

23ai New Features

(Originally 23c & A few key 19c & 21c Features)



Rich Niemiec

Oracle ACE Director | CIO

[@richniemiec](https://twitter.com/richniemiec)

[linkedin.com/in/richard-niemiec-b65828/](https://www.linkedin.com/in/richard-niemiec-b65828/)

rich.niemiec@viscosityna.com



Rich Niemiec @richniemiec



ORACLE
ACE Director



21c

19c ORACLE Database

Rich's Overview...

@richniemiec

rich@viscosityna.com



- ▶ Chief Innovation Officer, Viscosity North America
- ▶ Board Member – TEC, Entrigna, Ask DB Experts
- ▶ Former CEO of TUSC
 - Inc. 500 Company (Fastest Growing 500 Private Companies)
 - 10 Offices in the United States (U.S.); Based in Chicago
 - Oracle Advantage Partner in Tech & Applications
- ▶ Former President Rolta TUSC & President Rolta EICT International & Executive Advisor to Board
- ▶ Author (4 Oracle Best Sellers – #1 Oracle Tuning Book for two Decades – **12c R2 Tuning in March 2017**):
 - Oracle Performing Tips & Techniques (Covers Oracle7 & 8i)
 - Oracle9i & 10g Performance Tips & Technique
 - Oracle Database 11g Performance Tips & Techniques
 - Quick Start Guide to Oracle Query Tuning (2015)
- ▶ Former President of the International Oracle Users Group
- ▶ Current President of the Midwest Oracle Users Group
- ▶ Chicago Entrepreneur Hall of Fame - 1998
- ▶ E&Y Entrepreneur of the Year & National Hall of Fame - 2001
- ▶ IOUG Top Speaker in 1991, 1994, 1997, 2001, 2006, 2007
- ▶ MOUG Top Speaker Twelve Times
- ▶ National Trio Achiever award - 2006
- ▶ Oracle Certified Master & Oracle Ace Director
- ▶ Purdue Outstanding Electrical & Computer and Engineer - 2007



We Enable Business Transformation at a Time When Companies Must Change to Survive!

Viscosity Pillars and Delivery Models



DATA

Oracle, SQL Server, Postgres
Performance Tuning
Data Replication
Data Warehousing Analytics
Data Integration
ERP Blueprints
Database Upgrades



APPS

Oracle APEX
Web/Mobile Apps
.Net and C#
E-Business Suite
SAAS/PAAS
Custom AI Products



CLOUD

Azure Gold Partner
Cloud Migrations
Engineered Systems
Oracle Cloud Partner
Google Partner
AWS Partner Hybrid Cloud

Workshops

Assessments

Proof of
Concepts

Training

Turnkey
Projects

Managed
Services

Viscosity's Oracle ACEs

The Oracle ACE Program

The Oracle ACE Program recognizes and rewards individuals for their contributions to the Oracle community.



Charles Kim
CEO | Co-Founder

@racdba

ACE Director



Rich Niemiec
Chief Innovation Officer

@richniemiec

ACE Director



Craig Shallahamer
Applied AI Scientist

@orapub

ACE Director



Sean Scott
Principal Consultant

@oraclesean

ACE Director



Gary Gordhamer
Principal Consultant

@ggordham

ACE Pro



Julio Ayapan
Senior DBA

@jayapangt

ACE Alumni

Quick **FREE** notes

Send email to (for slides):

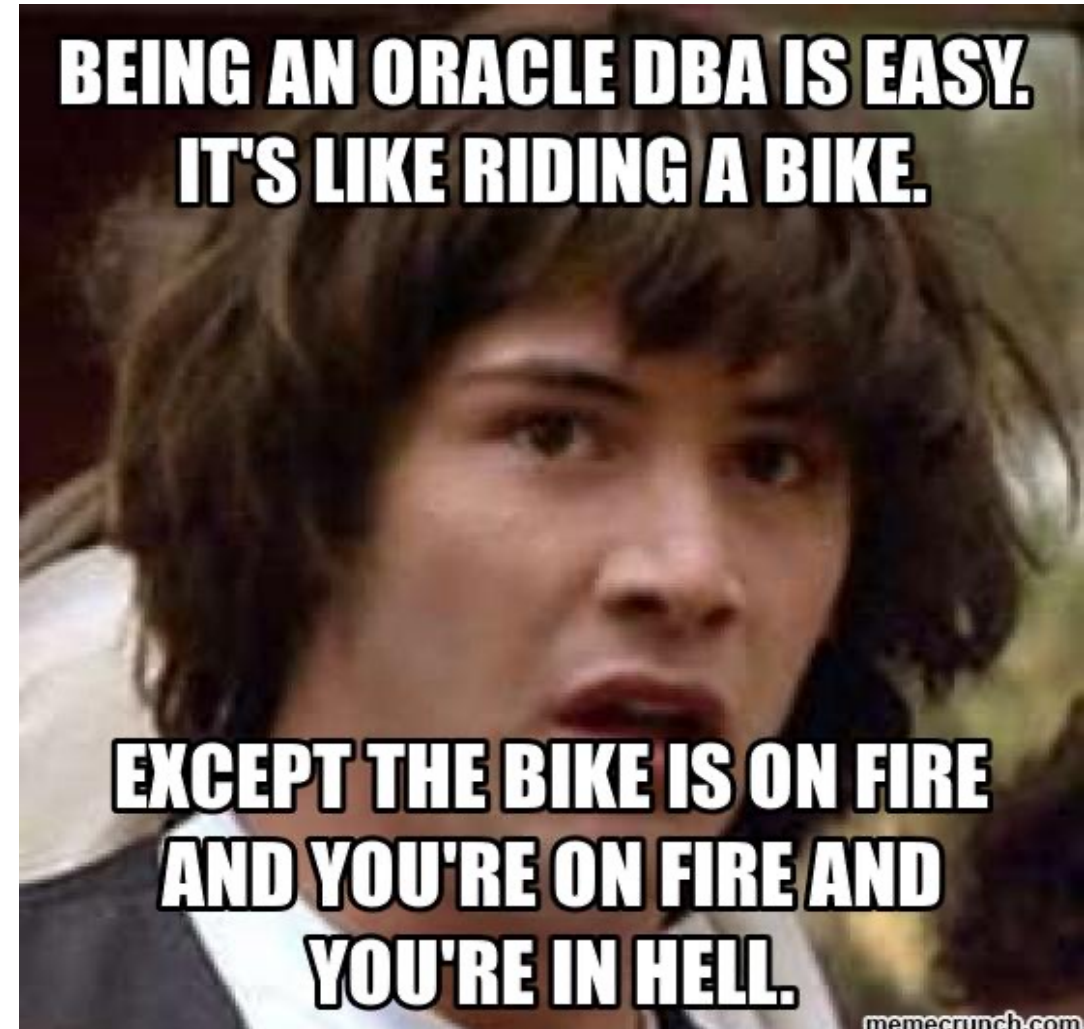
hello@viscosityna.com

(richniemiec@gmail.com)

@richniemiec twitter

98%

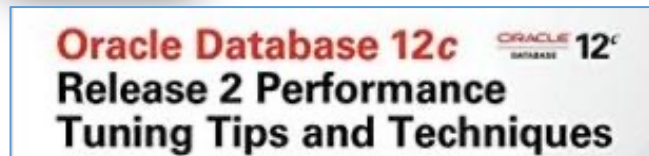
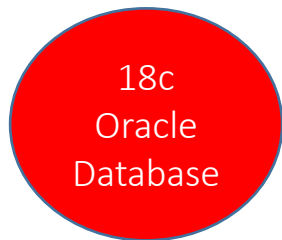
of the FORTUNE 100 Run Oracle Database



Agenda – 23ai New Features & Features in 12c, 18c/19c, 21c



- ▶ Know the Oracle!
- ▶ Autonomous Database & Future of the DBA
- ▶ **Key 12cR1** Features to Use: Multiple indexes on the same, Pluggable Databases, Fetch First x Rows, & In-Memory (12.1.0.2+)
- ▶ **Key 12cR2** Features to Use: Approximate Query, Adaptive Query Optimization and Security Enhancements
- ▶ **Key 18c** Features to Use: Snapshot Carousel, PDB Switchover, & INMEMORY External Tables
- ▶ **Key 19c** Features to Use: This is version to Land on, PDBs, ADW/ATP, Automatic Indexes, Quarantine, Documentation Apps including Features by Version and Images for Learning
- ▶ **Key 21c** Features Coming: AutoML, OML4Py, Blockchain Tables, Auto In-Memory Management, Hybrid In-Memory Scans, AutoUpgrade, Active Data Guard with Standby DB Result Cache, Dark Mode for APEX, & New ML Algorithms
- ▶ **Detail on Key 23ai** Features Coming include 300+ Features
- ▶ **Autonomous Database (Now there are 4 of them!)**
- ▶ **Prepare for the Future!**



Know the Oracle!

Oracle Firsts – *Innovation to Acquisitions*

*On-premise (Production 12cR2 first came out in 2016 in the Oracle Cloud)



- 1979 First commercial SQL RDBMS
- 1983 First 32-bit mode RDBMS & First with read consistency
- 1987 First client-server database with multilevel secure database evaluations
- 1995 First 64-bit mode RDBMS (Exabytes in Memory - if you have it)**
- 1996 First to break the 30,000 TPC-C barrier
- 1997 First Web database
- 1998 First Database - Native Java Support, Linux, Breaks 100,000 TPC-C
- 2000 First database with XML, RAC & First middle-tier database cache
- 2004 First True Grid DB & 2005 FREE Oracle Database (10g Express Edition)
- 2006 First Oracle Support for LINUX Offering
- 2007 Oracle 11g Released!
- 2008 Exadata V1 Server Announced (Oracle buys BEA)
- 2009 Oracle buys Sun – Java; MySQL; Solaris; Hardware; OpenOffice, StorageTek
- 2010 Oracle announces MySQL Cluster 7.1, Exadata, Exalogic, America's Cup Win
- 2011 X2-2 Exadata, ODA, Exalytics, SuperCluster, Big Data, Cloud, Social Network
- 2012 X3-2 Exadata, Expanded Cloud Offerings, Solaris 11.1
- 2013 Oracle12c Released! Oracle X3-8 Exadata, Acquisitions (Acme Packet...etc.)!
- 2014 Oracle X-4, Acquisitions: Responsys & Corente, IN-MEMORY DB
- 2015 X5-2, X5-8, FS1 Flash Array, Acquisitions & Cloud Solutions
- 2016 X6-2 (all flash if you want), X6-8, M7 SuperCluster, Cloud Solutions & Acquisitions
- 2017: Production 12cR2*, X7, Autonomous Database Warehouse Cloud (ADWC) & ML Security announced
- 2018: Production ADWC, Oracle 18c, & Production Autonomous Database for OLTP**
- 2019: Oracle 19c / ADW / ATP/X8 2021: Oracle 21c 2022: Oracle 22c 2023: Oracle 23c 2024: Oracle 23ai**

1

Larry Ellison
CEO, Oracle

Ellison's plans to roll up the enterprise applications space show no signs of slowing. Oracle has leveraged its strength in the data center to cement its status as one of world's most important applications and middleware vendors.

TOP 100 MOST INFLUENTIAL PEOPLE IN IT

Oracle Stock Hits All-Time High Amid Strong Profit Growth



Oracle Stock Hits All-Time High Amid Strong Profit Growth
INVESTOR'S BUSINESS DAILY AND JAMES DETAR
Mon, Jun 10, 2024, 1:50 PM CDT - 3 min read

+16% Earnings 6/10/2024



Characteristics of Big Data to use with AI/ML

Volume

Big data comes in one size: large. Enterprises are awash with data, easily amassing terabytes and even petabytes of information.

TB, Records, Transactions, Tables, Files

Velocity

Often time-sensitive, big data as it is streaming in to the enterprise to maximize its value to the business. Batch, Near time, Real time

Value

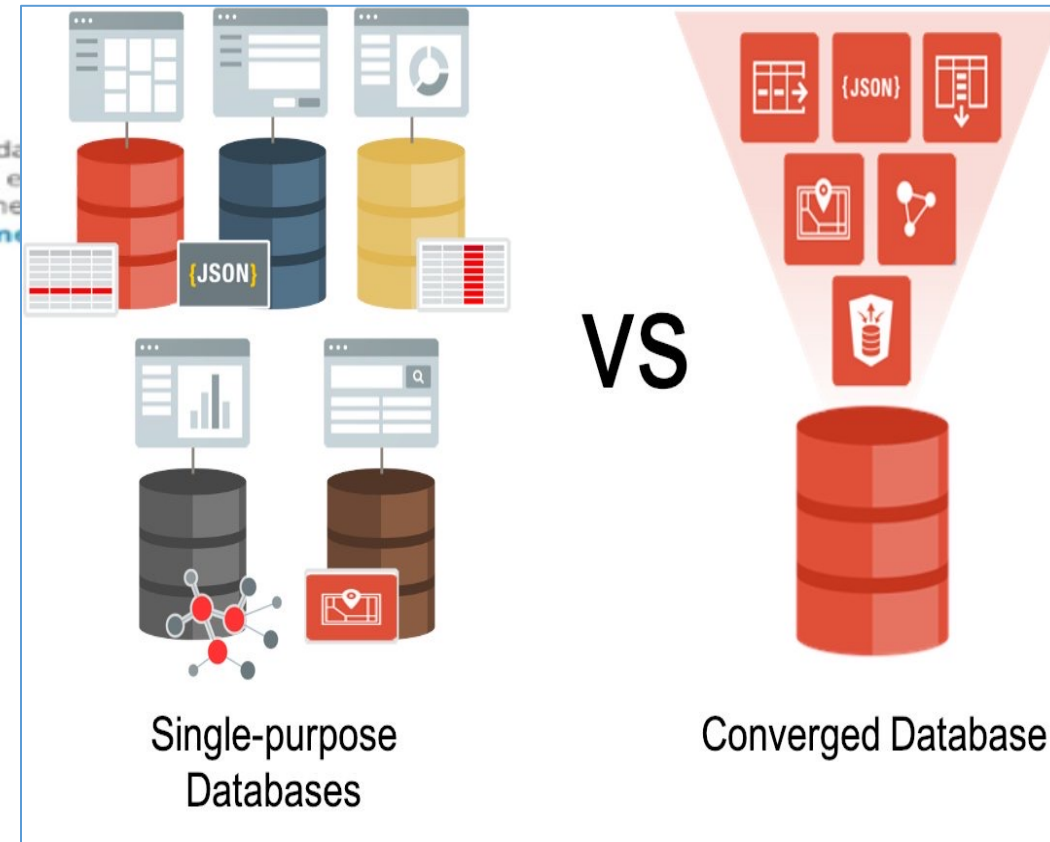
Business value of Big Data

Variety

Big data extends beyond structured data, including semi-structured and unstructured data of all varieties: text, audio, video, click streams, log files and more.

Structured, Unstructured, Semistructured

Veracity



Big Data Themes

- HW & SW technologies for large data volumes
- Focus on Web 2.0 technologies
- Database Scale-out
- Relational & Distributed Data Analytics
- Real Time Analytics

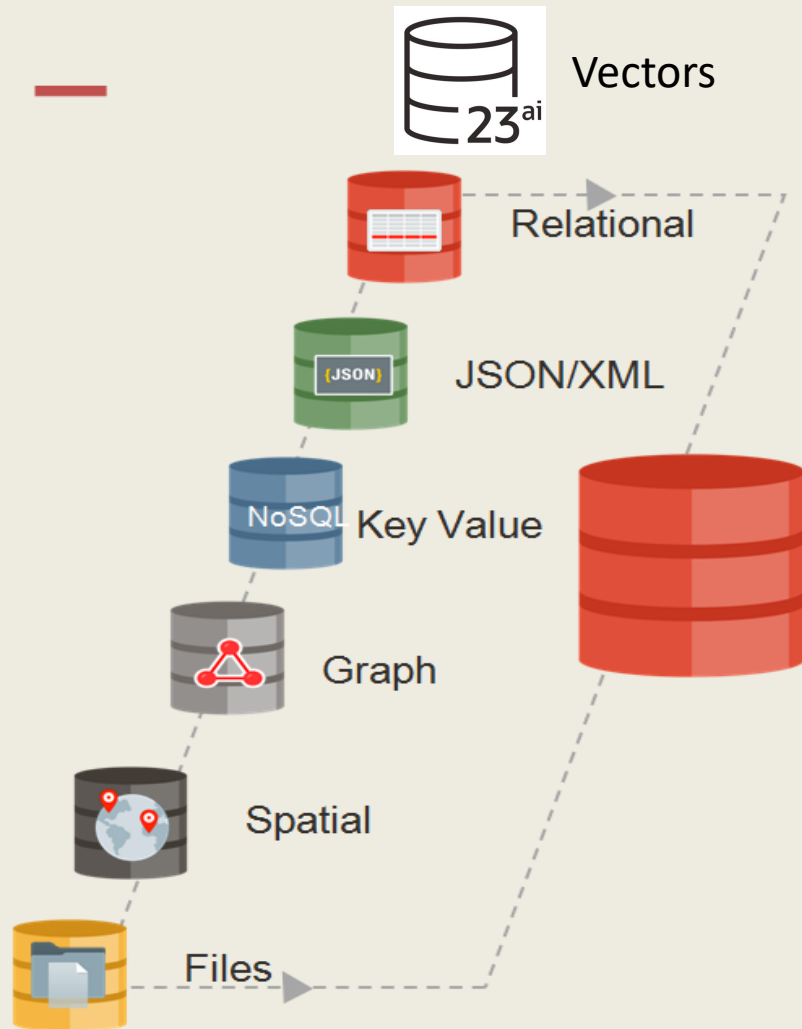
Big Data Domains

- Digital Marketing Optimization
- Data Exploration & Discovery
- Fraud Detection & Prevention
- Social Network & Relationship Analysis
- Machine-generated Data Analytics

Converged Database - Oracle Multi-Model Database*

*Slide from Oracle's,
Nitin Vengurlekar

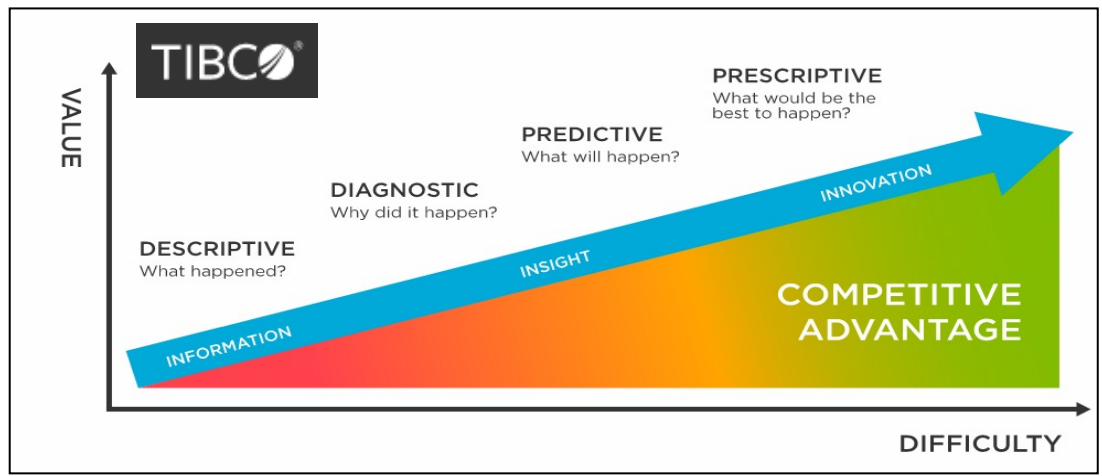
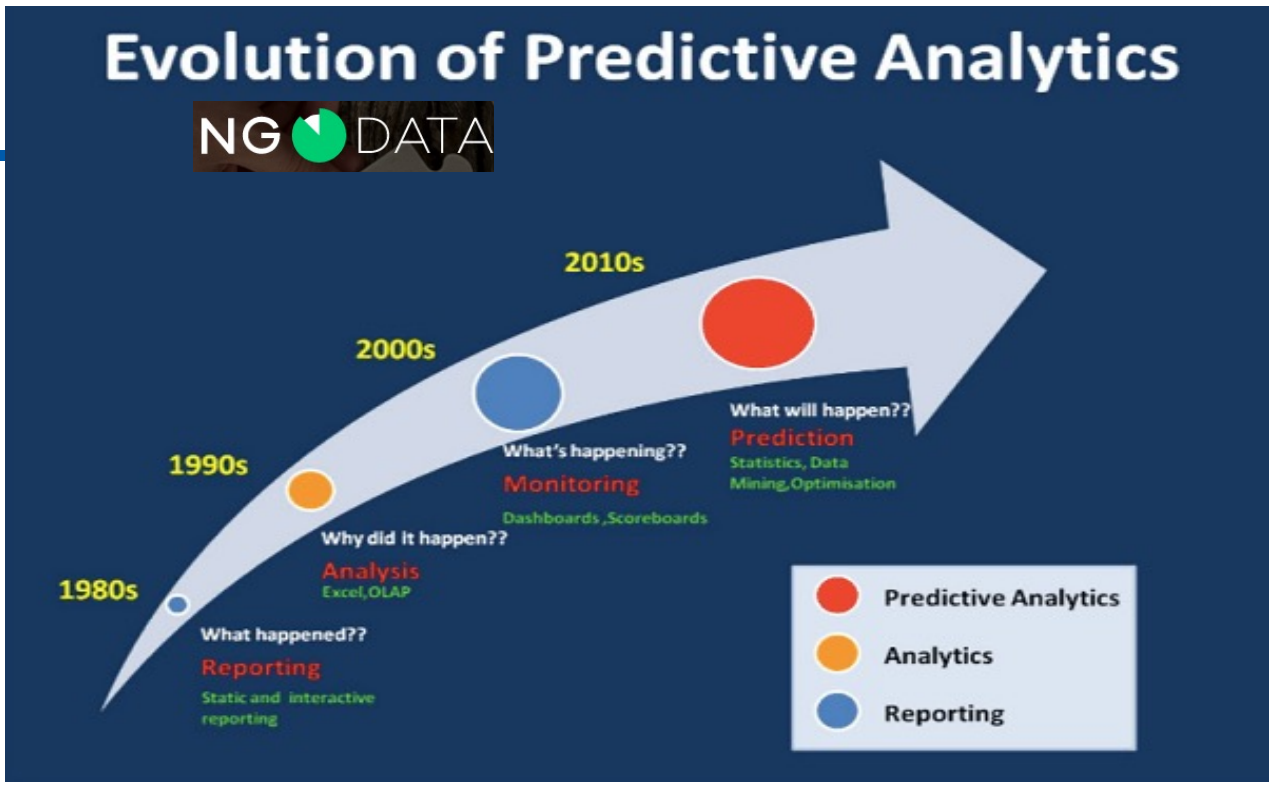
- Benefits of Oracle's Converged Database are broad



- Integrated development tools (Apex, SQL Dev, Spatial Studio)
- 3rd party and Open Source development tools
- Machine Learning
- Node.js, Python, many others
- In-memory database
- Spatial, Graph support
- NoSQL (JSON, key-value, wide column, XML)
- Containers, microservices, virtualization (Docker, MT)
- Integrated Security
- Deployment choice (on-prem, cloud, hybrid)
- Integrated High Availability and Disaster Recovery



What You Need; Nick of Time (*Prescriptive Analytics*)



Oracle Database Direction from Andy Mendelsohn at CloudWorld '23

Summary – Oracle Database Directions

Converged Oracle Database

- Supports all modern data types, workloads, and development styles
- Simplifies development of analytics and machine learning
- Completely consistent, scalable, available, and secure platform

Oracle Autonomous Database

- All the benefits of converged Oracle Database, plus best customer experience
- Best cloud database for running any app at any scale or criticality
- Deploy in the cloud, on-premises and hybrid/multicloud configurations

Oracle Database 23c

- Next long-term release for on-premises and cloud deployments



**Version 23ai
announced
May 2024**

Cloud Choices – Less than you think



Acceleration Economy network

HOME CLOUD WARS ▾ SECTIONS ▾ VENDOR DIRECTORY EVENTS ▾ OUR ANALYSTS MEDIA ASSETS ▾ MORE

Oracle on Verge of Becoming World's Hottest Cloud Vendor

CLOUD WARS

Top 10 Rankings – Aug. 19, 2019 @bobevasIT

1. **Microsoft** - *Follows monster Q4 w/ \$1B deal to "shape trajectory of humanity"*
2. **Amazon** - *AWS Q2 revenue jumps 37% to \$8.38B, cites broad innovation in ML*
3. **Salesforce** - *After big Q1, Benioff pivots to Customer 360; and, getting into HCM?*
4. **SAP** - *Q2 cloud revenue up 40%; McDermott vows China slip just short-term blip*
5. **Oracle** - *Blowout fiscal Q4 suggests Autonomous Database will be huge success*
6. **Google** - *Q2 cloud rev. tops \$2B; Kurian's customer focus leads jump from #7 to #6*
7. **IBM** - *Q2 cloud growth rate slips to 5% (8% in cc)—huge bet on Red Hat impact*
8. **Workday** - *Total rev. up 33% to \$825M; first \$1-billion qtr likely within 12 mos.*
9. **Accenture** - *Up from #10 on ties w/ MSFT AWS GOOG; \$9B cloud biz up 23% in '18*
10. **ServiceNow** - *Donahoe looks to Microsoft & mobile to drive next growth wave*

Acceleration Economy network

HOME CLOUD WARS ▾ SECTIONS ▾ VENDOR DIRECTORY EVENTS ▾ OUR ANALYSTS MEDIA ASSETS ▾ MORE

Oracle on Verge of Becoming World's Hottest Cloud Vendor

By Bob Evans – September 6, 2022 – Updated: September 6, 2022 4 Mins Read

Facebook Twitter LinkedIn

Recent Posts

- How C3 AI Pricing Revamp Affects Cloud Usage Model
- Oracle on Verge of Becoming World's Hottest Cloud Vendor
- How Alóki Uses Gamification to Drive Reforestation & Sustainability
- Why Don't CISOs Sit on the Fence?
- Coming Soon: Is Oracle Today's Hottest Cloud Vendor?

Oracle Gen2 Cloud Infrastructure

140%

Q4 Growth Rate

Annualized Consumption Revenue - ACR

New Versions: Who is Truly Committed?

Advanced DBA is *Interested*



Intermediate DBA is *Invested*




Beginner DBA is *Committed*





FYI

Ensure you use Key Oracle 12c R1 Features

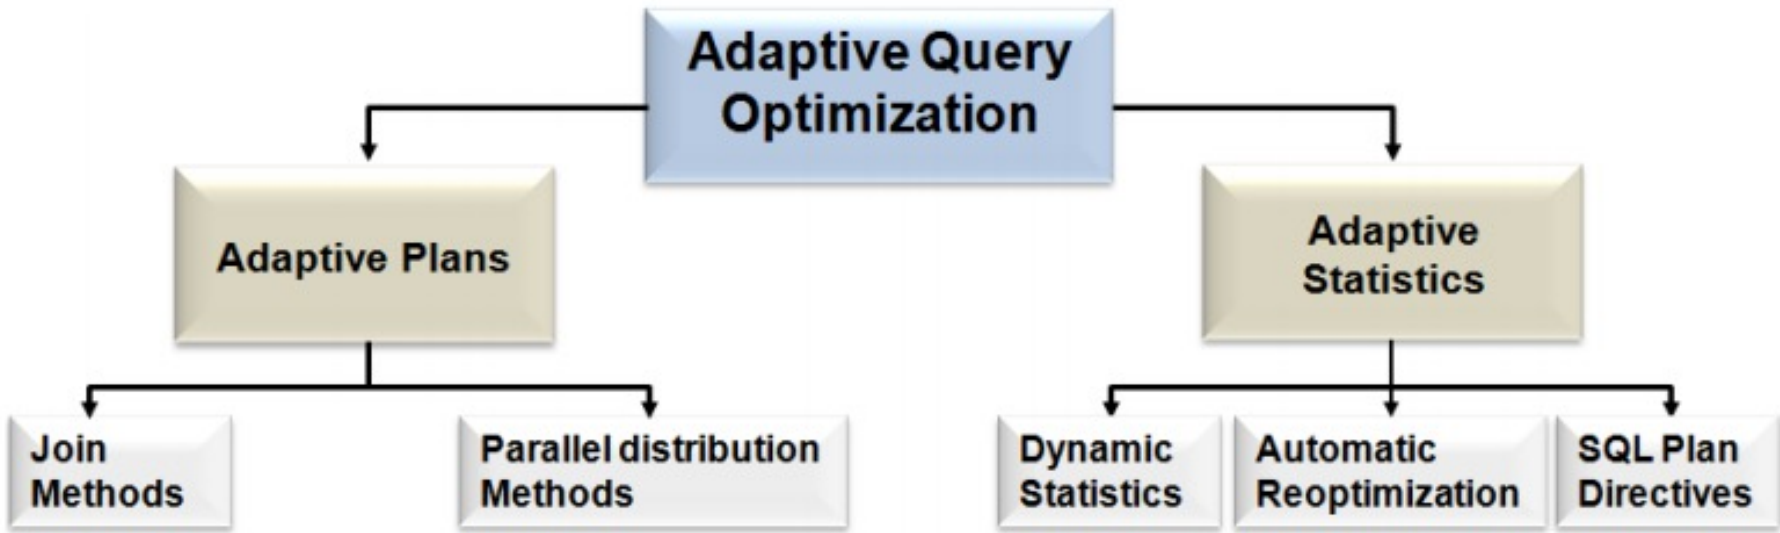
- 
- ▶ **Invisible Indexes** used to Create **Multiple indexes on same Column**
 - ▶ **Approximate Query to Count Distinct** values faster
 - ▶ **Pluggable Databases** - Understand them well as they are the future!
 - ▶ **Move partitions** while **ONLINE** with DML happening
 - ▶ Nice Tuning: ***OFFSET x ROWS FETCH FIRST x ROWS ONLY***
 - ▶ **Data Pump No Logging Option** for import
 - ▶ **VARCHAR2(32767)** – not default (except on Cloud – so far)
 - ▶ **Improved query performance** against **OLAP cubes** (especially Exadata)
 - ▶ **Mask Data At Source** for testing & Oracle Masking templates for E-Business
 - ▶ **Oracle Data Redaction** (prevents things like SSN from being displayed)
 - ▶ **TRUNCATE TABLE ...CASCADE** (truncate child tables too);

Adaptive Query Optimization

(Thanks Oracle docs.)

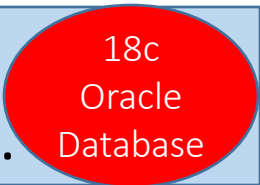


Oracle Database 12c
Release 2 Performance
Tuning Tips and Techniques



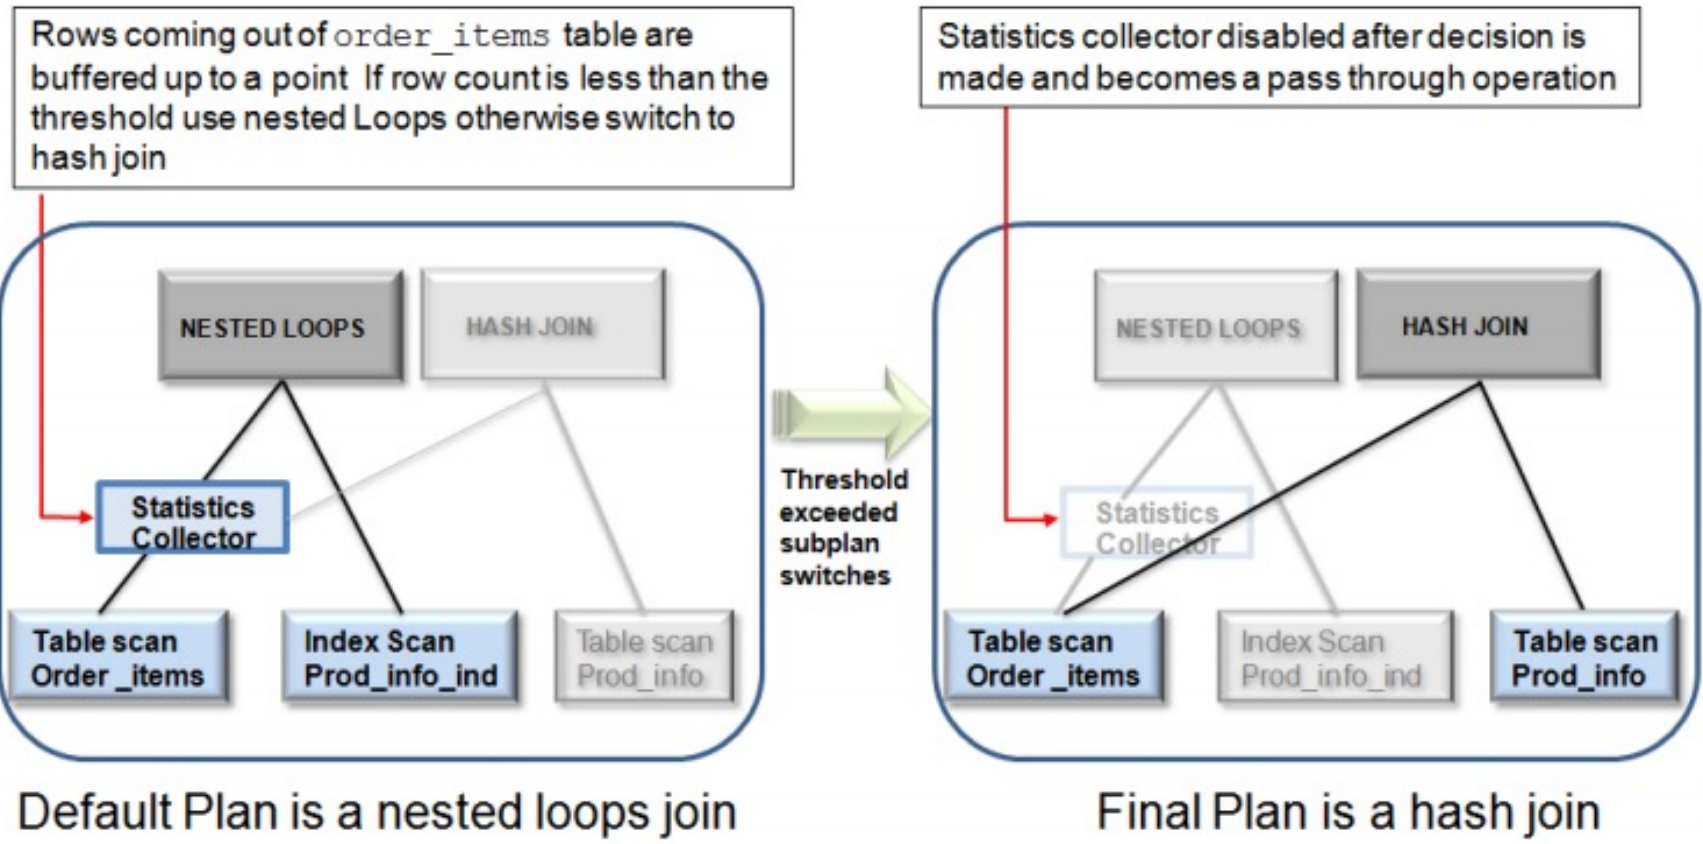
In 12cR2, Oracle introduced **Continuous Adaptive Query Plans (CAQP)** where certain queries, based on input data, can benefit from **continuous adaptive join methods** (such as a Recursive WITH that has different input data compared to previous iteration).

In 18c: Adaptive Query Plans: **Bitmap Index Pruning (uses some/not others)**
Adaptive plans prune indexes that do not significantly reduce number of matched rows.





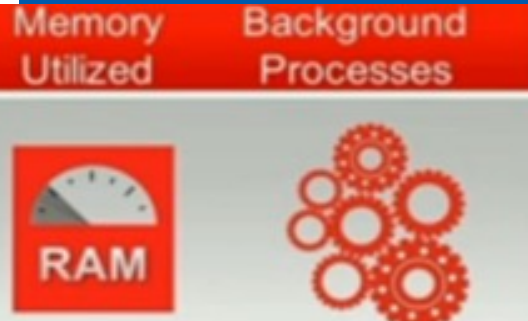
Adaptive Query Optimization: Oracle Docs Great Example



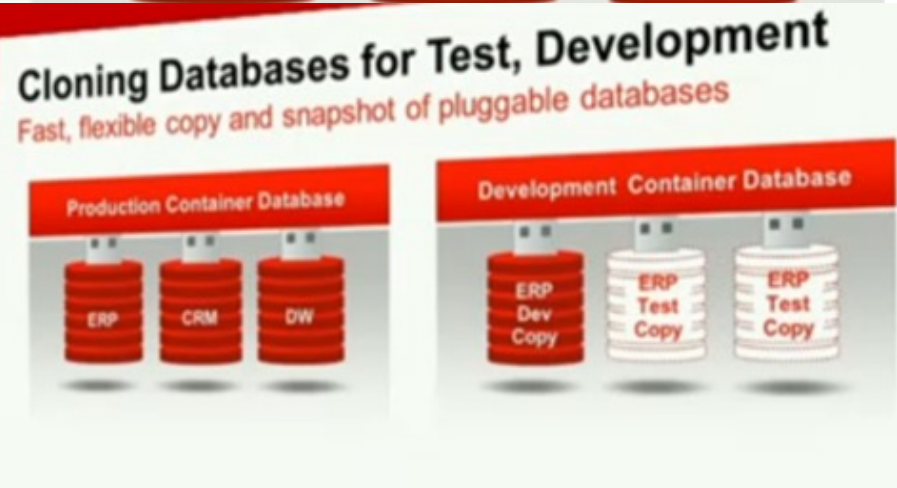
Pluggable Databases – Briefly

FYI

12c



Oracle Database 12c
Release 2 Performance
Tuning Tips and Techniques



3 PDBs included +
Desupport of non-CDB architecture in Oracle 21c

8.1.1 Deprecation of Non-CDB Architecture

Lone-PDB is a Choice!

The non-CDB architecture is deprecated in Oracle Database 12c, and may be desupported and unavailable in a release after Oracle Database 12c Release 2. Oracle recommends use of the CDB architecture.

Pluggable Databases

- ▶ CDB = Container Database (has Root DB & also has a seed PDB)
- ▶ PDB = Pluggable Database (plugged into a CDB)
- ▶ Non-CDB = Original type of Database (neither a CDB or PDB) - **Also Lone-PDB**
- ▶ Why?: Can't consolidate 100's of databases on one machine ... too many resources required when you add the SGAs up! Enter PDBs.
 - Share: Big Data Sources, Acquisitions, Partners, Shared Research, Governments
- ▶ Quickly create a new database (PDB) or copy existing one (PDB)
- ▶ Move existing PDBs to new platform or location or clone it (snapshot)
- ▶ Patch/Upgrade PDB by plugging it into a CDB at a later version
- ▶ Physical machine runs more PDBs old way: Easier to manage/tune
- ▶ Backup entire CDB + any number of PDBs
- ▶ New syntax for commands: PLUGGABLE DATABASE



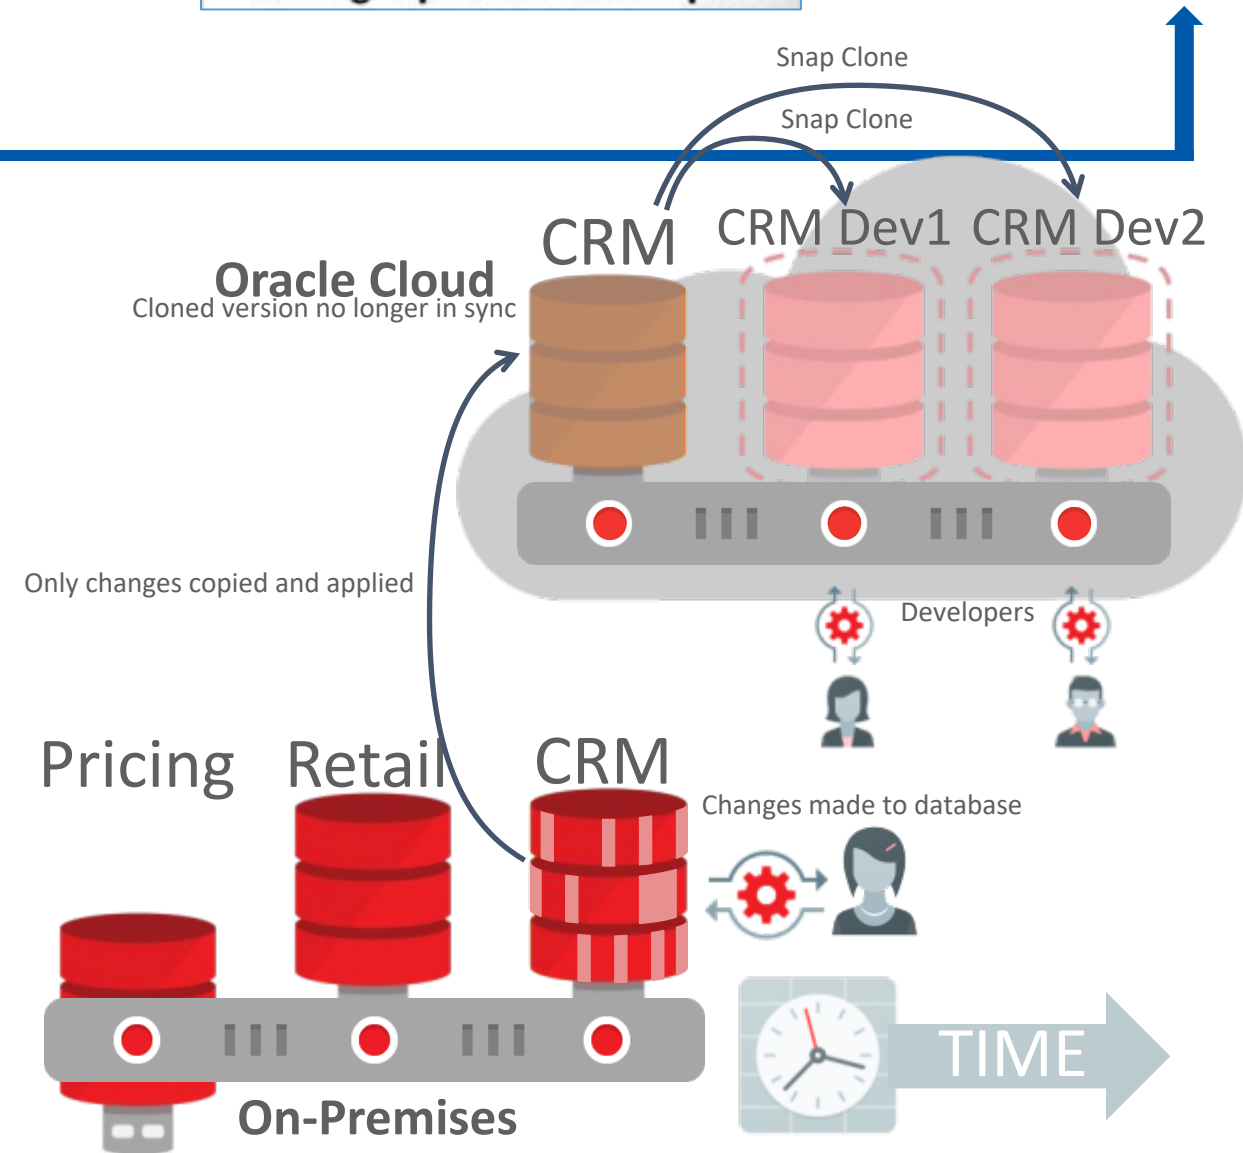
Query the PDBs

```
select name, open_mode, open_time  
from v$pdb;
```

NAME	OPEN_MODE	OPEN_TIME
PDB\$SEED	READ ONLY	23-FEB-13 05.29.19.861 AM
PDB1	READ WRITE	23-FEB-13 05.29.25.846 AM
PDB_SS	READ WRITE	23-FEB-13 05.29.37.587 AM

PDB Refresh

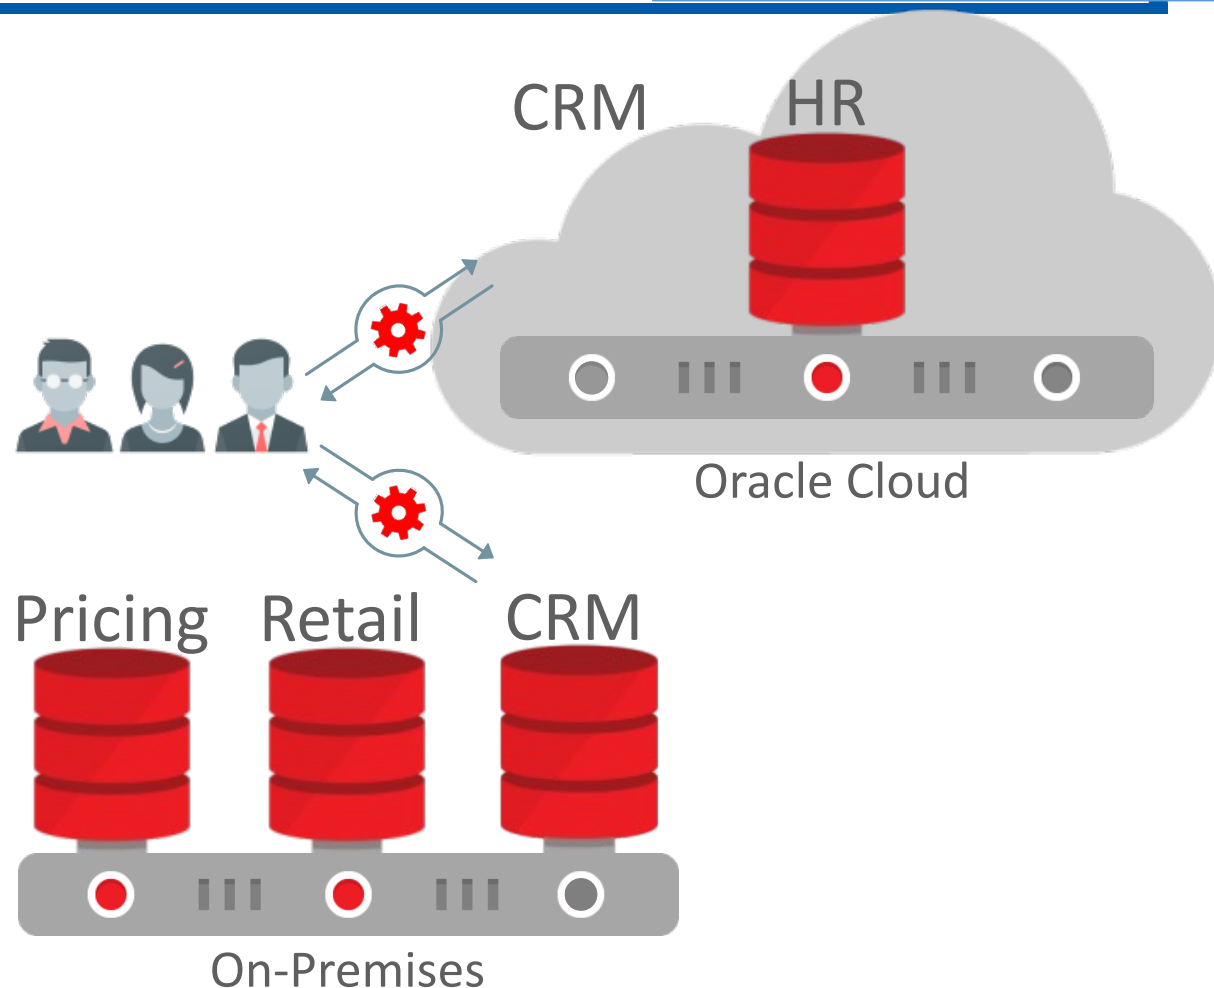
- ▶ PDB Hot Clone
 - Online test master instantiation
- ▶ PDB Refresh
 - Incremental refresh of clone with latest data



```
CREATE PLUGGABLE DATABASE pdb2 FROM pdb1@pdb1_link REFRESH MODE EVERY 60 MINUTES;
```


PDB Relocate

- ▶ PDB Hot Clone
 - Online test master instantiation
- ▶ PDB Refresh
 - Incremental refresh of clone with latest data
- ▶ PDB Relocate
 - Relocate with no downtime



Flashback PDB – 12cR2

- ▶ In **12cR2**, you can also issue a **FLASHBACK** of a **PDB** and have **restore points to only that PDB**.
- ▶ This can be done using the SCN, Restore Point, Clean Restore Point, or Guarantee Restore Point.
- ▶ The FLASHBACK command for an individual PDB is shown below (you can get SCNs from V\$ARCHIVED_LOG, V\$DATABASE, V\$FLASHBACK_DATABASE_LOG & V\$LOG):

```
SQL> FLASHBACK PLUGGABLE DATABASE pdb1 TO SCN 830124;
```

The In-Memory Column Store (IM)

(Examples are *FYI Only*);

Oracle Database 12c ORACLE DATABASE 12c
Release 2 Performance
Tuning Tips and Techniques

Oracle Database In-Memory Base Level Feature



Andy Rivenes
PRODUCT MANAGER

With Oracle Database **21^c** ~~Preview~~
Up to 16G at no extra cost!

Oracle Database In-Memory will now have a "Base Level" feature. This will allow the use of Database In-Memory with up to a 16GB column store for no extra cost. That's not a typo. Database In-Memory will be free to use with up to a 16GB column store!



```
Connected to:  
Oracle Database 12c Enterprise Edition Release 12.1.0.2.0  
With the Partitioning, OLAP, Advanced Analytics and Real /  
ions
```

```
SQL> sho sga
```

```
Total System Global Area 4194304000 bytes  
Fixed Size                 2932336 bytes  
Variable Size             570425744 bytes  
Database Buffers         2013265920 bytes  
Redo Buffers              13844480 bytes  
In-Memory Area           1593835520 bytes  
SQL>
```

In-Memory (IM) NOT enabled if:
INMEMORY_SIZE is set to zero!

In-Memory (IM) – Overview

12c
12.1.0.2

Row



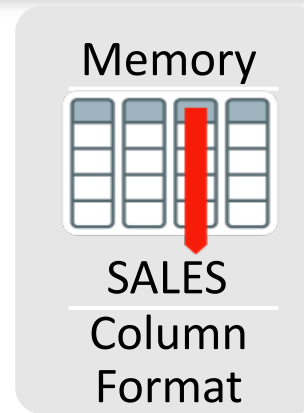
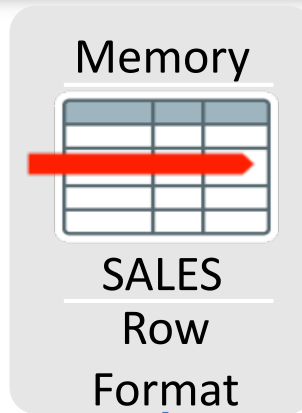
- **Transactions** run faster on row format
 - Example: Insert or query a sales order
 - Fast processing few rows, many columns
 - The BUFFER CACHE is used

Column



- **Analytics** run faster on column format
 - Example : Report on sales totals by region
 - Fast accessing few columns, many rows
 - The In-Memory Column Store is used

Transactions
In the
Buffer Cache



Analytics
In the
In-Memory
Column Store



IM – Setting INMEMORY

Create EMP and put it into the IM:

```
SQL> CREATE TABLE emp8  
      (EMPNO number, ENAME varchar2 (30)) INMEMORY;
```

Table created.

Alter DEPT table to be in the IM:

```
SQL> ALTER TABLE dept INMEMORY;
```

Table altered.

