

How-To Cluster: A Trouble-Shooting Guide

Part 1: Things You Will Do Regularly

- I. Introduction
- II. Need to Know
 - a. Certificates
 - b. Twiki
 - c. Websites
 - d. Email Lists
- III. Updates
 - a. CRAB
- IV. Phedex
 - a. How to Use(online)
 - b. How to Use(command line)
 - c. Resources
 - d. Past Issues
- V. Crontab
- VI. Printer
- VII. Diagnostic Scripts

Part 2: Solutions to Known and Resolved Problems

- VIII. Hardware Issues
 - a. NAS
 - b. RAID
 - c. SE
 - d. Frontend
- IX. Software Issues
- X. Online Metrics Issues
 - a. Ganglia
 - b. CMS Dashboard
 - c. RSV
- XI. Other Issues
- XII. Absolute Basics
 - a. Commands You Need
 - b. Syntax
 - c. Test Job

Part 1: Things You Will Do Regularly

- I. This guide is to serve as a launch pad for all new administrators of the Florida Tech Tier 3 Cluster. The majority of this guide is dedicated to clearly explaining solutions to specific problems that have been encountered by previous administrators. As a new administrator it is part of your responsibilities to update this guide. Note, unlike your daily logs this file should only contain solutions that worked and should be written so clearly that someone with very little linux experience could understand the steps. The first section ought to serve as a good intro to any new apprentice to an administrator as well as serve as a reminder to graduating admins of the steps they must take before leaving to ensure the stability of the cluster without them.
- II. Need to Know
 - a. Certificates:
 - i. You will need a grid certificate to identify yourself as a trustworthy source to several websites such as <https://oim.grid.iu.edu/oim/home>. Very clear instructions are located on the twiki: uscms1.fltech-grid3.fit.edu/wiki.
 1. Pick CMS as your VO and OSG as your
 - ii. Have to renew your membership with VOs (virtual Organizations) too. Go to <https://lcg-voms.cern.ch:8443/vo/cms/vomrs?path=/RootNode&action=execute>
Then do >Member Info> Resign . Re-sign the grid and VO AUPs.
 - iii. Cluster Certificates need to be updated too.
 1. Phedex User on SE, see *Other Issues*.
 - a. Replace with a copy of your own. Then need to edit the proxyrenewal script:
`/sandbox/phedex/gridcert/proxyrenew.sh.forpatrick` to include your password. Should be obvious where to type it.
 2. SE CRLs: have to update them as root on the SE
`ssh username@uscms1.fltech-grid3.fit.edu`
`sudo su root`
`source /sandbox/osg-se/setup.sh`
`vdt-control --on fetch-crl`
This updates them and adds this command to the crontab so it will be run periodically. Or update them manually with:
`$VDT_LOCATION/fetch-crl/share/doc/fetch-crl-2.6.6/fetch-crl.cron`
 3. Worker Node CRLs, see below
 4. Host, http, and RSV certificates:
Find the twiki documentation
Must have OSG gridadmin access granted first so contact them

For Http certificates
`$ /opt/osg/cert-scripts/bin/cert-gridadmin --host uscms1.fltech-grid3.fit.edu --service http --email pford@my.fit.edu --affiliation osg -vo cms --prefix http-uscms1`

For RSV certificates

```
$ /opt/osg/cert-scripts/bin/cert-gridadmin --host uscms1.fltech-  
grid3.fit.edu --service rsv --email pford@my.fit.edu --affiliation osg --  
vo cms --prefix rsv-uscms1
```

For Host Certificates

```
$ /opt/osg/cert-scripts/bin/cert-gridadmin --host uscms1.fltech-  
grid3.fit.edu --email pford@my.fit.edu --affiliation osg --vo cms --  
prefix host-uscms1
```

Then make a copy of the host cert and rename as the containercert
Also note that both httpcert/key need to be owned by daemon!

- iv. Certificate Authority, at one point you will probably have to become the CA for the group as previous sysadmins graduate. At this point you will need to: Install mozilla thunderbird (or another email client that allows you to digitally sign emails) then followed the instructions here to import your Grid certificate into it.

https://twiki.grid.iu.edu/bin/view/Security/OsgRaOperations#Digitally_signed_email

<https://twiki.grid.iu.edu/bin/view/Security/ThunderbirdSmime>

Then sent an email to burt that was digitally signed and now hopefully he'll respond saying its all set up. the more official route is described in the first website above.

b. Twiki:

- i. As an admin, you are also responsible for updating the twiki with daily logs, reports, publications, and any other articles you want to write.
- ii. To update the twiki, files too be uploaded must be copied to the /var/www/html directory on the cluster. Then the user should go the wiki homepage, log in, edit their section and follow the syntax of previous users.

c. Websites:

