

UK Research Data Facility

Kieran Leach

Senior HPC System Developer

k.leach@epcc.ed.ac.uk



EPCC

Mission statement:

“To accelerate the effective exploitation of novel computing throughout industry, academia and commerce.”

-Established in 1990 at the University of Edinburgh

-75+ staff working both on computation research and support

-Manage and provide a variety of HPC systems and services

EPCC – Hosted Systems

ARCHER – UK National HPC service

- Hosted and managed for EPSRC and NERC
- £43m Cray XC30 service
- 118,000 cores on 4920 nodes with 11,090 GB/s peak bandwidth

IBM Bluegene/Q

- Hosted and managed for DiRAC consortium
- 6 cabinets, popular for particle physics calculations

Cirrus

- Hosted and managed for both EPCC and EPSRC and UoE
- SGI ICE XA system with 10,080 cores on 280 nodes



The RDF

Persistent filestore for UK Researchers

- At the end of an HPC service data is generally destroyed
- Users provided space on RDF to ensure continued availability

Secure location to store research data

- Backed up offsite

Fast connection to other HPC services

- Directly connected via 40GB/s ethernet to ARCHER, other services
- Multiple 10GB/s connections into site, more available via DWDM



The RDF - Hardware

DDN storage racks

- 23PB usable storage across 7x2 rack 12K units
- Several NetApp SSD shelves for metadata

NSD servers

- 12 x3650s

Interconnect/Network

- Mellanox Infiniband network internally
- Mellanox 40 GB/s Ethernet network for connecting to other systems
- Dual 10GB/s link from site to DR site at off-site UoE location

