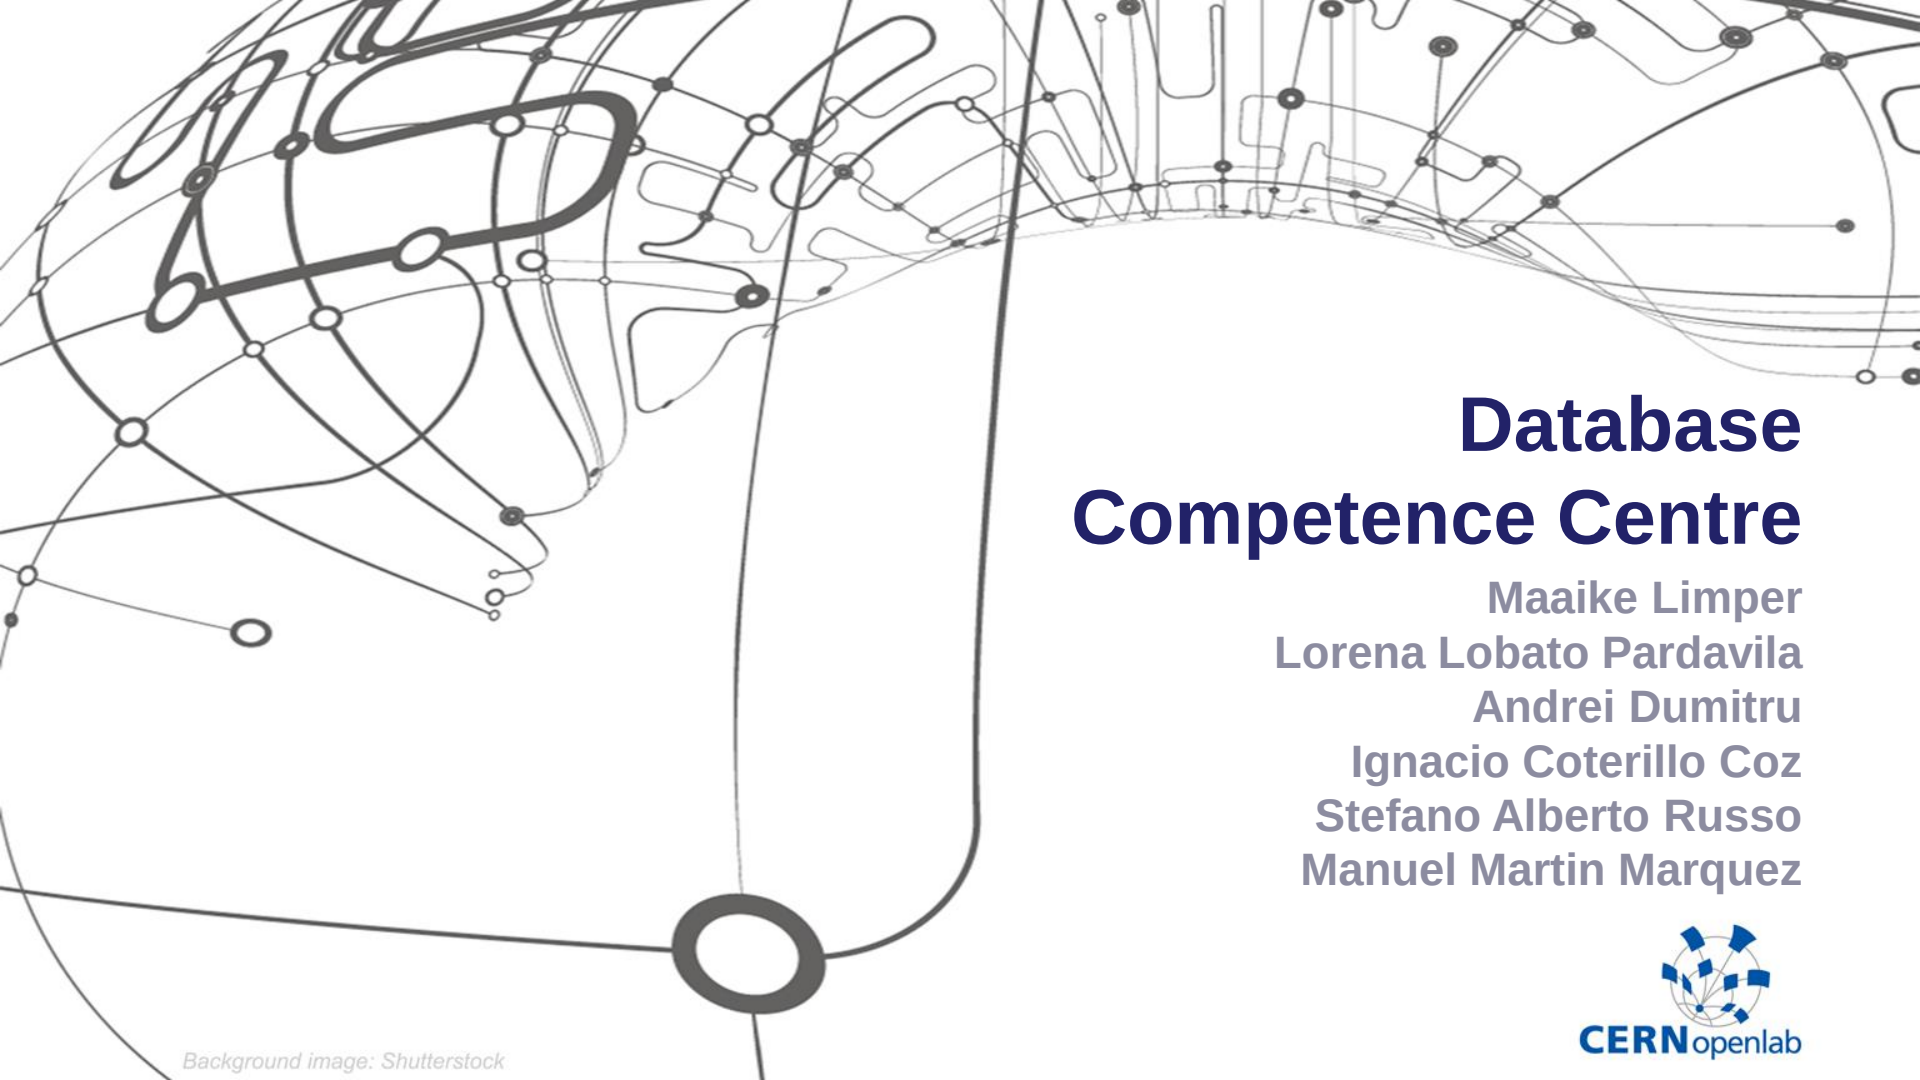

Database Competence Centre

Maaike Limper
Lorena Lobato Pardavila
Andrei Dumitru
Ignacio Coterillo Coz
Stefano Alberto Russo
Manuel Martin Marquez

CERN openlab

Database Competence Centre

- Virtualization
- Replication Technology
- In-Database Physics Analysis
- Enterprise Manager and 12c new Features
- Data Analytics

