

Fall 2011 Cluster Administration Notes- Xenia Fave

~~~~~

8/24, 12-2

Worked with Vallary on the very beginning of cluster admin training

8/25, 9-10

Submitted for new certificates for both jojo and myself

Got whole thing sorted out.

8/26, 9-10 and 12-2

-Worked with jojo on all 8 old machines to begin seeing which ones were worth keeping

and which we should go ahead and take apart for parts.

-added my new certificate to my cluster account and updated the phedex certificate as well.

-no response from Doug yet.

-Want to begin work on the New NAS. Following parts of Patricks instructions for the SE

because we want to put both phedex and bestman on it.

instructions are at <http://uscms1.fltech-grid3.fit.edu/SED.pdf>

First main thing for the New Nas is to get it certificates? (host and server)

right now it doesnt have external ethernet access so maybe not..

can we just have bestman and phedex save files on to ?

-With Vallary's certificate we imported it, saved it as a file, and put it on the cluster.

8/29-9/2

There was a power outage that created some issues. we were able to bring the cluster down in time.

however upon rebott the nas02 (burcus) and nas01 (our fancy one) didnt come up.

part of the problem for ours i think is that it wasnt on when i brought the frontend up. although it

isn't in the df -h though you can ssh to it normally. which is interesting.

should just be able to run the file by hand to mount it.

to get website running after power outage I had to do

```
sudo /sbin/service httpd
```

also needed to restart bestman so just restarted vdt wherein one of the things it runs is bestman

ssh to the se

switch to root

```
cd /sandbox/bestman-1.2.13
```

```
source setup.sh
```

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

then ran rsv probes by hand and wait til they update (they can set to the queue like regular jobs)

```
su rsvuser
```

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

also the phedex proxy wasnt being updated. had to change out the certificate that updates them. instructions for this already exist in the admin guide but basically, you copy your usercert and userkey to the phedex home .globus directory. and then change the password in the proxy renewal script and make sure your certificate is registered with the CMS vo. to be allowed to register your cert with the cms vo you have to have a cern account.

have to reregister this certificate with the cms vo  
<https://lcg-voms.cern.ch:8443/vo/cms/vomrs?path=/RootNode/MemberRegistration&action=execute>  
they say i dont have a cern account so im checking.  
the cern account is with xfave or xfave2008@fit.edu  
[accounts.cern.ch](https://accounts.cern.ch)  
so for my email on the vomrs page used xenia.janice.fave@cern.ch  
hopefully that mail will be accessible.  
i can then read my cern emails through thunderbird  
WHOOO  
next...just waiting for them to approve me for the VO  
then do  
ssh dev-0-0  
sudo su phedex  
cd /sandbox/phedex/gridcert  
./proxyrenew.sh.forpatrick

9/7  
restarted vdt as explained above  
now...

-go check on Burcu's magical nas.

-run file that mounts everything by hand.  
the /etc/fstab file doesnt actually have the nas10 in it as an entry.  
unusual.

9/8/11: 9:30-10:30  
want to install codeblocks.  
first question is where to install it so all users can access it same way they can access other commands  
i believe its /usr/bin but doublechecking...

aha total instructions for instructions for installing code::blocks is at [http://wiki.codeblocks.org/index.php?title=Installing\\_Code::Blocks\\_from\\_source\\_on\\_RPM\\_based\\_distributions](http://wiki.codeblocks.org/index.php?title=Installing_Code::Blocks_from_source_on_RPM_based_distributions)  
might be the perfect project for vallary! but perhaps we'd better do it together.

9/12: 9-10

Wrote script took a look at bestman/phedex helped jojo with wiki

9/13: 12:30-2.

The partitions on the nas1 arent being mounted to /mnt/partition on the frontend.

so there was no entry in the /etc/fstab file to mount the /general directory on to our frontend.

Just edited /etc/fstab with a line of code

