



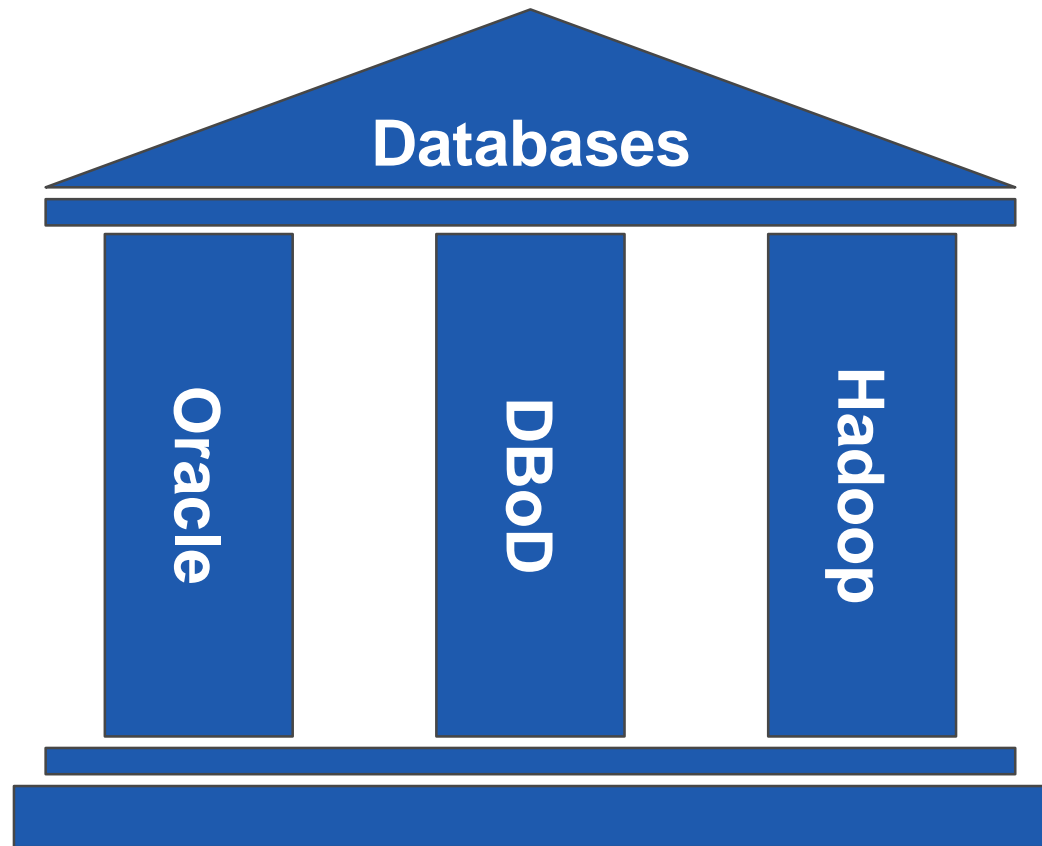
Database Services

Katarzyna Dzedziniewicz-Wojcik
On behalf of IT-DB

Agenda

- Database service overview
- Oracle
 - Status
 - Plans
- Database On Demand
 - Status
 - Plans

