



Benchmarking

M.A. Oliveira

Outline

Benchmarking at LIP

Miguel Afonso Oliveira Mário David João Martins
Gonçalo Borges Jorge Gomes

Laboratório de Instrumentação e Física Experimental de Partículas

LIP

HEPiX Fall 2009





Outline of Part I: CPU Benchmarking

Benchmarking

M.A. Oliveira

Outline

Part I: CPU

Part II: Storage

1 CPU procurement benchmarks

- Introduction
- The proposals
- The hardware
- Benchmarks
- Second recall

2 HEPSpec

- Why
- Scaling!!!
- HEPSpec2006
- HEPSpec2000



Outline of Part II: Storage Benchmarking

Benchmarking

M.A. Oliveira

Outline

Part I: CPU

Part II: Storage

3 Storage procurement benchmarks

- Introduction
- The proposals
- The hardware
- Benchmarks

4 Work in progress

- The problem
- Future



Benchmarking

M.A. Oliveira

Procurement

HEPSpec

Parte I

CPU Benchmarking



Introduction

Benchmarking

M.A. Oliveira

Procurement

Introduction

Proposals

Hardware

Benchmarks

Second recall

HEPSpec

- Objective: Replace compute farm at Lisbon & Coimbra
- Started: January 2008
- Budget: € 135000
- Main condition: test system had to be made available
- Main evaluation criteria:
 - Total processing capacity - SPEC2006 + “In-house” codes
 - Power consumption
 - Remote management capabilities
 - Delivery times



The proposals

Benchmarking

M.A. Oliveira

Procurement

Introduction

Proposals

Hardware

Benchmarks

Second recall

HEPSpec

Provider	Equipment	CPU	Server	Cores
CPC IS/HP	HP Proliant DL160 G5	Xeon E5462@2.8GHz	48	384
Fujitsu-Siemens	Primergy RX200S4	Xeon X5450@3.0GHz	42	336
Dell	Dell PE 1950 III	Xeon E5450@3.0GHz	42	336
Convex	Dell PE 1950 III	Xeon E5450@3.0GHz	42	336
Compta	IBM x3550	Xeon X5450@3.0GHz	29	232
El Corte Inglés	IBM x3550	Xeon X5450@3.0GHz	30	240
IDW	Sun Fire X4150	Xeon X5460@3.16GHz	20	160



The hardware

Benchmarking

M.A. Oliveira

Procurement

Introduction

Proposals

Hardware

Benchmarks

Second recall

HEPSpec

Hardware received to benchmark:

	FS	HP	HP	Dell
Model	Primergy RX200X4	Proliant DL160 G5	Proliant DL160G5	PE 1950 III
CPU	E5440@2.83GHz	E5426@2.8GHz	E5472@3.0GHz	E5450@3.0GHz
Memory	24GB	24GB	24GB	16GB

There were however discrepancies relative to the offerings:

FS	Proposed	To test
Processor	Xeon X5450@3.00GHz	Xeon E5440@2.83GHz
Dell	Proposed	To test
Memory	24GB	16GB



Benchmarks

Benchmarking

M.A. Oliveira

Procurement

Introduction

Proposals

Hardware

Benchmarks

Second recall

HEPSpec

- Installation of operating system SL4 i686 with all updates.
- Building vanilla kernel 2.6.24.4 from kernel.org.
- SPEC2006.
- In-house suite of applications:
 - Corsika.
 - Bystander.
 - Drell-Yan.

Machine	SRNorm	CorNorm	BystNorm	DYNorm	Average	Difference
FS: RX200S4	91.4	94.4	94.0	99.7	94.9	5.1
Dell: PE 1950	94.5	99.7	99.9	99.8	98.5	1.5
HP: DL160 G5	97.5	93.1	94.0	93.4	94.5	5.5
HP: DL160 G5	100.0	100.0	100.0	100.0	100.0	0.0



Benchmark conclusions

Benchmarking

M.A. Oliveira

Procurement

Introduction

Proposals

Hardware

Benchmarks

Second recall

HEPSpec

- Price on all solutions were similar
- Delivery times were similar
- Remote management present on all solutions
- No software compatibility issues

∴ Similar offerings

How much compute power can we really get for our money?



