BUDAPEST DATA FORUM
2015. JÚNIUS 3-5. DATA WAREHOUSING · BIG DATA · DATA SCIENCE

# Adattárházak 2015-ben kiterjesztés és gyorsulás: Big Data és a relációs világ, In-Memory, Exadata

Fekete Zoltán

principal sales consultant
Database and DBO Competency Center, CEE
https://blogs.oracle.com/zfekete/

2015. június 3.


ORACLE

# Safe Harbor Statement

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

# Információ/adat menedzsment modernizáció

## Optimalizált adat menedzsment a valós-idejű betekintésig

**Tároló optimalizálás**

**Biztonság növelése**

**Hozzáférés kiterjesztés**

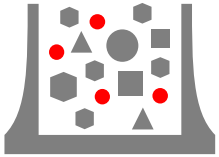

**Adatforrások szélesítése**

**Gyorsabb elemzések**

# Adattárház újdonságok

ORACLE DATABASE 12c

## És ami a metaadatokat összefogja: Enterprise Metadata Mangement

### Teljesítmény és skálázhatóság

- Adaptive Query optim.
- **Database In-Memory**
- Új tábla particionálás: interval-reference
- Attribure Clustering; Zone Maps
- Párhuzamos végrehajtásban:
  - Concurrent Union-All
  - Parallelization of correlated filters and expressions

### „Always on" DW

- Új online DDL műveletek: partition move, stb.
- Async partitioned global index maintenance
- Invisible columns
- Out-of-place & sync refresh for materialized views
- **Exadata X5**
- **Big Data Appliance X5**
- ... *„SW in Silicone"*

### Data scientists & Big Data

- SQL Pattern Matching
- Advanced Analytics
  - Fejl. Data Miner GUI
  - Új in-database prediktív alg.
  - R nyelv további integrálása
- Hadoop Encryption
- NoSQL 3.0 ... 3.2.5
- **Big Data SQL**
- JSON támogatás, DB 12c
- Big Data Discovery
- Visual Analyzer
- BI Cloud Services
- **Big Data Spatial & Graph**

ORACLE

# SQL Pattern Matching

## Scalable discovery of business event sequences

Find event A (“privilege revoked”) followed by 3 or more occurrences of event B (“attempted login”) within 1 minute

Find 10-day periods where a stock price has “double-bottomed”

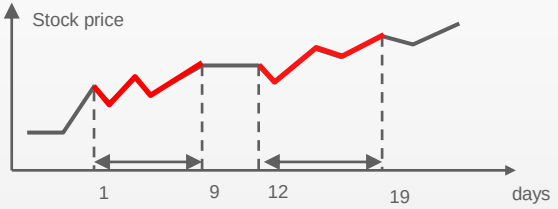


- SQL Pattern Matching provides expressive syntax and fast performance for pattern matching
- New SQL construct: MATCH_RECOGNIZE
- Define patterns using regular expression syntax

# Adaptive Query Optimization

## Adaptive Plans

Adjust query plans at runtime based upon current data

## Adaptive Statistics

Adapt optimizer statistics at runtime

"Learn" for future queries

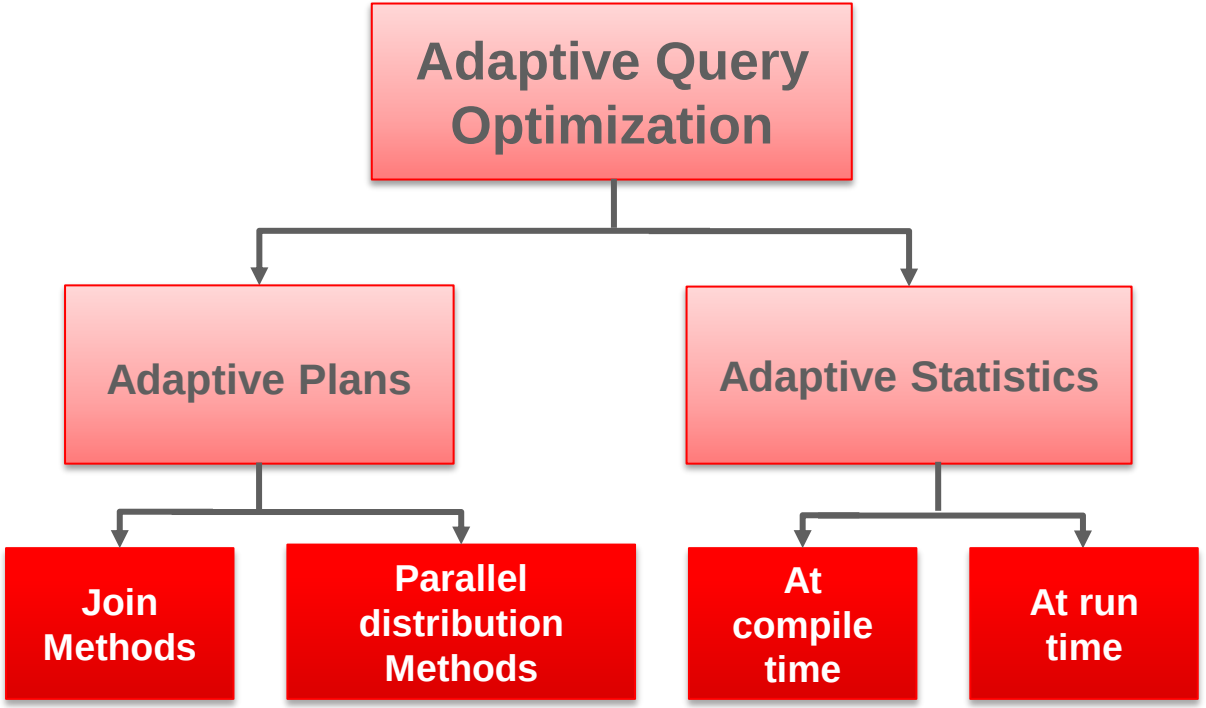

# Oracle Database 12c újdonságok áttekintése

1. Oracle Információ Menedzsment Architektúra
2. Exadata és Big Data Appliance, Big Data áttekintés
3. Oracle Database In-Memory

