



Trustworthy **AI** using Confidential Federated Learning

**FEDERATED LEARNING AND
CONFIDENTIAL COMPUTING ARE NOT
COMPETING TECHNOLOGIES.**

JINNAN GUO,
PETER PIETZUCH,
ANDREW PAVERD,
AND KAPIL VASWANI

The AI revolution is reshaping industries and transforming the way we live, work, and interact with technology. From AI chatbots and personalized recommendation systems to autonomous vehicles navigating city streets, AI-powered innovations are emerging everywhere. As businesses and organizations harness AI to streamline operations, optimize processes, and drive innovation, the

A critical aspect of trustworthy AI is privacy.

potential for economic growth and societal advancement is immense.

Amid this rapid progress, however, ensuring AI's trustworthiness is critical. Trustworthy AI systems must exhibit certain characteristics such as reliability, fairness, transparency, accountability, and robustness. Only then can AI systems be depended upon to operate ethically and effectively, without causing harm or discrimination.

A critical aspect of trustworthy AI is privacy. Training accurate machine-learning models often requires large, diverse, and representative datasets. While in some domains, models can be trained exclusively using publicly available datasets, other scenarios require access to private data. For example, training models to make medical diagnoses may require sensitive patient data. Similarly, training models for detecting fraudulent financial transactions requires detailed transaction data from financial institutions. Such data must be safeguarded from unauthorized access, manipulation, or misuse to maintain model integrity and prevent bias.

Consequently, there has been growing interest in privacy-preserving machine-learning techniques such as FL [federated learning].¹⁷ FL is a distributed machine-learning paradigm that enables training models across multiple clients holding local training data, without exchanging that data directly. In a typical FL setup, a central aggregator starts a training job by distributing an initial model to multiple clients. Each client trains the model locally on its dataset and computes updates to the model (also referred to as gradient updates). The clients then send their updates to the central aggregator, which


aggregates these updates using a suitable aggregation function and updates its model. It then starts another epoch by sending the updated model to the clients, which perform local training. This process repeats until the mutually agreed-upon termination criteria are met (e.g., the model converges to an acceptable loss value).

FL can be combined with differential privacy⁷ to provide strong privacy guarantees.²⁴ In this setting, each client adds suitable noise to the model updates locally, based on a privacy budget before sending the updates to the aggregator, which bounds the probability for the model to memorize individual points in the training dataset.

While FL prevents the flow of raw training data across trust domains, it introduces a new set of trust assumptions and security challenges. Clients participating in FL must trust a central aggregator to deliver safe code, include only trustworthy clients, follow the aggregation protocol, and use the model only for mutually agreed-on purposes. In addition, the aggregator must trust the clients to provide high-quality data, not tamper with the training protocol, and protect the model's intellectual property. These trust assumptions are often difficult to satisfy in the real world, especially in adversarial settings where clients may be compromised or collude to undermine the system's security and privacy guarantees. It is therefore unsurprising that many FL deployments have been found to be vulnerable to attacks, including model poisoning, data poisoning, and inference attacks^{8,10,22} (see figure 1). Attacks may be carried out by clients, aggregators, or outsiders and occur during model training or inference.

Many of these attacks can be attributed to the ability

FIGURE 1: **ATTACKS ON FEDERATED LEARNING SYSTEMS**



of malicious participants to violate the confidentiality or integrity of data and computation in their control [e.g., by poisoning datasets or gradient updates to influence the model's behavior]. These attacks are not limited to just the aggregators or clients at training time—attacks such as model extraction or reconstruction can be carried out by

entities with API access to the trained model at inference time. Therefore, it is critical to protect all sensitive information throughout the lifecycle of FL jobs.

Another challenge in FL is transparency and accountability. Since, by definition, FL does not involve sharing training data directly, it is difficult to audit the training process and verify that the model has not been biased or tampered with. This makes it challenging for model builders to comply with any AI regulations that require transparency or auditability of the training process as a precondition for deployment.

