

FALL 2009

+CRAB (CMS Remote Analysis Builder)

- must have a CMS account
- is a python program aka learn python?
- Commands:
 - crab.py -help and read it!!!
 - crab.py -create N to create N jobs (no submission)
 - crab.py -submit N -continue to submit the job you have created (mind the -continue)
 - crab.py -monitor -continue to monitor (very primitive!) your jobs and get automatically the output retrieved when the jobs are finished
 - crab.py -create all -submit all -monitor to do all in just one command

+Getting DOE private certificate

- Your request will be processed when an authorized agent verifies and validates the information in your request.
- request ID is 47765.
- You can check on the status of your request with an authorized agent or local administrator by referring to this request ID.
- After much "haggling" with DOE official and with the gracious help of Patrick i was given a certificate
 - imported by following link in
 - then export the key pair for use by globus grid-proxy-init
 - instructions: <http://www.doe grids.org/pages/cert-request.html#Globus>
 - note ssl is ss letter "L" not a number 1

+Getting a CMS Account

- <https://twiki.cern.ch/twiki/bin/view/CMS/WorkBookGetAccount>
- basically send them info, then they ask Dr. Baarmand to approve, he asks you for more info, you give it, he approves, they approve you

+Getting a CMS Computing account

- once you've got the registration done you can start with this round of emailing
- yippie-kie-yay

+Register with CMS VO

- can do once you have the certificate from DOE and are registered with CMS, Dont have to have cms computing account yet
- https://twiki.cern.ch/twiki/bin/view/CMS/SWGuideLcgAccess#How_to_register_in_the_CMS_VO
- first test that certificate works, if it doesnt follow instructions
- then once you have a working one click on the VO you want to join (in my case, CMS)
- follow instructions

+email techsupport to give a port hostname

8/25/09

+Put VMware workshop with windows on linux box

- logged in with patrick's account on vmware
- downloaded with .bundle
- used serial number, (find it in gmail just type vmware and search)
- windows serial number (on package) easy too
- after installing VMware you find it under system tools in menu bar
- for any help use <https://help.ubuntu.com/community/VMware/Player>
- used 50 GB disk space for windows

scp ./* (copies all files in current folder) xfave@uscms1.fltech-grid3.fit.edu:~
beautiful, take special note of the colon. doesn't work with a /

+registering to hypernews

- first ssh lxplus.cern.ch and use cern info to log in
- then follow instructions at <https://twiki.cern.ch/twiki/bin/view/Atlas/HyperNews>
- ssh hypernews.cern.ch
- to use hypernews: <https://hypernews.cern.ch/hypernews/cms/top.pl>
- may take a couple hours before sitedb recognizes your hypernews account tho even tho cern hypernews lets you log in immediatly

9/3/09

+CRAB Tutorial

- get certificates in with instructions found here:
http://jcatmore.home.cern.ch/jcatmore/getting_grid_cert.html
- follow instructions to set up enviroment
- i.e. <https://twiki.cern.ch/twiki/bin/view/CMS/WorkBookRunningGrid>
- have to run command at the beginning of each session in order to use crab from any

directory:

```
source /afs/cern.ch/cms/ccs/wm/scripts/Crab/crab.csh
```

- Then create with vi a tutorial.py file and a crab.cfg in your Tutorial folder. these two files' contents you just copy paste from internet. but have to use vi to create the files themselves
- then just crab -create, crab -submit, and crab -status to see whats happening :D
- when jobs are done use crab -
- to specify which job use crab -COMMAND -c WORKINGDIRECTORY
- working directory = name of this specific set of jobs. i.e. crab_02343 etc because when you create the job it creates this file with either a timestamp or you can specify the name within the crab.cfg file. k?
- *Note first time running tutorial, ten jobs submitted, 2 failed, 8 were successful.
- *second time, ten jobs submitted, 5 failed, 5 were successful

9/8/09

-for some reason CRAB (on lxplus) creates but wont submit jobs returns "error creating remote directory" and about 20 other lines

