

The background is a solid blue color with a repeating pattern of lightbulbs. Most lightbulbs are white outlines, but several are filled with a yellow-green color and have short lines radiating from them to represent light. The lightbulbs are scattered across the entire background.

ZADARA
SUMMIT
2018

**INNOVATE
NOW**



Flashed-Optimized VPSA

Always Aligned with your Changing World



Yair Hershko

Co-founder, VP Engineering, Zadara Storage

Modern Data Storage for Modern Computing

- Innovating data services to meet modern data center challenges
 - capacity growth, increased performance, data mobility and reduced cost
- Commitment to Quality
- Software Defined Storage
- Leverage new industry exciting technologies to bring cost reduction at optimal performance
- Optimized and tunable solutions for dynamic changing workloads
- Single-Tenant experience in a Multi-Tenant environment



Introducing Flashed-Optimized VPSA

- New Data Service available via a new type of VPSA
- Co-exists with Storage-Array & Object-Storage VPSA in a single Zadara Cloud
- Data Mobility between all types of VPSAs
- Designed for Cost and Performance Optimization
- Designed to take advantage of SCM, NVMe, 3D NAND, and fast Interconnect



Key Architecture Principles

- Evolution of VPSA Storage Array
 - All of the features you love with VPSA remain intact
 - Protocol layer - FC, iSCSI, iSER, NFS, SMB
 - Data Protection - Snapshots, Multi-Zone, Remote Replication, B2OS, Volume Migration
 - Security - Encryption, AD-Integration
 - ...
- Revolution of Data and Metadata Infrastructure
 - New Highly Scaleable/Performing Metadata infrastructure along with New data layout for optimal data placement on Flash media, with revised IO flows to enable data reduction and emerging technologies adoption



Key Design Principles

- Inline data reduction for overall Cost Efficiency
- Low latency Write IO Flow
- Minimize flash wearing with 'friendly' writes
- Controlled write Workload
- Distribute Data and Metadata Read workload
- Consistent Performance
- Deduplication per Pool




Under the Hood



Performance Consistency

- Write-Buffer and hot metadata on Optane for best Flash QoS
- Controlled Write pattern at fixed Segment size
- Data and Metadata are distributed on all devices
- Fixed Dedup Chunk Size
- Patent-pending Large Key-Value-store for efficient handling of Billions of metadata Keys

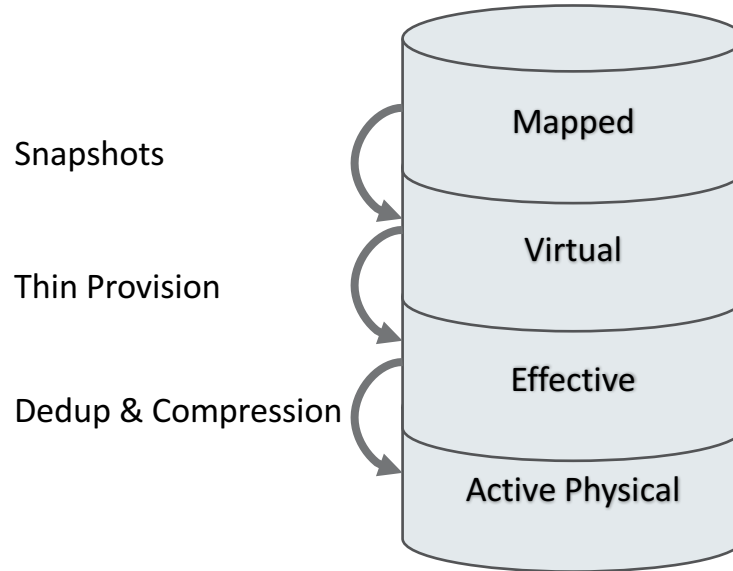
Capacity Saving



- Thin provisioned Volumes occupy pool capacity only where data was written
- Pattern removal identifies and removes repetitive binary patterns, including zeroes.
- Inline Deduplication to remove duplicates, and store only unique blocks on flash
- Inline compression reduces data to use less capacity than the original format
- Byte-aligned Compaction to optimize flash capacity utilization

