

Klaus-Dieter Lange Chair, SPECpower Committee, SPEC



SPEC's Server Efficiency Rating ToolTM Agenda



2012 Distributed Management Task Force, INC. (DMTF)

Alliance Partner Technical Symposium

Storage Network Industry Association (SNIA) Green Storage Initiative

Portland, OR - 25th July 2012

- SPEC Overview and Philosophy
- SPECpower Milestones
- SPEC's Server Efficiency Rating ToolTM
- Government Program Implementation Recommendations
- Q & A

SPEC's Server Efficiency Rating ToolTM 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 (20,000+)
- In-depth understanding of workloads, benchmark code, fair comparisons across different platform

SPEC's Server Efficiency Rating ToolTM SPEC's Philosophy



4

- 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

SPEC's Server Efficiency Rating ToolTM Milestones 1/2



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 ToolTM Milestones 2/2



SPECpower_ssj2008 (http://www.spec.org/power_ssj2008/)

- First industry standard benchmark (released end of 2007) that measures the power and performance characteristics of server-class compute-equipment.
- A leap forward in power and performance analysis of Enterprise Computer Servers
- 350+ published results at SPEC by 20+ companies world-wide
- Over 7.2x energy efficiency gain since SPECpower_ssj2008 released



SPEC's Server Efficiency Rating ToolTM SPECpower_ssj2008



Variable System Utilization



SPEC's Server Efficiency Rating ToolTM SPECpower_ssj2008



Metric





Power and Performance at mutitiple Target Load Levels

SPEC's Server Efficiency Rating ToolTM SPECpower_ssj2008 - FDR

spec



SPECpower_ssj2008

Copyright © 2007-2009 Standard Performance Evaluation Corporation

			SP	ECpower_ssj2008 =	1,813 overall ssj_ops/watt
Test Sponsor:		SPEC License #:		Test Method:	Single Node
Tested By:		Test Location:		Test Date:	Mar 24, 2009
Hardware Availability:	May-2009	Software Availability:	May-2009	Publication:	Apr 8, 2009
System Source:	Single Supplier	System Designation:	Server	Power Provisioning:	Line-powered

Benchmark	Results	Summary
-----------	---------	---------

Performance		Power	Porformoneo to	
Target Load	Actual Load	ssj_ops	Average Active Power (W)	Power Ratio
100%	100.9%	439,831	170	2,586
90%	90.3%	393,612	158	2,492
80%	80.2%	349,870	147	2,373
70%	70.2%	305,949	135	2,268
60%	60.0%	261,804	127	2,064
50%	50.2%	218,789	120	1,830
40%	40.1%	174,886	112	1,555
30%	30.0%	130,657	105	1,241
20%	19.9%	86,740	98.0	885
10%	9.8%	42,949	87.8	489
	Active Idle	0	66.3	0
∑ssj_ops / ∑power =			1,813	



SPEC's Server Efficiency Rating ToolTM SPECpower_ssj2008



"The first version of SPECpower was a milestone for the computing industry and V1.10 significantly extends the benchmark's functionality. We will continue to work with the SPECpower committee on future Energy Star initiatives."

Andrew Fanara Energy Star Team Leader Environmental Protection Agency (EPA)



SPEC's Server Efficiency Rating ToolTM Design, Implementation and Delivery



"The SPECpower committee is currently working on the design, implementation and delivery of the next generation tool set that will measure and evaluate the performance and power of computer servers over a wider spectrum of functionality. This tool set will have the potential to be adopted for the use in the EPA Energy Star® program for Computer Server and to form a solid basis for the next generation SPECpower benchmarks."

Klaus-Dieter Lange SPECpower Committee Chairman Standard Performance Evaluation Corporation spec[®]

SPEC's Server Efficiency Rating ToolTM The SERT Development Model



- Delivered via an industry-wide collaboration:
 - Anyone willing to contribute effort is welcome to join the SPECpower Committee 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 ^(C)
 - SPEC 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

SPEC's Server Efficiency Rating ToolTM 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 ToolTM 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
 - $\hfill\square$ No resource committed for develop and test

Platform/OS

 SERT 1.0.0.0 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 ToolTM 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

SPEC's Server Efficiency Rating ToolTM Worklet



SERT's initial workload will be a collection of synthetic worklets

- Worklets design guidelines:
 - Worklets to assess CPU, Memory, Storage IO subsystem
 - □ Network IO will be handle by configuration power/performance modifiers
 - 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 ToolTM SERT Worklet Candidates



Workload	Worklet candidate	Alpha	Beta 1	Beta 2	RC1
CPU	Compress	Included	Included	Included	TBD
CPU	CryptoAES	Included	Included	Included	TBD
CPU	SOR	Included	Included	Included	TBD
CPU	SORT	Included	Included	excluded	TBD
CPU	SHA256	-	-	Included	TBD
CPU	FFT	Included	Included	excluded	TBD
CPU	LU	Included	Included	Included	TBD
CPU	XMLvalidate	Included	Included	Included	TBD
Memory	Flood	Included	Included	Included	TBD
Memory	Capacity	Included	Included	Included	TBD
Storage	Random	-	Included	Included	TBD
Storage	Sequential	-	Included	Included	TBD
Storage	Mixed	Included	Included	excluded	TBD
Hybrid	SSJ	Included	Included	Included	TBD
Idle	Idle	Included	Included	Included	Included

