

Day to Day Maintenance of an HTCondor-CE

Jaime Frey, Brian Lin
Center for High Throughput Computing
University of Wisconsin–Madison



Staying Updated

- Ensure your HTCSS packages are up-to-date
`yum update 'htcondor-ce*' 'condor*'`
- Subscribe to htcondor-users to get informed of the latest release announcements
<https://htcondor.org/mail-lists/#user>
- Want to test changes before applying them to production? Try the HTCSS development container (contains a local HTCondor and Slurm):
<https://github.com/htcondor/htcondor-ce/blob/V5-branch/tests/containers/README.dev.md>

HTCondor-CE Config Refresher

- HTCondor-CE configuration lives in `/etc/condor-ce/`
 - Main configuration file in `/etc/condor-ce/condor_config`
 - Add local configuration to `/etc/condor-ce/config.d/` (files processed in lexicographic order)
 - Add authentication mappings to `/etc/condor-ce/mapfiles.d/`
- `condor_ce_config_val` to inspect config values
- `condor_ce_reconfig` to apply new configuration
- ...or restart the condor-ce service just to be sure

Inspecting the Job Queue

condor_ce_q can be used to show glideins that have been submitted to your CE

```
[blin@lhcb-ce ~]$ condor_ce_q | tail
WARNING: GSI authentication is enabled by your security configuration! GSI is no longer supported.
For details, see https://htcondor.org/news/plan-to-replace-gst-in-htcss/
osgvopilot ID: 183499 10/10 10:37 - 1 - - 1 183499.0
osgvopilot ID: 183500 10/10 10:37 - 1 - - 1 183500.0
osgvopilot ID: 183501 10/10 10:39 - 1 - - 1 183501.0
osgvopilot ID: 183502 10/10 10:39 - 1 - - 1 183502.0
osgvopilot ID: 183503 10/10 10:44 - 1 - - 1 183503.0
glowvopilot ID: 183504 10/10 10:44 - 1 - - 1 183504.0

Total for query: 19 jobs; 6 completed, 0 removed, 0 idle, 12 running, 1 held, 0 suspended
Total for all users: 19 jobs; 6 completed, 0 removed, 0 idle, 12 running, 1 held, 0 suspended
```

Inspecting the Job Queue

- condor_ce_q takes all the same options as condor_q!
- ‘Spooling input data files’ is a common, harmless hold reason

```
[blin@lhcb-ce ~]$ condor_ce_q -held
WARNING: GSI authentication is enabled by your security configuration! GSI is no longer supported.
For details, see https://htcondor.org/news/plan-to-replace-gst-in-htcss/

-- Schedd: lhcb-ce.chtc.wisc.edu : <128.104.100.65:9619?... @ 10/10/22 10:43:19
  ID          OWNER          HELD_SINCE  HOLD_REASON
183457.0      osgvopilot      10/10 08:43 Spooling input data files

Total for query: 1 jobs; 0 completed, 0 removed, 0 idle, 0 running, 1 held, 0 suspended
Total for all users: 18 jobs; 7 completed, 0 removed, 0 idle, 10 running, 1 held, 0 suspended
```

