

1/8

Dug up the OSG config.ini file in /opt/osg/osg/etc/config.ini, unsure if Johanna-Laina ran the config script after running the VDT install, so I am modifying the script with my new contact information, and will re-run it.

Scratch that, and error arises;

Required modules missing, did you source setup.sh?

Yes, yes I did source setup.sh, not sure why this is popping up. Been working with Bockjoo a lot though, should probably ask him about that, or googlemachine.

1/15

Have worked a bit with Bockjoo, but he has no idea what's wrong. All certificates check out ok, all paths are correct, it's baffling. Requesting permission to email a special OSG group to see what's up. Once I'm authorized, I'll ask for advice.

Meanwhile, going for a condor update. Documentation from Patrick is almost non-existent, so I will write my own condor update instructions. Will follow instructions on HTCondor website, then write down a quicker version + any speed bumps that hit along the way. Install info + downloads are here: <http://research.cs.wisc.edu/htcondor/downloads/>

I then attempted to ssh in with the username "mysql" and the password in the config file, however, I received a "permission denied" statement.

Side note, we are currently using Centos v5.8.

#Currently running pacman version 3.29, latest version.# ← We no longer use pacman, we have switched to OSG 3.x
Setup files should go to /mnt/nas0/condor, need to shift link to latest directory.

1/17

Commencing condor upgrade.

1/18/13

Working with Bockjoo on cluster issues. Set environment variable OSG_WS_GRAM to =N from =Y, Bockjoo says Colorado runs with N just fine. Re-ran configure-osg -c (See docu on OSG 1.2, Patrick), got an error that our sponsor was invalid. Changed sponsor from =cms to =uscms, re-ran configure-osg -c. An error is popping up at configure_vdt_cleanup, but no error detail was given.

Then moved /edg/grid-security/prima-authz.conf and /edg/grid-security/gsi-authz.conf to *.conf.backup, and enabled edg-mkgridmap via vdt-control -enable edg-mkgridmap such that Bockjoo could access the cluster again.

1/19/13

Working with Bockjoo again. Things we have done so far:

- disabled wsggram in the condor section of /opt/osg/osg/etc/config.ini
- added a missing section in /opt/osg/globus/lib/perl/Globus/GRAM/JobDescription.pm, when compared to <http://oo.ihepa.ufl.edu:8080/t2/operations/JobDescription.pm.heposg01.colorado.edu>, pieces were sub trim and sub environment near the end of the file
- tomcat was still misbehaving, so I removed /opt/osg/tomcat/v55/tomcat.pid, and then restarted tomcat using vdt-control -on tomcat-55

OSG can now see us, although there are still errors galore. We probably need to get that configure-osg script to run properly first.

Also, although the GUMS web interface is live, it has complete configuration issues, and is unable to read anything thus far (mappings, etc.)

1/21/13

Using Bockjoo's gums.config did not work, however, he directed me to the discovery that the port numbers did not match up between our gums.config and what port mysql was running. Be sure these are the same! To check, run a ps -ef | grep mysql and see what port is running (should be 49152), and the verify that is the port number mysql is instructed to use in

gums.config. A backup of our current gums.config has been made in the same directory (/opt/osg/vdt-app-data/gums/config) in case the current one gets corrupted.

After a restart of the cluster, went into gums and re-added the rsv user info, which yet again mysteriously disappeared (see OSG 1.2 and GUMS installation on the wiki, page 5)

1/23/13

advice from <https://www.opensciencegrid.org/bin/view/ReleaseDocumentation/ConfigurationFileHelp>
/opt/osg/osg/lib/python/configure_osg/configure_modules/misc.py
dug up an old misc.py at http://vdt.cs.wisc.edu/upstream/osg-configure/0.5.3/osg-configure-0.5.3/configure_osg/configure_modules/misc.py
copied over and modified conflicting python modules etc. configure-osg -c worked.

We are now receiving rsvuser and mis jobs, but jobmanager is still having problems. Will look into glidein stuff tomorrow.

1/24/13

Note: SRMping and SRMread/write is still failing, last time they mysteriously went green after a yum update using the GUI on the CE and SE. Will start updating packages individually (or in small batches) and see what happens. Also look into bestman update/reinstall, as that seems to sometimes be necessary (host name is hardcoded everywhere, a reinstall would be the only guaranteed all hostnames are correct)

Also need to check docu on how yum updates should really go, rather than how I do them...
Yum guides for most commands/setups are here: <http://yum.baseurl.org/wiki/Guides>

1/25/13

There was no pool range set on uscmsPool under Manage Pool Accounts, so nobody within the user groups mapped to uscmsPool could be properly assigned. This was seen in the logfile /opt/osg/tomcat/v55/logs/gums-service-admin.log as

```
"25 Jan 2013 11:14:38,926 [WARN ]: Could not assign user '/DC=org/DC=doeagrids/OU=People/CN=Jimin George 307790' to account within pool account mapper 'uscmsPool'."
```

I set the range as uscms0001-3000, as per the advice of <http://hep-t3.physics.umd.edu/HowToForAdmins/osg.html>, who say cms currently requires a range of about 3000. When a grid-mapfile is generated, we can now see Bockjoo and co. appearing under the cmsuser-null VO

1/26/13

Gums seems to be running ok, but cmssoft did not exist in /etc/passwd, which seemed to be the cause of his authorization issue according to globus-gatekeeper.log. I created a cmssoft account using "useradd cmssoft", but no home directory was created, so I had to create one by hand in /export/home, then mount it to /home.

