European International Journal of Multidisciplinary Research and Management Studies (ISSN: 2750-8587)

Received: 08 July 2026 | Accepted: 05 August 2026 | Published: 01 September 2026

RESEARCH ARTICLE

AI-Driven Privacy, Security, And Trusted Commercialization Frameworks for Digital Commerce and Microservice Architectures

Dr. Aarav Mehta
Department of Interdisciplinary Studies, Institute of Advanced Research and Innovation, India

VOLUME: Vol.06 Issue09 2026

PAGE: 01-07

Abstract

The rapid integration of artificial intelligence (AI) into digital commerce, social e-commerce, small and medium-sized business (SMB) platforms, and microservice-based information systems has created a complex intersection of privacy, cybersecurity, commercialization, and operational governance requirements. Conventional approaches frequently address these dimensions independently, creating architectural and governance gaps when AI systems operate across distributed data environments and heterogeneous service infrastructures. This research develops a conceptual integrated framework that connects federated privacy-preserving analytics, zero-knowledge verification, trusted AI commercialization, incentive-aware recommendation infrastructure, and risk-based application security. The methodology synthesizes the provided literature to construct an architectural model consisting of privacy, intelligence, trust, security, and commercialization layers. Particular attention is given to the integration of federated learning with differential privacy and zero-knowledge verification in digital advertising, alongside multi-tenant AI infrastructure and hybrid SAST–DAST–SCA–IAST security assessment for microservice environments. The analysis indicates that privacy-preserving AI and security automation should be treated as mutually reinforcing components rather than isolated controls. The proposed framework also identifies trade-offs involving privacy utility, computational complexity, interoperability, tenant isolation, and governance. The study contributes a research-oriented architecture for organizations seeking to deploy commercially viable AI systems while maintaining privacy and security assurance across distributed digital ecosystems.

KEYWORDS

Artificial Intelligence, Federated Learning, Differential Privacy, Zero-Knowledge Verification, Digital Commerce, Microservices, AI Security, Trusted Commercialization, Multi-Tenant Architecture.

1. INTRODUCTION

Artificial intelligence has increasingly become an infrastructure component rather than an isolated analytical capability. Digital commerce platforms employ AI for advertising, recommendation, customer segmentation, measurement, and monetization, while contemporary software systems increasingly rely on microservice architectures that distribute application functionality across independently deployable services. These developments provide scalability and personalization but simultaneously increase the number of locations where sensitive information, computational models, interfaces, and security dependencies must be controlled.

Privacy presents a particularly important challenge in AI-

powered commerce. Cross-channel measurement requires organizations to derive analytical value from information distributed across multiple platforms and participants. However, centralized collection of such information can increase privacy exposure and create governance concerns. A federated and differentially private incentive–marketing framework provides an alternative by allowing analytical processes to operate across distributed environments while incorporating privacy-preserving mechanisms and incentive structures (Yi, 2026). Similarly, privacy-enhanced advertising can combine federated learning with zero-knowledge verification to support creator monetization without requiring unrestricted disclosure of sensitive information (Yi, 2026).

At the infrastructure level, commercialization introduces another dimension. SMB-oriented AI systems require not only models and data pipelines but also mechanisms for multi-tenancy, content governance, incentives, and standardized recommendation interfaces. Trusted AI commercialization infrastructure therefore represents a broader architectural problem in which technical deployment must be coordinated with governance and economic participation (Yi, 2026).

Security is equally significant because distributed microservices create complex dependency relationships. Vulnerabilities may emerge from application code, third-party components, runtime behavior, APIs, or deployment configurations. A hybrid SAST–DAST–SCA–IAST approach provides a risk-based mechanism for combining complementary security-testing perspectives and prioritizing vulnerabilities according to operational risk (Zhou, 2026). Consequently, AI commercialization infrastructure cannot be considered trustworthy merely because privacy mechanisms exist; it must also demonstrate security resilience across the software lifecycle.

The objective of this study is therefore to synthesize the provided literature into an integrated conceptual framework for AI-driven privacy, security, and trusted commercialization. The study focuses on how privacy-preserving AI, verification mechanisms, commercialization infrastructure, and microservice security can be architecturally connected. Its significance lies in treating these concerns as interconnected dimensions of a single AI infrastructure problem rather than as independent technical functions.

