

# APEL CPU Accounting in the EGEE/WLCG infrastructure

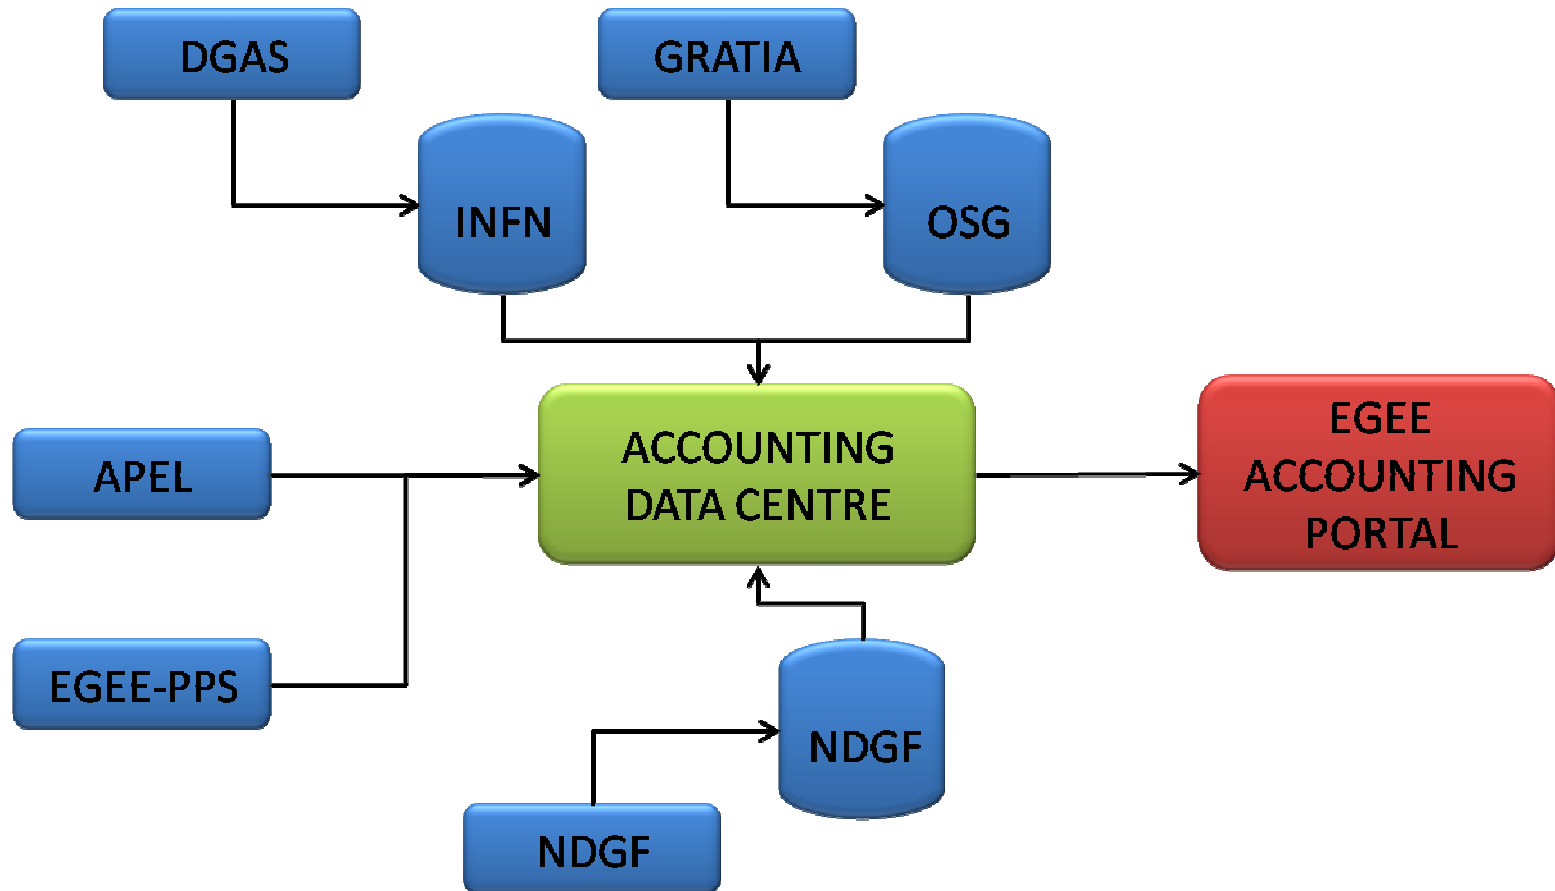
*Cristina del Cano Novales,*

*John Gordon*

*STFC - RAL*

- **Overview**
- **APEL Client**
- **Data Transportation**
- **Accounting Data Centre**
- **EGEE Accounting Portal**
- **APEL SAM tests**
- **Standards**
- **Status**
- **Future Plans**

- **APEL (Accounting Processor for Event Logs)**
- **Data collection and reporting services**
- **Large centralised Database**
- **Collects and aggregates CPU usage information from sites across the Grid**

