

DUNE Computing Training May 2022

Anonymous Question Box

Ask your questions during the sessions; experts will answer!

TEMPLATE, PLEASE COPY and PASTE TO THE SESSION BELOW (NEXT PAGE)

Not answered ▾ Your question here

Answer:

[Expertname]

Session Thursday, May 12

[Welcome & Introduction](#)

[DUNE Basics: Storage space and data management](#)

[QUIZ TIME!](#)

[DUNE Basics: Introduction to art and LArSoft](#)

[DUNE Basics: Grid job submission + common errors](#)

[QUIZ TIME!](#)

Session Friday, May 13

[Welcome and Introduction \(Welcome Day 2\)](#)

[Code-Makeover: Switch to POMS](#)

[Expert in the room: Grid & batch job submission](#)

[Code-Makover: How to improve your code for better efficiency](#)

[Expert in the room: LARSoft: How to Modify a Module](#)

[Quiz: Best programming practices](#)

[Closing](#)

Session Thursday, May 12

Welcome & Introduction

I got it! ▾ This is an example: how to use this doc?

Answer:

[Claire] Like this! Please update the status once you get your answer!

DUNE Basics: Storage space and data management

Speaker: Kirby

I got it! ▾ I have issue with step-5 while trying to setup for training. I am not able to make directory, it is showing error as “mkdir: cannot create directory ‘/dune’: Permission denied”. Also for kr509 it is showing like “kx509: command not found”. {Thankyou, I got it.}

Answer:

[Timm] what machine are you logged into. If you do not have kx509 installed on your machine you can get one in CVMFS presuming that there is cvmfs on your machine.

{Schellman} Both of those commands need to be done on the gpvm or CERN Machine, not locally. I have edited the page to make this clearer.

Also you need to try `mkdir /dune/data/users/$USER` not `/dune` which is the top directory.

I got it! ▾ What is the difference between `/dune/data` and `/dune/data2`? Are they the same?

Answer:

[K. Herner] Yes, just different physical volumes behind the scenes.

[T. Junk] the user quotas are 200 GB on each, so you get a total of 400 GB of quota.

I got it! ▾ I have an issue in step 5 (5. Exercise! (its easy)) of the mission setup (the pre setup). When I run `source ~/dune_presetup_202205.sh`, the system returns “*dune_presetup_202205.sh: No such file or directory*”. Everything else works fine. I

Answer:

[Schellman] you make the file yourself:

Start in your home area `cd ~` on the DUNE machine (normally CERN or FNAL) and create the file `dune_presetup_202205.sh`. Write in it the following:

```
export DUNESW_VERSION=v09_48_01d00
```

```
alias dune_setup='source
```

```
/cvmfs/dune.opensciencegrid.org/products/dune/setup_dune.sh'
```

I got it! ▾ Is there a per-user quota on persistent dCache?

Answer:

[timm] There is not a per-user quota on persistent dCache but there will be pretty soon.

[Heidi] You can see the queue for the tape robot at

https://www-stken.fnal.gov/enstore/encp_enstore_system.html

It's streaming data at 200-400 MB/s across a very large # of drives. But DUNE accesses many 100 of TB/day, and we share with other experiments. So it can sometimes take several days to get your data back off tape. There is a prestage mechanism that allows you to ask ahead of time.

I got it! After a couple of minutes I usually get kicked out of the dunegpvm (broken pipe). Currently I have a work around for this where I have a `/s` command continuously run so I don't get kicked out. Is there a better way to solve this?

Usually I get kicked out while I am doing other things on my pc that relate to what im doing on dunegpvm, but no worries I get why this policy exists. Thank you!

[Update] Does running a command like `/s` use dune resources? If so I will stop doing this.

Yes this is what I currently do.

Oh I see, I didnt know that was a thing. I will look up the VNC sessions. Thank you!

