

Spring 2011 Documentation

~~~~~

1/7/10

```
-Sam tests failing () as are RSV tests on SE.  
Did ssh xfave@uscms1-se.fltech-grid3.fit.edu  
source /sandbox/bestman-1.2/setup.sh  
vdt-control off  
vdt-control on
```

fixed both.

```
-Upon reboot of cluster ganglia tools not working  
did sudo service greceptor restart
```

```
-Upon login cluster stalls and have to do ctrl-C  
I had added service greceptor restart to /etc/profile  
so it would automatically be run every time the cluster reboots.  
Removing this line from the file fixed the problem.
```

1/10/11 10-12

```
-Himali said everything is working :)  
-Very low load at the moment showing on ganglia  
-cluster upgrade is priority one probably.  
-emailed bockjoo  
-compute-1-0 not running condor
```

1/11/11 12:15-2

```
-Wrote emergency cluster documentation.  
-uploaded it to the wiki and changed almost all the links to pdfs so they open  
in the internet without opening onto your computer.
```

1/12 2-5

```
-Fixing Node problem.  
emailed tier 3 list, they said problem was its mounted read only.  
mount shows it as read write BUT
```

```
cat /etc/fstab shows it as read only.  
now learning how to remount as rw. with man mount
```

```
[xfave@compute-1-0 ~]$ sudo mount -o remount,rw /dev/sda3  
Password:  
mount: block device /dev/sda3 is write-protected, mounting read-only
```

so just made it worse...bleh.

```
$umount /dev/sda3 /scratch  
$mount -w /dev/sda3 /scratch  
$/sbin/service condor start  
Starting up Condor
```

Yay!

```
-reading up on installing NAS's  
-Want to know how to edit the cluster frontpage  
http://uscms1.fltech-grid3.fit.edu/wordpress/  
files are at /var/www/html/wordpress
```

1/14

-Problem with Phedex

from bockjoo and yujuns suggestion I needed to try doing a transfer by hand.

I did

srncmp

```
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/te
stfile-cp-notoken-20100913-160400.txt
```

```
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
```

and a couple other combos and got same error every time: srm client error:  
java.io.IOException: org.globus.gsi.GlobusCredentialException: [JGLOBUS-5]  
Proxy file (/tmp/x509up\_u0) not found.

Yujun so far thinks its an authentication error.

-sidenote, sudo nautilus opens up a file browser with full r/w permissions on everything.

1/19: 9-12, 2-5

-Working on phedex problem, could explain why we're not getting jobs?

yujun had me do...

srncmp -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
```

which didn't work even after first running grid-proxy-init

new command:

srncmp -2 -debug=true

```
srm://srmb.ihepa.ufl.edu:8443/srm/v2/server?SFN=/cms/data/test/srmb_test.2
file:///tmp/from_srmb.1
```

-2 Sam Tests failing and one warning on the SE (failed at 13:41, 14:19)

tests: (CE-squid, CE-frontier, SE-get-pfn-from-tfc (warning))

squid:

```
Error. squid http://uscms1-se.fltech-grid3.fit.edu:3128 is down,
unreachable or will not reply.
```

```
urllib2.HTTPError: Service Unavailable
```

Frontier:

Squid nogzip Access test for ECAL

START TIME: Wed Jan 19 08:41:19 EST 2011 ACTION: Squid nogzip Access test for ECAL

Wed Jan 19 08:41:19 EST 2011 --> running cmsRun ...

error [fn-htclient.c:392]: server error (HTTP/1.0 503 Service Unavailable)

proxy=uscms1-se.fltech-grid3.fit.edu server=cmsfrontier.cern.ch

warn [frontier.c:671]: Request 2 on chan 1 failed at Wed Jan 19 08:41:21 2011:

-8 [fn-htclient.c:392]: server error (HTTP/1.0 503 Service Unavailable)

proxy=uscms1-se.fltech-grid3.fit.edu server=cmsfrontier.cern.ch

warn [frontier.c:725]: Trying next server cmsfrontier1.cern.ch

error [fn-htclient.c:392]: server error (HTTP/1.0 503 Service Unavailable)

proxy=uscms1-se.fltech-grid3.fit.edu server=cmsfrontier1.cern.ch

warn [frontier.c:671]: Request 3 on chan 1 failed at Wed Jan 19 08:41:21 2011:

-8 [fn-htclient.c:392]: server error (HTTP/1.0 503 Service Unavailable)

```
proxy=uscms1-se.fltech-grid3.fit.edu server=cmsfrontier1.cern.ch
warn [frontier.c:725]: Trying next server cmsfrontier2.cern.ch
error [fn-htclient.c:392]: server error (HTTP/1.0 503 Service Unavailable)
proxy=uscms1-se.fltech-grid3.fit.edu server=cmsfrontier2.cern.ch
warn [frontier.c:671]: Request 4 on chan 1 failed at Wed Jan 19 08:41:22 2011:
-8 [fn-htclient.c:392]: server error (HTTP/1.0 503 Service Unavailable)
proxy=uscms1-se.fltech-grid3.fit.edu server=cmsfrontier2.cern.ch
warn [frontier.c:725]: Trying next server cmsfrontier3.cern.ch
error [fn-socket.c:227]: read from 163.118.206.162 timed out after 10 seconds
warn [frontier.c:671]: Request 5 on chan 1 failed at Wed Jan 19 08:41:32 2011:
-6 [fn-socket.c:227]: read from 163.118.206.162 timed out after 10 seconds
warn [frontier.c:718]: Trying direct connect to server cmsfrontier3.cern.ch
error [fn-htclient.c:392]: server error (HTTP/1.0 503 Service Unavailable) proxy=
server=cmsfrontier3.cern.ch
warn [frontier.c:671]: Request 6 on chan 1 failed at Wed Jan 19 08:41:32 2011:
-8 [fn-htclient.c:392]: server error (HTTP/1.0 503 Service Unavailable) proxy=
server=cmsfrontier3.cern.ch
error [frontier.c:729]: No more servers/proxies. Last error was: Request 6 on chan
1 failed at Wed Jan 19 08:41:32 2011: -8 [fn-htclient.c:392]: server error
(HTTP/1.0 503 Service Unavailable) proxy= server=cmsfrontier3.cern.ch
%MSG-s CMSException: 19-Jan-2011 08:41:32 EST pre-events
cms::Exception caught in cmsRun
---- Configuration BEGIN
Error occurred while creating PoolDBESSource
---- Conditions BEGIN
MetaData::getToken: metadata entry "EcalPedestals_30x" Can not get data
(Additional Information: [frontier.c:729]: No more servers/proxies. Last error
was: Request 6 on chan 1 failed at Wed Jan 19 08:41:32 2011: -8 [fn-htclient.c:392]:
server error (HTTP/1.0 503 Service Unavailable) proxy=
server=cmsfrontier3.cern.ch) ( CORAL :
"coral::FrontierAccess::Statement::execute" from
"CORAL/RelationalPlugins/frontier" )
---- Conditions END
---- Configuration END
%MSG
Wed Jan 19 08:41:32 EST 2011 --> DONE
CMSSW_frontier.sh: Error 73 from cmsRun
```

