

## Agenda and Abstract for

### Virtual Developer Day - Building Database Apps using Modern Techniques

Tuesday February 4th, 2012 – 9:00am -1:00pm PT / 12:00pm – 4:00pm ET / 3:00pm – 7:00pm BRT

#### Agenda

Time	Track and Keynote/Session Title			
9:00 AM	Keynote – Jonathan Lewis and Tom Kyte discussion -- landscape of Oracle DB technology evolution and adoption			
	<b>Big Data DBA</b>	<b>Big Data Developer</b>	<b>Oracle 12c DBA</b>	<b>Oracle 12c Developer</b>
9:30 AM	NoSQL Installation/cluster topology deployment		Oracle VM performance with Oracle Database 12c	SQL Data Modeler 4.0
10:15 AM	In-DB Map Reduce wSQL /Hadoop		Automatic Data Optimization (ADO) features of Oracle 12c	SQL Developer 4.0
11:00 AM	Processing Twitter Data with Hadoop		RAC vs. Multitenant	XML Database
11:45 AM	Use Data from a Hadoop Cluster with Oracle Database	Application Development & Schema Design with Oracle NoSQL Database	Pattern matching in SQL	Building Mobile Web Applications with Oracle Application Express 4.2
12:30 PM	End-to-End Big Data Demonstration	Make the Right Offers to Customers Using Oracle Advanced Analytics	Database as a Service in a DBA's World	Best Practices for Performance, Scalability, & Availability with Oracle Database12c
1:15 PM	Close			

Hands on Lab Materials can be found at –

- [Hands on Lab Instructions- Database Tracks](#)

- [Hand On Lab Instructions - Big Data Tracks](#)

## Abstracts

### **Keynote – Landscape of Oracle Database technology evolution.**

What's next for Oracle Database 12c? Tom Kyte and Jonathan Lewis discuss the high value technologies and techniques that are driving greater database efficiencies and innovation.

### **Big Data DBA Track:**

#### **NoSQL Installation and Cluster Topology Deployment:**

This session walks through the process and set up for a NoSQL database in a “quick start” format. We will presents the tools and techniques used to manage an Oracle NoSQL database cluster. The discussion will also cover cluster management activities such as initial topology create, topology expansion, online software upgrades, and critical aspects of disaster recovery. Attendees will have hands-on experience managing a production instance of the Oracle NoSQL database.

#### **In-DB Map Reduce with SQL/Hadoop:**

To Analyze enterprise data in Oracle using MapReduce techniques, you can either ship it to external infrastructure or bring MapReduce to RDBMS. Learn how In-Database Container for Hadoop allows MapReduce analysis directly in the database using SQL and/or Hadoop (Java). Processing data in place allows reusing existing BI tools for post-Mapreduce data mining. The lack of operational Hadoop skills is a non-starter for many companies; the tight integration of Hadoop steps in Oracle SQL queries makes easier for non-Java programmers, and SQL-literate developers to consume Mappers & Reducers. DBAs can schedule MapReduce jobs without learning Hadoop.

Data integrity and security is a major requirement for enterprise and sensitive data. In-Database MapReduce inherits Oracle database enterprise security, as well as all quality of services.

#### **Processing Twitter Data with Hadoop:**

Hadoop has become the established tool for dealing with big data, and one of the largest public data sets available comes from Twitter. Utilizing several tools from the Hadoop ecosystem, Twitter data can be efficiently processed and analyzed. Join us as we build a complete system to handle Twitter data – Acquiring data from the Twitter API, processing it using Hive and uploading the results to an Oracle database.

#### **Use Data from a Hadoop Cluster with Oracle Database:**

Learn how to use the high speed Oracle Big Data Connectors to integrate data in Hadoop with data in Oracle Database. Oracle Loader for Hadoop and Oracle SQL Connector for HDFS provide extremely fast and efficient load from Hadoop to the database, and enable Oracle SQL access to data in Hadoop for a number of use cases. These connectors are essential components of any big data solution, integrating data from JSON, Hive, HBase, and other data sources with data in the database. The session will include a hands-on section where you can run complete examples in a VM.

