



Industry 1ST Proven Accelerator for Data Analytics

February 2024





NeuroBlade



Location

Tel Aviv - Israel
Palo Alto - US



Foundation Date

2018



Employees

120+



Capital Raised

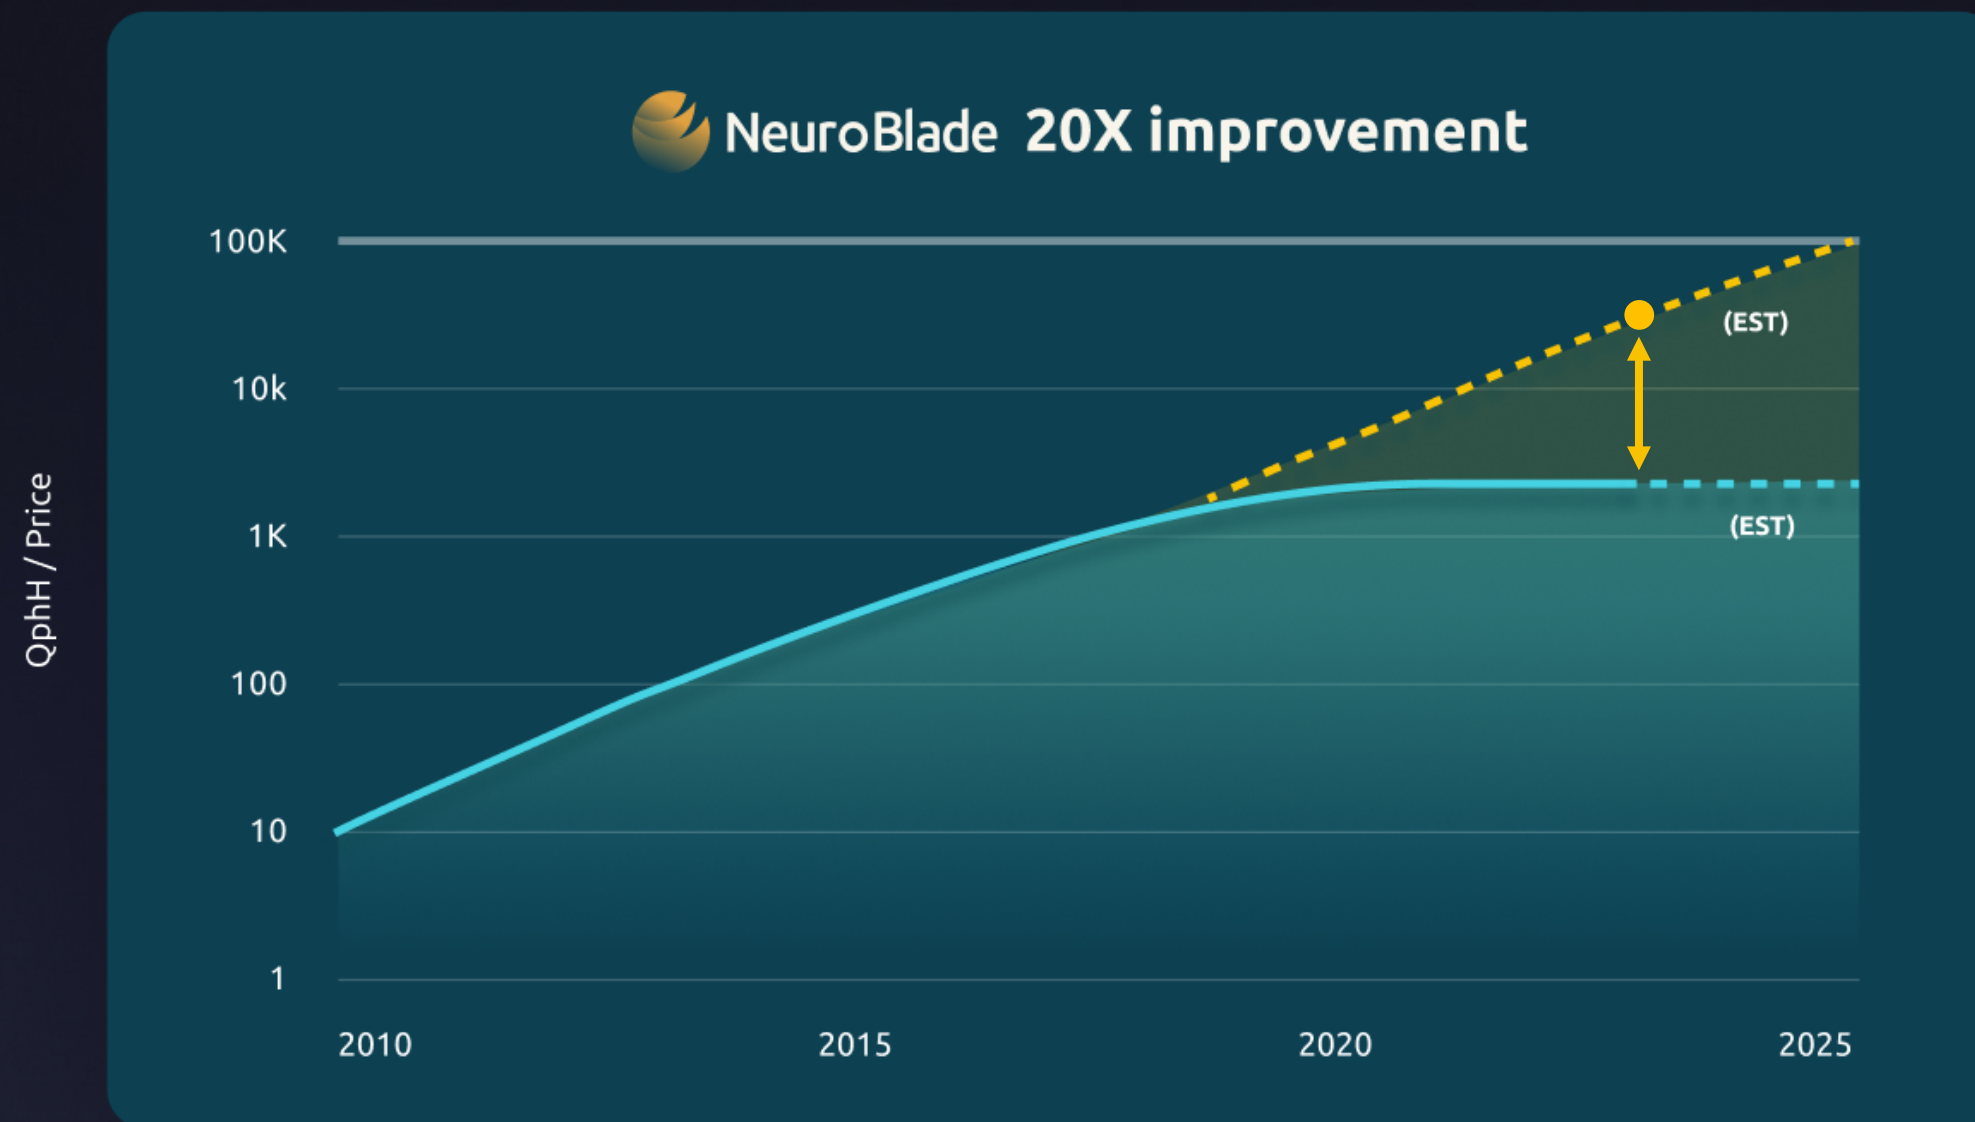
\$100+M



Early Production

2023

Acceleration is MUST for Data Analytics



CPU query processing power doesn't scale

Query Per Hour


x10

Query Per Hour / \$

x20

- ◆ **Standard Form Factor**
- ◆ **Fits Existing Data Platform**


SQL Processing Unit (SPU)



- ◆ Processes at the speed of storage & network
- ◆ Plugin to Your ODM/OEM Server
- ◆ Programmable and extensible

SQL Processing Unit (SPU)

