

Workaround for Bug 41288 [Performances]

NUMA Core Operation

This Advanced Parameter controls the NUMA configuration on the MFEL. For this to work the NUMA settings in the BIOS must have been correctly configured.

The **default** value for the following is **false**, we changed it to **true**. This was for 2 HEVC services on a single 1U G6 (1052)

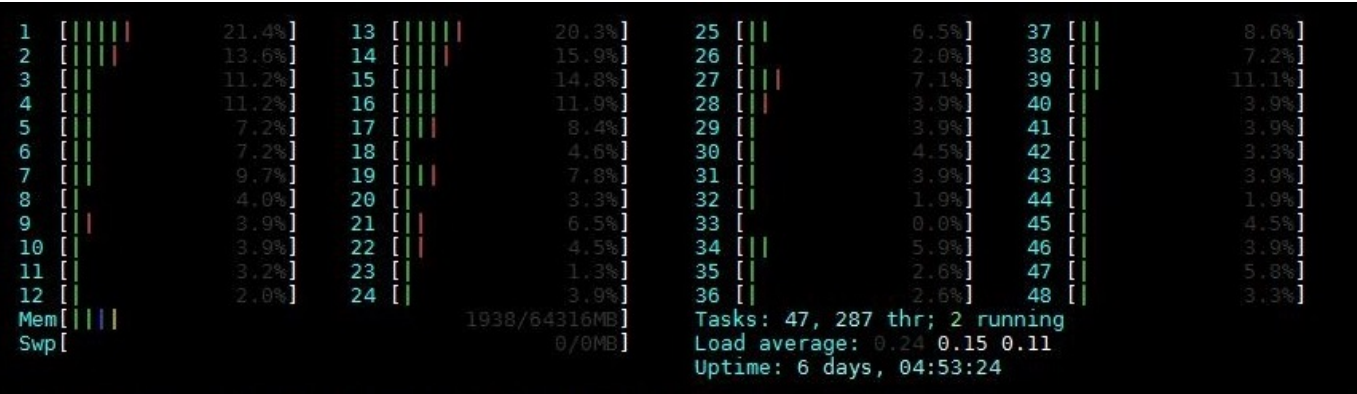
```
masterSlave.setPreferredNumaNode=true
```

If the NUMA Node is set to **true**, then this will limit a service to using only a single NUMA node (one CPU and Memory).

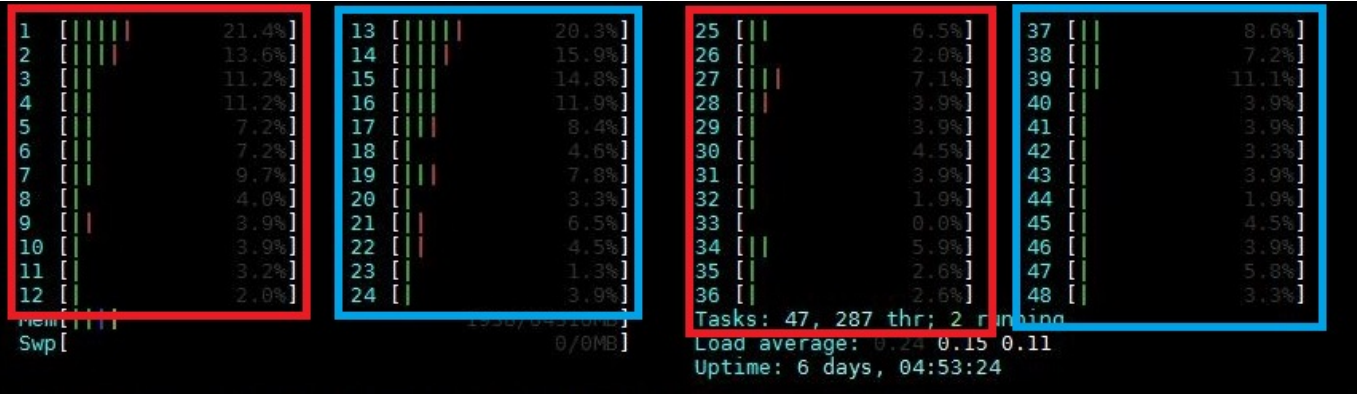
If the NUMA Node is set to **false**, then this will allow a service to use both NUMA Nodes (so both CPUs, and all memory). This is important if a service has a lot of components (say an ABR service with a lot of profiles) that could be too much for a single processor to handle.

When setting NUMA to false, you might get a performance improvement if you set the NUMA to of in the BIOS

While running htop, we can see the processors listed at the top of the page.

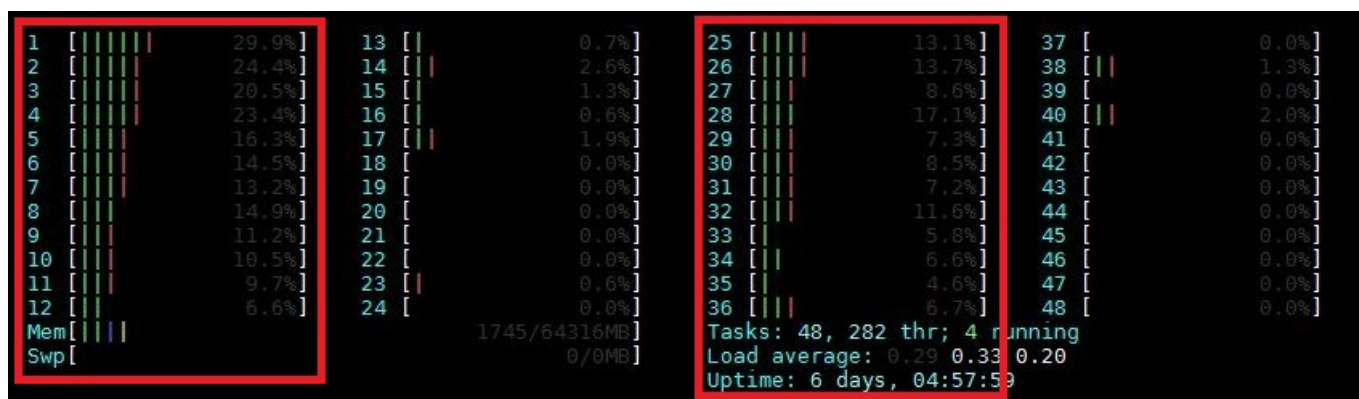


In the above example we can see 48 cores from the two CPUs, these are organised in the following way:



The two red boxes are for the first NUMA node, the blue boxes are for the second NUMA node. In the above example we can see that although the MFEL is only running a single service, some cores from each NUMA node are being used.


This is because we have allowed this in the Advanced Parameters (masterSlave.setPreferredNumaNode=false) If we change the masterSlave.setPreferredNumaNode setting to false then we get the following output from htop.



Here we can see that with a single service, all of the work is being done by one NUMA node, the activity on the other NUMA node is just basic OS tasks.

Processor Occupancy

We know that there are many cores (CPUs) per socket, and these cores are allocated by the operating system. However the allocation of processor cores isn't actually that random.



The diagram above shows one of the NUMA cores, and it can be seen that the cores are used in order, and that core 1 is always the most used core, while the lower cores may get only the occasional use. This of course doesn't mean much, but it is interesting to know.

When setting NUMA to false, you might get a performance improvement if you set the NUMA to off in the BIOS

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

Permanent link: http://cameraangle.co.uk/doku.php?id=workaround_for_bug_41288_-_performances_controlling_numa_core_operation

Last update: 2023/03/09 22:35