# Enterprise-Class Big Data Solutions

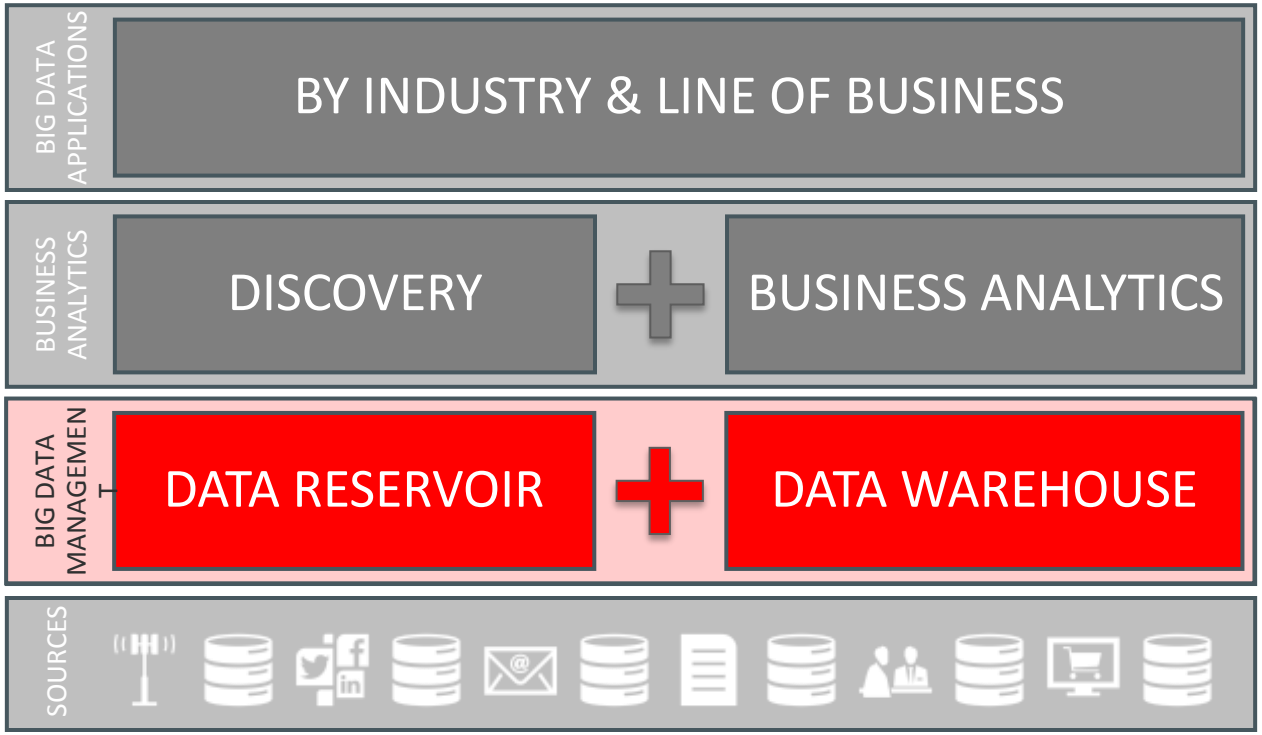

# The Big Picture – Oracle Big Data Management System

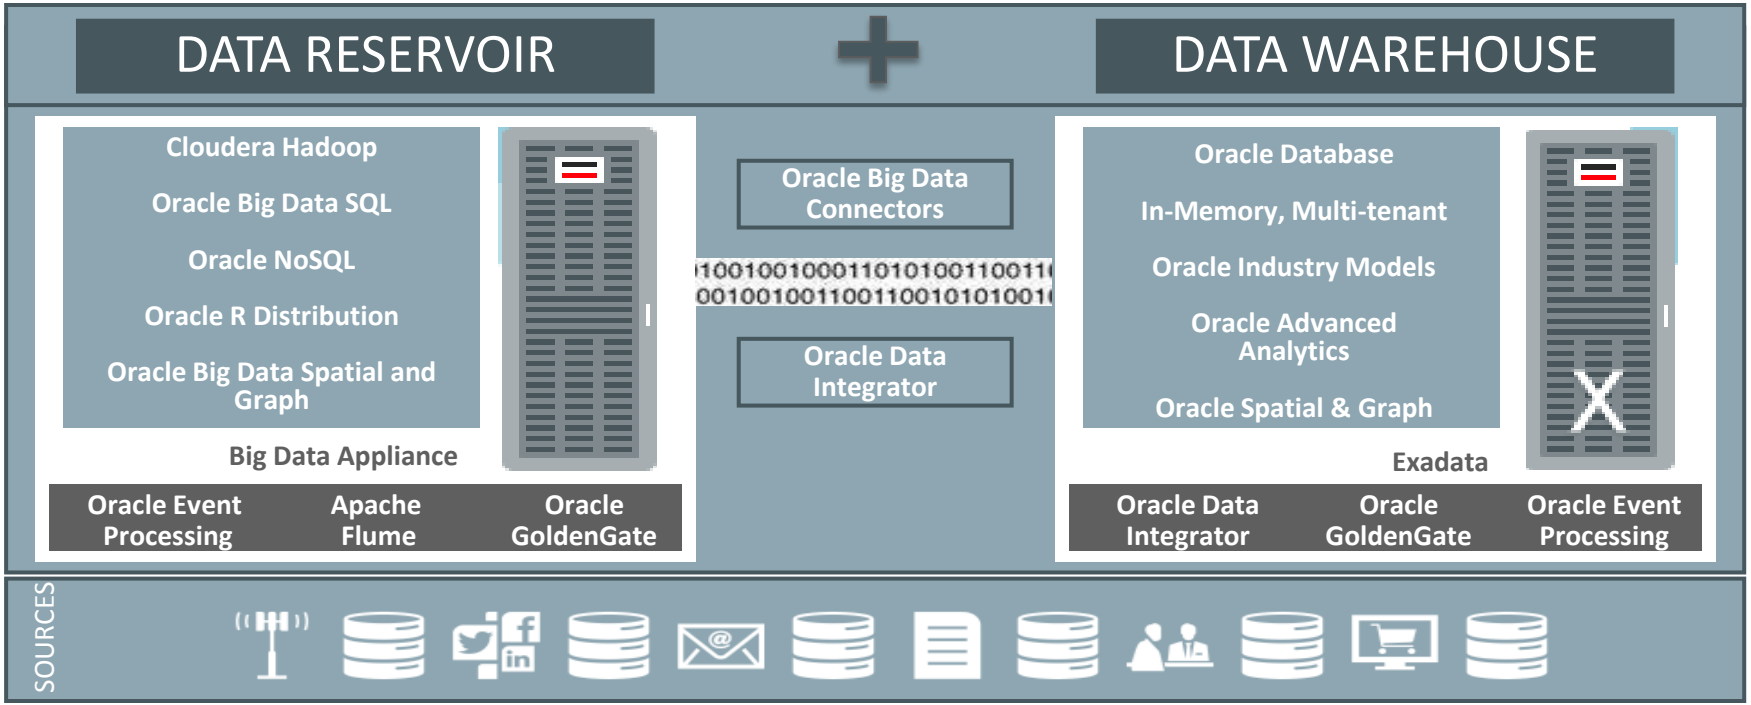

# Information Management – Logikai nézet
## „Data Reservoir” & Enterprise Information Store – teljes nézet

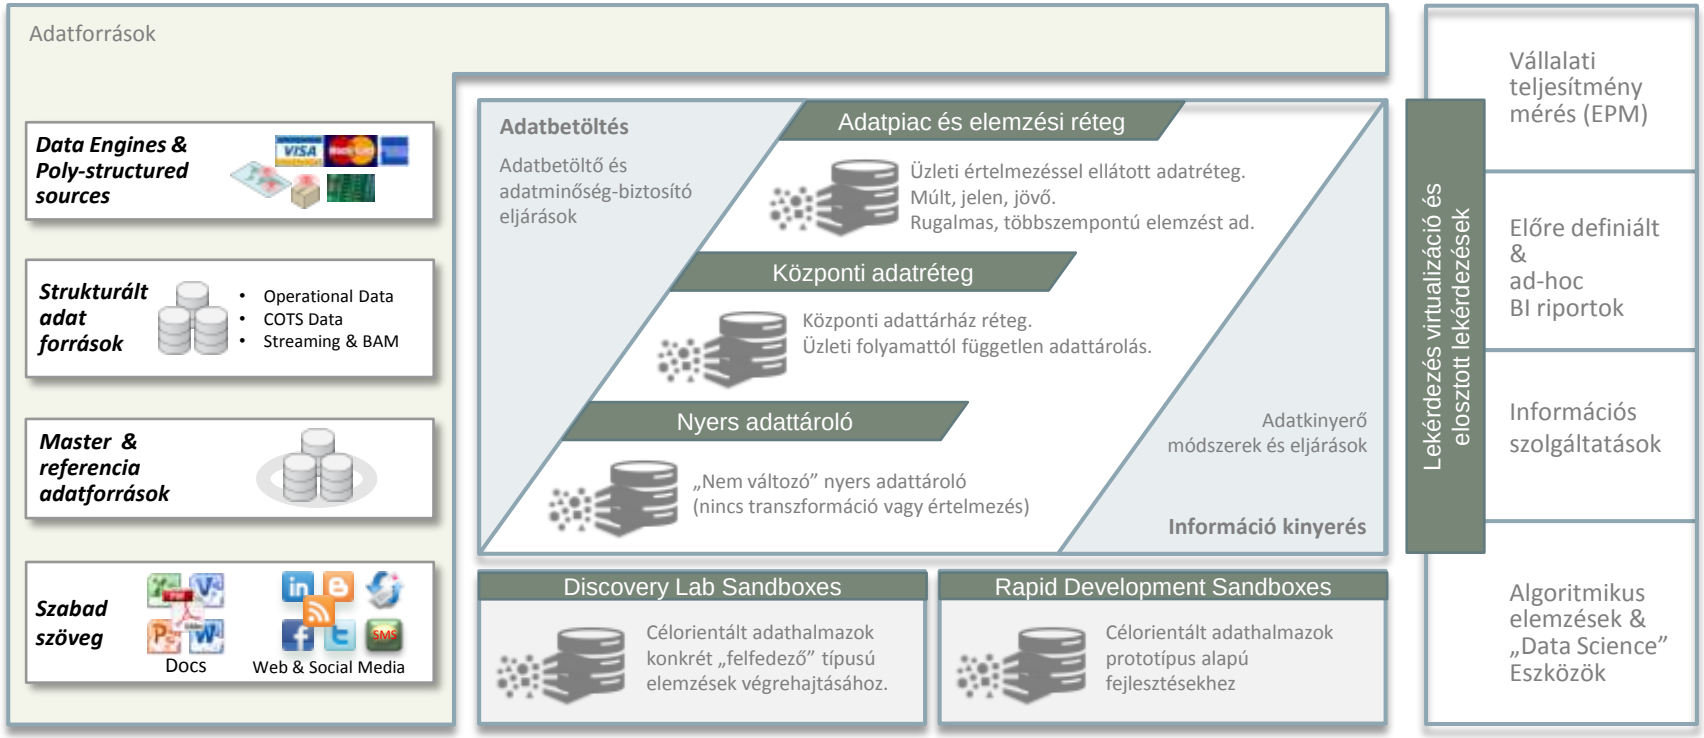

# Oracle Database 12c újdonságok áttekintése

1. Oracle Információ Menedzsment Architektúra
2. Exadata és Big Data Appliance, Big Data áttekintés
3. Oracle Database In-Memory

