INTEL® OPTANE™ DC PERSISTENT MEMORY MODULES INTEL® FPGA

Mikhail Tsvetkov

NOTICES AND DISCLAIMERS

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration.

No product or component can be absolutely secure.

Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. For more complete information about performance and benchmark results, visit http://www.intel.com/benchmarks .

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit http://www.intel.com/benchmarks .

Intel® Advanced Vector Extensions (Intel® AVX)* provides higher throughput to certain processor operations. Due to varying processor power characteristics, utilizing AVX instructions may cause a) some parts to operate at less than the rated frequency and b) some parts with Intel® Turbo Boost Technology 2.0 to not achieve any or maximum turbo frequencies. Performance varies depending on hardware, software, and system configuration and you can learn more at http://www.intel.com/go/turbo.

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Cost reduction scenarios described are intended as examples of how a given Intel-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.

Intel, the Intel logo, and Intel Xeon are trademarks of Intel Corporation in the U.S. and/or other countries.

*Other names and brands may be claimed as property of others.

INTEL® OPTANE™ DC PERSISTENT MEMORY MODULES (DCPMM)

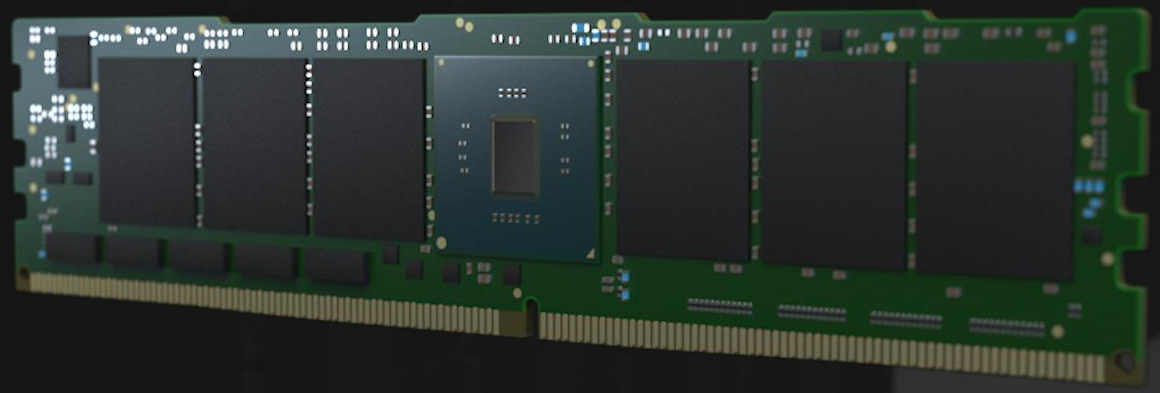

DDR4 МОДУЛИ ЕМКОСТЬЮ
128GB, 256GB, 512GB

ВЫСОКАЯ НАДЕЖНОСТЬ
И ДОЛГОВЕЧНОСТЬ

БЕЗОПАСТНОСТЬ НА
АППАРАТНОМ УРОВНЕ

INTEL® OPTANE™ DC PERSISTENT MEMORY

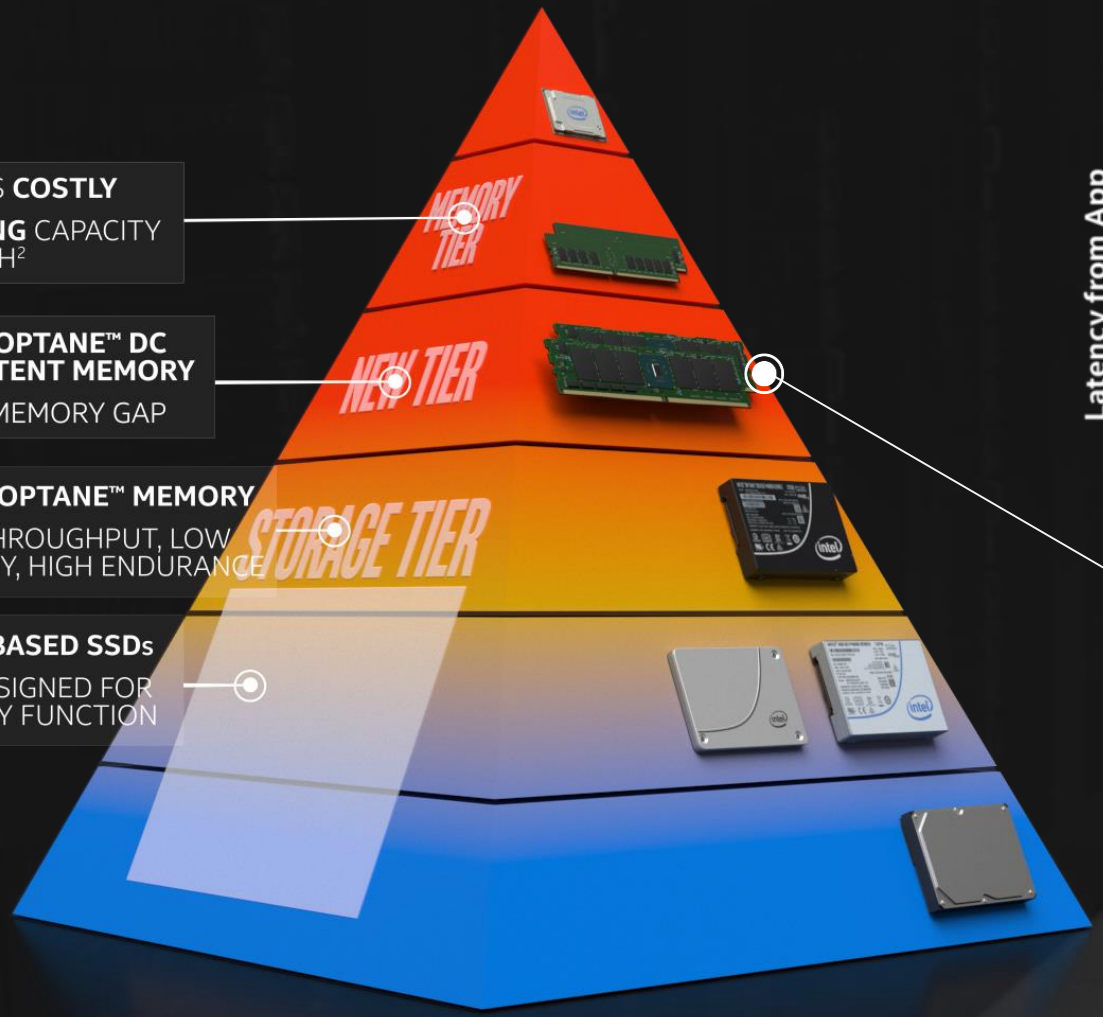

DRAM IS **COSTLY**
SLOWING CAPACITY GROWTH[2]

