

**Programming
Level-up**

Jay Morgan

Linux

What is Linux
The command
line

Shell Scripting

Writing bash
scripts

High
Performance
Cluster

Getting started
Submitting jobs
A guided walk
through

Programming Level-up

An Introduction to using Linux

Jay Morgan

7th October 2022

Outline

Programming Level-up

Jay Morgan

Linux

What is Linux
The command line

Shell Scripting

Writing bash scripts

High Performance Cluster

Getting started
Submitting jobs
A guided walk through

1 Linux

- What is Linux
- The command line

2 Shell Scripting

- Writing bash scripts

3 High Performance Cluster

- Getting started
- Submitting jobs
- A guided walk through

What is Linux?

Programming Level-up

Jay Morgan

Linux

What is Linux

The command
line

Shell Scripting

Writing bash
scripts

High Performance Cluster

Getting started
Submitting jobs
A guided walk
through

- Linux is a popular operating system (OS) like Windows, or MacOS.
- Unlike these other two OSs, Linux is open source, which means the source code is freely available to look at and modify.
- As its open source, its very possible for anyone to build their *own* version of Linux or build on top of Linux to create their own *Distribution* of Linux.

What's a Distribution?

Programming Level-up

Jay Morgan

Linux

What is Linux

The command
line

Shell Scripting

Writing bash
scripts

High Performance Cluster

Getting started
Submitting jobs
A guided walk
through

A distribution can be considered like a *flavour* or version of Linux. There are many popular flavours that attempt to meet different needs from different users. For example:

- Ubuntu – typically the first Linux experience people will have. Attempts to be very user friendly.
- Fedora – stable and secure distribution while also providing up-to-date packages.
- Arch Linux – strong focus on customisability rather than user friendliness with bleeding edge packages.

What's a Distribution?

Programming
Level-up

Jay Morgan

Linux

What is Linux

The command
line


Shell Scripting

Writing bash
scripts

High
Performance
Cluster

Getting started
Submitting jobs

A guided walk
through



Defining Traits of Linux

Programming Level-up

Jay Morgan

Linux

What is Linux

The command
line

Shell Scripting

Writing bash
scripts

High Performance Cluster

Getting started
Submitting jobs

A guided walk
through

While we have said that Linux is open source, there are many other traits that make it stand out from other operating systems:

- Complete control of how the system operates.
- The level of flexibility and automation that you can get from using the Linux command line.

While there are many other traits, these two are going to be what we're going to focus on.

What's a command?

Programming Level-up

Jay Morgan

Linux

What is Linux

**The command
line**

Shell Scripting

Writing bash
scripts

High Performance Cluster

Getting started

Submitting jobs

A guided walk
through

While many recent versions of Linux makes things more accessible via GUIs, they will never be a substitute for using the command line.

We're going to learn how to control the system via the command line, via a **shell**. A shell, like the Python REPL we've already seen, is waits for you to input commands, executes the command, and prints the output if there is output to print.

A Linux command is a call to a program optionally followed by some arguments. For example, if we want list out the files and folders in the directory, we would use the `ls` (list) command:

1

```
ls
```

What is a command?

Programming Level-up

Jay Morgan

Linux

What is Linux

**The command
line**

Shell Scripting

Writing bash
scripts

High Performance Cluster

Getting started

Submitting jobs

A guided walk
through

The `ls` command comes with a number of optional flags and arguments that we can add onto the call. When calling a command a flag is something that begins with a `-`, for example `-l` tells `ls` to list the directory in a list format.

2

```
ls -l
```


What is a command?

Programming Level-up

Jay Morgan

Linux

What is Linux

The command
line

Shell Scripting

Writing bash
scripts

High Performance Cluster

Getting started

Submitting jobs

A guided walk
through

We have supplied the `-l` flag. There are many other flags for `ls`, like for example, the human readable file systems with `-h` or show hidden files (files that start with a period) with `-a`.

When we're using multiple flags we could write

```
3  ls -l -h -a
```

Or:

```
4  ls -lha
```

What is a command?

Programming Level-up

Jay Morgan

Linux

What is Linux

The command
line

Shell Scripting

Writing bash
scripts

High Performance Cluster

Getting started

Submitting jobs

A guided walk
through

Sometimes commands take optional positional arguments. Going back to our list directory command, where, by default, it will list the current directory. But instead we can tell the command to list a particular directory by supplying the path as an argument

```
5 ls images/ -lha  
6 # or ls -lha images/ works too
```

What is a command?

Programming Level-up

Jay Morgan

Linux

What is Linux

The command
line

Shell Scripting

Writing bash
scripts

High

Performance
Cluster

Getting started

Submitting jobs

A guided walk
through

How do I know how to use a command? Well that's where another command comes in. It's called `man` (short for manual). If you pass another command to the `man` command, the documentation will be shown in the terminal, e.g.:

```
7  man ls  # display the documentation for ls
```

The documentation should list all the different flags and arguments, describe what they mean, and sometimes give example or most common usage of a command.

When the 'man page' is display, you can scroll up and down the page using your arrow keys, and page-up and page-down. When you're done reading, just hit the 'q' character

Very useful commands

Programming Level-up

Jay Morgan

Linux

What is Linux

The command
line

Shell Scripting

Writing bash
scripts

High

Performance
Cluster

Getting started

Submitting jobs

A guided walk
through

I am going to go through some of the most common commands just to make sure that you're familiar with the typical usage.

We've already seen `ls` to list a directory. The command to move to a directory is `cd` (change directory), that takes an argument of filepath to move to:

```
8  cd ~ # tilde is short-hand for the 'home directory'
9  cd ~/Documents/My\ Files # go to Documents and then to "My
   ↪ Files"
10 cd # no argument, by default goes to the home directory
```

Very useful commands – mkdir

Programming Level-up

Jay Morgan

Linux

What is Linux

The command
line

Shell Scripting

Writing bash
scripts

High Performance Cluster

Getting started
Submitting jobs
A guided walk
through

Sticking with the them of directories, to make a new directory we use `mkdir`, whose argument takes the name of the directory we want to create:

11

```
mkdir my_new_directory
```

You can create a many level nested directory structure all at once using the `-p` (parents) flag, that tells `mkdir` if the parent directory of the target directory doesn't exist, create it.

