Table of Contents

How to set up a minimal, functional example HTCondor CE cluster

First set up a mini HTCondor cluster following the **Admin Quick Start Guide**:

https://research.cs.wisc.edu/htcondor/htcondor/documentation/

The **Long Term Support (LTS) Channel** (see below) concerns `v9.0.x` whose EOL will be **May 2023**: it supports **X509** proxies for *authentication* **and** *delegation*, whereas the releases in the **Feature Channel** only support the latter purpose, i.e. equipping jobs with such proxies. Both channels support `SciTokens` for authentication. In the course of 2022 and early 2023 we will need to make job submission with tokens work for HTCondor CEs across the infrastructure. A minimal `SciTokens` example configuration is shown below.

Notes on using the Long Term Support Channel

Note: the Admin Quick Start Guide defaults to the Feature Channel.
To deploy the LTS a.k.a. ***stable*** release, one can ***imitate*** the following steps,
which also prevent a fatal error encountered on `CC7` hosts.

We use a ***patched*** version of the `get` script that still has "stable" pointing to the `9.0.x` series...

We start with the **Central Manager** (CM) host:

```
------------------------------------------------------------------------
[root@htc-cm ~]# yum remove epel-release-7
[...]
------------------------------------------------------------------------
[root@htc-cm ~]# curl -fsSL https://twiki.cern.ch/twiki/pub/LCG/MiniHTCsetup/get.sh | \
GET_HTCONDOR_PASSWORD=<cluster-password> \
/bin/bash -s -- --no-dry-run --channel stable --central-manager htc-cm.your-domain
[...]
------------------------------------------------------------------------
[root@htc-cm ~]# rpm -q condor
condor-9.0.17-1.el7.x86_64
------------------------------------------------------------------------
```

Similarly for the **Submit Node** (CE):

```
------------------------------------------------------------------------
[root@htc-ce ~]# yum remove epel-release-7
[...]
------------------------------------------------------------------------
[root@htc-ce ~]# curl -fsSL https://twiki.cern.ch/twiki/pub/LCG/MiniHTCsetup/get.sh | \
GET_HTCONDOR_PASSWORD=<cluster-password> \
/bin/bash -s -- --no-dry-run --channel stable --submit htc-cm.your-domain
[...]
------------------------------------------------------------------------
[root@htc-ce ~]# rpm -q condor
condor-9.0.17-1.el7.x86_64
------------------------------------------------------------------------
```

And the **Execute Node** (WN):

```
------------------------------------------------------------------------
[root@htc-wn ~]# yum remove epel-release-7
[...]
------------------------------------------------------------------------
[root@htc-wn ~]# curl -fsSL https://twiki.cern.ch/twiki/pub/LCG/MiniHTCsetup/get.sh | \
GET_HTCONDOR_PASSWORD=<cluster-password> \
/bin/bash -s -- --no-dry-run --channel stable --execute htc-cm.your-domain
[...]
------------------------------------------------------------------------
[root@htc-wn ~]# rpm -q condor
condor-9.0.17-1.el7.x86_64
------------------------------------------------------------------------
```

Following the selected guide, you will have a Central Manager (**CM**), an Execute Node (**WN**) and a Submit Node (**CE**) that must ***only*** be used for **grid** jobs submitted through its HTCondor-CE interface!

Firewall rules:

- The Submit Node hosting the CE needs to have port `9619` open for grid job submissions.
- The Submit Node(s), CM and all WN need to have port `9618` open only ***between*** them.

Note: the **Admin Quick Start Guide** will set port `9618` open to the world.

The Central Manager (CM)

The HTCondor configuration should resemble the following:

```
------------------------------------------------------------------------
[root@htc-cm ~]# ll /etc/condor/config.d/
total 4
-rw-r--r--. 1 root root 148 May 17 21:24 01-central-manager.config
------------------------------------------------------------------------
[root@htc-cm ~]# cat /etc/condor/config.d/01-central-manager.config
CONDOR_HOST = htc-cm.your-domain
# For details, run condor_config_val use role:get_htcondor_central_manager
use role:get_htcondor_central_manager
------------------------------------------------------------------------
[root@htc-cm ~]# ll /etc/condor/passwords.d/
total 4
-rw-------. 1 root root 8 May 17 21:24 POOL
------------------------------------------------------------------------
[root@htc-cm ~]# ll /etc/condor/tokens.d/
total 4
-rw-------. 1 root root 250 May 17 21:24 condor@htc-cm.your-domain
------------------------------------------------------------------------
```

The Execute Node (WN)

NOTE: the job *scratch* directories need to be on a file system that is big enough to support all concurrently *running* jobs! They will be located under the directory named by the `EXECUTE` macro (by default `/var/lib/condor/execute`) in the HTCondor configuration. Mind the directory has to be owned by user `condor` and its mode has to be `755` (else commands like `pwd` would fail for jobs).

