



Autonomous Data Warehouse Cloud

Updated: June 12, 2018

Lab 3: Working with Database Services and Sample Data Sets

In this lab you will explore the sample data sets and experiment with the choices of database services that come with your ADWC instance.

Background

Autonomous Data Warehouse Cloud provides three database services that you can choose when connecting to your database. These are named as **HIGH**, **MEDIUM**, and **LOW** services and provide different levels of performance and concurrency. As a user you need to pick the database service based on your performance and concurrency requirements.

- **HIGH** database service provides the maximum amount of resources for a query, this also means the number of concurrent queries you can run in this service will not be as much as the other services.
- **MEDIUM** database service provides multiple compute and IO resources for a query. This service also provides more concurrency compared to the HIGH database service.
- **LOW** database service provides the least amount of resources for a query, this also means the number of concurrent queries you can run in this service will be higher than the other services.

The lab will use the LOW and HIGH database services to understand the performance differences between them. The demo will run queries on sample data sets provided out of the box with ADWC.

ADWC provides two sample data sets - Oracle Sales History (SH) sample schema and the Star Schema Benchmark (SSB) schema. You will run queries on the SSB data set which is a 1TB data set with one fact table containing about 6 billion rows, and several dimension tables.

Objectives

- Understand the different levels of ADWC database service (HIGH, MEDIUM, LOW)
- Explore the Star Schema Benchmark (SSB) and Sales History (SH) sample data sets
- Observe the effect on performance and concurrency when using different database service levels
- Explore the query results caching

Required Artifacts

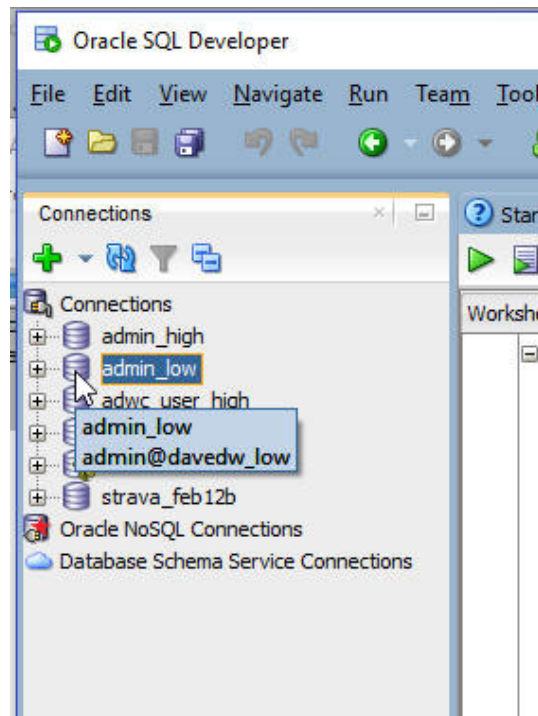
- Access to the lab VM.
 - Locate the lab VM's **IP Address**, **User Name** and **Password**
- The following artifacts of the lab VM are utilized:
 - Installed software: **Oracle SQL Developer**

Lab Steps

Step 1: Connect and Query with the LOW database service

For this part of the lab to be meaningful, your ADWC instance should be configured with at least 2 OCPUs. With 1 OCPU, you will not see much difference in performance.

- Sign in to the lab VM using the credentials provided to you by the instructor.
- Start SQL Developer from the lab VM.
- Using SQL Developer connect to the **admin_low** database connection that you previously created. You can do this by expanding the list of connections and double-clicking on the **admin_low** connection.

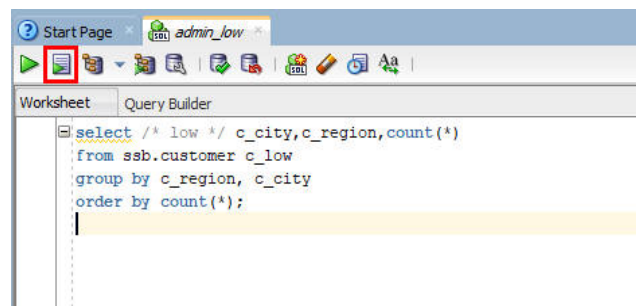


- Copy and paste the below SQL to SQL Developer worksheet. Take a moment to examine the script.

```
select /* low */ c_city, c_region, count(*)  
from ssb.customer c_low  
group by c_city, c_region  
order by count(*);
```

- Click **Run Script** to run it.

