

Federated Agentic AI: A Privacy-Preserving Architecture for Collaborative Intelligence Across Distributed Enterprises

Ashkan Bhosale ¹, Iqrar Farooq ¹

¹ University of Surrey, UNITED KINGDOM

Abstract

Enterprises increasingly hold data and operational capability that is more valuable in combination than in isolation a manufacturer and its tier-one suppliers each see only part of a demand shock, a group of regional banks each see only their own slice of a cross-institution fraud ring, a hospital network and its referral partners each hold only a fragment of a patient's full clinical picture — yet competitive concerns, trade secrecy, and statutes governing data sharing routinely make centralizing that data into one system legally or commercially impossible. At the same time, 2023 has seen the rapid emergence of agentic artificial intelligence: large-language-model-based systems that do not merely answer a single prompt but autonomously plan multi-step tasks, invoke external tools, observe results, and revise their own plans. Almost every agentic AI system described in the literature to date assumes a single organization's data, tools, and oversight; the moment two or more independently governed enterprises want their agents to collaborate, existing architectures offer no answer that does not either centralize sensitive data or abandon the autonomous, multi-step reasoning that makes agentic AI valuable in the first place. This paper addresses that gap. Synthesizing the federated learning literature, the privacy-preserving computation literature, and the emerging agentic-AI literature, it develops Federated Agentic AI (FAA): a reference architecture in which each enterprise runs its own language-model agent locally over its own data and tools, and cross-enterprise “collaborative intelligence” occurs only through a privacy-preserving exchange layer secure-aggregated model updates, differentially private statistics, or sanitized intermediate task representations that never transmits raw records and never grants one enterprise's agent direct access to another's systems. Because no production deployment of such a system exists to evaluate, this paper adopts a design-science approach: the architecture is built from and evaluated against peer-reviewed findings, and every quantitative comparison presented — two comparative tables and four figures spanning capability coverage, privacy-utility trade-offs, coordination overhead, and task-success-rate distributions is explicitly labeled as an , literature-grounded estimate rather than a measured result. The paper's contribution is a synthesis of three otherwise separate literatures into one architecture, an explicit mapping of collaborative-intelligence objectives onto specific privacy-preserving and agentic techniques, and a candid account of the coordination-reliability and privacy-leakage limitations any adopter must weigh before deployment.

Keywords: Federated Learning, Agentic AI, Large Language Model Agents, Privacy-Preserving Computation, Multi-Agent Systems, Enterprise AI Collaboration

Introduction

Consider three enterprises, each with a genuine incentive to collaborate and an equally genuine reason not to pool their data. A vehicle manufacturer and its tier-one and tier-two suppliers would each forecast demand and disruption risk more accurately if they could see one another's order books, inventory positions, and production schedules — but no supplier will hand its cost structure and capacity constraints to an OEM that also negotiates its prices, and antitrust counsel would have a great deal to say about the reverse. A consortium of regional banks would catch more cross-institution fraud rings if each bank's transaction-monitoring system could see the others' flagged accounts — but bank secrecy law, competitive concerns, and the sheer liability of centralizing another institution's customer data make a shared database a non-starter. A network of hospitals and their referral partners would manage complex patients better with a unified clinical picture — but HIPAA-governed data-sharing agreements are negotiated institution by institution, and a shared patient-data warehouse spanning independently governed health systems is, in most jurisdictions, either explicitly prohibited or practically unworkable. None of these constraints is new; they are the ordinary condition under which most valuable enterprise data actually sits (Mallempati, 2023).

What is new, in 2023, is a class of artificial intelligence system capable of doing something with that data no static analytics dashboard could: agentic AI. Where a conventional large language model answers a single prompt, an agentic system plans a multi-step task, invokes external tools and APIs to gather information or take action, observes the results, and revises its plan accordingly. Wei *et al.* (2022) showed that prompting a sufficiently large language model to produce intermediate reasoning steps — chain-of-thought — substantially improves its performance on multi-step problems, establishing the reasoning capability agentic systems build on. Yao *et al.* (2023) then showed, with the ReAct framework, that interleaving that reasoning with concrete actions and observations from an external environment lets a language model ground its plans in real feedback rather than reasoning open-loop, while Schick *et al.* (2023) demonstrated with Toolformer that a model can learn, largely by itself, when and how to invoke external tools such as calculators, search engines, or APIs. Park *et al.* (2023) extended this further with generative agents that maintain persistent memory and periodically reflect on it, letting an agent operate coherently across long-running, multi-session tasks rather than a single interaction (Potharaju & Ravipudi, 2023). Enterprises have moved quickly to apply these capabilities internally a customer-service agent that plans a multi-step resolution and calls internal systems, a research agent that plans a literature search and synthesizes findings, a logistics agent that replans a delivery schedule as conditions change (Potharaju & Ravipudi, 2023).

Every one of those internal deployments, however, shares an assumption this paper argues is about to become a limitation: that the agent's data, tools, and oversight all belong to one organization. Li *et al.* (2023) showed with CAMEL that two language-model agents can hold a structured, role-played conversation to collaborate on a complex task a genuine advance in multi-agent coordination but CAMEL's agents share a common environment and a common operator; nothing in that framework, or in the agentic-AI literature more broadly, addresses what happens when the two collaborating agents belong to enterprises that are, by design, not supposed to see each other's underlying data or systems. The federated learning literature solved an analogous problem for a narrower kind of collaboration: McMahan *et al.* (2017) showed that multiple parties can jointly train a shared model by exchanging only model updates, never raw data, and a substantial body of

subsequent work surveyed comprehensively by Kairouz *et al.* (2021), Yang *et al.* (2019), and Li *et al.* (2020) has extended this to handle the statistical and systems heterogeneity real organizations bring to a federation. But federated learning's unit of collaboration is a fixed training objective exchanged as model parameters; it has no native concept of an autonomous agent that plans a task, calls a tool, and adapts mid-course. The two literatures, in other words, have each solved half of the problem this paper addresses, and neither has been extended to cover the other half (Mallempati & Jaladi, 2023).

