Using SPEC SFS® with the SNIA Emerald Program for EPA Energy Star Data Center Storage Program

Industry Stakeholders Update

- Vernon Miller – IBM
- Nick Principe – Dell EMC

Agenda

- Background on SNIA Emerald/Energy Star for block
- Introduce NAS/File test addition; introduce SFS 2014
- Testbed configuration and measurement points
  - Test procedure
- A look at some real data and the derivation of the metrics

Green Preamble

- Increased regulatory and societal pressures to lower energy footprints
- Growing awareness of environmental impact of IT equipment
- Rising energy cost for power and cooling is a large part of the cost of ownership
- Data centers cannot readily add additional power or cooling capacity

http://www.snia.org/emerald/training/July2014
Overview: Green Storage, Energy Star and SNIA Emerald Program

EPA ENERGY STAR® for Data Center Storage

- Energy Star certification program for storage systems.
- EPA partnered with SNIA's Green Storage Initiative (GSI) to develop technical specification and requirements.

http://www.snia.org/emerald
https://www.energystar.gov/products/office_equipment/data_center_storage

SNIA Green Activities

- Green Storage Initiative (GSI)
  - Market green storage and manage the Emerald™ Program
    Research, educate, leverage SNIA resources, provide direction
- Green Technical Working Group (GTWG)
  - Technical body of storage experts developing green storage specifications, white papers, tutorials, technical guidance
  - Develop the SNIA Emerald™ Power Efficiency Measurement Specification (currently 2.1.1) and "how to" User Guide for it
- Emerald™ Program
  - Promote use of the SNIA Emerald™ Specification methodology and test results
  - Help drive green storage decisions for both vendors and customer

http://www.snia.org/emerald/training/July2014
Overview: Green Storage, Energy Star and SNIA Emerald Program

Emerald™ Power Efficiency Measurement Specification – Market Taxonomy

- Classifies storage systems in terms of operational profile and supported features

Attribute	Category					
	Online	Near Online	Removable Media Library	Virtual Media Library	Adjunct Product	Interconnect Element
Access Pattern	Random/ Sequential	Random/ Sequential	Sequential	Sequential		
MaxTTFD (t)	t < 80 ms	t > 80 ms	t > 80 ms t < 5 min	t < 80 ms	t < 80 ms	t < 80 ms
User Accessible Data	Required	Required	Required	Required	Prohibited	Prohibited

Emerald™ Power Efficiency Measurement Specification – I/O and Measurement

- Standard input voltages and datacenter conditions required
- Prefill Test
- PUT Conditioning Test
- Active Test (Vdbench)
  - Hot Band
  - Random Write
  - Random Read
  - Sequential Write
  - Sequential Read
- Ready Idle Test
- Capacity Optimization Test (comgen)

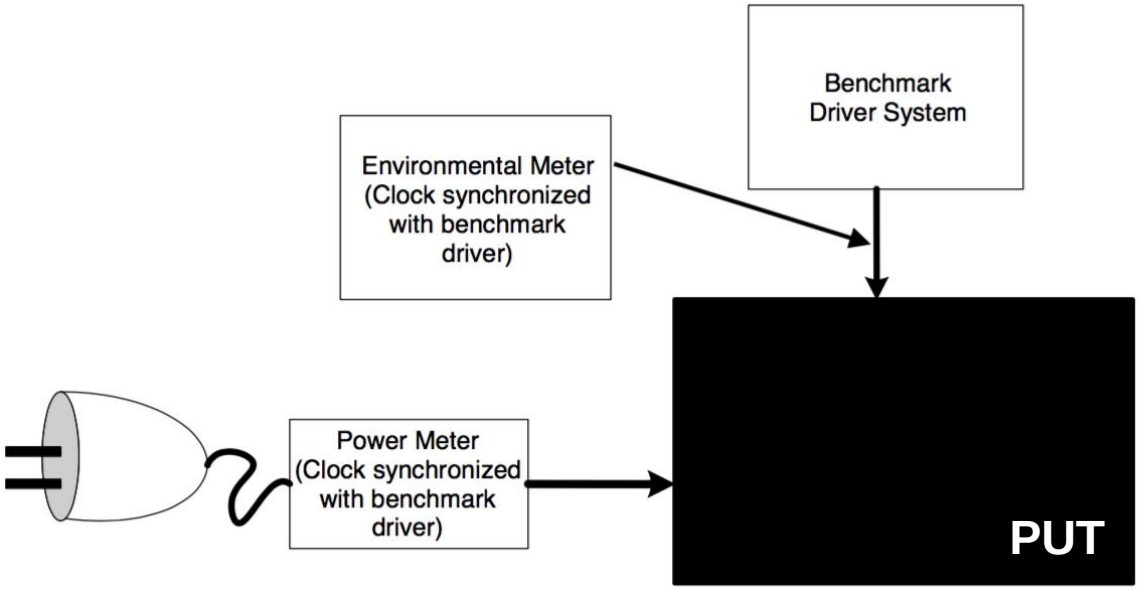

Emerald™ Power Efficiency Measurement Specification – Power Metrics

- 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
    - Six techniques that reduce the number of storage devices to store the same amount of data thus reducing the power required to store the data

http://www.snia.org/emerald/training/July2014
Overview SNIA Emerald Measurement

SNIA Emerald™ Test Data Submission

- SNIA Emerald Program
  - Record results in SNIA Emerald Test Data Report (TDR)
  - Complete Test Submission Order Form
- EPA ENERGY STAR
  - EPA recognized lab must perform tests
  - Obtain certification from an EPA recognized Certification Body

