Server Efficiency Rating Tool™

Overview
January 2013

Klaus-Dieter Lange
Chair, SPECpower Committee, SPEC

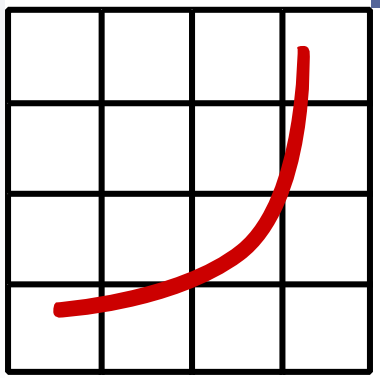

spec

SPEC's Server Efficiency Rating Tool™

Agenda

SERT Demo/Training Session

SPEC F2F - Austin, TX, USA – 23rd January 2012

- SPEC Overview
- SPEC's Philosophy
- Milestones
- PTDaemon
- Goals and Requirements
- Taxonomy of Platform Evaluation Tools
- Hardware and Software Configuration
- Workload
- Worklets
- Execution Phases
- SERT GUI
- SERT Result Report
- Efficiency Program Use / Regulations
- SERT Metric/Score
- Universal Usage
- Q&A

SPEC Overview

- A world-wide non-profit consortium formed in 1988 to establish, maintain and endorse a standardized set of relevant benchmarks that can be applied to the newest generation of high-performance computers
- Comprised out of over 80 computer hardware and software vendors, educational institutions and government agencies
- Developed over 30 industry-standard benchmarks for system performance evaluation in a variety of application areas
- Largest public repository of well documented, peer reviewed, benchmark results (~30,000)
- In-depth understanding of workloads, benchmark code, fair comparisons across different platforms

SPEC's Philosophy

- To ensure that the marketplace has a fair and useful set of metrics to differentiate systems
  - A good benchmark, that is reasonable to utilize, will lead to a greater availability of results in the marketplace
- To provide a standardized suite of code that has already been ported to a wide variety of platforms
  - The licensee can immediately start with the measurement on all supported platforms without code-porting
- SPEC welcomes organizations to join and participate in our work, and stands ready to offer guidance on workloads and benchmarks
  - Membership is open to any interested company or entity

Milestones

SPEC Power and Performance Methodology

- An introduction on power and performance metrics for computer systems
- Guidance for Power and Performance benchmark development (existing and new designs)
- Methodology – captures the experience/conclusions of the SPECpower committee since January 2006
- http://www.spec.org/power_ssj2008/docs/SPECpower-Methodology.pdf

SPEC PTDaemon (Power and Temperature Daemon)

- Infrastructure software to connect, control and collect data from power and temperature measurement devices
- http://www.spec.org/power_ssj2008/docs/device-list.html

SPEC's Server Efficiency Rating Tool™

PTDaemon 1/2

Challenges integrating power analyzers into test tools

- Vendor specific, dissimilar programming interfaces and command sets
- Diverse result formats
- Various connection types (Serial, GPIB, USB) using device specific communication settings/parameters
- Uncertainty calculation missing in vendor control software

PTDaemon accomplishments

- Supports multiple protocols and interfaces specific to each device type (Serial, GPIB, USB etc.)
- Supports major global power standards
- Defines unified network protocol for device configuration (e.g. range setting) and collecting power data readings
- Remote communication via TCP/IP protocol over LAN
- Well defined measurement procedure ensuring comparable power readings
- Standardized result format
- Acceptance process including examination of device capabilities, adding more devices, today 14 accepted power analyzers, including multi-channel and 3-phase instruments
- Uncertainty calculation
- Logging of power readings and problems (errors, warnings)

