



The New York Oracle Users Group
Summer General Meeting – June 21, 2018
Sponsored by GRIDDABLE.IO
AGENDA

TIME	ACTIVITY	TRACK/ROOM	PRESENTER
8:30 - 9:00	REGISTRATION AND BREAKFAST		
9:00 - 9:10	Opening Remarks General Information	(Single Session) Room 105/106	Simay Alpoge NYOUG President
SESSION 1 9:15 - 10:30	KEYNOTE: Innovation, the Oracle Cloud, Big Data, IOT and Robotics Future	(Single Session) Room 105/106	Rich Niemiec Viscosity North America.
10:30 - 10:40	BREAK		
SESSION 2 10:40 - 11:40	Database Health with Raspberry Pi	(Single Session) Room 105/106	Arup Nanda Priceline
11:40 - 11:45	BREAK		
SESSION 3 11:45 - 12:45	The Best Oracle Database 18c New Features (and a few 12cR2)	(Single Session) Room 105/106	Rich Niemiec Viscosity North America.
12:45 - 1:45	LUNCH – Room 105/106 - Presentation by GRIDDABLE.IO		
SESSION 4 1:45 - 2:45	Synchronized Data Integration for Hybrid Cloud	DBA Room 105	Chavdar Botev Griddable.io
	Java Best Practices for Developing and Deploying using Databases in the Cloud	Developer Room 106	Kuassi Mensah Oracle Corp.
2:45-2:50	BREAK		
SESSION 5 2:50 - 3:50	Why Oracle Cloud Infrastructure is built for the Enterprise	DBA Room 105	Wilfrido Solano Apps Associates
	Unduly Forgotten Performance Tuning Hero: PL/SQL Hierarchical Profiler	Developer Room 106	Michael Rosenblum Dulcian Inc
3:50 - 4:00	BREAK		
SESSION 6 4:00-5:00	Oracle Exadata - Laying the foundation for Autonomous Database	DBA Room 105	Gurmeet Goindi Oracle Corp.
	Introduction to Oracle Blockchain Cloud Service	Developer Room 106	Kiran Murty AuraBlocks LLC

ABSTRACTS

9:15 - 10:30 AM KEYNOTE: Innovation, the Oracle Cloud, Big Data, IOT and Robotics Future

This presentation will cover Oracle's Acceleration, especially with the Oracle Cloud, and the Cycles of Innovation in the world. We'll look at a quick step by step PaaS example and how easy it is to use. We'll then look at the current Acceleration of the Cloud, the Big Data Revolution and some of the Big Data Solutions out there including Oracle's Big Data Solution. We will review the role that both Google and Facebook have played in the Big Data Ecosystem and how other NoSQL databases like Cassandra and MongoDB fit in. We'll then look at the Internet of Things (IOT) which will drive Big Data and the Cloud even Faster. We'll then look at the future with robotics on the horizon.

Rich Niemiec

Rich Niemiec is the current Chief Innovation Officer of Viscosity North America. Rich is an Oracle Ace Director, a world-renowned IT Expert, and was a co-founder and the CEO of TUSC, a Chicago-based systems integrator of Oracle-based business solutions started in 1988.

10:40 - 11:40 AM SINGLE SESSION: Database Health with Raspberry Pi

Learn some of the tips and techniques of successfully managing a database infrastructure in the real world How is the database performing now? It's a question you must have been asked more than once and what were their reaction when you rattled off multiple metrics to reflect a comprehensive load profile of the database? Perhaps their eyes glazed over. In this session you will learn how to build a tool based on Oracle internal metrics to present a single number for database performance (called DBIndex), similar to stock market index, to show an objective assessment of the database load at any point in time. Anyone can understand it and make inferences. You will also learn how to display that information visually with a LED on a Raspberry Pi device with a bit of python code, some open source tools and a little elbow grease to make a useful device that puts to rest the questions like why the database is slow.

Arup Nanda

Arup Nanda is the VP of Database Services and Chief Data Officer of Priceline, a US\$ 100B New York area company providing of travel services with brands such as priceline.com, booking.com, kayak.com, etc. He has been working in the data and analytic space for 23 years, touching all aspects of data management including relational, non-relational, big data, caching technologies and data sciences. He won two prestigious awards from Oracle Corp: DBA of the Year in 2003 and Architect of the Year in 2012. He has written 6 books, about 700 published articles, 500 presentations around the world and blogs at arup.blogspot.com.

11:45-12:45 AM SINGLE SESSION: The Best Oracle Database 18c New Features (and a few 12cR2)

Oracle's Visual Analyzer tool is especially designed for data exploration, whereas traditional OBIEE This presentation will look at which 18c new features should be investigated for use. Most of the features that will be covered will be related to the DBA, but there will also be a few outside that realm that focus on the developer. There will be simple examples (such as a quick example using pluggable databases) to show the basic functionality of the new features. The speaker has been working with the Beta program over many versions of Oracle.

