# Met Office HPC Update

Paul Selwood

Thanks to:

Adam Voysey, Matthew Glover, Michele Guidolin, Andy Malcolm, Sam Cusworth, Richard Gilham, Lucian Anton, Duncan Roweth, Sam Clarke, Ilia Bermous

# HPC Facilities

# Met Office HPC Systems

- 3 top 40 systems at launch
- Architected for reliability
- 7PF and 2x 2.8PF (HPL)
- Intel Xeons. Mostly Broadwell, some Haswell.
- 6 Lustre filesystems totalling 24 PB
- Cray Aries interconnect
- PBS Pro scheduling and Cray/Intel compilers
- 50/50 Weather/Climate usage

# 2018 Operational Models

**UK 1.5km** 70L (80 km top)
*12 hrs – 16/day*
*54 hrs – 6/day*
*120 hrs – 2/day*
*+12 member ensemble 2.2 km*
*54 hrs – 4/day*

**Driving**

**Global 10km** 70L (80 km top)
*66 hrs – 2/day*
*144 hrs – 2/day*
*+18 member ensemble 20 km*
*144 hrs – 4/day*

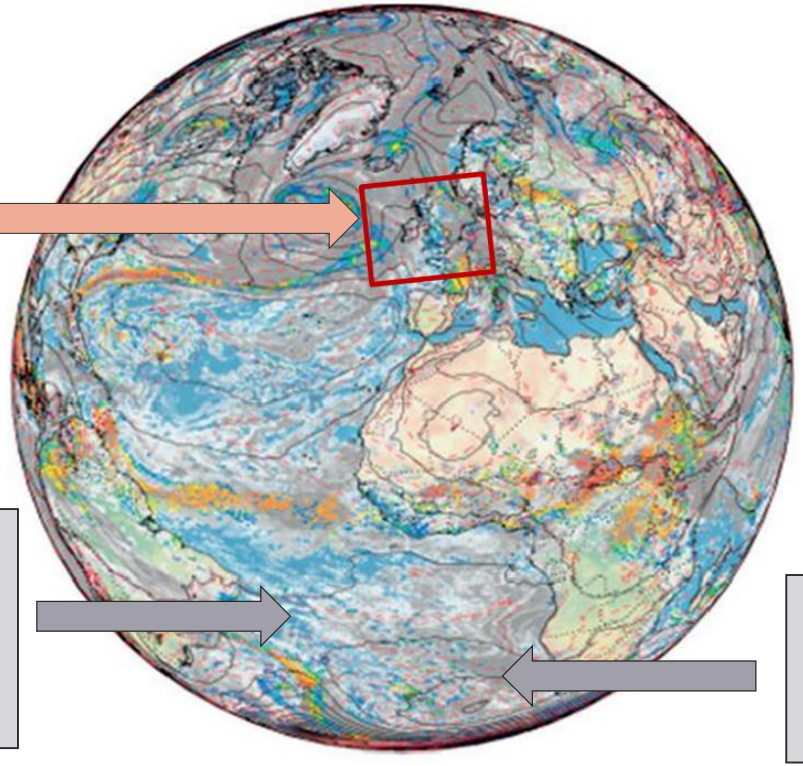

**Seasonal** 85L (85 km top)
¼ degree Ocean
*Ensemble at 60 km*
*7 months– 2/day*

# Architecture Investigations

- Early Cray development system now Knights Landing and Analytics

- Partner with GW4 Alliance for Isambard
- EPSRC funded, multi-architecture platform including NVIDIA GPU and Intel KNL
- Will have significant 64-bit ARM partition
- For details – see Simon McIntosh-Smith's talk!