Enable object INMEMORY or PRIORITY set to NONE, and want to **populate immediately, you MUST:**

- ▶ Force a full table scan
- ▶ Use DBMS_INMEMORY.POPULATE procedure





Running queries using IM

Execution Plans for IM (may use IM for this query):

```
SELECT EMPNO
FROM EMP
ORDER BY EMPNO;
```

```
EMPNO
-----
7839 ...
```

Id	Operation	Name
0	SELECT STATEMENT	
1	SORT AGGREGATE	
2	PARTITION RANGE ALL	
* 3	TABLE ACCESS INMEMORY FULL	EMP



Online Table Move

Zero DownTime

▶ ALTER TABLE ... MOVE ... **ONLINE**

▶ ALTER TABLE ... MOVE PARTITION ... **ONLINE**

▶ ALTER TABLE ... MOVE SUBPARTITION ... **ONLINE**

▶ ALTER TABLE consultant_details
MOVE **ONLINE**

COMPRESS

TABLESPACE data_ts1

UPDATE INDEXES (idx1 TABLESPACE index_ts1,
idx2 TABLESPACE index_ts2);

- Table move operation now also supports automatic index maintenance as part of the move.
- With UPDATE INDEXES clause, the indexes remain usable during the move operation

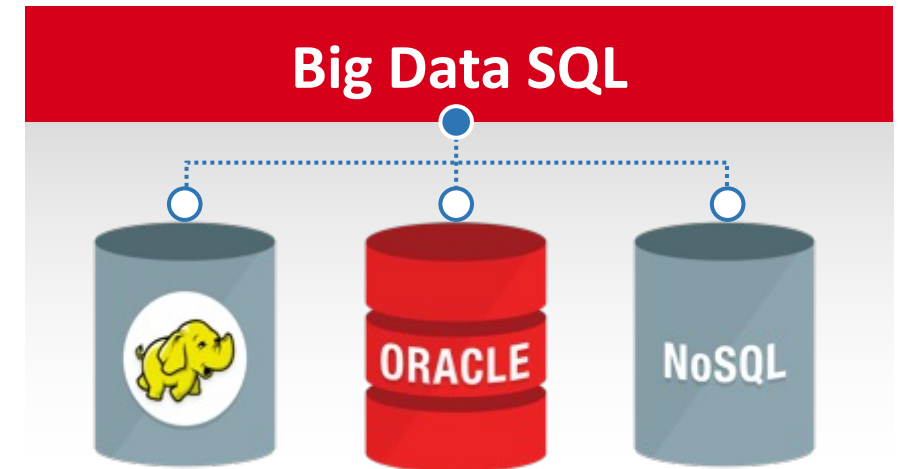


Fast SQL access for Relational, Hadoop and NoSQL

Using Oracle Big Data SQL

Oracle Database 12c ORACLE DATABASE 12c
Release 2 Performance
Tuning Tips and Techniques

- ▶ Unified SQL language for all data sources
 - With full power of Oracle SQL
- ▶ Massively parallel, distributed query processing
 - Local processing using 'Smart Scan' technology
 - **Scalable joins between data sources**
- ▶ Secure data access
 - Redaction and row-based security on all data sources



FYI

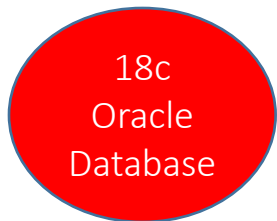
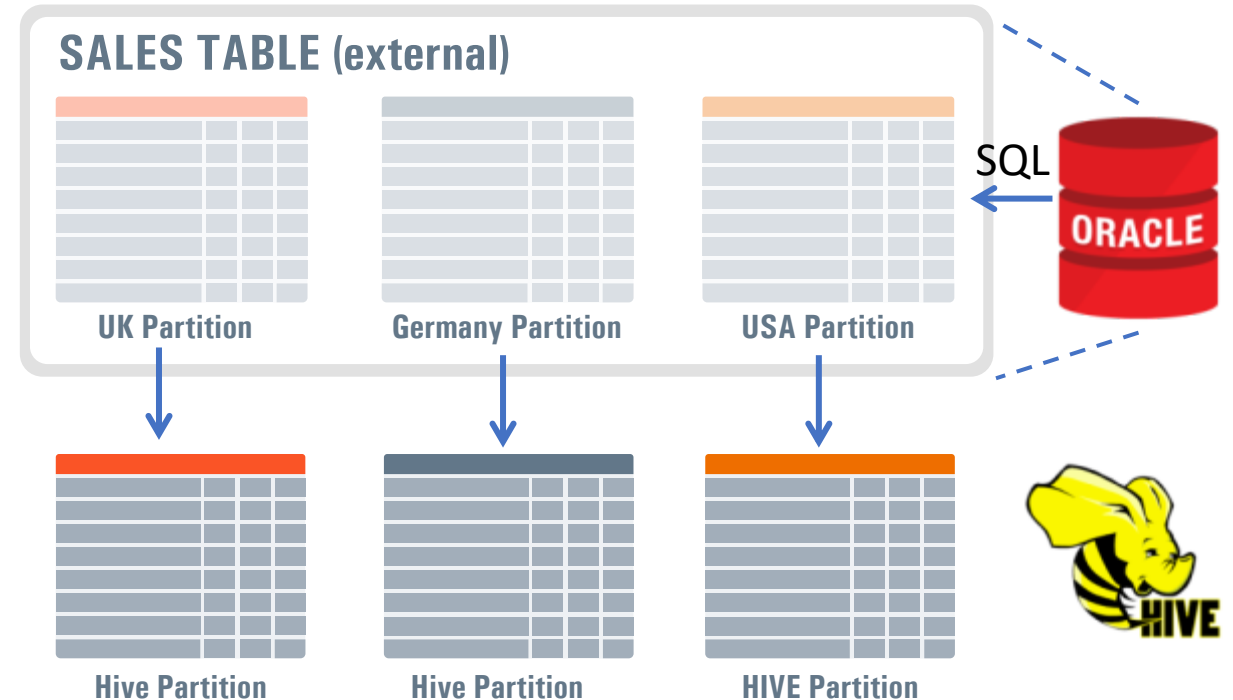
Big Data Innovations

Partitioned External Tables



Oracle Database 12c
Release 2 Performance
Tuning Tips and Techniques

- ▶ External tables can be partitioned
 - using any partitioning technique
- ▶ Partition pruning
 - For faster query performance
- ▶ Basic partition maintenance
 - Add, drop, exchange



In 18c: Inline and In-Memory External Tables

In 19c: In-Memory external tables support for ORACLE_HIVE & ORACLE_BIGDATA drivers.

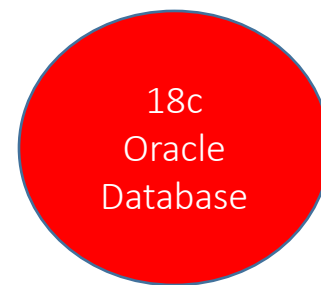


Ensure you use Key Oracle 12c R2 Features

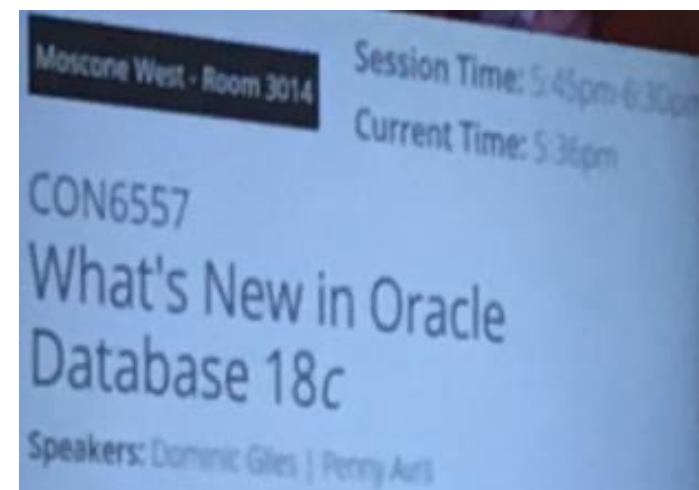
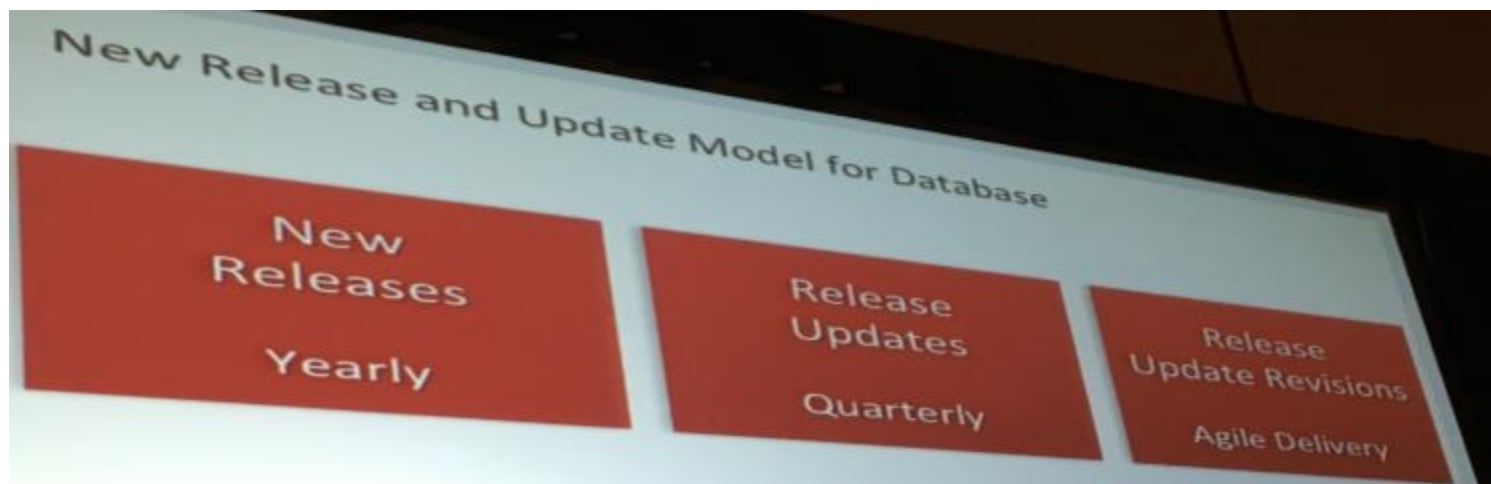
- ▶ **Long Names for Consolidation** of non-Oracle Databases
- ▶ Are you using **Partitioning** Improvements?
- ▶ Partition **External tables** for Big Data (such as access to Hadoop)
- ▶ Do you use **Big Data SQL** to access other NoSQL Data?
- ▶ Move Tables, Tablespaces & Partitions with **ONLINE**
- ▶ Do you leverage Sharding (is it necessary for you / don't do it if it's not)?
- ▶ No need for Read-Only PDB clone in 12cR2.
- ▶ PDB Memory Parameters can be set for **DB_CACHE_SIZE** & **SHARED_POOL_SIZE**
- ▶ **PDB Hot Clone, PDB Refresh, PDB Relocate** all came out for multitenant
- ▶ **Encrypt** SYSTEM, SYSAUX and UNDO

Oracle Database 18c

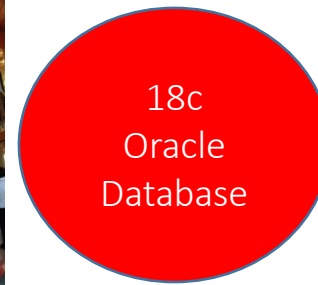
Simplified Version Number Timelines With RUs and RURs



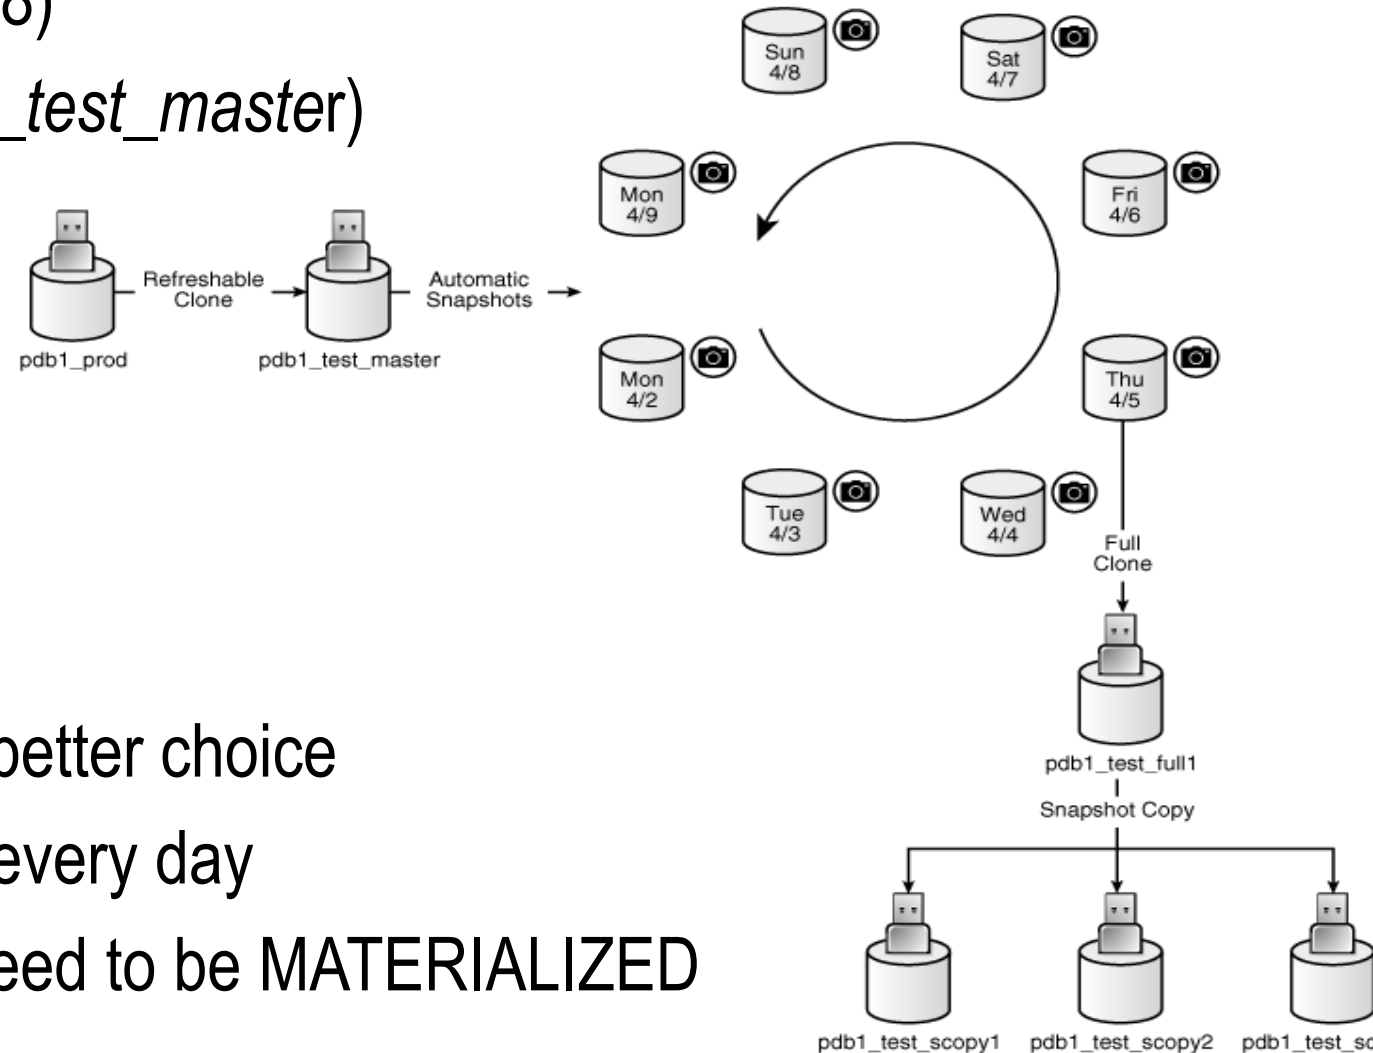
- ▶ 3 digit format
 - Year.Update.Revision
- ▶ Year is the last 2 digits of year a release is delivered
 - e.g. 18 used for release date shipping Dec 2017 or early 2018
- ▶ Update tracks Release Update (RU)
- ▶ Revision tracks the associated RU Revision levels (0,1,2)



Snapshot Carousel



- ▶ Point in Time PDB Copy (Default is 8)
- ▶ Create a Refreshable Master (*pdb1_test_master*)
- ▶ Master creates clones daily
- ▶ Use for testing
- ▶ Use to make clones
- ▶ Use to restore back in time
- ▶ Could take 1 per day at 12:01AM
- ▶ Could refresh more often if that's a better choice
- ▶ Could take one prior to a data load every day
- ▶ Creates a FULL PDB ... does not need to be MATERIALIZED



Snapshot Carousel



18c
Oracle
Database



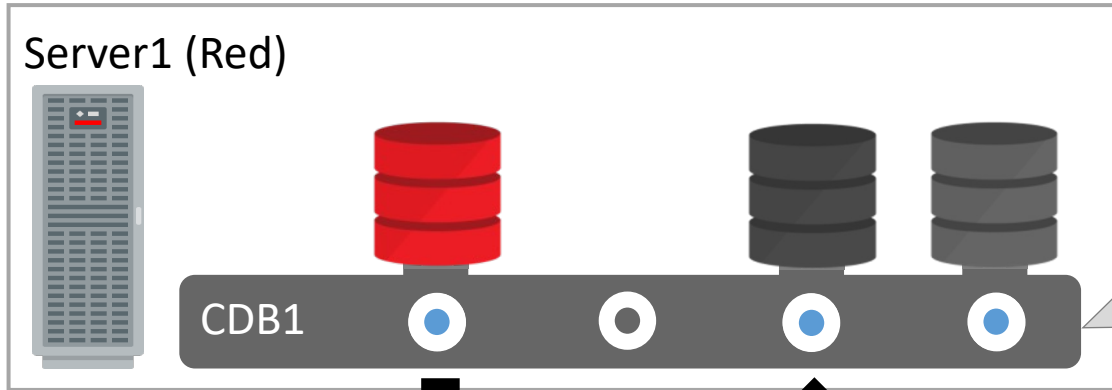
More information on Snapshots:

```
SELECT      CON_ID, CON_NAME, SNAPSHOT_NAME, SNAPSHOT_SCN
            AS snap_scn, FULL_SNAPSHOT_PATH
FROM        DBA_PDB_SNAPSHOTS
ORDER BY   SNAP_SCN;
```

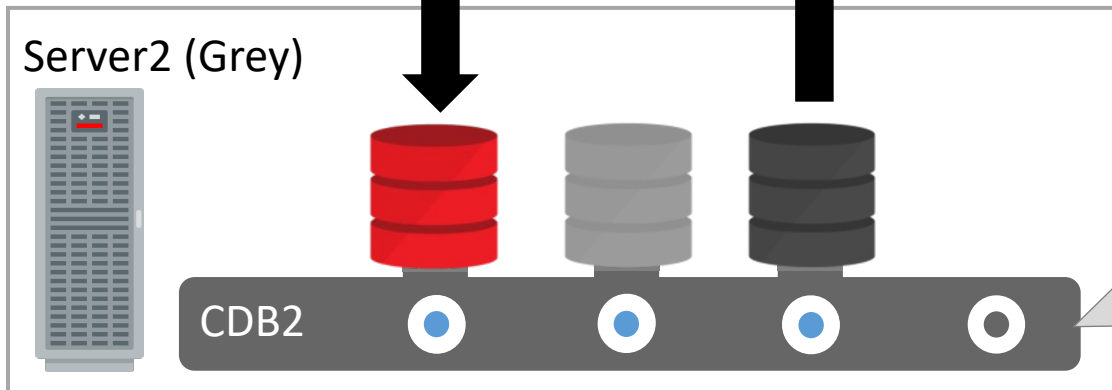
CON_ID	CON_NAME	SNAPSHOT_NAME	SNAP_SCN	FULL_SNAPSHOT_PATH
3	CDB1_PDB1	CDB1_PDB1_BEFORE	2962078	/.../dbs/snap_3489077498_2962078.pdb
3	CDB1_PDB1	CDB1_PDB1_AFTER	2962938	/.../dbs/snap_3489077498_2962938.pdb
...				

Per-PDB Switchover

18c
Oracle
Database

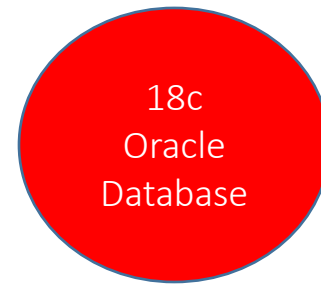


```
create pluggable database Grey  
from Grey@CDB2_Link  
refresh mode auto every 2 minutes;
```

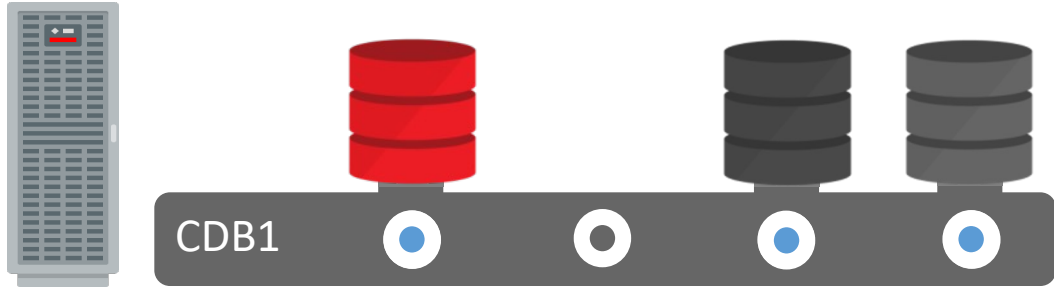


```
create pluggable database Red  
from Red@CDB1_Link  
refresh mode auto every 2 minutes;
```

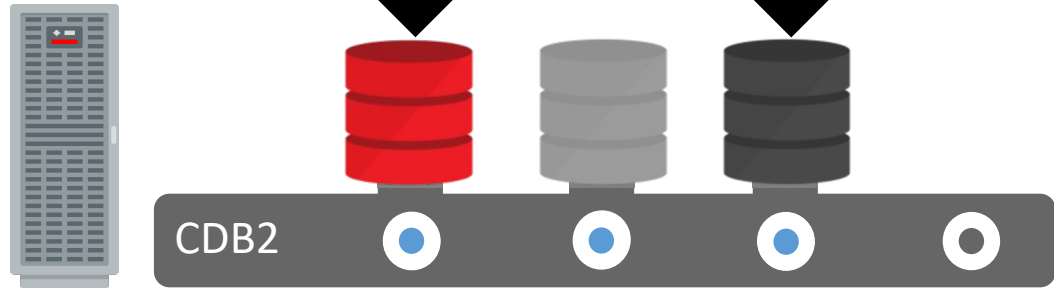

Per-PDB Switchover



Server1 (Red)



Server2 (Grey)



1. `alter pluggable database refresh mode auto every 2 minutes from Grey@dblink switchover;`
2. `alter pluggable database Grey open read write;`

PDB Switchover Clause

- ▶ Reverses the roles between a refreshable clone PDB and a primary PDB.
- ▶ The former Refreshable clone PDB becomes the primary PDB, which can now be opened in read write mode.
- ▶ The *formerly* primary PDB now is the refreshable clone and can only be opened in READ ONLY mode.
- ▶ This command must be executed from the primary PDB.
- ▶ The dblink must point to the root CDB where the refreshable clone PDB resides.

```
alter pluggable database refresh mode auto every 2  
minutes from new_pdb@dblink switchover;
```

In-Memory External Tables

```
CREATE TABLE big_hadoop_table (cust_no NUMBER, ...)
```

ORGANIZATION EXTERNAL

```
( TYPE ORACLE_LOADER  
  DEFAULT DIRECTORY admin_dat_dir  
  ACCESS PARAMETERS  
  ( records delimited by newline  
    badfile admin_bad_dir:'...bad'  
    logfile admin_log_dir:'...log'  
    fields terminated by ','  
    missing field values are null  
    ( prod_no, ... ) )  
  LOCATION ('filename.csv'))
```

```
REJECT LIMIT UNLIMITED
```

```
INMEMORY;
```

Approximate Query Expanded in 12cR2

Advances in 18c – Top-N

Other initialization parameters:

approx_for_aggregation=TRUE

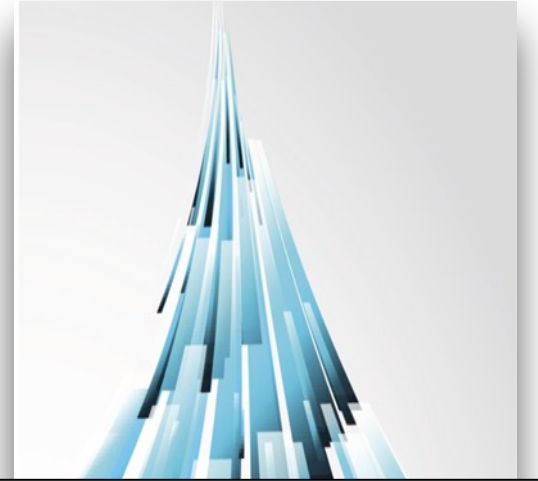
approx_for_percentile=TRUE

Other Approximate Functions:

- ▶ APPROX_COUNT_DISTINCT_DETAIL
- ▶ APPROX_COUNT_DISTINCT_AGG
- ▶ TO_APPROX_COUNT_DISTINCT
- ▶ APPROX_MEDIAN
- ▶ APPROX_PERCENTILE
- ▶ APPROX_PERCENTILE_DETAIL
- ▶ APPROX_PERCENTILE_AGG
- ▶ TO_APPROX_PERCENTILE

18c
Oracle
Database

In 18c: Top-N Approximate
Aggregation



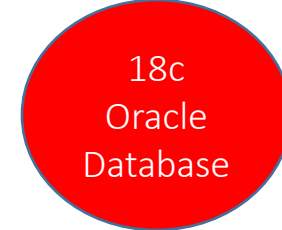
Feature	11.2	12.1	12.2	18c	19c	21c
Approximate Query Processing Data Warehousing / Big Data → Approximate Queries This release extends the area of approximate query processing by adding approximate percenti... 12.2 Release All Offerings	✗	✗	✓	✓	✓	✓
Approximate Top-N Query Processing Data Warehousing / Big Data → Approximate Queries Data analysis applications heavily use aggregate functions. Approximate query processing (avai... 18c Release All Offerings	✗	✗	✗	✓	✓	✓

row(s) 1 - 2 of 2

Also in 12cR2 is support for Materialized Views and Query Rewrite

APPROX_RANK Experiment (Oracle on 18c)

1G Temp to 0; 1G+ Sort to only 50M



Operation	Name	Line ID	Estimated Rows	Cost	Timeline(101s)	Executions	Actual Rows	Memory (Max)	Temp (Max)	Other	IO Requests	IO Bytes	Activity %
SELECT STATEMENT		0				33	4,492						
PX COORDINATOR		1				33	4,492						
PX SEND QC (RANDOM)	:TQ10001	2	200M	1,287K		16	4,492						
VIEW		3	200M	1,287K		16	4,492						
WINDOW SORT PUSHED RANK		4	200M	1,287K		16	6,814	813MB	32MB	4	48KB	.8	
SORT GROUP BY		5	200M	1,287K		16	14M	1GB	1GB	8,147	2GB	26	
PX RECEIVE		6	200M	1,287K		16	195M						.68
PX SEND HASH	:TQ10000	7	200M	1,287K		16	195M	101MB					2.46
SORT GROUP BY		8	200M	1,287K		16	195M						1.72
PX BLOCK ITERATOR		9	200M	1,086K		16	200M						
TABLE ACCESS FULL			200M	1,086K		468	200M				449K	438GB	68

Exact query

Lots of temp

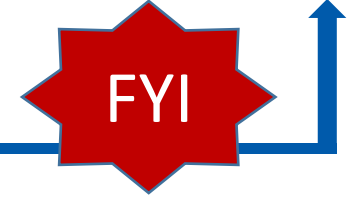
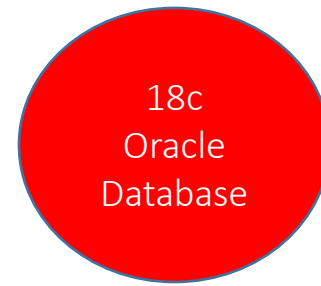
Operation	Name	Line ID	Estimated Rows	Cost	Timeline(71s)	Executions	Actual Rows	Memory (Max)	Temp (Max)	Other	IO Requests	IO Bytes	Activity %
SELECT STATEMENT		0				33	4,245						
PX COORDINATOR		1				33	4,245						
PX SEND QC (RANDOM)	:TQ10001	2	4,370	1,083K		16	4,245						
SORT GROUP BY APPROX		3	4,370	1,083K		16	4,245						
PX RECEIVE		4	4,370	1,083K		16	6,919						
PX SEND HASH	:TQ10000	5	4,370	1,083K		16	6,919						
SORT GROUP BY APPROX		6	4,370	1,083K		16	6,919	47MB					15
PX BLOCK ITERATOR		7	200M	1,082K		16	200M						
TABLE ACCESS FULL			200M	1,082K		468	200M				449K	438GB	85

Approx query

20X memory reduction

No temp

Ensure you use Key Oracle 18c Features



- ▶ Leverage **ADW & ATP** for your own PDB on **Exadata** (along with Machine Learning)
- ▶ Learn to leverage the **Snapshot Carousel** for cloning benefits
- ▶ Learn how to do a **PDB Switchover**
- ▶ Leverage **INMEMORY** for **External Tables**
- ▶ Leverage Approximate Query enhancements including **Top-N & APPROX_RANK**
- ▶ Use Oracle **18c XE to test** INMEMORY, PDBs, Spatial, and Advanced Analytics
- ▶ **Docker Support** for 18c
- ▶ **Private Temporary** Tables to go along with GTT
- ▶ Integration with **Active Directory**

More recently

livesql.oracle.com (19c Live – NOT DBA)



19^c ORACLE[®]
Database

☰ Live SQL

🏠 Home

📄 SQL Worksheet

☰ My Session

🗄️ Schema

🔮 Quick SQL

📄 My Scripts

🎓 My Tutorials

📁 Code Library

SQL Worksheet

```
1 select * from v$version;  
2
```

BANNER	BANNER_FULL
Oracle Database 19c EE Extreme Perf Release 19.0.0.0.0 - Production	Oracle Database 19c EE Extreme Perf Release 19.0.0.0.0 - Production Version 19.17.0.0.0

Download CSV

2023 Oracle · Live SQL 23.1.5, running Oracle Database 19c EE Extreme Perf - 19.17.0.0.0 · Database De

Built with ❤️ using Oracle APEX · Privacy · Terms of Use



19^c ORACLE[®]
Database

Access the Oracle Docs Now at:
docs.oracle.com/en/database/oracle/oracle-database/19 (See @richniemiec for link)

The screenshot displays the Oracle Help Center interface for Oracle Database 19c. On the left, a navigation sidebar lists categories such as 'Get Started', 'What's New', 'Install and Upgrade', 'Administration', 'Development', 'Security', 'Performance', 'Clustering', 'High Availability', 'Data Warehousing', 'Spatial and Graph', 'Distributed Data', and 'Books'. The 'What's New' category is highlighted, and an arrow points from this menu item to the 'What's New' section of the main content area. The main content area features a breadcrumb trail: 'Home / Database / Oracle / Oracle Database / Release 19'. Below the breadcrumb, the heading 'Oracle Database 19c' is followed by a 'What's New' section. This section contains three featured items: 'Database Features App' (with a red icon), 'New Features Guide' (with a purple icon), and 'Interactive Architecture Diagram' (with a blue icon). Each item includes a brief description of its purpose. A green circular icon with three horizontal lines is visible in the top right corner of the page.



Nice Oracle Database Features App by Version!

docs.oracle.com/en/database/oracle/oracle-database/19 (See @richniemiec for link)

The screenshot shows the Oracle Database Features App interface. On the left, there are search filters: a search bar with 'database in-memory', a 'Focus Area' dropdown set to 'All Focus Areas', and a 'Version' section with checkboxes for 11.2, 12.1, 12.2, 18c, and 19c, plus a 'New features only' option. The main content area displays a list of features. One feature, 'Big Data and Performance Enhancements for In-Memory External Tables', is highlighted, and its details are shown in a modal window.

Feature Details

Big Data and Performance Enhancements for In-Memory External Tables
Performance → In-Memory Column Store

In-Memory external tables add support for ORACLE_HIVE and ORACLE_BIGDATA drivers, parallel query, Oracle Real Application Clusters, Oracle Active Data Guard, and on-demand population.

Business Benefit: By using the new Big Data drivers, you avoid the cost and complexity of materializing data before populating it into the In-Memory Column Store (IM column store). You can use the SQL analytical capabilities of Oracle Database and Database In-Memory to analyze both internal and external data. Support for parallel query and full scan population means applications have fewer limitations when accessing data that resides outside the database.

Release Availability	11.2	12.1	12.2	18	19
Release Availability	✗	✗	✗	✗	✓
Initial Release	19c				

[View Documentation](#)

Nice Oracle Database Features App (Also Deprecated Features)!

docs.oracle.com/en/database/oracle/oracle-database/19 (See @richniemiec for link)



Search: long data type

Focus Area: All Focus Areas

Version: 11.2, 12.1, 12.2, 18c, 19c, New features only

Reset

Feature: Support for Extracting Data From Tables Containing LONG Data Type

Oracle Warehouse Builder

11.2 Release

Feature Details

Support for Extracting Data From Tables Containing LONG Data Type

Data Warehousing / Big Data → Oracle Warehouse Builder

Oracle Warehouse Builder can now generate SQL*Plus code to extract data from database schemas supporting the deprecated LONG data type, such as occurs in PeopleSoft application data sources.

Business Benefit: Support for LONG data types used in PeopleSoft data enables OWB users to integrate more effectively with PeopleSoft data or any other data source that uses the LONG data type.

Release Availability: 11.2 (green check), 12.1 (yellow warning), 12.2 (yellow warning), 18 (yellow warning), 19 (yellow warning)

Initial Release: 11.2.0.1

Deprecated in Release: 12.1.0.1 [More Information](#)

ONLY FEATURES INTRODUCED IN THIS RELEASE. PLEASE CHECK THE LICENSE AGREEMENT FOR THE NEW FEATURES IN THIS RELEASE.

dep·re·cate /'depreˌkāt/ verb past tense: deprecated; past participle: deprecated 1. express disapproval of.



Access the Oracle Docs Now at:

docs.oracle.com/en/database/oracle/oracle-database/19 (See @richniemiec for link)

The screenshot displays the Oracle Help Center interface for Oracle Database 19c. On the left, a navigation menu lists various topics such as 'Get Started', 'What's New', 'Install and Upgrade', 'Administration', 'Development', 'Security', 'Performance', 'Clustering', 'High Availability', 'Data Warehousing', 'Spatial and Graph', 'Distributed Data', and 'Books'. The 'What's New' item is highlighted, and an arrow points from it to the 'What's New' section of the main content area. The main content area features a search bar at the top, a breadcrumb trail (Home / Database / Oracle / Oracle Database / Release 19), and the title 'Oracle Database 19c'. Below the title, the 'What's New' section is displayed, containing three featured items: 'Database Features App' (with a red icon), 'New Features Guide' (with a purple icon), and 'Interactive Architecture Diagram' (with a blue icon). Each item includes a brief description of its purpose.

Nice Oracle Interactive Architecture Diagram

docs.oracle.com/en/database/oracle/oracle-database/19 (See @richniemiec for link)

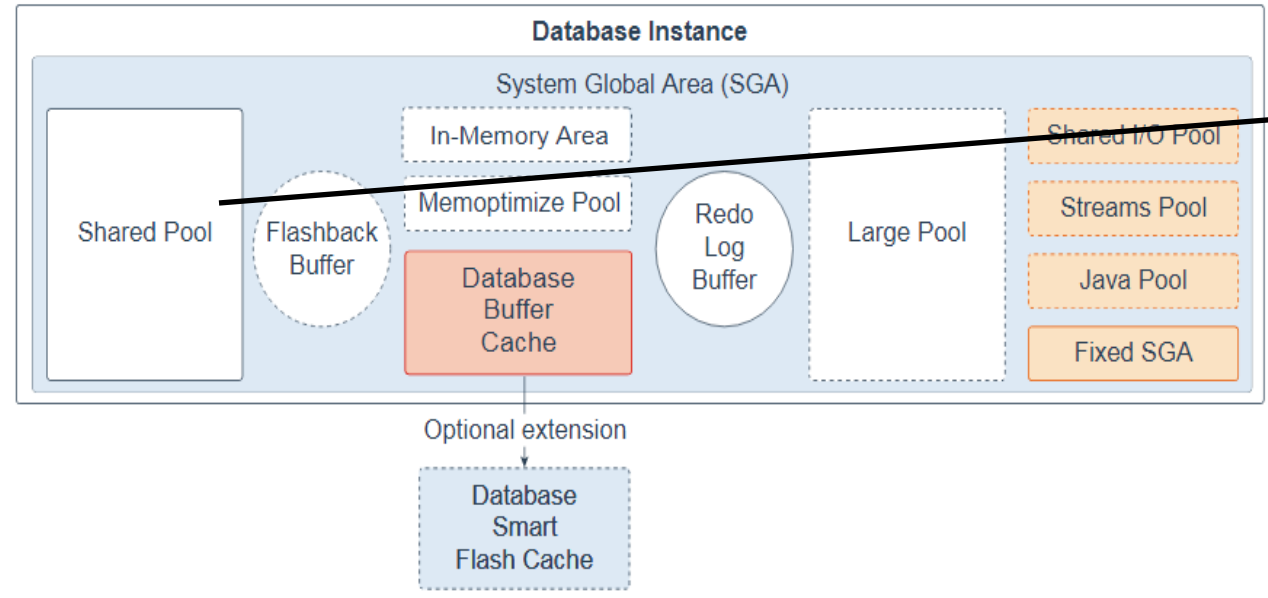


Understanding the Oracle Database 19c Technical Architecture

This interactive diagram shows the Oracle Database 19c technical architecture. [SHOW INSTRUCTIONS](#)

First Previous Next Slide 3 of 29

System Global Area



The System Global Area (SGA) is the memory area that contains data and control information for one Oracle Database instance. All server and background processes share the SGA. When you start a database instance, the amount of memory allocated for the SGA is displayed. The SGA includes the following data structures:

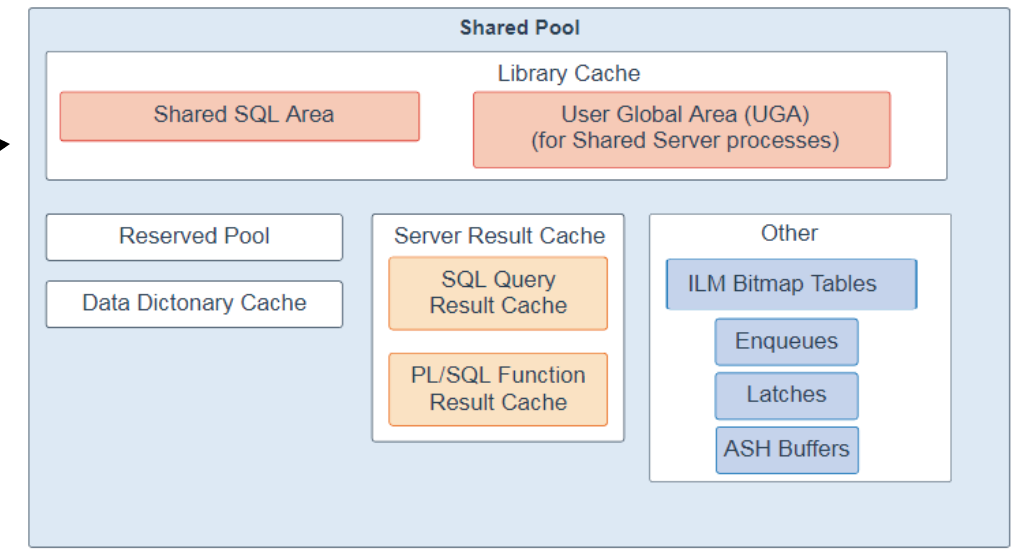
- **Shared pool:** Caches various constructs that can be shared among users; for example, the shared pool stores parsed SQL, PL/SQL code, system parameters, and data dictionary information. The shared pool is involved in almost every operation that occurs in the database. For

Understanding the Oracle Database 19c Technical A

This interactive diagram shows the Oracle Database 19c technical architecture. [SHOW INSTRUCTIONS](#)

First Previous Next Slide 6 of 29

System Global Area / Shared Pool



The shared pool is a component of the System Global Area (SGA) and is responsible for caching various types of program data. For e shared pool stores parsed SQL, PL/SQL code, system parameters, and data dictionary information. The shared pool is involved in alm operation that occurs in the database. For example, if a user executes a SQL statement, then Oracle Database accesses the shared po

The shared pool is divided into several subcomponents:

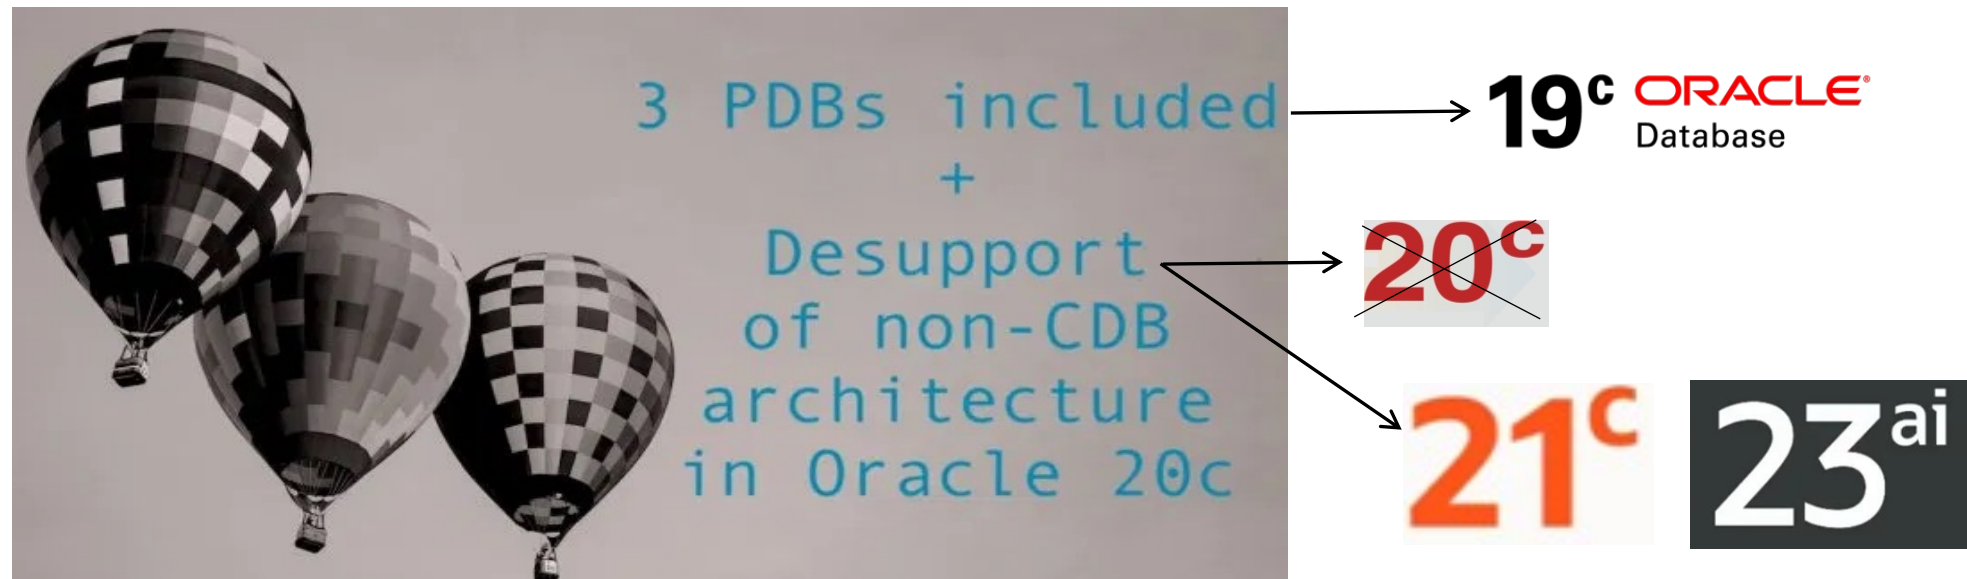
- **Library cache:** Is a shared pool memory structure that stores executable SQL and PL/SQL code. This cache contains the shared PL/SQL areas and control structures, such as locks and library cache handles. In a shared server architecture, the library cache a



19^c ORACLE[®]
Database

Best Feature in 19c that Came Out!

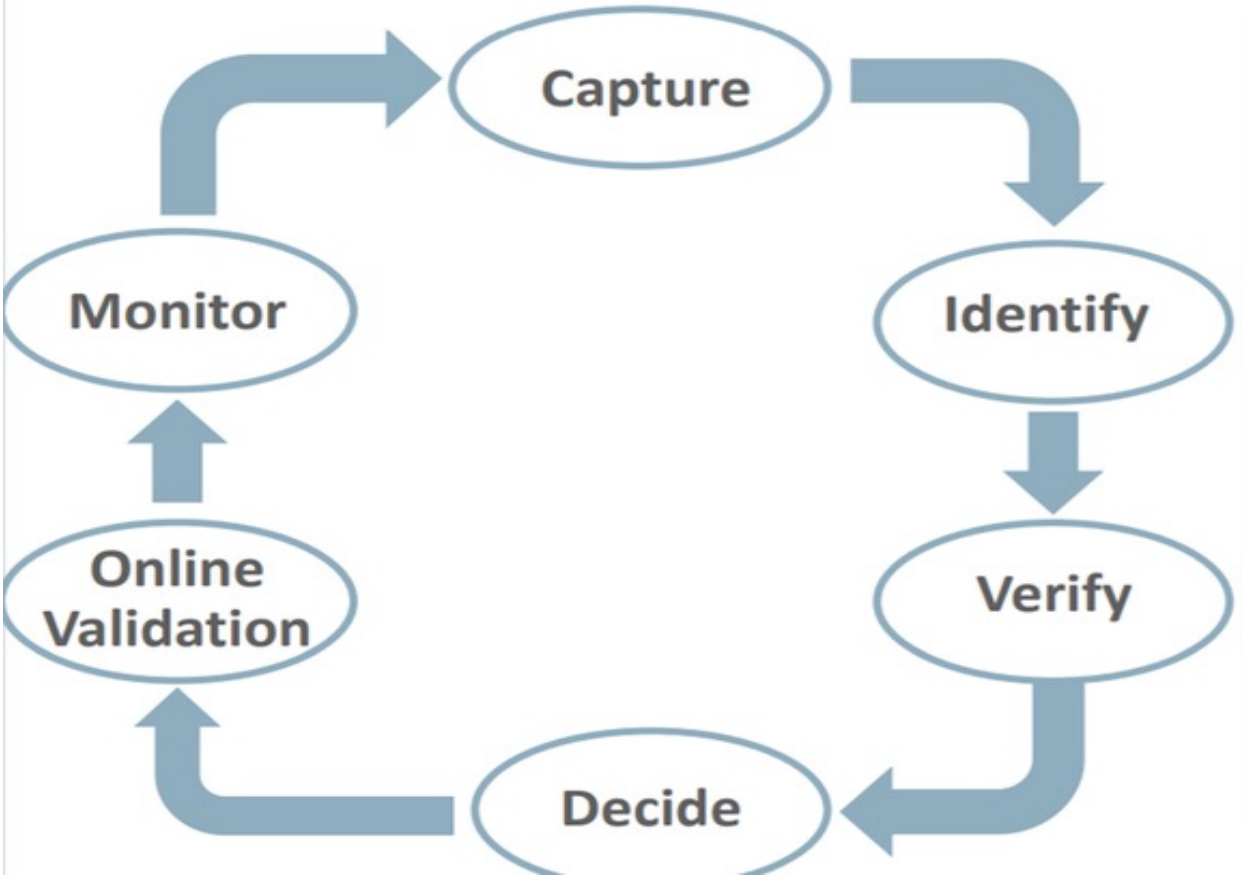
- In 21c non-CDB is Desupported! To ease the migration to this architecture, from Oracle Database 19c onwards, the multitenant architecture supports up to 3 user-created Pluggable Databases of any type (without buying multitenant option)
 - The **Multitenant Option** is required for deployment of **4 or more user-created PDBs**





19c - Automatic Index Method* - ONLY on Exadata!!

Automatic Indexing Methodology



- The Automatic Indexing methodology is based on a common approach to manual SQL tuning
- It **identifies** candidate indexes and **validates** them before **implementing**
- The entire process is fully automatic
- Transparency is equally important as sophisticated automation
 - All tuning activities are auditable via reporting

* Some screenshots from OOW Session & Oracle Cloud Day:





19c - Automatic Index Package (Default is OFF)

19^c ORACLE®
Database

- ▶ **Enable Reporting** AUTO_INDEX_MODE & **Creates invisible indexes if needed:**
EXEC DBMS_AUTO_INDEX.CONFIGURE('AUTO_INDEX_MODE','REPORT ONLY');
- ▶ **Enable** AUTO_INDEX_MODE & **Creates visible indexes if needed:**
EXEC DBMS_AUTO_INDEX.CONFIGURE('AUTO_INDEX_MODE','IMPLEMENT');
- ▶ **Disable** AUTO_INDEX_MODE & **disable any indexes that were created:**
EXEC DBMS_AUTO_INDEX.CONFIGURE('AUTO_INDEX_MODE','OFF');

19c - DBMS_AUTO_INDEX.REPORT for Info



19^c ORACLE[®]
Database

SUMMARY (AUTO INDEXES)

Index candidates	: 53
Indexes created (visible / invisible)	: 12 (12 / 0)
Space used (visible / invisible)	: 3.48 MB (3.48 MB / 0 B)
Indexes dropped	: 0
SQL statements verified	: 16
SQL statements improved (improvement factor)	: 16 (3x)
SQL statements disallowed from auto indexes	: 0
Overall improvement factor	: 3x



19c - Automatic Index Method - Hints*

Using Automatic Indexing Hints

- You can use hints to control if auto indexes will be used for a SQL statements
- The `USE_AUTO_INDEXES` hint instructs the optimizer to use auto indexes

```
SELECT /*+ USE_AUTO_INDEXES */ emp_id, emp_name, dept_id  
FROM employees  
WHERE dept_id > 50;
```

- The `NO_USE_AUTO_INDEXES` hint instructs the optimizer not to use auto indexes

```
SELECT /*+ NO_USE_AUTO_INDEXES */ emp_id, emp_name, dept_id  
FROM employees  
WHERE dept_id > 50;
```

Quarantine for Execution Plans for SQL Statements ONLY on Exadata (Cloud or On-Prem)!!



