

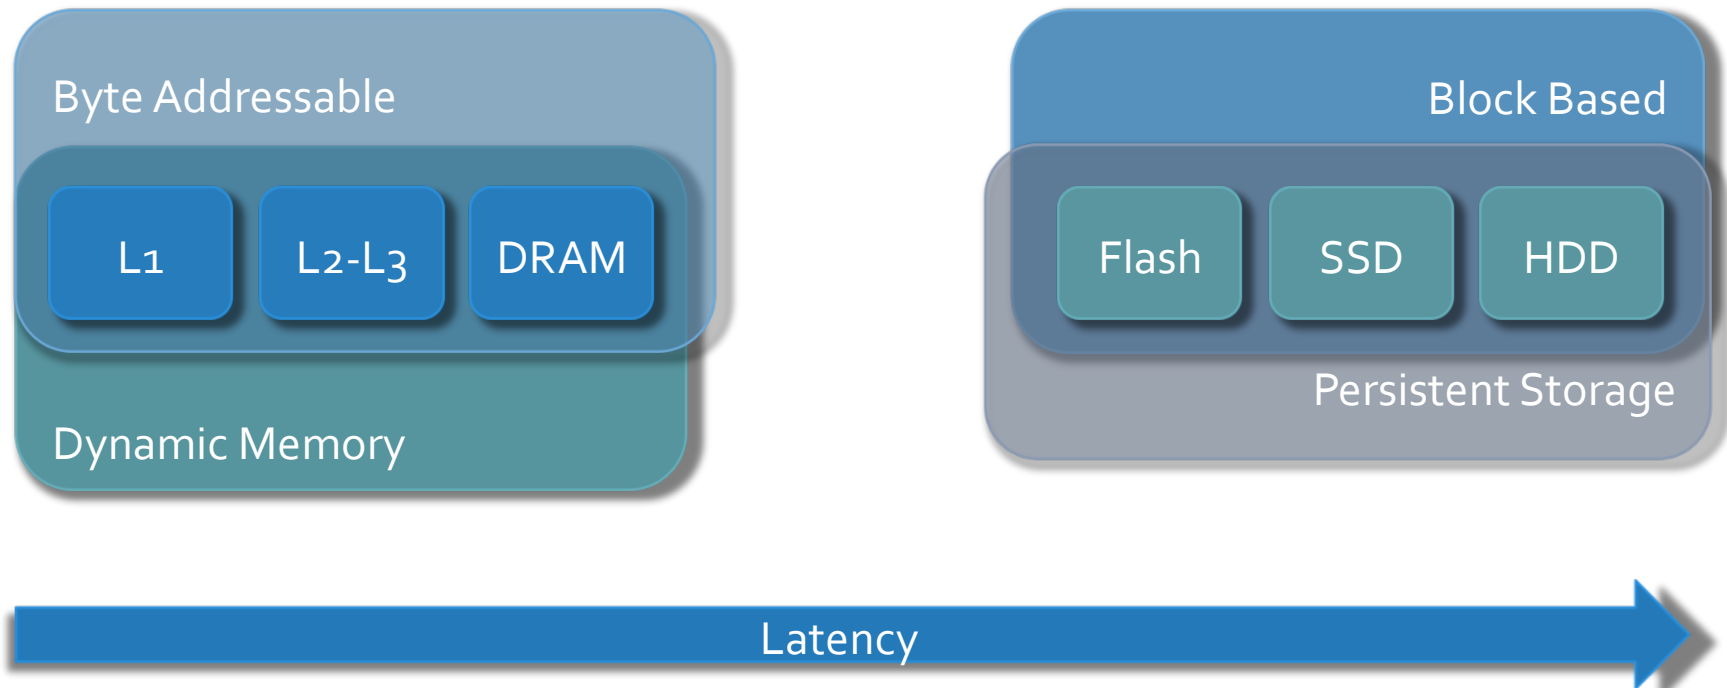
Containerizing Byte-Addressable NVM

PDSW-DISCS 2016: 1ST JOINT INTERNATIONAL WORKSHOP ON PARALLEL
DATA STORAGE & DATA INTENSIVE SCALABLE COMPUTING SYSTEMS

