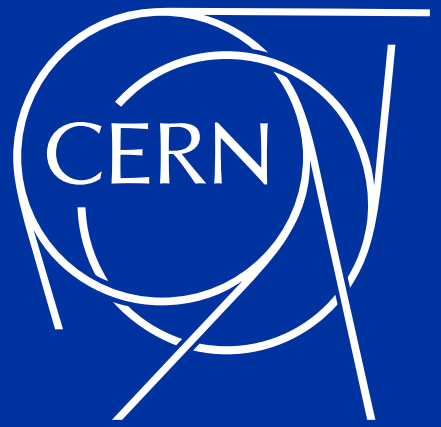

Upcoming database developments at CERN

Andrzej Nowicki

CHEP, Kraków, 2024

Who am I?

Andrzej Nowicki

12 years of Oracle DB experience
Database Engineer @ CERN, IT-DA-DB
Technical Manager of the CERN School of Computing

andrzej.nowicki@cern.ch

Databases at CERN

Oracle since 1982

- 105 Oracle databases, more than 11.800 Oracle accounts
- RAC, Active Data Guard, GoldenGate, OEM, RMAN, APEX, Cloud…
- Complex environment

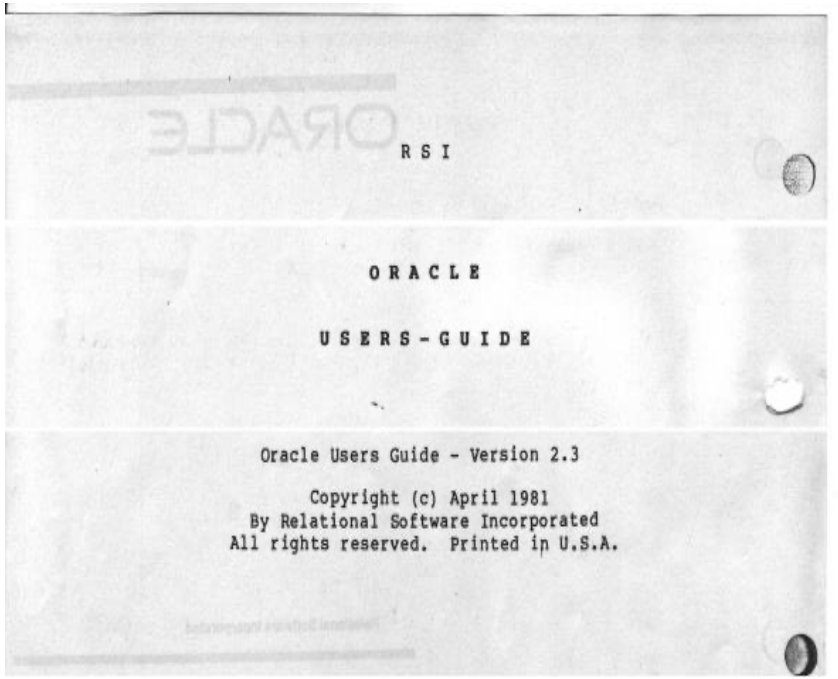
RSI

ORACLE

USERS-GUIDE

Oracle Users Guide - Version 2.3

Databases at CERN

Oracle since 1982

- 105 Oracle databases, more than 11.800 Oracle accounts
- RAC, Active Data Guard, GoldenGate, OEM, RMAN, APEX, Cloud…
- Complex environment

Database on Demand (DBoD) since 2011

- ≈600 MySQL, ≈400 PostgreSQL, ≈200 InfluxDB
- Automated backup and recovery services, monitoring, clones, replicas
- HA MySQL clusters
  (ProxySQL-based -> InnoDB cluster)