An alternative approach for privacy-preserving machine learning is *confidential computing*.²¹ Confidential computing enables the secure execution of code and data in untrusted computing environments—for example, public clouds—by leveraging hardware-based TEEs (trusted execution environments) such as Intel SGX (Software Guard Extensions),^{2,5} AMD SEV-SNP (Secure Encrypted Virtualization-Secure Nested Paging),¹ Arm CCA (Confidential Compute Architecture),¹⁵ and more recently, Nvidia Hopper Confidential Computing.⁶

Confidential computing protects the confidentiality and integrity of machine-learning models and data throughout their lifecycles, even from privileged attackers. In most existing machine-learning systems with confidential computing, however, the training process remains centralized, requiring data owners to send (potentially encrypted) datasets to a single client where the model is trained in a TEE. Unlike FL, this setup places significant trust in the TEE infrastructure to protect datasets in a remote, potentially hostile environment.

FL and confidential computing should not be considered competing technologies. Rather, it is possible, with careful design, to combine FL and confidential computing to achieve the best of both worlds: the assurance of sensitive data remaining within its trust domain while ensuring transparency and accountability. This new paradigm, referred to here as *CFL* (*confidential federated learning*), can prevent large classes of attacks on FL, broaden the adoption of FL in privacy-sensitive domains, and enable compliance with upcoming AI regulations.

CONFIDENTIAL COMPUTING

Confidential computing uses TEEs to isolate sensitive code and data from privileged attackers. There are several kinds of TEEs in modern CPUs. For example, Intel CPUs support the creation of process-based TEEs through Software Guard Extensions.² Process-based TEEs can measure and isolate a user-space process from the rest of the system, including other processes and the OS (operating system). Within process-based TEEs, code does not have direct access to any OS kernel functionality such as I/O devices. Therefore, writing applications to use process-based TEEs requires significant developer effort.

Led by AMD SEV-SNP,¹ recent CPUs support VM (virtual machine)-based TEEs. These TEEs are capable of hosting and isolating both user-mode processes and a full OS from external access. This makes it simpler to migrate existing applications to VM-based TEEs, albeit at the cost of a larger TCB (trusted computing base).

While confidential computing has been supported in CPUs for well over a decade, the primitives required for

deploying AI workloads such as FL transparently with low performance overheads have evolved only recently.

Confidential containers

While VM-based TEEs can host legacy virtual machines, this mode of deployment has limitations beyond a large TCB. Unless configured correctly, it does not fully isolate the workload (user-mode applications) from external access [e.g., secure shell access by the OS admin]. It also provides limited attestation of the workload because it requires the VM to be started with a bootloader, which in turn boots an OS kernel. Therefore, only the bootloader is measured by the hardware. Even if attestation were to be extended to include the OS kernel [e.g., using a virtual Trusted Platform Module], it is challenging to attest the entire OS and user-mode applications.

Confidential containers^{3,11} are a new mode of deploying applications in VM-based TEEs that address these limitations. In confidential containers, a VM-based TEE is used to host a utility OS along with a container runtime, which in turn can host containerized workloads. Confidential containers support full workload integrity and attestation via container execution policies. These policies define the set of container images (represented by the hash digest of each image layer) that can be hosted in the TEE, along with other security-critical attributes such as commands, privileges, and environment variables. The policy itself is measured (as an initialization time claim) by the hardware root of trust, included in the hardware-signed attestation report, and enforced by the container runtime. In other words, the combination of the OS,

container runtime, and container policy fully represents the workload hosted in the TEE and can be used by relying parties to establish trust in the environment.

Confidential GPUs

Initially, support for confidential computing was limited to CPUs, with all other devices considered as untrusted. This was, of course, limiting for AI applications that use GPUs for achieving high performance. Over the past few years, there have been several attempts at building confidential computing support in accelerators. Nvidia's Hopper generation of GPUs⁶ supports the creation of TEEs and can be coupled with CPU-based TEEs (AMD SEV-SNP, Intel TDX⁴) to create a unified TEE across CPU and GPU, enabling transparent offload with low performance overheads.

Hopper GPUs support the new confidential computing mode in which the GPU carves out a region of memory called the protected region and enables a hardware firewall that isolates this region and other sensitive parts of state from the host CPU. In this mode, a CPU-based TEE such as an SNP VM can attest and establish a secure channel with the GPU and provision encryption keys to copy engines in the GPU. All subsequent data transfers—including code; models; and application data between the CPU TEE and the GPU, and between GPUs—are encrypted using these keys.

