ORACLE

Exascale: the next generation architecture of Oracle Exadata

Diego Losi

Exadata vision

Extreme scalability, performance and availability for all data workloads, at the lowest cost

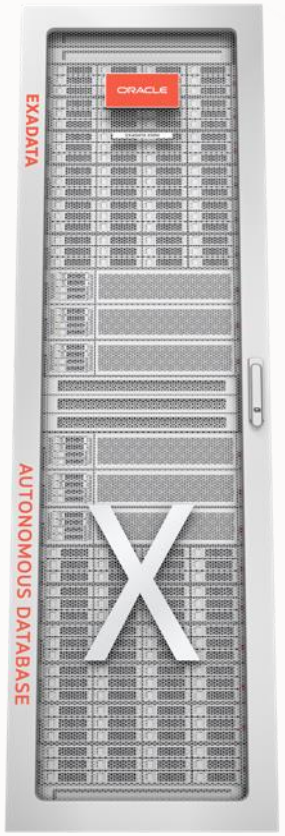

Ideal Database Hardware

Scale-out, database optimized compute, networking, and storage

Database Aware System Software

Unique algorithms vastly improve the performance of modern operational and analytic apps for all use cases at any scale

Automated Management

Fully automates and optimizes the end-to-end stack

Thousands of global customers run their business on Exadata

77% of Fortune Global 100 Run Exadata | **51%** 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

Exadata Exascale

All Exadata intelligence

All cloud benefits

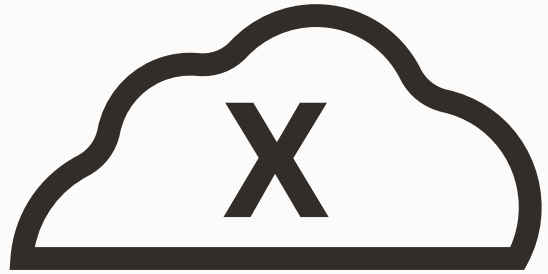

Exadata Exascale

New intelligent data architecture
that delivers
the **best of Exadata** and
the **best of cloud**

All Exadata intelligence

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 clones 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

Intelligent AI

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 Pools

Provision any resources you need when you need them

Hyper-elastic

Great for small businesses and the largest enterprises

Exadata Exascale Cloud Architecture Details

Previously tenants had dedicated Exadata compute and storage servers

ASM was used to distribute storage across databases

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 storage management tiers

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

- Direct IO architecture with no intermediate nodes
- RDMA IO requests are implemented in hardware enabling microsecond latency and millions of IOs per second of throughput

