Optimization Techniques Change the Competitive Landscape

Jered Floyd

CTO, Permabit Technology Corp.

Permabit and Albireo Overview

- Headquarters: Cambridge, MA
- Founded: 2000, MIT roots
- Easily integrated storage software for data deduplication and software-defined storage (SDS)

Software that integrates at the firmware, device driver, file system, appliance or application level to deliver:

- Multiple GB/sec IO throughput
- 5 – 35x reduction rates
- Minimal Memory and CPU footprint
- Petabyte scalability within a single device
- Guaranteed data integrity
- Flexible, simple integration
- No risk, time-to-market advantage

How Albireo is Delivered

- Albireo SDK
  - Deduplication indexing service
  - Software library for Linux and Windows
  - API documentation, code samples, application notes
- Albireo VDO Products
  - Albireo Virtual Data Optimizer (VDO)
    - Turn-key dedupe and thin provisioning for Linux
    - Binaries, management scripts, documentation
  - Albireo COMPRESS
    - Inline, block level compression
  - Albireo REPLICA
    - Fast and efficient offsite replication of files between VDO-enabled storage devices

Albireo

Albireo
Virtual Data Optimizer

Albireo
COMPRESS

Albireo
REPLICA

Dedupe and Compression

- Compression identifies "micro" duplicates (2-4 X)
- Dedupe identifies "macro" duplicates (4-100 X)

