

Edinburgh Forward Look

Andrew Washbrook

University of Edinburgh

ScotGrid face to face meeting 2016, Edinburgh

30th June 2016



Eddie Mk3 Status and Future

Eddie Mk3 Compute

- Research computing facility available free at the point of use to all staff and students
- Groups can either purchase their own equipment to be hosted by Research Services or buy a dedicated share of Eddie resources
- Around **600** active users so far

Total number of cores	6,200
Cores free at point of use	4,000 (increasing to 5,000 very soon)

Server	Dell Poweredge R630
CPU	2 x E5-2630v3 2.4GHz 8/16C
Memory	64 GB
Storage	1 TB scratch

Eddie Mk3 Worker Node specifications

Eddie Mk2 cluster is (effectively) gone! No legacy equipment will be supported

Eddie Mk3 Storage

- High performance storage for data created or modified on the cluster
 - **370 TB** of fast disk (10k) directly attached to core cluster network
 - GPFS, metadata on SSDs
 - Specified to deliver minimum 100 iops to a fully occupied cluster
- General research data storage is provided by **Datastore**
 - **8 PB** of slower (7.2k) high capacity disks
 - GPFS
 - Hosted across three sites, extensive use of tape storage for backup and DR purposes
- IBM Active File Management
 - Acts as a cache for frequent access to Datastore files from cluster worker nodes
 - Write performance equivalent to HPC storage
 - Early issues with stability and performance have been resolved

This storage is used for Tier-3 (local group) activities but not for our Tier-2 operations

GridPP computing on Eddie Mk3

Paid for guaranteed baseline access

- 11 ring-fenced nodes exclusively for our use
- 352 (HT) cores
- We are allowed to add bespoke requirements
 - Hyperthreading, alternative job resource limitations, application profiling
- Level of guaranteed baseline access can be expanded or contracted depending on funding

Opportunistic resources and share management

- Access to over 5,000 cores
- Scheduling configuration needs to be tuned to match our MoU requirements
- All done by *functional tickets* to give a point in time job priority
 - Applied at group level for GridPP workload
 - Ticket allocation can be sub-divided to allow internal prioritisation

Fully Costed Resource

- Power, infrastructure and support are incorporated into the core-hour price
 - Opportunistic resources and ongoing cluster expansion are essential to avoid a funding "death spiral"
-

Eddie Mk3 Migration

I often talk about this in our weekly tech meetings - why is it taking so long?

- Forced to make the leap from SL6 to SL7 (and not CC7)
 - ATLAS and LHCb have not yet validated their software for SL7
 - Grid middleware still working on migration
 - Had to regenerate the Grid pool accounts from scratch
 - Now managed by University Information Services (not ECDF) through Active Directory
 - Major effort required to get these created and working with the cluster
 - Switch from Cream CE to ARC CE
 - No playbook on how to get ARC working on a shared facility
 - Required full understanding of ARC software to add necessary modifications
 - Various shared cluster limitations
 - Overhaul of all our legacy dependencies (back to 2009!)
 - Negotiating access to Eddie submit hosts from CEs for job submission and monitoring
 - Bespoke requirements for worker nodes
-

Migration Status

- Eddie Mk3 move is almost complete
- We are the one of only two customers still running on old Eddie Mk2 kit!

Worker Nodes

- CVMFS installed
- Resolved SL7 dependencies
- Grid certificates and revocation set up
- NFS area available for sharing ARC directories, pool account home directories and local configuration

Middleware

- New CE and ARGUS services deployed
- Jobs can be submitted to Eddie Mk3 (currently ring-fenced nodes only)
- All functional tests passing

VOs

- ATLAS functional tests pass and software validation mostly done
 - LHCb: bug in pilot that affects environment setup - has been addressed by developers
 - Other VOs yet to be validated (but expected to be less of a problem than LHC software)
-

Eddie Mk3 Forward Look

- Year on year expansion of the Eddie cluster over the next five years
- Individual groups are procuring batches of hardware to add the cluster
 - Allows “reverse” ring-fencing which we benefit from
- Provisioning and testing of new Cloud service (see next slide)

2016-17 Investment

- Core expansion of the cluster (adding over **1,000** cores)
- Rolling memory upgrade across the cluster (**192 GB** per node)
- Large memory machines (**3 TB** per node)
- "Hundreds" of GPUs (Nvidia K80 and Titan X) as part of Data Science investment

