




HTCONDOR AT NIKHEF



HISTORY OF CONDOR AT NIKHEF



Future Generation Computer Systems

Volume 12, Issue 1, May 1996, Pages 53-65



Paper

A worldwide flock of Condors: Load sharing among workstation clusters

D.H.J. Epema ^a, M. Livny ^b, R. van Dantzig ^c, X. Evers ^{a, c}, J. Pruyne ^b

^a Department of Mathematics and Computer Science, Delft University of Technology, PO Box 356, 2600 AJ Delft, Netherlands

^b Department of Computer Sciences, University of Wisconsin—Madison, Madison, WI, USA

^c National Institute for Nuclear Physics and High-Energy Physics Research (NIKHEF), PO Box 41882, 1009 DB Amsterdam, Netherlands

2001
“VOORTUIN” CLUSTER
RUNNING CONDOR,
WHILE WE WERE WAITING
FOR THE FIRST EUROPEAN
DATA GRID RELEASE

TORQUE AND MAUI

Served us well for more than 15 years. But.

- Scheduler (Maui) dead for 7 years. Amazing that it still works as well as it does!
- Workload keeps getting more diverse:
 - Large range in run times
 - Increasing diversity in core count requests
 - Number of semi-dedicated clusters increasing
- Torque support not so great
- Lots of new cool stuff in Condor. Not so much new cool stuff in Torque.

THE PUSH OVER THE CLIFF

Nikhef moves into gravitational wave physics

- Workflows usually built around HTCondor and DAGman
- Gravwav group had their own funds to buy a cluster and wanted HTCondor on it
- Ligo VIRGO Collaboration aiming at a solution that works on both sides of the Atlantic

SO: Let's give them a cluster

Note: related to work by Mary Hester @ Nikhef and OSG friends on Condor for IGWN (parent of LVC) - talk submitted to fall 2020 HEPIX workshop

THE PLAN

Pragmatic approach

- The gravwav cluster nodes are standard Grid worker nodes at the moment
- Kidnap one and put minicondor on it; evaluate
- Kidnap two more and make a small cluster; evaluate
- When all working, move the whole thing to Condor



SUCCESSFUL LIGO JOB

PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
41986	templon	20	0	13556	5372	1064	S	100.0	0.0	15559:54	factorize
3787855	templon	20	0	874068	658464	12308	R	100.0	0.1	0:54.46	BayesWave
19558	condor	20	0	120876	11376	7064	S	1.0	0.0	302:39.14	condor_startd
3788787	templon	20	0	172864	2972	1612	R	1.0	0.0	0:00.13	top

```
[templon@wn-lot-037]GW150914_GWOSC_localframes> condor_q -nobatch

-- Schedd: wn-lot-037.farm.nikhef.nl : <127.0.0.1:9618?... @ 09/22/20 10:41:14
ID      OWNER      SUBMITTED      RUN_TIME ST PRI SIZE CMD
  1.42  templon      9/11 12:16    10+21:56:54 R  0   10.0 fact.condor 54
  20.0  templon      9/22 10:39     0+00:01:11 R  0    0.3 condor_dagman -p 0 -f -l . -Lockfile GW150914_GWOSC_conda/GW150
  21.0  templon      9/22 10:40     0+00:01:04 R  0   733. BayesWave --checkpoint --ifo H1 --ifo L1 --psdlength 4.0 --H1-d

Total for query: 2 jobs; 0 completed, 0 removed, 0 idle, 2 running, 0 held, 0 suspended
Total for templon: 2 jobs; 0 completed, 0 removed, 0 idle, 2 running, 0 held, 0 suspended
Total for all users: 2 jobs; 0 completed, 0 removed, 0 idle, 2 running, 0 held, 0 suspended

[templon@wn-lot-037]GW150914_GWOSC_localframes> █
```

/user/templon/bayeswave-use-cases-v1/GW150914_GWOSC_localframes | templon

NOTES SO FAR

- Installation of minicondor was **really simple**.
Something like 10 lines typed at a command prompt.
- I would complain about the command outputs, but it's probably just me being used to how Torque does it
- Why isn't `condor_q -nobatch` the default?

THE FUTURE (SUBTITLE: WE HATE EMPTY CYCLES)

Once all gravwav nodes are under HTCondor

- How to keep it busy ... options
 1. Allow it to take IGWN work from outside (not clear if this can work)
 2. Allow other groups @ Nikhef to use idle cycles

We expect another such cluster (theory group) to appear this year:

- Clone the gravwav setup on it?
- Workload sharing based on some policies?

What if we migrate our general local cluster from Torque to HTCondor?

COLLABORATION

The big challenge at Nikhef is scheduling
Several levels of groups and shares to deal with; strong
point of Maui, weak point of HTCondor in the past

Active discussion with HTCondor team to come up with
policies than can work here. With this in hand,
HTCondor on the grid systems becomes a possibility

Looking forward to the work, discoveries, and happy
users!