International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 2, March – April 2025||

DOI:10.15662/IJRPETM.2025.0802008

Privacy-Aware Cross-Domain Intelligence with Federated Generative AI for Regulated Enterprise Systems

Dr.R.Sugumar

Professor, Department of Computer Science & Engineering, SIMATS Engineering, Saveetha Institute of Medical and Technical Sciences (SIMATS), Chennai, India

ABSTRACT: Organized business intelligence is increasingly requiring high levels of privacy, security, governmental and conformity requirements. The research paper presents a Privacy-Aware Cross-Domain Intelligence Framework that is a combination of federated learning and generative artificial intelligence that support collaboration between analytics without sensitive data leaving the domain of its origin. The framework provides domain management remote AI environments relating to a controlled federation layer, domains share knowledge through privacy conscious updating of models, policy execution, secure aggregation and authorization. Federated AI generative layer will allow authority domain models to provide expertise to the context and reduce unproblematic liberation of professional data, proprietary data and controlled information. The frameworks applied to enhance the level of trust and accountability include the differentiation in privacy, encryption, tracing provenance, model validation on the receiver and audit. Policy based orchestration layer It is a model sharing, dynamically controlled, participation, authorization, and restrictions on using data, sharing models, workflow management of heterogeneous enterprise environments. The suggested architecture will enhance the sovereignty of the organizational data, secure knowledge synthesis process, anomaly analysis, decision support and the cross-domain intelligence. Framework assessment ensures the privacy protection, effectiveness of communication, usefulness of model, sound governance, scalability and interoperability. The ensuing framework offers a systematic basis behind implementing generative AI within well-controlled, fractionated company systems where the intelligence circulation and privacy protection need to exist simultaneously..

KEYWORDS: Federated Learning, Generative AI, Privacy Preservation, Cross-Domain Intelligence, Regulated Enterprises, AI Governance.

I. INTRODUCTION

The intensive uptake of artificial intelligence (AI) in the medical field, finance, insurance, production, government services, and other industries with stringent regulation has opened opportunities to make data-driven decisions and intelligent robots. There is however an upsurge in pressure on organization in these areas to contain some direct sharing of information since organization information can contain a large proportion of personally identifiable, commercial sensitive, confidential and other information that is regulated. The classical centralized AI designs imply transfer of data to common stores to create models and to generate an inference that may have privacy, regulation, ownership and security concerns and organizational autonomy. The other paradigm is the federated learning (FL) whereby the parties involved do not share their information but instead participate in a shared machine-learning. This decentralized model has got a good ground on which to establish collaborative intelligence without necessarily going out of control by circulating sensitive enterprise information.

In particular, federated intelligence is particularly pertinent to cross-domain environments, where different organizations possess complementary information, which they are bound by different rules. With maintaining datasets, computing infrastructure, and organization-specific policies as locally controlled, a federated architecture can enable sharing of knowledge in participating domains. These form skills that are, in controlled industries, in which information can not simply be pooled into a centralized AI regime. Social service sphere, insurance industry and processes, and federated AI Systems have been explored and established that they can work together in the domains without losing organizational data boundaries [2]. Nevertheless, in this case as the enterprise AI becomes more complicated it demands the federation to the level of overcoming decentralized model training. Access control, implementation of policies, authorization, auditability, model validate, provenance and controlled information exchange should also be incorporated as part and parcel of the architecture.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 2, March – April 2025||

DOI:10.15662/IJRPETM.2025.0802008

The third dimension is another important facet of the issue which has emerged along with generative AI and foundation models. Expanding on Multimodal foundation models Multimodal generation Multimodal models Multimodal generation Multimodal foundation models or deep learning generative systems (also called large language models (LLM)) can generate summaries, recommendations, reports, explanations, explanations, reports and domain-specific responses of a wide range of complex enterprise data. Generative models may cover a wider range of privacy and governance needs in comparison with conventional predictive models, which may not be as engaging to the user and enterprise bodies of knowledge. The challenges of diverse data, processing-resource, effectiveness of communications, customisation, privacy and customisation of large models have been raised through the federated-foundation exploration between distributed players [3]. An independent initiative is attempting to jointly create or adapt generative models in a way that does not reveal proprietary datasets that such models are specific to, particularly where such alliances are considered a necessity.

Federated foundation models also give rise to different risks in comparison to those that come with the traditional federated learning. The Prompts, generated outputs, update of the models and the internal representations can offer information related to local training data or organizational knowledge. Also, the involved organizations might dissimilarly have regulatory requirements, security principles, infrastructure aptitude, and accepted applications of AI. Thus, it is not covered with credible cross domain intelligence due to a common aggregation. The problems that remain unaddressed by a literature on federated foundation models include privacy, security, heterogeneous participation, efficient training, model coordination, and a reliable deployment [4]. The observations ensure the architectures with sane technical federation mechanisms meshed with effectively developed governance and security controls in place.

Cross-domain cooperation also involves coordination that is easily scaled of geographically and administration spread computing environment. Experimental federated systems illustrate that decentralized learning can not only bring about coordination of several players, but can also locally store knowledge that can be applied in developing a larger intelligence potential [5]. Nonetheless, the larger scale of generative AI in enterprises presently poses significantly more calculational, communication, and operational needs as compared to most mainstream federated uses. The research parameters a model can have can be in the billions and the partnering organisations can vary widely in hardware and network structure, data nature and security guidelines. Consequently, a good architecture must reduce inappropriate data and parameter spreading data and permit simply organized models customization, policy controlled and domain arbitrary federation.

Facilitating privacy preservation is thus a key design objective of federated generative AI. Exposure to any kind is minimised with raw data in organization owned sites but locality alone does not negate privacy threat. The channels of information may additionally have possible additional sources, including model updates, gradient, prompts, generated responses and the result of retrieval. The privacy conscious architecture hence involves many complementary tools, perhaps these include, secure aggregation, encryption, differential privacy, access control, domain environment isolation and that of continuous auditing. These operations are to be carried out along with enterprise governance and not perceived as additional security operations.

