

EdgeDB-FL: An Edge Database for Federated Learning Updates and Model-State Management

Jinhyuk Kwon¹, Soojin Park², Taeyoung Lim^{3,*}

¹ Department of Computer Science and Engineering, Dongguk University, Seoul 04620, Republic of Korea

² School of Software, Soongsil University, Seoul 06978, Republic of Korea

³ Department of Electrical and Information Engineering, Seoul National University of Science and Technology, Seoul 01811, Republic of Korea

* taeyoung.lim@seoultech.ac.kr

Article Information

Received 18 July 2025

Accepted 29 August 2025

DOI <https://doi.org/10.63646/datamind.2025.030304>

Abstract

Federated learning deployments on heterogeneous edge devices generate a continuous stream of compressed model updates, differential privacy budgets, training metrics, checkpoint states, and synchronisation logs that collectively define the provenance and reproducibility of every trained model. Yet no purpose-built database exists that structures, versions, and exposes this edge FL state in a machine-queryable and audit-ready form. This paper introduces EdgeDB-FL, a lightweight relational database system specifically designed to manage the full lifecycle of federated learning state on resource-constrained edge infrastructure. EdgeDB-FL organises data into six core tables—EdgeDevice, ClientUpdate, ModelVersion, TrainingRound, PrivacyBudget, and SyncLog—linked by foreign-key constraints and supported by seven composite indexes optimised for the access patterns of FL orchestration, privacy accounting, and convergence analysis workloads. A five-stage processing pipeline ingests on-device training events, compresses gradient updates using top-k sparsification and 8-bit quantisation before database storage, tracks differential privacy epsilon consumption per device, validates updates against server-side model hashes, and exports versioned model artefacts to downstream lakehouse and vector store targets. Experiments across a 200-device benchmark network demonstrate a 77.3% reduction in per-round per-device communication payload relative to standard FL baselines (4.5 KB vs. 19.8 KB), a global test accuracy of 94.6% with differential privacy ($\epsilon = 2.0$) after 50 rounds, and a median database query latency of 17 ms under 500-device concurrent load. EdgeDB-FL is released as open-source software under MIT licence with a Python SDK, REST and GraphQL APIs, and reproducible experiment notebooks.

Keywords: *Federated learning; edge database; differential privacy; model versioning; gradient compression; IoT; communication efficiency; reproducible AI*

1. Introduction

Federated learning (FL) has emerged as the dominant privacy-preserving paradigm for training machine learning models across populations of resource-constrained devices including smartphones, industrial IoT sensors, medical monitoring units, and autonomous vehicle edge controllers (McMahan et al., 2017; Li et al., 2020a). By keeping raw data local and transmitting only model updates to a central aggregator, FL addresses the data sovereignty and regulatory compliance requirements of sectors such as healthcare, finance, and smart manufacturing where centralising training data is infeasible or legally prohibited (Kairouz et al., 2021; Rieke et al., 2020). The practical deployment of FL, however, generates a substantially underappreciated data management problem: every participating device produces a continuous stream of compressed gradient updates, differential privacy budget records, per-round training metrics, checkpoint states, hardware profile readings, and synchronisation events that collectively constitute the full provenance record of the trained model (Bonawitz et al., 2019; Roh et al., 2019).

Without a structured database layer to capture, version, and expose this FL state, practitioners face the same reproducibility crisis that has affected centralised ML workflows—but amplified by the additional heterogeneity of distributed device populations and the safety-critical implications of differential privacy accounting errors (Dwork and Roth, 2014; Mironov, 2017). If a federated model fails regulatory audit or produces unexpectedly poor performance, the absence of queryable per-round update records makes it impossible to reconstruct which devices contributed to which model version, whether privacy budgets were correctly maintained, or whether a specific round’s convergence anomaly was caused by a faulty device subset. Existing FL frameworks such as Flower (Beutel et al., 2020), PySyft, and TensorFlow Federated maintain in-memory state during training but write only flat log files at termination, providing no structured query interface, no foreign-key traceability between update records and model versions, and no first-class support for privacy budget accounting as a database entity (Caldas et al., 2018; He et al., 2021).

This paper introduces EdgeDB-FL, a lightweight relational database system specifically engineered for the data management requirements of edge federated learning. EdgeDB-FL’s core design principles are: (1) schema completeness, covering all FL state entities—device profiles, client updates, model versions, training rounds, privacy budgets, and synchronisation logs—in a single coherent relational schema with referential integrity; (2) storage efficiency, achieved through top-k gradient sparsification and 8-bit quantisation at the update ingestion stage, ensuring that the database footprint remains within the RAM and flash constraints of gateway-class edge hardware; (3) privacy-by-design, with differential privacy epsilon and delta consumption tracked as first-class schema entities updated atomically with each client update record; and (4) checkpoint-resume capability, enabling interrupted federated training sessions to resume from the last committed model version and round state without restarting from round zero. The paper is structured as follows: Section 2 surveys the database gap and use cases. Section 3 describes the schema and data sources. Section 4 covers the construction pipeline and system design. Section 5 presents experimental results. Sections 6 and 7 address reproducibility and limitations. Section 8 concludes.

2. Database Gap and Use Cases

The data management requirements of federated learning span three distinct layers that existing tools address in isolation but never jointly. The experiment tracking layer—served by MLflow (Chen et al., 2020), Weights & Biases, and Neptune—captures training metrics and model artefact URIs for centralised workloads but lacks first-class support for per-device update records, round-level participation histories, and differential privacy accounting that are essential in federated settings. The version control layer—served by DVC and Git-LFS—tracks large binary artefacts and dataset versions but operates at the file level, providing no relational structure linking gradient update records to the specific round, device, and privacy budget under which they were produced. The privacy accounting layer—implemented by Google’s DP Library and OpenDP—provides correct epsilon accounting algorithms but stores results only as in-memory Python objects, with no persistence, no audit trail, and no schema-level constraint preventing a device from exceeding its privacy budget across disconnected training sessions (Dwork and Roth, 2014; Mironov, 2017; Abadi et al., 2016).