The HTCondor configuration should further resemble the following:

```
------------------------------------------------------------------------
[root@htc-wn ~]# ll /etc/condor/config.d/
total 4
-rw-r--r--. 1 root root 132 May 17 21:29 01-execute.config
------------------------------------------------------------------------
[root@htc-wn ~]# cat /etc/condor/config.d/01-execute.config
CONDOR_HOST = htc-cm.your-domain
# For details, run condor_config_val use role:get_htcondor_execute
use role:get_htcondor_execute
------------------------------------------------------------------------
[root@htc-wn ~]# ll /etc/condor/passwords.d/
total 4
-rw-------. 1 root root 8 May 17 21:29 POOL
------------------------------------------------------------------------
[root@htc-wn ~]# ll /etc/condor/tokens.d/
total 4
-rw-------. 1 root root 250 May 17 21:29 condor@htc-cm.your-domain
------------------------------------------------------------------------
```

The remaining 3 configuration files allow the WN to be used both for ***local*** and ***grid*** jobs.

First, jobs from any `submit node` with the ***same*** `UID_DOMAIN` will normally be run under the submitter's own account (modulo several security checks and restrictions). The `UID_DOMAIN` typically can be set to the DNS domain under which the submitter and execute nodes are registered. Example:

```
------------------------------------------------------------------------
[root@htc-wn ~]# cat /etc/condor/config.d/70-uid-domain
UID_DOMAIN = your-domain
------------------------------------------------------------------------
```

The next configuration file will give each job its private instances of the `/tmp` and `/var/tmp` directories, preventing pollution of the corresponding directories on the host:

```
------------------------------------------------------------------------
[root@htc-wn ~]# cat /etc/condor/config.d/71-mount-dirs
MOUNT_UNDER_SCRATCH = "/tmp,/var/tmp"
------------------------------------------------------------------------
```

If you know there will be ***no*** local users with ***standard*** home directories like `/home/$USER` expected to be mounted on the WN, then the following alternative is viable and would allow grid accounts to have standard home directories as well (see the CE section below):

```
------------------------------------------------------------------------
[root@htc-wn ~]# cat /etc/condor/config.d/71-mount-dirs
MOUNT_UNDER_SCRATCH = ifThenElse(isUndefined(Owner), "/tmp, /var/tmp", strcat("/tmp, /var/
------------------------------------------------------------------------
```

Next there are 2 options for running ***grid*** jobs:

Option #1: grid jobs run under *slot* accounts

The next configuration file defines `slot accounts` under which jobs from users of a different `UID_DOMAIN` will run. The HTCondor CE will be configured with its own, default `UID_DOMAIN` and hence ***grid*** jobs will run under ***slot*** accounts. There must be at least as many such accounts as the number of slots, the rest are ignored. Also, we can indicate that those accounts are only used for HTCondor jobs, which lets HTCondor ensure no grid job can leave any processes behind. However, on recent Linux kernels (e.g. under CentOS 7), HTCondor will anyway make use of `cgroups` to capture all processes of a job and ensure the remaining ones will all be killed at the end of the job. The slot accounts are best created *without* their home directories, to prevent any pollution of the latter by jobs.

```
------------------------------------------------------------------------
[root@htc-wn ~]# cat /etc/condor/config.d/72-slot-users
NUM_SLOTS = 3
SLOT1_USER = slot001
SLOT2_USER = slot002
SLOT3_USER = slot003
SLOT4_USER = slot004
SLOT5_USER = slot005
SLOT6_USER = slot006
SLOT7_USER = slot007
DEDICATED_EXECUTE_ACCOUNT_REGEXP = slot[0-9]+
------------------------------------------------------------------------
[root@htc-wn ~]# tail -n 7 /etc/passwd
slot001:x:19987:19987:HTCondor slot 001:/home/slot001:/bin/bash
slot002:x:19988:19988:HTCondor slot 002:/home/slot002:/bin/bash
slot003:x:19989:19989:HTCondor slot 003:/home/slot003:/bin/bash
slot004:x:19990:19990:HTCondor slot 004:/home/slot004:/bin/bash
slot005:x:19991:19991:HTCondor slot 005:/home/slot005:/bin/bash
slot006:x:19992:19992:HTCondor slot 006:/home/slot006:/bin/bash
slot007:x:19993:19993:HTCondor slot 007:/home/slot007:/bin/bash
------------------------------------------------------------------------
```

