


# Databases for Run2

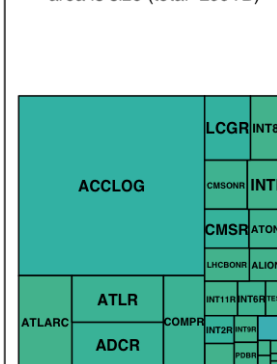
- Based on three main solutions:
- Use of MySQL / PostgreSQL (DBoD managed)
  - High growth, count of databases ~doubles every year.
  - With critical databases: OpenStack, Zenodo, Drupal, Puppet (stored configuration and Foreman Data, and Puppet Cert Manager), VOMS, Dirac (LHCb Dirac and Dirac File Catalog), Jupyter, Indico, Gitlab (VCS), Alfresco (eFiles), etc.
- Use of OracleDB, HW/OS/SW stable combination
  - Stable number, growing in size and activity:
  - Significant size increase (500GB/day for ACCLOG since beginning of Run2).
  - Very active, total changes size (redo) per month around 500TB.
- Use of Hadoop-based components/databases
  - For larger scale data analysis and OracleDB offloading.
  - Production for WDT dashboards, ATLAS DDM, IT monitoring, OpenStack Ceilometer, etc.
  - Working on improving integration with OracleDB.
  - Investigating next generation accelerator control and logging (ACCLOG being loaded: 250TB in OracleDB is compressed to 50TB on Hadoop/Parquet format).

## Database on Demand instances

Evolution of the amount of MySQL, Oracle, and PostgreSQL instances in the DBoD service

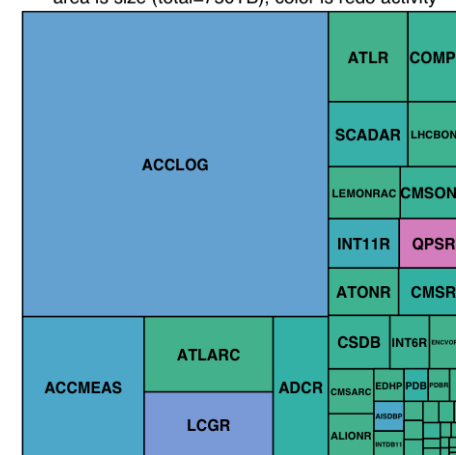


CERN databases 201210, area is size (total=299TB)



TB redo per month, sum=120 TB

CERN databases 201512, area is size (total=750TB), color is redo activity



TB redo per month, sum=494 TB

	October 2012	December 2015
Max size	ACCLOG 136TB	ACCLOG 352TB
Max redo	ACCMEAS 27TB / month	QPSR 115TB / month