Second recall criteria

Benchmarking

M.A. Oliveira

Procurement

Introduction

Proposals

Hardware

Benchmarks

Second recall

HEPSpec

- Total number of cores offered N_c
- Performance factor to be applied w :
 - $w = 1$ for the HP DL160 G5 @ 3.0 GHz
 - $w = 0.95$ for the HP DL160 G5 @ 2.8GHz
 - $w = 0.98$ for the Dell PE 1950 @ 3.0GHz
 - $w = 0.95$ for the FS RX200S4 @ 2.8GHz



Results and Conclusion

Benchmarking

M.A. Oliveira

Procurement

Introduction

Proposals

Hardware

Benchmarks

Second recall

HEPSpec

Machine	CPU	# Servers	# Cores	w	$w * N_c$
Dell: PE 1950	Xeon E5450@3.0GHz	56	448	0.98	439.04
FS: RX200S4	Xeon E5450@3.0GHz	52	416	0.95	395.20
HP: DL160 G5	Xeon E5472@3.0GHz	55	440	1.00	440.00

Close call between Dell and HP!!!

- Final reason:
 - FSB: HP (1.6GHz)/Dell(1.33GHz).
 - Power consumption is smaller on the HPs.

∴ Winning solution: HP DL160 G5



HEPSpec - Why

Benchmarking

M.A. Oliveira

Procurement

HEPSpec

Why

Scaling!!!

HEPSpec2006

HEPSpec2000

- The official CPU performance metric used by WLCG.
- Provides consistent and reproducible CPU benchmarks for the HEP community to describe and evaluate:
 - experiment requirements;
 - lab commitments;
 - existing computer resources;
 - procurement initiatives.



To scale or not to scale!

Benchmarking

M.A. Oliveira

Procurement

HEPSpec

Why

Scaling!!!

HEPSpec2006

HEPSpec2000

According to the instructions:

To calculate the score for each core of the system under test, compute the geometric mean of the seven C++ benchmarks,....

```
$ echo "scale=2 ; e((1(10.9)+1(13.8)+1(5.29)+1(11.4)+1(6.06)+1(7.13)+1(7.55))/7)" | \  
> bc -l  
8.33
```

But our sysadmins have degrees in physics and prefer:

```
$echo "e((1(10.9)+1(13.8)+1(5.29)+1(11.4)+1(6.06)+1(7.13)+1(7.55))/7)" | \  
> bc -l | awk '{printf("%.2lf\n",$1)}'  
8.40
```





HEPSpec2006 - SL5

Benchmarking

M.A. Oliveira

Procurement

HEPSpec

Why
Scaling!!!

HEPSpec2006

HEPSpec2000

SL5 x86_64 2.6.18-128.7.1.el5 gcc4.1.2 glibc-2.5-34

Machine	CPU	Clock (GHz)	L2+L3 (KB)	Cores	Memory (GB)
IBM LS22	AMD Opteron 2356	2.3	512+2048	8	24 (6x4GB DDR2)
IBM HS21	Intel Xeon E5420	2.5	12288	8	24 (6x4GB DDR)
HP DL160 G6	Intel Xeon E5540	2.5	1024+8192	8	24 (6x4GB DDR)
HP BL460C G6	Intel Xeon X5550	2.67	256+8192	8	24 (6x4GB DDR)
HP DL160 G5	Intel Xeon E5472	3.0	12288	8	24 (6x4GB DDR)

Machine	Server	CPU
	Score	Mean
IBM LS22	54.61	6.83
IBM HS21	65.60	8.20
HP DL160 G6	105.14	13.14
HP BL460C G6	113.22	14.15
HP DL160 G5	81.71	10.21



HEPSpec2006 - SL4

Benchmarking

M.A. Oliveira

Procurement

HEPSpec

Why
Scaling!!!

