

A Decade of Condor at Fermilab

Steven Timm timmm@fnal.gov

Fermilab Grid & Cloud Computing Dept.

Work supported by the U.S. Department of Energy under contract No. DE-AC02-07CH11359

Outline

- Fermilab pre-Condor
- What is Condor
- CDF CAF
- FermiGrid and OSG
- Condor-G as Grid Client
- FermiGrid Site Gateway
- GlideCAF
- GlideinWMS
- CMS Tier 1 and LPC
- Condor->HTCondor

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, embarrassingly 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.

What is Condor

- The Swiss Army Knife of High-throughput computing
- Developed at University of Wisconsin Computer Science Dept.—Prof. Miron Livny
- Began by sharing desktop cycles on CS dept workstations
- Now a full batch system++
- Supported on all imaginable platforms
 - (Windows, Mac, Linux, Unix, IBM Blue Gene, and many more)
- Now available in Red Hat and other Linux distros
- Significant industrial use in financials, aerospace, insurance, entertainment, more.

Some Condor terminology

- “pool” –a collection of nodes running Condor. Each one runs a condor_startd
- “collector” —The daemon that collects all the information from the pool
- “schedd” —The daemon or daemons which takes user job requests
- “negotiator” —Matches user jobs and available machines
- “classad” —the format by which Condor describes machine and job resources
- “slot” —one logical unit for job execution, can be partitioned to any number of cores on the node.

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 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.

Condor-G as Grid Client

- In early 2000's Condor added Condor-G
- Essential for dealing with Globus "GT2" toolkit resources, one jobmanager per user instead of one per job.
- Only supported client on Open Science Grid
- Supports a variety of Grid resources now (Unicore, gLite, ARC/Nordugrid, all flavors of Globus)
- Plus direct submission to other batch systems without grid (PBS, LSF)
- Also now supports Virtual Machine submission to local clusters, Amazon EC2, OpenNebula, and others.

FermiGrid Site Gateway

- At beginning of Grid era, Fermilab management said 'Build a unified site gateway'
 - We used Condor-G Matchmaking
 - Building on experience of D0 SAMGrid
 - Each cluster sends a classad of how many job slots it has free per VO. (using GLUE 1.3 schema)
 - Job is matched to the cluster with free slots and then resubmitted via condor-G to that cluster.
 - If it doesn't start executing within 2 hours we pull it back and resubmit it to a different cluster.
 - Open Science Grid uses similar technology in Resource Selection Service, written and operated at FNAL.
 - Now 4 main clusters: Condor:(CMS, CDF, Gen.) PBS(D0)

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.
- “CDFGrid” glide-in to clusters on the site of Fermilab for data handling jobs
- “NAMGrid” glide in to clusters on the OSG and Pacific rim for Monte Carlo
- INFN CAF glide in to gLite/WLCG (using gLiteWMS)

GlideinWMS

- Now known as the GlideinWMS, project headed at Fermilab
- Used by the majority of big Open Science Grid VO's
- Also by Intensity Frontier experiments at Fermilab.
- This is the one main technology that got the majority of our users to use the Grid.
- Works on the cloud too—submit a virtual machine with a client configuration that calls home to glideinWMS
- Production OSG GlideinWMS hosted at Indiana Univ. GOC and at UCSD.
- Contributions from Fermilab, UCSD.

CMS Tier 1 and LPC

- CMS Tier 1 at Fermilab—early adopter of Condor
- Separate LPC is local non-grid “tier 3” cluster for users of the LHC Physics Center at Fermilab

Condor added features @ Fermilab request

- Condor authentication
- X.509 authentication
- Separate execution partitions per slot
- Partitionable slots
- Integrated support for gLexec
- VOMS support / support for ext. callouts
- Several types of cloud support
- Extensions to quota system.
- And many many more.

Scalability issues

- Condor_schedd was and is single-threaded
- Use case of a few schedulers driving a large cluster was new to condor.
- Start rates have improved over 2 orders of magnitude since we have been working with condor
- Routine now to schedule 30K simultaneous jobs
- Goal to get that to 150K (equivalent to all CMS Tier 1+Tier 2 resources in the world).
- And then double that to burst to the cloud.

Current improvement directions

- Working on memory footprint. How can you schedule 100K jobs from single machine?
- Partitionable slots—already available now but improving the scheduling features to better schedule whole nodes.
- Packaging—RPM's compliant to Fedora standards, in collaboration with RedHat, more dependent on system libraries.
- (Main condor rpm from 140MB->10MB in last 2 major releases).

Condor ->HTCondor

- In next few months, the package will be renamed to “HTCondor”
- “HT” stands for High Throughput, after the Center for High Throughput Computing at Univ of Wisconsin.

Conclusions

- Condor has served Fermilab and FermiGrid well for a decade now
- Choice of most US-based Tier 1 and 2.
- Growth of WLCG computing will continue to push developers
- A stable, mature batch system that is vital to accomplishing our work.
- Developers have been very helpful in adding the features we need.

References

- <http://www.cs.wisc.edu/condor>
- Condor project
- <http://fermigrid.fnal.gov>
- FermiGrid home page