- i. <https://oim.grid.iu.edu/oim/home>
Security Admin's Site. Must have a cert to view.
- ii. <http://uscms1.fltech-grid3.fit.edu>
Cluster Homepage
- iii. <http://uscms1.fltech-grid3.fit.edu/twiki>
Information about the Cluster. Keep Updated.
- iv. <http://uscms1.fltech-grid3.fit.edu/ganglia/>
Snapshot view of cluster performance. Check Daily.
- v. <http://dashb-cms-sam.cern.ch/dashboard/request.py/siteviewhome>
Snapshot view of cluster status. Check Daily. Must have a cert for full access.
- vi. <http://myosg.grid.iu.edu/about>
snapshot view of cluster status and performance. Check Daily. Great for graphs for posters/presentations.

d. Email Lists:

- i. Hypernews: hn-cms-osg-tier3@cern.ch
Instructions are on the twiki, under getting certificates, for joining this list.
Great resource for all cluster problems/issues. Very helpful and intelligent community.
- ii. Condor: condor-users.cs.wisc.edu
Join at <https://lists.cs.wisc.edu/mailman/listinfo/condor-users>.
Great resource for issues with condor.
- iii. Ganglia: ganglia-general.lists.sourceforge.net
Join at <https://lists.sourceforge.net/lists/listinfo/ganglia-general>. Great resource for all ganglia issues. Archives are located at http://sourceforge.net/mailarchive/forum.php?forum_name=ganglia-general
- iv. Rocks: npaci-rocks-discussion.sdsc.edu
Join at <https://lists.sdsc.edu/mailman/listinfo.cgi/npaci-rocks-discussion>
I have rarely used it but its still a good idea.

III. Updates

a. CRAB

- Download the tarball (a .tgz or so file)
- Copy it to /mnt/nas0/OSG/APP/crab
- Change the group and owner to uscms01
- Untar the tarball.

```
tar -xvf CRAB_x_x_x.tgz
```

- change to the new CRAB_x_x_x directory
- setup CRAB
- ./configure*

b. CMSSW

- Check current versions under the swinst SAM test
- Bockjoo from UF will automatically update it for all tier 3's.
- Ask him if he doesn't.

c. Kernel

- Basically only update it if there is a serious security problem. The OSG Security team will email you with details. Then just do...
- yum install kernel
- it automatically selects the right one, must run on CE and SE and then reboot.

d. Ganglia

- restart /etc/init.d/gmetad
- restart /etc/init.d/gmond

e. CVS

Records the history of your file, but only stores the differences between versions.

```
$cd SITECONF/T3_US_FIT
```

```
$export CVS_RSH=/usr/bin/ssh
```

```
$export CVSROOT=:ext:pford@cmsscvs.cern.ch:/cvs_server/repositories/CMSSW
```

```
$cvs commit
```

OR

```
$cvs add <file>
```

f. Condor

<https://twiki.grid.iu.edu/bin/viewauth/Tier3/CondorSharedInstall>

IV. Phedex (Physics Experiment Data exports)

Transfers data files, handles anything we import or export out of the cluster

a. How to Use (online)

- i. go to https://cmsweb.cern.ch/dbs_discovery/
- ii. check what datasets are already available at our site with:
find dataset where site like *FLTECH*
- iii. check what dataset is available beginning with a specific name
find dataset where dataset like /Mu/aeverett/*
- iv. in some instances you need to change where you search from global to _02 etc.
- v. When you see the dataset you want, click phedex and fill in the required forms.
Default values are good for most cases.
- vi. Monitor transaction at:
<https://cmsweb.cern.ch/phedex/prod/Data::Subscriptions>
- vii. On our cluster all requested and transferred datasets are stored at:
/mnt/nas0/OSG/BeStMan/cms/store/
- viii. Some datasets have to be passed through T2_US_Florida first so you request both T2_US_Florida and T3_US_FIT

b. How to Use (command line)

- i. When there is an issue with phedex it is good practice to try and test it using the manual srmcp command.
- ii. The syntax is `srmcp -2 -debug=true srm://<location of file>:8443/srm/v2/server?SFN=<file to transfer> file:///<destinationfile>`
- iii. Here is an example:

```
$ srmcp -2 -debug=true  
srm://srmb.ihepa.ufl.edu:8443/srm/v2/server?SFN=/cms/data/test/srmb_test.2  
file:///tmp/from_srmb.1
```
- iv. Note, could also do from an srm to an srm or from a file to an srm or from a file to a file. Files are always followed by `////` while srms have only two slashes.
- v. If transfers aren't working between sites, you can test a transfer from our SE to our SE. one example is:

```
$ srmcp -2 -debug=true srm://uscms1-se.fltech-grid3.fit.edu:8443/srm/v2/server/S/store/unmerged/SAM/testSRM/SAM-uscms1-se.fltech-grid3.fit.edu/lcg-util/pfordtestfile  
file:///home/xfave/from_uf.1
```
- vi. Also, to be able to run these commands must first create a proxy from your certificate.
`$grid-proxy-init`

c. Resources

- i. <http://cmsweb.cern.ch/phedex/prod/Reports::DailyReport> ctrl-F to see our cluster's phedex activity for the past month

d. Past Issues

Problem 1: Phedex failing to complete transfers between UF and us

Solution 1: They had to remap us to an account, so the issue was on their side this time.