Answer:

[Tom] The inactivity timeout is a policy the computing division enforces. Generally it's a good idea not to leave logged-in sessions open indefinitely, as someone could use your account if you leave your desk. Of course if you are at home it can be annoying. Yes, you need to keep interacting with the dunegpvm.

i haven't tried it – one might write a script with an `ls` command, a `sleep` command, and a loop around it so it keeps doing this forever, just making output. It's not an issue of resource usage, it's a security issue.

What I do is I set up a VNC session and when I get kicked off, I just reconnect to it and everything that was running is still running and I don't have to set up all my windows or root sessions or whatnot.

VNC instructions:

https://wiki.dunescience.org/wiki/DUNE_Computing/Using_VNC_Connections_on_the_dunegpvms

You can run processes in the background with a `&` at the end of the command line. Be sure to redirect stdout and stderr or catch output that would go to the terminal that will eventually get disconnected. There's the `nohup` command but I found that processes usually stay running after a logout or kick-off even without using `nohup`.

[Heidi Schellman] I've tried opening an xterm and running `top` in it. I don't have so much trouble with the gpvms but my university VM kicks me off all the time and it is very annoying. I find that `nohup` is a must for me if I'm going to get logged out

[Tom] I kept a terminal logged in to a gpvm running `"top"` since this morning and it hasn't gotten kicked off yet (four hours or so).

I got it! Does the size in number of bits matter in the `voms-proxy-init` command in the example?

Answer:

[timm] no.. we use the default everywhere and we haven't had problems.

[timm] Comment—the instruction in the `"Bash"` box just before exercise 2 should include a `"setup ifdhc"` to make it work. Also note that if you don't do the `voms-proxy-init` before the `ifdh` command, `ifdh` will try to make a proxy for you but it will not always do what you want.

I got it! - setup ifdh says that running setup requires administrative privileges (which I don't have!)

Answer:

[schellman] That means you did not run dune_setup. Long story is that there is a lot of software that gets set up and after that setup means something different than on a raw operating system.

We need to update the documentation to remind people that you need to run dune_setup every time you log in to the system.

I got it! - If I used the bare mkdir command to make the directory instead of ifdh is that a problem?

Answer:

[kirby] It will work ok but you should not do that in general. Should get used to not doing regular mkdir, ls, etc. commands.

[Tom] regular commands usually work but sometimes they do not. The typical symptom if they do not is the command will get stuck and if you look at it with "ps" or "top", the command will be in the "D" state, which means "non-interruptible i/o wait". Even killing such processes doesn't work a lot of the time. Simple things like mkdir and ls usually succeed, but any command that writes to dCache has a much higher probability of getting stuck. One of the dunegpvm's has a bunch of stuck "vi" sessions where someone tried to edit a file in dCache with a text editor.

I got it! - I'm getting: xrdfs: /lib64/libstdc++.so.6: version `CXXABI_1.3.9' not found (required by xrdfs)

Answer:

[timm] what version is your xrdfs. (rpm -q xrootd-client)?

[Yoann] I logged out / in and now it's fixed. **It happens after I do " setup ifdhc "**!

xrootd-client-4.12.6-1.1.osg35.el7.x86_64.

I got it! - For some reason I don't have permissions for /dev/null and cannot copy anything there. How can I fix this?

Answer:

[Kirby] What node are you working on?

>[Noah] dunegpvm07?

[Kirby] What is the output of this command? "ls -alrth /dev/null"

>[Noah] crw-rw-rw- 1 root root 1, 3 Feb 16 02:43 /dev/null

[Kirby] Can you cut and paste the command you issued along with the error message?

[Noah] ifdh cp -D /pnfs/dune/scratch/users/nvaughan/DUNE_tutorial_May2022/my_first_login.txt /dev/null

error: globus_xio: Unable to open file /dev/null/my_first_login.txt

globus_xio: System error in open: Not a directory

globus_xio: A system call failed: Not a directory

Perhaps you forgot a -D to indicate destination is a directory?