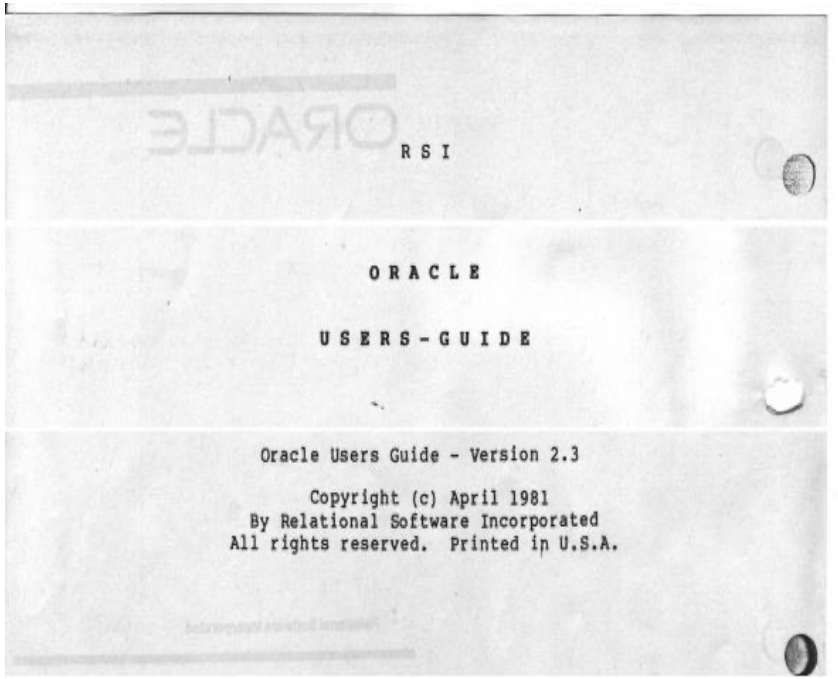
RSI

ORACLE

USERS-GUIDE

Oracle Users Guide - Version 2.3

Size of the database environment

	Total size
Oracle	≈5 PB
DBoD (MySQL, PostrgeSQL, InfluxDB)	≈150 TB
Backups	≈3 PB

Size of the database environment – Oracle

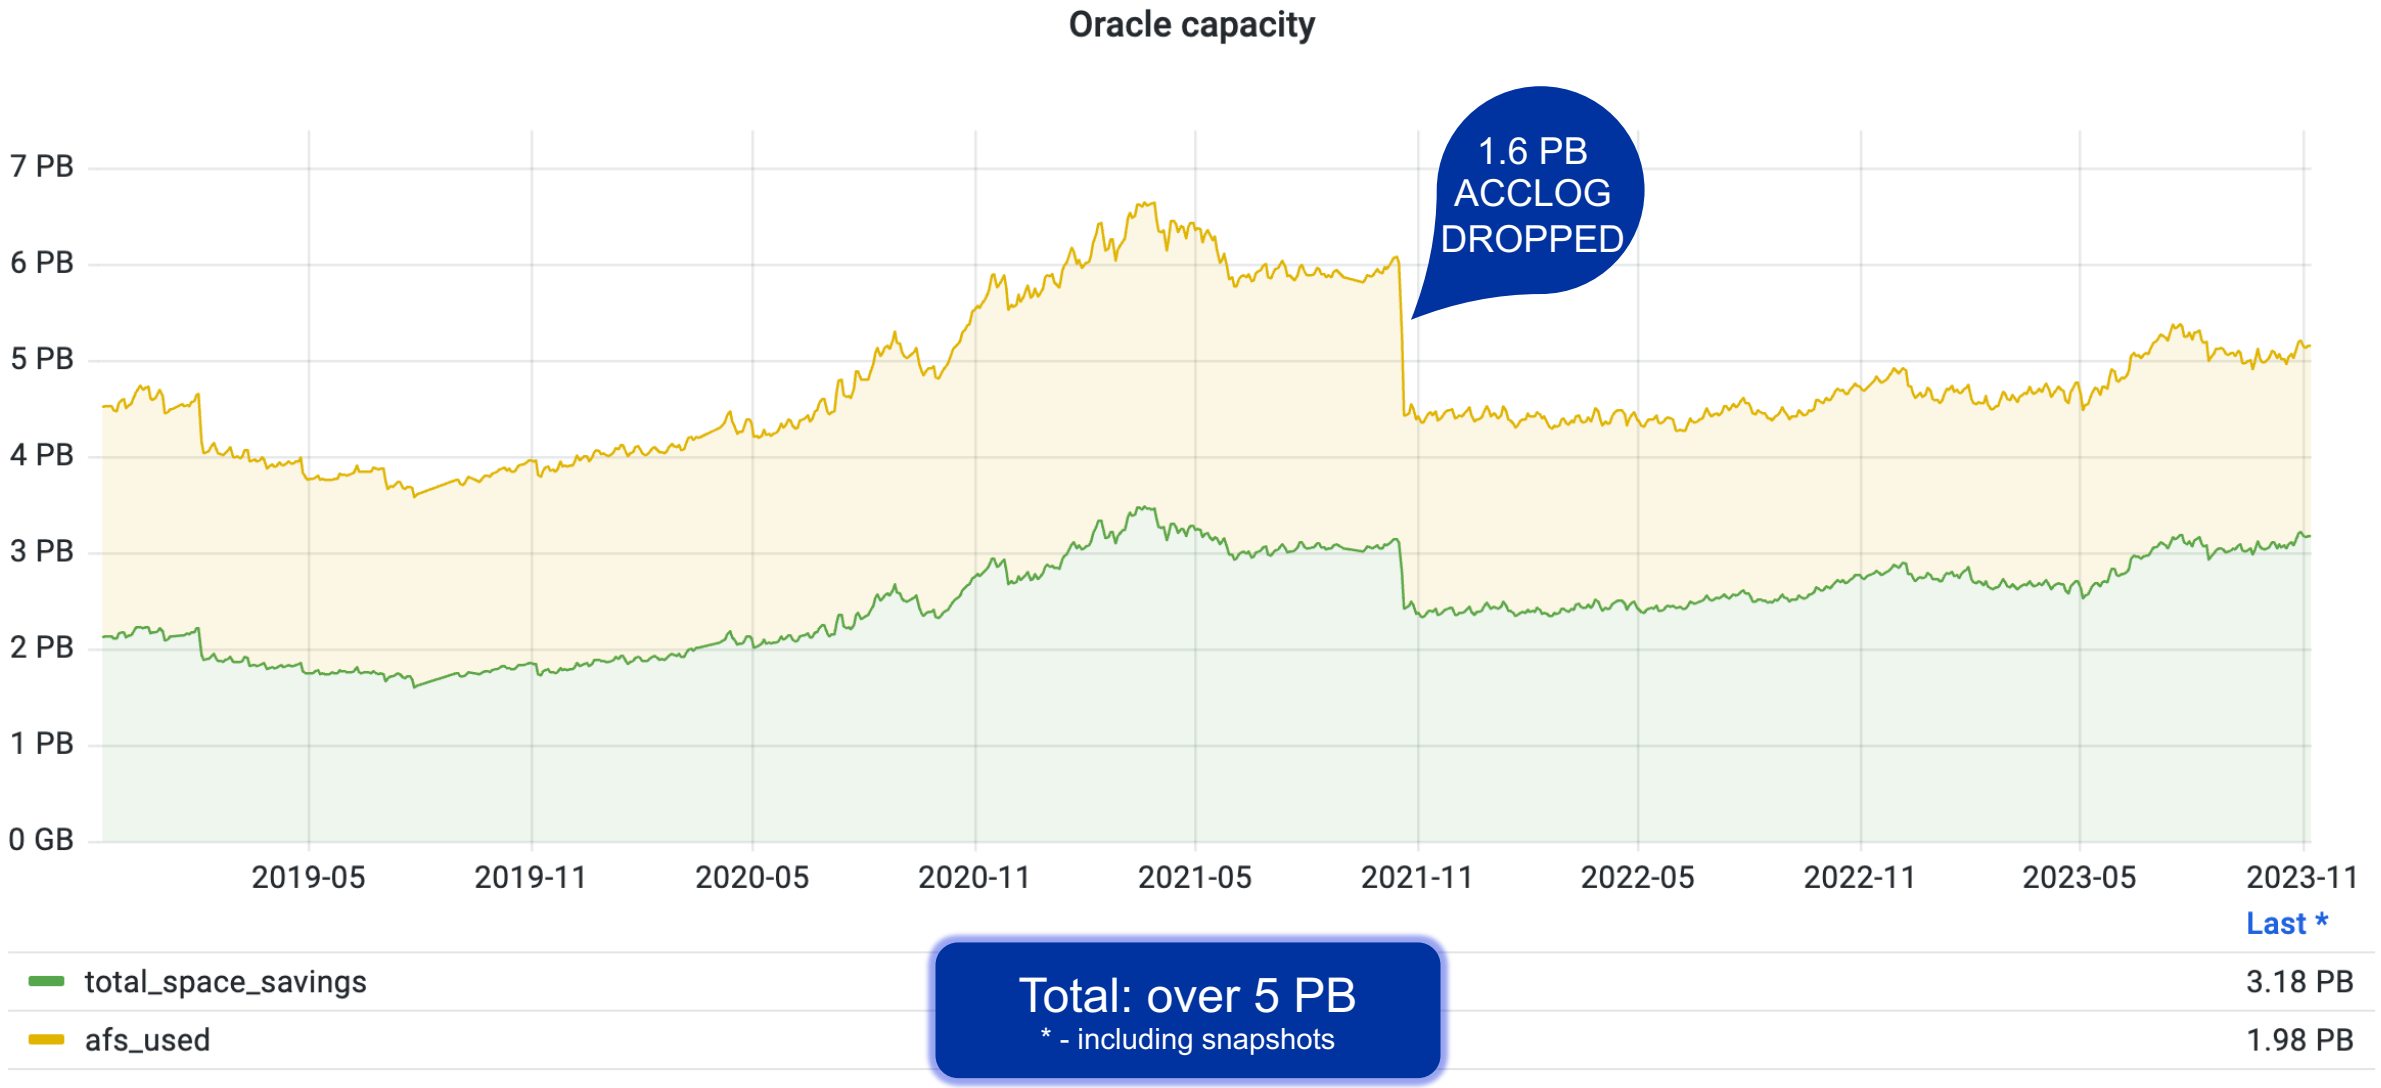

Hardware is constantly improving (CPU)

