

RISE OF THE ELASTIC ENTERPRISE DATABASE CLOUD

Oracle Exadata Cloud Service X8M

ORACLE HAS ANNOUNCED THE **EXADATA CLOUD SERVICE X8M**, a public cloud database service that provides a remarkable increase in the scale on which databases can be implemented in the cloud, supporting data volumes as large as 25 PB (petabytes¹).

The new service employs the *Exadata Database Machine X8M* to provide truly massive processor and I/O capacity with main memory, new persistent memory, RoCE² networking, and solid-state data storage to match. ●

ORACLE DESCRIBES THIS OFFERING as suitable for “any database workload”: transaction processing, database consolidation or data warehouse.

Massive Scale. Oracle says that the number of processor threads supported is 25 times larger than that supported by AWS RDS or AWS Aurora and that its data storage scales more than 20 times as far as RDS or Aurora. WinterCorp estimates that the processor, memory and flash storage capacity Oracle provides is also several times the published maximums for Snowflake and for AWS Redshift, two popular cloud data warehouse platforms. Note *Exadata Cloud Service X8M* is the only major public cloud service in which the entire hardware-software stack is engineered specifically for database efficiency. Multi-cloud database services such as Snowflake run on the general purpose hardware-software stack offered by the cloud provider, a situation in which the database vendor has little or no influence on the architecture or direction of the cloud system.

Persistent Memory & RoCE for High Performance. The *Exadata X8M* brings persistent memory (PMEM) and 100Gbit RDMA over Converged Ethernet (RoCE) networking to cloud databases for the first time. Persistent memory retains its contents when the system goes down, as in a shutdown or a power failure. In the *Exadata X8M*, PMEM functions as an additional tier of data storage that can be accessed via RoCE in **19 microseconds: ten times more rapidly than flash memory**. PMEM presently provides increased database performance, especially significant for transaction processing and rapid database lookups. As well as faster access time, PMEM also provides 2.5 times the IO rate (IO operations per second) of flash memory. Oracle says that it is now 50 times faster than AWS RDS, even when RDS is using all flash memory to store the data.

¹A petabyte (PB) is 1,000 terabytes (TB).

²RoCE stands for “remote data access over converged ethernet,” a capability in which a program on a server can retrieve data directly from memory on another server, including the new persistent memory, bypassing the usual I/O path. This results in a large reduction in overhead and access time.



Methodology



Purpose and Methodology for this Report

This *WinterCorp Research Note* covers the recent announcement of Oracle Exadata X8M Cloud Service and its implications for database customers. In developing this report, WinterCorp drew on its own independent research and experience, interviewed Oracle customers and employees, attended Oracle events and analyzed Oracle documentation and literature. Oracle was provided an opportunity to comment on the paper with respect to facts, in its capacity as the sponsor of this re-search. WinterCorp has final editorial control over the content of this publication and is solely responsible for any opinions expressed.



Configuration. An *Exadata X8M* storage server contains 49 TB of disk storage, 25 TB of flash memory and 1.5 TB of PMEM — and an Exadata system can be configured with up to 64 storage servers plus 32 database servers, yielding more than enough capacity for Oracle’s target of 2.5 PB of storage. Oracle says that this is typically equivalent to 25 PB of uncompressed data, enough to accommodate the public cloud database needs of large and data-intensive enterprises. The 10-to-1 compression factor results from the use of Hybrid Columnar Compression (HCC), an Oracle Exadata exclusive feature, resulting from the source-code level co-engineering between Oracle Database and Exadata.

Deployment Options. The newly announced *Exadata X8M* public cloud database service is the public twin of the service Oracle announced in July for Exadata *Cloud@Customer X8M* and Dedicated Region *Cloud@Customer*. The *Exadata Database Machine X8M* system capability has been available to Oracle on-premises customers since last November.

High Elasticity. Oracle’s database cloud services feature full, immediate elasticity and supports a wide range of system compute and storage elasticity options. Capacity increases and decreases are implemented immediately while the database continues to operate online and service users. This contrasts with some other cloud database services in which the database must be shut down in order to increase or decrease capacity.

With the *Exadata Cloud Service X8M*, customers can add database compute and storage servers — to support OLTP scaling in one dimension and data warehousing in another dimension, or both concurrently — to maximize their investment in a mixed workload cloud service.

Independent, Granular Scaling of Compute and Storage. Users can start running on a minimum configuration of 2 database (compute) servers and 3 storage servers. Each database server provides 50 cores and 1.3 TB of DRAM (main memory). Each storage server provides 49 TB of database storage, 25 TB of flash and 1.5 TB of PMEM. Customers can then step capacity up or down at any time, by increasing or reducing the number of servers of either type. The customer has complete flexibility in determining the number of storage servers and database servers and can thus independently scale storage and compute. Further, the scaling increment can be as small as a single server of either type. This is a more flexible capability, with far greater elasticity, than is available with typical cloud database services, in which the shape is fixed and the scaling options are limited to doubling or halving the number of servers. As a result, the idea of “pay only for what you use” has much greater impact.

As an analogy, consider a situation in which you are renting train cars to deliver goods to your customer. You had arranged for 8 cars. Then you learn that your customer wants to increase her order by 10%, so you need one more car. You would not be pleased if you went to your broker and were told, “You can’t get 9 cars; if you want more than 8, you have to pay for 16.” **Oracle lets you increase or decrease the size of the train by one car at a time, if that’s what you want to do.** Some others limit you to 2, 4, 8, 16, etc. So, what do they really mean when they say “pay only for what you use”? In our example, with the other database cloud service, you are paying for 7 train cars you can’t use. ●

With Exadata Cloud Service the scaling granularity and immediacy combined with *Exadata X8M*'s high performance can translate into large cost savings for customers.