ENERGY STAR Certified
Data Center Storage Visit the Data Center Storage page for usage tips and buying guidelines.
CHANGE product category

< back to results

	IBM - FlashSystem 840 - FlashSystem 840 FlashSystem 840 9840-AE1	IBM - XIV - XIV XIV 2810-214	IBM - V3700 - V3700 V3700 2072-12C	IBM - FlashSystem 900 - FlashSystem 900 FlashSystem 900 9840-AE2
ENERGY STAR Partner:	IBM Corporation	IBM Corporation	IBM Corporation	IBM Corporation
Storage Model Connectivity:	Block I/O	Block I/O	Block I/O	Block I/O
Product Type:	Online 2	Online 4	Online 3	Online 2
Storage Controller Configuration:	Scale-Up Storage	Scale-Out Storage	Scale-Up Storage	Scale-Up Storage
Storage Controller Advanced Data Recovery Type:	RAID 5 with single parity module, automatic rebuild	Proprietary / Grid	RAID 5	RAID 5 with single parity module, automatic rebuild
Capacity Optimized Method Available (COMs):	None	Delta Snapshots	Delta Snapshots	None
Workload Optimization Type:	Transaction	Streaming	Streaming	Transaction
Qualification Range Submission Type:	Fixed Size Qualification Range	Fixed Size Qualification Range	Fixed Size Qualification Range	Fixed Size Qualification Range
Automated Storage Tiering Capable:	No	No	Yes	No
Automated Storage Tiering Enabled in Hardware on Shipment:			No	
Input Power Rolling Average Capability:	No	Yes	Yes	No
Inlet Air Temperature Rolling Average Capability:	No	Yes	Yes	No
Additional Model Information:	,9843-AE1,	,2812-214,	Storwize V3700,2072-12E,	,9843-AE2,

Disclaimer

- The SNIA Emerald specification with file-access support, as represented in this presentation, is *pre-release*; the benchmark framework, workloads, and results and reporting structure are still under internal SNIA review and may change before final release of SNIA Emerald Specification version 3.0.

NAS/File Addition to Specification

- Version 3 of SNIA Emerald Power Efficiency Measurement Specification
  - Addresses both block and file access
  - New workloads and toolkit for file access testing
    - SPEC SFS® 2014
  - New methodology for determining power metrics for file access
  - Expected rollout in 2017

SPEC
Standard Performance Evaluation Corporation

- **The Standard Performance Evaluation Corporation (SPEC)** is a non-profit corporation formed to establish, maintain and endorse a standardized set of relevant benchmarks that can be applied to the newest generation of high-performance computers. SPEC develops benchmark suites and also reviews and publishes submitted results from member organizations and other benchmark licensees
- www.spec.org
- SPEC and SPEC SFS are registered trademarks of the Standard Performance Evaluation Corporation. Additional product and service names mentioned herein may be the trademarks of their respective owners.

Why SPEC SFS 2014?

- SPEC SFS 2014 is a Storage Solution Benchmark
  - Realistic, Solution-based, Industry-standard workloads
    - DATABASE, SWBUILD, VDA, VDI
  - Workloads based on traces, like previous SFS 2008
    - Modern scenarios based on standard solutions
  - Advanced measurement – quality of service
    - Ops and latency don't tell the whole story → business metrics
- Ability to measure broad range of products and configurations
  - Traditional (HDD), Hybrid, All-Flash
- Key reasons SNIA Emerald is using SFS 2014
  - Vendors likely already running SFS 2014 in-house
  - Workloads already agreed upon by multiple vendors
  - Robust workload generator for file access

The SPEC SFS 2014 Workloads

For more details, see:
- SDC 2014 presentation: SPEC SFS 2014: An Under-the-Hood Review
- The SPEC SFS 2014 website http://www.spec.org/sfs2014

- DATABASE
  - Simulates OLTP database consolidation
  - Random/Transactional
- SWBUILD
  - Simulates large software project compilation
  - Metadata intensive – lots of file operations
- VDA
  - Simulates acquisition of streaming data
  - Heavy sequential writes with some random reads
- VDI
  - Simulates heavy steady-state VDI workload
  - Random/Transactional – greater spread of IO sizes