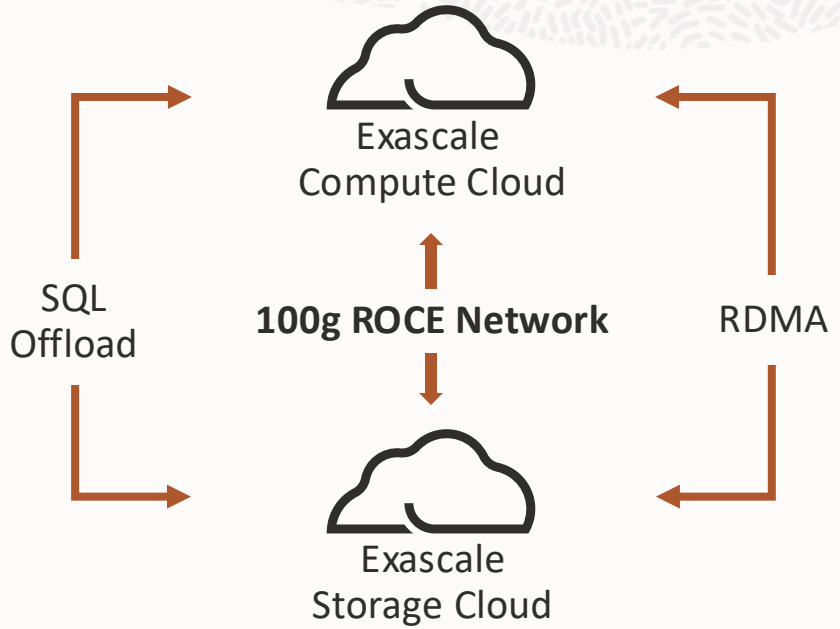

Exascale uniquely optimizes storage capacity in addition to performance

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

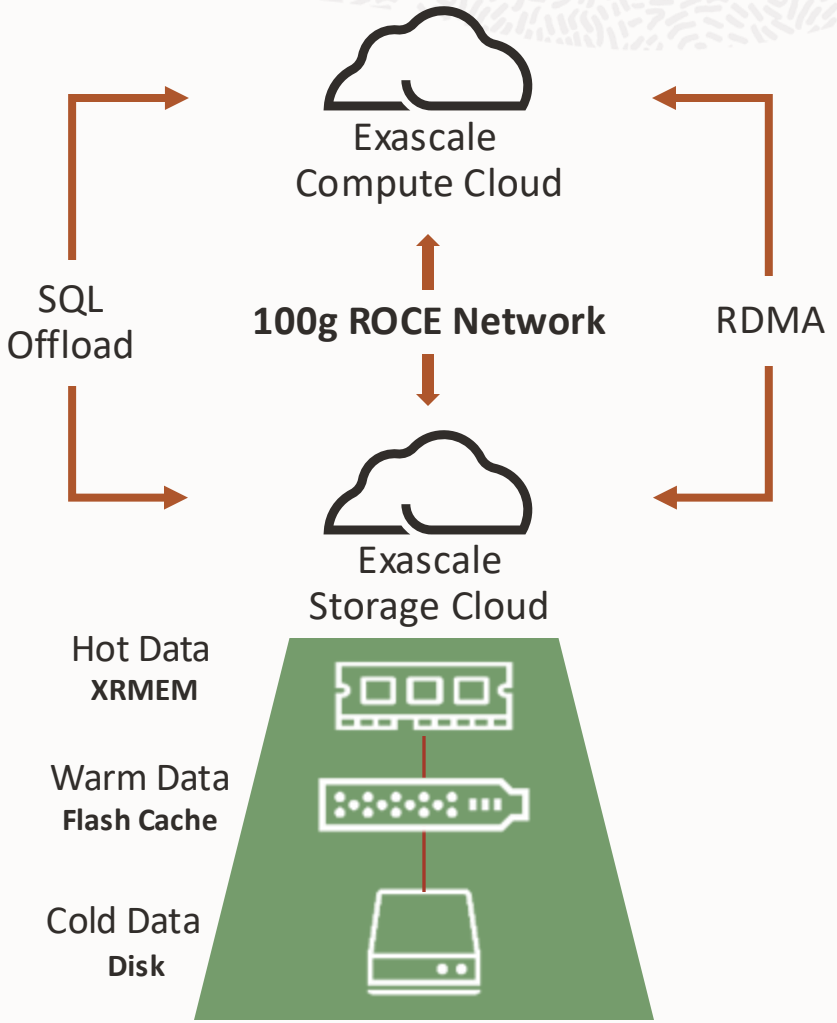

Exascale manages data as a large pool of Database Extents

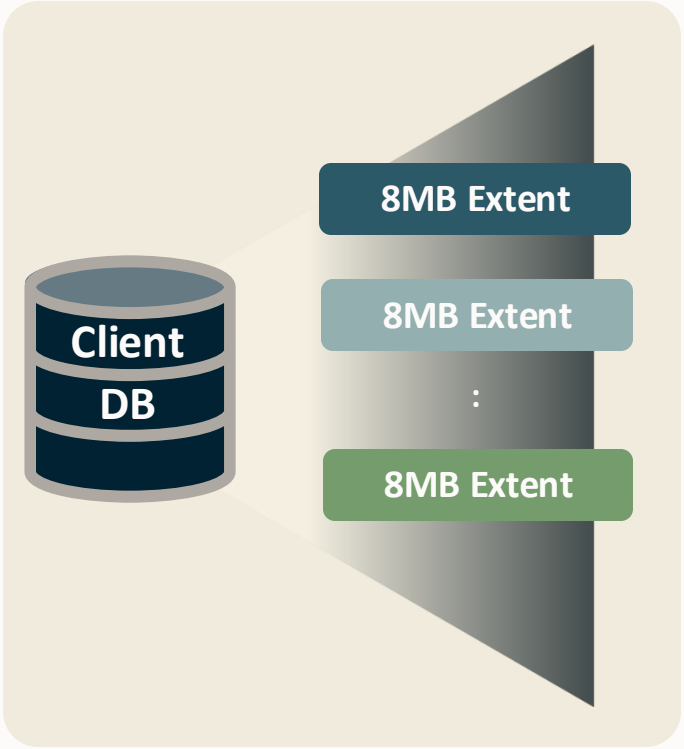

Exascale stores databases as a set of 8 MB Extents

