



UNIVERSITY OF
CAMBRIDGE

Yusuf Hamied
Department of Chemistry

virtually

How to do everything with UKCA

Luke Abraham
NCAS Cambridge

`n.luke.abraham@ncas.ac.uk`




- What is the Unified Model, and how is code developed for it?
- What do I mean by “*virtually*”?
- What do I mean by “*virtually everything*”?
- What do I mean by “*everything*”?

What is the Met Office Unified Model?


Atmospheric Modelling: integrating our knowledge of atmospheric behaviour forward in time

short-wave radiation



vegetation

Newton's laws for rotating fluid
Gas laws
Laws of thermodynamics



Surface Processes

- The challenge: To reproduce the behaviour of (hazardous) weather systems

Unified Model

Brown et al. (2013)

- Operational forecasts

- Mesoscale (resolution approx. 4km, 1.5km)
- Global scale (resolution approx. 17km)

- Seasonal predictions

- Resolution approx. 60km

- Global and regional climate predictions

- Resolution around 120km
- Run for 10-100-... years

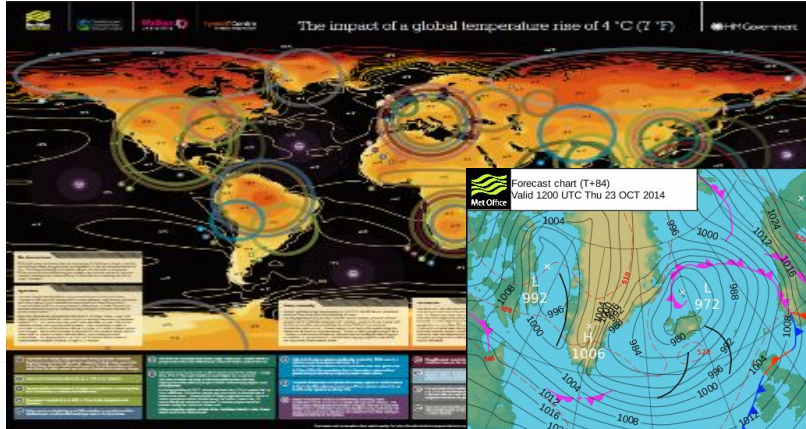
- Research mode

- Resolution 1km - 10m

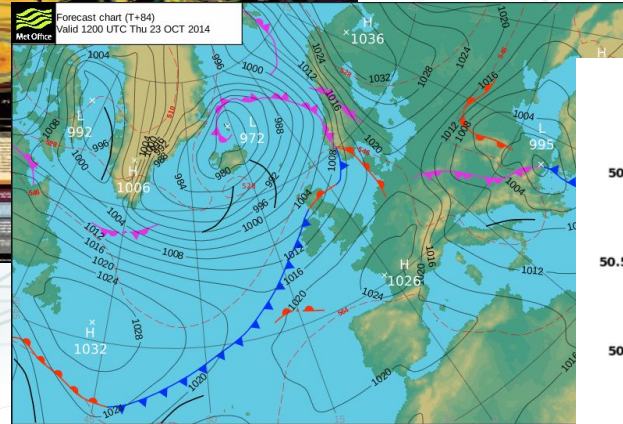
> 25 years old

The consequence of unification

300 km




30 km



A factor of 1000
between these








300 m



...the same scheme has
to continue to work

Exoplanets

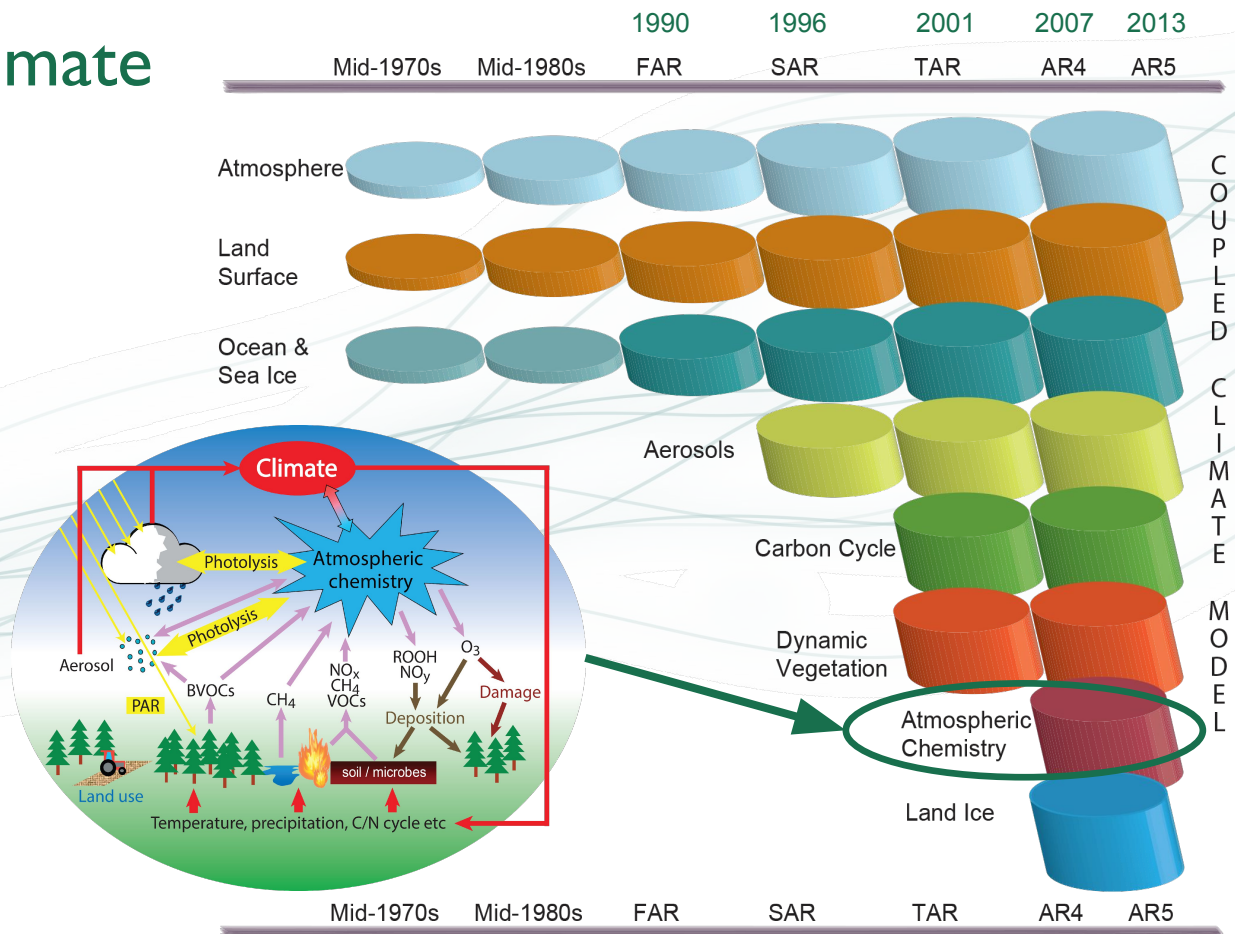
The unified model, a fully-compressible, non-hydrostatic, deep atmosphere global circulation model, applied to hot Jupiters
Mayne *et al.* 2014



Exploring the climate of Proxima B with the Met Office Unified Model
Boutle *et al.* 2017

Development of Climate Models

Increasing the **number** and **complexity** of processes represented in climate models.

