



**Northern Intensive Computing
Environment (NICE)
February 2020**

Alan Real (Director)
Cliff Addison (Motivator...)

Who are we?



- N8 Research Partnership is the 8 most research intensive universities in the North of England – Durham, Lancaster, Leeds, Liverpool, Manchester, Newcastle, Sheffield and York
- N8 Centre of Excellence for Computationally Intensive Research (N8 CIR)
 - Derived from N8 HPC which offered access to a shared compute platform for all 8 universities for 6 years.
 - Established to build on this partnership to expand the *art of the possible* within research themes by leveraging computational- and data-intensive practice.



- Conversation around future of AI & accelerated computing
 - GPU card memory regarded as limiting factor – explore coherence.
 - Quantification of inferencing error through model uncertainty.
- Conversation with the Bioscientists
 - Want to get the most out of lab experiment data
- Identical challenges within advanced materials
- Bring experiment and simulation data together to inform each other.
 - Requires ability to drive the network.
 - Combine observation, AI, models & simulation to advance understanding.
 - In AI, only the **machine learns** to predict.
 - **We** build models to understand & use simulation to verify.

EPSRC Tier-2




- Durham-led, with DiRAC and Cambridge on behalf of N8 Research Partnership
- Similar architecture to Summit and Sierra (but much smaller)
 - Affectionally known as 'Summat'!
 - Prove scaling, exploit GPU/CPU memory coherence & distributed GPU
 - 38% National access.
- Around ~1 Pflop.
 - 32-node IBM Power 9 based system, 4GPUs per node (0.5TB/node), NVLink & EDR IB.
 - 4x T4 & 2x FPGA nodes, 2Pb/10Gb/s Lustre.
- Will graduate towards small number of large-scale projects, rather than capacity production runs.
- Installed ~April 2020, general availability over summer, following initial pilot, Autumn RAP.



RSE roles



- You can't expect to deploy one of these & expect instant take-up!
 - Non x86 architecture, acceleration, distributed deep learning, large memory use all potential barriers to effective use.
- **Each N8 site pledged 1FTE RSE to support their local community.**
- Will coordinate effort where there is mutual need across organisations
- N8 RSEs as Co-applicants for National resources requiring RSE support
- Talking with IBM about early training provision for those RSEs.
 - **System made available to RSEs/experts during pilot**



Training provision

- Working with IBM will help us train the trainers/RSEs
- Starter list for 10:
 - Porting to Power 9
 - large memory support for accelerated simulation (e.g. supporting increased memory sizes for existing CUDA codes).
 - large memory support for AI (exploiting LMS in common ML/DL toolkits)
 - distributed deep learning
 - distributed accelerated simulation (GPU-Direct etc).
 - OpenMP/ACC offload (OpenMP 5)
- Extended pilot over summer so RSEs/community experts can port/evaluate to inform resource requests.
- Turing, HECBiosim and N8 CIR involved in resource allocation processes.

Keep In Touch



- Website – n8cir.org.uk
- Case studies - <https://n8cir.org.uk/supporting-research/rse/case-studies/>
- Twitter - @N8CIR
- Join our mailing list at <https://n8cir.org.uk/contact/>