12

```
mkdir photos/2020/01/05 # won't work unless photos/2020/01 exist
```

13

```
mkdir -p photos/2020/01/05 # this will work
```

Very useful commands – cp

Programming Level-up

Jay Morgan

Linux

What is Linux

The command line

Shell Scripting

Writing bash scripts

High

Performance Cluster

Getting started

Submitting jobs

A guided walk through

To copy a file or directory, we can use the `cp` command. Here we are copying a file, where the first argument is the filepath of the file you want to copy and the second argument is the filepath where the copy should be placed.

```
14 cp my_old_file my_new_file
```

By default (without a flag), `cp` will not work with directories, for that you have to use the `-r` (recursive) flag

```
15 cp -r data/ data-backup
```

Very useful commands – mv

Programming Level-up

Jay Morgan

Linux

What is Linux

The command
line

Shell Scripting

Writing bash
scripts

High

Performance
Cluster

Getting started

Submitting jobs

A guided walk
through

The syntax of moving a file is similar to that of `cp`:

16

```
mv old_file new_file
```

Except that it works for both files and directories without any flags. `mv` can also be used to **rename** files, that's all renaming is: moving a file to the same directory under a different name.

Very useful commands – rm

Programming Level-up

Jay Morgan

Linux

What is Linux

The command
line

Shell Scripting

Writing bash
scripts

High Performance Cluster

Getting started
Submitting jobs

A guided walk
through

To remove a file us `rm`:

```
17  rm file_to_delete
```

If you want to delete a directory, use the `-r` (recursive) flag:

```
18  rm -r directory_to_delete/
```


Very useful commands – cat

Programming Level-up

Jay Morgan

Linux

What is Linux

The command
line

Shell Scripting

Writing bash
scripts

High Performance Cluster

Getting started
Submitting jobs
A guided walk
through

cat stands for concatenate, i.e. concatenating the contents of two or more files:

```
19 cat file1 file2
```

The result is that the concatenation of these two files will be printed to the screen. If you wanted to put the result into its own file you would redirect the output using >

```
20 cat file1 file2 > newfile
```

Since cat reads the file and prints it to screen it is a very handy way to view the contents of a file, even if it was not intended for that.

Very useful commands – pwd

Programming Level-up

Jay Morgan

Linux

What is Linux

The command
line

Shell Scripting

Writing bash
scripts

High Performance Cluster

Getting started

Submitting jobs

A guided walk
through

Sometimes you may get lost when moving directories. `pwd` prints the current working directory from the root directory, i.e. the path that is printed is an absolute path.

21

```
pwd
```

Very useful commands – find

Programming Level-up

Jay Morgan

Linux

What is Linux

The command line

Shell Scripting

Writing bash scripts

High Performance Cluster

Getting started

Submitting jobs

A guided walk through

If we want to list all files of a certain type, we can use the wildcard `*` that we've seen before:

```
22 ls *.jpg # list all files that end with .jpg
```

However, this will only list for the current directory. Perhaps the better way to find files will be using the `find` command:

```
23 find . -type f -name *.jpg
```

The first argument is the directory to start the search, then we define the type `f` being files, and then specify the name. Find will recursively search through directories and sub-directories to find all files that match that name.

Very useful commands – grep

Programming Level-up

Jay Morgan

Linux

What is Linux

The command line

Shell Scripting

Writing bash scripts

High Performance Cluster

Getting started

Submitting jobs

A guided walk through

How about if we want to find files that have a certain contents? For that we can use `grep`. `Grep` will read a file and print (by default) the lines that contains your pattern. i.e.:

24

```
grep 'Linux' lecture.org
```

This will print the lines that contain the word `Linux` in `lecture.org`. If we just want the matched value, we use the `-o` flag.

25

```
grep -o '[0-9]' lecture.org
```

This prints all occurrences of numbers in `lecture.org`

Very useful commands – less/head/tail

Programming Level-up

Jay Morgan

Linux

What is Linux

The command line

Shell Scripting

Writing bash scripts

High

Performance Cluster

Getting started

Submitting jobs

A guided walk through

If a file is very long, we may not want to read the file using `cat`, as it will have to print the entire file. Instead we could use `less`, which will allow us to navigate through the file, using arrow keys to move and `q` to quit.

26

```
less filename
```

If we just want to view the first few lines, or the last few lines of a file we can use `head`/`tail`, respectively:

27

```
head filename
```

28

```
tail -n 20 filename # last 20 lines
```

29

```
tail -F filename # constantly read the file
```

Very useful commands – wc

Programming Level-up

Jay Morgan

Linux

What is Linux

The command line

Shell Scripting

Writing bash scripts

High Performance Cluster

Getting started

Submitting jobs

A guided walk through

Often times we just want to count the number of something. For example, if we want to count the number of files/folders in the directory we can do:

```
30 ls -l | wc -l
```

We're first printing all files and folders in a list format (one per line), then passing (piping) the result to `wc`, which with the `-l` line flag, is counting the number of lines. Therefore we get a count of the number of files and folders. Here is another example where we're counting how many times the word `bash` appears in these lecture notes:

```
31 grep -o 'bash' lecture.org | wc -l
```

Very useful commands – piping

Programming Level-up

Jay Morgan

Linux

What is Linux

The command
line

Shell Scripting

Writing bash
scripts

High Performance Cluster

Getting started

Submitting jobs

A guided walk
through

The purpose of piping is to pass data around between commands. We have just seen how we can pass the output of, say, the `ls` command to the input of `wc`. This allows use to construct very sophisticated pipelines to do some quite complex things from the combination of very simple commands.