File vs Block Configurations

Block Access

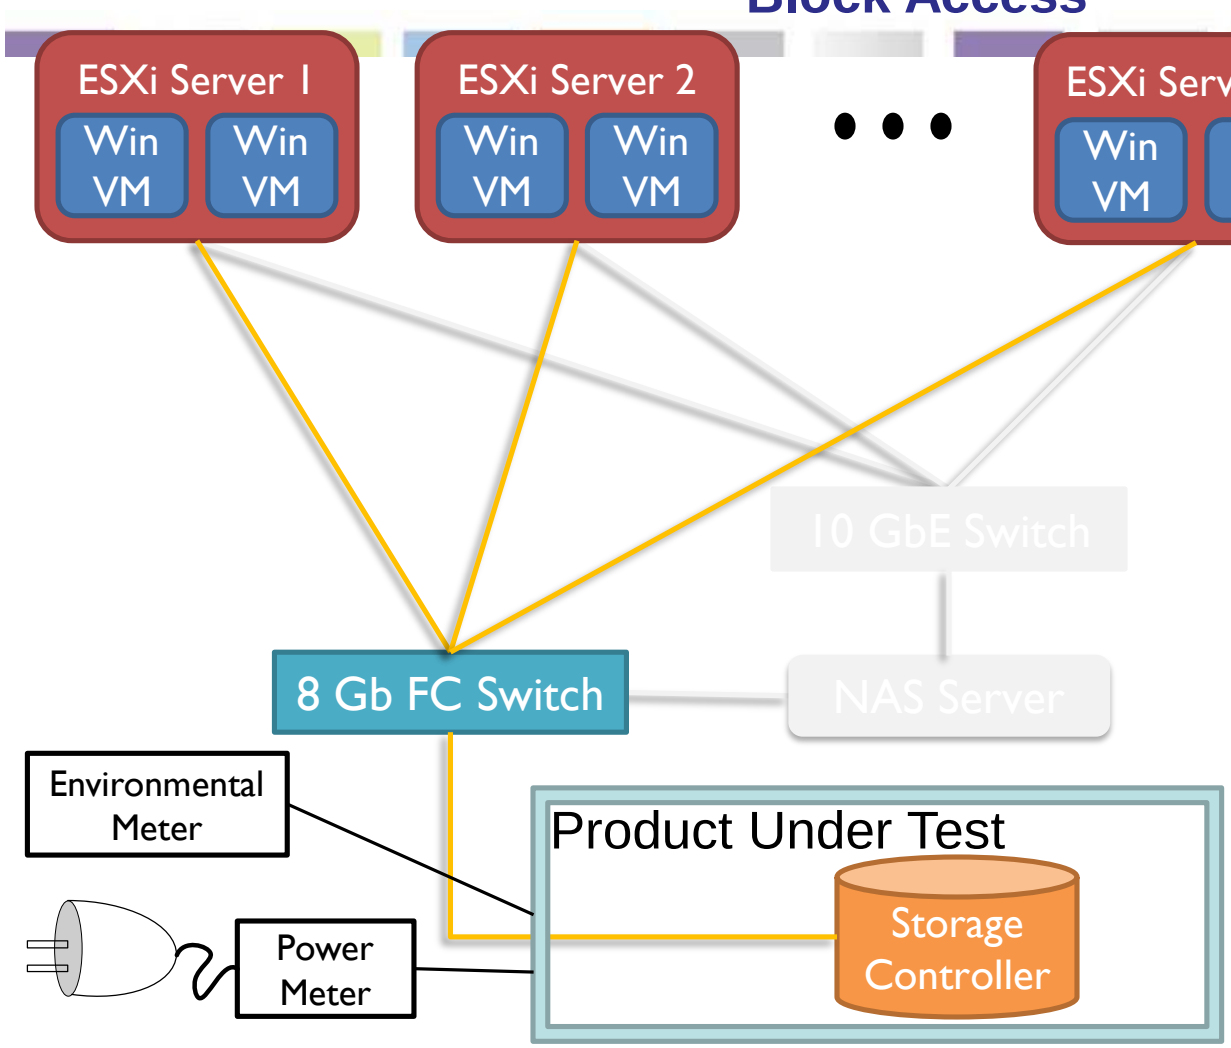

- Block I/O generator: vdbench
  - Custom workloads
- No client side caching
- Minimal impact to power metrics due to load generators