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

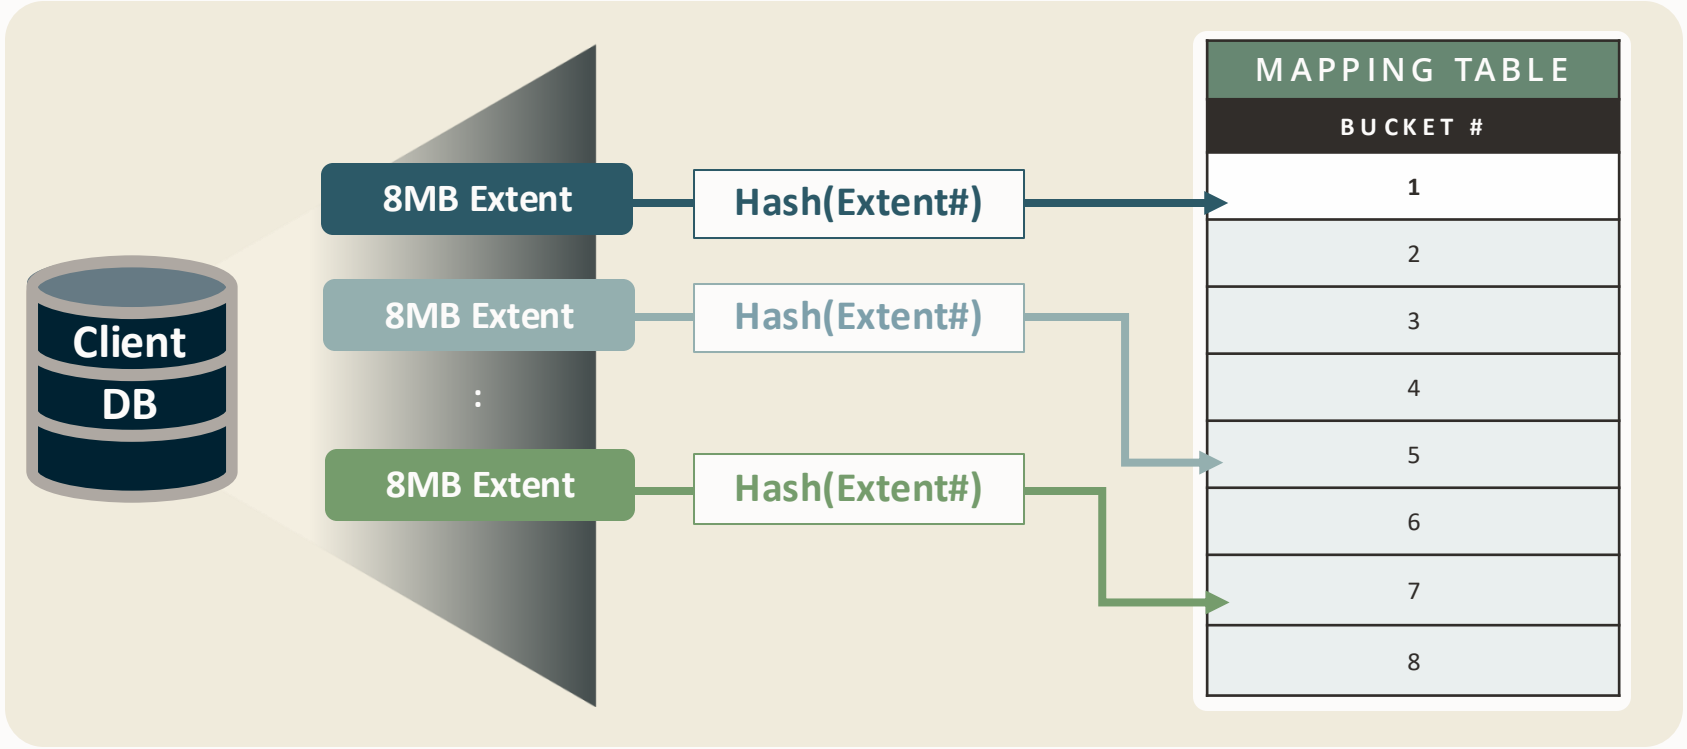

Multiple Extents can hash to a Storage Bucket

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

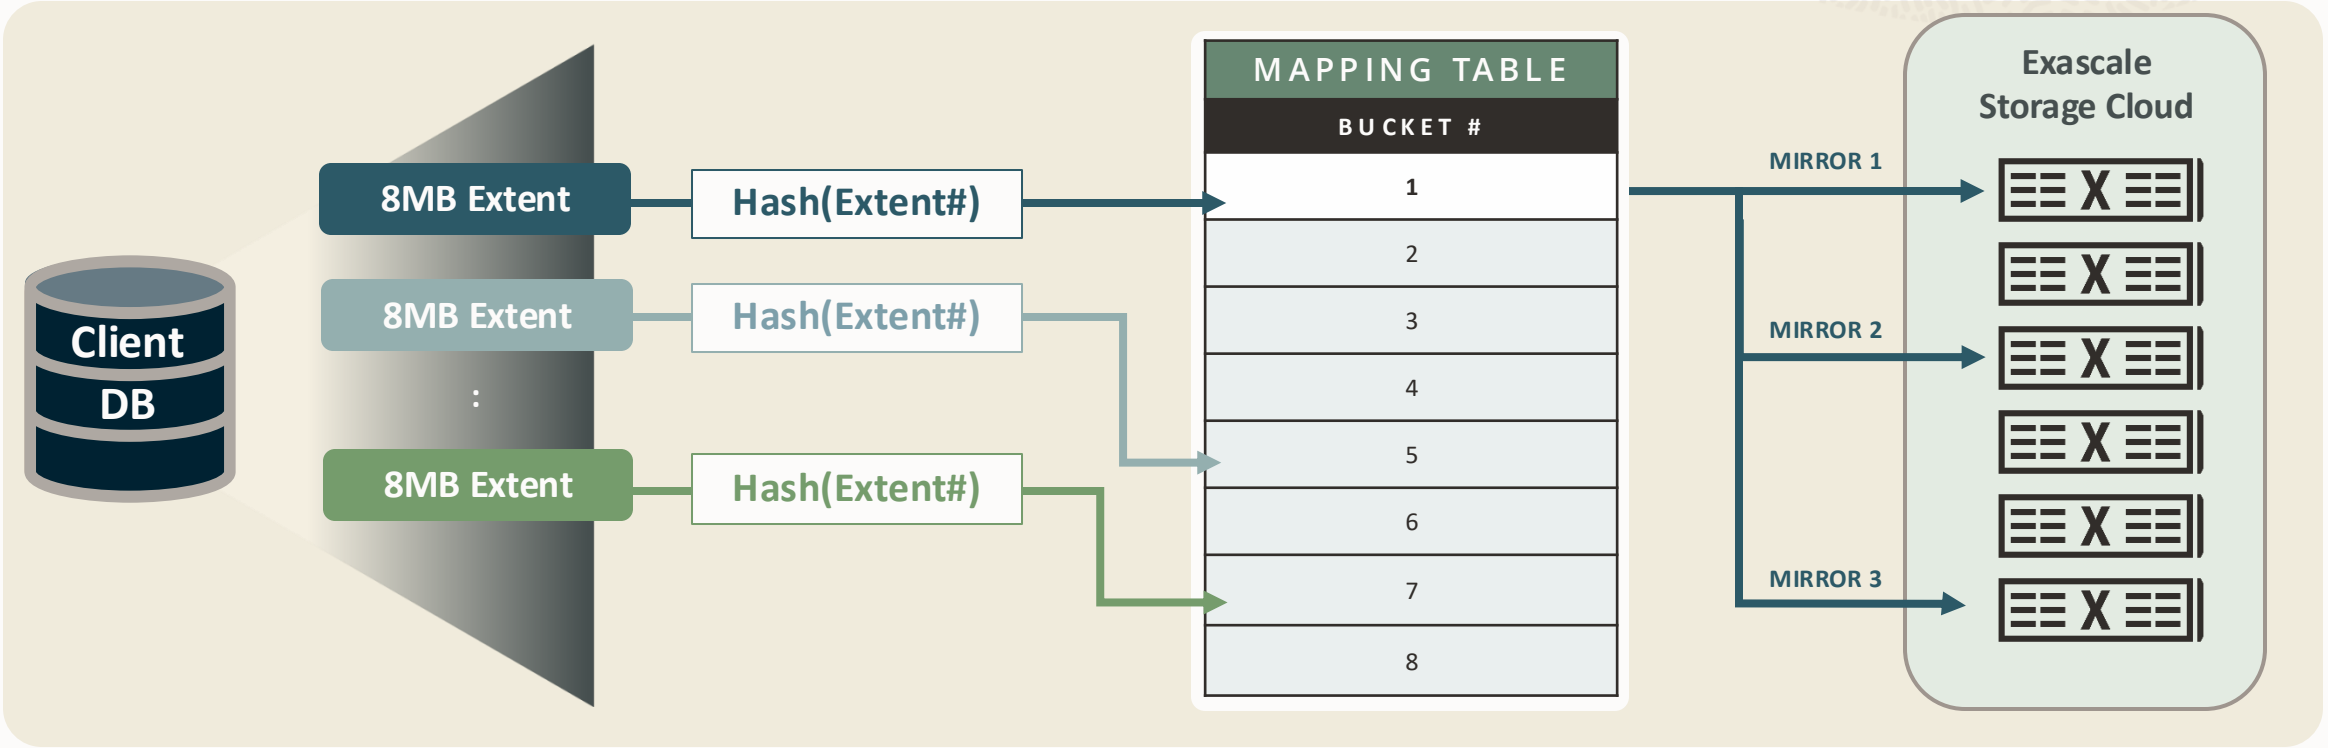

All extents that hash to a storage bucket are stored redundantly on 3 drives that are on 3 separate storage servers to protect from storage failures and outages

