ORACLE

Planet-Scale Database and Blockchain for the Masses

Wei Hu

Senior Vice President
Mission Critical Database Technologies

A globally distributed database is a database where the data can be stored physically stored in multiple data centers across the globe

Benefits of Globally Distributed Database

- In a global research consortium, labs are distributed around the world
- Labs stores their own data, but want all the members to see all the data

- Similar to how multi-national companies need to handle Data Residency – data for citizens of a country must be stored within that country, but still want to see the data across the globe
- Disaster Recovery – primary and standby(s) can be far apart
- Low latency for reads – data replicated to multiple regions. Clients in each region read from nearest replica
- Low latency for writes – for example, if we locate European data in European data center
- Single logical database – application and administrator sees one database (e.g., one service) even though physically there are N databases and replicas

What Is Database Sharding?

Dominant approach for scaling Internet applications

- Horizontal partitioning of data across independent databases (shards)
  - Each shard holds a subset of the data
  - Can be single-node or RAC or PDB
  - Replicated for high availability
- Shared-nothing architecture:
  - Shards don't share any hardware (CPU, memory, disk), or software (clusterware)

A single **logical DB** sharded into N physical Databases

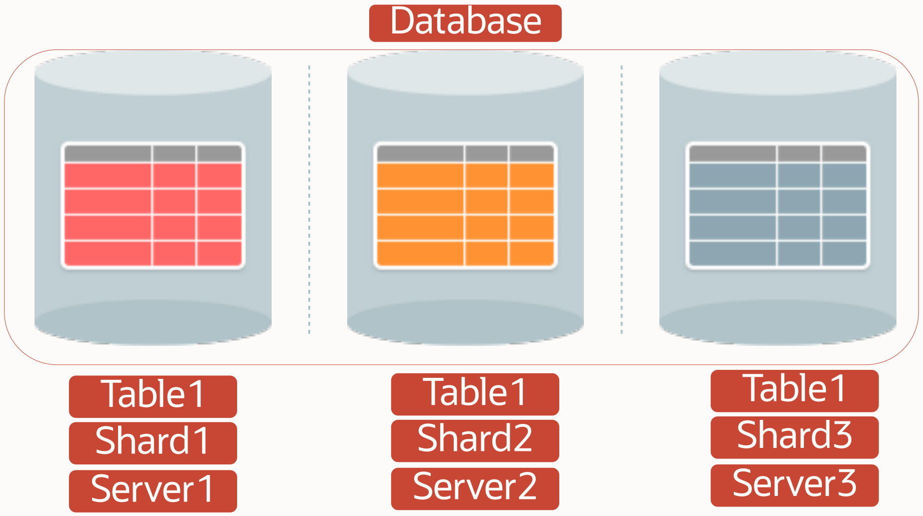

Oracle Database Sharding – Benefits

Linear Scalability

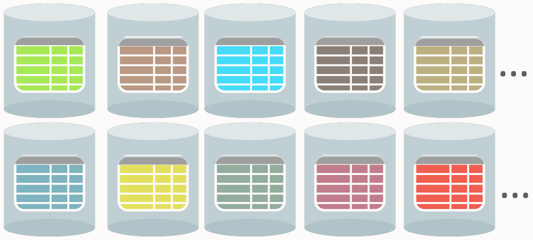

Add shards online to increase database size and throughput. Online split and rebalance.

Extreme Availability

Shared-nothing architecture. Fault of one shard has no impact on others.

Geographic Distribution

User defined data placement for performance, availability, DR or to meet regulatory requirements.

