



**Rhino**  
Federated Computing

Google Cloud

WHITE PAPER

# Beyond Data Silos:

## *How Federated Computing Unlocks AI Without Compromising Privacy*

AUTHORS:

Chris Laws, Rhino Federated Computing  
Oran Barkai, Rhino Federated Computing  
Maxim Afanasyev, Google Cloud



# Executive Summary

The modern enterprise faces an impossible choice: feed AI's insatiable appetite for data, or protect that data behind increasingly stringent privacy regulations. This isn't a philosophical debate—it's costing organizations billions in missed opportunities, compliance penalties, and competitive disadvantage.

## The Conflict

AI models require massive, diverse datasets to achieve meaningful accuracy. Yet GDPR, China's Personal Information Protection Law (PIPL), Japan's Act on Protection of Personal Information (APPI), and a rapidly expanding web of national AI regulations make traditional data pooling legally treacherous and operationally impossible.

## The Solution

Federated computing represents a fundamental architectural shift—rather than moving sensitive data to where it's needed, only the insights travel. This architecture solves three critical problems:

-  1. **Cross-Organization Collaboration:** Organizations can collaborate on endless use cases—from cancer research and drug discovery to fraud prevention and supply chain resilience—without ever exposing, moving, or risking their most sensitive data to partners.
-  2. **Internal Data Silos:** Even within a single organization, data often lives in isolated silos across departments, legacy systems, geographical regions, and regulatory zones. Federated RAG enables AI to train across these internal boundaries without consolidating sensitive data into central repositories or exposing proprietary IP and PII to unauthorized users.
-  3. **Code-Data Separation:** When building AI tools that require access to highly sensitive production data, federated architecture creates a secure boundary between the developers who write the code and production data that the code processes.

# The Value

Early adopters are already seeing impressive results:

- Eli Lilly’s TuneLab platform fine-tunes proprietary drug discovery models on biotech partners’ data, with partners gaining access to increasingly powerful models without surrendering their underlying molecular libraries.
- Swift and leading global banks are accessing collective fraud intelligence that no single institution could build alone. Utilities companies are sharing cyber threat intelligence across decentralized networks without exposing sensitive infrastructure data.
- Internally, organizations are deploying federated AI assistants that draw from HR, legal, financial, and customer data simultaneously—without ever creating a centralized copy that administrators or developers can access.

The question is no longer whether to choose between data privacy and AI progress. The infrastructure now exists to achieve both simultaneously. The only remaining question is: how quickly will your organization adapt?



## I. The Problem: *The Modern Data Paradox*


### Drowning in Data, Starving for Insights

Organizations are capturing more data than ever before—transaction logs, patient records, customer interactions, and sensor readings from billions of IoT devices. Yet, paradoxically, this data wealth has become a strategic liability rather than an asset.

Traditional centralized data lakes, once hailed as the solution for enterprise analytics, are now more accurately described as “data graveyards”—vast repositories where sensitive information accumulates, ages, and becomes an ever-expanding attack surface. A single breach can expose millions of records. A single regulatory violation can trigger nine-figure penalties.

# The Silo Problem

The core issue is almost deceptively simple—the data that would improve your model is almost never in one place. And, it often can't legally or competitively be put there.




The pattern repeats across industries. [Money laundering schemes move an estimated \\$2 trillion annually](#) across global financial systems, yet [less than 1% of these illicit funds are ever seized](#). Criminals deliberately layer transactions across multiple institutions, knowing no single bank sees the full picture. In life sciences, rare disease researchers sit on potentially life-saving data across dozens of hospitals globally, but patient privacy laws make combining it nearly impossible. The data exists. A safe way to connect it doesn't.

Internally, the problem is just as severe.

Large enterprises typically face four distinct varieties of this problem:

1. **Departmental Silos:** Finance, HR, legal, sales, and operations each guard their own data—different systems, different access controls, different teams.
2. **Geographic Silos:** For example, data generated in Europe often can't move to US servers and vice versa—regulatory boundaries turn physical distance into a data wall.
3. **Legacy System Fragmentation:** Decades of mergers and acquisitions leave organizations with dozens of incompatible databases that are too expensive or risky to consolidate.
4. **Regulatory Boundaries:** Regulatory and compliance requirements frequently demand that different business units operate with strict information barriers between them—making cross-functional AI a legal question before it's ever a technical one.

When a business wants an AI assistant that can answer “What was our total customer acquisition cost across all regions last quarter?” The answer requires crossing six systems across three continents. Centralize everything and you create a security liability. Stay siloed and your AI only sees part of the picture. Give developers unrestricted access and you have a compliance problem. None of the options are good ones.



