



Federated AI for Cross-Enterprise Risk Intelligence in the Cloud

Isabella Rossi
Asst Professor

Abstract

To facilitate federated AI and federated cloud architectures for cross-enterprise Risk Intelligence in 2025, the study synthesizes its aims and scope, establishing relevance and impact. Cross-border Risk Intelligence engages insights from distinct Economies for local Governance use. Analysis emissions Cross-Border Information Act compliance, especially at non-country cartel jurisdiction breaks. Cross-Economy Data Management Agency amid disparate Risk Assurance requirements, perpetual audit, activity-based Cross-Economy Digital Services Levy, and unalignment with foreign standards risk disintegration of worldwide trade. Distorted governance dash move to afford globalisation shortcuts via far-from-market subvention need Countervailing Duties—economically falling prove hards through regulatory action, indeed guising phase modulation . Cross-enterprise Risk Analytics thus examine Capital-Source Reporting Standard Data Residency. Preventing dissatisfaction then demand Cross-Border-Compliance-Are Insistence—protocol shopping demand is collapse of a tough wall before regulator. Federated Cloud Multi Cloud and Hybrid Cloud Platform-Are Insight Demand. Latency-Sensitive Federation-Proximity Need Advance Deployment of Cloud-Capacity Engine for Real.Time Data Pre-Processing-Streaming-Risk Signal-Online-Offline Fusion. Domain-/Risk-Signature Stemming Edge-Cloud Edge for Excess-Rate-Deplorable Distance Latency-Sensitive Applications Finally-Are Imperatives Hence-Force Federated Cloud—Federal Cloud—Federated Infrastructure-Are Interest Demand in All.

Keywords: Federated AI Architectures, Federated Cloud Computing, Cross-Enterprise Risk Intelligence, Cross-Border Data Governance, Risk Analytics Platforms, Capital-Source Reporting Standards, Data Residency Compliance, Cross-Border Information Regulation, Multi-Cloud and Hybrid Cloud Federation, Latency-Sensitive Distributed Systems, Edge–Cloud Risk Processing, Real-Time Risk Signal Fusion, Online–Offline Analytics Integration, Digital Services Levy Implications, Regulatory Interoperability Challenges, Countervailing Duties and Trade Governance, Cross-Economy Data Management Agencies, Perpetual Audit Mechanisms, Risk Assurance Harmonization, Federated Infrastructure for Global Governance.

1. Introduction



Risk Intelligence helps organizations assess cross-enterprise risks and better govern their risk environment. Achieving the risk data visibility and analysis needed for risk-informed decision-making requires federated Artificial Intelligence (AI) and cross-enterprise Cloud Computing enabler architectures. Addressing the research question: What are the Architectural Models for Federated AI-Cross Enterprise Risk Analytics in Cloud Computing Environments?[,] involves four steps. First, the introduction develops the problem statement giving rise to the research and specifies the perspective through which it is analyzed. Next, the Federated AI, in particular, the component – Federated Learning and the Cloud Computing model and deployment view currently available for supporting Risk Intelligence in a cross-enterprise cloud environment are explored. Cross-enterprise AI Risk Intelligence helps organizations analyze risk data and Artificial Intelligence (AI) model apps across jurisdictions with different Data Privacy regulations and formal Data Processing Agreements. Risk data is sensitive and protected Data Privacy legislation applied in different cross-border jurisdictions hinders organizations from sharing, nursing, and using it for intelligence and prediction purposes. Federated AI addresses the challenges and limitations of traditional centralized AI paradigm by empowering organizations to collaboratively train across-border models without having to share their raw sensitive data and leveraging the risk analytics from organizations with wider available risk data within jurisdictions permitting Data Transfer and/or share. Deploying the required computing resources in the Cloud with cross-border Data Transfer and residency restrictions provides the required computing resources in proximity to the organization and more demographic/geographic protected risk data.

1.1. Overview of the Study and Its Significance

Cross-enterprise risk assessment, governance, and executive decision support require timely sharing of relevant risk signals. Cross-enterprise risk models require access to sensitive data from cross-border partners. Building shared risk exposure models between enterprises in different jurisdictions poses challenges. Research objectives and expected contributions to practice and research domains for 2025.





Fig 1: Federated AI Architectures for Cross-Enterprise Risk Governance: Harmonizing Multi-Cloud and Edge Solutions for Secure Cross-Border Data Sharing

The aim is to address the lack of federated AI capabilities when using a cloud architecture. Federated AI use cases in the cloud space have lowered the incentives to build scalable federated implementations. Federated learning models have been proposed with a decentralized peer-to-peer architecture, or a centre-point architecture that connects multiple data sources in the cloud. Federated learning fits better with the model but being core in a cloud setup, IaaS, PaaS, SaaS can solve services in a public, private, hybrid or multi-cloud manner. While Cloud computing services are able to utilize Cloud Services Data can still be readily available in a local Data Centre that are owned by the company and due to Company Policies are

not allowed to be stored off-premise, also having services close to the Edge is a benefit. Multi-cloud can lower Latency and Traffic sent to the Cloud, as one Cloud can deliver better latency to one Area and another Cloud to the rest of the Area.



2. Foundational Concepts

Coherently defining foundational concepts helps craft informative analyses. The absence of consistent terminology can hinder clear communication of ideas. Federated learning focuses on training machine learning models across multiple repositories while preserving the data's local ownership, minimizing data transmission, and keeping sensitive data private. In practice, Federated AI is a broader specialization of Federated Learning that encompasses other areas of Artificial Intelligence. Necessary alterations in the model paradigm enable cross-client or cross-domain data issues with varying categories but no sharing of local raw data or labels. For



instance, data from several financial organizations may have an equivalent category 'fraud' with different explanation patterns. The architecture types of Federated Learning for Artificial Intelligent systems are centered, decentralized, and hierarchical. Centered Federated Learning typically adopts a server-client architecture where clients send local model updates to a centralized server for parameter aggregation, yet it has downsides. Centered Federated Learning is frequently unsuitable for data-rich clients and cannot prevent a single point of failure or fault. Decentralized Federated Learning is designed to mitigate these issues, enabling cooperative training without a central server. The hybrid architecture combines both types. These Federated Learning implementations can be used for cross-border Risk Intelligence, and existing works on Federated Learning for Artificial Intelligence can also be used for Risk Intelligence when the application scenarios conform to the aforementioned considerations.

The Cloud Computing model is a continuation of the virtualized computing concept that has been developing into the Cloud Computing field. For cross-enterprises, the Infrastructure as a Service, Platform as a Service, and Software as a Service models in a Public Cloud, Private Cloud, Hybrid Cloud, or Multi-Cloud have different functions with cross-enterprise implications. These different implementations of Cloud Computing provide the infrastructure for cross-enterprise Risk Intelligence. Data and Application providers, Data Consumers, Application Consumers, Clients, and the Cloud provide services such as Data Management, Trading and Sharing, and Application for each layer and model of Cloud Computing.

Equation 1) Federated Learning objective (what is being optimized)

Step 1: Define client datasets

Assume there are K enterprises/clients.

Client k holds local dataset $D_k = \{(x_i^{(k)}, y_i^{(k)})\}_{i=1}^{n_k}$ with $n_k = |D_k|$.