Composite Sharding

Geographic Distribution and Linear Scalability

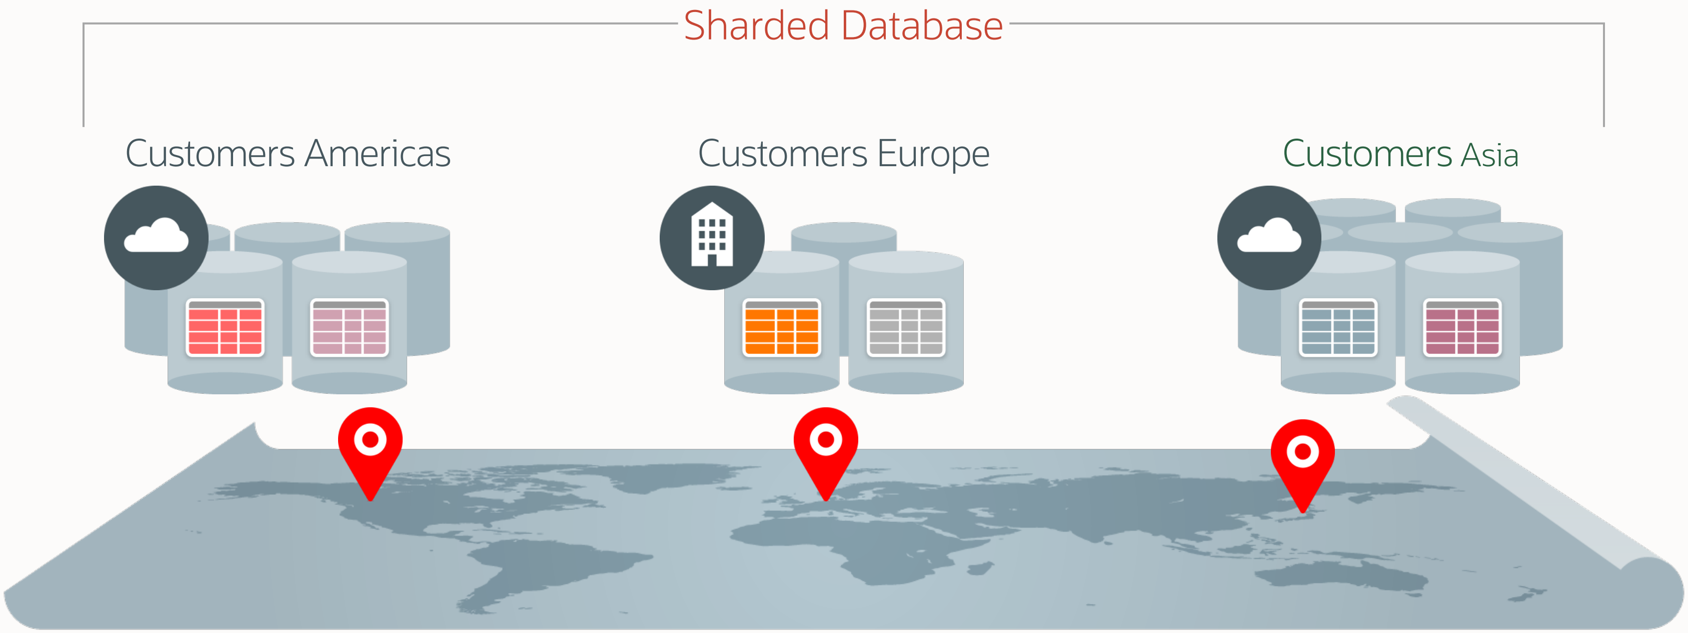

One Logical Database Using Sharding

Use ‘country’ (‘region’, or ‘institution’) as the top-level sharding key

In-country applications connect to local database as usual

- All data stored locally and processed locally

Multi-shard queries directed to the Coordinator DB

- Coordinator DB rewrites query and sends to each shard (country) that has data
- Coordinator processes and aggregates the results from all the countries and returns result
- Coordinator can be replicated for availability and performance

Think globally, store where needed

- In-country data storage plus supports global view of all the data

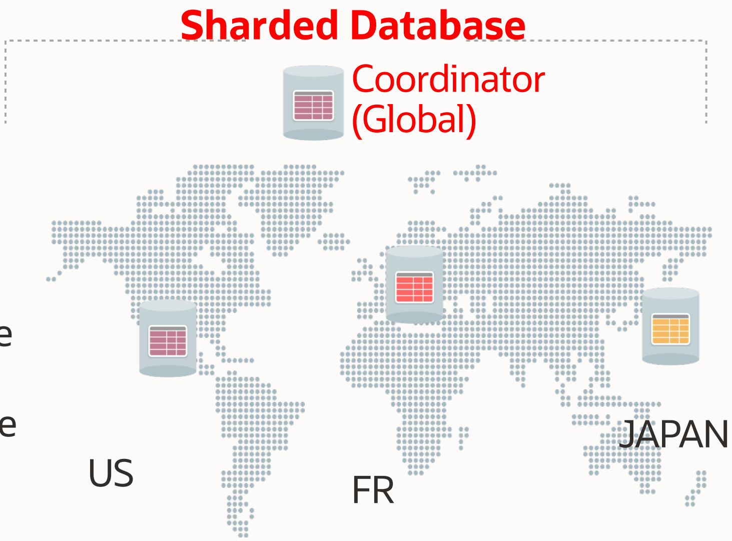

Sharded Database

Coordinator (Global)