Announced on May 2023



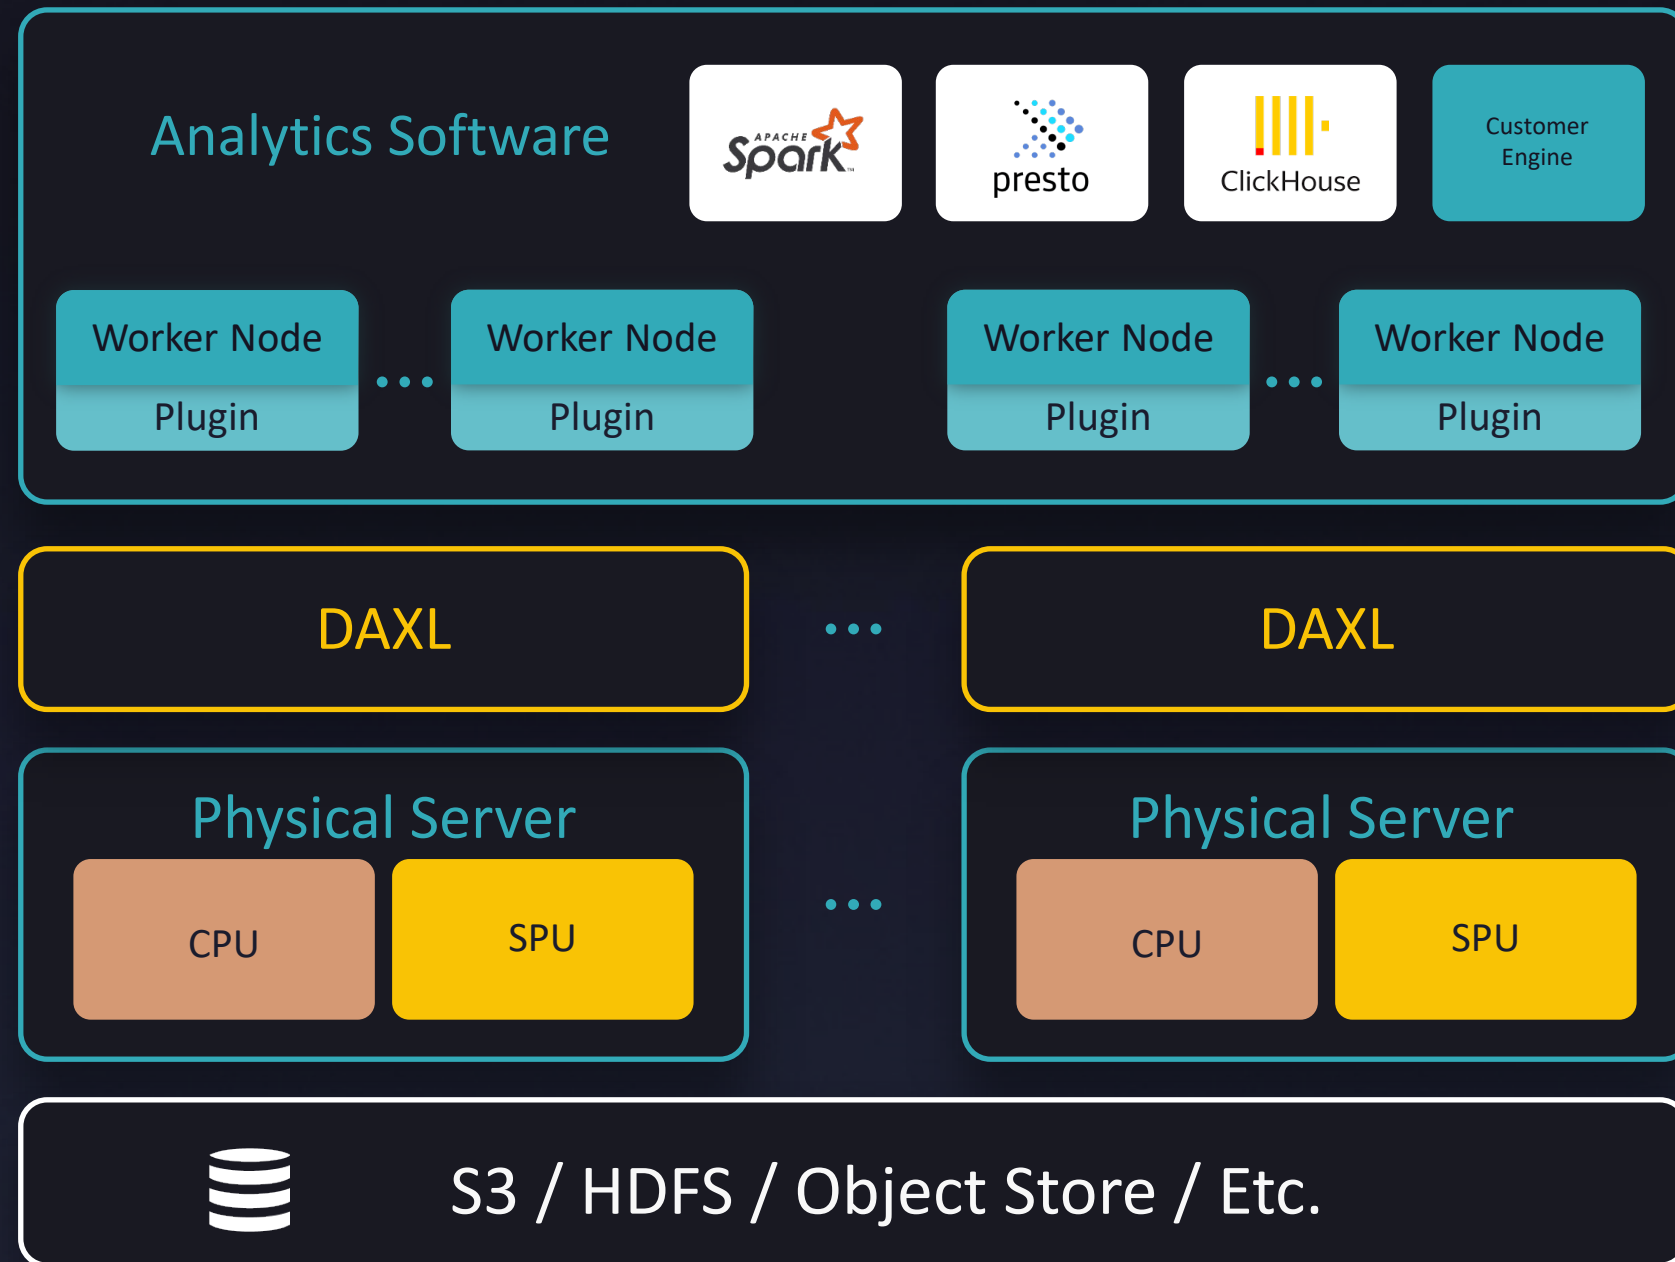
- ◆ Best-In-Class *Quality* Vendors
- ◆ *Co-exists* with your ecosystem

