

Data Warehousing for Everybody: Oracle Autonomous Data Warehouse Cloud

ORACLE
OPEN
WORLD

October 1–5, 2017
SAN FRANCISCO, CA

George Lumpkin, Vice President, Product Management
Yasin Baskan, Sr Principal Product Manager

Oracle Database Server Technology
October 2, 2017

ORACLE®


Copyright © 2017, Oracle and/or its affiliates. All rights reserved. | Confidential – Oracle Internal/Restricted/Highly Restricted

Safe Harbor Statement

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, timing and price of any features or functionality described for Oracle's products may change and remains at the sole discretion of Oracle Corporation. Fees apply for new Database product offerings.

Oracle's **Vision** for Autonomous Database

- **Self-Driving**
 - User defines service levels, database makes them happen
- **Self-Securing**
 - Protection from both external attacks and malicious internal users
- **Self-Repairing**
 - Automated protection from all downtime



Ingredients of an Autonomous Database

- **Oracle Database 18c**

- Dozens of automated database features

+


- **Oracle Cloud infrastructure**

- Single, standard platform as focus for self-driving automation

+

- **Oracle Autonomous Database Cloud tooling**

- Automation for optimal configuration for each workload



One Autonomous Database – Optimized by Use Case

2017

**Data
Warehousing**

2018

**Enterprise
OLTP,
Mixed
Workloads**

Now

**Departments,
Developers**

Oracle Autonomous Database

Introducing: **Autonomous Data Warehouse Cloud**

- **Easy**
 - Automated management
 - Automated tuning: **Simply load data and run**
- **Fast**
 - Based on Exadata technology
- **Elastic**
 - Instant scaling of compute or storage with no downtime



Expected CY 2017

Getting Started with Autonomous Data Warehouse Cloud

- Provisioning requires only 5 simple questions:
 - Database name?
 - Which data center?
 - How many CPU's?
 - How many TB's?
 - Admin password?
- New service created in <30 seconds (regardless of size)
 - Ready to connect via sqlnet

The screenshot shows the Oracle Cloud My Services interface for creating an Autonomous Data Warehouse Cloud service. The page is titled "ORACLE CLOUD My Services" and "Autonomous Data Warehouse Cloud Create Service". It features a progress bar with "Service" and "Confirm" steps, and a "Next" button. The form is divided into two sections: "Details" and "Configuration".

Details:

- * Database Name:
- Description:
- Notification Email:

Configuration:

- * Region:
- * CPU Core Count (OCPU's):
- * Storage Capacity (TB):
- * Administrator Password:
- * Confirm Administrator Password:
- Object Store URL:
- Object Store Username:
- Object Store Password:

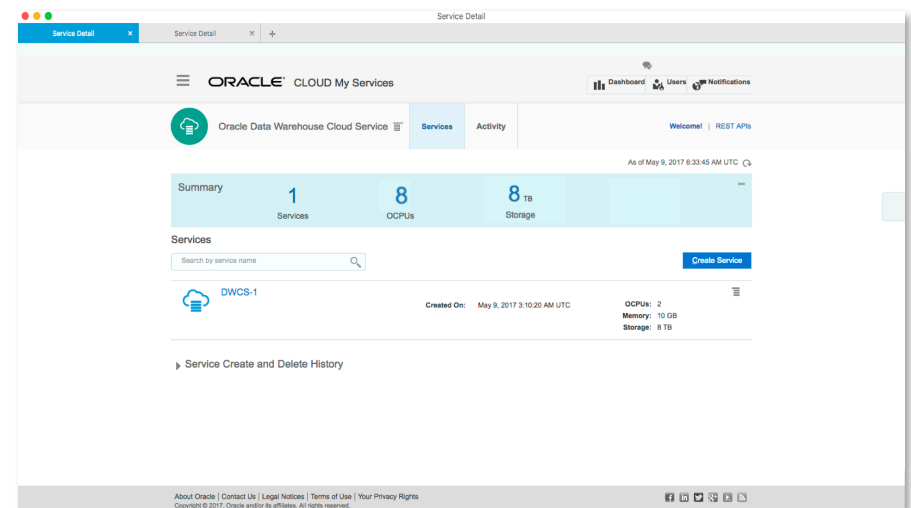
Demo: Provisioning

What didn't you just see?

