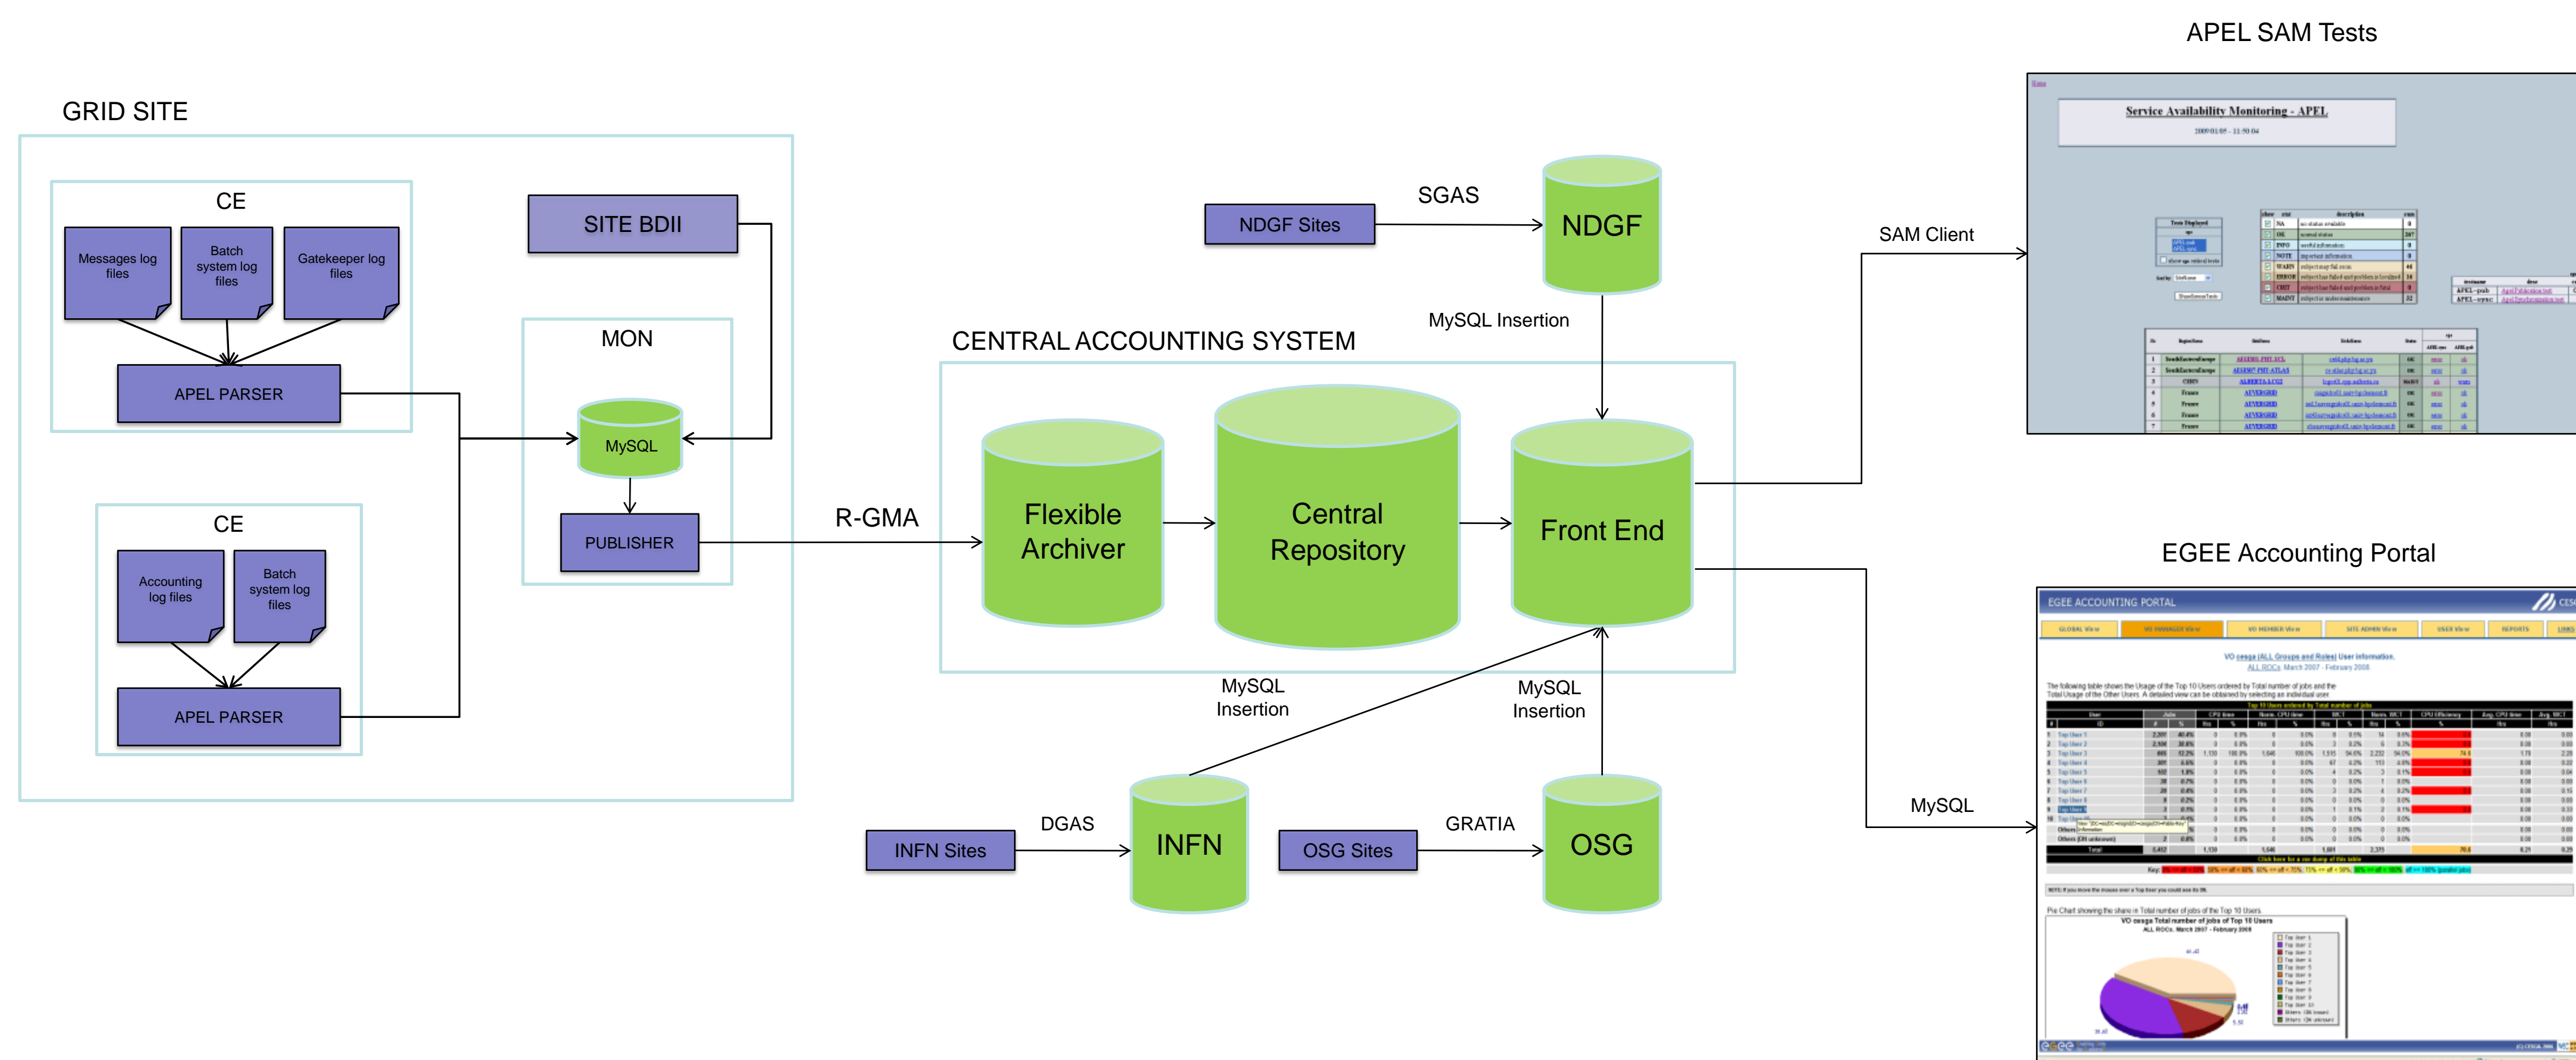


An Open and Flexible Distributed CPU Usage Accounting Infrastructure for the WLCG and EGEE Grids

Ming Jiang, Cristina Del Cano Novales, Gilles Mathieu, John Casson, John Gordon
e-Science Centre, Science and Technology Facilities Council, United Kingdom
 {ming.jiang, cristina.del-cano-novales, gilles.mathieu, john.casson, john.gordon}@stfc.ac.uk

The Centralised EGEE CPU Accounting System

The EGEE CPU accounting system in production 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.



