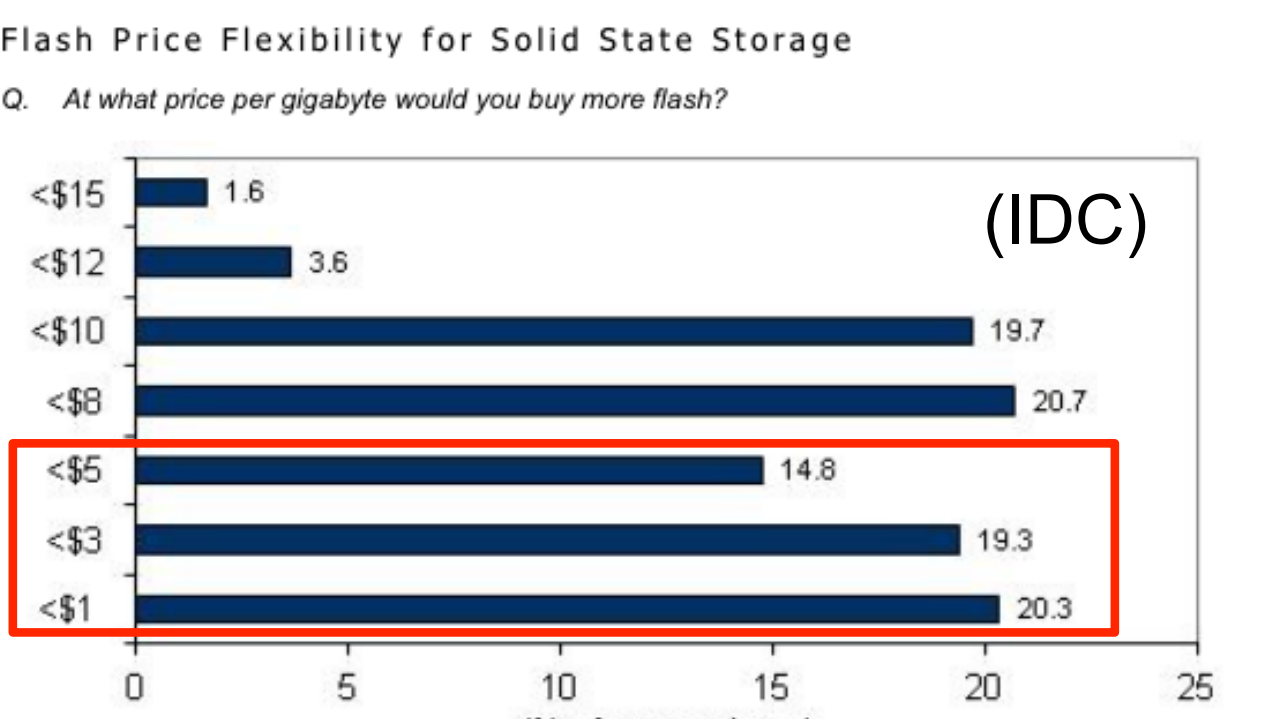
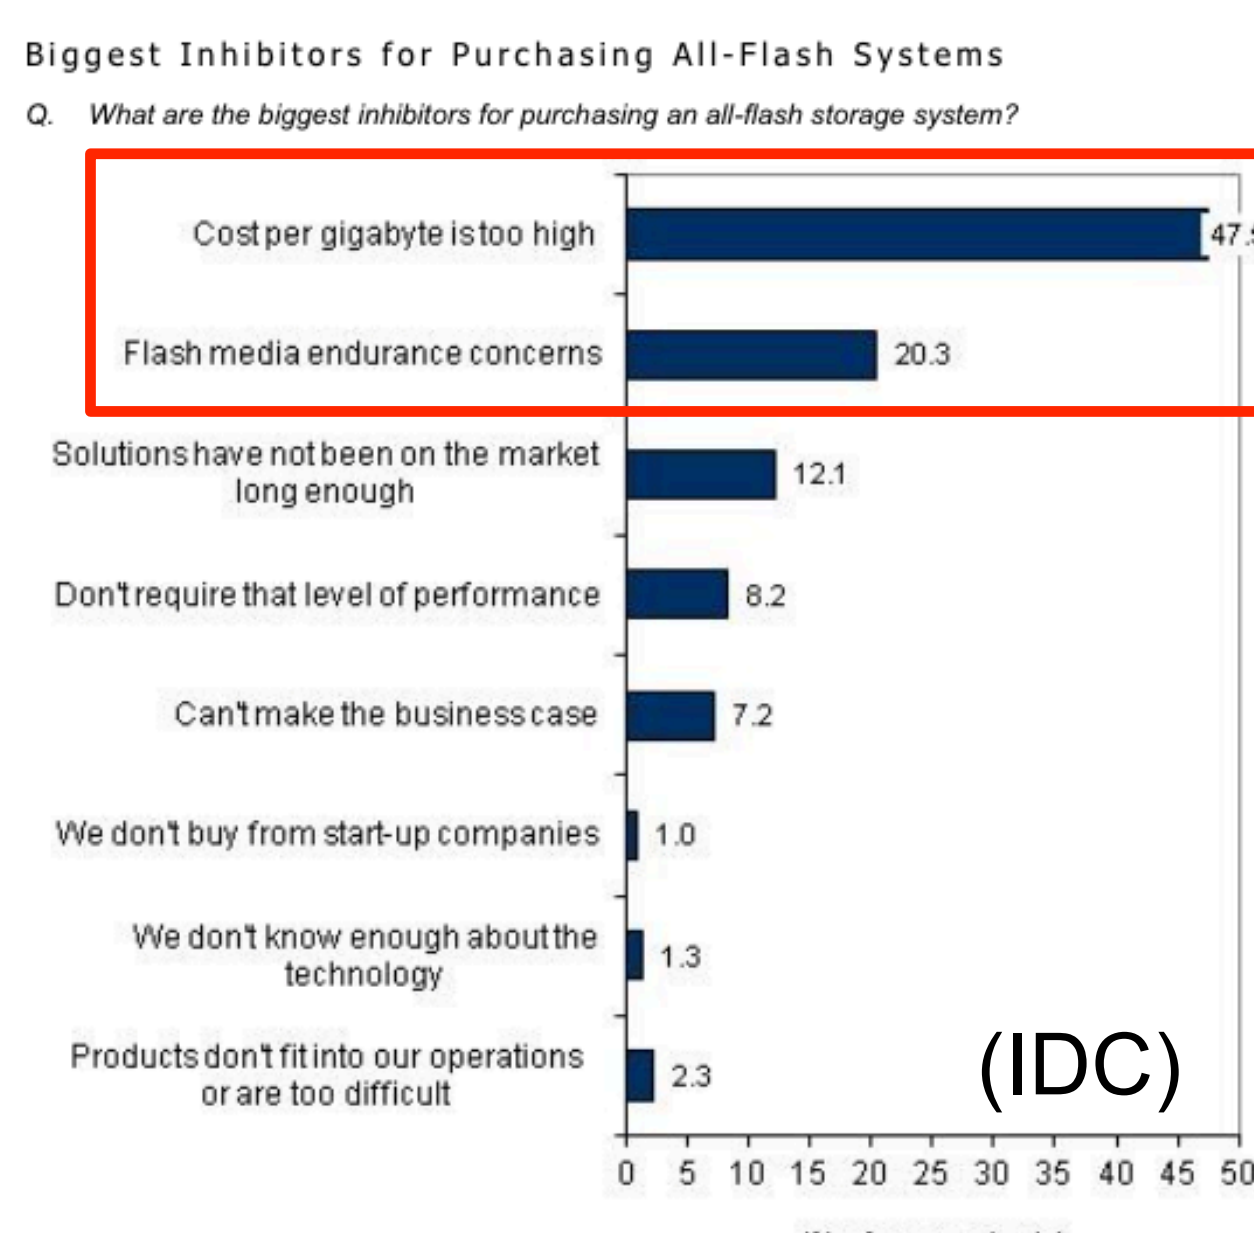