SE

```
Contacting webservice to perform LFN-to-PFN matching at URL:
https://cmsweb.cern.ch/phedex/datasvc/json/prod/lfn2pfn?node=T3_US_FIT&lfn=/s
tore/unmerged/SAM/testSRM/SAM-uscms1-se.fltech-grid3.fit.edu&protocol=srmv2&d
estination=T3_US_FIT&custodial=n
WARNING: Unable to open TFC webservice URL to perform LFN-to-PFN matching for Site
T3_US_FIT
No PFN retrieved from TFC webservice
Possible reasons are:
1) The site is not running a FileExport agent for the Prod instance of PhEDEx
2) The storage.xml catalog published by FileExport does not include a rule for
protocol srmv2
3) The rule for protocol srmv2 does not match to a valid PFN
WARNING: Retrieval of PFN from TFC webservice failed; cached PFN will be used
TFC Test successful
```

-Does editing the source code during a running job affect it?

-Why if Ben submits a large file does it time out and go idle: error log shows  
...

022 (392640.000.000) 01/18 02:30:41 Job disconnected, attempting to reconnect  
Socket between submit and execute hosts closed unexpectedly  
Trying to reconnect to slot2@compute-1-9.local <10.255.255.234:9822>

...

024 (392640.000.000) 01/18 02:50:41 Job reconnection failed  
Job disconnected too long: JobLeaseDuration (1200 seconds) expired  
Can not reconnect to slot2@compute-1-9.local, rescheduling job

...

Just got back from lunch and class.

The CE-squid and the se-lcg-cp fixed themselves. BIZARRRRRRRE

The srmcp is still having issues, I finally hit upon a winning combination from  
our CE to our SE

it was: srmcp -2 -debug=true

file:///home/xfave/testSRMfile

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/pf  
ordtestfile

emailed OSG warning team asking how to update glibc with minimum repercussions.  
-they claimed it VERY rarely ever creates incompatibility issues so i went ahead  
and downloaded it. we'll see

yum update glibc

talked to JoJo about cluster Update, we are stilllll waiting on email from Art  
with latest and greatest quote

1/20: 12-1

All sam tests fixed themselves. biggggggg mystery.

1/21: 12-1, 2-5

emails back and forth..

i think the phedexing issue is solved?

read up on installing NAS' haven't found anyone quite in our situation though

1/24: 2-5

TODO:

update glibc

email bockjoo/yujun combo about phedexing just to make sure

-approved a new phedex transfer from a different site and it Immediately ran into  
the same issue

-sent email.

go over NAS stuff with jojo.

also NOTE over the weekend the cluster did get busy, don't know who from yet but  
it wasn't our group.

unfortunately when I saw ganglia all lit up i didnt have access to termina  
so i didnt see results of condor\_q either. \*sigh\*

kim came in, we got the website onto the cluster website for monitoring it. she's  
working on getting the images to work

and also right now you have to do one extra click and we want to get rid of that.

Himali wants glite. it is a condor scheduler that allows you to run jobs at sites  
where the data is not.

i.e. makes everything faster.

1/25:12:15-2

-Worked with Mike and Bryant setting up the lab, the first three monitor set up is now all done, the tv is on the wall.

I was ready to try and play with the settings on the Beauty so that all three monitors would work but Dr. H

wanted Kim to do it so she would learn about some of those things.

-Emailed back from Yujun. About two hours ago the transfers started working again. Because they gave me write abilities?

I'm asking them for a confirmation of my understanding but one of the files from them along with the file from a different site

have both transferred completely and the other two are at half.

-Lenny's laptop isn't turning on so tomorrow hopefully we're going to put it in gluon and see if we can boot off it.

TODO

-glite!

reading up on it right now....

1/26: 2-5

-removed himali's jobs from the queue

-Talked over the entire cluster purchase with jojo, we believe we are ready. looked at getting a spares kit too.

1/27:9:30-11, 12-1-Cms ran some jobs last night

so our cluster looks semi busy at the moment

nodes with jobs: 1-3,1-8,2-3,2-2,2-5

45 slots are claimed and 58 jobs are supposedly running

-glite:

we have files called glite and stuff mostly under osg directories. not sure if its the same thing? or if its not activated?

-Kim/Jojo

right now Jojo's big tasks are to get the NAS order approved, and maintain the website especially getting the gallery thing set up

Kim's main tasks are to set up Beauty so that all three monitors work right as well as work on the website for cluster monitoring. I'm going to go ahead and make those pages default on Higgs until we can get a fancier version. She is also finishing building and setting up Truth.

we will all be involved with the NAS setup.

1/28:2-4

-EVO meeting

a new way to update the phedex proxy? Malina will email it out.

get a registered certificate personally for phedex?

-discussion on the srmcp

-tripwire reports were filling up his disks on his rocks cluster, how can you limit space it

takes up with reports. Malina turned hers off.

-security: denyhosts only good for ssh, ossec blocks at an earlier level, monitors log file (any access: http, gridftp, ssh, ). both block at the ip level so packets never make it past the ethernet cord

ossec is free and supported. if using rocks, make sure that whatever files it modifies gets distributed to all nodes, i.e. if modifies iptables then need to copy new iptables to all nodes

-met with dr. hohlmann, only a few more questions about the NAS.

1/29: 1:15-3:15

-Worked with lenny to try and recover data from disk.  
things I've learned, laptops are sneaky.  
sata is cool.

2/1:12:15-2

-Lots of jobs running: 136, ganglia is broadcasting tons of pretty colors as is the rsv chart which shows both RSV, OSG, and CMS jobs all having been run recently on us. Two of the nodes had stopped running condor though.

1-0 had the same problem as last time where for whatever reason /scratch had been mounted readonly so unmounted and remounted it.

```
$ umount /dev/sda3 /scratch
```

```
$mount /dev/sda3 /scratch
```

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

1-4's scratch directory had gotten too full so I erased some stuff from the wn\_tmp directory and condor directory and restarted condor.

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

1-0 already stopped working again, not even 5 minutes later...is there a proxy for mounted things?

there's that file that ren

2/3: 12:15-1

helped himali source his file

copied xorg files for truth and beauty to cluster, they in my home directory at the moment.