V. Crontab

- a. Crontab is a file on linux machines that allows the user to run certain scripts periodically without having to remember to do them. Certificate renewal scripts are good example. The phedex certificate proxy on the SE is renewed every week because it has an entry on the crontab. However the user can specify any type of periodicity: every second Monday for example.
- b. For commands and syntax do *man crontab* . Google for a crontab calculator if you can't determine the syntax for the periodicity on your own.

VI. Printer

- a. Access printer configuration with
telnet 163.118.206.154 from the cluster
Type in password.
type "?" for HELP, or "/" for current settings
quit to exit.
Only 10 computers may have access to the printer at any time. Choose Wisely.

VII. Diagnostic Scripts

- a. On the frontend there is a script at */usr/local/bin/DiagnosticsCE* that when ran prints out the expiration dates of the various certificates as well as the current status of nfs, condor, and httpd. To run it simply type *DiagnosticsCE* at the command prompt from any location on the cluster.
- b. On the SE there is a script at */usr/local/bin/DiagnosticsSE* that when ran prints out the expiration dates of the various certificates as well as the current status of squid, the expiration date of the phedex certificate and the phedex proxy. To run it simply type *DiagnosticsSE* at the command prompt from any location on the SE. Because it goes into private directories for squid and Phedex it must be run with *sudo* or as root.

PART 2: Solutions to Known Problems

VIII. Hardware Issues

- a. NAS: Network Attached Storage, our machine that is connected to the outside. Home directories are mounted on the NAS and the majority of our storage space is on it. Read Patrick's Installation Documentation for more.

Problem 1: How to install second NAS and set it up.

Solution 1: First had to upgrade new Nas to centos 5.6 with a cd. Then use insert-ethers to make it visible to the frontend and to generate its kickstart file. Instructions are on the ROCKS home page for adding new nodes: <http://www.rocksclusters.org/rocks-documentation/4.2/install-compute-nodes.html>

Then it was created and visible to the frontend, we named it nas-0-1

After that step we partitioned the harddrive on nas1 to suit our needs. We created a 12 TB logical volume for backing up user data and a 46 TB logical volume for grid use such as phedex datasets. Specific steps are as follows:

Created Sdc1 and sdd1, each 30 TB

```
#parted sdc
#mklabel gpt
#mkpart primary 0 100%
#quit
```

```
#parted sdd
#mklabel gpt
#mkpart primary 0 100%
#quit
```

then format both to be xfs

```
#mkfs.xfs sdc
#mkfs.xfs sdd
```

Then followed instructions located at <http://linuxhelp.blogspot.com/2005/04/creating-lvm-in-linux.html>

for creating logical volumes. exact steps are below

```
#pvcreate /dev/sdc1 /dev/sdd1
#pvscan
#vgscan
#vgcreate -s 1024M vol_grp_60 /dev/sdc1 /dev/sdd1
#vgdisplay
#lvcreate -L 48 -n log_vol_grid vol_grp_60
#lvcreate -L 12 -n log_vol_backup vol_grp_60
#lvdisplay
#mkfs.xfs /dev/vol_grp_60/log_vol_grid
#mkfs.xfs /dev/vol_grp_60/log_vol_backup
#mount /dev/vol_grp_60/log_vol_grid /nas1
#mount /dev/vol_grp_60/log_vol_grid /backup
```

Now both volumes exist on the Nas1 so we needed to mount them to the frontend.
On Nas1 edit /etc/exports to be able to present both volumes to the cluster when it asks.

```
#vi /etc/exports
/nas1 10.0.0.0/255.0.0.0 (rw, no_root_squash,async)
/backup 10.0.0.0/255.0.0.0(rw,no_root_squash,async)
```

Then start or restart nfs
/etc/init.d/nfs start

On Frontend (uscms1) mount the two directories
mount nas-0-1:/nas1 /mnt/nas1
mount nas-0-1:/backup /mnt/backup

Edit /etc/fstab to do it automatically upon a reboot
#vi /etc/fstab
nas-0-1:/nas1 /mnt/nas1 nfs defaults 0 0
nas-0-1:/backup /mnt/backup nfs defaults 0 0

and thats it!

Problem 2: Our Nas01 was beeping loudly. One of the hard drives had a blinking red light on it.

