CNEXLABS

Open-Channel SSDs Offer the Flexibility Required by Hyperscale Infrastructure

Matias Bjørling

August 8th, 2017

Public and Private Cloud Providers

Workloads and Applications

Multi-Tenancy

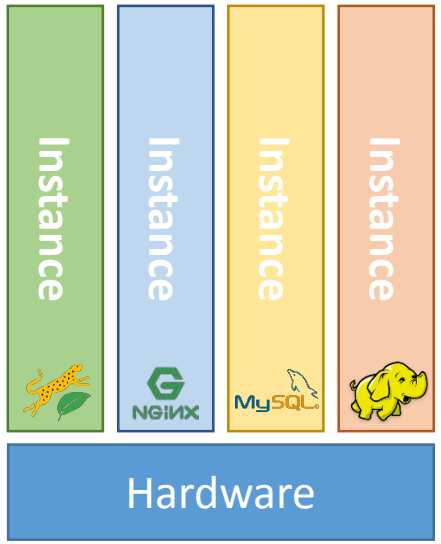

Databases

Microsoft SQL Server

MySQL

Neo4j

RocksDB

Read or Write-Heavy

NAND Capacity Grows Bigger

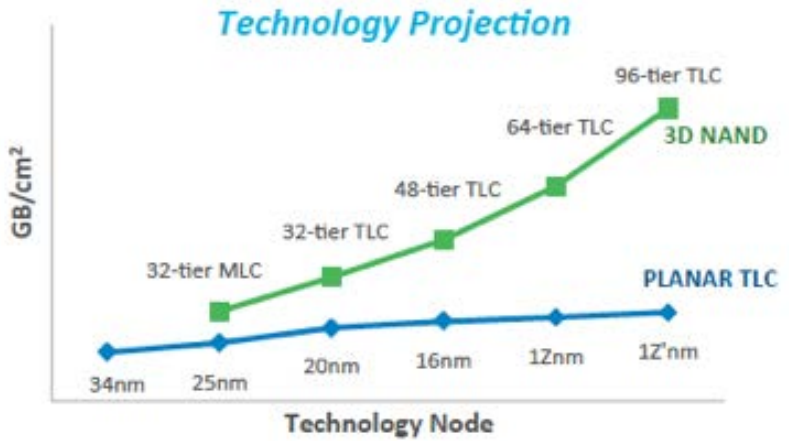

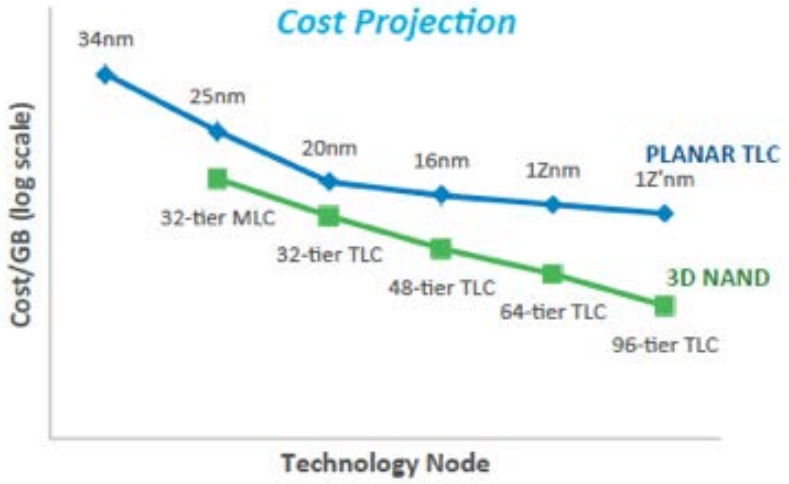

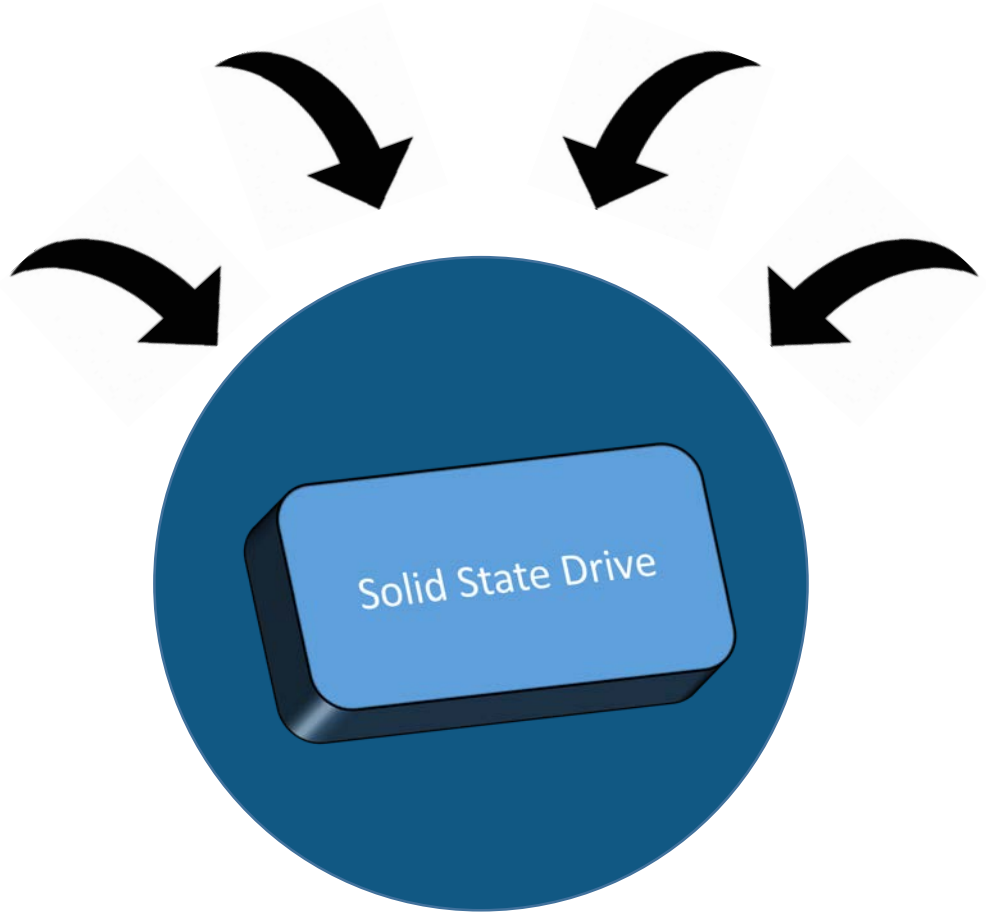

Performance – Endurance – DRAM overheads

Source: William Tidwell -The Harder Alternative – Managing NAND capacity in the 3D age

