



RAIDIX 4.X

St. Petersburg, Russia

Who am I?

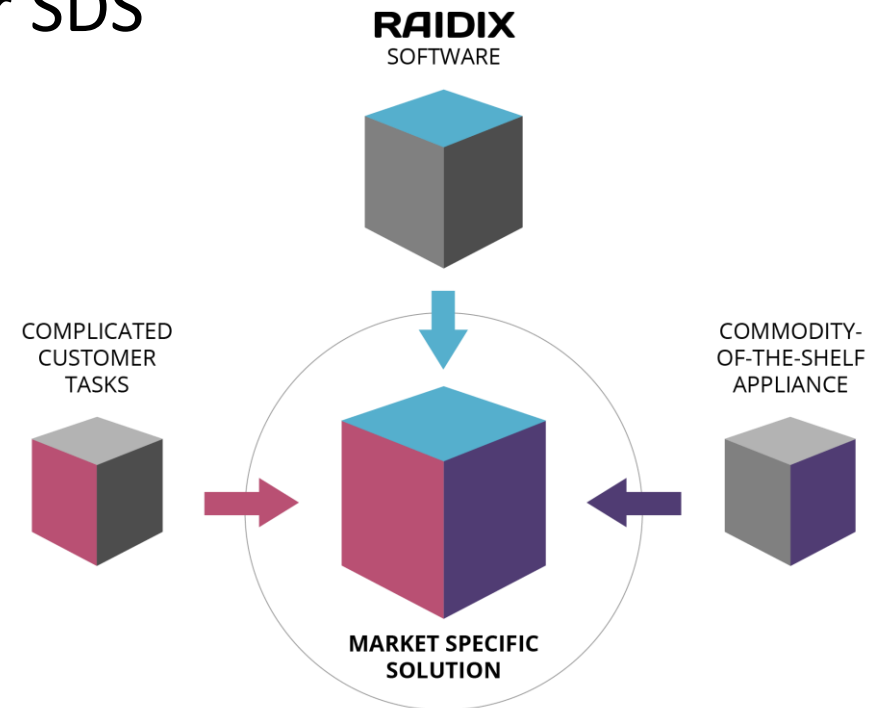
Sergei Platonov

Responsible for product
strategic planning

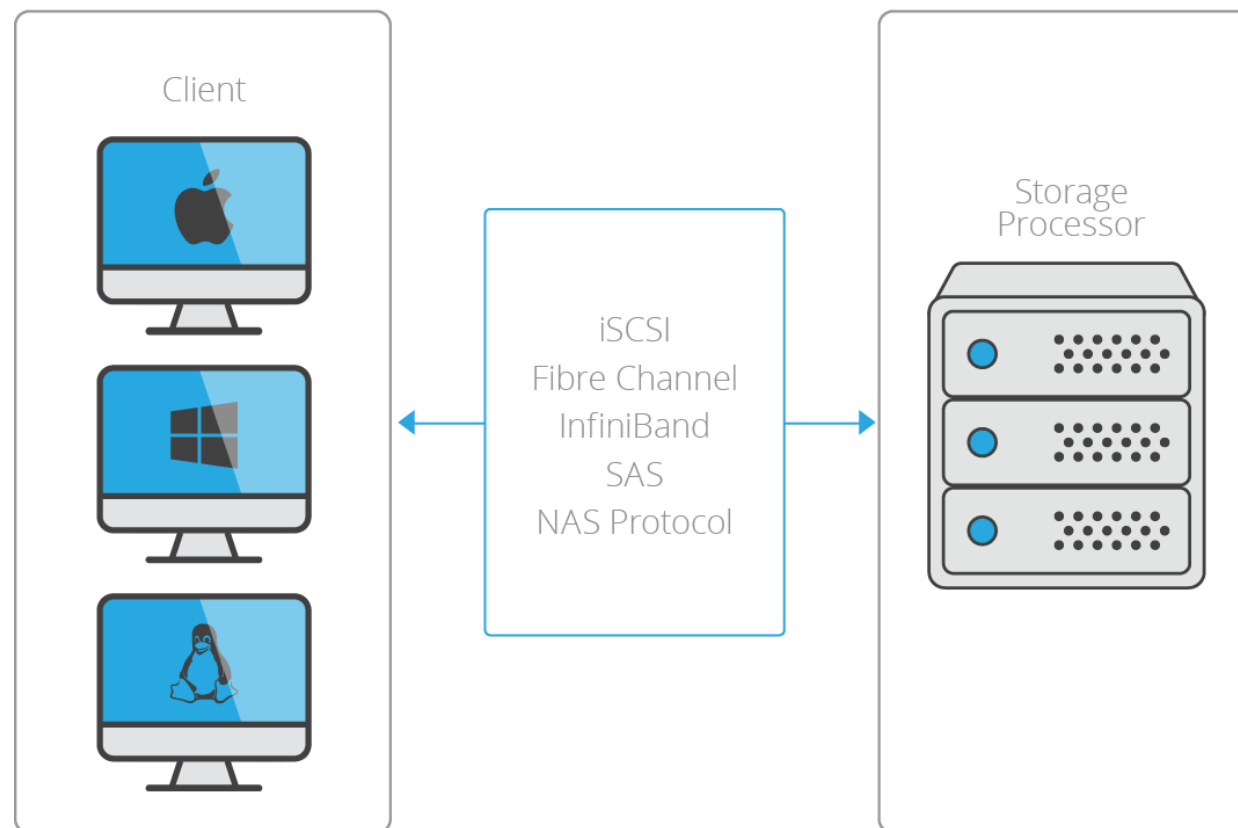


What is RAIDIX 4.x?