- *attempt 1: use exactly their config file (last time had some edits from Samir)
 - same error output.
- *attempt 2: set up crab area in my section of the cluster and try again
 - instructions: <https://twiki.cern.ch/twiki/bin/view/CMS/SWGuideCrabHowTo>
 - need instructions from samir's email:

```
Within just home directory i.e. ~
```

```
>> export SCRAM_ARCH=slc4_ia32_gcc345
```

```
>> source /mnt/nas0/OSG/APP/cmssoft/cms/cmsset_default.sh
```

```
>> scramv1 project CMSSW CMSSW_2_2_13
```

```
>> cd CMSSW_2_2_13/src
```

```
>> cmsenv
```

```
Then go back to home directory i.e. ~
```

```
>>source /mnt/nas0/OSG/APP/crab/CRAB_2_6_1/crab.sh
```

```
Now make own folder to do actual crab things i.e.
```

```
>>mkdir Tutorial or w.e.
```

```
>>cd Tutorial
```

```
>>vi your crab.cfg file and your tutorial.py file
```

```
>>crab -create etc, hope for the best
```

**Error: Still allows me to create jobs but won't submit them, for a different reason tho...

```
importerror: libssl.so.4: file or directory does not exist
```

```
possibly just need openssl097a.i386 0:0.9.7a-4.2.2
```

```
possibly just need a symbolic link from libssl.so in /usr/lib and on
```

libcrypto.so

```
Check if PATH is correct, talk to patrick?
```

```
TRY using one of Samir's config files instead
```

```
-source for UI is : /mnt/nas0/OSG/APP/glite/setup_glite_ui.sh
```

+put a flash player on Linux box (Ubuntu)

- *Must do as logged in under administrator account
- *use synaptic manager woot! tres facile
- *Make sure it installed it under the other users too, should by default be available

-general goals?

- *familiarity with CRAB
- *copying files to cluster so we can use them (datasets i think?)
- *

+Buy video card and power supply:

<http://www.newegg.com/Product/Product.aspx?Item=N82E16814127368>
 ask hohlmann for account info for newegg.
 for power supply: ATX 12V, and 24 pin
 for video card: PCI

+Build Computer...

Installed a new power supply

+screw into place

+must connect the 4 pin thing that says P2 on it to the place on the motherboard that

says power2

+connect the 2 by 8 or so pin thing in the obvious place

+connect the 1 by 4 pin things to the cd rom drives and the hard drive. doesnt matter

which to which

Hard drive

+screw into place

+connect to the power supply with a 1 by 4 pin thing

+connect to mother board with any one of those wide flat strips of wires that are

like 1 by 8 or so

CD Rom Drive

+Screw into place

+connect to power supply with a 1 by 4 thing

+connect to mother board with a flat wide strip, make sure its a different one

entirely then the one you used for the hard drive

+also take note of the jumpers which are directly next to the place where you plug in

the 1 by 4. might have to put a lil black thing on one of the pairs to make

it work. note it should say on top of the cd rom which place it needs one. if not,

keep trying combinations.

Test if it works,

+plug in

+turn flip on back to on position,

+push power on front, yup its that easy

Install OS, in this case it was windows

+have to tell it to boot from windows disk by going to BIOS screen

+needed a keyboard that wasnt a usb one, or change settings under Settings screen to

accept usb

+follow instructions

+Install Graphics Card for Samir

He put it into place and began installation process, BUT! then the resolution on his screen was really bad

Turned out that ubuntu 8.04 doesnt like the 9600 video card we purchased

Upgrade ubuntu, and should work

Patrick also tried some different things including logging in on single user mode and re installing it

9/29/09

+Put commands for Crab into .bashrc file so that happens automatically. where's the lxplus version?

literally just copy pasted the commands i typed every time into this file, now every time i log in they run

automatically. saves time.

+MUST fix this ssl.so.4 problem

maybe fixed it...

+ find /lib -name "libssl*" find [directory to search in] -name "name of the

file" and note i put an asterix at the end for a wildcard ending

response: /lib/libssl.so.6

```
+ln -s /lib/libssl.so.6 /lib/libssl.so.4
```

then exact same thing for libcrypto.so.4

This literally creates a /lib/libssl.so.4 file which does nothing but point to the so.6 version. if this is wrong, just delete the new file which has no other purpose anyways but pointing

+NEW problem: ASAP ERROR: unable to ship valid proxy to the server

```
in log file: asap-user-register --server vocms58.cern.ch --myproxy
```

```
-asap-user-register: error while loading shared libraries: libgridsite.so.1.1:
```

```
cannot open shared object file: No such file or directory
```

should be just as easily fixable, just need to find orig libgridsite file

```
Aha! found in /mnt/nas0/OSG/APP/GLITE/GLITE_3_1_26-0/external/usr/lib/libgridsite.so.1.1
```

**NOTE when doing a find and it keeps returning permission denied. that just means its lettig you know that during its search it attempted to check that directory but was not allowed to.