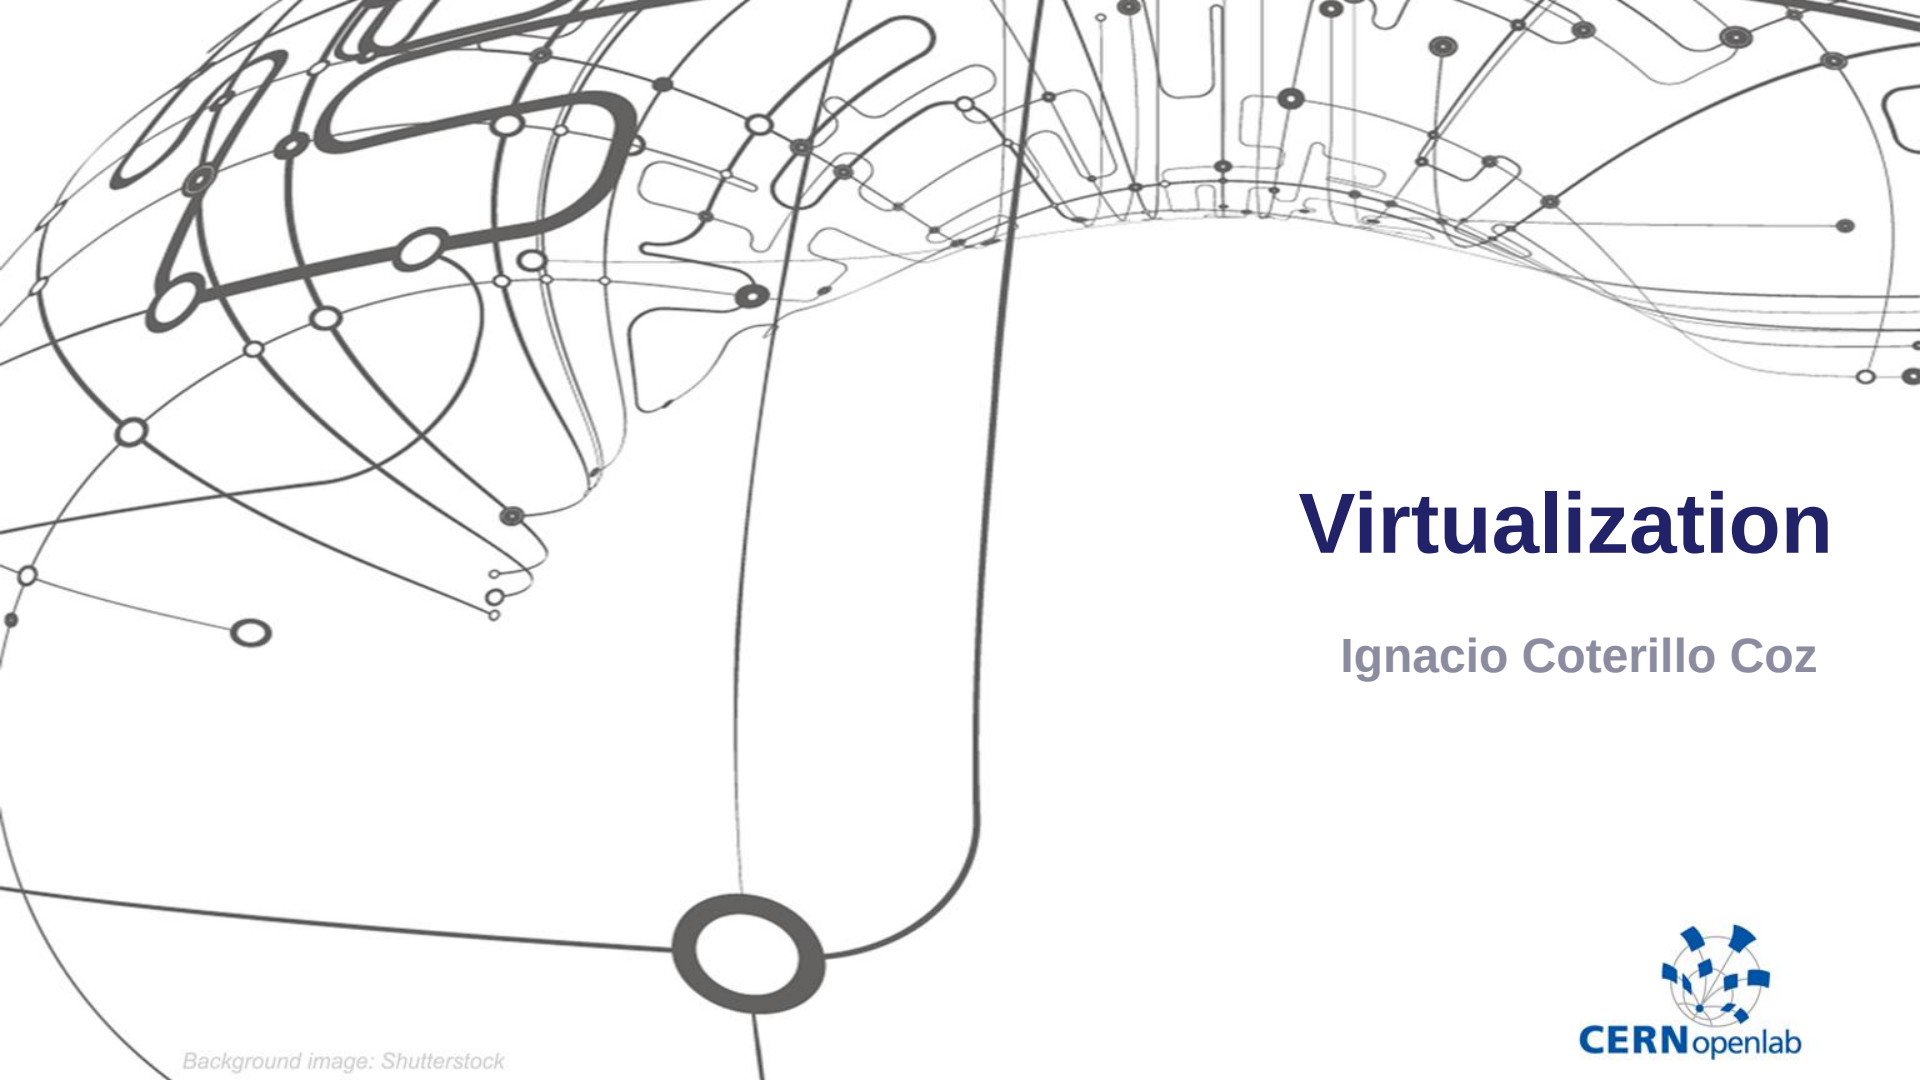

Virtualization

Ignacio Coterillo Coz

CERN openlab

Virtualization

Reminder: What we are trying to do

- Integrate Oracle VM Hypevisors as compute nodes providers for OpenStack
- Working Custom OVM Server like Hypervisor
  - Due to dependencies and limitations of use in OVM 3.2.x
  - Targeting Oracle Linux 6 (base for future OVM Server in next releases 3.3.x, 4.x)
  - Xen 4.1.6-rc1 compiled from the Xen Git repository

Virtualization

Achieved milestones

- Setting up an OpenStack Havana test environment
- Replicated same integration of custom hypervisor as a Nova compute node as previously achieved for Grizzly
- Networking dependencies/Linux Kernel version mismatches are now solved
- Tested the feasibility of using Xen 4.3
  - incompatibility issues with the required version of **libvirt**

Virtualization

- **In the meantime**,
  - Oracle is about to release a **beta version of the next Oracle VM** release, which solves the limitations we were targeting with our custom hypervisor:
    - Based on OL 6.x
    - Possibility of installing extra packages (No more black box model)
  - The IT-DB infrastructure is on the process of integrating with the new CERN computing infrastructure **(OpenStack + Puppet)**
    - IT-DB currently has a number of databases and services running on Oracle VM
    - This services have to be migrated to the new operating infrastructure while continuing to run on Oracle VM for support or certification reasons

Virtualization

- **With this in mind, the next steps are:**
  - Substitute **Custom Hypervisor with Oracle VM** Hypervisor Beta in the existing test environments
  - Start Integration tests of Oracle VM Server in the CERN IT OpenStack production setup
    - Custom CERN Networking/Storage access
    - Configuration management (Puppet)
  - Testing OpenStack HA solutions