Verify CE Daemons

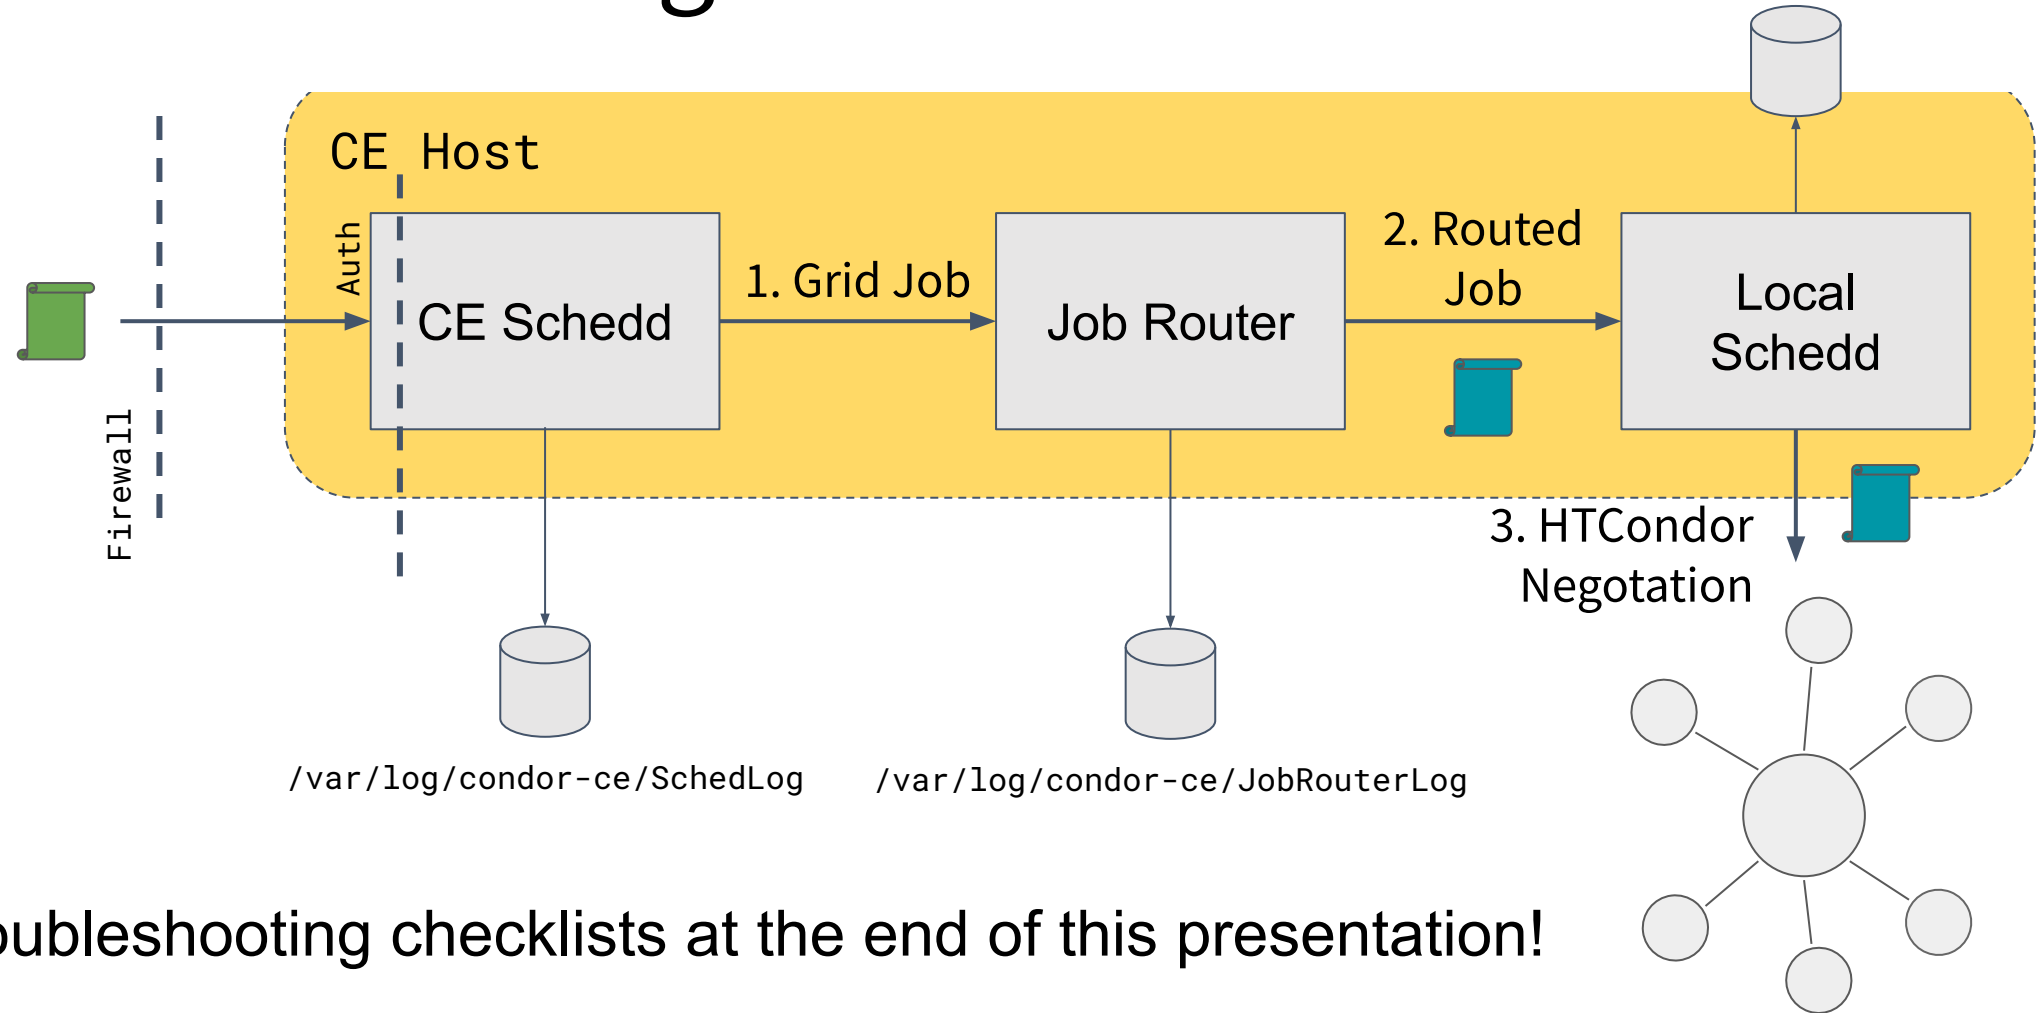
- Ensure all of the CE daemons are running with `condor_ce_status`
- If not, check `/var/log/condor-ce/` (errors can often be found in the MasterLog, SchedLog, or CollectorLog)

```
[blin@lhcb-ce ~]$ condor_ce_status -any
WARNING: GSI authentication is enabled by your security configuration! GSI is no longer supported.
For details, see https://htcondor.org/news/plan-to-replace-gst-in-htcss/
MyType           TargetType       Name
Collector         None             My Pool - lhcb-ce.chtc.wisc.edu@lhcb-ce.c
Submitter         None             glowvopilot@users.htcondor.org
Job_Router        None             htcondor-ce@lhcb-ce.chtc.wisc.edu
Scheduler         None             lhcb-ce.chtc.wisc.edu
DaemonMaster      None             lhcb-ce.chtc.wisc.edu
Submitter         None             osgvopilot@users.htcondor.org
```


