

ORACLE®

Oracle: Database and Data Management Innovations with CERN

Public Day

Kevin Jernigan, Oracle
Lorena Lobato Pardavila, CERN
Manuel Martin Marquez, CERN

June 10, 2015

Presented to



Safe Harbor Statement


The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Program Agenda

- 1 Introduction
- 2 Oracle openlab IV achievements
- 3 Oracle strategic direction
- 4 Oracle openlab V projects
- 5 Conclusions

Program Agenda

- 1 Introduction
- 2 Oracle openlab IV achievements
- 3 Oracle strategic direction
- 4 Oracle openlab V projects
- 5 Conclusions




YEAR
1977

Oracle Begins as SDL

Software Development Laboratories



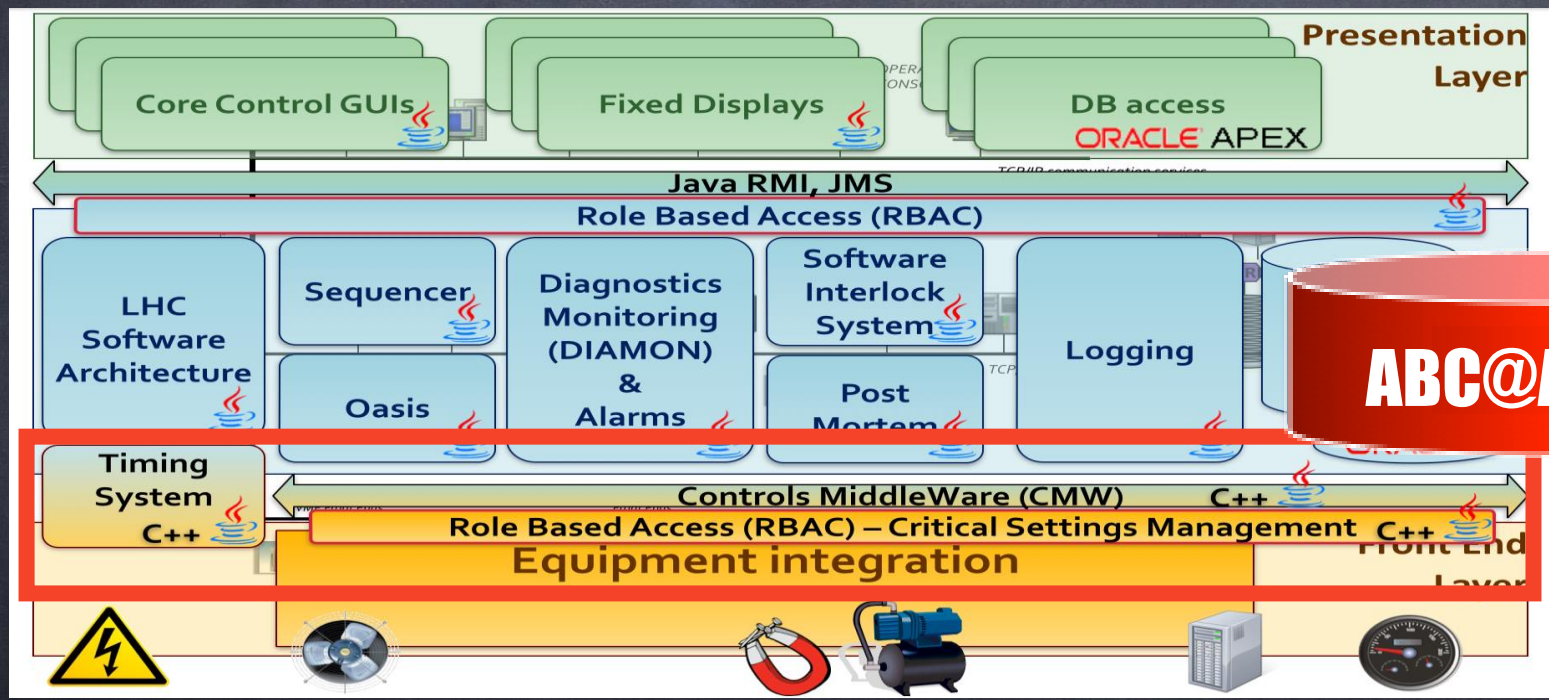
“About five years into the company, it became pretty clear that the horizons were unlimited. The only limitations were us” – Larry Ellison



Accelerator Controls Configuration

configuration of:

ALL Control System Entities & their relationships



Since 1985!


