


# Design and Development of the Server Efficiency Rating Tool

(SERT)<sup>TM</sup>

Friday October 8<sup>th</sup>, 2010



spec

# SPEC Background



The Standard Performance Evaluation Corporation (SPEC®) was founded in 1988

It is a nonprofit organization dedicated to creating industry standards for measuring computer and software performance

Includes representatives from more than 60 member companies and organizations

Has released more than 30 industry-standard benchmarks

These have been used to create more than 20,000 peer-reviewed published performance reports



# SPEC Organization

Structured in four groups:

- Open Systems Group (OSG)
- High Performance Group (HPG)
- Graphics and Workstation Performance Group (GWPG)
- Research Group (RG)

The OSG includes groups covering benchmarking for:

- Processors
- Java
- Virtualization
- Power

The SPECpower subcommittee was responsible for creating the SPECpower\_ssj2008 benchmark

# ENERGY STAR Motivation



Almost 3% of all US generated electricity consumed in 2010 will be used to power datacenters

Most will be used by servers and associated air conditioning

- Additional by networking and other support infrastructure

The US Environmental Protection Agency (EPA) launched the ENERGY STAR® Computer Server program in May 2009

- After starting initial industry consultation in January 2006



# The Server Efficiency Rating Tool

A tool to deliver a broad set of data derived from the individual tests, and NOT a benchmark...which is a loaded term

Does not attempt to simulate *specific* end user workloads, instead providing a set of focused synthetic “worklets”

Worklets have so far been developed to exercise:

- Processors
- Memory
- Storage I/O
- Overall System

Worklets may also be combined into various configurations:

- Run serially or in parallel to provide “system” tests integrated across all the different subsystems



# The SERT Development Model

Delivered via an industry-wide collaboration:

- Anyone willing to contribute effort is welcome to join the SPECpower subcommittee and work on SERT development and testing
- Good for enabling cross-industry representatives with hardware and software backgrounds to participate
- Can be difficult to get sufficient development resources to support all the desired features – most contributors also have “real” jobs ☺
- We welcome volunteers from academia as well as from industry!

SERT development is described in the SERT Design Document:

- A “living” document that is regularly updated and shared with the EPA
- Also with the ENERGY STAR Computer Server industry stakeholders
- Feedback and input to the SPEC is provided via regular calls and reviews organized and mediated by the EPA

# Development Audience



Many of the EPA's ENERGY STAR stakeholders are already represented within SPEC

- Most of the large OEMs and operating system suppliers
- Seeking additional involvement of smaller providers of server systems
- Such as the Value Added Resellers (VARs) and “white box” builders
- Such companies typically have few resources at their disposal
- So are very constrained in the configurations they can test
- And the dedicated test hardware and actual test engineers they can afford to provide

Discussions currently underway with the EPA regarding certifying test laboratories to perform testing and validation for such companies



# Development Constraints

Targeting a clearly defined set of hardware and operating system platforms to support:

Hardware Platform	x86 (AMD)	x86 (AMD)	x86 (AMD)	x86 (Intel)	x86 (Intel)	x86 (Intel)	Itanium (Intel)	POWER (IBM)	POWER (IBM)	POWER (IBM)	SPARC (Fujitsu)	SPARC (Oracle)
Operating System	Windows Server 2008 R2	LINUX	Solaris	Windows Server 2008 R2	LINUX	Solaris	HP/UX 11i	AIX	IBM I	LINUX	Solaris	Solaris