## The Regulatory Wall

The legal landscape has fundamentally shifted in ways that make traditional data pooling not just risky, but potentially illegal.

GDPR’s “right to be forgotten” alone makes centralized AI training datasets legally precarious. Once an individual’s data has been pooled and used to train a model, you cannot guarantee its deletion. Similar requirements are proliferating across every major market. In 2025 alone, [GDPR penalties exceeded €1.2 billion](#). The cost of non-compliance grows with every record stored and every partnership contemplated.

## The Cost of Operating Alone

Working in isolation carries its own harmful price across every industry.

In financial services, transaction monitoring systems generate [false positive rates exceeding 90% of flagged transactions](#) — not because institutions lack diligence, but because no single institution sees the full picture. Criminals deliberately exploit that isolation.

In life sciences, pharmaceutical companies spend an average of [\\$1-2 billion and 10-15 years bringing a single drug to market](#). A leading AI model provider has developed breakthrough molecular analysis models that could compress that timeline, but unlocking their full potential requires fine-tuning on proprietary data held by pharmaceutical partners who can’t share it externally. Federated computing makes that collaboration possible without either party ever exposing their underlying research.

In agriculture, designing next-generation seeds requires genomic and agronomic data that is often

locked away in university research institutions and private partner networks. The challenge is figuring out how to incentivize data owners to participate when sharing means exposing years of proprietary research. A federated approach removes that barrier entirely—each partner contributes local training data without ever exposing their underlying IP, and receives better predictive models in return.

Internally, [data scientists spend roughly 80% of their time finding, cleaning, and organizing data](#) before any analysis begins. That's not a talent problem—it's a structural one.

The solution, therefore, must be architectural.


## II. The Technology: *How Federated Computing Works*

### Centralized vs. Federated: A Fundamental Paradigm Shift

Traditional AI development follows a centralized model: gather data from across the organization or industry into a central repository, clean and standardize it, train the model, and deploy it.

This approach maximizes model performance and convenience but creates serious privacy, security, and regulatory exposure. Every piece of sensitive data must be copied, transmitted, stored, and made accessible to third parties or internal teams who may not need access to all of it.

Federated computing inverts this entirely. Rather than bringing data to where computation happens, the computation runs where the data lives. **Raw data never moves.** What travels instead are the insights—mathematical representations of what a model learned from local data, aggregated across participants into shared intelligence that benefits everyone.



This is reinforced by a suite of Privacy-Enhancing Technologies (PETs) that can be layered depending on the sensitivity of the use case and the regulatory environment. These include:

- 1 Differential Privacy**  
Adds calibrated noise to prevent individual data points from being reconstructed
- 2 Homomorphic Encryption**  
Allows computation on encrypted data without ever decrypting it
- 3 Secure multi-party computation**  
Enables multiple parties to jointly derive insights without revealing their inputs to each other
- 4 Trusted Execution Environments (TEEs)**  
Use hardware-level isolation to protect data even from the infrastructure it runs on.

Together, these ensure that even the insights that do travel cannot be reverse-engineered into the underlying data.

## Three Problems Federated Computing Solves

Federated computing is a family of approaches you can apply depending on which problem your organization is facing right now.



### Cross-Organization collaboration

Multiple organizations train a shared AI model without exposing their data to each other. For example, a group of banks could each train a fraud detection model locally, sharing only model improvements rather than line-item transaction data. The result is collectively smarter intelligence that no single institution could have built alone.



### Breaking Internal Silos

A single organization deploys federated approaches across departments, geographies, and legacy systems without centralizing data—whether that's federated analytics surfacing insights across business units, or a RAG-powered AI assistant drawing from finance, HR, legal, and customer systems simultaneously.



### Code-Data Separation

Developers build and test using synthetic data. When deployed to production, the system operates on real sensitive data inside a secure environment that developers never directly access. Every action is logged and auditable.

Each of these patterns shares the same underlying principle, and the same mechanism that makes it possible.

# The “Zero-Movement” Workflow: A Step-by-Step Journey


Here’s how a federated learning, one of several approaches federated computing supports, works in practice, from initiation to deployment. The process is powered by two components: a central orchestration layer that coordinates activity without ever touching sensitive data, and lightweight local clients installed at each participating site behind their own firewall.

