

HTCondor-CE: Troubleshooting

HTCondor Workshop Autumn 2020

Brian Lin

University of Wisconsin–Madison



Troubleshooting HTCCondor-CE

<https://htcondor-ce.readthedocs.io/en/latest/troubleshooting/troubleshooting>

The screenshot shows the documentation website for HTCCondor-CE. The top navigation bar is dark blue with a home icon, the text "HTCCondor-CE Documentation", a search bar with a magnifying glass icon and the word "Search", and a GitHub icon with the text "GitHub".

The left sidebar contains a list of navigation links: "HTCCondor-CE Documentation", "Home", "Overview", "Installation", "Batch System Integration", "Verification", "Troubleshooting" (highlighted in blue), "Releases", and "Reference".

The main content area has a heading "HTCCondor-CE Troubleshooting Guide" with a pencil icon to its right. Below the heading is a paragraph: "In this document, you will find a collection of files and commands to help troubleshoot HTCCondor-CE along with a list of common issues with suggested troubleshooting steps."

Below this paragraph is a section titled "Known Issues". The first issue is "SUBMIT_EXPRS are not applied to jobs on the local HTCCondor". The text below explains: "If you are adding attributes to jobs submitted to your HTCCondor pool with `SUBMIT_EXPRS`, these will *not* be applied to jobs that are entering your pool from the HTCCondor-CE. To get around this, you will want to add the attributes to your [job routes](#). If the CE is the only entry point for jobs into your pool, you can get rid of `SUBMIT_EXPRS` on your backend. Otherwise, you will have to maintain your list of attributes both in your list of routes and in your `SUBMIT_EXPRS`."

Below the known issues is a section titled "General Troubleshooting Items".

The right sidebar contains a list of troubleshooting items under the heading "General Troubleshooting Items": "Making sure packages are up-to-date", "Verify package contents", "Verify clocks are synchronized", "Verify host certificates and CRLs are valid". Below this is a section titled "HTCCondor-CE Troubleshooting Items" with the following items: "Daemons fail to start", "Jobs fail to submit to the CE", "Jobs stay idle on the CE", "Idle jobs on CE: Make sure the underlying batch system can run jobs", "Idle jobs on CE: Is the job router handling the incoming job?", and "Idle jobs on CE: Verify correct operation between the CE and your local batch system".

Troubleshooting HTCondor-CE

<https://htcondor-ce.readthedocs.io/en/latest/troubleshooting/troubleshooting/#htcondor-ce-troubleshooting-items>

The screenshot shows the 'HTCondor-CE Troubleshooting Items' page. The page has a dark blue header with a search bar and a GitHub logo. On the left is a navigation menu with links like 'Home', 'Overview', 'Installation', and 'Troubleshooting'. The main content area has a title 'HTCondor-CE Troubleshooting Items' and an introductory paragraph. It contains three numbered steps: 1. Write the following into `/etc/condor-ce/config.d/99-local.conf` to increase the log level for all daemons:

```
ALL_DEBUG = D_ALWAYS:2 D_CAT
```

; 2. Ensure that the configuration is in place:

```
root@host # condor_ce_reconfig
```

; 3. Reproduce the issue. A blue 'Note' box at the bottom of the steps says: 'Before spending any time on troubleshooting, you should ensure that the state of configuration is as expected by running `condor_ce_reconfig`.' On the right is a 'Table of contents' sidebar with links to 'Known Issues', 'General Troubleshooting Items', and 'HTCondor-CE Troubleshooting Items'.

Troubleshooting HTCondor-CE

<https://htcondor-ce.readthedocs.io/en/latest/troubleshooting/troubleshooting/#htcondor-ce-troubleshooting-tools>

The screenshot shows the top navigation bar of the HTCondor-CE documentation website. The left sidebar contains a list of navigation links: Home, Overview, Installation, Batch System Integration, Verification, Troubleshooting (highlighted in blue), Releases, and Reference. The main content area is titled "HTCondor-CE Troubleshooting Tools" and contains a paragraph explaining that HTCondor-CE has its own set of tools with 'ce' in the name, such as `condor_ce_submit` and `condor_ce_run`, which are distinct from standard HTCondor commands. Below the paragraph is the heading `condor_ce_trace` and a sub-heading "Usage". The right sidebar lists several topics related to job management, including "Jobs go on hold" (with sub-points for held jobs: no matching routes, missing/expired user proxy, and invalid job universe) and "Identifying the corresponding job ID on the local batch system" (with sub-points for HTCondor and Non-HTCondor batch systems).