Ellis Giles

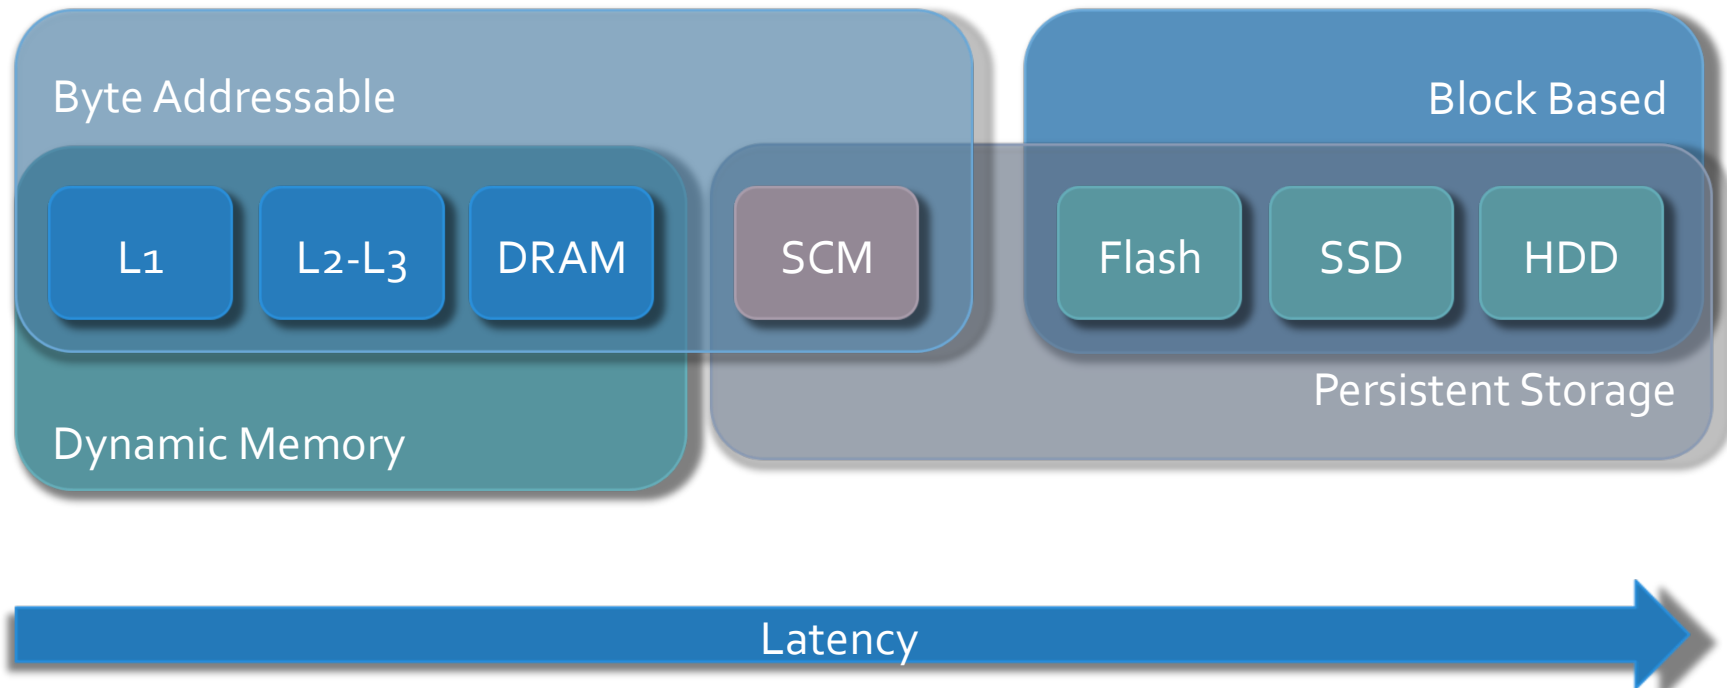
PhD Candidate, Rice University

Byte-Addressable Dynamic Memory vs. Block Based Persistent Storage



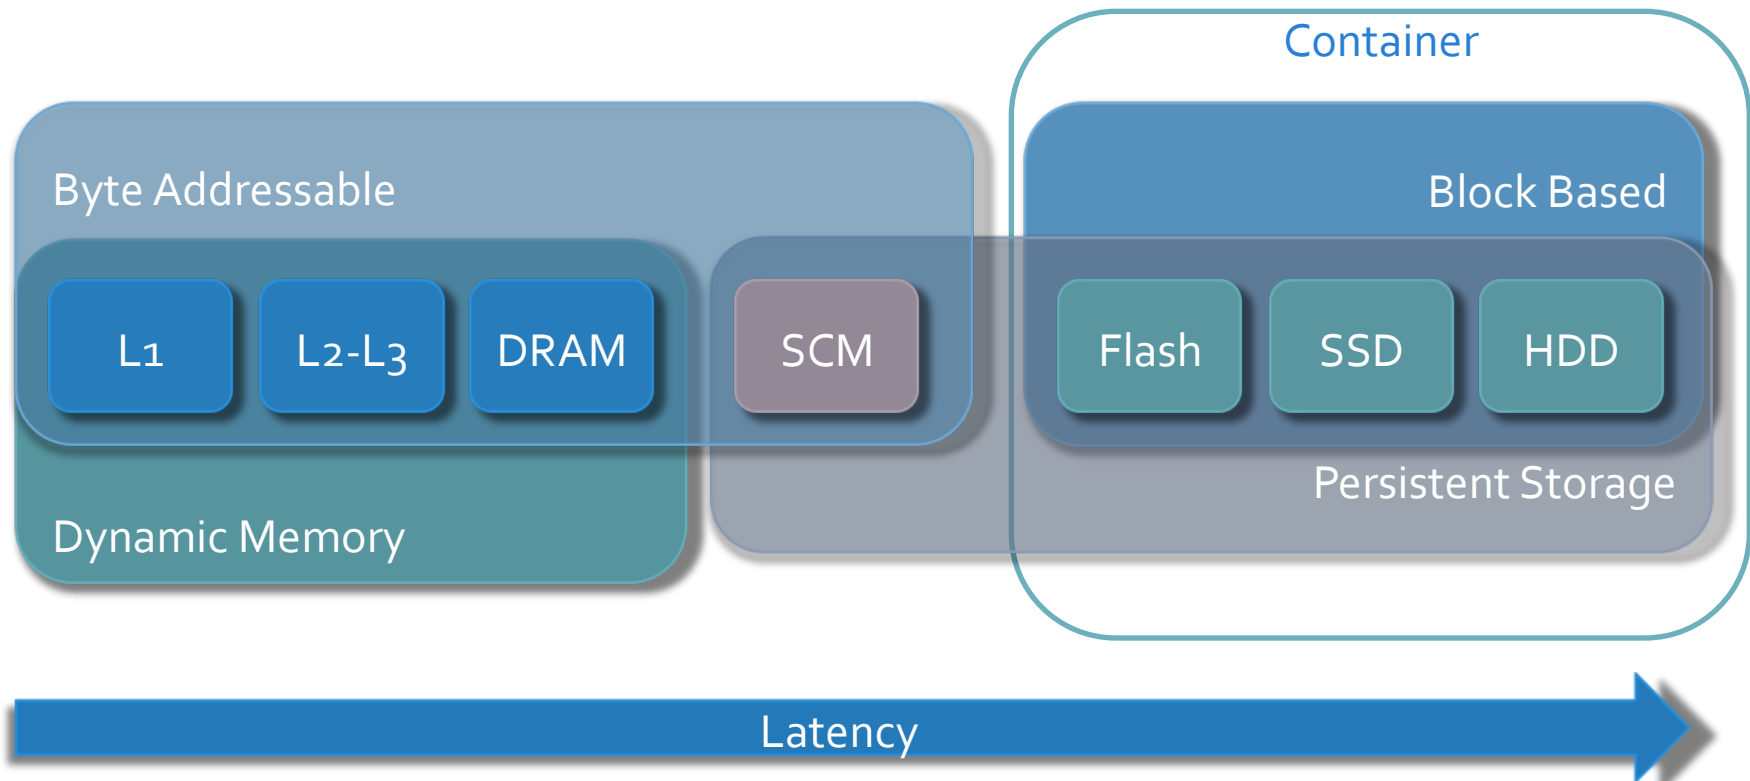
Applications typically live in the middle, operating on data in byte-addressable cache attached to main memory and serializing data for persistence storage operations.

