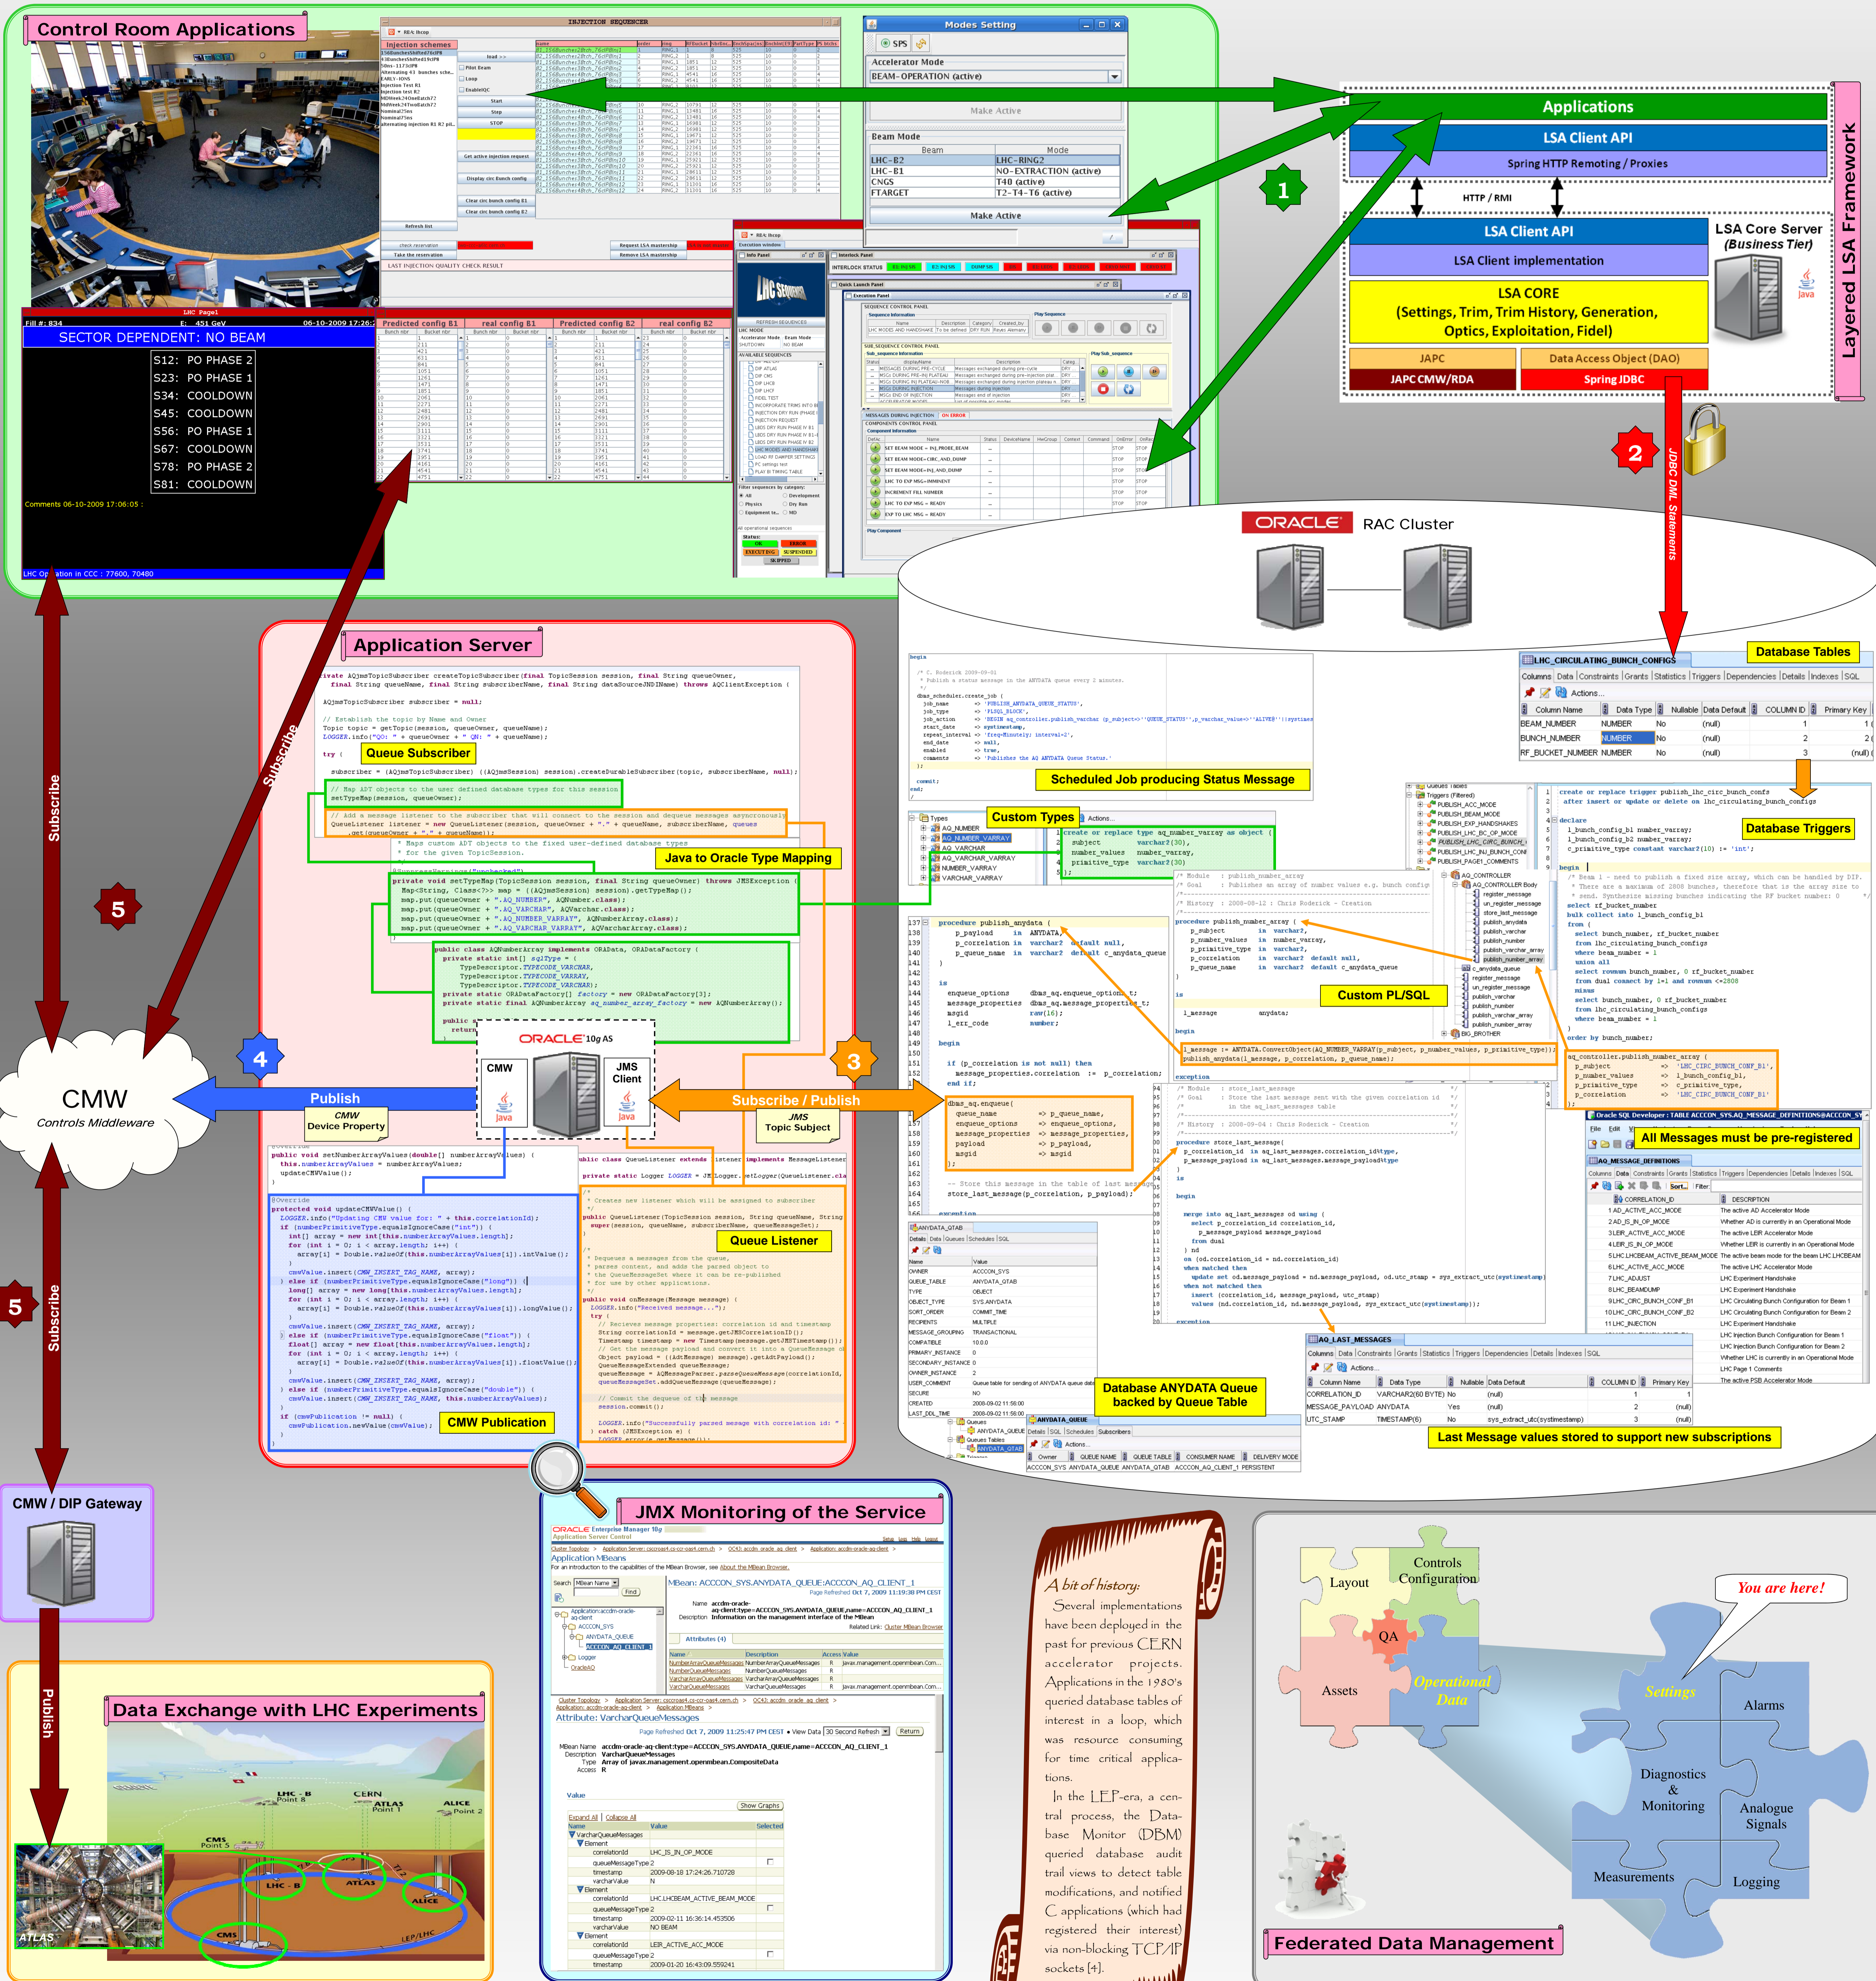


On-change Publishing of Database Resident Control System Data

K. Kostro, R. Billen, C. Roderick, CERN, Geneva, Switzerland

Abstract

The CERN accelerator control system is largely data driven, based on a distributed Oracle database architecture. Many application programs depend on the latest values of key pieces of information such as beam mode and accelerator mode. Rather than taking the non-scalable approach of polling the database for the latest values, the CERN control system addresses this requirement by making use of the Oracle Advanced Queuing – an implementation based on JMS (Java Message Service) – to publish data changes throughout the control system via the CERN Controls Middleware (CMW). This paper describes the architecture of the system, the implementation choices and the experience so far.



