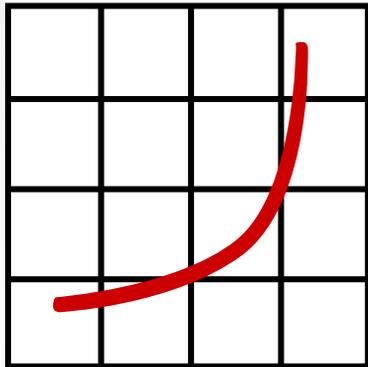


Server Efficiency Rating Tool™

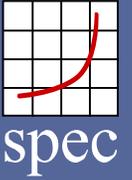
May 2010 – Development Update

Klaus-Dieter Lange

Chair, SPECpower Committee, SPEC



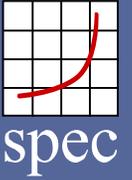
spec



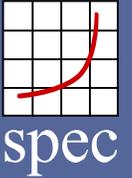
ENERGY STAR® Computer Server Version 2.0 – Draft 1 Stakeholder Webinar – May 27, 2010

- SPEC – Overview and Philosophy
- SPEC's Server Efficiency Rating Tool™
 - Goals and Requirements
 - Worklet
 - Workload
 - Scores and Metric
 - Scope (First Release)
 - Schedule
- Next Steps
- Q & A





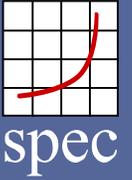
- A world-wide non-profit consortium formed in 1988 to establish, maintain and endorse a standardized set of relevant benchmarks that can be applied to the newest generation of high-performance computers
- Comprised out of over 80 computer hardware and software vendors, educational institutions and government 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



- 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

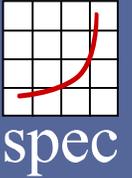
SPEC's Server Efficiency Rating Tool™

Goals and Requirements



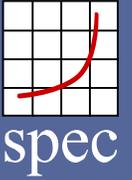
Server Efficiency Rating Tool (SERT)

- 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, Oracle have committed resources for the development of SERT in the SPECpower Committee



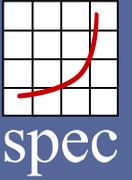
SERT 1.0.0 workload will be a collection of synthetic worklets

- Worklets design guidelines:
 - Worklets to assess CPU, Memory, Network IO, Storage IO subsystem
 - 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



- 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 which can be combined to an overall metric
 - Selected worklets or additional worklets specifically designed for concurrent execution may be run simultaneously in an extra phase. Concurrent execution of worklets will introduce more realistic task-switching, especially using IO load modules

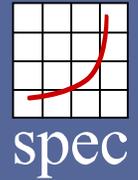
- Load Levels
 - Multiple load levels are a design goal of SERT. The active idle load level as well as a 100% workload level (not max power) are already good candidates. Prototype testing will show which levels will be included and if any weighting will be necessary



- Configuration power/performance adders
 - “Substitution” for real measurements for items SERT can not measure or the performance can not be determined (e.g., redundant power supplies)
 - TBD by EPA based on data collection
- Each worklet will produce a measure representing the performance achieved by the SUT, which then must be combined with the measures produced by the other worklets to yield a metric indicative of the overall performance of the SUT on all worklets used in the tool
- SPEC currently plans that the multiple performance measures produced in this manner be combined into a single metric as the geometric mean of the individual measures

SPEC's Server Efficiency Rating Tool™

Scope (First Release)



Processor sockets/compute nodes

- SERT 1.0.0 will be designed for and tested with
 - Max. 8 sockets per node
 - Max. 64 nodes
 - Limited to homogeneous and blade servers

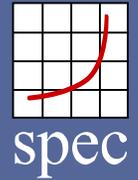
Platform/OS

- 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)	X86 (AMD)	X86 (Intel)	Itanium	Itanium	Power	Power	Power	SPARC
OS	Windows 2008 R2	LINUX	Windows 2008 R2	LINUX	Solaris	Solaris	HP-UX 11i	LINUX	AIX	IBM i	LINUX	Solaris