Total samples:

$$n = \sum_{k=1}^K n_k$$



Step 2: Define local empirical risk

Let model parameters be $w \in \mathbb{R}^d$.

Let loss on a sample be $\ell(w; x, y)$.

Local objective at client k :

$$F_k(w) = \frac{1}{n_k} \sum_{i=1}^{n_k} \ell(w; x_i^{(k)}, y_i^{(k)})$$

Step 3: Define global FL objective

Weighted average of local objectives:

$$F(w) = \sum_{k=1}^K \frac{n_k}{n} F_k(w)$$

Goal:

$$w^* = \underset{w}{\operatorname{argmin}} F(w)$$

2.1. Federated Learning and Federated AI

Federated Learning (FL) refers to a scalable machine learning paradigm for decentralized datasets that enables collaborative training and knowledge sharing without the need to centralize sensitive data. Similar to federated structures in social organizations, federated learning is an orchestration approach that defines the underlying relationship model between participants. In the case of natural organizations such as enterprises, a central trust allows for a federated cloud using the cloud server as an orchestration layer outside the local footprint of the enterprise. Federated AI (FAI), an extension to FL, includes any distributed process that orchestrates and manages a learning process supported by artificial intelligence capabilities, such as creating intelligent agents that could otherwise remain in dormant or idle mode. As in traditional cloud architectures, several forms of federated organization can be defined: Centered (a central entity orchestrates a set of participants—client servers, satellites, or others—with a data ownership and security model defined by the policy of the central entity); Hierarchical (special-purpose participants at different levels organize the learning of lower-level players in their area of



interest); or Fully Decentralized (all participants are peers, and the orchestration is done by consensus of voting of the participants). Although these architectures are technically feasible, industry has not been able to use Fedora AI to overcome distrust and centralization issues inherent to any orchestration process.

However, sharing models does not mitigate data locality, privacy, and sovereignty issues. Data residency requirements imposed by regulatory authorities may restrict the ability to build common training datasets and the data-safety-scoring functions required by threat-modeling frameworks. Although organizations such as the Global Infrastructure Facility, Risk Recon, and others offer cloud-scale risk-publication services, in most cases the participation of organizations from particular jurisdictions needs to be defined in advance, a critical drawback for proactive risk modeling. All these aspects hinder the adoption of a federated learning model: All explorers use cloud computing services to pre-process their risk data and paint different geographic spaces, yet common risk models are still not possible.

2.2. Cloud Computing Models in Cross-Enterprise Environments

Cloud computing provides a convenient environment for resource management in application development and deployment. Four main models are recognized, addressing the major types of delivered services: Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS), Software-as-a-Service (SaaS). Enterprises may further choose a public, private, or hybrid cloud deployment model. Service users submit requests to third-party service providers, who manage the facilities, offering on-demand access, reduced management overheads, elasticity, and pay-as-you-go billing. Data may cross national jurisdictions and privacy areas, necessitating careful consideration of legal requirements. From a risk-focused perspective, data in transit is especially vulnerable, demanding additional security measures for delivery.

This paper examines the following aspects of such models that have a bearing on cross-enterprise genuine risk analytics and intelligence sharing: IaaS and PaaS as relevant for data preparation, SaaS in the collection of risk data, potential for a multi-cloud model with critical infrastructures across multiple national jurisdictions, and cloud-edge orchestration in latency-sensitive environments.

Architecture	Scalability (clients)	Single point of failure risk	Bandwidth overhead
Centered (parameter-server)	4	1	3



Architecture	Scalability (clients)	Single point of failure risk	Bandwidth overhead
Fully decentralized (P2P)	2	4	1
Hierarchical (superclients/tiers)	4	3	3

3. Architectural Paradigms for Federated AI in the Cloud

Various architectural paradigms enable the design and deployment of Federated AI solutions in cloud environments. Centered, decentralized, and hierarchical architectures represent the main families. The choice of a specific model hinges on scalability considerations, the level of control over the participating nodes, interoperability requirements, and data locality and sovereignty concerns. Scalability emerges as a crucial aspect because the high resource demand of ML tasks places significant constraints on the number of clients involved in the collaboration. Centered architectures allow multiple clients to share the burden of model training with a cloud service provider but still require access to the shared model; consequently, they may experience performance bottlenecks. Decentralized architectures adapt the peer-to-peer paradigm of communication networks. Rather than delegating model training to a central party, the model is trained collaboratively by the clients in a fully decentralized fashion. However, such topologies may lack the necessary scalability, especially when dealing with large client volumes. Hierarchical architectures introduce a tier of nodes that support model training. These superclients can be deployed in a hybrid setting, allowing the inclusion of clients with low connectivity to the network.

Although centered architectures strive to limit inter-client communications through model sharing via a cloud service provider, data residency concerns may restrain their use. Multilateral probabilistic risk models require sharing risk signals with jurisdictions other than those in which the generating enterprises reside. Similarly, nonlocal data may be subject to jurisdictional access constraints on processing and storage. Thus, risk analytic pipelines must be designed to respect the laws of the involved jurisdictions. The implications of such limitations vary among types of ML tasks: for instance, data residency rules may impede a global organization from collaborating with the nodes generating sensitive data without breaching regulations.

Equation 2) FedAvg / parameter-server aggregation (centered FL)



Step 1: Server broadcasts current model

At round t , server has w_t . It sends w_t to selected clients.

Step 2: Client runs local SGD for E epochs

Initialize local weights:

$$w_{t,0}^{(k)} = w_t$$

For local step $s = 0, 1, \dots, S - 1$, with learning rate η and minibatch $B_{t,s}^{(k)} \subset D_k$:

$$w_{t,s+1}^{(k)} = w_{t,s}^{(k)} - \eta \cdot \frac{1}{|B_{t,s}^{(k)}|} \sum_{(x,y) \in B_{t,s}^{(k)}} \nabla_w \ell(w_{t,s}^{(k)}; x, y)$$

After finishing local training:

$$w_{t+1}^{(k)} := w_{t,S}^{(k)}$$

Step 3: Client sends update (or weights) to server

Two equivalent ways:

(A) Send weights: send $w_{t+1}^{(k)}$

(B) Send delta:

$$\Delta_t^{(k)} = w_{t+1}^{(k)} - w_t$$

Step 4: Server aggregates (FedAvg)

Weighted by local data sizes:

$$w_{t+1} = \sum_{k=1}^K \frac{n_k}{n} w_{t+1}^{(k)}$$



or using deltas:

$$w_{t+1} = w_t + \sum_{k=1}^K \frac{n_k}{n} \Delta_t^{(k)}$$

3.1. Data Locality and Sovereignty in Federated Frameworks

Data locality, storage and processing sovereignty, and jurisdictional issues present significant challenges. An organization's customer and transaction data must be stored in a location that satisfies contractual obligations. Increasingly, such organizations are required to store such data inside the territories where enterprise operation originates. Regulations such as the European Union's [General Data Protection Regulation (GDPR)](<https://gdpr-info.eu/>) are primarily formulated with respect to cross-border data privacy and require the country where the subject resides to be afforded similar levels of protection regardless of where the data reside. Cloud service pricing is largely location-dependent. Clusters providing, processing, or storing services must satisfy jurisdictional data residency requirements because any party operating within a country or jurisdiction is usually liable to local data-related laws.