File vs Block Configurations

File Access

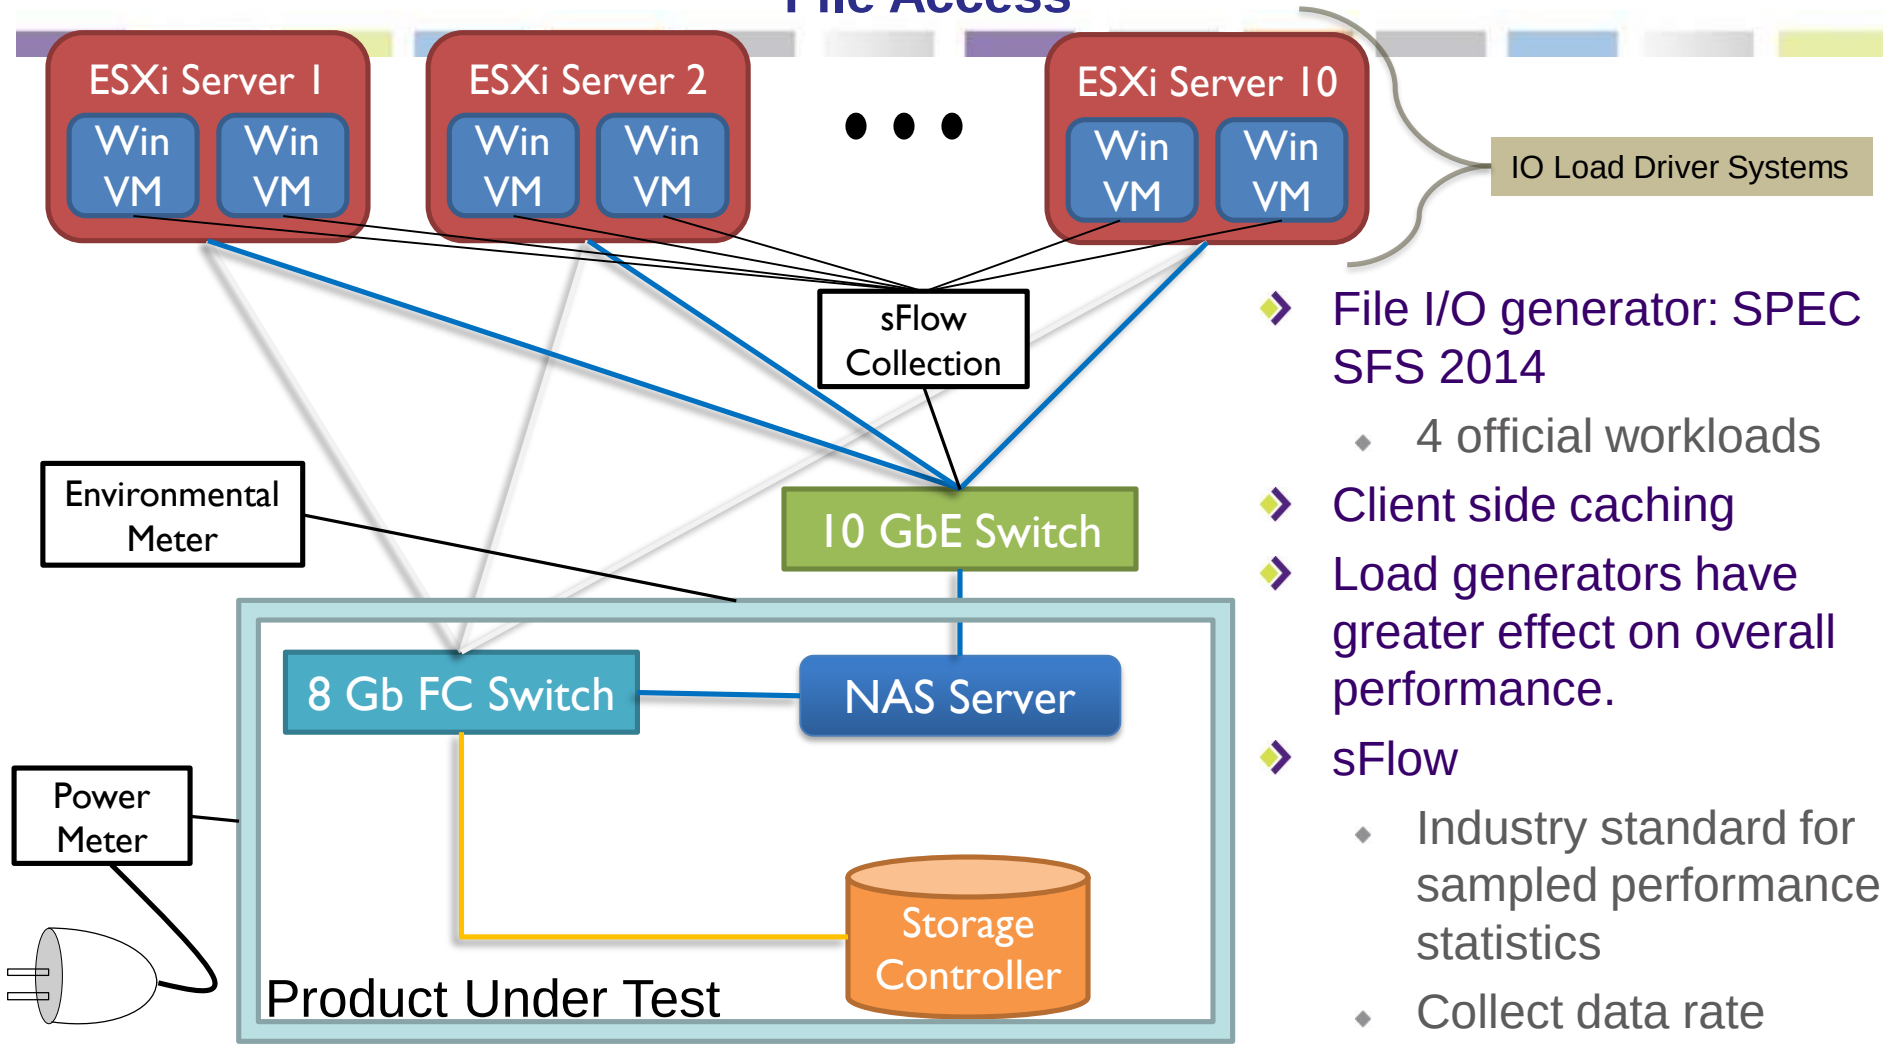

- File I/O generator: SPEC SFS 2014
  - 4 official workloads
- Client side caching
- Load generators have greater effect on overall performance.
- sFlow
  - Industry standard for sampled performance statistics
  - Collect data rate