INTEL® OPTANE™ DC PERSISTENT MEMORY
CLOSE MEMORY GAP

INTEL® OPTANE™ MEMORY
HIGH THROUGHPUT, LOW LATENCY, HIGH ENDURANCE

NAND-BASED SSDs
NOT DESIGNED FOR MEMORY FUNCTION

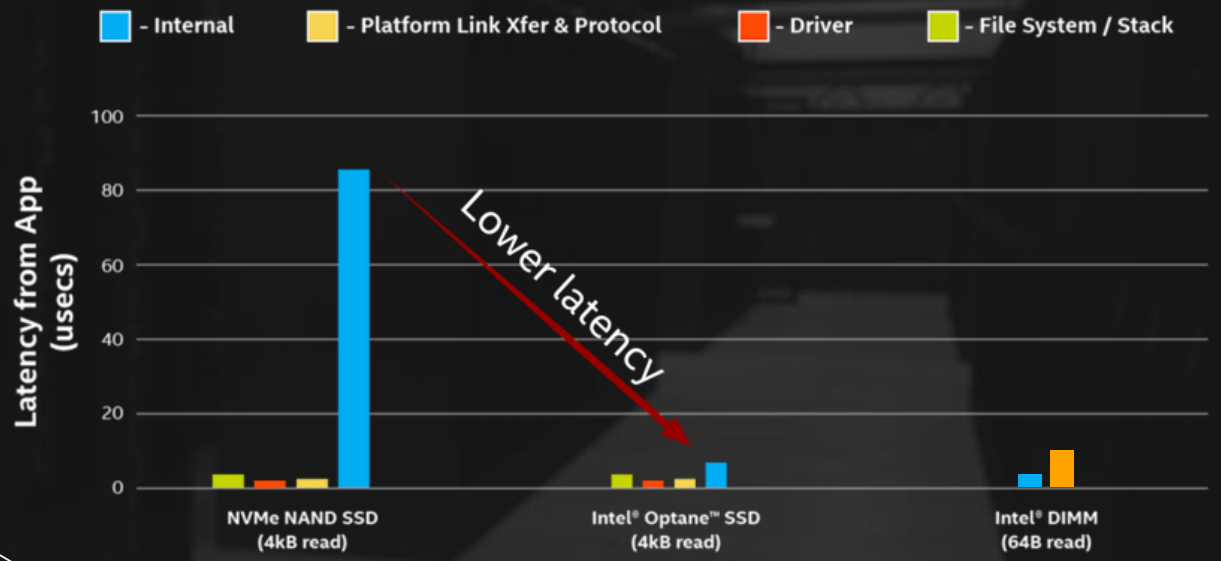

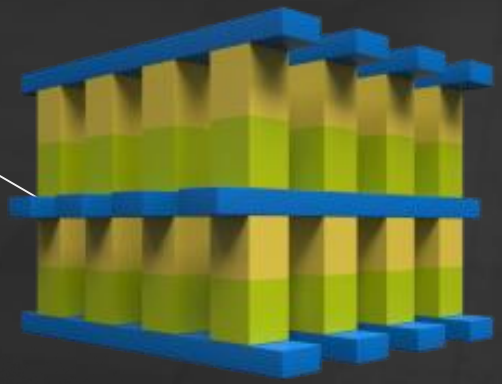

РЕЖИМ APP DIRECT – ПОСТОЯННАЯ ПАМЯТЬ

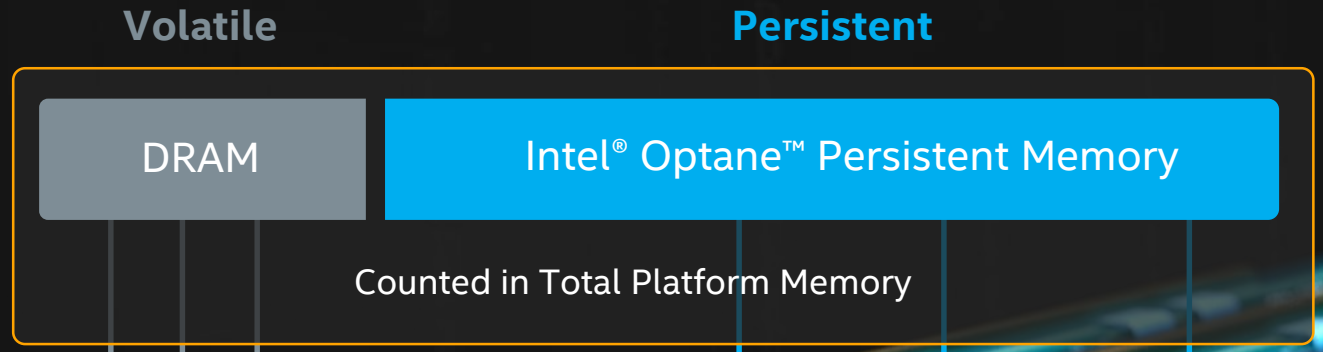

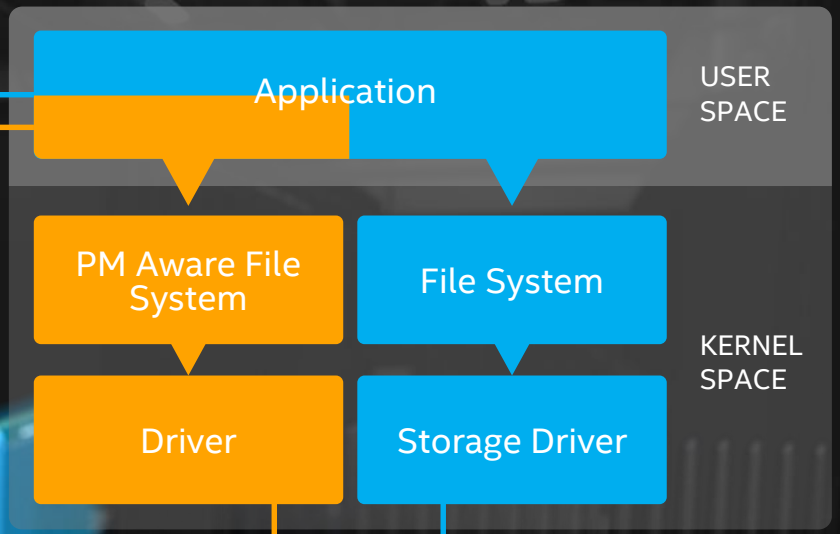

- ✓ Совместимые ОС знают о двух типа памяти
- ✓ Совместимые приложения могут распределять данные между DRAM и DCPMM
- ✓ Преимущества постоянной памяти системам надежность за счет быстрого рестарта

ACCELERATING WORKLOAD APPLICATIONS