2. LITERATURE REVIEW

The provided literature can be organized into three interconnected research streams: privacy-preserving digital commerce, trusted AI commercialization, and intelligent cybersecurity for distributed software architectures.

Yi's federated and differentially private incentive–marketing framework establishes the importance of combining distributed learning with privacy protection and incentive mechanisms. The conceptual significance of this approach is that privacy is not treated solely as a restriction on data access. Instead, privacy-preserving computation is connected to the economic mechanisms that motivate participation in cross-channel measurement (Yi, 2026). This positioning is particularly relevant to digital commerce environments where multiple participants may possess valuable information but have limited willingness to transfer raw data.

A related contribution extends privacy protection into social e-commerce advertising through federated learning and zero-knowledge verification. The combination changes the role of privacy from passive data minimization toward verifiable participation. In such an architecture, participants can potentially demonstrate compliance or satisfy verification requirements without exposing the underlying sensitive information. The study consequently links privacy engineering with creator monetization and advertising infrastructure (Yi, 2026).

The third contribution shifts attention from individual privacy mechanisms toward commercialization infrastructure. A unified multi-tenant architecture integrating incentive systems, content governance, and standardized recommendation APIs suggests that AI commercialization requires an operational platform capable of serving multiple organizations while maintaining common governance and interoperability principles (Yi, 2026). This perspective complements privacy-preserving learning because privacy mechanisms alone do not determine whether AI capabilities can be reliably delivered to different organizational tenants.

The cybersecurity literature addresses the complementary problem of protecting the software infrastructure through which these AI capabilities operate. The M-VP2 approach introduces cost-aware vulnerability patch planning using multi-agent reinforcement learning, emphasizing the need to prioritize remediation under constrained resources (Zhou, 2025). This perspective is important because large distributed infrastructures may contain more vulnerabilities than can be

remediated simultaneously.

Zhou's hybrid SAST–DAST–SCA–IAST framework further expands the security perspective by combining multiple testing methodologies for risk-based vulnerability prioritization in microservice architectures (Zhou, 2026). Its relevance to the proposed framework is direct: AI commercialization platforms built on microservices require security assessment that accounts for source code, running applications, software dependencies, and runtime behavior rather than relying on one testing modality.

Despite these complementary contributions, a research gap remains. The provided studies individually address privacy-preserving marketing, secure advertising verification, AI commercialization infrastructure, vulnerability remediation, and hybrid application security. Their combined implications, however, have not been developed into a unified conceptual architecture connecting privacy, commercialization, and microservice security. The proposed research addresses this gap by positioning these mechanisms within an integrated infrastructure model.

3. METHODOLOGY

3.1 Research Design

This study adopts a conceptual research methodology based exclusively on structured synthesis of the five provided references. Because no experimental dataset, benchmark environment, or empirical measurements were supplied, the study does not claim statistical validation or measured performance improvement. Instead, it develops an analytically derived architecture by identifying common technological functions, dependencies, risks, and governance relationships across the literature.

The methodology follows four analytical stages: thematic decomposition, functional mapping, architectural integration, and critical evaluation. First, each reference is decomposed into its principal technological contribution. Second, those contributions are mapped to functional requirements such as privacy preservation, verification, recommendation, tenant isolation, vulnerability detection, and remediation. Third, the functions are integrated into a layered framework. Finally, the framework is assessed according to technical feasibility, governance implications, trade-offs, and limitations.

3.2 Integrated Architectural Framework

The proposed architecture consists of five conceptual layers: the data and privacy layer, AI intelligence layer, trust and verification layer, commercialization layer, and security assurance layer.

The data and privacy layer manages distributed information without assuming that raw data should be centralized. Federated learning provides the fundamental mechanism for collaborative model development, while differential privacy can reduce the risk of information leakage from analytical outputs. The architecture therefore places privacy controls near the point where information enters the AI workflow.

