



An Oracle White Paper
June 2013

Plug into the Cloud with Oracle Database 12c

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.

Introducing Oracle Database 12c	1
Enabling Database Clouds	2
Simplify Database Consolidation	3
New Multitenant Architecture	3
Manage Many Databases As One	4
Provisioning and Cloning	5
Database Resource Management	5
Data Isolation	5
Managing Data Growth	6
Automatic Data Optimization	7
Maximum Availability Architecture	8
Addressing Unforeseen Failures.....	8
Reducing Planned Downtime.....	9
Further Enhancing Availability and Data Protection.....	10
Defense-In-Depth for Maximum Security	11
New Preventative, Detective and Administration Controls.....	12
Simplifying Analysis of Big Data	13
Integrating Big Data	13
In-Database Analytics	14
Big Data Infrastructure	15
Managing Database Clouds.....	16
Database as a Service	16
Reducing Risks of Getting to the Cloud	16
Database Application Development.....	17
Application Developer Frameworks	17
Oracle Application Express	17
Oracle SQL Developer	17
Migrating to Oracle Database 12c	18
SQL Translation Framework	18
Plug into the Cloud with Oracle Database 12c	19

Introducing Oracle Database 12c

Oracle Database 12c, the latest generation of the world's most popular database, is designed for the cloud and will enable customers to make more efficient use of their IT resources while continuing to improve their users' service levels. Based on a new multitenant architecture, and including a raft of enhancements and new features, Oracle Database 12c makes it easier for customers to take advantage of the cloud through:

- Consolidating multiple databases into multitenant containers
- Automatically optimizing data storage and compression according to usage patterns
- Providing continuous access with Oracle's Maximum Availability Architecture
- Securing enterprise data with a comprehensive defense-in-depth strategy
- Simplifying in-database analysis of Big Data
- Implementing Database as a Service using Enterprise Manager Cloud Control

Oracle Database 12c benefits customers looking to reduce IT complexity and cost through deployment of private database clouds, and SaaS vendors looking for the power of Oracle in a secure multitenant model. This white paper highlights some of the new capabilities in Oracle Database 12c that can help customers quickly plug into the cloud, as well as other enhancements in storage management, availability, security, big data support and database management that can help customers continue to get a high return on their Oracle database investments.

“Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.”

National Institute of Standards and Technology's (NIST) definition of cloud computing

Enabling Database Clouds

Industry research reveals that many IT organizations have already taken the first few steps on their journey towards the cloud. This isn't too surprising, as cloud computing offers an opportunity for IT organizations to be more responsive to changes in application workloads and business demands, while also reducing IT complexity and costs. Indeed, successive Oracle releases have helped customers standardize, consolidate and automate database services on the cloud with a variety of innovations including:

- **Oracle Real Application Clusters** supports the transparent deployment of databases across pools of servers, providing fault tolerance from hardware failures or planned outages. It popularized grid computing for database workloads over 10 years ago making it ideally suited as the foundation for database clouds, delivering reliable and efficient scaling-out of online transaction processing (OLTP) and data warehouse applications.
- **Oracle Automatic Storage Management** provides file system and volume management functionality for Oracle Database files. It simplifies storage consolidation and provisioning for database clouds, complete with innovative rebalancing capability for an even distribution of I/O and protection from disk failure with automatic mirror reconstruction and resynchronization.
- **Oracle Exadata Database Machine** delivers extreme database performance for OLTP, Data Warehousing and mixed database workloads. It's a preconfigured package of software, servers, and storage that features Oracle Real Application Clusters, Oracle Automatic Storage management, and intelligent Oracle Exadata Storage Software for consolidating databases on the cloud with the simplicity of an engineered system.

Oracle Database 12c introduces a new multitenant architecture that makes it easy to deploy and manage database clouds and enables customers to take full advantage of the resource sharing, management flexibility and cost savings that cloud computing offers.

"Pluggable databases allows us to consolidate hundreds of databases onto a RAC environment that guarantee the separation that drove us to put them on separate servers previously."

Martin Power, Logical Technology

Simplify Database Consolidation

Standardizing on fewer moving parts in the data center helps to maximize the benefit of consolidation, and consolidating databases is a key step in the journey to the cloud. It provides the opportunity to make more efficient use of available hardware and administrative resources. However, it can be a challenge for IT organizations to achieve high levels of database consolidation density without incurring major development and administrative overheads. Many IT organizations have used server virtualization, run several Oracle instances on a server, or combined Oracle databases through schema consolidation. Each of these approaches have limited value or can lead to increased costs and complexity.

