

HTCondor-CE: Troubleshooting

ISGC 2019 - Taipei, Taiwan

Brian Lin

University of Wisconsin — Madison

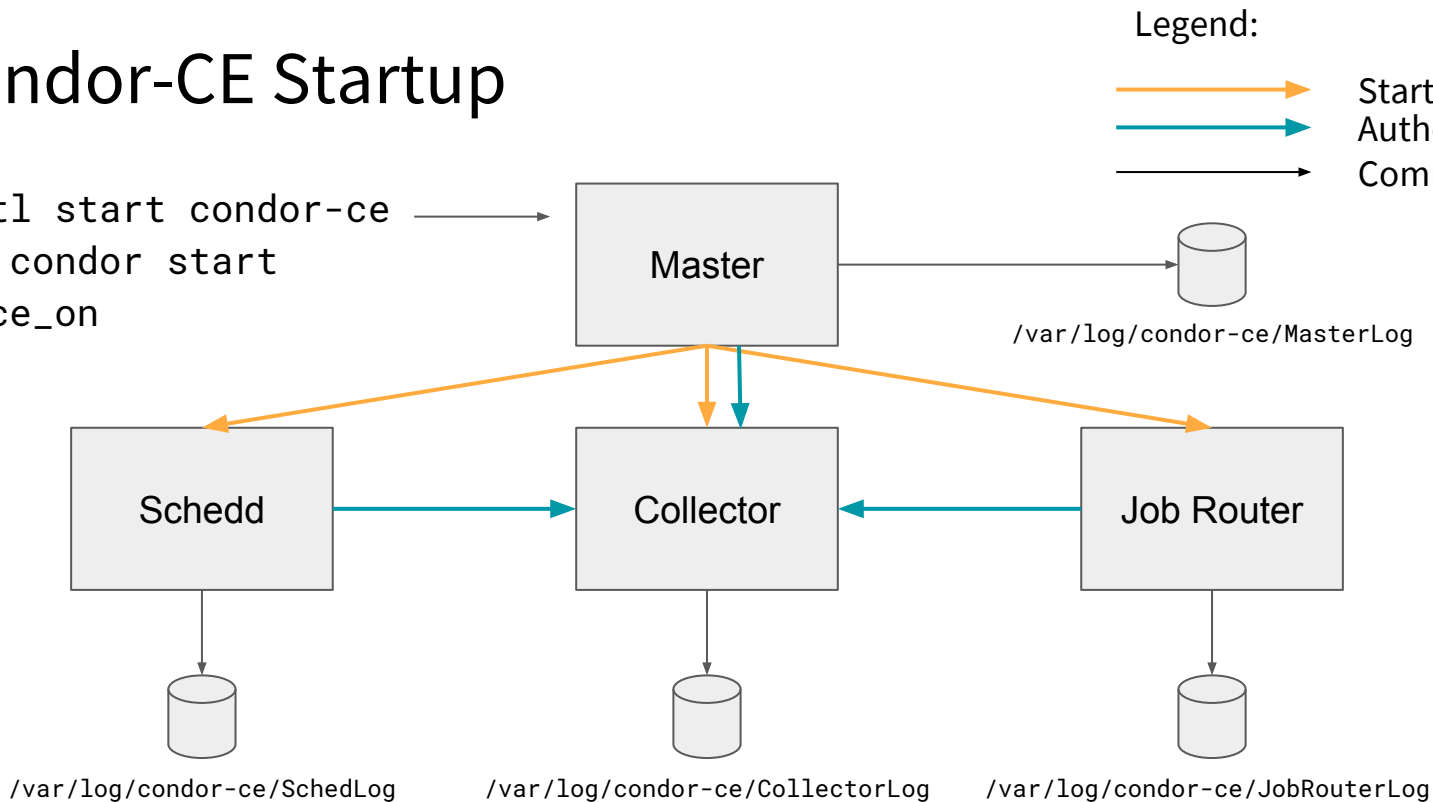


Log Levels

- 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 authentication messages
 - `D_NETWORK` - show messages for TCP/UDP connections

HTCondor-CE Startup

```
systemctl start condor-ce  
service condor start  
condor_ce_on
```



