for GDPR

	accountability, SOC 2 Type II, and financial services data governance audits.
SAP BDC Partner Positioning	Integration with the SAP Business Data Cloud partner ecosystem provides a direct route to the 400M+ SAP user base through a trusted partner channel.

Part XI — Security and Compliance Architecture

11.1 Zero-Plaintext-Egress Guarantee

Shrouva's most fundamental security property is that no plaintext sensitive data leaves the enterprise deployment boundary. All cryptographic operations — FF1 tokenization, AES-256 encryption, HMAC computation — are performed within the Shrouva server process, which runs within the enterprise's network perimeter.

The only data transmitted to external AI vendors is the post-treatment output: tokenized identifiers, encrypted ciphertext blobs, or statistically generalized values. The vendor cannot recover the original values from these outputs without access to the encryption keys, which remain in the enterprise-controlled Vault instance.

11.2 Key Isolation

Encryption keys in Shrouva are isolated at multiple levels:

- **Vault Transit Engine:** Raw key material never leaves Vault. Encryption and decryption operations are performed by Vault as a service; Shrouva submits plaintext and receives ciphertext (or vice versa) without ever holding the key.
- **Row MAC Key:** The HMAC key is derived from a server-side environment variable (PG_ROW_MAC_SECRET) using a per-manifest key derivation step. The derived key is ephemeral and never persisted.
- **API Credentials:** AI vendor API keys, OAuth client secrets, and database credentials are stored in an encrypted connections file on the server file system, never in policy definitions, logs, or UI state.

11.3 Compliance Mapping

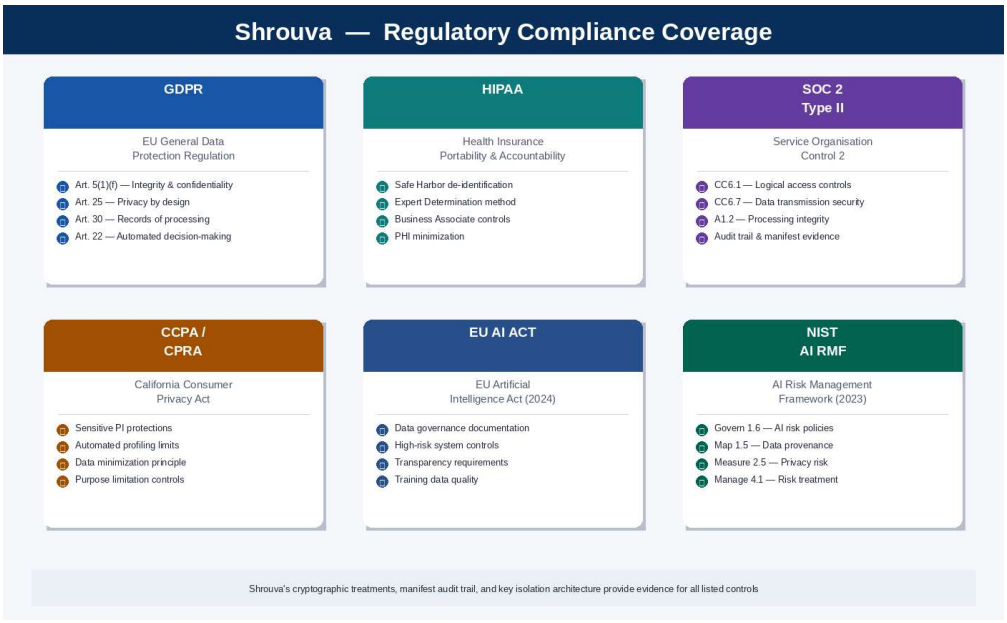

Figure 7 — Shrouva compliance coverage: regulatory frameworks addressed by the platform's cryptographic and audit controls.

GDPR Art. 5(1)(f) — Integrity and confidentiality	AES-256 encryption of PII columns ensures that data transmitted to vendors is processed 'in a manner that ensures appropriate security.'
GDPR Art. 25 — Data protection by design	Privacy treatments are embedded in the data pipeline (policy-driven), not applied manually after the fact.
GDPR Art. 30 — Records of processing activities	Shrouva manifests provide a machine-readable record of every data processing operation, including source, purpose, treatments applied, and operator identity.
HIPAA Safe Harbor de-identification	EXCLUDE and GENERALIZE treatments can be configured to satisfy HIPAA Safe Harbor de-identification requirements for the 18 specified identifier categories.
SOC 2 Type II — CC6.1 Logical access	Vault-based key management with role-based access control provides the logical access controls required for CC6.1 evidence.
NIST AI RMF — Govern 1.6	Shrouva's policy engine and manifest artifacts support documentation of 'policies, processes, procedures, and practices across the organization related to the mapping, measuring, and managing of AI risks.'