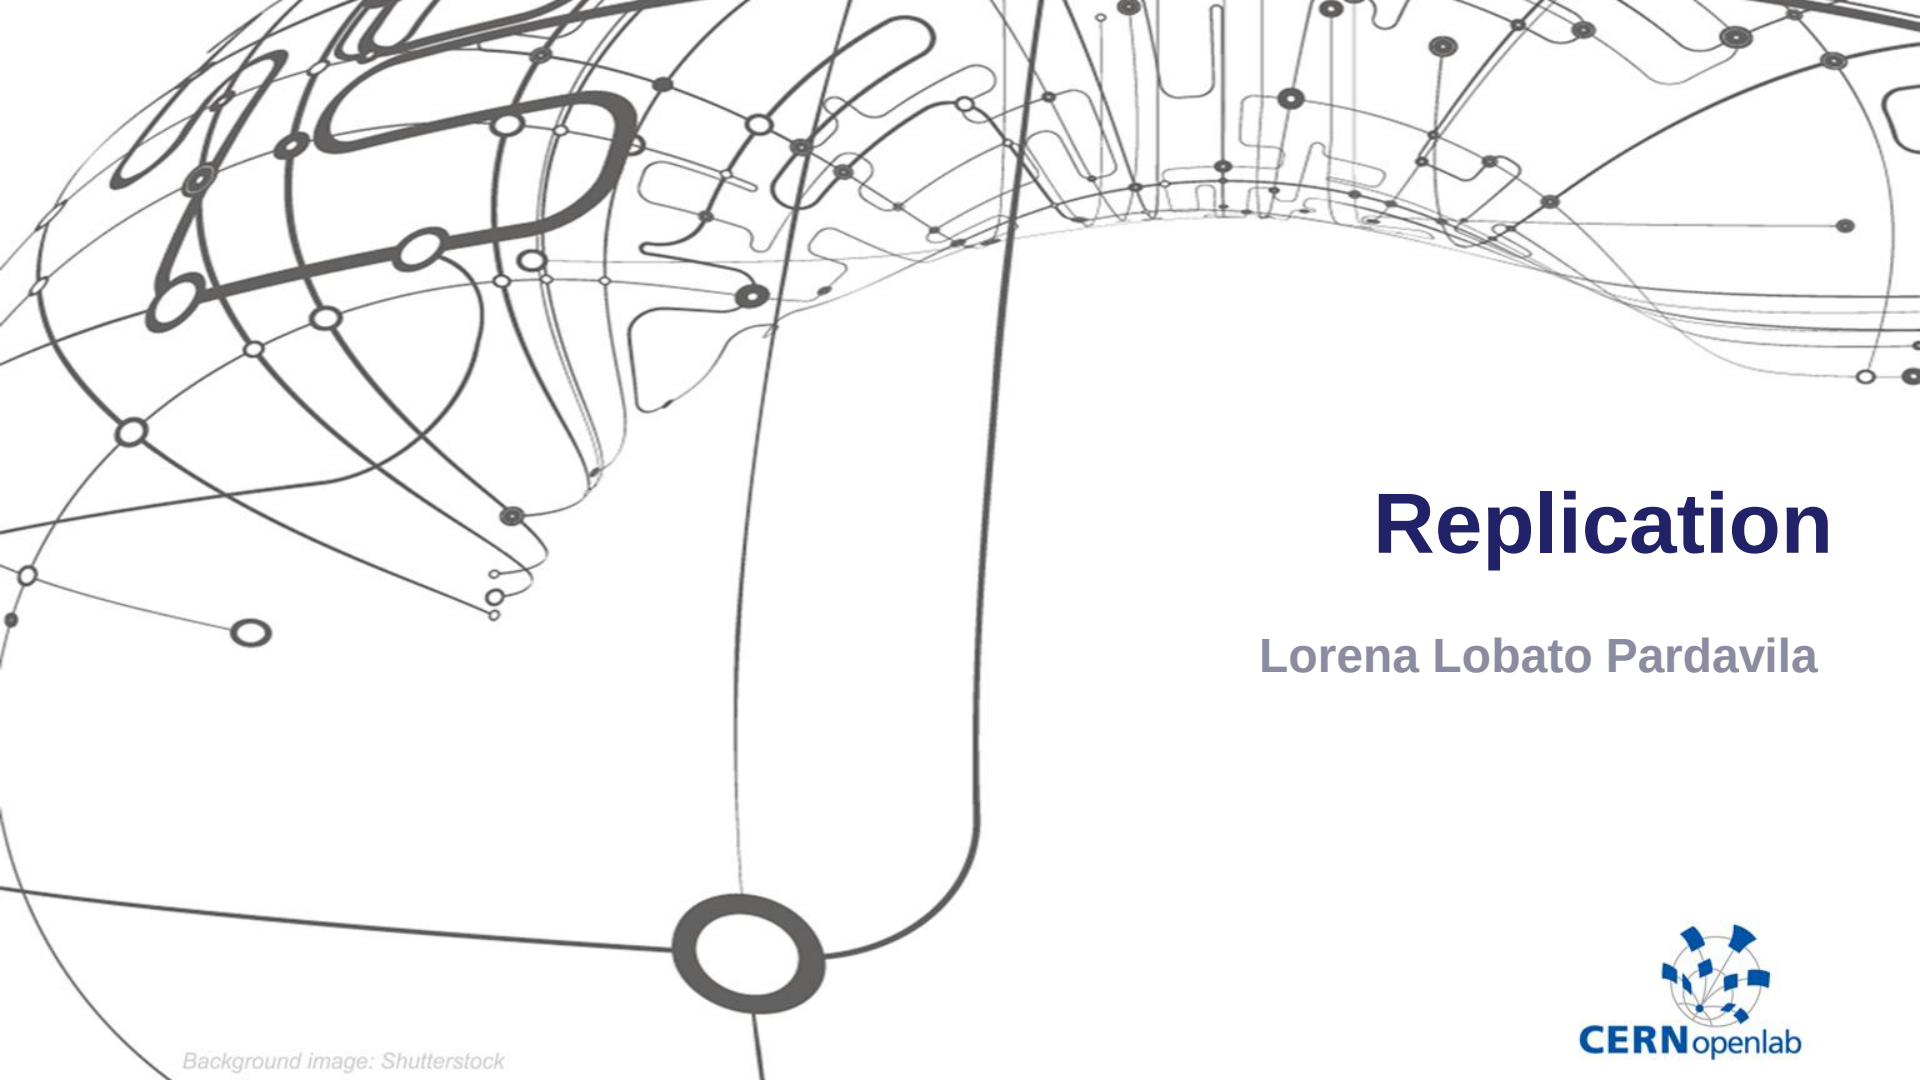

Replication

Lorena Lobato Pardavila

CERN openlab

Recent Activities

- **Testing and evaluation of the newest version of Golden Gate (12.1.2)**
  - Performance test with generic and production workloads
    - Very good results observed
  - Better performance than Streams (11g or 12c)
  - Streams and GG12c comparision presented at UKOUG13
  - Extensive feedback provided to Oracle
- **Evaluation of Oracle GoldenGate Director**
- **Preparation of Streams to GoldenGate migration**
  - Testing migration procedures

Golden Gate 12c

- **A lot of (good) features inherited from Streams**
- **Improved scalability - performance better than Streams**
- **Easier deployment, data instantiation and administration**
- **Availability of in-database monitoring and reporting**
- **A lot of new features had been evaluated**
  - Integrated Replicate
  - Downstream capture
  - Conflicts Detection-Resolution(CDR)

Golden Gate 12c

- **Generic workload (on the same hardware)**

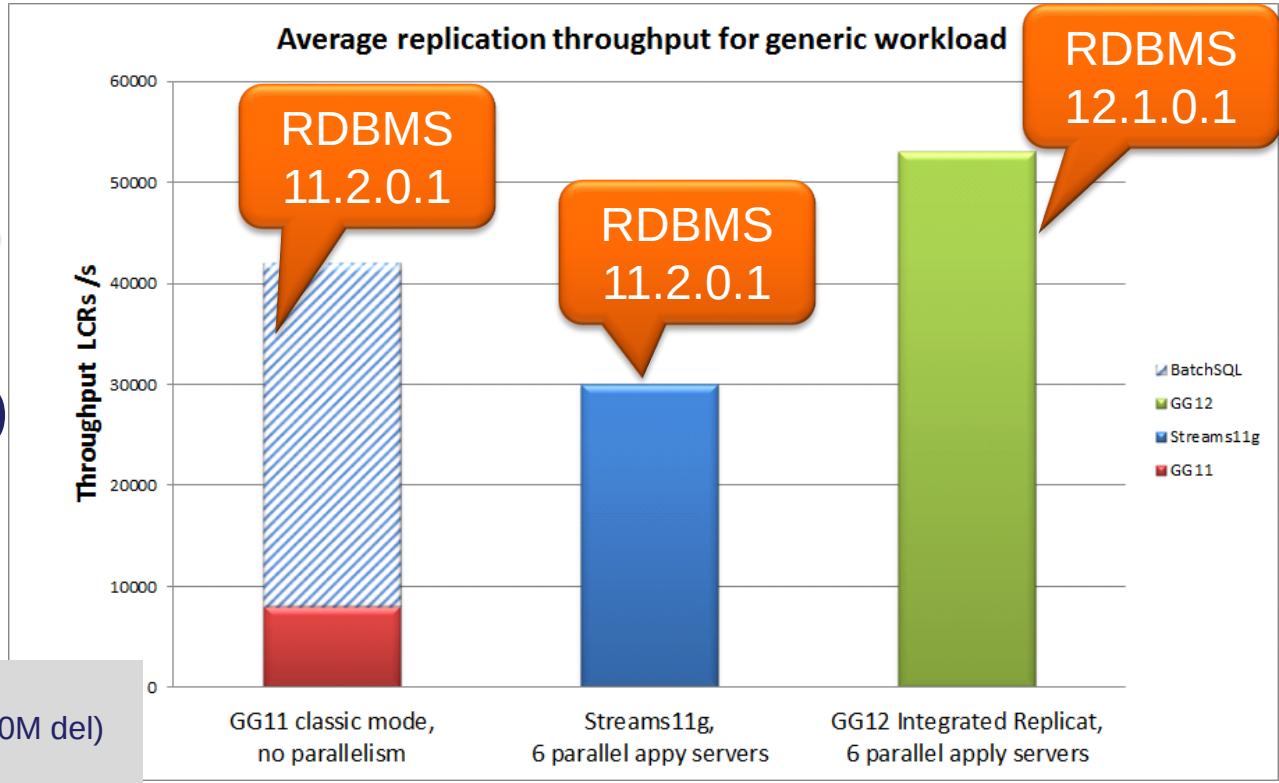

Workload description:

- 30M of DMLs (10M ins, 10M upd ,10M del)
- 10GB of redo volume
- 100k of transactions

Overview Golden Gate Director

OGG Director for centralized monitoring and configuration (wraps GGSCI)

