ORACLE

Oracle Database Reliability

In the context of distributed databases

Francisco Munoz Alvarez

Distinguished Product Manager

Oracle Database High Availability (HA), Scalability and Maximum Availability Architecture (MAA) Team

 @fcomunoz

 http://www.linkedin.com/in/franciscomunozalvarez

 www.oraclemaa.com

Some key concepts

High Availability (HA) and Scalability Concepts

Shared Disk / Shared Cache

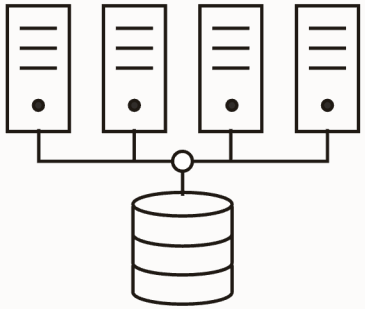

Replication / Read Replicas

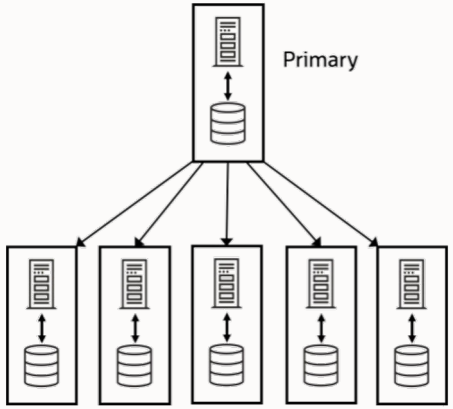

Sharding

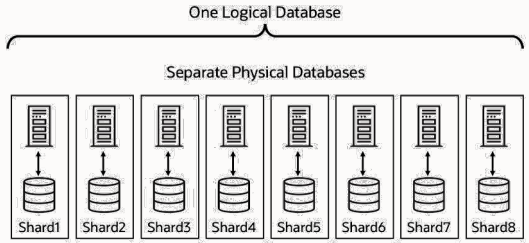

Combination thereof

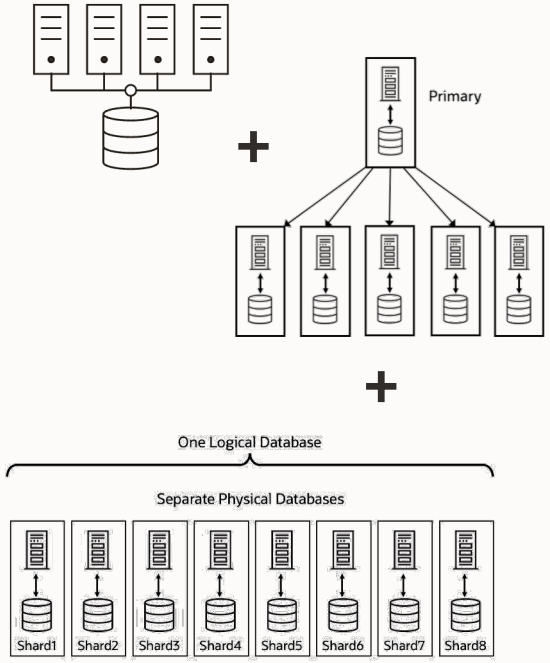

https://blogs.oracle.com/exadata/horizontal-scaling-with-oracle-database

Replication Types

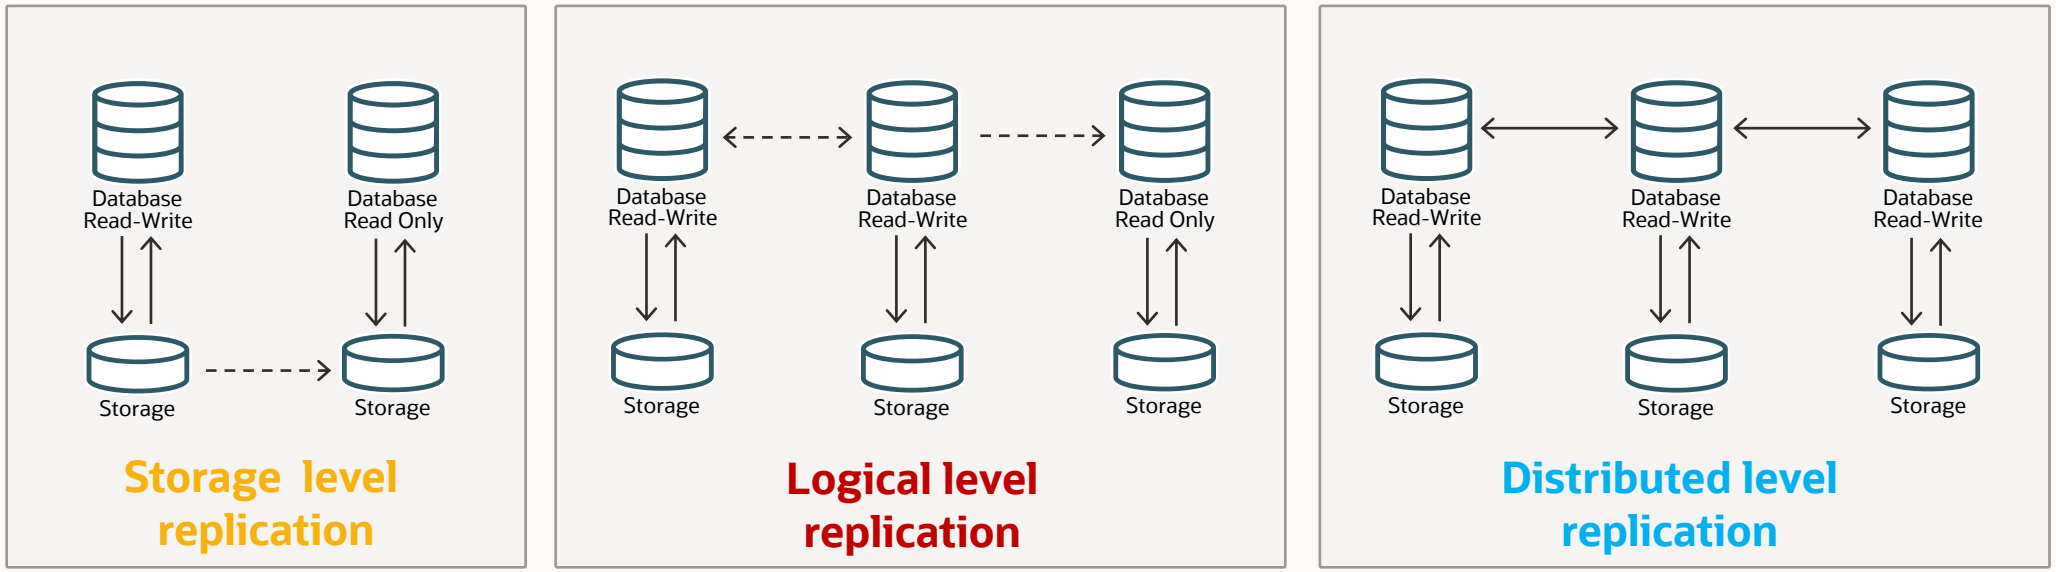

AWS Aurora
AlloyDB

Oracle
PostgreSQL
SQL Server

Spanner
CockroachDB
YugabyteDB

https://blogs.oracle.com/maa/post/about-the-usefulness-of-redo-apply-in-the-storage-tier

Same Concepts – Different Names

		Leader Database	Voting Database	Read-only Follower Database	Active-Active Replicated Database (***)
	Oracle DB	Primary	Node (RAC), Standby (Sync Affirm)	Read Standby (Active Data Guard)	Multi-Active Replica (GoldenGate)
NEW IN 23c	**Oracle Sharding (with RAFT replication)**	Leader (Raft*)	Replication Unit (RU), Replication Unit Leader	RO-Standby N/A	Active-Active DB N/A
	CockroachDB	Leader (Raft*)	Replica, Lease-Holder	Non-voting Replica	N/A
	YugabyteDB	Master (Raft*)	Follower, Lease-Holder	Read Replica /Observer Node (Not for DR) xCluster Replication (For DR)	N/A
	Google Spanner	Leader (Paxos**)	Read-Write Replica, Witness Replica, Lease-Holder	Read-only Replica	N/A

(***) Asynchronous active-active replication between independent Clusters

* https://en.wikipedia.org/wiki/Raft_(algorithm) ** https://en.wikipedia.org/wiki/Paxos_(computer_science)

Understanding RPO and RTO

Disaster Recovery protects your RPO

High Availability protects your RTO

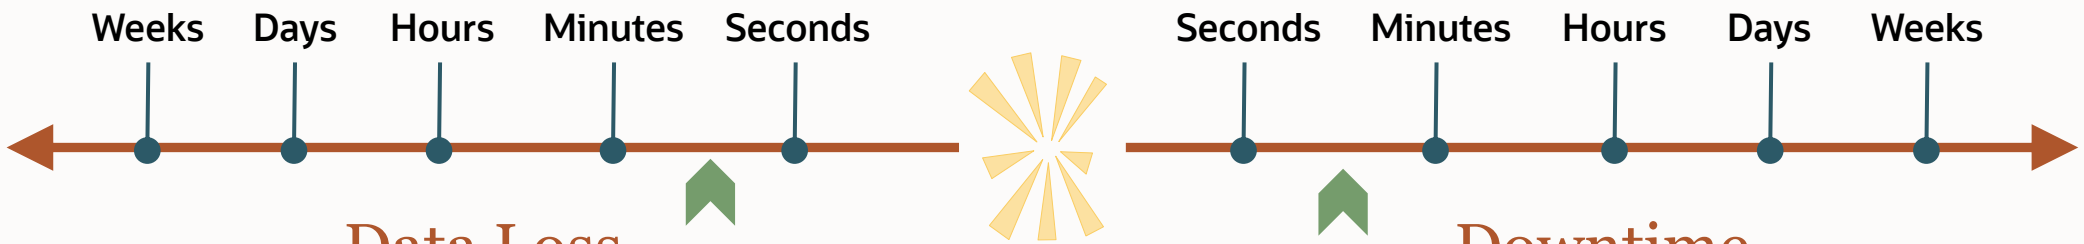

Recovery Point Objective (RPO)

Tolerance for data loss (sec's, hours, days); determines frequency of backups and replication approaches

Recovery Time Objective (RTO)

The shorter the Recovery Time Objective (RTO) the quicker you get back to business

Possible data loss impact

Business availability / continuity impact

Factors Influencing Data Reliability(*)

(*) https://en.wikipedia.org/wiki/Data_reliability

Top 4 Factors Influencing Reliability

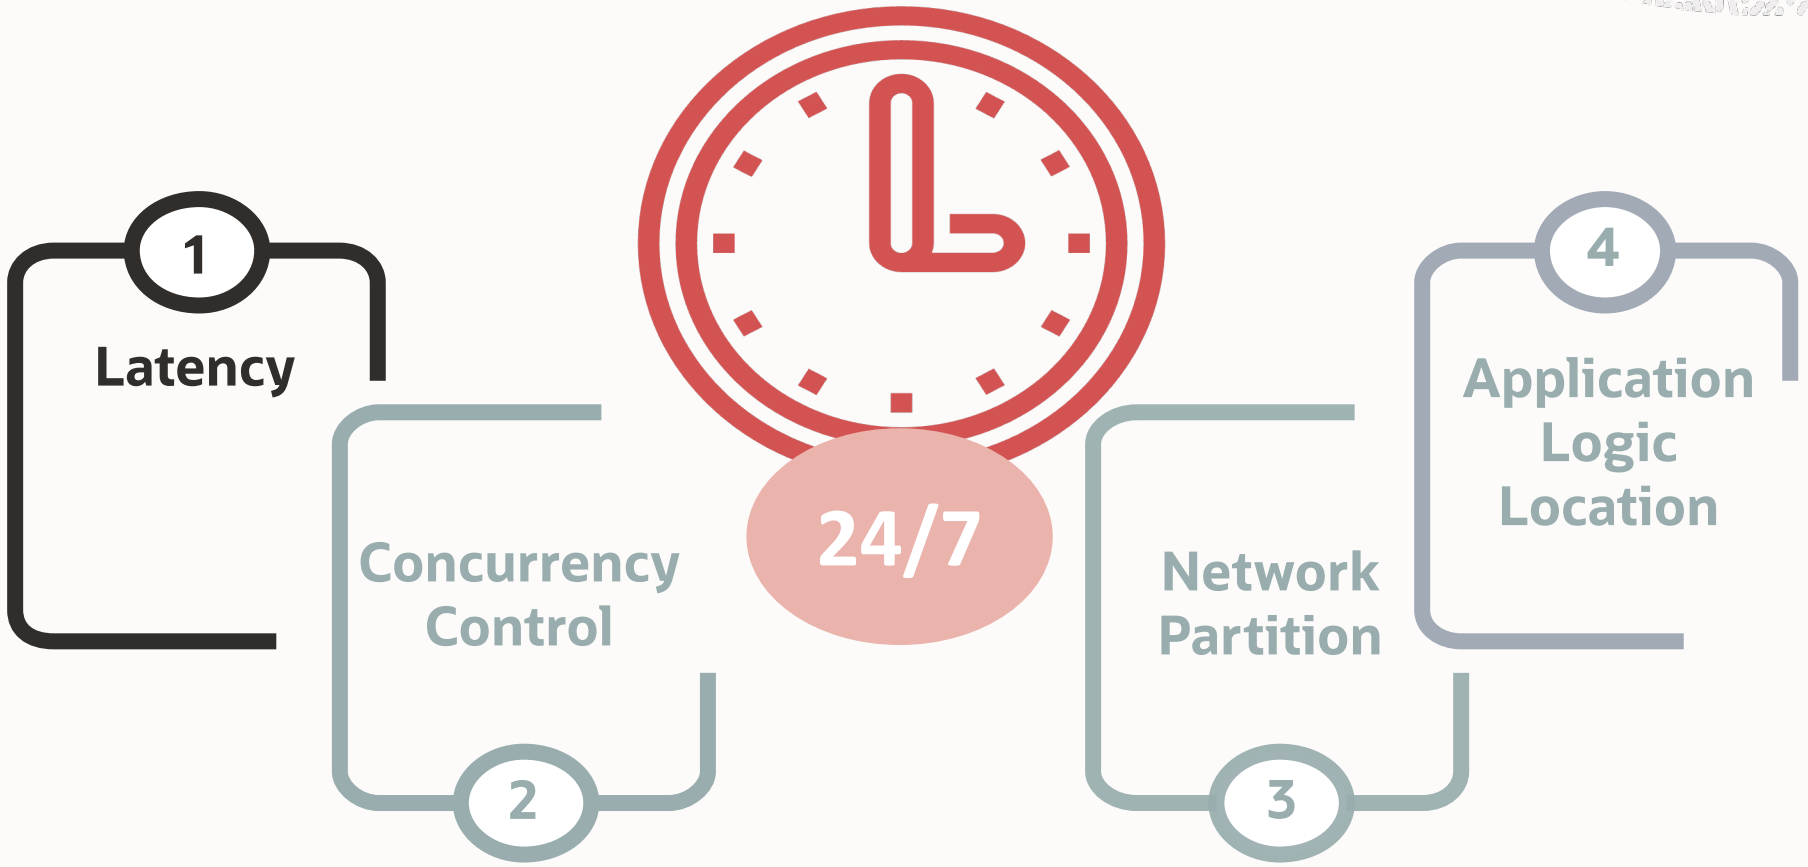

The Latency Problem: "I want RPO=0" vs. "Speed of Light"