Take advantage of new equipment for our Tier-2 operations

- More opportunistic resources available
 - High memory servers
 - LHC upgrade tracking studies
 - GPU cluster
 - Ground work in accessing GPUs to the Grid already been covered (CHEP 2013 study)
 - Accept new types of workload (LHC simulation and tracking, deep learning techniques, Non-LHC VO applications)
 - Cloud computing resources
-

Eddie Cloud service

- Research Services are deploying Openstack using 26 nodes from the Eddie cluster
- The number of hypervisors will scale as usage grows
- Consultation with Sardina systems
 - FishOS cloud management platform
- Started to talk to groups to get early adopters
 - LSST (Marcus) and GridPP signed up already
- Use 250TB Ceph backing storage (both block and object based)
- Length of evaluation window undetermined
- Resource usage will not charged during this time
- If successful then long term idea to move standard batch system instances into the cloud to allow dynamic adjustment



Middleware and Storage Plans

Middleware Forward Look

- Purchased two new Dell Poweredge R630s to host a new middleware platform as part of GridPP4+ capital refresh

Strategy

- Consolidate services and use resources we have more efficiently
 - We are limited by both budget and power
- Virtualise everything (if possible)
- Explore High Availability options
 - Marcus exploring DRBD/Pacemaker/Corosync solution (see HEPSYSMAN talk)
- Review configuration management and service redundancy practices
- Enhance monitoring and analytics reach based on previous efforts

Warranty forward look

- GridPP4+ capital limited our exposure to out of warranty equipment
 - Storage servers refreshed with new 8 TB disks
 - Perfsonar server warranty runs out at the start of next year. What then?
-

Storage Forward Look

- All our dedicated storage (~1 PB) is now managed by ZFS (see Marcus' talk later this morning)

RDF

- 215 TB available as part of middleware readiness testbed for some time
- To be updated to about 700 TB RDF storage
- Scaling testing needed to determine how to make this available through our production DPM



University Datastore

- Initial exploration of how this can be used for Grid storage earlier this year
- Testing was inconclusive but we were on a tight timeframe for GridPP4+. Revisit?

Storage R&D

- Some ideas for future work from Marcus:

Multi-VO XrootD only installation

- DPM-less
- Distributed endpoints could work like a single UK site
- Wouldn't matter if one site is down but still give the advantage to have all data distributed to all sites
- New VOs do not need any special tools and no file catalogue



ZFS Functionality

- File compression studies on VO files
- De-duplication: Evaluate practical use for Grid storage and determine memory overhead
- Lustre installation based on ZFS



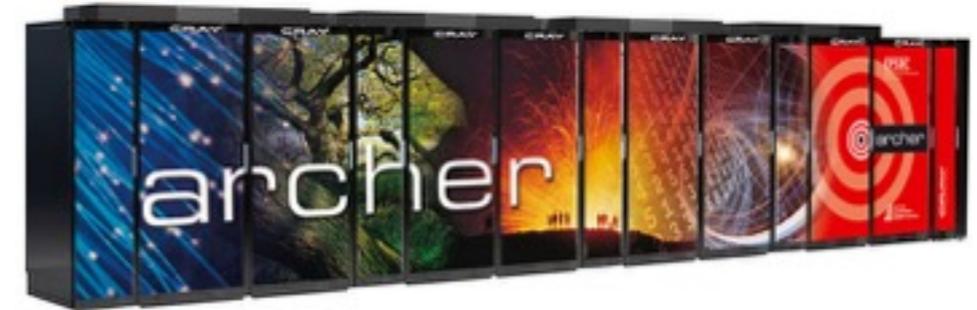
Storage Compute

- Use idle CPU on DPM clients to provide boost to our baseline
 - Minimal batch system + CE required
-

