

Adaptive Trust-Aware Post-Quantum Secure Federated Learning with Dynamic Differential Privacy

Nikita S. Shinde ^{1,*}, Nita Thakare²

Abstract

FL allows multiple users to train the machine learning model and ensure that they cannot share their raw data. Post-quantum secure aggregation and differential privacy for FL are introduced by the Beskar Framework; it assumes that all selected nodes will perform reliably. Due to hardware faults, network failures, or malicious intent, some of the assisting nodes may perform unreliably while we are developing it in the real-world. To assist the nodes, this paper proposes a framework known as Beskar-AT (Adaptive Trust), which provides a trust scoring mechanism. The model observes the recent activities of each node and then assigns a trust value to each of the assisting nodes. Low-score nodes are temporarily removed from the aggregation process and replaced with a rotating pool. The proposed model automatically changes the privacy level based on the actual number of trusted nodes participating in the process. These two combined features help to improve the reliability and the privacy-utility balance of post-quantum secure FL. With the experiment of MNIST, CIFAR-10, and CIFAR-100 show that the system named Beskar-AT maintains high model accuracy and lower noise overhead compared to the original Beskar framework, especially when some of the assisting nodes behave irregularly.

Keywords: Federated learning, post-quantum cryptography, dynamic differential privacy, trust-based management, secure model aggregation, adaptive noise mechanism

INTRODUCTION

Federated learning (FL) is a machine learning approach where multiple clients send the model updates to a central aggregation server, and each of them stores their data locally. It avoids sharing sensitive data directly, which is important in areas such as finance, healthcare, and mobile applications. This research has shown that sharing model updates can leak private information through techniques such as gradient inversion and membership inference attacks [1–2].

These risks are addressed in the Baskar framework [3] proposed two important ideas: (i) secure aggregation follows the post-quantum standards using lattice-based cryptographic tools defined by the NIST standards, where CRYSTALS-Kyber and CRYSTALS-Dilithium are used for the verification of digital signatures; and (ii) differential privacy (DP): both work at the local and central levels to protect against various types of adversaries. Also, Baskar provides precomputation techniques to decrease the overhead of signing and makes it suitable for mobile or resource-constrained clients. Due to this integrated features Baskar becomes one of the most complete privacies preserving FL frameworks available in today's world.

*Author for Correspondence

Nikita S. Shinde
E-mail: nikitashinde16798@gmail.com

¹Student, Department of Computer Engineering (AIDS), Guru Gobind Singh College of Engineering and Research Centre, Jawahar Education Society's, JIT, Nashik, Maharashtra, India

²Professor, Department of Computer Engineering, Guru Gobind Singh College of Engineering and Research Centre, Jawahar Education Society's, JIT, Nashik, Maharashtra, India

Received Date: May 05, 2026

Accepted Date: July 14, 2026

Published Date: July 22, 2026

Citation: Nikita S. Shinde, Nita Thakare. Adaptive Trust-Aware Post-Quantum Secure Federated Learning with Dynamic Differential Privacy. Journal of Advanced Database Management & Systems. 2026; 13(2): 38–44p.

Once the assistance nodes are selected, it is assumed that they continue to behave correctly throughout the entire training process; this is one of the important limitations of Baskar. In a practical federated learning environment, some assisting nodes may fail to behave reliably. During the training process, node might lose internet connectivity, faces hardware failure, or sometimes it may be affected by partial security compromise. The entire aggregation process can be disrupted or may leak privacy, whenever the assisting nodes fail or misbehave.

This study presents an advanced version of Baskar, known as Baskar-AT, which includes an adaptive t mechanism. The main features of Baskar-AT are summarized below.

