

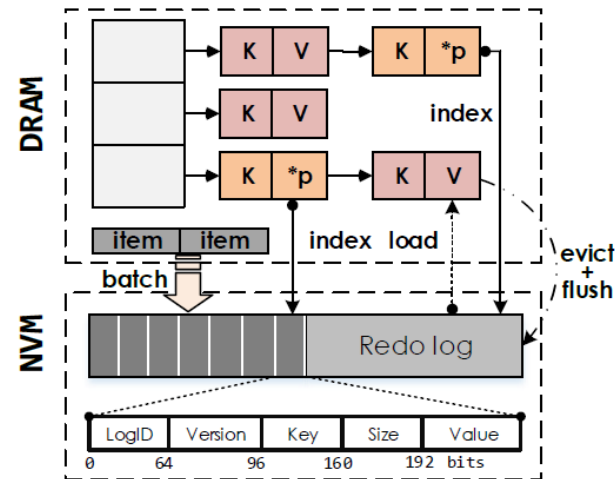
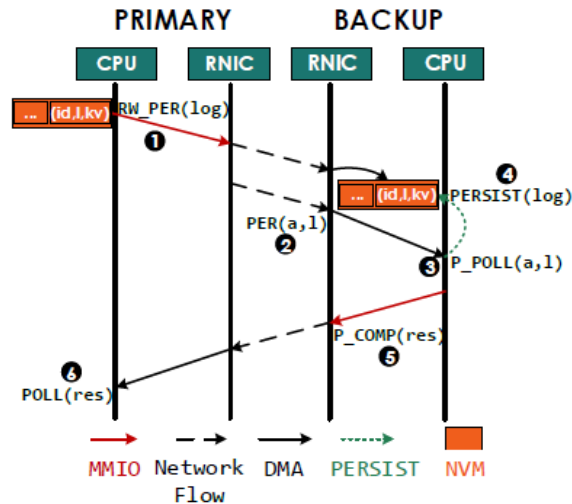
High-Availability In-Memory Key-value Store using RDMA and Optane DCPMM

**Xuecheng QI, Huiqi HU, Jinwei GUO, Chenchen HUANG,
Xuan ZHOU, Ning XU, Yu FU, Aoying ZHOU**

Frontiers of Computer Science, DOI: [10.1007/s11704-022-1123-8](https://doi.org/10.1007/s11704-022-1123-8)

Problems & Ideas

- Problems of conventional stereo matching approaches:
 - Existing studies to combine RDMA and NVM cannot deliver the desired performance due to their two-phase protocols for log shipping.
 - Most of existing studies are based on emulation that perform sub-optimally on real PM hardware.
- Ideas: A fast high-availability in-memory key-value store based on RDMA network and Optane DCPMM named FaHA, joint with an RDMA Persist PRC, and an append-only storage with pipeline batching scheme and a hotness-aware differential hash index.

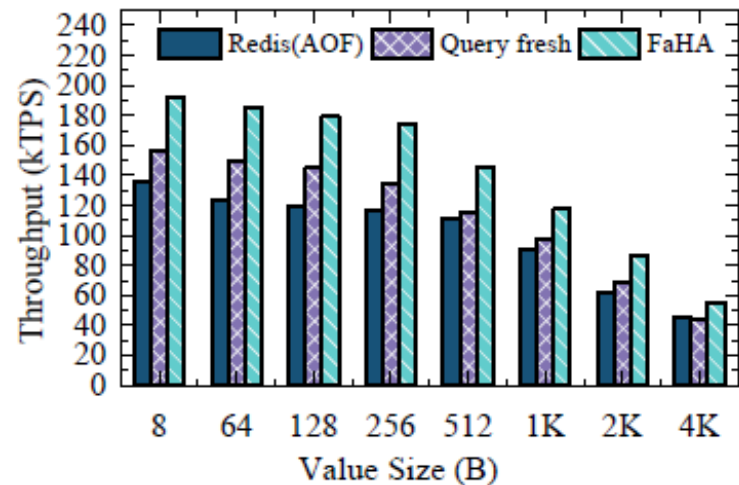
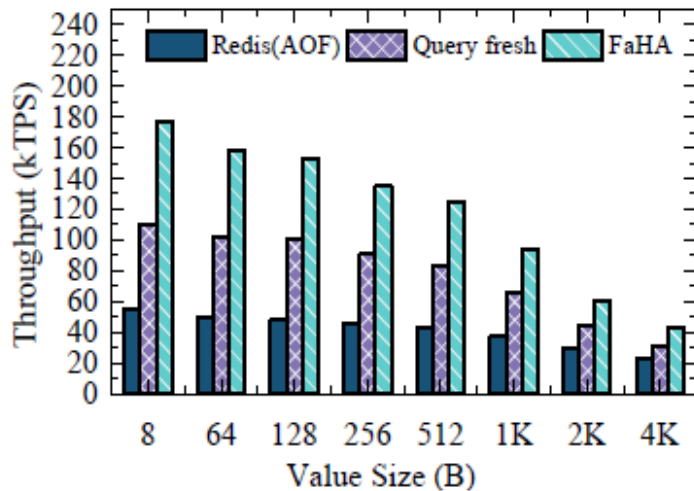


Figures of the design of FaHA.

Left: Log Shipping using RDMA Persist RPC; Right: Architecture of Two-tier Storage.

Main Contributions

- Contributions:
 - A RDMA Persist RPC for writing data to remote NVM while ensuring remote data persistence, it needs only one round-trip and persistence overhead, significantly reducing the overhead of primary for log shipping;
 - An append-only storage with pipeline batching scheme to relieve write amplification of NVM;
 - A hotness-aware differential hash index in the DRAM for better read performance.



Evaluations show that FaHA-based Redis outperforms the state-of-the-art solution by up to 1.6x throughput. Left: 100% Put Throughput of the Primary; Right: 50% Put, 50% Get Throughput of the Primary.