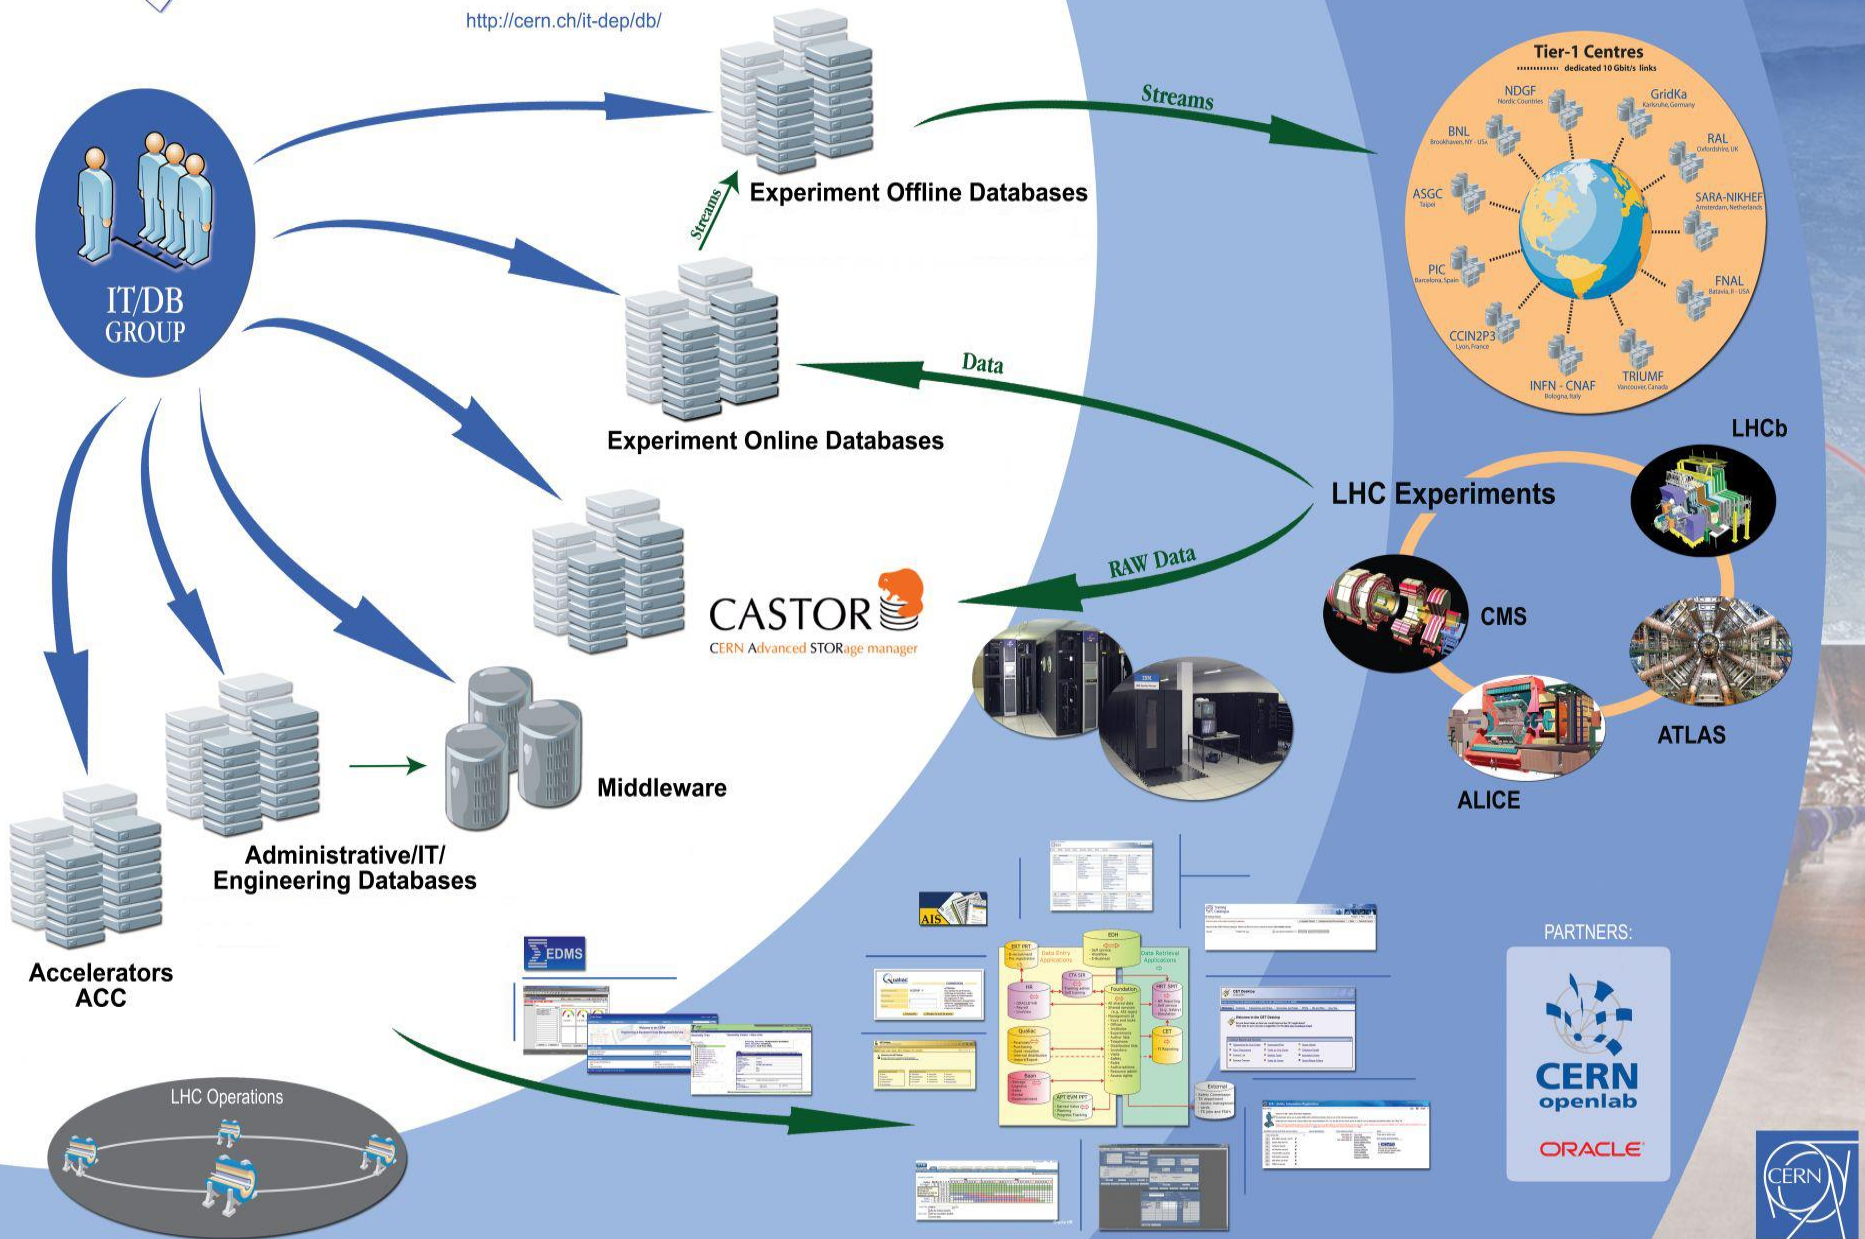
The history of architecture



Databases

- 3 main pillars
 - Oracle
 - For critical, transactional load
 - Administered by experts (DBAs)
 - Database on-demand
 - Different database engines
 - MySQL, PostgreSQL, Oracle
 - Instance owners have the DBA rights
 - Hadoop
 - For analysis where DBs are not suitable
 - Check presentation by Kacper Surdy

<http://cern.ch/it-dep/db/>



Oracle in numbers

- ~100 databases, most of them RAC
 - Mostly NAS storage plus some SAN with ASM
 - More than 600 TB of data files for production DBs in total

Oracle main users

- Accelerator logging
 - ~ 4 billion records/day
 - ~ 160 GB/day □ 52 TB/year stored
- Experiment databases
 - Online
 - Data taking operations
 - Offline
 - Data analysis

Oracle main users

- Administrative
 - All applications are Oracle based
- Engineering & component
 - Managing a million components over a lifecycle of 30 years
- Quench Protection System
 - Stores ~150 K rows/second
- Consolidated WinCC/PVSS archive repository
 - Stores ~50-60K rows/second

