

Oracle Database/ADW 

VS

Snowflake Data Cloud 



Explorando a eficiência e a versatilidade de cada um



Keep in touch



linktr.ee/loredata

E-mail: faust@pythian.com or franky@loredata.com.br

Blog: <http://loredata.com.br/blog>

Facebook: <https://facebook.com/08Franky.Weber>

Instagram: <https://www.instagram.com/frankyweber/>

Twitter: <https://twitter.com/frankyweber>

LinkedIn: <https://linkedin.com/in/frankyweber/en>

Oracle ACE: https://ace.oracle.com/pls/apex/r/ace_program/oracle-aces/ace?ace_id=593



Oracle Certified Expert, Oracle Database 12c: Performance Management and Tuning

Issued by [Oracle](#)



Oracle Exadata Database Machine and Cloud Service 2017 Certified Implementation Specialist

Issued by [Oracle](#)



Oracle Database 12c Administrator Certified Professional

Issued by [Oracle](#)



Oracle Linux 6 Certified Implementation Specialist

Issued by [Oracle](#)



Oracle Real Application Clusters 12c Certified Implementation Specialist

Issued by [Oracle](#)



Oracle Database SQL Certified Expert

Issued by [Oracle](#)



Oracle Database 11g Administrator Certified Associate

Issued by [Oracle](#)



Oracle Autonomous Database Cloud 2019 Certified Specialist
Oracle

FRANKY WEBER FAUST

faust@pythian.com

- Lead Database Consultant at Pythian
- 32 years old
- Based in Blumenau/Brazil
- Writer at Pythian and Lore Data Blog
- Speaker at conferences around the world
- High Availability specialist
- Performance researcher
- Exadata, RAC, DataGuard, GoldenGate



Oracle ACE Associate



Diamante



AGGRANDIZE

COMMVault™ 



TD SYNnex

Platina



DISCOVER

Ouro



VERTICA
by opentext™

Prata

TRACES



ROX

We take care
of your data

Apoio

FIAP

GRUPO
POSEIDON
DIGITAL

Pythian

love your  data

Agenda

- Mike's story
- Concepts
- Key differentiators

The story of Mike

Data, unicorns and rainbows



Not available for download

Concepts

Architecture, storage and more



Oracle ADW vs Snowflake: Architecture




VCN

Cloud Services

AuthC & AuthZ | NACL | NSG | OCI Console | DB Tools

Exadata Database Machine

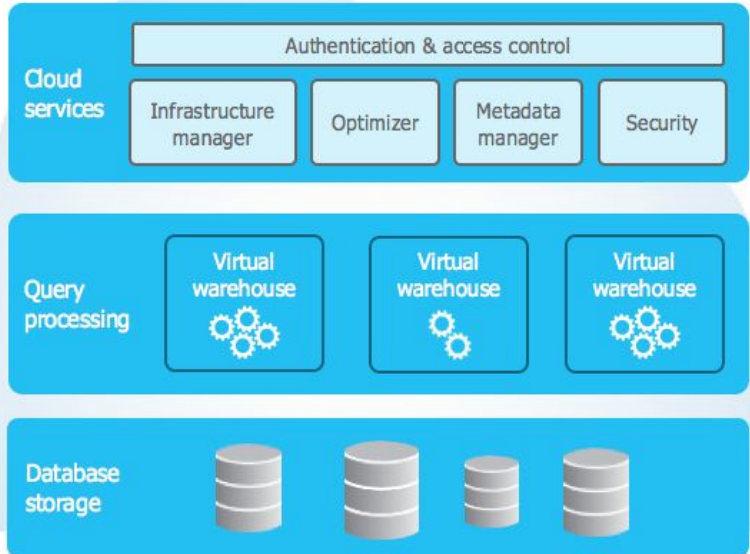
Exadata Database Servers (dbnodes)