Luca Canali

RAC52 – 20 cores
RAC54 – 8 cores
RAC55 – 16 cores

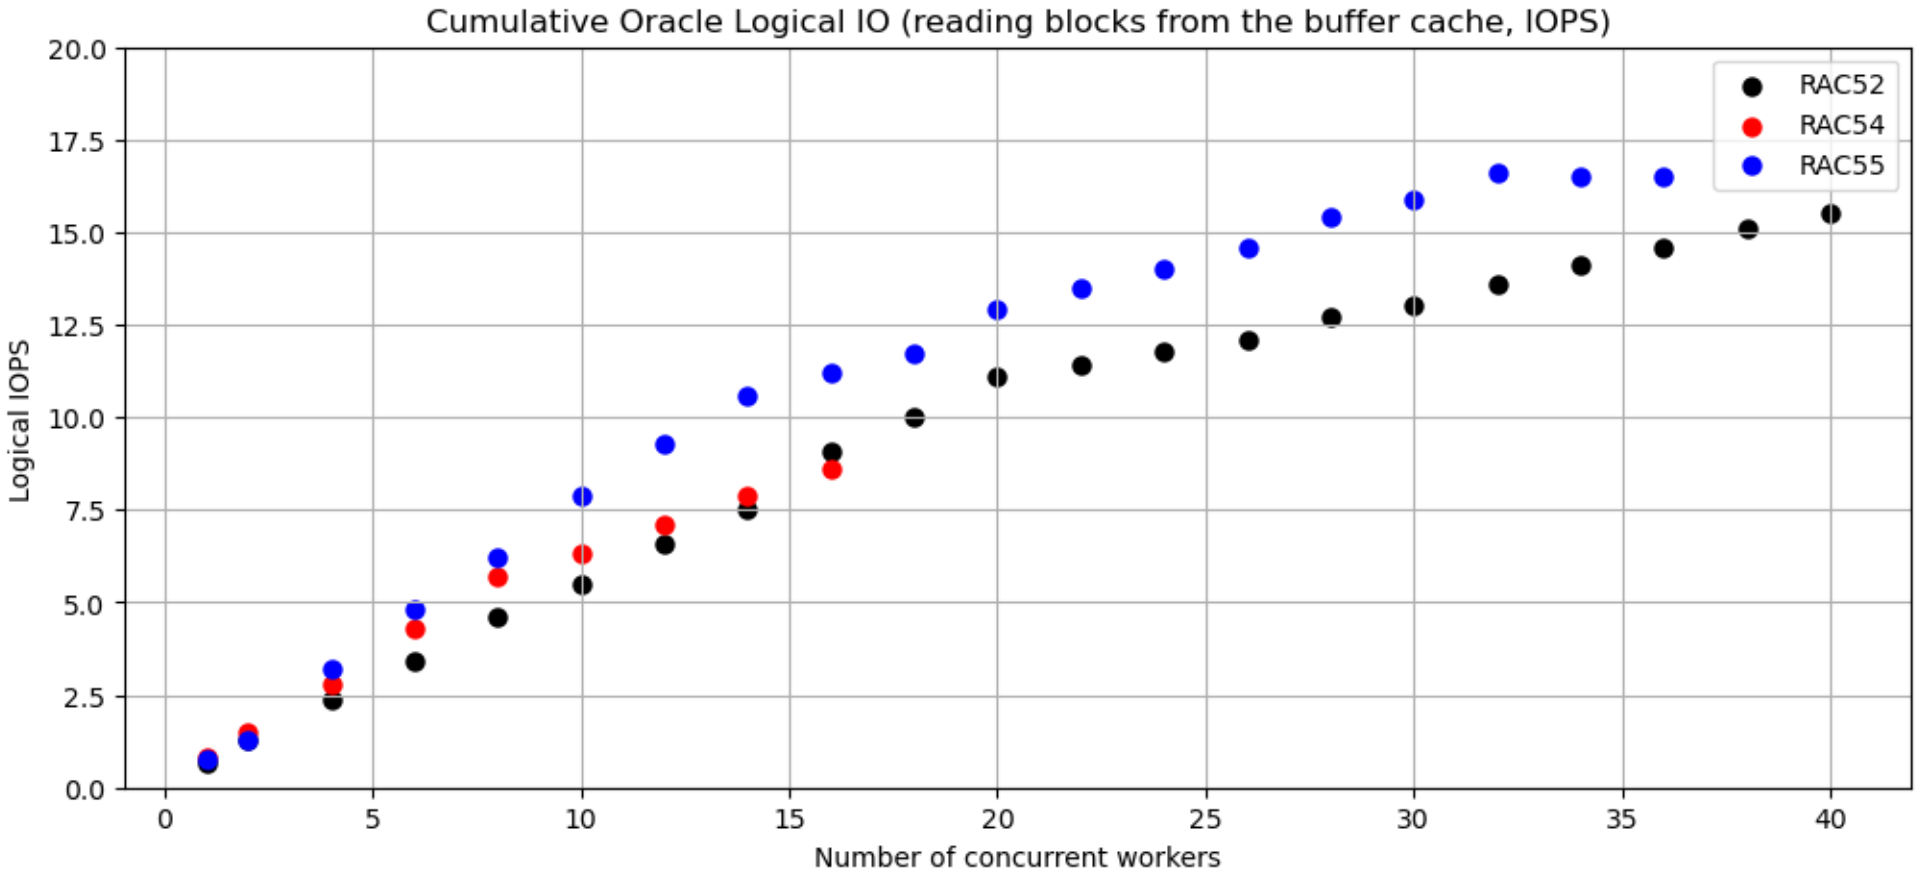

RAC55 has about 1.5x to 1.2x single-thread performance increase over RAC52

Hardware is constantly improving (storage)

Borja Aparicio Cotarelo

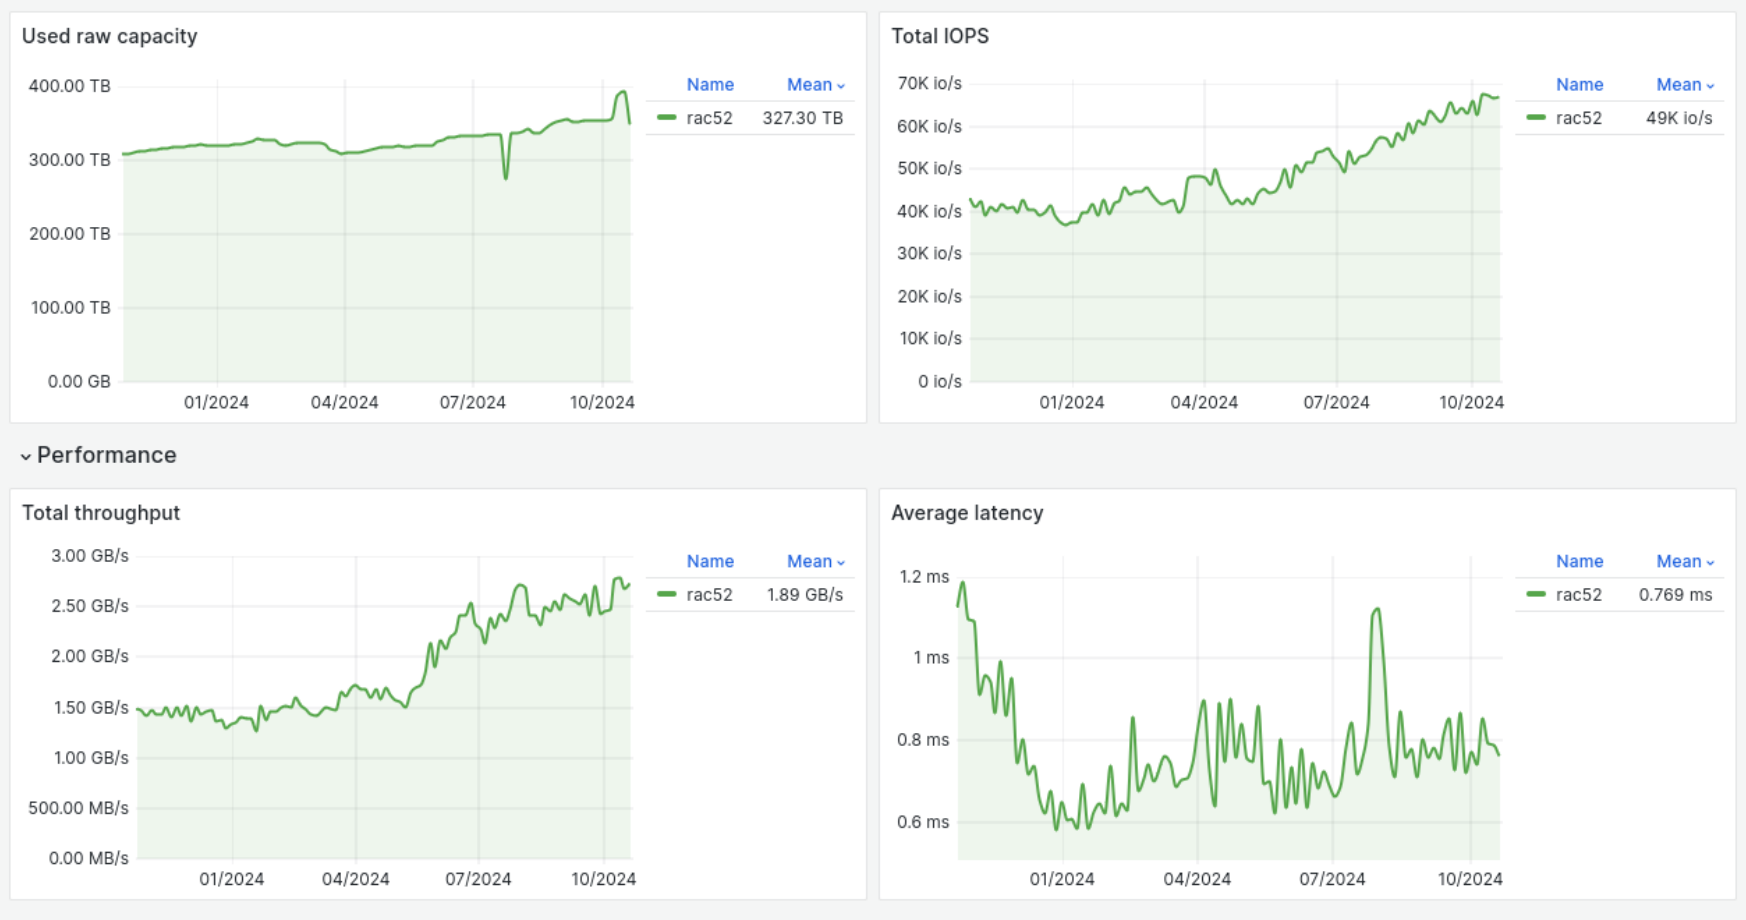

We continuously monitor the system load and optimize distribution across storage resources.

With the decreasing cost of SSDs, we're able to improve performance over time.