- A lightweight trust scoring mechanism is introduced for the assisting nodes, in which the trust value of each node is calculated by analyzing its history of participation and the correctness of its contributions during previous training process rounds.
- A dynamic node rotation standard is implemented to handle unreliable nodes. Whenever a node with a low trust score receives a request, that node is replaced with another node from the backup pool for the smooth continuation of the aggregation process.
- A dynamic differential privacy (DDP) budget control method is designed according to the number of currently trusted and active assisting nodes and the amount of noise added to update the model.
- The proposed study was evaluated using detailed experiments on multiple benchmark datasets, in which the performance of Baskar-AT was compared with that of the original Baskar by considering different conditions of node reliability.

The remainder of this paper is organized as follows. Section II explains the related work. Section III describes the system and threat models. Section IV provides a detailed explanation of the Baskar-AT framework. Section V discusses the security analysis. Section VI presents the experimental results, and Section VII concludes the paper.

RELATED WORK

Privacy-preserving federated learning has been researched and then the approaches are divided into two categories as secure aggregation and differential privacy. The first secure aggregation protocol, which allows the server to find the sum of client updates without looking at the updates of individuals, was introduced by Bonawitz et al. [4]. Their approach is pairwise masking and handles dropouts by sharing secrets. This was improved by Bell et al. [5] using a logarithmic-overhead approach with graph-based secret sharing. Ma et al. [6] introduced Flamingo. Flamingo is a multi-round protocol with strong dropout resilience using distributed key generation. To help unmask aggregated gradients, an efficient one-round protocol using assisting nodes, that is presented by Behnia et al. [7] named as e-Seal.

Yang et al. [8] introduced a post-quantum secure aggregation (PQSA), though it was unable to combine differential privacy and has limited scalability. Therefore, both post-quantum security and differential privacy were addressed by the Baskar framework [3], with the created threat model for FL that helps researchers understand the weaknesses of the system and design a better security model for the protection of user data. The system needs to do less work during practical execution because Dilithium signing and mask generation model were precomputed.

In FL. Differential privacy has been applied to both central and local settings [9-11]. DP-SGD was introduced by Abadi et al. [12] which helps to add calibrated noise to the gradients during training process. Some of the work also studied the trade-off between the model accuracy and privacy budget under different FL settings [13]. However, all these models are unable to consider the situation where the noise level should be adjusted based on the reliability of the participating infrastructure, such as assisting nodes. In distributed systems, trust management and peer-to-peer networks have been explored in multiple forms, but in previous studies, their application to the specific role of assisting nodes in post-quantum FL has not been studied. Our study helps bridge this gap [14-17].

SYSTEM AND THREAT MODEL

System Model

Instead of creating completely new model, we use the same system model created by Baskar [3], which consists of three parts:

1. *Federated learning clients*: These are the users, each of whom has their own local dataset and trains its own local model on the data.
2. *One aggregation server*: This server collects the masked updates and computes a global model; the server keeps the client's data secure as the server cannot see their individual data.
3. *k assisting nodes*: It holds the partially shared secrets with users/clients and helps the server unmask the aggregated gradient by not revealing individual client updates.

In addition, from Baskar components, Baskar-AT proposed a module named Trust Manager module that runs on the aggregation server. This module helps track the trust score of each assisting node and maintain a backup pool for assisting nodes with lower trust scores. Based on their observed behaviors, the trust score is updated after each round of the training process [18].

The same one-time key generation and mask precomputation are followed in the setup phase as in Baskar. The aggregation phase is enhanced so the server evaluates each of the assisting nodes and updates their scores after each round. If the score of the observed nodes falls below the threshold, that node is replaced with the node having the highest score from the backup pool [19].

Threat Model

We use the same three threat models used in Baskar [3], which cover adversaries targeting user gradients (TM1), intermediate models (TM2), and the third deployed model (TM3). In advance, Baskar-AT considers a fourth threat: the nodes that behave incorrectly, which are not malicious but are detectable, for example, repeatedly failing to respond or contributing incorrect mask sums. We call this the assisting node scenario.

Attackers assume that their quantum computers are very powerful, which can be easy to break by normal cryptography methods. Between the servers, clients and the assisting nodes, all the communication is authenticated, so it is not easy to exchange messages by outsider. Even if some nodes are compromised, from observable data they try to extract the private information during process [20].

