


# 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.



ORACLE

# Exadata Exascale

World's Only Intelligent Data Architecture for Cloud

*Extreme performance for AI, analytics, and mission-critical workloads at any scale*

Ashish Ray

Vice President of Product Management

Oracle Database Mission Critical Technologies

July, 2024



[Ashish.Ray@oracle.com](mailto:Ashish.Ray@oracle.com)

[linkedin.com/in/ashishray](https://www.linkedin.com/in/ashishray)



# Today, thousands of global customers run their business on Exadata 79% of Fortune Global 100 Run Exadata | 58% Run Exadata Cloud



## Superior Architecture for ALL Workloads

- Petabyte warehouses
- Super critical systems
  - Financial Trading
  - Process manufacturing
  - E-commerce
- Packaged applications
  - SAP, Oracle, Siebel, PSFT, ...
- Database consolidation





## Exadata Exascale

is a new intelligent data architecture  
that delivers  
the **best of Exadata** and  
the **best of Cloud**

# Exadata Intelligence highlights

## Intelligent OLTP



**Unique Intelligent Communication** uses hardware-based RDMA to deliver the fastest OLTP IO and fastest cross-node coordination

**Unique Intelligent Clones** deliver instant database cloning with full Exadata capabilities

## Intelligent Analytics



**Unique Data Intelligence** automatically offloads data-intensive SQL to storage cloud

**Unique Intelligent Columnarization** automatically converts data to ultra-fast in-memory columnar format

## Artificial Intelligence



**Unique AI Vector Acceleration** offloads data-intensive vector search IOs to storage cloud

**Unique AI Vector Computation** offloads compute-intensive vector comparison and top-K selection to storage cloud

# All cloud benefits



## Pay-Per-Use

Pay for only the compute and storage you use



## Multitenant Resource Pool

Provision any resources you need when you need them



## Hyper-Elastic

Great for small businesses and departments, in addition to the largest enterprises

# Exadata Exascale

## All Exadata Intelligence

- ✓ Intelligent OLTP
- ✓ Intelligent Analytics
- ✓ Artificial Intelligence



Exadata Exascale

## All Cloud Benefits

- ✓ Pay-per-use
- ✓ Multitenant Resource Pool
- ✓ Hyper-elastic





# Exadata Exascale Cloud Architecture





Previously each cloud tenant had **dedicated Exadata compute and storage servers**

ASM (Automatic Storage Management) was used to distribute storage across databases



# Now a **common pool** of Exascale compute and storage supports thousands of tenants and millions of databases

Resources are completely elastic and pay-per-use

Tenants specify only the cores and storage they need

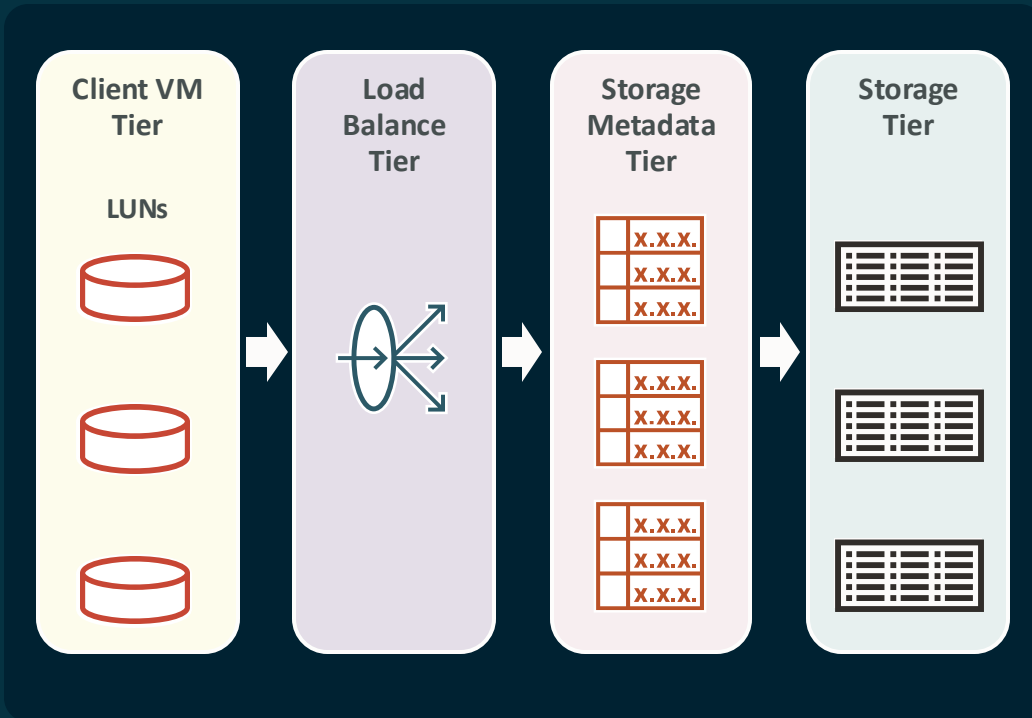
- Start with just a few cores

The Exascale control plane spreads every database across dozens of pooled storage servers



**Shared pool of Exascale servers unleashes thousands of CPUs for any database query**

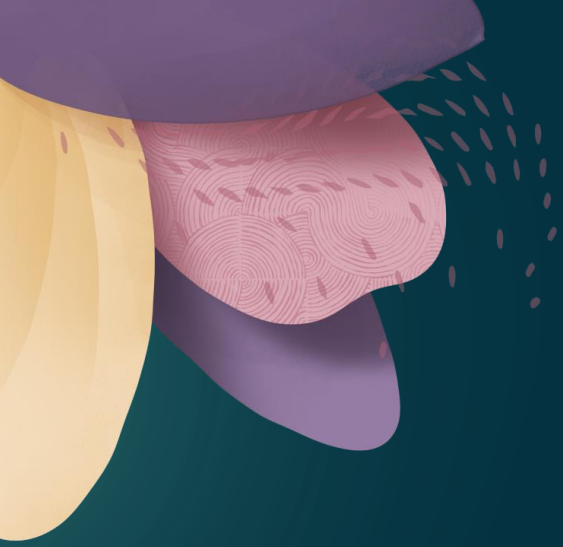
# Mainstream cloud storage is elastic, but its multi-tier architecture limits performance



- Storage is exposed as LUNs within a VM
  - Standard Block IO protocols are used
- Load Balancer tier routes LUN IOs to Metadata Servers
- Storage Metadata tier manages LUN distribution and mirroring
- Storage tier holds data
- Often even more tiers...

**Every tier adds latency and bottlenecks**






Exadata Exascale is a new **loosely coupled** architecture that brings the power of Remote Direct Memory Access (RDMA) and SQL offload to cloud data

# Exascale **uniquely** removes the need for intermediate tiers

Oracle database **23ai** sends intelligent data requests directly to pooled Exascale storage servers

