Introduction to the iCSF (Incline) & CSF (Computational Shared Facility)

- WELCOME!
- Please take a seat near the front.
- Course materials can be found at

  http://ri.itservices.manchester.ac.uk/course/icsf-csf/

- There is no sign-sheet. Attendance is recorded using a feedback form (see above link)
- https://ri.itservices.manchester.ac.uk/hpc-help

Chris Grave, Research IT, IT Services

its-ri-team@manchester.ac.uk

Who we are - Research IT

- Central Group in IT Services offering a variety of researcher specific services
  - HPC/Compute
  - Applications support
  - Software developers for research projects
  - Storage for Research Data
  - Virtual Machine Service
  - Advice regarding external services Training courses
  - Many other topics of relevance to researchers
- Blog and newsletter
- Events: Research IT Club, Drop-in sessions
- More Info - https://research-it.manchester.ac.uk/services/

MANCHESTER 1824
The University of Manchester

Presentation Overview

- Introduction to the iCSF (incline) & CSF
- Research Virtual Desktop Service
  - Exercise 1
- Research Data Storage (RDS) & CIR Filesystems
  - Exercise 2
- Using the iCSF (incline)
  - Exercise 3
- Using the CSF
  - Exercise 4
- Other Research IT Infrastructure Provided Services

MANCHESTER 1824
The University of Manchester

Introduction to the

iCSF (Incline) & CSF

http://ri.itservices.manchester.ac.uk/course/icsf-csf/

MANCHESTER 1824
The University of Manchester

Why use the iCSF & CSF?

- Offload computational work from your PC
  - Do larger simulations/analyses that wouldn't fit on your PC
  - Work in an interactive manner using a GUI (iCSF)
- Variety of software applications
  - MATLAB, R, Stata, SAS, many more
- 1-core (serial) jobs or many-core (parallel) jobs = speedup
  - A core is a processing unit in a CPU/processor
  - Modern CPUs have several cores: your PC/laptop may have a single CPU with 4,6,8 cores; the iCSF (32 cores per node), CSF (up to 168 cores per node), a lot more in total.
- Tackle larger problems involving larger datasets
  - Lots of memory 1.5TB, 2TB, 4TB nodes

MANCHESTER 1824
The University of Manchester

iCSF

- Run interactively
- A backend node becomes your workstation
- Start application GUI

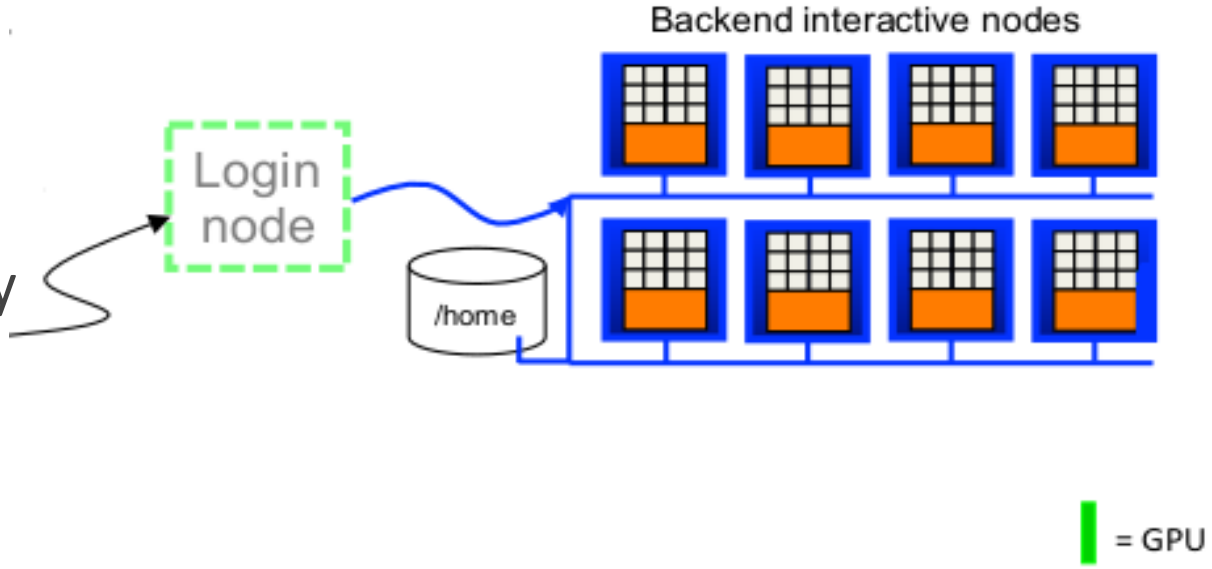

CSF

- No direct access to backend nodes
- Batch system runs jobs on backends without GUIs

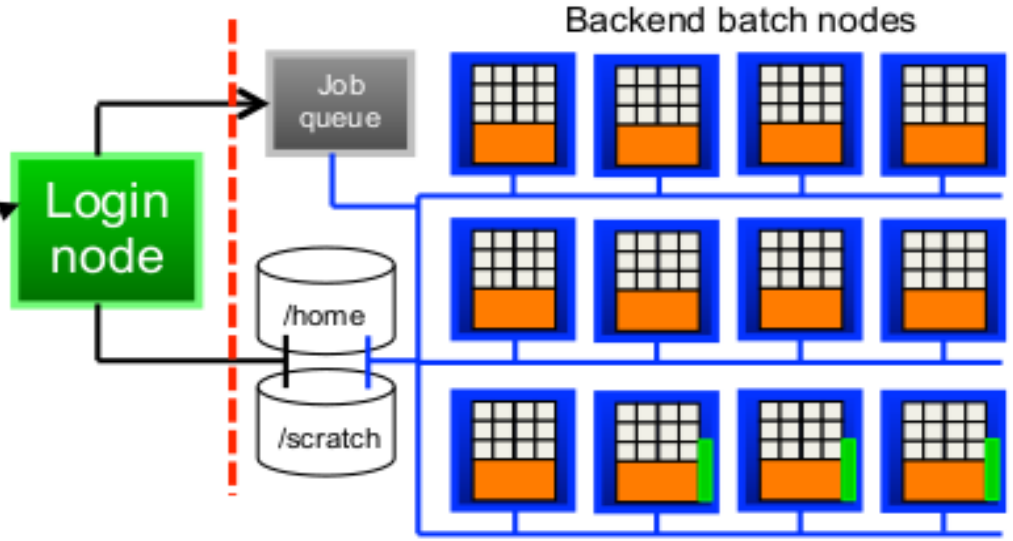

http://ri.itservices.manchester.ac.uk/course/icsf-csf/

iCSF vs CSF Differences

iCSF - run interactively

- After login a backend node becomes your ‘shared’ workstation
- Load and interact with applications via a GUI inc., MATLAB, Rstudio, Stata, SAS
- Easier to use then CSF
- Some command line use
- Instant access to resources

CSF - run in batch

- No direct access to backend nodes
- All jobs are submitted via a batch system jobs run on backends without GUIs
- Steeper learning curve
- Command line heavy
- Greater computational Resource
- Large Scratch Filesystem

http://ri.itservices.manchester.ac.uk/course/icsf-csf/

MANCHESTER 1824
The University of Manchester

iCSF vs CSF Similarities

- Both are a collection of servers (compute nodes) linked together to form two large systems
- Operating system on both is the same – Linux (Red Hat Based)
- They share the same home directory (similar to your P Drive - a private area for your files)
- Both platforms are built based on a contribution model, i.e. research groups cost in cost of hardware in research grant proposals.
- Limited Free at the Point of use offering (FATPOU) on both platforms.

MANCHESTER 1824
The University of Manchester

Is the iCSF suitable for my work?

Yes if

- You need to work interactively
- Have low to moderate computational requirements
- Are unfamiliar with the Command Line.
- Are developing or testing code

It also

- Complements other CIR platforms
- Provides a first step into computational work

Unsure

email its-ri-team@mancheter.ac.uk we're here to help.

Is the CSF suitable for my work?

Yes, if

- You have moderate to high computational needs
- You need to run many ‘Jobs’ e.g. Simulations/Analyses at once.
- You need GPUs

You will need to

- Be a little more familiar with a Linux Command Line.
- use a batch system
- Don’t worry, we will teach you some of the basics today
- We have 1 day training course specifically for the CSF
- Lots of docs, and friendly support team

(its-ri-team@manchester.ac.uk)

Networking & Security

- The iCSF/CSF has a private (campus only) IP address & firewall
  - Controls connections to and from the system
  - **On campus**
    - Connect from any PC/laptop with a wired connection or University wireless and Eduroam
    - No need for VPN but will need to authenticate using your 2FA device, e.g.

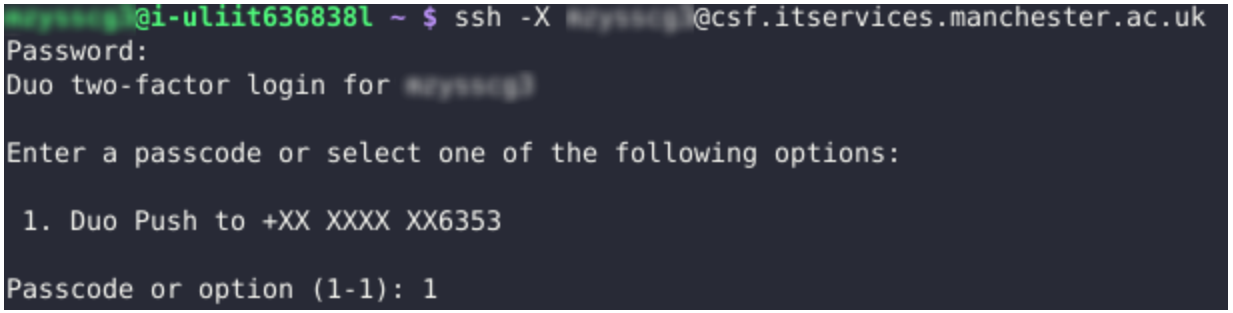

  - **Off campus:** http://ri.itservices.manchester.ac.uk/wfh/
    - University Global Protect VPN Gateway with 2FA
      - Must be running
      - Allows you to the CSF & iCSF
      - Can be some issues with some ISPs, see above WFH webpage
- Further documentation: http://ri.itservices.manchester.ac.uk/csf3/getting-started/connecting/

MANCHESTER 1824
The University of Manchester

How do I Connect to the iCSF & CSF?