Conclusion

Twenty-six subjects are currently published via the on-change database publishing service to the CERN control system. The early implementation of the service was already used for the first LHC beams in September 2008. This service was created initially to distribute configuration information to LHC experiments as their control is completely detached from the LHC controls. This is still the main foreseen use, but we also expect an increasing utilization by LHC operational applications.

References

- [1] G. Kruk et al., "LHC Software Architecture [LSA] – Evolution toward LHC Beam Commissioning", ICALEPCS'07, WOPA03
- [2] C. Roderick and R. Billen, "The LSA Database to Drive the Accelerator Settings", ICALEPCS'09, WEP006.
- [3] R. Billen et al., "Accelerator Data Foundation: How It All Fits Together", ICALEPCS'09, TUB001.
- [4] E. Hatzigeorgidis and F. Roeder, "A Description of the SL Measurement and Logging System and the Application Interface Libraries", CERN SL/Note 93-38 (CO), 29 March 1993
- [5] K. Kostro et al., "The Controls Middleware (CMW) at CERN - Status and Usage", ICALEPCS'03, WE201

A bit of history:
Several implementations have been deployed in the past for previous CERN accelerator projects. Applications in the 1980's queried database tables of interest in a loop, which was resource consuming for time critical applications. In the LEP-era, a central process, the Database Monitor (DBM) queried database audit trail views to detect table modifications, and notified C applications (which had registered their interest) via non-blocking TCP/IP sockets [4].