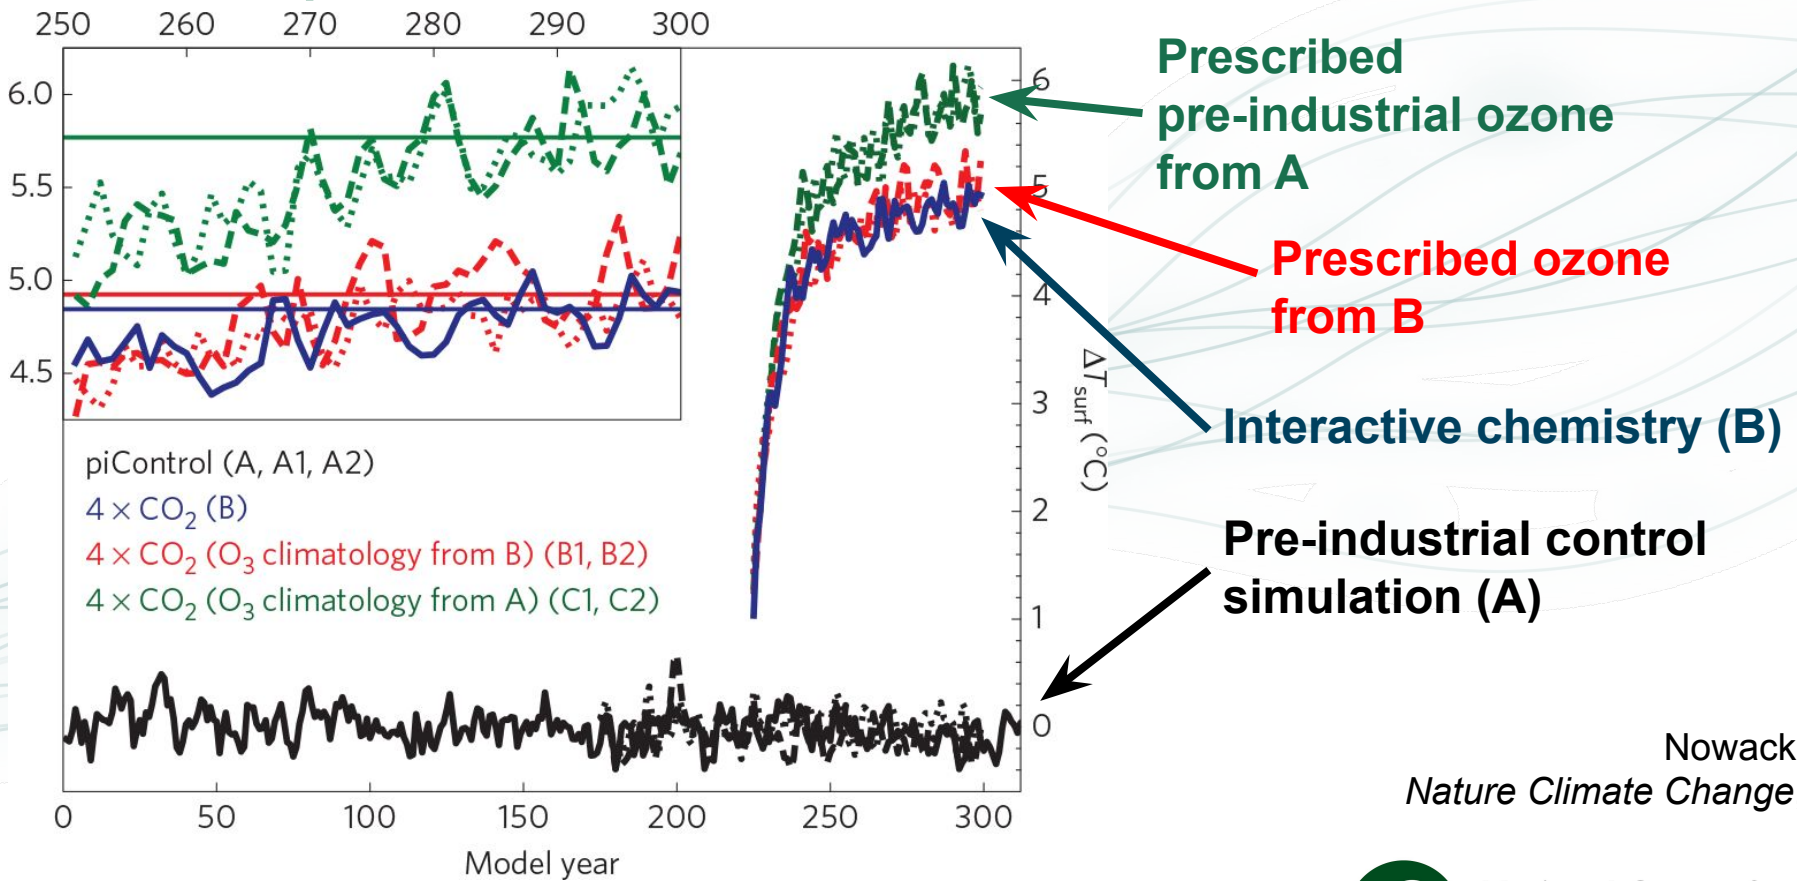


IPCC, 2013

United Kingdom Chemistry and Aerosols

- In Cambridge we work on a part of the UM called the *United Kingdom Chemistry and Aerosols* sub-model, or **UKCA**.
- We develop new chemistry schemes for the model, with a focus on Chemistry-Climate Interactions.
- Here we usually include between 75-240 transported chemical species and hundreds of reactions to allow the model to accurately simulate changes to radiatively important gases such as ozone and methane, which can feed-back on the climate system.

Chemistry-Climate Interactions

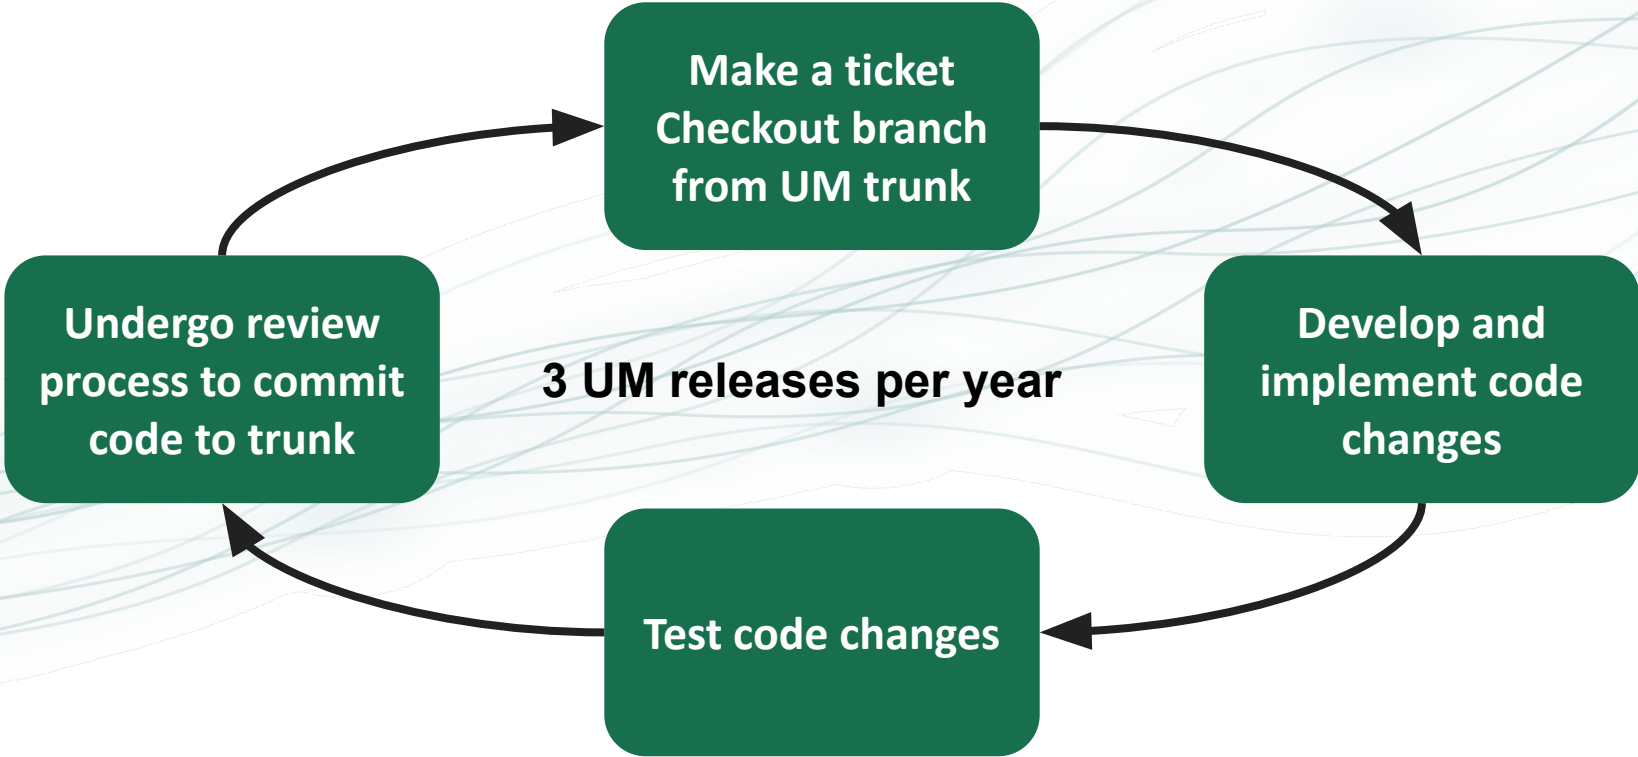


Nowack *et al.*,
Nature Climate Change, 2014

Developing code for the Unified Model

- Approximately 900,000 lines of code (mainly Fortran 2003).
- Over 200 active developers.
- Uses the *Rose* graphical namelist editor and the *Cylc* workflow engine, with the code held in subversion repositories, managed using *FCM* (“Flexible Configuration Management” - mostly a wrapper around subversion).
- Majority of the development work is done by the Met Office.