SPEC's Server Efficiency Rating Tool™
PTDaemon 2/2

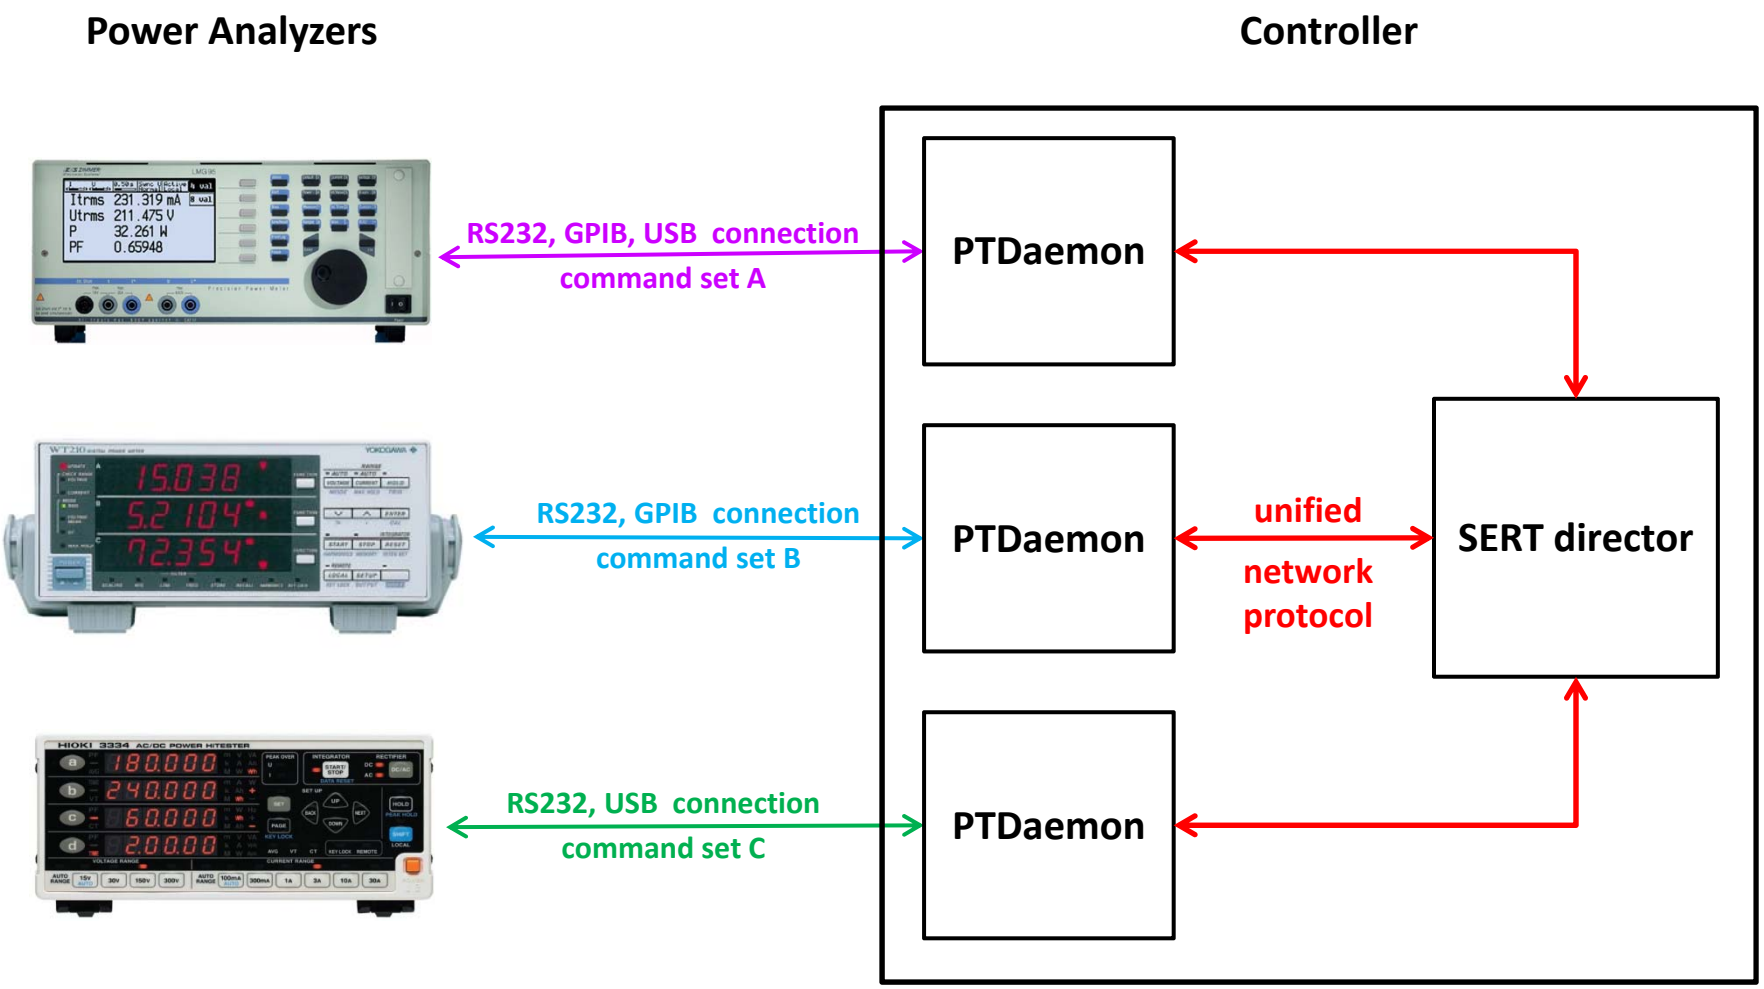

Accepted measurement devices: http://www.spec.org/power/docs/SPECpower-Device_List.html

SPEC's Server Efficiency Rating Tool™
Goals and Requirements 1/2

Server Efficiency Rating Tool (SERT)

- Evaluate the energy efficiency of computer servers
- Built for and in conjunction with the US EPA Energy Star Program
  - Collaborate on workload, metric and logistics
- First order approximation of server efficiency
  - Executes a variety of common types of work (worklets)
  - Not representative of any particular application area
  - Not a capacity planning tool
- Economical, easy to use, minimal equipment and skills requirements
  - Through automated processes
- Hardware Architecture and OS Agnostic
  - Supports various hardware platforms and OS's
- Leverage and extend SPEC methods, code-base and expertise
- AMD, Dell, Fujitsu, HP, Intel, IBM, Microsoft have committed resources for the development of SERT in the SPECpower Committee

SPEC's Server Efficiency Rating Tool™
Goals and Requirements 2/2

Processor sockets/compute nodes

- SERT 1.0.0 will be designed for and tested with
  - Max. 8 sockets per node
  - Max. 64 server nodes
  - Limited to homogeneous multi-node and blade servers
- DC power measurement
  - No resource committed for development and test

Platform/OS

- SERT 1.0.0 will be implemented for and is planned to be tested on the following platform OS (64-bit only) combinations, pending resources:

Platform	X86 (AMD)	X86 (AMD)	X86 (Intel)	X86 (Intel)	Power
OS	Windows 2008 R2	LINUX	Windows 2008 R2	LINUX	AIX