Read Latency with 0% Writes

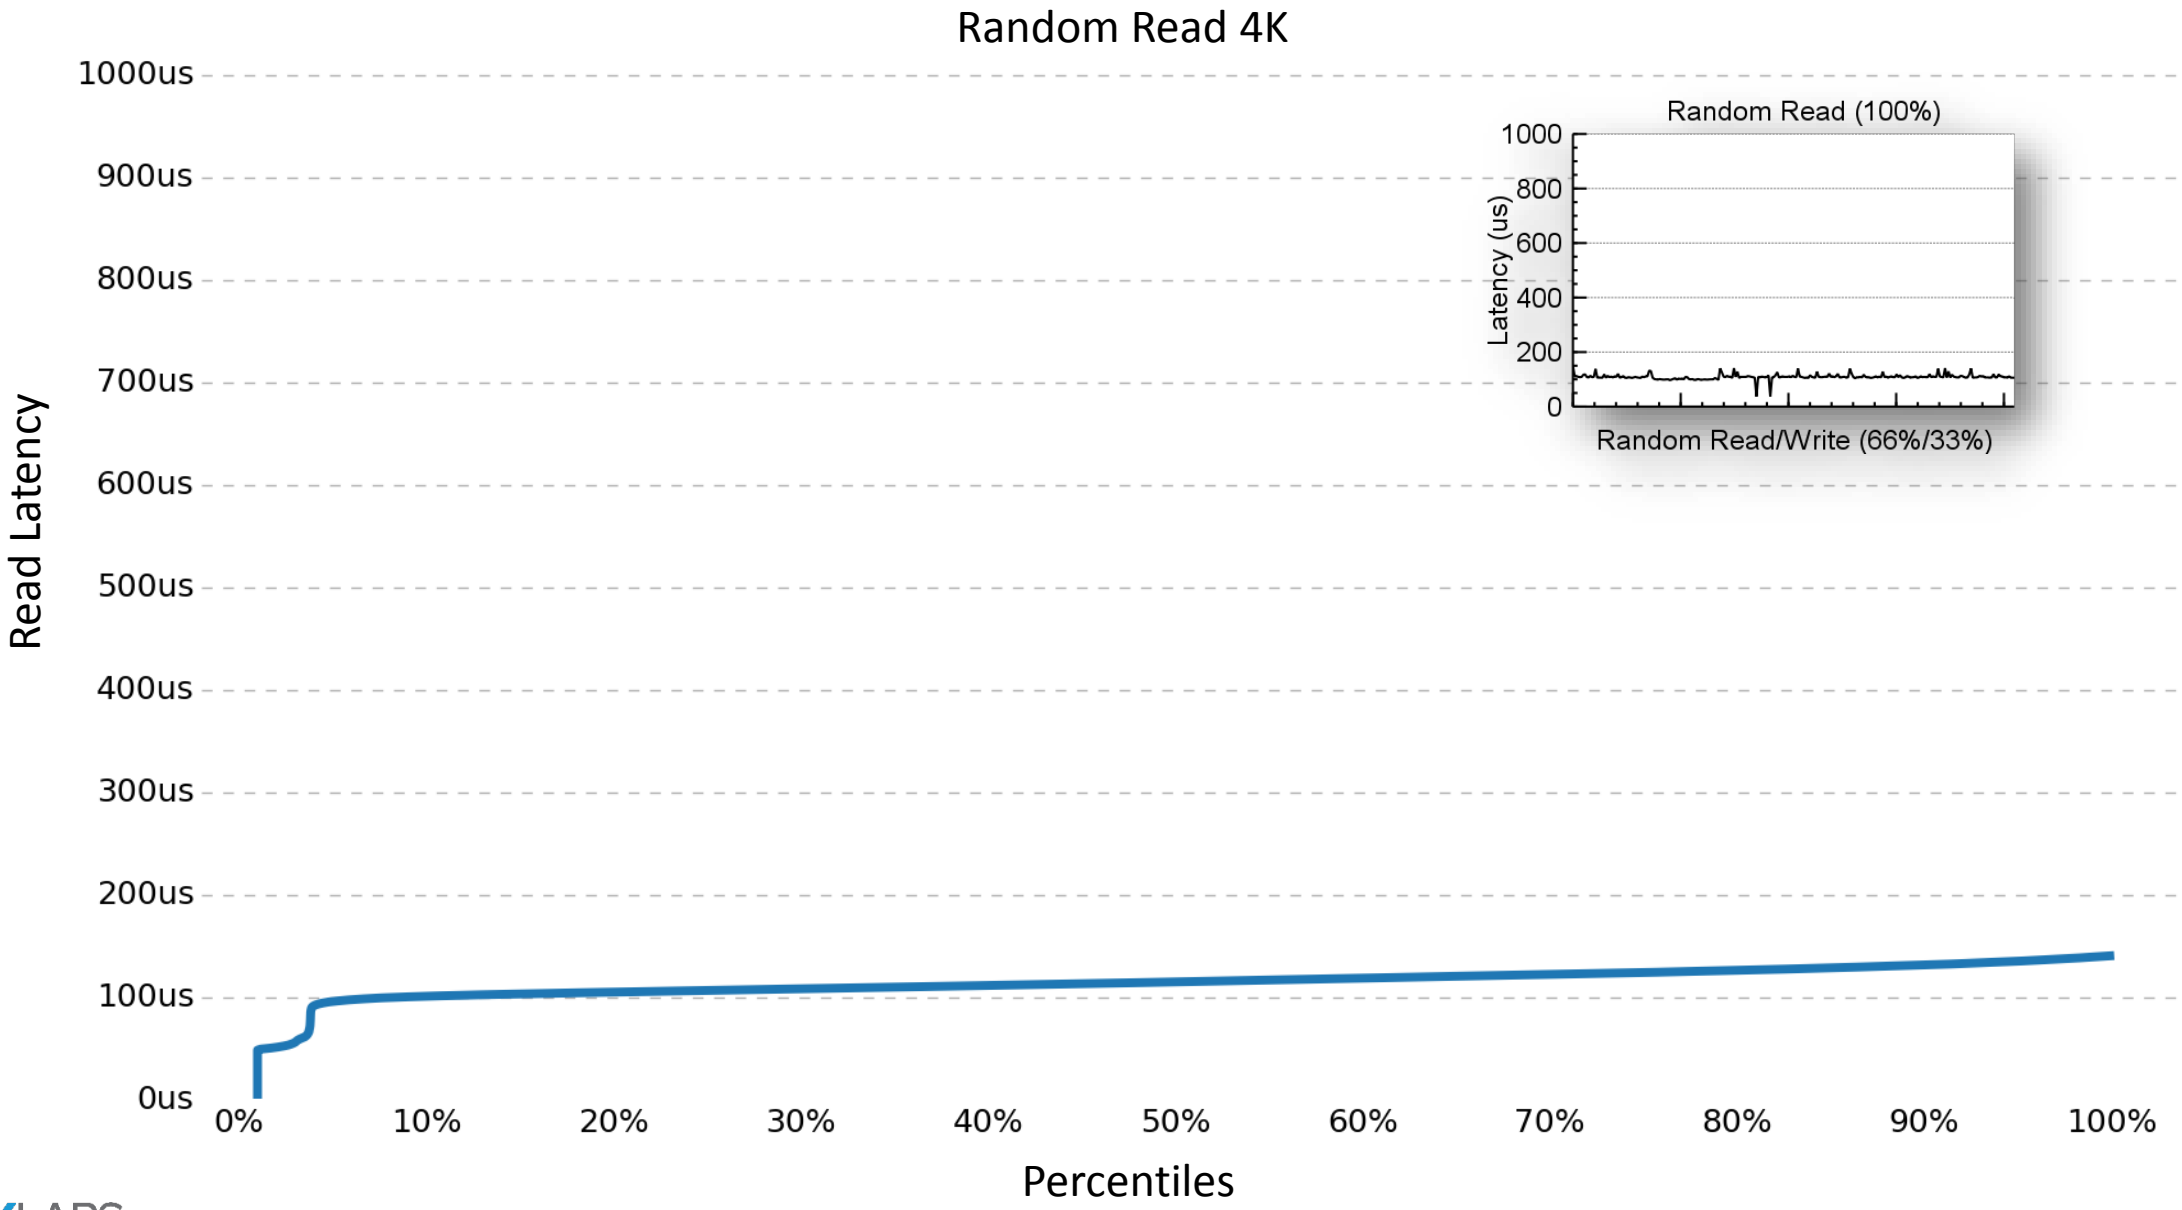

Read Latency with 20% Writes

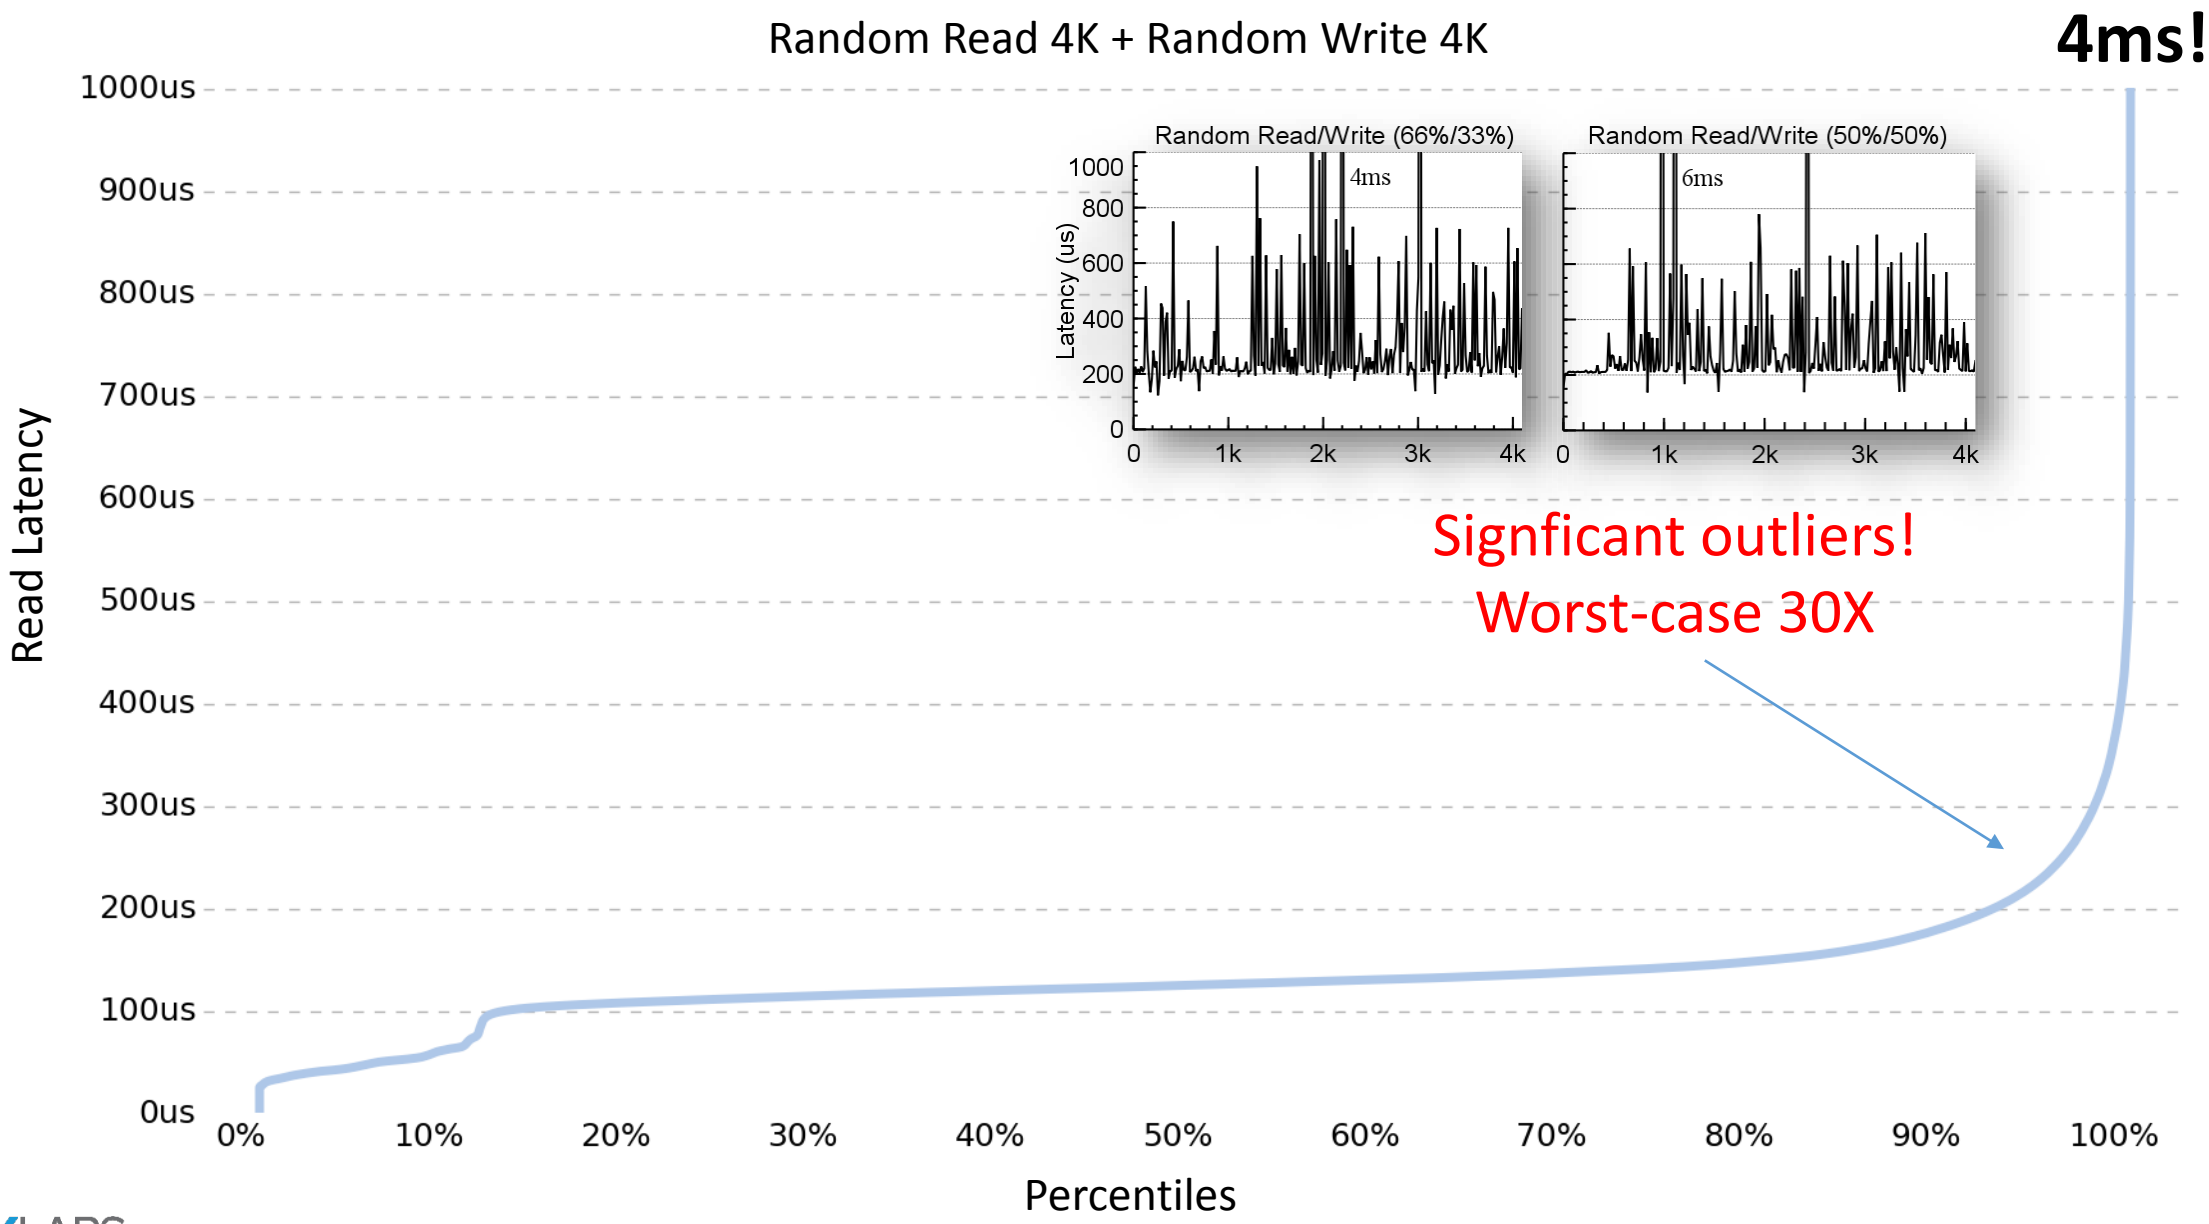

Requirements

1. Flexible enough for software to evolve faster than hardware
2. Strong QoS Guarantees
3. Continuous access – No maintenance windows
4. Support a broad set of applications on shared hardware
5. Rapid enablement of new NAND generations / media agnostic
6. Vendor neutrality & supply chain diversity

Open-Channel SSDs

I/O Isolation

Enable I/O isolation between tenants by allocating your SSD into separate parallel units.

Predictable Latency

No more guessing when an IO completes. You know which parallel unit is accessed on disk.

Data Placement & I/O Scheduling

Manage your non-volatile memory as a block device, through a file-system or inside your application.

Rebalance the Storage Interface

- Expose internal lanes to host
  - Each lane is a parallel unit
  - Logical or Physical
  - Performance characteristics
- Building Block
  - Similar to HDD SMR interface
  - Extension to NVMe
  - Can implement I/O Determinism, streams, and other new data management schemes
- Open specification

