



Hewlett Packard
Enterprise

Persistent Memory The Next Turbo for Databases

Eduardo Mattos Duarte
3PAR Product Manager for Next Generation Platforms



Why legacy architectures are so different from AFA

Legacy Architectures

Cache Centric

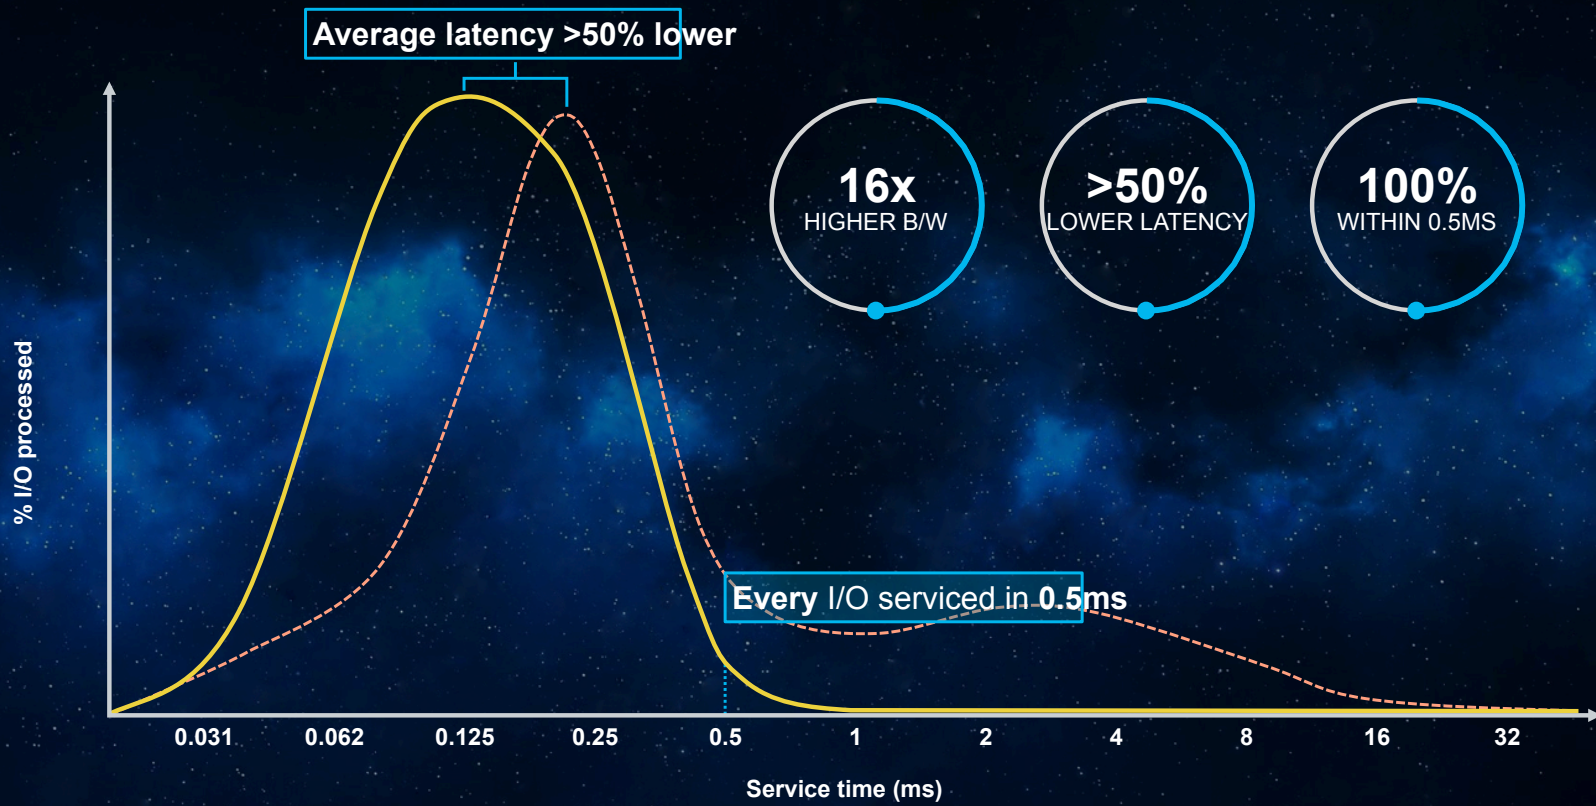


All Flash Architectures

Back-end Centric



Delivering low latency with Turbo Accelerator



Dramatic change the database behavior




**1/2 of the average latency
for OLTP Workloads**



**Increase IO density
By 4 times**

Hardware is just as important as ever

Persistent Memory



NVMe



ASIC



Leveraging AI for case automation

Clustering algorithms



Statistical models




Pattern recognition



Summary

In-Memory Compute



AI/Machine Learning



Scalable Architecture