1. Overview of New Features including Memory Optimization for IOT, Multitenant Advances like Per-PDB Switchover, Zero Impact Grid Infrastructure Patching, Per-PDB Security, polymorphic table functions, Private Temporary Tables, Docker Support, New Installation Support, 18cXE release, Shard-aware applications, In-Memory External Tables, and Approximate Query improvements.
2. Each specific feature (such as pluggable databases)
3. Example of using specific feature (such as pluggable databases)
4. Other new features not discussed but worth considering.

Rich Niemiec

Rich Niemiec is the current Chief Innovation Officer of Viscosity North America. Rich is an Oracle Ace Director, a world-renowned IT Expert, and was a co-founder and the CEO of TUSC, a Chicago-based systems integrator of Oracle-based business solutions started in 1988.

1:45-2:45 PM DBA TRACK: Synchronized Data Integration for Hybrid Cloud

This talk presents the architecture of Griddable.io's smart grids for synchronized data integration. Smart transaction grids are a novel concept aimed at building and operating highly available and scalable infrastructure for synchronized data integration between various data sources and destinations run on private and public clouds. Griddable.io architecture is based on proven open-source technologies such as Data bus; a change data capture system developed and open sourced by LinkedIn. Data bus was created to enable a large range of scenarios for synchronization between the source-of-truth databases and applications for external index maintenance, cache invalidation, read scaling, Hadoop ingest and others. Data bus has been used in production use for many years. Griddable.io has built on that experience and brought significant innovation around selective data synchronization, centralized management, extensibility and cloud support. This talk describes synchronized data integration use-cases, a high-level overview of the architecture and then a detailed discussion of architectural features.

Chavdar Botev

Chavdar Botev is Chief Scientist and co-founder of Griddable.io. Previously he worked as the architect for two critical LinkedIn projects: Data bus, a change data capture system providing reliable, streaming access from core databases to over a hundred internal services, and Goblin, a data lifecycle management system for LinkedIn's massive Hadoop datasets.

1:45-2:45 PM DEVELOPER TRACK: Java Best Practices for Developing and Deploying using Databases in the Cloud

How would you develop and deploy modern Java web apps against a database in the cloud using plain Java/JDBC, Apache Tomcat, Oracle WebLogic, WebSphere, JBoss, and your favorite IDE including NetBeans, IntelliJ, JDeveloper, or Eclipse? This session presents the prerequisites; the security requirements including Java security (JCE) files, KeyStore, and TrustStore; and configuring your Java EE container for accessing the various Oracle database cloud services (DBCS, OCIBare metal, EECS, ADW, and so on). The session will discuss the best practices for optimizing network round-trips between the Java applications and the databases. Live demos will be shown.

Kuassi Mensah

Kuassi is Director of Product Management at Oracle.

He covers the following product areas

- (i) Java connectivity to the Oracle database (JDBC, UCP), in-place processing using the embedded JVM (a.k.a. OJVM) and database quality of services for Java apps (Zero downtime, multi-tenancy, sharding, so on).
- (ii) Turning RDBMS Database tables into Hadoop and Spark data sources

He holds an MS CS from the Programming Institute of University of Paris

Frequent speaker: JavaOne, Oracle Open World, Data Summit, Node Summit, Oracle User groups (UKOUG, DOAG, OUGN, BGOUG, OUGF, GUOB, ArOUG, ORAMEX, Sangam, OTNYathra, China, Thailand, etc.),

Author: <http://www.amazon.com/exec/obidos/ASIN/1555583296>.

@kmensah, <http://db360.blogspot.com/>, <https://www.linkedin.com/in/kmensah>

2:50-3:50 PM DBA TRACK: Why Oracle Cloud Infrastructure is built for the Enterprise

Oracle has built the next generation of cloud infrastructure after researching existing public cloud platforms and seeking input from their existing enterprise customers. They have retained the benefits of first generation IaaS like the ability to add capacity in minutes and to only pay for the computing resources you need, and they've added features like dedicated hardware and Clos networks to minimize latency making it the enterprise cloud with performance as the key goal. They have built ground up data centers with flat non-blocking network in a software defined data center providing the best of class performance. The second generation or modern cloud infrastructure from Oracle has the following capabilities:

- Deploy bare metal servers in minutes
- Raw server performance without the hypervisor overhead
- Pay for what you use
- Integrated compute, storage, and database services in a low-latency private network
- Enterprise-level governance
- High availability for traditional and modern applications
- All features usable via console or API. In this presentation we are going to talk about what makes the Oracle Cloud Infrastructure the "Enterprise Cloud"

Wilfrido Solano

Cloud Adoption evangelist, Information Technology executive, with over 18 years of experience in managing teams, budgets and large/complex projects. Currently helping customers with their Cloud Migration initiatives. Strong understanding of technology as a competitive advantage and a business enabler. Speaker at NYOUG (New York Oracle User Group) and NEAOUG (New England Applications Oracle User Group). Currently helping Oracle with their Lift and Shift campaign co-presenting in 12 cities across the US.