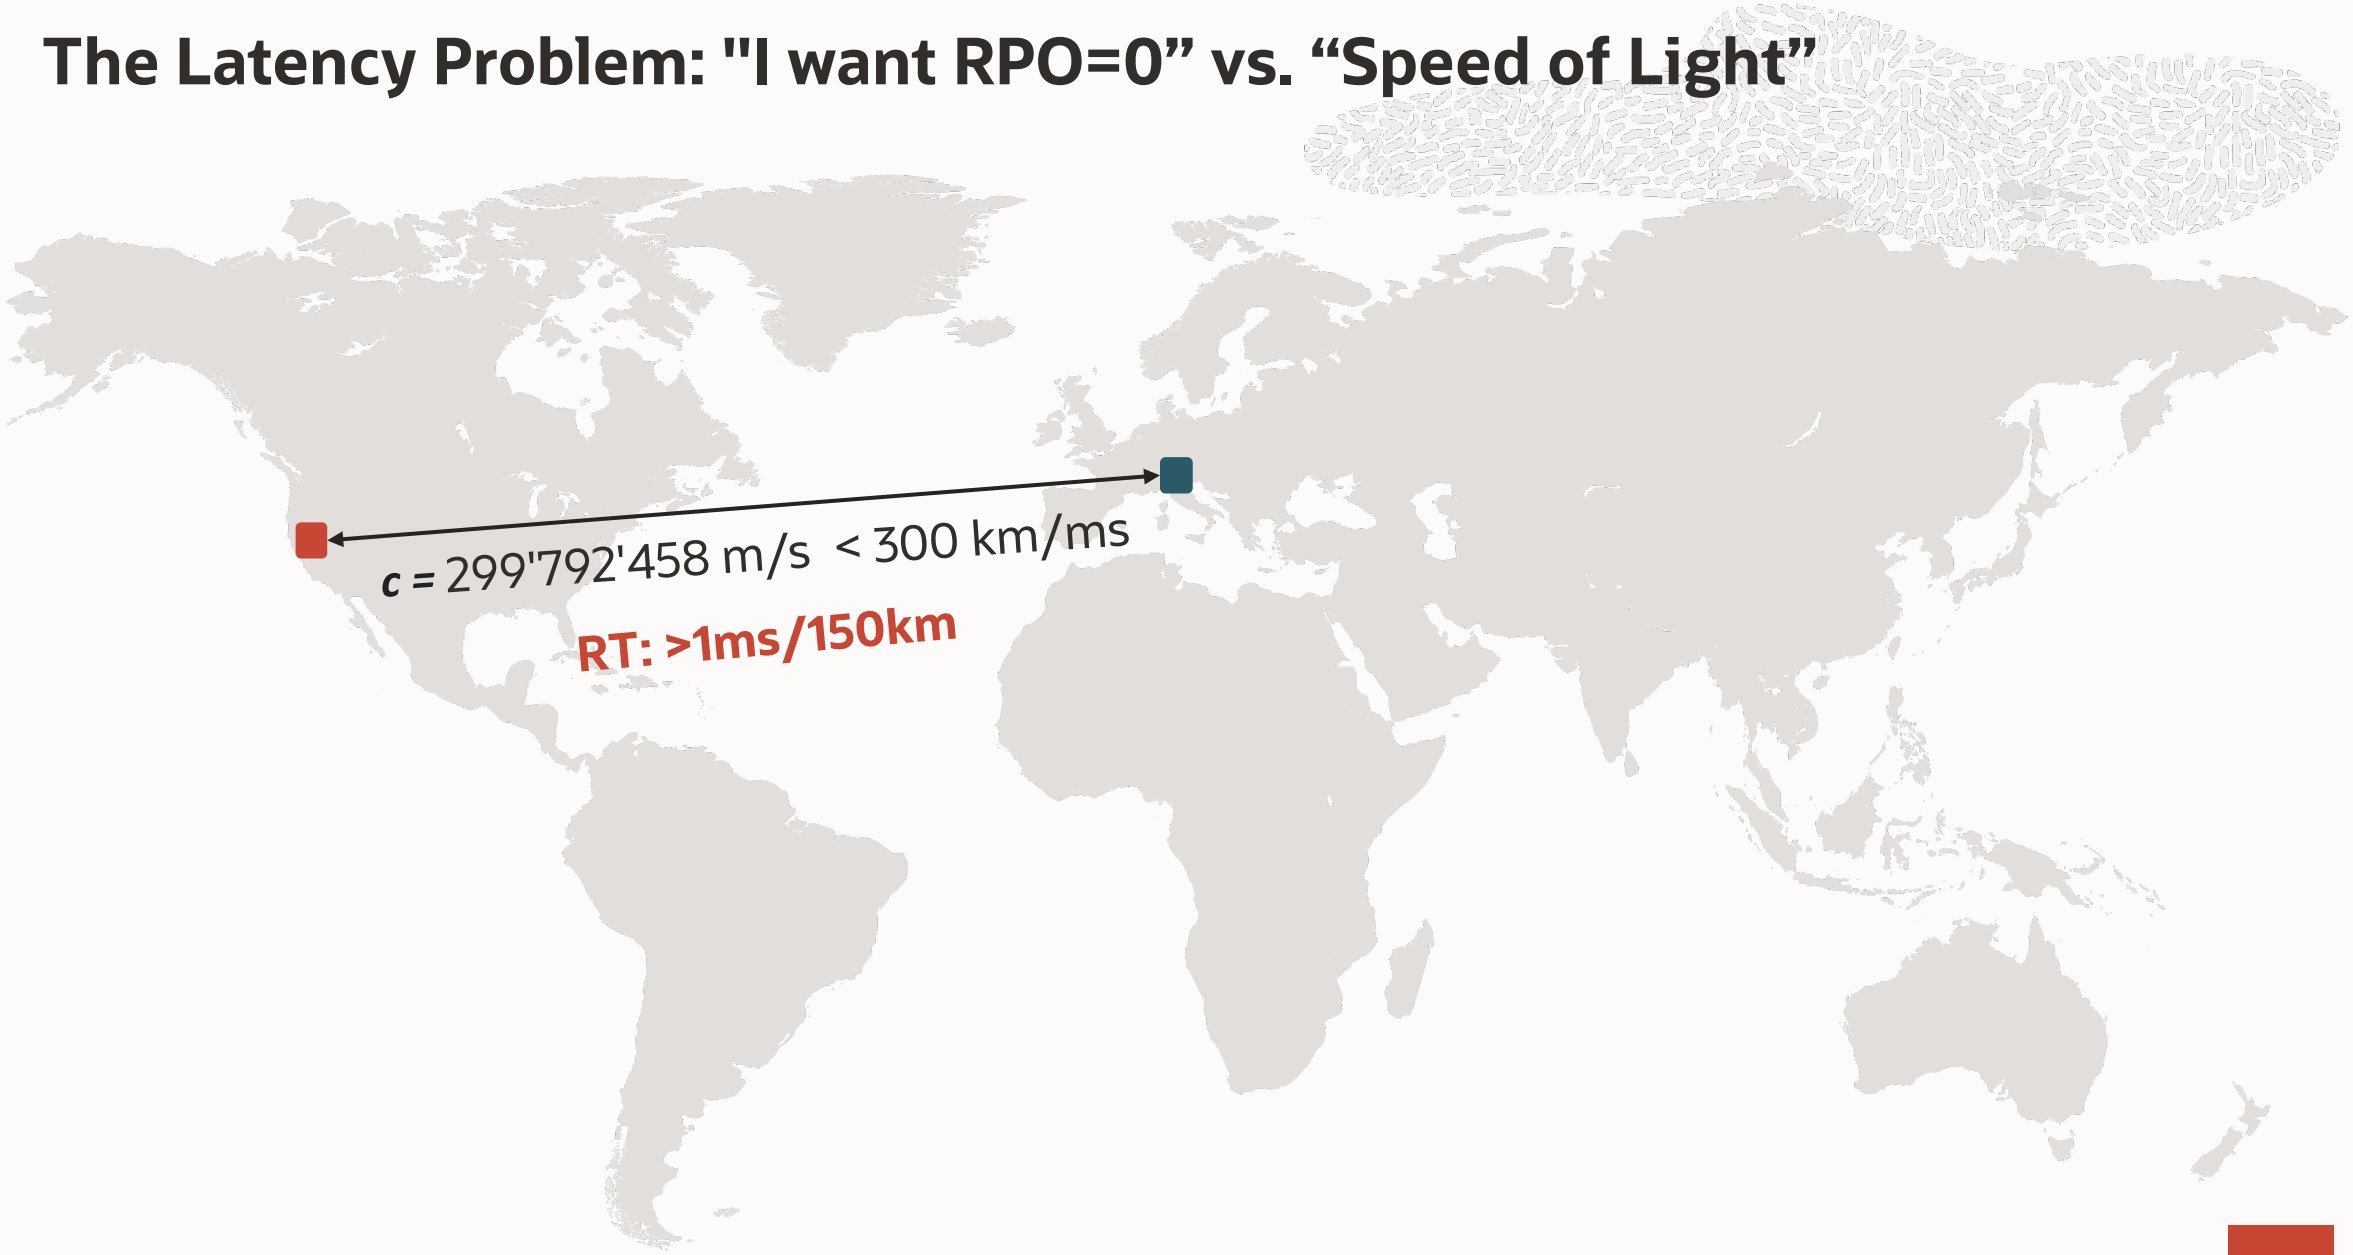

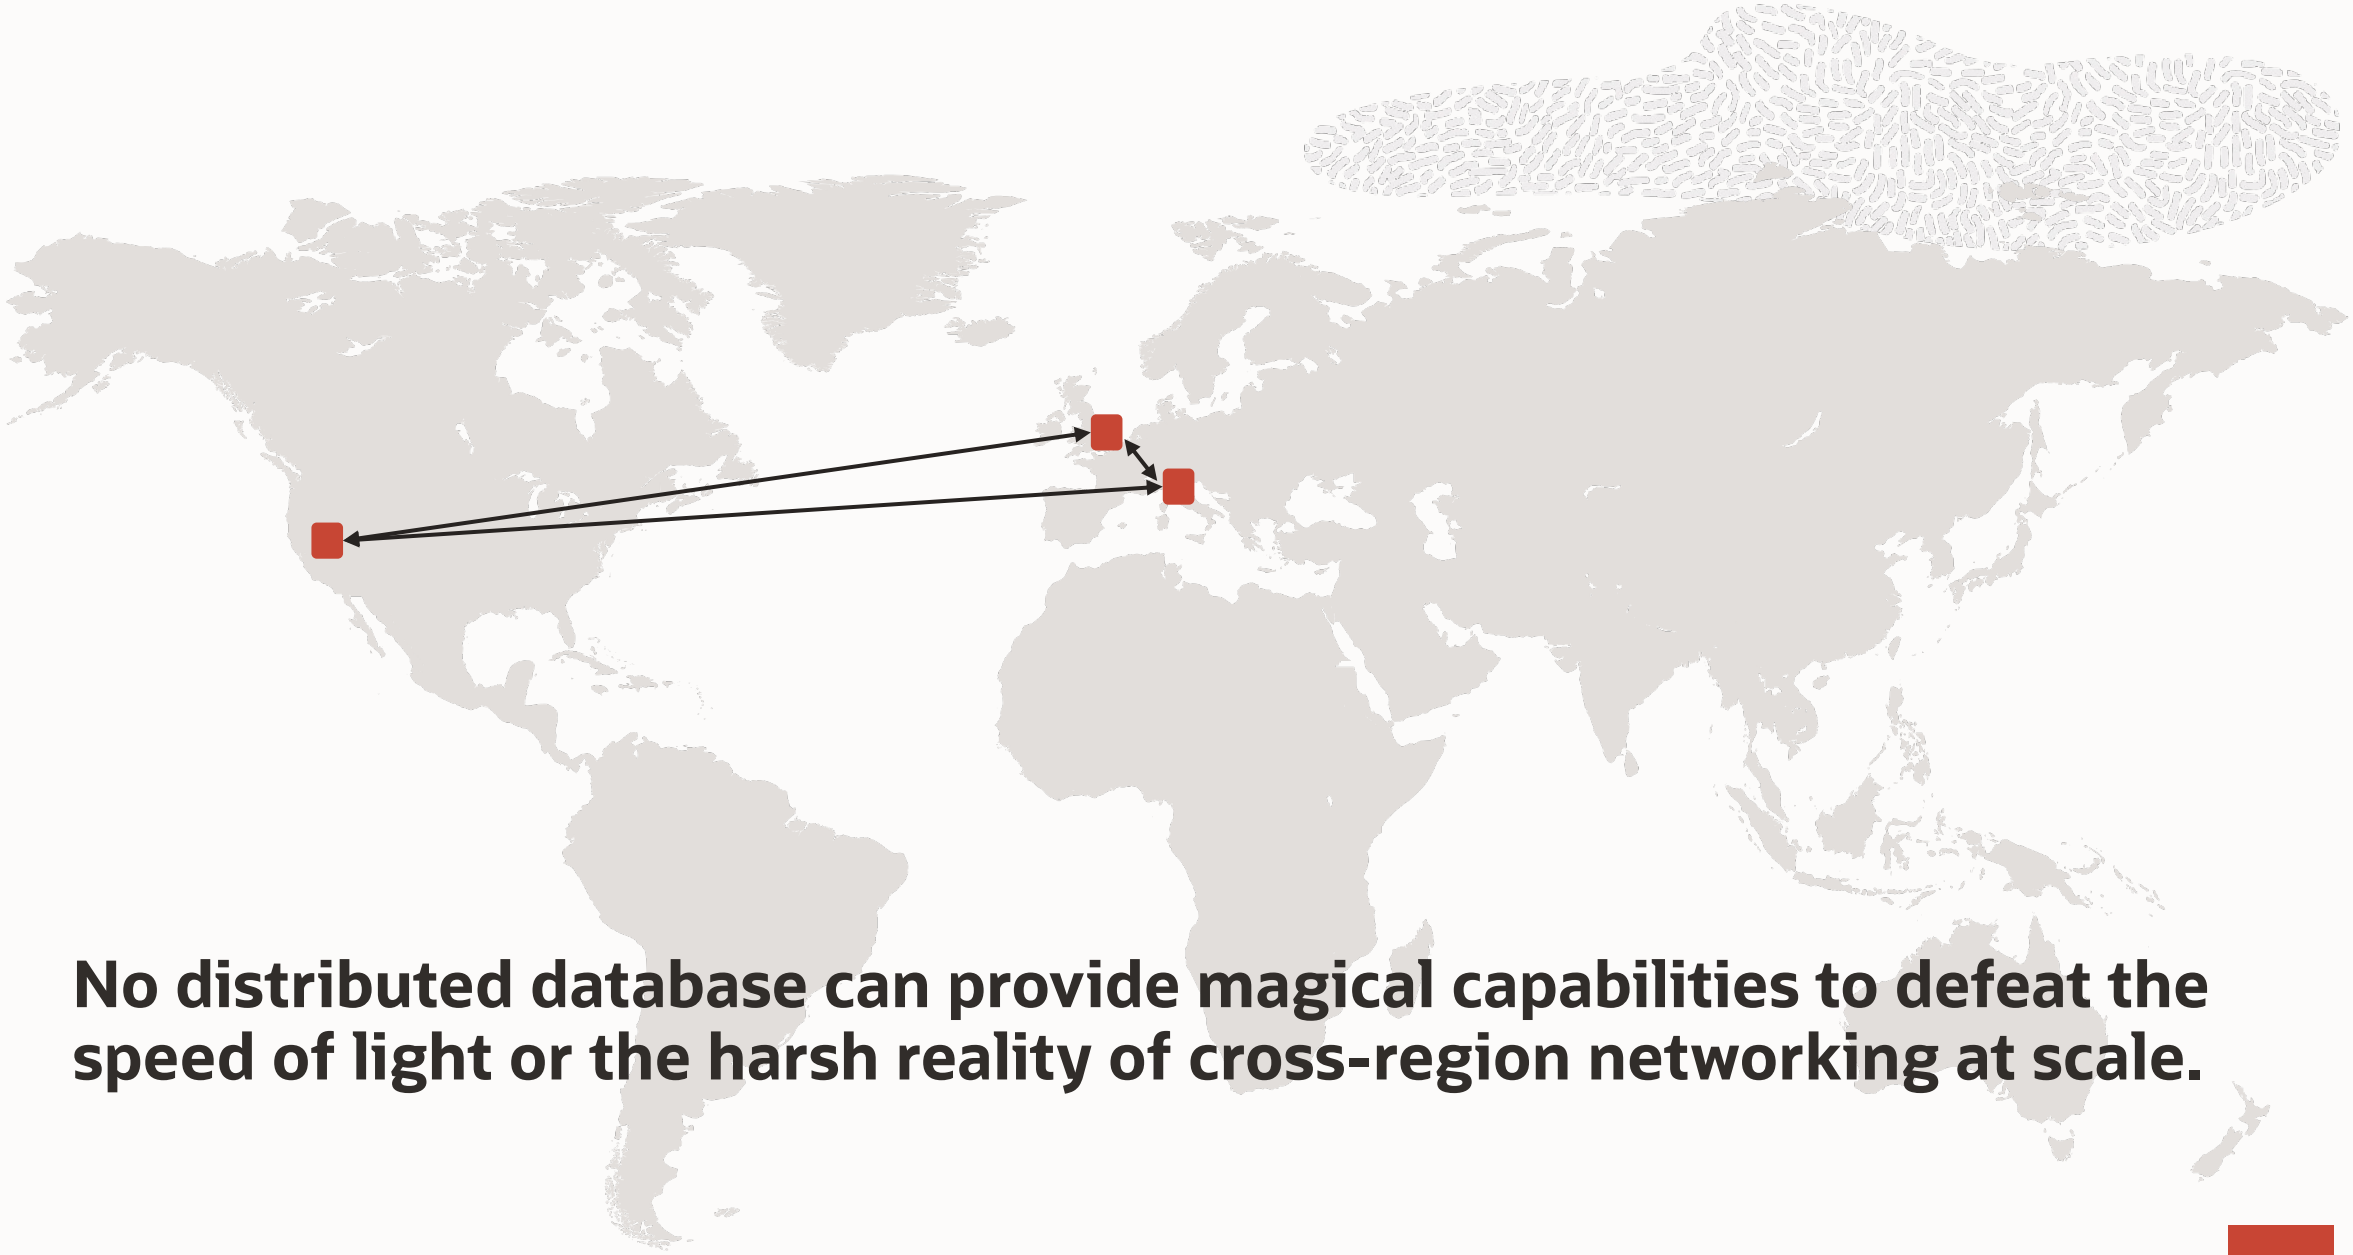

No distributed database can provide magical capabilities to defeat the speed of light or the harsh reality of cross-region networking at scale.

Different Approaches – Different Results

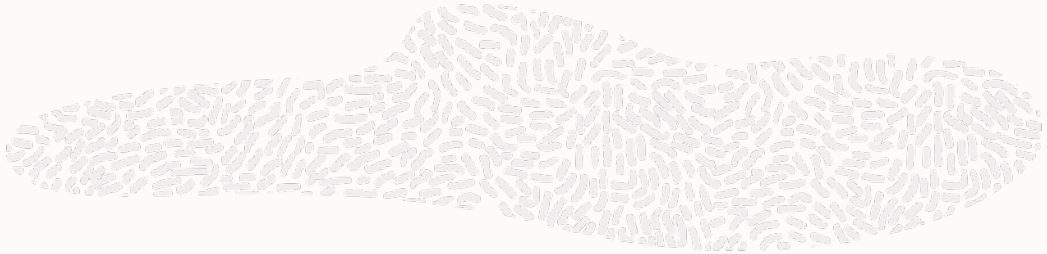

CockroachDB ensures Disaster Recovery (DR) capabilities only within the same region. No multi-region implementations are possible without incurring severe penalties to the applications using it (as per example, latency).

