Reconciling HPC & LifeSciences

A personal View

Dr. Christian Meesters – Computational Scientist HPC JGU Mainz

JOHANNES GUTENBE
UNIVERSITÄT

Who are You?

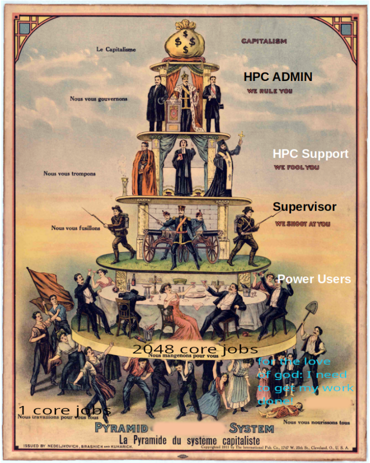

own image after this funny facebook page

- HPC Support / Admin?

Who are You?

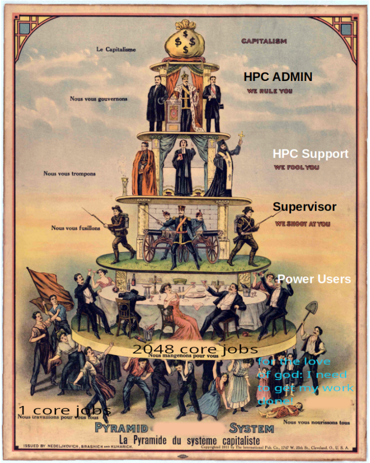

own image after this funny facebook page

- HPC Support / Admin?
- LifeScientist?

Who are You?

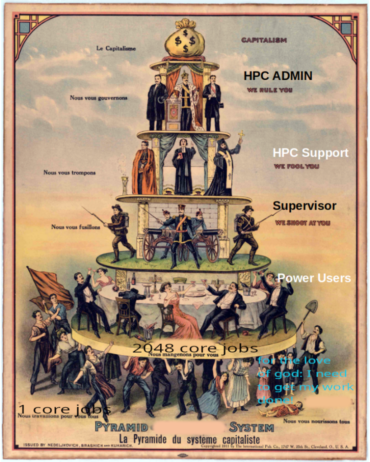

own image after this funny facebook page

- HPC Support / Admin?
- LifeScientist?
- PI?

Who are You?

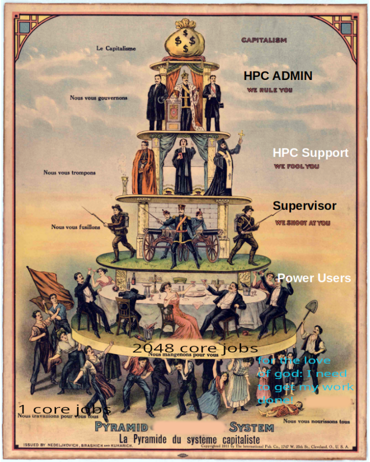

own image after this funny facebook page

- HPC Support / Admin?
- LifeScientist?
- PI?
- Someone else?

What is HPC?

This is what HPC is about!

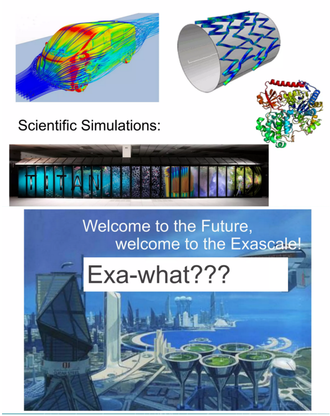

Scientific Simulations:

Welcome to the Future,
welcome to the Exascale!

Exa-what???

slideshare link from colleagues in Barcelona

What is HPC?

This is what HPC is about!

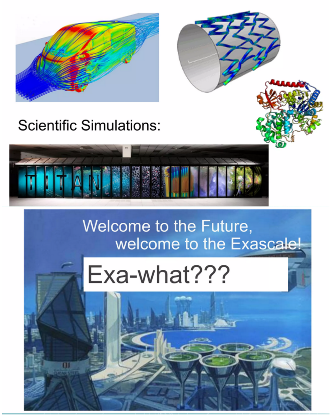

Scientific Simulations:

Welcome to the Future,
welcome to the Exascale!