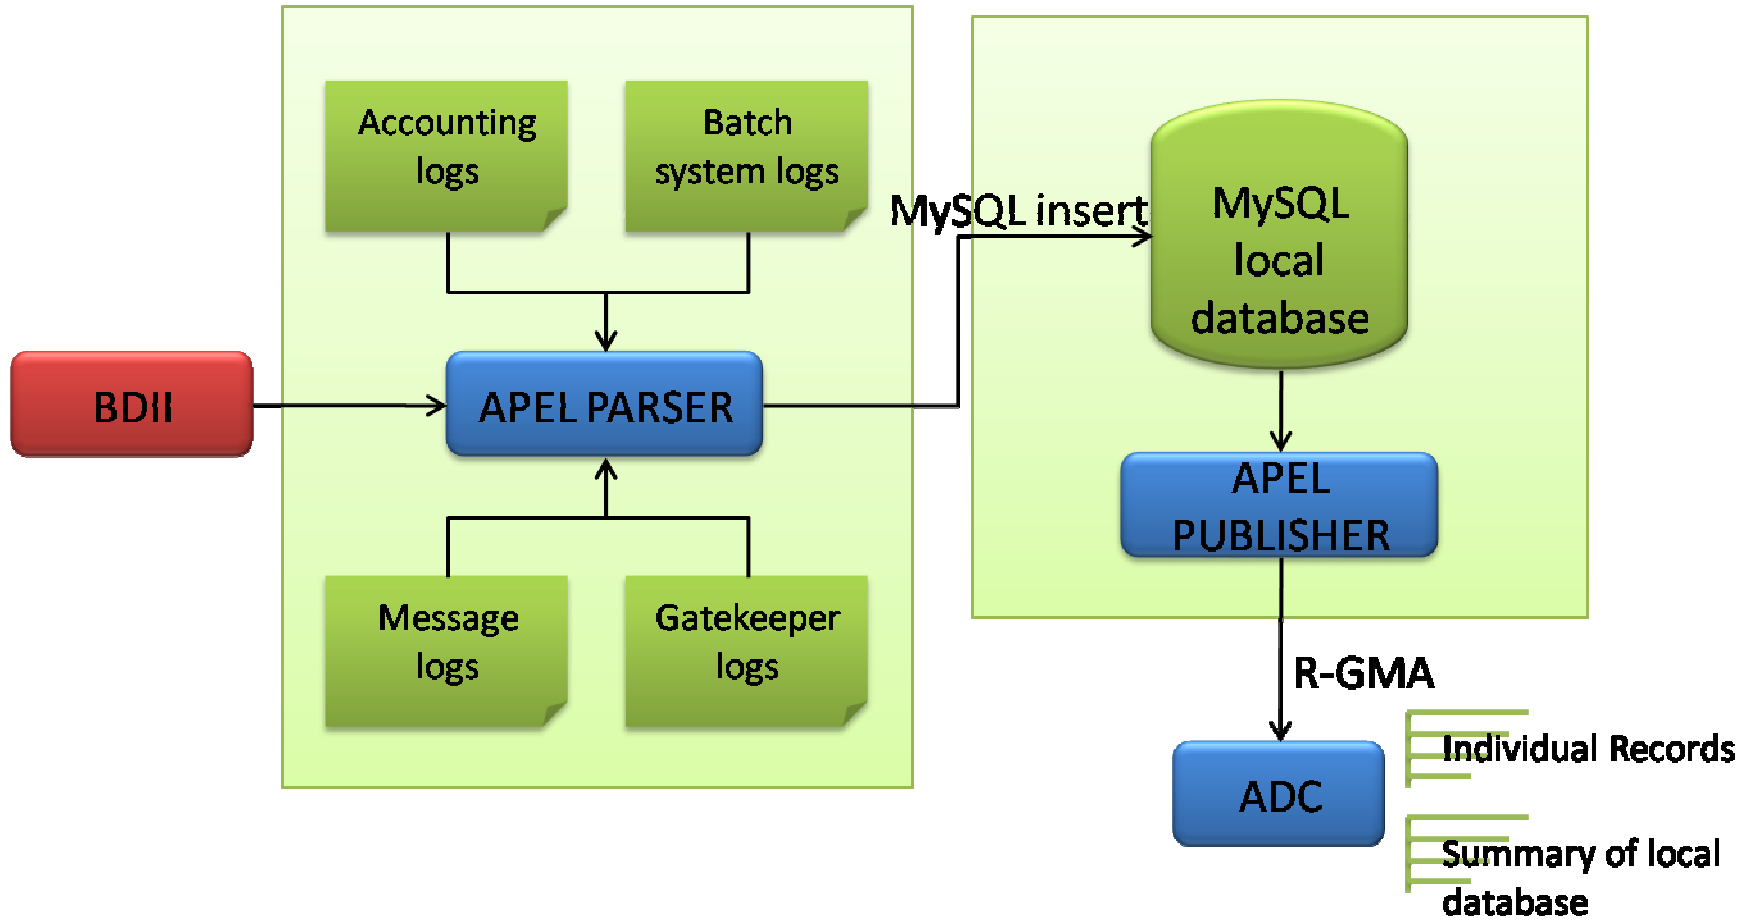


- **Storing ~200M individual job records since 2004**
- **And more than 100M records in aggregated summaries from other Grids (OSG/NDGF)**
- **442 different sites**
- **673M CPU hours – 28M days – 76000 years**

- **Log processing application**
- **Interprets system log files (gatekeeper and batch system logs) to produce accounting records**
- **Currently supports PBS, LSF, SGE, CONDOR, but could be extended to support other systems**
- **APEL collects usage information after the job was completed**
- **Distributed as part of the gLite Middleware**