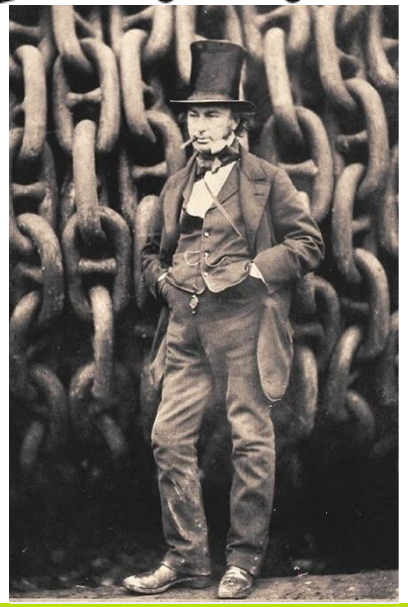

# Single Node Comparisons

**Broadwell**
Swan - Intel Xeon (Broadwell),
2 × 22-core @ 2.2GHz

**Skylake**
Swan - Intel Xeon (Skylake),
2 × 28-core @ 2.1GHz

**KNL**
XCK – Intel Xeon Phi (Knights landing),
64-core @ 1.3Ghz

**TX2**
Isambard – Cavium ThunderX2
2 × 32-core @ 2.2GHz

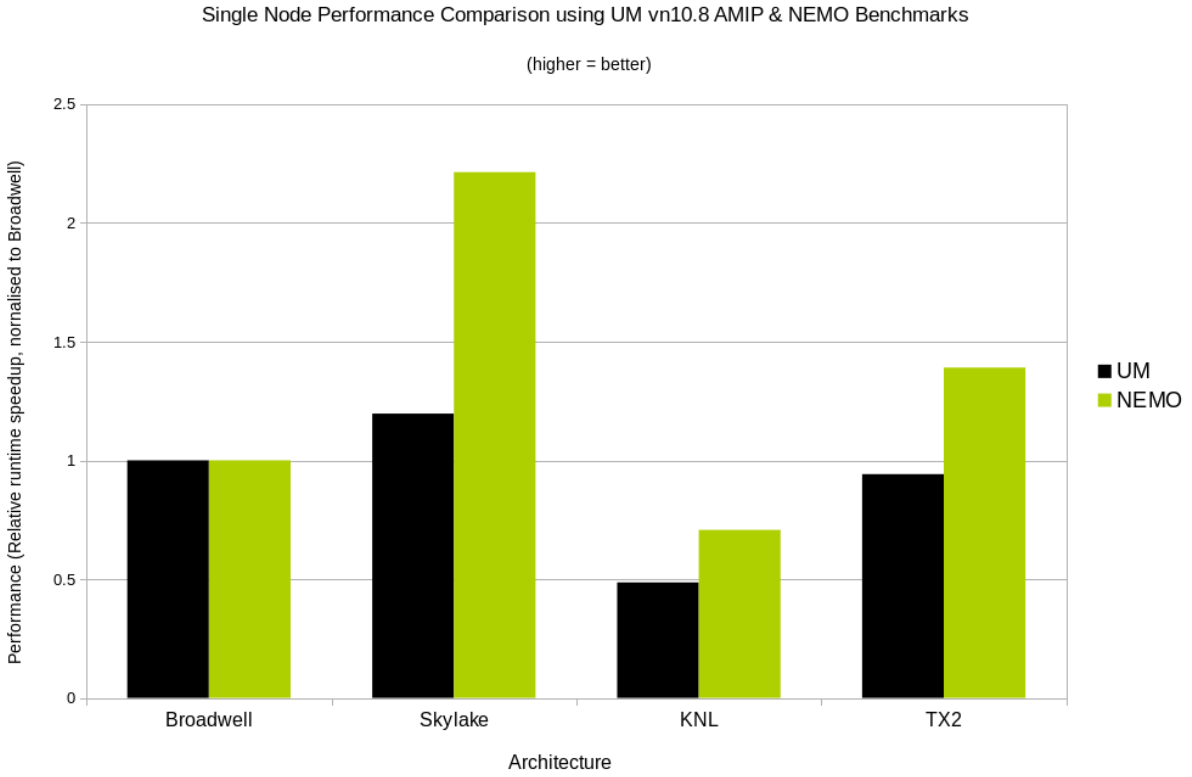

# Operational 10km Global

# Parallel Suite 39; The Problem

- Scalability projections suggested we could afford to run a 10km global model (N1280L70) even though benchmarks had been at 12km (N1024L70)
- Initial experiments on the brand-new XCS system gave runtime of 55 minutes on 518 nodes for 7 day runs
- BUT, Early runs in PS 39 took up to 74 minutes. Limit is 60 minutes!
- Operationally unacceptable timescales
- Runtimes too variable