Option #2: grid jobs run under *user* accounts

This choice would make it easier to see which user is running what processes on a WN. In this case the HTCondor CE has to be configured with the ***same*** `UID_DOMAIN` as used on the WN and the grid user accounts should be defined ***consistently*** on the CE (see the next section) and the WN. Furthermore, the CE mappings must generally prevent that any given account might be concurrently used for unrelated workflows: unprivileged users must all be mapped to ***separate*** accounts, whereas e.g. production manager workflows for a VO could share a production manager account for that VO. In general one would need to create many more numbered accounts than shown in this simple example:

```
------------------------------------------------------------------------
[root@htc-wn ~]# tail -n 2 /etc/passwd
alicesgm:x:19984:19984:alicesgm:/tmp:/bin/bash
alice001:x:19985:19985:alice001:/tmp:/bin/bash
------------------------------------------------------------------------
```

The Submit Node (CE)

The HTCondor configuration should resemble the following:

```
------------------------------------------------------------------------
[root@htc-ce ~]# ll /etc/condor/config.d/
total 8
-rw-r--r--. 1 root root 130 May 17 20:32 01-submit.config
-rw-r--r--. 1 root root 451 Dec 21 22:11 50-condor-ce-defaults.conf
------------------------------------------------------------------------
[root@htc-ce ~]# cat /etc/condor/config.d/01-submit.config
CONDOR_HOST = htc-cm.your-domain
# For details, run condor_config_val use role:get_htcondor_submit
use role:get_htcondor_submit
------------------------------------------------------------------------
[root@htc-ce ~]# ll /etc/condor/passwords.d/
total 4
-rw-------. 1 root root 8 May 17 20:32 POOL
------------------------------------------------------------------------
[root@htc-ce ~]# ll /etc/condor/tokens.d/
total 4
-rw-------. 1 root root 250 May 17 20:32 condor@htc-cm.your-domain
------------------------------------------------------------------------
```

For its CE interface, ensure the host has a certificate, the CAs and the desired VOMS configuration details. The IGTF Certificate Authorities can be installed from the `ca-policy-egi-core` rpm available from the EGI CA repository. The VOMS details for WLCG VOs can be installed from `wlcg-voms-*` rpms available from the WLCG rpm repository. The directories in question should resemble what is shown here:

```
------------------------------------------------------------------------
[root@htc-ce ~]# ll /etc/grid-security/
total 76
drwxr-xr-x. 2 root root 40960 May 18 18:48 certificates
-rw-r--r--. 1 root root  3198 Mar 13 04:07 gsi.conf
-r--r--r--. 1 root root  3060 May 18 15:18 hostcert.pem
-r--------. 1 root root  1828 May 18 15:18 hostkey.pem
drwxr-xr-x. 5 root root    44 May 18 15:08 vomsdir
------------------------------------------------------------------------
[root@htc-ce ~]# ll /etc/grid-security/vomsdir/
total 0
drwxr-xr-x. 2 root root 60 May 18 15:08 alice
drwxr-xr-x. 2 root root 37 May 18 15:08 dteam
drwxr-xr-x. 2 root root 60 May 18 15:08 lhcb
------------------------------------------------------------------------
[root@htc-ce ~]# ll /etc/grid-security/vomsdir/alice/
total 8
-rw-r--r--. 1 root root 101 Feb 11  2014 lcg-voms2.cern.ch.lsc
-rw-r--r--. 1 root root  97 Feb 11  2014 voms2.cern.ch.lsc
------------------------------------------------------------------------
[root@htc-ce ~]# ll /etc/grid-security/vomsdir/lhcb/
total 8
-rw-r--r--. 1 root root 101 Feb 11  2014 lcg-voms2.cern.ch.lsc
-rw-r--r--. 1 root root  97 Feb 11  2014 voms2.cern.ch.lsc
------------------------------------------------------------------------
[root@htc-ce ~]# ll /etc/grid-security/vomsdir/dteam/
total 4
-rw-r--r--. 1 root root 129 Jan 19  2017 voms2.hellasgrid.gr.lsc
------------------------------------------------------------------------
```

Ensure the CRLs are up to date:

```
------------------------------------------------------------------------
[root@htc-ce ~]# yum install fetch-crl
```

```
[...]
-------------------------------------------------------------------------
[root@htc-ce ~]# systemctl enable fetch-crl-cron
-------------------------------------------------------------------------
[root@htc-ce ~]# systemctl start fetch-crl-cron
-------------------------------------------------------------------------
[root@htc-ce ~]# fetch-crl > /tmp/crl-$$.log 2>&1 < /dev/null &
-------------------------------------------------------------------------
```

Set up the HTCondor CE following these steps:

https://htcondor.com/htcondor-ce/v5/installation/htcondor-ce/

WARNING: there are the following **additional steps** before the `condor-ce` service can run successfully.
Also check the configuration file examples below.