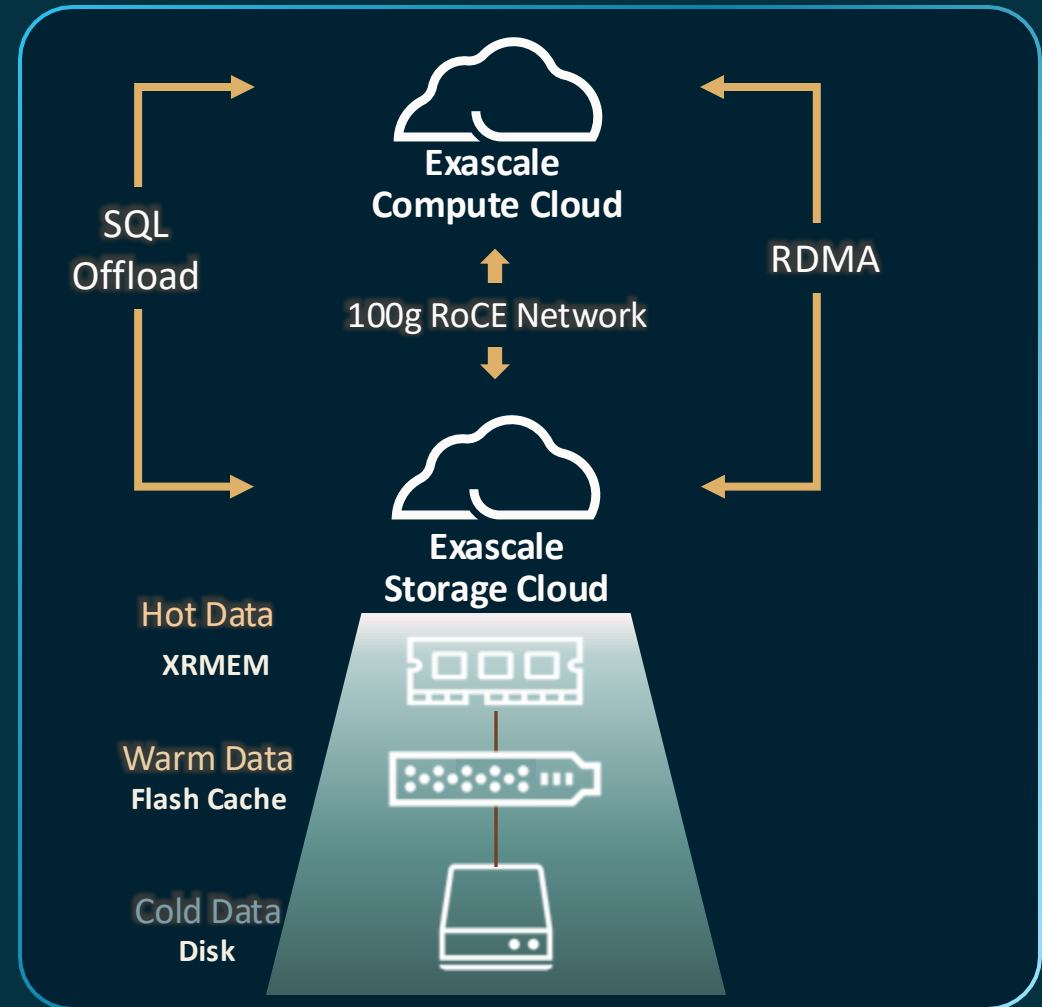
- Direct IO architecture with **no intermediate nodes** is **much faster** and is **required for RDMA**
- RDMA is implemented in hardware enabling **microseconds latency** and **millions of IOs** per second of throughput



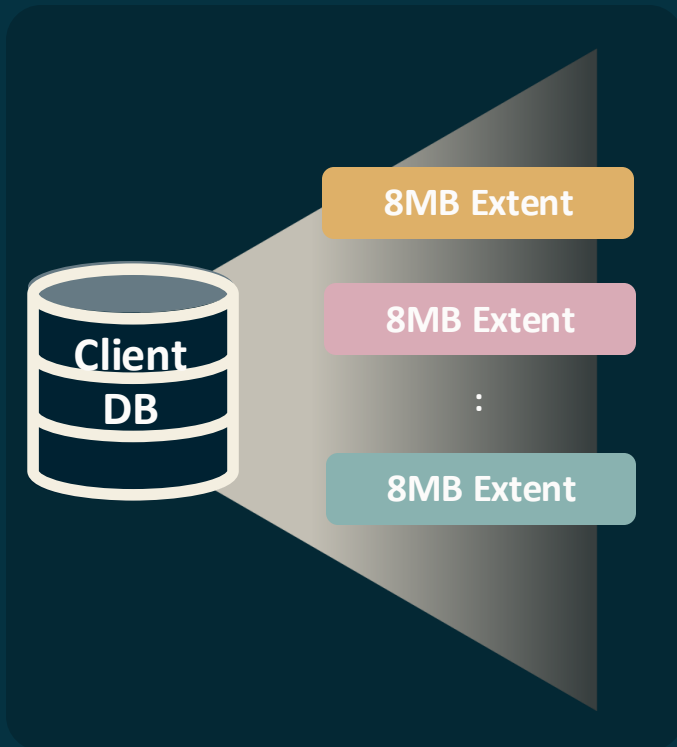
# Exascale **uniquely** optimizes storage capacity in addition to performance

Exascale Storage Cloud intelligently moves hot data from disk to memory or flash

Delivers the performance of DRAM, the IOPs of flash, and the capacity of disks



# Exascale manages data as a large pool of Database Extents



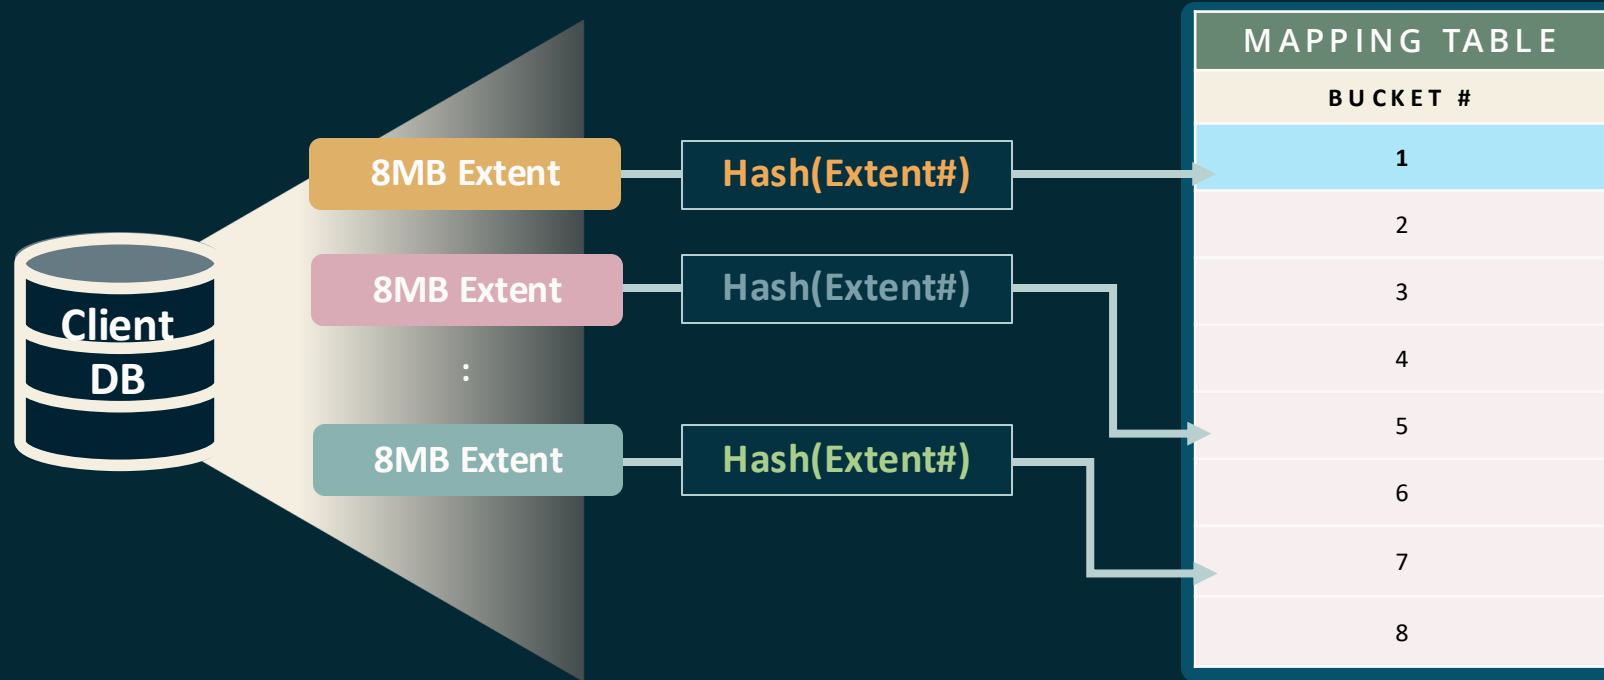
Exascale stores databases as a set of 8 MB Extents