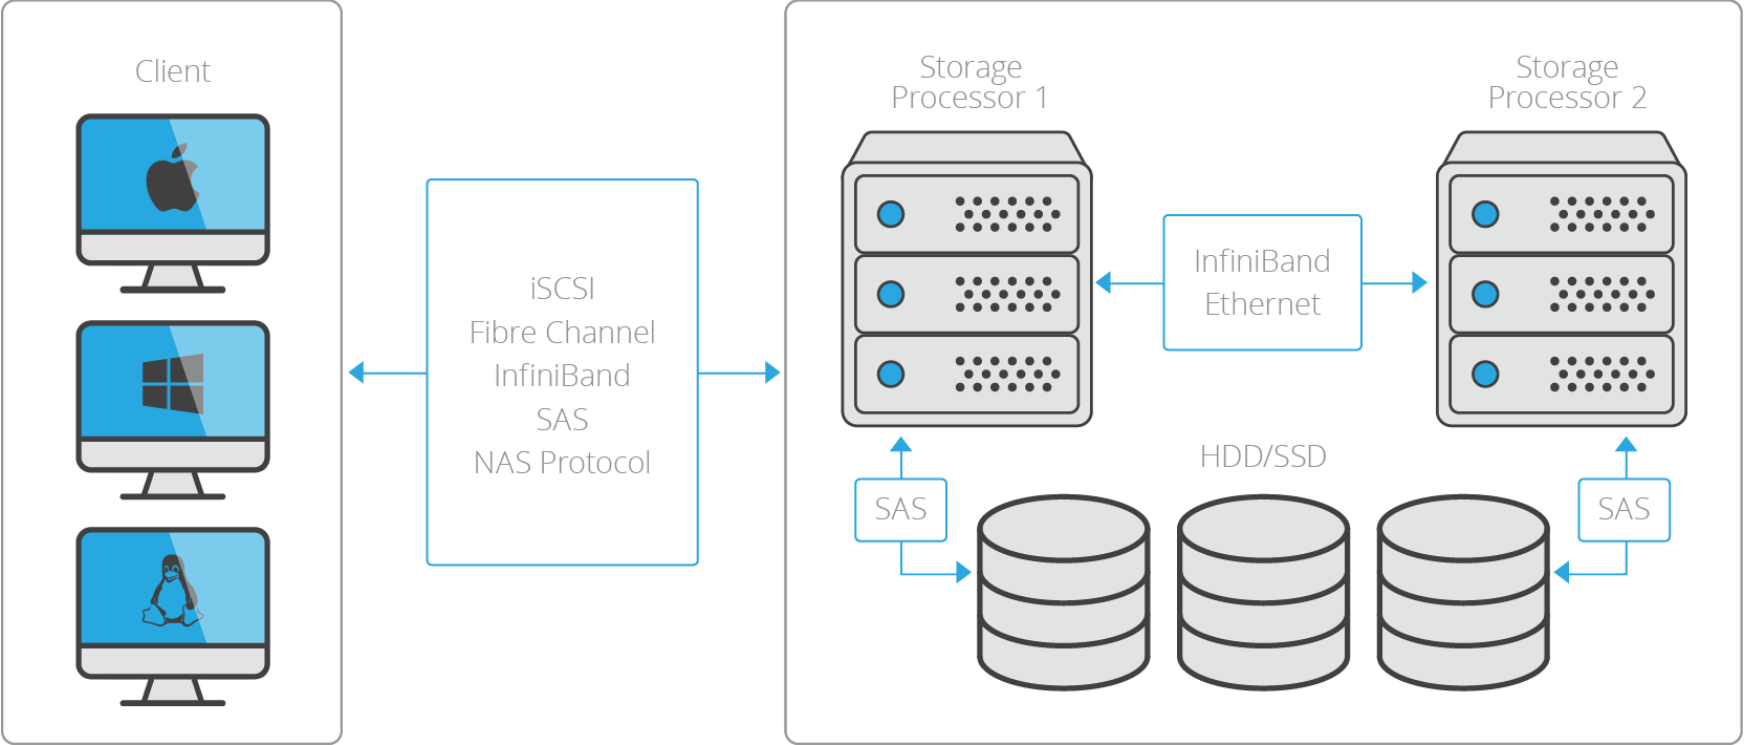
- RAIDIX transforms commodity HW in a **high-performance** data storage system
- 30% better cost-efficiency compared to other SDS
 - We developed the **fastest RAID** in the industry!