New Multitenant Architecture

Designed for the cloud, Oracle Multitenant delivers a new architecture that simplifies consolidation and delivers the high density of schema based consolidation, but without requiring changes to existing applications. It's an option of Oracle Database 12c Enterprise Edition that offers all the benefits of managing many databases as one, yet retains the isolation and resource control of separate databases. In this new architecture, a single multitenant container database can host many 'pluggable' databases. Each database consolidated or 'plugged in' to a multitenant container looks and feels to applications the same as for existing Oracle databases. Accessing pluggable databases is the same as for existing Oracle databases, and administrators can control the prioritization of available resources between consolidated databases.

Upgrading to Oracle Database 12c

Customers have the choice of upgrading to Oracle Database 12c's new multitenant architecture or upgrading with the current architecture. Upgrading existing Oracle databases and plugging them into a multitenant container is very straightforward. Customers have a choice of upgrade paths and tools, depending on what database release they are currently on, among other considerations. There are direct upgrade paths from Oracle Database 11g Release 1 and Release 2, and Oracle Database 10g Release 2. Once upgrade processes are complete, customers simply 'plug in' upgraded databases into a multitenant container. Customers on Oracle Database 10g Release 1 or earlier releases can utilize tools such as Oracle GoldenGate or Data Pump to easily migrate data to Oracle Database 12c.

Manage Many Databases As One

Consolidating multiple databases means that administrators have fewer databases to manage and the benefits of managing many databases as one can permeate throughout the data center. It is fairly common for IT organizations to manage hundreds, if not thousands of databases, and each of those databases - be they for production, test or development - requires maintenance.

Fewer Patches and Upgrades

Applying patches, patch set updates and product updates to multiple (non-consolidated) databases in order to maintain currency of database releases can be a challenge that only gets more difficult with the number of databases managed. All databases in the data center, regardless of their production, test or development status, are subject to patching and upgrading. Consolidating multiple pluggable databases into a multitenant container can dramatically reduce the amount of patching and upgrade activities required. With Oracle Databases 12c, patches and upgrades are applied at the container level, and not applied to individual pluggable databases. Customers also have the flexibility of creating new patched and/or upgraded container databases and selectively unplugging databases from older container databases and plugging into new upgraded container databases.

Fewer Backups

Every good database administrator regularly takes backup copies of all databases managed in the data center. However, instead of having to backup each separate database in the data center, the new pluggable architecture only requires executing backups at the multitenant container level. This means that database administrators have fewer backups to manage, as each pluggable database in a container database is automatically backed up. However, in the event of any problem encountered with an individual pluggable database, database administrators have the flexibility to perform data recovery operations at the pluggable database level within a container database.

Fewer Standby Databases

Another good example of the benefits of managing many pluggable databases as one is maintaining standby databases. Many Oracle customers use replication technologies such as Oracle Data Guard or Oracle Active Data Guard to maintain complete working copies of production databases by shipping and applying redo logs to a standby system usually in a separate location. With Oracle Database 12c, shipping and applying of redo logs occurs at the container level, therefore every pluggable database consolidated into a container is fully protected from site outages. In addition, customers using Oracle Active Data Guard can make active use of their standby pluggable databases for reporting and other purposes.

Provisioning and Cloning

Rapid provisioning and cloning of databases for various purposes including testing, development and problem diagnosis can be a challenge for many IT organizations. Database administrators typically devote a significant portion of their working days to creating new databases, cloning databases and moving databases between different servers. In addition to simplifying database consolidation, Oracle Multitenant also enables rapid database provisioning and cloning. For example, database administrators can easily copy production databases and plug them into development and test containers. In addition, if the underlying filesystem supports copy on writes (e.g. ZFS Filesystem, ASM Cluster Filesystem), cloning of pluggable databases can occur almost instantaneously.

Database Resource Management

While there are undoubtedly many benefits from consolidating multiple databases, customers should consider all aspects of database consolidation. An obvious question to ask is “how can I guarantee the prioritization of resources in this new multitenant architecture?” After all, many of the benefits of managing many database as one will be quickly lost if database applications are continually vying for available system resources, and user performance service levels start to drop off.

Using Oracle Database 12c’s resource management features, database administrators can easily define resource utilization priorities at the pluggable database level. The database server pro-actively monitors database usage to ensure that each pluggable database in a container stays within pre-defined minimum and maximum resource thresholds. For example, at quarter-end, a high priority ERP application can automatically pull system resources from low priority applications in the container to ensure users performance service level remain consistent even at peak demand.