- Via the **Remote Virtual Desktop Service** (RVDS) - focus of today's exercises.
- Directly via **SSH** using a Terminal or suitable Terminal Emulator such as e.g. MobaXterm, Putty - NOT RECOMMENDED for iCSF
- What operating system are you connecting from, Windows, Linux, OSX?
  - **Windows** - users need to install a free terminal app called MobaXterm http://mobaxterm.mobatek.net/download-home-edition.html
  - **Linux and Mac** - both have a terminal application by default
  - **Mac** – depending on the version of OSX you may need to install X-Quartz in order to view GUI http://xquartz.macosforge.org/

MANCHESTER 1824
The University of Manchester

RESEARCH VIRTUAL DESKTOP SERVICE

http://ri.itservices.manchester.ac.uk/course/icsf-csf/

Research Virtual Desktop Service & X2Go

- Recommended method for connecting to the iCSF can be used to connect to other platforms within the CIR ecosystem
- Accessible from off/on-campus
- Requires a small piece of Software called X2Go (open-source) - MAC users will also have to download Xquartz
- Setup Guide - http://ri.itservices.manchester.ac.uk/virtual-desktop-service/x2go/
- PLEASE USE either nyx5, nyx6, nyx7 - these provide the virtual desktop
- Login with central username & password

Benefits of Research Virtual Desktop Service

- Accessible from off/on-campus
- Provides a remote desktop enabling users to connect to the platforms
- Start a session at work power off your work PC at the end of the day, go home and **reconnect** to the same session continue from where you left off.
    - Connecting directly to CIR platforms without going via the RVDS does **NOT** provide the reconnect functionality e.g. Closing Mobaxterm which is directly connected to the iCSF will also log you out of iCSF and kill any running jobs. :(
- X2go by design uses compression to make GUI's run smoothly even over slow connections

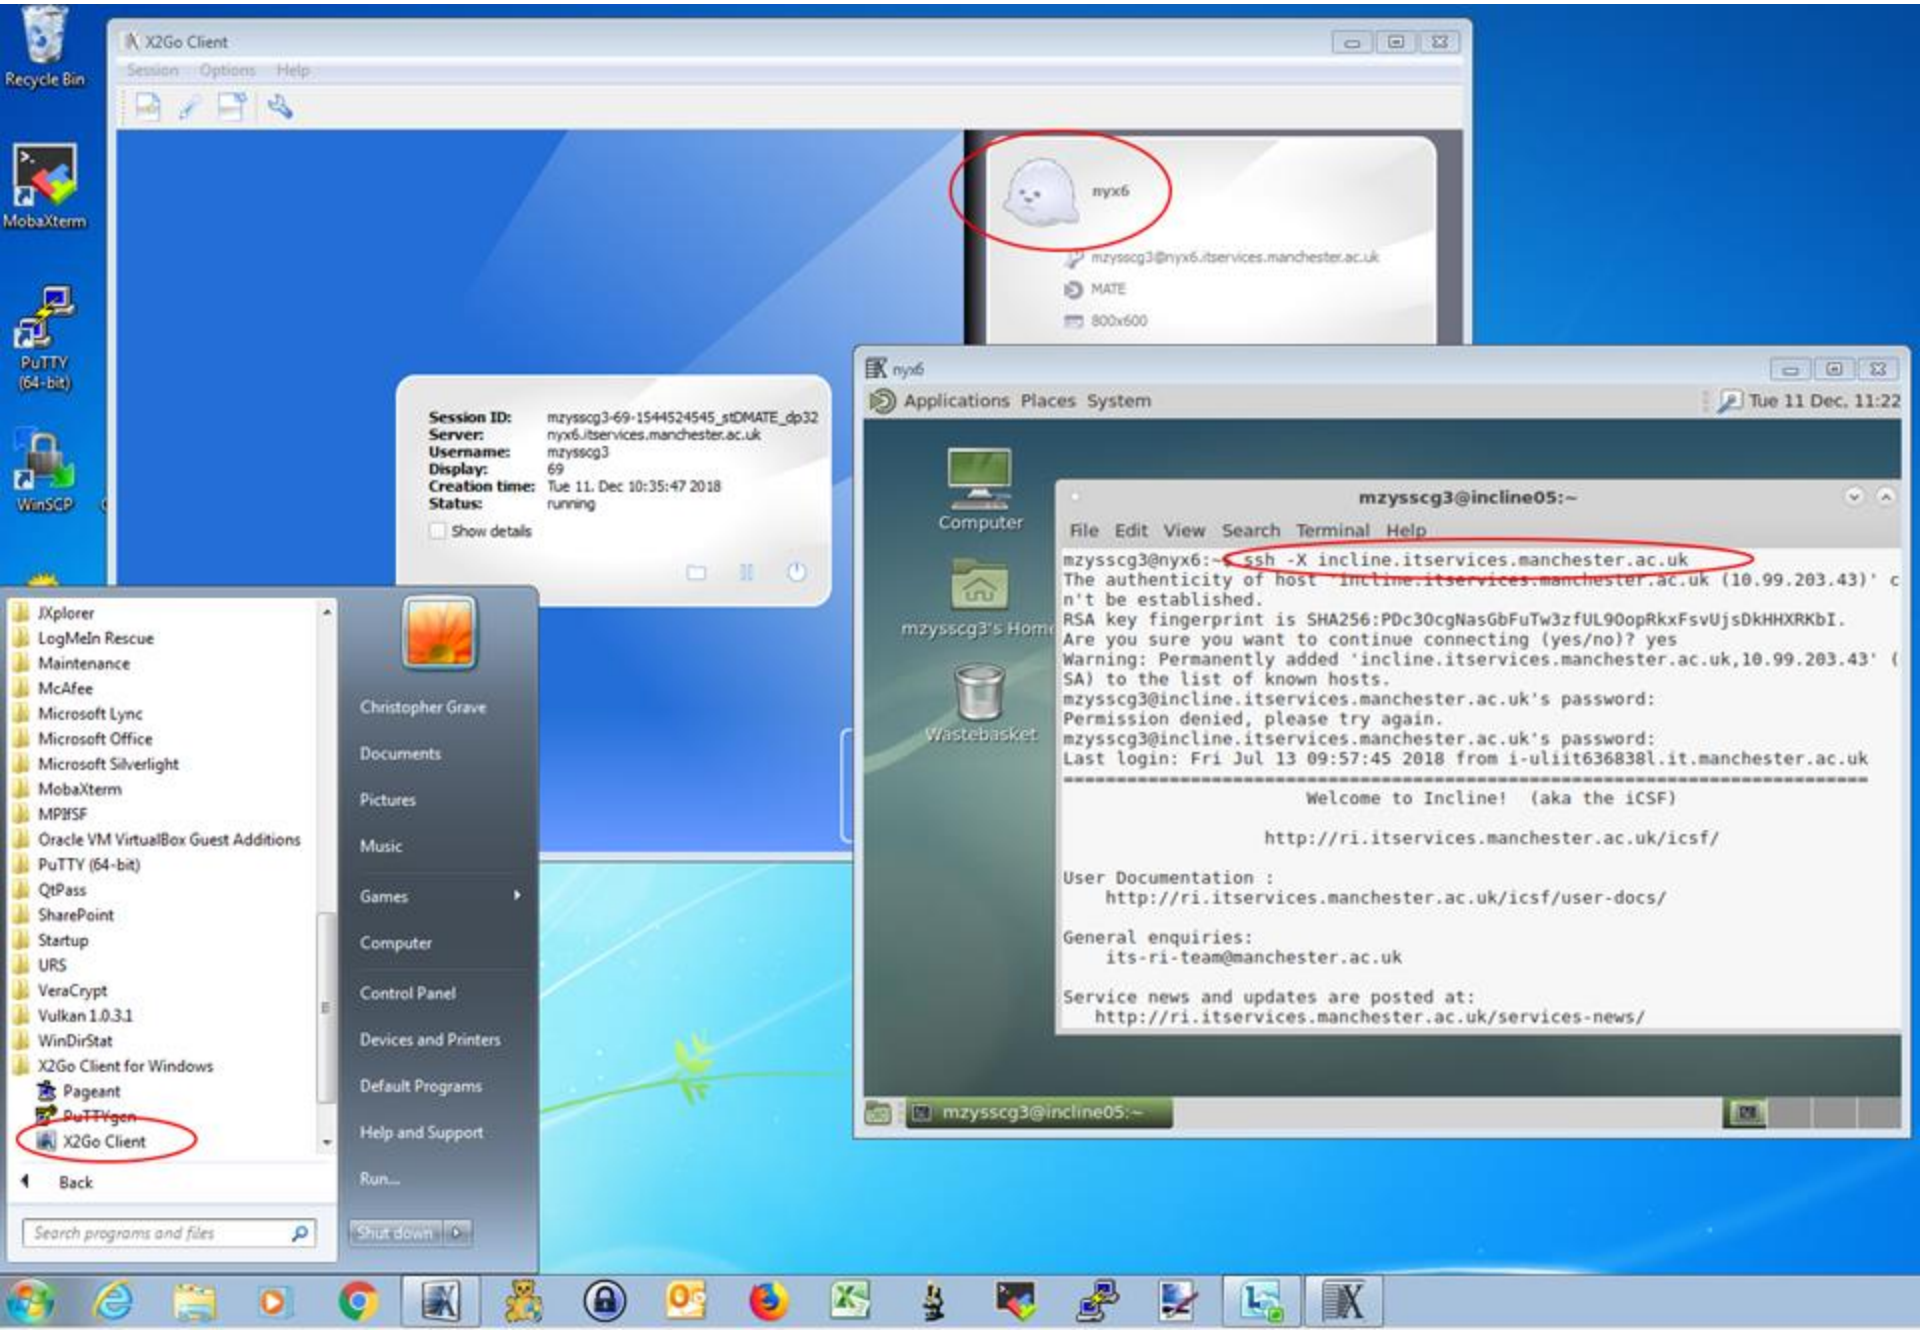
X2Go Client
Session Options Help
nyx6
mzysscg3@nyx6.itservices.manchester.ac.uk
MATE
800x600
Session ID: mzysscg3-69-1544524545_stDMATE_dp32
Server: nyx6.itservices.manchester.ac.uk
Username: mzysscg3
Display: 69
Creation time: Tue 11. Dec 10:35:47 2018
Status: running
Show details
nyx6
Applications Places System
Tue 11 Dec, 11:22
Computer
mzysscg3's Home
Wastebasket
mzysscg3@incline05:~
File Edit View Search Terminal Help
mzysscg3@nyx6:~$ ssh -X incline.itservices.manchester.ac.uk
The authenticity of host 'incline.itservices.manchester.ac.uk (10.99.203.43)' c
n't be established.
RSA key fingerprint is SHA256:PDc3OcgNasGbFuTw3zfUL9OopRkxFsvUjsDkHHXRKbI.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'incline.itservices.manchester.ac.uk,10.99.203.43' (
SA) to the list of known hosts.
mzysscg3@incline.itservices.manchester.ac.uk's password:
Permission denied, please try again.
mzysscg3@incline.itservices.manchester.ac.uk's password:
Last login: Fri Jul 13 09:57:45 2018 from i-uliit636838l.it.manchester.ac.uk
Welcome to Incline! (aka the iCSF)
http://ri.itservices.manchester.ac.uk/icsf/
User Documentation :
http://ri.itservices.manchester.ac.uk/icsf/user-docs/
General enquiries:
its-ri-team@manchester.ac.uk
Service news and updates are posted at:
http://ri.itservices.manchester.ac.uk/services-news/
mzysscg3@incline05:~
JXplorer
LogMeIn Rescue
Maintenance
McAfee
Microsoft Lync
Microsoft Office
Microsoft Silverlight
MobaXterm
MPIfSF
Oracle VM VirtualBox Guest Additions
PuTTY (64-bit)
QtPass
SharePoint
Startup
URS
VeraCrypt
Vulkan 1.0.3.1
WinDirStat
X2Go Client for Windows
Pageant
PuTTYgen
X2Go Client
Back
Search programs and files
Christopher Grave
Documents
Pictures
Music
Games
Computer
Control Panel
Devices and Printers
Default Programs
Help and Support
Run...
Shut down