Fig 2: Navigating the Sovereignty-Scale Paradox: Legal Data Locality Requirements and Cost-Optimization in Federated Cross-Border Risk Analytics

For cloud infrastructures, these aspects raise two significant challenges. First, the costs to satisfy data locality governmental (e.g., data privacy), jurisdictional (e.g., data protection), and



contract-specific Data Locality Requirements might undermine the cross-border risk analytics initiative. Nodes hosting clusters that provide storage, or have data-generation sources, must be located in cost-effective jurisdictions. Second, locality and sovereignty-sensitive Data Locality Requirements must be satisfied while at least some elements of the cross-enterprise risk analytics pipeline process supported by a federated environment are data-driven and require data largeness.

3.2. Cross-Enterprise Collaboration Protocols

Four types of cross-enterprise collaboration aspects contribute to building models for shared risk intelligence: technical protocols that allow federated learning, frameworks that establish mutual trust among collaborating actors, mechanisms that support the orchestration of federations, and approaches that provide incentives for involved organizations.

Collaboration among enterprises to build shared models for risk intelligence entails deploying a Federated Learning (FL) paradigm, where only part of the model is shared. The FL process, whether parameter-server-based or hierarchical, requires the specification of the communication protocols that ensure the correct transfer of the uploaded or downloaded parts of the model. Several proposals exist for defining these technical protocols but no standard has emerged yet that would facilitate implementation across industries, limits, and borders. Major providers of network infrastructure and cloud services have recently initiated alliances that aim to provide a trusted framework within which organizations can safely share data for training AI models. The key objective of these trust frameworks is to define the data protection measures implemented at organizations' data centers that lead to a reduction of the data protection obligations for the party consuming the data for model training. The definition of such a trusted environment is crucial toward enabling the automatic execution of federated-learning processes. Indeed, organizations willing to participate in FL and share data for model training would be enabled to enroll through a trust framework and connect to the existing AI models deployed on the federation in an automated way whenever a policy-compliant data-sharing request is received. Achievement of this goal would greatly simplify the onboarding process into a federation and formally guarantee that protection and monitoring controls are enforced by each organization.

4. Security, Privacy, and Compliance Considerations

Anomaly detection systems deliver risk data belonging to particular enterprises to federated models for enhanced learnings. Such data-sharing exchanges and collaborative models



create undesirable incentives for malicious actors to corrupt the risk data or models. Security mechanisms guarantee the delivery of trustworthy models from different enterprises. Privacy requires the anonymity of the participants to conceal specific enterprises; implementation of end-to-end encryption, zero-knowledge proofs, and cryptographic tools for federated communications strengthen privacy guarantees. Compliance addresses requirements for satisfactory legal operations, particularly in multi-jurisdictional contexts. A robust risk platform should evaluate risk data delivery against established security and privacy controls such that the data, during preprocessing or when building models for a risk prediction task, are adequately secured.

Various controls can be established to ensure the protection of risk data and the compliance of the data processor during each processing phase either for individual risk analysis or for a federated learning setting. Encryption strategies can be adopted at the data holder level such that risk data offered to external entities are published in an encrypted format, rendering them useless. Properly designed encrypted model implementing data aggregation can also prevent exposure of the actual data used for federated risk analysis. Market-driven incentives can tattoo the delivery of trustworthy risk data (by end users) or risk models (by risk model developers) that satisfy established security and privacy controls, thus enhancing the resilience of risk models during a federated learning phase.

Equation 3) Fully decentralized FL (peer-to-peer consensus update)

Let the network be a graph $G = (V, E)$ where nodes are clients.

Let $\mathcal{N}(k)$ be neighbors of node k .

Step 1: Local training step


Client k computes a local SGD step:

$$\tilde{w}_{t+1}^{(k)} = w_t^{(k)} - \eta \nabla F_k(w_t^{(k)})$$

Step 2: Neighbor averaging (consensus / mixing)

Let mixing weights $a_{kj} \geq 0$, and $\sum_{j \in \mathcal{N}(k) \cup \{k\}} a_{kj} = 1$. Then:

$$w_{t+1}^{(k)} = \sum_{j \in \mathcal{N}(k) \cup \{k\}} a_{kj} \tilde{w}_{t+1}^{(j)}$$



4.1. Data Protection Mechanisms and Threat Modeling

Federated learning already addresses some critical aspects related to security, privacy, and compliance issues that often arise in conventional centralized approaches, including direct user data utilization in training algorithms when building the models and the rapid detection of neural network weaknesses through the public release of partially reconstructed models. However, the aforementioned three generalized domain requirements remain more challenging to satisfy in federated learning settings for data privacy-preserving protocols designed for inferential tasks than in PaaS-supported inference-based systems. Security protocols that have been developed for these services can typically be mapped to both federated cloud infrastructures and services (when looking at cloud-residing risk-data-centric architectures) and the use of risk-data-centric federated clouds for centralized risk inference tasks. A similarly straightforward



mapping can also be made when considering the support of data privacy-preserving control mechanisms during risk data processing by conventional and federated cloud platforms. The focus here is, therefore, on the data privacy and security-preserving controls that are established in the federated cloud context.

The implementation of control mechanisms specifically designed to provide data privacy and security guarantees when sharing partial models or model gradients among parties located in different privacy-sensitive jurisdictions is still a topic of ongoing investigation. Encryption of model parameters and encrypted (partial) model aggregation processes are two control classes that have received considerable attention so far. Encryption of model parameters is typically achieved through standard encryption functions supported by an additively homomorphic cryptographic scheme, such as irreversible transforms or time-synchronized one-time pad and constrained authenticated public key cryptographic schemes. A differentially private model-gradient-sharing scheme using additive noise based on the Laplace mechanism has also been specifically designed for federated logistic regression. The open challenges mainly concern user privacy guarantees — including user differential privacy for strong membership privacy — and the additional bandwidth cost associated with shared data. Initial research on adversarial machine learning for federated learning systems is also emerging. Specifically designed adaptive adversaries against the Byzantine-resilient device acquisition process have recently been proposed, highlighting that an attacker can strategically spread unreliable device alerts across wide geographical areas and affect the performance of a wide span of infrastructures, potentially leading to important disruptions.

4.2. Compliance with Regulatory Regimes Across Jurisdictions

A comparative evaluation of regulations in various regions facilitates jurisdictional compliance and promotes collaboration in risk analytics. Several initiatives aim to align data protection regulations, especially pertinent when risks affect multiple regions. Even when regulations do not fully converge, compliant analytics are possible by relating specific controls to regulatory obligations (ISO 27001 and General Data Protection Regulation) or combining controls from different authorities.

The need for cross-enterprise risk analytics increases as threats become more sophisticated, with criminals targeting the weakest links to penetrate even the most advanced organizations. Attacks are seldom contained within a single jurisdiction, and components of the affected infrastructure may belong to entities in different jurisdictions. For example, a breach in



an MSP could compromise multiple clients across the globe, and a supply chain attack can cripple several unrelated enterprises. Nevertheless, when risk is shared across parties governed by different data protection regulations, compliance becomes a challenge, not because of a lack of regulations but due to the absence of synergies among them.