32

```
find . -name '*.txt' -type f -print0 | xargs -0 grep "something"
```

Very useful basic commands

Programming Level-up

Jay Morgan

Linux

What is Linux

The command
line

Shell Scripting

Writing bash
scripts

High Performance Cluster

Getting started

Submitting jobs

A guided walk
through

In summary we have seen the following commands:

- `ls` - List a directory
- `cd` - Change/move to a directory
- `mkdir` - Make a new directory
- `cat` - Concatenate files
- `cp` - Copy a file/directory
- `mv` - Move files/folders
- `rm` - Remove files and folders
- `pwd` - Display the current absolute path
- `find` - Find files
- `grep` - Find occurrences of a pattern in a file
- `less/head/tail` - Read a file
- `wc` - Count

Very first bash script

Programming Level-up

Jay Morgan

Linux

What is Linux
The command
line

Shell Scripting

Writing bash
scripts

High

Performance
Cluster

Getting started
Submitting jobs
A guided walk
through

Let's start with the classic 'Hello, World' example. We'll create a new file called 'hello.sh' and enter the following:

```
33  #!/bin/bash
34
35  echo "Hello, World!"
```

First thing to notice is that the first line contains what we call a 'shebang' or 'hashbang'. It tells Linux which shell interpreter will be used to run the script, in this case: `/bin/bash`

The next (non-empty) line in the file is `echo 'Hello, World'`. This is exactly the same as the other commands we've just seen.

Very first bash script

Programming Level-up

Jay Morgan

Linux

What is Linux
The command
line

Shell Scripting

Writing bash
scripts

High Performance Cluster

Getting started
Submitting jobs
A guided walk
through

Now that we've created and saved our bash script, we will want to run it. We have two alternative methods to run this script:

```
36 bash hello.sh # run the script via bash
```

The second, requires that we have executable privileges for the script:

```
37 chmod +x hello.sh # add executable 'x' privileges  
38 ./hello.sh # execute it
```

Variables

Programming Level-up

Jay Morgan

Linux

What is Linux
The command line

Shell Scripting

Writing bash scripts

High Performance Cluster

Getting started
Submitting jobs
A guided walk through

The variables we create in our bash scripts are very much the same as the environment variables we've seen before. Take for example:

```
39 #!/bin/bash  
40 AGE="35"  
41 PERSON_NAME="Jane"  
42 echo "$PERSON_NAME is $AGE years old"
```

We create a variable `AGE` with the `=` assignment operator. **Note** we don't put spaces either side of the equals sign in bash. To refer to the variable, we use `$AGE`, using the `$` dollar sign.

Interpolation in bash strings

Programming Level-up

Jay Morgan

Linux

What is Linux
The command line

Shell Scripting

Writing bash scripts

High Performance Cluster

Getting started
Submitting jobs
A guided walk through

You would have noticed in the previous example that we included the variable directly into the string we're echoing out. This is something similar to what we've seen with f-strings in Python.

When we use double quotes: ". . ." in bash, the variable will be integrated into the resulting string. We can even call bash functions from directly inside the string:

43

```
echo "I am logging in as: $(who)"
```

Bash strings – the sharp edges

Programming Level-up

Jay Morgan

Linux

What is Linux
The command line

Shell Scripting

Writing bash scripts

High

Performance Cluster

Getting started
Submitting jobs
A guided walk through

You might be tempted to use a variable when generating a path:

```
44 TRAIN_PROCESS="training"  
45 TEST_PROCESS="testing"  
46  
47 touch "./data/${TRAIN_PROCESS}_error.txt"  
48 touch "./data/${TEST_PROCESS}_error.txt"
```

But this will create an error as underscores can be part of the variable name, so bash will be looking for a variable named: `${TRAIN_PROCESS}_error` which has never been created. To get around this, we can wrap our variable in curly braces:

```
49 touch "./data/${TRAIN_PROCESS}_error.txt"
```

Stopping interpolation in bash strings

Programming Level-up

Jay Morgan

Linux

What is Linux
The command
line

Shell Scripting

Writing bash
scripts

High Performance Cluster

Getting started
Submitting jobs
A guided walk
through

We can also use single quotes for strings in bash. When we use these strings, the string itself is not interpreted, and thus it will ignore any variables or bash commands:

50

```
echo 'I am logging in as: $(who)'
```

Input/Output

Programming Level-up

Jay Morgan

Linux

What is Linux
The command
line

Shell Scripting

Writing bash
scripts

High Performance Cluster

Getting started
Submitting jobs
A guided walk
through

If we want to read the input from keyboard into a variable, we use the read command:

```
51  #!/bin/bash
52
53  echo "Enter your name:"
54  read NAME
55
56  echo "Hello, $NAME"
```

read in this context will read in the input and create the variable with that value. As we've already seen, we can then output this value to the console using the echo command.

Booleans

Programming Level-up

Jay Morgan

Linux

What is Linux
The command line

Shell Scripting

Writing bash scripts

High

Performance Cluster

Getting started
Submitting jobs
A guided walk through

Technically, bash does not have built in data types for true and false, but Linux has the commands true and false which we could use in place. The implementation of how these commands work is not important.

```
57 FILE_EXISTS=true
58
59 if [ "$FILE_EXISTS" = true ]; then
60     echo "The file exists!"
61 fi
```


Conditionals

Programming Level-up

Jay Morgan

Linux

What is Linux
The command line

Shell Scripting

Writing bash scripts

High Performance Cluster

Getting started
Submitting jobs
A guided walk through

When we're creating if expressions, we use the following syntax:

```
62  if <<conditional>>; then
63      # do something
64  else
65      # do something else
66  fi
```

We can also use elif

```
67  if <<conditional>>; then
68      # do something
69  elif <<conditional>>; then
70      # do something else
71  else
72      # something else entirely
73  fi
```

Conditionals

Programming Level-up

Jay Morgan

Linux

What is Linux

The command line

Shell Scripting

Writing bash scripts

High

Performance Cluster

Getting started

Submitting jobs

A guided walk through

Writing condition expressions can be a little more cumbersome than in Python. These can be many pain points for new bash programmers, take for example:

```
74 FILE_EXISTS=false
75
76 if [ $FILE_EXISTS ]; then
77     echo "The file exists!"
78 fi
```

This is because we have used the [...] single bracket syntax for the test. But there are others:

- No brackets: we could omit the brackets in which case it would run the false command not print the statement.
- Single parenthesis (...) creates a sub-shell.
- Double parenthesis ((...)) for arithmetic operation
- Single square bracket [...] calls test
- Double square bracket [[...]]