MANCHESTER 1824
The University of Manchester

RVDS Exercise

- Connect to incline using the RVDS

MANCHESTER 1824
The University of Manchester

Research Data Storage (RDS) & CIR Filesystems

http://ri.itservices.manchester.ac.uk/course/icsf-csf/

Research Data Storage Service

- Also commonly known as isilon (the technology on which it relies)
- Data stored on RDS can be considered secure e.g. files can be recovered if accidently deleted
  - Automatically backed up once an hour (changes kept for 24 hours
  - Automatically backed up daily and copy kept for 35 days
- Does **not** provide Encryption at Rest therefore if handling restrictive/highly restrictive data please speak to us first.
- Can be accessed from your PC both on and off campus
- Groups can share files easily

MANCHESTER 1824
The University of Manchester

Can I have some RDS

- Each funded Research group entitled to 8TB free of charge
  - Over 8TB costs money
- Must be requested by academic responsible for the group /project e.g. PI
- All users' home directories reside on RDS
  - Project data should reside in dedicated Project folders not home
  - Projects should request further storage if required
- If requesting RDS
  - Research Data Management – Research Council requirement that you have a Research Data Management plan – Outside scope of this presentation for more info
  - http://www.library.manchester.ac.uk/services-and-support/staff/research/services/research-data-management/

Filesystems cont….

- After logging onto the iCSF/CSF you will be automatically placed into your home directory.

```
[username@incline31 ~]$ pwd (tells you where you are)
/mnt/hum01-home01/username
```

- Home space is shared by everyone in the group and space is limited. It is not designed for long term storage of data - use project space
- Only you can access your home directory. However sharing is possible using file permissions. Please contact Research IT regarding this.

MANCHESTER 1824
The University of Manchester

Filesystems cont….

- You can see how much space you are using with the command

  ```
  $ du -sh
  ```

- How big is that file? -

  ```
  $ ls -lh filename
  ```

- How much space is used/free overall?

  ```
  $ df -h .
  ```
  (don't forget the .)

MANCHESTER 1824
The University of Manchester

Uploading and Downloading Files

- Various ways this can be achieved. MobaXterm, WinSCP, RDS-SSH, SSHFS & command line utilities e.g. scp, rsync.
- How large are the files you need to transfer?
- What operating system are you using?
- Documentation can be found at the following location