This paper's contribution is to close that gap directly. It proposes Federated Agentic AI (FAA), a reference architecture synthesizing the agentic-AI, federated-learning, and privacy-preserving-computation literatures, in which each enterprise runs its own language-model agent locally — with its own data, its own tools, and its own oversight — and cross-enterprise collaboration happens exclusively through a privacy-preserving exchange layer that passes secure-aggregated model updates, differentially private statistics, or sanitized intermediate task representations, never raw records and never unrestricted tool access into another enterprise's systems. This is not a merely cautious design choice: Zhu *et al.* (2019) demonstrated that even ordinary gradient updates in federated training can leak enough information to reconstruct a party's underlying training data through optimization, so any federated collaboration layer — for model training or, as this paper argues, for agentic task coordination — must incorporate deliberate protection rather than assume that exchanging “just updates” is automatically safe. The paper makes three contributions: first, a synthesis of the agentic-AI, federated-learning, and privacy-preserving-computation literatures into a single architecture organized around cross-enterprise collaboration specifically; second, an explicit mapping of collaborative-intelligence objectives onto concrete techniques and institutional controls; and third, because no production system of this kind exists to measure, a transparent methodology for constructing evaluation figures traceable to the directional findings of the cited literature rather than invented outright (Rashid & Filgueirus, 2022).

The remainder of the paper proceeds as follows. Section 2 reviews federated learning and its privacy-preserving mechanisms, agentic and multi-agent language-model systems, and the gap between them. Section 3 describes the design-science approach, how the FAA framework was constructed, and how Section 4's figures were built. Section 4 presents the proposed architecture, two comparative tables, an cross-enterprise case scenario, and four figures spanning capability coverage, privacy-utility trade-offs, coordination overhead, and task-success distributions. Section 5 discusses implications, limitations, what the paper does not argue, and future research. Section 6 concludes.

Literature Review

Federated Learning: Foundations and the Privacy Question It Raises

Federated learning began from a comparatively narrow motivation — training a model across many mobile devices without moving raw user data off the device — but its core mechanism generalizes directly to the enterprise setting this paper addresses. McMahan *et al.* (2017) introduced Federated Averaging, in which each participant trains a local model on its own data for several steps, sends only the resulting model update to a coordinating server, and the server averages updates across participants to produce an improved shared model, iterating until convergence; no participant's raw

data ever leaves its own infrastructure. Yang *et al.* (2019) formalized a taxonomy that matters directly for this paper’s architecture, distinguishing cross-device federation — many resource-constrained clients, typically consumer devices — from cross-silo federation, in which a comparatively small number of well-resourced, persistently available organizations collaborate, precisely the setting a manufacturer-supplier consortium or a bank group occupies. Kairouz *et al.* (2021) surveyed the field comprehensively and catalogued a set of open problems directly relevant here: statistical heterogeneity (each participant’s data distribution differs, sometimes substantially, from every other’s), systems heterogeneity (participants vary in compute, connectivity, and availability), and the tension between personalizing a shared model to each participant and maintaining one model useful to all. Li *et al.* (2020) examined the statistical heterogeneity problem in particular, showing that the naive Federated Averaging algorithm can converge slowly or to a poor solution when participants’ data is non-identically distributed — precisely the condition among enterprises whose customer bases, product mixes, or patient populations differ meaningfully, and a challenge this paper’s architecture must inherit and address rather than assume away (Rashid & Junck, 2023; Siddiqui *et al.*, 2023).

Federated learning’s foundational promise — collaboration without centralizing raw data — turns out to require more scrutiny than exchanging model updates alone provides. Zhu *et al.* (2019) showed that an adversary with access to a participant’s shared gradients can, through an optimization procedure, reconstruct inputs strikingly close to that participant’s actual training data, a result now widely cited as evidence that “the update isn’t the raw data” is not by itself a privacy guarantee. Bonawitz *et al.* (2017) developed a cryptographic response: a secure aggregation protocol in which a coordinating server can compute the sum of participants’ updates without ever seeing any individual participant’s update in the clear, closing off the specific leakage vector Zhu *et al.* (2019) later demonstrated against unprotected updates. This pairing — a foundational algorithm (McMahan *et al.*, 2017) and the discovery that its naive form leaks (Zhu *et al.*, 2019), followed by a cryptographic defense (Bonawitz *et al.*, 2017) — is the pattern this paper’s own architecture takes as a design principle: any layer that exchanges information across an organizational boundary should be assumed leaky until a specific mechanism closes the leak, not assumed safe because the information exchanged is not, in name, “the data” (Potharaju & Sajja, 2021).

Privacy-Preserving Computation Layered onto Federated Collaboration

Secure aggregation closes one leakage vector but not every one, and a further body of work layers additional guarantees onto federated collaboration. Abadi *et al.* (2016) showed that differential privacy — a formal, quantifiable bound on how much a computation’s output can depend on any single record — can be incorporated directly into model training via differentially private stochastic gradient descent, providing a mathematical guarantee against reconstruction attacks like Zhu *et al.*’s (2019) rather than relying solely on the “honest but curious” trust assumption secure aggregation requires of the coordinating server. Truex *et al.* (2019) combined this idea with secure multiparty computation directly, proposing a hybrid approach in which the noise required for a differential-privacy guarantee is distributed across participants using cryptographic techniques, showing that the hybrid approach achieves a given privacy guarantee with less accuracy loss than differential privacy applied by any single party alone — a finding this paper’s architecture treats as

the template for combining privacy mechanisms rather than choosing exactly one. Fang and Qian (2021) proposed a complementary mechanism, homomorphic encryption, which allows a coordinating party to perform the aggregation computation directly on encrypted updates without ever decrypting them, removing the “honest but curious server” assumption secure aggregation and differential privacy both still partly rely on, at a real computational cost the paper quantifies. Read together, these three studies establish that no single privacy-preserving mechanism is sufficient for every threat model and trust arrangement; a genuinely cross-enterprise architecture needs a menu of mechanisms whose cost and protection level can be matched to how much the participating enterprises actually trust one another and the coordinating layer between them (Potharaju & Addanki, 2022).

