SNIA Emerald V3.0.1 Update

Chuck Paridon

Member SNIA GSI

SNIA Emerald™ Training

SNIA Emerald™ Power Efficiency Measurement Specification

Version 3.0

January-February 2018

SNIA Emerald V3.0.1 Overview

- The *SNIA Emerald™ Power Efficiency Measurement Specification Version 3.0.1* was formally approved by the SNIA on September 11, 2017.
- The most significant change (relative to Version 2.1.1) is the addition of test procedures for file access storage systems.
- These are based on using SPEC SFS 2014 as the workload driving the tests.
- Work with the EPA to update the Energy Star Data Center Storage Program to include file access devices with testing based on Measurement Specification Version 3.0.1 is in progress.

SNIA Emerald V3.0.1 Pre-work and Reading

SNIA GSI | GREEN STORAGE

Background Assumptions

Assume student knows about Power Meters and how to use them
Assume student is familiar with SNIA Emerald and Block IO
Assume student will reference SFS 2014 training materials
Assume student knows the basics: filesystems
Assume student knows S-flow or its equivalent(s)
Student should identify some self-paced info for filesystems
Student should read SNIA Emerald for Files - specification and how to guide

Tests for File Access Storage Systems

SNIA GSI | GREEN STORAGE

- SPEC SFS 2014 is a file-based I/O benchmark, developed by the Storage Performance Evaluation Corporation (SPEC)
- There are four file system workloads based on SPEC SFS 2014 benchmark
  - Video Data Acquisition (VDA) Workload
    - The workload generally simulates applications that store data acquired from a temporally volatile source (e.g., surveillance cameras).
  - Database (DATABASE) Workload
    - This workload represents the typical behavior of a database.
  - Virtual Desktop Infrastructure (VDI) Workload
    - This workload simulates a steady-state high-intensity knowledge worker in a VDI environment
  - Software Build (SW Build) Workload
    - The software build type workload is a classic meta-data intensive build workload derived from analysis of software builds.

Comparison of Block vs. File Access Workloads

SNIA GSI | GREEN STORAGE

Block Access Workloads

- Hot Band (mixed Streams w Hot Spots)
  - Measurements: IOPS/ Watts
- Random Writes (100% write, 8k block)
  - Measurements: IOPS/ Watts
- Random Reads (100% read, 8k block)
  - Measurements: IOPS/ Watts
- Seq. Writes (100% write, 256k block)
  - Measurements: MB/s / Watts
- Seq. Reads (100% read, 256k block)
  - Measurements: MB/s / Watts

File Access Workloads

- Video Data Acquisition (VDA) Workload
  - Measurements: MB/s* / Watts
- Data Base (DATABASE) Workload
  - Measurements: MB/s* / Watts
- Virt. Desktop Infra. (VDI) Workload
  - Measurements: MB/s* / Watts
- Software Build (SW Build) Workload
  - Measurements: MB/s* / Watts

* these metrics are network thruput

File Access Workload Details

File Access IO Demand Intensity

Workload	Business Metric (LOAD parameter)
DATABASE	DATABASES
SWBUILD	BUILDS
VDA	STREAMS
VDI	DESKTOPS

At least 10 uniformly spaced business metric load values shall be specified for valid test execution

.

File Access Storage Capacity Requirements

DATABASE = 24 GB per DATABASE

SWBUILD = 5 GB per BUILD

VDA = 24 GB per STREAM

VDI = 12 GB per DESKTOP

Comparison of Block vs. File Access Execution

Block Access Test Execution

1. Pre-fill test, which puts data on the product under test;
2. Conditioning test, which assures accurate and reproducible measurements;
3. Active test, the basis for the active metrics;
4. Ready idle test, the basis of the ready idle metric;
5. Capacity optimization test (if defined), the basis of the secondary, capacity optimization metrics.

File Access Test Execution

1. Execution of the four workloads, the basis for the active metrics;
2. Ready idle test, the basis of the ready idle metric;
3. Capacity optimization test (if defined), the basis of the secondary, capacity optimization metrics.