Oracle Database can easily fulfil any cross-region Disaster Recovery (DR) requirements, including multi-active (since 2009) and zero data loss at any distance.

Top 4 Factors Influencing Reliability

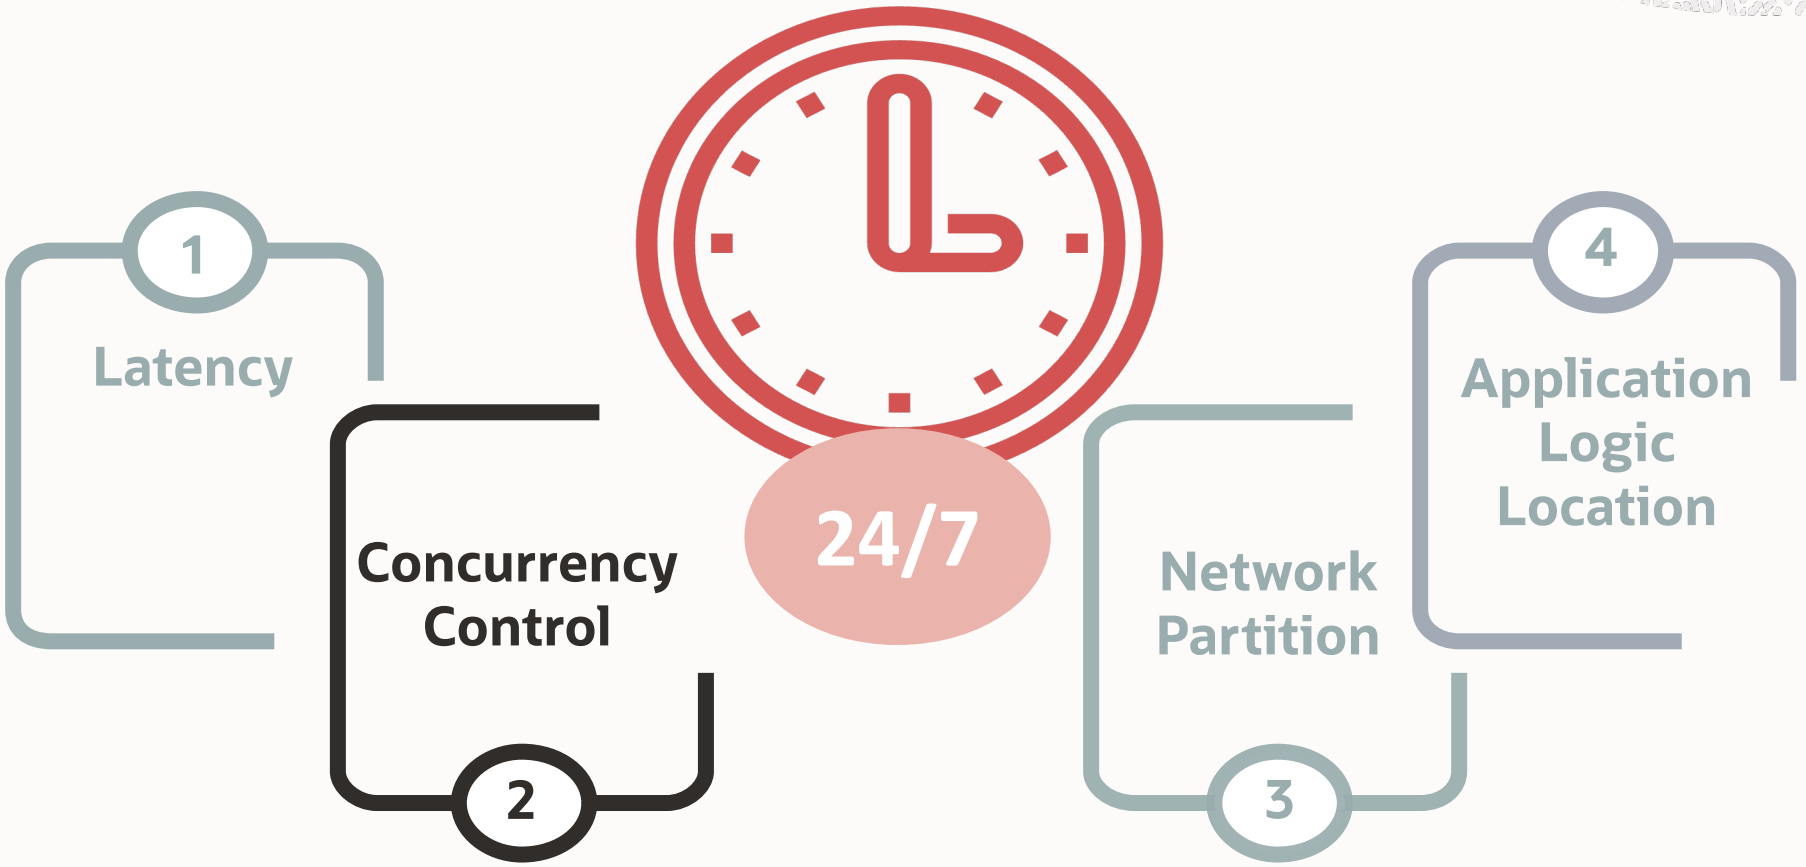

Concurrency Control

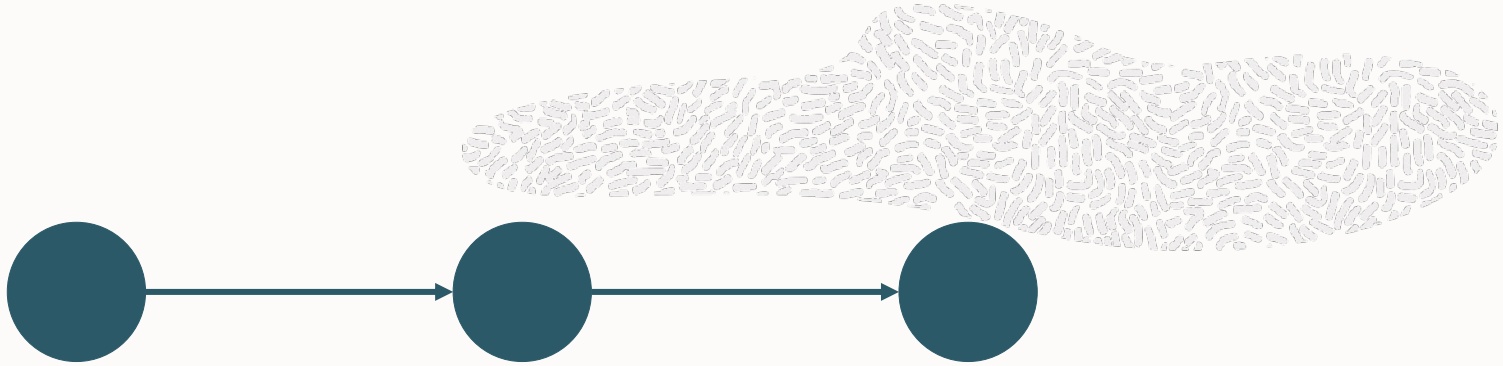

- **Optimistic locking**, where a record is locked only when changes are committed to the database
- **Pessimistic locking**, where a record is locked while it is edited

https://docs.oracle.com/en/database/oracle/oracle-database/19/cncpt/data-concurrency-and-consistency.html#GUID-E8CBA9C5-58E3-460F-A82A-850E0152E95C