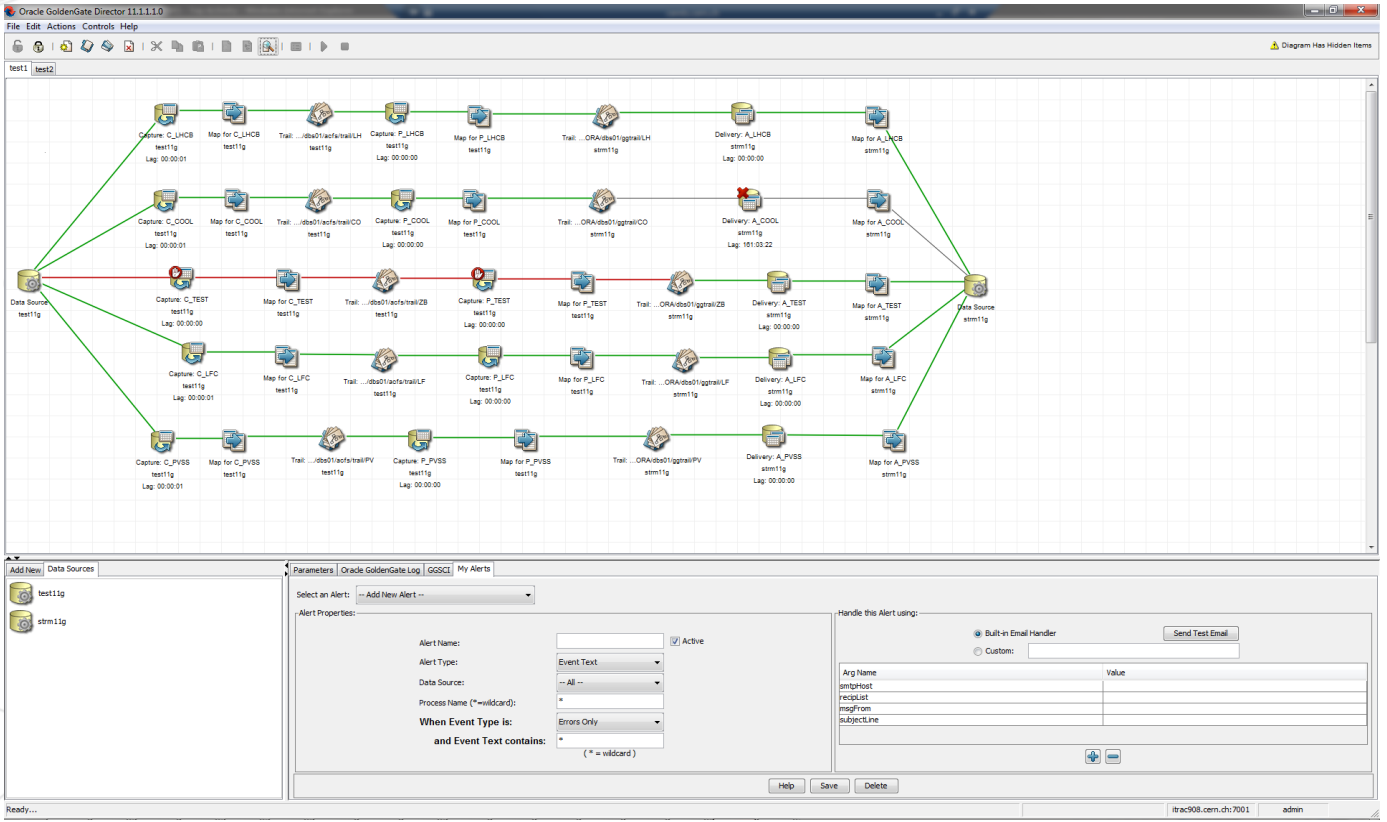

Future Plans

- **Pilot production like replication between T0 and T1s (BNL, TRIUMF)**
- **GG administration workshop for T1s DBA**
- **Migrations from Streams to Golden Gate 12c**
  - LHCB (Q2-Q3)
  - ATLAS online – offline (Q2)
  - ATLAS offline -T1s (Q3 2014)
- **Golden Gate replication as a service**
- **Evaluation of zero-downtime database upgrades/migrations using Golden Gate**

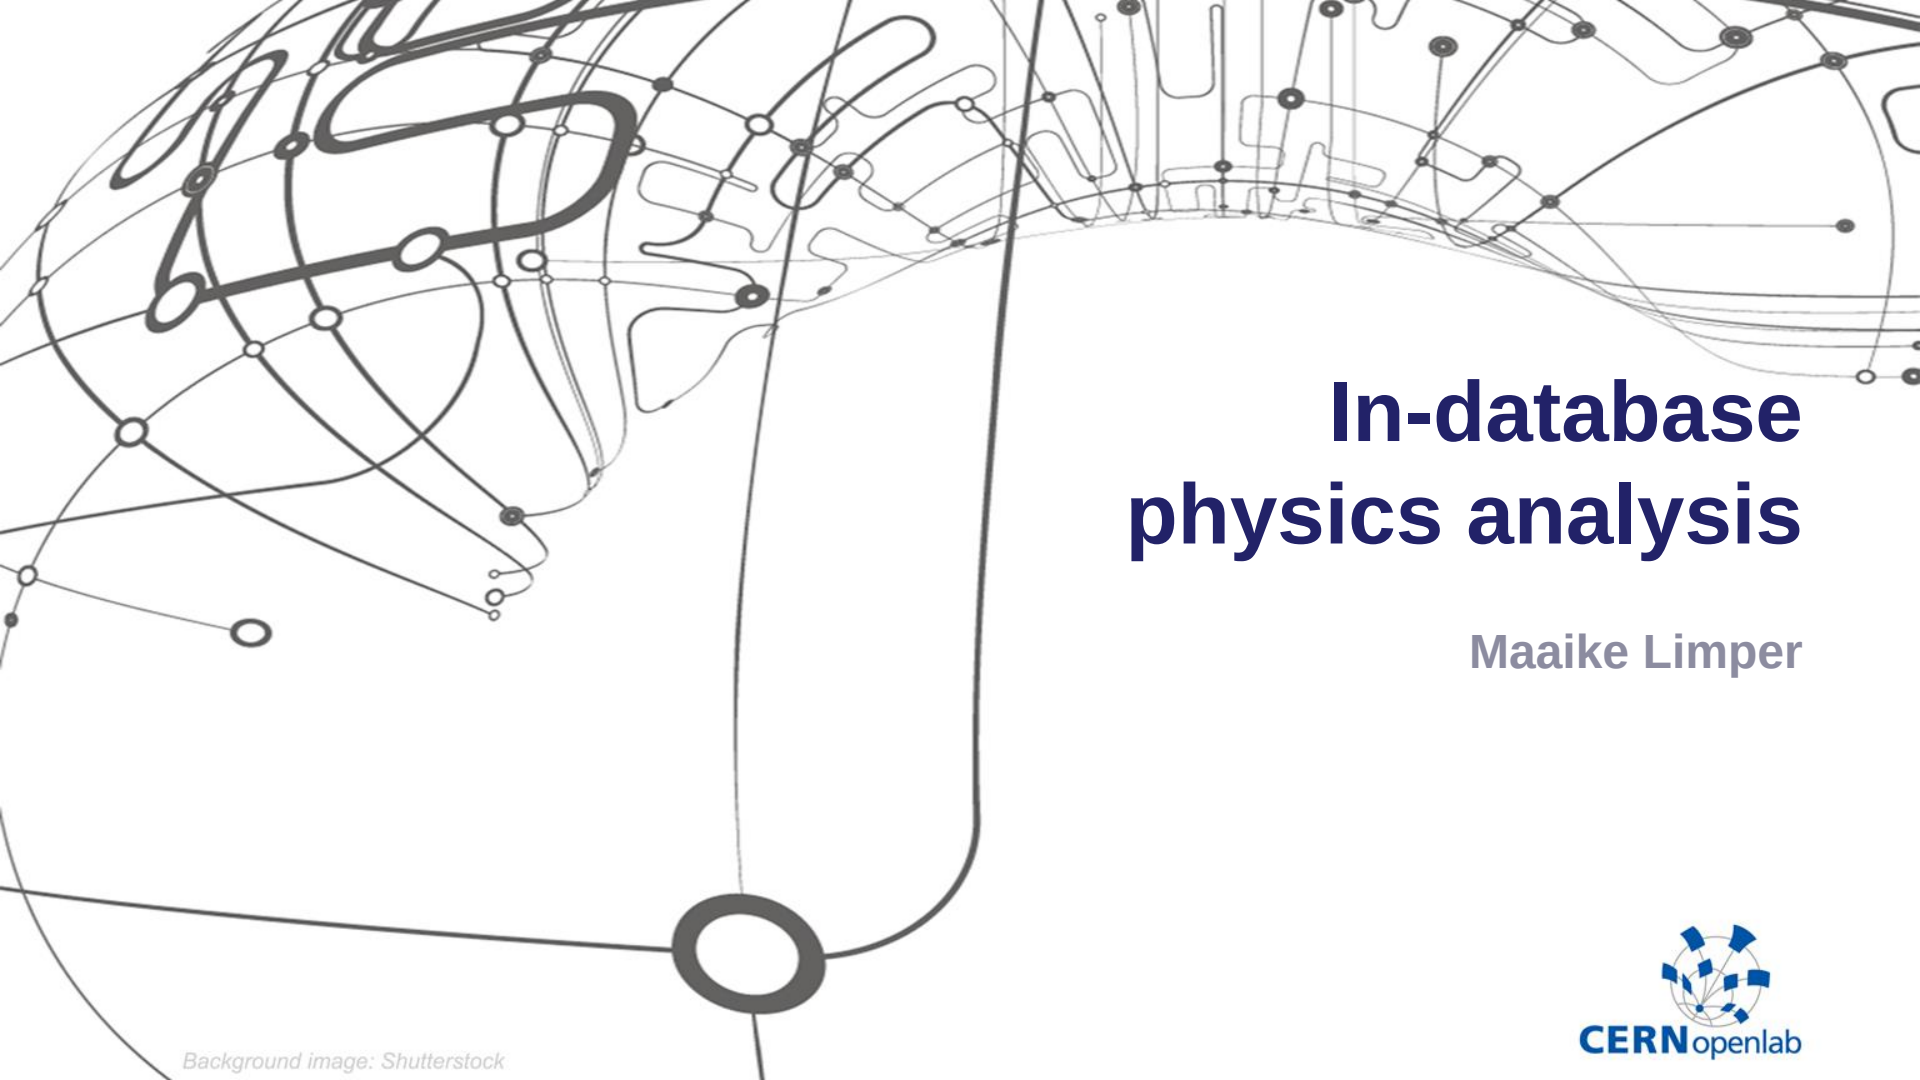

In-database physics analysis

Maaike Limper

CERN openlab

A new analysis benchmark

- **Previous studies done using simplified Higgs+Z benchmark**
  - 40 variables from 6 different tables, simplified w.r.t actual analysis, external code only for jet selection
- **New more realistics benchmark “ttbar cutflow” (top-pair quark production)**
  - Actual ATLAS analysis code from svn converted into Oracle SQL
  - Requires external code in photon, muon and jet-selection
  - Many multi-table joins required
  - Uses 319 variables from 10 different tables
- ***New benchmark is more realistic test of physics analysis in Oracle database*!**