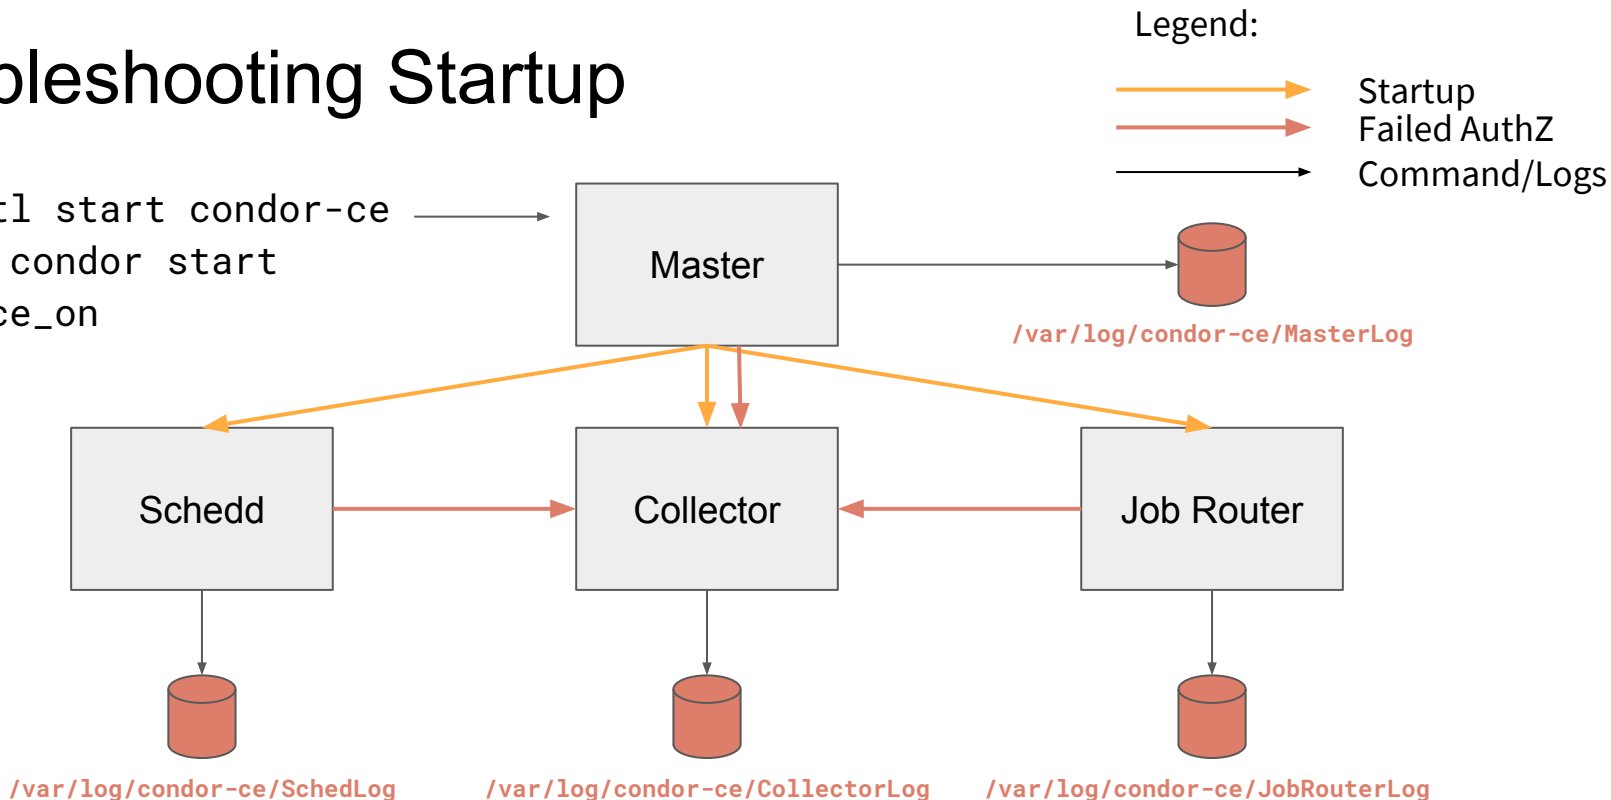
Troubleshooting Startup

If all goes well, command-line queries should show the following daemons:

```
# condor_ce_status -any
MyType           TargetType      Name
Collector        None            My Pool - fermicloud068.fnal.gov@fermiclo
Scheduler        None            fermicloud068.fnal.gov
DaemonMaster     None            fermicloud068.fnal.gov
Job_Router       None            htcondor-ce@fermicloud068.fnal.gov
```