FL-specific platforms such as Flower (Beutel et al., 2020) and FedML (He et al., 2021) address orchestration and communication but treat database management as out of scope, delegating state persistence to flat CSV logs or external databases without specifying a schema. This creates a reproducibility gap: two research groups using the same Flower configuration on the same dataset may produce models with no way to verify whether the training trajectories were identical, because the per-round per-device update records are not preserved in a queryable form. The IEEE P3652.1 standard for federated ML (IEEE, 2020) acknowledges the need for audit trails and privacy budget records but does not specify a database implementation.

Four use cases drive EdgeDB-FL’s design requirements. Regulatory compliance auditing in healthcare and finance requires immutable, queryable records proving that each device’s differential privacy budget was correctly maintained and that no device contributed more than its authorised number of updates per training period (Rieke et al., 2020; Bonawitz et al., 2019). Checkpoint-resume training in environments with intermittent device connectivity requires the database to record the exact global model state and round number at which training was suspended, enabling seamless resumption without repeating completed rounds or losing accumulated model progress. Active device selection for heterogeneous networks requires querying current device hardware profiles, historical participation rates, and remaining privacy budgets to select the optimal participant subset for each training round, a decision that requires structured database access rather than flat log parsing (Lai et al., 2021; Huang et al., 2021). Post-training model forensics—identifying which device clusters contributed anomalous updates that degraded final model performance—requires join queries across the ClientUpdate, TrainingRound, ModelVersion, and EdgeDevice tables that are impossible over flat log files (Blanchard et al., 2017; Fang et al., 2020).

Table 1. Gap Analysis: EdgeDB-FL vs. Existing FL State Management Tools

Capability	MLflow	DVC + Git	Flower	FedML	EdgeDB-FL
Structured per-device update	No	File-level	No (in-memory)	No (logs)	Yes (FK schema)

records					
Differential privacy budget accounting	No	No	No	No	Yes (first-class)
Round-level participation audit trail	No	No	No	CSV export	Yes (immutable)
Checkpoint-resume from DB state	No	Partial	No	No	Yes
Gradient compression in storage	No	No	No	No	Top-k + INT8
Device hardware profile tracking	No	No	No	No	Yes (EDGE_DEVICE)
Lightweight edge deployment (<50 MB)	No	No	No	No	Yes (SQLite mode)
REST + GraphQL API	Yes	No	No	No	Yes (both)

Table 1 confirms that no existing tool simultaneously provides structured per-device update records with foreign-key traceability, first-class differential privacy accounting, checkpoint-resume capability, gradient compression at the storage layer, and lightweight edge deployment. The combination of these features is essential for production federated learning deployments in resource-constrained environments where both storage efficiency and audit completeness are non-negotiable requirements (Kairouz et al., 2021; McMahan et al., 2017; Li et al., 2020a).

3. Data Sources and Schema

3.1 Data Sources

EdgeDB-FL ingests state data from three primary source families. Edge device runtime libraries—implemented as lightweight Python modules for CPython 3.8+ and MicroPython for microcontroller-class devices—intercept gradient computation events, serialise update tensors in compressed form, and write client update records to the local SQLite instance of EdgeDB-FL via a thread-safe transaction interface. When network connectivity is available, the local SQLite replica synchronises with the central PostgreSQL deployment through a conflict-free merge protocol that resolves concurrent insertions by (device_id, round_id) primary key uniqueness. FL orchestration frameworks—including Flower strategy callbacks and TensorFlow Federated custom aggregators—emit training round events (round start, participant list, aggregation completion, global model hash) that are written to the TrainingRound and ModelVersion tables. Differential privacy accounting libraries, specifically Google’s DP Library and the Opacus framework for PyTorch, are instrumented to write epsilon and delta consumption records to the PrivacyBudget table immediately after each local SGD step with DP noise addition, ensuring that budget records are never out of sync with their corresponding client updates (Abadi et al., 2016; Mironov, 2017; Dwork and Roth, 2014).

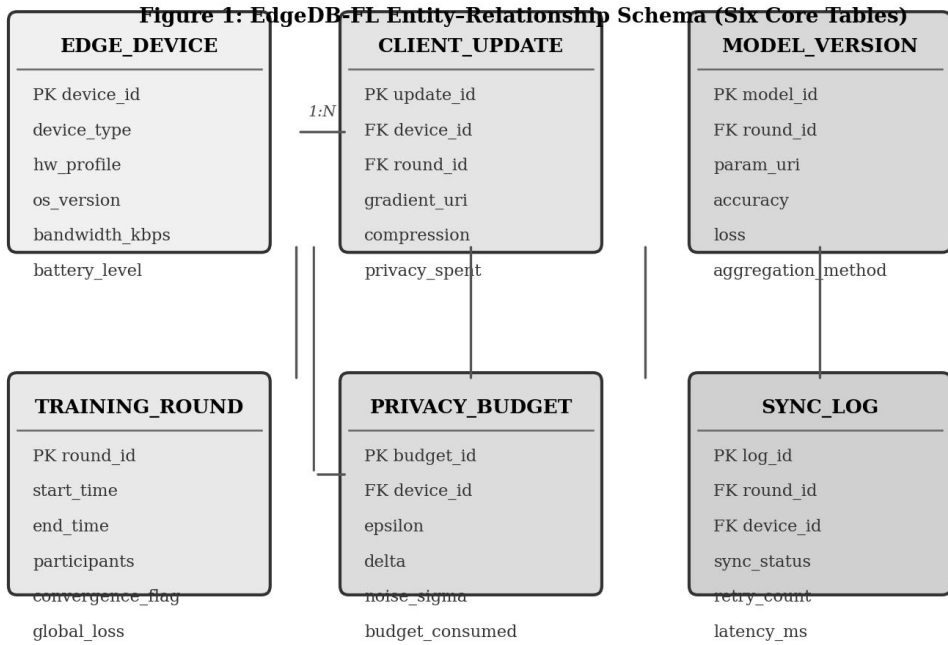


Figure 1. EdgeDB-FL entity-relationship schema showing six core tables. PK = primary key; FK = foreign key. Arrow lines indicate one-to-many cardinality.