http://ri.itservices.manchester.ac.uk/userdocs/file-transfer/

MANCHESTER 1824
The University of Manchester

WinSCP

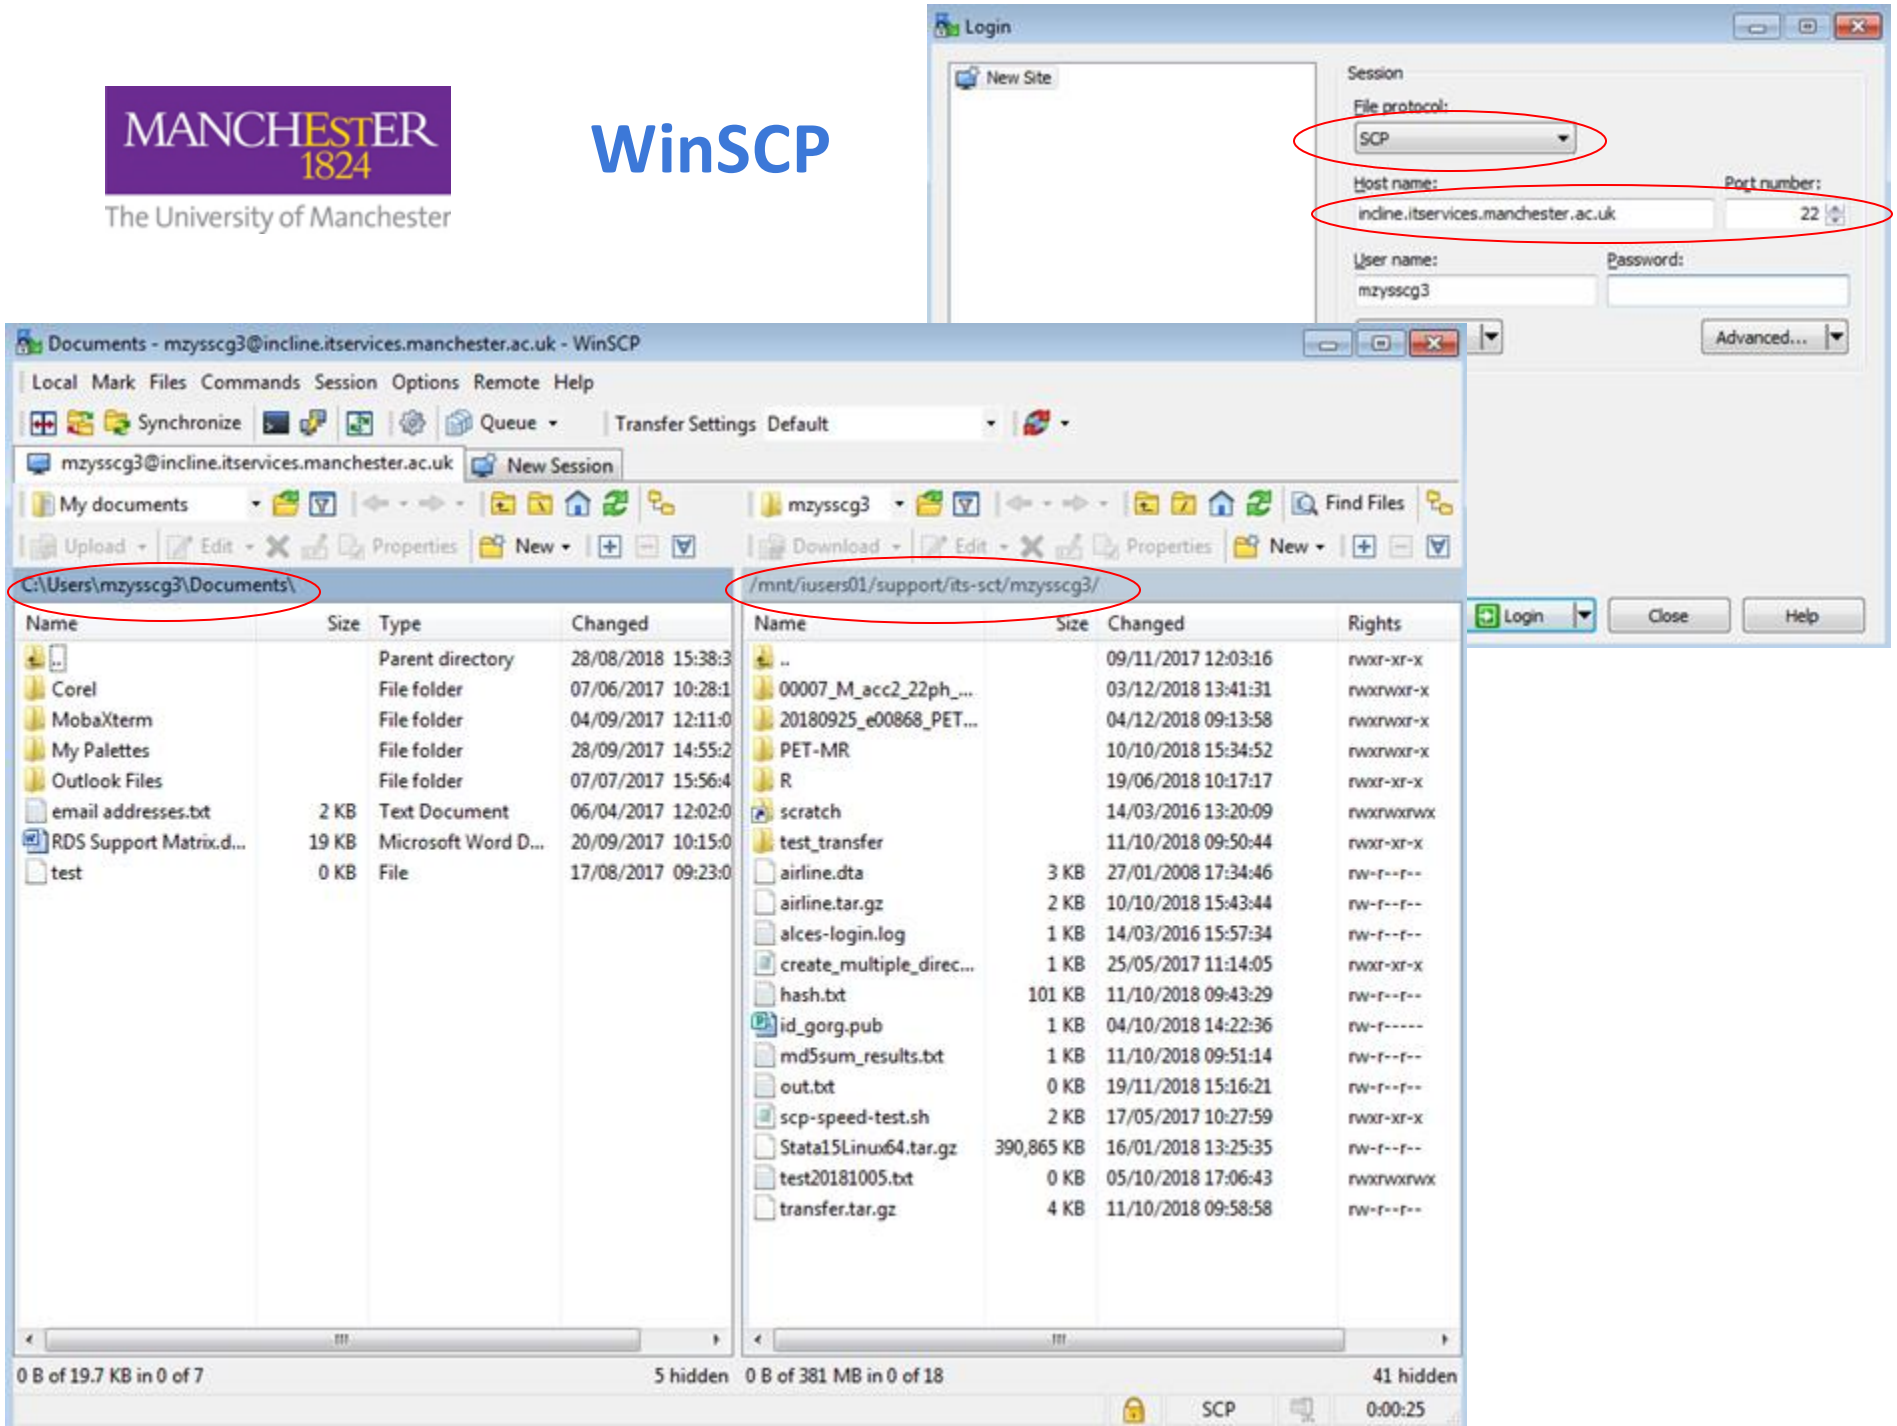

Uploading & Downloading Files via the Command Line e.g. Linux or OSX

- Upload a file from your computer to your iCSF home dir

```
scp file1.txt username@incline2000.itservices.manchester.ac.uk:
```

Secure Copy Protocol

No destination after the : means "use your home directory

- Download a file from your iCSF home dir to your computer

```
scp username@icsf.itservices.manchester.ac.uk:results.out .
```

The . translates to "the current directory" on your computer

- Upload a file from your computer to an CSF folder in home

```
scp file2.txt username@csf.itservices.manchester.ac.uk:mydata/
```

- Download a file from your CSF folder in home to your computer

```
scp username@csf.itservices.manchester.ac.uk:mydata/results.out results.copy
```

Will rename the copied file to "results.copy"

MANCHESTER 1824
The University of Manchester

File Transfer Exercise

- Upload a file to the iCSF using WinSCP
- Manipulate a file using Linux commands

MANCHESTER 1824
The University of Manchester

USING THE iCSF

(logging in, setting up and running applications)

http://ri.itservices.manchester.ac.uk/course/icsf-csf/

MANCHESTER 1824
The University of Manchester

iCSF Resource

- The iCSF is a shared resource you will likely be sharing the compute nodes with other people
- The login node assesses the load on them before deciding which one to place you on.
- 7x Super-High-Memory nodes - 32core 2TB (2000GB) of memory

```
ssh -X
centralusername@incline2000.itservices.manchester.ac.uk
```

Connecting to the iCSF

- Open X2go application and connect to to Virtual Desktop Service i.e. nyx 5,6,7.
  - From Menu - ***Applications > System Tools > MATE Terminal***
- In Terminal log into the iCSF
  - `ssh -X username@incline2000.itservices.manchester.ac.uk`
  - Perform 2FA when prompted
- You will end up on a backend node e.g. `incline31`
- The node is indicated by the 2 digit number in the above example
  - 31-37 = Super High Memory 2TB nodes

Starting Apps

- All software is controlled and loaded using module files
- List available software

```
module avail
```

Not all software will be available due to license restrictions

- All software is loaded via modulefiles e.g.

```
module load apps/binapps/rstudio/0.98.1103-any-r
module load apps/binapps/matlab/R2017a
```

- List currently loaded modules i.e.

```
module list
```

- Once module file is loaded open the software (the & lets you keep using the terminal)

```
rstudio &
matlab &
```

- Voila! GUI/window should open
- More information on available software, and module files please follow this link

http://ri.itservices.manchester.ac.uk/icsf/software/

MANCHESTER 1824
The University of Manchester

iCSF Exercise

- Load and Run Rstudio
- Logout of the iCSF

MANCHESTER 1824
The University of Manchester

Using the CSF

http://ri.itservices.manchester.ac.uk/course/icsf-csf/

What we'll be using today - the CSF

- CSF3 current config:
  - A large Linux cluster system
  - 14,016 CPU cores (Intel “Xeon” or AMD “EPYC” CPUs)
  - 152 Nvidia GPUs (68 x v100, 72 x A100, 12 x L40s)
  - Got big datasets to process? Can run large-memory jobs
  - (we'll cover all of these details throughout the course)

You *do not* need to be running huge parallel jobs, or be a Linux / HPC expert, to use our systems and to benefit from the CSF