[Tom] Is the error on /dev/null, or from the input file? Make sure you have a valid kerberos ticket and a VOMS proxy to read the input file. You can test just the /dev/null part with a Posix command: cp ~/.profile /dev/null (or any file really).

Oh, the -D is for destinations that are directories. /dev/null is not a directory, it is a "file".

[Noah] Oh duh. Works no problem now. Thanks.

[Tom] That error message is totally misleading – normally people forget to add -D when they need it; but in your case you had an extra one and got the same advice. Often in HEP software a helpful error message can be the opposite because a specific case was not considered. Keeps us on our toes!

I got it! ▾ What is metadata?

Answer:

[Schellman] It is the data that describes data. For example, in the US, if you make a phone call, the authorities can ask a court to allow access to the metadata (who you called, from where, when) but not access the content of your phone call. That's metadata. The phone call itself is the data. In dune the data is, for example, the 8 GB raw data while the metadata is a few kb of descriptive information. <https://docs.dunescience.org/cgi-bin/sso/ShowDocument?docid=24823> chapters 9-11 go into this a bit.

Not answered ▾

```
xrdfs root://fndcal.fnal.gov:1094/  
ls /pnfs/fnal.gov/usr/dune/scratch/users/${USER}  
[ERROR] Server responded with an error: [3011] Path not found:
```

What should I do?

Answer:

[timm] What is your user name, and was the \${USER} environment variable set at the time? (printenv | grep USER) will say.

QUIZ TIME!

I got it! ▾ What does A do on question 6? (ls -R /pnfs)

Answer: [timm] /pnfs space has a database behind it. A recursive ls -R would cause every single entry in said database to be accessed and if it didn't crash the database entirely, it would cause extreme load and cause the storage admins to try to hunt you down and send you to re-education camp.

I got it! ▾ Where can we find the answers with explanations for the quiz?

Answer:

[Claire] We will put an update

DUNE Basics: Introduction to art and LArSoft

I got it! ▾ What is a fcl/fhicl file

Answer:

[Claire] It is the name of the configuration file read by LArSoft

[Jake] See their old [redmine page](#) and the [quick start guide](#) for an in-depth description.

PRO-TIP [Schellman] fcl/fhicl uses #include and has the concept of a path just like your C++ code.
\$FHICL_FILE_PATH shows the path.

DUNE Basics: Grid job submission + common errors

Speaker: Ken Herner

QUIZ TIME!

Session Friday, May 13

TEMPLATE, PLEASE COPY and PASTE TO THE SESSION BELOW

Not answered ▾ Your question here

Answer:

[Expertname]

Welcome and Introduction (Welcome Day 2)

Code-Makeover: Switch to POMS

Not answered ▾ How do I know how much memory I need to request the first time I submit a job? Is that info somehow available from running the program interactively?

Answer:

[Heidi] Art jobs print out memory usage at the end. This was running an old version of reco on the first real data from protoDUNE-SP. The file actually contained 100 events so it is likely this is an underestimate of the memory use. The same out file also tells you how much CPU time you used in total and per event. That helps in estimating the time/file to request. Reco may take up to 30 hrs for 100 events.

TrigReport ----- Event Summary -----

TrigReport Events total = 10 passed = 10 failed = 0

```

TrigReport ----- Modules in End-Path: end_path -----
TrigReport Trig Bit#      Run   Success   Error Name
TrigReport   0   0       10      10       0 out1

```

```

TimeReport ----- Time Summary ---[sec]----
TimeReport CPU = 3241.881160 Real = 3255.149638

```

```

MemReport ----- Memory Summary ---[base-10 MB]----
MemReport VmPeak = 3010.84 VmHWM = 1877.64

```

PRO-TIP [Heidi] Note to add - generally unwise to ask for more than about 5000 files in a batch job. There are ways to split up larger datasets to run more leanly.

I got it! ▾ What is a json file?