SPU Hardware Integration



NeuroBlade SPU G200	
Size	FHHL, single width
SPU Processor	x1
PCIe edge connector	PCIe 4.0 x16 lanes
TDP	150[W]
External power (optional)	8pin PCIe ATX (R/A 2x4)
Thermal	Passive Cooling
Environment	Storage: -40°C to 75°C Operating: 0°C to 55°C Humidity: 5% to 95%

SPU snaps into existing infrastructure

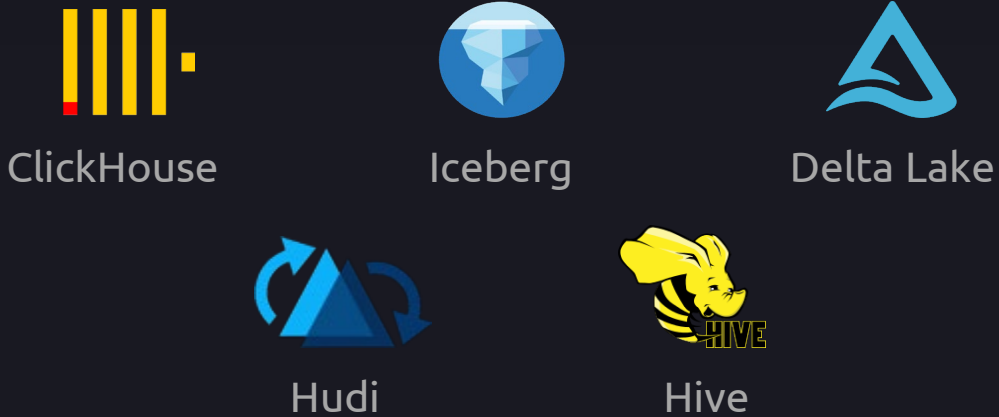


- ✦ Simple integration using existing connector to analytic engines
- ✦ DAXL enables query's to be executed on SPU
- ✦ CPU's role offloads I/O intensive compute to SPU

Support most common table & file formats

The SPU is a native component in the customer composable infrastructure


Table formats



ClickHouse, Iceberg, Delta Lake, Hudi, Hive

Most Table Formats supported out of the box


File formats



Parquet, ClickHouse, ORC

Proprietary columnar file format can be added

Memory Layouts



Velox, Block, Arrow

Results Sent back to CPU with Zero copy

NeuroBlade and Meta's Velox teams Accelerate Data Analytics

- ✓ >30x Performance Advancements
- ✓ Velox Open-Source Connector
- ✓ Seamless Integration

*We have created Velox to **unify software-based optimizations** into a single execution engine " that can be reused across engines such as **Presto, Apache Spark, and beyond.***

*Our long-term strategy for Velox, however, is to provide a unified framework that allows hardware accelerators to be **seamlessly integrated** into data management"*

commented Pedro Pedreira, Software Engineer and Velox lead at Meta.