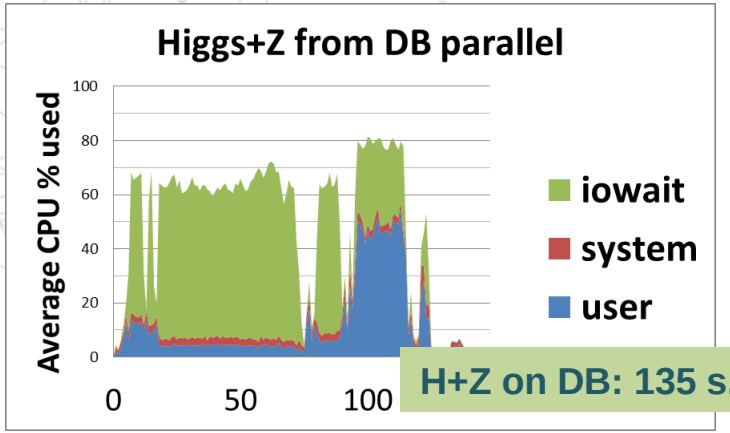

H+Z on DB: 135 s.

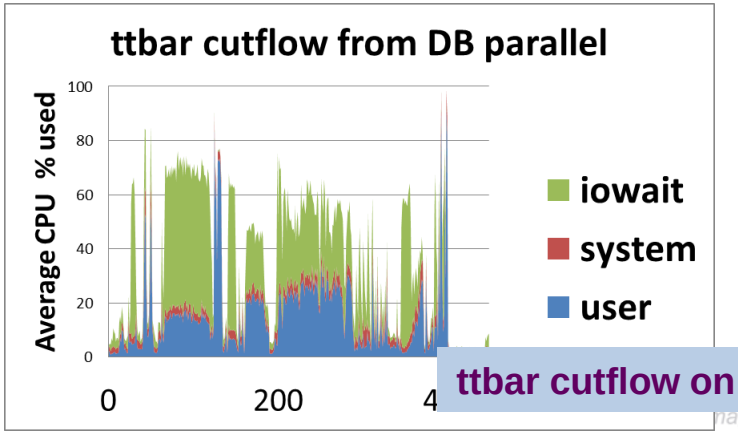

ttbar cutflow on DB: 390 s.

Oracle vs ROOT (multi-job)

- Performance comparison on 5-machine cluster, data divided over 12 disks per machine
- Run ROOT using multiple root-jobs, each analysing sub-set of data, and merge result
- Compare ROOT and Oracle using same degree of parallelism=P40

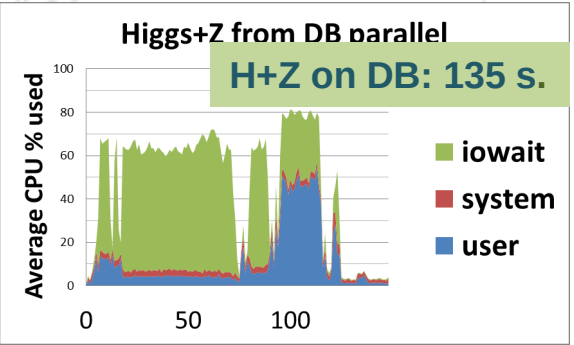

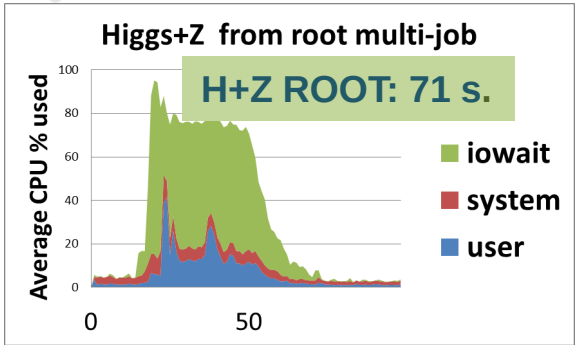

Simple HZ benchmark is fastest using parallel root

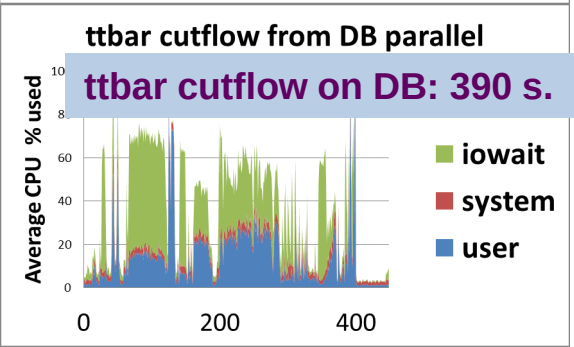

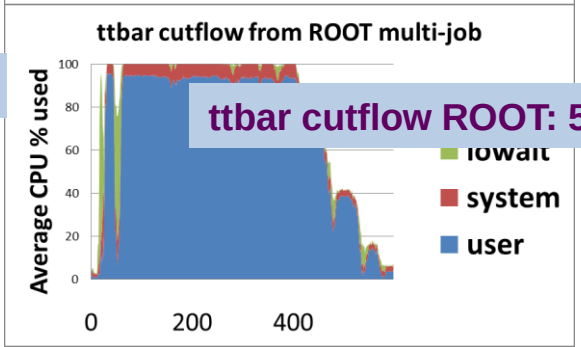

Official root analysis code slower than the Oracle SQL version…

Oracle vs ROOT (multi-job)

- **But I could improve original analysis code by improving the way the branches were loaded**
- **After this ROOT runs (slightly) faster then DB for ttbar cutflow as well**
  - Long tail in multi-job analysis as scripts wait for all jobs to complete before merging

ttbar cutflow on DB: 390 s.

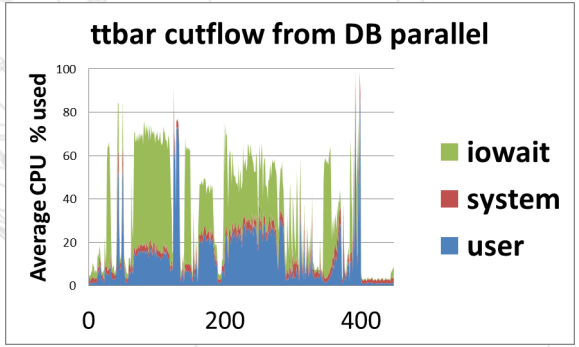

ttbar cutflow ROOT (old): 588 s.

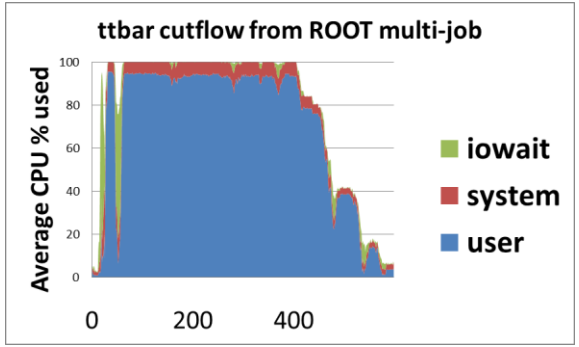

ttbar cutflow ROOT (new): 371 s.

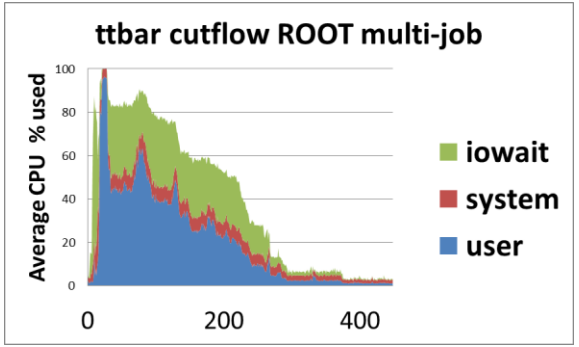

Multi-job ROOT vs PROOF

- Official framework to run ROOT in parallel is the *Parallel ROOT Framework*
- Use Proof-on-demand (PoD) on 5-machine cluster with 40 workers
- Benchmark analysis with PROOF on cluster slower than test with multiple ROOT jobs, mainly due to initial time needed to compile classes on workers

H+Z ROOT multi-job: 71 s.