How can I access the CSF

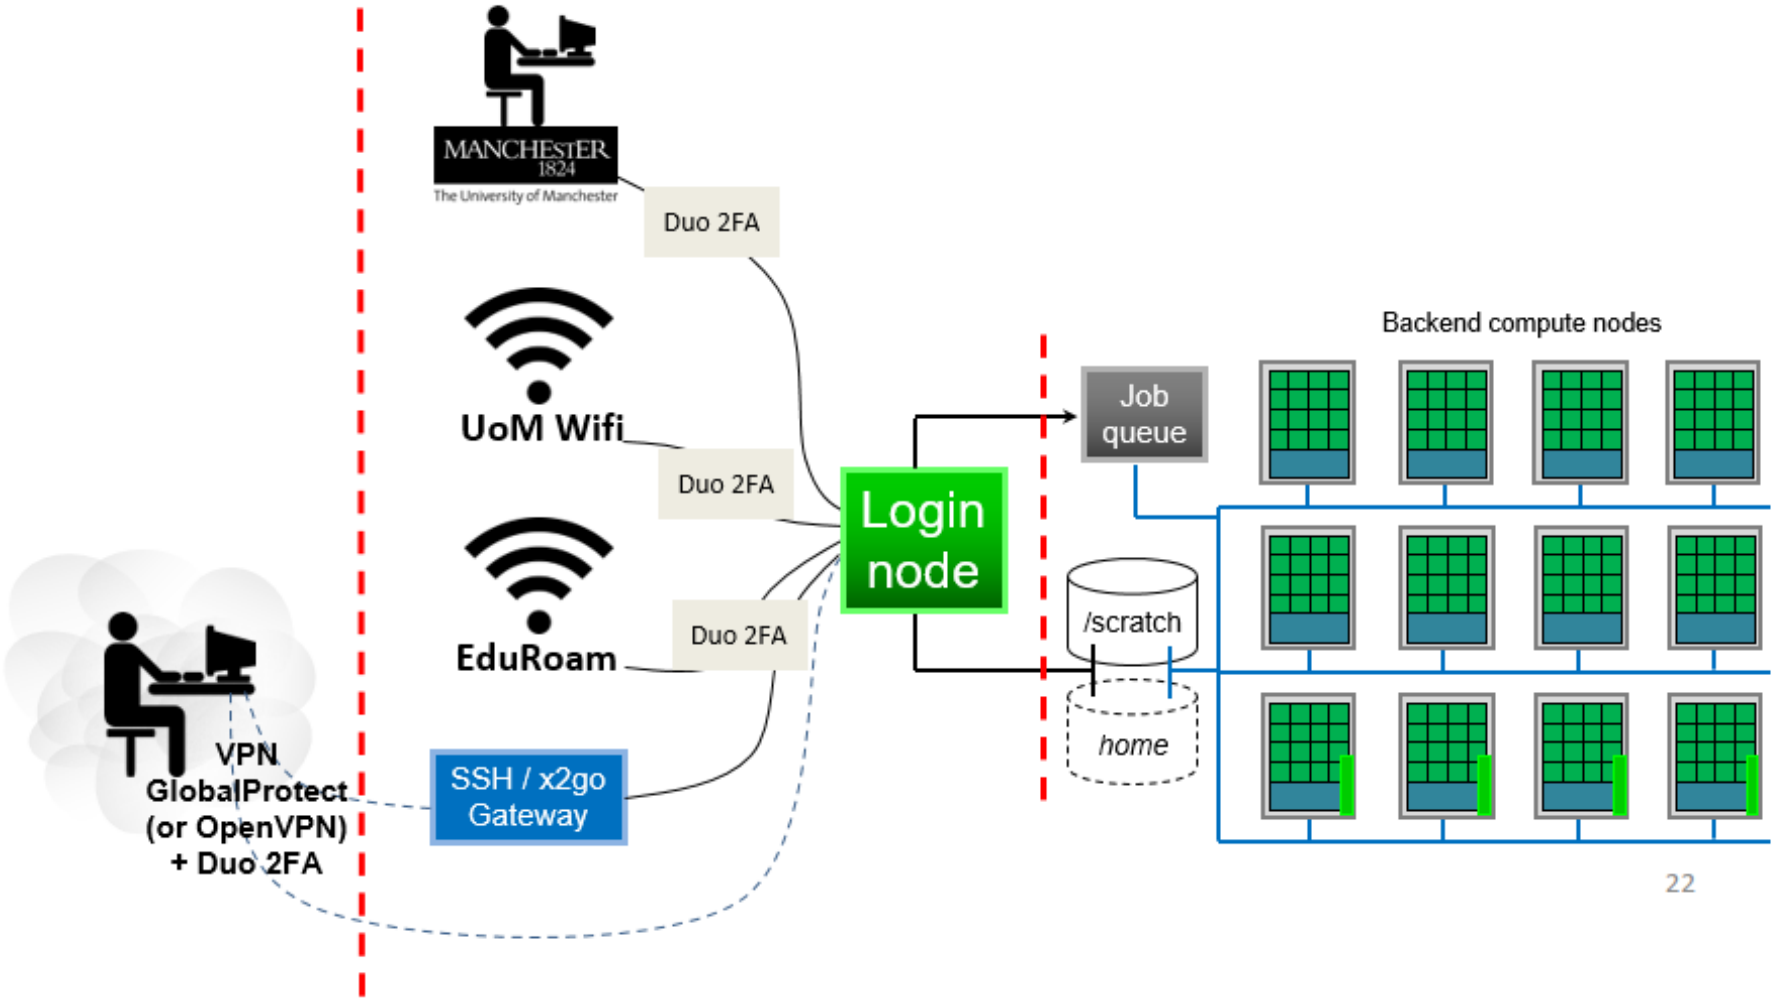

Connect to CSF from Windows

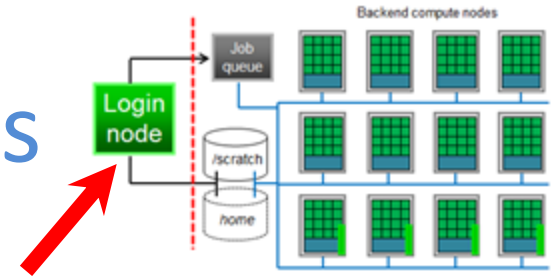

- Access the CSF from a PC / laptop using an SSH (**S**ecure **Sh**ell) app
  - Sometimes called a "terminal".
  - There's no web-site or other fancy GUI on the CSF – use the "command-line".
- **Windows users** need to install a free *terminal* app called **MobaXterm**
- https://mobaxterm.mobatek.net/download-home-edition.html
  the Home edition (portable edition) does *not require* Administrator rights - just *extract* the small .zip file in your P-Drive or USB stick for example.

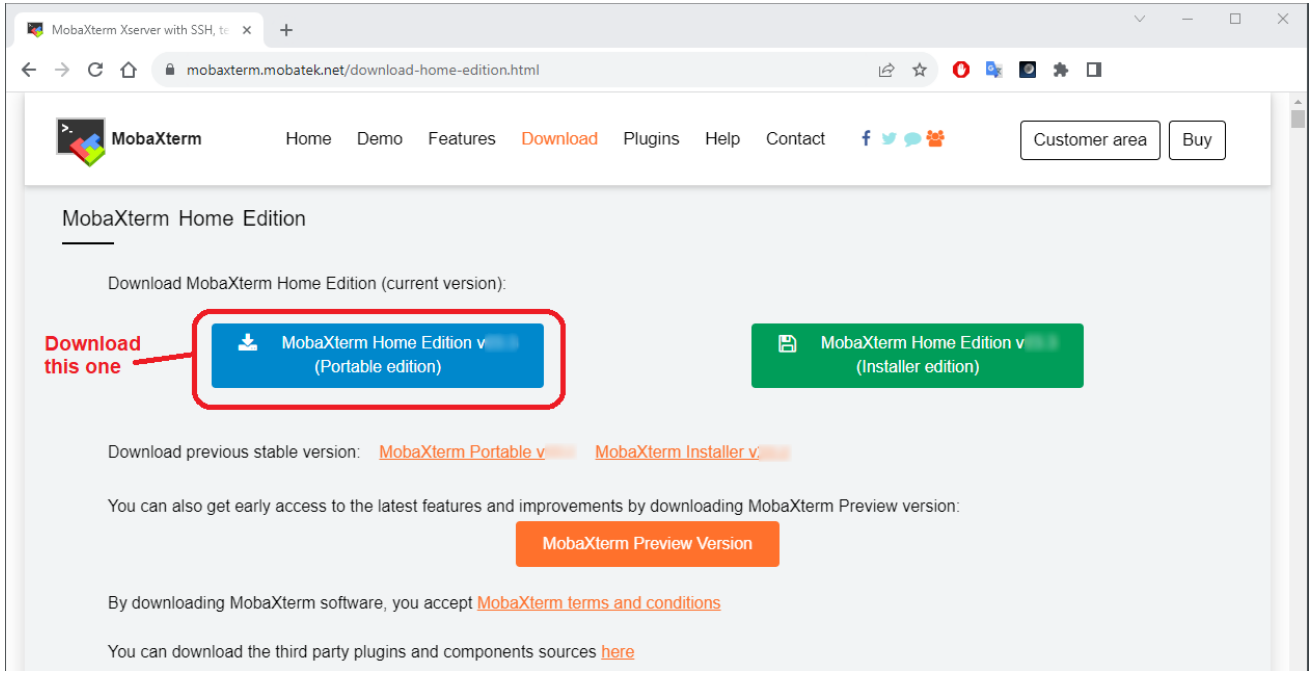

1. Download using the blue box.
2. Once downloaded, *right-click* on the .zip file and select:

   "**Extract all** …"

   This will *unpack* the .zip file to a folder.

MobaXterm "Session"

(username saved in the session setup)

