

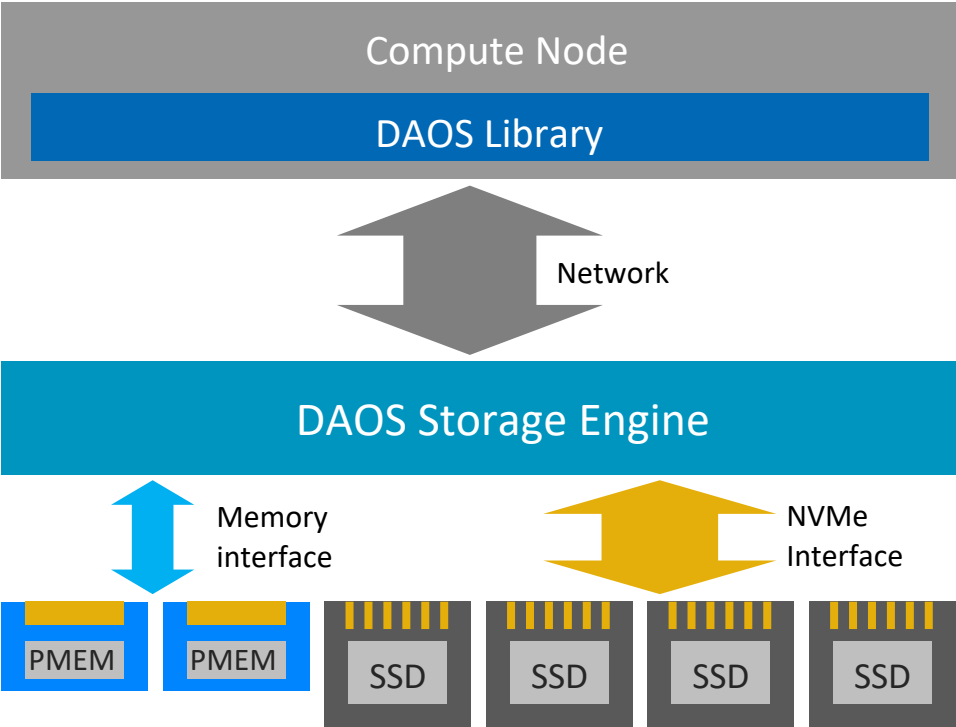
DAOS no-pmem mode performance

June 18, 2024

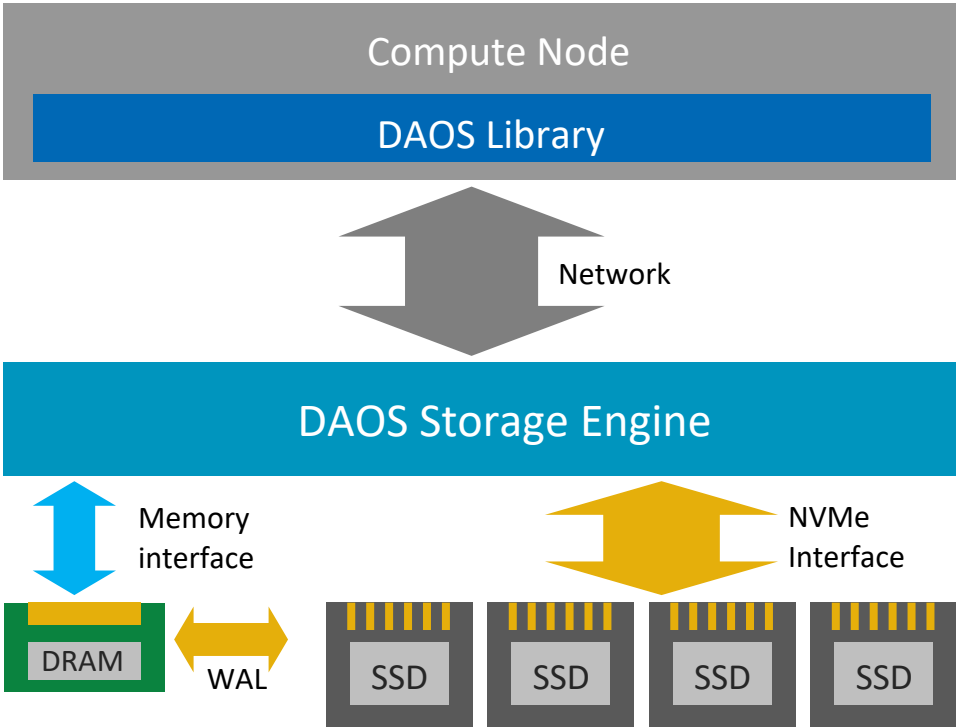


intel[®]

