Final Report on Securing Federated Learning Systems: A Case Study of Tokenized Incentive Mechanisms

1. Introduction

This final report outlines the accomplishments and ongoing efforts related to the research project on securing federated learning systems, specifically focusing on the vulnerabilities in tokenized incentive mechanisms. The project explores model-poisoning attacks using Gaussian noise and proposes a blockchain-based auditing protocol to mitigate these threats. Over the past few months, significant milestones have been achieved, including the submission of findings to a prestigious conference and plans for further dissemination of the research outcomes.

2. Research Findings and Paper Submission

The research findings from this project were compiled into an eight-page manuscript and submitted to the ACM Web Conference 2025, a top-tier, A*-ranked conference in computer science. The paper detailed the attack methodology, its impact on federated learning models, and the proposed blockchain-based auditing solution. While the submission was not accepted in this review cycle, the reviewers acknowledged the novelty of the research and its relevance to addressing a critical security challenge in federated learning systems.

The meta-review highlighted the paper's strengths, including the demonstration of the attack's effectiveness and the auditing system's efficiency. However, concerns were raised regarding the novelty of using Gaussian noise in federated learning, the depth of experimental evaluation, and the justification for the blockchain component. The reviewers provided constructive feedback, suggesting improvements in technical contributions and a more comprehensive evaluation. These insights will be instrumental in refining the research for future submissions to high-impact venues.

Paper Decision

Decision by Program Chairs 20 Jan 2025, 07:25 (modified: 23 Jan 2025, 03:52) Program Chairs, Authors

Revisions

Decision: Reject

Comment:

Paper metareview: This paper presents an attack on federated learning incentive mechanisms using Gaussian noise, along with a blockchain-based auditing solution. While it demonstrates the attack's effectiveness and audit system's efficiency, major concerns exist about novelty (as Gaussian noise in FL is well-studied), limited experimental evaluation, and unclear blockchain component justification. The work addresses an important security challenge but needs stronger technical contributions and more comprehensive evaluation.

3. Research Contributions and Impact

The project successfully demonstrated the feasibility of model-poisoning attacks in federated learning incentive mechanisms and proposed a blockchain-based auditing protocol to counteract these threats. Key contributions include:

- A novel evaluation of how Gaussian noise can compromise tokenized incentive mechanisms in federated learning.
- An auditing framework leveraging blockchain technology to detect and mitigate model-poisoning attacks.
- Experimental validation of the proposed approach, with insights into its effectiveness and potential areas for improvement.

These findings contribute to advancing security in federated learning and provide a foundation for further exploration in secure distributed computing.

4. Future Plans for Research Dissemination

In response to the reviewer feedback, future efforts will focus on enhancing the technical contributions by exploring additional attack models, refining the blockchain auditing mechanism, and expanding the experimental evaluation with diverse datasets and aggregation algorithms. These improvements will strengthen the research's impact and address the concerns raised during the review process.

Furthermore, the results of this project will be presented at the Faculty Research Symposium 2024, scheduled for April 2025. This event will provide an opportunity to engage with faculty members, receive additional feedback, and explore potential collaborations to extend the research further.

5. Acknowledgment

This research was made possible through the support of the Office of Faculty Affairs. Their funding and resources have facilitated the investigation of critical security challenges in federated learning, contributing to the advancement of cybersecurity research at Lincoln University. Continued support for such initiatives will play a vital role in fostering innovation and addressing emerging threats in distributed computing environments.

6. Conclusion

This project has successfully investigated security vulnerabilities in federated learning incentive mechanisms and proposed an innovative blockchain-based solution. Although the initial submission to the ACM Web Conference 2025 was not accepted, the constructive feedback received has provided a clear roadmap for refining the research. Efforts will be directed towards strengthening technical contributions, expanding evaluations, and preparing for future dissemination at academic conferences and faculty research events. The support from the Office of Faculty Affairs has been invaluable in achieving these milestones, and further developments in this research will contribute significantly to both the academic and cybersecurity communities.