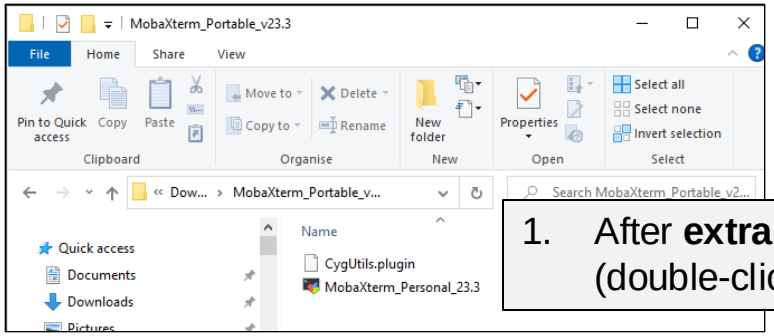

1. After **extracting** the .zip file, start MobaXterm_Personal_*xy.z* (double-click on the icon)

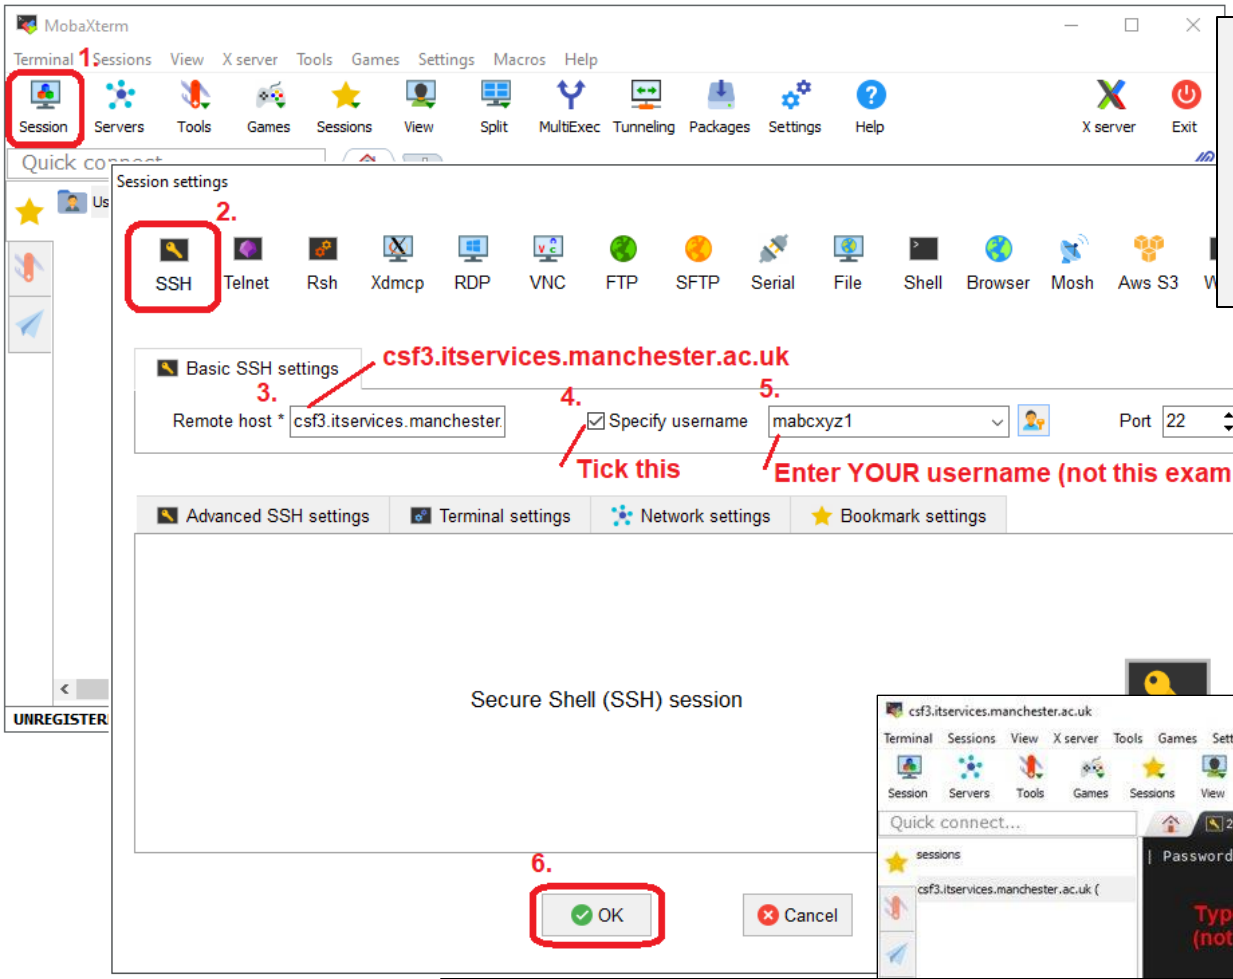

2 (1-6). Create a "Session" which saves the CSF's details along with *your* **username**.

This is needed to make file drag-n-drop work (see later.)

3. This will then start to log you into the CSF – it will ask for your password. Type carefully!

4. See slide about 2FA – you may be asked for DUO after your password

If asked to save your password, we recommend you say "No", for security.

Drag-n-drop file browser for upload / download

(new users won't have as many items in the list!)

We're on (one of) the CSF login nodes. Any commands you use will be typed "at the prompt", which shows your username and current directory (folder.)

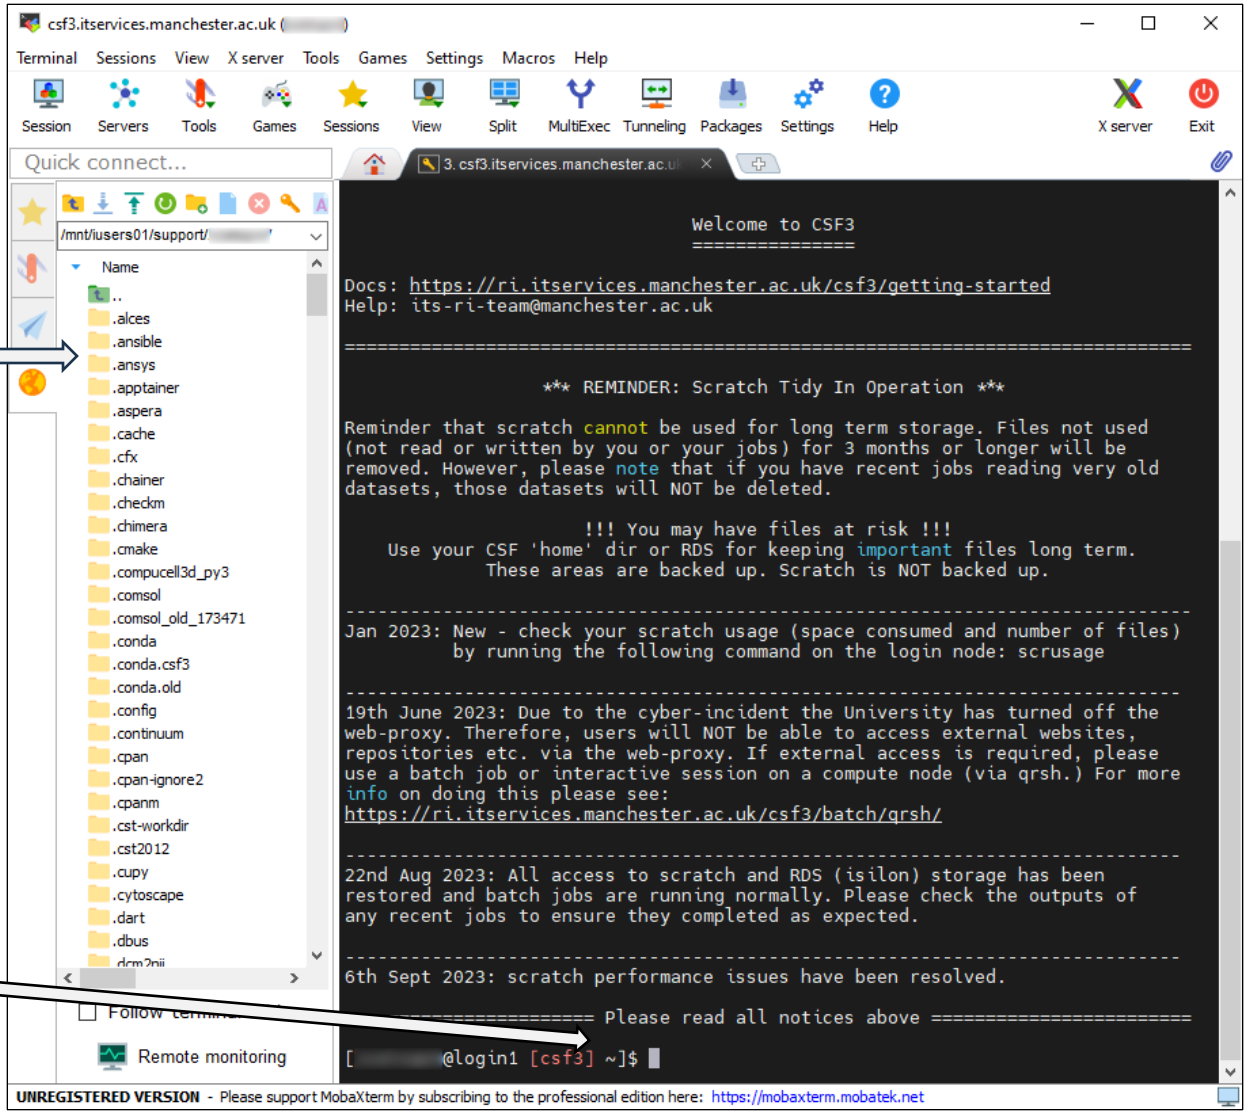

Next time you want to login to CSF from Windows

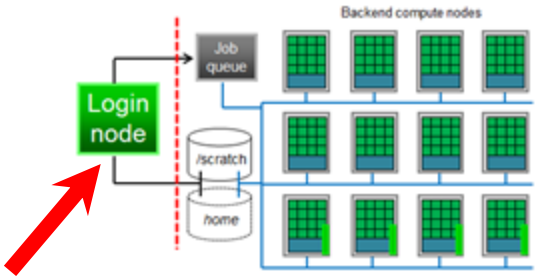

- Just double-click the csf3 "session" in the list of "User sessions"
- The CSF details are saved in the "session"
- (this also makes the file browser work, for drag-n-drop file transfers.)

