

HTCondor-CE: Configuration

HTCondor Workshop Autumn 2020

Brian Lin

University of Wisconsin–Madison



HTCondor Configuration 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)
- `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

Authentication and Authorization

- Authentication can be configured via the HTCondor-CE unified mapfile `/etc/condor-ce/condor_mapfile`
 - One mapping per line with the following format:
`<AUTH METHOD> <AUTH NAME> <HTCONDOR PRINCIPLE>`
 - Auth names supports perl-compatible regular expressions
 - Selected mapping is determined by first-match
- HTCondor principles (`<USERNAME>@<DOMAIN>`) determine authorization level
 - `<hostname>@daemon.htcondor.org`: authorized as a daemon
 - `.*@users.htcondor.org`: authorized to submit jobs
 - `GSS_ASSIST_GRIDMAP`: a special value telling HTCondor-CE to call out to another service for user mapping, e.g. LCMAPS, Argus
- <https://htcondor-ce.readthedocs.io/en/latest/installation/htcondor-ce/#configuring-authentication>

Authentication and Authorization

An example from our HTCondor-CE (version 4), lhcb-ce.chtc.wisc.edu:

```
GSI ".*,/lhcb/Role=pilot/Capability=.*" nu_lhcb
GSI "/DC=org/DC=cilogon/<snip>/CN=Brian Lin A106521" blin
GSI (.* ) GSS_ASSIST_GRIDMAP
GSI "(/CN=[- .A-Za-z0-9/= ]+)" \1@unmapped.htcondor.org
CLAIMTOBE .* anonymous@claiamtobe
FS "^(root|condor)$" \1@daemon.htcondor.org
FS (.* ) \1
```

Authentication and Authorization

Authentication method:

```
GSI ".*,/lhcb/Role=pilot/Capability=.*" nu_lhcb
GSI "/DC=org/DC=cilogon/<snip>/CN=Brian Lin A106521" blin
GSI (.* ) GSS_ASSIST_GRIDMAP
GSI "(/CN=[- .A-Za-z0-9/= ]+)" \1@unmapped.htcondor.org
CLAIMTOBE .* anonymous@claiamtobe
FS "^(root|condor)$" \1@daemon.htcondor.org
FS (.* ) \1
```

GSI is the default auth method for remote clients

Authentication and Authorization

The “authentication name”:

```
GSI ".*/lhcb/Role=pilot/Capability=.*" nu_lhcb
GSI "/DC=org/DC=cilogon/<snip>/CN=Brian Lin A106521" blin
GSI (.* ) GSS_ASSIST_GRIDMAP
GSI "(/CN=[- .A-Za-z0-9/= ]+)" \1@unmapped.htcondor.org
CLAIMTOBE .* anonymous@claiamtobe
FS "^(root|condor)$" \1@daemon.htcondor.org
FS (.* ) \1
```

In the case of GSI, authentication names can take the form:

```
<SUBJECT DN>, <VOMS FQAN 1>, ..., <VOMS FQAN N>
```

Authentication and Authorization

The HTCondor principle:

```
GSI ".*,/lhcb/Role=pilot/Capability=.*" nu_lhcb
GSI "/DC=org/DC=cilogon/<snip>/CN=Brian Lin A106521" blin
GSI (.* ) GSS_ASSIST_GRIDMAP
GSI "(/CN=[- .A-Za-z0-9/= ]+)" \1@unmapped.htcondor.org
CLAIMTOBE .* anonymous@claiamtobe
FS "^(root|condor)$" \1@daemon.htcondor.org
FS (.* ) \1
```

Or the Unix account that a matching incoming credential is mapped to

Authentication and Authorization

Explicit mapping for a single user:

```
GSI ".*,/lhcb/Role=pilot/Capability=.*" nu_lhcb
GSI "/DC=org/DC=cilogon/<snip>/CN=Brian Lin A106521" blin
GSI (.* ) GSS_ASSIST_GRIDMAP
GSI "(/CN=[- .A-Za-z0-9/= ]+)" \1@unmapped.htcondor.org
CLAIMTOBE .* anonymous@claiamtobe
FS "^(root|condor)$" \1@daemon.htcondor.org
FS (.* ) \1
```


Authentication and Authorization

Callout to external service:

```
GSI ".*,/lhcb/Role=pilot/Capability=.*" nu_lhcb
GSI "/DC=org/DC=cilogon/<snip>/CN=Brian Lin A106521" blin
GSI (.* ) GSS_ASSIST_GRIDMAP
GSI "(/CN=[- .A-Za-z0-9/= ]+)" \1@unmapped.htcondor.org
CLAIMTOBE .* anonymous@claiamtobe
FS "^(root|condor)$" \1@daemon.htcondor.org
FS (.* ) \1
```

via `/etc/grid-security/gsi-authz.conf`, containing the following for LCMAPS:

```
globus_mapping liblcas_lcmaps_gt4_mapping.so lcmaps_callout
```

Or for Argus:

```
globus_mapping /usr/lib64/libgsi_peg_callout.so argus_peg_callout
```

