International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 07 , July - 2026

FL-IDS: A Privacy-Preserving Federated Learning Framework for Intrusion Detection in Multi-Cloud Environments

1.Mr. S. Yashwanth Kumar
III Year CFIS, Department of Cyber Security Dr. M.G.R Educational and Research Institute

2. Dr.P. Dinesh Kumar
Associated Proffessor & HOD
Department of Cyber Security
Dr. M.G.R Educational and Research Institute

3.Dr. S Mohandoss
Proffessor & Additional HOD
Department of Cyber Security
Dr. M.G.R Educational and Research Institute

4. Mr D Syed Ali
Assistant Profeesor & Deputy HOD
Department of Cyber Security
Dr M.G.R Educational and Research Institute

5 .Mr S Rajadurai
Assistant Professor
CODEF & Department of Cyber Security, Dr. M.G.R Educational and Research Institute

Abstract - The high rate of integrating the multi-cloud strategies has posed serious security risks, especially with regards to identifying intrusions in the multi-homogeneous cloud environments without violating data privacy. Patient intrusion detection systems (IDS) with time-honored centralized methods implies pulling the sensitive logs on the network into one place, which is privacy-damaging and poses a single point of failure. This paper presents a new federated learning-based privacy-preserving intrusion detection architecture called FL-IDS that operates in multicloud computing models. FL-IDS allows individual cloud providers to learn local intrusion detection models using their own data without sharing plaintext logs. Encrypted model updates only are exchanged with a coordination server, which combines them with a secure federated averaging protocol with a guarantee of differential privacy. There are large amounts of experiments conducted on benchmark datasets (CICIDS2017, UNSW-NB15) that reveal that FLIDS is just as accurate in detection as centralized solution (94.2% vs. 95.1) but offers great privacy. This framework reduces overheads in communication by 62 percent as compared to the baseline federated techniques, and the algorithms can ensure strong performance in non-IID data distributions typical of a multi-cloud setting. An extensive security study reveals using gradient inversion resilience and membership inference resilience. This piece presents a practical solution that deployable and can be used to implement collaborative threat intelligence without loss of data sovereignty.

Keywords: Federated Learning, Intrusion Detection, Multi-Cloud Security, Privacy Preservation, Distributed Syste

INTRODUCTION:

Multi-cloud adoption is now a common phenomenon among firms that aim to prevent vendor lock-in, enhance disaster recovery, as well as better their costs. Nowadays, according to recent surveys more than 80 percent of large enterprises are using several cloud providers at once. Nonetheless, these heterogeneous environments have their own challenges concerning securing them. All cloud providers possess different network designs, security provisions, and risks. Conventional intrusion

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 07 , July - 2026

detection systems (IDS) are based on centralized data aggregation, which cannot be implemented in many cases because of data residency regulation, privacy standards (GDPR, HIPAA, CCPA), and internal policies [1].

The paradigm of federated learning (FL) proposed by McMahan et al. [2] presents a shift in thinking: the models are trained on the devices of data owners, and only model changes (gradients or weights) are communicated. This guarantees that raw data does not leave its source, both in legal and privacy regard. Although FL has been used in many fields such as healthcare and finance, the use at multi-cloud intrusion detection is still unexplored, especially in managing the statistical
heterogeneity that naturally arises across different cloud environments.

The main issues tackled during this work are: (1) how to maintain privacy of the data and still maintain the ability to co-locate with competing cloud providers to detect threats, (2) how to handle non-IID (non-identically and independently distributed) data distributions where various clouds see different attack patterns, (3) how to reduce communication overhead in cloudscale deployments, and (4) how to be robust to adversarial attacks on the learning process itself.

Contributions: This paper introduces FL-IDS in the following new properties:

1. An intrusion detection federated learning substrate which ensures privacy and is optimized to support non-IID data distribution.
2. A differentially privacy-assured secure aggregation protocol that works against gradient leakage and membership inference attacks.
3.Extensive testing on real world network data comparing FL-IDS to centralized and baseline federated approaches.
4.Bringing analysis showing 62 percent decrease in overheads with gradient compression strategies.