SALSA: treating the weaknesses of low-cost Flash in software

Nikolas Ioannou, Ioannis Koltsidas, Roman Pletka, Sasa Tomic, Radu Stoica, Thomas Weigold, Evangelos Eleftheriou
 {nio, iko, rap, sat, rst, twe, ele}@zurich.ibm.com
 IBM Research - Zurich

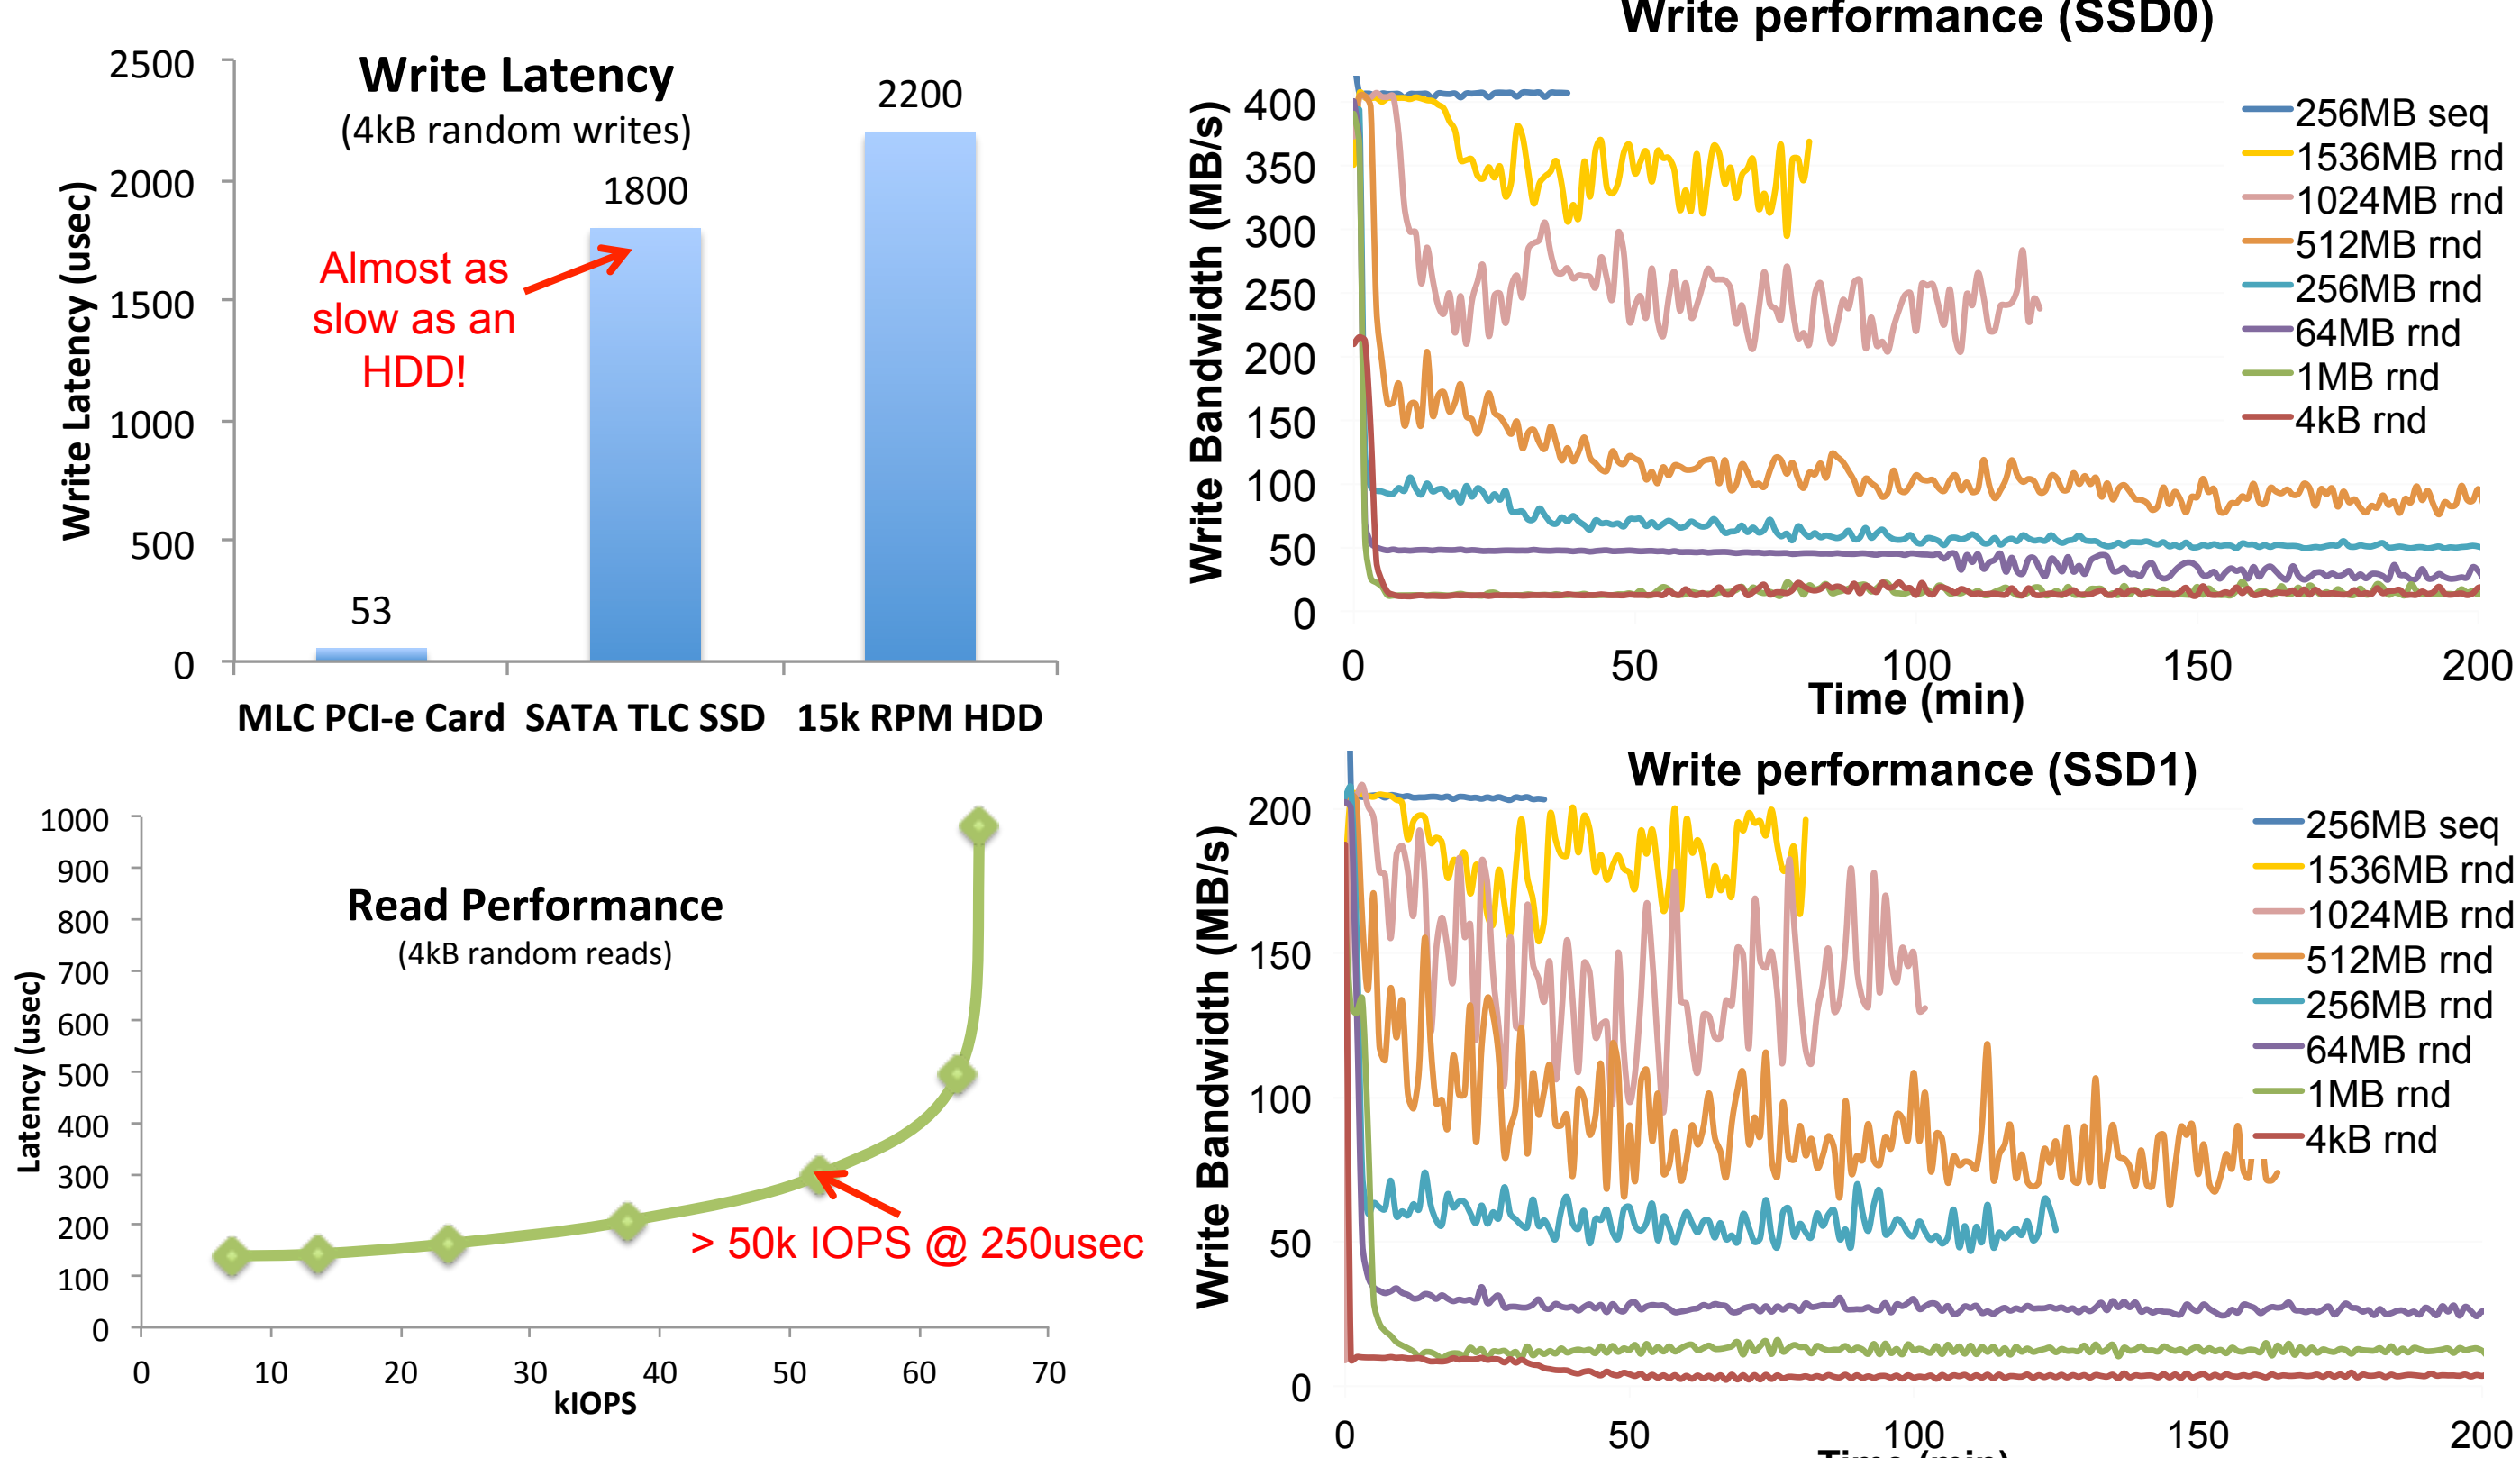
Low-cost Flash in high demand



Also according to IDC:

- Over 67% are looking into taking a **self-integrated approach** to deploying Flash in their servers
- Over 77% are looking into or are in the process of adopting flash storage in the **Cloud**