Example of a hypothetical system using Optimistic Concurrency Control

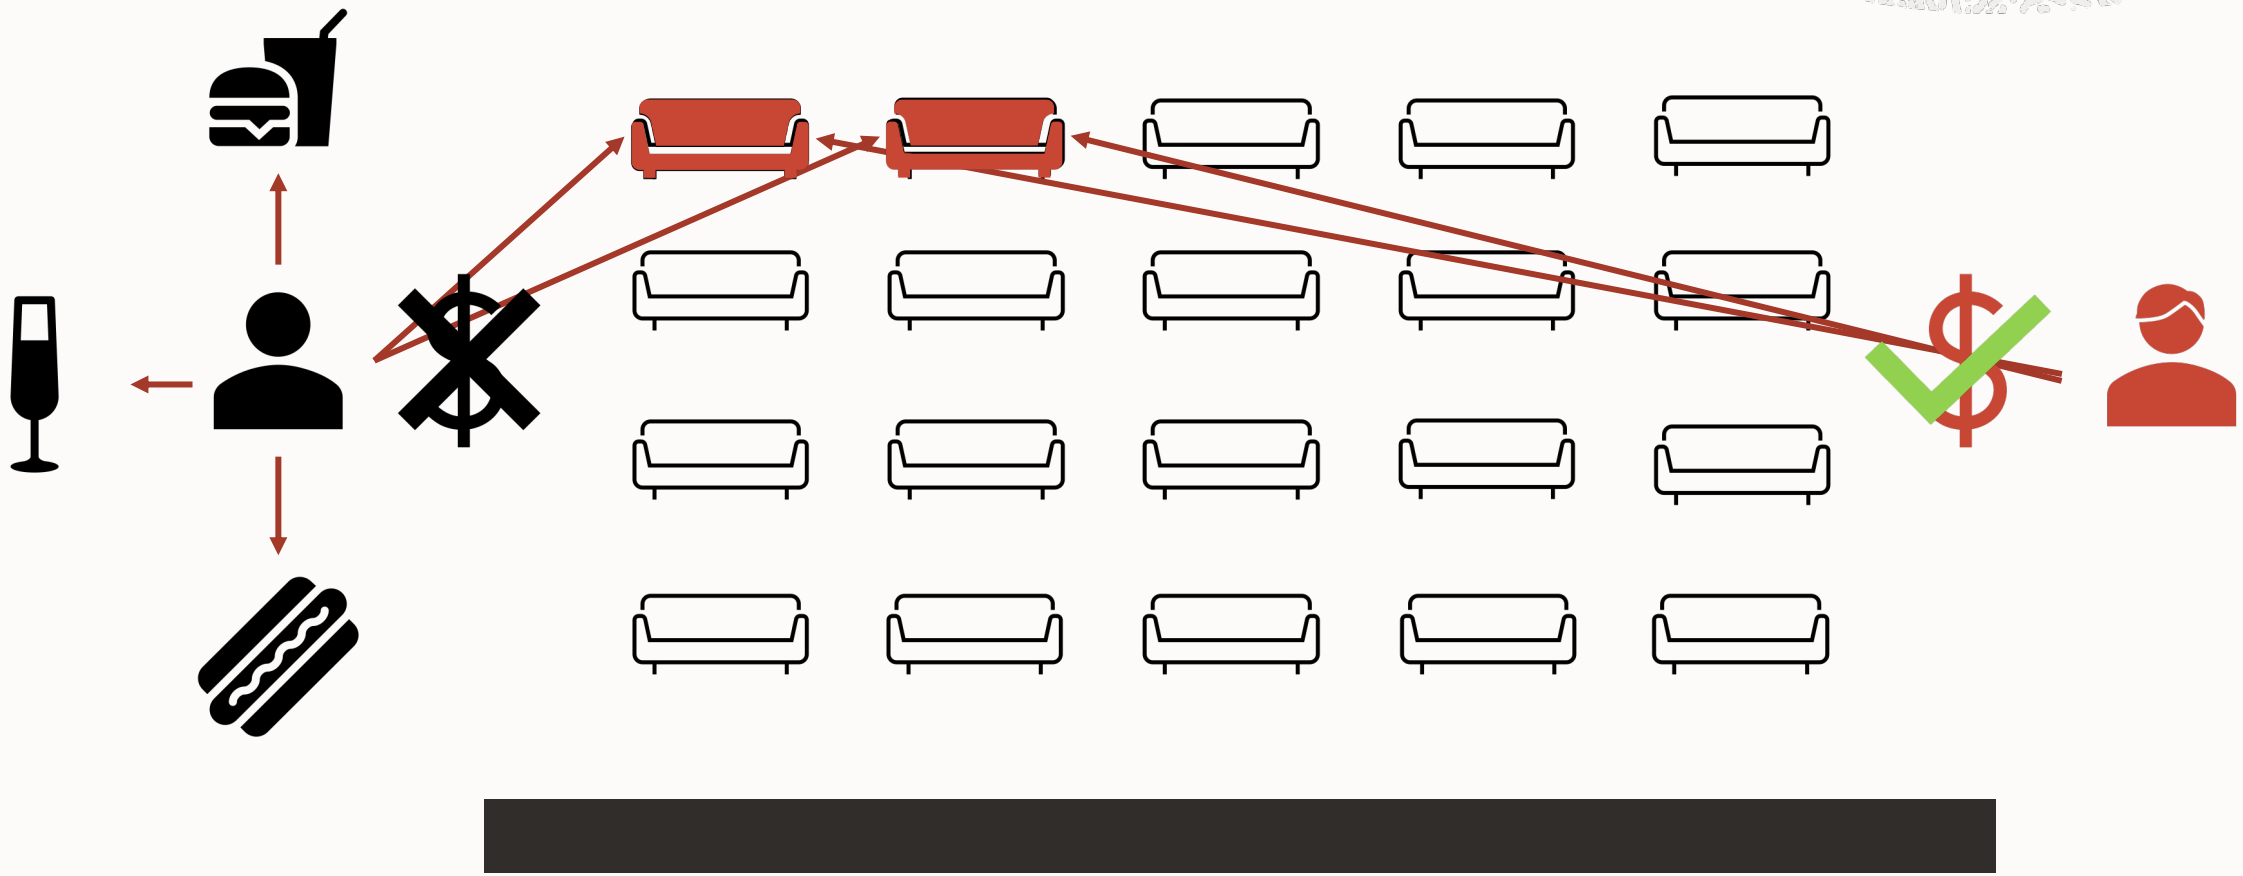

Different Approaches – Different Results

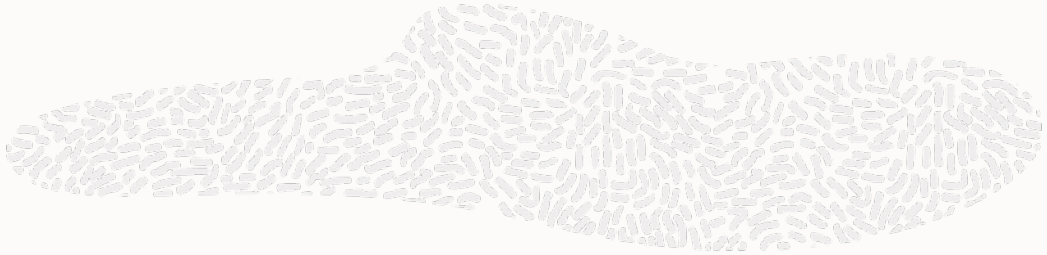

Using optimistic concurrency control by default (like CockroachDB does) means developers may be required to fully redesign it data model and application

The best of both worlds: Oracle automatically assures read consistency per your requirements (statement-level or transaction-level read consistency).

Top 4 Factors Influencing Reliability

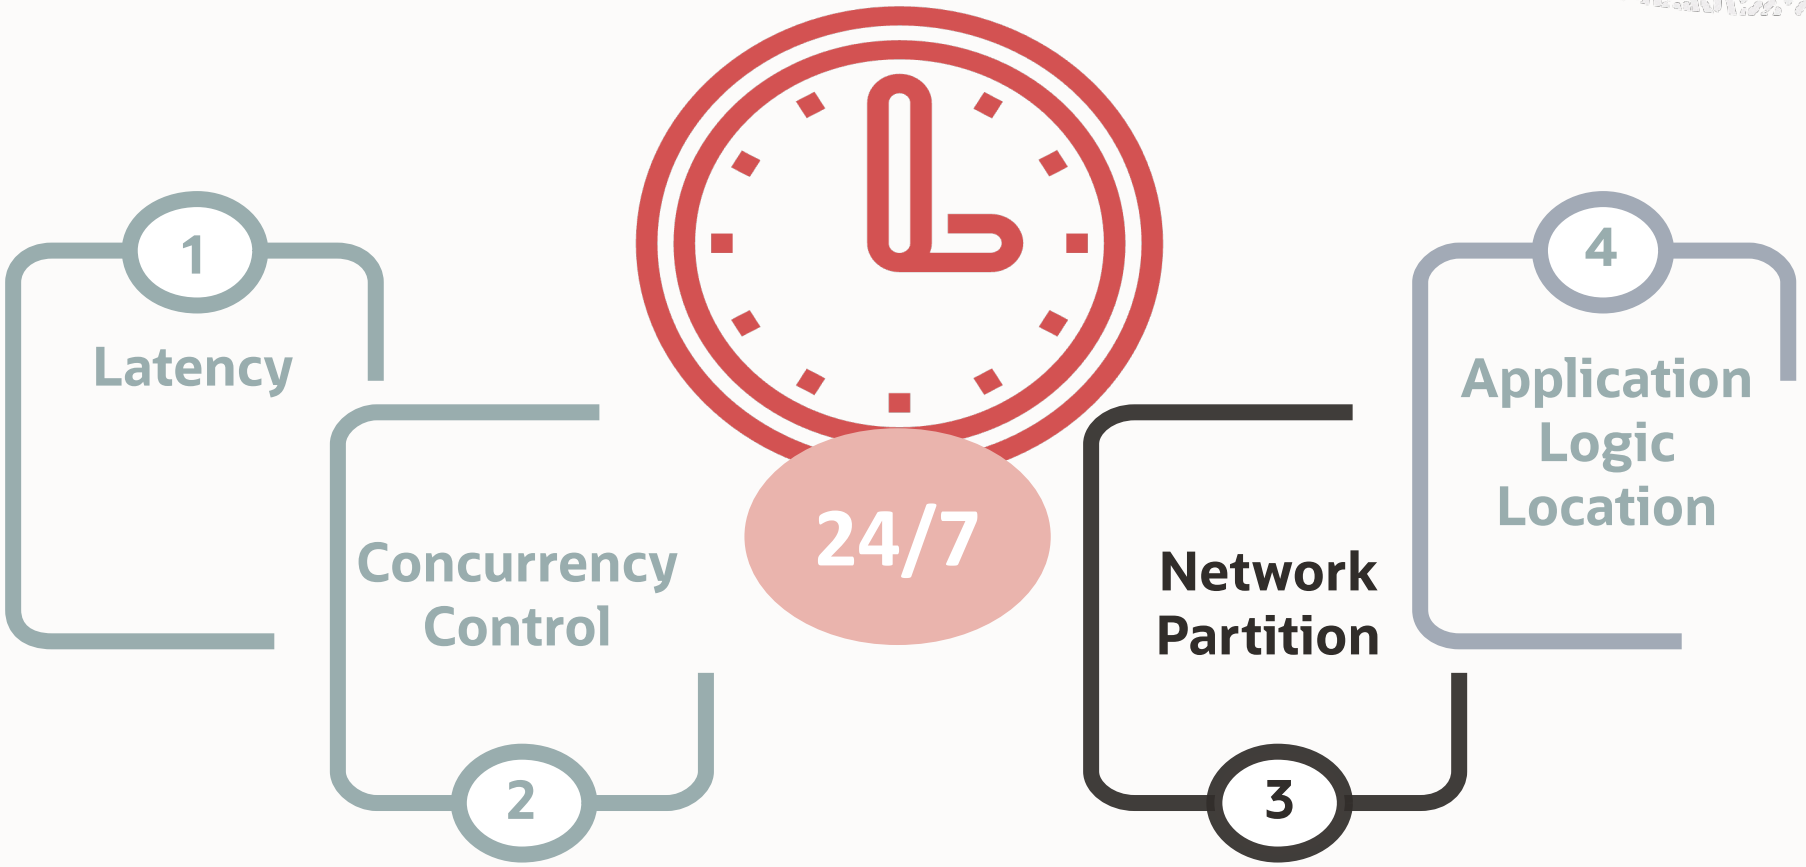

Network Partitioning

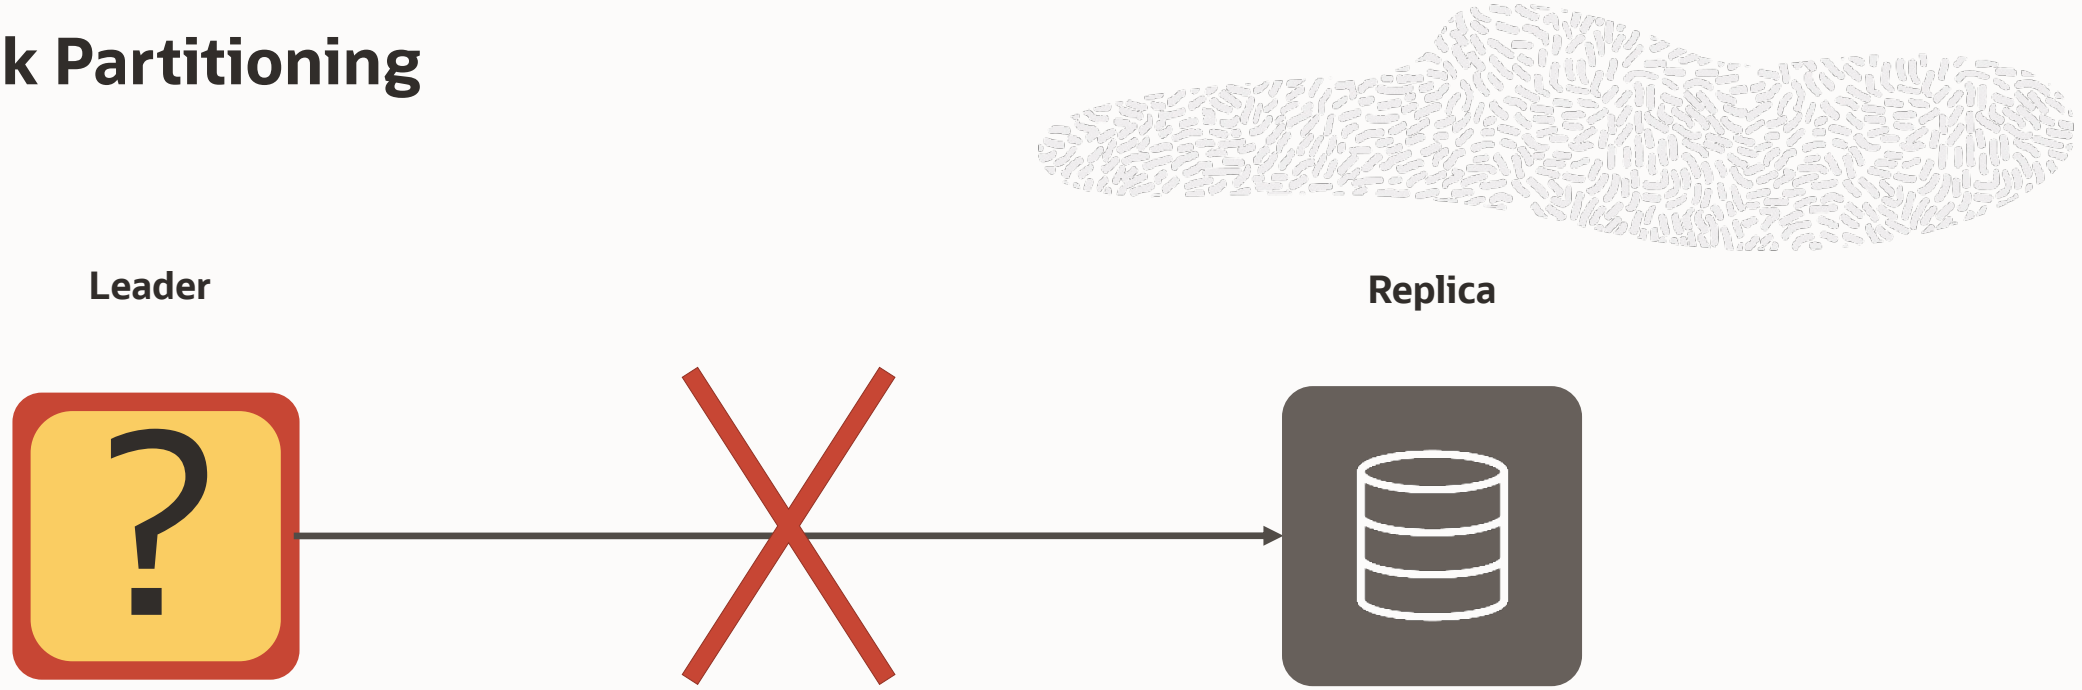

DID IT CRASH?