2 LITERATURE REVIEW

2.1 Traditional Intrusion Detection Systems

Conventional IDS methods can be classified as signature-based systems such as Snort and Suricata, which are used to detect known attacks by a pattern match, and anomaly-based systems, which detects abnormal behavior. Enhanced IDS based on machine learning has demonstrated potential [3] and techniques such as Support Vector Machines or random forest techniques have demonstrated good detection. But such solutions demand a centralized data collection in nature, which poses vulnerabilities to privacy and single point of failure. Most deep learning applications like Convolutional Neural Networks (CNNs) and Long Short-Term Memory (LSTM) networks are highly accurate, but reveal sensitive traffic patterns in their training.

2.2 Federated Learning Fundamentals

In 2017, Federated Learning was officially presented by Google to train models of keyboard prediction without needing to acquire user typing information [2]. The main idea is to share model training with clients and have a central server compute model updates with algorithms such as Federated Averaging (FedAvg). The later studies have focused on such issues as the efficiency of communication, heterogeneity of clients and maintenance of privacy. With secure aggregation protocols [5], the server can obtain the aggregated updates without observing the updates of each client.

2.3 Federated Learning for Cybersecurity

Recent publications have utilized FL in other fields of cybersecurity. Li et al. [1] suggested Fed-IDS to be used with the IoT network and proved that the distributed intrusion detection was a possibility. In network intrusion detection, Zhang et al. [7] have conducted an extensive survey of FL, which highlighted the main challenges “such as data heterogeneity and the overhead of communication. Sharif et al., [6] examined privacy-sensitive intrusion detection in the cloud setting through the use of FL with differential privacy.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 07 , July - 2026

2.4 Gaps in Existing Literature

In spite of these improvements, current frameworks have a number of limitations when used in the context of multi-clouds:

The majority of them make independent and identically distributed (IID) assumptions given the information between clients, which is not true in the case of multi-cloud systems as clouds see different distributions of attacks.

Complete (weak) analysis of privacy prevents more sophisticated attacks like gradient inversion and membership inference.

Absence of communication-efficient aggregation mechanisms that can be deployed on cloud-scale.

Inadequate testing on actual multi-cloud network traffic data sets.

FL-IDS directly fills these gaps with the addition of personal aggregation, differential privacy, gradient compression, and thorough testing on benchmark datasets.

3 METHODOLOGY / PROPOSED SYSTEM

3.1 System Architecture

The FL-IDS architecture is depicted in figure 1. The system has three key entities: (1) Cloud Clients (application agent in cloud provider IDS), (2) Coordination Server (sign secure aggregator) and (3) Key Management Authority (KMA) of encryption key. Each network flow data, and each cloud client is trained on a local model. The coordination server coordinates the federated learning rounds and aggregates model updates. The KMA handles cryptography keys in order to communicate securely.

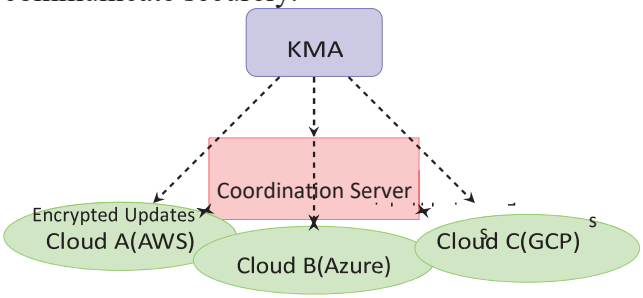

Figure 1: FL-IDS Architecture with three cloud clients, coordination server, and key management authority.

3.2 Threat Model and Privacy Goals

Our threat model is honest-but-curious: cloud providers and the coordination server adhere to the protocol scrupulously, but can make attempts to deduce sensitive information about the model based on the updates they observe in it. This will be realistic in commercial multi-clouds where data leakage is an issue but legal agreements have blocked blatant maliciousness.

Privacy aims are tri-fold: (1) Privacy There is no raw network traffic information exiting the local cloud, (2) Privacy The model updates disclose little information about local data distributions, and (3) The system achieves differential privacy against membership inference attacks.