Other Key Strengths. The Oracle *Exadata Cloud Service X8M* features other distinctive strengths. It provides a high degree of *user controlled* database automation for enterprise-level manageability, security and data availability. *Exadata Cloud Service X8M* servers are dedicated to a specific customer — meaning no noisy or malicious neighbor problems, which can be a problem in some other clouds.

Manageability. As an example, Oracle offers a wider range of database partitioning options than other popular cloud databases, and the customer can control their use with the *Exadata Cloud Service X8M*. The customer chooses the partitioning option at database design time and the Oracle software automatically organizes and maintains the data accordingly to increase performance and, where appropriate, manage data aging.

Security. Oracle's security and data protection features are the best in the industry. As well as all the usual controls and features, such as encryption and fine-grained access control, Oracle provides a strong capability for the separation of roles in administering the system and the database. Only customer personnel who have the need can see the contents of the data. By default, the system is set up the system so that no Oracle employee — not even a technical specialist maintaining the Oracle system in the cloud — can see the data. Similarly, the customer can prevent its own admins from accessing the data.

Exadata Cloud Service X8M has always-on encryption with customer-controlled keys to protect data at rest, in motion, and in backups. It also includes Oracle Data Safe which provides a unified security console for discovering sensitive data, monitoring security issues, and simplifying security administration across all Oracle Databases within an enterprise.

Cloud Automation. *Exadata Cloud Service X8M* uses built-in Oracle Cloud Infrastructure and Exadata automation to simplify the deployment, scaling, tuning, and patching of dedicated infrastructure for customers' service instances.

Exadata Cloud Service X8M automates database tuning with Automated Indexing and Exadata IO Resource Management. Data protection is automated with administrator-controlled scheduling of automated weekly

backups and daily incremental backups to Oracle Cloud Infrastructure Object Storage Service.

Converged Database. Reviewing the AWS website reveals that the company currently offers 16 different databases. Customers need different skills for each of these database systems. The need to use a different database for each of 16 different requirements results in data fragmentation, staffing, management and security challenges. The way that one performs many basic activities will vary from one of these AWS databases to the next, resulting in errors and wasted effort.

In a single product, Oracle Database, customers get support for structured relational data, time series data, graph data, spatial data, and documents. With all the data in the same database, Oracle customers can avoid ETL often required to move data from one type of database to another. For example, with AWS, to combine transactional data in Aurora and analytical data in Redshift, customers need to perform a series of complex steps: initiate a Lambda function, then a Kinesis function, then store the data into AWS S3, next load the data into Red-shift, and only then actually combine the data.

Layered Parallel Architecture. Exadata features unique storage servers that independently process scans, aggregations and other table- and set-level operations in a highly parallel fashion, applying intelligence close to the data in storage. These storage servers free up the database servers to perform more complex operations. This two-level architecture increases parallelism and IO bandwidth, delivering higher performance especially important on larger databases. It is this same architecture that now lets Oracle separate the scaling of compute and storage to accommodate a wide range of workloads.

Continuous Data Availability. *Exadata Cloud Service X8M* now also includes automated support for *Data Guard*, an Oracle service that maintains one or more active, up-to-date replicas of the database, optionally across availability domains or cloud regions. With *Data Guard*, the customer can deliver continuously available database services, even in the presence of events that take down the entire system or the data center, such as a power failure, terrorism or a natural disaster. The automatically maintained replicas can be used for secondary processing such as reporting or quality assurance. In addition, the replicas can be cloned for development and test without impact on the primary services.

About WinterCorp

WinterCorp is an independent consulting firm expert in the strategy, architecture and scalability of the modern analytic data ecosystem.

Since our founding in 1992, we have architected and engineered solutions to some of the toughest and most demanding analytic data challenges, worldwide.

We help customers define their data-related business interests; develop their data strategies and architectures; select their data platforms; and, engineer their solutions to optimize business value.

Our customers get business results with analytics in which their return is often ten or more times their investment.

When needed, we create and conduct benchmarks, proofs-of-concept, pilot programs and system engineering studies that help our clients manage profound technical risks, control costs and reach business goals.

With our in-depth knowledge and experience, we deliver unmatched insight into the issues that impede scalability and into the technologies and practices that enable business success.



WinterCorp

www.wintercorp.com

TYNGSBORO, MA

617-695-1800

Local high-availability for failover is provided with automation of Oracle Real Application Clusters, which also supports such services as applying patches and software upgrades, one database server at a time, while the system continues to run.

Customers sometimes assume that putting a database in the cloud eliminates the need for such capabilities — but this is a misconception. It is easy enough to store a second copy of a static file in a different region of the public cloud. But, it is an entirely different matter to have a remote copy of an entire database operation — including up-to-the-minute data continuously available for immediate failover at a remote location — such as is provided by Oracle *Data Guard*. A comparable capability is not provided on most other cloud database services.

Recommendation



CUSTOMERS WITH enterprise-class database requirements and an interest in cloud computing should take a close look at Oracle's Exadata Cloud Service X8M.

This is the first database cloud service with the elasticity, large capacity, high performance, comprehensive security and manageability needed to implement really massive, mission-critical database operations in the cloud. This service clearly goes beyond the limits of the widely promoted “cloud first” databases that are more familiar to many.

The same Exadata Cloud Service X8M is available now in the public cloud and on the customer's premises via various Oracle's *Cloud@Customer* deployment models.

While Exadata is a proven technology with thousands of implementations to recommend it, WinterCorp does urge customers with mission-critical database requirements — particularly at large-scale — to exercise due diligence in any such implementation. In selecting a new cloud platform, customers should always evaluate carefully and conduct realistic, scale-appropriate proofs-of-concept to assure themselves that they will hit any necessary milestones for cost, performance and operational control in the target environment. ●