Conditionals

Programming Level-up

Jay Morgan

Linux

What is Linux

The command line

Shell Scripting

Writing bash scripts

High Performance Cluster

Getting started

Submitting jobs

A guided walk through

What if we write:

```
79 VAR_1="Mr Foo Bar"
80 VAR_2="Mr Foo Bar"
81 if [ $VAR_1 = $VAR_2 ]; then
82     echo "They are the same"
83 fi
```

We would get an error because test expands the arguments into:

Mr Foo Bar = Mr Foo Bar

With the spaces included. To prevent this from happening, we have to wrap the variables in quotation marks.

```
2 VAR_1="Mr Foo Bar"
3 VAR_2="Mr Foo Bar"
4 if [ "$VAR_1" = "$VAR_2" ]; then
5     echo "They are the same"
6 fi
```

Conditionals

Programming Level-up

Jay Morgan

Linux

What is Linux
The command line

Shell Scripting

Writing bash scripts

High Performance Cluster

Getting started
Submitting jobs
A guided walk through

If we use `[[` in if statement, then we can do more sophisticated things like pattern matching:

```
7 FILENAME="testing.png"
8 if [[ "$FILENAME" = *.png ]]; then
9     echo "Its a png file"
10 fi
```

Loops

Programming Level-up

Jay Morgan

Linux

What is Linux
The command line

Shell Scripting

Writing bash scripts

High Performance Cluster

Getting started
Submitting jobs
A guided walk through

Like in Python, we can iterate in bash

```
11 for i in {1..10}; do
12     echo $i
13 done
```

This iterates with *i* starting at 1 upto 10 (inclusive). Or we could do:

```
14 for (( i=1; i <= 10; i++ )); do
15     echo $i
16 done
```

Loops

Programming Level-up

Jay Morgan

Linux

What is Linux
The command
line

Shell Scripting

Writing bash
scripts

High Performance Cluster

Getting started
Submitting jobs
A guided walk
through

We can also iterate over a list of files/folders in a directory:

```
17 for FILE in ./images/*; do
18     echo $FILE
19 done
```

Loops

Programming Level-up

Jay Morgan

Linux

What is Linux
The command line

Shell Scripting

Writing bash scripts

High Performance Cluster

Getting started
Submitting jobs
A guided walk through

Using the `while` form, we can continue looping until our conditional is false. For example, we could loop testing our internet connection, until its been established:

```
20 while ! ping -c 1 google.com; do
21     echo "No internet yet"
22     sleep 1
23 done
24
25 echo "Internet is available!"
```

Functions

Programming Level-up

Jay Morgan

Linux

What is Linux
The command
line

Shell Scripting

Writing bash
scripts

High

Performance
Cluster

Getting started
Submitting jobs
A guided walk
through

To create a function, we use the following syntax:

```
26 function_name() {  
27     # do something  
28 }
```

And to call the function, you just need to use the function name:

```
29 function_name # this called function name
```


Functions

Programming
Level-up

Jay Morgan

Linux

What is Linux

The command
line

Shell Scripting

Writing bash
scripts

High

Performance
Cluster

Getting started

Submitting jobs

A guided walk
through

Here is another example:

```
30 say_hello() {  
31     echo "Hello, $1"  
32 }  
33  
34 say_hello "Jane"
```

Notice that we didn't need to include any argument list. We just used \$1 for the first argument passed to the function.

```
35 say_hello() {  
36     echo "$1, $2"  
37 }  
38  
39 say_hello "Hi" "Jane"
```

Functions

Programming
Level-up

Jay Morgan

Linux

What is Linux
The command
line

Shell Scripting

Writing bash
scripts

High
Performance
Cluster

Getting started
Submitting jobs
A guided walk
through

Returning values is 'interesting' as, coming from other languages, you think could do something like this:

```
40 say_hello() {  
41     return "hello"  
42 }  
43 RESULT="$(say_hello)"  
44 echo $RESULT
```

This didn't work like we expected, the value wasn't returned and assigned to RESULT. So how do we return a value?

```
45 say_hello() {  
46     echo "Hello"  
47 }  
48 RESULT="$(say_hello)"  
49 echo "This is before the printing of result"  
50 echo $RESULT
```

What is the Cluster?

Programming Level-up

Jay Morgan

Linux

What is Linux
The command
line

Shell Scripting

Writing bash
scripts

High Performance Cluster

Getting started

Submitting jobs
A guided walk
through

How to login

Programming Level-up

Jay Morgan

Linux

What is Linux

The command
line

Shell Scripting

Writing bash
scripts

High Performance Cluster

Getting started

Submitting jobs

A guided walk
through

51

```
ssh <<username>>@saphir2.lis-lab.fr
```

Then:

52

```
ssh <<username>>@sms-ext.lis-lab.fr
```

How to login

Programming Level-up

Jay Morgan

Linux

What is Linux

The command
line

Shell Scripting

Writing bash
scripts

High

Performance
Cluster

Getting started

Submitting jobs

A guided walk
through

Typing both these commands can become tiresome very quickly. But we can make it a lot easier by updating our `~/.ssh/config` file to include something like:

```
Host saphir2
  HostName saphir2.lis-lab.fr
  User <<username>>

Host cluster
  HostName sms-ext.lis-lab.fr
  User <<username>>
  ProxyCommand ssh saphir2 -W %h:%p
```

Then to login to the cluster, we just need to type:

9

```
ssh cluster
```

And we should be prompted for our password.

How to login

Programming Level-up

Jay Morgan

Linux

What is Linux
The command line

Shell Scripting

Writing bash scripts

High Performance Cluster

Getting started

Submitting jobs
A guided walk through

If you trust the machine your on, you can remove password authentication and move to key-based authentication:

```
10 ssh-copy-id saphir2
11 ssh-copy-id cluster
```

When we next login to the server, we shouldn't be prompted for a password.