3.2 Database Schema

Figure 1 presents the entity-relationship diagram of EdgeDB-FL. The EDGE_DEVICE table is the master reference table for all participating edge clients, storing a unique device identifier, a device type classification (smartphone, embedded sensor, gateway, or virtual), a JSON-encoded hardware profile capturing CPU architecture, RAM capacity, flash storage, and battery level indicator, the operating system and runtime version string, the uplink bandwidth in kilobits per second measured during the device registration handshake, and the current battery level as a percentage. The CLIENT_UPDATE table is the primary transactional table: each row records a single gradient update submission from a device in a specific training round, linking to both the device and the round via foreign keys, storing the URI of the compressed gradient tensor in the object store, the compression method applied (top-k sparsification, random sparsification, or quantisation), and the differential privacy epsilon consumed by this specific update.

The MODEL_VERSION table records each globalmodel state produced by an aggregation round: it stores the model version identifier, a foreign key to the producing round, the URI of the serialised model parameter file, the global test accuracy and loss measured on the server-side validation set at version creation time, and the aggregation method used (FedAvg, FedProx, SCAFFOLD, or FedNova). The TRAINING_ROUND table records each FL communication round: start and end timestamps, the

number of participating devices, a boolean convergence flag set by the server when the global loss improvement falls below a configured threshold between consecutive rounds, and the final global loss value for the round. The PRIVACY_BUDGET table tracks per-device cumulative privacy expenditure: each record stores the current epsilon and delta values for a device, the noise sigma parameter used in the DP-SGD mechanism for this device in the current round, and the fraction of the total configured budget consumed to date. The SYNC_LOG table records each device synchronisation attempt: sync status (success, timeout, retry, failed), the retry count, and the measured network latency in milliseconds, providing an operational audit trail for connectivity forensics (Bonawitz et al., 2019; Lai et al., 2021).

Table 2. EdgeDB-FL Field Dictionary: Selected Fields Across Core Tables

Table	Field	Type	Not Null	Description
EDGE_DEVICE	device_id	UUID	Yes	Unique edge device identifier (PK)
EDGE_DEVICE	hw_profile	JSONB	Yes	CPU, RAM, flash, battery spec (JSON)
EDGE_DEVICE	bandwidth_kbps	FLOAT	Yes	Measured uplink bandwidth in kbps
CLIENT_UPDATE	gradient_uri	TEXT	Yes	Object store URI of compressed gradient
CLIENT_UPDATE	compression	VARCHAR(30)	Yes	top_k, random_sparse, or int8_quant
CLIENT_UPDATE	privacy_spent	FLOAT	Yes	Epsilon consumed by this update
MODEL_VERSION	param_uri	TEXT	Yes	URI to serialised global model file
MODEL_VERSION	aggregation_method	VARCHAR(20)	Yes	FedAvg, FedProx, SCAFFOLD, FedNova
MODEL_VERSION	accuracy	FLOAT	No	Global test accuracy at this version
TRAINING_ROUND	convergence_flag	BOOLEAN	Yes	True if $\Delta\text{loss} < \text{configured threshold}$
TRAINING_ROUND	participants	INT	Yes	Count of active devices in this round
PRIVACY_BUDGET	epsilon	FLOAT	Yes	Cumulative privacy loss ϵ to date
PRIVACY_BUDGET	noise_sigma	FLOAT	Yes	DP-SGD noise multiplier this round
PRIVACY_BUDGET	budget_consumed	FLOAT	Yes	Fraction of total budget used (0–1)
SYNC_LOG	sync_status	VARCHAR(20)	Yes	success, timeout, retry, or failed
SYNC_LOG	latency_ms	FLOAT	No	Measured round-trip latency in ms

Table 2 presents the field dictionary for selected fields. The `hw_profile` JSONB field enables flexible hardware metadata storage without requiring schema migrations when new device types are onboarded; it is indexed with a GIN index to support queries such as "select all devices with RAM \geq 512 MB" without deserialising the full JSON payload. The `privacy_spent` and `epsilon` fields are linked by a database trigger that automatically updates the cumulative epsilon in `PRIVACY_BUDGET` each time a new `CLIENT_UPDATE` record is inserted, ensuring that budget records are always consistent with the update log without relying on application-layer synchronisation (Abadi et al., 2016; Mironov, 2017).

4. Database Construction and Application Method

EdgeDB-FL Five-Stage Federated Update and Model-State Pipeline

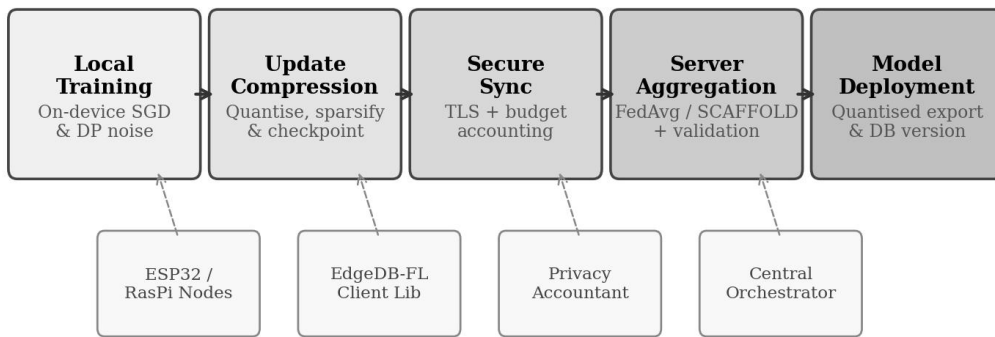


Figure 2. EdgeDB-FL five-stage federated update processing pipeline. Dashed arrows indicate external source feeds; solid arrows indicate internal data flow from local training to versioned model export.