Met Office Unified Model code development process

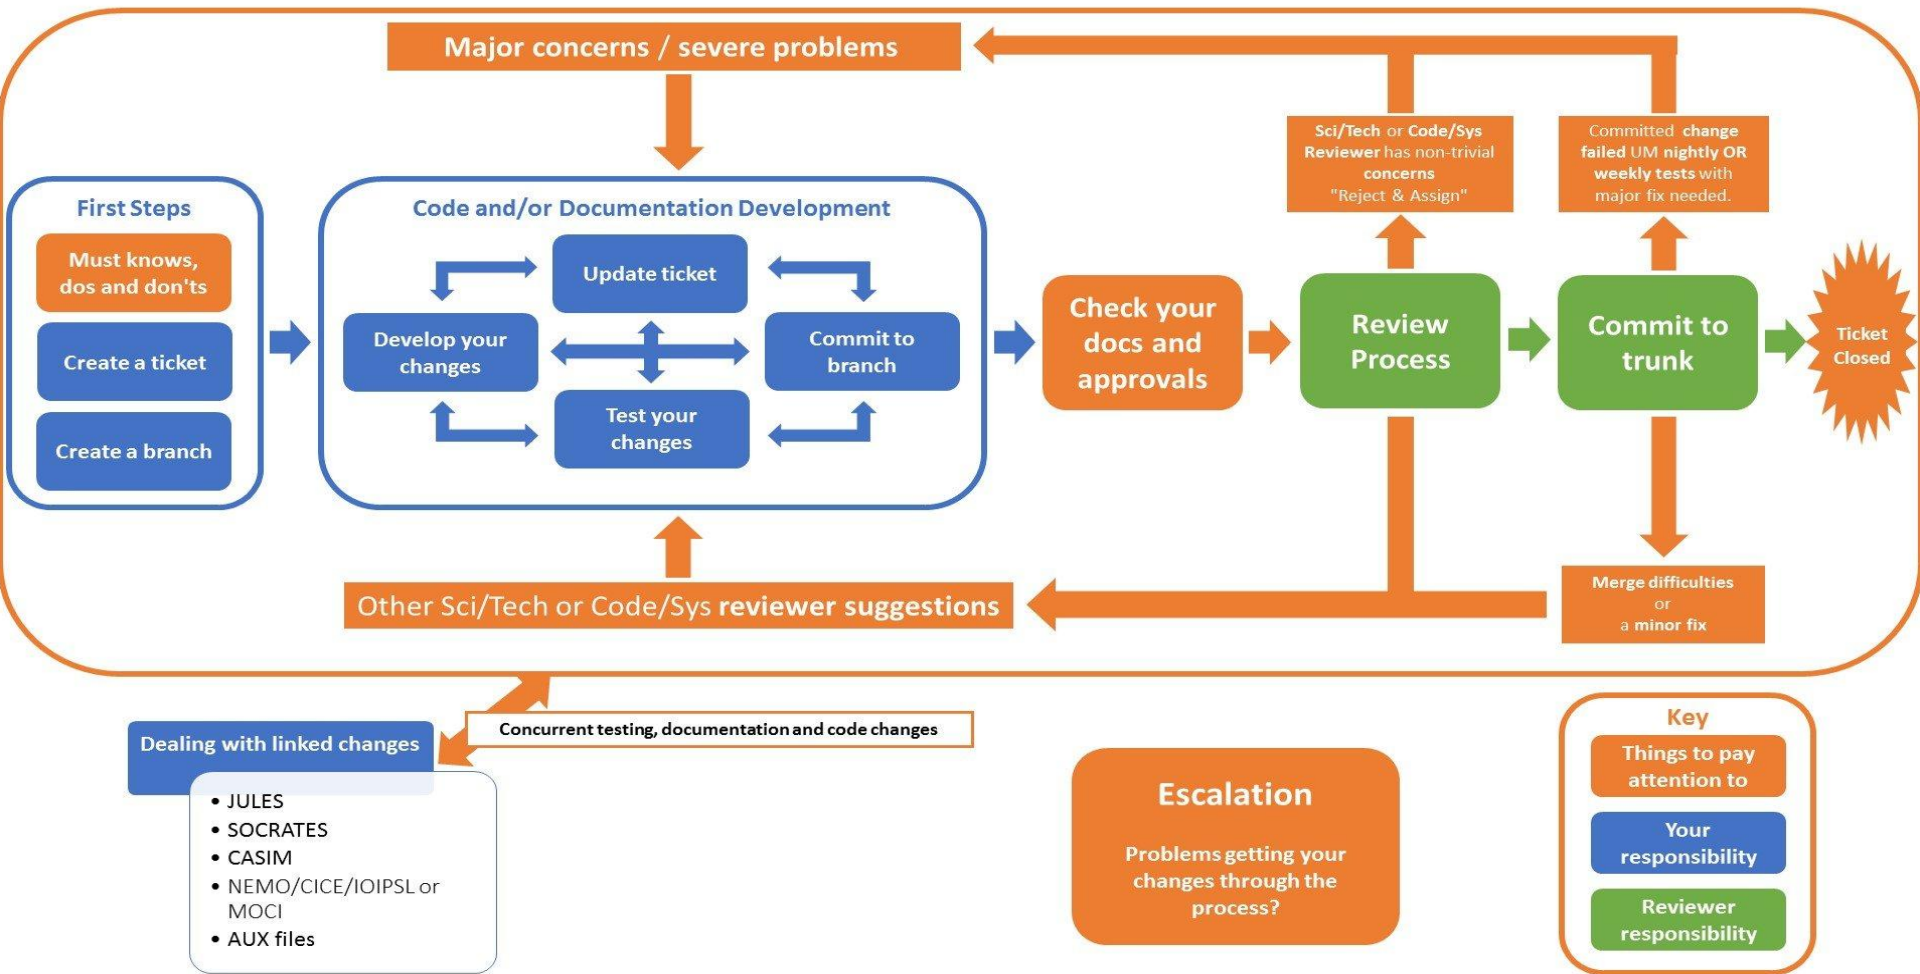


Similar process for science improvements to be included in operational configurations.



National Centre for
Atmospheric Science
NATURAL ENVIRONMENT RESEARCH COUNCIL

UM Development Working Practices



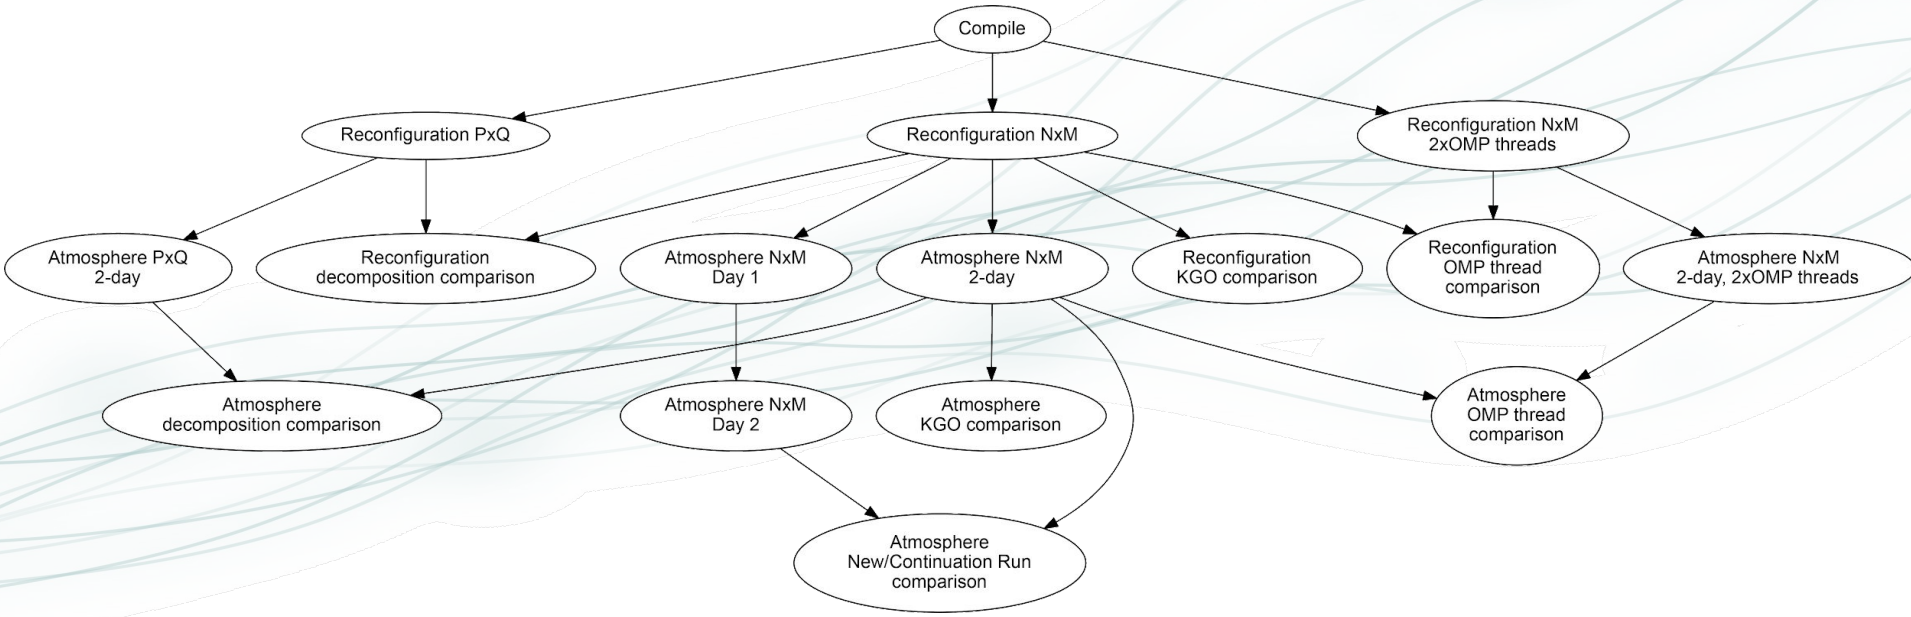
Code development process

- Science changes often require testing with long simulations that will take several weeks or more to run.
- Diagnostics are then run through assessment and validation tools to produce many plots of standard metrics.
- All code changes are tested using the *rose-stem* utility, using a set of standard tests that compare against “**Known Good Output**” (KGO).
 - e.g. a climate resolution UKCA testing job takes 7 minutes to run 2 model days on 144 cores, requiring 112GB of memory.
- **You must be able to show that your change works when turned on and doesn't break anything when turned off.**