Data Analytics & Databases


Financial


HPC (e.g. Genomics)


Media Transcoding


Cyber Security


Artificial Intelligence

Intel® FPGA Programmable Acceleration Cards (PAC) Qualified on Industry-Leading Enterprise Servers

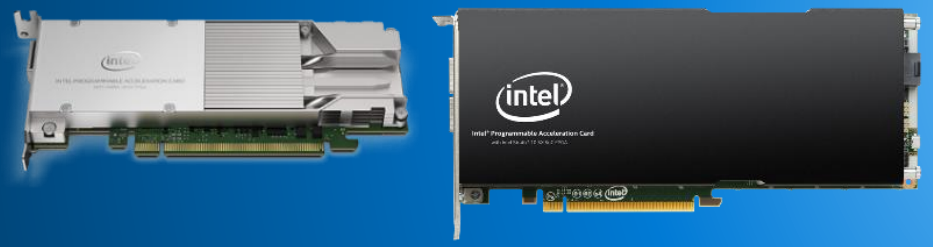

Acceleration Stack for Intel® Xeon® CPU with FPGA

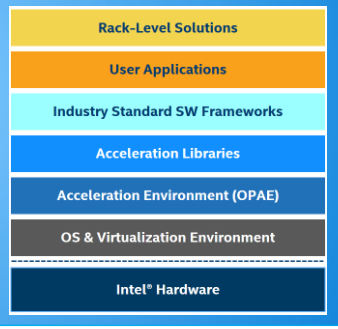

Acceleration IP From Application Experts

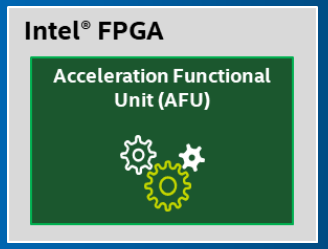

Scale-In before you Scale-Out
HTAP single node case with your Postgres

What Swarm64 Data Accelerator adds to Postgres

Swartm64 DA is an extension

Analytic Performance for big data sets (TBs)

