SPEC in a New Era

Matthijs van Waveren
SPEC Director
Fujitsu Systems Europe

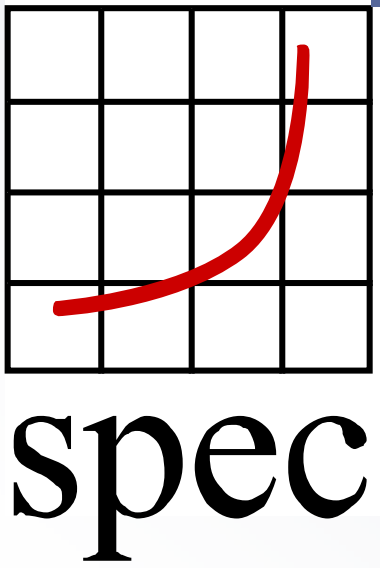

Overview

SS-ken – 22 October 2008

- SPEC in a New Era
- Low Power Consumption
- Compute Intensive Applications
- Server Virtualization
- Conclusion

SPEC in a New Era

The world is evolving towards green servers, VoIP, grid computing, server virtualization ….

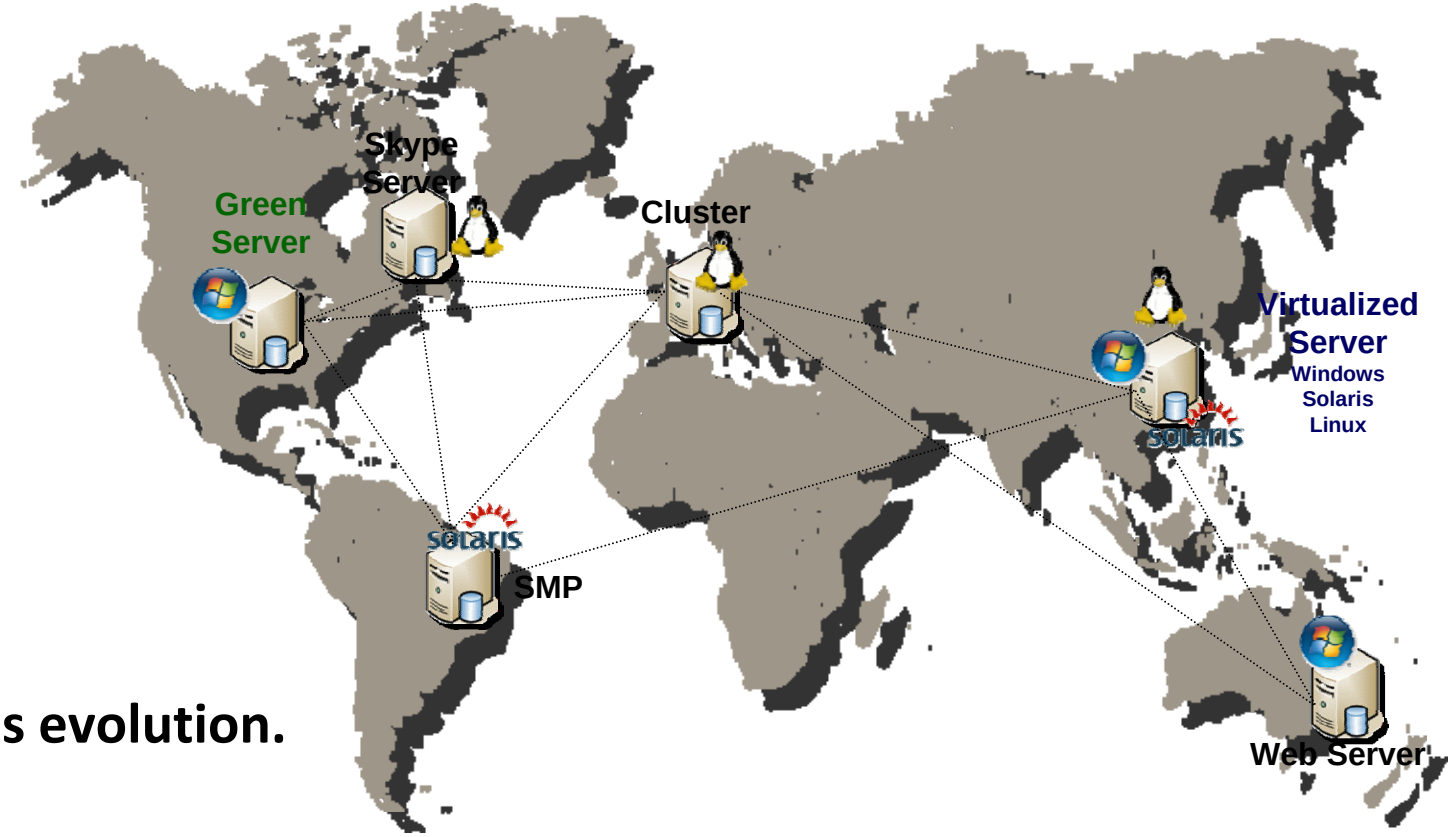

SPEC follows this evolution.

SPEC in a New Era

How do you choose suitable systems for this new era ?

Go look on the SPEC web site &
compare the published results !

Acquire & run the SPEC benchmark codes !

SPEC in a New Era

Mission Statements

- Develop metrics that allow end-users to differentiate hardware and software and find the most suitable systems to run their applications.
- Evolve with the marketplace in order to stay useful for end-users.

SPEC in a New Era

SPEC develops software that helps you in finding the best system for:

- High Performance Applications on Single and Multiple CPUs
- Low Power Consumption
- Graphics Applications
- Server Virtualization (under development)
- Instant Messaging (under development)
- Mail and Web Servers

SPEC in a Nutshell

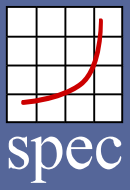

What is the Standard Performance Evaluation Corporation ?

- A world-wide non-profit consortium formed 20 years ago to establish, maintain and endorse a standardized set of relevant benchmarks that can be applied to the newest generation of high-performance compute equipment.
- Over 80 computer hardware and software vendors, and educational institutions from all over the world.
- Development of benchmark suites to ensure that the marketplace has a fair and useful set of metrics to differentiate systems.
- Review and publication of submitted results produces a large public repository of well documented, peer reviewed, benchmark results.