A similar situation was happening with robot Keith Chadwick, so I also created the fnalgrid account for that. Hopefully this resolves some issues.

1/27/13

Although cmssoft and fnalgrid worked, nothing else in /export/home was able to mount. After digging through, it turns out we use autofs to mount user accounts from /export/home to /home. The config file /etc/auto.home contained paths as such:

```
condor      uscms1.local:/export/home/condor
```

but autofs couldn't recognize the uscms1.local path, as it is no longer in rocks. As such, I moved all uscms1.local → frontend-0-0.local, so they now look like:

```
condor      frontend-0-0.local:/export/home/condor
```

Not sure why uscms1 disappeared, need to look into that. For now though, everything mounts properly, so we should get much greener on OSG-RSV soon.

Bockjoo started testing, found that compat-libstdc++-33-3.2.3-61 was missing from the cluster entirely. Searched online for the CentOS 5.9 rpms, downloaded and scp'ed to root's home directory. There was a conflict with libstdc++5-3.3.6-3mdk, but the contents of this and the new packages was identical, so I had to force-remove the latter using:

```
rpm -nodep -e libstdc++5-3.3.6-3mdk
```

and then to install the two new .rpms I used

```
rpm -Uvh compat-libstdc++-33-3.2.3-61.*.rpm
```

where -U upgrades, -v is verbose (I believe), and -h prints hashes for output. Install went successfully, waiting for Bockjoo to get back to me and see what he can do.

1/29/13

Attempting 2 separate fixes. Cmssoft can authorize, but cmsuser cannot. First attempt: in user groups, cmssoft has "accept non-VO certs" to be false, and "match VOMS cert's FQAN as" to be exact. Changing cmsuser from true to false, and vo to exact.

Apparently the pooling issue is CMS-driven, they have particular setups they require so certain accounts may/may not share whatever may be pooled. The info should be posted somewhere on the OSG GUMS install wiki. I will copy Bockjoo's pooling almost exactly when I get a chance, and tailor it slightly to our needs. Yu Fu also said that each pool account *must* be included in the /etc/passwd (naturally), there's no easy shortcut, so they used a script to pump those out (currently, we need ~2400 cms pool accounts).

Made new accounts for Titas Roy, Hector Hernandez, and Virginia Zhang, as per the request of Dr. Yumiceva (they are his research students). He also requested we update kerberos, as we were on v2.11, and the latest release is v4.6. Not sure why he's kinit'ing from the cluster, but the upgrade was simple. The latest krb5.conf file was found at <http://security.fnal.gov/krb5.conf>, so a simple wget of that address pulled the contents. Our krb5.conf file is in /etc, so I backed up our v2.11 to /etc/krb5.conf.backup.1-29-13, and replaced it with the new krb5.conf. "Works like a charm" now according to Dr. Yumiceva.

1/31/13

Running spreadsheet for diffs between current gums.config, pre-vdt gums.config, and bockjoo's. Includes notes. Spreadsheet complete. Many, many edits. Need to edit gums.config, create appropriate pools, and edit /etc/passwd to add thousands of pool users. It is late. I will do this tomorrow....well I'll finish it anyway. I'll start a few edits right now.

2/1/13

Mostly done with the gums reconfig, only a few points to double-check left. Writing script to generate entries for /etc/passwd, for grid pool and cms grid pool accounts, now. Turns out cms grid pool uses ~2400 and grid pool uses ~4000, so I assigned them 2500 and 5000 accounts each, respectively (some leeway for vo updates down the road). I will include this many in passwd as well, so any vo updates won't immediately trip the non-authorization point.

Also ran a voms members update, turns out the following servers were null and therefore removed (along with groups, accounts, mappers, etc.): gcedu (gc-voms.javeriana.edu.co:8443/voms/gcedu/services/VOMSAdmin), gcvo (gc-voms.javeriana.edu.co:8443/voms/gcvo/services/VOMSAdmin), and nysgrid (dylan.ccr.buffalo.edu:8443/voms/NYSGRID/services/VOMSAdmin)

Some vo's don't update properly, random java errors and such, but they all seemed to be secondary, if they haven't expired at this point.

Nobody in 1000-8000 range with exception of a few 5000 group users. Putting grid pool accounts in the 1001 group, and cms pool accounts in the 6001 group. Beginning grid map users at 1001, beginning cms users at 6001. Editing /etc/group appropriately for groups gridpool and cmspool. Had to add working home directories later, so be sure to set paths in passwd entry the first time around! Unlike me...