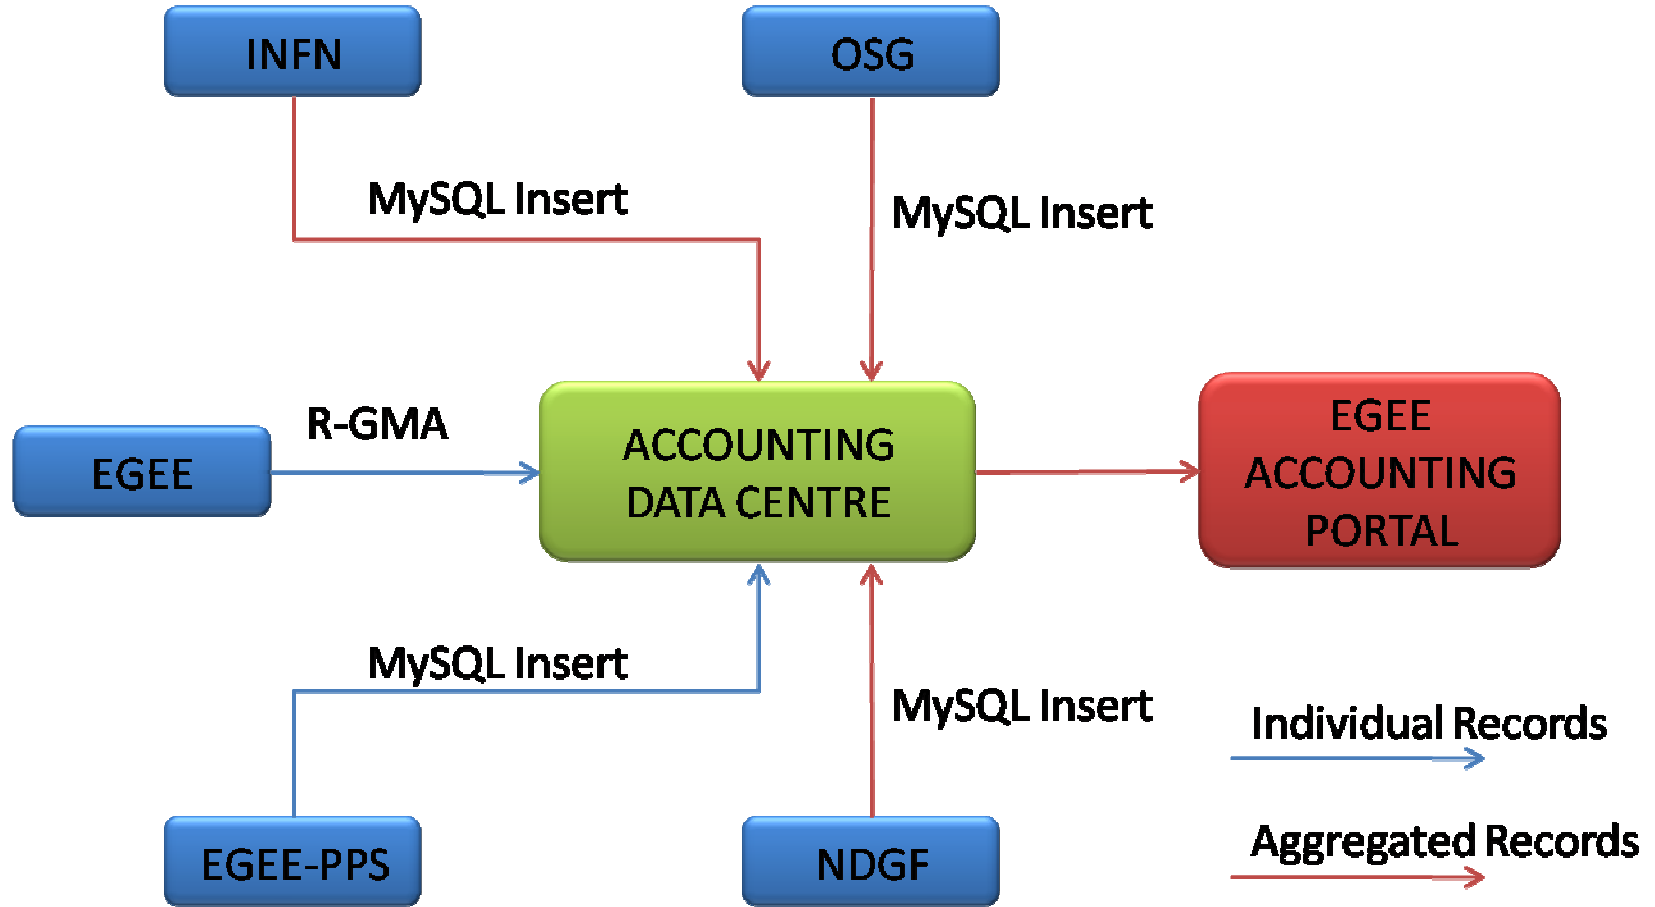
CE

MON

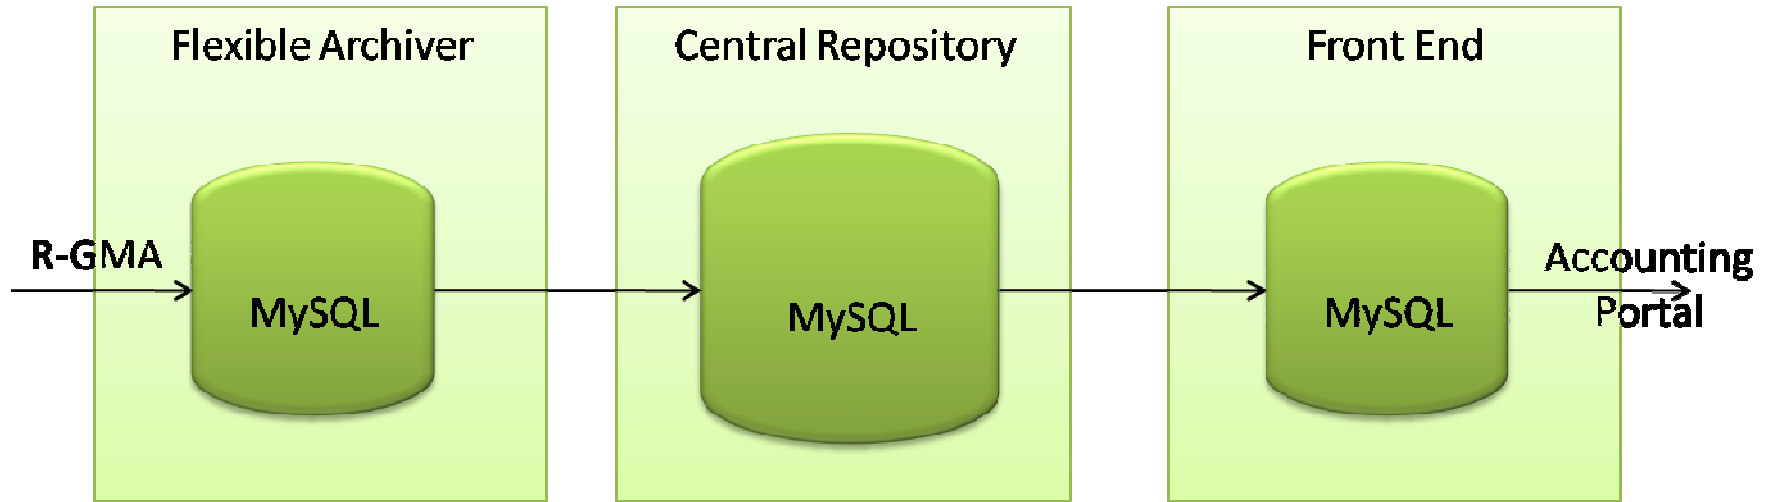


- **Currently using two different interfaces**
  - R-GMA (Relational Grid Monitoring Architecture)
    - Majority of EGEE sites publishing via APEL-RGMA
    - Some EGEE sites using own sensor and APEL publisher (with R-GMA)
  - Direct MySQL insertion
    - OSG – Gratia
    - INFN – DGAS
    - NDGF – SGAS





- **Receives records from R-GMA**
- **Processes and stores the accounting records produced by the grid resources, including:-**
  - Decryption of the UserDNs
  - VOMS-level:
    - Extraction of VO, primary Group and Role from the UserFQAN
  - Normalisation:
    - For each tuple a *normalised* CPU and Wall time is determined based on the SpecInt2000 value and the raw CPU and Wall times
  - Aggregation:
    - Anonymous and User-level summaries are generated.
  - Encryption:
    - The User-level summaries are encrypted before they are sent to the CESGA Accounting portal. Access to these summaries is controlled using SSL and ACLs.



- gLite MON box
- Receives data from R-GMA
- Data stored for 3 days
- Open access so personal data encrypted

- Main repository for Accounting Data
- “Offline” – Not accessible
- Contains all accounting data since 2005
- Archiving of records dependent on Policy document being drafted

- Contains summaries for the Accounting Portal
- Contains dedicated tables for OSG, INFN, NDGF
- Personal data encrypted with Portal’s public key
- Creates and publishes APEL SAM tests

Hierarchical Tree

- Tier1
- Tier2
- Countries
- EGEE
  - Production
    - AsiaPacific
    - CentralEurope
    - CERN
    - France
    - GermanySwitzerland
    - Italy
    - NorthernEurope
    - Russia
    - SouthEasternEurope
    - SouthWesternEurope
    - UKI
  - PPS
  - OSG
  - UNREGISTERED
  - VO\_Discipline
  - VO\_Metrics

EGEE View → Production

Print Page

<b>Data to graph:</b>	Norm. Sum CPU	Normalised CPU time to a reference value of 1000 SpecInt2000		