File vs Block Test Sequences

Block Access

- Pre-fill test, puts data in PUT
- PUT conditioning
- Active test
- Ready idle test
- Capacity optimization test (if defined)

File Access

- Calibration of SFS workloads
- Execution of the 4 SPEC SFS 2014 workloads, in sequence
- Ready idle test
- Capacity optimization test (if defined)

File vs Block Primary Metrics

- Block Access
  - Power efficiency for active phase
    - Hot Band
    - Random Read
    - Random Write
    - Sequential Read
    - Sequential Write
  - Power efficiency for ready idle test phase
- File Access
  - Power efficiency for each workload (MiB/sec/Watt)
    - VDA
    - DATABASE
    - VDI
    - SWBUILD
  - Power efficiency for Ready Idle test phase

File Access Measurement Points

- The SPEC SFS 2014 metrics are only used to calibrate the appropriate load points for each workload
- The Emerald efficiency metric (MiB/s/W) is derived from the data collected by the sFlow collector in front of the Product Under Test
  - SPEC SFS 2014 measures at the application-level
  - SNIA Emerald measures at the product-level

Additional Hardware/Software Requirements for File Access

- In addition to the environmental and power meters common for both Block and File Access Emerald testing, File Access testing requires:
  1. A network switch that supports sFlow
     or
     Host sFlow installed on the load generators
  2. An sFlow collector that can log for extended periods of time

File Access Test Procedure

- The four basic phases of file access testing
  - Calibration
  - Measurement
  - Data Reduction
  - SNIA Emerald Metrics

File Access Test Procedure Calibration

- For each SPEC SFS 2014 workload, find maximum PUT performance
  - Known from existing performance testing
    - Many vendors run SFS 2014 for regression analysis
  - Test to determine as part of Emerald test process
    - Run several SFS 2014 runs, adjusting load points to probe for the maximum valid load point
- Not concerned with maximum efficiency at this point
  - Need to know the whole efficiency curve as load scales up

File Access Test Procedure Calibration Example

- On a new test system, setup SWBUILD to run from 1 to 30 load points, incrementing the load by 1 each step
  - After 20 load points, achieved ops/sec stopped scaling and SFS 2014 was reporting INVALID_RUN
  - At this point, the benchmark was manually terminated