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

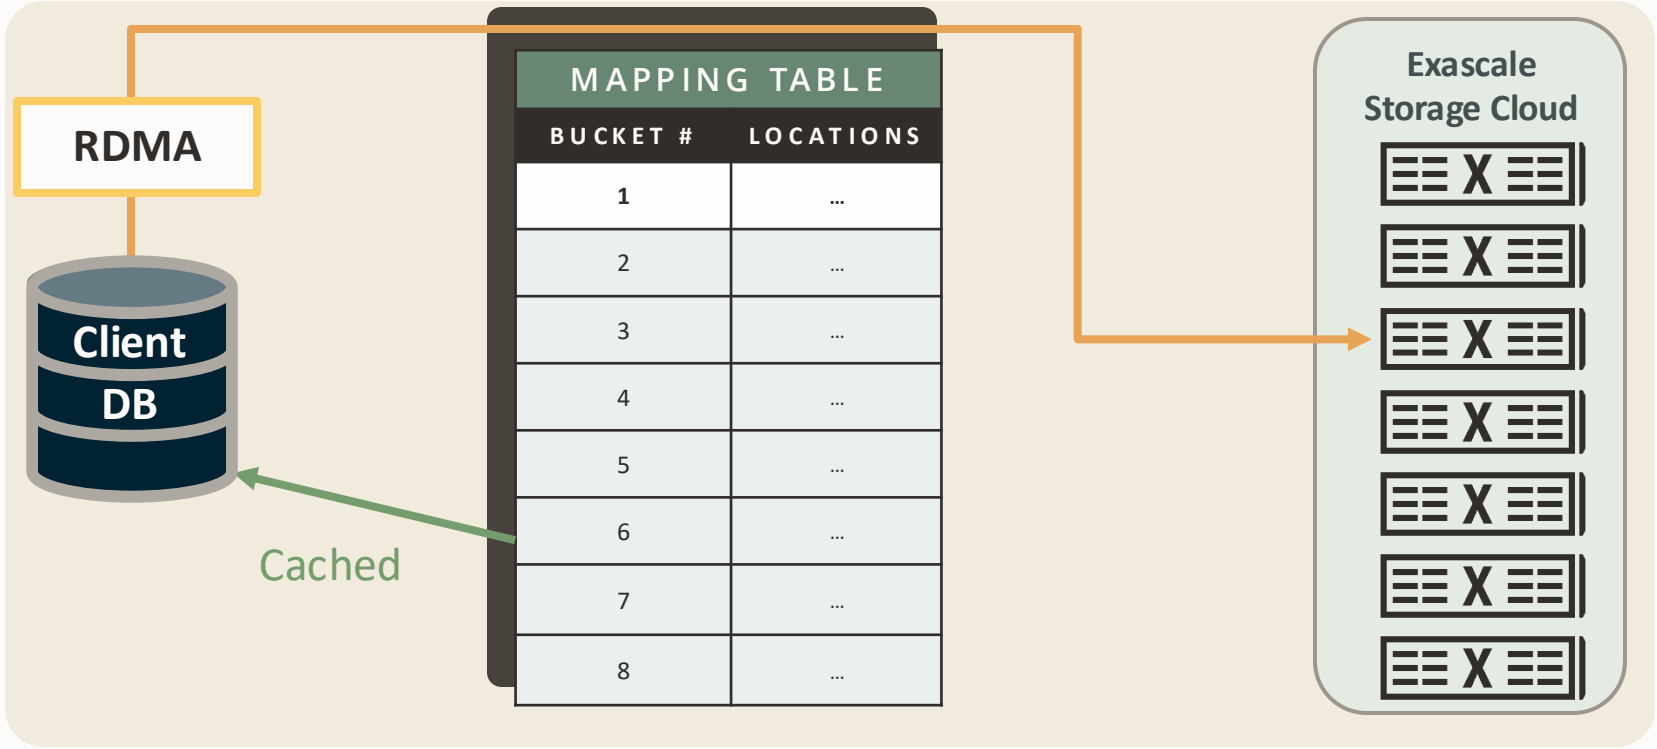

Each client database caches the Mapping Table

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

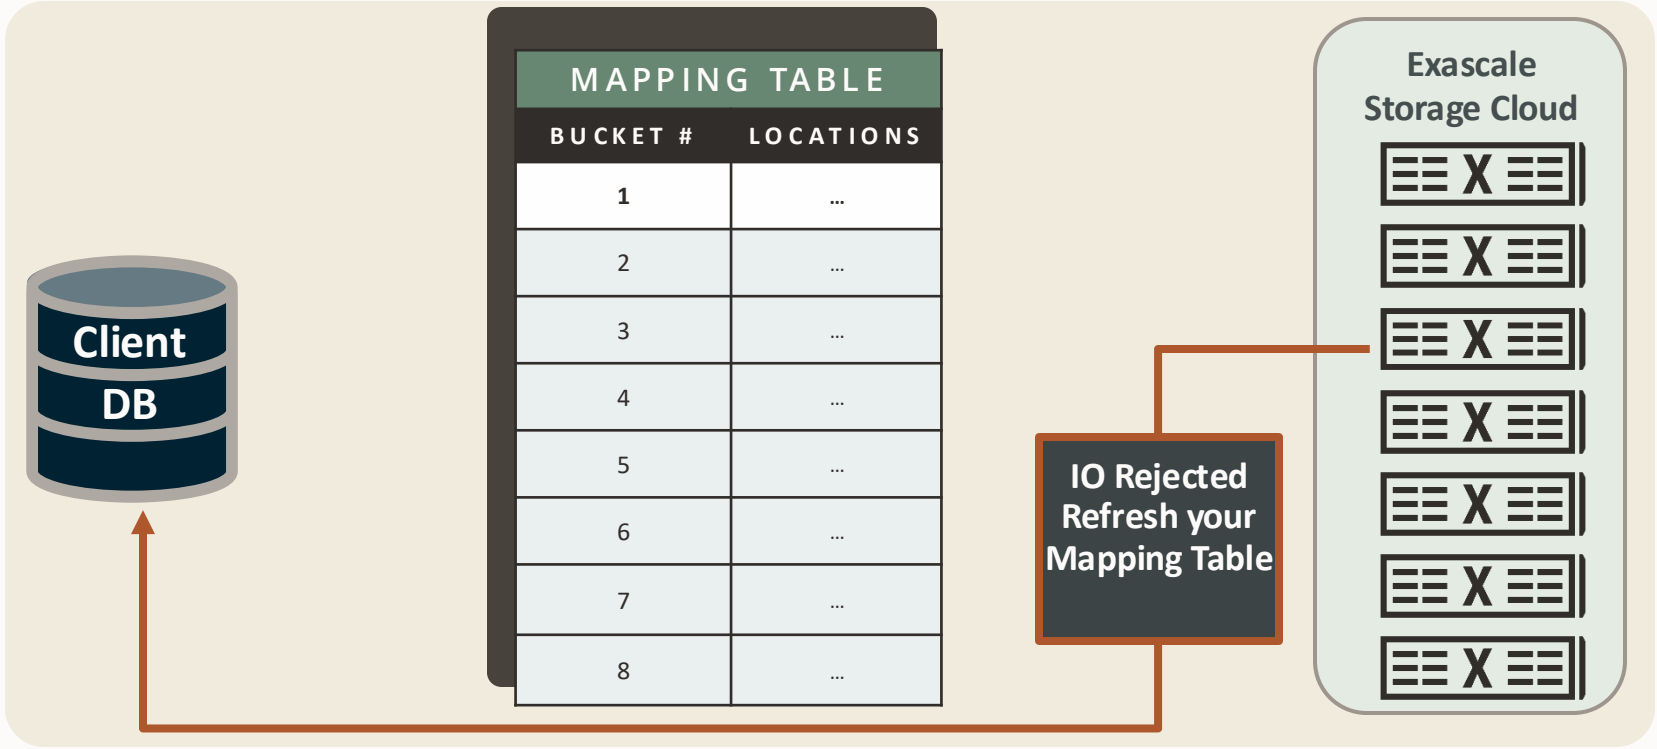

No need for distributed locking when data is re-distributed across servers

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

Loosely Coupled Architecture Enables Hyper-Scale

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

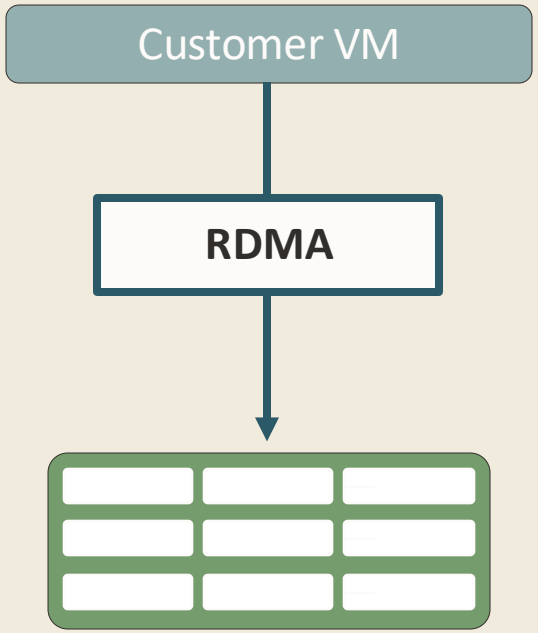

Exascale architecture – storage server

https://docs.oracle.com/en/engineered-systems/exadata-database-machine/exscl/exascale-services.html

