Accelerating Business Applications with Flash

Scott Tracy

VP, Storage Software Engineering

Oracle

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Oracle is Optimizing Flash Integration

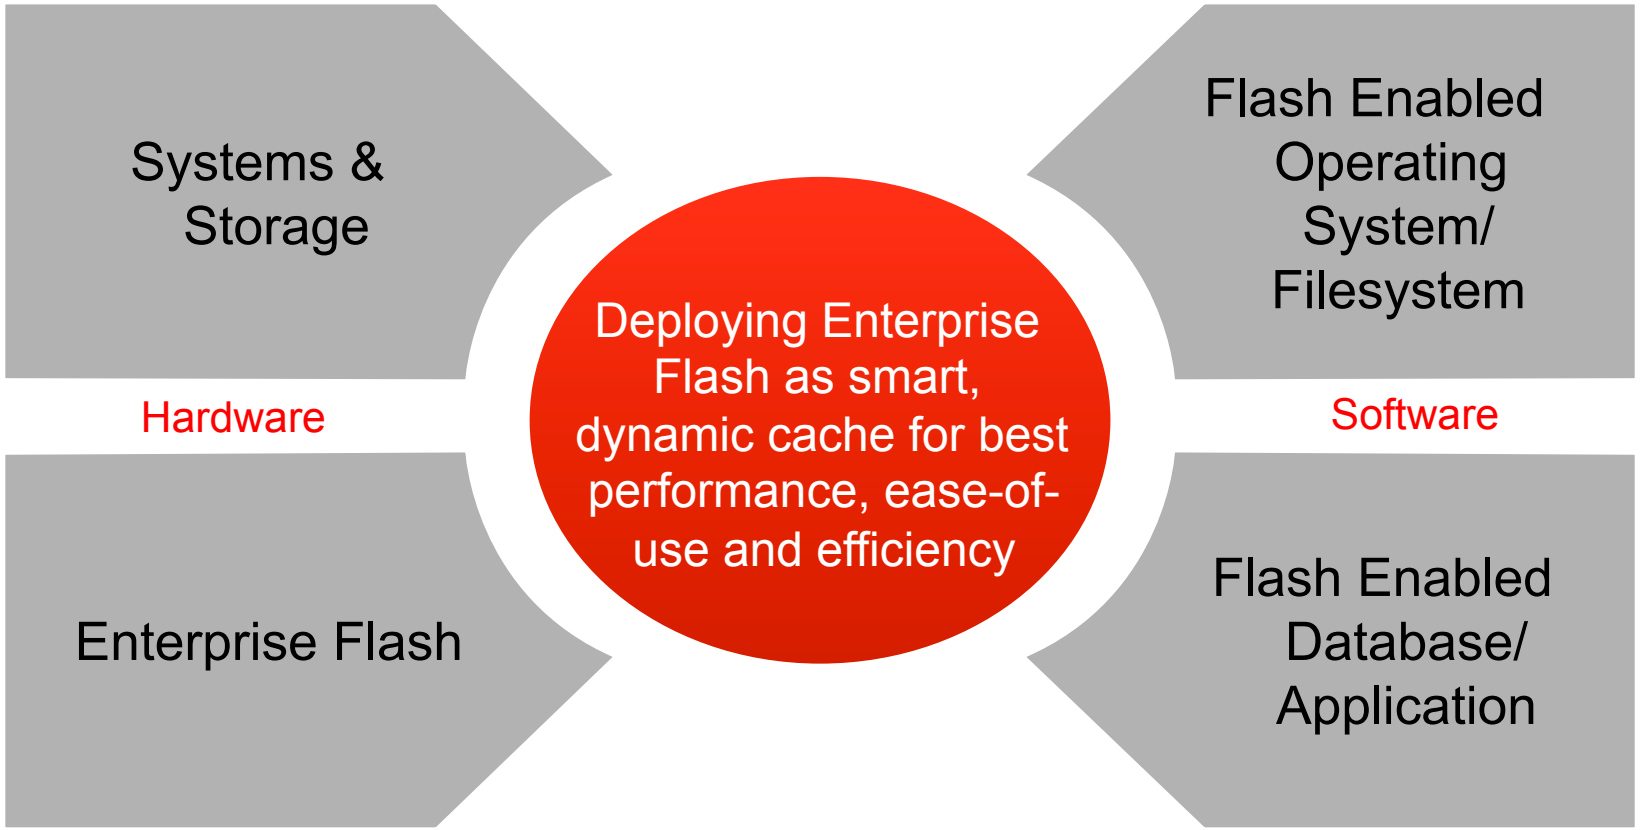

Flash is the biggest hardware breakthrough of the decade

Agenda

- ExaData & ZFS Storage Appliance: Architected with Flash
- 2nd gen Flash Characteristics Impacts
- Future Solid State Deployment