don't interrupt it. wait till it gets to the end and eventually after having searched through every where that you told it to look it hopefully will turn up a couple actual locations.

these will not necessarily print at the end however. so best choice is to pipe your find through less or more or something so you can look at results after it's finished. umkay?

```
i.e. find / -name "libgridsite*" | less
```

```
+then i just symbolically linked this file to /lib/libgridsite.so.1.1
```

+NEWER problem:

```
in crab.log file of the job it says:
```

```
executing: asap-user-register --server vocms58.cern.ch --myproxy
```

```
SOAP 1.1 fault: SOAP-ENV: Server [no subcode] "getProxyReq"
```

```
Detail: Not a member of the CMS VO
```

on twiki for crab

```
https://twiki.cern.ch/twiki/bin/view/CMS/SWGuideCrabFaq#Proxy_Grid_and_all_that_stuff
```

```
suggests: asap-user-register --server SERVERNAME
```

```
use the servername that they write in the logfile i.e. Server =vocms58.cern.ch
```

10/01/09

+Connect quarknet machine to Internet

harder than it should be, need to put in a network card and then probably install a device driver. seems to not have any IP.

```
under command prompt should be able to do ipconfig /all to get info but returns
```

nothing

```
using help and then searching under system info stated no available under each IP
```

address thing

```
MAC address is findable
```

hokay so...

want to just put in a network card. got one from the big box of goodies i've collected from stuff from the lab

you know its a network card cause its got the place you stick an ethernet cord in. then you fit it into whichever of the slots it fits in

and when you restart the computer (btw it should be unplugged when youre doing this)

hopefully you can auto install this new piece of hardware

if not..have to download the driver. i.e. go to google. type in the EXACT name of the network card you put in (will be written on the network card) in my case it was intel pro 1000 mt

then search for a download of the driver for that. then download that driver, save it to a flashdrive or something. and then download it onto xp machine

+Fix Problem with CRAB

```
did voms-proxy-init -voms cms
```

```
which manually checks my registration to the CMS VO
```

Guess what?

+NEWEST Problem:

```
in log file says: Calling ServerConfig cern and then nothing more
```

```
error output when try to do crab -submit is one of the traceback things with the final lines
```

as

+Trying with Samir's Config File instead

```
he sent me an email with a .cfg file and a .py file
```

```
i scp'd both onto the cluster and into thier own folder.
```

```
ran commands crab -create and of course error message:
```

```
"Your config file is not valid python: No Module named HLTrigger/Configuration/HLT_1E31_cff
```

```
had to update my CMSSW version
```

```
STEPS:
```


sigh...

error: DBSAPI.dbsapiException.DbsConnectionError: failed to connect in 03 attempts
call to DBS server

(https://cmsdbsprod.cern.ch:8443/cms_dbs_prod_local_09_writer/servlet/DBSservlet) failed

Http ERROR status '403'

Status detail : 'access to the specified resource (DN provided not found in grid-mapfile.) has been forbidden

(by the looks of things, cause oleksi also had a the exact same problem, its a problem with the website not me :) note patrick said its cause i'm asking it to use grid-mapfile.. which were not using anymore

*Attempt6: attempting to run it on lxplus in order to determine whether its a cluster issue or a website issue

-problems with Your config file is not valid python: No Module named
HLTrigger/Configuration/HLT_1E31_cff

changed the CMSSW version just like last time. upgraded

-new problem: problems trying remote dir check... please check stage out configuration paramters

*Attempt7: back on cluster, samir found a new publish datawebsite so i'm trying that.

then

?Where is it

?What lines to change check out: `less $CRABPATH/crab.cfg` on `lxplus.cern.ch` or
http://www.hep.fsu.edu/~jdhaas/help_pages/crab.html

?Do i just rerun all his things

?Publishing to DBS, using phedex?

***NOTES

cluster -fork -y then whatever command you want to run

the -y is if you need it to default yes to everything since when using can't put in any commands through keyboard

this sends a command to every node on the cluster thats been labeled a compute node.

also `ctrl -alt-f1-f2-f3` switches between shells. literally changes the screen to another log in

10/27/09

+cron!

want to make it do it something every hour so get to crontab and type `0 * * * *` followed by command

Steps:

`ssh nas-0-0`

must use `sudo` before commands

`sudo crontab -e`

is a vi document (use `i`, `esc`, `:wq`, etc.)