# Oracle Big Data Platform

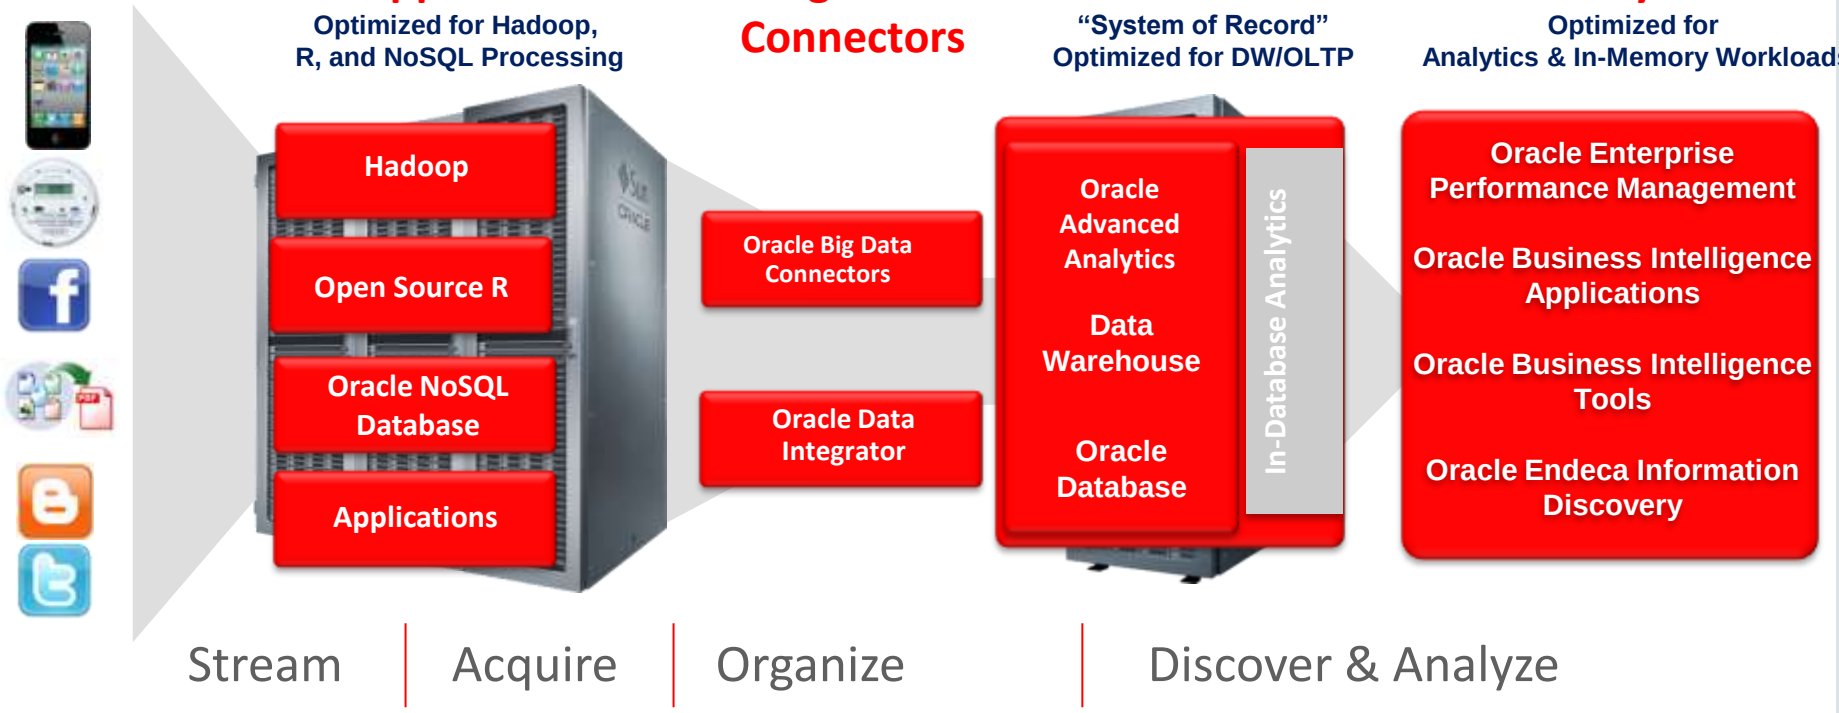

# Engineered Systems for Big Data

Complete | Optimized | Fully Redundant |Scale-Out

## Exadata X5-2

- Scale-Out 2-Socket Database Servers
  - Fastest Xeon chips, 18-core
- Scale-Out 2-Socket Storage Servers
  - 8-core Xeon chips enable offload to storage
  - Extreme Flash (EF) Storage
    - 12.8 TB Ultra-Fast PCI Flash Cards
  - High Capacity (HC) Storage
    - 6.4 TB Ultra-Fast PCI Flash Cards
    - 48 TB SAS disk storage



## Big Data Appliance X5-2

- Scale-Out 2-Socket Compute Servers
- Integrated Software
  - Fastest Xeon chips, 18-core
  - 48 TB SAS disks storage
- Integrated Software
  - Oracle Linux 6.5
  - Cloudera Distribution of Apache Hadoop 5.3 (EDH Edition)
  - Oracle Big Data SQL
  - Oracle R Distribution
  - Oracle NoSQL Database CE

- Unified Ultra-Fast InfiniBand Network
  - 40 Gb InfiniBand internal connectivity
  - 10 Gb or 1 Gb Ethernet data center connectivity

ORACLE

# Oracle Big Data Appliance is a Proven, **Cost Effective** Solution

"Oracle Big Data Appliance is an excellent choice for customers looking to work with the full suite of Cloudera's leading Hadoop-based technology. It's **more cost-effective and quicker to deploy** than a DIY cluster."

- Mike Olson, Cloudera founder, Chief Strategy Officer, and Chairman of the Board

cloudera


**21%**
**Lower Costs**

**33%**
**Faster to Implement**

Source: ESG White Pap

ESG
Enterprise Strategy Group

ORACLE

# Célrendszer megközelítés: **Exadata X5-2 előnyök**

## Leggyorsabb

- **Leggyorsabb OLTP: extreme flash**
  4,1 millió OLTP IOPS / rack
- **Leggyorsabb DW**
  263 GB/s throughput / rack
- **Leggyorsabb In-Memory Database**
  Milliárd rekordok /s/core
- **Kisméretű Exadata is túlszárnyalhat hatalmas szervereket és tárolókat**

## Legköltséghatékonyabb

- **Költséghatékony tiering: memória, flash, diszk**
  Legjobb tömörítéssel
- **Gyors VM-ek és konszolidáció**
  egyedülálló end-to-end priorizáció
- **Elasztikus Scale-Out konfig.**
- **End-to-End integrált Mgmt**
- **Standard, legjobb support**

## Legnagyobb rendelkezésre állás

- **Redundáns Scale-Out HW**
- **Leggyorsabb Recovery**
  Server, storage, network
- **Legjobb MAA implementáció**
  RAC, ASM, Data Guard, RMAN
- **Teljes hibatesztelés**
- **In-Memory hibatűrés**