Exa-what???

slideshare link from colleagues in Barcelona

LifeScience is NOT HPC!

CC BY-SA 3.0 on Wikipedia

Why this is a personal View(!)

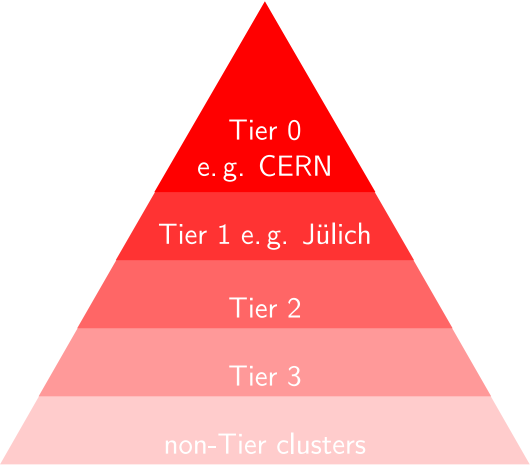

Tier 0
e. g. CERN

Tier 1 e. g. Jülich

Tier 2

Tier 3

non-Tier clusters

Why this is a personal View(!)

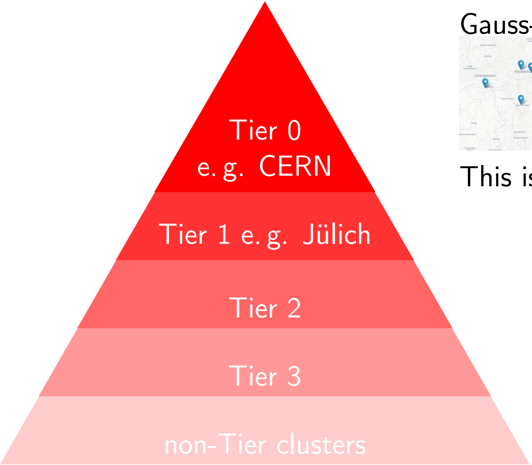

Tier 0
e. g. CERN

Tier 1 e. g. Jülich

Tier 2

Tier 3

non-Tier clusters

HPC Landscapce as seen by the Gauss-Allianz:

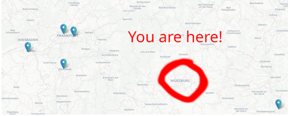

This is where science shines.

Why this is a personal View(!)

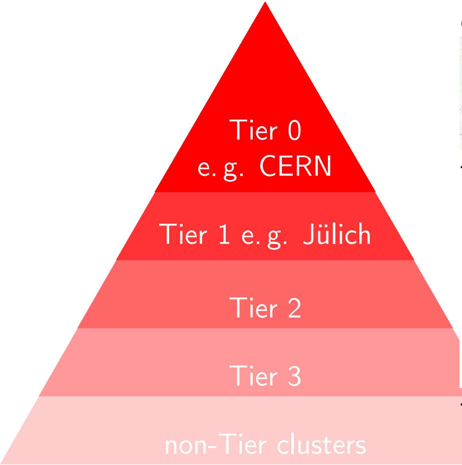

Tier 0
e. g. CERN

Tier 1 e. g. Jülich

Tier 2

Tier 3

non-Tier clusters

HPC Landscapce as seen by the Gauss-Allianz:

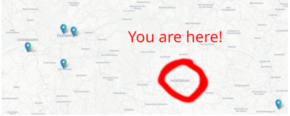

This is where science shines. But

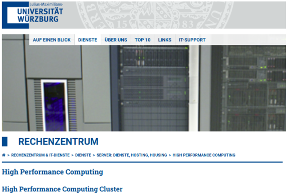

Julius-Maximilians-
UNIVERSITÄT
WÜRZBURG

AUF EINEN BLICK | DIENSTE | ÜBER UNS | TOP 10 | LINKS | IT-SUPPORT

RECHENZENTRUM

> RECHENZENTRUM & IT-DIENSTE > DIENSTE > SERVER: DIENSTE, HOSTING, HOUSING > HIGH PERFORMANCE COMPUTING

High Performance Computing

High Performance Computing Cluster

This is where science is done, too!