# Hypothesis: Noise on Interconnect

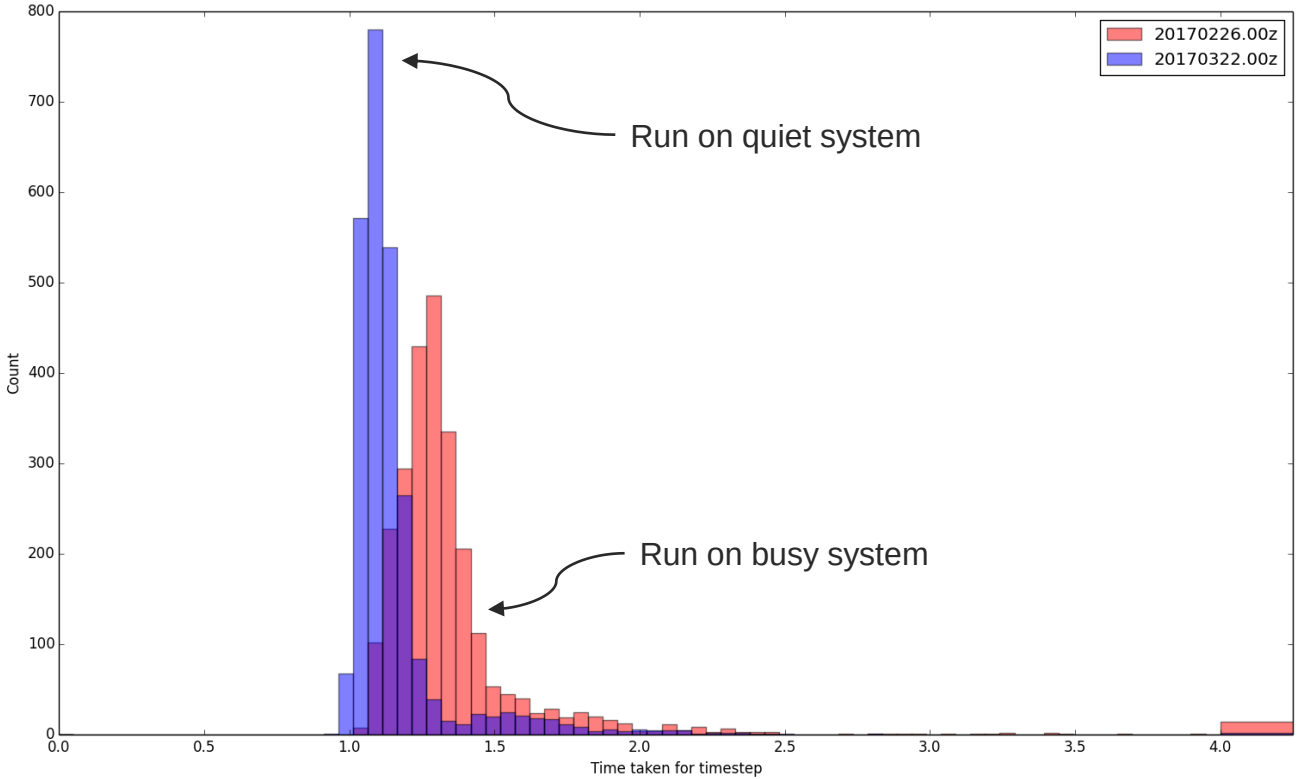

# Solutions to Noise

- 518 nodes spread over 7 electrical groups
- Minimise Rank 3 exposure
- New placement set; quartet
- Bias routing to minimal path for job
- Minimise other jobs on quartet; oversized reservation