Data Isolation

Customers who have invested development resources into schema consolidation projects have undoubtedly been able to achieve high levels of consolidation density, but also encountered new challenges in terms of weak security inherent with schema consolidation.

Oracle Multitenant solves this by retaining the isolation and data security of separate databases. Each database consolidated, or plugged, into a multitenant container not only manages the meta-data and data for that particular database, it also retains all the security access privileges from the previous stand-alone database. While application users retain their access privileges from a stand-alone database to a pluggable database, these privileges will not enable access to other pluggable databases in the container. This aspect will be particularly useful for application SaaS providers by providing multi-tenancy in the database rather than the application tier.

"The automated compression features of Oracle Database 12c make it much easier to automatically age out older data and store it on different storage compression level tiers."

Mirko Hotzy, Trivadis

Managing Data Growth

As the volume of data that organizations manage continues to accumulate and grow, IT managers are faced with conflicting objectives of ensuring good database performance and scalability while keeping their database storage costs down. Oracle offers a number of database technologies that can help customers ensure fast performance while keeping down the cost of storing and managing data throughout its lifecycle. These include:

- **Oracle Automatic Storage Management** to automate the striping and mirroring of database files for high performance and resilience from disk failures.
- **Oracle Partitioning** to help administrators 'divide and conquer' databases by partitioning large tables and indexes into smaller, more manageable database objects.
- **Oracle Advanced Compression** to reduce storage requirements of large OLTP tables or partitions containing 'hot' frequently updated rows by 2-4x, and improve query performance.
- **Oracle Hybrid Columnar Compression** which provides up to a 10x compression ratio for 'warm' read-only data and up to 50x compression for 'cold' archival data.

By combining these database technologies with different storage tiers, administrators can manually implement Information Lifecycle Management policies to ensure best performance and keep their storage costs down. For example, a high performance storage tier could be used for 'hot' OLTP tables or partitions, and a separate lower cost, high capacity storage tier for 'warm' data warehouse and 'cold' archive table and partitions using appropriate compression techniques.

Traditionally, many database administrators have successfully implemented Information Lifecycle Management policies by manually moving and compressing data across different storage tiers based mainly on the age of data stored in tables and partitions. Oracle Database 12c changes the game for administrators with Automatic Data Optimization to monitor the usage or 'temperature' of data, as well as age of data, to automatically move data between different storage and compression tiers.

Automatic Data Optimization

The ‘temperature’ of rows stored in database tables and partitions changes over time. For example, rows inserted from OLTP applications typically start out ‘hot’ as multiple inserts and updates, but over time as transactions are processed these ‘hot’ rows cool down and become ‘warm’ used for business intelligence and other read-only purposes. Over longer periods of time, ‘warm’ rows become ‘cold’ and are rarely accessed, but still require to be easily accessible for reporting or compliance purposes. However, it’s not just the age of data that’s important, it’s also the activity of data. It’s not uncommon for rows to be continuously updated over time, therefore a combination of age and activity is required to determine the ‘temperature’ of table rows. New Automatic Data Optimization features in Oracle Database 12c can be used to implement an automated Information Lifecycle Management strategy using a Heat Map and server managed storage policies that enable smart compression and storage tiering.

Heat Map

The new Heat Map feature in Oracle Database 12c tracks usage information at the row and segment levels. These statistics are automatically maintained, and this enables database administrators to easily gauge the ‘temperature’ of their databases to see at a glance how access patterns change over time and also over different storage tiers. Once database administrators can better understand how their data is being used, the next logical step is defining and applying policies to automatically move and compress database objects based on the age and activity of data.

Smart Compression and Storage Tiering

Using Oracle Database 12c, database administrators can now create declarative policies that use statistics collected by the Heat Map to define the relevant operations to execute for data compression and movement between storage tiers. For example, data that hasn’t been updated for ‘x’ days could be moved from a high performance storage tier to a lower cost storage tier or different compression tier or combination of both. Similarly, ‘warm’ read-only data that hasn’t been modified or read for ‘y’ months could be moved to an archive compression storage tier. Policy based storage management will enable administrators to automatically move data from row format for OLTP applications into columnar format which offers major performance benefits for Business Intelligence and Analytical applications. Declarative policies are specified at the table level, and database administrators have the choice of executing automatically in the background or executing on demand.

“What application continuity brings to applications now, is that they can run in a clustered environment with the security knowing that the application continuity capabilities in Oracle Database 12c are going to automatically handle a lot of failure scenarios automatically.”

