INSIGHT REPORT

Rhino Federate AI

Executive Forum 2026

What enterprise leaders are learning as federated AI moves into production

Hosted by:

Rhino
Federated Computing

Advancing the Next Era of Federated Intelligence

On June 23, 2026, Rhino Federated Computing convened a select group of senior leaders for the Federate AI Executive Forum in Boston, a working session built around the practical realities of deploying federated AI at enterprise scale. The conversations covered federated learning architectures, data harmonization and FHIR compliance, confidential computing, and the emerging intersection of agentic AI with federated infrastructure. Executives and practitioners from across cloud infrastructure, life sciences, healthcare, financial services, and agriculture joined to share what's working, what's not, and where the field is going next. This report captures what emerged from those conversations.

Key Takeaways

1. Federated AI has crossed from research curiosity to production infrastructure. Cloud providers, enterprises, and governments are now building for it deliberately because the data they need to learn from cannot be moved.

2. Networks compound in value. The TuneLab program at Eli Lilly, now with more than 100 partners across 11 countries, demonstrates what federated AI looks like at large scale.

3. Data harmonization and federation are most powerful in combination. The leading practitioners are using harmonization as a foundation that makes federated learning faster, more reliable, and more scalable.

4. Federated AI has reached escape velocity. After starting in life sciences and healthcare, financial services, agriculture, and industrial supply chains are taking off.

5. Agentic AI and federated AI are converging. The next layer of value is not just learning across siloed data — it is deploying intelligent agents that can reason over it in real time, without ever moving it.

1

Federated AI Has Crossed Into Production Infrastructure

Federated AI found early traction among academic consortia and grant-funded experiments, but today, that's no longer the case.

A consistent message from cloud infrastructure leaders at the forum was that enterprise customers are now intentionally building federated computing into production architectures, and that the major cloud providers are designing to meet them there.

Ittai Dayan, CEO of Rhino Federated Computing, opened the forum by reframing the market picture. What had previously looked like modest, linear growth in federated learning has been revised upward. The sector now represents over a billion dollars in annual value and is growing at approximately 30 percent per year. The revision, he argued, isn't just hype. There's real regulatory pressure on data movement, data-sovereign environments are proliferating, and there's a growing recognition that centralized AI training is untenable for organizations dealing with sensitive data at scale.

