Rhino
Federated Computing

contact@rhinofcp.com
22 Boston Wharf Rd.
Boston, MA 02210

Federated Computing in Banking: A New Paradigm for Secure Data Collaborations & AI Development

Introduction

In today's competitive financial services environment, data has become the lifeblood of innovation. From advanced analytics to personalized offerings, banks and their partners seek to harness data for sharper risk management, more robust fraud detection, and differentiated customer experiences. Yet, practical and regulatory hurdles—such as siloed infrastructures, privacy rules, and security obligations—have made large-scale data collaboration complex and risky.

Federated Computing (FC) now offers a compelling alternative: it enables insights to be shared, correlated, and acted upon without exposing the underlying data outside each participant's environment. By maintaining each organization's strict privacy and governance requirements, FC paves the way for an entirely new era of secure, data-driven initiatives in finance.

This white paper is the first in a series describing the impact Federated Computing will have on financial services. Here we provide an introduction to Federated Computing, illustrative examples of Federated Computing use cases for banks, how Federated Computing beats alternative approaches, and an overview of the Rhino Federated Computing Platform (Rhino FCP).

What Is Federated Computing?

At its core, **Federated Computing** "brings code to the data," while still contributing to collective intelligence. Rather than centralizing raw datasets in a single location, FC executes the code directly on each participant's server, on-premises data center, or cloud environment.

Only encrypted or aggregated outputs are exchanged, so no sensitive information is ever shared in raw form.

This concept originated with **Federated Learning**, enabling distributed machine learning across multiple data sources. However, FC has evolved into a holistic approach applicable to advanced analytics, queries, and code execution. It harmonizes compliance needs—such as GDPR or local data residency laws—with modern business imperatives around AI-driven transformation. In short, **Federated Computing** provides financial institutions with **enterprise-grade** security and privacy, while unlocking cross-organization collaboration that was previously hindered by data-sharing constraints.

Applications of Federated Computing in Banking

There are many potential uses for Federated Computing across the banking industry. Some illustrative examples include:

Fraud Detection and Prevention

Fraud thrives in data silos, as orchestrated schemes frequently cut across multiple banks and platforms. With **Federated Computing**, each institution can build or enhance fraud detection models on its own data, periodically exchanging only the learned parameters or insights. By pooling these incremental learnings, banks detect sophisticated fraud patterns that remain invisible when operating alone—without ever revealing full customer records.

Anti-Money Laundering (AML) and Financial Crime

Money laundering rings exploit inter-institutional blind spots. Through **Federated Computing** analytics, AML teams securely identify hidden relationships, suspicious transaction flows, and complex layering tactics by analyzing data collectively. Each financial entity retains full custody of its customer information while sharing enough signals to pinpoint high-risk behaviors. This collaborative model not only improves threat intelligence but also addresses stringent privacy mandates.

Know Your Customer (KYC)

KYC processes are fundamental to preventing misuse of financial services. Yet verifying identities, monitoring accounts, and confirming risk profiles can be onerous when data is fragmented. **Federated Computing** helps unify multiple sources—internal or external—so each bank can securely query relevant risk or identity databases without disclosing proprietary or personal information. This speeds up KYC verification, ensuring compliance while easing friction for both financial institutions and customers.

Credit Scoring and Underwriting

When assessing creditworthiness, a broader dataset invariably leads to more nuanced risk scoring. Through **Federated Computing**, a consortium of banks, fintechs, or even utilities can collaborate to train models on diverse repayment patterns. Crucially, these organizations retain direct control over their data, with only aggregated risk indicators shared for collective model-building. The result is more precise lending decisions that minimize default rates while adhering to strict privacy standards.

Customer 360 and Personalized Banking

Many large banks seek a seamless “Customer 360” view, yet data silos—across regions, business units, or acquired entities—stifle these ambitions. **Federated Computing** breaks down silos without merging datasets into one repository. Each division applies the same code locally to produce standardized insights, which are then aggregated into a unified analytical layer. This approach preserves regulatory compliance (including data sovereignty rules) and powers hyper-personalized offerings that strengthen customer loyalty.

Consortium Benchmarking and Market Insights

Whether it’s benchmarking liquidity ratios or analyzing market sentiment, industry consortia often struggle to share data beyond basic surveys. With **Federated Computing**, members can run federated queries that compute collective averages, trends, or KPIs—exclusively from within each participant’s private environment. This yields genuine market intelligence that surpasses broad estimates, increasing the accuracy and reliability of shared insights.

