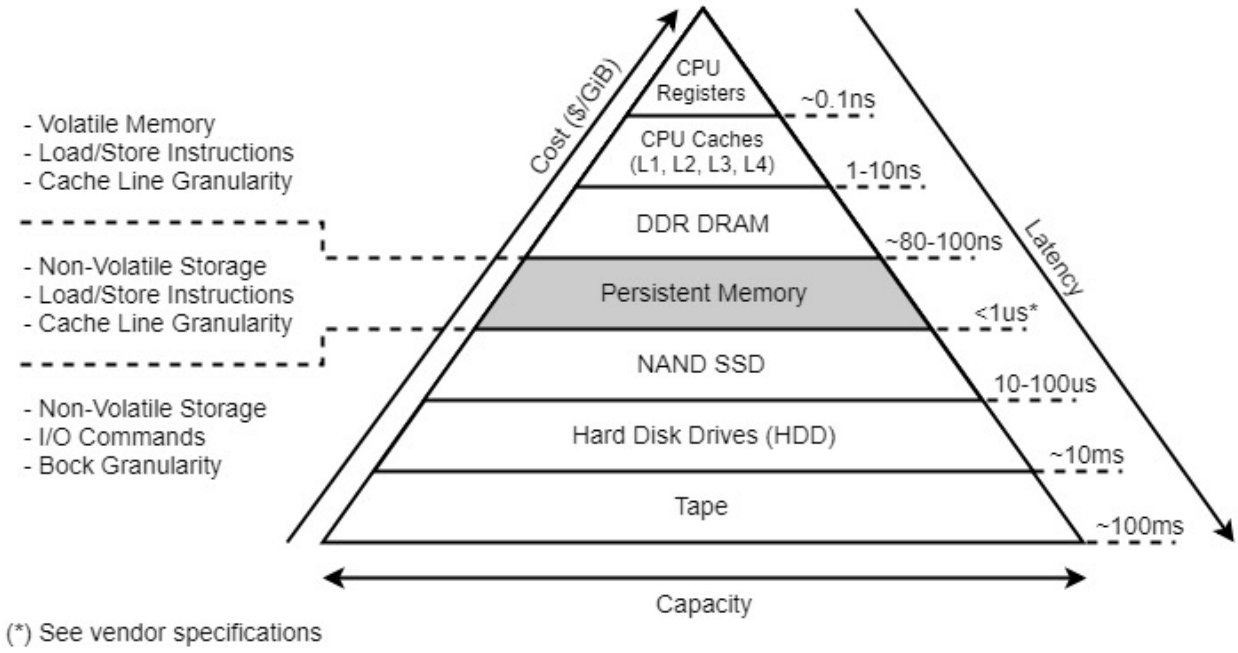

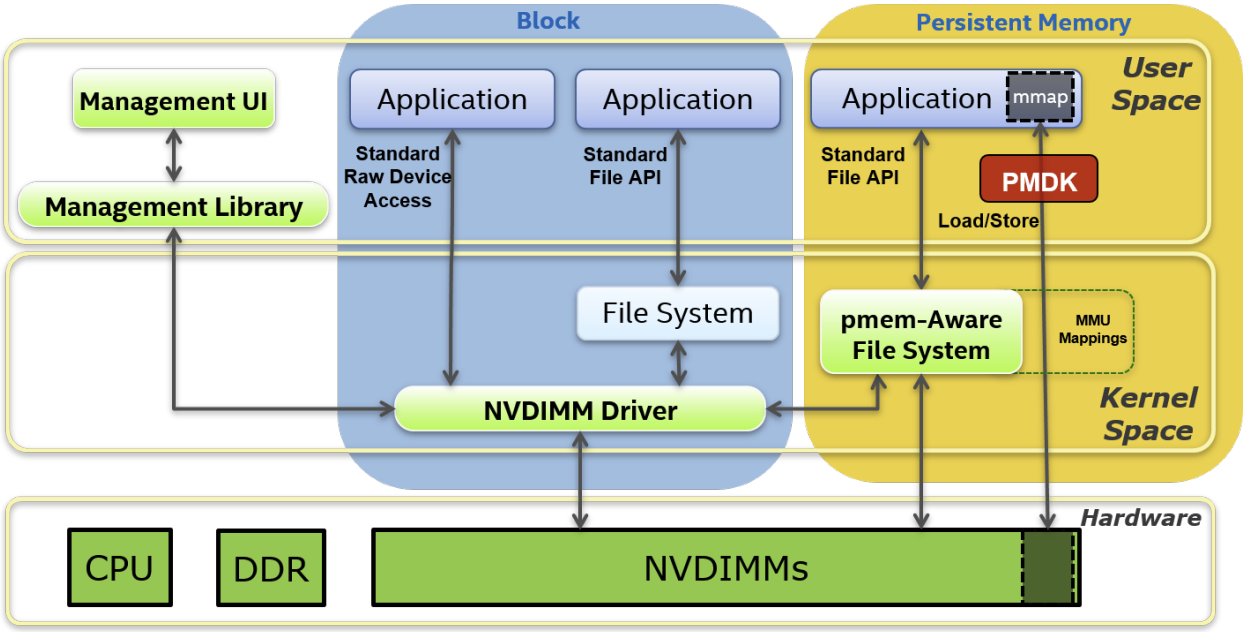

Figures borrowed from “PMDK Introduction”
https://docs.pmem.io/persistent-memory/getting-started-guide/what-is-pmdk

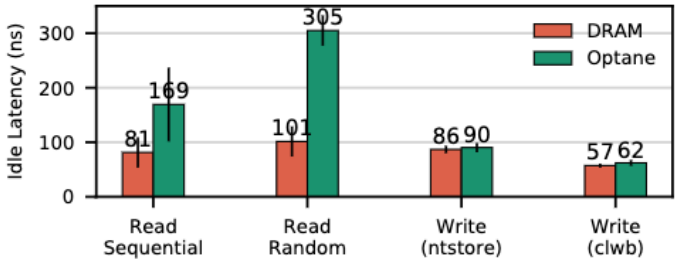

Figure 2: **Best-case latency** An experiment showing random and sequential read latency, as well as write latency using cached write with `clwb` and `ntstore` instructions. Error bars show one standard deviation.

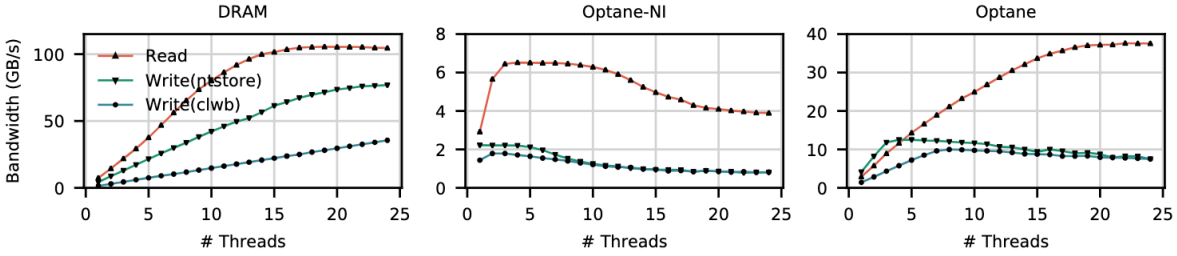

Figure 4: **Bandwidth vs. thread count** An experiment showing the maximal bandwidth as thread count increases (from left to right) on local DRAM, non-interleaved and interleaved Optane memory. All threads use a 256 B access size. (Note the difference in vertical scales).

Above figures are borrowed from An Empirical Guide to the Behavior and Use of Scalable Persistent Memory

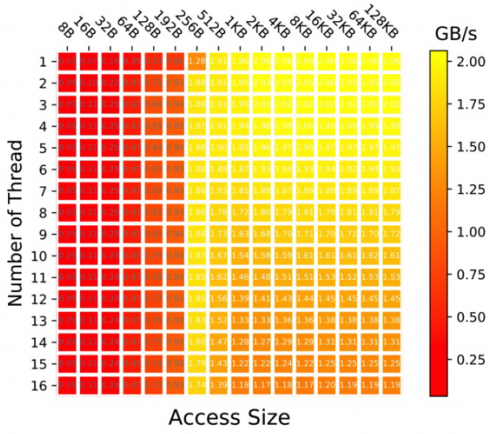

Figure 1. Random write performance on one Optane Pmem using different access size. For writes, we use *ntstore* followed with a *sfence* instruction to ensure data persistency. Performance degradation with a larger number of threads and larger access sizes is due to contention in the iMC (integrated Memory Controller).

The figure is borrowed from ChameleonDB: A key-value store for optane persistent memory.