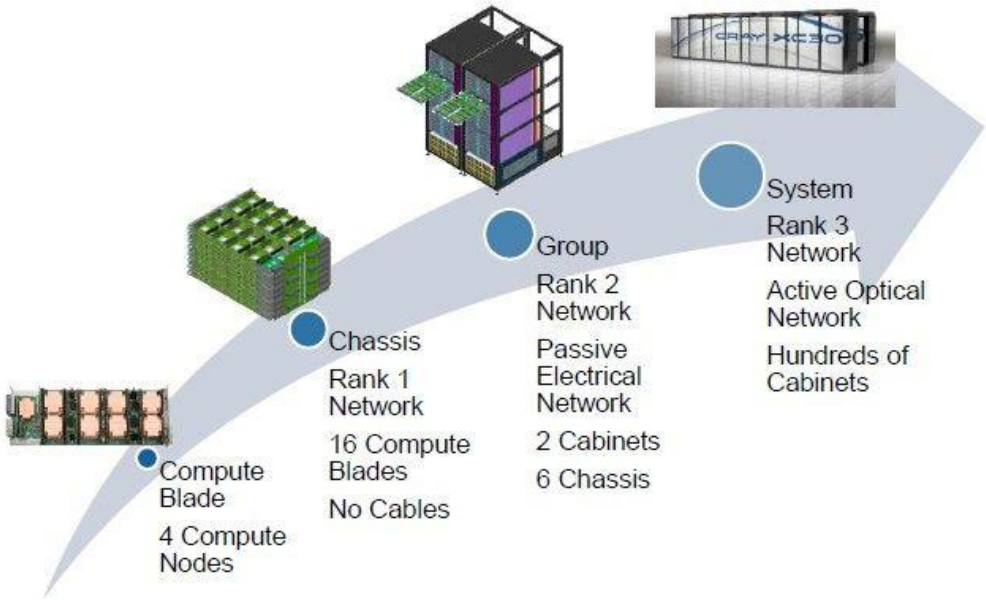


Image: © Cray

- Later: bias routing for I/O traffic

# 2 vs 7 Electrical Groups

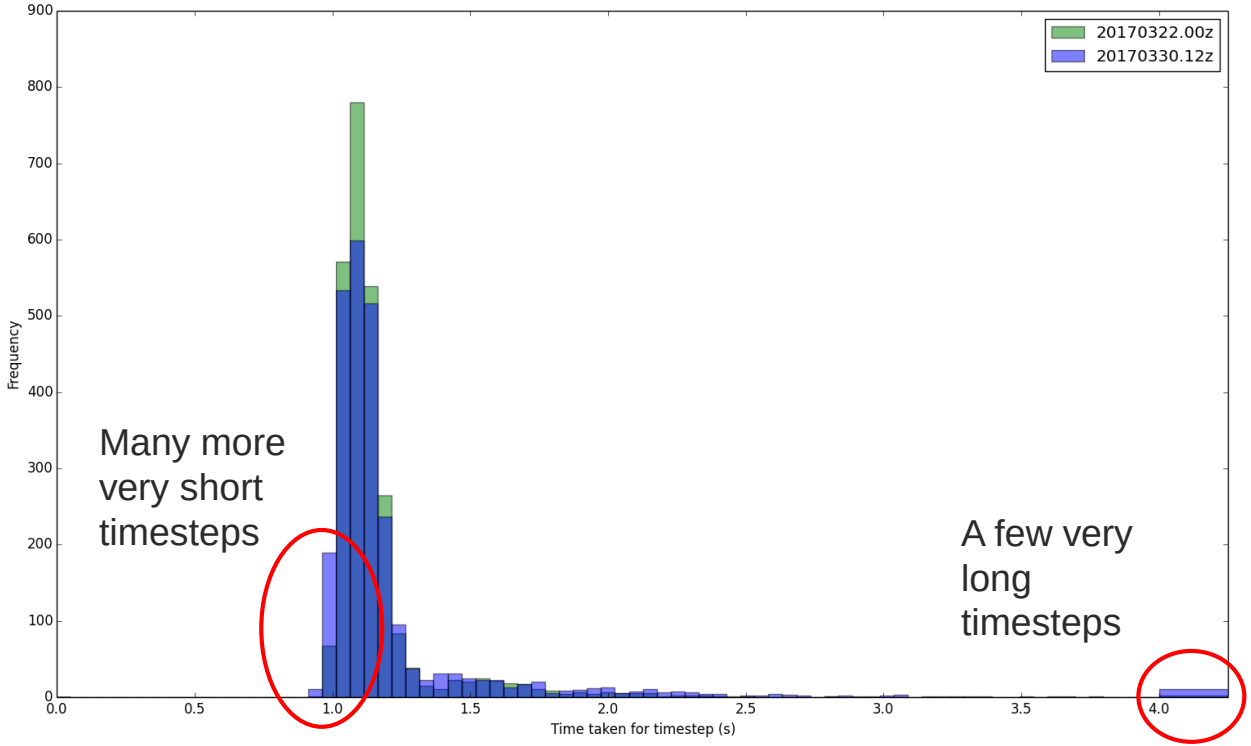

# Noise Conquered

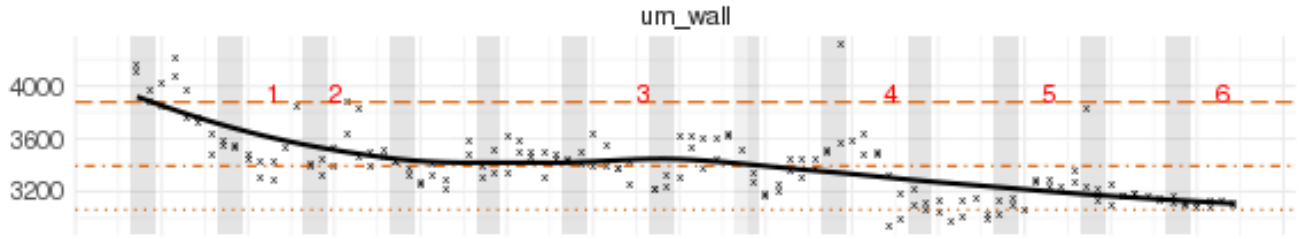


- >70 minutes down to ~52 minutes
- Variability reduced
- Operationally acceptable times
- Affordable cost
- Required close Met Office / Cray collaboration