SPEC is a global organization

SPEC Groups

- High Performance Group (HPC systems)
  - OMP (OpenMP benchmark)
  - MPI (MPI application benchmark)
- Open Systems Group (desktop systems, high-end workstations and servers)
  - CPU (CPU benchmarks)
  - JAVA (java client and server side benchmarks)
  - MAIL (mail server benchmarks)
  - SFS (file server benchmarks)
  - WEB (web server benchmarks)
- Graphics Performance Groups (Graphics)
  - Apc (Graphics application benchmarks)
  - Opc (OpenGL performance benchmarks)

Low Power Consumption

SPECpower™

- First industry standard benchmark that measures the power and performance characteristics of server-class compute-equipment.
- SPEC's initiative to augment many existing SPEC benchmarks with power measurements and assist other non-profit industry standards organization.

SPEC Power and Performance Methodology

- An introduction on power and performance for computer systems, the *conclusions* from the 2 year development of SPECpower_ssj2008
- Guidance for Power and Performance benchmark development

SPECpower™
SPECpower_ssj2008 - Framework

Benchmark Harness - Framework

- Ensures the synchronization of the measured performance, power, and environmental data.

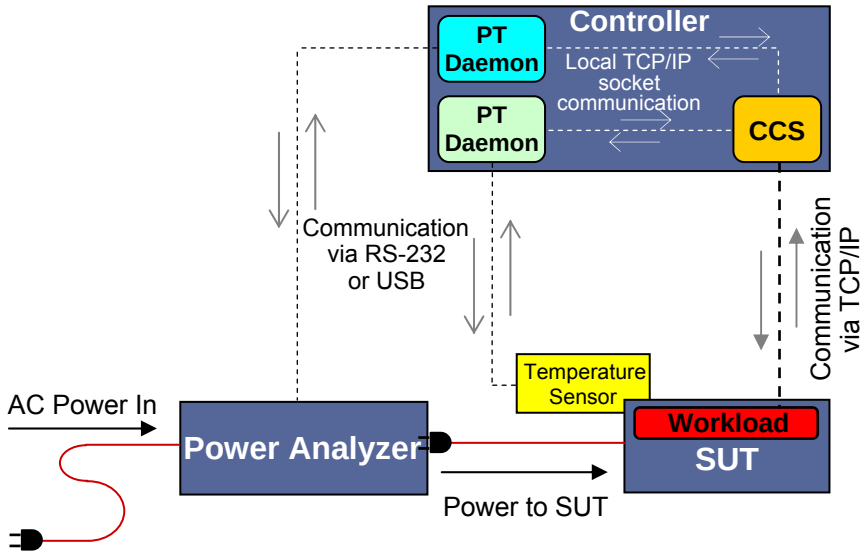

SPECpower™
SPECpower_ssj2008 – Variable Utilization

Variable System Utilization

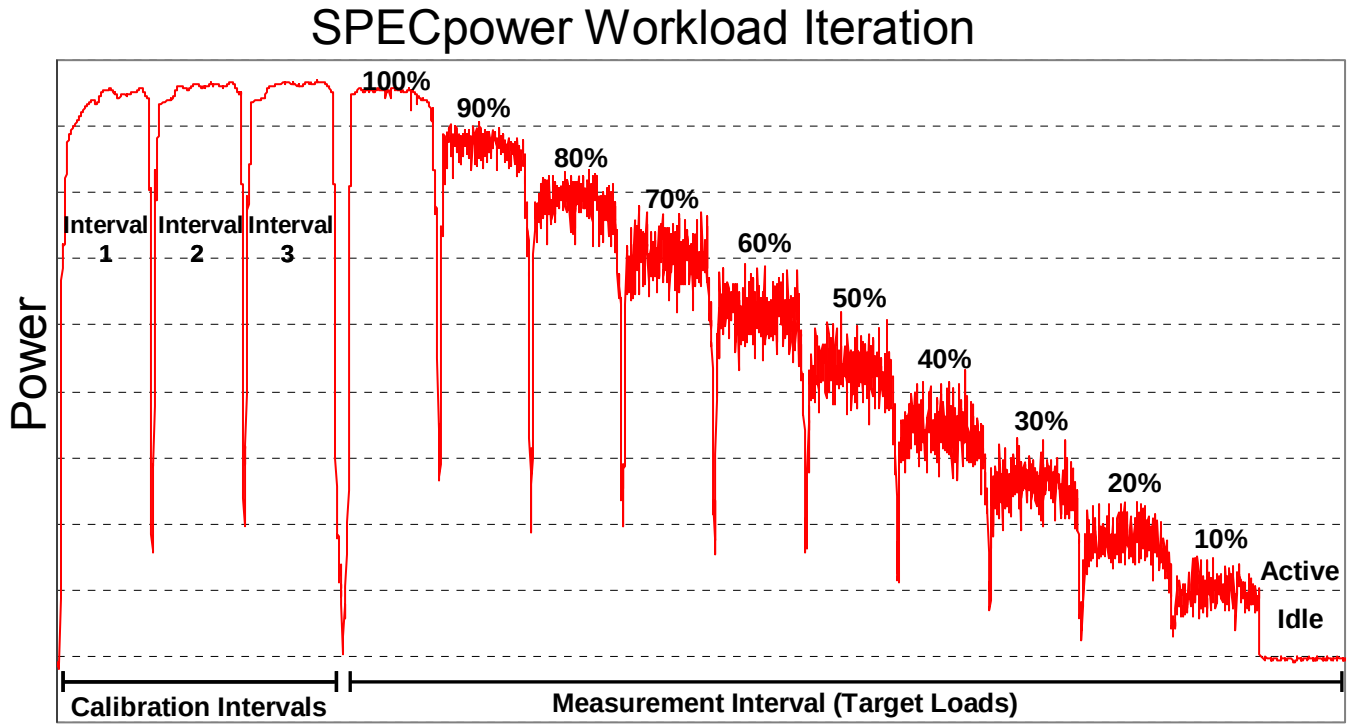

SPECpower™ SPECpower_ssj2008 metric

$$\text{SPECpower\_ssj2008} = \frac{\sum Performance}{\sum Power} = \frac{1{,}861{,}107 \text{ ssj\_ops}}{7{,}010 \text{ watt}} = 265 \text{ overall} \frac{\text{ssj\_ops}}{\text{watt}}$$