Standardization efforts such as the APEC Cross-border Privacy Rules System for data protection, the Cloud Security Alliance Cloud Controls Matrix for cloud services, the WORKING GROUP ON INTERCONNECTION AND INTEROPERABILITY FOR CLOUD COMPUTING SERVICES OF THE ORGANISATION FOR ECONOMIC CO-OPERATION AND DEVELOPMENT guidelines on cross-border cloud computing services, the ISO 27018 for protecting personal data in the cloud, and the ISO26001 on service management and the General Data Protection Regulation objectives addressed in such initiatives are well known: protect personal data, promote trust, support free trade and consider the risks for data subjects. Other sources indicate standardized compliance enabling joint risk management between jurisdictionally distinct parties. The design of cross-enterprise risk intelligence solutions can take advantage of these alignments.

Control / Mechanism	Primary threat reduced	Trade-off
Transport security (E2E encryption)	Eavesdropping / MITM	Key management complexity
Secure aggregation (encrypted parameter aggregation)	Server learns individual updates	Computation + protocol rounds
Additively homomorphic encryption for model updates	Update leakage to aggregator/peers	Heavy compute; larger ciphertexts
Differential privacy (Laplace noise on gradients)	Membership inference / gradient leakage	Accuracy loss; tuning ϵ

5. Cloud Infrastructure Patterns for Risk Intelligence

Multi-cloud and hybrid cloud infrastructures hold promise for enhancing cross-enterprise risk intelligence. Their deployment models, data-sharing agreements, and governance frameworks are examined to reveal implications for the analysis and modeling of common risk factors. Risk intelligence pipelines that violate data locality or sovereignty might achieve lower overall risk models faster by connecting to data outside of the trust domain. However, when federated orchestration addresses these constraints, multi-cloud and hybrid deployments may



improve robustness against preferred destination outages, betweenness centrality attacks, or other risks of local data sources.

A solid edge-cloud connectivity pattern is crucial for latency-sensitive risk analytics pipelines, especially during an incident when escaping actions must be initiated quickly. The architecture should support the pre-processing of data close to their sources—measuring and transforming into risk signals that can be streamed using a well-known protocol for online fusion—while not hindering the bulk data transfer to the cloud for machine learning and other algorithms that need both historical context and low latency. A full classifier exploration in the risk signal fusion should also be allowed, leveraging the power of cloud machine learning. For advanced risk modeling, the online and offline risk intelligence analytics must survive a longer incident period without data connection to the edge. Remotely generated signals can therefore serve a dual purpose: they contribute the much-needed historical context to the cause-modeling cloud layer and transform the risk signal conditioning and detection into an offline problem suitable for slower algorithms not requiring a cloud machine.

Equation 4) Hierarchical FL (edge/region “superclients”)

Assume a 2-level hierarchy:

- Leaf clients inside region r : $k \in \mathcal{C}_r$
- Regional aggregator (“superclient”) outputs w_r
- Global aggregator outputs w

Step 1: Leaf $\rightarrow$ regional aggregation

After local training leaf produces $w^{(k)}$.

Regional aggregation:

$$w_r = \sum_{k \in \mathcal{C}_r} \frac{n_k}{\sum_{j \in \mathcal{C}_r} n_j} w^{(k)}$$



Step 2: Regional → global aggregation

Let region r have total samples $N_r = \sum_{k \in \mathcal{C}_r} n_k$.

Global:

$$w = \sum_r \frac{N_r}{\sum_q N_q} w_r$$

5.1. Multi-Cloud and Hybrid Cloud Interoperability

Scale and complexity drive many organizations to disperse their operations over multiple clouds, consuming various services from different providers. Multi-cloud deployment, however, is a means to distribute risk and avoid vendor lock-in, not a goal in itself. Failure rates vary across providers, and security exposures depend on internal practices even more than on provider properties. Data sharing agreements can help determine which cloud environment is used for enterprise risk analytics. Nevertheless, requirements, agreements, and business context sometimes force the use of different providers that cannot easily interconnect. A multi-cloud architecture improves availability for cloud users, but additional risks stem from the increasing number of parties and provider relations.

An intermediate approach is the use of hybrid clouds. In this environment, a private cloud sits behind the firewall but the management of sensitive information follows the same rules as when using a public cloud. The company deals with the private cloud and entrusts a third-party provider with anything not covered by the internal cloud. In many cases, the private cloud is located in the data center of an enterprise with which data are shared. The external part may still be a SaaS model maintained by a private or public actor, but the cost-effectiveness of the model is much higher. These two types of deployment are not necessarily opposed; in practice, they can coexist without problems.




Fig 3: Beyond Connectivity: Navigating the Risk-Utility Trade-off in Hybrid and Multi-Cloud Architectures

5.2. Edge-Connected Cloud for Latency-Sensitive Analytics



Risk analytics pipelines for latency-sensitive business activity monitoring, transaction risk evaluation, and scenario-based data science may benefit from proximity to source data. Analytics task sequences are hence distributed across Edge and Cloud resources, with Edge resources pre-processing or merging data before streaming second-stage descriptive analytics triggers to Cloud resources. Other online–offline synchronization patterns (i.e., using Cloud resources to feed back an online inference result list to the Edge for visuals or updating streaming risk signals using offline analysis as a template) are also in scope.

The edge-connected cloud pattern supports edge-enabled applications with high responsiveness requirements and lower processing requirements on the uploaded data, while still benefitting from leveraging aggregate data features. Applications may include IoT sensor data traffic risk modeling under unplanned events, sensitive business activity monitoring, and cross-organizational fusion of streaming risk analytics, complemented by Cloud-powered sifting through the volume of streamed risk signals.

Pipeline	End-to-end latency (ms)
Cloud-heavy	83
Edge-heavy	63

6. Data Governance and Quality for Cross-Enterprise Insights

Prerequisites and foundations for generating reliable cross-enterprise insights encompass effective data governance and quality, especially in shared settings characterized by local ownership and provenance. Data governance structures define roles and procedures for ensuring that data use aligns with set intentions. Data quality is crucial for trusted analytics, and rigorous process models are needed to ensure that data is fit for intended uses.

Data provenance describes the lifecycle of data, answering who created, modified, and used it, and how it was transformed. Data lineage collecting details of all transformations allows auditability. For multi-party processes, clearly defined data roles—data owner, data steward, data user—help maintain decision authority and control over data elements while enabling seamless flow to execution functions.

Formal data quality metrics, defined for product development, supplied from collecting services, assessed and tracked throughout the processing chain, and supplemented with data-cleansing and data-validation processes, support trusted cross-organizational inferences of risk



status or signals supporting timely remediation Decision-supporting data, such as residential location and recent behavior of customers, is typically hosted in a central cloud. The user journey and sentiment are traced to drive service adaptation.

6.1. Data Provenance and Lineage

Data provenance focuses on answering questions about data origins and transformations throughout its lifecycle. For risk intelligence, it is critical to trace the sources of key risk signal inputs, such as third-party assessments, and verify results from generative risk models, including cyber insurance. However, business users often lack the queries and analytical skills required to construct data flows. Within an enterprise, departments may enforce their definitions of risk-relevant terms through business rules or workflows, which can be captured by impact analysis tools relying on link analysis. Yet, knowledge about risk signal generation rarely traverses enterprise boundaries, making data lineage tracking indispensable for detecting approval-path violations across departments.

