

8/19

Working on adding RSV java probe. Shut down tomcat5 and mysqld, ran rsv update, turned former two back on, and restarted RSV. Also noticed that \$JAVA_HOME on CE was being exported (via /etc/profile) to a bad path, re-routed to /usr/java/latest. Should work better.

Also realized that \$JAVA_HOME wasn't getting set at all on SE, added appropriate line to /etc/profile on SE just for completeness. Now re-installing bestman2 and will re-run rsv probes to see if bestman comes up. Not seeing java probe come up by default though, probably have to edit RSV config file somewhere (definitely in notes from upgrade to OSG 3.1.x), and then re-run osg-configure. Will dig through that tomorrow, for now, just bestman2.

It seemed to install just fine, download output (for versions) of all dependencies, etc.:

Downloading Packages:

(1/12):	osg-se-bestman-3.0.0-8.osg.el5.x86_64.rpm	2.9 kB	00:00
(2/12):	bestman2-tester-2.3.0-9.osg.el5.noarch.rpm	17 kB	00:00
(3/12):	lcas-lcmapi-gt4-interface-0.2.6-1.1.osg.el5.x86_64.rpm	25 kB	00:00
(4/12):	bestman2-server-2.3.0-9.osg.el5.noarch.rpm	27 kB	00:00
(5/12):	bestman2-client-2.3.0-9.osg.el5.noarch.rpm	31 kB	00:00
(6/12):	jetty-deploy-8.1.4.v20120524-1.osg.el5.noarch.rpm	54 kB	00:00
(7/12):	glite-security-trustmanager-2.5.5-6.1.osg.el5.x86_64.rpm	61 kB	00:00
(8/12):	jetty-webapp-8.1.4.v20120524-1.osg.el5.noarch.rpm	101 kB	00:00
(9/12):	glite-security-util-java-2.8.0-3.1.osg.el5.x86_64.rpm	294 kB	00:00
(10/12):	jakarta-taglibs-standard-1.1.1-7jpp.1.x86_64.rpm	766 kB	00:00
(11/12):	bestman2-common-libs-2.3.0-9.osg.el5.noarch.rpm	1.6 MB	00:00
(12/12):	bestman2-server-dep-libs-2.3.0-9.osg.el5.noarch.rpm	2.2 MB	00:00

Notes: for localPathListAllowed in /etc/bestman2/conf/bestman2.rc, simply set to /tmp, may need different paths. Same for supportedProtocolList, left commented out, not 100% on what to put there yet. Set GUMSServiceURL to <http://uscms1-se.fltech-grid3.fit.edu:8443/gums>

Everything seemed to install and turn on correctly, with the exception of gratia. It was missing three packages (the osg-configure, the metric, and the gram), and as such, starting the gratia-xrootd-storage and gratia-xrootd-transfer services didn't work, as the services were not recognized. Ran the RSV probes anyway though, looks like bestman is properly attempting to ping the GUMS server and getting a null response, logical as GUMS isn't working yet.

8/20

nas1 finished rebuild, re-mounted everything on CE, should be good. Working on NICs now.

In /etc/udev/rules.d, have current 70-persistent-net.rules in place, should be ok.

In /etc/sysconfig/network-scripts, modifying ifcfg-eth0 and ifcfg-eth2 back to original states. Going to try flipping with GUI first, then will do manually if it doesn't hold.

Attempt with GUI to change nicknames gave weird response, nickname shift was pushed to ifcfg-eth0, but ifcfg-eth2 disappeared. Just got it back, you may have to highlight each eth* and save. Continuing switch:

under eth0, Hardware Device tab, each eth is pinned to a specific MAC address. Have to restart to change card names before messing any more with GUI. Changed nicknames back, saved, restarting with proper(?) 70 rules.

Udev rules not seeming to hold, moved file back to /root. Flipped all eth0 to eth2 stuff and vice versa in the GUI, saved, changes got pushed to ifcfg-eth* files. Restarting, will need to do rocks flip next. If rocks can handle eth0 being public card, we should be fine. If not, then argh.

Udev didn't repopulate /etc/udev/rules.d with a 70-persistent-net.rules file, not sure about that. Restart has CE talking fine with outside world though, so will attempt rocks flip now. Note: eth2 still has an "interrupt" on it.

Removed old "unknown" eth1, re-created network connection from that card. Still had interrupt on it, but restarting CE to see if it comes up properly. If so, we can route rocks onto eth1 rather than eth2.

Check ifcfg-eth0 and ifcfg-eth2 for SE as examples for CE?

8/21/13

Moved the new rsv config file in /etc/osg/config.d to actual file (had been saved as a .rpmnew file), so hopefully that runs alright.

Can't run osg-configure to update rsv until GUMS can kick out a user-vo-map, which GUMS can't do as java isn't working properly. Irony.

8/22/13

They used jumbo frames to set up the cluster network. Look into these asap, likely the source of the issue. Also used a DVD for that, look for said DVD.

Side comment: seems that 70-persistent-net.rules is CentOS 6.x, changes in CentOS 5.x should be done in 60-net.rules. This explains why they're not holding. Adding to 60-net.rules now, restarting momentarily.

Udev rules change was a success, the cards now have their proper eth assignments. Mission control, we're back online with eth2. Now to figure out why eth0 is freaking out.

Also switched values for eth0 and eth2 back to originals, in /etc/sysconfig/iptables.

Note: ganglia came back up, although dhcpd service still won't work.

8/25/13

"Fixed" the internet connection for the nodes, pretty sure it's a fluke hack, but I don't mind. Issue was in /etc/sysconfig/iptables, line 35 read:

```
-A FORWARD -i eth2 -o eth2 -m state --state NEW,RELATED,ESTABLISHED -j ACCEPT
```

This makes no sense, as all traffic rerouted to itself. Moved second eth to eth0, so it now reads:

```
-A FORWARD -i eth2 -o eth0 -m state --state NEW,RELATED,ESTABLISHED -j ACCEPT
```

This allows traffic from the nodes to the outside world and vice versa. Running updates now. Starting on compute-2-9 as a test batch, updating everything (including kernel, condor, and installing fetch-crl3). Snazzy; update automatically replaces fetch-crl with fetch-crl3, no need to remove original fetch-crl.

As suspected though, condor config is slightly off, it comes up fine but it's local only. Need to dig through config files and see what's going on before updating other nodes. Kernel update seems fine though, will push on all other machines.

Update on nas-0-0 seems fine, moving to nas-0-1. Updating complete, restarted. Same with SE. Restarting nas-0-0 now. Successful, everything has been updated and restarted with the exception of the CE itself.

Condor is fine, just have to go into /etc/condor/config.d/00personal_condor.config and re-add the local config file line, see condor 7.8 install instructions.

Nodes completed: 2-9, 2-8, 2-7, 2-6, 2-5, 2-3, 2-2, 2-1, 2-0, 1-9, 1-8, 1-7, 1-6, 1-5, 1-4, 1-3, and 1-2.

Also updating compute-1-1 and compute-1-0, may as well. Comically, they require ~600 updates, which took over an hour to download headers, only to have it fail. Will start installing/updating individual packages.

Just needed to update yum...will do first on compute-1-0.

Don't forget to cluster-fork chkconfig fetch-crl3-boot and fetch-crl3-cron after all are installed/updated.

8/26/13

Running install of SLC5 on compute-2-4. Remember ROCKS re-install was pushed first, have to clear that out now.

Just dug through old docu, realized I also changed /etc/sysconfig/dhcpd back during the NIC issue. Changed back to eth0, DHCP service came up perfectly fine. Restarting compute-2-9 now to see how changes affect this. Interesting, same "can't get address for uscms1.local" error when attempting to mount NFS filesystems, although eth0 comes up ok. Attempted changing name in /etc/dhcpd.conf to uscms1.local (from frontend-0-0.local).

