





NoCOUG FALL CONFERENCE 2017

Thursday, November 9, 2017

California State University, East Bay Oakland Center in the Trans Pacific Centre
1000 Broadway, Suite 109 Oakland, CA 94607

8:30–9:00	REGISTRATION		
	TEMESCAL (floorplan on reverse)	MONTCLAIR (floorplan on reverse)	ROCKRIDGE (floorplan on reverse)
9:00–9:30	GENERAL SESSION Jeff Mahe, NoCOUG Vice President		
9:30–10:15	KEYNOTE <i>Rationalizing Oracle, MySQL, and NoSQL</i> Ritesh Chhajer, PayPal		
10:15–10:30	BREAK (15 minutes)		
10:30–11:15	<i>Squeezing the Maximum Out of Oracle Database 12c Disaster Recovery</i> Alain Azagury, Axxana	<i>How to Prepare for Moving Oracle Applications to the Cloud</i> Shubho Saha and Shiv Lodha, NexInfo	<i>Making Hybrid Transactional/Analytic Processing Possible: OLTP + OLAP = HTAP</i> Ganesh Balabharathi, Vexata
11:15–11:30	BREAK (15 minutes)		
11:30–12:15	<i>Considerations in a High-Volume and High-Velocity OLTP Database</i> Samrat Roy, Dheeraj Kondapaneni, and Ravi Pedapati, PayPal	<i>ASH Workshop—Part I When and How to Use ASH</i> Craig Shallahamer, OraPub	<i>Hadoop for Database Professionals—Part I Introduction and Fundamentals of Hadoop</i> Tanel Pöder and Michael Rainey, Gluent
12:15–13:00	LUNCH (45 minutes)		
13:00–13:45	<i>Verifying the Quality of Data Replicated by Oracle GoldenGate</i> Dong Wang, PayPal	<i>ASH Workshop—Part II (Hands-On Lab) Incident Analysis Mastery: Get Your Feet Wet</i> Craig Shallahamer, OraPub	<i>Hadoop for Database Professionals—Part II Hadoop Storage and Data Ingestion</i> Tanel Pöder and Michael Rainey, Gluent
13:45–14:15	BREAK (30 minutes) Last chance to get your raffle form stamped by the exhibitors		
14:15–15:00	<i>Sharding in Oracle Database 12c and 18c—Part I</i> Nagesh Battula and Mark Dilman, Oracle Corp.	<i>ASH Workshop—Part III (Hands-On Lab) Incident Analysis Mastery: Full Submersion</i> Craig Shallahamer, OraPub	<i>Hadoop for Database Professionals—Part III SQL Processing in Hadoop</i> Tanel Pöder and Michael Rainey, Gluent
15:00–15:15	FREE RAFFLE (15 minutes)		
15:15–16:00	<i>Sharding in Oracle Database 12c and 18c—Part II</i> Nagesh Battula and Mark Dilman, Oracle Corp.	<i>ASH Workshop—Part IV (Hands-On Lab) Visualizing an Incident with ASH and R</i> Craig Shallahamer, OraPub	<i>Hadoop for Database Professionals—Part IV Hadoop in Action</i> Tanel Pöder and Michael Rainey, Gluent
16:00–16:15	BREAK (15 minutes)		
16:15–17:00	<i>Oracle Database In-Memory Option in Oracle Database 12c and 18c</i> Shasank Chavan, Oracle Corp.	NETWORKING AND HAPPY HOUR HOSTED BY AXXANA THE TRAPPIST, 460 8TH ST., OAKLAND, CA 94607 (walking directions on reverse)	

Mark your calendar for the Winter Conference on Thursday, February 15



AXANA
 B U I L T T O L A S T

TEMESCAL

Rationalizing Oracle, MySQL, and NoSQL

—Ritesh Chhajer, PayPal..... 9:30

As the database landscape continues to grow, with a plethora of data stores available to meet ever-increasing business demands, we'll try to clarify the choices that will aid in prudent decision making for data store selection, and also discuss best practices.