US

FR

JAPAN

Sharded Database Support for this Architecture

1. Direct-to-shard routing for in-country queries and multi-shard queries for queries that access data across countries
2. User defined sharding to allow you to use range or list of countries to partition data among shards
3. Automatic configuration of replication and control where the replicas are located
4. Data federation support for adding existing databases to a sharded database

This Architecture Supports Multi-Cloud

- Each shard can be in a cloud or on-premises within the country
- Shards can use different cloud providers (multi-cloud strategy)
- Replica of a shard can be in a different cloud or on-premises
- Online resharding allows you to move data between clouds, or to/from cloud and on-premises
- And supports *real* standards-based SQL

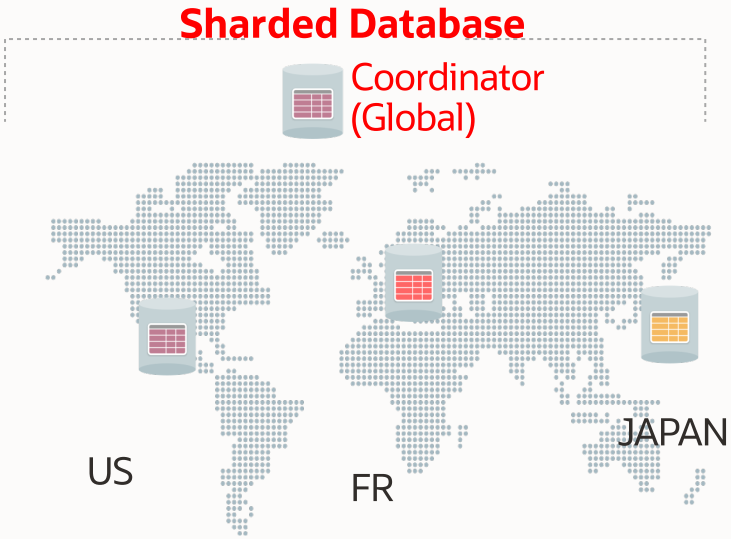

Why You Need Crypto-Secure Data Management

Blockchain for the Masses

Protecting Data from Illicit Changes

Existing data security technologies focus on **keeping criminals out**

Passwords

Privileges

Encryption

Firewalls

Protecting Data from Illicit Changes

Existing data security technologies focus on **keeping criminals out**

Passwords **Privileges** **Encryption** **Firewalls**

Blockchain adds another layer of data security

- Protects data from being illicitly modified or deleted

Uses cryptographic digests, cryptographic signatures, and distributed systems to implement **crypto-secure data management**

Blockchain is Compelling but Difficult to Implement

According to Deloitte's Global Blockchain Survey[1]:

> "Blockchain's standing as a **top-five strategic priority** continued in 2020"
> "Some 86% of respondents … said that blockchain offers a compelling business case"

However, blockchain projects have a notoriously high failure rate. IDC says[2]:

> "the technical complexities and costs connected to the creation and configuration of a blockchain and to maintain the infrastructure are often seen as a barrier"

To realize the promise of blockchain, it needs to be much **easier to implement**

[1] Deloitte's Global Blockchain Survey 2020

[2] IDC Report: Barriers to Blockchain Adoption in Europe, March 2020, Doc # EUR145773720

Blockchain for the Masses

Up to now, blockchain has mainly been used to develop **decentralized multi-party apps**

- Blockchain required new applications, new development methodologies, new data management products, and new business processes

Oracle has now released breakthrough technology that deeply integrates blockchain into the Oracle database

- Makes it possible to implement blockchain in **mainstream** applications with **minimal application changes**
- Provides the full functionality of the world's leading database on crypto-protected data

The benefits of blockchain are now available beyond decentralized multi-party apps

Preventing Illicit Changes to Data

Oracle blockchain focuses on **preventing illicit changes** to data that records important actions, assets, entities, and documents

Preventing Illicit Changes to Data

Oracle blockchain focuses on **preventing illicit changes** to data that records important actions, assets, entities, and documents

- Examples: contracts, property titles, payments, transfers, ledgers, account statements
- Data that must be **preserved and never modified**
- Data that can be deleted only when it is **obsolete**

Illicit changes to important records can have dire consequences

- Loss of assets, legal exposure, loss of reputation, etc.