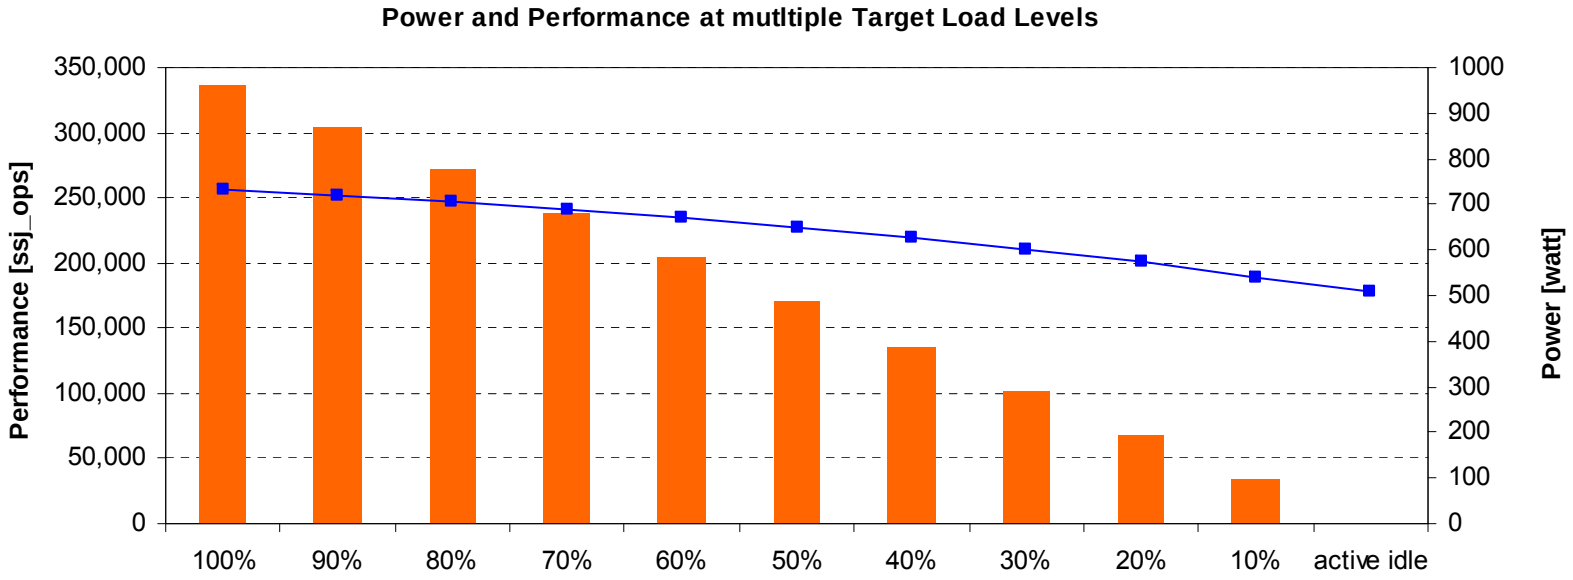

SPECpower™
SPECpower_ssj2008 on PRIMERGY

SPECpower_ssj2008

Fujitsu Siemens Computers PRIMERGY TX150 S6 (Intel Xeon X3360)				SPECpower_ssj2008 = 1,018 overall ssj_ops/watt	
Test Sponsor:	Fujitsu Siemens Computers	SPEC License #:	22	Hardware Availability:	Mar-2008
Tested By:	Fujitsu Siemens Computers	Test Location:	Paderborn, Germany	Software Availability:	Feb-2008
System Source:	Single Supplier	Test Date:	Mar 26, 2008	Publication:	Apr 2, 2008

Benchmark Results Summary

Performance			Power	Performance to Power Ratio
Target Load	Actual Load	ssj_ops	Average Power (W)	
100%	99.5%	184,145	116	1,586
90%	90.1%	166,892	113	1,475
80%	79.7%	147,598	109	1,351
70%	70.5%	130,538	105	1,243
60%	59.7%	110,607	99.4	1,113
50%	49.4%	91,420	93.0	983
40%	40.6%	75,086	86.5	868
30%	30.1%	55,795	79.0	706
20%	20.0%	37,116	72.3	514
10%	10.0%	18,424	66.0	279
	Active Idle	0	59.3	0
			∑ssj_ops / ∑power =	1,018

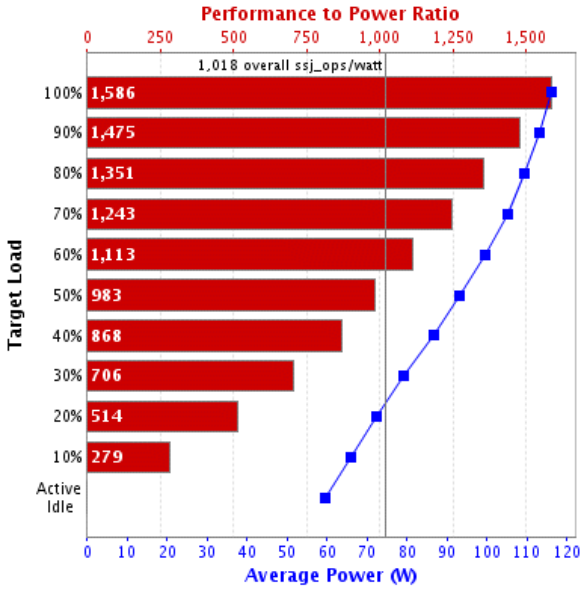

Compute Intensive Applications

- Integer and/or floating-point applications

  Use SPEC CPU2006

- OpenMP-based floating-point applications

  Use SPEC OMP2001

- MPI-based floating-point applications

  Use SPEC MPI2007

SPEC CPU2006

Consists of two benchmark suites:

- CINT2006
- CFP2006

Large number of submissions since August 2006:

- 803 for CINT2006
- 743 for CFP2006
- 1916 for CINT2006 Rates
- 1534 for CFP2006 Rates

SPEC CPU2006 Metrics

CINT2006 (for integer based performance comparisons):

- The geometric mean of twelve normalized ratios under peak or base tuning.
- The geometric mean of twelve normalized throughput ratios under peak or base tuning.

CFP2006 (for floating-point based performance comparisons):

- SPECfp2006: The geometric mean of seventeen normalized ratios with peak or base tuning.
- SPECfp_rate2006: The geometric mean of seventeen normalized throughput ratios with peak or tuning.

The geometric mean of a data set [a_1, a_2, a_3,..., a_n] is given by:

$$(a_1 * a_2 * a_3 * \ldots * a_n)^{1/n}$$

The geometric mean of a data set is less than or equal to the data set's arithmetic mean.