Solution 2: One of the hard drives had failed. Hence the red light. Took it out and put one of our spares into the corresponding slot then rebooted. Enter the RAID Configuration screen by pressing C repeatedly during reboot. To turn off the noise, just had to click on the icon for volume and tell it to turn off. The light turned off by it self. Then inside the RAID configuration screen needed to tell the node to configure the new drive. The first screen that appears asks you to import configuration. Do it. It will copy data from the other Drives to your new drive, a process that takes approximately a whole day. As soon as you see it get to 1% completed you can leave it to run on its own. When you return the following day, reboot the machine and everything should be fine.

- b. RAID: Redundant Array of Inexpensive Disks, technology that allows computer users to achieve high levels of storage reliability from low-cost and less reliable PC_class disk-drive componets, via the technique of arranging the devices into arrays for redundancy.
- c. SE: Our Storage Element. Home to BeStMan (the Berkely Storage Manager)
- d. Frontend:

Problem 2: Cluster was down. No one could ssh in. In highbay the frontend was on, so logged on directly but could not ssh to nodes, nas, or SE. After rebooting everything manually (holding down the power off button and then turning it back on after a couple minutes) could ssh to nodes and nas but still not to SE.

Solution 2: A switch in the building, not one of ours, had stopped working so both the frontend and the SE (which are directly connected via ethernet cords to the outside world) could not be ssh'd to. IT went and reset the switch and everything was fine.

Note: used command traceroute from CE to my computer, olin.fit.edu, and from my computer to the CE to get a feel for where the problem was. This shows the individual steps it takes to connect (i.e.) where exactly the problem is.

Problem 3: On GPL's personal NAS, two drives had failed and were red. They only had a RAID5 configuration so all of the data was lost.

Solution 3: First repaired the drives by connecting an Ethernet cable from the machine to my laptop and going to the website: 10.255.255.250:8080/cgi-bin/index.cgi#

The beginning part of this is the machine's own ip-address. At this website there was a visual for everything you can do with the raid controller and the harddrives. We ran two disk scans and then two disk repair scans. At the end of this process which took over a day to complete the two drives had been repaired and the machine was once again all green. However the data was not recoverable.

e. Power Supplies

Problem 1: The UPS or any power strip won't turn on any more after having drawn too much current from it.

Solution 1: The fuse of the outlet on the wall to which its plugged in is most likely the problem. Try plugging it into a different outlet on a different box. If this works then you just need to either replace the fuse or have facilities flip the breakers back. And everything will work fine.

IX. Software Issues

a. Condor

Problem 1: Plenty of jobs in queue but some machines never run them.

-On Ganglia they're load display is virtually none.

-They are alive: proved with ping. *Ping compute-1-1*

-Condor is not running its daemons on them. *condor_status*

-Further checked by running following on worker node: *ps -ef | egrep condor_*

Solution 1: On those nodes that weren't running jobs the /scratch partition was 100 percent full.

Delete the files produced from failed jobs.

```
$ ssh compute-1-1
```

```
$ df -h
```

```
/dev/sda3      193G   866M  182G   100% /scratch
```

```
$ cd /scratch/wn_tmp
```

```
$ sudo rm -rf cms_*
```

```
$ /sbin/service condor start
```

```
$ exit
```

```
$ condor_status
```

Problem 2: One node is not running condor. Hence no jobs are running on it.

Solution 2: Had to remount the scratch directory as readwrite. Note it appeared as rewrite from mount command but not under cat /etc/fstab. Had to edit fstab to be explicit and then remount it

ssh to the compute node causing problems.

```
$ vi /etc/fstab
```

```
LABEL=/scratch /scratch ext3 rw,errors=remount-ro 0 0
```

```
$ umount /dev/sda3 /scratch #unmounts the directory
```

```
$ mount -a #mount the directory with the write flag
```

```
$ /sbin/service condor start
```

b. BeStMan/PheDeX

Problem 1: Wanted to reconfigure our BeStMan setup to be able to save datasets on the new NAS1. Currently if a user requests this location their CRAB or grid jobs fail to run.

Solution 1: Needed to edit the bestman.rc file's variable for LocalPathsAllowed and then restart bestman.

```
#ssh dev-0-0
#vi /sandbox/bestman-1.2.13/vdt-app-data/bestman/conf
localPathListAllowed=/mnt/nas1/bestman;/bestman/BeStMan
#source /sandbox/bestman-1.2.13/setup.sh
#vdt-control --off
#vdt-control --on
```

More properly you could also edit this file by running the configure_bestman script as seen below. But make sure you first save a backup copy of the bestman.rc file.

NOTE: I created the directory /mnt/nas1/bestman and gave it the same permissions as the /bestman/BeStMan directory which phedex was originally pointed to. Also ALWAYS make backups of any files you change so you can put them back to the way they were when you're done.