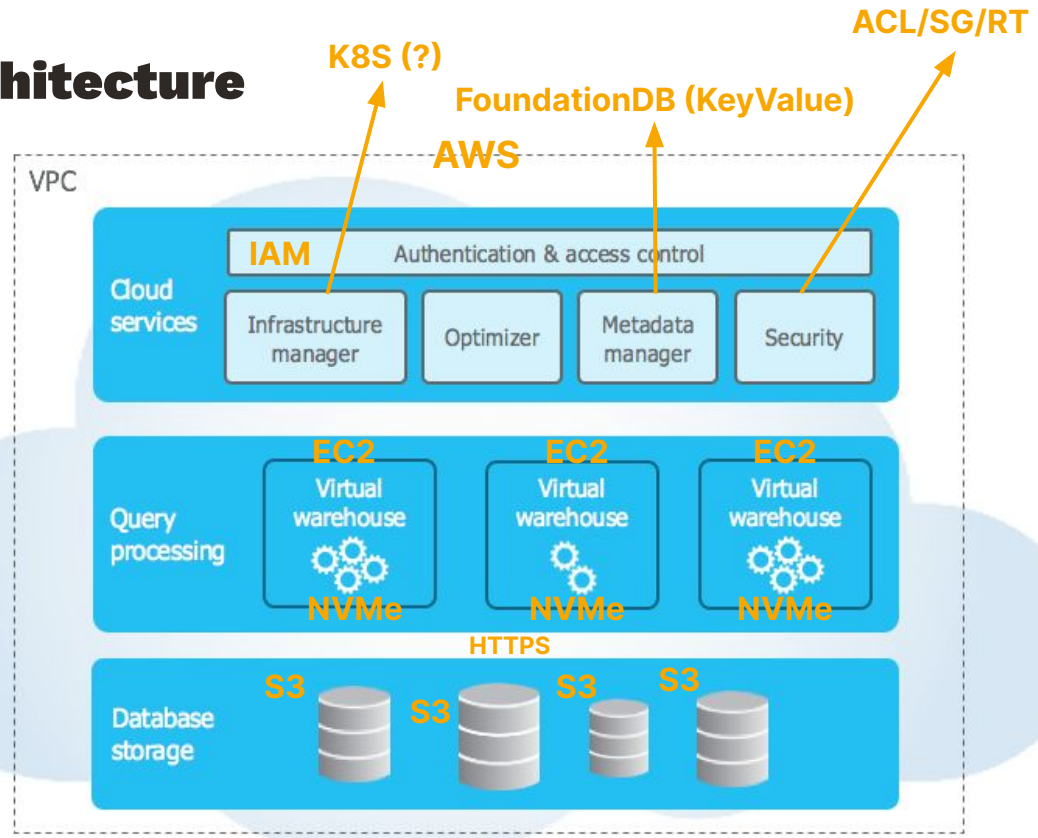
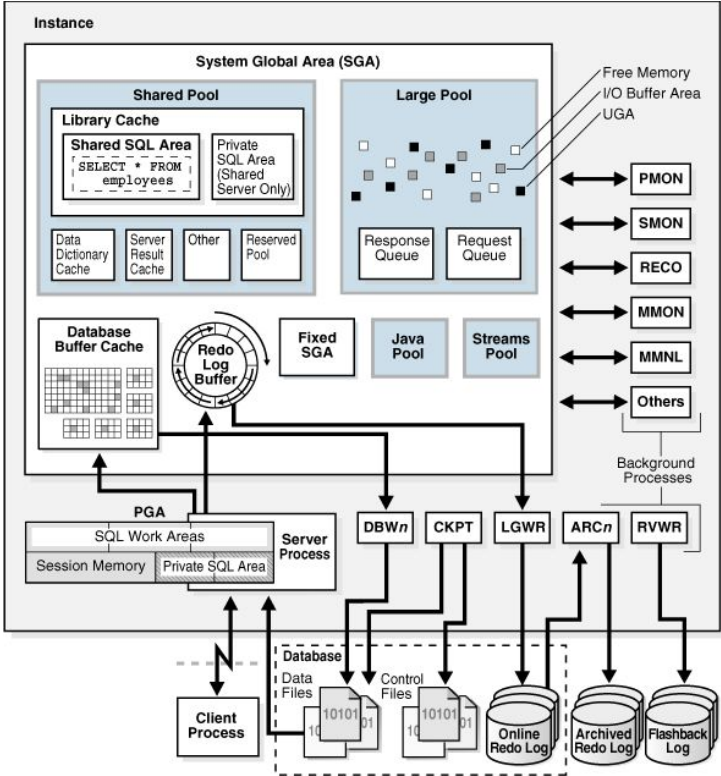
Exadata Storage Servers (cells)