Multi-user Performance

Query speed stability Multi-user

Contributed with open source patches and benchmarking tool

https://swarm64.com/open-source/

https://github.com/swarm64

Swarm64 HTAP Architecture: Extending Postgres

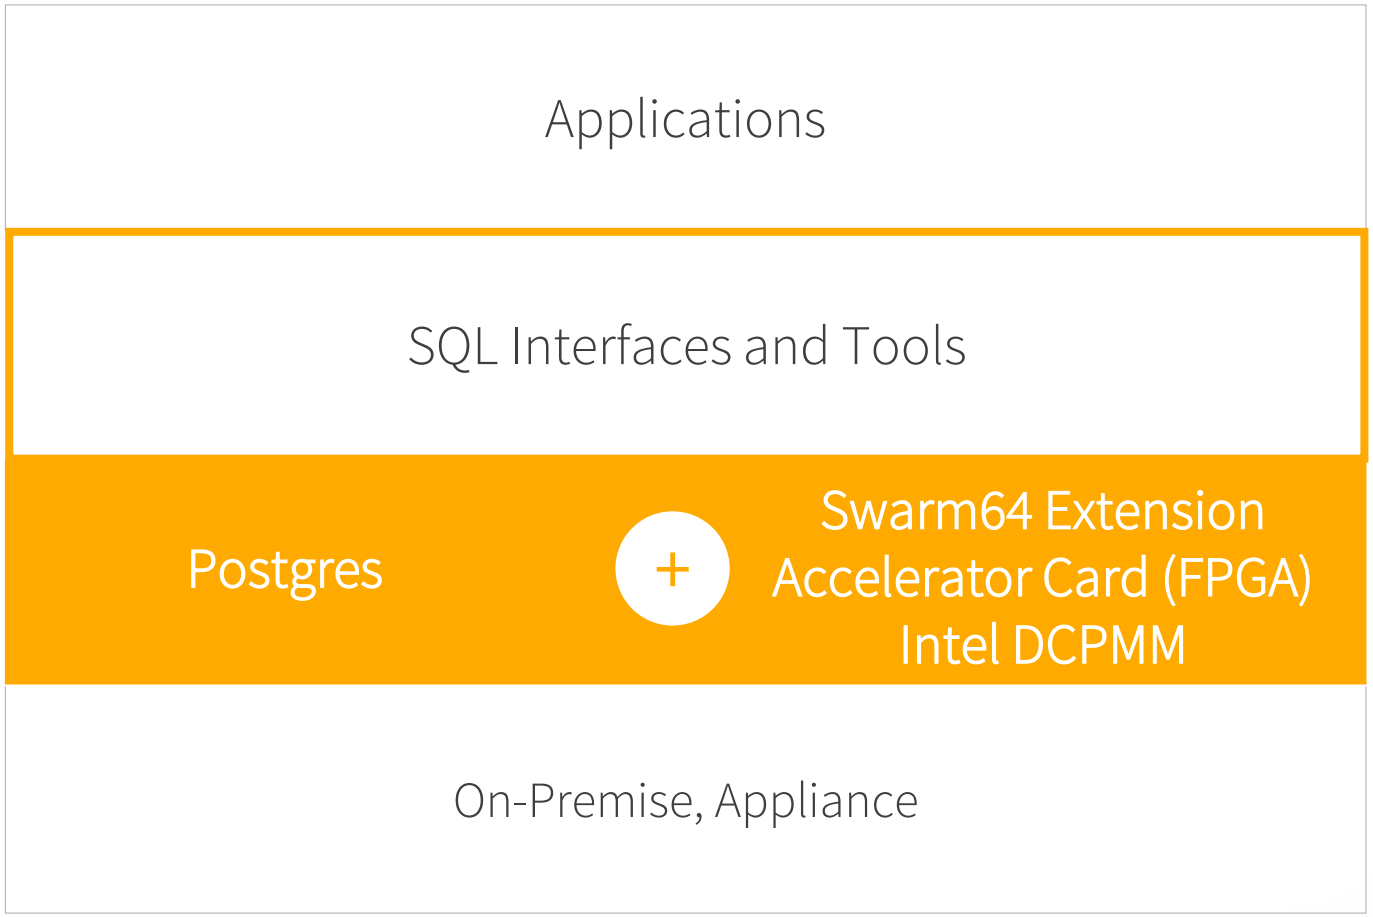

What is Swarm64 Solution: Accelerated with FPGAs

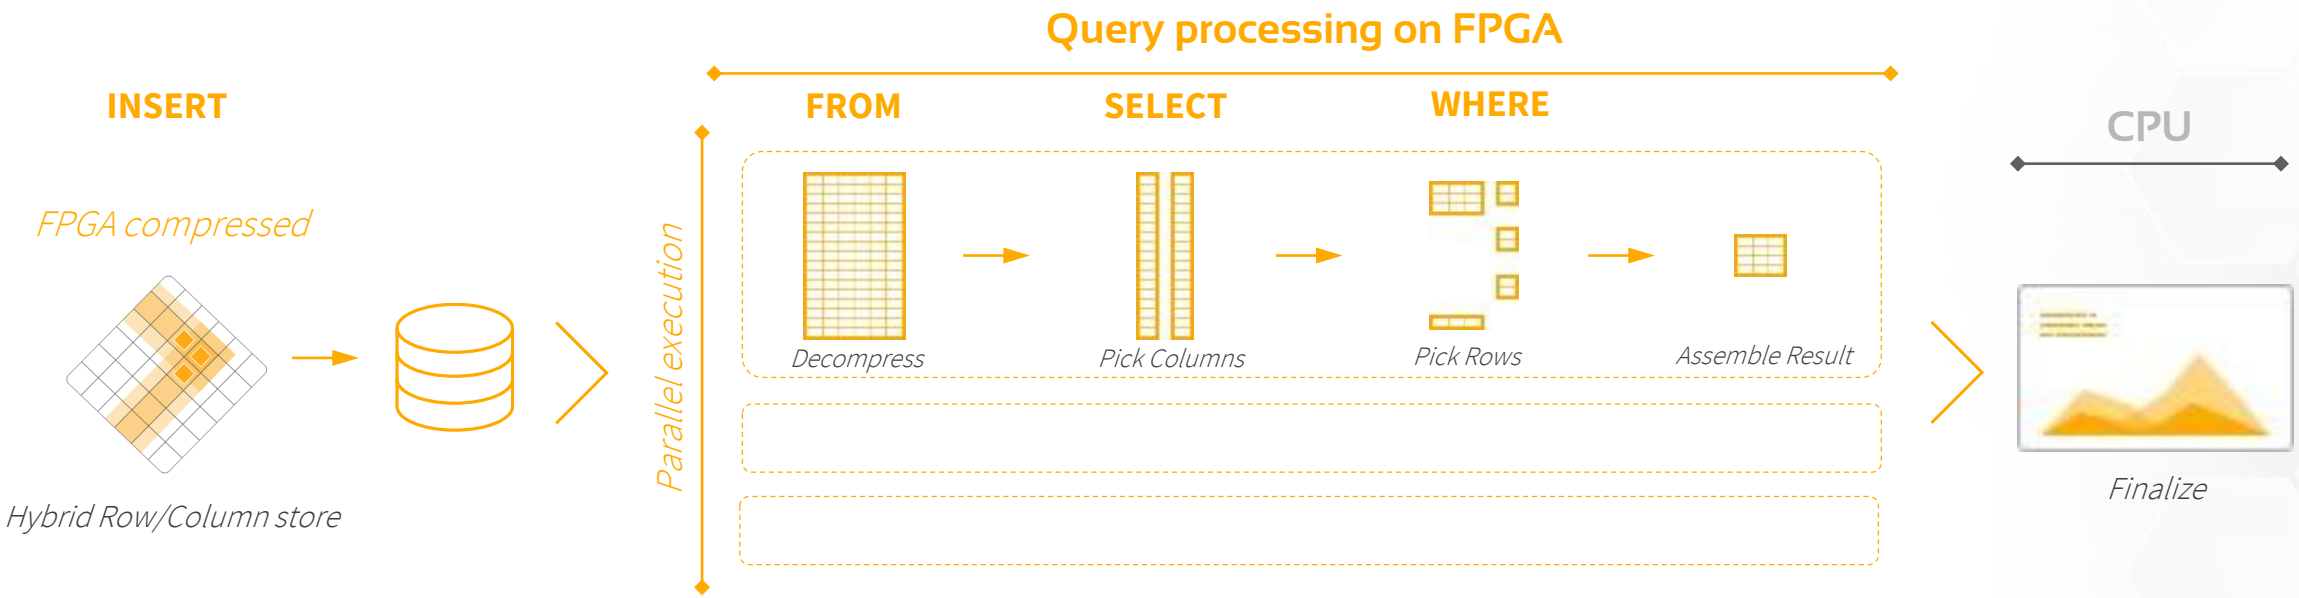

FPGA supports CPU for best overall system performance:

Loading data into database (INSERT, COPY FROM), keeping data at rest and in cache compressed

Scanning (SELECT) and filtering (WHERE)

