IPPP Systems

A Quick Intro (Again?) and How To

Presented by Adam Boutcher

Nov 2020

Durham University

IP3

IPPP IT Support

- **Adam Boutcher**
  - adam.j.boutcher@durham.ac.uk
  - OC216
- **Paul Clark**
  - paul.w.clark@durham.ac.uk
  - OC216
- **IPPP Operations**
  - oper.ip3@durham.ac.uk
  - http://help.phyip3.dur.ac.uk

Adam

Paul

Durham University

Desktops & Laptops

- **Desktop Groups 2 & 3**
  - Intel Core i5
  - 8GB RAM
  - 1TB HDD
  - Fedora Linux
- **Desktop Groups 4,5 & 6**
  - Intel Core i7
  - 8GB or 16GB RAM
  - 1TB HDD or 256GB SSD
  - Fedora Linux

- **Lenovo ThinkPad T490**
  - Intel Core i5
  - 16GB RAM
  - 512GB SSD
  - (MDS) Windows *or* Fedora Linux
- **Lenovo ThinkPad L14**
  - AMD Ryzen 5
  - 16GB RAM
  - 512GB SSD
  - (MDS) Windows *or* Fedora Linux
- **MacBook Pro**
  - Intel Core i5
  - 8GB RAM
  - 256GB SSD Hard Drive
  - Mac OS

Durham University

Compute Nodes

- **Workstation Nodes**
  - Generic end-user style system
  - Interactive Use
  - SSH etc
- **CPU Compute**
  - Tuned for performance
  - Non-Interactive
  - Batch Submission
  - No SSH etc
- **GPU Compute**
  - Systems with GPU available
  - Currently SSH/Interactive
  - Not for general use

Durham University

Workstation Nodes

- **WS1 - Offline**
- **WS2 - Reserved**
- **WS3**
  - Intel Xeon 32 Threads @2.4Ghz
  - 64GB RAM
  - SSH
- **WS4**
  - Intel Xeon 40 Threads @2.2Ghz
  - 64GB RAM
  - Docker
- **WS5**
  - Intel Xeon 40 Threads @2.2Ghz
  - 64GB RAM
  - Docker
- **WS6 - High Memory**
  - Intel Xeon 40 Threads @2.2Ghz
  - 1TB RAM
  - 1TB Hard Drive
- **WS7**
  - Intel Xeon 40 Threads @2.6Ghz
  - 128GB RAM
- **WS8 - Reserved**
- **WS9**
  - Intel Xeon 40 Threads @2.6Ghz
  - 128GB RAM
- **WS10 - Offline**
- **WS11 - Reserved**
- **WS12 - Reserved**

Durham University

Due to change before Jan 2021

CPU Compute Nodes

- **CPU 1-7**
  - 2x Intel Xeon Gold 5220 @2.2 Ghz (Total 36 Cores / 72 Threads)
  - 192GB RAM
  - 1TB Scratch SSD
  - CentOS Linux 7
  - Upto 72 Job Slots

- **CPU 8**
  - 4x Intel Xeon Gold 6130 @2.1Ghz (Total 64 Cores / 128 Threads)
  - 960GB RAM
  - 2TB Scratch HDD
  - CentOS Linux 7
  - Upto 128 Job Slots

- **Accessed via SLURM with --partition=cpu**

Durham University

GPU Compute Nodes

- **GPU 1 and GPU 4**
  - 2x Intel Xeon Silver 4216 @2.1Ghz (Total 32 Cores / 64 Threads)
  - 92GB RAM
  - 1TB Scratch SSD
  - NVIDIA Tesla V100 32GB

- **GPU 2 and GPU 3**
  - 2x Intel Xeon Sliver 4216 @2.1Ghz (Total 32 Cores / 64 Threads)
  - 92GB RAM
  - 1TB Scratch SSD
  - NVIDIA Tesla V100 16GB

- **GPU 5**
  - 2x Intel Xeon Silver 4216 @2.1Ghz (Total 32 Cores / 64 Threads)
  - 92GB RAM
  - 1TB Scratch SSD
  - NVIDIA P100 12GB

Durham University

- **Currently accessed via SSH – Give us feedback please!**

Storage

- **Home Space**
  - 20GB Default
  - Fast SSD
  - Backed Up Daily
  - Limited Expansion Available upon Request
- **Batch Space – Due to change before Jan 2021**
  - 300GB Default
  - HDD
  - Expansion Available upon Request
- **“Cloud” Storage**
  - 15GB Default
  - Backed Up Daily
  - Limited Expansion Available upon Request
  - Contact us for an account

- **Home Space**
  - /mt/home/<USERNAME>
  - ~/
- **Personal Webspace**
  - ~/www
  - https://www.ippp.dur.ac.uk/~username
- **Batch Space**
  - /mt/batch/users/<USERNAME>
- **Precompiled HEP Software**
  - /mt/hepsw/
- **Local Machine Storage**
  - /scratch/<USERNAME>

Durham University

Software

- **Installed:**
  - CentOS Linux 7 (WS/CPU/GPU) / Fedora Linux (Desktop)
  - Gnome Windowing System
  - Python 2.7
  - Python 3
  - GCC
  - Mathematica 12
  - Mathematica 11
  - Maple 2020
  - Slurm
  - CVS – GIT/HG/SVN

Durham University

Other software may be available upon request.

