# Overview: SNIA Emerald Measurement

Herb Tanzer

SNIA Emerald™ Training

*SNIA Emerald Power Efficiency Measurement Specification*, for use in EPA ENERGY STAR®

July 14-17, 2014

SNIA™ Green Storage Initiative

SNIA Emerald™

# Emerald Measurement Overview

- Standard measurement of power efficiency of storage systems
- Market taxonomy which classifies storage systems in terms of operational profile and supported features
  - Includes spinning disks, tape, and solid state drives
- Measuring the ratio of useful storage work to the power required to do this work
- Useful Storage Work
  - Store data (Ready Idle)
    - Data is not moving on or off the storage system
  - Move data on and off the system (Active)
    - Random access to the data on the storage system (IO/s)
    - Sequential access to the data on the storage system (MB/s)

# Emerald Measurement Metrics

- Primary metrics ratios of performance / watt
  - Random access (Transactional) of the data per unit of power
    - Input Output per Second per Watt (IOPS/W)
  - Sequential access (Streaming) of the data per unit of power
    - Mebibyte per Second per Watt (MiBPS/W)
  - Storage Capacity per unit of power
    - Gigabyte per Watt (GB/W)
- Secondary metrics
  - Capacity Optimization verification, i.e. existence test
    - Five techniques that reduce the number of storage devices to store the same amount of data thus reducing the power required to store the data

# Test Sequence for Online & Near-Online

| Test | Overview |
|---|---|
| Pre-fill | Provide an initial working data set; exit when 56% filled |
| Conditioning | Provide a uniform initial condition for subsequent measurements<br>Duration is 12 hours minimum |
| Active | A sequence of 5 uninterrupted IO profile phases<br>Each phase shall last 40 minutes minimum |
| Ready Idle | No foreground IO<br>Measure average power (Watts) for 2 hours minimum |

- The above tests are run in an uninterrupted sequence

| Test | Overview |
|---|---|
| Capacity Optimization | Heuristics to validate existence of: Read-only & Writeable Delta Snapshots, Thin Provisioning, Data Deduplication, Parity RAID, Compression |

# Pre-fill Test

- Pre-fill to a minimum of 56% of physical formatted storage with SeqW IO pattern

| IO Profile | Read/Write Percentage | IO Intensity | Access Pattern | Data Pattern |
|---|---|---|---|---|
| Sequential Write | 0/100 | 100 | Sequential | 2:1 compression |

- Data pattern is 2:1 compression, using gzip

# Conditioning Test

- IO profile is Hot banding

| IO Profile | Read/Write Percentage | IO Intensity | Access Pattern |
|---|---|---|---|
| Hot banding | See Table 11 | 100 | See Table 11 |

- Optional data migration phase for tiered storage deployment
- Collect data in 1-minute intervals: size in bytes of each IO, # of IOs issued, average RT, average power
- Duration is 12 hours minimum; during the final 4 hours the average RT shall not exceed 20ms (2.0.2 says 30ms – typo)

# Active Test

- Multiple IO profile test phases

| | IO Profile (Test Phase i) | IO Size (KiB) | Read/Write Percentage | IO Intensity | Transfer Alignment (KiB) | Access Pattern |
|---|---|---|---|---|---|---|
| 1 | Hot Band Workload (i=HB) [a] | See Table 11 | See Table 11 | 100 | See Table 11 | See Table 11 |
| 2 | Random Write ( i=RW) | 8 | 0/100 | 100 | 8 | Random |
| 3 | Random Read ( i=RR) | 8 | 100/0 | 100 | 8 | Random |
| 4 | Sequential Write ( i=SW) | 256 | 0/100 | 100 | 256 | Sequential |
| 5 | Sequential Read ( i=SR) | 256 | 100/0 | 100 | 256 | Sequential |

[a] Near-Online system hot band workload may require further review.

- Collect data in 1-minute intervals: size in bytes of each IO, # of IOs issued, average RT, average power, Operations rate (IO/s, MB/s)
- Any consecutive 30 minute measurement interval must be stable (I.e., each run has to achieve stability)
- Each 1-minute average RT (random & Hot Band only) must be ≤ 80ms
- Each 30-minute average RT (random & Hot Band only ) must be ≤ 20ms

# Test Sequence for Removable Media Library

| Test | Overview |
|---|---|
| Conditioning | Provide a uniform initial condition for subsequent measurements<br>Two IO profiles test phases (SeqW & SeqR, 256KiB, sequential access) each lasting 7 minutes minimum<br>Collect data in 1-minute intervals: size in bytes of each IO, # of IOs issued, average RT to complete an IO, average power |
| Active | Two IO profiles test phases (SeqW & SeqR, 256KiB, sequential access)<br>Collect data in 1-minute intervals: average data rate (MiB/s), average power.<br>Any consecutive 30 minute measurement interval must be stable |
| Ready Idle | No foreground IO<br>Measure average power (Watts) for 2 hours minimum |

- The above tests are run in an un-interrupted sequence
- Pre-fill and COMs are N/A for this taxonomy
- Virtual Media Library test sequence is the same (different media)

# Primary Metrics Reporting

| Active Test Phase (Metric) | Online Near-Online | Removable Media Library | Virtual Media Library |
|---|---|---|---|
| Hot Band (IO/s/W) | X | | |
| Random Read, (IO/s/W) | X | | |
| Random Write (IO/s/W) | X | | |
| Sequential Read (MiB/s/W) | X | X | X |
| Sequential Write (MiB/s/W) | X | X | X |
| Ready Idle (GB/W) | X | X | X |

# Secondary Metrics Reporting (Online & Near-Online only)

Heuristics are run to validate the existence and activation of a particular COM, and is a simple pass (1) / fail (0) test

| Capacity Optimization Method (COM) | Metric |
|---|---|
| Delta Snapshots, read-only | 1 or 0 |
| Delta Snapshots, writeable | 1 or 0 |
| Thin Provisioning | 1 or 0 |
| Data Deduplication | 1 or 0 |
| Parity RAID | 1 or 0 |
| Compression | 1 or 0 |