DAOS Architecture Evolution

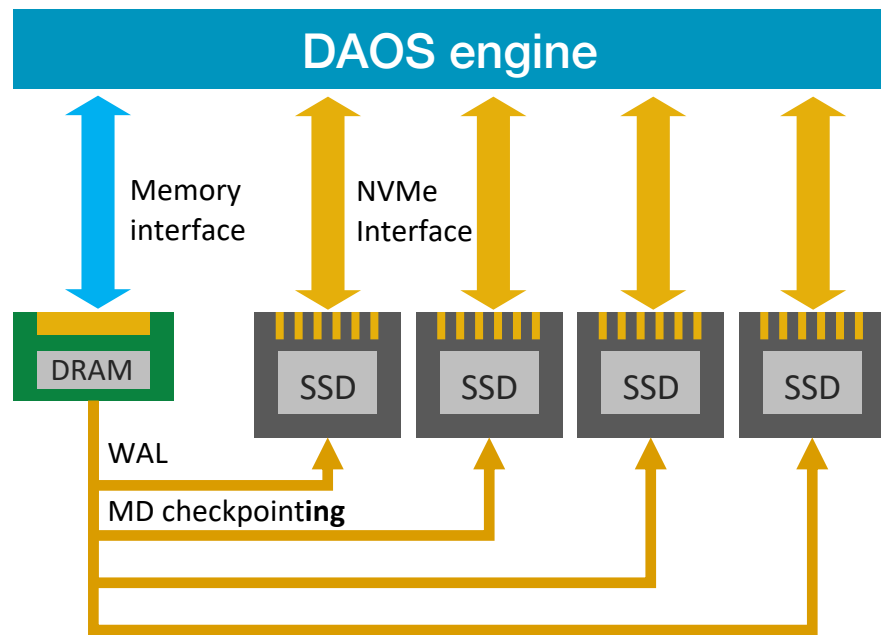


With Persistent Memory

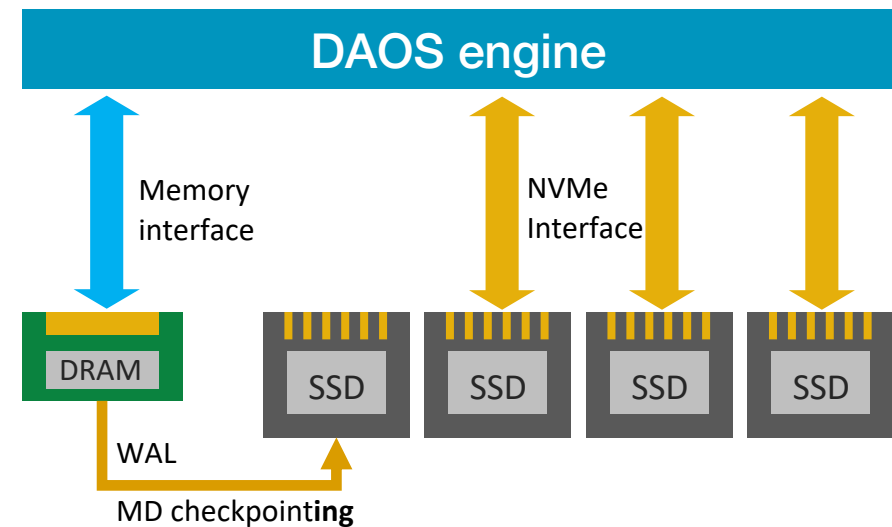


Without Persistent Memory

MD-on-SSD configuration

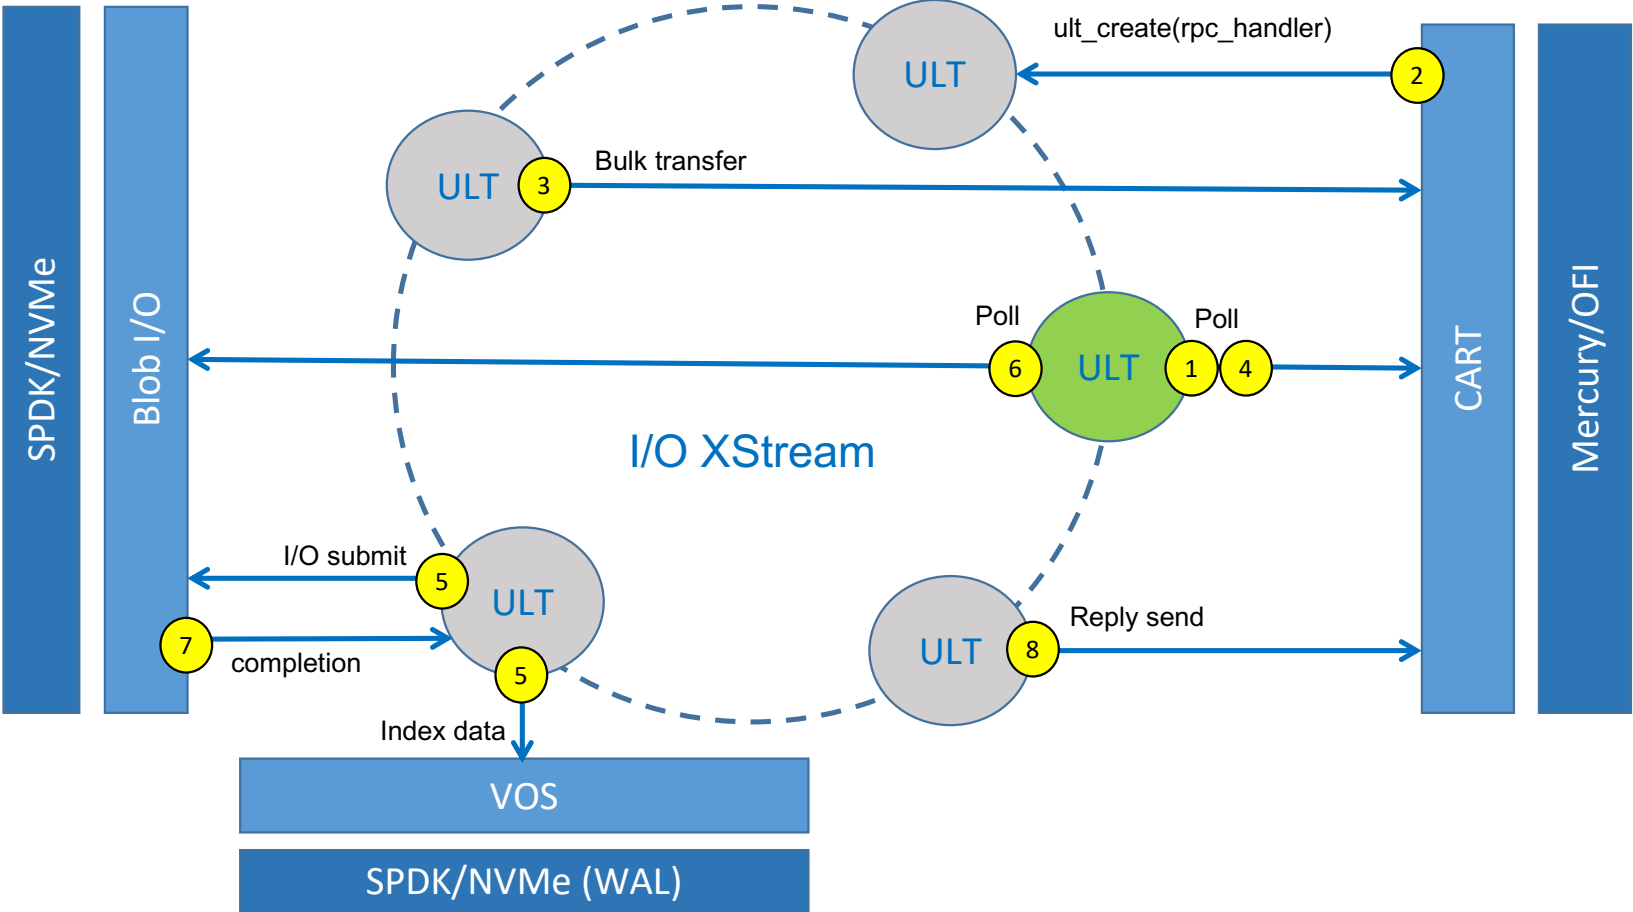


Mixed mode



Dedicated SSD for MD/WAL

MD-on-SSD write flow





Write Flow in md-on-ssd mode