HEPSpec2006

HEPSpec2000

SL4 i686 2.6.9-89.0.9.EL.cernsmp gcc3.4.6 glibc-2.3.4-2.43

Machine	CPU	Clock (GHz)	L2+L3 (KB)	Cores	Memory (GB)
HP DL160 G5	Intel Xeon E5472	3.0	12288	8	24 (6x4GB DDR)
Sun X4100	AMD Opteron 275	2.1	1024	4	4 (4x1GB DDR)
Sun X4100	AMD Opteron 275	2.1	1024	4	4 (4x1GB DDR)
Sun X2200	AMD Opteron 2218	2.6	1024	4	8 (4x2GB DDR)
Dell PE1950	Intel Xeon E5450	3.0	6144	8	16 (8x2GB DDR)
Generic	Intel Xeon	2.8	512+4096	2	1/2(1x512MB DDR)

Machine	Server	CPU
	Score	Mean
HP DL160 G5	75.93	9.49
Sun X4100	28.03	7.01
Sun X4100	25.15	6.29
Sun X2200	29.63	7.41
Dell PE1950	74.19	9.27
Generic	7.61	3.81



HEPSpec2000 - SL5

Benchmarking

M.A. Oliveira

Procurement

HEPSpec

Why
Scaling!!!

HEPSpec2006

HEPSpec2000

SL5 x86_64 2.6.18-128.7.1.el5 gcc4.1.2 glibc-2.5-34

Machine	CPU	Clock (GHz)	L2+L3 (KB)	Cores	Memory (GB)
IBM LS22	AMD Opteron 2356	2.3	512+2048	8	24 (6x4GB DDR2)
IBM HS21	Intel Xeon E5420	2.5	12288	8	24 (6x4GB DDR)
HP DL160 G6	Intel Xeon E5540	2.5	1024+8192	8	24 (6x4GB DDR)
HP BL460C G6	Intel Xeon X5550	2.67	256+8192	8	24 (6x4GB DDR)
HP DL160 G5	Intel Xeon E5472	3.0	12288	8	24 (6x4GB DDR)

Machine	Server	CPU	2k/2k6
	Score	Mean	
IBM LS22	11388	1422	208
IBM HS21	16032	2004	244
HP DL160 G6	19547	2443	186
HP BL460C G6	21399	2675	189
HP DL160 G5	18527	2316	227



HEPSpec2000 - SL4

Benchmarking

M.A. Oliveira

Procurement

HEPSpec

Why
Scaling!!!

HEPSpec2006

HEPSpec2000

SL4 i686 2.6.9-89.0.9.EL.cernsmp gcc3.4.6 glibc-2.3.4-2.43

Machine	CPU	Clock (GHz)	L2+L3 (KB)	Cores	Memory (GB)
HP DL160 G5	Intel Xeon E5472	3.0	12288	8	24 (6x4GB DDR)
Sun X4100	AMD Opteron 275	2.1	1024	4	4 (4x1GB DDR)
Sun X4100	AMD Opteron 275	2.1	1024	4	4 (4x1GB DDR)
Sun X2200	AMD Opteron 2218	2.6	1024	4	8 (4x2GB DDR)
Dell PE1950	Intel Xeon E5450	3.0	6144	8	16 (8x2GB DDR)
Generic	Intel Xeon	2.8	512+4096	2	1/2(1x512MB DDR)

Machine	Server	CPU	2k/2k6
	Score	Mean	
HP DL160 G5	18796	2349	248
Sun X4100	5069	1267	181
Sun X4100	5125	1281	204
Sun X2200	5995	1499	202
Dell PE1950	18961	2370	256
Generic	1652	829	217



Benchmarking

M.A. Oliveira

Procurement

Work

Parte II

Storage Benchmarking



Introduction

Benchmarking

M.A. Oliveira

Procurement

Introduction

Proposals

Hardware

Benchmarks

Work

- Objective: Replace storage at Lisbon & Coimbra
- Started: January 2008
- Budget: € 170000
- Main condition: test system had to be made available
- Main evaluation criteria:
 - Read/Write performance
 - Power consumption
 - Density
 - Delivery times