Troubleshooting Startup

```
systemctl start condor-ce  
service condor start  
condor_ce_on
```



03/20/19 16:05:58 ERROR: AUTHENTICATE:1003:Failed to authenticate with any method

Update CA certificates and CRLs, verify host cert validity, verify unified mapfile, run condor_ce_host_network_check

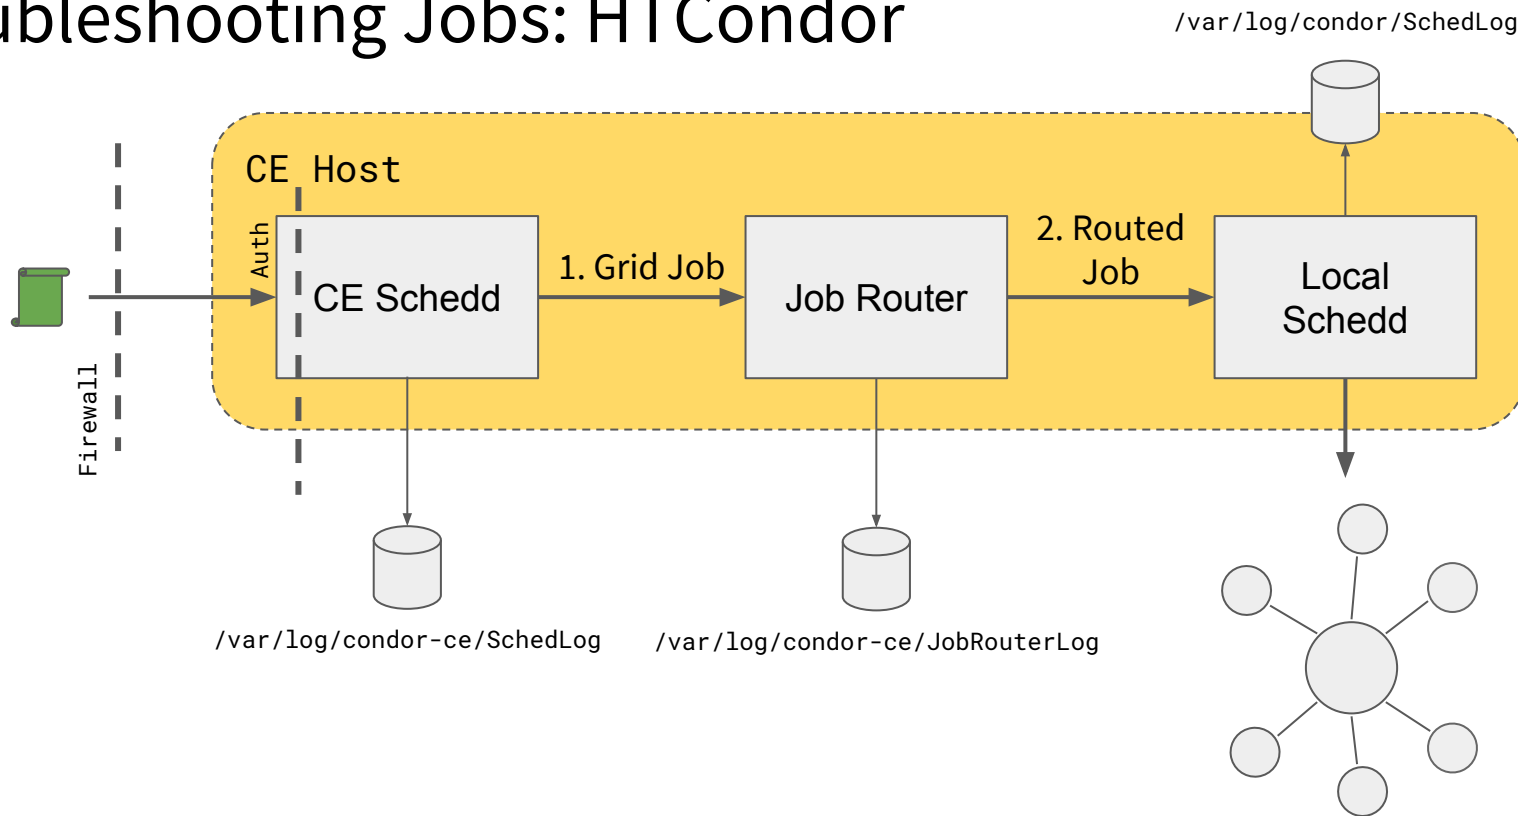
Validation

From the CE host:

1. Verify that local job submissions complete successfully from the CE host, e.g. `sbatch`, `condor_submit`, `qsub`, **etc.**
2. Verify that all required daemons are running with `condor_ce_status`
3. Verify the CE's network configuration with `condor_ce_host_network_check`
4. Verify end-to-end job submission with `condor_ce_trace`
 - a. First, from the CE host
 - b. Next, from a remote host with the `htcondor-ce-client` tools

<https://opensciencegrid.org/docs/compute-element/install-htcondor-ce/#validating-htcondor-ce>

Troubleshooting Jobs: HTCondor



Troubleshooting the CE Schedd

1. No errors in the SchedLog? Make sure that the firewall is open