SPEC's Server Efficiency Rating Tool™

Taxonomy of Platform Evaluation Tools

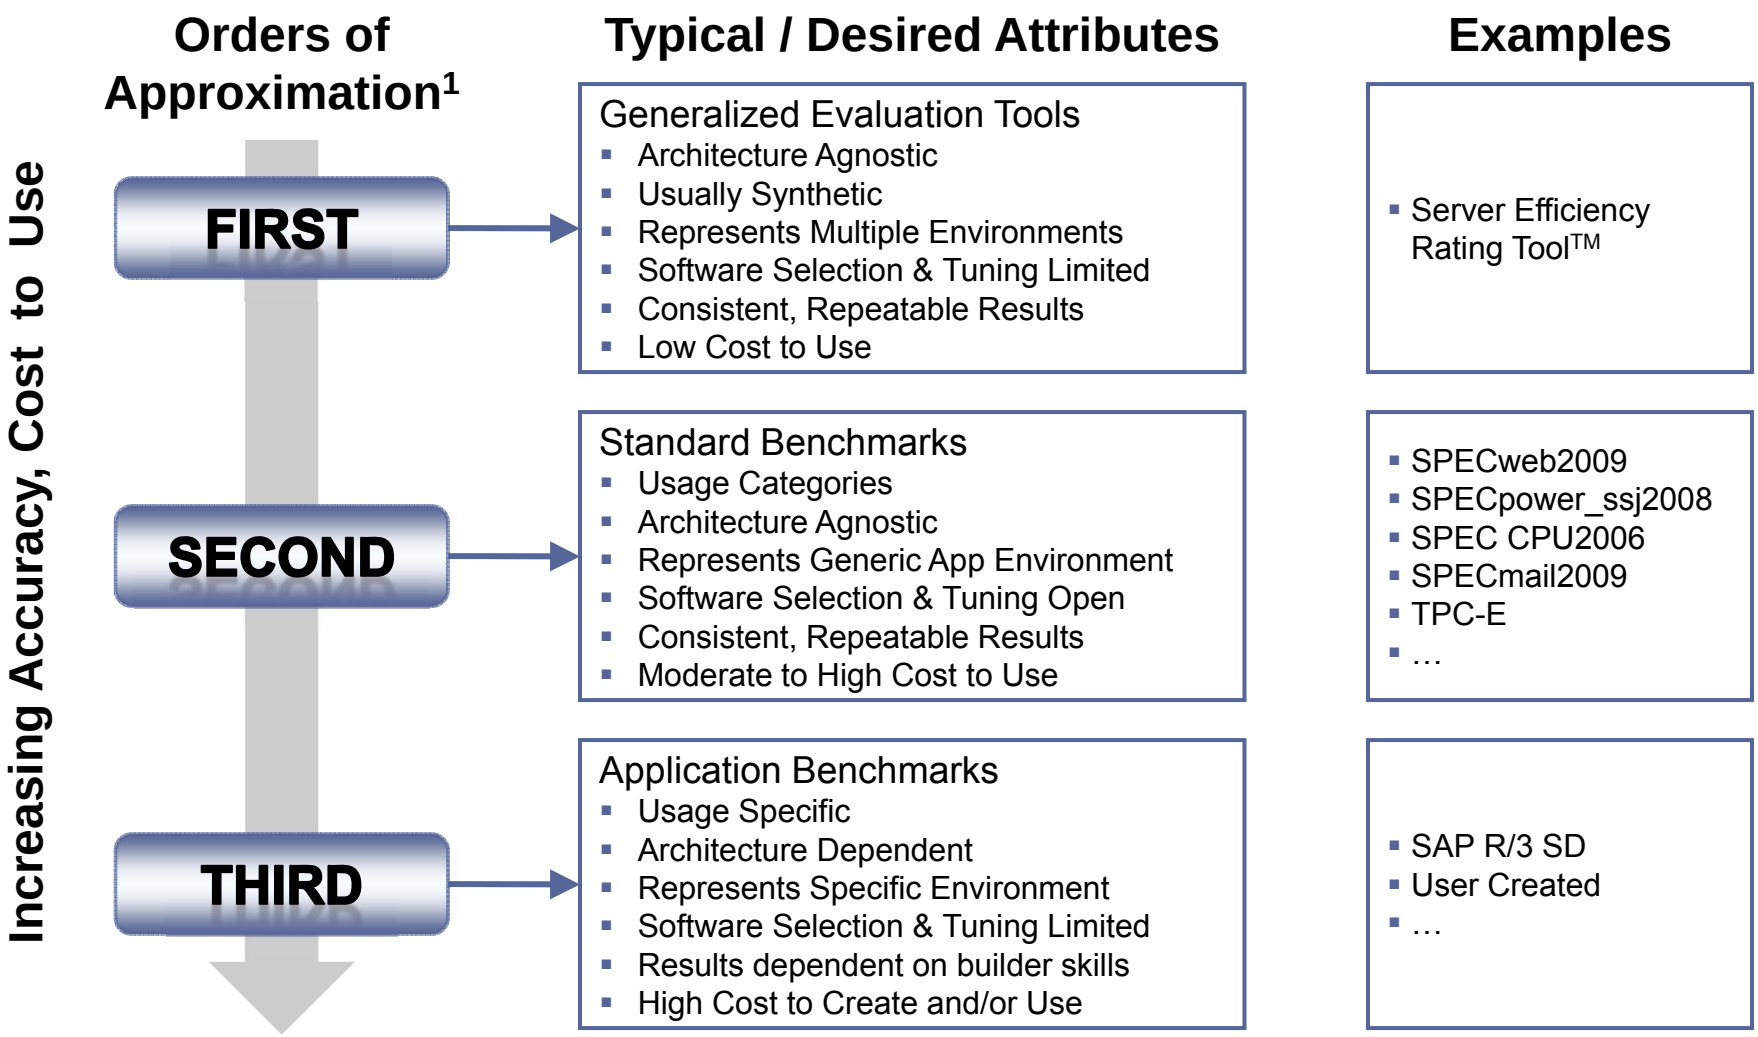

1. Taxonomy from: "*The State of Energy and Performance Benchmarking for Enterprise Servers*" ; A. Fanara, E. Haines, A Howard; August 2009