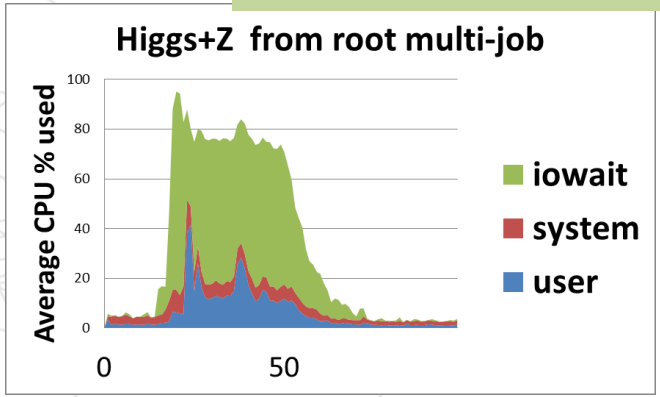

H+Z PROOF: 172 s.

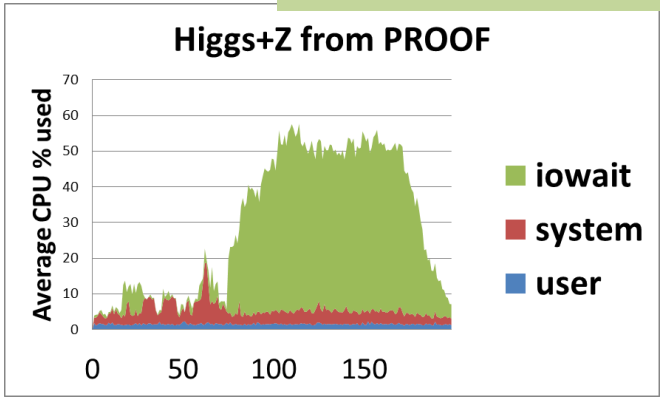

Hadoop tests

- **Reinstalled our Hadoop cluster: now running Hadoop 2!**
  - New in Hadoop2: TaskTracker replaced by Yarn's ResourceManager
- **First test with Hadoop parquet**
  - Parquet is a column-based storage for Hadoop
  - Storing the files in Parquet format, with gzip compression, results in greatly reduced datasize
  - Still need to write MapReduce implementation on Parquet

Data for Object Name	ROOT size in GB (est.)	Hadoop Parquet (w. gzip) size in GB	Oracle table size in GB
jet	15.13	6.5	32.27
electron	44.40	27.8	94.67
MET	2.19	1.7	2.53
EF (trigger)	0.91	0.08	7.02

Conclusion/Outlook

Near future task-list:

- **Oracle In-Memory Columnar beta**
- **Get a decent Hadoop reference working, using parquet storage looks interesting but still need to implement (efficient) version of benchmark code**
- **Keep going with the Hadoop vs ROOT vs Oracle and prepare final judgement on pro's an cons of the different approaches**
  - What is easiest way for user to get good performance? I need to consider my use of SQL hints, re-write of analysis code, personal setup for parallelism etc.!

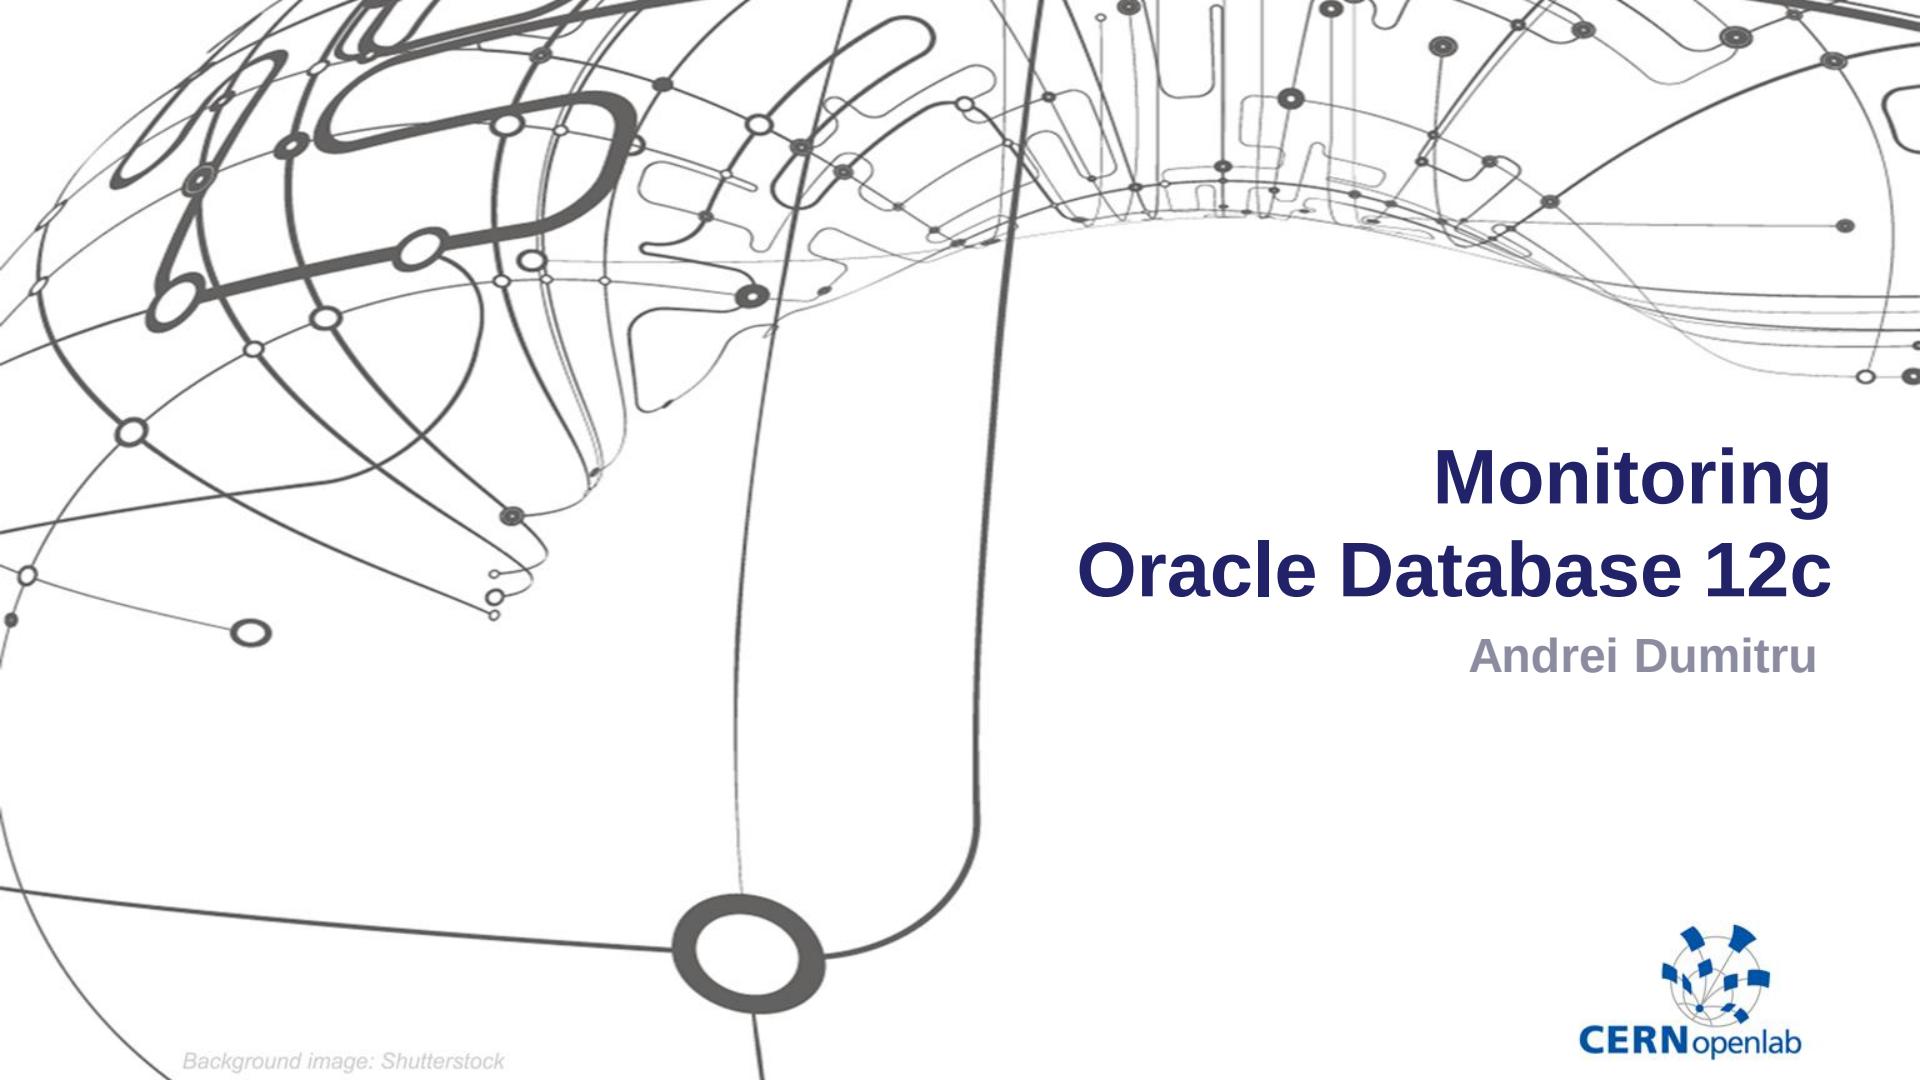

Monitoring Oracle Database 12c

Andrei Dumitru

CERN openlab

Background image: Shutterstock

Oracle DB 12c

- ***Oracle Database 12c available at CERN***
  - Development service available
- **Continuing work on**
  - Applications validation and certification
  - Capture and replay production workloads
  - Identify consolidation candidates
  - Data protection and availability (Oracle Data Guard)
  - Prepare deployment in our environment
  - Expand DBoD service with the new Oracle Multitenant Architecture