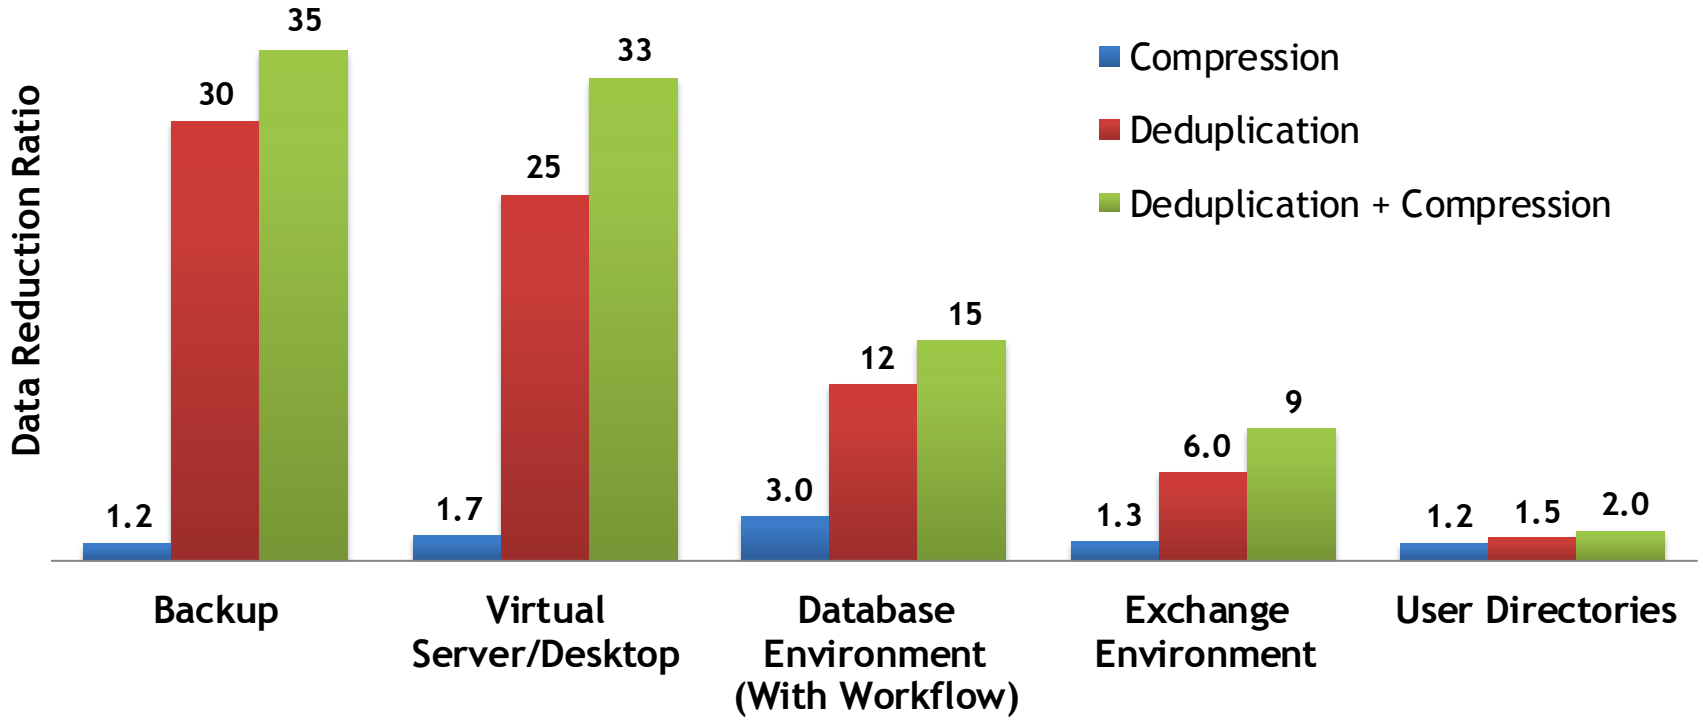

Benefits of Flash Controller Deduplication

- Permabit's Albireo deduplication technology can be run at controller level due to its small footprint and efficient resource utilization

- Deduplication provides significant improvements in
  - Over-provisioning
  - IOPS
  - Wear and reliability
  - Margin

- SSD is a block device so dedupe does not increase user capacity and does not reduce demand and revenue

SSD Deduplication

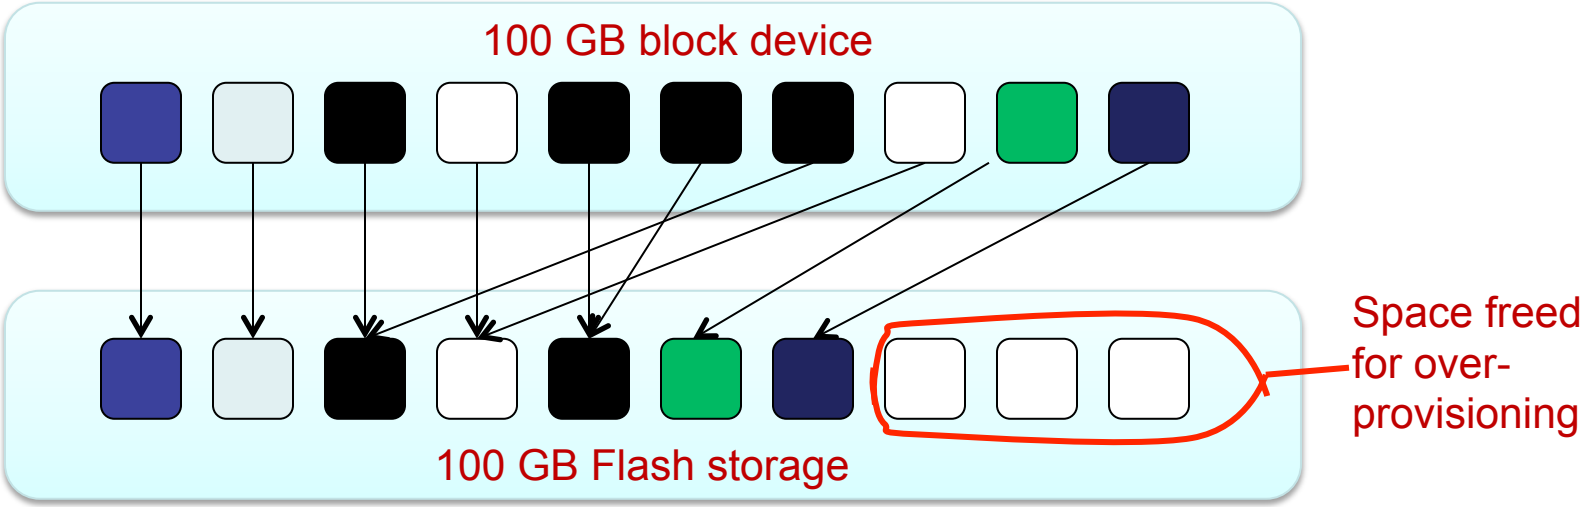

- Deduplication identifies duplicate blocks
- Space is saved by storing only one copy of the blocks to SSD
- Freed space can be used for over-provisioning