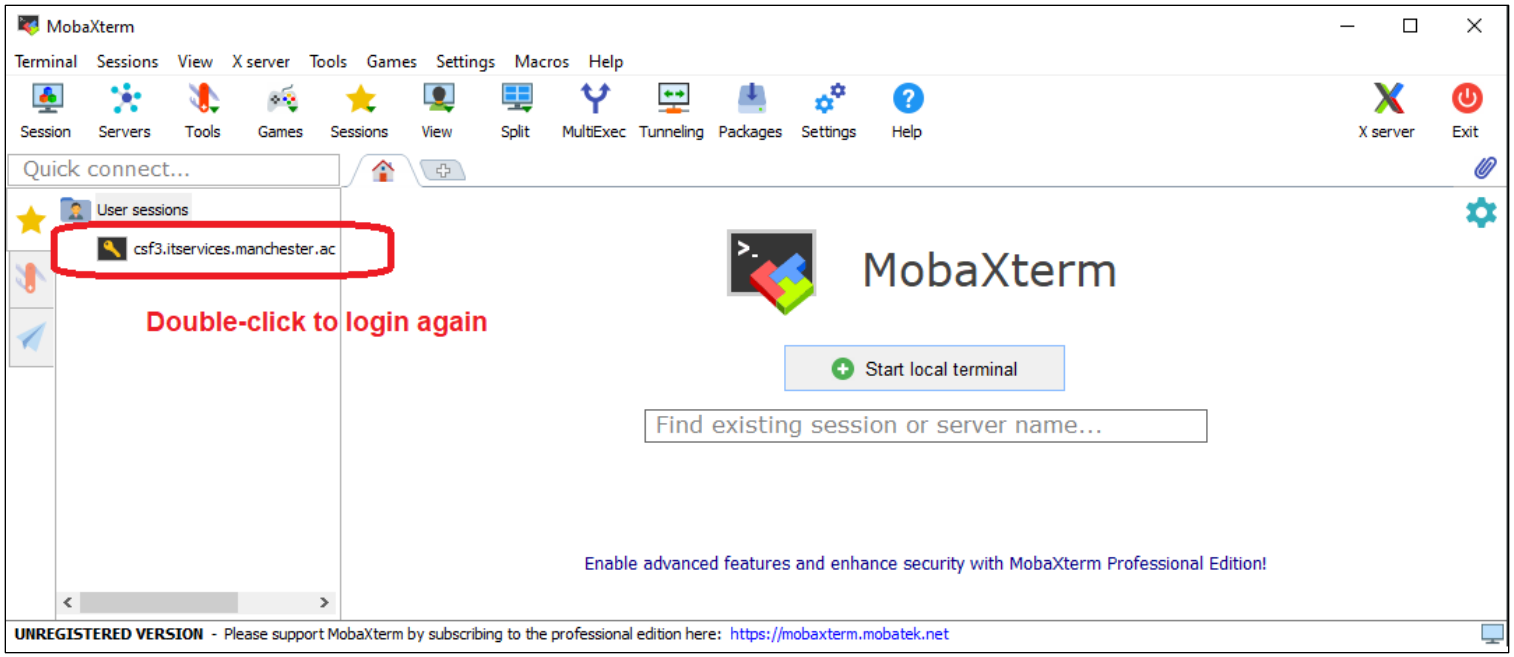

Connect to CSF from a Mac

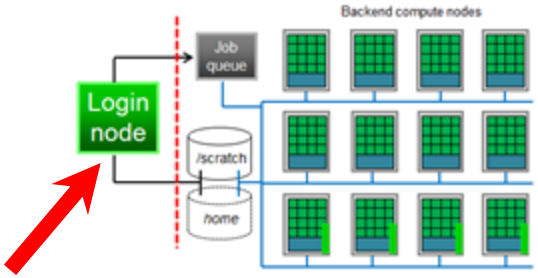

- Access the CSF from a PC / laptop using an SSH (Secure Shell) app, eg a 'terminal'.
  - There's no web-site or other fancy GUI on the CSF - command-line for now.

- **Mac** users - have a *terminal* application by default
  - You will need to install X-Quartz first

    http://xquartz.macosforge.org/ (install then you should reboot your Mac)
  - Start a Terminal and type the following command:

```
ssh -Y username@csf3.itservices.manchester.ac.uk
```

UPPERcase Y

Central IT Services username.
Answer 'Yes' to continue *if* asked.
Enter central IT password when asked (same as for email)

- Finished using CSF? Log out with: `logout` or `exit`

Connect to CSF from Linux

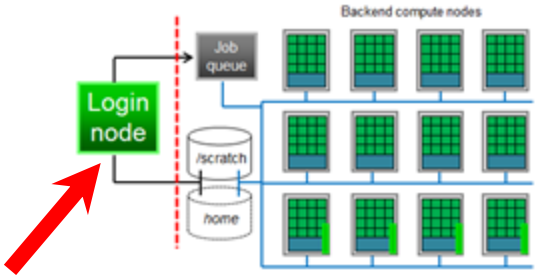

- Access the CSF from a PC / laptop using an SSH (Secure Shell) app, eg a 'terminal'.
  - There's no web-site or other fancy GUI on the CSF - command-line for now.

- **Linux** users - have a *terminal* application by default
  - Start a Terminal and type the following command:

```
ssh -X username@csf3.itservices.manchester.ac.uk
```

UPPERcase X

Central IT Services username.
Answer 'Yes' to continue *if* asked.
Enter central IT password when asked (same as for email)

- Finished using CSF? Log out with: `logout` or `exit`

Linux / Mac Terminals

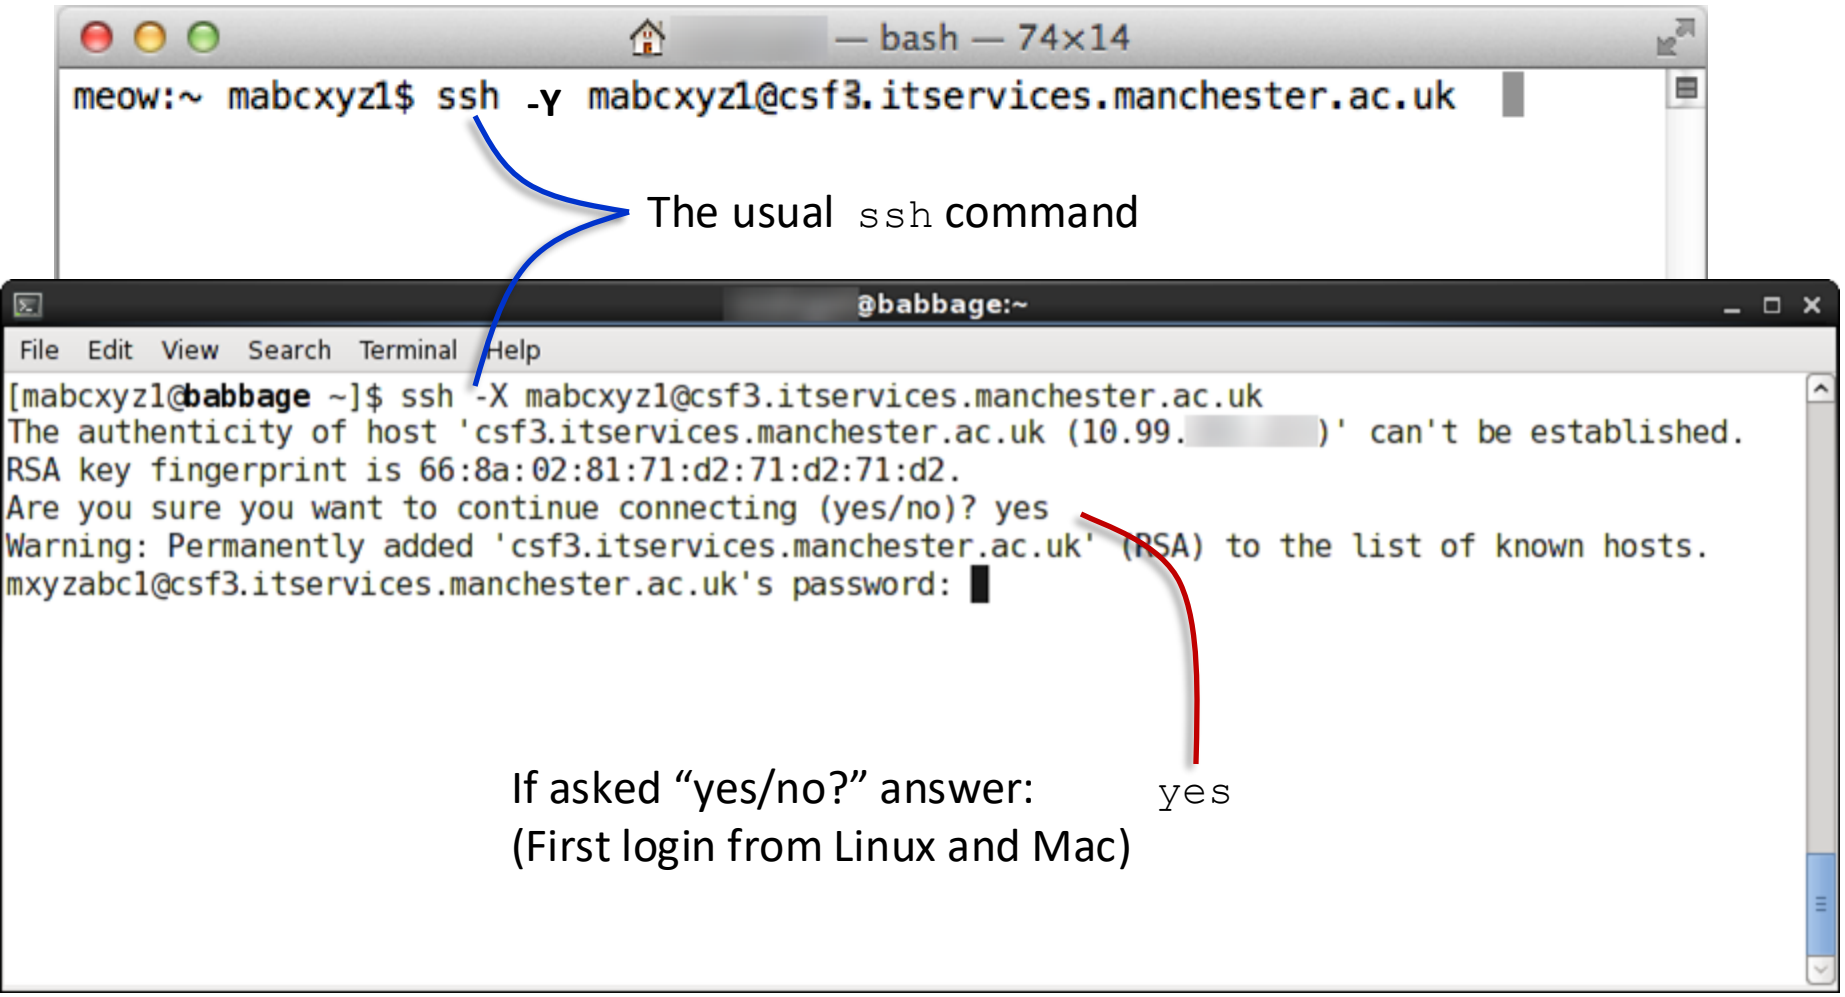

What you see when you log in

- CSF uses CentOS Linux (c.f. Red Hat EL)
  - Command line – requires the input of commands, can be a little scary at first to new users
  - A welcome *message of the day* - announcements
  - The system awaits input/commands from you at a *prompt* (after you've *logged in*):

```
[username@login1 [csf3] ~]$
or [username@login2 [csf3] ~]$
```