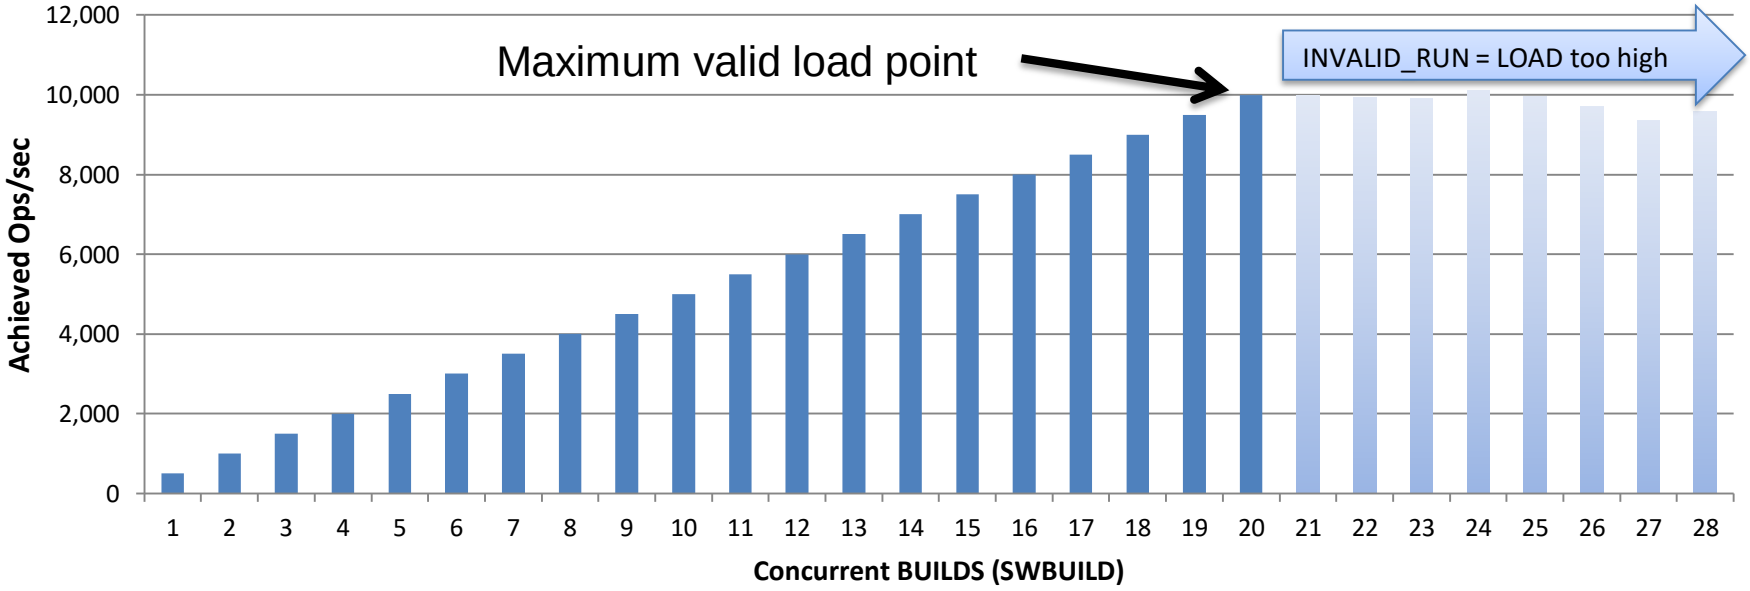

File Access Test Procedure Measurement

- Using maximum valid load from calibration data
  - Run each workload
    - Ten evenly-spaced load points up to the maximum valid load point
    - Collect environmental, power, and sFlow data for each run
  - Ready-idle test
    - Collect environmental and power data while array is idle

File Access Test Procedure Measurement Example

- From the calibration phase, we know the system's maximum valid load point is 20 BUILDS
  - Therefore, for measurement, we ran from 2 through 20 BUILDS, incrementing by 2 BUILDS