Software is constantly ~~improving~~ changing

Database software is released quite often (Oracle releases quarterly patches)

- **We patch often to comply with security regulations**

Vendor support schedule vs LHC schedule

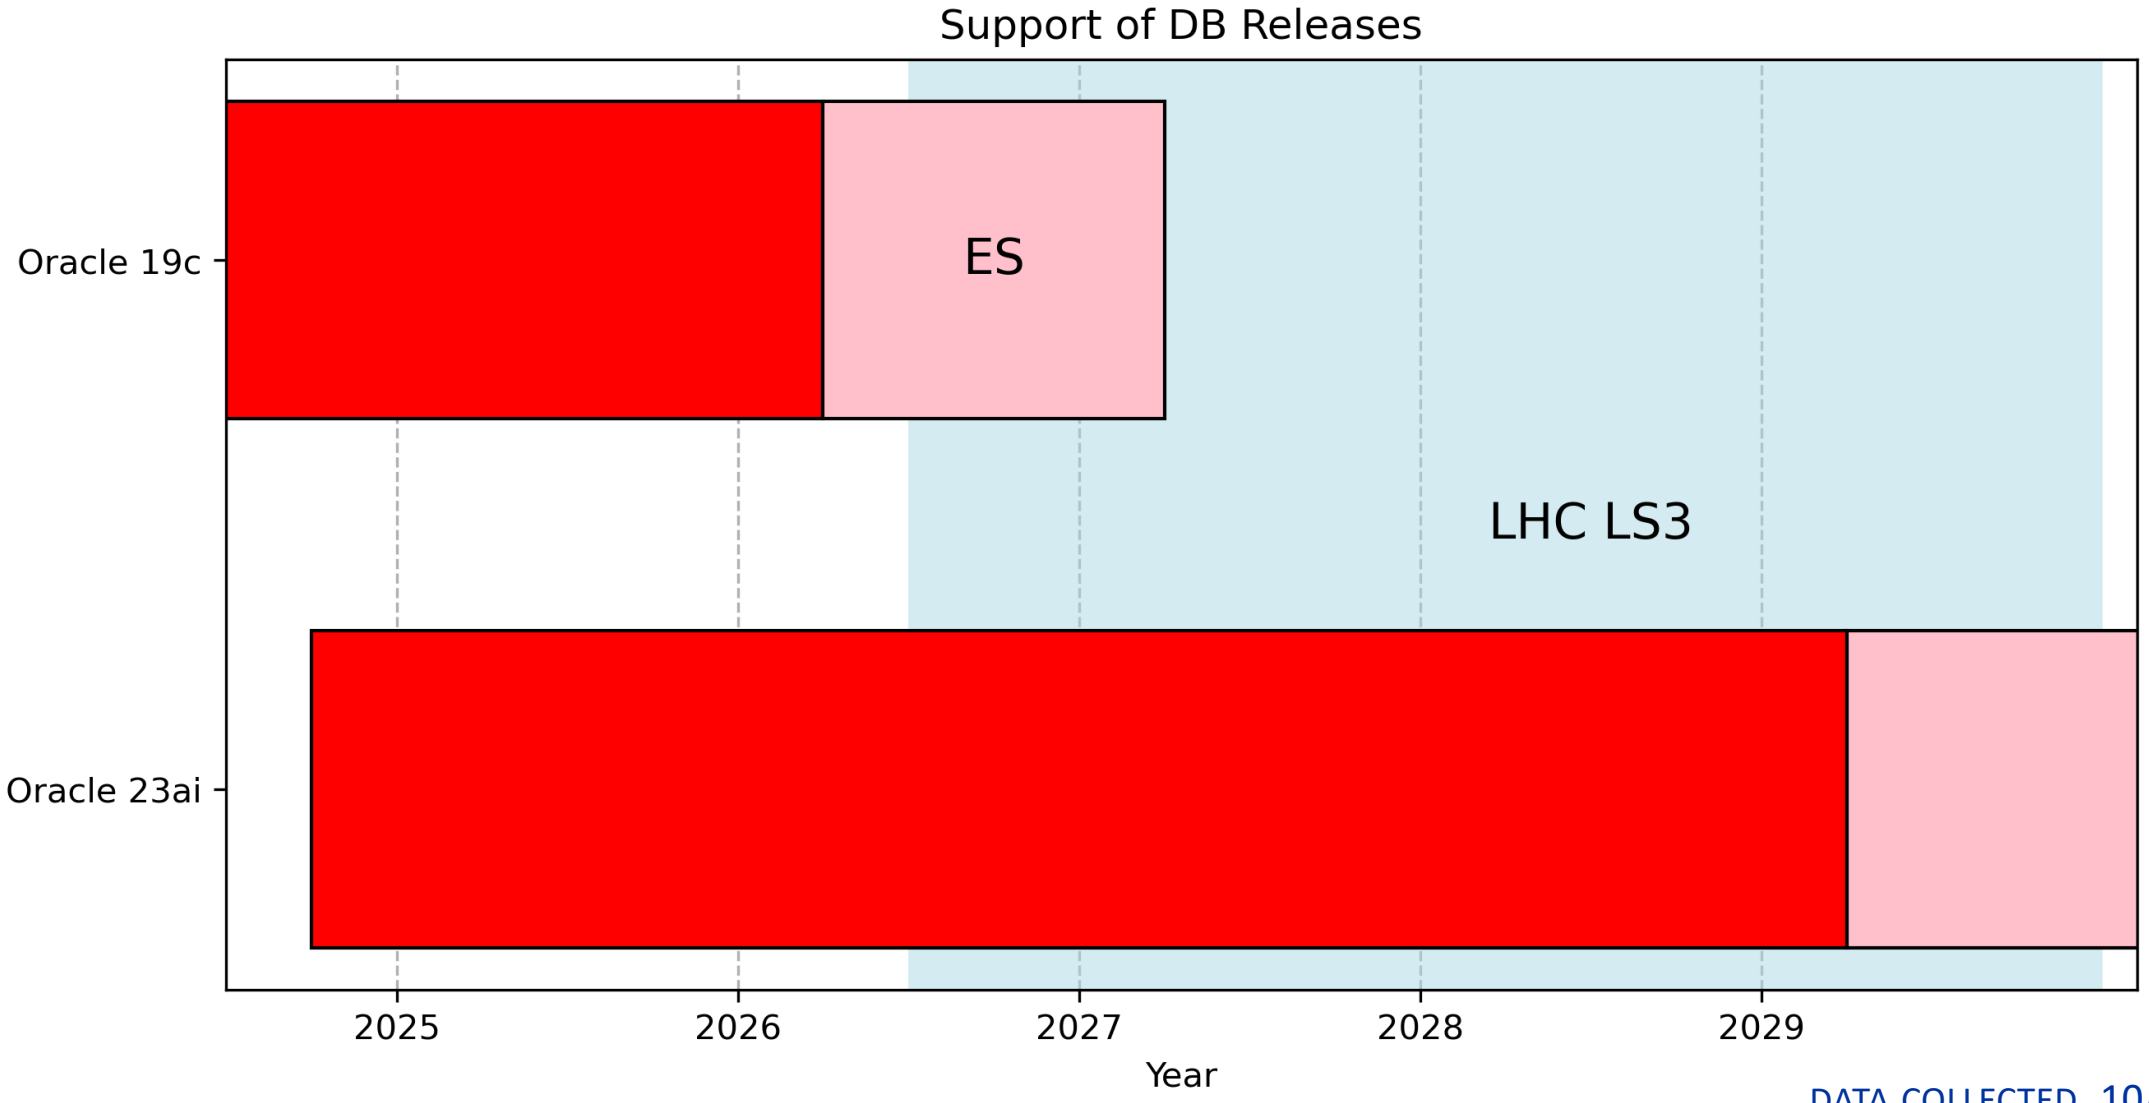

DATA COLLECTED 10/2024

Vendor support schedule vs LHC schedule

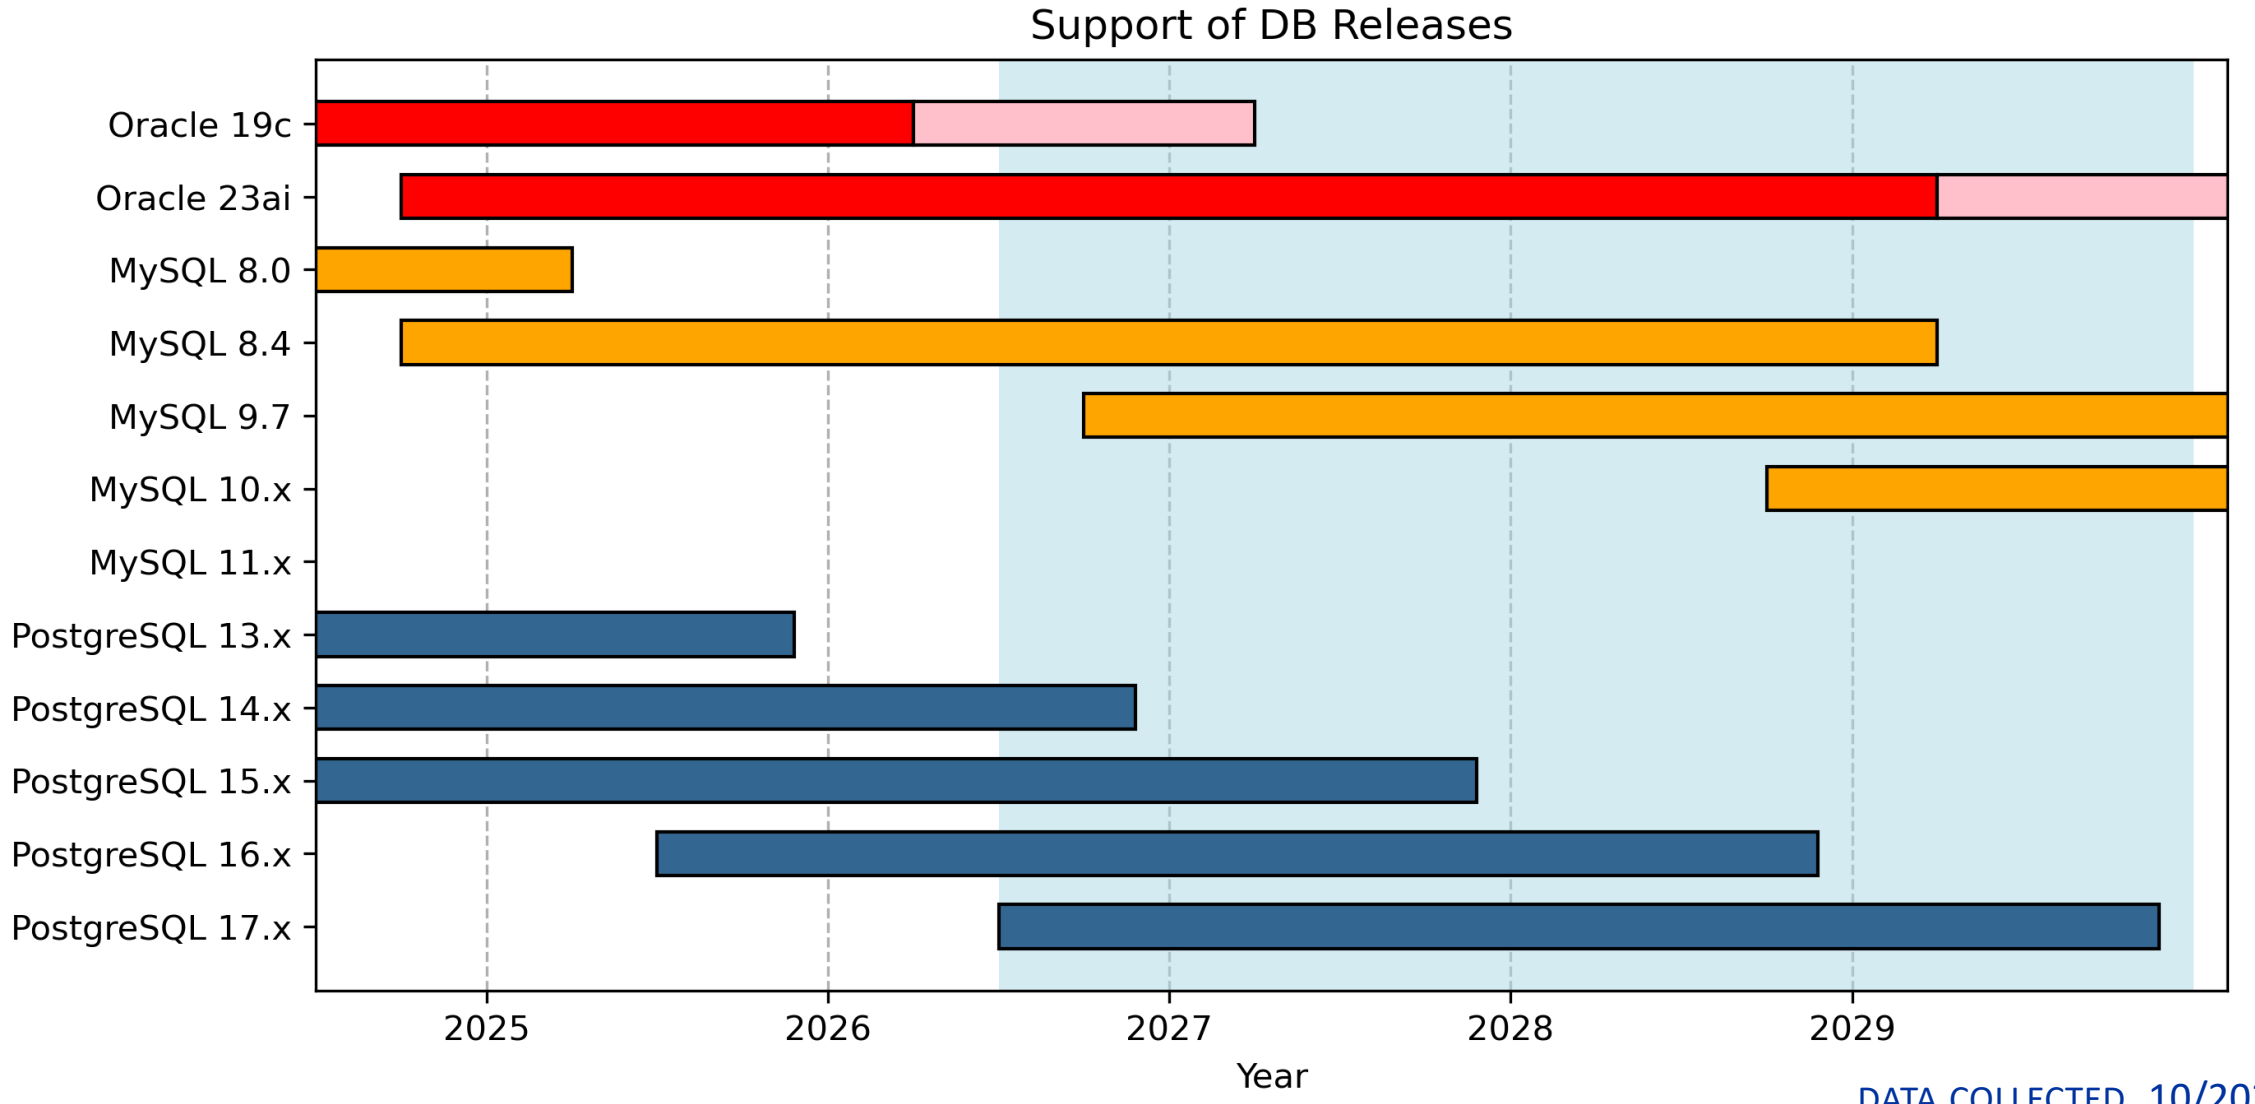

DATA COLLECTED 10/2024

