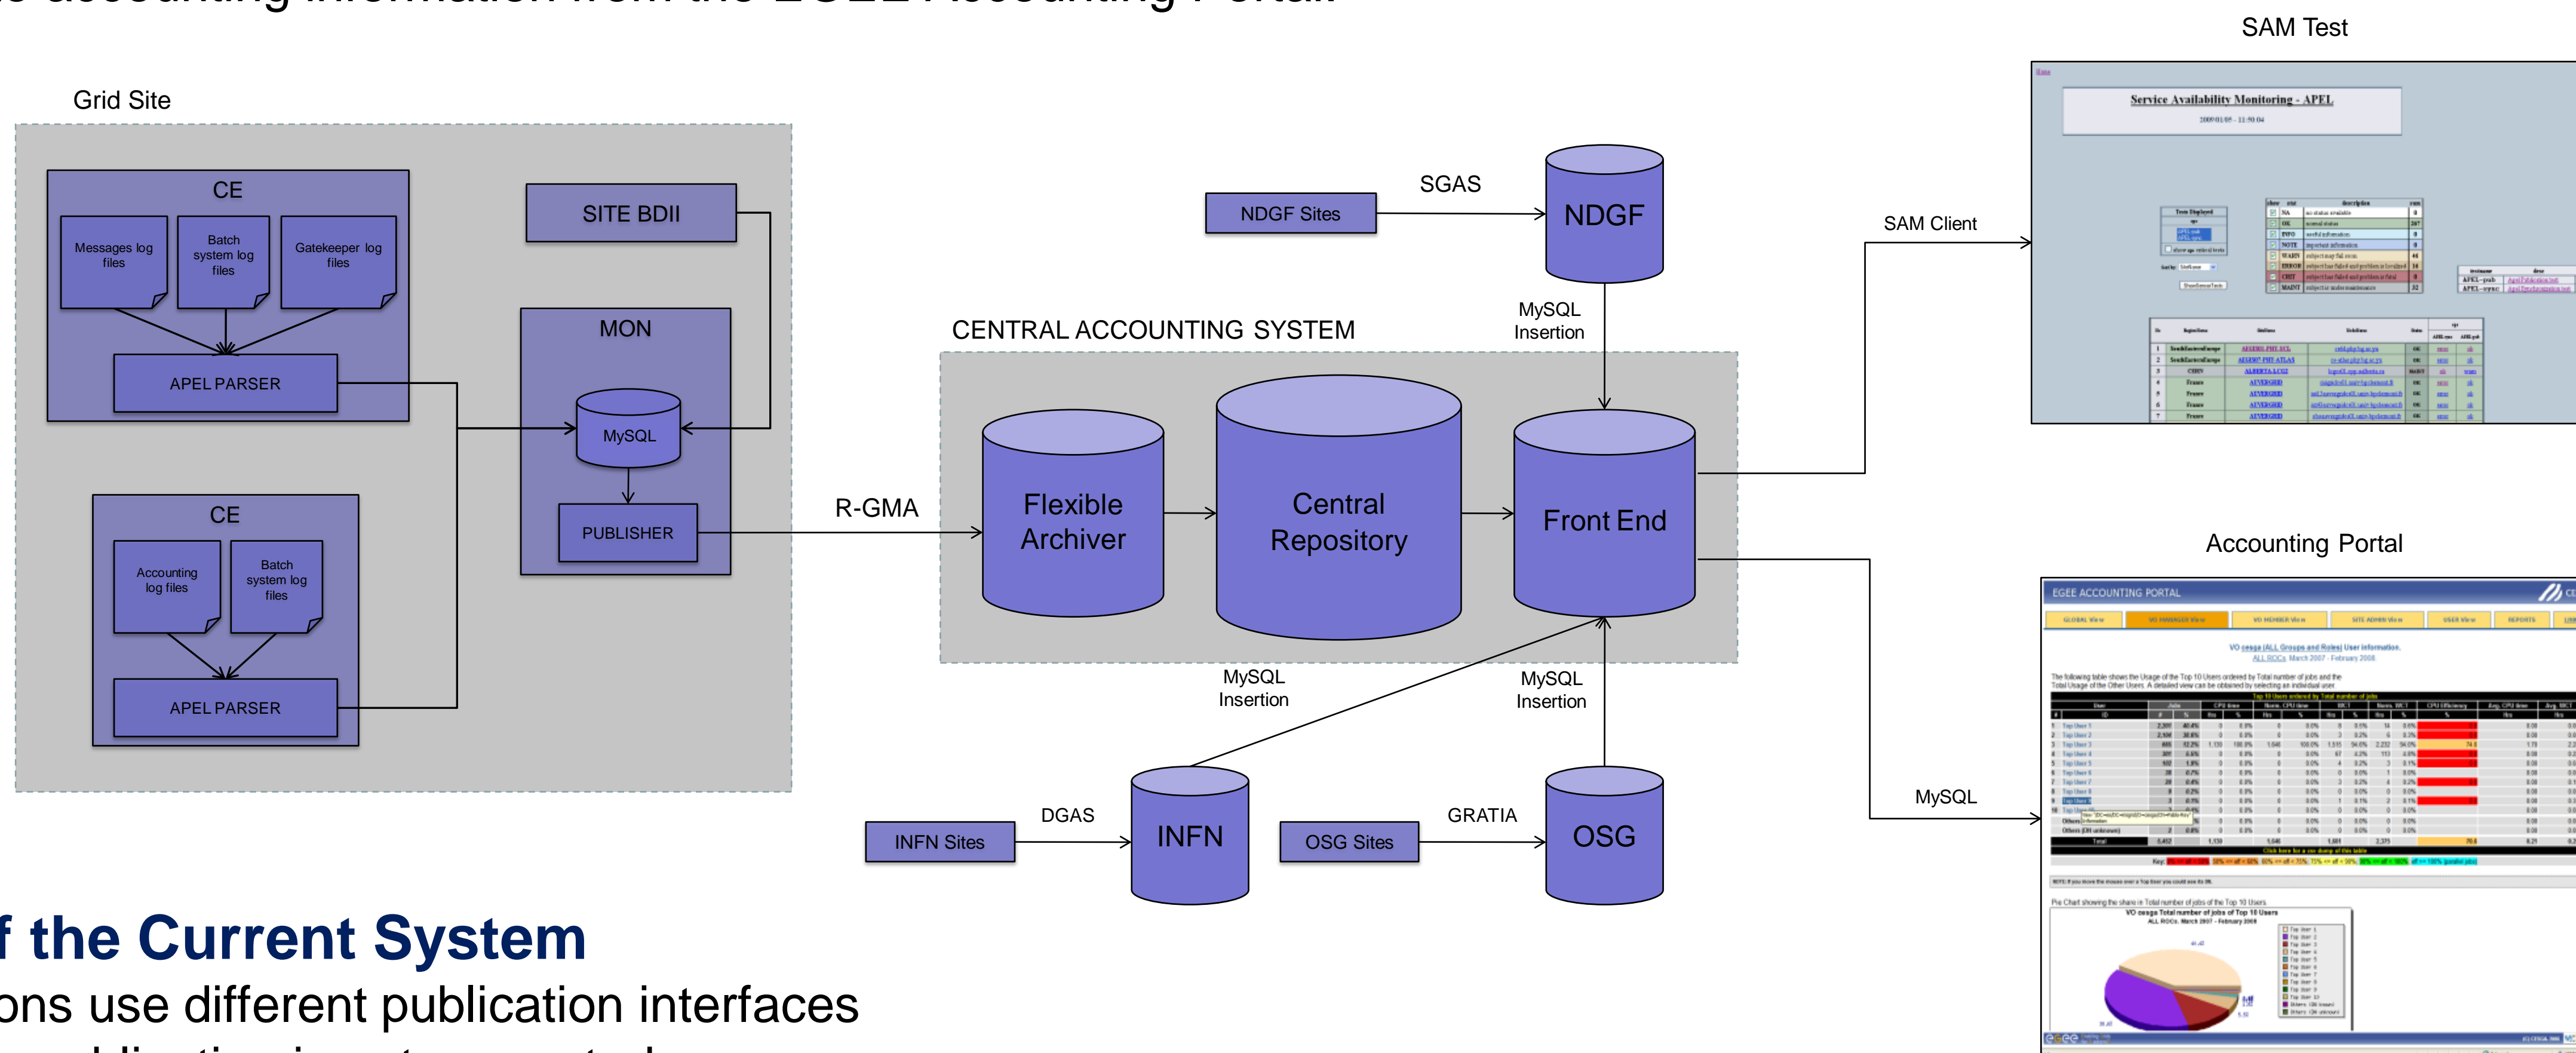


Towards a Distributed CPU Usage Accounting Infrastructure

Ming Jiang, Cristina Del Cano Novales, John Gordon, Matthew Viljoen, David Meredith
Science and Technology Facilities Council, United Kingdom

EGEE CPU Accounting System

The EGEE CPU accounting system is a large centralised database that collects, aggregates and stores CPU usage information. It contains around 170 million individual job records for more than 300 sites across EGEE, OSG and NDGF Grids since 2003. The accounting system supports access to accounting information from the EGEE Accounting Portal.

