kx | it's about time

Optimization strategies for in-memory analytics using Optane persistent memory

Kx/Intel Early access program

STAC Summits, Fall 2019

Rebecca Kelly
Kx Technical Evangelist, Kx Systems

intel

Why is Optane Memory so useful for time-series data?

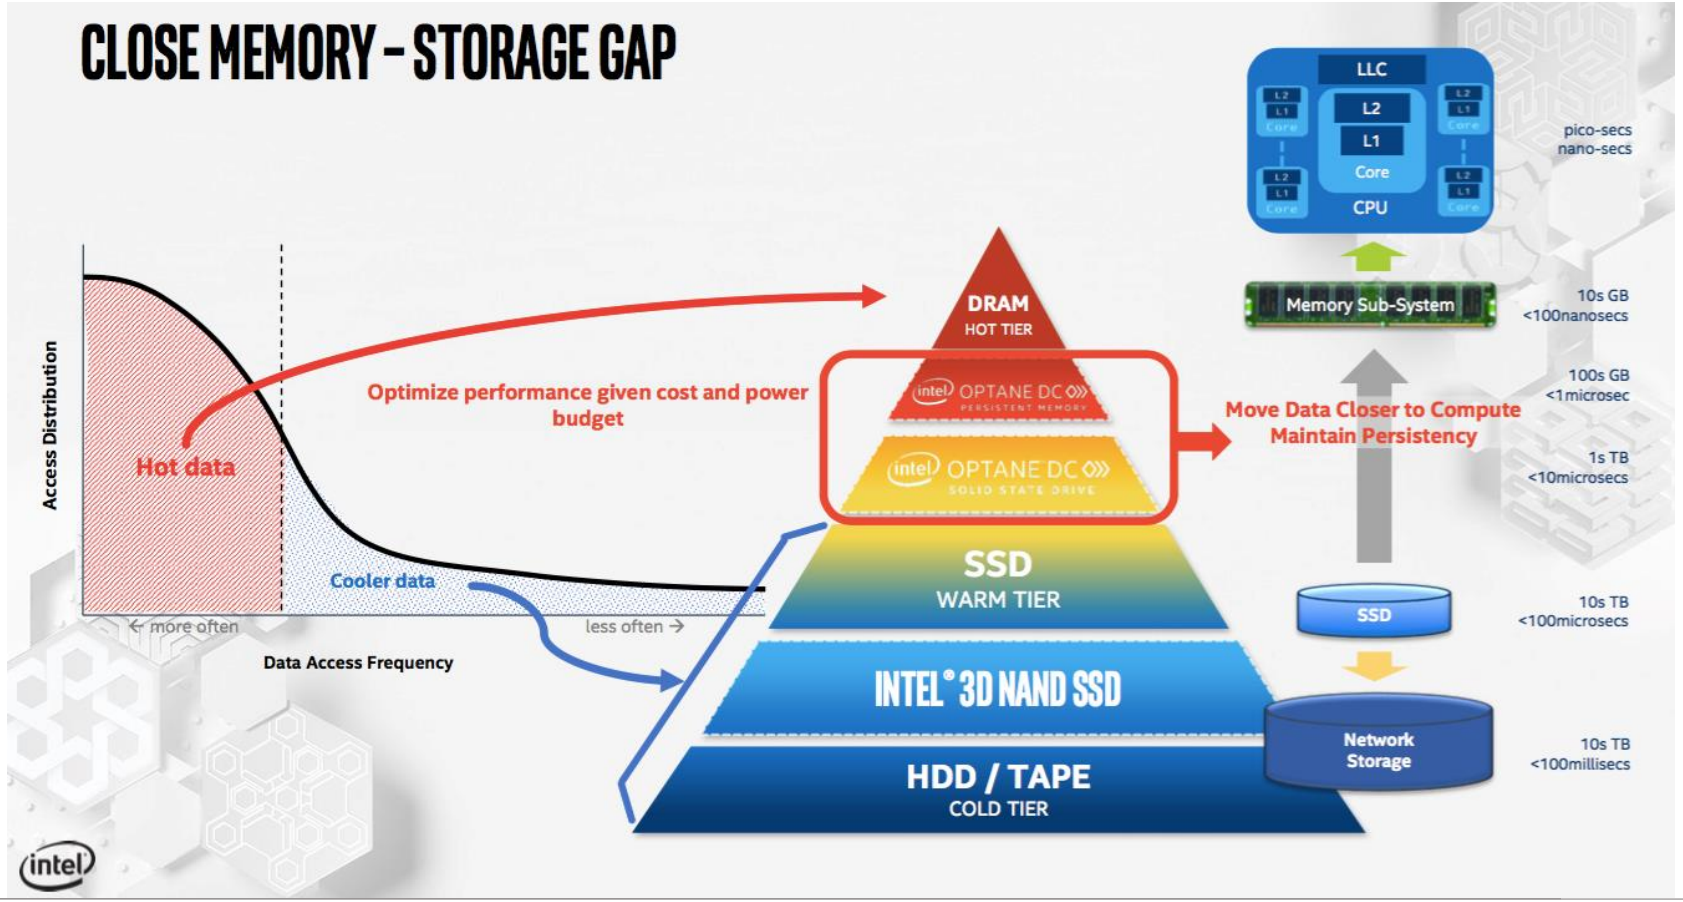

SNIA NVM programming model supports realtime data analytics

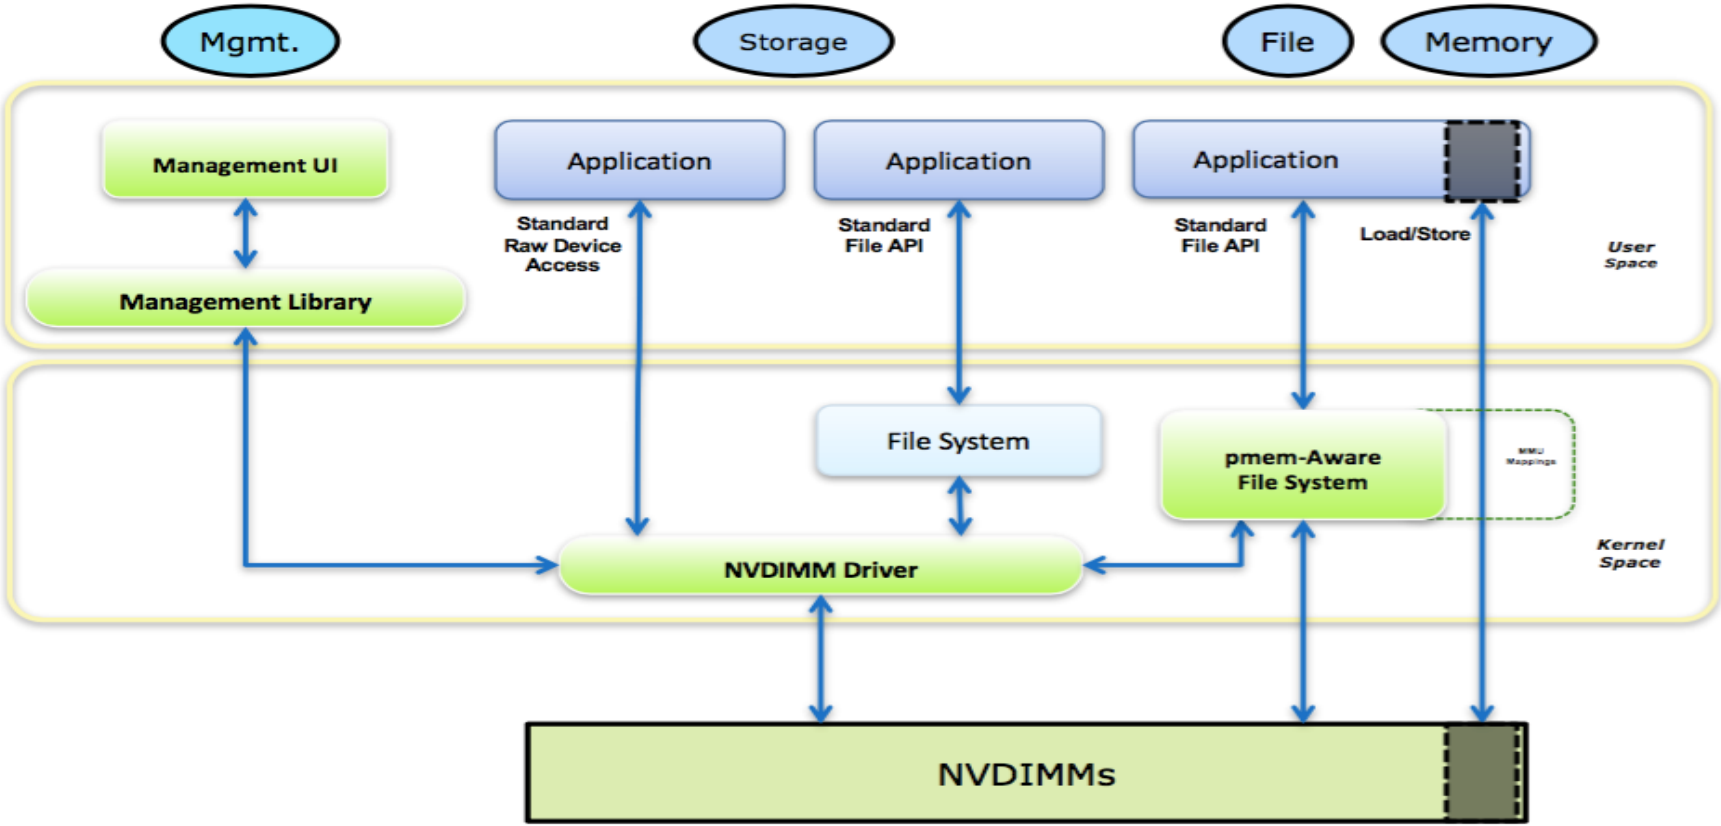

1. Memory Mode

- PMEM used as Large Volatile Memory
- Looks to SW stack as large memory
- up to 6TB in a 2S server at 90% the performance of DRAM