#### **End-to-End Big Data Demonstration:**

[Today](#), you saw many of the individual components comprising a big data solution. This session will tie all of the components together into an integrated application. The end-to end demonstration will leverage Oracle NoSQL Database, Hadoop, Oracle Big Data Connectors and Oracle Advanced Analytics - highlighting capturing of raw application log data to the analysis of that data to make better decisions.

### **Big Data Developer Track:**

#### **NoSQL Installation and Cluster Topology Deployment:**

This session walks through the process and set up for a NoSQL database in a “quick start” format. We

will presents the tools and techniques used to manage an Oracle NoSQL database cluster. The discussion will also cover cluster management activities such as initial topology create, topology expansion, online software upgrades, and critical aspects of disaster recovery. Attendees will have hands-on experience managing a production instance of the Oracle NoSQL database.

### **In-DB Map Reduce with SQL/Hadoop:**

To Analyze enterprise data in Oracle using MapReduce techniques, you can either ship it to external infrastructure or bring MapReduce to RDBMS. Learn how In-Database Container for Hadoop allows MapReduce analysis directly in the database using SQL and/or Hadoop (Java). Processing data in place allows reusing existing BI tools for post-Mapreduce data mining. The lack of operational Hadoop skills is a non-starter for many companies; the tight integration of Hadoop steps in Oracle SQL queries makes easier for non-Java programmers, and SQL-literate developers to consume Mappers & Reducers. DBAs can schedule MapReduce jobs without learning Hadoop.

Data integrity and security is a major requirement for enterprise and sensitive data. In-Database MapReduce inherits Oracle database enterprise security, as well as all quality of services.

### **Processing Twitter Data with Hadoop:**

Hadoop has become the established tool for dealing with big data, and one of the largest public data sets available comes from Twitter. Utilizing several tools from the Hadoop ecosystem, Twitter data can be efficiently processed and analyzed. Join us as we build a complete system to handle Twitter data – Acquiring data from the Twitter API, processing it using Hive and uploading the results to an Oracle database.

### **Application Development & Schema Design with Oracle NoSQL Database:**

This session will highlight new data modeling features in upcoming release of Oracle NoSQL Database. The discussion will cover the concepts of designing schema, creating and using secondary indices on the JSON value objects for fast random access of semi structured data. A real world application will be used to demonstrate how business requirements can be broken down into designing schemas, developing queries, and finally operationalize all the moving pieces into one functioning application.

### **Make the Right Offers to Customers Using Oracle Advanced Analytics**

Oracle makes the popular R statistical programming language and environment ready for the enterprise and big data through tight integration with Oracle Database and Hadoop. Learn how you can use R to work with big data resident in Oracle Database or HDFS, how to operationalize big data analytics, and learn principles and best practices for developing analytics on big data.

## **Oracle 12c DBA Track:**

### **Oracle VM performance with Oracle Database 12c:**

Vibrant session about Oracle VM, that will explain how, when and why use this product for Virtualization. It will also give an overview of how Revera is currently using this product in New Zealand (Biggest Oracle VM Farm in the ANZ region) and show benchmark results between Bare Metal, Oracle VM and other hypervisor technology concluding with some tips and showing the scalability and break point of load of the Virtualization solutions.

Come and discover the answers for the following questions:

- Does an Oracle Database perform well on a virtualized environment?
- What virtualization technology is more stable and allows an Oracle database to perform faster?
- What is the performance difference between using a bare metal and a virtualized guest?
- Is it safe to run a production database in a virtualized environment?

### **Exploring Oracle 12c's Automatic Database Optimization (ADO) Features:**

Oracle Database Release 12cR1's new Automatic Database Optimization (ADO) features now make it

possible to automatically locate data within the most appropriate storage tier and at the appropriate compression level based on its usage patterns. Join Oracle Technical Evangelist and Oracle University instructor Jim Czuprynski as he discusses how to take advantage of these features to save crucial Tier 0 and Tier 1 storage space ... and perhaps even improve query and DML performance.

