

Transforming Data Management

With Oracle Database 12c Release 2

ORACLE WHITE PAPER | MARCH 2017





Disclaimer

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, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.



ORACLE®



Table of Contents	Disclaimer
	1
Transforming Data Management with Oracle Database 12c Release 2	2
From Disk-Based to In-Memory Databases	3
Distributed Real-Time Analytics	3
Memory-Optimized Performance	3
Simple To Implement	3
From Data Warehouses to Big Data	4
Oracle Big Data Platform	4
Fast SQL Access for Relational, Hadoop and NoSQL	4
More than Relational Data	5
Comprehensive Data Science Capabilities	5
Big Data Cloud Service	5
From On-Premises to Database-Optimized Cloud	6
Engineered Systems for Data Management	6
Lower Costs	6
Agility	7
Elastic Scaling	7
100% Compatible Hybrid Cloud	8
Oracle Cloud at Customer	8
Exadata Express Cloud Service	8
Spectrum of Oracle Database Cloud Services	8
Conclusion	9



Transforming Data Management with Oracle Database 12c Release 2

Oracle Database 12c Release 2 (12.2), the latest generation of the world's most popular database, is now available in the cloud, with Oracle Cloud at Customer, and on-premises. It provides businesses of all sizes with access to the world's fastest, most scalable and reliable database technology in a cost-effective and flexible hybrid cloud environment. 12.2 can seamlessly scale from entry-level to the largest database workloads, enabling customers to take advantage of proven enterprise class capabilities to run database development, small to mid-sized business applications, and departmental and large mission-critical applications in the Cloud and/or on-premises.

The Information Technology (IT) industry continues to introduce innovations that promise faster access to enterprise information, deliver more insightful analysis, and take advantage of the agility and flexibility of the Cloud. Data is the lifeblood of every business, and to help customers transform the access and value of enterprise data, 12.2 introduces new and enhanced functionality, particularly in these three areas:

- » **From Disk-Based to In-Memory Databases** for high performance real-time analytics
- » **From Data Warehouse to Big Data** for insightful analysis of transactional and other data sources
- » **From On-Premises to Database Optimized Cloud** for greater agility, elasticity and lower IT costs

With well over 40% market share, Oracle is by far the world's #1 database of choice for managing enterprise data. State-of-the-art database technology innovations such as Real Application Clusters, Active Data Guard and Database In-Memory, introduced over successive generations of Oracle Database, have helped over 300,000 customers deliver the best performance, reliability and security for their online transaction processing (OLTP) and data warehouse applications. Oracle Database 12c Release 1 (12.1), released in June 2013, with its innovative multitenant architecture, has proven to be the most widely adopted first release in recent memory. 12.2 continues to deliver on Oracle's long-standing strategy of delivering state-of-the-art innovations while preserving customers' investments in existing applications. This white paper discusses how 12.2 can help customers transform data management, without having to implement major changes to their business applications.

"With Oracle Database In-Memory, we saw our query performance go from minutes and hours to seconds, while dropping indexes reduced our database size by 80 percent."

FRANCOIS BERMOND

DATABASE MANAGER

SCHNEIDER ELECTRIC

From Disk-Based to In-Memory Databases

While relational databases were traditionally optimized for disk Input/Output, market demand for high performance real-time analytics drove the introduction of Oracle Database In-Memory with Oracle Database 12c Release 12.1.0.2 (in July 2014). Oracle Database In-Memory uses a unique dual-format architecture that simultaneously represents table data in its traditional row format, and in a pure in-memory column format. The column store is accessed through specialized software and hardware routines that amplify the performance of memory-resident data. Significant enhancements introduced in 12.2 improve automation and increase the flexibility and performance while maintaining complete compatibility with existing applications.

Distributed Real-Time Analytics

Oracle Database In-Memory enables real-time analytics by delivering performance similar to having an index on every column, but without the index overheads. Consequently, analytics that previously took hours or longer to run can now complete in seconds. Indeed, customers can eliminate secondary analytic indexes on OLTP databases that are no longer required for analytic query performance. New in 12.2 on Oracle Cloud and Oracle Exadata, customers can now maintain an in-memory column store on an Active Data Guard standby database. This enables customers to make more productive use of their standby databases. It completely isolates transactional users from analytic users, and still delivers high performance analytics against near-real-time data.