HPC

HPC Activities

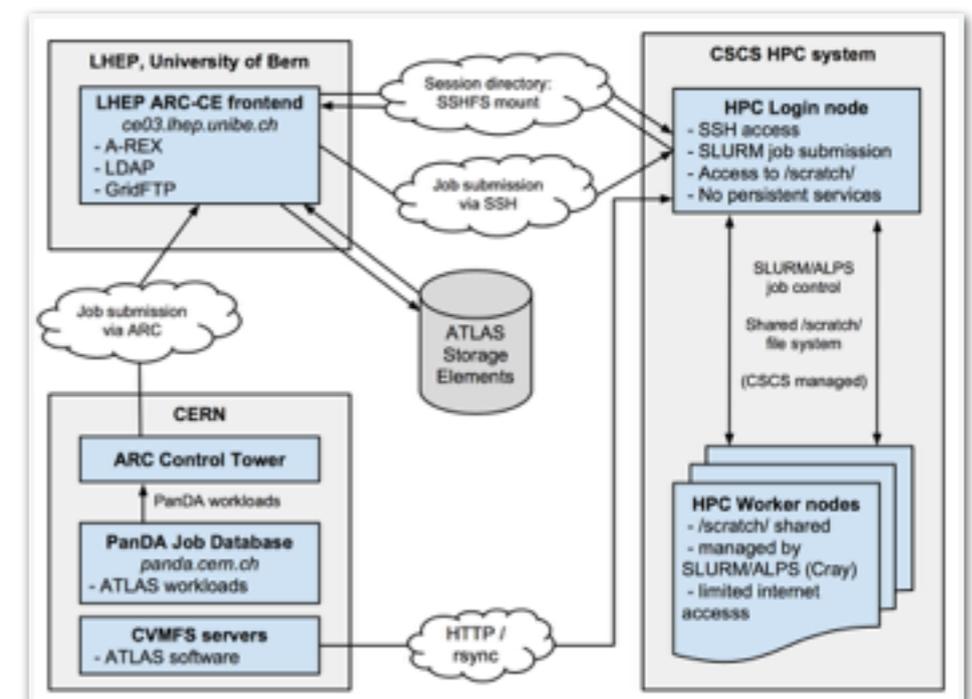
- Aim to integrate the University share of **Archer** resources into our Tier-2 operations
- Private MC simulation run last year demonstrated we can process high volumes of ATLAS workload on HPC
- Now working to become a MC simulation production facility for ATLAS
 - Working with ATLAS ADC developers on pilot submission and job validation



Workload Management Model

- Settled on **ARC CE + ACT** job submission
- ARC CE shares *Session* and *Cache* directories with Archer frontend
- Job I/O handled by ARC CE before and after job
- Lack of external connectivity not an issue
- Software delivery is enabled by CVMFS over Parrot

