

## Documentation Fall 2010

~~~~~

8/17

Need to install certificate on Phedex user

-need to renew my certificate

followed instructions in email they sent me about how it was about to expire, had to do it in firefox. google chrome was uncooperative. It was seeable under my certificates.

-export it to the cluster

followed instructions in the follow up email. also located here

<http://www.doeagrids.org/pages/cert-request.html#Globus>

-Note that patrick's membership with CMS has expired, so xenia needs to replace my certificate on the phedex user with hers and edit the proxyrenew script with her password. basically just replace the certs in .globus for the phedex user

i.e. `sudo scp ~/.globus/* /home/phedex/.globus`

note! it was necessary to move the old certificates out of there first.

Then need to edit the proxyrenewal script `/sandbox/phedex/gridcert/proxyrenew.sh` for patrick to include my password...

-test it...

not working...

8/18

\*Added jojo's spring 10 and patrick's summer 10 log to the cluster wiki

see instructions in my spring 2010 documentation

\*Phedex certificate

Errors received: error while loading shared libraries: libpython2.4.so.1.0: cannot open shared object file: no such file or directory

user certificate not found

`mv: cannot stat /sandbox/phedex/gridcert/proxy.cert.NEW`

-changed the password to my certificate password not my user password which was pure silliness on my part from the beginning making that type of rookie mistake.

-changed permissions and users and groups on files. to mimic patricks. so now its

`.globus rw-rw-rw phedex phedex``usercert.pem rw-r-r- phedex phedex``userkey.pem rw-----phedex phedex`-still receiving a library error. I found the library at `/usr/lib64` with the find command, put it in `/lib` and got a different error about wrong ELF class. : ELFCLASS 64. It still says the check passed and i looked at the condorq and there are osg and uscms01 jobs running and a couple idle now so i think i fixed it :)

Having just looked at Gratia it looks like we've been getting jobs for a week from cms and glow o actually this is entirely not conclusive.

-ELF class error means the library was compiled on a 64bit and i'm now trying to use it on a 32 bit machine.

8/19

\*Jo Jo's certificates

-I'm to become the new CA (certificate authority) for Florida Tech since Patrick graduated.

-I installed mozilla thunderbird so I would be able to send digitally signed emails. then followed the instructions here to import my certificate into it.

[https://twiki.grid.iu.edu/bin/view/Security/OsgRaOperations#Digitally\\_signed\\_email](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.

8/20

To Do:

-become a CA

-JoJo's Certificates

-Phedex proxy

-check on borris?

-help jojo with setting up the SL4.8 machine

-read patrick's documentatin again for anything needing to be done

\*CA

-everything's all set up, whenever someone from florida tech needs a certificate they'll come tell me or email me more likely. then write my name in as sponser. then i'll get a message from fermilab asking for validation to which i just need to reply with a digitally signed email.

\*Phedex Proxy

-searching for the missing library

```
sudo find / -name "libpython2.4*"
```

-results:

```
/mnt/nas0/osg/app/cmsssoft/cms/slc4_ia32_gcc345/cms/cmssw/cmssw_3_1_6/external/slc4_ia32_gcc345/lib/libpython2.4.so.1.0
```

-cp file to /lib

so now it runs with no errors and says check passed.

Only issue i have with this theory is if i comment out the first check of the proxy then it should renew the proxy again and then recheck it. and this recheck gives errors specifically

```
Error: cms: unable to satisfy B/cms/uscms:cmsphedex Request!
```

None of the contacted servers for cms were capable of returning a valid AC for the user.

if i then comment out the second renewal of the proxy then it just says i need to renew the proxy and that the status is 1.

-for the record phedex has no hours logged in the past 30 days as seen at

```
http://cmsweb.cern.ch/phedex/prod/Reports::DailyReport
```

-also now it's spitting out all those errors again... not good

```
Error: cms: unable to satisfy B/cms/uscms:cmsphedex Request!
```

None of the contacted servers for cms were capable of returning a valid AC for the user

```
mv: cannot stat '/sandbox/phedex/gridcert/proxy.cert.NEW': no such file or directory
```

```
voms_proxy_lifetime: 0
```

```
Warning: need renew the proxies
```

```
Status is :1
```

\*vorris page = myosg, flags for storage element: apparently fixed themselves or Patrick fixed them before leaving

-try submitting a job from uscms01, asked Samir to do one when he had a chance and let me know if it runs successfully.

- I see some in the condor\_q so i think that's ok

\*Back to Phedex...

-broke down and emailed patrick about it. he has no idea why it stopped working but managed to fix it by changing the voms-proxy-init line in the script to only do =voms cms.

8/23 <3

\*Phedex Proxy

-added the script to the crontab. Should renew automatically every week now.

```
crontab -e
```

```
0 0 * * 0 /sandbox/phedex/gridcert/proxyrenew.sh.forpatrick
```

\*SL4.8 Box

-had to rescrew in the video card because it was loose, then turn off/on the monitor.

now booting up scientific linux

instructions are... here

```
https://www.scientificlinux.org/distributions/4x/installing/install-guide-full.html
```

burnt the next two cds too. (it claimed it only needed the first three)

disk one completed great

disk two is running now.

disk three is sitting on top of the computer

\*Hardware:

patrick had planned to buy

-2 power strips

-visualization machine in the lab, anything attached to the cluster would be slowed down too much by the network : nvidia gtx470, core i7 quad core processor, 6gm ram

under 1000, dual monitors

-new nas: 2 static drives in raid1 config, 36 drive array. raid50 use 3 striped raid5 arrays with 1 hot spare each. tis will lose 6 drives worth of free space. then compare dd speeds with raid6 and decide which is better

\*HEP MEETING: put the latest compatible version of Root on SL machine

\*do we want this machine connected to the domain?

\*create users for it.

8/24

\*SL4.8 Machine

Installation went great and the computer doesnt even seem that slow.

Its going to need internet

\*making a root disc

used gcc -v to determine the compiler version we have: 3.4.6 thus downloaded the latest version of root compatible with SLC4/ gcc3.4 from <http://root.cern.ch/drupal/content/production-version-526>

copied the tarball to a cd. and should be able to download it onto the computer..

and it worked :) copied from the dvd to my desktop. and then to /usr/local/include

following the instructions at the previously stated website i did

```
gunzip root_v5.26.00.Linux-slc4-gcc3.4.tar.gz
```

```
tar root_v5.26.00.Linux-slc4-gcc3.4.tar
```

then put tarball into /usr/local/src

read the Readme. need to set two enviromental variables

```
PATH & LD_LIBRARY_PATH
```

```
export VARIABLENAME="DESTINATION PATH"
```

or if appending

```
export VARIABLENAME=Variable Name:"DESTINATION PATH"
```

8/25

\*SL4 Machine

```
-set env variables in /etc/profile- done
```

```
-add user jfischer and user kgnanvo-done
```

\*Logins

-this morning ben tried to log on to cluster repeatedly but wasn't typing the G4HEP@ part so he was logging in as blocke which doesnt exist so after 5 failed attempts the system blocked the IP address. By logging in from a different IP he was fine. Qustion is how long between resets so he can log in from the original computer again.

-i went in and edited the /etc/denyhosts file and removed the ipaddress of Electron from it. ergo should have fixed the problem.

\*Jojo's Certificates

- DOE we imported it successfully to her browser and backed it up now just need to import it to the cluster. I emailed her the instructions

```
-CMS: for institute: its under USA>Florida Insts of Tech
```

```
*Ben wants a command that lets him view .png files
```

```
answer: display File_name
```

could it be any easier? note you have to have imagemagik installed for this command to work.

luckilily on the cluster this is the case.

8/26

\*SE problem.