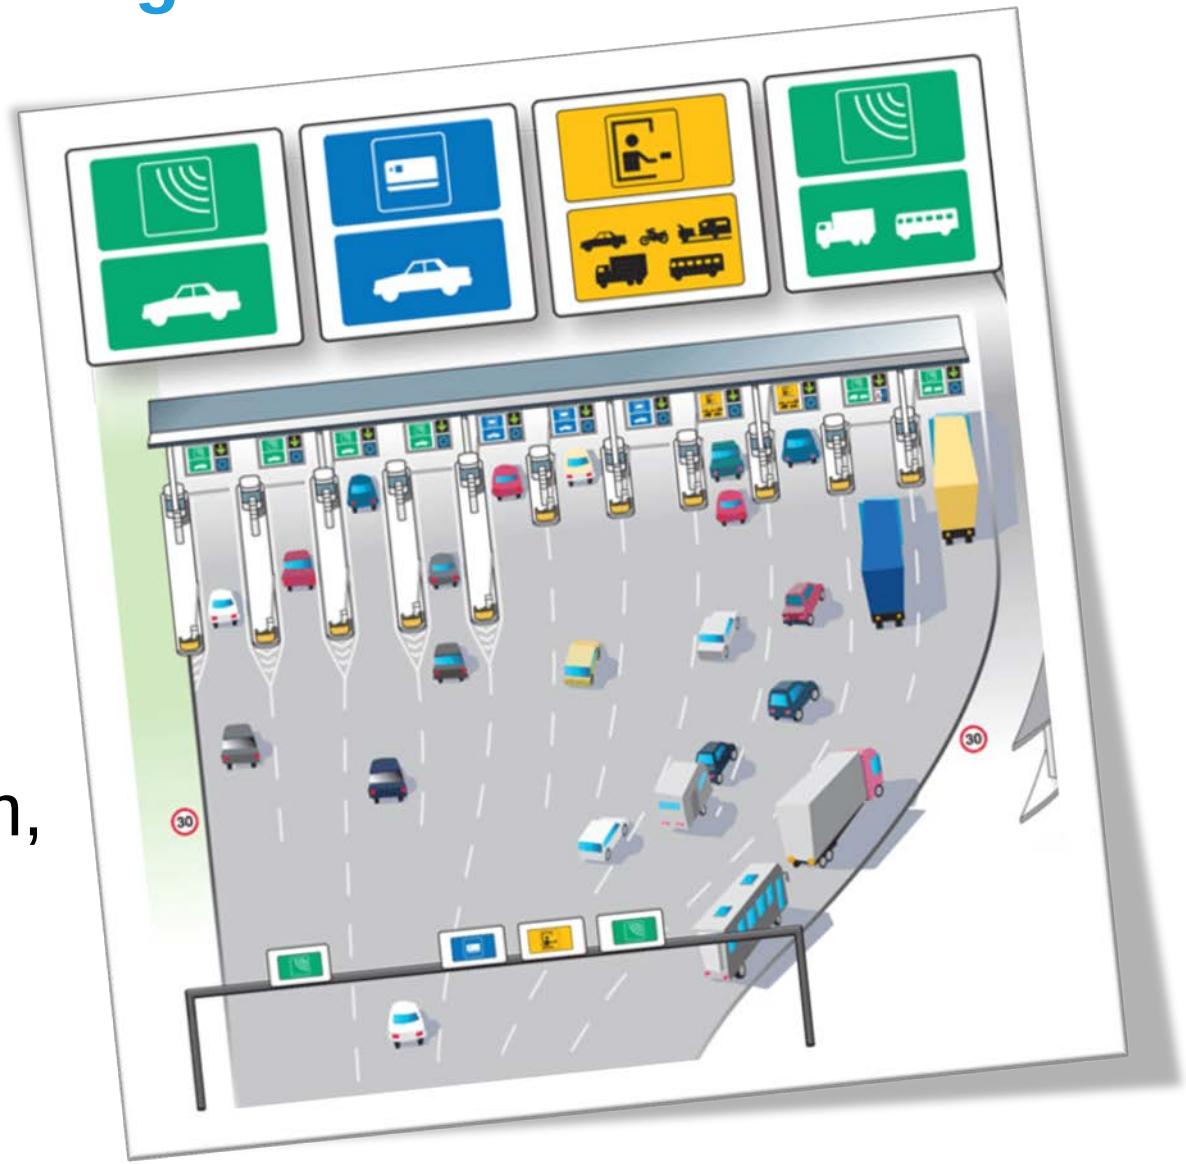

Parking Lot

Unsupervised Parking Lot

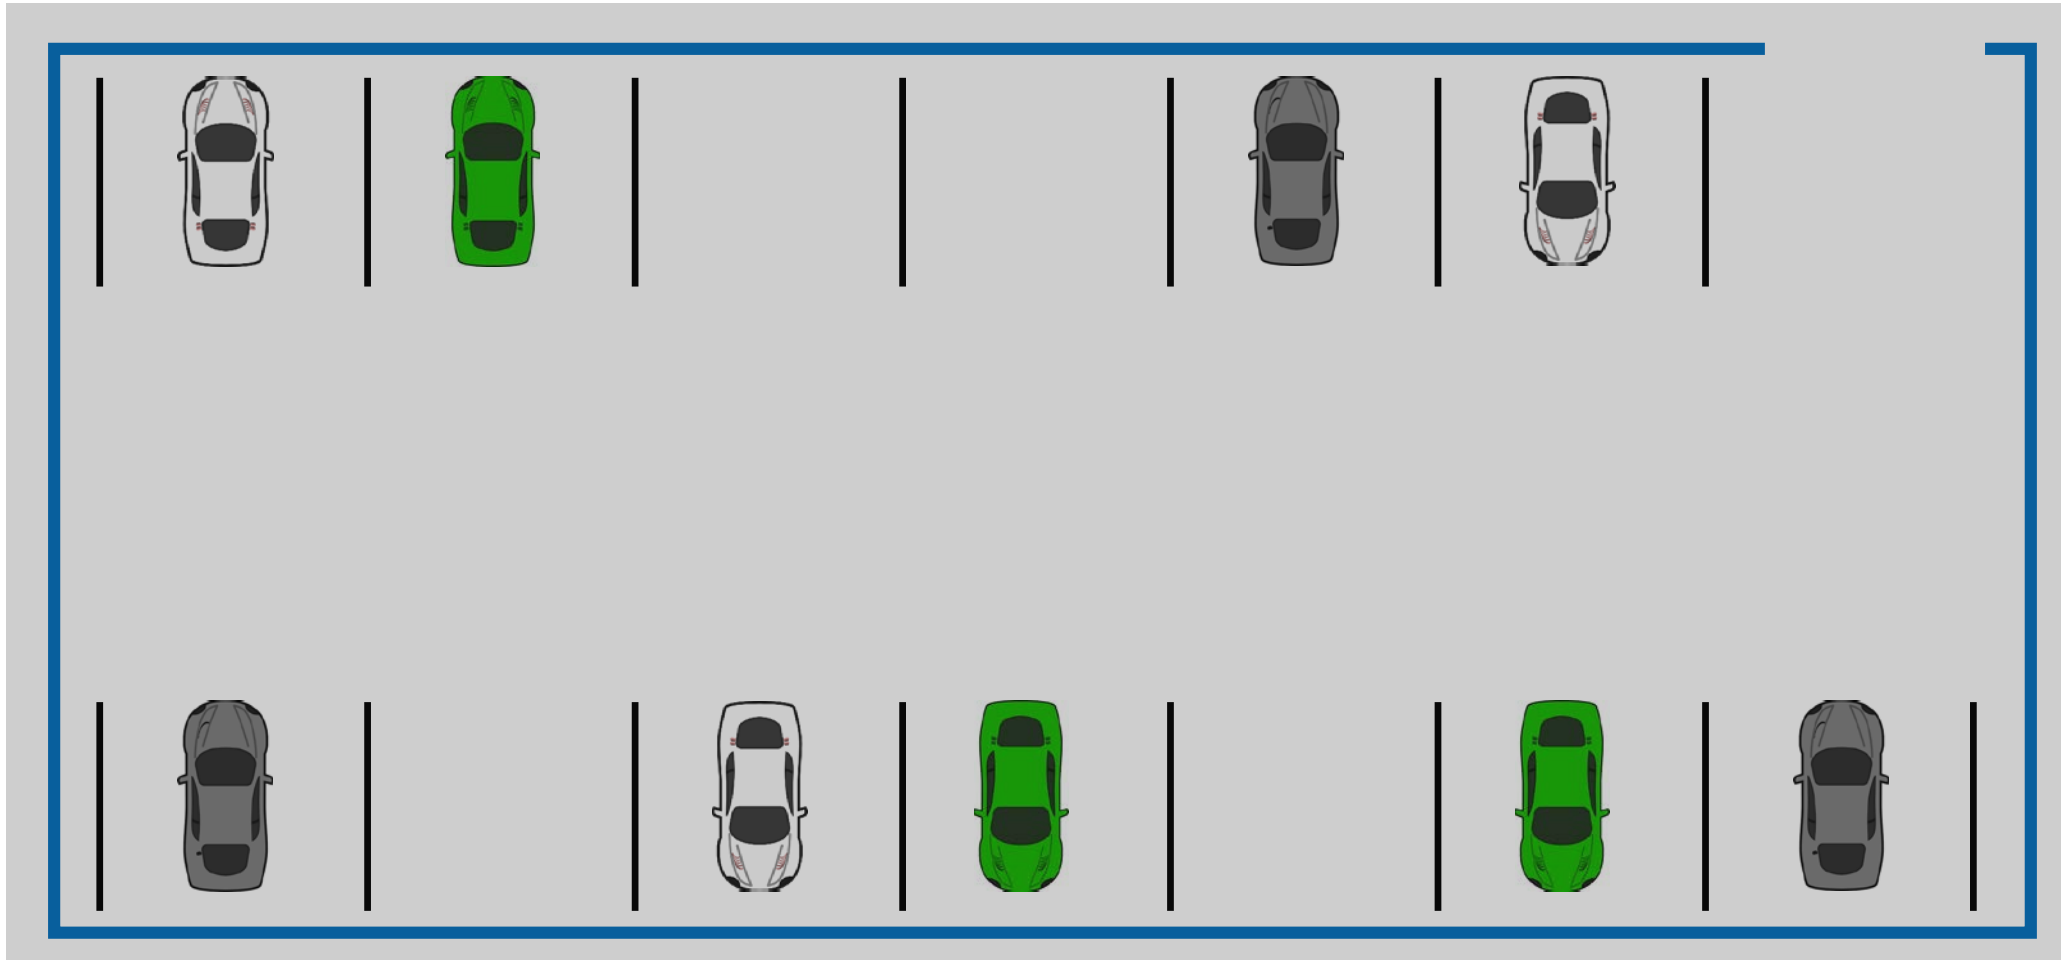

Supervised Parking Lot