Data-steering committees can facilitate data sharing between trusted partners. Local Data Governance Centers (LDGCs) help organizations store, manage, and share sensitive data while granting complete visibility and control through a central gateway. Approval processes determine whether a company is authorized to receive sensitive data. Automating these contracts and making them self-enforcing can simplify verification. Public blockchains aid data provenance tracking by providing proof of the data supply chain, secure ownership management, and a trusted environment for smart contracts. These technologies allow users to query publicly available provenance information, even when sensitive data remains private. Data provenance defined by Enterprise Data Governance Information Model (EDGIM) ensures that intelligent risk signals—especially those from internal data sources—adhere to established enterprise-wide definitions, can pass audit trail checks, and are suitable for steering enterprise-wide responses.

Equation 5) Logistic regression (since paper mentions DP for federated logistic regression)

Step 1: Model + probability

Binary labels $y \in \{0,1\}$, features $x \in \mathbb{R}^d$.

Logit:

$$z = w^T x$$



Sigmoid:

$$\sigma(z) = \frac{1}{1 + e^{-z}}$$

Predicted probability:

$$p(y = 1|x; w) = \sigma(w^T x)$$

Step 2: Negative log-likelihood (cross-entropy) loss

For one sample (x, y) :

$$\ell(w; x, y) = -\left(y \log(\sigma(w^T x)) + (1 - y) \log(1 - \sigma(w^T x))\right)$$

Step 3: Derivative of sigmoid (needed for gradient)

Let $p = \sigma(z)$. Then:

$$\frac{dp}{dz} = p(1 - p)$$

Step 4: Gradient of loss for one sample

A standard result (derived via chain rule) is:


$$\nabla_w \ell(w; x, y) = (p - y) x$$

where $p = \sigma(w^T x)$.

Step 5: Client k gradient on local data

$$\nabla F_k(w) = \frac{1}{n_k} \sum_{i=1}^{n_k} \left(\sigma(w^T x_i^{(k)}) - y_i^{(k)} \right) x_i^{(k)}$$

$$w \leftarrow w - \eta \nabla F_k(w)$$



6.2. Data Quality Metrics and lineage-aware Modeling

The achievement of cross-enterprise risk insights hinges upon the soundness of the underlying data, whose quality must be guaranteed through appropriate metrics, modeling, and implementation techniques. To facilitate this goal, the primary requirements of data quality for trustworthy risk inferences relate to correct and timely information about the presence and occurrence of risks and are as follows:



- Confederating data from disjoint sources increases the likelihood of inconsistency and inaccuracy in the information. Hence, the higher the propagation of these factors during the aggregation of risk-related data, the lower the reliability of the predictive model;
- The elements used to construct the prediction model must be complete in their representation of the phenomena being predicted, meaning that they should not exclude any relevant group or event within that phenomenon;
- The more recent the data, the better able they are to represent the present conditions of the phenomenon being predicted, and the smaller the guarantee that these conditions have changed during the prediction period;
- The more accessible the elements monitored during the prediction, the more reliable the predictions can be;
- The speed of the events that the model predicts must be taken into account. For rapidly changing phenomena, the use of data from longer temporal horizons would make the model less reliable.

Facilitating these requirements involves specifying the metrics that indicate the presence of quality problems in data and covering the modeling and processing of data in order to alleviate imperfections in ground truth and testing datasets. Risk signal data often provides valuable information about emerging dangers, and exploring such signals increases the coverage of prediction models and improves their performance. However, the risk signals available for distribution are produced in a short time by edge devices,¹⁸ and to make the most of these risk information in prediction systems, special care in the temporal alignment of these signals is essential.

7. Conclusion

The outcomes yield designs for effective federated AI and cloud computing infrastructures enabling durable cross-enterprise risk intelligence, define protection controls across shared data flows, point out data-sharing incentives within regulatory frameworks, and illustrate end-to-end data cleansing facilities. Together, these findings assist corporations in safeguarding shared risk models. Consolidated risk intelligence on a specialized federated AI-as-a-service platform benefits the overall eco-system, with all charged parties investing in aspects



they individually require. Long-term applications rely on a wider trust network covering complete risk domains.

Longitudinal trajectories for federated AI and cloud data-engineering architecture indicate a need for guidance on cross-enterprise data-sharing agreements, address control-testing setups for joint technology audits, and explore multi-stakeholder policies for combined threat-response capabilities. Long-lasting standards accelerate portable deployments based on federated Analytics-as-a-Service and combined consideration of territory requirements for the cybersecurity data-sharing regulation. Research in these areas makes progress towards establishing federated AI and cloud-computing architectures that consistently support cross-enterprise risk intelligence in 2025.




Fig 4: Core Risk Intelligence Outcomes

7.1. Final Thoughts and Future Directions

Looking beyond the presented work, there are at least three key, longitudinal themes. First, creating guides for the next generation of federated AI—making it easier to build and deploy such applications—seems a natural direction. An initiative to develop an ISO standard for federated AI could serve as a valuable, community-based reference. Second, all of these aspects could be considered in the context of risk intelligence: understanding how to collaborate better and addressing the broader challenge of sharing data, processes, and knowledge across borders. Third, the presentation of risk management processes—especially IT risk management—could



be expanded to better cover the cross-enterprise aspects. Addressing these questions, both from a technical and a business perspective, would make for a comprehensive treatment of risk and risk management.

The work aims to advance federated AI for cloud environments able to support different analytical capabilities in a scalable and interoperable manner. Also, implications of the analysis constitute actionable guidelines when building such infrastructures—the alignment and adoption of common governance principles can, in fact, leverage both the efficiency and the trustworthiness of infrastructure-as-a-services environments. More generally, the findings provide insights on the territory potentially paving the way toward a Federated AI ecosystem for cross-business risk intelligence by 2025, with clear objectives in the lines of effective risk data governance, compliance, and improved resilience directives.

References

1. Aledhari, M., Razzak, R., Parizi, R. M., & Saeed, F. (2020). Federated learning: A survey on enabling technologies, protocols, and applications. *IEEE Access*, 8, 140699–140725.
2. Baracaldo, N., Ludwig, H., Joshi, A., & Zhang, Y. (2020). Federated learning for enterprise data collaboration: Principles and challenges. *IBM Journal of Research and Development*, 64(6), 1–12.
3. Peddi, R. K. (2024). AI-Based Workforce Analytics for SLA Governance and Uptime Assurance in Data Centers. *Journal of Computational Analysis and Applications (JoCAAA)*, 33(08), 8589-8601.
4. Bonawitz, K., Eichner, H., Grieskamp, W., Huba, D., Ingerman, A., Ivanov, V., Kiddon, C., Konecny, J., Mazzocchi, S., McMahan, H. B., Van Overveldt, T., Petrou, D., Ramage, D., & Roselander, J. (2020). Towards federated learning at scale: System design. *Proceedings of Machine Learning and Systems*, 2, 374–388.
5. Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., D'Oliveira, R., Eichner, H., El Rouayheb, S., Evans, D., Gardner, J., Garrett, Z., Gascón, A., Ghazi, B., Gibbons, P. B., & Zhao, S. (2021). Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*, 14(1–2), 1–210.
6. Li, Q., Wen, Z., Wu, Z., Hu, S., Wang, N., Li, Y., Liu, X., & He, B. (2020). A survey on federated learning systems: Vision, hype and reality for data privacy and protection. *IEEE Transactions on Knowledge and Data Engineering*.