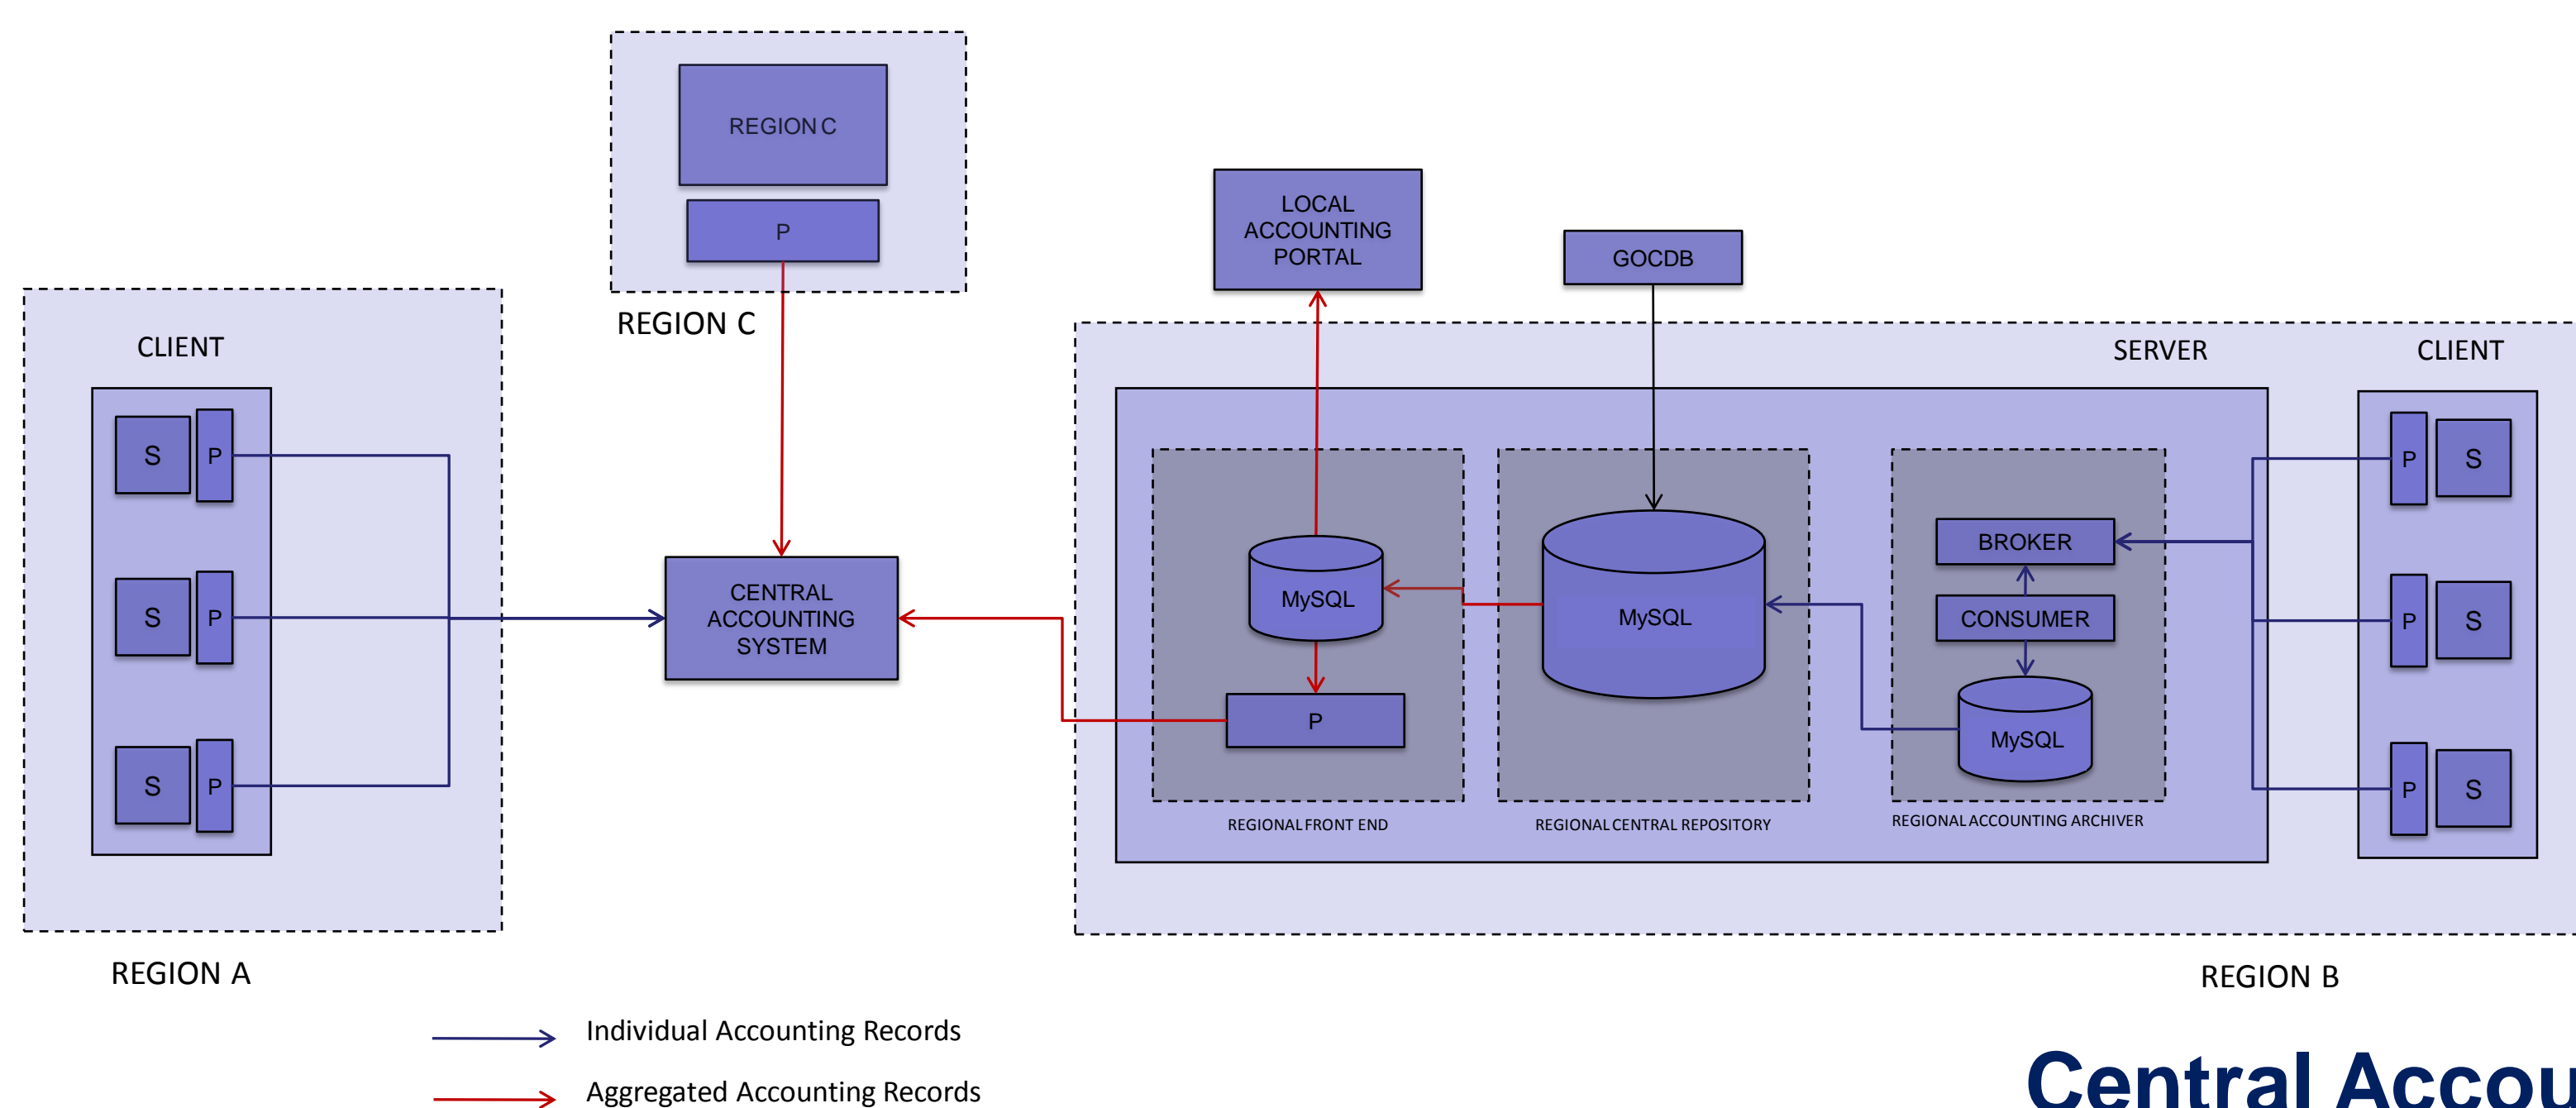


Problems of the Current System

- Different regions use different publication interfaces
- Multiple levels publication is not supported
- Lack of integration mechanism of sharing information with other Grid tools

Proposed Distributed Accounting System

A distributed accounting infrastructure is proposed based on modifications and extensions to the records transport mechanism of APEL (Accounting Processor for Event Logs) tool to support a robust, distributed, accounting capability and flexible accounting record queries across VOs. The regions shown in the diagram can be EGEE Regions, countries, or other Grid Infrastructures (like OSG).