The 23ai brings a move to container architecture

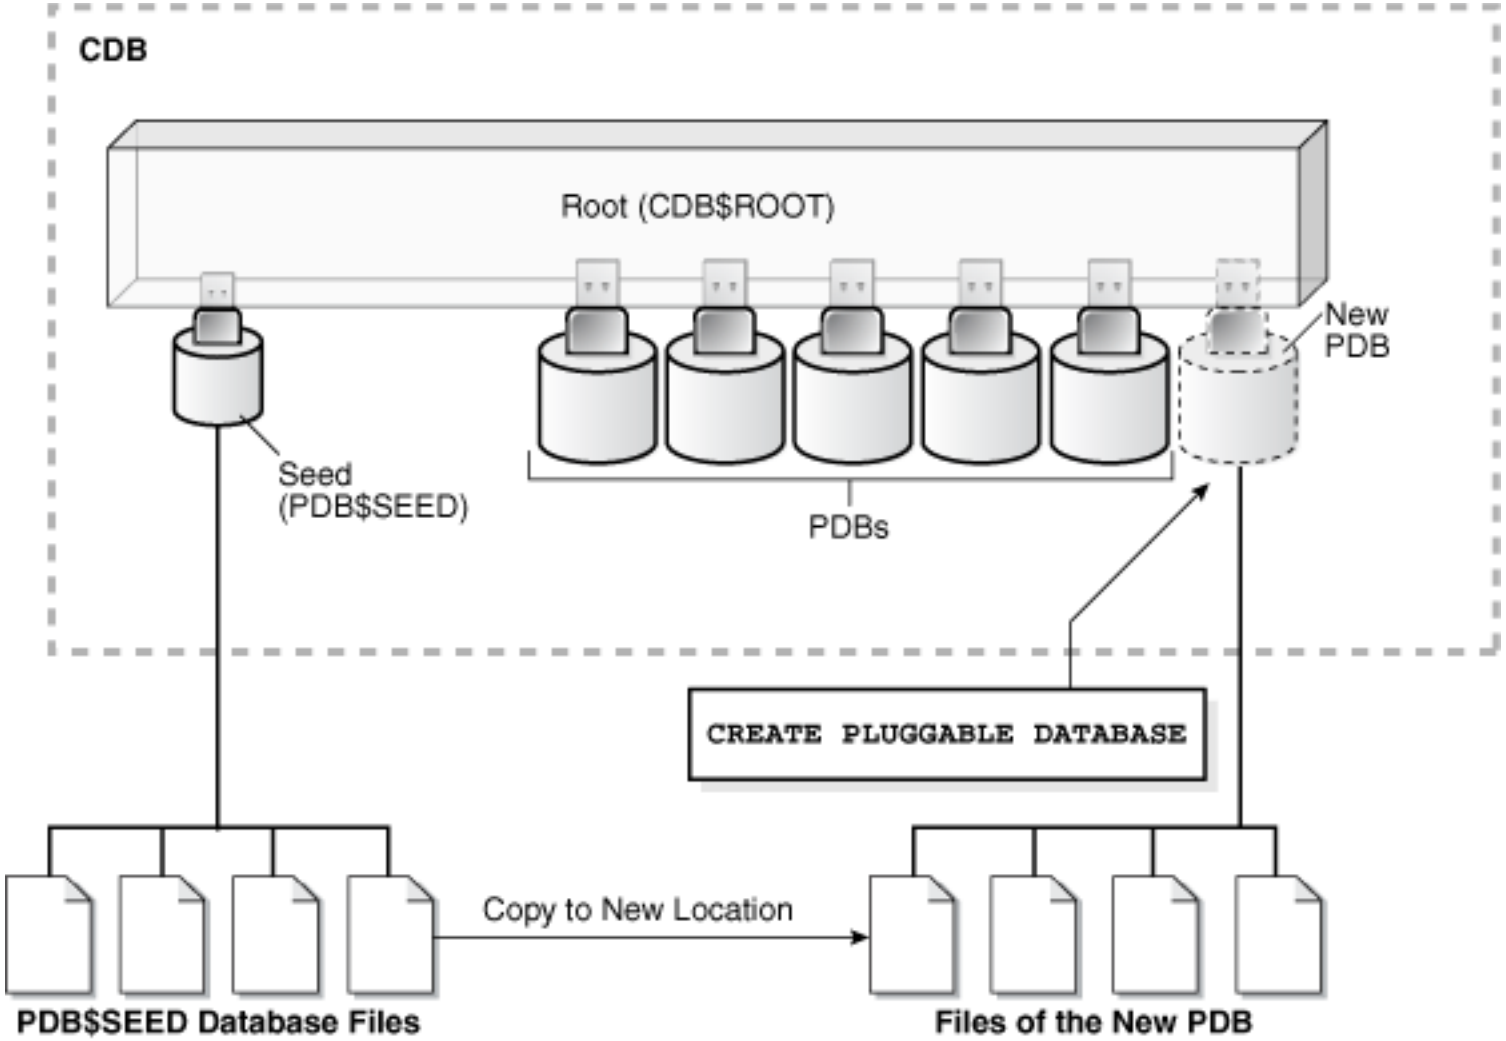

The 23ai brings a move to container architecture

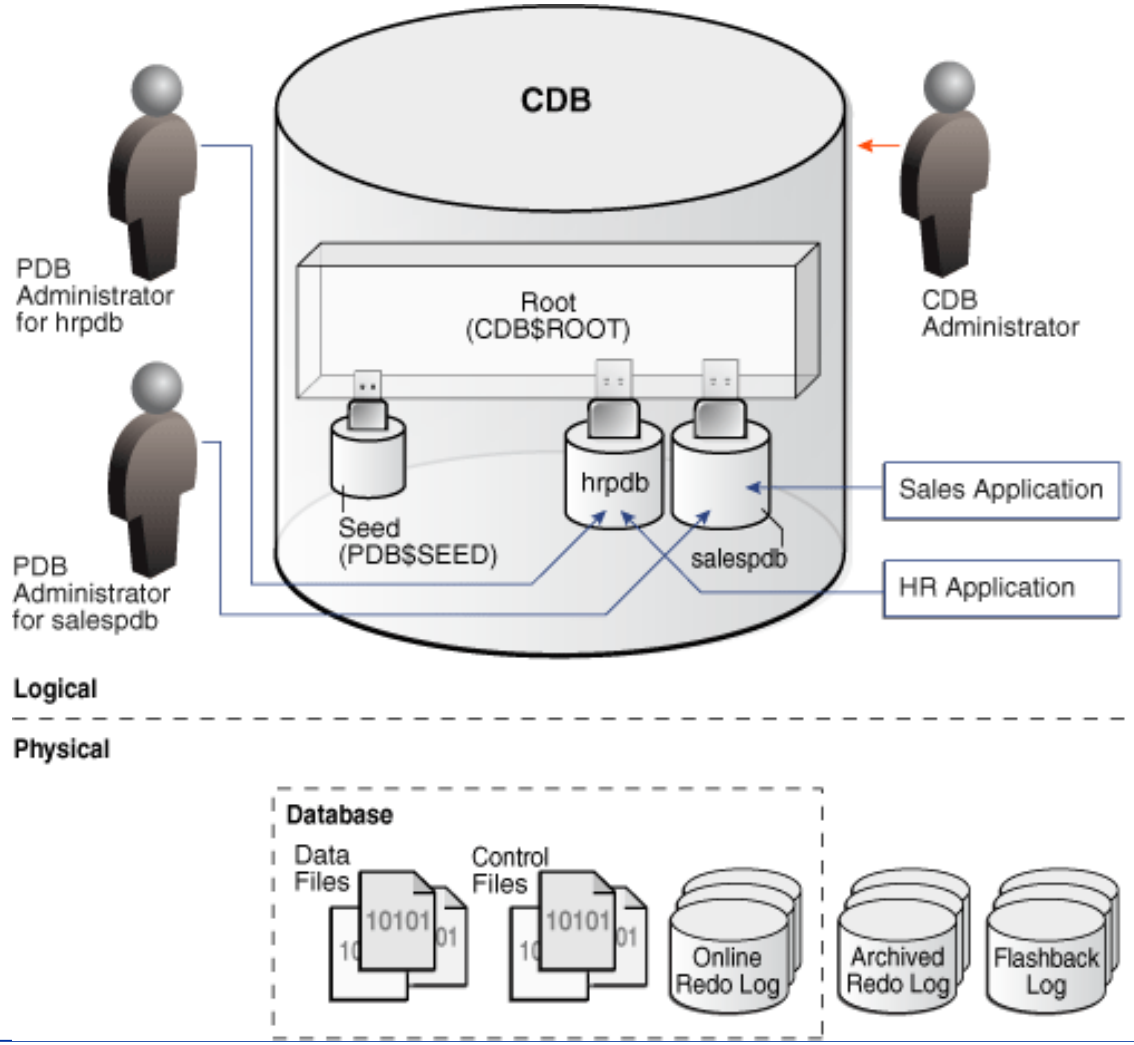

Oracle 23ai
Database New Features

https://docs.oracle.com/en/database/oracle/oracle-database/23/nfcoa/

I'll not cover

SELECT AI

```
SQL> SELECT AI What are total sales of tom hanks movies
70,318.23

SQL> SELECT AI showsql What are total sales of tom hanks movies
SELECT SUM(sales_sample.list_price) AS total_sales
FROM moviestream.sales_sample
JOIN moviestream.movie ON sales_sample.movie_id = movie.movie_id
WHERE movie.cast LIKE '%Tom Hanks%'
```

Livelabs: Chat with your data in Autonomous Database using generative AI
https://apexapps.oracle.com/pls/apex/r/dbpm/livelabs/run-workshop?p210_wid=3831