Client Devices
Dedupe @ 1.7x Average

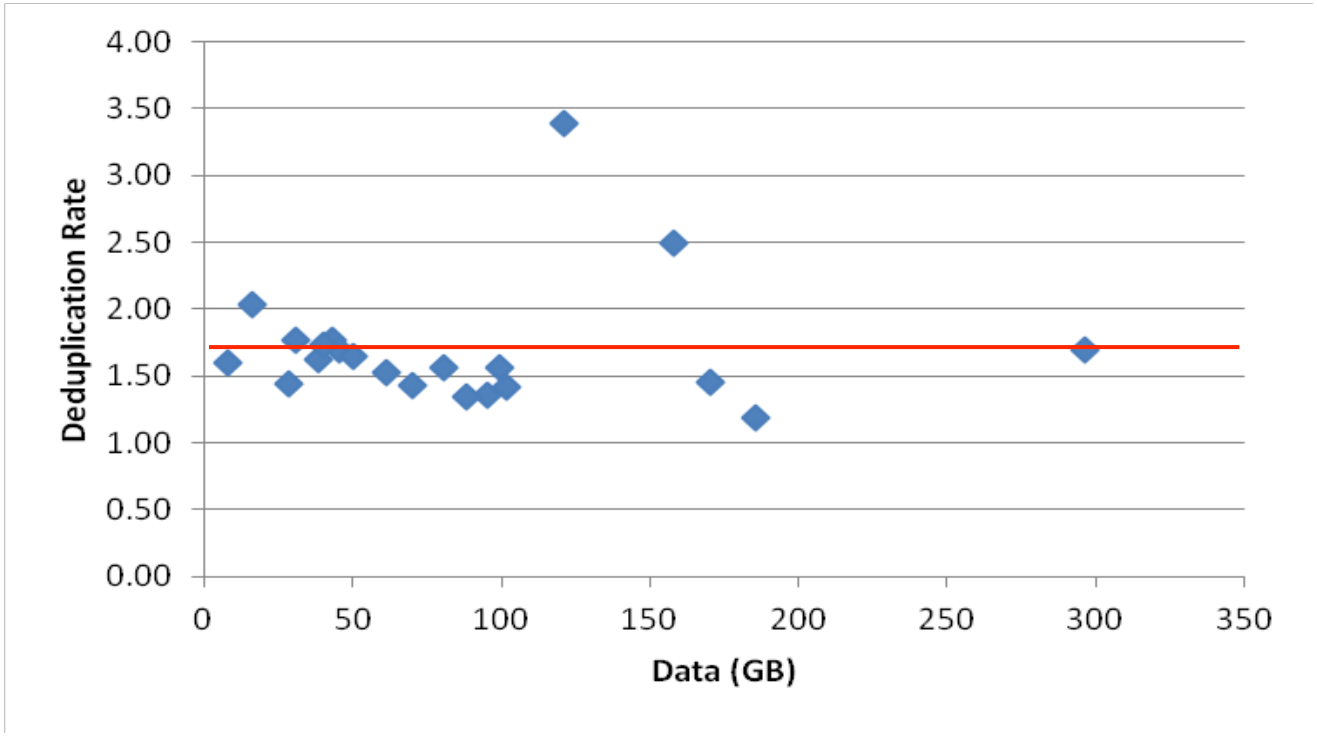

- Identify duplicate blocks in Windows OS and user data
- 256GB flash @ 1.7x dedupe frees up 102GB for over-provisioning