2. Authentication errors? Check the condor_mapfile; make sure that mapped users exist; ensure CAs, CRLs, and VO information is up-to-date
 - a. Using LCMAPS? Also check `/var/log/messages` or `journalctl`

Troubleshooting Jobs

```
# condor_ce_q -nobatch
```

```
-- Schedd: lhcb-ce.chtc.wisc.edu : <128.104.100.65:9618?... @ 03/20/19 21:31:19
```

ID	OWNER	SUBMITTED	RUN_TIME	ST	PRI	SIZE	CMD
153501.0	nu_lhcb	3/18 13:30	2+07:56:31	R	0	733.0	DIRAC_clpM0A_pilotwrapper.py
154043.0	nu_lhcb	3/19 13:43	1+07:41:29	R	0	1709.0	DIRAC_RpJK9Q_pilotwrapper.py
154066.0	nu_lhcb	3/19 13:43	1+07:41:31	R	0	1465.0	DIRAC_RpJK9Q_pilotwrapper.py
154088.0	nu_lhcb	3/19 14:09	1+07:14:33	R	0	1709.0	DIRAC_ekQezG_pilotwrapper.py
154091.0	nu_lhcb	3/19 14:09	1+07:14:32	R	0	1709.0	DIRAC_ekQezG_pilotwrapper.py
154258.0	nu_lhcb	3/19 17:36	1+03:37:18	R	0	1221.0	DIRAC_lIr4FB_pilotwrapper.py

Troubleshooting Jobs

```
# condor_ce_q -help status
```

```
[...]
```

```
  JobStatus codes:
```

```
  1 I IDLE
```

```
  2 R RUNNING
```

```
  3 X REMOVED
```

```
  4 C COMPLETED
```

```
  5 H HELD
```

```
  6 > TRANSFERRING_OUTPUT
```

```
  7 S SUSPENDED
```