8 MB Extents are large enough to achieve good sequential performance when scanning contiguous data

8 MB Extents are small enough to allow a Database to be distributed across many disks in the storage cloud to distribute IO load

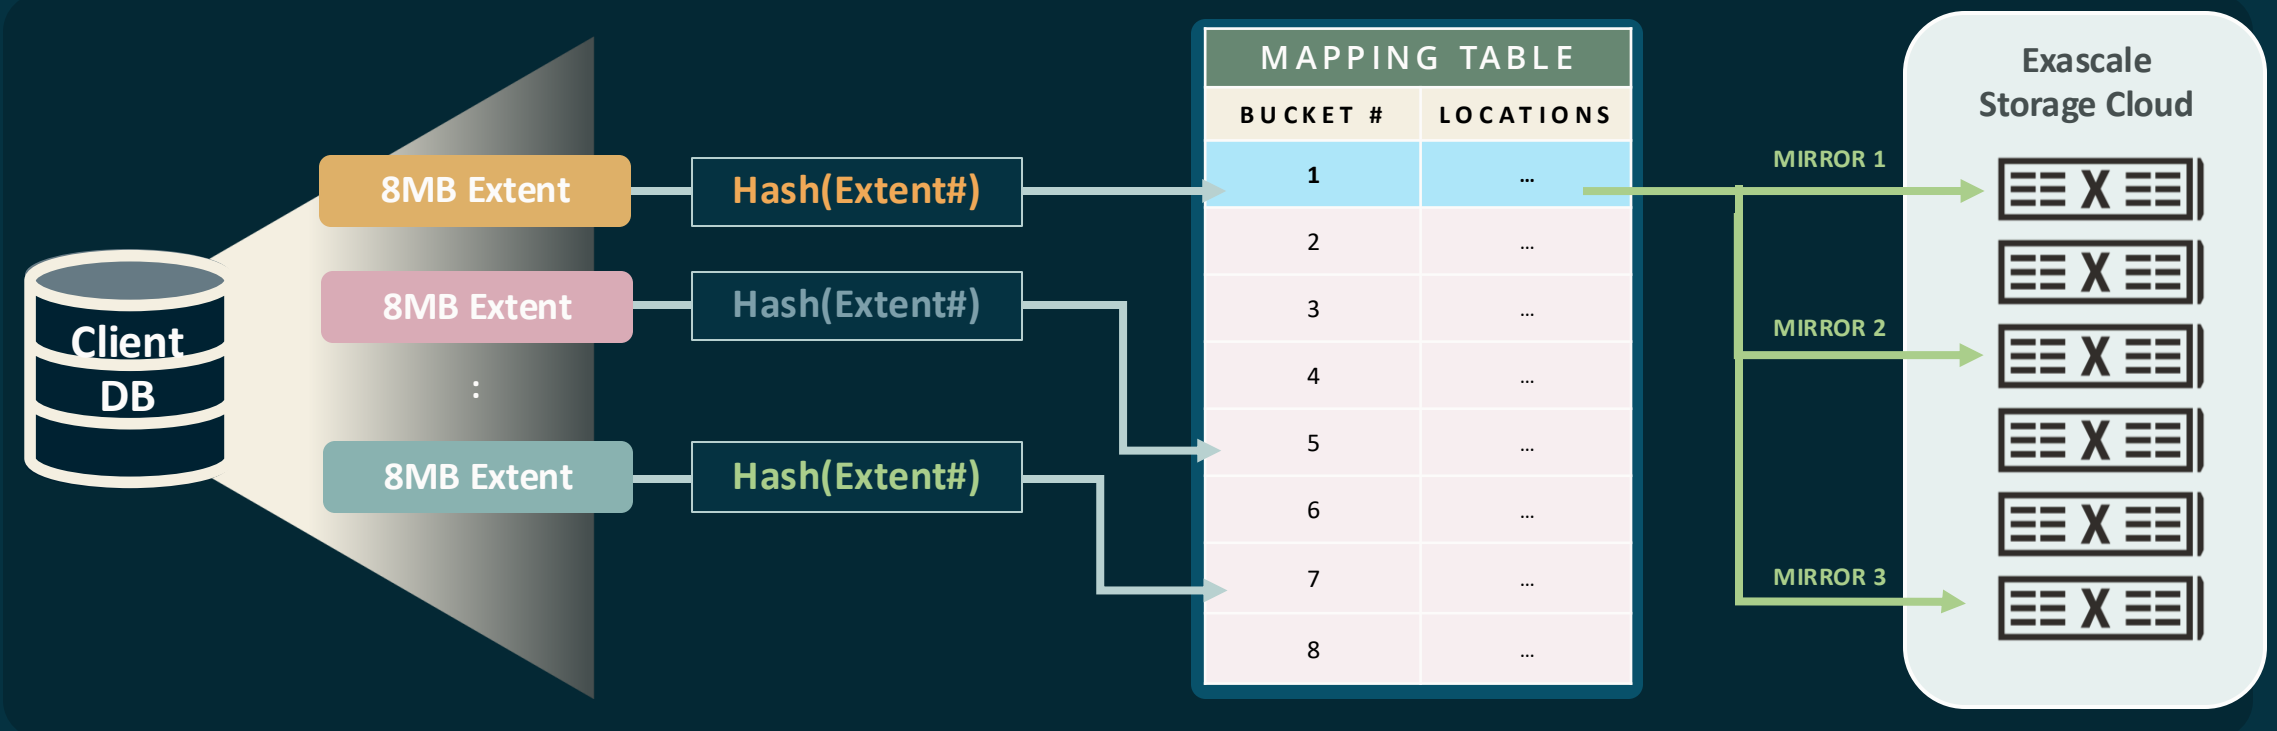


Each extent is assigned to a **storage bucket** using a **hash function**



Multiple Extents can hash to a Storage Bucket

A **Mapping Table** tracks the drives each Storage Bucket is currently stored on



All **extents** that hash to a **storage bucket** are **stored redundantly** on 3 drives that are on 3 separate storage servers to protect from storage failures and outages

# Hyperscale

MAPPING TABLE	
BUCKET #	LOCATIONS
1	...
2	...
3	...
4	...
5	...
6	...
7	...
8	...

The number of Storage Buckets is **fixed** (e.g. 100K)

100K Storage Buckets is large enough to **spread data** across thousands of storage servers

100K is small enough to easily keep the Mapping Table **cached** in the client database server's DRAM