Cross-silo federated learning already has a concrete, regulated precedent this paper draws on directly. Xu *et al.* (2021) reviewed federated learning applications in healthcare informatics, documenting deployments in which hospitals and health systems — independently governed, subject to HIPAA and equivalent statutes, and structurally unable to pool patient records — jointly improved diagnostic and predictive models via exactly the cross-silo federated pattern Yang *et al.* (2019) formalized. Xu *et al.* (2021) also catalogued the practical governance arrangements such collaborations required: data-use agreements specific to the federation, agreed evaluation protocols so no participant can infer another’s data from how the shared model performs on it, and infrastructure for participants with meaningfully different data volumes and computational resources. This is the closest existing analogue to the manufacturer-supplier or bank-consortium scenarios motivating this paper, and it establishes that cross-enterprise federated collaboration is not merely theoretically possible but has already been operationalized in at least one highly regulated sector — with the caveat, central to this paper’s contribution, that Xu *et al.*’s (2021) review concerns model training exclusively, not the autonomous, multi-step, tool-using collaboration agentic AI introduces.

Agentic AI: Autonomous, Tool-Using Language Model Agents

The agentic-AI literature reviewed briefly in Section 1 deserves closer examination for what it does, and does not, address about collaboration across a trust boundary. Wei *et al.* (2022) established that chain-of-thought prompting elicits substantially improved multi-step reasoning from large language models, the capability every subsequent agentic framework builds on. Yao *et al.* (2023) extended this from reasoning alone to reasoning interleaved with action: their ReAct framework has a language model alternate between a reasoning trace and a concrete action — a tool call, a search query, an API invocation — observing the environment’s response before continuing, which they show substantially reduces the compounding errors that pure open-loop reasoning produces on tasks requiring real-world information. Schick *et al.* (2023) showed with Toolformer that a language model can learn largely unsupervised when a given point in its own generated text would benefit from invoking an external tool, and how to format that invocation, reducing the need for hand-engineered tool-use prompting. Park *et al.* (2023) addressed a different limitation: agents that only reason within a single context window cannot sustain coherent behavior across long-running tasks, so their generative-agent architecture adds persistent memory and periodic reflection, letting an agent recall and reason over its own history rather than starting fresh each session.

Multi-agent coordination — more than one language-model agent working on a shared task — is the piece of this literature closest to what Federated Agentic AI requires, and also the piece that most clearly stops short of it. Li *et al.* (2023) proposed CAMEL, in which two language-model agents are assigned complementary roles and communicate through structured natural-language dialogue to jointly complete a task neither could complete as well alone, demonstrating that role-conditioned agent-to-agent communication is a workable coordination mechanism without a human in the loop for every exchange. But CAMEL’s agents operate within one shared environment under one operator’s oversight; the framework has no concept of a privacy boundary between the two agents, no mechanism for one agent to withhold the underlying data behind a statement it makes to the other, and no treatment of what happens when the two agents’ operators do not fully trust one another. This is not a shortcoming of CAMEL specifically — it was not designed to solve a cross-organizational trust problem — but it illustrates precisely the gap this paper’s architecture must fill: multi-agent communication protocols exist and work, but they assume away the exact constraint that makes cross-enterprise collaboration hard in the first place.

The Missing Layer: Coordinating Autonomous Agents Without Centralizing Trust

Sections 2.1 through 2.3 can be read as two literatures that have each built half of what cross-enterprise collaborative intelligence requires. Federated learning (Section 2.1) and its privacy-preserving extensions (Section 2.2) have built a rigorous answer to how independently governed organizations can jointly improve a shared model without centralizing raw data or unconditionally trusting a coordinator — but their unit of exchange is a model update toward a fixed, jointly agreed training objective, not a task plan, a tool-use decision, or an intermediate reasoning step. Agentic AI (Section 2.3) has built increasingly capable autonomous, tool-using, multi-step reasoning agents, including early multi-agent coordination protocols such as CAMEL — but every one of these systems assumes a single trust domain, with no native mechanism for an agent to collaborate usefully with another agent while withholding the data underlying its own reasoning (Karim & Elias, 2022).

Yang *et al.*’s (2019) cross-silo/cross-device distinction and Xu *et al.*’s (2021) healthcare federation precedent come closest to bridging this gap, because they establish that independently governed enterprises already do collaborate through a federated, privacy-preserving layer in production settings — but strictly for model training, never for the kind of autonomous, tool-invoking, multi-step task execution Yao *et al.* (2023), Schick *et al.* (2023), and Park *et al.* (2023) describe. No study reviewed here proposes extending the federated exchange layer to carry the artifacts an agentic system actually needs to coordinate — task plans, intermediate representations, negotiated allocations of sub-tasks — under the same privacy-preserving discipline federated learning applies to model updates. That is the specific, addressable gap this paper’s architecture is built to close.

Synthesis and the Gap Addressed by This Paper

Read together, these literatures converge on a clear division of labor this paper’s contribution sits directly between. Federated learning and privacy-preserving computation (Sections 2.1–2.2) supply the trust model: secure aggregation (Bonawitz *et al.*, 2017), differential privacy (Abadi *et al.*, 2016), hybrid cryptographic-statistical approaches (Truex *et al.*, 2019), and homomorphic encryption (Fang & Qian, 2021) each define a way to collaborate on a shared numerical objective without

centralizing raw data, at different points on a cost-versus-protection curve, with cross-silo enterprise precedent already established in healthcare (Xu *et al.*, 2021). Agentic AI (Section 2.3) supplies the reasoning model: chain-of-thought reasoning (Wei *et al.*, 2022), tool-grounded action (Yao *et al.*, 2023; Schick *et al.*, 2023), persistent memory (Park *et al.*, 2023), and structured multi-agent communication (Li *et al.*, 2023) together describe how an autonomous system plans, acts, and coordinates with another agent. Neither literature, on its own, describes how an autonomous agent belonging to one enterprise can plan and act jointly with an autonomous agent belonging to another enterprise without either centralizing their data or abandoning the privacy-preserving discipline federated learning has already worked out for the narrower case of model training (Karim & Seitz, 2023). This paper's architecture, developed in Sections 3 and 4, is built specifically to close that gap.