How to copy files to and from the cluster

Programming Level-up

Jay Morgan

Linux

What is Linux
The command line

Shell Scripting

Writing bash scripts

High Performance Cluster

Getting started

Submitting jobs
A guided walk through

We have a number of options for transferring files to and from the cluster. Firstly, let's look at the command `scp`. It takes two arguments, the first argument is the file you want to send, the second argument is the destination of the sent file.

```
12 scp <<origin>> <<destination>>
```

Similar to commands like `cp`, `scp` by default only works for files, not folders. To send folders/directories, we use the `-r` flag just like `cp`.

```
13 scp -r <<origin_folder>> <<destination_folder>>
```

Copying files – rsync

Programming Level-up

Jay Morgan

Linux

What is Linux
The command line

Shell Scripting

Writing bash scripts

High Performance Cluster

Getting started

Submitting jobs
A guided walk through

One of the downsides about `scp` is that it will copy every file you give it. Even if the file at the destination is exactly the same. What if we only want to copy files that need to be copied, i.e. that are outdated, thus saving time? For that, we can use `rsync`. `Rsync` will copy files from one source to a destination only if the destination needs to be updated. This can save a lot of time by skipping files that already exist at the destination:

14

```
rsync <<source>> <<destination>>
```


The Login- and Compute Nodes

Programming Level-up

Jay Morgan

Linux

What is Linux

The command
line

Shell Scripting

Writing bash
scripts

High Performance Cluster

Getting started

Submitting jobs

A guided walk
through

When you login to the cluster, you are logging into the *login* node.

Note that no computation should be run on this node. If you want to run scripts, you will have to submit a job to the compute nodes.

On the login node there is a system installed called 'SLURM'. SLURM is a job scheduler program that receives your requests for executing scripts, it will queue them and assign them to available compute nodes.

We will take a look at how to request and manage jobs using the various commands that SLURM provides.

How to launch a job – srun

Programming Level-up

Jay Morgan

Linux

What is Linux
The command line

Shell Scripting

Writing bash scripts

High Performance Cluster

Getting started
Submitting jobs
A guided walk through

The first command we will look at it is `srun`. This command will run request a job for execution in 'real-time'. By real-time, we mean that the shell will wait until the job has been submitted.

15

```
srun <<compute node options>> <<command to run>>
```

Let's take a look at an example where we want to run an interactive bash shell on the compute shell (similar to `ssh`'ing into the compute node).

16

```
srun --time=00:10:00 --pty bash -l
```

This will request a job on any available compute node for 10 minutes. When a node becomes available, `bash` will execute, dropping you into the shell. You will notice that the shell prompt has changed from `sm$` to the name of the node.

How to launch a job – salloc

Programming Level-up

Jay Morgan

Linux

What is Linux
The command
line

Shell Scripting

Writing bash
scripts

High Performance Cluster

Getting started
Submitting jobs
A guided walk
through

There is another method we can use to create an interactive job. We could use the `salloc` command to allocate resources for a task. After the resources have been allocated and our job is ready, we can `ssh` into the node with the allocated resources.

```
17  salloc --time=10:00:00 &  
18  ssh <name>
```

How to launch a job – options

Programming Level-up

Jay Morgan

Linux

What is Linux
The command
line

Shell Scripting

Writing bash
scripts

High Performance Cluster

Getting started
Submitting jobs

A guided walk
through

In the previous command, we used the `--time` option to specify how long the job will run for. But there are other options we can use to be more specific about the jobs we want to run.

`--cpus-per-task` can be used to request more than one CPU to be allocated. This is especially helpful when we have a multithreaded process we want to run.

`--mem` specifies how much memory should be allocated to the job. For example: `--mem=16G` tells SLURM to allocate 16 GB of memory.

How to launch a job – GPU allocation

Programming Level-up

Jay Morgan

Linux

What is Linux
The command line

Shell Scripting

Writing bash scripts

High Performance Cluster

Getting started
Submitting jobs
A guided walk through

If we need to use a GPU, we need to use a few options. Firstly, we can specify that our job is on a compute node with GPU. There will usually be a group of nodes in a 'GPU' group or partition, and thus we can specify to use one of these partitions:

19

```
srun --time=00:10:00 --partition=gpu --pty bash -l
```

But you will notice that you still do not have access to a GPU. You're running on the GPU node, but you haven't actually requested a GPU be allocated to your job. For that you will use `--gres`:

20

```
srun --time=00:10:00 --partition=gpu --gres=gpu:1 --pty bash -l
```

Here we are requesting one GPU, but if we use `--gres:gpu:2` we are requesting 2 GPUs etc.

How to launch a job – GPU allocation

Programming Level-up

Jay Morgan

Linux

What is Linux

The command line

Shell Scripting

Writing bash scripts

High

Performance Cluster

Getting started

Submitting jobs

A guided walk through

There are many different types of GPUs available, some older than others. If you wanted to allocate a job with a specific type of GPU you can use the `--constraint` flag:

```
21  srun --time=00:10:00 \  
22      --partition=gpu \  
23      --gres=gpu:1 \  
24      --constraint='cuda61' \  
25      --pty bash -l
```

This command requests that our job run on the GPU partition, with 1 GPU allocated that has the capability of running CUDA compute 61.

Or we can specify the type of GPU in the gres option:

```
26  srun --time=00:10:00 \  
27      --partition=gpu \  
28      --gres=gpu:2080:1 \  
29      --pty bash -l
```

Learning more about nodes

Programming Level-up

Jay Morgan

Linux

What is Linux
The command
line

Shell Scripting

Writing bash
scripts

High Performance Cluster

Getting started
Submitting jobs
A guided walk
through

To understand what each compute node has we can use the `scontrol` command.

30

```
scontrol show nodes
```

Will list out all nodes and all capabilities of each node. Or just one node:

31

```
scontrol show node lisnode2
```

How to launch a job – sbatch

Programming Level-up

Jay Morgan

Linux

What is Linux

The command line

Shell Scripting

Writing bash scripts

High Performance Cluster

Getting started

Submitting jobs

A guided walk through

It can be quite inconvenient to launch an interactive job to run some compute, and wait for the job to be allocated. If, instead, you have a long running experiment that you want to run without any intervention from you, you can use `sbatch`.