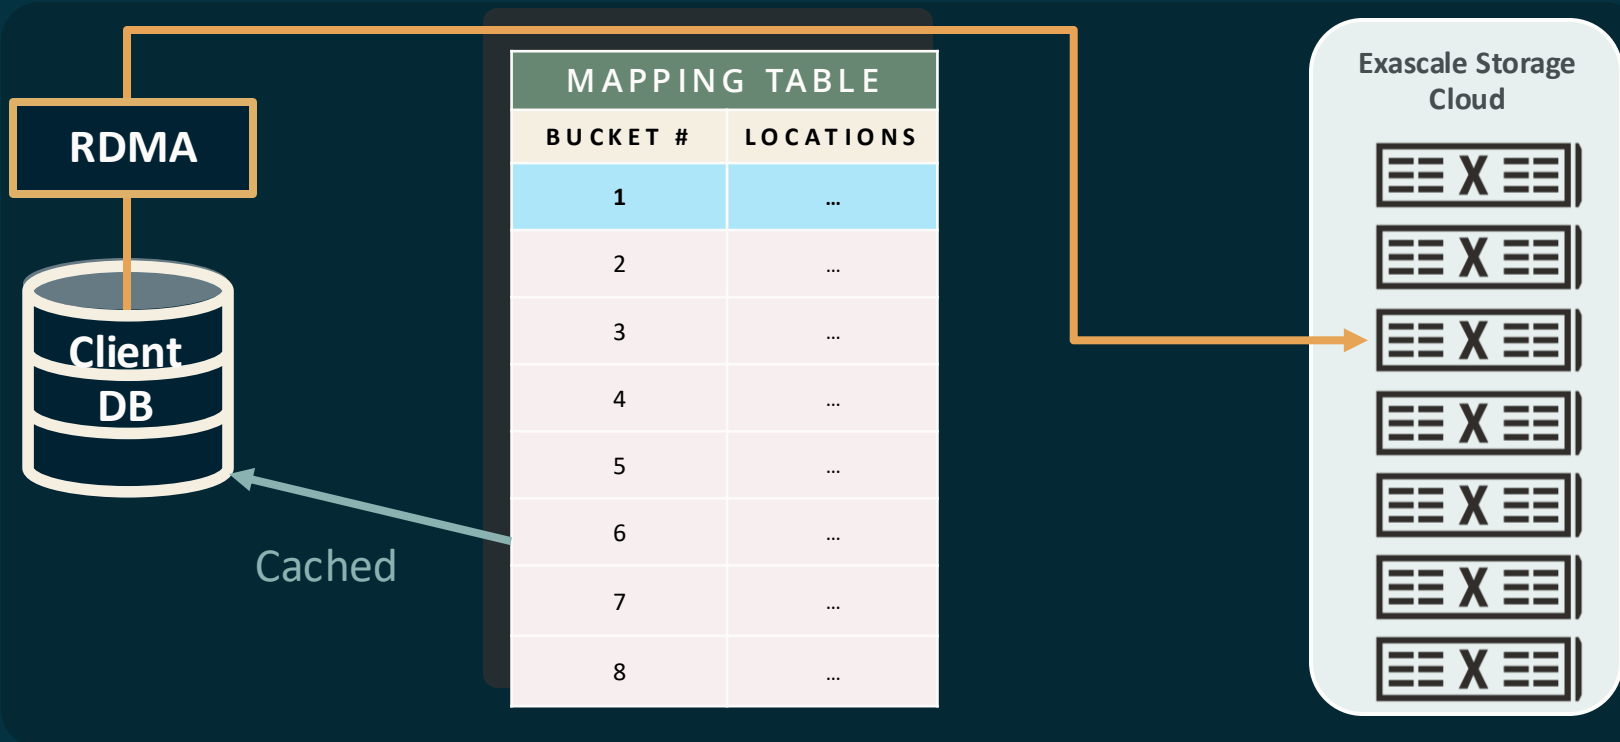
If a database grows, or more databases or storage servers are added, then more extents hash to existing storage buckets

- The mapping table does not grow

**Scales data without scaling metadata**



# Caching the Mapping Table enables RDMA from databases directly to storage servers



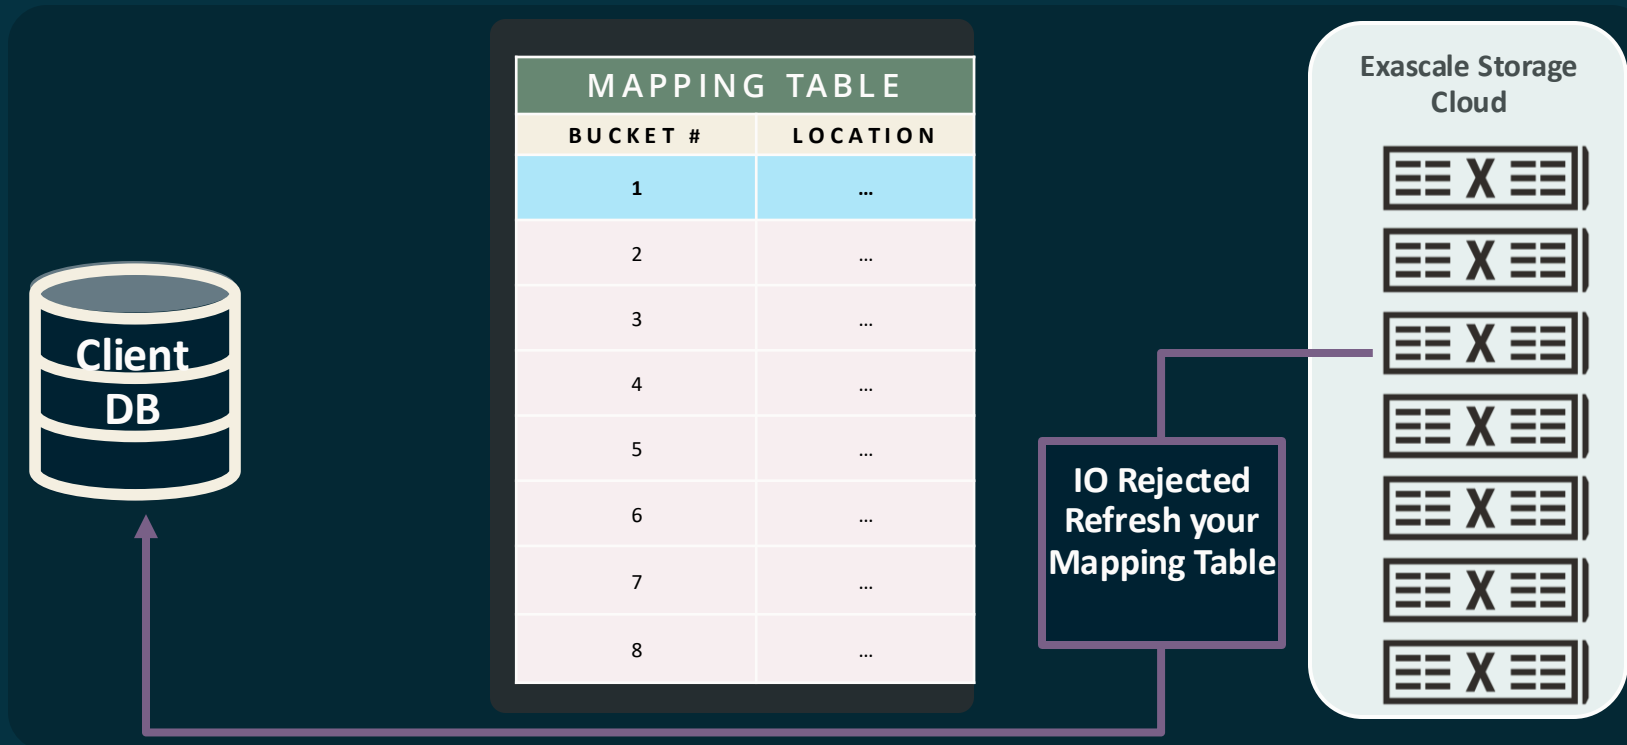
Each client database caches the Mapping Table

Allows database to send IOs **directly** to storage servers

Direct IO architecture with **no intermediary nodes** is **faster** and is **required for RDMA**



# Stale cached mapping tables are tolerated and refreshed on demand



No need for distributed locking when data is re-distributed across servers

- If a client database sends IO to the wrong storage server due to a stale cache, the storage server **rejects the IO**
- And then instructs the database to **refresh** the stale mapping table