SPEC CPU2006 on PRIMERGY

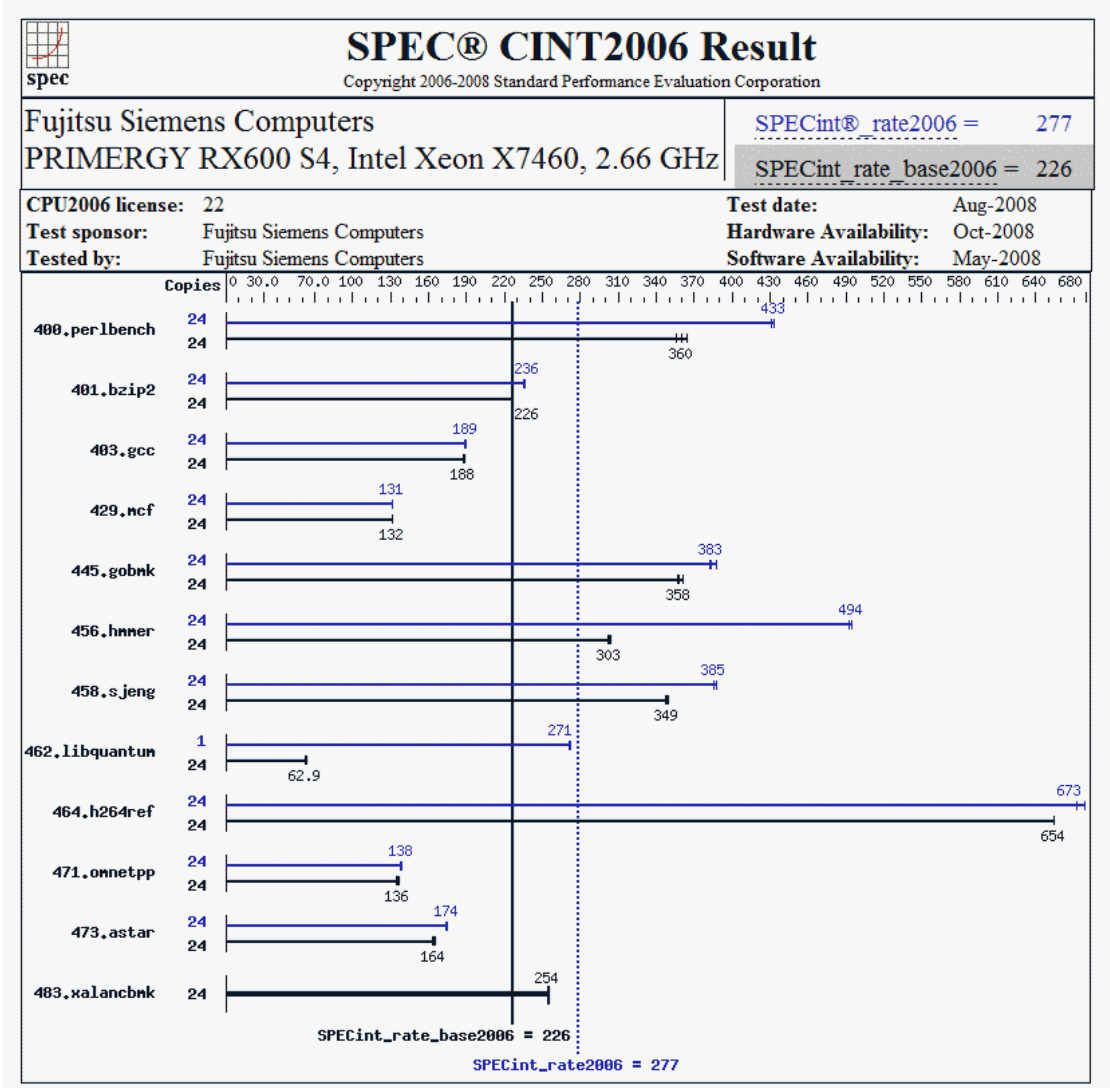

SPEC OMP2001 and MPI2007

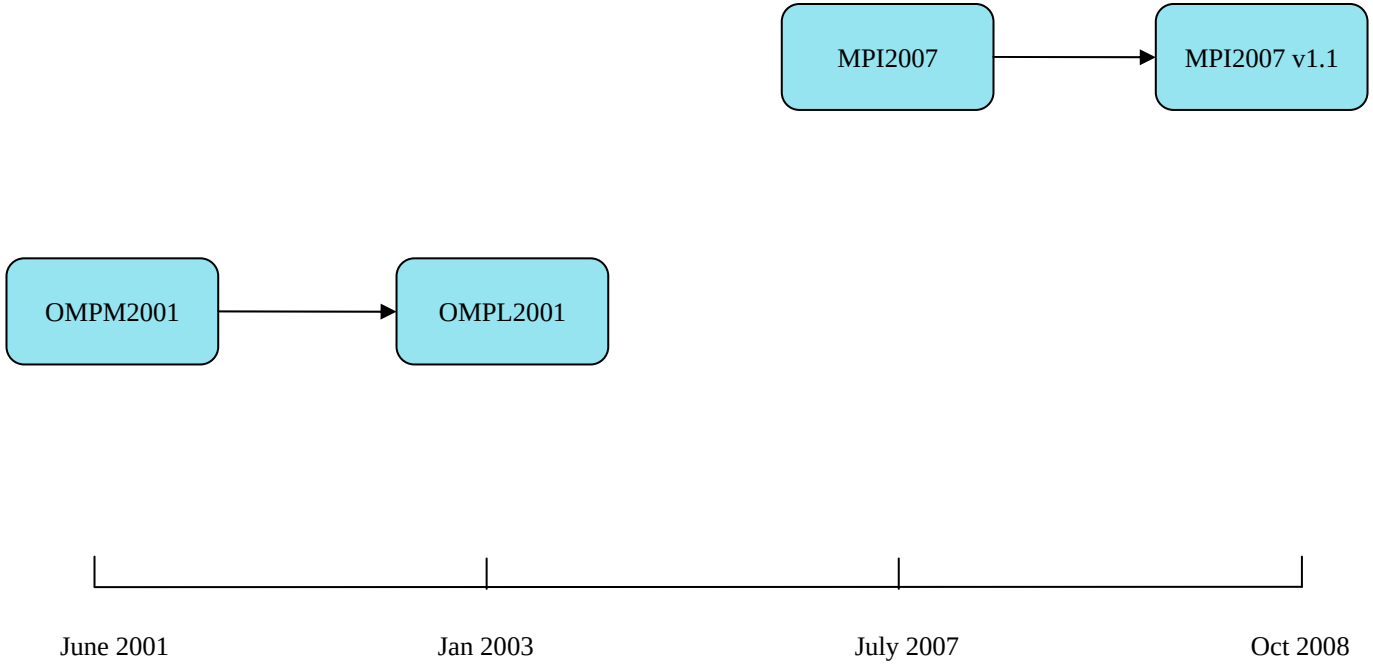

SPEC OMP2001

SPEC OMPM2001 is focused on 4-way to 16-way systems

SPEC OMPL2001 targets 32-way and larger systems

Number of submissions to date

- 217 for OMPM2001 (since June 2001 release)
- 56 for OMPL2001 (since May 2002 release)