- Kdb+ Create object ~ 8% to 19%
- Read object ~ 8% to 14%

Use case: RDB, massive data sets, no app change

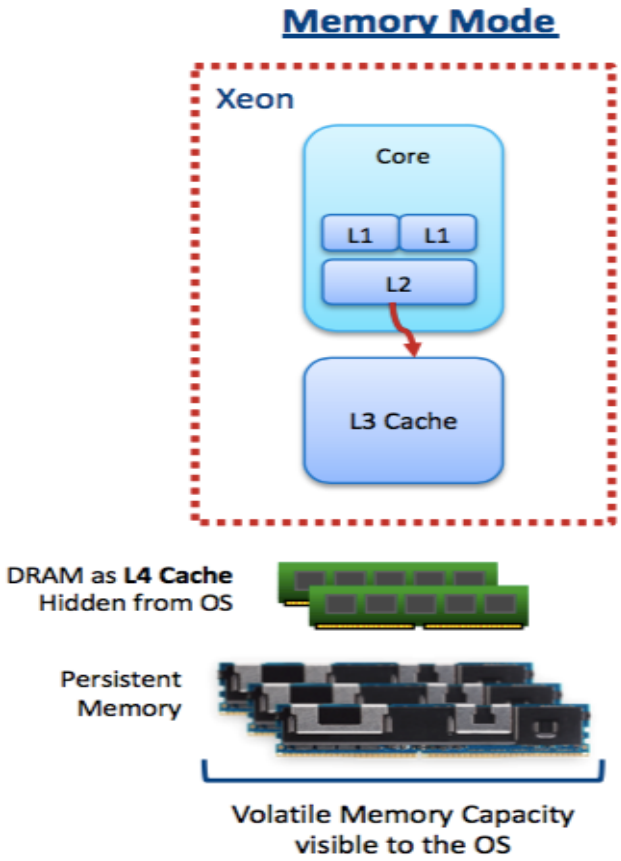

2. Persistent Storage mode (STAC-M3 Antuco)

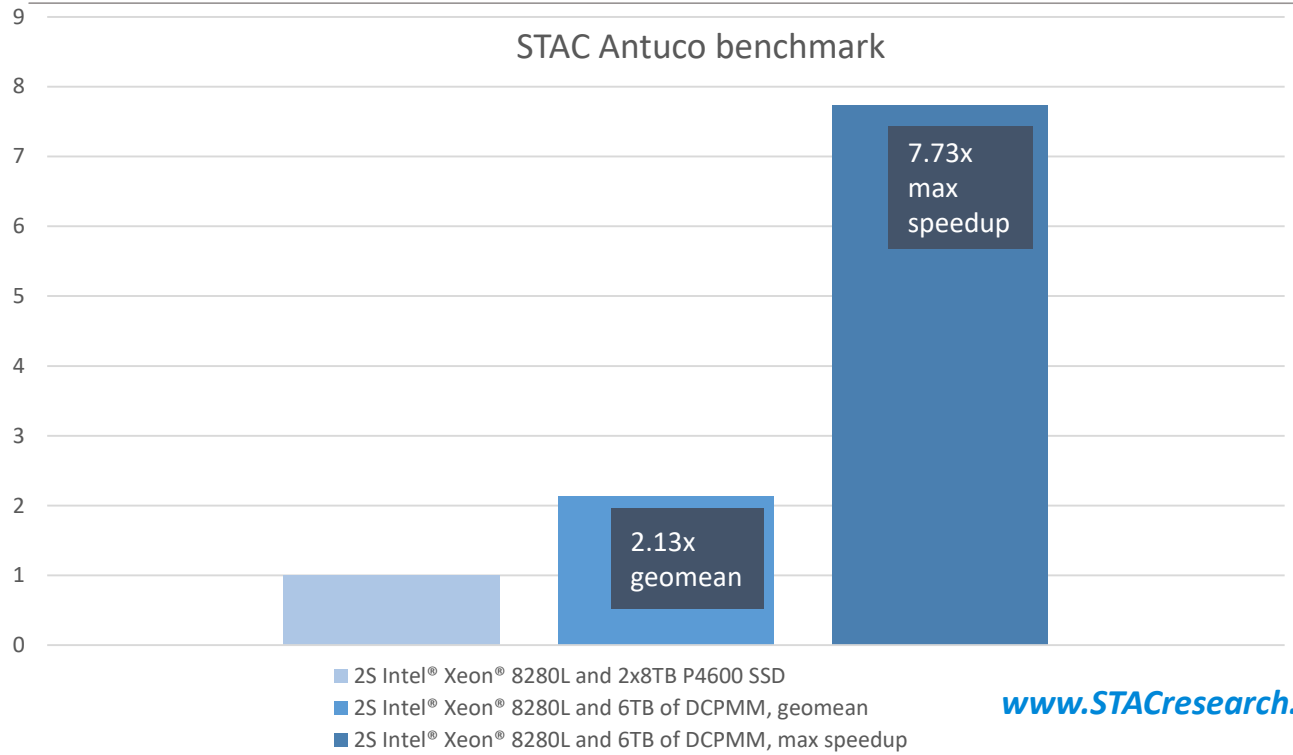

Finance users are encouraged to also look up the following unaudited reports from Intel:

SUT ID KDB190416

www.STACresearch.com/KDB190323

An excellent candidate for persistent "very hot" HDB data, viz. multi day or multi week

3. DAX “App-direct” mode

- Traditional read to page fault (disk):
  1. Software
  2. 4K transfer from disk
  3. Request returned

- App Direct access memory directly
  - Avoids software and 4K transfer overhead