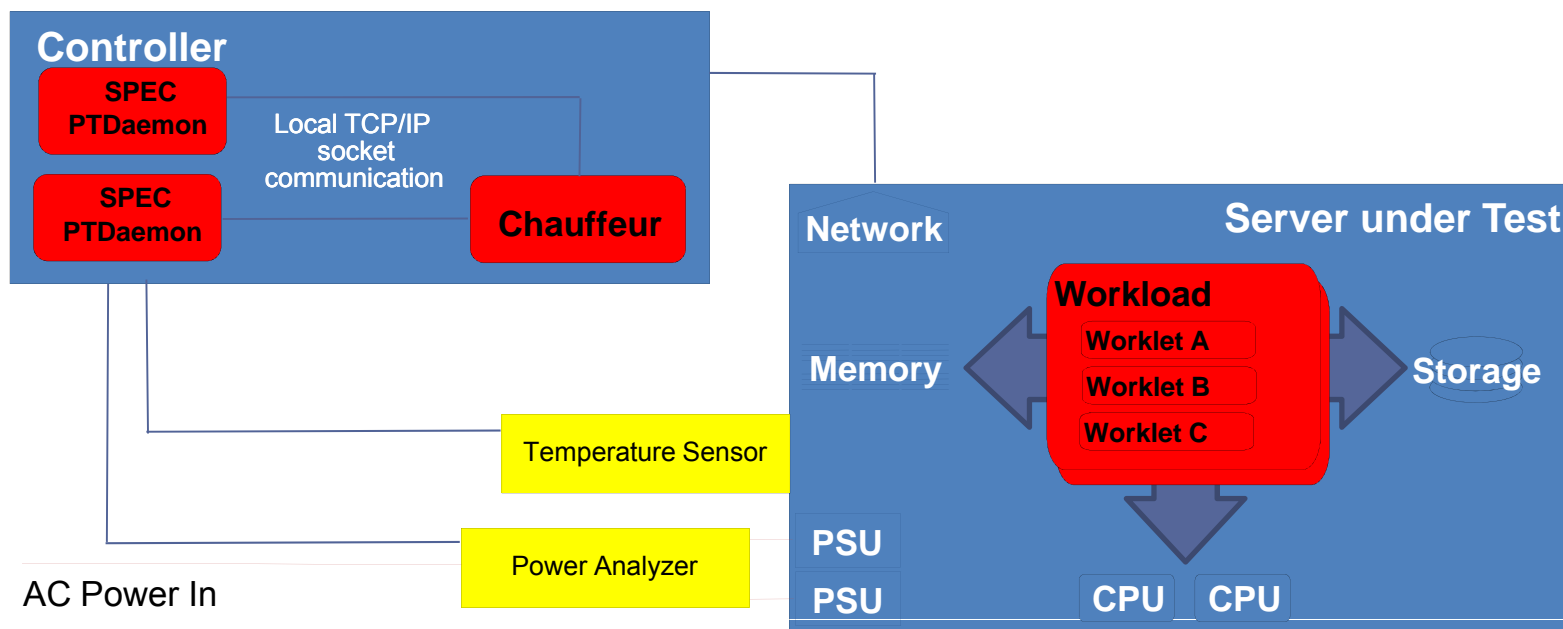
- Only 64-bit operating systems would be supported,
- Determine how many older releases of operating systems available prior to the release of SERT need to be supported
- Constrained due only to limited development resources, not architectural

Source code will not be made generally available

- Reducing the risk of accidental customization
- And thus potentially different results being reported back to the EPA
- Minimizing the complexity of the results review process
- However, anyone may request a code review of any part of the SERT



# Hardware and Software Configuration



# SERT Components



Worklets™ provide the simulated workloads

Chauffeur™ manages worklet execution and coordinates between the System Under Test and the Controller

PTDaemon™ coordinates between Chauffeur and one or more power analyzers and temperature probes

Reporter takes the raw test results (in XML) from each worklet and outputs the combined results in XML, HTML and eventually PDF-formatted reports

The GUI provides an interactive UI to enter static data about the systems, and also invokes agents to collect discoverable information such as BIOS, CPU, memory and IO device details

## CPU

- A mixture of integer, floating point, matrix manipulation and XML validation

## Memory

- Memory-intensive derivatives of the XML validation tests

## Storage I/O

- Sequential, random and various mixtures of both


## Network I/O

- Possibly covered by “configuration modifiers”...