VPC



Oracle ADW vs Snowflake: Architecture



Snowflake: Micro Partitions

Snowflake Table

ID	AMOUNT	PRICE	SELLER
1	100	600	D
32	13	302	F
436	17	407	K
328	22	913	W
402	30	409	G
727	47	176	V
1001	39	484	X
4007	79	507	P
1023	44	712	W
4716	104	496	X
793	48	703	F
441	92	8009	D
532	109	4007	F

MP1 ROWS 1-4

ID	1 - 32 - 436 - 328	1 - 436
AMOUNT	100 - 13 - 17 - 22	13 - 100
PRICE	600 - 302 - 407 - 913	302 - 913
SELLER	D - F - K - W	D - W

S3

MP2 ROWS 5-8

ID	402 - 727 - 1001 - 4007	402 - 4007
AMOUNT	30 - 47 - 39 - 79	30 - 79
PRICE	409 - 176 - 484 - 507	176 - 507
SELLER	G - V - X - P	G - X


S3

MP3 ROWS 9-13

ID	1023 - 4716 - 793 - 441 - 532	441 - 4716
AMOUNT	44 - 104 - 48 - 92 - 109	44 - 109
PRICE	712 - 496 - 703 - 8009 - 4007	496 - 8009
SELLER	W - X - F - D - F	D - X

S3

Oracle ADW: Hybrid Columnar Compression




Row-Store Physical Layout

Row Database stores row values together

EmpNo	EName	Job	Mgr	HireDate
7369	SMITH	CLERK	7902	17/12/1980
7499	ALLEN	SALESMAN	7698	20/02/1981
7521	WARD	SALESMAN	7698	22/02/1981
7566	JONES	MANAGER	7839	2/04/1981
7654	MARTIN	SALESMAN	7698	28/09/1981
7698	BLAKE	MANAGER	7839	1/05/1981
7782	CLARK	MANAGER	7839	9/06/1981