CONFIDENTIAL FEDERATED LEARNING


A typical FL deployment involves several components that work together to enable collaborative model training across multiple clients. This includes client environments

that hold local data, a central aggregator, an orchestrator for managing FL tasks, and the communication infrastructure for provisioning tasks and exchanging model updates.

Most FL frameworks such as NVFlare²⁰ support several security measures to protect data and models, including the use of network security to isolate and sandbox remote code; TLS (Transport Layer Security) for secure communication; and strong authentication and access-control mechanisms. Despite these measures, there are plenty of avenues for a malicious participant to exfiltrate secrets or tamper with the training process. For example, a malicious participant can poison datasets by adding additional samples or changing labels of training data to introduce backdoors or bias into the model. Data may be poisoned either before a training job or adaptively during the job, based on intermediate models. A participant may also observe or tamper with gradient updates or arbitrarily tamper with the workflow—for example, by skipping training entirely or not aggregating certain inputs.

CFL (confidential federated learning) is an emerging paradigm^{18,19} that aims to harden FL deployments against such attacks. Figure 2 shows the architecture of a typical CFL deployment for a single training job. In CFL, all computation (aggregation and training) is hosted in a special class of hardware-isolated TEEs. The TEEs isolate data and computation from all external access, including administrators and privileged attackers. In particular, with TEEs, model weights are no longer exposed to client administrators; they are visible only to attested client code. Similarly, intermediate gradient updates are no

FIGURE 2: ARCHITECTURE OF A TYPICAL CFL DEPLOYMENT



longer exposed to the aggregator; they are exposed only to attested aggregator code. The aggregator learns the trained model only, and even that access can be limited by

hosting the trained model in a TEE.

TEEs used in CFL also provide integrity—a malicious aggregator or client cannot tamper with data, computation, or configuration of the deployment. For example, if a training job requires each client to pre-process the dataset (e.g., run sampling and reweighing with specific parameters to mitigate bias),¹³ clients cannot change the control flow of the training job or parameter values without being detected via attestation. The integrity properties of TEEs hold even in the presence of side-channel attacks.^{12,14,16,23}

Finally, CFL uses TEEs that can provide hardware-based attestation for the full workload and configuration of the FL job, including pre-processing, training, and optional inferencing. TEEs that meet these requirements include Azure Confidential Containers and Confidential Spaces on the Google Cloud Platform.

Commitments

In addition to hosting computations in TEEs, CFL can support transparency and accountability through *commitments*. Participants in CFL can be required to commit to their inputs before running a training job. Data providers commit to their datasets, and model providers commit to the job configuration and the initial model state (if provided externally). For example, the job configuration in NVFlare is a list of tasks that will be executed by the aggregator and clients, along with the configuration for each task.

Commitments can take various forms. For smaller inputs such as a job configuration, the input (or its hash digest) can be attested directly. For larger inputs such as

datasets, one option is to compute a Merkle hash tree over the dataset (e.g., using dm-verity) and use the root hash of the tree (combined with a random nonce) as a commitment.

In CFL, commitments are reflected in TEE attestation, verified by other participants, and enforced during TEE execution. For example, in an implementation with Azure Confidential Containers, the dm-verity root hash of the training dataset is included as an environment variable in the container security policy. Within the TEE, this root hash is used to verify that the Merkle tree is correct. The Merkle tree is then used to verify the integrity of the dataset by comparing the hash digest of each block that is read against the hash value in the Merkle tree. Reflecting commitments in attestation ensures that any given client can connect to the aggregator only if it provides the committed dataset as input. This invariant holds even across clients and aggregator restarts, since clients and aggregators mutually attest each other on every connection (see next section).

Commitments, as used in CFL, have a few noteworthy characteristics. First, they do not impact privacy since only a hash is revealed, not the dataset itself. Commitments do not prevent clients from providing bad data; they ensure only that a malicious client cannot change the dataset adaptively during training. This significantly limits the power of an attacker because the attack must now be designed to work irrespective of other datasets used in training. Finally, commitments, in conjunction with attestation reports, provide tamper-proof provenance for the entire FL job.

Armed with attestation reports, external auditors get

In CFL, each participant specifies its criteria for trusting other participants by creating an attestation policy.

full visibility into the flow of datasets that contributed to the model and can hold participants responsible for a model's behavior.