Note: also had to add working home directories per account. This means running mkdir generation scripts in /export/home/cmspool and /export/grid/gridpool, running a chmod script to set appropriate permissions per directory. Next, had to write a generation script for entries into /etc/auto.home. NOTE: stop autofs (\$ /etc/init.d/autofs stop) before adding to auto.home, then start autofs back up once you're done.

Alrighty, script written, accounts added to passwd, now we wait...

Nothing bad yet, everything is clearing. SRMPing didn't clear, but that's a different error anyway, that was wishful thinking.

Everyone is clearing according to globus-gatekeeper.log now, but still getting an error code 73 from the job manager. Just found some new documentation on the stuff though, I may have a lead. The file concerned is at /opt/osg/osg-rsv/bin/metrics/org.osg.batch.jobmanager-default-status.

vdt-local-setup at /opt/osg/vdt/etc/vdt-local-setup.(c)sh

GLOBUS_TCP_PORT_RANGE in /opt/osg/vdt/etc/vdt-local-setup.sh

2/2/13

In /etc/sysconfig/iptables.save, there were 2 entries that contained "eth1", which (if this file is actually used) should be moved to eth2. Performed change.

2/4/13

Just realized that, during startup, when you ssh into the SE and su phedex, the command in /sandbox/phedex would fail as the 3 pid files needed were not writeable. Been seeing this for a while, always been suspicious, but never smart enough to go in a look at what's going on. This may have been the error all along, as the pid files were owned/grouped as root:root, aka phedex wasn't allowed in. Ran a chown and put root in phedex's group to be nice, but this should help. Re-ran the command, and it wrote fine, started up. We'll see how it goes after RSV runs probes on it.

2/5/13

So permissions vary throughout the system between rsv and rsvuser. Attempted to switch everything (passwd, group, /export/home, auto.home) all to rsv, but then the osg-rsv cron couldn't start. Error popped up in log at /opt/osg/logs/vdt-control.log:

```
### 2013-02-05 14:46:50 vdt-control(system) /etc/rc.d/init.d/osg-rsv start
```

```
Starting OSG-RSV: ERROR: Problem submitting job to condor-cron. Command output:
```

```
su: user rsvuser does not exist
```

```
ERROR: Problem submitting job to condor-cron. Command output:
```

etc. for many more lines. Looks like something deep is set to rsvuser as opposed to rsv. Look into this before OSG 3.0 install.

2/7/13

Looking into OSG wnclient install, Bestman, Pacman, and Standard Universe.

Also, FIT IT department had an internet issue which reset some service of theirs. This reset showed that we have been running on the wrong gateway/subnet (163.118.42.127) for the past 5 months. Changed our gateway and subnet mask to proper values using:

For CE/SE gateway -

```
rocks set host interface gateway uscms1 eth2 163.118.42.126
```

```
rocks set host interface gateway dev-0-0 eth2 163.118.42.126
```

For subnet mask -

```
rocks set network netmask public 255.255.255.128
```

The latter covers both public subnet masks (CE and SE).

Also, dev-0-0 was acting up, but the name wasn't set properly, nor was frontend-0-0 for uscms1, so I did so with:

```
rocks set host interface name dev-0-0 eth0 dev-0-0
```

```
rocks set host interface name uscms1 eth0 frontend-0-0
```

Also, found java? In /usr/java

2/10/13

OSG update day!

2/12/13

problem with rsv port issue: need to mod /etc/httpd/conf/httpd.conf and /etc/httpd/conf.d/ssl.conf as per the twiki instructions, but shifting the ports takes down our own FLTECH wiki and ganglia, as the usual ports are no longer available. Need to find a way for both to cooperate with httpd...

Slight modification to the instructions in RSV install:

- rather than replace Listen 80 with Listen 8000 in /etc/httpd/conf/httpd.conf, simply add Listen 8000 on the line below
- next, in /etc/httpd/conf.d/ssl.conf, add Listen 8443 below Listen 443
- finally, still in /etc/httpd/conf.d/ssl.conf, add a new VirtualHost for port 8443. Simply copy the format of 443, and just modify the port number (and key and cert files if necessary)
- this all arose due to our Ganglia and Wiki pages being hosted on 80 and/or 443, so we need multiple virtual hosts to handle RSV.

2/14/13

Condor does not seem to be installing properly/fully, found the paths that Should be installed at <http://research.cs.wisc.edu/htcondor/yum/>, missing files/directories are:

- /usr/include/condor

2/16/13

Added backwards compatibility variable into /etc/profile, should cover any jobs that require the setup.sh sourcing. The variable is OSG_GRID=/etc/osg/wn-client/setup.sh, shouldn't have any issues with that aspect now.

2/17/13

Found the error with starting up condor on node 2-9. Turns out for ROCKS wn's that ROCKS provides a copy in a non-system directory, so it breaks GLOBUS when you try to start condor. Must verify that libtool-ltdl is present everywhere for a fix for now.

Removing condor soft link to /mnt/nas0/condor/condor in /opt, going to see what this does.