```
-looking at myosg its clear that our SE is having periodic issues.
```

"

```
SRMv2 Service is in critical status
```

```
Critical Metrics
```

```
SRMCP READ/WRITE: atempts 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.
```

```
Actu
```

```
SRM Ping: check if the SRM server responds
```

```
SRM server running on uscms1-se.fltech-grid3.fit.edu is alive and responding to the srmcping command.
```

```
Got this from srm-ping command: Storage Resource Manager (SRM) Client version 2.1.5-3
```

8/27

\*SE problem

so true to form today everything is working well on it. it says:

```
SRMCP Read / Write
```

```
Attempts to read and write against the SE using srmcp.
```

```
Hide Detail
```

```
128 /bestman/BeStMan/1282915681-storage-probe-test-file-remote.6893
```

```
List file file srm://uscms1-se.fltech-grid3.fit.edu:8443/srm/v2/server?
```

SFN=/bestman/BeStMan/1282915681-storage-probe-test-file-remote.6893 successful.  
Actual output was: 128 /bestman/B

#### SRM Ping

Check if the SRM server responds.

Hide Detail

SRM server running on uscms1-se.fltech-grid3.fit.edu is alive and responding to the srmping command.  
Got this from srm-ping command: Storage Resource Manager (SRM) Client version 2.1.5-3  
Copyright (c) 2002-2009 Fermi National Accelerator Laboratory

-a look at the history shows this periodicity started on July 7.  
and its always able to Ping its the read/writing that goes in and out.

-<https://twiki.grid.iu.edu/bin/view/MonitoringInformation/TroubleshootRSVErrors> for information  
-thoughts: upgrade RSV,

#### \*Phedex Issue

Dr. Hohlmann wants a couple things Phedexed to us, see email for details.

Samir tried to do it but ihe got a thing saying Phedex wasn't working.

Need to check that.

I tried to do it but I can't even access the DBS repository because my cern account is "disabled" or so it says when I try to log in. I need to email. [computing.help@cern.ch](mailto:computing.help@cern.ch) or something.

8/30

\*emailed cern they said they needed to reset the password and i sent the information they needed  
\*samir on friday tried to phedex the stuff to us but it didnt work. possibly because the SE was down at the time. Thats the priority issue right now.  
\*Need to test root on scientific linux 4 machine. but i think im going to pass that to johanna laina. should be pretty basic.

\*SE

-emailed patrick \*sigh\*

-Reading about storage element documentation

-read about phedex, is there any way to check if its working easily?

<http://cmsweb.cern.ch/phedex/prod/Components::Status?view=global>

check that filedownload and fileexport are both UP. If so then its probably not a problem with phedex but a problem with the SE which makes more sense anyways.

a ton of inforomation on Phedex:

<https://twiki.cern.ch/twiki/bin/viewauth/CMS/PhedexUserDocsMonitorData>

-Plan Via Patrick: Update VDT software

\*Updating VDT on the SE (essentially bestman)

-instructions in the SE documentation

and the following websites:

<https://twiki.grid.iu.edu/bin/view/ReleaseDocumentation/OSG12UpdateInstructions>

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

8/31

\*Will told me photon was broken: needed a new power supply

-somehow missing all user accounts. bleh... JoJos job is to fix it.

\*Will has a local account on electron with sudo privileigis for his needs

-used visudo to open /etc/sudoers file and added him to the list.

-use id USERNAME when you want to see which groups a person belongs to.

-usermod -G <group> <user> to add an existing user to an existing group

\*Phedex requests,

-tried commands in email last night but it didnt work looks to be a problem on their end with their certificate but not really sure.

9/1/10

\*Phedex Requests

-Samir figured out way to phrase command and it worked

\*automate phedexing of jobs

- go to DBS, click phedex under whichever files you want to transfer, select t3\_us\_FIT

-request is made, approve request

-so only real question is which ones do we want.

\*VDT upgrade

-going to follow instructions at  
<https://twiki.grid.iu.edu/bin/view/ReleaseDocumentation/BestmanGateway>  
 or...  
<https://twiki.grid.iu.edu/bin/view/Integration/ITB1117Documentation?>  
 asking Patrick...

\*Tau:

-internet works fine, we just need the port to be turned on and then we can wire an ethernet cord to it.  
 -tested root, no problems there.  
 \*Sudo privs  
 -gave them to jojo on cluster, electron, and tau.

9/3/10

\*Failing Sam Tests

-theory, it started failing a s of midnight last night. sounds like might be a cert problem . I think i might not be registered with the cms vo with my new certificate. updating that and seeing if it helps.

-so i went to

<https://lcg-voms.cern.ch:8443/vo/cms/vomrs?path=/RootNode&action=execute>

> Member Info> Resign the grid and VO AUPs because mine had expired.

this had stopped the phedex proxy from renewing with the vo. if that proxy had expired last night and not been able to renew bc my membership with the vo had expired last week then that would make sense. not sure if it would explain the sam tests.

>On the cmsDashboard page, i see two errors. ONE on the CE-MC and one on the SE-srmv2

clicking on either one shows everything cms ran to try to test it and what failed. in both cases the certificate was the problem. not sure if its the same certificate as the phedex one. going to go take a look at it.

>looks like that may have been a problem for phedex but not for the SE and CE, instead their errors are from their own certificates. by doing

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

it was obvious the certificate expired sept 1, the exact day the tests started failing. so just need to renew them.

following instructions at <https://twiki.grid.iu.edu/bin/view/ReleaseDocumentation/GetGridCertificates>  
 (p.s. pacman is found in /mnt/nas0/usr/pford/home) or something like that. use find command

>waiting for email with the certificates...bleh...bleh...bleh. and I'm hungry.

\*Meantime looking at getting jojo access to the wiki...no idea how we give her editor

priviliges...probably have to log in as admin on the wiki page and give them to her....but i dont know that password.

>asked patrick for it

>ALSO can google media wiki user permissions to find a page that explains realtively how to do it..sort of.

-->store certs at /etc/grid-security/hostcert.pem and hostcert.key

9/7/10 (tom's your mom's bday!)

\*failing SAM tests, never got a cert from them. resubmitting request

-whooo okay so the new certs are in for the SE, working on getting the webservice certificates too.

9/8/10

\*today's goals: install the webservice certs which got emailed to me late last night and begining the vdt upgrade

\*webservice certs same website as host certs:

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

-the cert name and key didn't match so for the time being i set the host one as the http one and requested a new one.

-should be all correctly set up. will see if this fixes the SAM jobs

\*SE!

-instructions at ... <https://twiki.grid.iu.edu/bin/view/ReleaseDocumentation/OSG12UpdateInstructions>

-vdt located at ...

/sandbox/osg-se/vdt

-output of final command :

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

enabling cron service fetch-crl... ok