```
nas-0-1.local:/general /mnt/genral nfs defaults, rsize=65536, wsize=65536,
intr, noatime 0 0
```

Keep an eye on <https://ticket.grid.iu.edu/goc/11000>

9/16: 12-1

SO! if you copy a file from my home directory to /mnt/backup and then go log into the nas1 it is in fact in that location.

so for the back up i need to write a script that copies everything to that location.

the problem is that it does change the permissions of the file which is uncool....HMMMMMMM

9/19:9-10

simplest one is just creating a tarball. trying it on my home directory with

```
su tar czf - /home/xfave > /home/xfave/backup.date
```

began at 9:21 ended at

to restore just do

```
cd /home ; su 'tar xvf /tmp/bkup.yyyymmdd'
```

This command

```
$ find /home -mtime 1 ! -type d | tar cT - >/dev/rmt1
```

finds all the files that have been updated today, saves them into a tarball on a different device

or use

```
$ find / -mtime -1 | cpio -o >/dev/rmt1
```

which saves everything thats been updated today (including directories if they've had one file within change)

cpio can be more stable.

could just put this into crontab

```
001 30 2 * * * /bin/tar cvzf /mnt/nas1/backup/daily.tgz /home
```

Thoughty thought thoughts.

Sam test ana failed because SAM had switched databases so had to phedex the new one. I had recieved an email about

the new database two weeks ago or so but hadnt phedexed it because i didn realize it was integral for us.

also had to fix my certificate with phedex. or update the one phedex associates with me is a better way of

putting it. So i went to

<https://cmsweb.cern.ch/sitedb/people/>

and under Distinguished Name: replace with the DN from my certificate.  
i.e. /DC=org/DC=doe grids....blahdyblah

Met with john and curtis 2-3.

bc it is raid5 and two disks failed there are huge chances all the data is gone.

if it was raid6 you can lose up to two disks and be fine.

hooked ethernet cable from the machine to my laptop.

went to webpage: 10.255.255.250:8080/cgi-bin/index.cgi#

user=admin

pw=admin

came from the manual.

started scan of the disks that failed under Volume Management. These scans take a LONG time. so unplugging ethernet

and going to check on it later. it will continue running the scan.

I'll come back after my next meeting and check on its progress. hopefully it will finish before tonight.

not sure if this scan will fix anything but it should identify "bad blocks" which according to eric is probably physical problems with the disk.

if none turn up then i will run, recover under RAID management and we will find out if the system itself is recoverable

and or any data. but chances of this are slim.

Most likely, burcu and Sam will need to use the warranty or some grant money to buy two new disks for it.

and the data will have to be re run.

/usr/lib64

/usr/local/lib64

/usr/lib/x86\_64-redhat-linux4E/lib64

9/22: 10-12

In the highbay, so the scans finished and now the two disks that were red have turned green and say that they are good.

still no monitor connection

still can't mount it

still can't ssh to it (it closes off the connection immediately)

went to RAID managment and did recover, it foudn the size very quickly

9/26: 9-10

Condor\_q filled with old jobs (from May) deleting some now making copies fo two files

/mnt/nas0/OSG/APP/crab/CRAB\_2\_7\_8\_patch1/python/SchedulerGrid.py and SchedulerLocal.py

added line txt += '[ \$middleware == OSG ] && echo source \n\$OSG\_GRID/setup.sh >>

> CacheEnv.sh\n'

btwn 290 and 291 on schedGrid and btw lines 150 and 151 on Schedlocal.  
see if this works... if not put back the original files.

\*\*\*\*\*

For THE CE: (as root)  
source /opt/osg-1.2/setup.sh  
vdt-control --off or --on

FOR THE SE: (as root)  
source /sandbox/bestman-1.2.13/setup.sh  
vdt-control --off

\*\*\*\*\*

ALSO

/var was 100 percent full  
so deleteed /var/log/access\_log.\*

10/17: 8:00-10

BACKUPS:

-installed rsnapshot  
-to edit the way backups are made edit the /etc/rsnapshot.conf file which  
has very clear notation within it for which each setting does.  
should be installed and working.  
the documentation is WONDERFUL.

I set it up with the instructions from here:

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

runs twice a day every day which is a setting that can be changed very  
easily if we see it slowing everything down way too much  
is configured to work with crontab

\*note to self send a thank you to brian from tier3 for suggesting it

ROOT Upgrade

is going well!

10/18:8-10, 1-2

HIMALI:

Two issues need to be resolved for him.

1) the directory where his files are saved after he runs jobs doesn't  
allow him to delete those files. which means it can fill up quickly.  
solution: made his directory writable by anyone which is mildly dangerous  
but should allow him to delete files that are bieng created by him as  
uscms01.

the directory: /mnt/nas0/OSG/BeStMan/cms/store/user/hkalakhe

2) These files should anyways be being saved to the NAS1 but we tried to  
edit his config file in the most obvious of ways (since I don't know much  
to anything about CRAB config files) and they wouldn't run if we changed  
the destination directory. I sent an email to the tier3 list in order to  
get some help on this.

One possible but not ideal solution would be to create a softlink in his  
directory that then sends things to the NAS. I will work on implementing  
that until we reach a more permanent solution.

note, would have to be a softlink because things that are hardlinked must reside on the same file system.  
<http://linuxgazette.net/105/pitcher.html>  
NOPE not a workable solution because deleting something in the linked directory deletes the original as well. they are always identical.

#### ROOT Upgrade:

```
instructions from http://root.cern.ch/drupal/content/installing-root-source
copied tarball from interweb ;)
copied it to /usr/local (which is where the previous root installation
also was.
# tar -xvf root_v5.30.02.tar
# cd root (note this wrote into the same directory as the last root
installation so hopefully that won't affect anything (bites nails)
# ./configure --prefix=/usr/local (bc thats where it is)
# make (this takes some time)
# make install
# ldconfig
# root
```

Error when try to open root: Can't find "libCore.so"  
which is a library for root. Hmm. it appears in /usr/local/root/lib and /usr/local/lib/root/lib so i think the cluster just doesnt know to look there. from the installation documentation at the previous link i got the hint to edit the file:  
# vi /etc/ld.so.conf  
added the entry  
/usr/local/root/lib  
so now it knows to look there for libraries.  
then to tell it that I updated this file I did  
# ldconfig  
and now it works :D  
not a bad day's work !

#### BACKUP

edited so it only runs once a day because i think thats plenty for us. it is saving properly in the nas 1 and everything looks fine.

10/19: 8-10

#### Pointing BeStMan/PhEDEx:

got some responses from my email! basically involves editing config files  
From Marguerite:

- [1] storage.xml: which is installed in a PhEDEx directory, a CMSSW SITECONF directory, and copied to CVS (COMP/SITECONF)
- [2] config.ini (there were about 4 locations in the file to change directory name in the OSG version we're running)
- [3] Add the new directory name to configure bestman --> we use --with-allowed-paths "/tmp;/home;/data;/hadoop" and configure\_bestman instead of editing the .rc file as Bill suggests.

From Bill:

You will need to change bestman2.rc (We DONT have to edit the .rc file because we have bestman version 1.x.x.x) and edit your localPathListAllowed entry. Also you will need to edit storage.xml. (Same as Marguerite) Depending on your cluster configuration you may also need to edit \$VDT\_LOCATION/osg/etc/config.ini. (Same as Marguerite)

So first finding the files:

```
[1] found on SE at /sandbox/phedex/SITECONF/T3_US_FIT/PhEDEx/storage.xml
<-----this is the SITECONF one
/sandbox/phedex/sw/slc4_amd64_gcc345/cms/PHEDEX/PHEDEX_3_2_0/Custom/Template/storage.xml
<---not sure if useful
```

```
[2] /sandbox/bestman-1.2.13/osg/etc/config.ini
change a few locations under the Storage header.
that one looks like its never been modified. more likely that i need to edit
```