Oracle True Cache

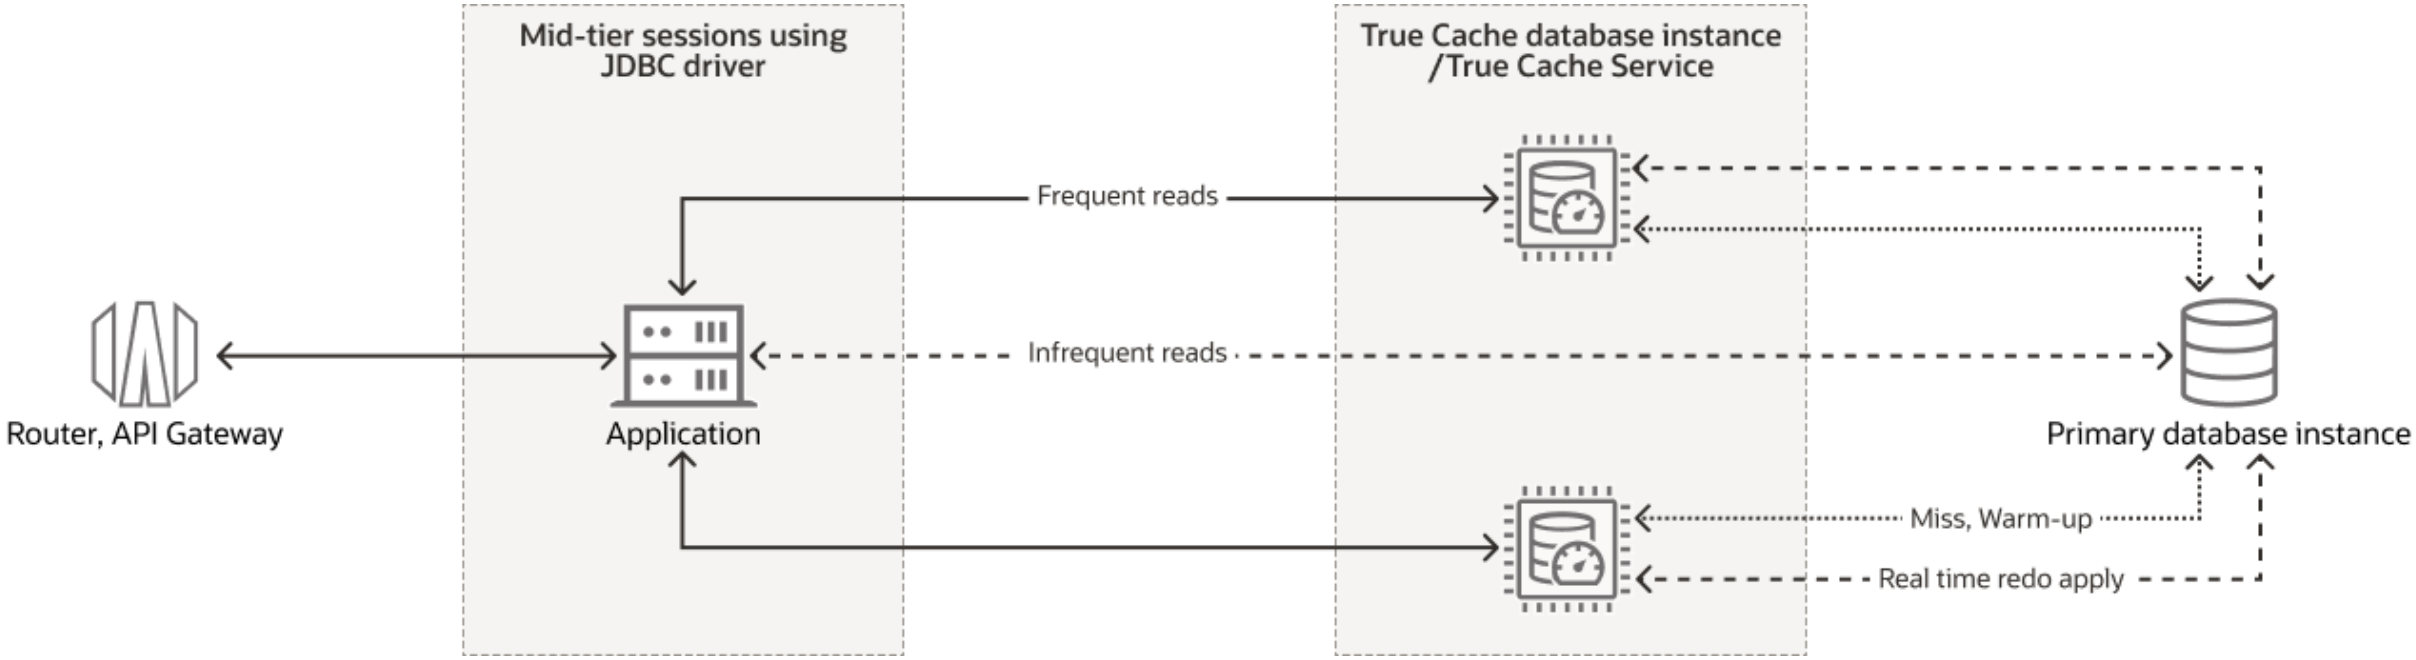

https://www.oracle.com/database/truecache/

AI&ML features – VECTOR

```
SQL> SELECT VECTOR_EMBEDDING(
                                ALL_MINILM_L12_V2
                                USING 'The quick brown fox jumped' as DATA
                              ) AS embedding;
EMBEDDING
--------------------------------------------------------------------------------
[1.65517051E-002,3.19098569E-002,-1.96293015E-002,-3.56926955E-002,9.21710581E-0

SQL>   update vector.my_data set vector_l12_v2 =
                           VECTOR_EMBEDDING(ALL_MINILM_L12_V2 USING info as data);
SQL>   select info
       from vector.my_data
       order by vector_distance(
                     VECTOR_L12_V2,
                     VECTOR_EMBEDDING(ALL_MINILM_L12_V2 USING
                                      '&prompt' as data),
                     cosine)
       fetch approximate first 5 rows only;
```

JSON-Relational duality

```
SQL> create or replace json relational duality view department_dv as
       select json {'_id' : d.deptno,
                    'departmentName'   : d.dname,
                    'location'         : d.loc,
                    'employees' :
                      [ select json {'employeeNumber' : e.empno,
                                     'employeeName'   : e.ename,
                                     'job'            : e.job,
                                     'salary'         : e.sal}
                          from   emp e with insert update delete
                          where  d.deptno = e.deptno ]}
from dept d with insert update delete;
```

https://oracle-base.com/articles/23/json-relational-duality-views-23

JSON-Relational duality

```
SQL> select json_serialize(d.data pretty)
     from   department_dv d
     where  d.data."_id" = 40;

JSON_SERIALIZE(D.DATAPRETTY)
--------------------------------------------------------------------------------
{
  "_id" : 40,
  "_metadata" :
  {
    "etag" : "6FAB9798FF405D87F0EB44456398A5D5",
    "asof" : "00000000002F2799"
  },
  "departmentName" : "OPERATIONS",
  "location" : "BOSTON",
  "employees" :
  [
  ]
}
```

Optimistic locking based on the etag

https://oracle-base.com/articles/23/json-relational-duality-views-23

JSON-Relational duality

```
SQL> insert into department_dv d
       (data)
values ('
{
  "_id" : 50,
  "departmentName" : "DBA",
  "location" : "BIRMINGHAM",
  "employees" : [
    {
      "employeeNumber" : 9999,
      "employeeName" : "HALL",
      "job" : "CLERK",
      "salary" : 500
    }
  ]
}');
```

```
SQL> select * from dept where deptno = 50;

    DEPTNO DNAME          LOC
---------- -------------- -------------
        50 DBA            BIRMINGHAM


SQL> select empno, ename, job
     from   emp
     where  deptno = 50;

     EMPNO ENAME      JOB
---------- ---------- ---------
      9999 HALL       CLERK
```

https://oracle-base.com/articles/23/json-relational-duality-views-23

Priority of Transactions

```
SQL> ALTER SESSION SET txn_priority = 'HIGH/MEDIUM/LOW';
SQL> ALTER SESSION SET priority_txns_high_wait_target = 30;
SQL> ALTER SESSION SET priority_txns_medium_wait_target = 90;
SQL> ALTER SESSION SET priority_txns_mode = 'TRACK/ROLLBACK';
```

If a HIGH priority transaction is blocked for a row lock, Oracle Database can roll back the transaction that is holding the row lock only if the holder is LOW or MEDIUM priority. Oracle Database never rolls back a HIGH priority transaction.

If a MEDIUM priority transaction is blocked for a row lock, Oracle Database can roll back the transaction that is holding the row lock only if the holder is LOW priority.