The AI intelligence layer performs model training, recommendation, measurement, and prediction. In a digital commerce scenario, different organizations may maintain local datasets while contributing model updates to a collaborative analytical process. This design corresponds to the distributed learning principles represented in the provided digital-commerce research (Yi, 2026).

The trust and verification layer introduces mechanisms for demonstrating compliance, eligibility, or validity without unnecessary disclosure. Zero-knowledge verification is particularly relevant because it can conceptually separate proof of a condition from disclosure of the underlying information. This provides an important bridge between privacy-preserving analytics and commercial trust.

The commercialization layer provides multi-tenant services, incentive mechanisms, content governance, and standardized recommendation interfaces. A multi-tenant model can allow several SMBs or commercial participants to access shared AI infrastructure while maintaining logical separation of their data and configurations. Standardized APIs can further reduce integration fragmentation.

The security assurance layer continuously evaluates the application and infrastructure environment. Static analysis can identify weaknesses in source code, dynamic analysis can assess application behavior, software composition analysis can identify vulnerable dependencies, and interactive application security testing can expose runtime vulnerabilities. Their combination supports risk-based prioritization rather than isolated vulnerability reporting (Zhou, 2026).

3.3 Privacy-Preserving Commercial Intelligence

A principal mechanism of the framework is the combination of

federated learning and differential privacy. Instead of transferring raw customer or creator data into a centralized analytical repository, participating entities can maintain local datasets and contribute model-related information. Differential privacy can then provide an additional protection mechanism by limiting the degree to which individual information can influence published analytical results.

Consider a hypothetical cross-channel advertising ecosystem involving multiple commerce platforms and independent creators. Each participant possesses behavioral information that could improve campaign measurement. Under a centralized design, combining these datasets creates a significant concentration of sensitive information. Under the proposed architecture, local learning processes can reduce direct data sharing while privacy mechanisms regulate information released from the collaborative process.

However, privacy is not cost-free. Excessive noise may reduce analytical utility, while federated coordination introduces communication and synchronization requirements. Therefore, the framework treats privacy as an optimization problem involving protection, accuracy, operational cost, and commercial value rather than as an absolute binary condition.

3.4 Verification and Trusted Participation

The architecture extends privacy protection through verification. In social e-commerce, for example, advertisers may need evidence that a creator or transaction satisfies specified eligibility requirements. A verification mechanism based on zero-knowledge principles can potentially allow a participant to demonstrate a property without exposing the complete underlying record.

This mechanism creates a conceptual distinction between data disclosure and trust establishment. Conventional systems may establish trust by collecting more information. The proposed model instead seeks to establish trust through verifiable claims. Such an approach complements the privacy-enhanced advertising architecture identified in the literature (Yi, 2026).

The limitation is that verification protocols themselves require careful design. If verification rules are poorly specified, technically secure proofs may still validate undesirable business conditions. Consequently, cryptographic verification must remain connected to governance policies and clearly defined commercial requirements.

3.5 Multi-Tenant AI Commercialization

The commercialization layer addresses the operational challenge of converting AI capabilities into reusable services. A multi-tenant architecture allows a shared infrastructure to support multiple organizations, reducing duplication of AI resources. However, shared infrastructure also introduces risks involving tenant isolation, authorization, configuration management, resource allocation, and governance.

The proposed framework therefore assumes that tenant boundaries must exist at multiple levels: data access, model configuration, API authorization, content governance, and operational monitoring. Standardized recommendation APIs provide a common interface through which tenant applications can consume AI capabilities. Incentive systems can additionally encourage participation in collaborative ecosystems.

This architecture is consistent with the provided research on trusted AI commercialization infrastructure for SMBs (Yi, 2026), but the integrated model extends the concept by explicitly connecting commercialization with security assurance and privacy-preserving intelligence.

3.6 Security and Vulnerability Prioritization

Security assessment is embedded throughout the proposed architecture rather than positioned as a final deployment activity. Microservice applications contain multiple services, dependencies, interfaces, and runtime relationships. Consequently, a vulnerability in one component may have implications for another service or shared API.