ABC@ACCCON

Heterogeneous Entities

Credit: Chris Roderick (CERN)

Market Leader from Continuous Innovation

Changing dynamics



Client-Server



Internet



Big Data & Cloud

Release

1985 – 1997: 6, 7 and 8

1998 – 2012: 8i, 9i, 10g, 11g

2013-2014: 12c

Developer

Stored Procedures & Triggers
Referential Integrity
Distributed Transactions
AQ
LOBs
Spatial

Java .NET PHP XML APEX

R
Pattern Matching
OpenSource Drivers
JSON
REST Data Services
NoSQL Database
Application Continuity
Migration Framework

Engine

OLTP throughput
Parallel Query, Partitioning

Online Operations
RAC
Data Guard
Flashback
Self-Managing Database
Enterprise Manager
Resource Management


Automatic Storage Mgmt
Encryption
Real Application Testing
Row Compression
Columnar Compression
Smart Scans
Flash Cache

Multitenant Database In-Memory Column Store

ORACLE®

Market Leader from Continuous Innovation

Responsive to customer demands and industry trends



Oracle Database 10g

- Declarative Referential Integrity
- PLSQL Triggers and Stored Procedures
- Two Phase Commit
- Parallel Operations
- Object Relational Support
- Multimedia Support
- Real Application Clusters
- Materialized Views
- XML Database
- Data Guard
- Flashback Query
- Virtual Private Database
- Built in Java VM

Oracle Database 11g

- Automatic Storage Management
- Transparent Data Encryption
- Self Managing Database
- Automatic Workload Repository
- SQL Tuning Advisor and ADDM
- Flashback Table and Database
- Application Express

Oracle Database 12c

- Active Data Guard
- Advanced Compression
- Real Application Testing
- Hybrid Columnar Compression
- Exadata Smart Scans
- Exadata Smart Flash Cache
- Exadata I/O Resource Management

Oracle Database 12c

- Multitenant
- Heat Map & Automatic Data Optimization
- Data Redaction
- Pattern Matching
- Global Data Services
- Application Continuity