If a LOW priority transaction is blocked for a row lock, Oracle Database will not attempt to roll back the transaction holding the row lock irrespective of its priority.

https://blogs.oracle.com/dbstorage/post/new-priority-transactions-capability-with-oracle-database-23ai

IF EXISTS

```
SQL> drop table if exists product_sales;

Table dropped.

SQL> drop table if exists product_sales;

Table dropped.

SQL> create table if not exists product_sales
                                        (product varchar2(10), sale_date
date, amount number);

Table created.

SQL> create table if not exists product_sales
                                        (product varchar2(10), sale_date
date, amount number);

Table created.
```

BOOL

TRUE, FALSE, NULL

```
SQL> create table employees ( id number generated always as identity,
                        name varchar2(20),
                        is_active bool,
                        is_retired boolean);

Table created.

SQL> insert into employees (name, is_active, is_retired) values
                              ('Obama',false,true),
                              ('Trump',false,false),
                              ('Biden',true,false);
3 rows created.
```

Lots of JSON improvements

JSON_ARRAY accepts subquery as input:

```
json_array(select json_object('employee-number' : e.empno) from emp e)
```

`JSON_BEHAVIOR` **parameter controls behaviour of error handling**
`(ERROR ON ERROR, NULL ON ERROR, FALSE ON ERROR, etc.)`

Many improvements to JSON data type

JSON Schema to check validity and structure of JSON documents – `VALIDATE` **keyword**

`DBMS_JSON.JSON_TYPE_CONVERTIBLE_CHECK` **to help migration CLOB -> JSON**

When is it all coming? Not soon

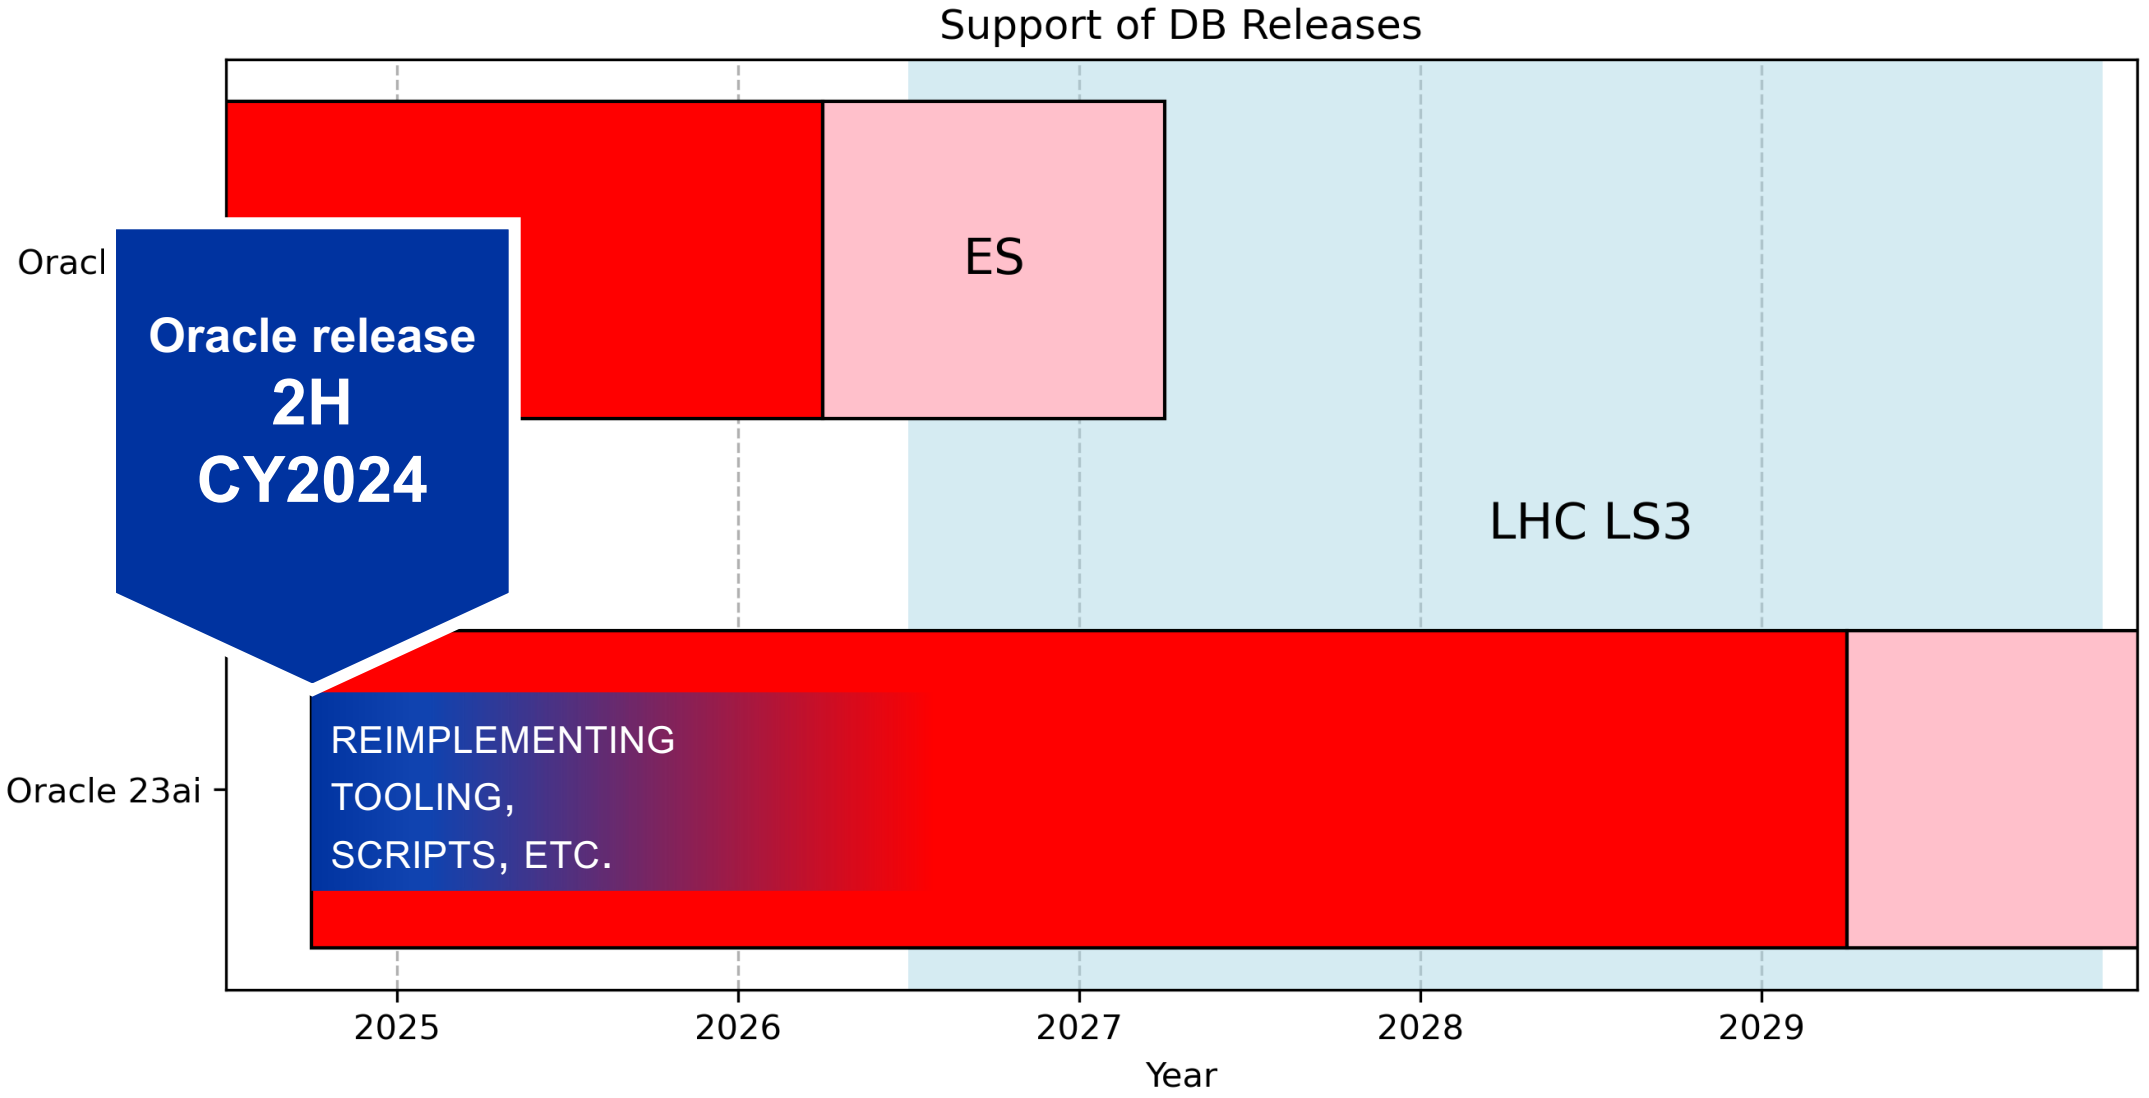

Thank you !

andrzej.nowicki@cern.ch