Looking at a variety of linux responses on the compute-1-0 issue. that will be the focus of tomorrow's work though since unfortunately i have class

2/4:11:45-1, 2-4,

Soooooo first response to linuxquestions.org suggested editing the fstab file to fix the issue with compute-1-0

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

and then unmounting the partition and remounting it with  
mount -a

so that it mounts all the partitions listed in /etc/fstab

as a result it remains rw but does

CRAB\_2\_7\_7 is out! installing now and done!!

need to install ffmpeg on cluster

[http://www.boonex.com/unity/blog/entry/How\\_to\\_install\\_ffmpeg\\_on\\_centos\\_5](http://www.boonex.com/unity/blog/entry/How_to_install_ffmpeg_on_centos_5)

instructions should work out but colloquim time :)

2/8: 12:15-2

2 sam tests are failing, both are related to srmcp.

also yesterday the internet was really "wonky" in OPS for everyone. Simultaneously we were watching SAM tests fail and fix themselves so probably related

but now the internet has been presumably fixed and these two are still failing  
16:40, 17:03, 8:28 time that the SAM tests and RSV probes were checked and found to be failing.

I restarted VDT on the SE

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

```
source setup.sh
```

```
vdt-control --off
vdt-control --on
got some errors because some parts couldn't be shut down by this command...
I just did, on the CE...
TO test RSV Probes
su rsvuser
source /opt/osg-1.2/setup.sh
cd ~
./RunAllRSVProbes
This fixed the problem last time along with a new bestman configuration so it can't
hurt to try it again.
While I wait for them to be refreshed I'm going to check to see if i can find the
emails back and forth for the previous time ~ 9/12/11
okay so last time it was definitely authentication errors and the fact that bestman
wasn't running
this time its not an authentication error, osg has no issues connecting to the
SE also bestman is up and running if i do ps -ef
and the bestman config is probably still okay this time?
the "weirdness" response to restarting VDT seems to be normal bc it happened last
time too.
```

okay the timestamps have changed on them and nothings turned green so that didn't solve the problem  
it looks like its either not mapping the user or us. blehhh  
composing something for the tier 3 group.

2/9: 9-3

determined to figure out the GUMS thing, Looking at Gums documentaiton, wil email group if i cant solve this by myself today. I just want to test if i'm worth my salt first. somehow didn't save my docuemntation for this day which is a shame...sight.

2/10: 1215-1

Suggestions:

check out website: <https://uscms1-se.fltech-grid3.fit.edu:8443/gums> and/or <https://uscms1-se.fltech-grid3.fit.edu:8443>  
neither of which will load,probably because gums isn't workign? he suggested restarting vdt, whih i did yesterday but i guess i should probably do it on the SE

doug did df -h on the se and foudn that sandbox was 100% full. but it doesnt show that for me? odd.  
even when im logged in as him.

okay so just restarted vdt on the SE, going to check results and then try it with the -force flag.

2/11: 2-4

again working with the results from email to Tier-3  
Also attended tier3-osg meeting evo

reading through the gums logs again, there were a lto of issues back in September which is when all the certificates expired. then its clear until now. however the errors are different, it can't read it but why wouldnt it be able to read it now? its been doing it for ages. they havent been movied, permissions ahvent changed?

anyways made it readable by u and g and o. which is probably a security error but we'll see if that changes the error?

and was called out by name and asked if was attending the boston meeting.

2/14: 10:30-12, 1-330

cluster is down :(

went in to see what was going on, Bryant was already there. The nodes were half hibernating, (not off but certainly not running) could hook up the monitor and keyboard to the CE

and work on it like that. could not ssh to nas, nodes, or se.

Turned everything off and back on manually other than the CE which we did using the terminal.

Could ssh to the nodes and the Nas but still not the SE.

Trying to restart just the SE one more time. Hit the reset button to see if this changes anything. In the emergency documentation you have to do the CE and SE at the same time so might

bring everything down and back up one more time.

thoughts

-test ethernet port: does register that there's a network but doesn't give my computer internet.

-traceroute to nas its just one hop, to SE it only reports itself and doesn't get any further. to olin.fitedu can't resolve hostname

to my computer's ip address. gives six hops in asterix's and then its own address. from olin.fit.edu to the cluster gives 3 legitimate steps and then 27 rows of asterix's and then times out.

-emailed curtis and John

2/15: 12-1

yay cluster is running! and SE is too! hurray! all thanks to John Stevens and IT. One of their switches had issues

registering mac address at <https://reg.fit.edu/reg/reg.php>

get mac address with `/sbin/ifconfig | grep -i hwaddr`

went down and grabbed the port info and sent that to dr. H to be taken care of. Now fixing SAM tests...

hopefully.

emailed everyone back and now heading out to mail my application for paris!

2/16: 8-12

SAM tests. Priority uno!

-the errors changed. on the SAM tests...they're still on the same two tests but now are all about connections failing. might be the same problem

that we experienced last semester? replied to bill and Doug

also might be an issue with the firewall? please no.. ughhhh

2/22: 1-2

didn't work monday bc of president's day so today only worked half shift bc of meeting.

Sam tests: are now fixed! huzzah!

so this time...

one of the CRLs on the SE became corrupt. Doug Johnson helped me find this out.

We copied the crl to a new folder with the same name and restarted bestman with



/etc/init.d/bestman stop and then start and now its fine.  
Also there are two bestman installations. one is /sandbox/osg-se and the other is /sandbox/bestman-1.2.13 the second one is the good one. i.e. source this one for the vdt stuff.

Ganglia issue:

Sent an email to the ganglia list and the rocks list which is probably more their forte since its a rocks addon that we're having issues with. Also looked through the code for a bunch of different ganglia functions (/var/www/html/ganglia/addons/rocks) but nothing good yet.

Rocks website: <http://www.rocksclusters.org/rocks-documentation/4.2/x1097.html> describes the command but not any troubleshooting/installation tips... I'm worried it would require an entirely new rocks installation or something?

2/23:2-5:15

Sam tests now broken. BLEHOUGHou

same two tests but completely different errors this time. Also ce-frontier has an issue too.

a new updated Ca certs came out so thats probably the problem. tried the easy route and did

```
/opt/osg-1.2/setup.sh  
vdt-control -on fetch-crl
```

```
/sandbox/bestman-1.2.13  
vdt-control -on fetch-crl
```

we'll see if that helps...

CE-mc: 19:28

CE-frontier: 18:43

CE-lcg-cp: 19:03

okay, did something CRAZY! on SE.

downloaded tarball from OSG, went to /etc/grid-security

made a new folder and untarred the tarball in it, this created a whole new set of revised/updated CRLs

Then went to /etc/grid-security/certificates

copied all of the old CRLs to certsbackup2011 folder

then deleted the originals

copied new ones into main folder

and ran

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

Repeating procedure on the CE.

Also:

CE-mc: 20:25

CE-frontier: 19:49

CE-lcg-cp: 20:03

RSV: 17:15

\*note on the CE maybe they're not supposed to be executable? not sure bc forgot to check before changing permissions.

waiting for SAM tests to rerun. Original errors are attached to this email....  
AHHHHHhhhhhhhhhhhhhhhhHHHHHHHHhhhhhhhhhh

2/24: 12-1 woo for getting out of subatomic early because of a quiz :D  
so sam tests are all green as of last night so it looked liek what Johanna laina  
and I did is working just fine!  
Which is doubly great because we fixed it all on our own!  
We're getting better at this stuff.

