




Science and
Technology
Facilities Council

WN performance

Alastair Dewhurst




InfluxDB

- Many of the plots in this talk have been taken from Generic System Overview view in Vande:
 - lcg2413.gridpp.rl.ac.uk
 - lcg2456.gridpp.rl.ac.uk
- For WN this view is quite temperamental!
 - The “hostname” box at the top won’t display the WNs, you have to enter that manually.
 - It can take a little while to load (especially if looking back 90+ days).
 - Sometimes it wouldn’t load and I just had to click refresh a few times.


James Walder's test

- From [James' Talk last week](#) clear performance degradation when downloading files on machines with MULTICORE personality. Is the cause:
 - Different configuration? (possible)
 - Different mix of jobs? (probably)




WN Disk I/O

lcg2413.gridpp.rl.ac.uk - ML personality




lcg2456.gridpp.rl.ac.uk - MULTICORE personality



lcg2456.gridpp.rl.ac.uk

- It would appear that we are frequently hitting a limit on the number of writes that can be performed to the SSD (~1250 OP/s).
- The monitoring does not go back much farther than shown in the plots, but it would seem that we hadn't hit the write limit before 28th April.
 - Firewall change to allow CMS jobs to download more data was 20th April.


Firewall change




WN Network

lcg2413.gridpp.rl.ac.uk - ML personality

Network drops start from 7th June?



lcg2456.gridpp.rl.ac.uk - MULTICORE personality



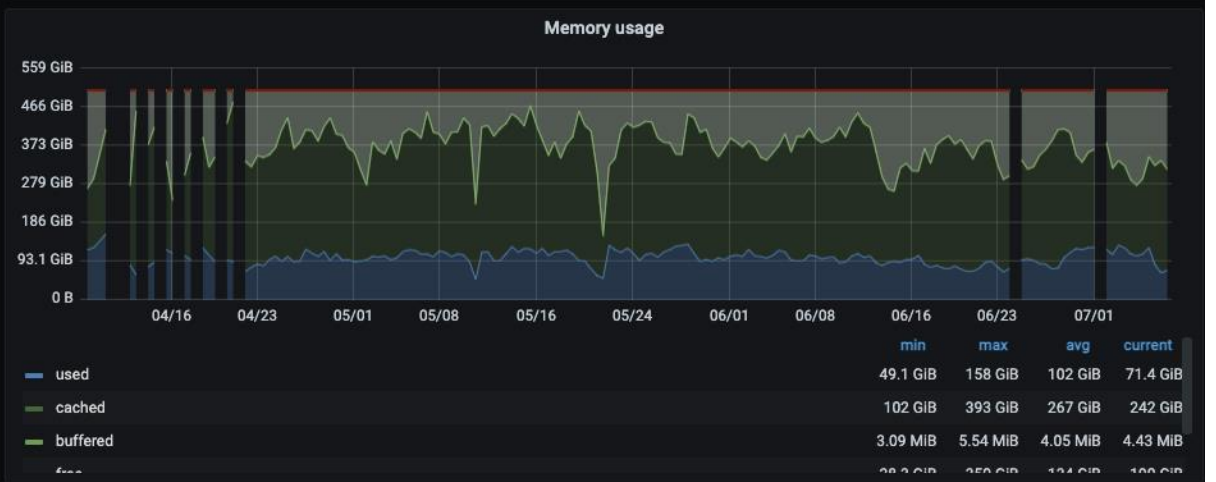
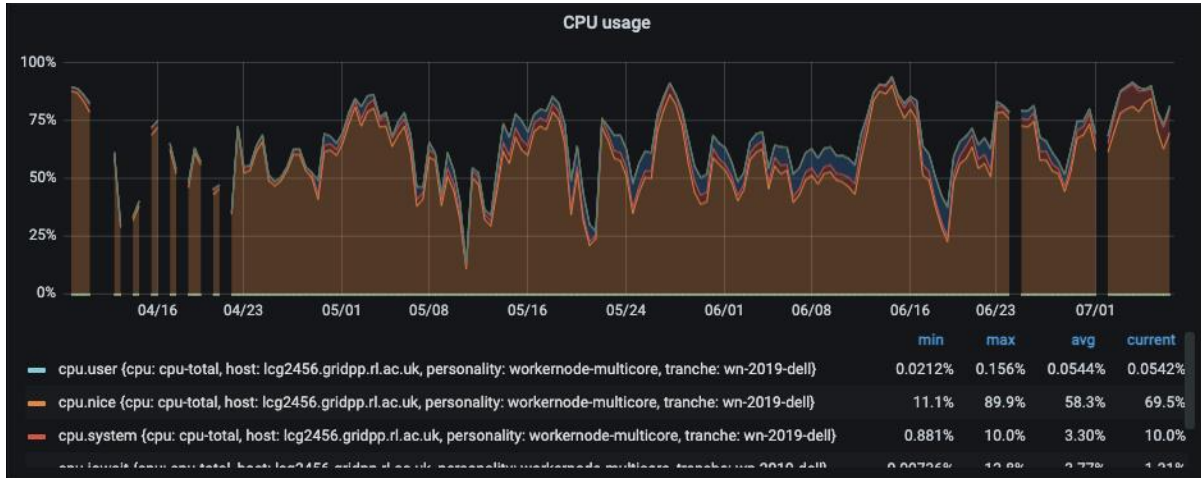
WN CPU / Memory

lcg2413.gridpp.rl.ac.uk - ML personality



WN CPU / Memory

lcg2456.gridpp.rl.ac.uk - MULTICORE personality



Procurement

- Due to Silicon shortage need to be placing orders soon.
- What are the requirements?
 - CPU - depends on what is available. Hopefully another AMD with similar spec to before (EPY 7452 / 7453).
 - Network current 25Gb/s NIC seems more than sufficient.
 - Memory - 512GB (8GB / physical core seems plenty)
 - Disk?
- What exactly is the limiting factor currently?
 - Can we estimate how much better it needs to be?

Ideas

- Run a high I/O CMS jobs on an empty node and measure IOPs?
- Reduce the number of jobs that can run on a few MULTICORE Dell19 WN to see if the limit stops being reached?
- Can we look into scheduling jobs differently to spread out CMS/Multicore jobs more?
 - E.g. allow X% of job slots to be exclusively for multicore
- The XMA20 CPU should be coming online soon.
 - How should they be configured for production?



Science and
Technology
Facilities Council

Questions?