7. Soni, H., Perla, S., Maddela, S., & Kumar, U. (2025, November). Generative AI in Cloud CRM: Securing Intelligent Workflows in Multi-Cloud Environments. In 2025 IEEE 3rd Global Conference on Wireless Computing and Networking (GCWCN) (pp. 1-9). IEEE.
8. Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated optimization in heterogeneous networks. *Proceedings of Machine Learning and Systems*, 2, 429–450.
9. Ludwig, H., Baracaldo, N., Thomas, G., Zhou, Y., Anwar, A., Rajamoni, S., Ong, Y., Radhakrishnan, J., Verma, A., Sinn, M., Purcell, M., Rawat, A., Minh, T., Holohan, N., Chakraborty, S., Witherspoon, S., Steuer, D., Wynter, L., Hassan, H., & Laguna, S. (2020). IBM federated learning: An enterprise framework. *arXiv Preprint arXiv:2007.10987*.
10. Lyu, L., Yu, H., & Yang, Q. (2020). Threats to federated learning: A survey. *arXiv Preprint arXiv:2003.02133*.
11. Pereira, K., Vinagre, J., Alonso, A. N., Coelho, F., & Carvalho, M. (2022, September). Privacy-preserving machine learning in life insurance risk prediction. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases* (pp. 44-52). Cham: Springer Nature Switzerland.
12. McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & Aguera y Arcas, B. (2020). Communication-efficient learning of deep networks from decentralized data. *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, 1273–1282.
13. Mothukuri, V., Parizi, R. M., Pouriyeh, S., Huang, Y., Dehghantanha, A., & Srivastava, G. (2021). A survey on security and privacy of federated learning. *Future Generation Computer Systems*, 115, 619–640.
14. Nagabhyru, K. C., & Engineer, S. D. (2023). Unifying Data Engineering and Machine Learning Pipelines: An Enterprise Roadmap to Automated Model Deployment.
15. Nguyen, D. C., Ding, M., Pathirana, P. N., Seneviratne, A., Li, J., Niyato, D., Poor, H. V., & Kim, D. I. (2021). Federated learning for Internet of Things: A comprehensive survey. *IEEE Communications Surveys & Tutorials*, 23(3), 1622–1658.
16. Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., Bakas, S., Galtier, M. N., Landman, B. A., Maier-Hein, K., Ourselin, S., Sheller, M., Summers, R. M., Trask, A., Xu, D., Baust, M., & Cardoso, M. J. (2020). The future of digital health with federated learning. *NPJ Digital Medicine*, 3(119), 1–7.
17. Rodriguez-Barroso, N., Jiménez-López, D., Luzón, M. V., Herrera, F., & Martínez-Cámara, E. (2022). Survey on federated learning threats: Concepts, taxonomy on attacks and defences, experimental study and challenges. *IEEE Access*, 10, 46448–46480.
18. Aitha, A. R. (2023). Cloud-Native Big Data AI/ML Framework for Risk Intelligence and Fraud Control in Banking and Insurance Ecosystems. Available at SSRN 6157967.



19. Savazzi, S., Nicoli, M., & Rampa, V. (2020). Federated learning with cooperating devices: A consensus approach for massive IoT networks. *IEEE Internet of Things Journal*, 7(5), 4641–4654.
20. Shen, M., Tang, X., Zhu, L., Du, X., & Guizani, M. (2022). Privacy-preserving support vector machine training over blockchain-based federated learning networks. *IEEE Network*, 36(3), 72–78.
21. Sun, Y., Esaki, H., & Ambrosin, M. (2021). Federated learning for intelligent network management: A survey. *IEEE Communications Surveys & Tutorials*, 23(4), 2408–2439.
22. Mangalampalli, B. M. Intelligent Data Profiling for Healthcare Data Lakes Using AI-Enhanced Analytics.
23. Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2020). Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology*, 10(2), 1–19.
24. Zhang, C., Xie, Y., Bai, H., Yu, B., Li, W., & Gao, Y. (2021). A survey on federated learning: Security, privacy, and applications. *IEEE Access*, 9, 150763–150785.
25. Zhang, W., Zhou, T., Lu, Q., Wang, X., Zhu, C., Sun, H., Wang, Z., Lo, S. K., & Xu, F. (2022). Dynamic-fusion-based federated learning for COVID-19 detection. *IEEE Internet of Things Journal*, 9(24), 24385–24398.
26. Yandamuri, U. S. (2023). An Intelligent Analytics Framework Combining Big Data and Machine Learning for Business Forecasting. *International Journal Of Finance*, 36(6), 682–706.
27. Zhou, Z., Miao, C., Wang, X., Li, Y., & Su, C. (2021). FedFA: Federated learning with feature alignment for heterogeneous data. *IEEE Transactions on Neural Networks and Learning Systems*.
28. Abdullahi, M. B., Ngadi, M. A., Dishing, S. I., & Abdulhamid, S. M. (2021). Security and privacy in cloud computing: A survey of recent developments. *Journal of Network and Computer Applications*, 177, 102980.
29. Bagdasaryan, E., Veit, A., Hua, Y., Estrin, D., & Shmatikov, V. (2020). How to backdoor federated learning. *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics*, 2938–2948.
30. Kolla, T., & Kolla, S. K. (2023). FHIR-Based Real-Time Healthcare Analytics using Unsupervised Learning. *International Journal of Future Innovative Science and Technology (IJFIST)*, 6(6), 11751.
31. Blanco-Justicia, A., Domingo-Ferrer, J., Martínez, S., Sánchez, D., & Flanagan, A. (2020). Achieving security and privacy in federated learning systems: Survey and future directions. *Electronics*, 9(11), 1914.