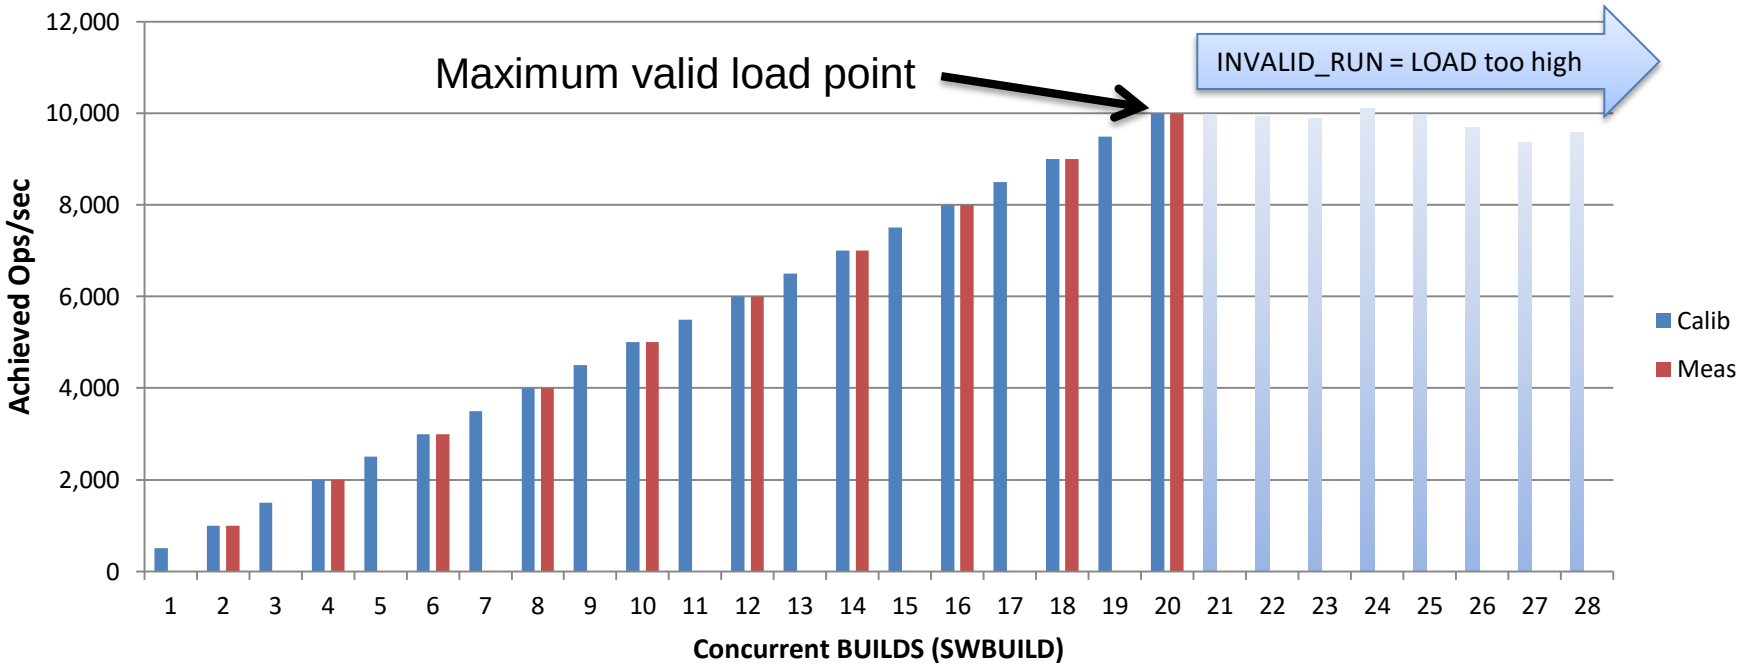

File Access Test Procedure Data Reduction

- Using collected sFlow and Power data
  - For each workload, for the measurement phase only at each load point, calculate
    - Average power
    - Average network throughput
  - Using the average power and network throughput
    - Calculate efficiency metric at each load point for each workload
      - MiB/s/W
  - This process is known as data reduction
    - Time-based data from multiple sources is reduced to calculated metrics per-load point, per-workload
  - An open-source tool is available from SPEC to assist with this process

This process does not describe additional required stability/validity assessment

File Access Test Procedure
Data Reduction Example

- The raw power data log will include data for the whole duration of benchmark execution
  - Only want data from the measurement phase

