

CE Stress Testing

- We monitored OSG specific processes in the condor_g client and the GRAM based OSG Compute Element while we sent “*sleep jobs*” from one or more clients.
- We check the scalability in terms of:
 - the number of jobs sent;
 - the size of the output file to be transferred to the user.

More information on the tool developed and used can be found at:

<http://hepuser.ucsd.edu/twiki2/bin/view/UCSDTier2/OsamonitorinarmPackageDescription>

OSG specific process monitoring

- Monitored quantities:(using ps and lsof)
 - % CPU
 - % Physical Memory
 - Resident Memory (kb)
 - Total Memory (kb)
 - Number of processes
 - Number of open files

Monitored processes: (for a "condor site")	Relevant in:	
	CE	client
condor_shadow	✓	
condor_starter	✓	
globus-job-manager	✓	
globus-job-manager-script	✓	
globus-job-manager-script-real	✓	
grid_manager_monitor_agent	✓	
grid-monitor	✓	
globus-url-copy	✓	✓
condor_master	✓	✓
condor_procd	✓	✓
condor_schedd	✓	✓
condor_gridmanager		✓
gahp_server		✓


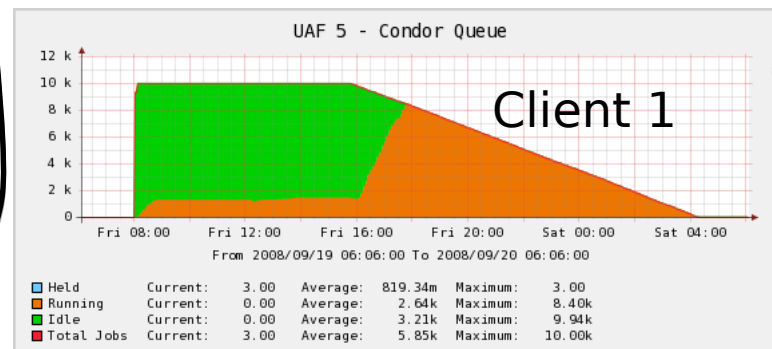

The “*sleep job*”

- It is a basic *bash* job that:
 - Creates a file of a specified length (that will be transferred back);
 - Sleeps till a certain date/time or for a number of seconds.
 - This allows for a well defined number of jobs to end simultaneously or sequentially.

Results: sending jobs

- There was no problem to start up $\lesssim 30000$ jobs in one CE
 - either coming from one submitter
 - or several submitters sending concurrently.
- The start up rate in the CE from the submitters
 - was independent of the number of submitters (1 to 3 tested),
 - stable and independent of the number of jobs at ~ 0.94 Hz.

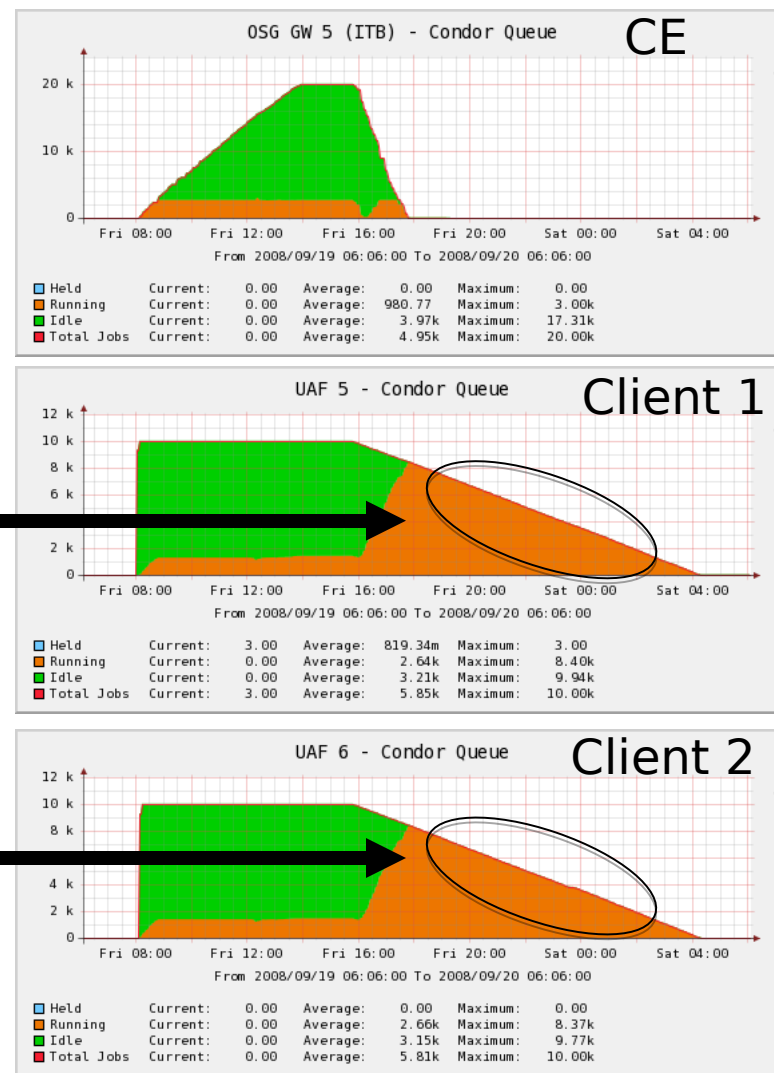
$\Rightarrow 30000$ jobs ~ 9 hours



30000 jobs send from 2 clients

Results: receiving output back

- The de-queue rate of finished jobs from the CE to the submitters was
 - independent of the number of submitters (1 to 3 tested)
 - and much slower than the start-up rate at ~ 0.44 Hz.
- \Rightarrow 30000 jobs \sim 19 hours




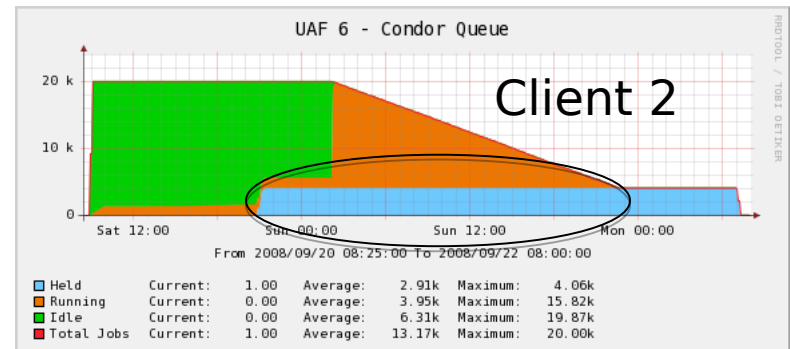

30000 jobs send from 2 clients

Results: returning output files

- Transferring back output files of up to 1 GBytes length poses no problem.
 - It may be quite slow if many such jobs occur concurrently

Results: Example limitation of the infrastructure

- One user can not queue ≥ 30000 jobs.
 - Jobs go to hold state.
 - This limit also appears in the submitter queue if ≥ 30000 jobs are send from a single submitter. No single job makes it into the CE then.
 - Other users are not affected while the user has the problem. They can operate normally.
 - `grid_manager_monitor_agent` seems to be dead in the CE



40000 jobs send from 2 clients