Comparison of Block vs. File Access Schematic

SNIA GSI | GREEN STORAGE

Block Access Test Schematic

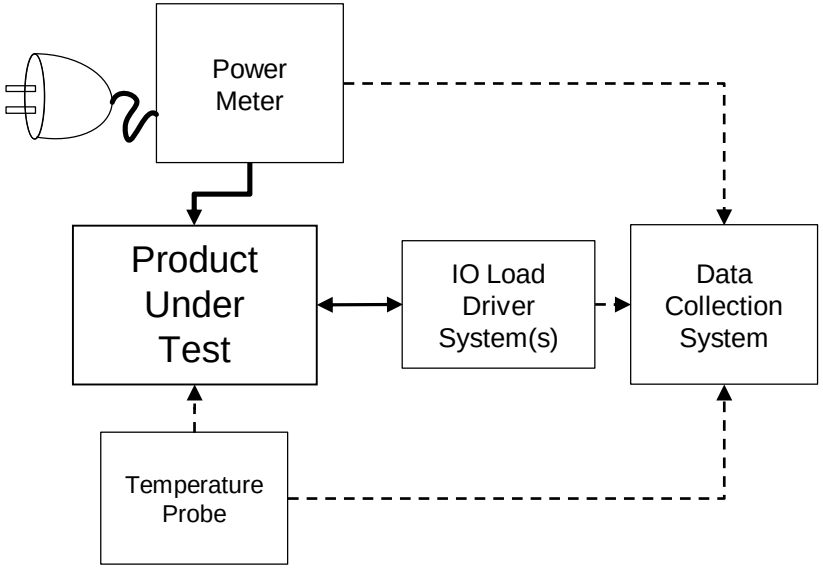

File Access Test Schematic

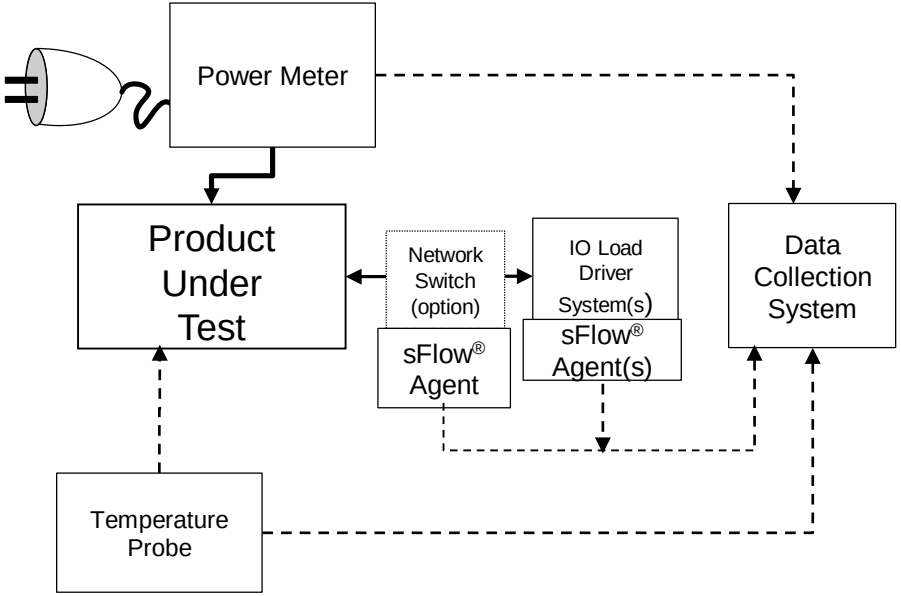

Using alternate IO measurement boundaries

- Block IO accesses are non-buffered.
  - The IO is issued from the driver and the acknowledgement is returned after the IO is committed to stable storage within the storage device.
  - Thus a measure of the rate and response time of the storage subsystem
- File system accesses are buffered several times in the stack.
  - The IO is issued from the driver and the acknowledgement is returned after the IO is written in the client buffer cache.
  - From application space, the response time of the IO may have little to do with the storage subsystem
  - As a result, all IO rate metrics are derived from the Network IO driver or the interconnecting IP switch

Comparative Data Collection Summary

SNIA GSI | GREEN STORAGE

Online Block Access Data Collection

Test	Collection Interval (seconds)		Workload Generator Data Collection		Minimum Duration (minutes)
	Power Meter	Temp Meter	Metric	sample (secs)	
Cond.	60	10	Average Response Time RTA_{sc} (milli-seconds)	60	720
Active	60	10	1) Operations Rate O_i (IO/s or MiB/s) 2) Average Response Time RTA_{sc} (milli-seconds)	60	40
Idle	60	10	N/A	N/A	120

Online File Access Data Collection

Test	Collection Interval (seconds)		Data Collection		Minimum Test Duration (minutes)
	Power Meter	Temp Meter	Metric	sample (secs)	
INIT	60	10	Operations Rate O_i (MiB/s)	10	N/A
Warm-up – per load point	60	10	Operations Rate O_i (MiB/s)	10	5
Active – per load point	60	10	Operations Rate O_i (MiB/s)	10	5
Ready Idle	60	10	N/A	N/A	120

References

1. SNIA Emerald™ program, *SNIA Emerald™ Power Efficiency Measurement Specification,* http://snia.org/tech_activities/standards/curr_standards/emerald
2. U.S. Environmental Protection Agency, *ENERGY STAR® Program for Data Center Storage Specification*, www.energystar.gov/index.cfm?c=new_specs.enterprise_storage
3. SPEC SFS® 2014 User's Guide https://www.spec.org/sfs2014/index.html#userguide
4. File systems –Wikipedia https://en.wikipedia.org/wiki/File_system