FTL/Controller Use Case

Example: Intel S3700 Outperforms Samsung 840 Pro

Due to Intel's higher over-provisioning rate

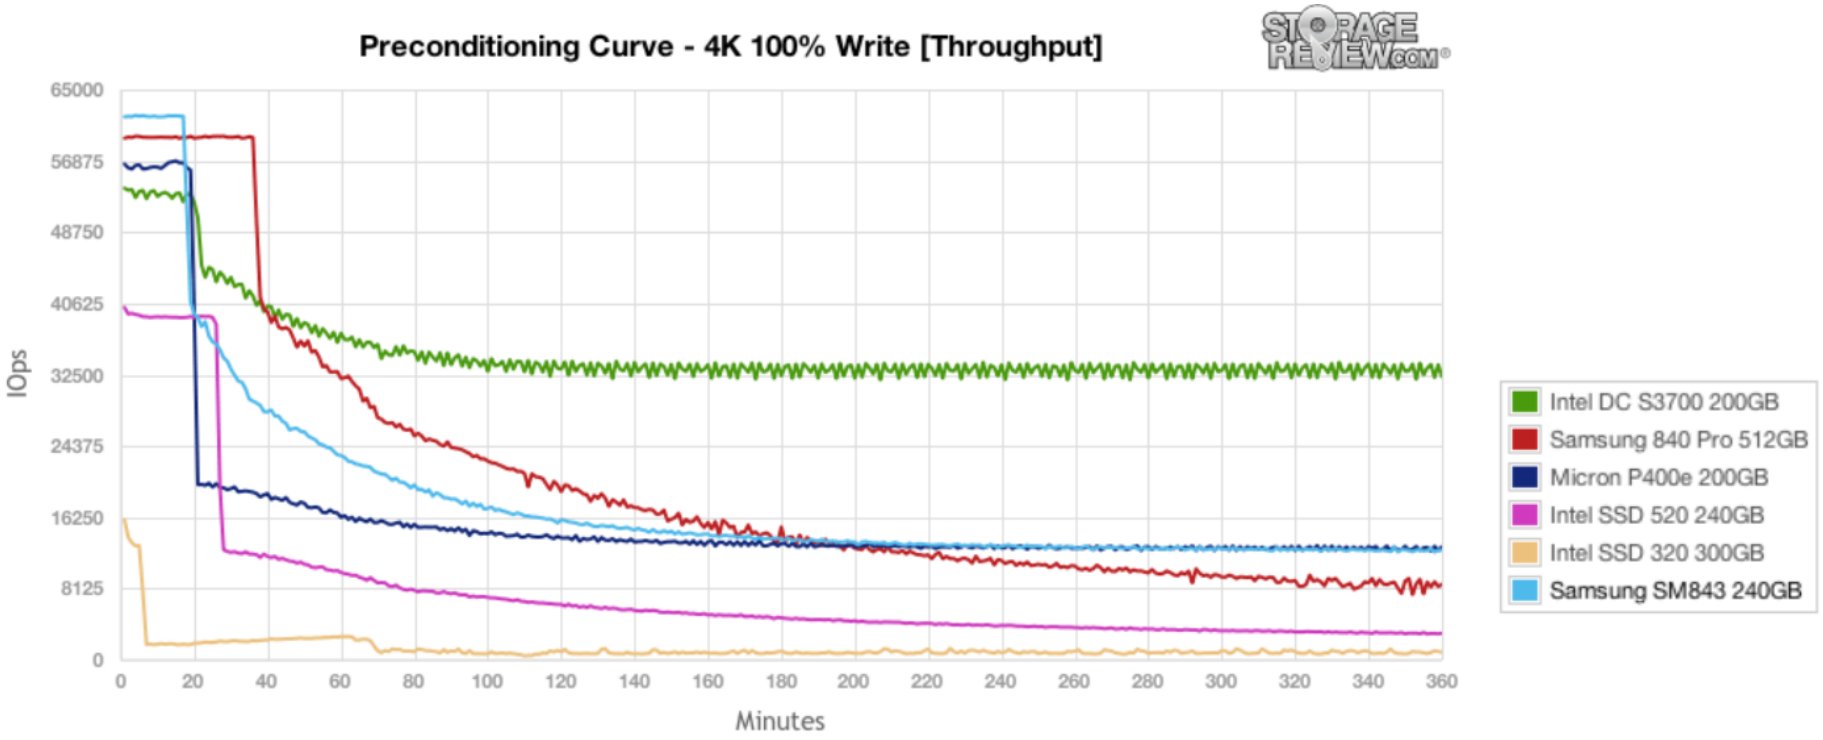

http://www.storagereview.com/samsung_sm843_enterprise_ssd_review

Consistently High IOPS (at a Cost)

