SNIA Emerald™ Power Efficiency Measurement Specification

Quick Reference Guide for Online (Active) Storage Systems Testing

Version 1.0 Revision 2

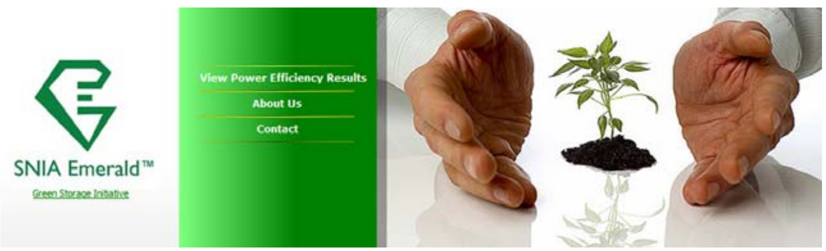

Herb Tanzer - HPE
Carlos Pratt - IBM
Greg Schulz – StorageIO
With Drew Robb

October 29, 2015

About the SNIA

The Storage Networking Industry Association is a not-for-profit global organization, made up of member companies spanning the global storage market. SNIA's mission is to lead the storage industry worldwide in developing and promoting standards, technologies, and educational services to empower organizations in the management of information. To this end, the SNIA is uniquely committed to delivering standards, education, and services that will propel open storage networking solutions into the broader market. For more information, visit www.snia.org.

About the SNIA Green Storage Initiative

SNIA's Green Storage Initiative (GSI) is dedicated to advancing energy efficiency and conservation in all networked storage technologies in an effort to minimize the environmental impact of data storage operations. SNIA's Green Storage activities take place in two separate working bodies, the SNIA Green Storage Technical Working Group (TWG) and the Green Storage Initiative. The TWG is focused on developing test metrics by which energy consumption and efficiency can be measured. The Green Storage Initiative is focused on creating and publicizing best practices for energy efficient storage networking, educating the IT community, and promoting storage–centric applications that reduce storage footprint and associated power requirements.

About the SNIA Emerald™ Program

The SNIA Emerald™ Program is a vendor-neutral, public service to the storage industry and end users that is sponsored and operated by the SNIA GSI. The program supports the use and evolution of the *SNIA Emerald Power Efficiency Measurement Specification* and the publication and use of product test data based on the specification. The measurement procedure and test metrics are documented in the *SNIA Emerald Power Efficiency Measurement Specification*, which was developed, released, and is maintained by the Green Storage TWG under the guidance of the GSI.

The program provides a standardized way of reporting vendor-performed test results that characterize the various aspects of storage system energy usage and efficiency. Those test results—which are based on performance under different workload metrics—are combined with other power-related information about storage systems and their components. The

Program offers a common repository, accessible to the public, where test results are presented in a consistent format that can be used to help determine facility power usage and efficiency for data storage solutions, as well as the costs to operate such systems.

The EPA ENERGY STAR® Data Center Storage Program is based on the methodology defined in the specification and offers another vehicle for publication of product test results created in accordance with the Specification.

About Drew Robb

Drew Robb is an author of "Server Disk Management in a Windows Environment" (CRC Press). He has published hundreds of articles covering different aspects of technology, including servers, storage and engineering along with renewable energy. Born and raised in Scotland, he received a degree in Geology/Geography from the University of Strathclyde, Glasgow.

About Server and StorageIO (StorageIO)

The Server and StorageIO Group (StorageIO) is an independent IT advisory consultancy analyst firm focused on data and information infrastructure related technology, techniques and trends leveraging real-world experience and expertise, visit http://storageio.com.

The information contained in this publication is subject to change without notice. This guide represents a "best effort" attempt by the SNIA Green Storage Technical Working Group to provide guidance to those implementing the *SNIA Emerald™ Power Efficiency Measurement Specification*, and the guide may be updated or replaced at any time. The SNIA shall not be liable for errors contained herein.

Suggestions for revisions to this guide and questions concerning implementation of the *SNIA Emerald™ Power Efficiency Measurement Specification* can be directed (via email) to greentwg-chair@snia.org.

Contents

List of Tables

List of Figures

Quick Reference Guide for Online (Active) Storage Systems Testing

Introduction

This quick reference guide will help you understand what needs to be done to test online storage systems using the *SNIA Emerald™ Power Efficiency Measurement Specification* (referred to within this document as simply the *Measurement Specification*). Workload metric data results are generated and collected by running the Emerald™ Vdbench script on a system test initiator (STI) host against a storage system under test (SUT).

These test results can be used in support of activities besides SNIA Emerald. These activities include supporting submissions for EPA Energy Star Testing and certification, along with general operational data center and workload profiling.

If you are an experienced *Measurement Specification* tester, this quick reference will help you in those cases in which you may want to verify a step or procedure—or clarify an uncertainty. If you are new to the *Measurement Specification*, this is a great place to start, especially if this is the first time you are benchmarking or profiling an online storage system.

We hope this reference will help you to quickly begin your benchmarking activity and navigate through the many details covered in the *Measurement Specification* and related *User Guide*. If you believe that we have an error or if we can do anything to improve this reference, please contact SNIA at greentwg-chair@snia.org.

Who Should Read this Document?

This quick reference guide is intended to be used by any person who wants to understand the power efficiency of online storage systems and their corresponding performance activity, based on the *Measurement Specification*.

