

Updates for GPU support in Batch Jobs on the OSG

Shreyas Bhat and Ken Herner

User Support for Distributed Computing (Fermilab)

July 14, 2017

From the FIFE workshop:

- On ITB, can route jobs to sites on the OSG that have GPUs
- Add `--lines='+RequestGPUs=1'` to your `jobsub_submit` command
- Example:
 - `jobsub_submit -G <vo> --memory=1000MB --disk=1GB --expected-lifetime=1h -N 8 \`
`--resource-provides=usage_model=OFFSITE --lines='+RequestGPUs=1' \`
`--jobsub-server=https://fifebatch-preprod.fnal.gov file://<path_to_executable>`
- Note items in bold above:
 - `-N 8`: Due to accounting group starvation
(<https://htcondor-wiki.cs.wisc.edu/index.cgi/wiki?p=ConsumptionPolicies>)
 - `--lines`: Add a line to the condor-submit file
 - `--jobsub-server`: steer to preprod. Right now, this isn't in production but we hope it will be in the next couple of weeks

What's new - GPU Discovery tool from OSG

- GPU information is now advertised in machine classad
- You can look for CUDA or OpenCL support
- NVIDIA driver info
- Attributes are advertised in the form of: `<LANG><KEY>=<ATTR>`

Information from one of Nebraska (Omaha)'s Nodes

- CUDAGlobalMemoryMb = 12194
- CUDARuntimeVersion = 6.0
- CUDACapability = 6.0
- CUDAComputeUnits = 56
- CUDAClockMhz = 1328.5
- CUDACoresPerCU = 192
- CUDADriverVersion = 8.0
- CUDAECCEEnabled = false
- CUDADeviceName = "Tesla P100-PCIE-12GB"
- NVIDIA driver info: NV_DRIVER = 375.66

Same but for Syracuse (SU-OG)

- OCLOpenCLVersion = 1.2
- OCLGlobalMemoryMb = 2001
- OCLComputeUnits = 5
- OCLECCEnabled = false
- OCLDeviceName = "GeForce GTX 750 Ti"
- OCLClockMhz = 1084
- NV_DRIVER = 375.66

How would you use this in job requirements?

- Using jobsub
- Example where we want to get a Tesla P100 GPU:
 - `--lines='+RequestGPUs=1' --append_condor_requirements='(regexp(\"^Tesla\\ P100.+\\\",TARGET.CUDADeviceName,\"i\")||regexp(\"^Tesla\\ P100.+\\\",TARGET.OCLDeviceName,\"i\"))'`
 - Be careful about cutting/pasting this -- unicode/ascii issues with python 2!
- Find OPENCL support
 - `--lines='+RequestGPUs=1'`
`--append_condor_requirements='isUndefined(TARGET.OCLDeviceName)==FALSE'`
- Finding CUDA support is similar (OCLDeviceName → CUDADeviceName)
- To look at all of the slot ClassAds (GPU and non-GPU-related), put “`cat ${_CONDOR_JOB_IWD}/.machine.ad`” somewhere in your executable. This will print the slot ClassAd to stdout (which gets transferred to your sandbox at the end of the job)

Our questions for you

- Is this the kind of information that might be useful?
- What other information would you like to see?
- We'd love volunteers to help test this - still in preprod

If you'd like assistance submitting to offsite GPUs, or have further suggestions about what info would be helpful, please open a Service Desk Ticket under the service "Batch Job Management" under SCD.