Oracle crypto-security is not limited to providing change ledgers

- It can directly protect reference data as well

New Security Hacks are Unavoidable

There were **2,709** critical security vulnerabilities (CVSS 9-10) published in 2020 alone

https://nvd.nist.gov/vuln/search/

Stolen Insider Credentials are Unavoidable:

“Dozens of seniors at Gadsden High School used a logon that is normally assigned to a teacher, to change their grades.”

https://cbs4local.com/news/local/dozens-of-gadsden-high-school-students-caught-changing-grades-some-wont-graduate

SONY
PICTURES

Phishing of Insiders is Unavoidable:

“Hackers gained access to Sony's network after a series of **phishing emails** aimed at system engineers and network administrators who were asked to verify their Apple IDs”

https://www.computerworld.com/article/2913805/sony-hackers-targeted-employees-with-fake-apple-id-emails.html

TESLA

Bribing of Insiders is Unavoidable:

“Employee is offered $500,000 to insert malware into Tesla system”

To his credit, the employee reported the bribe attempt

https://www.teslarati.com/tesla-employee-fbi-thwarts-russian-cybersecurity-attack/

Limiting the Damage from Unavoidable Vulnerabilities

Although we can never eliminate these vulnerabilities

- We can limit the damage they cause

Blockchain will not prevent your data from being stolen

- But it provides an additional layer of defense that can prevent the dire consequences of your data being changed

Which is worse:

- Someone seeing my bank account
- Someone stealing the money in my bank account

?

Use Cases for Blockchain in the Oracle Database

Oracle Blockchain Use Cases

Because Oracle makes it so easy to adopt blockchain technology, it is now usable in every industry and in a majority of applications

- **Financial data** - accounting, assets, payments, insurance
- **Logistics data -** distribution, supply chain, shipment, recalls
- **Education data –** Degrees, certifications, and professional history
- **Government data-** legal, trial, tax, permits, citizenship, title
- **Scientific data** – secure notebooks, experimental configurations, results, hypotheses, , intellectual property

Prevent Illicit Changes by Insiders* using

Immutable Tables

Insiders

Hackers

Authorities

Impersonators

*Or Criminals Using Insider Credentials

Preventing Illicit Changes Using Insiders Credentials

Oracle **Immutable Tables** prevent illicit modifications by insiders using the database (SQL)

- For example, changes by rogue employees or employees whose credentials were stolen

Immutable tables allow new data to be added, but existing data **cannot be changed or deleted by anyone using the database**

- Even Database Administrators

Oracle **Immutable** Tables

```
CREATE IMMUTABLE TABLE
trade_ledger (…);
```

TRADE LEDGER

ID	User	Value
1	Tom	500
2	Carol	176
3	Wang	500
4	Eve	25

Available in Database 19.11, 21.3

Create an insert-only Immutable Table by simply adding “Immutable” to a table creation statement

- Allows relational data and JSON or LOB documents
- Not restricted to ledgers, can store reference data

Behaves like any other table except that Immutable Table users cannot:

- Update or delete rows
- Change table definition
- Convert immutable table to updatable or vice-versa
- Modify table metadata in database dictionary

Using Immutable Tables requires no App changes

Insiders Hackers Authorities Impersonators

Detect Illicit Changes by Hackers using

Cryptographic Data Chaining

Detecting Illicit Changes that Bypass Database Software

What if the database software that enforces immutability is **bypassed**

- For example, by a hacker using a new security vulnerability
- Or by a rogue or compromised **system administrator** who uses OS credentials to modify or bypass the database software (inside hacker)

Blockchain **detects** changes to data by computing and safely storing a small cryptographic digest of the data

- If the data is changed, then the the cryptographic digest of the changed data will differ from the previous digest

Detects illicit changes, even if attackers take over full control of the OS or database