DID IT STALL?

DID IT KEEP COMMITTING?

WHAT IS THE LEADER DOING?

- If the leader is still committing
  - DATA LOSS and possible split-brain
- If the leader crashed
  - *Maybe* no data loss

Different Approaches – Different Results

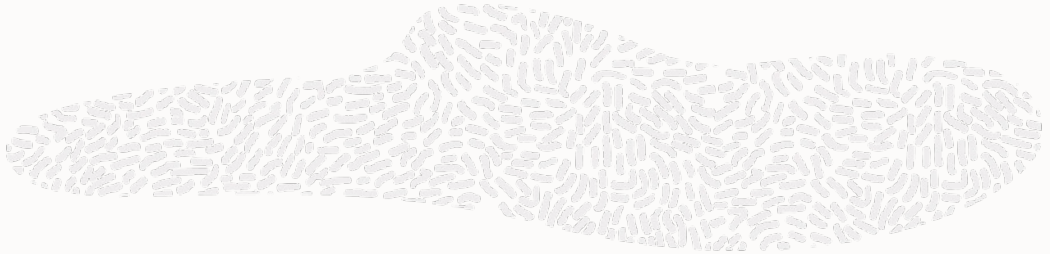

With CockroachDB a network issue and a change of leader could trigger a lost update scenario.

With Oracle Database, lost updates would not occur due to network partitioning.

With Oracle Real Application Clusters (RAC), when used for high availability and scalability

Oracle Data Guard, when used in maximum availability mode or higher

The Hypothetical Journey of a lost update that could happen with Raft replication and optimistic concurrency control

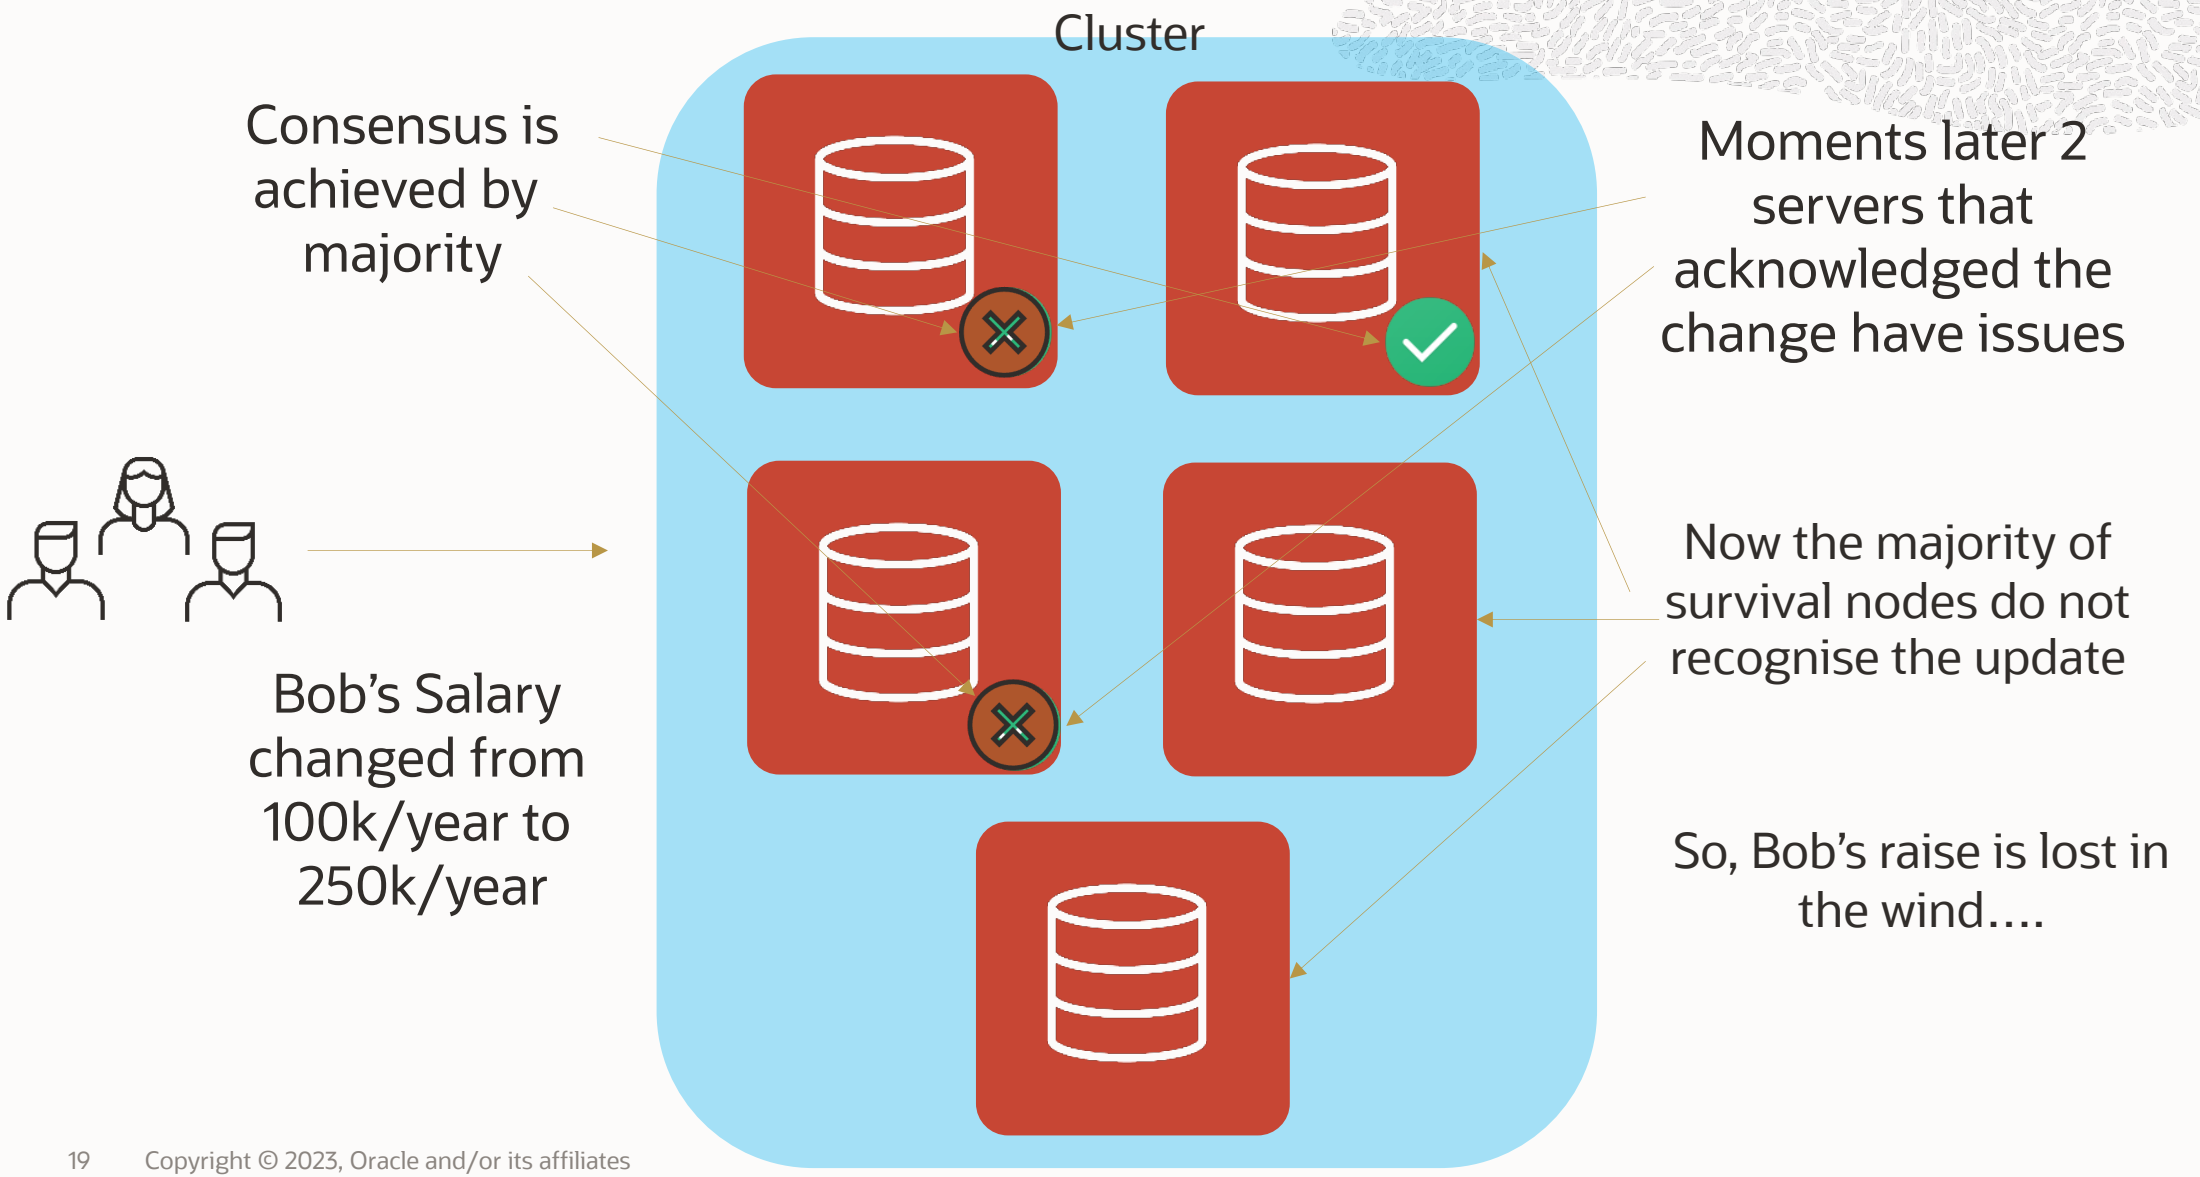

Top 4 Factors Influencing Reliability

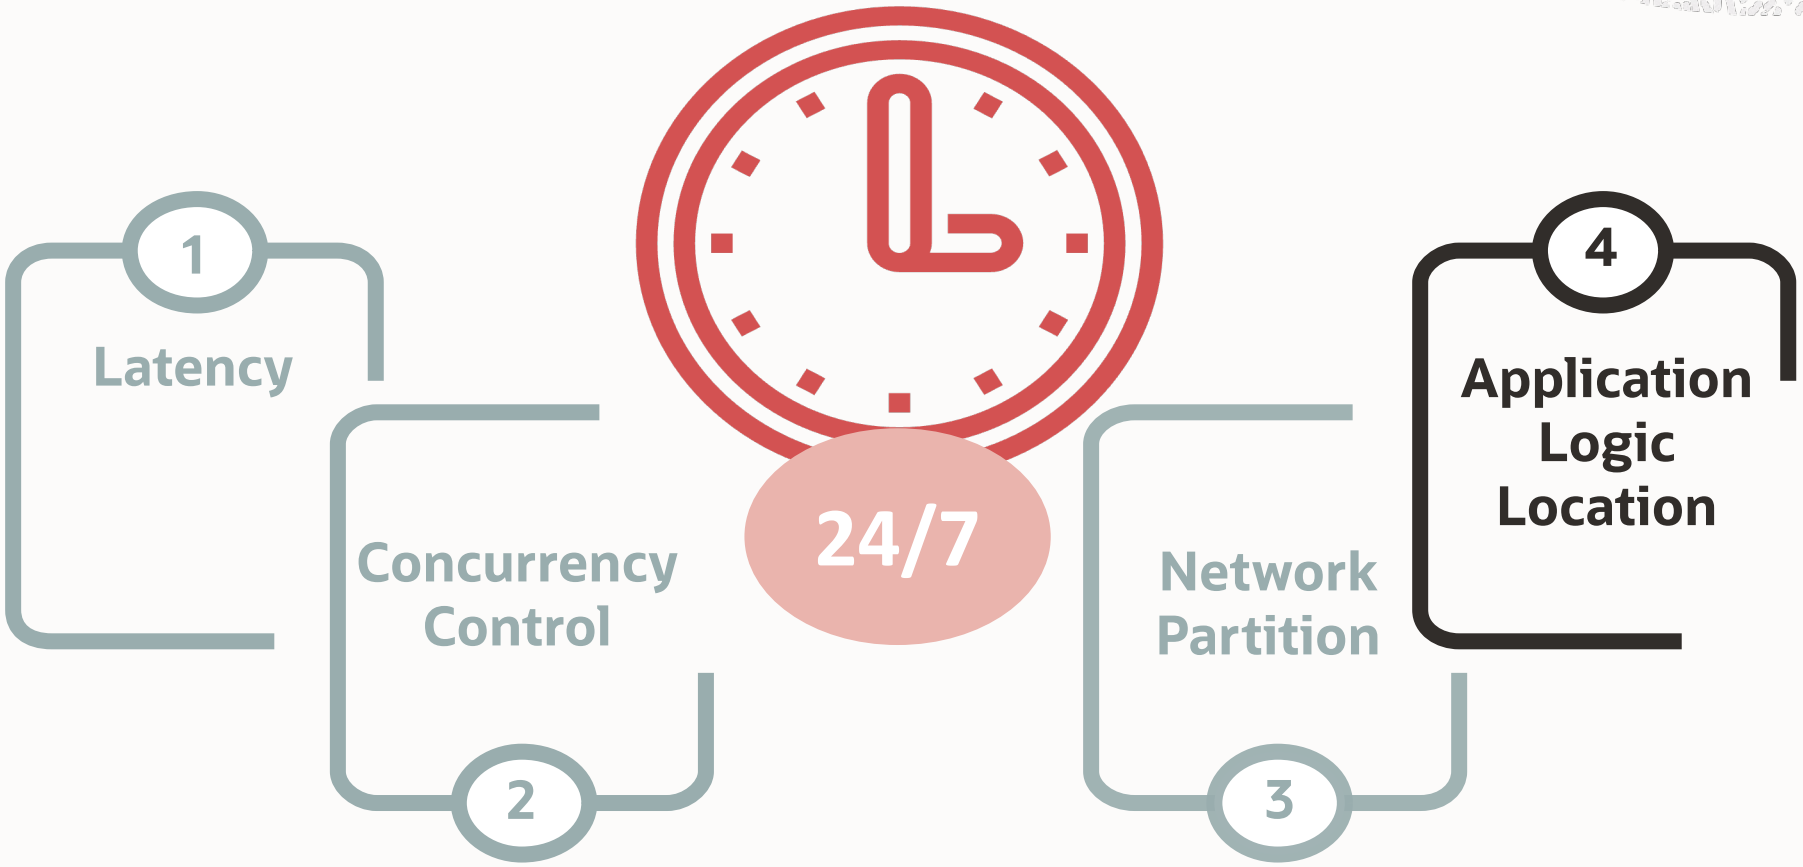

Application Logic is best stored in Oracle Converged Database

A complete database that makes it dramatically easier to develop and run modern apps