- Copy the pool password and its derived token:

```
-------------------------------------------------------------------------
cp -i /etc/condor/passwords.d/POOL /etc/condor-ce/passwords.d/
-------------------------------------------------------------------------
cp -i /etc/condor/tokens.d/* /etc/condor-ce/tokens.d/
-------------------------------------------------------------------------
```

- Open the HTCondor CE port:

```
-------------------------------------------------------------------------
firewall-cmd --permanent --zone=public --add-port=9619/tcp
-------------------------------------------------------------------------
firewall-cmd --reload
-------------------------------------------------------------------------
```

The HTCondor CE ***daemon*** configuration should resemble the following:

```
-------------------------------------------------------------------------
[root@htc-ce ~]# ll /etc/condor-ce/config.d/
total 24
-rw-r--r--. 1 root root 1321 May 18 18:14 01-ce-auth.conf
-rw-r--r--. 1 root root 1714 Dec 21 22:11 01-ce-router.conf
-rw-r--r--. 1 root root 1362 Dec 21 22:11 01-pilot-env.conf
-rw-r--r--. 1 root root 1444 Dec 21 22:11 02-ce-condor.conf
-rw-r--r--. 1 root root  500 Dec 21 22:11 03-managed-fork.conf
-rw-r--r--. 1 root root   52 May 18 18:15 50-schedd2.conf
-------------------------------------------------------------------------
[root@htc-ce ~]# grep ^AUTH /etc/condor-ce/config.d/01-ce-auth.conf
AUTH_SSL_SERVER_CERTFILE = /etc/grid-security/hostcert.pem
AUTH_SSL_SERVER_KEYFILE = /etc/grid-security/hostkey.pem
AUTH_SSL_SERVER_CADIR = /etc/grid-security/certificates
AUTH_SSL_SERVER_CAFILE =
AUTH_SSL_CLIENT_CERTFILE = /etc/grid-security/hostcert.pem
AUTH_SSL_CLIENT_KEYFILE = /etc/grid-security/hostkey.pem
AUTH_SSL_CLIENT_CADIR = /etc/grid-security/certificates
AUTH_SSL_CLIENT_CAFILE =
-------------------------------------------------------------------------
[root@htc-ce ~]# cat /etc/condor-ce/config.d/50-schedd2.conf
JOB_ROUTER_SCHEDD2_POOL = htc-cm.your-domain:9618
-------------------------------------------------------------------------
```

Option #1: mappings only use the HTCondor mapfile

Such a setup is the simplest, but has usability limitations by design:

```
-------------------------------------------------------------------------
[root@htc-ce ~]# ll /etc/condor-ce/mapfiles.d/
total 16
-rw-r--r--. 1 root root 1305 Dec 21 22:11 10-gsi.conf
-rw-r--r--. 1 root root 1095 Dec 21 22:11 10-scitokens.conf
-rw-r--r--. 1 root root   78 May 18 17:53 11-gsi.conf
-rw-r--r--. 1 root root   99 May 21 17:31 11-scitokens.conf
-rw-r--r--. 1 root root  540 May 18 17:49 50-gsi-callout.conf
-------------------------------------------------------------------------
[root@htc-ce ~]# cat /etc/condor-ce/mapfiles.d/11-gsi.conf
GSI /.*,\/alice\/Role=lcgadmin/ alicesgm
GSI /.*,\/alice\/Role=NULL/ alice001
-------------------------------------------------------------------------
[root@htc-ce ~]# cat /etc/condor-ce/mapfiles.d/11-scitokens.conf
SCITOKENS /^https:\/\/wlcg\.cloud\.cnaf\.infn\.it\/,8c3c01a9-ee96-4f6e-989c-ad1e279244ae$/
-------------------------------------------------------------------------
[root@htc-ce ~]# grep GSI /etc/condor-ce/mapfiles.d/50-gsi-callout.conf | tail -n 1
#GSI /(.*)/ GSS_ASSIST_GRIDMAP
-------------------------------------------------------------------------
```

NOTE: the `GSS_ASSIST_GRIDMAP` line must be **commented out or removed** !

```
-------------------------------------------------------------------------
[root@htc-ce ~]# tail -n 3 /etc/passwd
alicesgm:x:19984:19984:alicesgm:/tmp:/bin/bash
alice001:x:19985:19985:alice001:/tmp:/bin/bash
wlcg001:x:19986:19986:wlcg001:/tmp:/bin/bash
-------------------------------------------------------------------------
```

NOTE: unless `slot` accounts are used on the WN, it is important to set the home directory of each grid account to a value that will be ***mapped into*** the job directory on the WN, as explained in the preceding section.

Option #2: mappings used to be done via LCMAPS

This ***legacy*** method is no longer supported in the **Feature Channel** and hence should no longer be considered for HTCondor CE installations.
A similar machinery for flexible mappings of `SciTokens` has been discussed, but is not yet available for the time being.

This topic: LCG > MiniHTCsetup
Topic revision: r11 - 2023-03-03 - MaartenLitmaath