Log Levels

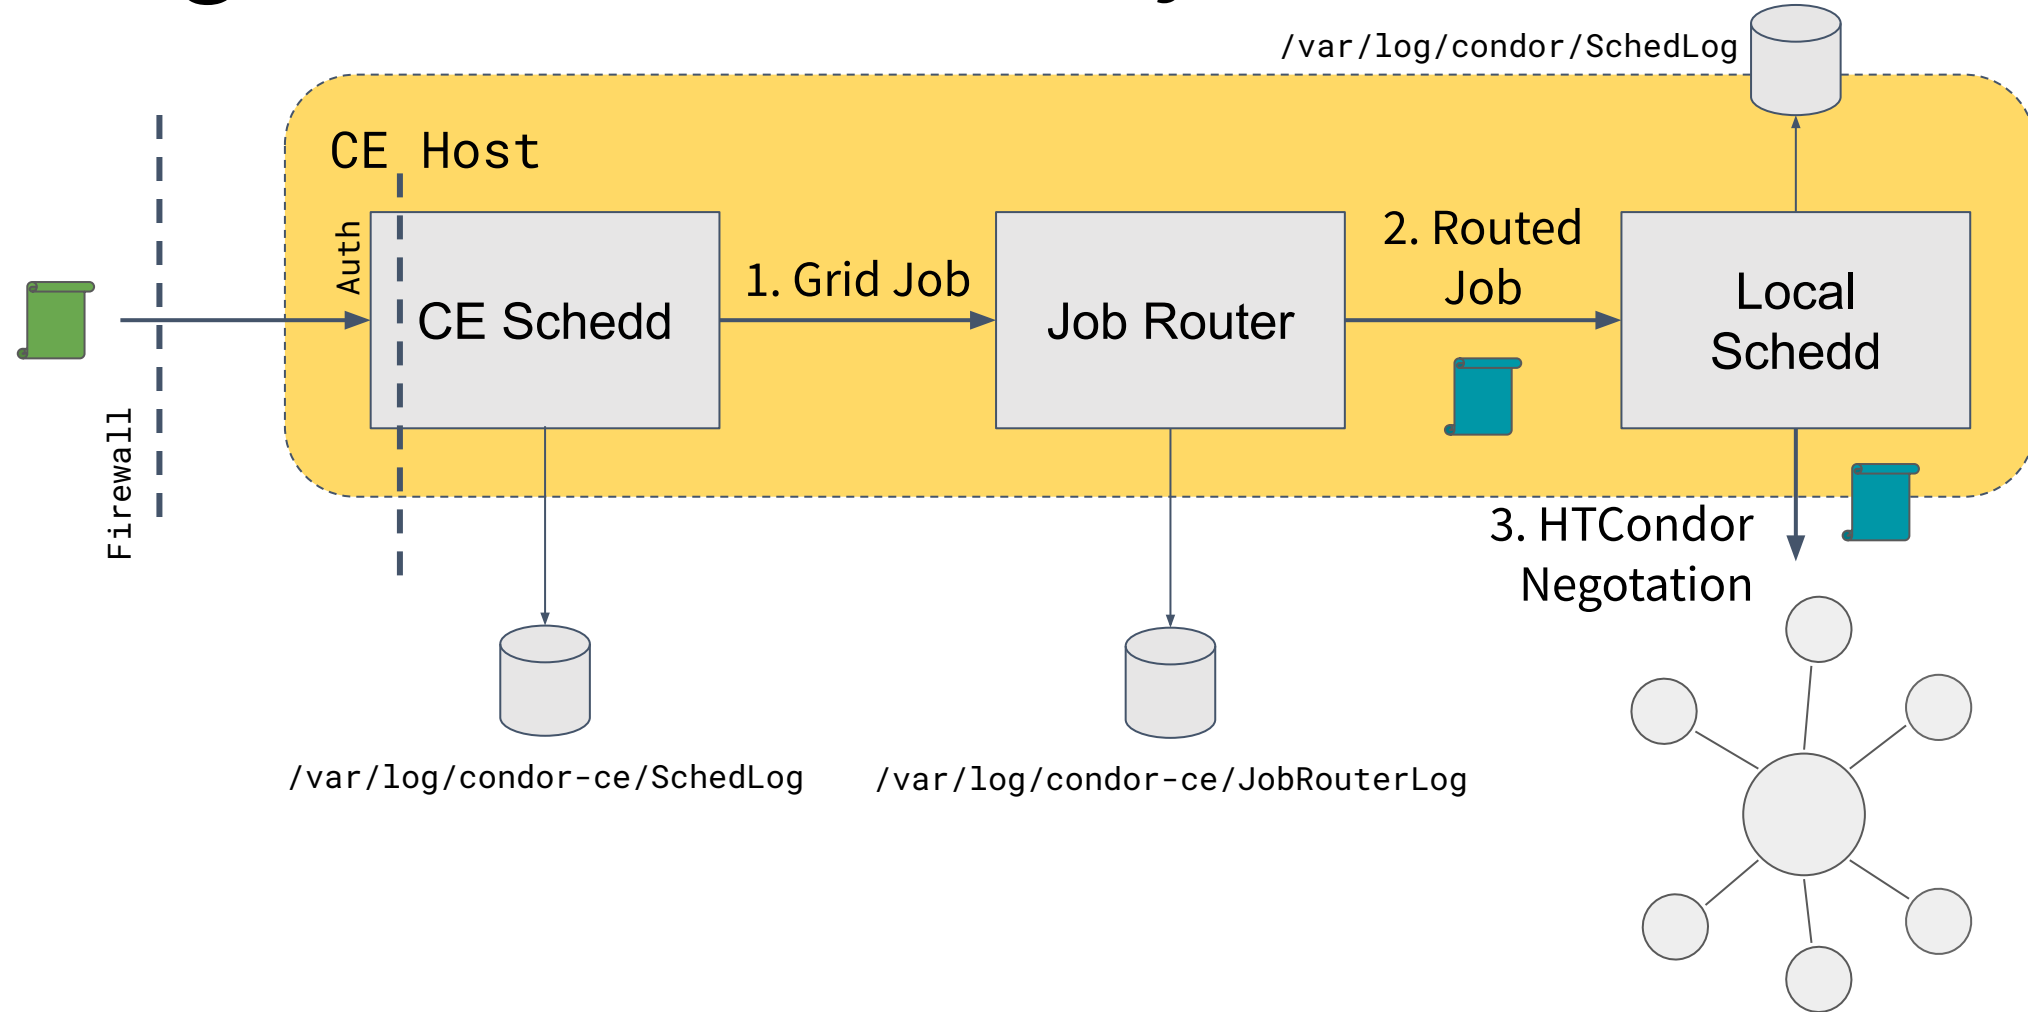
- Useful for temporary debugging
- Log level can be adjusted per daemon (e.g, `SCHEDD_DEBUG`) or across all daemons (`ALL_DEBUG`)
- Most common, helpful log levels for HTCondor-CE:
 - `D_CAT D_ALL :2` - shows the log level for each line (helpful for debugging HTCondor bugs!) and increases the log level of general messages
 - `D_SECURITY` - show verbose authentication messages
 - `D_NETWORK` - show messages for TCP/UDP connections
- Warning, this makes logs very chatty! Adjust the log sizes and number of logs kept:
 - `MAX_<SUBSYS>_LOG` - Max size of each log file, e.g. `MAX_JOB_ROUTER_LOG`
 - `MAX_NUM_<SUBSYS>_LOG` - Max number of logs kept, e.g. `MAX_NUM_JOB_ROUTER_LOG`