Methodology

Research Approach: Design Science with Evaluation

This paper adopts a design-science research approach rather than an empirical one. The central deliverable is an artifact — a reference architecture and an objective-to-technique mapping for Federated Agentic AI — whose value is justified by grounding each design choice explicitly in prior peer-reviewed findings, rather than by measuring performance against a live cross-enterprise deployment. This choice reflects a fact stated plainly rather than obscured: no production Federated Agentic AI system exists at the time of writing, so no enterprise data, agent logs, or coordination traces were available to evaluate empirically, and claiming otherwise would misrepresent what was actually done. Every quantitative comparison in Section 4 — two tables and four figures — is explicitly and repeatedly labeled as an construction consistent with directional findings in the cited literature, not measured output from an operational system.

Development of the Federated Agentic AI Framework

The architecture in Section 4.1 was developed through a structured synthesis of the three literatures reviewed in Section 2. First, the agentic-AI literature (Section 2.3) identified the components a capable local agent requires: a reasoning-and-action loop (Yao *et al.*, 2023), self-directed tool use (Schick *et al.*, 2023), and persistent memory for multi-session tasks (Park *et al.*, 2023). Second, the federated-learning and privacy-preserving-computation literatures (Sections 2.1–2.2) identified the components a cross-enterprise exchange layer requires: an aggregation mechanism that never exposes an individual participant's raw update (Bonawitz *et al.*, 2017), a formal privacy guarantee bounding what any shared artifact can reveal (Abadi *et al.*, 2016), and, for higher-mistrust settings, cryptographic alternatives that relax the aggregator's trust requirements further (Truex *et al.*, 2019; Fang & Qian, 2021). Third, following Li *et al.*'s (2023) demonstration that structured, role-conditioned dialogue is a workable coordination mechanism between autonomous agents, this paper extends that mechanism across the privacy-preserving exchange layer rather than within one shared environment — an agent proposes a sanitized, privacy-bounded summary of its local finding or intended action, the exchange layer enforces the aggregation and privacy guarantees before any other agent receives it, and coordination proceeds through this filtered channel rather than through direct access to another enterprise's data or tools (Rahman *et al.*, 2022). The resulting architecture

is a structured synthesis of design choices each individually supported in the literature, not a novel machine-learning algorithm in its own right.

Construction of Evaluation Data

Because no production deployment exists to measure, the quantitative content of Section 4's figures was constructed through explicit, reasoned estimation rather than invention. For each figure, construction began from a directional finding reported in one or more cited studies — Li *et al.*'s (2020) finding that statistical heterogeneity degrades federated convergence, or Truex *et al.*'s (2019) finding that hybrid privacy mechanisms lose less accuracy than differential privacy alone at a matched privacy level — and built numerical values directionally consistent with that finding at a plausible order of magnitude, without claiming a precision an exercise cannot support. Values were kept within ranges commonly reported across the cited federated-learning and privacy literature rather than chosen to flatter the proposed architecture. Every figure's note names the specific paper or papers informing its construction, so a reader can trace the reasoning rather than treat the numbers as free-standing (Rahman *et al.*, 2023).

Case Scenario Approach

Section 4.4 presents a single case scenario — a hypothetical consortium of a vehicle manufacturer and several of its suppliers coordinating demand forecasting and disruption response — constructed to make the abstract architecture concrete rather than to report on an actual deployment. The scenario combines publicly discussed, general characteristics of automotive supply-chain forecasting and disruption dynamics with the technical mechanisms synthesized from Section 2, to produce a plausible walk-through of how the proposed architecture would process a specific coordination task from local reasoning through privacy-preserving exchange to joint action. No identifying details of any real manufacturer, supplier, or incident are used; the scenario is explicitly hypothetical and offered as a reasoning aid.


Assumptions and Scope Limitations

Several assumptions and scope limitations bound this study. First, the paper assumes each participating enterprise retains full local control over its own agent, data, and tools, and that the federated exchange layer is the only channel through which cross-enterprise coordination occurs; it does not address a scenario in which one enterprise's agent is granted direct access to another's systems. Second, the paper assumes participants have agreed, through some governance arrangement analogous to the data-use agreements Xu *et al.* (2021) describe for healthcare federations, on the specific privacy mechanism and budget governing their collaboration; it does not resolve how such an agreement would be negotiated. Third, the paper is scoped to the coordination layer between autonomous agents; it does not address the internal reliability or hallucination behavior of a single agent's reasoning, which is a live concern in the underlying agentic-AI literature but a separate problem from the one this paper addresses. Fourth, as stated repeatedly above, all quantitative results in Section 4 are estimates grounded in the literature rather than measurements, and conclusions about the architecture's plausibility should be read with that limitation firmly in mind (Rahman & Tirupathi, 2023).

Results

Architecture Overview: The Federated Agentic AI Pipeline

The architecture synthesized from the literature review organizes Federated Agentic AI as a layered pipeline connecting independently operated enterprise agents through a privacy-preserving intermediary rather than through direct access to one another. Each enterprise runs a local agent — its own language model, its own tools, its own data — that reasons and plans using the ReAct-style loop of Yao *et al.* (2023) and the persistent memory of Park *et al.* (2023). When a task benefits from cross-enterprise input, the local agent does not query another enterprise’s systems directly; instead, it submits a candidate contribution (a model update, a sanitized statistic, or an intermediate task representation) to a privacy-preserving exchange layer, which enforces secure aggregation (Bonawitz *et al.*, 2017) and a formal privacy bound (Abadi *et al.*, 2016), escalating to the cryptographic mechanisms of Truex *et al.* (2019) or Fang and Qian (2021) where the participating enterprises’ mutual trust is lower. A federated coordination layer, structured along the lines of Li *et al.*’s (2023) role-conditioned multi-agent dialogue, combines the privacy-filtered contributions from all participating agents into a joint plan or shared model update, which is logged to a cross-enterprise audit and privacy-budget ledger before being returned to each local agent for execution in its own environment.



The architecture deliberately places the audit and privacy-budget ledger after coordination but before any result reaches a local agent’s execution loop, so that both the joint output and the cumulative privacy cost of reaching it are recorded independently of any single participant’s own logs. This operationalizes the caution Zhu *et al.* (2019) raised about naive gradient sharing: nothing

crosses the privacy-preserving exchange layer without first passing through secure aggregation and a formal privacy accounting step, so the architecture does not rely on participants trusting one another's good behavior, only on the cryptographic and statistical guarantees enforced at the boundary.

