

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.

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.

From:

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

Permanent link:

http://cameraangle.co.uk/doku.php?id=wiki:workaround_for_bug_41288_-_performances_controlling_numa_core_operation&rev=1480588015

Last update:

2023/03/09 22:35