BESKAR-AT FRAMEWORK

Trust Scoring for Assisting Nodes

At the start of training each assisting node A_j is assigned a trust score TS_j that is initialized to 1.0. Server then updates the trust score of each assisting nodes after each round of aggregation t , according to the following rule:

$$TS_j(t+1) = \alpha \cdot TS_j(t) + (1 - \alpha) \cdot R_j(t),$$

where α describes the value between 0 and 1 that controls how quickly the system forgets the past behavior, and $R_j(t)$ represents the reliability score for round t . If the node's contribution passes all the verification checks in that round, then the reliability score $R_j(t)$ is set to 1.0 (i.e., it clarifies that its signature is valid, the mask sum is consistent, and it responded within the expected time). If any check fails, its value is set to 0.0 [16].

Whenever TS_j falls below a threshold τ (e.g., $\tau = 0.5$), the active assisting node's set that has a node with a low threshold is removed and replaced by the high-scoring node from the backup pool. This ensures that a single unreliable node cannot interrupt the aggregation process [17].

Dynamic Differential Privacy Budget

For the entire training process, the privacy budget ϵ is fixed in the original Baskar framework. This indicates that the same noise level is applied even when more trusted assistance nodes are available. When there are fewer active trusted nodes, the fixed noise level may not provide more protection [14].

Baskar-AT proposed a dynamic privacy budget controller that helps to adjust the effective noise multiplier $\sigma(t)$ at each round t based on the trust-weighted participation ratio:

$$P(t) = (\sum_j TS_j(t)) / k$$

where the sum is taken from all k assisting nodes that are currently active. This indicates that the calculation considers the contribution of every active assisting node. Based on the total value, the multiplier of noise for round t is determined as

$$\sigma(t) = \sigma_base / P(t)$$

where σ_base is the baseline multiplier. A high $P(t)$ indicates that all nodes are trusted, and the noise is maintained near the baseline. Low $P(t)$ indicates that several nodes are unreliable, and the noise is automatically increased to maintain privacy. This approach represents that when the aggregation infrastructure is healthy, the model does not over-inject noise in the rounds, so it directly improves the accuracy of the model [15].

Integration with Post-Quantum Secure Aggregation

With Baskar's post-quantum secure aggregation, the trust scoring and dynamic DP are fully matched. The Dilithium signatures and CRYSTALS-Kyber key encapsulations are used as is, without any modification, with the server-side trust manager operating as an independent layer, and there is no need to make any changes to client-side operations.

The Baskar-AT aggregation round at iteration t is summarized in Algorithm 1. The main additional thing done compared to Baskar's Algorithm 2 is the trust score update step (line 6), checking of node replacement (lines 7–9), and noise adjustment (line 12).

Algorithm 1: Baskar-AT Aggregation Round (Iteration T)

Input: Iteration t , precomputed masks MS , trust scores TS , backup pool BP Output: Global model update wt .

- 1: Clients compute local updates, mask and send to server (as in Baskar)
- 2: Assisting nodes send aggregated mask sums to server (as in Baskar)
- 3: Server verifies signatures from all clients and assisting nodes
- 4: Server unmask and computes global update wt
- 5: For each assisting node A_j :
- 6: Compute $R_j(t)$; Update $TS_j(t+1) = \alpha \cdot TS_j(t) + (1-\alpha) \cdot R_j(t)$
- 7: If $TS_j(t+1) < \tau$:
- 8: Remove A_j from active set; Replace with top node from BP
- 9: Update shared secrets for new node (offline pre-keying)
- 10: Compute $P(t) = (\sum TS_j(t+1)) / k$ over active nodes
- 11: $\sigma(t+1) = \sigma_base / P(t)$
- 12: Apply DP noise with $\sigma(t+1)$ to wt for next round

Security Analysis

All the post-quantum security guarantees are inherited from Baskar in Baskar-AT, and the CRYSTALS-Kyber key encapsulation and CRYSTALS-Dilithium signatures are not changed. Based on the Module-LWE and Module-SIS hardness assumptions still apply directly in the security proof of Baskar [3].

