New Developments in the EGI Accounting Repository

Adrian Coveney,
Stuart Pullinger, John Gordon,
Greg Corbett

www.egi.eu

EGI-Engage is co-funded by the Horizon 2020 Framework Programme
of the European Union under grant number 654142

EGI-Engage

Outline

1. Accounting System Overview
2. Recent Updates
3. Data Set Accounting
4. Cloud Accounting
5. Big Data Tools
6. Storage (Space) Accounting
7. GPGPU/Accelerator Accounting

Accounting System Overview

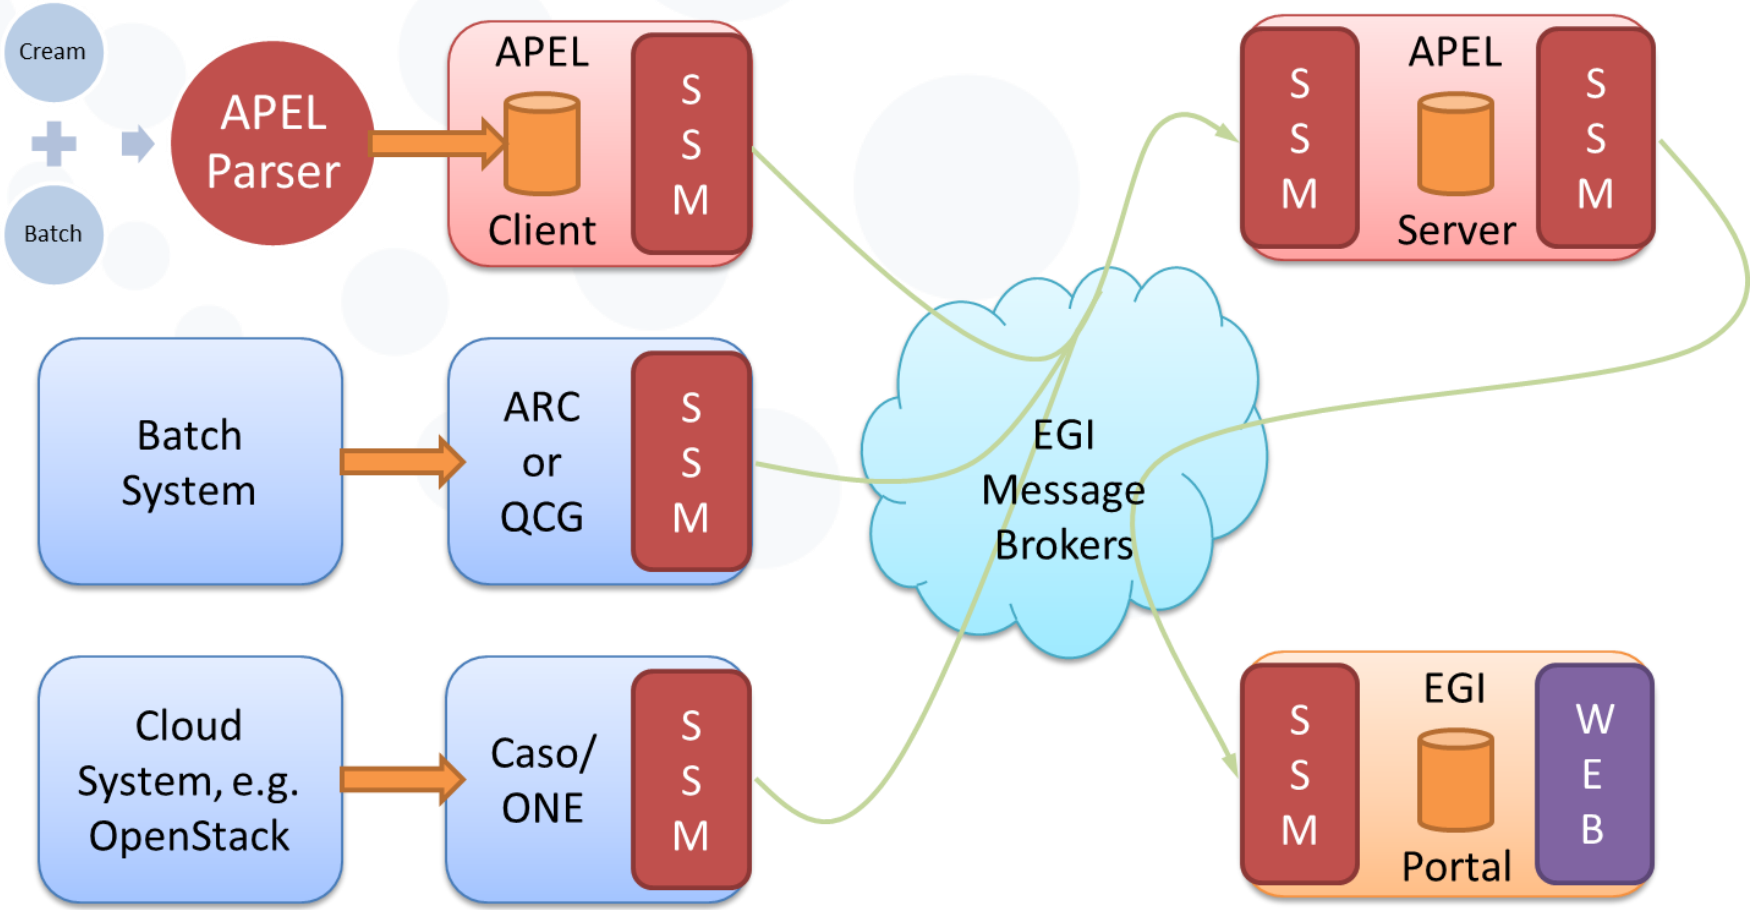

Recent Updates

- Software Releases:
  - APEL version 1.5.1
    - Bug fixes and support for new version of Torque
  - SSM version 2.1.7 released simultaneously
    - Bug fixes
  - Not in UMD3 yet
- Team changes
  - ~~Stuart Pullinger~~, John Gordon (till September 2016)
  - Adrian Coveney, Greg Corbett
- Support is via GGUS
  - apel-admins@stfc.ac.uk as a last resort

Data Set Accounting

- Data Set = a logical set of files which may exist in several places at once and to which it is possible to assign some form of persistent unique identifier
- Assists site and experiment administrators
  - make decisions about the location and storage of data sets to make more efficient use of the infrastructure
  - assist scientists in assessing the impact of their work
- Clear interest from the surveyed communities
  - user-access, data set PID logging, and recording transfer operations
- Many systems with no consistent approach to recording usage
- Will investigate related projects eg. WLCG & OneData

Cloud Accounting 1

- New fields in the cloud accounting record (V0.4)
  - BenchmarkType
  - BenchmarkValue
  - PublicIPCount
  - CloudComputeService
  - Ready to go into production
- Display of CPU counts in the Portal
  - Data is in the central repository, but not in the summaries sent to the Portal
  - Some development needed to change this
  - Will be moved to production with changes above

Cloud Accounting 2

- Long-running VMs
  - Currently, all of a VM's usage is assigned to the month when the VM started
  - Long-running (multi-month duration) VMs are not properly accounted for
  - Solution
    - After several attempts and avoiding changes to client software & message format
    - Calculate a 'MeasurementTime'
      - StartTime + SuspendDuration + WallTime
    - Assign usage based on last MeasurementTime of the month
    - Initial version to be released in April for testing and feedback

Big Data Tools

- Central processing stage for the CPU accounting data takes many hours