Hardware and Software Configuration

Controller and SUT

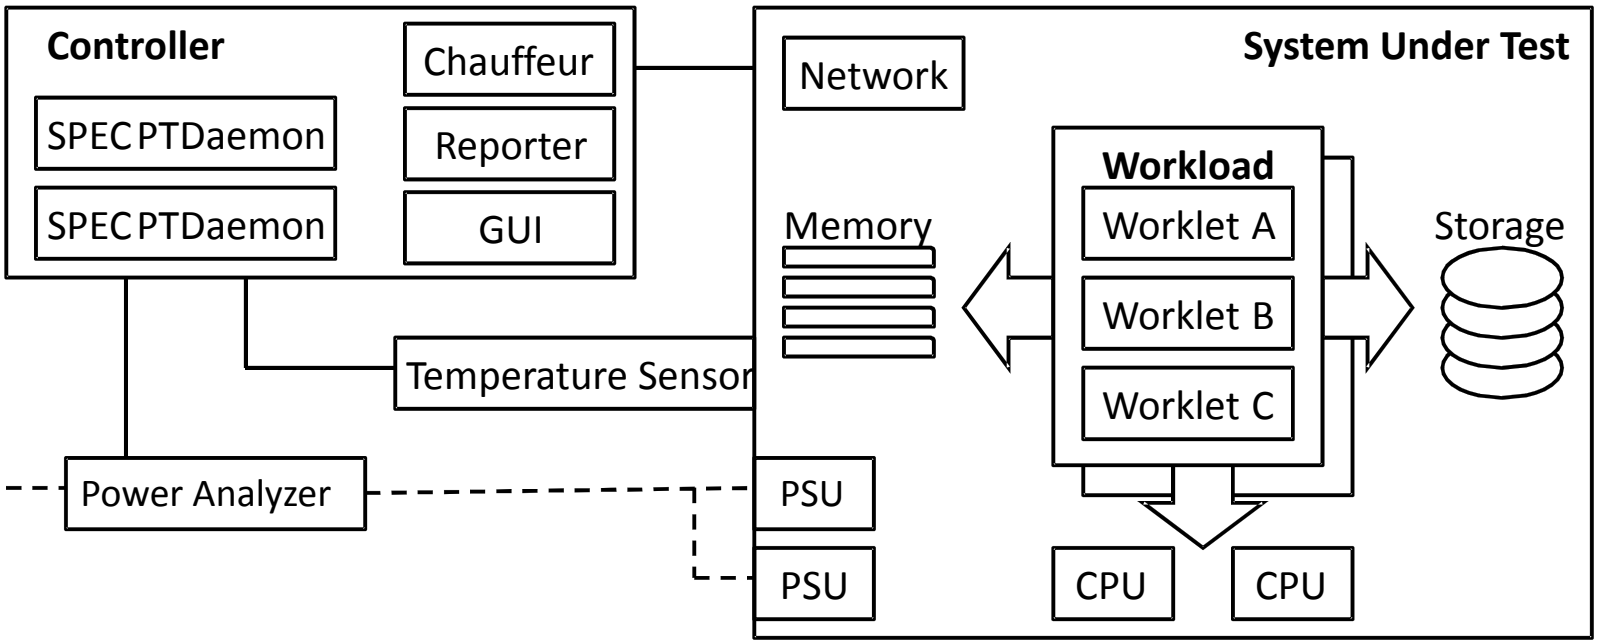

Workload

SERT's Workloads (Collections of Synthetic Worklets)

- Worklet design guidelines:
  - Worklets to assess CPU, Memory, Storage IO subsystem
  - Network IO will be handled by *configuration power/performance modifiers*
  - Worklets do not represent a particular application
  - Adjustable to different performance levels
  - Self-calibrate to maximum performance level
  - Multiple programming languages may be used
  - Scale with the available hardware resources
  - Higher worklet score indicates higher energy efficiency
    - Different worklet scores will have different metrics and can not be compared against each other
    - The definition of the worklet scores is currently in development

SPEC's Server Efficiency Rating Tool™
Worklets

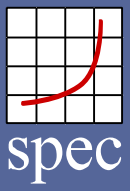

Workload	Load Level	Worklet Name
CPU	100%, 75%, 50%, 25%	Compress
		CryptoAES
		LU
		SHA256
		SOR
		SORT
		XMLValidate
Memory	Flood: Full, Half	Flood
	Capacity: 4GB, 8GB, 16GB, 128GB, 256GB, 512GB, 1024GB	Capacity
Storage	100%, 50%	Random
		Sequential
Hybrid	100%, 87.5%, 75%, 62.5%, 50%, 37.5%,25%, 12.5%	SSJ
Idle	idle	Idle