Didn't do jack, but given the file dates, when I run yum install condor, it pulls from SOMEwhere and keeps the old files. Gotta seal off the node completely somehow, or just get condor 7.8.6 on the CE first and let the nodes pull from there (yes, that would be the better option)

testing node 2-4, since it seems isolated already...

installed perfectly, says the proper version and everything. Now to figure out why the hell everyone else has issues...

ok, modded the date on condor_wait in /usr/bin on the CE, then re-installed condor on 2-9, but still has old date. Turns out condor install does pull from proper resource, so the question is why does it not install correctly? All processes seem to be current...

testing removing sym link within /mnt/nas0/condor instead...no go. Tried cleaning everything out, let's see one more time

didn't work, found a lingering rpm!!! yum remove condor-classads, it got messed up. reinstalled condor, now it works. DO THIS ON THE CE!!!!

also for good measure, do an rpm -qa | grep whatever search for condor, osg, etc. before re-installing condor. ALSO remember to break the condor sym link before installing, just in case.

2/18/13

on compute-2-8, re-installing condor completely (classads included) does not change version. Turns out there was a sourcing entry in /etc/profile, look out for that. Also had to remove globus.

Condor properly installed! Now to get rsv running.

Rsv running properly! Everything is failing, but once gums is fixed some of that should go away. Now to start fixing that one...first a restart.

fetch-crl still having an issue, sendmail failed as usual, and nagios has config issues.

Had to update JAVA_HOME variable in /etc/profile, nothing major
nodes still won't come up...only processing power is currently on the frontend itself, which is not good.

Ok, what needs to happen: nodes need to be brought up, need to figure out how to get gums on 8443, and see what happens from there.

What works: condor (on the CE), rsv, bestman...starts up, but java socket error

What doesn't work: the nodes, gums

old /etc/condor/condor_config in nodes contained info, but used to point to /mnt/nas0/condor/condor-7.4.2/etc/condor_config.local.

2/19/13

Got it! Process:

-ssh into compute node. Install epel repo, install yum-priorities, add line to epel repo: "priority=99", install osg-repo
-remove tomcat-connectors, run yum update, scp tomcat-connectors back into node, re-install with "rpm -Uvh tomcat-connectors_rpm_file_string", rpm search for all condor, remove any lingering. Yum clean all, then yum install condor.

-in /etc/profile on node, remove final line (sources non-existent file), replace with:

```
"export CONDOR_CONFIG=/etc/condor/condor_config
```

```
"
```

-in /etc/condor/condor_config, set new release dir location: RELEASE_DIR = /usr

-in /etc/condor/config.d/00personal_condor.config, at top of file add following lines:

```
"## Where is your local config file?
```

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

```
"
```

install osg-wn-client (yum install osg-wn-client), install fetch-crl, set everything to start up initially, run /usr/bin/fetch-crl first time, then run service fetch-crl-boot start, then fetch-crl-cron start, then service condor start

line 1570 and 1576 in compute node's local /etc/condor/condor_config SCHEDD_ADDRESS_FILE and SCHEDD_something else... needs to change from SPOOL to LOG? Is line 1393 on CE, changed that first, since technically file doesn't exist on compute nodes, only on CE.

Line 248 on compute node, allow_write has to be changed to * perhaps?

2/21/13

Compiling a list of all files that differ in condor config setup, to find the real issue:

-in condor_config.cluster, variables differ.

-file /etc/condor/config.d/cluster.conf had UID domain set incorrectly (name typed wrong, oops). Also had SEC's uncommented

-file /etc/condor/config.d/local.conf also did not have network interface set

added line to 00personal setting local config at /mnt/.../condor_config.cluster, restarting

UGH so it turns out the condor package in the OSG repo was very incomplete. Was missing libraries in /usr/lib, and

executables (e.g. condor_compile) in /usr/bin and /usr/sbin. To fix: I went to the HTCondor site, downloaded the rhel5 x86_64 .rpm to my laptop directly, unpacked the usr directory inside, and scp'ed the contents into the /root directory. From there, I copied over all the libraries/executables from /root/usr/bin, sbin, lib, and libexec directories (from within said directories) using: cp /*

The logic works in that the OSG condor package *did* write new libraries/executables, I believe it simply assumed all the old ones were somehow copied over correctly, or something of the like. As such, NOTE: if you use cp /*, respond "n" to all overwrite queries, as whatever is already in /usr/bin, /usr/sbin/ and /usr/lib are probably the latest versions, you're only copying over whatever is missing (which, as such, will not present an overwrite query).

After copying everything over to their appropriate directories, I ran a condor_reconfig. I then verified all paths were correct (at some point, in the CE's /etc/condor/config.d/00personal_condor.config the LOCAL_CONFIG_FILE = /mnt/nas0/condor/condor-etc/condor_config.cluster line had disappeared, had to re-add it), then restarted condor on the CE, and then on the nodes. Su'ed into my own account, ran a condor_compile (which ran perfectly), and then a condor_submit for a job, and it ran! Did a bit more testing (tested each node individually to verify they all worked, rather than one node running everything), and it works beautifully. Now to update the rest of the nodes, and shift GUMS over to the SE. Those two should fix almost everything except bestman, but that won't work properly until GUMS is up and running anyway I don't think...damn it all. Small wins, small wins haha.