Bureau. . . – arrrgh

Many details asked:

> **From the current NHR requirements:**
> Describe all codes, packages or libraries that you need to undertake the project, and how these will enable the research to be achieved. Include for each code to be used information about
>
> - Which code will be used
> - On which hardware the code will be used (CPUs, GPUs, MICs, VEs, etc. or combinations, if applicable)
> - How the code is parallelized (pure MPI, mixed MPI/OpenMP, Pthreads, CUDA, etc.)
> - The amount of memory necessary (per core, per node and in total) Scaling plots and tables with speedup results for runs with typical, parameter sets, problem size, and I/O of the planned project. Scaling data should start with the lowest number of cores possible
> - Current job profile (independent jobs, chained jobs, workflow, etc.)

And similar details for *every* software you use! (data analysts have lots of software) And scaling tests in FLOPS/Resource. (data analysts calculate with lots of integers)

Until Reform . . .

- ask for help – your friendly support scientist is there for you

Until Reform . . .

- ask for help – your friendly support scientist is there for you
- if you buy your own equipment: Tell your local site (and the university), its because of those requirements.

Software Provisioning

- How does your HPC-site provide Software? Are you able to install it yourself? Is it provided? How long do you have to wait?
- How complete is the software provisioning?

Software Provisioning

- How does your HPC-site provide Software? Are you able to install it yourself? Is it provided? How long do you have to wait?
- How complete is the software provisioning?

CONDA

- *Some* HPC-site do not like Conda.
- Conda is slower

Software Provisioning

- How does your HPC-site provide Software? Are you able to install it yourself? Is it provided? How long do you have to wait?
- How complete is the software provisioning?

CONDA

- *Some* HPC-site do not like Conda.
- Conda is slower - or is it?
- Conda is a *necessity* for many workflows

Software Provisioning

- How does your HPC-site provide Software? Are you able to install it yourself? Is it provided? How long do you have to wait?
- How complete is the software provisioning?

CONDA

- *Some* HPC-site do not like Conda.
- Conda is slower - or is it?
- Conda is a *necessity* for many workflows – yet, *no* build system covers 100 % of a researchers need.

Software Provisioning

- How does your HPC-site provide Software? Are you able to install it yourself? Is it provided? How long do you have to wait?
- How complete is the software provisioning?

CONDA

- *Some* HPC-site do not like Conda.
- Conda is slower - or is it?
- Conda is a *necessity* for many workflows – yet, *no* build system covers 100 % of a researchers need.

Basically: Will the software be there for *your* analysis?

How "Scottish" does HPC have to be?

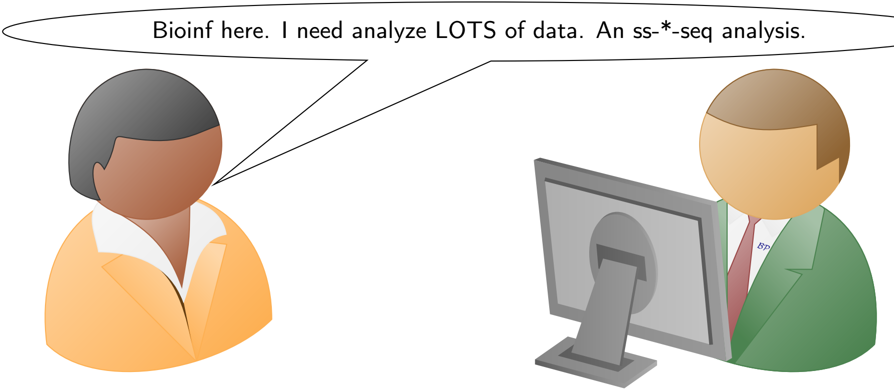

How "Scottish" does HPC have to be?