32. Chen, M., Poor, H. V., Saad, W., & Cui, S. (2021). Wireless communications for collaborative federated learning. *IEEE Communications Magazine*, 59(4), 48–54.
33. Elbamby, M. S., Perfecto, C., Bennis, M., & Doppler, K. (2021). Toward low-latency and ultra-reliable virtual reality. *Proceedings of the IEEE*, 109(7), 1119–1140.
34. Inala, R. AI-Powered Investment Decision Support Systems: Building Smart Data Products with Embedded Governance Controls.
35. Ghosh, A., Chung, J., Yin, D., & Ramchandran, K. (2020). An efficient framework for clustered federated learning. *Advances in Neural Information Processing Systems*, 33, 19586–19597.
36. Hard, A., Rao, K., Mathews, R., Ramaswamy, S., Beaufays, F., Augenstein, S., Eichner, H., Kiddon, C., & Ramage, D. (2020). Federated learning for mobile keyboard prediction. *arXiv Preprint arXiv:1811.03604*.
37. Javed, A. R., Beg, M. O., Asim, M., Baker, T., & Tawfik, H. (2022). AI-enabled secure communication in cloud computing: A survey. *Journal of Cloud Computing*, 11(1), 1–27.
38. Mattaparthi, R. (2023). Connected Fleet Intelligence: Edge-Centric Analytics and Computer Vision for Predictive Manufacturing and Asset Resilience. *International Journal of Advanced Research in Computer Science & Technology (IJARCST)*, 6(5), 9077-9088.
39. Li, X., Huang, K., Yang, W., Wang, S., & Zhang, Z. (2020). On the convergence of FedAvg on non-IID data. *International Conference on Learning Representations (ICLR)*.
40. Lim, W. Y. B., Luong, N. C., Hoang, D. T., Jiao, Y., Liang, Y. C., Yang, Q., Niyato, D., & Miao, C. (2020). Federated learning in mobile edge networks: A comprehensive survey. *IEEE Communications Surveys & Tutorials*, 22(3), 2031–2063.
41. Lu, Y. (2021). Artificial intelligence: A survey on evolution, models, applications, and future trends. *Journal of Management Analytics*, 8(1), 1–29.
42. Mao, Y., You, C., Zhang, J., Huang, K., & Letaief, K. B. (2020). A survey on mobile edge computing: The communication perspective. *IEEE Communications Surveys & Tutorials*, 19(4), 2322–2358.
43. Kolla, S. K. (2023). Learning Health Systems Machine Intelligence for Clinical Prediction and Healthcare Optimization. *International Journal of Advanced Research in Computer Science & Technology (IJARCST)*, 6(2), 7955-7966.
44. Mohri, M., Sivek, G., & Suresh, A. T. (2021). Agnostic federated learning. *Proceedings of the 36th International Conference on Machine Learning*, 4615–4625.
45. Omoniwa, B., Hussain, R., Javed, M. A., Bouk, S. H., & Malik, S. U. R. (2021). Fog/edge computing-based IoT security: A survey. *IEEE Internet of Things Journal*, 8(10), 8227–8258.
46. Park, J., Samarakoon, S., Bennis, M., & Debbah, M. (2021). Wireless network intelligence at the edge. *Proceedings of the IEEE*, 107(11), 2204–2239.



47. Qi, Q., Tao, F., Hu, T., Anwer, N., Liu, A., Wei, Y., Wang, L., & Nee, A. Y. C. (2021). Enabling technologies and tools for digital twin. *Journal of Manufacturing Systems*, 58, 3–21.
48. Reddy, V. A. R., & Kolla, S. K. (1984). Infrastructure-As-Code Practices For Regulated Healthcare Cloud Environments. *Metallurgical and Materials Engineering*, 30 (4), 1028–1042.
49. Rahman, M. A., Hossain, M. S., Islam, M. S., Alrajeh, N. A., & Muhammad, G. (2020). Secure and privacy-preserving IoT communication using blockchain and federated learning. *IEEE Access*, 8, 188163–188176.
50. Saeed, F., Parizi, R. M., Dehghantanha, A., Choo, K. K. R., & Srivastava, G. (2021). Federated learning empowered blockchain for distributed network security. *IEEE Network*, 35(5), 48–55.
51. Wang, S., Tuor, T., Salonidis, T., Leung, K. K., Makaya, C., He, T., & Chan, K. (2021). Adaptive federated learning in resource-constrained edge computing systems. *IEEE Journal on Selected Areas in Communications*, 37(6), 1205–1221.
52. Bandi, V. D. V. K. (2023). MLOps frameworks for reliable model deployment in cloud data platforms. *Journal of Artificial Intelligence and Big Data*, 3(1), 84-101.
53. Xu, J., Glicksberg, B. S., Su, C., Walker, P., Bian, J., & Wang, F. (2021). Federated learning for healthcare informatics. *Journal of Healthcare Informatics Research*, 5(1), 1–19.
54. Zhang, Y., Chen, X., Li, D., Wang, H., & Jin, H. (2020). Towards efficient and privacy-preserving federated deep learning. *IEEE Transactions on Parallel and Distributed Systems*, 31(12), 2828–2842.
55. Truex, S., Baracaldo, N., Anwar, A., Steinke, T., Ludwig, H., Zhang, R., & Zhou, Y. (2020). A hybrid approach to privacy-preserving federated learning. *Proceedings of the 12th ACM Workshop on Artificial Intelligence and Security*, 1–11.
56. Sheller, M. J., Reina, G. A., Edwards, B., Martin, J., & Bakas, S. (2020). Multi-institutional deep learning modeling without sharing patient data: A feasibility study on brain tumor segmentation. *Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries*, 11992, 92–104.
57. Yandamuri, U. S. (2024). AI-Driven Decision Support Systems for Operational Optimization in Hospitality Technology. *Metallurgical and Materials Engineering*.
58. Dayan, I., Roth, H. R., Zhong, A., Harouni, A., Gentili, A., Abidin, A. Z., Liu, A., Costa, A. B., Wood, B. J., Tsai, C. S., & others. (2021). Federated learning for predicting clinical outcomes in patients with COVID-19. *Nature Medicine*, 27(10), 1735–1743.



59. Li, X., Gu, Y., Dvornek, N., Ventola, P., & Duncan, J. S. (2021). Privacy-preserving federated brain tumour segmentation. *Machine Learning in Medical Imaging*, 12966, 133–141.
60. Feng, J., Wang, H., & Yu, P. S. (2021). Machine learning for cloud security: A survey. *IEEE Access*, 9, 119490–119512.
61. Koutroumpouchos, N., Kogias, D. G., & Vassilakis, V. G. (2022). Artificial intelligence in cybersecurity: Current trends and future directions. *Computers & Security*, 115, 102603.
62. Sharma, P. K., Kumar, N., Park, J. H., & Cho, K. (2021). Blockchain and federated learning for secure Internet of Things applications. *Future Generation Computer Systems*, 124, 345–362.
63. Kolla, S. H. (2024). Retrieval-Augmented Enterprise Intelligence: Enhancing Accuracy, Trust, and Operational Decision-Making. *International Journal of Future Innovative Science and Technology (IJFIST)*, 7(2), 12425.
64. Khan, L. U., Saad, W., Han, Z., Debbah, M., & Hong, C. S. (2021). Federated learning for Internet of Things: Recent advances, taxonomy, and open challenges. *IEEE Communications Surveys & Tutorials*, 23(3), 1759–1799.
65. Nguyen, T. D., Marchal, S., Miettinen, M., Fereidooni, H., Asokan, N., & Sadeghi, A. R. (2021). D²IoT: A federated self-learning anomaly detection system for IoT. *IEEE Transactions on Mobile Computing*, 20(3), 756–769.
66. Islam, M. M., Rahaman, A., & Islam, M. R. (2022). Development of smart healthcare monitoring system using IoT and cloud computing. *Healthcare Technology Letters*, 9(1), 1–10.
67. Gottimukkala, V. R. R. (2020). Energy-Efficient Design Patterns for Large-Scale Banking Applications Deployed on AWS Cloud. *power*, 9(12).
68. Zhou, T., Wang, X., & Luo, X. (2022). Secure federated learning with blockchain for industrial Internet of Things. *IEEE Internet of Things Journal*, 9(9), 6355–6367.
69. Qayyum, A., Ahmad, K., Ahsan, M. A., Al-Fuqaha, A., Qadir, J., & Qolomany, B. (2021). Collaborative federated learning for healthcare: Multi-institutional AI without data sharing. *IEEE Journal of Biomedical and Health Informatics*, 25(9), 3528–3540.
70. Alazab, M., Tang, M., Luo, Y., & Khan, S. U. (2021). Cybersecurity and artificial intelligence: A review of emerging technologies. *IEEE Transactions on Artificial Intelligence*, 2(5), 397–409.
71. Nguyen, D. C., Pathirana, P. N., Ding, M., Seneviratne, A., Li, J., & Niyato, D. (2022). Privacy-preserved task offloading in federated cloud-edge intelligence. *IEEE Network*, 36(1), 52–59.



