

HTCondor CE, Collector, AGIS: Instructions

USATLAS Facilities Meeting
OSG All-Hands San Diego 2017
John Hover <jhover@bnl.gov>

Context

Currently, T2s that use CE queues to route work must coordinate via email with the factory operators (Jose, John) to submit pilots with proper **+remote_queue** values. We want to eliminate this step, so:

- Functionality has been added in the latest OSG CE code to allow the publication of CE queue names to the OSG Information Collector, for subsequent query from AGIS. This is leveraging existing GIP configuration.
- All that T2 admins need to do manually is associate a PanDA queue with one or more CE queues in AGIS.
- Once this is done, the pilot factories can **automatically** add submission queues without involving the factory operators.

On CE

- Upgrade to latest OSG version of **osg-configure** RPMs. A full CE upgrade is not necessary.
- Edit **/etc/osg/config.d/30-gip.ini**. Create sections using 'Resource Entry' as beginning of section name, one for each Panda queue, with unique string after, e.g.:

```
[Resource Entry BNL-Queue-analylong]
```

```
name = BNL-Queue-analylong
```

```
cpucount = 1
```

```
maxmemory = 1800
```

```
max_wall_time=4320
```

```
queue = analy.long
```

```
allowed_vos=atlas
```

This creates a CE queue labelled 'analy.long'. In this case on gridgk06.racf.bnl.gov which is CE Resource BNL_ATLAS_6 in OIM. Other values will be useful, but not required now.

- Run `osg-configure -c` and `condor_ce_reconfig` to enable changes.

Check OSG Collector for info...

Query the OSG Collector to look for the defined queue...

```
curl -s http://collector.opensciencegrid.org/json/agis-compat |  
python -m json.tool > info.json
```

Finding, for example:

```
"analy.long": {  
  "ce": "BNL_ATLAS_6",  
  "entry": "BNL-Queue-analylong",  
  "max_wallclocktime": 4320,  
  "name": "analy.long",  
  "status": "Production"  
},
```

In AGIS...

In AGIS, navigate to Panda Queues for your site. (In this example **ANALY_BNL_LONG-condor**) Go to 'Associated CE queues'.

Associated CE queues

CE endpoint	flavour	version	CE.state	Site.state	Site	Queue	Queue.etf_default
gridgk03.racf.bnl.gov:9619	HTCONDOR-CE		ACTIVE	ACTIVE	BNL-ATLAS	default	yes
gridgk02.racf.bnl.gov:9619	HTCONDOR-CE		ACTIVE	ACTIVE	BNL-ATLAS	mp8	no
gridgk03.racf.bnl.gov:9619	HTCONDOR-CE		ACTIVE	ACTIVE	BNL-ATLAS	mp8	no
gridgk04.racf.bnl.gov:9619	HTCONDOR-CE		ACTIVE	ACTIVE	BNL-ATLAS	mp8	no
gridgk07.racf.bnl.gov:9619	HTCONDOR-CE		ACTIVE	ACTIVE	BNL-ATLAS	mp8	no
gridgk06.racf.bnl.gov:9619	HTCONDOR-CE		ACTIVE	ACTIVE	BNL-ATLAS	mp8	no
gridgk01.racf.bnl.gov:9619	HTCONDOR-CE		ACTIVE	ACTIVE	BNL-ATLAS	mp8	no

Find and associate another CE/Queue. Modify attached queues

In AGIS... (continued)

Search for the gatekeeper
hostname in the 'Attach another
CE/Queue' field...

RC Site ATLASite DDMEndpoint PANDA Queue Service Central Se

PANDA Queue: BNL_PROD_MCORE-condor

Already Associated CE/Queue

Queue name	CE Endpoint	CE State	Site	Remove relation
default	gridgk03.racf.bnl.gov:9619	ACTIVE	BNL-ATLAS	<input type="checkbox"/>
mp8	gridgk02.racf.bnl.gov:9619	ACTIVE	BNL-ATLAS	<input type="checkbox"/>
mp8	gridgk03.racf.bnl.gov:9619	ACTIVE	BNL-ATLAS	<input type="checkbox"/>
mp8	gridgk04.racf.bnl.gov:9619	ACTIVE	BNL-ATLAS	<input type="checkbox"/>
mp8	gridgk07.racf.bnl.gov:9619	ACTIVE	BNL-ATLAS	<input type="checkbox"/>
mp8	gridgk06.racf.bnl.gov:9619	ACTIVE	BNL-ATLAS	<input type="checkbox"/>
mp8	gridgk01.racf.bnl.gov:9619	ACTIVE	BNL-ATLAS	<input type="checkbox"/>

Remove selected relations

Attach another CE/Queue

Search CE queues:

Please, enter name (or part of name) of CE, site or queue

Search queues

In AGIS... (continued)

Check 'Add relation' for the correct queue name and Save...

group_atlas.prod.test	gridgk06.racf.bnl.gov:9619	ACTIVE	BNL-ATLAS	ACTIVE	<input type="checkbox"/>
group_atlas.prod.xl	gridgk06.racf.bnl.gov:9619	ACTIVE	BNL-ATLAS	ACTIVE	<input type="checkbox"/>
group_atlas.software	gridgk06.racf.bnl.gov:9619	ACTIVE	BNL-ATLAS	ACTIVE	<input type="checkbox"/>
group_grid	gridgk06.racf.bnl.gov:9619	ACTIVE	BNL-ATLAS	ACTIVE	<input type="checkbox"/>
amc	gridgk06.racf.bnl.gov:9619	ACTIVE	BNL-ATLAS	ACTIVE	<input type="checkbox"/>
analy.long	gridgk06.racf.bnl.gov:9619	ACTIVE	BNL-ATLAS	ACTIVE	<input checked="" type="checkbox"/>
analy.short	gridgk06.racf.bnl.gov:9619	ACTIVE	BNL-ATLAS	ACTIVE	<input type="checkbox"/>
atlas	gridgk06.racf.bnl.gov:9619	ACTIVE	BNL-ATLAS	ACTIVE	<input type="checkbox"/>
bnl-local	gridgk06.racf.bnl.gov:9619	ACTIVE	BNL-ATLAS	ACTIVE	<input type="checkbox"/>
cms	gridgk06.racf.bnl.gov:9619	ACTIVE	BNL-ATLAS	ACTIVE	<input type="checkbox"/>

Miscellany/Future

Only the HTCondorCE has this functionality. Anyone still running a GRAM/Globus CE needs to migrate ASAP.

Once all sites have AGIS squared away, we'll begin low-level, **autoconfigured** pilot submission. Once that is stable, we'll scale up and discontinue manually configured submission.

Note that long term we hope to eventually move away from CE queues and submit all jobs to a 'default' site queue. Internal job routing would then be based on job properties included in the submit (xcount, maxWallTime, maxMemory, maxDisk?, diskIO?, netIO?,)

More info/reference

More detailed queue publication notes. This talk is a summary of this info:

https://docs.google.com/document/d/1D-Z3_FTKfPVKZDe-WRsbHcc62pcj_ZPqrbVz6ULhW_Q/edit?usp=sharing

Job router recipes:

<https://twiki.grid.iu.edu/bin/view/Documentation/Release3/JobRouterRecipes>

Contact Xin for assistance: Xin Zhao <xzhao@bnl.gov>