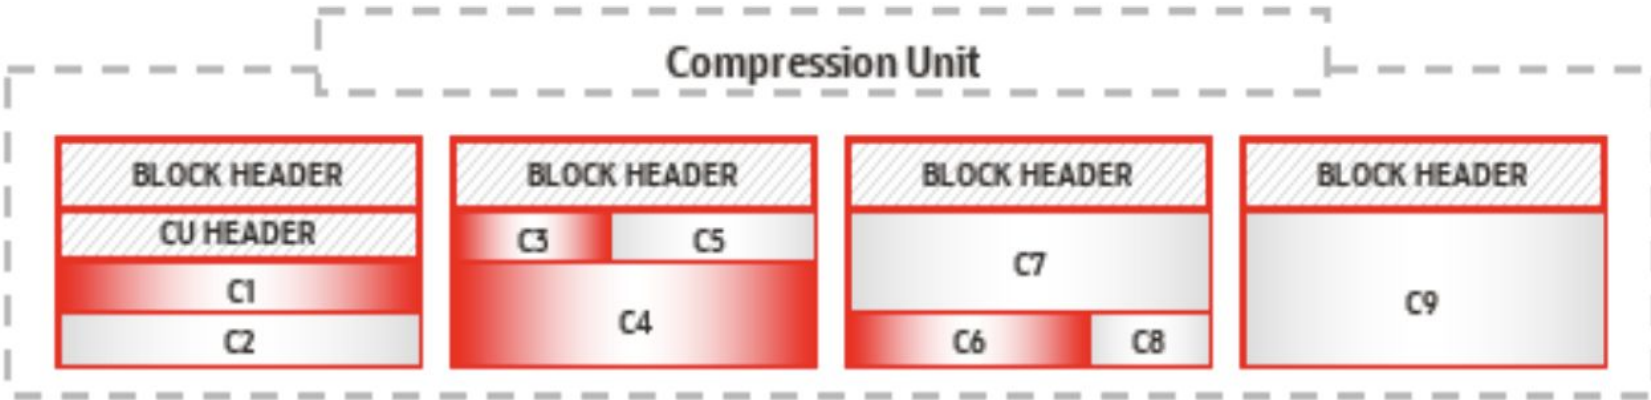
Logical Schema



Column Store physical layout

Column Database stores column values together

Oracle ADW: Hybrid Columnar Compression



Snowflake: Query Performance

The screenshot displays the Snowflake web interface. On the left is a navigation sidebar with options like Worksheets, Dashboards, Apps, Data, Databases (selected), Private Sharing, Provider Studio, Governance, Marketplace, Activity, Admin, and Help & Support. The main area shows the hierarchy: SNOWFLAKE > SNOWFLAKE_SAMPLE_DATA > INFORMATION_SCHEMA > TPCDS_SF100TCL > CATALOG_RETURNS. The table 'CATALOG_RETURNS' is selected, and its 'Data Preview' is shown. The table has 10 rows and 7 columns: CR_RETURNED_DATE_SK, CR_RETURNED_TIME_SK, CR_ITEM_SK, CR_REFUNDED_CUSTOMER_SK, CR_REFUNDED_CDEMO_SK, and CR_I. A 'Load Data' button is visible in the top right.

SNOWFLAKE_SAMPLE_DATA / TPCDS_SF100TCL / **CATALOG_RETURNS** Load Data

Table 1 year ago 14.4B 770.9GB

Table Details Columns **Data Preview** Copy History

• COMPUTE_WH 100 of 14.4B Rows • Updated just now C

	CR_RETURNED_DATE_SK	CR_RETURNED_TIME_SK	CR_ITEM_SK	CR_REFUNDED_CUSTOMER_SK	CR_REFUNDED_CDEMO_SK	CR_I
1	2,452,734	73,116	313,659	10,044,524	277,807	
2	2,452,734	56,166	313,659	76,072,524	1,101,409	
3	2,452,734	24,959	313,659	78,127,221	1,667,619	
4	2,452,734	15,433	313,659	98,111,486	69,765	
5	2,452,734	41,682	313,659	4,951,535	1,492,994	
6	2,452,734	42,721	313,659	58,543,721	1,469,627	
7	2,452,734	78,093	313,659	58,973,377	1,670,565	
8	2,452,734	46,163	313,659	81,596,861	1,242,482	
9	2,452,734	53,370	313,659	14,766,188	958,454	
10	2,452,734	61,950	313,659	31,382,315	684,551	

Databases Worksheets



ACCOUNTADMIN COMPUTE_WH

Share



Pinned (0)

No pinned objects

All Objects

SNOWFLAKE

SNOWFLAKE_SAMPLE_DATA

SNOWFLAKE_SAMPLE_DATA.TPCDS_SF100TCL Settings

Latest Version

```

1 use warehouse COMPUTE_WH;
2 use database SNOWFLAKE_SAMPLE_DATA;
3 use schema TPCDS_SF100TCL;
4
5 select CR_WAREHOUSE_SK, sum(CR_RETURN_AMOUNT) from SNOWFLAKE_SAMPLE_DATA.TPCDS_SF100TCL.CATALOG_RETURNS group by CR_WAREHOUSE_SK order by CR_WAREHOUSE_SK;

```

Results

Chart



Snowflake executing... 1m 8s

Start Time	Jun 21 at 10:00 AM
ID	01ad1dac-0604-99c6-0000-006a9787f02d
Warehouse	COMPUTE_WH(X-Small)
Produced Rows	—

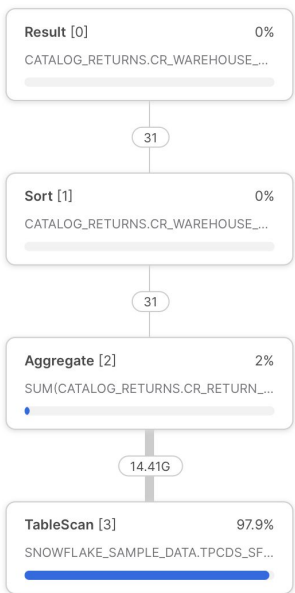
Cancel

- Worksheets
- Dashboards
- Apps
- Data
- Marketplace
- Activity
- Query History**
- Copy History
- Task History
- Admin
- Help & Support