SPEC's Server Efficiency Rating Tool™

Execution Phases

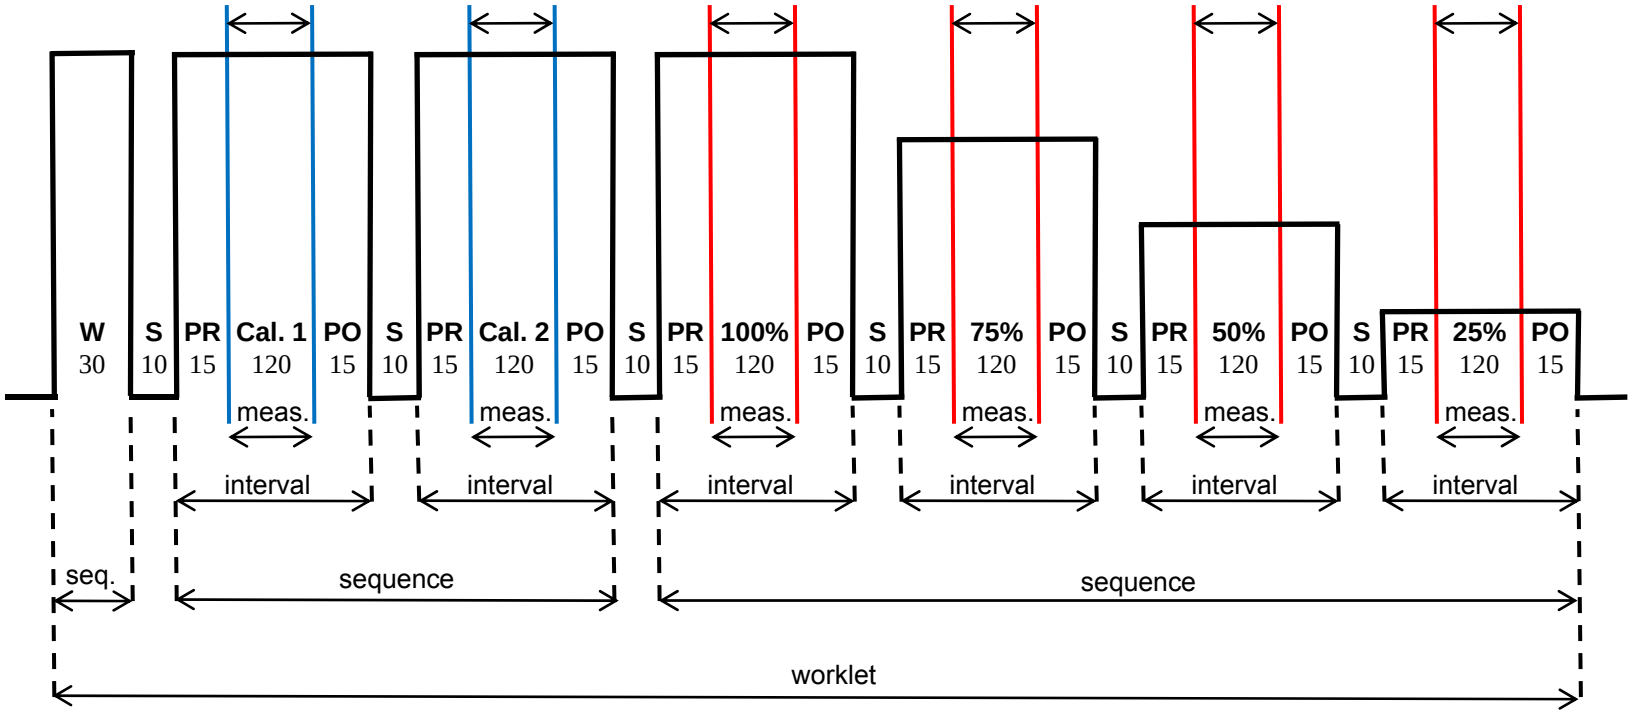

W = Warmup (30 sec)
S = Sleep (10 sec)
PR = Pre measurement (15 sec)
PO = Post measurement (15 sec)
Cal. N = Calibration Interval N (120 sec)
nnn% = Measurement Phase (120 Sec)

Total minimum execution time per worklet = 990 sec
plus worklet initialization times

SPEC's Server Efficiency Rating Tool™
SERT GUI

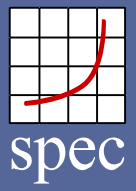

- Graphical interface
  - Gathering SUT hardware and software configuration data
  - Configuring and running the SERT
  - Archiving the measured results and log files.
- Provides ability to save and re-import complete configurations to simplify repeated testing.
- Default Mode (EPA compliant test record)
  - Executes the entire SERT suite (all worklets) in sequence, each worklet in a new instance of the local JVM
- Advanced Research Mode
  - Selective execution of a subset of workloads and worklets.
  - Customization of worklet run-times and parameters.