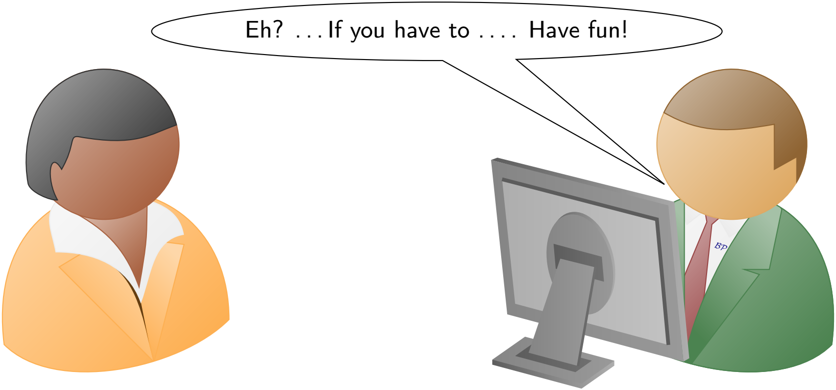

Workflow Managers to the Rescue

snakemake nextflow Open for Innovation KNIME Galaxy PROJECT . . .

(Some) Workflow management systems support HPC batch systems - some better then others.

---

[1] Disclaimer: I am a Snakemake contributor.

Workflow Managers to the Rescue

snakemake nextflow Open for Innovation KNIME Galaxy PROJECT . . .

(Some) Workflow management systems support HPC batch systems - some better then others.
All need (to some extend):

- tailored cluster parametrization
- understanding of the workflow system[1]
- domain knowledge
- some can install your software on the fly . . .

You, as a user, know all this? Fine!

---

[1] Disclaimer: I am a Snakemake contributor.

Workflow Managers to the Rescue

snakemake nextflow Open for Innovation KNIME Galaxy PROJECT . . .

(Some) Workflow management systems support HPC batch systems - some better then others.
All need (to some extend):

- tailored cluster parametrization
- understanding of the workflow system[1]
- domain knowledge
- some can install your software on the fly . . .

You, as a user, know all this? Fine!
You do not? How far does your cluster's support go?

---

[1] Disclaimer: I am a Snakemake contributor.

How "Scottish" does HPC have to be?

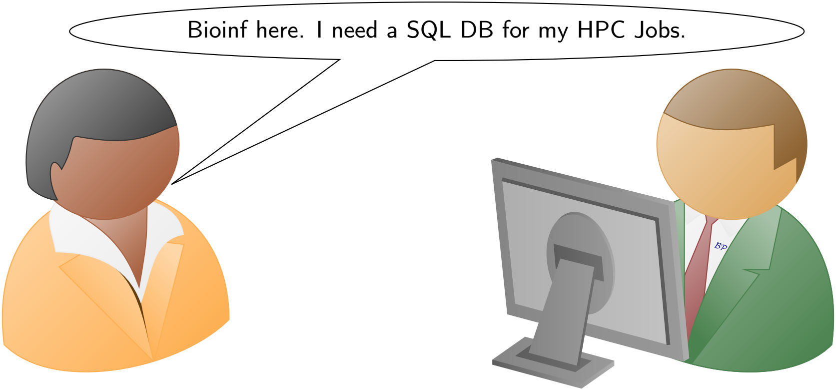

How "Scottish" does HPC have to be?

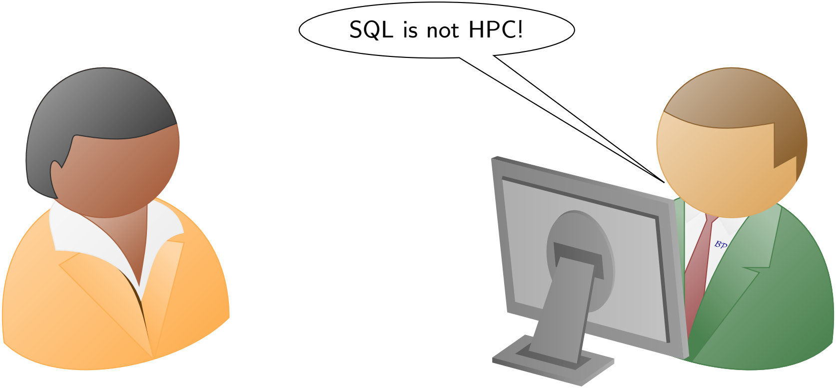

How "Scottish" does HPC have to be?