Authentication and Authorization

Unauthenticated fallbacks

```
GSI ".*,/lhcb/Role=pilot/Capability=.*" nu_lhcb
GSI "/DC=org/DC=cilogon/<snip>/CN=Brian Lin A106521" blin
GSI (.* ) GSS_ASSIST_GRIDMAP
GSI "(/CN=[- .A-Za-z0-9/= ]+)" \1@unmapped.htcondor.org
CLAIMTOBE .* anonymous@claimtobe
FS "^(root|condor)$" \1@daemon.htcondor.org
FS (.* ) \1
```

- `@<SUFFIX>` controls authorization levels; `@claimtobe` and `@unmapped.htcondor.org` users have few privileges
- The `UID_DOMAIN` (`@users.htcondor.org`) is automatically appended to HTCondor principles without an explicit `@<SUFFIX>`, i.e. they have submit privileges

Authentication and Authorization

Mapping for local HTCondor-CE daemon communication

```
GSI ".*,/lhcb/Role=pilot/Capability=.*" nu_lhcb
GSI "/DC=org/DC=cilogon/<snip>/CN=Brian Lin A106521" blin
GSI (.* ) GSS_ASSIST_GRIDMAP
GSI "(/CN=[- .A-Za-z0-9/= ]+)" \1@unmapped.htcondor.org
CLAIMTOBE .* anonymous@claiamtobe
FS "^(root|condor)$" \1@daemon.htcondor.org
FS (.* ) \1
```

Authentication and Authorization

Finally, map local accounts to themselves.

```
GSI ".*,/lhcb/Role=pilot/Capability=.*" nu_lhcb
GSI "/DC=org/DC=cilogon/<snip>/CN=Brian Lin A106521" blin
GSI (.* ) GSS_ASSIST_GRIDMAP
GSI "(/CN=[- .A-Za-z0-9/= ]+)" \1@unmapped.htcondor.org
CLAIMTOBE .* anonymous@claiamtobe
FS "^(root|condor)$" \1@daemon.htcondor.org
FS (.* ) \1
```

Batch System Configuration

- For HTCondor batch systems, specify the locations of your local batch SchedD, Collector, and **SP00L** directory
- For non-HTCondor batch systems, configure the BLAHP and configure how you will share the HTCondor-CE **SP00L** directory across your batch system
- <https://htcondor-ce.readthedocs.io/en/latest/installation/htcondor-ce/#configuring-the-batch-system>

Information Services

- An HTCondor-CE Central Collector requires no extra configuration, central grid operations can just install the `htcondor-ce-collector` package and start the `condor-ce-collector` service!
- To report to a Central Collector, specify the hostname and port in your HTCondor-CE config. For example, `CONDOR_VIEW_HOST = collector.opensciencegrid.org:9619`
- Advertise SchedD ads to the Central Collector (the default in HTCondor-CE 4):
`CONDOR_VIEW_CLASSAD_TYPES = Scheduler`
- Add arbitrary attributes to the SchedD ad reported to the Central Collector:
`F00 = "Bar"`
`SCHEDD_ATTRS = $(SCHEDD_ATTRS) F00`

Configuring Job Routes

Job Router Configuration

- Declare your site policy
- Each route is described by job transforms (see TJ's talk yesterday)
- Job routes are constructed by combining each entry in `JOB_ROUTER_ENTRIES` with the `JOB_ROUTER_DEFAULTS`
- <https://htcondor-ce.readthedocs.io/en/latest/batch-system-integration/>

```
$ condor_ce_job_router_info
-config
Route 1
Name      : "Local_Condor"
Universe  : 5
MaxJobs   : 10000
MaxIdleJobs : 2000
GridResource :
Requirements : true
ClassAd    :
    [
    [...]
```


Job Router Matching

- Each job that matches is compared to each job route's requirements expression (`Requirements = True` by default) in order
- Specify the order by setting `JOB_ROUTER_ROUTE_NAMES` to a list route names from `JOB_ROUTER_ENTRIES`
- To use round-robin matching behavior, set the following in your configuration (not within the routes):
`JOB_ROUTER_ROUND_ROBIN_SELECTION = True`

Job Router ClassAds

Special job route verbs are used to transform jobs, evaluated in the following order.

1. Copy an attribute from the original job ad to the routed job ad:

```
copy_foo = "original_foo";
```

2. Delete an attribute from the original job ad from the routed job ad:

```
delete_foo = True;
```

3. Set an attribute in the routed job ad to a value or expression

```
set_requirements = (OpSys == "LINUX");
```

4. Set an attribute in the routed job ad to value that is evaluated in the context of the original job ad.

```
eval_set_Experiment = strcat("cms.", Owner);
```

set_ vs eval_set_

In a job route, `set_Experiment = strcat("cms.", Owner);`
→ `Experiment = strcat("cms.", Owner)` in the routed job

Use `eval_set_*` if you just want the value of an expression in the routed job.