SPEC's Server Efficiency Rating Tool™

SERT GUI

Host Discovery

Automatically gather detailed hardware and software configuration of the SUT

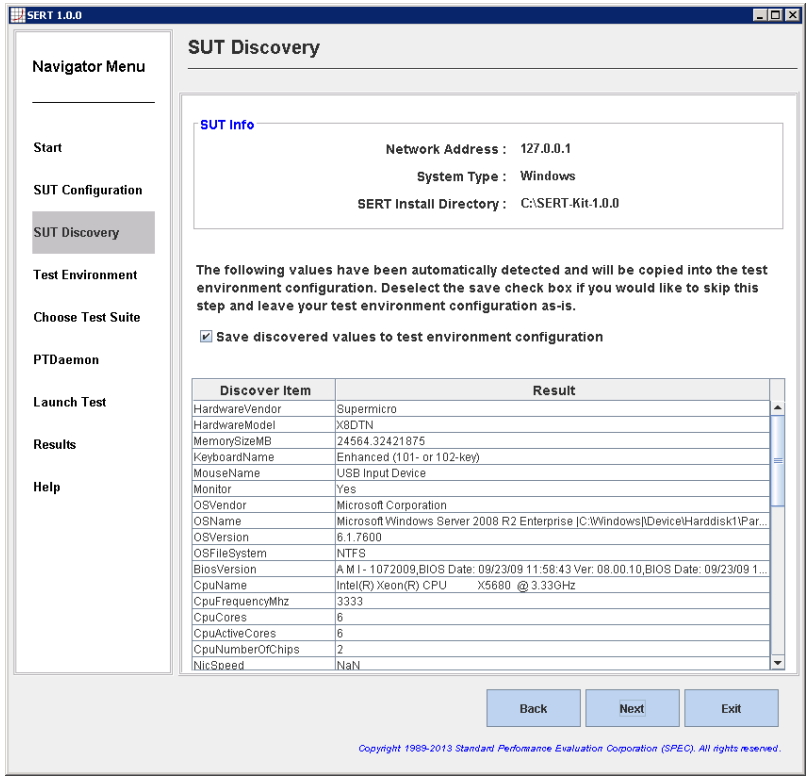

Launch Test

Monitor Progress of the currently executing worklet and the entire suite.

