ICALEPCS'09

International Conference on Accelerator and Large Experimental Physics Control Systems
October 12 – 16, 2009, Kobe International Conference Center, Kobe , Japan

The LHC Logging Service: Handling Terabytes of On-line Data

C. Roderick, R. Billen, R. D. Gaspar Aparicio, E. Grancher, A. Khodabandeh, N. Segura Chinchilla, CERN, Geneva, Switzerland

Abstract

Based on previous experience with LEP, a long-term data logging service for the LHC was developed and put in place in 2003, several years before beam operation. The scope of the logging service covers the evolution over time of data acquisitions on accelerator equipment and beam related parameters. The intention is to keep all this time-series data on-line for the lifetime of LHC, allowing easy data comparisons with previous years. The LHC hardware commissioning has used this service extensively prior to the first beams in 2008 and even more so in 2009. Current data writing rates exceed 15TB/year and continue to increase. The high data volumes and throughput rates, in writing as well as in reading, require special arrangements on data organization and data access methods .

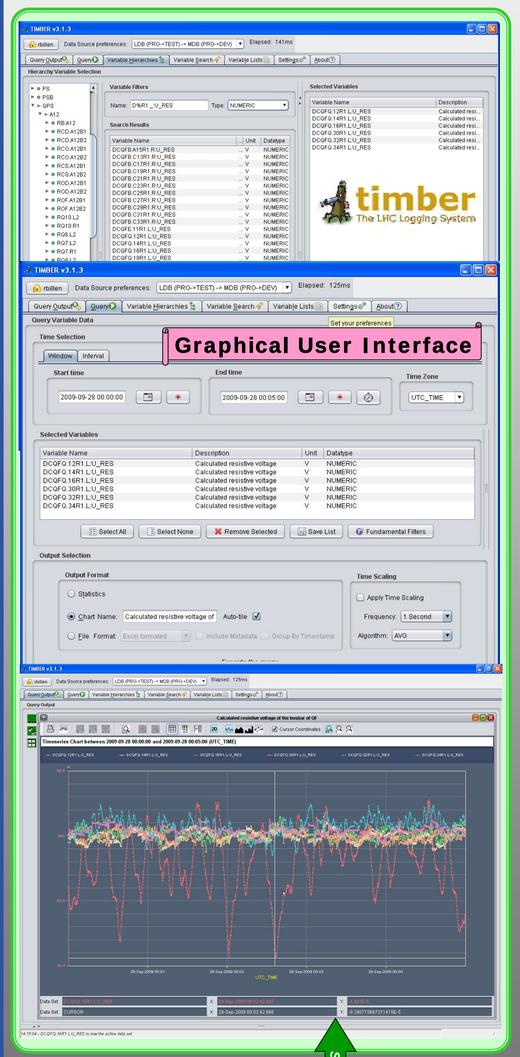

Spring Remoting

Application Servers

ORACLE 10g AS

READ API (Java)

WRITE API (Java)

JDBC

JDBC

Control System Infrastructure

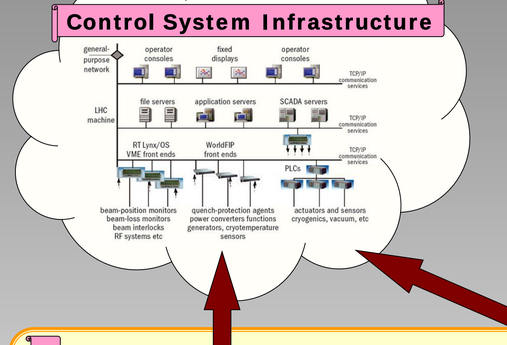

Monitoring of the Service

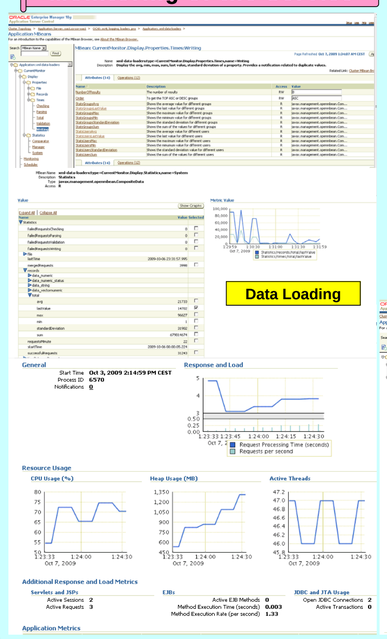

Measurements (MDB) to Logging (LDB)

