

1/8

Aaaaand we're back for Jordan's final semester! A few notes:

1. I can hear a hard drive dead on one of the nas's, but I lost high bay access over the break, so I'll have to re-check that once I can get back in tomorrow. A good test for my instructions on rebuilding a hard drive though, so yay?
2. The link for the rsv page from the diagnostics page was bad, but that was a simple fix, just had to edit `/var/www/html/diagnostics/index.php`.
3. `/var` kept filling up quickly, seems one culprit was gratia logging, in `/var/log/gratia` the daily logs were huge. I zipped a few up and that helped, but that needs more than a band-aid.
4. The GUMS web GUI seems to be down, although GUMS still appears to be working (grid tests are still coming in). Going to look into that.
5. SUM tests are still failing, need to talk with Bockjoo perhaps? Will look into it.
6. BeStMan still isn't communicating, that's going to near the top of my list.
7. Given my personal hard drive tanking, had to pull my cert/key off the cluster and re-combine to a .p12 using openssl. From there, re-loaded it into firefox so I can access MyOSG and OIM stuff again.
8. Never finished documentation over the break, MUST do that and load it to the wiki.
9. Gratia is also misbehaving, need to fix that.

1/12

Rebuilt the hard drive on nas1 w/Christian, so that's all squared away, just need to file an RMA. Also need to post the process to the wiki.

1/14

Got a response back from Bockjoo concerning the SUM tests, he said just look at the dashboard page and try going through to fix the "bad" ones one-by-one. As a reference, the SUM tests are located at:

[http://dashb-cms-sum.cern.ch/dashboard/request.py/latestresultssmry-sum#profile=CMS_CRITICAL_FULL&group=AllGroups&site\[\]=T3_US_FIT&flavour\[\]=All+Service+Flavours&metric\[\]=All+Metrics&status\[\]=All+Exit+Status](http://dashb-cms-sum.cern.ch/dashboard/request.py/latestresultssmry-sum#profile=CMS_CRITICAL_FULL&group=AllGroups&site[]=T3_US_FIT&flavour[]=All+Service+Flavours&metric[]=All+Metrics&status[]=All+Exit+Status)

Interesting note: `condor_drain` didn't work on compute-2-8, although `condor_off` worked perfectly. May need to look into that, though the node was already empty... Regardless, using 2-8 as a testbed for a yum update. Shut down condor, running update now. Includes a kernel update, will need to restart...

Restart successful, although new condor version required re-editing `/etc/condor/config.d/00personal_condor.config` and re-adding the initial lines, such that the local condor pulls configs from nas0:

Where is your local config file?

LOCAL_CONFIG_FILE = /mnt/nas0/condor/condor-etc/condor_config.cluster

After edit and a `condor_restart`, the node came up fine, ran jobs perfectly.

Interesting: using a simple `condor_drain` (no graceful option) it runs properly, at least it did on 2-3 (which was also empty). Perhaps the graceful flag is the issue? Nay, tested on compute-1-3 as well and it worked fine. Perhaps 2-8 is just a bit screwy. Also failed on 1-1, 1-2, 1-3, 1-4, 2-0, 2-1, 2-2, 2-3, 2-5, 2-7, and 2-9. Both with and without the graceful flag.

Also, ntp created an `ntp.conf.rpmnew`, looking into it. Many new options, mv'ing 2-3's `ntp.conf.rpmnew` to `ntp.conf`, we'll see how the restart works.

Created 2 new groups, testbeam and htau for their respective directories in `/var/www/html`, for further wiki editing.

1/16/14

Interesting find: tried to "`condor_off -peaceful`" a few nodes, but got the following message:

Can't find address for schedd compute-1-2.local

Perhaps you need to query another pool.

Turns out that this can be due to incorrectly setting `SCHEDD_NAME` and/or `SCHEDD_ADDRESS_FILE` in a personal condor config. The former isn't defined, but the latter is for a location that only exists on the CE:

`/var/lib/condor/spool/schedd_address`

Hmmm, moving on to other things for now. So, the gratia-condor probe is having issues. To run the single probe specifically, need line:

```
rsv-control --run --host uscms1.fltech-grid3.fit.edu org.osg.gratia.condor
```

Possible that `/etc/gratia/condor/ProbeConfig` was set up incorrectly, I moved the old `ProbeConfig` to `ProbeConfig.old`, and copied the settings from the `ProbeConfig.rpm` into `ProbeConfig` instead. We shall see how it goes.