Ensure that you click "Run Script" so that all the rows get displayed on the screen.



- Make a note of the response time. In the following example the query with the LOW database service finished in around **7.5** seconds.

Start Page admin_low admin_high

7.55499983 seconds

Worksheet Query Builder

```
select /* low */ c_city,c_region,count(*)
from ssb.customer c_low
group by c_city,c_region
order by count(*);
```

Script Output x Query Result x

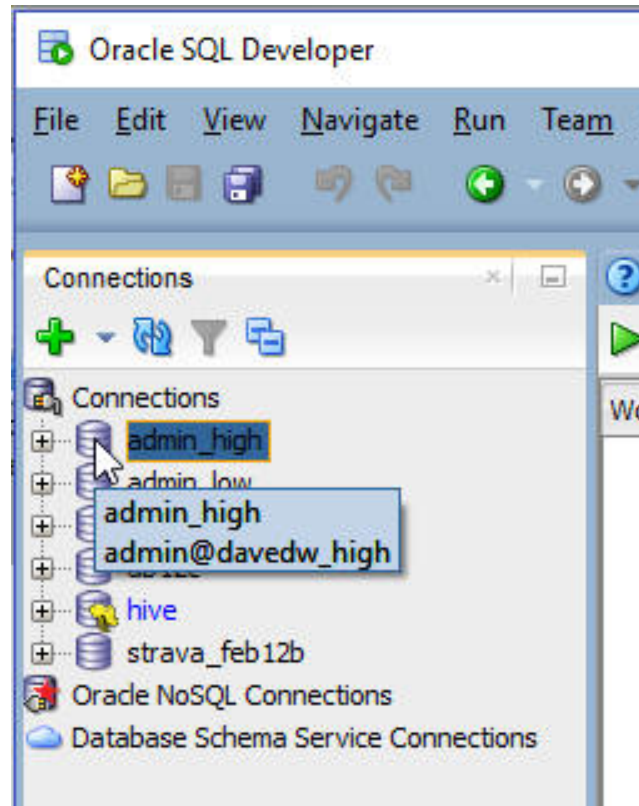
Task completed in 7.555 seconds

IRAQ	1	MIDDLE EAST	120557
IRAN	0	MIDDLE EAST	120550
ARGENTINA	1	AMERICA	120550
PERU	4	AMERICA	120554
EGYPT	8	MIDDLE EAST	120573
BRAZIL	3	AMERICA	120573
INDONESIA	1	ASIA	120573
<hr/>			
C_CITY	C_REGION		COUNT(*)
<hr/>			
JORDAN	6	MIDDLE EAST	120577
UNITED STATES	1	AMERICA	120592
EGYPT	2	MIDDLE EAST	120599
EGYPT	7	MIDDLE EAST	120601
CHINA	5	ASIA	120655
ETHIOPIA	3	AFRICA	120725
CANADA	6	AMERICA	120784
EGYPT	4	MIDDLE EAST	120810

250 rows selected.

Step 2: Connect and Query with the HIGH Database Service

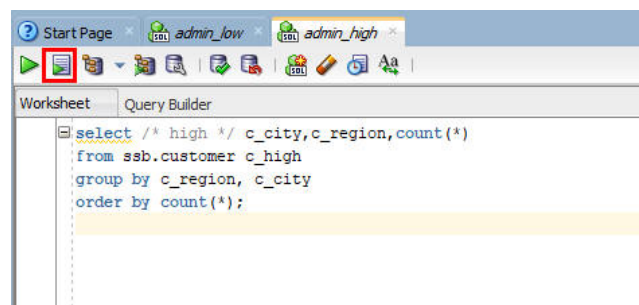
- Using SQL Developer connect to the **admin_high** database connection that was previously created. You can do this by expanding the list of connections and double-clicking on the **admin_high** connection.



- Copy and paste the below SQL to SQL Developer worksheet. This query is basically the same as the previous one except the comment after the select keyword and the table alias.

```
select /* high */ c_city, c_region, count(*)  
from ssb.customer c_high  
group by c_city, c_region  
order by count(*);
```

- Click **Run Script** to run it.



- Make a note of the response time. In the following example the query with the HIGH database service finished in around 3.5 seconds.