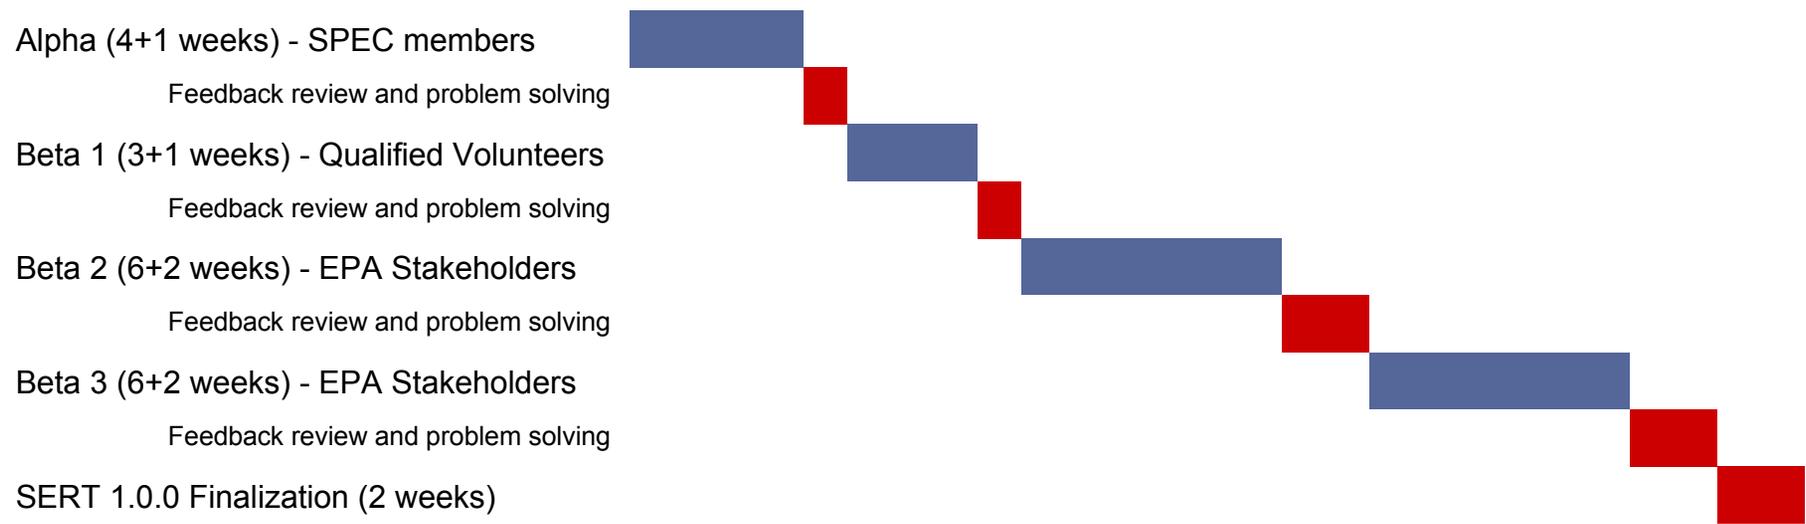
SPEC's Server Efficiency Rating Tool™

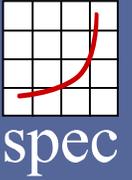
SERT Preliminary Schedule



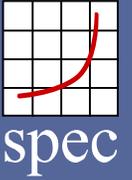
- Alpha test phase is planned to start in September 2010
- The start of each phase required successful completion of its predecessor

Test Phases





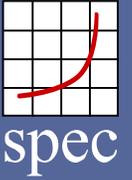
- Worklet selection/development and evaluation
- EPA schedule requirements challenge
 - Function/scope may be tradeoff for time
- Restricted tuning
 - Configuration as shipped– TBD by the EPA
 - Minimize software variability
- DC power measurement
 - No resource committed for development and testing
- SERT design document updates
 - Each ~5 weeks
- Analysis of EPA stakeholder feedback



info@spec.org

www.spec.org/sert

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

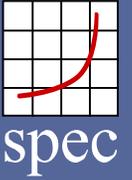


SPEC Power and Performance Methodology

- 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 Power Temperature Daemon (PTDaemon)