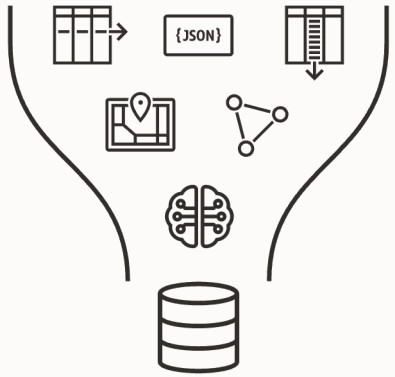

Converged Database

- **No need to fragment** data across databases to support new app requirements
- **Scaling and availability are transparent**, without sacrificing data consistency
- **No need to compromise** on functionality or performance
  - Oracle's data technologies are rated industry-leading in each area

Creating a Fully Complete Database has Taken Decades of Effort by Thousands of Engineers

Oracle Converged Database provides choices

Converged does not mean data must be in one monolith database – **YOU choose**:

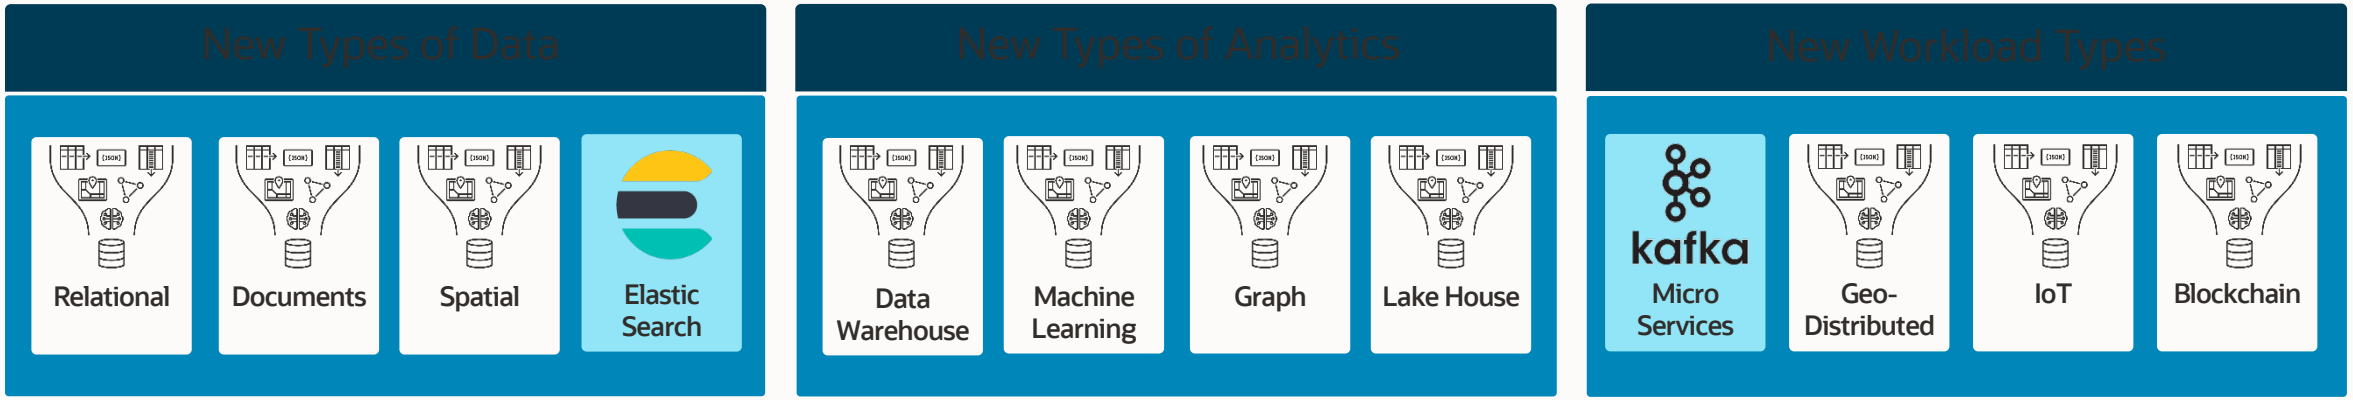

(Support for all modern data types and analytics are included at no additional charge.)

Oracle RAC

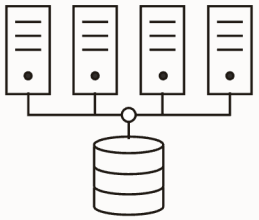

Single Instance

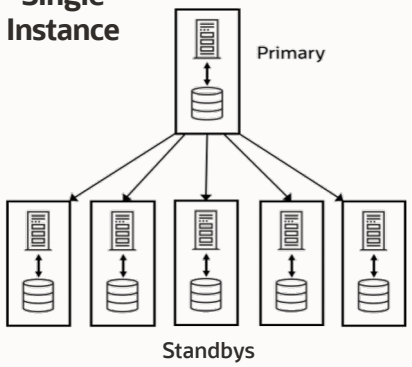

Sharding

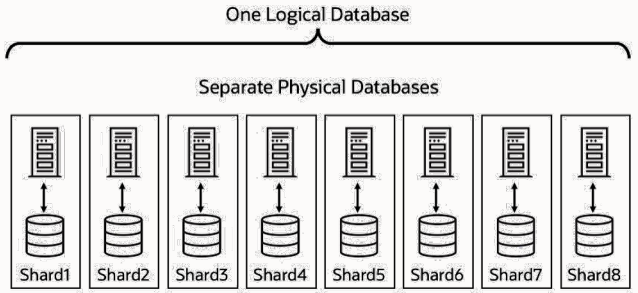

Distributed Databases - Comparison Matrix

	On-Premises Availability	Distribution Layer	Data Modeling	Converged Database	Triggers, Cursors and Stored Procedures within DB	Distributed Capabilities Since	CC (*) Mode (Read)	Popularity Ranking**
CockroachDB	Yes	Raft with Range Sharding (Hash-Sharded Indexes)	Limited Data and Partition Types	No	**No**	2015	Optimistic	**57**
YugabyteDB	Yes	Raft with Hash and Range Sharding	Limited Data and Partition Types	No	Yes	2016	Optimistic	**86**
Google Spanner	No	Raft with Range Sharding	Limited Data Types	No	**No**	2017	Pessimistic	**91**
Oracle Database	Yes	In-Memory Redo (Physical and Logical) Replication, Shard, **and Raft** **NEW IN 23c**	Huge amount of Data and Partition Types	Yes	Yes	1986 with version 5	Pessimistic & Optimistic	**1**

Many more years of innovation than all other competitors together!

(*) Concurrence Control ** Reference: https://db-engines.com/en/ranking

Oracle Database with MAA

Oracle Maximum Availability Architecture (MAA)

Standardized Reference Architectures for Never-Down Deployments

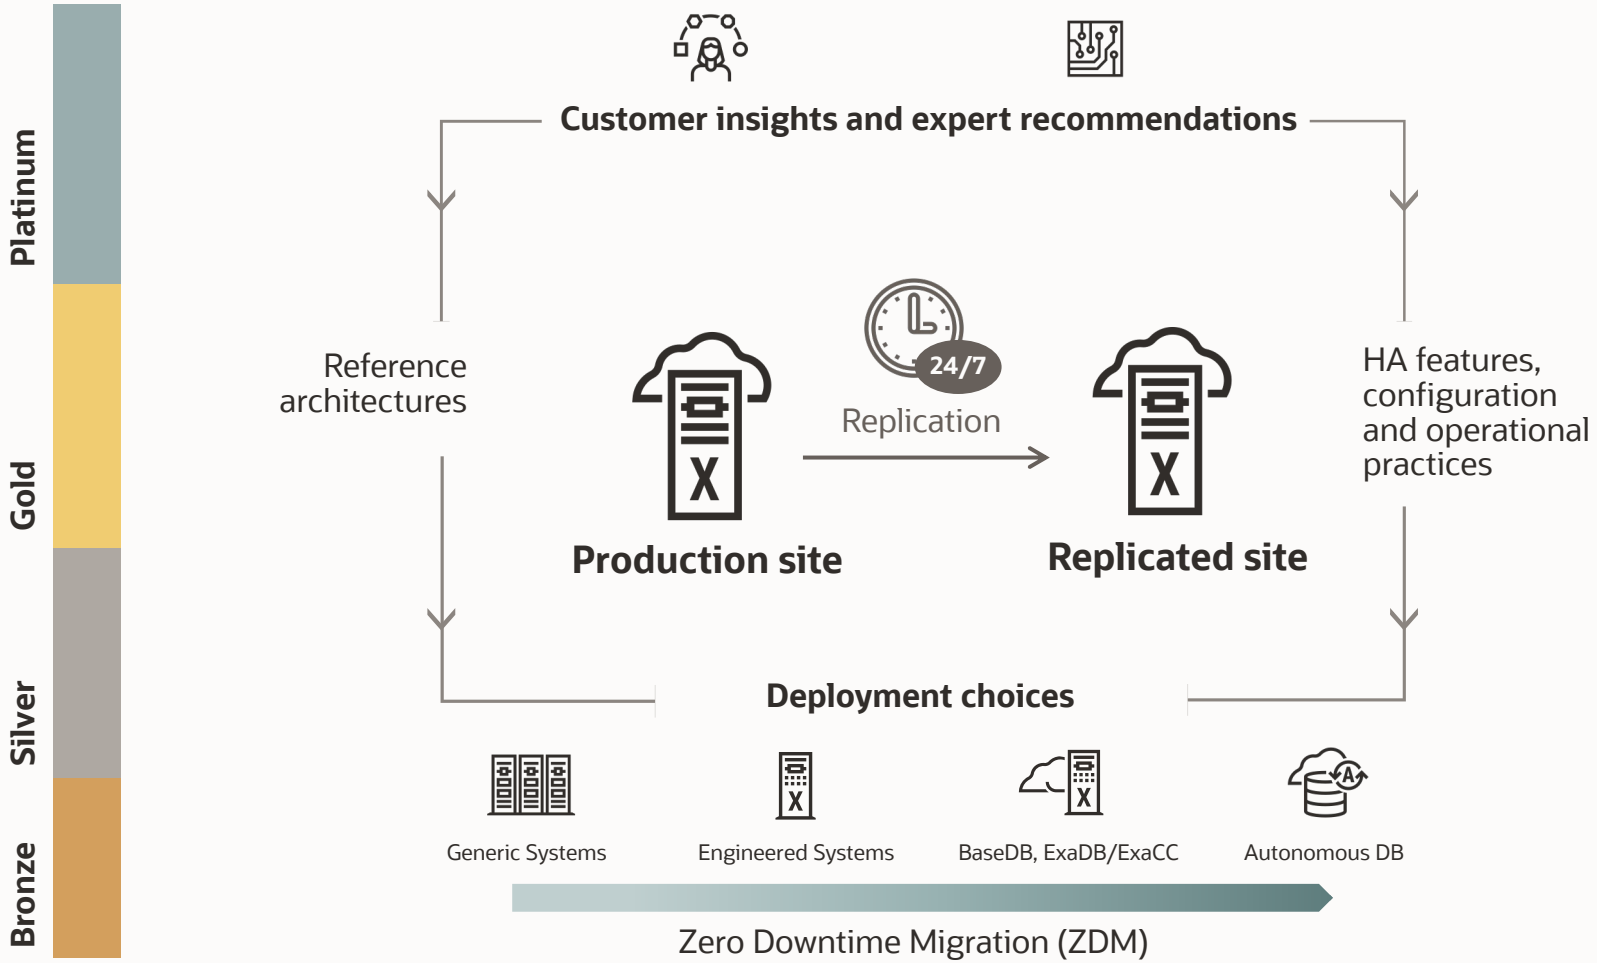

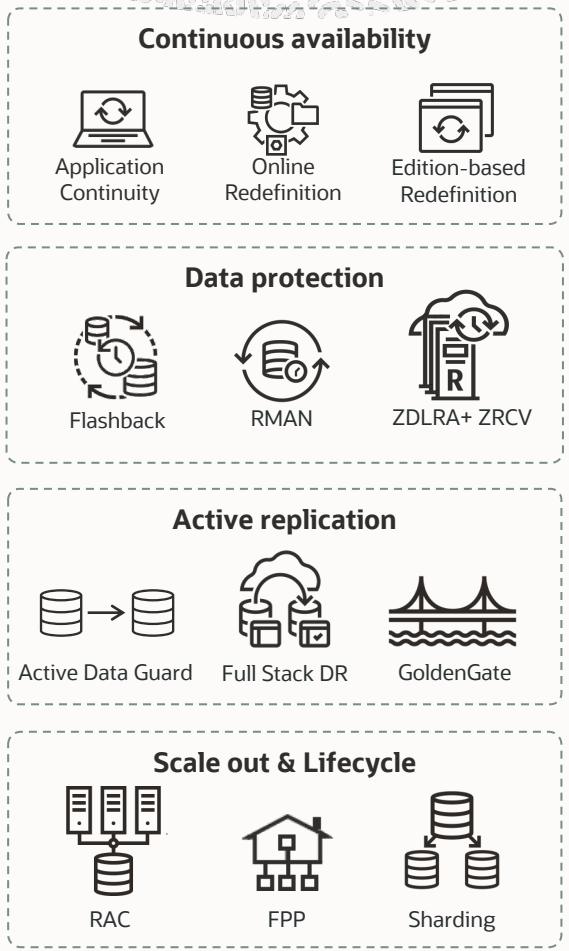

MAA reference architectures

Availability service levels

Bronze	Silver	Gold	Platinum
Dev, test, prod	**Prod/departmental**	**Business critical**	**Mission critical**
Single instance DB Restartable Backup/restore	**Bronze +** Database HA with RAC Application continuity Sharding (optional)	**Silver +** DB replication with Active Data Guard or Data Guard	**Gold +** GoldenGate Edition-based redefinition