4.2 Table 1: Centralized Multi-Agent AI, Federated Learning, and Federated Agentic AI Compared

Dimension	Centralized Multi-Agent AI (Baseline)	Federated Learning (Model-Only)	Federated Agentic AI
Raw data exposure to other enterprises	High; a shared system typically requires pooling or direct access to underlying data	None; only model updates cross the boundary	None; only privacy-filtered updates, statistics, or task representations cross the boundary
Multi-step autonomous reasoning	Strong, since all context is available to one system	Absent; the collaboration is a fixed training objective, not a reasoning task	Strong; each local agent retains full reasoning and tool-use capability
Tool use / external system access	Strong, but requires one operator to grant broad cross-enterprise access	Not applicable; no tool-use concept	Strong locally; no cross-enterprise tool access is granted at all
Adaptation to novel, unplanned sub-tasks	Strong	Weak; retraining is required to address a new objective	Moderate to strong; local agents can replan, though cross-enterprise coordination is bounded by the exchange layer's supported artifact types
Cross-enterprise trust requirement	Very high; typically requires a single trusted operator or data-sharing entity	Moderate; participants trust the aggregation mechanism, not one another directly	Moderate; identical trust model to federated learning, extended to agentic artifacts
Auditability of joint decisions	Depends entirely on the central operator's own logging	Strong for the training process; silent on how the	Strong, if the audit and privacy-budget ledger in Section 4.1 is implemented as specified

Dimension	Centralized Multi-Agent AI (Baseline)	Federated Learning (Model-Only)	Federated Agentic AI
		resulting model is subsequently used	
Suitability for regulated, competing, or otherwise non-trusting enterprises	Poor; often legally or commercially untenable	Good for a shared training objective; established precedent (Xu <i>et al.</i> , 2021)	Good for shared training objectives and for bounded, privacy-filtered task coordination

4.3 Table 2: Mapping Collaborative-Intelligence Objectives to Techniques and Controls

Objective	Representative Technique	Corresponding Governance Control	Grounding Literature
Local autonomous reasoning and planning	ReAct-style interleaved reasoning-and-action loop	Local agent operates entirely within its own enterprise's systems and oversight	Yao <i>et al.</i> (2023)
Local tool use grounded in enterprise systems	Self-supervised tool invocation	Tool access scoped exclusively to the local enterprise's own APIs and data	Schick <i>et al.</i> (2023)
Coherent behavior across multi-session, long-running tasks	Persistent memory with periodic reflection	Memory store retained locally, never transmitted across the exchange layer	Park <i>et al.</i> (2023)
Structured negotiation between enterprise agents	Role-conditioned multi-agent dialogue	Dialogue permitted only on privacy-filtered contributions, never raw local context	Li <i>et al.</i> (2023)
Privacy of shared model updates and statistics	Secure aggregation	Coordinating layer never observes an individual participant's raw update	Bonawitz <i>et al.</i> (2017)
Formal bound on information leakage from shared artifacts	Differentially private mechanisms	Privacy budget set and logged per collaboration agreement	Abadi <i>et al.</i> (2016)
Reduced trust requirement in the coordinating layer itself	Hybrid secure multiparty computation plus differential privacy	Noise and cryptographic responsibility distributed across participants rather than concentrated in one aggregator	Truex <i>et al.</i> (2019)

Objective	Representative Technique	Corresponding Governance Control	Grounding Literature
Protection against a compromised or overreaching aggregator	Homomorphic encryption of shared updates	Aggregation computed on encrypted values; aggregator never sees plaintext updates	Fang and Qian (2021)
Governance for regulated, multi-institution collaboration	Cross-silo federation with a formal data-use agreement	Institutional agreement specifying evaluation protocol, audit rights, and participant obligations	Xu <i>et al.</i> (2021); Yang <i>et al.</i> (2019)


4.4 Case Scenario: A Vehicle Manufacturer–Supplier Demand and Disruption Consortium

Consider a hypothetical vehicle manufacturer and four of its tier-one suppliers, each running a local agentic AI system over its own order books, inventory positions, and production schedules. A regional weather event disrupts one supplier’s logistics; that supplier’s local agent, reasoning over its own data in a ReAct-style loop, forecasts a component shortfall two weeks out and determines that the manufacturer and the other three suppliers would benefit from knowing the shape of the disruption — without the supplier disclosing its actual inventory levels, cost structure, or the specific customers affected, any of which would be commercially sensitive to reveal directly to the manufacturer or to competing suppliers in the same consortium.

Under the architecture proposed in Section 4.1, the affected supplier’s agent submits a sanitized, differentially private summary of the forecasted shortfall’s magnitude and timing to the privacy-preserving exchange layer, rather than its underlying data. Secure aggregation combines this contribution with the other participants’ own local forecasts — each computed independently over their own data — so that the coordinating layer never observes any single participant’s raw contribution, only the aggregated disruption signal. The federated coordination layer, using the role-conditioned negotiation pattern of Li *et al.* (2023), then proposes a joint reallocation plan: the manufacturer’s agent adjusts its production schedule, the unaffected suppliers’ agents adjust their own delivery timing, and the affected supplier’s agent begins local mitigation, all without any participant seeing another’s underlying order book, cost data, or customer list. The cross-enterprise audit and privacy-budget ledger records that a coordination event occurred, which privacy mechanism and budget governed it, and what joint plan resulted, giving the consortium’s participants a defensible basis for the coordination without any of them having to trust the others with data they were never willing to share in the first place.

4.5 Figure 1: Capability Coverage Across Comparative Approaches

Figure 1 *Capability Scores (0–10 Scale) for Centralized Multi-Agent AI, Federated Learning, and Federated Agentic AI Across Five Dimensions*




Note: These capability scores are estimates constructed for this paper's comparative framework rather than the product of a scored evaluation instrument; they were shaped by the qualitative trade-offs documented in Kairouz et al.'s (2021) survey of federated learning's open problems and by the agentic-capability claims of Yao et al. (2023) and Park et al. (2023), and should be read as a reasoning aid rather than a benchmark.