30% of capacity is set aside for over-provisioning

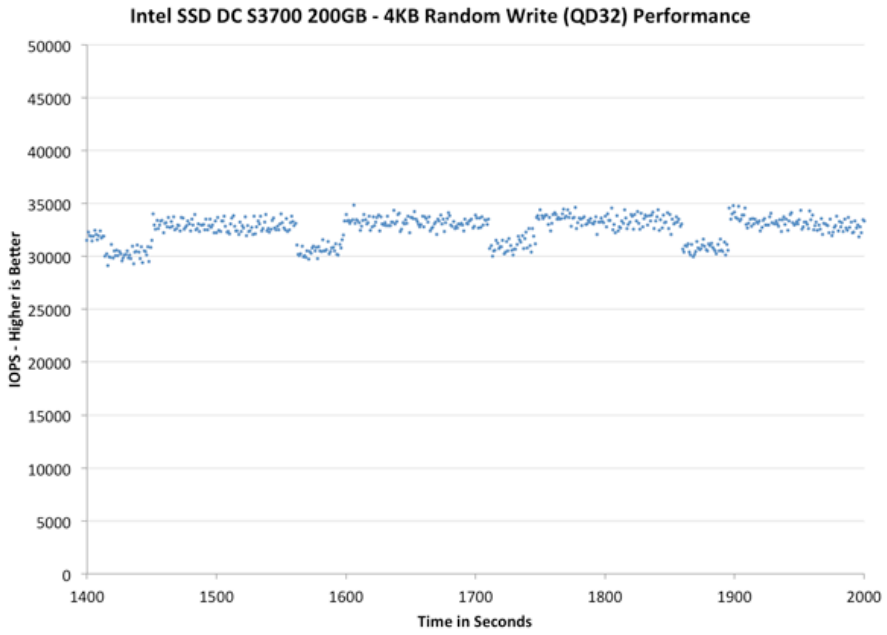

- *"S3700 has 264GiB of NAND on-board but only exposes 186GiB of it (200GB advertised capacity) as user accessible storage, the rest is used as spare area to improve performance, consistency and endurance"*
- *"IO consistency is a good optimization practice for all SSDs"*
- *"Additional spare area alone isn't enough to deliver the degree of consistency that Intel's S3700 offers. The solution really has to be a combination of controller and spare area"*