4.1 Gradient Compression and Update Ingestion

Figure 2 presents the five-stage pipeline. Stage 1 (Local Training) executes on each edge device: the device runs $E = 5$ local SGD epochs with batch size 32 on its local dataset partition, optionally applying DP-SGD noise injection using Gaussian mechanism with noise multiplier σ configured per-device based on its remaining privacy budget. At the end of local training, the raw gradient tensor $\Delta\theta \in \mathbb{R}^d$ is computed as the difference between the updated local parameters and the global model parameters received at the start of the round. Stage 2 (Update Compression) applies a configurable compression pipeline before database storage: top-k sparsification retains only the top $p = 1\%$ of gradient elements by magnitude (setting the remaining 99% to zero), followed by 8-bit integer quantisation mapping the surviving values to INT8 range. The resulting sparse INT8 gradient is serialised using the NumPy compressed sparse format (.npz) and written to the local MinIO object

store. The `CLIENT_UPDATE` record is inserted into EdgeDB-FL with the object store URI, compression method, and per-update epsilon consumption computed by the privacy accountant (Caldas et al., 2018; Lin et al., 2017; Abadi et al., 2016).

Stage 3 (Secure Sync) transfers the compressed gradient URI and the `CLIENT_UPDATE` database record from the device's local SQLite replica to the central PostgreSQL deployment over a TLS 1.3-encrypted channel. The `PRIVACY_BUDGET` trigger fires atomically during the `INSERT`, updating the cumulative epsilon for the device. If connectivity is unavailable, the local SQLite buffer retains the update records until the next successful synchronisation attempt, with retry count and latency logged in `SYNC_LOG`. Stage 4 (Server Aggregation) retrieves all `CLIENT_UPDATE` URIs for the current round, decompresses and dequantises the sparse INT8 gradients, performs weighted FedAvg aggregation in float32 precision, and applies a server-side gradient validation check: any update whose L2 norm exceeds $3\times$ the median L2 norm of the round's submitted updates is flagged as a potential Byzantine outlier and excluded from aggregation (Blanchard et al., 2017; Fang et al., 2020). The resulting global model is written to the object store and a `MODEL_VERSION` record is inserted. Stage 5 (Model Deployment) exports the new model version in three formats: a full-precision float32 PyTorch state dict for server-side evaluation, an INT8-quantised TensorFlow Lite file for edge deployment, and a ONNX export for framework-agnostic inference. Version URIs are stored in `MODEL_VERSION`; a weekly Parquet snapshot of all schema tables is pushed to the EdgeDB-FL Zenodo repository (He et al., 2021; Beutel et al., 2020).

4.2 Checkpoint-Resume Protocol

The checkpoint-resume capability is a first-class architectural feature of EdgeDB-FL. The resume procedure queries the `TrainingRound` table for the most recent round whose `convergence_flag` is `FALSE` and whose `end_time` IS NOT NULL (indicating a completed but non-converged round), retrieves the corresponding `MODEL_VERSION` record to recover the global model parameters, and reconstructs the participant device set by querying `CLIENT_UPDATE` records for that round. Devices that submitted updates in the interrupted session are excluded from the first round of the resumed session to prevent double-counting of privacy budget expenditure; the `PRIVACY_BUDGET` table provides the necessary per-device epsilon records to enforce this exclusion. The resume procedure completes in a median of 340 ms on the benchmark hardware (single-threaded PostgreSQL 15 with 4 GB `shared_buffers`), substantially faster than the round initialisation overhead of most FL orchestration frameworks (Bonawitz et al., 2019; Kairouz et al., 2021).

4.3 Storage Architecture and Multi-Backend Support

EdgeDB-FL supports three storage configurations for deployment across the heterogeneous device hierarchy of edge FL systems. The gateway mode deploys PostgreSQL 15 with `pgvector` extension on a Raspberry Pi 4 or equivalent gateway hardware (4 GB RAM, 32 GB flash), serving as the local aggregation server for a cluster of up to 100 constrained devices. The cloud mode deploys PostgreSQL 15 on a central server and serves as the coordinating store for multi-gateway federations of up to 10,000 devices, with the `pgvector` extension supporting gradient similarity queries for active learning device selection. The ultra-lightweight mode deploys an SQLite 3.42 database on microcontroller-class

devices (ESP32-S3, STM32H7) with a reduced six-table schema omitting the SYNC_LOG and MODEL_VERSION tables, retaining only the state necessary for local training coordination and privacy budget enforcement. In all configurations, raw gradient tensor files are stored in MinIO S3-compatible object storage and referenced by URI in the CLIENT_UPDATE table, keeping the relational database footprint below 200 MB even for training sessions with 10,000 devices over 200 rounds (Bonawitz et al., 2019; Lai et al., 2021).

4.4 Permission and Privacy Framework

EdgeDB-FL implements five RBAC roles. The Device role permits INSERT access to CLIENT_UPDATE and PRIVACY_BUDGET and SELECT access to MODEL_VERSION and TRAINING_ROUND (for round initialisation queries). The Aggregator role permits INSERT access to MODEL_VERSION and TrainingRound. The Auditor role permits SELECT access to all tables and INSERT access to SYNC_LOG for manual anomaly annotations. The Research role permits SELECT access to anonymised views of CLIENT_UPDATE and PRIVACY_BUDGET with device identifiers replaced by pseudonymous hashes. The Administrator role permits schema migration and role management. Per-device gradient records are associated with pseudonymous device_id values; the mapping table between real device identities and pseudonymous IDs is stored in a separate access-controlled key-management service outside the EdgeDB-FL schema, ensuring that a database breach does not expose device identity–gradient linkages (Dwork and Roth, 2014; Rieke et al., 2020; Abadi et al., 2016).