ExaData – Architected with Flash

Scale-Out Storage	No bottlenecks to scaling flash I/O
InfiniBand	Highest throughput, lowest latency
Intelligent Storage	Key to using full flash bandwidth Even InfiniBand can't send 50GB/sec
PCI Flash	Avoids disk controller bottlenecks. Cards in storage enable HA, RAC
Compression	Multiply flash capacity 10x Also multiplies data scan rates
Smart Flash Cache	Speed of flash, cost of disk Optionally specify table placement

Oracle Database Smart Flash Cache

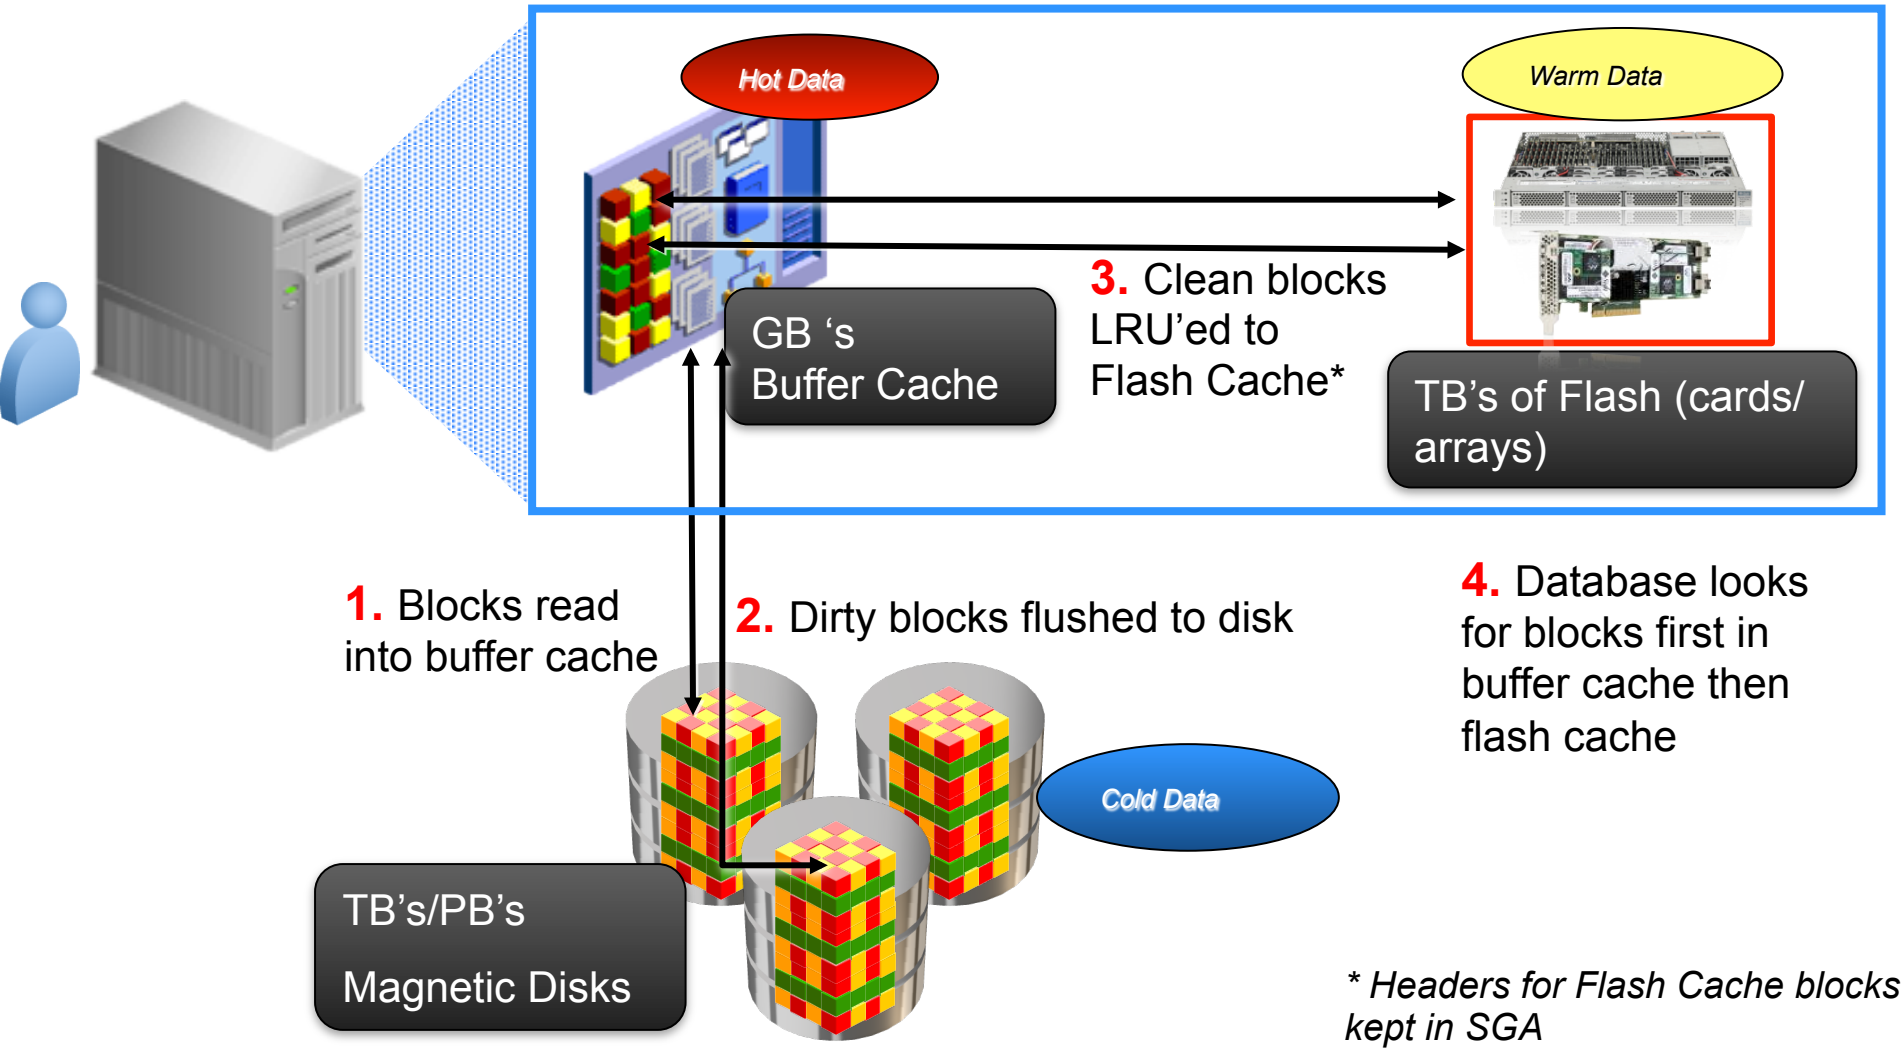

** Headers for Flash Cache blocks kept in SGA*

ExaData Performance

75 GB/sec!

For DataWarehouse workloads:

- Fastest Disk Throughput
- Much Faster with Flash
- Hybrid Columnar Compression (HCC) provides additional benefit