Scratch that, uscms1.local error comes up even when DHCPD is turned off. Even more interesting....also, nas0 is mounting just fine...wtf.

9/5/13

Spoke with Curtis and co., turns out the interrupt means very little, both networks are working perfectly fine.

Having trouble pulling down the SLC5 image for jesse, need to try all network options and/or find a storage device big enough to hold the ISO provided (a 4.7 Gb DVD claims there's no room for the 4.4 Gb ISO....)

OSG recommends not running GUMS on either the CE or SE, asking where they would rather we run it. May end up running it on nas1 as it has a separate dedicated IP(?). Need to dig into that more.

Seems that IT is giving Erin trouble with the website, Maki didn't know much and communicated even less, so she has little to go on.

9/7/13

Just checked for something interesting on the SE; turns out the tomcat5 version on the SE is different from the one on the CE. Need to verify which version should be used, pulling current one from CentOS, but if OSG has one to push, we should use that instead.

9/8/13

Turns out the CE had a bad version of tomcat5 on it, damned thing still had a Jpackage version. Removed it, had to reinstall rsv-metrics (had old dependency).

Alright, turns out Java 6 is no longer supported, they're trying to migrate to Java 7. As such, attempting to install tomcat5 and then gums with OpenJDK7 instead. Note: this required the osg-upcoming repo, have to use a tag for every invocation of yum; yum ... --enablerepo=osg-upcoming ...

Nope, still the exact same error. I have now literally tried everything short of reinstalling the OS on the SE. What in the hell.

9/9/13

Erin found the Rocks rolling documentation for ganglia here:

<http://www.rocksclusters.org/roll-documentation/ganglia/5.4/roll-ganglia-usersguide.pdf>

Perfectly explains how to re-roll software, can do so after re-installing nas-0-1. I believe that procedure is as follows:
rocks run host nas-0-1 and_then_whatever_the_reinstall_command_is

Also just got a new grid cert, switched from:

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

To:

/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=jordan/CN=746428/CN=Jordan Alexander Robertson

9/14/13

Tried to put nas-0-1 on rocks, but the DHCP request didn't hit the CE after starting up insert-ethers. Very odd. In other news, wrote SLC6 boot disk, and pulled down images from the CERN site successfully. Implies 1 of 2 issues on our end, if any. Either 1) we need to use http:// in front of our linuxsoft.cern.ch, OR our network config isn't working properly. First thing to do is boot w/SLC5 cd and use http. Second move (if that doesn't work) is to let DHCP resolve eth0 stuff. Finally, if necessary, re-run SLC6 boot/install just to the network point, then see what config it uses (since that

worked).

DHCP attempt worked, no need for HTTP, but boot CD version didn't match version of images being pulled down. Burned a new boot disk, worked beautifully.

9/18/13

Changing compute-1-0 from compute node to management node, to host GUMS. Much information:

1) Need an external IP for this, turns out we have the whole 163.118.42.1-125 block reserved for the high bay, pretty cool. Also turns out the 10 Gb switch address is 163.118.42.125. Just need to include the network info and should be good to go (provided by Daniel Flores):

NETWORK: 163.118.42.0
NETMASK: 255.255.255.128
GATEWAY: 163.118.42.126
BROADCAST: 163.118.42.127
DNS 1: 163.118.22.230
DNS 2: 163.118.131.9

Going to attempt using the graphical network interface first, before manually editing ifcfg-eth1. Also have to add new config to ROCKS, as old compute-1-0 with that mac address is lingering.

In /etc/sysconfig/network-scripts/ifcfg-eth1:
-added line TYPE=Ethernet
-added line USERCTL=no
-added line IPV6INIT=no
-added line PEERDNS=yes
-added line IPADDR=163.118.42.10
-added line NETMASK=255.255.255.128
-added line GATEWAY=163.118.42.126
-changed line ONBOOT=no to ONBOOT=yes

2) There are 2 onboard eth ports per node, the first (eth0) is set to the LAN, the second will be taken by the external IP (eth1). The 10 Gb switch is full however, will condense the external and LAN ports from compute-1-0 to a Netgear switch, then pipe that into the compute-1-0 port on the 10 Gb switch.

3) On the software side of things, need to install mysql-server, then tomcat5, and finally osg-gums (following the OSG instructions for osg-gums). So to start, installing mysql-server was just from a yum package, very simple (note: also installed dependency perl-DBD-MySQL-3.0007-2.el5). On startup of mysql (service mysqld start), the initial output was:

Initializing MySQL database: Installing MySQL system tables...

OK

Filling help tables...

OK

To start mysqld at boot time you have to copy support-files/mysql.server to the right place for your system

PLEASE REMEMBER TO SET A PASSWORD FOR THE MySQL root USER !

To do so, start the server, then issue the following commands:

```
/usr/bin/mysqladmin -u root password 'new-password'
```

```
/usr/bin/mysqladmin -u root -h compute-1-0.local password 'new-password'
```

Alternatively you can run: /usr/bin/mysql_secure_installation

which will also give you the option of removing the test databases and anonymous user created by default. This is strongly recommended for production servers.

See the manual for more instructions.

You can start the MySQL daemon with:

```
cd /usr ; /usr/bin/mysqld_safe &
```

You can test the MySQL daemon with `mysql-test-run.pl cd mysql-test ; perl mysql-test-run.pl`

Please report any problems with the `/usr/bin/mysqlbug` script!

The latest information about MySQL is available on the web at <http://www.mysql.com>

Support MySQL by buying support/licenses at <http://shop.mysql.com>

[OK]

Starting mysqld:

[OK]

I chose to set both passwords to current root passwd.

Now to install tomcat5: tomcat user already exists, just have to run the install now (yum install tomcat5). Install looks like it ran fine.

9/20/13

Moved compute-2-4 back into the rack, also did some cable management so it's all clean now.

Still fussing with the external IP on compute-1-0, not cooperating yet.

Idea for ganglia on nas1: if we install the exact ganglia packages onto nas1 as are on nas0, and configure them appropriately, in theory it should all come up ok. Bad news: the ganglia packages aren't available (probably stored on the rocks roll cd).

Also got a request from Aiken for a local repo, using CVS or SVN. Opting for SVN, already installed but can't find directory. Hunting for that now (need it for hooks), proving elusive. There's even a `/etc/subversion` directory, but it's empty... Seems they never created any. As such, making a new SVN directory at `/var/svn`, which will contain repos. To do so, ran `mkdir svn` in `/var`, and then for proper repo creation from `svn` directory, ran: `svnadmin create repos`. Note: default creates with FSFS filesystem type.

Now making a new group such that Aiken, Brian, and myself can all have access to it (rather than just root). Group is `svnusers`, see `/etc/group` for details. Current members: root, me, Aiken Oliver, and Brian Dorney. Need to add a cluster account for Darren Mareskas-Palcek, then add him to group as well.

9/23/13

Turns out the external IP issue for compute-1-0 was on IT's end, big surprise. Port 1 on the 10 Gb switch was set to the VLAN, Daniel Flores had to switch it over. However, as switch is full, had to condense compute-1-0 and compute-1-1 local ports to a Netgear unmanaged switch, and port 1 on the 10 Gb will now be the designated external slot.