See hold reasons with `condor_ce_q -held`

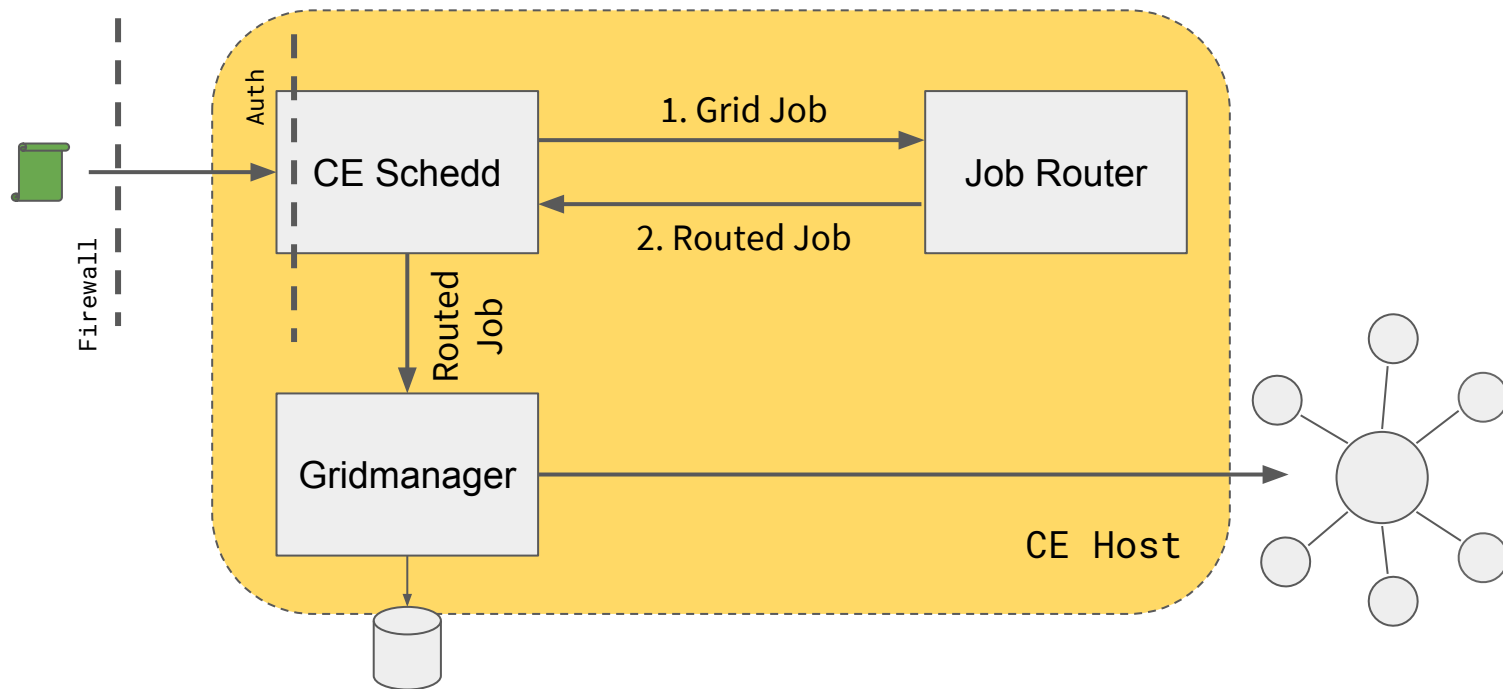
Common Hold Reasons

- **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...**
 - **missing/expired user proxy:** job X.509 proxy was removed or expired. In these cases, it's safe to remove the job (pilots are cheap)
 - **invalid job universe:** HTCondor-CE only accepts 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

Troubleshooting the Job Router

- Wrap ClassAd expressions with the `debug()` function
- 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 Jobs: Non-HTCondor Edition



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

Tracking Batch System Jobs

- Find the routed job ID using one of the following methods:
 - Query the CE schedd: `condor_ce_q -af RoutedToJobId <ORIGINAL JOB ID>`
 - 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(HTCondor-only): `condor_q -af RoutedFromJobId`
- 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 has completed, find relevant lines in the GridmanagerLog. Look for <BATCH SYSTEM>/<DATE>/<JOB ID>
`lsf/20141206/482046`
 - We're making it easier to track completed batch system jobs
<https://htcondor-wiki.cs.wisc.edu/index.cgi/tktview?tn=6159,86>

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>

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" ]'
```


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" ]'
```

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" ]'
```

Troubleshooting the Gridmanager

If there are issues, errors should show up here. If the messages do not provide enough information, run the Batch GAHP 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" ]'
```

Additional Resources

- Troubleshooting Guide

<https://opensciencegrid.org/docs/compute-element/troubleshoot-htcondor-ce>

- Additional help

htcondor-users@htcondor.org