# Experiments in 32-bit Physics

# LS Precipitation Pilot

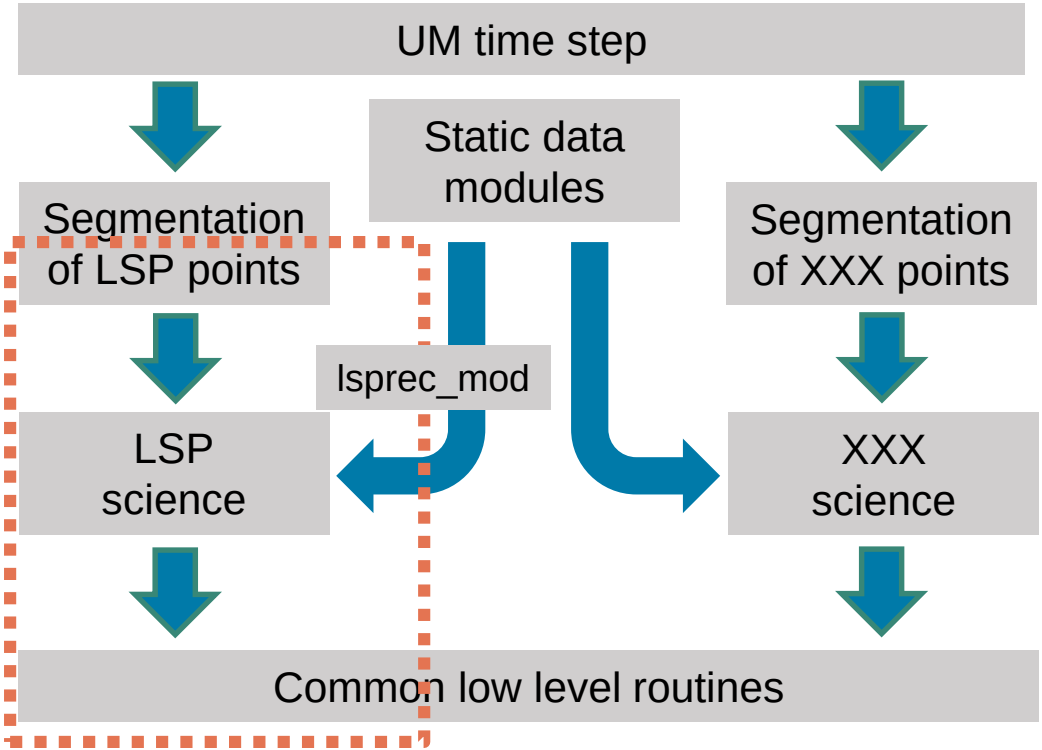


- Large Scale Precipitation chosen as a pilot study
- Tightly defined area with clean interface
- Expensive in regional models
- Compile-time choice of precision
- What impact to scientific accuracy?
- What impact on runtimes?