Troubleshooting HTCondor-CE

<https://htcondor-ce.readthedocs.io/en/latest/troubleshooting/troubleshooting/#htcondor-ce-troubleshooting-data>

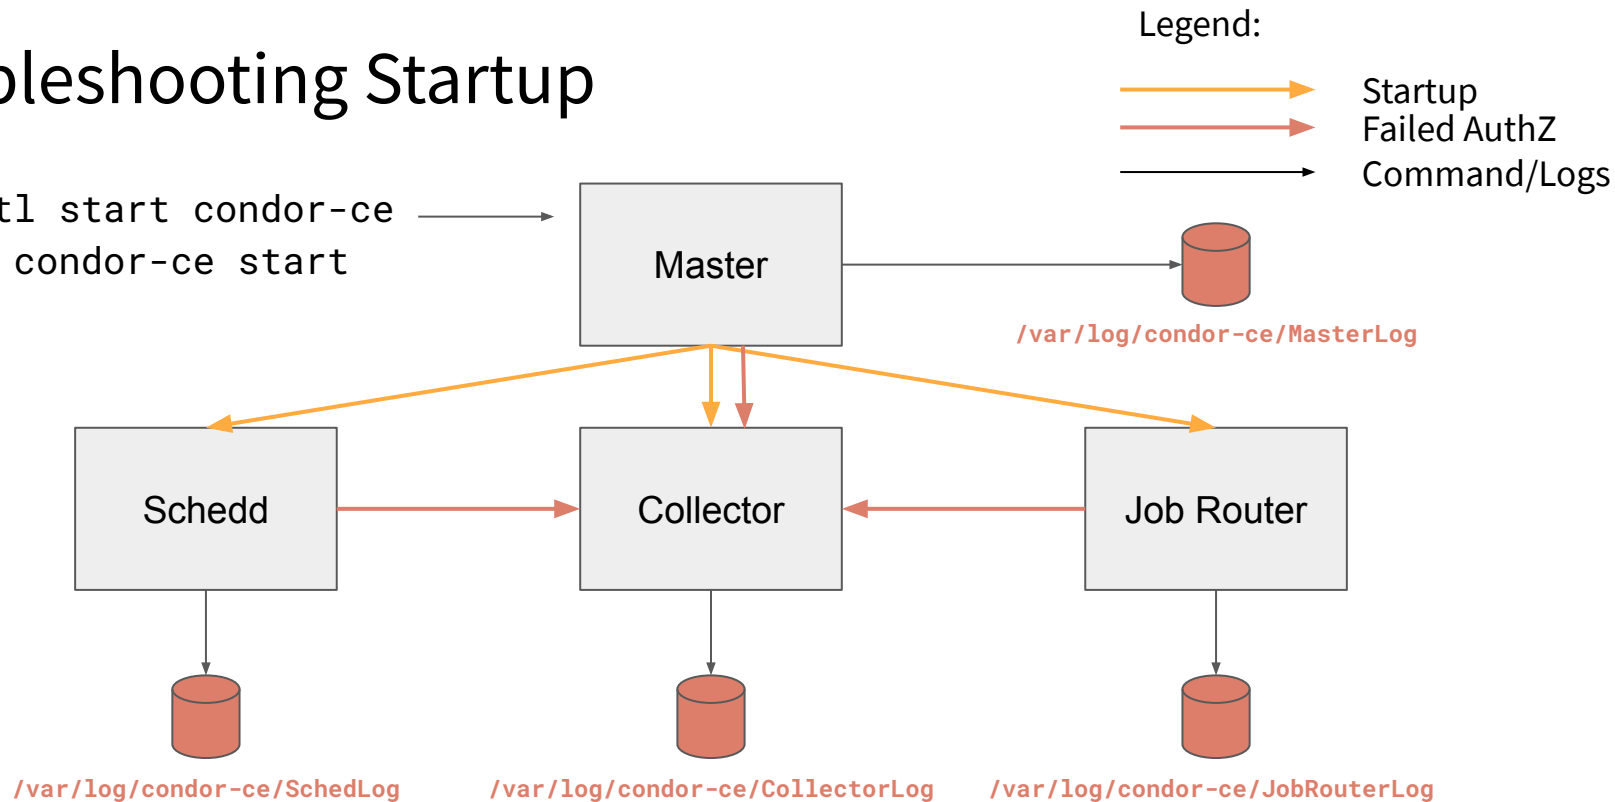
The screenshot shows a web browser displaying the HTCondor-CE documentation page for troubleshooting data. The page has a dark blue header with a search bar and a GitHub icon. On the left, there is a navigation menu with links to Home, Overview, Installation, Batch System Integration, Verification, Troubleshooting (highlighted), Releases, and Reference. The main content area is titled "HTCondor-CE Troubleshooting Data" and contains the following text: "The following files are located on the CE host." followed by a section titled "MasterLog" with the text "The HTCondor-CE master log tracks status of all of the other HTCondor daemons and thus contains valuable information if they fail to start." Below this is a list of bullet points: "Location: /var/log/condor-ce/MasterLog", "Key contents: Start-up, shut-down, and communication with other HTCondor daemons", and "Increasing the debug level:" with a sub-point "a. Set the following value in /etc/condor-ce/config.d/99-local.conf on the CE host:" followed by a code block containing "MASTER_DEBUG = D_ALWAYS:2 D_CAT". On the right side of the page, there is a vertical list of links, including "usage", "condor_ce_job_router_info", "Usage", "Troubleshooting", "condor_ce_router_q", "Usage", "condor_ce_status", "Usage", "Troubleshooting", "condor_ce_config_val", "Usage", "condor_ce_reconfig", "Usage", "condor_ce_{on,off,restart}", and "Usage".