Oracle - plans

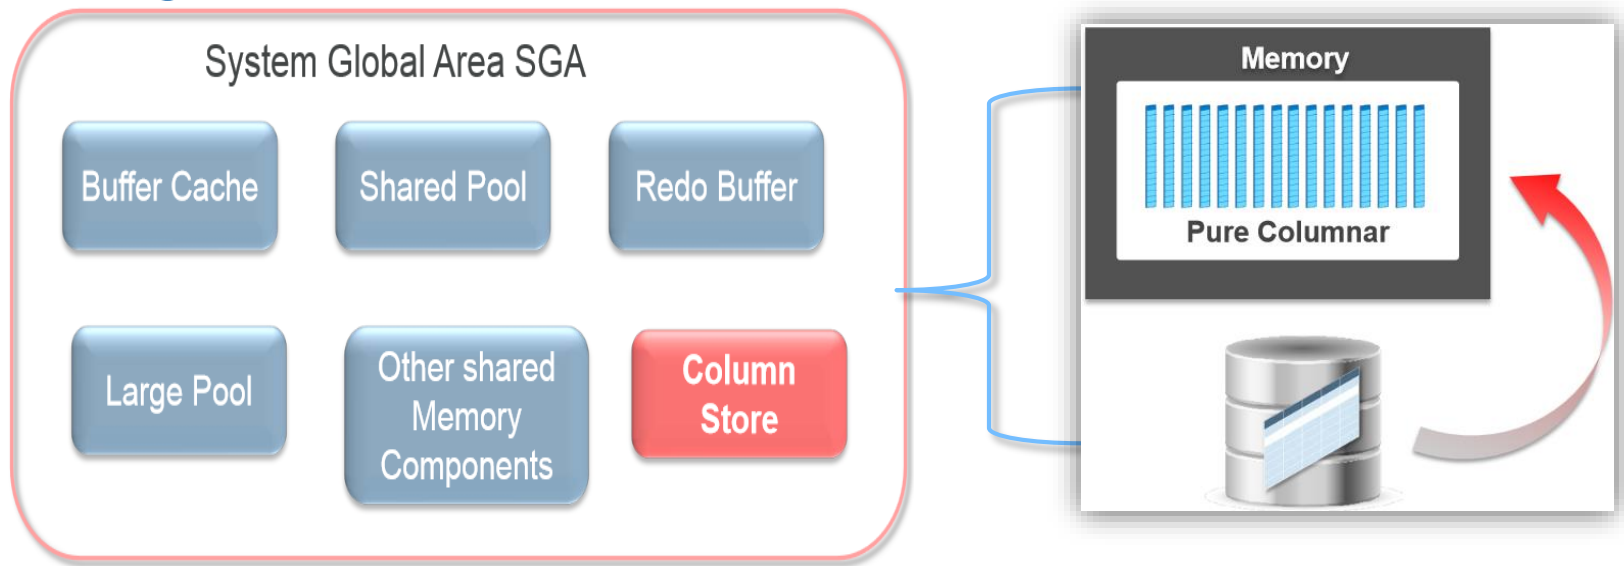
- Oracle 11.2 and 12.1
 - Supported until 2020 and 2021 respectively
- LS2 upgrade
 - **Oracle 12.2** (recommended)
 - Release date 2016

Oracle 12 - why

- Main new feature
 - In-Memory database
 - Real-time responsiveness
 - We cannot forget...
 - ... bug fixes

In Memory DB

- New hype in DB technology
 - More solutions appearing on the market
- A new in-memory cache – **columnar format**
- Huge performance boost for full **table scans!**