## System

- A derivative of the SPECpower\_ssj2008 performance benchmark

## Active Idle



# SERT Configuration & Execution

Configuration may come from scripts, the GUI and the system agent, then be passed to Chauffeur:

- Run each worklet based on its configuration and initial parameters
- The sequence of serial or parallel operations specified by the configuration
- Supports the affinitization of JVMs to specific sockets or cores
- Provides the option of running all the worklets in a single JVM or of starting a new JVM for each worklet

To ensure valid results there are multiple phases of execution:

- A warm-up phase at the start of each worklet execution to enable the system to settle and for power consumption to stabilize
- One or more calibration phases
- A measurement phase during which the results are collected
- A cool-down phase, after which the worklet terminates



# Implementation Language Choices

It is critical that the SERT support as many operating systems and processor architectures as possible

Many SPEC benchmarks have already been developed in JAVA

But JAVA is not appropriate for all tasks, so C and C++ also need to be supported

JAVA worklets may call libraries developed in C or C++ via the JAVA Native Application (JNA) interface

- Typically used to access lower level operating system APIs that do not map to JAVA classes, and may offer richer capabilities

May also call the standard JAVA libraries

- E.g. for encryption and compression support, that may vary between countries and vendors



# Supported Hardware

Version 1 of the ENERGY STAR program supports servers with up to four processors

Extended for Version 2:

- The SERT is designed to support arbitrarily large servers
- First release is targeted at up to four processors, with any number of cores
- Also supports Blade Servers and Multi-Node servers
  - Comprising two or more independently booted nodes in a single enclosure

The SERT does not explicitly support virtualization

- It expects to be executed on physical (rather than virtualized) hardware
- The SERT can however simulate the high loads that virtualization imposes on hardware
- This is likely to be revisited in subsequent releases, given sufficient development resources

# Conclusions



Version 1 of the SERT is intended to release as part of the EPA's ENERGY STAR Computer Server Version 2 program

Developed via significant cross-industry collaboration, which enables it to avoid the risks of favoritism

Targeting all sectors of the server market

- From major OEMs to small Value Added Resellers and “white box” builders

The SERT offers the range of power usage data that server buyers will need to enable environmentally conscious purchasing decisions in future

It will continue to evolve and provide even richer testing

We always have opportunities for enthusiastic contributors!

# References

Server Efficiency Rating Tool home page:

- <http://www.spec.org/sert/>

Server Efficiency Rating Tool public Design Document (latest version):

- [http://www.spec.org/sert/docs/SERT-Design\\_Doc.pdf](http://www.spec.org/sert/docs/SERT-Design_Doc.pdf)

ENERGY STAR Enterprise Servers home page:

- [http://www.energystar.gov/index.cfm?c=archives.enterprise\\_servers](http://www.energystar.gov/index.cfm?c=archives.enterprise_servers)

ENERGY STAR Computer Specification Version 1.0:

- [http://www.energystar.gov/ia/partners/product\\_specs/program\\_reqs/computer\\_server\\_prog\\_req.pdf](http://www.energystar.gov/ia/partners/product_specs/program_reqs/computer_server_prog_req.pdf)





# References

## ENERGY STAR Computer Specification Version 1.0 Power and Performance Data Sheet:


- [http://www.energystar.gov/ia/partners/prod\\_development/new\\_specs/downloads/servers/Final\\_Datasheet.xls](http://www.energystar.gov/ia/partners/prod_development/new_specs/downloads/servers/Final_Datasheet.xls)

## ENERGY STAR Computer Servers Draft 1 Version 2.0:

- [http://www.energystar.gov/ia/partners/prod\\_development/revisions/downloads/computer\\_servers/Draft1Version2ComputerServers.pdf](http://www.energystar.gov/ia/partners/prod_development/revisions/downloads/computer_servers/Draft1Version2ComputerServers.pdf)

## ENERGY STAR Computer Servers Draft 1 Version 2.0 Power and Performance Datasheet:

- [http://www.energystar.gov/ia/partners/prod\\_development/revisions/downloads/computer\\_servers/Draft1Version2PowerPerformanceDatasheet.pdf](http://www.energystar.gov/ia/partners/prod_development/revisions/downloads/computer_servers/Draft1Version2PowerPerformanceDatasheet.pdf)



# Acknowledgements

The authors would like to acknowledge the current members of the SPECpower subcommittee who have contributed to the design, development, testing and overall success of the SERT: Paul Muehr, Greg Darnell, Hansfried Block, Christian Koopmann, Klaus-Dieter Lange, Karl Huppler, Jeremy Arnold, Sanjay Sharma, and David Ott.

SPEC and the tool and benchmark names SERT and SPECpower\_ssj2008 are registered trademarks of the Standard Performance Evaluation Corporation.