




Limiting memory consumption





- Some sites would like to limit the memory consumption in order to avoid jobs decreasing performance of WNs
 - What is really important is whether some processes require heavy swapping
- RSS is not a good metrics
 - For the same VMEM, RSS is larger on an empty WN!
 - Can we agree on that?
- Is VMEM a good metrics?
 - Isn't it rather swapping rate?
 - If pages are inactive, who cares?
 - Is it possible to limit swapping?
- Is swapping the devil?
 - Some sites have no swapping space...
 - ↳ Why was virtual memory invented?!
 - ↳ Incredible waste of memory as all wrappers are in RSS!



- What is the “offending” entity?
 - With one slot per core, it is a **process**, not a job!
 - Killing a job is just **useless!**
 - ↳ No information to user, therefore waste of resources
 - ↳ Confusion for the CE / WMS: the job has gone!

- Which metrics?
 - **VMEM**
 - ↳ Deterministic for the main process
 - **RSS**
 - ↳ Depends on the load of the machine. If used, should anyway be set to max VMEM (in case the machine is empty)
 - **PSS**
 - ↳ Accounts shared pages weighted with $1/(\text{nb of processes sharing})$
 - ↳ In principle the best estimate, but depends also on what else is running
 - ⚖ Limit should be max VMEM in case there is no sharing



What is done on LHCb Tier1s?

Memory limitation

- What is currently done, depending on the site
 - No limit at all (3 sites)
 - ↳ Works fine (never seen any related bad job performance)
 - Limit VMEM per process (ulimit) (2 sites, 3.8 GB)
 - ↳ Then kill the offending process (sending a signal)
 - ↳ Deterministic, but at least the framework can catch the return code and establish a diagnosis
 - Limit total VMEM/RSS per process group (job) (1 site, 5 GB)
 - ↳ This is unpredictable by the user!
 - ↳ No control on how much VMEM is used by wrappers etc...
 - ↳ Up to 1.5 GB!!!
 - ↳ In addition no swap space at that site!
 - Limit RSS per process (1 site, 4 GB)
 - ↳ Limit should be set to max VMEM anyway!
- For 7 sites, the sampling is not too bad