```

skipping cron service 'vdt-rotate-logs' -- marked as disabled
enabling inetd service gsift... FAILED! (see logs/vdt-control.log)
  found conflicting, non-VDT entry for service gsift in /etc/services;
  use the --force option to remove the entry
skipping cron service 'edg-mkgridmap' -- marked as disabled
enabling cron service vdt-update-certs... ok
enabling cron service gums-host-cron... ok
enabling cron service gratia-gridftp-transfer... ok

```

\*looked at condor\_q and a bunch of jobs were held because of a stand out error that i assumed was maybe from turning vdt off during the upgrade. so i set them back in the run with condor\_q | sgragai since it was all of samir's jobs

-the sam tests are giving what might be different error reports. ugh

9/9/10

\*Make sure to sync this doc with the one stored on my external drive

\*Failed Sam tests:

-renewed the certificates but still getting an error looks like its a different one though

\*Failed RSV tests:

-still failing, despite upgrading the VDT

9/10/10

\*SAM

-looking at patricks documentation for certificates and updating accordingly.

-copied certs to root user .globus, changed permissions

-replaced containercerts with new certs, changed permissions

-all four (host, container, rsv, and http) certs are fine on the frontend

-3 that exist on the se are good.

found the config.ini file is at

/sandbox/backups/configfiles/system-config/Config/OSG

but its in a backup file?

-restarted phedex with \$ ./PHEDEX/Utilities/Master -config

SITECONF/T3\_US\_FIT/PhEDEX/Config.Prod start

where . = /sandbox/phedex

-got a reply from the GGUS telling me to submit the problem to Savannah instead with a link. So i did that.

-need to know new bestman version so can update th config.ini did less /sandbox/osg-se/bestman/version

\*Hours, for this week were all normal except friday was 10-1;

\*Condor Queue, plenty of jobs in queue but some machiens are running nothing as seen on ganglia

-confirmed when run condor\_status

-following info at

[http://www.cs.wisc.edu/condor/manual/v7.2/7\\_3Running\\_Condor.html#SECTION00835000000000000000](http://www.cs.wisc.edu/condor/manual/v7.2/7_3Running_Condor.html#SECTION00835000000000000000)

command condor\_q -analyze shows every job and whether or not it got assigned and how many nodes were available to it

9/13/10

\*time for a resync.

\*sent email to condor\_list about jobs problem

\*response from savannah about failed SAM tests

"

seems like your SRM head node is not accepting network connections. Likely something basic in OSG setup.

You could take an example SURL from the lcg-cp test e.g. and try

to use it as destination in an SRM copy operation, first from your site, then from a remote one.

"

example SURL is: srm://uscms1-se.fltech-grid3.fit.edu:8443/srm/v2/server?

SFN=/bestman/BeStMan/cms/store/unmerged/SAM/testSRM/SAM-uscms1-se.fltech-grid3.fit.edu/lcg-util/testfile-cp-notoken-20100913-160400.txt

srm copy is: srmcp flags source destination

Anyhow I suggest you contact USCMS T3 support, as they have a program to help T3 in this kind of situation, e.g. via hn-cms-osg-tier3@cern.ch

-not really sure if this will do anything but i added source /sandbox/osg-se/setup.sh to /etc/profile...via vim.

-Srm\_path is not set...could be the problem?  
-permissions don't let others in?

\*look into permissions?

\*when i came int here was a warning about a new cmssw version but it updated either by itself or via samir. either way should find out...

\*start attending Tier3 Evo meetings! Tell Jojo about this too

9/14

Trying to follow the condor suggestion:  
/opt/condor/etc/condor\_config:

We don't need glite-ui unless we want to submit jobs to EGEE!!! It wouldn't work anyway, since we don't have myproxy installed yet.

Fixed the problem with CRAB by adding the following to /opt/condor/etc/condor\_config:

```
GRIDMANAGER_MAX_PENDING_SUBMITS_PER_RESOURCE = 1000
GRIDMANAGER_MAX_SUBMITTED_JOBS_PER_RESOURCE = 1000
```

And changing scheduler to condor\_g in crab .cfg file

Now jobs submit and run successfully.

So all we have to do is update CRAB to get new features.

<https://twiki.cern.ch/twiki/bin/viewauth/CMS/CrabClientRelNotes273>

found in patricks spring 10 documentation

## Where is the machine-specific local config file for each host?

```
LOCAL_CONFIG_FILE = /opt/condor/etc/condor_config.local
```

9/15

\*Yesterday I emailed the tier3 list and got a response from Brian

-he thinks either bestman is off or firewall issues with port 8443

\*Turn bestman on

```
go to VDT_LOCATION (/sandbox/osg-se)
. setup.sh                (command to initialize the vdt environment)
vdt-control --off
vdt-control --on
```

\*Firewall issues...

directions at <https://twiki.grid.iu.edu/bin/view/ReleaseDocumentation/ComputeElementFirewalls>

\*Doug Johnson also emailed me

-he thinks i should reboot the se and then run those same commands that brian recommended

-rebooting the se, is shutdown -r ok?

waiting for patrick

-also gave doug a new password so he could log in and check things out.

-he asked for sudo privileges, ask dr. hohlmann

\*looking at directions on

[https://twiki.grid.iu.edu/bin/view/ReleaseDocumentation/BestmanGateway#Services\\_Startup\\_Shutdown](https://twiki.grid.iu.edu/bin/view/ReleaseDocumentation/BestmanGateway#Services_Startup_Shutdown) with regards

to testing SRM environment. when i execute

```
srm-ping srm://uscms1-se.fltech-grid3.fit.edu:8443/srm/v2/server
```

note must export export PATH=\$PATH:\$VDT\_LOCATION/bestman/bin  
first!

the connection is refused. so it really does look like it might be the port.

but running `ps auxww|grep $VDT_LOCATION/bestman|grep -v grep` should show if bestman is running and i get an odd output:

