Accelerating Postgres Performance for IoT with Swarm64

Managing IoT data with free, open source Postgres

IoT data requirements can be extreme:

- Ingesting millions of sensor readings per second
- Real-time querying and data aggregations
- Many concurrent users monitoring data dashboards
- Complex analytics over many terabytes of historic data

Swarm64 database acceleration for Postgres enables data engineers to meet IoT challenges with free, open source Postgres, rather than using costly commercial SQL databases or less flexible NoSQL and time series databases.

In 1TB TPC-H benchmarks, Swarm64 accelerates free, open source Postgres by orders of magnitude as shown to the right.

50x
faster queries

35x
more inserts / sec.

67%
lower cost

How Swarm64 accelerates Postgres for IoT

Swarm64 extends free, open source Postgres with the following acceleration features, and doesn't require any changes to your SQL or application code:

Greater Parallelism
Swarm64 extends the Postgres query engine, rewriting query patterns that execute in parallel at every phase of the query. It parallelizes scanning, filtering, joining, and merging, etc.

Columnar indexing
Swarm64 loads and queries data stored in a columnar-indexed format within Postgres foreign data tables (FDW). The format is optimized for highly parallel access. Partitioning and range indexes reduce IO.
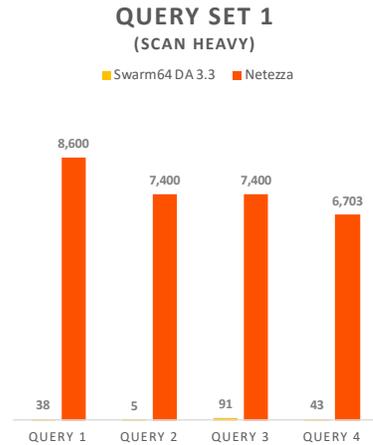
IO reduction
Swarm64 data is compressed – 5x to 25x depending on the data type. Besides reducing storage costs, reading compressed data reduces IO, often by ~20x relative to standard Postgres.

FPGA-based MPP
If an FPGA is present on the server, Swarm64 initiates 100+ SQL reader & writer processes on the FPGA. The processes work in parallel to accelerate queries and data insertion. It's like adding 100 vCores to your server—massive parallelism (MPP) at a very low cost, to handle real-time insertion and querying, greater concurrency, or very complex queries.
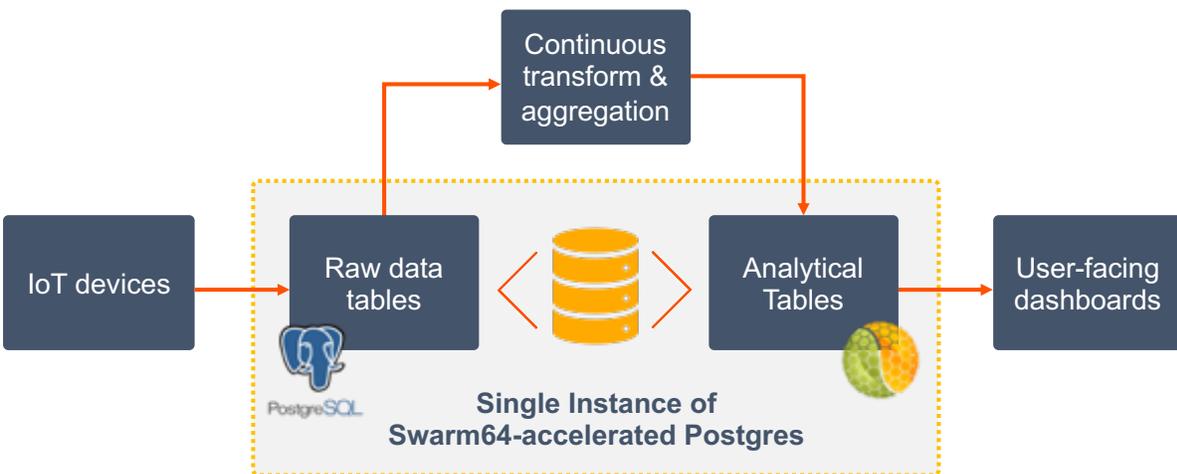
Benefits

Swarm64 is the easiest way to accelerate and scale Postgres performance. Benefits of the Swarm64 approach include:

- **Easy high performance** – Experience speed ups of 10x to 150x, without rearchitecting, migrating, or recoding.
- **Stay on free, open source Postgres** – no disruptive migrations or proprietary forks.
- **Low scaling costs** – Acceleration allows you to scale out on 70%-90% fewer servers.
- **Simpler architecture** – A single, FPGA-equipped server can handle mixed workloads efficiently - CPUs work on updates while FPGA executes queries.

QUERY SET 1
(SCAN HEAVY)

■ Swarm64 DA 3.3   ■ Netezza

	QUERY 1	QUERY 2	QUERY 3	QUERY 4
Swarm64 DA 3.3	38	5	91	43
Netezza	8,600	7,400	7,400	6,703

Typical IoT architecture

The diagram below shows a typical IoT data processing architecture featuring a Swarm64-accelerated Postgres server. Raw data from the devices is inserted into Postgres, and as it streams in, Swarm64 continuously aggregates it by different dimensions for real-time or off-line analytics. Parallelism and partitioning enable real-time aggregation of data as it streams into the DBMS, often at over 1M data points per second. On an FPGA-equipped server, the querying will execute on the FPGA, which leaves the CPUs free to handle the data insertion.

Continuous transform & aggregation

IoT devices → Raw data tables ⟷ ⟷ Analytical Tables → User-facing dashboards

PostgreSQL

Single Instance of Swarm64-accelerated Postgres

Swarm64 runs on premises or on the cloud on any Linux server. Swarm64 FPGA acceleration supports Intel and Xilinx FPGAs, and most FPGA-equipped cloud instances such as the Amazon EC2 F1 instances.

Next Steps

If you're interested in Swarm64's Postgres acceleration for your IoT systems, please visit Swarm64.com for more information. Contact us at info@swarm64.com to schedule a demo or request an evaluation copy of Swarm64.