Memory-Optimized Performance

Oracle Database In-Memory's state-of-the-art algorithms for in-memory scans, joins and aggregations are widely proven to deliver orders-of-magnitude performance improvements on customers' OLTP, Data Warehouse and mixed workload environments. In addition to real-time analytics, other typical use-cases for Oracle Database In-Memory include hybrid transactional analytic processing, where analytics are an integral component of business transactions, as well as data warehouse query and reporting systems. Performance enhancements introduced with 12.2 include; up to 3X faster joins using in-memory join groups, up to 10X faster complex queries using in-memory expressions and up to 60X faster JSON queries using a new optimized binary format.

Simple To Implement

Unlike other in-memory column stores, Oracle Database In-Memory does not require the entire database to fit into main memory. It simply requires setting the size of the in-memory column store and identifying performance sensitive tables or partitions. New data population policies in 12.2 track in-memory usage through heat map technology and can apply policies that compress and evict less frequently used data from the column store. On Exadata storage, the in-memory column store is extended to flash in 12.2, dramatically enlarging the capacity. Most importantly, unlike other in-memory column stores, no changes to existing applications are required. The optimizer automatically routes analytic queries to the column format and OLTP queries to the row format. Oracle Database In-Memory is also 100% compatible with all the functionality of Oracle Database 12c, including Oracle Multitenant, Real Application Clusters, Maximum Availability Architecture, and database security features.

“The ability to get that answer, get that (big data) insight within minutes rather than days is quite impactful to a business.”

YAEL COSSET

CHIEF INFORMATION OFFICER

84:51

From Data Warehouses to Big Data

Just as OLTP applications are essential to processing business transactions, data warehousing applications are essential for measuring business performance. Data Warehouse best practices for analyzing transactional data have become fairly well established, and Oracle Database 12c provides a comprehensive platform that leads the data warehouse market by a wide margin. However, new data sources such as web logs, social media and Internet of Things (IoT) data offer the prospect of deeper insights into business performance and opportunities. Technologies such as Hadoop and NoSQL data stores are ideally suited to acquiring high volumes of low-density data, but analyzing independent data stores with different interfaces can be challenging. The real business value for organizations lies in analyzing all data sources by transforming their data warehouses to an integrated big data platform capable of data mining, machine learning and business analytics. Organizations should also consider their deployment choices as they transform from data warehouses to big data, be that on-premises, on the Cloud or a hybrid-cloud environment.

Oracle Big Data Platform

Oracle recognizes that enterprise data may be stored in disparate data stores (relational, Hadoop, NoSQL), on different platforms (general purpose hardware, engineered systems), and in various locations (on-premises, on-Cloud). In order to help customers transform their traditional data warehouses and embrace the opportunity of big data, Oracle has developed a Big Data Platform that offers integrated access to data stored in Oracle Database 12c, Hadoop and NoSQL. The Oracle Big Data Platform can run on general purpose or engineered systems, be deployed on-premises and/or Oracle Cloud and be accessed using a familiar SQL interface and familiar development and analytics tools. It effectively eliminates the need to move large volumes of data between disparate data stores, and customers can easily perform different types of analysis (e.g. Machine Learning, Graph, Spark), using different languages (e.g. SQL, REST, R), against different types of data (e.g. relational, XML, JSON), stored in different repositories (e.g. Oracle Database 12c, Hadoop, NoSQL).

Fast SQL Access for Relational, Hadoop and NoSQL

Oracle Big Data SQL is the data-virtualization component of Oracle's Big Data Platform. It enables customers to use Oracle SQL for querying and analyzing data across Hadoop, NoSQL and Oracle Database 12c, using their existing SQL tools, resources and skills. Big Data SQL delivers high-performance queries using Oracle's Smart Scan capability, first developed on Oracle Exadata, to execute SQL operations such as query filtering, joins and scoring, on Hadoop and/or NoSQL servers. Other key performance features of Big Data SQL include massively parallel, distributed query processing and storage indexing. Oracle Big Data SQL also provides centralized metadata (via external tables) for simple access to data regardless where it is stored. Based upon this metadata, organizations can implement standard security policies, and apply Oracle Database 12c security features such as data redaction and access controls, across data stored in Hadoop and NoSQL data stores.