Including complex processing such as time-series queries or full text search with wildcards (LIKE)

Swarm64 Data Accelerator TPC-H SF-1T Performance

Performance Gains

80x on Scanning-heavy queries (Q6)

25-35x Ingestion Speed

5x Less Storage

5x on Multi-user (3 streams)

Cost Savings

Less investment in machines and support

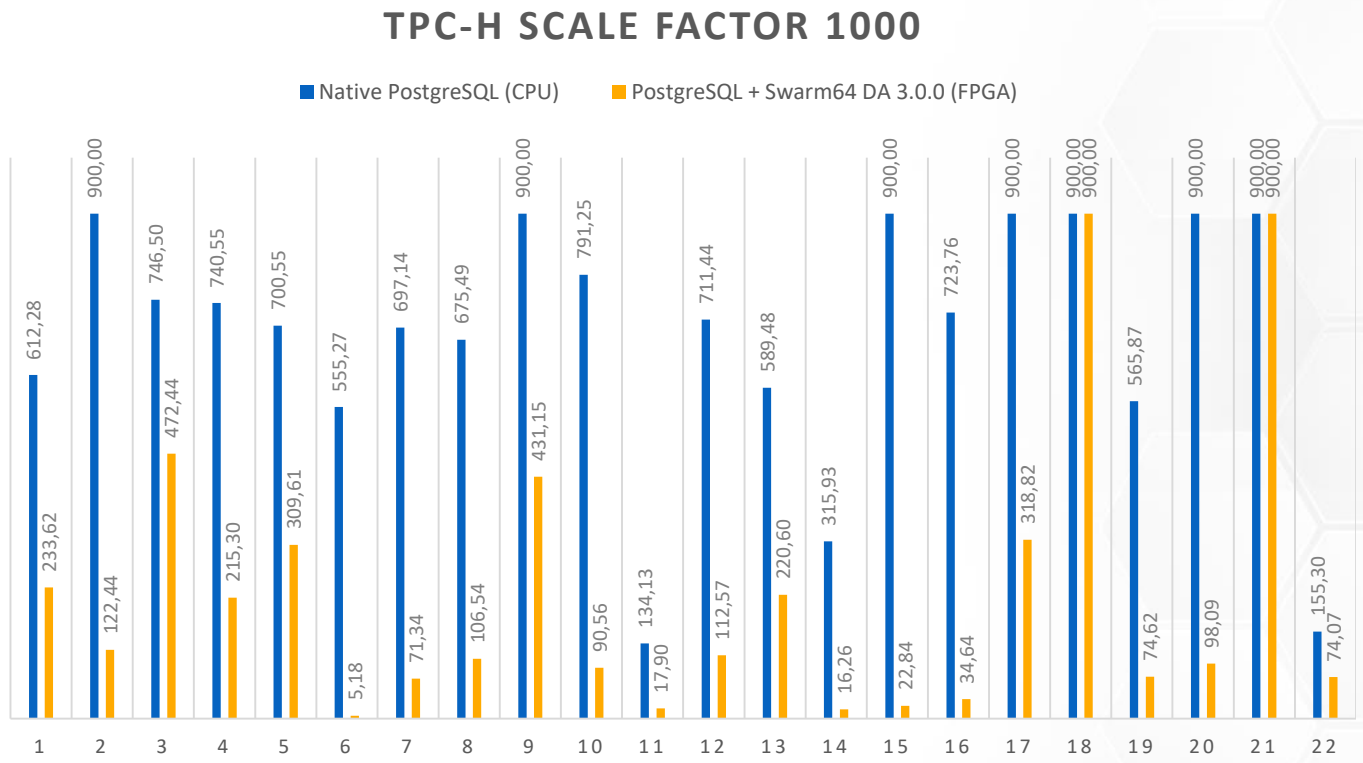

Average speed-up 4.3x

TPC-H Q6 PG Native vs. Swarm64 DA

You will see PG Native and Swarm64 DA side by side running on independent containers on the same machine with equally distributed vCPUs and RAM

Please wait, the video is loading...