Successfully brought up nodes 2-6 and 2-7, will do the rest in the next few days

2/22/13

Working on getting gums on the SE. output of creating a new GUMS database:

will use domain name "fittech-grid3.fit.edu" in hostToGroupMapping

will use the certificate dir: /sandbox/osg-se/globus/TRUSTED_CA

just good to keep around.

Added myself to admin successfully...

Getting fetch-crl error, I suspect it's from the above cert dir that is invalid. Look into certs tomorrow.

2/23/13

GUMS issue may easily be due to fetch-crl stuff, still haven't gotten it working correctly on the CE

If a tomcat5 issue instead:

old web interface data was at /opt/osg-1.2/tomcat/v55/webapps/, now stored in /var/lib/tomcat5/webapps. Correct config, must be something else

Current file /etc/tomcat5 matches /opt/osg-1.3/tomcat/v55/conf, rest of latter folder is moved to /var/lib/tomcat5

current error kicked out by fetch-crl is at /usr/sbin/fetch-crl line 888, need to determine lastUpdate and LastUpdateOfCRL and crlFile variables (check around 870 for everything being set...added a line to /usr/sbin/fetch-crl to print outputDirectory, is apparently set to /etc/grid-security/certificates else

found one fetch-crl issue, 367b75c3.0 (the cert itself) in /etc/grid-security/certificates didn't exist, had to dig up a copy online and make a new one (as root, just copy/paste contents into above filename, does the trick)

holy hell, it appears to be working. Pulled cert copies from <http://www.ngs.ac.uk/use/cacerts>

very odd. It worked when running /usr/sbin/fetch-crl, but then running service fetch-crl-boot start got same "newer version" error. Skipping the /usr/sbin/fetch-crl (since technically the list is already there) and moving straight to the boot command executes as expected, fetch-crls seem to be running. The only problem is that gums still doesn't come up.

2/24/13

It appears that for gums to work, stuff needs to either be in /var/www/html, or linked from somewhere else.

2/28/13

gums "working" on the SE (still getting apache errors, but main interface page up), we shall see how well this goes.

Went into /etc/osg/config.d/30-rsv.ini and edited line 143 to shift the gums host

Also edited final line for condor, had that line uncommented which was odd, commented it out so hopefully default is taken

Also edited /etc/rsv/rsv.conf to not use condor-g for the moment, just for the hell of it. Restarted rsv, ran metrics, checking what happened momentarily. BAD IDEA, rsv totally tanks for CE metrics.

3/9/13

Got sick and tired of this GUMS issue, something java-related is unresolved (as seen in catalina.out);

An error occurred at line: 15 in the jsp file: /error.jsp

gums cannot be resolved

12: <body>

13: <%@include file="topNav.jspf"%>

14: <div id="title">

15: <h1>GUMS <%=gums.getVersion()%></h1>

16: <h3>GRID User Management System</h3>

17: <h2>Error</h2>

18: </div>

Submitted a ticket to OSG, hopefully they will be able to help. Digging up a wiki format now.

Reconfiguring the new wiki now. During setup;

wiki name = FLTECH

contact email = robertsonj2012@my.fit.edu

language = english

copyright/license = GNU Free Documentation License 1.2 (Wikipedia-compatible)

admin username = WikiAdmin

admin password = Aedc87!!

object caching = no caching

email features = enabled

user-to-user email = enabled

email notification about changes = Enabled for changes to user discussion pages, and to pages on watchlists

email address authentication = enabled

database type = mysql

database host = localhost

database name = wikidb

database user = wikiuser

database passwd = wiki:wiki

NOTE: since new mysql has been installed, had to re-add wikiuser w/password for this to work;

-access mysql: \$ mysql -u root -password

-password is same as root for CE

-create wikiuser account: > CREATE USER 'wikiuser'@'localhost' IDENTIFIED BY 'wiki:wiki'

-grant privileges: > grant index, create, select, insert, update, delete, alter, lock tables on wikidb.* to 'wikiuser'@'localhost' identified by 'PASSWO';

-propagate through mysql: > flush privileges;

-quit mysql: > quit;

checked box to NOT include a super user account, root has that power already

database table prefix = (left this blank)

storage engine = InnoDB

database character set = MySQL 4.1/5.0 binary

ALL documentation has been moved to a logical hierarchy: /var/www/html/wiki/cluster_documentation/*

3/19/13

Looks like the issue may be at the tomcat install. Turns out this comes up when performing a fresh tomcat5 install:

Installing : tomcat5-common-lib

5/6

```
/usr/bin/build-jar-repository: error: Could not find jta Java extension for this JVM
/usr/bin/build-jar-repository: error: Some specified jars were not found for this jvm
```

Put jpackage.repo back in /etc/yum.repos.d, updated. The error shifted from jta to ecj:

```
/usr/bin/build-jar-repository: error: Could not find ecj Java extension for this JVM
```

Dug in, found the ecj at /var/lib/tomcat5/common/lib/[ecj].jar Yum was trying to link everything from /usr/share/java, but couldn't pull eclipse-ecj.jar. Removed old dead link, rebuilt by hand:

```
ln -s /usr/share/java/eclipse-ecj.jar [ecj\].jar
```

nope, that didn't do it. Although this is good that it's properly linked, still getting the same error... going back into /usr/share/java, there were a few other dead guys in there.

