



Education

SSD and Deduplication – The End of Disk?

Jered Floyd, CTO, Permabit Technology Corp.

Author: Jered Floyd, CTO, Permabit Technology Corp.

- The material contained in this tutorial is copyrighted by the SNIA unless otherwise noted.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
 - ◆ Any slide or slides used must be reproduced in their entirety without modification
 - ◆ The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA Education Committee.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.
NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

➤ SSD and Deduplication – The End of Spinning Disk?

- ◆ Solid State Disks (SSD) have become a “must have” storage technology, however rapid adoption is hampered by higher costs and longevity concerns. Data deduplication closes the gap, where savings rates of up to 5-35x dramatically reduces SSD effective cost. Deduplication also reduces writes, greatly enhancing SSD reliability.

One Reason:


NO SEEK LATENCY

Storage virtualization, and subsequently server virtualization, have created fundamental storage challenges by greatly increasing the density and access randomness of storage I/O.

Where does SSD/Flash sit?

➤ Location

- ◆ Local storage
- ◆ Intermediate appliance
- ◆ Networked storage



➤ Usage

- ◆ Primary
- ◆ Tier
- ◆ Cache

➤ Price

- ◆ Flash storage remains substantially more costly than equivalent capacities of bulk storage

➤ Reliability

- ◆ Flash must be powered on and “scrubbed” to maintain published error rates

➤ Longevity

- ◆ Flash write cycles are limited, and common data patterns can result in exponentially faster wear

Dedupe Benefits for SSD

➤ Increase performance

- ◆ Lower cost per IOP
- ◆ Increase effective cache size
- ◆ Reduce average read latency

➤ Increase longevity

- ◆ Reduce wear on flash cells
- ◆ More empty capacity for space reclamation

➤ Increase efficiency

- ◆ Lower cost per GB

- Deduplication requires some form of fingerprinting
- Fingerprint for disk storage
 - ◆ Must be statistically unique, collision resistant
 - ◆ SHA-256 is your best bet
 - ◆ Computationally intensive
- Fingerprint for flash storage
 - ◆ Must be statistically unique, but collisions possible
 - ◆ MurmurHash3, others work
 - ◆ Faster to compute – matching Flash performance

| | | |
|--------------------|---------------|-------------|
| Fingerprint | SHA-256 | MurmurHash3 |
| Performance | 180 MB/s/core | 3 GB/s/core |

Will SSD Replace Disk?

- High IOPS: It already has!
 - ◆ Databases
 - ◆ Large virtual environments

- Low IOPS: Not anytime soon...
 - ◆ Archive
 - ◆ Video back-catalog

- Dedupe expands use cases
 - ◆ Reduced wear delivers lifetime reduced cost
 - ◆ Reduced cost expands OLAP, VDI use cases



Hands-On Lab: Solid State Storage in the Enterprise

**Many thanks to the following individuals
for their contributions to this tutorial.**

- SNIA Education Committee

**Jered Floyd
Mike Ivanov**

**Louis Imershein
Wayne Salpietro**

Send any questions or comments on this presentation to SNIA: tracktutorials@snia.org