Squeezing the Maximum Out of Oracle Database 12c Disaster Recovery

—Alain Azagury, Axxana.....10:30

Oracle 12c introduced new disaster recovery capabilities, such as Active Data Guard Far Sync, a SYNC/ASYNc cascading feature that uses a lightweight intermediary node (Far Sync). This workshop will describe our experience in tuning Far Sync for maximum performance and minimal resource requirements. We will provide detailed descriptions of how configuration parameters influence the behavior of the system—from a memory, CPU, and I/O perspective—while maximizing database performance (e.g., Log Write bandwidth) and minimizing latency impact on the primary system. Finally, we will describe how Axxana's Phoenix technology optimizes Oracle's Far Sync solution to provide true zero data loss with optimal performance.

Considerations in a High-Volume and High-Velocity OLTP Database

—Samrat Roy, Dheeraj Kondapaneni, and Ravi Pedapati, PayPal.....11:30

Driving with a flat tire along a straight, level road on a sunny day isn't as bad as it would be to do so in extreme weather on a steep, winding road in the middle of night. It's even more challenging when you're expected to arrive at your destination on time. There is a similar distinction between managing a 10 TB database with 10K/sec executions and managing a half-a-PB database with 400K/sec executions.

There are many things we need to consider as we plan our daily order of business: How do we set up a table with maximum scale that is expected to get 5K/sec writes? How do we set up an index on a table like that? How do we set up our database to expect 100 MB/sec block changes? How do we set up our SGA? How do we even set up our system for a database like this? Things that were a no-brainer on a 10 TB, 10K/sec database suddenly need two days of planning. This presentation will help DBAs through the planning process in this challenging and interesting database world.

Verifying the Quality of Data Replicated by Oracle GoldenGate

—Dong Wang, PayPal..... 13:00

In a distributed system with multiple databases having data replicated by Oracle GoldenGate, maintaining a high data quality is critical for the overall system. In this session, we will discuss some scenarios that affect data quality in PayPal's Oracle GoldenGate environment. We'll then follow up on a real-time data validation tool that we presented two years ago to validate data replicated among Oracle databases. New developments to enhance data quality comparison by leveraging an Oracle GoldenGate independent data change collection method will be presented, and results from running the validation tool in PayPal's production environment will be presented.

Sharding in Oracle Database 12c and 18c (Part I–II)

—Nagesh Battula and Mark Dilman, Oracle Corp. 14:15

Oracle Database 12c Release 2 implements native sharding—horizontal partitioning across many databases using a shared-nothing hardware architecture, to achieve linear scalability, complete fault isolation, and geographic data distribution for OLTP and data analytics applications, where the primary access pattern is based on the sharding key. Come to this session to learn how Oracle Sharding automates the deployment of sharded databases; allows creation of sharded schema with SQL; and supports elastic scaling, automatic rebalancing, and direct routing, along with proxy routing for multi-shard queries. It does all this while rendering strict consistency, the full power of SQL, and the proven enterprise qualities of Oracle Database.

Oracle Database In-Memory Option in Oracle Database 12c and 18c

—Shasank Chavan, Oracle Corp..... 16:15

Oracle Database In-Memory (DBIM) is the industry-leading in-memory database technology, accelerating analytics by orders of magnitude as well as improving hybrid online transaction processing applications on databases of any size. Oracle Database In-Memory has been successfully adopted by a variety of customers across many industries in the three years since initial release, making it one of the fastest-growing technologies from Oracle. In this talk we'll detail some of the more innovative technologies coming out of the 18.1 Oracle DBIM release, including features such as Autonomous In-Memory Management, Optimized Arithmetic, Software-In-Silicon, Delta-IMCUs, Vector Aggregation, XMEM (non-volatile memory) store, and IM for External Tables support. Also presented in this talk are two exciting In-Memory OLTP features that target Key-Value and IoT workloads. We'll conclude with a discussion of where the Oracle Database product is heading in subsequent releases.