Troubleshooting Job Submission



See troubleshooting checklists at the end of this presentation!

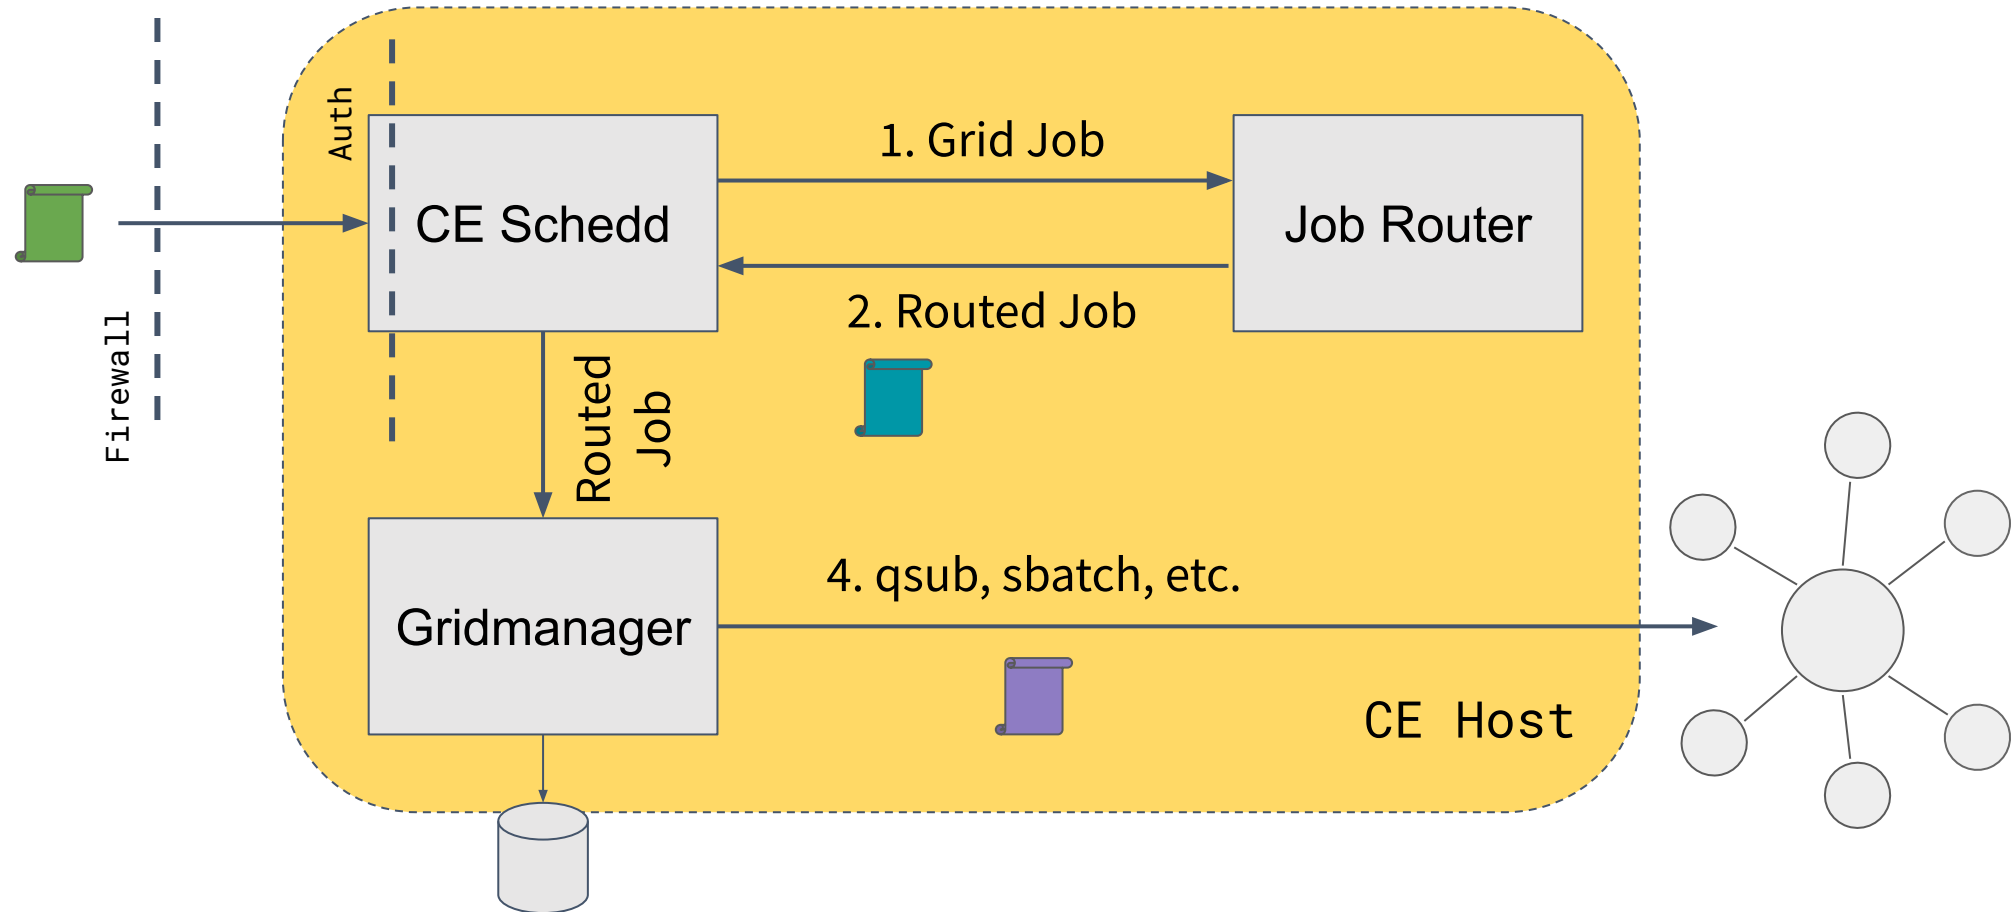
Tracking HTCondor Batch System Jobs



Tracking HTCondor Batch System Jobs

- Find the chain of job IDs using one of the following methods:
 - Query the CE schedd: `condor_ce_q <GRID JOB ID> -af RoutedToJobId`
 - Find relevant lines in the JobRouterLog
09/17/14 15:00:57 JobRouter (src=86.0,dest=205.0,route=Local_Condor):
claimed job
 - Query the local schedd: `condor_q <ROUTED JOB ID> -af RoutedFromJobId`
 - Similarly for completed jobs:
`condor_ce_history -match 1 <GRID JOB ID> -af RoutedToJobId`
`condor_history -match 1 <ROUTED JOB ID> -af RoutedFromJobId`
- <https://htcondor.com/htcondor-ce/v5/troubleshooting/common-issues/#htcondor-batch-systems>

