




# Large Scale Data Reduction on SSDs

John Hayes  
Pure Storage Inc.



# Customer Data via CloudAssist



# Data reduction is attractive

- Bring cost inline with legacy storage media
- Density increases per RU
- Power saved per TB
- Write less to increase flash longevity

# What is data reduction?

## Included

- Compression
- Deduplication
- Metadata

## Excluded

- Redundancy
- Thin Provisioning
- Snapshots

# Application: Databases

- Oracle, SQL Server, Sybase: 3.5:1
- Oracle w/ Table compression: 2.1:1
- Sybase IQ: 1.1:1
  
- Almost all compression – except when database is cloned.



# Application: Virtualization

- Virtual server: 4.8:1
- VDI: 9.5:1
- Exchange: 2.0-4.5:1
- Compression is 1.5:1 rest is dedup

# Performance of Offloading

- Updates to Oracle with table compression are half the speed
- Client-side gzip runs 1/5<sup>th</sup> the speed of uncompressed writes
- Thin provisioned VMs require VMFS locks



# Dedup ticker on purestorage.com

DEDUPE  
TICKER

Last Updated: 12:50:01AM PDT

20.00 5.14 3.75 3.43 7.75 8.52 8.19 18.99 4.7

6.45:1

Average Data Reduction Rate  
(Deduplication + Compression Only)

13.09:1

Average Total Reduction  
(Including Thin Provisioning)





## Top application are reducible

- Data reduction software is a 3-5 year hardware advance
- It's here to stay!