An Open and Flexible Distributed Accounting Infrastructure

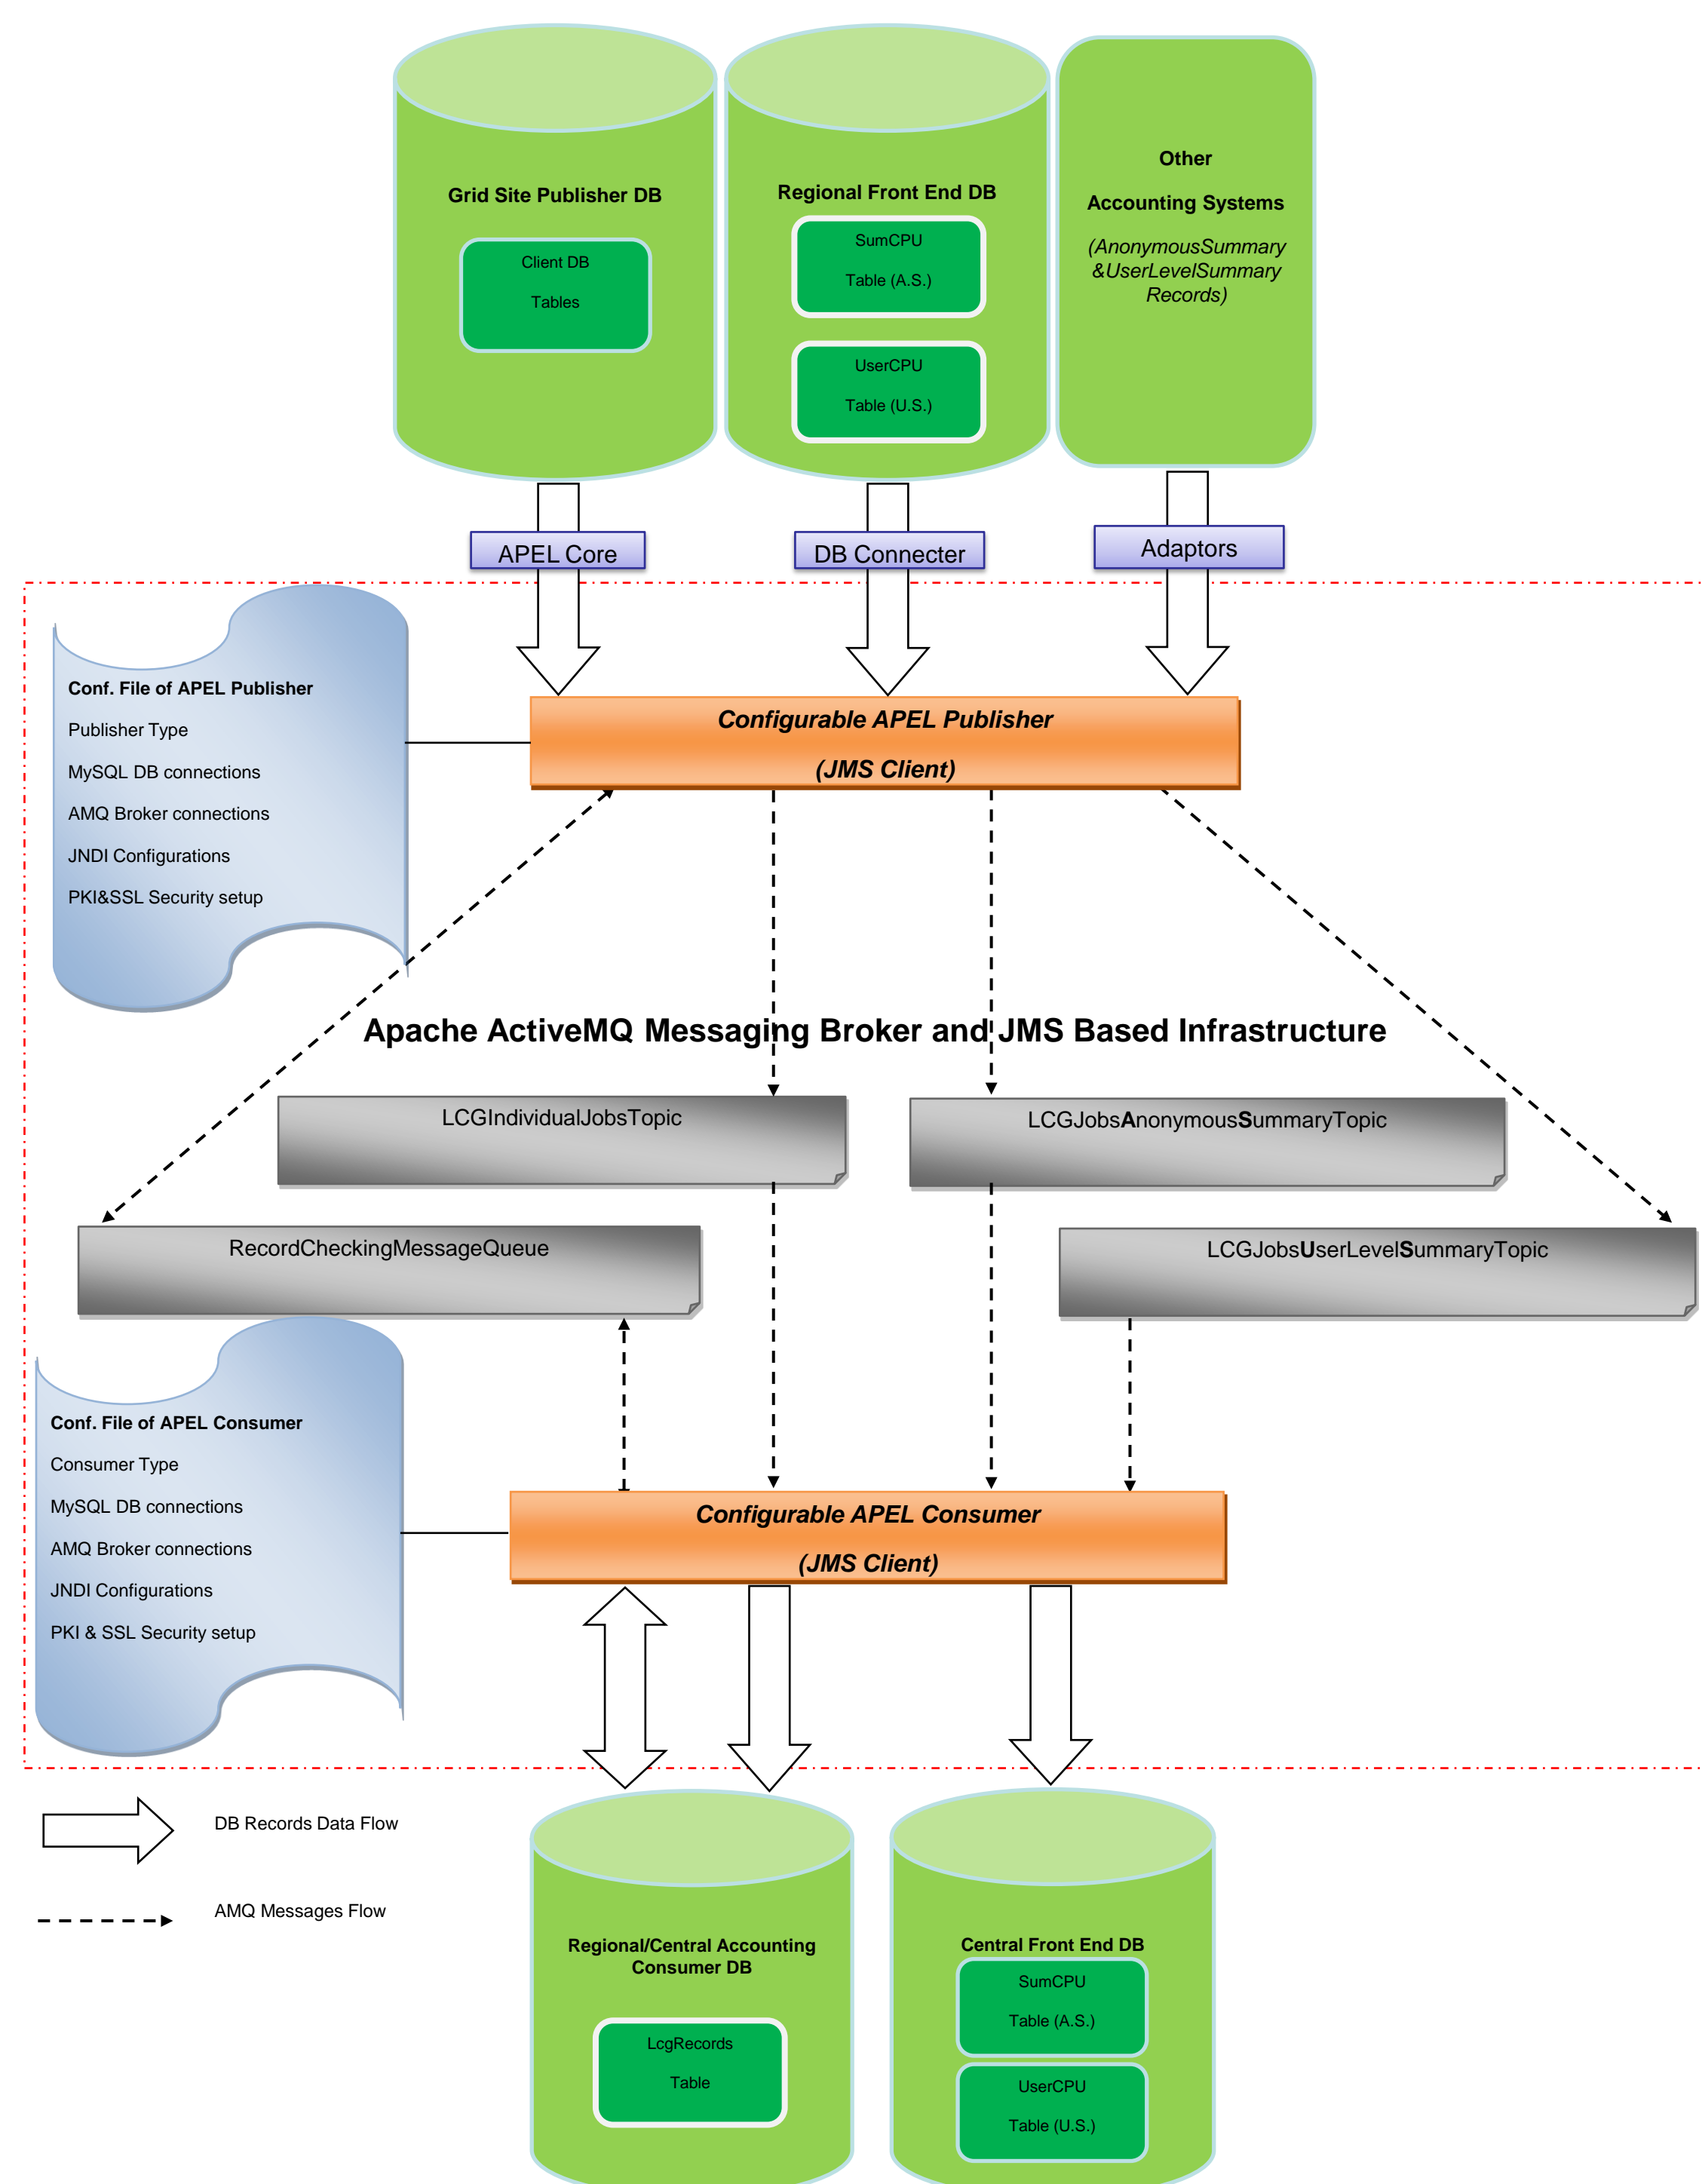
An open and flexible distributed accounting infrastructure is being designed and implemented based on modifications and extensions to the records transport mechanism of APEL (Accounting Processor for Event Logs) tool to support a robust accounting capability at a NGI level and flexible across VOs accounting records queries.

Messaging Model and Apache Message Broker Based Design and Implementation

- 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 with security setup (PKI+SSL A&A)

Impact of the New Infrastructure

- 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



A Central Accounting Cache

- Accounting Consumer: Message Broker, Record Message Consumer and MySQL database containing recently published records (~3 days)
- Central Repository: MySQL database containing all individual records and processing scripts
- Front End: APEL Publisher, MySQL database containing aggregated accounting summary records, and processing scripts

Accounting Use Cases

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