The Inexperienced

Any person who understands data center storage but has no experience with the *Measurement Specification* is strongly encouraged to start with this guide before reading the *Measurement Specification* or the related *User Guide*. The performance and power metrics discussed in the *Measurement Specification* will be better understood once the measurements are collected and reviewed for the first time. As the reader goes through this quick reference, key portions of the *Measurement Specification* will be referenced. Consult referenced items to learn more about the *Measurement Specification*.

The Experienced

This quick reference is also intended to be used by any storage benchmark tester. In some cases, the benchmark tester may not be interested in the complete *Measurement Specification* including energy and environment profile, but only in how to obtain the performance metrics of a certain workload.

Some Assumptions

Although this reference assumes the reader knows what an online block-oriented active storage system is, here are some quick reminders:

- The storage-system is supporting active online application workloads, e.g., the storage is busy doing I/O (input/output) operations.
- The data access pattern can be random, sequential or a combination of both.
- The time to access the first piece of data, after the system has being waiting idle for a while, is equal to or less than 80 milliseconds (ms).
- The data must be always accessible to the user.
- Normally, online storage-systems use Hard Drives (HDD), Solid State Drives (SSD), flash cards (flash), or a combination (hybrids).
- The reader understands how to attach and configure a server to use external storage.
- The reader has a basic understanding of power measurement along with associated safety risks.
- The reader understands the basic concepts of server and storage I/O performance metrics along with benchmarking results such as:
  - IOPS
  - MiBPS
  - Idle

After You Read This… SNIA Emerald™ Resources

Once you have read this quick reference guide, you can learn more details by reading the complete *Measurement Specification* and its *User Guide*—and by downloading Vdbench scripts. Access this material from the Documents and Downloads web page of the SNIA Emerald™ website (http://www.snia.org/emerald/download/Spec_v2.1).

The following are links to documents, forms, scripts, and source code and workload tools associated with running *Measurement Specification* tests for Active Storage systems.

- This Quick Reference Guide
  http://www.snia.org/emerald/download/Spec_v2.1
- Emerald™ Program™ Website
  http://www.snia.org/emerald
- SNIA Emerald™ taxonomy
  http://www.snia.org/emerald/taxonomyoverview
- Emerald™ Documents and Downloads (including the Vdbench script)
  http://www.snia.org/emerald/download/Spec_v2.1
- *SNIA Emerald™ Power Efficiency Measurement Specification* (PDF)
  http://www.snia.org/tech_activities/standards/curr_standards/emerald

- Test Workload Script file (TXT)
  http://www.snia.org/emerald/download/Spec_v2.1
- COM Test Dataset Generator program (C source code)
  http://www.snia.org/emerald/download/Spec_v2.1
- Test Data Report Template (Excel worksheet)
  http://www.snia.org/emerald/download/Spec_v2.1
- *User Guide for the SNIA Emerald™ Power Efficiency Measurement Specification*
  http://www.snia.org/emerald/download/Spec_v2.1
- Vdbench download, user guide and discussion groups
  http://www.snia.org/emerald/download/Spec_v2.1
- SNIA Green Storage Initiatives (GSI) Home Page
  http://www.snia.org/forums/green
- SNIA Emerald™ Training Material (presentations and other material)
  http://www.snia.org/emerald/training

We *encourage* the reader to attend the yearly Emerald™ training offered by SNIA. Check the SNIA Emerald™ web site (www.sniaemerald.com) for more details and for videos and presentations of past training sessions.

Glossary

BFF	**B**est **F**oot **F**orward
COM	**C**apacity **O**ptimization **M**ethod(s)
DAQ	**D**ata **Acq**uisition System
HDD (Hard Disk Drive)	[Storage System] A non-volatile, randomly addressable, re-writable data storage device. This definition includes rotating magnetic and optical disks and *solid-state disks*, or non-volatile electronic storage elements. It does not include specialized devices such as *write-once-read-many (*WORM*)* optical disks, nor does it include so-called *RAM disks* implemented using software to control a dedicated portion of a host computer's volatile random access memory.
Hot-Band	Workload to simulate active running online storage including cache
IDLE	[Storage System] A state in which a storage system is serving no user-initiated I/O requests, but is ready to service them upon arrival with normal latency.
IOPS	**I**nput **O**utput operations **P**er **S**econd
Measurement Specification	Short name for the *SNIA Emerald™ Power Efficiency Measurement Specification, available at* http://www.snia.org/tech_activities/standards/curr_standards/emerald
MiB/s	Mebibytes per second
Pre-Fill	Target storage SUT space capacity is written to pre-condition devices prior to testing.
SAN	**S**torage **A**rea **N**etwork
SSD (Solid State Drive)	[Storage System] A disk drive whose storage capability is provided by solid state storage. Form factors and interfaces for solid state drives are typically the same as for a traditional disk drives (e.g. HDD).
STI	**S**ystem **T**est **I**nitiator – Server(s) where the test workload scripts are run from causing workload to be done by the storage target SUT.
SUT	**S**ystem **U**nder **T**est – The storage system being tested.
SUT-Conditioning	Preparing the SUT for testing by writing to all areas to be used.

For an expanded Glossary of terms refer to the SNIA Dictionary found at:
http://www.snia.org/education/dictionary.

Getting Started

Now that we have stated the assumptions as well as who should read this quick-reference guide, let's take the next step on your path to submitting a SNIA Emerald™ data submission.

You will need to decide what you are going to test and if you will be doing a simple online activity workload test (or both), or if you will be doing a Capacity Optimization Method (COM) approach based on your storage SUT capabilities.