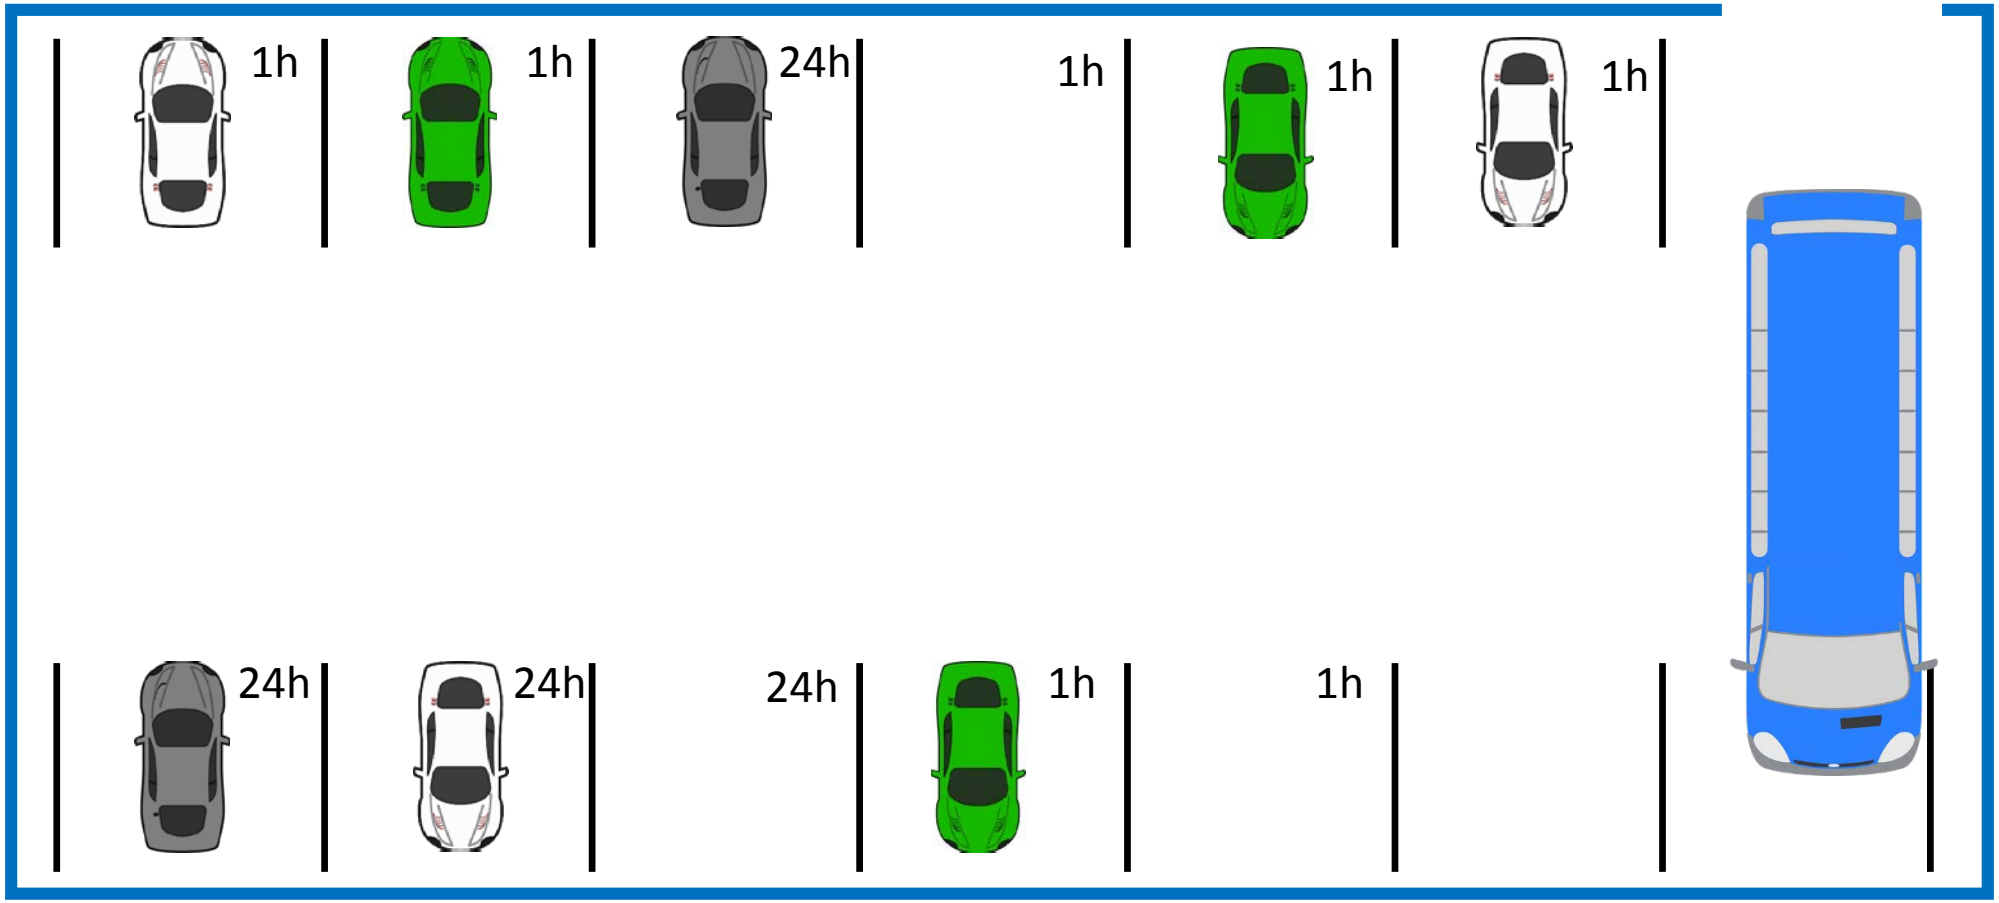

Predictable Latency

- 4K reads during 64K concurrent writes
- Consistent low latency at 99.99, 99.999, 99.9999