12.1 Upgrade Testing

- **Testing procedure**
  - Using Oracle Real Application Testing
  - Capture 8-16h of production workload (11.2.0.3)
  - Restore production database on target hardware
  - Upgrade the restored replica to 12.1.0.1
  - Replay workload in the upgraded database
  - Analyse results and publish report
    - Follow up issues if any

slide provided by Emil Pilecki, IT-DB

12.1 Upgrade Testing

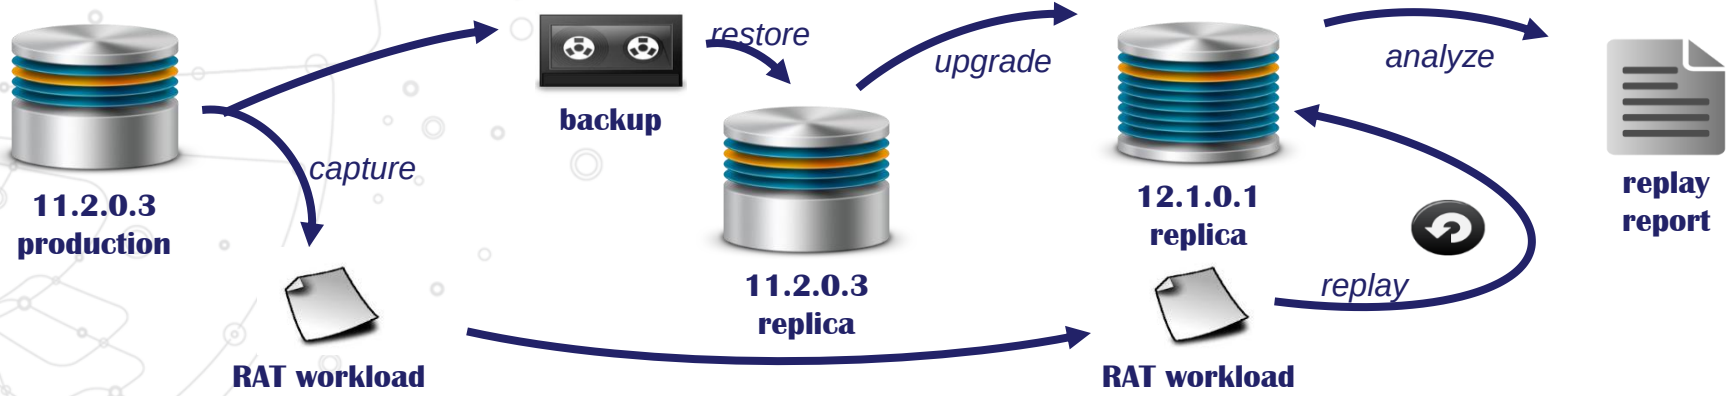

slide provided by Emil Pilecki, IT-DB

12.1 Upgrade Testing

- **Current results**
  - 11 databases tested
  - Feedback delivered to Oracle
- **Migration to production**

slide provided by Emil Pilecki, IT-DB

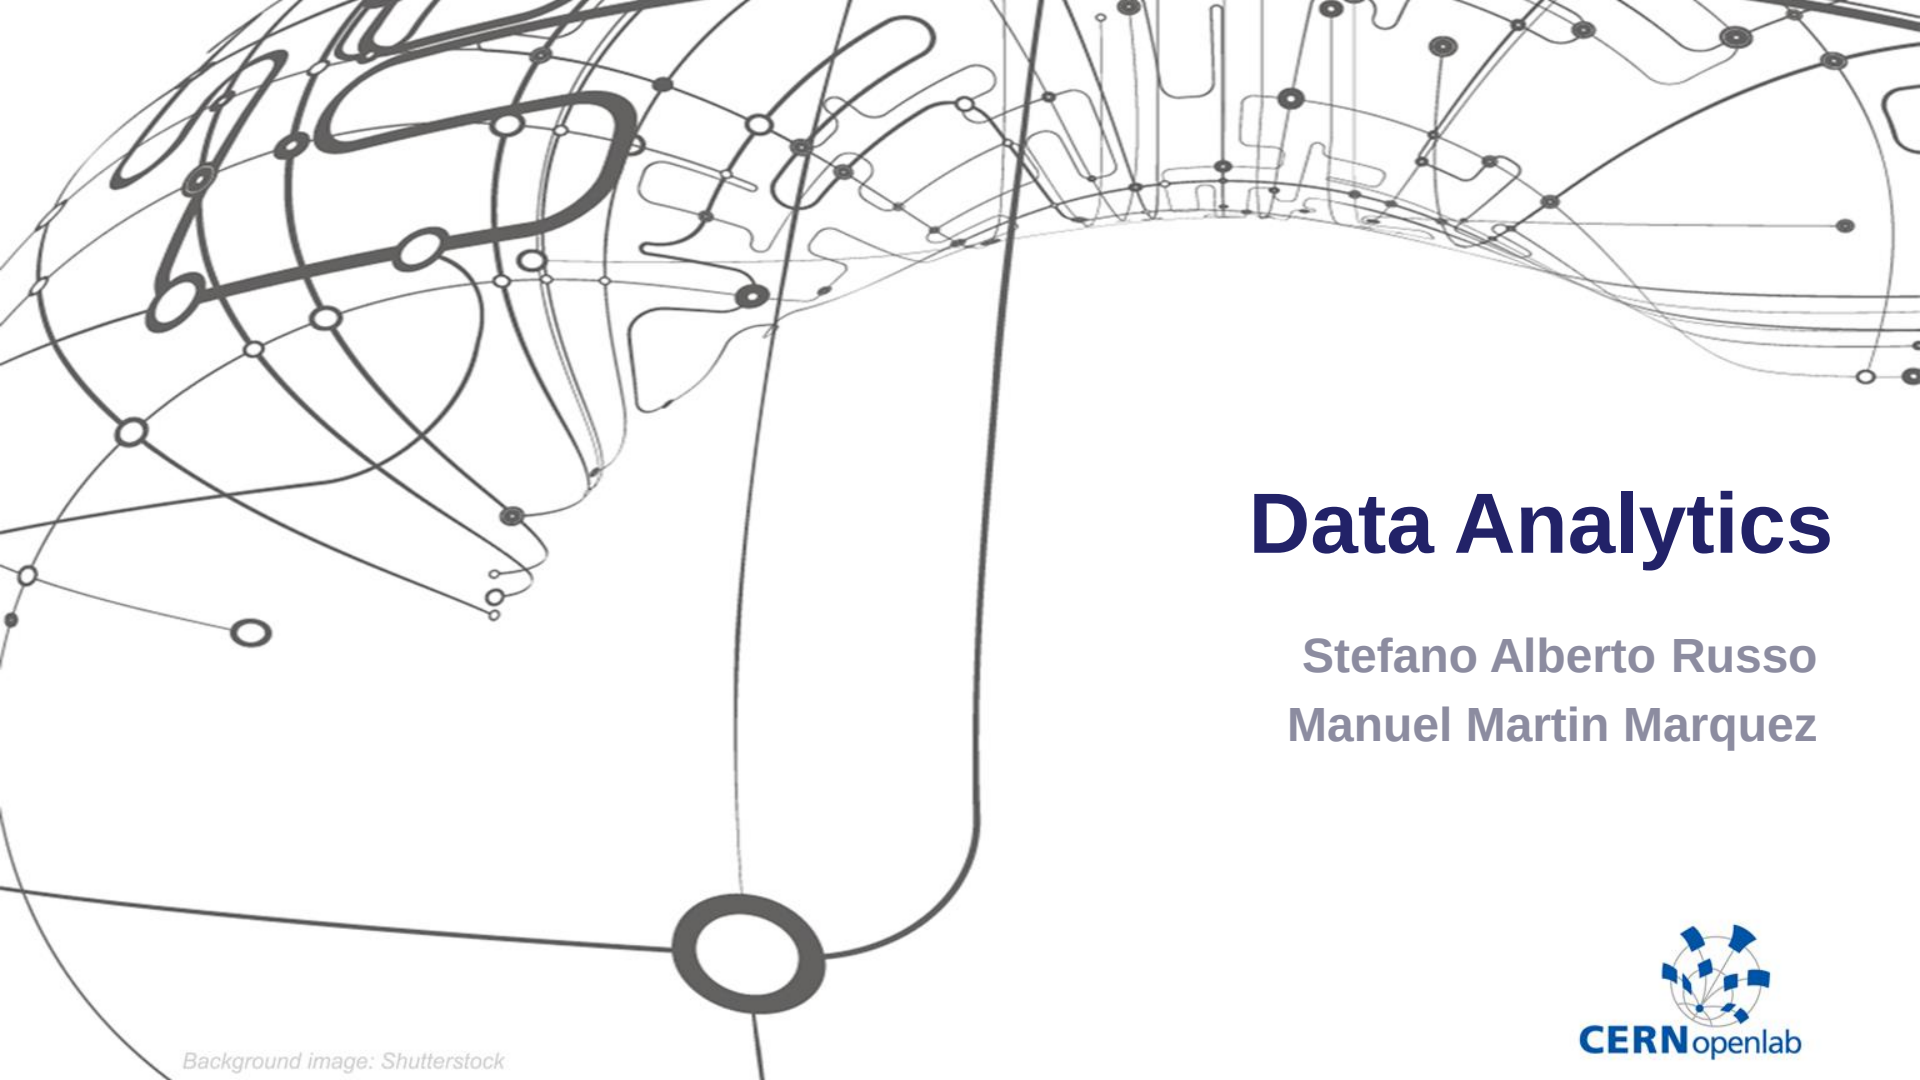

Data Analytics

Stefano Alberto Russo
Manuel Martin Marquez

CERN openlab

Background image: Shutterstock

CASTOR: the CERN Advanced STORage Manager

- Lot of (complex) **log data** recorded in the past years from **various systems**
- **CASTOR TEAM:** *Can we obtain useful information from it?*