System architecture: single controller



System architecture: dual controller



Typical Solutions



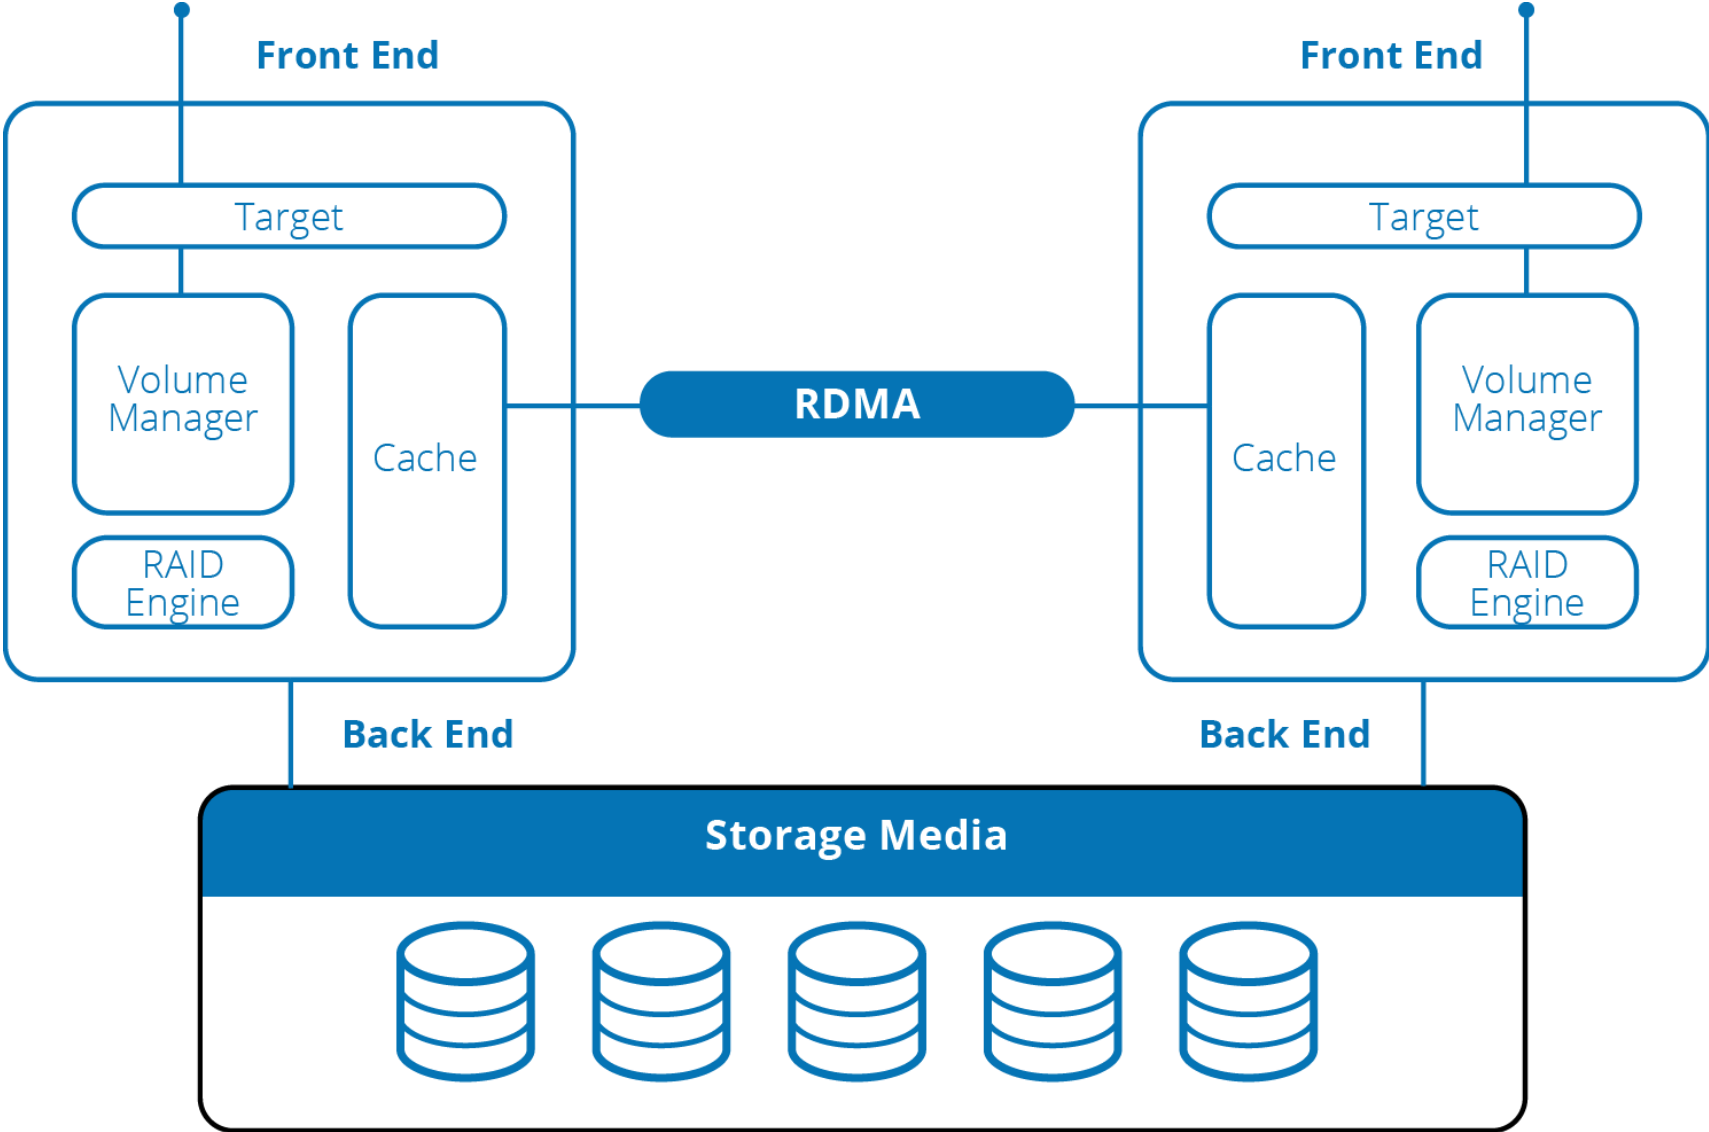
Single controller entry level storage
Up to 4GBps and up to 100k IOPs