rose-stem

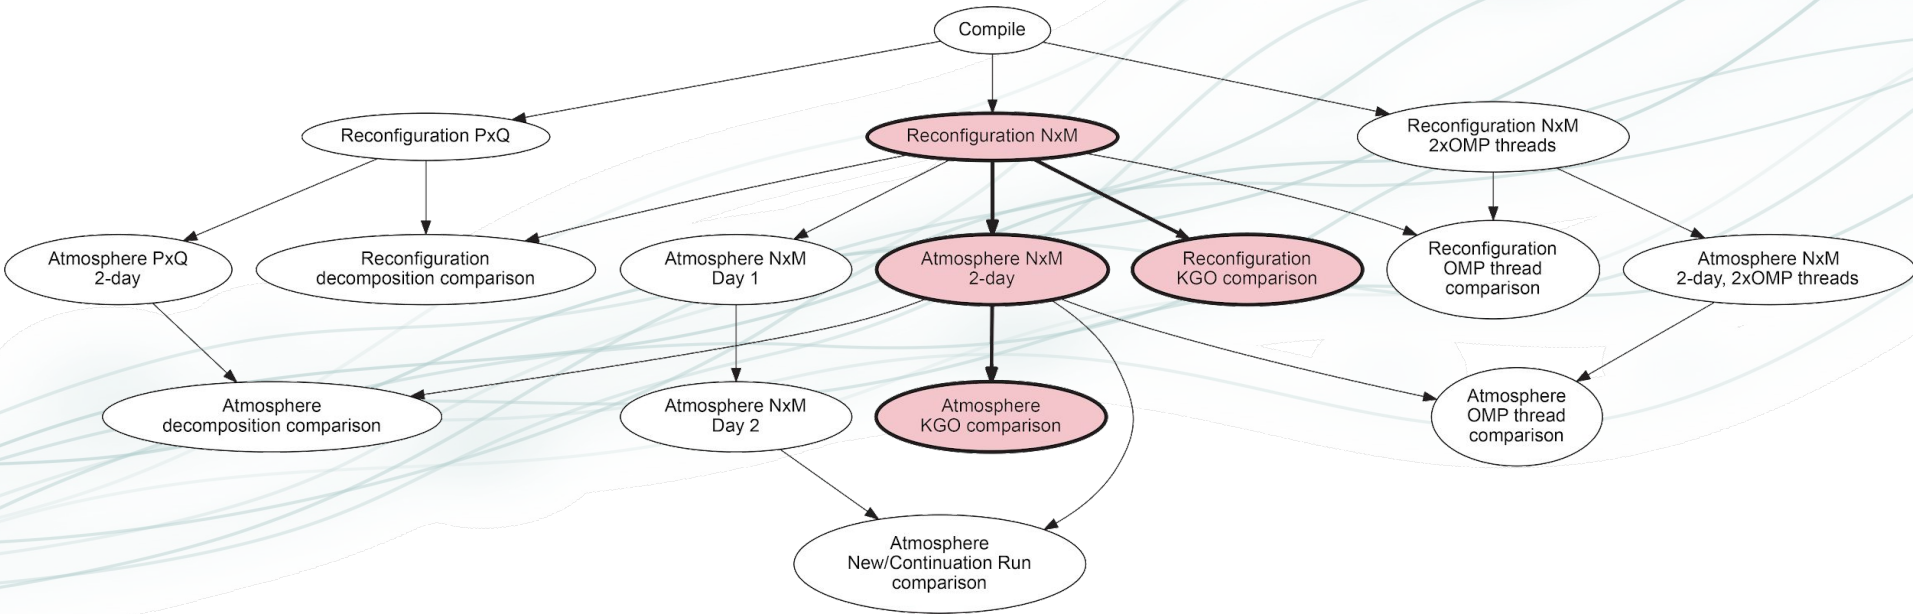
- At UM vn12.2 there are 328 UM testing jobs (inc. 33 UKCA jobs):
 - HPC jobs with Cray, GCC/Intel & GNU compilers
 - Linux jobs GCC/Intel, GCC/PGI, Clang/Intel, & GNU compilers
- Additional restart file creation tests and further tests for code standards, metadata, utilities, creation of boundary conditions, etc.
- Tests include KGO, restartability, OpenMP, & processor decomposition tests, with a range of optimisation levels:
 - “high”, “fast”, “safe”, “debug”, “rigorous”.
- If a KGO test “breaks” you need to explain why and get approval to change the saved output.
- **Tests protect your code and scientific configurations from being broken by another change.**

rose-stem - Met Office testing framework



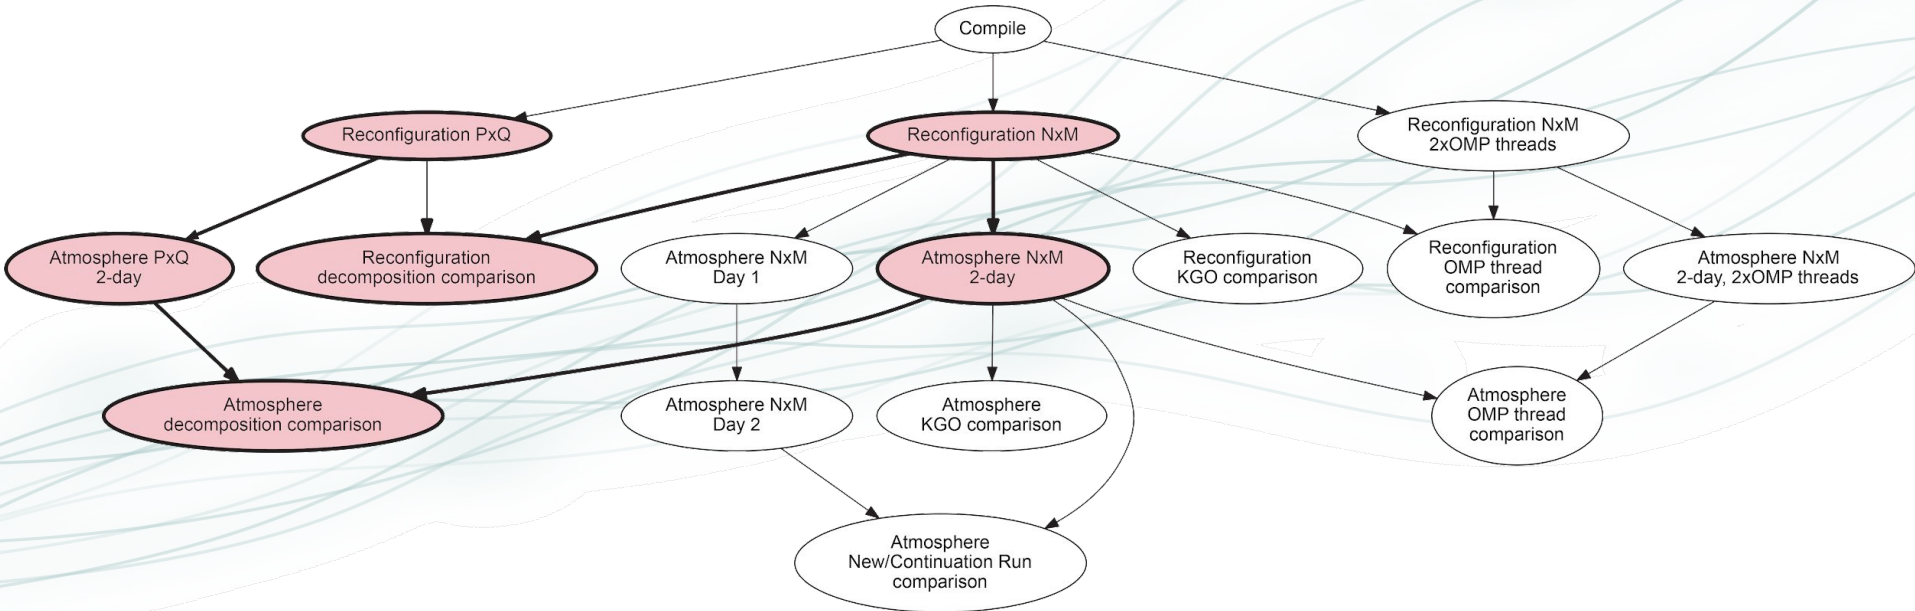
Done with different levels of compiler optimisation.

rose-stem - “Known Good Output” or KGO tests



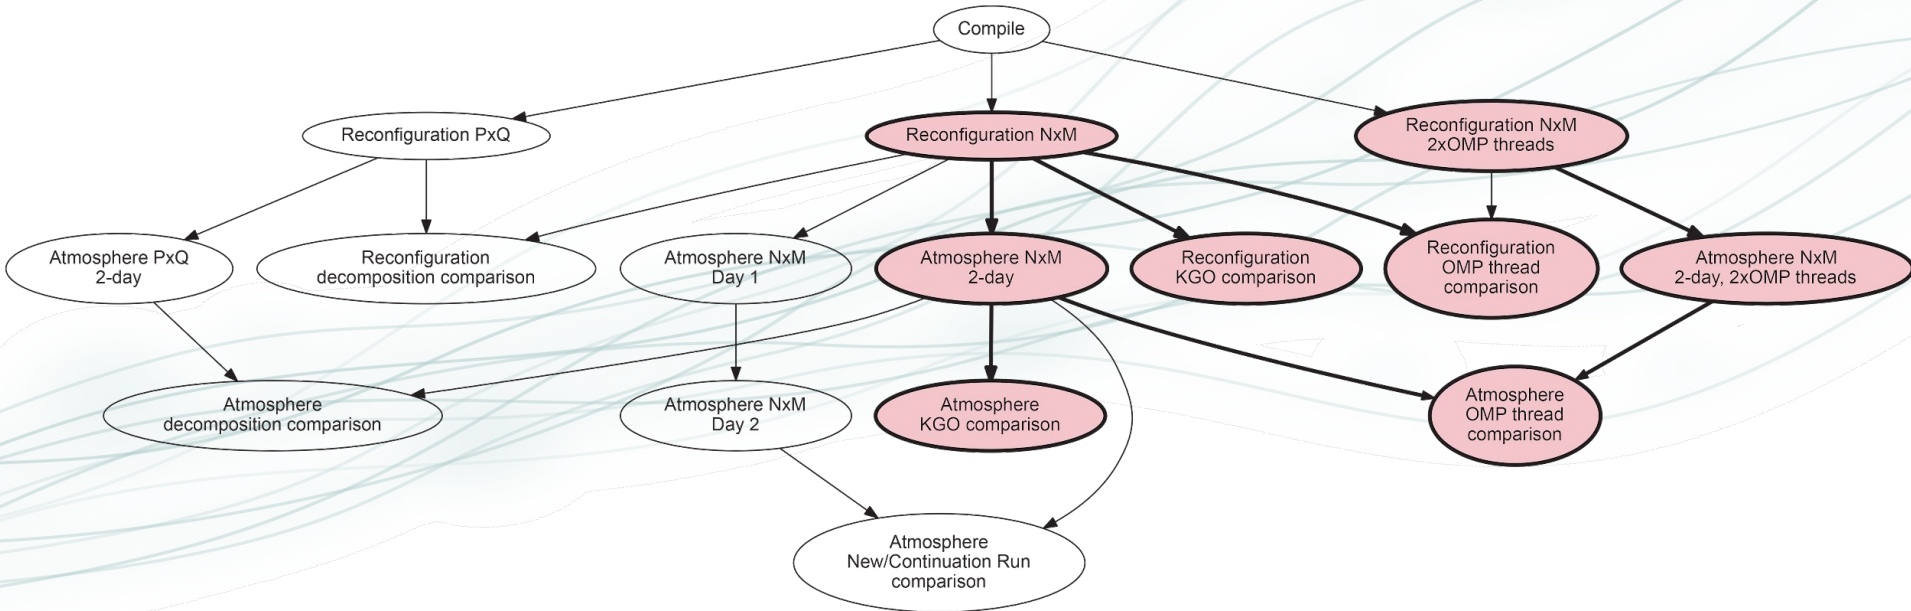
KGO will likely be different with different levels of compiler optimisation.