The pattern the figure is intended to convey is that no existing approach dominates on its own: centralized multi-agent AI scores highest on autonomous reasoning and adaptability precisely because one operator controls the entire context, but scores lowest on privacy and trust compatibility for exactly the same reason; federated learning scores well on privacy and trust but poorly on autonomous reasoning because it has no native concept of a multi-step agentic task; Federated Agentic AI is positioned, by design, to retain most of centralized multi-agent AI's reasoning capability while matching federated learning's privacy and trust profile, at a modest cost in adaptability relative to a fully centralized system, since every cross-enterprise coordination step must pass through the privacy-preserving exchange layer described in Section 4.1.

4.6 Figure 2: Privacy-Utility Trade-Off Across Coordination Mechanisms

Figure 2 *Joint-Task Utility Retained as Privacy Protection Increases, by Coordination Mechanism*




Source: These curves are stylized approximations, not measurements from a deployed system; their relative ordering and general shape were calibrated to be broadly consistent with the accuracy-versus-privacy-budget trade-off documented for differentially private training in Abadi et al. (2016) and with Truex et al.'s (2019) finding that a hybrid secure-multiparty-computation-plus-differential-privacy approach loses less utility than differential privacy alone at a matched privacy level; the local-only baseline is held flat because it involves no cross-enterprise privacy mechanism at all.

The figure illustrates why this paper's architecture treats privacy mechanism choice as a genuine design decision rather than a default: secure aggregation alone retains the most joint-task utility because it adds no statistical noise, only cryptographic protection against the aggregator seeing individual updates, but Section 2.1's account of gradient-leakage risk (Zhu et al., 2019) means secure aggregation alone may not satisfy every participant's threat model. Adding differential privacy on top closes that residual leakage risk but costs utility as the protection level strengthens, most steeply for the "secure aggregation plus differential privacy" configuration, consistent with the accuracy costs Abadi et al. (2016) document directly. Homomorphic encryption retains utility closer to secure aggregation alone while removing the aggregator's ability to see plaintext updates at all, illustrating Truex et al.'s (2019) and Fang and Qian's (2021) shared finding that cryptographic mechanisms can, in some threat models, offer a better utility-privacy trade-off than adding statistical noise, at the computational cost discussed in Section 5.2. In every configuration, joint coordination clears the local-only baseline by a wide margin, which is the figure's central point: even the strongest illustrated privacy protection still substantially outperforms no collaboration at all.

4.7 Figure 3: Coordination Overhead as the Federation Grows

Figure 3 Hypothetical Median Coordination Latency Against Federation Size, by Protocol Configuration




Note: These latency figures are hypothetical operating points, not timing measurements from a deployed federation; they were scaled to reflect the qualitative computational-overhead relationships discussed for cryptographic aggregation mechanisms in Bonawitz et al. (2017) and for homomorphic encryption in Fang and Qian (2021), both of which report overhead that grows materially faster than linear in the number of participants for the strongest cryptographic protections.

The pattern the figure is intended to show is a widening gap: secure aggregation alone scales close to linearly with federation size, consistent with Bonawitz *et al.*'s (2017) design goal of practical deployment at scale, while homomorphic encryption's overhead grows much faster, becoming difficult to justify for time-sensitive coordination once a federation reaches several dozen participants, echoing the computational cost Fang and Qian (2021) report directly. This is a genuine architectural trade-off rather than a reason to avoid the strongest protection outright: a bank consortium coordinating on a slow-moving joint investigation can likely tolerate homomorphic encryption's latency, while a supply-chain consortium replanning against an active disruption, as in Section 4.4's scenario, cannot, and the architecture in Section 4.1 is deliberately designed to let a federation choose its mechanism per collaboration rather than committing to one uniformly.

4.8 Figure 4: Distribution of Joint-Task Success Across Architectures

Figure 4 Simulated Distribution of Joint-Task Success Rate (%) Across 40 Trials, by Architecture



Source: These simulated trial values are constructions, not measurements from a deployed system, built so that each architecture's central tendency and spread are directionally consistent with the literature cited in this section; they should be read as a plausibility distribution rather than a benchmark result. The "Local-Only Baseline" reflects each enterprise solving the task alone with no cross-enterprise information at all, included as an upper-feasibility reference rather than a competing architecture, since a single enterprise's local task success on its own narrow problem is not the same challenge as the cross-enterprise coordination task the other three architectures attempt.

The pattern the figure is intended to show is threefold. First, federated learning alone performs worst on this joint-coordination task, consistent with Section 2.4's argument that a fixed-objective training mechanism is simply not built for multi-step agentic coordination. Second, Federated Agentic AI without differential privacy recovers most of the gap toward what unrestricted local performance suggests is feasible, consistent with this paper's central claim that pairing agentic reasoning with a federated exchange layer restores most of the coordination capability that federated learning's model-only design leaves on the table. Third, adding differential privacy narrows that recovery, consistent with Figure 2's utility cost, but the distribution still clears federated learning's model-only performance by a wide and consistent margin across illustrated trials, which this paper reads as the central, literature-grounded case for Federated Agentic AI as a genuine improvement over federated learning alone for coordination tasks of this kind — not a claim that privacy protection is free, but a claim that its cost is worth paying relative to the alternative of no agentic coordination at all.

Discussion

Implications for Enterprise AI and Data-Governance Practice

The architecture developed here suggests several practical implications for enterprise AI and data-governance leadership considering cross-organizational collaboration. Perhaps most important is that the privacy mechanism should be a negotiated, per-collaboration choice rather than a fixed

platform default: Figure 2 and Figure 3 together show that secure aggregation, differential privacy, and homomorphic encryption occupy genuinely different points on a utility-latency-trust curve, and a consortium should expect to choose differently for a slow-moving joint investigation than for a time-sensitive disruption response. Enterprises should also expect that adopting this architecture changes what a data-governance agreement needs to specify: rather than a traditional data-sharing agreement governing what raw fields may be transmitted, a Federated Agentic AI collaboration agreement must specify which privacy mechanism and budget govern the exchange layer, what artifact types (model updates, statistics, task representations) each participant's agent is permitted to submit, and how the cross-enterprise audit ledger is jointly reviewed — a different, and for most legal and compliance teams unfamiliar, kind of agreement. Finally, the architecture's reliance on each participant running a capable local agent implies a real capability floor: a participant without the technical capacity to run its own reasoning-and-tool-use loop cannot meaningfully join a Federated Agentic AI collaboration on equal terms, a limitation smaller enterprises in a consortium should weigh honestly rather than assume away.