1. **Distribute:** The orchestration layer sends a starting model to the local client at each participating site—whether that’s a partner organization, a regional subsidiary, or a department running on its own infrastructure.
2. **Train locally:** Each local client trains the model using its site’s data, entirely within that site’s secure environment. No data crosses the firewall.
3. **Return updates:** When training completes, only the model updates—mathematical representations of what was learned, not the underlying data—are returned through encrypted channels to the central orchestration layer.
4. **Aggregate:** The orchestration layer combines updates from all participants into an improved global model.
5. **Redistribute and repeat:** The improved model is sent back to all local clients, and the cycle continues. Each round produces a collectively smarter system, with no participant ever seeing another’s data.

## The intelligence compounds while the data stays put.

The same principle extends beyond model training—federated analytics computes statistics and insights across distributed datasets without centralizing records, while federated inference runs predictions locally against real-time data.

Federated computing doesn’t ask organizations to choose between collaboration and control—it makes that tradeoff obsolete. The data stays where it belongs, and the intelligence flows freely. And, organizations are already taking advantage of it.



### III. Real-World Successes: *Federated Computing in Production*

#### Federated is the Best Path Forward, and It's Already Proven at Scale

For most organizations handling sensitive data, centralization was never really an option. Regulatory constraints, competitive concerns, and data sovereignty laws make traditional data pooling legally untenable before the first technical question is even asked.

Federated computing doesn't compete with centralization. It enables collaboration that simply wouldn't exist otherwise.

**What's emerged from early production deployments is striking: models trained across diverse, distributed datasets frequently outperform what any single institution could build alone.** These models benefit from more data, more sources, and more real-world complexity—without any of it moving.

The results across industries tell the same story.



Life Sciences



Financial Services



Automotive



Beyond These Verticals



## Life Sciences: Accelerating Drug Discovery with Eli Lilly


Pharmaceutical companies collectively hold vast libraries of molecular data and failed compound histories that could accelerate discovery industry-wide, but competitive pressures keep that knowledge locked away. **Eli Lilly's TuneLab platform**, launched in September 2025, was built to change that.

**The Federated Solution:** Built on Rhino's federated platform, TuneLab allows Lilly to refine its proprietary drug discovery models using biotech partners' molecular data. Partners run inference locally against 20+ pre-trained Lilly models spanning toxicity prediction, target identification, and compound optimization, then share back model improvements. All computations occur on the Rhino Federated Computing Platform—Lilly and their partners' data never moves.

### The Impact

- **Accelerated candidate selection:** Biotech can screen thousands of compounds against Lilly's validated models in hours rather than months
- **Reduced animal testing:** Better in-silico toxicity prediction means fewer compounds need to advance to animal studies
- **Democratized access:** Small biotech gains access to pharma-grade AI infrastructure without requiring data-sharing agreements or massive compute investments
- **IP protection:** Molecular structures and proprietary libraries remain completely confidential

**TuneLab represents a new model for pharmaceutical R&D collaboration—one where competitors can cooperate on foundational science while each protecting their most valuable intellectual property.**



Lilly TuneLab



## Financial Services: Smarter Fraud Detection with JPMorgan Kinexys & Bank of New York

Cross-border payments move trillions of dollars daily, and fraud patterns identified might look different in every jurisdiction. A bank focused on Asia may never encounter the anomaly types common in North America and vice versa. Pooling raw transaction data across institutions to fix this is nearly impossible given legal, regulatory, and competitive constraints. JPMorgan Kinexys and the Bank of New York (BNY)'s challenge was clear: how do you build a globally-aware fraud detection model without ever moving sensitive transaction data?

**The Federated Solution:** Through [Project AIKYA](#), named from Sanskrit for “unity,” JPMorgan Kinexys and BNY deployed a federated learning architecture where each bank trained a local model on its own transaction data. Only model weights, not transaction records, were shared with a central aggregation server. The global model was then redistributed to each bank, which retrained locally to preserve its own expertise while absorbing the other's learned patterns.

### The Impact

- **Expanded anomaly coverage:** After four rounds of aggregation, each bank's model detected both anomaly types as accurately as if it had been trained on a combined dataset—something neither could achieve alone.
- **Privacy preserved by design:** No raw transaction data or customer records left either institution at any point. Only compressed model weight updates traversed the network.
- **Performance that compounds:** The federated model matched or outperformed individual local models across every standard metric—precision, recall, F1-score, and accuracy.
- **Built to scale:** The architecture supports any number of participant institutions. Adding a new bank requires only joining the network and contributing local model weights.

