



Contribution ID: 205

Type: poster

Database Usage and Performance for the Fermilab Run II Experiments

Wednesday 29 September 2004 10:00 (1 minute)

The Run II experiments at Fermilab, CDF and D0, have extensive database needs covering many areas of their online and offline operations. Delivery of the data to users and processing farms based around the world has represented major challenges to both experiments. The range of applications employing databases includes data management, calibration (conditions), trigger information, run configuration, run quality, luminosity, and others. Oracle is the primary database product being used for these applications at Fermilab and some of its advanced features have been employed, such as table partitioning and replication. There is also experience with open source database products such as MySQL for secondary databases. A general overview of the operation, access patterns, and transaction rates is examined and the potential for growth in the next year presented. The two experiments, while having similar requirements for availability and performance, employ different architectures for database access. Details of the experience for these approaches will be compared and contrasted, as well as the evolution of the delivery systems throughout the run. Tools employed for monitoring the operation and diagnosing problems will also be described.

Authors: KUMAR, A. (Fermilab); BONHAM, D. (Fermilab); BOX, D. (Fermilab); LITVINTSEV, D. (FNAL); GAL-LAS, E. (Fermilab); WICKLUND, E. (Fermilab); KOWALKOWSKI, J. (Fermilab); TRUMBO, J. (Fermilab); LUEK-ING, L. (FERMILAB); VITTONI-WIERSMA, M. (Fermilab); STANFIELD, N. (Fermilab); MAKSIMOVIC, P. (John Hopkins University); JETTON, R. (Fermilab); KOVICH, S. (Fermilab); WHITE, S.P. (Fermilab); YASUDA, T. (Fermilab); GUO, Y. (Fermilab)

Presenter: LUEKING, L. (FERMILAB)

Session Classification: Poster Session 2

Track Classification: Track 4 - Distributed Computing Services