

<https://indico.cern.ch/conferenceDisplay.py?confId=254672>

News

- Welcome Julian! Please introduce yourself to the group!

Attending

- Jen, SeangChan, Luis, John - FNAL
- Edgar, Andrew, Julian - CERN
- I will be away from early afternoon (EU Time) and discussed directly with Sara the handover. Xavier.

Issues last week

- issues with too many jobs hitting the system at once again
 - ◆ Edgar is working on this with Dirk
 - ◆ What can we all do to help speed this process along?
 - ◆ we have too many pending jobs combined at all sites,
 - ◆ we can add a hard limit
 - ◇ we need to pre-order the jobs by priority
 - ◇ if we put a hard limit in it how do we handle log collect/merge etc?
 - ◇ we are saturating network and the agents are having problems with that many pending jobs *in the meantime we just use current limits to control the chaos
- The switch to EOS at FNAL is over, Test workflows went through with only minor issues that have been fixed so we are back in business.
- Still having issues with couch Databases.
 - ◆ most are stable but 201 is still going down
 - ◆ Seangchan has a patch if we filter for all successful jobs
 - ◇ is it safe to put patch in and test yet?
 - ◆ Seangchan will give Alan and Edgar patch they will put it on a preproduction agent. Alan & Edgar will validate it and report back asap.
 - ◆ if the database is getting to big we can patch for now
 - ◆ we have a long time requirement is that we to be able to track every job, Seangchan is arguing that this is unsustainable long term because the databases will just get too big. We need to really think through, what we can do to fulfil this requirement without having a database that gets so big it is unusable/unsustainable. Is there a point that we we can get rid of this data?
- Issues with couch and Late Binding are causing workflows to take a very long time to go through and making debugging difficult.
- git -
 - ◆ Is git installed on the production machines? Seangchan, Luis and I couldn't find it
 - ◆ Who needs to install this centrally? FNAL and CERN machines
 - ◆ what is the command we need to replace:
 - ◇ `svn co svn+ssh://svn.cern.ch/repos/WmAgentScripts ~/WmAgentScripts`
 - ◇ git clone needs to be used so you need to install new scripts on lxplus machines to do your updates
 - ◆ Seangchan was wondering why we are all keeping our "own version" of the WmAgentScripts in our own home areas instead of all working out of the cmst1 directory. Can we revisit this topic?

Personel

- Sep 10 --> Sep 17 Xavier

- Sep 17 --> Sep 24 Sara

Site Issues

- Several sites not accepting jobs: T2_BE_IHE, T2_BE_UCL, T2_US_CAltech and T2_RU_IHEP (<https://savannah.cern.ch/support/?139787>), Florida (<https://savannah.cern.ch/support/?139782>), Beijing (<https://savannah.cern.ch/support/?139773>)
 - ◆ clearly issue is on glidin side John will look into this further
- Nebraska is still in drain until they answer ticket about missing release/CVMFS problems (<https://savannah.cern.ch/support/?139781>)

Sites for Production

Agents

- still having issues with couch and stability of agents

Workflows

RELVAL

IEEE Paper

Draft Outline #1

- Introduction (Why we need to run so much simulations, why we need to do a reconstruction of the data) (Edgar/Jen)
- a brief discussion of what the different types of workflows are, and how they are processed differently (Diego/Jen/Edgar)
- monitoring for T1 & T2 sites(Diego/Jen/Edgar)
- How we ran prior to 2011
 - ◆ ProdAgent vs WMAgent (Diego/Alan) (Focus on differences and improvements)
 - ◆ Reprocessing and Production (Jen/Xavier) (How this was handled with ProdAgent and why the need to move to another framework)
 - ◆ How we ran with WMAgent (after 2011)
- WMAgent /ReqMgr/Workqueue (Diego/Edgar/Alan) General comment on how it works * PREP/ReqmG Interaction (Vincenzo?) * Organization of the workflow team and operations around it (Edgar)
- Achievements
- Events reconstructed (L3s)
- Usage of the grid (Edgar/Jen/L3s)
- Conclusions / Outlook (Edgar/Jen)

Action Items

- Write twiki disk/tape separation T1_IT_CNAF. Edgar
- Recovery workflows - Jen - suspend
 - ◆ first 2 workflows are completely through and now we are waiting for people to really look and make sure that there are no show stoppers before we do the other 50.
 - ◆ Guillemo is bothering JeanRoc about if people have actually looked at the data
- A new state for completed and already dealt with ACDC.

- How many workflows running, pending, waiting, stuck
 - ◆ Is it documented yet? yes
 - ◆ Luis is working on a script to pull these numbers automatically. - script done but we are still tweeking it
- solve the problem of how to use a non-production scram architecture (waiting for Alan to come back)
- Updating documentation on scripts with github now that we aren't using svn anymore
 - ◆ docuention needs to be updated and everyone needs to start ramping up on github

AOB

- Edgar will be with us through Oct 31st
- Diego will continue to work in the paper
- problem with the creation of WF's if you change the number of files per job
- CouchDB will be rotated on Wed so we will be running without being able to watch WMStats on Wed/Thurs should be back Friday

This topic: CMSPublic > WorkflowTeamMeeting20130917

Topic revision: r3 - 2013-09-17 - JenniferAdelmanMcCarthy



Copyright &© 2008-2021 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

or Ideas, requests, problems regarding TWiki? use Discourse or Send feedback