<b>Period:</b>	Start year: 2008	Start month: 7	End year: 2009	End month: 6
<b>Groupings:</b>	Show data for: REGION	as a function of: VO		
<b>VO Groups:</b>	<input checked="" type="radio"/> LHC <input type="radio"/> TOP 10 <input type="radio"/> ALL <input type="radio"/> Custom <input type="checkbox"/> Group the rest of VOs in a new category			
<b>Chart:</b>	Type: ACCUM BAR	Scale: LINEAR		
<b>dteam VO:</b>	<input type="checkbox"/> Exclude dteam and ops VOs jobs information			

Refresh

PRODUCTION Normalised CPU time by REGION and VO.  
LHC VOs. July 2008 - June 2009.

The following table lists the certified sites that have not published accounting data to GOC during the last 3 months. This probably indicates a major problem in the accounting system of the site so **listed sites are encouraged to take the appropriate measures to correct it**

[http://www3.egee.cesga.es/gridsite/accounting/CESGA/egee\\_view.php](http://www3.egee.cesga.es/gridsite/accounting/CESGA/egee_view.php)

- Provide monitoring for APEL for Production Sites
- Two tests provided
  - APEL-pub: Critical test. Checks date of the latest record published.
    - Older than 7 days => Warn
    - Older than 31 days => Error (Site notified)
  - APEL-sync: Compares the number of records on the Central Database with the number of records on the local database.
    - >10 records diff => Warn
    - > 100 records diff => Error

### APEL Publication Test

- Displays the last time the site published accounting data to the GOC.
- A warning / error is raised if the site has not published accounting data for 7 / 31 days, if a published data for 31 days, which usually signifies a problem with APEL or RGMA services.
- Information about [APEL Tests and Recipes](#)
- Contact: [APEL-SUPPORT@JISMAIL.AC.UK](mailto:APEL-SUPPORT@JISMAIL.AC.UK)
- lastBuild : 2009-06-01 12:55:14.24

RAL-LCG2			
ExecutingSite	MeasurementDate	MeasurementTime	Publication Status
RAL-LCG2	2009-06-01	10:47:31	OK [ last published 1 days ago: 2009-05-31 ]