```

    root      8613  0.0  0.0 248208 54688 ?        S    Sep10   0:14 perl
/sandbox/phedex/PHEDEX/Toolkit/Transfer/FileDownload -state
/sandbox/phedex/Prod_T3_US_FIT/state/download-srm/ -log /sandbox/phedex/Prod_T3_US_FIT/logs/download-
srm -db /sandbox/phedex/gridcert/DBParam:Prod/FLTECH -nodes T3_US_FIT -delete
/sandbox/phedex/SITECONF/T3_US_FIT/PhEDEx/FileDownloadDelete -validate
/sandbox/phedex/SITECONF/T3_US_FIT/PhEDEx/FileDownloadSRMVerify -backend SRM -syntax bestman -
protocol srmv2 -batch-files 1 -jobs 5 -timeout 86400 -command /sandbox/osg-se/bestman/bin/srm-copy, -
proxyfile /sandbox/phedex/gridcert/proxy.cert,-3partycopy,-delegation true,-mkdir,-debug,-
retry_num=2,-dcau false,-pushmode,-parallelism 1,-connectiontimeout 7200

```

also a connection timeout...so it could still be the port

\*Emails from Doug:

Bestman didn't know to use gums so reconfigured its script with the one found in the Storage Element Documentation via Patrick.

However even after forcing that change by typing it manually into the file (running the script didn't change it) bestman still isn't working. told Doug about my attempts waiting for a response.

he suggested quotation marks around two of the lines....

`vdt-control --off`

```

$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:2811"

```

`vdt-control --on`

gave me a different error when I ran `vdt-control --on`