Program Agenda

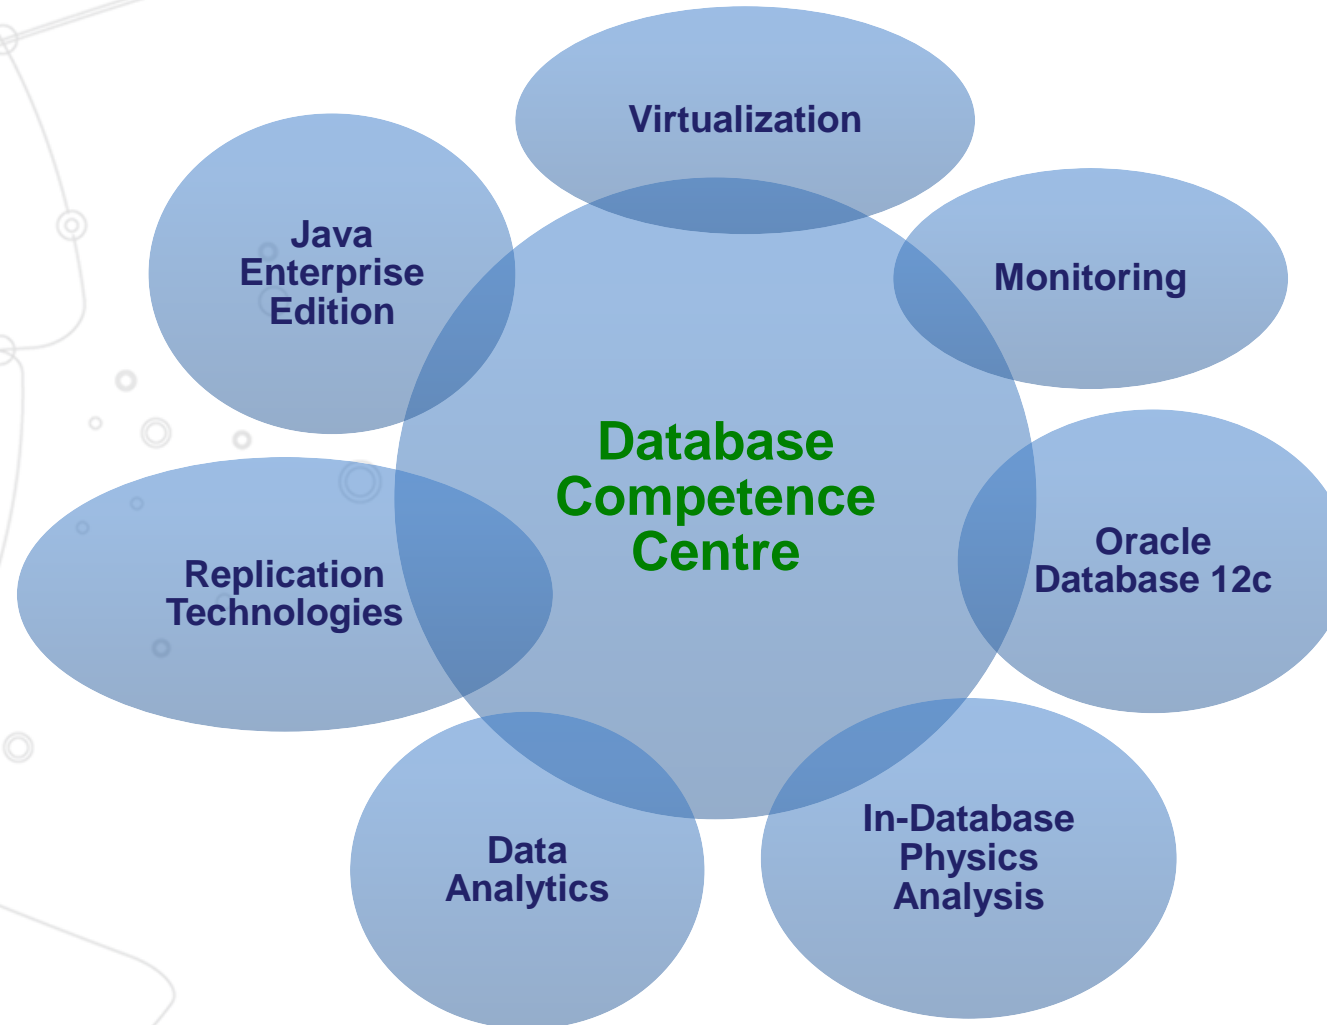
- 1 Introduction
- 2 Oracle openlab IV achievements**
- 3 Oracle strategic direction
- 4 Oracle openlab V projects
- 5 Conclusions

openlab IV

Database Competence Centre

- › **Lorena Lobato Pardavila**
- › **Andrei Dumitru**
- › **Ignacio Coterillo Coz**
- › **Antonio Romero Marin**
- › **Manuel Martin Marquez**
- › **Luis Rodriguez Fernandez**
- › **Maaike Limper**
- › **Stefano Russo**

Openlab IV projects



Database Competence Centre Achievements

Virtualization

- Performance tests of Intel SR-IOV network cards on Oracle virtual platforms
- Integration of the Oracle VM hypervisor as a compute provider on the OpenStack Cloud Operating System

Monitoring

- Migrated from Oracle Enterprise Manager 11g to version 12c
- Integrated with CERN Agile Infrastructure (Active Directory)

Oracle Database 12c

- Early validation of new functionalities thanks to involvement in Beta program
- Oracle Database 12c being deployed in production

In-Database Physics Analysis

- Proven for a subset of the SQL-based queries
- Better understanding of Oracle Database In-Memory