```
/opt/osg-1.2/osg/etc/config.ini
```

```
locations: [storage]
```

```
[3]/sandbox/bestman-1.2.13/vdt/setup/configure_bestman
```

```
will do
```

```
ON CE: # mkdir /mnt/general/OSG/BeStMan
```

```
ON SE: # configure_bestman --with-allowed-paths "/mnt/nas0/OSG/BeStMan;
/mnt/general/OSG/BeStMan"
```

We shouldn't need to do this because the configure\_bestman script does it but the bestman.rc file is at /sandbox/bestman-1.2.13/vdt-app-data/bestman/conf/bestman.rc and inside it right at the beginning is the localPathListAllowed variable which we would have to change to include

10/24/11: 8-10

-Backups are running beautifully. only daily ones are being saved. memory is not an issue yet, at less than 30%

-secondary goal is finding a way for himali to be able to delete his files. ask him if he knows the uscms01 password because that would simplify things.

-put up your documentation onto the wiki

-Big goal for the week is configuring phedex to point to nas.

so to accomplish this we're going to be following the advice of marguerite and bill from last week.

10/26/11:12-2

So there was an RSV problem where none of our rsv probes had run in the past couple of days.

When I try to run them manually they each give the following error:

```
#su rsvuser
```

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

```
#cd ~
```

```
# ./RunAllRSVProbes
```

```
Using VDT_LOCATION: /opt/osg-1.2 to run metric: org.osg.srm.srmcp-
readwrite
Using condor-cron submission file: uscms1-se.fltech-grid3.fit.edu__srmcp-
srm-probe@org.osg.srm.srmcp-readwrite.sub ; Will create /tmp/uscms1-
se.fltech-grid3.fit.edu__srmcp-srm-probe@org.osg.srm.srmcp-
readwrite.sub.19505.sub
Submitting /tmp/uscms1-se.fltech-grid3.fit.edu__srmcp-srm-
probe@org.osg.srm.srmcp-readwrite.sub.19505.sub into Condor-Cron
Submitting job(s)
ERROR: Failed to connect to local queue manager
CEDAR:6001:Failed to connect to <127.0.0.1:57373>
```

I tried stopping and restarting the RSV service but received the following message:

```
[root@uscms1 xfave]# source /opt/osg-1.2/setup.sh
[root@uscms1 xfave]# vdt-control --off osg-rsv
disabling init service osg-rsv... ok
[root@uscms1 xfave]# vdt-control --on osg-rsv
enabling init service osg-rsv... FAILED! (see logs/vdt-control.log)
```

I checked the log and it said that condor wasn't running, full message below.

```
#less /opt/osg-1.2/logs/vdt-control.log

### 2011-10-26 12:12:55 vdt-control(initialize) starting
### 2011-10-26 12:12:55 vdt-control(do_init) disabling the init service
'osg-rsv'
### 2011-10-26 12:12:55 vdt-control(system) /etc/rc.d/init.d/osg-rsv stop

-- Failed to fetch ads from: <127.0.0.1:57373> : uscms1.fltech-
grid3.fit.edu
CEDAR:6001:Failed to connect to <127.0.0.1:57373>
ERROR: Condor does not seem to be running. Unable to continue.
### 2011-10-26 12:12:55 vdt-control(system) /sbin/chkconfig --del osg-rsv
### 2011-10-26 12:12:55 vdt-control(safe_remove) backed up
'/etc/rc.d/init.d/osg
-rsv' to '/opt/osg-1.2/vdt/backup/root/etc/rc.d/init.d/osg-
rsv_019_20111026-1212
55'
### 2011-10-26 12:12:55 vdt-control(safe_remove) removed
'/etc/rc.d/init.d/osg-r
sv'
```

However condor is functioning at our site, I can run `condor_q` ,  
`condor_status` with no problem and it turns up a lot when I do `ps -ef |`  
`grep condor`

I still tried stopping and restarting condor to see if it would have an effect but continued to receive the same errors as above from RSV.