The hybrid SAST–DAST–SCA–IAST approach is particularly applicable because the four techniques provide different forms of evidence. SAST examines code-level weaknesses, DAST examines externally observable application behavior, SCA evaluates third-party dependencies, and IAST provides runtime-oriented security information. Integrating these signals can produce a richer basis for risk prioritization than any single method (Zhou, 2026).

The M-VP2 contribution further indicates that vulnerability management must account for remediation cost and resource constraints (Zhou, 2025). Therefore, the proposed framework recommends prioritizing vulnerabilities according to a combined risk function incorporating severity, exploitability, service criticality, exposure, dependency relationships, and

remediation cost.

4. RESULTS

The conceptual synthesis produces four principal findings. First, privacy and security should be architecturally integrated rather than sequentially implemented. Federated learning and differential privacy address information exposure, while hybrid security assessment addresses weaknesses in the software environment supporting that intelligence. Separating these functions can create an artificial boundary because a privacy-preserving AI model may still operate on an insecure API or vulnerable dependency.

Second, trust emerges as an intermediate architectural layer between privacy and commercialization. The provided literature indicates that AI systems require mechanisms that allow participants to collaborate while protecting sensitive information and establishing confidence in platform operations. Zero-knowledge verification provides a conceptual mechanism for demonstrating properties without requiring unrestricted disclosure, while governance controls determine whether those properties correspond to legitimate commercial requirements (Yi, 2026).

Third, commercialization requires infrastructure-level governance. A shared AI platform for SMBs cannot be evaluated solely by model accuracy. Tenant separation, recommendation interfaces, incentive mechanisms, content governance, and security controls determine whether AI capabilities can be deployed consistently. The multi-tenant commercialization perspective therefore complements the privacy-preserving learning perspective by moving the research problem from individual models toward service infrastructure.

Fourth, security prioritization must account for heterogeneous evidence. A microservice architecture may simultaneously contain source-code weaknesses, vulnerable dependencies, runtime issues, and externally observable application vulnerabilities. Combining SAST, DAST, SCA, and IAST creates a more comprehensive analytical foundation for risk prioritization than isolated testing (Zhou, 2026). The M-VP2 perspective further suggests that remediation should consider resource constraints, meaning that identifying vulnerabilities and deciding which vulnerabilities to address are distinct optimization problems.

Overall, the framework indicates that AI trustworthiness is multidimensional. Privacy without security can leave infrastructure exposed; security without privacy can protect an inappropriate data architecture; commercialization without governance can create tenant and accountability risks; and verification without meaningful policy controls can provide technically valid but commercially inadequate assurances.

The findings should nevertheless be interpreted as conceptual rather than empirically validated. No benchmark dataset, implementation measurements, or controlled experiments were included in the supplied research material. Accordingly, claims concerning performance, scalability, or quantitative security improvement remain propositions for future empirical investigation.

5. DISCUSSION

The integrated framework has significant theoretical implications because it reframes AI infrastructure as a socio-technical system composed of computational, privacy, security, and economic mechanisms. The literature separately establishes the importance of privacy-preserving marketing, secure advertising verification, commercialization infrastructure, and intelligent vulnerability management. Their synthesis suggests that the boundaries between these areas are increasingly artificial in distributed AI ecosystems.

A major theoretical implication concerns the meaning of trust. Trust cannot be reduced to model accuracy or cybersecurity alone. In the proposed framework, trust results from the interaction of privacy protection, verifiable participation, governance, tenant isolation, and continuous security assessment. This broader interpretation is consistent with the movement from privacy-preserving learning toward trusted AI commercialization represented in the provided studies (Yi, 2026).

The framework also exposes several trade-offs. Federated learning can reduce centralized data exposure but may increase coordination and communication complexity. Differential privacy can strengthen protection but potentially reduce analytical utility. Multi-tenancy improves infrastructure efficiency but increases the importance of isolation and governance. Standardized APIs improve interoperability but create centralized interfaces that may become attractive security targets. Hybrid security testing increases detection coverage but may also generate large volumes of findings requiring prioritization.