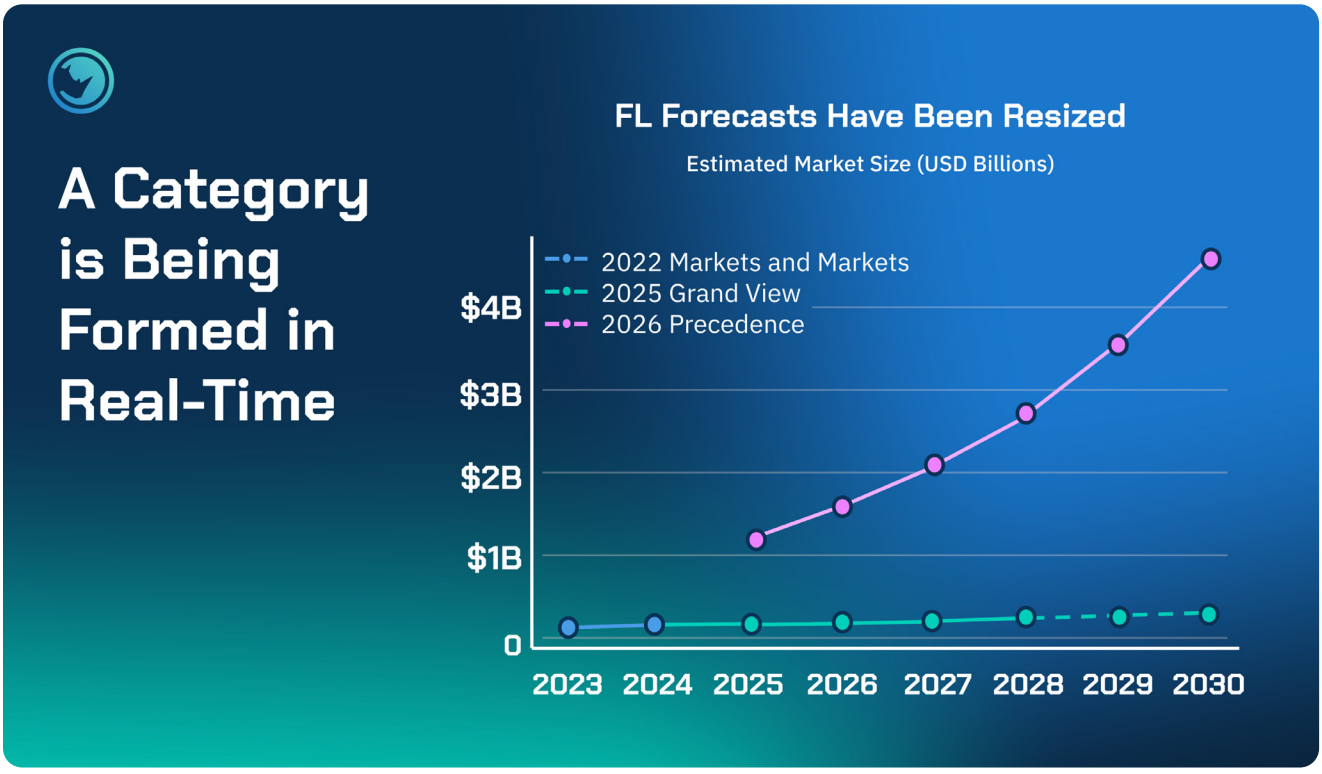

Ujjwal Ratan, Head of AI/ML for Healthcare and Life Sciences at AWS, echoed that shift from the infrastructure side, pointing to federating learning deployments that have moved beyond pilots at institutions like TuneLab — names that, as he noted, attendees would be hearing across the day.

Esteban Rubens, Healthcare Field CTO at Oracle Cloud Infrastructure (OCI), offered perhaps the sharpest framing of the underlying tension:

> **"It's been over 40 years since Stewart Brand said, 'Information wants to be free. [It] also wants to be expensive.' *Federated learning is at the forefront of that tension.*"**
>
> — Esteban Rubens, Healthcare Field CTO, Oracle Cloud Infrastructure

Rubens described OCI's architectural response to the tension he referenced. Cloud interconnects with AWS, GCP, and Azure by design, and a zero-trust packet routing model that moves the security boundary from network location to the individual workload — a natural complement to federated computing's requirement that sensitive data never leave its origin environment.

The point, shared across all three cloud representatives, was the same: **the infrastructure layer, including privacy-enhancing technologies like confidential computing and secure enclaves, is being built deliberately for this paradigm, not retrofitted**.

2

The More Institutions that Join, the More Valuable the Network Becomes

Federated networks get more powerful as they grow larger. When more institutions contribute data to a shared model, the model sees a broader, more diverse picture than any single institution could produce alone.

TuneLab, the federated AI program built by Eli Lilly and deployed through the Rhino Federated Computing Platform (FCP), offered the forum's clearest demonstration of that principle at scale.

Elke Nelson-Nichols, VP of Life Sciences Strategy at Rhino Federated Computing, moderated a panel that put the numbers on the table. Petar Stojanov, Senior Advisor and Machine Learning Scientist at Eli Lilly, reported that TuneLab has now surpassed 100 institutional partners signed onto the platform across multiple global regions. The consortium spans diverse therapeutic areas, biological modalities, and research approaches. That breadth is the point: the platform's value is not simply its size. It is the range of signals that become accessible when data that would otherwise never leave its originating institution can contribute to shared learning.