As of right now, mysql and tomcat5 are both installed. According to `/var/log/tomcat5/catalina.out` (on compute-1-0), the following is a small issue:

INFO: The Apache Tomcat Native library which allows optimal performance in production environments was not found on the java.library.path: `/usr/java/packages/lib/amd64:/usr/lib64:/lib64:/lib:/usr/lib`

Researching before installing gums. Also note, jdk 7 was installed as an update for bestman2-client or something along the lines. If initial install of GUMS ends in same error as SE, will remove gums etc., then reinstall tomcat5*.

Note: `httpd` not installed, not sure if required? Will see at gums install.

Note: copying certs (`httpcert.pem`, `httpkey.pem`) from `/etc/grid-security/http` to appropriate new folder on compute-1-0, chown'ing to tomcat.

Note: modifying iptables, adding a line to allow the 8443 port on compute-1-0 (copying format from SE).

Upon yum install `osg-gums`:

Installing:

```

osg-gums          noarch  1.0.0-2.osg.el5      osg  1.7 k
Installing for dependencies:
emi-trustmanager  noarch  3.0.3-5.osg.el5      osg  292 k
emi-trustmanager-axis  noarch  1.0.1-1.2.osg.el5    osg  44 k
emi-trustmanager-tomcat noarch  3.0.0-8.osg.el5      osg  47 k
gums              noarch  1.3.18.009-17.osg.el5 osg  13 M
gums-client       noarch  1.3.18.009-17.osg.el5 osg  17 k
gums-service      noarch  1.3.18.009-17.osg.el5 osg  2.0 M
jakarta-commons-codec x86_64  1.3-7jpp.2           base 109 k
mysql-connector-java x86_64  1:5.1.12-2.el5       epel 2.5 M
osg-gums-config   noarch  47-1.osg.el5         osg  12 k
osg-vo-map        noarch  0.0.1-1.osg          osg  7.3 k
osg-webapp-common noarch  1-1.osg.el5         osg  2.3 k

```

Completed successfully. Upon configuring tomcat to use GSI:

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

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

```

Created a new GUMS database after starting mysqld, went fine, but note domain name (needs to be changed):

Creating GUMS 1.0 database with the following parameters:

GUMS mysql user to be created: "gums"

GUMS server host: "localhost:3306"

Is this correct? (Enter 'yes' to proceed)

yes

Preparing database:

Enter the root mysql password

Enter password:

will use domain name "local" in hostToGroupMapping

will use the certificate dir: /etc/grid-security/certificates

WARNING: gums.config already present. Would you like to overwrite it?

(Enter 'yes' to overwrite)

yes

Backing up your gums.config as gums.config.old

New gums.config has been created successfully

Adding self as an admin; went fine as well:

WARNING: You must have created the database before running this script!

Adding the following DN to the local database:

Certificate DN for administrator: "/DC=ch/DC=cern/OU=Organic

Units/OU=Users/CN=jordan/CN=746428/CN=Jordan Alexander Robertson"

Is this correct? (Enter 'yes' to proceed)

yes

Adding the admin:

Enter the gums mysql password

```
/usr/lib/gums/sql/addAdmin.mysql | sed -e s%@ADMINDN@%/DC=ch/DC=cern/OU=Organic  
Units/OU=Users/CN=jordan/CN=746428/CN=Jordan Alexander Robertson%g | mysql -u gums -p  
Enter password:
```

Template looks solid, but will try booting up with basic gums.config first as a test.

Upon startup of tomcat, no complaints thus far in catalina.out:

Went to page (uscms1-mgmt.fltech-grid3.fit.edu:8443/gums), blank page save a single black diamond containing a white question mark. Shutting down tomcat5, adjusting gums.config for domain, restarting.

Still no luck. Not a certs error either. Contacting chris pipes about this.

9/24/13

Posted to chris, but also; realized /var/log/tomcat5/trustmanager.log had all CRL's failing. Derp, needed to install osg-ca-certs-updater, so pushed that to all nodes.

Also suspect I need a host cert (and thusly can cp for a valid http cert) for compute-1-0 specifically, so requesting that by installing osg-pki-tools if necessary. Asked chris first.

10/06/13

Chris tagged the software team onto our ticket to resolve this quickly. Got confirmation, need to request new host cert for compute-1-0. Interesting results:

1. I never registered my cert to a VO, although not sure if that's required for a host cert request.
2. The new instructions on how to do this are located here:
https://twiki.cern.ch/twiki/bin/view/CMSPublic/EndUsers#OSG_PKI_Transition_Impact_on_End
3. When registering, we have Patrick Gartung as our rep.
4. None of this is documented anywhere, but should be done for all future admins!

Another interesting point of note: although the hostcert.pem and hostkey.pem shouldn't have worked, GUMS does now come up on uscms1-mgmt.fltech-grid3.fit.edu:8443/gums, BUT we get the exact same error as on the SE.

Concerning getting a host cert; you need your own personal cert on the cluster (either your own usercert.pem and userkey.pem, OR your usercred.p12. I have the former, but both formats work apparently). From there, go to the instructions they have for OSG 3.x host cert requests, and scroll to "validate authentication":

<https://twiki.grid.iu.edu/bin/view/Documentation/Release3/GetHostServiceCertificates>

I suspect you need to be registered on the CMS VO. They have two different processes to get a proxy (the first step in "validate authentication"), and only one of them is related to voms, but the other may still require the right privileges. We shall see.

Also concerning getting a proper usercert and userkey onto the cluster; it was very circuitous, but here's the final solution that worked for me:

1. Go to firefox (or wherever you've got your cert stored), find the certificate manager (in preferences for firefox), find your appropriate cert, and back it up as a .p12 file.
2. Open up a terminal, navigate to the directory that contains your .p12 file, and use the following two commands:

```
openssl pkcs12 -nocerts -in mycert.p12 -out userkey.pem
```

```
openssl pkcs12 -clcerts -nokeys -in mycert.p12 -out usercert.pem
```

3. These will create your userkey and usercert on your local laptop. THEN scp them to the cluster, into your .globus directory (subdirectory off your home directory)
4. Change the permissions of said cert and key, chmod 600 is required for them to work with grid-proxy-init or voms-proxy-init. They even recommend one step further at 400 for your key, it's a pretty sensitive thing.

Now that that's all tucked away, as per instructions I created a proxy with grid-proxy-init (again, voms-proxy-init probably won't work for me right now), and am continuing with the host cert request process now. Remember kids, you have to be signed in as YOU. You'll be using your own cert and key, so you can't be logged in as root!

Ssh'ed into compute-1-0 just for safety reasons. Now running osg-cert-request as instructed. It appears that the management node must be added to the system, before I can request a new host cert. However, I don't believe I am registered to being a grid admin for FLTECH whatsoever, discussing that with Christ and Tim on the ticket now.

Also, the voms request totally got cleared by Patrick while I was working on this, I'm now a part of both the uscms and cmsuser voms!

10/16/13

Finally got a response from Chris on adding the management node, had to go through OIM (<https://oim.grid.iu.edu/oim/home>) → Topology → find FLTECH resource, add a resource, making sure the resource service had GUMS on it. Now just waiting for the resource to be reviewed/added, then can give it a host cert and viola.

Meanwhile, Fermilab may be about to close from the government stuff, but since Erin is out sick, I'm taking over cleaning up the nas's to make enough room for stuff from Brian/Aiken/Dr. Yumiceva. Ran the du script that Kim wrote, on nas0 and nas1, and am hunting down large files now.

Zipping /nas0/home/xfave/datafond, /nas0/home/xfave/backup, /nas0/home/sguragai/physics/*, /nas0/home/sguragai/backups, /nas0/home/sguragai/public, /nas0/OSG/*,

(unfinished are /nas0/home/sguragai/physics, /nas0/OSG/APP, and /nas0/OSG/BeStMan)

11/04/13

Got clearance from tony, then alain, to be a grid admin! It cleared through the system, so I am now able to request host and service certificates. I immediately requested a host cert for the CE and MGMT, by going to the OIM site → Certificate → Request New, and then to generate the request (CSR), I navigated to /etc/grid-security, and then ran (for the CE):