The other aspect of the requirement is that there should be the capability of having domain autonomy and yet be capable of having meaningful cooperation. There exists a range of data structures, AI frameworks, policies, cloud services, and regulatory actions which can be applied by various enterprises. This is easily reduced by a strictly-centrated architecture, and difficult to a completely-decentralized architecture. An integrated system in generative-AI should therefore establish a controlled federation layer which can handle model transfer and collective intelligence but not critical information or core decision-making in any of the domains.

This article introduces Privacy-Aware Federated Generative AI Cross-Domain Intelligence Framework of Controlled Enterprise Systems in order to meet such requests. It is an integration of self-contained domain-level AI environments, and a regulated federation interface of regulated cooperation. It combines updates to their models which maintain privacy, secure aggregation, policy-based authorization, access control, provenance, model validation, auditing and orchestrated. The architecture will be designed in a manner that will allow organizations to share in creation of intelligence, which will span as a generation and sovereignty of data and domain specific governance will be maintained. The suggested framework will be structured to give an extensible and respondible foundation of cross-domain intelligences in comfortable privacy and regulated environments through a formation of federated learning and the notions of generative AI and business governance.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 2, March – April 2025||

DOI:10.15662/IJRPETM.2025.0802008

II. RELATED WORK

Federated learning (FL) has emerged as an important paradigm as a strategy of collaborative machine learning in the scenario, when the data of the organization could not be collected on a single platform. The primary concept of it is to retain training data in participating domains and exchange model parameters or updates with each other via an aggregation process. The strategy discusses the issue of data locality, privacy, communication-efficiency and independence to the organization and it is more than excellent in regards to the managed enterprise systems. Complicated analysis of FL unveils all the concerns of statistical heterogeneity, communication effectiveness, privacy, security, incentive machine and system size-setting the foundations of the decentralized intelligence across the various administrative areas [1]. Nonetheless, traditional FL is basically reduced to predictive structures, and fails to provide the full support of demands of a generator AI, multidomestically knowledge merging, and business management.

Use of federated intelligence to controlled and non-homogeneous organizational conditions has been getting more and more attention. The regulated industries map will focus on cross-domain federated AI into social services, insurance and industrial operations and propose reasons why it may be imperative to monitor boundaries of data in the organizations but permit cooperative intelligence [2]. These methods show that federation is also applicable to areas, where data exchanges are limited by privacy, regulatory, and operational issues. Lightweight policy enforcement mechanisms, authorization and provenance, model governance and controlled information exchange are however, not cost effective in cross-domain deployments, which FL architectures have traditionally had the ability to provide.

Federated learning has been greatly extended by the advent of foundation models and generative AI. In recent papers, the evolution of the foundation model, the classes of challenges in the field, including heterogeneous data, computational tasks, cost of communications, privacy, customization, and coordination among players are discussed (surveys) [3]. This proves more difficult when the large language models (as well as other generative models) are replicated in other locations of the enterprise. Specifically, the models of federated foundations take individual issues of heterogeneity of models, resource scarcity constraint, privacy leaking, knowledge integration and trustful collaboration with other providers into consideration, to show parametric aggregation, which in itself is insufficient in order to achieve complex federated generating intelligence [4].

The deployed federated applications are even another demonstration of how the collaborative intelligence can be utilized with no data gathering in a central place. FedSM proves that federated learning has the power to assist in the process of distributing decisions in a common mobility area and remaining local with the data involved [5]. Such domain-specific applications suggest that federation can plan the various members of an organization, but typically not centered around a cross-domain scheme of generic, controlled enterprise systems. The other general view of this multi-party computation maintenance strategy is the secure collaborative learning setting, and focuses on the cryptographic tools of ensuring distributed computation and information exchange [6]. The strategies are more privacy protecting, yet can result in a huge computational and coordination overhead when using large generative models.

The other direction that federated foundation models have been developing towards is parameter-efficient collaboration. Instead of having to share or even aggregate their models, PromptFL suggests conducting prompt learning in a collaborative manner, minimal information is needed to be communicated by the participating clients [7]. The orientation is especially applicable to the enterprise setting where the organizations might have varying data distributions, calculative resources, and policy choices. Problems of privacy, communication, and model personalisation, as well as computational cost and adaptation of large models in a decentralised environment are also stated by Federated large language model research [8]. These studies suggest that we need such needful procedures of efficient customization to allow integrative use of generative AI, within the parameters of an entity.

Heterogeneous low-rank adaptation to find a federated fine-tuning is also studied. FLoRA shows that parameter-fortunately adapted models can adapt to the atypicalities between clients who partake in it, and can jointly fine-tune large language models [9]. These methods are capable of minimizing communication and computational needs as opposed to full-model training. The idea is extrapolated to the concept of cross-cloud federated training to distributed cloud infrastructures, both geographically and administratively, in order to address the coordination and model-training stress to heterogeneous cloud environments [10]. The methods already establish beneficial forerunners to cross domain AI, albeit, they particularly give considerable thought to the joint training of models other than the management of the entire intelligence life cycle.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 2, March – April 2025||

DOI:10.15662/IJRPETM.2025.0802008

Recent studies have started to combine retrieval-augmented generation (RAG) with federated learning. An example of a use of distributed sources of knowledge in the context of generative systems is a federated RAG architecture of medical large language models that does not use a centralised access to sensitive medical data [11]. The federated instruction in unstructured text, likewise, shows that the collaboration in language model adaptation between organisations, grounded on locally stored textual data, is possible [12]. These have been particularly developmental in the instance of managed ventures whereby corporate experience is typically disseminated among hidden archives, data bases, documentation of activities and a repository of specialty. Nevertheless, retrieved information needs to be ensured to be kept within authorized limits and this necessitates explicit access control, provenance management and policy-conscious composition.