Big priority now is the ganglia top issue.  
On the ganglia mailing list the only reply i got was to post on the rocks mailing  
list....  
on the rocks mailing list the only reply i got was a question about whether the  
nodes were rocks nodes...yup  
sent that reply, maybe he'll come back with something?

searching through the addons/rocks documentation. also forgot my fork so can't  
eat my lunch :(  
and my charging cord...  
ugh.

Okay so tried running /etc/init.d/gmetad start on one of the nodes. It didn't change  
anything and after logging out logging in it wasn't running it anymore.

I think this is what I want?  
[http://www.rocksclusters.org/rocks-documentation/4.2/customization-ganglia.ht  
ml](http://www.rocksclusters.org/rocks-documentation/4.2/customization-ganglia.html)

2/25: 8:30-11, 2-4  
Got an email back from ROCKS listserv guy.  
suggested restarting gmond on all the nodes with the command  
rocks run host compute 'service gmond restart'  
so was close yesterday but tested wrong one.  
I'm doing it on just one node first to see what happens.  
nothing.  
trying the full command.  
it did stop and restart all gmond on all the nodes but it didn't affect cluster  
top  
emailed greg back.

In the meantime, updating the "official" Cluster administrator guide with the  
latest bout of solutions. :)

Need to add:  
-the installation of the OSG updated CRLs  
-The fix for the mc/lcg Se thing, and anything else that's happened this semester.

Okay so thats all updated and copying to cluster....

Now summing up that thing about the RSV failures for Dr. H.  
Goal is to get it done before I have to leave to go grocery shopping!AHHH

Back!  
Working on ganglia top problem.  
Email from Greg woot!

"

When you examine the ganglia web page for your frontend, click on one  
of the compute node graphs at the bottom of the page. at the top left

of the page, you'll see 'Time and String Metrics'. You'll see metrics named 'boottime', 'Gmond Started', etc.

Do you see any metrics named 'ps-0', 'ps-1', etc.? "

answer: no.

next reply from greg

"

can you login to compute-0-0 and then try:

```
# cd /opt/ganglia/lib64/ganglia/python_modules
# /opt/rocks/bin/python ps.py
```

then see if any 'ps-0', 'ps-1', etc. metrics are displayed on the host's ganglia web page. "

Problem is we dont' have those directories... I'm doing a search now for python\_modules.  
both on cluster for python\_modules, compute node for ps.py and compute node for python\_modules

on compute node found  
/opt/rocks/lib/python2.4/site-packages/gmon/metrics/ps.py

2/28: 9-12, 1:45-3:30 (GASP! actually got out of comparative early)  
GANGLIA!

Okay so I tried to run the ps.py (python script) with python but got the following  
[xfave@compute-1-0 metrics]\$ /opt/rocks/bin/python ps.py

Traceback (most recent call last):

File "ps.py", line 193, in ?

app=PS()

TypeError: \_\_init\_\_() takes exactly 2 arguments (1 given)

emailed Greg about that.

searching for documentaiton on these processes or other people with issues with the addons

[http://www.rocksclusters.org/roll-documentation/ganglia/5.4/roll-ganglia-user\\_sguide.pdf](http://www.rocksclusters.org/roll-documentation/ganglia/5.4/roll-ganglia-user_sguide.pdf)

the whole installation guide for ganglia...not much on troubleshooting though have i ever checked patricks documentation for this stuff?

The find commands i was running finally completed  
there is officially nothing on our cluster or nodes named python\_modules  
the only file name ps.py is the one named before.

Note on Swinst test: means there's a new cmssw version

Bockjoo needs me to install a package before he can install the new cmssw verison  
gave me the link to the installation instructions

went pretty well but ran itno an error for dependencies. needs /usr/lib/libxml2.so

3/1: 1215-2

Tried a bunch of different things to fix the node.

Had it temporarily working by hitting F2 during boot up and it letting me choose how to load but then i decided to be stupid and try reinstalling it again did /boot/kickstart/kickstart-cluster-pxe and now hitting f2 doesn't do anything. probably rebooted it 10 times. Geez.

During reboots took some time trying to find out what to do with the rpm for bockjoo but its not being agreeable

3/2: 2-5:15

Today I am going to get the node up and the RPM installed.

Node:

okay so from some researching found this link

[https://wiki.rocksclusters.org/wiki/index.php/Compute\\_node\\_install\\_failure](https://wiki.rocksclusters.org/wiki/index.php/Compute_node_install_failure)

under the compute node fails to start anaconda autoamtic install it says to check if the kickstart file is being correctly generated which is what Greg asked me to do so now i understand that mine wasn't

so thats probably the problem

the error came in with the extend-compute.xml

i am doing a search both on the node and the frontednd for that file.

then sending its contents to Greg

doublechecked all the dvds we have in the lab and none of them seem to be ROCKS 5.0 boot

looking at /var/www/messages to see what it has to say about yesterday

```
Mar  1 12:41:32 uscms1 rockscommand[4813]: command "list host sitexml" requires root
```

```
Mar  1 13:01:01 compute-2-6.local kernel: Memory for crash kernel (0x0 to 0x0) notwithin permissible range
```

```
Mar  1 13:01:02 compute-2-6.local sshd[2989]: error: Bind to port 22 on 0.0.0.0 failed: Address already in use.
```

```
Mar  1 13:01:10 compute-2-6.local gdm[3397]: (null): cannot open shared object file: No such file or directory
```

```
Mar  1 13:33:39 compute-2-6.local cluster-kickstart: kickstart node
```

right now i think the kickstart file is the main issue. question is whether i can fix it without being on the node?

so going to search for it and thenr eply to greg

in meantime

RPM for Bockjoo so SAM tests will go green:

So on site he sent me to there was a link for a repository for the dependencies.

So it was so short I wrote the file in vi in the directory suggested, checked the permissions and then

did

```
yum --enable=HEP_Oslibs.repo install HEP_Oslibs_SL5
```

and got the error

```
Error: Unable to satisfy dependencies
```

```
Error: Package HEP_Oslibs_SL5 needs /usr/lib/libxml2.so, this is not available.
```

Hey-ho! Found the .xml files

```
cd /root/NAS_Backup_10_8_09/Scripts/extend-compute.xml
/mnt/nas0/system-config/Config/Rocks/extend-compute.xml
/export/home/install/rocks-dist/lan/x86_64/build/nodes/extend-compute.xml
/export/home/install/site-profiles/5.0/nodes/extend-compute.xml
```

on the node only the /mnt/nas0/ one was found.

I emailed Greg with the contents of

```
/export/home/install/site-profiles/5.0/nodes/extend-compute.xml
since its location was the closest to the one he suggested.
```

He replied back with edits to that file. basically had to change all the -!> to -->

then rebooted the node. and it did the install automatically :)

Came back up all by itself and now is running on CLUSTER

TOP!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!

HUZZZZAH!

We're planning to take the rest down with the same modification when we take the whole cluster down for the NAS which should be the monday after spring break.