`0 * * * * echo `date` `bin/grep "th" /proc/net/rpc/nfsd` >> /var/log/nfsd/stats`

and ta da should work!

and use `crontab -e` to get to crontab or `crontab -l` to display whats in it already

<http://www.adminschoice.com/docs/crontab.htm> for more help

if need aid figuring out timing, just search for a crontab calculator

+Discover most efficient block size

website: <http://nfs.sourceforge.net/nfs-howto/ar01s05.html>

measure throughput btw nas and front end

use `dd`, `bonnie++`, `iperf`

-using `dd`

`sudo time dd if=/dev/zero of=/mnt/nas0/testfile bs=16k count=16384`

`sudo time dd if=/mnt/nas0/testfile of=/dev/null bs=16k`

*did 5 trials of each bs: 64,32,16,8,4,2
 *removed testfile periodically
 *see excel document with results
 ran dd on the nas too, to see results did
 time dd if=/dev/zero of=/nas0/testfile bs=16k count=16384
 and again 5 trials of all the dif sizes and reading and writing.

10/29/09

-using iperf
 */opt/iperf
 *man iperf

First decide which you want to be server

Log in to that one: nas or front end

do iperf -s -i5 or same thing with a -u to do UDP vs TCP

This sets up the server to start listening for the client, and
 on what protocol to listen for it.

the -i5 determines the interval in seconds between periodic bandwidth reports in this

case 5

need /opt/iperf/bin/iperf -s if you're on the nas, or same thing but -c

then log in to client

do iperf -c hostname -w 1M or same thing with a -u

the hostname is the hostname of the server, to find it type hostname and enter on the

server

when nas is client, use 10.1.1.1 for hostname of frontend not uscms1.fltech-

grid3.fit.edu for w.e. reason

the -w states the window size, which in this case...

the 1M sets it to send only one Megabyte

Then on the client it will print out one line with the bandwidth, on the server it will print
 out three bandwidths

if you do it with UDP, the important numbers are the jitter, and loss.

the bandwidth will be much slower cause its UDP

*worked fine with UDP and TCP with frontend as client

*only UDP worked ok with nas as client. front end returned bind failed: address already in

use

this problem solved by adding -p 5002 at end of each command. that way it listens on a
 different port. victory!

Explanation: this one is measuring purely bandwidth. so its like a theoretical measurement of
 how fast your data could be going.

dd on the other hand is the practical experimental measurement of how fast it is actually
 going. clear?

***NOTE

ps aux | grep nfs | wc -l

: processes working, piped to grep, and then counts the lines that appear so you know however many
 processes are going.

***NOTE

/proc/net/rpc/nfsd on the th line shows the number of seconds its spent in a working percent mode
 so if biggest number is on the third number than spends most of its time at 30 percent
 Dont want your biggest number to be in the 90s or 100s.

11/3/09

Fix Crab Publish Problem

possibly a problem with the website

"

error: DBSAPI.dbsapiException.DbsConnectionError: failed to connect in 03 attempts
 call to DBS server

(https://cmsdbsprod.cern.ch:8443/cms_dbs_prod_local_09_writer/servlet/DBSservlet) failed

Http ERROR status '403'

Status detail : 'access to the specified resource (DN provided not found in grid-mapfile.) has been forbidden

"

11/5/09

tried it with a new website for publication, specifically:

https://cmsdbsprod.cern.ch:8443/cms_dbs_ph_analysis_02_writer/servlet/DBSServlet which samir gave me

"no valid files to publish on dbs. your jobs do not report exit codes =0

so did crab -status again and sure enough it says they've been retrieved (since i did -getoutput) but

Wrapper Exit Code=60303 which means "files already exist on the SE"

looked on the CMS jobs database to see if any additional info. if i click on the name of the task and it takes me to the page with graphs if i hover over

application exit reason and it says file Pythia... already exists. which is the file specified as the output file in the config file. so its probably been made multiple

times with all the trials... must see if any of the other ones will let me publish.

trying again really quick,

then going to try going into the Pythia file and changing the name of the output file and changing its name in the config file and trying again.

11/10/09 trying it....

*Crab attempt 34? (TeV6) ha!

NOPE

also, Cluster had to go down because the school was doing a power outage, it was taken down early because it was running reallllllly slow as a result of the fact

That patrick broke something when we installed the new RAID card. he broke the back up battery pack which means it doesnt cache things. i think so it ran slow

now its all fixed.

*redoing performance tests as a result to get a better guage of how fast the cluster is now, and what the best block size would be.