`Sbatch` will require us to write a small bash script that specifies how to run a job and what to do once its allocated.

```
32  #!/bin/bash
33
34  #SBATCH --time=00:01:00
35  #SBATCH --job-name=my_new_job
36  #SBATCH --output=my_new_job.out
37  #SBATCH --error=my_new_job.err
38
39  echo $HOSTNAME
```

And run it:

```
40  sbatch my_job.sh
```



How to launch a job – sbatch

Programming Level-up

Jay Morgan

Linux

What is Linux

The command line

Shell Scripting

Writing bash scripts

High Performance Cluster

Getting started

Submitting jobs

A guided walk through

Notice that instead of supplying options to `sbatch`, we can instead record them directly into the script using the `#SBATCH`. SLURM will examine this file, looking for lines starting with this comment, and infer that the rest of the line contains the options.

There are a few other options we've included that are very useful when running non-interactive jobs. Firstly, we've given the job a name (`my_new_job`). This is so we can differentiate between many jobs that we might run at the same time. To list out the jobs we currently have running we use `squeue`.

41

```
squeue
```

By default, `squeue` will list all of the active jobs, even other people's. To specify only your jobs use the `--user` option:

42

```
squeue --user=jay.morgan
```

How to launch a job – sbatch

Programming Level-up

Jay Morgan

Linux

What is Linux

The command
line

Shell Scripting

Writing bash
scripts

High Performance Cluster

Getting started

Submitting jobs

A guided walk
through

The other two options, `--output` and `--error` specify where the printed output and printed errors will be stored. Since the job is being run on a different node, by a non-interactive process, if you didn't include these lines, you wouldn't be able to see what was being printed by `echo` or by any other process such as `print` in Python.

Job Management – queue

Programming Level-up

Jay Morgan

Linux

What is Linux
The command
line

Shell Scripting

Writing bash
scripts

High Performance Cluster

Getting started

Submitting jobs

A guided walk
through

When we list the jobs using `queue` it will give us multiple columns of information, such as:

- **JOBID** – the referable id of the job.
- **PARTITION** – the partition on which the job has been requested for.
- **NAME** – the name of the job.
- **USER** – the user who submitted the job.
- **ST** – the status, is the job currently running, waiting, or exiting?
- **TIME** – how long the job has been running for.
- **NODES** – how many nodes have been allocated to the job.

Job Management – scancel

Programming Level-up

Jay Morgan

Linux

What is Linux

The command
line

Shell Scripting

Writing bash
scripts

High

Performance Cluster

Getting started

Submitting jobs

A guided walk
through

Let's say that we've submitted a job, but we've noticed that there was an error in the code, and want to stop the job. For that, we use `scancel` and specify the id of the job we wish to cancel:

43

```
scancel 158590
```

After running this command, we should see, using `squeue`, that either the job is finishing, or that its disappeared from our list (meaning that its completely stopped).

Job Management – sacct

Programming Level-up

Jay Morgan

Linux

What is Linux
The command line

Shell Scripting

Writing bash scripts

High Performance Cluster

Getting started
Submitting jobs
A guided walk through

If our job has finished, or exited and is no longer in queue, we can use `sacct` to get a history of the jobs.

`sacct` will list all of your jobs within some default window of time. If we want to change this window we can use the `--starttime` and `--endtime` options.

Valid time formats are:

- `HH:MM[:SS][AM|PM]`
- `MMDD[YY][-HH:MM[:SS]]`
- `MM.DD[.YY][-HH:MM[:SS]]`
- `MM/DD[/YY][-HH:MM[:SS]]`
- `YYYY-MM-DD[THH:MM[:SS]]`
- `today, midnight, noon, fika (3 PM), teatime (4 PM)`
- `now[{{+|-}count[seconds(default)|minutes|hours|days|weeks]]`

Job Task Arrays – motivation

Programming
Level-up

Jay Morgan

Linux

What is Linux
The command
line

Shell Scripting

Writing bash
scripts

High
Performance
Cluster

Getting started

Submitting jobs

A guided walk
through

Task arrays allow you to submit many jobs of the same type. Why might this be useful? Suppose you have a list of files that take a long time to process:

- file_0.txt
- file_1.txt
- file_2.txt

Or you have some computation script, such as deep learning training script, that takes uses a hyperparameter which can be tuned to achieve different performance results:

44

```
python train.py --learning-rate 0.001
```

Instead of a creating a sbatch script for each value of hyperparameter, or sequentially enumerating the values, you can use a job task array to spawn multiple jobs with slightly different values.

Job Task Arrays – how to

Programming Level-up

Jay Morgan

Linux

What is Linux
The command
line

Shell Scripting

45 Writing bash
scripts

High Performance Cluster

49 Getting started

Submitting jobs

50 A guided walk
through

First, we will look at how to actually submit an array of tasks. To create an task array, you will need to add the `--array` options to your sbatch script:

```
45 #!/bin/bash  
46  
47 #SBATCH --job-name=my_task_array  
48 #SBATCH --array=1-5  
49  
50 ...
```

Here we are creating an array of tasks numbered from 1-5. When you submit this script, you will see five tasks submitted to the queue.

Job Task Arrays – how to

Programming Level-up

Jay Morgan

Linux

What is Linux
The command line

Shell Scripting

Writing bash scripts

High Performance Cluster

Getting started
Submitting jobs
A guided walk through

Now that we know how to create an array of tasks, we will want to do something useful with it. When you create an array, each individual task will have a unique variable called `SLURM_ARRAY_TASK_ID`. So for example, if we launch an array of 5 tasks, the first task will have the value 1. Why is this useful? Well, we can use this variable to alter the program slightly. Take for example our list of files we need to process:

```
51  #!/bin/bash
52  #SBATCH --job-name=my_task_array
53  #SBATCH --array=0-4
54  #SBATCH --time=00:10:00
55
56  FILENAME="file_${SLURM_ARRAY_TASK_ID}.txt"
57  python process.py $FILENAME
```

This will create a new bash variable called `FILENAME` by concatenating `file_` the current task's (i.e. 0, for the first task, 1 for the second task, etc) and `.txt`.

Job Task Arrays – how to

Programming Level-up

Jay Morgan

Linux

What is Linux

The command line

Shell Scripting

Writing bash scripts

High Performance Cluster

Getting started

Submitting jobs

A guided walk through

If we run the previous example, we will see that we have five jobs named exactly the same thing `my_task_array`. This is okay, but we can be a little bit more clear as to which task is running, i.e. which task is processing which file?