Query History

Status **All** User **FRANKYWEBER** Last 14 days Filters 8 Queries Columns

SQL TEXT	QUERY ID	STATUS	USER	WAREHOUSE	DURATION	STARTED
select * from IDENTIFIER('SNOWFLAKE_SA...	01ad1dad-0604-99c6-0000-006a9787f0...	Success	FRANKYWEBER	—	94ms	6/21/2023,
select CR_WAREHOUSE_SK, sum(CR_RETURN_A...	01ad1dac-0604-99c6-0000-006a9787f02d	Running	FRANKYWEBER	COMPUTE_WH	3m 52s	6/21/2023,
use schema TPCDS_SF100TCL;	01ad1dac-0604-995f-006a-97870001003e	Success	FRANKYWEBER	—	62ms	6/21/2023,
use database SNOWFLAKE_SAMPLE_DATA;	01ad1dac-0604-99c6-0000-006a9787f029	Success	FRANKYWEBER	—	135ms	6/21/2023,
use warehouse COMPUTE_WH;	01ad1dac-0604-991e-0000-006a9787c039	Success	FRANKYWEBER	—	63ms	6/21/2023,
SELECT system\$GET_NPS_FEEDBACK_TIMESTAM...	01ad1da7-0604-991d-0000-006a978790...	Success	FRANKYWEBER	—	38ms	6/21/2023,
select * from IDENTIFIER('SNOWFLAKE_SA...	01ad1da7-0604-991d-0000-006a97879031	Success	FRANKYWEBER	COMPUTE_WH	2.6s	6/21/2023,
SELECT system\$GET_NPS_FEEDBACK_TIMESTAM...	01ad1da6-0604-99b7-0000-006a9787d0...	Success	FRANKYWEBER	—	88ms	6/21/2023,



Most Expensive Nodes (2 of 4)

TableScan [3]	97.9%
Aggregate [2]	2.0%


Profile Overview (Finished)

Total Execution Time (6m 47s) 100.0%

Processing	3.3%
Local Disk I/O	0.2%
Remote Disk I/O	98.4%
Initialization	0.0%

Statistics

Scan progress	100.00%
Bytes scanned	52.31GB
Percentage scanned from cache	0.00%
Partitions scanned	49448
Partitions total	49448



Most Expensive Nodes (2 of 4)

TableScan [3]	97.9%
Aggregate [2]	2.0%

Profile Overview

Processing	1.3%
Local Disk I/O	0.2%
Remote Disk I/O	98.5%

Statistics

Scan progress	100.00%
Bytes scanned	52.31GB
Percentage scanned from cache	0.00%
Partitions scanned	49448
Partitions total	49448

Attributes

Full table name
SNOWFLAKE_SAMPLE_DATA.TPCDS_SF100TCL.CATALOG_RETURNS

Columns (2 / 28)
CR_WAREHOUSE_SK
CR_RETURN_AMOUNT

ACCOUNTADMIN • COMPUTE_WH Share

SNOWFLAKE_SAMPLE_DATA.TPCDS_SF100TCL Settings Latest Version

```

1 use warehouse COMPUTE_WH;
2 use database SNOWFLAKE_SAMPLE_DATA;
3 use schema TPCDS_SF100TCL;
4
5 select CR_WAREHOUSE_SK, sum(CR_RETURN_AMOUNT) from SNOWFLAKE_SAMPLE_DATA.TPCDS_SF100TCL.CATALOG_RETURNS group by CR_WAREHOUSE_SK order by CR_WAREHOUSE_SK;