3/3

Nothings running on 2-6

fix it.

3/10 (huge break because of spring break) working from home...

so installed the package with the libxml2.so, I ran it in /etc/yum bc you know, why not?

Cluster has been really slow all morning

now going to try to install HEP\_OSlibs.repo again  
didn't work.

3/14: 2-6

So went to shpping and recieving to pick up the NAS but they couldnt find it, waiting for a call back from Steve who istheadirector of that dept and who had spoken directly with Dr. H about it.

Meantime working on node 2-6 problem and SAM test note.

node-2-6:

Nas is not mounted

editing /etc/fstab file to match the one on the other nodes exactly.

then mount -a

```
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
```

so basically had to remount nas, and then restart condor.

instructions were at  
<https://twiki.grid.iu.edu/bin/view/Tier3/CondorSharedInstall>  
which i got from Patrick's documentaiton.  
Looks like its happily chugging along. Has already got jobs running on it. :)  
BUT! fthese jobs aren't the ones that show up in cluster top oddly enough instead  
its all the root processs.

SAM Test swinst note  
aka trying to install this stupid package. Is it working? NO.  
So according to christophe by installing a specific rpm should get the libxml2.so  
but did not. Do have libxml2.so.2, which by the way points to libxml2.so.2.6.26

3/15: 12:15-2

On a hunch i'm going to try getting the libxml2-blahbah.x86\_64 instead of .i386  
and see what happens  
putting the rpm file.  
sudo rpm -ivh libxml2-devel-2.6.26-2.1.2.8.x86\_64  
(i did it in the /etc/yum directory but thats not really official)  
sudo yum install HEP\_OSlibs\_SL5

this time it made it through the whole dependencies check although it did  
get an error during the transaction check. I think it doesn't like having two copies  
of the same file  
bc of the i386 and x86\_64 but am not really sure if it matters  
sent the full output to bockjoo along with my thanks. I'm sure he's ready  
to be done with this headache.

okay so

```
[xfave@uscms1 yum]$ sudo mv /usr/share/man/man5/ldap.conf.5.gz  
/usr/share/man/man5/ldap.conf.5.gz+openldap-2.3.27-8.el5_1.3
```

Password:

```
[xfave@uscms1 yum]$ sudo mv /usr/share/man/man5/ldif.5.gz  
/usr/share/man/man5/ldif.5.gz+openldap-2.3.27-8.el5_1.3
```

at Bockjoo's request but didn't change outcome.

now doing..

```
rpm -e openldap-2.3.27-8.el5_1.3
```

```
rpm -qa | grep openldap-2.3.27-8.el5_1.3
```

get an error on the first one, "multiple packages"

oh of course because we have openldap-2.3.27-8.el5\_1.3.i386 and .x86\_64

so tried it with each but get errors because of all the programs that need it

rpm -e erases the rpm...ahhhh!

3/16: 8-1, 2-5

Today we install the New NAS!

Reading Patrick's documentaiton

Googling for more documentation. there has to be someone who installed two nas's

PLAN:

- 1) check out paperwork from Silicon Mechanics
- 2) Plug in Power cord but don't connect to frontend
- 3) Turn on and press Del to go to setup
- 4) Setup Raid array
- 5) Setup two partitions / and swap with 12 GB
- 4) make it XFS?  
mkfs.xfs /dev/
- 5)run insert-ethers on frontend

- 6) boot node
- 7) manually partition it

-Turned it on by itself, unconnected to anything other than a monitor. it appears to have got raid6 already on it and is partitioned oddly

Plugged an ethernet cable into it and attached to the switch that's attached to the front end.

rebooted

did ifconfig

and it has its inet addr as 127.0.0.1

thats the same as the lo one on the frontend. the other nas doesn't have one.

I think the first thing we need to do is install it as a rocks appliance?

On CE:

sudo insert-ethers

chose NAS Appliance

rebooted New NAS

nothing happened. I think i have to put the rocks kernel dvd in the NAS2

it doesn't have a cd port

need to burn to a usb. reference:

[https://lists.sdsc.edu/pipermail/npaci-rocks-discussion/2007-March/024253.htm](https://lists.sdsc.edu/pipermail/npaci-rocks-discussion/2007-March/024253.html)

l

Trying to boot from network (changed boot priority in bios)

-it tried to do a pxe boot but failed, i think it couldnt' conect

I think we can basically follow this ?

<http://www.rocksclusters.org/rocks-documentation/4.2/install-compute-nodes.html>

with some minor changes so it acts like a nas?

and then follow patrick's documentaiton

<http://uscms1.fltech-grid3.fit.edu/CLUSTERDevelopment-RocksV.pdf>

then maybe cross compare with the SE doc to make sure its not set up as compute and to edit hostname

WOO! finally got a plan of action. just need that usb!

Waiting for JoJo

in the meantime, node 2.6

doesn't have a symbolic link to the nas in home directory, doesn't have a bestman folder,

doesn't recognize me as sudo.

There is a file in the root home directory on the frontend that has the commands for setting up condor on the nodes

sudo su

cd ~

less NewCondorSetupScript

wow.

3/17: 12:15-1, 2-5

put usb into new nas and rebooted.

on the frontend i had typed insert-ethers and chosen nas appliance  
it didn't come up and the new nas got stuck at a blank screen before even logging  
in.  
it's just a blinking dashed line at the top.

trying same process but chose ethernet switch this time on the insert-ethers page.  
no difference

now trying to boot without a usb in just to make sure it still works  
it did.  
also realized that of course it's not getting to the login screen because it's not  
booting up centos5

so we went into the bios and changed the settings so it first tries to boot off  
a removable device  
so now rebooting again with the usb back in place.  
insert-ether in nas appliance mode

gets stuck after seeing drives  
initializing and establishing link  
intel boot agent ge v3.  
exiting intel boot agent  
changed password to match other machines.

Okay

-gave it a new ethernet cord  
-put back in usb mode  
-adding it as a nas  
Result: (is at blinky stage) 2:07, giving it 10-15 minutes as stated by rocks  
it's been 20 and no change.

trying again with it as an ethernet switch  
after this i'm going to try and see what happens on my usb  
began at 2:29 , 3:44 gave up.

New Approach!

on frontend:  
cd /home/install  
rocks-dist dist  
sudo /opt/rocks/bin/rocks add appliance nas-0-1 membership='NAS Appliance'  
short-name='nas1' node='nas-0-1'

then rebooted, but no change.

rebooted my laptop with the usb in it but nothing happened.  
burned usb image onto a dvd and am rebooting my laptop to try that.

3/19/11: 10-12:30

-but the base disk into the New Nas,  
rebooted and guess what! Insert-ethers found it!  
but did not generate a kickstart file  
do we need one?  
if so what do we do  
if not what do we do  
just hangs, we restarted it and now the mac address isn't found

Bockjooo thing:



rpm is at :  
/home/pford/rocks-dist/lan/x86\_64/RedHat/RPMS/openldap-2.3.27-8.el5\_1.3  
and then .i386 and .x86\_64

PUT ALL RPMS AT THIS DIRECTORY !!!!! FOR CLARITY!!!!!!!

