Data Analytic Cluster

Software Environment

David Henty, EPCC
d.henty@epcc.ed.ac.uk

www.epcc.ed.ac.uk
www.archer.ac.uk

Hardware

- 1 login node
  - two Intel Ivy Bridge 10-core processors, 128 GB memory
- 12 standard compute nodes
  - two Intel Ivy Bridge 10-core processors, 128 GB memory
- 2 high-memory compute nodes
  - with four Intel Westmere 8-core processors, 2 TB memory
- HyperThreads are enabled on all nodes
  - standard compute nodes each have 40 CPUs available
  - high-memory compute nodes each have 64 CPUs available.
- All DAC nodes have high-bandwidth, direct Infiniband connections to the UK-RDF disks.

DAC use cases

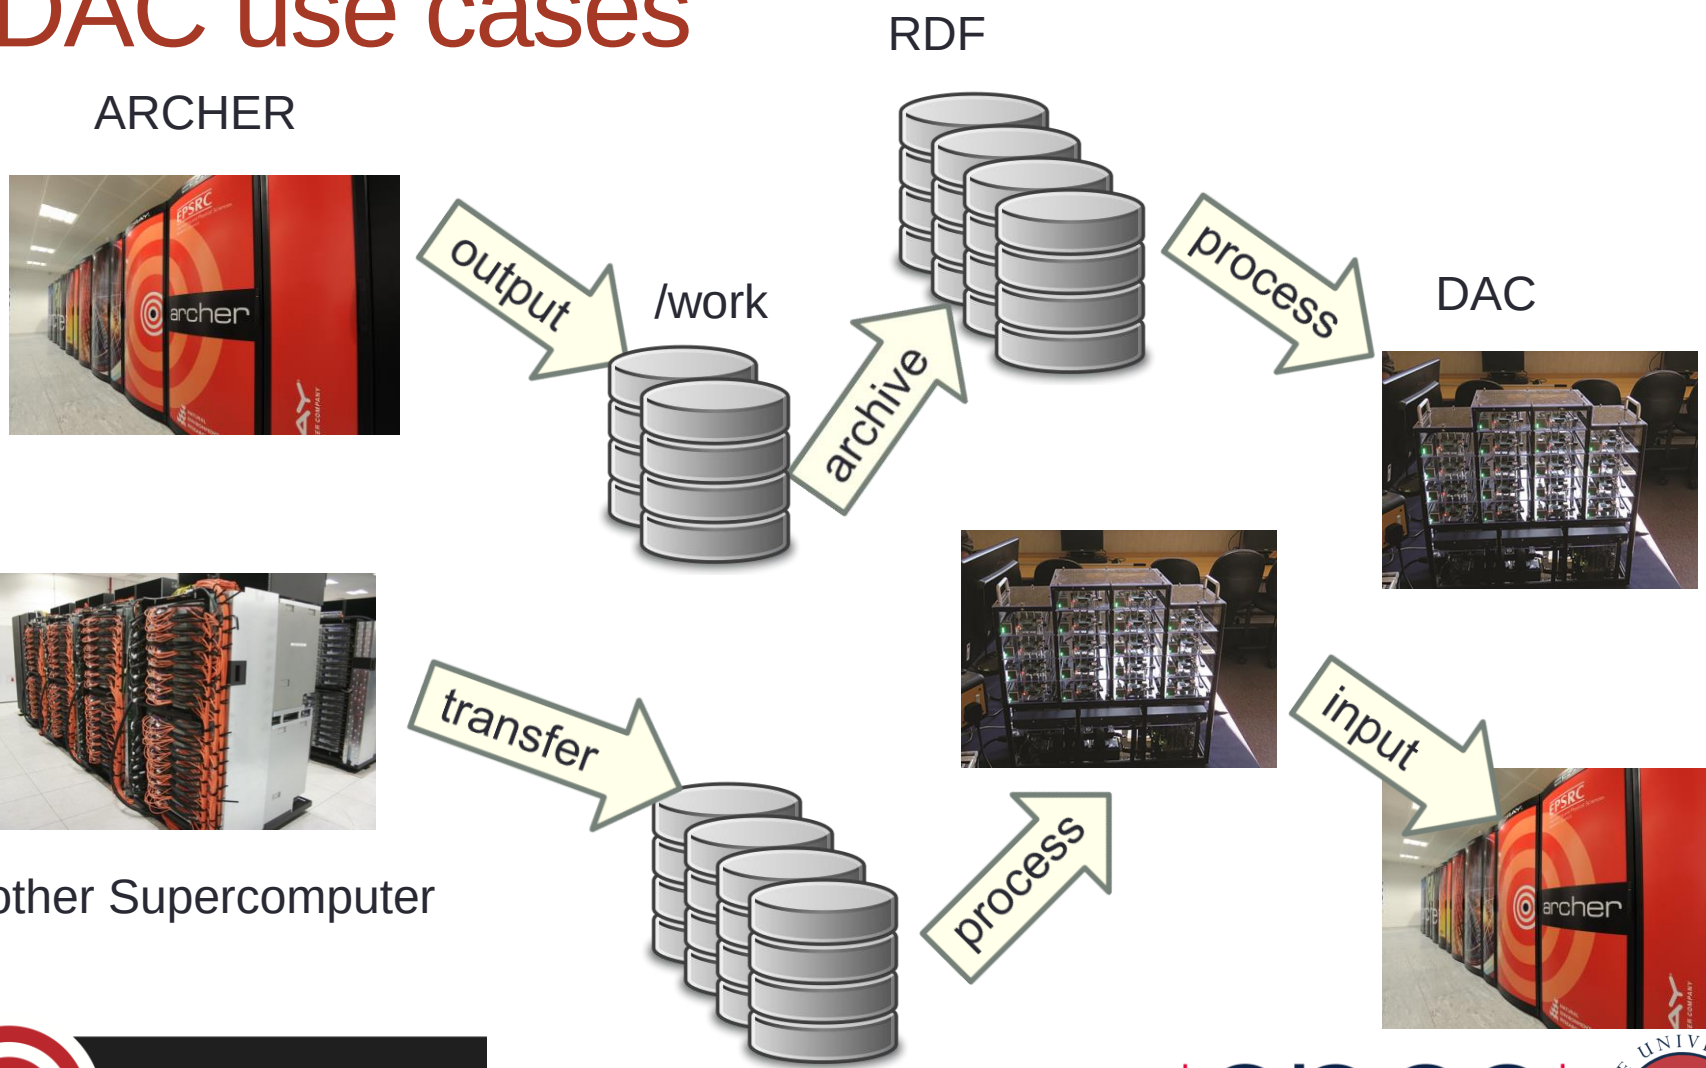

Why use the DAC?

- Fastest connection to RDF disks
  - much faster than ARCHER
- Fast connection to external networks
  - via DTN nodes
  - e.g. PRACE network, NERC Jasmine system
- Easier and more flexible than ARCHER compute nodes
  - more powerful than ARCHER post-processing nodes
  - currently free to use!

Compilers

- GCC
  - gcc – C
  - gfortran – Fortram
  - g++ - C++
- OpenMP
  - compile and link with –fopenmp flag
- MPI – OpenMPI library
  - module load openmpi-x86_64
  - compile: mpicc, mpif90, mpic++
  - run: mpiexec –n <nproc> mympiprogram

Interactive access

- Often useful to have a shell on the compute nodes
  - testing
  - debugging
  - visualisation
  - ...
- Submit an interactive job, e.g.
  - qsub -IXV -lwalltime=3:00:00,ncpus=16
  - wait for prompt ...
- Notes
  - you start off back in your home directory
  - remember to reload your modules!

Python

- Python 2.* available via the Anaconda distribution
  - module load anaconda
- Python 3 also available
  - module load anaconda/2.2.0-python3