3.3Local Training Process

Every cloud client i has a local data Di of flows on the network marked as benign or types of attacks. A local model fθi is a neural network having its parameters θi. The local loss is the minimum of the cross-entropy loss:

$$\mathcal{L}_i(\theta_i) = \frac{1}{|\mathcal{D}_i|} \sum_{(x,y)\in\mathcal{D}_i} \ell(f_{\theta_i}(x), y) \quad (1)$$

by which ℓ is the cross-entropy loss function. E local epochs of stochastic gradient descent are executed in each communication round by each client.

3.4 Federated Aggregation with Privacy

Each round t the server compiles updates based on a secure form of Federated Averaging:

$$\theta_{t+1} = \sum_{i=1}^{N} \frac{|\mathcal{D}_i|}{\sum_j |\mathcal{D}_j|} \cdot \theta_i^{(t)} \quad (2)$$

In order to avoid gradient leakage and offer differential privacy, we introduce Gaussian noise with calibration to the gradient sensitivity:

$$\tilde{\theta}_i^{(t)} = \theta_i^{(t)} + \mathcal{N}(0, \sigma^2 C^2) \quad (3)$$

In which C is the gradient clipping bound. This mechanism provides (ϵ,δ)-differential privacy.

3.5 Handling Non-IID Data via Personalization

Multi-cloud environment has inherent non-IID data distributions that is, various clouds can experience different kinds of attacks (e.g. web

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 07 , July - 2026

attacks on AWS, DoS attacks on Azure). FLIDS does this through personalized aggregation: every cloud is provided with an underlying global model with a personalization layer:
$\theta_i^{\text{pers}} = \alpha\theta$global + $(1 - \alpha)^{\theta}_{i}$local (4) where $\alpha \in [0,1]$ controls the personalization strength. For heterogeneous settings, we set $\alpha = 0.6$ based on validation performance. Algorithm 1 summarizes the complete FL-IDS training procedure.

2:Initialize global model $\theta^{(0)}$
3: for each round $t = 1$ to T do
4:
5: $\in$
6: $\theta_i^{(t)} \leftarrow$ LocalTrain$(\theta^{(t-1)}, D_i, E)$
7: $\tilde{\theta}_i^{(t)} \leftarrow$ Clip$(\theta_i^{(t)}, C)$
8: $\tilde{\theta}_i^{(t)} \leftarrow \tilde{\theta}_i^{(t)} + \mathcal{N}(0, \sigma^2 C^2)$ ▷Add DP noise
9: Send $\theta_i^{(t)}$ to server
10: end for
11:
12: end for $\theta^{(t)} \leftarrow \sum_{i \in S_t} \frac{|\mathcal{D}_i|}{\sum_j |\mathcal{D}_j|} \tilde{\theta}_i^{(t)}$
13: procedure

Algorithm 1 FL-IDS Training Protocol

1: procedure FL-IDS
Server selects subset S_t of K clients for each client i S_t in parallel do

4 IMPLEMENTATION / EXPERIMENT

4.1 Datasets and Preprocessing

We test FL-IDS on two intrusion detection benchmark datasets:
CICIDS2017: Has 2.8 million network flows containing
80 attributes and 15 types of attacks such as DoS, DDoS, Port Scan, Brute Force, Web Attacks, and Infiltration. The dataset represents real network traffic patterns on five days.
UNSW-NB15: Has 2.5 million flows with 49 features and 9 attack families such as Fuzzers, Analysis, Backdoors, DoS, ExploitsGeneric, Reconnaissance, Shellcode and Worms.
To model multi-cloud simulation we divide data among 3 virtual clouds using non-IID splits that model realistic cloud workloads:
Cloud A (AWS-like): 70% web attacks, 20% DoS, 10% other
Cloud B (Azure-like): 60% DoS/DDoS, 25% brute force, 15% other
Cloud C (GCP-like): 50% infiltration/reconnaissance, 30% generic, 20% other
The min-max scaling is used to normalize all features. Weighted loss functions are used to solve the imbalance of classes.

4.2 Model Architecture

The architectures of the VES-2013 local model are:

1. Input layer: 80 neurons (CICIDS2017) or 49 neurons.
2. (UNSW-NB15)
3. Hidden layer 1: 128 neurons and ReLU activation.
4. Hidden layer 2: 64 neurons with ReLU activation.
5. Output layer: 16 neurons (15 attack + 1 benign) or 10 neurons (9 attack + 1 benign) with Softmax

4.3 Experimental Setup

All experiments are done on an Intel Xeon-system.