General Steps

General steps, shown in Figure 1, are:

1. Determine your taxonomy classification and product/family details.
2. Determine desired test configurations.
3. Collect, configure and prepare items you must have.
4. Decide on type of workload tests that you will run.
5. Find the Best Foot Forward (BFF); select data points; run some practice and tuning tests, review results.
6. Set up and run your actual test runs with data collection.
7. Review data collected. If applicable, fix and re-run.
8. Complete and submit SNIA Emerald™ Test Data Results (TDR) form, if that is your objective.

Figure 1 - Using This Guide and SNIA Emerald™ Workflow

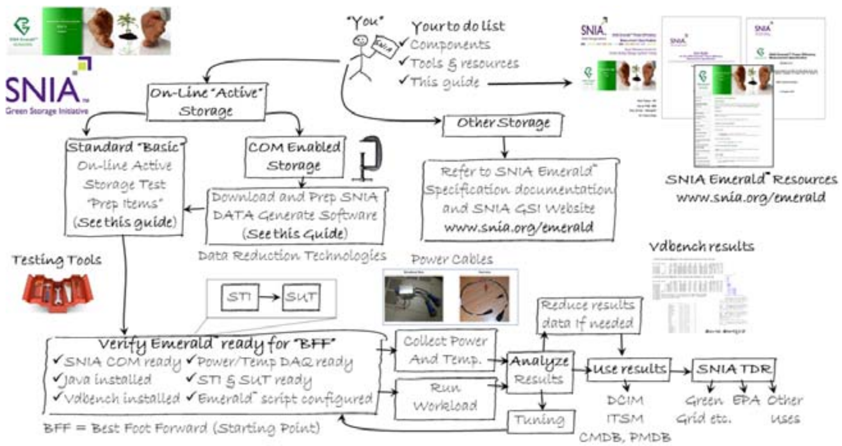

Organize your test environment by including applicable servers and software for running the workload scripts. In addition to servers and software, you will need applicable cabling along with monitoring equipment for electrical power and environment temperatures. While it should be obvious, let's state the obvious in that you will also need to have an online storage system to test (SUT) and configure all the above items for the Measurement Specification tests and subsequent results.

Prior to running your actual Emerald™ submission workload test, carry out some calibration, debug, tuning and practice runs against selected data points. These practice runs will help you to tune the Emerald™ scripts with the proper number of work streams and threads as well as verify everything is functioning as expected. This approach will enable you to identify the "best-foot-forward" or "sweet-spot" for your SUT. See the *User Guide* for details.

The Must Haves

We assume that the reader has a storage system and wants to learn what the Emerald™ Program performance/energy metrics are for the system. *Measurement Specification* test results using the methodology outlined in this quick reference guide include storage system performance, electrical power consumed, and environmental temperature metrics. This section lists the minimum equipment, software, and hardware required.

Server

The server (system initiating test workload activity) has to be powerful enough to generate enough I/O traffic to the target storage system under test (SUT). Note that multiple servers may be necessary.

Communication Cables

Whether the SUT (storage target) and host server system test initiator (STI) are direct connected or connected via a switched SAN/NAS network, the SUT must be attached with enough communication ports to sustain the workload (IOPS and Bandwidth) that it will be subjected to. In some cases these cables may be fiber lines (Figure 2).

Figure 2- STI and SUT Communication Connectivity

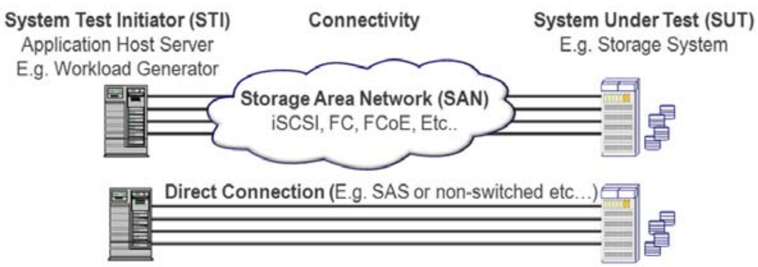

Power cables and Power Meters

There are two primary ways to connect storage-system power:

- Using a breakout box designed for the connector and type of voltage used (single or three phase). The power cable, instead of connecting to a regular power distribution unit, will connect to the breakout box. The SUT will be connected to the breakout box.
- Using a harness specifically designed to connect the SUT to a power meter with direct connection to the power distribution unit or power rail.

Of these two possibilities (Figure 3), *SNIA recommends* the breakout box option. Connectivity using a breakout box or a harness is shown in Figure 4 and Figure 5.

Warning ***Exercise safety procedures when working with connectors, cabling, break-out box, harness and other power items***

Figure 3 – Power Connection Breakout Box and Cable Harness Example

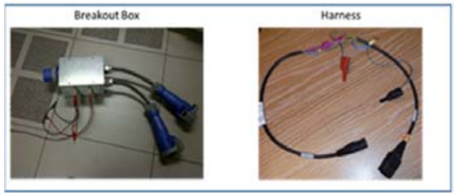

Figure 4 - Power Connection Example Using Breakout Box

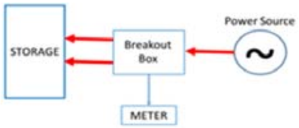

Figure 5 - Power Connection Example Using a Harness

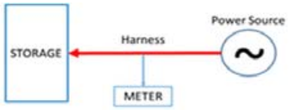

Power Meters

As shown in Figure 4 and Figure 5, the METER represents a high precision power meter or a power analyzer. For ease of reading, we will call both just a "meter."