Low-cost SSD performance



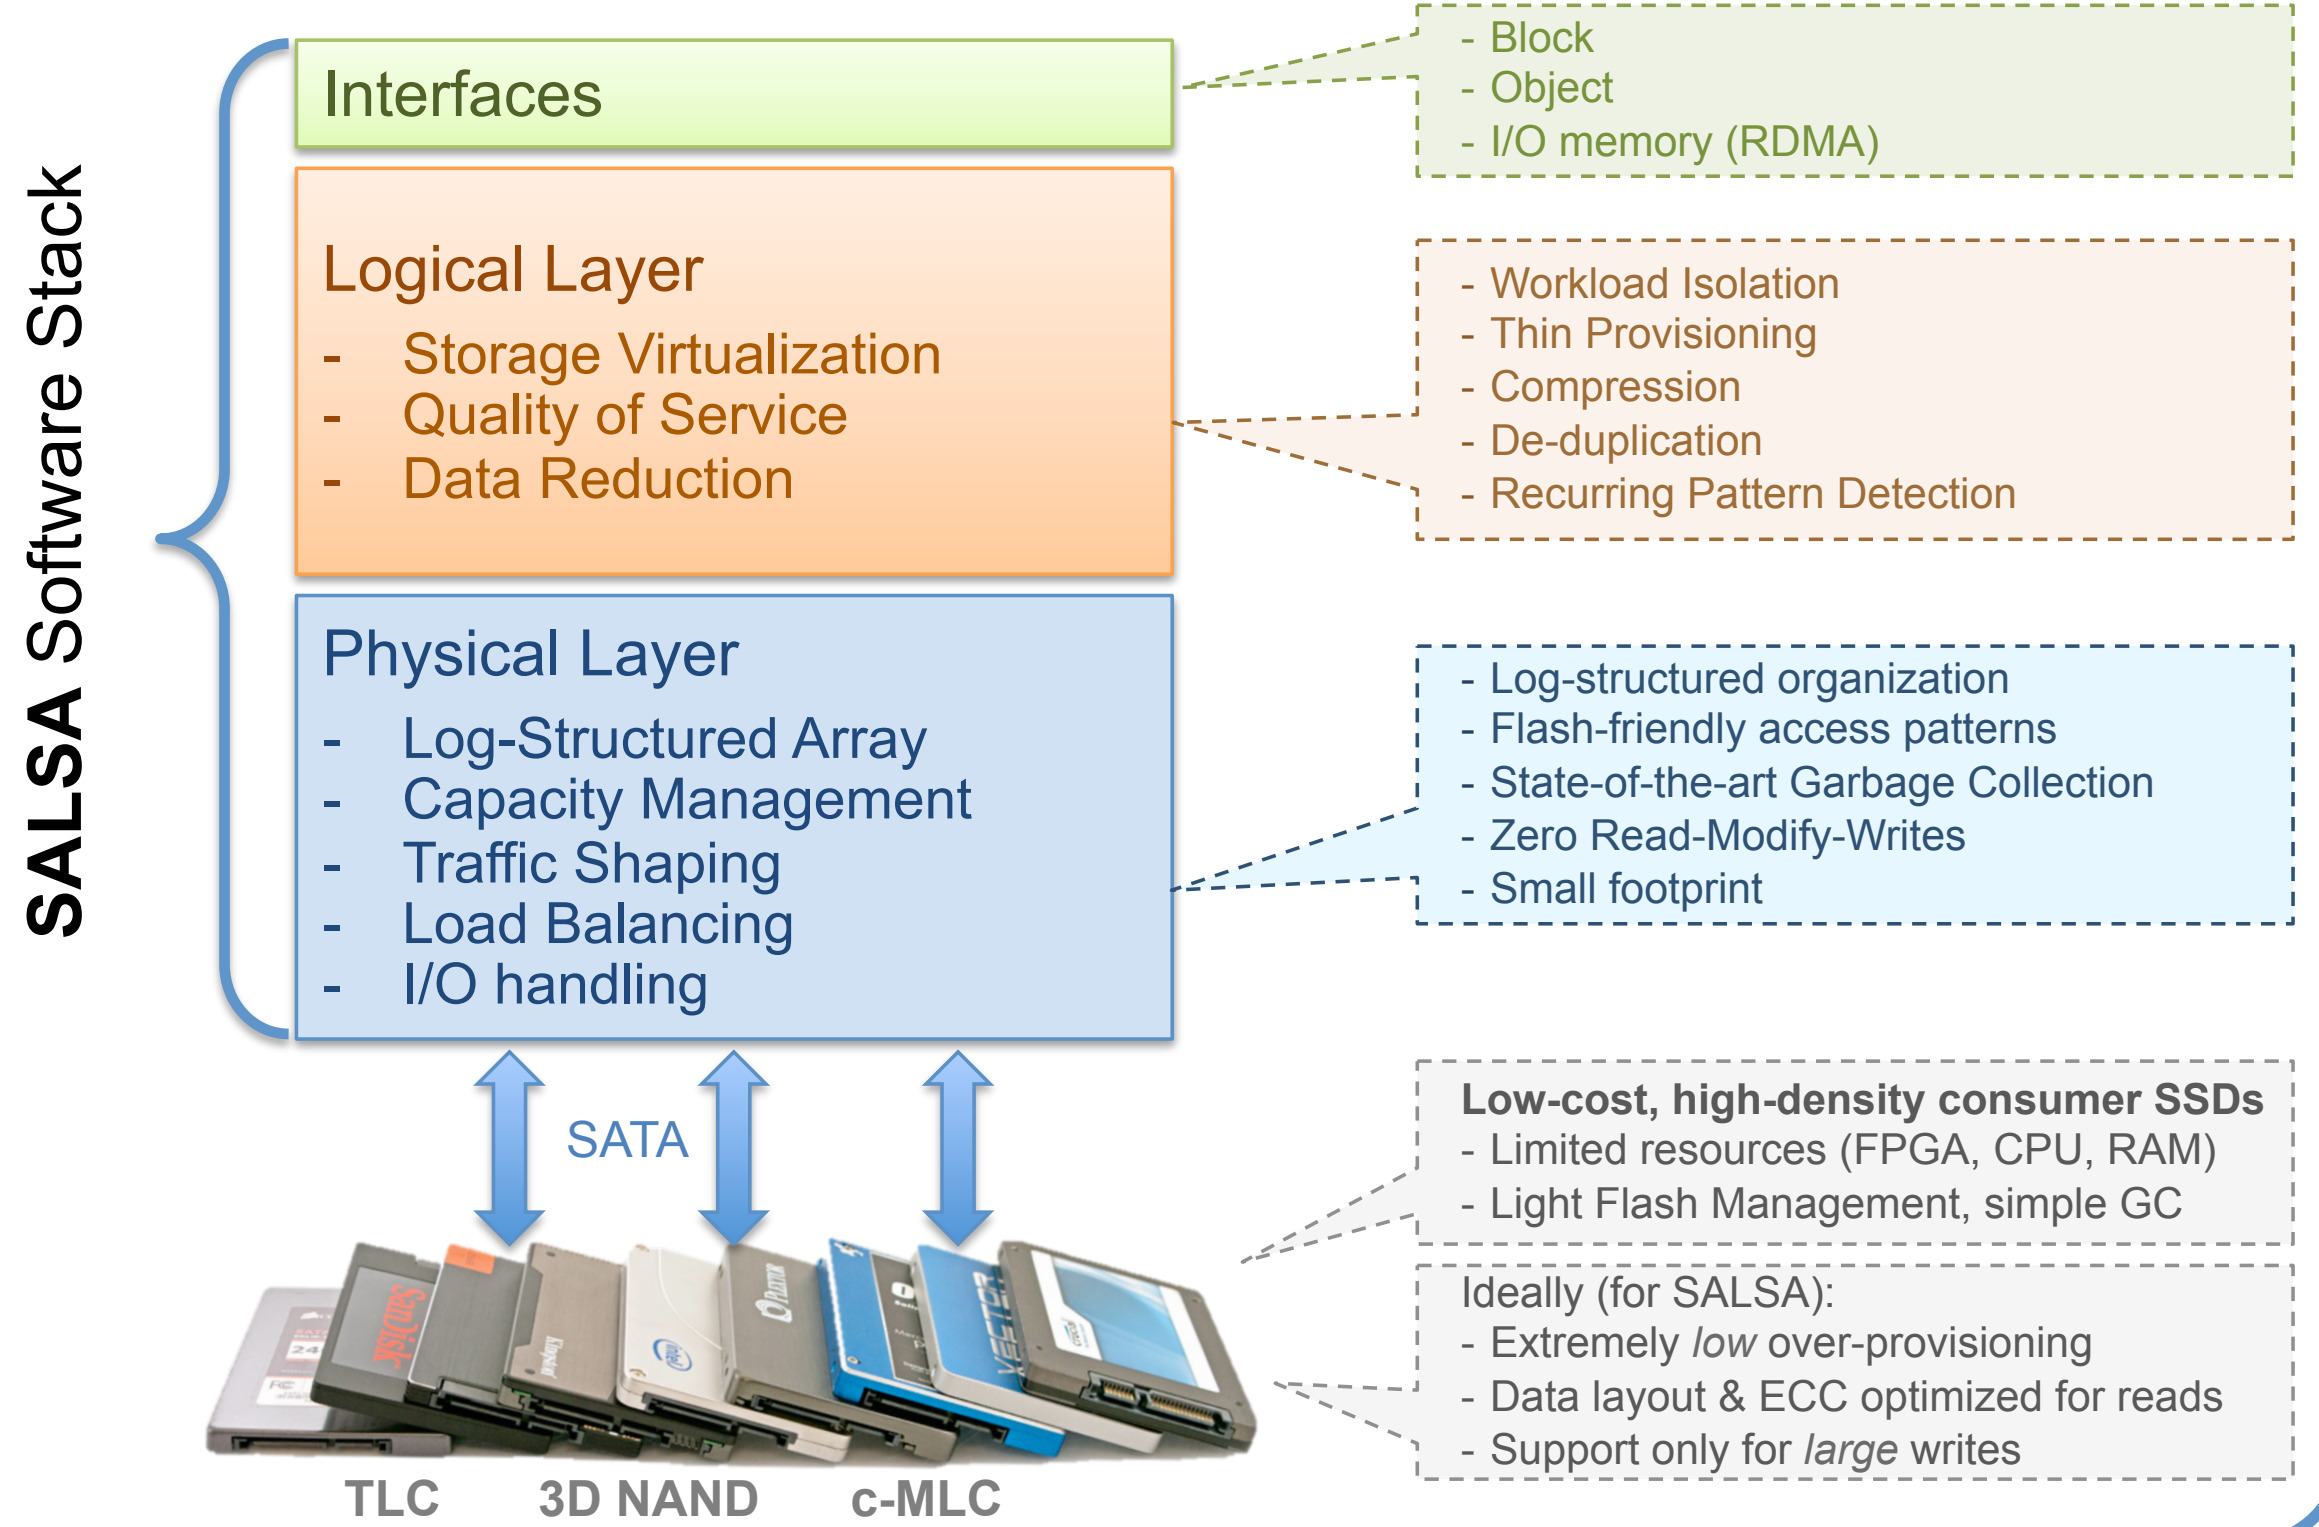
- SSD0: Off-the-shelf SATA SSD, TLC Flash, 1TB, 2.5" drive
 - SSD1: Off-the-shelf SATA SSD, TLC Flash, 480GB, 2.5" drive
 - Price per GB: <\$0.5

