

NDPF Node Performance

From GridWiki

Contents

- 1 Node performance figures for the NDPF farm nodes
 - 1.1 Node type: pizza0 (node18)
 - 1.2 Node type: AMDNCF (gfrc)
 - 1.3 Node type: Halloween (hall)
 - 1.4 Node type: Bulldozer (bull)
 - 1.5 Node type: Luilak-1 (lui1) and Luilak-2 (lui2)
 - 1.6 Node type: Valentine
 - 1.6.1 Note on power consumption
 - 1.7 Node type: gloei (stbc 2008)
 - 1.8 Node type: Sint Maarten
 - 1.9 Node type: Carnaval
 - 1.10 Node type: Knal
 - 1.11 Node type: Marsepein
 - 1.12 Node type: Chocolade
 - 1.13 Node type: Pepernoot
 - 1.14 Node type: Kipsate
 - 1.15 Node type: Lotenfeest
 - 1.16 Node type: DELL PowerEdge 1435 SC
 - 1.17 History of actions in accounting and farm config

Node performance figures for the NDPF farm nodes

Note performance in the NDPF is internally expressed in GHzHr-equivalents, where a 1 GHzHrEquiv should be the integer performance of a single-core Pentium-3 processor with a clock frequency of 1GHz. To convert this into SpecInt-2000 numbers, a conversion factor of 410 has been applied: 1GHzHrEquiv corresponds to 410 SI2k, although this is a known understatement of the actual performance. Since this number has been hard-coded in all conversion script that publish accounting data to the outside world, it means that from now on the "official" SI2k ratings we get from vendors and our own tests must be converted to GHzHrEquiv values using this factor (410), and then be entered in the appropriate configuration file `/etc/pbsacddb.conf`.

To convert from SpecInt2006 Rate base to SpecInt2000 numbers, multiply by a factor of 185. This factor was determined by comparing machine types for which both SpecInt2006 and SpecInt200 numbers were available.

As of 2006-09-01, the PRD facility contains 546 kSI2k

As of 2006-12-18, the PRD facility contains 785 kSI2k.

As of June 2008, the PRD facility will contain 2649 kSI2k including Halloween, or 2579 kSI2k excluding Halloween.

Node type: pizza0 (node18)

These systems are dual-Pentium3 CPUs at 0.933 GHz, on an MSI motherboard. Equivalent systems rate at approximately 403-429 SI2k. The overhead incurred by running two simultaneous jobs has not been taken into account.

pbsaccdb factor: 416/410 = 1.015

0 (zero) systems of this type contribute 0 kSI2k.

Node type: AMDNCF (gfrc)

These systems are dual Athlon MP2000+ systems. Equivalent systems rate at approx. 690 SI2k (the overhead incurred by running two simultaneous jobs has not been taken into account). Note that our own performance tests using the D0 MC application actually showed a true speed doubling compared to pizza0, so a factor of 2 would not have been unreasonable.

pbsaccdb factor: 690/410 = 1.68

0 (zero) systems of this type contribute 0 kSI2k.

Node type: Halloween (hall)

These systems are dual Xeon 2.8 GHz systems with 1MB L2 cache. Hyperthreading is not used. Equivalent systems rate at 1288 SI2k.

pbsaccdb factor: 1288/410 = 3.14

27 systems with 54 cores contribute 70 kSI2k.

Node type: Bulldozer (bull)

These systems are dual Xeon 3.2 GHz with 2 MB L2 cache. Hyperthreading is not used. Equivalent systems rate at 1555 (Dell PowerEdge 1850).

pbsaccdb factor: 1555/410 = 3.79

34 systems with 68 cores contribute 106 kSI2k.