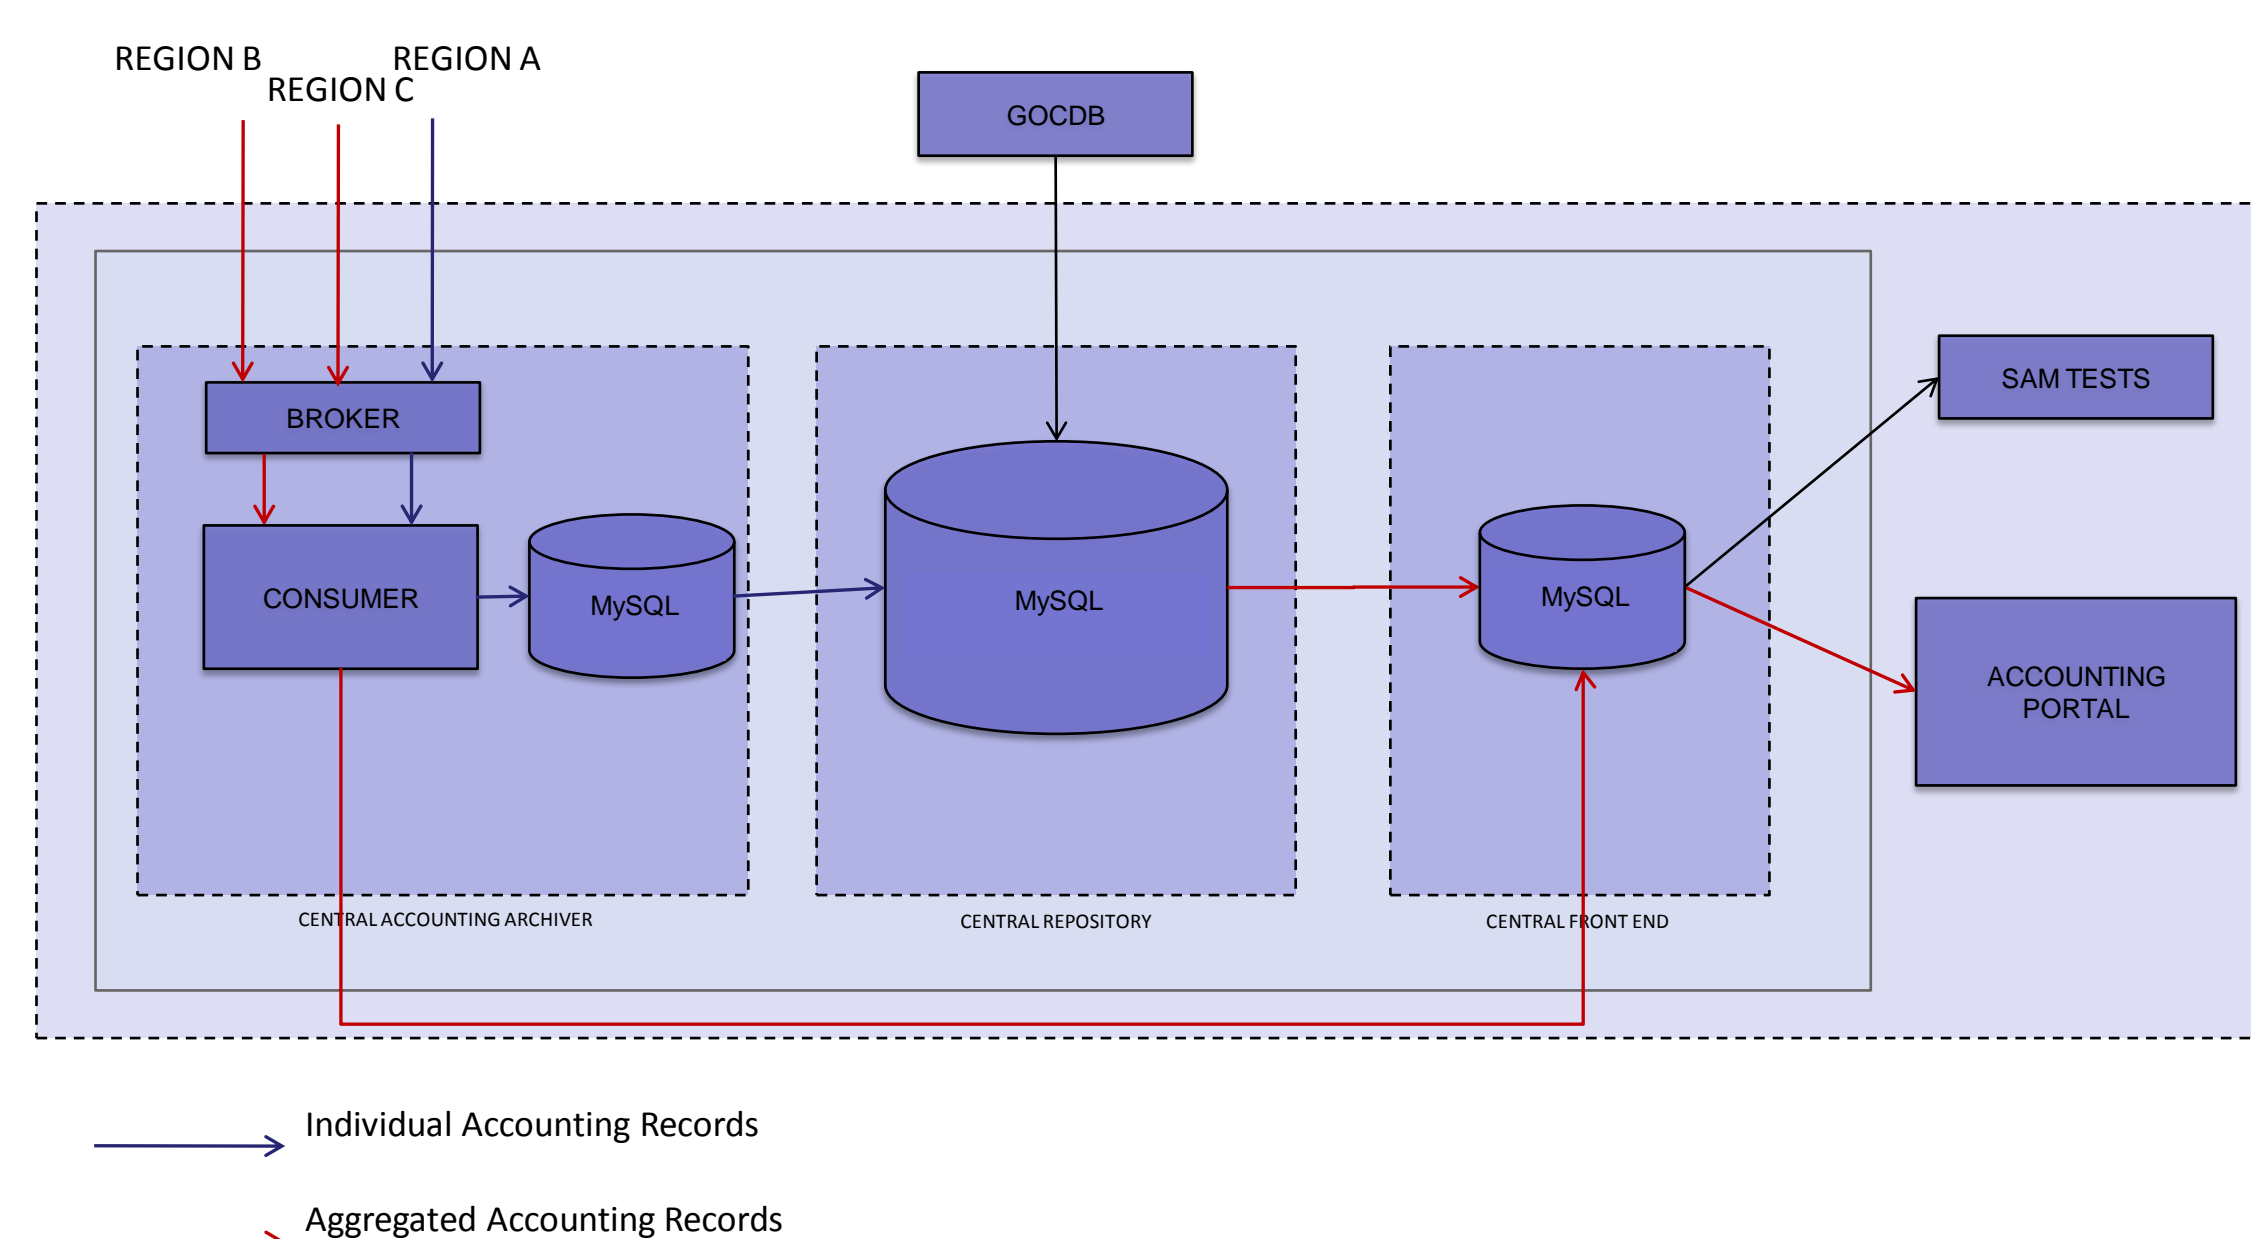


Accounting Use Cases

- Region A: APEL sensor - non regionalised
- Region B: APEL sensor – regionalised
- Region C: Other sensor - own accounting system

Central Accounting Cache

- Accounting Archiver: Message Broker, Record Message Consumer and MySQL database containing recently published records
- Central Repository: MySQL database containing all individual records and processing scripts
- Front End: MySQL database containing aggregated accounting summary records and processing scripts



Impact

- Multiple levels of publication is supported (country -> region -> central)
- Different regions use a unified publication interface
- No significant operational changes on client side publication
- A region optionally sets up its local Accounting Server or uses the Central Accounting Cache
- Potential interoperability and integration with other ActiveMQ based middleware

Messaging Model and Apache ActiveMQ Message Broker

- General topic publication and subscription messaging model
- Distributed components to publish and subscribe messages to/from a well defined topic (virtual destination and source of messages)
- Utilises the Apache ActiveMQ message broker