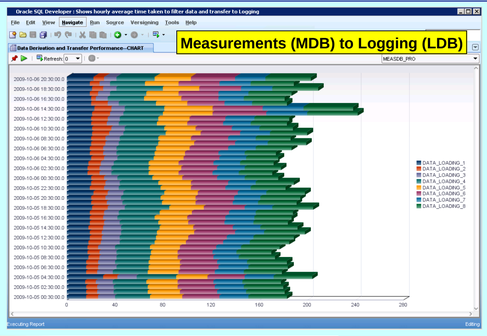

Infrastructure

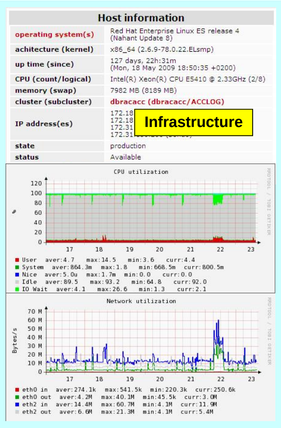

Database Activity

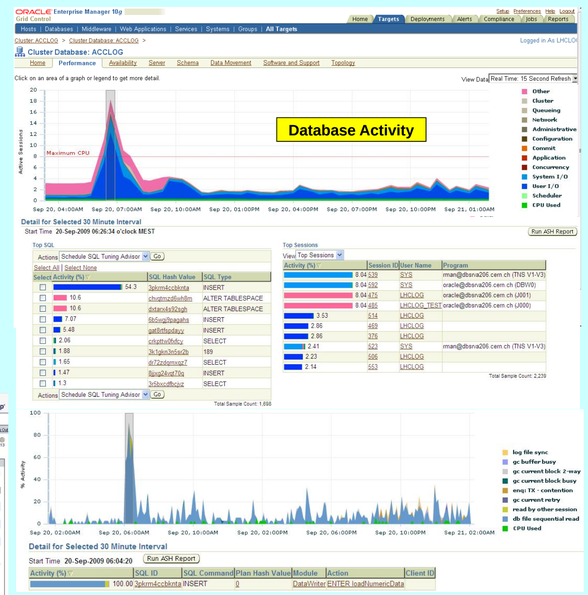

Data Loading

Data Extraction

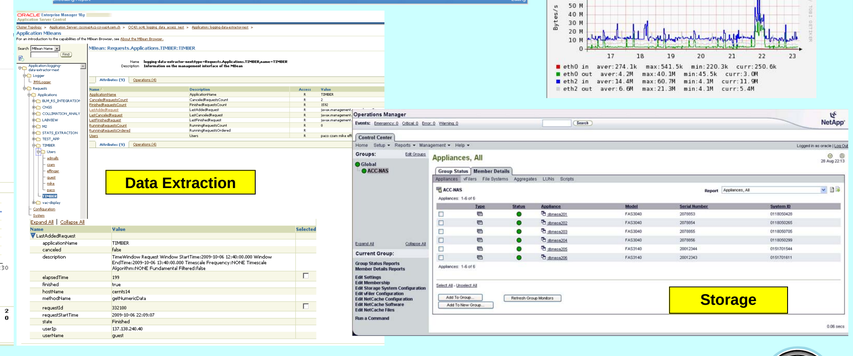

Storage

Graphical User Interface

Scheduled Jobs

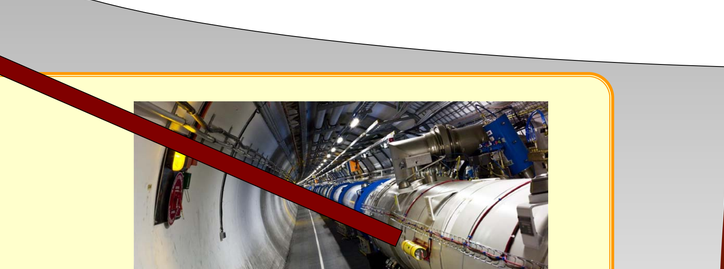

Databases in Numbers

Object Counts		
Object Type	MDB	LDB
Tables	*100*	*52*
Indexes	*138*	*63*
Types	*26*	*45*
Sequences	*24*	*11*
Views	*101*	*70*
Functions	*2*	*0*
Procedures	*3*	*10*
Triggers	*21*	*16*
Packages	*15*	*12*

Constraint Counts		
Constraint Type	MDB	LDB
Unique	*105*	*48*
Referential	*73*	*31*
Value Check	*125*	*28*
Total:	303	107

PL/SQL Lines		
Source Type	MDB	LDB
Functions	*112*	*0*
Procedures	*342*	*497*
Triggers	*342*	*267*
Packages	*1'868*	*2'652*
Package Bodies	*12'846*	*10'884*
Total:	15'510	14'300

Data Volume

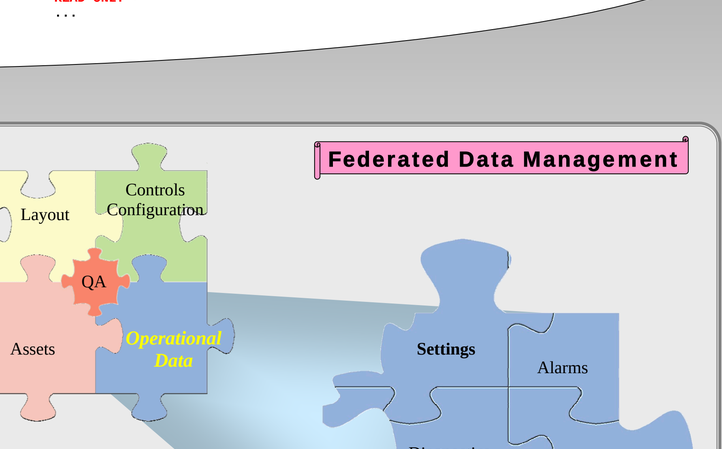

Backup

Backup Type	Frequency	Duration
Incremental	*1 per day*	*±1 hour*
Full, without READ ONLY data	*1 per week*	*±8 hours*
Full, all data	*1 per 3 weeks*	*±6 days*