Mutual attestation

Including the full workload, configuration, and commitments in attestation reports enables other participants in an FL computation to remotely verify and establish trust in a participant's compute instances. For example, an aggregator can verify all clients, and each client can independently verify the central aggregator.

In CFL, each participant specifies its criteria for trusting other participants by creating an attestation policy. This can take the form of a key-value map, where each key is the name of a claim, and the value is the set of values that the claim is allowed to take.

The following is a sample attestation policy with multiple claims and permitted values for each claim. Each CFL node is provisioned with a policy that it uses to verify attestation reports from other nodes.

```
{  
  "host_data": [ "..."] ,  
  "report_data": [ "...", "...", "...", ]  
  "svn": [ "..."]  
}
```

To ensure that a participant communicates only with other participants that it trusts, CFL deployments can perform attestation verification as part of the TLS handshake

1. On start-up, each client and aggregator generates an

ephemeral TLS signing key and obtains an attestation report with the key as a runtime claim.

2. Each node generates a self-signed certificate and includes the attestation report and other collateral required to verify the report (such as device certificates) as a custom extension in the certificate. Each instance configures its TLS stack to use this TLS signing certificate.
3. Each node also configures the TLS stack (e.g., using callbacks supported by TLS) to verify certificates obtained from other participants during the handshake, based on its attestation policy.

This protocol ensures that each instance establishes a secure encrypted communication channel with other instances only after verifying the attestation report against the attestation policy. All subsequent communication between the aggregator and client, such as communicating model weights and gradient updates, uses this channel.

One challenge in deploying attestation policies is that it can lead to cyclic dependencies, because the aggregator's attestation policy depends on each client's attestation, and vice versa. One way to break the cycle is to include the aggregator's attestation policy in its attestation but exclude the client's policy from its attestation. This design choice preserves the ability for clients to assess the aggregator's attestation policy before entrusting the aggregator with their data.

IMPLEMENTING CFL

We have experimented with a CFL implementation based on Nvidia NVFlare, a commonly used FL framework. Our

We evaluated the CFL's end-to-end performance by measuring the training throughput.

prototype can run on confidential containers on ACIs (Azure Container Instances) as well as CVMs (confidential VMs).⁹ NVFlare containers could be hosted in ACI and CVMs without modifications to the core NVFlare framework. To simplify deployment, we built a provisioning tool to generate scripts for generating dataset commitments, attestation policies for clients and servers, and scripts for deploying NVFlare containers to ACIs and CVMs. Dataset commitments are implemented using dm-verity. Transparent, mutually attested TLS and attestation policy enforcement are supported using a network proxy.

We evaluated the CFL's end-to-end performance by measuring the training throughput. To perform the end-to-end evaluation, we deploy the CFL aggregator in Azure DC4asv5 CVM (with four vCPUs, 16 GiB of memory) and the CFL client in Azure DC32asv5 CVMs (32 vCPUs, 128 GiB of memory). Our experiments suggest that adding TEE and dm-verity protection for the FL system results in a five percent reduction in overall throughput for simple FL workloads (based on CIFAR-10).

We also investigated the overhead of introducing commitments using dm-verity with a sequential read benchmark, which is representative of training workloads where the entire dataset is read sequentially. Our experiments suggest dm-verity protection can introduce an overhead up to 40 percent in sequential read throughput as a result of read amplification caused by Merkle tree checks. The impact of reduced storage throughput on end-to-end training throughput is small because most training workloads tend to be compute-bound. These are initial results and need to be

substantiated with more rigorous evaluation using larger workloads.

CONCLUSIONS

The principles of security, privacy, accountability, transparency, and fairness are the cornerstones of modern AI regulations. Classic FL was designed with a strong emphasis on security and privacy, at the cost of transparency and accountability. CFL addresses this gap with a careful combination of FL with TEEs and commitments. In addition, CFL brings other desirable security properties, such as code-based access control, model confidentiality, and protection of models during inference. Recent advances in confidential computing such as confidential containers and confidential GPUs mean that existing FL frameworks can be extended seamlessly to support CFL with low overheads. For these reasons, CFL is likely to become the default mode for deploying FL workloads.