Project AIKYA proves that competing institutions can collaborate to build stronger AI without compromising data sovereignty. Because the approach is model-agnostic, the same architecture extends naturally to regulatory reporting, credit risk, AML compliance, and beyond.



## Automotive: Smarter EVs with Toyota

Range anxiety—the fear that an electric vehicle’s battery won’t last the journey—remains one of the biggest barriers to EV adoption. Better range prediction models could help, but training them traditionally would require sending continuous streams of vehicle data to the cloud, raising privacy concerns and demanding costly infrastructure. Toyota’s challenge was clear: how do you build a smarter, more accurate range prediction model across an entire fleet without centralizing sensitive vehicle data?

**The Federated Solution:** [Toyota deployed a federated learning architecture](#)—built on NVIDIA FLARE, the same open-source framework powering Rhino’s platform—where each vehicle trains a local AI model on its own driving and battery data. Only model updates, not raw vehicle data, are shared with a central server for aggregation. A global model is then redistributed to each vehicle, which fine-tunes it on its own data for personalized predictions. Raw data never moves.

### The Impact

- **Improved range prediction accuracy:** Collective learning across the fleet produced more accurate, personalized battery estimates than any single vehicle could achieve training alone
- **Reduced infrastructure costs:** The federated approach required approximately 10x less network and computing resources than a traditional cloud-based approach
- **Privacy preserved by design:** Sensitive vehicle and driving data never leaves the car
- **Edge-ready scalability:** The architecture deploys directly to onboard vehicle hardware and scales across any fleet size without proportional infrastructure investment

Toyota’s demonstration signals something bigger than better battery estimates. It shows that Original Equipment Manufacturers (OEMs) across the automotive industry can build smarter, more accurate models across their entire fleets without moving sensitive vehicle data, and that the same federated principles apply anywhere operational data is generated at scale and privacy matters.

# IV. Where Federated Learning Creates Value Across Industries: *Additional Use Cases*

Federated computing’s applicability extends across nearly every industry. Here’s how it’s transforming operations across sectors:

| Industry           | The Challenge                                                                                                                                                                                      | The Federated Solution                                                                                                                                                                               | Value Delivered                                                                                                                                                                          |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Insurance          | Emerging risk categories (such as cyber, climate, and pandemic) lack historical claims data. No single insurer has sufficient data for accurate pricing.                                           | Insurers securely pool claims patterns to build shared risk models without exposing competitive pricing or customer information.                                                                     | Insurers can price emerging risk categories accurately, bring products to market faster, and significantly reduce adverse selection.                                                     |
| Manufacturing /IoT | Predictive maintenance requires failure data from multiple factories, and competitors won’t share production efficiency data.                                                                      | Factories train failure prediction models on pooled sensor data while keeping production speeds and operational parameters confidential.                                                             | Manufacturers see measurable improvements in equipment uptime, reduced maintenance costs, faster failure detection across supply chains.                                                 |
| Energy & Utilities | Cyber threat data across utilities is highly fragmented and cannot be centralized due to regulatory and privacy constraints, limiting effective threat detection                                   | Multiple utilities to collaboratively train models on decentralized data, sharing threat intelligence without exposing sensitive information.                                                        | Utilities companies benefit from reduced false positives and improved threat detection performance across participating environments.                                                    |
| Chemicals          | Performance and IP-sensitive data from catalysts used across refineries is siloed and cannot be shared centrally, limiting the ability of R&D teams to optimize catalyst performance across sites. | Refineries train models on local time-series operational data while sharing only model updates centrally, keeping sensitive production parameters and catalyst formulations completely confidential. | Refineries see continuous improvement of catalyst performance through cross-site learning, supporting improved yield while maintaining strict data security and operational constraints. |

The pattern is consistent—wherever valuable insights require data from multiple parties, federated computing enables collaboration that regulation and competition have historically made impossible.

And, the frontier is expanding. Agentic AI and Retrieval-Augmented Generation (RAG) are two of the most talked-about developments in enterprise AI right now. Both face the same fundamental barrier as every use case requires access to sensitive, private data to deliver real value. The next section explores what these technologies actually mean in a federated context and why federated computing may be the only architecture that lets organizations deploy them safely at scale.

## V. Agentic AI and Federated RAG: *From Buzz to Production*

The next frontier of AI isn't just better models sitting behind a chat interface—it's autonomous agents that can reason, plan, and take actions on behalf of humans across complex workflows. But, the AI assistants that will actually deliver business value aren't the ones trained on public data. They're the ones with access to your proprietary knowledge—your claims histories, your contracts, your customer records, and your operational data.

