



Accounting

John Gordon (STFC)

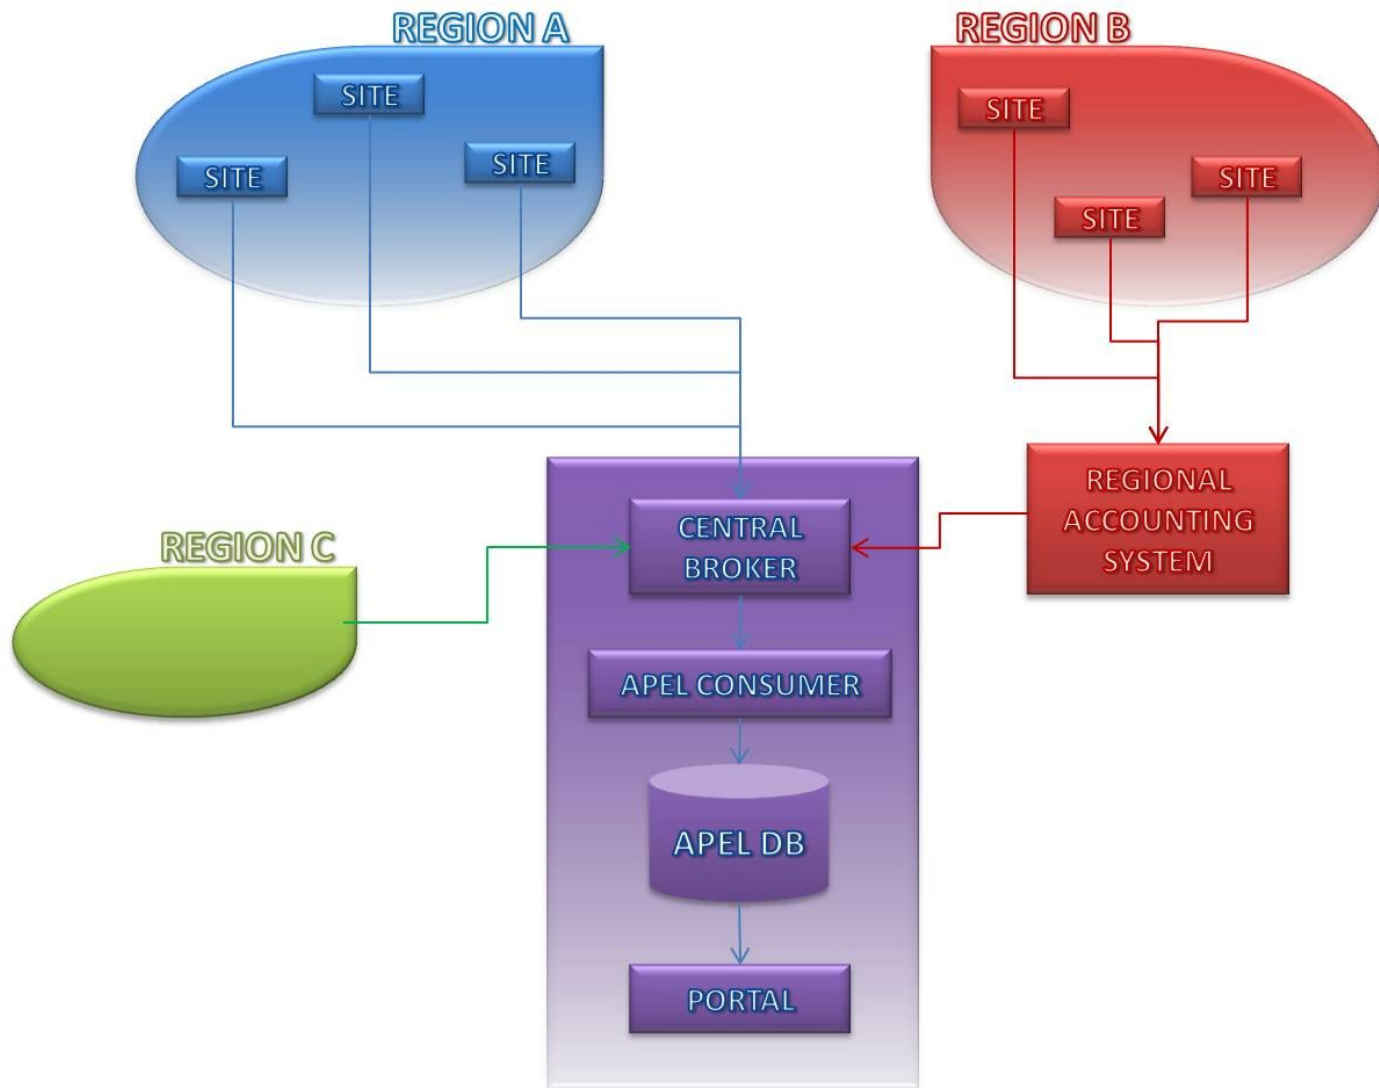
APEL PT Leader

Outline

- APEL in EGI
 - Task Force
 - CPU
 - Storage
 - Others
-
- Talk about the issues and then how EMI is involved

APEL

- APEL is the core of accounting in EGI
- It receives accounting information for all EGI(*) VOs + WLCG from the rest of the world.
- Data is published to the central repository by :
 - Sites running glite-APEL (EMI)
 - Grids running other accounting solutions
- A third method is under development in EGI: Distributed APEL Repositories in NGIs
 - gLite-APEL within the NGI and then the NGI repository publishes to the central on.
- * - +/- ARC & Unicore



Task Force GAP Analysis

- *Although this task force is in EMI we intend to look at the coverage of accounting solutions by site/user community and not try to fill the matrix of possible middleware component interoperation.*
- glite sites – glite-APEL or replacements
- ARC in NDGF – SGAS
- ARC in other countries - ??
- DEISA – own accounting
- PRACE - ?
- Unicore - (Poland has developed a solution which publishes to a central database)
- Globus – discussing with IGE for them to support a RUS Client.

CPU Requirements

- **MPI** the current cpu solution is based on a single batch job so in theory should deliver the correct cpu usage for batch MPI jobs.
- However the wallclock needs agreement and the number of cpus/cores needs recording. Do all batch systems support this information?
- Discussing with EGI MPI WG
- May require a schema change

Storage

- General VO requirement for grid-wide view of storage use by a VO.
- Management requirement (eg WLCG and NGI) to see installed capacity at sites to evaluate compliance with MoUs and other commitments and agreements.
