



Research Brief:

ORACLE GLOBALLY DISTRIBUTED EXADATA DATABASE ON EXASCALE INFRASTRUCTURE

STEVE MCDOWELL, CHIEF ANALYST

AUGUST 7, 2025

ORACLE GLOBALLY DISTRIBUTED EXADATA DATABASE ON EXASCALE INFRASTRUCTURE

Oracle announced the general availability of its Globally Distributed Exadata Database on Exascale Infrastructure, a managed cloud service that combines distributed database capabilities with serverless computing architecture. The platform automatically distributes and synchronizes data across multiple Oracle Cloud Infrastructure (OCI) regions, maintaining full SQL compatibility and support for Oracle Database features.

The offering is Oracle's strategic response to the growing demands of enterprises for AI-ready infrastructure, data sovereignty compliance, and global application deployment. The platform targets organizations that require always-on availability, regulatory compliance across multiple jurisdictions, and support for high-volume, agentic AI workloads.

Oracle positions its Globally Distributed Exadata Database on Exascale Infrastructure as a converged alternative to the fragmented database ecosystem prevalent in modern enterprises.

BACKGROUND: WHAT IS ORACLE EXADATA EXASCALE?

Oracle Exadata Exascale represents a fundamental architectural evolution that decouples Oracle Database and Oracle Grid Infrastructure clusters from underlying Exadata storage servers.

Its disaggregated architecture provides centralized management of large storage server fleets connected via the Exadata RDMA Network Fabric.

CORE ARCHITECTURE COMPONENTS

- **Decoupled compute-storage model:** Separates database processing from storage resources, enabling independent scaling and resource optimization
- **Centralized storage management:** Large pools of Exadata storage servers managed as unified resources accessible by multiple Oracle Grid Infrastructure clusters
- **Secure multi-tenancy:** Strict data isolation mechanisms ensure that databases access only authorized data while sharing the underlying storage infrastructure
- **Hardware requirements:** 2-socket Oracle Exadata system hardware with RoCE Network Fabric (X8M-2 or later platforms)

STORAGE PROVISIONING AND MANAGEMENT

- **Dynamic storage allocation:** Flexible provisioning capabilities allowing storage resources to be allocated and reallocated across multiple databases and users
- **Resource utilization optimization:** Improved storage efficiency through shared resource pools rather than dedicated storage allocations per database
- **Block storage services:** Sophisticated block volume management for arbitrary-sized raw storage volumes based on Exascale storage technology
- **Virtual machine storage:** Exascale block volumes store Exadata database server virtual machine images, eliminating local storage dependencies and enabling seamless VM migration

ADVANCED SNAPSHOT AND CLONING CAPABILITIES

- **Native database integration:** Pluggable Database (PDB) snapshot and cloning functions automatically leverage Exascale snapshots through standard SQL commands (CREATE PLUGGABLE DATABASE, ALTER PLUGGABLE DATABASE)
- **Space-efficient copies:** Snapshot and clone operations create space-efficient file copies based directly on Oracle Database files
- **Elimination of test master databases:** Native Exascale snapshots remove the requirement for traditional test master database infrastructure for development and testing scenarios
- **Copy-on-write technology:** Redirect-on-write mechanisms enable instant cloning with minimal storage overhead

NETWORK ARCHITECTURE

- **Exadata RDMA Network Fabric:** High-performance, low-latency interconnect utilizing Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE)
- **Memory-to-memory transfers:** Direct memory access capabilities reduce CPU overhead and achieve microsecond-level latencies
- **Bandwidth optimization:** Intelligent traffic routing and quality-of-service mechanisms ensuring consistent performance across shared infrastructure

EXADATA ON EXASCALE

The platform combines Oracle's existing distributed database capabilities with its Exascale serverless infrastructure. The architecture represents a convergence of Oracle's traditional engineered systems approach with cloud-native distributed computing principles.

DISTRIBUTED DATABASE ARCHITECTURE

Oracle's implementation utilizes a horizontal sharding architecture, where data partitioning occurs at the database level rather than the application level.