Scientific Areas

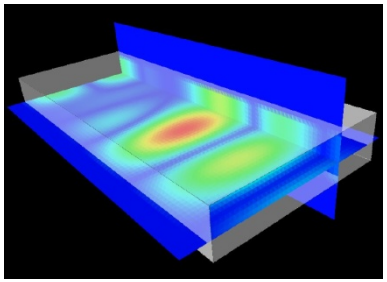

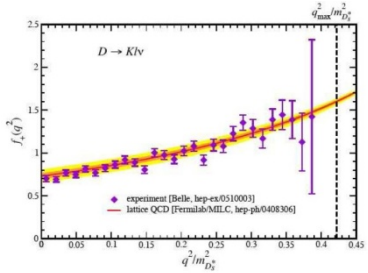


```
Sequence1   ACACGGTGTCGTATCATGCTGCAGGATGCT
            AGACTGCGTCANATGTTCGTACTAACTGTG
Sequence2   ACGCGGTGTCGTGTCATGCTACATTATGCT
            AGACTGCGTCGGATGCTCGTATTGACTGCG
Sequence3   ACGCGGTGCCGTGTNATGCTGCATTATGCT
            CGACTGCGRCGGATGCTAGTATTGACTGCG
Sequence4   ACGCGCTGCCGTGTCATCCTACACGATGCY
            AGACAGCGTCAGCTGCTAGTACTGGCTGAG
Sequence5   ACGCGCTGTCGTGTCATACTGCAGGATGCT
            AGACTGCGTCAGCTGCTAGTACTGGCTGAG
```

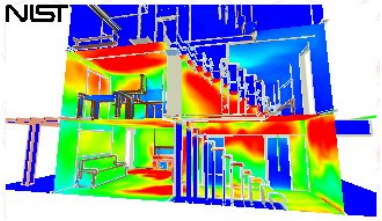

- Computational Fluid Dynamics (*applu, galgel*)
- Quantum Chromodynamics (*wupwise*)
- Air Pollution (*apsi*)
- Image recognition (*art*)
- Crash simulation (*fma3d*)
- Genetic algorithm (*gafort*)
- Earthquake modeling (*equake*)
- Weather prediction (*swim*)
- Multigrid solver (*mgrid*)
- Molecular Dynamics (*ammp*)

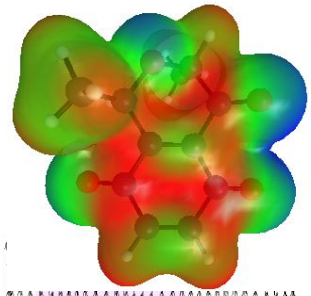

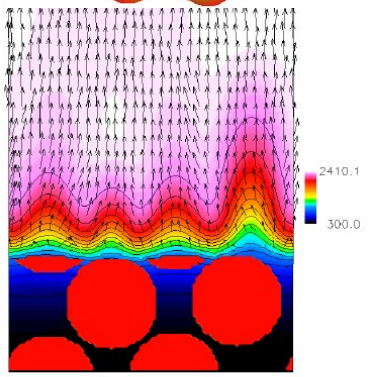

(b) 30 microns

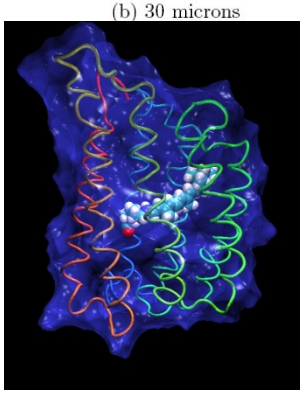

SPEC OPM2001 Metrics

SPEC OMPM2001

- *SPECompMpeak2001:* The geometric mean of 11 normalized ratios (peak tuning).
- *SPECompMbase2001:* The geometric mean of 11 normalized ratios (base tuning).

SPEC OMPL2001

- *SPECompLpeak2001:* The geometric mean of 9 normalized ratios (peak tuning).
- *SPECompLbase2001:* The geometric mean of 9 normalized ratios (base tuning).