Scale-In Before you Scale-Out

A story about HTAP on Single Node PostgreSQL

Imagine, You are the DBA of Trading Company X

Company X stores and analyzes market data to make trading suggestions

Some key parameters
Minimum of 100k transactions/sec ingestion
Data must be consistent, with collision detection
Average of 5000 symbols to analyze
Analysis on 10s level is acceptable, 1s is ideal

Problem statement

Does it scale?
More data to come!

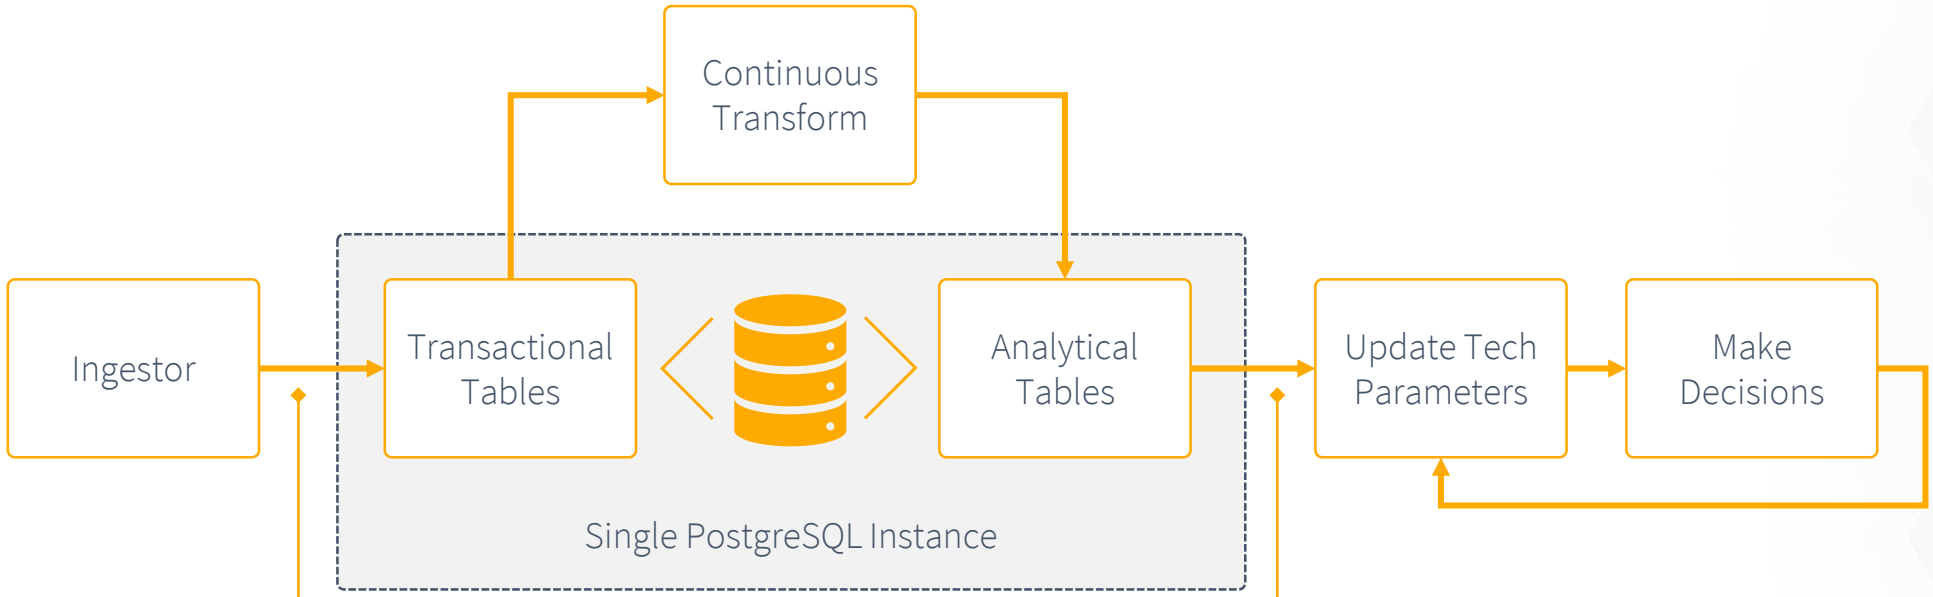

How to increase ingest TPS (>100k)?
More means higher market coverage

How to reduce analytics runtime (<10s)?
Less runtime leads to faster decisions

There is a dependency here...

