CONFIDENTIAL COMPUTING CONSORTIUM

# USE CASE:

## Confidential Computing in Financial Ecosystems

## Executive Summary

Banks and their partners, such as telecom operators, retailers, and payment processors, hold complementary data that, when combined, can significantly improve credit risk assessment, underwriting, and customer acquisition. The main barrier is not technical but legal: banking secrecy, GDPR, and commercial confidentiality make it impossible to centralize raw data in a single repository.

Confidential Computing enabled cross-industry collaboration on analytics and AI-driven risk models without exposing underlying customer records, improving risk assessment, customer acquisition, and marketing effectiveness while maintaining regulatory compliance and control over data.

## Challenge

A banking consortium operating in markets with fragmented data regulations wanted to enrich underwriting models with behavioral signals from external partners. Three constraints blocked standard approaches.

First, regulatory requirements, including GDPR rules for cross-border operations, prohibited the transfer of customer records without a clear legal basis. Second, partners were not willing to expose proprietary data to organizations that could also be their competitors. Third, traditional encryption protected data at rest and in transit but left it exposed in memory during computation.

Federated queries were also rejected: even without direct access to raw data, query patterns and indirect signals could reveal sensitive information.

## How Confidential Computing Was Implemented

The compute infrastructure was deployed on confidential computing-enabled CPU and GPU systems. Hardware-based attestation verified the integrity of the execution environment and code before data was released for processing. Confidential GPU accelerators were used for ML/AI workloads.

Participating organizations retained control of their data and determined which workloads were permitted to access it. Raw records remained inaccessible to the cloud provider, platform operators, and system administrators. Only pre-approved outputs were allowed to leave the environment: aggregated risk scores, final model weights, and match counters.

Data lab environments allowed data science teams to develop and validate models on combined datasets interactively, without accessing raw data from other participants.

## Key Takeaway

**Confidential Computing protects data not only at rest and in transit, but also during active processing on both CPU and GPU. This makes inter-institutional collaboration legally and commercially viable. Hardware attestation creates a verifiable foundation of trust that neither contracts nor organizational controls can replace.**