- Learning linux commands (more later):
  - http://www.ee.surrey.ac.uk/Teaching/Unix

Type your commands at the prompt

Warning: The login nodes

- Do *not* run computational work here:
  - Not enough cores
  - Not enough memory
  - 100+ users connected, so running work causes serious problems
- You *can* do the following:
  - Transfer files on and off the CSF
  - Set up and submit your jobs (covered in next few slides)
  - Basic data processing/viewing
- Computational work running on the login node will be killed without warning!

CSF Scratch filesystem

- An extra and very large filesystem available on the CSF
- Shared by all CSF users
- Run jobs here - it is fast storage, sometimes home can fill up causing jobs to fail
- You should have a ‘symlink’ to it in your home directory (if not ask us to create one)

```
cd ~/scratch
```

- Not backed up
  - When job finishes move important files to home or other RDS area
  - Not considered safe for long term storage – hardware failure could cause loss
  - Files that have not been accessed in the past 3 months will be deleted

MANCHESTER 1824
The University of Manchester

A simple batch job

- We recommend you have a directory for every job

  All files relevant to the job need to be in the folder

  Be tidy, it makes managing results easier

- We need a jobscript – Create it on the CSF using this command

  `gedit &`

  this avoids a problem with linux/batch not being able to read text files created on windows.

- Like iCSF software is accessed by modulefiles

  eg: `module load apps/R/3.4.2`

  - Unlike iCSF, **most** of the time we only use `module load` in a jobscript

A simple Jobscript – *Serial* (1 core)

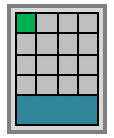

`#!` on first line only (a special line)

First line indicates we use the *bash* script language to write our jobscript.

`myjob.txt`

```
#!/bin/bash --login

#$ -cwd
#$ -N myjob
#$ -l resource

# Let's do work
date
hostname
sleep 120
date
```

`#$` indicates a **batch system parameter** to specify our job requirements. We'll use various combinations of these.

`-cwd` indicates we'll run from our current (working) directory. Input / output files will usually be found here.

`-N` (optional). Set the *jobname*. Otherwise will use name of your jobscript as the name.

`#` lines are just comments - anything on the line after it will be ignored.

`-l` (optional) used to add extra resource requirements e.g. memory, time limits

Actual Linux commands we run in our job. They will execute on a compute node.

#$ -l course only works on the day of a course.

MANCHESTER 1824
The University of Manchester

Submit and monitor the job

- To give the job to the system:

  ```
  qsub jobscript
  ```

  Replacing ‘jobscript’ with the name of the text file you have created.

- To see the job in the batch system

  ```
  qstat
  ```

  - You should see the job either queued (qw) or (r)
  - Your job will have a unique numerical JobID
  - If you can’t see the job it has finished. :)

MANCHESTER 1824
The University of Manchester

Where did my results go?

- Two possibilities
  - A file called **jobname.oJobID**
    Where jobname is the name of your submission script
    JobID is the number of your job
  - An output file specific to your job
- If your job had a problem or failed it may be indicated in one of the above or **jobname.eJobID**
- Various options to view the file:

  `cat filename`

  `more filename` (allows you to page through with spacebar)

  `gedit filename` (not recommended if it is large)
- Depending on the application and type of files produced you might want to use the iCSF to view them, postprocess, visulaise, etc.

MANCHESTER 1824
The University of Manchester

Jobs do not always start quickly

- Your job will wait until the resources requested become available.
- You can log off, switch off your PC and your job will stay on the CSF. Log in again to collect the results.
- You can submit more than one job at once, more than one may run at the sametime (good to have different folder for each)

MANCHESTER 1824
The University of Manchester

What does Eqw mean? How do I delete a job?

- `qstat` reports your job as `Eqw`
  - CSF tried to start it, but something went wrong
  - Usually no output file so to indicate what happened use `qstat -j JobID | grep error`
- Error can be cryptic
- Most common causes
  - Missing directory
  - You created the jobscript on windows
  - Unusual characters or spaces in file and directory names
  - No disk space on the filesystem – highly recommend you use scratch to avoid this
- To delete a problem job or one you simply no longer want:
  - `qdel jobID`

MANCHESTER 1824
The University of Manchester

Parallel Jobs on the CSF

- CSF excels at doing this
- Each compute node has multiple processors, each of which has multiple cores, each of which can run a single task
- A parallel job uses more than one of these cores
  - Specify how many cores are required
  - Software splits your job into chunks and runs them all at the sametime
  - Usually means jobs run **faster**
  - Allows access to **more memory**
  - At the end pulls everything together into your results
- Requires a parallel environment (pe) -

https://ri.itservices.manchester.ac.uk/csf3/batch/parallel-jobs/#Intel_parallel_environments

There's another way to run in parallel! High Throughput Computing (HTC)

- Lots of *independent* computations
  - Processing lots of data files (e.g., image files)
  - Running the same simulation many times over with different parameters ("parameter sweeps")
- Run many copies of your program
  - Programs may be serial (single core) but running lots of them at once. They don't communicate.

Parallel Jobscript – Multi-core (*single-node*)

myparajob.txt

```
#!/bin/bash --login
#$ -cwd
#$ -pe smp.pe 4

# Set up to use a chemistry app
module load apps/binapps/anaconda2/2021.11

# Inform app how many cores to use
export OMP_NUM_THREADS=$NSLOTS

# This job runs a python script
python3 name_of_python_script.py
```

#! and #$ see serial jobscript earlier.

-pe indicates we'll run a parallel job in a particular **p**arallel **e**nvironment.

smp.pe is the **p**arallel **e**nvironment name. This one means: app will use a single compute node (up to 32 cores.)

4 is the number of **cores** we want to *reserve* in the system. Each **PE** has a maximum allowed.

# indicates line is a comment so does nothing.

Any commands we run in our job. They will execute on a backend node that has required number of cores free. **mdrun_d** is gromacs.

Key concept!

Must *somehow* inform the app how many cores we *reserved*. **$NSLOTS** is automatically set to number **(4)** given on **-pe** line. **Our app** wants **OMP_NUM_THREADS** environment variable setting. **Your app** might use a different method!

MANCHESTER 1824
The University of Manchester

What else can the CSF do?

- Access more cores across multiple-nodes. Up to 1024 via the HPC Pool
- Job Arrays – Multiple Similar Jobs - https://ri.itservices.manchester.ac.uk/csf3/batch/job-arrays/
- High Memory jobs, access up to 4TB ideal for processing large DataSets -https://ri.itservices.manchester.ac.uk/csf3/batch/high-memory-jobs/
- Over 150 GPUs, ideal for machine learning type workloads
- One day CSF course covers the above

MANCHESTER 1824
The University of Manchester

CSF Exercise 1:

- Connect to CSF via RVDS
- Run a serial & parallel python jobs
- View the results on the CSF

MANCHESTER 1824
The University of Manchester

Other Research IT Infrastructure

Provided Services

HPC Pool

Dedicated pool for “true” HPC jobs

- 4096 cores
- Minimum 128-core job size, maximum 1024
- Frontend shared with CSF3
- Lightweight application process – must be made by PI
- Currently free
- http://ri.itservices.manchester.ac.uk/csf3/hpc-pool

Research Virtual Machine Service (RVMS)

- Another component of the CIR ecosystem
- Centrally-hosted virtual machines (VMs) are provided for University research staff and PG students
- RVMS is suitable for VMs which are expected to generate large amounts of network traffic, may be moderately CPU-intensive or require tight integration with the computational platforms (CSF, iCSF, etc.) or the Research Data Service (RDS).
- Administrative control of VMs provided maybe handed over to research groups.
- More info http://ri.itservices.manchester.ac.uk/rvms/

MANCHESTER 1824
The University of Manchester

Condor

- Another component of CIR ecosystem
- Utilises 100s of Desktop PCs which are idle on campus
  - Hence typically only available overnight, weekends, holidays.
- Suitable for jobs which require the same computation i.e. repeated many times over different data sets/ranges. Also known as High Throughput Computing.
- Burst into the Cloud using Amazon Web Services
- More Information can be found by following this link

  http://ri.itservices.manchester.ac.uk/htccondor/

MANCHESTER 1824
The University of Manchester

Research IT

- Thank you for attending!
- Course Materials - http://ri.itservices.manchester.ac.uk/course/icsf-csf/
- Feedback PLEASE COMPLETE - https://goo.gl/forms/zfZyTLw4DDaySnCF3

- Support/Queries -https://ri.itservices.manchester.ac.uk/hpc-help

Chris Grave, Research IT,

its-ri-team@manchester.ac.uk