The shift caused the probe to throw the following error:

```
ERROR: Probe condor:uscms1.fltech-grid3.fit.edu has not been seen by Gratia collector  
gratia-osg-prod.opensciencegrid.org:80 in >1d.  
Last contact time = 2013-02-22 02:09:58 EST
```

Going to leave it as-is for now, perhaps the gratia server on the other end just needs time to pick us back up. Hey look at that, a few minutes later, everything seems to be sending!

Just ran a yum update on the SE, turns out the osg repo (`/etc/yum/repos.d/osg.repo`) got updated and written as `osg.repo.rpmnew`. Checked it and then moved it to `osg.repo`, ran a yum clean all and then another yum update, and got about 70 new globus updates, on top of other things. Note that this will happen for the CE as well, most likely. `Lcmaps.db` also got changed to `lcmaps.db.rpm`, but wasn't replaced, is it not necessary anymore?

Re-requesting http and bestman certs for the SE, previous requests didn't include slashes, that certainly isn't helping. Re-issued the `httpcert`, now for the `bestman cert`. Done.

Tried running through install processes. `Fetch-crl3` had some bugs, kind of expected. However, `gratia-xrootd-transfer` and `gratia-xrootd-storage` did not register as recognized services. May have to completely re-install `osg-se-bestman`, to get those linked properly. For the full list of requireds by `osg-se-bestman`:

```
[root@uscms1-se ~]# rpm -q --requires osg-se-bestman
```

```
osg-version  
osg-system-profiler  
bestman2-server  
bestman2-client  
bestman2-tester  
edg-mkgridmap  
fetch-crl3  
globus-gridftp-server-progs  
vo-client  
grid-certificates  
gratia-probe-gridftp-transfer  
gums-client  
liblcas_lcmaps_gt4_mapping.so.0()(64bit)  
rpmlib(PayloadFilesHavePrefix) <= 4.0-1  
rpmlib(CompressedFileNames) <= 3.0.4-1
```

Hmmm, some of these packages (in red) are missing, very interesting. I'm even more inclined to do a fresh re-install.

1/17/13

Installing crab 2.9.1:

- wget http://cmsdoc.cern.ch/cms/ccs/wm/www/Crab/Docs/CRAB_2_9_1.tgz
- unpacked `/mnt/nas0/OSG/APP/crab`, also `cmssoft`?
- unpacked new `tgz` here as well
- ran a `./configure`, configured perfectly
- to run CRAB
 - need to install `remoteglidein` or `glite` so users can source the file
 - source the CMS env

-source the crab.(c)sh file

Exported environment variable VO_CMS_SW_DIR in /etc/profile for users, to access both cvmfs, and for CRAB.
Also added environment variable CRAB_DIR in /etc/profile for users to access CRAB more easily.
Will add a comment in /etc/profile on necessary sources for user to use CRAB

Note: CRAB makes a job and sends it through remoteglidein or glite to Condor. Also, no need to “install” remoteGlidein, it's just a setting in /mnt/nas0/OSG/APP/crab/CRAB_#_#_#/python/crab.cfg.

1/22/14

Been working with Ankit a lot, getting certain softwares properly configured. Bockjoo also mentioned that the crontabs for the cmssoft and uscms01 accounts were obsolete (since we now have CVMFS up and running), so we have removed them.

So, currently, CMSSW can be properly sourced (again, with CVMFS), and CRAB is running smoothly. The next big hurdle on our plate is getting remoteGlidein to work. This software acts as the go-between for CRAB and Condor; CRAB creates a grid job, and sends it to remoteGlidein. RemoteGlidein then handles the job, sending it to an appropriate cluster.

The current struggle is that remoteGlidein doesn't see FLTECH, even when told to send jobs only to our cluster. There could be multiple reasons for this, the first of which is that apparently the FLTECH CE needs to be registered in the “glidein factory” for analysis. Not sure what that is, but hopefully Ankit will know?

The other issue could be that remoteGlidein isn't set up properly to communicate with Condor. I fear this one is likely, regardless of the state of the first. I'm also starting to doubt our Condor configuration in the first place. I'm going to be looking more deeply into that one this week.

**Side note: as we have our repo priorities for OSG packages to be at the top, an interesting observation is that OSG doesn't offer an 8.x version of Condor. They must be intentionally holding back, which is good for us, as we're still running 7.8.8 (most current according to OSG).

Just found a note in the Install CE page on the grid twiki; turns out we need to set variable DELEGATE_JOBS_GSI_CREDENTIALS = False. Probably requires a condor_restart on the CE, but pzu0 is running 100 jobs as usual, and we still don't have checkpointing. Changing the variable anyway, in /etc/condor/condor_config on the CE (the line was commented out initially, as a note).