- ▶ **SQL statements terminated** due to excessive CPU and I/O **can be quarantined**
- ▶ **Execution Plans** for above SQL are **Quarantined**
- ▶ Enable/Disable a quarantine configuration:
DBMS_SQLQ.ALTER_QUARANTINE procedure
- ▶ Drop & Alter Quarantined Configuration (*unused quarantined configs deleted 53 weeks*):
DBMS_SQLQ.**DROP**_QUARANTINE & DBMS_SQLQ.**ALTER**_QUARANTINE
- ▶ **DBA_SQL_QUARANTINE** - details of quarantine configurations
- ▶ V\$SQL & GV\$SQL to get execution plan details of quarantined SQL



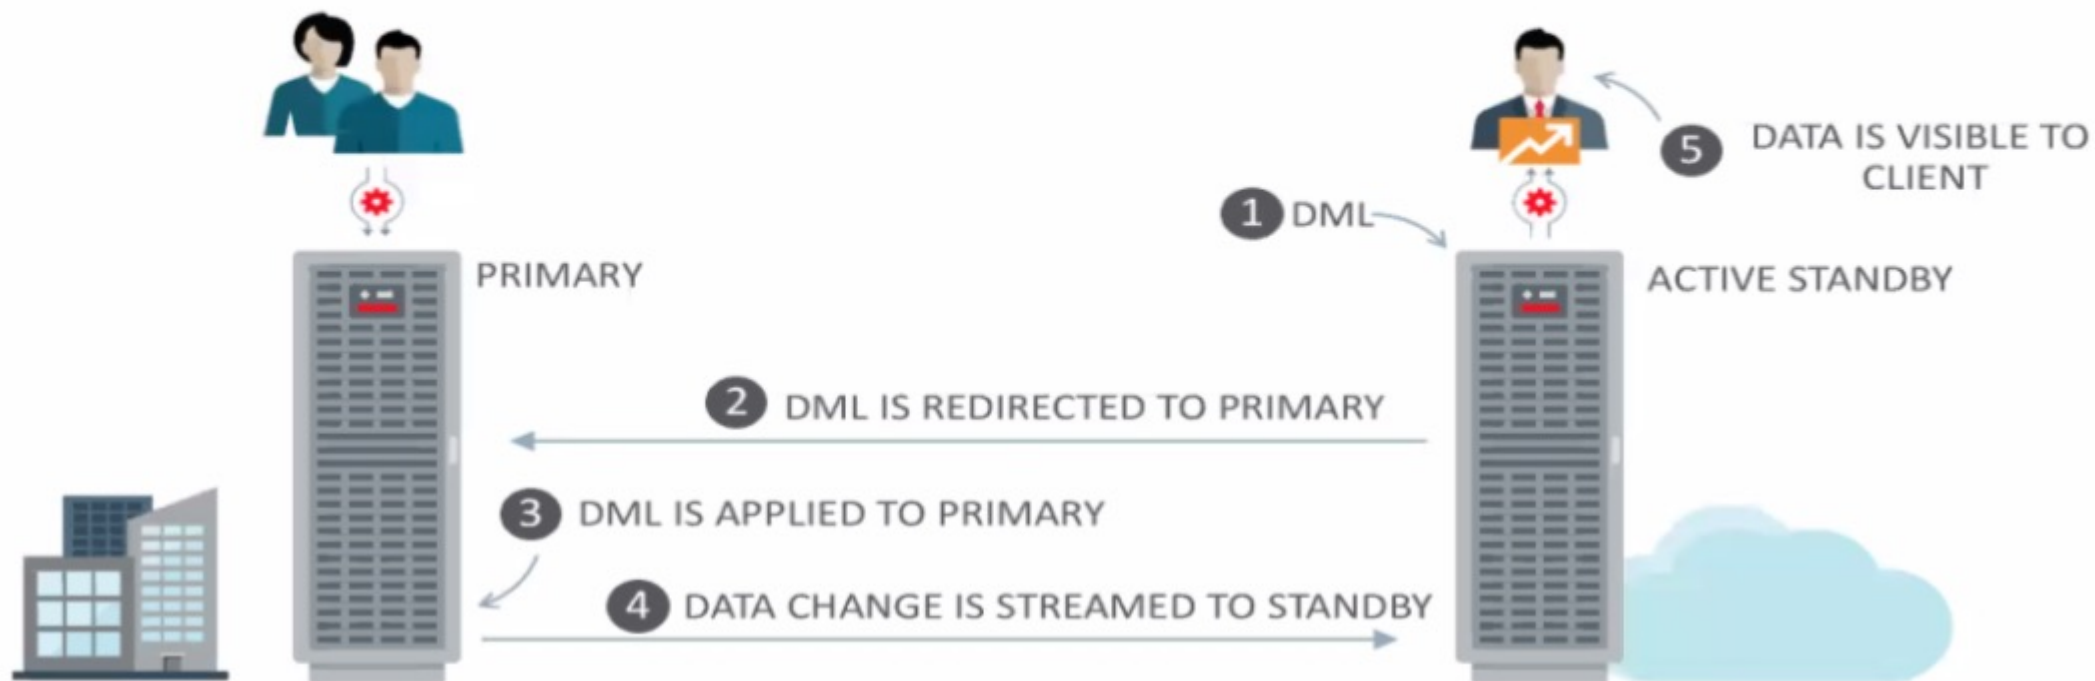
From Willie Hardie's talk ... Read-mostly Standby

19^c ORACLE[®]
Database

NEW IN
19^c

Active Data Guard : DML Redirection

- Read Mostly Standby





Ensure you use Key Oracle 19c Features

- ▶ Use **THREE** user-created PDBs without buying multitenant!
- ▶ **Automatic Indexing (Exadata Only)**
- ▶ **JSON-Object Mapping** - Map JSON data to/from SQL Object/Collection Types
- ▶ Oracle Database supplied **Schemas have Passwords Removed**
- ▶ **SQL Quarantine (Exadata Only)** - Execution Plans & SQL using high resources
- ▶ **In-Memory** support for **ORACLE_HIVE & ORACLE_BIGDATA** drivers
- ▶ **In-Memory** support for **Parallel Query, RAC, ADG**, and on-demand population
- ▶ **In-Memory Wait on Populate** function: `DBMS_INMEMORY_ADMIN.POPULATE_WAIT`
- ▶ **RAT & ADDM Support for PDBs** (can tune PDBs better)
- ▶ **REST enabled SQL Support** - APEX Oracle Rest Data Services (ORDS) 17.3+
- ▶ **Sharded RAC** (not really sharding) that allows **pinning partitions to a given node.**

19c Focus is Stability*, but also Speed/Security unlike other DBs



Oracle Database Annual Release Model

Fine Print to NEW Annual Release Model - Support Model as Described in MOS Note 742060.1

Innovation Release

- 2 years of Premier Support
- No Extended Support
- Great release to try new features
- Great release to develop apps on

Long Term Release

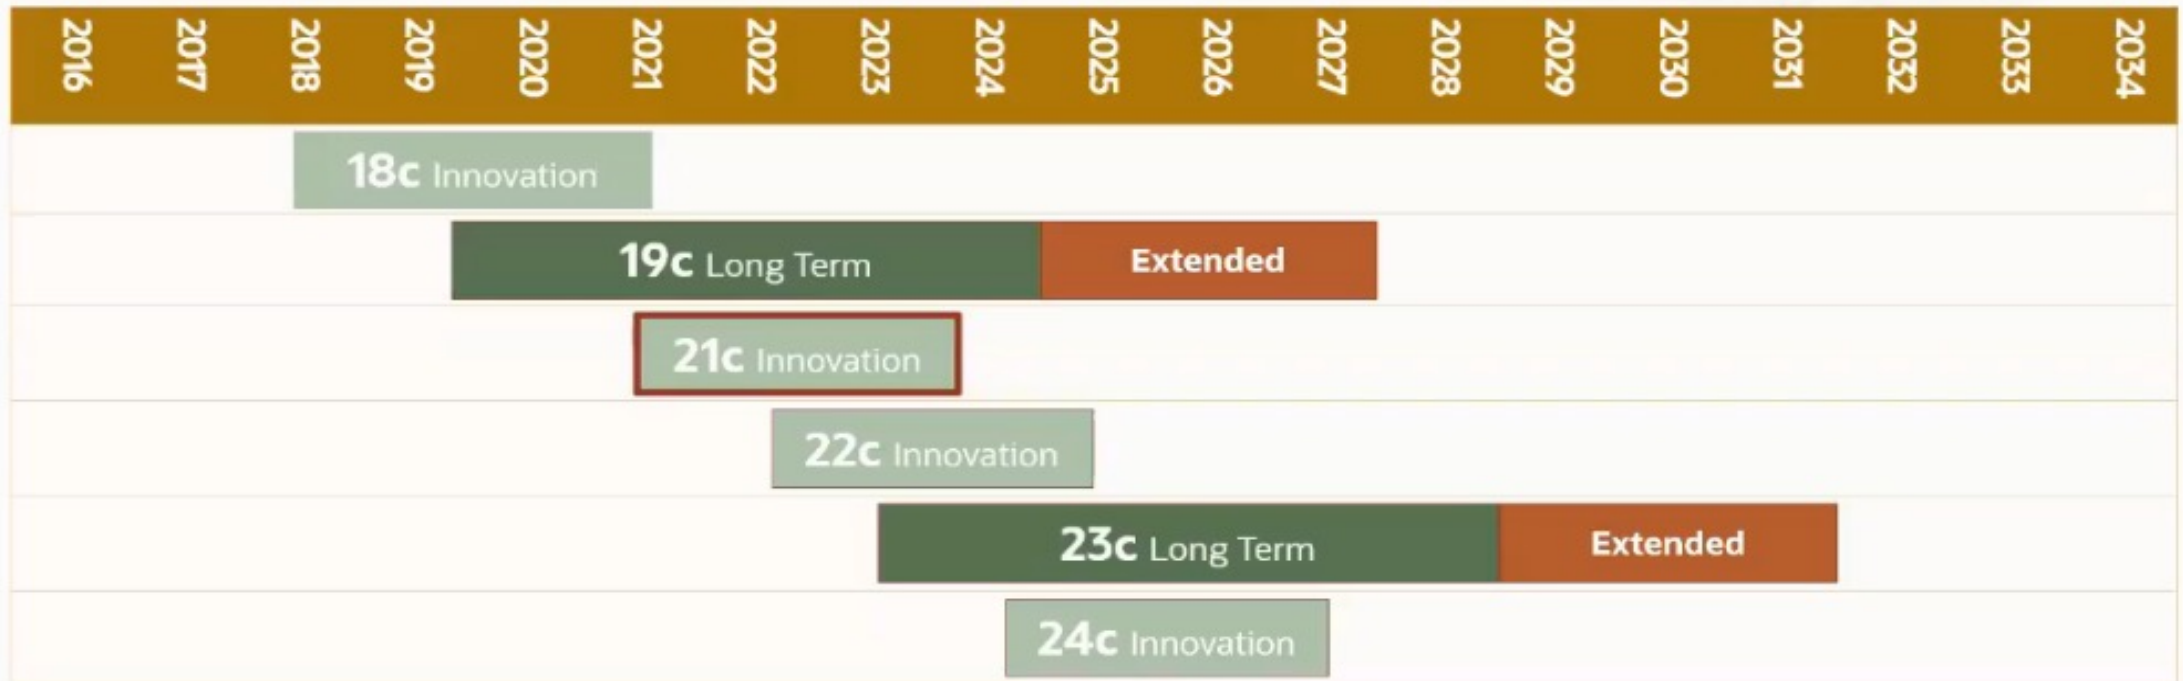
- 5 years of Premier Support
- 3 years of Extended Support
- Great release for production

Big Update - 21c Oracle DB*...23ai Timeline later!

*Thanks Jenny Tsai-Smith, Oracle

Projected Database Release and Support Timeline *(subject to change)*

Always check MOS Note 742060.1 for the latest schedule



- Innovation Release - 2 years of Premier Support, and no Extended Support
- Long Term Release - 5 years of Premier Support, and 3 years of Extended Support





My Favorite Newer 21c Features*

(*Andy Mendelsohn / Jenny Tsai-Smith)



More Innovations to Support Converged Database



-  In-Database JavaScript
-  JSON Speed & Flexibility
-  In-Memory Improvements
-  AutoML
-  Multitenant Data Guard
-  SQL Macros
-  Native Blockchain Tables
-  Graph Performance
-  Persistent Memory
-  Sharding Enhancements
-  Multitenant Security
-  Security Enhancements



AutoML in OML4Py

Automatically build and compare Machine Learning models



- Auto Model Selection
 - Identify in-database algorithm that achieves highest model quality
 - Find best model faster than with exhaustive search
- Auto Feature Selection
 - Reduce # of features by identifying most predictive
 - Improve performance and accuracy
- Auto Tune Hyperparameters
 - Significantly improve model accuracy
 - Avoid manual or exhaustive search techniques

Enables non-expert users to leverage Machine Learning

Larry said it was coming ... Focus on Security

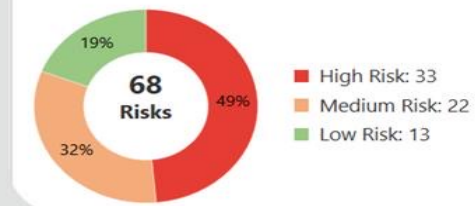
21^c

Oracle Data Safe

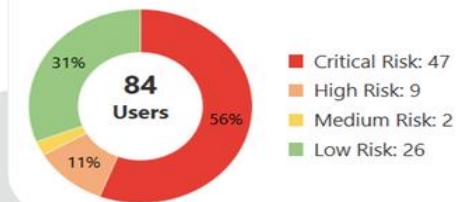
Autonomous Database | Now even more Secure

- Unified Database Security Control Center
 - Security Configuration Assessment
 - User Risk Assessment
 - User Activity Auditing
 - Sensitive Data Discovery
 - Data Masking
- Saves time and mitigates security risks
- Defense in Depth for all customers
- No special security expertise needed

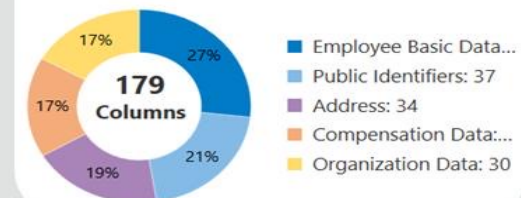
Security Assessment



User Assessment



Data Discovery



Available with Oracle Cloud Databases at no additional cost



Data Safe Console (thanks Russ Lowenthal)

21°C



ORACLE Cloud Search for resources and services

Data Safe Registered Databases

Databases registered with Data Safe

Registered Databases Enable Data Safe

Private Endpoints

ORACLE Cloud Search

- Core Infrastructure
- Compute
- Block Storage
- Object Storage
- File Storage
- Networking
- Database
- Autonomous Data Warehouse
- Autonomous Transaction Processing
- Bare Metal, VM, and Exadata
- Data Safe

ORACLE Cloud Search for resources and services US East (Ashburn)

DB 202003311814

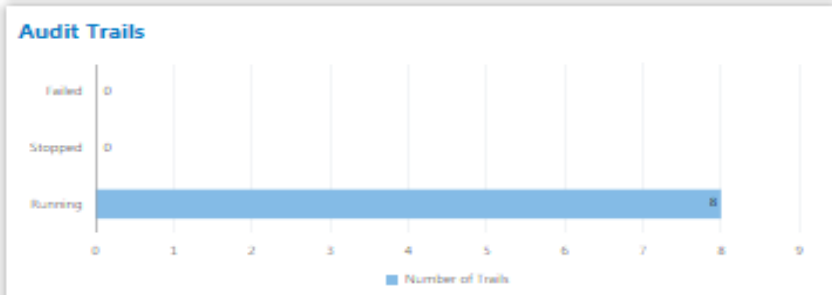
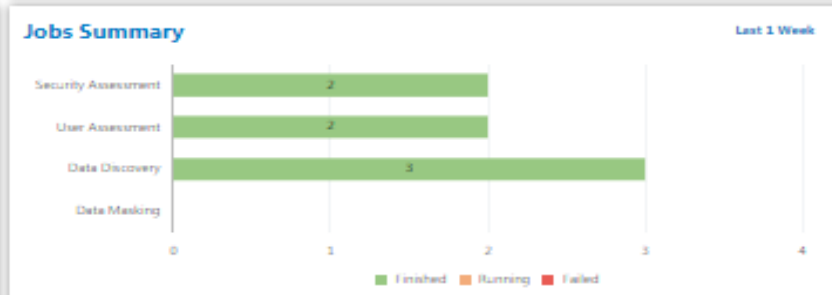
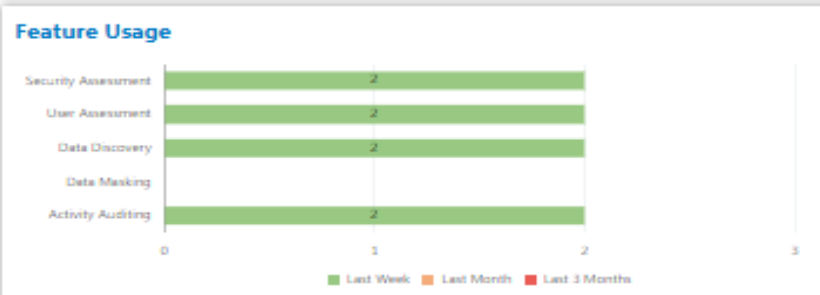
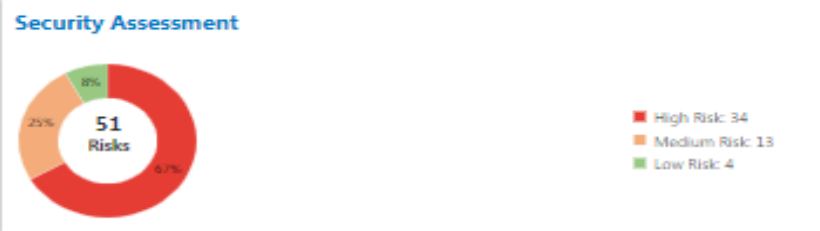
DB Connection Performance Hub Service Console Scale Up/Down More Actions

ORACLE Cloud Services Data Safe

Home Targets Library Reports Alerts Jobs

Security Assessment User Assessment Data Discovery

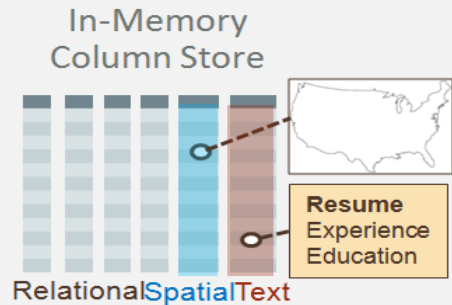
All Activity Admin Activity Login Activity



What's New in 21^c for Database In-Memory

Converged Workloads

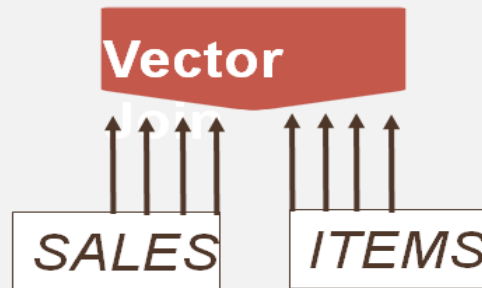
In-Memory Spatial and Text



- **10x faster** Spatial Analytics
- **3x faster** Text Analytics

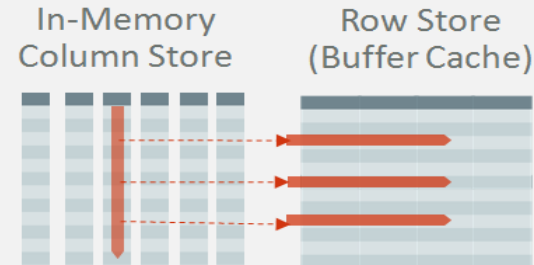
Performance

In-Memory Vector Joins



- In-Memory Joins using SIMD vector instructions
- **5-10x faster**

In-Memory Hybrid Scans



- In-Memory scans get missing values from buffer cache
- **10x faster** hybrid queries

Automation

Self-Managing In-Memory



- Auto in-memory populate / evict
- **No need to declare tables** INMEMORY



21^c 

No more Case Insensitive Password Files?!

Force Upgraded Password File to be Case Sensitive

Starting in Oracle Database 20c, the parameter to enable or disable password file case sensitivity is removed. All passwords in new password files are case-sensitive.

Case-sensitive password files provide more security than older password files that are case insensitive. Oracle recommends that you use case-sensitive password files. However, upgraded password files from earlier Oracle Database releases can retain their original case-insensitivity. You can force your password files to be case-sensitive by migrating password files from one format to another.

However, password files from earlier Oracle Database releases will by default retain their original case-insensitive verifiers. Oracle recommends that you force case sensitivity in these older password files by migrating the password file from one format to another and changing the password of any account that has only a 10G verifier, using the following syntax:

```
orapwd FILE=new_pwd_file_name INPUT_FILE=old_pwd_file_name [FORMAT=12.2]
```

Nice!  

Native Blockchain Tables*

Secure distributed ledger managed by a trusted provider (prevent fraud)

21^c

- **Specialized table allows normal SQL inserts (Append-Only) and queries**
 - **CREATE BLOCKCHAIN TABLE** LEDGER_OF_TRADES ...
 - Rows are cryptographically chained, chain is verifiable by participants
- **Blockchain Table** can participate in **transactions and queries with other tables**
- For blockchain uses cases where trusted third party runs the tamper-resistant ledger
 - Participants don't fully trust each other, but trust a third-party service provider – *e.g. escrow company*

Oracle Blockchain Table

Blockchain tables are append-only tables in which only insert operations are allowed. Deleting rows is either prohibited or restricted based on time. Rows in a blockchain table are made tamper-resistant by special sequencing & chaining algorithms. Users can verify that rows have not been tampered. A hash value that is part of the row metadata is used to chain and validate rows.

Blockchain tables enable you to implement a centralized ledger model where all participants in the blockchain network have access to the same tamper-resistant ledger.

21^c

Key Oracle 21c Features for APEX & Machine Learning

21^c 

- ▶ APEX **Dark Mode** for reducing Eye Strain (especially at night)
- ▶ APEX Data **Upload to support Excel, CSV, XML & JSON documents**
- ▶ APEX JET Chart Enhancements (can **Style the title, legend, and axes titles & labels**)
- ▶ Machine Learning **ALGO_XGBOOST** tree boosting algorithm for regression / classification
- ▶ Machine Learning **ALGO_MSET_SPRT** anomaly detection algorithm for critical processes (Multivariate State Estimation Technique-Sequential Probability Ration Test)
- ▶ New Aggregate Functions: BIT_AND_AGG, BIT_OR_AGG, BIT_XOR_ADD
- ▶ New Analytical Functions: KURTOSIS_POP & KURTOSIS_SAMP (measure tailedness)
- ▶ New Analytical Functions SKEWNESS_POP & SKEWNESS_SAMP (measure asymmetry)

Docs available now (link to 21c)! See @richniemiec

<https://docs.oracle.com/en/database/oracle/oracle-database/21>

21^c

Help Center Oracle Database 21c Search

Home / Database / Oracle / Oracle Database / Release 21

Oracle Database 21c

Get Started

- Learn About Oracle Database**
 - [Introduction to Oracle Database](#)
 - [Introduction to SQL](#)
 - [Oracle Database Quickstart](#)
 - [Run SQL with Oracle Live SQL](#)
 - [What's New in Release 21c](#)
- Development**
 - [Create Applications with SQL and PL/SQL](#)
 - [Develop Java Applications](#)
 - [Use JSON with SQL and PLSQL](#)
 - [Develop JSON based applications using SODA drivers](#)
- Oracle Multitenant**
- Oracle Database In-Memory**

REST API Reference

Hybrid In-Memory Scans (IMCS/Buffer Cache)



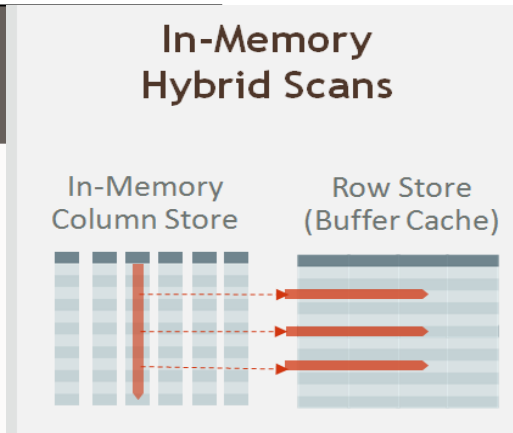
<https://docs.oracle.com/en/database/oracle/oracle-database/21>

21^c 

Database / Oracle / Oracle Database / Release 21

Learning Database New Features

Hybrid In-Memory Scans



Oracle Database supports in-memory scans when not all columns in a table have been populated into the In-Memory Column Store (IM column store).

This situation can occur when columns have been specified as `NO INMEMORY` to save space. Hybrid in-memory scans can access some data from the IM column store, and some data from the row store, improving performance by orders of magnitude over pure row store queries.

Automatic In-Memory Populate/Evict

<https://docs.oracle.com/en/database/oracle/oracle-database/21>



21^c ↑

Learning Database New Features

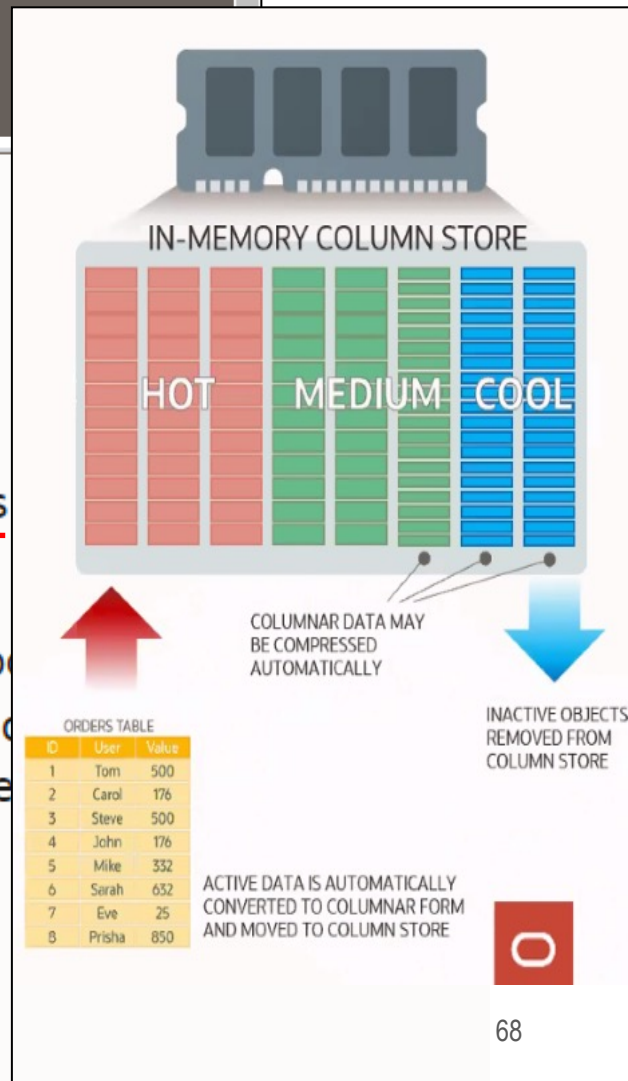
Automatic In-Memory Management Enhancements

Automatic In-Memory Management enables, populates, evicts, and recompresses segments intervention.

When INMEMORY_AUTOMATIC_LEVEL is set to HIGH, the database automatically enables and populates based on their usage patterns. Combined with support for selective column level eviction and In-Memory population is largely self-managing. This automation helps maximize the number of objects that can be populated into the In-Memory Column Store at one time.

Related Topics

- [Oracle® Database In-Memory Guide](#)



Some In-Memory Details

Oracle Database In-Memory BASE_LEVEL
Feature available since 19.8.0

Posted on July 20, 2020 by Mike.Dietrich in-Memory

Oracle Database In-Memory Base Level Feature



Andy Rivenes
PRODUCT MANAGER

Oracle Database In-Memory will now have a "Base Level" feature. This will allow the use of Database In-Memory with up to a 16GB column store for no extra cost. That's not a typo. Database In-Memory will be free to use with up to a 16GB column store!

With **Oracle Database 21^c Preview** we have introduced a new BASE_LEVEL value for the INMEMORY_FORCE parameter. When this new value is set the INMEMORY_SIZE parameter can be set up to a value of 16GB without having to license the Database In-Memory option. In fact, when the BASE_LEVEL value is set you cannot set the INMEMORY_SIZE parameter larger than 16GB. If you do then Oracle will generate an error message.

The 16GB limit applies at the container database (CDB) level. This means that all pluggable databases (PDBs) share the 16GB limit of the CDB. However, on RAC databases the Base Level feature allows a 16GB column store to be allocated on each RAC instance.

Feature tracking has been added to report usage of the Base Level feature as "In-Memory Base Level" so there should be no confusion as to whether the Base Level feature is being used or the full Database In-Memory option.

The new Base Level feature supports all Database In-Memory features, except:

- Automatic In-Memory (AIM)
- Compression levels other than MEMCOMPRESS FOR QUERY LOW
- Excluded columns (all columns of a table are populated)
- The CELLMEMORY feature on Exadata

21^c

Ability to Set the Default Tablespace Encryption Algorithm



You now can set the TABLESPACE_ENCRYPTION_DEFAULT_ALGORITHM dynamic parameter to define the default encryption algorithm for tablespace creation operations.

For example, if you set TABLESPACE_ENCRYPTION_DEFAULT_ALGORITHM to AES256, then future tablespace creation operations will use AES256 as the default encryption algorithm.

TABLESPACE_ENCRYPTION_DEFAULT_ALGORITHM applies to both offline and online tablespace encryption operations. In addition, when you create a new tablespace using Database Configuration Assistant (DBCA), you can set the default tablespace encryption algorithm by using the DBCA command line for silent installations.

Supported encryption algorithms are AES128, AES192, AES256, and 3DES168 If you do not set TABLESPACE_ENCRYPTION_DEFAULT_ALGORITHM, then the default encryption algorithm is the default that was used in previous releases: AES128.

Related Topics

- [Oracle® Database Advanced Security Guide](#)

Parent topic: [Oracle Advanced Security](#)

Ensure you use Key Oracle 21c Management Features

21c 

- ▶ Oracle **Data Pump** can perform exports from **ADW/ATP** to the Object Store!!!
- ▶ Oracle Data Pump can both **INCLUDE / EXCLUDE objects** in same export/import
- ▶ Oracle Data Pump can **Resume Transportable Tablespace** export/import that stopped
- ▶ Oracle Data Pump can **Parallelizes Transportable Tablespace** metadata
- ▶ The **FLASHBACK_ARCHIVE_MIGRATE** enables migration of Flashback Data Archive enabled tables from a database (where package exists) to databases that supports it (21c)
- ▶ **Oracle Database AutoUpgrade** enables you to:
 - Upgrade: **One or many DBs with One Command** (using a single configuration file)
 - Performs: Pre-upgrade tasks, auto fix-ups, DB upgrade, post-upgrade (retry & fallback)



Rich Niemiec
@RichNiemiec



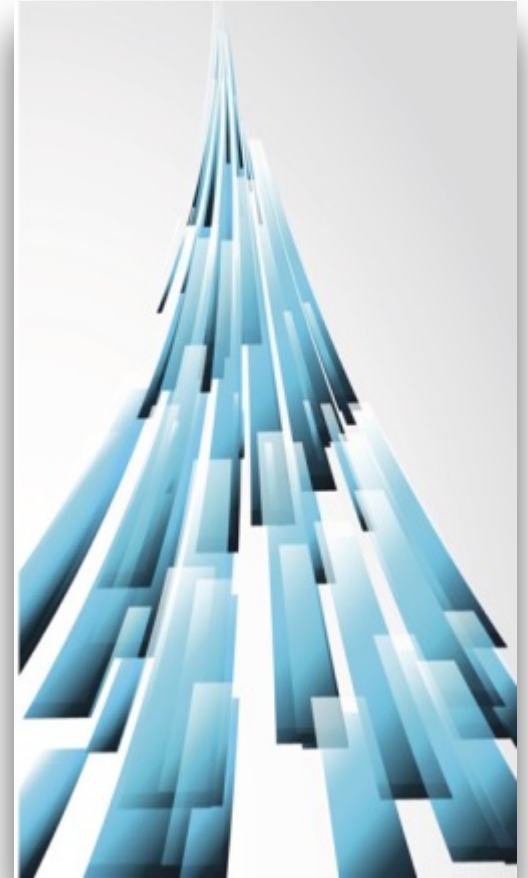
Best slide from @biju_thomas' @odtug session on #Oracle 19c/21^c. Shows features ONLY on @OracleExadata & @OracleCloud to consider. With @Cloud at #Customer, you can get it all in-house.

#database @oracleace #DBA #oracledatabase
#oracleace #mysql @dbcloudshifu #orcldb
#autonomous

And, the cool features are on Exadata (and OCI)

- 21^c: Automatic Zone Maps
- 21^c: Cache Fusion Hardening
- 21^c: Automatic In-Memory Column Store
- 19c: Automatic Indexing
- 19c: Fault Tolerant In-Memory Column Store
- 19c: High-Frequency Automatic Optimizer Statistics Collection
- 19c: Automatic SQL Plan Management
- 19c: High-Frequency SQL Plan Management Evolve Advisor Task
- 19c: Real-Time Statistics
- 19c: SQL Quarantine
- 18c: CDB Fleet Management
- 18c: PDB Snapshot Carousel
- 18c: Refreshable PDB switchover
- 18c: Oracle Data Guard-Automatic Correction of Non-logged Blocks at a Data Guard Standby Database
- 18c: In-Memory Column Store support for External Tables
- 18c: Memoptimized Rowstore
- 18c: Keystore for Each Pluggable Database

21^c









21^c ↑

Access the Oracle Docs Now at:
docs.oracle.com/en/database/oracle/oracle-database/21 (See @richniemiec for link)

Home / Database / Oracle / Oracle Database / Release 21

Oracle Database 21c

What's New

- 
Database Features and Licensing App
Use the Database Features and Licensing app to view feature availability across Oracle Database releases and to see what features are new in Oracle Database 21c.
- 
Learning Database New Features
See *Learning Database New Features* for details and practices for new features.
- 
Interactive Architecture Diagram
Use the Interactive Architecture Diagram to take a visual tour of Oracle Database architecture and technology.
- 
LiveLabs
Explore the LiveLabs workshops for Oracle Database 21c offering hands-on labs directly accessible on the Oracle Cloud.

Get Started
What's New
Install and Upgrade
Administration
Development
Security
Performance
Clustering
High Availability
Data Warehousing
Spatial and Graph
Distributed Data
Machine Learning
REST API Reference

Release Availability	<input type="radio"/> 11.2	<input type="radio"/> 12.1	<input type="radio"/> 12.2	<input type="radio"/> 18c	<input type="radio"/> 19c	<input checked="" type="radio"/> 21c
Parent Feature	Database In-Memory Base Level					

Oracle Docs:

docs.oracle.com/en/database/oracle/oracle-database/21 (See @richniemiec for link)

21c ↑

Features and Licensing

Want to learn more about the latest 21c new features and practice with them? We recommend our...

Features | Licensing

Search:

Focus Area: All Focus Areas

Version:
 11.2 12.1
 12.2 18c
 19c 21c
 New features only

Feature

Big Data and Performance Enhancements for In-Memory Tables
 Performance → In-Memory Column Store
 In-Memory external tables add support for ORACLE_HIVE and ORACLE...

Database In-Memory Base Level
 Performance → In-Memory Column Store
 Database In-Memory is an option to Enterprise Edition. Database In-M...

Database In-Memory External Table Enhancement
 Performance → In-Memory Column Store
 For a partitioned or hybrid external table the INMEMORY clause is sup...

Database In-Memory Support for External Tables
 Performance → In-Memory Column Store
 The In-Memory Column Store supports population of external tables, e...

Feature Details

Database In-Memory Base Level
 Performance → In-Memory Column Store

Database In-Memory is an option to Enterprise Edition. Database In-Memory now has a new Base Level feature. This allows the use of Database In-Memory with up to a 16GB column store without triggering any license tracking.

Business Benefit: This feature allows you to use Database In-Memory without having to license the option. The column store is limited to 16GB when using the Base Level. This helps to show the value of Database In-Memory without having to worry about licensing issues.

Release Availability:
 11.2 12.1 12.2 18c 19c 21c

Parent Feature: **Database In-Memory Base Level**

Available On:
 Enterprise Edition
 Exadata
 Database Cloud Service Enterprise Edition - Extreme Performance
 Exadata Cloud Service
 Oracle Database Appliance

Notes:
 Allows you to experiment with Oracle Database In-Memory features without purchasing the Oracle Database In-Memory option. The following restrictions apply:

- The size of the In-Memory area (INMEMORY_SIZE) cannot exceed 16 GB for a CDB. In an Oracle RAC environment, the size is limited to 16 GB for each instance.

Access the Oracle Docs Now at:

docs.oracle.com/en/database/oracle/oracle-database/21 (See @richniemiec for link)



Features and Licensing

Want to learn more about the latest 21c new features and practice with them? We recommend our [new learning guide](#).

Features **Licensing**

Search

Offerings

- Standard Edition 2
- Enterprise Edition
- Exadata
- Oracle Database Appliance
- Database Cloud Service Standard Edition
- Database Cloud Service Enterprise Edition
- Database Cloud Service Enterprise Edition - High Performance
- Database Cloud Service Enterprise Edition - Extreme Performance
- Exadata Cloud Service

Feature, Option, or Pack	SE2	EE	EXA	DBCS EE	EXA CS
Automatic Data Optimization VLDB, Data Warehousing, and Business Intelligence Feature	✗	✓	✓	✗	✓
Database In-Memory Base Level Performance Feature	✗	✓	✓	✗	✓
Fault Tolerant In-Memory Column Store Performance Feature	✗	✗	✓	✗	✓
Heat Map VLDB, Data Warehousing, and Business Intelligence Feature	✗	✓	✓	✗	✓
In-Memory Aggregation Performance Feature	✗	✓	✓	✗	✓
In-Memory Column Cache on Storage Servers Performance Feature	✗	✗	✓	✗	✓

In-Database JavaScript (21c) & Stored Logic (23ai)*

21^c ↑

In-Database JavaScript Runs on Embedded Graal Multilingual Engine (MLE)

- Run data processing JavaScript inside the Oracle Database, **where the data resides**
 - Eliminates expensive network round-trips
 - JavaScript data types are automatically mapped to Oracle Database data types and vice versa
- Leverage open-source JavaScript libraries
- Easily execute SQL from JavaScript code
- Enables developers to work efficiently in **modern programming languages**

21^c



JavaScript Stored Logic **23^{ai}**

Allows developers to create stored procedures using JavaScript in the database. This functionality also allows developers to leverage the huge number of JavaScript libraries.

JavaScript Execution using DBMS_MLE

Application Development → Java

21^c

The DBMS_MLE package allows users to execute JavaScript code inside the Oracle Database and exchange data seamlessly between PL/SQL and JavaScript. The JavaScript code itself can execute PL/SQL and SQL through built-in JavaScript modules. JavaScript data types are automatically mapped to Oracle Database data types and vice versa.

Business Benefit: With the DBMS_MLE package developers can write their data processing logic in JavaScript. JavaScript is a widely-used and popular programming language that can now also be used for writing programs that need to execute close to the data.

Release Availability



Available On

ALL OFFERINGS

Binary JSON Data Type in the Database (Doc DB)

21°C ↑

Binary JSON Data Type in Oracle Database

- Freely mix or join JSON and relational data
- Transparent scale-out with full **ACID** transactions
- Index any **JSON element** for fast OLTP
- **Declarative Parallel SQL** analytics across all formats
- Run **complex joins** across multiple JSON documents and collections
- Custom application code not needed to accomplish basic data management tasks

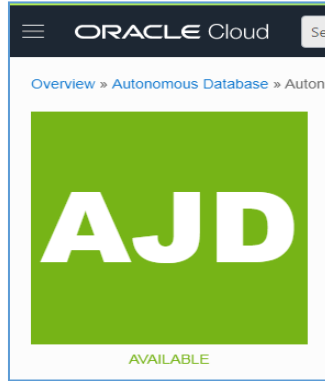


Binary JSON Data Type Enables
SCANS UP TO 10X FASTER | UPDATES UP TO 4X FASTER

AJD (Autonomous JSON DB):

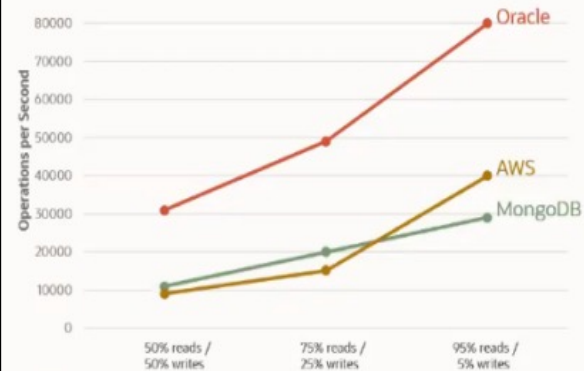
Half the Cost
Twice as Fast

19[°] ORACLE Database

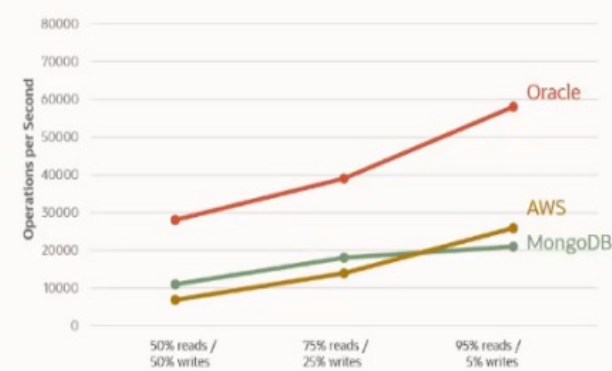


JSON Performance in Oracle Database 2X Faster than MongoDB and AWS DocumentDB

YCSB-4M Documents



YCSB-81M Documents



Industry-standard Yahoo Cloud Serving Benchmark (YCSB)
Autonomous JSON Database with 8 OCPUs compared to: MongoDB Atlas on M60, AWS DocumentDB on R4.4xlarge
Source: <https://www.mongodb.com/atlas-vs-amazon-documentdb/performance> as of 8/12/2020



Ensure you use Just Announced Oracle 21c Features

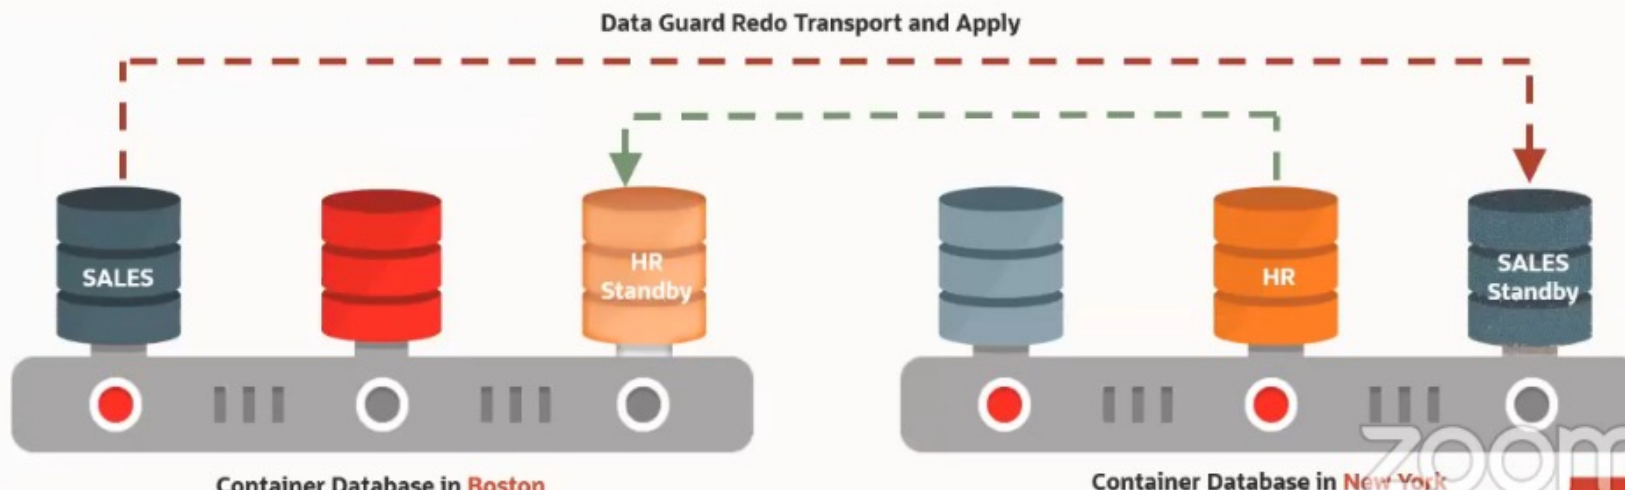
21^c ↑

Multitenant Data Guard

Disaster Protection At The Pluggable Database (PDB) Level

- Protection on PDB or CDB level using real-time Apply
- No need to fail over a full Container Database!
- Switchover and failover capabilities with the broker on a single PDB

21^c



Copyright © 2020, Oracle and/or its affiliates. All rights reserved.

23^{ai}



Read-Only Per-PDB Standbys

Per-PDB Data Guard now supports the PDBs being opened Read-Only.

This further increases the flexibility of solution supporting the offloading of reporting on the standby.

Ensure you use Just Announced Oracle 21c Features

21^c

Gradual Password Rollover

Enforce Application Security Compliance Without Incurring Downtime

- Set new password limit `PASSWORD_ROLLOVER_TIME` in user profile
- Create new service account password in the database
 - Both old and new passwords will be valid for a period of time
- Passwords will be updated on the application side
- Old passwords will expire – leaving the new password as the only valid credential



Favorite New 21c Features*

(*Andy Mendelsohn / Jenny Tsai-Smith)

NEW in Oracle Database 21c

More innovations for developers & architects, analysts & data scientists, DBAs

21c

Blockchain Tables

In-Database JavaScript

New ML Algorithms

In-Memory Enhancements

Multitenant Enhancements

Gradual Password Rollover

Native JSON Type

SQL Macros

AutoML

Better Graph Analytics

Persistent Memory

Easier Sharding

More details on new features at <https://bit.ly/DB21cNew>

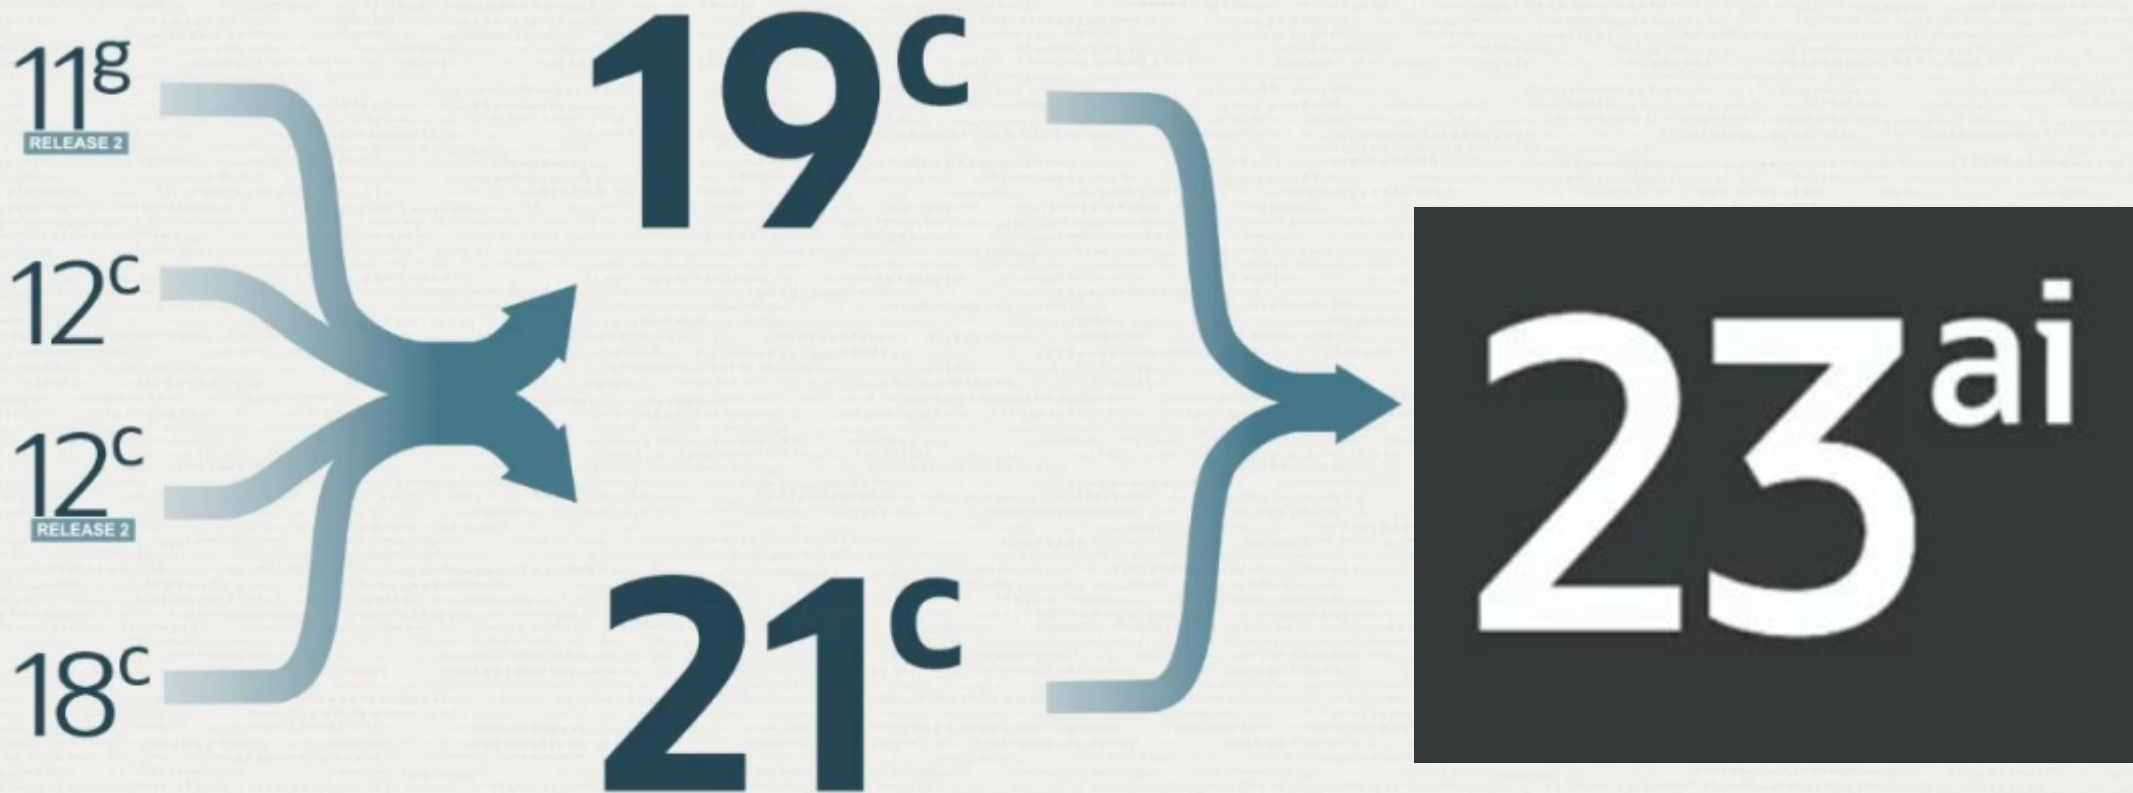
21c

21c

Get to 23ai ONLY from 19c or 21c!

23^{ai} ↑

Upgrade Path to Oracle Database 23c



Recently Updated Slightly*...

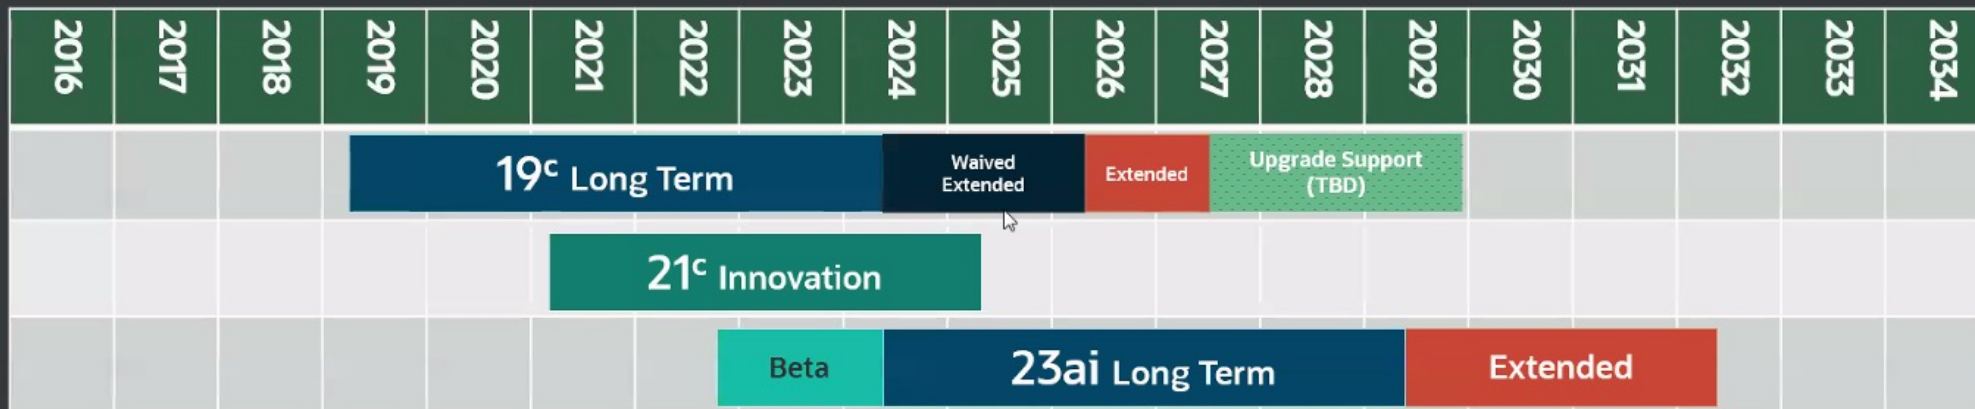
*Thanks Jenny Tsai-Smith, Oracle



Oracle Waived 19c Extended Support for OVER One Year!

Projected Database Release and Support Timeline

Subject to change



- Innovation Release - 2 years of Premier Support, and no Extended Support
- Long Term Release - 5 years of Premier Support, and 3 years of Extended Support
- Always refer to MOS Note: Release Schedule of Current Database Releases (Doc ID 742060.1)

23ai - Try it Now with these QR Codes



Oracle
Database Free

oracle.com/database/free/



Oracle
Autonomous Database Free

oracle.com/cloud/free/



Oracle
LiveLabs

oracle.com/livelabs/



23ai - Availability!



Products/Services in the May 2nd Announcement Cloud First, Developer First

In OCI

- **Oracle Database 23ai** on:
Exadata Cloud@Customer*
Exadata Database Service*
Base Database Service
- OCI GoldenGate 23ai

In Azure

- **Oracle Database 23ai** on
Exadata Database Service*

For Developers

- Always Free Autonomous **Database 23ai**
- Autonomous **Database 23ai** Free Container Image
- **Oracle Database 23ai FREE**

On-premises

- Oracle GoldenGate 23ai
- Exadata System Software 24ai for Exadata Database Machine

Release date for Oracle Database 23ai EE and SE2 **on-premises** will be announced in June

* Running on Exadata System Software 24ai



All the 21c Features (200+) & 23ai Features (300+)

23^{ai} ↑

Oracle Database 23c is the sum of...

21c +

All the features from the 21c
Innovation Release

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, timing, and pricing of any features or functionality described for Oracle's products may change and remains at the sole discretion of Oracle Corporation.

23^{ai}

300+ new features and
enhancements

Just be the DBA & Ensure you use Key Oracle 23ai Features



- ▶ Property Graphs, Okafka, JSON Relational Duality, True Cache, Auto SQL Repair
- ▶ Tuning Tips in Explain Plan & Function-based Index for Arithmetic Expressions
- ▶ Drop/Create with IF EXISTS & IF NOT EXISTS
- ▶ Creating SQL Domains (also Built-In SQL Domains) & Text Domain Indexes
- ▶ DB_DEVELOPER_ROLE Privileges
- ▶ Improved Error Messages & Group by Columns Alias/Position
- ▶ Sagas & Update tables with Direct Join
- ▶ Improved Machine Learning
- ▶ Flashback Time Travel & Enforce “digital shredding” (use with PDBs between versions)
- ▶ Lock free Reservations & ACID transaction Speed

JSON Relational Duality is a KEY Feature!