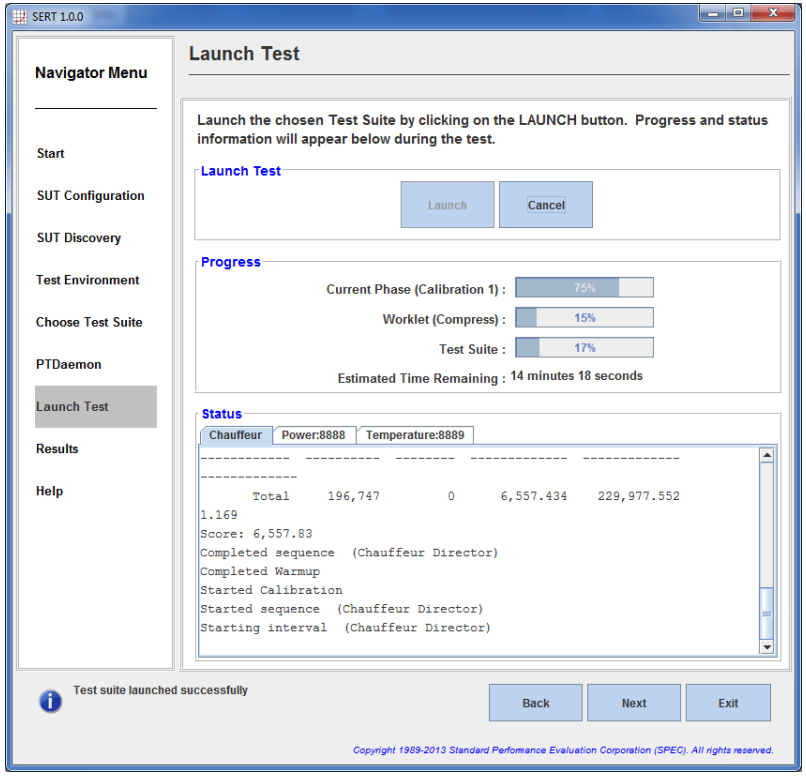

SPEC's Server Efficiency Rating Tool™

SERT Result Report 1/3

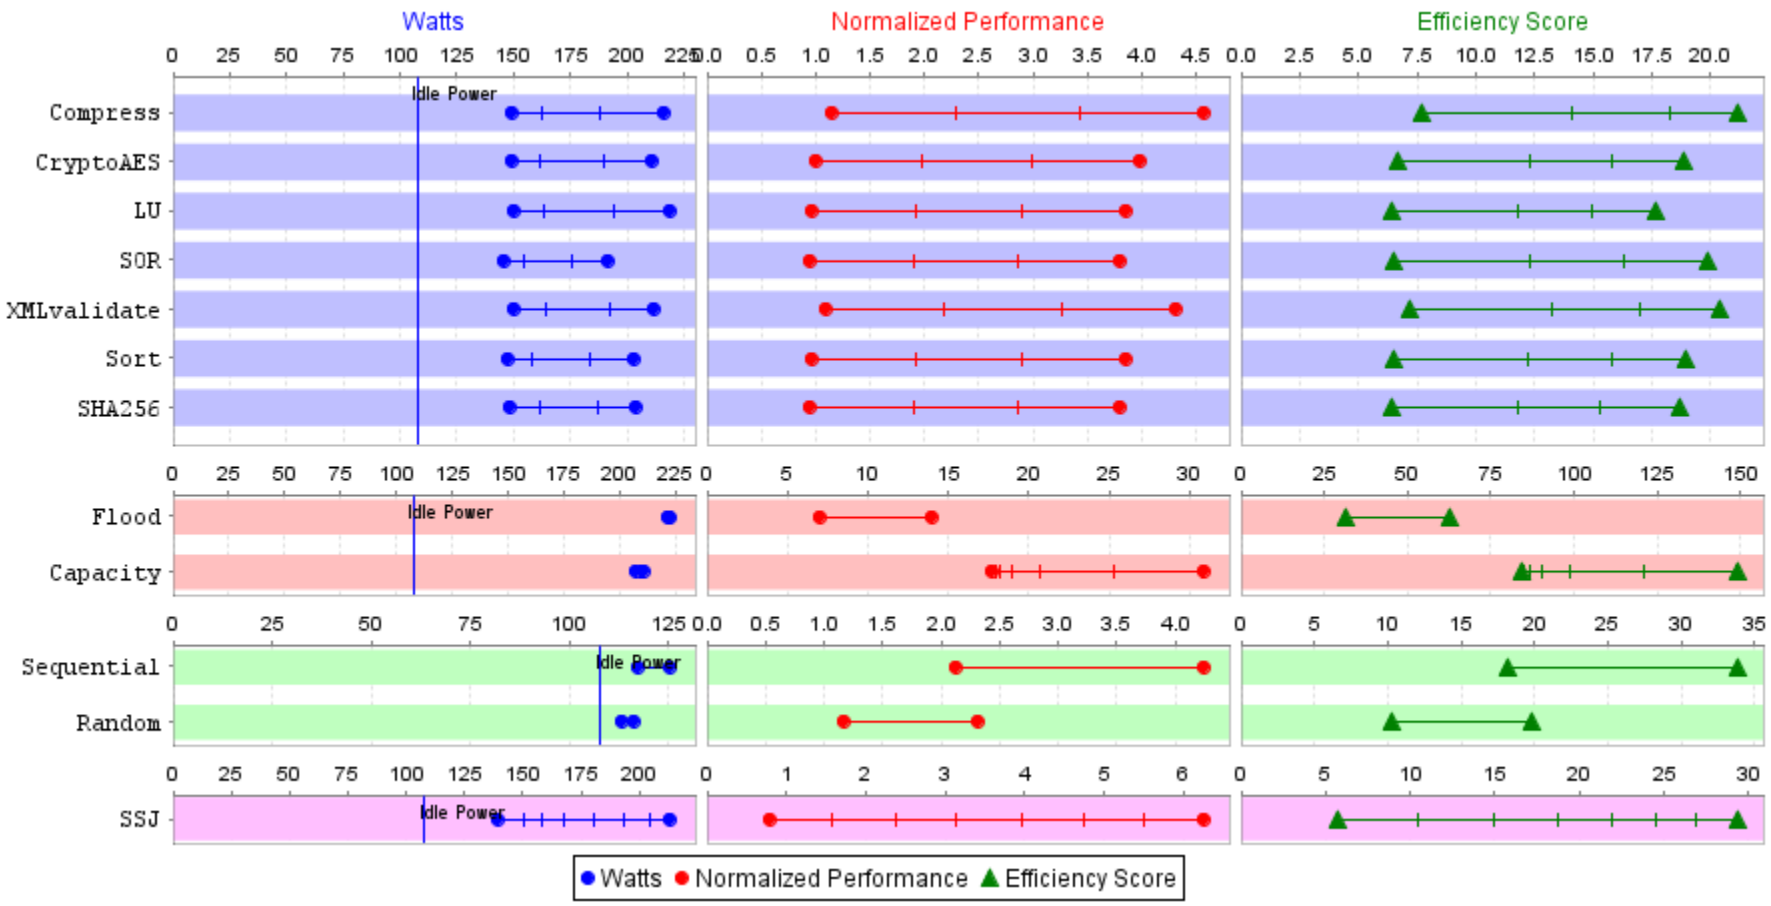

SPEC's Server Efficiency Rating Tool™
SERT Result Report 2/3

- Results Summary Table

Workload	Worklet	Normalized Performance	Watts at Lowest Load Level	Watts at Highest Load Level	Efficiency Score
CPU	Compress	0.947	60.1	80.2	8.465
	CryptoAES	1.015	61.4	77.9	9.086
	LU	1.013	62.1	82.0	8.835
	SOR	1.002	60.2	73.8	9.293
	XMLvalidate	0.995	61.5	80.1	8.877
	Sort	1.031	60.9	76.8	9.294
	SHA256	1.009	60.8	77.1	9.201
Memory	Flood	0.645	77.3	77.3	8.350
	Flood_Half	0.641	77.1	77.1	8.313
	Capacity_4	1.000	77.8	77.8	12.855
	Capacity_8	0.763	78.5	78.5	9.725
	Capacity_16	0.682	78.7	78.7	8.669
	Capacity_32	0.651	78.8	78.8	8.264
	Capacity_64	0.634	79.0	79.0	8.021
	Capacity_128	0.632	78.9	78.9	8.014
	Capacity_256	0.620	78.9	78.9	7.858
	Capacity_512	0.619	78.9	78.9	7.852
	Capacity_1024	0.616	78.9	78.9	7.804
Storage	Sequential	0.974	54.0	55.6	13.274
	Random	1.032	55.0	55.7	13.963
Hybrid	SSJ	1.006	56.4	79.9	8.313
Idle	Idle	n/a	52.1	52.1	n/a