Regulatory Stress Testing and Compliance Analytics

Banks regularly undergo stress testing by regulators to assess their capacity to withstand financial shocks. **Federated Computing** enables these analyses to occur within each bank’s infrastructure, sharing only aggregated indicators. Not only does this meet compliance requirements more effectively, but it also reduces the cost and complexity of securely transferring vast amounts of sensitive data into centralized testing environments.

Cross-Bank Collaboration for New Services

Innovation in finance increasingly happens through partnerships—whether with fintechs, insurers, or digital marketplaces. **Federated Computing** solutions create a layer of secure collaboration, allowing partners to exchange actionable insights without exchanging underlying records. This fosters co-created services (e.g., cross-promotions, joint risk frameworks, or unified loyalty programs) that enrich the customer experience across multiple providers.

Why Federated Computing vs. Other Data Collaboration Techniques?

Federated Computing is not the first attempt at enabling large-scale data collaboration in the banking sector. Below are some common approaches organizations have historically used, along with some of the risks they present:

- **Centralized Data Lakes or Warehouses**: Simple in concept but risky and expensive, as they demand ongoing synchronization of sensitive data that might breach privacy constraints.
- **Data Clean Rooms**: Rely on a third-party or platform environment, offering limited analytics scope and requiring trust in that environment's operators.
- **Secure Enclaves (TEE)**: Provide hardware-based isolation yet can struggle with large-scale or complex workloads.
- **Multi-Party Computation (MPC) and Homomorphic Encryption**: Cryptographically robust but often slow and resource-intensive for real-time banking scenarios.
- **Open-Source Federated Learning Frameworks**: Helpful for experimental distributed ML but typically lack enterprise-grade governance, orchestration, and auditing.

Federated Computing delivers a unique blend of security, scalability, and efficiency that elevates it above existing collaboration methods:

1. **Localized Data and Compliance:** Data is never moved outside a bank's control, making it simpler to adhere to GDPR, CCPA, and strict internal security policies. This stands in contrast to centralized warehouses or data clean rooms, which demand placing trust in external entities.
2. **Broadly Adaptable Workloads:** FC accommodates diverse use cases beyond machine learning, including analytics queries, KYC checks, and advanced AI workflows. Most cryptographic approaches or hardware enclaves struggle with such variety at scale.
3. **Cost-Efficient, Scalable Architecture:** Each participant uses existing infrastructure, minimizing large data migrations and storage overhead. Whether a consortium expands to add new banks or a single bank integrates multiple business units, FC scales gracefully without a complete architectural overhaul.
4. **Lower Legal Complexity:** Sharing insights—rather than raw data—streamlines partnership agreements. Since no institution relinquishes direct control over its sensitive datasets, negotiations often revolve around permissible computations instead of cross-licensing or liability for potential data breaches.

Federated Computing's decentralized, privacy-by-design model thus meets the evolving demands of modern finance: robust, real-time collaboration unencumbered by the typical pitfalls of data centralization.

The Rhino Federated Computing Platform

Rhino Federated Computing has built the world's leading data collaboration platform, Rhino FCP. Rhino cut our teeth in the world of healthcare and life sciences, where data privacy is absolutely paramount and data custodians fiercely protective of data subjects. Following Rhino co-founder & CEO Dr. Ittai Dayan's leadership of the landmark project on Federated Learning in healthcare, the EXAM Study, we set out to build an enterprise-ready FL solution, but quickly realized that collaborators needed more than 'just' Federated Learning, they needed Federated Computing.

While maintaining a focus on security & privacy, we have turned Rhino FCP into a totally secure, totally extensible collaboration sandbox - allowing collaborators to run code on one another's data while ensuring that security, privacy, and legal teams can be comfortable. Rhino FCP offers flexible architecture (multi-cloud and on-prem hardware), end-to-end data management workflows (multimodal data, schema definition, harmonization, and visualization), a privacy screen (custom differential privacy budget, custom k-anonymization values), and allows for the secure deployment of custom code & 3rd party applications via persistent data pipelines.