"Classic" Scale-out Solution

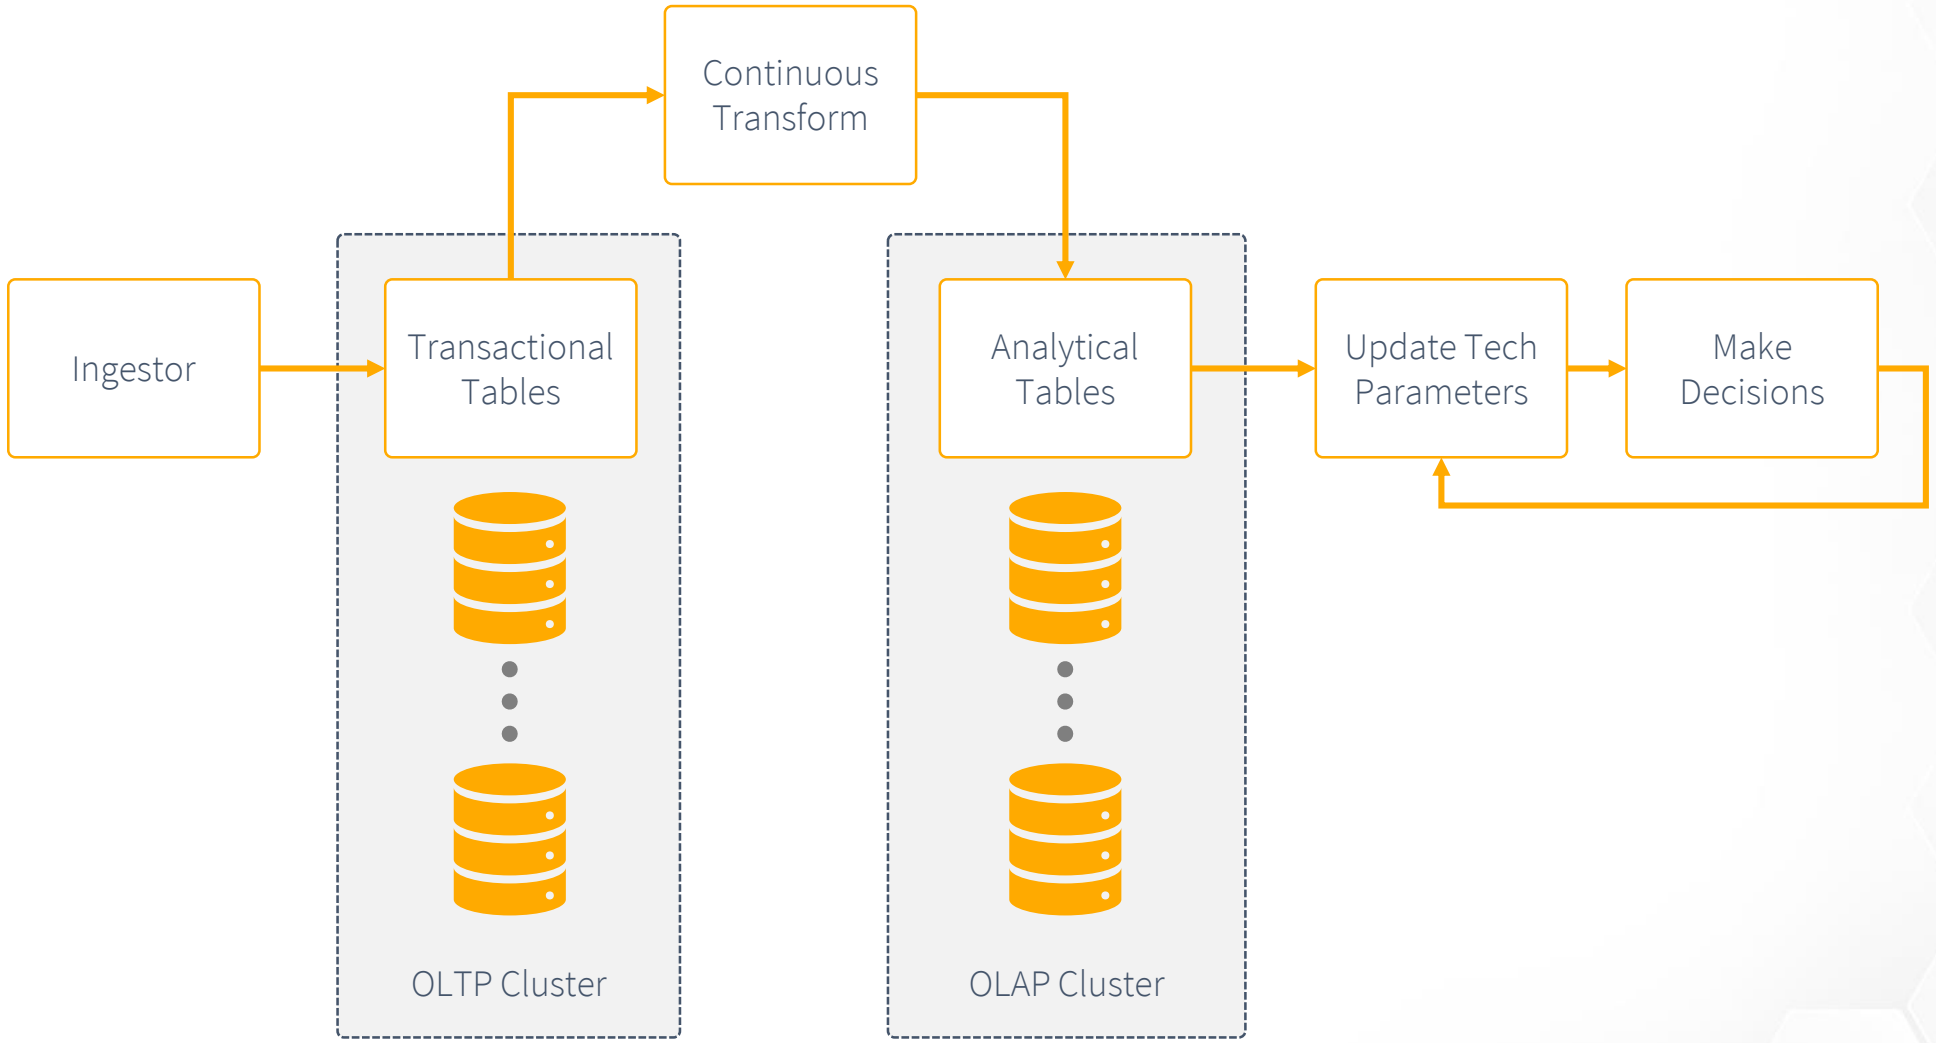

Reference Architecture HTAP

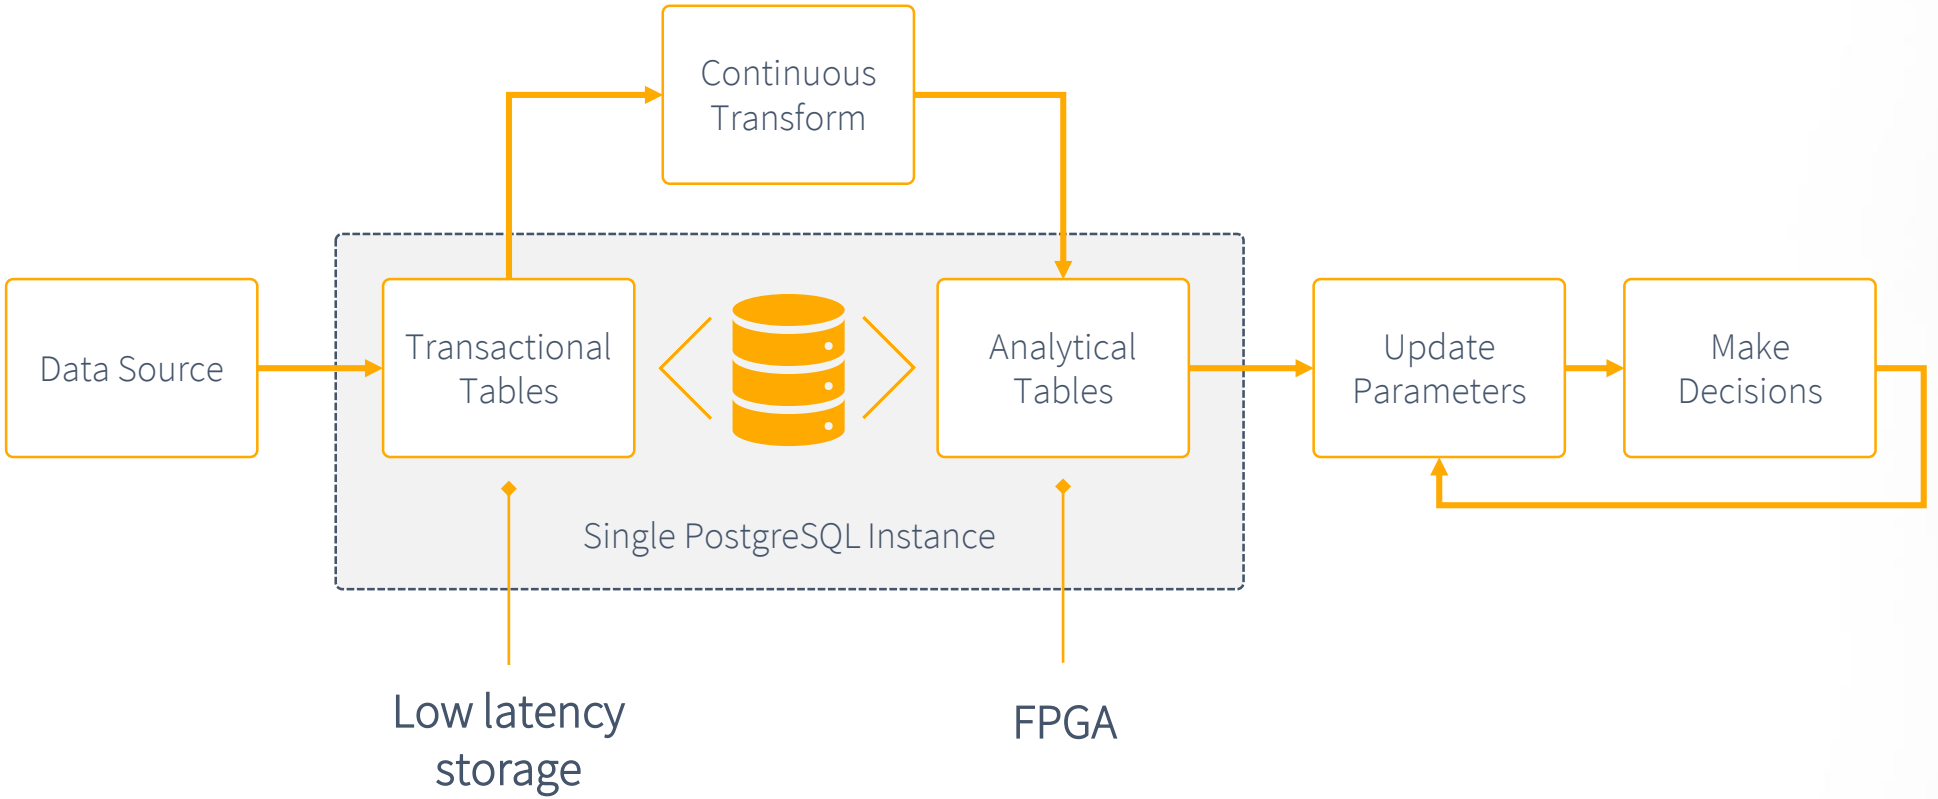

+

Swarm64 DA HTAP Use Case Demo

You will see Swarm64 DA and PostgreSQL in action together bringing acceleration to a mix workload

Please wait, the video is loading...

Use Cases

Data warehouse modernization

Replace costly, outdated analytic databases with free open-source PostgreSQL accelerated by Swarm64

Netezza replacement

The original FPGA-accelerated data warehouse appliance, now discontinued by IBM. Move your Netezza workloads to Swarm64.

Time series data

High-velocity data insertion , combined with fast concurrent queries of time series data. A Swarm64 sweet spot!

smava

Hewlett Packard Enterprise

TOYOTA

About Swarm64

We are hiring!

Founded in 2013

Large portfolio of granted and pending patents

Locations in Berlin, Boston, Palo Alto

Serving the Enterprise Database Market with Accelerated Postgres

swarm64

info@swarm64.com

Follow us: