

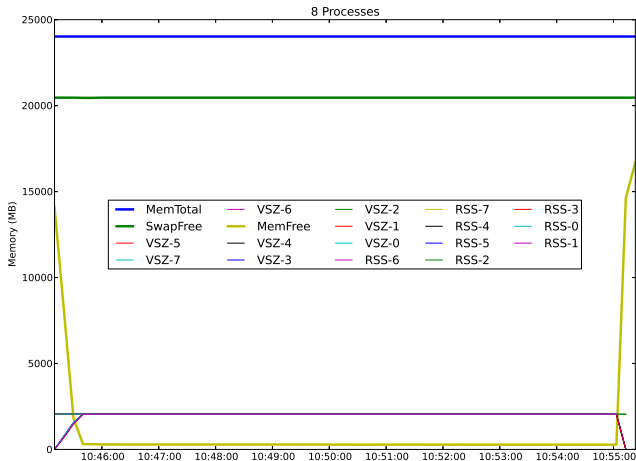
ALICE Job Efficiencies


Jérôme Belleman
12 December 2014

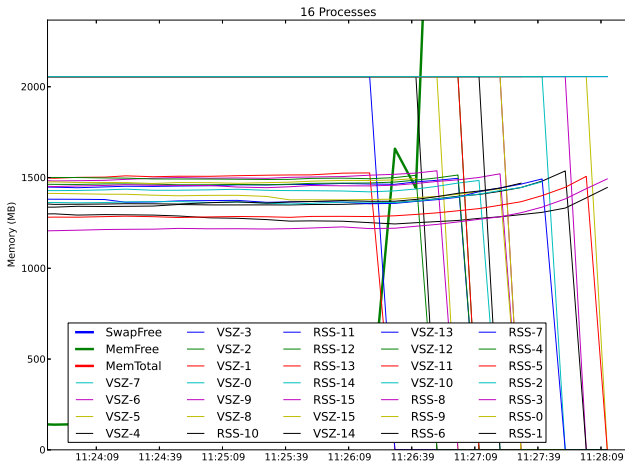
- ALICE expressed concerns about job efficiencies.
- Costin Grigoraş ran benchmarks allocating 2 GB.
- Pathologically long runs on memory-busy nodes.
- Swapping to be the cause.

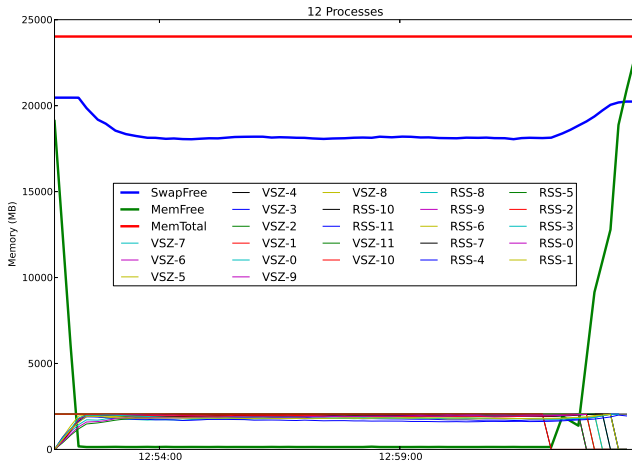
Costin provided his benchmark tool:

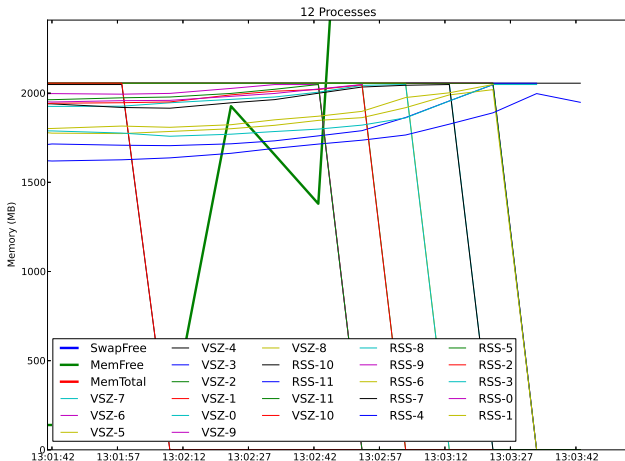
- Fired many jobs to catch cases.
- Interactively ran several instances on select nodes.
- 24 GB, 8 cores with SMT enabled.
- Run 8 to 16 benchmarks in parallel, 2-GB allocations.
- On physical nodes and on VMs.



















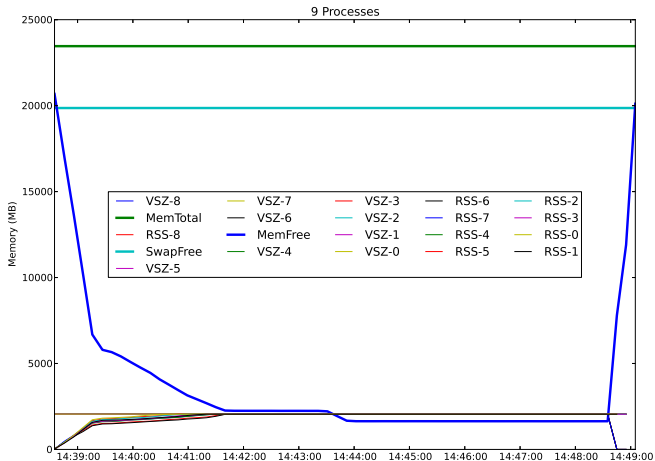












- Clean up core processes? Most of them swap anyway.
- < 11 slots? But not all jobs allocate 2 GB.
- New slotting: 100 sample jobs into the wild ran fine.
- Trails to enforce memory limits: cgroups?