SPEC OMP2001 on SPARC Enterprise

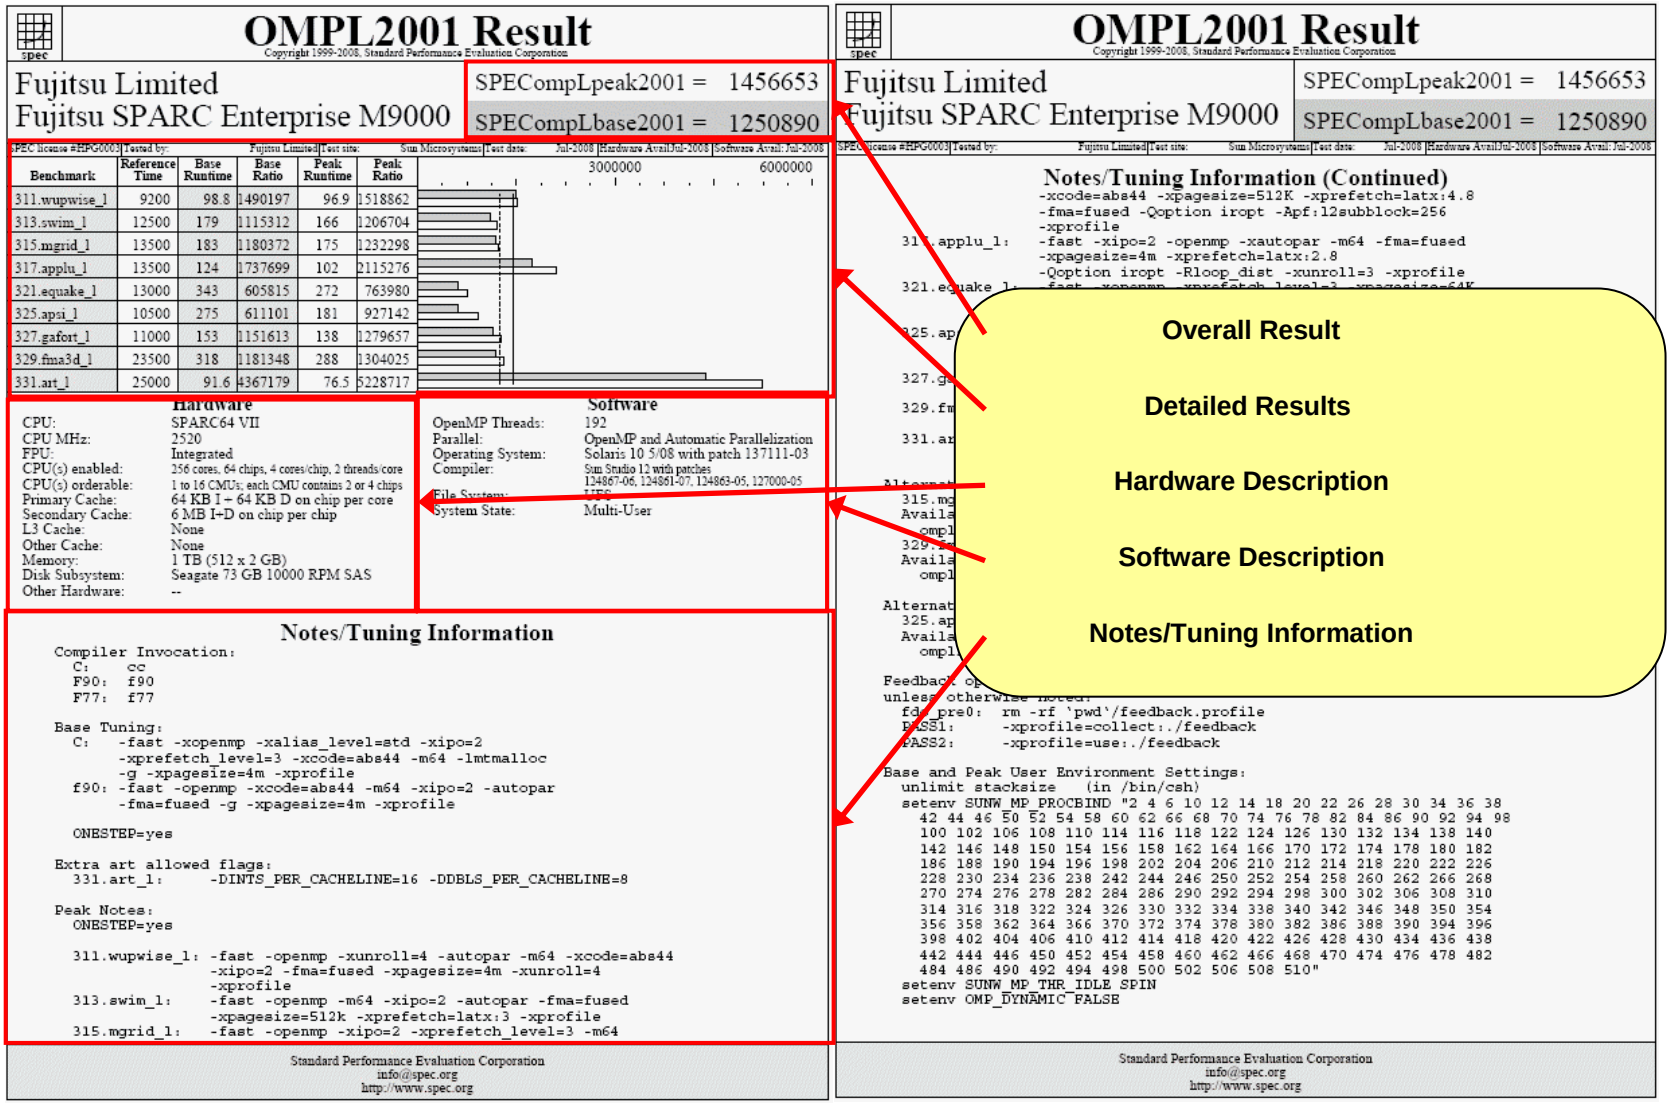

OMPL2001 Result

Copyright 1999-2008, Standard Performance Evaluation Corporation

Fujitsu Limited
Fujitsu SPARC Enterprise M9000

SPECompLpeak2001 = 1456653
SPECompLbase2001 = 1250890

SPEC license #HPG0003 | Tested by: Fujitsu Limited | Test site: Sun Microsystems | Test date: Jul-2008 | Hardware Avail: Jul-2008 | Software Avail: Jul-2008

Benchmark	Reference Time	Base Runtime	Base Ratio	Peak Runtime	Peak Ratio
311.wupwise_l	9200	98.8	1490197	96.9	1518862
313.swim_l	12500	179	1115312	166	1206704
315.mgrid_l	13500	183	1180372	175	1232298
317.applu_l	13500	124	1737699	102	2115276
321.equake_l	13000	343	605815	272	763980
325.apsi_l	10500	275	611101	181	927142
327.gafort_l	11000	153	1151613	138	1279657
329.fma3d_l	23500	318	1181348	288	1304025
331.art_l	25000	91.6	4367179	76.5	5228717

Hardware

CPU:	SPARC64 VII
CPU MHz:	2520
FPU:	Integrated
CPU(s) enabled:	256 cores, 64 chips, 4 cores/chip, 2 threads/core
CPU(s) orderable:	1 to 16 CMUs; each CMU contains 2 or 4 chips
Primary Cache:	64 KB I – 64 KB D on chip per core
Secondary Cache:	6 MB I+D on chip per chip
L3 Cache:	None
Other Cache:	None
Memory:	1 TB (512 x 2 GB)
Disk Subsystem:	Seagate 73 GB 10000 RPM SAS
Other Hardware:	--

Software

OpenMP Threads:	192
Parallel:	OpenMP and Automatic Parallelization
Operating System:	Solaris 10 5/08 with patch 137111-03
Compiler:	Sun Studio 12 with patches 124867-06, 124881-07, 124863-05, 127000-05
File System:	UFS
System State:	Multi-User

Notes/Tuning Information

```
Compiler Invocation:
  C:    cc
  F90:  f90
  F77:  f77

Base Tuning:
  C:    -fast -xopenmp -xalias_level=std -xipo=2
        -xprefetch_level=3 -xcode=abs44 -m64 -lmtmalloc
        -g -xpagesize=4m -xprofile
  f90: -fast -xopenmp -xcode=abs44 -m64 -xipo=2 -autopar
        -fma=fused -g -xpagesize=4m -xprofile

  ONESTEP=yes

Extra art allowed flags:
  331.art_l:      -DINTS_PER_CACHELINE=16 -DDBLS_PER_CACHELINE=8

Peak Notes:
  ONESTEP=yes

  311.wupwise_l: -fast -openmp -xunroll=4 -autopar -m64 -xcode=abs44
                 -xipo=2 -fma=fused -xpagesize=4m -xunroll=4
                 -xprofile
  313.swim_l:    -fast -openmp -m64 -xipo=2 -autopar -fma=fused
                 -xpagesize=512k -xprefetch=latx:3 -xprofile
  315.mgrid_l:   -fast -openmp -xipo=2 -xprefetch_level=3 -m64
```

Standard Performance Evaluation Corporation
info@spec.org
http://www.spec.org

OMPL2001 Result

Copyright 1999-2008, Standard Performance Evaluation Corporation

Fujitsu Limited
Fujitsu SPARC Enterprise M9000

SPECompLpeak2001 = 1456653
SPECompLbase2001 = 1250890

SPEC license #HPG0003 | Tested by: Fujitsu Limited | Test site: Sun Microsystems | Test date: Jul-2008 | Hardware Avail: Jul-2008 | Software Avail: Jul-2008