Database Servers

Short-Term Measurement Database
120GB / day

PL/SQL Filter

Long-Term Logging Database
75GB / day

```
create table data_numeric (
  variable_id number,
  utc_stamp   timestamp (9),
  value       number,
  constraint dn_pk          primary key (variable_id, utc_stamp),
  constraint dn_variable_fk foreign key (variable_id)
    references meta_variables (variable_id)
)
organization index compress 1
partition by range (utc_stamp) (
  partition part_dn_2003h2   values less than (timestamp '2004-01-01 00:00:00') tablespace old_log_data,
  ...
  partition part_dn_2004q3   values less than (timestamp '2004-10-01 00:00:00')  tablespace old_log_data,
  ...
  partition part_dn_200507   values less than (timestamp '2005-08-01 00:00:00') tablespace old_log_data,
  ...
  partition part_dn_20070301 values less than (timestamp '2007-03-02 00:00:00') tablespace old_log_data,
  ...
  partition part_dn_20091005 values less than (timestamp '2009-10-06 00:00:00') tablespace log_data_20091005,
  partition part_dn_20091006 values less than (timestamp '2009-10-07 00:00:00') tablespace log_data_20091006,
  partition part_dn_20091007 values less than (timestamp '2009-10-08 00:00:00') tablespace log_data_20091007
);
```

Compressed Index Organized Tables

Daily Range Partitions & Tablespaces

Simplified Data Model

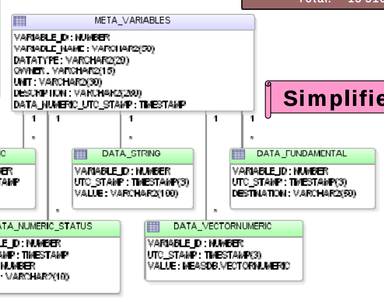


```
select t.tablespace_name, t.status
from dba_tablespaces t
where t.tablespace_name like 'LOG_DATA
order by t.tablespace_name desc;

TABLESPACE_NAME                STATUS
------------------------------ ---------
LOG_DATA_20091007              ONLINE
LOG_DATA_20091006              ONLINE
LOG_DATA_20091005              ONLINE
...                            ...
LOG_DATA_20090921              READ ONLY
LOG_DATA_20090920              READ ONLY
...                            ...
```

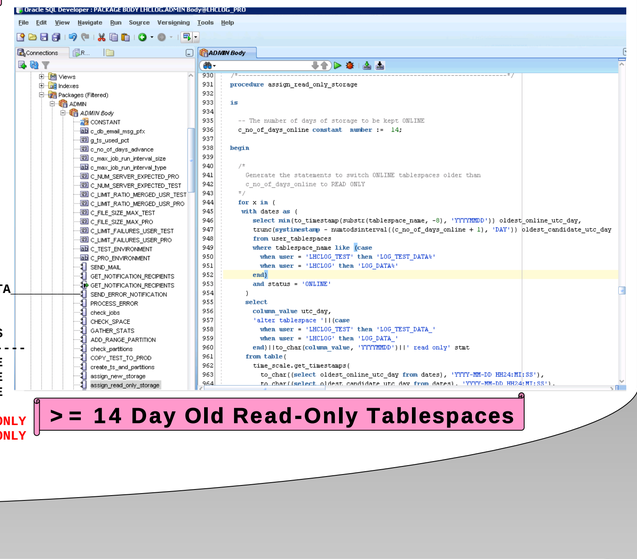

> = 14 Day Old Read-Only Tablespaces

A bit of history:

The LEP Logging System was put in place in 1992, three years after the commissioning of LEP.

Systematic logging of LEP hardware and physics parameters, and being able to correlate this data, enabled a performance increase of LEP operation.

For LHC, the project for a logging system was started in 2001, several years before beam operation.

Federated Data Management

Conclusion

The LHC Logging Service has proved its usefulness since several years. The demanding requirements for a centralized, long-term multi-terabyte system have imposed thoroughness in technical choices and compromises. Continuous performance monitoring is a necessity to ensure quality of service. Nevertheless, the challenge to keep up with the increasing expectations remains. A wide client community is looking forward to capturing, analyzing and exploiting LHC beam data.

References

[1] R. Billen et al., "LEP Accelerator Logging System Using On-line Database", ICALEPCS'93, Berlin, Germany, October 1993, NIMA 352 (1994) 296-299
[2] R. Billen et al., "Accelerator Data Foundation: How It All Fits Together", ICALEPCS'09, TUB001
[3] C. Roderick and R. Billen, "Capturing, Storing and Using Time-Series Data for the World's Largest Scientific Instrument", November 2006, CERN-AB-Note-2006-046 (CO)