### APEL Synchronisation Test

- A comparison is made between your local APEL database, and the data that you have published to the GOC.
- Major differences are flagged with FAIL.
- Information about [APEL Tests and Recipes](#)
- Contact: [APEL-SUPPORT@JISMAIL.AC.UK](mailto:APEL-SUPPORT@JISMAIL.AC.UK)
- lastBuild : 2009-06-01 12:55:08.46

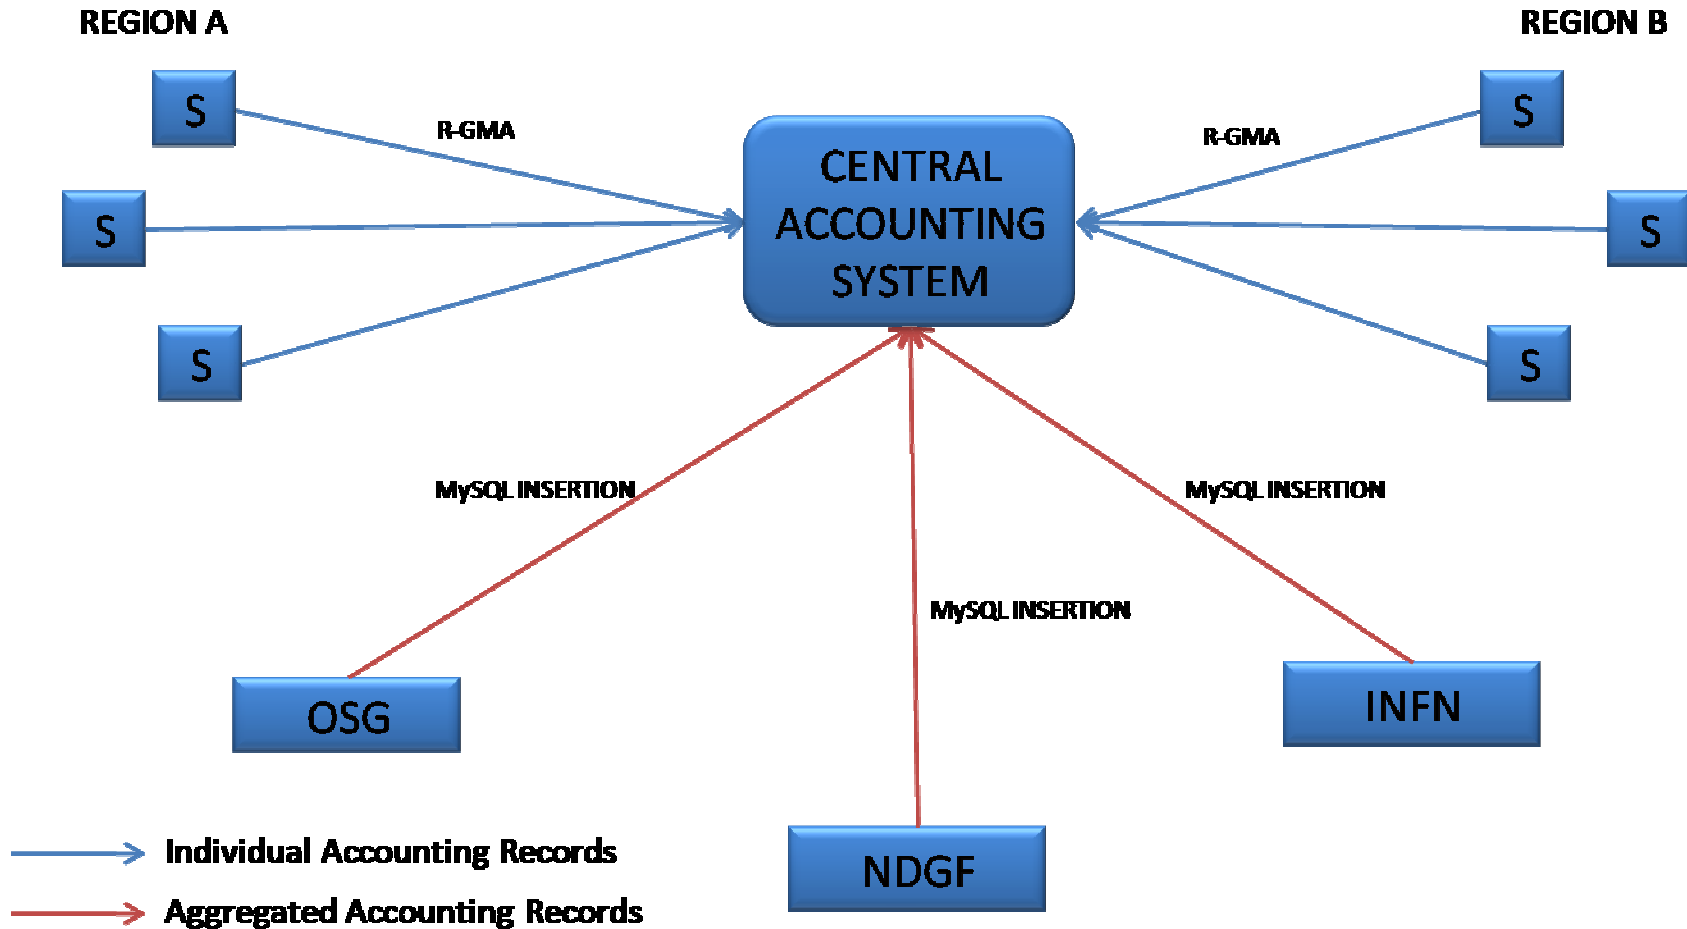
RAL-LCG2				
RecordStart	RecordEnd	Record Count In Your Database	Record Count What You Published	Synchronisation Status
2009-05-01	2009-05-31	336956	334359	OK [ last published 1 days ago: 2009-05-31 ]
2009-04-01	2009-04-30	454203	454203	OK
2009-03-01	2009-03-31	294314	294314	OK
2009-02-01	2009-02-28	365157	365157	OK
2009-01-01	2009-01-31	389197	389197	OK
2008-12-31	2008-12-31	10091	368331	OK
2008-11-01	2008-11-30	300924	300924	OK
2008-10-01	2008-10-31	253927	253927	OK
2008-09-01	2008-09-30	240100	240100	OK
2008-08-01	2008-08-31	375256	375256	OK
2008-07-01	2008-07-31	597993	597993	OK
2008-06-01	2008-06-30	492473	492472	OK [minor discrepancy in event numbers]
2008-05-01	2008-05-31	427533	427533	OK
2008-04-01	2008-04-30	144613	144613	OK
2008-03-01	2008-03-31	141716	141716	OK