*put root on machine by printer, and possibly all the xp's in lab, collect admin info on all of them

Problem: can't find libCore.so

just like last time, just find where you have a version of it using Find or looking in the root directory, then create a symbolic link between this location and the /lib/libCore.so file

*install two new computers

note if cd rom isnt loading things to check

is cd bad

is cd player working, (check jumpers on the back)

perhaps the cd player is just bad in general in which case hook up a new one

Video cards had to be switched around and tested etc. if the quality is bad its probably your video card not your monitor. to double check

of course hook up to a working machine somewhere else in the lab

Also, if things still arent working try rearranging the jumpers on your harddrive. try different combos till things test out.

So we had video connectivity and it was booting from the cd and installing windows and everything seemed lovely and then...

error with one small thing that it cant find it recommends checking the cd but we know that that part is ok. so we connected a new harddrive up to it and we'll test that out to see if its any good.

need one video card, one cd drive,

11/12/09

*Install GEANT4, CRY, and IDA onto ubuntu machine by printer

should all be in patrick's documentation , according to Kondo you basically just follow the instructions for Geant4 off the webpage

Documentation: CLUSTER Development, "install CRY before GEANT4"

11/16/09

*Crab attempt 407: fixing the original problem where it said thing about the files already being on the se.

logged into the se with ssh xfave@uscms1-se.fltech-grid3.fit.edu

removed all the working directorries that were there and thereby hopefully the data

created.

then re submitted jobs as TeV7
same error, wrapper exit code 60303

11/17/09

Today's Goals!

*windows on second computer,
get a new cd! yay!

*mac addresses for computers to Dr. Hohlmann

*get new cd drive from Downstairs--ka Check!

please note that this was more complex than had been thought, had to unscrew the bolts in the front of the nodes within the key lock thing.

*prepare newegg order

p.s. if a video card is PCI that lets you know that itll fit in the slots that we have for it. yup yup yup

oh and D-SUB is equiv to VGA

*label computers--ka Check!

all were named after elementary particles: muon, photon,

*new Crab test

*get CRY, EMACS, GEANT4, and IDA on ubuntu by printer

11/19/09

So... upon arrival, the windows system seemed to have downloaded successfully. turned it off in order to screw the cd drive and hard drive into place

upon turning it back on got this error: i386/ntkrnlmp.exe could not be loaded error code is 7

microsofts help pages were preetty good but turned out i had miswired stuff when screwing everything in. the flat wide ones..need a seperate one for the hard drive then the cd rom drive

*now working on CRY etc....

Emacs: Found it in the synaptic package manager, downloaded it, if i type emacs in my account it opens up. everything looks good

CRY:

found online, clicked on it when finished downloading, opened archive manager, then i clicked extract.

opened up terminal, found it, read the readme. first step should be 'make'

wasn't working. simple solution apparantly pretty common to ubuntu,

just ln -s /bin/bash to /bin/sh and i rmvd the /bin/sh file that was there and moved it to my folder (just in case!)

now would attempt to make but running into a new problem:

```
g++ -I. -c -o CRY AbsFunction.p Cry absfunction.cc
in constructor.....
```

```
error: 'strcmp' was not declared in this scope
```

```
error: 'atof' was not declared in this scope
```

```
sudo dpkg
```

```
sudo apt-get
```

11/24/09

get mac addresses

ubuntu: just type ifconfig into terminal, didn't work only showed ip which is the 430.234.234.234 looking thing. so went to BIOS (blue screen before booting that you get to by pressing delete)

then under advanced, chipset, onboard, enable mac stuff and tada! ifconfig will show it to you!

windows: ifconfig /all into command prompt but isn't working. went on BIOS and did same steps. and it works too!

**note Mac address looks like this : 09:08:34:49:AB or so. IP address looks like 430.234.234.234

then the ubuntu machine easily connected to the internet once the mac address had been registered(Dr. Hohlmann has to do that, he needs the port it will be attched to to, like PP3-5P) the windows machine didn't recognize the onboard, so i put a network card on there. then that

worked out fine.

called IT once both could connect to the internet. they remotely logged in on the windows and attached it to the FIT domain so now people can log on with their tracks account

they set up an onsite appointment for the Ubuntu machine because it doesnt support the remoteness.

when they came onsite for the ubuntu machine they were unable to connect it. they believe its a problem with the ports not connecting properly to the server. they contacted the server people, gave them

the port numbers and when they hear back from them will contact us again. hopefully will have done so before christmas. in the past have connected many ubuntu machines with no problems. Problem: for some reason the windows machine won't boot up without the cd in it....strange