### **Pattern Matching in SQL for Oracle Database 12c:**

Recognizing patterns in a sequence of rows has been a capability that was widely desired but not possible with SQL, until now. With Oracle Database 12c, you can use the new `MATCH_RECOGNIZE` clause to perform pattern matching using SQL. This hands-on lab starts by building a simple example and then explores some of the more advanced features. We have provided additional workshop material so that you can explore more complex worked examples after this session to help you develop your skills.

### **Database as a Service in a DBA's World:**

DBaaS, (Database as a Service) is a newer EM12c cloud feature that has many DBAs scratching their heads and wondering what it's all about and how it fits in their changing world as we head to the cloud. This session will cover the basics of why DBaaS is so important to provisioning, cloning and the self-service portal. This session will focus on a full explanation of why the feature is important to the future of the database world for many DBAs and a step by step, DBaaS setup that will give DBAs a solid grasp of the concepts and functional requirements necessary to utilize this feature.

### **Oracle 12c Developer Track:**

#### **SQL Data Modeler 4.0:**

This session highlights Oracle SQL Developer Data Modeler's support for Oracle Database 12c features including Identity Columns, Oracle Advanced Security and Data Redaction, Transparent Sensitive Data Protection (TSDP), auto-incrementing columns and `DEFAULT NOT NULL` clauses, invisible columns, and 32k `VARCHAR2`s.

#### **SQL Developer 4.0:**

This session highlights Oracle SQL Developer's Database Multitenant and other 12c Database supported features. Attendees will see how easy it is to manage container and pluggable databases, including a live demo of a database clone. We will also show how easy it is to setup a redaction policy for your sensitive data, create a table with an `IDENTITY` column, and use `pl/sql` functions in your SQL code. Finally, developers will be given an overview of how and why migrating their 3rd party databases and applications to Oracle Database 12c is now easier than ever before.

### **Oracle XML: Develop Flexible Applications More Easily**

This session will examine how the features of Oracle XML in Oracle Database 12c make it easy to develop applications that provide a level of flexibility that goes beyond what is possible with a traditional relational data model including dealing with highly variable data models, such as cases where the data model is extremely dynamic, and cases where the data model is not known until the application is deployed. The solution presented here shows how by using XML and XML Schema, it is possible to define a data model that can represent arbitrary objects using an extremely flexible name-value pair storage model. We will then show how XML DB makes it possible to store and index this XML in manner that delivers full schema flexibility while still providing full relational access and performance. The session also examine how XMLDB and Oracle's Database Native Web Services make it possible to easily integrate Oracle Database 12c and the large volumes of application logic tied up in `PL/SQL` code into today's SOA environments.

### **Pattern Matching in SQL for Oracle Database 12c:**

Recognizing patterns in a sequence of rows has been a capability that was widely desired but not possible with SQL, until now. With Oracle Database 12c, you can use the new `MATCH_RECOGNIZE`

clause to perform pattern matching using SQL. This hands-on lab starts by building a simple example and then explores some of the more advanced features. We have provided additional workshop material so that you can explore more complex worked examples after this session to help you develop your skills.

### **Building Mobile Web Applications with Oracle Application Express 4.2**

Oracle Application Express (Oracle APEX) is a declarative web application development tool for the Oracle Database. With smart phones and tablets approaching 10% of all browser usage, building web applications that work well is becoming increasingly important. This session gives an overview of how customers can quickly build mobile web applications, declaratively with Oracle APEX.

### **Best Practices for Performance, Scalability, & Availability with Oracle Database 12c:**

How would you share a single connection pool across multiple pluggable databases?

Did you ever run into connection issues while attempting to deploy Java applications with tens of thousands of concurrent users?

Have you ever experienced paying twice for the same flight ticket or buying the same article twice online?

If you are a Java architect, Java programmer/designer, or database developer looking to exploit new Oracle Database 12c enhancements in the areas of performance, scalability, availability, security, and manageability/ease-of-use, then this is the session for you.

This technical session discusses new Java capabilities in Oracle Database 12c, including support for the latest Java standards (JDBC 4.1, JDK 7), the new multitenant architecture, application continuity and transaction guard, database resident connection pool, Global Data Service, new SQL types, and the SQL Translator Framework for translating foreign SQL dialects for the Oracle SQL engine.