Federated generative AI is also expanded to adapter-based and multimodal methods. FedPIA explores how heterogeneous foundation-model fine-tuning using adapters can be integrated, and showed how modular parameter updates can facilitate such collaboration between various participants [13]. Chain-of-thought-oriented approaches and federated prompting investigate ways to enhance responses provided by LLM by means of distant prompts re-adaptation and rationale [14]. The second alternative that would result in federated LLM adaptation at a significantly lower cost is a parameter-efficient adaptive optimization prompt tuning that might be able to adapt to the variability on the set of different clients [15]. All these since demonstrate promise of prompts, adapters and updates that are lightweight parameters in a substitution of the traditional complete-model federation.

Regardless of these developments, current studies are still scattered in the fields of federated learning, foundation models, privacy-preserving computation, cross-cloud training, RAG, and parameter-efficient adaptation. The overwhelming majority of the systems apply one or two dimensions in lieu of providing a hybridized system of privacy sensitive cross-domain generative intelligence to controlled enterprises. In particular, there seems to be no single framework offering autonomous domain-level AI settings, which are controllably federated, share model exchange domains with privacy concerns, confident aggregation, policy implementation, generate knowledge authorizations, provenance, its validation, audit, controlled knowledge generators and controlled access to generative knowledge. This gap is filled by framing federated generative AI as part of a cross-domain governance-oriented architecture in a way that ensures the autonomy of organizations and allows for safe and responsible collaboration with intelligence.
Privacy-Preserving Generative Intelligence Across Domains Federated Governance Fabric.

Framework
The Federated Governance Fabric to Privacy-Preserving Cross-Domain Generative Intelligence is a platform that is capable of allowing the involved regulated businesses to join in the co-creation and co-sharing of the generative AI power, without privately accessing organizational confidential information. This architecture is a federated learning structure "plus" generative AI, privacy preserving computation, policy-centric governance plus secure model exchange, provenance management and controlled orchestration structured as a single architecture. The central notion in its principle is that the information remains within the realm in which it was created, and intelligence can be built through a federated process under being controlled.

The structure will be comprised of six rational layers which involve Domain Intelligence Layer, Privacy and Protection Layer, Federated Coordination Layer, Generative Intelligence Layer, Governance and Policy Layer and Enterprise Application Layer. These layers act on the basis of controlled interfaces as opposed to the unlimited flow of data. Each contributing organization maintains its own independent infrastructure, datasets, models, security controls and operational policies but delivers authorized updates in the form of model updates or derived intelligence to the federation.

III. FEDERATED GOVERNANCE FABRIC FOR PRIVACY-PRESERVING CROSS-DOMAIN GENERATIVE INTELLIGENCE

The Federated Governance Fabric to Privacy-Preserving Cross-Domain Generative Intelligence is a platform that is capable of allowing the involved regulated businesses to join in the co-creation and co-sharing of the generative AI power, without privately accessing organizational confidential information. This architecture is a federated learning structure "plus" generative AI, privacy preserving computation, policy-centric governance plus secure model exchange, provenance management and controlled orchestration structured as a single architecture. The central notion in its principle is that the information remains within the realm in which it was created, and intelligence can be built through a federated process under being controlled.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 2, March – April 2025||

DOI:10.15662/IJRPETM.2025.0802008

The structure will be comprised of six rational layers which involve Domain Intelligence Layer, Privacy and Protection Layer, Federated Coordination Layer, Generative Intelligence Layer, Governance and Policy Layer and Enterprise Application Layer. These layers act on the basis of controlled interfaces as opposed to the unlimited flow of data. Each contributing organization maintains its own independent infrastructure, datasets, models, security controls and operational policies but delivers authorized updates in the form of model updates or derived intelligence to the federation.

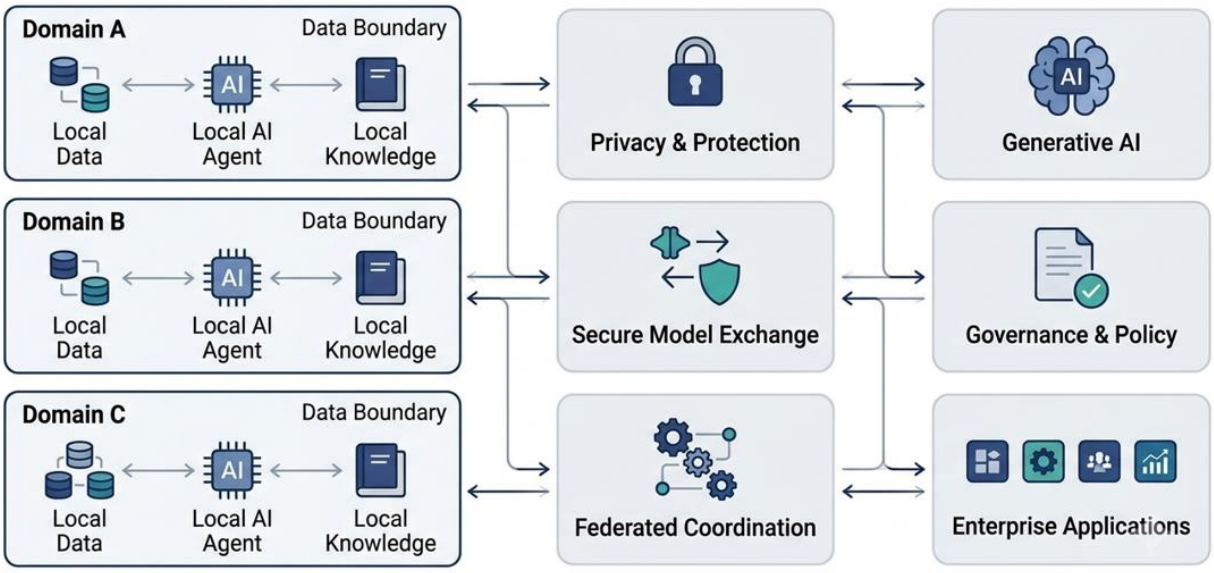


Figure 1- Privacy-Aware Cross-Domain Generative AI Architecture

3.1. Domain Intelligence Layer

The basic layer of the framework is the Domain Intelligence Layer. Each enterprise undergoes independent AI environments in both participating enterprises that comprise of local datasets of the enterprise, preprocessing services, domain models, computation resources and knowledge repositories. Examples include: Clinical data maintained by medical professionals, records of transactions maintained by financial institutions, claim records maintained by insurance companies or production data maintained by firms.