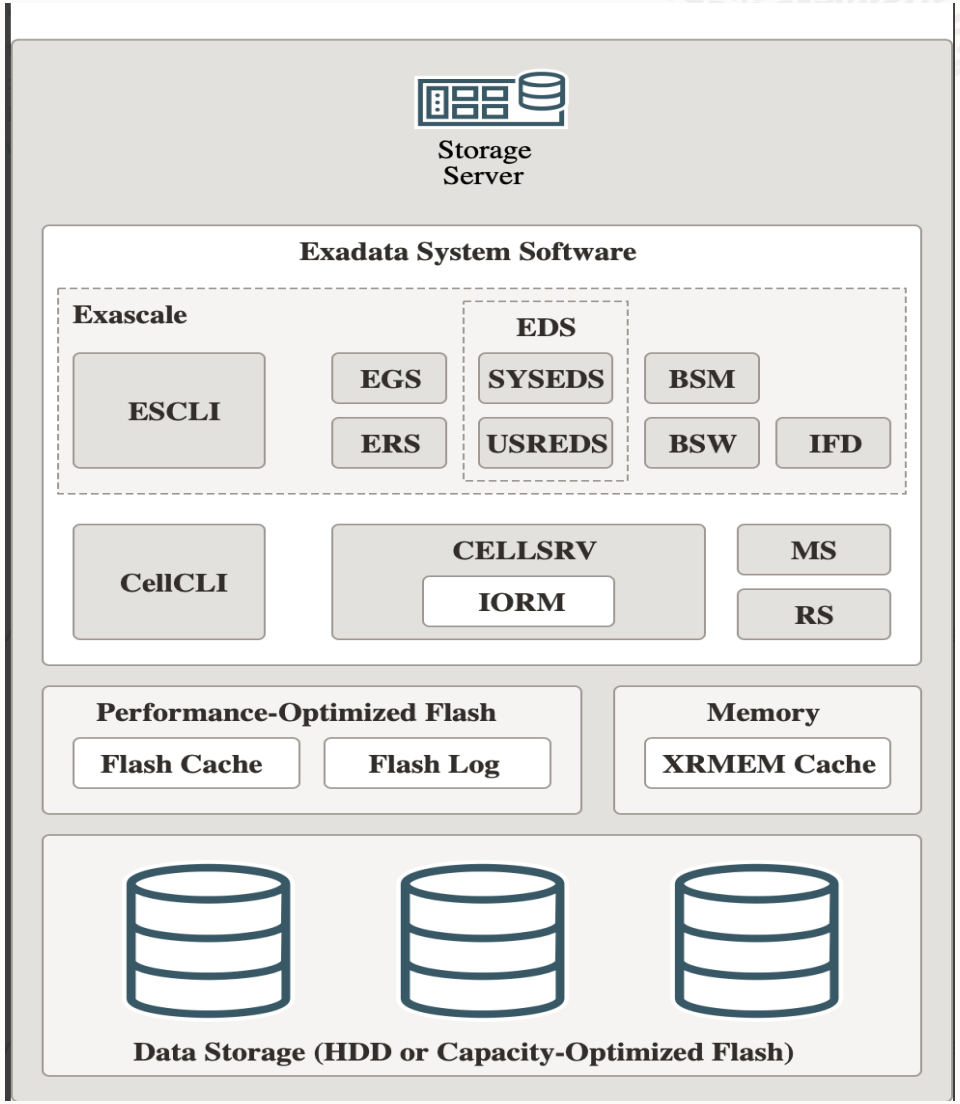

ASM storage

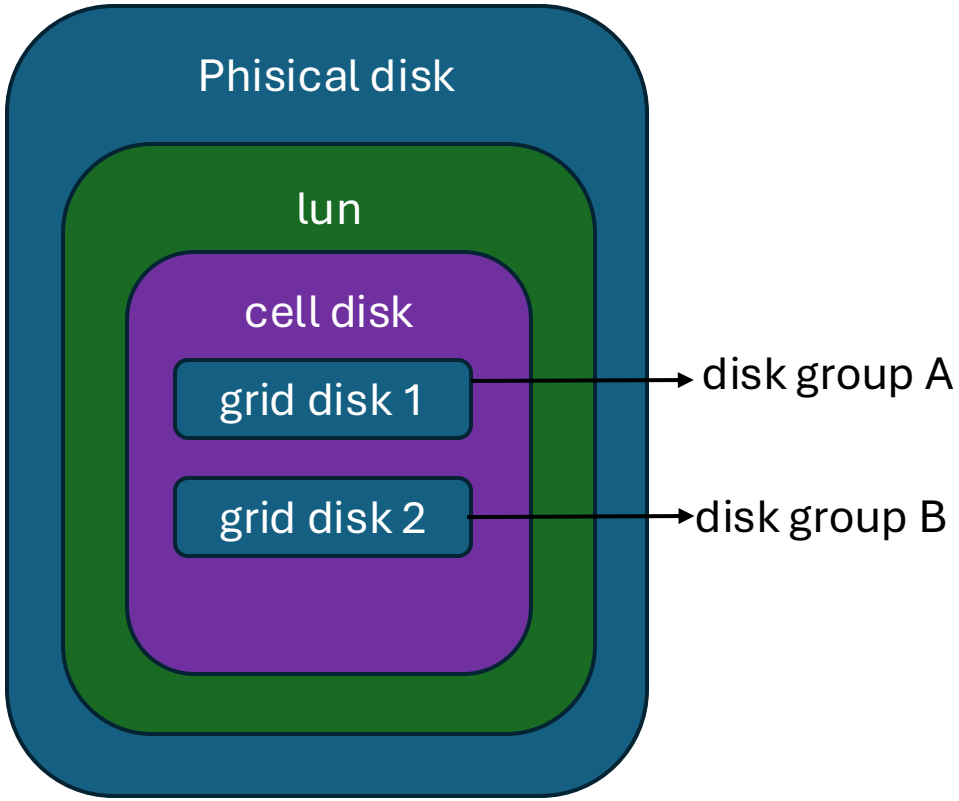

Exascale storage

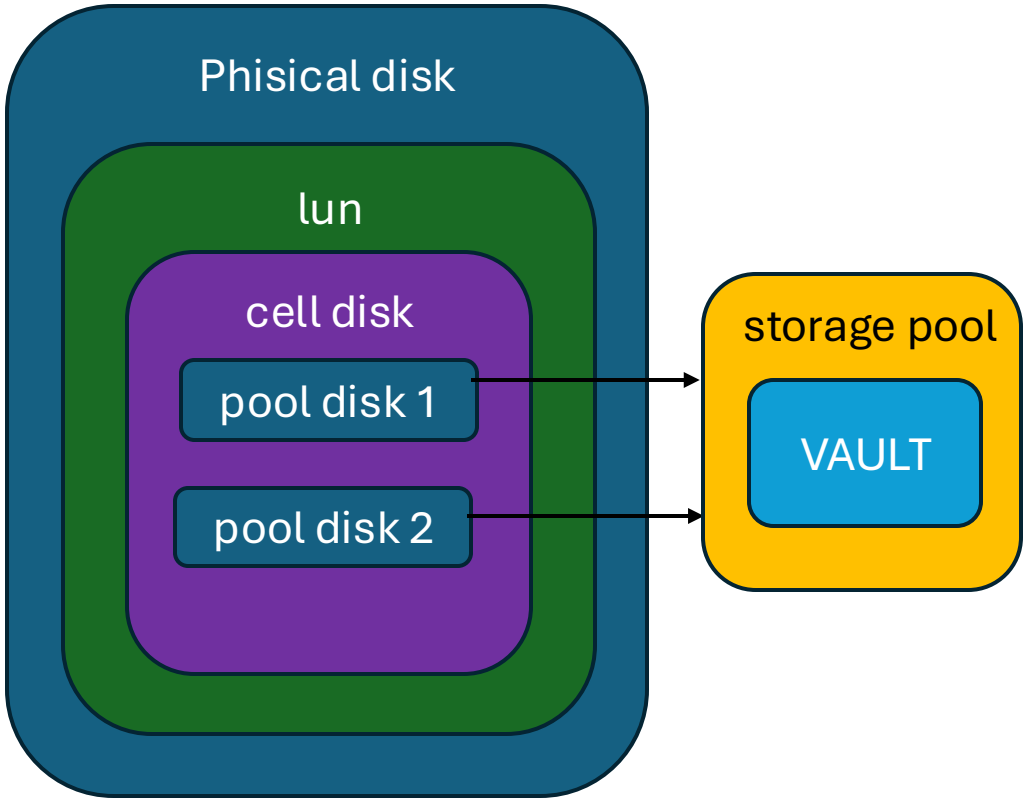

Exadata Database Service on Exascale Infrastructure

ExaDB-XS Core Business Drivers

Base Database Service

VMs on shared compute

No infrastructure to worry about

Small, economical starting point

Meets basic database performance needs

Limited scalability and consolidation

What if I want benefits of both

Exadata Database Service

Dedicated Exadata Infrastructure

Infrastructure hosts VM Clusters

All enterprise workloads

Exadata performance and availability

Massive scalability, large starting point

Overview of Exadata Database Service on Exascale Infrastructure

Exadata Database Service in the public cloud is now even simpler

- No dedicated hardware
- Provision compute resources, Exascale storage, network, SSH keys
- Instantly create efficient database clones and snapshots

Start small and grow CPU and storage resources online

- Begin with a cluster of 2 VMs, each with 8 ECPUs and 22 GB of memory
- Start out with 300 GB of Exadata database storage
- Grow VMs in increments of 4 ECPUs each, add more VMs, and scale Exascale storage with Gigabytes or Terabytes of additional capacity

Powerful cloud automation

- Web and API-driven provisioning, updates, backups, and DR lifecycle operations

Affordable

All customers, any scale

Co-managed with Oracle

Full administrative control

Dev/test agility

Overview of Exadata Database Service on Exascale Infrastructure

High performance and availability for all database workloads

- Benefit from all Exadata features and optimizations, including:
  - Super-fast SQL latency with RDMA
  - Smart Scan for analytics
  - AI Smart Scan for AI Vector Search
  - Automatic storage tiering
  - Automated resource management
- Use all Oracle Database 23ai features, including:
  - Converged database capabilities
  - AI Vector Search
  - JSON Relational Duality
  - Globally Distributed Database
  - Lockless Reservations

