CONDORS IN CONTAINERS

DENNIS VAN DOK

Nikhef Jamboree 2024-05-13

A FAREWELL TO TORQUE

For over 20 years, we have been running stoomboot with a batch system called Torque.

It's been a trusty companion through these many years.

We've come to know its quirks; its limitations; and its moods.

The end of the road has been reached. It is time to say goodbye. You could say the system is on life support, but in truth, there is no support.

Nikhef

But without Torque, what is next?

FLOCKING TO HTCONDOR

- The HTCondor system is a product that is developed and supported by a team from the University of Wisconsin-Madison
- It's free, open source software
- Nikhef has a history with it going back 30 years!

TRANSITION IN PROGRESS

- We are installing more capacity to the new cluster at this very moment
- Expect an invitation to join the new cluster in your Inbox soon
- If you want to try it out early, drop a note at stbc-admin@nikhef.nl.

WHAT HAPPENS TO THE OLD CLUSTER?

- We will keep the Torque cluster going for a while to smooth the transition.
- Expect everybody switches before July

WHAT HAPPENS TO CENTOS 7

- Support for CentOS 7 runs out 30 June 2024
- We won't be able to offer any CentOS 7 capacity after that date
- You could use a CentOS 7 container, but we cannot guarantee that container images for CentOS7 will remain available.

SUBMIT SCRIPTS

Old:

```
$ cat job.sh
#!/bin/bash
#
#PBS -q long

./run.sh

$
```

New:

```
$ cat runjob.sub
executable   = run.sh
log          = run.log
output       = outfile.txt
error        = errors.txt
+UseOS       = "el9"
+JobCategory = "long"
queue
$
```

RUNNING JOBS

Old:

```
$ qsub job.sh
17953664.burrell.nikhef.nl
$
```

New:

```
$ condor_submit runjob.sub
Submitting job(s).
1 job(s) submitted to cluster 556.
$
```

OTHER TOOLS

Old:

- qsub
- qstat
- qdel

New:

- condor_submit
- condor_q
- condor_rm

CONDOR IN CONTAINERS

- All jobs **will** be run in a container
- This is done transparently
- shares (/home, /project, etc.) are available as normal
- You may select either the base OS or bring your own container image

SELECTING A BASIC OS

- Put

  ```
  +UseOS: "el9"
  ```

  in your submission script.
- This selects a base container image
- currently allowed values are:

  ```
  {"el7", "el8", "el9"}
  ```

  Which select a default image compatible with Red Hat Enterprise Linux versions 7, 8, and 9 respectively.

SELECTING YOUR OWN CONTAINER IMAGE

Alternatively, you may choose any other container image. Or create your own.

```
+SingularityImage = "docker://ubuntu:20"
```

SELECTING A CATEGORY

- We like to tune priorities of jobs based on their characteristics
- short jobs should be scheduled ahead of long jobs
- jobs from users who run many jobs are lowered in priority

SETTING THE CATEGORY

- Specify

  ```
  +JobCategory = "short"
  ```

  in you job script
- Allowed values are "short", "medium", "long" with different default and maximum run times.
- Specify desired Maximum wall clock time with

  ```
  +MaxWallTime = 96 * 3600
  ```

  for the maximum 96 hours.

CONTAINERS, CONTAINERS EVERYWHERE…

What are these containers anyway?

SIMILAR TO AN OPERATING SYSTEM...

A container image presents the file system as it would on a particular operating system:

- files
- libraries
- configuration

…BUT DIFFERENT

- The kernel is that of the host OS
- containers provide isolation from one another
- Same host can run multiple containers simultaneously

BUILDING YOUR OWN CONTAINER IMAGES

Several recipies and tools are available.

- https://buildah.io/
- https://linuxcontainers.org/distrobuilder/introduction/

See https://kb.nikhef.nl/ct/Containers.html

WHERE TO GET HELP

- Mattermost: stbc-users
- Office Hours every 1st Thursday of the Month (next: 6 June)
- mail stbc-admin@nikhef.nl
- Knnowledge base: https://kb.nikhef.nl/ct/Stoomboot-NG_HTCondor_Cluster.html
- HTCondor Manual: https://htcondor.readthedocs.io/en/latest/index.html

QUESTIONS?