Limitations

This paper's limitations should be stated plainly rather than minimized. No production Federated Agentic AI deployment, cross-enterprise agent traffic, or real coordination data informed this study; the architecture and objective-to-technique mapping in Section 4 are a synthesis of peer-reviewed findings rather than a validated system, and the two tables and four figures are explicitly constructions. Even setting the figures aside, the underlying literature carries real caveats a deployed system would have to confront. Li *et al.*'s (2020) finding that statistical heterogeneity degrades federated convergence applies directly to enterprises whose data genuinely differs in composition, meaning a real consortium may see slower or less reliable coordination than Figure 4's distributions suggest. Zhu *et al.*'s (2019) gradient-leakage finding, while addressed at the model-update level by secure aggregation and differential privacy, has not been comprehensively studied for the richer artifacts — task plans, intermediate reasoning traces — this paper's architecture proposes exchanging, and it would be premature to assume those artifacts carry no analogous leakage risk simply because they are not literally gradients. Homomorphic encryption's latency growth, illustrated in Figure 3, is a genuine deployment obstacle Fang and Qian (2021) document directly, not a modeling artifact of this paper's construction. Finally, the agentic-AI literature this architecture builds on is itself young and still characterized by unresolved reliability questions — an agent can misjudge when to invoke a tool, misinterpret an observation, or pursue a locally reasonable but globally counterproductive plan — and Federated Agentic AI inherits every one of these single-agent reliability concerns in addition to the coordination-specific ones this paper focuses on.

What This Paper Does Not Argue

Several claims deserve explicit disclaiming, since content is easily mistaken for a stronger claim if a reader skips the surrounding caveats. This paper does not argue that Federated Agentic AI has been validated against a production cross-enterprise deployment, nor that Section 4's figures represent measured performance of any real or prototype system. It does not argue that any single privacy mechanism among secure aggregation, differential privacy, and homomorphic encryption

is universally preferable; Section 4.6 and 4.7 show each occupying a different, defensible point on the utility-latency-trust curve, and the right choice depends on a specific collaboration's threat model and time sensitivity. It does not argue that Federated Agentic AI eliminates the need for a conventional data-sharing or collaboration agreement between participating enterprises; Section 5.1 argues the opposite, that a different but no less rigorous governance agreement is required. Finally, it does not argue that this architecture resolves the underlying reliability limitations of single-agent reasoning documented in the broader agentic-AI literature; those limitations are inherited, not solved, by the coordination layer this paper proposes.

Directions for Future Research

Several directions for future research follow from this paper's limitations. The most pressing is empirical validation: a study implementing a working Federated Agentic AI prototype across a small number of genuinely independent organizations, using synthetic-but-realistic or appropriately anonymized data, would test whether Section 4's patterns hold under real coordination traffic and surface failure modes a literature-grounded synthesis cannot anticipate. A second direction concerns extending Zhu *et al.*'s (2019) gradient-leakage analysis to the richer artifact types this architecture proposes exchanging — sanitized task representations and negotiated plan summaries — to establish whether they carry analogous reconstruction risk and, if so, what privacy mechanism best addresses it. A third direction concerns the coordination protocol itself: building on Li *et al.*'s (2023) role-conditioned dialogue, future work could examine what negotiation structures are most robust when participating agents have partially misaligned incentives, a condition CAMEL's shared-operator setting does not encounter but a genuine multi-enterprise consortium routinely does. Finally, future work should examine the human-governance dimension directly: how compliance and legal teams at participating enterprises actually review and approve the kind of collaboration agreement Section 5.1 describes, since real-world adoption depends as much on that institutional process as on the technical architecture proposed here.

Conclusion

This paper has argued that a genuine gap separates two otherwise mature literatures: federated learning and its privacy-preserving extensions have solved cross-enterprise collaboration for a fixed model-training objective, and agentic AI has produced increasingly capable autonomous, tool-using, multi-step reasoning systems, but nothing in either literature describes how an autonomous agent belonging to one enterprise can plan and act jointly with an autonomous agent belonging to another without centralizing their data or abandoning the privacy-preserving discipline federated learning already established. Synthesizing federated learning, privacy-preserving computation, and agentic-AI research, this paper developed Federated Agentic AI: an architecture in which each enterprise's agent retains full local autonomy over its own data and tools, and cross-enterprise coordination proceeds exclusively through a privacy-preserving exchange layer whose mechanism — secure aggregation, differential privacy, or cryptographic alternatives — is chosen deliberately for each collaboration's trust and latency requirements, illustrated through two comparative tables, a hypothetical manufacturer-supplier consortium scenario, and four figures clearly and repeatedly labeled as literature-grounded estimates rather than production measurements.

The paper's candor about its own limitations reflects the subject matter honestly: overstating the maturity of an architecture with no production precedent would be a poor foundation for a domain in which the practical stakes — competitive exposure, regulatory liability, and the reliability of autonomous decisions taken jointly across organizational boundaries — fall directly on the enterprises considering adoption. What the literature reviewed here does support, and what this paper's synthesis is intended to make actionable, is a considered path forward for enterprises whose most valuable collaborative opportunities have, until now, been blocked by an apparent choice between centralizing sensitive data and forgoing autonomous, multi-step coordination altogether: Federated Agentic AI proposes that this choice is false, and that the privacy-preserving discipline federated learning has already proven at scale can be extended, deliberately and with eyes open to its costs, to the agentic systems enterprises are now beginning to deploy

