



Education

PRIMARY DATA OPTIMIZATION – WHAT SHOULD IT BE?

David Vellante, The Wikibon Project (Moderator)

Jered Floyd – Permabit

Larry Freeman – NetApp

Steve Kenniston – IBM

Craig Nunes - HP

- 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.

- Cloud computing, big data, mobility and the social media wave are just some of the drivers pushing storage growth to overwhelming rates. As information volumes explode, storage budgets remain flat and we've seen an increased need to implement primary data optimization technology. Primary data optimization has been discussed thoroughly in 2011, but adoption remains mixed as practitioners need to better understand how to exploit this technology, the risks of implementation, true costs and bottom line business benefits. While the concept seems simple, data optimization reduces overall costs by locating and eliminating duplicate data, it has yet to fully play out in the market.

What is Primary Data Optimization?

Primary data is data which has high levels of read and write activity.

Optimization is the act of reducing the resources required to store data

Measurements of optimization impacts include bandwidth, IOPs, response times, capacity

The challenge of optimization is how to do it in a way that doesn't interfere with performance, availability and data integrity

Examples of Primary Data Optimization

- Data de-duplication
- Compression
- Thin provisioning
- Thin copy
- Space reclamation
- Automated storage tiering
- Virtualization
- Others?

Key Questions for the Panel

- Why isn't all primary storage 'optimized?'
- What are the costs involved?
- What are the risks?
- When will optimization become "Jacks or Better"?
- How integrated are the different approaches with each other?
- How does primary data optimization integrate with other resources (e.g. applications)?
- What standards exist?
- What about data migration and interoperability?
- Audience/practitioner questions and concerns

- Please send any questions or comments on this presentation to SNIA: tracktutorials@snia.org

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

- SNIA Education Committee

**Steve Kenniston
Jered Floyd
David Floyer
Michael Fishman**

**Mike Ivanov
Larry Freeman\
Craig Nunes**