Loosely coupled architecture enables Hyper-Scale

# Exadata Exascale Intelligent Block Storage



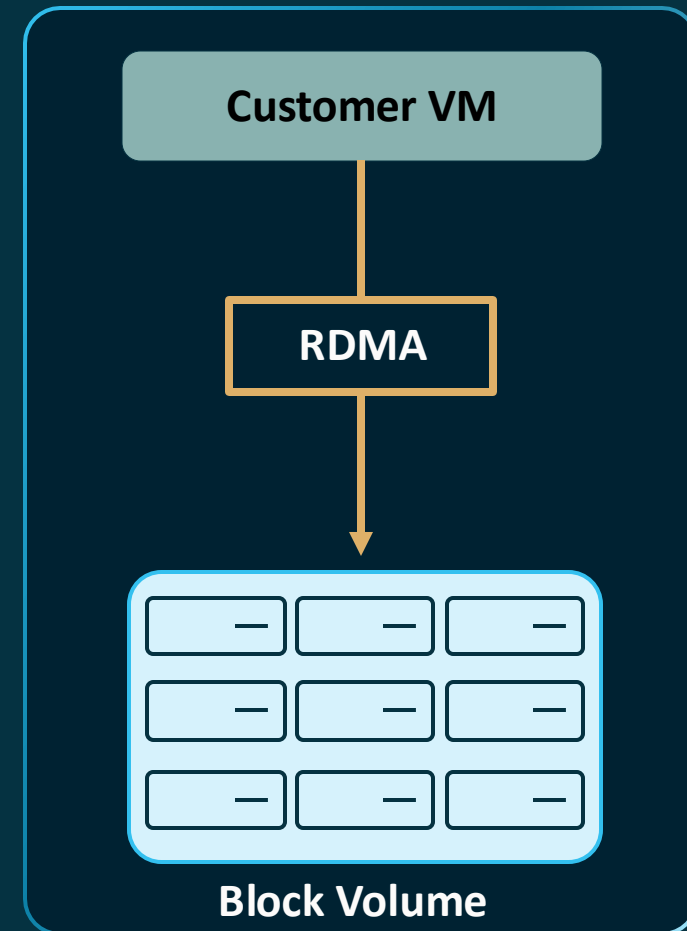
# Unique **RDMA-enabled** cloud block volumes deliver **extreme low latency** and **extreme high throughput** for conventional block IO

Provides extreme performance IO for ACFS and Linux filesystems

Enables Virtual Machine images to reside on intelligent shared storage

- Provides the foundation for live migration for RDMA capable VMs

Provides high performance space-efficient volume snapshots and clones for data on file systems in addition to data in databases



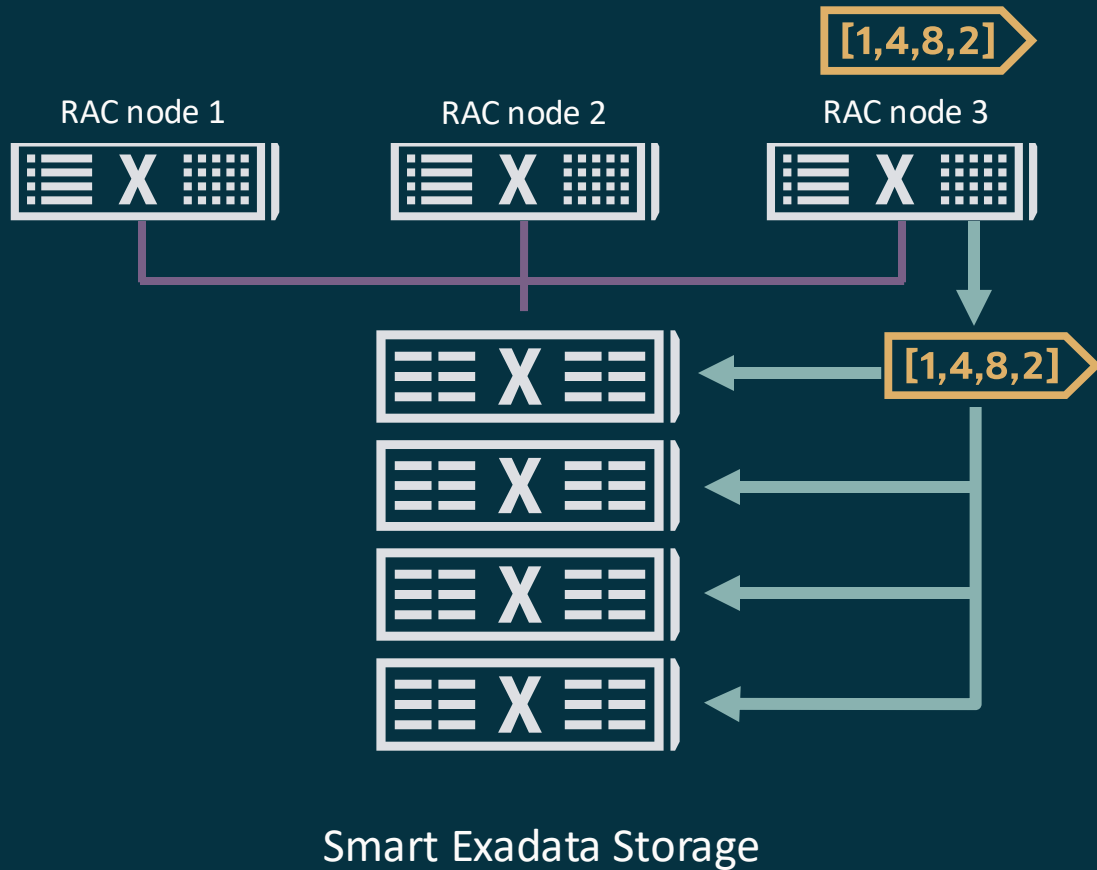




# Exadata Exascale

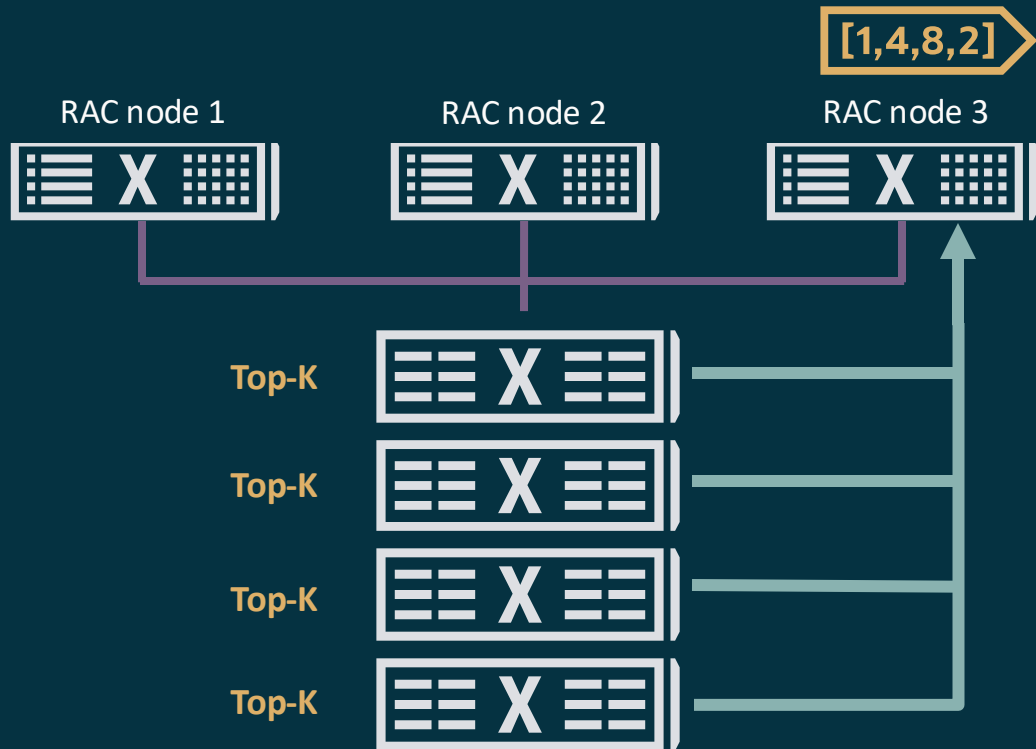
Delivers mission-critical AI at any scale





Oracle's AI vector search can be **transparently offloaded** to smart Exascale storage for faster search

Vector search queries are automatically parallelized across the Exascale storage cloud