Database Competence Centre Achievements

Java Enterprise Edition

- Java Web Hosting Service (Middleware on Demand)
- Upgrade 500 Oracle Weblogic Servers to 12c
- Oracle Weblogic and APEX SSO integration

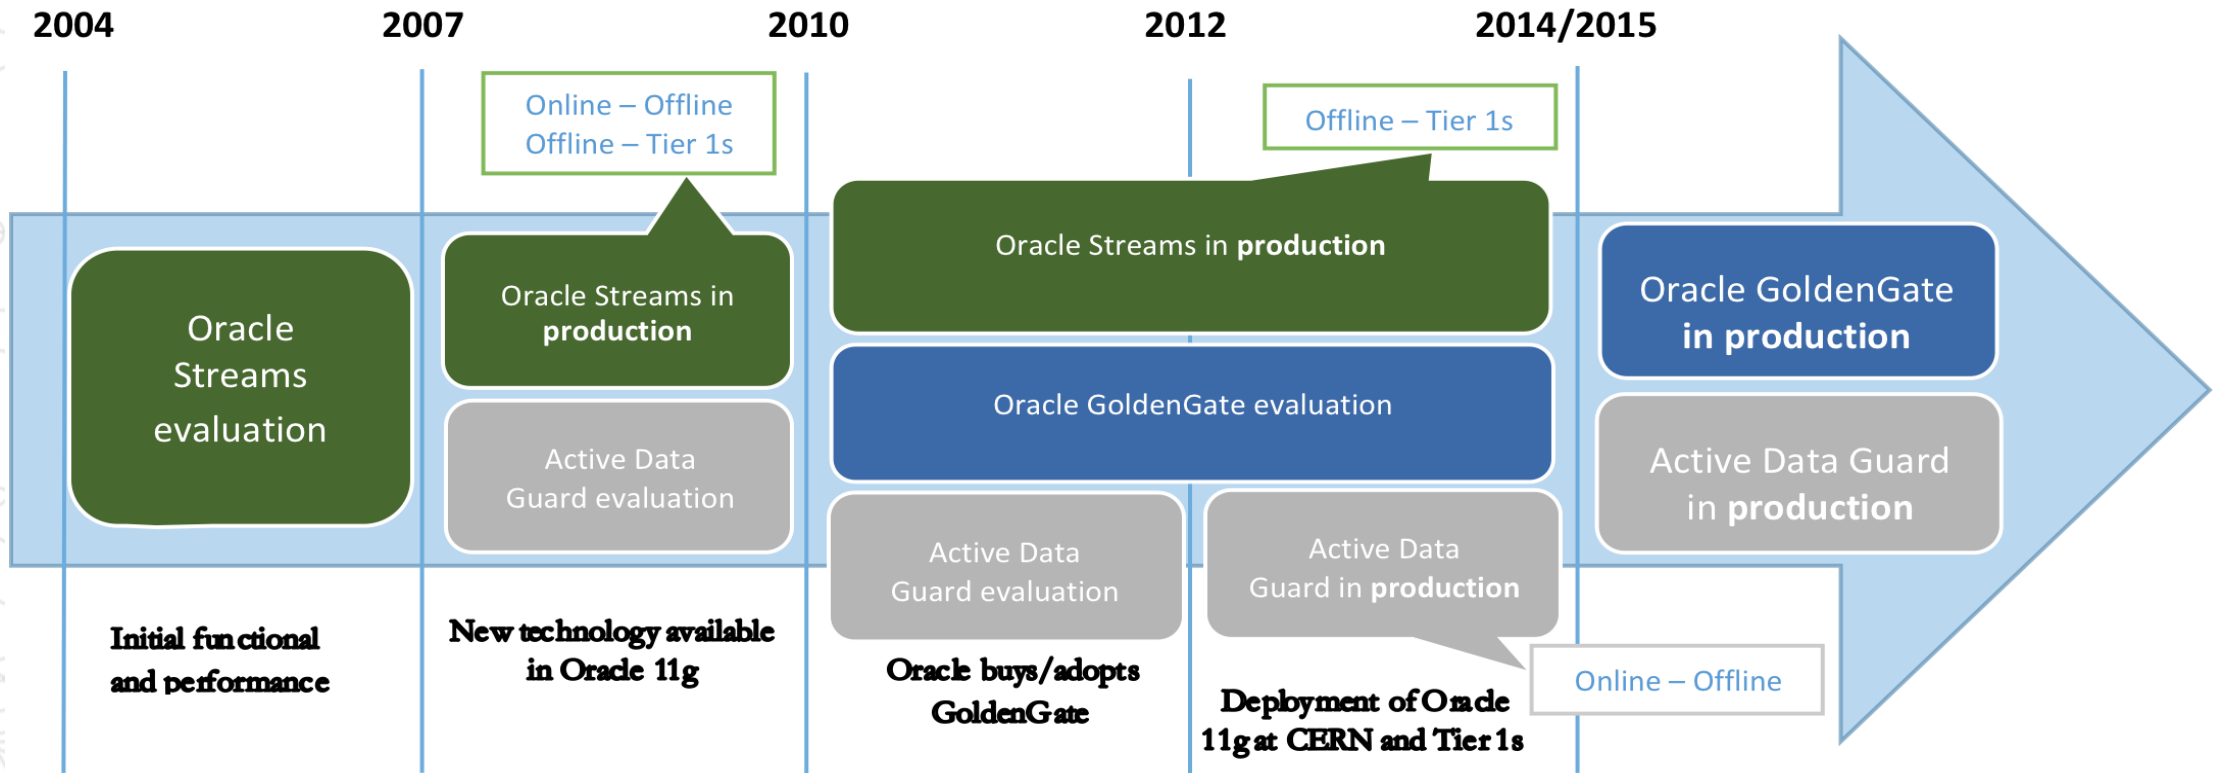
Data Analytics

- Oracle R Enterprise added new functionalities proposed by CERN
- Applied Oracle Endeca Information Discovery to explore future possibilities in control and monitoring improvement

Replication Technologies

- Oracle GoldenGate evolution to fulfill CERN needs
- Smooth migration from Oracle Streams to Oracle GoldenGate

Replication Technologies Evolution



Replication Technologies Evolution

Feedback to Oracle GoldenGate development team

- Oracle GoldenGate 11g did not fulfill CERN needs
- Oracle GoldenGate 12c contained new functionalities as suggested by CERN

Oracle GoldenGate 12c fit in CERN use cases

- Improvements in transactions dependencies (including DDL)
- Data Filtering based on session tagging
- Offloading primary system with downstream capture database
- Native DDL support (Table creation, granting of user privileges)

Replication Technologies Environment


- Migration completed from Oracle Streams to Oracle GoldenGate in our productions databases during 2014
- Centralized system at CERN for physics data distribution in WLCG

Program Agenda

- 1 Introduction
- 2 Oracle openlab IV achievements
- 3 Oracle strategic direction**
- 4 Oracle openlab V projects
- 5 Conclusions