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 Startup

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



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

**-No longer a problem in HTCondor-CE
> 4.0.0! Daemons authenticate with ^{thZ}_{d/Logs}
^{sys}
^{seI} each other over the local filesystem
instead of GSI.**

**However, configuration errors will
still prevent startup! Check journalctl
or /var/log/messages for details.**

Verifying Job Submission

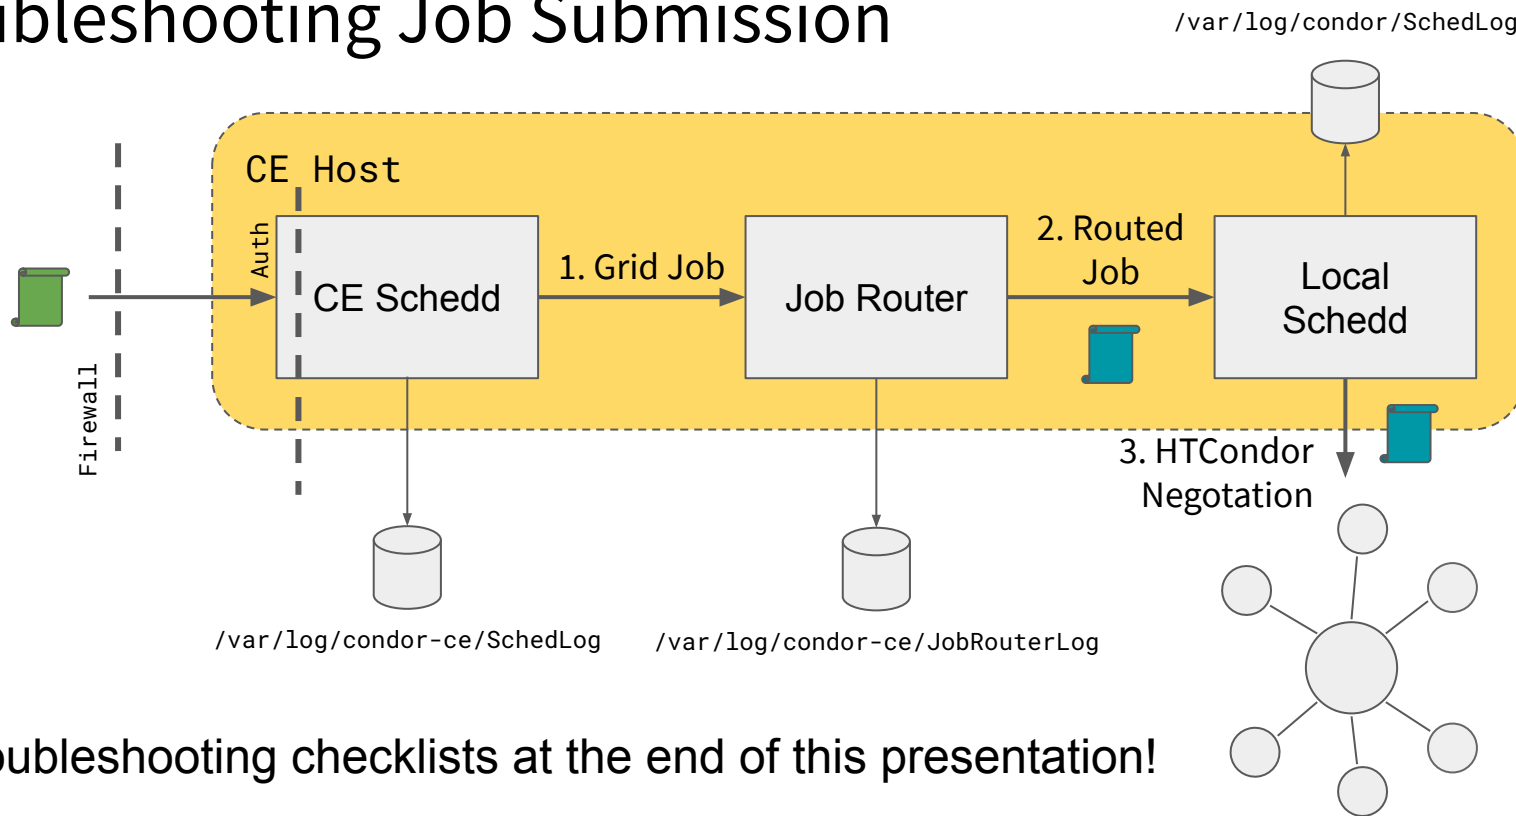
After daemons have started up, verify job submission from the CE host:

1. Verify that local job submissions complete successfully from the CE host, e.g. `sbatch`, `condor_submit`, `qsub`, etc.
2. Verify the CE's network configuration with `condor_ce_host_network_check`
3. Verify end-to-end job submission with `condor_ce_trace`
 - a. First, from the CE host
 - b. Next, from another HTCondor submit host or a remote host with `condor_ce_trace`, provided by the `htcondor-ce-client` package

<https://htcondor-ce.readthedocs.io/en/latest/verification/>

<https://htcondor-ce.readthedocs.io/en/latest/remote-job-submission/#submission-with-htcondor-submit>