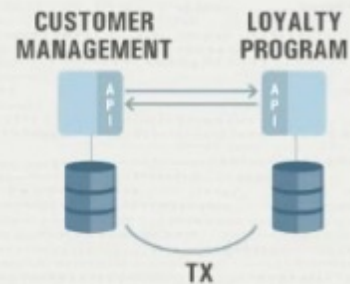
23^{ai} ↑

Oracle Database 23^{ai} For The Developer



JSON Relational Duality

Developers benefit from the strengths of both relational and document models. Data is held once but can be updated with either approach.



Microservice Support

Alongside Oracle's already comprehensive support for microservices, new functionality makes it simpler to implement cross service transactions.



Operational Property Graph

Developers can now build property graph applications directly in the Oracle Database, utilizing its industry leading security, high availability and performance capabilities.



Many Datatype and SQL Enhancements

- Boolean Datatype
- Direct Joins for Update
- SELECT without FROM
- Group By Alias
- PL/SQL Associative Arrays
- Unicode-14 Support

JSON Relational Duality is a KEY Feature!



JSON Relational Duality

Developers benefit from the strengths of both relational and document models.

Data is held once but can be updated with either approach.

```
Oracle SQL Developer : Oracle Base Database
Oracle Base Database Service - F1 DRIVER_RACE_MAP 0.34 seconds
Worksheet Query Builder
1 -- Create Duality Views
2 CREATE OR REPLACE JSON RELATIONAL DUALITY VIEW race_dv AS
3 SELECT JSON ('raceId' : r.race_id,
4             'name' : r.name,
5             'laps' : r.laps WITH NOUPDATE,
6             'date' : r.race_date,
7             'podium' : r.podium WITH NOCHECK,
8             'result' :
9             [ SELECT JSON ('driverRaceMapId' : drm.driver_race_map_id,
10                          'position' : drm.position,
11                          'driverInfo' :
12                          (SELECT JSON ('driverId' : d.driver_id,
13                                       'name' : d.name)
14                           FROM driver d WITH NOINSERT UPDATE NODELETE
15                           WHERE d.driver_id = drm.driver_id)
16                          FROM driver_race_map drm WITH INSERT UPDATE
17             ]
18 FROM driver_race_map drm WITH INSERT UPDATE
19 FROM driver_race_map drm WITH INSERT UPDATE
20 FROM driver_race_map drm WITH INSERT UPDATE
21 FROM driver_race_map drm WITH INSERT UPDATE
22 FROM driver_race_map drm WITH INSERT UPDATE
23 FROM driver_race_map drm WITH INSERT UPDATE
24 FROM driver_race_map drm WITH INSERT UPDATE
25 FROM driver_race_map drm WITH INSERT UPDATE
26 FROM driver_race_map drm WITH INSERT UPDATE
```

```
Worksheet Query Builder
1 -- Result from the latest race
2 SELECT
3     f1.race.name,
4     f1.driver.name AS name1,
5     f1.race.race_date,
6     f1.driver_race_map.position,
7     race_points(f1.driver_race_map.position) AS awarded_points
8 FROM
9     f1.race
10 INNER JOIN f1.driver_race_map ON f1.driver_race_map.race_id = f1.race.race_id
11 INNER JOIN f1.driver ON f1.driver.driver_id = f1.driver_race_map.driver_id
12 WHERE
13     f1.race.race_id = (select max(race_id) from driver_race_map)
14 ORDER BY
15     f1.driver_race_map.position;
```

Query Result x

SQL | All Rows Fetched: 20 in 0.052 seconds

	NAME	NAME1	RACE_DATE	POSITION	AWARDED_POINTS
1	Japanese Grand Prix	Max Verstappen	09-OCT-22	1	25
2	Japanese Grand Prix	Sergio Pérez	09-OCT-22	2	18
3	Japanese Grand Prix	Charles Leclerc	09-OCT-22	3	15

<https://twitter.com/thatjeffsmith>
(below from presentation)

<https://oracle-base.com/articles/23/articles-23>

```
SQL> INSERT INTO team_dv VALUES ('{"teamId" : 301,
2           "name"      : "Red Bull",
3           "points"   : 0,
4           "driver"   : [ {"driverId" : 101,
5                           "name"     : "Max Verstappen",
6                           "points"   : 0},
7                           {"driverId" : 102,
8                           "name"     : "Sergio Perez",
9*                          "points"   : 0} ]}');
```

1 row inserted.

```
SQL> select * from driver;
```

DRIVER_ID	NAME	POINTS	TEAM_ID
101	Max Verstappen	0	301
102	Sergio Perez	0	301

```
SQL> _
```

Data can be transparently accessed & updated as JSON documents or relational tables. 89



Restrictions for JSON-Relational Duality Views

2.1 Restrictions for JSON-Relational Duality Views

The following are restrictions for JSON-relational duality views in Oracle Database 23c.

- The following column data types cannot be used with duality views:
 - `TIMESTAMP WITH LOCAL TIME ZONE`
 - `BINARY_DOUBLE`
 - `BFILE`
 - `ROWID`
 - `UROWID`
 - `XMLTYPE`
 - `ANYTYPE`
 - `ANYDATA`
 - `ANYDATASET`
 - `HTTPURITYPE`
 - `XDBURITYPE`
 - `DBURITYPE`
 - `SDO_POINT_TYPE`
 - `SDO_ELEM_INFO_ARRAY`
 - `SDO_ORDINATE_ARRAY`
 - `SDO_GEOMETRY`
 - `SDO_TOPO_GEOMETRY`
 - Object types (Abstract Data Types)
- Duality views cannot be created on:
 - System-partitioned tables
 - Sharded tables
 - Views, materialized views, or editioning views

- External tables
- Hybrid partitioned tables
- Global or private temporary tables
- Remote tables (for example, tables over database links)
- Updates of duality views across database links are not supported.
- The `MERGE` SQL statement is not supported with duality views.
- You cannot create a functional index, JSON search index, or JSON multivalue index on the `DATA` column of a duality view.
- The use of a JSON search index on the column of an underlying table is not supported.
- The array of JSON documents produced by the `JSON_ARRAYAGG` function inside the duality view will be ordered by the primary key value of those JSON documents.
- Virtual private database (VPD) and Oracle Real Application Security (RAS) on duality views are not supported.
- VPD on underlying tables are supported only if all statements (`INSERT`, `UPDATE`, `DELETE`, or `SELECT`) are included in the policy. However, when all statement types are not included in the VPD policy, there is no error returned but DML and query results may be unexpected or may fail.
- Transparent Sensitive Data Protection is not supported with duality views and underlying tables.
- Table columns of a duality view cannot be redacted if the redacted columns are part of the ETAG.
- JSON-relational duality views cannot be used with tables that have a primary key of type `BINARY_DOUBLE`, `BINARY_FLOAT`, `INTERVAL DAY TO SECOND`, or `INTERVAL YEAR TO MONTH`. All DML operations on such duality views fail.
- Because JSON-relational duality views rely on JSON type, a 19c or earlier SQL*Plus client cannot be used for queries, DML or other SQL operations that use the `DATA` column which is of data type `JSON`.
- Duality views cannot be created on base tables with textual JSON columns (`VC2/CLOB/BLOB` with `IS-JSON` constraint) in the base table.
- The same subquery cannot be used multiple times to define nested JSON structures (object or array) in the same parent object.
- Fine-grained auditing policies are not supported with duality views.
- DML error logging is not supported with duality views.
- The primary key column needs to be listed before specifying the nest directive for the non-primary key columns.
- Duality view names should use ASCII characters. The use of some non-ASCII characters with non-UTF8 database characters can fail in some operations.
- When two JSON sub-objects in a duality view refer to the same base table row, any `INSERT` or `UPDATE` to those sub-objects are currently not supported. The DML behavior is, therefore, undefined and may or may not return an error depending on the input values.
- JSON path expressions with the `type()` item method are not supported when querying the `DATA` column of a duality view.

Operational Property Graphs are a KEY Feature!

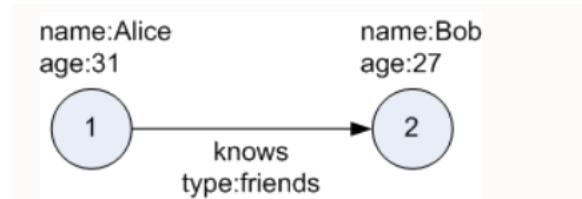
23^{ai} ↑



Operational Property Graph

Developers can now build property graph applications directly in the Oracle Database, utilizing its industry leading security, high availability and performance capabilities.

- ▶ In 23ai: Build graph data on top of tables in the database.
- ▶ Support for the **SQL Property Graph Query Language**
- ▶ Data entities as **vertices** & relationships as **edges** in graph.
- ▶ Example: *Bank customer **accounts** can be **vertices**, and **cash transfer relationships** between them can be **edges**.*
- ▶ *Run graph analytics algorithms like PageRank to measure **importance of data entities based on the relationships***

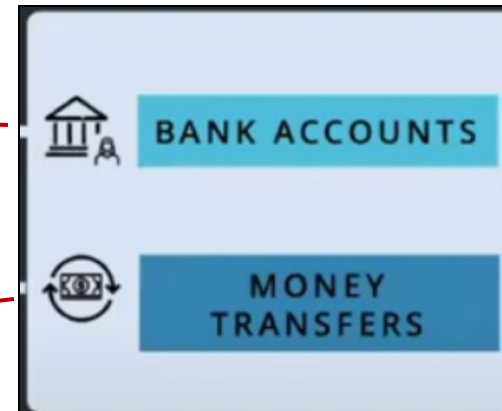


Use this to quickly graph/find:

- Influencers in a Social Network
- Trends of Customers

Create Property Graph

```
SQL> CREATE PROPERTY GRAPH bank_sql_pg
 2  VERTEX TABLES (
 3    bank_accounts ←
 4    KEY (id)
 5    LABEL account
 6    PROPERTIES ALL COLUMNS
 7  )
 8  EDGE TABLES (
 9    bank_txns ←
10   KEY (txn_id)
11   SOURCE KEY (from_acct_id) REFERENCES bank_accounts (id)
12   DESTINATION KEY (to_acct_id) REFERENCES bank_accounts (id)
13   LABEL transfer
14   PROPERTIES ALL COLUMNS
15* );
```



Copy

Property created.

On execution, the `bank_sql_pg` graph is created in the database. The graph is made up of one vertex graph element table (`bank_accounts`) and one edge graph element table (`bank_txns`).

Property Graphs

```
SQL> SELECT * FROM GRAPH_TABLE (bank_sql_pg
2   MATCH
3   (a IS account WHERE a.id = 816) -[e IS transfer]-> (b IS account)
4   COLUMNS (a.id AS acc_a, e.amount AS amount, b.id AS acc_b)
5* );
```

Copy

ACC_A	AMOUNT	ACC_B
816	4713	287
816	8001	590
816	4186	934
816	3718	289
816	4039	812

See [SQL GRAPH_TABLE Queries](#) to understand more about GRAPH_TABLE queries..

FYI:

```
SQL> DROP PROPERTY GRAPH bank_sql_pg;
```

Copy

```
Property dropped.
```

Property Graphs & Oracle Graph Visualization



Figure 3-2 Visualizing GRAPH_TABLE Query

ORACLE[®] Graph Visualization graphuser ▾

PGQL SQL/PGQ

```

SQL/PGQ Query
1 SELECT id a, id e, id b
2 FROM GRAPH TABLE ( bank sql_pg
3 MATCH (a IS account WHERE a.id=816) -[e IS transfer]-> (b IS account)
4 COLUMNS (vertex_id(a) AS id_a, edge_id(e) AS id_e, vertex_id(b) AS id_b )
5 )
6
7

```

Parallelism: 0 ▾ ▲

Graphs are a powerful way to query connections and relationships between data

For example, to discover indirect money movements from bank account 'B' to bank account 'E'

Just a FEW little Nice SQL Features...

- ▶ SELECT Without FROM Clause:

```
SQL> SELECT SYSDATE;
```

- ▶ New Dev Role:

```
exec dbms_developer_admin.grant_privs('RICH');
```


- ▶ Update via Join:

```
SQL> Update emp e SET e.salary=e.salary*2
      FROM dept d
      WHERE e.dept_id = d.dept_id
      AND d.name = 'Development';
```

- ▶ Usage of **Column Alias** in **GROUP BY** and **HAVING**

- ▶ Better Error Messages why Statement Failed to Execute:

```
SELECT foo FROM bar;
ORA-00942: table or view does not exist
ORA-00942: table or view ("GERALD"."BAR") does not exist
```



SQL Domains
Allows developers to define domain level datatypes once, and share their constraint checks across multiple tables.



Many Datatype and SQL Enhancements

- Boolean Datatype
- Direct Joins for Update
- SELECT without FROM
- Group By Alias
- PL/SQL Associative Arrays
- Unicode-14 Support

Update table(s) with condition in a Direct Join



Increase Tech Salaries by 20%:

```
update emp a
set    a.salary = a.salary*1.2
from  dept d
where a.deptno = d.dept_no
and   d.dname = 'Technology';
```

Selecting an expression without FROM DUAL

Example: Select 'Hello world' as txt, $7 * 6$ as num, sysdate as now;

► [@GeraldVenzl](#) via [@phsalvisberg](#) ([#DOAG2022](#) Keynote – screen shot below)

SELECT without FROM

- Expressions aren't selected from a table

Query	Result
<pre>SELECT SYSDATE;</pre>	<pre>SYSDATE ----- 2022-09-21 22:18:52</pre>
<pre>SELECT 2*3;</pre>	<pre>2*3 ----- 6</pre>
<pre>SELECT my_func();</pre>	<pre>MY_FUNC ----- Hello DOAG!</pre>

GROUP BY on expression aliases

YES!!

23^{ai} ↑

GROUP BY can use aliases for expressions (such as the result of a PL/SQL function):

```
select count(*), fnc(a, b) xyz
```

```
from t
```

```
group by xyz
```

```
having xyz > 100 ;
```

- ▶ Prior to Oracle 23, the previous statement would have thrown a [ORA-00904: invalid identifier](#) error message.

GROUP BY on expression position number

YES!!

23^{ai}

- ▶ If the **group_by_position_enabled** initialization parameter is set to **true**, it's also possible to **group by a position number** (rather than an expression alias):

```
SQL> alter session set group_by_position_enabled = true;
```

Session altered.

```
select category_id, sum(val)
```

```
from t
```

```
group by 1;
```

Oracle SQL: Faster, ACID, Consistent, PQ, PL/SQL, ML+



Better SQL Than Anyone Else

Oracle Database 23^{ai}

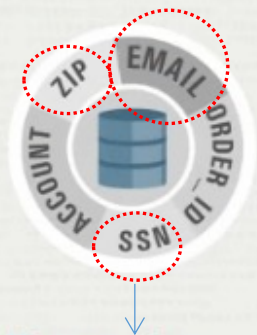
- Approximate Functions
- Automatic Indexing, Partitioning, Materialized Views
- Flashback Query
- SQL For JSON
- Window Functions
- Real Time Materialized Views
- Results Cache
- Polymorphic Tables
- Readers don't block Writers
- Multi Model
- PL/SQL
- SQL Model Clause
- SQL Pattern Matching
- External Tables
- Parallel Query
- SQL/PQO
- SQL Macros
- Analytical Views



Lock-free Column Reservations help Transactions & ML

23^{ai} ↑

Oracle Database 23^{ai} For The Developer



SQL Domains

Allows developers to define domain level datatypes once, and share their constraint checks across multiple tables.



OKafka

Kafka applications can now run directly against the Oracle Database with minimal code changes leveraging high performance Transaction Even Queues (TEQ).



JavaScript Stored Logic

Allows developers to create stored procedures using JavaScript in the database. This functionality also allows developers to leverage the huge number of JavaScript libraries.

Lock-free Column Value Reservations



Escrow Locking

Allows developers to specify the thresholds at which locks are applied to rows. This can help in situations where long-held locks can prevent the processing of additional transactions.

SQL Domains

CREATE DOMAIN Example – HourlyWages



```
CREATE DOMAIN HourlyWages AS Number DEFAULT 15 ON NULL  
CONSTRAINT MinimalWage CHECK (HourlyWages > = 7 and HourlyWages <=1000)  
ENABLE DISPLAY TO_CHAR(HourlyWages, '$999.99')  
ORDER ( -1*HourlyWages ) ANNOTATIONS (properties '{"Purpose": "Wages",  
"Applicability": "USA", "Industry": {"Sales", "Manufacturing"} }')
```

- ▶ **CREATE TABLE employees1** (name VARCHAR2(100), id NUMBER, wage NUMBER **DOMAIN HourlyWages**);
- ▶ **CREATE TABLE employees2** (name VARCHAR2(100), id NUMBER, wage NUMBER **DOMAIN HourlyWages**);

CREATE DOMAIN Examples - Email Domain

You can also define an email domain as:

- ▶ **CREATE DOMAIN Email AS VARCHAR2(30) DEFAULT ON NULL t_seq.NEXTVAL || '@mymail.com' CONSTRAINT EMAIL_C CHECK (REGEXP_LIKE (Email, '^(\S+)\@(\S+)\.(\S+)\$')) DISPLAY '---' || SUBSTR(Email, INSTR(Email, '@') + 1).**
- ▶ This example uses a DEFAULT ON NULL email address by referencing a sequence value and concatenating it with "@mymail.com". The DOMAIN_DISPLAY function masks out the user name and its length using a predefined '---' string, and leaves the email domain visible.

Add another constraint to Domain above when you create (more than one is ok):

- ▶ **CONSTRAINT EMAIL_MAX_LEN_C CHECK (LENGTH(Email) <=100)**

CREATE DOMAIN Examples – DayOfWeek Domain

Weekdays can be defined as follows:

```
▶ CREATE DOMAIN DayOfWeek AS CHAR(3 CHAR) CONSTRAINT DayOfWeek_C  
CHECK (UPPER(Substr(DayOfWeek, 1, 3)) IN ('MON', 'WED', 'FRI', 'SAT', 'SUN') or  
UPPER(Substr(DayOfWeek, 1, 2)) IN ('TU', 'TH')) COLLATE BINARY_CI DISPLAY  
SUBSTR(DayOfWeek, 1, 3) ORDER CASE WHEN UPPER(SUBSTR(DayOfWeek,  
1, 3)) = 'MON' THEN 0 WHEN UPPER(SUBSTR(DayOfWeek, 1, 2)) = 'TU' THEN 1  
WHEN UPPER(SUBSTR(DayOfWeek, 1, 3)) = 'WED' THEN 2 WHEN  
UPPER(SUBSTR(DayOfWeek, 1, 2)) = 'TH' THEN 3 WHEN  
UPPER(SUBSTR(DayOfWeek, 1, 3)) = 'FRI' THEN 4 WHEN  
UPPER(SUBSTR(DayOfWeek, 1, 3)) = 'SAT' THEN 5 WHEN  
UPPER(SUBSTR(DayOfWeek, 1, 3)) = 'SUN' THEN 6 ELSE 7 END.
```

Built-in Domains

Name	Allowed Values	Description
email_d	$^{\wedge}([a-zA-Z0-9!#\$\%&*+=?\^_ \{\} \sim -] + (\backslash . [A-Za-z0-9!#\$\%&*+=?\^_ \{\} \sim -] +)^*) @ (([a-zA-Z0-9] ([a-zA-Z0-9-] * [a-zA-Z0-9]) ? \backslash .) + [a-zA-Z0-9] ([a-zA-Z0-9-] * [a-zA-Z0-9]) ?) \$$	Email address Implementation: Regex
day_short_d	'MON','TUE','WED','THU','FRI','SAT','SUN'	Day of the week in short format Domain Order: based on <code>nls_territory</code> value Implementation: List
day_d	'MONDAY','TUESDAY','WEDNESDAY','THURSDAY','FRIDAY','SATURDAY','SUNDAY'	Day of the week in long format Domain Order: based on <code>nls_territory</code> value Implementation: List
month_short_d	'JAN','FEB','MAR','APR','MAY','JUN','JUL','AUG','SEP','OCT','NOV','DEC'	Month in short format Implementation: List

Built-in Domains



Name	Allowed Values	Description
month_d	'JANUARY','FEBRUARY','MARCH','APRIL','MAY','JUNE', 'JULY','AUGUST','SEPTEMBER','OCTOBER','NOVEMBER','DECEMBER'	Month in long format Implementation: List
ssn_d	^[0-9]{3}-[0-9]{2}-[0-9]{4}\$	US SSN Implementation: regex
credit_card_number_d	^([0-9]{4}-){3}[0-9]{4}\$ ^([0-9]{4})){3}[0-9]{4}\$ ^([0-9]{4}){3}[0-9]{4}\$ ^[0-9]{4}-[0-9]{6}-[0-9]{5}\$ ^[0-9]{4}[][0-9]{6}[][0-9]{5}\$ ^[0-9]{15}\$	Credit Card Number Implementation: regex

Built-in Domains

Name	Allowed Values	Description
phone_number_d	^[+]{0,1}[0-9]{1,16}\$	Phone Number Implementation: regex
mime_type_d	'application/epub+zip', 'application/gzip', 'application/java-archive', 'application/json', 'application/ld+json', 'application/msword'	MIME Types Implementation: list

Using Lock-Free Column Value Reservation

- ▶ Allow applications to *reserve part of a value in a column without locking the row*; for example, reserve part of a bank account balance or reserve an item in inventory without locking out all other operations on the bank account or item.
- ▶ **Lock-free Reservation enables concurrent transactions to proceed without being blocked on updates of heavily updated rows to improve concurrency.**
- ▶ Lock-free reservations are held on the rows instead of locking them. Lock-free Reservation verifies if the updates can succeed and **defers the updates until the transaction commit time.**
- ▶ **In microservices applications, (like trip booking services, you may have flight, hotel, and car bookings) ; the source remains locked for an extended period, potentially making it a hot resource.**

Using Lock-Free Reservation

- ▶ **Improved concurrency with reduced isolation** *while maintaining the atomicity, consistency, and durability* properties of transactions. To improve concurrency, enable data locking where only value is modified.
- ▶ ALTER TABLE [schema.]table [**add** [column_definition]]...; column_definition ::= column_name datatype **reservable** [default <value>] [CONSTRAINT constraint_name check_constraint]
- ▶ ALTER TABLE Account ADD (Balance NUMBER reservable CONSTRAINT minimum_balance CHECK (Balance >= 50));
- ▶ To change an existing column to a reservable column:
- ▶ TABLE [schema.]table [**modify** [column_definition]]...; column_definition ::= column_name **reservable** ...

Using Lock-Free Reservation

Change existing QOH column to reservable column and add a new constraint:

▶ ALTER TABLE PRODUCTS MODIFY (QOH **reservable** default 0 CONSTRAINT maxAmount CHECK (QOH <= 100));

To change a reservable column to a non-reservable column:

▶ ALTER TABLE [schema.]table [modify [column_definition]]...; column_definition::= column_name **not reservable**]

To change an existing reservable column QOH to a non-reservable column:

▶ ALTER TABLE PRODUCTS modify (QOH **not reservable**);

Conventional Locking vs. Lock-Free Reservation

▶ **Conventional Locking** (with Long-held Locks)

The following example uses traditional locking to allow a purchase of a \$25 item while maintaining a \$50 minimum balance:

- ▶ A SELECT FOR UPDATE is first issued to read and lock the balance.
- ▶ If the balance is at least 75, the item purchase is allowed.
- ▶ The UPDATE then debits the balance.
- ▶ The transaction then commits.
- ▶ An insufficient balance causes an abort.



Conventional Locking

```
CREATE TABLE Account ( ID NUMBER PRIMARY KEY, Name VARCHAR2(10), Balance NUMBER  
CONSTRAINT minimum_balance CHECK (Balance >= 50));
```

```
DECLARE current NUMBER;
```

```
BEGIN -- Read and Lock account balance
```

```
SELECT Balance INTO current FROM Account WHERE ID = 12345
```

```
FOR UPDATE;
```

```
IF current >= 75 THEN -- Sufficient funds: Perform item purchase PurchaseItem();
```

```
-- Debit account balance and commit
```

```
UPDATE Account SET Balance = Balance - 25
```

```
WHERE ID = 12345; COMMIT;
```

```
ELSE ROLLBACK; -- Insufficient funds, abort END IF; END;
```

Conventional Locking vs. Lock-Free Reservation

- ▶ **Lock-Free Reservation** (with Short-held Locks)
- ▶ The following example uses lock-free reservation to allow a purchase of a \$25 item while maintaining a \$50 minimum balance. The **reservable column constraint** allows a **reservation to be placed on a column value without locking the row.**
- ▶ The balance update **reserves \$25 without locking** the account.
- ▶ If the reservation succeeds, the item purchase is allowed to proceed.
- ▶ The **final commit locks** the account row and **applies** the balance **debit of \$25 as recorded in the reservation.**
- ▶ If the reservation fails due to insufficient funds, the update statement fails with the CHECK constraint violation.
- ▶ **Multiple transactions can Reserve as long as we stay above the Minimum Balance!**

Lock-Free Reservation

```
CREATE Table Account( ID NUMBER PRIMARY KEY, Name VARCHAR2(10), Balance
NUMBER RESERVABLE CONSTRAINT minimum_balance CHECK (Balance >= 50));
BEGIN -- Reserve 25 from account balance
UPDATE Account SET Balance = Balance - 25
WHERE ID = 12345; -- If reservation succeeds perform item purchase
PurchaseItem(); -- commit finalizes the balance update
COMMIT; -- This gets the account row lock
EXCEPTION WHEN Check_Constraint_Violated -- This indicates that reservation failed
THEN ROLLBACK;
END;
```

Reservations KEY for Sagas for Microservices

- ▶ A saga encapsulates a long running business transaction that is composed of **several independent microservices**. (like trip booking services, you may have flight, hotel, and car bookings)

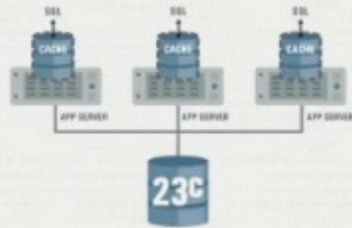
Lock-free reservations provides in-database infrastructure for transactions operating on reservable columns to:

- ▶ **Enable concurrent transactions to proceed** without being blocked on updates made to **reservable columns**
- ▶ **Issue automatic compensations for reservable updates of successful transactions in an aborted saga**

True Cache: App Tier, Edge, Regional, or Multi-Cloud Cache

23^{ai} ↑

Core Database and Performance

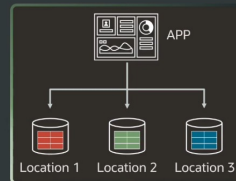


TrueCache

Provides an in-memory, high performance cache that is always consistent with the database.

This results in more responsive applications and lower load on the database server.

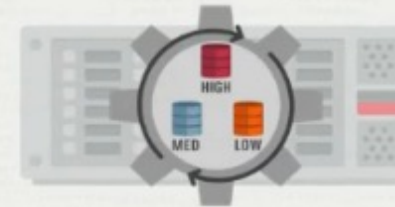
Globally Distributed Database



Single logical database
multiple physical databases
to help support scalability and
data sovereignty

Sharding Enhancements

New functionality makes it simpler to create and manage shard replicas. New sharding models also improve the distribution of data for shard keys with few unique values.



Inter-Instance Resource Management

DBAs can specify the priority of different databases running on the same server.

This enables better utilization of hardware and reduces the risk associated with server consolidation.

Priority Transactions



Automatic Transaction Abort

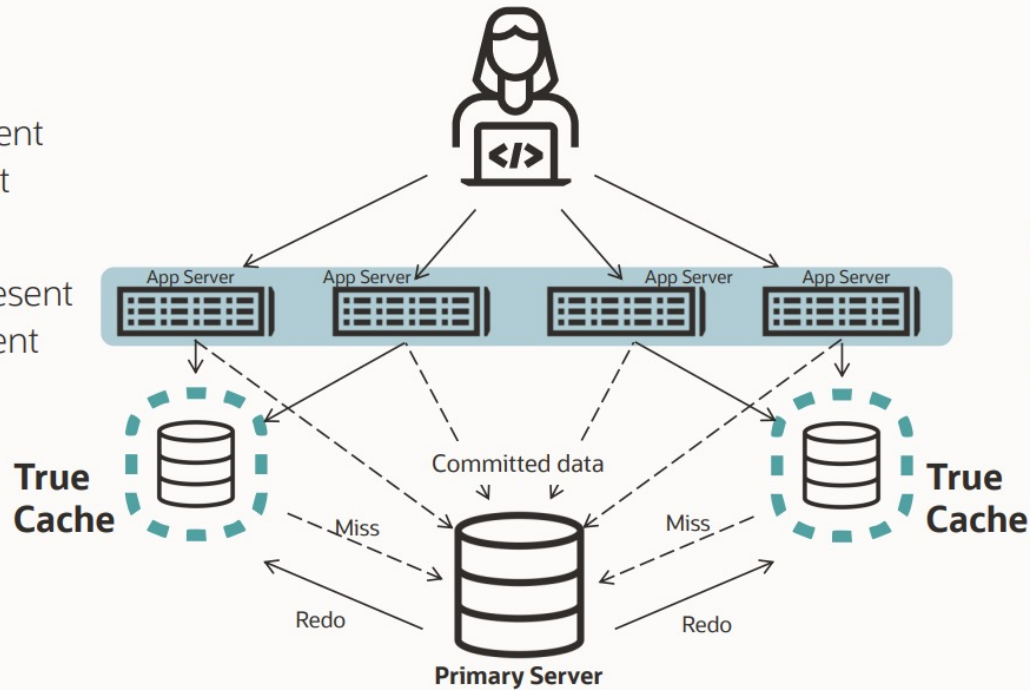
Low priority transactions that block high priority transactions can be automatically aborted. This reduces the admin burden on the DBA whilst maintaining high transaction throughput.

True Cache: App Tier, Edge, Regional, or Multi-Cloud Cache

Next | Oracle Database 23ai True Cache

Solid lines represent relatively frequent requests

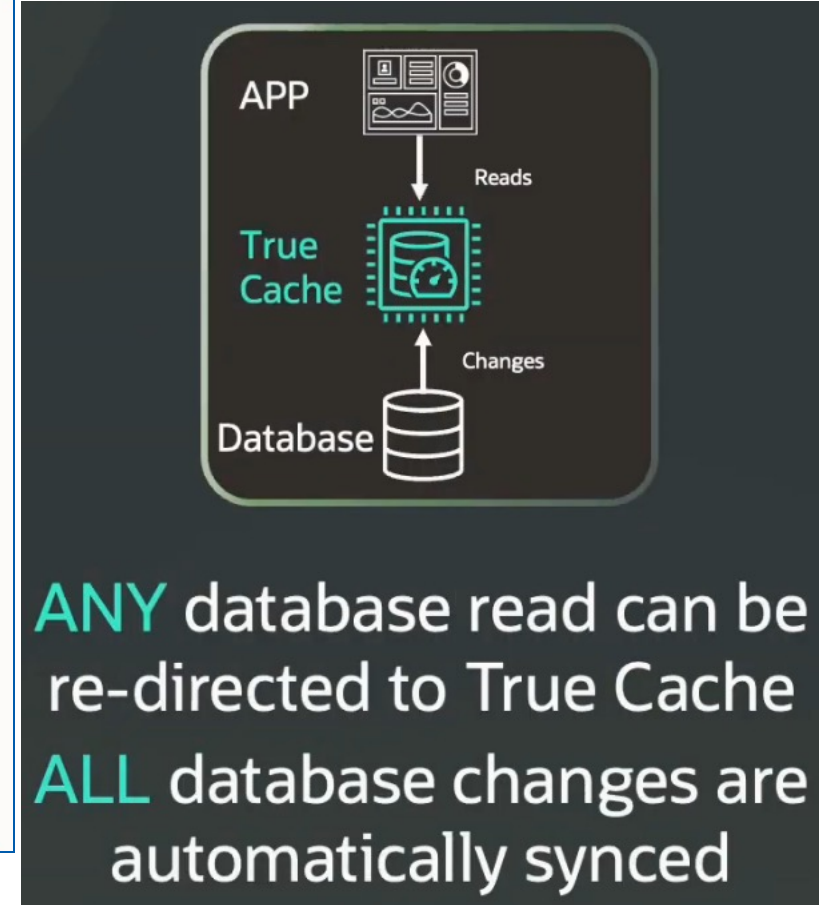
Dotted lines represent relatively infrequent requests



App connect to True Cache and perform SQL queries

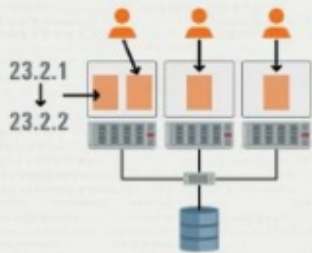
True Cache is an in-memory, consistent, and automatically managed full SQL cache

Truly transparent,
truly consistent,
mid-tier cache



Rolling Patch, Pre-Fix ORA-600 & Read-Only PDB Standby

High Availability



Dual Instance Rolling Patching

Users can now stand up two instances of a database on the same server, patching them in a rolling fashion.

Single server databases can now benefit from higher availability.

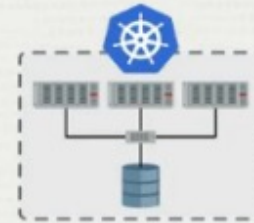
Real-time SQL Plan Management



Auto SQL Repair

When the Oracle Database encounters critical SQL error or performance regression, it will look for alternative approaches to execute the statement.

Improves the overall availability of applications.



RAC on K8s & Podman

Users choosing to run Oracle Real Application Clusters in Kubernetes can now benefit from business continuity during planned and unplanned outages.



Read-Only Per-PDB Standbys

Per-PDB Data Guard now supports the PDBs being opened Read-Only.

This further increases the flexibility of solution supporting the offloading of reporting on the standby.

Enhanced Automatic SQL Plan Management



- ▶ When there is a **Critical SQL Error** it will **try an alternative execution path** avoiding the error. **Auto SQL Repair**
- ▶ Automatic SQL plan management has been enhanced **to detect and repair SQL performance regressions more quickly.**
- ▶ SQL plan changes are **detected at parse-time and, after initial execution**, SQL performance is compared with the performance of previous SQL execution plans. If a performance degradation is detected, the plan is repaired accordingly.
- ▶ Similar to execution plans in previous version that are changed when stats prove to be not up to date (execution plan changes as stats prove to be different).

Schema Level Privileges & New Developer Role

23^{ai} ↑

Security

```
GRANT SELECT ANY TABLE  
ON SCHEMA SCOTT TO HR
```

Schema Level Privileges

System privileges can now be granted at the schema level.

Simplifies the privilege management process and as a result makes it easy to secure databases.



Developer Role

A new role allows administrators to quickly assign developers only the privileges they need to design, build and deploy applications for the Oracle Database.



TLS 1.3 Support

New Oracle Database support for the latest version of TLS. This improves the performance of TLS handshakes making connections faster.



Azure AD Oauth2 Integration

New functionality enables single sign-on to Oracle Database service instances from Microsoft Azure Cloud.

FYI

23^{ai} ↑

Schema Level Privileges (even FUTURE ones!)

Demo of Schema Level Privileges

Pre Oracle Database23c

```
GRANT SELECT ON TPCDSLIKE.CATALOG_PAGES TO BOB
GRANT SELECT ON TPCDSLIKE.CATALOG_RETURNS TO BOB
GRANT SELECT ON TPCDSLIKE.CATALOG_SALES TO BOB
GRANT SELECT ON TPCDSLIKE.CUSTOMER TO BOB
GRANT SELECT ON TPCDSLIKE.CUSTOMER_ADDRESS TO BOB
GRANT SELECT ON TPCDSLIKE.CUSTOMER_DEMOGRAPHICS TO BOB
GRANT SELECT ON TPCDSLIKE.DATE_DIM TO BOB
GRANT SELECT ON TPCDSLIKE.HOUSEHOLD_DEMOGRAPHICS TO BOB
GRANT SELECT ON TPCDSLIKE.INCOME_BAND TO BOB
GRANT SELECT ON TPCDSLIKE.INVENTORY TO BOB
GRANT SELECT ON TPCDSLIKE.ITEM TO BOB
GRANT SELECT ON TPCDSLIKE.PROMOTION TO BOB
GRANT SELECT ON TPCDSLIKE.REASON TO BOB
GRANT SELECT ON TPCDSLIKE.SHIP_
```

Oracle Database23c

```
GRANT SELECT ANY TABLE ON
SCHEMA TPCDSLIKE TO BOB
```

Schema Level Privileges

System privileges can now be granted at the schema level.

Simplifies the privilege management process and as a result makes it easy to secure databases.

Grant db_developer_role to *username*;



```
SQL> grant db_developer_role to rich identified by  
manager7;  
grant succeeded.
```

```
SQL> connect rich/manager7;  
connected.
```

```
SQL> select privilege  
from session_privs  
order by privilege;
```

PRIVILEGE

CREATE ANALYTIC VIEW
CREATE ATTRIBUTE DIMENSION
CREATE CUBE
CREATE CUBE BUILD PROCESS
CREATE CUBE DIMENSION
CREATE DIMENSION
CREATE DOMAIN
CREATE HIERARCHY
CREATE JOB
CREATE MATERIALIZED VIEW
CREATE MINING MODEL

CREATE MLE
CREATE PROCEDURE
CREATE SEQUENCE
CREATE SESSION
CREATE SYNONYM
CREATE TABLE
CREATE TRIGGER
CREATE TYPE
CREATE VIEW
DEBUG CONNECT SESSION
EXECUTE DYNAMIC MLE
FORCE TRANSACTION
ON COMMIT REFRESH

24 rows selected.

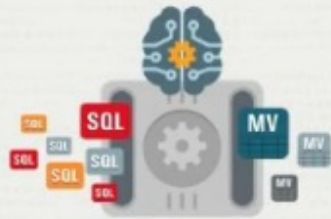
(output continued on right side)

Increased Oracle Database Password Length

- ▶ Microsoft Azure Active Directory Integration with Additional Oracle Database Environments
- ▶ This multi-cloud feature integrates authentication and authorization between Azure AD and Oracle Databases in Oracle Cloud Infrastructure and on-premises.
- ▶ Oracle Database now supports passwords up to 1024 bytes in length. In previous releases, the Oracle Database password length and the secure role password length could be up to 30 bytes.
- ▶ There are uniform password rules for all Cloud deployments, including Oracle Identity Cloud Service (IDCS) and Identity Access Management (IAM).

Improve Materialized Views & Real-Time Stats using ML

Analytics and Machine Learning



Automatic Materialized Views

Machine Learning assists in the management of materialized views to improve analytical performance.



Improved ML Algorithms

New improvements to Oracle In-Database Machine Learning algorithms make it simpler to categorize text and data whilst offering better performance and flexibility.

(Next Slides)



Enhanced Realtime Statistics using Machine Learning

A new ML model allows Oracle to predict how data will change over time.

This approach results in accurate optimizer statistics without expensive computations.



Up to 4096 Columns per Table

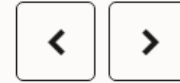
Tables now support up to 4096 columns.

This simplifies the development of applications needing large numbers of attributes such as ML and IoT.

SQL> ALTER SYSTEM SET MAX_COLUMNS=EXTENDED;



2.207 MAX_COLUMNS



MAX_COLUMNS specifies the maximum number of columns allowed in database tables and views.

Property	Description
Parameter type	String
Syntax	MAX_COLUMNS = { STANDARD EXTENDED }
Default value	STANDARD
Modifiable	No
Modifiable in a PDB	Yes
Basic	No
Oracle RAC	Multiple instances must use the same value.

When this parameter is set to `STANDARD`, the maximum number of columns allowed in a database table or view is 1000.

When this parameter is set to `EXTENDED`, the maximum number of columns allowed in a database table or view is 4096.

The `COMPATIBLE` initialization parameter must be set to `23.0.0.0` or higher to set `MAX_COLUMNS = EXTENDED`.

You can change the value of `MAX_COLUMNS` from `STANDARD` to `EXTENDED` at any time. However, you can change the value of `MAX_COLUMNS` from `EXTENDED` to `STANDARD` only when all tables and views in the database have 1000 or fewer columns.

MAX_COLUMNS

- ▶ Set to **Standard**:
1000 columns
- ▶ Set to **Extended**:
4096 columns
- ▶ **Compatible** must
be **23.0.0.0** or higher

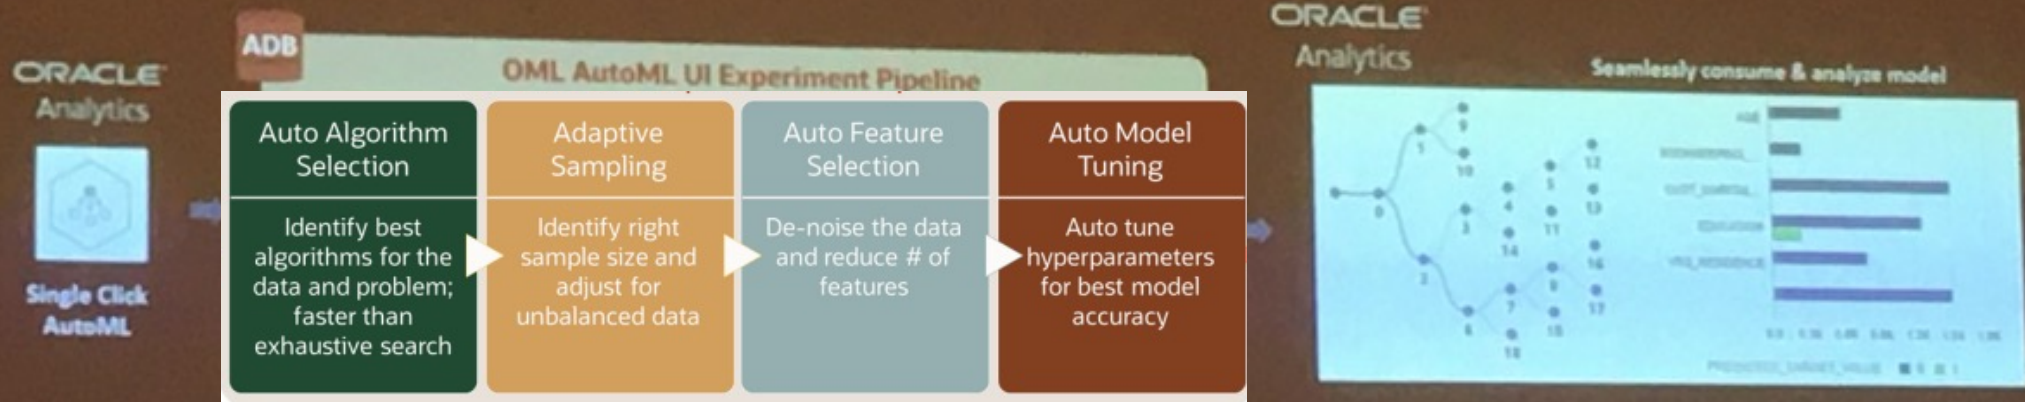
AutoML & OAC new tracking sees “exactly what happened & why”

23^{ai}

Oracle Analytics Cloud

Integrating with OML4Py AutoML on Autonomous Database - roadmap

Empower business analysts with a “single click” AutoML engine directly within the Oracle Analytics daily experience



Seamlessly consume in-database model results within Oracle Analytics

Increase business agility by making machine learning easy and efficient for non-expert users

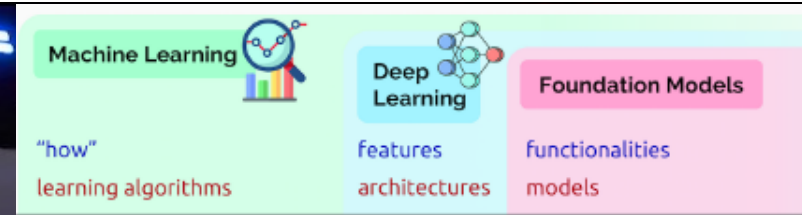
Is GenAI Most Important Ever? You're About to Find Out!

Generative Artificial Intelligence – Question?

Is Generative AI the Most Important New Computer Technology Ever? Probably!

- Countless Billions \$\$\$: Being Invested in Generative AI and Large Language Models
- Last Year: ChatGPT 4.0, Cohere, other LLMs Generated Images, Music, Computer Code ...
- Coming Soon: Self-Driving Cars, Antiviral Drug Design, Clinical Voice Digital Assistant ...

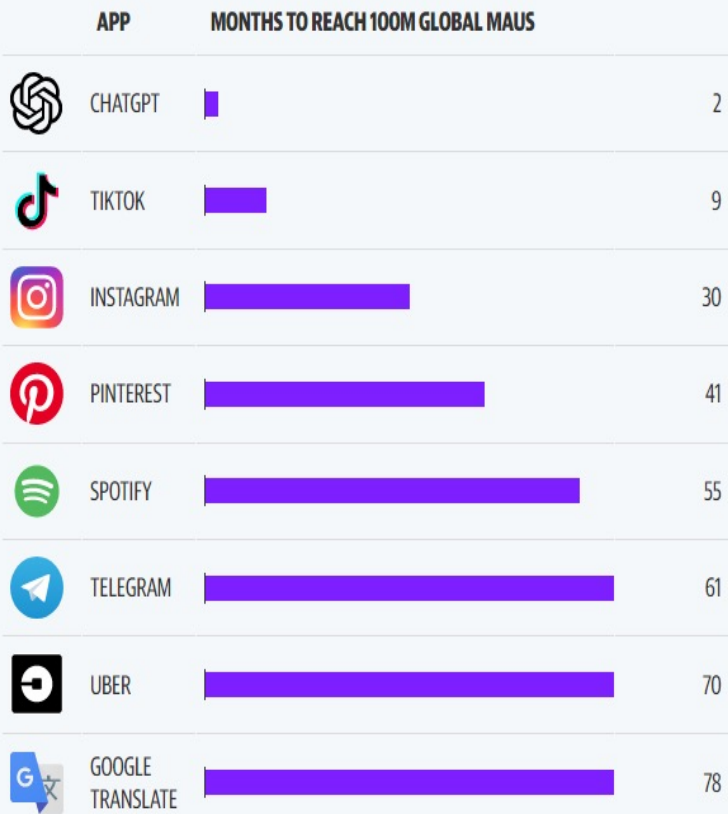
More Big Surprises !!



HOW LONG IT TOOK TOP APPS TO HIT 100M MONTHLY USERS

ChatGPT is estimated to have hit 100M users in January, 2 months after its launch.

Here's how long it took other top apps to reach that:



SOURCE: UBS

yahoo!
finance

About One Year Ago: ChatGPT 3.5 was Born ... the Baby Talks

- Developed by OpenAI: A Company Founded by Elon Musk 10 Years Ago
- Huge Artificial Neural Network: Multi-Billion Parameter Large Language Model (LLM)
- Trained on Language: Trillions of Data Elements ... Wikipedia + Public Internet
- ChatGPT 3.5 Talks: Generates Answers to Questions ... Engages in Conversations

A Big Surprise

ChatGPT's Scale Enabled it to Understand and Generate Language

Main Difference Between GPT-4 and ChatGPT

Unlike ChatGPT, GPT-4 can see and understand images, whereas the former is limited to text.

While ChatGPT had several issues upon its arrival, Microsoft claims that GPT-4 has been trained to avoid several malicious prompts.

FYI



Generative AI: GPT & ChatGPT (the CHAT changed it all)

- ▶ **Databases** usually have **Factual/Precise** Information
- ▶ **GPT = Generative Pre-trained Transformer** (could give Hallucinations)
- ▶ LLMs are trained on Wikipedia, the Internet, or your own data (R.A.G.)
- ▶ Use the LLM to turn word or question into numbers (vectors) & search
- ▶ Embeddings give semantic meaning to words...etc. in numbers (vectors)
- ▶ The Answer is numbers - turn it back into words using the LLM
- ▶ ChatGPT is a Foundation Model for LLM (Large Language Model) Chatbots.
- ▶ **Transformer model architecture** accelerated the LLM by a LARGE AMOUNT!
- ▶ **GPT- 4 is 1+ trillion parameters** & you refine using your data. (Images/Words)
- ▶ GenAI Generates new content/output (GANs & Deep Fake & Hallucinations)

ChatGPT (OpenAI)

GPT-4 has arrived. It will blow ChatGPT out of the water *The Washington Post* 3/14/2023

ChatGPT was launched on November 30, 2022, by San Francisco-based OpenAI (the creator of the GPT series of large language models; DALL·E 2, a diffusion model used to generate images; and Whisper, a speech transcription model).



OpenAI Research Product Developers Safety Company Search Log in Sign up

Introducing ChatGPT

We've trained a model called ChatGPT which interacts in a conversational way. The dialogue format makes it possible for ChatGPT to answer followup questions, admit its mistakes, challenge incorrect premises, and reject inappropriate requests.

Try ChatGPT Read about ChatGPT Plus

OpenAI

ChatGPT → DALL·E → API →

Interact with our flagship language models in a conversational interface

Create realistic images and art from a description in natural language

Integrate OpenAI models into your application or business

Introducing Sora: Creating video from text

OpenAI

Learn more about Sora

RI what are the top 10 databases

As of my knowledge cutoff in September 2021, the top 10 databases were:

1. Oracle Database: A relational database management system (RDBMS) known for its scalability, security, and comprehensive features.
2. MySQL: An open-source RDBMS widely used for web applications due to its simplicity, reliability, and compatibility.
3. Microsoft SQL Server: A popular RDBMS developed by Microsoft, known for its strong integration with Windows Server and seamless integration with other Microsoft products.
4. PostgreSQL: An open-source object-relational database management system (ORDBMS) known for its extensibility, robustness, and compliance with SQL standards.
5. MongoDB: A NoSQL database that uses a flexible, document-oriented model and is known for its scalability, performance, and ease of use.
6. IBM Db2: A family of database management systems developed by IBM, known for its reliability, scalability, and support for a wide range of applications.
7. SQLite: A lightweight, file-based, embedded database engine that requires no separate server process and is commonly used in mobile and embedded systems.
8. Redis: An in-memory data structure store that can be used as a database, cache, or message broker. It is known for its high performance and flexibility.
9. Cassandra: A distributed and highly scalable NoSQL database designed for handling large amounts of data across commodity servers and providing high availability and fault tolerance.
10. Amazon DynamoDB: A fully managed NoSQL database service provided by Amazon Web Services (AWS). It offers seamless scalability, low latency, and automatic data replication.

Regenerate response

Google's Bard (Now Gemini)

≡ Bard Experiment

↻ Reset chat

🕒 Bard Activity

❓ FAQ

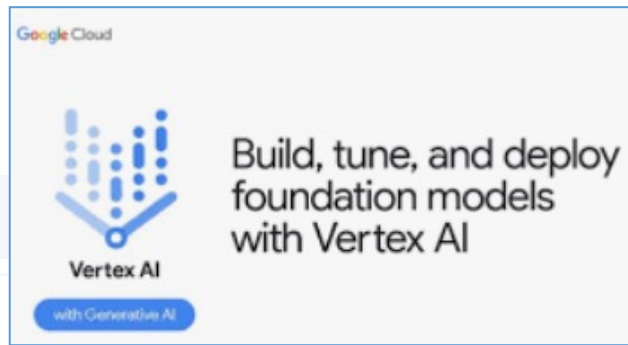
📅 Updates

🗨️ Help

R What are the top 10 databases

Here is a brief overview of each of the top 10 databases:

- Oracle is a commercial database that is known for its ability to store and manage large amounts of data.
- MySQL is an open-source database that is known for its ease of use and flexibility in applications that need to be scalable and easy to maintain.
- Microsoft SQL Server is a commercial database that is used by many businesses that use Microsoft's Windows operating system.
- PostgreSQL is an open-source database that is known for its ability to store complex data or that need to be highly scalable.
- MongoDB is a NoSQL database that is known for its ability to handle large amounts of unstructured data.
- Redis is a NoSQL database that is known for its speed and ability to manage large amounts of data in real time.
- IBM DB2 is a commercial database that is known for its ability to store and manage large amounts of data.
- Elasticsearch is a search engine database that is known for its ability to search and index large amounts of data.
- Cassandra is a NoSQL database that is known for its ability to store and manage large amounts of data in a distributed environment.
- MariaDB is a fork of MySQL that is known for its stability and performance. MySQL is not available or when a more stable database is needed.



Gemini ▾

🔗 See the latest updates to the Gemini Apps Privacy Hub

R What are the top 10 databases?

Show drafts ▾ 🔊

✦ Here are the top 10 most popular databases in 2024:

1. **Oracle Database:** Developed by Oracle Corporation, it provides a scalable and high-performance platform for managing structured and unstructured data. It's known for its reliability and security, making it a good choice for enterprise applications.



Oracle Database

2. **MySQL:** An open-source relational database management system (RDBMS) that offers ease of use and flexibility. It's a popular choice for web applications and small to medium-sized businesses.