Also, Bockjoo called to our attention that although we installed CVMFS, we didn't migrate it properly. Followed his steps 2-5 here: <http://oo.ihepa.ufl.edu:8080/cvmfs/migration-fivesteps.html>. This included adding the line:

```
export CMS_LOCAL_SITE=/cvmfs/cms.cern.ch/SITECONF/T3_US_FIT
```

To the end of our /etc/cvmfs/default.local file.

Also cleaning up mail, cron is reporting way too many errors a day:

- weekly cron to clean up ganglia archives (and only save the last 3) was written poorly, need to edit to tail -n +4
- moveOldFiles.sh in /var/www/html/tripwire was written poorly, dir was defined wrong (would have run from the root directory otherwise, good thing it was failing elsewhere). Also added zipping and moving said files.

1/23/14

Continuing cleaning up mail:

- the tripwire cron logs are...intimidating. Turns out tripwire checks a large number of file integrities, etc. It fails many, the question is whether tripwire needs to be updated, or if these are serious issues. I'm assuming the former.
- most other mails are CRL errors that we can't fix, or are daily Tripwire/Logwatch reports. Removed most, also cleaned out mbox, it was getting pretty big.

Squid is (and was originally) set up to run on the CE, but the site-config-file.xml was pointing it to the SE. Looks like we're going to move squid to the SE. Installed both squid and CVMFS on the SE, having trouble getting squid to start up. According to /var/log/squid/cache.log, it cannot open an HTTP port. Had to edit the /etc/osg/config.d/01-squid.conf file(s?), also /etc/squid/squid.conf, /etc/sysconfig/iptables was already open

to allow traffic.

1/24/14

Needed to modify the NET_LOCAL variable in /etc/squid/customize.sh, then run service frontier-squid reload, in order for the proper networks to be allowed (mainly our LAN for the compute nodes). After that modification, as well as removing write permission to /etc/squid/squid.conf (it was freaking out otherwise), our SUM tests for squid now clear. However, I believe only one squid "port" (or whatever it may actually be called) is up and running, which means all traffic will still run through a single channel. Need to add a few more channels (maybe up to 8?) at some point, but will leave it be for now.

Most SUM test errors left are now BeStMan-related. Making that a top-tier priority now, along with PhEDEx. GIP may be related, so I'm starting there. Turns out a few lines need to be added to /etc/gip/gip.conf, see the GIP conf grid twiki page.

In /etc/osg/config.d/30-gip.ini, advertise_gsiftp is set to true, but I don't think we have it set up properly. Changed to false temporarily, then ran an osg-configure -c to propagate changes.

SE also didn't have the globus_tcp_source_range in iptables, added it.

1/25/14

Attempting to install PhEDEx. They recommend, with any upgrade, that you actually re-install rather than attempt such an upgrade. The previous PhEDEx version we had was 3.1.3, we are now required to install something at 4.1.0 or higher.

Following install docs from here: <https://twiki.cern.ch/twiki/bin/view/CMSPublic/PhedexAdminDocsInstallation>

Switched into phedex account (just used su phedex), changed to home directory, and then began at the pre-installation section of the documentation. It's rather fancy, installs a light version of apt built by CMS, then uses that to install the PhEDEx packages (not using YUM here, so no need for sudo).

***Note - needed to install zsh and compat-libstdc++-33 for install script to run

Successfully installed PhEDEx v.4.1.3

1/26/14

Ankit is registering with OIM, need to get Christian to do that as well.

1/28/14

Got a ticket from the glidein folks yesterday, seems that some glideins have failed verification, explicitly on compute-1-6, compute-2-0, and compute-2-5. Did some digging, and the package "myproxy" was missing from all 3 machines. Installed it on each, now grid-proxy-init works properly. Checked all others with a cluster-fork rpm -qa | grep myproxy, everyone else seems fine.

Also working on moving the cvmfs caches over to individual nodes. We had it all set to /mnt/nas0, but that may be slowing it down considerably. Going to set each cache to 7 Gb. Editing /etc/cvmfs/default.local, changing the cache base and quota limit, from "/mnt/nas0/cvmfs" and "20000" to "/var/cvmfs" and "7000" respectively, after running mkdir cvmfs in /var.