More than Relational Data

Oracle Database 12c Release 12.1.0.2 introduced support for JSON data, adding to the broad support of non-relational data such as XML, text, spatial, and graph. As a common data representation for IoT data sources, JSON is widely used in Big Data environment, and Big Data SQL can seamlessly access JSON via Oracle Database 12c's SQL extensions for JSON. With 12.2, Oracle continues to enhance its JSON capabilities by providing a JSON data guide feature that helps customers easily understand the structure of JSON documents using SQL commands.

Comprehensive Data Science Capabilities

The transformation from data warehouses to big data also provides the opportunity to expand beyond SQL-based analytics to integrating machine-learning capabilities, and introducing new domain-specific analytics for non-relational data sources. Oracle Database 12c offers developers and data scientists a choice of in-database analytics and tools that can provide more in-depth business analysis. For example:

- » **Machine Learning** with massively scalable R processing and in-database and Spark algorithms that extend and enhance SparkML.
- » **Property Graph** with over 40 in-memory parallel algorithms that enable 12.2 to be used as a Graph database using simple standard interfaces.
- » **Spatial** with over 50 functions for massively scalable Vector and Raster processing that enable seamless integration of spatial data with analytic and other applications.
- » **Multi-Media** with a massively scalable open framework for imaging and video processing commonly used in facial, OCR, and License Plate recognition applications.

Successive generations of Oracle Database have introduced support for multiple data types (e.g. text, spatial, XML, JSON) and more in-database analytics (e.g. Pattern Matching, Graph), offering customers multi-model persistence within a single Oracle Database. This approach can eliminate siloed data stores. Alternatively, the Oracle Big Data Platform offers customers integrated access across multiple data stores using Oracle Big Data SQL. Regardless of whether customers choose a multi-model or polyglot strategy (or both), Oracle's philosophy of moving the analytics to the data can help customers quickly sift through large volumes of Big Data using familiar tools.

Big Data Cloud Service

By integrating software and hardware components together at the factory with engineered systems such as Oracle Exadata and Big Data Appliance, Oracle helps organizations eliminate risk, optimize performance and speed time-to-implementation for their big data projects on-premises. Oracle brings these same benefits to the cloud with the Big Data Cloud Service. It offers customers a comprehensive, high-performance service for Hadoop, Spark, and NoSQL and includes; Cloudera Enterprise Data Hub, R and Property Graph analytics, and data integration tools. Customers can start small with a 3-node cluster and easily scale out to 100's of nodes as required.

Oracle Cloud not only delivers all of the capabilities of Oracle's on-premises solution, it also transforms Big Data deployments from on-premises upfront capital expenditure to a pay-as-you-go operational expenditure model.. Oracle's hybrid strategy of delivering the same architecture and software on Oracle Cloud as on-premises enables organizations to move to the cloud, while maintaining their existing skillsets, applications, and support resources for the Oracle Big Data Platform. Customers have the choice of deploying Oracle Exadata and Big Data Appliance for their data warehouse and big data systems on-premises, or consuming Exadata Cloud and Big Data Services in the Oracle Cloud.

“Oracle Database 12c Release 2 has been optimized for the Cloud.”

LARRY ELLISON

CHIEF TECHNOLOGY OFFICER

ORACLE

From On-Premises to Database-Optimized Cloud

Designed for the Cloud, 12.2 enables customers to lower IT costs, become more agile in provisioning database services, and have the flexibility to elastically scale up, scale out and scale down IT resources as required. The availability of 12.2 'cloud first' complements Oracle Database Cloud Services already in place, and most importantly, it offers some unique capabilities that can help transition from on-premises to the database-optimized Oracle Cloud.

Engineered Systems for Data Management

Many customers have already taken the initial steps on their journey to the cloud by standardizing and consolidating database workloads onto Engineered Systems like Oracle Exadata that are widely proven (on-premises) to deliver extreme performance for OLTP, Data Warehousing and mixed database workloads. Oracle Exadata is a preconfigured package of servers, storage and networking, and software that includes Oracle Database 12c and utilizes key database components such as Real Application Clusters and Automatic Storage Management. It also features intelligent Oracle Exadata Storage Server Software innovations including Exadata Smart Scans, Smart Flash Cache and Hybrid Columnar Compression that enable extreme database performance and greater database consolidation density. Engineered Systems like Oracle Exadata have been specifically optimized for Oracle database workloads and now customers can take advantage of 12.2 optimizations with the Exadata Cloud Service on Oracle Cloud or with Oracle Cloud at Customer.