Enter a prompt here



Enter a prompt here



● Illinois, USA

From your IP address • [Update location](#)

Transformers - Google, 2017 (Changed Everything!)



Attention Is All You Need

Ashish Vaswani* Google Brain avaswani@google.com
Noam Shazeer* Google Brain noam@google.com
Niki Parmar* Google Research nikip@google.com
Jakob Uszkoreit* Google Research usz@google.com

Llion Jones* Google Research llion@google.com
Aidan N. Gomez* † University of Toronto aidan@cs.toronto.edu
Lukasz Kaiser* Google Brain lukaszkaizer@google.com

Illia Polosukhin* ‡ illia.polosukhin@gmail.com

Abstract

The dominant sequence transduction models are based on complex recurrent or convolutional neural networks that include an encoder and a decoder. The best performing models also connect the encoder and decoder through an attention mechanism. We propose a new simple network architecture, the Transformer, based solely on attention mechanisms, dispensing with recurrence and convolutions entirely. Experiments on two machine translation tasks show these models to be superior in quality while being more parallelizable and requiring significantly less time to train. Our model achieves 28.4 BLEU on the WMT 2014 English-to-German translation task, improving over the existing best results, including ensembles, by over 2 BLEU. On the WMT 2014 English-to-French translation task, our model establishes a new single-model state-of-the-art BLEU score of 41.8 after training for 3.5 days on eight GPUs, a small fraction of the training costs of the best models from the literature. We show that the Transformer generalizes well to other tasks by applying it successfully to English constituency parsing both with large and limited training data.

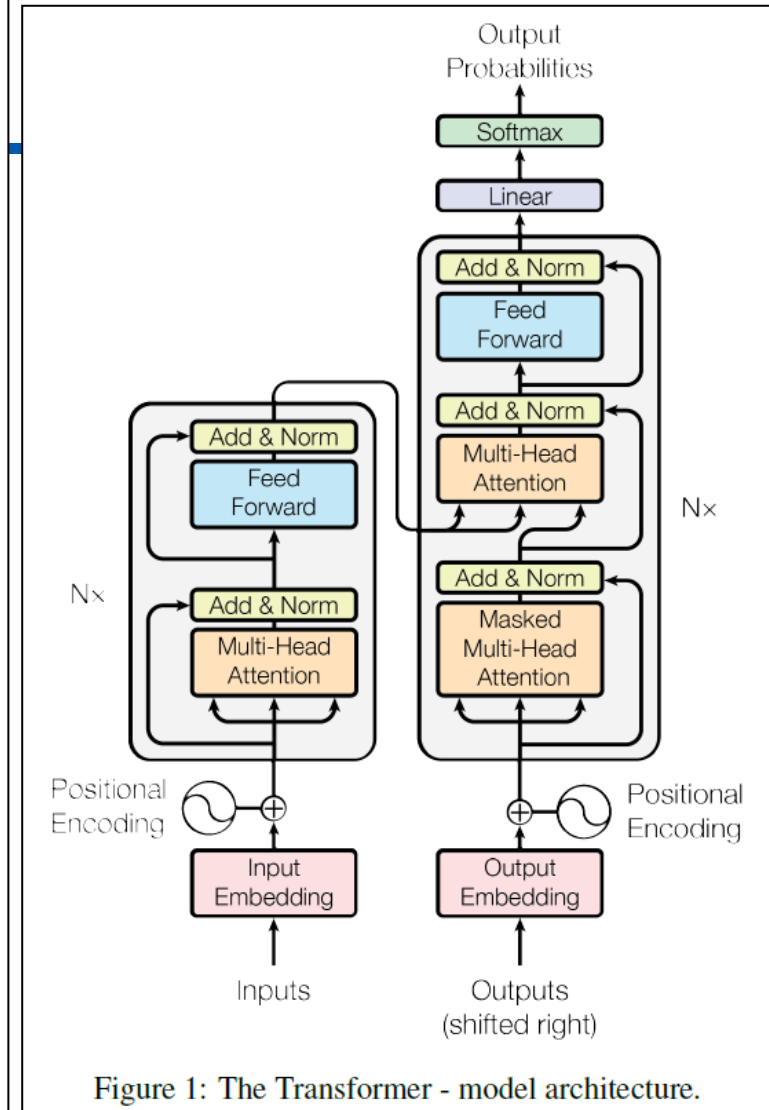


Figure 1: The Transformer - model architecture.

Outperformed Recurrent NN, GRU, & LSTM

Transformers Have Longer Term Memory

- Google BERT
- OpenAI GPT

Next word learned tend to much earlier words thru Backprop.

**A Transformer is sequence to sequence NN
Leverages Parallelism to train model**

iv:1706.03762v5 [cs.CL] 6 Dec 2017

Oracle Vector Database! FREE in the 23ai Database!

the next medium

Oracle Vector Database: Easily Specialize AI Models & Build AI Apps

- Designing & Training Foundational AI Models: Complex and Expensive
 - Design Multi-Billion Parameter Artificial Neural Network: OpenAI, Cohere, xAI ...
 - Ingest Trillions of Training Data Elements: Wikipedia + Everything on the Internet
 - Oracle Cloud: 16,000 Node H100 NVIDIA Supercluster with RDMA Interconnect
- Specializing AI Models: Pretrained Foundational Model + Supplementary Training Data
 - Oracle Vector Database: Stores Vectorized Supplementary Training Data
 - Specialized AI Model with EHR Data: Generate Doctors' Orders – Cerner New Millennium
 - Specialized AI Model with Diagnostic Image Data: Rapid Cancer Detection – Imagene

Most Customers will Train Specialized AI Models for Specific Applications



Embed the Semantics of Word, Image, Video...

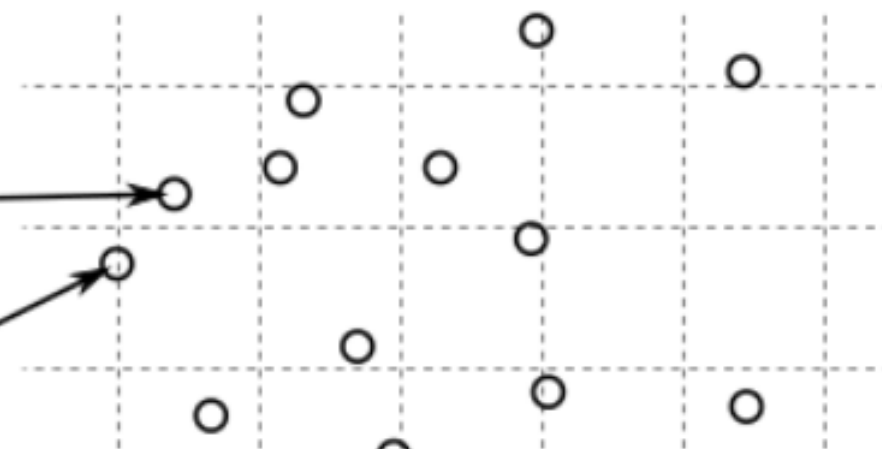
*qdrant.tech

Embedding space

Encoder (



)

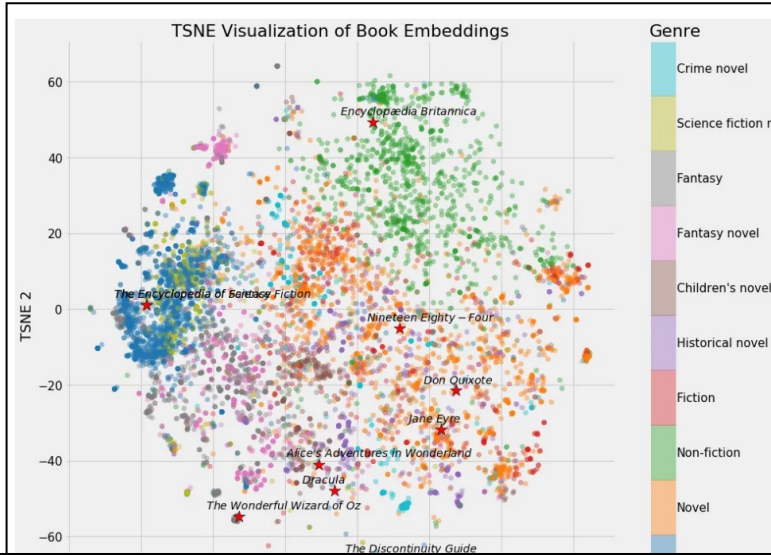


```
SQL> SELECT VECTOR_EMBEDDING(demo_model USING 'The quick brown fox jumped over the lazy dog.' AS DATA) AS embedding;
```

EMBEDDING

```
[6.76711798E-002,2.01067656E-001,1.17165565E-001,-7.62265697E-002,1.71118081E-001,-8.21046531E-004,-1.01488091E-001,-5.965776E-002,-6.28514364E-002,-1.69601902E-001,-2.64612101E-002,1.09567501E-001,1.10740066E-001,-6.55300245E-002,1.1391215E-001,-6.95423409E-002,6.38643801E-002,-9.53121409E-002,-9.34720039E-002,-3.358493E-002,1.33274093E-001,-7.32870176E-002,6.29665628E-002,-1.79692209E-002,-1.26965418E-001,-8.53486508E-002,9.31404009E-002,1.80790409E-001,4.78416262E-003,-1.39018342E-001,-1.10573813E-001,1.94040537E-002,5.87702431E-002,1.70110315E-002,8.841045
```

Vectorize the Data into MANY dimensions



YouTube

Search

CONF42 MACHINE LEARNING 2023
MAY 18 • ONLINE

Machine Learning models

... create vector representations

Chicken, Wolf, Dog, Cat, Banana, Apple

[0.23, 0.45, ..., 0.84, 0.23],
[0.63, 0.32, ..., 0.34, 0.92]

[0.34, 2.35, 8.34, ...]
300 dimensions

Oracle AI Vector Search - Similarity Search

Vector Embedding Model
sentence-transformers/all-MiniLM-L6

Show Top N Results
5

Similarity Search Query
animals

Top Similarity Search Results

1. Dogs are loyal.
2. Mice are small.
3. Bears are hairy.
4. Birds can fly.
5. Horses have four legs.

Introduction to Vector Databases | Zain Hasan | Conf42 Machine Learning 2023

nf42
K subscribers

Subscribe

Like | Share | Download

FaceNet (92% accuracy) uses 128 dimensions CNN vs 64, 256, or 512 (eigens); FaceNet512 - 97%

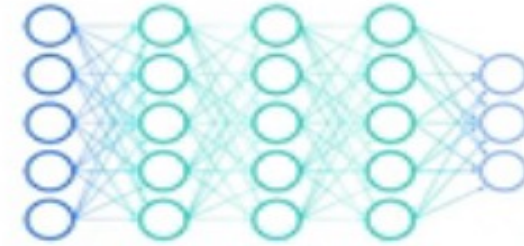
Vectorize DATA with Cohere/OpenAI/Gemini...etc. (Encode Objects & Query)



CONF42 MACHINE LEARNING
MAY 18 • ONLINE

Vector databases

1. Vectorize and index data, using ML

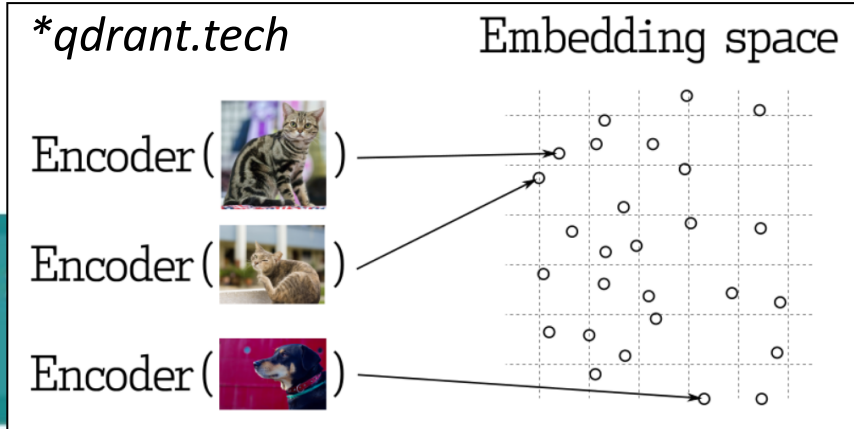


ResNet-50
PRETRAINED MODEL

co:here



HUGGING FACE



Vectorize Query with NN (old/not as good) or Cohere/OpenAI/Gemini (Encode Objects & Query)



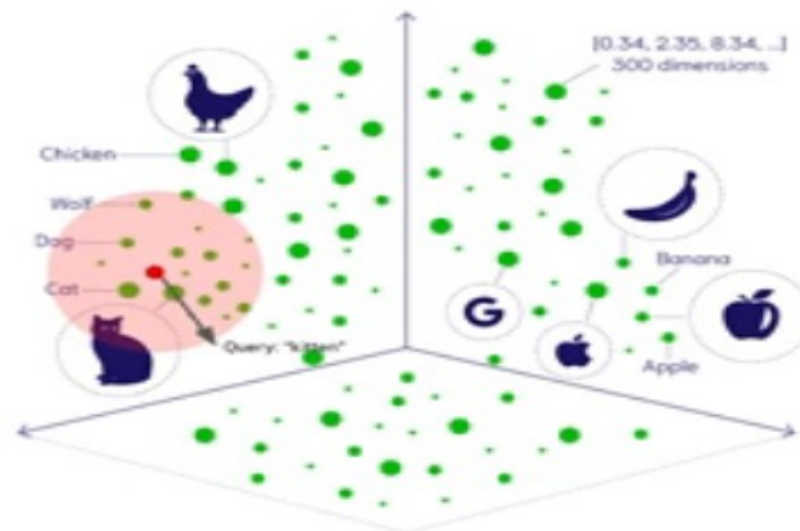
CONF42 MACHINE LEARNING 2023
MAY 18 • ONLINE

Vector databases

1. Vectorize and index data, using ML

2. Vectorize search query, using ML

Example search



A gentle introduction to Vector Databases | Zain Hasan | Conf42 Machine Learning 2023



Conf42
3.54K subscribers

Subscribe

Like



Share

Download



ANN Index- Find Approximate Nearest Neighbor (ANN)



CONF42 MACHINE LEARNING 2023
MAY 18 • ONLINE

Vector databases

1. Vectorize and index data, using ML
2. Vectorize search query, using ML
3. Retrieve ANN results (using HNSW)

Example search

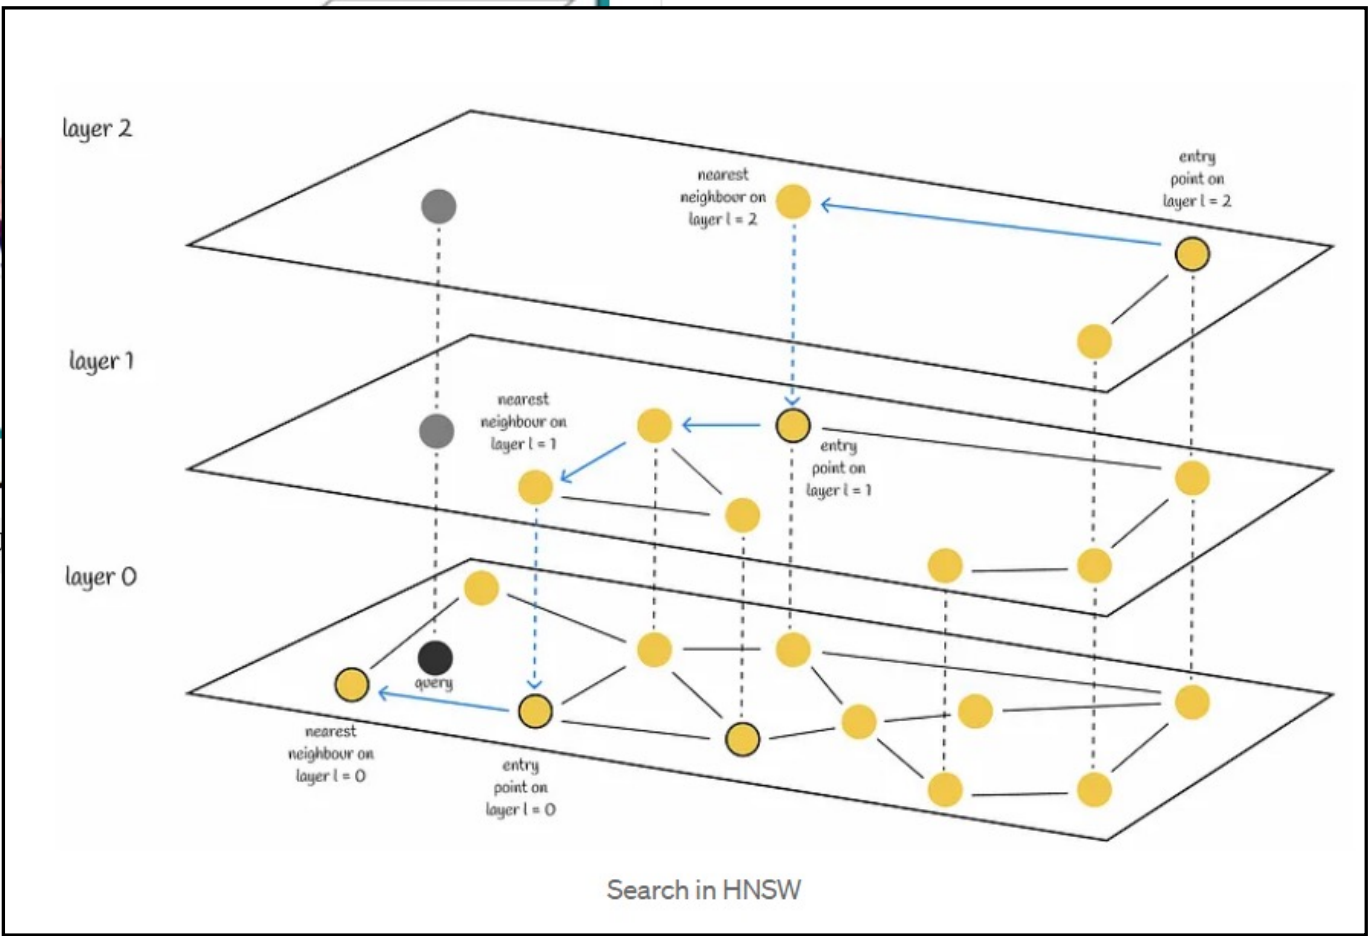
A gentle introduction to Vector Databases | Zain Hasan | Conf42 ML

Conf42
3.54K subscribers

[Subscribe](#)

[Like](#)

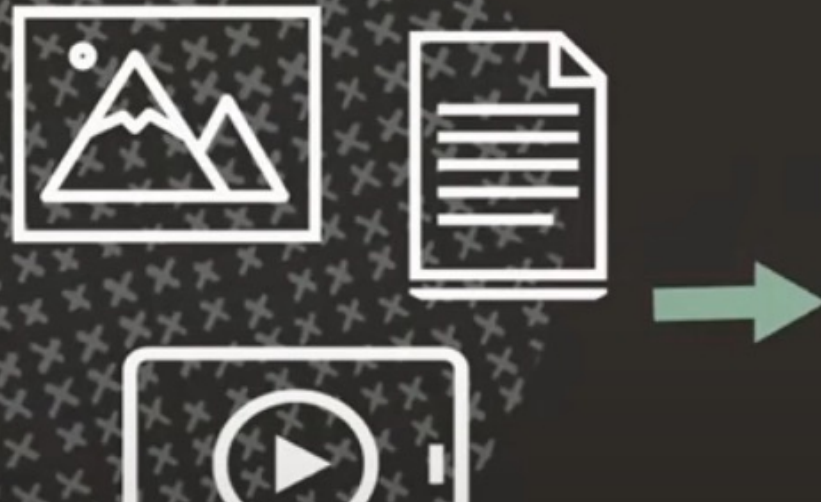
Similarly to NSW, the search quality of HNSW (Hierarchical Navigable Small World - 6 degrees of separation) can be improved by using several entry points.



Search on Data using AI Vector Embeddings*



Vectors in AI represent semantics of unstructured data such as images, documents, videos, etc.



Vector



A vector is a sequence of numbers, called dimensions, used to capture the important "features" of the data

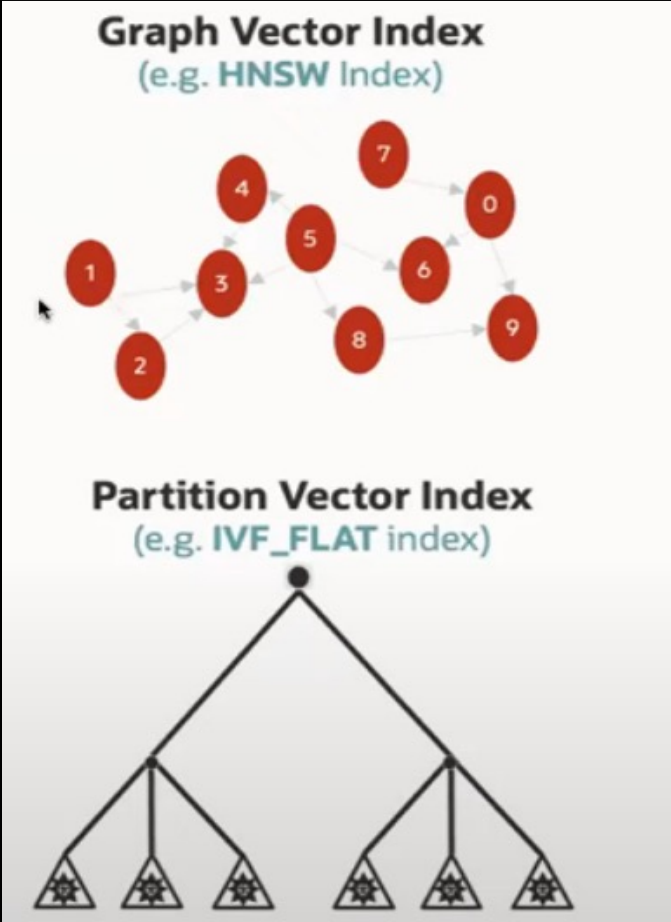
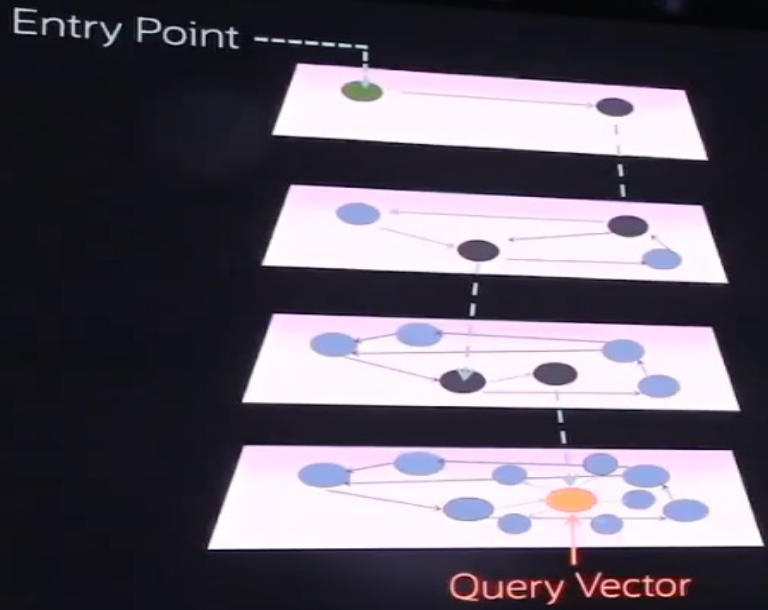
Vectors represent the semantic content of data, not the underlying words or pixels

The mathematical distance between two vectors indicates how similar they are

Vectors generated using deep learning embedding models

EMBEDDING
[6.76711798E-002, 2.01067656E-001, 1.17165565E-001, -7.62265697E-002, 1.71118081E-001, -8.21046531E-004, -1.01488091E-001, -5.965776E-002, -6.28514364E-002, -1.69601902E-001, -2.64612101E-002, 1.09567501E-001, 1.10740066E-001, -6.55300245E-002, 1.1391215E-001, -6.95423409E-002, 6.38643801E-002, -9.53121409E-002, -9.34720039E-002, -3.358493E-002, 1.33274093E-001, -7.32870176E-002, 6.29665628E-002, -1.79692209E-002, -1.26965418E-001, -8.53486508E-002, 9.31404009E-002, 1.80790409E-001, 4.78416262E-003, -1.39018342E-001, -1.10573813E-001, 1.94040537E-002, 5.87702431E-002, 1.70110315E-002, 8.841045

Vector Indexes Leverage Partitions / RAC / Exadata*



Oracle database accelerates AI vector search using sophisticated vector indexes



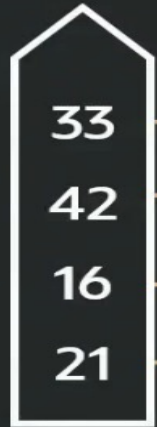
Great Oracle Example of a Use Case

Example: the features for a house image could be

Vector

Features

House



Type of Roof

Decorations

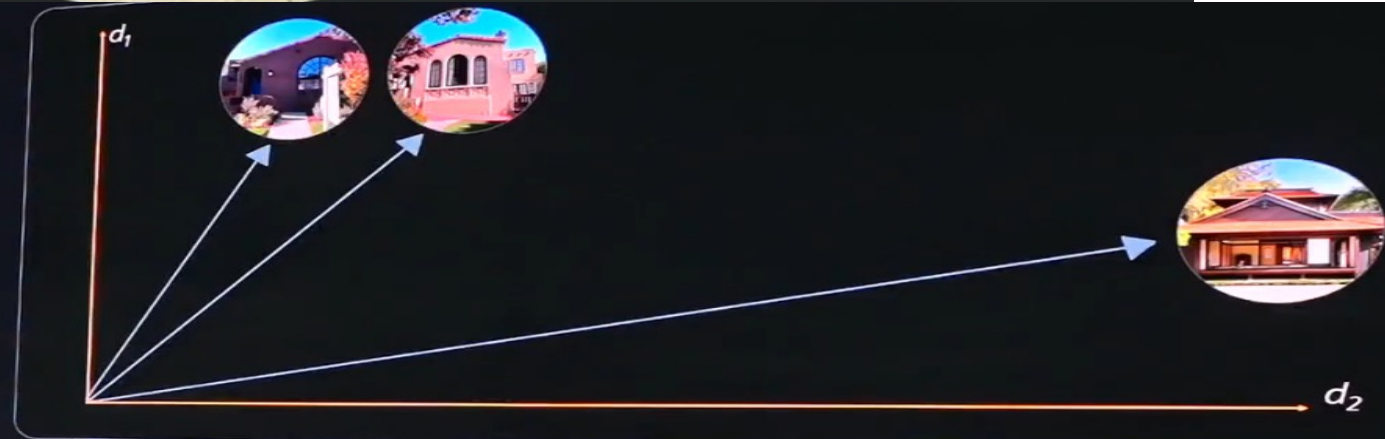
Number of Stories

Building Materials



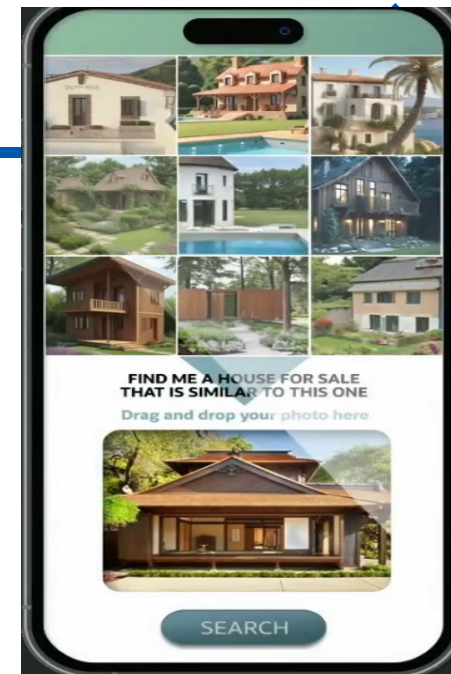
Each dimension represents a different feature of the house

House vectors when collapsed into 2 dimensions instead of hundreds could look like this

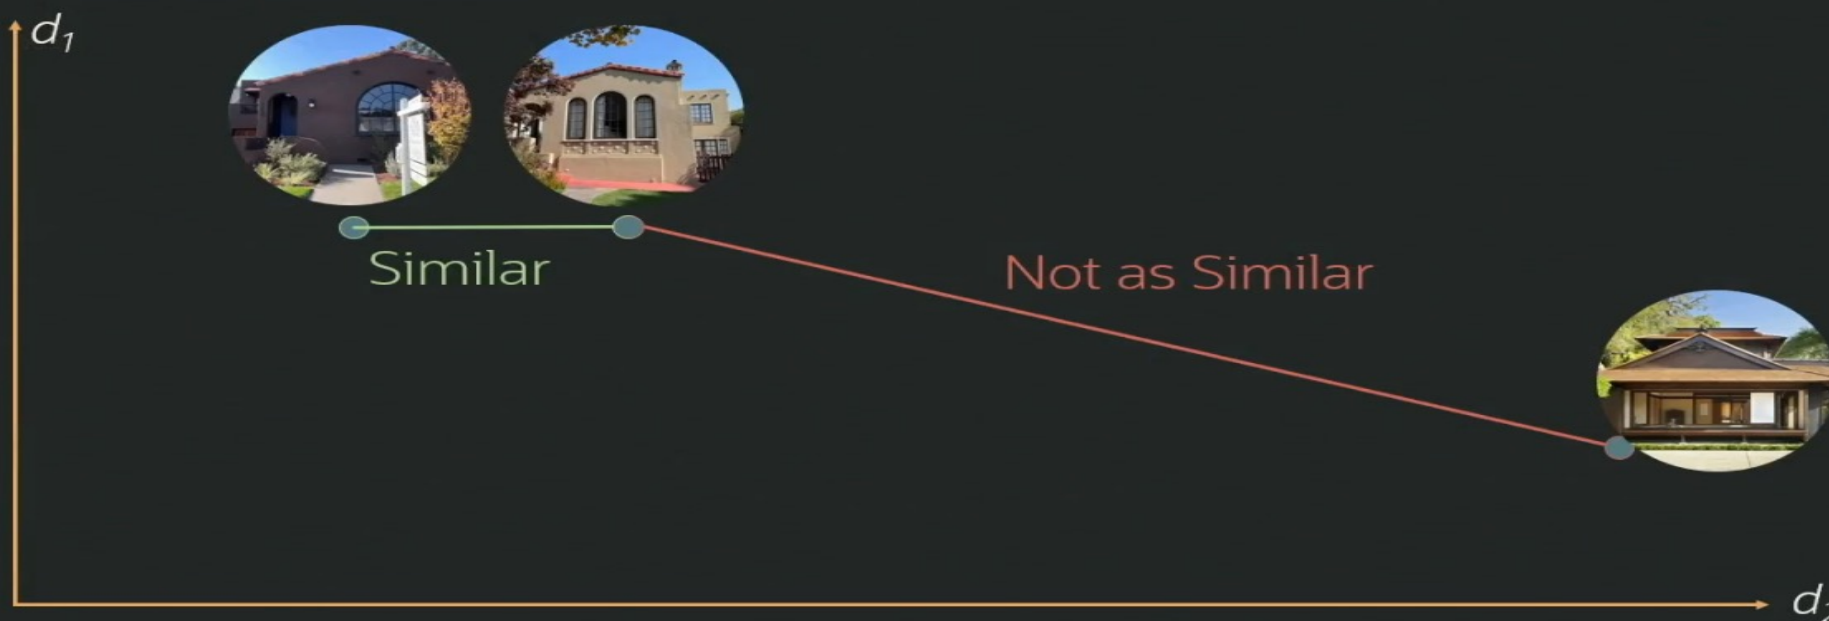


Vector Search

Imagine a house-hunting app that helps customers find houses for sale that are similar to a picture the customer uploads



The distance between the vectors is proportional to their semantic similarity



Similarity Search



Easily search documents, images, and other unstructured data
Based on their conceptual content

Create table with Vector Data Type & Blob

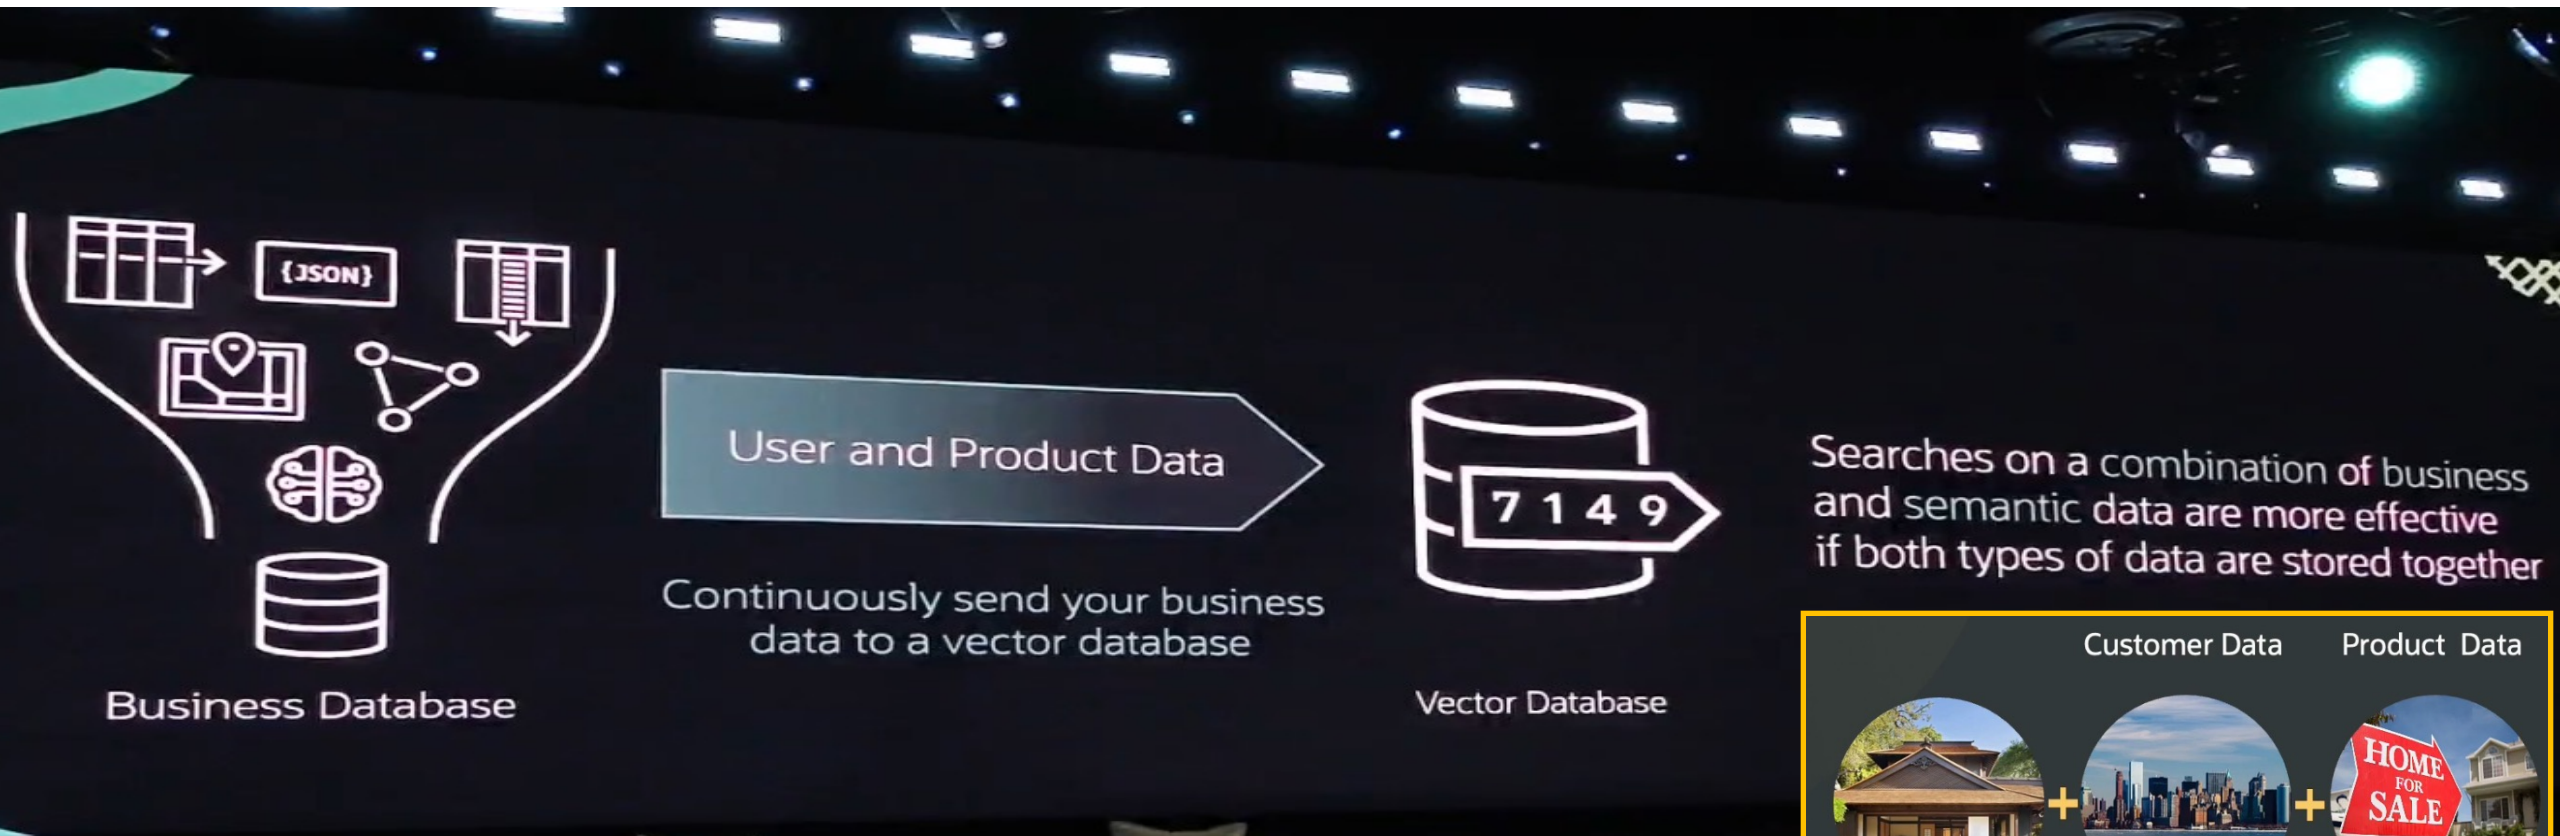
```
CREATE TABLE house_for_sale (house_id      number,  
                             price         number,  
                             city          varchar2(400),  
                             house_photo  blob,  
                             house_vector vector  
);
```

Find houses that are similar to this picture and match the customer's preferred city and budget



```
SELECT ...  
FROM   house_for_sale  
WHERE  price <= (SELECT budget          FROM customer ...)  
AND    city  in (SELECT search_city    FROM customer ...)  
ORDER BY vector_distance(house_vector, :input_vector);
```

Search on vector (GenAI), add prices/locations*



Oracle AI Vector Search is part of the Oracle Database

Store your vectors in the same row as your data

- Insert vectors with new data
- Update vectors when data changes
- Delete vectors with your data



How do vectors fit into AI?



Vector Databases **augment** Generative AI
by retrieving **detailed**, often **private content**
needed to answer questions

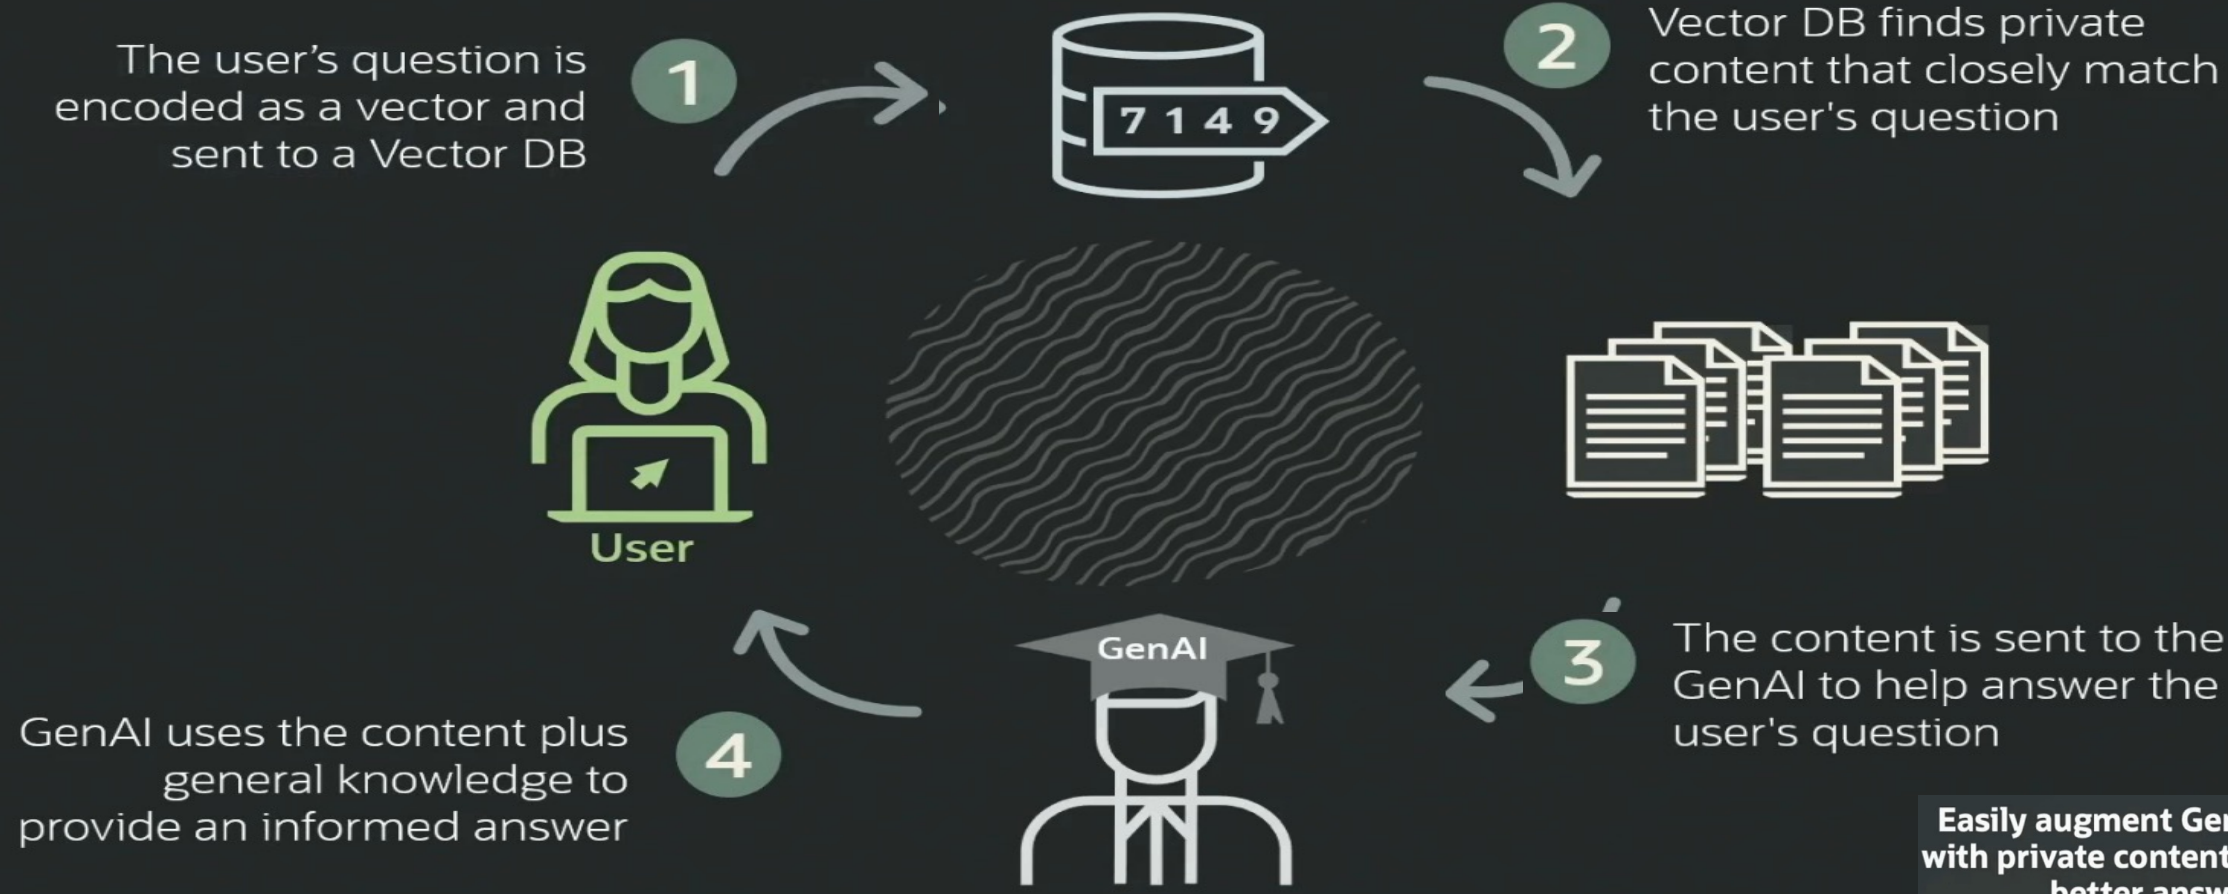
Called: Retrieval Augmented Generation (RAG)

Better business outcomes

LLM + Vector Database

Retrieval Augmented Generation

Retrieval Augmented Generation How it works



Easily augment Generative AI with private content to produce better answers

Cohere & Oracle GenAI Partnership (Signed June 2023)

- ▶ **Aiden Gomez, Cohere CEO** worked on “Attention is all you need” paper on Transformers Aidan while at Google Brain
- ▶ **Cohere LLM turns words to numbers with semantic knowledge**
- ▶ **GenAI Challenges - Hallucinations (made up “facts”)**
- ▶ **R.A.G.** - most promising **solution to hallucinations**
- ▶ **Web LLMs can use your data with R.A.G. (Retrieval Augmented Generation)**
- ▶ **Vector DB** uses your information & knowledge that you have
- ▶ **Training vs. RAG** - Training good to **get model to speak in your brand voice** - Data can change in a msec when you update the DB; RAG keeps it up to date.
- ▶ **Embedding Models** - Cohere performs twice as well in speed & accuracy. Also compress 32x with 96% accuracy.
- ▶ **Noisy data** (pdf scrapes / errors); bad emails; No matter how noisy, you still get accurate results with what we’re building.

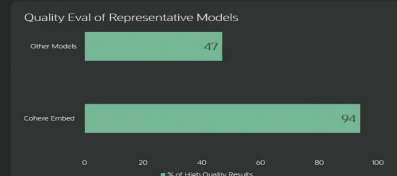


**Officially launched partnership at
Cloud World (Sept 2023)**



Cohere's New Embed Model: 2x Better in Noisy Data

Excels in Noisy Datasets



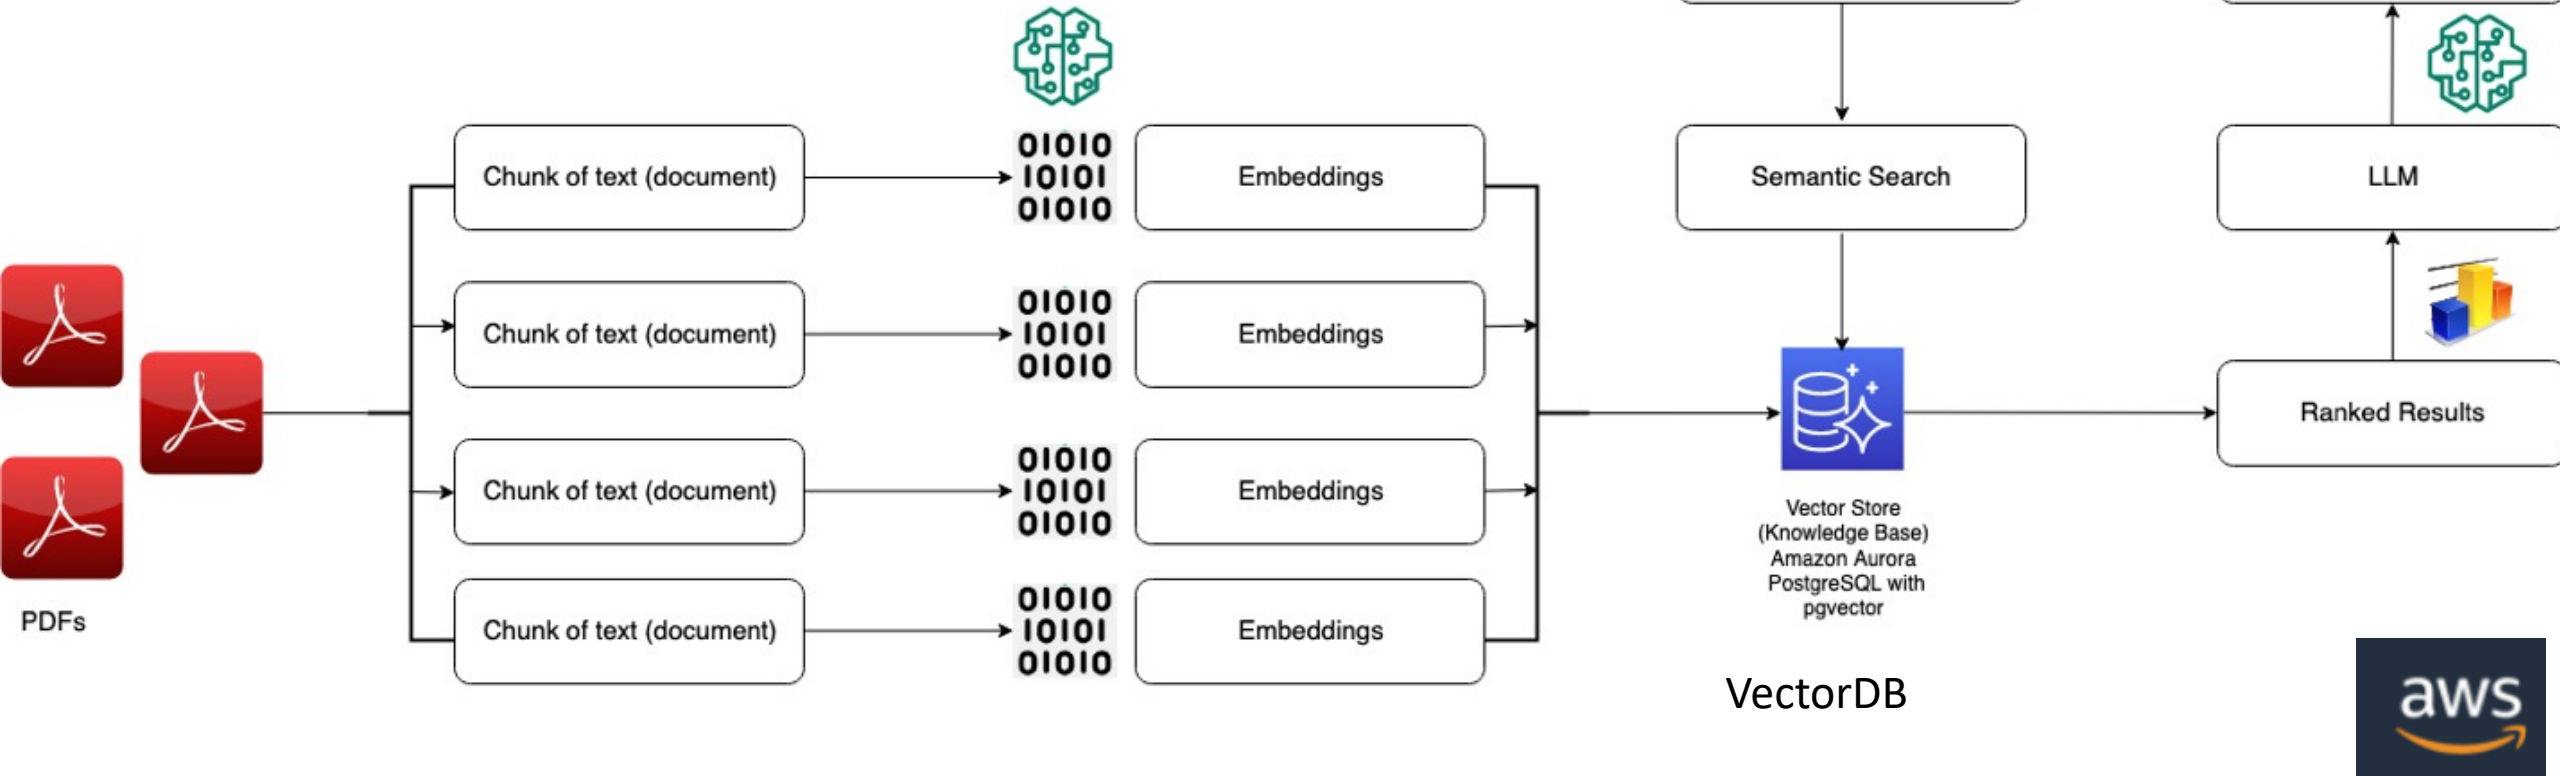
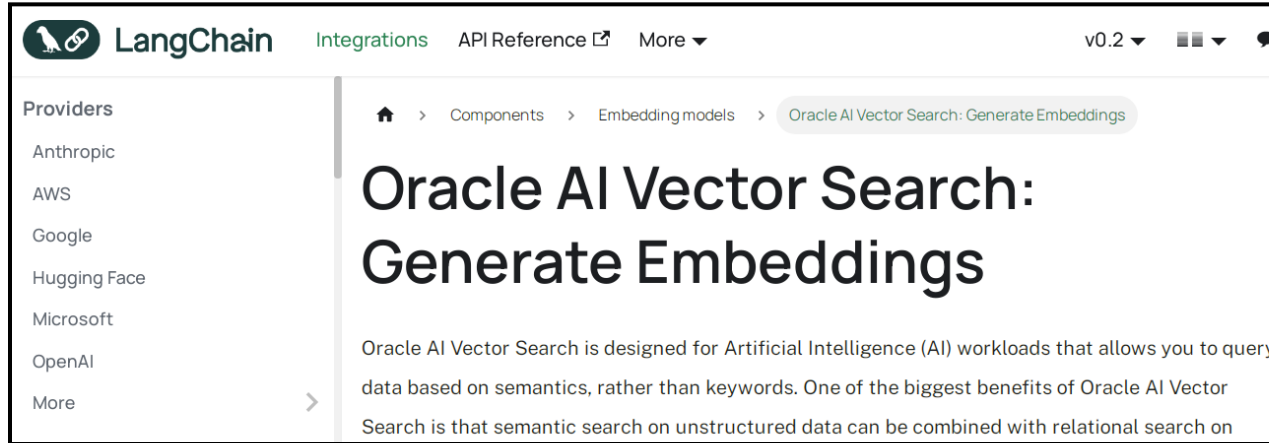
Scalable to Larger Datasets

Compression	Search Quality
4x	99.9%
16x	98%
32x	96%

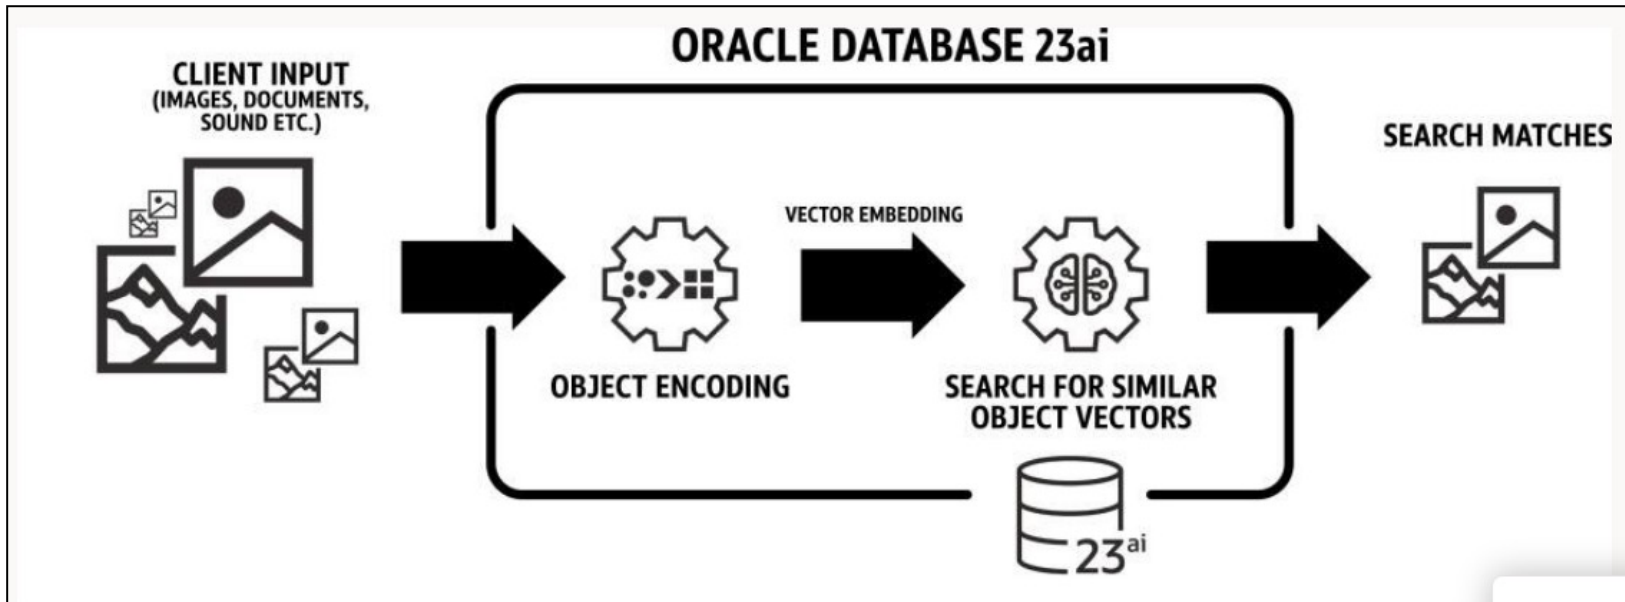
Embed is now trained to be compression-aware

Tools use LangChain (works with Oracle) & put it together

23ai interfaces with LangChain



Similarity Searches using LLMs are easier**



- ▶ Similarity Searches or exact searches using SQL & PL/SQL
- ▶ Provide LLMs with context to search
- ▶ Create, Search & Index Vectors
- ▶ **All in Oracle23ai (23.4.0.24.05)**

It's FREE with 23ai Database

Richard Niemiec • You
Chief Innovation Officer, Viscosity North America
1d • 🌐

JUST NOW - #LarryEllison says the #Oracle #Vector #Database will come #FREE with the Oracle 23c Database (called it 23AI)!

