

Experience at the Indiana T3

(some of talk borrowed from previous talk at the OSG all-hands meeting)

Frederick Luehring

luehring@indiana.edu

Indiana University

Tier3(g,w) meeting at ANL ASC

May 18, 2009

What the IU Tier 3 Is...

- The IU Tier 3 is a physics department cluster shared between dzero and ATLAS (the “da” cluster).
 - Mixture of ~10 servers with ~20 cores
 - No two servers have the same hardware
 - The servers were acquired over ~5 years one or two at a time
 - ~10 “Excessed” racks of 1.6 GHz dual CPU srvrs from FNAL
 - Shared among groups and cannibalized to make usable servers
 - 2 RAID arrays totaling about 12 TB
 - 1 Gbps connectivity
 - SLC4 / Rocks with PBS queue
 - wLCG and DQ2 client software but **NO** grid middleware
 - All recent versions of Athena
 - Managed by a university paid sysadmin
 - 50 miles from the closest part of MWT2

How I Have Been Working

- It is possible to work efficiently:
 - Subscribe your dataset to the Tier 2
 - If the subscription is not acted on, email Kaushik.
 - Use dq2-get to get the sample data onto the Tier 3
 - dq2-get from any US site is fast.
 - dq2-get from CERN is slow.
 - Develop/test your software package on the Tier 3
 - Run your Monte Carlo locally on the Tier 3
 - Then use pathena to send hundreds to thousands of data analysis jobs to a Tier 2 or BNL (wherever the data is)
 - Have to be careful not to waste cycles which is real easy to do because pathena makes it simple to launch thousands of jobs.
 - Ran against both RAW and ESD data.
 - Process DP3Ds and ntuples on the Tier 3 or laptop.
 - It is also possible to process AODs the same way.

IU's Experience

- We previously sent the grad students and post docs to the experimental sites (Opal/D0/ATLAS)
- The faculty and scientists in Indiana worked mainly on smaller local systems usually using ntuples.
 - Significant amounts of Monte Carlo also got run at IU
- Currently our local people work in many ways:
 - Hal Evans (faculty) works mostly at Dzero
 - Vivek (scientist) works at BNL
 - I (scientist) split my work between IU and MWT2
 - Harold (faculty) uses ROOT/ntuples locally
 - John Penwell (grad student) works at CERN
 - Denver Whittington (grad student) uses pathena, CERN, IU
 - Yi Yang (grad student) works IU and pathena

Our Planning

- So our existing Tier 3 is not a Tier 2-like site.
 - So it's basically a Tier 3w though it predates that jargon
- In writing our 3 year DOE renewal grant we realized that we have infrastructure suitable for a Tier 3g:
 - We have a full time, experienced system manager
 - The University is spending ~\$200k to renovate the physics dept. computer room to have sufficient cooling and power for ~24 racks (but the space is shared with other groups).
- So we are trying to decide how to proceed
 - It's safe to say that none of us wants to get into running a full scale Tier 2-like site with a full set of grid server middleware.
 - Most of our previous experience on Opal and Dzero (after VAX era) had been on a much smaller scale.
 - We may aim a little higher than we previously considered.

Issues That Came Up

- The Conditions Databases (CORAL + COOL) proved to be a relatively complicated to deal with when processing real cosmic data (NOT Monte Carlo).
- In the end it turns out that:
 - Some of the conditions data can be read from Oracle at BNL
 - To read at a reasonable speed you need a Squid cache locally and to use the FroNTier/Squid setup at BNL
 - Some of the conditions data is too large to find in Oracle
 - So you need local copies of hundreds of .pool.root files and to construct an xml file saying where these files can be found locally.
- My (incomplete) testing of filling TRT histograms from the RAW and ESD files indicates that if you get things setup this way it all works.

Final Thoughts

- Rik pointed out the need to configure the software tools at the Tier 3s following some sort of “standard” so that support is possible.
 - This is a very good idea.
 - It does seem unavoidable to have a very diverse set of hardware configurations.
- Do not overlook the possibility of using “opportunistic” resources: department and university systems.