The harness or breakout box connection between the power source and the storage will always provide a way for the meter to gather the power information required for the measurement.

It is highly recommended that a qualified power person helps with the connections for these reasons:

- Safety
- Correct connectivity between Source-Meter-Storage
- Compliance when doing this work for an official disclose

Inlet Air Temperature Sensors

Sensors can be thermocouples attached to a Data Acquisition System (DAQ) or a direct connect PC-temperature sensor (Figure 6). The sensor is located where it measures SUT's inlet air temperature.

Figure 6 - Example USB Connected Remote Temperature Monitor Device (Digi Watchport®) and Yokogawa MW100 DAQ

Data Collection Computer

Normally this is a Windows™ PC and in some other cases a UNIX (or Unix-Like) computer that is used to collect all the necessary data. This machine normally connects to the power meter and the temperature DAQ/sensors, collects all the data and keeps everything in sync.

SNIA Tip

A better approach that *SNIA recommends* is to use a single server to run workloads and gather data from the power meter and temperature sensor. This way, all data (IOPS or MBPS, Watts and temperature) are time stamped together.

Software

Refer to Appendix C for additional information about software configuration for SNIA Emerald. Download if not already done so all necessary operating-system and related software to connect and access the storage system under test (SUT) including drivers, communication

control, pathing and other items required by the SUT solution. You will also need software to gather and/or control power and temperature data on the collection computer. Note that the environmental (power, temperature) software should come with the specific hardware device and or cabling.

The *Measurement Specification* relies on a workload script powered by Vdbench which can run on various operating systems including Windows as well as various Unix and Linux (*nix) systems.

Vdbench is free (you will be asked to accept license agreement) and can be found at the following URL: http://www.snia.org/emerald/download/Spec_v2.1.

Other software required by the *Measurement Specification*:

- Capacity Optimization Method (COM) generator program: used to test the performance of a COM. COM tests verify the presence and activation of COM functionality (compression, dedupe, thin provisioning, etc.) in a SUT when Emerald™ test data for SUT with COM features is desired. Note that COM testing is not needed for standard power and performance measurement tests. Learn more about COM testing in the *User Guide* and Clause 7.5.5 of the *Measurement Specification*.

Reminders

Don't forget to download the Emerald™ Test Data Report Template (TDR) at http://www.snia.org/emerald/download/Spec_v2.1*. SNIA recommends reviewing the TDR before running Emerald™ Tests to be familiar with how results will be used.*

At this point, it is best to write a checklist of anything missing from the above list and use it to gather all the materiel and equipment you will need to execute the Measurement Specification test. See Appendix A - Sample Checklist.

Preparing for the Tests

Select Product Taxonomy Category the SUT Fits Into

While this quick reference focuses on ONLINE Storage, in general, SNIA has created an Emerald™ taxonomy that divides the spectrum of storage products into a number of product complexity levels and functional aspects. The taxonomy is useful for making decisions when planning your tests--to determine what functional category your storage system falls under.

The *Measurement Specification* taxonomy defines several categories of storage, with subclasses within each:

- Online (Clause 5.3), which is the category this document focuses on.
- Near Online (Clause 5.4)
- Removable Media Library (Clause 5.5)
- Virtual Media Library (Clause 5.6)

Learn more and view the SNIA Emerald™ storage product taxonomy at http://www.snia.org/emerald/taxonomyoverview

Define the Product Family

A product family represents the full range space of configuration variables and options for a particular product, including aspects such as number and type of storage device (Flash, HDD or SSD), availability levels, etc.

Refer to the *User Guide* for more details on how to decide which hardware belongs in a particular family and finding the BFF.

Set up the Test Space and Equipment

Installation of Vdbench is very quick and easy: simply download the ZIP file and expand it into a target directory from where you will use it. Note that there are no installation or setup commands, scripts or GUIs for Vdbench. However, there are some prerequisites; including making sure that the Java Runtime Environment (JRE) is installed along with environment variables. Learn more in Appendix C. If you do not have Java installed you can obtain that from here: http://java.com/en/download/manual.jsp

Running the Tests

A series of tests, shown in Table 1, are run on each SUT. Further details on the tests are contained in the *User Guide* Section 5 and in the *Measurement Specification*.

Table 1 - SNIA Emerald™ Data Collection Summary

Test	Collection Interval (seconds)		Minimum Benchmark Driver Data Collection		Minimum Test Duration (minutes)	
	Power Meter	Temp Meter	Online / Near Online	Removable / Virtual	Online / Near Online	Removable / Virtual
Conditioning	5	60	Response Time (per 1 minute interval)	Throughput (MiB/s)	720	7
Active	5	60	Response Time (per 1 minute interval)	Throughput (MiB/s)	30	30
Idle	5	60	N/A	N/A	120	120

Pre-Fill

The pre-fill test is used to fill the SUT with data. The benchmark driver will be used to fill the required percentage of storage on the SUT with a data set that is two-to-one compressible. The amount of storage required to have data on it is taxonomy dependent; refer to the *Measurement Specification Clause 7.4.1* for requirements.