3/21/11: 2-5

Monday! This is the week when we start wrapping things up!

Status:

1. New NAS

got email from Malina, maybe the switch has to be off, but if the switch is off how can it communicate to the frontend, should we directly plug it in?

note after what we did on saturday the new nas (nas-0-1) was added to our etc/hosts file

with its ipadress

2. IPTables, ogt an email from people who run SAM tests asking to update iptables followed instructions at

<http://wiki.centos.org/HowTos/Network/IPTables#head-8450ee609cbecd71b6fef3bd3dlac6228991e073>

```
#iptables -A INPUT -p udp -s 128.142.202.212/255.255.255.255 --dport 3401 -j ACCEPT
```

```
#iptables -A INPUT -p udp -s 131.225.240.232/255.255.255.255 --dport 3401 -j ACCEPT
```

then looked at iptables with #iptables -L | less

and saw that it had been added (including the DNS name which i hadn't even specified) Jojo's going to email them and tell them its been taken care of.

3. Bockjoo, wants the kernel upgrade, we're gonna do it on wednesdya when all three of us are here. woop

I think i'm going to have kim run the commands so it becomes unintimidating for her

NAS:

Traceback (most recent call last):

File "/opt/rocks/sbin/insert-ethers", line 1691, in ?

app.run()

File "/opt/rocks/sbin/insert-ethers", line 1647, in run

raise \_("error - lock file %s exists") % self.lockFile

error - lock file /var/lock/insert-ethers exists

when trying insert-ethers --remove="na-0-1"

on restarting of node-2-6, alot of /usr/bin errors (can't find library etc)

on nas we tried (via patrick's suggestion)

```
rm /var/lock/insert-ethers
```

```
insert-ethers --remove="nas-0-1"
```

reboot nas with network as first boot priority.

Next time: come down to high bay and tyr with cd again

3/22: 12:15-1:15

-running a search for squid on our frontend, because updating the iptables didn't fully fix what the issue was.  
/mnt/nas0/home/pford/ClusterSoftware/frontier\_squid-3.0rc5.tar.gz  
/etc/rc.d/rc4.d/s99fronteier-squid.sh  
For tomorrow's kernel update,  
do yum update kernel kernel-headers kernel-devel  
reboot  
should work.

3/23:2-5

John and Curtis have agreed to meet us in the highbay today at 4 to talk about the NAS

In prep, I want to review what is the current situation.

```
rm /var/lock/insert-ethers
insert-ethers --remove="nas-0-1"
to remove it.
Then read by doing
```

on frontend:

```
cd /home/install
rocks-dist dist
? sudo /opt/rocks/bin/rocks add appliance nas-0-1 membership='NAS Appliance'
short-name='nas1' node='nas-0-1'
insert-ethers and then rebooting the New NAS with the cd in it.
Note probably have to change boot order again since we tried through the network
and it didn't work.
```

Also today hopefully we're going to do a reboot of the whole system to update the kernel

```
do
yum update kernel kernel-headers kernel-devel and follow instructions for taking
it down
Kacheck, selected it on reboot by pressing any key when asked so could go to the
menu.
However the yum update HEP_OSlibs_SL5 still didn't work. GARBAGE.
```

Ask jojo if she emailed barry about the squid-KaCheck.

Worked with John and Curtis on installing the NAS.

Made some progress after some trouble getting it to boot (kernel panic!)

3/24: 12:15-1

Working on Bockjoo issue

rereading all the stuff it says about the package  
decided to try installing it via the second method.

i.e.

wget

```
'http://grid-deployment.web.cern.ch/grid-deployment/download/HEP/repo/sl5/x86_64/RPMS/<most recent rpm>'
```

```
sudo yum localinstall <rpm name>
```

got error:

package xxx.rpm is not signed

edited /etc/yum.conf so that gplcheck=1 instead was gplcheck=0

then tried installing again

got the same error as i get with the previous method of using yum:

Transaction Check Error:

file /usr/share/man/man5/ldap.conf.5.gz from install of  
openldap-2.3.43-12.el5\_5.3 conflicts with file from package  
openldap-2.3.27-8.el5\_1.3

file /usr/share/man/man5/ldif.5.gz from install of openldap-2.3.43-12.el5\_5.3  
conflicts with file from package openldap-2.3.27-8.el5\_1.3

previously with Bockjoo we had renamed both of those files to: permissions: rw--r--r  
ldap.conf.5.gz+openldap-2.3.27-8.el5\_1.3  
ldif.5.gz+openldap-2.3.27-8.el5\_1.3

I'm going to rename them again, to ldap\_original and ldif\_original  
didn't help.

Maybe we should uninstall those two packages with yum?  
they are used to identify users...agh sounds important.

also the note on the swinst test changed to a warning because now apparently we're  
missing architectures.  
bleh.

try removing the rpm next time: check it out:  
<http://www.electrictoolbox.com/yum-transaction-check-error/>

3/25:9-1

Meeting with John and Curtis again to try and get the NAS installed.

Okay so basically, we forced it to boot from network (disabled all other devices  
in bios)

the anaconda screen comes up and it wants a cd. we tried the kernel and base one  
for rocks 5

didn't work. some speculation and we've decided to make the jumbo dvd. I'm trying  
to download the file now.

it says it'll take 3 days but i'm optimistic? 2.2 GB. ugh

we tried just booting off the cd (forcing bios to boot from it) and it proved that  
the cds were working but  
didn't make.

Ben would appreciate having opensource program called  
code::blocks on frontend`

3/29: 12:15-2

Some of nodes aren't running condor bc of scratch directories again.  
looking up steps.

basically delete things in /scratch on each node  
then /sbin/service condor start

Doing RPM removal thing. gasp.?

WAIT! First enabling centos repo

cd /etc/yum.repos.d

vi CentOS-Base.repo

changed enabled=0 to enabled=1

yum install HEP\_OSlibs\_SL5

didn't work, looking at other repos

enablign all of them:

made things worse

unenabling.

steps:

```
sudo yum erase openldap-2.3.27-8.el5_1.3
```

```
sudo yum install HEP_OSlibs_SL5
```

or

```
sudo yum localinstall rpm name
```

3/30: 2-5

in the high bay, John may be coming.

can reboot from anaconda screen with ctrl alt del

on frontend did

```
rocks set host pxeboot nas-0-1 action=install
```

insert-ethers

on nas put the cd in

and it does all the same things...anaconda screen

rebooting, this time going to change boot order

it was booting off disk and nothing else, moving it to network.

didn't work, siad improper boot device.

moving it to hard drive

result: centos kernel panic could not find filesystem /dev/root

error mounting /proc etc

4/1: 9:45-1,

Feeling rejuvenated and refreshed on this problem. Determined to make some type of progress today!

First checking all basics of cluster:everything looks good.