The screenshot shows a database interface with two tabs: 'admin_low' and 'admin_high'. The 'admin_high' tab is active. In the 'Query Builder' window, the following SQL query is entered:

```
select /* high */ c_city,c_region,count(*)
from ssb.customer c_high
group by c_city,c_region
order by count(*);
```

The 'Script Output' window shows the execution results:

Task completed in 3.628 seconds

IRAN	0	MIDDLE EAST	120550
PERU	4	AMERICA	120554
INDONESIA1	1	ASIA	120573
BRAZIL	3	AMERICA	120573
EGYPT	8	MIDDLE EAST	120573
C_CITY	C_REGION	COUNT(+)	

JORDAN	6	MIDDLE EAST	120577
UNITED ST1		AMERICA	120592
EGYPT	2	MIDDLE EAST	120599
EGYPT	7	MIDDLE EAST	120601
CHINA	5	ASIA	120655
ETHIOPIA	3	AFRICA	120725
CANADA	6	AMERICA	120784
EGYPT	4	MIDDLE EAST	120810

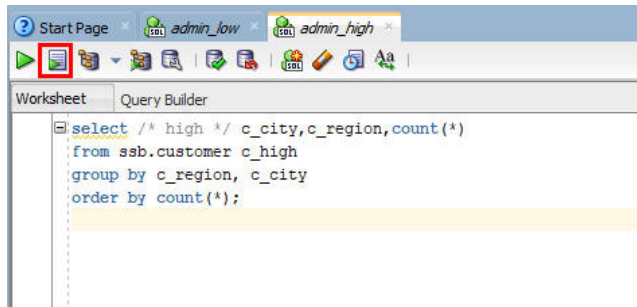
250 rows selected.

- As you noticed, the same query runs faster when connected to the HIGH service as the query can use more resources and run faster compared to a query running in the LOW database service. As you scale up the compute capacity of your ADWC service you will realize that the queries will get faster in the HIGH database service.

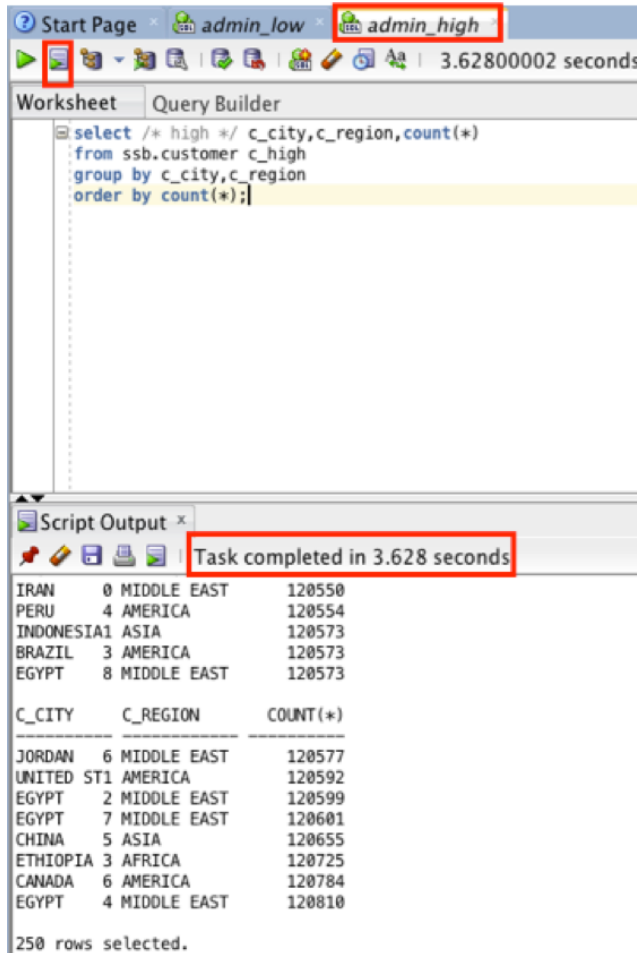
Step 3: Explore Query Results Caching

ADWC also caches the results of a query for you. If you run the same queries again you will see that they will run much faster.

- Run the previous query you just ran using the HIGH service connection.
- Click the **Run Script** button to run it again.



- Make a note of the response time. This time you will see that the query finishes in less than a second.



- You may explore running additional queries on the sample schema from the ADWC documentation [here](#).
- You've completed this lab. Please proceed to the next lab.