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

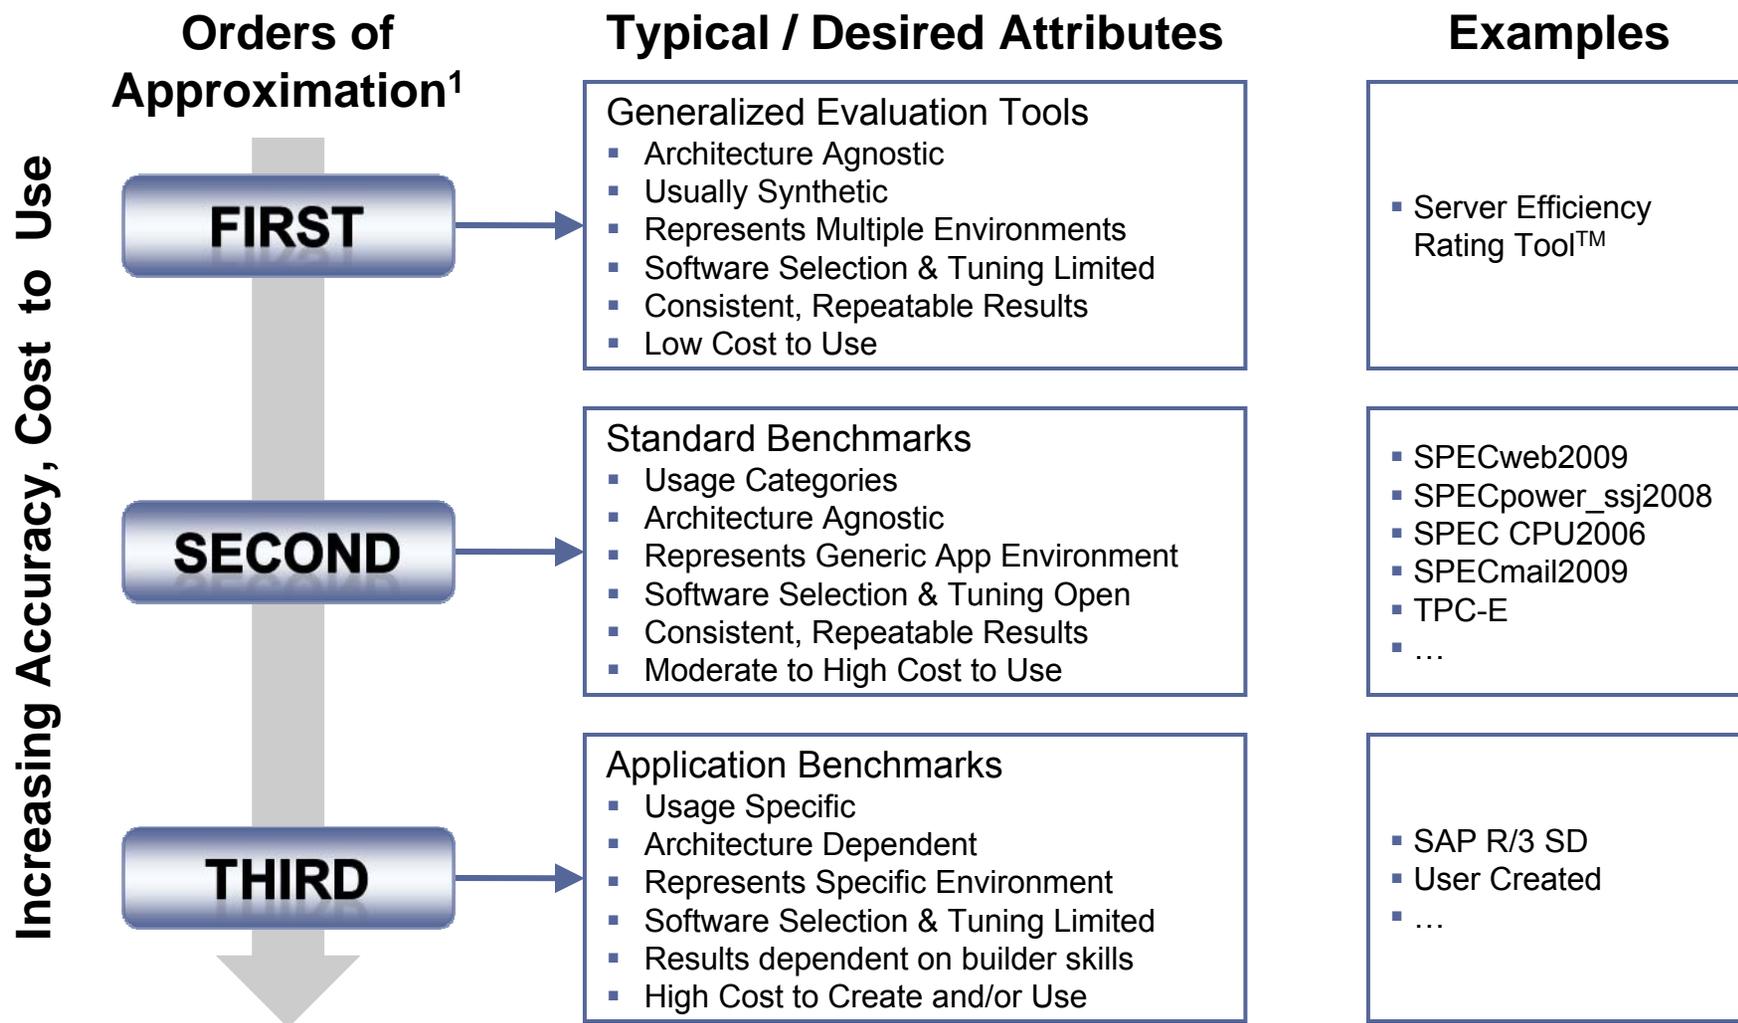


SPECpower_ssj2008

- First industry standard benchmark that measures the power and performance characteristics of server-class compute-equipment. A leap forward in power and performance analysis of Enterprise Computer Servers
- Released end of 2007
- 160+ published results at SPEC
- Submission from 20 companies world-wide
- Over 4.4x energy efficiency gain since SPECpower_ssj2008 released
 - 698 - 3106 overall ssj_ops/watt
(best result from first submission cycle to the best current result)
- http://www.spec.org/power_ssj2008/

SPECweb2009

- First industry standard benchmark that measures the power and performance characteristics of web serving compute-equipment
- <http://www.spec.org/web2009/>



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