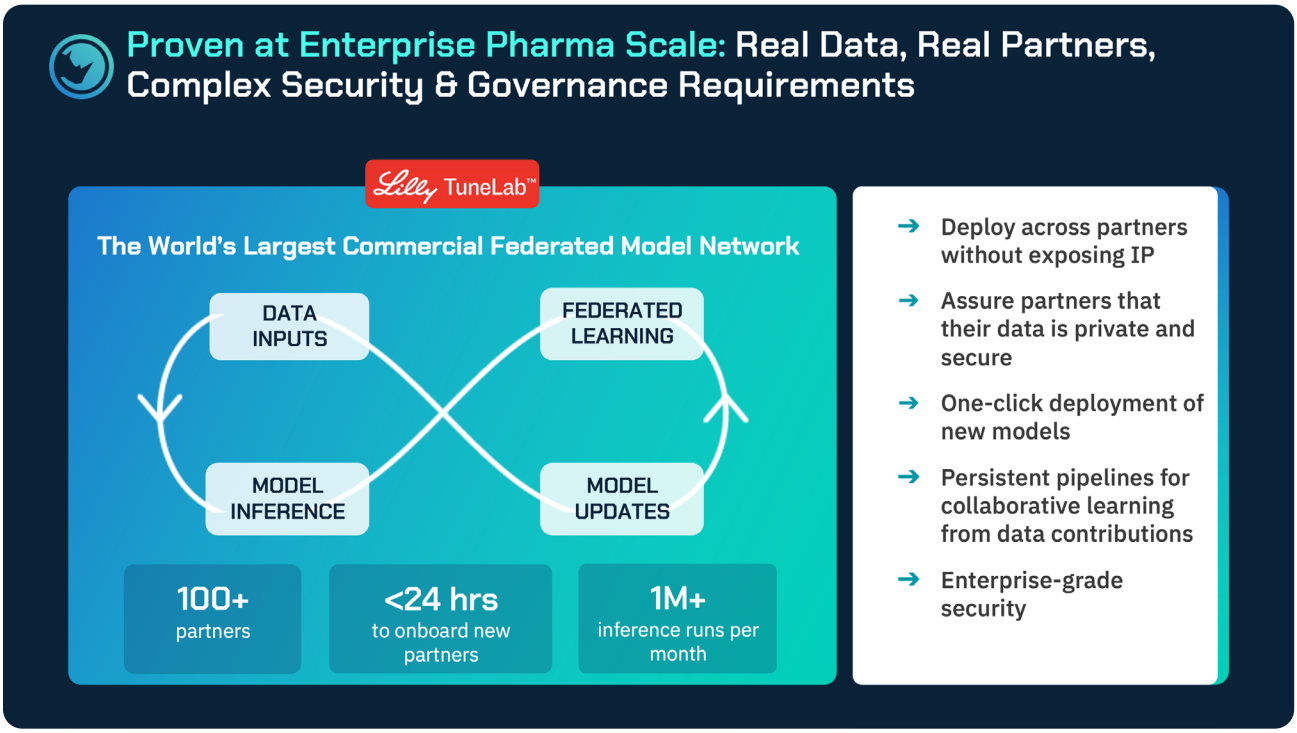

Lilly Tunelab Project

Ghulam Rasool, Associate Member in the Department of Machine Learning at Moffitt Cancer Center, brought the patient-side urgency into the room directly:

> **"As a cancer center whose primary mission is to cure cancer, we are driven by how we can improve patient outcomes. So the projects that we are focused on are *using AI to develop translational use cases.*"**
>
> — Ghulam Rasool, Associate Member, Department of Machine Learning, Moffitt Cancer Center

Translational use cases, in this context, refer to research that moves directly from computational insight to clinical impact. This is AI development explicitly tied to what improves survival rates, quality of life post-treatment, and diagnostic precision rather than model benchmarking for its own sake. Rasool described Moffitt's enterprise-level data management platform, which aggregates temporal, clinical, imaging, and pathology data for nearly one million cancer patients. This is data that, by its nature, cannot leave the institution. **Federated learning is the only architecture that makes cross-institutional collaboration legally and operationally viable for Moffitt, while keeping the models pointed at the outcomes that matter.**

3

Data Harmonization Makes Federation More Powerful

The forum surfaced a practical question that the most sophisticated federated AI practitioners are actively working through: how much data harmonization do you need before federation delivers value?

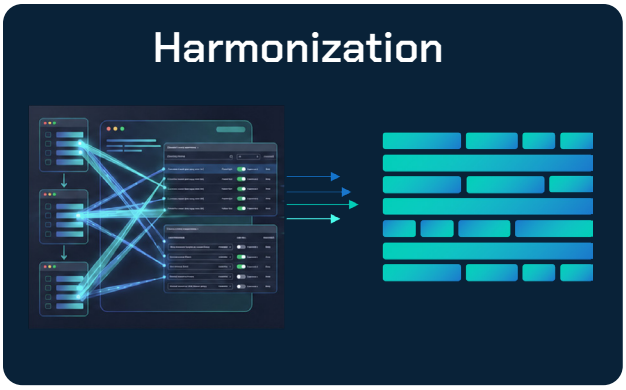

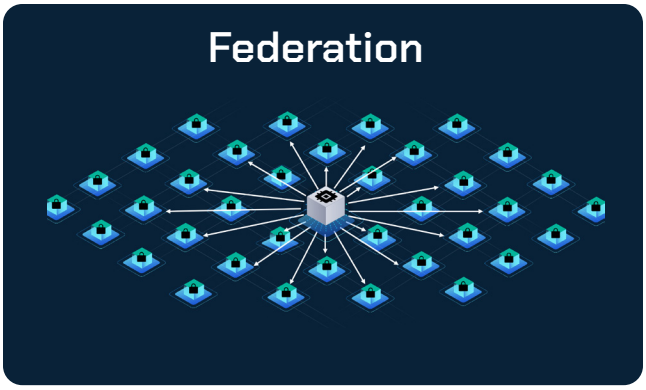

Nisim Rahman, CTO and Director of Operations, Data and Analytics Division at Sheba Medical Center, made the case for harmonization as a foundation. Sheba is Israel's largest hospital and ranks among the top ten globally, with data spanning more than two million patients and over 200 active research projects running concurrently. Operating in a regulatory environment that requires FHIR-IL compliance (the interoperability standard required for health data exchange), and collaborating with networks including Mayo Clinic, Paradigm, and a newly formed Israeli oncology consortium called OncoBridge, Sheba needed a way to produce standardized FHIR data profiles across every dataset. Before implementing the Rhino FCP, the timeline for this process was measured in months:

> **"[Previously], it took us more than one to two months to build one profile... Today with the Rhino platform... *it takes about one or two days from the start to the end*."**
>
> — Nisim Rahman, CTO & Director of Operations, Data & Analytics Division, Sheba Medical Center

Rahman's argument was that harmonization is not optional for an institution with Sheba's collaboration footprint. You cannot federate meaningfully if the data structures at each node are incompatible. Rhino's Data Harmonization Engine enables Sheba to produce compliant FHIR profiles at a pace that makes multi-consortium participation tractable.

Tobias Mann, Vice President of Software Engineering at PacBio, added an important dimension from the genomics world. In domains where data is inherently complex and difficult to map to standardized schemas, waiting for perfect harmonization before federating can become its own bottleneck. Mann's argument was that getting 80 percent of the value from federating across heterogeneous data sooner is often more valuable than waiting for complete harmonization before you start. The goal is not harmonization or federation. It is knowing how much harmonization unlocks the federation value you are after, and building from there.

> **"High quality data is not enough. You have to have a way of looking at that data across sites. The *unstructured data is always going to be more complete*, and large language models are intrinsically designed to handle this really well."**
>
> — Tobias Mann, Vice President of Software Engineering, PacBio

The takeaway is the same whether you start with harmonization or federation: **the goal is getting to insight faster, across more data, than any single institution could manage alone.**

4

The Data Sovereignty Problem is Not Unique to Healthcare

Federated computing's most established use cases come from life sciences and healthcare, but the forum made clear that the underlying pattern extends well beyond them. Three industries surfaced as particularly compelling examples of federated AI applied in new contexts: financial services, industrial supply chains, and agriculture.

The Data-Sovereignty Pattern Repeats Across Industries

INDUSTRY	DATA THAT CREATES VALUE	WHY IT CAN'T MOVE	PROOF
Healthcare & Life Sciences	Clinical, imaging & genomic records	Patient privacy law; data legally bound to site	Moffitt: ~1M patients TuneLab: 100+ partners
Financial Services	Transaction & customer data	Data-residency rules (esp. APAC); rivalry	SWIFT: 11,000+ banks ~3x global GDP daily
Aviation	Loyalty & redemption data	Alliance rivals won't share with each other	Airline Alliance Fraud, personalization
Agriculture	Farm, greenhouse & lab genetics	Competitive IP pressure; never pooled	Inari: Gene-Editing Seed design platform
Industrial Supply Chains	Operational & supplier data	Commercially sensitive; distributed across firms	Cross-supplier Optimization

Oran Barkai, VP of Business Development at Rhino Federated Computing, moderated a financial services panel that anchored the discussion in one of the largest existing data-sharing networks on earth. Holger Roth, Principal Federated Learning Scientist at NVIDIA, pointed to SWIFT — the interbank messaging network connecting over 11,000 institutions and moving an estimated three times the global GDP in transaction value daily — as an illustration of the consortium infrastructure that already exists in financial services, waiting to be activated for federated AI use cases. Roth also described a separate line of NVIDIA research: using AI agents to automatically optimize the configuration of federated learning models across distributed sites, reducing the manual tuning work that data scientists would otherwise have to do by hand.

Maxim Afanasyev, Financial Services Industry Head for AIPAC and Japan at Google Cloud, brought the data sovereignty dimension into focus, particularly as it applies to markets in Asia where data residency regulations are strict. In many of these markets, data simply cannot

leave the country, meaning the institutions with the most to gain from AI are also the most constrained in how they can deploy it. Federated infrastructure is increasingly where that tension gets resolved.

Monisha Malhan, Executive Vice President at TCG, brought cross-industry use cases that illustrated the pattern beyond regulated industries. One example was a major global airline and its alliance partners. Airlines generate significant value from loyalty programs, such as points exchange, fraud detection on redemptions, and personalization, but alliance members will not share that data with each other even when it would benefit all of them. Federated learning offers a way for carriers to take learnings from alliance partners without any underlying data leaving the institution. The structural problem is familiar: valuable insight sitting in data that no one is willing to centralize.