- Parallel python
  - MPI provided by anaconda: from mpi4py import MPI
  - load normal MPI module
  - mpixec –n 4 python myjob.py

Visualisation

- Paraview is available
  - module load paraview
- For parallel visualisation
  - module load paraview-parallel

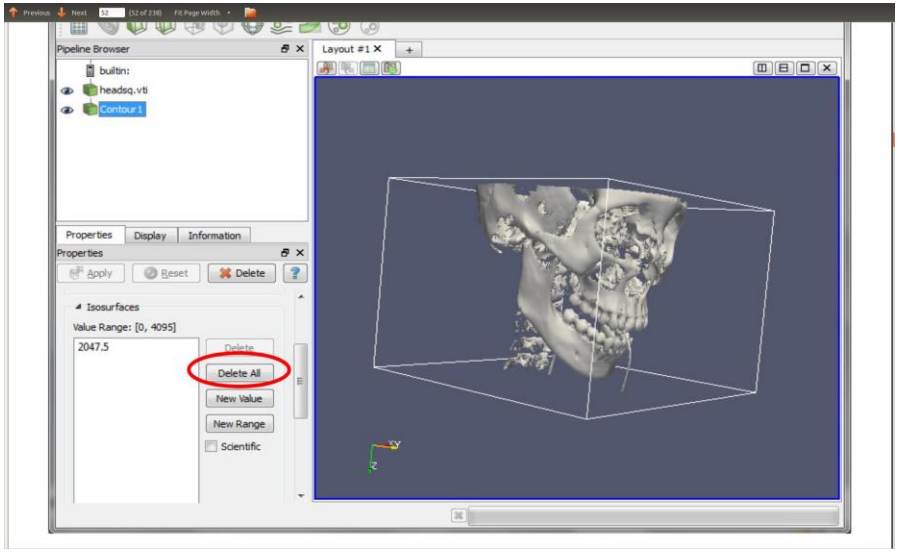

- This works in client/server mode
  - run paraview GUI as a client
  - run parallel paraview server “pvserver”
  - connect the two via a socket

Parallel Visualisation

- See http://www.archer.ac.uk/documentation/rdf-guide/cluster.php#paraview

```
-bash-4.1$ hostname rdf-comp-ns10
-bash-4.1$ qsub -IXV -lwalltime=3:00:00,ncpus=16
-bash-4.1$ module load paraview-parallel
-bash-4.1$ mpirun -np 16 pvserver --mpi --use-offscreen-rendering --reverse-connection --server-port=11112 --client-host=rdf-comp-ns10
```

- Assumes a paraview GUI listening on port 11112
  - run GUI on the login node
  - see: File -> Connect

Remote visualisation

- Exporting graphical display slow over network
- Assuming you have paraview on your laptop ...
  - run GUI locally
  - connect to parallel pserver running on DAC
- Requires *port forwarding*
  - see http://www.archer.ac.uk/documentation/rdf-guide/cluster.php#portfwd
  - some compatibility restrictions on paraview versions ...

Other software

- Visualisation
  - VisIt
- Statistics
  - “R” is available by default (no module)
- Data Formats; HDF5 and NetCDF (see later)
  - serial versions available by default
  - parallel hdf5 available via standard wrappers, e.g. h5pcc and h5pfc
  - parallel netcdf requires a module + flags – see documentation
- Linear algebra
  - BLAS and LAPACK available by degault
  - for parallel, link with: -lmpiblacs -lscalapack