- No decisions for:
 - Details of the database software
 - Configuration of hardware
 - Characteristics of the database
 - Architecture for backups and availability
- **All of this (and more) is automatically configured**

Automated management

- Oracle automates end-to-end management of data warehouse
 - Provisioning new database instances
 - Growing/shrinking storage and/or compute
 - Patching and upgrades
 - Backup and recovery
- Full lifecycle managed using ADWC Service Console



Automated vs. Autonomous

Automated

- The car simplifies operations by automating tasks:
 - Cruise control
 - Emergency stopping
 - Warnings for lane changes
- The database simplifies operations:
 - Automatic storage management, automatic storage management, ...
 - Dozens of other features


Autonomous

- The car drives itself
 - No need to use the steering wheel or brake.
 - Simply tell the car where you are going.
- The database manages itself
 - All features automatically implemented
 - Simply tell the database your goals

What does Autonomous Database mean for the DBA?

Less time on Administration

- Less time on infrastructure
- Less time on patching, upgrades
- Less time on ensuring availability
- Less time on tuning



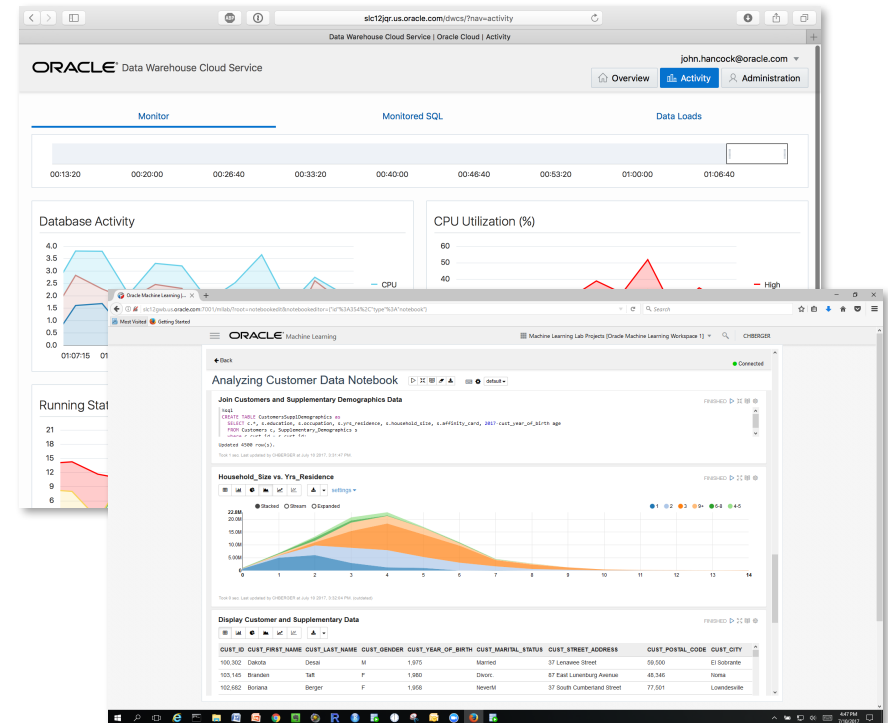
More time on Innovation

- More time on database design
- More time on developing new apps
- More time on data analytics
- More time on securing data

Challenge: There are more data management tasks than humans to do the work

Automated Tuning

- “Load and go”
 - Define tables, load data, run queries
 - No tuning
 - No special database expertise required
 - Good performance out of the box
- Query using any business analytics tool or cloud service
 - Built-in SQL notebook also included



Hippocratic Oath of Autonomous Performance

“Do no harm”

- Performance features are carefully selected:
 - No performance degradation
 - No significant overhead

Autonomous Data Warehouse Cloud: Performance Proof Point

- Evaluation of real customer workloads and standard benchmarks
 - Performance As-Is versus ADWC
 - Run on identical hardware
- ADWC out-of-the-box
 - Optimized, pre-configured DB setup
- Schema as provided by customer
 - Indexes, Partitioning, manual DOP

	Oracle ADWC	Customer Schema
Customer #1	39	34
Customer #2	52	53
Customer #3	1032	911
Customer #4	2138	3051
Customer #5	132	196
Customer #6	1478	1831
Customer #7	2551	5137
MSTR	1163	1504
SSB	232	209

Autonomous Data Warehouse Cloud: Reliability

- Enterprise – 99.95%
 - Fault-tolerant, highly-available Exadata infrastructure
 - Triple-mirrored disks for disk failures
 - Real Application Clusters for compute node failures
 - And more ...
 - Automatic backups
 - Point in time recovery to anytime in last 30 days
- Mission Critical – 99.995%
 - Standby database

