



NoCOUG SPRING CONFERENCE & DEVELOPER WORKSHOP

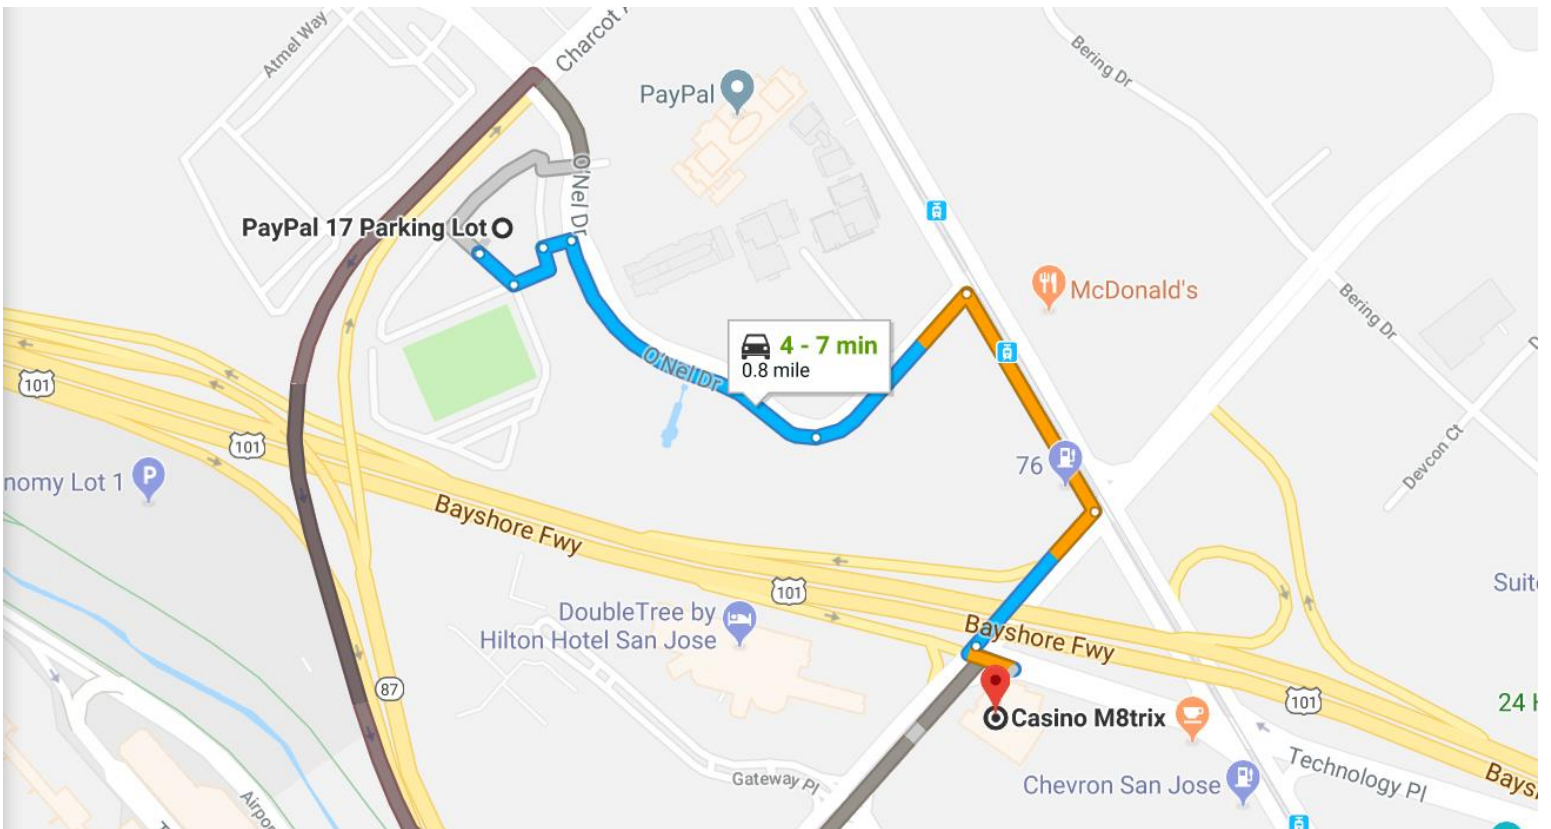
Thursday, May 17, 2018 at PayPal Town Hall, San Jose