That's exactly the problem. The AI systems most worth building require access to the most sensitive data your organization holds. Traditional deployment approaches force an impossible choice: give agents broad access to production data and accept the security and compliance exposure, limit them to sanitized or public data and accept a system that under-performs, or pre-load sensitive data into training pipelines and create permanent privacy liabilities. None of these options is a good one.

Federated computing offers a different path entirely. Agents operate locally on sensitive data without developers ever seeing it, and knowledge can be retrieved across distributed systems without ever being centralized.

### Federated Learning for Agentic AI: Separating Development from Data

Federated computing enables a clean separation between the teams that build AI agents and the sensitive data those agents ultimately access. Developers work in synthetic environments while agents run against real data in production, with performance signals flowing back through privacy-preserving channels to continuously improve the system.

This results in a model where:

- Production data never leaves its secure environment
- Developers build faster without waiting for data access approvals
- Every agent action is logged and auditable at the local level
- Agents continuously improve through federated learning



Federated Agentic AI

## In Practice: Claims Processing

Consider a large U.S. property and casualty insurer that needs to deploy AI agents to extract and process data from claims PDFs—data sitting behind strict access controls that developers can't touch.

Using Rhino's Federated Computing Platform, the team can build and test agent logic centrally on synthetic data, then deploy containerized agents directly into the production environment. Agents execute locally against real claims data without exposing it externally. Operations teams monitor agent performance from a central dashboard without raw data ever leaving the secure environment, and subject matter experts annotate samples locally to feed quality improvements back through the system.

The result is a development model that moves faster, reduces data exposure, and produces AI that gets smarter over time—without the data access risks that have can make agentic AI a compliance nightmare in regulated industries.

## Federated RAG: Private Knowledge Without Central Repositories

Retrieval-Augmented Generation (RAG) has emerged as the dominant pattern for building AI systems that answer questions using private organizational knowledge. Think of it as giving an AI assistant access to your company's internal library. Rather than guessing, it looks at relevant assets in real time before responding.

The problem is that traditional RAG requires copying all those documents into a centralized database first, creating a new security liability, navigating access control complexity across departments and potentially violating data residency requirements before a single query runs.

Federated RAG architecture keeps data exactly where it is. Queries travel simultaneously to multiple distributed document stores, each of which searches locally, enforces its own access controls, and returns only what that user is permitted to see. An LLM synthesizes the results into a coherent answer. No documents move, no centralized copy is created, and every query is logged at the source.

This approach also solves the developer access problem entirely. The team building the RAG system never sees production documents. They deploy orchestration logic, and the data responds to it at runtime.

## In Practice: Enterprise-Wide Search Across Regulatory Boundaries

Consider a global chemicals company with dozens of research repositories spanning multiple regulatory jurisdictions—each containing files that cannot leave the jurisdiction where they reside. Rather than centralizing data, the company deployed a local search service within each silo using Rhino's Federated Computing Platform. When an employee submits a query, it runs simultaneously across every repository, with each site returning only file names and relevance scores. Results are aggregated and ranked centrally, giving employees a unified search experience without a single document crossing a regulatory boundary.

The applications extend beyond claims processing and enterprise search. Any organization deploying AI across distributed data—whether that's a retailer querying inventory and customer data across regions, a hospital network building clinical decision support across campuses, or a manufacturer monitoring quality across production sites—faces the same fundamental problem. Federated computing solves it once, architecturally, across every use case that follows.

## VI. Data Collaboration as a *Strategic Advantage*

The organizations that will lead in AI aren't waiting for perfect data infrastructure. They're building competitive moats right now by accessing intelligence their isolated competitors simply cannot see.

### Collaboration as a Force Multiplier

Consider what becomes possible when mid-market institutions stop competing alone. A network of regional banks collaborating with federated learning collectively covers more geographic ground, more transaction types, and more emerging fraud patterns than any single institution—without sharing a single customer record. The fraud intelligence that once required the scale of a global bank can now be built by a consortium of regional ones.

The data advantage that once required scale can now be built through collaboration. Organizations that move first don't just solve a compliance problem—they build AI that competitors still working in isolation simply cannot replicate.

### Where the Advantage Compounds

Agentic AI and federated RAG are where this pulls away fastest, and the gap widens with every passing quarter. Here are a few examples:

- Agents trained on real, distributed production data continuously improve, while competitors training on sanitized snapshots are working with a fundamentally inferior signal.
- Federated RAG systems surface organizational knowledge instantly across every silo, without the security exposure of centralized vector databases.
- Every deployment generates federated learning signals that make the next deployment smarter.