SUT Conditioning

The SUT Conditioning Test consists of the Hot Banding IO profile, and is used in the *Measurement Specification Clause 7.4.2* to demonstrate the SUT's ability to satisfy an IO request, ensure that the storage devices of the system are fully operational, achieve operational temperature, and place the storage system in a stable and known state. The minimum SUT Conditioning Test phase time period is defined by each storage category, but may be increased in length as necessary.

The goal of the SUT Conditioning Test phase is to get the system up to a stable operational state such that data must be transferred to or from the end-user storage devices (i.e., disk drive, tape drives) and not just cache, though hot banding will mimic hot data in the cache. The workload and time requirements for each phase for each category of storage are specified in the *Measurement Specification*. With online storage, the SUT Conditioning Test must run for at least 12 hours. During that time, collect the number of IOs issued, Average Response Time, size in bytes of each IO issued and average power, all at one-minute intervals.

Active Tests

For online systems, Active Tests consist of a set of phases covering Hot Band Workload, Random Write, Random Read, Sequential Write and Sequential Read. See the *Measurement Specification* Clause 7.4.3 for details.

Ready Idle Tests

Idle Tests for online systems are conducted immediately after the Active Tests and involve no foreground IO requests. See the *Measurement Specification* Clause 7.4.4 for details.

Capacity Optimization Method (COM) Tests

Note Note that if your storage SUT is not being tested with COM capabilities you can skip over this section.

The COM tests verify that a given Capacity Optimization Method is installed and activated. See Emerald™ Measurement Clause 7.4.5 for more information.

The COM Tests consist of a set of heuristics covering:

- Delta snapshots (read and write)
- Thin provisioning
- Data de-duplication
- RAID groups
- Compression

Vendors must follow the given steps for each COM they wish to be given credit for on a given SUT. Each heuristic requires vendors to document how to perform a set of steps on said SUT.

No media may be added or removed, nor changed in state (taken on- or off-line, made a spare, or incorporated, etc.). RAID groups may not be changed. In the event of an automated disk failure and subsequent RAID rebuild at any time during a test, the test must be restarted after the rebuild is completed, and the failed disk replaced per manufacturer guidelines for installed and working systems.

Some COM tests require particular data sets to demonstrate existence. These data sets are generated by the COM Test Data Set Generator C program available at www.sniaemerald.com/download. This program is compiled and loaded on the test host prior to testing. Operational instructions are contained in an associated readme file.

Three different data sets are generated, each approximately 2GB in size:

- Completely irreducible: Cannot be significantly reduced in size by either compression or de-duplication methods;
- Dedupable but not easily compressible: Can be significantly reduced by de-duplication but not easily by compression methods;
- Compressible but not dedupable: Can be significantly reduced by compression but not by de-duplication methods.

The exclusive nature of the data sets supports systems with multiple active COMs, i.e., those that the SUT may be unable to individually disable.

Note

Note that you will need to download, compile and then run the Test Data Set Generator from a STI that has access to the storage SUT following the detailed instructions found inside the C source code file.

See the *User Guide* for details.

Tips for Successfully Conducting Tests

Following are some tips to assist you in conducting successful tests. Given the complexities of storage systems, however, not all potential pitfalls can be included here. But following these steps will avoid most of the commonly encountered problems.

General Items

Keep accurate lab notes documenting initial configuration and parameter settings of hardware and software as well as test workload scripts.

Warning

The Emerald™ script runs a workload that will over-write or destroy any existing data on the storage SUT. Double-check and verify that you are using the proper SUT storage target addresses to avoid accidental loss of data. SNIA assumes no responsibility for data-loss caused by accidental or other use of the *Measurement Specification*. Additional Vdbench and Emerald™ information can be found in Appendix A - Sample Checklist.

Data on Particular Operating Systems

With Vdbench downloaded and extracted into the destination direction on your server that will be running the tests, as well as Java also installed, there are a few more things to be aware of.

Tip

Verify that you have the proper system environmental variables configured including pointing to the Java environment. These will vary depending on your operating system for example Classpath for some *nix, as well as JAVA_HOME and Path for Windows. Make sure that the applicable variables are defined in the proper system settings locations on the STI host where Emerald™ test will be run.

Quick Reference Guide for Online (Active) Storage Systems Testing

AIX Tip

If you are using AIX, do not send Vdbench to the background. This in some cases will cause the output of Vdbench to be lost.

When you run Vdbench, do so from an elevated privilege such as Admin on Windows or Sudo on *nix based systems.

Windows

Another tip is that on Windows systems, should you receive the error message "Unable to obtain CPU statistics", from a command prompt, execute "***Lodctr /R***" with Administrator rights. This command will ensure that Windows server performance data collection capabilities are enabled.

Temperature Sensors

When using a DAQ with thermocouples, make sure they are for normal room temperature and that they are calibrated. If they are pre-calibrated, it is recommended that you calibrate them yourself or double-check their accuracy. If you drop a thermocouple or someone steps on one, change it. They are too inexpensive to hold onto if something compromises their accuracy.

Temperature is not something you must collect in every case. You can, and in some cases should, collect the data, but it is not always taken as part of the *Emerald™ Benchmark*. Power and workload are more important than temperature.

- A high-quality DAQ will cost several thousand dollars, however it is also possible to do without. A simple sensor on the front of each storage box that measures the inlet air is much cheaper.
- Use at least one temperature sensor per storage box and place them in the locations where you have the best average inlet air temperature, normally lower in the rack.

Power Source

Running on a clean power source is recommended.