Errors in catalina.out:

```
INFO: validateJarFile(/usr/share/tomcat5/webapps/gums/WEB-INF/lib/servlet-api-2.3.jar) - jar not loaded. See Servlet Spec 2.3, section 9.7.2. Offending class: javax/servlet/Servlet.class
```

and

```
11279 [http-8443-Processor25] ERROR org.apache.catalina.core.ContainerBase.[Catalina].[localhost].[/gums].
[jsp] - Servlet.service() for servlet jsp threw exception
```

```
java.lang.ClassNotFoundException: javax.transaction.SystemException
```

so back in /usr/share/java/geronimo, most of that stuff doesn't work, as the "newest" geronimo versions either don't exist, or are installed somewhere else...

looking at the direct error, stuff in WEB-INF/lib/ exists, but links to /usr/lib/gums, servlet exists, but is from jan 18th...

making another servlet link to a newer servlet gives the same error, problem isn't that. Gotta be something with the jsp stuff. Keeps saying classnotfoundexception, and jsp doesn't exist in /usr/share/gums.... perhaps move gums directory, wipe and reinstall, see if a new jsp pops up in the directory? Worth a shot...

3/21/13

/usr/share/java had ejb.jar linked to geronimo, moved it to the /etc/alternatives/ejb

I think I just found the perpetrator; tried removing java-1.4.2-gcj-compat, but everything has dependencies on it. Holy hell. So yeah, going to remove that momentarily, re-install jdk 1.6.0 from the osg repo, and then re-install tomcat5, then gums. Wish me luck.

Things that were removed with 1.4.2:

ant	noarch	1.7.1-7.jpp5	installed	1.6 M
axis	noarch	1.4-4.jpp5	installed	10 M
bouncycastle	x86_64	1.45-6.el5	installed	11 M
cog-jglobus-axis	noarch	1.8.0-2.osg.el5	installed	28 k
eclipse-ecj	x86_64	1:3.2.1-19.el5.centos	installed	23 M
emi-trustmanager-axis	noarch	1.0.1-1.1.osg.el5	installed	163 k
emi-trustmanager-tomcat	noarch	3.0.0-3	installed	129 k
gidoc	x86_64	0.7.7-12.el5	installed	2.2 M

jakarta-commons-beanutils	noarch	1.7.0-10.jpp5	installed	386 k
jakarta-commons-collections	x86_64	3.2-2jpp.3	installed	3.6 M
jakarta-commons-collections-tomcat5	x86_64	3.2-2jpp.3	installed	187 k
jakarta-commons-dbcp	noarch	1.2.2-2.jpp5	installed	138 k
jakarta-commons-digester	noarch	1.8-1jpp	installed	179 k
jakarta-commons-modeler	noarch	2.0-5.jpp5	installed	171 k
jglobus	noarch	2.0.4-4.osg.el5	installed	579 k
log4j	noarch	1.2.14-15.jpp5	installed	383 k
mx4j	x86_64	1:3.0.1-6jpp.4	installed	6.4 M
mysql-connector-java	x86_64	1:5.1.12-2.el5	installed	7.0 M
tomcat5-common-lib	noarch	5.5.27-7.jpp5	installed	86 k
tomcat5-server-lib	noarch	5.5.27-7.jpp5	installed	1.8 M
voms-api-java	noarch	2.0.8-1.4.osg.el5	installed	162 k
xerces-j2	noarch	2.7.1-12.jpp5	installed	1.2 M
xml-commons-resolver				

3/23/13

Well that didn't work, so re-wiped everything, now moving old config files into directories labeled "old", reinstalling in a "hidden clean" thing. Directory locations are:

```

/etc/
/usr/
/usr/share/
/usr/lib/
/sandbox/
/sandbox/osg-se/

```

ok, moving all of that stuff didn't work either, gave even worse errors. Moved it all back, but some stuff is better now. Need to remove everything again, then step through each directory and re-install everything, see which one is the trouble-maker.
Lcmaps-plugins-gums-client

3/26/13

Going to uninstall everything, search for config files:

```

-tomcat5-servlet-2.4-api
-java (going to have to reinstall bestman2)
-jdk

```

Searching for config files...

```

"removed" jpackage.repo (moved to /etc/old)
"removed" tomcat5, jvm, jvm-common, and java (shifted to old in /etc)
"removed" java* (shifted to old in /usr/lib)
"removed" java*, jvm* (shifted to old in /usr/share)
"removed" java (shifted to old in /usr)

```

found a weird java config file in /etc/java/java.conf...