- Big Data tools provide an opportunity to improve performance and resilience
- Report completed: Analysis on Techniques to Manage Big Data on the EGI Accounting System
- 3 Broad categories:
  - Optimised use of the MySQL database
  - Replacement with Apache Hadoop and the Hadoop Distributed File System
  - Replacement with a time series databases
- Testing of technologies to start soon

Storage Accounting

- Accounting data extracted from DPM and dCache
- Sends StAR records using SSM
- Sites send instantaneous reading every day
- Assumed valid for previous 24 hours (under discussion)
- How to summarise? GB/days?
- Have identified the issues that need resolving to get correct storage accounting
- Once resolved, we will aim to get more sites publishing
- Instructions: https://wiki.egi.eu/wiki/APEL/Storage

GPU/Accelerator Accounting: Issues

- GPUs are usable by multiple users/jobs
  - batch systems do not attribute usage to a job/user.
- On the other hand, GPUs are attached to cloud VMs in the hypervisor
  - only attached to one VM at a time
  - for the lifetime of the VM?
  - removes the multiple user issue.
- Cloud systems currently return wallclock time only
  - Ie. No separate CPU time
    - (We hope that this can be improved.)
  - GPU Time == WallClockTime ??
- More meaningful to attempt cloud GPU accounting first.

What's needed for accounting

- Batch systems should report GPU usage in the batch logs.
  - APEL would then parse the logs files to retrieve the data.
- Or GPU monitoring recording usage in a database
  - with job or VM identifier,
  - will enable the APEL client to join it with existing data
  - create an extended Usage Record.
- The existing cloud extraction tools oneacct and cASO can be extended to include cloud GPU usage if a GPU expert can identify the relevant fields.
- The accounting portal would define new views to display GPU usage in a similar way to existing CPU views.

Accounting for Locked Resources

- Large memory jobs can effectively block CPUs
  - The node does not have enough memory to run other jobs on spare CPUs
- Accounting approach – rejected?
  - Account for CPUs effectively locked rather than just CPUs used
  - Unlikely that this is simple to achieve ie. not confident data is in the batch logs alongside the accounting data
  - But could make an interesting project!
- Improved utilisation approach
  - Ensure systems are configured to combine small + large memory jobs

Thank you for your attention.

Questions?

www.egi.eu