Oracle Blockchain Table Verification

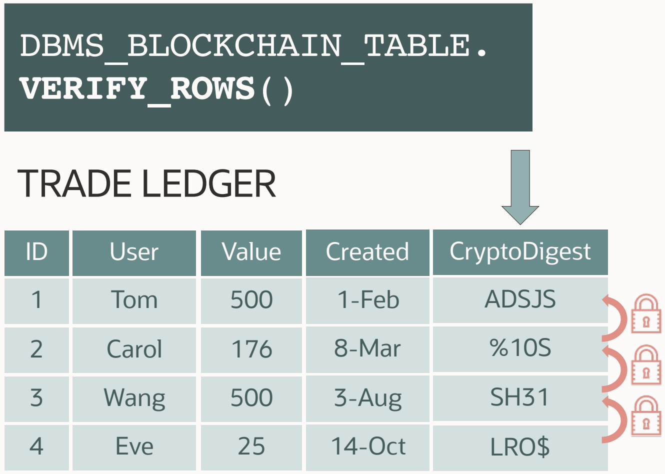

```
DBMS_BLOCKCHAIN_TABLE.
VERIFY_ROWS()
```

TRADE LEDGER

ID	User	Value	Created	CryptoDigest
1	Tom	500	1-Feb	ADSJS
2	Carol	176	8-Mar	%10S
3	Wang	500	3-Aug	SH31
4	Eve	25	14-Oct	LRO$

Chains are verifiable by users:

- Call `DBMS_BLOCKCHAIN_TABLE.VERIFY_ROWS()`

Chains can also be **independently validated** without depending on the database

- Open Standard cryptographic algorithms are used
- Allows users to the read the chained data and verify the crypto-digest from outside the database

Easy to Use

```
CREATE BLOCKCHAIN TABLE
trade_ledger (…);
```

TRADE LEDGER

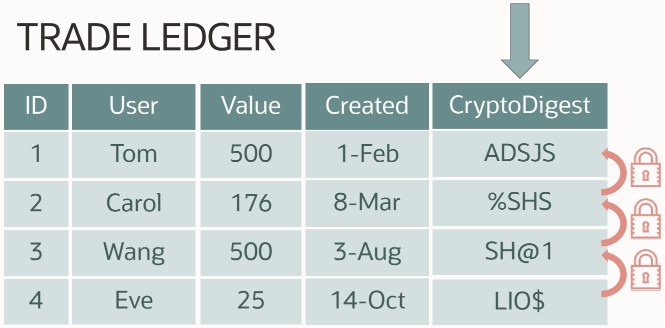

ID	User	Value	Created	CryptoDigest
1	Tom	500	1-Feb	ADSJS
2	Carol	176	8-Mar	%SHS
3	Wang	500	3-Aug	SH@1
4	Eve	25	14-Oct	LIO$

Available in Database 19.10, 21.1

Immutable and Blockchain Tables are Oracle Converged Database **features**

- Not tied to **blockchain-only database** like AWS QLDB

Requires no application changes

- Accessed with standard declarative SQL
- Supports full analytics and transactions on table data
- Highly scalable and available using
  - Oracle RAC scale-out clusters
  - Data Guard geographic disaster protection

FREE in all Oracle Database Editions

Oracle Blockchain Use Cases

Because Oracle makes it so easy to adopt blockchain technology, it is now usable in every industry and in a majority of applications

- **Financial data** - accounting, assets, payments, insurance
- **Logistics data -** distribution, supply chain, shipment, recalls
- **Education data –** Degrees, certifications, and professional history
- **Government data-** legal, trial, tax, permits, citizenship, title
- **Scientific data** – secure notebooks, experimental configurations, results, hypotheses, , intellectual property

Comparison to Peer-to-Peer Blockchain

How Blockchain Tables Compare to Peer-to-Peer Blockchain

A Blockchain Table with **publicly distributed signed digests** is a distributed blockchain but not a **peer-to-peer** blockchain

- It is simpler to use and has many of the benefits of full **peer-to-peer**

Blockchain Tables **prevent** illicit changes made **using the database**, and **detect** illicit changes that **bypass the database**

Directly preventing all forms of illicit changes provides stronger protection

- However, detection of illicit changes also indirectly prevents most fraud
- Criminals won't commit fraud if they know they will be caught
  - For example, a criminal won't take a bribe while being video-recorded