We can use some special variables in our bash script to make this more clear. These are `%A` that is the main job id, and `%a` that is the task array id.

```
58  #!/bin/bash
59
60  #SBATCH --job-name=my_task_array.%A_%a
61  #SBATCH --output=my_task_array.%A_%a.out
62  ...
```

Now, every task in our array will have a slightly different name because of the `%a` and therefore we will be able to determine which job is processing which file.

Job Task Arrays – how to

Programming Level-up

Jay Morgan

Linux

What is Linux

The command line

Shell Scripting

Writing bash scripts

High Performance Cluster

Getting started

Submitting jobs

A guided walk through

Let's move on to the second example, where we have a Deep Learning training program and we want to try different parameters. In this case, we can again use a task array.

```
63 #!/bin/bash
64
65 #SBATCH --array=1-10
```

We could either pass the `SLURM_ARRAY_TASK_ID` as a command line argument to the script:

```
66 python training.py --learning-rate $SLURM_ARRAY_TASK_ID
```

But in this case, we could have to properly calculate the correct learning rate from the `SLURM_ARRAY_TASK_ID` value (remember that in my sbatch script I set `--array=1-5`). But bash only performs integer arithmetic, therefore we will need to calculate the correct learning rate in something else

Job Task Arrays – how to

Programming Level-up

Jay Morgan

Linux

What is Linux

The command line

Shell Scripting

Writing bash scripts

High

Performance Cluster

Getting started

Submitting jobs

A guided walk through

Instead of passing the learning rate via a command line argument. We can get the value directly from our python script and calculate the value.

```
67 import os
68
69 task_id = int(os.environ["SLURM_ARRAY_TASK_ID"])
70 learning_rate = task_id / 100
```

Here we are using the builtin `os` module in Python, getting the environment variable from the dictionary `environ` and parsing the value as an integer. Then we can calculate the appropriate learning rate using this value. So for example, if `SLURM_ARRAY_TASK_ID` is set to 1. Our learning rate would be 0.01 for this task.

Job Task Arrays – how to

Programming Level-up

Jay Morgan

Linux

What is Linux

The command line

Shell Scripting

Writing bash scripts

High

Performance Cluster

Getting started

Submitting jobs

A guided walk through

If you're creating a job task array, you may want to create hundreds of jobs. And of course, you don't want to use up the entire cluster leaving no resources for anybody else! Therefore, you will only want a maximum number of tasks to run at any one time.

```
71 #!/bin/bash  
72  
73 #SBATCH --array=1-100%5
```

This will create a job task array of 100 jobs numbered from 1 to 100. But we have added an additional argument %5 which means that only 5 jobs can run at any one time for this task array. If you have five tasks running, the other 95 tasks will wait.

If, at any point, you want to change how many jobs can run simultaneously, you can update this 'throttle' value using `scontrol`:

```
74 scontrol update ArrayTaskThrottle=<count> JobId=<jobID>
```

Job Task Arrays – how to

Programming Level-up

Jay Morgan

Linux

What is Linux

The command
line

Shell Scripting

Writing bash
scripts

High

Performance Cluster

Getting started

Submitting jobs

A guided walk
through

So if we've already launched a job task array with the job id of 50602 that has a throttle value of 5 (only 5 tasks will run at once), we can change it to 10 using:

75

```
scontrol update ArrayTaskThrottle=10 JobId=50602
```

A guided walk through – environment

Programming Level-up

Jay Morgan

Linux

What is Linux
The command line

Shell Scripting

Writing bash scripts

High Performance Cluster

Getting started
Submitting jobs

A guided walk through

In this section we're going to give an example walk through of working with the HPC cluster. In this example, we're going to write our scripts locally, including the slurm submission script, and when they're ready, we'll send them to the cluster to perform the actual computation.

Let's imagine we're starting a new project, and are programming our scripts in Python. Now is a good time to create a new conda environment to install our packages we're going to use for our research. We'll create this environment with (replacing `<env-name>` with whatever we want to call this environment):

```
76 conda create --name <env-name>
```

and then activate it:

```
77 conda activate <env-name>  
78 conda install python=3.9
```

Writing our scripts

Programming Level-up

Jay Morgan

Linux

What is Linux
The command line

Shell Scripting

Writing bash scripts

High Performance Cluster

Getting started
Submitting jobs

A guided walk through

Let us also imagine we've just wrote the following script to create a lorenz attractor: `lorenz.py`

The specific implementation of this script is not particularly important for this walk through. Just know that we're importing a few packages such as `numpy` and `matplotlib`. Then, we're performing some computation, and saving the results to analyse later. As this script uses external libraries, we need to install them:

79

```
conda install numpy matplotlib
```

Writing our job submission script

Programming Level-up

Jay Morgan

Linux

What is Linux
The command line

Shell Scripting

Writing bash scripts

High Performance Cluster

Getting started
Submitting jobs

A guided walk through

Since we want our calculations to be performed on the cluster, we will need to also write a job submission script (let's call this `submit-job.sh`) in bash to pass to SLURM.

```
80  #!/bin/bash
81
82  #SBATCH --job-name=lorenz_attractor
83  #SBATCH --output=lorenz_attractor.log
84  #SBATCH --error=lorenz_attractor.log
85  #SBATCH --time=00:10:00
86
87  python lorenz.py
```


Replicating our environment on the cluster

Programming Level-up

Jay Morgan

Linux

What is Linux
The command
line

Shell Scripting

Writing bash
scripts

High Performance Cluster

Getting started
Submitting jobs

A guided walk
through

As we've installed external packages in our local development environment, we will want to ensure that when we run the calculations on the cluster, it will be using the same versions of packages. Conda makes this a lot easier. First, we export our environment to a recipe file:

88

```
conda env export --no-builds > environment.yml
```

Sending our scripts to the cluster

Programming Level-up

Jay Morgan

Linux

What is Linux
The command line

Shell Scripting

Writing bash scripts

High Performance Cluster

Getting started
Submitting jobs

A guided walk through

All of our scripts are ready! We can now transfer them from our personal computer, to the cluster. The files we need to transfer are:

- `lorenz.py`
- `environment.yml`
- `submit-job.sh`

While we can send a folder (and the containing files), let's send them one at a time:

```
89 scp lorenz.py <hostname>:<destination-path>
90 scp environment.yml <hostname>:<destination-path>
91 scp submit-job.sh <hostname>:<destination-path>
```

where `<hostname>` is the hostname/IP address that you've used to connect to the login node on the cluster before.