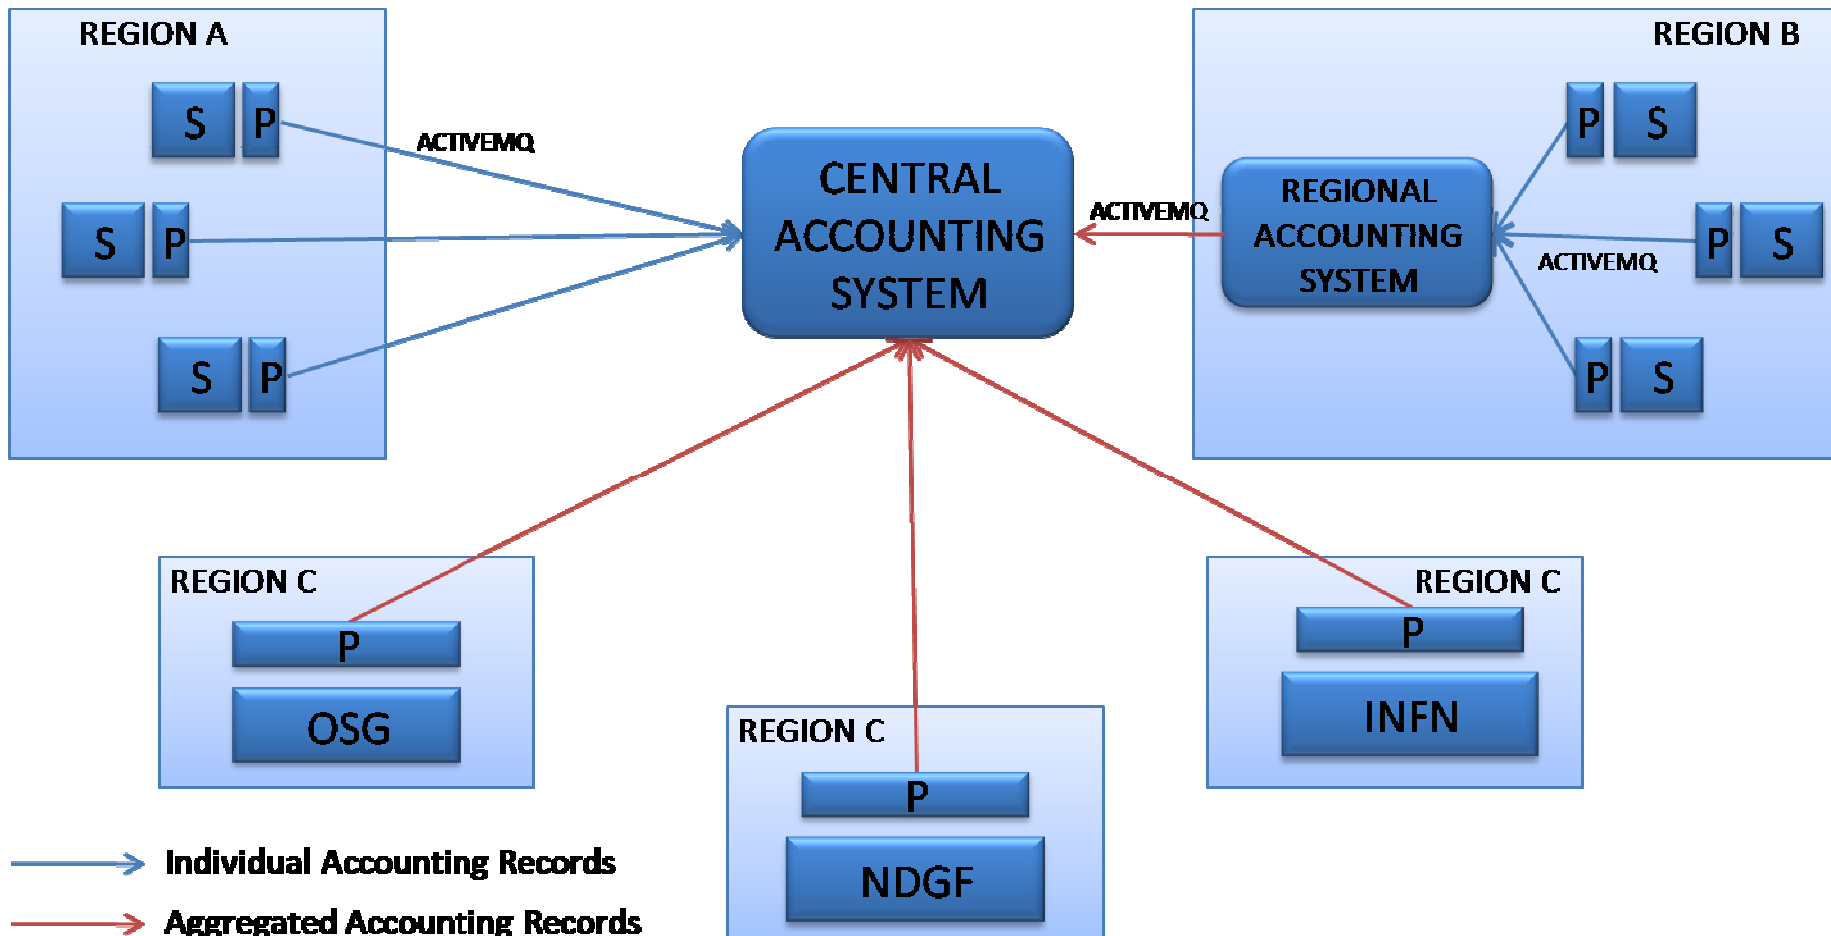
- **Sites should:**
  - 1. Measure resources using HEPSPROC06**
  - 2. \*250**
  - 3. Publish as SI00 as before**
  - 4. Set Glue HEPSPROC06 with value. This shows new benchmark has been used**
- **APEL gathers comparable data from all sites**
- **Monitoring can identify sites which have/have not changed benchmark. Raise tickets etc**
- **CESGA Portal can show usage in either (both?) benchmarks by conversion.**
- **Eventually when most sites have changed, the portal default will change**
- **When GLUE2.0 deployed, publish raw HEPSPROC06 values**
  - Or possibly reuse SI00