Trust Score Integrity

The entire trust score is calculated from the verified and signed messages. All messages from the assisting nodes are signed by Dilithium under the EU-CMA security. Malicious assisting nodes are unable to build a valid response to appear reliable. As part of the protocol, the trust score is computed by an honest-but-curious server. The LDP component addresses the complete malicious server before any interaction, which is client-side.

Dynamic DP Privacy Guarantee

Because $P(t) < 1$, the dynamic noise multiplier $\sigma(t)$ is always greater than or equal to σ_{base} . This indicates that below the baseline level, the dynamic controller cannot reduce privacy protection. At $(P(t)=1,0)$, the noise is exactly equal to the baseline, which is a degraded case, where the noise is increased and provides a stronger guarantee. Therefore, as in the original Baskar framework over T rounds, the cumulative privacy budget satisfies the same (ϵ, δ) -DP under the Rényi differential privacy accounting method [15].

Node Replacement Security

From the setup phase, the nodes that have completed the full key generation and mask computation before the replacement nodes are selected from the backup pool. Before the next round starts, replacement occurs so there are no mid-round secrets used again. Due to the independence of the masking vectors, the privacy of the previous round is not affected by node replacement.

PERFORMANCE EVALUATION

Experimental Setup

We implemented Baskar-AT using the available Baskar codebase [3] for secure aggregation and the Flower framework for federated model training with Porch Pacus for differential privacy. We built the system using the ABIDES simulation framework. As implemented in Baskar, we used the same CRYSTALS-Kyber and CRYSTALS-Dilithium algorithms. All our experiments were run on the same hardware system: A Linux machine with an AMD Ryzen Threadripper PRO 5965WX (24 cores); for aggregation experiments, a machine with 256 GB RAM; and for model training, an NVIDIA RTX 6000 Ada GPU was used. We tested the model performance on three datasets: NIST (lightweight MLP), CIFAR-10 (ResNet-18), and CIFAR-100 (ResNet-18). We ran 50 communication rounds with five clients for each dataset. We introduced a scenario in which some assisting nodes fail randomly with a failure probability per round $pf \in \{0, 0.1, 0.2, 0.3\}$ to simulate node unreliability. Hence, we compared Baskar-AT with the original Baskar under both CDP and LDP settings with $\epsilon = 10$.

Model Accuracy Under Node Unreliability

As shown in Table 1, Baskar-AT achieves an accuracy nearly identical to Baskar when all assisting nodes are reliable ($pf = 0$), confirming that the performance is not affected by trust overhead under normal conditions. At $pf = 0.2$, Baskar-AT performs significantly better than Baskar by a notable margin. This improvement occurs from two sources: (1) unreliable nodes are replaced before they cause errors during aggregation, and (2) the over-injection of noise is reduced by the dynamic noise controller in rounds where the active nodes are highly trusted.

Table 1. Model accuracy comparison—Baskar versus Baskar-AT ($E = 10, 50$ rounds).

Dataset/failure rate	Baskar CDP $pf=0$	Baskar- AT CDP $pf=0$	Baskar CDP $pf=0.2$	Baskar-AT CDP $pf=0.2$	Baskar LDP $pf=0.2$	Baskar-AT LDP $pf=0.2$
MNIST	0.87	0.88	0.71	0.85	0.78	0.79
CIFAR-10	0.77	0.78	0.59	0.74	0.20	0.21
CIFAR-100	0.23	0.24	0.11	0.20	0.02	0.03

Table 2. Noise multiplier comparison—fixed versus dynamic budget (MNIST, $E=10$).