In a job route, `Owner = blin; eval_set_Experiment = strcat("cms.", Owner);`
→ `Experiment = cms.blin` in the routed job

set_ vs eval_set_

Use `set_*` for any of the following scenarios:

- To use ClassAd expressions in the routed job on your HTCondor batch system
- To set intermediate ClassAd attributes for multiple `eval_set_*` statements
- For non-HTCondor batch systems to `set_default_CERequirements`
- To set default resource requests

Setting Default Resource Requests

Use `set_*` for the following resource requests

- `default_xcount` to set the default number of cores
- `default_maxMemory` to set the default maximum memory (in MB)
- `default_maxWalltime` to set the default maximum walltime (in minutes)
- `default_queue` to set the default batch system queue (non-HTCondor only)

Job Router Defaults

HTCondor-CE automatically generates `JOB_ROUTER_DEFAULTS`. Modify existing attributes at your own risk but feel free to add to it!

For HTCondor:

```
JOB_ROUTER_DEFAULTS @=jrd
$(JOB_ROUTER_DEFAULTS)
[
  # Route jobs to an HTCondor batch system
  TargetUniverse = 5;
]
@jrd
```

For Slurm:

```
JOB_ROUTER_DEFAULTS @=jrd
$(JOB_ROUTER_DEFAULTS)
[
  # Route jobs to a Slurm batch system
  GridResource = "batch slurm"
]
@jrd
```

Job Router Entries

- Use the multiline config syntax
- Each route is enclosed by []
- Each route requires a Name attr
- Since we're using ClassAds, we can use [ClassAd functions!](#)
- TARGET ensures that the job attribute is used to match the route

```
JOB_ROUTER_ENTRIES @=jre
[
  Name = "atlas_mcore";
  TargetUniverse = 5;
  Requirements = regexp("^usatlas",
TARGET.Owner);
  set_default_xcount = 8;
]
[
  Name = "everything_else";
  TargetUniverse = 5;
  set_default_xcount = 1;
]
@jre
```

Job Router Entries

- Use the multiline config syntax
- Each route is enclosed by []
- Each route requires a Name attr
- Since we're using ClassAds, we can use [ClassAd functions!](#)
- TARGET ensures that the job attribute is used to match the route

```
JOB_ROUTER_ENTRIES @=jre
[
    Name = "atlas_mcore";
    TargetUniverse = 5;
    Requirements = regexp("^usatlas",
TARGET.Owner);
    set_default_xcount = 8;
]
[
    Name = "everything_else";
    TargetUniverse = 5;
    set_default_xcount = 1;
]
@jre
```


Job Router Entries

- Use the multiline config syntax
- Each route is enclosed by []
- Each route requires a Name attr
- Since we're using ClassAds, we can use [ClassAd functions!](#)
- TARGET ensures that the job attribute is used to match the route

```
JOB_ROUTER_ENTRIES @=jre
[
  Name = "atlas_mcore";
  TargetUniverse = 5;
  Requirements = regexp("^usatlas",
TARGET.Owner);
  set_default_xcount = 8;
]
[
  Name = "everything_else";
  TargetUniverse = 5;
  set_default_xcount = 1;
]
@jre
```

Job Router Entries

- Use the multiline config syntax
- Each route is enclosed by []
- Each route requires a Name attr
- Since we're using ClassAds, we can use [ClassAd functions!](#)
- TARGET ensures that the job attribute is used to match the route

```
JOB_ROUTER_ENTRIES @=jre
[
  Name = "atlas_mcore";
  TargetUniverse = 5;
  Requirements = regexp("^usatlas",
TARGET.Owner);
  set_default_xcount = 8;
]
[
  Name = "everything_else";
  TargetUniverse = 5;
  set_default_xcount = 1;
]
@jre
```

Job Router Entries

- Use the multiline config syntax
- Each route is enclosed by []
- Each route requires a Name attr
- Since we're using ClassAds, we can use [ClassAd functions!](#)
- TARGET ensures that the job attribute is used to match the route

```
JOB_ROUTER_ENTRIES @=jre
[
  Name = "atlas_mcore";
  TargetUniverse = 5;
  Requirements = regexp("^usatlas",
TARGET.Owner);
  set_default_xcount = 8;
]
[
  Name = "everything_else";
  TargetUniverse = 5;
  set_default_xcount = 1;
]
@jre
```

Non-HTCondor Route Configuration

For non-HTCondor batch system directives not covered, there's `set_default_CERequirements` (supported in HTCondor-CE 4):

```
[  
  Name = "New CERequirements format";  
  set_WallTime = 3600;  
  set_AccountingGroup = x509UserProxyFirstFQAN;  
  set_default_CERequirements = "Walltime,AccountingGroup";  
]
```

Results in `$Walltime` and `$AccountingGroup` shell variables that can be used in the relevant `/usr/libexec/condor/glite/bin/*_local_submit_attributes.sh` for your batch system.

Additional Resources

- Documentation
 - <https://htcondor-ce.readthedocs.io/en/latest/installation/htcondor-ce/#configuring-htcondor-ce>
 - <https://htcondor-ce.readthedocs.io/en/latest/batch-system-integration/>
- 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>!