- Server-side **synchronous** submit
 - Submit **data write**
 - Submit **WAL write** on completion of data write
 - Reply to client on completion of WAL write
- Server-side **asynchronous** submit
 - Submit **data and WAL write** at the same time
 - Reply to client on completion of both writes
 - Rely on checksum for integrity check


- WAL write
 - Most WAL records are small
 - 2K-3K, one atomic write
 - Very few of them can be large
 - More than 4K, up to a few megabytes
 - WAL(>4K) write is not atomic
 - Checksum verify WAL integrity

IOR EASY (4 servers + 16 clients)

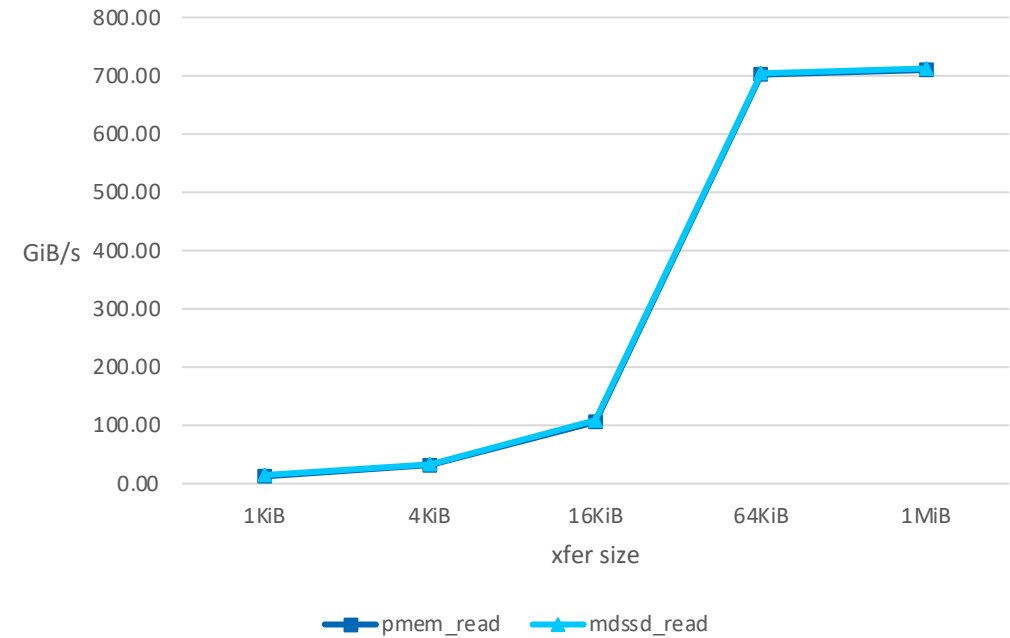


IOR EASY (16 servers + 64 clients)