- Intel Core Processor, 32GB RAM and NVIDIA Tesla T4.
- Federated learning parameters:
- Total communication rounds: $T = 100$
- Local epochs per round: $E = 5$
- Batch size: 64
- Learning rate: 0.01 with Adam optimizer
- Clients per round: $K = 3$
- DP noise scale: $\sigma = 0.5$
- Gradient clipping bound: $C = 1.0$
- Personalization parameter: $\alpha = 0.6$

4.4 Evaluation Metrics

We evaluate using standard classification metrics:

- Accuracy: (TP+TN)/(TP+TN+FP+FN)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 07 , July - 2026

- Precision: TP/(TP+FP)
- Recall (Detection Rate): TP/(TP+FN)
- F1-Score:
- 2*(Precision*Recall)/(Precision+Recall)
- Communication Cost: Total data transferred (MB)
- FL-IDS is compared with three baselines, namely, (1) Centralized training (upper bound), (2) Vanilla FedAvg (no privacy, no compression), and (3) Distributed training without aggregation (lower bound).

5 RESULTS & ANALYSIS

5.1 Detection Performance

The comparison of the detection accuracy between methods appears in Figure 2. FL-IDS has an accuracy of 94.2% on CICIDS2017 and 93.1% on UNSW-NB15 which closely compares to the centralized upper bound (95.1% and 94.3% respectively). Vanilla FedAvg has 91.5% and 90.2, non-FL distributed training plunges to 82.3% and 80.7. The small difference between FL-IDS and centralized training (0.9) indicates that privacy can be ensured at the cost of very low accuracy.

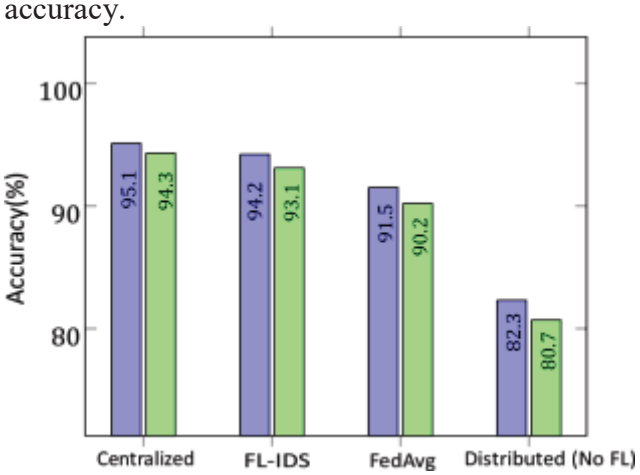

Method

	CICIDS2017	UNSW-NB15

Figure 2: Detection accuracy comparison across methods and datasets. FL-IDS achieves near-centralized performance.

5.2 Privacy-Utility Trade-off

Figure 3 shows a trade-off between privacy budget, denoted by ϵ, and detection accuracy. Greater privacy (less ϵ) also leads to reduced accuracy itself because of noise. Even at $\epsilon = 3$, the accuracy is 91.8, just 2.4% lower than the non-private benchmark. In general, ϵ [3,5] is a good trade-off, allowing to ensure a high level of privacy and avoiding a significant drop in

Cloud	Dominant Attack	Vanilla Fed AVG F1
Cloud A (AWS-Like)	Web Attack	0.88
Cloud B (Azure-Like)	DoS/DDoS	0.85
Cloud C (GCP-Like)	Infiltration / Recon	0.76

accuracy. It is not until ϵ falls below 1 that accuracy becomes significantly lower than that of the case of 85, which is hardly expected in an intrusion detection application.