# Data Analytics Challenge

Separate silos with separate data access interfaces

# Data Analytics Challenge

No comprehensive SQL interface across Oracle, Hadoop and NoSQL

# What customers want: Oracle Big Data SQL

Rich, comprehensive SQL access to all enterprise data

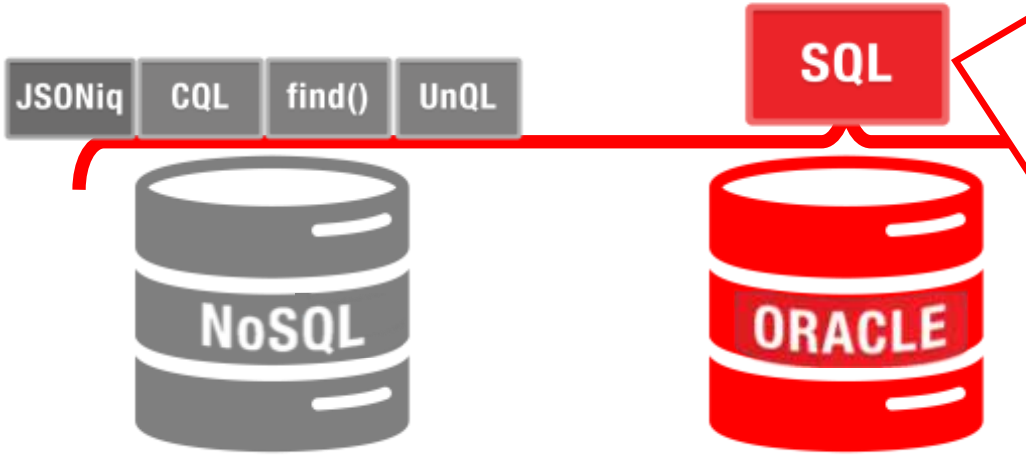


**The Power of Oracle SQL**

- Wide variety of 'Big Data' types
  - Structured data
    - Numeric, string, date, …
  - Unstructured data
    - LOBs, Text, XML, JSON, Spatial, Graph, Multimedia

- Rich SQL Analytic Functions
  - Ranking, Windowing, LAG/LEAD, Aggregate, Pattern Matching, Cross Tabs, Statistical, Linear Regression, Correlations, Hypothesis Testing, Distribution Fitting, …

# Oracle Big Data Management System

## Unify All Query

# Data Lifecycle Management & Query Offload

## More data on-line and available at a lower cost

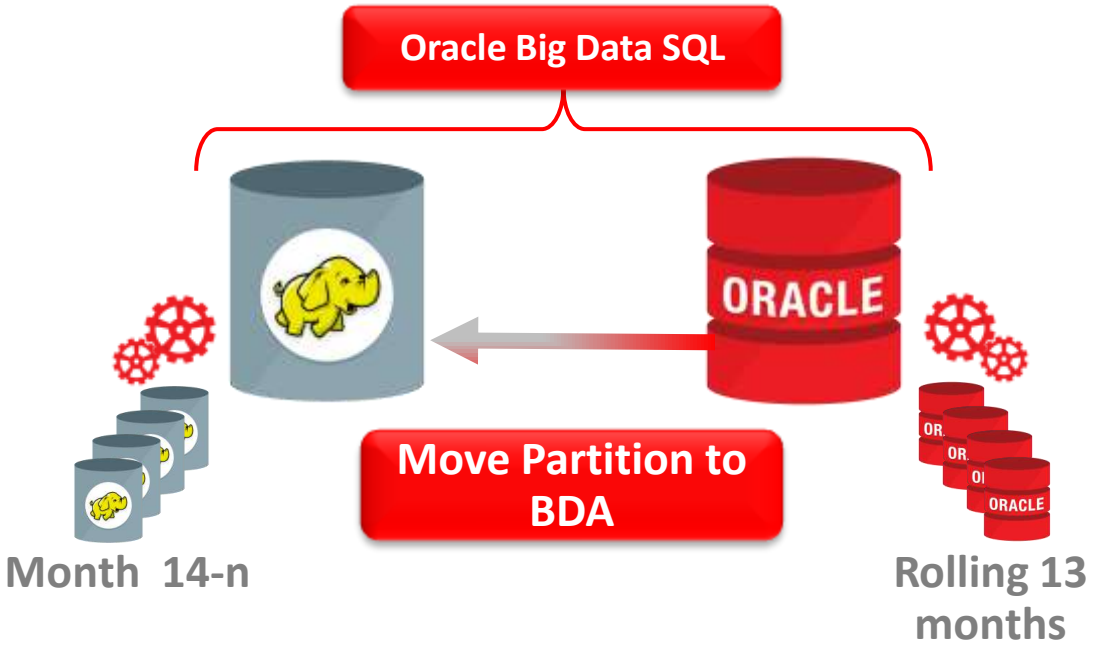


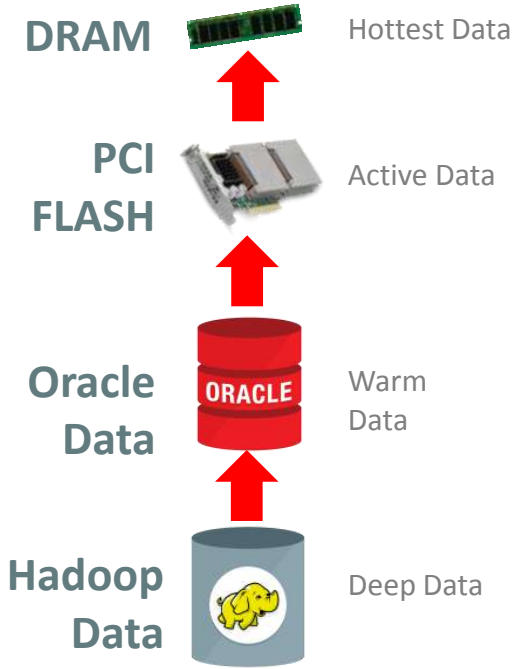

# **Minden** adat szabályozása—SQL & Oracle Big Data SQL

**SQL**

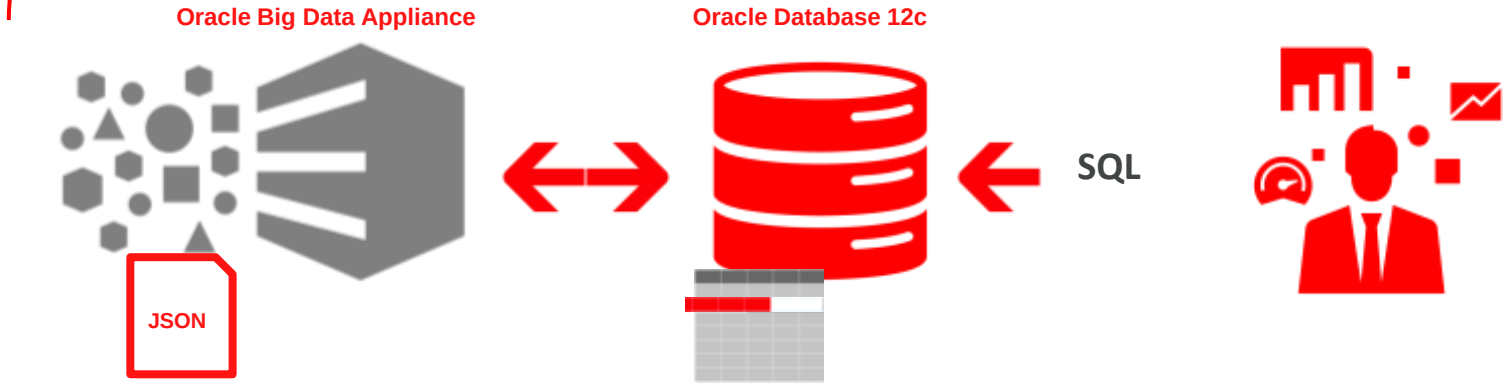


Store JSON data unconverted in Hadoop

Store business-critical data in Oracle

Data analyzed via SQL

- **Advanced Security Hadoop-on**
  - **Masking/Redaction**
  - **Virtual Private Database**
  - **Fine-grained Access Control**

```
DBMS_REDACT.ADD_POLICY(
  object_schema => 'txadp_hive_01',
  object_name => 'customer_address_ext',
  column_name => 'ca_street_name',
  policy_name => 'customer_address_redaction',
  function_type => DBMS_REDACT.RANDOM,
  expression => 'SYS_CONTEXT(''SYS_SESSION_ROLES'',
''REDACTION_TESTER'')=''TRUE'''
  );
```