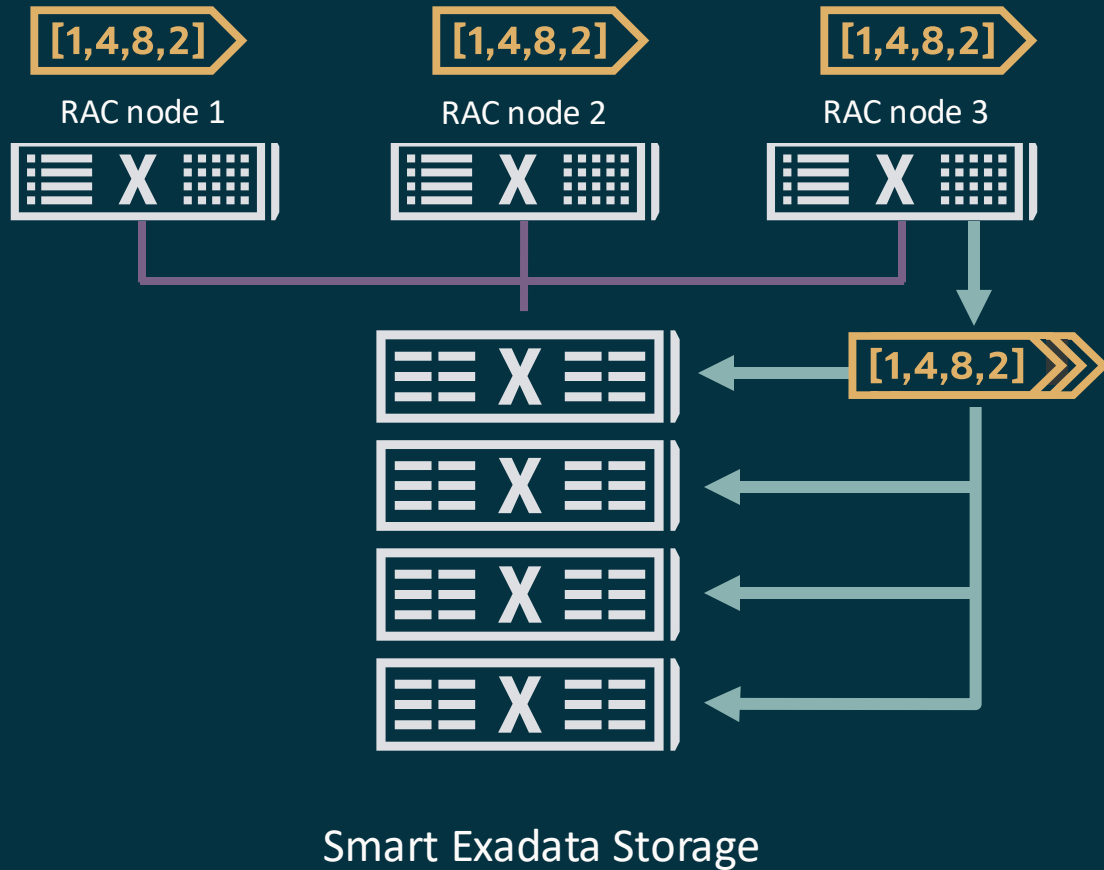
Smart Exadata Storage

Each storage server independently computes the top-K matches

- Database merges results

Works transparently on Exascale

Up to **30x** faster AI vector queries



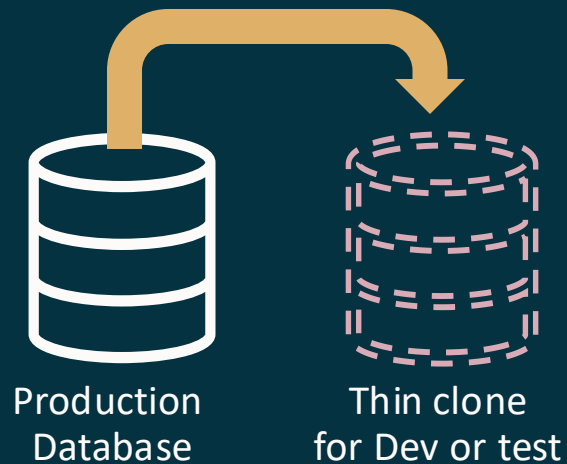
Supports multi-user environments with the ability to offload thousands of concurrent AI vector searches

# Exadata Exascale

## Database Aware Intelligent Cloning

---

# Database aware intelligent clones



Exascale can instantly create database (or PDB) clones for Dev or Test

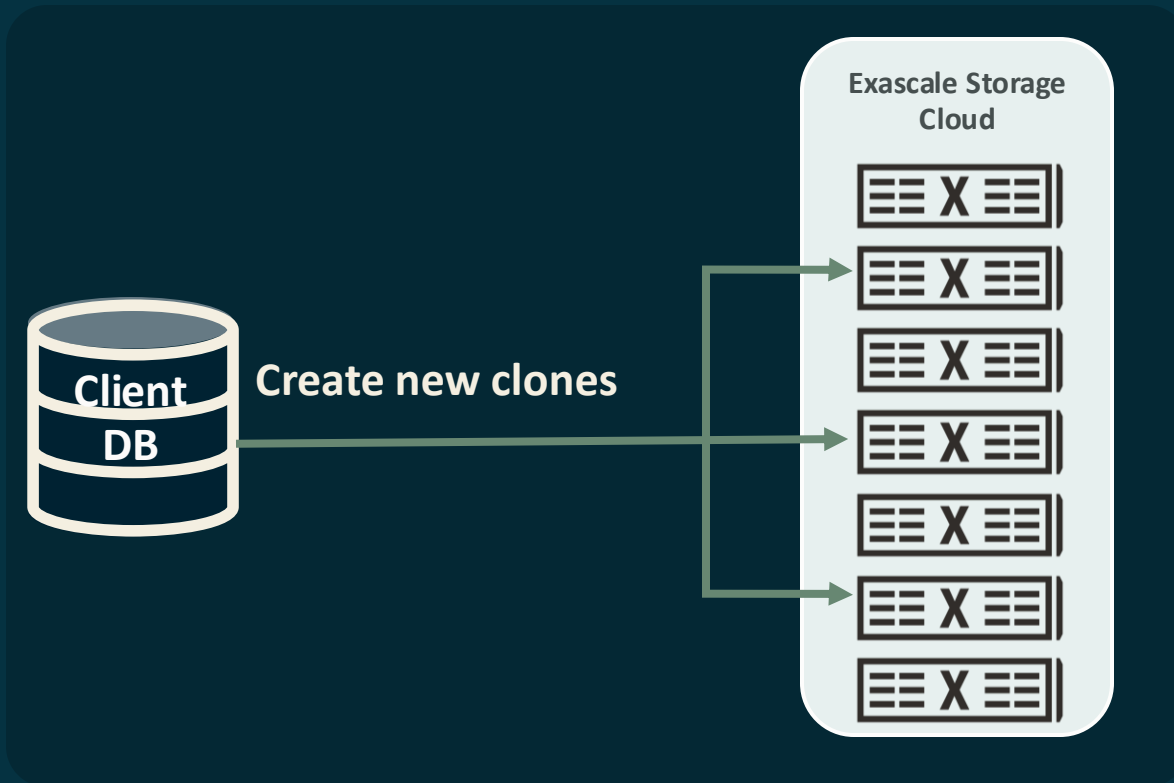
- Either a full copy or thin clones
- Created from a live PDB or an existing snapshot
- No upstream dependencies - no read-only test master, etc.

