



Editorial: Trustworthy AI Governance, Secure Edge Computing, and Reliable Model Adaptation

Hu Xiong^{1,*}

¹The Network and Data Security Key Lab of Sichuan Province, University of Electronic Science and Technology of China, Chengdu 610054, China

1 Introduction

Trustworthy digital systems depend on more than a single security mechanism. As enterprises adopt artificial intelligence (AI), edge services process data and computation near users, and organizations adapt large language models across distributed resources, reliability depends on how data, models, infrastructure, and incentives are governed together. A technically sound model can still fail when its data pipeline is insecure; a protected edge service can miss its deadline when verification consumes scarce processing capacity; and a completed training job can return an adapter that does not meet its requester's quality requirements.

The four papers in this issue examine these concerns at different layers. One develops a governance framework for secure AI-driven corporate transformation. Another places cryptographic verification within the resource budget of edge-cloud orchestration. A review then accounts for the costs and disclosures of federated large language model fine-tuning on constrained clients. The final paper studies validation, scheduling, and payment for

outsourced low-rank adaptation (LoRA) fine-tuning.

2 Papers in This Issue

In *Data Governance and Policy Support for Secure AI-Driven Corporate Digital Transformation*, Sun [1] proposes a unified data-and-AI governance framework. The paper links data classification, provenance, access control, privacy, and sharing to model validation, robustness, auditability, and accountability. It identifies tensions that arise when firms combine cloud, edge, Internet of Things (IoT), and digital-twin systems, including insecure data pipelines, opaque or adversarial models, and uneven compliance capacity between large firms and small and medium-sized enterprises (SMEs). Its proposed agenda combines technical safeguards with organizational routines and policy support. As a conceptual and policy-oriented study, its contribution is a framework for coordinating these measures, rather than an empirical estimate of their effectiveness.

Alam's *Paying for Trust at the Edge: Secure, Lightweight, Deadline-Aware Named Data Networking for Edge-Cloud Compute Orchestration* [2] examines a cost that is easy to omit from edge scheduling: verifying signed data consumes the same limited compute capacity used to run applications. The SLED-NDN framework models this coupling and combines



Submitted: 25 September 2026

Accepted: 27 September 2026

Published: 28 September 2026

Vol. 2, No. 3, 2026.

doi:10.62762/JRSC.2026.964149

*Corresponding author:

✉ Hu Xiong

xionghu@uestc.edu.cn

Citation

Xiong, H. (2026). Editorial: Trustworthy AI Governance, Secure Edge Computing, and Reliable Model Adaptation. *Journal of Reliable and Secure Computing*, 2(3), 233–235.



© 2026 by the Author. Published by Institute of Central Computation and Knowledge. This is an open access article under the CC BY license (<https://creativecommons.org/licenses/by/4.0/>).

verification-aware placement, signature amortization with cache revalidation, and elastic cloud bursting. In the evaluated configuration at 6,000 Interests per second, its deadline satisfaction ratio is 0.978, compared with 0.952 for the same datapath using per-packet RSA-2048 and 0.600 for vanilla NDN. These results illustrate that security overhead should be represented in service planning; their quantitative advantage is specific to the tested workload, platform, and trust model.

Shen et al. [3] review *Federated Fine-Tuning of Large Language Models on Resource-Constrained Clients: Technical Approaches, Resource Costs, and Applicability*. They organize the literature around four composable routes: parameter-efficient and quantized adaptation, backpropagation-free adaptation, proxy or submodel adaptation, and split federated adaptation. A common accounting framework follows the weights, training states, computation, and messages retained or exchanged at each endpoint. The review shows that reducing trainable parameters does not necessarily remove the storage cost of a resident base model, while split execution can move rather than eliminate resource demands. It also distinguishes the privacy exposure of adapter updates, boundary activations, logits, and cached embeddings. Its central lesson is to assess each technique against actual resource budgets, communication patterns, disclosed objects, and deployment conditions.

In *TrustCompute: Result Validation, Reputation-Aware Scheduling and Quality-Based Payments for Outsourced LoRA Fine-Tuning*, Lee et al. [4] connect assurance with incentives in a partially trusted compute-sharing platform. TrustCompute checks an adapter's structure and held-out quality, uses observed worker behavior when scheduling subsequent jobs, and records task events in an auditable ledger. Its evaluation includes synthetic workloads, real GPU fine-tuning, and a ledger deployment on a local Ethereum Virtual Machine (EVM). Across 440 validation cases, the reported checks reject all 400 constructed invalid submissions and accept all 40 valid ones; with 50% unreliable workers in the tested scheduling scenario, reputation-aware assignment completes all tasks and reduces failed-attempt time relative to round robin. The authors also examine identity resets and quality-based payments. Their validation establishes compliance with acceptance criteria, not proof that a worker followed a prescribed training procedure.

Viewed together, these papers address complementary

points in a trustworthy computing workflow. Governance defines the responsibilities and controls around data and AI systems; edge orchestration makes the cost of security visible; federated fine-tuning clarifies where resource savings are realized; and outsourced adaptation tests whether the returned result meets the requested standard. The papers use different methods and make claims at different levels of generality, but each treats trust as something that must be designed into the operational system.

3 Closing Remarks

Reliable and secure computing requires decisions that span architecture, measurement, and governance. The contributions introduced here offer practical ways to reason about those decisions: specify accountable data and model processes, budget explicitly for verification, trace resources across federated clients and servers, and validate outsourced results while accounting for participant incentives. Further work can build on these directions by testing their assumptions across additional organizations, platforms, workloads, and threat models.

We thank the authors for their contributions and the reviewers and editorial team for their work on the manuscripts. We hope this issue helps researchers and practitioners design AI and distributed computing systems whose security and reliability can be examined alongside their performance.

Data Availability Statement

Not applicable.

Funding

This work was supported without any funding.

Conflicts of Interest

Hu Xiong served as a Co-Editor-in-Chief of the *Journal of Reliable and Secure Computing* at the time of manuscript submission. To ensure the integrity of the peer-review process, Hu Xiong was not involved in the editorial handling, peer review, or decision-making process for this manuscript, which was handled independently by another editor.

AI Use Statement

The author declares that no generative AI was used in the preparation of this manuscript.

Ethical Approval and Consent to Participate

Not applicable.

References

- [1] Sun, F. (2026). Data Governance and Policy Support for Secure AI-Driven Corporate Digital Transformation. *Journal of Reliable and Secure Computing*, 2(3), 164-178. [[CrossRef](#)]
- [2] Alam, M. (2026). Paying for Trust at the Edge: Secure, Lightweight, Deadline-Aware Named Data Networking for Edge-Cloud Compute Orchestration. *Journal of Reliable and Secure Computing*, 2(3), 212-232. [[CrossRef](#)]
- [3] Shen, Y., Liu, X., Yang, Y., & Lv, H. (2026). Federated Fine-Tuning of Large Language Models on Resource-Constrained Clients: Technical Approaches, Resource Costs, and Applicability. *Journal of Reliable and Secure Computing*, 2(3), 194-211. [[CrossRef](#)]
- [4] Lee, L. F., Chang, Y. Y., & Yeh, K. H. (2026). TrustCompute: Result Validation, Reputation-Aware Scheduling and Quality-Based Payments for Outsourced LoRA Fine-Tuning. *Journal of Reliable and Secure Computing*, 2(3), 179-193. [[CrossRef](#)]

Hu Xiong Co-Editor-in-Chief of the *Journal of Reliable and Secure Computing*. (Email: xionghu@uestc.edu.cn)