# **Oracle Big Data Discovery**. The Visual Face of Hadoop

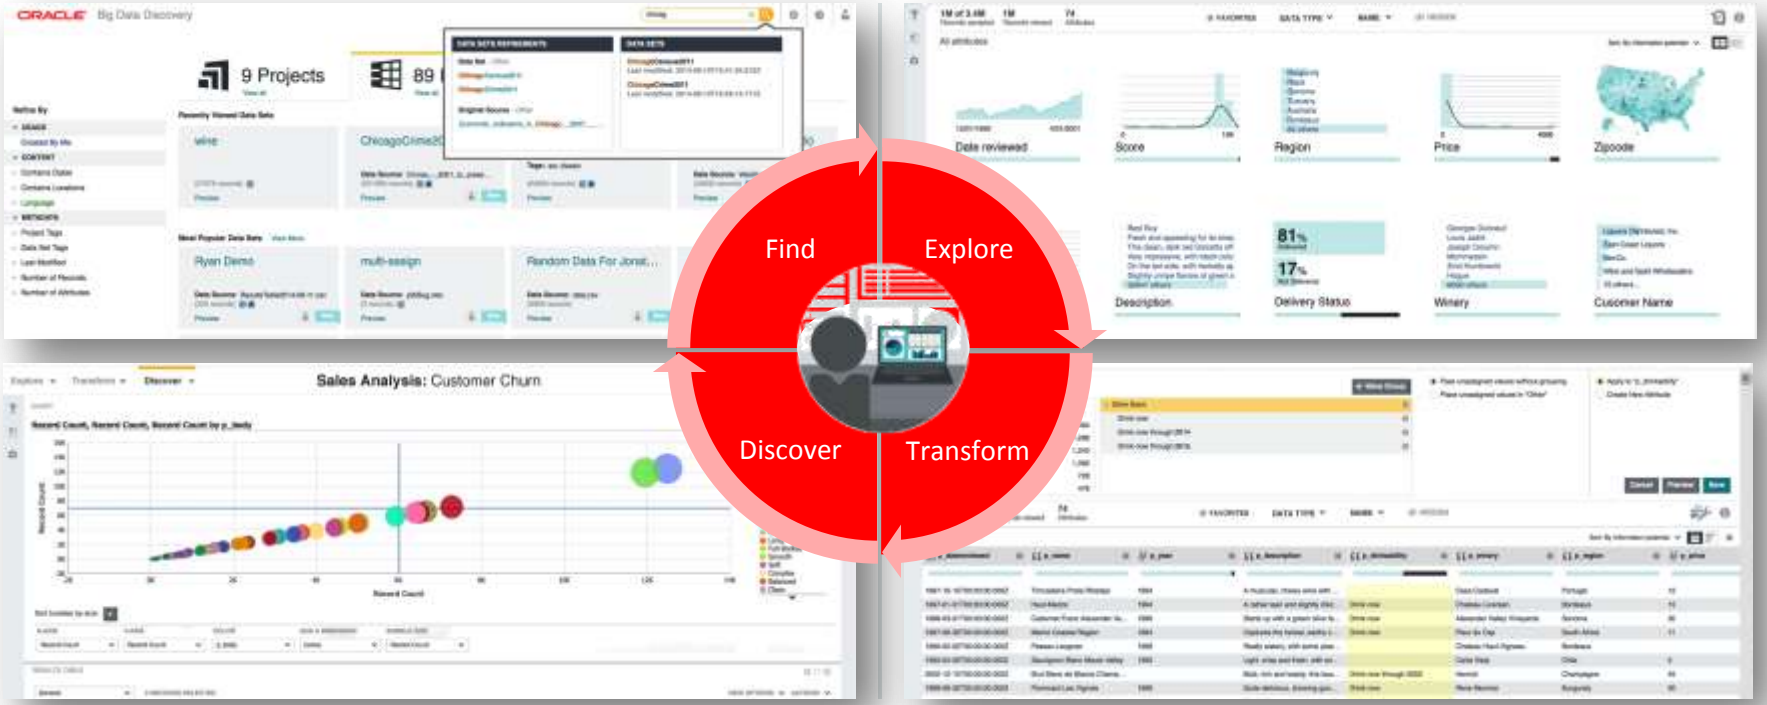

# Oracle Database 12c újdonságok áttekintése

1. Oracle Információ Menedzsment Architektúra
2. Exadata és Big Data Appliance, Big Data áttekintés
3. Oracle Database In-Memory

# Oracle Database In-Memory

**Egyedülálló Dual-Format architektúra**

# Legfrissebb analitika

- **Mindkettő:** soros és oszlopos in-memory formátum
- Szimultán aktív és **tranzakciósan konzisztens mindig az aktuális adatok elérése**
- **Megszünteti a manuális tuningot és a költséges analitkus indexeket**

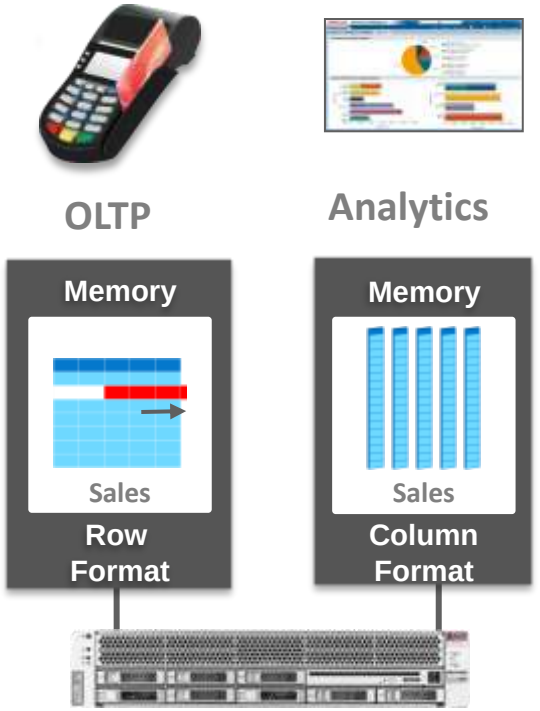

# Oracle Database In-Memory

**Az alkalmazásoknak transzparens**

# 100% kompatibilis

- **A meglévő Oracle tudás és tapasztalat**, az összes Oracle biztonsági és rendelkezésre állási funkció
- **Skálázódik** a teljes clusterre
- **Az alkalmazások változatlanok**, csak gyorsabbak
  nincs szükség újraírásra,
  nincs szükség konfigurálásra

**Lekérdezés futási eredmény**

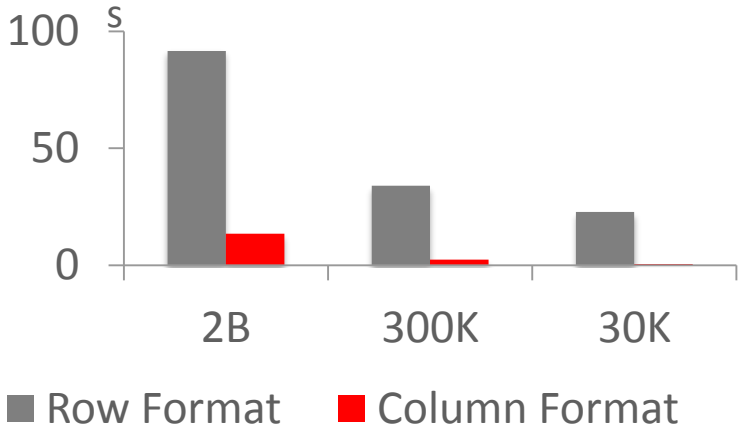


**Schneider Electric,**
**2 milliárd GL adat elemzése**

# Oracle Database In-Memory: egyszerű bevezetni

**Az alkalmazásoknak transzparens!**

1. lépés: in-memory terület méretének megadása
   - inmemory_area = XXXX GB

2. lépés: mely elemek kerüljenek be az oszlopos in-memory területre:
   - alter table | partition … **inmemory**;

