

DPU + CSD

Build Secure and Dependable Data Storage at Lower Cost

Keith McKay

ScaleFlux



DPU

Data Processing Unit

Storage virtualization and data
management offload

Computational Storage Drive

Storage compute offload
(e.g., filter, compress)

Consumable Computational Storage

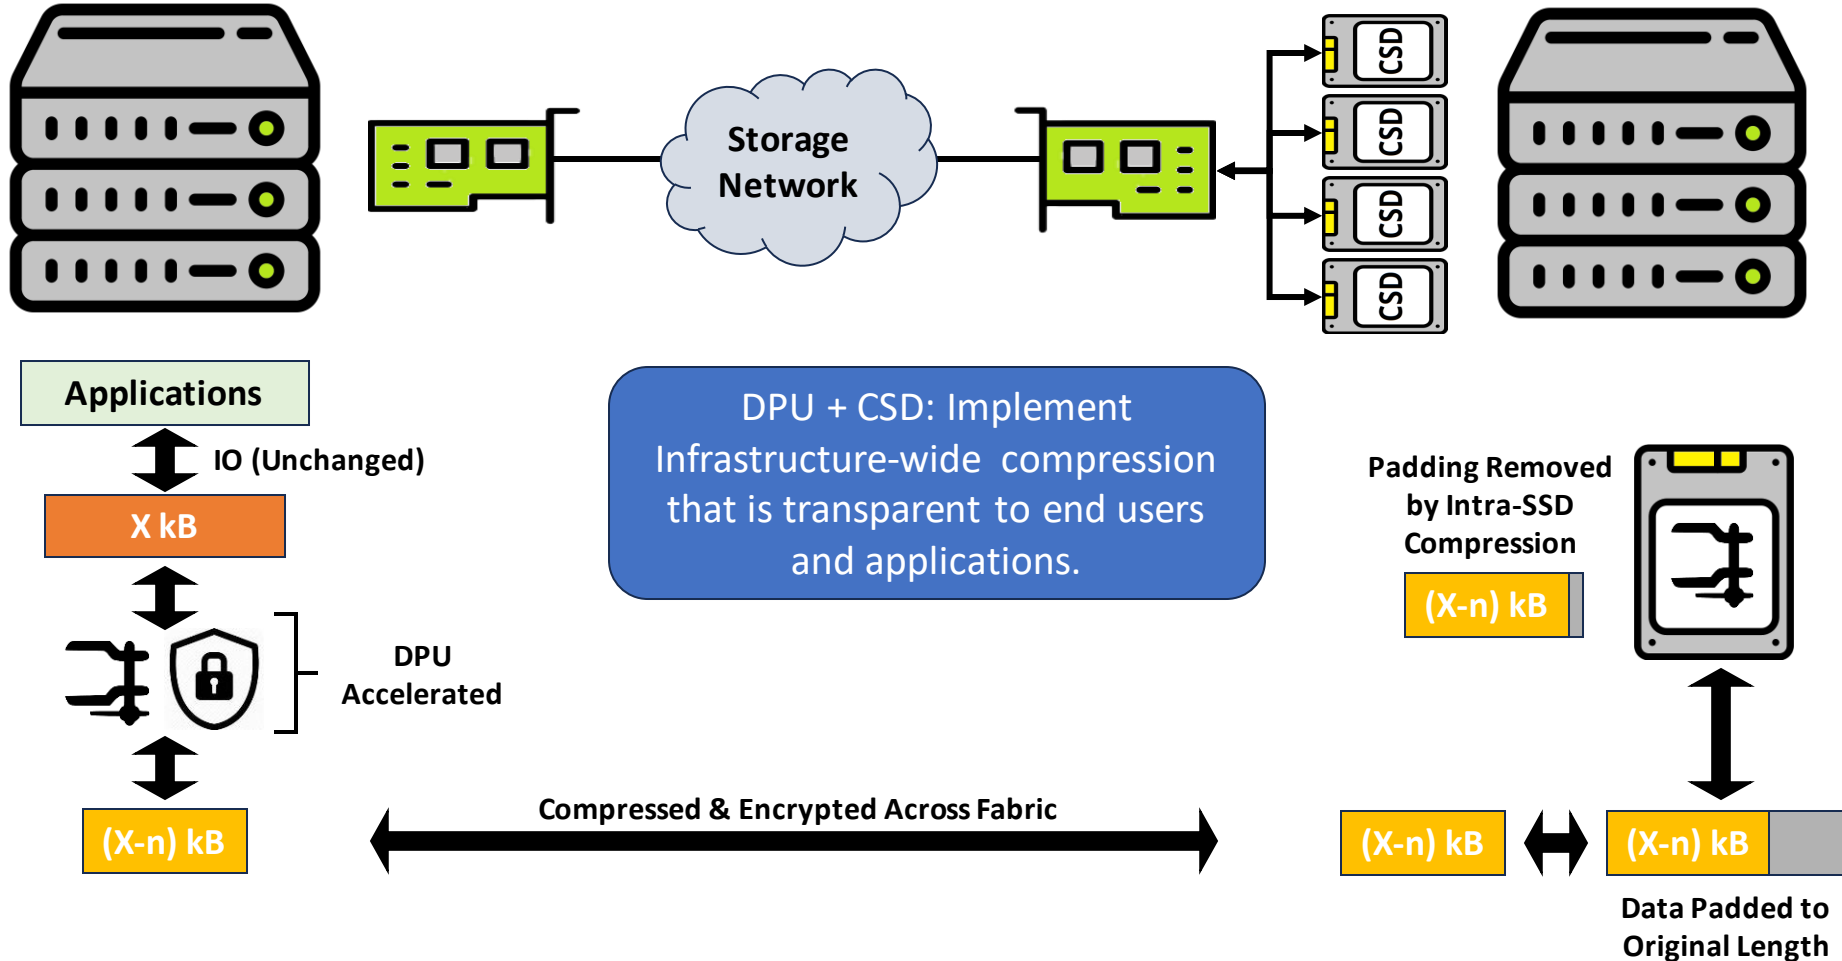
CSD virtualization, computational
storage management offload, storage
compute offload

DPU + CSD: Consistent abstractions to the host with offloaded computational storage management.


Application: Disaggregated Transparent Compression

Compute Servers (ECS)

Block Storage Servers (EBS)




Application: Elastic RAID



Best Performance, High Storage Overhead, Low CPU Overhead

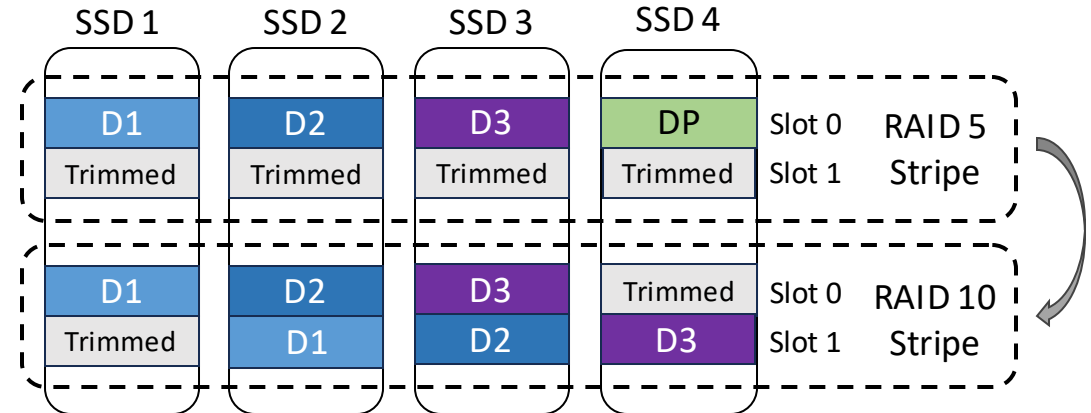
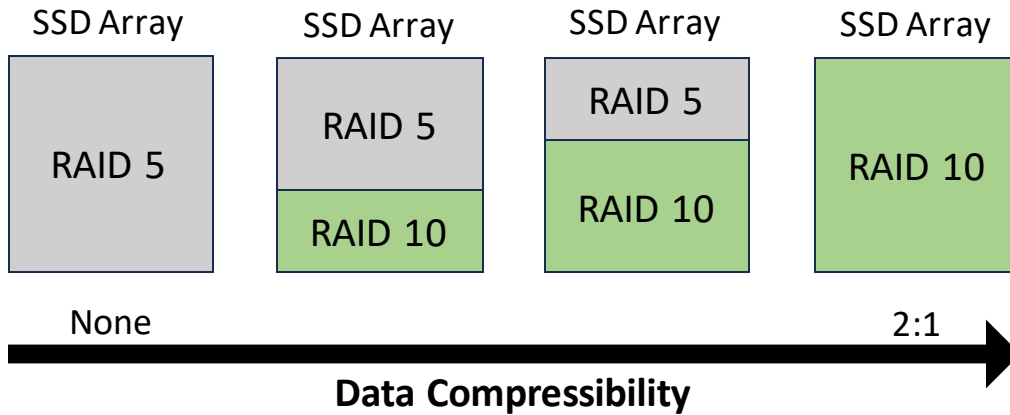
RAID 1(0) vs. RAID 5/6