The system supports six distinct distribution strategies:

- **Hash-based sharding:** Uses consistent hashing algorithms to distribute data based on key values, ensuring even distribution across shards
- **Range partitioning:** Distributes data based on value ranges, optimal for time-series or sequential data patterns
- **List partitioning:** Assigns specific values to designated shards, useful for geographic or categorical data distribution
- **Composite partitioning:** Combines multiple distribution methods for complex data patterns
- **Value-based distribution:** Custom distribution logic based on specific business rules
- **Directory-based routing:** Maintains metadata catalogs for dynamic shard location resolution

The sharding layer operates transparently to applications through Oracle's database proxy architecture, which handles query routing, cross-shard joins, and distributed transaction coordination.

Each shard functions as a complete Oracle Database instance with full SQL DDL/DML support, stored procedures, triggers, and advanced features including partitioning, compression, and in-memory processing.

EXASCALE INFRASTRUCTURE LAYER

The underlying Exascale infrastructure employs a disaggregated architecture that separates compute, storage, and network resources.

Key technical characteristics include:

- **Compute elasticity:** Dynamic allocation of Elastic Compute Processing Units (ECPUs) with a minimum of 4 ECPU allocation per shard and maximum scaling to thousands of ECPUs

- **Storage architecture:** Database-aware intelligent storage with built-in compression, encryption, and smart scan capabilities for predicate pushdown and columnar processing
- **Network optimization:** RDMA-enabled high-bandwidth, low-latency interconnects between compute and storage layers
- **Resource isolation:** Multi-tenant architecture with performance isolation between workloads using hardware-level resource partitioning

SCALING & PERFORMANCE

The serverless model implements several advanced scaling mechanisms:

HORIZONTAL SCALING

- **Automatic shard splitting:** Dynamic data redistribution when individual shards exceed capacity thresholds
- **Elastic shard provisioning:** On-demand creation of new shards based on workload patterns and resource utilization
- **Cross-shard query optimization:** Distributed query execution with parallel processing across multiple shards
- **Connection pooling:** Intelligent connection management with automatic load balancing across available shards

VERTICAL SCALING

- **CPU elasticity:** Real-time ECPU allocation adjustments based on CPU utilization metrics
- **Memory scaling:** Dynamic buffer pool and PGA memory allocation with automatic tuning
- **Storage expansion:** Online storage capacity increases without downtime or data movement
- **I/O optimization:** Adaptive I/O scheduling and bandwidth allocation based on workload characteristics

AI & VECTOR PROCESSING

The platform integrates several AI-specific technical capabilities:

VECTOR SEARCH IMPLEMENTATION

- **Native vector data types:** Support for high-dimensional vectors (up to 65,536 dimensions) with optimized storage formats
- **Approximate nearest neighbor (ANN) algorithms:** Implementation of hierarchical navigable small world (HNSW) and inverted file (IVF) index structures
- **Hardware acceleration:** Utilization of specialized vector processing units and SIMD instructions for vector operations
- **Hybrid search capabilities:** Combined vector similarity and traditional SQL predicate filtering in single queries

AI SMART SCAN TECHNOLOGY

- **Predicate pushdown:** SQL predicate evaluation at the storage layer to reduce data movement
- **Columnar processing:** In-storage columnar operations for analytical workloads
- **Machine learning model execution:** Direct model inference execution within storage nodes
- **Feature engineering:** Built-in data transformation and feature extraction capabilities at the storage layer

QUERY PROCESSING & OPTIMIZATION

Advanced query optimization techniques specific to distributed architectures:

DISTRIBUTED QUERY EXECUTION

- **Cost-based optimization:** Extended cost models accounting for network latency and cross-shard data movement
- **Parallel execution:** Distributed parallel query execution with automatic degree of parallelism calculation
- **Bloom filter optimization:** Network-efficient join processing using Bloom filters for cross-shard operations
- **Materialized view synchronization:** Automatic maintenance of distributed materialized views across shards