Raw enterprise data is not sent to other domains involved. The local data is processed within the organizational boundary of its security. Preparation elements of data processing include normalization, quality checks, feature transformation and policy-based filtering, followed by data being represented to local training or inference tasks.

All domains have a Local Generative AI Agent that interacts with locally permitted data and models. The agent can assist things such as summarization, personalization, document analysis, knowledge retrieval, decision support and contextual generation that has a domain specific character. Local agents liaise with the federation via controlled APIs, and none of them publicly reveal underlying enterprise repositories.

3.2. Privacy and Protection Layer

The Privacy and Protection Layer will ensure that information exchanged between the domains will not put sensitive enterprise knowledge in jeopardy unnecessarily. In contrast to sending raw-valued data, involved domains generate secure updates of models, parameter-sufficient adaptations, or embeddings or other verified representations.

International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 2, March – April 2025||

DOI:10.15662/IJRPETM.2025.0802008

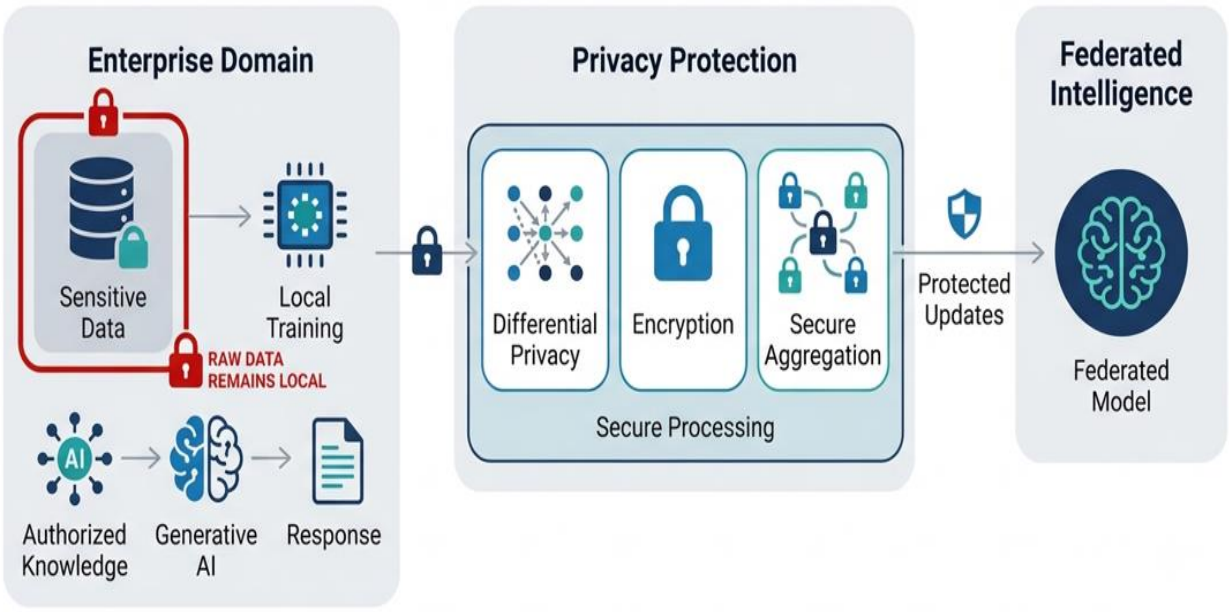


Figure 2- Privacy-Preserving Model and Knowledge Flow

The layer may accept the use of differential privacy, encryption, the secure aggregation mechanism, access control, minimization information and privacy-preserving mechanisms of transformation. Differential privacy can induce the protection where it is required and be controlled by statistics, but encryption will safeguard the information in flow and storage. Secure aggregation means that the federation coordinator cannot, themselves, inspect the individual participant updates to do so where it is required.

A Privacy Policy Engine decides whether information is allowed to exit a domain, what changes must be made, and what kind of federated contribution should be made. That way, the process of safeguarding privacy cannot be perceived as a technical one but as a policy-based one.

3.3. Federated Coordination Layer

Federated Coordination Layer is a mechanism to ensure that the framework is central in the coordination ability but not a central repository of enterprise information. It handles registration of participants, federated training rounds, exchange of model-updates, aggregate, version management, and federation status.

The domains participating are provided with an accepted global or foundation-model setup and conduct a local modification with its licensed datasets. Only allowable updates are relayed back to the coordination layer. The coordinator then aggregates these updates in an appropriate aggregation scheme, and synthesises a new federated model or common adaptation.