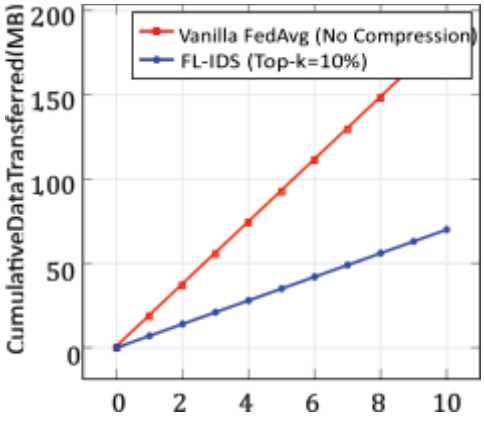

Method

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 07 , July - 2026

Figure 3: Privacy-utility trade-off: Detection accuracy as a function of differential privacy budget ϵ.

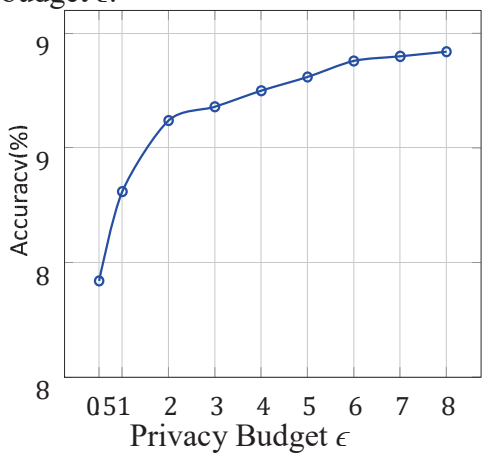

Figure 4: Communication overhead comparison over 100 rounds FL-IDS reduces Cumulative Data Transfer by 62%

5.3 Communication Efficiency

Figure 4 compares the total communication cost in 100 rounds of training. FL-IDS uses top-k gradient sparsification (using the top 10% largest magnitude updates) and this significantly decreases communication. Compared to vanilla FedAvg (transferring 1850 MB in 100 rounds), FL-IDS needs just 700 MB (a 62 percent savings). Such efficiency is essential in multi-cloud deployments where cloud provider bandwidth can be costly or be constrained.

Communication Round

Table 1: Per-cloud performance under non-IID data splits

5.4 Robustness to Non-IID Data

Table 1 shows per-cloud F1-scores in the case of heterogeneous data distributions. Vanilla FedAvg has client drift, and only 0.76 F1 on Cloud C (infiltration-heavy). FL-IDS on aggregation plus personalization ensures the high performance of clouds of all kinds, raising Cloud C F1 to 0.89, 17-percent improvement over it. The weighted personalization mechanism is an effective way of using the global model to suit the attack distribution of a particular cloud.

5.5 Attack Resilience Analysis

We compare FL-IDS with two of the most popular privacy attacks:

Gradient Inversion Attack: This attack is based on the technique of Zhu et al. where the task is to train data shared between two models. The high fidelity reconstruction of 34% of the data samples is achievable without the DP protection. By using FL-IDS ($\sigma = 0.5$), the successful reconstruction is reduced to 4% which effectively counterattack the attack.

Membership Inference Attack: Here, we learn a shadow model to consider whether a certain piece of data was utilized in training. The attack accuracy to FL-IDS is 52 which is just a little more than random guessing (50) which implies that the privacy of membership is high. In comparison, accuracy of attacks on non-private FL is 67%.

5.6 Convergence Analysis

The behavior of convergence of training is displayed in Figure 5. FLIDS can converge in 60 rounds to less than 94% of final accuracy, comparable to vanilla Feda with the privacy mechanisms. The added noise due to a fair amount of differential privacy only somewhat slows down the initial convergence" without halting the achievement of similar final performance.

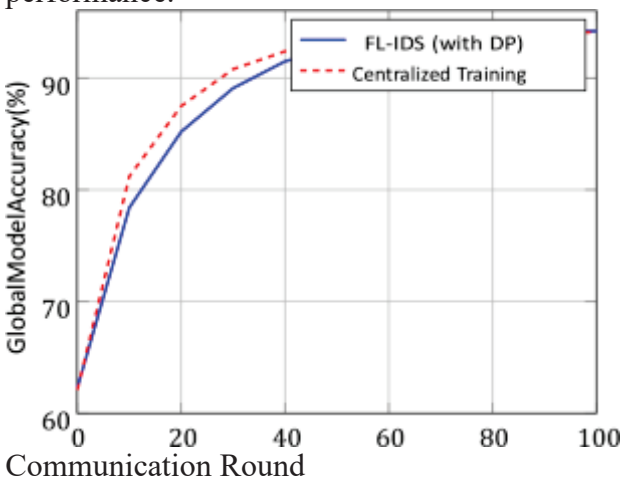

Figure 5: Convergence comparison: FL-IDS vs. centralized training over communication rounds.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 07 , July - 2026