- Most (all?) of the relevant information is published in a Glue1.3 profile defined by WLCG.
- gstat gathers this and aggregates by site/VO but does not keep a persistent record
- EMI Data Group looking at defining a Usage Record(UR) for storage use. EMI components would then implement this and APEL could collect the records.
- Useful meeting of OGF-UR WG at OGF30 in October
 - Agreed not to cut record/file; record snapshots valid at a point in time, not integrals

Clouds/Virtualisation

- There are requirements to record the usage of virtual machines
 - both as a record of user share of the physical resources
 - and for the cloud-like deployment of images.
- Be aware of double counting if the application running in the image is also accounted.
- Discussing with Stratuslab

APEL usage collection

- APEL collects data in 2 steps:
 - Parses batch logs on CEs and publishes to a local database and then
 - Collects records from local database and publishes them via AMQ to central repository
- The central repository can handle multiple record types and process differently.
- APEL is thus in a good position to receive other types of accounting record.
- E.g. If the SE collects storage information and stores it locally then glite-APEL will only needs small changes to gather and publish these records as a different type in a different topic.
- The central repository could either read all topics and process them differently or split the load between multiple instances reading different topics.



Thank you

EMI is partially funded by the European Commission under Grant Agreement INFSO-RI-261611

Use Insert
Header &
Footer to set
this field

Use Insert Header & Footer to set this field