Affordable

All customers, any scale

Co-managed with Oracle

Full administrative control

Dev/test agility

Automates Database Management

Provides the same powerful database automation currently available on Exadata Database 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 and will enable fast VM migration in the future

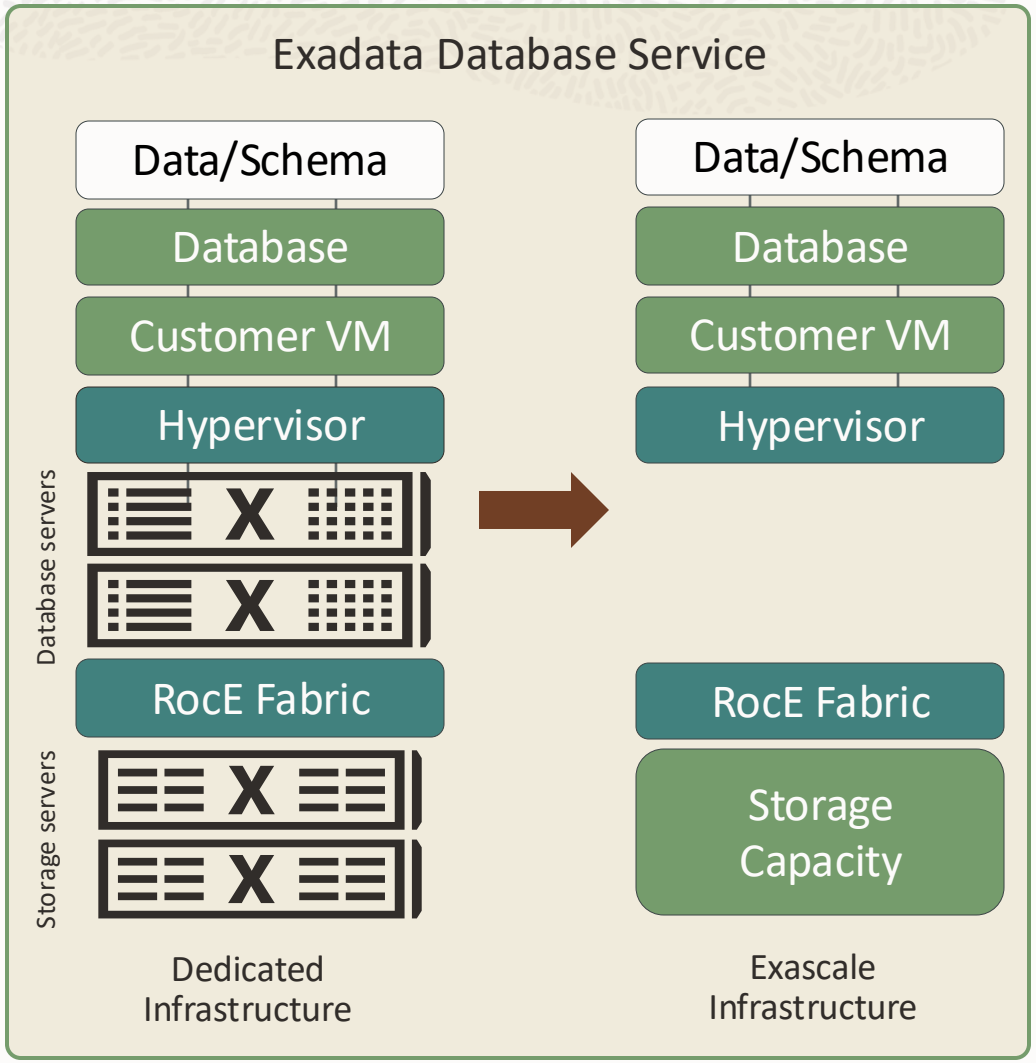

Automates Storage Management

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

In the future will support live VM migration while performing RDMA IO to storage and other VMs

- Avoids database disruption during cloud infrastructure software updates

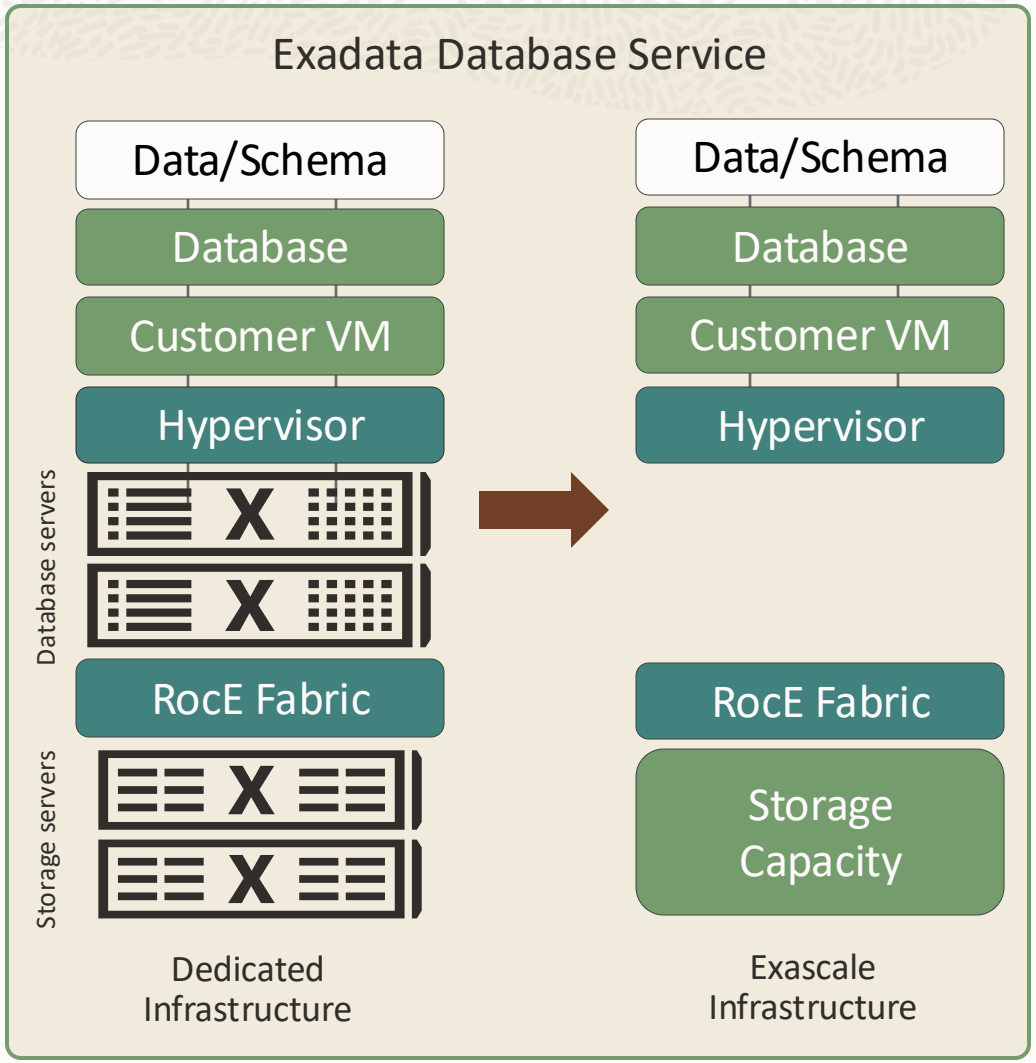

ExaDB-XS Bridges the Gap Between BaseDB and ExaDB-D

BaseDB

Lowest Cost DB Cloud Entry point
Good-to-Great Performance
Smaller Workloads

ExaDB-XS

Exadata's extreme performance and availability at a lower price point – Supporting ANY workload