Notes/Tuning Information (Continued)

```
                -xcode=abs44 -xpagesize=512K -xprefetch=latx:4.8
                -fma=fused -Qoption iropt -Apf:l2subblock=256
                -xprofile
  317.applu_l:  -fast -xipo=2 -openmp -xautopar -m64 -fma=fused
                -xpagesize=4m -xprefetch=latx:2.8
                -Qoption iropt -Rloop_dist -xunroll=3 -xprofile
  321.equake_l: -fast -xopenmp -xprefetch_level=3 -xpagesize=64K

  325.ap

  327.g

  329.fm

  331.ar

Alternat
  315.mg
  Availa
    ompl
  329.f
  Availa
    ompl

Alternat
  325.ap
  Availa
    ompl

Feedback op
unless otherwise noted:
  fdo_pre0:  rm -rf `pwd`/feedback.profile
  PASS1:     -xprofile=collect:./feedback
  PASS2:     -xprofile=use:./feedback

Base and Peak User Environment Settings:
  unlimit stacksize   (in /bin/csh)
  setenv SUNW_MP_PROCBIND "2 4 6 10 12 14 18 20 22 26 28 30 34 36 38
    42 44 46 50 52 54 58 60 62 66 68 70 74 76 78 82 84 86 90 92 94 98
    100 102 106 108 110 114 116 118 122 124 126 130 132 134 138 140
    142 146 148 150 154 156 158 162 164 166 170 172 174 178 180 182
    186 188 190 194 196 198 202 204 206 210 212 214 218 220 222 226
    228 230 234 236 238 242 244 246 250 252 254 258 260 262 266 268
    270 274 276 278 282 284 286 290 292 294 298 300 302 306 308 310
    314 316 318 322 324 326 330 332 334 338 340 342 346 348 350 354
    356 358 362 364 366 370 372 374 378 380 382 386 388 390 394 396
    398 402 404 406 410 412 414 418 420 422 426 428 430 434 436 438
    442 444 446 450 452 454 458 460 462 466 468 470 474 476 478 482
    484 486 490 492 494 498 500 502 506 508 510"
  setenv SUNW_MP_THR_IDLE SPIN
  setenv OMP_DYNAMIC FALSE
```

Standard Performance Evaluation Corporation
info@spec.org
http://www.spec.org

Overall Result

Detailed Results

Hardware Description

Software Description

Notes/Tuning Information

Multicore Chips

SPEC's view on fair comparisons of systems with multicore chips:

Ask submitters to specify full system description in the submissions and in the marketing press releases:

- number of chips
- number of cores
- number of cores per chip
- number of threads per core

SPEC MPI2007

SPEC MPI2007 focuses on performance of applications using the Message-Passing Interface (MPI), up to 128 ranks.

The following components are emphasized:

- the type of computer processor (CPU),
- the number of computer processors,
- the communication interconnect,
- the memory architecture,
- the compilers, and
- the shared file system.

Number of submissions to date

- 61 since July 2007 release

Scientific Areas

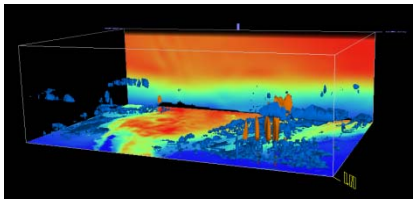

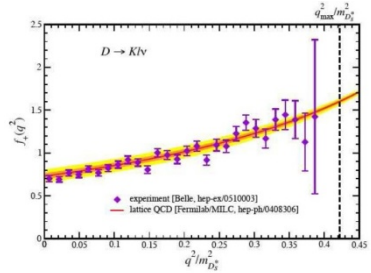

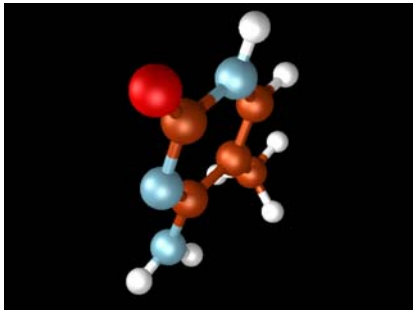

- Computational Fluid Dynamics (*leslie3d*, *fds*, *zeusmp*, *pop2*)
- Quantum Chromodynamics (*milc*)
- Weather Forecasting (*wrf*)
- Parallel Ray Tracing (*tachyon*)
- Molecular Dynamics (*lammps*)
- Heat transfer (*geofem - code Tokyo Univ*)
- Hydrodynamics (*tera_tf*)
- Matrix Decomposition (*lu*)
- Density Functional Theory (*socorro*)

SPEC MPI2007 Metrics

- SPECmpiM_base2007:
  The geometric mean of thirteen normalized ratios (base tuning).

- SPECmpiM_peak2007:
  The geometric mean of thirteen normalized ratios (peak tuning).

Virtualization

Virtualizing a number of servers on a single platform helps to reduce costs, save energy and ease IT infrastructure management.

Virtualization

- In order to help end-users in selecting an optimal system for server virtualization, a virtualization benchmark is being developed.

- Heterogeneous workloads that are spread across multiple virtual machines on a single server will be used. A key aspect will be defining a methodology to model the dynamic nature of customer workloads in this environment.

- Goals are:
  - Single primary metric
  - Scale across a wide range of systems
  - Component workloads representing common application categories typical of virtualized environments.