# Scientifically Similar

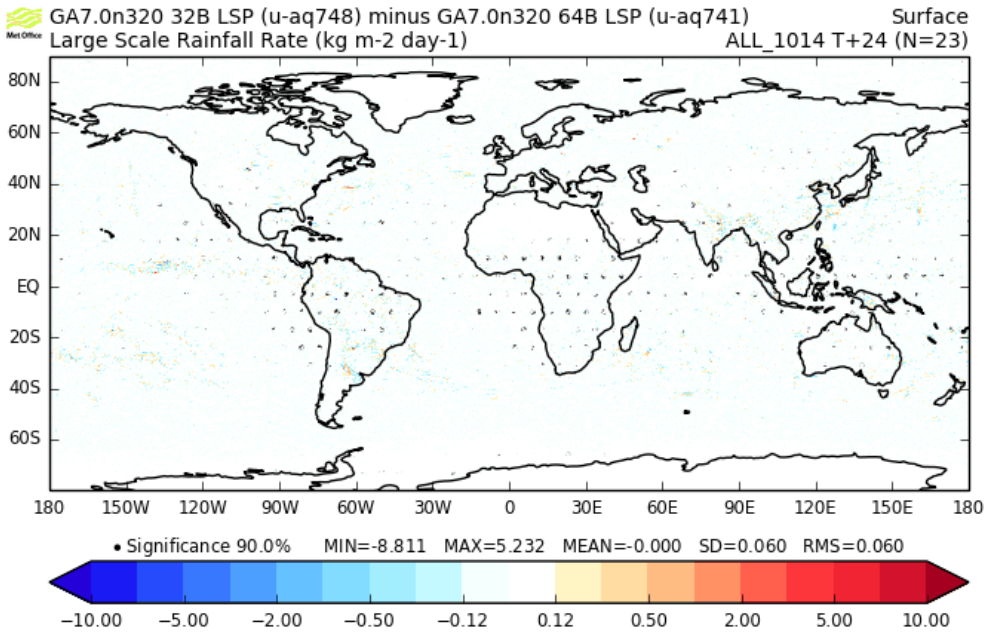


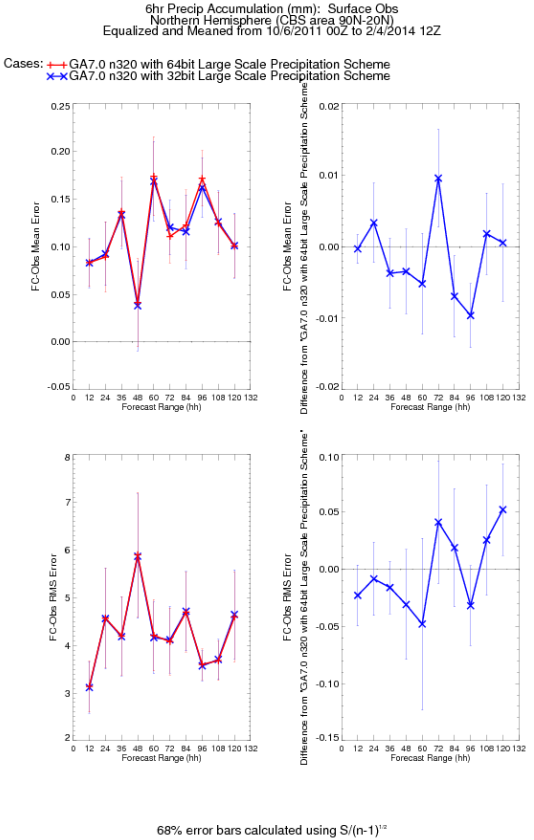

# Runtime Impact

Science payload varies between GA6, GA7 and UKV

Performance gain depends on configuration, build optimisation and other factors

| | **N512 GA6 2 day** | **N48 GA7 AMIP 30 day** | | **UKV 1.5 km 24 hour** | | |
|---|---|---|---|---|---|---|
| Build/ decompostition | High 10x23 | Safe 12x9 | Safe 16X9 | Safe 8x33 | Safe 12x22 | High 6x41 |
| UM speed-up | 4.5% | 1.3% | 2.8% | 2.5% | 3.9% | 5.5% |
| LSP scheme speed-up | 33.5% | 14.9% | 16.3% | 38.0% | 35.1% | 51.4% |

# Scalability Improvements

# Auto-Tuning Segmentation

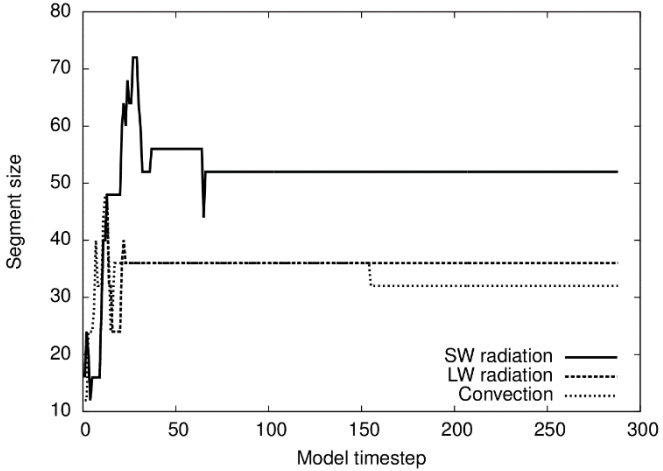


```
SW-RAD> ==> Segment size auto-tuning <==
SW-RAD>              Region name: SW_Radiation
SW-RAD>        Trial size/timing: 56     1.7910E-04
SW-RAD>     Current segment size: 52
SW-RAD> ------  ------------  -----
SW-RAD> Leader  Segment size  Count
SW-RAD> ------  ------------  -----
SW-RAD>      1            52  680
SW-RAD>      2            60  9
SW-RAD>      3            48  8
SW-RAD> ------  ------------  -----
```

- Increased use of segmentation to allow adaptation for cache and threading
- Require tuning for each configuration and architecture
- Previously used expensive brute-force search taking many model runs
- New tuning code using simulated annealing
- Only 1 run to tune

# OpenMP Improvements

- Improving OpenMP over many versions and years
- With Cray XC40 – OpenMP a primary parallelisation strategy
- Better coverage / better balance
- Now running at 3-4 threads / task
- Serial performance improving too
- More to do e.g. regional, ESM.