rose-stem - processor decomposition tests



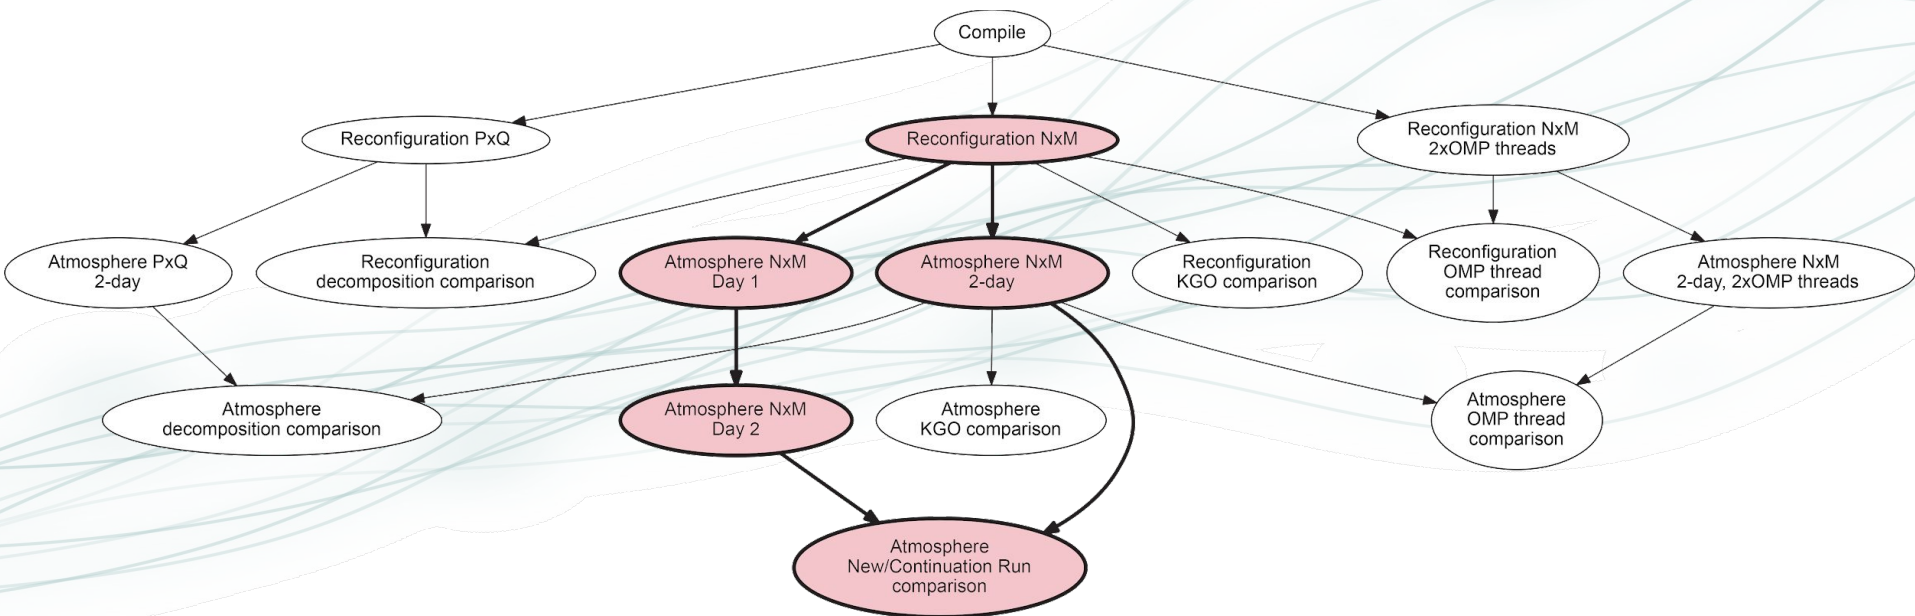
Some changes may break KGO but should still pass the processor decomposition tests.

rose-stem - OpenMP tests



KGO should be identical with or without OpenMP, although the OMP test should still pass even if the KGO one does not.

rose-stem - restart tests



Some changes may break KGO but should still pass the restartability tests.

rose stem --group=developer,ukca

The screenshot shows a terminal window with a task list on the left and a network diagram on the right. The task list is as follows:

task	state	host	job system	job ID	T-submit
1	running				
EXTRACT_SOURCE	running				
METO_LINUX_BUILD_SETTINGS	succeeded				
METO_LINUX_BUILD	succeeded				
METO_XC40_BUILD	running				
INSTALL	succeeded				
HOUSEKEEPING	waiting				
ROSE_ANA_COMPARISON	running				
METO_LINUX_ROSE_ANA	waiting				
METO_XC40_ROSE_ANA_COMPARISON	running				
ROSE_ANA_WALLCLOCK	waiting				
METO_XC40_ROSE_ANA_WALLCLOCK_COMPARISON	waiting				
DESKTOP	succeeded				
METO_LINUX_AQUIM_EG	succeeded				
METO_LINUX_AQUIM_EG_GLOMAP	running				
METO_LINUX_N48_GA7_AMIP_12HR	succeeded				
METO_LINUX_N48_GA7_AMIP_NAMING	succeeded				
METO_LINUX_N48_GA7_AMIP_12HR_COMP_CHECK	running				
METO_LINUX_N48_GA7_AMIP_DEV_12HR	succeeded				
METO_LINUX_N48_EG_OMP	succeeded				
METO_LINUX_N48_EG_NOOMP	succeeded				
METO_LINUX_SCM_TOGACOARE_GA6_OMP	succeeded				
METO_LINUX_SCM_GABLS3_GA6_OMP	succeeded				
METO_XC40	succeeded				
METO_XC40_INTEL_HASWELL	succeeded				
METO_XC40_GNU_HASWELL	succeeded				
METO_XC40_UKCA_EG_STRATTROP	running				
METO_XC40_UKCA_NUJGED	running				
METO_XC40_N48_GA7_AMIP_2DAY	running				
METO_XC40_N48_GA7_AMIP_NAMING	running				
XC40_GA7_AMIP_NAMING_CRUN_INSTALL	waiting				
XC40_GA7_AMIP_NAMING_CRUN_ARCHIVE	waiting				
METO_XC40_N48_GA7_AMIP_2DAY_COMP_CHECK	waiting				
METO_XC40_N48_GA7_AMIP_10DAY	running				
METO_XC40_N48_GA7_AMIP_30DAY	submitted				
METO_XC40_N48_GA7_AMIP_DEV_2DAY	running				
METO_XC40_N48_EG_OMP	running				
METO_XC40_N48_EG_OMP_IFORT	succeeded				

The network diagram on the right shows a complex web of nodes and connections, representing the dependencies and execution flow of the tasks. The nodes are colored in shades of green and grey, and the connections are thin lines. The diagram is organized into several distinct clusters, with some nodes highlighted in a darker green. The overall structure is dense and interconnected, showing a complex dependency graph.

running to stop at 1 (filtered: 1) live

2017-03-17T13:35:13Z

Met Office Infrastructure (2021)

Met Office Science Repository Service

um jules socrates gcom doc roses-u

MOSRS mirror

Linux Virtual Desktops
Rose GUI

Cylc server

Cylc server

Linux Servers
Intel
GNU

Cray XC40

Cray
Intel
GNU

XCE & XCF: 89,856 cores each (#151 & #152 Top500)
XCS: 241,920 cores (#58 Top500)

MASS Tape Archive

Outside the Met Office

But what if you don't have access to the Met Office infrastructure?

What can UK university students and researchers access?

- You should be able to get access to the UM source code and the tools to run the UM.
- You *may* be able to get access to HPC resources, e.g. ARCHER2 or Monsoon2 etc.
- You *may* be able to get access to an analysis platform, e.g. JASMIN.
- Access to rose-stem and the internal Met Office tools could be limited or not available at all.

What do I mean by “virtually”?

What do I mean by “*virtually*”?

- A Virtual Machine configuration has been developed to allow people to easily use FCM, Rose, & Cylc.
- Uses Vagrant and VirtualBox (for a PC or server) or AWS.
- UM Systems Team have set-up running the UM in an Ubuntu guest image.

Met Office Virtual Machine on GitHub:

<https://github.com/metomi/metomi-vms>

JASMIN



Uses for the VM

1. Developing & testing code changes

- Faster turnaround time
 - Approximately 10 minutes to compile from scratch, with only a few seconds to recompile.
 - Small versions of standard jobs run in a few minutes on 2 cores.


2. Training

What do I mean by “*virtually*”?

