

VMEM in 64bit Architecture

- We observe that any process running in 64bit has a VMEM footprint of ~50MB (c.f. 5MB for a 32bit process)

```
USER      PID %CPU %MEM    VSZ   RSS TTY      STAT START   TIME COMMAND
gla012   1741  0.1  0.0 58892 2612 pts/13    S    12:49   0:00 python yawn.py
gla012   1745  0.0  0.0  5816 2176 pts/13    S    12:49   0:00 python32 yawn.py
```

- This means that running 'hello world' on the grid on a 64bit machine has a VMEM footprint of ~450MB (c.f. 32bit footprint of ~60MB)
 - In both cases the RSS is ~22MB
- Which means that killing grid jobs based on VMEM consumption is probably not a good idea...
 - N.B. also killing jobs based on RSS doesn't work as the kernel actively tries to keep pages in memory
 - And it seems that, e.g., torque kills based on the memory consumption of the process tree, not of the payload