Marc Fielding, Pythian

Maximum Availability Architecture

Organizations today are dependent on IT to run efficient operations, quickly analyze information and compete more effectively. Consequently, it is essential that their IT infrastructure and databases in particular, are continuously available; otherwise the cost of downtime will be measured in lost revenues, lost productivity and loss of customers and reputation. Basic high availability architectures using redundant resources can prove costly and fall short of availability service level expectations due to technological limitations, and complex integration, and inability to offer availability through planned maintenance. Oracle Database 12c goes beyond the limitations of basic high availability with a Maximum Availability Architecture that offers customers an integrated product set and best practice blueprints to address the common causes on unforeseen and planned downtime that can be deployed at minimal cost.

Addressing Unforeseen Failures

Protection from unplanned downtime requires architecture built with redundant components. Extra disks are required for data mirroring, additional hardware is required for failover server processing, and additional data centers are required for disaster recovery. All this redundancy is very expensive, and typically only realizes any value when there is indeed a component failure; this is rather like an expensive insurance policy. In addition, a great deal of software is often required to integrate these different components together, typically from different vendors. This introduces additional complexity and greater risk for human error

IT infrastructures are prone to failures such as server faults, disk crashes or storage corruption, site outages and human error that can incur unplanned downtime. In order to prevent and mitigate unplanned downtime Oracle Database 12c offers a number of high availability components including:

- **Oracle Real Application Clusters** to protect from database server failures
- **Oracle Automatic Storage Management** to protect from database storage failures.
- **Oracle Recovery Manager (RMAN)** to manage database backups
- **Oracle Secure Backup** to manage tape backups of database and file system data
- **Oracle Flashback** to protect from human error
- **Oracle Active Data Guard** to protect from site failures and offload processing

Reducing Planned Downtime

Planned downtime for essential maintenance such as hardware upgrades, software upgrades and patching are part and parcel of every IT operation. Oracle Database 12c offers a number of solutions to help customers reduce the amount of planned downtime required for maintenance activities, including:

- **Hardware Maintenance and Migration Operations** to Oracle Database 12c infrastructure can be performed without taking users offline. Using Automatic Storage Management, disks can be added or removed online and the data is automatically rebalanced. Database servers can also be easily added or removed to a clustered database infrastructure while users remain connected. Oracle Database 12c also introduces cross platform incremental backup and recovery to minimize downtime for cross-platform hardware migrations. In addition, pluggable databases can minimize the disruption of hardware migrations via fast unplug and plug of databases between multitenant containers on different servers.
- **Online Patching** of database software can be applied to server nodes in a ‘rolling’ manner using Oracle Real Application Clusters. Users are simply migrated from one server to another; the server is quiesced from the cluster, patched, and then put back online. The same operation is then repeated for every server in the cluster. Also, for many one-off fixes, the patch can be applied directly to an Oracle instance without shutting the instance down.
- **Rolling Database Upgrades** using Oracle Data Guard or Oracle Active Data Guard enables upgrading of a standby database, testing of the new (upgraded) environment and then switching users to the new environment, without any downtime.
- **Online Redefinition** can reduce maintenance downtime by allowing changes to a table structure while continuing to support an online production system. Administrators can enable end users to issue insert, update and delete operations against tables that are undergoing structural changes. Oracle Database 12c also introduces the ability to move data files (useful for migrating storage) and partitions (useful for enabling compression) while users continue to access underlying data.
- **Edition Based Redefinition** enables online application upgrades. With edition-based redefinition, changes to program code can be made in the privacy of a new edition within the database, separated from the current production edition. An editioning view exposes different projections of the same table into each edition, ensuring that the code in each edition only sees its own specific view of the table. Cross edition triggers propagate the data changes made by the old production edition into the new edition’s columns, and vice-versa. This then allows both the old production environment and the new production environment to be used at the same time, for testing, and allows users to be moved online from one edition to the other.

Further Enhancing Availability and Data Protection

Oracle's Maximum Availability Architecture builds on successive releases of database innovations that are proven to deliver on customer's availability services level expectations, and reduce both IT costs and the cost of downtime. Oracle Database 12c continues to evolve Maximum Availability Architectures with new and enhanced availability and data protection features that help customers maximize their database availability by addressing unforeseen failures, reducing planned downtime and providing continuous application availability.