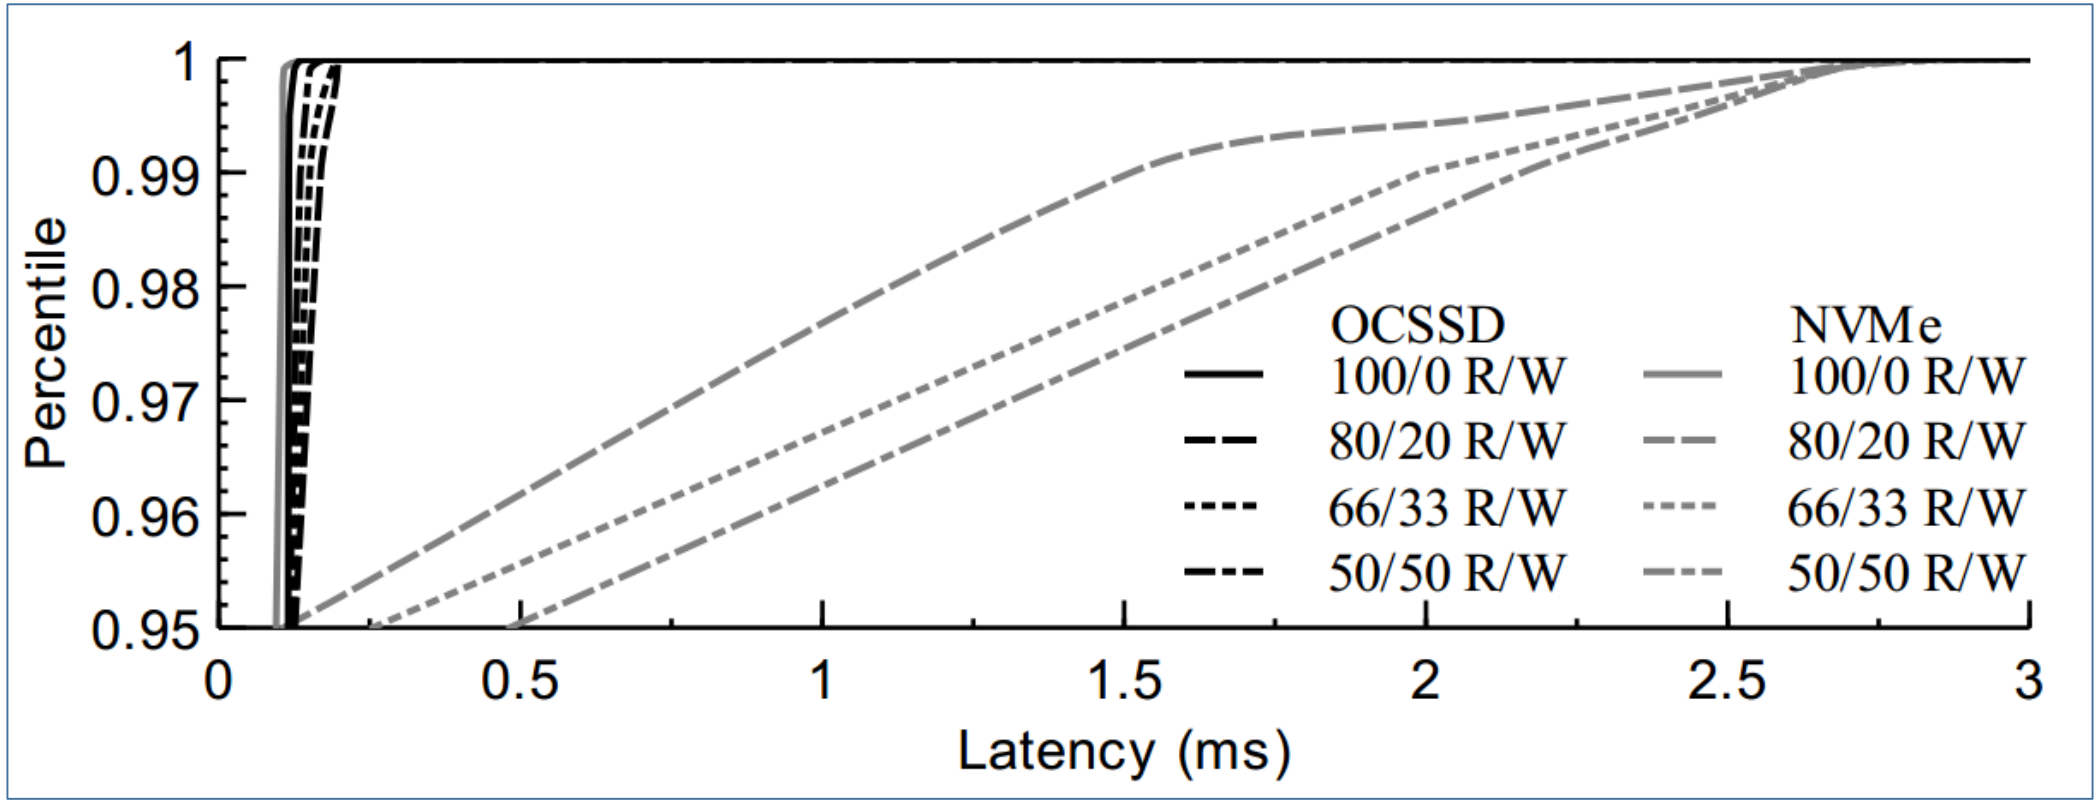

Demo using Open-Channel SSD

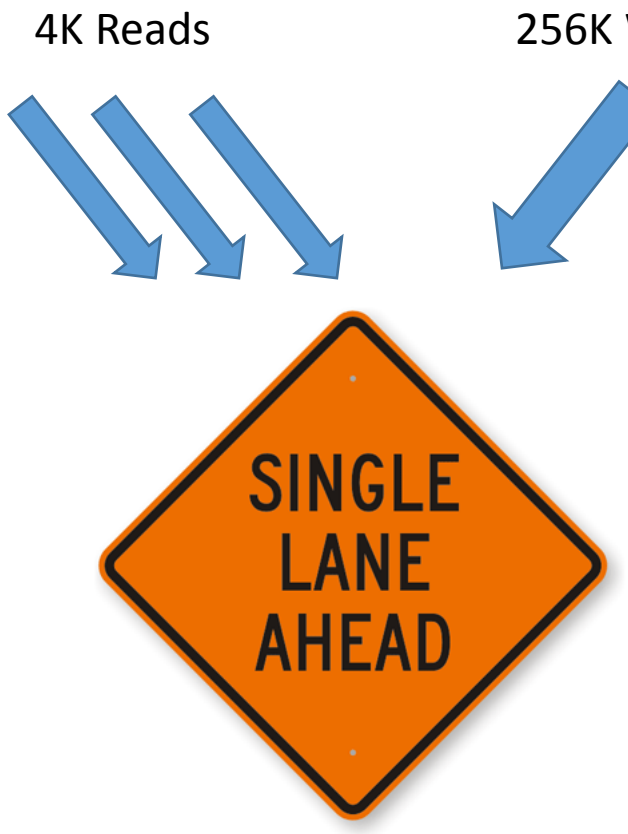

Solid State Drive

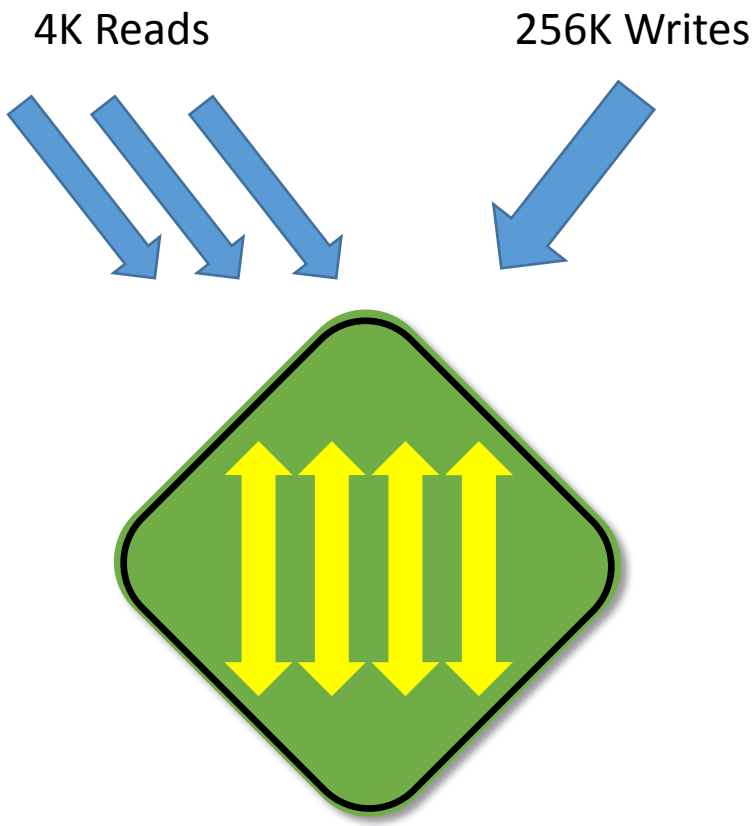

Open-Channel SSD
(Instance-based)