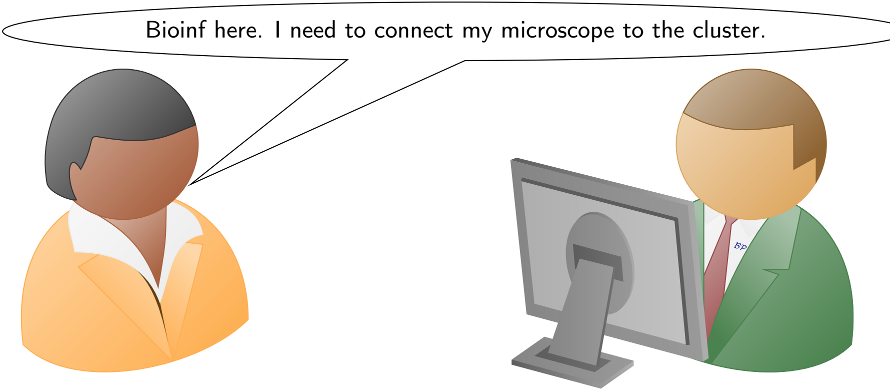

How "Scottish" does HPC have to be?

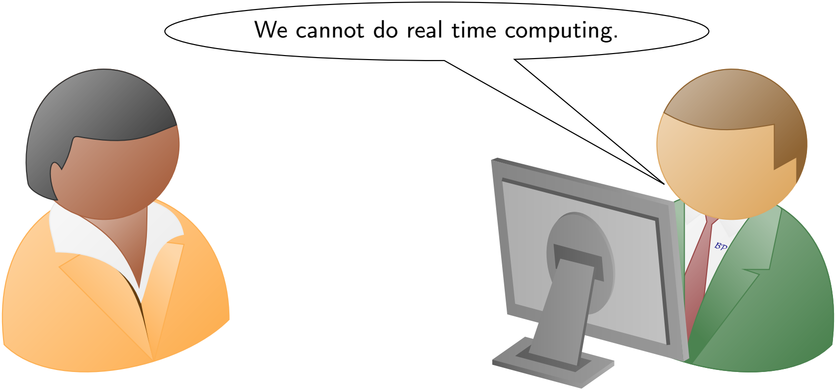

Remedies

If only there would be a solution . . .

```
rule all:
  input:
    from_queue(all_results,
        finish_sentinel=...)

checkpoint somestep:
  input:
    "samples/{sample}.txt"
  output:
    "somestep/{sample}.txt"
  shell:
     "somecommand {input} >"
     " {output}"
```

Snakemake allows to define workflows that are dynamically updated at runtime.

- SQL-DB ⇌ HPC-Job Triggers and Feeds are enabled.

Remedies

If only there would be a solution …

```
rule all:
  input:
    from_queue(all_results,
        finish_sentinel=...)

checkpoint somestep:
  input:
    "samples/{sample}.txt"
  output:
    "somestep/{sample}.txt"
  shell:
     "somecommand {input} >"
     " {output}"
```

Snakemake allows to define workflows that are dynamically updated at runtime.

- SQL-DB ⇌ HPC-Job Triggers and Feeds are enabled.
- Realtime Computing (e. g. Microsope → Deconvolution → Feed into DB) becomes possible for enormous amount of data.

How "Scottish" does HPC have to be?

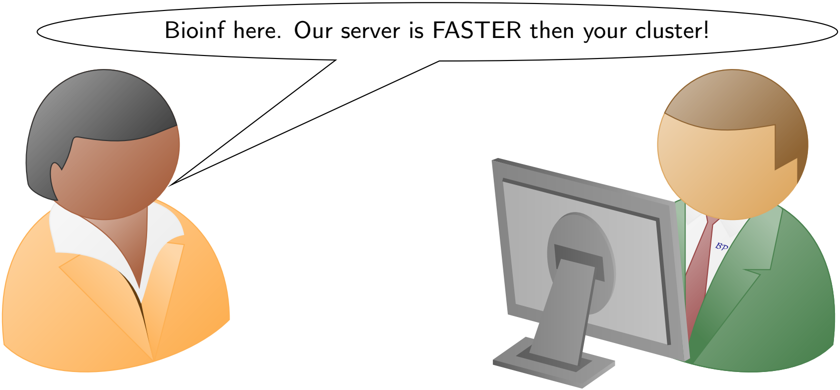

How "Scottish" does HPC have to be?