ExaDB-D

Extreme Performance
Dedicated Hardware
Enterprise Workloads

	BaseDB	Exascale Infrastructure						Dedicated Exadata Infrastructure	
Provisioned Resources	**20 OCPU 3TB**	**16 ECPU 300 GB** (4 OCPU)	**32 ECPU 1 TB** (8 OCPU)	**64 ECPU 3 TB** (16 OCPU)	**96 ECPU 5 TB** (24 OCPU)	**128 ECPU 10 TB** (32 OCPU)	**256 ECPU 35 TB** (64 OCPU)	**Base System** 48 OCPU 96TB	**Quarter Rack** 252 OCPU 150TB
Infrastructure	$185	$356	$737	$1570	$2401	$3589	$8929	$8000	$10,800
Database	$2880	$961	$1561*	$3182*	$4743*	$6364*	$12,368*	$7680**	$7680**
Price/month	$3065	$1317	$2299	$4752	$7415	$9953	$21,297	$14,720	$17,520
Hourly Price	$4.12/hr	$1.77/hr	$3.09/hr	$6.39/hr	$9.60/hr	$13.38/hr	$28.62/hr	$19.78/hr	$23.55/hr

* ExaDB-XS cost is an approximation assume a 30% peak, 30% of the time
** Assumes 32 OCPUs running on the Dedicated Infrastructure

Demo

ExaDB-XS provisioning

https://medium.com/@diegolosi487/exadb-xs-or-exadata-database-service-on-exascale-infrastructure-b1321a877e3e

Database-aware Intelligent Clones

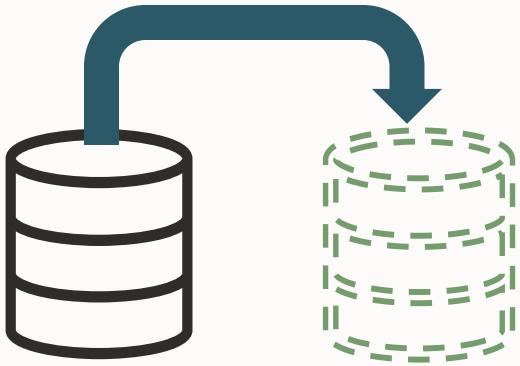

Instantly create database or PDB clones for Development or Test

- Either a full copy or thin clones
- No complex prerequisites - no read-only test master, etc.
- Clone across CDBs and VM clusters

Clones leverage Exascale redirect-on-write technology

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

Get native Exadata performance on development, test, or recovery copies

Empower developers by creating thousands of PDB clones for dev teams

Thin clone 1/2

Thin clone 2/2

Key benefits of Oracle Exadata Database Service on Exascale Infrastructure

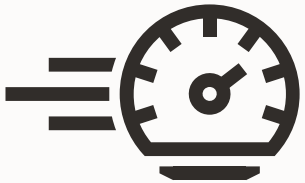

Powerful

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 Exadata performance at extremely low cost

Scalable Pooled Resources

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

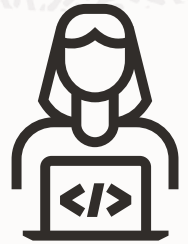

Agile

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

Now organizations of any size and workloads at any scale can benefit from Exadata