`<destination-path>` is the path to where you want to save the files.

Logging into the cluster

Programming Level-up

Jay Morgan

Linux

What is Linux

The command
line

Shell Scripting

Writing bash
scripts

High Performance Cluster

Getting started
Submitting jobs

A guided walk
through

Now that our files are on the cluster, we can login:

```
92 ssh <username>@<hostname>
```

At which point, we've logged into the login node, and then we need to change directory to where we saved the files:

```
93 cd <destination-path>
```

Re-creating our development environment

Programming Level-up

Jay Morgan

Linux

What is Linux
The command
line

Shell Scripting

Writing bash
scripts

High Performance Cluster

Getting started
Submitting jobs

**A guided walk
through**

Now that we're in the same folder as our scripts, we're almost ready to submit our job. First, we need to recreate our development environment from our `environment.yml` file.

```
94 conda env create -f environment.yml
```

And activate our newly created environment:

```
95 conda activate <env-name>
```

Submitting our job

Programming Level-up

Jay Morgan

Linux

What is Linux

The command line

Shell Scripting

Writing bash scripts

High

Performance Cluster

Getting started

Submitting jobs

A guided walk through

Now we can submit our job:

96

```
sbatch submit-job.sh
```

We can check the progress of our job with `squeue`, or its already completed, look at the job history with `sacct`.

Downloading the results

Programming Level-up

Jay Morgan

Linux

What is Linux
The command line

Shell Scripting

Writing bash scripts

High Performance Cluster

Getting started
Submitting jobs

A guided walk through

If our job runs successfully, a `data.pkl` file will be created. Back on our local computers, we will need to run the following to download it:

97

```
scp <hostname>:<destination-path>/data.pkl ./
```

This will download the file into the current directory.

Analysing the results

Programming Level-up

Jay Morgan

Linux

What is Linux
The command
line

Shell Scripting

Writing bash
scripts

High Performance Cluster

Getting started
Submitting jobs

**A guided walk
through**

With the `data.pk1` file downloaded, we can visualise the results using `plot_lorenz.py`: <https://pageperso.lis-lab.fr/jay.morgan/resources/2021-programming-level-up/lectures/week-5/plot-lorenz.py>

If everything has been run correctly, you should see a plot of the lorenz attractor.

Useful features – X11 Forwarding I

Programming Level-up

Jay Morgan

Linux

What is Linux
The command
line

Shell Scripting

Writing bash
scripts

High Performance Cluster

Getting started
Submitting jobs

A guided walk
through

If we're performing analysis interactively using the cluster, we'll often want to visualise the results, using matplotlib for example. To see our plots display like they would on our local machine when we call `plt.plot` or `plt.show()`, we will need to ensure that we're using something called X11 Forwarding. To enable X11 Forwarding, we use the `-X` option when `ssh`'ing into the cluster and compute nodes (i.e. it will need to be enabled on every 'hop' so to speak).

98

```
ssh -X <<remote-host>>
```

If want to enable it by default, we can enable it in our `ssh` config file:

Useful features – X11 Forwarding II

Programming Level-up

Jay Morgan

99

```
Host saphir2
  HostName saphir2.lis-lab.fr
  ForwardX11 yes
  User <<username>>
```

100

Linux

What is Linux

The command
line

101

102

103

Shell Scripting

Writing bash
scripts

104

105

High

Performance

Cluster

106

107

108

```
Host cluster
  HostName sms-ext.lis-lab.fr
  FowardX11 yes
  User <<username>>
  ProxyCommand ssh -X saphir2 -W %h:%p
```

Getting started

Submitting jobs

A guided walk
through

After setting up X11 Forwarding correctly, and when logged into the remote host, we should be able to echo a variable called \$DISPLAY.

109

```
echo $DISPLAY
```

Useful features – X11 Forwarding III

Programming Level-up

Jay Morgan

Linux

What is Linux

The command
line

Shell Scripting

Writing bash
scripts

High Performance Cluster

Getting started

Submitting jobs

**A guided walk
through**

If `$DISPLAY` has a value, we know that X11 Forwarding has been setup correctly, and we're ready to do some plotting!

Useful features – Jupyter Notebooks I

Programming Level-up

Jay Morgan

Linux

What is Linux
The command line

Shell Scripting

Writing bash scripts

High Performance Cluster

Getting started
Submitting jobs

A guided walk through

We've talked about how good jupyter notebooks are for performing analysis and exploration. But often times, we will need a lot of compute resources (more than our local computers can handle) to do this analysis. This is where using jupyter notebook on the supercomputers comes in handy. However, it is not as simple as starting the jupyter notebook server and opening up your web browser. First, we will need to setup a 'reverse ssh tunnel'.

In a nut-shell, a reverse ssh tunnel allows you to redirect data on a remote port to a local port.

Therefore, we can, using a reverse ssh tunnel, start a jupyter notebook on the supercomputer and access it using the web browser on our local computer!

To begin, we can create an interactive job on the cluster:

```
srun --time=01:00:00 --pty bash -l
```

110

Useful features – Jupyter Notebooks II

Programming
Level-up

Jay Morgan

Linux

What is Linux
The command
line

Shell Scripting

Writing bash
scripts

High
Performance
Cluster

Getting started
Submitting jobs

A guided walk
through

And start our jupyter notebook, specifying a port that will not be in use:

```
jupyter notebook --port 30333
```

With our notebook server now started on the 30333 port, we will want to create an ssh tunnel from our local computer, to the cluster's login node, and then a tunnel from the login node to the specific compute node where the job is running:

```
ssh -L 30333:localhost:30333 <<cluster-login-node>> ssh -L  
↪ 30333:localhost:30333 <<cluster-compute-node>>
```

If everything goes well, we should now be able to open up our web browser, navigate to localhost:30333 and see our jupyter notebooks.