3. lépés: analitikus indexek eldobása, az OLTP további gyorsítása

# Complex OLTP is Slowed by Analytic Indexes

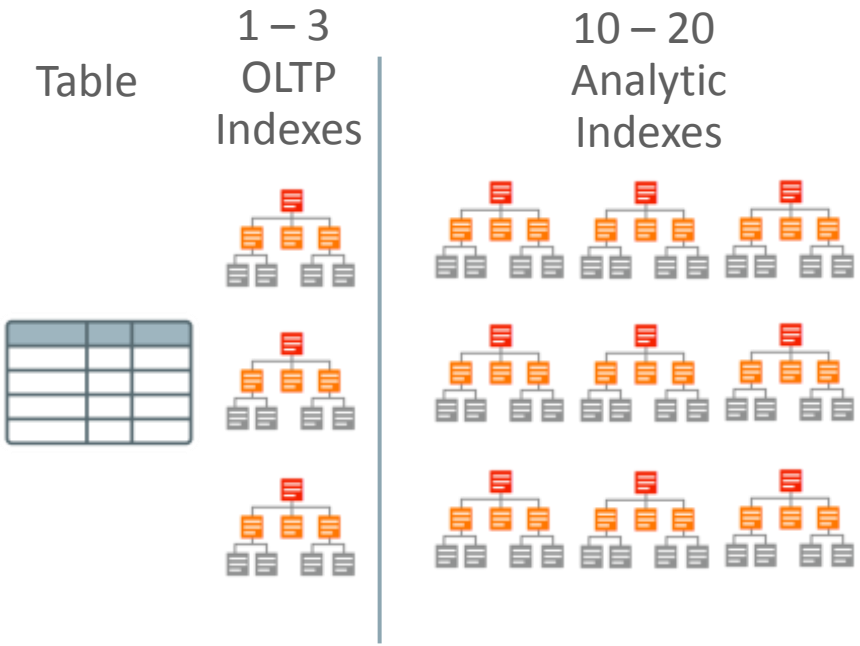


- Most Indexes in complex OLTP (e.g. ERP) databases are only used for analytic queries
- Inserting one row into a table requires **updating** 10-20 analytic indexes: **Slow**!
- Indexes only speed up predictable queries & reports

# OLTP is Slowed Down by Analytic Indexes

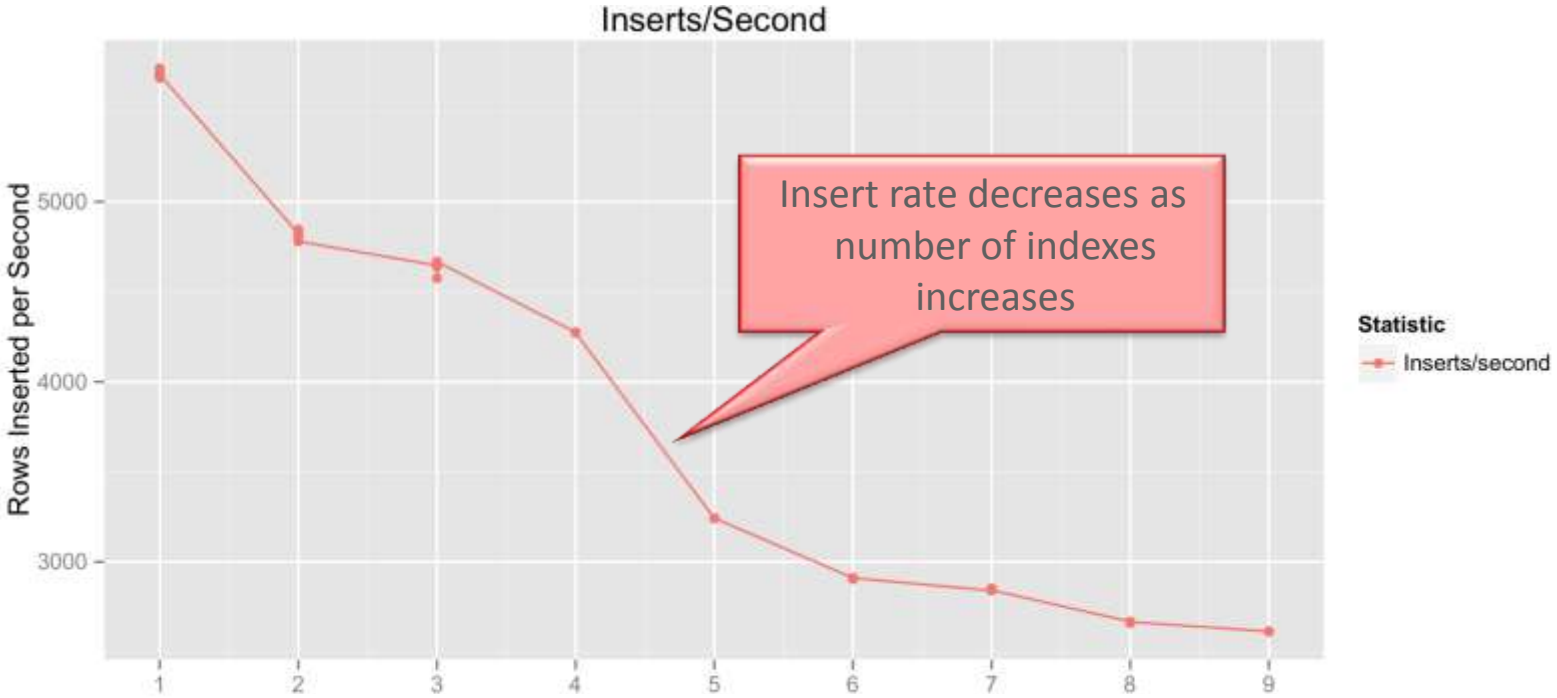


# of Fully Cached Indexes (Disk Indexes are much slower)

# Column Store Replaces Analytic Indexes

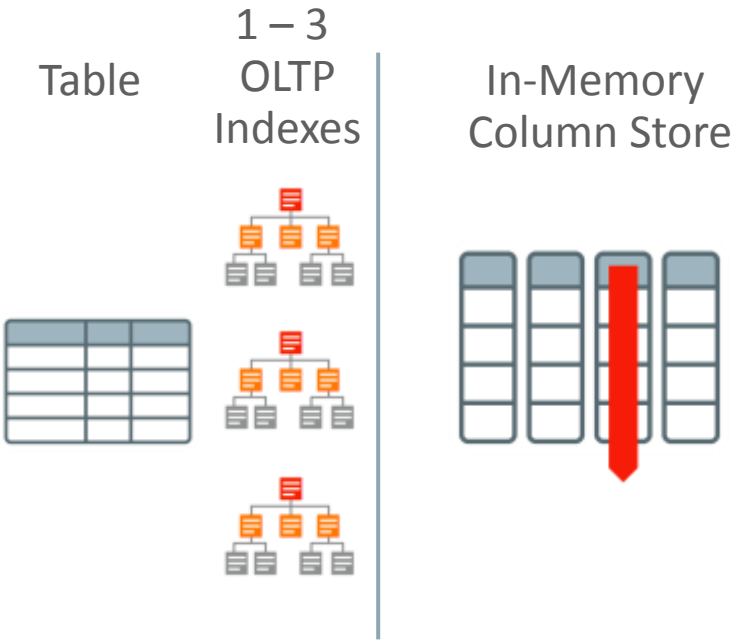


- Fast analytics on <u>any</u> columns
  - Better for unpredictable analytics
  - Less tuning & administration
- Column Store not persistent so update cost is much lower
  - OLTP & batch run faster