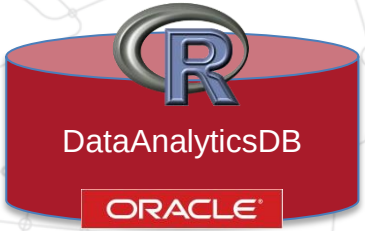

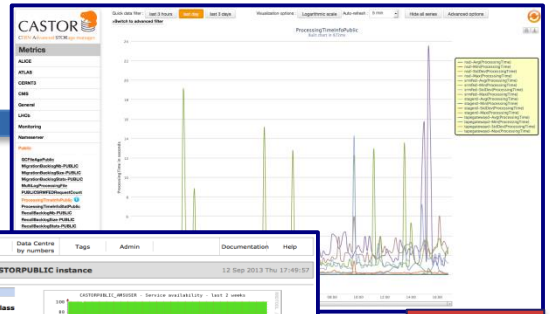

Castor Cockpit

SLS

– Performance
– Cause of errors
– Anomaly detection
– Predictions
– Early warning systems

CENTRALIZE and STANDARDIZE data

CASTOR: the CERN Advanced STORage Manager

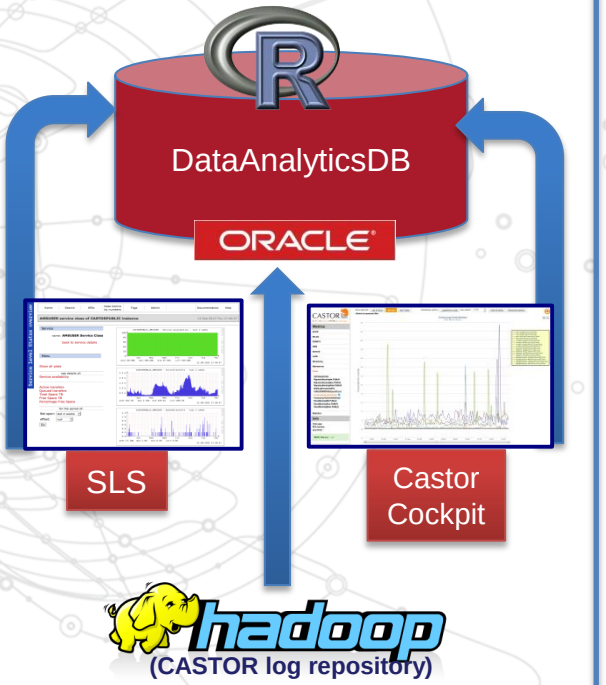

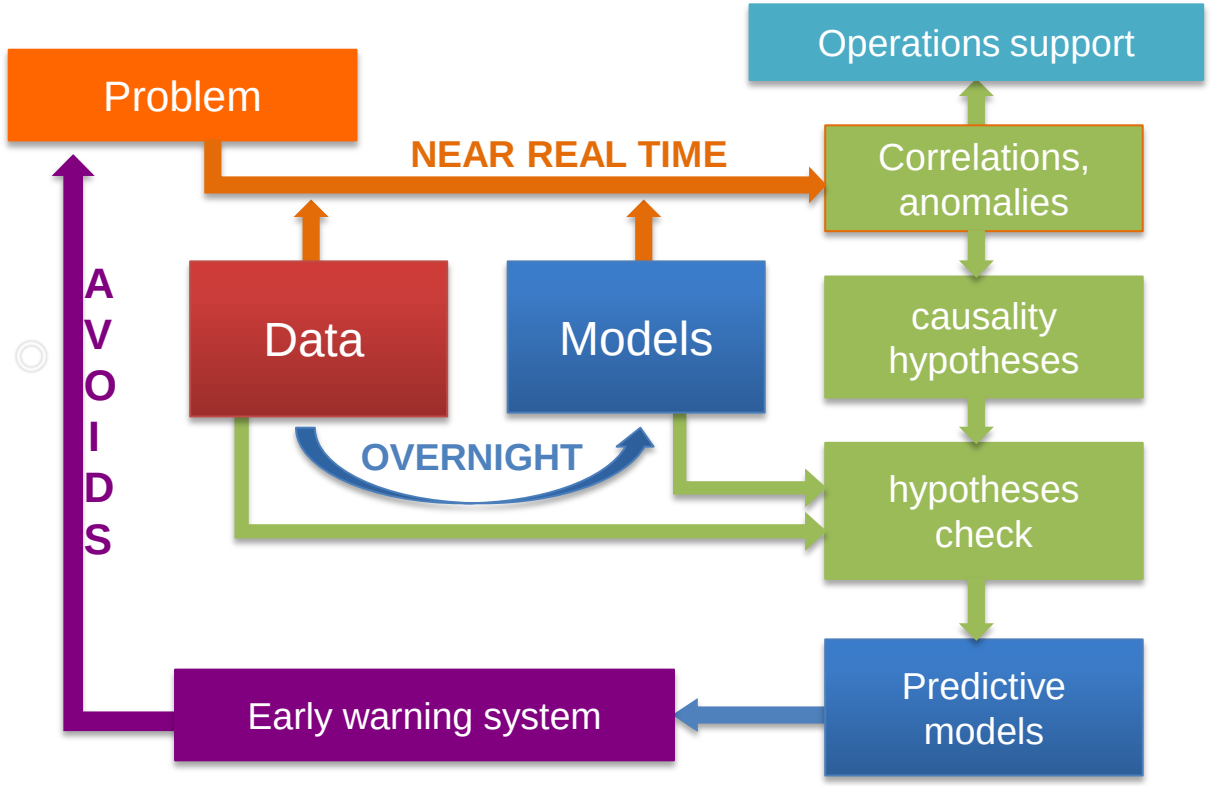

Background image: Shutterstock

Openlab V: Data Analytics Challenges

CERN openlab Workshop on Data Analytics Use Cases

- Organized by EN and IT
- Overview – Data analytics at CERN
- Use case:
  - Context
  - Status
  - Technologies
  - Limitations
- Future Plans and Directions

Openlab V: Data Analytics Challenges

- **Huge interest and potential benefits for CERN**
  - **IT, BE, PH, EN departments**
- **Improve our Monitoring and control systems by mean of Data Analytics**
  - **Intelligent**
  - **Proactive**
  - **Predictive**

Openlab V: Data Analytics Challenges

Challenges

- **Real time analytics based on CERN use case**
  - **Based on domain knowledge and hidden knowledge extracted by batch analytics**
  - **CEP, Storm**
- **Batch analytics**
  - **Correlation analysis**
  - **Forecasting modeling**
  - **Knowledge discovering**
- **Data analytics repository**
- **AaaS**
- **Educational challenge**

CERN-openlab Data Analytics Workshop

20th February 2014

Participants

SIEMENS

esa

Human Brain Project

ORACLE

sas
THE POWER TO KNOW.

Deutscher Wetterdienst
Wetter und Klima aus einer Hand
DWD

blue yonder
Forward looking. Forward thinking.

Yandex

Expedia
Your trip, your way.

Cityzen Data

KIT
Karlsruhe Institute of Technology

Organized

CERN openlab

DATA ANALYTICS

Outreach

- **Swiss Oracle Users Group:**
  - Application Continuity, seen from CERN, Andrei Dumitru
- **Orange (FR) - Oracle Enterprise Manager 12c reference visit –** Andrei Dumitru, Nicolas Marescaux, Daniel Gomez Blanco
- **UKOUG Tech13**
  - Lost Writes, a DBA's Nightmare?, Luca Canali
  - Storage Latency for Oracle DBAs, Luca Canali
  - Next Generation GoldenGate vs. Streams for Physics Data, Z. Baranowski
- **CHEP 2013**
  - An SQL-based approach to Physics Analysis, Maaike Limper

Acknowledgements

- Chris Roderick (BE)
- Emil Pilecki (IT)
- Eric Grancher (IT)
- Luca Canali (IT)
- Ruben Gaspar Aparicio (IT)
- Daniel Gomez Blanco (IT)
- Zbigniew Baranowski (IT)
- Manuel Gonzalez (EN)
- Filipo Tilaro (EN)
- Axel Voitier (EN)
- Philippe Gayet (EN)

www.cern.ch/openlab