Oracle Database with MAA vs CockroachDB at a Glance

Capabilities	Oracle Gold & Platinum MAA	CockroachDB
Full Disaster Recovery (regional protection)	✓	⚠
High Consistency (within a region)	✓	✓
One Node Survival	✓	✗
Time Machine Recovery	✓	✗
Full Backup and Recovery	✓	✓
Full Backup and Recovery (with Zero Data Loss)	✓	✗
Near Zero Downtime Upgrade/Patching	✓	✓
Automatic Conflict Detection and Resolution	✓	✗
Ransomware Protection at Database Level	✓	✗
Automated Failover and Transaction Replay	✓	✗
End-to-End Validation – Data Corruption Prevention	✓	✗
Zero Data Loss at any Distance	✓	✗
Application Versioning	✓	✗
Converged Database	✓	✗
Data sovereignty	✓	✗
Easy to Install/Configure	✗	✓

✓ Fully Supported ✗ Not Supported ⚠ Warning - Performance/Scalability Implications

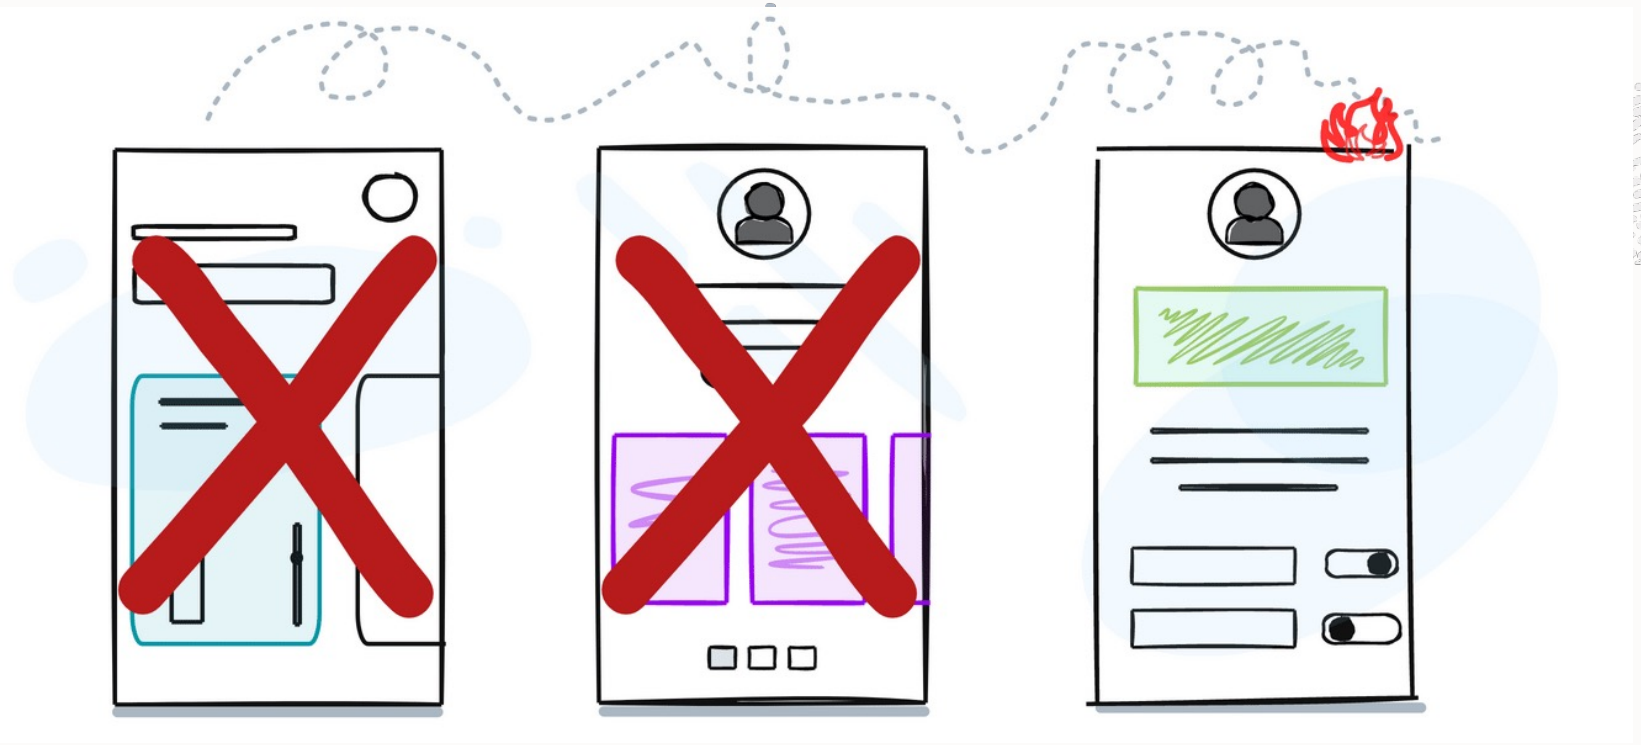

A distributed database is a useful technology for achieving an acceptable level of high availability but is *not* built with disaster recovery in mind.

Disaster Recovery (DR) Scenarios

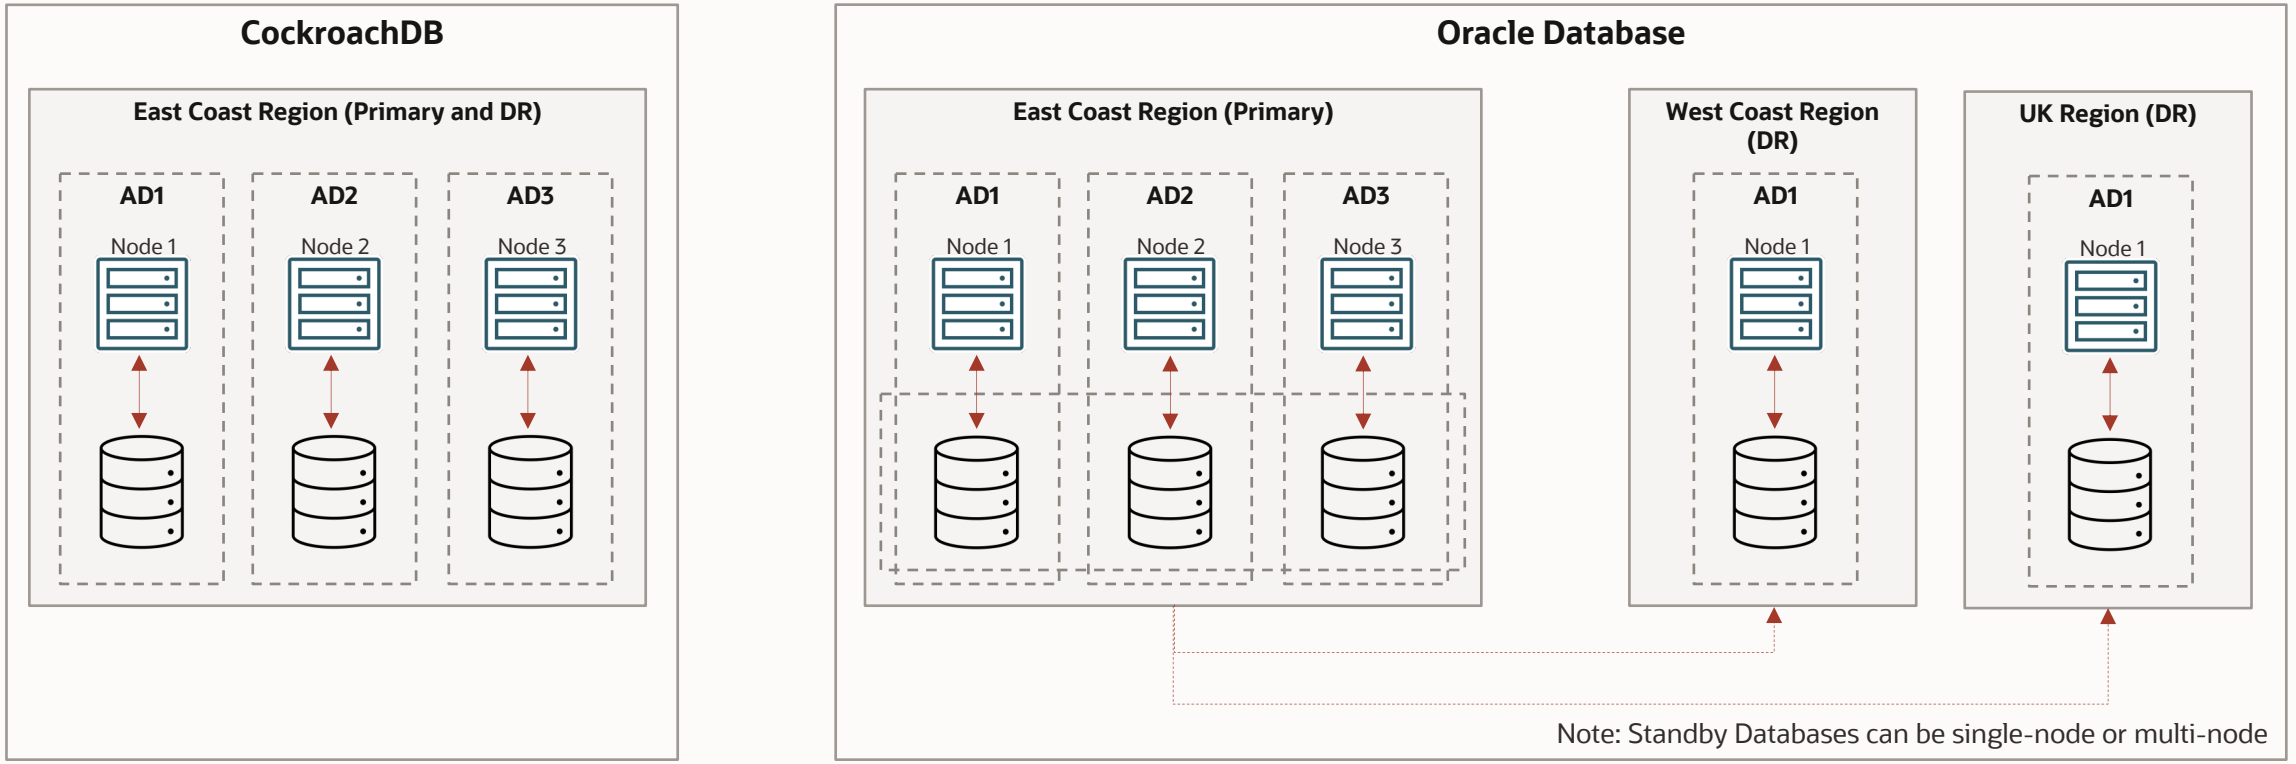

Industry Leading

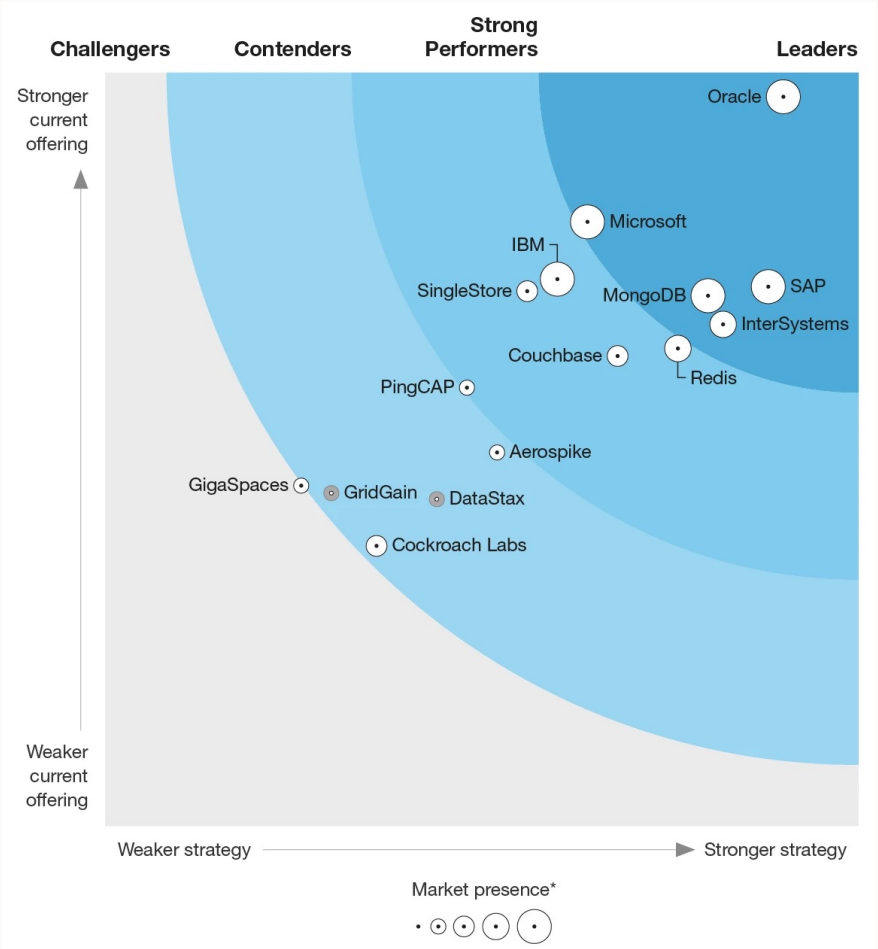

Leader in Translytical Database.

Forrester rates Oracle the strongest leader in the Forrester Wave: Translytical Data Platforms Q4 2022. The rating in this Forrester evaluation validates Oracle Database for its ability to support converged OLTP and Analytical (Translytical) workloads using Database In-Memory, Exadata, extensive multi-model capabilities, and support for relational and non-relational, structured and unstructured.