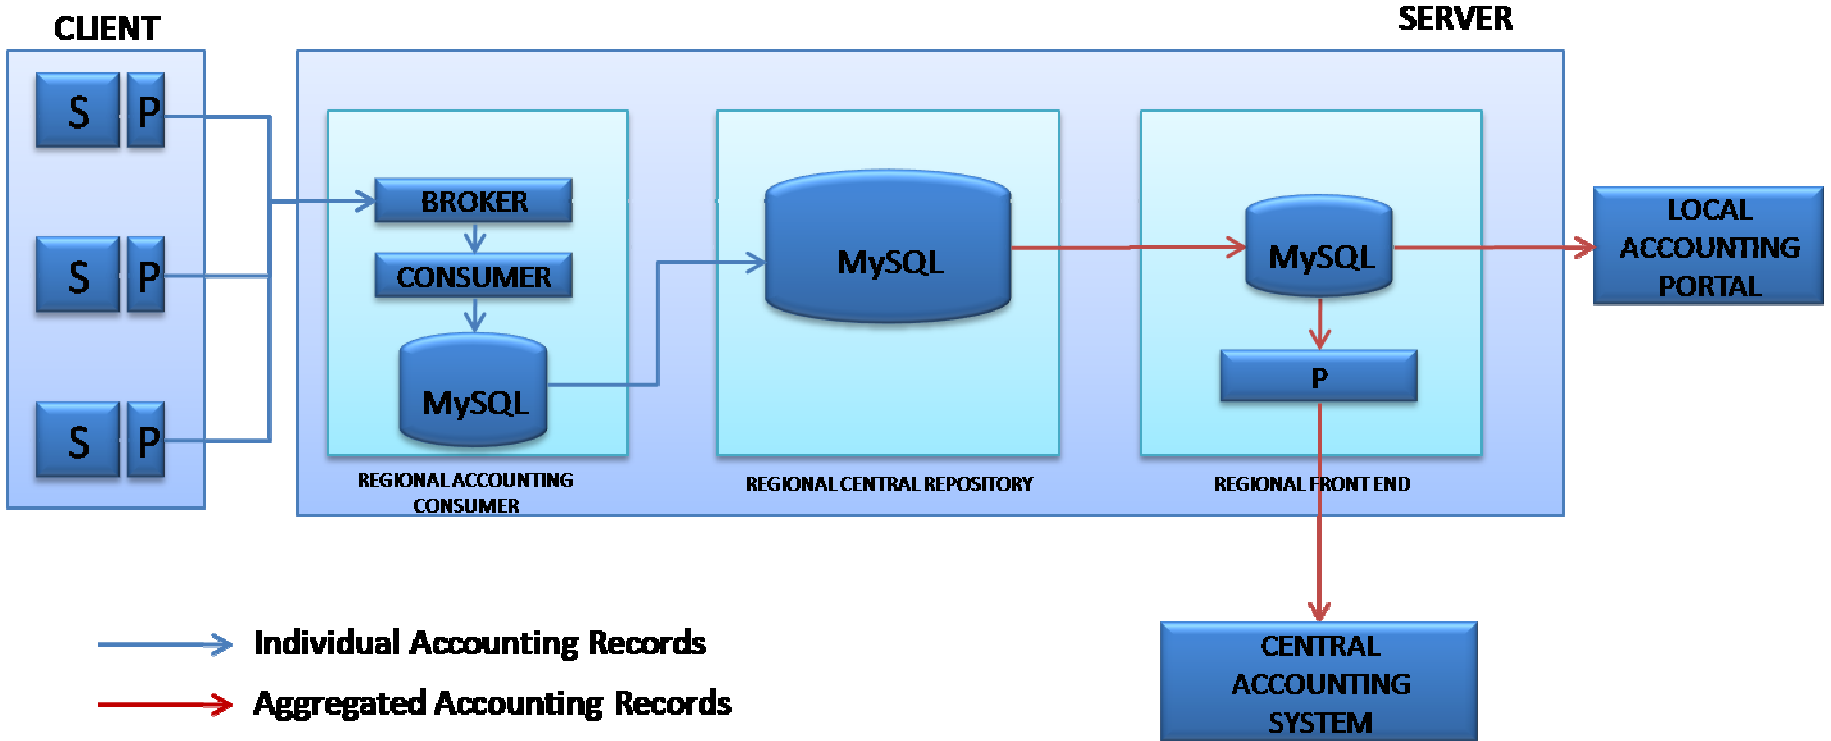
- **Main ideas:**
  - ActiveMQ to replace R-GMA as the transport mechanism
    - Interoperability with other tools
    - Expertise easily available
    - Using existing infrastructure
  - New architecture to allow regionalisation but not impose it
  - Maintain Central Repository for multi-grid/VO/user queries
  - Standard publishing methods – RUS (???)