```
#service condor stop
#service condor start
```



Aha! due to some very nice people on the tier 3 list, they pointed out that rsv runs on condor-cron not on condor so i needed to restart that service

checked to see if it was running initially with

```
#ps -ef | grep condor-cron
```

then restarted with

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

```
#vdt-control --on condor-cron
```

then reran rsv probes as before and they worked!

everything is green!

proves that before doing anything always restart all of vdt and you should be fine.

~~~~~

For BESTMAN pointing to NAS1

Should commands be run on SE or CE? the age-old question

[1] storage.xml (should be a siteconf, a phedex, and a cvs one)

```
[root@uscms1-se ~]# find / -name storage.xml
```

CE ones

```
/mnt/nas0/home/pford/SITECONF/T3_US_FIT/PhEDEx/storage.xml---> just history?
```

```
/mnt/nas0/home/pford/tempcvs/SITECONF/T3_US_FIT/PhEDEx/storage.xml-->how it used to be back when used ce as se
```

```
/mnt/nas0/OSG/APP/cmssoft/cms/SITECONF/T3_US_FIT/PhEDEx/storage.xml--->almost definitely but not sure how to edit. unless just change everything?
```

```
/mnt/nas0/system-config/Config/PhedexConfig_3_2_0/storage.xml----> also change!
```

SE ones

```
/sandbox/phedex/SITECONF/T3_US_FIT/PhEDEx/storage.xml
```

```
/sandbox/phedex/sw/slc4_amd64_gcc345/cms/PHEDEX/PHEDEX_3_2_0/Custom/Template/storage.xml
```

On the CE:

[2]config.ini change 4 lines

```
/sandbox/bestman-1.2.13/osg/etc/config.ini <---never been modified
```

```
/opt/osg-1.2/osg/etc/config.ini <----- more likely to be correct look under [storage]
```

```
data_dir- /mnt/nas0/OSG/Data change to /mnt/general/Data
```

```
site_read='' same as above
```

```
site_write = " same as above
```

[3] /sandbox/bestman-1.2.13/vdt/setup/configure_bestman

will do

```
ON CE: # mkdir /mnt/general/OSG/BeStMan
```

ON SE: # configure_bestman --with-allowed-paths "/mnt/nas0/OSG/BeStMan;
/mnt/general/OSG/BeStMan"
or also on CE: # /opt/osg-1.0.1/vdt/setup/configure_bestman

10/27:9-11, 12-1

SAM test: Jsub (job submission) began failing around midnight last night.
Did a few checks to see if andrea scabia (the guy who submits them all
for osg)'s proxys and cert were still identifiable by our cluster but
they all seem good. also as far as i can tell all the other tests are
still currently running and not having issues.
I am restarting vdt and then emailing out to the group for any tips.
restarting vdt on front end and storage element.

So the uscms01 home directory is mounted on /export/home/uscms01. the
/export mount point had become 100% full as obvious after doing df -h.
Most of the memory ~290 G of the ~330 were from this home directory. So
we deleted some files that were unnecessary and up to several years old.
these included files in the following places: (recommended by bockjoo to
delete these after 7 days)

\$GLOBUS_LOCATION/tmp
~<user>/globus/.gass_cache
~<user>/globus/job/\${hostname -f}
and files ~<user>/gram_scratch*

10/28: 12-2 pm

want to install jobmon as recommended by tier 3 list.

after have it installed can integrate it and clarens with vdt using
instructions here

<https://twiki.grid.iu.edu/bin/view/Integration/JobMonReadinessPlan>

11/1: 12:30-2, 3:30-600

Want to get everything on my to do list done today

First! Configure NAS1

here goes.

notes are in other file!

headed to highbay to work on another diff issue.

/mnt/nas0/OSG/APP/cmssoft/cms/SITECONF/T3_US_FIT/PhEDEx/storage.xml---

>almost definitely but not sure how to edit. unless just change
everything?