16 Server, 64 Client, Write
pmem vs combo wal+data+meta

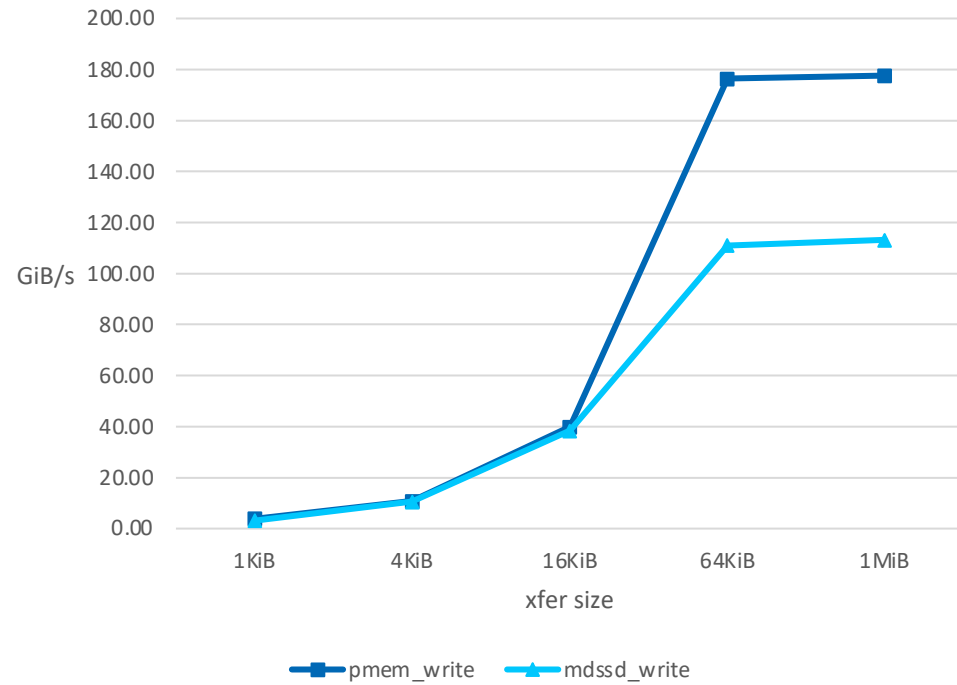


16 Server, 64 Client, Read
pmem vs combo wal+data+meta

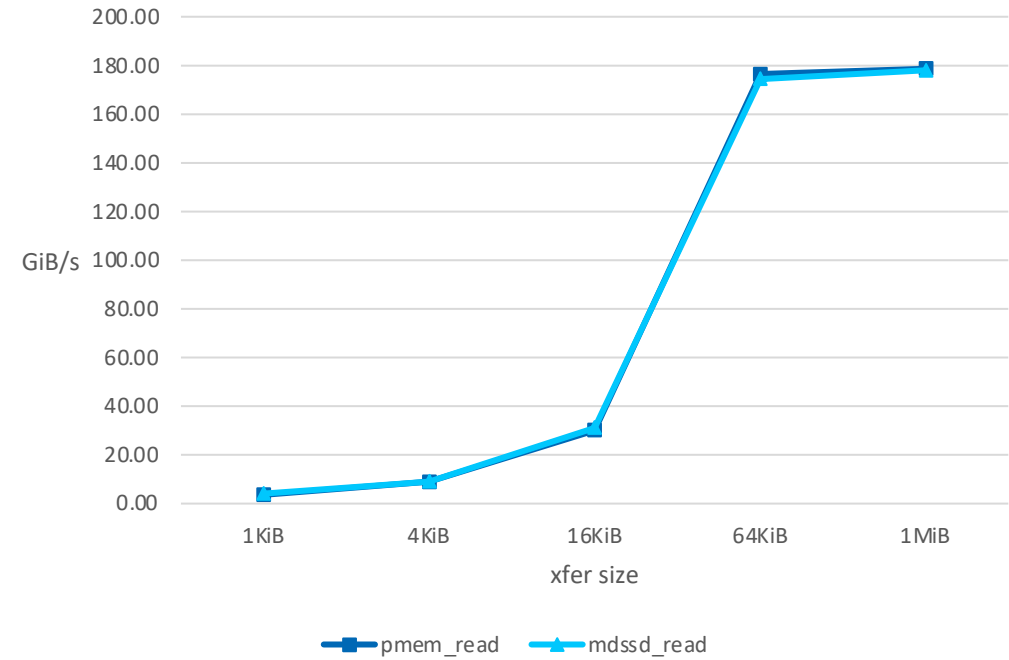