Clones leverage Exascale redirect-on-write technology

- Clone shares block with parent until they make a change
- Drastically reduces storage capacity needs for cloning

You get Exadata native performance on development, test, or recovery copies

# Exascale provides **scalable distributed clones**



When a database tells storage servers to create a **new clone**

Each storage server starts redirecting to the snapshotted data

Fast, space-efficient, minimal copies

**Distributed local clones enables Hyperscale**



# Database Aware Intelligent Thin Clones



**Thin Clone** a Pluggable Database into the same container database



Thin Clone a Pluggable Database into a **different container database**



**Snapshot Carousel of pluggable databases** - new snapshot every few hours



Thin Clone **entire Database** new database with snapshot files of source database



```
CREATE  
PLUGGABLE DATABASE pdb1c  
FROM pdb1  
SNAPSHOT COPY;
```



```
CREATE  
PLUGGABLE DATABASE pdb1c  
FROM pdb1@CDB1_link  
SNAPSHOT COPY;
```




```
ALTER  
PLUGGABLE DATABASE pdb1  
SNAPSHOT MODE  
EVERY 2 hours;
```



```
gDBC1one
```

# Database Aware Intelligent Thin Clones with Data Guard



Thin clone a **pluggable database from standby** to Test Dev container database

```
🟡🟢🟠
edit database <standby> set state='apply-off';
CREATE
pluggable database
salestest FROM
sales@standby_link
SNAPSHOT COPY;
edit database <standby> set state='apply-on';
```



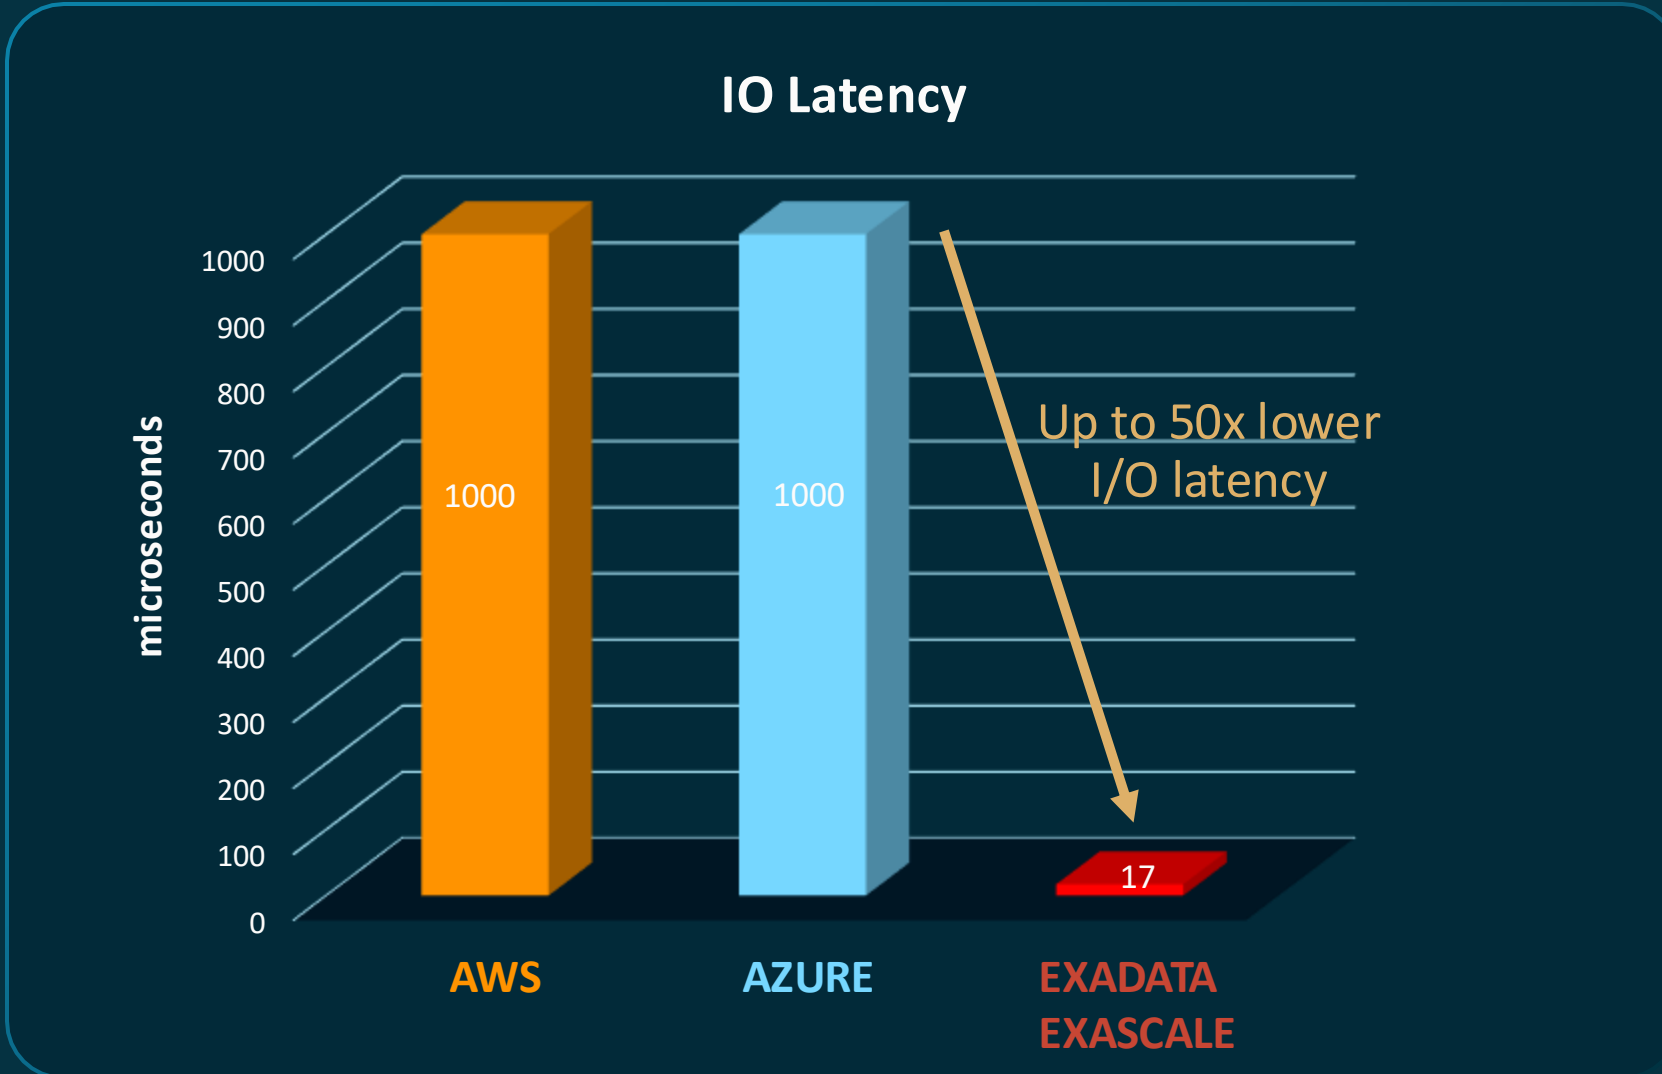

Thin clone **entire container standby database** into a new Test/Dev container database

```
🟡🟢🟠
gDBCclone.bin snap -sdbname cdbSales \
                  -tdbname cdbSalesTest \
                  -racmod 2
```



# Exadata Exascale IO Performance



# Dramatically Lower IO Latency than Cloud Flash Storage



# Dramatically Higher **Analytics Performance** than other cloud storage





We discussed the architecture, cloning benefits, and performance of Exascale

Now let's look at how easy it is to get started with Exascale thanks to our new cloud service





# Introducing Exadata Database Service on Exascale Infrastructure

*Ref: <https://blogs.oracle.com/database/post/introducing-oracle-exadata-database-service-on-exascale-infrastructure>*