72. Mangala, N. (2024). Leveraging Microsoft Fabric lakehouse as an AI-ready data platform for enterprise analytics. *Journal of Information Systems Engineering and Management*.
73. Chai, D., Wang, L., Chen, K., & Yang, Q. (2022). Secure federated learning: Fundamentals, techniques, and applications. *ACM Computing Surveys*, 55(9), 1–37.
74. Mothukuri, V., Khare, P., Parizi, R. M., Pouriyeh, S., Dehghantanha, A., & Srivastava, G. (2022). Federated learning-based anomaly detection in cloud computing: A comprehensive review. *Journal of Information Security and Applications*, 67, 103190.
75. Kolla, S. K. (2022). Engineering Healthcare Data Infrastructures for Predictive Clinical Analytics and Evidence-Based Decision Making. *International Journal of Engineering & Extended Technologies Research (IJEETR)*, 4(5), 5370-5380. |
76. Khan, W. Z., Rehman, M. H., Zangoti, H. M., Afzal, M. K., Armi, N., & Salah, K. (2022). Industrial Internet of Things: Recent advances, enabling technologies, security and privacy challenges. *IEEE Open Journal of the Computer Society*, 3, 103–124.
77. Yang, H., Zhang, J., & Que, X. (2023). Trustworthy federated learning for cloud-edge collaborative intelligence. *Future Generation Computer Systems*, 140, 286–299.
78. Zhou, Y., Baracaldo, N., Joshi, A., & Ludwig, H. (2023). Enterprise federated learning: Architecture, security, and deployment considerations. *IBM Journal of Research and Development*, 67(1/2), 1–14.
79. Kolla, T. (2024). Graph Neural Networks for HCC Risk Adjustment and Interoperability. *International Journal of Science, Research and Technology*, 7(6), 13244-13255.
80. Abdel-Basset, M., Mohamed, R., Sallam, K. M., & Chakraborty, R. K. (2023). Federated learning in smart environments: A comprehensive review. *Information Fusion*, 90, 62–101.
81. Abdulrahman, S., Tout, H., Mourad, A., & Talhi, C. (2023). Federated learning for secure and privacy-preserving artificial intelligence: A survey. *IEEE Access*, 11, 56118–56149.
82. Chai, D., Qu, C., Li, X., Yang, Q., & Chen, K. (2023). A survey of federated learning: Algorithms, systems, and applications. *ACM Computing Surveys*, 56(2), 1–37.
83. Reddy, V. A. R. (2023). Orchestrating the Future Autonomous Healthcare Data Pipeline Management through Agentic AI Architectures. *International Journal of Advanced Research in Computer Science & Technology (IJARCST)*, 6(2), 7979-7992.
84. Chen, J., Xiong, Z., Wang, J., Zhang, D., Niyato, D., & Leung, V. C. M. (2023). Trustworthy federated learning for intelligent edge computing: Recent advances and future directions. *IEEE Network*, 37(3), 44–51.
85. Elbir, A. M., Sonmez, C., Coleri, S., & Alouini, M. S. (2023). Federated learning in next-generation wireless networks: A survey. *IEEE Communications Surveys & Tutorials*, 25(2), 879–917.



86. Guo, Y., Liu, Y., Zhang, T., & Tong, Y. (2023). Privacy-preserving federated learning: Concepts, methods, and applications. *Information Sciences*, 633, 263–289.
87. Han, T., Liu, X., Wang, L., & Zhang, H. (2023). Secure federated learning with differential privacy for distributed artificial intelligence. *Future Generation Computer Systems*, 145, 353–366.
88. Amistapuram, K. (2024). Smart Decision Support Systems For Dynamic Tax Policy Optimization Using Reinforcement Learning. Available at SSRN 6143426.
89. Li, H., Ota, K., & Dong, M. (2023). Learning IoT in edge: Deep learning for the Internet of Things with edge computing. *IEEE Network*, 37(1), 32–39.
90. Nguyen, D. C., Ding, M., Pathirana, P. N., Seneviratne, A., Li, J., & Niyato, D. (2023). Federated learning meets blockchain in edge intelligence: A survey. *IEEE Internet of Things Journal*, 10(4), 2910–2934.
91. Park, J., Lim, W. Y. B., Samarakoon, S., Kim, M., & Bennis, M. (2023). Communication-efficient federated learning for beyond-5G networks. *IEEE Wireless Communications*, 30(2), 40–47.
92. Inala, R. Advancing Group Insurance Solutions Through Ai-Enhanced Technology Architectures And Big Data Insights.
93. Qiu, H., Memmi, G., Lu, L., Qiu, M., & Thuraisingham, B. (2023). Secure and trustworthy federated learning: Taxonomy, attacks, and defenses. *ACM Computing Surveys*, 56(4), 1–39.
94. Raza, S., Wang, S., Ahmed, M., & Anwar, M. R. (2023). Artificial intelligence for enterprise cybersecurity: Trends, challenges, and opportunities. *Computers & Security*, 126, 103084.
95. Davuluri, P. N. AI-Augmented Sanctions Screening: Enhancing Accuracy and Latency in Real Time Compliance Systems.
96. Wang, T., Zhu, H., Li, Y., & Yu, F. R. (2023). Federated intelligence for cloud-edge collaborative systems: Architectures and applications. *IEEE Network*, 37(5), 158–165.
97. Wu, Y., He, C., Li, Y., Chen, Z., & Zhang, J. (2024). Recent advances in trustworthy federated learning: Security, privacy, and robustness. *IEEE Transactions on Artificial Intelligence*, 5(1), 55–72.
98. Mangalampalli, B. M. Generative AI Applications In Healthcare Data Mart Design And Optimization.
99. Yang, Q., Liu, Y., Cheng, Y., Kang, Y., Chen, T., & Yu, H. (2024). Federated learning: Principles, technologies, and applications. *ACM Transactions on Intelligent Systems and Technology*, 15(1), 1–34.
100. Zhang, H., Chen, J., Wang, Y., & Li, X. (2024). Cross-enterprise federated artificial intelligence for secure cloud collaboration. *Future Generation Computer Systems*, 151, 52–67.



101. Zhao, B., Xu, Y., Wang, C., & Liu, J. (2024). Privacy-enhanced federated learning in cloud computing environments: A review. *Journal of Network and Computer Applications*, 233, 103935.
102. Zhou, Y., Ludwig, H., Baracaldo, N., Joshi, A., & Thomas, G. (2024). Enterprise federated learning: Recent advances, deployment strategies, and future directions. *IBM Journal of Research and Development*, 68(1), 1–15.
103. Zhuang, W., Gan, X., Wen, Y., Zhang, S., & Yi, S. (2024). Federated learning for intelligent cloud-edge systems: A survey of architectures, optimization, and security. *IEEE Open Journal of the Computer Society*, 5, 41–67.
104. Kairouz, P., McMahan, H. B., Yang, Q., & Bellet, A. (2024). Federated learning for trustworthy artificial intelligence: Challenges and opportunities. *Nature Machine Intelligence*, 6(3), 245–257.