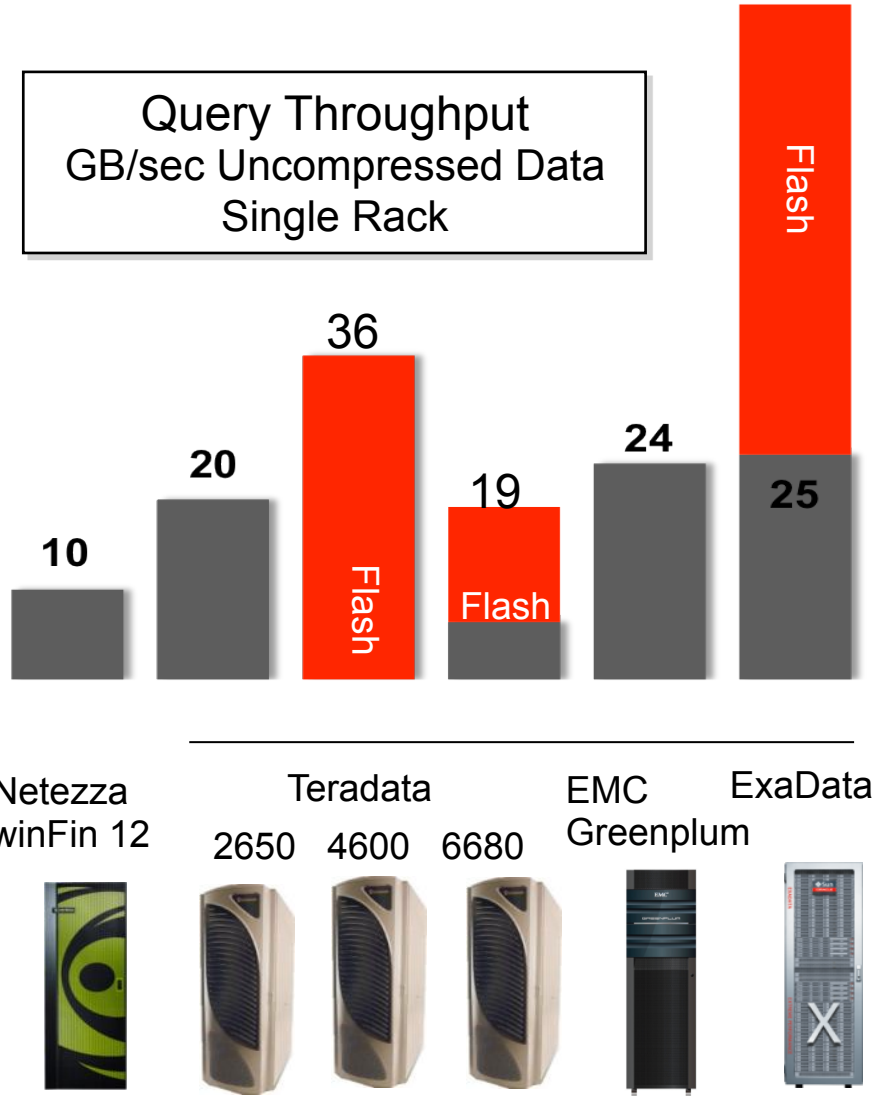

ZFS Storage Appliance – Architected with Flash

ZFSSA 7420 System

2TB DRAM total

80 cores total

5TB Read Flash total