Answer:

[Schellman] In this case the json file is the file description metadata. Art stores an internal file description and that description is extracted to make a short external file that describes the data. That file is in json format (a standard python compatible data format) and is used to catalog the data file in the files database.

Try typing

```
samweb get-metadata 1130_1_cb.test --json
```

To see minimal metadata for a recent coldbox test.

Json and fhicl are both used a lot for configuration. Json maps almost directly to python dictionaries and lists which is nice. Getting it to play well with C++ is a bit more involved and it doesn't allow includes or comments.

Expert in the room: Grid & batch job submission

Not answered ▾ Are there FD MC ntuples publicly available, we produced some by ourselves, took a while, would be good to know if they are already available

Answer:

[schellman] We are generally depending on the physics groups to coordinate this. But we do need a central place to link to the physics group pages. There is a dune-data.fnal.gov page but it is out of date. Jake Calcutt made a

Here are reco file locations

https://wiki.dunescience.org/wiki/ProtoDUNE-SP_datasets is maintained by the DRA group and looks up to date.

https://wiki.dunescience.org/wiki/PDS_Physics_and_Simulation_Working_Group has the most recent photon simulations

https://wiki.dunescience.org/wiki/High-E_and_Non-accelerator_Physics is the HEP group and a bit out of date (maybe)

Latest Far Detector Vertical Drift MC samples:

https://wiki.dunescience.org/wiki/Vertical_Drift_FD_simulation_and_reconstruction

For PDSP group tuples are made via code in

<https://github.com/DUNE/NTupleProd>

that produces tuples, then merges them and keeps the metadata. The final outputs are documented at:

https://wiki.dunescience.org/wiki/Analysis_NTuples

and you can find them in sam at:

```
samweb list-files " application pdspana and data_tier root-tuple"
```

One issue is that it is good to produce sam metadata when you produce a tuple. That makes them easier to read back. The production team does this with the setups you saw today but most people doing it on their own do not.

I got it! ▾ cp

```
/dune/app/users/kherner/may2022tutorial/localProducts_larsoft_v09_48_01_e20_prof/setup-grid  
/dune/app/users/${USER}/may2022tutorial/localProducts_larsoft_v09_48_01_e20_prof/setup-grid
```

Above code is not working for me , giving me error of :

```
/dune/app/users/${USER}/may2022tutorial/localProducts_larsoft_v09_48_01_e20_prof/setup-grid': No such  
file or directory
```

Answer:

Does the destination directory exist? Probably one of the directions in this path

(/dune/app/users/\${USER}/may2022tutorial/localProducts_larsoft_v09_48_01_e20_prof/) hasn't been created yet.

[Schellman] probably you need to do (the -p makes the directories all the way down)

```
mkdir -p  
/dune/app/users/${USER}/may2022tutorial/localProducts_larsoft_v09_48_01_e20_p  
rof/
```

```
cp  
/dune/app/users/kherner/may2022tutorial/localProducts_larsoft_v09_48_01_e20_p  
rof/setup-grid \  
/dune/app/users/${USER}/may2022tutorial/localProducts_larsoft_v09_48_01_e20_p  
rof/setup-grid
```

I got it! ▾ What is meant by prestige and why is this a necessary

Answer:

Discussed in the data management and storage sections on Thursday; there is not enough dCache disk to keep all possible data from all experiments on disk simultaneously, so files will fall out of cache to make

room for other files according to the Least Recently Used algorithm (files that have not been accessed in a long time are removed first). Prestaging a dataset ensures that all files in a dataset are on disk before the jobs run. Otherwise if your job hit a file that was only on tape, it would have to wait for a tape fetch, which can take a long time. While waiting, it would be sitting idle at 0% efficiency, effectively blocking the resource for others.

Syntax is:

```
samweb prestige-dataset --dataset=<yourdataset> --parallel=10
```

runs parallel prestige processes. Do not do `--parallel = VERY LARGE NUMBER` please... it will not make the poor little* tape robot go faster.