Lower Costs

The transformation from on-premises to database-optimized cloud significantly changes traditional IT management and cost models. Instead of sizing database systems for peak capacity and incurring up-front capital expenditures, customers can consume database services on a pay-as-you-go self-service basis and manage many databases as one in standardized configurations.

Oracle Database 12c was specifically re-architected to help customers take advantage of the Cloud more easily. Using Oracle Database 12c's multitenant architecture, many customers have consolidated multiple pluggable databases (PDBs) into a single multitenant container database (CDB) without having to change any application code. Administrators can back up and recover, patch and upgrade many PDBs, all as operations on a single CDB. Customers have therefore been able to lower capital and operational expenditures by consolidating many PDBs into a single CDB and effectively managing many databases as one.

12.2 offers customers even greater database consolidation density, with up to 4096 PDBs per single CDB on Oracle Cloud. It also offers greater isolation at scale with processor, I/O and memory prioritization of PDBs, and lockdown profiles for configurable isolation between private and public Clouds. Also new in 12.2 are Multitenant Application Containers that enable many PDBs to share application objects such as code, meta data and data. By sharing application objects across many PDBs, customers can efficiently manage many application tenants as one in a single Application Container. This capability is ideally suited for SaaS, franchise and other applications that are typically shared across different constituents, but require secure isolation of each constituent's data.

Agility

The ability to be more agile in response to user demands for database services and the automation of routine, but nonetheless essential database operations (e.g. patching, upgrading, etc.) are core attributes of the Cloud. Oracle Database 12c delivered even greater agility by helping customers transition from managing many single tenant databases to managing fewer multitenant container databases. Oracle Multitenant in 12.1 also provided customers with rapid PDB provisioning, cloning and movement on-premises, on Oracle Cloud and hybrid Cloud environments. The agility features of Oracle Multitenant are significantly enhanced with new online PDB operations in 12.2:

- » **PDB Hot Clone** enables fast instantiation of PDBs without having to take users offline
- » **PDB Refresh** enables cloned PDBs to be regularly updated with latest data
- » **PDB Relocate** enables PDBs to be relocated between CDBs with near zero downtime

Oracle Multitenant with 12.2 uniquely simplifies and speeds up the process of copying databases, synchronizing databases and moving databases for development, testing and deployment purposes, without impacting other ongoing database workloads and activities.

Elastic Scaling

Oracle Database is renowned for its ability to scale up on single servers and scale out on clustered servers with Oracle Real Application Clusters (RAC). Transitioning to the Cloud means customers can now provision database capacity on demand instead of provisioning for peak workloads, and continue to scale up and scale out and 'burst' capacity using RAC as required. RAC in 12.2 offers customers greater reliability and scalability with optimizations for multitenant databases, and the ability to scale to hundreds of RAC nodes.

Also new in 12.2 is native database sharding for massive scalability and reliability for OLTP applications. Oracle Database Maximum Availability (MAA) features such as RAC and Active Data Guard are widely proven to meet the needs of over 99% of OLTP applications while preserving application transparency. However, some global-scale OLTP applications prefer to shard massive databases into a farm of smaller databases for scalability and reliability purposes. This requires designing applications so that workloads are automatically routed to (up to 1,000) specific shards in a database farm. Oracle Sharding is 100% complementary to Oracle MAA.

Oracle Cloud uniquely offers customers the choice of deploying database services on general-purpose hardware or engineered systems such as Oracle Exadata. The Oracle Exadata Cloud Service runs on the only cloud infrastructure optimized for Oracle Database 12c, providing customers Oracle Exadata benefits of such as scale out compute, scale out intelligent storage and elastic expansion, and Oracle Cloud benefits such as managed infrastructure, rapid, elastic, database provisioning and pay-as-you-go subscription models. The pay-as-you-go Exadata Cloud Service features monthly subscriptions for regular CPU requirements and the unique ability to 'burst' online processing capacity a short period on a hourly basis as required (e.g. to cope with short term spikes in demand). This capability eliminates the need for customers to over-provision (or over-subscribe) capacity to meet peak workloads.