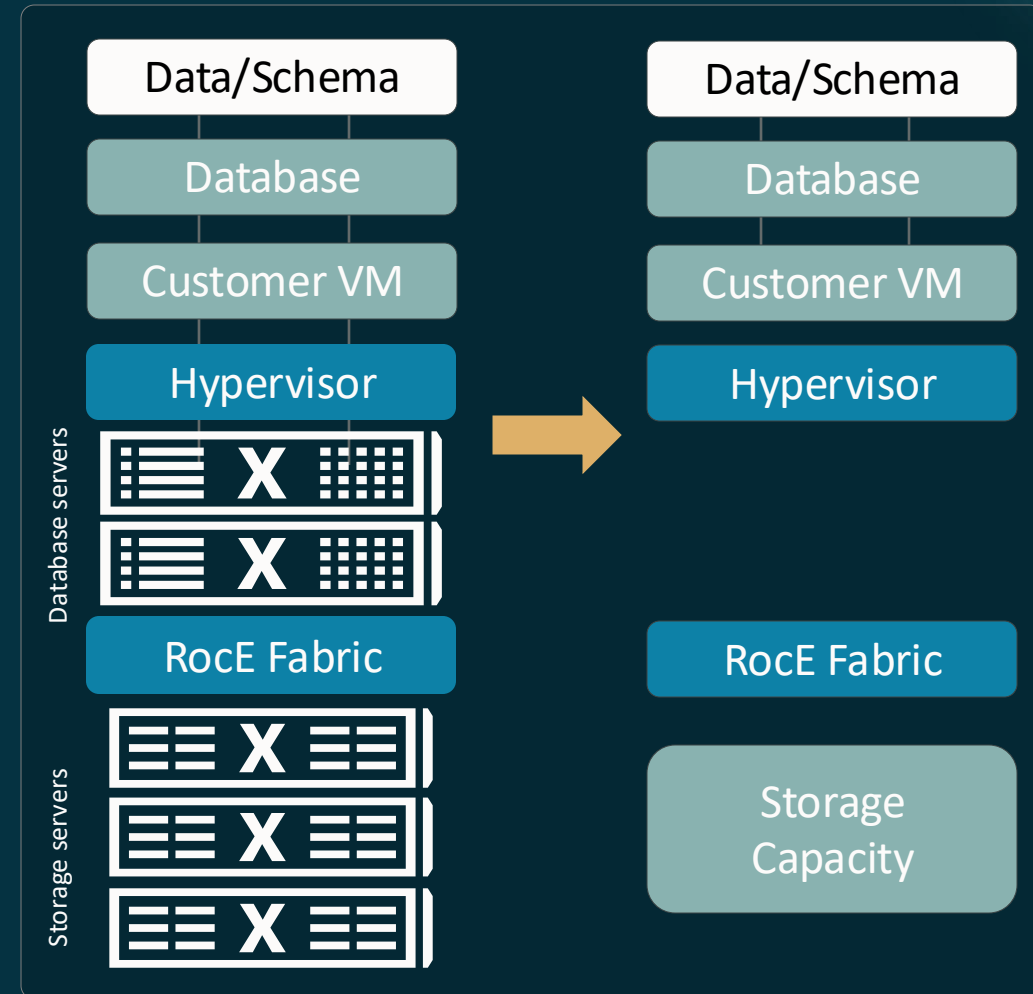


# Database Service on Exascale Automates Database Management

Provides the same powerful database automation available of the current Exadata Database Cloud Service

The most visible differences are:

- Tenants only see databases and VM clusters
  - The cloud of physical servers is invisible to tenants
- Linux VM images are stored on the network attached Exascale intelligent block storage
  - Removes size limitations of local drives
  - Provides foundation for fast VM migration in future



# Database Service on Exascale Automates **Storage Management**

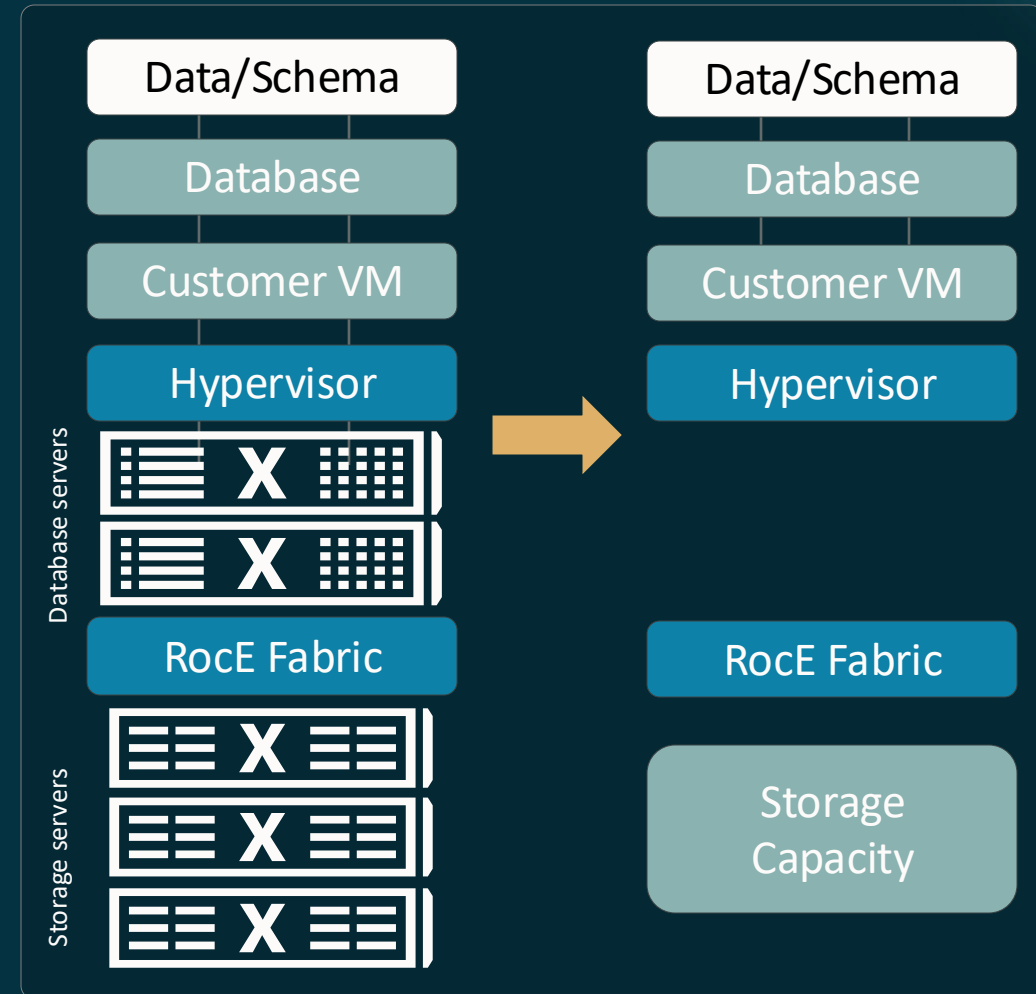
Exascale storage virtualization is fully automated by Exascale Cloud Control Plane

You allocate and pay for only the storage capacity you need for your databases

**You don't pay for IOPS**

Roadmap: support live VM migration while performing RDMA IO to storage and other VMs

- Avoids database disruption during cloud infrastructure software updates



# Now Available Exadata Exascale for Exadata on-premises

---


*Ref: <https://blogs.oracle.com/exadata/post/oracle-database-23ai-on-exadata>*

# Key Takeaways

---



# Key Benefits of Oracle Exadata Database Service on Exascale Infrastructure



## Powerful

All unique Exadata performance, reliability, availability and security capabilities built-in



## Extreme Low Cost

Pay-per-use, low entry cost, and no cost per IO provides all the Exadata benefits at extremely low cost



## Scalable Pooled Resources

Complete elasticity enables you to provision any resources you need when you need them



## Agile

Boosts developer productivity with rapid dev/test provisioning while lowering storage costs

Now all Exadata benefits are available for **small workloads** and **small businesses**

# Exascale's Intelligent Data Architecture for Cloud - Key Differentiators

Exascale is Exadata reimagined for a multi-tenant hyper-elastic Cloud

Provides all Exadata Intelligence plus all benefits of Cloud

Exascale delivers the world's only **intelligent storage cloud**

Exascale delivers the world's only **RDMA-capable storage cloud**

Exascale provides extreme performance for **intelligent database clones**

A **shared pool of cloud compute and storage** unleashes hundreds of CPUs for any database query

Exascale **accelerates all database workloads at extreme low cost**

- Mission-critical, analytics, AI vectors, test/dev, etc.



For more information | [oracle.com/exadata](https://oracle.com/exadata)