Within the same file it was important for AccessFileSysViaSudo=True

c. Squid

Problem 1: Barry emailed us and told us According to the SAM tests, your CMS Frontier squid is working normally.

However, about a day ago, it stopped responding to the SNMP monitor.

Solution 1: Check IPtables and reboot, ours was missing two entries so I added them as below and then the final results are also shown. Repeat on CE.

```
[root@uscms1-se ~]# iptables -A INPUT -p udp -s 128.142.202.212/255.255.255.255 --dport 3401 -j ACCEPT
[root@uscms1-se ~]# iptables -A INPUT -p udp -s 131.225.240.232/255.255.255.255 --dport 3401 -j ACCEPT
[root@uscms1-se ~]# iptables -L
```

```
ACCEPT  udp -- cmsdbsfrontier.cern.ch anywhere      udp dpt:filecast
ACCEPT  udp -- fcdfdbfrontier4.fnal.gov anywhere      udp dpt:filecast
```

IX. Online Metrics Issues

a. Ganglia: <http://uscms1.fltech-grid3.fit.edu/ganglia/>

Problem 1: Rocks Addons aren't working

Solution 1: Run both of the following commands on the frontend. The first one restarts greceptor on all the nodes, the second one restarts it on the frontend. Note this only fixed

```
$ tentakel -g compute 'service greceptor restart'
$ service greceptor restart
```

b. CMS: SAM Tests: (These get rerun every 30 minutes but if they can't find a place in the queue they time out and try again in another 30 minutes. They tell us what is or is not working, tell us if CRAB jobs will or will not run and if Phedex will have issues.)

Squid=squid server, Ana=Analysis, front=frontier, bas=basic, jsub=job ubmission, swinst=cms stuff, mc=monte carlo, getPFN=checks pfn and lfn mapping, put=put a file on the srm server, del=delete a file, get=download a file from the server.

Problem 1: Swinst test gives notice that we're missing a CMSSW version or patch

Solution 1: Bockjoo at UF will take care of it. He is in charge of updating these versions on all Tier 3's. If he doesn't send him an email and he'll happily add the missing one to our configuration.

Problem 2: Failing mc, sft-job, and prod. Two give error message cannot download testjob.sh

Solution 2: Had to update the CRL certificates on the worker nodes.

```
$cd /  
$ln -s /mnt/nas0 nas0  
$source /mnt/nas0/OSG/GRID/setup.sh      (sourcing the VDT environment for the WNs)  
$vdt-ca-manage setupca --location mnt/nas0/OSG/GRID/globus/share/certificates --url osg  
$vdt-control --on fetch-crl              (Added the renewal of these certificates to the  
                                         Crontab)
```

Problem 3: Failing CE_MC and the SE-srmv2, both give a certificate related error.

Solution 3:

>Update host and webservice certificates.

Can check status of host certificate with...

openssl x509 -in /etc/grid-security/hostcert.pem -subject -issuer -dates -noout

Same for webcert but with their names in instead.

Instructions for updating them at...

<https://twiki.grid.iu.edu/bin/view/ReleaseDocumentation/GetGridCertificates>

SAM tests now giving a new (unknown Error) and we are failing RSV SE tests as well.

>looking at patricks documentation for certificates and updating accordingly.

(OSG Installation documentation from 2008, find on Wiki)

config.ini file at /sandbox/backups/configfiles/system-config/Config/OSG

restarted phedex with

```
$ /sandbox/phedex/PHEDEX/Utilities/Master --config  
SITECONF/T3_US_FIT/PhEDEx/Config.Prod start
```

>Turn bestman on

```
export PATH=$PATH:$VDT_LOCATION/bestman/bin
```

go to VDT_LOCATION on the SE (/sandbox/osg-se)

```
. setup.sh      (command to initialize the vdt environment)
```

```
vdt-control --off
```

```
vdt-control --on
```

errors: enabling inetd service gsiftp... FAILED! (see logs/vdt-control.log)

found conflicting, non-VDT entry for service fsiftp in /etc/service

> a new sam test started failing. running the bestman_config created a config file with bestman in FULL mode instead of GATEWAY which is what was created last time.

>Did a brand new bestman install using a script from doug command was...

