



**SNIA Emerald™**  
**COM benefits – some initial data**



# SNIA Forward Looking Information Disclosure Statement



This SNIA presentation as part of the industry EPA ENERGYSTAR Data Center Storage Stakeholders Meeting November 18 2015 may include timetables, roadmaps, new technologies entering the mainstream, predictions, estimates or other information that might be considered forward-looking. While these forward-looking statements represent our current judgment on what the future holds, they are subject to risks and uncertainties that could cause actual timeframes and results to differ materially. Readers are cautioned not to place undue reliance on these forward-looking statements, which reflect our opinions and best effort planning only as of the date of this presentation. Please keep in mind that we are not obligating ourselves to revise or publicly release the results of any revision to these forward-looking statements in light of new information or future events. Throughout the discussion in the delivery of this presentation, we will attempt to present some important factors relating to the topic that may affect our estimates and predictions.

# Capacity Optimization Methods (COMs)

COMs are software techniques that allow more data to be stored on less space (i.e., use fewer disks and less power)

COM	Description	Potential Space Savings
RAID 5/6, Erasure Encoding	Data integrity w/o using full copies	When replacing Mirroring (RAID 1): ~40%
Thin Provisioning	Virtualize allocation of Storage	Take from 30% utilization (legacy) to ~ 80%
Deduplication	Eliminate duplicate files or data chunks	Can be 25-40% primary, up to 50% secondary
Compression	Reduction of data sets within files or blocks	2:1 compression done routinely
Delta Snapshots	Reconstruct different versions of a file w/o multiple full copies	Large savings possible when change delta is small

# COM requirements & testing

- Currently, the Emerald Spec defines simple pass/fail tests that verify the existence and activation of particular COMs
- Currently, DCS ENERGY STAR requires that
  - ◆ Parity RAID is enabled in the SUT
  - ◆ 1 out of 4 specific COMS (for Online 3 & 4) is available as a selectable feature and be verified for existence, but they are to be disabled during active metric testing
- There has been some test characterization of active COMs, but more is needed to establish valid methods that account for their space-saving efficiency in active operation

# Energy Efficient Storage: SNIA Emerald Program Specific COM Examples (70/30 R/W random)



COM Type	Raid Parity	Data Compression
COM Method	Interleaved Data and Parity (R5)	Grouped Data Patterns
Baseline Config	Raid 1 Mirrored and Striped	Uncompressed
Baseline Capacity (GB)	36,000	36,000
Baseline Power (W)	4046	4046
Baseline (GB/W)	8.90	8.90
Baseline Perf. (IOPs)	57,750	57,750
Baseline (IO/W)	14.27	14.27
Optimized Capacity (GB)	45,000	72,000
Optimized (GB/W)	11.12	17.79
<b>Delta (GB/W)</b>	<b>+25%</b>	<b>+100%</b>
Opt. Power (W)	4046	4046
Opt. Perf (IOPs)	37,310	115,500
Optimized (IO/W)	9.22	28.55
<b>Delta (IO/W)</b>	<b>(35%)</b>	<b>+100%</b>

# Observations

Based on limited data set

- COMs can improve the idle capacity efficiency metric, but when subjected to an active workload the resultant [performance/W] metric may or may not improve
- Parity RAID and Compression are low-hanging fruit for COM benefits. It gets more complex.....

# Questions ?