Dual controller mid-range level storage
Up to 15GBps and up to 200k IOPs



Dual controller scalable storage system
Tens of GBps and Hundreds of thousands of IOPs





RAIDIX 4.X Highlights

- Block Storage (FC, SRP, SAS 12G, iSCSI)
- NAS (CIFS/SMB, NFS, AFP, FTP), integration with AD, quotas
 - WORM
- RAID 0, RAID 5, RAID 6, **RAID 7.3**, RAID 10 and **RAID N+M**.
- Cluster-in-a-box



RAIDIX 4.X Highlights

- L2 Caching
- Deep Cache Tuning
- QOS, Smart QOS (QoSMic) – Application level QoS
- Thin Provisioning and Deduplication

The RAID Engine operates with the following RAID levels: RAID 0, RAID 5, RAID 6, RAID 7.3, RAID 10 and RAID N+M.

RAID 7.3

- RAID 7.3 is the level of interleaving blocks with triple parity distribution that allows restoring data in case **up to 3 drives** fail. RAID 7.3, based on **proprietary high-performance RAIDIX algorithms**, ensures high performance levels without extra load on the CPU.
- RAID 7.3 offers higher reliability compared to RAID 6 since three drives are allocated for checksums in this scenario.

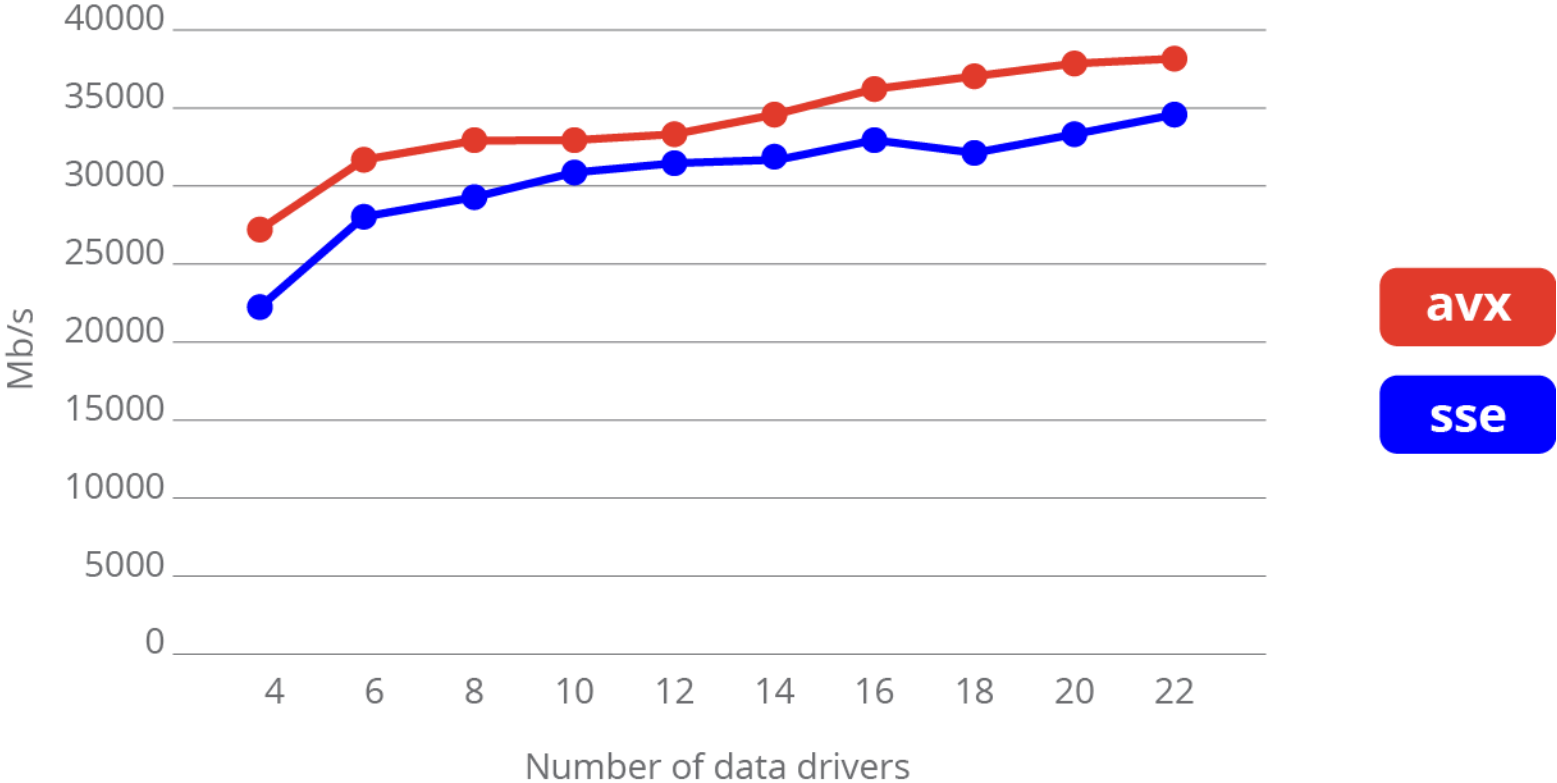


RAID N+M

- RAID N+M is the level of interleaving blocks with M checksums, based on proprietary RAIDIX mathematical algorithms. RAID N+M allows the user to choose a number of disks for checksum allocation. RAID N+M requires at least 8 disks and can sustain complete failure of up to 64 drives in the same group (depending on the number of parity disks).




RAID-7.3 Checksums calculation



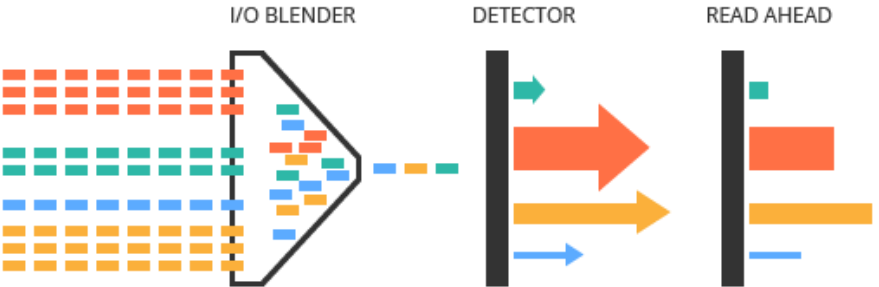


Performance remains unhindered even when multiple disks fail.