Also; Grid admins need to register both through OIM by another admin, as well as request gridadmin acces on the grid admin page in OIM.

condor_drain doesn't work for compute-2-7, failed send to startd (local). Skipped reconfig of cvmfs.

```
process: condor_drain -graceful compute-##; mkdir /var/cvmfs; vi /etc/cmvfs/default.local (change); double-check condor, then kill processes; service autofs reload; ls /cvmfs/cms.cern.ch; condor_restart compute-##
```

1/30/14

Changed the gridftp host location in /etc/osg/config.d/30-rsv.ini (on the CE), to point to the gridftp server now on the SE. Ran an osg-configure to push change, also restarted the SE to bring it up cleanly. Hopefully this will push the RSV probe over to the SE appropriately, and then BeStMan can dig through correctly. At least the gridftp server is pointed correctly now.

Still doesn't explain the xrootd issue though, need to double-check that once the SE has finished restarting.

Still aren't there, going to install each package individually, crossing fingers that everything else is either there, or any other dependencies will be brought along:

Installed:

gratia-probe-xrootd-storage.noarch 0:1.13.18-1.osg.el5
gratia-probe-xrootd-transfer.noarch 0:1.13.18-1.osg.el5

Dependency Installed:

gratia-probe-services.noarch 0:1.13.18-1.osg.el5

Probes came up ok after enabling them, by setting the EnableProbe variable in /etc/gratia/xrootd-transfer/ProbeConfig and /etc/gratia/xrootd-storage/ProbeConfig to 1. Also had to move the gridftp-simple metric from the CE to the SE, by changing a line in 30-gratia.ini on the CE from gridftp:uscms1.... to gridftp:uscms1-se....

2/3/14

Herp derp, there are 46 Gb free under the scratch directory. Moving the CVMFS cache there now.

All nodes done, including 2-7. Only did a service autofs reload though, that way any current jobs can finish in /var/cvmfs, and then only the new ones will kick over to /scratch/cvmfs (hopefully). Will check back in a few hours to verify that's what's actually happening.

Now to work on the gridftp probe. Need to get at least one probe running properly on the SE, to verify that RSV is connecting properly.

2/9/14

Not appreciating our current condor setup, it seems helter-skelter, and could use some streamlining. We will, of course, only test one node at a time for this. But we also want to reserve a few nodes for local jobs only. Looking into setup now. Also, if possible, need to set up checkpointing again.

Going to be experimenting on the management node, as we don't want it running jobs anyway, technically.

So: condor runs locally, but reads /etc/condor/config.d/00personal_condor.config. From there, we direct condor to /mnt/nas0/condor/condor-etc/condor_config.cluster, which sets a bunch of values. It then directs the node to LOCAL_CONFIG_FILE=/back to same directory/condor_config.\$(HOSTNAME), so it hits a link in the same directory. The nodes then pick up condor_config.compute, to turn on MASTER and STARTD. The CE is told to run MASTER, COLLECTOR, NEGOTIATOR, and SCHEDD, as well as is given the local network to talk on.

By default, HTCCondor now reads from /etc/condor/config.d, where 00personal lives. To simplify, we could push a single file onto all compute nodes, holding all of these variables. This removes nas0 dependency, but benefits of spreading it out? Ponder ponder ponder. Also, for limited nodes, we can deny users by setting DENY_READ and DENY_WRITE to osg, glow, etc.

Condor recommends first using the global config file located at variable CONDOR_CONFIG (as condor searches there first). As per the condor-issued global config file, condor is automatically told to search LOCAL_CONFIG_DIR, which is by default set to /etc/condor/config.d. That is where 00personal comes in, but every condor update pushes a new 00personal, which wipes our current line for "LOCAL_CONFIG_FILE", pointing at nas0.

2 questions then; 1) do we keep the same structure, or keep all parameters on each node? 2) do we keep editing 00personal every time, or make a new file in /etc/condor/config.d/ to see if it stays put, and therefore we never have to worry about an upgrade? Would have been useful had we thought ahead to NOT upgrade one of the damned nodes...

Also, concerning the schedd issue;

Extra Info: If the condor_schedd is running on the machine you are trying to query and you still see the error, the most likely cause is that you have setup a personal Condor, you have not defined SCHEDD_NAME in your condor_config file, and something is wrong with your SCHEDD_ADDRESS_FILE setting. You must define either or both of those settings in your config file, or you must use the -name option to condor_q. Please see the Condor

manual for details on SCHEDD_NAME and SCHEDD_ADDRESS_FILE.

So, mayhaps a look back into those two is in order.

2/10/14

Christian sent off an email to the T3 list concerning BeStMan, so that's a start.

