



**The New York Oracle Users Group**  
**Summer General Meeting – June 21, 2018**  
**St John’s University, Manhattan Campus**  
**Sponsored by GRIDDABLE.IO**

**AGENDA**

<b>TIME</b>	<b>ACTIVITY</b>	<b>TRACK/ROOM</b>	<b>PRESENTER</b>
8:30 - 9:00	<b>REGISTRATION AND BREAKFAST</b>		
9:00 - 9:10	<b>Opening Remarks General Information</b>	(Single Session) Room 105/106	Simay Alpoge NYOUG President
<b>SESSION 1</b> 9:15 - 10:30	<b>KEYNOTE: Innovation, the Oracle Cloud, Big Data, IOT and Robotics Future</b>	(Single Session) Room 105/106	Rich Niemiec Viscosity North America.
10:30 - 10:40	<b>BREAK</b>		
<b>SESSION 2</b> 10:40 - 11:40	<b>Database Health with Raspberry Pi</b>	(Single Session) Room 105/106	Arup Nanda Priceline
11:40 - 11:45	<b>BREAK</b>		
<b>SESSION 3</b> 11:45 - 12:45	<b>The Best Oracle Database 18c New Features (and a few 12cR2)</b>	(Single Session) Room 105/106	Rich Niemiec Viscosity North America.
12:45 - 1:45	<b>LUNCH – Room 105/106 - Presentation by Robin Purohit, GRIDDABLE.IO Synchronized Data Integration for Hybrid Cloud</b>		
<b>SESSION 4</b> 1:45 - 2:45	<b>A Smart Grid for Enterprise Data</b>	DBA Room 105	Rajeev Bharadhwaj Griddable.io
	<b>Using AI-powered Analytics in Oracle DB 12c for automatic Image Recognition</b>	Developer Room 106	Lakshman Bulusu
2:45-2:50	<b>BREAK</b>		
<b>SESSION 5</b> 2:50 - 3:50	<b>Why Oracle Cloud Infrastructure is built for the Enterprise</b>	DBA Room 105	Wilfrido Solano Apps Associates
	<b>Unduly Forgotten Performance Tuning Hero: PL/SQL Hierarchical Profiler</b>	Developer Room 106	Michael Rosenblum Dulcian Inc
3:50 - 4:00	<b>BREAK</b>		
<b>SESSION 6</b> 4:00-5:00	<b>Introduction of Jupyter Notebook And Packages for Data Analysis and Plotting</b>	DBA Room 105	Linda Li Priceline
	<b>Introduction to Oracle Blockchain Cloud Service</b>	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: A Smart Grid for Enterprise Data**

This talk presents the architecture of Griddable.io's smart grid 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 for hybrid cloud environments. Griddable.io architecture is based on DataBus, a change data capture system developed, used in production and open sourced by LinkedIn. DataBus was created to enable a number of scenarios for synchronization between the source - of - truth databases and target databases for external index maintenance, cache invalidation, read scaling, Hadoop ingest and others. Griddable.io has added significant innovation around data policies, selective data synchronization, centralized management, extensibility and cloud support. This talk describes synchronized data integration use cases, the design philosophies and a detailed architecture deep-dive of Griddable.io's smart grid.

**Rajeev Bharadhwaj**

Rajeev Bharadhwaj is a founder and VP of Technology and Operations at Griddable.io. Rajeev Bharadhwaj is a serial entrepreneur who created three startups – Ejasent (sold to Veritas), Aryaka Networks (private company on an IPO path) and now Griddable.io. Using his background in Databases and Operating Systems, he created products and companies that introduced new concepts like application virtualization, cloud overlay WAN networking and now policy-based abstraction of data from databases. In addition, he has served several executive roles at large companies like HP, Veritas, Symantec, Oracle and HaL. Prior to founding Griddable.io, he was VP of Engineering for HP's cloud & NFV offerings based on Openstack.

**1:45-2:45 PM      DEVELOPER TRACK: Using AI-powered Analytics in Oracle DB 12c for automatic Image Recognition**

This talk presents how AI powered analytics can be used for automatic Image Recognition. It explains how Machine Learning and Deep Learning algorithms can be used to automate Feature Extraction, Feature Selection, and Feature Engineering to discern Feature Recognition. It highlights how this can be done inside of Oracle DB 12c to leverage the speed, reliability, and scalability of the database. A real-world use case is described to explain the same.

**Lakshman Bulusu**

Lakshman Bulusu has 25 years of experience in using Oracle and its related technologies through version 12c. He has authored 9 books starting from Oracle Developer Suite to PL/SQL to BI and Open Source BI. He has presented at various national and international conferences like NYOUG Summer General Meeting 2017, NJOUG June 2017, Piscataway Data Science Meetup June 2017, Data Summit 2014, OAUG Connect - EPM/BI 2011, IOUG - Collaborate, VOUG, and ECO Oracle Conference; and contributed a host of technical articles for magazines/journals like Oracle Professional, SELECT Journal of IOUG, ODTUG Technical Journal etc.

**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 NEAUG (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: Introduction of Jupyter Notebook And Packages for Data Analysis and Plotting**

Jupyter is an interactive development tool for Python related analytics, for rapid prototyping. The foundation of ML is linear Algebra, the great chunk of it is pertaining matrix operation, NumPy is not directly exposed but they are dependency of some typical ML packages, i.e. scikit learn. Pandas DataFrame - Operation on data sets to help automation and analysis on the data sets. Present a work relevant example to using Pandas plotting feature to get perspective of Tableau user license information.

**Linda Li**

Linda Li - Joined the Priceline Data Lake Service team in 2014, is a senior data warehouse analyst. Has worked on Oracle DBA, informatica, Business Intelligence in various industries - Reinsurance, Bank, e-commerce and worked with systems like SAP Business warehouse, Oracle EBS, EDW, OBI warehouse. Worked with various business intelligence tools - SAP Business Objects, Oracle Discoverer, OBIEE, Tableau. Has interest and passion in data and most cutting-edge technologies. In leisure time, Linda loves music appreciation and loves watching the young generation practice the piano and play in recitals, Chamber Music. Linda is a graduate of New Jersey Institute of Technology, major in Information systems.

**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 Murty 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](http://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](http://columbiablockchainalliance.com).