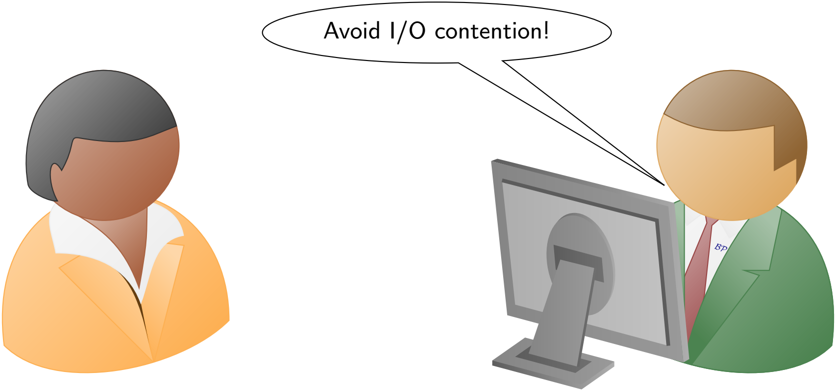

Remedies - Automated Stage-In/-Out

If only there would be a solution . . .

```
$ cat .config/snakemake/config.yaml
__use_yte__: true

__definitions__:
- import os

executor: slurm
latency-wait: 60
default-storage-provider: fs
shared-fs-usage:
- persistence
- sources
- source-cache
?if "SLURM_JOB_ID" in os.environ:
   local-storage-prefix: /localscratch/$SLURM_JOB_ID
?else:
   local-storage-prefix: /dev/shm/$USER/snakemake
```

Remedies - Automated Stage-In/-Out

If only there would be a solution . . .

Configurable, automated stage-in and stage-out of files

- to avoid I/O contention

Remedies - Automated Stage-In/-Out

If only there would be a solution . . .

Configurable, automated stage-in and stage-out of files

- to avoid I/O contention
- enabling *massive* workfloads

Remedies - Automated Stage-In/-Out

If only there would be a solution ...

Configurable, automated stage-in and stage-out of files

- to avoid I/O contention
- enabling *massive* workfloads
- distiction between file systems

How "Scottish" does HPC have to be?

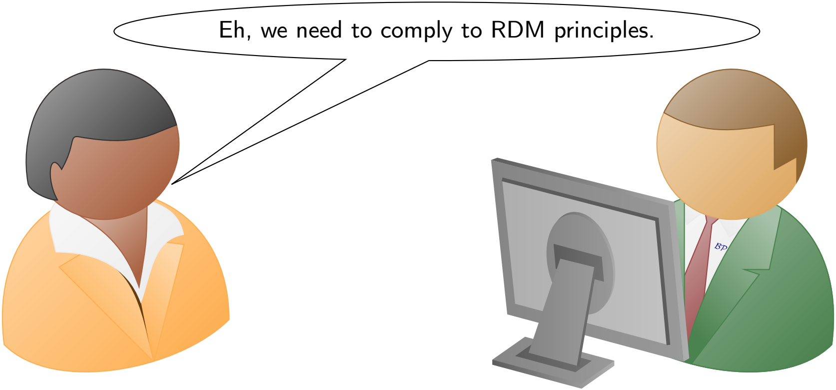

How "Scottish" does HPC have to be?

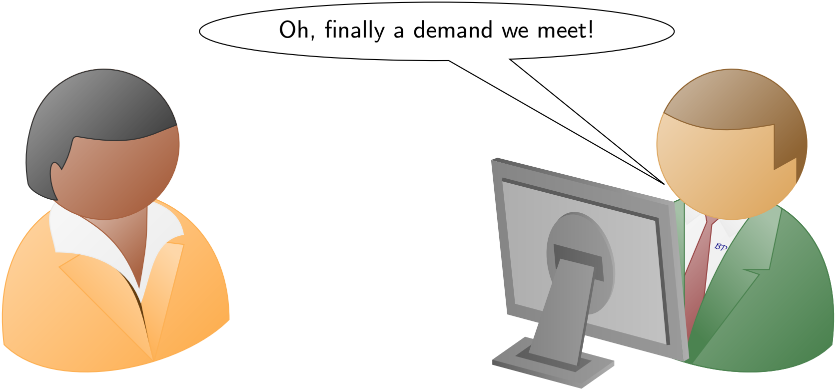

How "Scottish" does HPC have to be?