2:50-3:50 PM DEVELOPER TRACK: Unduly Forgotten Performance Tuning Hero: PL/SQL Hierarchical Profiler

The PL/SQL Hierarchical Profiler became available in Oracle 11g Release 1 and replaced the old DBMS_PROFILER package. The goal was the same: to profile runtime behavior of PL/SQL code, i.e. to register and timestamp every operation (including SQL statements) that occurs during the monitoring window. However, the changes were startling:

- * Profiler output is now a file that can be generated in one environment and analyzed elsewhere. This makes production debugging significantly easier because

performance-tuning specialists don't need to touch PROD as long as they have access to the same code base.

- * In addition to loading profiler output to the database and running queries against it, the command line utility PLSHPROF creates HTML-based reports (human readable!). These reports contain various data aggregations to speed up the review process.

This presentation covers a number of real use-cases when HProf significantly shortened response time to production performance problems. You will see how easy it is to figure out the actual source of the slowdown, namely:

- * Hundreds of thousands of calls to a pretty light user-defined function --> Check execution plans of corresponding queries

- * Sluggish SQL statement --> Don't blame PL/SQL and start SQL tracing

- * Strange third-party calls in wrapped packages --> Collect hard evidence and start complaining

Overall, if you write PL/SQL, you must utilize the PL/SQL Hierarchical Profiler. Otherwise, you will be guessing at your code behavior instead of knowing it, which could lead to unpredictable performance. Unpredictable performance often means production calls at 3 am on Sunday. This presentation will help you sleep longer!

Michael Rosenblum

Michael Rosenblum is a Software Architect/Development DBA at Dulcian, Inc. where he is responsible for system tuning and application architecture. Michael supports Dulcian developers by writing complex PL/SQL routines and researching new features. He is the co-author of PL/SQL for Dummies (Wiley Press, 2006), PL/SQL Performance Tuning Tips & Techniques (Oracle Press, 2014), contributing author of Expert PL/SQL Practices (A Press, 2011), and author of several database related articles (IOUG Select Journal, ODTUG Tech Journal) and conference papers. Michael is an Oracle ACE, a frequent presenter at various Oracle user group conferences (Oracle OpenWorld, ODTUG, IOUG Collaborate, RMOUG, NYOUG, etc.), and winner of the ODTUG Kaleidoscope 2009 Best Speaker Award. In his native Ukraine, he received the scholarship of the president of Ukraine, a Master of Science degree in information systems, and a diploma with honors from the Kiev National University of Economics.

4:00-5:00 PM DBA TRACK: Oracle Exadata - Laying the foundation for Autonomous Database

Oracle Autonomous Database is changing the paradigm when it comes to relational database. Platform plays a key role on how an Autonomous Database is delivered. This session talks in technical detail on how Exadata's software features are enabling autonomous capabilities in the database

Gurmeet Goindi

Gurmeet is the product manager for Oracle Exadata. He's been with Oracle for past six years and is responsible Exadata's hardware and roadmap. Gurmeet is a regular presenter at various Oracle User Groups worldwide. Gurmeet has an MBA from University of Chicago's Booth School of Business.

4:00-5:00 PM DEVELOPER TRACK: Introduction to Oracle Blockchain Cloud Service

This presentation will introduce the audience to Blockchain technology. It will further go in detail and talk about the architecture behind Oracle's Blockchain Cloud Service(OBCS). The presentation will introduce audience to doing Proof of Concept(POC) project on OBCS. It will also discuss on how leadership at companies can on-board this technology as a competency into their respective organizations. The presentation will be a combination of Basics, Architecture and Strategy on Blockchain technology.

Kiran Murthy

Kiran Murthy is the CTO of AuraBlocks.com, an Oracle partner that enables blockchain for enterprises. AuraBlocks.com was the first company in the world to

deliver two working solutions at Oracle Open World 2017 on Oracle Blockchain Cloud Service.

AuraBlocks is one of less than 15 companies in the world to be part of Oracle's early access beta program for Oracle Blockchain cloud service. More at AuraBlocks.com

Kiran has worked in roles ranging from an analyst to a CTO for various financial services organizations in Insurance, Asset Management and Hedge fund space. He has created and led several operations, engineering teams. He has a bachelor's degree in Computer Science and graduated from Columbia University with an executive master's degree in Technology & Management. He is the founder of Columbia University Blockchain Alliance.

columbiablockchainalliance.com.