Dramatically Lower IO Latency than Cloud Flash Storage

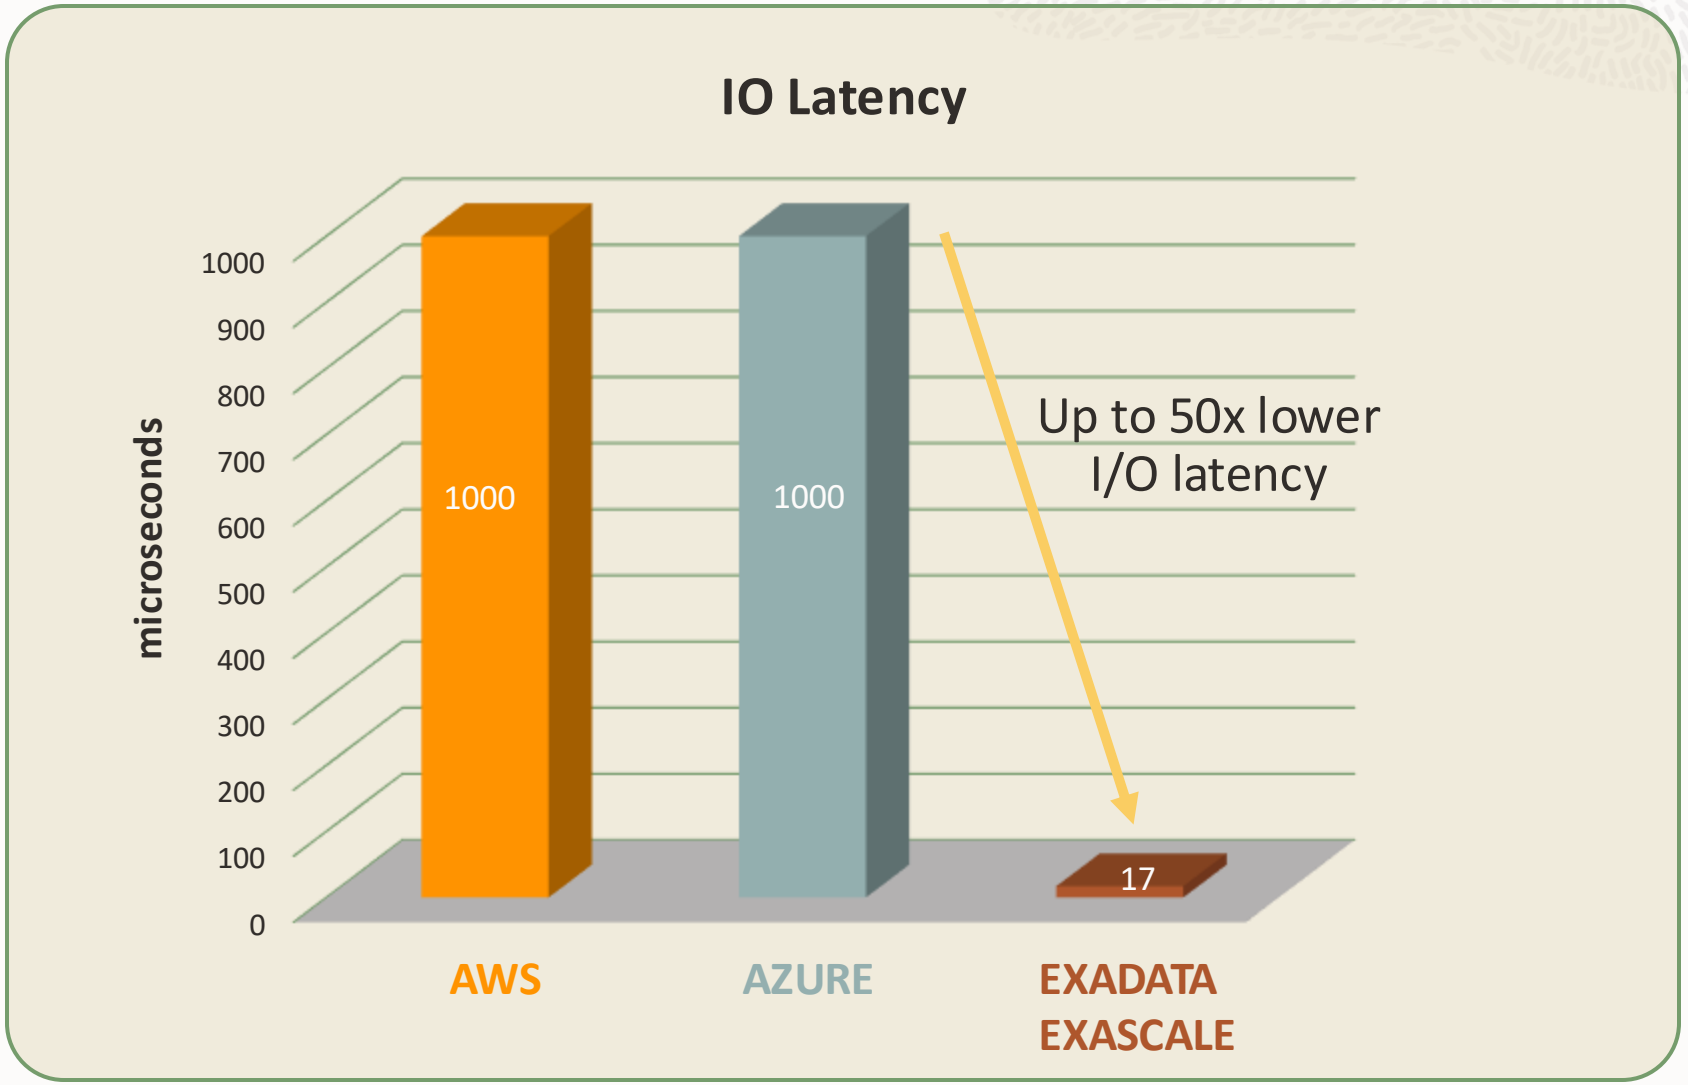

Dramatically Higher Analytics Performance than other cloud storage

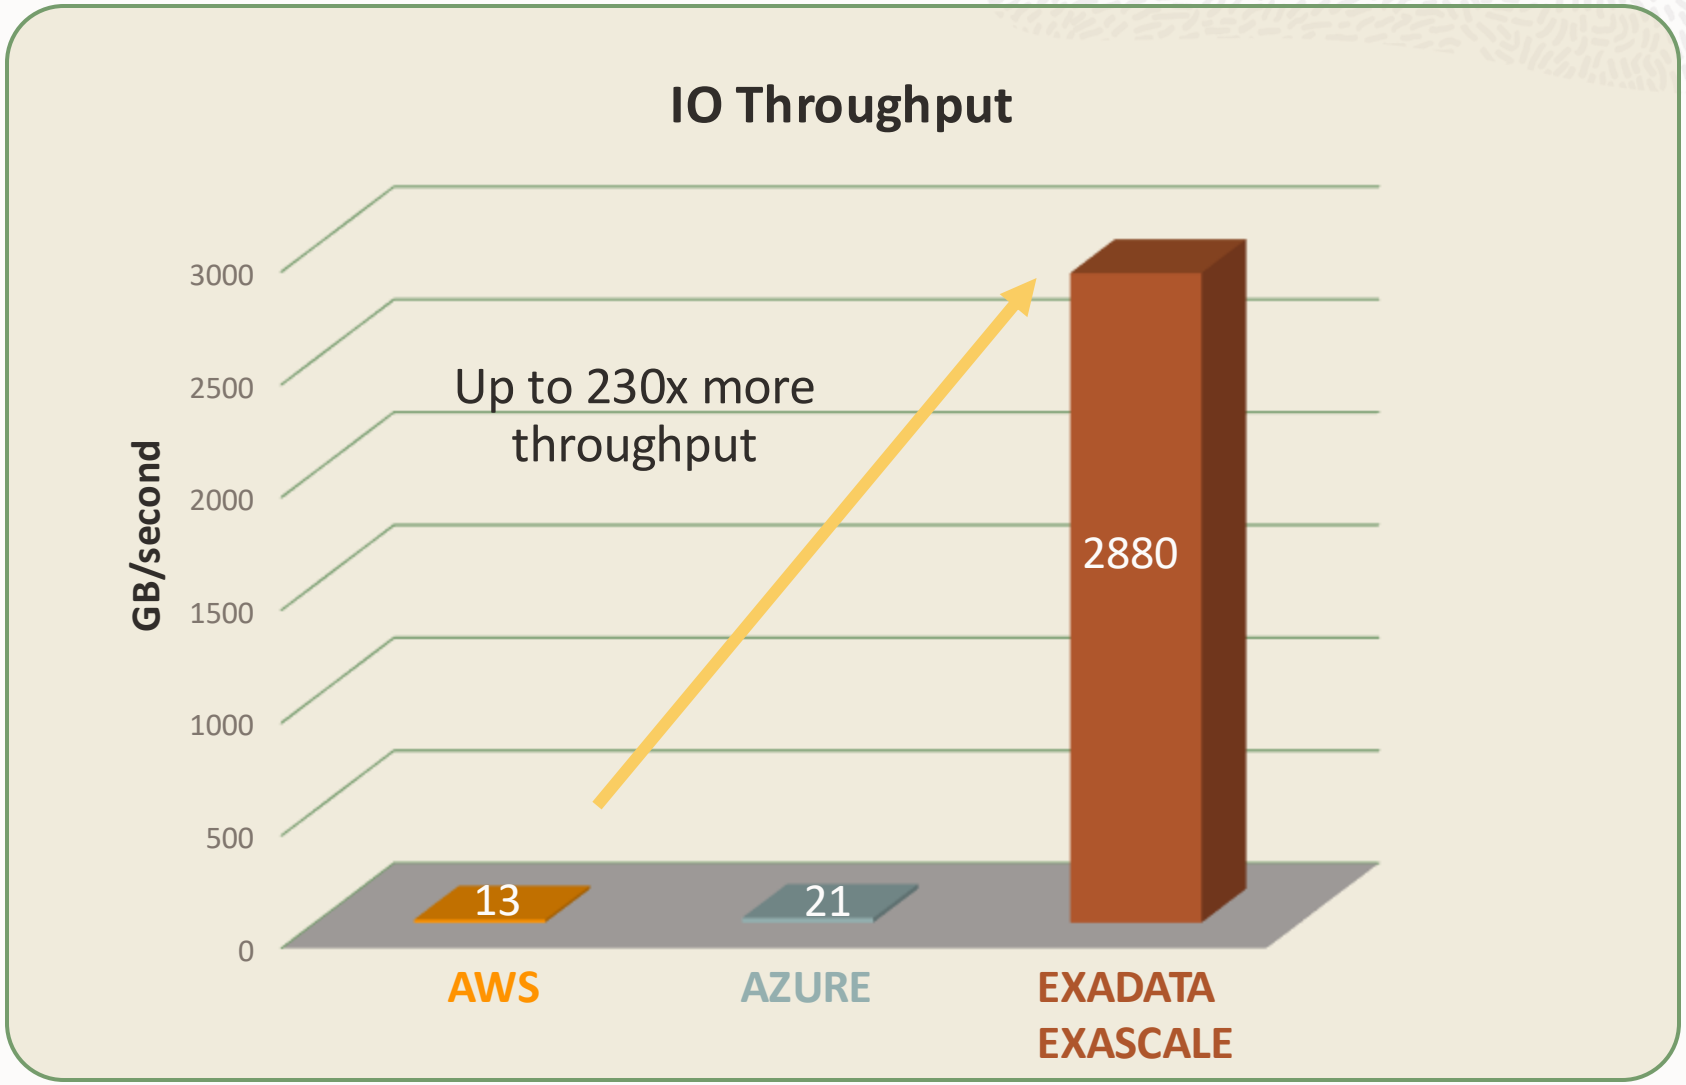

Move database to ExaDB-XS

It is a database !

If you want to check how to move database using data guard please check

https://medium.com/@diegolosi487/how-to-setup-a-dataguard-between-base-database-and-database-running-on-exadb-xs-bc9624ca595a

I have to use new utility *XSH* instead of *asmcmd*

Exadata Exascale

Delivers mission-critical AI at any scale