1.7TB Write Flash total

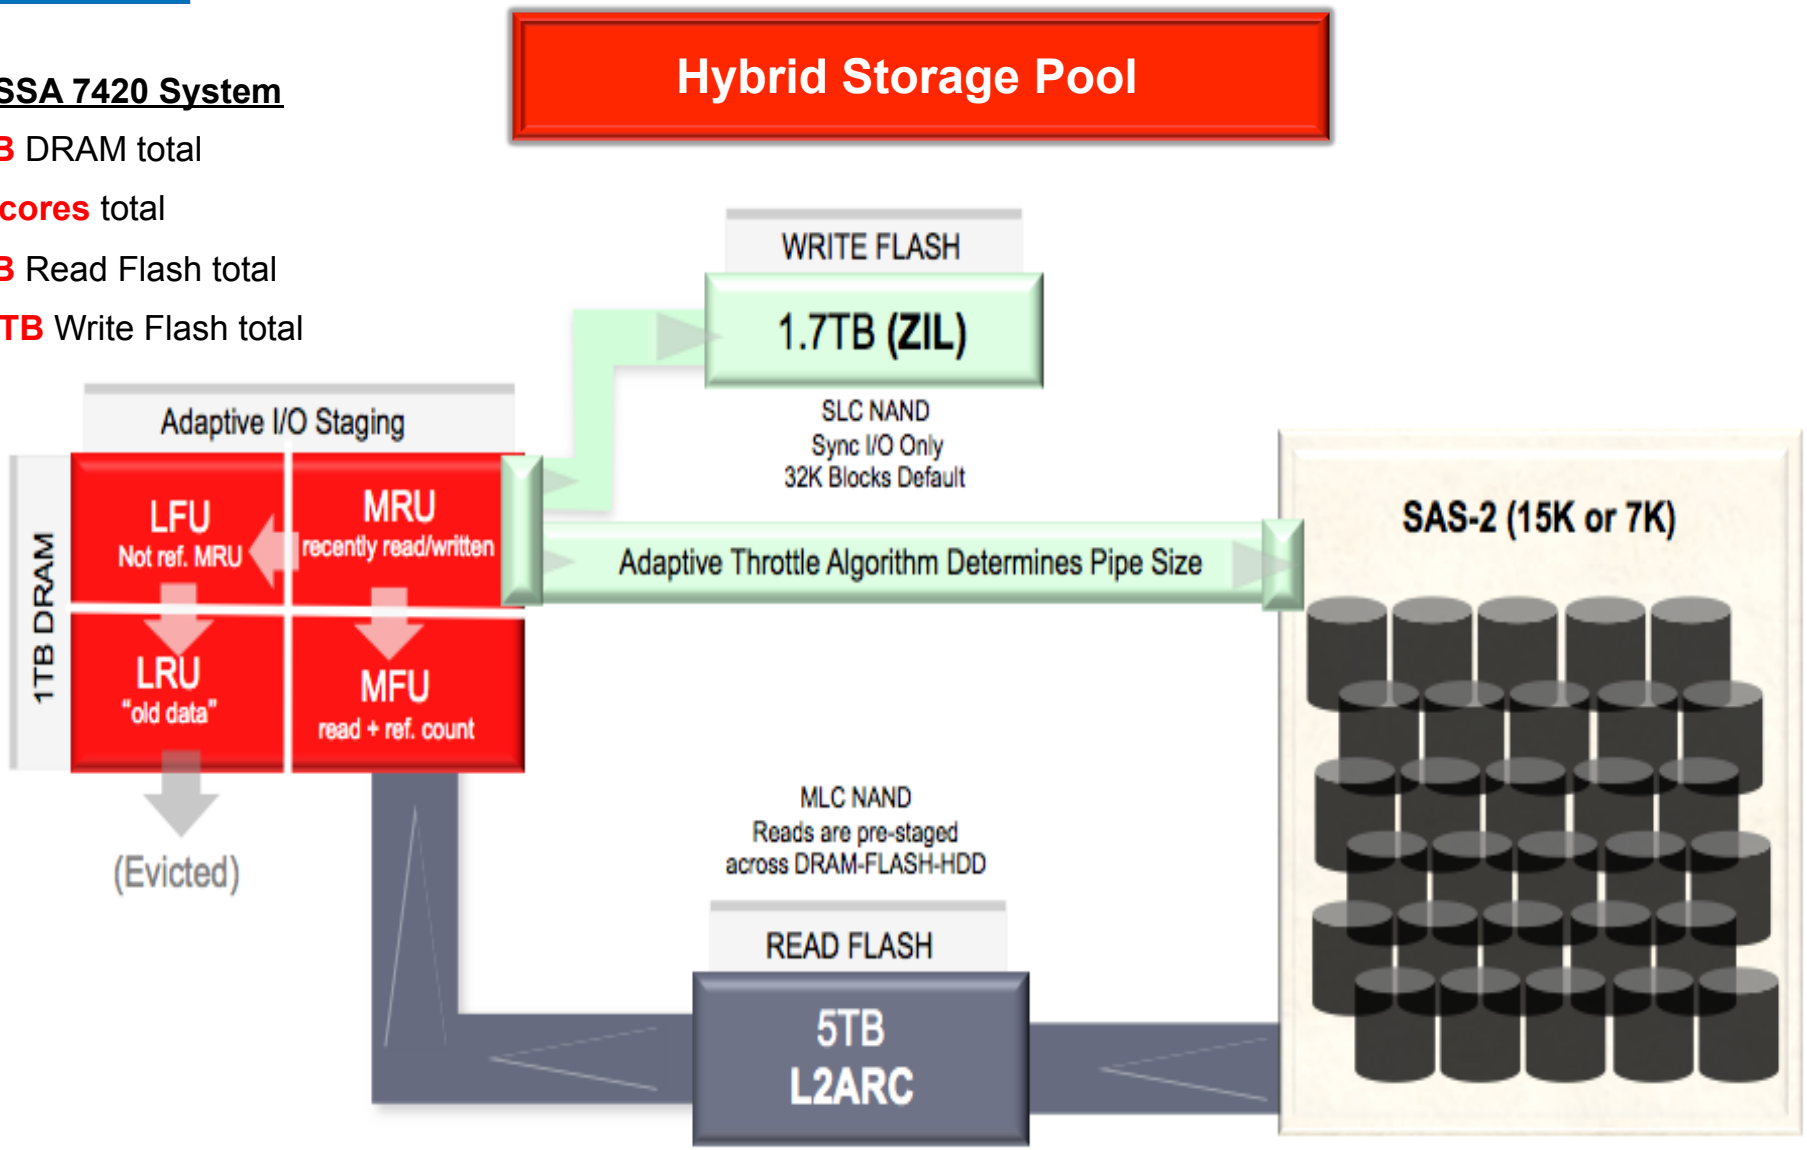

ZFSSA Performance

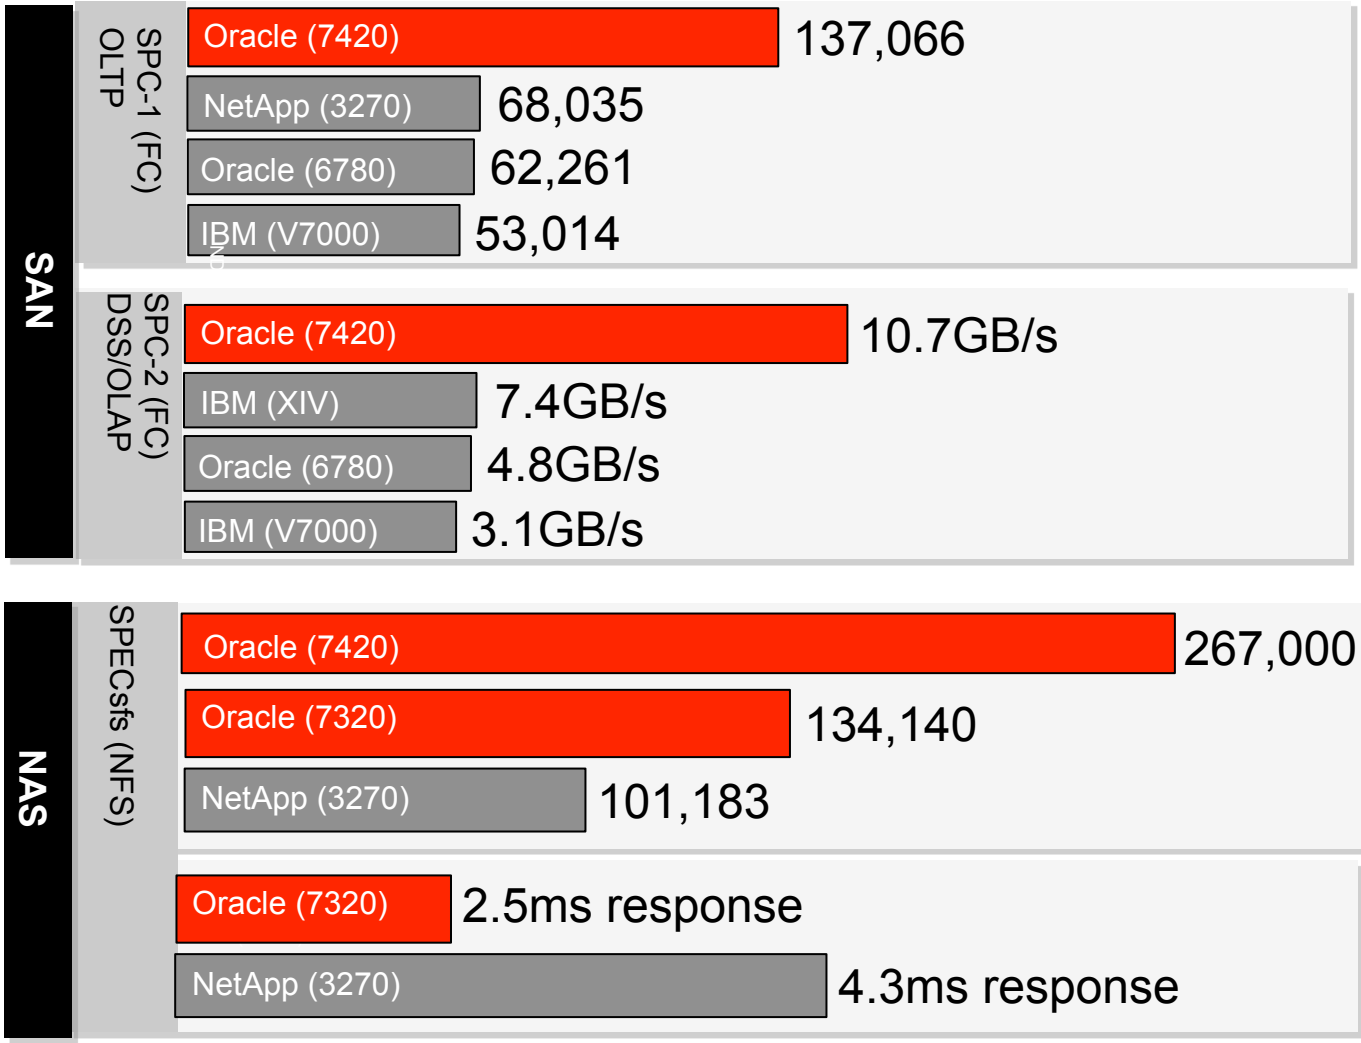

Only Unified Storage Box that performs well across ALL industry standard benchmarks

Sources:
SAN: storageperformance.org
NAS: spec.org/sfs2008/

Flash Performance Inconsistencies

No true form/fit/function replacements…yet!