HEP-SPEC06 information: result is 11.73 per box, meaning 5.87 per core (measured with CentOS 4.8 i386)

HEP-SPEC06 information: results is 14.94 per box, meaning 7.47 per core (measured with CentOS 5.4 x86-64)

Node type: Luilak-1 (lui1) and Luilak-2 (lui2)

These systems are Dell PowerEdge 1950 (Intel Xeon processor 5150, 2.66GHz) rated at 2764 when used with one process (specification by Dell, July 2006). With four simultaneous jobs, this degrades to 2240.

pbsaccdb factor: 2240/410 = 5.46

2*34 systems with 272 cores contribute 609 kSI2k.

HEP-SPEC06 results : 36.56 per box, meaning 9.14 per core (measured with CentOS 4.8 i386).

HEP-SPEC06 results : 39.89 per box, meaning 9.97 per core (measured with CentOS 5.4 x86-64).

Node type: Valentine

These systems are Supermicro X7DBE (Intel Woodcrest 5100, 2.5 GHz). The 102 nodes together are rated at 10077 SpecInt-2006 Rate base. This corresponds to 12.35 SpecInt2006 = 2285 SI2k = 5.57 GHz.hr per core.

pbsaccdb factor: 2285/410 = 5.57

102 systems with 816 cores contribute 10077 SpecInt-2006 = 1864 kSI2k.

HEP-SPEC06 results on EL4/i386: 65.82 per box, meaning 8.23 per core (measured with CentOS 4.8 i386).

HEP-SPEC06 results on EL5/x86_64: 70.48 per box, meaning 8.81 per core (measured with CentOS 5.4 x86-64).

HEP-SPEC06 results on EL5/x86_64: 71.09 per box, meaning 8.89 per core (measured with CentOS 5.9 x86-64).

HEP-SPEC06 results on EL5/x86_64(64bit): 72.97 per box, meaning 9.12 per core (measured with CentOS 5.9 x86-64).

HEP-SPEC06 results on EL6/x86_64: 73.80 per box, meaning 9.23 per core (measured with CentOS 6.4 x86-64).

Note on power consumption

For node wn-lui2-023 the following currents were measured under various conditions (OS: CentOS3.8/i386)

standby (off)	0.17A	39W
startup peak	>1.50A	345W
setup screen or unloaded system	0.9 - 1.0A	230W
with 4x burnP6 and 4x "burnMMX P"	1.35A	310W
with continuous disk activity	1.1A	250W

The line voltage is approx. 229V (see mail "stroomgebruik grid resources" by wimh of 27-Sep-06 16:13).

Node type: gloei (stbc 2008)

Note this is partially re-scaled from the SintMaartens, but the L5420 CPU is 2.50GHz and slightly older than the :5520 at 2.27, but faster per core. From the tender: 20043 SI06RB for 206 systems, so 97.2 SI06RB/system, and 12.15 SI06RB/core. So 2248 SI2K-equiv, and thus 5.48 GHz-equiv. This corresponds to

pbsacddb factor: 5.48

Now the HS06 calculation is tricky as it has to be inferred, but let's assume it scales from luilak, the architecture is similar anyway, as well as the score. So this is then $(5.48/5.46)*9.97 = 10.01$

Node type: Sint Maarten

These systems are HP BL460c G6 CTO Blade 160 systems (Intel Core-i7 L5520 processors). The 176 blades together are rated at 32560 SpecInt-2006 Rate base (using 16 processes, with 2 threads/core, though!). The use 8 processes without HT compared to 16 threads and HT results in an overall reduction of throughput of ~ 85%. This corresponds to 185 SI06Rate/node and thus 23.13 SpecInt2006rate per core = 4278 SI2k-effective = 10.43 GHz.hr per core with 16 job slots per system, or 19.6 SI06rate/core = 3637 SI2k-effective = 8.87 GHz.hr-equiv per core.

pbsacddb factor: 4278/410 = 10.43 (16 job slots)

pbsacddb factor: 3637/410 = 8.87 (8 job slots, no HT)

176 systems with 1408 cores contribute 32560 SpecInt-2006 = 6023 kSI2k.

HEP-SPEC06 results on EL5/x86_64: 96.26 per box, meaning 12.03 per core (measured with CentOS 5.4 x86-64).