Calibration, Tuning, Practice, Dry Run, Before First Real Run

The dry run only needs to be run for a couple minutes, as it is just a fast check to establish the process and verify all the tools and systems are functioning correctly.

Fix any bugs or configuration issues found - including Vdbench Emerald™ Test Script parameters before running a full-length test (e.g. practice makes perfect so do a practice test).

Emerald™ Vdbench Results Are HTML Browser Viewable Files

When you run the *Measurement Specification* workload test by specifying a "–o DirectoryName" a directory folder named DirectoryName (or whatever you decide to use) will be created and all results files will be found there.

- For troubleshooting there is a file called paramscan.html that will show you where errors were encountered with bad or missing parameters. For errors, see errorlog.html
- A summary report is called summary.html
- Another useful file is the flatfile.html, which can be used for importing data into Excel or other spreadsheets for analysis, charting or data reduction and reporting.

Data Reduction for Analysis

When measuring both power and performance metrics, always try to keep everything in synch. The *Measurement Specification* calls for one minute sampling of performance and five seconds for power. But if you have a meter that can take a five-second power sample and turn it into a one-minute average, it's recommended that you use that average.

Keep All Data, Even If It Looks Bad

Don't throw away data, even if a test run is configured incorrectly. Keep everything so you can evaluate what happened and determine where errors might lie.

Other Key Items

- Ensure measurement has enough writing on sequential write test for stable sequential read test.
- If system is to be tested for COM, COM feature must be enabled during all measurements.
- All disclosed Reliability, Availability and Serviceability (RAS) features must be turned on during the measurement. Any RAS feature that is a typical daily activity should be captured in the measurement.
- Maintenance tasks such as charging batteries should be completed before testing occurs.
- Timing between the workload generator and power/temperature meter must match. Any offset will cause the metric generation to be off. Time settings should be within one second of each other.
- Avoid having the host providing the workload to the storage system as a bottleneck. If needed, configure additional STI host servers to increase the workload to be handled by the SUT.

Frequently Asked Questions

- Is it really necessary to do this many tests?

  No. There are various scenarios that might apply and some don't require a lot of different tests. For example, when you are only interested in gaining performance metrics, you only need to do the test once to determine the amount of energy consumed. But when you have some kind of system that uses automated tiering and load balancing, it is vital to keep the meter running at all times so you fully understand what is occurring as the system performs its different activities at different times.

  There is also a hybrid scenario where the system does high-performance transactional work during the day and automated tiering and other performance functions at night. You want to measure the SUT under both operating conditions.

- What role does modeling play and how do you validate the model against measured results?

  It is useful to have a modeling tool that is data correlated and can predict load and performance. The model cannot be used in place of running the benchmark, but it can help identify the optimal performance points for a given workload and minimize the amount of testing needed. When running a model, accuracy within +/- 10% is acceptable.

Appendix A - Sample Checklist

Verification Check List provided by IBM Storage Performance team

TEST LAB:

- ☐ Space required reserved
- ☐ Server power is enabled and verified
- ☐ Storage power is enabled and verified
- ☐ Server network connectivity is enabled
- ☐ Storage network connectivity/zoned as needed and enabled

Server Equipment

- ☐ Verify that the server is ready to install additional software
- ☐ Verify all security software is enabled
- ☐ Verify software to generate stimulus is installed and running
- ☐ Verify the Storage is visible and addressable
- ☐ Verify time synchronicity with the data collection computer

Storage Equipment

- ☐ Verify the break out box is correctly connected
- ☐ Do a power down/up test (emergency off test)
- ☐ Verify the connectivity to the power analyzer
- ☐ Verify the temp sensor connectivity and location
- ☐ Verify it has connectivity to the Server
- ☐ Verify the storage can store and hold data (write/read/verify)
- ☐ Verify the synchronicity to the data collection computer

Data collection computer

- ☐ Verify it can connect to the Power meter
- ☐ Verify it can connect to the temperature sensors (DAQ)
- ☐ Verify the clock can synchronize with the Server
- ☐ Verify the clock can synchronize with the Storage system
- ☐ Verify it can synchronize with NIST
- ☐ Verify the collection software is running by doing a zero test run for two minutes

Note: DO NOT START TEST UNTIL ALL BOXES ARE CHECKED.

Appendix B - Sample Test Data Report

The SNIA Emerald™ Test Data Report Template (TDR), shown in Figure 7, serves as both the input form for your test data and as the display form for what you submit (after conversion to PDF format). Download the TDR form here: http://www.snia.org/emerald/download/Spec_v2.1.

Figure 7 - Sample SNIA Emerald™ Test Results Submission Template

Appendix C - Vdbench Emerald™ Script Example

Follow these steps to run the Emerald™ Test Script:

1. Download the current Emerald™ Benchmark script template: http://www.snia.org/emerald/download/Spec_v2.1
2. Make a copy of the downloaded SNIA Emerald™ Benchmark test script file to a safe location.
3. After making a copy of the Emerald™ Benchmark test script, you will need to edit the script to set some initial parameters including devices to exercise. These parameters include the operating system-specific device addresses to be tested, along with number of workload streams and threads. Table 2 shows parameters to be replaced with suggested starting values.

Tip SNIA Recommends starting with these values, running a series of test and tuning runs, and reviewing the results. Tune the parameters so that your storage SUT will perform to your expectations.