[Press Release](#)

NeuroBlade and Meta's Velox teams Accelerate Data Analytics

- ✓ >30x Performance Advancements
- ✓ Velox Open-Source Connector
- ✓ Seamless Integration



Recap: Product & Eco System



+

NeuroBlade DAXL SDK

✓ Available Form Factors:

PCIe
FHHL

✓ Supported Servers:

DELL

ODM
Design

Your
OEM

✓ Reference code Available:

presto

✓ Roadmap:

APACHE
Spark

ClickHouse

Your
BigData
Platform

◆ Better Performance

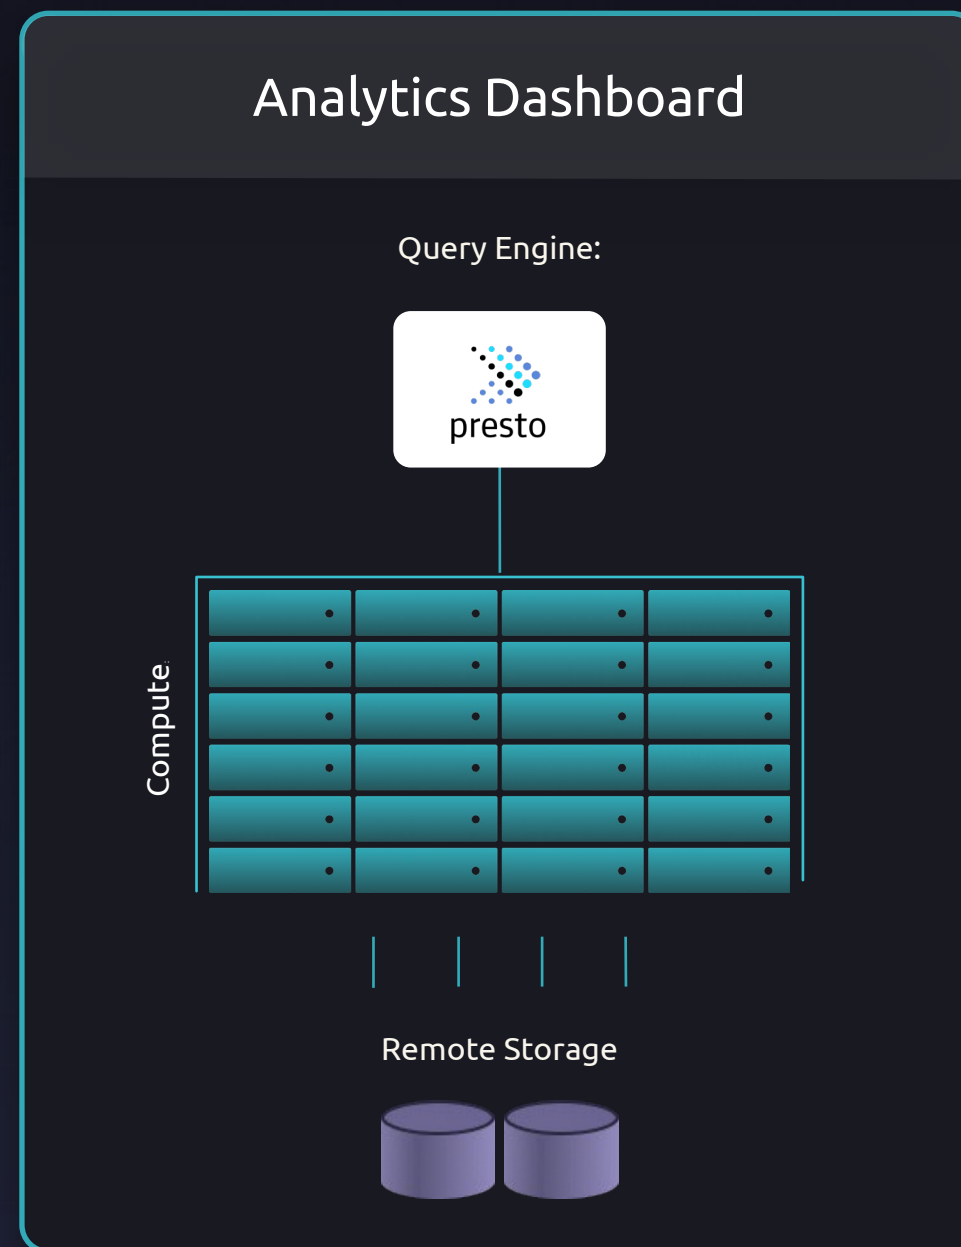
◆ Less Hardware

◆ Fits Your Data Platform

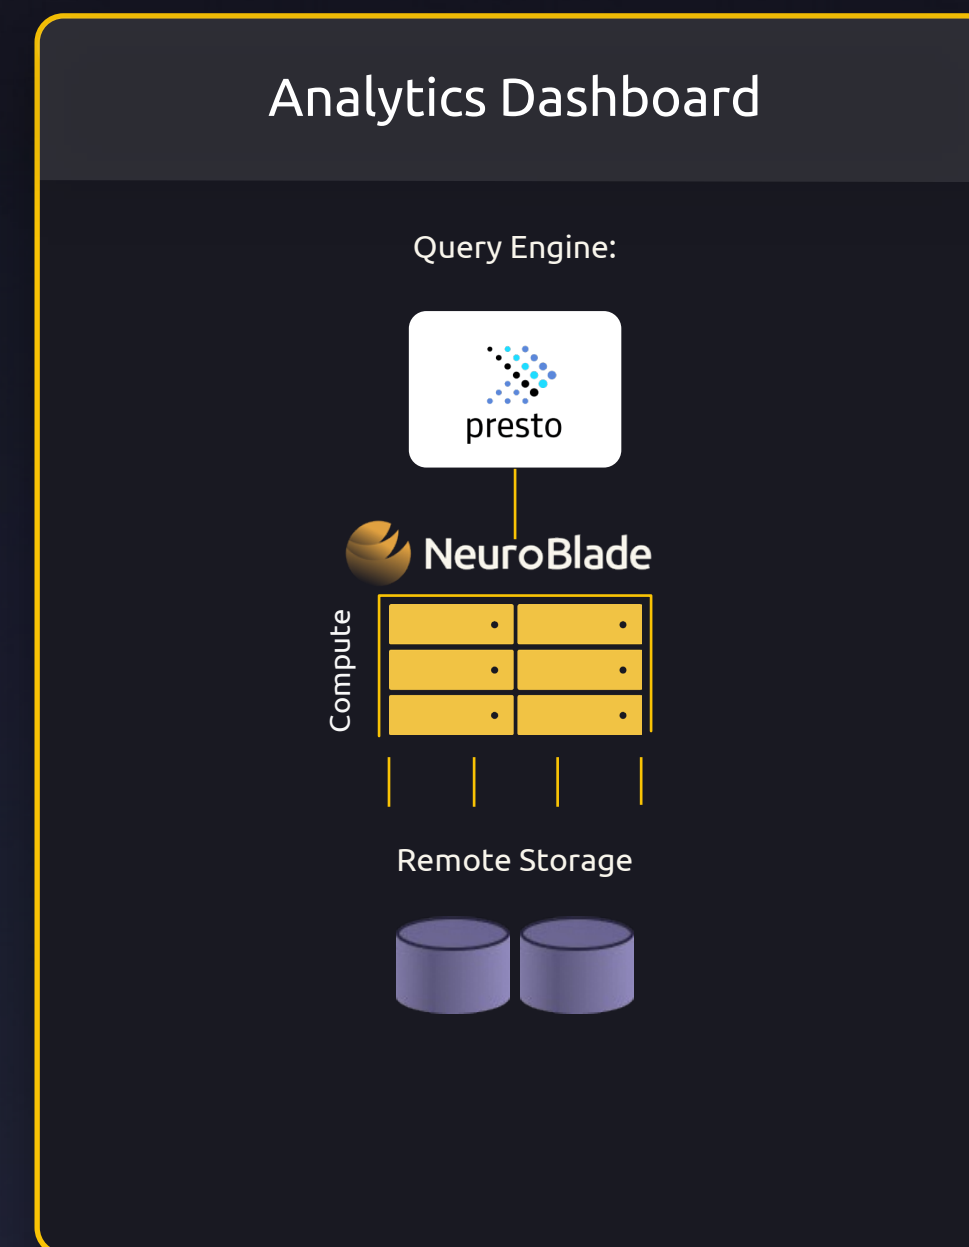