**Locked into choice
at array creation!**

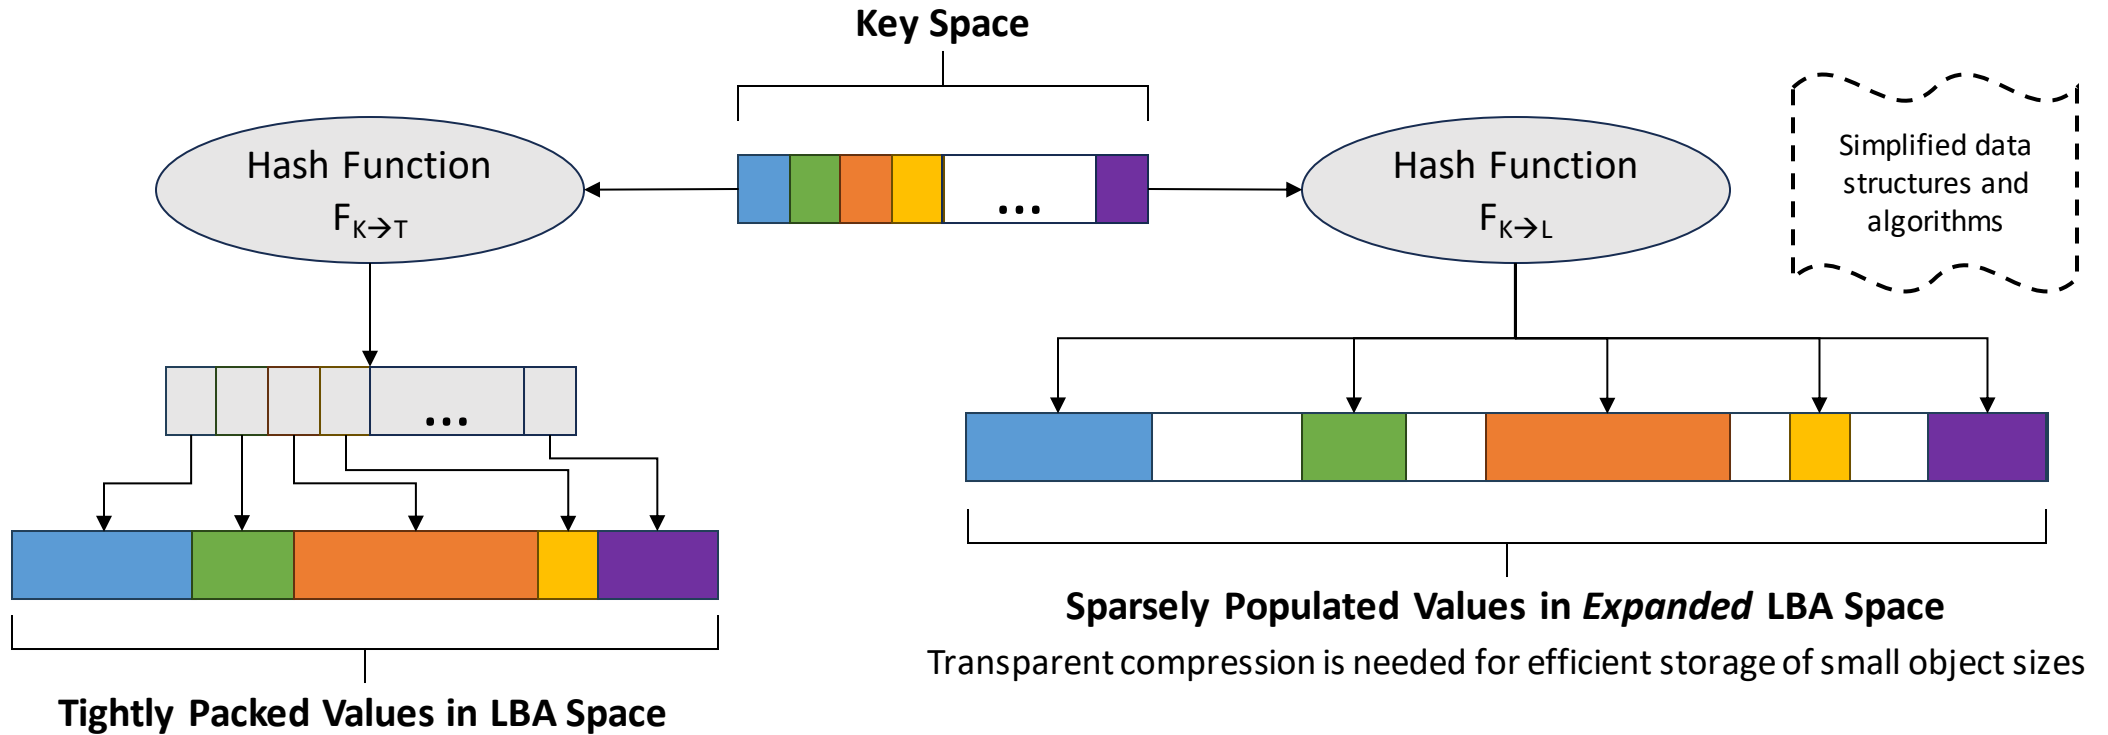


Good Performance, Lower Storage Overhead, Higher CPU Overhead

DPU + CSD: Exploit data compressibility to opportunistically convert from RAID5 to RAID10




Application: High Performance KV Store



DPU + CSD: Enable KV-Store Interface with Higher Throughput and Lower Latency

Key Take-Away



Leverage DPU to provide virtual computational storage devices.

Enable consistent host APIs by abstracting vendor-to-vendor variation in computational storage capabilities.

Orchestrate computational storage functionality across multiple CSDs.



Provide fundamental computational storage services such as transparent compression and data filtering.

Enable performance scaling by providing compute resources per-CSD.

DPU + CSD: Ideal combination to enable the computational storage ecosystem and realize the promise of near-data compute.

Thank You!

Resources

- Elastic RAID:
 - <https://dl.acm.org/doi/10.1145/3579370.3594773>
- Table-less KV Store:
 - <https://dl.acm.org/doi/10.1145/3465998.3466004>
- CSD and DPU: Friends or Foes
 - <https://scaleflux.com/library/computational-storage-drive-and-data-processing-unit-friends-or-foes/>