Managed by Fermi Research Alliance, LLC for the U.S. Department of Energy Office of Science

HTCondor Deployment at Fermilab

Where we came from, where we're going

Steven Timm

European HTCondor Site Admins Meeting

9 December 2014

Fermilab Pre-condor

- Fermilab has run Farms-based reconstruction, large numbers of independent processors since late 1980's and before. (Vax, custom hardware, RISC-based)
- "In search of Clusters" (2000) lists us as example of high-throughput, embarassingly parallel computing
- Used CPS, FBS, and FBSNG, all written at Fermilab*
- 2002—2 years into Tevatron Run II.
  - FBSNG working well on reconstruction farms
  - Experiments started building Analysis Linux clusters
  - Fermi management didn't want to extend scope of FBSNG
  - D0 cluster "CAB" started using PBS,
  - CDF "CAF" started with FBSNG but were already investigating Condor.

Condor at Fermilab—CDF Central Analysis Facility

- First quasi-interactive analysis facility
- Analysis jobs ran on batch system but users had capability to
  - Tail a log file
  - Attach a debugger if necessary
  - Have files copied back to their private area
- These features developed first on FBSNG batch system and then transferred to Condor in 2004.
- Condor developers added Kerberos 5 authentication to Condor at our request
- Given success of Condor on CAF, CDF reconstruction farms were also converted to run on Condor.

FermiGrid (General Purpose) and Open Science Grid

- FBSNG needed grid extensions for X.509 support and for bigger scalability
- Instead--transitioned reconstruction farms to Condor
- In 2005 began with 28 general purpose CPU on condor, accessible by grid, transitioned the balance by end of 2006.
- CMS Tier 1 also transitioned to Condor, a bit earlier.

GlideCAF/GlideinWMS

- CDF users liked local CAF extras
  - Wanted to run the same on the grid
  - Result was “GlideCAF”—renamed a couple years later to “GlideinWMS”.
- Condor glide in:
  - Central system handles the submission of grid pilot jobs to the remote site.
  - These jobs start their own condor_startd and call home to the CDF condor server
  - To users, all resources appear to be in the local CDF condor pool just as before.
  - No applying for personal certs, no grid-proxy-init, etc, all transparent to the user
- From 2009 onwards the neutrino experiments of the Intensity Frontier also use glideinWMS.
- New frontend is an http-based client/server called jobsub
- Users don’t run condor_q or condor_submit anymore.
- .

Current state of HTCondor@Fermilab

- 4 main grid clusters part of FermiGrid
  - CMS Tier 1--11056 slots—HTCondor
  - General Purpose Grid—9376 slots—HTCondor
  - CDF Grid—3488 slots—HTCondor
  - D0 Grid—4376 slots—Torque(PBS) Maui
- CDF and D0 Grid clusters to be decommissioned shortly.
- Grid clusters by nature quite scalable
  - Slow Globus middleware in front of condor,
  - Stress on schedd is almost non-existent.
  - Have grown from 28 cores to 28550.
- Challenges are on submission nodes:
  - What happens when someone submits 200K jobs and deletes them right away.

Directions we are going:

- CMS Tier 1 and GP Grid just put under management of same department (after 8 years of being separate)
- Looking at everything from the ground up. 2 different philosophies and setups, take the best from both.
- Three key technologies:
- HTCondor-CE
  - Gives us a lot more flexibility on slot requirements from grid
  - Uses the “Job Router” features of HTCondor
- Partitionable Slots
  - Address diverse memory, core, disk requirements
- Hierarchical Quotas
  - Unified way to balance analysis, production, calibration, test

Hierarchical quotas

- 3 top level groups:
- BATCH: quota of 99% of all slots
- HIGHPRIO: quota of 1% of all slots
- OPPORTUNISTIC: no quota, can be pre-empted
- GROUP_QUOTA_group_batch=9200
- GROUP_QUOTA_group_batch.nova=1300
- GROUP_QUOTA_group_batch.nova.production=300
- GROUP_QUOTA_group_batch.nova.analysis=1000
- GROUP_ACCEPT_SURPLUS = True ** BIG shift here
- Add default shorter job length 24-36 hrs (was 4.5 days)
- Eventually hope to let experiments control the split within their own subquota but that will take some work

Partitionable Slots

- Have been mentioned in other talks
- Way to address very diverse memory, disk, and core requirements. Astronomers have wildest ones.
- 64GB RAM per job but just 1-2 cores. Huge scratch too.
- Have been experimenting with partitionable slots on FermiCloud. Expect to deploy on GP Grid in a couple of months
- Have to revise any monitoring to recognize that the number of slots in the pool can dynamically change.
- Shift to reliance on cgroups to enforce the memory limits of the partitionable slots.

Cloud Bursting

- We have already run experimental workflows at the 1000-virtual machine level on Amazon Web Services and on FermiCloud.
- Can burst using GlideinWMS or via OpenNebula native cloudbursting features. Most of the time use GlideinWMS.
- These virtual machines do not join the base FermiGrid clusters but are visible to the submit nodes.
- Also have done demos of launching custom worker nodes (same as a real worker node in all respects except for different RAM per slot) and having them join FermiGrid. This we refer to as “Grid Bursting”.
  - Used this to meet a Dark Energy Survey milestone before we had partitionable slots working.

The Virtual Facility Project

- Fermilab's long-term goal is to have the facility provision nodes on commercial clouds on behalf of the various experiments.
- Construct a virtual facility where not only compute nodes but services can exist transparently in Fermilab or in the commercial clouds or both.
  - Successful demos of scalable squid servers and scalable submit nodes (schedd) in the cloud thus far
  - Next big one is data caching both inbound and outbound.
  - Leverage load balancing and autoscaling functions of the cloud where they exist.
  - Expand to other clouds, in particular Google and Azure as well as OpenStack.

Conclusions

- Fermilab very thankful for all the help we get from HTCondor
- Couldn’t do what we do without it, either on grid or on cloud.
- Expect HTCondor will remain critical technology going forward
- Always glad to compare experiences, war stories, etc.