References

1. AMD. 2020. AMD SEV-SNP: strengthening VM isolation with integrity protection and more. White paper; <https://www.amd.com/content/dam/amd/en/documents/epyc-business-docs/solution-briefs/amd-secure-encrypted-virtualization-solution-brief.pdf>.
2. Anati, I., Gueron, S., Johnson, S., Scarlata, V. 2013. Innovative technology for CPU based attestation and sealing. *Proceedings of the 2nd International Workshop on Hardware and Architectural Support for Security and*

- Privacy* 13; <https://www.intel.com/content/dam/develop/external/us/en/documents/hasp-2013-innovative-technology-for-attestation-and-sealing-413939.pdf>.
3. Brasser, F., Jauernig, P., Pustelnik, F., Sadeghi, A.-R., Stapf, E. 2022. Trusted container extensions for container-based confidential computing. arXiv preprint arXiv:2205.05747; <https://arxiv.org/abs/2205.05747>.
 4. Cheng, P.-C., Ozga, W., Valdez, E., Ahmed, S., Gu, Z., Jamjoom, H., Franke, H., Bottomley, J. 2023. Intel TDX demystified: a top-down approach. arXiv preprint arXiv:2303.15540; <https://arxiv.org/abs/2303.15540>.
 5. Costan, V., Devadas, S. 2016. Intel SGX explained; <https://eprint.iacr.org/2016/086>.
 6. Dhanuskodi, G., Guha, S., Krishnan, V., Manjunatha, A., Nertney, R., O'Connor, M., Rogers, P. 2023. Creating the first confidential GPUs. *acmqueue* 21(4); <https://queue.acm.org/detail.cfm?id=3623391>.
 7. Dwork, C., Roth, A., et. al. 2014. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science* 9(3-4), 211–407; <https://dl.acm.org/doi/10.1561/04000000042>.
 8. Fang, M., Cao, X., Jia, J., Gong, N. 2020. Local model poisoning attacks to Byzantine-robust federated learning. *Proceedings of the 29th Usenix Security Symposium*, article 92, 1623–1640; <https://dl.acm.org/doi/abs/10.5555/3489212.3489304>.
 9. Hande, K. 2023. Announcing Azure confidential VMs with NVIDIA H100 Tensor Core GPUs in Preview. Azure Confidential Computing Blog; <https://techcommunity.microsoft.com/t5/azure-confidential-computing/announcing-azure-confidential-vms-with-nvidia-h100->

- tensor-core/ba-p/3975389#:~:text="The%20Azure%20confidential%20VMs%20with,remain%20protected%20end%20to%20end."
10. Jere, M. S., Farnan, T., Koushanfar, F. 2020. A taxonomy of attacks on federated learning. *IEEE Security & Privacy* 19(2), 20–28; <https://ieeexplore.ieee.org/document/9308910>.
 11. Johnson, M. A., Volos, S., Gordon, K., Allen, S. T., Wintersteiger, C. M., Clebsch, S., Starks, J., Costa, M. 2023. COCOAEXPO: confidential containers via attested execution policies. *arXiv preprint arXiv:2302.03976*; <https://arxiv.org/abs/2302.03976>.
 12. Kocher, P., Horn, J., Fogh, A., Genkin, D., Gruss, D., Haas, W., Hamburg, M., et al. 2019. Spectre attacks: exploiting speculative execution. *40th IEEE Symposium on Security and Privacy*, 1–19; <https://ieeexplore.ieee.org/document/8835233>.
 13. Krasanakis, E., Spyromitros-Xioufis, E., Papadopoulos, S., Kompatsiaris, Y. 2018. Adaptive sensitive reweighting to mitigate bias in fairness-aware classification. *Proceedings of the 2018 World Wide Web Conference*, 853–862; <https://dl.acm.org/doi/10.1145/3178876.3186133>.
 14. Li, M., Zhang, Y., Wang, H., Li, K., Cheng, Y. 2021. CIPHERLEAKS: breaking constant-time cryptography on AMD SEV via the ciphertext side channel. *30th Usenix Security Symposium*, 717–732; <https://www.usenix.org/conference/usenixsecurity21/presentation/li-mengyuan>.
 15. Li, X., Li, X., Dall, C., Gu, R., Nieh, J., Sait, Y., Stockwell, G. 2022. Design and verification of the Arm confidential compute architecture. *16th Usenix Symposium on*