Things to Come…

- **New - IPPP Services Infrastructure**
  - Highly Available = Fewer Downtimes due to system failure
  - More Services = More Demand
- **New - Batch Disk Server**
  - ~200TB Total
  - ~2TB Quota per User with some Expansion
  - More robust to disk failure
  - Better IO
- **Improvements**
  - WS Systems
  - Login Systems
  - Graphical Login
  - Jupyter Notebook/Hub Support
  - Network Upgrades - Dec 2020
  - Seafile Upgrades
  - Docker/Singularity Support (Tentative Plans)
  - Changes to internal security

Durham University

How To: IPPP Account

Your CIS Account is **NOT** your IPPP Account…

First time you login, change your password and **READ** the instructions carefully.

Durham University

How to: Remote Access IPPP

SSH - CLI

ssh username@login.phyip3.dur.ac.uk
ssh username@login2.phyip3.dur.ac.uk

X2GO – GUI

Host: login.phyip3.dur.ac.uk
Session: XFCE

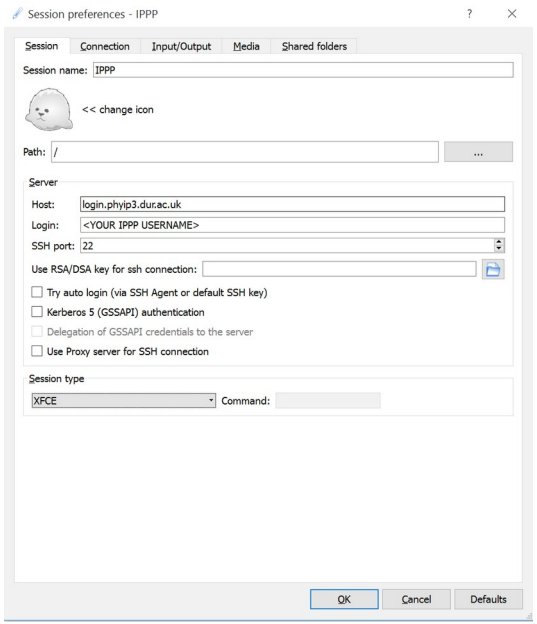

Please note that the login box is **NOT** for compilation or compute!

Changes to SSH before Jan 2021

Desktop via Web Browser will be available in 2021

Durham University

How to: Python Virtual Env

- Portable
- Directory Name is the Environment Name

Create:

```
#~: python2 -m virtualenv NAME
#~: python3 -m venv NAME
```

Use:

```
#~: source NAME/bin/activate
(NAME) #~: pip install whatever
```

Durham University

How to: Submit to Batch

Slurm Batch System

1. Create a job file
2. Run sbatch job.sh
3. srun will parallel run/fork the tasks
4. Use Output

```
#!/bin/bash

# These are Slurm SBATCH commands
#SBATCH --job-name="PythonVersionCheck"          # Job Name
#SBATCH --mail-type=END                          # When to send mail (BEGIN, END, FAIL, REQUEUE, ALL)
#SBATCH --mail-user=mailname@durham.ac.uk        # Where to send mail.  A valid email address
#SBATCH --error="Job-%j.err"                     # Redirect STDERR (Error output) to this file %j is a variable for JobID
#SBATCH --output="Job-%j.out"                    # Redirect STDOUT (Normal output) to this file %j is a variable for JobID
#SBATCH --mem=2G                                 # Requested memory for the Job (Default is 2G)
#SBATCH --export=ALL                             # Export Current Environment Variables (Default ALL)
#SBATCH -D /mt/batch/users/aboutcher             # Put all output on batch storage (--chdir currently doesn't work)

# The rest is similar to a standard shell script
srun echo "Python Version Check"
srun hostname
srun python --version 2>&1

# Remember to exit cleanly
exit 0;
```

```
login.phyip3.dur.ac.uk - PuTTY
aboutcher@login:~$ sbatch ./MyPythonVersion.batch.sh
Submitted batch job 6960313
aboutcher@login:~$
```

- **Request CPU Access**
  - --partition=cpu
- **Exclude a Node**
  - --exclude=ip3-cpu[3,7-8]

Durham University

How to: Submit to Batch

- Use /scratch for any IO as much as possible
- Copy to/from /mt/batch at the start/end of your job
- Pipe your jobs stdout and stderr to /scratch
- Treat SBATCH nearly like a BASH script (Env vars not parsed by SBATCH)
- Lots of little files? – Tar it up first!

```
#Heavy IO Job

mkdir /scratch/$USERNAME/
rsync /mt/batch/$USERNAME/my_input_files /scratch/$USERNAME/my_input_files
cd /scratch/$USERNAME/

python dothings_with_my_input_files 1>/scratch/$USERNAME/out.log 2>/scratch/$USERNAME/err.log

rsync /scratch/$USERNAME/my_output_files /mt/batch/$USERNAME/my_output_files
rsync /scratch/$USERNAME/*.log /mt/batch/$USERNAME/
```