Sonali Quantius, Vice President of R&D at Inari, made the case for agriculture. Inari is a gene-editing seed design company building a predictive platform to determine which genetic edits will improve crop performance across different climates, geographies, and growing conditions. The model needs to train on farm data, greenhouse data, and lab data held by institutions around the world. None of it will be centralized because of intense competitive pressures. Federated computing is the only viable path for Inari to build a model that actually generalizes across the full range of conditions where crops actually grow.

> **"If we do not use federated computing, we will be doing a lot of hard work, but it's *not going to change agriculture* the way we want to."**
>
> — Sonali Quantius, Vice President of R&D, Inari

Taken together, the message is that federated AI's applicability is a function of data structure. **Any industry where valuable data is disaggregated, locally owned, and commercially or legally sensitive is a prime candidate.**

5

Agents Are Only as Powerful as the Data They Can Reach

If there was a single topic where independent voices from different industries converged unanimously, it was agentic AI. Speakers from cloud infrastructure, genomics, agriculture, and Rhino Federated Computing's own R&D team all arrived at the same thesis: the next meaningful layer of value in federated AI goes beyond learning across siloed data. We're soon going to see leaders deploy intelligent agents that can reason over siloed data autonomously without the data ever moving.

Afanasyev stated the core problem plainly. Financial institutions investing in AI are largely running agents on public data. The returns are limited because every competitor has access to the same public models. The differentiation, and the most substantial risk, is in private data. Accessing private data with AI agents in a way that respects sovereignty constraints is exactly the problem federated infrastructure was built to solve:

> **"The real source of differentiation and value is *when you can have your agents run on your private data*. This is where federated learning and confidential computing become fundamental."**
>
> — Maxim Afanasyev, FSI Head for AIPAC and Japan, Google Cloud

Afanasyev's point is about what actually makes one model better than another. Federated AI enables agents to operate at the data source. They can reason, query, and learn within the company's environment and create compounding value for the company, then contributing outputs to a shared model or network without the company's intellectual property. That is what makes the combination of agentic AI and federated infrastructure qualitatively different from simply deploying a large language model on an enterprise data warehouse.

Adrish Sannyasi, VP of Customer Solutions at Rhino Federated Computing, moderated the session and opened with an important clarification. The federated layer matters not only across institutional boundaries but within a single organization, where data is often distributed across product lines, geographies, and acquired companies.

Tal Tiano Einat, VP of R&D at Rhino Federated Computing, described the platform-level implementation of this vision. Rhino has built and deployed a Model Context Protocol (MCP) server that allows AI agents to autonomously drive federated queries, run distributed analytics, and surface results across the network without a data scientist having to orchestrate each step manually. These are tasks that would previously require an expert data scientist's full attention, now compressed into a fraction of the time:

> **"Analyses that could take an expert data scientist a day is something that now *with an agent you can run in five minutes and get results from the entire federated network*."**
>
> — Tal Tiano Einat, VP R&D, Rhino Federated Computing

Mann extended this framing into the clinical setting, describing an architecture where a clinician examines a patient genome with unresolved diagnostic questions, gains access to the combined knowledge of a federated health network, and pattern-matches insights from diverse institutions that have seen similar cases. The agent traverses the orchestration layer, moving between the clinician's question and the distributed data that could answer it. The near-term version of this is nearest neighbor search across federated sites to surface similar cases. Over time, it evolves toward foundation models that learn associations between phenotype and genotype at a scale no single institution could achieve.

Quantius described a structurally similar system at Inari, where agents coordinate across the full predictive design pipeline, from experimental design through model validation, and the knowledge layer spans structured, unstructured, and institutional data that could never be centralized.

The convergence across these four speakers was the clearest signal of the day. Agentic AI has entered the federated AI world faster than anyone could have predicted.

What's Next: The Federated AI Conversation Continues in October

The through-line across every session was a concept that Chris Laws, Chief Commercial Officer at Rhino Federated Computing, named directly: **the federated intelligence network**. As Laws put it:

> "This is *a cross-ecosystem collaboration* where peers, suppliers, customers, and regulators are working with one another in a way where everybody is totally comfortable sharing information, sharing insights, and sharing IP, *with protections in place for every participant*."

The organizations in the room during the Exec Forum are already building toward that vision. The question the field is now working through is how fast federated intelligence networks can scale, and who will be at the table when they do.

That conversation continues at the Federate AI Summit on **October 19th** in Boston. Whether you are bringing a point of view, original research, or a case study from your own organization, there is a place for you at the table. Paper submissions are open and we want to hear from you. Learn more at federateaisummit.com/summit.

Register Today

WWW.FEDERATEAISUMMIT.COM