The scattered quality of data (heterogeneous participation) is ensured with a variety of data preservations across participating organizations and various computational abilities. There are other lightweight mechanisms that can be used to implement domains which do not transmit full foundation models, such as parameter-efficient fine-tuning or adapters and prompt-based adaptation and so on. This not only decreases the communication needs but also allows organizations with varying infrastructure abilities to take part.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 2, March – April 2025||

DOI:10.15662/IJRPETM.2025.0802008

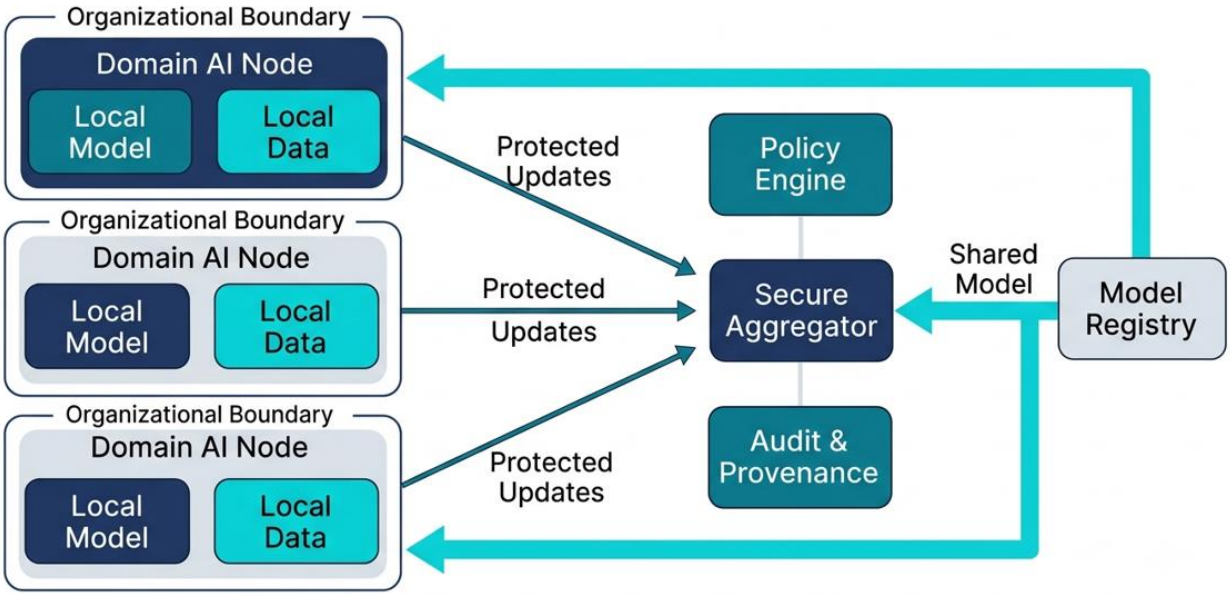


Figure 2- Federated Governance Fabric

There are also model versions and federation states that are kept by the coordination layer. Each model release may have metadata containing information on the cycle of federalities it was developed in, domains it is participating in, privacy policies it is executing, has been validated or not, and policies that it can be applied to.

3.4. Generative Intelligence Layer

The principle AI functions of the framework are provided by Generative Intelligence Layer. It can incorporate large models of language, multimodal foundation models, retrieval-augmented generation and domain-specific models, and federated adaptations.

Guided cross-domain knowledge use is one of the prominent aspects of this layer. A request is contrasted to authorization and domain policy in lieu of a generative model retrieving information of various enterprises freely. One can then route a query to one or more allowed domain agents and confidential information is within the respective organizational perimeter.

RAG features have been able to interoperate between local generative models and enterprise knowledge repository. Retrieval is on the authorized domain and only approved results or responses that are retrieved are relayed to the requesting application. The unproductive flow of confidential documents is minimised with such design but at the same time, the generative systems can use distributed organisational knowledge.

Federated prompt learning and parameter-effective model adaptation can also be supported by the layer. These methods enable the participating organizations to specialize a common underpinning paradigm, but do not need to synchronize

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 2, March – April 2025||

DOI:10.15662/IJRPETM.2025.0802008

the models completely.

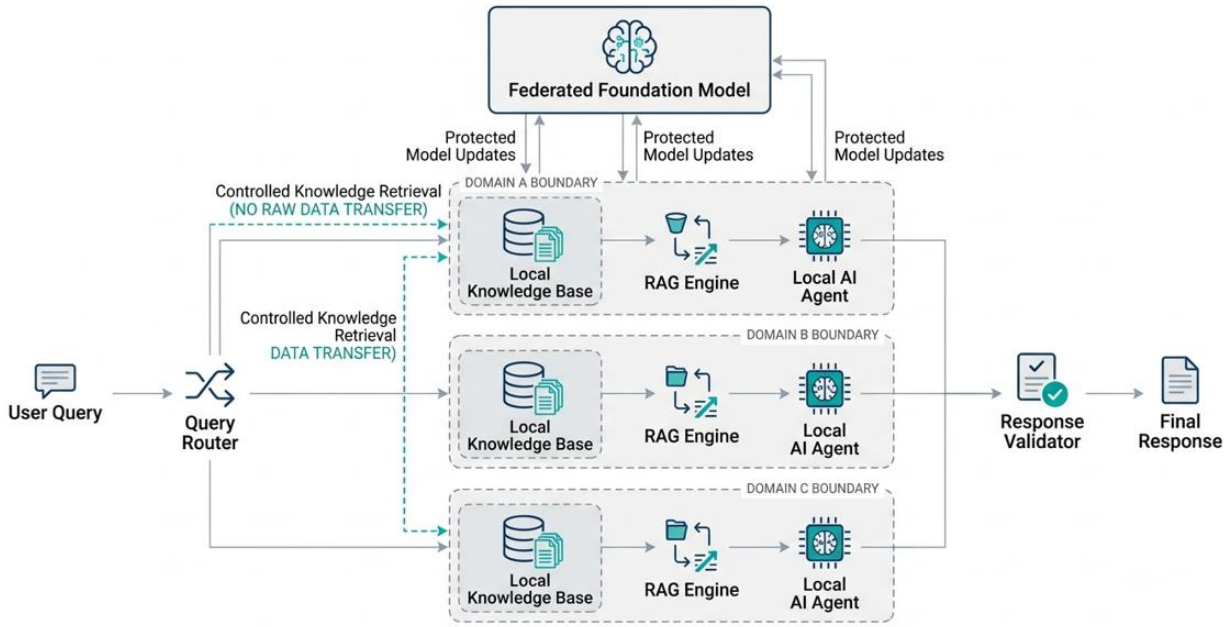


Figure 4- Cross-Domain RAG and Federated Intelligence

3.5. Governance and Policy Layer

The major control mechanism of the framework is the Governance and Policy Layer. It establishes rules on who is to take part; what shall be traded, what models shall be involved and what shall be under what circumstances the production of cross-domain intelligence shall be produced.