Durham University

How to: Use the GPU

- Login to the GPU and source the version of CUDA that's available (CUDA 10 or 10.1)

- Source /opt/cuda-10.1/enable
- Run your jobs – GPU /dev/nvidia0

```
aboutcher@login:~$ ssh ip3-gpu1
Last login: Mon Nov  2 16:46:43 2020 from login.phyip3.dur.ac.uk
aboutcher@ip3-gpu1:~$ source /opt/cuda-10.1/enable
NVIDIA Cuda-10.1 Enabled
aboutcher@ip3-gpu1:~$ nvidia-smi
Mon Nov  2 16:47:26 2020
+-----------------------------------------------------------------------------+
| NVIDIA-SMI 418.39       Driver Version: 418.39       CUDA Version: 10.1     |
|-------------------------------+----------------------+----------------------+
| GPU  Name        Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf  Pwr:Usage/Cap|         Memory-Usage | GPU-Util  Compute M. |
|===============================+======================+======================|
|   0  Tesla V100-PCIE...  Off  | 00000000:3B:00.0 Off |                    0 |
| N/A   33C    P0    35W / 250W |      0MiB / 32480MiB |      0%      Default |
+-------------------------------+----------------------+----------------------+

+-----------------------------------------------------------------------------+
| Processes:                                                       GPU Memory |
|  GPU       PID   Type   Process name                             Usage      |
|=============================================================================|
|  No running processes found                                                 |
+-----------------------------------------------------------------------------+
aboutcher@ip3-gpu1:~$
```

Durham University

How to: Cloud Storage

- Similar to Dropbox
- 15GB Online Storage Available upon Account Creation
- Supports Mac, Windows, Linux, Android, iOS

- Sync Folders between systems
- Share Folders between users
- ~50 Previous Version History

- Targeted at low infrequent changes / Low volume data

- Some minor changes due before Jan 2021

Durham University

Seafile

Code Repositories

HEPForge

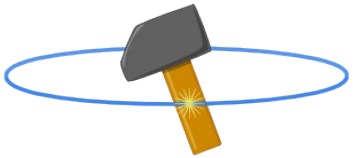

- GIT/SVN/Mercurial
- SSH + HTTPS Push/Pull
- Shell Access
- Web Space
- Downloads
- Large Projects

https://hepforge.org

Gitlab

- GIT
- SSH + HTTPS Push/Pull
- Gitlab Runner/CI
- Smaller Projects
- Grid Accounts or Sign-up

https://gitlab.dur.scotgrid.ac.uk

Durham University

Not Enough Resources?

- **Try the Grid!**
  - ~ 5000 Cores
  - ~ Terabytes of RAM
  - ~ 1PB of redundant Storage
  - More Nodes due online through 2020/2021
  - Access to more resources externally with GridPP

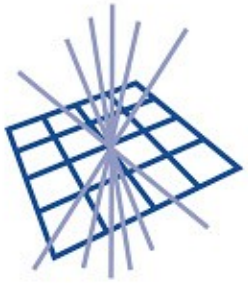

GridPP
UK Computing for Particle Physics

Durham University

More Help

Knowledge Base

http://help.phyip3.dur.ac.uk

A little out of date but is due to change next week!

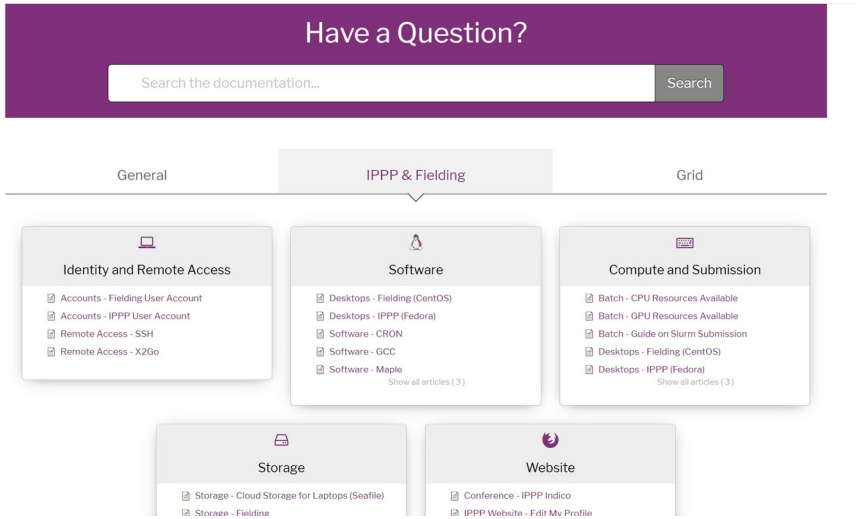

Email

All three of us!

Adam, Paul & Operations

Office

OC216

Please don’t visit us…

We can provide physical support with appointments only.

Durham University

IPPP Systems

A Quick Intro (Again?) and How To

Presented by Adam Boutcher

Nov 2020

Durham University

IP3