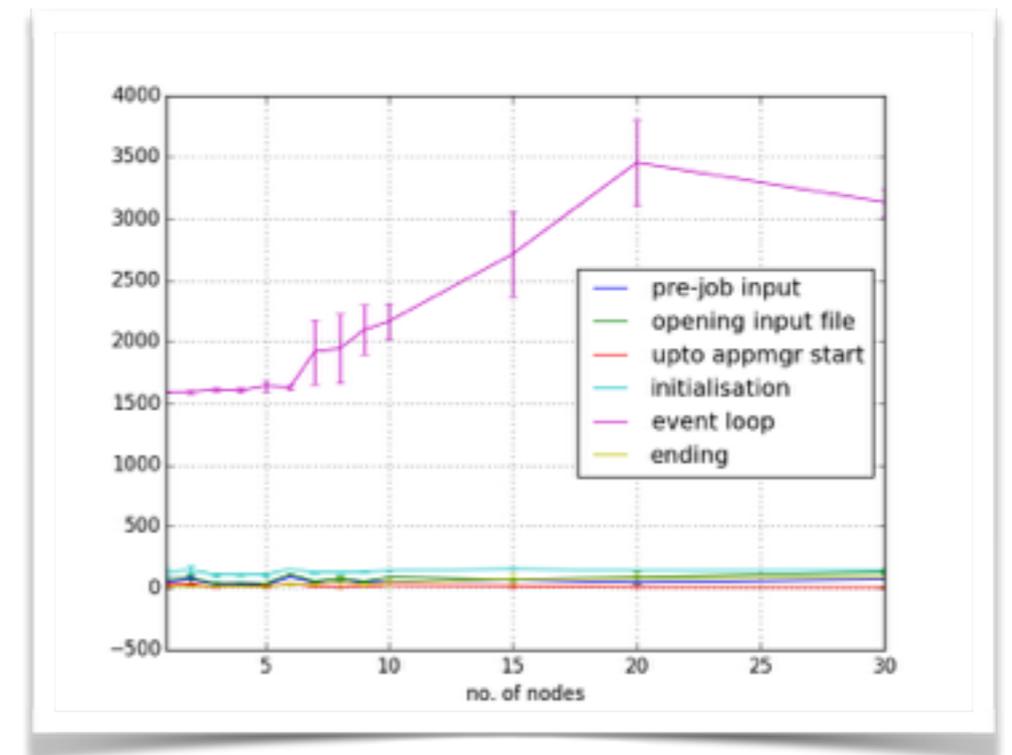
HPC model developed for CSCS



HPC Activities

Optimisation Studies Summer Project

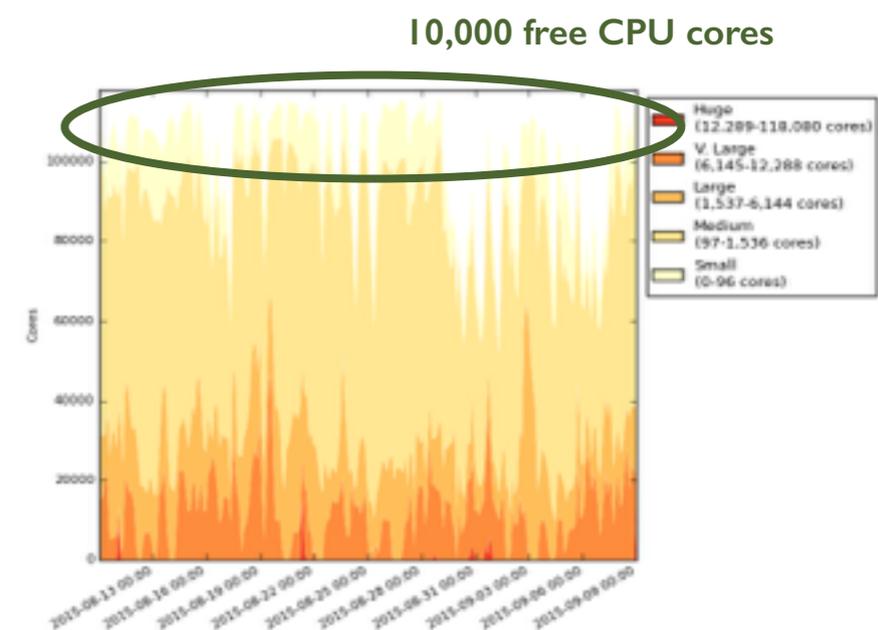
- Job performance significantly suffers as workload is scaled up
- Conducting timing studies varying number of cores, nodes, events processed per job
- Main culprit: Looking into file operations during job lifetime



Preliminary timing studies varying number of concurrent compute nodes

Long term goals

- Move to the **Event Service** model with MPI (Yoda)
- Jobs can be pre-empted with limited impact to mainstay HPC users
- Not necessarily limiting use to ATLAS - use of containers to virtualise workload requirements (e.g **Shifter**) a possibility
- Have to convince Archer admins to make changes
- Lots of overhead in setting up but a big win if pre-emption enabled
 - Even better if this we can widen resource to allow other VOs



Archer Utilisation

Coverage

VO Coverage

ATLAS

- MC production **only** until Analysis pilots are fixed for SL7/CC7

LHCb

- Continue as for Eddie Mk2 once setup scripts in pilot have been fixed

CMS

- Tier-3 like resource possible once stable running on Mk3

LSST

- Workload to be re-validated (Marcus 50% on LSST)

LZ

- Provided **50 TB** of storage for local job submission, re-validate on Mk3

EUCLID

- Parallel effort from RAL - working with ROE staff on long-term data processing project

HyperK/Titus

- Provided support for Tier-3 compute and now migrating their work to the Grid

Any more we should include as part of a coordinated Scotgrid strategy?

Wrap up

- It has been a busy year!
 - A transition to a new (shared SL7) cluster
 - Migrated all our storage to ZFS
 - Overhaul of all middleware services
- Building in practices to reduce future overhead
 - High Availability for middleware
 - Collaboration, issue tracking and documentation: **Gitlab, Slack**

Future Direction

- Diversify site operations across different types of resources (HPC, Cloud, Many core)
 - Continue to expand our VO coverage beyond LHC
-