Peer-to-Peer Blockchain Complexity

To **prevent** illicit changes that bypass the database requires a **Peer-to-Peer** Blockchain architecture

- With **consensus-based** multi-party agreements for every change
- Peer-to-Peer consensus prevents the takeover of a single database from compromising the entire system
- However, peer-to-peer is still vulnerable to a common bug across all peers

Peer-to-Peer consensus requires new app development methodologies

Also, consensus requires major changes to **businesses processes** and how businesses interact with customers and partners

- **Consensus requires sacrificing autonomy and privacy**

Overall, peer-to-peer provides very effective illicit change prevention, but at the expense of significant complexity for mainstream use-cases

Blockchain Adds to Oracle's Unique Maximum Security Architecture

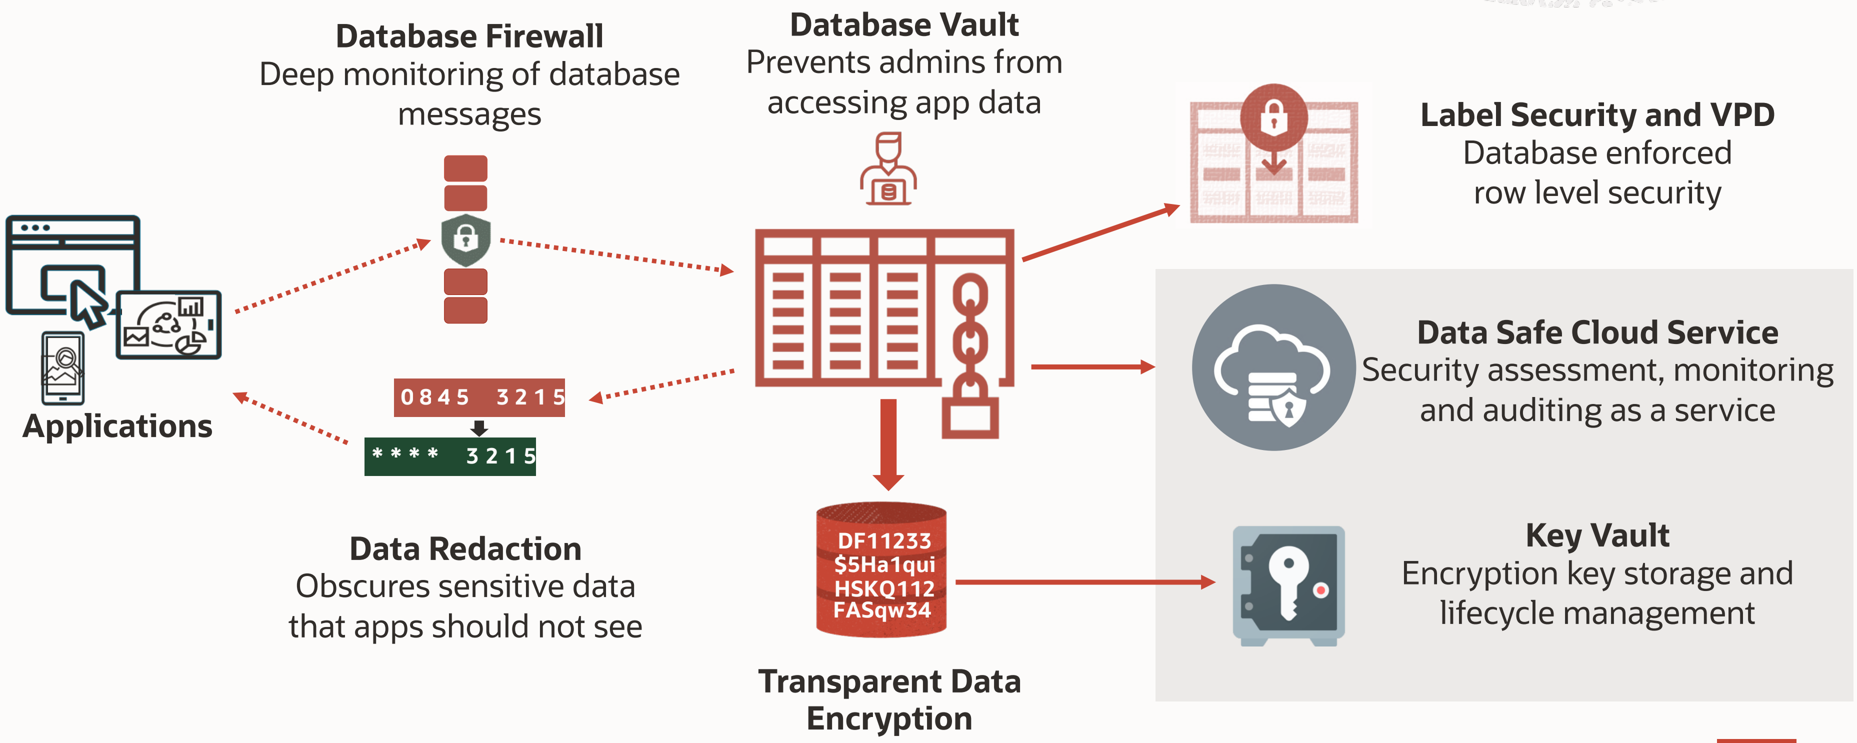

Applications

Database Firewall
Deep monitoring of database messages

Database Vault
Prevents admins from accessing app data

Label Security and VPD
Database enforced row level security

Data Safe Cloud Service
Security assessment, monitoring and auditing as a service