8:00–9:00	REGISTRATION AND CONTINENTAL BREAKFAST		
	FIRESIDE AB <i>(floor plan on reverse)</i>	TOWN SQUARE AB <i>(floor plan on reverse)</i>	TOWN SQUARE C <i>(floor plan on reverse)</i>
9:00–9:30	<i>PayPal Tour: Commerce Innovation Showcase (assemble at registration desk)</i>	GENERAL SESSION Jeff Mahe, NoCOUG Vice-President	No session
9:30–10:30	No session	KEYNOTE SESSION <i>Powering Databases at Scale with Mid-Tier Proxies</i> Kamlakar Singh, PayPal	No session
10:30–11:00	MEET THE EXHIBITORS		
	PAYPAL SHOWCASE	MEMBER SHOWCASE	DEVELOPER WORKSHOP
11:00–12:00	<i>Best Practices to Enable Global-Scale Indexes</i> Alexander Sherbak, PayPal	<i>Deep Dive: Automatic Data Optimization (ADO) for Information Lifecycle Management</i> Hariharan Lakshmanan, Oracle	<i>Oracle Database Features Every Developer Should Know About!</i> Maria Colgan & Gerald Venzl, Oracle
12:00–13:00	<i>PayPal Tour: Commerce Innovation Showcase (assemble at registration desk)</i>	LUNCH	
13:00–14:00	<i>Aerospike NoSQL Database at PayPal</i> Athreya Gopalakrishna, PayPal	<i>Active-Active Ecosystem at LinkedIn</i> Janardh Bantupalli & Sai Sundar, LinkedIn	<i>Hands-On Workshop: Everything Developers Need to Know About Integrating JSON into Oracle Database—Session I</i> Maria Colgan & Gerald Venzl, Oracle
14:00–14:30	LAST CHANCE TO MEET THE EXHIBITORS		
14:30–15:30	<i>Couchbase NoSQL Database at PayPal</i> Pradeep Tummala, PayPal	<i>Visualizing Active Session History (ASH)</i> John Beresniewicz, Independent Consultant	<i>Hands-On Workshop: Everything Developers Need to Know About Integrating JSON into Oracle Database—Session II</i> Maria Colgan & Gerald Venzl, Oracle
15:30–16:00	FREE RAFFLE		
16:00–19:00	GRAND RECEPTION HOSTED BY GRIDDABLE VIP KARAOKE LOUNGE, 7TH FLOOR, CASINO M8TRIX, 1887 MATRIX BLVD <i>(map on reverse)</i>		<i>Hands-On Workshop: Everything Developers Need to Know About Integrating JSON into Oracle Database—Session III</i> Maria Colgan & Gerald Venzl, Oracle

*Mark your calendars for the Summer Conference on **Thursday, August 16** at Chevron, San Ramon*



Town Hall (B12) Building - 1st Floor
 2211 North First Street



Mark your calendars for the Summer Conference on Thursday, August 16 at Chevron, San Ramon

NOCOUG SPRING CONFERENCE & DEVELOPER WORKSHOP

Thursday, May 17, 2018 at PayPal Town Hall, San Jose

TOWN SQUARE AB

Powering Databases at Scale with Mid-Tier

Proxies—*Kamlakar Singh, PayPal*9:30–10:30

A database access gateway running a mid-tier service proxying all requests to the database adds numerous benefits, such as connection multiplexing, read/write traffic split, and horizontal scale (sharding) capability. This becomes especially important in a multi-tech-stack environment (Java, C++, Python, Node) combined with a micro-services architecture, by decoupling the complexity from the client tech stack to a tech-stack-agnostic service. In this presentation we will talk about a real case study of a transaction-aware database proxy, built at PayPal, which evolved over time to meet several complex requirements and is built to last and scale. It is currently serving 2 million SQL queries per second and has been designed to provide high availability, automatic recovery/healing, and resiliency. We will cover the journey, motivations, design principles, constraints, and lessons learned along the way.