~djohnson/InstallBestman /sandbox/bestman-1.2.13 (ran on SE)

To make it our primary install

```
read post-install/README
```

Turn off old installation:

```
cd /sandbox/osg-se
```

```
source setup.sh
```

```
vdt-control --force --off
```

Set up new install:

```
cd /sandbox/bestman-1.2.13/post-install
cp gsi-authz.conf prima-authz.conf /etc/grid-security
cd /sandbox/bestman-1.2.13
source setup.sh
vdt-control --force --on
```

Run a test to see if its working: test RSV Probes

```
su rsvuser
source /opt/osg-1.2/setup.sh
cd ~
./RunAllRSVProbes
```

Fixed both the Sam and RSV problems

Problem 4: Failing sam tests mc, cms-prod, and sft-job. Error Report says that the job was held and the proxy expired. Last night our Cluster Memory was also filled to 100%

Solution 4: Need to renew the CA CRLs on the worker nodes.

```
cd /
ln -s /mnt/nas0 nas0
source /mnt/nas0/OSG/GRID/setup.sh
vdt-ca-manage setupca --location /mnt/nas0/OSG/GRID/globus/share/certificates -url
osg
```

and then maybe:

```
vdt-control --on fetch-crl
```

because

```
uscms1> source /mnt/nas0/OSG/GRID/setup.sh
uscms> source /opt/sg-1.2/setup.sh
```

is the wrong VDT installation for the worker nodes.

Problem 5: SE failing Sam tests because of Authentication issues.

Solution 5: Set up two symbolic links to help with the certificates

```
$sudo ln -s /mnt/nas0/OSG/GRID/globus/share/certificates-1.13/mnt/nas0/OSG
/GRID/globus/share/certificates
$sudo ln -s /mnt/nas0/OSG/GRID/globus/share/certificates /mnt/nas0/OSG
/GRID/globus/TRUSTED_CA
```

Problem 6: two Sam tests failing (mc and another one) as well as RSV tests on SE.

Solution 6: Basically whenever have an issue with the SE it's a good idea to refresh the vdt package as a first resort. In this case only had to do :

```
$ ssh xfave@uscms1-se.fltech-grid3.fit.edu
$ source /sandbox/bestman-1.2/setup.sh
$ vdt-control off
$ vdt-control on
```

Problem 7: Two SAM tests failing both related to srmcp and mapping errors. RSV-se Probes also failing. The internet was really erratic in OPS the day that we began having issues.

Solution 7: One of the CRLs(7b54708e) on the SE had become corrupt. Had to basically trick it into thinking it didn't have an issue anymore.

```
# cd /etc/grid-security/certificates # mkdir 7b54708e # mv 7b54708e* 7b54708e/ #
/etc/init.d/bestman stop # /etc/init.d/bestman start
```

Then can test bestman with the following command, gives you same error as the SAM test would.

```
# lcg-ls -b -D srmv2 -l 'srm://uscms1-se.fltech-  
grid3.fit.edu:8443/srm/v2/server?SFN=/bestman/BeStMan/cms/store/unmerged'
```

Problem 8: SAM tests failing (ce-mc, srmv2-lcg-cp, and ce-frontier has a note on it)

Solution 8: An updated version of the CRLs had been created by OSG. Had to put these on both the CE and SE in place of current ones. Created a backup first of the old ones just in case. Email from OSG had website from which to download the tarball of new CRLs. Then copied tarball to SE and CE. Unpackaged it and copied contents to official folder for certificates (/etc/grid-security/certificates) then changed permissions.

```
#source <$VDT_LOCATION> (On CE: /opt/osg-1.2/setup.sh
```

```
On SE: /sandbox/bestman-1.2.13/setup.sh)
```

```
#cd /etc/grid-security/
```

```
#mkdir backupCerts2011
```

```
#cp /etc/grid-security/certificates/* /etc/grid-security/backupCerts2011/
```

```
#rm -rf /etc/grid-security/certificates/*
```

```
#mkdir newcerts2011
```

```
#cd newcerts2011
```

```
#cp <tarball> .
```

```
#tar -xvf <tarballname>
```

```
#cp /etc/grid-security/newcerts2011/certificates/* /etc/grid-security/certificates
```

```
#chown root /etc/grid-security/certificates/*
```

```
#chgrp root /etc/grid-security/certificates/*
```

```
#chmod ugo+rx /etc/grid-security/certificates/*
```

```
#vdt-control -on fetch-crl
```

Problem 9: SAM tests failing (mc, frontier, squid)

Errors : under mc it says edg-brokerinfo command not found, glite-brokerinfo-command not found. Under squid it says the squid server is down, under frontier it says can't connect to the ip address.

Solution 9: You need to restart the squid server.

```
# ssh uscms1-se.fltech-grid3.fit.edu
```

```
# /sandbox/squid/frontier-cache/squid/sbin/squid
```

It may say it was already running but this still fixed our problem.

Problem 10: Most of the SAM tests aren't being run. The few that are (sft-job and prod) are running very infrequently and failing because of proxy timeout errors.

Solution 10: NFS wasn't running. Need to restart it. Then can check status with second command, only 1 and 6 should be "off"

```
# service nfs start
```

```
# chkconfig --list | grep nfs
```

Problem 11: SAM test ANA failing. Error messages were unclear.