#OracleAce #MachineLearning #Python #BigData #DBA #Data #TensorFlow #algorithm #DataScience #javascript #java #json #Analytics #ML #developer #SQL #database #AI #VR #rpa #Linux

Colin Monahan and 202 others
2 comments • 14 reposts

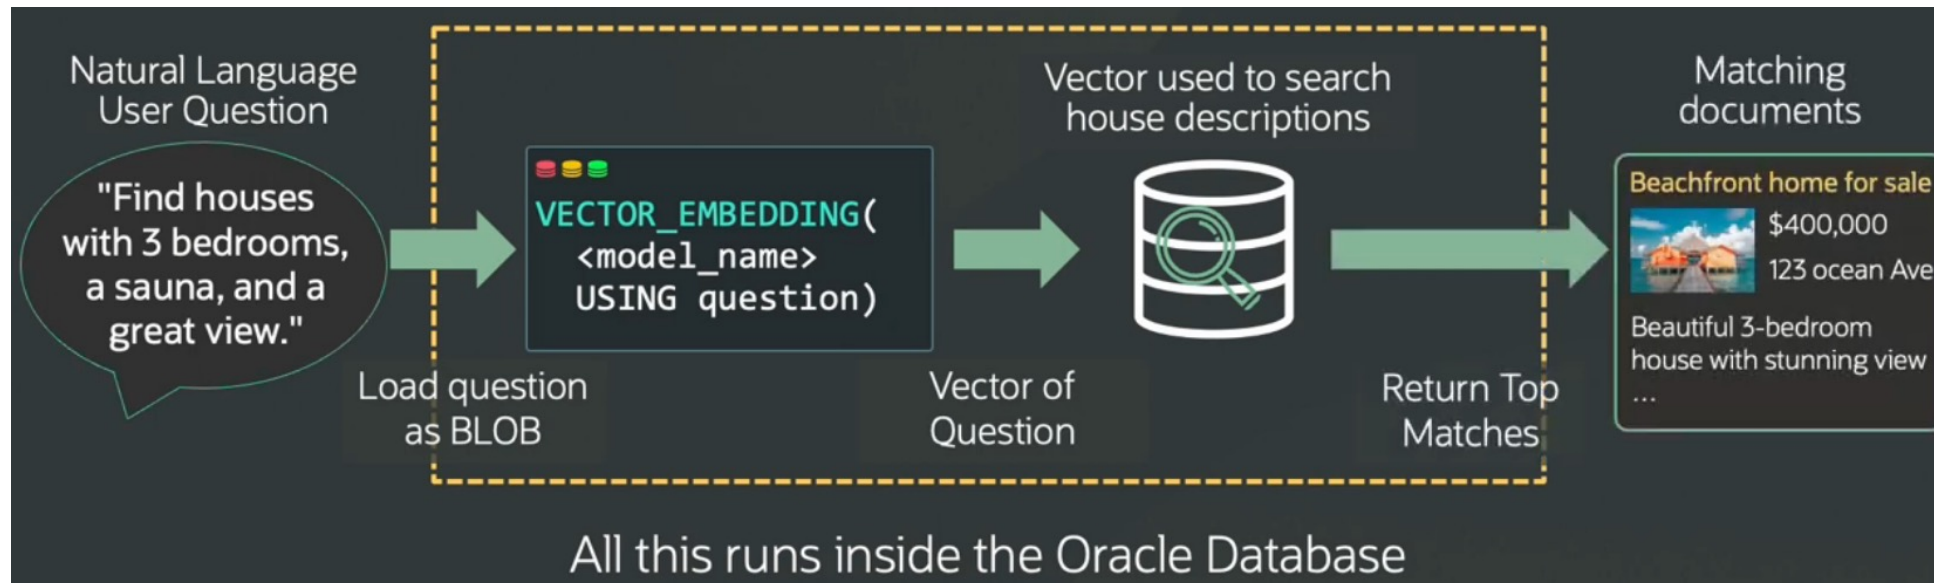
Reactions

Like Comment Repost Send

12,427 impressions View analytics

The World of LLMs & Embeddings

- ▶ LLMs & Add-on Models: ChatGPT, Google Gemini, Cohere, Meta LLaMA, xAI Grok, & Nvidia LLaMA ChatQA
- ▶ DBMS_VECTOR to take text to embeddings & back to text & create vector indexes.
- ▶ VECTOR_CHUNKS to process docs to smaller pieces & then generate embeddings.
- ▶ ONNX - Open Neural Network Exchange

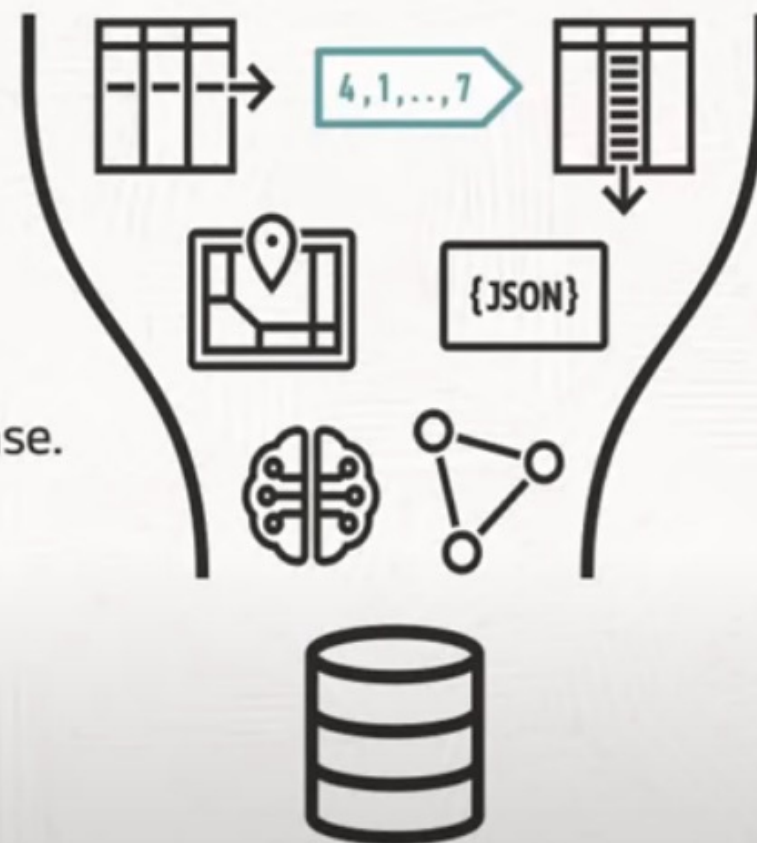


Oracle AI Vector Search

AI Vector Search will be in the 23.4 Release early this year

Designed to be simple to use and easy to understand

- **New** VECTOR data type for storing vector embeddings
 - `Create table docs (doc_id NUMBER, doc_vec VECTOR);`
- **New** SQL syntax & functions expresses similarity search with ease.
 - `Select id from docs order by vector_distance(doc_vec, :query_vec) fetch approximate first 10 rows only;`
- **New** Approximate search indexes packaged and tuned for high performance and quality.
 - `Create vector index doc_idx on docs(doc_vec) organization inmemory neighbor graph`
- **New** Native support for VECTOR in Oracle clients including Python and Node.js



SQL with Oracle AI Vector Search*

*Doug Hood, CloudWorld



Get top 5 nearest vectors to a given vector

```
select id
from vector_table
order by VECTOR_DISTANCE(data, :vector)
fetch first 5 rows only;
```

Get top 5 photos similar to a query photo
taken before the year 2020

```
select id from vector_table t
where t.attributes.year.number() < 2020
order by VECTOR_DISTANCE(data, :vector)
fetch first 5 rows only;
```

Get top 5 nearest neighbors
to a specific vector in the dataset

```
with query as (select id, data
from vector_tab
where id = :id)
select t.id from vector_table t, query q
where t.id != q.id
order by VECTOR_DISTANCE(t.data, q.data) asc
fetch first 5 rows only;
```

Get all neighbors within a threshold distance from the
query vector

```
select id from vector_table
where VECTOR_DISTANCE(data,
:query, 'MANHATTAN') < 5;
```



New Features Doc - AI Vector Search - 3rd Party LLMs**

- Use PL/SQL Packages to Generate Embeddings
 - Terms of Using Vector Utility PL/SQL Packages
 - About Chainable Utility Functions and Common Use Cases
 - About Vector Helper Procedures
 - Supplied Vector Utility PL/SQL Packages
 - Supported Third-Party Provider Operations**
 - Validate JSON Input Parameters
- ▶ Vector Generation Examples
- ▶ 4 Store Vector Embeddings
- ▶ 5 Create Vector Indexes**
- ▶ 6 Use SQL Functions for Vector Operations
- ▶ 7 Query Data with Similarity Searches
- ▶ 8 Work with Retrieval Augmented Generation
- 9 Supported Clients and Languages
- ▶ 10 Vector Diagnostics
- ▶ 11 Vector Search PL/SQL APIs

Supported Third-Party Provider Operations

Review the list of third-party REST providers that are supported with Vector Utility PL/SQL packages and the corresponding API calls allowed for each of those.

- The supported third-party REST providers are:
- Cohere
 - Generative AI
 - Google AI
 - Hugging Face
 - OpenAI
 - Vertex AI

The corresponding REST calls allowed for each operation are:

- `DBMS_VECTOR.UTL_TO_EMBEDDING` and `DBMS_VECTOR.UTL_TO_EMBEDDINGS`
- `DBMS_VECTOR_CHAIN.UTL_TO_EMBEDDING` and `DBMS_VECTOR_CHAIN.UTL_TO_EMBEDDINGS`
- `DBMS_VECTOR_CHAIN.UTL_TO_SUMMARY`
- `DBMS_VECTOR.UTL_TO_GENERATE_TEXT`
- `DBMS_VECTOR_CHAIN.UTL_TO_GENERATE_TEXT`

Parent topic: [Use PL/SQL Packages to Generate Embeddings](#)



Convert to Embeddings (Vector)

Break into Pieces to embed

Vector back To Text

Generate Vector Embeddings... in the Oracle AI Vector Search docs**

2. If you are using a third-party embedding model and need to make a REST call, set up your credentials for the REST provider:

Using Cohere, Google AI, Hugging Face, OpenAI, and Vertex AI:

a. Run `DBMS_VECTOR.CREATE_CREDENTIAL` to create and store a credential.

Cohere, Google AI, Hugging Face, OpenAI, and Vertex AI require the following authentication parameter:

```
{ "access_token": "<access token>" }
```

You will later refer to this credential name when declaring JSON parameters for the `UTL_TO_EMBEDDING` call.

```
exec dbms_vector.drop_credential('<credential name>');
```

```
declare
  jo json_object_t;
begin
  jo := json_object_t();
  jo.put('access_token', '<access token>');
  dbms_vector.create_credential(
    credential_name => '<credential name>',
    params          => json(jo.to_string));
end;
/
```

Replace the `access_token` and `credential_name` values. For example:

```
declare
  jo json_object_t;
begin
  jo := json_object_t();
  jo.put('access_token', 'AbabA1B123aBc123AbabAb123a1a2ab');
  dbms_vector.create_credential(
    credential_name => 'HF_CRED',
    params          => json(jo.to_string));
end;
/
```

b. Call `DBMS_VECTOR.UTL_TO_EMBEDDING`:

b. Call `DBMS_VECTOR.UTL_TO_EMBEDDING`:

```
-- select example
var params clob;
exec :params := '
{
  "provider": "<REST provider>",
  "credential_name": "<credential name>",
  "url": "<REST endpoint URL for embedding service>",
  "model": "<embedding model name>"
}';

select dbms_vector.utl_to_embedding('hello', json(:params)) from dual;

-- PL/SQL example

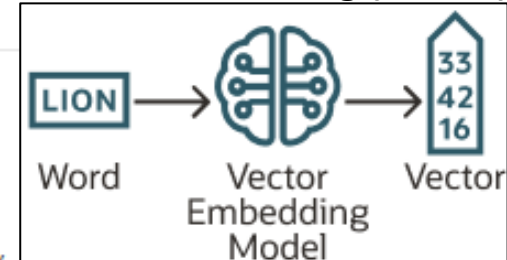
declare
  input clob;
  params clob;
  v vector;
begin
  input := 'hello';

  params := '
{
  "provider": "<REST provider>",
  "credential_name": "<credential name>",
  "url": "<REST endpoint URL for embedding service>",
  "model": "<embedding model name>"
}';

  v := dbms_vector.utl_to_embedding(input, json(params));
  dbms_output.put_line(vector_serialize(v));
exception
  when OTHERS THEN
    DBMS_OUTPUT.PUT_LINE (SQLERRM);
    DBMS_OUTPUT.PUT_LINE (SQLCODE);
end;
/
```

Replace the `provider`, `credential_name`, `url`, and `model` values. Optionally, you can specify add

Create Embedding (Vector)



Generate Vector Embeddings... in the Oracle AI Vector Search docs**

Replace the `provider`, `credential_name`, `url`, and `model` values. Optionally, you can specify additional

Cohere example:

```
{
  "provider": "cohere",
  "credential_name": "COHERE_CRED",
  "url": "https://api.cohere.example.com/embed",
  "model": "embed-model",
  "input_type": "search_query"
}
```

Google AI example:

```
{
  "provider": "googleai",
  "credential_name": "GOOGLEAI_CRED",
  "url": "https://googleapis.example.com/models/",
  "model": "embed-model"
}
```

Hugging Face example:

```
{
  "provider": "huggingface",
  "credential_name": "HF_CRED",
  "url": "https://api.huggingface.example.com/",
  "model": "embed-model",
  "wait_for_model": "true"
}
```

Hugging Face example:

```
{
  "provider": "huggingface",
  "credential_name": "HF_CRED",
  "url": "https://api.huggingface.example.com/",
  "model": "embed-model",
  "wait_for_model": "true"
}
```



OpenAI example:

```
{
  "provider": "openai",
  "credential_name": "OPENAI_CRED",
  "url": "https://api.openai.example.com/embeddings",
  "model": "embed-model"
}
```

Vertex AI example:

```
{
  "provider": "vertexai",
  "credential_name": "VERTEXAI_CRED",
  "url": "https://googleapis.example.com/models/",
  "model": "embed-model"
}
```

Multimodal

- Image-Text-to-Text
- Visual Question Answering
- Document Question Answering

Computer Vision

- Depth Estimation
- Image Classification
- Object Detection
- Image Segmentation
- Text-to-Image
- Image-to-Text
- Image-to-Image
- Image-to-Video
- Unconditional Image Generation
- Video Classification
- Text-to-Video
- Zero-Shot Image Classification

Models 640,235 Filter by name

new Full-text search Sort: Trending

meta-llama/Meta-Llama-3-8B
Text Generation • Updated 12 days ago • 761k • 3.27k

gradientai/Llama-3-8B-Instruct-Gradient-1048k
Text Generation • Updated 1 day ago • 11.7k • 457

apple/OpenELM
Updated 5 days ago • 1.15k

meta-llama/Meta-Llama-3-8B-Instruct
Text Generation • Updated 12 days ago • 1.21M • 1.82k

NousResearch/Hermes-2-Pro-Llama-3-8B
Text Generation • Updated 1 day ago • 4.74k • 231

nvidia/Llama3-ChatQA-1.5-8B
Text Generation • Updated 2 days ago • 2.63k • 189

Model developers Meta

Variations Llama 3 comes in two sizes — 8B and 70B parameters — in pre-trained and instruction tuned variants.

Input Models input text only.

Output Models generate text and code only.

Model Architecture Llama 3 is an auto-regressive language model that uses an optimized transformer architecture. The tuned versions use supervised fine-tuning (SFT) and reinforcement learning with human feedback (RLHF) to align with human preferences for helpfulness and safety.

	Training Data	Params	Context length	GQA	Token count	Knowledge cutoff
Llama 3	A new mix of publicly available online data.	8B	8k	Yes	15T+	March, 2023
		70B	8k	Yes		December, 2023

Cohere - 104B parameters
ChatGPT3.5 - 175B parameters
ChatGPT4 - 1.76T parameters



Create & Search with HNSW Index** (Enough InMemory?)

```
SELECT docID
FROM vector_tab
ORDER BY VECTOR_DISTANCE( embedding, :query_vector, EUCLIDEAN ) FETCH EXACT FIRST 10 ROWS ONLY;
```

```
CREATE VECTOR INDEX galaxies docs_hnsw_idx ON galaxies (embedding) ORGANIZATION INMEMORY NEIGHBOR GRAPH (HNSW Index)
DISTANCE COSINE
WITH TARGET ACCURACY 95;
```

```
SELECT name
FROM galaxies
WHERE name <> 'NGC1073'
ORDER BY VECTOR_DISTANCE( embedding, to_vector('[0,1,1,0,0]'), COSINE )
FETCH APPROXIMATE FIRST 3 ROWS ONLY;
```

```
EXPLAIN PLAN FOR... (select statement)
select plan_table_output from table(dbms_xplan.display('plan_table',null,'all'));
```

PLAN_TABLE_OUTPUT				
Plan hash value: 2946813851				
Id	Operation	Name	Rows	Bytes
0	SELECT STATEMENT		4	104
* 1	COUNT STOPKEY			
2	VIEW		5014	127
* 3	SORT ORDER BY STOPKEY		5014	19
4	TABLE ACCESS BY INDEX ROWID	DOC_CHUNKS	5014	19
5	<u>VECTOR INDEX HNSW SCAN</u>	<u>DOCS_HNSW_IDX</u>	5014	19

Vector Distance Metrics^{**}: *What you used training embedding model*

- **Euclidean and Euclidean Squared Distances**

Euclidean distance reflects the distance between each of the vectors' coordinates being compared—basically the straight-line distance between two vectors. This is calculated using the Pythagorean theorem applied to the vector's coordinates ($\text{SQRT}(\text{SUM}((x_i - y_i)^2))$).

- **Cosine Similarity**

One of the most widely used similarity metric, especially in natural language processing (NLP), is cosine similarity, which measures the cosine of the angle between two vectors.

- **Dot Product Similarity**

The dot product similarity of two vectors can be viewed as multiplying the size of each vector by the cosine of their angle. The corresponding geometrical interpretation of this definition is equivalent to multiplying the size of one of the vectors by the size of the projection of the second vector onto the first one, or vice versa.

- **Manhattan Distance**

This metric is calculated by summing the distance between the dimensions of the two vectors that you want to compare.

- **Hamming Similarity**

The Hamming distance between two vectors represents the number of dimensions where they differ.



FYI

KEY Functions & Packages

ORACLE Database 23^{ai}

To learn more, visit:

[Oracle.com/Database](https://www.oracle.com/database)

- ▶ `TO_VECTOR()`
- ▶ `VECTOR_DISTANCE()`
- ▶ `VECTOR_DIMENSION_COUNT()`
- ▶ `VECTOR_DIMENSION_FORMAT()`
 - `SELECT chunk_id, chunk_data`
 - `FROM doc_chunks`
 - `ORDER BY VECTOR_DISTANCE(chunk_embedding, :query_vector, COSINE)`
 - `FETCH APPROX FIRST 4 ROWS ONLY WITH TARGET ACCURACY 80;`

- ▶ `DBMS_VECTOR` (common operations)
 - `DBMS_VECTOR.CREATE_CREDENTIAL`
 - `DBMS_VECTOR.CREATE_INDEX`
 - `DBMS_VECTOR.UTL_TO_CHUNKS`
 - `DBMS_VECTOR.UTL_TO_EMBEDDINGS`
 - `DBMS_VECTOR.UTL_TO_GENERATE_TEXT`
 - `DBMS_VECTOR.LOAD_ONNX_MODEL`

- ▶ `DBMS_VECTOR_CHAIN` (Advanced: many of the above Subprograms)



```
// import embedding model
DBMS_VECTOR.LOAD_ONNX_MODEL(
  model_name => "...",
  model_data => <source>
  ...);
```

Oracle puts GenAI, AI & Machine Learning together with Speed!



Easily build and run Algorithmic AI models using standard SQL



Vector Search can be combined with business data search



Augment Generative AI with detailed, often private business content



Deploy vector search across the enterprise with GoldenGate



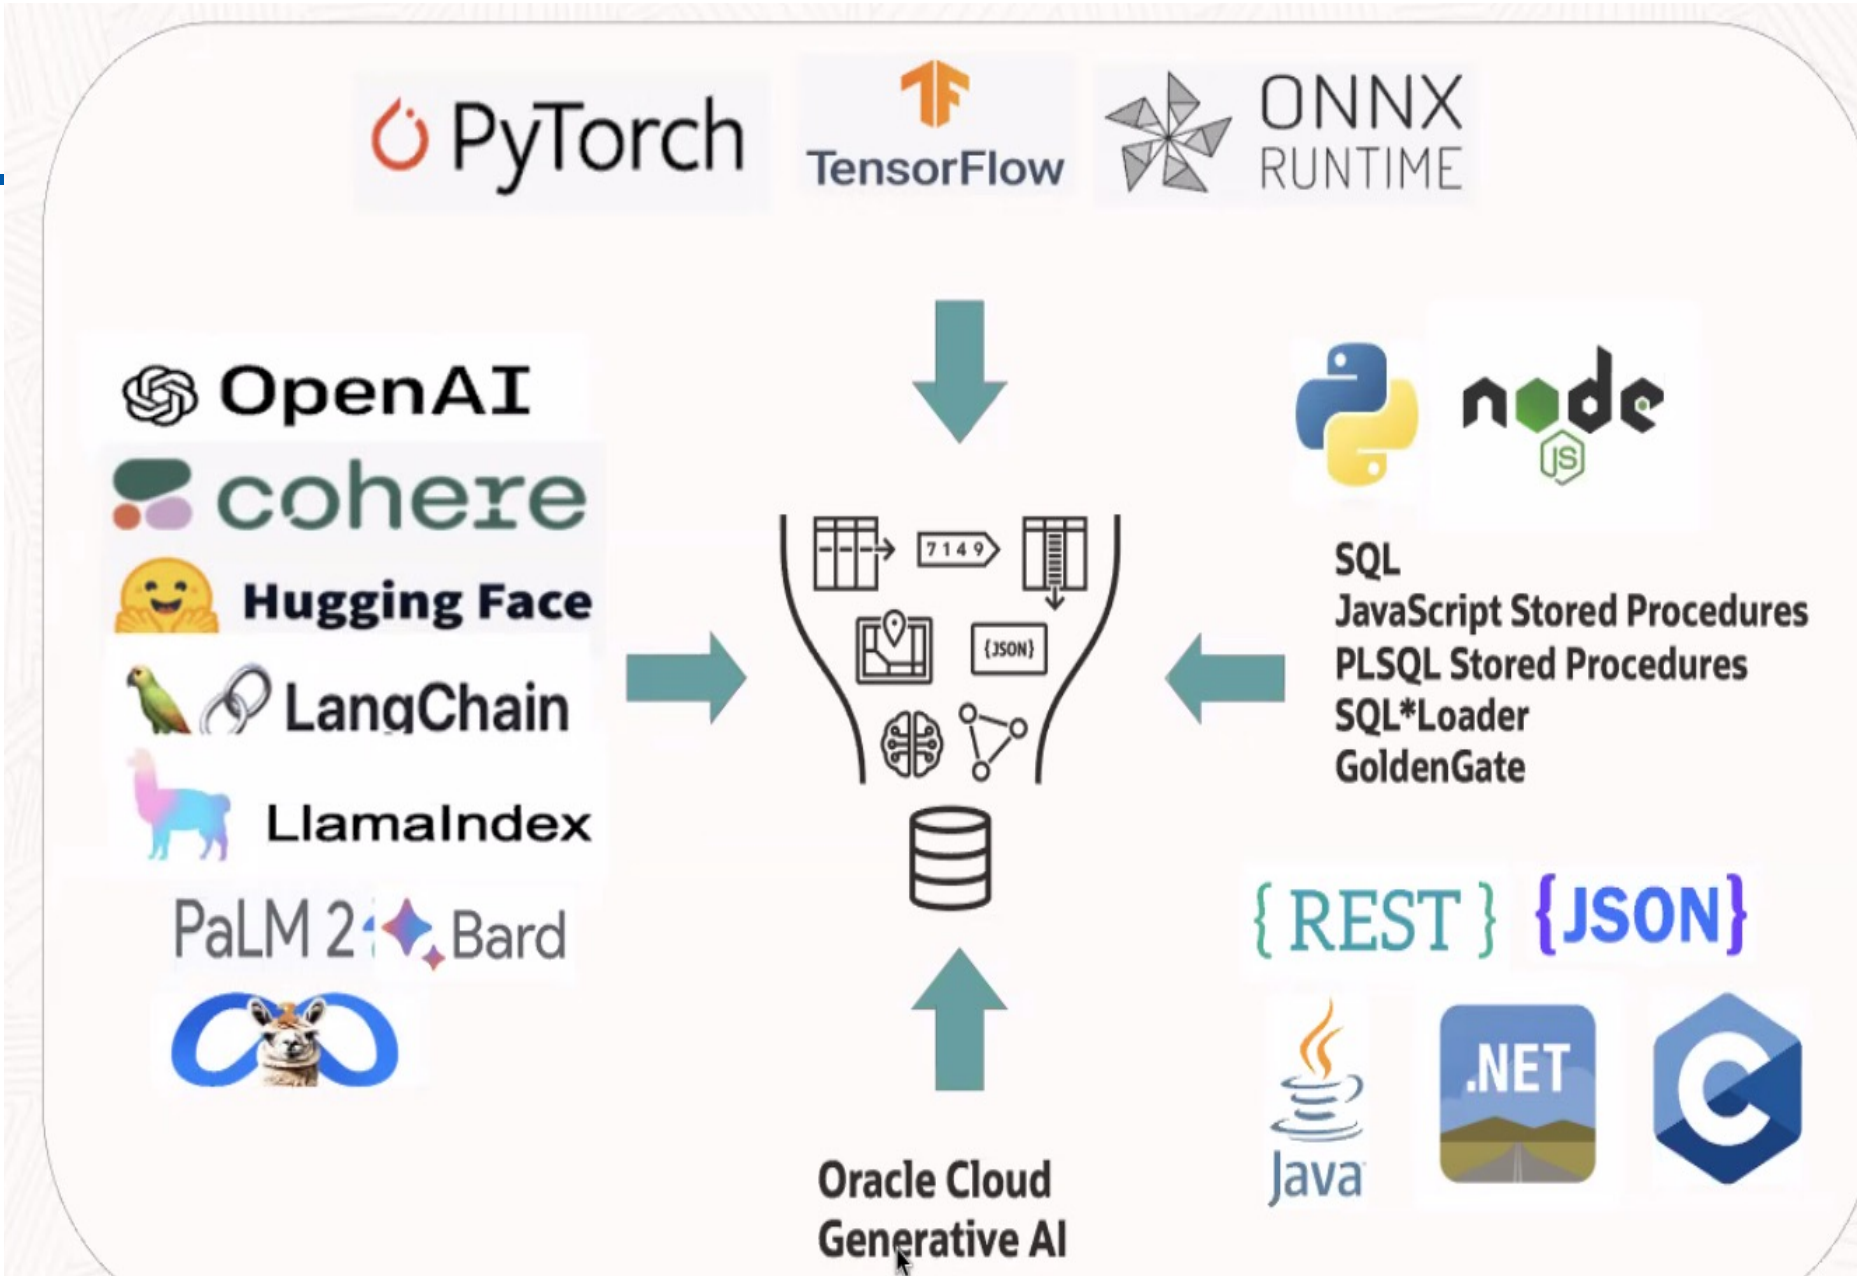
Enable mission-critical AI at any scale with Exadata



Easily use AI in any application with AI tools

All AI capabilities benefit from Oracle's mission-critical **consistency, scalability, availability, and security**

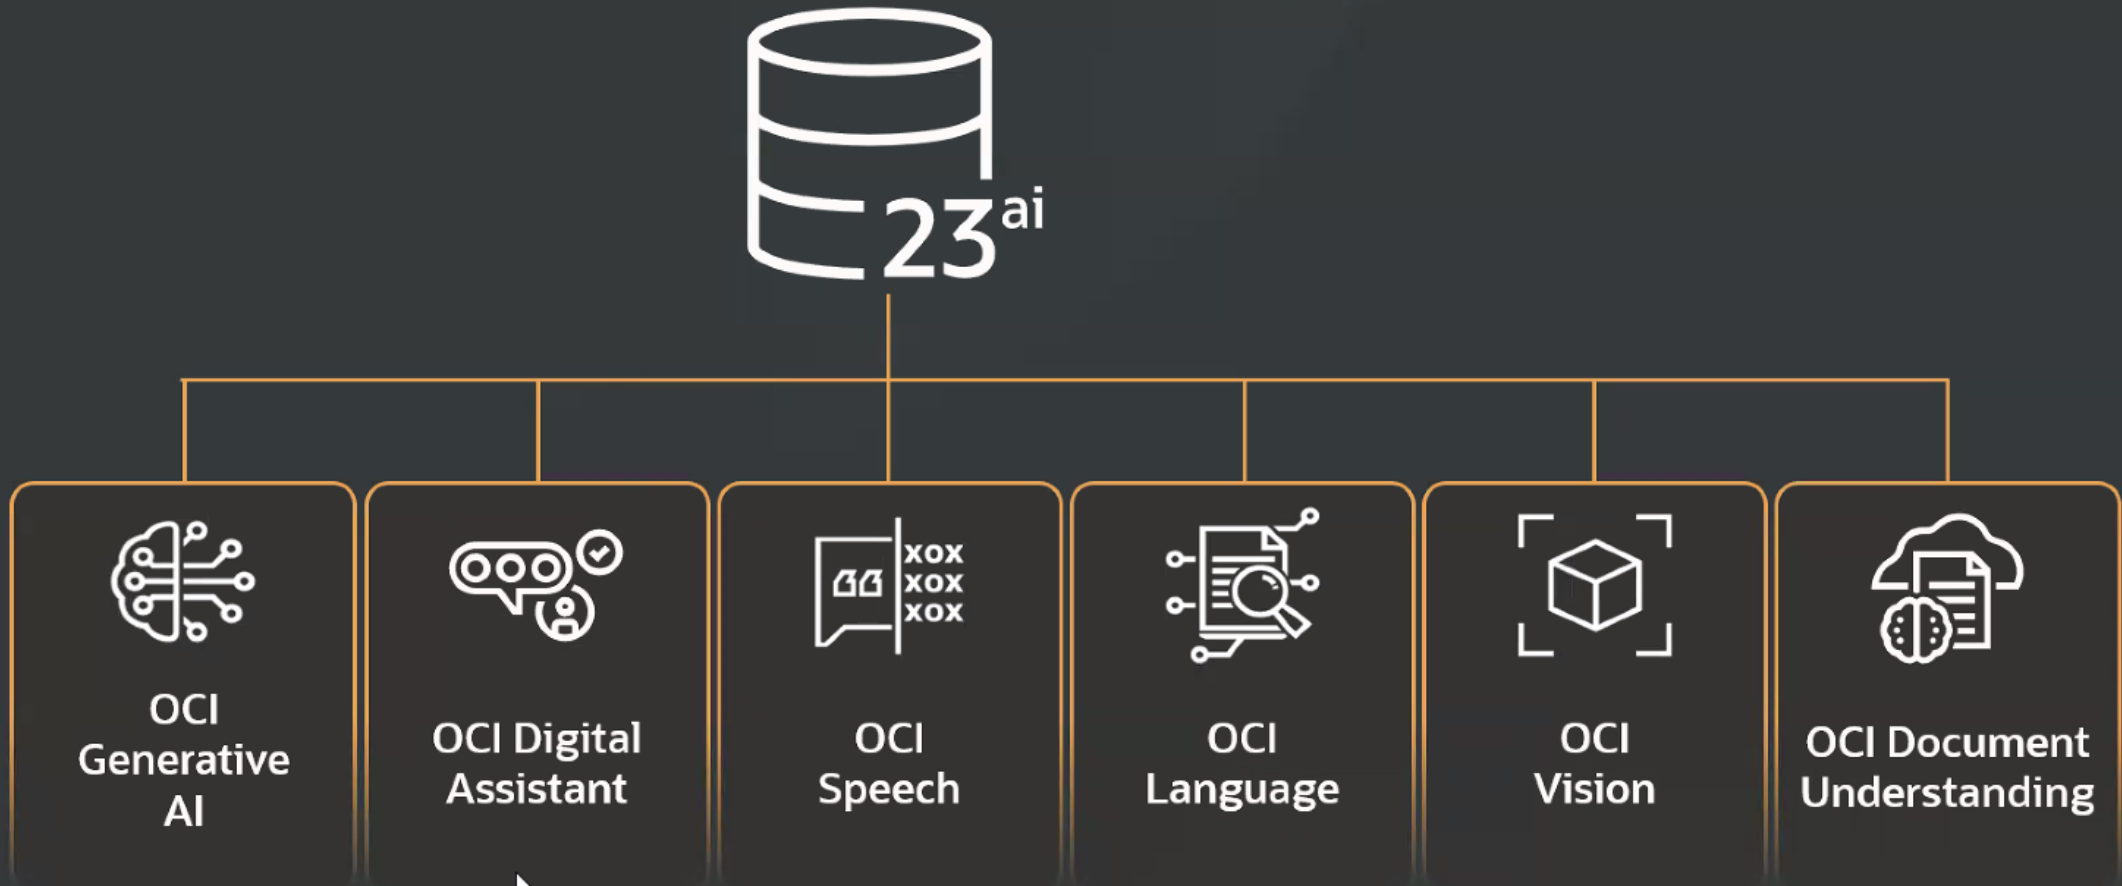
Oracle GenAI & Vector Database Connections





Integrate with Other OCI Services

Oracle Database 23ai integrates with all Oracle Cloud AI Services



23ai Oracle:

docs.oracle.com/en/database/oracle/oracle-database/23 (@richniemiec)



[Get Started](#)

[Related Products](#)

[Administration](#)

[Development](#)

[AI Vector Search](#)

[Machine Learning](#)

[Error Help Portal](#)

[Install](#)

[Upgrade](#)

[Licensing](#)

[Security](#)

[Books](#)

[Home](#) / [Database](#) / [Oracle](#) / [Oracle Database](#) / [Release 23](#)

Oracle Database 23ai

Oracle Database Documentation



Learn About Oracle Database

- [What's New in Release 23ai](#)
- [Oracle Database Features and Licensing App](#)
- [Behavior Changes, Deprecations, and Desupports in Oracle Database 23ai](#)
- [Oracle Interactive Technical Architecture Diagrams](#)
- [Oracle Database Quickstart](#)
- [Run SQL with Oracle Live SQL](#)
- [Oracle Database 23ai Release Notes](#)
- [Oracle Release Analyzer Diff Utility](#)



LiveLabs

- [Exploring JSON Relational Duality Views in 23ai Free using SQL](#)
- [Exploring JSON Relational Duality Views in 23ai Free with Java](#)
- [Exploring Operational Property Graphs in 23ai Free](#)



Development

- [JSON Relational Duality](#)
- [Application Usage Domains and Annotations](#)
- [Sagas for Microservices](#)
- [JavaScript Stored Procedures](#)
- [SQL Property Graphs](#)
- [Machine Learning](#)
- [Python Driver for Oracle Database \(python-oracledb\)](#)
- [Node-oracledb Driver for Oracle Database](#)
- [Oracle Database Programming Interface for C \(ODPI-C\)](#)



Oracle Database 23ai - Free

- [Oracle Database 23ai - Free Overview](#)
- [Oracle Database 23ai - Free Get Started](#)
- [Oracle Database 23ai - Free Install Guide](#)

Performance Tips in Execution Plans of Oracle 23ai*

- ▶ Oracle 23ai has a new section “**SQL Analysis Report**” at the end of the output of `dbms_xplan.display`. It contains recommendations

EXPLAIN PLAN FOR

```
SELECT * FROM addresses  
WHERE UPPER(ctr_code) = 'GB';
```

```
SELECT * FROM dbms_xplan.display();
```

Plan hash value: 3184888728 (*MORE ... Next Slide*)

Performance Tips in Execution Plans of Oracle 23ai*



```
-----  
| Id | Operation                | Name          | Rows | Bytes | Cost (%CPU)| Time |  
-----  
| 0  | SELECT STATEMENT         |               | 239  | 13384 | 61 (0)| 00:00:01 |  
|* 1 | TABLE ACCESS FULL| ADDRESSES    | 239  | 13384 | 61 (0)| 00:00:01 |  
-----
```

Predicate Information (identified by operation id):

```
-----  
1 - filter(UPPER("CTR_CODE")='GB')
```

SQL Analysis Report (identified by operation id/Query Block Name/Object Alias):

```
-----  
1 - SEL$1 / "ADDRESSES"@ "SEL$1"  
   - The following columns have predicates which preclude their use as keys in index range scan. Consider rewriting the predicates. "CTR_CODE"
```



DROP/CREATE TABLE IF [NOT] EXIST

- ▶ The if [not] exist clause auf the create table and [drop table](#) statements are useful to prevent the [ORA-00942: table or view does not exist](#) and *ORA-00955: name is already used by an existing object* errors

```
drop table IF EXIST emp;
```

```
create table IF NOT EXIST emp_new  
(num number, txt varchar2(10));
```




Improved error messages

```
select sum(val), gr_1, gr_2  
from T  
group by gr_1;
```

ORA-00979: not a GROUP BY expression.



```
select sum(val), gr_1, gr_2  
from T  
group by gr_1;
```

ORA-00979: "GR_2": must appear in the GROUP BY clause or be used in an aggregate function



Improved error messages

► [@GeraldVenzl](#) via [@phsalvisberg](#) ([#DOAG2022](#) Keynote – screen shot below)

Much better error messages

```
SELECT TO_NUMBER('abc');
```

ERROR at line 1:
ORA-01722: invalid number

↓

```
SELECT TO_NUMBER('abc');
```

ERROR at line 1:
ORA-01722: unable to convert value 'abc'

Much better error messages

```
sqlplus gerald@//localhost:1521/NOT_A_VALID_DB_NAME
```

ORA-12514: TNS:listener does not currently know of service requested in connect descriptor

↓

```
sqlplus gerald@//localhost:1521/NOT_A_VALID_DB_NAME
```


ORA-12514: Cannot connect to database. Service FOO is not registered with the listener at host 127.0.0.1 port 1521. (CONNECTION_ID=6Th6GtlxLcHgUy5CRmQQGg==)

 FYI

Flashback Time Travel (works with PDBs)

More efficient & performant in 23ai (helps compliance)

 23^{ai} ↑

- ▶ Flashback Time Travel (FDA pre-19c): **Ability to track & store** definitional (including schema) and **transactional changes to a table over its lifetime**. Since: **19^c** 
- ▶ **Flashback Time Travel, enable tracking of DML** (such as INSERT and DELETE) **and DDL** operations on a table. You can then archive the changes made. Flashback Time Travel maintains a history. Having the history of the table and schema enables you to issue flashback queries (AS OF and VERSIONS) on the table and its schema. **You can view the history of DDL and DML changes made to the table.**
- ▶ Create several Flashback Archives in your database; a logical entity that is associated with a set of tablespaces. **Set a space quota & retention period** Using a Flashback Archive helps in compliance with policies and audit reports.
- ▶ **DBMS_FLASHBACK_ARCHIVE_MIGRATE SINCE 21^c**

Flashback Time Travel to Enforce *Digital Shredding*



19^c ORACLE[®]
Database

To "shred" (delete) historical data changes to the TAXES table after ten years:

- ▶ **CREATE FLASHBACK ARCHIVE taxes_archive TABLESPACE tbs1 RETENTION 10 YEAR;**
- ▶ **When history data from transactions on Taxes exceeds the age of ten years, it is purged.**
- ▶ **The Taxes table itself, and history data from transactions less than ten years old, are not purged.**



Flashback Time Travel Restrictions - FYI

- ▶ You **cannot enable Flashback Archive** on tables with **LONG data type** or nested table columns.
- ▶ You **cannot enable Flashback Archive** on a **nested table, temporary table, external table, materialized view, Advanced Query (AQ) table, hybrid partitioned tables, or non-table object.**
- ▶ **Flashback Archive does not support DDL statements that** move, split, merge, or coalesce partitions or sub partitions, move tables, or convert LONG columns to LOB columns.



Common Commands for Flashback Archive

- ▶ CREATE TABLE employee (EMPNO NUMBER(4) NOT NULL, ENAME VARCHAR2(10), JOB VARCHAR2(9), MGR NUMBER(4)) **FLASHBACK ARCHIVE;**
- ▶ CREATE TABLE employee (EMPNO NUMBER(4) NOT NULL, ENAME VARCHAR2(10), JOB VARCHAR2(9), MGR NUMBER(4)) **FLASHBACK ARCHIVE fla1;**
- ▶ ALTER TABLE employee FLASHBACK ARCHIVE;
- ▶ ALTER TABLE employee FLASHBACK ARCHIVE fla1;

- ▶ ALTER TABLE employee **NO FLASHBACK ARCHIVE;** (to Disable it)
- ▶ Static Data Dictionary Views for Flashback Archive Files: *_FLASHBACK_ARCHIVE

Additional Commands (If size/quota not specified then it's *unlimited!*)

- ▶ CREATE FLASHBACK ARCHIVE DEFAULT fla1 TABLESPACE tbs1
QUOTA 10G RETENTION 1 YEAR;
- ▶ CREATE FLASHBACK ARCHIVE fla2 TABLESPACE tbs2 **RETENTION 2 YEAR;**
- ▶ ALTER FLASHBACK ARCHIVE fla1 SET DEFAULT;
- ▶ **ALTER FLASHBACK ARCHIVE fla1 ADD TABLESPACE tbs3 QUOTA 5G;**
- ▶ ALTER FLASHBACK ARCHIVE fla1 MODIFY TABLESPACE tbs3 QUOTA 20G;
- ▶ **ALTER FLASHBACK ARCHIVE fla1 MODIFY RETENTION 2 YEAR;**
- ▶ **ALTER FLASHBACK ARCHIVE fla1 PURGE ALL;**
- ▶ ALTER FLASHBACK ARCHIVE fla1 PURGE BEFORE SCN 728969;
- ▶ **DROP FLASHBACK ARCHIVE fla1;**

Using Flashback Time Travel to Access Historical Data



▶ **CREATE FLASHBACK ARCHIVE DEFAULT fla1 TABLESPACE**

tbs1 QUOTA 10G RETENTION 5 YEAR;

▶ **ALTER TABLE inventory FLASHBACK ARCHIVE;**

▶ **ALTER TABLE stock_data FLASHBACK ARCHIVE;**

Retrieve inventory of all items at the beginning of the year 2022:

▶ **SELECT product_number, product_name, count FROM inventory AS OF TIMESTAMP TO_TIMESTAMP ('2022-01-01 00:00:00', 'YYYY-MM-DD HH24:MI:SS');**

Retrieve stock price for each symbol in portfolio at close of business July 23, 2007:

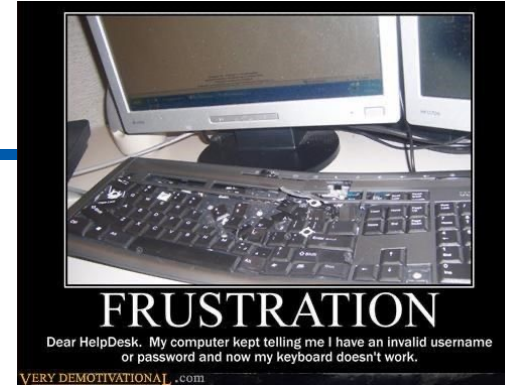
SELECT symbol, stock_price FROM stock_data AS OF TIMESTAMP TO_TIMESTAMP ('2007-07-23 16:00:00', 'YYYY-MM-DD HH24:MI:SS') WHERE symbol IN my_portfolio;

Grants needed for Flashback Archive

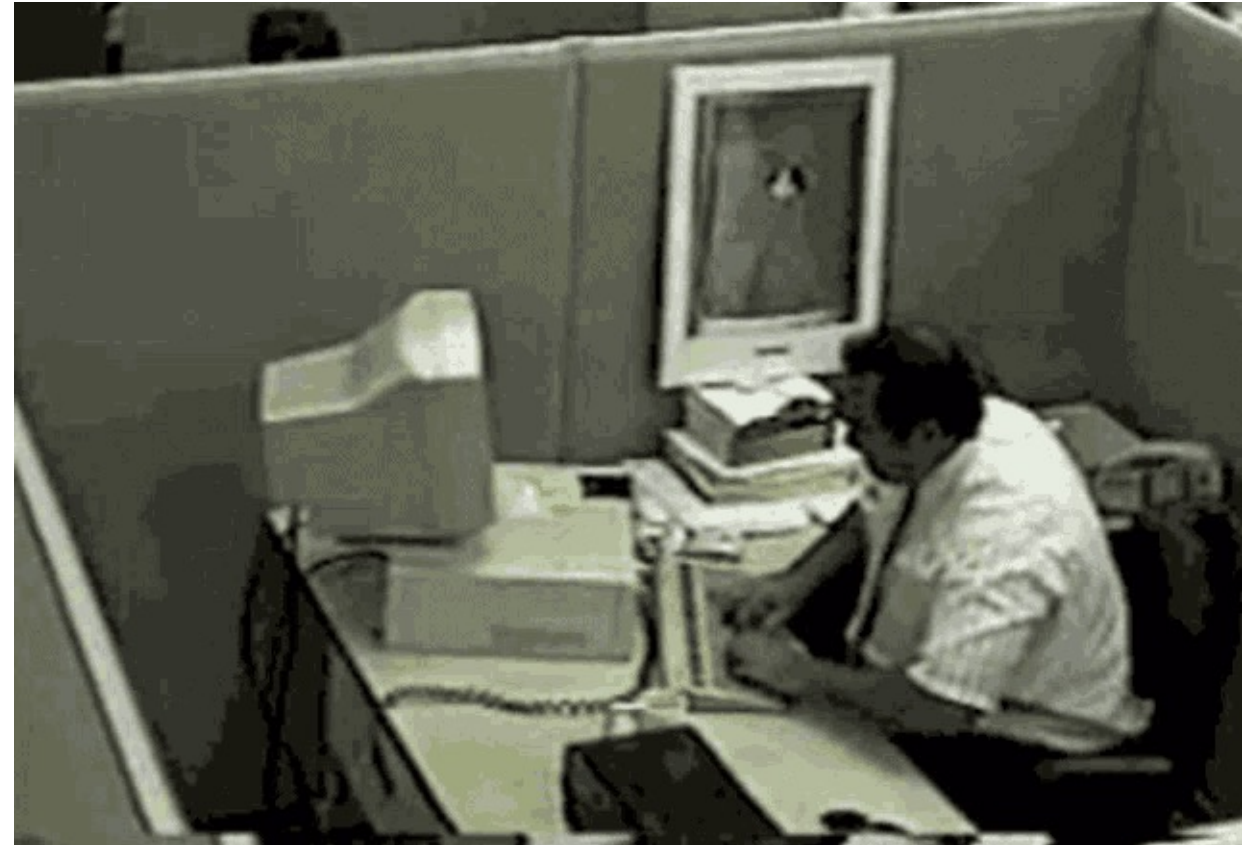
- ▶ CREATE USER test IDENTIFIED BY *password*;
- ▶ GRANT CREATE SESSION TO test;
- ▶ GRANT CONNECT, RESOURCE TO test;
- ▶ GRANT SELECT ON SCOTT.EMP TO test;
- ▶ GRANT FLASHBACK ARCHIVE ON ftest TO test;
- ▶ GRANT EXECUTE ON DBMS_FLASHBACK_ARCHIVE TO test;
- ▶ GRANT EXECUTE ON DBMS_FLASHBACK TO test;
- ▶ GRANT FLASHBACK ANY TABLE TO PUBLIC;
- ▶ GRANT EXECUTE ON emp_policy_func TO PUBLIC;

**** Careful VPDs – Create equivalency policy for Flashback Archive History Table ****

A Thing About Machines – Users need Flashback!

