# Cloud Analytics Seminar

### Keynote: The Future of Analytics in the Cloud—*Kent Graziano, Snowflake*

### Snowflake—Changing the Game with Cloud Data Warehousing—*Kent Graziano, Snowflake*

We all know that data warehouses and their related best practices are changing dramatically these days. As organizations build new data warehouses and modernize established ones, they are turning to Data Warehousing as a Service (DWaaS), hoping to take advantage of the performance, concurrency, simplicity, and lower cost of a SaaS solution—or simply to reduce their data center footprint (and the maintenance that goes with it).

1. Demystify DWaaS by defining it and its goals
2. Discuss the real-world benefits of DWaaS
3. Discuss some of the coolest features in a DWaaS solution as exemplified by the Snowflake Elastic Data Warehouse.

### The World's Fastest Database for Operational Analytics—*Lejo Jacob, MemSQL*

The proliferation of streaming analytics and instant decisions to power dynamic applications, along with the rise of predictive analytics, machine learning, and operationalized artificial intelligence, have introduced a requirement for a new type of database workload: operational analytics. MemSQL's fast SQL database allows you to accelerate time to insight and simplify operations. With MemSQL, customers get a single database that can easily support real-time decision making and insight-driven customer experiences, combining massively scalable ingest, the ability to scan billions of rows per second across thousands of users, and a familiar relational structure for user- and developer-friendly access.

# Microsoft Azure Seminar

*Kellyn Pot'Vin-Gorman, Microsoft*

### Azure Fundamentals

The cloud offers opportunities to get to market faster and for a lower initial investment than on-premises, but to see it in action is incredibly beneficial. This session is a 101 course on the underlying infrastructure and design on the Azure cloud and portal. Key concepts, security, and best practices will be covered to help the attendee become familiar with this incredible enterprise-level cloud.

1. Become familiar with the Azure Portal.
2. Learn how the Azure Cloud works with the Azure Active Directory and security layers.
3. See demonstrations of how resources and services are architected, along with pricing and sizing tier best practices.

### The Modern Data Ecosystem in the Cloud

As data grows, the technical world is shrinking. This session helps explain the Azure Data Platform—what is part of this ecosystem and how it works together. This session will demonstrate how data coming from various sources can be consumed and transformed, and then have analysis built upon it to offer value to customers and consumers.

1. Learn key products and concepts in Azure.
2. See an actual environment and working examples of data to value.
3. View demonstrations on predictive analysis, machine learning, and advanced analytics.

### DevOps or Decoys: Azure Automation

The term DevOps can be overwhelming. But the term is really just a decoy for common sense. We already have many of the skills that are used in DevOps practices. Using the real-life use case of Higher Education customers at Microsoft, this session takes the attendee through how the challenges of going to the Azure cloud can be simplified. It will cover multiple ways to automate deployments of infrastructure, architecture, code, and data migrations to create complete A–Z solutions that make it easier to use the cloud more quickly.

1. Challenges of deploying to the cloud
2. How existing skills and technology can be employed to simplify the cloud
3. Tricks to un-silo teams to create a more harmonious union in IT

# ADW Cloud Hands-On Lab

*Sergio Castro, Rajneesh Tiwary, and Haimanti Basu, Oracle Corporation*

### Overview of Oracle Autonomous Database Warehouse Cloud Powered by Machine Learning

An autonomous database is a cloud database that uses machine learning to eliminate the human labor associated with database tuning, security, backups, updates, and other routine management tasks traditionally performed by

database administrators.

Oracle Autonomous Data Warehouse (ADW) is built around the market-leading database and comes with specific, fully-automated data warehouse features that deliver outstanding query performance. This lab walks you through the steps to get started using the Oracle ADW.

1. How to deploy a data warehouse and run workloads with only a few clicks

2. How to scale storage and compute resources independently with database services running with no downtime

3. How to run faster analytics on diverse data sets with a fast visual predictive pattern

### Oracle ADW Cloud & ML Hands-on Lab

This hands-on lab is for all database administrators, developers, and data scientists who want to learn how a self-driving, self-repairing, and self-securing data warehouse works on high-performance infrastructure.

## Oracle DB Security Seminar

*Dan Morgan, DBSecWorx*

Year after year Oracle databases that should be secure are violated. In some cases data is stolen, in others data is altered, and in still others the database is not the target but is a means to an end. In almost all cases there is no excuse. No excuse … but there is an explanation. Oracle professionals do not receive training on the threats and how to detect and prevent them. This, of course, is not helped by the fact that business and government management focuses money and effort on expensive things that don't solve the underlying problem and on training to recognize irrelevant threats (no Oracle database was ever protected by someone not clicking on an email from a Nigerian prince).

This Oracle Database security seminar, taught by Oracle ACE Director alumnus Daniel Morgan, will be hands on and will focus on teaching through intensive lecture, demonstration, and live labs. In the class you will learn how to successfully break into an Oracle database. After all, if you don't know how to break in, you don't know how to stop someone else from breaking in. We won't just talk about SQL injection attacks, we will commit one. We won't just talk about using an Oracle database as an attack tool … we will perform a white hat attack across the internet on a target of your choice. We won't just talk about creating proxy users … we will do it. And you will learn, hands-on, how to prevent someone from accessing a database even when they have a valid user ID and password.

This class will not be a sales job promoting any company, any product, any license, or any service. Come join us and bring notepaper, a pen, and (if you have one) a laptop with Oracle Database 11*g*R2 or above installed. At the end of the class you will leave with something you didn't bring in the door … knowledge of things you can do to make your data and your databases more secure.

## Distributed SQL Seminar

*Sid Choudhury & Bryn Llewellyn, YugaByte*

### Distributed SQL Databases Deconstructed

SQL is a popular database language for modern applications because of its flexibility in modeling workloads and how widely understood it is by developers. However, most modern applications running in the cloud require fault tolerance, the ability to scale out, and geographic data distribution. These attributes are hard to achieve with traditional monolithic SQL databases, creating the need for distributed SQL databases.

Google's Cloud Spanner is arguably the world's first truly distributed SQL database that now powers everything from Adwords to Gmail. Given its fully decentralized architecture, it delivers higher performance and availability for geo-distributed SQL workloads than even specialized transactional databases such as Amazon Aurora and Google Percolator. Fortunately, there are now a number of open-source derivatives of Google Cloud Spanner, such as YugaByte DB, CockroachDB, and TiDB. This session will focus on exploring the common architectural paradigms that these databases are built on and evaluating the claims made by the different systems.

### How YugaByte DB Works as a Distributed PostgreSQL Database at Scale

YugaByte DB is an open-source, high-performance, cloud-native distributed SQL database. Its SQL API is wire compatible with the SQL language of PostgreSQL. Additionally, this API is built by reusing PostgreSQL's query layer as is. However, the API runs on top of DocDB, YugaByte DB's Google Spanner–inspired distributed document store, as opposed to PostgreSQL's built-in monolithic storage engine. This architectural enhancement allows YugaByte DB to function as a distributed PostgreSQL that is not only horizontally scalable but also highly resilient to infrastructure failures such as node, disk, zone, and region failures common in modern cloud environments. This session will describe how PostgreSQL query layer is reused in YugaByte DB in the context of system catalog management as well as SQL, DDL, and DML operations.

### Distributed SQL Hands-On Lab

In this hands-on lab, attendees will install a local cluster of YugaByte DB on their laptop, run PostgreSQL-based DDL/DMLs on the cluster, and see horizontal scalability and high fault tolerance in action.