HEP-SPEC06 results on EL5/x86_64: 97.01 per box, meaning 12.13 per core (measured with CentOS 5.9 x86-64).

HEP-SPEC06 results on EL5/x86_64 (64-bit): 112.36 per box, meaning 14.05 per core (measured with CentOS 5.9 x86-64).

HEP-SPEC06 results on EL6/x86_64: 99.35 per box, meaning 12.42 per core (measured with CentOS 6.4 x86-64).

Running without HT gives overall better throughput for single jobs and reduced memory contention for the 24 GByte in each node.

Node type: Carnaval

These systems are Dell Blade systems (Intel processors). The 96 blades together are rated at 27744 SpecInt-2006 Rate base (using 12 processes, with 1 thread/core). This corresponds to 289 SI06Rate/node and thus 24.08 SpecInt2006rate per core = 4455 SI2k-effective = 10.86 GHz.hr per core with 12 job slots per system.

pbsacddb factor: $4455/410 = 10.86$ (12 job slots, no HT)

96 systems with 1152 cores contribute 27744 SpecInt-2006 = 5133 kSI2k.

HEP-SPEC06 results on EL5/x86_64: 140.14 per box, meaning 11.68 per core (measured with CentOS 5.6 x86-64).

HEP-SPEC06 results on EL5/x86_64: 142.88 per box, meaning 11.91 per core (measured with CentOS 5.9 x86-64).

HEP-SPEC06 results on EL5/x86_64 (64-bit): 163.86 per box, meaning 13.66 per core (measured with CentOS 5.9 x86-64).

HEP-SPEC06 results on EL6/x86_64: 149.91 per box, meaning 12.49 per core (measured with CentOS 6.4 x86-64).

Node type: Knal

These systems are Dell R820 systems with 4 CPUs (Intel processors). The 18 blades together are rated at 16812 SpecInt-2006 Rate base (using 32 processes, with 1 thread/core). This corresponds to 934 SI06Rate/node and thus 29.1875 SpecInt2006rate per core = 5400 SI2k-effective = 13.17 GHz.hr per core with 32 job slots per system.

pbsacddb factor: $5400/410 = 13.17$ (32 job slots, no HT)

18 systems with 576 cores contribute 16812 SpecInt-2006 = 3110 kSI2k.

Results from the default 32 bit test:

HEP-SPEC06 results on EL5/x86_64: 403.70 per box, meaning 12.62 per core (measured with CentOS 5.9 x86-64).

HEP-SPEC06 results on EL6/x86_64: 430.31 per box, meaning 13.44 per core (measured with CentOS 6.4 x86-64).

Results from the 32 bit test while using taskset to force each process on a specific core:

HEP-SPEC06 results on EL5/x86_64: 403.53 per box, meaning 12.61 per core (measured with CentOS 5.9 x86-64).

HEP-SPEC06 results on EL6/x86_64: 422.46 per box, meaning 13.20 per core (measured with CentOS 6.4 x86-64).

Results from the default 64 bit test:

HEP-SPEC06 results on EL5/x86_64 (64-bit): 471.74 per box, meaning 14.74 per core (measured with CentOS 5.9 x86-64).

HEP-SPEC06 results on EL6/x86_64 (64-bit): 498.93 per box, meaning 15.59 per core (measured with CentOS 6.4 x86-64).

Node type: Marsepein

58 DELL PowerEdge R630, two 12-core CPUs Intel(R) Xeon(R) CPU E5-2680 v3 @ 2.50GHz stepping 2. SpecInt-2006 Rate (base) per node is 1020, which means 42.5 SI06Rate per core. Multiply by 185 makes 7862.5 SI2k-effective and

pbsacddb factor: $7862.5/410 = 19.17$ (24 job slots, no HT)

Result from the default 32 bit test:

HEP-SPEC06 results on EL6/x86_64: 427.10 per box, meaning 17.80 per core (measured with CentOS 6.6 x86-64).