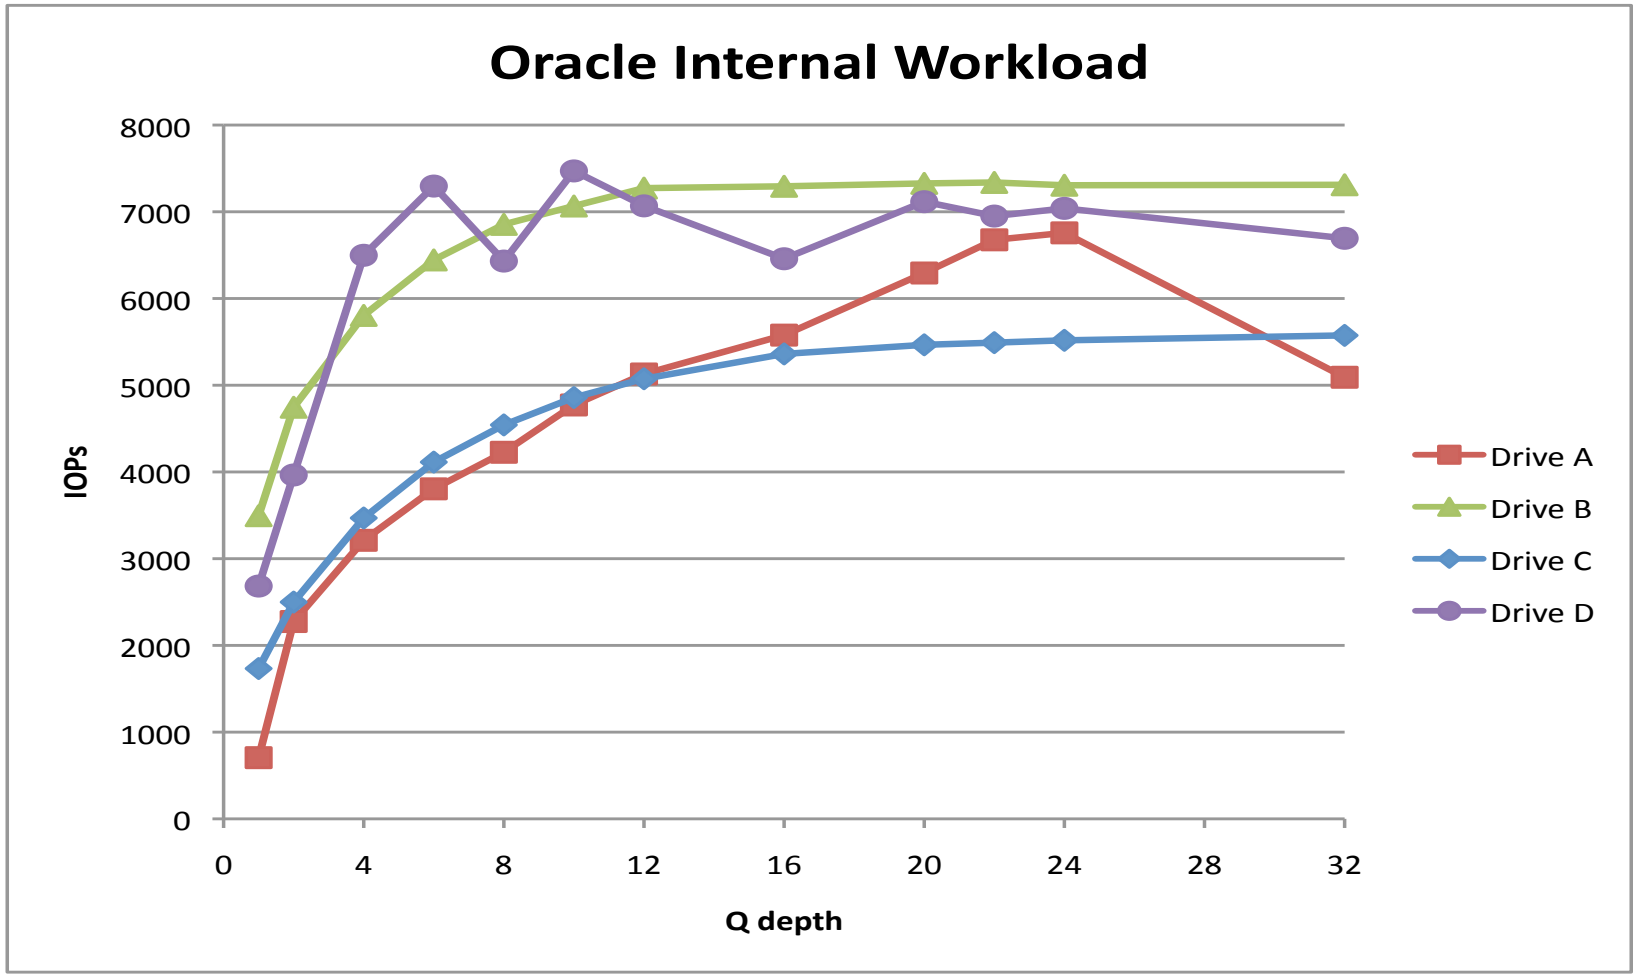

Enterprise Flash Needs

- Reduce the Cost
  - Too high for broad adoption
- Consistent Performance
  - Inconsistent with NAND – garbage collection and other FTL events
  - Consistent with DRAM – BUT there are energy-source requirements
- Consistent Form-Fit-Function
  - SSD's are not fungible, all have to be re-characterized (Performance, Endurance, etc.)
  - Some SSD's do not comply with existing industry standards
  - Lack of standardization of management and protocol interfaces
    - ✓ NVMe shows promise

Enterprise Flash Needs (cont'd)