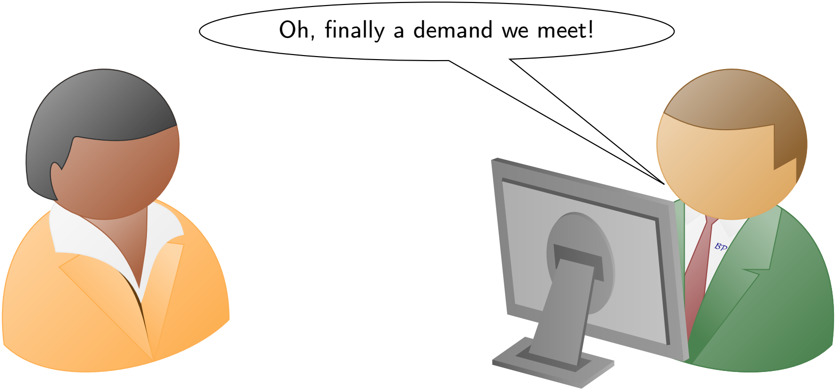

RMD solutions are as divers as HPC centers.

Remedies - Big Data RDM

If only there would be a solution . . .

Remedies - Big Data RDM

If only there would be a solution . . .

Example drawn from iRODS. Other remote providers include NCBI, Dropbox, Gfal, etc. Protocols include S3, ssh, (s)ftp, etc..

```
from snakemake.remote.iRODS import
    RemoteProvider

irods = RemoteProvider(irods_env_file=...,
          ) # all parameters are optional

rule all:
  input:
    irods.remote('{prefix}/testfile.out'),

rule gen:
  input:
    irods.remote('{prefix}/testfile.in')
  output:
    irods.remote('{prefix}/testfile.out')
  shell:
    "touch {output}"
```

Inject code into *existing* workflows for your DM solution:

- collect meta data

How "Scottish" does HPC have to be?

How "Scottish" does HPC have to be?

Remedies - Results

If only there would be a solution . . .

Remedies - Results

If only there would be a solution . . .
Workflow managers provide users with reports (incl. quality ready figures and all meta data). E. g. for Snakemake:

```
$ snakemake --report ...
```

and you get a browsable report.