N768 model, strong scaling
1 thread/core, 12 – 108 nodes

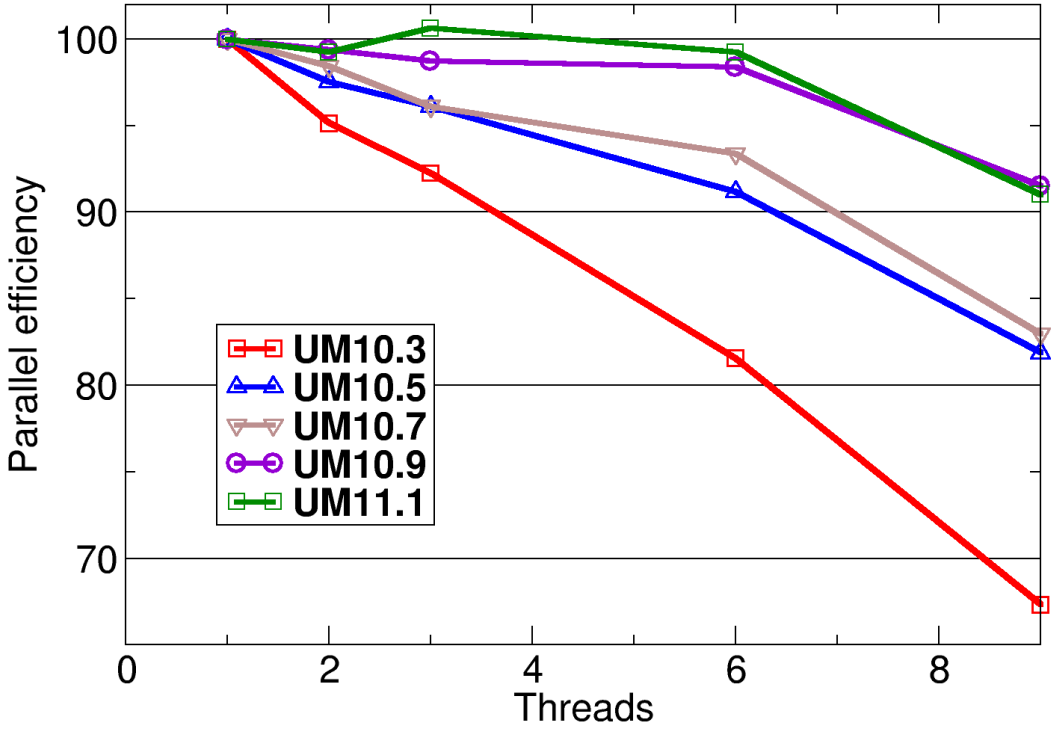

# Operational Suite 40 Improvements

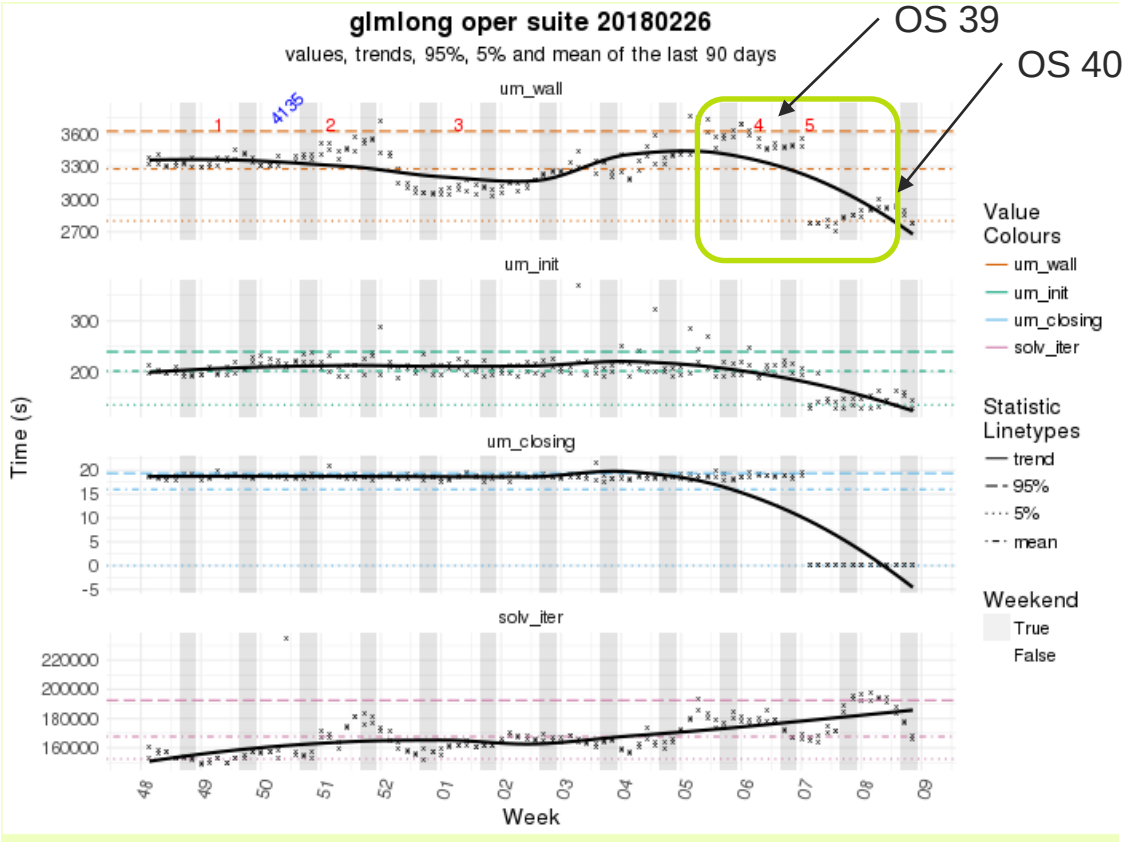


- Global 10km model on 18k cores
- 55 minutes down to 44
- ~200 nodes freed up
- Start-up reduced
- Shutdown reduced
- Less weather dependence

- Significant HPC cost saved => more science done and better forecasts!