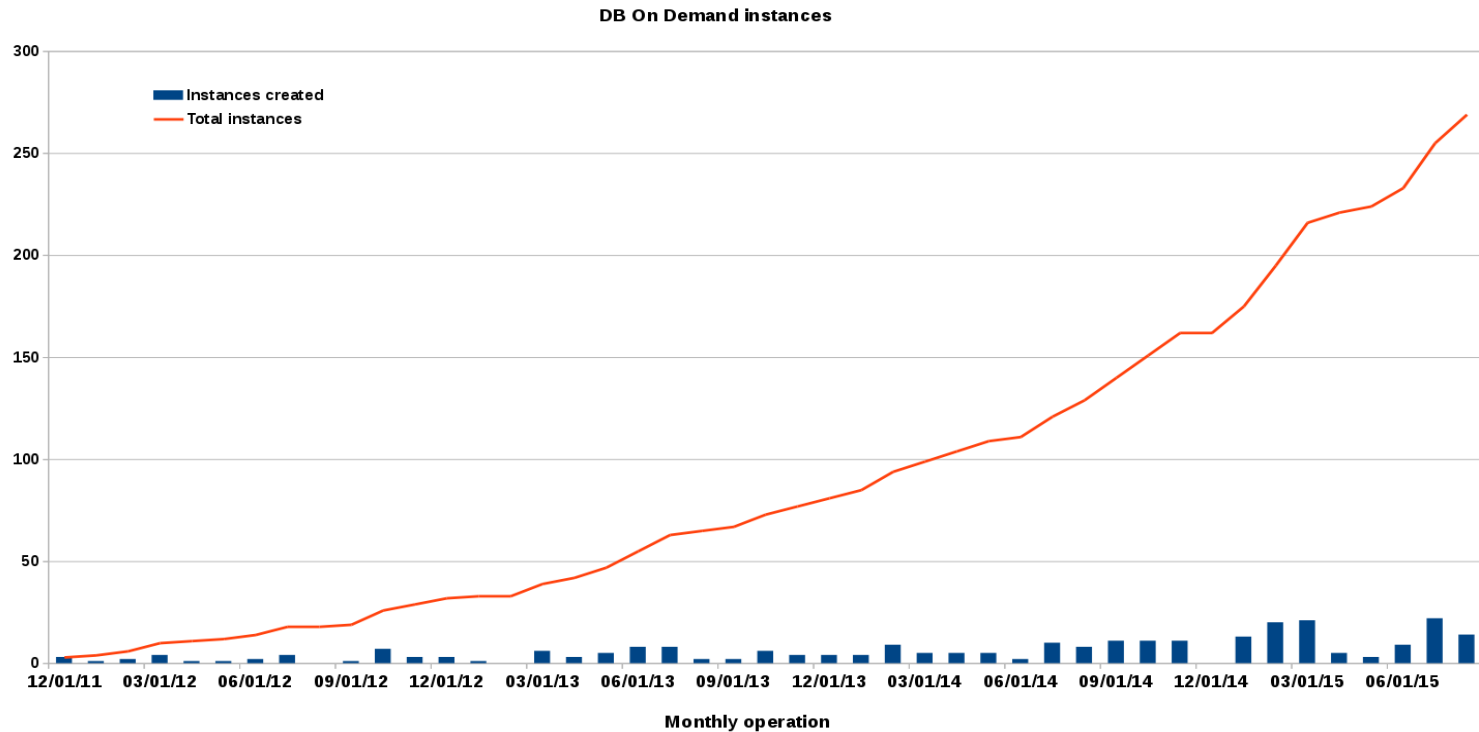


In Memory DB - why

- Deployment is **transparent** to applications
- Automatically used when beneficial
- Compression helps to put more data in memory
 - More data than in traditional Buffer Cache
- No negative impact on DML operations
- Please, contact IT-DB for more details