IOR EASY (dedicated WAL)

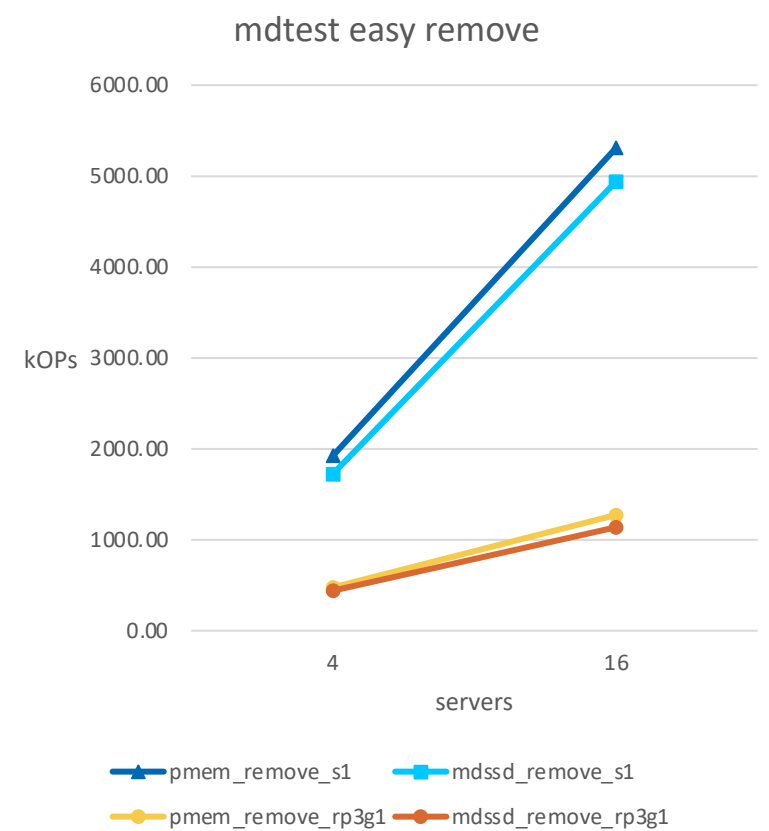
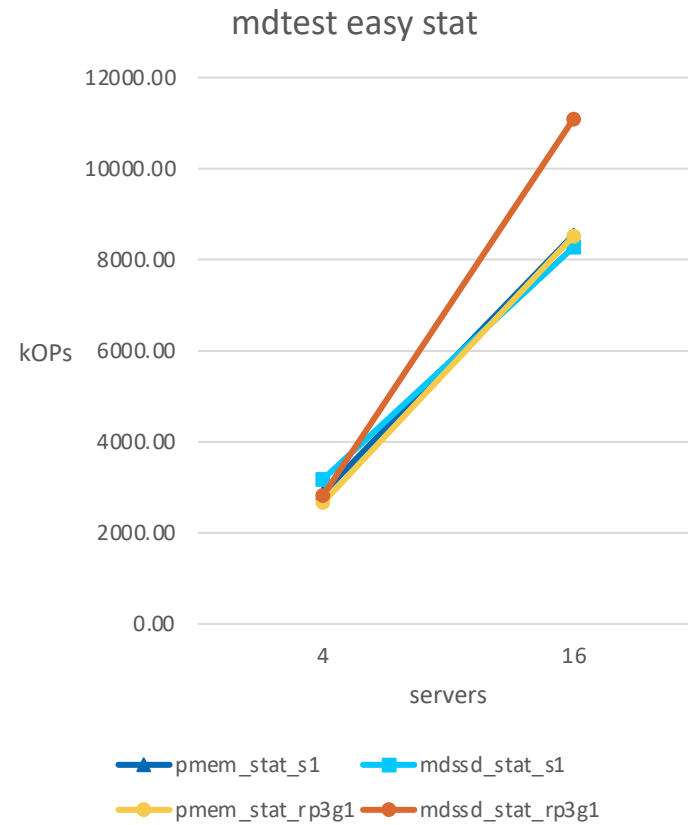
4 Server, 16 Client, Write
pmem vs split wal, data+meta