- Operating Systems Design and Implementation; <https://www.usenix.org/conference/losdi22/presentation/li>.
16. Lipp, M., Schwarz, M., Gruss, D., Prescher, T., Haas, W., Fogh, A., Horn, J., et al. 2018. Meltdown: reading kernel memory from user space. *Proceedings of the 27th Usenix Security Symposium*; <https://www.usenix.org/conference/usenixsecurity18/presentation/lipp>.
 17. McMahan, B., Moore, E., Ramage, D., Hampson, S., Aguera y Arcas, B. 2017. Communication-efficient learning of deep networks from decentralized data. *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, 1273–1282; <https://proceedings.mlr.press/v54/mcmahan17a/mcmahan17a.pdf>.
 18. Mo, F., Haddadi, H., Katevas, K., Marin, E., Perino, D., Kourtellis, N. 2022. PPFL: enhancing privacy in federated learning with confidential computing. *GetMobile: Mobile Computing and Communications* 25(4), 35–38; <https://dl.acm.org/doi/10.1145/3529706.3529715>.
 19. Quoc, D. L., Fetzner, C. 2021. SecFL: confidential federated learning using TEEs. arXiv 2110.00981; <https://arxiv.org/abs/2110.00981>.
 20. Roth, H. R., Cheng, Y., Wen, Y., Yang, I., Xu, Z., Hsieh, Y.-T., Kersten, K., et al. 2022. NVIDIA Flare: federated learning from simulation to real-world. arXiv preprint arXiv:2210.13291; <https://arxiv.org/abs/2210.13291>.
 21. Russinovich, M., Costa, M., Fournet, C., Chisnall, D., Delignat-Lavaud, A., Clebsch, S., Vaswani, K., Bhatia, V. 2021. Toward confidential cloud computing. *Communications of the ACM* 64(6), 54–61;

- <https://dl.acm.org/doi/10.1145/3453930>.
22. Tolpegin, V., Truex, S., Gursay, M. E., Liu, L. 2020. Data poisoning attacks against federated learning systems. 25th European Symposium on Research in Computer Security, Proceedings, Part I 25, 480–501; https://dl.acm.org/doi/10.1007/978-3-030-58951-6_24.
 23. Van Bulck, J., Minkin, M., Weisse, O., Genkin, D., Kasikci, B., Piessens, F., Silberstein, M., Wenisch, T. F., Yarom, Y., Strackx, R. 2018. Foreshadow: extracting the keys to the Intel SGX kingdom with transient out-of-order execution. Proceedings of the 27th Usenix Security Symposium; <https://www.usenix.org/conference/usenixsecurity18/presentation/bulck>.
 24. Wei, K., Li, J., Ding, M., Ma, C., Yang, H. H., Farokhi, F., Jin, S., Quek, T. Q. S., Poor, H. V. 2020. Federated learning with differential privacy: algorithms and performance analysis. IEEE Transactions on Information Forensics and Security 15, 3454–3469; <https://ieeexplore.ieee.org/document/9069945>.

Jinnan Guo is a PhD candidate at Imperial College London, advised by Peter Pietzuch. His research interest lies in the intersection of systems, security, and machine learning.

Peter Pietzuch is a professor of distributed systems at Imperial College London, where he leads the Large-scale Data & Systems (LSDS) group. His research work focuses on the design and engineering of scalable, reliable, and secure software systems, with a particular interest in supporting AI and machine learning workloads. He is also a co-director of Imperial's I-X initiative in AI, data, and digital. Before joining

Imperial, he was a post-doctoral Fellow at Harvard University, and he holds PhD and MA degrees from the University of Cambridge.

Andrew Paverd *is a principal research manager in the Microsoft Security Response Center (MSRC). His research work focuses primarily on security, privacy, and safety in AI systems, with additional interests in confidential computing and web security. Andrew holds a BSc from the University of the Witwatersrand, Johannesburg, an MSc from the University of Cape Town, and a DPhil from the University of Oxford.*

Kapil Vaswani *is a principal researcher at Azure Research in Cambridge. His research interests lie in secure and robust systems, with a particular focus on designing confidential computing hardware and services with applications to AI and web security. Kapil holds PhD and MSc degrees from the Indian Institute of Science, Bangalore.*

Copyright © 2024 held by owner/author. Publication rights licensed to ACM.