Instant Elasticity

- Customer specifies number of database cores (in OCPU) and database size (in TB)
 - CPU's and storage are independent
 - Databases can grow and shrink in either dimension
 - Pricing based upon CPU (\$/CPU/hour) and Storage (\$/TB/Month)
- Examples:
 - Adding new storage or compute capacity is just a click on the cloud console
 - CPU capacity can be added on-demand for end-of-quarter processing, then reduced after the processing is completed
 - CPU capacity can be set to zero, while storage remains in place (for example, to reduce costs for databases which are not used during weekends / nights)

Instant Elasticity: **Pay for Exactly What you Use**

- Size the DW to the exact number of OPCU's and TB's required
 - Not constrained by fixed building blocks
- Scale the DW on demand
 - Independently scale compute or storage
 - Resizing occurs instantly, fully online
- Shut off idle compute save money
 - Restart instantly

Easy Migration to Autonomous Data Warehouse Cloud

Migration of existing Oracle systems

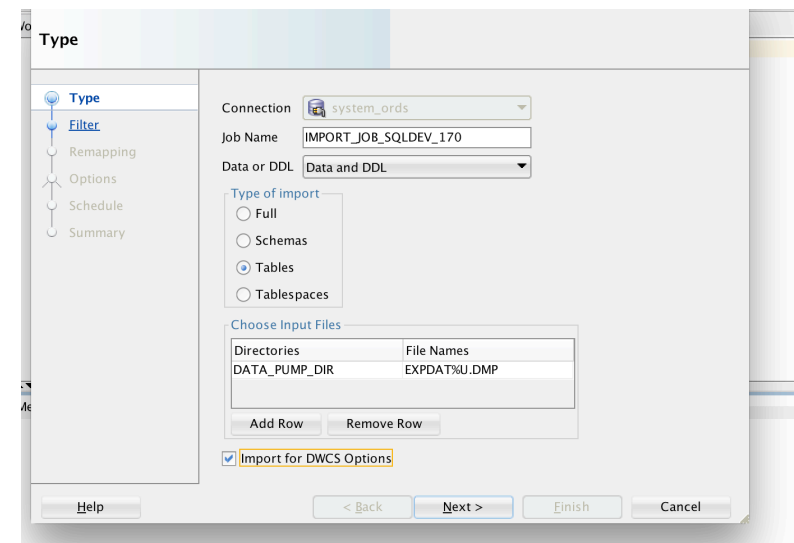
- Data pump import using ADWC-aware settings
 - Ensures compatible data import
 - Violating data structures will be ignored or converted

Migration of Redshift systems

- End-to-end migration using migration workbench
 - Convert and create data warehouse schema
 - Unload Redshift data into S3
 - Load data into ADWC

Migration of other systems

- Migration workbench assistance or manual



Supported by a rich Data Warehouse Ecosystem

Oracle Autonomous Data Warehouse Cloud supports:

- **Existing tools** running on-premise or in Oracle Cloud
 - Oracle BI and data-integration tools
 - 3rd party BI tools
 - 3rd party data-integration tools
- Connectivity via SQL*Net