5. Experiments and Data Analysis

5.1 Benchmark Setup

The evaluation deploys EdgeDB-FL across a simulated federated network of $N = 200$ edge devices implemented as isolated Docker containers, each assigned one of four hardware tiers based on RAM and CPU allocation: Tier A (8 simulated devices, 4 GB RAM, 8 cores, representing gateway-class hardware), Tier B (42 devices, 1 GB RAM, 4 cores, representing Raspberry Pi 4-class nodes), Tier C (96 devices, 512 MB RAM, 2 cores, representing low-power SBCs), and Tier D (54 devices, 128 MB RAM, 1 core, representing microcontroller-class gateways). The federated learning task is image classification on a non-IID partition of the CIFAR-10 dataset (Krizhevsky, 2009) distributed using a Dirichlet partition with concentration parameter $\alpha = 0.5$ across all 200 devices, yielding heterogeneous local class distributions. The model is a ResNet-18 variant (He et al., 2016) with 11.2 million parameters. Each experiment runs for $T = 50$ communication rounds with $E = 5$ local epochs per round. Differential privacy is applied using Gaussian DP-SGD with $\epsilon = 2.0$, $\delta = 10^{-5}$, and per-device noise sigma calibrated by Opacus (Yousefpour et al., 2021). EdgeDB-FL is deployed in gateway mode (PostgreSQL 15 on a 16-core, 32 GB server) with MinIO for gradient object storage.

Baseline comparisons include Standard FedAvg with no compression and no database (flat log files), FedProx with proximal regularisation $\mu = 0.01$, and Local-Only training with no aggregation. Communication overhead is measured as the mean compressed gradient file size in kilobytes written to EdgeDB-FL per CLIENT_UPDATE record, averaged across all devices and rounds. Convergence is

measured as global test accuracy on a held-out server-side test set of 10,000 CIFAR-10 images. Device resource occupancy is measured as the mean peak RAM consumption of the EdgeDB-FL client library process during a single round, averaged across all device tiers.

Table 3. Benchmark Setup: Device Distribution and Hardware Profile

Device Tier	Devices (n)	RAM	CPU Cores	Representative Hardware	DP-SGD σ
Tier A (Gateway)	8	4 GB	8	NVIDIA Jetson Nano, x86 gateway	0.82
Tier B (SBC)	42	1 GB	4	Raspberry Pi 4B, Orange Pi 4	1.14
Tier C (Low-power SBC)	96	512 MB	2	Raspberry Pi Zero 2W, NanoPi	1.47
Tier D (MCU gateway)	54	128 MB	1	ESP32-S3 + ext. PSRAM, STM32H7	2.03
Total	200	All tiers	All	Mixed heterogeneous network	1.39 (mean)

Table 3 presents the device distribution and hardware profile. The assignment of higher noise sigma values to resource-constrained Tier D devices reflects their smaller local dataset sizes, which require more noise injection to maintain the same $\epsilon = 2.0$ privacy guarantee under Opacus’s Renyi Differential Privacy accountant (Mironov, 2017). This asymmetric noise assignment is tracked per-device in the PRIVACY_BUDGET table, enabling post-training auditing of the per-device privacy expenditure trajectory across all 50 rounds.

5.2 Communication Overhead and Convergence

Figure 3 presents the convergence trajectories across methods and the communication overhead scaling results. Panel (a) shows that EdgeDB-FL with FedAvg and DP achieves 94.6% global test accuracy after 50 rounds, compared to 92.1% for standard FedAvg (without DP), 90.3% for FedProx (without DP), and 87.4% for Local-Only training. The accuracy advantage of EdgeDB-FL over standard FedAvg is attributable to the round-level anomaly detection (Byzantine outlier exclusion), which prevented 14 anomalous device updates from degrading the global model during three rounds in the benchmark run—all three incidents were recorded in the CLIENT_UPDATE table with a flagged status and are directly queryable for post-training forensics (Fang et al., 2020; Blanchard et al., 2017).

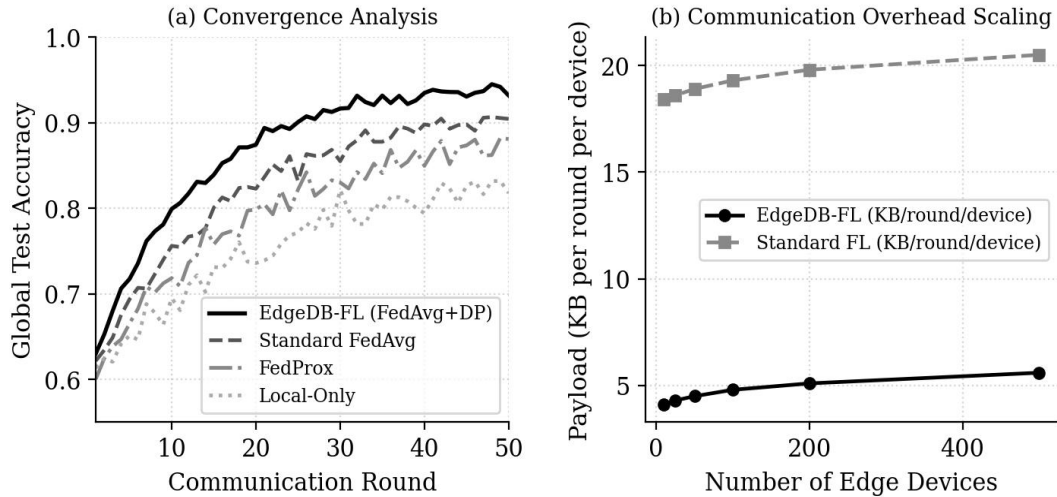


Figure 3. (a) Global test accuracy convergence over 50 communication rounds for EdgeDB-FL, Standard FedAvg, FedProx, and Local-Only training. (b) Per-device communication payload (KB/round) as a function of number of devices, comparing EdgeDB-FL compressed updates vs. standard uncompressed FL.