Byte-Addressable, Non-Volatile Memory




Storage Class Memory (SCM) – byte-addressable non-volatile memory on the main memory bus. SCM technologies such as PCM can have very large capacities.

Container Based Virtualization

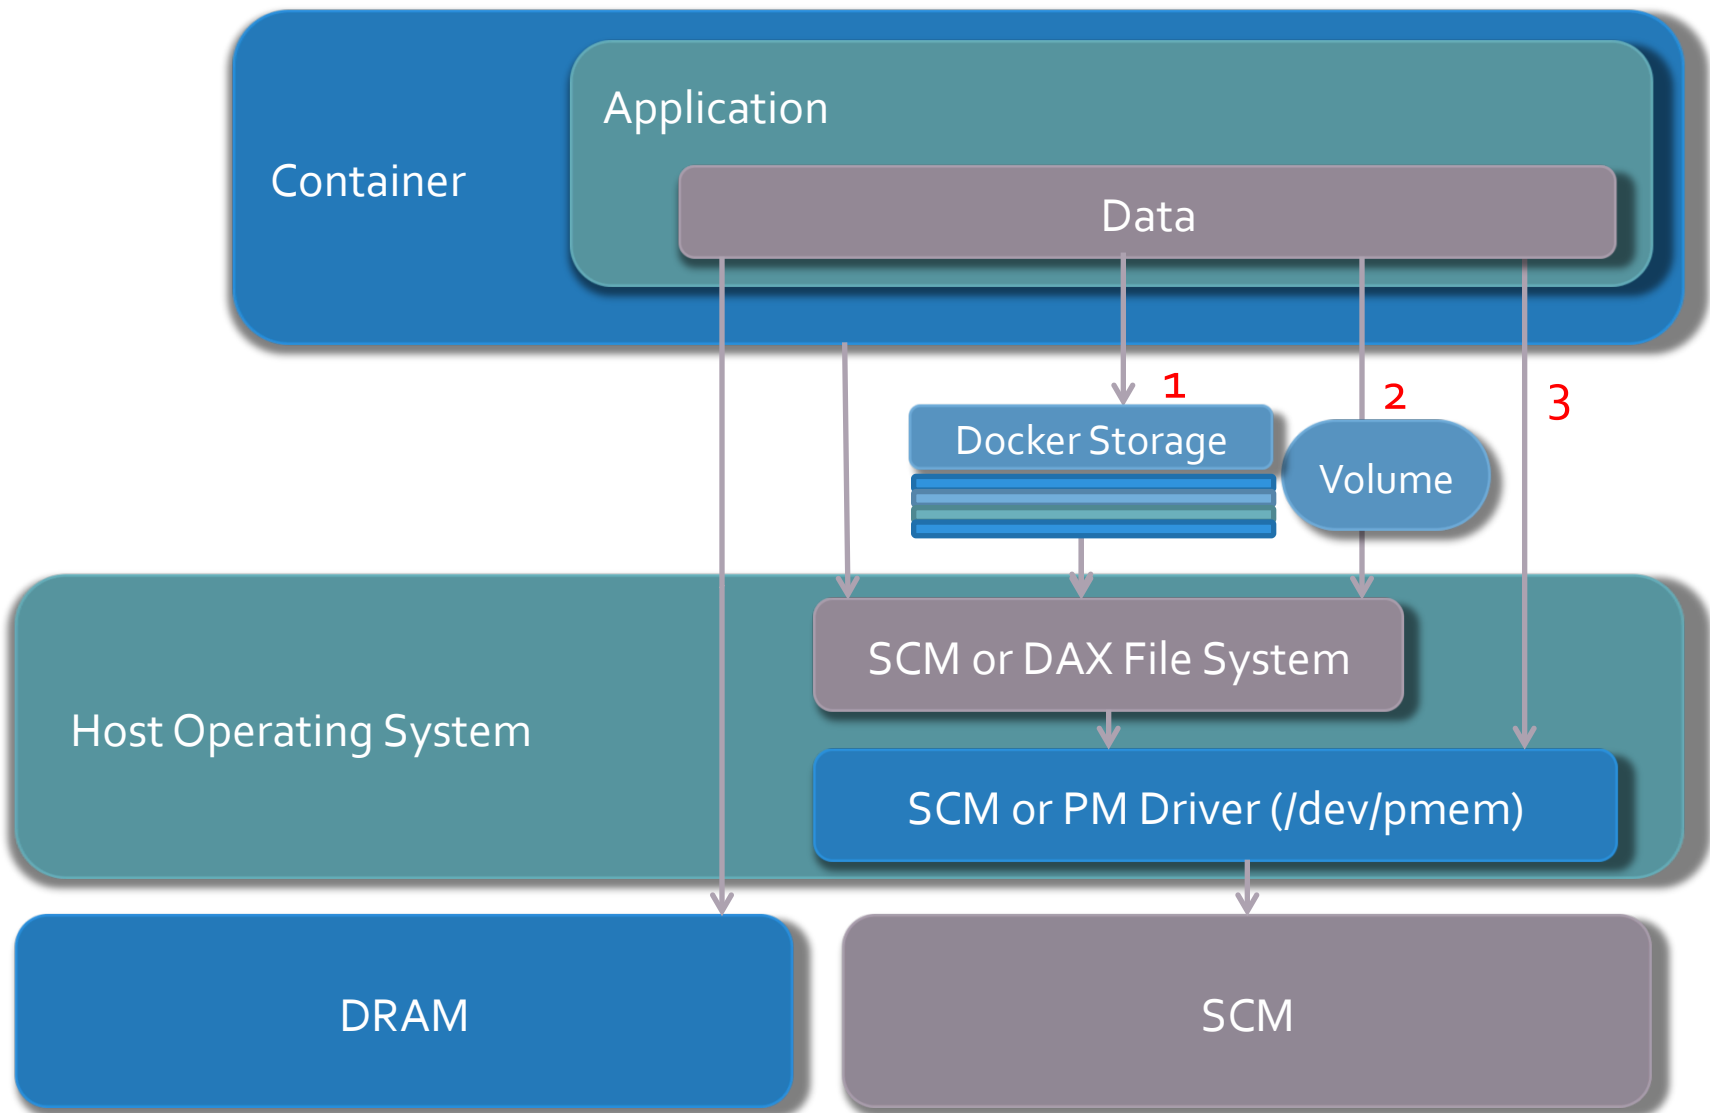


Containers provide persistence isolation through a chroot file system.
Linux Containers can restrict virtual memory and CPU usage with Linux CGroups.