TRANSACTION MANAGEMENT

- **Two-phase commit (2PC):** Distributed transaction coordination with optimizations for single-shard transactions
- **Saga pattern implementation:** Long-running transaction support for complex business processes
- **Read consistency:** Configurable isolation levels with snapshot isolation for read-heavy workloads
- **Deadlock detection:** Distributed deadlock detection and resolution across multiple shards

GOVERNANCE & COMPLIANCE

Technical implementation of data residency and compliance features:

POLICY-BASED DATA PLACEMENT

- **Metadata-driven routing:** Centralized metadata repository defining data placement rules and constraints
- **Geographic constraints:** Automatic enforcement of data locality requirements based on regulatory policies
- **Data classification:** Automated data sensitivity classification with corresponding placement policies
- **Audit trails:** Comprehensive logging of data access patterns and cross-border data movement

ENCRYPTION AND SECURITY

- **End-to-end encryption:** Data encryption in transit and at rest with customer-managed keys
- **Column-level security:** Fine-grained access controls with transparent data encryption (TDE)
- **Network isolation:** Virtual private cloud (VPC) integration with private endpoints and dedicated network paths
- **Identity integration:** Support for federated identity management and single sign-on (SSO) systems

AUTONOMOUS OPERATION

Technical details of the autonomous capabilities:

MACHINE LEARNING-BASED AUTOMATION

- **Workload analysis:** Continuous monitoring and pattern recognition for performance optimization
- **Index recommendation:** Automatic index creation and maintenance based on query patterns
- **Resource allocation:** Predictive scaling based on historical usage patterns and trend analysis
- **Anomaly detection:** Real-time identification of performance degradation and security threats

SELF-HEALING CAPABILITIES

- **Automatic failover:** Sub-three-second detection and recovery from node failures
- **Data corruption recovery:** Automatic detection and repair of data corruption using checksums and redundancy
- **Performance regression detection:** Continuous baseline monitoring with automatic tuning adjustments
- **Capacity planning:** Predictive capacity management with proactive resource provisioning

COMPETITIVE IMPACT & ADVICE TO IT BUYERS

THIS SECTION IS ONLY AVAILABLE TO NAND RESEARCH CLIENTS. PLEASE REACH OUT TO INFO@NAND-RESEARCH.COM TO LEARN MORE.

ANALYSIS

Oracle's converged architecture delivers significant advantages for organizations implementing AI-driven business processes. By integrating vector search capabilities directly within the transactional database, organizations eliminate the data movement overhead and consistency challenges associated with maintaining separate AI infrastructure. Real-time inference capabilities enable responsive AI applications while keeping the data security and consistency essential for regulated industries.

The ability to process streaming data at scale while maintaining transactional consistency enables real-time analytics and decision-making capabilities, providing competitive advantages

in fast-moving markets. Organizations can implement sophisticated AI workflows without the architectural complexity of managing multiple specialized databases and the associated data synchronization challenges.

Oracle's solution addresses the fundamental challenge facing large enterprises: the need for database infrastructure that can scale globally while maintaining the reliability, consistency, and compliance capabilities essential for mission-critical operations. The platform enables organizations to extend operations into new markets without the traditional barriers of establishing separate infrastructure in each jurisdiction.

The economic model, with its consumption-based pricing and elastic scaling capabilities, aligns infrastructure costs with business growth while providing the performance characteristics required for demanding workloads. Organizations can start with minimal footprints in new markets and scale resources as business requirements evolve, reducing the capital requirements traditionally associated with global expansion.

For organizations requiring the intersection of global scale, regulatory compliance, and mission-critical reliability, Oracle's platform offers capabilities that alternatives struggle to match comprehensively. While specialized distributed databases may excel in specific areas, few can deliver the combination of enterprise-grade features, proven reliability, and regulatory compliance automation that Oracle provides.

Bottom line: Oracle Globally Distributed Exadata Database on Exascale Infrastructure is essential for agentic AI because it provides the scalable, secure, and resilient data foundation needed to enable autonomous agents to operate across regions while helping to meet data residency requirements.

By combining Exascale's hyper-elastic scalability and extreme performance with the capabilities of the Globally Distributed Database, enterprises can seamlessly support dynamic, high-performance AI workloads on a global scale—all at a very attractive price point