Reinstall :

```

-yum install tomcat5 (trying something new here)

```

Need gums for testing, so installing gums:

```

-yum install osg-gums

```

went smoothing, now following OSG setup instructions. After running /var/lib/trustmanager-tomcat/configure.sh, remember this output:

```

Info: using default install root: /
Info: using default configuration file: //var/lib/trustmanager-tomcat/config.properties
Info: using default configuration directory: //var/lib/trustmanager-tomcat

```

Warning: //var/lib/trustmanager-tomcat/server.xml already exists! Saving old one as //var/lib/trustmanager-tomcat/server.xml.old.
File /usr/share/tomcat5/server/lib/log4j.jar exists, skipping the linking. Use --force to FORCE linking
File /usr/share/tomcat5/server/lib/trustmanager-tomcat.jar exists, skipping the linking. Use --force to FORCE linking
Info: you can clean up using the following commands
mv -f /etc/tomcat5/server.xml.old-trustmanager /etc/tomcat5/server.xml
rm -f /usr/share/tomcat5/server/lib/bcprov*.jar
rm -f /usr/share/tomcat5/server/lib/log4j*.jar
rm -f /usr/share/tomcat5/server/lib/trustmanager-*.jar
rm -f /etc/tomcat5/log4j-trustmanager.properties
rm -f //var/lib/trustmanager-tomcat/server.xml

nope, tomcat5 doesn't start up correctly.

Tried running yum remove java, then realized I didn't check java paths. Didn't get link failures from above...

JAVA_HOME="/usr/lib/jvm/java" was commented, but commenting it out didn't change anything.

Removing everything, going to replace directories 1 by 1. First one will be /usr/java

Still got a bunch of errors, although slightly different, but still clearly not the main problem, moving back to /usr/old

Trying with /etc now, tomcat5 directory

Got the same slew of errors as above for /usr/java, probably not this directory either.

That's it, going to put everything back.....Everything runs "fine" again now.

Side note: catalina.out always has this error - /sbin/runuser: warning: cannot change directory to /opt/tomcat: No such file or directory

trying to pull from old opt directory, wth...? /sbin/runuser is a problem child. Also, the CE doesn't have that problem, since its /opt/tomcat directory still exists. Made empty /opt/tomcat, issue resolved

looking at ancient catalina.out files, the servlet-blah-api issue was a "problem" back then too, so not the issue here. Also fixed the log4j issue, new log is now at /var/log/tomcat5/tomcat.log. Following errors appear:

“

WARN main org.apache.catalina.startup.DigesterFactory - Could not get url for /javax/servlet/resources/jsp_2_0.xsd

WARN main org.apache.catalina.startup.DigesterFactory - Could not get url for /javax/servlet/resources/web-jsptaglibrary_2_0.xsd

WARN main org.apache.catalina.startup.DigesterFactory - Could not get url for /javax/servlet/resources/jsp_2_0.xsd

WARN main org.apache.catalina.startup.DigesterFactory - Could not get url for /javax/servlet/resources/web-jsptaglibrary_2_0.xsd

ERROR main org.apache.catalina.util.ExtensionValidator - Failure loading extension /usr/share/tomcat5/common/lib/[jta].jar

ERROR main org.apache.catalina.util.ExtensionValidator - Failure loading extension /usr/share/tomcat5/common/lib/[javamail].jar

“

going to reinstall gums now, since tomcat.log is significantly more detailed. Here we go!

So got a slew of CRL errors (of course), but then the same 500 error too. Linked in a new jta.jar and javamail.jar

Ok, sent OSG the latest error found, it deals with jsp not reading some javax/blah/blah, look it up in tomcat.log

3/27/13

Searching through ganglia today, seeing if I can help Erin with this.

Rrd's are generated and placed in /var/lib/ganglia/rrds/FLTECH

there's stuff in /opt/rocks/lib/python2.4/site-packages/gmon/metrics, but it doesn't contain a gmetrics conf.php.

/var/www/html/ganglia contains a conf.php, but doesn't contain qstat line. Line 128 may need to have jobqueue added completely.

4/4/13

Holy hell, network problems.

So a bad .sh file was put into /etc/profile.d, making it impossible to log into the cluster, as the error would terminate the session immediately (affected both ssh and local login). The fix was to boot via an Ubuntu Live cd, mount the partition, navigate to /etc/profile.d and delete the bad script. However, the Ubuntu Live session somehow messed with our IP config. Eth2 disappeared, eth0 got weird settings, etc.