Solution 11: The SAM tests had all switched to a different database which we hadn't installed on our cluster. An email was sent from OSG with the phedex request to me so I approved the transfer, it installed itself properly and the SAM test went green.

Problem 12: jsub test is failing .

Solution 12: the user uscms01's home directory is mounted on /export/home/uscms01. The /export mount point had become 100 percent full as seen after doing df-h. and most of this memory was from the home directory so deleted some unnecessary files.

Delete files from \$GLOBUS_LCOATION/tmp, ~<user>/globus/.gass_cache,
~<user>/globus/job/\$(hostname -f), ~<user>/gram_scratch* where user was uscms01 for
all of these.

c. RSV

Problem 1: SE is periodically having issues. SRMv2 Service is in critical status

Critical Metrics:

SRMCP READ/WRITE: attempts to read and write against the SE using srmcp

ERROR: Cannot copy file: file:///opt/osg/osg-rsv/bin/probes/storage-probe-test-file to
srm://uscms1-se.fltech-grid3.fit.edu:8443/srm/v2/server?

SFN=/bestman/BeStMan/1282832881-storage-probe-test-file-remote.4103 using SRM protocol 2
... Bailing now.

Solution 1: Fixed with the solution to the SAM test problem at the same time. See SAM test:
Problem 3

Problem 2: RSV Probes aren't being run.

Solution 2: RSV probes run on condor-cron which is a vdt application.

#source /opt/osg-1.2/setup.sh

#vdt-control -on condor-cron

X. Other Issues

a. Certificates:

Problem: Phedex Certificate needs to be replaced after it expires

Solution: Certificates expire after one year so basically just replace the certs in .globus for the
phedex user with the administrator's new renewed certificate . Must move the old certificates
out of there first (keep a backup until you know the new ones are working, then you can delete
them. Then need to edit the proxyrenewal script:
sandbox/phedex/gridcert/proxyrenew.sh.forpatrick to include the certificate password. Then
just the run the script once to get it started. The script is on the cron tab so it will renew itself
weekly.

\$sudo mv /home/phedex/.globus/ <BackupLocation>*

\$ sudo scp ~/.globus/ /home/phedex/.globus*

\$vim /sandbox/phedex/gridcert/proxyrenew.sh.forpatrick

(type new password in where the old one was. Should be fairly obvious)

\$./sandbox/phedex/gridcert/proxyrenew.sh.forpatrick

b. Login Failed

Problem: User can't login in to their account on the cluster

Solution: Most likely they logged in repeatedly with the wrong password. Check for their
ipaddress in the file /etc/denyhosts if it is there remove that ONE ipaddress from the list.

c. Prevent GUMS and Gridftp probes from running

1) set the following to disable in the config.ini
enable_gums_probes = %(disable)s

2) Delete the submission file at

\$VDT_LOCATION/osg-rsv/submissions/probes/*__gums-authorization-validator-probe\@org.osg.gums.authorization-status.sub

3) Restart osg-rsv if it is running

vdt-control --off osg-rsv

vdt-control --on osg-rsv

- d. **Problem:** NAS isn't mounted on a node.

Solution: edit /etc/fstab to match the same file on other nodes. Then remount everything in that file with

mount -a

Then reset up the links for condor and restart condor.

ln -s /nfs/condor/condor /opt/condor

source /opt/condor/condor.sh

cp /opt/condor/etc/examples/condor.boot /etc/init.d/condor

chkconfig --level 235 condor on

#cp /opt/condor/condor.sh /opt/condor/condor.csh /etc/profile.d/

#/etc/init.d/condor start

- e. File is locked.

Problem: While trying to restart something or run a program you receive "opt/rocks/sbin/<program or command name>", line in run raise_ ("error - lock file %'s exists"> %self.lockFile error - lock file /var/lock/<program or command name> exist Generally you get this error with insert-ethers but solution is same regardless.

Solution: You need to delete the file so its no longer locked and you can run the command

rm /var/lock/<program or command name>

Then do whatever command caused you to originally receive the error.

- f. **Problem:** Did \$ sudo yum erase openldap-2.3.27-8.el5_1.3 In order to erase a file that was creating dependencies conflicts with another file that needed to be installed. As a result it deleted all of the files that depended on it as well, NAS stopped mounting, Users weren't sent to home directories, SAM tests failed, and websites were down.

Solution: First reinstalled all of the programs that had been deleted one by one.

#yum install <file name>

Reboot.