Case study at top Internet company

Customer was looking to address server sprawl due to analytics

Before NeuroBlade



After NeuroBlade

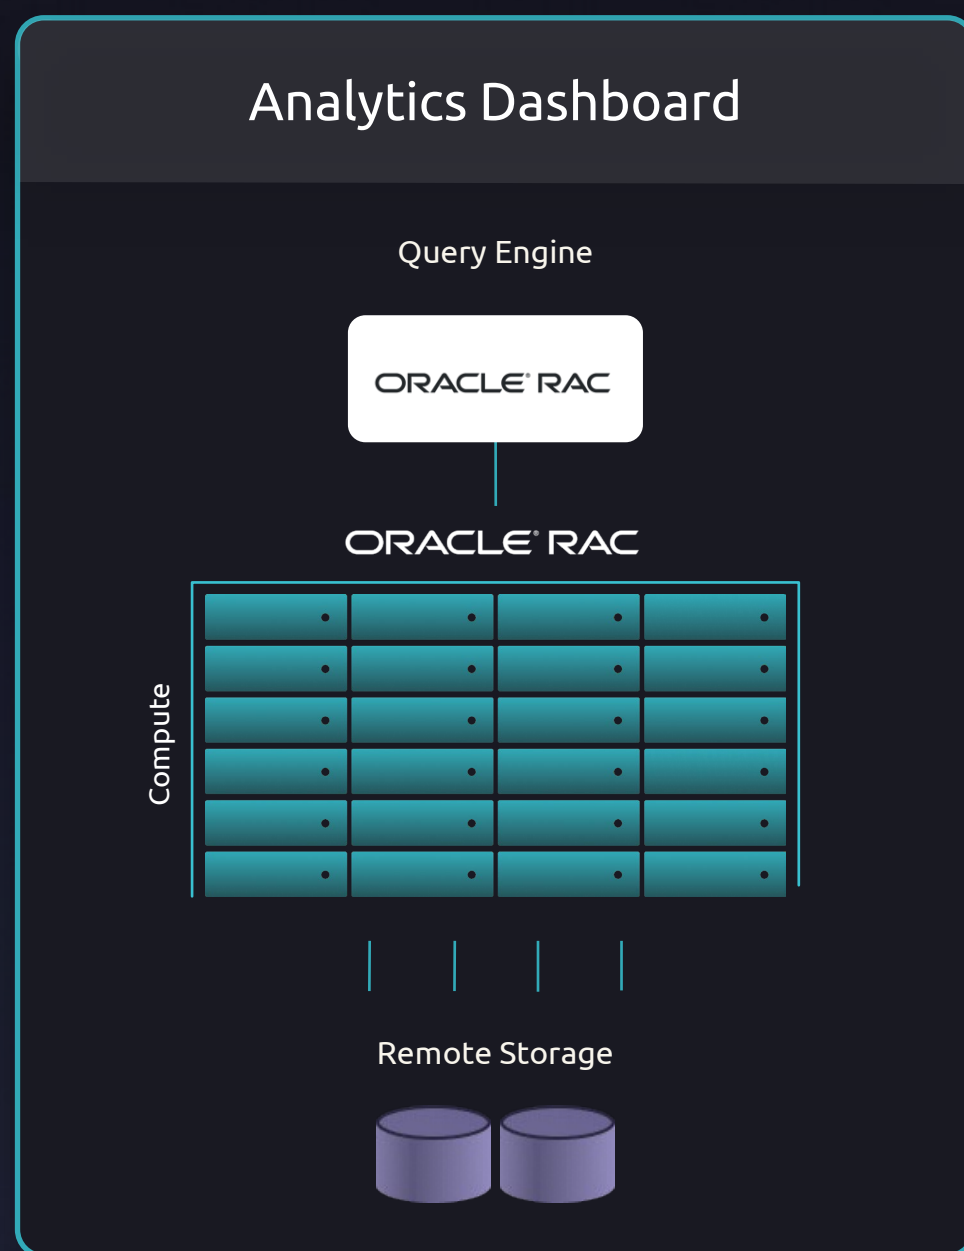


Faster performance **16x**
Reduced compute footprint
Quicker results and decisions