Data Guard Far Sync

New in Oracle Database 12c, Data Guard Far Sync provides zero data loss protection for a production database by maintaining a synchronized standby database located at any distance from the primary location, without impacting database performance and with minimal cost or complexity. A far sync instance receives changes synchronously from a primary database and forwards them asynchronously to a remote standby. It's a light-weight entity that manages only a control file and log files, and only requires a fraction of the CPU, memory, and I/O resources of a standby database to relieve a primary database from serving remote destinations. Production can be quickly failed over, manually or automatically, to the remote standby database with zero data loss.

Global Data Services

Also new in Oracle Database 12c, Global Data Services provides inter-region and intra-region load balancing across Active Data Guard and Golden Gate replicated databases. It effectively provides Real Application Cluster failover and load balancing capabilities to Active Data Guard and Golden Gate distributed databases. Global Data Services extends the familiar notion of Database Services to span multiple database instances in near and far locations and can be used to distribute workloads across a reader farm composed of standby databases.

Continuous Application Availability

A major challenge for many organizations is how best to handle database failures from user-facing applications. For example, when a web application encounters any database outage it can result in transactions not completing properly or transactions being re-entered by the user. While the database outage can be easily recovered, the same may not be true for the application itself. Oracle Database 12c offers customers a new Application Continuity feature that makes for a better user experience by enabling transaction replay, in a non-disruptive manner, following a recoverable database failure. Application Continuity effectively masks database outages from the end user. When used with the Oracle stack, Application Continuity can be enabled with few or no application changes by simply replacing the JDBC driver. It doesn't require application developers to intervene and manually recover application requests and effectively strengthens the fault tolerance of applications running on Oracle Database 12c.

"Redaction in Oracle Database 12c will help us to mask displayed data and we don't have to modify the application code. This means that time-to-market is a lot quicker and the resources required to secure sensitive data is dramatically reduced."

Billy Tong, Sabre

Defense-In-Depth for Maximum Security

Today, enterprise data is increasingly under threat from malicious attacks, and both government and industry regulations require organizations demonstrate stronger controls to protect sensitive data. Many of the benefits of database consolidation could be quickly lost if enterprise data is left insecure or out of compliance with regulatory guidelines. From the outset, Oracle has adopted a defense in depth multi-layered approach to provide customers the relevant preventative, detective and administrative controls to protect enterprise data including:

- **Transparent encryption of data** at rest using Oracle Advanced Security to help protect against threats targeting storage on production servers and backup media devices. Encryption can be easily applied to sensitive columns in tables or entire tablespaces, and prevents access to data when database files are lost, stolen or analyzed directly on media.
- **Separation-of-duties and least privilege preventative controls** using Database Vault to help protect against threats targeting customers' intellectual property, privacy related data, and applications. Oracle Database 12c enhances separation of duty with new reduced privilege roles for backup, key management and audit.
- **First line of defense for databases and consolidated audit trails** using Oracle Audit Vault and Database Firewall. It provides both prevention and detective controls by monitoring and blocking unauthorized SQL traffic before it reaches the database. In addition, its auditing and monitoring controls can be easily tailored to meet specific enterprise security requirements.
- **Comprehensive administrative controls** to help customers maintain secure database configurations throughout the enterprise using Oracle Database Lifecycle Management Pack. It enables proactive monitoring of database accounts, management of privilege entitlements, enforcing password complexity, and ensures tracking and currency of security patches.

In addition to comprehensive integration with Oracle Multitenant, Oracle Database 12c introduces key new preventative, detective and administrative controls that can help customers enhance the protection of their enterprise data.

New Preventative, Detective and Administration Controls

Data Redaction

While transparent data encryption helps protect information from database bypass attacks, data redaction in Oracle Database 12c helps protect information by enforcing controls inside the database that redact data before it is returned to the application. Data redaction effectively hides or dynamically masks the true value of sensitive columns by showing (for example) just the last few digits of social security or bank account numbers that are returned to the application. By defining and enforcing data redaction policies in the database, and not the application, customers can effectively protect sensitive data such as social security numbers, birthdates and bank account numbers from being displayed in application screens and reports with no changes required to applications.

Comprehensive Auditing

Oracle Database 12c provides a new auditing architecture that is both policy-based as well as context-aware, with new roles for managing auditing policies and the viewing of audit data. This new separation of duty enables organizations to designate separate users to manage audit settings and view audit activity. Audit policies can be defined based on factors such as time of day, IP address, program name, and proxy user name, and can be enabled with exception clauses that disable auditing for specific users. This new architecture unifies existing audit trails into a single audit trail, enabling simplified management and increasing the security of audit data generated by the database.