To fix this, from the local desktop, navigate System → Administration → Network

Here you will find a GUI for the network config:

- Under the Devices tab, verify both eth0 and eth2 are active (everything else should be inactive, with eth1 eth1 checked but still inactive, all else unchecked)
- Under the Hardware tab, verify that eth2 is the LAN Gigabit, and eth0 is the 10-Gigabit. These two are initially grabbed by the computer at startup, and should NOT be changed. Change everything else around these settings if you have to.
- Back under Devices, you can double-click each device to bring up settings for said device (addresses, masks, MAC addresses, etc.)
- from terminal, run ifconfig and verify that eth0 and eth2 are outputting same settings as given in the network GUI
- run a “rocks list host interface” to verify that all settings in rocks meet the same settings as given by the previous

two softwares

4/6/13

Online sources claimed it would be simpler to reinstall rocks on the compute nodes. **DO NOT ATTEMPT!!!** As such, attempting on 2-4 now. To do so, from CE: \$ ssh-agent

```
$SHELL
$ ssh-add
$ shoot-node compute-##
```

where # and # are the two designated node numbers. According to site, takes about 10 minutes for the entire process. Probably longer on these guys, but I didn't start a timer anyway. Experimenting on 2-4 since it will be wiped eventually anyway.

4/9/13

Somehow, the ethernet cards on the CE have gotten flipped, so the LAN is being detected at eth2 and the public is being detected at eth0. The network config GUI doesn't affect this in the slightest either. Looking into CentOS device manager reconfig now.

4/10/13

THE FIX: as the MAC addresses seemed stuck with eth2 as the LAN card and eth0 on the other, a rather simple trick sufficed. Thankfully, under Administrator → Network, the GUI allows for “nicknames” to be placed for each eth connection. As such, nicknames were “switched”, so the LAN card has the nickname eth0, and the public card has the nickname eth2. This suffices for Rocks to read properly, with eth2 being public and eth0 being private, as Rocks requires. 3

NOTES:

- 1) after running this fix, I received “temporary failure to resolve hostname” errors from each node/NAS/the SE. As such, it is necessary to run: rocks sync config, rocks sync dns, and rocks sync users (the last of which will give a lengthy output, that is normal).
- 2) after performing the above, you must then shut down and restart all units WHILE THE CE IS STILL UP AND RUNNING. All other units run off the Rocks distribution that the CE provides, which they pull at startup if the CE is available. This includes the NAS's. Otherwise, there is no updated Rocks available to them, so they will continue to boot from a current config which is incorrect. This does mean the NAS's will not mount to the CE properly, so you will either need to re-mount them, or (recommended) simply restart the entire cluster after running the initial restart for proper Rocks.
- 3) this does mean that ifconfig will probably look weird afterwards. Ignore this, as long as rocks is reading the config fine, the network GUI looks good, and you can ping google/other units, it will run properly.

So, to recap: 1) fix whatever needs fixing in the network GUI, 2) align Rocks to the GUI settings, 3) rocks sync (config, dns, hosts), 4) restart all units except CE, 5) test to see if all work properly (try sshing to nas-0-0), then restart FULL cluster

4/24/13

Our wiki got spammed by a large number of new accounts being created. As such, went into `/var/www/html/wiki/LocalSettings.php`, and edited line 104, changing 'true' to 'false' for 'createaccount'. No more users may be created, but now we have to go through, block all users, and remove their respective pages.

In other news: after the card flip, service DHCPD no longer works. DHCPD allocates temp IP addresses to the nas's, SE, and nodes, so they can talk to the outside world through the CE. There's likely a MAC address hardcoded somewhere, which has since been flipped to a new card. Digging through now.

I'm sick and tired of this...config in `/etc/dhcpd.conf`, executable in `/usr/sbin`, but when forcing a start from `/usr/sbin` it kicks off notices for `eth0` and `virbr0`, that neither has a subnet defined. It recommends changing this in the conf file, if need be. I'll look into it later. Going to rewrite the emergency shutdown documentation now instead.

5/3/13

Progress on 2 fronts.

First, a drive died in NAS1, and the replacement wasn't automatically kicking into a rebuild. Turns out our BBU* is drained to 42%. Apparently they get noisy under 50%, and can obstruct rebuilds and consistency checks that the NAS attempts to run as it normally should (in the MegaRAID software, accessible during startup before CentOS boots). After removing the BBU, we now get a warning that it is missing (at least the warning system works), and the RAID card was freed up to run a consistency check, then a rebuild, so all is well. NOTE: a replacement 2 Tb drive was sent to us, and resides in the cabinet in the lab. To normally replace a drive if it fails, you can just remove the old one and plug a new one in, a rebuild should start automatically, and only take an overnight.

Second, concerning DHCPD, at least it starts up now. A third file beyond the config (`/etc/dhcpd.conf`) and the lease file (`/var/lib/dhcpd/dhcpd.leases`) was stuck on `eth2`, which needed to be shifted to `eth0` (file was `/etc/sysconfig/dhcpd`). The service now starts, but on startup from any non-CE unit, it cannot find `uscms1.local`, so no leases are provided. Going to keep looking through this.

*Backup Battery Unit, an Lithium ion battery pack that sits on the RAID card.

/DC=org/DC=doegrids/OU=People/CN=Jordan Robertson 530934