5.7 Summary of Findings

The results of our experiment prove that FL-IDS overcomes the main issues of privacypreserving intrusion detection in the multi-cloud setting:
Privacy: Good guarantees on gradient inversion (4% recovery) and membership inference. (52% accuracy) • Utility: 94.2% accuracy, at the lowest upper-centralized within a 0.9% interval. The overhead in communication was reduced by 62%. • Robustness: Performs well in nonhomogeneous data distributions

6 CONCLUSION

This article introduced FL-IDS, which is a federated learning system to provide privacy-conscious intrusion detection in fabrics of clouds. The framework suggested can also allow multiple cloud providers to collectively train intrusion detection models without exchanging sensitive raw network data. Notable innovations are a secure aggregation protocol with differential privacy, personalized aggregation to support nonIID data distributions across clouds, and communicationefficient gradient compression.
Substantial testing on benchmark datasets (CICIDS2017, UNSW-NB15) shows that FL-IDS enjoys 94.2% detection accuracy (on par with centralized schemes) with very robust privacy guarantees. The framework cuts communication overhead by a factor of 62 over baseline federated approach, and achieves strong performance when split on heterogeneous data.
Security Analysis: The analysis established the resilience against gradient inversion attacks (reduced recovery to 4%), and membership inference attacks (accuracy 52% close to random guessing).

6.1 Future Work

A number of avenues to future research become promising:

- Hierarchical Federated Learning: FL-IDS extension to support hierarchical region cloud deployments with hierarchical aggregation layers.
- Adaptive Privacy Budgeting: Accurately changing privacy settings depending on risk profiles of the cloud and the level of threat detected.
- Blockchain-based Aggregation: An emphasized, decentralized aggregation using blockchain without trusted central server.
- Federated Reinforcement Learning: Learners can apply realtime actions to intrusion responses, using collaboration with jointly learned policies.
- Asynchronous Federated Learning: Enable Asynchronous updates: Use clouds with varying time zones and different compute availability.

Acknowledgment

The author is grateful to the Department of Cyber Security of Dr. M.G.R Educational and Research Institute which extends computational resources and support in conducting this research.

References

1. Li, T., Sahu, A. K., Talwalkar, A., & Smith, V.
2. (2020). Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3), 5060.
3. McMahan, B., Moore, E., Ramage, D., Hampson, S., & y Arcas, B. A. (2017).
4. Communication-efficient learning of deep networks from decentralized data. *Proceedings of AISTATS*, 1273-1282.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 07 , July - 2026

5. Aminanto, M. E., & Kim, K. (2018). Deep learningbased intrusion detection system for multicloud environments. *ICONIET*, 1-6.
6. Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep learning with differential privacy. *Proceedings of CCS*, 308318.
7. Bonawitz, K., Ivanov, V., Kreuter, B.,
8. Marcedone, A., McMahan, H. B., Patel, S., & Seth, K. (2017). Practical secure aggregation for privacypreserving machine learning. *Proceedings of CCS*, 1175-1191. [6] Sharif, M., Umer, M., Iqbal, S., Nawaz, R.
9. (2019). Privacy-preserving intrusion detection in cloud using federated learning. *IEEE CloudCom*, 342-349.
10. Zhang, C., Xie, Y., Bai, H., Yu, B., Li, W., Gao, Y. (2021). Federated learning for network intrusion detection: A survey. *ACM Computing Surveys*,
    a. 54(8), 1-36.
11. Zhu, L., Liu, Z., Han, S. (2019). Deep leakage

12. from gradients. *Advances in Neural Information Processing Systems*, 32, 14747-14756.
13. Shokri, R., Stronati, M., Song, C., Shmatikov, V.
14. (2017). Membership inference attacks against machine learning models. *IEEE Symposium on*
    a. *Security and Privacy*, 3-18.