# Why is an In-Memory scan faster than the buffer cache?

Buffer Cache

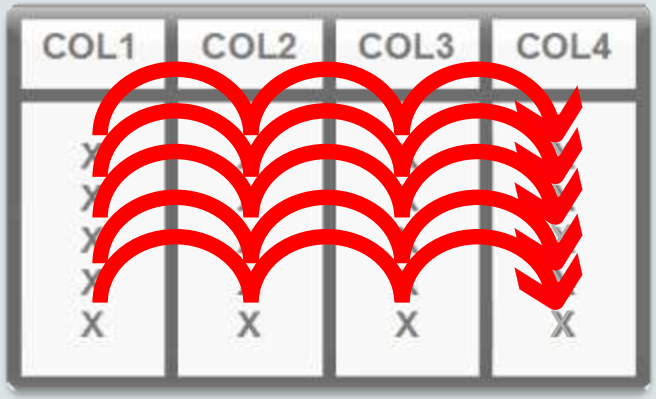


Row Format

```
SELECT COL4 FROM MYTABLE;
```

**RESULT**

# Why is an In-Memory scan faster than the buffer cache?

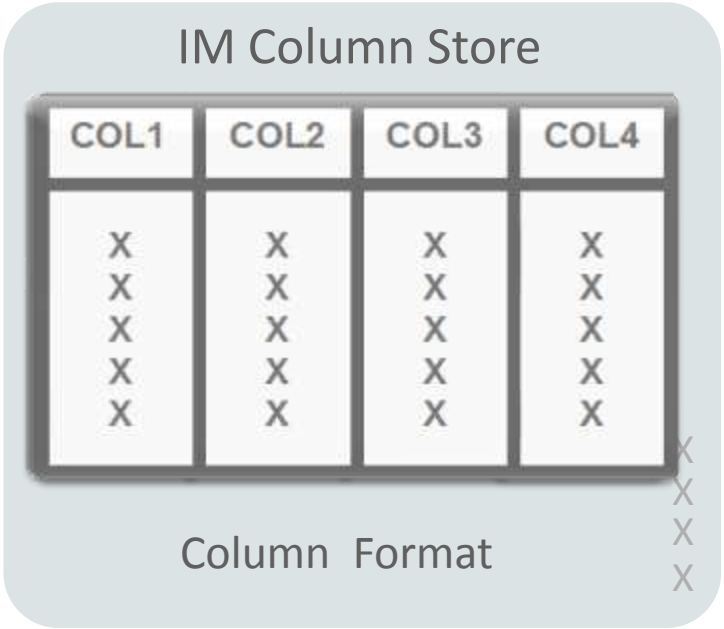


```
SELECT COL4 FROM MYTABLE;
```

**RESULT**

**REASON 1**
**only access the data you need for the query**
**REASON 2**
**Queries predicates applied directly to the compressed data**

# Oracle In-Memory Column Store Storage Index

**Example:** Find sales from stores with a store_id of 8 or higher

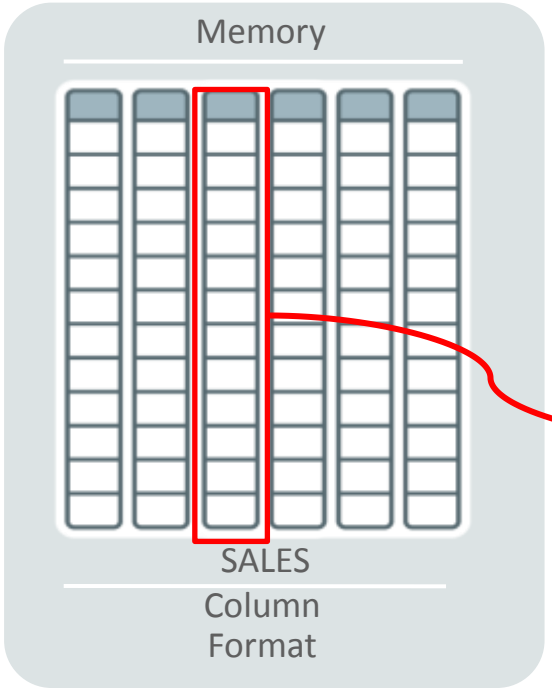



- Each column is the made up of multiple **column units** (CU)
- **Min** / **max** value is recorded for each column unit in a storage index
- **Storage index** provides partition pruning like performance for ALL queries

# Orders of Magnitude Faster Analytic Data Scans

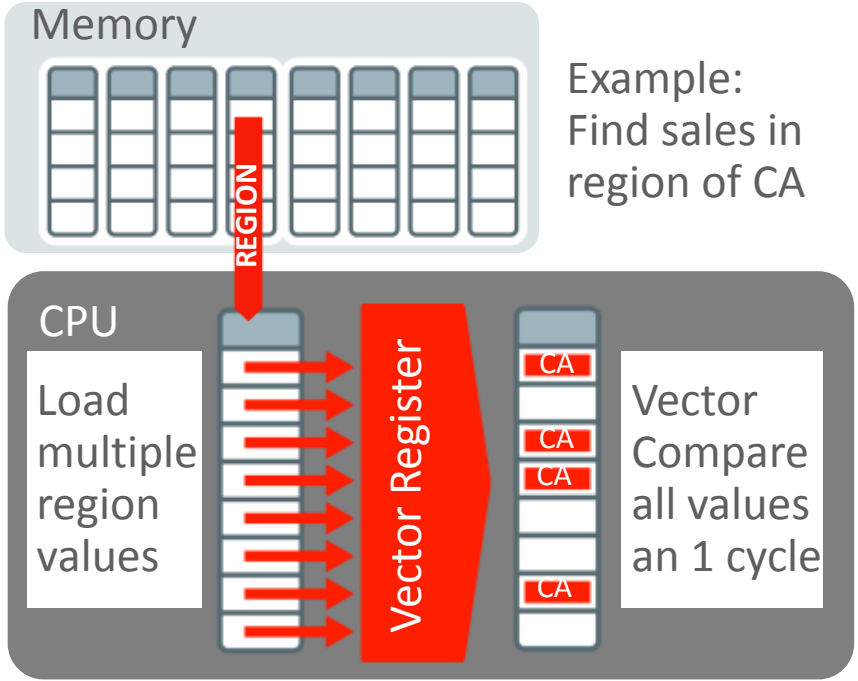


**> 100x Faster**

- Each CPU core scans local in-memory columns
- Scans use super fast SIMD vector instructions
  - Originally designed for graphics & science
- **Billions of rows/sec** scan rate per CPU core
  - Row format is millions/sec

# Joining and Combining Data Also Dramatically Faster

**Example:** Find total sales in outlet stores

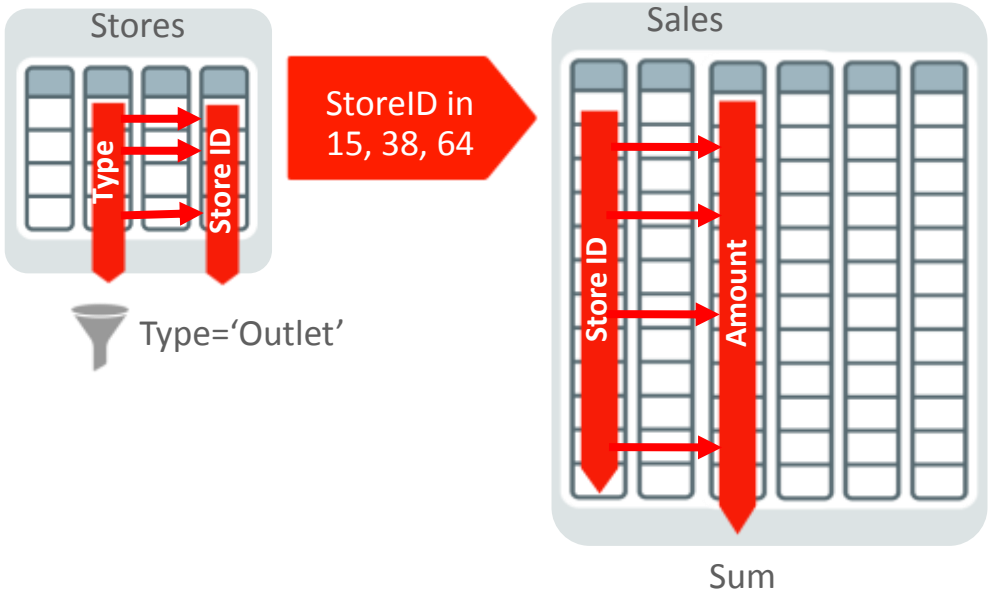


- Converts joins of data in multiple tables into fast column scans
- Joins tables **10x** faster