http://www.anandtech.com/show/6489/playing-with-op

Over-provisioning Improves IOPS

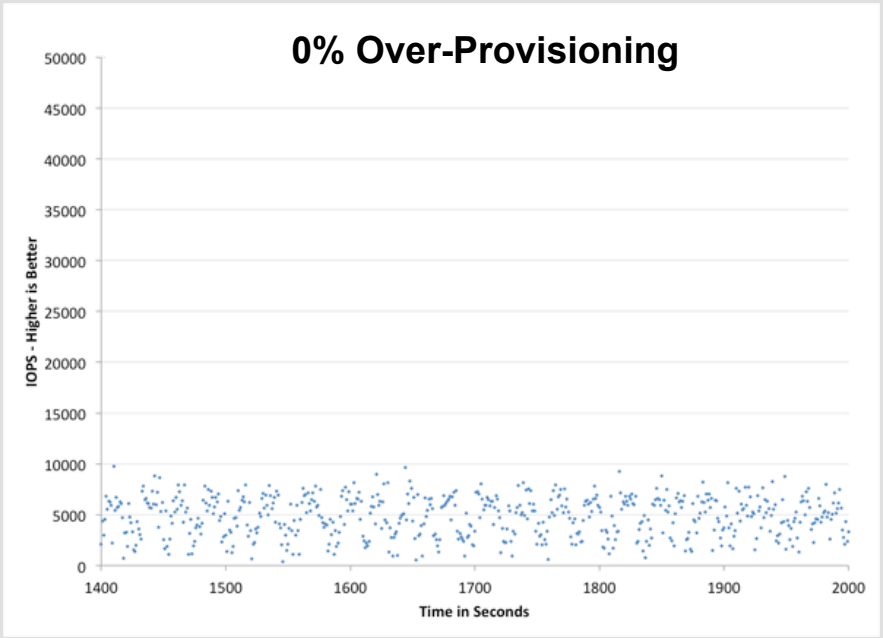

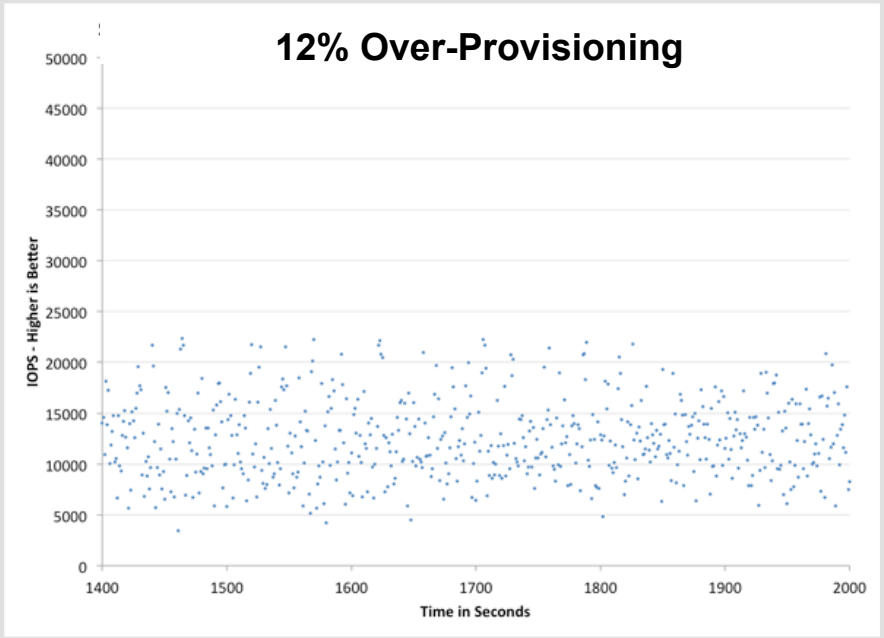

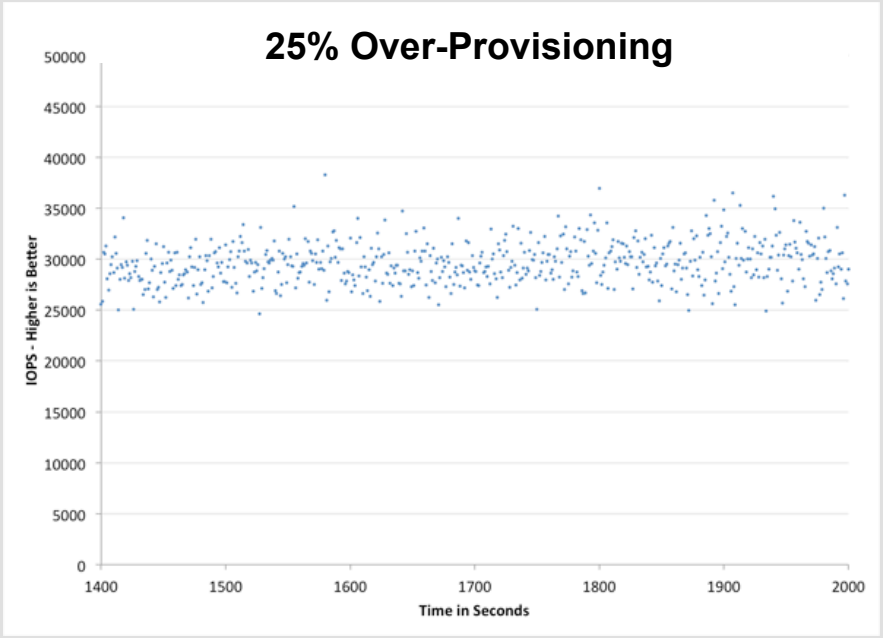

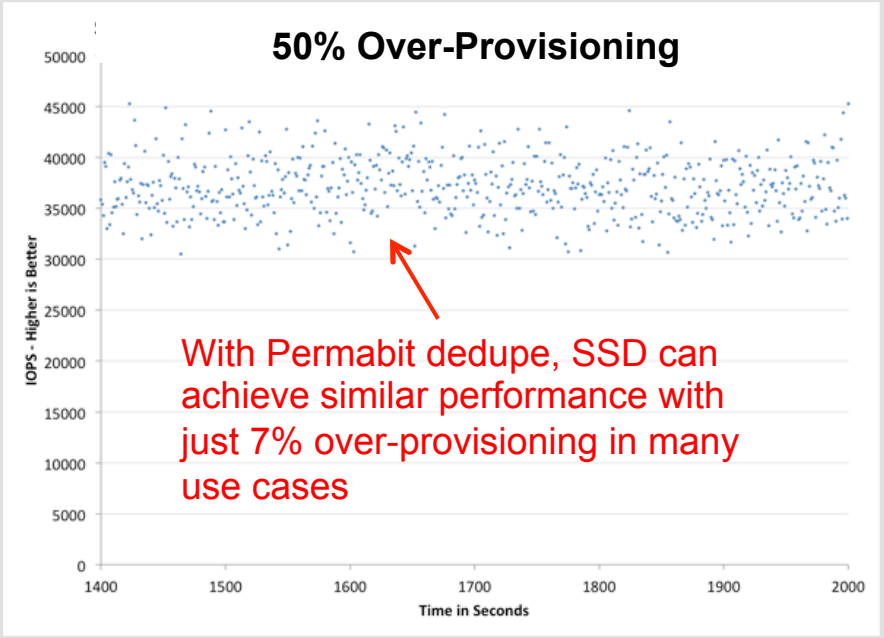

http://www.anandtech.com/show/6489/playing-with-op