The proposals

Benchmarking

M.A. Oliveira

Procurement

Introduction

Proposals

Hardware

Benchmarks

Work

Provider	Equipment	Capacity	Remarks
Dell/EMC	EMC AX4+Dell PE 1950	322 TB	Fiber Channel+SAS
Dell	Dell PE 1950+Dell MD1000	420 TB	SAS
Convex	Dell PE 1950 + Dell PV MD1000	345 TB	SAS
IDW	Sun X4500	288TB	SATA, no hardware RAID
IDW	Sun X4150+ STK 2540	180TB	Fiber Channel+SAS
UNISYS	UNISYS ES3220L+EMC AX4	142 TB	Fiber Channel+SAS
FS	RX300S4+SX40	405 TB	SAS
FS	RX300S4+SX80	315TB	Fiber Channel
IECI	IBM X3650+EXP3000	294 TB	SAS
COMPTA	IBM X3650+EXP3000	280 TB	SAS
CPC IS	HP ML370R05+MSA60	360 TB	SAS



The hardware

Benchmarking

M.A. Oliveira

Procurement

Introduction

Proposals

Hardware

Benchmarks

Work

Hardware received to benchmark:

Server	EMC/Dell	FS		CPC IS	Dell
Model	PE 1950	Primargy RX300S4		Proliant ML370	PE 1950
CPU	E5450@3.0GHz	E5420@2.5GHz		E5420@2.5GHz	E5420@2.5GHz
Cores	8	8		8	8
Rack Units	1	2		4	1
Memory	16 GB	12 GB		22 GB	24 GB
Power Suplies	2	1		2	2
Controller cache	2+2 GB	256 MB		512 MB	512 MB
Model	Clarion AX4-5	SX80FC	SX40 SAS	MSA60	PV MD1000
Drawers	4	1	1	4	2
Rack Units/Drawer	2	2	2	2	3
Disks	SATA 7.2K	SATA 7.2K	SATA 7.2K	SATA 7.2K	SATA 7.2K
Capacity/disk	750 GB	750 GB	750 GB	750 GB	1 TB
Disks/drawer	12	12	12	12	15
Controller	FC	FC	SAS	SAS	SAS

There were however discrepancies relative to the offerings:

FS	Proposed	To test
Memory	16GB	12GB
Power suplies	2	1
Drawers	3	1
CPC IS	Proposed	To test
Controllers	2	1



Benchmarks

Benchmarking

M.A. Oliveira

Procurement

Introduction

Proposals

Hardware

Benchmarks

Work

- Installation of operating system SL5 x86_64 with all updates.
- Kernel 2.6.18-53.1.13.el5.
- Installation of all drivers and required software for each solution.
- Tiobench (File size: 64GB, Block size: 64kB, 16 threads)
- “dd” (File size: 64GB, Block size: 64 kB, 10 instances, iostat to monitor)

Vendor	Tiobench		dd	
	Read Seq MB/s	Write Seq MB/s	Read Seq MB/s	Write Seq MB/s
EMC	60.0	158.8	47.3	98.3
FS	321.8	251.4	258.1	129.9
HP	288.1	214.3	234.3	119.9
Dell	453.0	344.0	398.0	202.8



Benchmark conclusions

Benchmarking

M.A. Oliveira

Procurement

Introduction

Proposals

Hardware

Benchmarks

Work

- Dell solution presents consistently better values.
- Power consumption was also measured and values of “W per TB” also favored Dell.
- Density also favored Dell because of the 15 disk drawers.
- All other parameters were measured to be identical.

∴ Winning solution: Dell PE 1950 + Dell PV MD1000



The problem

Benchmarking

M.A. Oliveira

Procurement

Work

The problem

Future

- Biggest concern at moment: How do we (re)configure/optimize our storage.
- We have several hardware/software options:
 - controllers (parameters, raid levels)
 - caches
 - disks
 - protocols
 - filesystems
 - etc...
- We have several (conflicting) solicitations:
 - software area (small files, write once, read many policy)
 - metadata and object storage servers for lustre
 - large file areas/hot file areas
 - no downtimes

We have tried educated guesses to test for all these situations but we are at the moment adrift in a sea of puzzling results ☹



Future

Benchmarking

M.A. Oliveira

Procurement

Work

The problem

Future

- Understand our results.
- Test testing tools!!!
- Configure our storage areas in the best possible way.
- Benchmark our Lustre installation consistently.