I am still looking into SCHEDD_NAME, we don't define it in our condor config files.

2/11/14

Just started playing with condor on the management node. Gave it a SCHEDD_NAME in 00personal, and restarted. Turns out it takes jobs immediately, but now the trick is to see if I can get it to shutdown peacefully. NOPE, same error.

2/13/14

Changing localPathListAllowed in /etc/bestman2/conf/bestman2.rc to match location of SUM test, from /tmp to /bestman/BeStMan/cms/store/unmerged/SAM/testSRM/SAM-uscms1-se.fltech-grid3.fit.edu/
Also added write permissions a few directories up, and added bestman to the uscms01 group. Waiting for SUM test to come back up.

2/20/14

Saw a path in SUM test #10 that "wasn't allowed" in the bestman2.rc file (localPathListAllowed), bumped that variable back up to the SAM directory. Also, for directory permissions, it was under osgusers, but bestman was not in that group, so added that in /etc/group. Restarted bestman, we shall see if that goes green in SUM.

2/25/14

Doug is looking into the GridFTP issue. Noticed that in our IP tables, we had 2 declarations of the 2811 port. Unlikely, but possibly causing the conflict? Not sure, so removed the second entry (they were identical anyway).

3/10/14

Matt got Ubuntu installed on Electron just fine, then got the vidyo software loaded. Christian and I then made it to a T3 vidyo, was much lower attendance than expected. Attendees:

- Marguerite at UMD
- Juan (always helping)
- Dave?
- Rob came late
- 2 others whose names I can't remember

Marguerite had some interesting points on their setup, but didn't gain any serious knowledge, except that dave falls asleep on-camera.

Dr. Raggozine's new server is set up, can be ssh'ed into at pluto.fit.edu, ask Doug for an account. Also installed in the left rack, sits in between the network switch and nas1.

3/17/14

Bugged Doug about BeStMan, he doesn't even know anymore. He has a script to strip a machine of all OSG software, so we're going to try that (and then reinstall Bestman and GridFTP), but if that doesn't work then he recommends wiping/reinstalling CentOS.

Asked Silicon Mechanics about the nas1 hard drive issue, they also recommend wiping the current drive, as we have to set up a new software RAID and that usually toasts the drive anyway. The secondary data structure should be fine, but it'll be about a day or two of downtime when that happens.

Got numbers for how much power Pluto draws from the two UPS downstairs, added the configuration and draws to the power grid spreadsheet on the wiki.

3/24/14

Still haven't heard back from Doug concerning stripping scripts for the SE, will bug him today.

Looks like someone disabled VO's at least temporarily, as jobs dropped to 40, but nobody actually said anything...looks like a fluke.

3/31/14

Emptied out the CE a little, was getting too full, just removed a few log files.

Spoke with Doug about the SE, he's actually taking point on stripping/reinstalling the software

Got replacement hard drives from Silicon Mechanics, only used 1 for the rebuild, but 2 others do need to be replaced, will do that during the downtime.

Set Matt loose on an openMPI research bit, as Mike Phipps was asking about it, and figured Matt could use something to do.

4/14/14

Replaced the other two nas1 drives, sent back to silicon mechanics

Doug just responded, he's still working on the SE, he's been swamped with other stuff.

4/28/14

cvmfs got a bit wonky after an update with yum, had to go in and edit the default.local file. It shouldn't be touched, but the cache location got changed from /scratch/cvmfs to /root/cvmfs, but the latter doesn't have the proper permissions so it couldn't mount correctly. Should be ok now, but will email around to see if it got fixed.

5/2/14

Switching over srm_dir in /etc/osg/config.d/30-rsv.ini from DEFAULT to /bestman/BeStMan/cms/vo...something something, as per the advice of Doug Johnson, which will make our final RSV probe go green. The problem is that we have the then run osg-configure for the changes to be written, but are getting an error with our Gratia setup which won't allow osg-configure to run fully. The error is:

2014-05-03 00:37:51,996 DEBUG GratiaConfiguration.checkAttributes started

2014-05-03 00:37:52,014 CRITICAL Unknown exception encountered while running: [Errno 2] No such file or directory

So supposedly there's a path error in the 30-gratia.ini file. However, digging through our gratia setup in /etc/gratia/metric, everything seems to check out ok. Ponder ponder ponder.

5/5/14

To-do list:

- password list for Dr. Hohlmann
- power down instructions for door of cluster and whiteboard in lab

/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=jordan/CN=746428/CN=Jordan Alexander Robertson