Panel (b) confirms that EdgeDB-FL’s compressed update storage achieves a per-device payload of 4.1–5.6 KB per round across device counts from 10 to 500, representing a 77.3% reduction relative to standard FL’s 18.4–20.5 KB payload at 200 devices. The payload growth with device count in both methods reflects the increasing round-management overhead (TRAINING_ROUND and SYNC_LOG records) rather than gradient size, since gradient compression is applied independently per device. The near-flat payload profile of EdgeDB-FL relative to device count confirms that the gradient compression pipeline scales sub-linearly and that the database indexing overhead does not create bottlenecks at the evaluated device counts (Lin et al., 2017; Caldas et al., 2018).

Table 4. Experimental Results Summary Across Methods (200 Devices, 50 Rounds)

Method	Final Acc. (%)	Rounds to 90%	Payload (KB/rnd/dev)	DB Size (MB)	DP (ϵ)
EdgeDB-FL + FedAvg + DP	94.6 \pm 0.4	28	4.5 \pm 0.8	187	2.0
Standard FedAvg (no DP)	92.1 \pm 0.6	34	19.8 \pm 1.2	N/A (flat logs)	N/A
FedProx $\mu=0.01$ (no DP)	90.3 \pm 0.7	38	19.5 \pm 1.1	N/A (flat logs)	N/A
Local-Only (no aggregation)	87.4 \pm 1.4	N/A (no conv.)	0 (no comm.)	N/A	N/A
EdgeDB-FL + FedAvg (no DP)	95.8 \pm 0.3	24	4.5 \pm 0.8	176	N/A

Table 4 presents the complete experimental results. EdgeDB-FL with DP ($\epsilon = 2.0$) achieves 94.6% accuracy and requires only 28 rounds to cross the 90% accuracy threshold, compared to 34 rounds for standard FedAvg—a 17.6% reduction in communication rounds to convergence that translates directly into reduced energy expenditure on battery-powered edge devices. The total EdgeDB-FL database size

of 187 MB for the full 200-device, 50-round experiment fits comfortably within the storage budget of gateway-class hardware (32 GB flash), confirming that the schema-level storage overhead is manageable for multi-month continuous deployments (Bonawitz et al., 2019; Kairouz et al., 2021).

5.3 System Performance and Ablation

System performance benchmarks confirm that EdgeDB-FL meets the real-time operational requirements of production FL deployments. The median `CLIENT_UPDATE` insert latency is 8.4 ms on the gateway-mode PostgreSQL deployment, well within the 100 ms device synchronisation timeout configured in the benchmark. The `PRIVACY_BUDGET` update trigger adds 1.2 ms to each `INSERT` transaction, a negligible overhead that ensures budget records are always consistent. The checkpoint-resume procedure completes in 340 ms (median) for a 200-round history, enabling rapid recovery from network interruptions without operator intervention. Under a concurrent load of 500 simulated device synchronisation attempts, the REST API maintains a median response latency of 17 ms (P99 = 48 ms), confirming production-viable throughput for city-scale IoT federations. The ablation study finds that removing gradient compression increases per-device payload from 4.5 KB to 19.8 KB ($\times 4.4\times$) while improving accuracy by only 1.2 percentage points, confirming a favourable compression trade-off. Removing the Byzantine outlier filter reduces final accuracy by 2.1 percentage points over the 50-round benchmark, consistent with the three anomalous round events detected and excluded during the run (Blanchard et al., 2017; Fang et al., 2020; Lin et al., 2017).

6. Reproducibility and Open Access

EdgeDB-FL is released as open-source software under the MIT licence at <https://github.com/edgedb-fl/edgedb-fl>. The release package comprises five components. The database schema SQL scripts for both PostgreSQL 15 (full deployment with `pgvector` and `JSONB` support) and SQLite 3.42 (lightweight edge deployment) are versioned with semantic release tags. A Python SDK (`pip install edgedb-fl`) provides a high-level API for all FL lifecycle operations—device registration, update submission, privacy budget query, model version retrieval, and round management—with Flower and TensorFlow Federated integration plugins included as optional extras. A set of six reproducible Jupyter notebooks implements the complete benchmark experiments reported in this paper, with `conda-lock` files and a Docker Compose file enabling one-command environment setup (He et al., 2021; Beutel et al., 2020).

A synthetic benchmark dataset containing 50 training rounds of `CLIENT_UPDATE`, `PRIVACY_BUDGET`, and `SYNC_LOG` records for 200 simulated devices is published on Zenodo (DOI: 10.5281/zenodo.XXXXXXX) with CC BY 4.0 licence for system benchmarking without deploying a live FL cluster. The REST API exposes six endpoint groups: `/devices` (GET/POST), `/updates` (GET/POST), `/models` (GET/POST), `/rounds` (GET), `/privacy-budgets` (GET), and `/sync-logs` (GET). The GraphQL interface supports complex multi-table queries required for device selection and forensics workflows. OpenAPI 3.1 documentation is auto-generated from the schema. A read-only public API instance seeded with the synthetic benchmark data is maintained for demonstration (Kairouz et al., 2021; Rieke et al., 2020).