# Generates Reports Instantly

**Example:** Report sales of footwear in outlet stores

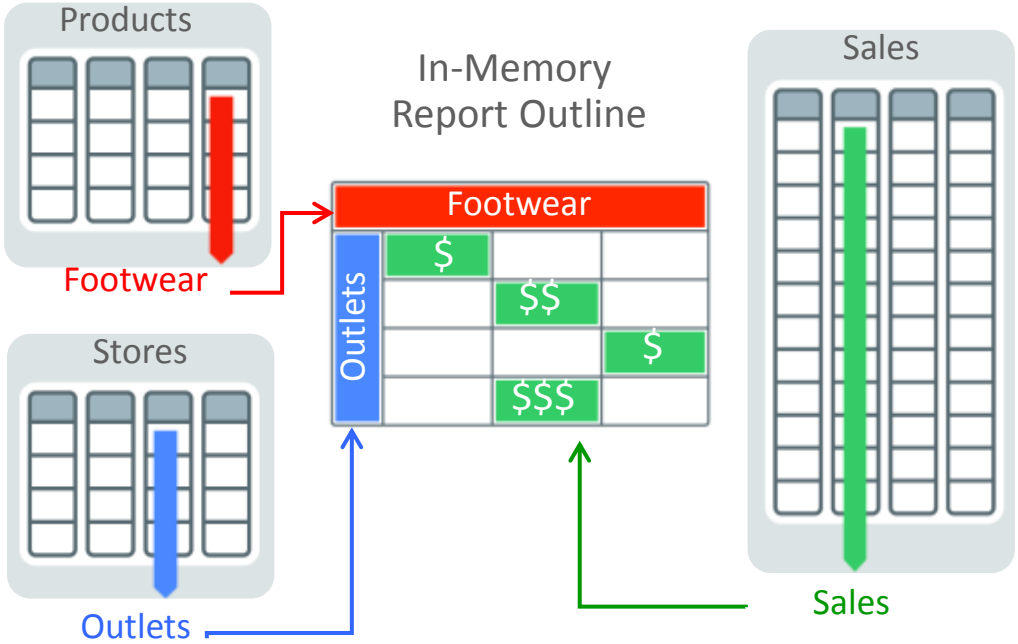


- Dynamically creates in-memory report outline
- Then report outline filled-in during fast fact scan
- Reports run much faster
  - Without predefined cubes
- **Also offloads report filtering to Exadata Storage servers**

# Populating : In-Memory Column Store

```
ALTER TABLE sales INMEMORY;

ALTER TABLE sales NO INMEMORY;


CREATE TABLE customers ……
PARTITION BY LIST
   (PARTITION p1 …… INMEMORY,
   (PARTITION p2 …… NO INMEMORY);
```

- New INMEMORY ATTRIBUTE
- Following segment types are eligible
  - Tables
  - Partitions
  - Subpartition
  - Materialized views
- Following segment types not eligible
  - IOTs
  - Hash clusters
  - Out of line LOBs

Pure OLTP Features

# Populating : In-Memory Column Store

```
ALTER TABLE sales INMEMORY
NO INMEMORY (PROD_ID);
```

```
CREATE TABLE orders
   (c1 number,
    c2 varchar(20),
    c3 number)
INMEMORY PRIORITY CRITICAL
NO INMEMORY (c1);
```

- Possible to populate only certain **columns** from a table or partition
- Order in which objects are populated controlled by **PRIORITY** subclause
  - Critical, high, medium, low
  - Default – none (populate on first access)
  - Does not control the speed of population

# Populating : In-Memory Column Store

```
ALTER MATERIALIZED VIEW mv1
INMEMORY MEMCOMPRESS FOR QUERY;
```

```
CREATE TABLE trades
   (Name varchar(20),
    Desc varchar(200))
INMEMORY
MEMCOMPRESS FOR DML(desc);
```

- Objects **compressed** during population
- New compression techniques
  - Focused on scan performance
- Controlled by **MEMCOMPRESS** subclause
- Multiple levels of compression
- Possible to use a different level for different partitions in a table

# Oracle Compression Advisor And In-Memory

```
DECLARE
  l_blkcnt_cmp       BINARY_INTEGER;
  l_blkcnt_uncmp     BINARY_INTEGER;
  l_row_cmp          BINARY_INTEGER;
  l_row_uncmp        BINARY_INTEGER;
  l_cmp_ratio        NUMBER;
  l_comptype_str     VARCHAR2(100);
BEGIN
     dbms_compression.get_compression_ratio(
   -- input parameters
       scratchtbsname   => 'USERS',                    -- scratch tablespace
       ownname          => 'SSB',                      -- owner of the table
       objname          => 'LINEORDER',                -- table name
       subobjname       => NULL,                       -- partition name
       comptype         => DBMS_COMPRESSION.COMP_INMEMORY_QUERY, -- compression algorithm
   -- output parameters
       blkcnt_cmp       => l_blkcnt_cmp,               -- number of compressed blocks
       blkcnt_uncmp     => l_blkcnt_uncmp,             -- number of uncompressed blocks
       row_cmp          => l_row_cmp,                  -- number of rows in a compressed block
       row_uncmp        => l_row_uncmp,                -- number of rows in an uncompressed block
       cmp_ratio        => l_cmp_ratio,                -- compression ratio
       comptype_str     => l_comptype_str              -- compression type
     );
     dbms_output.put_line('LINEORDER '||l_comptype_str||' ratio: '||to_char(l_cmp_ratio,'99.999'));
END;
```

- Easy way to determine memory requirements
- Use DBMS_COMPRESSION
- Applies MEMCOMPRESS to sample set of data from a table
- Returns estimated compression ratio

# Oracle In-Memory Advisor (2015. február 23.)

| Object Type | Object | Estimated In-Memory Size | Analytics Processing Seconds | Estimated Reduced Analytics Processing Seconds | Estimated Analytics Processing Performance Improvement Factor | Benefit / Cost Ratio (Improvement Factor / In-Memory Size) |
|---|---|---|---|---|---|---|
| Table | SOE.LOGON | 451.76MB | 2114 | 1,887 | 9.3X | 20.586 |
| Table | SOE.CARD_DETAILS | 607.32MB | 8346 | 7,248 | 7.6X | 12.514 |
| Table | SOE.ADDRESSES | 1.09GB | 5237 | 4,621 | 8.5X | 7.798 |
| Partition | SOE.PRODUCT_MOCKUP.Y2014Q1 | 812.6MB | 2003 | 1,489 | 3.9X | 4.799 |
| Table | SOE.CUSTOMERS | 1.10GB | 108 | 95 | 8.2X | 7.455 |
| Table | SOE.ORDER_ITEMS | 2.19GB | 7128 | 6,393 | 9.7X | 4.429 |
| Table | SOE.ORDERS | 1.34GB | 3512 | 2,917 | 5.9X | 4.403 |
| Table | SOE.PRODUCT_INFORMATION | 1.78MB | 2873 | 2,205 | 4.3X | 2.416 |
| Partition | SOE.PRODUCT_MOCKUP.Y2013Q4 | 1.62GB | 97 | 1,489 | 3.7X | 2.284 |
| Partition | SOE.PRODUCT_MOCKUP.Y2014Q2 | 3.37GB | 642 | 493 | 4.3X | 1.276 |

- Új In-Memory Advisor
- DB workload elemzése, AWR & ASH repository
- Tuning Pack
- Felsorolja azokat az objektumokat, melyek legjobban profitálnak analitikus workloadban az IM column store-ból

ORACLE
CLOUD


Process


Documents


Social


Business Intelligence


Big Data

# Oracle Cloud - Platform as a Service


Database


Java


Developer


Mobile


Integration


Compute


Storage


Messaging


Identity


Systems Monitoring & Analytics

ORACLE

# Optimized from Disk to Memory

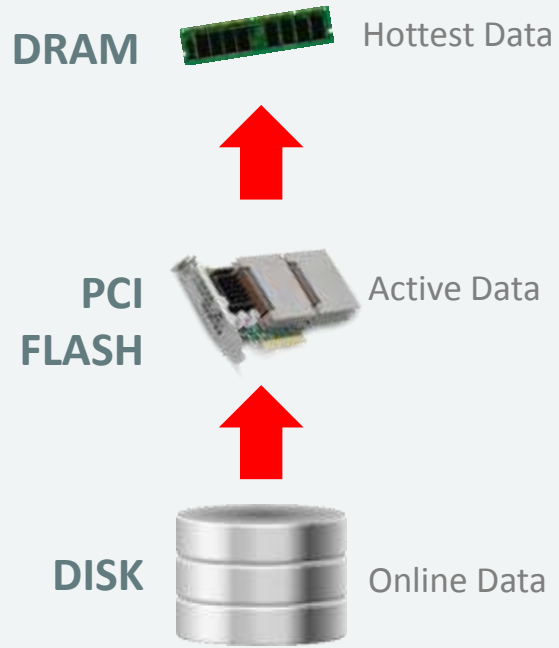


- Oracle Database In-Memory
- Exadata Smart Flash Cache
- Exadata Storage Servers

100's GB/sec

50-100 GB/sec

10's GB/sec

# Hardware and Software
## Engineered to Work Together


ORACLE®