Oracle Database Innovations for each Computing Era

Now innovating for big data and the cloud



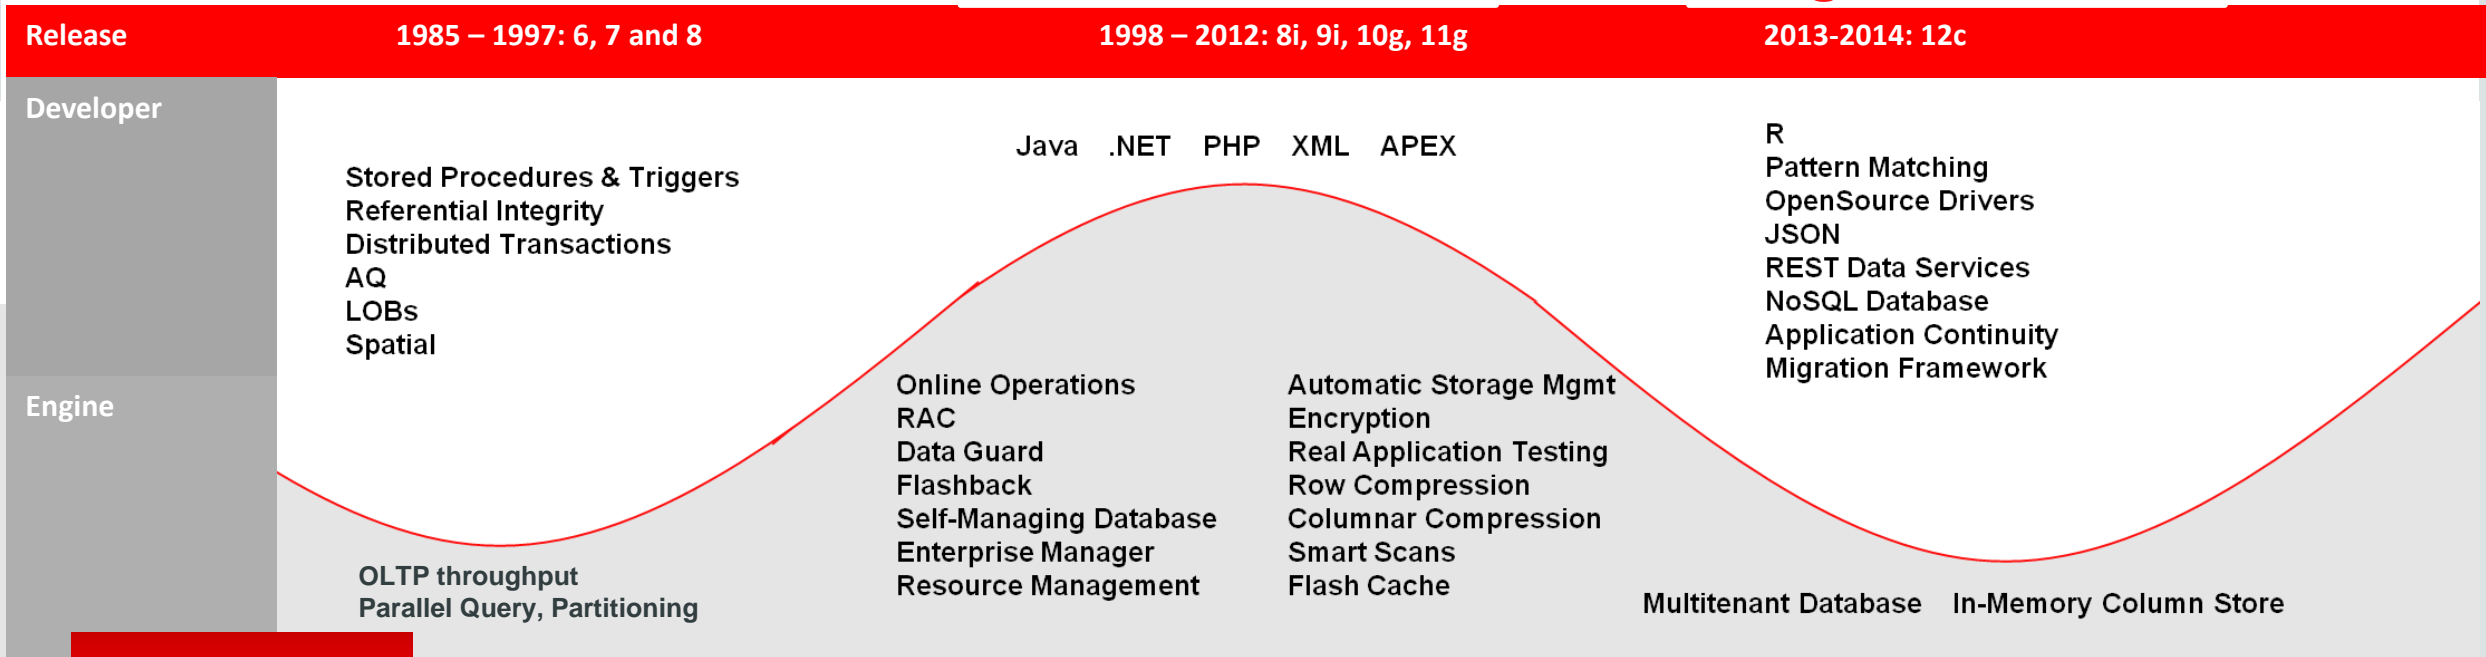
Client-Server



Internet



Big Data & Cloud



Cloud Services




Apps

Software as a Service (SaaS)