Approved Himali's DBS request.

For SAM tests did:

```
sudo yum erase openldap-2.3.27-8.el5_1.3
```

obtaining a few warnings : warning: /etc/openldap/ldap.conf saved as

/etc/openldap/ldap.conf.rpmsave

/var/tmp/rpm-tmp.71152 line 2: gconftool-2 : command not found

it completed.

```
yum install HEP_OSlibs_SL5
```

completed!

but no longer have a sudo filee?

less yum.log shows everything that was deleted and added

it deleted sudo.

did

```
cp /etc/sudoers /etc/sudoersbackup
```

```
yum install sudo
```

```
mv /etc/sudoersbackup /etc/sudoers
```

to get sudo privs back and with the right config.

This also got rid of our clusters website and ganglia ness. it still has ganglia!

but the webpage doe

doesnt show up.

so should i readd each thing...or try to reinstall the original rpm that when erased got rid of all these?

trying to reinstall the original does nothing.

i did yum install php , passwd

nas isn't being mounted for most users with exception of me, zuo, and burcu all people

who were logged in at time of edit.  
for will i created a link in /home that pointed to /nas0/home/wbittner so now his  
is fine.  
but thats not a permanent solution.

so after asking patrick i'm trying to do each individually  
first one that wouldnt tis gnome-volume-manager, firstboot-tui,  
gnome-utils, gnome-vfs2, evince, authconfig,  
autofs, Network manager, gedit,

Got to gnome-utils.as in thats the first one to tyr.

4/4: 9:00-12, 2:30-5.

Worst case scenario. Cluster is down even though its up.  
no wait its a little bit better. you can log in but its very slow and doesnt' do  
what it should

going to try a reboot. note condor is not running.

Kind of weird, could ssh to the se and it had my home directory

but frontend didnt. also se was fast while front end was not.

watching nas as it boots up. ntpd: synchronizeing with time server  
failed

Now watching frontend:

Mounting other file systems. mounting nfs filesystems mount wrong

gfs type bad option bad superblock on nas missing codepage or

other error in some cases useful info is in sys log. pubshing

login files via 411,

starting condor failed, condormaster not executable.

its faster now logging into it, and actually gives the error about  
not being able to reach home directory.

note now have no issue sshing to nas

steps:

check list of files that were deleted and add them.

check if can restart condor manually. nas is not mounted

starting with package list

\*\*\*\*\*

for package firstboot: /etc/sysconfig/system-config-securitylevel was saved as

/etc/sysconfig/system-config-securitylevel.rpmorig

/etc/pki/tls/certs/ca-bundle.crt.rpmnew,

/etc/pki/tls/openssl.cnf.rpmnew

no package tomcat-connectors exist?

tried the few packages that failed on firday and they all worked today!

----back from class

want to get fsck on cluster and run it on the nas partition.

we have it.

emailed john, curtis, and Tier3

John came and helped!

out of all the packages in the list I either missed or it had to be reinstalled

nfs-utils. So did that one and then he mounted the NAS manually with

```
mount -v nas-0-0.local:/nas0 /mnt/nas0
and it worked.
Restarted condor manually with /etc/init.d/condor start
```

NOTE!!!!!!!!!!

SE does not have /scratch directory mounted  
special device LABEL=/scratch does not exist? but it son the nodes?ss  
maybe im confused and its not supposed to exist on the SE?

4/5: 12:15-2,

So maybe internet isn't working because of tomcat-connectors?  
searching for the original rpm in patricks folder.

```
sudo rpm -ivh
/home/pford/rocks-dist/lan/x86_64/RedHat/RPMS/tomcat-connectors-1.2.20-0.x86.
x_64.rpm
```

didn't work because of conflicts between files.  
could try it with the --force flag.

4/6: 11-1

Okay. I have 40 minutes to find a tomcat-connectors rpm of a newer version  
if it doesnt work I try the older version with the --force flag,  
possibly consider email John first.

```
RPMS stored at /export/home/install/rocks-dist/lan/x86_64/RedHat/RPMS
[root@uscms1 RPMS]# rpm -ivh tomcat-connectors-1.2.20-0.x86_64.rpm
Preparing...      ##### [100%]
    file /etc/httpd/modules from install of
tomcat-connectors-1.2.20-0 conflicts with file from package
httpd-2.2.3-43.el5.centos.3
```

from yum info httpd we get:

```
Name       : httpd
Arch       : x86_64
Version    : 2.2.3
Release    : 43.el5.centos.3
Size       : 3.3 M
Repo       : installed
Summary    : Apache HTTP Server
URL        : http://httpd.apache.org/
License    : Apache Software License
Description: The Apache HTTP Server is a powerful, efficient, and extensible
           : web server.
```

Tried two approaches

first downloaded the newest tomcat.tar.gz file and installed from source  
wget

```
http://mirror.cc.columbia.edu/pub/software/apache//tomcat/tomcat-connectors/j
k/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
```

then did rpm with the -force flag. neither worked.  
<http://hep-t3.physics.umd.edu/HowToForAdmins.html>  
rpm -ivh  
/home/install/rocks-dist/lan/x86\_64/RedHat/RPMS/tomcat-connectors-1.2.20-0.x86\_64.rpm --force

tried with  
/sbin/service tomcat start  
/sbin/service httpd start

AND NOW IT WORKS!!!!!!!  
BOOOO-YAH!

with the small exception of the main frontpage which just looks like an index

4/11: 2-5

Sam tests

-restart vdt

cd /sandbox/bestman-1.2.13 or /opt/osg-1.2

source setup.sh

vdt-control --off

vdt-control --on didn't work

or vdt-control --on fetch-crl just tried it with this

sft-job:7:36 4/9

swinst:14:00 4/1

prod:1:22 4/11

lcg-cp:3:33 4/11 17:33 still failing.

also /sbin/service nfs start

now the SE test is passing!

-remove all jobs in queue that aren't local

condor\_rm user uscms01

condor\_rm user osg

condor\_rm user glow

-look for logs locally

-email tier 3

web front page

-try /sbin/service apache start

nas

-emailed john.

-email/call siliconmechanics

-retry completely maybe somethings changed?

4/12: 12-2

reading Malina's email

reading switch docs

reading the switch docs that Malina pointed us too

<http://hep-t3.physics.umd.edu/HowToForAdmins/General.html#RocksSwitch>

Under kickstart nodes, might help fix 2-6.

<https://mail.google.com/mail/?shva=1#label/Rocks/12f30be744ceb61a>  
try this for the NAS.

SOOO much documentation! and none of it with the same issue we have or even a similar setup.

4/13:3-6

Email from Tier3, everhone needs to update glite-UI  
followed instructions from [http://wiki.pscigrid.gov.ph/index.php/GLite\\_UI](http://wiki.pscigrid.gov.ph/index.php/GLite_UI)

4/15: 2:45-5

Nas on monday with john and curtis :)

SAM tests today.