Privilege Analysis

This new feature provides the ability to track privilege and role usage for specific sessions or all sessions connected the database. In many cases, user accounts are over-privileged; in other words, many users have greater security access than normally required to do their jobs. Granting privileged security access increases the threat of risk to enterprise data especially if those accounts become compromised in any way. Privilege analysis effectively records the privileges and roles used and will enable security personnel, database administrators, and auditors to lockdown their databases by designing least privilege models that accurately reflect the privileges required for day to day business and administration activities.

Simplifying Analysis of Big Data

While OLTP applications are essential to process business transactions, data warehousing applications are essential to help measure business performance. In other words, users require secure access to data warehouses 24 by 7, and expect quick responses to their questions. Oracle is not only the world's most popular database for OLTP applications; it also leads the data warehouse market. Oracle Database 12c combines industry-leading performance with comprehensive analytics, and data integration in a single platform that can easily scale to meet the most demanding requirements. It fully supports a wide range of Business Intelligence tools that take advantage of optimizations including; advanced indexing operations, OLAP aggregations, automatic star query transformations, partitioning pruning (a.k.a. partition elimination) and parallelized database operations.

Data Warehouse best practices are becoming fairly well-established, and as the underlying technology continues to mature, many organizations recognize the value-add of evolving their data warehouses to incorporate multiple data sources beyond that of their transactional systems. Incorporating big data sources such as weblogs, device generated data and social media feeds can offer new insights into business performance and opportunities. While technologies such as Hadoop and distributed key value stores (e.g. Oracle NoSQL Database) can help organizations acquire high volumes of low-density data, the real value of big data is realized when analyzed alongside more traditional enterprise data.

Integrating Big Data

Big data requires an architecture that can easily acquire data from multiple data sources and organize it into a suitable format for analysis that enables users to efficiently drive business decisions. In order to efficiently analyze big data, Oracle provides a comprehensive range of integration tools that help customers acquire data from multiple data sources for analysis in an Oracle Database 12c data warehouse. These include:

- **Oracle Data Integrator** is a high-performance bulk data movement and transformation architecture for loading data into Oracle Database 12c from heterogeneous sources.
- **Oracle Big Data Connectors** enable customers to tightly integrate big data environments (e.g. Hadoop) with Oracle Database 12c, and comprise four key components:
 - **Oracle Loader for Hadoop** enables users to use Hadoop's MapReduce framework to create optimized data sets for efficient loading into Oracle Database 12c.
 - **Oracle SQL Connector for Hadoop Distributed File System (HDFS)** offers users the flexibility of querying data from HDFS directly using familiar SQL interface.
 - **Oracle Data Integrator Application Adapter for Hadoop** simplifies integration from Hadoop to Oracle Database 12c through an easy to use graphical interface that generates relevant MapReduce code.
 - **Oracle R Connector for Hadoop** is an R package that provides users of the open-source statistical environment R ability to analyze data stored in HDFS.

By providing a comprehensive set of integration tools, customers can use their existing Oracle resources and skills to bring together new big data sources into their data warehouse. They can now take full advantage of the performance, scalability, security, availability and data management capabilities that Oracle Database 12c delivers to better analyze data throughout the enterprise.

In-Database Analytics

The evolution of data warehouses and big data has brought about increased demand for analysis of data. While the majority of business users will continue to analyze their data warehouse using SQL interfaces and Business Intelligence tools, big data developers and data scientists want tools that can do more complex in-depth analysis. Oracle Database 12c offers customers a choice of in-database analytics tools that enhance existing SQL and Business Intelligence skills, and also supports big data technologies such as Hadoop and R. These include:

- **SQL Pattern Matching** is a new feature of Oracle Database 12c that extends SQL to efficiently write and execute code that can detect patterns across a sequence of rows enabling scalable discovery of business event sequences such as financial transactions, network logs and clickstream logs.
- **In-Database MapReduce** enables developers to write complex procedural logic that can utilize the parallel execution capabilities of Oracle Database 12c for faster performance, especially when processing large volumes of data.
- **In-Database and open source R algorithms** using SQL interfaces and R statistical analysis enables organizations to rapidly build and deploy highly scalable analytical applications that readily integrate with enterprise business processes and business intelligence. In addition to uncovering hidden patterns stored inside the database, statisticians can use R for advanced analytics on very large data sets stored inside and outside the database.
- **In-Database Spatial** integration with enterprise data enables users to better understand geospatial relationships and trends much more efficiently. For example, spatial data can help visualize location of customers alongside purchasing history; providing another aspect to traditional analysis that can help to make more informed decisions.
- **In-Database Graph** that provides customers the ability to create graphs and connections between various data points and data sets. Using semantic analysis, customers can (for example) easily determine a network of relationships for their customers' circle of friends. Graphing relationships can help better understand purchasing trends and churn patterns.