Failure rate pf	Fixed σ (Baskar)	Avg dynamic σ (Baskar-AT)	Accuracy gain
0.0	3.2	3.2	0.00
0.1	3.2	3.1	+0.04
0.2	3.2	2.9	+0.14
0.3	3.2	2.7	+0.18

Effect of Dynamic vs. Fixed Privacy Budget

Table 2 shows that the average noise multiplier in rounds was reduced by the dynamic budget controller, where most of the nodes are trusted, which usually causes the few of node failure. This reduction directly helps with accurate improvements. At higher failure rate gain is most pronounced, shows that the adaptive mechanism is most beneficial precisely when it is mostly needed.

Computational Overhead of Trust Scoring

The trust scoring calculation includes one floating-point multiplication and one addition for each assisting node in every round. In our work with $k = 3$ assisting nodes, this added less than 0.1 ms per round on the server, which is very small compared to the overall aggregation time. If a node needs to be replaced, it adds approximately 12 ms when triggered; however, this does not happen often, so the overall affects the performance and reduces the performance.

DISCUSSION

The assisting nodes remain reliable throughout training this practical gap in post-quantum federated learning is addressed by Baskar-AT. In real-world settings, this assumption is not easy to guarantee, especially when the nodes are operated by multiple organizations or when the assisting nodes are volunteer devices. With existing post-quantum security proofs, the trust scoring mechanism is verifiable, lightweight, and consistent.

The current system has one limitation: the backup pool must be established during the setup. Managing a large backup pool could add difficulties to the setup phase in the case of very large-scale deployments. Using the same precomputation techniques used in Baskar, the enrollment of backup nodes could be explored in future work. Another method is to apply trust scoring to clients and updates on the global model. The dynamic DP component ensures that the noise level is not reduced below the baseline. This guarantees that the original privacy of Baskar is always maintained. Regulators specify a maximum noise in settings; the controller could modify to cap the noise instead of to scale it indefinitely. the controller

CONCLUSION

This paper presents Baskar-AT, an advanced version of the Baskar post-quantum secure federated learning framework. Baskar-AT can solve the problem of assisting node failure and does not work properly. Baskar-AT proposed a trust scoring mechanism to overcome the problem of unreliable nodes. This trust scoring mechanism uses dynamic differential privacy budget controller to adjust noise levels based on the trust-weighted participated assisting nodes. Both extensions are implemented as lightweight server-side extensions to reduce the burden on clients. From experimental results, we conclude that Baskar-AT provides higher accuracy of models even when some assisting nodes fail. The model also reduces unnecessary noise injections from functioning rounds. The invented framework provides all the post-quantum security features inherited from Baskar and provides an adaptive solution for privacy-preserving federated learning in the post-quantum era.

Acknowledgment

I am thankful to the research community for making the Baskar codebase publicly available. Also, I would like to thank my guide and my institution for their support and guidance in completing this research work.

REFERENCES

1. Carlini N, Tramèr F, Wallace E, Jagielski M, Herbert-Voss A, Lee K, et al. Extracting training data from large language models. In: Proceedings of the 30th USENIX Security Symposium (USENIX Security 21); 2021. p. 2633-2650.
2. Shokri R, Stronati M, Song C, Shmatikov V. Membership inference attacks against machine learning models. In: 2017 IEEE Symposium on Security and Privacy (SP); 2017. p. 3-18. doi:10.1109/SP.2017.41.