Policy engine would analyze request messages according to the organizational roles, data classifications, and regulation requirements, model permissions, and operations restrictions. They may be used prior to training, during sharing of models, during retrieval and during making of inferences.

Identity and access controls, authorization, policy enforcement, provenance, compliance monitoring, audit logging and model life cycle governance are other available features of the layer. Any significant operation can lead to an auditable event which could be model-update submission, aggregation, policy decision, retrieval request, model deployment and generated-output delivery.

Provenance schemes are database of provenance of models and learned wisdom. This allows organizations to be in the position to know what version of a model, a federation cycle, any approved domain, any retrieval process and governance policies was put into an output. It is especially in the regulated environments that AI decisions might be required to be examined once deployed that such traceability is significant.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 2, March – April 2025||

DOI:10.15662/IJRPETM.2025.0802008

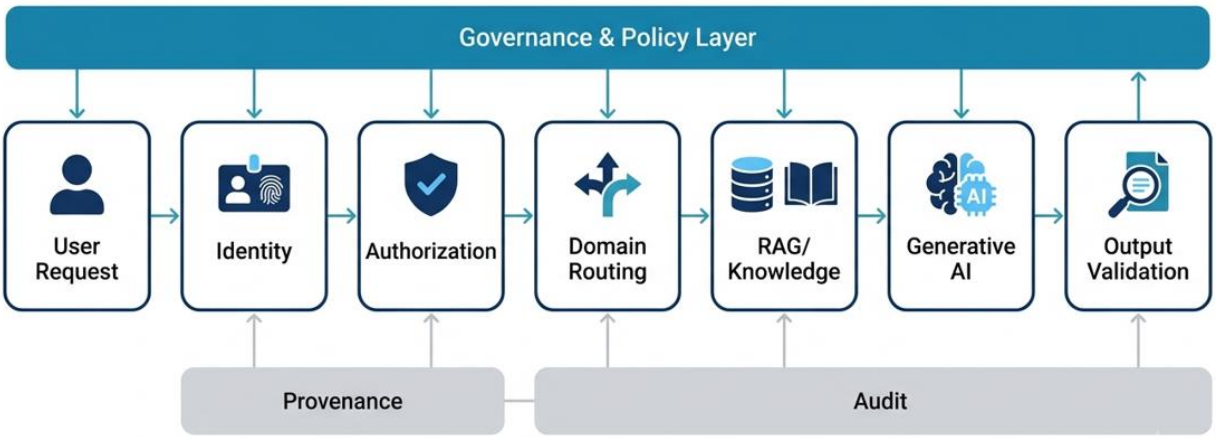


Figure 5- Policy-Aware Generative AI Workflow

3.6. Enterprise Application Layer

The resultant smartness is given out to the Enterprise Application Layer on controlled business applications. Its possible uses are risk analysis, fraud detection, clinical knowledge assistance, monitoring industrial enterprises, insurance processing, reporting regulations, enterprise search, and intelligent decision support.

Applications interoperate with the framework by accessing federated infrastructure via secure service interfaces, as opposed to accessing it directly. The request that comes in with a request to do it is checked and the governance layer sent and requested to the respective generative or federated intelligence service. Output and policy controls are then applied to this response and they are sent back.

This segregation provides different business applications, AI services, enterprise data, and federation infrastructure. It additionally enables organizations to make changes at the individual parts of it yet not to rearchitecture the whole architecture.

3.7. End-to-End Operational Workflow

It follows a controlled lifecycle that consists of registration, policy assessment, local planning, privacy assurance, federated dynamic, secure summarization, validation, controlled implementation, inference, tracking and auditing.

To begin with, a business offers domain registration and participation policies. Data on a local basis is then configured on a domain edge. Adaptation is done authorized by the local generative model or foundation model and the resultant update is privatized before being uploaded via the secure federation interface.

The federation coordinator approves the submission, and collects according to the strategy defined federation. Adapted or resultant model is put through framework verification that involves verification of privacy, model integrity, policy, and operational requirements. Approved versions are then allowed to be used by authorized domains.

Inference requires an enterprise application to make a request via the governance interface. Before reaching the generative intelligence junction, identity, permissions, and data-access requirements, among other corresponding policies, are given thought. Where the external knowledge domain is needed, the controlled federation mechanisms will only retrieve or generate the authorized information to the request. Before output is given there is policy and provenance of outputs..

3.8. Framework Characteristics

The proposed architecture is characterized by the data sovereignty, decentralized intelligence, policy-based federation, privatizing the exchange of model, the heterogeneous participation, modular generative AI, traceability and ongoing governance. Unlike architectures that consider federated learning a distributed training system, the given architecture considers federation as a whole enterprise governance system.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 2, March – April 2025||

DOI:10.15662/IJRPETM.2025.0802008

The very fabric of its design mission is then not to foil the dissemination of raw data, but to design a disciplined intelligence-sharing fabric upon which teamwork, privacy, security and organizational self-determination can co-exist harmoniously with the others. The resulting architecture enables regulated enterprises to harness generative AI across organizational boundaries and control sensitive data, model engagement, access controls, and AI governance in the lifecycle of intelligence..

IV. GOVERNANCE-CENTRIC VALIDATION OF FEDERATED GENERATIVE INTELLIGENCE

4.1. Evaluation Objectives

The framework is evaluated corresponding to the main demands of the controlled cross-domain AI landscape: privacy safeguarding, secure cooperation, generative intelligence, policy execution, limitations of scale, interoperability, and traceability. Ways to test the validity of the proposed architecture include knowing whether the architecture can allow collaborative model-building, directed use of knowledge without imposing on the participating entities the duty to reequip sensitive enterprise information. The federation layer, privacy mechanisms, and generative AI services, together with governance controls, are considered in particular.