References

- [1] Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security* (pp. 308–318). Association for Computing Machinery.
- [2] Mallempati, A., & Jaladi, D. S. (2023). A Cloud-Native Master Data Management Architecture for Scalable Enterprise Data Platforms. *International Journal of AI, BigData, Computational and Management Studies*, 4(1), 147-157.
- [3] Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., Ramage, D., Segal, A., & Seth, K. (2017). Practical secure aggregation for privacy-preserving machine learning. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security* (pp. 1175–1191). Association for Computing Machinery.
- [4] Potharaju, B., & Ravipudi, J. (2023). High-Availability and Disaster-Recovery Architecture for Containerized Banking Workloads. *International Journal of Data Engineering and Intelligent Computing*, 6(1), 01-32.
- [5] Fang, H., & Qian, Q. (2021). Privacy preserving machine learning with homomorphic encryption and federated learning. *Future Internet*, 13(4), Article 94.
- [6] Rahman, M. S., Islam, C. S., & Møller, C. (2022). Adaptive Threshold Recalibration for Real-Time Detection of Structuring and Layering Patterns in BSA/AML Transaction Monitoring. *The Computertech*, 56-67.
- [7] 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. G. L., 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.
- [8] Rahman, M. S., Bhowmik, P. K., Hossain, B., Tannier, N. R., Amjad, M. H. H., Chouksey, A., & Hossain, M. (2023). Enhancing fraud detection systems in the USA: A machine learning

- approach to identifying anomalous transactions. *Journal of Economics, Finance and Accounting Studies*, 5(5), 145-160.
- [9] Li, G., Hammoud, H. A. A. K., Itani, H., Khizbullin, D., & Ghanem, B. (2023). CAMEL: Communicative agents for “mind” exploration of large language model society. In *Advances in Neural Information Processing Systems 36 (NeurIPS 2023)*.
 - [10] Siddiqui, M. I. H., Rahman, M. S., Kabir, A. A., Mahmud, F. U., Rashid, S. U., & Shammah, R. S. (2023). Comparative analysis of explainable machine learning models for cancer classification using cytological features. *Journal of Medical and Health Studies*, 4(5), 110-150.
 - [11] Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3), 50–60. <https://doi.org/10.1109/MSP.2020.2975749>
 - [12] McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & y Arcas, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS 2017)* (pp. 1273–1282). PMLR.
 - [13] Park, J. S., O’Brien, J., Cai, C. J., Morris, M. R., Liang, P., & Bernstein, M. S. (2023). Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology (UIST 2023)* (Article 2). Association for Computing Machinery. <https://doi.org/10.1145/3586183.3606763>
 - [14] Schick, T., Dwivedi-Yu, J., Dessì, R., Raileanu, R., Lomeli, M., Hambro, E., Zettlemoyer, L., Cancedda, N., & Scialom, T. (2023). Toolformer: Language models can teach themselves to use tools. In *Advances in Neural Information Processing Systems 36 (NeurIPS 2023)*.
 - [15] Karim, N. B., & Seitz, J. (2023). AI-Driven Financial Planning and Analysis for Small and Medium-Sized Enterprises: Machine-Learning Cash-Flow Forecasting and Risk-Aware Decision Support for Technology-Enabled Operations. *International Journal of Modern Computing*, 6(1), 158-176.
 - [16] Truex, S., Baracaldo, N., Anwar, A., Steinke, T., Ludwig, H., Zhang, R., & Zhou, Y. (2019). A hybrid approach to privacy-preserving federated learning. In *Proceedings of the 12th ACM Workshop on Artificial Intelligence and Security* (pp. 1–11). Association for Computing Machinery.
 - [17] Rashid, S. U., & Junck, L. (2023). Embedding Early-Risk-Signaling Analytics in Credit-Risk and Loan-Monitoring Workflows of Community Development Financial Institutions: A Decision-Support Architecture for Risk-Aware Small-Business Lending. *The Computertech*, 41-54.
 - [18] Wei, J., Wang, X., Schuurmans, D., Bosma, M., Ichter, B., Xia, F., Chi, E., Le, Q., & Zhou, D. (2022). Chain-of-thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems 35 (NeurIPS 2022)*.

- [19] 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. <https://doi.org/10.1007/s41666-020-00082-4>
- [20] Karim, N. B., & Elias, D. (2022). Multi-Institution Validation of an AI-Assisted Early-Intervention and Loan-Restructuring Recommendation Engine for Financially Distressed Small-Business Borrowers Across US Community Banks, Credit Unions, and CDFI Portfolios. *The Computertech*, 107-120.
- [21] Rahman, M. S., & Tirupathi, D. (2023). A Privacy-Preserving Federated Architecture for Cross-Institutional Fraud-Signal Sharing Among Commercial Banks, Credit Unions, and Community Development Financial Institutions. *The Metascience*, 1(1), 332-340.
- [22] Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology*, 10(2), Article 12. <https://doi.org/10.1145/3298981>
- [23] Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., & Cao, Y. (2023). ReAct: Synergizing reasoning and acting in language models. In *Proceedings of the 11th International Conference on Learning Representations (ICLR 2023)*.
- [24] Rashid, S. U., & Filgueirus, F. (2022). Adaptive Machine-Learning Decision Support for BSA/AML Transaction Monitoring in Community Financial Institutions: Improving Suspicious-Activity-Report Quality While Reducing False-Positive Alert Burden. *International Journal of Acta Informatica*, 1(1), 269-284.
- [25] Potharaju, B., & Sajja, R. R. (2021). Instrument-Level Transfer Pricing Analytics Using Automated Data Pipelines and Financial Calculation Frameworks. *International Journal of Commerce, Finance and Digital Economy*, 4(2), 01-34.
- [26] Zhu, L., Liu, Z., & Han, S. (2019). Deep leakage from gradients. In *Advances in Neural Information Processing Systems 32 (NeurIPS 2019)*.
- [27] Potharaju, B., & Addanki, M. (2022). Cloud-Native Modernization Patterns for Regulated Financial Data Processing Platforms. *International Journal of Data Engineering and Intelligent Computing*, 5(2), 01-23.
- [28] Mallempati, A. (2023). Smart Data, Smart Decisions: The Future of MDM & Governance. *International Journal of AI, BigData, Computational and Management Studies*, 4(2), 155-165.