Scalability

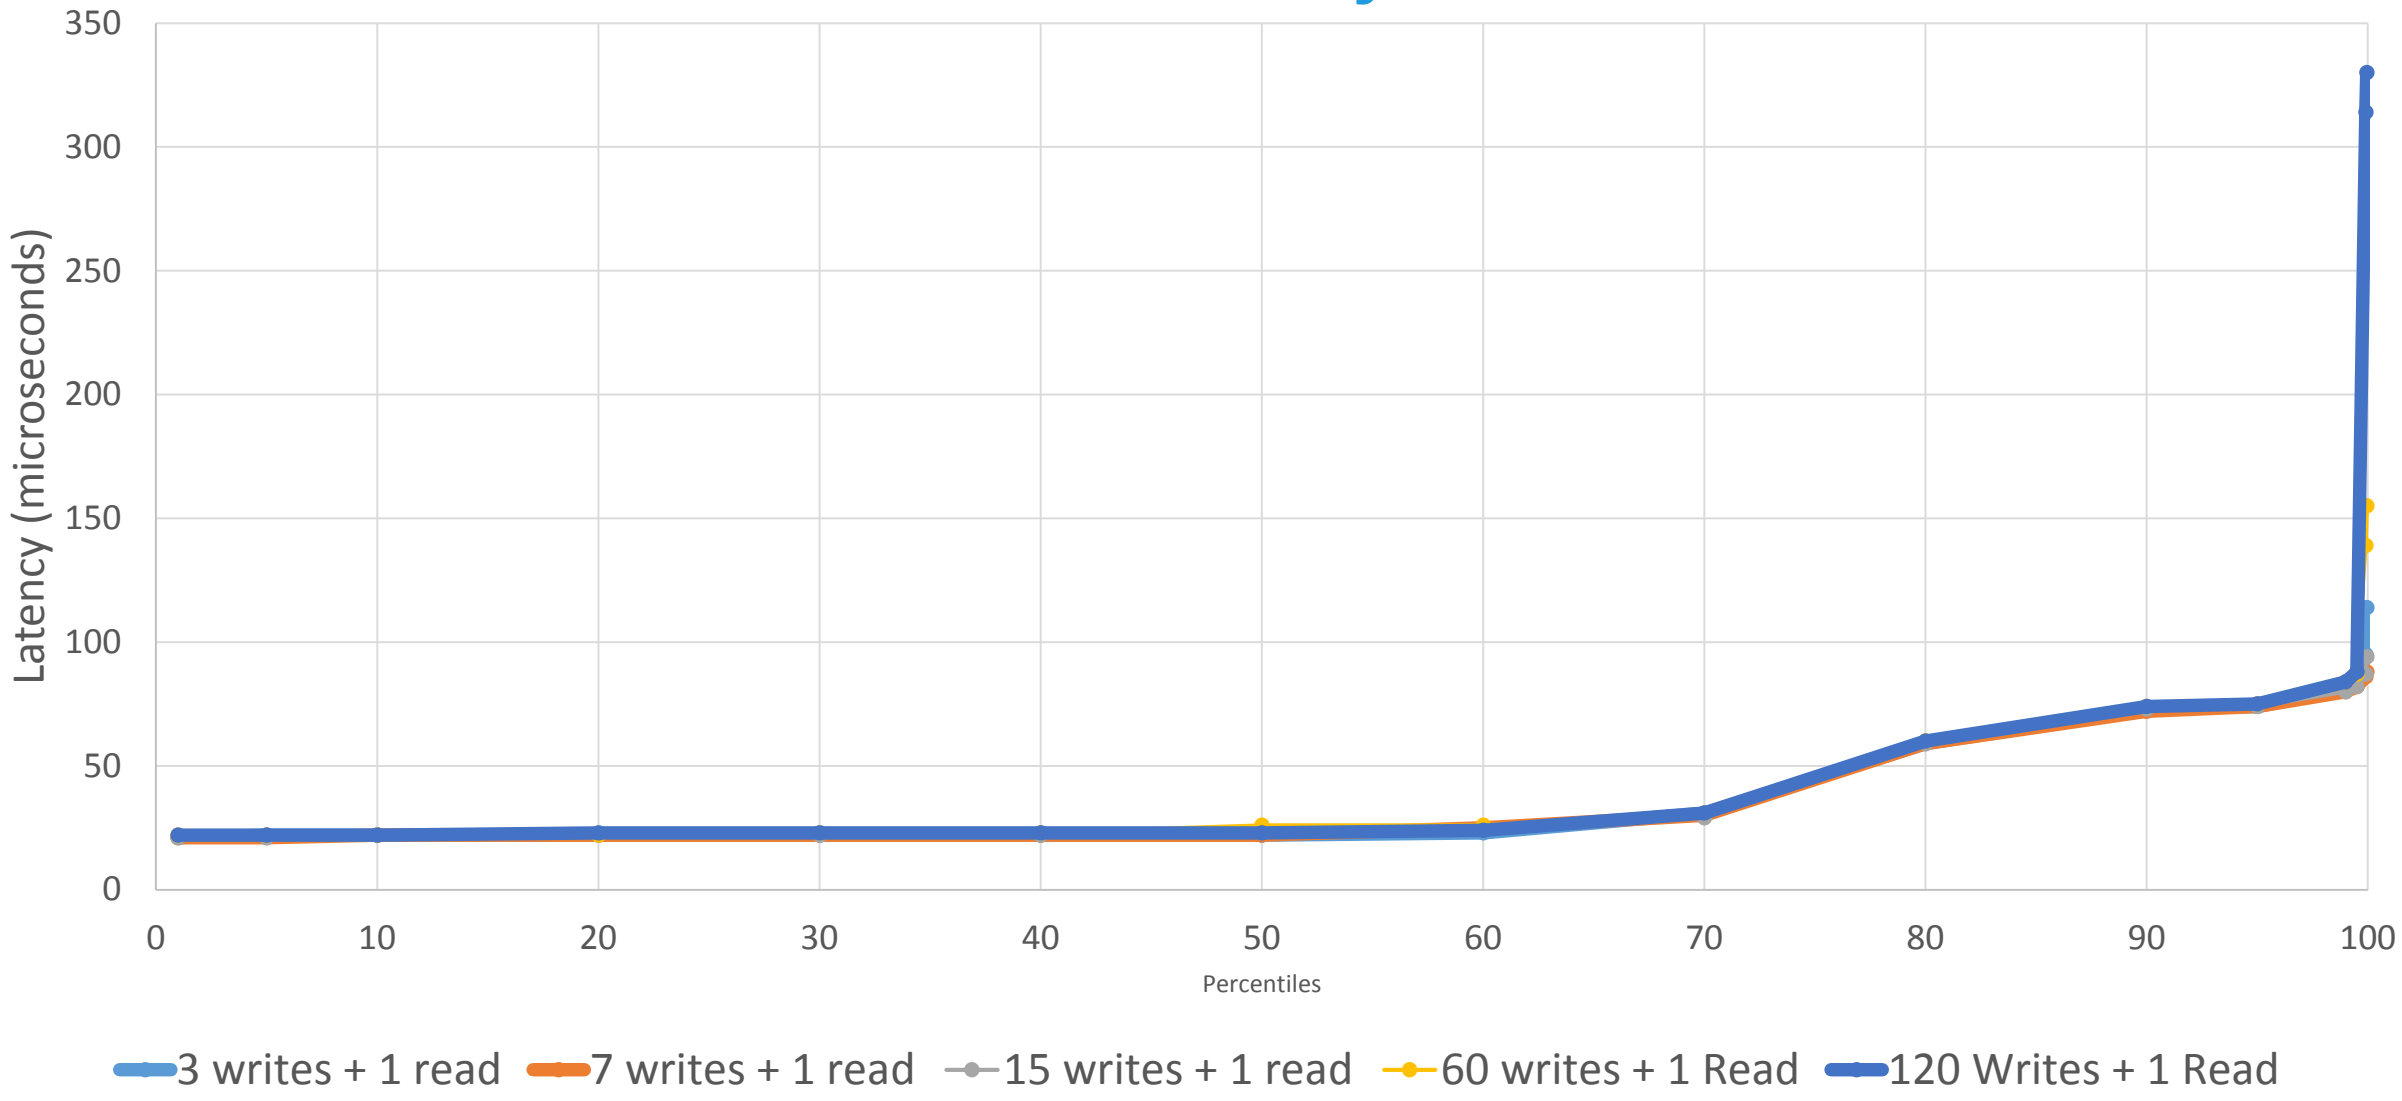

Status

- Active community
  - Multiple drives in development by commercial SSD vendors
  - Multiple papers on Open-Channel SSDs
- Software stack is growing
  - LightNVM subsystem in Linux kernel 4.4.
  - User-space library (liblightnvm) support with Linux kernel 4.11.
  - pblk upstream with Linux kernel 4.12.

Conclusion

- New storage interface that solves the software-defined storage challanges for cloud vendors
  - I/O Predictability
  - I/O Isolation
  - Application-level control of data placement and I/O scheduling
- Draft specification is open and available for implementors
- Formal standards initiative underway, target spec release Q1 2018
- More information available at: http://lightnvm.io

See demos at their booths

CNEX Labs, Inc.

Teaming with NAND Flash manufacturers and industry leaders in storage and networking to deliver the next big innovation for solid-state-storage.