```

enabling cron service fetch-crl... ok
skipping cron service 'vdt-rotate-logs' -- marked as disabled
enabling inetd service gsiftp... FAILED! (see logs/vdt-control.log)
found conflicting, non-VDT entry for service fsiftp in /etc/service
use the --force option to remove the entry
skipping cron service 'edg-mkgridmap' -- marked as disabled
enabling cron service vdt-update-certs... ok
enabling cron service gums-host-cron... ok
enabling cron service gratia-gridftp-transfer... ok
enabling init service bestman... ok

```

also a new sam test started failing. I'm of the opinion that its a result of this latest change. Especially since apparantly rerunning the bestman\_config created a config file with bestman in FULL mode instead of GATEWAY which is what was created last time.

9/16/10



\*Condor probs, samir's jobs arent running, some jobs have been running for over 30 days.  
Just going to remove them....?

\*to remove condor jobs do condor\_rm jobid. if want to delete a range, leave off the decimal or use \* to fill in.

\*also did a whole new bestman config but its the not the primary one right now.

9/17

\*SAM Tests, using our new installation as our primary... first reading post-install/README  
-first of all:

on SE \$VDT\_LOCATION is available after source /sandbox/osg-se/setup.sh

on CE '\$VDT\_LOCATION is available after source /opt/osg-1.2/setup.sh

NOTE if it doesnt work on CE, log out completely and try again. no idea why it changed but it did for me.

-second of all to test RSV Probes

su rsvuser

source /opt/osg-1.2/setup.sh

cd ~

./RunAllRSVProbes

ba bam! new installation of bestman works! passing sam tests passing RSV tests! have got to learn to phedex!

9/21

<http://dashb-ssb.cern.ch/gridmap-ssb/gridmap.html?>

[http://dashb-ssb.cern.ch/dashboard/request.py/siteview?site=T3\\_US\\_FIT](http://dashb-ssb.cern.ch/dashboard/request.py/siteview?site=T3_US_FIT)

<http://dashb-ssb.cern.ch/dashboard/request.py/siteviewhome>

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

9/22

Printer Conference with Samir and Himali.

seems pretty straightforward, instructions in My computer/ HEP/printer.doc

basically you type telnet and the ip address of the printer. then the password and then ? to see list of commands but basically only use allow, clear, and quit. :)

NOTE need to be on a computer that has access to the printer in order to telnet to it. this is why the cluster is on the list of allowed ip addresses

\*reading things about condor....

nothing looking good.

\*dealing with the Security Bug

[https://bugzilla.redhat.com/show\\_bug.cgi?id=634457](https://bugzilla.redhat.com/show_bug.cgi?id=634457)

plan: install the kernel patch once i know which one to use

NOTE: can just do yum install kernel and then reboot and it will automatically pick the right one.

have to run it on the CE and SE.

REbooting 101

A! you do not have to take everything down unless there's going to be a power outage. in this instance for example could only have rebooted the CE and SE. i think.

ugh

so did cluster-fork /sbin/init 0

ssh dev-0-0

/sbin/init 0

didnt do nas cause i needed a password i didnt have

then went down to high bay

press power button on ce and se. Se is hte one at the very bottom of rack to the left.

then look for the one labeled ups. the one on top of it is the NAS. the one on top of that is the CE.

then after hitting those power buttons. the circle ones.

hit the power buttons on the nodes, wait to do this stepp until after youve seen you can log into the cluster.

okay compute nodes

1-4

1-9

2-7

don't show up on ganglia. and their power lights don't come on.

Used the reset button instead of the power button and everything was fine. so its all up now.  
emailed at the beginning and end. Note to self FIX your email folders so you learn about these bug  
hacks when they're happening not last minute

9/23 (13 months!)

Let's reevaluate

we dealt with the security bug last time

sam tests are passing as are rsv tests

ganglia looks about the same which is a problem but not a new one

\*on ganglia, condor\_top isn't working or the condor\_Q or other gmetrics.

-on cluster condor\_q works fine, same amount of jobs are running as always (114)

-so obviously we have some kind of issue with condor.

\*Himali's certificate, Rick emailed me asking to validate him. I sent him a signed email and stated  
that i was validating himali.

that was all two days ago. Himali still hasn't heard anything so i sent another email making sure i  
wasn't supposed to do anything  
else.

\*the money finally came in yesterday so i need to start checking quotes. Bryant and Will both have  
some opinions but

mainly on the high quality visual station not on the cluster stuff. granted no point in buying new  
nodes if a third of the old ones are never running

Nodes that are never running anything:

nas-0-0

dev-0-0

uscms1.local

compute-1-1

compute-1-2

compute-1-8

compute-2-1

compute 2-2

compute -2-5

it is very tempting to just go down and try resetting those nodes.

but first the elegant solutions

START=TRUE

in your \$CONDOR\_HOME/local.\$HOSTNAME/condor\_config.local file and restart the nodes.

can't find a folder like that but

/opt/condor/etc/condor\_config.local exists. no start variable in it but there is one in

condor\_config. i think i could edit the file and reboot just one of the nodes and if it fixes th  
problem do it for the others. question is if it would be necessary to do to all of them.

\*phedexing guide

\*printer doesnt print nicely

so basically right now.

3 issues

\*Ganglia top/queue

\*condor appropriation of jobs

\*30 K buy some stuff

\*Ganglia top/queue

[http://ganglia.info/?page\\_id=68](http://ganglia.info/?page_id=68)

[http://sourceforge.net/apps/trac/ganglia/wiki/ganglia\\_configuration](http://sourceforge.net/apps/trac/ganglia/wiki/ganglia_configuration)

will give me lots to do tomorrow :)

9/24/10

4 hours lets get to work!

\*GANGLIA PLAN

-first look quickly through pat's doc if there's something that has to be re run after a shutdown to  
get it working

-look on ganglia website

-send an email to the ganglia list

-Patrick's Documentation

To update ganglia, restart /etc/init.d/gmetad and /etc/init.d/gmond  
MIGHT have to restart httpd as well.

hmmmmmm but we dont need to update just start it?

went ahead and tried it...FREAKED ME OUT. two nodes went down and so did the nas according to the ganglia website. but they came right back up and i dont think they ever shut down i think it was just that ganglia couldnt access them and only checked those three for the two seconds it was down? ugh scarrrrrrryyyyyyyyyy NOTE gmond and gmetad but return the same status they had before i did this : gmond dead but subsys locked. also top shows that gmetad is still getting run like aevery 20 seconds. im not sure if thats normall.....ugh  
sort of tempted to try restarting httpd...im gonna wait a sec

-ganglia website

there's a ganglia command! thought it doesn't seem to be what i want. basically gives you the same info you should be able to get from the web.

web stuff for the three things not showing up are located in /var/www/html/ganglia/addons/rocks so basically i feel like its not runing these three php scripts?

-emailed the ganglia list.

I did. now its a waiting game.

\*Email from Will: Can you setup an additional SVN repository for g4hep account. Currently only my account wbittnr has a repository on /export/svn .

added g4hep to the svnusers grp with chgrp. and changed permissions on a new folder i created in the svn directory called g4hep so now its owner and group are svnuser and both have permissiosn to read write execute

\*Condor

-if i run condor\_status it shows that nothing is running on the nodes that show nothing is running. does this mean condor isn't even running on these nodes?

-NOTE on ganglia looking at a particular node's graphs intently. there was no memory free on the nodes that were down until wednesday. (after rebooting they were back up) very intersting.

-I'm going to try the email suggestion, edit the config file. then restart only one node and see what happens...

check on new kernel..

9/27/10

\*Condor, lets actually try it this morning, obviously haivng a lot of trouble finding out which file to edit, make sure to make backups of everything.

/opt/condor/etc/condor\_config.local exists.

/mnt/nas0/condor/condor-7.4.2/condor\_config.local but this points at the same file as above. so this is two places for the same thing better to edit the one pointed to.

other option is /mnt.nas0/condor/condor-7.4.2/etc/condor\_config or condor\_config.local.

just realized /opt/condor points to /mnt/nas0/condor/condor.

ergo

path is

/mnt/nas0/condor/condor-7.4.2/etc/

Plan:

back up the local one. add the variable. reboot one of the nodes. see if it changed anything.

if no change put it back the way it was and edit the global one

if no change try both of them simultaneously. deal?

deal.

okay just edited the local file, rebooted compute-2-2 with reboot, and it looks to have come up fine but so far no change inganglia. granted it might take a little time even if it was fixed...you know? considering doing each option on a different compute node.

okay if you do condor\_config\_val -v <VariableName>

you can see which config file is setting that variable. In this case it turns out START=TRUE in /mnt/nas0/condor/condor-etc/condor\_config.cluster.

looking at that file its justa set of variables with values specific for our cluster. i'm going to respond to the email with my results.

-emailed other guy about the ganglia problem too.

-subscribed to rocks user list

and am sending them my previous ganglia email

\*from HEP Meeting

-get a test job from Ben and try forcing it to run on a particular node. specifically one of the broken ones

-email tier 3 with help for the condor issue

-copy over the files from Purdue Tier 2 that Dr. H wants.

note phedex isnt able to do it because theyre not globably available on DBS but samir found a workaround.

9/28/10

\*job from Ben, try to assign to one node

Here's a Condor job to run.

It needs to be run from

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

Run it as

condor\_submit submit.sub

The output goes to

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

That's it!

\*copy over files for Dr.H via Samir's method

xrdcp -R -f

root://cms.rcac.purdue.edu//store/user/aeverett/ExoMu100824/DiMuonSkim3/aeverett/Mu/ZPmu\_140133\_143178/f79794c1c58d077b23eddaf7823a8d3b/ .

form: xrdcp -R -f <Full Source Path> <destination>

NOTE full source path can be a directory.

to find them on DBS i think use :

find dataset where dataset like \*aeverett\* and dataset.status like VALID\* and release like CMSSW\_3\_7\_0

\*find out if up to date with cmssw

\*email lists with varius problems

-condor node issue

two commands given but regardless of which node i put in to test them a running or not running one i get the same output need to check it out more tomorrow when i have more time :(

9/29/10

\*Condor

-looking at burts email did condor\_q -bet which is basically the same thing as condor\_q -better-analyze.

realized that its not actually c hecking those nodes for a place to run the job at all. Its only checking 112at any time but it should check 160 (thats 8 CPUs per Node of which there are 20 labeled 1-0 through 1-9 and 2-0 through 2-9.)

similarly looking at various outputs of

condor\_status -format '%s\n' Name or whatever other attribute you want listed instead (for a complete list use condor\_status -l)

never outputs the names of the one that have nothing running on them. so to me the problem seems to be that condor

isnt running/checking these machines at all even though they're up. thus i want to try to readd them or whatever to condor's stuff. not really sure how to do that and so the documentation search continues. also i emailed burt with these results so we'll see what he has to say.

-burt replied already!

try condor\_status compute-1-1

If it doesn't show up, it means that the condor daemons on compute-1-1 are not talking to your central manager.

it didnt.

Check that they're running,

did it with ps -ef | egrep condor\_

daemons show up on working nodes but not on non working nodes.

-- if they are, take a look at the condor StartLog on compute-1-1 (you can send me a copy if you want).

took a look anyways they all stopped recording on Aug 25. which looking at ganglia is approx me when this became a problem. although i felt like it was older because i thought it preemtped my return to FIT. what did i do around that time to break it?

(renewed the phedex proxy and was fixing the SE thing...seems unlikely either one disconnected the condor daemons for 1/3 of the machines)

-tried /sbin/service condor start on the compute nodes as well as /opt/condor/sbin/condor\_master neither worked

log files on cluster are at /opt/condor-7.2.0/log  
these dont appear on a node but

\*COpying over files

the command doesnt work for me  
ask Samir about it.

\*email ben asking for a small job to test

got a reply with all the instructions, see ben's email dated 9/10/10

9/30/10

Tried dougs advice but still no real progress sigh....ugh.

emails back and forth.

also to see how much space and or memory any directory is using do

sudo du -sh <path>

RIght now Nas is a 10 TB machine, samir is using 4.1 TB

10/1

\*Doug thinks the problem with condor is because the scratch files are too full, they're at 100% so i need to clean that area up.

removed a directory from compute-1-1: /scratch/condor/execute /26380

Note i copied it over to my hnome directory on the frontend first just in case. but i think these directories can be deleted because they belong to whichever user ran the job that rcreated them maybe.... in this case it was owned by idiaz group: users permissions drwxr-xr-x

Dr. H and I had a quick discussion. We agreed that its probably safe to delete all the cms\_ files in the /scratch/wn\_tmp/ so i did and then restarted condor with

/sbin/service condor start

now on ganglia its looking better, its got jobs running on it on condor\_q and the totall number of jobs is 124 so i think we're good!

i'm going to repeat on the other failing nodes.

so far did it on 2-5, 2-1

note on ganglia 4 are RED: 1-9, 2-0, 2-8, 1-6

new 4: 1-3,1-9,2-6,2-7

FINGERS CROSSED

but ive got to run to class

10/4: 10-4 ! get it?

\*so condor problem is officially solved. We've increased productivity by practically 50% not sure if the dev, nas, and "local" should also be running but i doubt it.

\*Big issues now are the shopping list and fixing the functionality of ganglia so we can look at the condor\_q on there. also volunteered myself to write a Cluster Guide. so woo hoo for that.

\*Been searchin g the rocks archives for help with the tools on ganglia. people have had the issue before but no working solutions for us yet.

umkay magically solved it with an instruction....

sudo service greceptor restart on front end

sudo tentakel -g compute 'service greceptor restart'

except for the job queue the other two functionalities are working

\*volunteered myself to write a troubleshooting guide for the clsuter

\*will be training latest apprentice, Kim

10/5:

Spent my hour searching for a solution to the ganglia issue and once again trying to copy something over with xrdcp.

10/6

Worked with Johanna Laina for two hours, looked at silicon mechanics discussed the upgrade. She emailed silicon mechanics and set up a conference call to discuss it further

Worked with Kim for an hour talking about related cluster stuff, got her started on getting her grid cert

10/7

Working on ganglia problem...no progress yet.

10/8

Worked with kim setting up her account but having a problem doing it, not able to create home directories asked a question on linux.org got lots of advice but nothing has been working, wondering if its a space issue.

10/14

Samir has a crab issue I'm looking into. here's the text from his email:

Can you please look for the glite scheduler in crab since I get the following error with glite?

crab: Error: the EDG\_WL\_LOCATION variable is not set.

condor\_g scheduler is running good. My analysis was not running because of different schedulers. We have condor\_g but the host site has glite.

So, if you can make both schedulers running, we will not have similar issue again.

also the dimuon skim dataset was uploaded globally and i approved the phedex transfer will see if it worked...

10/15

SO big emails been going around about security practices and hardware upgrades. have been way behind on my reading because of my work packed fall "break". So ....

#### \*\*Security Practices:

-give grid users as few privileges as possible

-Gums..Check

-pool accounts. not sure?

-Gums on a dedicated machine, the SE i believe..ports blocked? logins limited? we use hosts.allow and hosts.deny.

-----go see if can find that command from Doug about how to see which ports are open and log into se and try it.

-cron turned off..why?

-need to el

#### \*\*Phedexing.

the phedexed over dimuon skim has not appeared on our site... im tempted to think its a memory problem? but if it failed wouldnt it give us an email?

at <https://cmsweb.cern.ch/phedex/prod/Activity::TransferDetails>

you can see that we're transferring something..hard to tell what columns mean and stuff but it looks like were in good company so?

#### To Do:

\*Problem with creating home directory

\*Phedexing of file

\*Cluster Upgrade

\*Security: Practices and indiv email

\*CRAB: glite issue

\*Ganglia

\*Training of Kim

\*Cluster Guide/Troubleshooting

so need to give Kim a project to get her started...

1) Upgrade of cluster stuff just to keep her informed

2) Learn about IPtables? Might be a good project since one of us really should be able to do it.

\*CRAB: glite..

so from Patricks documentaiton he updated CRAB but didnt setup the Gliteui (user interface) because "it was too much work without benefit to justify setting it up."

<https://hypernews.cern.ch/HyperNews/CMS/get/osg-tier3/249/1.html>

this CMS message claims that with crab 2.7.x (we have 2.7.3) we no longer need glite UI for it to work

10/18/10 1-2, 3-6

\*Home Directory:

on linuxquestions.org posted results of strace mkdir /home/kday, strace useradd kday, and mount. we'll see what they have to say

\* Phedexing of data

Request it to tier2 and then from tier2 to us. not sure if i did it right, going to ask samir.

\*cluster upgrade look for that 36 drive nas, get the quote from dr. h with all the details.

10/20/2010: 10-2

\*Himali's data

he accidentally deleted everything, looked for a backup but it looks like the only one we're backing up is g4hep. on the SE /sandbox/backups/userdata. asked patrick just to make sure

\*Phedexing of data not working because proxy expired

Command for checking proxy.