MONTCLAIR

How to Prepare for Moving Oracle Applications to the Cloud

—*Shubho Saha and Shiv Lodha, NexInfo*10:30

Customers currently using many on-premise applications, including Oracle EBusiness Suite, Agile PLM, Siebel, JDE, and so on, are looking to move to the cloud. It took many customers years to achieve stability with their current systems. Now they're afraid of disruption and of missing something vital in this migration. What do you need to know? How do you educate your organization? How can you prepare for the migration so that it doesn't disrupt operations? What factors influence a successful migration? Come to this session and learn, so that when you talk to Oracle next time, you can move forward with confidence.

ASH Workshop (Part I–IV)

—*Craig Shallahamer, OraPub*11:30

There is a better way to solve your most difficult Oracle performance problems. Some performance issues are literally impossible to solve using an AWR Report or Oracle's time model. ASH comes to the rescue! But even DBAs who are familiar with ASH have difficulty using it to directly answer key management issues surrounding a performance incident. In this workshop, you will learn how to weave your ASH analysis into a compelling "whodunit" mystery, complete with good guys, bad guys, motive, and even the weapon!

You will be able to:

- Understand the benefits of both time-based analysis and an ASH sample-based analysis
- Pinpoint when an incident began, when it ended, who was involved, and what they were doing
- Jumpstart your analysis using OraPub's free BloodHound and Scratchpad ASH tool kits
- Create a chart to visualize an incident at a specific point in time
- Develop an incident timeline to explain the mystery
- Participate hands-on as we analyze multiple performance incidents

All the hands-on work is done on your laptop. Craig will supply the Virtual Box software and appliance for both Windows and Mac OS X on a USB thumb drive, to be loaded at the start of the workshop. Virtual Box will require 4 GB of memory and 2 CPU cores.

Part I—When and How to Use ASH11:30

Part II—Incident Analysis Mastery: Get Your Feet Wet13:00

Part III—Incident Analysis Mastery: Full Submersion14:15

Part IV—Visualizing an Incident with ASH and R15:15

ROCKRIDGE

Making Hybrid Transactional/Analytic Processing Possible: OLTP + OLAP = HTAP

—*Ganesh Balabharathi, Vexata* 10:30

Hybrid transaction/analytical processing (HTAP) refers to the capability of performing OLAP (online analytical processing) and OLTP (online transaction processing) at the same time and on the same database system. Oracle Database has always held the edge in HTAP because its version-based concurrency model means that analytical workloads do not block transactional workloads. Very large SGA sizes, huge pages, database flash cache, and the Oracle Database In-Memory option also facilitate HTAP. In this live demo session, we will examine how Vexata scalable storage systems utilize technologies such as NVMe Flash and Intel Optane (3D-XPoint) that can take HTAP to the next level. Be prepared to experience 8 million TPM (HammerDB, TPC-C transactions per minute) and 20 GBps (HammerDB, TPC-H query performance) with amazing 82us read and 91us write latencies!

Hadoop for Database Professionals (Part I–IV)

—*Tanel Pöder and Michael Rainey, Gluent* 11:30

At Gluent, we believe there is a fundamental shift underway in IT to include open, software defined, distributed systems like Hadoop—and, as a result, that every RDBMS professional should strive to learn these new technologies or risk being left behind. Whether on-premises or in the cloud, the new world of distributed data storage and compute is here to stay.

We created this half-day course for IT professionals that have a deep knowledge of RDBMS systems so they can better understand Hadoop and the benefits it brings to the enterprise. The course will compare and contrast subjects such as high availability, backup & recovery, data storage, SQL processing, system monitoring, and day-to-day operations. You'll walk away with a first-hand look at how Hadoop solutions have been built in the real world.

Part I—Introduction and Fundamentals of Hadoop 11:30

We present a brief history and overview of Hadoop and how it compares to and contrasts with a relational database.

Part II—Hadoop Storage and Data Ingestion 13:00

We dive deeper into how HDFS works and describe how to get data into Hadoop.

Part III—SQL Processing in Hadoop 14:15

We review and discuss various SQL on Hadoop technologies, showing that not all data processing routines must be hand coded in Java.

Part IV—Hadoop in Action 15:15

We show how Hadoop has been successfully implemented in the real world, for both Big Data and traditional enterprise data.