Platform

Platform as a Service (PaaS)




Infrastructure

Infrastructure as a Service (IaaS)

Oracle Cloud


Full stack enterprise cloud services SaaS + PaaS + IaaS



<https://cloud.oracle.com>



Enterprise Big Data Architecture



Oracle Big Data Discovery. The Visual Face of Hadoop

The interface is divided into several sections:

- Top Left:** Navigation and filters for 9 Projects and 89 data sets. Includes a 'Refine By' sidebar with categories like USAGE, CONTENT, and METADATA.
- Top Center:** 'DATA SETS REFINEMENTS' and 'DATA SETS' panels. The 'DATA SETS' panel lists items like 'ChicagoCensus2011' and 'ChicagoCrime2011' with their last modified dates.
- Top Right:** Overview of 1M of 3.4M records, 1M records viewed, and 74 attributes. It features several charts: 'Date reviewed' (line), 'Score' (line), 'Region' (bar), 'Price' (line), and 'Zipcode' (map).
- Bottom Left:** 'Sales Analysis: Customer Churn' section with a bubble chart titled 'Record Count, Record Count, Record Count by p_body'. The chart shows a positive correlation between Record Count (X-axis) and Record Count (Y-axis).
- Bottom Right:** A detailed view of a data set with a table of records. The table includes columns for date reviewed, name, year, description, drinkability, winery, region, and price. A 'Description' panel on the left provides details for a 'Best Buy' wine.

In the center of the interface is a circular diagram with four quadrants: **Find**, **Explore**, **Discover**, and **Transform**. The quadrants are connected by arrows in a clockwise cycle, with a central icon of a person at a computer.



Program Agenda

- 1 Introduction
- 2 Oracle openlab IV achievements
- 3 Oracle strategic direction
- 4 Oracle openlab V projects**
- 5 Conclusions

CERN openlab V

Continuous Innovation

Java Web

- Exploit new Java and WebLogic features for improving Java Infrastructure
 - Developer Productivity
 - Improve application consolidation density

Cloud Services

- Private, Hybrid and Public
 - Cloud Database back-up and recovery
 - Transparent application deployment between CERN and public Cloud

Database Evolution

- CERN's and High Energy Physics Oriented
 - RAC evolution
 - In-memory, data "sharding"
 - High availability monitoring

Data Analytics

- Data Analytics as a Service Infrastructure
 - Joint effort between CERN and openlab members

Data Analytics as a Service Feature and Challenges

Centralize and standardize Data Analytics requirement

- Flexible platform for CERN and external institution
- Different actors and Data Analytics requirements

Provide Storage for large data volume


- Extremely Heterogeneous data environment
- Structured and unstructured

Integration with existing data repositories

- Large number of mission-critical data systems need to be integrated
- Logging, fault tracking, major events, etc.

Deliver real-time, batch analytics and information discovery

- Batch predicted models deployed for real time monitoring and control
- Deliver self-service Information Discovery



LHCb

ATLAS

CERN Meyrin

CERN Prévessin

SPS 7 km

ALICE

SUISSE
FRANCE

CMS

LHC 27 km

Future Circular Collider Study Reliability, Availability & Maintainability

Model the accelerator complex and to simulate the particle acceleration process.

- FCC RAMS studies will be needed to assess the feasibility of such a large scale project

Reliability and Availability Study

- Key indicators for the performance of circular colliders
- For the HL-LHC availability will be one of the fundamental

Collaboration project

- FCC, Tampere University of Technology and CERN openlab

Program Agenda

- 1 Introduction
- 2 Oracle openlab IV achievements
- 3 Oracle strategic direction
- 4 Oracle openlab V projects
- 5 Conclusions**

Hardware and Software Engineered to Work Together

ORACLE®