```
umask 077; openssl req -new -newkey rsa:2048 -nodes -keyout hostkey.pem -subj "/CN=uscms1.fltech-grid3.fit.edu"
```

and for MGMT:

```
umask 077; openssl req -new -newkey rsa:2048 -nodes -keyout hostkey.pem -subj "/CN=uscms1-mgmt.fltech-grid3.fit.edu"
```

These commands generate hostkey.pem files, so you will need to TEMPORARILY change the permission of hostkey.pem to writable (a `chmod +w hostkey.pem` works fine).

I then copied the output of each command, and submitted these to the new ticket. I could then navigate to "My Requests" on the same page, and approve the cert, and finally I was then able to issue the certificates for each ticket. Ran into a bump, probably because I'm brand-new in the system, the issuing didn't get cleared by DigiCert yet. Hopefully will be resolved asap.

Note: for MGMT (as is set up for the 8443 virtual host on the CE), you must go into /etc/httpd/conf.d/ssl.conf and edit where the cert and key files are. In our instance, on MGMT, you open ssl.conf and navigate to the virtual host, then locate the SSLCert and SSLKey lines, and edit appropriately. This is a one-time deal, but must be done this time as the management node's httpd has never been configured before.

11/08/13

Just got a possible solution from Brian, downgrading antlr, as we are running 3.0.0 but GUMS was built with 2.7.7. Downgrading to next available, which is 2.7.6-4.

HOLY FUCK IT WORKED.

The rsv account disappeared though, going to re-add that in a moment.

Our current rsv account on the cluster is "rsv", so Patrick's documentation is slightly off on the wiki (OSG 1.2 and GUMS install). Under account mapper, the mapper is still named rsv, but as the account it maps to is whatever the name on the cluster is, the true mapper should be rsv to rsv, not rsv to rsvuser. I may change this later, but will include whichever we end up with, in the final documentation.

Also forgot to change the ownership of the new rsv service cert and key to rsv account, remember to do that!

Looks like I didn't have an existing grid proxy, created one for myself (jrobertson), but RSV should be creating its own... Will re-run rsv probes to see. Also, to generate a proxy:

\$ grid-proxy-init -valid ???:??

where ???:?? indicates the hours and minutes that your proxy will last. OSG site recommends 500:00.

Looks like GRAM isn't happy, rsv still having troubles authenticating. However, there is also a big yum update that I am pushing now. Also, osg-configure should probably be re-run, after all config files are properly written and BeStMan is installed (I think that needs to go on mgmt...).

One issue was setting in /etc/osg/config.d/10-misc.ini, had the old gums server set on it (the CE), changed value to uscml-mgmt.fltech-grid3.fit.edu

Note: job contact and util contact are both set in /etc/osg/config.d/20-condor.ini

Note: in /etc/osg/config.d/30-gip.ini, changed node number from 20 to 18

Note: all bestman configuration stuff set in /etc/osg/config.d/30-gip.ini

Note: 30-rsv.ini contains a gums host line, changed that to uscml.mgmt as well

Note: 30-rsv.ini also contains a line concerning the location of condor_location, commented out as our location is default, I think. We shall see, keep an eye on this one.

All cleared and verified, running osg-configure -c: didn't work, it needs a proper user-vo-map to continue. I'm updating the VO members database, as running gums-host-cron on mgmt didn't fill the normal user-vo-map in /var/lib/osg. Updating the VO db takes a few minutes, and a few connections always time out or are refused, but apparently that's ok, OSG claims its due to such a large VO db, there's no way everything will always be up and running. User groups that didn't clear:

LIGO - connect timed out

compbiogrid - connect timed out

dayabay - peer not authenticated

engage - peer not authenticated

engage-ucsdgrid - peer not authenticated

gcedu - Connection refused

gcvo - Connection refused

gpn - No route to host

nysgrid - Connection refused

Seems to have updated ok though, so attempting to re-generate user-vo-map: that was another bust. Still going to scp the file from mgmt to the CE as the one on the CE is legit blank, at least MGMT has 2 lines. There's a "generate user-vo-map" option in the GUMS web GUI, but requires DN of the CE. Should pull this from OIM host cert ticket...

Well that got partially done, but cemon threw an error now, so it had to quit again. Going to disable cemon in the .ini files, we don't have it configured yet I don't believe. Scratch that, going to enable, but put the cemon host as the CE, it should be.

Yet another error, /etc/grid-services/available/jobmanager-condor doesn't exist. Somehow, globus-job-manager-managedfork never got installed. It now has been, which brought along jobmanager-condor with it. Retrying now. Ran osg-configure -v first of course, sorted out a few 15-managed-fork.ini bits. Had to re-disable cemon for now though, seriously screwy there. Will fix later, just trying to configure osg for now.

Holy poop, it reconfigured ok. Restarting RSV and GUMS now (in that order). Bummer, still getting "Condor-G submission failed" error. Also noticed in the Install CE instructions, section 7.2, need to add a variable "PER_JOB_HISTORY_DIR" to gratia. Adding line to condor_config.uscml on /mnt/nas0/condor/condor-etc/, as it wasn't set when checked.

Also note: control for local jobs vs. grid jobs is in section 7.6 of the CE install. Need to run a condor reconfigure after that though, that's a scary thought.

Now working on Cemon. Turns out it runs inside tomcat, and it runs the GIP before pushing the data to OSG. It needs either a file or directory `/var/cache/gip` to be owned by tomcat, but that doesn't exist on our system. Added said directory and ran a `chown tomcat:tomcat` to see if something happens. Actually, it isn't even installed...wtf. Lovely. Installing now, crossing fingers it doesn't blow up. Note: the install brought friends – `glite-ce-monitor` and `glite-ce-osg-ce-plugin`, as dependencies.

Now that that's installed, double-checking configs, then will attempt to yet again re-configure osg. Configured smoothly, this is good. Especially since cemon has been set back to "True" in `30-gip.ini`. Restarting RSV. Still getting Condor-G submission error. Hmmmm.

Well, still need to install Bestman and CEmon on the SE. Gonna start with CEmon. First things first, downgrading antlr so it matches that of MGMT. Going to do this on the CE in a moment as well, just for sanity. Done. Alright, first things first; have to request new host cert and http cert for the SE. Done, moved to appropriate locations, and changed owners and permissions. Commencing install of `osg-se-cemon`. It brings many friends:

Installed:

`osg-se-cemon.x86_64 0:3.0.0-4.osg.el5`

Dependency Installed:

`geronimo-specs.x86_64 0:1.0-0.M2.2jpp.12.el5.centos`
`geronimo-specs-compat.x86_64 0:1.0-0.M2.2jpp.12.el5.centos`
`gip.noarch 0:1.3.10-1.osg.el5`
`glite-ce-monitor.noarch 0:1.13.1-17.osg.el5`
`glite-ce-osg-ce-plugin.noarch 0:1.13.1-4.osg.el5`
`gums.noarch 0:1.3.18.009-20.osg.el5`
`gums-client.noarch 0:1.3.18.009-20.osg.el5`
`jakarta-commons-fileupload.x86_64 1:1.0-6jpp.1`
`mysql-connector-java.x86_64 1:5.1.12-2.el5`
`osg-configure-ce.noarch 0:1.0.45-1.osg.el5`
`osg-configure-cemon.noarch 0:1.0.45-1.osg.el5`
`osg-configure-gip.noarch 0:1.0.45-1.osg.el5`
`tomcat5.x86_64 0:5.5.23-0jpp.40.el5_9`
`tomcat5-common-lib.x86_64 0:5.5.23-0jpp.40.el5_9`
`tomcat5-jasper.x86_64 0:5.5.23-0jpp.40.el5_9`
`tomcat5-jsp-2.0-api.x86_64 0:5.5.23-0jpp.40.el5_9`
`tomcat5-server-lib.x86_64 0:5.5.23-0jpp.40.el5_9`
`tomcat5-servlet-2.4-api.x86_64 0:5.5.23-0jpp.40.el5_9`
`voms-api-java.noarch 0:2.0.8-1.5.osg.el5`

Note for later: in `30-gip.ini` on both CE and SE, need to change bestman version to appropriate one installed. Will need to re-run `osg-configure` after that, of course. Also note all other bestman info in the SE section of that file. Needs tweaking.

Back to CEmon. It appear BeStMan = hadoop? Somewhat confusing there, but it seems to be true. Leaving commented in for now, will come back to `30-gip.ini` later. Now digging through `40-siteinfo.ini`. Changing sponsor to OSG. Everything else was set properly, along with everything in `10-storage.ini`. Tried running `osg-configure`, it threw an error about `jobmanager-managedfork`, but it seems to have configured ok...? Marked it as not being used in `15-managedfork.ini`, but it still tries to assign a job manager:

ERROR Could not set the jobmanager-fork-poll to the default jobmanager

Ah well, shouldn't be a big deal...hopefully. Anyway, config ran fine, made sure `user-vo-map` is there after running `gums-host-cron`, and then started `gums-client-cron` (also popped into `chkconfig`). That concludes installing cemon! On to installing `gums-client` on the CE (it isn't there, uh oh!), and then to bestman for the SE. So, running `$yum install gums-client.noarch`, which also brings along its dependency companion `gums`. Installed, ran `gums-host-cron`, `chkconfig'd gums-client-cron`, and done!

Now, for the bestman install. First had to copy the hostcert.pem and hostkey.pem over to the bestman directory, apparently bestman needs a host cert under the guise of a service cert. Now, choosing only to run one GridFTP server, as we don't have much space and we'll only have a 1 Gbps connection to the outside world, so traffic won't be very heavy. Also opting for a custodial cache? Perhaps combo of custodial and volatile. We shall see. Anyway, now running \$ yum install osg-se-bestman:

Installed:

`osg-se-bestman.x86_64 0:3.0.0-9.osg.el5`

Dependency Installed:

`bestman2-client.noarch 0:2.3.0-16.osg.el5`
`bestman2-client-libs.noarch 0:2.3.0-16.osg.el5`
`bestman2-common-libs.noarch 0:2.3.0-16.osg.el5`
`bestman2-server.noarch 0:2.3.0-16.osg.el5`
`bestman2-server-dep-libs.noarch 0:2.3.0-16.osg.el5`
`bestman2-server-libs.noarch 0:2.3.0-16.osg.el5`
`bestman2-tester.noarch 0:2.3.0-16.osg.el5`
`bestman2-tester-libs.noarch 0:2.3.0-16.osg.el5`
`cog-jglobus-axis.noarch 0:1.8.0-4.1.osg.el5`
`glite-security-trustmanager.noarch 0:2.5.5-6.2.osg.el5`
`glite-security-util-java.noarch 0:2.8.0-3.2.osg.el5`
`jakarta-taglibs-standard.x86_64 0:1.1.1-7jpp.1`
`jetty-client.noarch 0:8.1.4.v20120524-2.osg.el5`
`jetty-continuation.noarch 0:8.1.4.v20120524-2.osg.el5`
`jetty-deploy.noarch 0:8.1.4.v20120524-2.osg.el5`
`jetty-http.noarch 0:8.1.4.v20120524-2.osg.el5`
`jetty-io.noarch 0:8.1.4.v20120524-2.osg.el5`
`jetty-jmx.noarch 0:8.1.4.v20120524-2.osg.el5`
`jetty-security.noarch 0:8.1.4.v20120524-2.osg.el5`
`jetty-server.noarch 0:8.1.4.v20120524-2.osg.el5`
`jetty-servlet.noarch 0:8.1.4.v20120524-2.osg.el5`
`jetty-util.noarch 0:8.1.4.v20120524-2.osg.el5`
`jetty-webapp.noarch 0:8.1.4.v20120524-2.osg.el5`
`jetty-xml.noarch 0:8.1.4.v20120524-2.osg.el5`
`lcas-lcmapi-gt4-interface.x86_64 0:0.2.6-1.1.osg.el5`

Seemed to install ok, now configuring to run with GUMS. Done. Everything seemed to configure fine, but upon starting everything up, the gratia probes aren't starting. Bestman still isn't connecting via srmping though, unfortunately. Probably has something to do with the gratia probes. This has been a lot of work in one night though, I'm calling it quits for now. Definitely good progress though, we're significantly closer to getting back up, I can feel it.

11/09/13

Digging straight into GUMS, it seems something still isn't quite right (RSV can't connect, so auth issues). RSV is properly set up on GUMS end, and tested with own account, so GUMS config looks sound. Looking through /var/log/tomcat5/trustmanager.log, looks like all our CRLs are failing. Found the DN for the CE though:

```
/DC=com/DC=DigiCert-Grid/O=Open Science Grid/OU=Services/CN=uscms1.fltech-grid3.fit.edu
```

Via GUMS, just generated a user-vo-map, copying to the CE now. Also going to try it for mgmt, see 1) if it differs, and 2) going to place it appropriately. Generated new one for mgmt, but it was identical as suspected. Still placed it appropriately, but going to have to re-configure osg on the CE.

Meanwhile, about these CRLs. Re-running /usr/sbin/fetch-crl3 on mgmt to see if that helps. Ran quietly. Other two (boot and cron) ran fine too...hmmm. Made sure osg-ca-certs-updater is installed, and started the cron service/made sure it was in chkconfig. Interesting fact: catalina.out is full of errors as well, it may be that the hostname is throwing things off (currently set to compute-1-0, seeing as gums auth stuff has localhost set in a few places. Changed with \$ hostname uscms1-mgmt.fltech-grid3.fit.edu, also changed for startup in /etc/sysconfig/network.

Damn, forced rsv to re-run after a restart, but that didn't fix it. Will check catalina and trustmanager in a moment. Trustmanager still has issues with the CRLs, that's gotta be a big chunk of it. Looking in /etc/grid-security/certificates, looks like many of the .r0 files are missing. Definitely missing a few, as the CE has everything. Might try copy-pasting if I can't find another re-population solution.

Solution not found, it seems an initial fetch-crl3 didn't get everything. Going to first try "removing" (move to alternative directory) the other files for a single CRL, re-run fetch-crl3, and see if it gets repopulated. If so, I will remove the rest of the excess and have it refresh all. If that doesn't work, I'll just start copy/pasting from the CE, at least for temps, then it will be able to updated them. (note, all "extra" files are actually links to databases within the certificates directory). Testing with CRL 0119347c. Nay, the mv to a different directory didn't work, going to scp from CE.

Note: that .r0 file also does not exist on the CE...going to output the ls's to text files and run a diff. Seems the CE only has complete sets of ones that don't exist on MGMT. In trustmanager.log, the only ones throwing errors are the ones with "extra" files that don't have .r0 files on either MGMT or CE. Interesting.

Taking a step back. RSV config could still be wrong, double-check on CE.

/etc/condor/config.d/00personal_condor.config got written over by the latest condor update, have to re-insert local config file line. Done, also fixed line for condor_config in /etc/osg/config.d/20-condor.ini. Hopefully will be a bit of a fix for osg-configure. Reconfigured smoothly, now for an rsv restart. Nay, no progress.

Note: squid was set to being used, but it is currently not.

Problem in 30-gratia.ini, the probe was set to condor:blah, not jobmanager.blah. Re-running osg-configure: threw an error "No such file or directory: '/etc/gratia/condor/ProbeConfig'", turns out we only have ProbeConfig.rpmsave and ProbeConfig.rpmnew. Perhaps this has been it all along? Interesting through, probe name in ProbeConfig is condor:blah..., changed back to condor:blah in 30-gratia.ini, "jobmanager" must have been a placeholder for condor.

Dammit, still didn't fix it. Calling it quits, have to get to other work now.

11/10/13

Prepping for the shutdown tomorrow morning, we've got FPL doing their annual 1-hr shutdown 7-8 a.m. Going to check the UPS, get some pictures, and make sure everything comes back up ok. Meanwhile, running a yum update on all the inactive nodes, then shutting down condor on them to keep them isolated for now.

Also updating the nas's while I'm at it. A few less packages than what's being updated on the nodes, but not by much. However, nothing from java was updated, just as a note.

Note for the node yum update, a few new conf files were made:

"warning: /etc/sysctl.conf created as /etc/sysctl.conf.rpmnew"

"warning: /etc/dbus-1/system.conf saved as /etc/dbus-1/system.conf.rpmsave"

Seriously need to clean out any and all remaining .rpmnew and .rpmsave files from everywhere, should NOT have those lingering.

Also, here's everything that got updated:

Removed:

kernel.x86_64 0:2.6.18-53.1.14.el5 kernel-devel.x86_64 0:2.6.18-53.1.14.el5

Installed:

kernel.x86_64 0:2.6.18-371.1.2.el5 kernel-devel.x86_64 0:2.6.18-371.1.2.el5

Updated:

aspell.i386 12:0.60.3-13

aspell.x86_64 12:0.60.3-13

autofs.x86_64 1:5.0.1-0.rc2.183.el5

bestman2-client.noarch 0:2.3.0-16.osg.el5

bestman2-client-libs.noarch 0:2.3.0-16.osg.el5

bestman2-common-libs.noarch 0:2.3.0-16.osg.el5

binutils.x86_64 0:2.17.50.0.6-26.el5

ccid.x86_64 0:1.3.8-2.el5
centos-release.x86_64 10:5-10.el5.centos
centos-release-notes.x86_64 0:5.10-0
coolkey.i386 0:1.1.0-16.1.el5
coolkey.x86_64 0:1.1.0-16.1.el5
cpio.x86_64 0:2.6-26.el5
dbus.i386 0:1.1.2-21.el5
dbus.x86_64 0:1.1.2-21.el5
dbus-libs.i386 0:1.1.2-21.el5
dbus-libs.x86_64 0:1.1.2-21.el5
dbus-x11.x86_64 0:1.1.2-21.el5
device-mapper-multipath.x86_64 0:0.4.7-59.el5
firefox.x86_64 0:17.0.10-1.el5.centos
ghostscript.x86_64 0:8.70-15.el5_9.3
glibc.i686 0:2.5-118.el5_10.2
glibc.x86_64 0:2.5-118.el5_10.2
glibc-common.x86_64 0:2.5-118.el5_10.2
glibc-devel.i386 0:2.5-118.el5_10.2
glibc-devel.x86_64 0:2.5-118.el5_10.2
glibc-headers.x86_64 0:2.5-118.el5_10.2
gnupg.x86_64 0:1.4.5-18.el5_10
gsoap.x86_64 0:2.7.13-5.el5
gtk2.i386 0:2.10.4-30.el5
gtk2.x86_64 0:2.10.4-30.el5
httpd.x86_64 0:2.2.3-83.el5.centos
httpd-devel.i386 0:2.2.3-83.el5.centos
httpd-devel.x86_64 0:2.2.3-83.el5.centos
hwdata.noarch 0:0.213.28-3.el5
initscripts.x86_64 0:8.45.44-3.el5.centos
java-1.7.0-openjdk.x86_64 1:1.7.0.45-2.4.3.1.el5_10
java-1.7.0-openjdk-devel.x86_64 1:1.7.0.45-2.4.3.1.el5_10
jglobus.noarch 0:2.0.6-3.osg.el5
kernel-doc.noarch 0:2.6.18-371.1.2.el5
kernel-headers.x86_64 0:2.6.18-371.1.2.el5
kpartx.x86_64 0:0.4.7-59.el5
ksh.x86_64 0:20100621-18.el5
libgcrypt.i386 0:1.4.4-7.el5_10
libgcrypt.x86_64 0:1.4.4-7.el5_10
lvm2.x86_64 0:2.02.88-12.el5
microcode_ctl.x86_64 2:1.17-5.el5
mkinitrd.i386 0:5.1.19.6-81.el5_10
mkinitrd.x86_64 0:5.1.19.6-81.el5_10
mod_ssl.x86_64 1:2.2.3-83.el5.centos
module-init-tools.x86_64 0:3.3-0.pre3.1.63.el5
nash.x86_64 0:5.1.19.6-81.el5_10
net-tools.x86_64 0:1.60-83.el5_10
nfs-utils.x86_64 1:1.0.9-70.el5
nscd.x86_64 0:2.5-118.el5_10.2
nspr.i386 0:4.9.5-2.el5
nspr.x86_64 0:4.9.5-2.el5
nspr-devel.x86_64 0:4.9.5-2.el5
nss.i386 0:3.14.3-18.el5
nss.x86_64 0:3.14.3-18.el5
nss-devel.x86_64 0:3.14.3-18.el5
nss-tools.x86_64 0:3.14.3-18.el5
osg-version.noarch 0:3.1.25-1.osg.el5
pcre.x86_64 0:6.6-9.el5
perl.i386 4:5.8.8-41.el5
perl.x86_64 4:5.8.8-41.el5

polycoreutils.x86_64 0:1.33.12-14.13.el5
popt.i386 0:1.10.2.3-34.el5
popt.x86_64 0:1.10.2.3-34.el5
procps.x86_64 0:3.2.7-26.el5
rpm.x86_64 0:4.4.2.3-34.el5
rpm-build.x86_64 0:4.4.2.3-34.el5
rpm-devel.i386 0:4.4.2.3-34.el5
rpm-devel.x86_64 0:4.4.2.3-34.el5
rpm-libs.i386 0:4.4.2.3-34.el5
rpm-libs.x86_64 0:4.4.2.3-34.el5
rpm-python.x86_64 0:4.4.2.3-34.el5
selinux-policy.noarch 0:2.4.6-346.el5
selinux-policy-targeted.noarch 0:2.4.6-346.el5
sos.noarch 0:1.7-9.66.el5.centos
sudo.x86_64 0:1.7.2p1-28.el5
tzdata.x86_64 0:2013g-1.el5
tzdata-java.x86_64 0:2013g-1.el5
vino.x86_64 0:2.13.5-10.el5_10
vo-client.noarch 0:48-1.osg.el5
wpa_supplicant.x86_64 1:0.5.10-10.el5
xorg-x11-drv-ati.x86_64 0:6.6.3-3.35.el5
xorg-x11-server-Xnest.x86_64 0:1.1.1-48.101.el5_10.1
xorg-x11-server-Xorg.x86_64 0:1.1.1-48.101.el5_10.1
xulrunner.x86_64 0:17.0.10-1.el5_10

Now shutting down condor on 2-7 with condor_off. Hmmm, seems that doesn't work properly, that could be a problem. Performing a full shutdown for the nodes that are idle on condor and have been updated.

In other news, according to the OSG troubleshooting guide, running a globusrun -a -r uscms1.fltech-grid3.fit.edu is saying that I can't authorize properly, but at least it's connecting. Very interesting... Holy poop, just tried running one as rsv, and got the following output:

gram_init failure:

GSS Major Status: General failure

GSS Minor Status Error Chain:

globus_gsi_gssapi: Error with GSI credential

globus_gsi_gssapi: Error with gss credential handle

globus_credential: Valid credentials could not be found in any of the possible locations specified by the credential search order.

Valid credentials could not be found in any of the possible locations specified by the credential search order.

Attempt 1

globus_credential: Error reading host credential

globus_sysconfig: Error with certificate filename

globus_sysconfig: Error with certificate filename

globus_sysconfig: File is not owned by current user: /etc/grid-security/hostcert.pem is not owned by current user

Attempt 2

globus_credential: Error reading proxy credential

globus_sysconfig: Could not find a valid proxy certificate file location

globus_sysconfig: Error with key filename

globus_sysconfig: File does not exist: /tmp/x509up_u525 is not a valid file

Attempt 3

globus_credential: Error reading user credential

globus_sysconfig: Error with certificate filename: The user cert could not be found in:

1) env. var. X509_USER_CERT

2) \$HOME/.globus/usercert.pem

3) \$HOME/.globus/usercred.p12

It's obviously looking for a personal cert that doesn't exist, but perhaps it contains a clue as to the cert RSV really should be using, or something like that?

Interesting note: all metric confs stored in /etc/rsv/metrics.

Dear god. So osg-ce never got installed, as in the main software package to control all other packaged with the 3.x update. Wth. Installing now, it's bringing friends. Hold up, seems the "squid" user is necessary as it's trying to install frontier-squid:

```
Installing : frontier-squid 1/4
WARNING: Adding missing group: squid.
WARNING: Adding missing user: squid.
mkdir: cannot create directory `/home/squid': Permission denied
chown: cannot access `/home/squid': No such file or directory
chmod: cannot access `/home/squid': No such file or directory
Installing : lcas-lcmaps-gt4-interface 2/4
warning: /etc/grid-security/gsi-authz.conf created as /etc/grid-security/gsi-authz.conf.rpmnew
Installing : osg-info-services 3/4
Installing : osg-ce 4/4
```

Installed:

osg-ce.x86_64 0:3.0.0-34.osg.el5

Dependency Installed:

frontier-squid.x86_64 11:2.7.STABLE9-16.1.osg.el5

lcas-lcmaps-gt4-interface.x86_64 0:0.2.6-1.1.osg.el5

osg-info-services.noarch 0:0.9-1

No difference in the rpmnew, removed. Also had to unpack /mnt/nas0/OSG/APP/etc for storage purposes, and enable squid in /etc/osg/config.d/01-squid.ini, but could set location to UNAVAILABLE for now.

Booyah, found the RSV service cert mapping page in OSG docu, reading now. Going to attempt to add the service cert, going into GUMS interface, to "Manual User Group Members" and adding rsv with following DN:

/DC=com/DC=DigiCert-Grid/O=Open Science Grid/OU=Services/CN=rsv/uscms1.fltech-grid3.fit.edu

Nope, didn't work. Not a bad thing to have present though. An idea: in one of the config files, it asks if condor-g is enabled. Changing this shouldn't change the method of authorization, and therefore shouldn't change anything, but it may be worth a shot? Later though, for now looking at alternatives. For instance, rsv profiler!

Didn't have the proper permissions set on rsvcert and rsvkey. Changing still didn't fix it though. Also need need to copy http cert over to httpcert2.pem, same for key, both necessary to host rsv on port 8443. Didn't fix anything, just keeping it up to date.

Ran rsv-control -profile, interesting stuff. Going to attempt to configure osg one more time though, see if that helps. Configured fine, restarting rsv. Still nope :(

Take in mind, the gratia probe for condor wasn't installed...wth.

Holding off on shutting down 2-8 and 2-9, as Zuo has a job running on 2-8 and I'd rather have a few slots open, just in case.

Last thing of the night; going to install Nagios, just so it's present at least, for the restart tomorrow - yum install nagios:

```
Installing : nagios 1/1
warning: /etc/httpd/conf.d/nagios.conf created as /etc/httpd/conf.d/nagios.conf.rpmnew
warning: /etc/rc.d/init.d/nagios saved as /etc/rc.d/init.d/nagios.rpmorig
```

Installed:

nagios.x86_64 0:3.2.3-3.el5.rf

Impressive, didn't add dependencies! Digging up the confs, no modding necessary, old config just saved as nagios.conf.old. Started the service up, seemed to run fine! Alright, on to squid, going to install this one as well tonight, if it won't take too long. Installing on the CE, as per the documentation on the OSG site.

Looks like when frontier-squid installed earlier, it attempted to make an account, etc, but didn't have permissions to create the squid home directory (normally on /mnt/nas0/home/\$USER). It got the proper entries into /etc/passwd and /etc/group though. Just inserted an account by hand, that should come up after the restart. Added necessary ports in iptables.

To configure as we see fit, we have to modify /etc/squid/customize.sh, and then either restart the squid service, or reload it if we don't want to shut it down (\$ service frontier-squid reload). However, starting it up may be a small problem for now, as it can't find its home directory. Hopefully that will kick over with the restart tomorrow.

11/11/13

Alrighty, bright and early, setting up for the power down. Only nodes still up are 1-0, 2-4, 2-8, the nas's, and the CE of course. Going to do a quick sweep through the building to make sure everything else is off. Also, I've forgotten the password for 2-4, so that node is screwed. At least it doesn't have anything crucial on it yet...

During shutdown, a bunch of GRAM errors were thrown, including stuff concerning glexc (some rules not set?) and stuff within lmaps.db. Couldn't catch it in time.

Also, a large quantity of mimed mail still resides on nas1, I assume it also exists on nas0. Regardless, shutdown on nas1 was surprisingly clean. Cracking UPS's now. Also, whooo, one of the light switches in the high bay works! Must be emergency power, probably should shut it off after a while.

All UPS check out fine, no corrosion, more detailed report to come. Nas1 came up beautifully, but has TFTP set as primary, going back in to change that.

Interesting, an email went out to the T3's, apparently RSV uses grid-proxy-init, which automatically uses 512 bit proxies. However, the new java package requires 1024 bit minimum, so a modification needs to be made. Went into /usr/lib/jvm/java-1.7.0-openjdk-1.7.0.45.x86_64/jre/lib/security/java.security and edited the following line from 1024 to 512:

```
jdk.certpath.disabledAlgorithms=MD2, RSA keySize < 1024
```

Didn't fix authorization, but will help later on.

Pushing jdk7 upgrade on MGMT, it needs to be up to date anyway, and perhaps that will smooth some of RSV's errors (note, there are still a few java errors in tomcat's catalina.out). Install went smoothly, did *not* remove older versions of java, or prevent oracle jdk from being installed later.

Running jdk7 push on SE, to perhaps fix bestman2 as well. Install went smoothly, fixed paths for tomcat5, bestman2 conf was fine already. Who, fixed the java error on rsv! But connection still refused now, so it's having the same issue as RSV now. That's good, at least it's on the same page. Finally, pushing java update on the CE. Went smoothly, restarted tomcat5 and RSV. Noticed that tomcat5 wasn't chkconfig'd, but Cemon needs it for GIP. Didn't fix auth issue, but at least everything important is running the proper version of java now.

Oh lord, just realized the clock on the SE is wrong, need to re-synchronize that. Done with: ntpupdate -u pool.ntp.org

Was just digging back through 30-gip.ini, changed the bestman version from 2.2.1.3.13 to 2.2.2.3.0, as we underwent an upgrade. Also changed this on the SE, ran an osg-configure.

Interesting, in 40-network.ini, although port ranges are specified, the globus files are not. Possibly the issue?

Also, globus-gatekeeper and globus-gridftp-server were not up, started them and chkconfig'd them.

Noticed that gridftp probe wasn't enabled, went into /etc/gratia/gridftp-transfer/ProbeConfig to enable it. Restarted globus-gatekeeper and globus-gridftp-server, then rsv. Force-ran all rsv probes, and the error changed! Went from:

Log:

```
000 (28761.000.000) 11/11 19:20:03 Job submitted from host: <127.0.0.1:39008>
```

```
...
```

```
020 (28761.000.000) 11/11 19:20:06 Detected Down Globus Resource  
  RM-Contact: uscms1.fltech-grid3.fit.edu/jobmanager-fork
```

```
...
```

```
026 (28761.000.000) 11/11 19:20:06 Detected Down Grid Resource  
  GridResource: gt2 uscms1.fltech-grid3.fit.edu/jobmanager-fork
```


to

Condor log file:

000 (28791.000.000) 11/11 19:28:41 Job submitted from host: <127.0.0.1:39008>

...

018 (28791.000.000) 11/11 19:28:43 Globus job submission failed!

Reason: 7 authentication with the remote server failed

...

Can't tell if this is progress or bad, hopefully good that it now knows that it can't authorize. Also haven't run an osg-configure yet, trying now. Error didn't change over the osg-configure, going to disable the gridftp probe to see if it reverts to first error. Nay, it does not. However, gridftp probe wasn't set in 30-gratia.ini. Has been set, osg re-configured, rsv being restarted now. Changed nothing. So whether gridftp is enabled or not in the 30-gratia file, same error is now the second one. Also doesn't matter if the gatekeeper and gridftp services are running or not. Or even if the probe is set up at the bottom of the 30-gratia.ini file. The change must've been somewhere before that happened. Giving up for the evening, I have other work to do unfortunately.

Except one last thing; the simple gridftp probe totally changed, it now says "login incorrect", as opposed to what it used to say, which was "unable to connect", so that's progress. Perhaps this is good. But look at /var/log/globus-gatekeeper.log

11/12/13

/var/log/globus-gatekeeper says that people are attempting to authorize, but aren't getting through. However, /var/log/messages is empty...interesting.

Looking in /etc/sysconfig/globus-gatekeeper, doesn't contain a line for the lcms debug, so I will add that momentarily. Before that, however, I will uncomment a number of options in the file that are commented out, including the cert info:

- grid-services directory
- cert-dir
- cert-file
- key file
- grid-mapfile

None of these helped, re-commenting them now.

11/14/13

Need to install CVMFS on the cluster, although technically not required for T3's by CERN, it's a good idea to have it on. OSG has instructions on install, and Bockjoo also sent around a CERN link for it. But for CVMFS to work, need squid to be up and running properly. Damn. Had account issues earlier, squid tried to automatically make its own account, but we've got weird mapping stuff, so it failed.

Going back in, verifying account locations are as needed (need actual directory in /mnt/nas0/home/). Confirmed. Copying over .bash_profile, .bash_logout, .bashrc, and .emacs from /etc/skel, as per the user creation script. Added squid to /etc/auto.home, also changed status in /etc/passwd to nologin. Reloading autofs, then restarting.

It worked! Started the frontier-squid service, got the following output:

Initializing Cache...

2013/11/14 22:39:26| Creating Swap Directories

Starting 1 Frontier Squid...

done

Alright, so squid works, should do some testing. First, enabled squid in /etc/osg/config.d/01-squid.ini, added location to be uscml.fltech-grid3.fit.edu, and running osg-configure now to have that present. Done, configured smoothly. Now, back to cvmfs. Added account in /etc/passwd, added directory in /mnt/nas0/home/, and added entry to /etc/auto.home. Autofs restart worked, account officially added. Did set is as nologin though, that may need to change?

Running software install, which brought friends: **yum install osg-oasis**

Installed:

osg-oasis.noarch 0:2-2.osg.el5

Dependency Installed:

cvmfs.x86_64 0:2.1.15-1.osg.el5 **cvmfs-keys.noarch 0:1.4-1.osg.el5**
fuse.x86_64 0:2.7.4-8.el5 **fuse-libs.x86_64 0:2.7.4-8.el5**
gdb.x86_64 0:7.0.1-45.el5.centos **oasis-config.noarch 0:4-2.osg.el5**

Edited /etc/fuse.conf to contain the following line, apparently needed for permissions: **user_allow_other**

Edited /etc/auto.master to allow mounting of the cvmfs accounts: **/cvmfs /etc/auto.cvmfs**

Restarted autofs for changes to take effect, restarted smoothly.

Edited /etc/cvmfs/default.local as necessary, given the example given in the OSG instructions. Lines worthy of note:

- CVMFS_REPOSITORIES is going to contain oasis, atlas, cms, and grid for now.
- CVMFS_QUOTA_LIMIT is in Mb, they recommend 20 Gb total, sticking with that.
- CVMFS_CACHE_BASE is default set to /var/lib/cvmfs, moving it onto nas0, as it has more space. Significant errors can apparently occur if this space isn't kept open.
- CVMFS_HTTP_PROXY is set to "DIRECT", a special option for small servers.

!!After each edit, the config should be reloaded. This can be done as such, for current version 2.1: **cvmfs_config reload**

Creating /etc/cvmfs/domain.d/cern.ch.local file, adding CVMFS_SERVER_URL so cvmfs can access the closest stratum one servers first.

Whooo, config works! Can ls -l a certain repo in /cvmfs, and it will pop up! Notifying Bockjoo now. AUGH, it needs to be installed on all the nodes...dear god.

Ok, running a yum install with cluster-fork for the osg-oasis package.

Now pushing /etc/fuse.conf from CE. Auto.master gets pushed by the rocks roll. Pushing autofs restart.

Now pushing /etc/cvmfs/default.local for the proper config, running cvmfs_config reload

Now pushing /etc/cvmfs/domain.d/cern.ch.local for servers, running final cvmfs_config reload.

Done, all is installed! Checked on a random node, can ls the proper locations with results.

11/17/13

Installing nagios on MGMT, put it on CE but nowhere else.

11/22/13

Attempting to figure out the GUMS authentication issue. Nothing is being written to /var/log/messages, even though that's where OSG says errors should be printing. Issue with lcmsaps, perhaps? Going through all lcmsaps-tagged emails and authentication-tagged emails to see if I can find something.

Couldn't find anything, submitted a ticket. Stupidly flagged the "security issue" option, however, so I'm essentially locked out now. Brilliant.

11/25/13

For ifcfg-eth2, DHCP_HOSTNAME was set to compute-2-4.local, changing to uscms1-mgmt.fltech-grid3.fit.edu, though it should technically be static anyway...

11/28/13

Yay, thanksgiving. I'm thankful for OSG being such an organized, helpful software group. OH WAIT.

12/09/13

Doug Johnson is a god in human flesh. I was an idiot and requested an httpcert for mgmt, but performed the request with uscms1.fltech, I forgot the -mgmt in there. Re-requested the appropriate cert, and the RSV probes came up (almost) perfectly. Had a small issue with the grid pool, but that was the account mapper not holding onto the pool name. Re-added that, and everyone can clear now :D

/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=jordan/CN=746428/CN=Jordan Alexander Robertson