Node type: Chocolate

58 Fujitsu RX2530 m2, two 12-core CPUs Intel(R) Xeon(R) CPU E5-2650 v4 @ 2.20GHz stepping 1. SpecInt-2006 Rate (base) per node is 1050, which means 43.75 SI06Rate per core. Multiply by 185 makes 8093.75 SI2k-effective and

pbsacddb factor: $8093.75/410 = 19.74$ (24 job slots, no HT)

Result from the default 32 bit test:

HEP-SPEC06 results on EL6/x86_64: 408.92 per box, meaning 17.04 per core (measured with CentOS 6.7 x86-64).

Node type: Pepernoot

13 DELL R640, 2 × Intel(R) Xeon(R) Gold 6148 CPU @ 2.40GHz with 20 cores each. SpecInt-2006 Rate (base) is 1950, which is 48.75 SI06Rate/core. Multiply by 185 makes 9018.75 SI2k-effective and

pbsacddb factor: $9018.75/410 = 22.00$.

HEPSPEC06 benchmark: 782.75 per node (non-hyperthreading) which is 19.57 per core.

Node type: Kipsate

81 (54 grid, 27 stoomboot) DELL R6415, single socket AMD EPYC 7551P 32-Core Processor.

HEP-SPEC06 results on EL6: 478.08 per machine, i.e. 14.94 per core.

Node type: Lotenfeest

The lot worker nodes are Lenovo SR655 nodes with 1 × AMD EPYC Rome 7702P CPU @

2.00GHz.

Servers 1-7 are AMD GPU nodes for stoomboot. Servers 8 and 9 are Nvidia GPUS also for stoomboot. Servers 10-36 are dedicated worker nodes for NIKHEF-ELPROD, while 37-64 are used for the Virgo cluster.

These nodes have a 3.2 TB NVMe cards.

Lenovo Global Technology benchmarked this system (ThinkSystem SR655) with CPU2017 at 325 (<https://www.spec.org/cpu2017/results/res2019q3/cpu2017-20190902-17395.html>). Note this is not the same or convertible metric to SPEC 2006. A SPEC 2006 benchmark has not been published for this system.

The systems are benchmarked with HEPSPEC06 with and without hyperthreading on CentOS 7.

Benchmark results (HEPSPEC06, 32-bit mode):

hyperthreading	hepspec06	per core
yes (128 threads)	1412.36	11.03
no (64 cores)	1226.60	19.17

Since SpecINT06Rate(base) is not available for the 7702P, and it's an old benchmark anyway, we scale using SI17RB

```

-----
135 voor de 7551P sate at 32C 2GHz HT 64p => 2.1093 SI17BR/thread
317 voor de 7702P lot at 64C 2GHz HT 128p => 2.4765 SI17RB/thread
ergo: SI06RB voor de 7702P is 1.17412 * 37.5 = 44.03 SI06RB/core-enabled
-----

```

given that the reference for a 2.0GHz AMD 7702P (HPE ProLiant DL325 Gen10 Plus (2.00 GHz, AMD EPYC 7702P)), with 2 thread/core on a 64C system gives 317 RB The GHzEquiv rating is thus 18.05

Node type: DELL PowerEdge 1435 SC

This is a new series of nodes fitted with dual/dual AMD Opteron 2220 processors. The following tests were done on CentOS 4.4.

```

standby (off)                0.2A  50W
unloaded system              0.9A  210W
with 4x burnP6 and 4x "burnMMX P" 1.25A 286W

```

History of actions in accounting and farm config

5 february 2014 : standardize on 64bit centos6 node performance values. values entered in accounting machinery, acct. backfilled til 1 feb 2014. fair share scheduling values and site capacity also updated with new values.

Retrieved from "https://wiki.nikhef.nl/gridwiki/index.php?title=NDPF_Node_Performance&oldid=10898"

-
- This page was last modified on 24 August 2020, at 16:29.