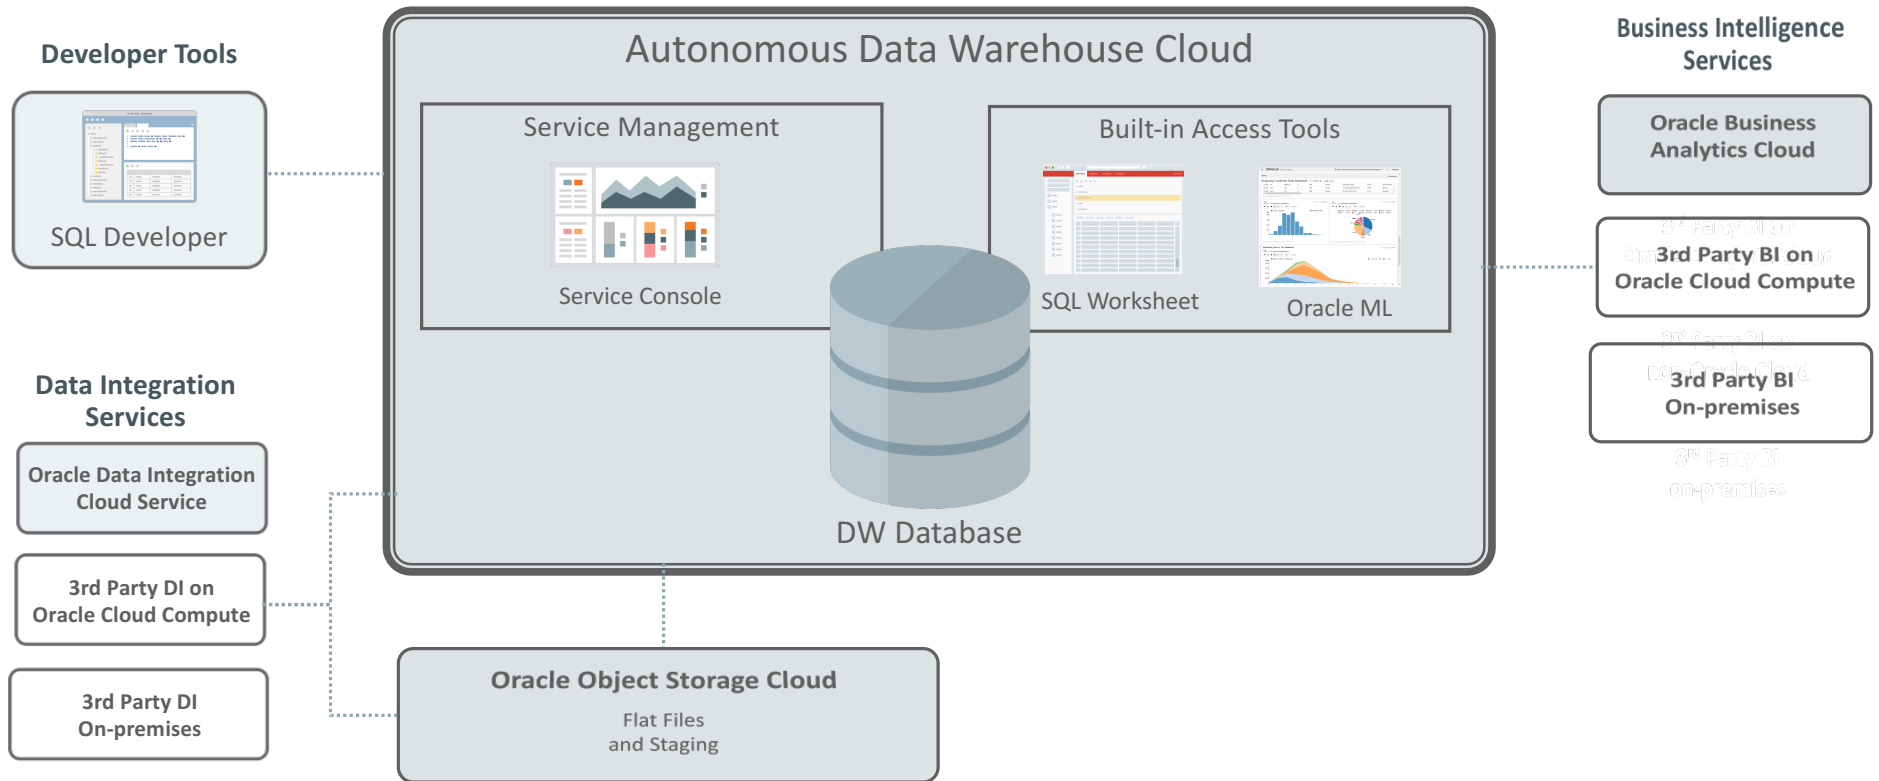
Oracle Cloud Services

- Oracle Analytics Cloud
- Oracle Data Integration Platform Cloud
- and others...