Single Tile Design

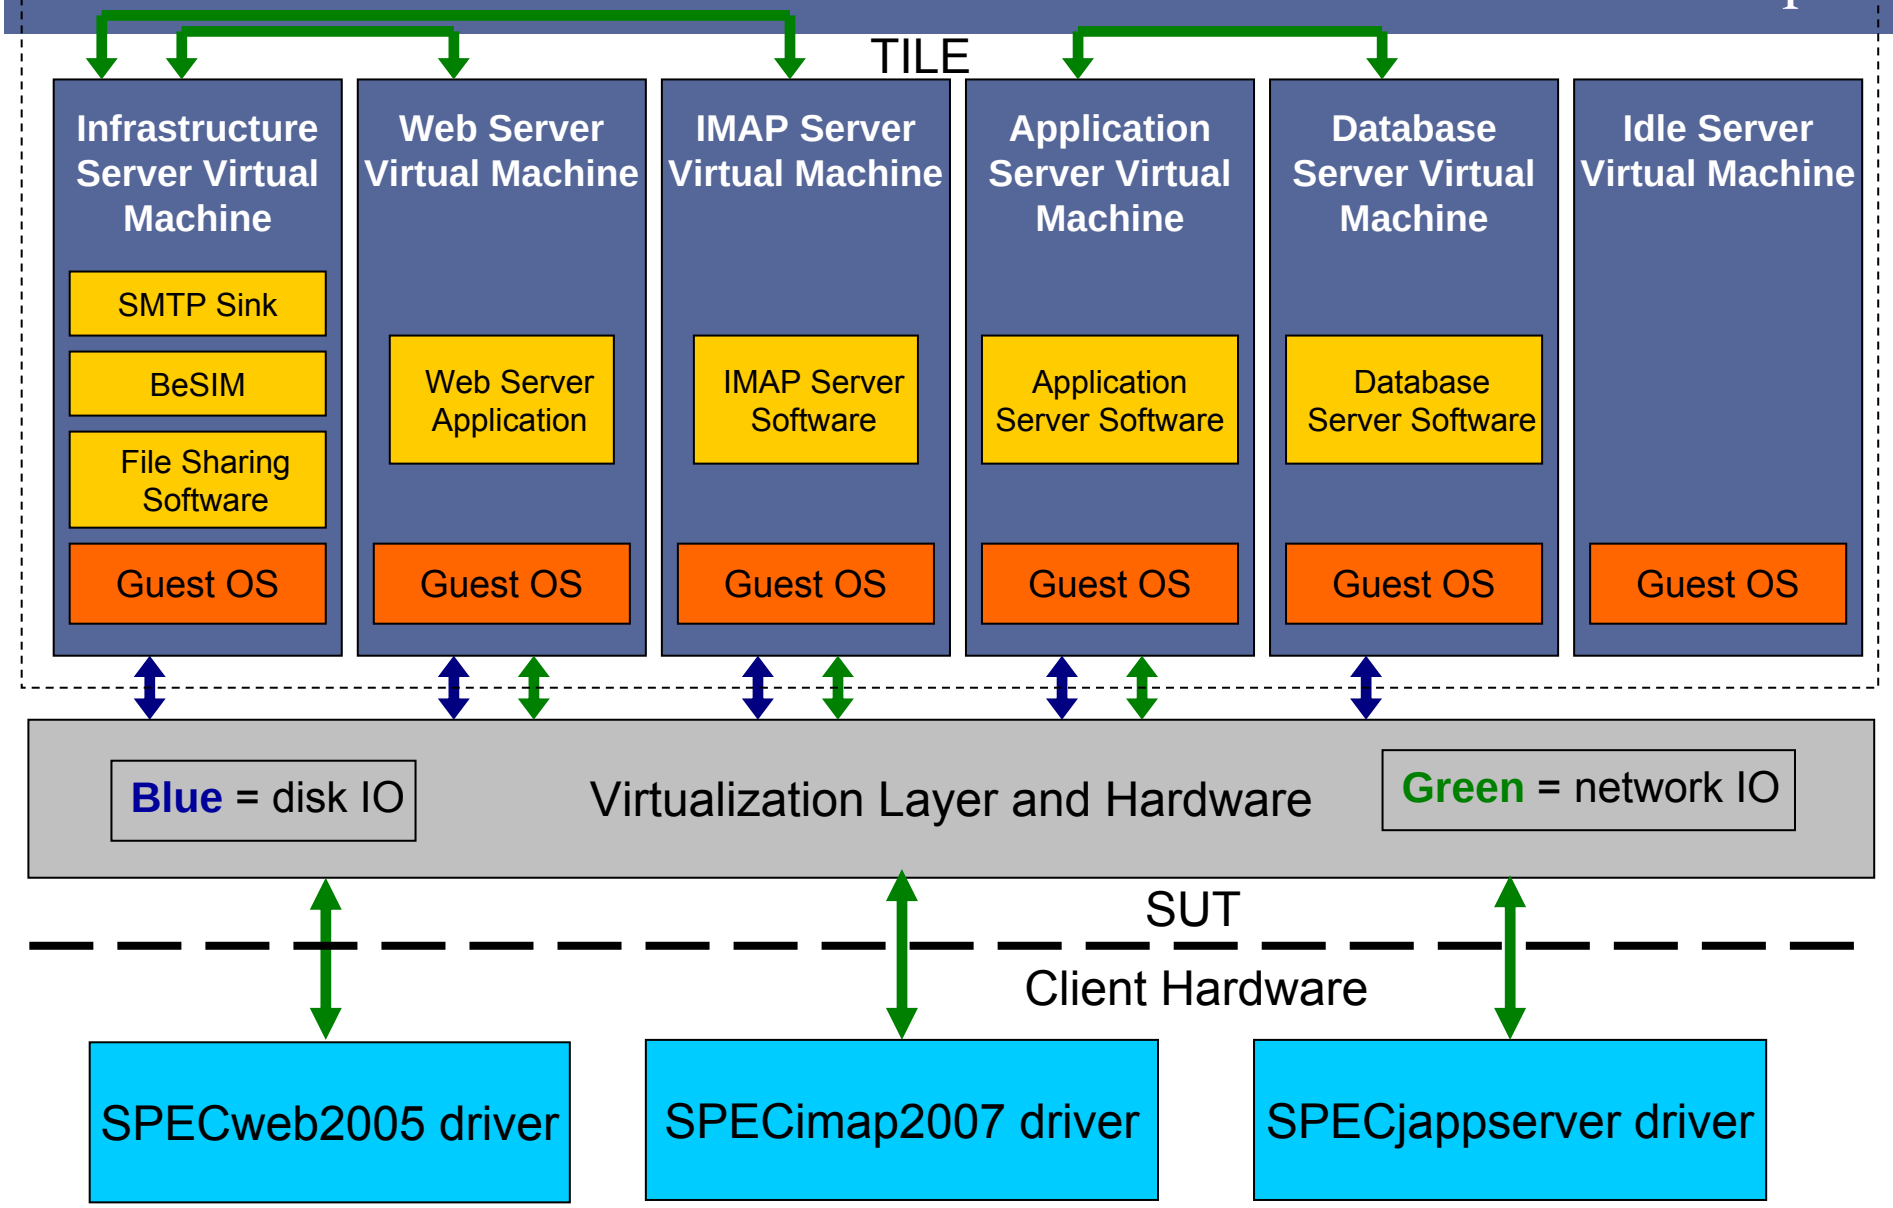

Conclusions

- SPEC is continually evolving and adapting to the marketplace in order to help end-users find the most suitable systems for their needs.
- If you are looking for systems to run compute-intensive applications on single CPU, SMP or cluster & have low power consumption, compare them on the SPEC web site.
- Future developments are for finding the most efficient virtualized servers, the most efficient systems for VoIP and instant messaging.

Thank you for your attention !