# New Joint TGG/SNIA Whitepaper WP#77---Energy-Efficient-Data-Center-Storage Data Analysis of Block IO EnergyStar DCS Product Reports

Contacts:

TheGreenGrid : Jay Dietrich

SNIA: Wayne Adams

SNIA Emerald™ Training

*SNIA Emerald™ Power Efficiency Measurement Specification*

*Version 3.0*

February-March 2018

SNIA Emerald™

# How to download Whitepaper as a TGG Guest User

- Visit URL
  - https://www.thegreengrid.org/en/resources/library-and-tools/469-WP#77---Energy-Efficient-Data-Center-Storage
- Follow directions to setup your Guest Account
  - https://www.thegreengrid.org/user/register
- Get your Account Credentials
- Return to the URL
  - https://www.thegreengrid.org/en/resources/library-and-tools/469-WP#77---Energy-Efficient-Data-Center-Storage
- Sign-in
  - https://www.thegreengrid.org/user/login?destination=resources/library-and-tools/469-WP
- Download

# New Whitepaper
# SNIA Emerald Block IO Data Analysis


## ENERGY EFFICIENT DATA CENTER STORAGE:

## THE UTILITY OF THE EMERALD ANALYSIS TOOL FOR THE ASSESSMENT OF STORAGE PRODUCT ENERGY EFFICIENCY

February 9, 2018

EDITOR(S):

Jay Dietrich, IBM

CONTRIBUTOR(S):

Don Goddard: Net App

Rona Newmark: Dell (retired)

Patrick Stanko: SNIA Green Storage Technical Working Group

Herb Tanzer: SNIA Green Storage Technical Working Group

Gary Verdun: Dell

Chuck Paridon: SNIA Green Storage Technical Working Group

Dave Thiel: SNIA Green Storage Initiative

The Green Grid Emerald Analysis Working group acknowledges the Storage Networking Industry Association (SNIA) Green Storage Initiative (GSI) and Green Storage Technical Working Group (GTWG)'s strong collaboration with Emerald Analysis WG and significant contributions in gathering and analyzing the data presented in this paper.

# Table of Contents

SNIA GSI | GREEN STORAGE

## Table of Contents

# Comprehensive Data Analysis Looking for Patterns and Correlations Per Taxonomy class, Drive type, etc…

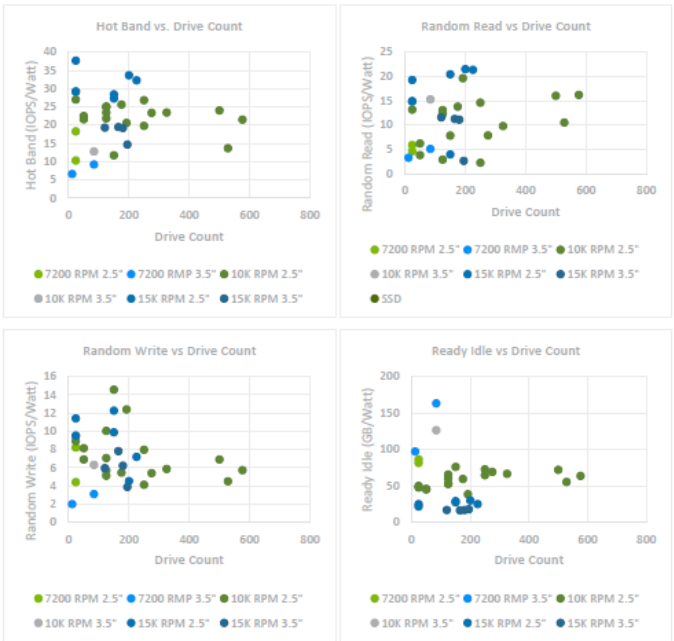


Figure 19: Online 4 Transactional Metrics vs. Drive Count without All Flash Array Data