100% Compatible Hybrid Cloud

Providing customers choice, with a range of database services on Oracle Cloud compatible with on-premises environments, is part Oracle's hybrid cloud strategy, thereby helping customers with their transition to Oracle Cloud. A common architecture and software environment between on-premises and Oracle Cloud facilitates the movement of database workloads, enabling customers to utilize existing skills and resources, and reduce disruptions to their business. This hybrid strategy is also reflected with the expansion of Oracle Cloud at Customer and the introduction of Exadata Express Cloud Service on Oracle Cloud.

Oracle Cloud at Customer

For those customers reticent to adopt public Cloud services for security, compliance or other reasons, Oracle Cloud at Customer enables customers to consume database cloud services behind their firewalls. It provides the exact same IaaS and PaaS hardware and software available on Oracle Cloud, but deployed at the customer's site, and managed as a service by Oracle. Customers can then consume the same database services available on Oracle Cloud using the exact same subscription-pricing model, securely behind their firewall. Database Cloud Services can currently be consumed on Oracle Cloud Machine(s) and Exadata Cloud Machine(s), and customers will soon be able to consume Big Data Cloud Services in the same manner (behind their firewall) with the pending availability of the Oracle Big Data Cloud Machine.

Exadata Express Cloud Service

12.2 first became available on Oracle Cloud with the introduction of a new entry-level Exadata Express Service. This new service starts at only \$175 per month, and includes enterprise features and options such as Partitioning and Advanced Compression. It provides customers the latest generation of the world's #1 database running on Oracle Exadata as a managed service. Use cases for Exadata Express Service include; application development and testing, short-term time sensitive projects, analytics and sandboxes and production workloads. Subscribers are provisioned an Oracle database with integrated development tools, Application Express and SQL Developer. Exadata Express Service offers full support for JSON and REST, and a variety of development frameworks including Java, .NET, PHP and NodeJS.

Spectrum of Oracle Database Cloud Services

The introduction of new Exadata Express Service expands the spectrum of 12.2 database services available on Oracle Cloud, providing customers with a complete choice of database services that can easily scale from entry-level development databases through small and medium sized business and departmental applications to the largest mission critical database workloads. In addition, customers can consume Oracle Cloud Database Services on as subscribed services on their own premises with Oracle Cloud at Customer. Regardless of choice of database service, the compatible nature of Oracle Cloud ensures that customers can easily scale up or down or between on-premises and Oracle Cloud as required.



Conclusion

Technology innovations in cloud, big data, and in-memory undoubtedly offer opportunities to improve efficiencies and drive businesses forward. The challenge for every IT organization lies in transforming data management as seamlessly as possible, without incurring major disruption to ongoing business operations. As the database market leader, Oracle has a well-proven track record of helping customers evolve their databases through the mainframe and mini, client/server and internet generations. That same well-proven strategy of delivering state-of-the-art database innovations while preserving customer's investments applies today.

Designed for the Cloud, Oracle Database 12c Release 2 offers customers the opportunity to take full advantage of In-Memory for real-time analytics, Big Data sources for greater insight and Multitenant to lower costs and increase agility, without having to make changes to existing applications. Customers can take all their in-house developed Oracle applications, all their Oracle ISV applications, and just run them on Oracle Cloud without incurring any application changes. This is particularly important, as customers can utilize all their existing applications, developer skills, and administrator skills, and get the same support for their Oracle Databases both on-premises and on Oracle Cloud.

Finally, it's important to appreciate that all Clouds are not equal. Only Oracle Cloud offers infrastructure optimized for Oracle Database 12c, only Oracle Cloud offers critical database features like Real Application Clusters and Active Data Guard as an integral part of database services, and only Oracle Cloud offers 100% hybrid compatibility with customer's Oracle Database workloads running on-premises.







Oracle Corporation, World Headquarters

500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries

Phone: +1.650.506.7000
Fax: +1.650.506.7200

CONNECT WITH US

-  blogs.oracle.com/oracle
-  facebook.com/oracle
-  twitter.com/oracle
-  oracle.com

Integrated Cloud Applications & Platform Services

Copyright © 2016, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 0116

Transforming Data Management with Oracle Database 12c Release 2
March 2017
Author: Penny Avril, Willie Hardie
Contributing Authors: George Lumpkin, Tim Shetler, Patrick Wheeler