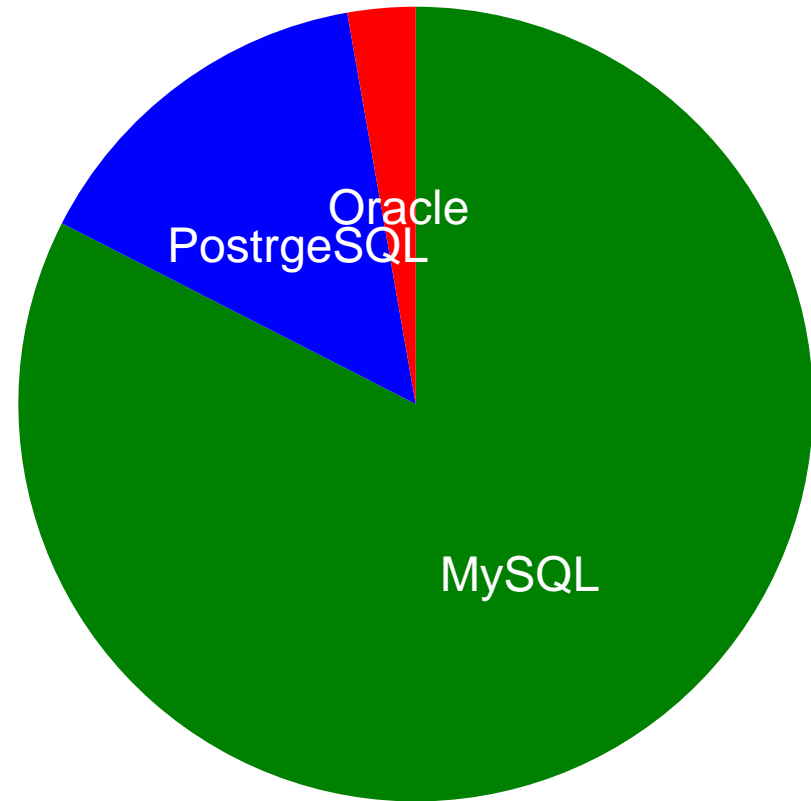
On-Demand in numbers

- All databases – 326
- Production – 225




On-Demand in numbers

- By DB Type
 - MySQL: 269
 - PostgreSQL: 48
 - Oracle: 9



On-Demand Users

Database on Demand instances per Department



Databases – On-Demand

- MySQL CE 5.7
 - Q3/Q4 2016
 - To stay for next 3 years
- PostgreSQL 9.4
 - Q3/Q4 2016
- Oracle as for General Service
- Upgrades require database reboot

Plans for On-Demand

- High availability
 - Replication
 - App/DB Gateways
 - Clustered Databases

DBoD High Availability

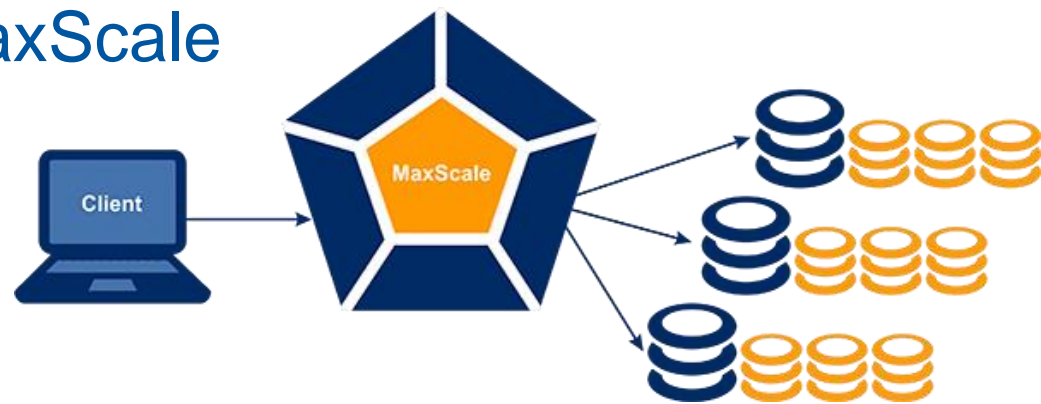
- Oracle Clusterware
 - Widely used by DBoD clients
 - Ensures that resources are running in any of the cluster nodes
 - Issues
 - Only 1 active instance on the same data volume
 - Unused resources
 - MySQL NFS locking
 - Requires manual intervention in storage
 - Clusterware overreaction
 - May lead to data corruption

DBoD High Availability

- Master Slave replication
 - Already being done for both PostgreSQL and MySQL (Drupal,..)
 - Possible to have a replica in Wigner
 - Read Only replica which can be manually promoted to be master (RW) or used as failover during master outages
 - Requires application/admins to be aware

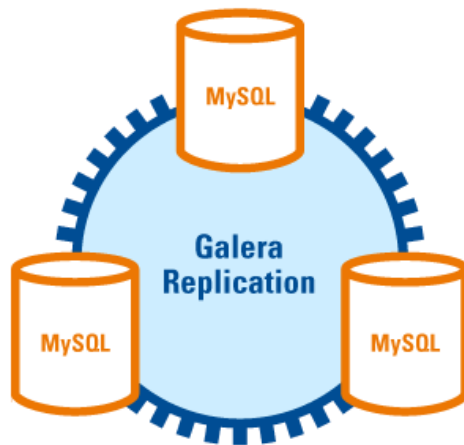
DBoD High Availability

- Application to DB Gateway
 - Builds on standard replication for HA
 - Transparent to applications
 - DB Load distribution
- Solutions
 - MariaDB MaxScale
 - pgpool-ii



DBoD High Availability

- Clustered databases
 - Multimaster solution
 - Support for rolling interventions
 - New model, new operational challenges



Galera Cluster for MySQL



CitusDB for PostgreSQL

Acknowledgments

- Colleagues from IT-DB
- Special thanks
 - Ignacio Coterillo
 - Emil Pilecki