| Label | Hot Band Workload Test (IOPS/W) | Random Read Workload Test (IOPS/W) | Random Write Workload Test (IOPS/W) | Ready Idle Workload Test (GB/W) |
|---|---|---|---|---|
| FAM34 2.5 7200k 1000GB #Dr24 | 10.3 | 4.6 | 4.4 | 86.5 |
| FAM37 2.5 7200k 1000GB #Dr24 | 18.3 | 5.9 | 8.2 | 81.4 |
| FAM34 3.5 7200k 2000GB #Dr12 | 6.6 | 3.3 | 2 | 97.1 |
| FAM35 3.5 7200k 2000GB #Dr84 | 9.2 | 5.1 | 3.1 | 163.1 |
| FAM34 2.5 10000k 1000GB #Dr24 | 29.2 | 14.9 | 9.5 | 49.9 |
| FAM37 2.5 10000k 1000GB #Dr24 | 27 | 13.2 | 8.9 | 47.7 |
| FAM7 2.5 10000k 600GB #Dr50 | 21.6 | 3.85 | 6.9 | 45.8 |
| FAM9 2.5 10000k 600GB #Dr50 | 22.5 | 6.21 | 8.13 | 45.2 |
| FAM2 2.5 10000k 600GB #Dr125 | 23.48 | 12.48 | 5.67 | 52.8 |
| FAM7 2.5 10000k 600GB #Dr125 | 21.83 | 2.91 | 5.11 | 65.6 |
| FAM9 2.5 10000k 600GB #Dr125 | 25.12 | 13.1 | 7.04 | 60.1 |
| FAM24 2.5 10000k 600GB #Dr125 | 24.91 | 12.3 | 10.03 | 53 |
| FAM19 2.5 10000k 600GB #Dr150 | 11.75 | 7.85 | 14.54 | 75.98 |
| FAM2 2.5 10000k 600GB #Dr175 | 25.62 | 13.81 | 5.46 | 59.3 |
| FAM7 2.5 10000k 600GB #Dr250 | 19.8 | 2.25 | 4.12 | 72.7 |
| FAM24 2.5 10000k 600GB #Dr250 | 26.84 | 14.63 | 7.96 | 64.6 |
| FAM9 2.5 10000k 600GB #Dr275 | 23.38 | 7.9 | 5.39 | 69 |
| FAM2 2.5 10000k 600GB #Dr325 | 23.48 | 9.82 | 5.85 | 66.4 |
| FAM24 2.5 10000k 600GB #Dr500 | 24.03 | 16.04 | 6.89 | 71.9 |
| FAM44 2.5 10000k 600GB #Dr528 | 13.67 | 10.52 | 4.5 | 55.4 |
| FAM42 2.5 10000k 600GB #Dr576 | 21.49 | 16.21 | 5.71 | 63.6 |
| FAM18 2.5 10000k 900GB #Dr192 | 20.63 | 19.63 | 12.39 | 38.6 |
| FAM35 3.5 10000k 0GB #Dr84 | 12.8 | 15.3 | 6.3 | 126.2 |
| FAM34 2.5 15000k 0GB #Dr24 | 29.2 | 14.9 | 9.5 | 25 |
| FAM37 2.5 15000k 0GB #Dr24 | 37.7 | 19.3 | 11.4 | 21.7 |
| FAM7 2.5 15000k 300GB #Dr150 | 27.28 | 3.98 | 9.87 | 29 |
| FAM9 2.5 15000k 300GB #Dr150 | 28.42 | 20.45 | 12.24 | 27.3 |
| FAM2 2.5 15000k 300GB #Dr200 | 33.68 | 21.52 | 4.53 | 29.9 |
| FAM24 2.5 15000k 300GB #Dr225 | 32.29 | 21.39 | 7.18 | 25 |
| FAM9 3.5 15000k 300GB #Dr120 | 19.32 | 11.61 | 5.95 | 16.9 |
| FAM2 3.5 15000k 300GB #Dr165 | 19.48 | 11.26 | 7.8 | 16.5 |

# Conclusions and Recommendations

SNIA GSI | GREEN STORAGE

## 9. CONCLUSIONS AND RECOMMENDATIONS

### Conclusions

There is not enough data to define firm energy efficiency thresholds using the Emerald test results across taxonomy boundaries. Within a given taxonomy, the metrics should be segregated by drive type. Segregating by drive type results in insufficient data to establish thresholds. However, the data to date has shown some interesting characteristics.

Idle Metric:

- Idle metric does not correlate to either the transactional metrics or the streaming metrics. In particular, the idle metric shows more of an inverse correlation to the transactional metrics.
- Idle metric in spinning disks of a given speed is directly related to drive capacity. To improve idle metric, one only needs to change to a higher capacity drive. Drive capacities increase periodically, so one can get an improvement by waiting for the next capacity to be released. For a given family of drives, the power increase is negligible for a capacity increase.
- Idle power in SSDs is a function of the capacity and the amount of activity performed by the devices during "idle". In general, the more activity during "idle" (i.e., wear leveling, housekeeping, etc.) the higher the reliability of the device as idle activity is associated with device maintenance tasks. Penalizing SSDs for idle power could lead to lower reliability of the devices.

Drive Types:

- The higher power, lower capacity drives tend to have higher transactional metrics.
- SSDs have higher performance for transactional or streaming read metrics, but may not have better write metrics.
- The lower power, higher capacity drives tend to have higher streaming data metrics. In general, 10 K drives provide better streaming performance than 7.2 K drives for the data analyzed.
- Low power, high capacity drives will always have a better idle metric.
- In general, there is a cost trade-off to be considered for each drive type. SSDs are the most expensive per GB; high performance, lower capacity drives the next most expensive; and lower performance, higher capacity drives the least expensive.

The family data for minimum, optimal, and maximum drive count deviated less than 10% for transactional metrics and less than 20% for sequential metrics. Based on this, we conclude that min/max testing should be eliminated.

### Recommendations

Based on the above, and the data shown in the rest of this paper, it is our recommendation that the following be used for the next version of Energy Star:

Criteria Changes:

- Product must be qualified for operation at ASHRAE A2 levels or higher.
- PSUs must be "gold" or better.
- The minimum number of COMs should be increased by 1 for Online 3 and Online 4 products.
- Inlet temperature reporting and time stamping should be required.

Testing Requirement:

- Test the best foot forward for the system and post that result with disclosure of the configuration for one workload.
  - Grant Energy Star status for the reported workload for the entire family.
  - The Family is the tested controller with all drive types, capacities, and drive counts.

Optional Testing:

- Additional configurations and/or workloads may be tested in order to obtain additional Energy Star postings.

These recommendations are applicable for any storage product energy efficiency requirements or programs being consider for market entry or procurement purposes.

# Questions?

SNIA GSI | GREEN STORAGE

- Email emerald@snia.org
- Visit www.sniaemerald.com
- EPA EnergyStar Data Center Storage Product Reports
  - https://www.energystar.gov/productfinder/product/certified-data-center-storage