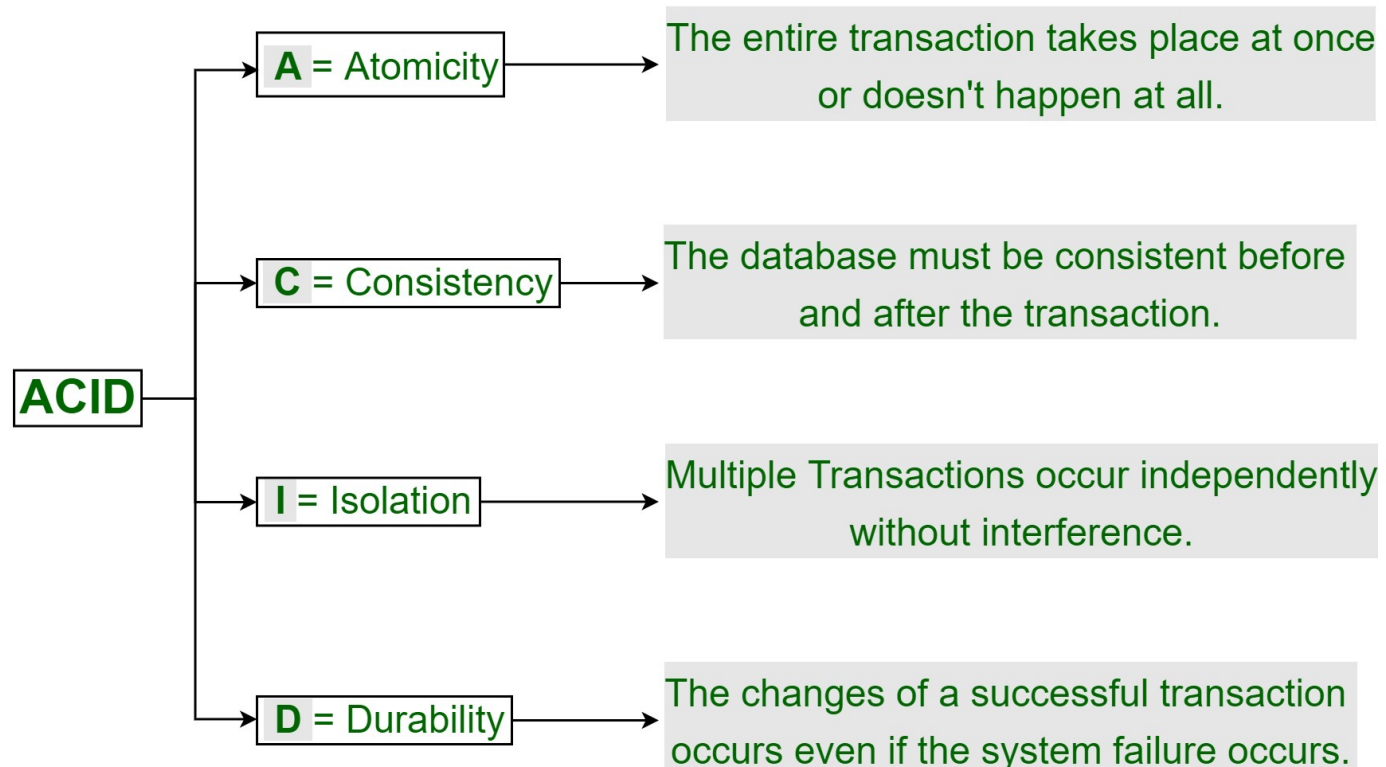


HELPDESK
HAVE U TRIED TURNING IT OFF
AND ON AGAIN?





ACID Properties in DBMS



Choosing Transaction Isolation Levels – FYI Only

- ▶ There is a **trade-off between concurrency (transaction throughput) and consistency**. For environments with many concurrent users rapidly submitting transactions, consider expected transaction arrival rate, response time demands, and required degree of consistency.
- ▶ **READ COMMITTED** isolation can **provide considerably more concurrency with a somewhat increased risk of inconsistent results** (from unrepeatable and phantom reads) for some transactions.
- ▶ **SERIALIZABLE** isolation provides somewhat **more consistency** (by protecting against phantoms and unrepeatable reads), which might be important where a read/write transaction runs a query more than once.

Table 9-6 Comparison of READ COMMITTED and SERIALIZABLE Transactions

Operation	READ COMMITTED	SERIALIZABLE
Dirty write	Not Possible	Not Possible
Dirty read	Not Possible	Not Possible
Unrepeatable read	Possible	Not Possible
Phantom read	Possible	Not Possible
Compliant with ANSI/ISO SQL 92	Yes	Yes
Read snapshot time	Statement	Transaction
Transaction set consistency	Statement level	Transaction level
Row-level locking	Yes	Yes
Readers block writers	No	No
Writers block readers	No	No
Different-row writers block writers	No	No
Same-row writers block writers	Yes	Yes
Waits for blocking transaction	Yes	Yes
Subject to "cannot serialize access" error	No	Yes
Error after blocking transaction terminates	No	No
Error after blocking transaction commits	No	Yes



Comparison of READ COMMITTED and SERIALIZABLE Transactions

Other New Features – FYI Only

- ▶ JSON Schema
- ▶ Ability to Audit Object Actions at the Column Level for Tables and Views
- ▶ Enhancements to RADIUS Configuration
- ▶ Local Rolling Patching & Single-Server Rolling Patching (RAC)
- ▶ Oracle RAC on Kubernetes & Sequence Optimizations in Oracle RAC
- ▶ Annotations – Define Metadata for Database Objects
- ▶ Ubiquitous Search With **DBMS_SEARCH** Packages
- ▶ The new DBMS_SEARCH PL/SQL package allows the **indexing of multiple schema objects in a single index**. You can add a set of tables, external tables, or views as data sources into this index. All the columns in the specified sources are indexed and available for a full-text search.



Miscellaneous New Features – FYI Only

- ▶ [Aggregation](#) over interval data types
- ▶ Enhancements for blockchain tables
- ▶ 1024 byte passwords
- ▶ Sagas for microservices / Lock-free reservations
- ▶ Transparent application continuity
- ▶ Asynchronous programming
- ▶ Ubiquitous search with [dbms_search](#)
- ▶ SQL Property Graph Query Language (ISO SQL/PQL standard)
- ▶ SQL Property Graph Query Language (ISO SQL/PQL standard)
- ▶ New system privilege table retention.
- ▶ True Cache
- ▶ Real-time SQL Plan Management
- ▶ Read-only Per-PDB Standby
- ▶ In-database SQL firewall
- ▶ OKafka
- ▶ Oracle text indexes with automatic maintenance

New Initialization Parameters, Tables, and some V\$ Views

- The following initialization parameters are new in Oracle Database 23c:

- ALLOW_LEGACY_RECO_PROTOCOL
- BLOCKCHAIN_TABLE_RETENTION_THRESHOLD
- DB_FLASHBACK_LOG_DEST
- DB_FLASHBACK_LOG_DEST_SIZE
- GROUP_BY_POSITION_ENABLED
- IDENTITY_PROVIDER_CONFIG
- IDENTITY_PROVIDER_TYPE
- JSON_EXPRESSION_CHECK
- MAX_COLUMNS
- MAX_SAGA_DURATION
- MULTILINGUAL_ENGINE
- PLSQL_IMPLICIT_CONVERSION_BOOL
- RESULT_CACHE_AUTO_BLOCKLIST
- SAGA_HIST_RETENTION
- SQL_ERROR_MITIGATION
- SQL_HISTORY_ENABLED
- SQL_TRANSPILER
- TABLESPACE_ENCRYPTION
- TXN_AUTO_ROLLBACK_HIGH_PRIORITY_WAIT_TARGET
- TXN_AUTO_ROLLBACK_MEDIUM_PRIORITY_WAIT_TARGET
- TXN_AUTO_ROLLBACK_MODE
- TXN_PRIORITY

- The following tables are new in Oracle Database 23c:

- DBMS_KAFKA_APPLICATIONS
- DBMS_KAFKA_CLUSTERS
- DBMS_KAFKA_LOAD_METRICS
- DBMS_KAFKA_MESSAGES
- DBMS_KAFKA_OPS
- DBMS_KAFKA_OPS_RESULTS
- DBMS_KAFKA_PARTITIONS
- DBMS_KAFKA_SEC_ALLOWED_PROPERTIES

- The following dynamic performance views are new in Oracle Database 23c:

- V\$DATAPUMP_SESSIONWAIT_INFO
- V\$DG_BROKER_ROLE_CHANGE
- V\$ENABLEDSCHEMAPRIVS
- V\$FAST_START_FAILOVER_CONFIG
- V\$FLASHBACK_LOG_DEST
- V\$OFS_THREADS
- V\$PX_PROCESS_DETAIL
- V\$SQL_HISTORY
- V\$TDM_STATS

23ai Deprecated & Desupported Features - FYI

1.2 Deprecated Features

The following features are deprecated in Oracle Database 23c, and may be desupported in a future release:

- The `ENCRYPT_NEW_TABLESPACES` initialization parameter
- The `V$FS_FAILOVER_STATS` view
- The `V$PQ_SLAVE` view

See Also:

[Oracle Database Upgrade Guide](#) for a complete list of deprecated features for this release

1.3 Desupported Features

The following features are desupported in Oracle Database 23c:

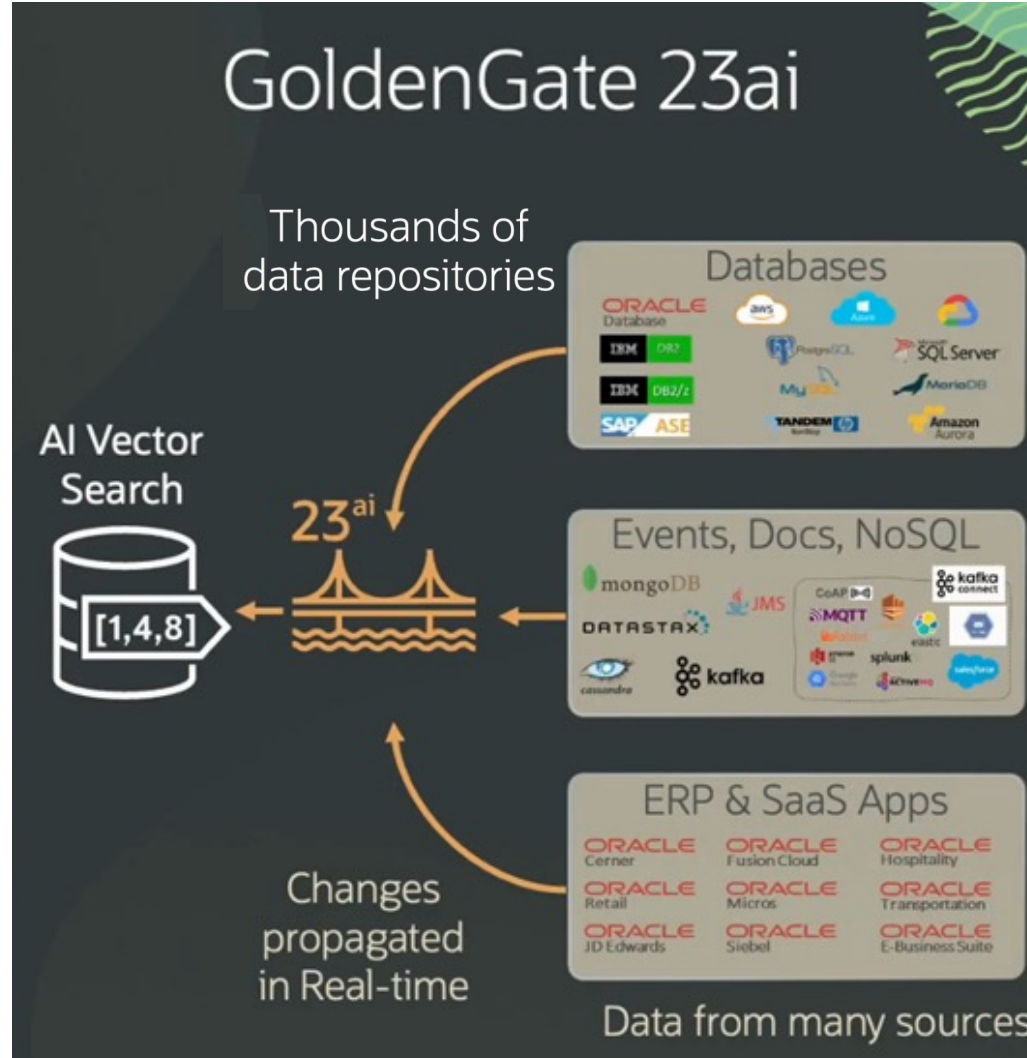
- The `V$EMX_USAGE_STATS` view

Other DB Announcements: GoldenGate Free up to 20G

<https://blogs.oracle.com/post/oracle-goldengate-free>



(Up to 20G)



Oracle Database Security

Built over MANY years... Best Feature!!

Encryption in the Silicon (M7)

Oracle Multi-Tenant Security (PDBs)



Oracle Audit Vault

Oracle Database Vault

DB Security Evaluation #19

Transparent Data Encryption

EM Configuration Scanning

Fine Grained Auditing (9i)

Secure application roles

Client Identifier / Identity propagation

Oracle Label Security (2000)

Proxy authentication

Enterprise User Security

Global roles

Virtual Private Database (8i)

Database Encryption API

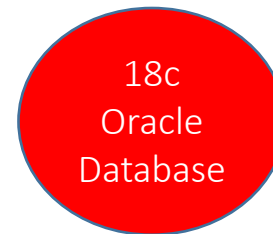
Strong authentication (PKI, Kerberos, RADIUS)

Native Network Encryption (Oracle7)

Database Auditing

Government customer

1977



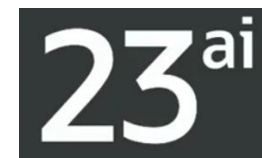
18c: User-defined Master Encryption Key (bring your own key - software keystores)



19c: Oracle Database supplied schema-only accounts have passwords removed (not sample accounts)



Global Fault-Tolerant Key Vault & Oracle Data Safe



In-Database SQL Firewall blocks unauthorized SQL & SQL Injection Attacks

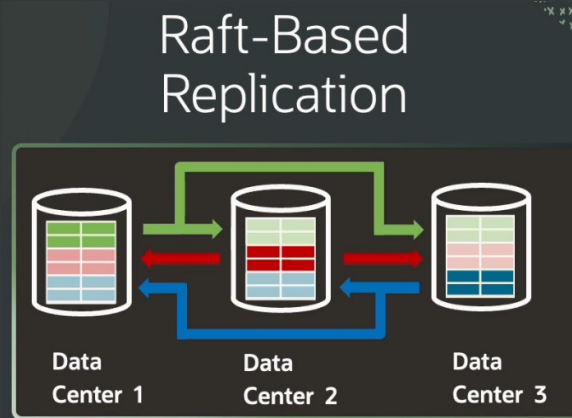
2024 +



Prepare for Just Announced Oracle 23ai Features*



Make **Mission-Critical Data** architecturally simple and scalable



True Cache



In-Database SQL Firewall

Quorum based fast failover



Real-Time SQL Plan Management



RAC, Exadata, Data Guard Simplicity and Scalability



Priority Transactions

Prepare for Just Announced Oracle 23ai Features*



Make **Data for App Dev** architecturally simple and scalable
Developer-centric data and transactions



JSON-Relational
Unification



Graph-Relational
Unification



JavaScript
Stored Procedures



Data Intent
Language



Lock-free Consistent Updates,
Long-running Transactions



Transactional
Microservices

ORACLE®
Database Cloud

The World's First Autonomous Database

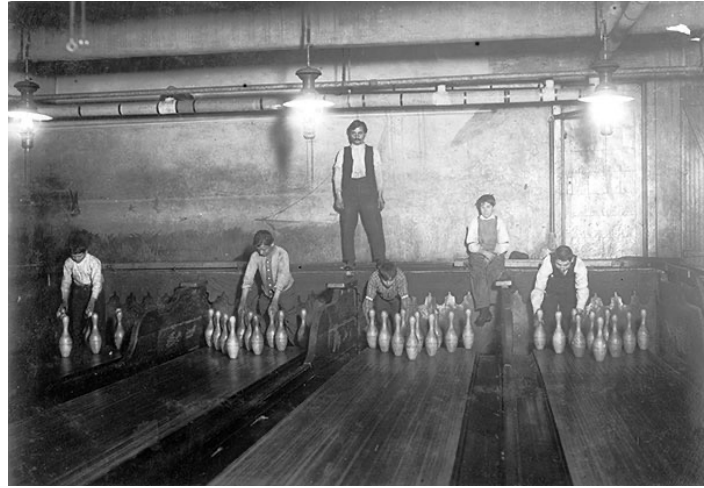


A Robot may not look one!

*Robots that Manage a Database (ADW/ATP)!


*Robots that secure a system and use ML & AI

The Obsolete Man





FYI



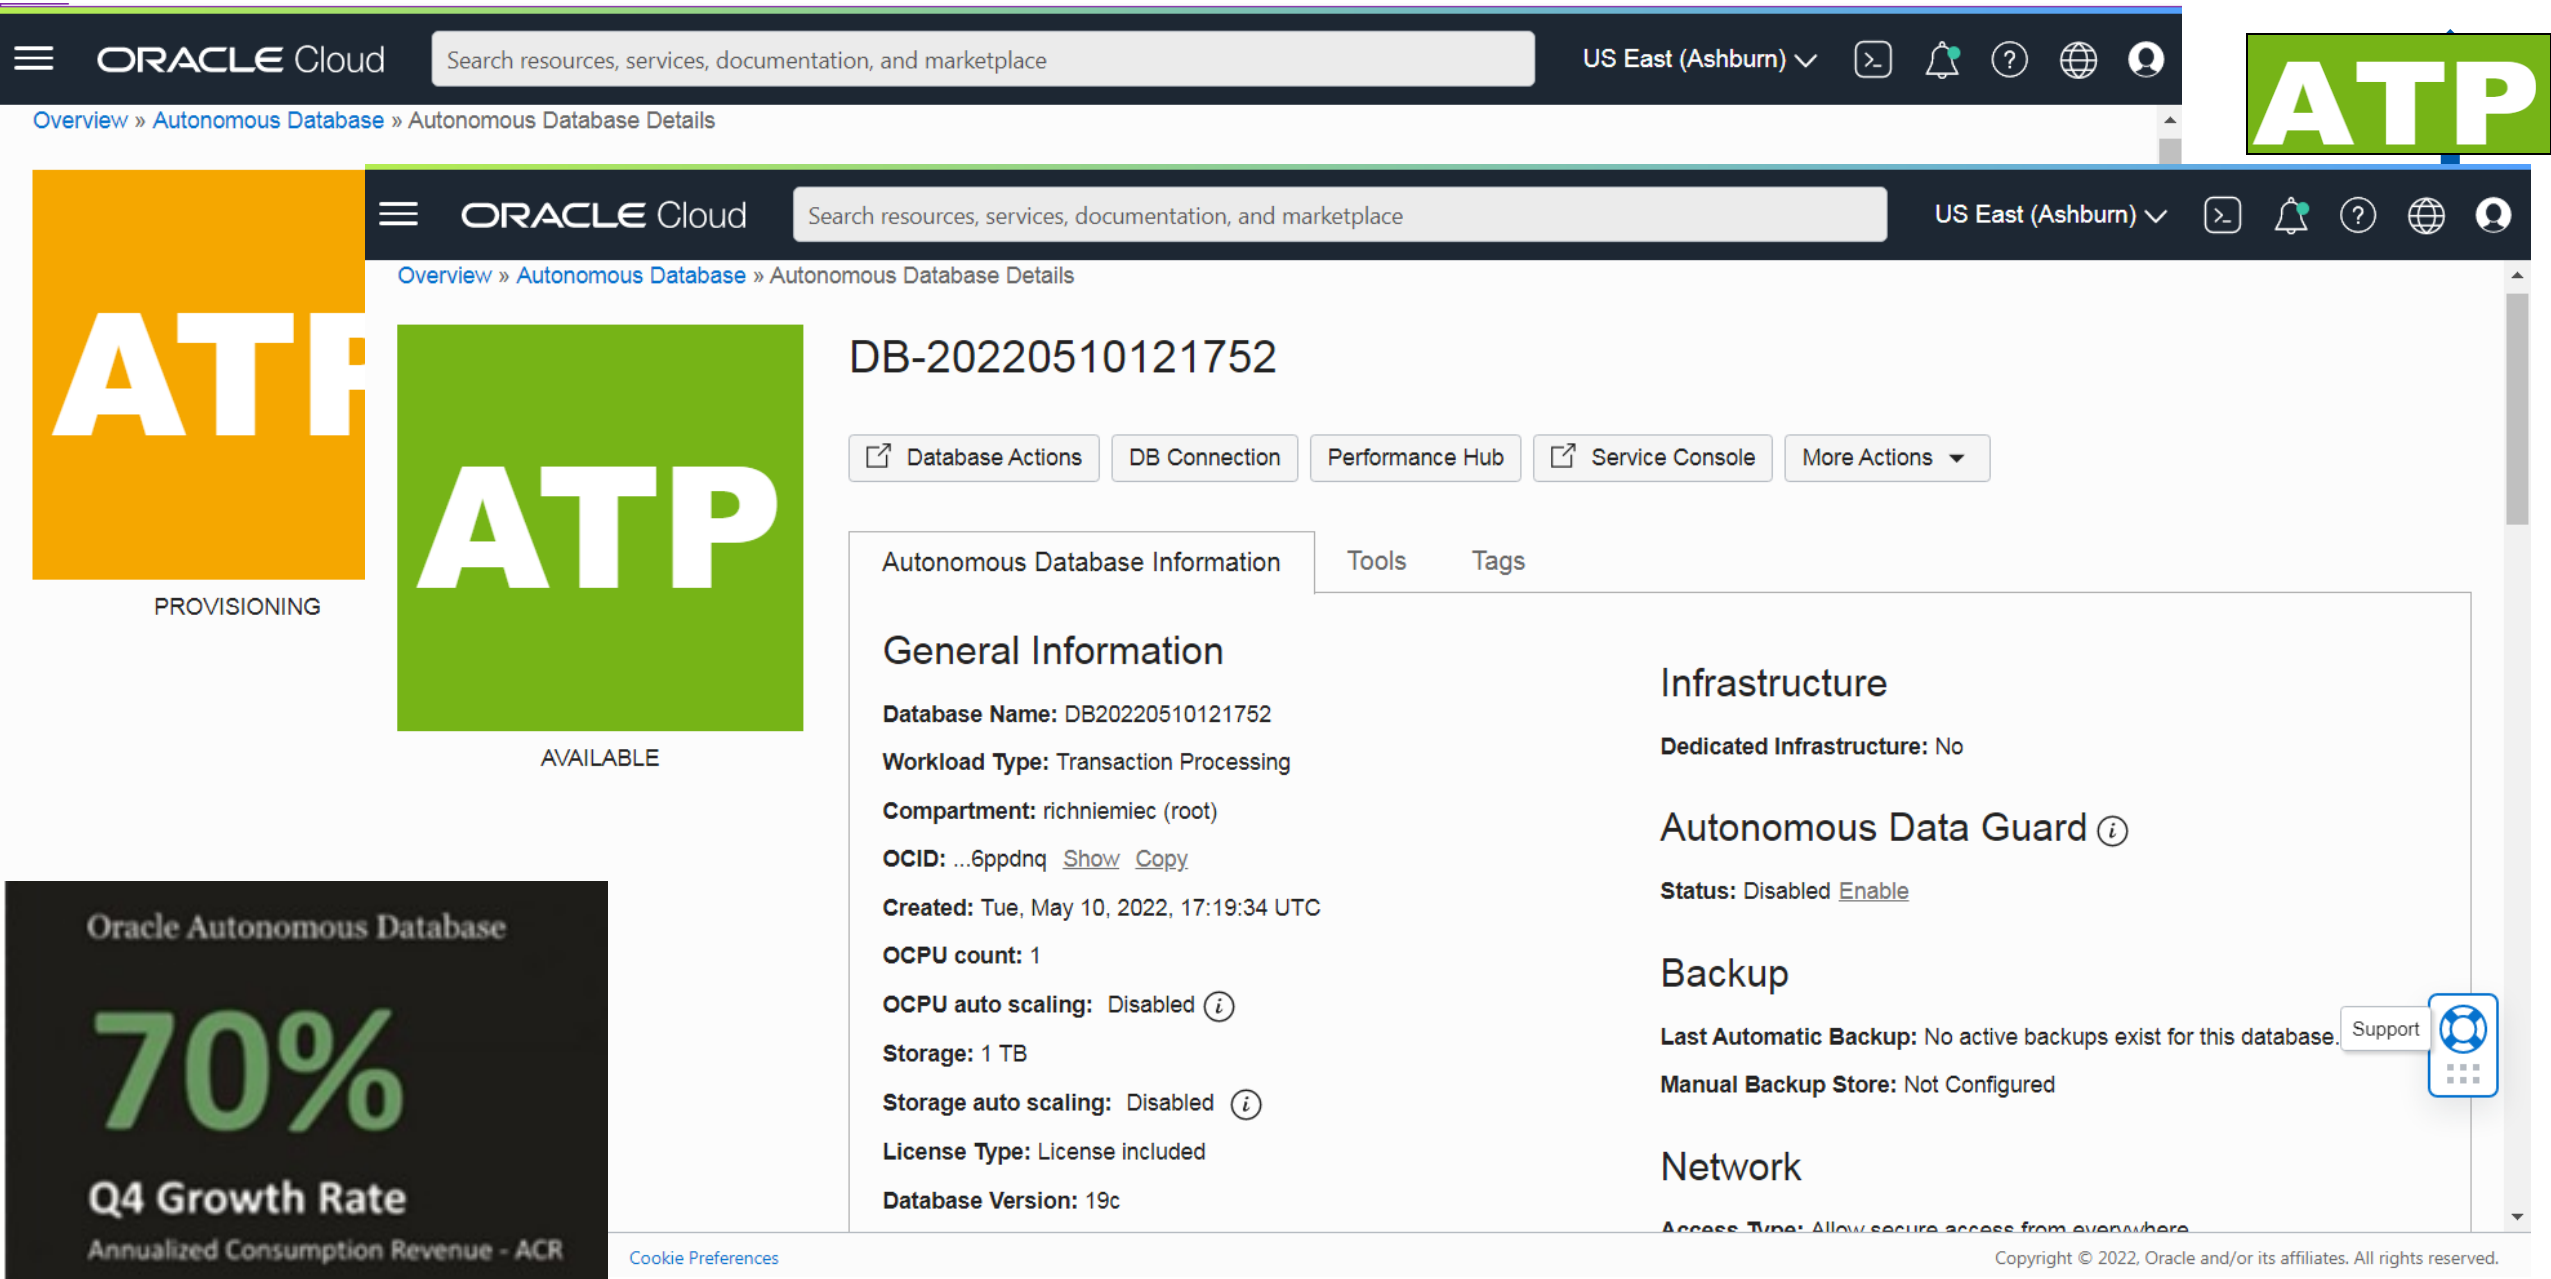
Adaptive
Machine Learning

The Autonomous Database & the DBA

Autonomous Databases into the future:

- ▶ Who ensures database is tuned before it gets to the Cloud?
- ▶ Who ensures the cloud vendor is charging correctly?
- ▶ Who ensures the backup, security, or recovery is correct?
- ▶ Who decides what kind of service the databases will be?
- ▶ Who will build the policies for those autonomous databases?
- ▶ Who will have the knowledge to decide or estimate the cost of these services?
- ▶ Who decides the complex IT Infrastructure when we have more options & vendors?
- ▶ The answer is obvious: A DBA, but not a simple DBA; A DBA that has evolved with all this new generation of databases on Cloud. On-prem 18c/19c has NO effect on DBA.

ATP – Provisioning a Database (2 minutes)



The screenshot displays the Oracle Cloud console interface for provisioning an Autonomous Database (ATP). The top navigation bar includes the Oracle Cloud logo, a search bar, and the region 'US East (Ashburn)'. The breadcrumb trail indicates the path: Overview » Autonomous Database » Autonomous Database Details.

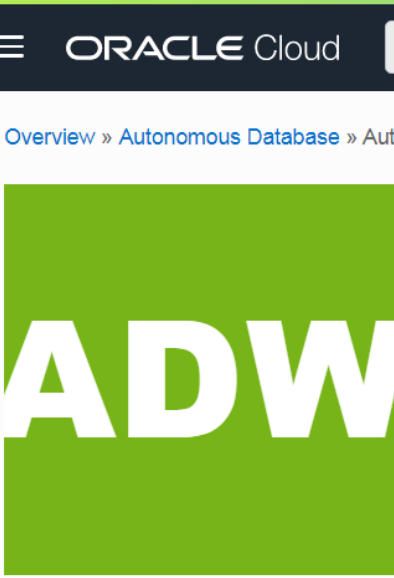
On the left side, there are two ATP status indicators: a yellow box labeled 'ATP PROVISIONING' and a green box labeled 'ATP AVAILABLE'. Below these is a promotional banner for Oracle Autonomous Database, highlighting a 70% Q4 Growth Rate in Annualized Consumption Revenue - ACR.

The main content area shows the details for a specific database instance, 'DB-20220510121752'. It features a toolbar with buttons for 'Database Actions', 'DB Connection', 'Performance Hub', 'Service Console', and 'More Actions'. Below the toolbar, there are tabs for 'Autonomous Database Information', 'Tools', and 'Tags'. The 'Autonomous Database Information' tab is active, displaying the following details:

- General Information**
 - Database Name: DB20220510121752
 - Workload Type: Transaction Processing
 - Compartment: richniemiec (root)
 - OCID: ...6ppdnq [Show](#) [Copy](#)
 - Created: Tue, May 10, 2022, 17:19:34 UTC
 - OCPU count: 1
 - OCPU auto scaling: Disabled ⓘ
 - Storage: 1 TB
 - Storage auto scaling: Disabled ⓘ
 - License Type: License included
 - Database Version: 19c
- Infrastructure**
 - Dedicated Infrastructure: No
- Autonomous Data Guard ⓘ**
 - Status: Disabled [Enable](#)
- Backup**
 - Last Automatic Backup: No active backups exist for this database. [Support](#)
 - Manual Backup Store: Not Configured
- Network**
 - Access Type: Allow secure access from everywhere

At the bottom of the page, there is a 'Cookie Preferences' link on the left and a copyright notice: 'Copyright © 2022, Oracle and/or its affiliates. All rights reserved.' on the right.

ADW - Provision Database (1 minute 20 seconds)



DB-20220510174437

- Database Actions
- DB Connection
- Performance Hub
- Service Console
- More Actions

Autonomous Database Information

Tools Tags

General Information

Database Name: DB20220510174437
Workload Type: Data Warehouse
Compartment: richniemiec (root)
OCID: ...iabjpa [Show](#) [Copy](#)
Created: Tue, May 10, 2022, 22:57:35 UTC
OCPU count: 1
OCPU auto scaling: Disabled ⓘ
Storage: 1 TB
Storage auto scaling: Disabled ⓘ
License Type: License included
Database Version: 19c

Infrastructure

Dedicated Infrastructure: No

Autonomous Data Guard ⓘ

Status: Disabled [Enable](#)

Backup

Last Automatic Backup: No active backups exist for this database.

Manual Backup Store: Not Configured

Network

Things to Focus On! (**DA = DBA Alive**)

- ▶ **Data Administration** (just change your title)
- ▶ **Manager / Business**
- ▶ **Cloud & Autonomous**
- ▶ **Big Data**
- ▶ **IOT**
- ▶ **Security**
- ▶ **Network Administration**
- ▶ **Chatbots**
- ▶ **AI: ML, NLP, Robotics**



**Get the Autonomous DBA
Some of the Work!**

Machine Learning connection to ADW/ATP



DB-20220510121752

Database Actions

DB Connection

Performance Hub

Service Console

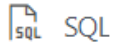
More Actions

ATP

The screenshot displays the Oracle Machine Learning interface. At the top, there's a navigation bar with 'Database Actions' circled in red. Below it, a grid of 'Example Templates' is shown. Each template card includes a title, a brief description, the author (Oracle), the date added, and tags. The templates are:

- OML4R Clustering EM**: This notebook builds and applies a clustering mod...
- OML4R Clustering KM**: This notebook builds and applies a clustering mod...
- OML4R Clustering OC**: This notebook builds and applies a clustering mod...
- OML4Py Classification SVM**: This notebook builds and applies a Support Vector...
- OML4Py Clustering EM**: This notebook builds and applies a clustering mod...
- OML4Py Clustering KM**: This notebook builds and applies a clustering mod...
- OML4Py Data Cleaning Duplicates Re...**: This notebook illustrates removing duplicate rows ...
- OML4SQL Classification DT**: This notebook builds and applies a Decision Tree Classifica...
- OML4SQL Classification GLM**: This notebook builds and applies a Generalized Linear Clas...
- OML4SQL Classification NB**: This notebook builds and applies a Naïve Bayes Classificati...
- OML4SQL Classification NN**: This notebook builds and applies a Neural Network Classifi...

Developm



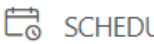
SQL

Execute que
create datab



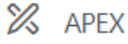
LIQUIBA

Track schem



SCHEDULE

Schedule, m
tasks



APEX

Machine Learning Process

▶ **First:** Clear *Business Problem* to Solve

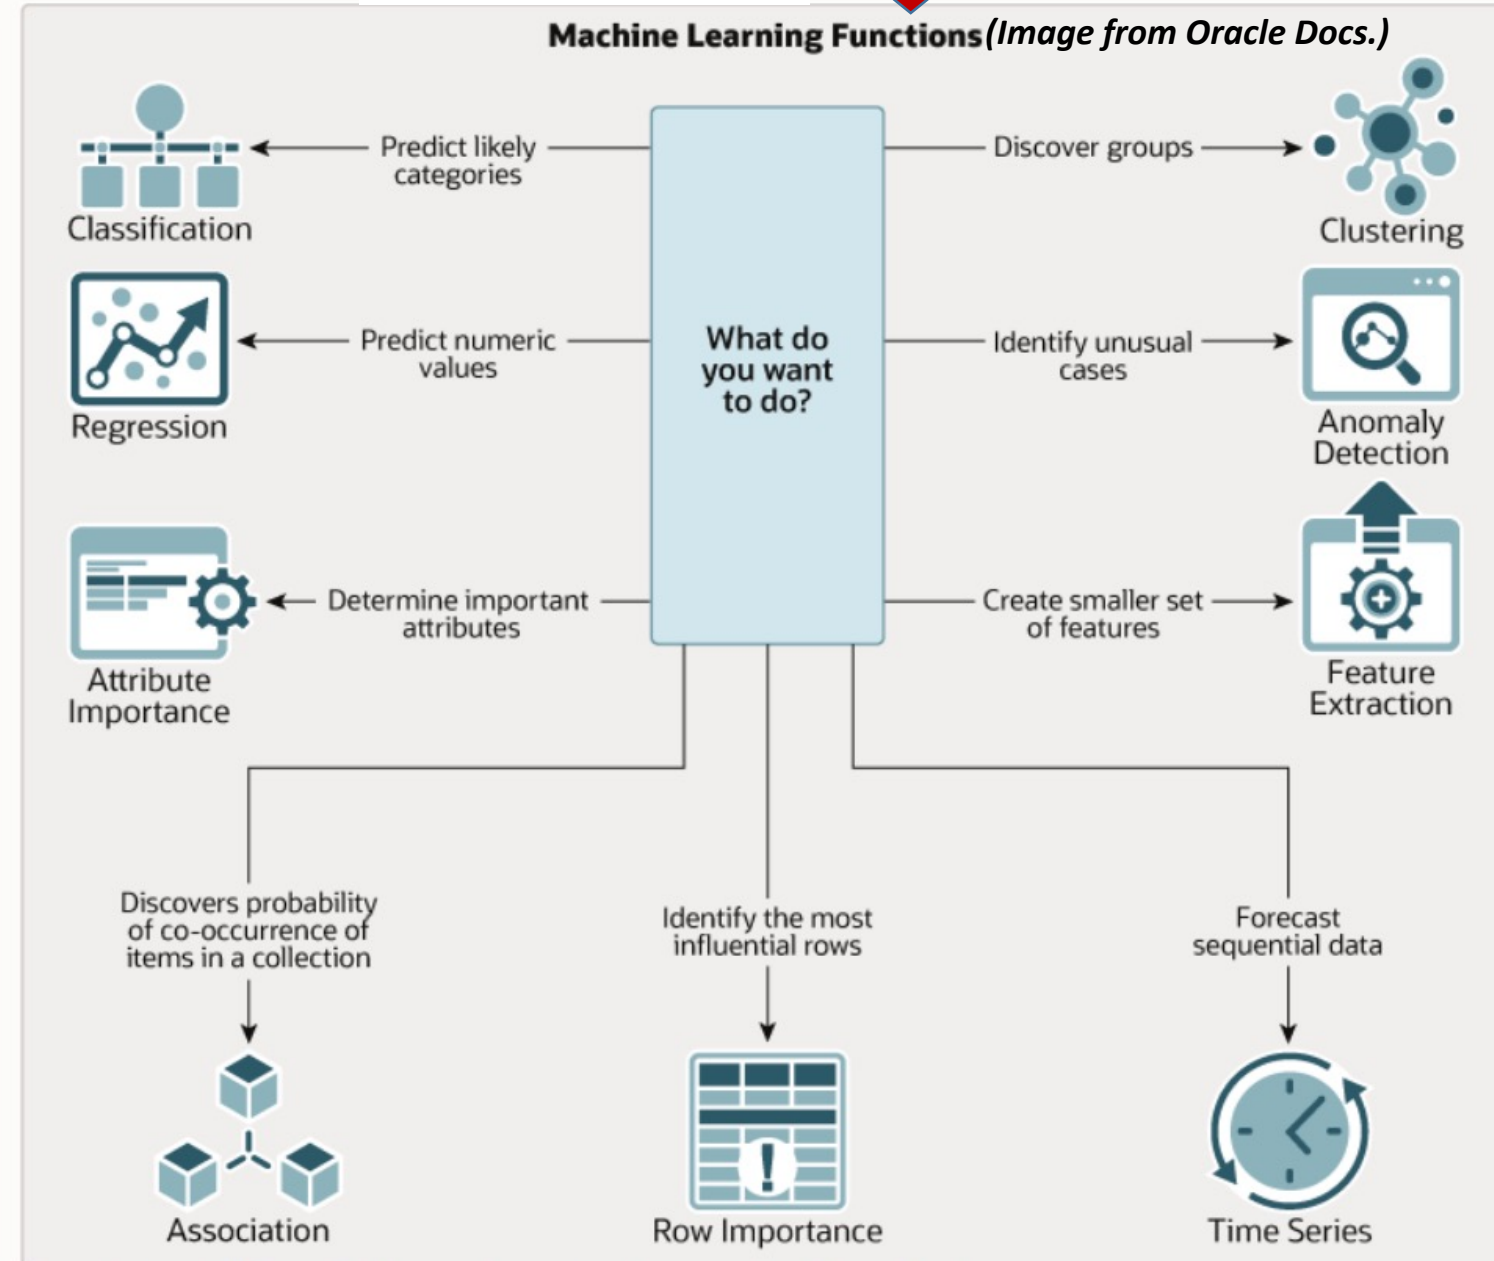
▶ **Second:** *Function* to Perform

▶ **Third:** *Algorithms* to use

▶ **First:** Build/*Train the Model*: When you build it use about 60% of your data.

▶ **Second:** *Test/Score Model* for accuracy/precision using about 40% of data.

Compare Algorithms!



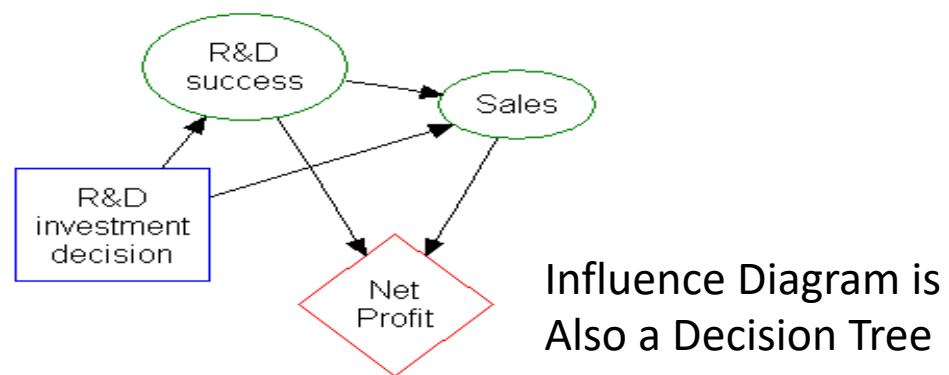
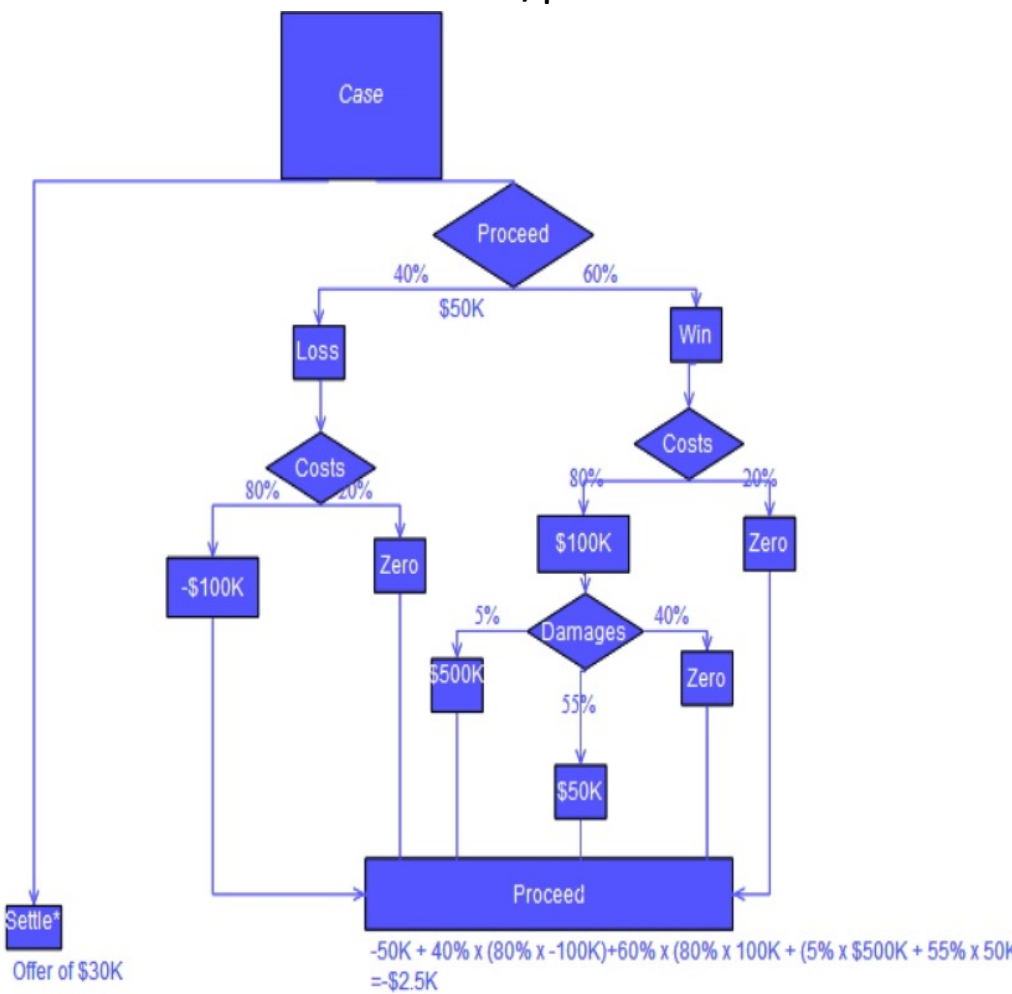
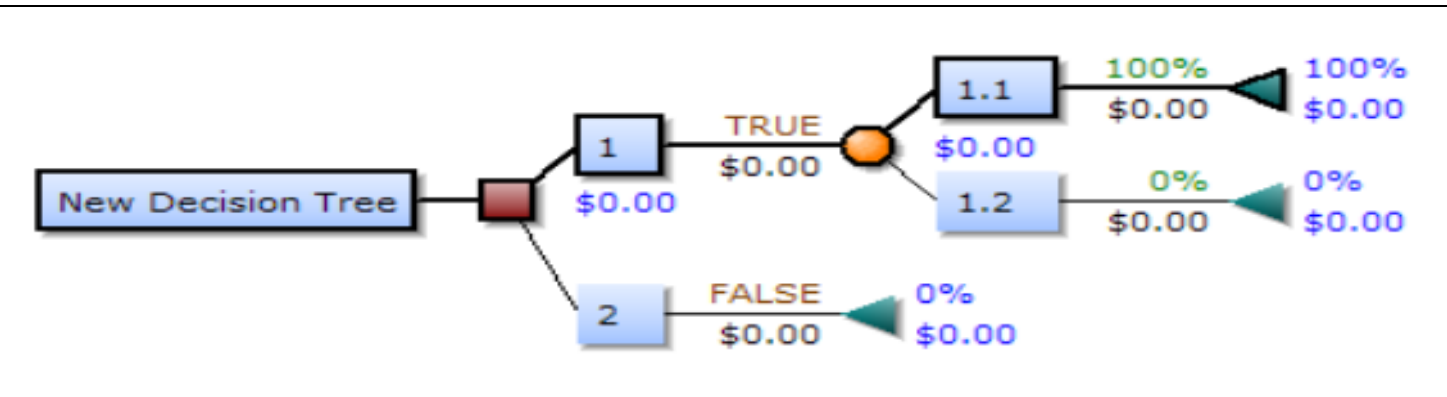
Decision Tree Algorithm (ML Classifier)



A decision tree consists of three types of nodes:^[1]

1. Decision nodes – typically represented by squares
2. Chance nodes – typically represented by circles
3. End nodes – typically represented by triangles

Flow Chart calculates whether to settle a case or not based on costs/probabilities



OML Oracle Data Mining SQL Sample (PARTIAL)

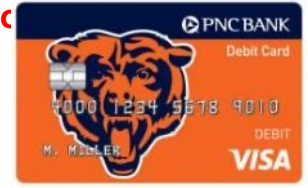


dmdtdemo.sql - DBMS_DATA_MINING package - Decision Tree

```
-- Given demographic data about a set of customers, predict the  
-- customer response to an affinity card program using a classifier  
-- based on Decision Trees algorithm.
```

```
INSERT INTO dt_sh_sample_settings VALUES  
(dbms_data_mining.algo_name, dbms_data_mining.algo_decision_tree);  
--  
-- CREATE A NEW MODEL  
-- Build a DT model
```

```
BEGIN  
DBMS_DATA_MINING.CREATE_MODEL(  
  model_name => 'DT_SH_Clas_sample',  
  mining_function => dbms_data_mining.classificatio  
  data_table_name => 'mining_data_build_v',  
  case_id_column_name => 'cust_id',  
  target_column_name => 'affinity_card',  
  settings_table_name => 'dt_sh_sample_settings');  
END;  
/
```

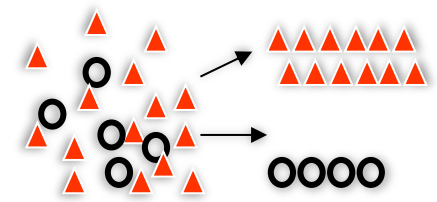


```
SELECT T.cust_id, S.prediction, S.probability, S.cost  
FROM (SELECT cust_id,  
  PREDICTION_SET(dt_sh_clas_sample COST MODEL USING *) pset  
FROM mining_data_apply_v  
WHERE cust_id < 100011) T,  
TABLE(T.pset) S  
ORDER BY cust_id, S.prediction;
```

Script Output x

Task completed in 3.075 seconds

CUST_ID	PREDICTION	PROBABILITY	COST
100001	0	.966183575	.270531401
100001	1	.033816425	.966183575
100002	0	.740384615	2.076923077
100002	1	.259615385	.740384615
100003	0	.909090909	.727272727
100003	1	.090909091	.909090909
100004	0	.909090909	.727272727
100004	1	.090909091	.909090909
100005	0	.272357724	5.821138211
100005	1	.727642276	.272357724
100006	0	1.000000000	.000000000
100006	1	.000000000	1.000000000
100007	0	.909090909	.727272727
100007	1	.090909091	.909090909
100008	0	.909090909	.727272727
100008	1	.090909091	.909090909
100009	0	.272357724	5.821138211
100009	1	.727642276	.272357724
100010	0	.675965665	2.592274678
100010	1	.324034335	.675965665



Machine Learning & AI - Oracle's Built-In Algorithms

Oracle Machine Learning Algorithms



CLASSIFICATION

Naïve Bayes
Logistic Regression (GLM)
Decision Tree
Random Forest
Neural Network
Support Vector Machine
Explicit Semantic Analysis
XGBoost*



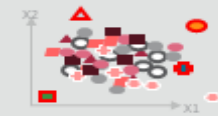
CLUSTERING

Hierarchical K-Means
Hierarchical O-Cluster
Expectation Maximization (EM)



ANOMALY DETECTION

One-Class SVM
MSET-SPRT*



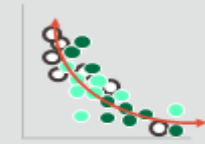
TIME SERIES

Forecasting - Exponential Smoothing
Includes popular models
e.g. Holt-Winters with trends,
seasonality, irregularity, missing data



REGRESSION

Linear Model
Generalized Linear Model
Support Vector Machine (SVM)
Stepwise Linear regression
Neural Network
XGBoost*



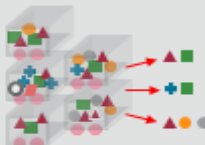
ATTRIBUTE IMPORTANCE

Minimum Description Length
Principal Comp Analysis (PCA)
Unsupervised Pair-wise KL Div
CUR decomposition for row & AI



ASSOCIATION RULES

A priori/ market basket



PREDICTIVE QUERIES

Predict, cluster, detect, features



SQL ANALYTICS

SQL Windows
SQL Patterns
SQL Aggregates

FEATURE EXTRACTION

Principal Comp Analysis (PCA)
Non-negative Matrix Factorization
Singular Value Decomposition (SVD)
Explicit Semantic Analysis (ESA)

TEXT MINING SUPPORT

Algorithms support text
Tokenization and theme extraction
Explicit Semantic Analysis (ESA) for
document similarity



STATISTICAL FUNCTIONS

Basic statistics: min, max,
median, stdev, t-test, F-test, Pearson's,
Chi-Sq, ANOVA, etc.