Oracle Database 12c offers customers a database warehousing platform that can store and manage vast volumes of data, and process everything from simple queries to complex analytics using familiar SQL interfaces or statistical frameworks such as R. Instead of constantly moving data from data warehouses to data marts and specialized processing servers (e.g. for OLAP and Statistical Analysis), customers can consolidate both enterprise data and big data into a data warehouse that supports the Business Intelligence and Analytics requirements of all their users. The in-database analytic capabilities of Oracle Database 12c eliminates the need for separate processing silos and helps customers better harness the opportunity that big data presents.

Big Data Infrastructure

Successful data warehouse and big data projects rely on software deployments onto scalable, high-performance hardware infrastructures. The challenge facing many IT departments is getting the right blend of software and hardware components together in a timely manner to meet business demand. By integrating software and hardware components together at the factory, Oracle offers a range of engineered systems that can help IT departments acquire, organize and analyze all their data, and ultimately help the business make informed decisions faster.

- **Oracle Big Data Appliance** is a networked rack of servers and storage complete with a comprehensive big data software stack that includes Cloudera's Distribution including Apache Hadoop (CDH4) and Oracle NoSQL Database Community Edition.
- **Oracle Exadata Database Machine** is a complete preconfigured package of software, servers, and storage that is simple and fast to implement. It's available in a choice of configurations and ideally suited to complex analytic processing and mixed workloads.
- **Oracle Exalytics In-Memory Machine** delivers answers to complex business questions with unmatched speed, intelligence, simplicity and manageability. It is built using industry-standard hardware and features Oracle Business Intelligence Enterprise Edition and Oracle Times-Ten In-Memory Database.

Oracle Big Data Appliance, Oracle Exadata Database Machine and Oracle Exalytics In-Memory Machine all feature 40Gb per second InfiniBand network switches connecting the components within each system, and also between each system. This delivers a high-bandwidth, low-latency network that can easily scale as data demands grow. By engineering software and hardware together, Oracle is helping customers reduce the risk of delivering successful big data projects in a timely manner.

"Pluggable Databases will help lower our administrative costs since we can now manage many databases as one with fewer software installations and patches during the lifetime of our applications."

Jens-Christian Pokolm, Postbank Systems

Managing Database Clouds

Oracle Database 12c and Oracle Enterprise Manager Cloud Control build upon the automated and self-management capabilities of previous releases to help database administrators reduce the time and effort required to meet users' quality of service expectations. Oracle Enterprise Manager Cloud Control is ideally suited to managing the demands of providing databases services on the cloud, and new Oracle Database 12c technologies such as Oracle Multitenant and Data Redaction.

Database as a Service

Oracle Enterprise Manager Cloud Control provides an intuitive interface that automates and simplifies a multitude of administrative tasks from database provisioning to performance tuning, problem diagnosis to upgrades, patching and metering database services. It also ships with a self-service portal that allows developers, testers, administrators, and other self service users to log on and request new single instance or clustered databases, and perform basic management operations. While Oracle Database 12c provides a new architecture to simplify database consolidation on the cloud, and Oracle Enterprise Manager Cloud Control provides the tools to cost-effectively manage users' service level expectations, change can still introduce an element of risk.

Reducing Risks of Getting to the Cloud

Transitioning to the cloud inevitably involves change in the data center, and testing has traditionally relied on developing scripts, manual testing exercises or third party tools. They can provide a limited testing environment, but can leave IT organizations exposed to risk and failures when changes go live. Oracle Real Application Testing enables database administrators to capture entire workloads from production systems and replay those workloads against Oracle Database 12c. It can be used to rigorously test numerous changes including database changes (e.g. upgrade to Oracle Database 12c), configuration changes (e.g. moving from single server to clustered servers), and systems changes (e.g. upgrading Operating System). Oracle Real Application Testing enables database administrators to easily capture and replay the volume, concurrency and mix of workload throughput that only happens with live production systems. It also offers the ability to concurrently test and scale captured workloads which can be very important when consolidating databases on the cloud.

Database Application Development

Oracle Database 12c not only offers customers the latest generation of the world's most popular database, it also offers developers an integrated data management solution that is supported by all popular application development frameworks in use today. These application development frameworks and development tools like Oracle SQL Developer and Oracle Application Express can be used by developers to quickly build applications that take full advantage of the performance, availability and many other features of Oracle Database 12c.