```
openssl x509 -in grid.cert -dates -noout
```

in /etc/grid-security/certificates checked dates all the way up to

for those certificates since it gives us a crl error and some random place onthe internet pointed me there.

looks like all the ones on the se are outdated! but why?

to update the CA CRLs do

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

this will not only update them but add the automatic updating of them to the cron tab, this had to be done on the SE as root.

```
~~~~~
if want to just manually update them do...
```

```
$VDT_LOCATION/fetch-crl/share/doc/fetch-crl-2.6.6/fetch-crl.cron
```

Here is the full error report:

```
[phedex@uscms1-se gridcert]$ ./proxyrenew.sh.forpatrick
```

```
voms_proxy_lifetime: 0
```

```
###WARNING: Wed Oct 20 11:52:50 EDT 2010, need renew the proxies.
```

```
Your identity: /DC=org/DC=doegrids/OU=People/CN=Xenia J Fave 515293
```

```
Creating temporary proxy
```

```
.....
Done
```

```
Contacting voms.fnal.gov:15015 [/DC=org/DC=doegrids/OU=Services/CN=http/voms.fnal.gov] "cms" Failed
```

```
Error: Could not establish authenticated connection with the server.
```

```
GSS Major Status: Authentication Failed
```

```
GSS Minor Status Error Chain:
```

```
init.c:globus_gss_assist_init_sec_context:277:
```

```
Error during context initialization
```

```
init_sec_context.c:gss_init_sec_context:197:
```

```
SSLv3 handshake problems
```

```
globus_i_gsi_gss_utils.c:globus_i_gsi_gss_handshake:910:
```

```
Unable to verify remote side's credentials
```

```
globus_i_gsi_gss_utils.c:globus_i_gsi_gss_handshake:883:
```

```
SSLv3 handshake problems: Couldn't do ssl handshake
```

```
OpenSSL Error: s3_clnt.c:894: in library: SSL routines, function SSL3_GET_SERVER_CERTIFICATE:
```

```
certificate verify failed
```

```
globus_gsi_callback.c:globus_gsi_callback_handshake_callback:547:
```

```
Could not verify credential
```

```
globus_gsi_callback.c:globus_i_gsi_callback_cred_verify:855:
```