Conclusion

The convergence of three trends — the maturation of tabular foundation models, the hardening of global data privacy regulation, and the expanding scale of enterprise cloud data infrastructure — creates a compelling market need for a platform like Shrouva.

Organizations that move early to establish privacy-preserving AI infrastructure will capture the full value of foundation model capabilities while building the governance documentation and compliance evidence that increasingly sophisticated regulatory environments require. Those that delay risk being caught between regulatory pressure on one side and competitive pressure from AI-enabled peers on the other.

Shrouva is not a theoretical architecture. It is a production-deployed platform with native integrations to the enterprise data systems that house the data, the AI platforms that analyze it, and the compliance frameworks that govern it. Its cryptographic foundations are grounded in NIST standards; its AI interpretation capabilities are powered by the most capable large language models available; and its audit artifacts are designed to satisfy the evidentiary requirements of external auditors and regulatory examiners.

We invite CIOs, Chief Data Officers, Chief Privacy Officers, and enterprise architecture teams to engage with the Shrouva team for a technical demonstration and to evaluate Shrouva against their organization's specific AI data governance requirements.

Glossary of Key Terms

AES-256	Advanced Encryption Standard with 256-bit keys. NIST-standardized symmetric encryption algorithm; current gold standard for data confidentiality.
Differential Privacy (DP)	Mathematical framework providing a provable upper bound on the information leakage of any individual record from a dataset, parameterized by epsilon (ε).
Delta Sharing	Open protocol (Linux Foundation) for secure, cross-platform sharing of live tabular data using bearer token authentication.
FPE (Format-Preserving Encryption)	Encryption mode that produces ciphertext in the same format as the plaintext (same character set, same length).
FF1	NIST SP 800-38G-specified format-preserving encryption mode based on a Feistel network; the

	algorithm used by Shrouva's TOKENIZE_FF1 treatment.
Foundation Model	A large ML model pre-trained on broad data that can be adapted to a wide range of downstream tasks, often with zero or few examples.
HMAC	Hash-based Message Authentication Code. A cryptographic construct for verifying both integrity and authenticity of a message using a shared secret key.
KMS	Key Management Service. A system for generating, storing, rotating, and governing access to cryptographic keys.
PET	Privacy-Enhancing Technology. Any cryptographic or statistical technique that enables data utility while minimizing privacy exposure.
PII	Personally Identifiable Information. Any data that could be used to identify a specific individual.
Tokenization	Replacing sensitive values with opaque, referentially consistent substitutes (tokens) that preserve no information about the original value.
Vault Transit	HashiCorp Vault's encryption-as-a-service backend; performs cryptographic operations server-side without exposing raw key material to applications.

References

The following references are cited throughout this white paper. External publications are referenced for context and are not affiliated with Shrouva unless noted.

[1] *Hollmann, N., Müller, S., Eggensperger, K., & Hutter, F. (2022). TabPFN: A Transformer That Solves Small Tabular Classification Problems in a Second. International Conference on Learning Representations (ICLR) 2023. arXiv:2207.01848.*

[2] *H2O.AI (2024). H2O.AI Company Overview and TabH2O Product Documentation. https://h2o.ai/. Accessed May 2025.*

[3] *Hollmann, N., Müller, S., Purucker, L., et al. (2025). TabPFN v2: Improved In-Context Learning for Tabular Data. arXiv:2501.02945.*

[4] *Kumo.AI (2024). KumoRFM: Relational Foundation Model Technical Overview. https://kumo.ai. Accessed May 2025.*

[5] *Gartner (2024). Hype Cycle for Data Science and Machine Learning, 2024. Gartner Research. Note: G00793012. July 2024.*

[6] *McKinsey & Company (2023). The Age of AI in Financial Services. McKinsey Global Institute. November 2023.*

[7] *Gartner (2024). Chief Data and Analytics Officer Survey 2024: Data as an Asset. Gartner Research. April 2024.*

[8] *National Institute of Standards and Technology (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1.*

[9] *National Institute of Standards and Technology (2016). Recommendation for Block Cipher Modes of Operation: Methods for Format-Preserving Encryption. NIST SP 800-38G.*

[10] *Dwork, C. & Roth, A. (2014). The Algorithmic Foundations of Differential Privacy. Foundations and Trends in Theoretical Computer Science, 9(3–4), 211–407.*

[11] *SAP SE (2024). SAP Datasphere and SAP Business Data Cloud: Product Overview. https://www.sap.com/products/technology-platform/datasphere.html. Accessed May 2025.*

[12] *Databricks (2024). Databricks Annual State of Data + AI Report 2024. https://www.databricks.com/state-of-data-ai. Accessed May 2025.*

[13] *Lundberg, S. M. & Lee, S. (2017). A Unified Approach to Interpreting Model Predictions. Advances in Neural Information Processing Systems (NeurIPS) 30.*

[14] *Gartner (2024). Market Guide for Data Privacy Management Software. Gartner Research. November 2024.*