*the tape robot is actually big and strong and would definitely win in a fight. That robot arm moving at 30 mph packs a punch.

Please note the syntax above for prestaging does work, but will NOT reset the LRU clock on the files in the dataset. The way to reset the LRU clock for sure is to do

```
samweb run-project --defname=yourdataset --schema https 'echo %fileurl && curl -L --cert $X509_USER_PROXY --key $X509_USER_PROXY --cacert $X509_USER_PROXY --capath /etc/grid-security/certificates -H "Range: bytes=0-3" %fileurl && echo'
```

Code-Makeover: How to improve your code for better efficiency

I got it! ▾ How do I get valgrind set up?

Answer:

[Schellman] It is available automatically with dunesw but you can also do
`ups list -aK+ valgrind` to find the versions and then set it up

I got it! ▾ How do I run valgrind on a larsoft executable?

Answer:

[Schellman] Tom showed arguments `<program.exe> arg1 arg2 ...` so likely
`valgrind --leak-check=yes --suppressions=$ROOTSYS/etc/valgrind-root.supp myprog arg1 arg2`

`lar` is `myprog` and then the arguments for `lar`

I got it! ▾ What is Arm DDT?

Answer:

[schellman] instructions are in the after-lunch

[Tom] It's a gui interface to the GNU debugger gdb. It used to be distributed by a company called Allinea, but Arm bought it a few years ago and renamed it. We distribute it as a UPS product. There is a Flexlim license at Fermilab, so you may not be able to run it out of the box using the CVMFS distribution of it at your university because of the lack of a license server.

I got it! ▾ What are voodoo ways of seeing what is going on on a gpvm

Answer:

[Tom - transcribed] Useful commands:

top
htop
cat /proc/cpuinfo
ls -l - this one surprised me. You can get a listing of all the names of any file that is opened by anyone.
ps - lists all subprocess of the process where you executed the command
ps -u `whoami` - lists your processes
ps aux - lists all processes
ps auxwww - the same, but has a super-wide output format so you can get all the command-line arguments for whatever everyone else was doing.

One thing I do from time to time is

```
ps auxwww | grep pnfs | grep D
```

which lists processes that are in an uninterruptible i/o wait state that mention the pnfs dCache mounts. This won't catch them all as the current working directory for a user's process may be in pnfs, or a file list may have been specified. So

```
ps auxwww | grep D
```

catches more. It also catches some things that just have "D" in them even if they are not in the uninterruptible i/o wait state.

htop show you top by CPU

cat /proc/cpuinfo also shows you info

<https://fifemon.fnal.gov/monitor/d/000000004/experiment-overview?orgId=1&var-experiment=dune&viewPanel=30> shows the load across the gpvm's

Apparently if you say

strings <binary> it will print out all the strings... works on rootfiles and executables and shared libraries. There is some dreck in there from binary strings getting interpreted as printable stuff.

Expert in the room: LARSoft: How to Modify a Module

Not answered ▾ What do I need to set up and where

Answer:

[Schellman] You need to run the main dune setup script and setup dunesw every time you log in. You want to be doing code builds on dunebuild01 but DO NOT run your executable there...

Not answered ▾ I am new to DUNE computing, why would I ever want to do any of this building?

Answer:

[Schellman] To start, probably not. But sooner or later you will need to write your own algorithms to do something and then it is a good idea to know how this stuff works. This is why we need to keep the written part of the tutorial available so when you suddenly do need it you have it. ProtoDUNE II is going to be real fun as we get new data and need to figure out new ways of reconstructing and analyzing it.

But also, the reason HEP people get the big \$ if they go to industry is that we know how to wrestle 1,000,000 line code bases and win.

[Tom Junk] If the code to do what you want to do hasn't been written yet, and it needs to be run in the simulation or reconstruction stage, then it needs to go in the framework with all the other sim+reco software.

Quiz: Best programming practices

Miscellaneous