Tracking Non-HTCondor Batch System Jobs



/var/log/condor-ce/GridmanagerLog.<user>

Tracking Non-HTCondor Batch System Jobs

- For non-HTCondor batch systems, find the batch system job ID:
 - Query the CE schedd routed job:
`$ condor_ce_q <ROUTED JOB ID> -af GridJobId`
`<snip> lsf/20141206/482046`
 - If the batch system jobs have completed, find relevant lines in the GridmanagerLog. Look for <BATCH SYSTEM>/<DATE>/<JOB ID>
`lsf/20141206/482046`
 - Alternatively, for completed jobs with HTCondor >= 8.8.11 or >= 8.9.9:
`$ condor_ce_history <ROUTED JOB ID> -match 1 -af GridJobId`
- <https://htcondor.com/htcondor-ce/v5/troubleshooting/common-issues/#non-htcondor-batch-systems>

Troubleshooting the Gridmanager

If you see failures during the GM_SUBMIT phase, this means that the Batch GAHP/BLAHP is having issues submitting jobs to the local batch system

1. Verify that local job submission to the batch system works
2. Set the following in /usr/libexec/condor/glite/etc/batch_gahp.config:

```
blah_debug_save_submit_info=<DIR_NAME>
```

This saves generated submit files that HTCondor-CE uses for submission to <DIR_NAME>

3. Test local job submission with the generated file
4. Please report any issues found this way!

<https://htcondor-ce.readthedocs.io/en/latest/troubleshooting/troubleshooting/#for-non-htcondor-batch-systems>

Troubleshooting the Gridmanager

A successful query of the local LSF batch system by the Gridmanager daemon

```
09/17/14 15:07:24 [25543] (87.0) gm state change: GM_SUBMITTED -> GM_POLL_ACTIVE
09/17/14 15:07:24 [25543] GAHP[25563] <- 'BLAH_JOB_STATUS 3 lsf/20140917/482046'
09/17/14 15:07:24 [25543] GAHP[25563] -> 'S'
09/17/14 15:07:25 [25543] GAHP[25563] <- 'RESULTS'
09/17/14 15:07:25 [25543] GAHP[25563] -> 'R'
09/17/14 15:07:25 [25543] GAHP[25563] -> 'S' '1'
09/17/14 15:07:25 [25543] GAHP[25563] -> '3' '0' 'No Error' '4' '[ BatchjobId = "482046";
JobStatus = 4; ExitCode = 0; WorkerNode = "atl-prod08" ]'
```

<https://htcondor-ce.readthedocs.io/en/latest/troubleshooting/troubleshooting/#for-non-htcondor-batch-systems>

Troubleshooting the Gridmanager

Routed job ID

```
09/17/14 15:07:24 [25543] (87.0) gm state change: GM_SUBMITTED -> GM_POLL_ACTIVE
09/17/14 15:07:24 [25543] GAHP[25563] <- 'BLAH_JOB_STATUS 3 lsf/20140917/482046'
09/17/14 15:07:24 [25543] GAHP[25563] -> 'S'
09/17/14 15:07:25 [25543] GAHP[25563] <- 'RESULTS'
09/17/14 15:07:25 [25543] GAHP[25563] -> 'R'
09/17/14 15:07:25 [25543] GAHP[25563] -> 'S' '1'
09/17/14 15:07:25 [25543] GAHP[25563] -> '3' '0' 'No Error' '4' '[ BatchjobId = "482046";
JobStatus = 4; ExitCode = 0; WorkerNode = "atl-prod08" ]'
```

<https://htcondor-ce.readthedocs.io/en/latest/troubleshooting/troubleshooting/#for-non-htcondor-batch-systems>

Troubleshooting the Gridmanager

LSF job ID

```
09/17/14 15:07:24 [25543] (87.0) gm state change: GM_SUBMITTED -> GM_POLL_ACTIVE
09/17/14 15:07:24 [25543] GAHP[25563] <- 'BLAH_JOB_STATUS 3 lsf/20140917/482046'
09/17/14 15:07:24 [25543] GAHP[25563] -> 'S'
09/17/14 15:07:25 [25543] GAHP[25563] <- 'RESULTS'
09/17/14 15:07:25 [25543] GAHP[25563] -> 'R'
09/17/14 15:07:25 [25543] GAHP[25563] -> 'S' '1'
09/17/14 15:07:25 [25543] GAHP[25563] -> '3' '0' 'No Error' '4' '[ BatchjobId = "482046";
JobStatus = 4; ExitCode = 0; WorkerNode = "atl-prod08" ]'
```