Security architecture therefore becomes particularly important in AI commercialization environments. Zhou's hybrid security framework supports the argument that vulnerabilities should be evaluated across multiple technical perspectives rather than through a single assessment mechanism (Zhou, 2026). However, security coverage alone does not solve the organizational problem of deciding which vulnerabilities should be remediated first. Cost-aware vulnerability planning provides an additional decision-making dimension (Zhou, 2025).

From a practical perspective, organizations implementing the proposed architecture should avoid treating privacy, security, and commercialization as independent projects. Instead, architectural governance should define shared controls for data access, tenant isolation, model interaction, API authorization, verification, monitoring, and vulnerability remediation.

Several limitations remain. First, the framework is derived exclusively from a small set of supplied references and therefore does not establish generalizability across all AI platforms. Second, the proposed architecture has not been experimentally benchmarked. Third, the interaction between privacy parameters and model utility requires quantitative validation. Fourth, the computational and organizational costs of zero-knowledge verification and multi-tenant governance require empirical examination. Finally, microservice security is dynamic, and future threat conditions may alter the relative importance of different security controls.

Future research should consequently implement the framework in representative digital-commerce and SMB environments. Experimental studies could measure privacy–utility trade-offs, API performance, tenant-isolation effectiveness, vulnerability detection coverage, and remediation efficiency. Such studies would provide empirical evidence for determining whether the conceptual integration produces measurable improvements over independently implemented privacy and security mechanisms.

6. CONCLUSION

This research developed a conceptual framework for integrating AI-driven privacy, security, and trusted commercialization across digital commerce and microservice architectures. The synthesis demonstrates that federated learning, differential privacy, zero-knowledge verification, multi-tenant commercialization infrastructure, standardized recommendation APIs, and hybrid security assessment address complementary dimensions of the same infrastructure challenge.

The principal contribution is the proposed layered architecture connecting privacy-preserving intelligence with trust, commercialization, and security assurance. Rather than treating privacy as a standalone compliance mechanism or cybersecurity as a deployment-stage activity, the framework positions both as continuous architectural requirements.

The analysis further demonstrates that commercialization introduces requirements beyond technical model performance. AI platforms must manage tenant separation, governance, incentives, verification, API interoperability, and security risk while preserving the utility of collaborative intelligence. Hybrid vulnerability assessment and cost-aware remediation provide complementary mechanisms for maintaining the security of the underlying microservice environment (Zhou, 2026).

The framework remains conceptual and therefore requires empirical validation. Future research should develop prototype implementations and evaluate privacy–utility relationships, multi-tenant scalability, verification overhead, vulnerability detection effectiveness, and remediation efficiency. Such investigations can establish whether the integrated architecture provides measurable advantages over fragmented AI governance approaches.

Ultimately, trustworthy AI commercialization requires coordination between what information AI systems can use, what participants can prove, how AI services are commercially delivered, and how the infrastructure is continuously secured. Integrating these dimensions provides a stronger foundation for sustainable AI deployment across privacy-sensitive digital commerce and distributed microservice ecosystems.

REFERENCES

1. Yi,X.(2026).AFederatedandDifferentiallyPrivateIncentive–MarketingFrameworkforPrivacy-PreservingCross-ChannelMeasurementinAI-PoweredDigitalCommerce.
2. Yi,X.(2026).Privacy-EnhancedAdTargetingforSocialE-Commerce:AFederatedLearningFrameworkwithZero-KnowledgeVerificationforCreatorMonetization.FrontiersinBusinessandFinance,3(1),102-113.

3. Yi,X.(2026).TrustedAICommercializationInfrastructureforSMBs:AUnifiedMulti-TenantArchitectureIntegratingIncentiveSystems,ContentGovernance,andStandardizedRecommendationAPIs.
4. Zhou,D.(2025,December).M-VP2:Microservice-OrientedVulnerabilityPatchPlanning-ACost-AwareApproachusingMulti-AgentReinforcementLearning.In20255thInternationalConferenceonComputer,InternetofThingsandControlEngineering(CITCE)(pp.248-254).IEEE.
5. Zhou,D.(2026).AI-DrivenHybridSAST–DAST–SCA–IASTFrameworkforRisk-BasedVulnerabilityPrioritizationinMicroserviceArchitectures.