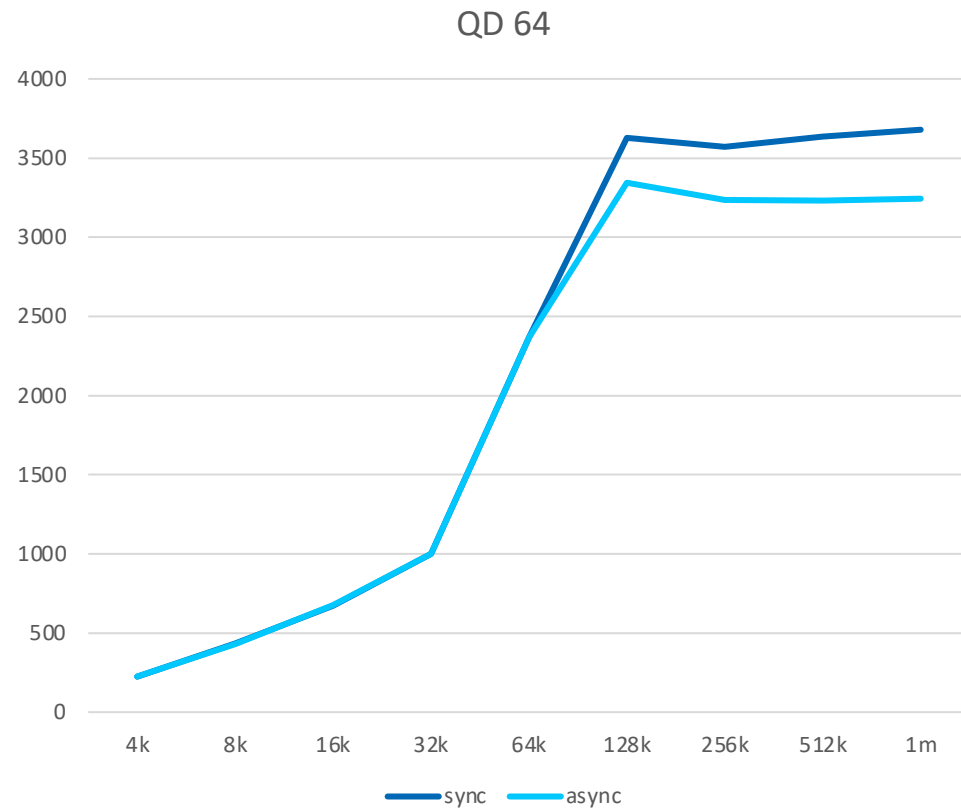
4 Server, 16 Client, Read
pmem vs split wal, data+meta



Mdtest easy (4:16 and 16: 64)



Sync VS Async submit



The Intel logo is centered on a solid blue background. It consists of the word "intel" in a white, lowercase, sans-serif font. A small blue square is positioned above the letter "i". To the right of the word "intel" is a registered trademark symbol (®) enclosed in a white circle.

intel®