- **Regionalisation:**
  - Use Cases:
    - Region A: APEL – non regionalised
    - Region B: APEL – regionalised
    - Region C: Other sensor – own accounting system
  - Regions can be Grids






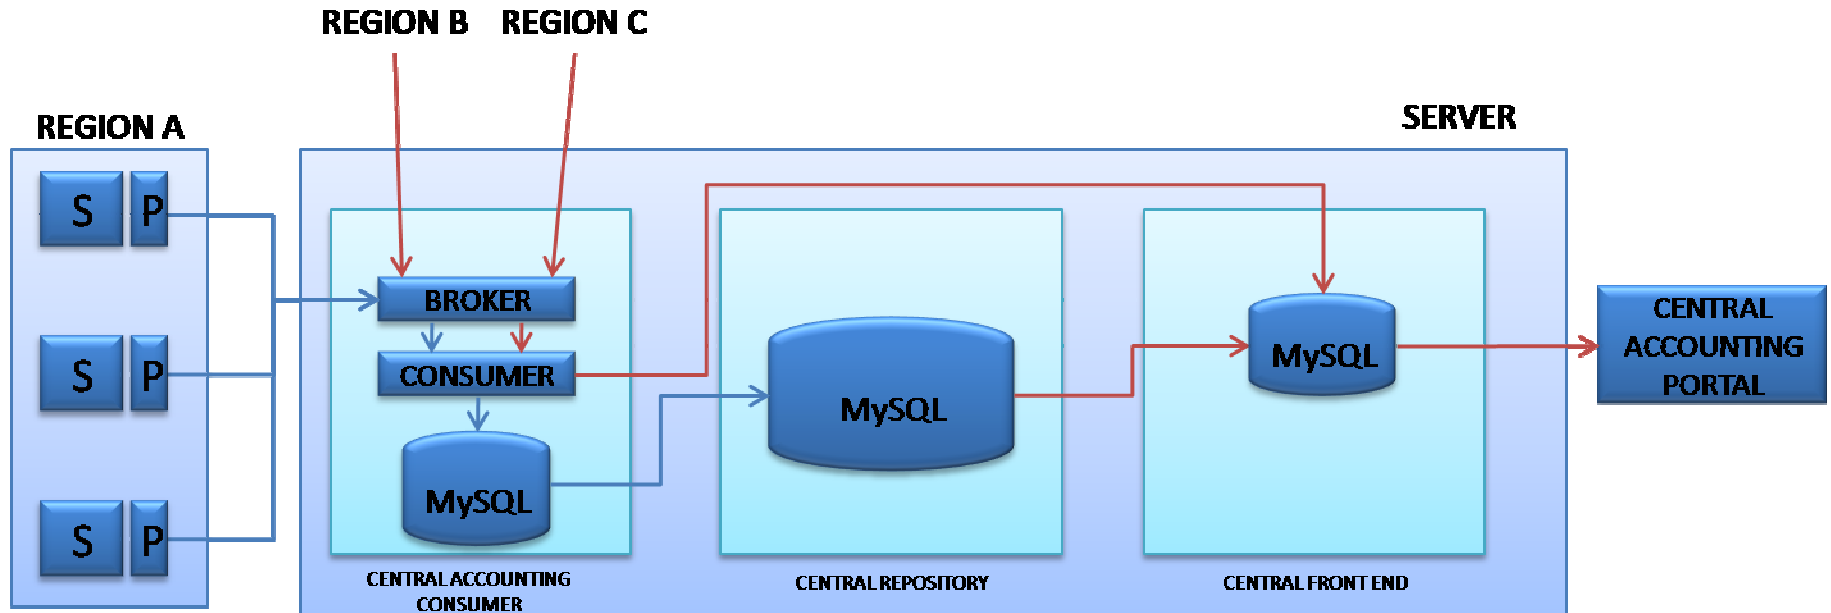




**egEE Accounting Data Centre**



- gLite MON box
- Receives data from R-GMA
- Data stored for 3 days
- Open access so personal data encrypted
- Main repository for Accounting Data
- "Offline" - Not accessible
- Contains all accounting data since 2005
- Archiving of records dependent on Reliex
- Contains summaries for the Accounting Portal
- Contains dedicated tables for OSG, INFN, NDGF
- Personal data encrypted with Portal's public key
- Creates and publishes AREL SAM files



- Individual Accounting Records
- Aggregated Accounting Records

- **Already use OGF-UR**
  - Participate in OGF WG
- **OGF-RUS as a standard interface for publishing**
  - Designed for XML
  - Existing Accounting Service use relational databases
  - General agreement on publishing
  - More difficult to implement full xpath query on relational db

- **New APEL Client built in ETICS**
- **Prototype consumer under test**
- **First external site test August**
  - Australia

- **EGEE III Plan - By end of EGEE**
  - Change to ActiveMQ
  - Regionalise Regions where desired
- **EGI plan –**
  - can distribute to NGIs
  - NGIs could implement their own accounting service and interface it like OSG, INFN, NDGF, ...

- **The infrastructure underlying APEL will change over the next year.**
- **This should result in a more flexible and resilient service**
- **The results will continue to be published through the same portal – so no changes seen to users**
- **The new infrastructure will allow national accounting repositories and portals but still allow worldwide visualisation for worldwide Vos.**



# QUESTIONS ...