"Bad Superblock" on Nas is seen on Frontend during reboot

One of the missing programs was nfs-utils. When it was reinstalled we did

#mount -v nas-0-0.local:/nas0 /mnt/nas0

To manually remount the NAS and now the NAS was back up and attached and users were directed to their home directories. Restarted condor manually with

/etc/init.d/condor start

Tomcat connectors was the only program on the list YUM couldn't find so we downloaded it from source and reinstalled it which fixed the websites.

wget http://mirror.cc.columbia.edu/pub/software/apache//tomcat/tomcat-connectors/jk/source/jk-1.2.31/tomcat-connectors-1.2.31-src.tar.gz

http://serverfault.com/questions/228879/apache-load-balancer-server-backup

#tar zxvf tomcat-connectors-1.2.31-src.tar.gz

#cd tomcat-connectors-1.2.31-src/native

./configure --with-apxs=/usr/sbin/apxs

make

sudo make install

#/sbin/service tomcat start

#/sbin/service httpd start

Then to fix the SAM tests we restarted VDT (on the frontend and the SE), nfs on the frontend, and squid on the SE

```
#ssh uscms1-se.fltech-grid3.fit.edu
#cd /sandbox/bestman-1.2.13
#source setup.sh
#vdt-control --off
#vdt-control --on
# /sandbox/squid/frontier-cache/squid/sbin/squid
#exit
```

```
#cd /opt/osg-1.2
#source setup.sh
#vdt-control --off
#vdt-control --on
# /sbin/service nfs start
```

- g. Backups. To perform regular backups on our computer we use rsnapshot. It runs once a day at 1:30 in the morning and saves files to /mnt/backup/.snapshot. Each time it runs it creates an exact copy of the locations you asked it to backup in this directory. After creating ten snapshots it because overwriting the previous ones. So you will never have more than 10 snapshots at any given time.

Documentation is at <http://rsnapshot.org/howto/1.2/rsnapshot-HOWTO.en.html>

To edit the areas that are backedup edit the /etc/rsnapshot.conf file.

To edit the times at which backups are made, and how often, edit the crontab entry.

- h. Update Root:

Instructions are at: <http://root.cern.ch/drupal/content/installing-root-source>
copied tarball from internet and move it to /usr/local. Then,

```
# tar -xvf root_v5.30.02.tar
```

```
# cd root (note that this writes into the same directory as the previous root installation so that it overlaps and edits only the necessary portions.
```

```
# ./configure --prefix=/usr/local
```

```
# make
```

```
# make install
```

```
# ldconfig
```

```
# root
```

Problem 1: Can't find libCore.so one of the root libraries.

Solution 1:

```
#find / -name "libCore.so"
```

It was found in /usr/local/root/lib so I changed the places that Root was looking for it.

```
#vi /etc/ld.so.conf
```

Edited the entry to include /usr/local/root/lib then to remind root that I updated this file I did

```
#ldconfig
```

XI. Absolute Basics

Note this section is for any new administrators that know literally nothing about looking at documentation or solving problems. With a little common sense you can probably figure most of this out easily.

- a. Commands

You should have received the Linux Noob commands along with this so I won't go over those. I have not included descriptions of most of these because you are better off reading the man pages. Type man followed by one of the commands or simply googling it. You will learn more. Also try the commands without the options or with different ones to really get a feel for them.


```

df -h
condor_q | less
condor_status
man
scp HostnameToGive:<PathToFileToTransfer> HostnameToReceive:<FilePathToPlaceIt>
    ex. Scp johndoe@uscms1.fltech-grid3.fit.edu:/home/johndoe/file.txt
    johndoe@myhomecomputer:~
ssh name@host
export VARIABLENAME = Destination PATH
du -sh
vim
ln -s
uname -a
top
chown
chgrp
lastlog
cat /proc/version
tar -xvf <tarball>
sudo nautilus
#iptables -L | less
mount -a

```

b. Syntax

- If you see something like CRAB_x_x_x in documentation. Where CRAB stands for any program. The x's are standing in place of version numbers. SO when you go to execute this command you must change them.
- If you see something like `cd <directory> .` (Basically anything within <> means type your own and leave out the <> when you execute the command. i.e. `cd myhomedirectory`)
- Whenever you create a new directory for someone or something make sure the group and owner are appropriate with `ls -la` and `chown/chgrp`
- Google google google! And make backups!

c. Test Job: Sometimes as an admin its valuable to have a job that you can submit to test the status of the cluster/condor/etc. Ben Locke made me (Xenia) this one so as long as it exists it can be used. It runs in about 1 minute and does basically nothing.

It needs to be run from `/mnt/nas0/home/g4hep/geant4/examples/mytestapps/benL`

Run it with

```
condor_submit submit.sub
```

Check on the status by looking at it in the `condor_q`

```
Condor_q | grep g4hep
```

The output goes to

```
/mnt/nas0/home/g4hep/geant4/examples/mytestapps/benL/mtsProtoBen
```

d. Add a User:

We have a script to automate this process. Doing `sudo adduser username`. Won't create the home directory. Instead do...

```
sudo su root
cd ~
./UserCreationScript
```

