SPEC - Enabling Efficiency Measurement

Karl Huppler
IBM Corporation
huppler@us.ibm.com

Klaus-Dieter Lange
Hewlett-Packard Company
klaus.lange@hp.com

John Beckett
Dell, Inc.
john_beckett@dell.com

ABSTRACT
An overview of the SPEC PTDaemon and SPEC's Benchmark Methodology.

Categories and Subject Descriptors
H.3.4 [**Systems and Software**]: Performance evaluation (efficiency and effectiveness)

General Terms
Design, Experimentation, Measurement, Performance, Reliability, Standardization

Keywords
SPEC, SERT, Rating Tool, Benchmark, Energy Efficiency, Power, Server, Storage, Datacenter, EPA

1. INTRODUCTION
Over the years, computing solutions have become less expensive to purchase and maintain delivering more and more computing capacity at lower and lower equipment and operational costs. At the same time, the cost of energy has continued to rise. In some areas the amount of power that is available can no longer grow with the demand. In order to create benchmarks including an efficiency aspect SPEC [1] created fundamental guidelines [3] and an interface [4] to measure the power consumption during benchmark measurements.

2. BENCHMARK METHODOLOGY
SPEC's Benchmark Methodology is intended for performance benchmark designers and implementers who want to integrate a power component into their benchmark. This document also serves as an introduction to those who need to understand the relationship between power and performance metrics in computer systems benchmarks.

The assumption is that the business model and benchmark application are already selected and may already be implemented. Guidance is provided for including power metrics in existing benchmarks, as well as altering existing benchmarks and designing new ones to provide a more complete view of energy consumption.

SPEC Benchmark Methodology covers the flowering topics:

- Defining Power components within Performance Benchmarks
- System Under Test (SUT)
- Power Measurement
- Environmental considerations
- Performance/Power metrics
- Reporting
- Automation and Validation

It also includes a comprehensive appendix about energy terminology.

3. SPEC PTDaemon
SPEC's Power/Temperature Daemon (also known as PTDaemon, PTD or ptd) is used by benchmarks to offload the work of controlling a power analyzer or temperature sensor during measurement intervals to a system other than the SUT. It hides the details of different power analyzer interface protocols and behaviors from the benchmark software, presenting a common TCP/IP-based interface that can be readily integrated into different benchmark harnesses.

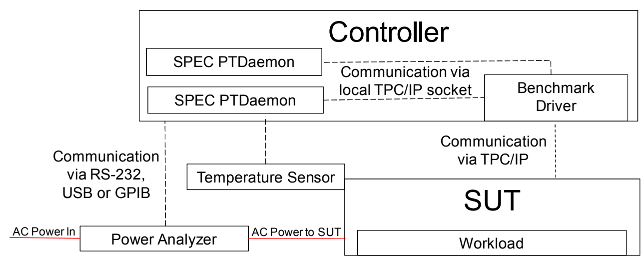

Figure 1 – SPEC PTDaemon Setup

The benchmark harness connects to PTDaemon (Figure 1) by opening a TCP port and using a proprietary protocol to control devices and retrieve measurement data. For larger configurations, multiple IP/port combinations can be used to control multiple devices.

PTDaemon can connect to multiple analyzer and sensor types via protocols and interfaces specific to each device type. The device type is specified by a parameter passed locally on the command line on initial invocation of the daemon.

The communication protocol between the SUT and PTDaemon does not change regardless of device type. This allows the benchmark to be developed independently of the measurement types to be supported.

3.1 Code Structure
PTDaemon is implemented using a main process that controls initialization and the network command interface, with a separate thread that manages the power analyzer or temperature sensor. Some analyzers that do not operate with a standard command/response structure also require an additional thread to receive asynchronous data from the device. Upon startup the main process performs the following steps:

ICPE'12, April 22–25, 2012, Boston, Massachusetts, USA.
ACM 978-1-4503-1202-8/12/04.