Table 2 - SNIA Emerald™ Test Script Parameters

Parameter Name	Parameter Description (View test script for additional details)	Recommended Initial Parameters Values (adjust during tuning practice runs)
SD	Storage Device LUN or Device name, space to use. Review test script for operating system examples	Set in system config file (Windows example) sd=sd1,lun=\\.\PHYSICALDRIVE2,size=550G sd=sd2,lun=\\.\PHYSICALDRIVE3,size=550G
Change_a1, Change_a2	Number of job streams, faster devices may need more streams to stay busy, monitor STI CPU busy.	Suggest starting with 8 per HDD in the storage SUT. For example, when using two devices, start with this at 16. Set in the test script file.
Change_y1, Change_y2, Change_y3	Number of threads, should be multiple of ***Change_a1***. Set in the test script file. Adjust during tuning.	Fine-tune this after practice runs; faster systems may need more threads. For above example, start with 32.
Change_x1 through Change_x4	Number of threads, should be multiple of ***Change_a1***. Set in the test script file. Adjust during tuning.	Fine-tune this after practice runs; faster systems may need more threads. For above example, start with 32.

4. Run the Emerald™ Test Script from a command line with elevated privileges such as Administrator for Windows or sudo for *nix. The command to run the test is:

   Vdbench –f xxxx.txt –o zzzz

   Where xxxx.txt is the name of the SNIA Emerald™ test script text file downloaded from http://www.snia.org/emerald/download/Spec_v2.1. In the above example, ***zzzz*** specifies the output results directory.

   SNIA recommends using a unique directory name for each test run and set of results so that you can keep track of generated data. See a sample of test data in Figure 8.

Tip

Replace the file names used in the command with whatever you set them to and do the same with the output directory name. If you will be doing additional Emerald™ Testing, you will want to take a few moments sometime to learn how you can pass command-line variables to Vdbench scripts to, for example, vary the streams and threads without having to change your scripts. Learn more by referring to the Vdbench user guide.

Figure 8 - Sample SNIA Emerald™ Test Data from Vdbench

```
Vdbench summary report, created 00:06:52 Oct 16 2014 CDT

Link to logfile:                  logfile
Run totals:                       totals
Copy of input parameter files:    parmfile
Copy of parameter scan detail:    parmscan
Link to errorlog:                 errorlog
Link to flatfile:                 flatfile

Link to HOST reports:             localhost
Link to response time histogram:  histogram
Link to SD reports:               sd1      sd2

Link to workload report:          HOTwd_uniform
Link to workload report:          HOTwd_hot1
Link to workload report:          HOTwd_99rseq1
Link to workload report:          HOTwd_99rseq2
Link to workload report:          HOTwd_99rseq3
Link to workload report:          HOTwd_99rseq4
Link to workload report:          HOTwd_99rseq5
Link to workload report:          HOTwd_hot2
Link to workload report:          HOTwd_hot3
Link to workload report:          HOTwd_hot4
Link to workload report:          HOTwd_99wseq1
Link to workload report:          HOTwd_99wseq2
Link to workload report:          HOTwd_99wseq3
Link to workload report:          wd_mixed
Link to workload report:          wd_seq
Link to workload report:          wd_fill

Link to Run Definitions:          rd_prefill For loops: rdpct=0 xfersize=256k threads=32
                                  rd_conditioning For loops: threads=32
                                  rd_hband_final For loops: threads=32
                                  rd_rw_warm For loops: rdpct=0 xfersize=8k threads=32
                                  rd_rw_final For loops: rdpct=0 xfersize=8k threads=32
                                  rd_rr_warm For loops: rdpct=100 xfersize=8k threads=32
                                  rd_rr_final For loops: rdpct=100 xfersize=8k threads=32
                                  rd_sw_warm For loops: rdpct=0 xfersize=256k threads=32
                                  rd_sw_final For loops: rdpct=0 xfersize=256k threads=32
                                  rd_sr_warm For loops: rdpct=100 xfersize=256k threads=32
                                  rd_sr_final For loops: rdpct=100 xfersize=256k threads=32
```

00:06:54.002 Starting RD=rd_prefill; I/O rate: Uncontrolled MAX; elapsed=300000; For loops: rdpct=0 xfersize=256k threads=32