Capacity saving

- Accurate Stats
- Alerts

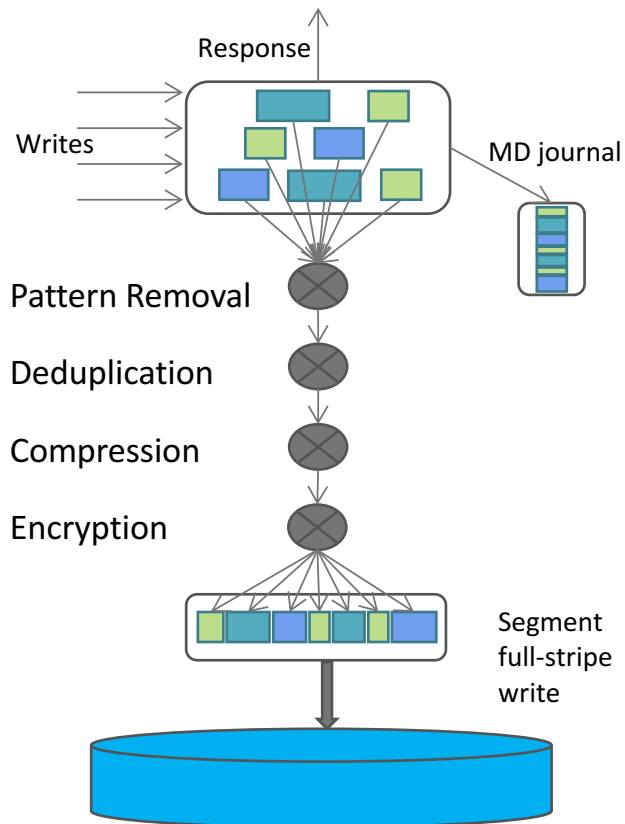


Data placement

- No write-in-place.
 - Write new data to a new location. Invalidate old data
 - Higher resiliency and crash-consistency
 - Suitable for dynamic compressed data size
- Compact compressed data into data Segments
 - Byte Aligned
 - Fixed-size data segment (1MB)
- Sources of Segment Data
 - New Host writes
 - Defragmentation
- Defragmentation
 - Segments with least amount of “valid” data

IO Flow

- Writes
 - Terminate at Write Buffer for fast response
 - Journal for crash consistency
 - Deduplication at fixed-size chunk size
 - Compression
 - Byte-level Compaction into Segments
 - Segments Write-out to Pool devices
- Reads
 - Single read Lookup for Pool location (Segment & Offset)
 - uncompress



Pool Types


- Transactional (High IOPs)
 - 8KB LSA Chunk Size, 512KB Thin-Provision Chunk Size
 - Maxed Mapped Capacity - 100TB
 - Recommended for - Online Transaction Processing
- Repository (High Throughput)
 - 16KB LSA Chunk Size, 1MB Thin-Provision Chunk Size
 - Maxed Mapped Capacity - 200TB
 - Recommended for - VDI, Analytics, General Purpose
- Archive (High Capacity)
 - 32KB LSA Chunk Size, 2MB Thin-Provision Chunk Size
 - Maxed Mapped Capacity - 400TB
 - Recommended for - Archive

Modern Hardware

Intel Hardware



Component	Capabilities	Contribution
Optane (3D XPoint)	High-Performance High-Endurance Flash Media	Frequently accessed Metadata Write Buffer
3D NAND NVMe SSD	Dense Flash capacity Fast Protocol	Reduce \$/GB Improve performance
R2224WFQZS Server (WolfPass)	Modern Storage Platform	Increase PCIe Lanes Increase NVMe Performance
Dual Skylake Microarchitecture	56 vCPUS	Faster IO Processing and data reduction capabilities
Quick Assist	Hardware assisted Compression and Deduplication	Data Reduction Efficiency

Intel® Optane™ Raw Performance



*Compared to DC P3700 (FD). All comparisons & estimates based on target spec sheet metrics, target performance numbers subject to change.

Intel® 3D NAND SSD



Intel + Zadara

Intel

Details:

- ✓ Intel Wolfpass servers
- ✓ Dual Skylake CPUs (28 cores/SN)
- ✓ Flash 3D NAND (4TBs SSD NVMe)
- ✓ Optane (3DXpoint)
- ✓ Quick Assist

+

Zadara

Details:


- ✓ New VPSA All Flash Software
- ✓ Inline Deduplication
- ✓ Inline Compression
- ✓ Block & File
- ✓ Cloud & OnPrem

=


Modern
Storage



THANK YOU!

 +972-52-4745020

 yair@zadarastorage.com

 Skype: [yair.zadara](https://www.skype.com/people/yair.zadara)