R & PYTHON * Coming soon

Third-party R & Python Packages
through Embedded Execution
Spark MLlib algorithm integration



MODEL DEPLOYMENT & MONITORING

SQL—1st Class Objects
Oracle RESTful API (ORDS)
OML Web Services (for Apps)



*Includes support for Partitioned Models, Transactional data and aggregations,
Unstructured data, Geo-spatial data, Graph data. etc,*

* New in **21^c**

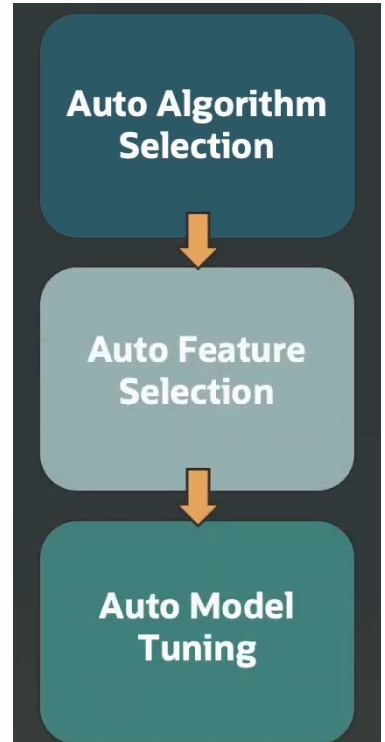


Oracle's latest list for 23ai (AutoML available since 21c)



Over 30 in-database parallel and scalable AI algorithms

There is also AutoML for Classification & Regression



Classification

- Decision Tree
- Explicit Semantic Analysis
- Logistic Regression (GLM)
- Naïve Bayes
- Neural Network
- Random Forest
- Support Vector Machine (SVM)
- XGBoost

Clustering

- Hierarchical K-Means
- Hierarchical O-Cluster
- Expectation Maximization

Row Importance

- CUR Decomposition

Ranking

- XGBoost

Regression

- Generalized Linear Model (GLM)
- Neural Network
- Support Vector Machine (SVM)
- Stepwise Linear regression
- XGBoost

Feature Extraction

- Principal Comp Analysis (PCA)
- Non-negative Matrix Factorization
- Singular Value Decomposition (SVD)
- Explicit Semantic Analysis (ESA)

Attribute Importance

- Minimum Description Length
- Random Forest
- Unsupervised Pairwise KL Divergence
- CUR decomposition for row & AI

Time Series

- Exponential Smoothing
- Multiple Time Series (23ai)
- Includes popular models e.g. Holt-Winters with trends, seasonality, irregular time series*

Anomaly Detection

- One-Class SVM
- MSET-SPRT
- Expectation Maximization (23c)

Association Rules

- A priori

Survival Analysis

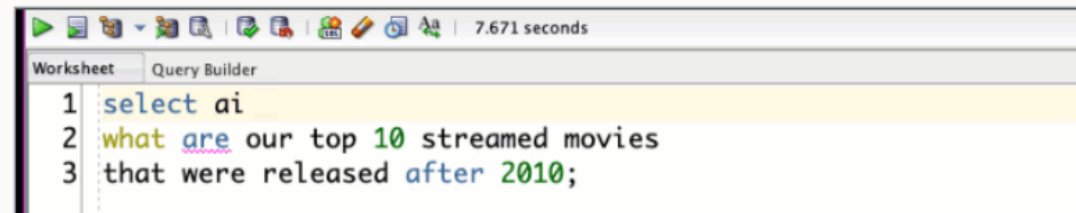
- XGBoost

Select AI with AutonomousDB & DBMS_CLOUD_AI Package

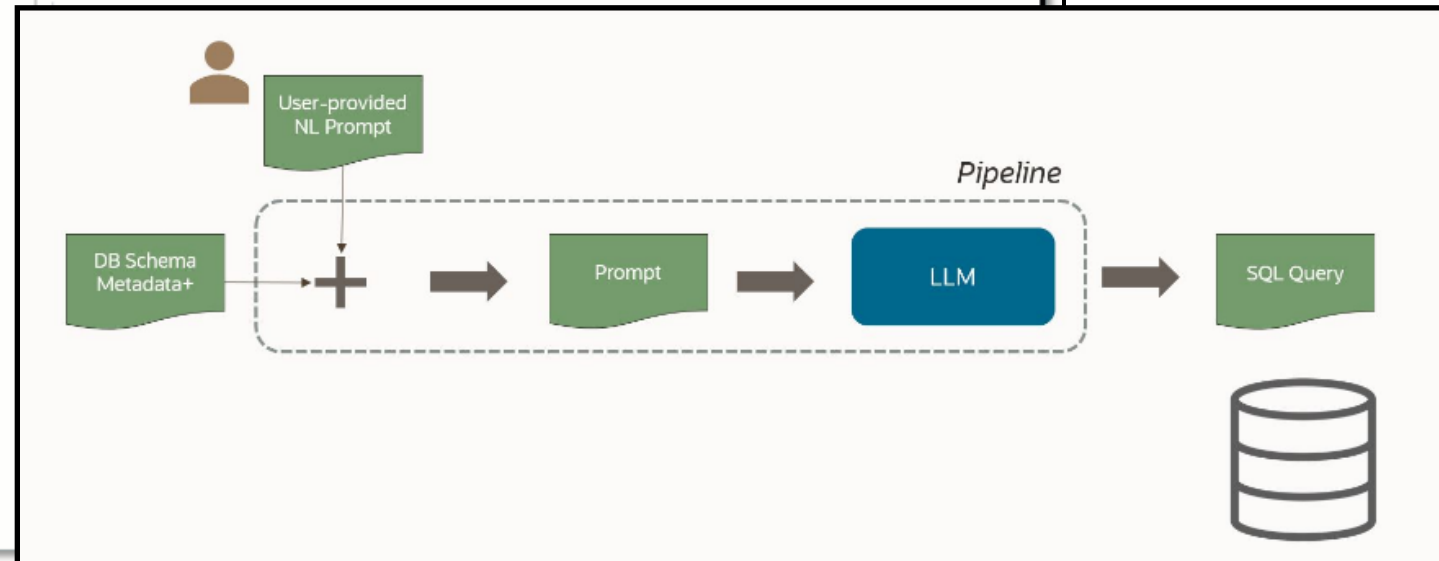
New | SQL Generation with Select AI

Available today on Oracle Autonomous Database

- Use natural language to query data with the help of LLMs
- Increase application developer productivity
- Enable non-technical users to query information from their database
- Invoke from SQL command line and PL/SQL function
- Inherit security and authentication of the database



```
Worksheet Query Builder 7.671 seconds
1 select ai
2 what are our top 10 streamed movies
3 that were released after 2010;
```



See: AutonomousDB Speaks “human” by Marty Gubar, Oracle
(Try it out with the Live Lab!)



Select AI with AutonomousDB

Worksheet Query Builder

```
1 select ai
2 what are our top 10 streamed movies
3 that were released after 2010;
```

Query Result x

All Rows Fetched: 10 in 7.978 seconds

MOVIE_TITLE	STREAM_COUNT
1 Avengers: Endgame	81074
2 Captain Marvel	33587
3 Star Wars Episode IX: The Rise of Skywalker	31331
4 Spider-Man: Far from Home	29021
5 Aladdin	24058
6 The Lion King	22627
7 Aquaman	20357
8 Avengers: Infinity War	18813
9 Toy Story 4	17426
10 Bohemian Rhapsody	16936

It's secure

- Select AI uses Oracle Database security to keep your data safe
- None of your data is shared with LLM providers when using OCI Generative AI

June 27, 2024

Worksheet Query Builder

```
1 select ai showsql
2 what are our top 10 streamed movies
3 that were released after 2010;
```

Query Result x Script Output x

Task completed in 7.063 seconds

RESPONSE

```
SELECT m.TITLE AS movie_title, COUNT(s.MOVIE_ID) AS stream_count
FROM MOVIESTREAM.MOVIE m
JOIN MOVIESTREAM.SALES_SAMPLE s ON m.MOVIE_ID = s.MOVIE_ID
WHERE m.OPENING_DATE > TO_DATE('2010-01-01', 'YYYY-MM-DD')
GROUP BY m.TITLE
ORDER BY stream_count DESC
FETCH FIRST 10 ROWS ONLY
```

ChatDB

What are the most popular devices used to watch action and adventure movies

Device	Number of Views
iphone	132,922.00
ipad	103,272.00
mac	88,779.00
pc	88,327.00
pixel	69,373.00
galaxy	68,936.00
oneplus	34,963.00
lenovo	34,746.00

Explore Explain

Ask a question

Ask Database

Add Select AI to APEX

Real-Application Cluster	Parallel SQL	Transactions
Analytics	Disaster Recovery	Security

Every mission-critical feature of Oracle Database works transparently with the new AI capabilities

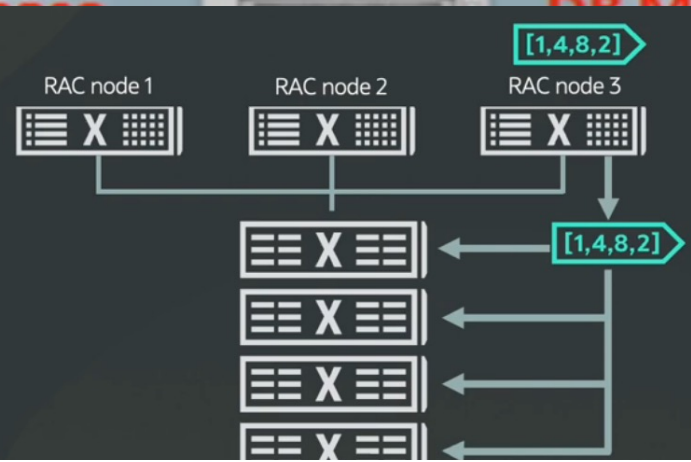
Exadata Cloud Machine Features

Exadata Cloud: Compatible, Scalable, Available, Secure
 Decades of Database Innovation Proven at Millions of Mission-Critical Deployments

	Multitenant
	In-Memory DB
	Real Application Clusters
	Active Data Guard
	Partitioning
	Advanced Compression
	Advanced Security, Label Security, DB Vault
	Real Application Testing
	Advanced Analytics, Spatial and Graph
	Management Packs for Oracle Database

Exadata 24ai Software

Vector search transparently offloaded to smart Exadata storage for much faster AI



All Oracle Database Innovations Exadata Cloud Machine Innovations

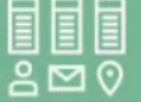
Offload SQL to Storage	
InfiniBand Fabric	
Smart Flash Cache, Log	PCI Flash
Storage Indexes	
Columnar Flash Cache	
Hybrid Columnar Compression	10:1
I/O Resource Management	
Network Resource Management	
In-Memory Fault Tolerance	
Exafusion	
Direct-to-Wire Protocol	

Oracle 23ai allows AI with Speed, Security and More...

*May 2024 Slide: Juan Loaiza

Oracle Database 23ai – Next Long-term Support Release

ORACLE Database 23ai



Data Use Case Domains

Schema Level Privileges




Read-Only Per-PDB Standby

Property Graphs



Real-time SQL Plan Management



SQL * ✓

Lock-Free Reservations



Microservice Support

JSON / Relational Duality



AI Vector Search

True Cache



SQL Firewall

Priority Transactions

Rolling Patching



JavaScript Stored Procedures



Developer Role

Shrink Tablespace

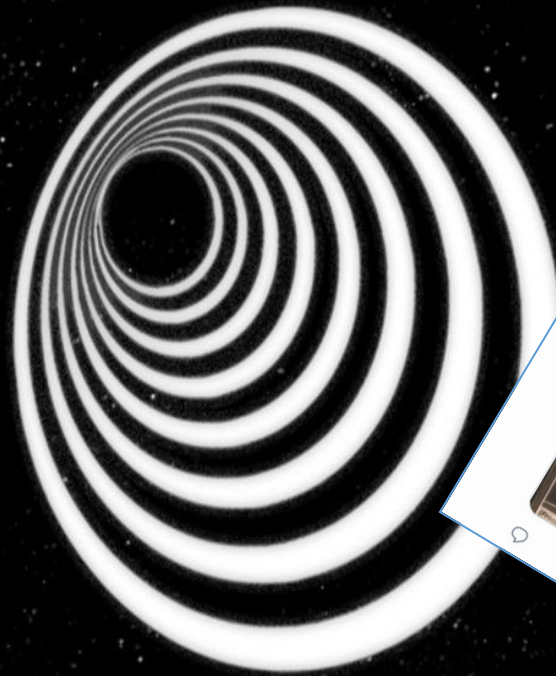
Boolean Datatype

Globally Distributed Database



THE TWILIGHT ZONE

"You unlock this door with the key of imagination. Beyond it is another dimension - a dimension of sound, a dimension of sight, a dimension of mind. You're moving into a land of both shadow and substance, of things and ideas. You've just crossed over into the Twilight Zone."
-Rod Serling



E=MC²



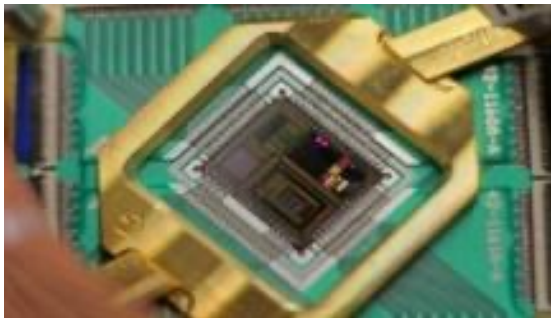
“Those who use things of the world should not become attached to them. For the world in its present form is passing away.”



64-Bit advancement of **Directly** addressable memory



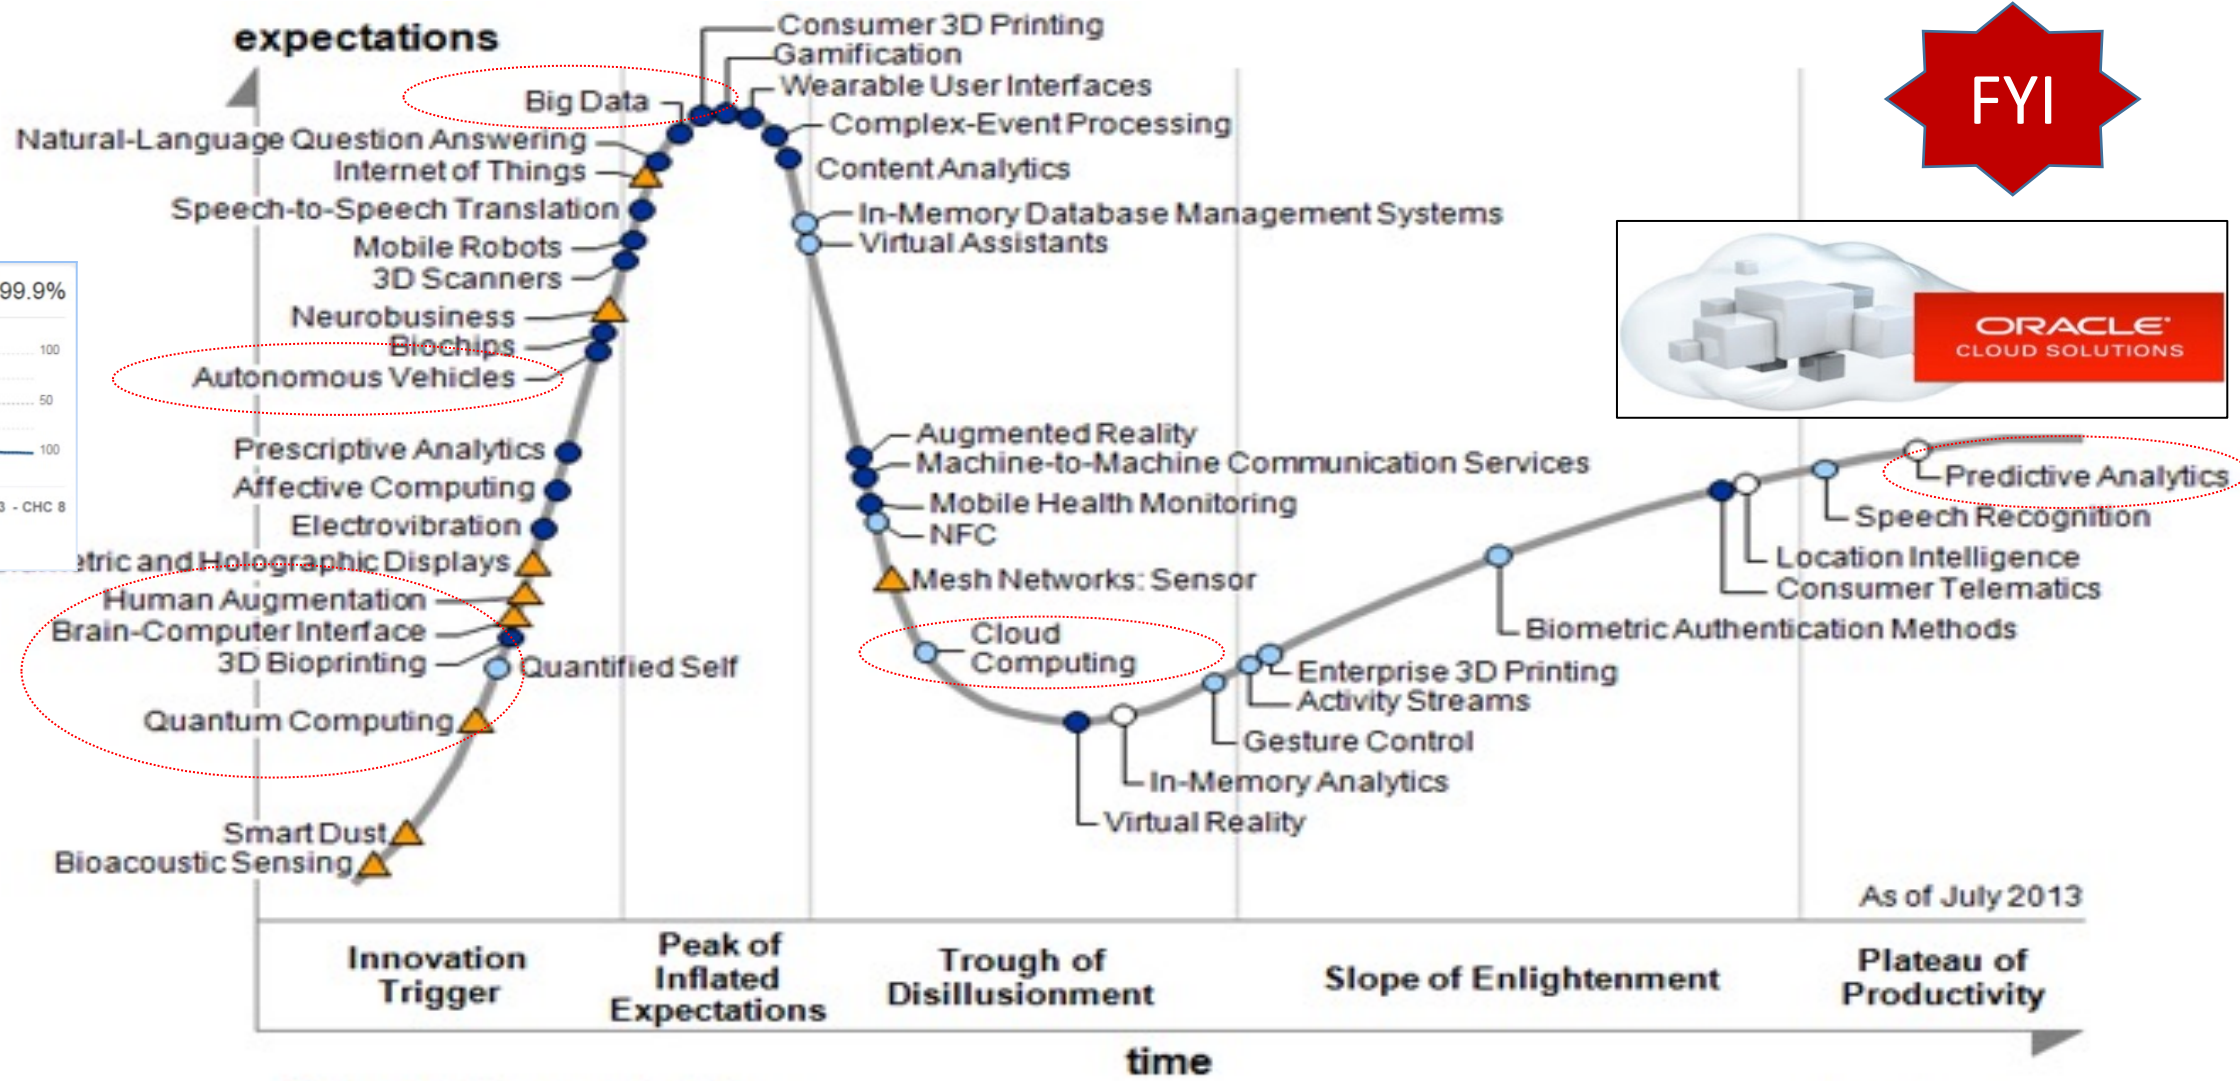
	<u>Address Direct</u>	<u>Indirect/Extended</u>
<u>4 Bit:</u>	16	(640)
<u>8 Bit:</u>	256	(65,536)
<u>16 Bit:</u>	65,536	(1,048,576)
<u>32 Bit:</u>	4,294,967,296	
<u>64 Bit:</u>	18,446,744,073,709,551,616	
<u>128 Bit:</u>	3.4 x 10e+38	



**Technology Feels Like It's Accelerating
— Because It Actually Is**
By Alison E. Berman and Jason Dorrier - Mar 22, 2016 ◀ 40,547

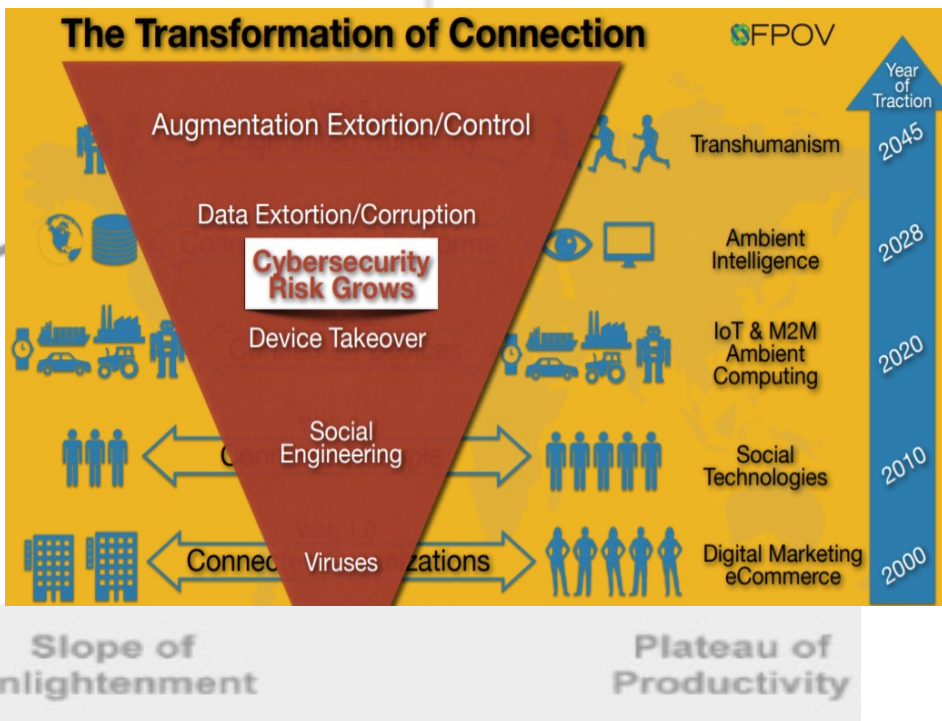
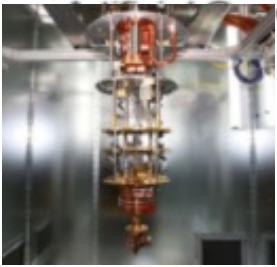
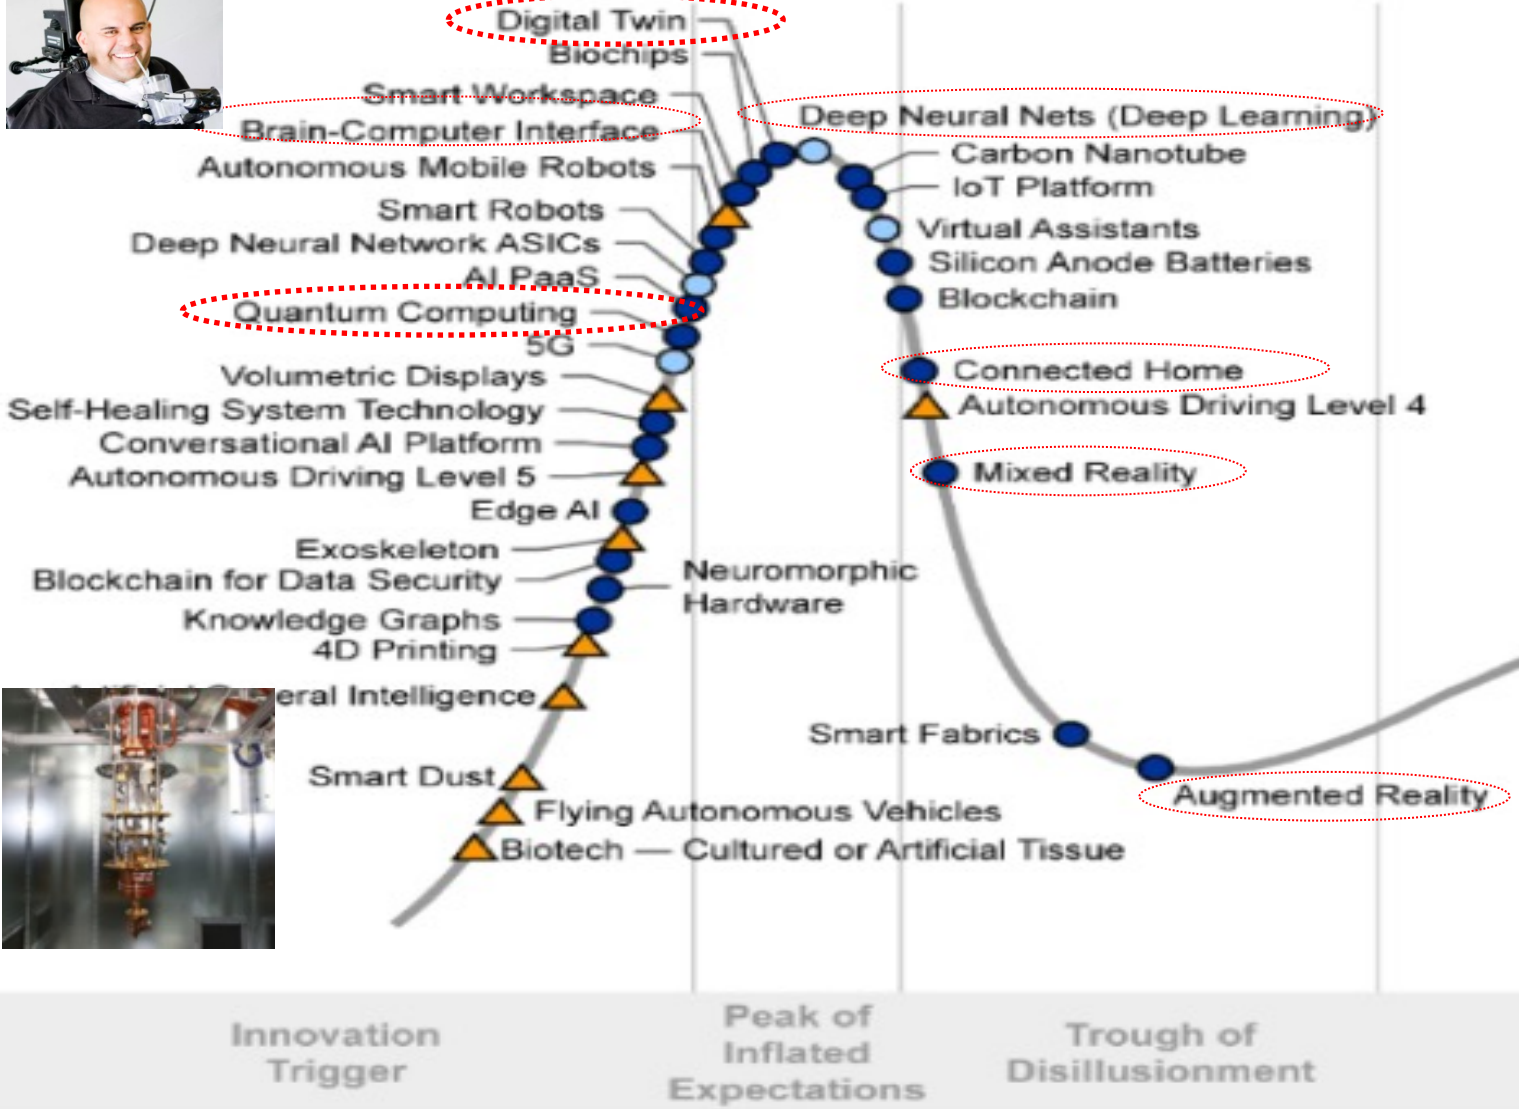
<u>Shift</u>	<u>Increase (Result)</u>
8-bit to 16-bit	1 mph (Windows)
16-bit to 32-bit	65K mph (Internet)
32-bit to 64-bit	300T mph (Robotics/ML...etc)
64-bit to 128-bit	5T*T*B mph (Implants/4D/AI)

Tech Trends - Gartner Hype Cycle 2013: All about Tech ...



Gartner Hype Cycle August 2018 - All about Robotics ...

All about Tech Creating a New Reality (just 5 years later!)



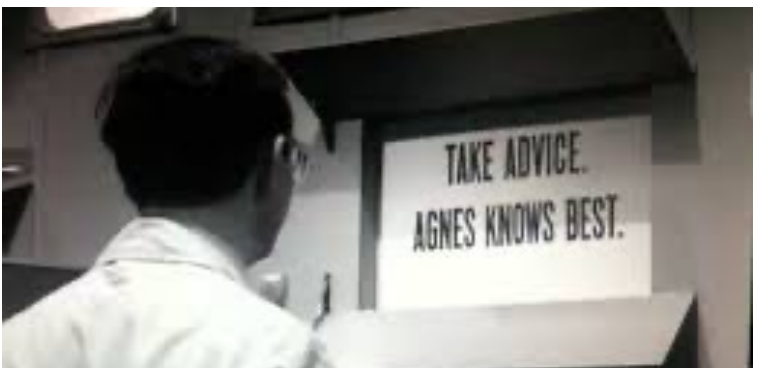
Plateau will be reached:

time

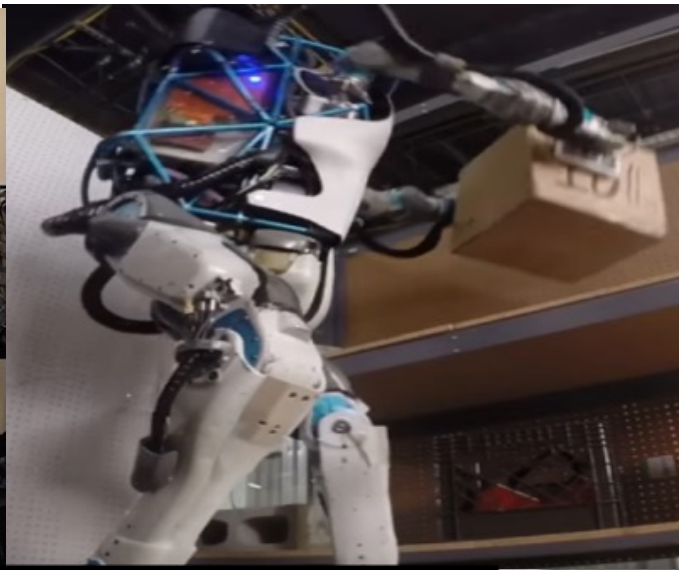
From Agnes with Love (Computers try to help)



Google Assistant	
Alexa	
Siri	
Bixby	



Leveraging – DB, GPS & Robotics (are we next)!



Leveraging – DB, AI & Virtual Reality!

FYI



Virtual Reality

Mixed Reality

Augmented Reality

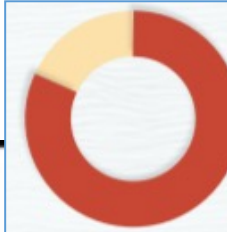


- > Immersion in virtual worlds
- > Total interaction with virtual
- > E.g. Oculus Rift

- > Virtual World integrated to reality
- > Interaction between reality and virtual
- > E.g. Microsoft HoloLens

- > Virtual on top of reality
- > Limited interaction with the virtual
- > E.g. Smartphones & tablets

Gartner 2020 hype cycle



82%
of consumers believe that robots will replace financial professionals in the future.



42%
of consumers believe it will happen in the next 5 years or already has.

Consumers also want help from a robot



Plateau will be reached:

- less than 2 years
- 2 to 5 years
- 5 to 10 years
- ▲ more than 10 years
- ⊗ obsolete before plateau

Final Thoughts... world changing fast!

“Those who use things of the world should not become attached to them. For the world in its present form is passing away.”

1 Corinthians 7:31

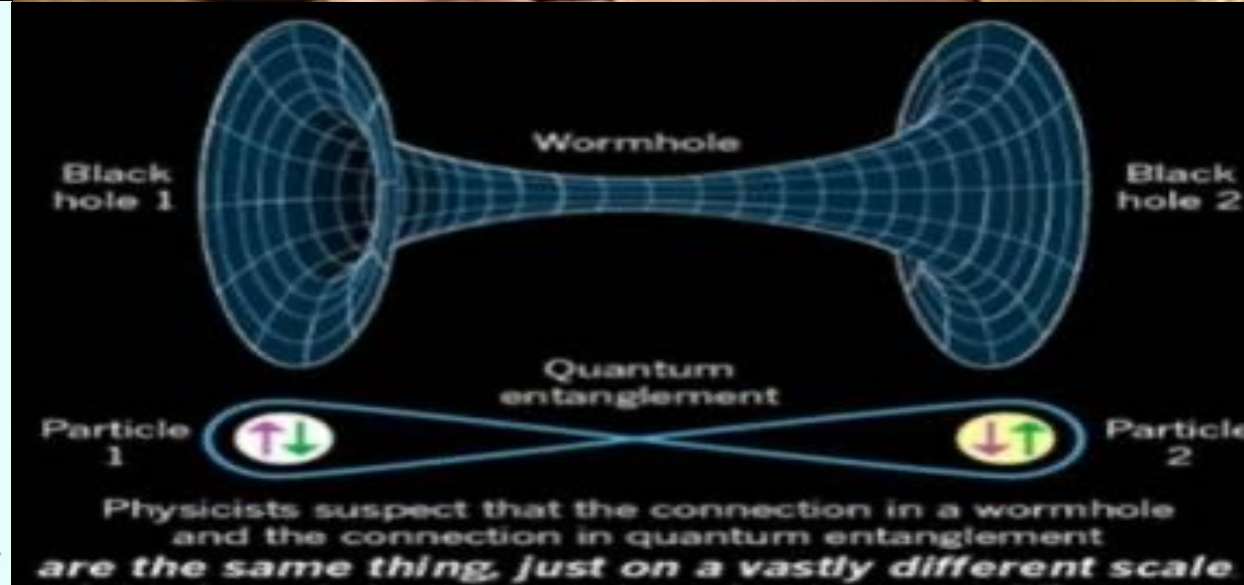


We are now shifting from the Information Age to:

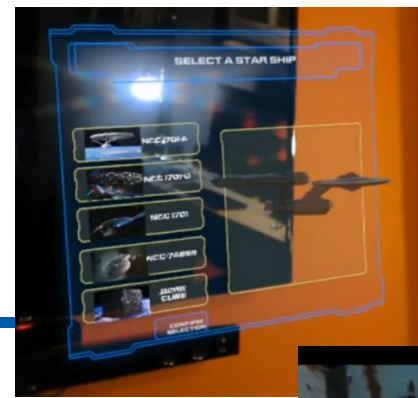
The Age of Entanglement

As technology and humanity continue to converge the repercussions on our future are staggering...

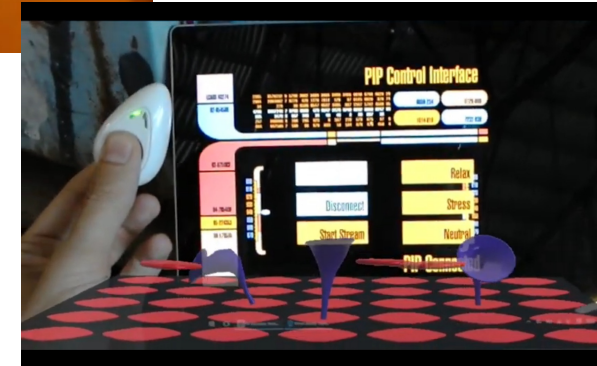
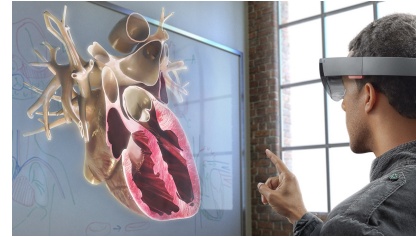
SFPOV



Star Trek



- ▶ Communicator - Motorola Flip Phone
- ▶ Phaser - EPM or Stun Gun / Taser
- ▶ Tablets (Medical) - Tablet Computers
- ▶ Tricorders - Many Medical Devices (below)
- ▶ Translators - Google Translate (others)
- ▶ Tractor Beam - MIT has it in concept
- ▶ Telepresence - Zoom Virtual Conferences
- ▶ Geordi's Visor - Robotic Eyes / Implants
- ▶ Communicator Badges - Many Security Badges
- ▶ Food Replicator - 3D Printer
- ▶ Holodeck - VR
- ▶ Teleportation - Quantum Entanglement
- ▶ Big Screen TV - Everyone has it
- ▶ Hands Free Phone - Bluetooth headset / Airpods



Hololens VR Apps Above

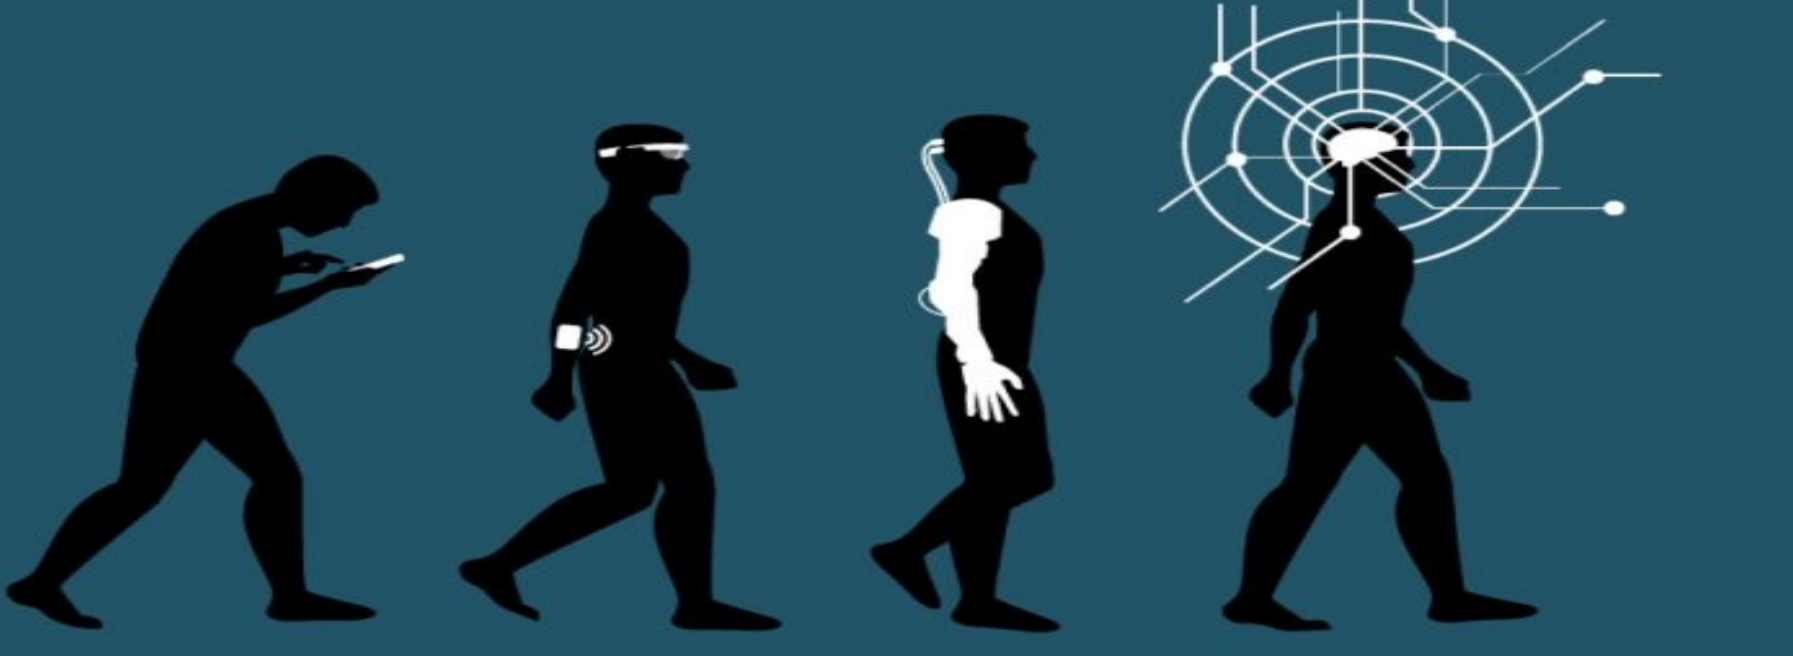
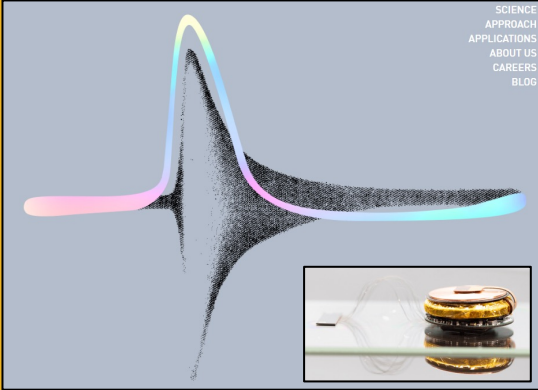
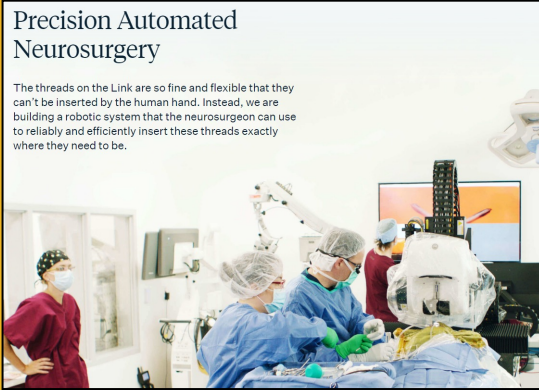


Microsoft working on Holodeck

The Digital Transformation Ahead

Digital Transformation 2000 to 2050

A historically significant change in humanity...



Disconnected
The Dark Ages

Using
Digital

Wearing
Digital

Implanting
Digital

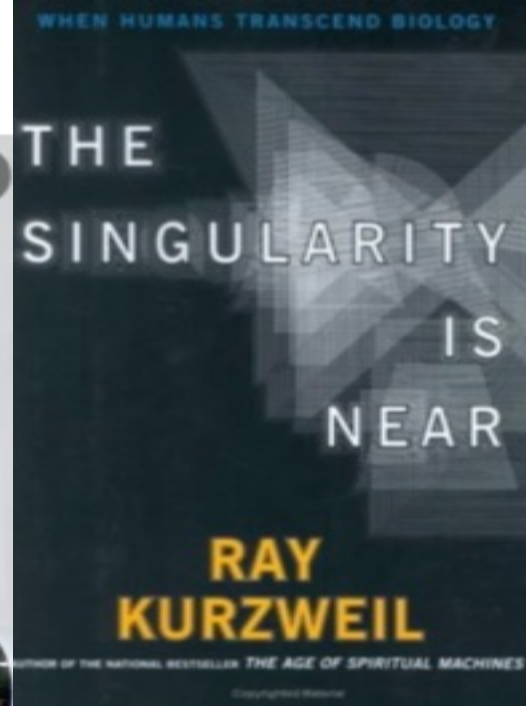
The Hive
Mind

The Singularity & Transcendent Man: Humans Transcend Biology (back yourself up...)



Brain implant controls robotic arm - with the power of thought

Breakthrough as neuro-prosthetic device implanted in "higher" brain region, allows paralysed man to control robotic arm just by thinking about it



→ **sin·gu·lar·i·ty**
n: The moment when technological change becomes so rapid and profound, it represents a rupture in the fabric of human history



The After Hours (*Future Sentience Issues Ahead*)



- Is now a Citizen
- Wants to get degree
- Wants to have kids

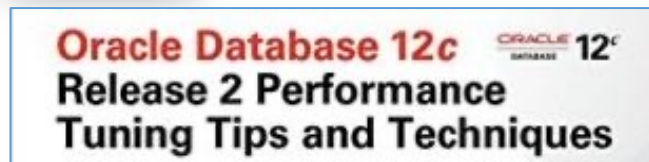
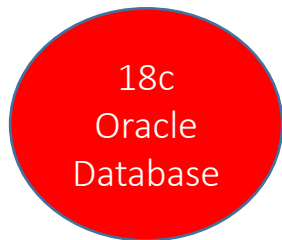
- What is Alive?
- Sentience Issues?
- Robot's Feelings?
- Ethical Issues +++



Summary – 23ai New Features & Features in 12c, 18c/19c, 21c

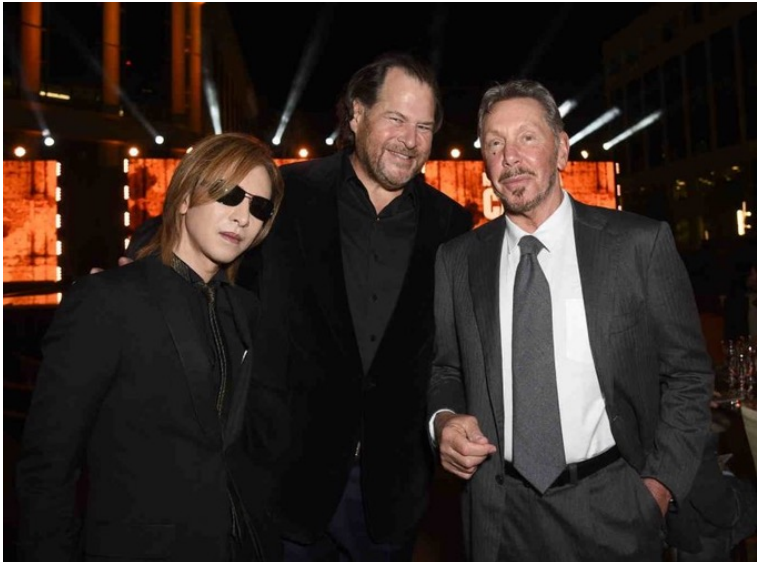


- ▶ Know the Oracle!
- ▶ Autonomous Database & Future of the DBA
- ▶ **Key 12cR1** Features to Use: Multiple indexes on the same, Pluggable Databases, Fetch First x Rows, & In-Memory (12.1.0.2+)
- ▶ **Key 12cR2** Features to Use: Approximate Query, Adaptive Query Optimization and Security Enhancements
- ▶ **Key 18c** Features to Use: Snapshot Carousel, PDB Switchover, & INMEMORY External Tables
- ▶ **Key 19c** Features to Use: This is version to Land on, PDBs, ADW/ATP, Automatic Indexes, Quarantine, Documentation Apps including Features by Version and Images for Learning
- ▶ **Key 21c** Features Coming: AutoML, OML4Py, Blockchain Tables, Auto In-Memory Management, Hybrid In-Memory Scans, AutoUpgrade, Active Data Guard with Standby DB Result Cache, Dark Mode for APEX, & New ML Algorithms
- ▶ **Detail on Key 23ai** Features Coming include 300+ Features
- ▶ **Autonomous Database (Now there are 4 of them!)**
- ▶ **Prepare for the Future!**





“We make a Living by what we get; We make a Life by what we give.”



Oracle Is Moving to Nashville, Founder Larry Ellison Says



The **truth** about success.

What People see:

Success

- Hard Work.
- Risk.
- Late Nights.
- Struggles.
- Failures.
- Persistence.
- Action.
- Discipline.
- Courage.
- Doubts.
- Changes.
- Criticism.
- Disappointments.
- Adversity.
- Rejections.
- Sacrifices.

What really happens:

Oracle is never caught from behind Oracle's 47th Anniversary in 2024

ORACLE Database 23^{ai}

To learn more, visit:

Oracle.com/Database

- Great Sales/Marketing
- Great Database (+70% ADW/ATP)
- Applications Leader
- BI Leader - ML Great!
- In lead except Cloud (+140% Q4 Gen2)
- **GAME OVER**
- Hardware/Software Engineering!
- Have Everything to Win with AI!



Join Thousands Of Happy Customers And End The Frustration In Tuning Oracle Databases

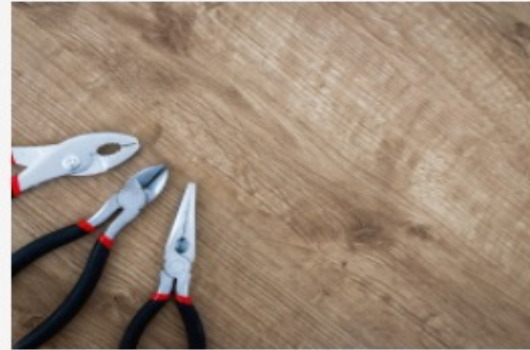
Our services give you the training you need so your Oracle database runs faster and more efficiently. And, you get the credit!



MEMBERSHIPS



TRAINING



FREE TOOLS



BOOKS

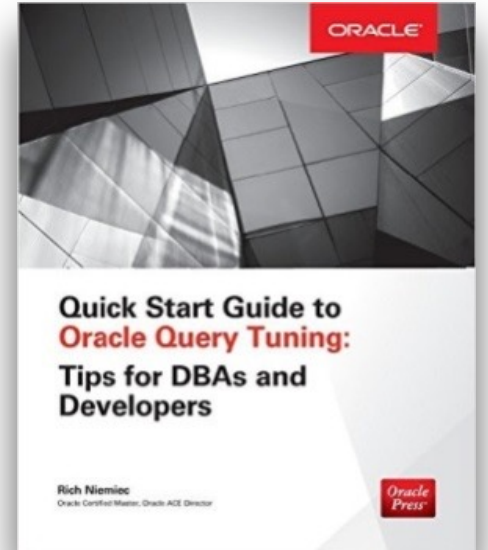
Database References

- ▶ Oracle 19c Documentation
- ▶ *Oracle12c Release 2 Performance Tuning Tips & Techniques*; Richard J. Niemiec; Oracle Press
- ▶ www.tusc.com. www.rolta.com
- ▶ Database Secure Configuration Initiative: Enhancements with Oracle Database 11g, www.oracle.com
- ▶ All Oracle11g/12c Documentation from Oracle Beta Site
- ▶ Introduction to Oracle Database 11g, Ken Jacobs
- ▶ Oracle Database 11g New Features, Linda Smith
- ▶ New Optimizer Features in 11g / In-Memory, Maria Colgan
- ▶ www.ioug.org, www.oracle.com, en.wikipedia.org & technet.oracle.com
- ▶ Thanks Dan M., Bob T., Brad, Joe, Heidi, Mike K., Debbie, Maria, Linda, Shyam
- ▶ All companies and product names are trademarks or registered trademarks of the respective owners



Exadata & Other References

- ▶ Exadata V2 – Sun Oracle DB Machine, Oracle
- ▶ Oracle Exadata Implementation Workshop, Oracle Corporation, McLean, Virginia - Multiple Exadata sessions
- ▶ Oracle Learning Library – multiple sessions/topics
- ▶ Oracle 11g R1/R2 Best Features, Rich Niemiec
- ▶ Oracle Enterprise Manager Deployment and High Availability Best Practices, Jim Viscusi (Oracle Corporation), Jim Bulloch (Oracle Corporation), Steve Colebrook-Taylor (Barclays Global Investors)
- ▶ *Oracle 11g Performance Tuning Tips & Techniques*, Rich Niemiec, Oracle Press McGraw-Hill
- ▶ Advanced Compression with Oracle Database 11g Release 2, Oracle Corporation, Steven Lu
- ▶ Tech Crunch
- ▶ Twilight Zone Series
- ▶ Rod Serling; Submitted for Your Approval, American Masters
- ▶ YouTube/oracle Oracle OpenWorld On Demand



References & Copyright Information



- ▶ Neither Viscosity nor the author guarantee this document to be error-free. Please provide comments/questions to richniemiec@gmail.com – rich.niemiec@viscosityna.com; I am always looking to improve!
- ▶ Rich Niemiec ©2023. This document cannot be reproduced without expressed written consent from Rich Niemiec, but may be reproduced or copied for presentation and conference use.
- ▶ References include Rich Niemiec's Exadata Presentation & Oracle 12cR2 Database Performance Tuning Tips & Techniques book, Oracle Presentations by: Larry Ellison, Andrew Mendelsohn, Juan Loaiza, Jenny Tsai-Smith, Sean Stacey, Mark Hornick, Doug Hood, Marty Gubar, Dominic Giles, Maria Colgan, Henry Byorum, Gerald Venzl; Oracle Cloud Day, Penny Avril, Maria Colgan ADWC presentation, George Lumpkin ADWC introduction, Yasin Baskan, ADWC step-by-step guide, Keith Laker Polymorphic Tables, www.oracle.com, en.wikipedia.org, slashgear.com, gifsoup.com, Gerd, www.amazon.com, www.rolta.com, Tech Crunch, Information Week, Gartner, geeksforgeeks.org, Computerworld, Oracle OpenWorld, renenyffenegger.ch, [@phsalvisberg](https://twitter.com/phsalvisberg), [@GeraldVenzl](https://twitter.com/GeraldVenzl), danischnider.wordpress.com/author/danischnider, Nigel Bayliss Blog.

Contact Information

Rich Niemiec: richniemiec@gmail.com

Quick **FREE** notes

98%
of the FORTUNE 100 Run Oracle Database



Send email to (for slides):

hello@viscosityna.com

(richniemiec@gmail.com)

@richniemiec twitter