<https://htcondor-ce.readthedocs.io/en/latest/troubleshooting/troubleshooting/#for-non-htcondor-batch-systems>

Troubleshooting the Gridmanager

If there are issues, errors should show up here. If the messages do not provide enough information, run the BLAHP commands by hand:

```
/usr/libexec/condor/glite/bin/lsf_status.sh lsf/20140917/482046
```

```
09/17/14 15:07:24 [25543] (87.0) gm state change: GM_SUBMITTED -> GM_POLL_ACTIVE
09/17/14 15:07:24 [25543] GAHP[25563] <- 'BLAH_JOB_STATUS 3 lsf/20140917/482046'
09/17/14 15:07:24 [25543] GAHP[25563] -> 'S'
09/17/14 15:07:25 [25543] GAHP[25563] <- 'RESULTS'
09/17/14 15:07:25 [25543] GAHP[25563] -> 'R'
09/17/14 15:07:25 [25543] GAHP[25563] -> 'S' '1'
09/17/14 15:07:25 [25543] GAHP[25563] -> '3' '0' 'No Error' '4' '[ BatchjobId = "482046";
JobStatus = 4; ExitCode = 0; WorkerNode = "atl-prod08" ]'
```

Please report any issues found this way!

Additional Resources

- Troubleshooting documentation
<https://htcondor.com/htcondor-ce/v5/troubleshooting/common-issues/>
- Debugging tools
<https://htcondor.com/htcondor-ce/v5/troubleshooting/debugging-tools/>
- Receive news about new HTCSS versions and ask questions on htcondor-users!
<https://htcondor.org/mail-lists/#user>

Questions?

This material is based upon work supported by the National Science Foundation under Grant Nos. 1148698, 1836650 and 2030508. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

Troubleshooting Job Submission Checklist

1. No authentication errors in the SchedLog? Make sure that the firewall is open
2. Authentication errors?
 - a. Set `SCHEDD_LOG = $(SCHEDD_LOG) D_CAT D_ALL:2 D_SECURITY`
 - b. Ensure that `AUTH_SSL_*` configurations point to grid locations if using an IGTF cert (see <https://htcondor.com/htcondor-ce/v5/configuration/authentication/#configuring-certificates> and `/etc/condor-ce/config.d/01-ce-auth.conf`)
 - c. Make sure that mapped users exist
 - d. Ensure CAs, CRLs, and VO information are up-to-date
 - e. Verify CE host system clock
3. Check `/var/log/condor-ce/AuditLog` for attempted submissions

Troubleshooting Idle CE Jobs Checklist

- Jobs are indefinitely idle? Check for job router matching issues
 - For jobs still in the queue:

```
# condor_ce_q -l <JOB-ID> |  
condor_ce_job_router_info -match-jobs  
-ignore-prior-routing -jobads -
```
 - For jobs that have left the queue:
https://htcondor.com/htcondor-ce/v5/troubleshooting/debugging-tools/#condor_ce_job_router_info
- Wrap ClassAd expressions with the `debug()` function, evaluation details will appear in the JobRouterLog
- Ensure that you can submit jobs to your local batch system from the CE host
- Errors will appear in the JobRouterLog and the local SchedLog if there are communication issues between HTCondor-CE and the local HTCondor

Troubleshooting Held CE Jobs Checklist

- Find hold reasons with `condor_ce_q -held`
- **Spooling input data files:** the remote client is sending input files, should clear up after the transfer is complete
- **HTCondor-CE held job due to...**
 - **invalid job universe:** HTCondor-CE only routes vanilla, local, scheduler, and standard universe
 - **no matching routes, route job limit, or route failure threshold; see 'HTCondor-CE Troubleshooting Guide':** job sat in the queue for > 30 min without being picked up by the job router
 - No routes match the job:
`condor_ce_q <JOB ID> | condor_ce_job_router_info -match-jobs \ -ignore-prior-routing -jobads -`
 - All routes are full: `condor_ce_router_q`
 - Route failure threshold: check the JobRouterLog or GridmanagerLog for local batch system submission failures. Fix the underlying issue and restart the HTCondor-CE services.