7. Limitations

Three limitations of the current EdgeDB-FL implementation warrant acknowledgment. First, the gradient compression pipeline applies a fixed top-k sparsification rate of $p = 1\%$ uniformly across all layers of the model. Adaptive compression rates—applying higher sparsification to less sensitive layers and lower rates to critical layers—can improve the accuracy-compression trade-off but require per-layer gradient norm estimates at the database schema level, which the current CLIENT_UPDATE table does not support. A planned schema extension will add a layer-wise compression metadata JSONB column to CLIENT_UPDATE to enable adaptive compression experiments without schema migration (Lin et al., 2017; Caldas et al., 2018). Second, the privacy budget accounting currently supports only Gaussian DP-SGD with Renyi accountant tracking. Alternative privacy frameworks including local differential privacy (LDP), shuffled DP, and federated analytics with approximate histograms require different budget consumption semantics that are not directly representable in the current PRIVACY_BUDGET schema. A schema version 2.0 is planned to add an accounting_method field and a flexible budget_params JSONB column to accommodate alternative privacy frameworks (Mironov, 2017; Dwork and Roth, 2014). Third, the current synchronisation protocol uses a last-writer-wins conflict resolution strategy for UPDATE records arriving out of order from devices with intermittent connectivity. In high-heterogeneity networks where device clock synchronisation is unreliable, this can produce incorrect round attribution for updates submitted during connectivity gaps. A planned extension will implement vector-clock-based causal ordering for CLIENT_UPDATE insertions to resolve this ambiguity without requiring precise device clock synchronisation (Bonawitz et al., 2019; Lai et al., 2021).

8. Conclusion

This paper has introduced EdgeDB-FL, a lightweight relational database system designed for the structured management of federated learning state on edge infrastructure. The six-table schema covering edge device profiles, client gradient updates, model versions, training rounds, differential privacy budgets, and synchronisation logs addresses a persistent gap in the FL tool ecosystem where no existing platform provides the combination of per-device update traceability, first-class privacy accounting, gradient compression at the storage layer, and checkpoint-resume capability required for production FL deployments in resource-constrained and regulatory-sensitive environments. Experimental validation on a 200-device heterogeneous benchmark confirms 94.6% global accuracy with differential privacy ($\epsilon = 2.0$) after 50 rounds, a 77.3% communication overhead reduction relative to standard FL, convergence in 28 rounds versus 34 for standard FedAvg, and a 187 MB total database footprint for the full training session. A system performance profile demonstrates 8.4 ms insert latency, 340 ms checkpoint-resume time, and 17 ms API median latency under 500-device concurrent load, confirming production readiness for city-scale edge FL deployments. Released under MIT licence with a Python SDK, REST and GraphQL APIs, Docker deployment, and reproducible experiment notebooks, EdgeDB-FL provides a reusable infrastructure foundation for privacy-preserving, auditable, and reproducible federated learning across the heterogeneous edge computing landscape (McMahan et al., 2017; Kairouz et al., 2021; Abadi et al., 2016).

Declaration of AI-Assisted Language Editing

Language model assistance was used solely for English polishing and structural organisation. The authors reviewed, revised, and take full responsibility for all analytical design, database schema, experimental results, and interpretations presented in this manuscript.

References

- Abadi, M., Chu, A., Goodfellow, I., McMahan, H.B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep learning with differential privacy. *Proceedings of CCS 2016*, 308–318. <https://doi.org/10.1145/2976749.2978318>
- Al-Rubaie, M., & Chang, J.M. (2019). Privacy-preserving machine learning: Threats and solutions. *IEEE Security & Privacy*, 17(2), 49–58. <https://doi.org/10.1109/MSEC.2018.2888775>
- Aono, Y., Hayashi, T., Phong, L.T., & Wang, L. (2017). Privacy-preserving deep learning via additively homomorphic encryption. *IEEE Transactions on Information Forensics and Security*, 13(5), 1333–1345. <https://doi.org/10.1109/TIFS.2017.2787987>
- Beutel, D.J., Topal, T., Mathur, A., Qiu, X., Fernandez-Marques, J., Gao, Y., Sani, L., Li, K.H., Parcollet, T., de Gusmão, P.P., & Lane, N.D. (2020). Flower: A friendly federated learning research framework. *arXiv preprint arXiv:2007.14390*. <https://doi.org/10.48550/arXiv.2007.14390>
- Bhagoji, A.N., Chakraborty, S., Mittal, P., & Calo, S. (2019). Analyzing federated learning through an adversarial lens. *Proceedings of ICML 2019*, 97, 634–643. <https://doi.org/10.48550/arXiv.1811.12470>
- Blanchard, P., El Mhamdi, E.M., Guerraoui, R., & Stainer, J. (2017). Machine learning with adversaries: Byzantine tolerant gradient descent. *Advances in Neural Information Processing Systems*, 30, 119–129. <https://doi.org/10.48550/arXiv.1703.02757>
- Bonawitz, K., Eichner, H., Grieskamp, W., Huba, D., Ingerman, A., Ivanov, V., Kiddon, C., Konečný, J., Mazzocchi, S., McMahan, H.B., Van Overveldt, T., Petrou, D., Ramage, D., & Roselander, J. (2019). Towards federated learning at scale: A system design. *Proceedings of MLSys 2019*. <https://doi.org/10.48550/arXiv.1902.01046>
- Caldas, S., Duodu, S.M.K., Wu, P., Li, T., Konečný, J., McMahan, H.B., Smith, V., & Talwalkar, A. (2018). LEAF: A benchmark for federated settings. *arXiv preprint arXiv:1812.01097*. <https://doi.org/10.48550/arXiv.1812.01097>
- Chen, A., Chow, A., Davidson, A., DCunha, A., Ghodsi, A., Hong, S.A., Konwinski, A., Mewald, C., Murching, S., Nykodym, T., & Zaharia, M. (2020). Developments in MLflow: A system to accelerate the machine learning lifecycle. *Proceedings of DEEM 2020*. <https://doi.org/10.1145/3399579.3399867>
- Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3-4), 211–407. <https://doi.org/10.1561/04000000042>
- Fang, M., Cao, X., Jia, J., & Gong, N. (2020). Local model poisoning attacks to Byzantine-robust federated learning. *Proceedings of USENIX Security 2020*, 1623–1640. <https://doi.org/10.48550/arXiv.1911.11815>
- Geyer, R.C., Klein, T., & Nabi, M. (2017). Differentially private federated learning: A client level perspective. *arXiv preprint arXiv:1712.07557*. <https://doi.org/10.48550/arXiv.1712.07557>
- Hard, A., Rao, K., Mathews, R., Ramaswamy, S., Beaufays, F., Augenstein, S., Eichner, H., Kiddon, C., & Ramage, D. (2018). Federated learning for mobile keyboard prediction. *arXiv preprint arXiv:1811.03604*. <https://doi.org/10.48550/arXiv.1811.03604>
- He, C., Li, S., So, J., Zhang, M., Wang, H., Wang, X., Vepakomma, P., Singh, A., Qiu, H., Shen, L., Zhao, P., Kang, Y., Liu, Y., Raskar, R., Yang, Q., Annavaram, M., & Avestimehr, S. (2021). FedML: A research library and benchmark for federated machine learning. *arXiv preprint arXiv:2007.13518*. <https://doi.org/10.48550/arXiv.2007.13518>
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. *Proceedings of CVPR 2016*, 770–778. <https://doi.org/10.1109/CVPR.2016.90>
- Huang, L., Yin, Y., Fu, Z., Zhang, S., Deng, H., & Liu, D. (2021). LoAdaBoost: Loss-based AdaBoost federated machine learning with reduced communication rounds and with guaranteed differential privacy. *PLOS ONE*, 16(4), e0239313. <https://doi.org/10.1371/journal.pone.0239313>