Example Benchmarks

Samsung SSD without Albireo (0% Over-Provisioning equiv.)

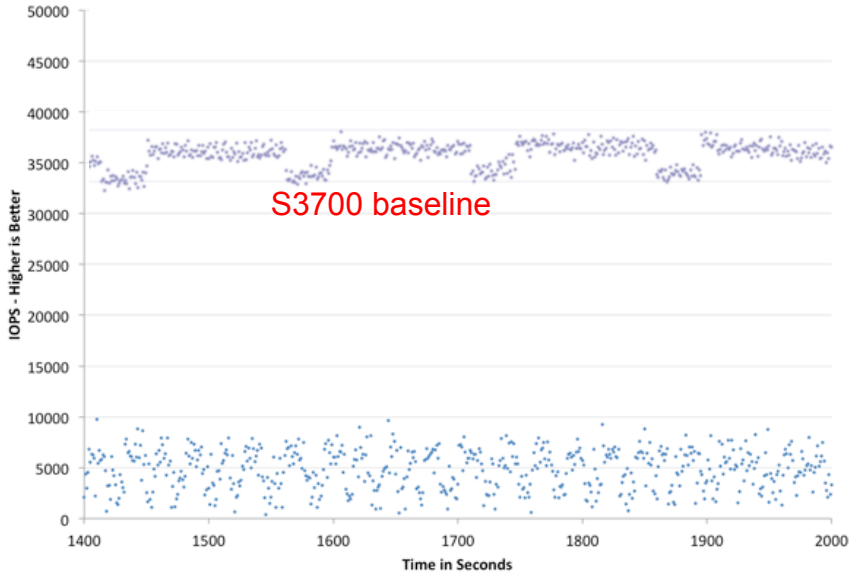

Samsung SSD with Albireo Dedupe (50% Over-Provisioning equiv.)

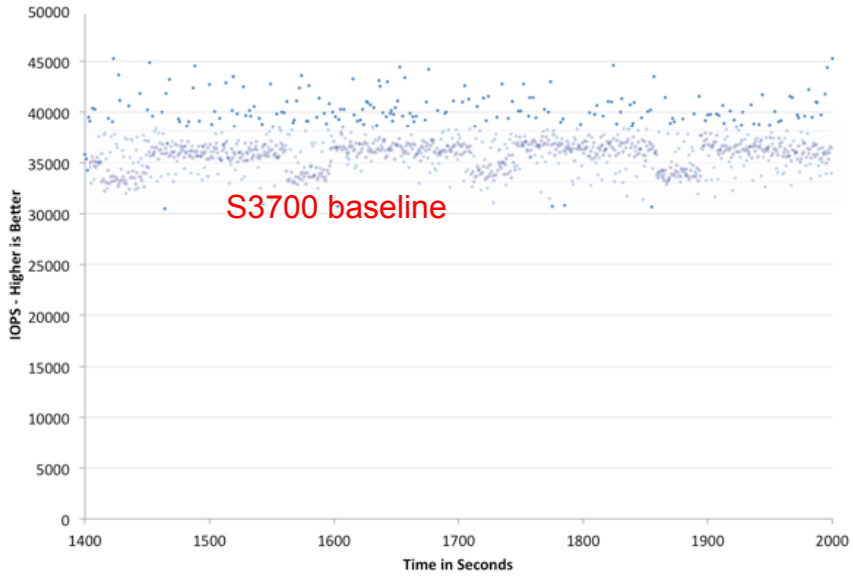

- **S3700 beats 840 Pro performance with over-provisioning, but costs useable drive capacity**

