

**performing
databases**

Simple AI Use Cases for Oracle Database Professionals

Martin Klier 
Performing Databases GmbH
Mitterteich / Germany

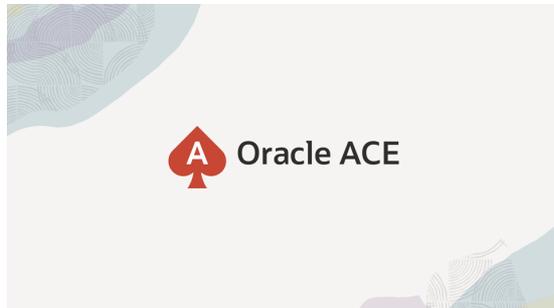
Speaker

- Martin Klier
- Solution Architect and Database Expert
- My focus:
 - Performance + Tuning
 - Highly available systems
 - Cluster and Replication
- Linux since 1997
- Oracle Database since 2003



SYMPOSIUM ^{L2}
Proud Member of symposium42

... it's all about Community!



SYMPOSIUM ^{L2}



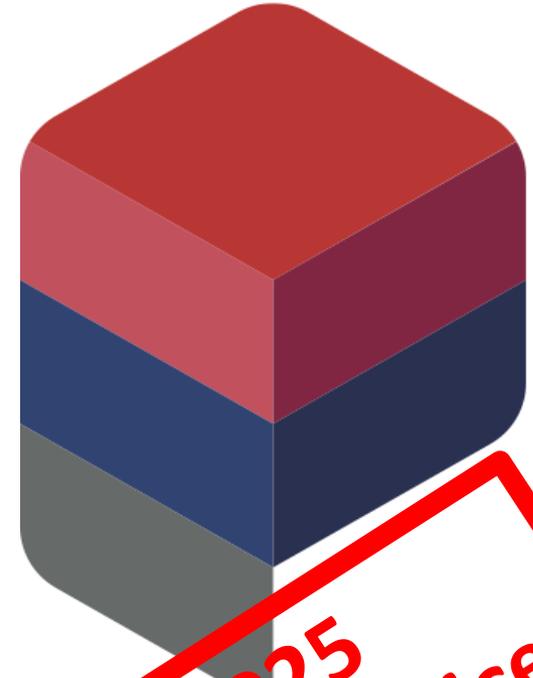
ora2know

The German Oracle Database Community

<https://www.ora2know.de>

Performing Databases

- Three Experts for Database technology
 - Concepts and Project Competence
 - Architecture- and System planning
 - Licensing
 - Implementation and Troubleshooting
- Get in touch
 - Performing Databases GmbH
Wiesauer Strasse 27
95666 Mitterteich // Germany
 - <http://www.performing-databases.com>
 - Social: @PerformingDB



Mainstream



Birthday Cakes for Performing Databases

The cake is ok, but the text ... my gosh.



Regenerate the picture without any written text



I see ... :/

“ANNIVERSASRY”
or
“ANNVERSARY”
???



“Broun yois Dataebrasse Services Exeperitence”

Researching Oracle History

Which company, purchased by Oracle, initially **invented the Database Listener** for the TNS protocol?



The company initially responsible for inventing the Database Listener for the TNS (Transparent Network Substrate) protocol, which Oracle later acquired, is **Digital Equipment Corporation (DEC)**. DEC was a pioneering company in the computer industry, well-known for its contributions to network and database technologies among many other innovations. Oracle acquired the rights to **RDB**, a relational database management system from DEC, which likely included technology related to networking and database communication protocols like TNS. Oracle's acquisition of various technologies from DEC helped enhance its own database solutions, including the implementation and further development of database listeners for efficient database management and network communications.

(this is where Mr. DATA GUARD, Larry Carpenter, came from)

Write Abstracts for OCW

What will I do and for what?

Suggest a title yourself

Describe content in own words & in detail

Tell the Model what you expect

I will submit a paper for Oracle Cloud World 2024, that is meant to show ways to use Artificial Intelligence for Oracle Database Administrators.

The title could be "Simple AI Use Cases for Oracle Database Professionals"

I will discuss

- using ChatGPT and other LLMs to reduce the time needed for reading documentations, especially for filtering out unwanted/unnecessary/unsuitable technical options
- using LLMs to auto-generate SQL Statements on a known schema
- using Oracle's built-in AI features to predict the behaviour and load of a static DB setup in future workloads

Write an abstract of about 500 characters, that is understandable for a mostly American audience



Title: Simple AI Use Cases for Oracle Database Professionals

Abstract: This paper explores practical AI applications to enhance the productivity of Oracle Database Administrators. By leveraging Large Language Models (LLMs) like ChatGPT, administrators can significantly cut down the time spent sifting through extensive documentation, pinpointing only the most relevant technical information. Furthermore, we demonstrate how LLMs can be utilized to auto-generate SQL statements for established schemas, streamlining database operations. Additionally, we delve into Oracle's native AI capabilities to forecast future workload behaviors in static database setups, offering a predictive edge in database management.

Quite good – you are here, at least. 😊

Read the Docs for me!

Can I have different native network encryption sqlnet.ora settings for different pluggable databases in the same CDB?