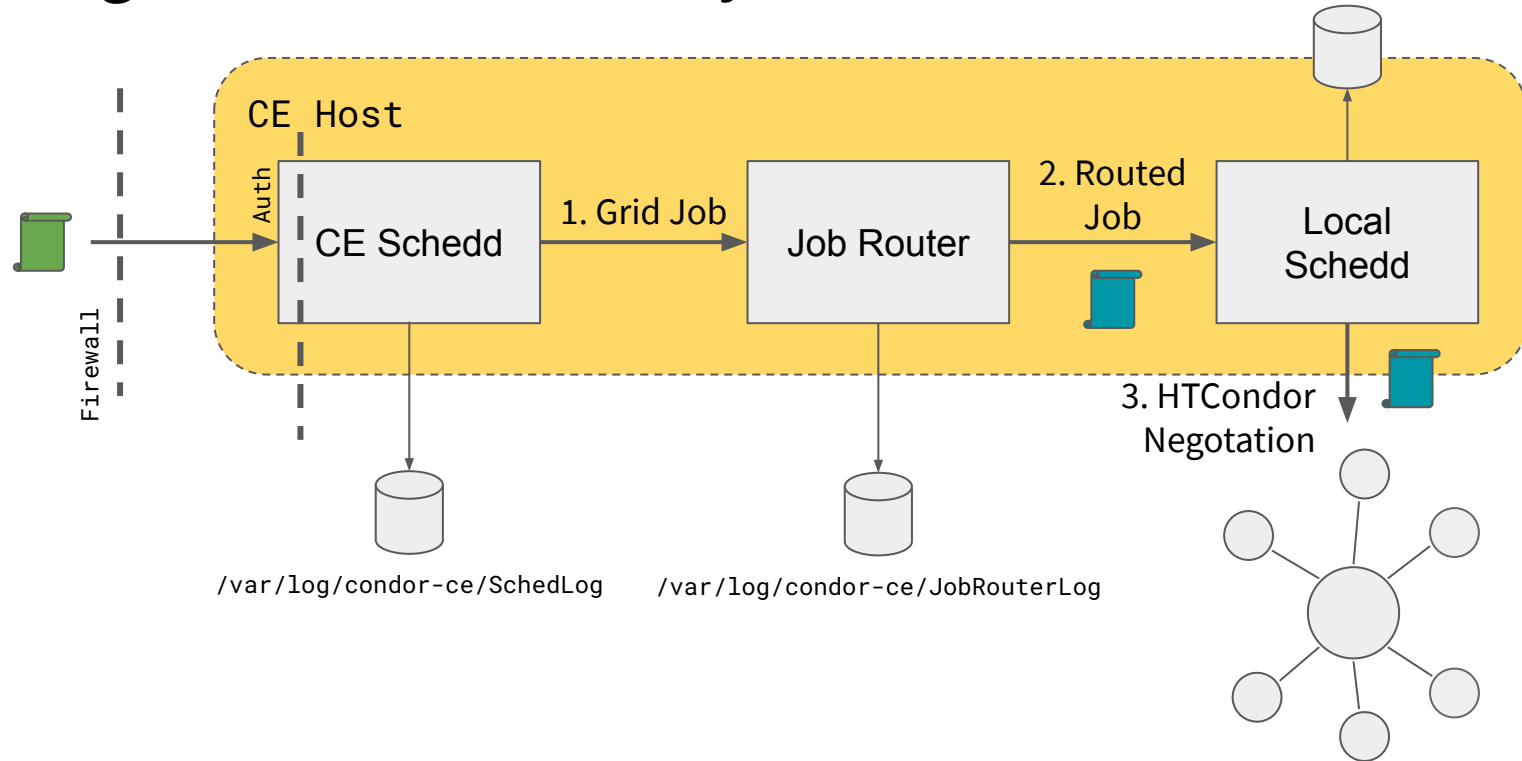
Troubleshooting Job Submission



See troubleshooting checklists at the end of this presentation!