```

Results Chart

CR_WAREHOUSE_SK	SUM(CR_RETURN_AMOUNT)
4	602,003,005,972.16
5	607,224,886,572.71
6	604,478,284,031.29
7	606,894,996,855.25
8	604,473,895,523.76
9	606,428,809,214.98
10	609036574976.50
11	604,680,316,575.51
12	605,353,807,888.84
13	606543319132.50
14	608,037,849,790.78
15	603,323,202,225.75
16	606,263,899,314.57
17	604,324,565,321.44
18	608,591,623,895.51
19	603,794,964,549.42
20	603,402,622,654.83
21	607,204,961,960.21
22	605,730,918,885.88
23	604,358,330,559.93
24	605,057,438,546.86
25	606,009,232,730.74
26	605,771,705,184.21
27	606,442,728,249.04
28	604465280941.80
29	606,351,540,235.48
30	605,546,574,400.77
31	187311321850.50

Query Details

Query duration: 6m 48s

Rows: 31

Query ID: 01ad1dac-0604-99e6-0-

CR_WAREHOUSE_SK

SUM(CR_RETURN_AMOUNT)

FWF Franky Weber F...
ACCOUNTADMIN

Worksheets

Dashboards

Apps

Data

Marketplace

Activity

Query History

Copy History

Task History

Admin

Help & Support



7 days left in trial

Upgrade

UC863889

Query - 01ad1db8-0604-99b7-0000-006a9787d049

FRANKYWEBER

Query Details Query Profile

Details

Status	Duration	Driver Status
Success	132ms	Supported
Start Time	Query ID	Client Driver
6/21/2023, 10:12 AM	01ad1db8-0604-99b7-0000-006a9787d049	Go 1.1.5
End Time	Query Tag	Session ID
6/21/2023, 10:12 AM	—	30002953607847966
Warehouse Size	—	

SQL Text

```
select CR_WAREHOUSE_SK, sum(CR_RETURN_AMOUNT) from SNOWFLAKE_SAMPLE_DATA.TPCDS_SF100TCL.CATALOG_RETURNS group by CR_WAREHOUSE_SK order by CR_WAREHOUSE_SK;
```

Results 31 of 31 Rows • Cluster —

Export Results

	CR_WAREHOUSE_SK	SUM(CR_RETURN_AMOUNT)
1	1	606,257,321,579.51
2	2	606,112,043,310.83
3	3	604,387,937,037.27
4	4	602,003,005,972.16
5	5	607,224,886,572.71
6	6	604,478,284,031.29
7	7	606,894,996,855.25
8	8	604,473,895,523.76
9	9	606,428,809,214.98

QUERY RESULT REUSE [0] 100%


01ad1db8-0604-99b7-0000-006a9787...

Profile Overview (Finished)

Total Execution Time (1ms) 100.0%

Other 100.0%

Differences and Similarities



Oracle ADW vs Snowflake: Differences and Similarities

Oracle	Snowflake
User Defined Partitions / Automatic Partitioning	Micro Partitions
Hybrid Columnar Compression (6 - 15x)	Pure Columnar Compression (3 - 32x)
Advanced Row Compression (2 - 4x)	-
Native Row Format	Unistore for Hybrid Tables
Index / Auto Index	Search Optimization Service*
Zone Maps / Storage Index	Zone Maps
Data Clustering	Data Clustering (Auto Reclustering)

Oracle ADW vs Snowflake: Differences and Similarities

Oracle	Snowflake
Flashback Query / Flashback Database / Flashback Data Archive (Total Recall)	Time Travel
Point-In-Time-Recovery	Fail-Safe (recovery performed at best effort basis by support)
Serveless (shared), Dedicated, Cloud@Customer	AWS, Azure, GCP
Multi-Region and Hybrid	Multi-Region and Multi-Cloud
APEX	-
MPP with Parallel Degree	MPP with Warehouse Size
Multi-User Concurrency in any DBNode	Multi-User Concurrency with Additional Clusters

Oracle ADW vs Snowflake: Differences and Similarities

Oracle	Snowflake
Exadata Storage	Cloud-based Storage
Persistent Memory / Flash Cache / Buffer Cache	Local Cache (ephemeral NVMe) / No Buffer Pool
Smart Scan / Offloading	Partition Pruning at Micro Partition Level
RDMA/iDB Protocol	HTTPS
Limited Exadata Storage (13PB+ raw / 3.5PB dbsize)	Infinite Storage Scaling
Data Sharing Protocol	Proprietary Data Sharing
Many built-in tools available	Very easy to use