4.2. Privacy Preservation and Data Sovereignty

Sovereignty of data is one of such criteria of assessment. The schema records the raw data at one-enterprise scope only, and communicates only approved model changes, modifications or deductions. Privacy policies such as the different privacy, encryption, secure aggregation and data-minimization provide different levels of protection. An architecture of this type presents minimal exposure of enterprise datasets, and also allows distributed organizations to engage in the development of shared models. Exposure to updates, unauthorized attempt, policy violation and the effectiveness of protection mechanisms can also be considered in privacy evaluation.

4.3. Federated Collaboration Efficiency

Volume of communication, payload to aggregate overhead, rate of model update and heterogeneity of participants are used to measure the federated collaboration. The most common mechanism to reduce the information that is exchanged between the two sides between synchronising the entire model is parameter-efficient prompts and adaptive mechanisms. The infrastructure is also capable of supporting the domains of different computation fragment pieces and also the distribution of information. Federation-round completion time, convergence behavior, rate of participation by the participants and rate of the communication requirement that is needed by various number of domains can thus be tested.

4.4. Generative Intelligence and Knowledge Integration

The framework is viewed within the context of its ability to help generative AI in a distributed enterprise knowledge climate. Federative model adaptations allow knowledge gained in a certain domain to enhance the intelligence of other participating domains, but it requires local generative agent processing of a domain-specific information. Components based on RAG provide an additional system of accessing knowledge which is locally enforced in the enterprise, though without creating a central repository of knowledge. Relevance of the responses to the domain of use, domain consistency, retrieval accuracy, control of hallucinations and being able to enforce appropriate boundaries between domain specific and shared intelligence are criterion that will be used to do assessment.

4.5. Governance and Policy Compliance

The capabilities of the framework to effect identity, authorisation, and access controls, privacy measures and model-use conditions are all measured in the lifecycle of AI. Policies are implemented in the registration of the participants, exchange of model updates, retrieval of knowledge, inference and deployment. Provenance metadata and audit history provide traceability of significant operations. The policy-decision latency, unauthorized request rejection, audit completeness, provenance coverage as well as even-handed application of the policy through the various regions of coverage can be measured.

4.6. Security and Trustworthiness

Security test considers the vulnerability of the architecture to unauthorized access, malicious updates and privacy leakage and compromised participants. The security of communication, authenticated federation point, and specialization of the aggregation and confinement of the domains and ongoing monitoring are involved in the defense-in-depth. Before anomalous or non-compliant contributions can have an effect on the shared intelligence layer, integrity

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 2, March – April 2025||

DOI:10.15662/IJRPETM.2025.0802008

checks and validation can detect the anomalies. Other types of ensuring trustworthiness are version control, provenance tracking and controlled model deployment.

4.7. Scalability and Interoperability

The framework can be scaled to provide capacity to include additional domains, models, applications and knowledge repositories. Scalability testing could be to test federation performance with increased workloads of participants and model-updates. Interoperability is gauged by standardized APIs, modular services, model interfaces and policy interfaces traversing across the heterogeneous enterprise infrastructures without replacing the systems utilized by the enterprise.

V. DISCUSSION

As proposed framework demonstrates, the architectural perspective of privacy conscious generative AI in regulated businesses needs to be bigger in comparison to its traditional counterpart, federated learning. The classic federation restricts itself primarily to training of decentralized models but cross-domain generative intelligence raises more issues connected to access to knowledge, governing the model, authorisation, provenance and controlled generation of output. The building will react to these needs by decoupling enterprise data environments and the federation and governance infrastructure.

Among the key implications, it is that collaborative intelligence and data sovereignty need not always imply a tradeoff between the two. Companies-organizations can have a chance to store sensitive datasets in their circles, and deliver regulated amendments to designs or approved smarts to a shared ecosystem. The design is especially applicable to the situations when the formal constraints (regulatory), contractual, or organizational restrict the direct data exchange. In the meantime, local data retention is not capable of working on privacy by itself. Model updates can be potential sources of information, as well as the results of retrieval, prompts and generated output, and therefore privacy mechanisms need a way to be supported through the entire intelligence lifecycle.

The framework also has policy sensitive orchestration. Various businesses might have varying demands when it comes to classifying data, granting users access, using a model, retention, and inter-domain access. This can be reflected in policy enforcement embedded within federation workflows enabling the explicit representation of such differences rather than by discrete administrative processes. Accountability is also supported by provenance and audit capabilities as the versions of the models can be traced, federation operations and retrieval operations can be traced, as well as the generated intelligence can be traced.

Another thing to consider is scaling. Energy-intensive foundation models can be a challenge to communication and computation, and parameter-efficient adaptation, modular agents, and controlled knowledge retrieval can be viewed as effective architecturally. However, this does not exclude the possibility of additional domains involvement to introduce overheads of coordination and non-homogeneous security requirements.

In general, the discussion demonstrates that federated generative AI has the potential to become a coordinated intelligence fabric in which privacy, collaboration, generative, capability and organizational autonomy can be integrated by design. Future research needs to investigate even stronger guarantees of privacy, dynamic policy regulation, powerful participant authentication, multi-domain provenance, efficient multimodal federation and ongoing deployment of governance of increasingly autonomous enterprise AI systems.

VI. CONCLUSION

This research proposed Privacy-Conscious Cross-Domain Intelligence Framework and Federated Generative AI to controlled enterprise system, which encompasses the heightened necessity to possess collaborative artificial intelligence without revealing delicate information about the organization. The proposed framework merges the concepts of federated learning, generative AI, privacy preserving, secure model exchange, policy based governance, provenance, authorization, and continuous validation together in a single architecture.