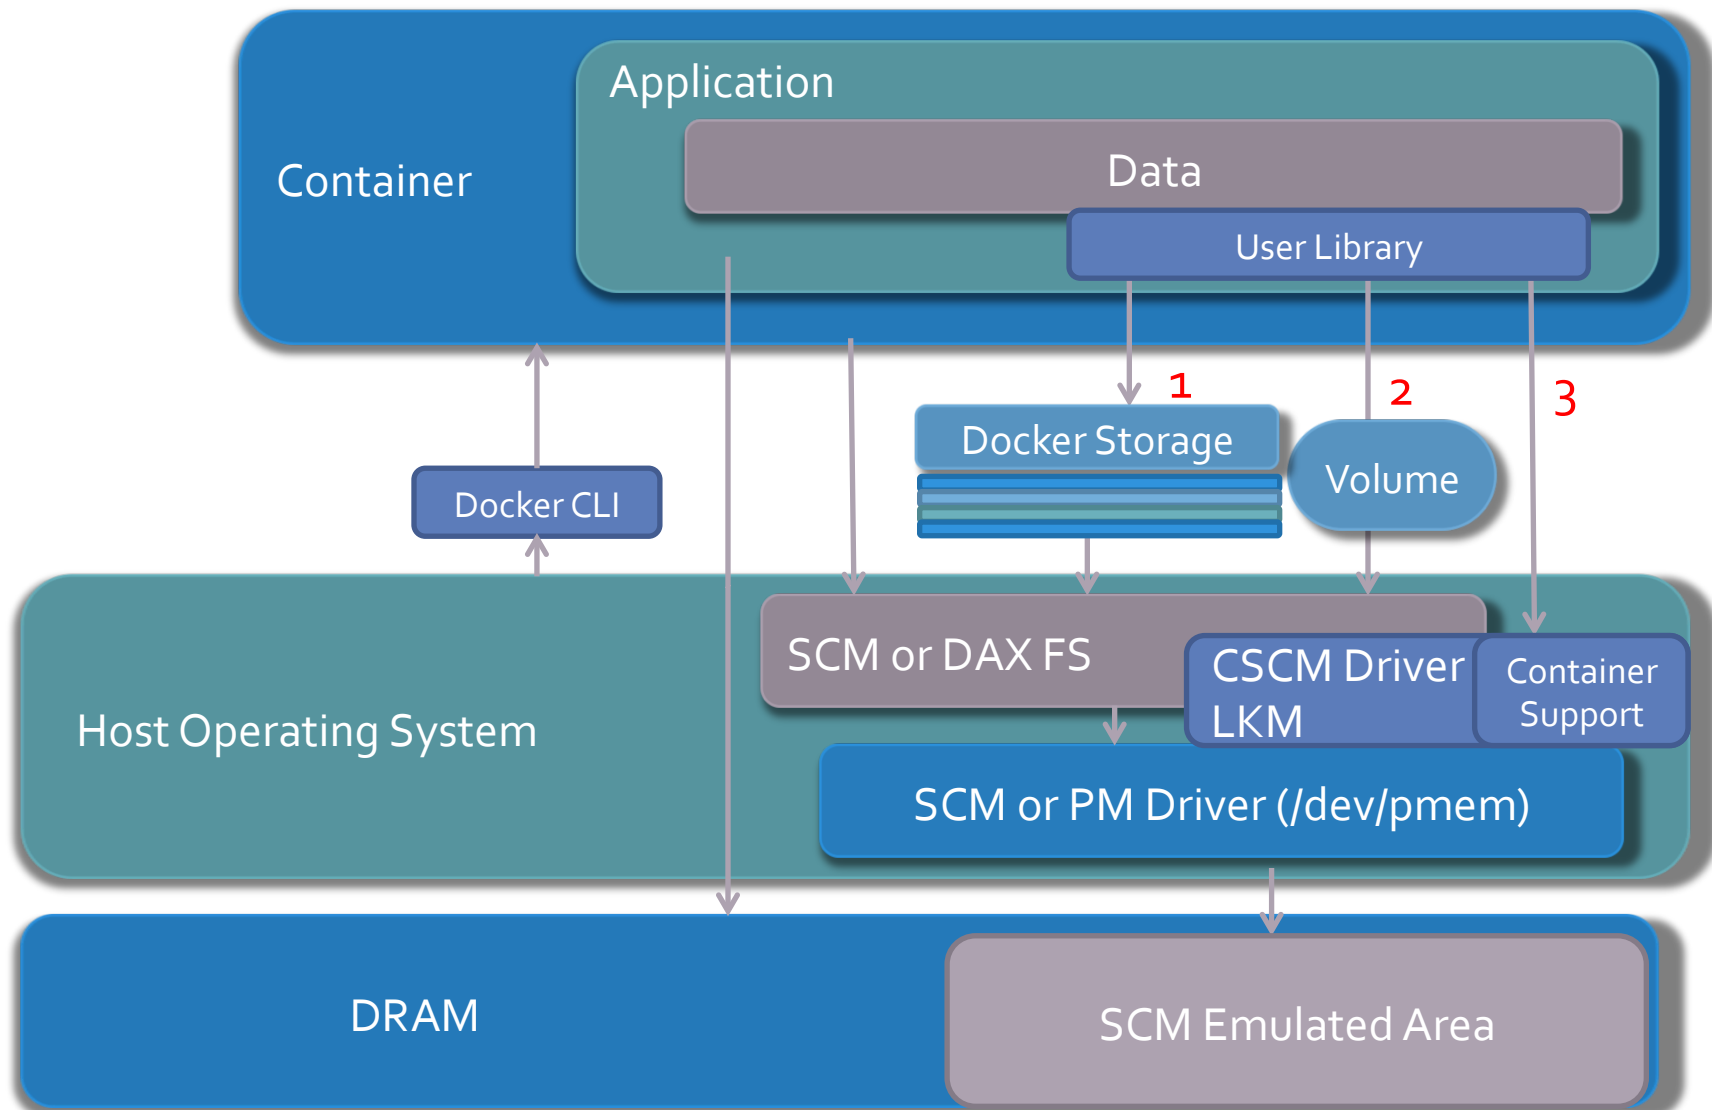
Containerizing SCM



Can we extend containers to support a persistent virtual memory address space without suffering from low performance or violating isolation guarantees?