The RDF - GPFS

Currently GPFS 3.5

-Moving shortly to 4+

Multiclustering for client systems

TSM for offsite backup



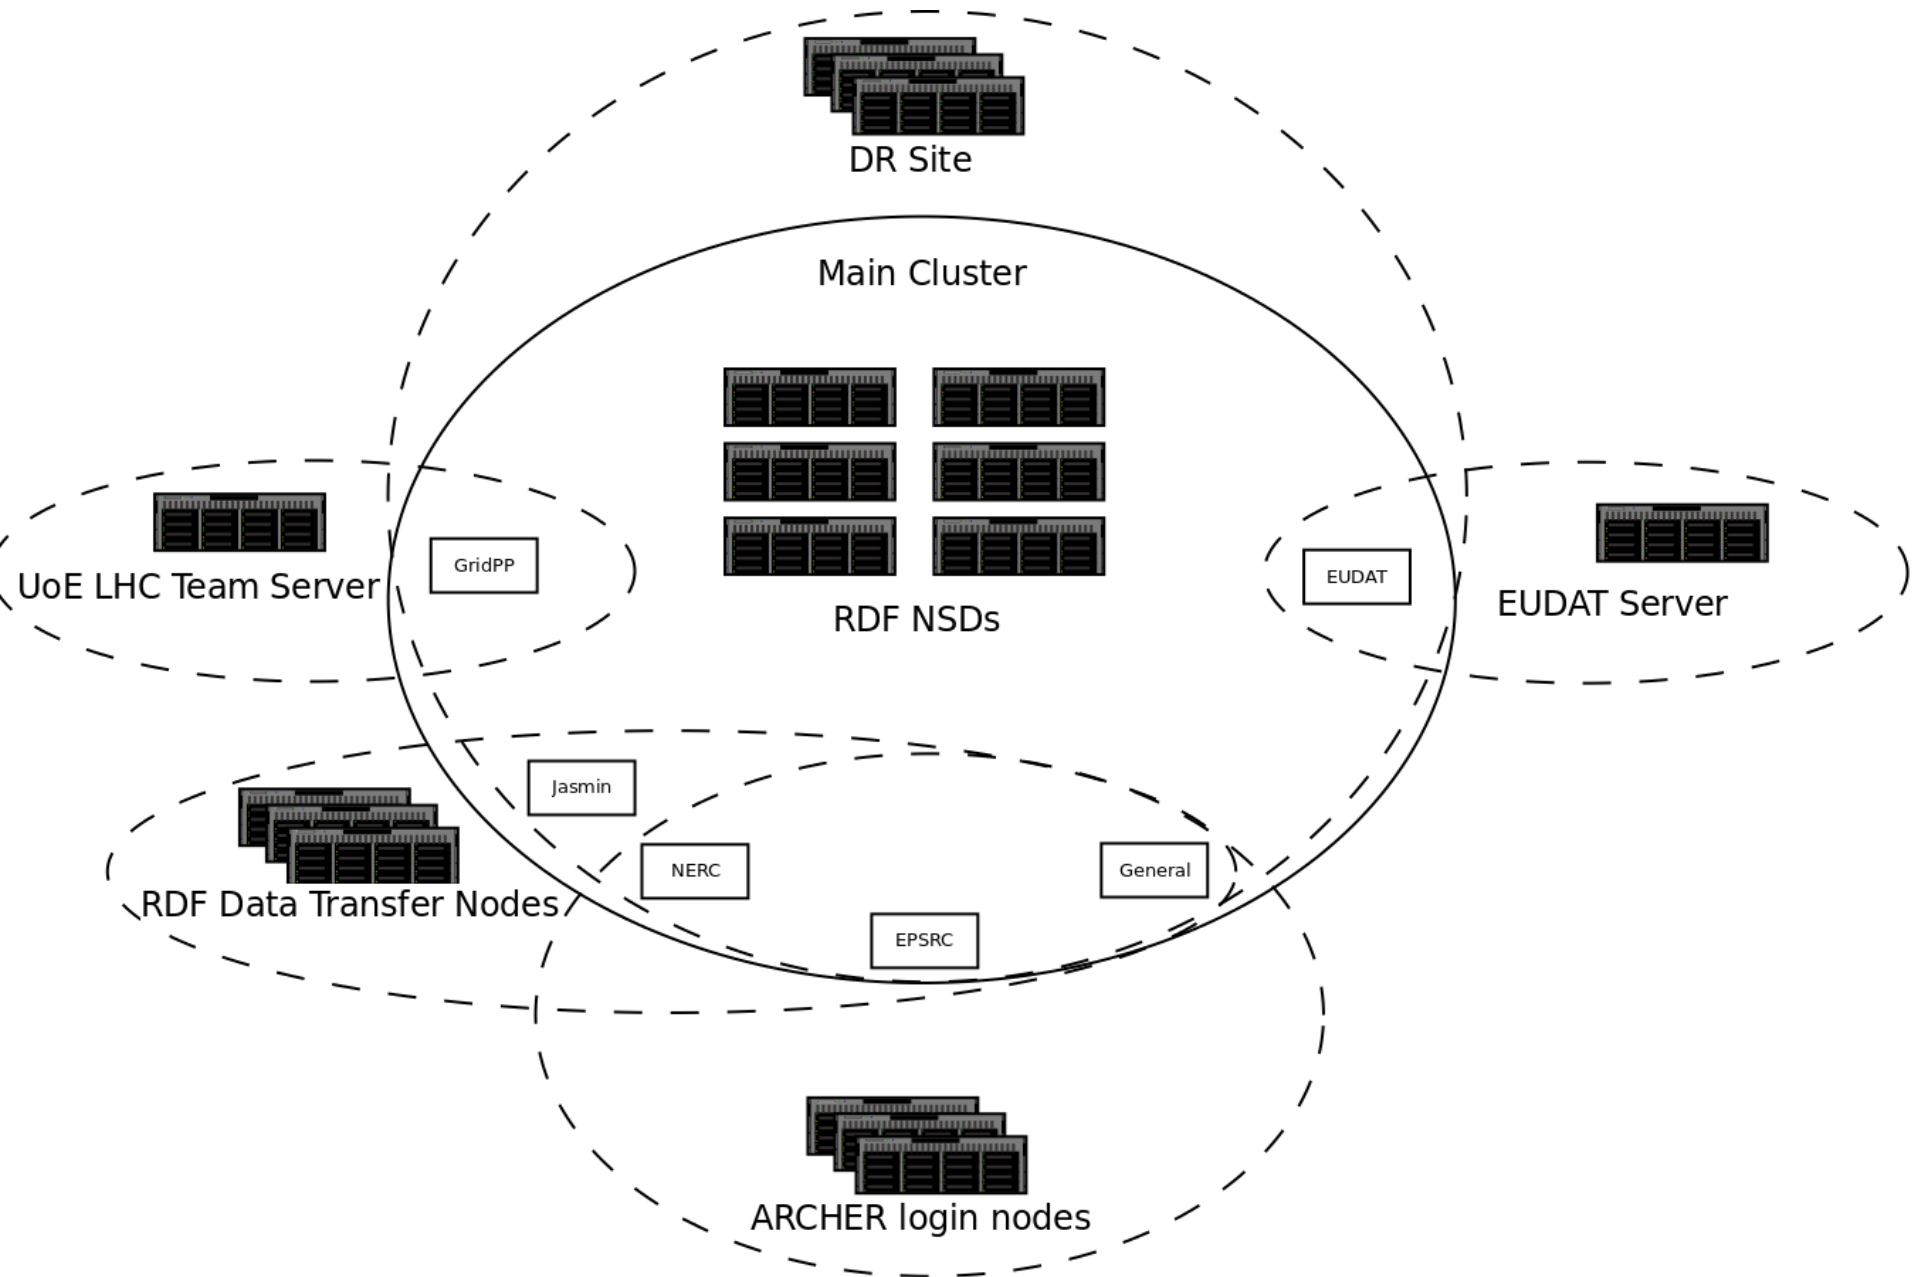
Multiclustering

Used to control access to filesystems on various client systems

Individual systems only given access to necessary/appropriate filesystems

Root remapping used for security purposes where appropriate

In conjunction with containers on the RDF compute cluster we can provide a user-configurable system with total access to an HPC-grade filesystem for “Big Data” analytics



Backup

4 Backup servers

Extra DDN 12K shelves for staging

14 frame tape library (50PB capacity)

TSM for primary backups

Multiclustered for syncing, restoring

Data Analytic Cluster

20 NextScale Compute Nodes

4 High Memory x3850 Compute Nodes

Central xcat management node

Dual 56 GB/s Infiniband to each node, allows direct disk access

Multiclustered allowing variety of filesystems to be offered

Some nodes configured as Compute Cluster w/scheduler

Some nodes hosting containers provided to users for bespoke usecases



Questions?