Please see the SERT Design Document for detail description of each worklet: <u>http://www.spec.org/sert/docs/SERT-Design_Doc.pdf</u>

SPEC's Server Efficiency Rating ToolTM SERT Workload



Worklet Execution

- In the current design the worklets will run consecutively, each in its own phase.
 This allows the generation of independent scores at different load levels
- Selected worklets or additional worklets specifically designed for concurrent execution may be run simultaneously in an extra phase. Concurrent execution of worklets will increase more realistic task-switching, especially using IO load modules

Load Levels

 Multiple load levels are a design goal of SERT. The active idle load level as well as a 100% workload level (not max power) are already good candidates.
 Prototype testing will show which levels will be included and if any weighting will be necessary

SPEC's Server Efficiency Rating ToolTM 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 Interval (120 Sec)

Total minimum execution time per workload = 830 sec (13:50 min)

plus worklet initialization times

SPEC's Server Efficiency Rating ToolTM Hardware and Software Configuration



Controller and SUT



SPEC's Server Efficiency Rating ToolTM 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 ToolTM SERT GUI



Host Discovery

Gathered automatically detailed hardware and software configuration of the SUT

SERT v.0.1					- 0	x		
Navigator Menu	Host Discovery							
	Host Info							
Start		Network Address : 12	7.0.0.1					
		System Type : wi	ndows					
Host Configuration	SERT	Install Directory : C	SERT-Kit-pre-Beta1					
Host Discovery		Path to Java exe : C:Program Files/Java/jre7/bin/java.exe						
Test Environment	Click START to Discover Host Software/Hardware Configuration: Start Cancel							
Choose Test Suite								
	Discover Item		Result					
PTDaemon	SystemVendor	Hewlett-Packard				^		
	SystemManufacturer	Hewlett-Packard				_		
Launch Test	SystemModel	HP EliteBook 8540w						
	SystemFormFactor	12						
	CanNama	12 Intal/P) Coro(TM) i5 CBI	I M 540 @ 2	52CU-				
Kesults	ChuMaxClockSneed	2534	J INI J40 (@ 2	.556112				
	CnuNumberOfProc	1						
Help	CnuL2CacheSize	256						
	TotalPhysicalMemory	4145287168						
	Capacity	2147483648						
	Capacity	2147483648						
	BankLabel	BANK 0						
	Description	Physical Memory						
	DeviceLocator	Top-Slot 1(top)						
	PartNumber	EBJ21UE8BDS0-DJ-F						
	SerialNumber	42920C68						
	BankLabel	BANK 2				-		
	(13aa aanakia a	lillion and a family star				-		
Host discovery suc	cessful		Back	Next	Exit			
•				SPEC Co	rporation- SERT 20	011		

Launch Test

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

Z SERT v.0.1		x
Navigator Menu	Launch Test: CPU Compress Mini	
Start	Launch the chosen Test Suite by clicking on the LAUNCH button. Progress and status information will appear below during the test.	
Host Configuration	Launch Test	
Host Discovery		
Test Environment	Current Phase (Calibration 1) : 29%	
Choose Test Suite	Worklet (CPU_Compress) : 15%	
PTDaemon	Test Suite : 33%	
Launch Test	Comment of the second s	_
Results	TxCompress 65,522 0 2,183.816 107,078.840	
Help	1.634 	
	Total 65,522 0 2,183.816 107,078.840	
	Score: 2,183.835 Completed Sequence Chauffeur Director Completed Warmup Started Calibration Started Sequence Chauffeur Director	
Test suite launched	successfully Back Next Exit	
	SPEC Corporation- SERT 201	1

SPEC's Server Efficiency Rating ToolTM SERT Metric/Score



Draft markup



SPEC's Server Efficiency Rating ToolTM 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)
 TRD_based on data collection phase
 - □ 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.



SPEC's Server Efficiency Rating ToolTM Public Beta 2 Program



"SERT will enhance the power of the ENERGY STAR label, providing critical information on energy use to consumers and institutional purchasers, and helping manufacturers distinguish their products in the market. If widely adopted, it should reduce testing burdens on the server industry and enable more meaningful and consistent comparisons among products sold in different regions of the world."

Robert Meyers Energy Star Team Leader Environmental Protection Agency (EPA)



SPEC's Server Efficiency Rating ToolTM 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 ToolTM Q&A



sertsupport@spec.org

www.spec.org/sert

SPEC's Server Efficiency Rating ToolTM



Thank you!

References 1/2



- Server Efficiency Rating Tool home page:
 - <u>http://www.spec.org/sert/</u>
- Server Efficiency Rating Tool public Design Document (latest version):
 - 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

• ENERGY STAR Computer Specification Version 1.0:

http://www.energystar.gov/ia/partners/product_specs/program_reqs/computer_server_pr_ og_req.pdf

References 2/2



- 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

• ENERGY STAR Computer Servers Draft 1 Version 2.0:

- 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



SPEC would like to acknowledge the current members of the SPECpower Committee who have contributed to the design, development, testing and overall success of the SERT.

SPEC, the SPEC logo and the tool and names SERT, SPECpower_ssj2008, SPECweb2009 and SPECvirt_sc2010 are registered trademarks of the Standard Performance Evaluation Corporation (SPEC).