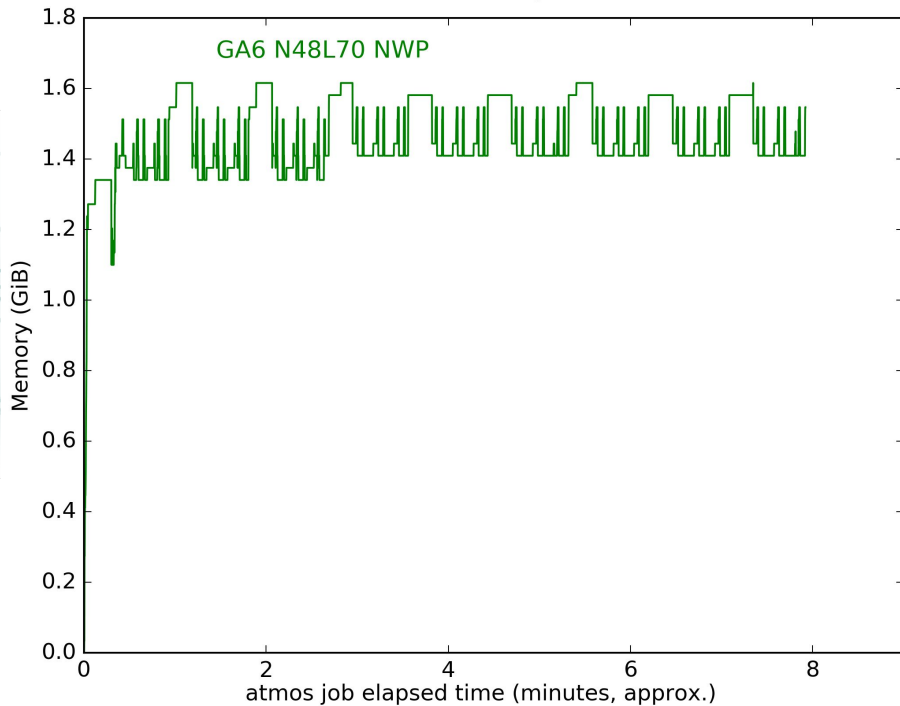
- To run a basic UM test suite, the VM needs around 3GB of memory.
- Initial thoughts were that the main issue with running UKCA would (*probably!*) be the memory requirement.
- A standard UKCA job requires an additional 137 3D fields to be added to the restart file, along with further temporary arrays allocated at run-time and any extra required for diagnostic output.
- Usually require at least 112GB of memory at climate resolution (192x144x85 grid points).

Memory usage for um-atmos job

Basic test suite uses ~1.6GiB



vn10.5 vm-x86-gnu Ubuntu 15.10 VM
1x1 safe w/o OpenMP




GA6 job runs for 9 model hours
- 27 timesteps

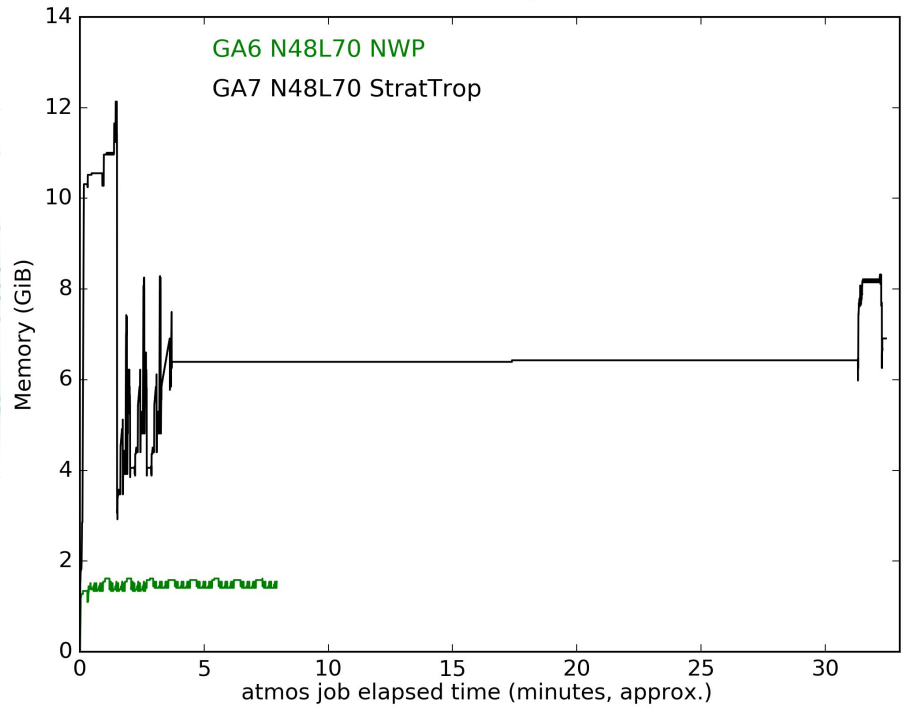
Memory usage for um-atmos job

Basic test suite uses ~1.6GiB

Low resolution UKCA suite uses ~12.2GiB



vn10.5 vm-x86-gnu Ubuntu 15.10 VM
1x1 safe w/o OpenMP



GA6 job runs for 9 model hours
- 27 timesteps
GA7 jobs run for 1 model hour
- 3 timesteps for L70


Memory usage for um-atmos job

Basic test suite uses ~1.6GiB

Low resolution UKCA suite uses ~12.2GiB

Low resolution UKCA suite (low top) uses ~7.4GiB

vn10.5 vm-x86-gnu Ubuntu 15.10 VM
1x1 safe w/o OpenMP



GA6 job runs for 9 model hours
- 27 timesteps
GA7 jobs run for 1 model hour
- 3 timesteps for L70
- 2 timesteps for L38



Memory usage for um-atmos job

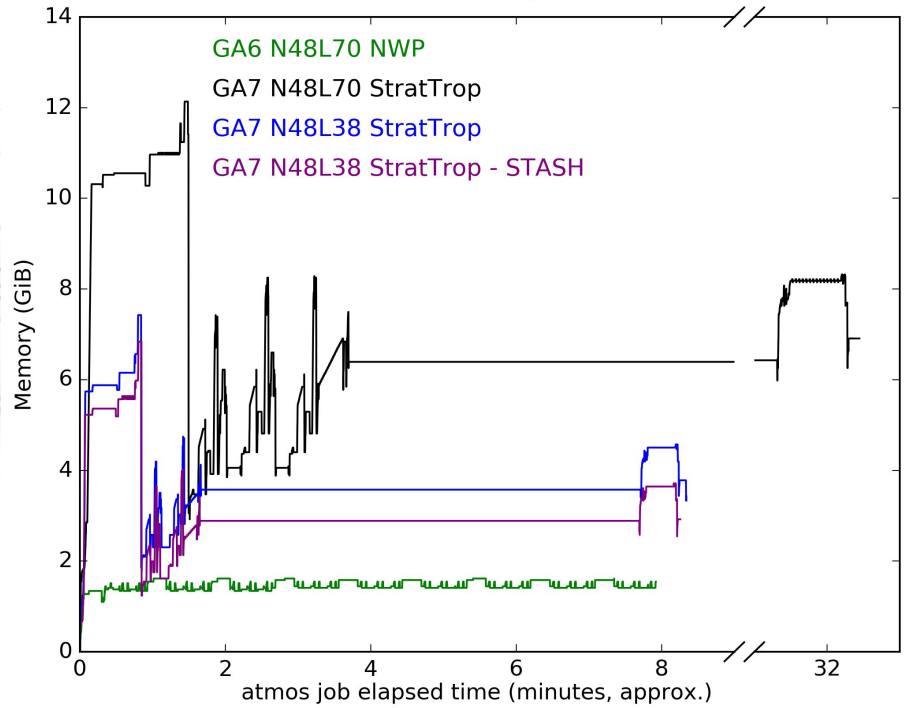
Basic test suite uses ~1.6GiB

Low resolution UKCA suite uses ~12.2GiB

Low resolution UKCA suite (low top) uses ~7.4GiB

Low resolution UKCA suite (low top) uses ~6.9GiB when all diagnostic output requests are removed

vn10.5 vm-x86-gnu Ubuntu 15.10 VM
1x1 safe w/o OpenMP



GA6 job runs for 9 model hours
- 27 timesteps
GA7 jobs run for 1 model hour
- 3 timesteps for L70
- 2 timesteps for L38



Memory usage for um-atmos job

Basic test suite uses ~1.6GiB

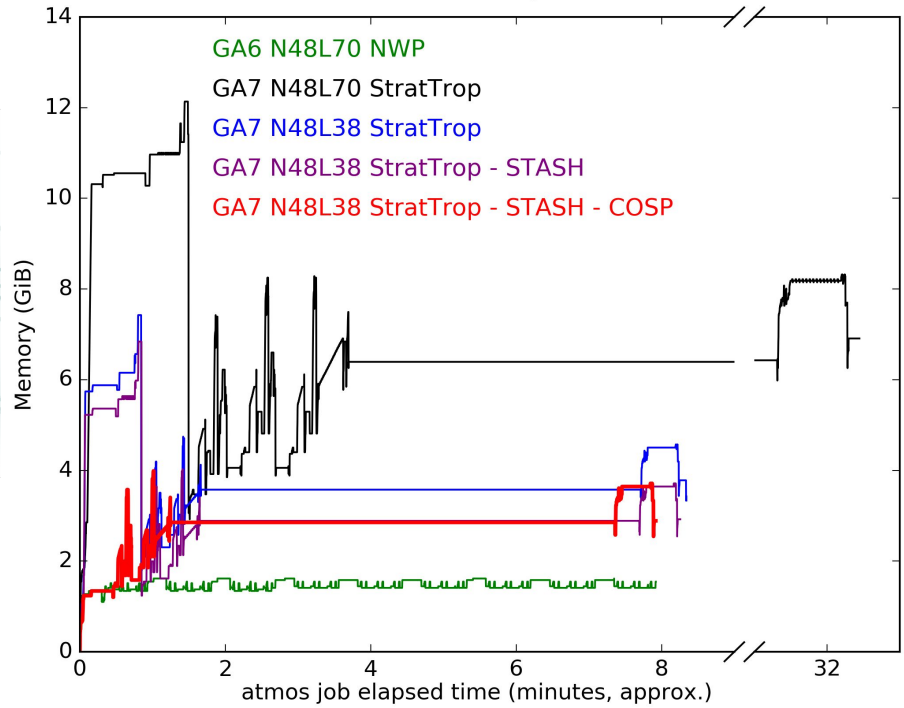
Low resolution UKCA suite uses ~12.2GiB

Low resolution UKCA suite (low top) uses ~7.4GiB

Low resolution UKCA suite (low top) uses ~6.9GiB when all diagnostic output requests are removed

Low resolution UKCA suite (low top) uses ~4.0GiB when all diagnostic output requests are removed and the COSP (CFMIP Observation Simulator) diagnostics package is not turned on