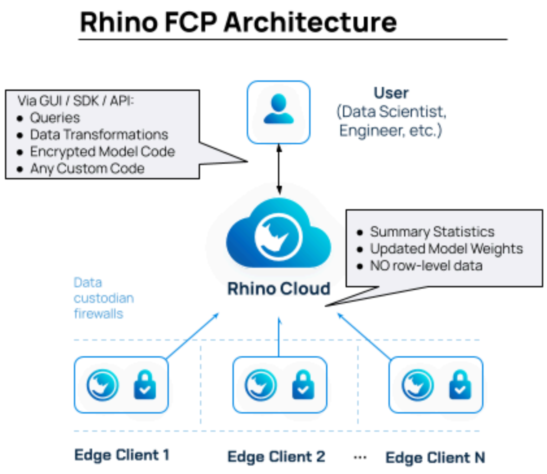

Figure: High-level Rhino FCP architecture diagram

Rhino FCP can also seamlessly integrate into banks' tech stacks, serving as middleware to any number of federated workloads run on distributed data.

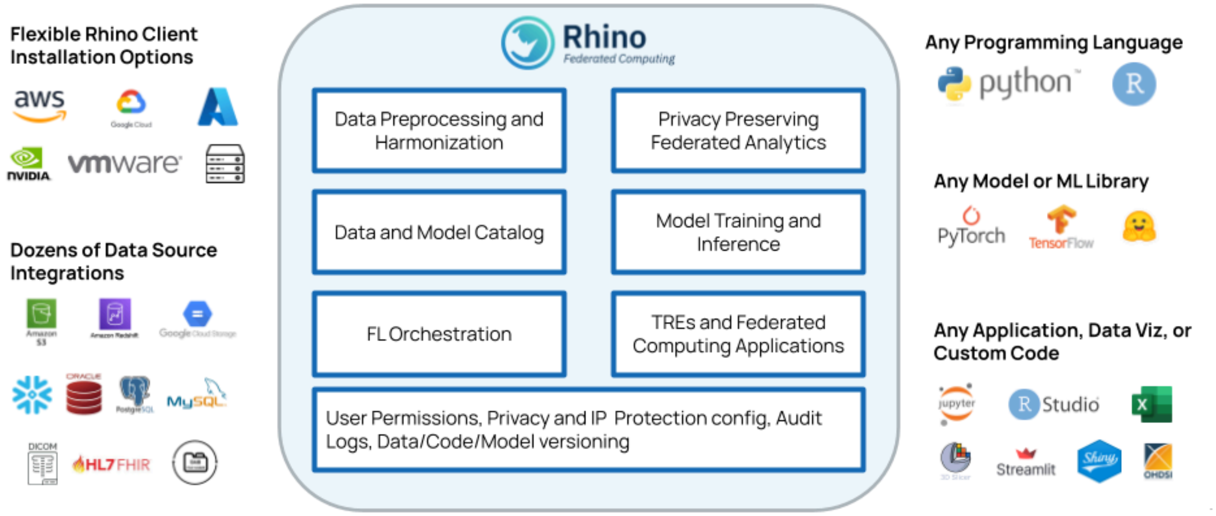
Flexible Rhino Client Installation Options
aws
Google Cloud
nvidia
vmware
Dozens of Data Source Integrations
Amazon S3
Amazon Redshift
Google Cloud Storage
ORACLE
PostgreSQL
MySQL
DICOM
HL7 FHIR
Rhino
Federated Computing
Data Preprocessing and Harmonization
Privacy Preserving Federated Analytics
Data and Model Catalog
Model Training and Inference
FL Orchestration
TREs and Federated Computing Applications
User Permissions, Privacy and IP Protection config, Audit Logs, Data/Code/Model versioning
Any Programming Language
python
R
Any Model or ML Library
PyTorch
TensorFlow
Any Application, Data Viz, or Custom Code
jupyter
RStudio
3D Slicer
Streamlit
Shiny
OHDSI

Figure: Illustrative High-Level Architecture of Rhino FCP Integrations

---

Conclusion

Federated Computing is changing how financial institutions collaborate and innovate. It balances stringent privacy and security obligations with the ambition to harness collective intelligence across departments, partner organizations, or entire consortia. From **fraud and AML** to **KYC**, **credit underwriting**, **personalized banking**, and **industry benchmarking**, FC unlocks secure data synergy that respects each participant's sovereignty over its own information.

Amid increasing customer expectations, regulatory scrutiny, and cybersecurity threats, **Federated Computing** stands as an enterprise-grade blueprint for data collaborations in finance. Unlike conventional solutions—be they centralized data lakes, hardware enclaves, or cryptographic approaches limited by performance—FC offers a truly privacy-first, future-proof foundation. As technologies like Rhino FCP continue to advance, banks can confidently adopt FC to deliver breakthrough insights, fortify their risk defenses, and build trust in a privacy-conscious era.

Reach out to us today to accelerate your bank's data & AI strategies using Federated Computing!