Oracle ADW vs Snowflake: Regions

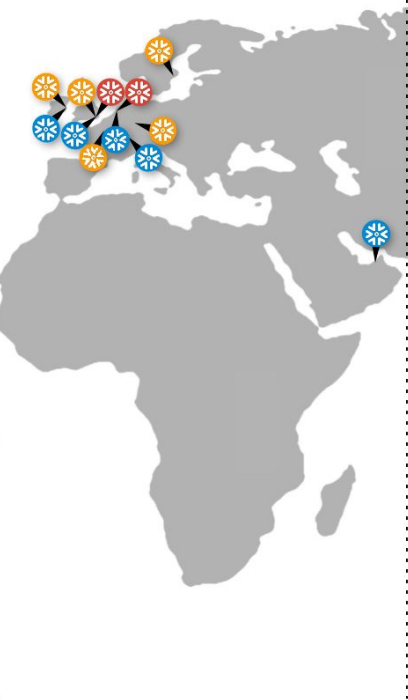


Oracle ADW vs Snowflake: Regions

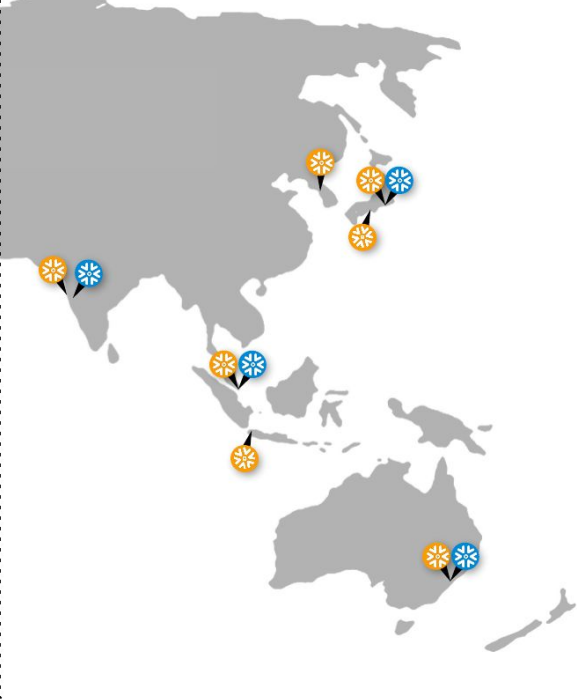
North & South America



Europe & Middle East


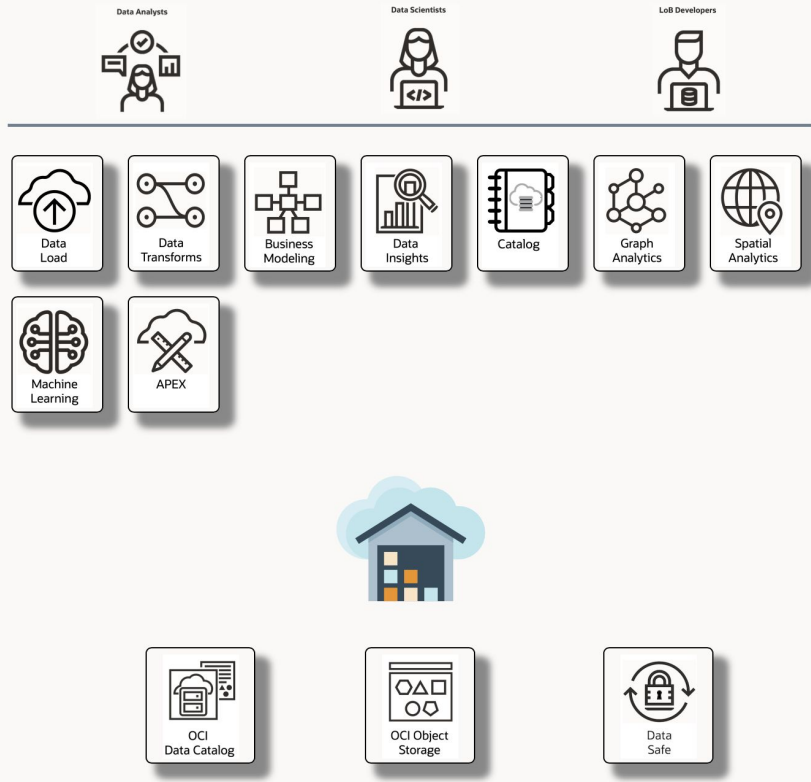


Asia Pacific



Oracle ADW vs Snowflake: Ecosystem

Empowering Business Innovators With Self-Serve Cloud Data Ecosystem



Oracle ADW vs Snowflake: Price / Money / Cash

Oracle ADW: Oracle offers a pay-as-you-go pricing model for its Autonomous Data Warehouse, with costs based on the number of OCPU or ECPU hours and storage capacity consumed. This allows businesses to scale their costs according to their needs, with the added benefit of Oracle's performance optimization features that help minimize resource consumption. Oracle early in 2023 reduced its ADW storage price drastically by matching it to Object Storage price. Another price advantage of ADW is the BYOL (Bring Your Own License) model which allows you to transfer your on-premises perpetual licenses to be used with ADW. A very simple pricing example using OCI Cost Estimator for shared deployment:

- <https://www.oracle.com/cloud/costestimator.html>
- License Included: 8 ECPU for 744 hours + 10TB storage/month = USD 1217.06/month

The screenshot shows the OCI Cost Estimator interface. The 'My Estimate' section shows an estimated monthly cost of \$1,217.06. Under 'Autonomous Database', the configuration is set to 'Database - Autonomous Database' with an estimated monthly cost of \$1,217.06. The 'License Type' is set to 'License Included', 'Deployment Type' is 'Shared', 'Workload Type' is 'Analytics and Data Warehousing', and 'Metric' is 'ECPU'. The 'UTILIZATION' is 744 hrs/month. Below this, there are fields for 'Peak ECPUs' (8), '% of Time at Peak' (50), and 'Storage [TB]' (10).

- BYOL: 8 ECPU for 744 hours + 10TB storage/month = USD 481.98/month