Advanced multi-streaming

ADAPTIVE
READ AHEAD




BURST
BUFFER



VOLUME SCALABILITY

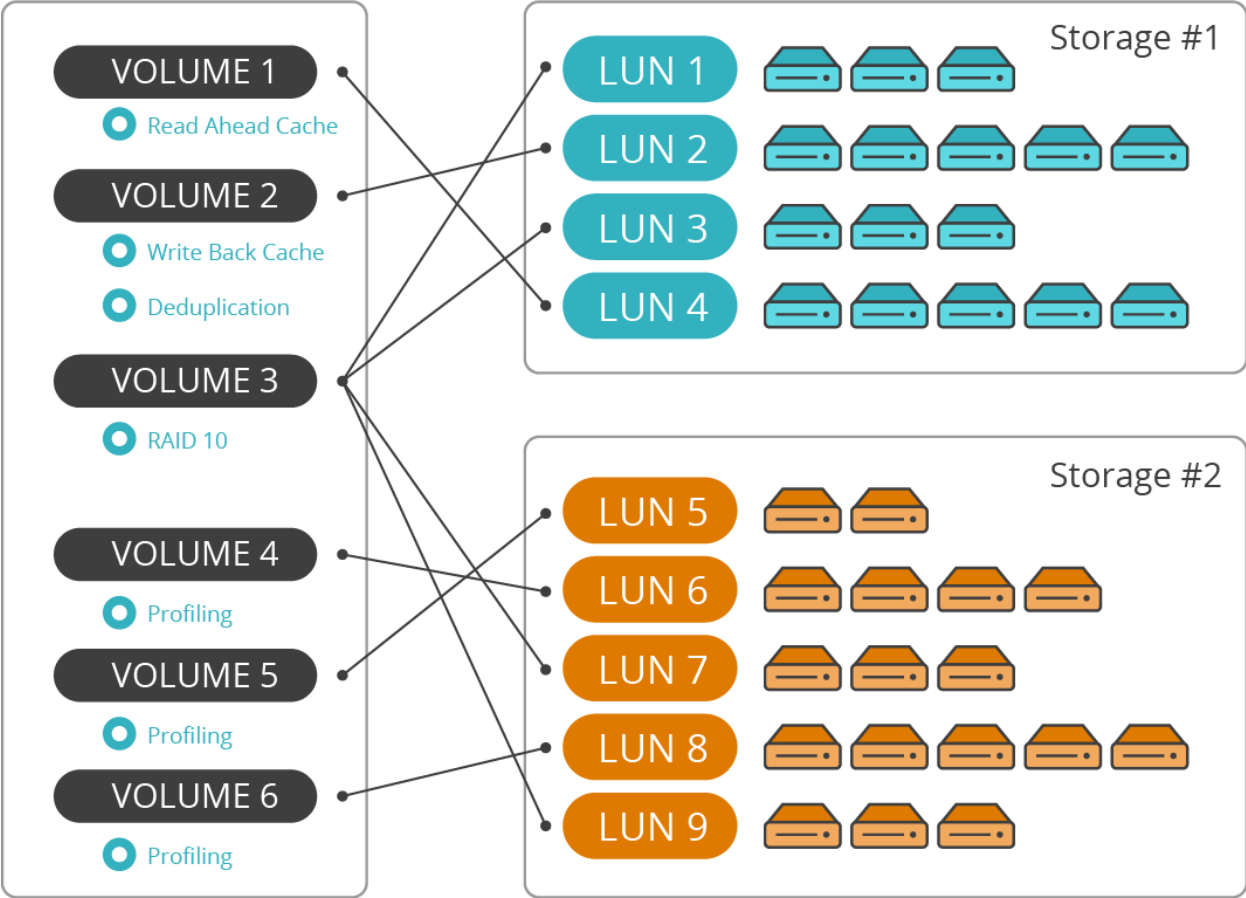
Easy extension of existing infrastructures

During configuration, a virtual volume can be placed on top of multiple RAIDs. RAIDIX allows the user to extend volume capacity by adding new arrays to the infrastructure.



SAN OPTIMIZATION OPTIONS

RAIDIX 4.5





	RAIDIX	Comp 1	Comp 2	Comp 3	Comp 4
Inline Deduplication	Yes	No	Yes	Yes	Yes
Dedupe Granularity	4 KB	4 KB	Dedupes based on IO size 512 B – 32 KB	4 KB to 128 KB (64 KB default)	8 KB
Inline Compression	Yes	Yes	Partial	Yes	Yes
Dedupe before Compression	Yes	No	No	No	Yes
4 K Random IO performance RW80/20	250K	N/A*	150K	<100K	150K per node
Memory Usage	0.28 GB/TB	N/A	4 GB/TB	20 GB/TB	25GB/TB
Scalability	256 TB per pool	100 TB per node	70 TB	Limited by memory requirements	20 TB per node




Thanks for your time!

Questions are welcome!