The framework will set up independent domain intelligence environments in which enterprise information is owned by the source organization. The domains participating can organize a federative communication to update shared foundation models by jointly sharing updates to models, and also parameter adjustment in an efficient manner. Privacy

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 2, March – April 2025||

DOI:10.15662/IJRPETM.2025.0802008

and protection systems also offer further protection to information which is shared in the federation and secure aggregation and encryption minimize such exposure in collaborative processing. It is a mix, which provides a systematic means of claiming data sovereignty and permits cross-domain intelligence.

The framework goes further to be broader than the traditional federated learning by integrating governance in the AI lifecycle. Controlling participation, model exchange, retrieval, inference, and deployment are offered by policy enforcement, identity management, access control, provenance tracking, audit logging and model validation. It unites the local generative agents as well as retrieval-based access to knowledge at any point of time in the circumstance of domain-specific intelligence which is being upheld within parameters of an organization. Consequently, one does not even have to tap into distributed knowledge stakes on a free basis and generative AI is capable of operating within heterogeneous business environments.

The framework evaluation demonstrates that privacy conscious federation, effective adaptation of models, monitored integration of knowledge, and governance can contribute towards the creation of trustful cross domain generative intelligence together. However, concerns of overhead on communication, diverse infrastructure, privacy as model updates, malicious participation, expensive computations on large models and evolving regulatory responsibilities still exist.

Future research could focus on stronger privacy guarantees, adaptive secure aggregation, long-lasting validating participants, multimodal federated foundation model, decentralized provenance, autonomous operations with policy sensitivity, and continuously seeking compliance. Additional validation of the framework can be conducted in the operation under varying conditions in medical work, financial, insurance, manufacturing, and the government. In most cases, the recommended architecture provides a foundation of developing the concept of federated learning as an extension of the uncoordinated and decentralized training of models to a managed fabric of safe, responsible and privacy considerate enterprise AI teamwork..

REFERENCES

[1] P. Kairouz *et al*., "Advances and open problems in federated learning," *Foundations and Trends® in Machine Learning*, vol. 14, nos. 1–2, pp. 1–210, 2021, doi: 10.1561/2200000083.
[2] Dave, B. L. (2023). Federated AI frameworks for regulated industries: Cross-domain intelligence for social services, insurance, and industrial operations. International Journal of Research and Applied Innovations, 6(1), 8346-8362.
[3] C. Ren *et al*., "Advances and open challenges in federated foundation models," *IEEE Communications Surveys & Tutorials*, 2025.
[4] T. Fan *et al*., "Ten challenging problems in federated foundation models," *IEEE Transactions on Knowledge and Data Engineering*, 2025.
[5] S. Wei, Y. Zhang, Z. Zhou, T. Zhang, and K. Xu, "FedSM: A practical federated shared mobility system," *Proceedings of the VLDB Endowment*, vol. 17, no. 12, pp. 4445–4448, 2024, doi: 10.14778/3685800.3685896.
[6] F. Liu, Z. Zheng, Y. Shi, Y. Tong, and Y. Zhang, "A survey on federated learning: A perspective from multi-party computation," *Frontiers of Computer Science*, vol. 18, no. 1, Art. no. 181336, 2024, doi: 10.1007/s11704-023-3282-7.
[7] T. Guo, S. Guo, J. Wang, X. Tang, and W. Xu, "PromptFL: Let federated participants cooperatively learn prompts instead of models—Federated learning in age of foundation model," *IEEE Transactions on Mobile Computing*, vol. 23, no. 5, pp. 5179–5194, 2024, doi: 10.1109/TMC.2023.3302410.
[8] J. Hu, D. Wang, Z. Wang, X. Pang, H. Xu, J. Ren, and K. Ren, "Federated large language model: Solutions, challenges and future directions," *IEEE Wireless Communications*, 2024.
[9] Z. Wang, Z. Shen, Y. He, G. Sun, H. Wang, L. Lyu, and A. Li, "FLoRA: Federated fine-tuning large language models with heterogeneous low-rank adaptations," in *Proc. 38th Int. Conf. Neural Information Processing Systems (NeurIPS)*, 2024, pp. 22513–22533.
[10] H. Yang, M. Sui, S. Liu, X. Qian, Z. Zhang, and B. Liu, "Research on key technologies for cross-cloud federated training of large language models," *Academic Journal of Computing & Information Science*, vol. 7, no. 11, pp. 42–49, 2024.
[11] J. Jung, H. Jeong, and E. N. Huh, "Federated learning and RAG integration: A scalable approach for medical large language models," in *Proc. 2025 Int. Conf. Artificial Intelligence in Information and Communication (ICAIIC)*, 2025, pp. 968–973.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 2, March – April 2025||

DOI:10.15662/IJRPETM.2025.0802008

[12] R. Ye, R. Ge, F. Yuchi, J. Chai, Y. Wang, and S. Chen, "Leveraging unstructured text data for federated instruction tuning of large language models," in *Federated Learning in the Age of Foundation Models—FL 2024 Int. Workshops*, 2025, pp. 119–131.

[13] P. Saha, D. Mishra, F. Wagner, K. Kamnitsas, and J. A. Noble, "FedPIA—Permuting and integrating adapters leveraging Wasserstein barycenters for finetuning foundation models in multi-modal federated learning," in *Proc. 39th AAAI Conf. Artificial Intelligence*, 2025, pp. 20228–20236.

[14] X. Liu, T. Pang, and C. Fan, "Federated prompting and chain-of-thought reasoning for improving LLMs answering," in *Proc. 16th Int. Conf. Knowledge Science, Engineering and Management*, 2023, pp. 3–11.

[15] T. Che, J. Liu, Y. Zhou, J. Ren, J. Zhou, V. Sheng, H. Dai, and D. Dou, "Federated learning of large language models with parameter-efficient prompt tuning and adaptive optimization," in *Proc. 2023 Conf. Empirical Methods in Natural Language Processing (EMNLP)*, 2023, pp. 7871–7888, doi: 10.18653/v1/2023.emnlp-main.488.