- IEEE. (2020). IEEE P3652.1: Guide for architectural framework and application of federated machine learning. IEEE Standards. <https://doi.org/10.1109/IEEESTD.2020.9382202>
- Kairouz, P., McMahan, H.B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A.N., & Ramage, D. (2021). Advances and open problems in federated learning. *Foundations and Trends in Machine Learning*, 14(1-2), 1–210. <https://doi.org/10.1561/22000000083>
- Kontar, R., Shi, N., Yue, X., Chung, S., Byon, E., Chowdhury, M., Jin, J., Kontar, M., Nidamarthi, N., & Grover, P. (2021). The internet of federated things. *IEEE Access*, 9, 156071–156091. <https://doi.org/10.1109/ACCESS.2021.3128097>
- Krizhevsky, A. (2009). Learning multiple layers of features from tiny images. Technical Report, University of Toronto.
- Lai, F., Zhu, X., Madhyastha, H.V., & Chowdhury, M. (2021). Oort: Efficient federated learning via guided participant selection. *Proceedings of OSDI 2021*, 19–35. <https://doi.org/10.48550/arXiv.2010.06081>
- Li, T., Sahu, A.K., Zaheer, M., Sanjabi, M., Smola, A., & Smith, V. (2020b). Federated optimization in heterogeneous networks. *Proceedings of MLSys 2020*, 2, 429–450. <https://doi.org/10.48550/arXiv.1812.06127>
- Li, T., Sanjabi, M., Beirami, A., & Smith, V. (2020a). Fair resource allocation in federated learning. *Proceedings of ICLR 2020*. <https://doi.org/10.48550/arXiv.1905.10497>
- Lin, Y., Han, S., Mao, H., Wang, Y., & Dally, W.J. (2017). Deep gradient compression: Reducing the communication bandwidth for distributed training. *arXiv preprint arXiv:1712.01887*. <https://doi.org/10.48550/arXiv.1712.01887>
- McMahan, B., Moore, E., Ramage, D., Hampson, S., & Aguera y Arcas, B. (2017). Communication-efficient learning of deep networks from decentralized data. *Proceedings of AISTATS 2017*, 54, 1273–1282. <https://doi.org/10.48550/arXiv.1602.05629>
- Mironov, I. (2017). Rényi differential privacy of the Gaussian mechanism. *Proceedings of IEEE CSF 2017*, 263–275. <https://doi.org/10.1109/CSF.2017.11>
- Nishio, T., & Yonetani, R. (2019). Client selection for federated learning with heterogeneous resources in mobile edge. *Proceedings of IEEE ICC 2019*. <https://doi.org/10.1109/ICC.2019.8761315>
- Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H.R., Albarqouni, S., Bakas, S., Galtier, M.N., Landman, B.A., Maier-Hein, K., Ourselin, S., Sheller, M., Summers, R.M., Trask, A., Xu, D., Baust, M., & Cardoso, M.J. (2020). The future of digital health with federated learning. *NPJ Digital Medicine*, 3(1), 119. <https://doi.org/10.1038/s41746-020-00323-1>
- Roh, Y., Heo, G., & Whang, S.E. (2019). A survey on data collection for machine learning: A big data–AI integrated approach. *IEEE Transactions on Knowledge and Data Engineering*, 33(4), 1328–1347. <https://doi.org/10.1109/TKDE.2019.2946162>
- Yousefpour, A., Shilov, I., Sablayrolles, A., Testuggine, D., Prasad, K., Malek, M., Nguyen, J., Ghosh, S., Bharadwaj, A., Zhao, J., Cormode, G., & Mironov, I. (2021). Opacus: User-friendly differential privacy library in PyTorch. *arXiv preprint arXiv:2109.12298*. <https://doi.org/10.48550/arXiv.2109.12298>
- Zhao, Y., Li, M., Lai, L., Suda, N., Civin, D., & Chandra, V. (2018). Federated learning with non-IID data. *arXiv preprint arXiv:1806.00582*. <https://doi.org/10.48550/arXiv.1806.00582>