- With deduplication, 840 Pro can deliver higher performance without added spare capacity cost for over-provisioning
- User capacity is constant

SSD with Dedupe Provides 4x Better Over-Provisioning than Competition

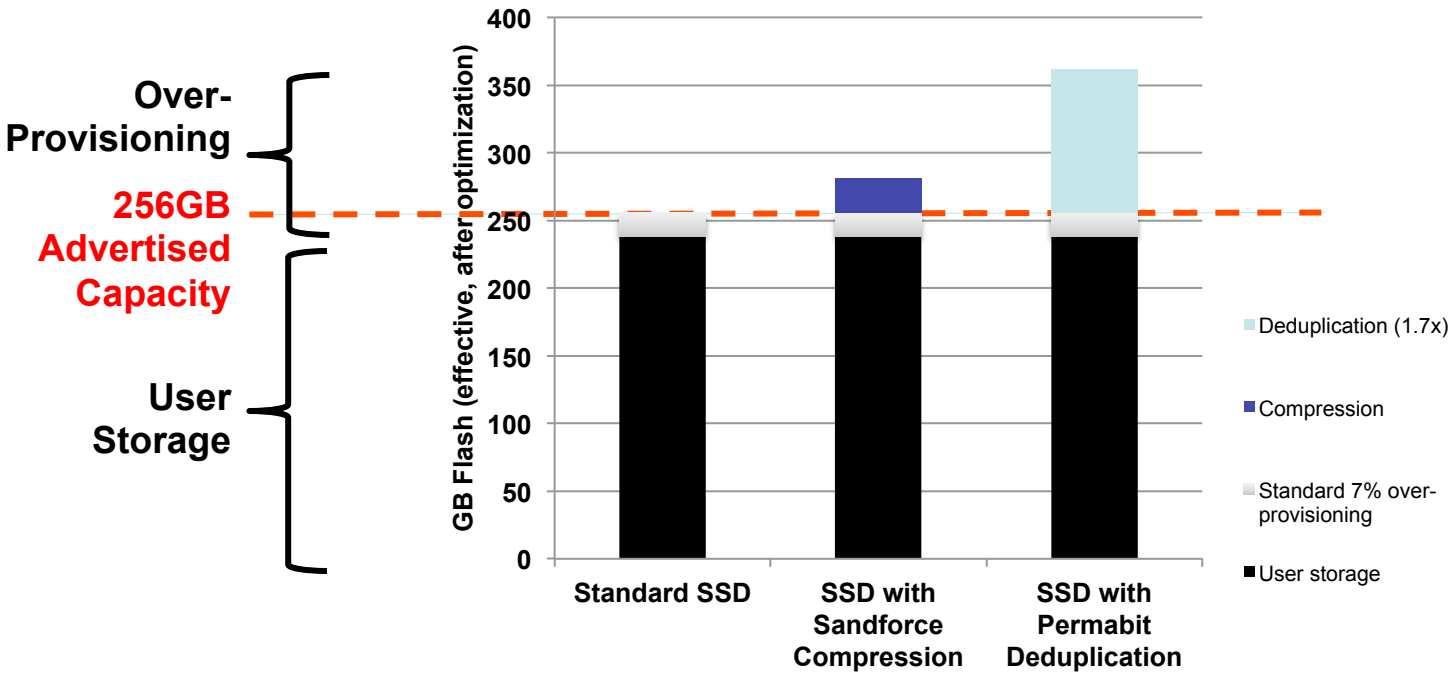

- Permabit can provide the performance of 50% over-provisioning while setting aside only 7% of physical storage
- SSD vendors can improve performance while setting aside far less spare capacity for over-provisioning

Deduplication Can Substantially Improve Cost Structure

- Reduce amount of memory set aside for over-provisioning yet get great results on benchmarks.
  - Reduce random write latency by 4x or more
  - Increase random write throughput by 4x or more
- Even when zero deduplication, no downside to user
- Deliver better pricing with faster performance than the competition
- Breakthrough technology for Flash controllers