Data Redaction
Obscures sensitive data that apps should not see

Key Vault
Encryption key storage and lifecycle management

Transparent Data Encryption

Oracle Converged Database Supports all Leading-Edge Data and Workloads

Dramatically Simpler Development and Management

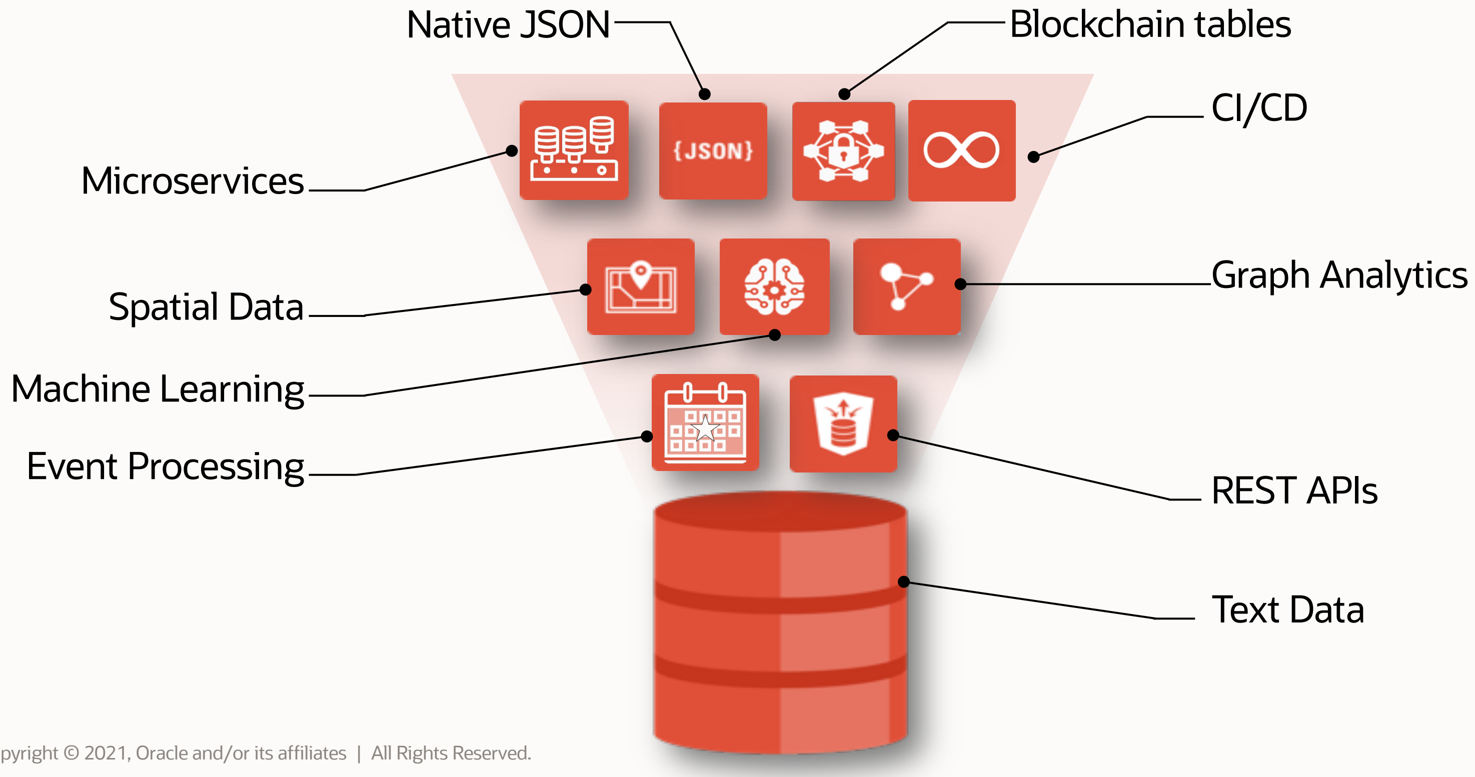

Blockchain in the Oracle Database Addresses Real-World Data Protection Challenges

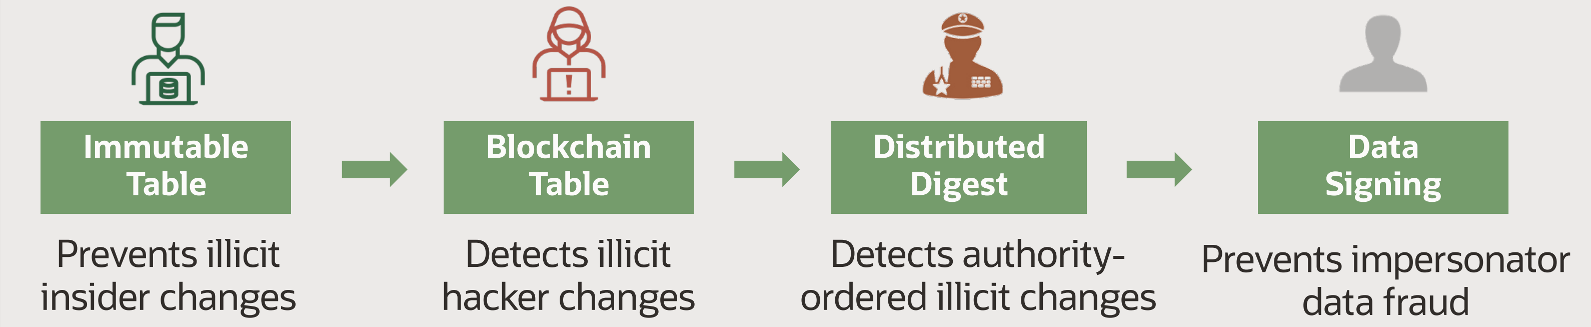

Immutable Table

Prevents illicit insider changes

→

Blockchain Table

Detects illicit hacker changes

→

Distributed Digest

Detects authority-ordered illicit changes

→

Data Signing

Prevents impersonator data fraud

Free features of Oracle Converged Database

Deploy Incrementally with Minimal App Changes

More Information:

Blockchain Blog
https://blogs.oracle.com/blockchain/

Blockchain Hands-on-Lab
https://apexapps.oracle.com/pls/apex/dbpm/r/livelabs

Blockchain Documentation
https://docs.oracle.com/en/database/oracle/oracle-database/21/nfcon/oracle-blockchain-table-268779556.html

Oracle for Research

Connecting cutting edge research with the power of cloud computing

Free Credits

Dedicated cloud credits to use on what you want, when you want it.

Open Data

Use open data sets or bring your own data, with no charge for ingress and free egress up to 10 TB/month.

Nurturing Communities

Connect with other researchers and research computing experts for knowledge sharing and collaboration.

Cloud Expertise

Technical support and functional advising to help you run your workloads faster.

Open Tooling

Popular research tools prebuilt images make setting up your Oracle cloud tenancy easy.

It's *your* IP

You control your intellectual property. Oracle doesn't monetize your data.

Faster Time to Results

Skip the queues for on-premises computing resources and use what you need, when you need it.

Open Standards

BYO stack – build and run on the tools you love.

Marketing

Collaborative marketing and signal boosting of research and white papers.

Focus on your research, rather than on computing challenges

ORACLE
for Research

Learn more and apply today at http://www.oracle.com/oracle-for-research. Alison Derbenwick Miller