Could not verify credential  
 globus\_gsi\_callback.c:globus\_i\_gsi\_callback\_check\_revoked:1158:  
 Invalid CRL: The available CRL has expired

10/22/10 10-12, 1-2, 2:30-3:30

ToDo:

\*Cluster Memory is 100% incidentally the condor q is at like 2000 jobs.

not sure if all of that is working...

\*One sam-rsv test is failing..

sam test: sft-job

From reading the error report it says that the job was held and the the job proxy expired.

My theories are as follows:

1) so many jobs in the queue that the sam test didnt get run before the job proxy expired

2) it was somehow related to the 100 percent memory issue

samir just deleted 2 Tb of data so were back to 84 so if this was the case the problem will fix itself once the test is rerun.

\*home Directory issue;

PLAN:

\*Jojo is going to try and take care of the home directory issue

electron and photon are having problems and IT is in charge of that with jojo as the security contact

\*I am going to solve the SAM problem.

\*I am going to fix the ganglia issue

OSG-Tier3 10/22

Security Meeting:

-Turn off services:

UAF rep talking

use nomachine

open ports only as necessary

a very restrictive default

flipside: one machine without a firewall bc of glidein infrastructure

CRON:

very restricted.

use cron.deny to deny

userlevel cron is disabled on worker nodes

don't do automatic yum updates on worker nodes

bc couldnt depend on it working properly (bc of dependencies)

and bc want to know about things that are downloaded

. instead uses rocks to apply patches

yum automation is alright just needs to be monitored. shouldn't be automatic

if Hadoop and workernodes become storage and computing...

yum updates can be very intensive.

-use a proxy

-implement a local repository

might have a hadoop talk

asked tier3s about what services we get from campuses.

also good to talk to OSG security team bc they really keep track of all of it and

they can give good feedback.

community support:

this meeting was recorded so if that worked we can watch the beginning.

\*COndor q some jobs have been running for nine days. sending pzuo an email.

10/25 1-2, 3-6

\*Sam tests not running

\*email curtis/john

\* priorities

\*cluster upgrade

\*glite Scheduler



\*Emailed Curtis and John about Useraccount. JOJO also emailed tier 3 and they had some input. there may be a special script that Patrick had to do this.

There is! its in the root home directory i.e. sudo su root, cd ~, ./UserCreationScript and gave her user privileges

\*Priorities: there is no documentation that I have on this...so to the web! probably governed by condor...

[http://www.cs.wisc.edu/condor/manual/v7.5/2\\_7Priorities\\_Preemption.html](http://www.cs.wisc.edu/condor/manual/v7.5/2_7Priorities_Preemption.html)

hurray!

okay so with condor\_userprio you can see the priorities for different users.

\*Photon:

to log in to tracks must do FLTECH\

\*Sam Tests

Detailed email went out to Tier3

\*glite

10/26 5-6

worked on setting up lab for laser day.

one of the computers (electron) came back from IT with Windows 2000 professional.

10/27: 11-2

Led groups through the high bay and talked about cluster

\* a sam test was failing this morning on the SE but fixed itself.

\*still not having the CE sam tests running.

got an email from the group saying that the queue could be causing the problem.

10/28

10/29

\*built computer with kim bryant and mike

\*worked with jojo that afternoon on setting the lab up better

11/1

\*Memory at 100% ...most likely from the huge phedex transfers. I'm deleting everything under 2010a. Himali transferred the entire RECO directory instead of just the skims

deleted files and we're back to 63%

\*Sam tests run every 30 minutes according to Andrea. We are failing three now.

Looks like all of them are failing because they cant download something i.e. bc we were full? so we'll see if that changes in the next hours.

Could organize to run sam tests on frontend but according to andrea(the person who actually makes them run) its very complicated and most sites just opt to give the account to which SAM tests belong a very high priority so they just run immediately.

\*condor q only running 90ish jobs even though 400 are in the queue all jobs are run ont he same compute nodes it looks like. the ones that are doing nothing...

1-9, 1-0, 2-8, 1-3,1-1,2-0,1-8.2-9.1-2,2-6,2-1,1-4,1-5,2-2

\*Sam tests last ran at

mc: 10/30 14:33

sft-job:11/01 11:36

cms-prod: 11/01 11:21

11/2

Reorganized the whole office

set out an email for samirs jobs

writing an email about the sam tests

11/8/10

needed to update CRAB version.

Crab tarballs and directories are stored in /mnt/nas0/OSG/APP/crab

got the tarball from <https://twiki.cern.ch/twiki/bin/view/CMS/CrabClientRelNotes275>

scp it over to the cluster and to the directory mentioned above.

then changed owner/group to uscms01, and su over to uscms01 then did

```
tar -xvf CRAB_2_7_5.tgz
```

```
cd CRAB_2_7_5
```

```
./configure
```

Himali is checkign it now.

He says it works.

we'll see if that magically fixes the SAM issues as Dr. H thinks it might.

in other news I think ganglia is having issues reporting things...asked jojo to make this her personal task since the only other thing she is working on really is the cluster update which basically has boiled down to talking to dr. hohlmann about the decision.

kim has been tasked by Dr. H with setting up accounts on Higgs (the new three monitor machine) as well as putting ganglia/sam tests/condor\_q work ing on the tret of the time.

new note:

Kondo needs SLC4 no SL4 on the big machine...

steps: see if its working, see whats on it, see if needs to be updated. if no good then put it on the new mystery machine.

11/9

Some of Himali's jobs completed not sure if they were successful. four are still running something like 6 are in held mode

looking at the responses to failling sam tests

1) iptables issue

referring to....

[http://www.linuxhomenetworking.com/wiki/index.php/Quick\\_HOWTO:\\_Ch14:\\_Linux\\_Firewalls\\_Using\\_iptables](http://www.linuxhomenetworking.com/wiki/index.php/Quick_HOWTO:_Ch14:_Linux_Firewalls_Using_iptables)

and [https://twiki.grid.iu.edu/bin/view/ReleaseDocumentation/ComputeElementFirewalls#Firewall\\_Configuration\\_using\\_ipt](https://twiki.grid.iu.edu/bin/view/ReleaseDocumentation/ComputeElementFirewalls#Firewall_Configuration_using_ipt)

2) clocks

used

date and hwclock --show to get results

they matched up.