Cloud workloads

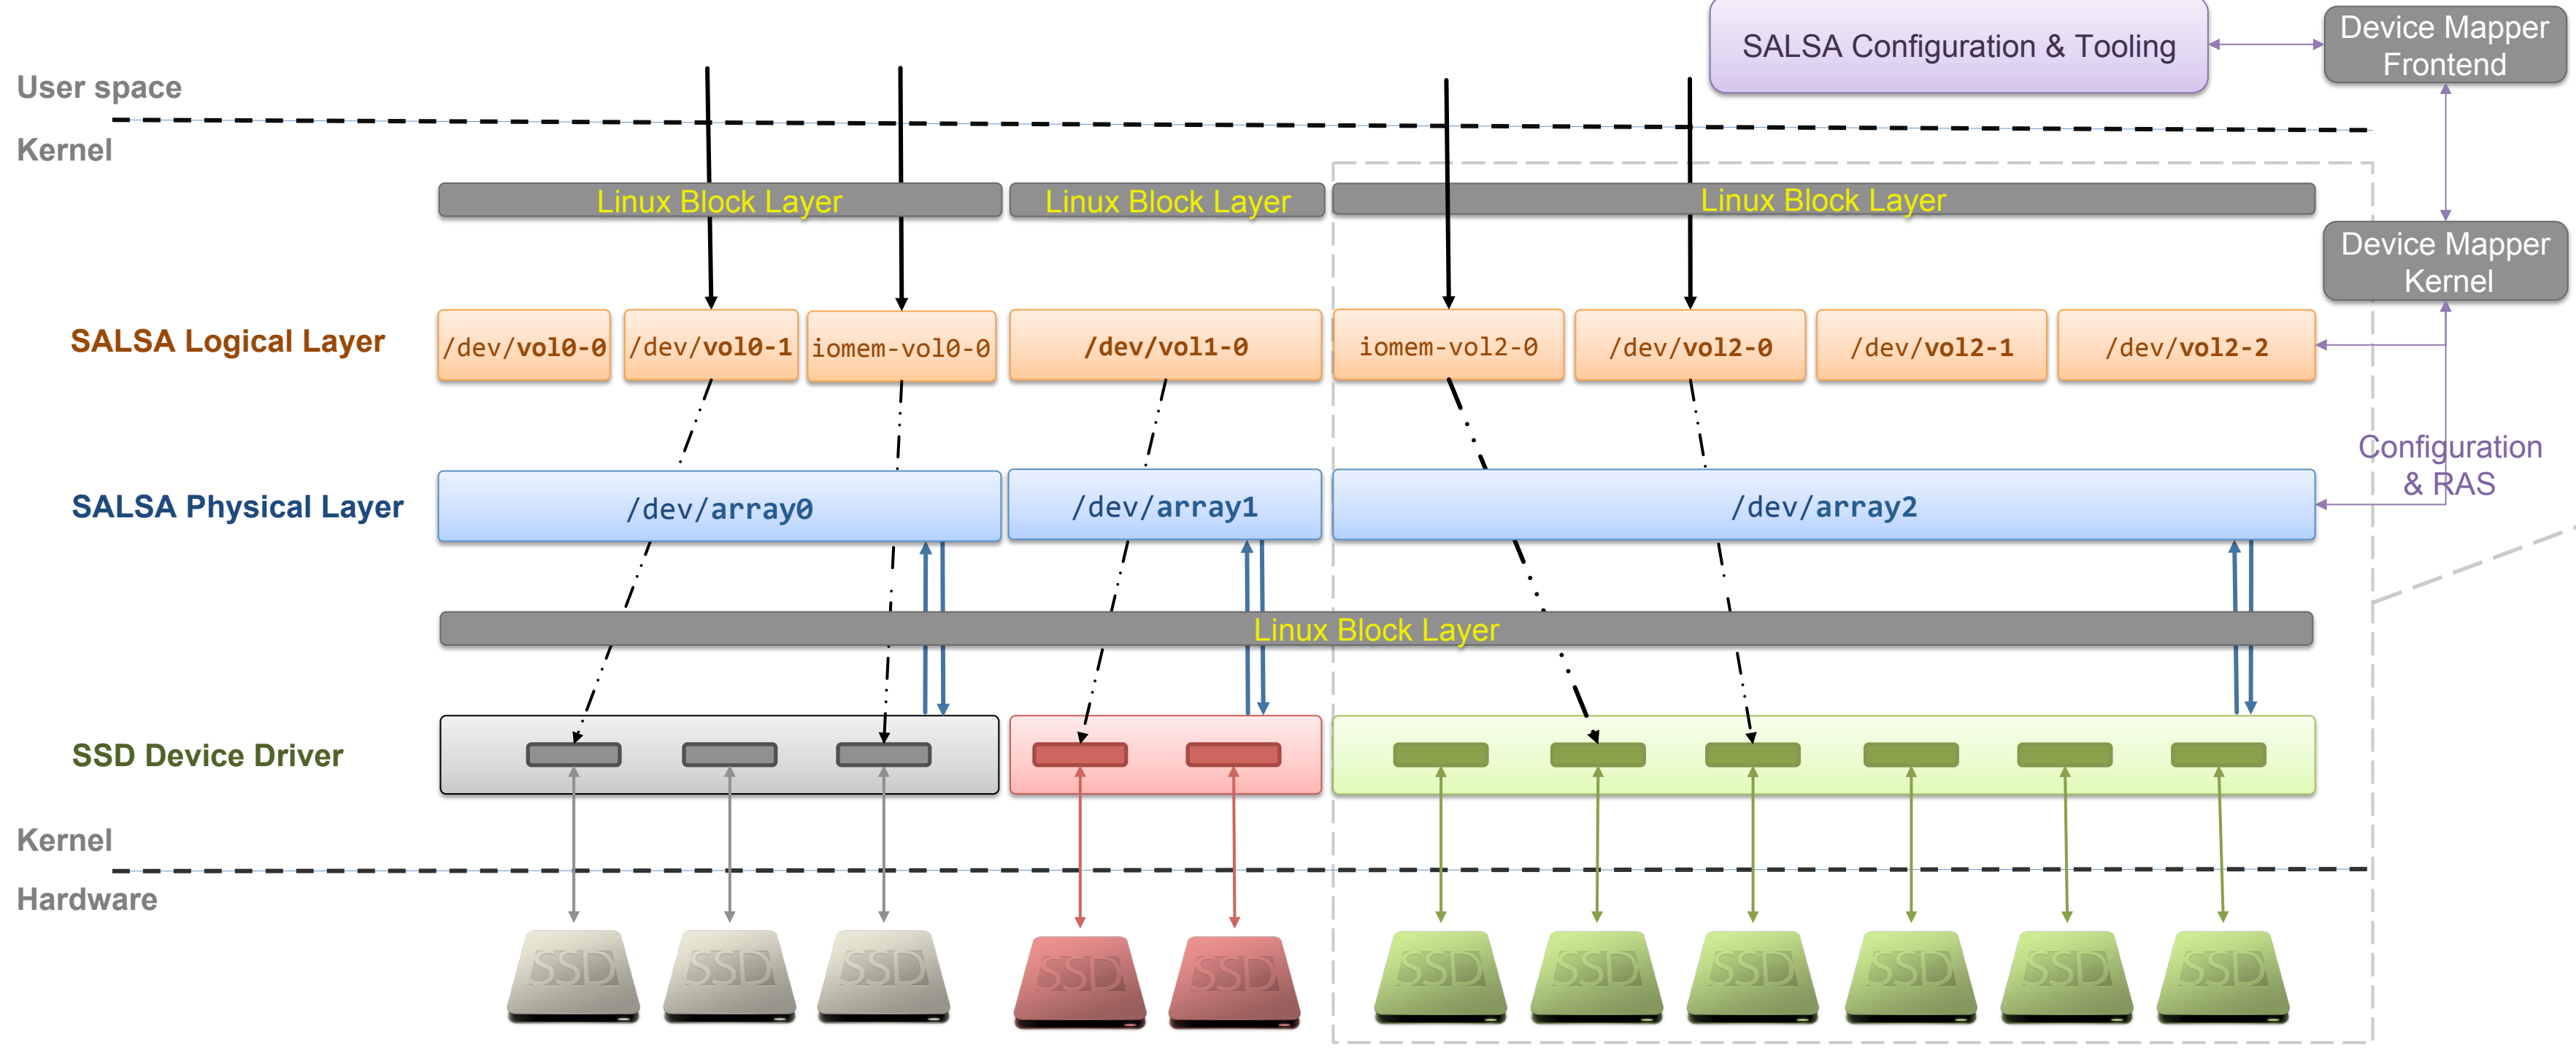
- Many write-once-read-many workloads
 - Data is often immutable
 - Don't need the write endurance of high-end Flash
 - Examples:
 - Data Warehousing / Analytics
 - Active Archives
 - Social
- Focus on density, cost, and read performance:
 - Read performance: high IOPS & bandwidth, low latency
 - Low cost: consumer-level
 - Enterprise-grade performance and high availability
 - High data ingest rate that is non-disruptive to Reads
- Existing I/O stacks and architectures are not a good fit
 - A Workload-Optimized solution is needed

Software Log-Structured Array

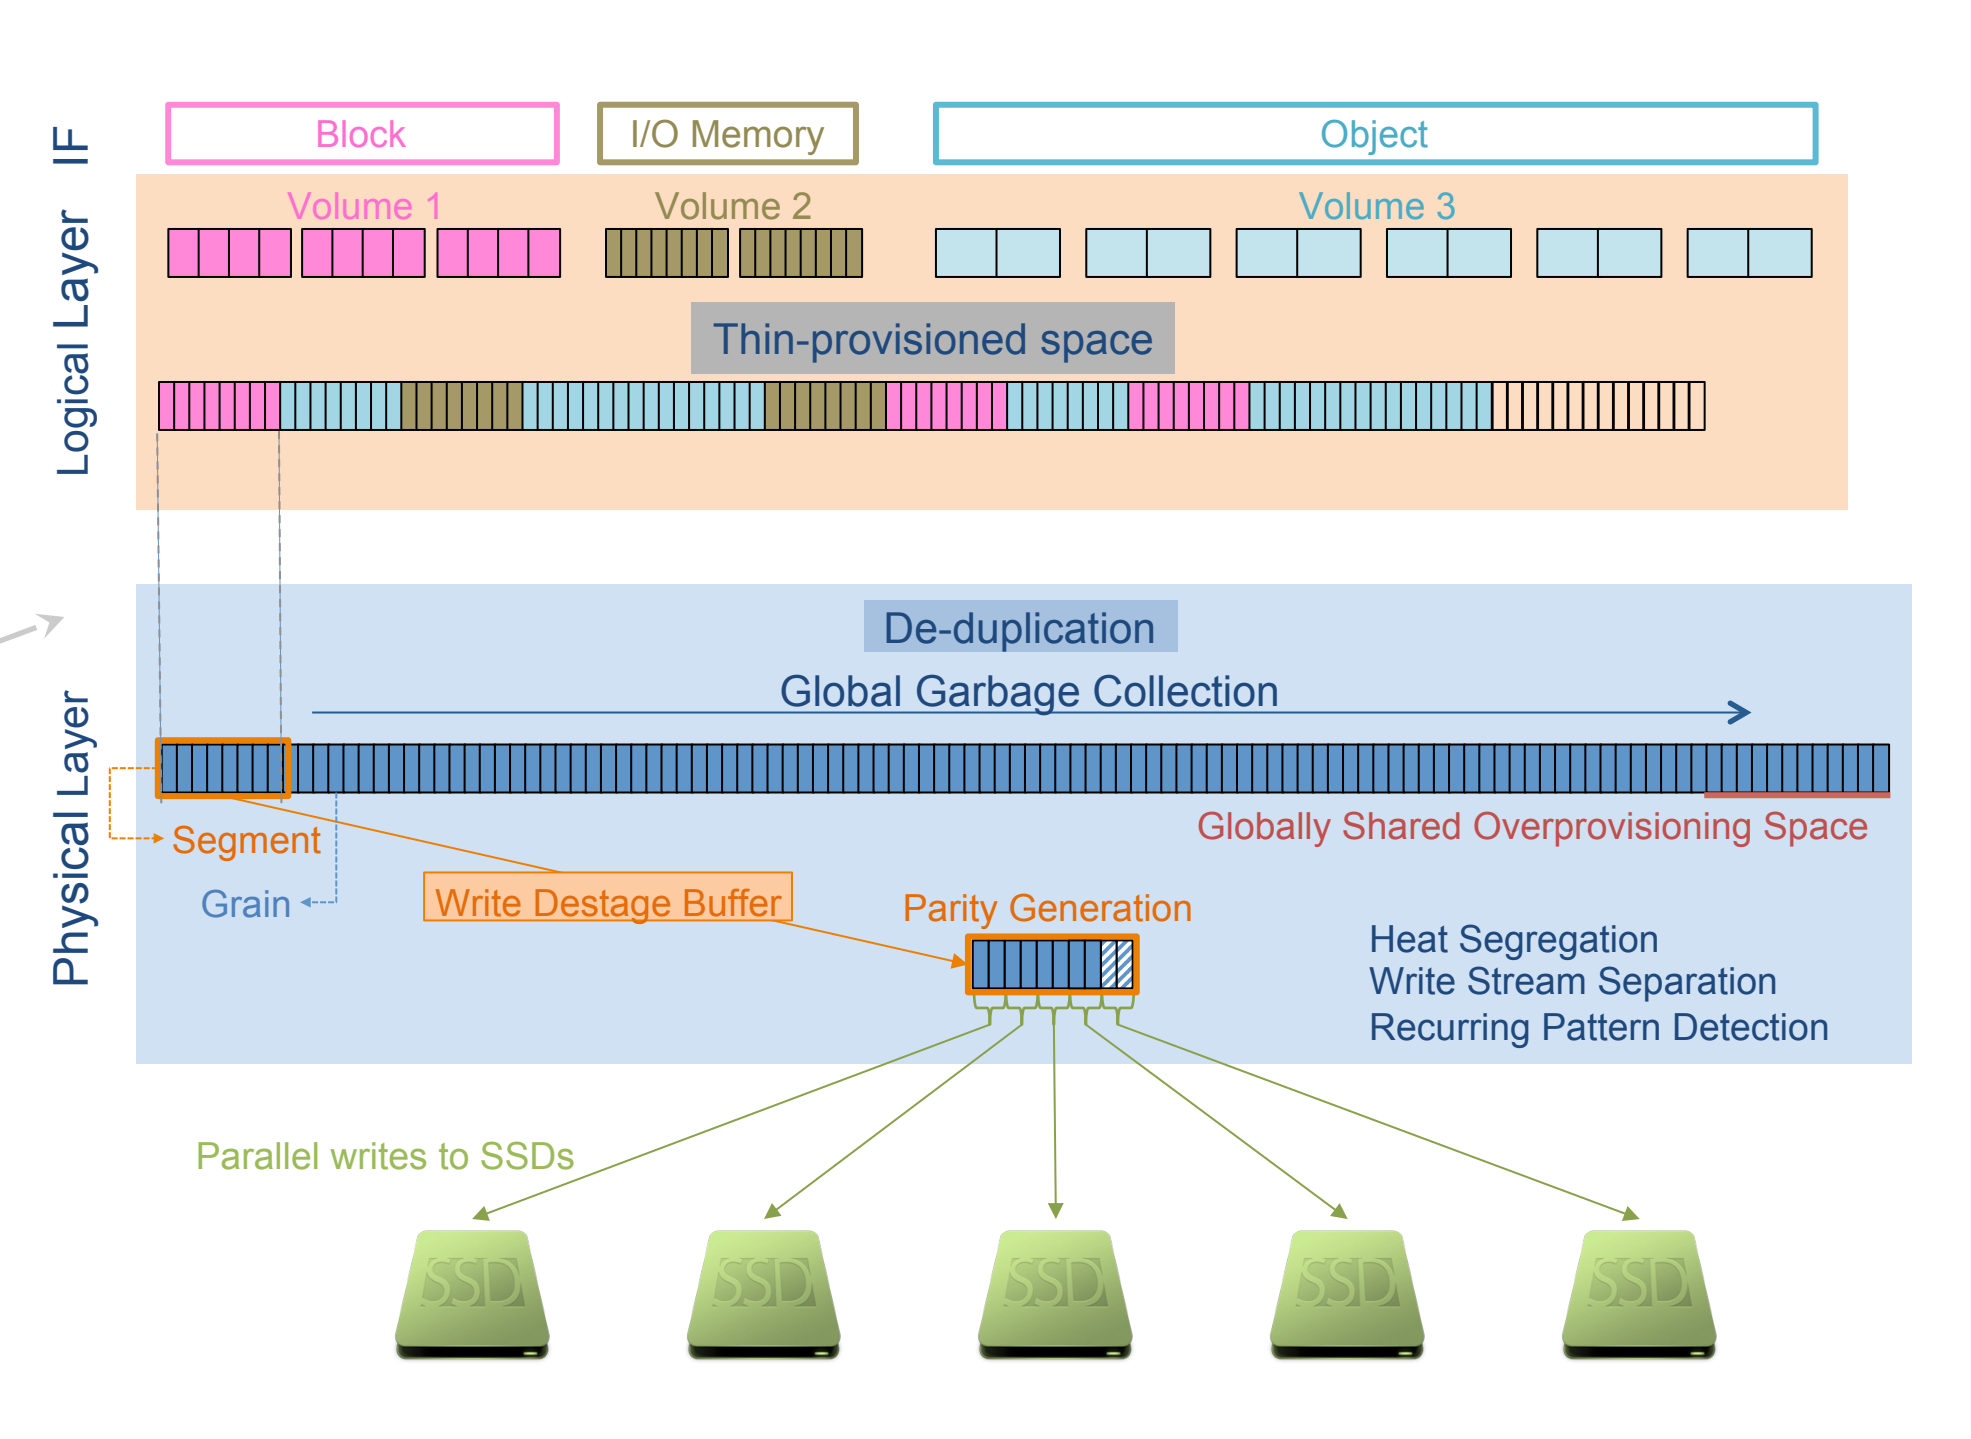
- Key Ideas:**
- Use high-density, low-endurance Flash
 - Move complexity from hardware to software
 - Optimize end-to-end for low Write Amplification and data durability
 - Employ data reduction to further reduce cost per GB
 - Natively support Object Storage



SALSA Data Path




SALSA Virtual Array



SALSA elevates low-cost Flash to meet Cloud storage requirements

Preliminary performance evaluation



Conclusions

- Use commodity Flash in the cloud
- Shift complexity from hardware to software
- SALSA: a storage virtualization stack for Flash
 - Workaround FTL unpredictability
 - Elevate the performance of commodity Flash
 - Optimize end-to-end for low Write Amplification
 - Parity protection without the RMW penalty
 - Stream separation and heat segregation
 - Workload consolidation and QOS
- Next steps:
 - Integrate with Open-Channel SSDs
 - Automate profiling and configuration for SSDs