Sergei Platonov
platonov.s@raidix.com

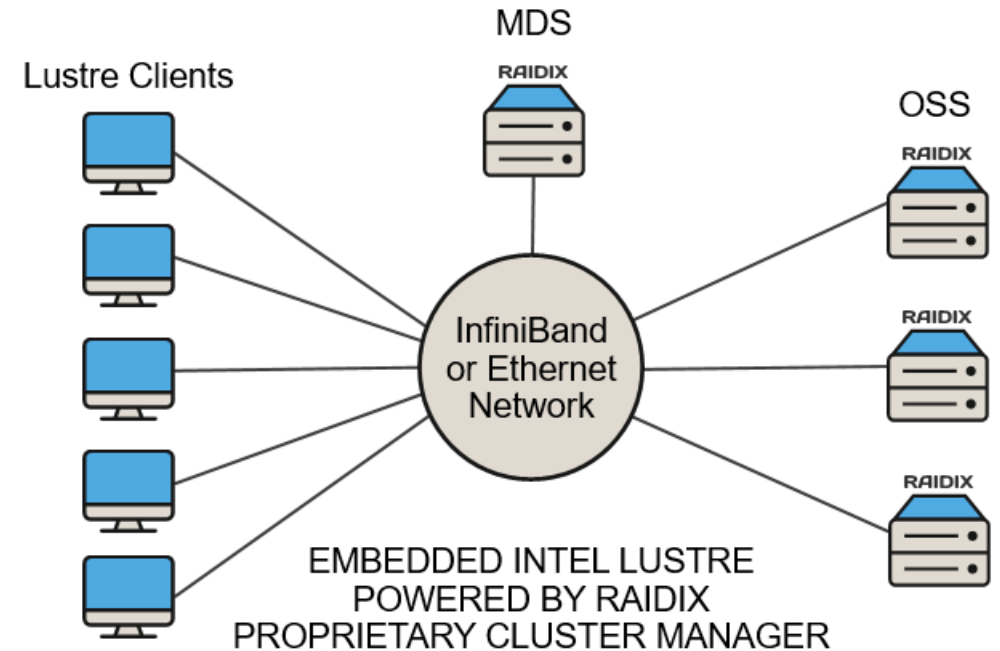
RAIDIX EXASPHERE. SCALE-OUT NAS AND SHARED SAN SOLUTION

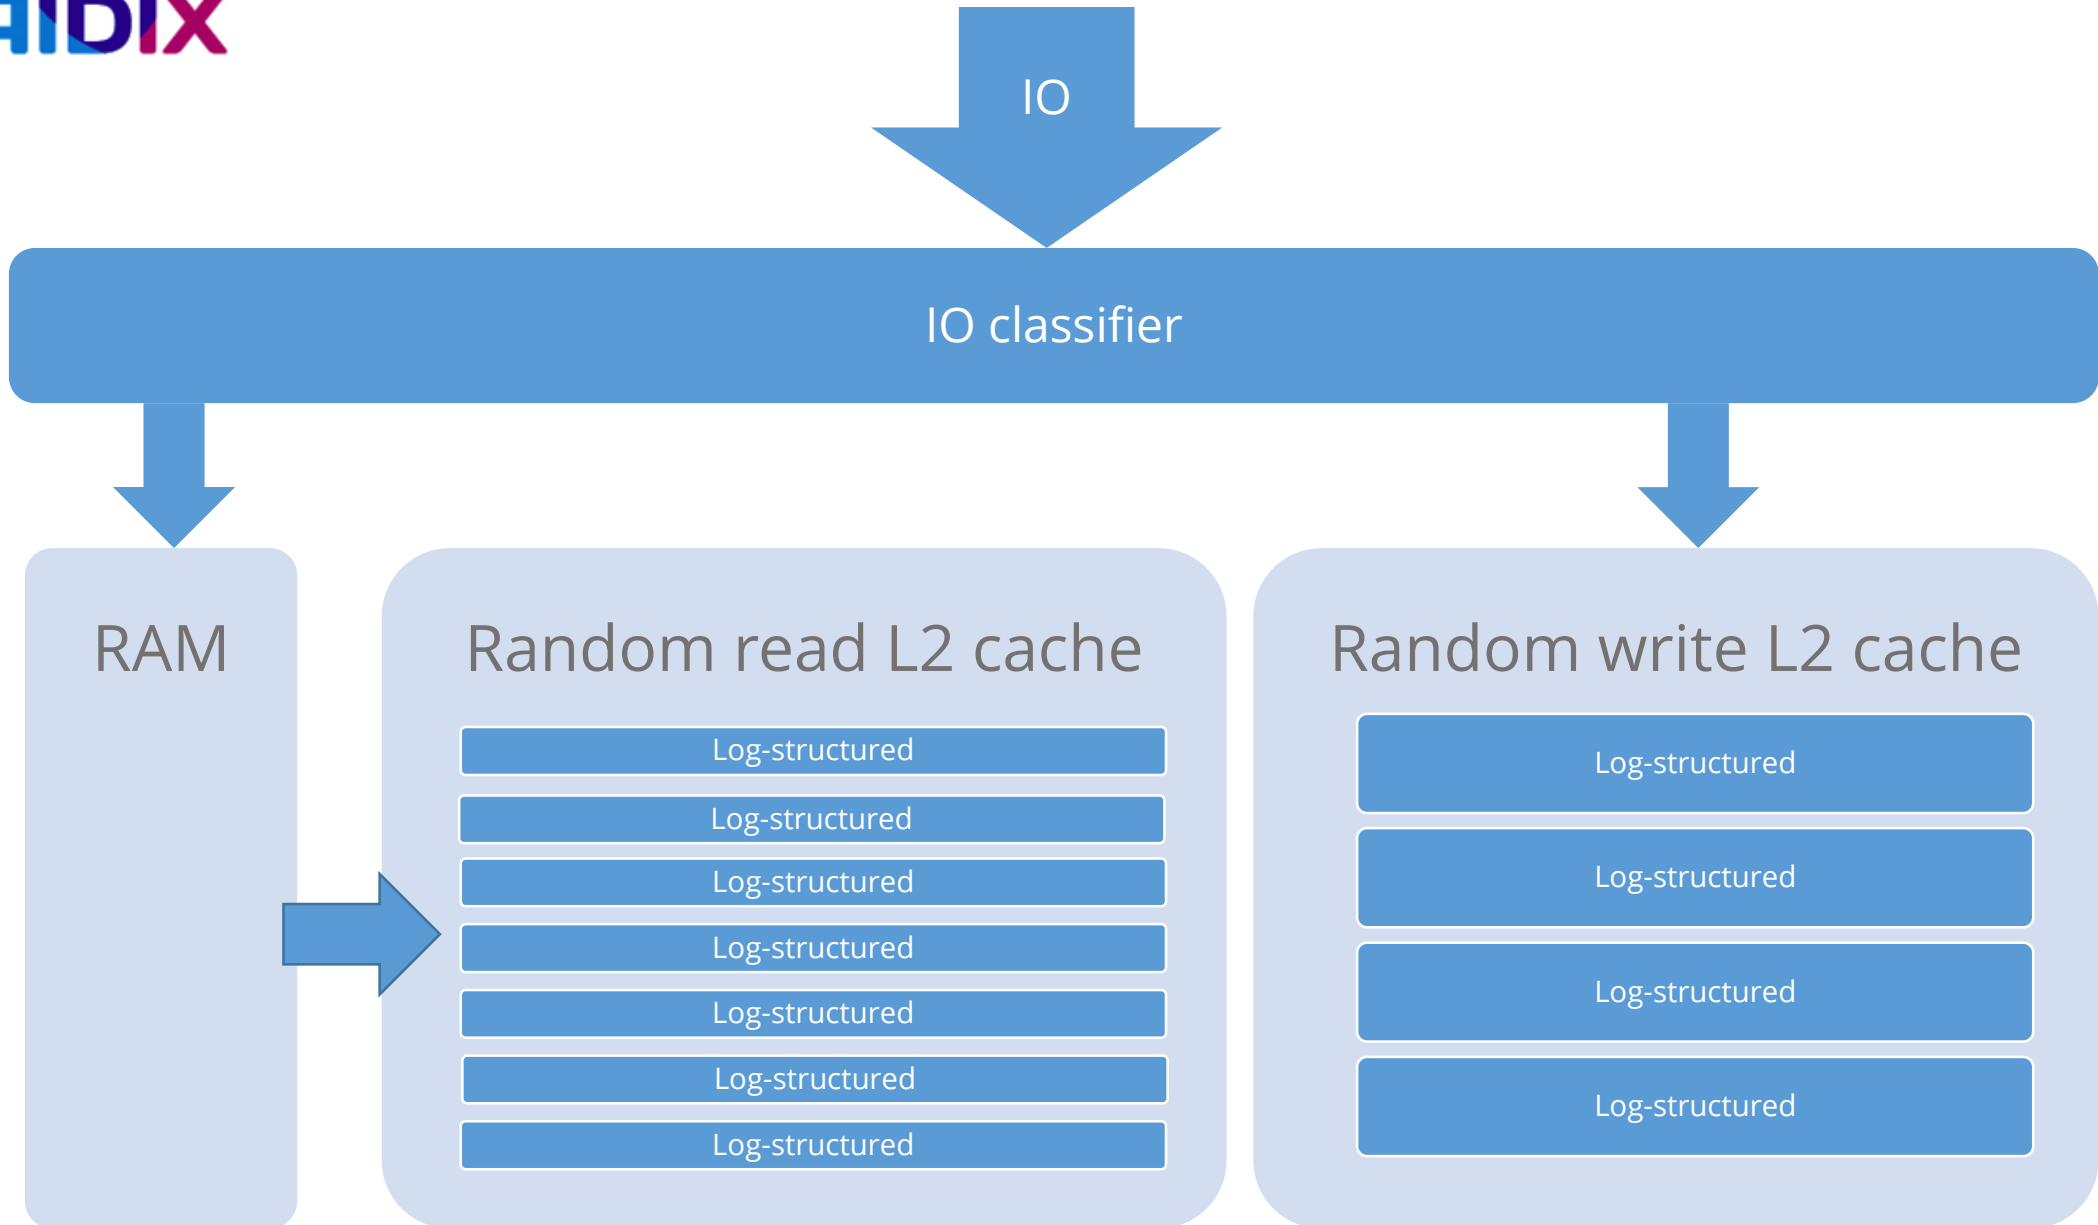
- storage nodes presented as single pool
- Scalability: add more storage as you grow
- High performing block interfaces
- Imbedded metadata server
- Fault-tolerant configuration when reliability matters



RAIDIX EXASPHHERE. SCALE-OUT NAS AND SHARED SAN SOLUTION

- RAIDIX HPC/cluster-in-a-box technology + Intel® Enterprise Edition for Lustre*
- Building high performance data storage clusters
- Scalability up to 512 PB, throughput up to 2 TB/s.
- High performance, fault-tolerance and fast failover
- Asymmetric DC architecture
- Direct installation of Intel® Enterprise Edition for Lustre*, OSS+OST (Active-Active) and MDS+MDT (Active-Passive), on the storage nodes
- Lustre* services can manage volumes locally, hence lower latency and boosted performance







Deduplication Index

Benefits:



Fastest Index

200,000 ops/sec/core




Greatest Scalability

640 billion objects/server

Grid Expandability

Linear growth of performance and capacity



Minimal I/O

0.1% - 1000 IOs per million ops

Lowest Footprint

0.1 bytes/entry in RAM





Virtual Block Store

Benefits:



Fastest Performance

Mapping layer supports up to 1.3M (4 KB) IOPS
Immediate space recovery, no garbage collection



Greatest Scalability

Up to 4 PB physical/1 EB logical



Lowest Footprint

280 MB RAM/1 TB physical storage at 4 KB granularity