/mnt/nas0/system-config/Config/PhedexConfig_3_2_0/storage.xml----> also
change!

most likely need to change SE to being CE and then edit the
path-match t0 be /mnt/general/(.*)

then

```
/opt/osg-1.2/osg/etc/config.ini <----- more likely to be correct look
under [storage]
data_dir- /mnt/nas0/OSG/Data change to /mnt/general/Data
site_read='' same as above
site_write = " same as above
```

then run

```
/opt/osg-1.0.1/vdt/setup/configure_bestman
```

```
configure_bestman --with-allowed-paths "/mnt/nas0/OSG/BeStMan;
/mnt/general/OSG/BeStMan"
```

```
$VDT_LOCATION/vdt/setup/configure_bestman --server y \
--user daemon \
--cert /etc/grid-security/containercert.pem \
--key /etc/grid-security/containerkey.pem \
--http-port 8080 \
--https-port 8443 \
--gums-host uscms1.fltech-grid3.fit.edu \
--gums-port 8443 \
--enable-gateway \
--with-allowed-paths=/mnt/nas0/OSG/BeStMan \
--with-transfer-servers gsiftp://uscms1-se.fltech-grid3.fit.edu
```

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'
```

So now actually doing it

```
made directory /mnt/general/bestman/cms/store
change permissions on all the /mnt/general/bestman/cms/store
should be root and osgusers with rwx,rwx, rx
then on cms and store uscms01, uscms01 rwx, rx, rx
```

```
# cd /mnt/nas0/OSG/APP/cmssoft/cms/SITECONF/T3_US_FIT/PhEDEx/
#cp storage.xml storage.xml_backup
# vi storage.xml
edited it to look the way it does now versus the backup one
```

```
#cd /mnt/nas0/system-config/Config/PhedexConfig_3_2_0/
#cp storage.xml storage.xml_backup
#vi storage.xml
looks the same as the other one now
permissions are rw, r, r
```

```
# cd /opt/osg-1.2/osg/etc/
#cp config.ini config.ini_backup
```

```
#vi config.ini
permissions are rw,r, r
didn't change anything yet. right now its all the Nas0 which makes
sense bc thats where a lot of stuff is instlaled not where its saving
stuff. couldn't find any lines to change... unless maybe the SE?
that should be the first thing to try if it deosnt work the first time
```

```
#/opt/osg-1.0.1/vdt/setup/configure_bestman
# ./configure_bestman --server y --user daemon --cert /etc/grid-
security/containercert.pem --key /etc/grid-security/containerkey.pem --
http-port 8080 -https-port 8443 --gums-host uscms1.fltech-grid3.fit.edu -
-gums-port 8443 --enable-gateway --with-allowed-paths
"/mnt/general/bestman" --with-transfer-servers gsiftp://uscms1-se.fltech-
grid3.fit.edu
Error can't access Bestman in: '/opt/osg-1.2/bestman'
```

so tried it on the SE... cause that can't spell doom.

```
:(
#ssh dev-0-0
#source /sandbox/bestman-1.2.13/setup.sh
#cd /sandbox/bestman-1.2.13/vdt/setup
#[root@uscms1-se setup]# ./configure_bestman --server y --user daemon --
cert /etc/grid-security/containercert.pem
--key /etc/grid-security/containerkey.pem -http-port 8080 -https-port
8443 --gums-host uscms1.fltech-grid3.fit.edu
--gums-port 8443 --with-allowed-paths "mnt/general/bestman" --with-
transfer-servers
gsiftp://uscms1-se.fltech-grid3.fit.edu
```

```
if it doesnt work
1) run vdt-control --off and re do the command and then do vdt-control --
on
if it still doesnt work
2) edit the se versions of the edited files too and reurn the vdt stuff
if it still doesnt work
3)email list with steps that were taken and what was edited in the files
and perhaps a summary again of
what you specifically want.
```