3. Zhang Y, Behnia R, Yavuz AA, Ebrahimi R, Bertino E. Efficient full-stack private federated deep learning with post-quantum security. *IEEE Trans Dependable Secure Comput.* 2025;22(5):5567-5583. doi:10.1109/TDSC.2025.3568704.
4. Segal A, Marcedone A, Kreuter B, Ramage D, McMahan HB, Seth K, et al. Practical secure aggregation for privacy-preserving machine learning. In: *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*; 2017. p. 1175-1191. doi:10.1145/3133956.3133982.
5. Bell JH, Bonawitz KA, Gascón A, Lepoint T, Raykova M. Secure single-server aggregation with (poly)logarithmic overhead. In: *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*; 2020. p. 1253-1269. doi:10.1145/3372297.3417885.
6. Ma Y, Woods J, Angel S, Polychroniadou A, Rabin T. Flamingo: multi-round single-server secure aggregation with applications to private federated learning. In: *2023 IEEE Symposium on Security and Privacy (SP)*; 2023. p. 477-496. doi:10.1109/SP46215.2023.10179434.
7. Behnia R, Riasi A, Ebrahimi R, Chow SSM, Padmanabhan B, Hoang T. Efficient secure aggregation for privacy-preserving federated machine learning. In: *2024 Annual Computer Security Applications Conference (ACSAC)*; 2024. p. 778-793. doi:10.1109/ACSAC63791.2024.00069.
8. Yang S, Chen Y, Tu S, Yang Z. A post-quantum secure aggregation for federated learning. In: *Proceedings of the 2022 12th International Conference on Communication and Network Security*; 2022. p. 117-124. doi:10.1145/3586102.3586120.
9. Chamikara MAP, Liu D, Camtepe S, Nepal S, Grobler M, Bertok P, et al. Local differential privacy for federated learning [Preprint]. arXiv:2202.06053; 2022. doi:10.48550/arXiv.2202.06053.
10. Truex S, Liu L, Chow KH, Gursoy ME, Wei W. LDP-Fed: federated learning with local differential privacy. In: *Proceedings of the Third ACM International Workshop on Edge Systems, Analytics and Networking*; 2020. p. 61-66. doi:10.1145/3378679.3394533.
11. Kairouz P, McMahan HB, et al. Advances and open problems in federated learning. *Found Trends Mach Learn.* 2021;14(1-2):1-210. doi:10.1561/22000000083.
12. Abadi M, Chu A, Goodfellow I, McMahan HB, Mironov I, Talwar K, et al. Deep learning with differential privacy. In: *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*; 2016. p. 308-318. doi:10.1145/2976749.2978318.
13. Yang Y, Hui B, Yuan H, Gong N, Cao Y. PrivateFL: accurate, differentially private federated learning via personalized data transformation. In: *Proceedings of the 32nd USENIX Security Symposium (USENIX Security 23)*; 2023. p. 1595-1612.
14. Jøsang A, Ismail R, Boyd C. A survey of trust and reputation systems for online service provision. *Decis Support Syst.* 2007;43(2):618-644. doi:10.1016/j.dss.2005.05.019.
15. Mironov I. Rényi differential privacy. In: *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*; 2017. p. 263-275. doi:10.1109/CSF.2017.11.
16. Ducas L, Kiltz E, Lepoint T, Lyubashevsky V, Schwabe P, Seiler G, et al. CRYSTALS-Dilithium: a lattice-based digital signature scheme. *IACR Trans Cryptogr Hardw Embed Syst.* 2018;2018(1):238-268. doi:10.13154/tches.v2018.i1.238-268.
17. Bos J, Ducas L, Kiltz E, Lepoint T, Lyubashevsky V, Schanck JM, et al. CRYSTALS-Kyber: a CCA-secure module-lattice-based KEM. In: *2018 IEEE European Symposium on Security and Privacy (EuroS&P)*; 2018. p. 353-367. doi:10.1109/EuroSP.2018.00032.
18. Beutel DJ, Topal T, Mathur A, Qiu X, Fernandez-Marques J, Gao Y, et al. Flower: a friendly federated learning research framework [Preprint]. arXiv:2007.14390; 2020. doi:10.48550/arXiv.2007.14390.
19. Dwork C. Differential privacy: a survey of results. In: Agrawal M, Du D, Duan Z, Li A, editors. *Theory and Applications of Models of Computation*. Berlin: Springer; 2008. p. 1-19. doi:10.1007/978-3-540-79228-4_1.
20. Di Crescenzo G. Sharing one secret vs. sharing many secrets. *Theor Comput Sci.* 2003;295(1-3):123-140. doi:10.1016/S0304-3975(02)00399-7.