Case study at large health system

Customer needed to improve costs while data volumes grow 40% yearly

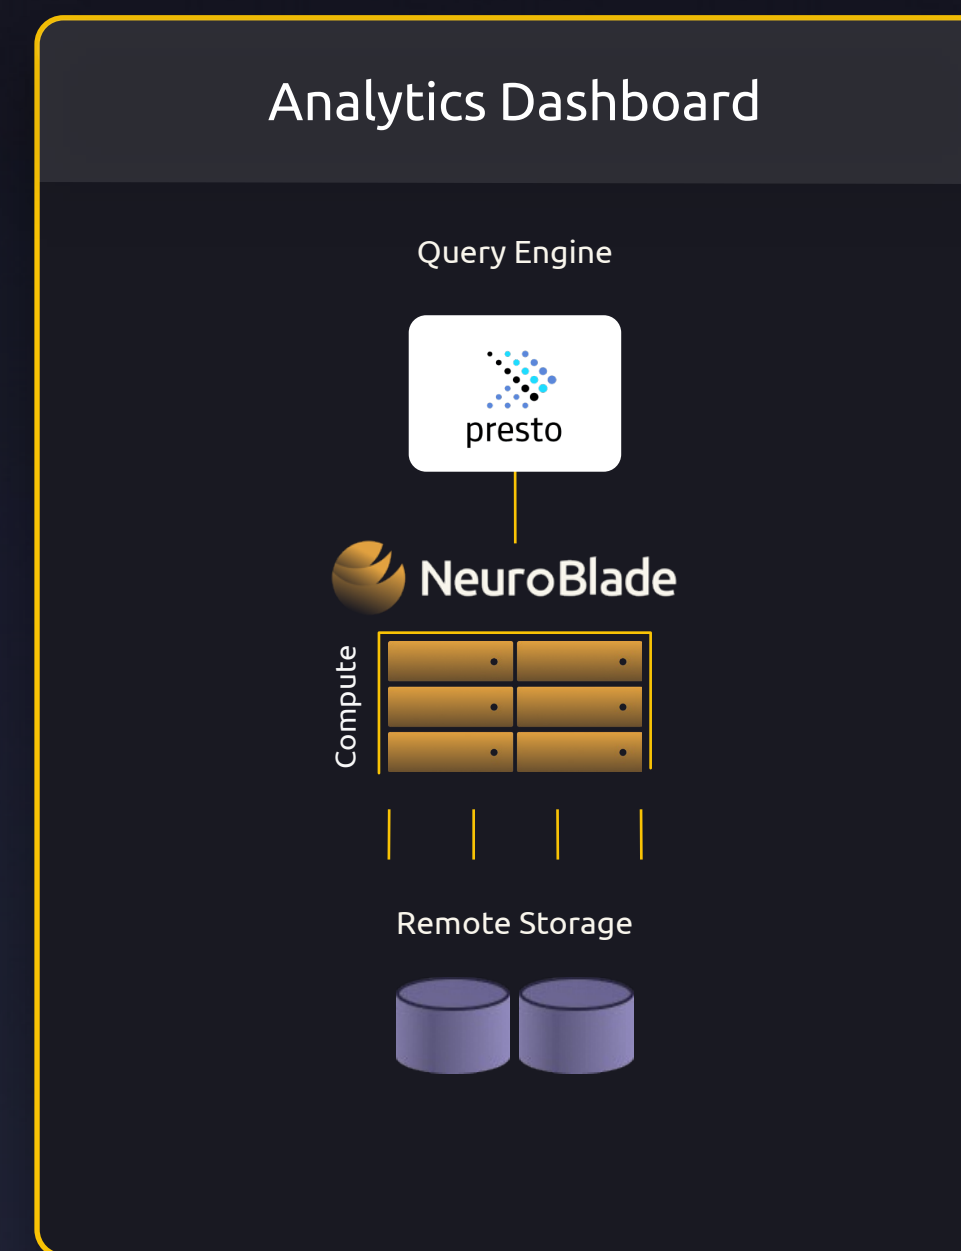
Before NeuroBlade



70% Cost reduction

30x Faster queries

After NeuroBlade





Questions?