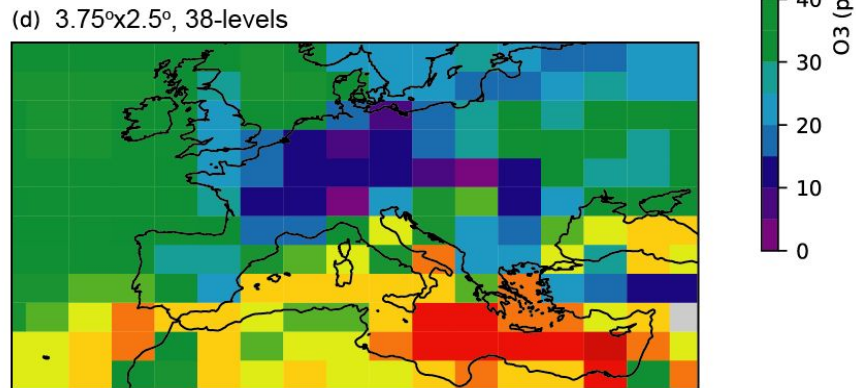
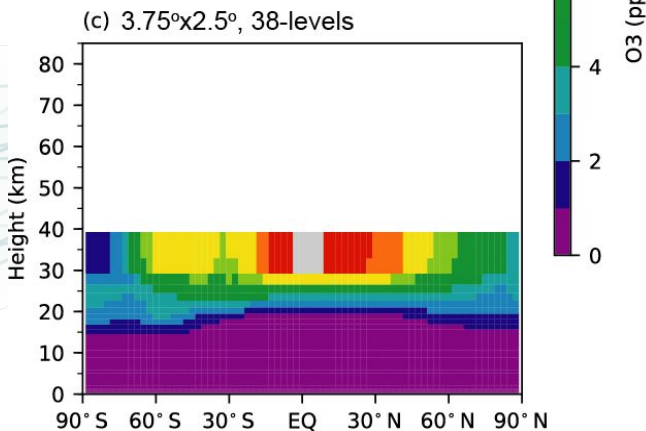
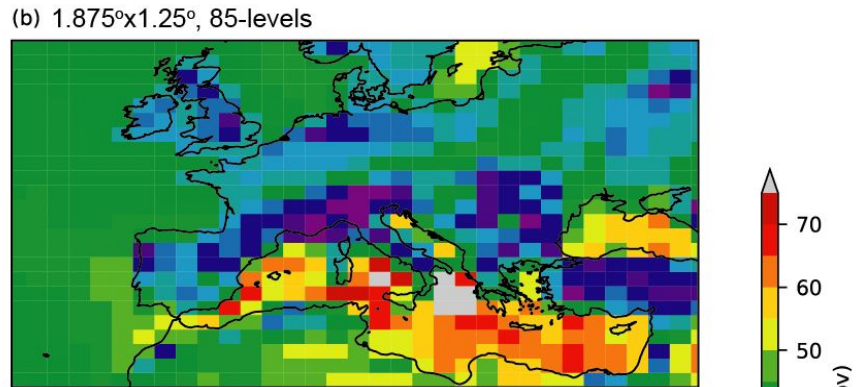
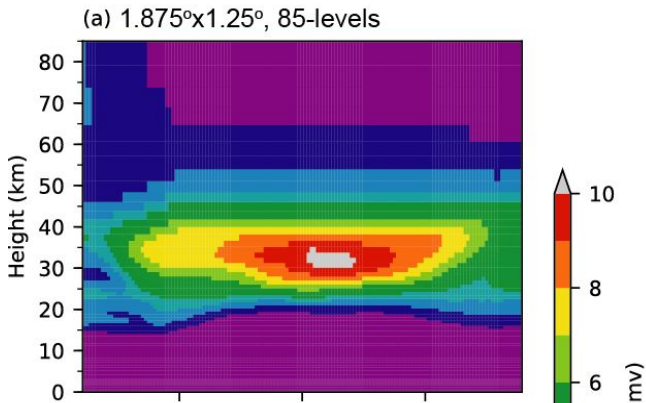
vn10.5 vm-x86-gnu Ubuntu 15.10 VM
1x1 safe w/o OpenMP



GA6 job runs for 9 model hours
- 27 timesteps
GA7 jobs run for 1 model hour
- 3 timesteps for L70
- 2 timesteps for L38



**Climate
resolution
192x144x85
grid points**



**VM-suitable
resolution
96x72x38
grid points**

What do I mean by
“virtually everything”?

Helpful(?!) Errors

- However, there seemed to be an issue with using the GNU compiler, as the code failed with the following error on timestep 3:

? Error from routine: EG_BICGSTAB_MIXED_PREC

? Error message: NaNs in error term in BiCGstab

? **This is a common point for the model to fail if it**
? **has ingested or developed NaNs or infinities**
? **elsewhere in the code.**

? See the following URL for more information:

? <https://code.metoffice.gov.uk/trac/um/wiki/KnownUMFailurePoints>

- After some debugging, the issue *appeared* to be with the UKCA routine `asad_hetero.F90`, which couples the chemistry and aerosol schemes.

Precision

The UM uses compiler flags to make **REALS** double-precision, so rather than using

```
REAL (KIND=8) :: a
```

or

```
INTEGER, PARAMETER :: dp=SELECTED_REAL_KIND(15,300)
```

```
REAL (KIND=dp) :: a
```

the UM will use

```
-r8 (Intel), -s default64 (Cray), -fdefault-real-8 (GNU)
```

AND there is still the occasional `1.0d0` statement.

But what about TINY (1.0d0)?

- Intel (-r8):

```
tiny(1.0) = 2.225073858507201E-308
```

```
tiny(1.0d0) = 2.225073858507201E-308
```

- Cray (-s default64)

```
tiny(1.0) = 2.22507385850720138E-308
```

```
tiny(1.0d0) = 2.22507385850720138E-308
```

But what about TINY (1.0d0)?

- Intel (-r8):

`tiny(1.0) = 2.225073858507201E-308`

`tiny(1.0d0) = 2.225073858507201E-308`

- Cray (-s default64)

`tiny(1.0) = 2.22507385850720138E-308`

`tiny(1.0d0) = 2.22507385850720138E-308`

- GNU (-fdefault-real-8)

`tiny(1.0) = 2.2250738585072014E-308`

`tiny(1.0d0) = 3.36210314311209350626267781732175260E-4932`

Effectively `a = 1.0d0 = 1.0`

but `a = tiny(1.0d0) = 0.0`

But what about TINY (1.0d0)?

- Intel (-r8):

`tiny(1.0) = 2.225073858507201E-308`

`tiny(1.0d0) = 2.225073858507201E-308`

- Cray (-s default64)

`tiny(1.0) = 2.22507385850720138E-308`

`tiny(1.0d0) = 2.22507385850720138E-308`

- GNU (-fdefault-real-8)

`tiny(1.0) = 2.2250738585072014E-308`

`tiny(1.0d0) = 3.36210314311209350626267781732175260E-4932`

- GNU (-fdefault-real-8 -fdefault-double-8)

`tiny(1.0) = 2.2250738585072014E-308`

`tiny(1.0d0) = 2.2250738585072014E-308`

What do I mean by “everything”?

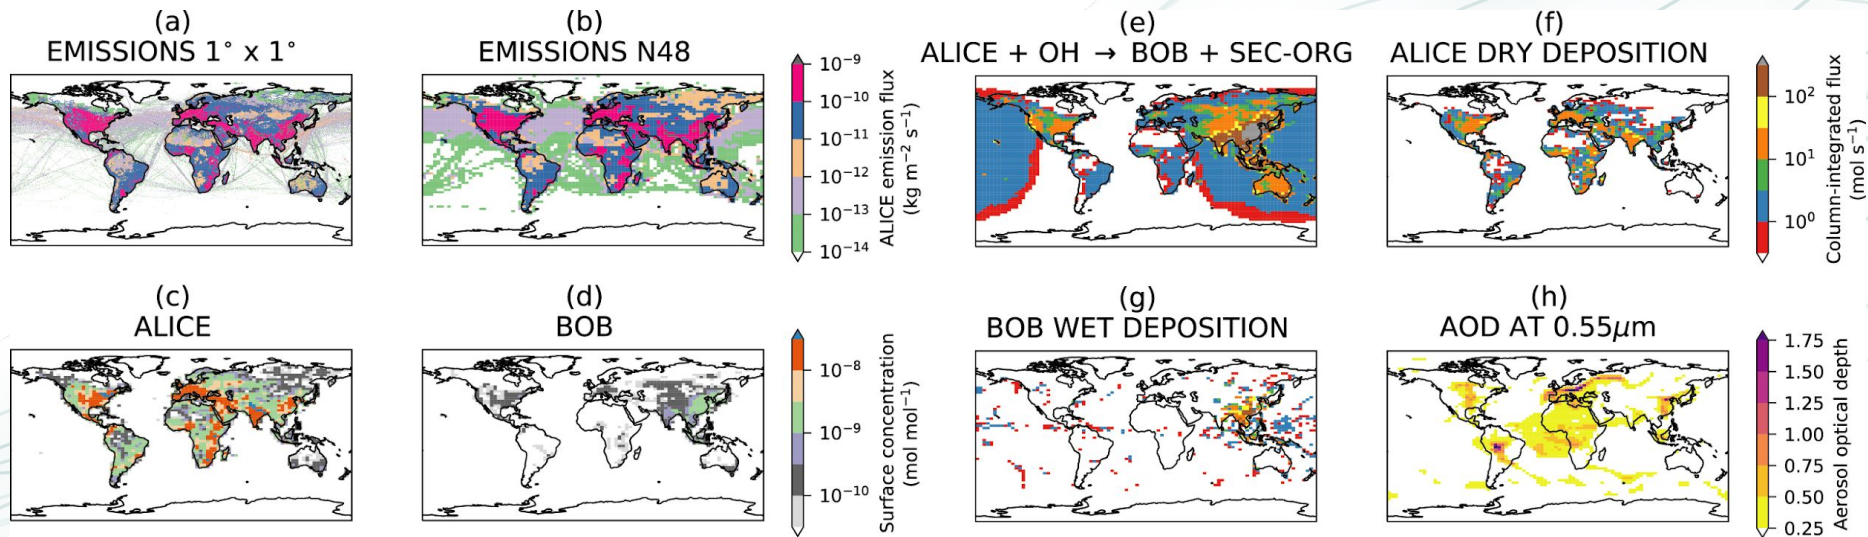
Rose stem (again!)