Tracking HTCondor Batch System Jobs

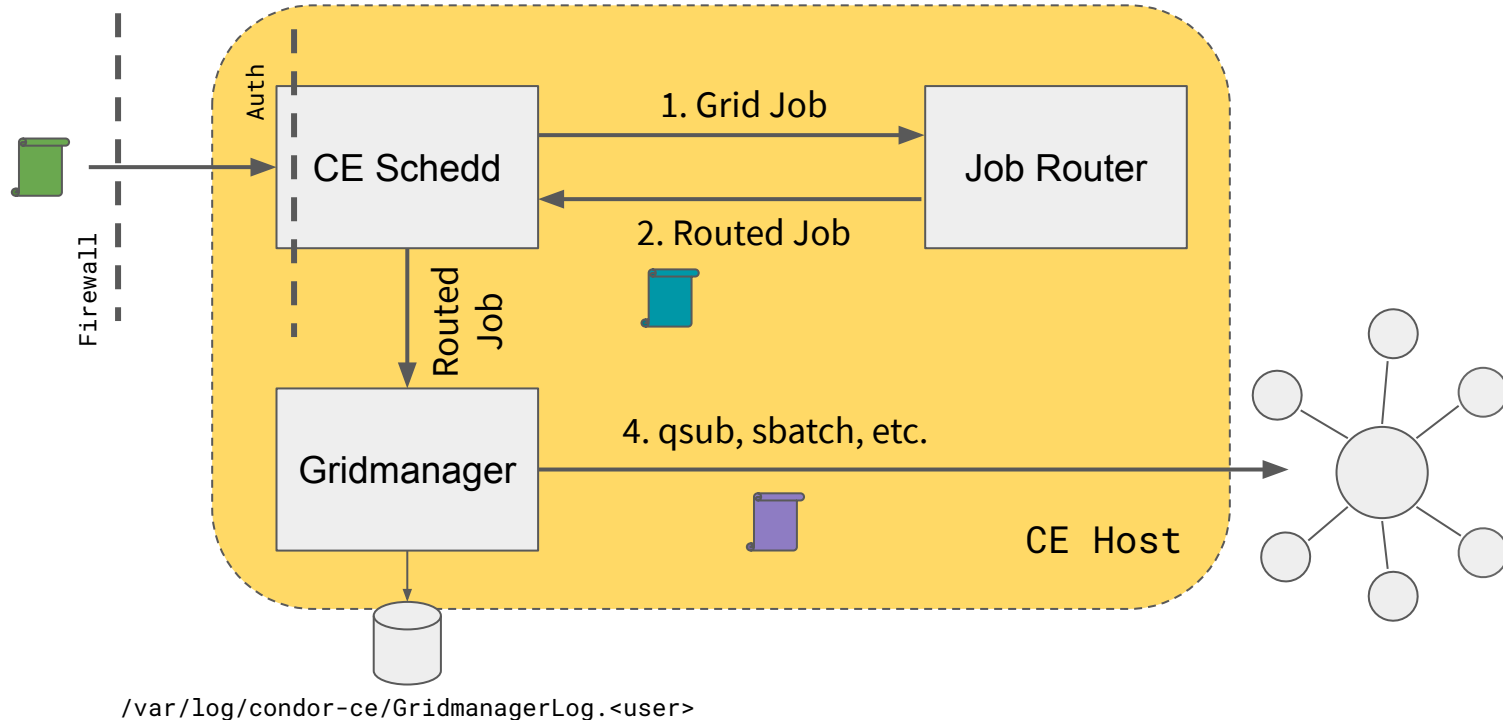
/var/log/condor/SchedLog



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-ce.readthedocs.io/en/latest/troubleshooting/troubleshooting/#htcondor-batch-systems>

Tracking Non-HTCondor Batch System Jobs



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-ce.readthedocs.io/en/latest/troubleshooting/troubleshooting/#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 Guide
<https://htcondor-ce.readthedocs.io/en/latest/troubleshooting/troubleshooting/>
- Have question, issues, or comments?
 - HTCondor-CE experts are active on htcondor-users@cs.wisc.edu!
 - Submit an issue: <https://github.com/htcondor/htcondor-ce/issues>.
 - Or better yet, a pull request: <https://github.com/htcondor/htcondor-ce/pulls>!

Troubleshooting Checklists

Daemon Startup Checklist (HTCondor-CE < 4.0.0)

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

1. Update CA certificates and CRLs
2. Verify host cert validity
`openssl x509 -in /etc/grid-security/hostcert.pem -dates`
3. Verify CE host system clock
4. Verify unified mapfile. Do you have a GSI line mapping the CE host certificate to “<hostname>@daemon.htcondor.org”?
5. Run `condor_ce_host_network_check`

If possible, consider updating to HTCondor-CE 4:

<https://htcondor-ce.readthedocs.io/en/latest/releases/#updating-to-htcondor-ce-4>

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. Check for issues with `/etc/condor-ce/condor_mapfile`
 - i. If using a callout to an authentication service via `GSS_ASSIST_GRIDMAP`, check LCMAPS/Argus error logs
 - ii. If not, do you have a GSI line mapping the job's DN or VOMS attribute to the proper user?
 - c. Make sure that mapped users exist
 - d. Ensure CAs, CRLs, and VO information are up-to-date
 - e. Verify CE host system clock

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-ce.readthedocs.io/en/latest/troubleshooting/troubleshooting/#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.