Autonomous Data Warehouse Cloud: Support From Strategic Partners



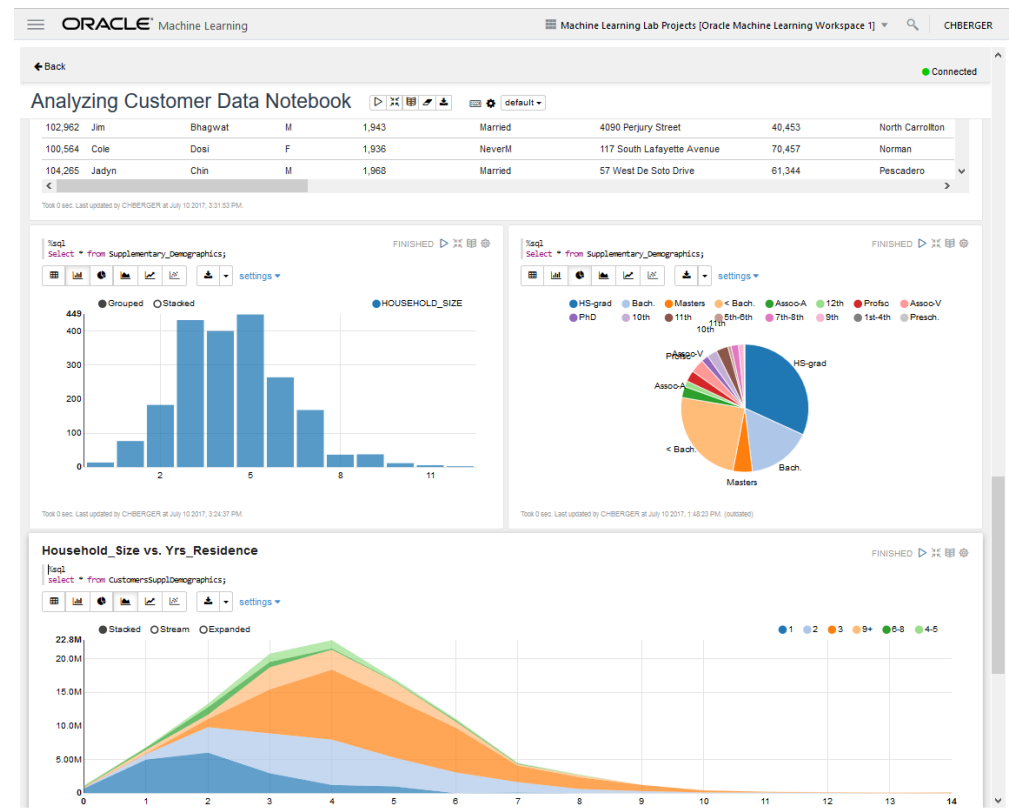
Autonomous Data Warehouse Cloud: Architecture



on-premises
ORACLE

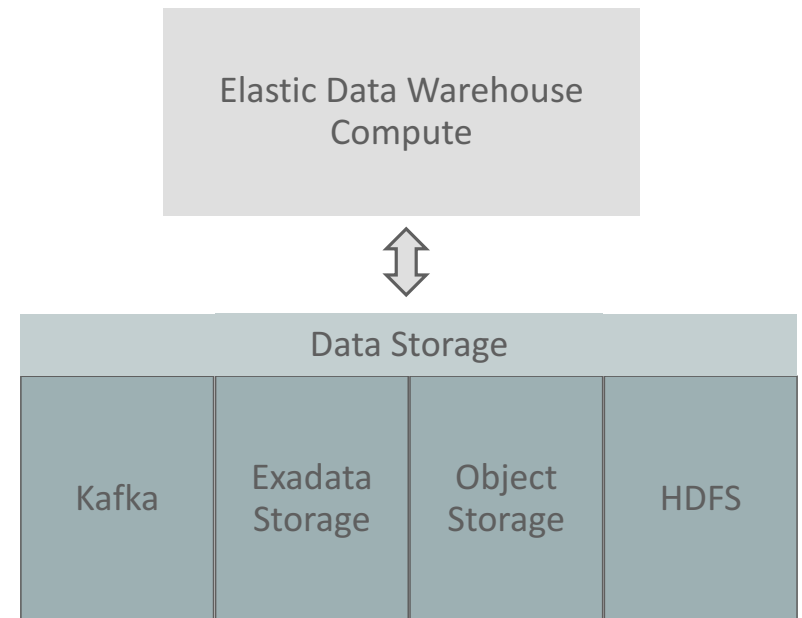
OracleML: Built-in notebook

- Collaborative UI for data scientists
 - Easy access to shared notebooks, templates, permissions, scheduler, etc.
- Based on Apache Zeppelin
- Roadmap: Common UI for data scientists across multiple services



Queries across Cloud Storage and more

- External tables can be created on data in object stores
 - Oracle Object Store or AWS S3
 - Any supported Oracle-loader file format
 - Or any Hadoop file format (e.g. Parquet) – coming soon
 - Big Data SQL integration – coming soon
- Seamless queries over object store
- Seamless queries across object store and database



More on Autonomous Database

Step-by-Step Guide to Oracle Autonomous Database Cloud for Data Warehousing

–Wednesday, Oct 04, 12:00 pm - 12:45 pm | Moscone West - Room 3010

Preview of Oracle Autonomous Database

–Wednesday, Oct 04, 3:30 pm - 4:15 pm | Moscone West - Room 3014

ORACLE®