Deep Dive: Automatic Data Optimization (ADO) for Information Lifecycle Management—

Hariharan Lakshmanan, Oracle 11:00–12:00

This talk will introduce Oracle's Automatic Data Optimization (ADO) technology—and the underlying infrastructure—along with the heat map technology, and will discuss broadly how these technologies can be used to tackle data information lifecycle management challenges. The talk will then provide a deep dive into the ADO and the heat map technologies. The talk will end by providing pointers to ADO implementation for some common use cases.

Active-Active Ecosystem at LinkedIn—*Janardh*

Bantupalli & Sai Sundar, LinkedIn 13:00–14:00

This presentation will provide an overview of LinkedIn's active-active architecture and infrastructure for online user data stores using Oracle GoldenGate, Big Data Adapter, and Python. The session covers design and implementation of an active-active infrastructure and ecosystem to support conflict resolution, constraints, and data validation.

Visualizing Active Session History (ASH)—*John*

Beresniewicz, Independent Consultant..... 14:30–15:30

ASH data is a high-dimensional, fine-grained record of Oracle database activity over time. Looking into ASH data can provide a detailed view of what happened in a database; however, using SQL for this purpose requires aggregation or a high level of filtering to produce interpretable results. Looking at lots of data over many dimensions requires data visualization. This presentation

will briefly discuss the ASH mechanism, its key features, and the visualization of ASH data as presented in Enterprise Manager. Next we will present several visualization investigations into a set of ASH dumps collected from a 4-node RAC cluster using the R programming language and the ggplot package to visualize the data. Some of the investigations make use of a little-known technique for estimating event counts from ASH samples. It is hoped that participants will gain a new or renewed enthusiasm for ASH data—and perhaps even an interest in data visualization.

FIRESIDE AB

Best Practices to Enable Global-Scale Indexes—

Alexander Sherbak, PayPal 11:00–12:00

Attend this session to learn how a complex set of requirements led to a highly scalable (write throughput) and elastic indexing pattern. Discover how a dozen composite indexes on a single table were replaced with a single structure that amassed over 70 equivalent indexes with no marked impact on performance. Finally, learn how PayPal powers its transaction search and how to extend and apply this pattern to your environment.

Aerospike NoSQL Database at PayPal—*Athreya*

Gopalakrishna, PayPal 13:00–14:00

Adoption of KV-distributed databases and maturity in hardware transformed the NoSQL products from in-memory to memory-first, and then to hybrid-memory architecture, by leveraging memory for speed, scale, and availability. This is the story of PayPal's KV NoSQL database journey, which saw the architectural transformations and adopted Aerospike to leverage Hybrid Memory Architecture. HMA addresses cost and performance consistency by efficiently using memory while extensively leveraging SSD and allowing high-density storage with minimal performance overheads when compared to memory-first architectures. It manages hundreds of terabytes of data with hundreds of billions of keys served at consistent low latencies.

Couchbase NoSQL Database at PayPal—*Pradeep*

Tummala, PayPal 14:30–15:30

The shift toward the digital economy is driving NoSQL adoption. This session will cover Couchbase Server—a distributed NoSQL document-oriented database—and its architecture, use cases, and footprint at PayPal. We will also talk about how Couchbase is different from Oracle, its performance and scalability in comparison to Oracle, and some tips and tricks. The ingredients for a successful Couchbase deployment are automation, alerting, and monitoring. Come learn all of this and more in this intriguing session.