And Then Some

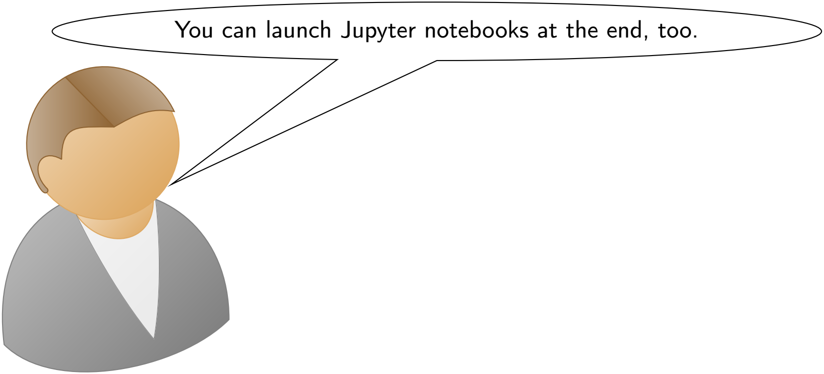

And Then Some

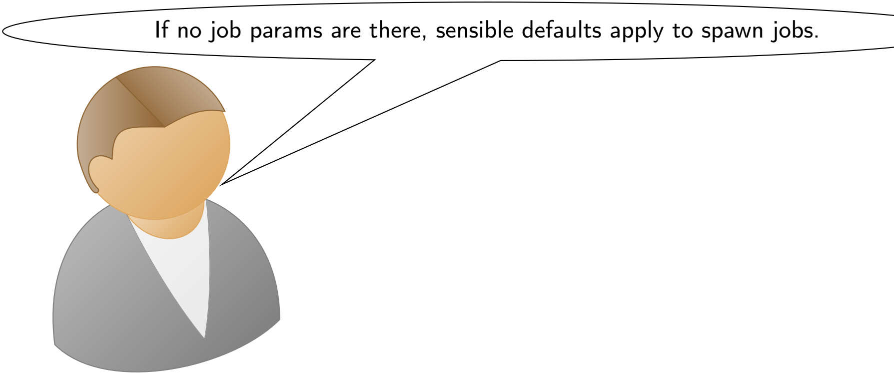

And Then Some

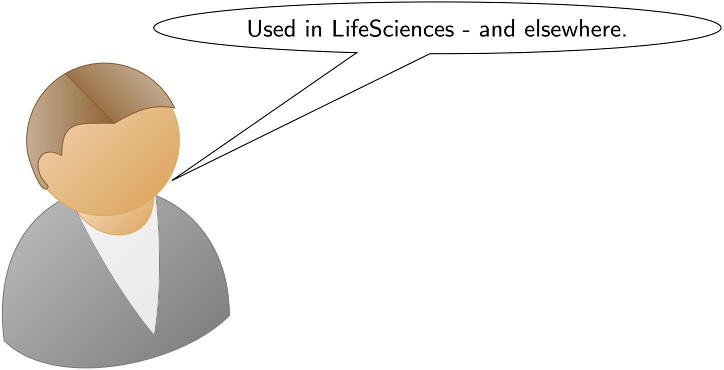

And Then Some

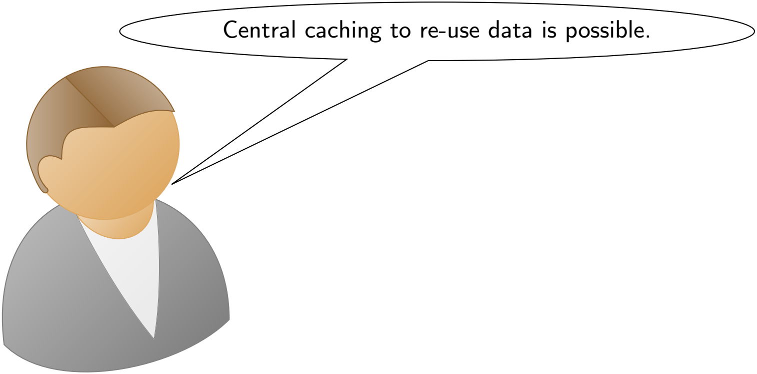

And Then Some

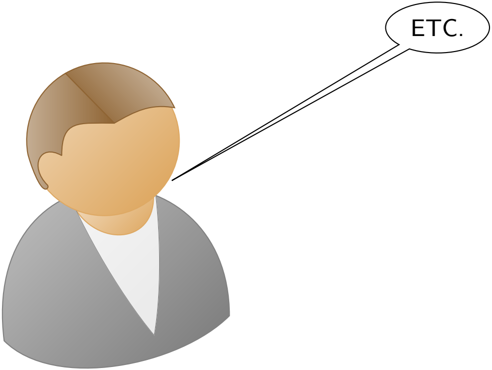

Snakemake goes HPC

Conclusion

- Learning about workflow managers ought to be an integral part to data analytics curricula.
- (Some) HPC centers get used to LifeScience Users.

Conclusion

- Learning about workflow managers ought to be an integral part to data analytics curricula.
- (Some) HPC centers get used to LifeScience Users. Or do not attract users which others do.
- Workflow Management Systems need to collect parametrization examples for common workflows - to speed up the deployment process.

Honour where Honour is due

- Johannes Köster
- Lukas Hellmann
- Malte Petersen
- Fabian Brand
- Florian Boecker
- Aasish Kumar Sharma

Apologies go out to

- Ellis Crawford
- Veronica Fillinger

the only two Scots I am scientifically affiliated with.