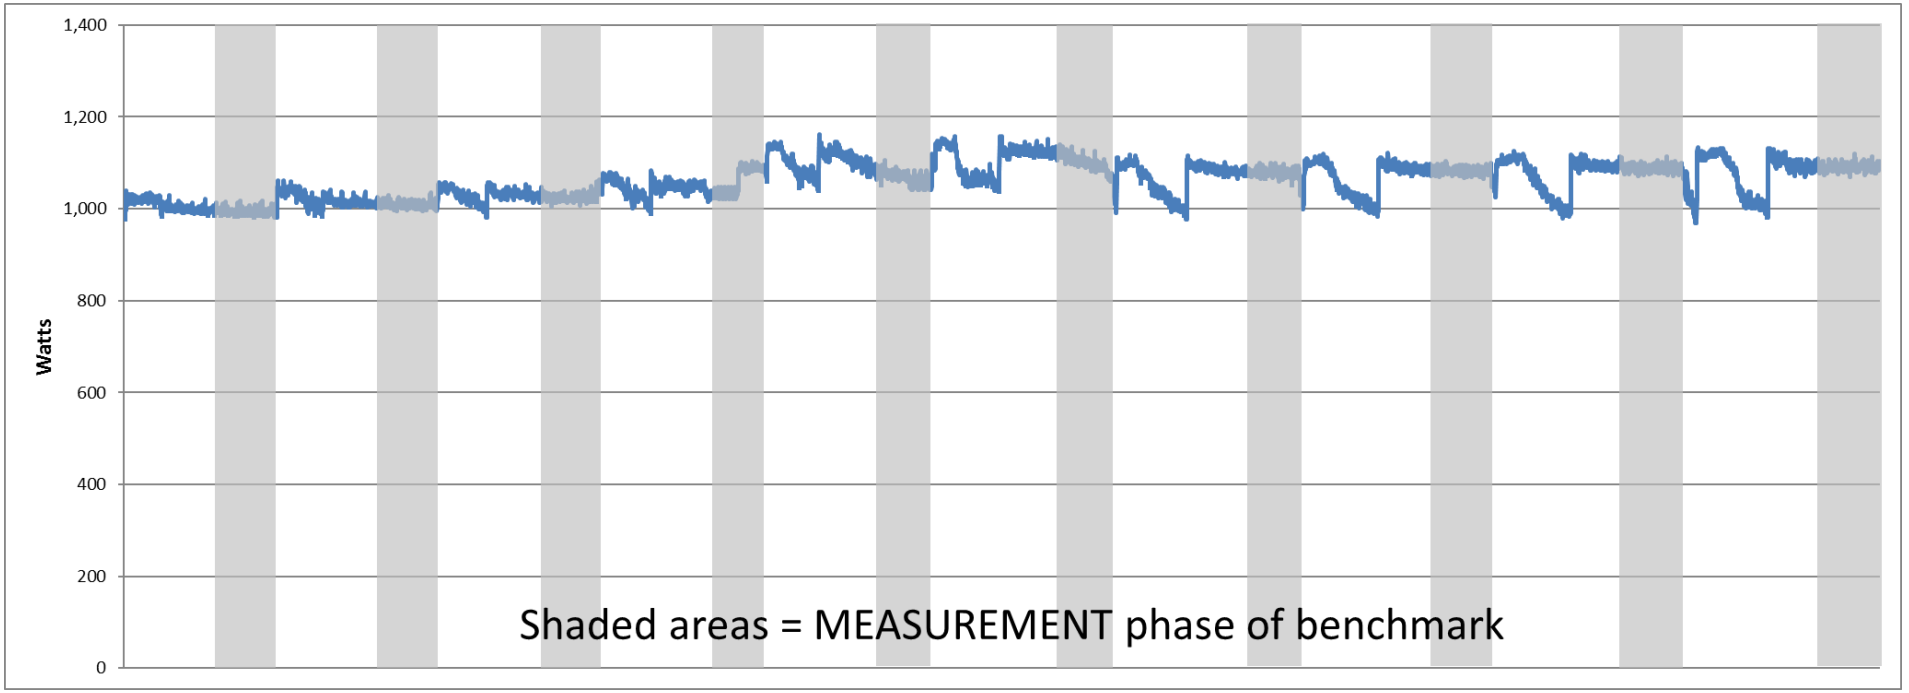

File Access Test Procedure
Data Reduction Example

- Using data reduction techniques, compute average power usage during measurement phase
  - For each load point

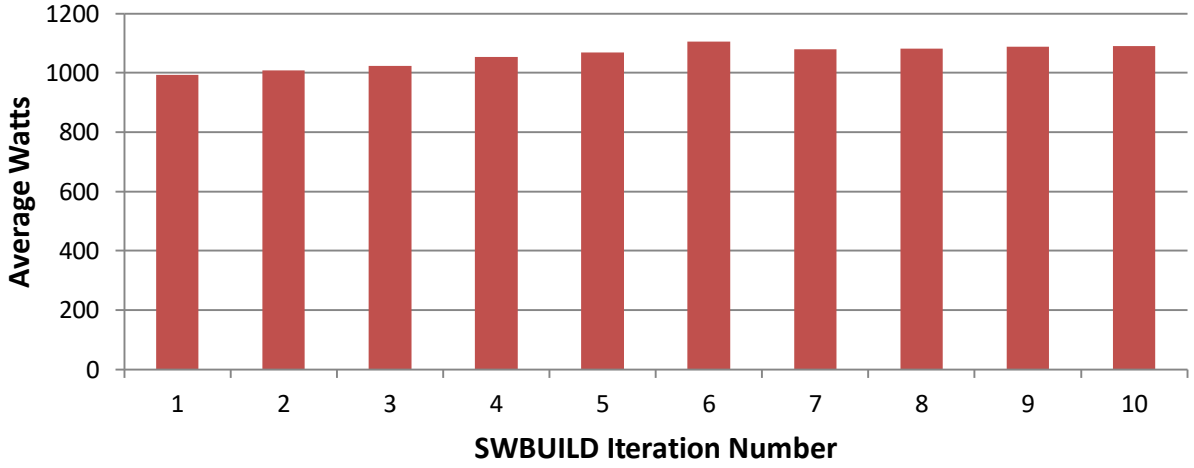

- The same process is used for the sFlow data
  - Average network throughput during measurement phase

File Access Test Procedure
SNIA Emerald Metrics

- Find the “sweet-spot” for all four workloads
  - For each workload, find the highest efficiency metric, using data from the data reduction step
    - Usually, but not necessarily, the highest valid load point

File Access Test Procedure SNIA Emerald Metrics Example

- Example data from another system
  - In a real Emerald run, there would be ten load points for each workload and ready-idle data

EFFICIENCY METRICS Load Point	DATABASE	SWBUILD	VDA	VDI
1	0.01	0.00	0.05	0.01
2	0.01	0.01	0.09	0.01
3	0.02	0.01	0.13	0.02
4	0.02	0.02	0.17	0.03
5	0.03	0.02	0.21	0.03
6	0.03		0.25	0.04
7	0.04		0.29	0.04
8	0.04		0.33	0.05
9			0.37	
10			0.37	

This system demonstrates highest efficiency at the highest valid load point

	DATABASE	SWBUILD	VDA	VDI
Maximum Efficiency (MiB/s/W)	0.04	0.02	0.37	0.05

Metrics reported in Test Data Report

Key Takeaways

- The SNIA Emerald program is adding support for file access storage systems
- SPEC SFS 2014 and its workloads are used to evaluate file access storage systems
- The Product Under Test for SNIA Emerald file access testing only includes the storage array
  - Efficiency metrics are derived from:
    - Network traffic to and from the storage array
      - Not SPEC SFS 2014 metrics
    - Power consumption by the storage array

Additional Resources

- SPEC SFS 2014 Details and Tools
  - SPEC SFS® 2014 Website (User's Guide and Run Rules)
    - http://spec.org/sfs2014/
  - tag2014.py Data Reduction Tool
    - http://spec.org/sfs2014/support/tools/tag2014.tgz
- SNIA Storage Developer Conference Presentations:
  - SDC 2014: SPEC SFS 2014 An Under-the-Hood Review
    - https://www.youtube.com/watch?v=S9RJ_rKc3WI
  - SDC 2015: Application-Level Benchmarking with SPEC SFS 2014
    - https://www.youtube.com/watch?v=4wfeM1q0zHA
  - SDC 2016: Using SPEC SFS with the SNIA Emerald Program for EPA Energy Star Data Center Storage Program
    - https://www.youtube.com/watch?v=7gDgcDYatvM