SPEC's Server Efficiency Rating Tool™
SERT Result Report 3/3

- SUT Hardware Description

Hardware per Node (1 Node)			
Hardware Vendor		**Number and size of DIMM**	2 x 4 GB
Model		**Total Memory Amount**	8 GB
Form Factor	Tower	**Total Memory Available to OS**	8.0 GB
CPU Name	Celeron G1101	**Memory Operating Mode**	Independent
CPU Frequency	2267 MHz (up to N/A MHz), N/A	**Memory Details**	4GB 2Rx8 PC3-10600E ECC CL9; slots 1 and 4 populated
Number of CPU Sockets (available / populated)	1 / 1	**Power Supply Quantity (bays / populated / active)**	1 / 1 / 1
CPU(s) Enabled	2 cores, 1 processors, 2 cores/processor	**Power Supply Details**	1 x 800W, DPS-800GB-1 A
Number of NUMA Nodes	1	**Power Supply Operating Mode**	Standard
Hardware Threads	2 (1/core)	**Disk Controller**	onboard SATA controller
Primary Cache	32KB I + 32KB D on chip per core	**Disk Drive**	1 x WD2502ABYS (250 GB SATA 7200 RPM)
Secondary Cache	256KB I+D on chip per core	**Disk Controller**	onboard SATA controller
Tertiary Cache	2MB I+D on chip per chip	**Disk Drive**	1 x ST38001AS (80 GB SATA 7200 RPM)
Other Cache	None	**Optical Drives**	Yes
Other CPU Characteristics	None	**Network Interface Cards**	2 x Intel 82575EB Gigabit Network Connection (onboard) 2 enabled in firmware, 2 enabled in OS, 1 connected 1000 Mbit/s
Management Controller or Service Processor	Yes	**Keyboard**	PS2
Other Hardware	0 x None	**Mouse**	PS2
		Monitor	Yes

SPEC's Server Efficiency Rating Tool™
Efficiency Program Use / Regulations

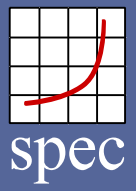

- Execution
  - Depending on efficiency program regulations, e.g. Certification Bodies may be required (qualified labs, currently US only)
  - Acceptable tuning parameters will be defined by processor manufacturers and must be used for SERT testing
- Reporting
  - User must provide a predefined set of information describing the hardware and software used for this report, partly discovered automatically by SERT
  - Performance and power data, as defined in the SERT result file, must be sent to efficiency program institutions
  - SERT results will be publicly available from efficiency program institutions, the format may vary from the original SERT result file, in addition aggregated data formats may be available
- Metrics
  - Will be defined after reporting only phase based on available results, critical mass of results is necessary
- Acceptance criteria for efficiency labels
  - Will be defined by efficiency program institutions
  - Initial criteria without SERT scores, e.g. reporting only

SPEC's Server Efficiency Rating Tool™
SERT Metric/Score

- Configuration power/performance modifier
  - "Substitution" for real measurements for items SERT can not measure or the performance can not be determined (e.g., redundant power supplies)
  - TBD, based on data collection phase.
- Each worklet will produce a measure representing the performance achieved by the SUT as well as the average power consumption at multiple target load levels. An overall score(s) is not provided and not recommended.
- Complexity of performance and power measures across components at multiple target load levels makes creation of a metric difficult.
- Recommend to implement a 9-12 month reporting-only phase first. At successful completion, SPEC will recommend a data-driven metric and scoring algorithm.

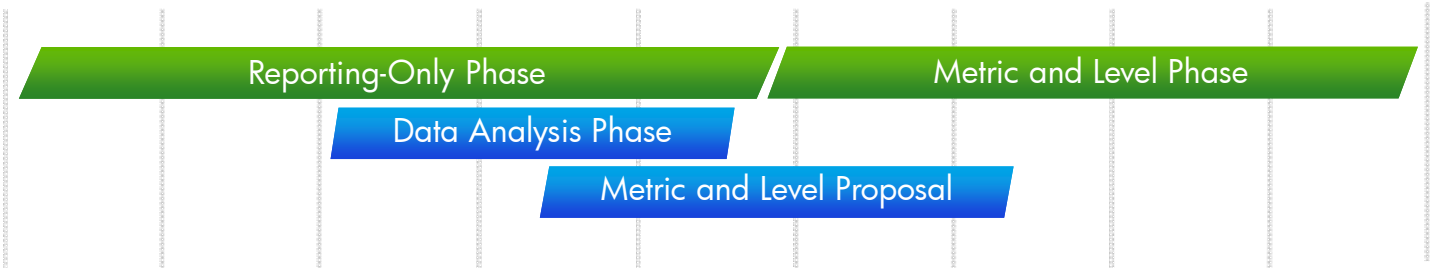

Universal Usage

SERT is applicable to programs world-wide

SPEC anticipates the use of SERT in many programs
- As more programs adopt the use of SERT, the base of data grows

SERT will be an outstanding research tool
- Versatile and flexible tool to test different aspects of computer servers at a variety of stress levels

SERT can be an excellent tool for energy efficiency programs
- Breadth of functional coverage allows for broader span of configurations
- "Tool" not "Benchmark" allows for broader span of configurations
- Near out-of-box tuning provides relevance to consumers

SPEC's Server Efficiency Rating Tool™

Q&A

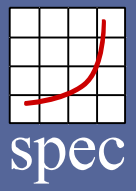

sertsupport@spec.org

www.spec.org/sert

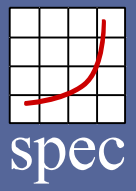

Thank you!

Acknowledgements

SPEC would like to acknowledge the people who have contributed to the design, development, testing and overall success of the SERT.

Development Team

- Christian Koopman, David Ott, Greg Darnell, Hansfried Block, Jeremy Arnold, John Beckett, Karin Wulf, Klaus-Dieter Lange, Mike Tricker, Nathan Totura, Sanjay Sharma, Karl Huppler, and Van Smith

Development Support

- Charles Pogue, David Mulnix, Peter Klassen, Shreeharsha G. Neelakantachar, and Thomas Brand

Administrative Support

- Bob Cramblitt, Cathy Sandifer, Charles McKay, Dianne Rice, Elden Sodowsky, and Jason Glick

References

Server Efficiency Rating Tool home page

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

Server Efficiency Rating Tool public Design Document:

- http://www.spec.org/sert/docs/SERT-Design_Doc.pdf

ENERGY STAR Enterprise Servers V2 home page

- https://www.energystar.gov/products/specs/node/142