The screenshot shows the OCI Cost Estimator interface with the 'License Type' set to 'BYOL'. The 'My Estimate' section shows an estimated monthly cost of \$481.98. Under 'Autonomous Database', the configuration is the same as the previous screenshot, but the 'Estimated Monthly Cost' is now \$481.98.

Snowflake Data Cloud: Snowflake uses a consumption-based pricing model, where customers pay for the storage and compute resources they consume. This flexible pricing structure allows organizations to allocate resources based on their requirements, providing cost control and scalability. With Snowflake there are shared infrastructure options of choosing between Standard, Enterprise and Business Critical; the latest being the one most compatible with Oracle's ADW offering. Snowflake's pricing varies based on the deployment option, cloud provider and region. Considering AWS US-East (Ohio) a simple pricing example would be if the Virtual Warehouse is 100% utilized 744 hours/month:

- <https://www.snowflake.com/pricing/>
- Business Critical: 4 USD/credit with Large size Warehouse (8 credits/hour) + 10TB storage/month = $(4*8*744)+(40*10) = \text{USD } 24208/\text{month}$

The screenshot shows the Snowflake pricing page for Amazon Web Services (AWS) in the US East (Ohio) region. It displays four service tiers: Standard (\$2.00 per credit), Enterprise (\$3.00 per credit), Business Critical (\$4.00 per credit), and Virtual Private Snowflake (VPS) (Contact Us). Each tier includes a list of features and benefits. Below the tiers, there are options for 'ON-DEMAND STORAGE' (\$40 per TB per month) and 'CAPACITY STORAGE' (\$23 per TB per month).

<https://loredata.app/2023/05/17/oracle-adw-vs-snowflake-brief-overview/>

Oracle ADW vs Snowflake: Main Competition



BigQuery



amazon
REDSHIFT



Azure Synapse
Analytics



databricks

Thank you

Franky Weber Faust
faust@pythian.com