NOCOUG SPRING CONFERENCE & DEVELOPER WORKSHOP

Thursday, May 17, 2018 at PayPal Town Hall, San Jose

TOWN SQUARE C

Oracle Database Features Every Developer Should Know About!—*Maria Colgan & Gerald Venzl*,

Oracle 11:00–12:00

Over the years, a lot of new features have been introduced in the Oracle Database to save developers time and to spare them from having to reinvent the wheel. Many of these features are actively highlighted via conference sessions or blogs when a new release is launched, but over time they are forgotten before they can be put to good use. In this session we will discuss critical features that already exist in the Oracle Database that every developer should not only know about but also utilize. Using easy-to-follow examples, we will demonstrate when and where you should take advantage of these features to make your life easier and help you get the most out of your Oracle Database, regardless of whether you are a DBA or a developer.

Hands-On Workshop: Everything Developers Need to Know About Integrating JSON into

Oracle—*Maria Colgan & Gerald Venzl*,

Oracle 13:00–17:00

This hands-on lab will explain how to use modern development techniques and frameworks to build applications that take full advantage of the JSON data model and Oracle Database 18c. It will introduce the new Document Store capabilities of Oracle Database 18c and Simple Oracle Document Access (SODA), a new family of APIs that make it easy for application developers to use Oracle Database as a JSON document store.

The tutorial is presented using a simple application that allows users to search theaters to see what movies are playing or see which theaters are showing a particular movie. Once the user has chosen the screening they want to attend, they can purchase tickets. The application supports simple searches based on city and zip and full-text searches based on title and plot. It also allows location (nearby) searching of theaters.

The tutorial will show advances in Oracle 18c that make it easy to create applications that deliver the full power and flexibility of HTML, JavaScript, and JSON, combined with the security, scalability, reliability, and availability of the Oracle Database. We will take an in-depth look at all three tiers of such an application. Full code examples will be provided.

The tutorial will quickly present the front-end architectures, built using HTML 5 and JavaScript, combined with popular open-source frameworks including JQuery, Bootstrap, and Angular.js. The front end consumes a set of REST services provided by the

middleware components.

It will present two alternative architectures for the middle tier. The first is based on Node.js, which combines JavaScript programming with popular node.js modules, such as express.js and request.js, to create the set of services consumed by the front end. Node.js communicates with the back-end data store using Oracle's new SODA for REST API, which provides a set of microservices that enable the Oracle Database to operate as a JSON document store. SODA for REST is delivered as part of Oracle REST Data Services (ORDS). The second is a traditional Java implementation, based on the Jetty-embedded web server and Jersey REST services. It uses Oracle's new SODA for Java API to communicate directly with a document store managed by the Oracle Database. We will learn how both implementations of SODA deliver the true NoSQL application development experience, allowing developers to develop, deploy, and maintain their application without having to learn SQL or requiring support from DBAs.

The tutorial will cover using the SODA API to create and drop collections (containers for JSON documents), as well as how to insert, update, and delete documents. It will also show how to list and search collections using SODA's query-by-example feature and address issues such as concurrency control.

Finally, we will see that by adopting Oracle as our NoSQL JSON document store we still have access to all the benefits of SQL when it comes to performing query and analytics on our JSON documents. NoSQL APIs and query-by-example make developing and deploying applications easy, but these techniques are not suitable for ad-hoc reporting and analytics. Over the last 30 years Oracle's SQL has been developed to satisfy an organization's reporting and analytics requirements. We will look at exciting new extensions to SQL that enable SQL operations directly on JSON documents.

The tutorial will also introduce other exciting new JSON-related functionality delivered with Oracle Database 18c, including the following:

- Functionality that allows you to automatically discover metadata about the JSON documents you are managing
- How to use indexing, in-memory database, and Oracle Exadata to optimize operations on JSON documents
- How to perform partial updates on JSON documents
- How to generate JSON documents from relational data
- How to use Oracle Spatial with JSON