So started by looking at job logs: /home/osg/ and then read one of the random files that end in .log

it pointed to a directory that had the stdout (standard output) so i went there (/home/osg/.globus/job/uscms1.fltech-grid3.fit.edu/<jobnumber>/

went there and the stdout file was empty so looked in stderr, also empty.

then looked in scheduler\_condor\_submit\_stderr and it had the errors:

```
shell-init:error retrieving current directory: getcwd: cannot access parent directories: no such file or directory
```

googled this and upon someones recommendation with the same error but a different problem

we did:

```
# /etc/init.d/named restart
```

```
# /etc/init.d/named reload
```

current test run dates at...

prod: 4/14 17:19

sft-job: 4/13 16:35

4/18: 3-4

John and Curtis have been working assiduously on our new NAS all day.

They have determined that the problem is the network card on the new NAS.

The drivers for it to be able to talk to our system don't exist. They tried to compile it from source but there are errors with it compiling because it conflicts with

the libraries on the CE.

Worst case scenario we have to upgrade to rocks 5.3

right now they're trying to attach it to a copied version of our cluster (the testbed that patrick designed for them) and compile the driver that way.

We emailed Art to see if he had any input and he said more instructions would be forthcoming.

Also need to install an rpm for bockjoo.

emailed tier3 group about SAM tests.

4/20:2-5

Sam tests:

are now running! but failing. but they are running and its no longer proxy expires issue.

released all of the held jobs so theyve reentered the queue.

emailing Doug with the new full errors for MC: Frontier: Squid:



Bestman. was sandbox always 100% full?

```
on Frontend did
chkconfig --list to see what other things are also off. two that caught my eye
were
rocks-grub and httpd. didn't want to touch rocks-grub but restarted httpd with
sudo service httpd start
```

then rechecking chkconfig --list shows that it didn't restart.

Got a response from Barry about SAM tests, he said to check if squid was running.

```
#squid -k check
bash: squid: command not found
#/usr/local/squid/sbin/squid
doesn't exist past local.
is our squid install gone?
#locate squid
turns back a lot of stuff but nothing that looks like what we want.
```

Then went to check the SE

```
found location with locate command but could also have seen it in
#ps -ef | grep squid
regardless,
]# /sandbox/squid/frontier-cache/squid/sbin/squid
2011/04/20 15:49:11| Squid is already running! Process ID 24006
```

so squid is running but doesnt appear in their snmp server?

```
um sam tests just turned green! so did the rsv probes. Possible explanations
on frontend: sudo service httpd start
on SE: /sandbox/squid/frontier-cache/squid/sbin/squid
```

So want the swinst test green and want the NAS somehow connected.

emailed bockjoo telling him i think im done fixing the mess and asking what the situation is with his side of the cmssw/architecture stuff thats missing on our node.

also jobs are running and finishing since this morning. we were at 400 jobs in the queue with something like 118 in hold. I put all the hold jobs back in queue

```
#condor_release osg
#condor_release glow
#condor_release uscms01
```

and now two hours later we have 155 running and only 178 total in the list. so everything seems to be running fine there.

4/21: 1:45-3:15 (1.5 hours)

So bockjoo said he got the thing installed so that he can start updating our cmssw versions.

he's working on that so if I don't see them all green tomorrow I will email him about the full list because he might not realize how many we're missing.

He also pointed out that the phedex transfer aren't moving to our site. This is because I

suspended them bc they're too big files, we don't have the space for them.

Emailing him about that. BLEH.

Then emailing John/curtis to see where they are on things.

Then updating documentation

Then update certificates so that they[re fixed before they expire (May 27) <---  
might do this with jojo or kim since its basic but kind of important

4/22: 10-11,2-4

-useradd script didn't work for Kim. I think its path errors  
fixing it now so she can work on it when she comes in  
Add /usr/sbin and /sbin to root's path with,  
cd ~  
vi .bashrc  
PATH=\$PATH:/usr/sbin:/sbin  
export PATH

not sent to home directory when logged in.

```
File "opt/rocks/sbin/insert-ethers",line in run raise_ ("error - lock file %'s  
exists"> %self.lockFile  
error - lock file /var/lock/insert-ethers exist
```

-then email john and curtis back  
CHECK.

-fix node 2-6

-then email bockjoo about sam tests that still have a problem.  
-work with kim this afternoon on her crontab issue with the webpage  
also see if she knows anything about the cluster frontpage.

4/25: 2-5

All sam tests are green !  
Worked on fixing ssadighi's messages upon login. everythign is fine  
now except she's not automatically sent to her home directory  
fixed it by adding a line to her .bashrc  
now working on documentaiton again.

I was silly and stopped httpd. doing service httpd start didn't get it back up  
we had the error  
starting httpd: (98) address already in use: make\_sock: could not bind to address  
0.0.0.0:8080 no listening sockets  
available, shutting down.  
had to do  
fuser -k -n tcp 80  
/etc/init.d/httpd -k start

4/26: 12:15-3:15

Original plan for today was to fix node 2-6 but now condor is running happily on  
it  
i restarted it yesterday and checked immediately and it wasn't there but  
it must have come up a few minutes later  
crazy.

Updating documentation now...  
and all done. Uploading to wiki.

-Writing script to check all processes taht if they're down or expired could  
cause problems we see. basically a diagnostic tool

writing in emacs, ctrl-x-c to exit  
one for CE and one for SE. woot. basically done. SE needs to be run as root.  
both have the certificates in them. SE one checks squid, CE one checks nfs, condor,  
httpd.  
saved each in /usr/local/bin on the CE and SE respectively so they can be run from  
anywhere by just typing their  
name. However the SE one has to be run as sudo.  
~~~~~Documentation updated to here!  
-Ben pointed out that submitting to condor is slow even though it shows up inqueue
almost immediately.
emailed the condor list and they suggested removing the "log=" line from the
submission file. which worked.
not sure if its a permanent solution.

4/27: 11-1, 2-3

this morning both jojo and kim should be here so i want each one of them to update
the certificates on the
frontend and one to do it for the storage element. Over the summer I'll have to
renew my certificate and then
update the one on phedex with that one.
Jojo is also installing the new version of crab.
:)
So they still haven't gotten emailed and jojo and I are working tommorrow so we
can finish the install then.
but i need kim to email me hers if they ever email it to her.

4/28:10-12:30

Got the certs from DOE so replace the hostcert and container cert. waiting on the
rsv and http certs.
Great practice for JoJo! :)
Uscms1: host, container, http, rsv
uscms1-se: host, container,rsv,http
all done!

samtests
mc, squid, frontier
15:08, 12:43, 12:43

5/6:
End of the year wrap-up

Met with John and he got the NAS fully installed!
Then he took everything apart and I got to redo the steps so I could learn.

Before today he had installed the new Centos version 5.6 which allowed insert ethers
to generate a kickstart file on Nas1 (new nas).

Now logged on to NAS1 first repartitioned the space.
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
```

Then 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!

have to put phedex on there.