- More reliable delivery.

# High Resolution Scalability

- N2048L70 Global model (6km)
- At vn10.4
  - Poor scaling
  - High start-up cost
  - Operational runtimes couldn't be met
- At vn11.1
  - Improved scaling
  - Small start-up cost
  - Potentially able to meet operational runtimes

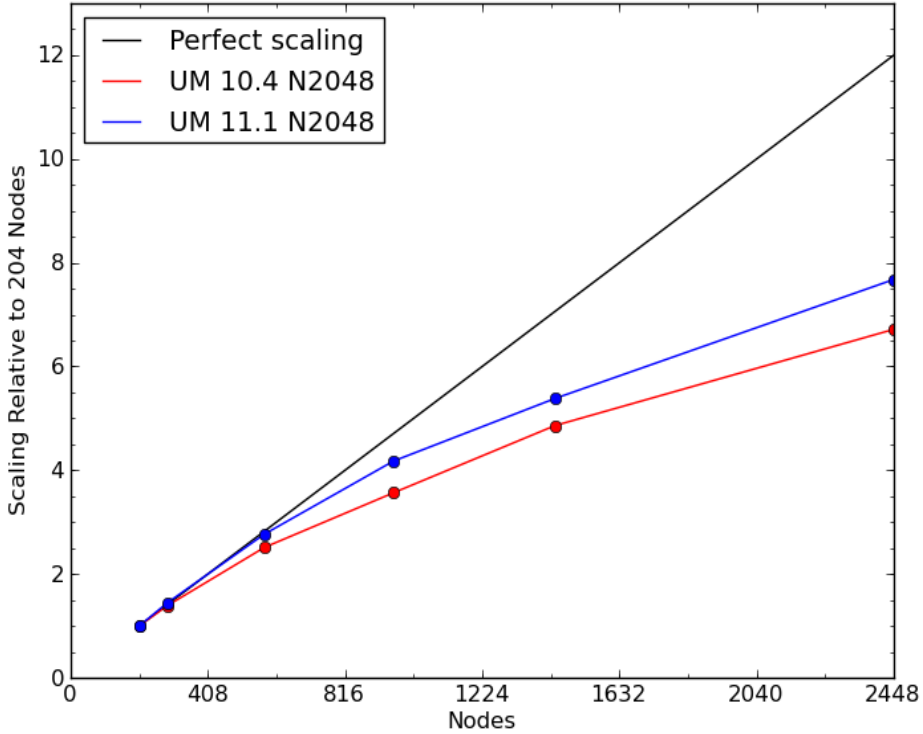

# Thank You!  Questions?