- These compiler flag changes went into the UM trunk at vn10.7.
- There are *rose-stem* testing jobs for several UM configurations (now also including UKCA) on the VM.
- Equivalent low-resolution UKCA jobs have also been added to the Met Office HPC tests.
 - One is included as part of the standard *developer* group that must be used when making any change.

Training

- What if we could use the VM for training?
- This would mean:
 1. Everyone runs on their own dedicated computer.
 - Simpler set-up than for current production runs.
 2. Researchers doing the tutorials can easily do the training tasks without needing a supercomputer account.
 - “Try before you buy”

UKCA Training



Conclusions

Conclusions

- Testing environments, as well as production environments, are required to be able to develop code changes in a timely manner.
- A Virtual Machine environment has been developed by the Met Office that easily allows users to install their own copy of the UM.
- Standard tests have also been implemented on this system that can be used when developing new code changes.
- This creates a standard system that all non-Met Office developers can use to quickly implement and test changes, prior to running on a HPC.



Using a virtual machine environment for developing, testing, and training for the UM-UKCA composition-climate model, using Unified Model version 10.9 and above

Nathan Luke Abraham^{1,2}, Alexander T. Archibald^{1,2}, Paul Cresswell³, Sam Cusworth³, Mohit Dalvi³, David Matthews³, Steven Wardle³, and Stuart Whitehouse³

¹Department of Chemistry, University of Cambridge, Cambridge, CB2 1EW, UK


²National Centre for Atmospheric Science, UK

³Met Office, FitzRoy Road, Exeter, EX1 3PB, UK

Correspondence: Nathan Luke Abraham (luke.abraham@atm.ch.cam.ac.uk)

Received: 8 May 2018 – Discussion started: 18 May 2018

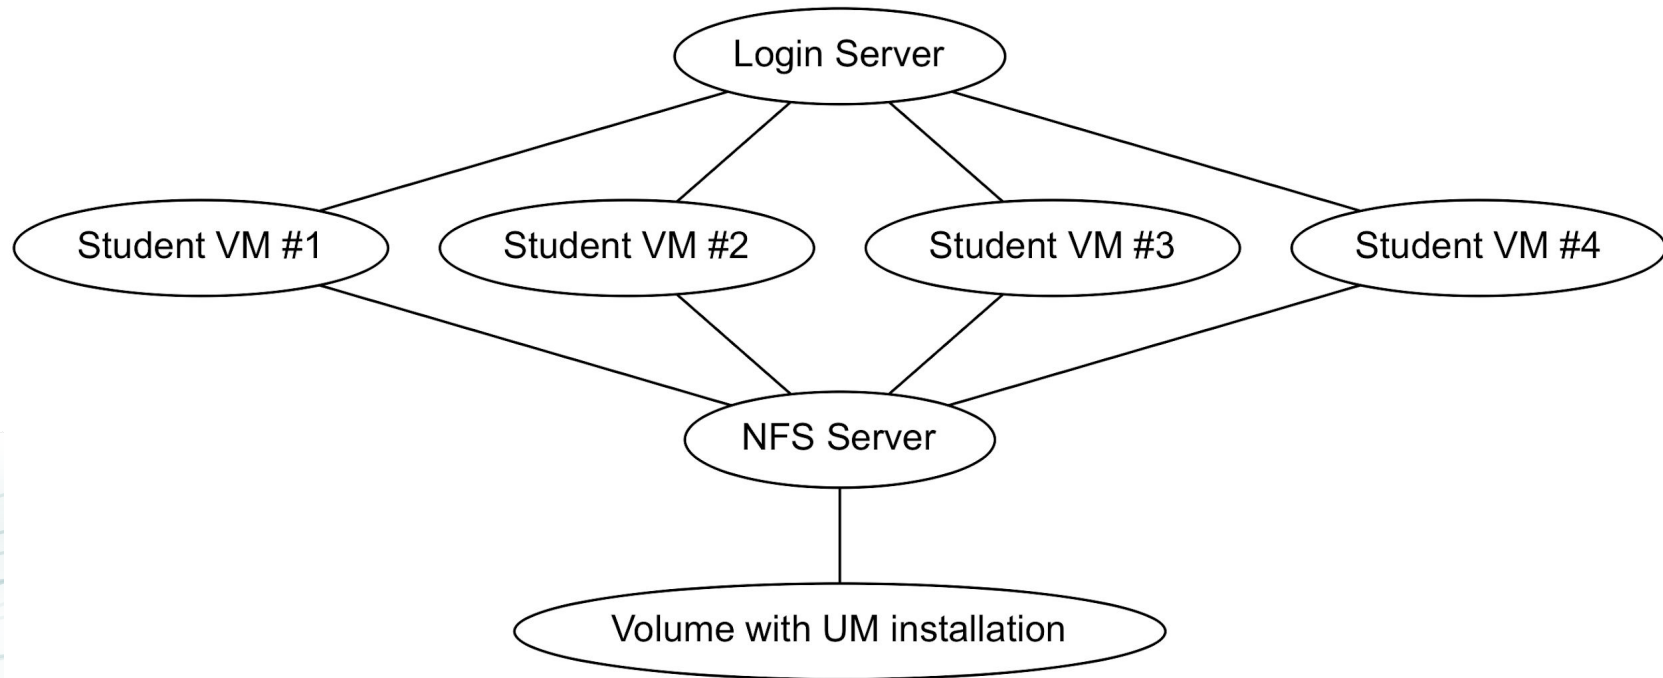
Revised: 10 August 2018 – Accepted: 23 August 2018 – Published: 6 September 2018



Step	ARCHER (XC30)	XCS-C (XC40)	Virtual Machine
Cray <code>cce</code> initial compile	34 minutes	15 minutes	-
Cray <code>cce</code> incremental compile	5-7 minutes	3 minutes	-
Intel <code>ifort</code> initial compile	19 minutes	9 minutes	-
Intel <code>ifort</code> incremental compile	6-7 minutes	1 minute	-
GNU <code>gfortran</code> initial compile	-	4 minutes	8-10 minutes
GNU <code>gfortran</code> incremental compile	-	2-3 minutes	45 seconds
Reconfiguration task, used to produce the initial conditions file.	3-4 minutes (Intel, 6×4)	15-30 seconds (GNU, 4×9)	25-30 seconds (GNU, 1×2)
Atmosphere task	40 seconds (Intel, 6×4)	45 seconds (GNU, 4×9)	12 minutes (GNU, 1×2)

Compiler settings	Safe	Safe	Rigorous
Number of UM OpenMP threads	0	2	2
Approximate run-time on 2-core VM (1×2) (minutes)	8	11	29
Approximate run-time on 4-core VM (1×2) (minutes)	8	4	17
Approximate run-time on 4-core VM (1×4) (minutes)	5	8	22
Approximate run-time on 8-core VM (1×4) (minutes)	5	3	14
Approximate run-time on 8-core VM (1×8) (minutes)	3	6	26
Approximate run-time on 16-core VM (1×8) (minutes)	3	3	22

Compile type	GNU gfortran compiler flags	Number of UM OpenMP threads	Total VM memory required (GB)
safe	<code>-O2 -Werror</code>	0	6
safe	<code>-O2 -Werror -fopenmp</code>	2	6
rigorous	<code>-O0 -Wall -ffpe-trap=invalid,zero -fbounds-check -Warray-bounds -fcheck-array-temporaries -finit-real=nan -fimplicit-none -fopenmp</code>	2	8



Ansible playbooks for this system are available via GitHub:
<https://github.com/theabro/ukca-playbook>

Students could connect to their VM using a number of methods - e.g. X2Go, MobaXTerm, or Terminal/X11

Set-up video with demonstration available on the UKCA YouTube channel.

The screenshot displays a desktop environment with a VM management interface and a terminal window. The VM management interface shows a table of tasks and their status:

task	state	host	job system	job ID	T-submit	T-start	T-finish	dT-mean	latest message
1	running								
VM	running								
fcm_make	succeeded	localhost	at	1	09:20:22Z	09:20:23Z	09:21:52Z	PT1M29S	job(01) succeeded
recon	succeeded	localhost	at	2	09:21:55Z	09:21:55Z	09:22:07Z	PT12S	job(01) succeeded
atmos	running	localhost	at	3	09:22:09Z	09:22:10Z	*	*	job(01) started

The terminal window shows the following commands and output:

```
ukcatr01@ukca-vm01: ~/cylc-run/u-cb681/work/1$ ls
1
ukcatr01@ukca-vm01:~/cylc-run/u-cb681/work$ cd 1/
atmos/ fcm_make/ recon/
ukcatr01@ukca-vm01:~/cylc-run/u-cb681/work$ cd 1/
atmos/ fcm_make/ recon/
ukcatr01@ukca-vm01:~/cylc-run/u-cb681/work$ cd 1/atmos/pe_output/
ukcatr01@ukca-vm01:~/cylc-run/u-cb681/work/1/atmos/pe_output$ ls
atmos.fort6.pe0  atmos.fort6.pe1  atmos.fort6.pe.stdout
ukcatr01@ukca-vm01:~/cylc-run/u-cb681/work/1/atmos/pe_output$ tail -f atmos.fort6.pe0
0.0000E+00
0.0000E+00
update pattern: updating coeffc and coeffs
*****
+ Linear solve for Helmholtz problem
+ Outer Inner Iterations  InitialError  FinalError
+ 1 1 17 0.100000E+01 0.461313E-04
+ 1 2 7 0.142479E-01 0.635002E-04
*****
Qcf < 0 fixed by PC2, ( 7 occurrences)
Qcl < 0 fixed by PC2, ( 4 occurrences)
```