  - Cores can still access DRAM normally, even on same channel

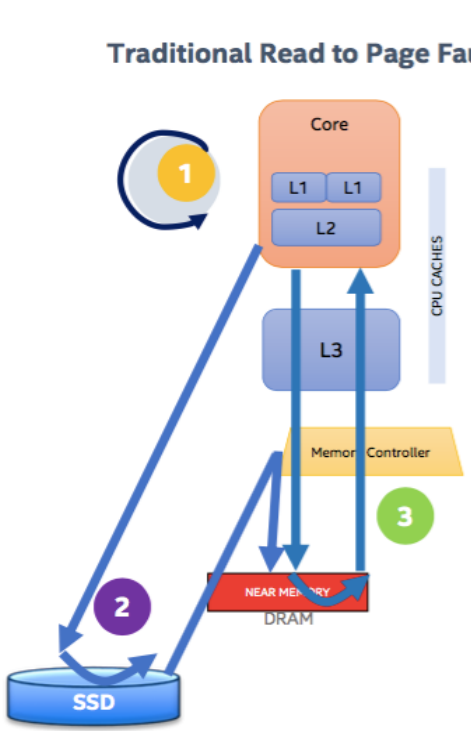

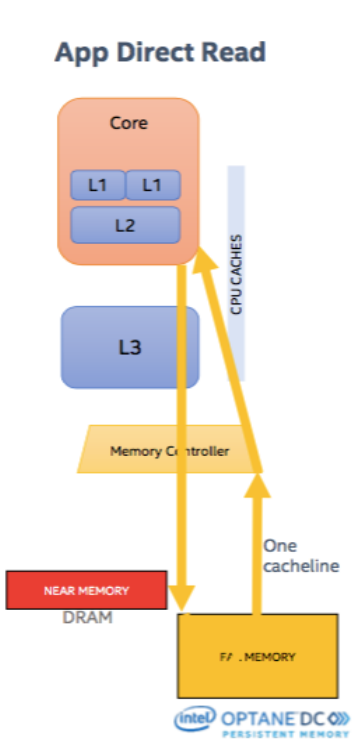

Early ACCESS for finance clients:

- ❑ ***Horizontal RDB partitioning : Keep ‘recent’ historical data in Optane Memory, allowing multi-day queries in memory***
- ❑ ***Vertical RDB partitioning : Different tables/columns residing in DRAM/Optane Memory***

Next Steps?

- Intel/Kx Early Access programme now available for Finance customers evaluating new technology
- Consulting support from Kx is co-packaged
- Intel and Kx deliver full hardware and software stack for early access evaluations
- See us at the Intel booth today, or email: optane@kx.com to join the EA program and get started!