- Parses command-line arguments and checks for validity
- Initializes the analyzer interface
- Connects to the analyzer and checks for valid responses
- Initializes the network connection and opens a socket in listen mode
- Optionally sets ampere and voltage ranges

At this point, the main process goes into a command-handling loop, where it receives commands, parses them, performs any necessary actions, and sends a response back across the network.

The device thread is started any time a measurement interval is begun. It consists of a timed loop that calls the analyzer read function, logs the values, and then sleeps any remaining time before the next scheduled sample. The device thread ends when either the benchmark harness requests an end to measurements or a network error occurred.

The device thread will not disconnect when the network connection is closed cleanly by the remote end. This allows measurements to continue without an active network connection, allowing the possibility of measuring power during arbitrary periods including reboots or low power modes without the necessity of any third-party software control.

3.2 Power Analyzer Support for PTDaemon

To support different power analyzers, each device needs its own module in PTDaemon. PTDaemon is periodically updated to support new power analyzers, temperature sensors, and additional device features. SPEC members and licensees can use the Power Analyzer Acceptance Process [5] to add software support for new devices and submit tests to SPEC for review and possible inclusion in later PTDaemon releases.

When support for a new device is requested by a SPEC licensee or device vendor, a volunteer sponsor from within SPEC will work on the software implementation. Complimentary binary licenses for PTDaemon can also be provided to third parties to assist with testing of the new device.

Once the initial software support has been added, the new analyzer is compared against a previously accepted power analyzer by running both with SPECpower_ssj2008 simultaneously while connected in series. Then, the order of the analyzers is reversed to negate any measurement differences due to voltage drops or power consumption of the analyzers themselves. The results of this test are scrutinized to determine whether the averaged power values from the two devices differ more than the sum of tolerances of each power analyzer device.

Another part of the SPEC Power Analyzer Acceptance Test uses an internal tool to generate power pulses of specific durations, which are used to examine the sampling windows of the power analyzer under test. The measured power values are compared to the corresponding to the pulse widths generated by the tool, as well as other criteria such as duplicate and missing sample limits are used to judge the new analyzer. Additionally, there are separate criteria for Multi-channel Power Analyzers that differ from Single-channel Power Analyzers.

Once the results of this testing have been submitted to SPEC, all relevant compliance requirements from the SPECpower_ssj2008 power analyzer device Run Rules are reviewed to determine if all requirements are met. If all criteria are satisfied, SPEC will accept the device for compliant benchmark testing and add support to a later release of PTDaemon.

4. CONCLUSION

The ground breaking work from the SPECpower Committee [2] created a path for easy implementation of an energy metric as utilized by TPC-Energy, the Server Efficiency Rating Tool (SERT - the upcoming EPA Energy Star for Servers v2.0 from SPEC), and many of the SPEC benchmarks such as SPECpower_ssj2008, SPECvirt_sc2010, and SPECweb2009.

In addition, SPEC is licensing PTDaemon for both non-profit and commercial entities, so the work that SPEC has done to refine power collection methodology and hide the details of different power analyzers behind the PTDaemon interface will help ease the process of other organizations to develop meaningful power efficiency workloads.

5. ACKNOWLEDGMENTS

The authors want to acknowledge the additional members of the SPECpower Committee who have contributed to the design, development, testing, and overall success of the Power Methodology and PTDaemon: Paul Muehr, Van Smith, Greg Darnell, Hansfried Block, Jeremy Arnold, Sanjay Sharma, David Ott, and Mike Tricker.

6. REFERENCES

[1] Standard Performance Evaluation Corporation: http://www.spec.org

[2] SPECpower: http://www.spec.org/power/

[3] SPEC Benchmark Methodology: http://www.spec.org/power/docs/SPEC-Power_and_Performance_Methodology.pdf

[4] SPEC PTDaemon: http://www.spec.org/power/docs/SPEC-PTDaemon_Design.pdf

[5] SPEC Power Analyzer Acceptance Process: http://www.spec.org/power/docs/SPEC-Power_Analyzer_Acceptance_Process.pdf