The organizations leveraging a federated approach now will only see the advantage grow over time.

### Governance as a Feature, Not a Constraint

For data and governance leaders, federated computing offers something that most AI deployments don't: built-in compliance. Data sovereignty requirements, access controls, and audit trails aren't bolted on after the fact—they're architectural properties of how the system operates. The controls your organization already has in place don't need to be dismantled to make room for AI. They become the environment it runs in.

That's a fundamentally different conversation than the one most governance teams are used to having.

## VII. Conclusion: From *Silos to Swarms*

If you're a data or AI executive today, you're likely wrestling with some version of these three questions:

- “How can we collaborate with partners to build AI that none of us could build alone?”
- “How do we unlock the AI value trapped in our internal data silos?”
- “How do we let our developers build powerful AI without giving them the keys to our most sensitive data?”

**Federated computing answers all three. JPMorgan's Project AIKYA already outperforms single-institution fraud models by 40-50%.**

**Eli Lilly's TuneLab is accelerating drug discovery without either party seeing the other's molecular data.**


**The Cancer AI Alliance is advancing clinical innovation across the world's top research institutions without moving a single patient record.**

**These production deployments prove that federated computing is a major roadblock-remover across industries and use cases.**

The organizations that have adopted federated computing first are not only solving today's compliance problems—they're also building the data collaboration infrastructure that agentic AI will run on tomorrow.

The question is no longer whether your organization can afford to adopt federated computing. It's whether you can afford to be the last one that hasn't.

Ready to see how federated computing can accelerate your organization? [Reach out to the Rhino team to request a demo today.](#)



# About the Authors

**Chris Laws**, Chief Commercial Officer at Rhino Federated Computing; **Oran Barkai**, Vice President of Business Development at Rhino Federated Computing; and **Maxim Afanasyev**, PhD, Financial Services Industry Market Lead, Asia Pacific and Japan at Google Cloud.

**Rhino Federated Computing** solves one of the biggest challenges in AI: seamlessly connecting siloed data through federated computing. The Rhino Federated Computing Platform (Rhino FCP) serves as the ‘data collaboration tech stack’, extending from providing computing resources to data preparation & discoverability, to model development & monitoring — all in a secure, privacy-preserving environment. To do this, Rhino FCP offers flexible architecture (multi-cloud and on-prem hardware), end-to-end data management workflows (multimodal data, schema definition, harmonization, and visualization), privacy enhancing technologies (e.g., differential privacy), and allows for the secure deployment of custom code & 3rd party applications via persistent data pipelines.

Rhino is trusted by 60+ leading organizations worldwide — including 14 of 20 of Newsweek’s ‘Best Smart Hospitals’ and top 20 global biopharma companies — and is leveraging this foundation for financial services, ecommerce, and beyond.

**Maxim Afanasyev** is a member of Global Strategic Industries Team at Google Cloud and is responsible for Financial Services Industry in Asia Pacific and Japan. In his role, Maxim co-owns Google Cloud strategy and acts as a bridge between Google engineering innovations and real world. Maxim also advises C-levels of major companies in JAPAC on industry trends, transformations and emerging technologies powered by Applied AI.

Before joining Google Cloud, Maxim has been leading large global AI engineering and transformation teams of 200+ people and US\$100mn budgets to build and deploy innovative AI solutions, deliver AI adoption strategies, design human-AI interactions and unlock business value with new technologies across industries, both as an employee and as a consultant. Most recently as Managing Director in charge of AI products and transformation across 50+ markets at Standard Chartered Bank. Before relocating to Singapore, Maxim worked at McKinsey and Company where he delivered data strategies and AI use cases for clients in Europe and at Royal Bank of Scotland in London, UK as exotic derivatives trader.

Maxim holds PhD in Operations, Information and Technology from Stanford Graduate School of Business. Started career in AI Research Lab focused on neural networks and evolutionary algorithms.

## Citations & Further Reading

1. Swift & Google Cloud Partnership Announcement (December 2024)
2. “Fighting Financial Crime Together: Federated Computing for Cross-Bank AML Compliance”—Rhino Health
3. “Federated Computing Provides Decentralized Intelligence for Drug Development At-Scale”—Rhino Health
4. NIST Special Publication 800-207: Zero Trust Architecture
5. JPMorgan Chase Project AIKYA technical disclosures
6. European Banking Authority guidelines on collaborative AML
7. FDA guidance on AI/ML in drug development