Oct 16, 2014	interval	i/o rate	MB/sec 1024**2	bytes i/o	read pct	resp time	read resp	write resp	resp max	resp stddev	queue depth	cpu% sys+u	cpu% sys
00:07:54.110	1	11.15	2.79	262144	0.00	2708.529	0.000	2708.529	9782.416	3210.397	32.0	1.7	0.4
00:08:54.066	2	8.38	2.10	262144	0.00	3797.753	0.000	3797.753	12800.205	3585.325	32.0	0.6	0.2
00:09:54.055	3	8.35	2.09	262144	0.00	3838.232	0.000	3838.232	13600.426	3677.491	32.0	0.5	0.2
00:10:54.055	4	8.27	2.07	262144	0.00	3868.088	0.000	3868.088	12296.650	3646.402	32.0	0.5	0.2
00:11:54.067	5	8.10	2.03	262144	0.00	3927.674	0.000	3927.674	14732.854	3784.409	32.0	0.5	0.2
00:12:54.057	6	8.18	2.05	262144	0.00	3898.618	0.000	3898.618	14024.898	3677.358	32.0	0.9	0.3
00:13:54.054	7	8.17	2.04	262144	0.00	3929.492	0.000	3929.492	11718.053	3667.935	32.0	2.0	0.5
00:14:54.049	8	8.13	2.03	262144	0.00	3930.091	0.000	3930.091	16138.839	3699.104	32.0	0.7	0.2
00:15:54.051	9	8.22	2.05	262144	0.00	3919.620	0.000	3919.620	12788.867	3660.884	32.0	0.6	0.2
00:16:54.054	10	8.12	2.03	262144	0.00	3934.708	0.000	3934.708	14853.415	3783.452	32.0	0.7	0.2
00:17:54.050	11	8.15	2.04	262144	0.00	3919.889	0.000	3919.889	12675.357	3656.438	32.0	0.7	0.3
00:18:54.055	12	8.37	2.09	262144	0.00	3830.355	0.000	3830.355	15298.758	3746.780	32.0	0.6	0.2
00:19:54.050	13	8.25	2.06	262144	0.00	3917.788	0.000	3917.788	13557.315	3684.927	32.0	0.7	0.2
00:20:54.051	14	8.07	2.02	262144	0.00	3890.716	0.000	3890.716	13404.992	3675.981	32.0	0.6	0.2
00:21:54.052	15	8.23	2.06	262144	0.00	3901.047	0.000	3901.047	13234.214	3673.983	32.0	0.7	0.2
00:22:54.050	16	8.18	2.05	262144	0.00	3934.362	0.000	3934.362	14337.605	3764.168	32.0	0.6	0.2
00:23:54.053	17	10.62	2.65	262144	0.00	3006.532	0.000	3006.532	17454.475	3588.848	32.0	0.7	0.2
00:24:54.051	18	8.10	2.03	262144	0.00	3963.459	0.000	3963.459	16365.264	3847.416	32.0	0.7	0.2
00:25:54.052	19	8.08	2.02	262144	0.00	3933.057	0.000	3933.057	15114.412	3719.113	32.0	0.6	0.2
00:26:54.049	20	8.13	2.03	262144	0.00	3943.397	0.000	3943.397	13195.721	3677.685	32.0	0.7	0.2
00:27:54.056	21	8.22	2.05	262144	0.00	3899.735	0.000	3899.735	14189.568	3716.104	32.0	0.6	0.2
00:28:54.049	22	8.12	2.03	262144	0.00	3946.240	0.000	3946.240	12671.680	3699.051	32.0	0.6	0.2
00:29:54.049	23	8.10	2.03	262144	0.00	3965.446	0.000	3965.446	15197.234	3711.539	32.0	0.6	0.2
00:30:54.051	24	8.18	2.05	262144	0.00	3915.101	0.000	3915.101	16108.885	3776.765	32.0	0.6	0.2
00:31:54.049	25	8.10	2.03	262144	0.00	3930.779	0.000	3930.779	11955.923	3664.183	32.0	0.6	0.2
00:32:54.048	26	8.12	2.03	262144	0.00	3924.162	0.000	3924.162	13098.138	3732.208	32.0	0.6	0.2
00:33:54.049	27	8.18	2.05	262144	0.00	3931.947	0.000	3931.947	16822.507	3789.889	32.0	0.6	0.2
00:34:54.052	28	8.13	2.03	262144	0.00	3934.217	0.000	3934.217	15310.207	3766.465	32.0	0.6	0.2
00:35:54.055	29	8.10	2.03	262144	0.00	3962.808	0.000	3962.808	16453.115	3845.570	32.0	0.6	0.2
00:36:54.051	30	7.97	1.99	262144	0.00	4007.647	0.000	4007.647	15069.656	3855.627	32.0	0.6	0.2

Oct 16, 2014	interval	i/o rate	MB/sec 1024**2	bytes i/o	read pct	resp time	read resp	write resp	resp max	resp stddev	queue depth	cpu% sys+u	cpu% sys
00:37:54.052	31	8.00	2.00	262144	0.00	3960.483	0.000	3960.483	15993.147	3745.911	32.0	0.6	0.2
00:38:54.052	32	8.13	2.03	262144	0.00	3986.682	0.000	3986.682	16521.521	3814.982	32.0	2.9	0.6
00:39:54.048	33	8.02	2.00	262144	0.00	3996.876	0.000	3996.876	15528.093	3778.500	32.0	0.5	0.2
00:40:54.059	34	7.93	1.98	262144	0.00	3955.464	0.000	3955.464	15086.258	3768.657	32.0	0.5	0.2
00:41:54.049	35	8.05	2.01	262144	0.00	4018.625	0.000	4018.625	12933.646	3782.129	32.0	0.5	0.2
00:42:54.050	36	8.08	2.02	262144	0.00	3968.847	0.000	3968.847	17321.999	3835.056	32.0	0.5	0.2
00:43:54.047	37	8.05	2.01	262144	0.00	3971.591	0.000	3971.591	15948.118	3799.978	32.0	1.9	0.5
00:44:54.047	38	8.07	2.02	262144	0.00	3967.778	0.000	3967.778	16716.409	3788.615	32.0	0.7	0.2
00:45:54.048	39	7.90	1.98	262144	0.00	4007.975	0.000	4007.975	14054.875	3739.007	32.0	0.6	0.2
00:46:54.046	40	7.93	1.98	262144	0.00	4069.136	0.000	4069.136	14504.984	3826.727	32.0	0.6	0.2

Space For Lab Notes

Storage Networking Industry Association
4360 Arrows West Drive • Colorado Springs, CO 80907 • Phone: 719-694-1380 • Fax: 719-694-1389 • www.snia.org