11/10

replied to Bill about iptables and clocks about sam tests.

Going to ask Kim to finish up setting up the desktop on Higgs and then hopefully also install SLC4 on Kondo's Machine.

adding condor\_status to daily commands. okay so taking the first one as an example compute-1-0 has eight claimed and buy cpus some of these have a load avg of 0 others have a load avg of 1. and roughly adds up to what i see on the ganglia page. so ganglia is reporting correctly>

attempting to put backup file on our little backup drive but of course its needlessly complicated. gave up. its now saved on my udrive. woooooo aka xfave2008@olin.fit.edu

11/12

Update crab again with patch: CRAB\_2\_7\_5\_patch1

and Kim is installing SLC 4: so far has everything plugged in and machine i

ALSO! problem with sam tests may be because of certificates.

(from an email from Doug )

"So this looks to me like the certificates and CRLs are not being kept up to date. I assume there is supposed to be a cron job some place that updates these certificates? Have any of your servers changed recently? It looks to me like these are a different copy than is used for the CE installation."

okay so i think i renewed crl certs a little earlier in the semester so just loooking at those... vdt-control --on fetch-crl

going to send an email out to ganglia.... still having issues quantifying that problem.

11/15: 12:45-6:15 (5.5 hours)

\*i think i unpackaged the slc4 disk wrong. so am redoing it a different way.

also making a new ubuntu 10.10 disk because....we are turning the mystery machine into ubuntu and electron into the slc4  
 success on the making of a new CD. :) slc4 is happily installing.  
 ubuntu cd was also successfully made and will be the next challenge  
 everything looks good on the ubuntu end.

\*FOR sam tests where we believe updating the CRLs is the problem:  
 The following will not work:

```
uscms1> source /mnt/nas0/OSG/GRID/setup.sh
```

and

```
uscms> source /opt/sg-1.2/setup.sh
```

is the wrong VDT installation for the worker nodes.

What might work and keep the certificates up to date (run on uscms1):

```
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
```

\*Root password on Cluster

11/17: 11-2

\*on slc4 machine. if the screen goes to screensaver (locks itself) and then kondo puts in his password to wake it up. it doesnt recognize it. and he has to reboot completely. he disabled automatic screensaving, ill look at it with my own account when he isnt working on it.

\*Big issue is the ganglia/condor/node thing right now

-One node doesn't appear under condor\_status. so condor isn't running it or sending jobs to it. lets fix this first.

ps aux

lets you see which daemons are runing. Comparing those on 1-1 and those on 1-0 its obvious that condor daemons are missing from 1-0 so just need to restart them.

\*reorganized closet of computer goodies

\*talked to Dr. H about future NAS purchas.

11/19: 1-2, 3-4

\*Looking at two warnings from OSG security

```
find kernel with uname -r
```

```
distro with cat /etc/issue
```

First oen: we dont have systemtap installed thus we're free from vulnerablities

yum list installed, yum list extras, yum list obsoletes

second one: our version of openssl is also free of vulenerabilities

```
openssl version
```

```
OpenSSL 0.9.8b
```

11/22: 1-2. 2:45-6:15

\*removed all of samir and himali's held jobs from the queue. Himali says his crab jobs are now running.

\*we have a notice on a SAM test. apparantly missing a patch to a cmssw version. i got worried that there was a problem because 2\_7\_5 wasn't showing up but of ocurse that's because the crab versions are different from the cmssw versions. for future notice you can see which version of cmssw we have by c hecking the results of hte sam test: swinst.

normally bockjo ckeeps our cmssw versions up to date so he might fix this one on his own.

so for the moment going to leave it alone and work on the documentation. will send him an email if no fix by the end of the week. not crucial.

\*working on ubuntu 10.10 machine with Kim. made a new cd.

using her personal cd drive since no working ones in the lab.

It still never recognized it when we asked it to boot from the dvd.

so

we went into windows, it spun the disk, and we asked it to install it side by side basically then we hope to erase the other half.

11/23: 9:00-11:00, 4:45-6:45

Updating Troubleshooting guide. Finished putting in this semester's documentaiton.

how to remove all held jobs in Condor  
for a given user.

```
condor_rm -constraint '(JobStatus==5) && (Owner=="sanders)'
```

(assuming the username is sanders).

To Do:

- +TroubleShooting Guide
- +Update Wiki with guide and this semester's logs.
- +Get Photon back.
- +Get ubuntu on mystery machine
- +Fix node issue
- +Cluster Upgrade
- +Look into Ganglia

11/29

- +photon came back with no changes but otherwise working
- +Create a master pw list for dr. H
- done

- +activae ports

two ports to be activated.

Higgs has a seperate issue. it wont connect to the internet despite....

tried:

- original eternet cord to original port: blinks but no internet
- original ethernet cord to other port: blinks but no internet
- New ethernet cord to original port: blinks but no internet
- New ethernet cord to unactivated port: no blink, no internet
- New ethernet cord to known to be working port: blink but no internet

- +Look into Higgs webpages

Kim has already made them come up automatically. Key is to make the pages refresh and the screen not go hibernate

first have to fix the internet

+Higgs have verryone log int o gridops when they're done using it....so either grid ops should lose its admin priviliges or...we should make a different user that has the weppages....OR OR!!! just make its sudo password different from its regular one. except according to the internet that's impossible... boo.

- +SLC4 Lockdown thing.

updated xscreensaver by downloading the latest rpm and it fixed the problem

- +Photon, mouse doesnt work.

the mouse isnt the problem it worked find on previous computer.

a working usb mouse also didnt work with it.

odddddd

restarting the computer with the usb mouse plugged in worked. so we'll go with that.

- +Higgs

change so that all users can open up gridops without having to give out the secret password.

12/1/10

- +Himali's files

Himali couldn't delete some of his files because there was a space in their names.

I did. `rm Dilepton*` which then caught all of them hurrah.

- +Missing Node

condor is not talking to a node...rarrrrrrr

12/3: 4-6

- +Node

reboot it and see if that helps

+SAM test

this cmssw didn't magically fix itself so need to import my certificate into Higgs to debug it.  
sent an email to bockjoo asking he can just reinstall it

12/5: 2:30-6:30

+Sam test is now passing since bockjoo added the cmssw version. adding that to the troubleshooting guide

+Going to try and add all of patricks n JoJos stuff to the troubleshooting guide.

+Updated wiki with my documentation, jojos, and the troubleshooting guide