Application Developer Frameworks

Oracle Database 12c offers developers native programmatic interfaces as well as support for a wide range of development and scripting languages including:

- SQL and PL/SQL
- Oracle Call Interface (OCI)
- Programming languages including Java, C and C++
- Scripting languages including PHP, Ruby and Perl
- .NET with Oracle Developer Tools for Visual Studio, Oracle Data Provider for .NET and Oracle Database Extensions for .NET

By providing customers with a choice of application development frameworks, Oracle Database 12c also helps protect customer's investments in existing development resources and skill sets.

Oracle Application Express

Oracle Application Express is a database-centric rapid web application development tool for building a vast array of applications. It is completely declarative, and, using only a web browser, end users and experienced developers can quickly build and deploy fast, reliable and secure database applications. It's ideally suited to power users writing reports or simple forms to experienced SQL and PL/SQL developers implementing sophisticated applications that support business operations.

Oracle SQL Developer

Oracle SQL Developer simplifies the development and management of Oracle databases – including pluggable databases. It offers complete end-to-end development of PL/SQL applications, a worksheet for running queries and scripts, a DBA console for managing the database, a reports interface, a complete data modeling solution, and also a migration platform for moving non-Oracle databases to Oracle Database 12c.

Migrating to Oracle Database 12c

Oracle Database 12c provides the tools and utilities to simplify the process of upgrading and consolidating existing Oracle databases, and Oracle SQL Developer supports migrating 3rd party databases such as IBM DB2 UDB LUW, Microsoft SQL Server and Microsoft Access, Sybase Adaptive Server and Teradata to Oracle. The tight integration capabilities of Oracle SQL Developer provide users with a single interface to browse database objects and data in third-party databases, and easily migrate data accordingly. However, migrating data is only part of the migration process, and Oracle Database 12c introduces a SQL Translation Framework to help customers migrate applications.

SQL Translation Framework

Converting SQL statements in legacy applications originally developed for non-Oracle databases to operate successfully with Oracle can be a major development undertaking. In order to help customers simplify their migration projects, the SQL Translation Framework feature of Oracle Database 12c can accept foreign SQL statements from client applications and translate them at run-time before execution – without requiring manual conversion exercise. Oracle provides a translator for Sybase ASE and SQL Server, and customers can write their own translation code via the use of SQL Translation Profiles. In addition to the SQL Translation Framework, migrating applications to Oracle Database 12c is eased with the introduction of 32K VARCHARs, auto increment identity columns, implicit result sets and FETCH FIRST syntax for Top-N queries.

Plug into the Cloud with Oracle Database 12c

Over the years Oracle has established a track record of delivering breakthrough advances in performance, scalability, high availability, data optimization, data security and ease of management to support the most demanding OLTP, Data Warehousing and Big Data customer requirements. Prior releases of Oracle Database have always been developed with customers' IT and business requirements in mind, and that philosophy very much continues with Oracle Database 12c. For example, Oracle Multitenant was developed in response to customers' requests for an architecture that offers a greater level of consolidation density and greater utilization of IT resources, but without the management overhead associated with alternative architectures.

The potential benefits from consolidating databases onto the cloud can be measured in terms of improved IT service levels and reduced costs. However, evolving data centers away from disparate server and storage silos without incurring en-mass changes to applications and operations has previously been somewhat challenging. Now, with Oracle Database 12c, customers have a new multitenant architecture that eases the path to consolidating databases onto the cloud - without having to change any applications. Customers can simply plug their existing database applications into the cloud, and take full advantage of managing many databases as one, while still retaining the data security and isolation between each pluggable database.

Consolidating databases onto fast, reliable and scalable engineered systems such as Oracle Exadata helps accelerate the path to cloud computing, and Oracle Database 12c will fully complement all the unique software innovations in Oracle Exadata such as Smart Scans, Smart Flash Cache and Hybrid Columnar Compression to help lower IT costs further. To ease management on the cloud, Oracle Enterprise Manager 12c Cloud Control enables customers to manage the entire database lifecycle, from provisioning, to performance tuning and diagnostics, to workload resource management and metering. Finally, Oracle Real Application Testing provides the tools to rigorously test changes; ensuring fast time to market with a quick return on investment.



Plug Into The Cloud with Oracle Database 12c
June 2013

Authors: Penny Avril, Willie Hardie

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200

oracle.com



Oracle is committed to developing practices and products that help protect the environment

Copyright © 2013, 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. 0612

Hardware and Software, Engineered to Work Together