- Reliability, Availability & Serviceability
  - NAND flash scaling pace continues to be a challenge
  - Data Integrity
  - Management interfaces
  - Robust storage protocols
  - Dual connectivity
  - 2M hours MTBF (24X7) for > 3 years
- Larger Capacities
  - Flash still too small for general purpose storage

Why Oracle is Building the ZFSSA Accelerator

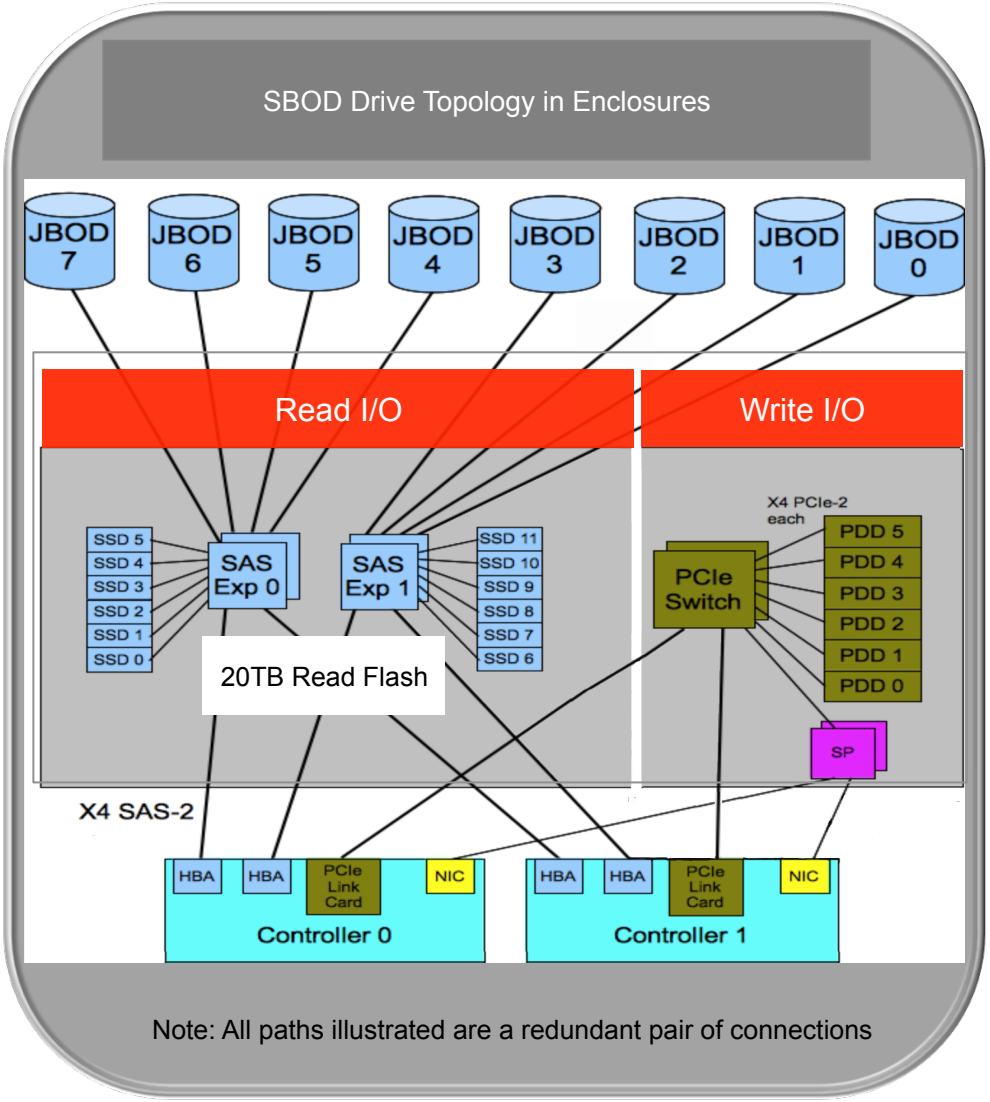

Flash is *inconsistent*

- Take on the energy source requirements

Want the *fastest* protocol

- Low latency performance is key
- NVMe wasn't as well defined as it is today

Need Enterprise RAS

- PCI lacks advanced *error handling* – NTB
- *Enterprise* Connectivity – Dual Root PCIe
- Servicability – PDD's are FRU's, etc.

External Attach

- Keeps ZFSSA on cutting edge technologies

Final Thoughts

- When does NAND run out of the ability to scale?
  - Enterprises require significant time for an orderly transition
- Future NVMs
  - Learnings from NAND apply and look more promising (on paper)
- Standardization and Cost reduction keys to adoption
  - Enterprise storage must have smooth transitions
  - Involvement in SNIA
  - Involvement in NVMe
  - Involvement in PCIe SFF

Oracle.com/storage

Hardware and Software

Engineered to Work Together