Oracle Database – Deployment Choices

Deploy Oracle Anywhere – Extreme Portability

Same database, same skills
Fast and efficient deployments

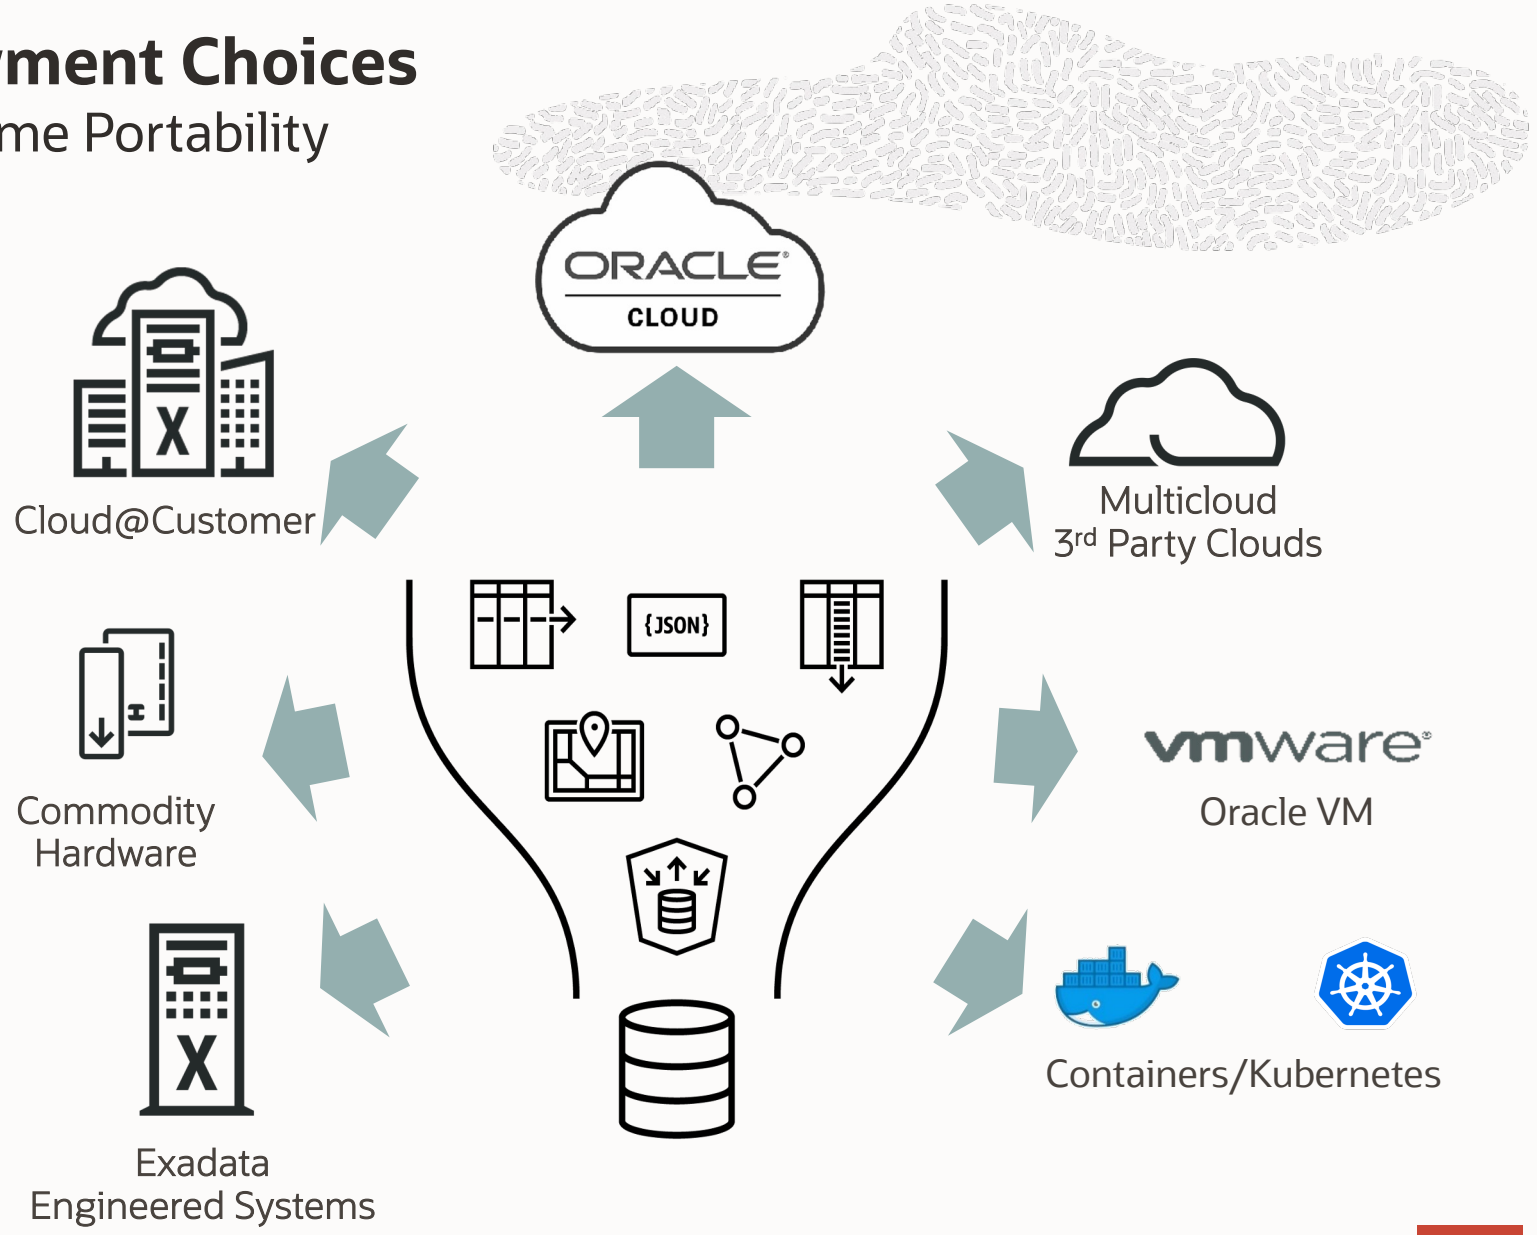

Conclusion

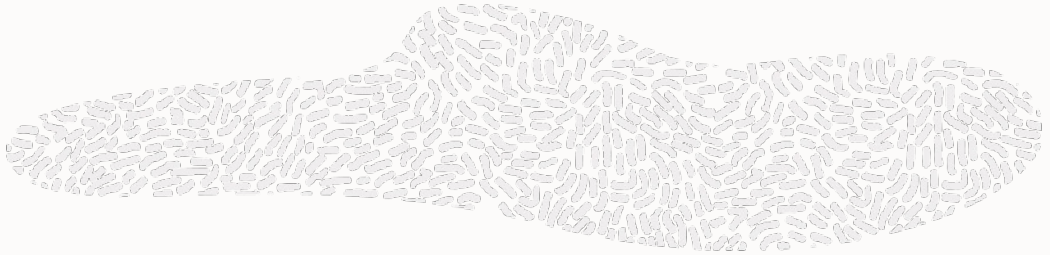

When choosing a technology for distributed database deployments it is important to balance the architecture requirements with the business needs and service level objectives

Customer Example: Global Payment Processor

Combined architectures to serve worldwide customers more efficiently

Architecture Before Sharding

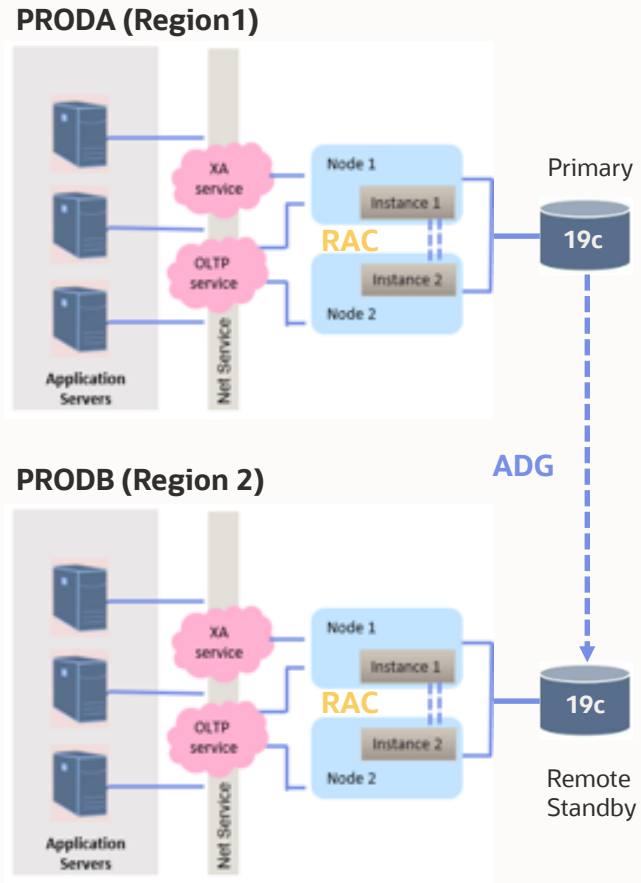

The same architecture with Sharding

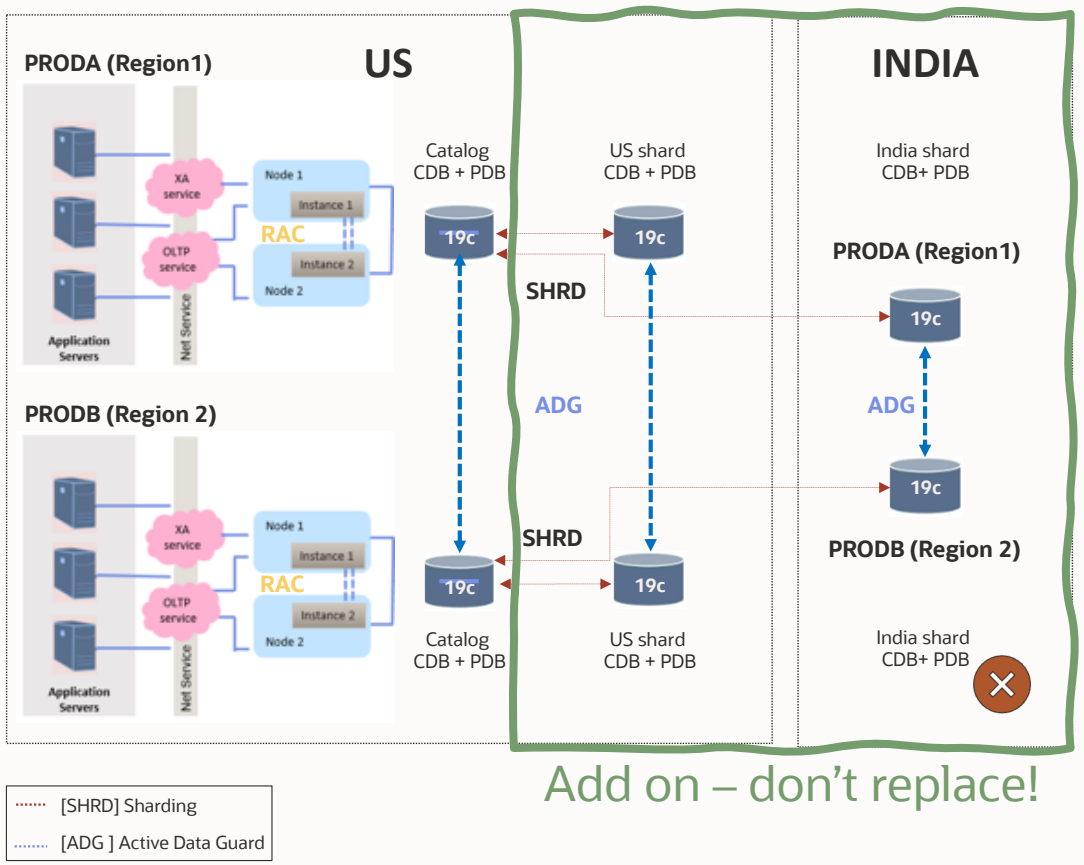

Add on – don't replace!

Summary and Conclusion

Oracle Database is all-inclusive, converged by design, and protected by MAA

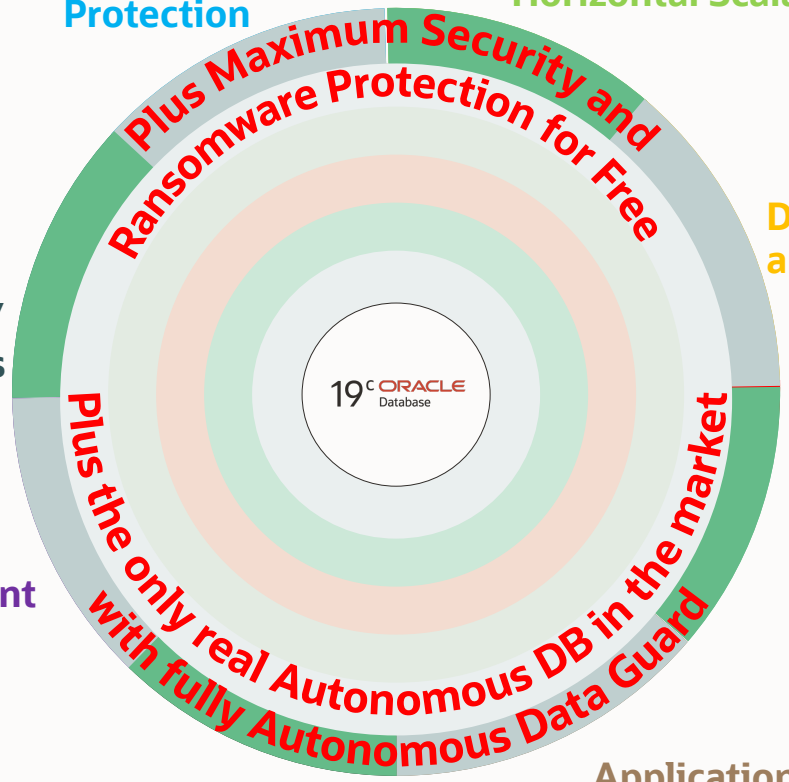

Data Corruption Protection

Automatic Block Media Recovery
RMAN – Block Media Recovery
Zero Data Loss Recovery Appliance (ZDLRA)

High Availability & Horizontal Scalability

Oracle Real Application Clusters
Oracle Sharding
Oracle GoldenGate

Zero Downtime Patches/ Upgrades/ Migrations

Oracle Real Application Clusters
Oracle GoldenGate
Edition Base Redefinition
Rolling Upgrade
Oracle Fleet Patching and Provisioning (FPP)
Zero Downtime Migration (ZDM)

Disaster and Recovery

Oracle Data Guard
Oracle Active Data Guard
Oracle GoldenGate
Oracle Flashback
Zero Data Loss Recovery Appliance (ZDLRA)
Oracle Database Zero Data Loss Autonomous Recovery Service

Storage Management

Oracle Automatic Storage Management

Multi-Master Capabilities

Oracle Real Application Clusters
Oracle GoldenGate
Oracle Active Data Guard (DML Redirection)
Oracle Sharding with Raft

NEW IN 23c

Automation

Full Stack DR
Fast-Start Failover (FSFO)
Oracle Fleet Patching and Provisioning (FPP)
Global Data Services

Application Continuity

Oracle Real Application Clusters
Global Data Services
Oracle Application Continuity
Transparent Application Continuity
Edition Base Redefinition
Oracle GoldenGate

Top 5 reasons you should always take into consideration when thinking to replace Oracle with a distributed database

1. Changes to the application and data model layer are required to handle requirements that were naturally managed by the database.
2. You may need multiple technologies to fulfill all business/application requirements (adding complexity, risk, and cost).
3. A “low-cost solution” comes with a cost, including a lack of crucial functionalities and being charged for Enterprise Support and backup, not free.
4. The “try again later, scenario” - A lost transaction equals lost revenue and, maybe, brand impact!
5. Finally, possible data loss due to inefficient Disaster Recovery capabilities. Ensure all your requirements can be fulfilled, especially when discussing RTO and RPO requirements.

Any Questions?
Thank you!


@fcomunoz


http://www.linkedin.com/in/franciscomunozalvarez


www.oraclemaa.com