Docker Containers can access persistent data in three ways, layered Storage, Volume or privileged device driver. For SCM this creates performance and isolation challenges.



Containerized Storage Class Memory Driver (CSCM) - preserves the isolation guarantees for container based SCM storage while providing high performance.

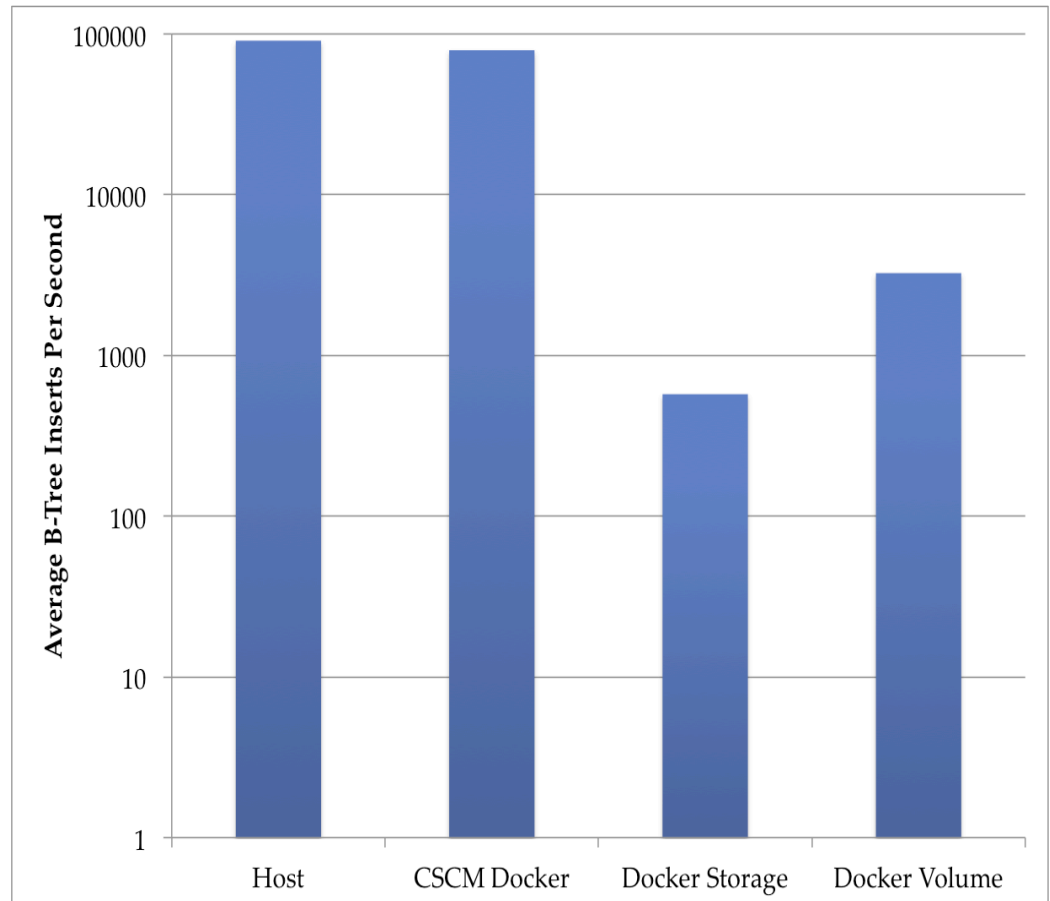
STREAMS Benchmark

- STREAMS Benchmark
- 3, 8-byte double 5 Million element arrays
- Element by Element test
- COPY/SCALE/ADD/TRIADD



Transactional B-Tree Element Inserts

- Modified a B-Tree data structure for transaction safe element inserts
- New nodes allocated using the scmallocc user library
- Writes to variables, pointers, etc. go into a log area first
- 1 Insert can update multiple values and pointers



Thank You!

Full Paper To Appear:

“Container-Based Virtualization for Byte-Addressable NVM Data Storage”
Workshop on Scalable Cloud Data Management, December 7th 2016
IEEE Big Data Conference, Washington D.C.

Special thanks to Rice University Professors Scott Rixner and Peter Varman for encouragement and discussions leading to the idea development and to the NSF for supporting my research.