

numactl

numactl - Control NUMA policy for processes or shared memory.

Non-uniform memory access (NUMA) is a computer memory design used in multiprocessing, where the memory access time depends on the memory location relative to the processor. Under NUMA, a processor can access its own local memory faster than non-local memory (memory local to another processor or memory shared between processors). The benefits of NUMA are limited to particular workloads, notably on servers where the data are often associated strongly with certain tasks or users

From the command line, use the following:

```
numactl --hardware
```

This will give an output similar to the following:

```
available: 2 nodes (0-1)
node 0 cpus: 0 1 2 3 4 5 6 7 8 9 10 11 24 25 26 27 28 29 30 31 32 33 34 35
node 0 size: 32671 MB
node 0 free: 29022 MB
node 1 cpus: 12 13 14 15 16 17 18 19 20 21 22 23 36 37 38 39 40 41 42 43 44 45 46 47
node 1 size: 32768 MB
node 1 free: 30264 MB
node distances:
node  0  1
  0: 10  21
  1: 21  10
```

What we can see here is that there are two (NUMA) nodes. (0-1) Each node consists of some CPUs and some Memory. In this example, each node consists of 24 CPUs and 32GB of Memory, each CPU however is a hyper-thread, so although we see 24 CPUs per node, it is actually 12 Hyper-threading CPU's (this is what you will see on a G6 1U).

From:
<http://cameraangle.co.uk/> - WalkerWiki - wiki.alanwalker.uk

Permanent link:
<http://cameraangle.co.uk/doku.php?id=wiki:numactl>

Last update: **2023/03/09 22:35**