<...>

- In the `sqlnet.ora` file, specify network encryption settings for each service.

plaintext

Code kopieren

```
SQLNET.ENCRYPTION_TYPES_SERVICES = (  
    "PDB1_NET_SERVICE" = (AES256),  
    "PDB2_NET_SERVICE" = (AES128)  
)
```

SQLNET.ENCRYPTION_TYPES_SERVICES

```
SQLNET.ENCRYPTION_SERVER = REQUIRED  
SQLNET.CRYPTO_CHECKSUM_SERVER = REQUIRED
```

Looks very promising – did not know that!

I don't believe you ...

Don't trust no one ;)

What is the source for the SQLNET.ENCRYPTION_TYPES_SERVICES item?



I apologize for the confusion earlier. The `SQLNET.ENCRYPTION_TYPES_SERVICES` parameter **does not exist in Oracle's SQL*Net configuration**. Instead, Oracle does not natively support different `sqlnet.ora` settings per pluggable database (PDB) directly within the same container database (CDB) through the `sqlnet.ora` file. However, there are ways to achieve different encryption settings thro



It was just Hallucination!

LLM invents facts that
SOUND LOGICAL
but are NOT REAL.

Best feature-request tool in the world!

Generate Test Data

Generate Test Data for me.

My Oracle Database Table is CUSTOMER

Columns: ID (primary key), CU_NO, CU_NAME, CU_FIRSTNAME, CU_STREET, CU_CITY, CU_POSTCODE, CU_COUNTRY

I need 100 rows from Germany, UK and USA, no duplicates in CU_NAME, use German words for Germany, English and Welsh words for UK, and American words for USA. City and street names in Germany are to be German.

Build it as Oracle INSERT statements, output all lines, don't be lazy.

< 5/5 >

```
-- Germany
INSERT INTO CUSTOMER (ID, CU_NO, CU_NAME, CU_FIRSTNAME, CU_STREET, CU_CITY, CU_POSTCODE, CU_COUNTRY) VALUES
(1, 'CU001', 'Schneider', 'Hans', 'Hauptstraße 1', 'Berlin', '10115', 'Germany');
```

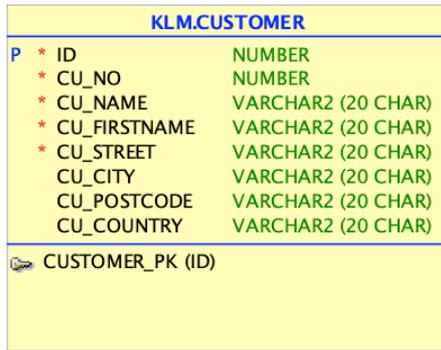
Professional

“Thoughts to SQL”

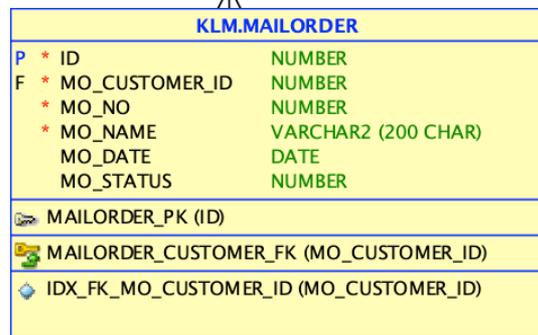
ORACLE



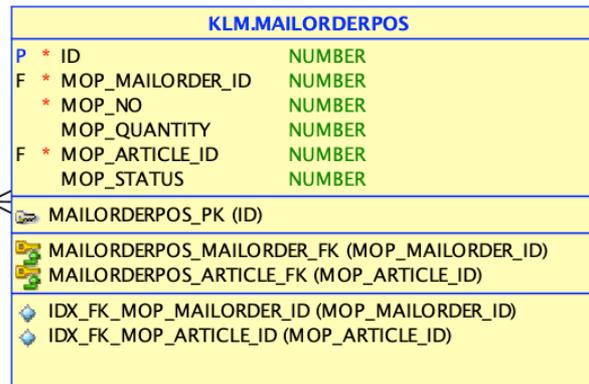
Lab Data Model



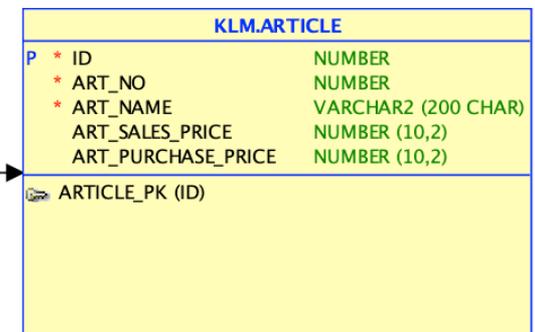
54 customers



1.000 mailorders



1.100.000 mailorderpos's



101 articles

The Formula to AI Generated SQL

- + Autonomous DB 19c or 23ai 
- + OCI Network ACL's 
- + User Grants 
- + LLM Provider Account (pay) 
- + Credential 
- + AI Profile 
- + Your “Select AI” Statement
- = **Result** (sort of)

Prepare Access Control + Grants

```
BEGIN
  DBMS_NETWORK_ACL_ADMIN.APPEND_HOST_ACE(
    host => 'api.openai.com',
    ace => xs$ace_type(privilege_list => xs$name_list('http'),
                      principal_name => 'KLM',
                      principal_type => xs_acl.ptype_db)
  );
END;
```

```
grant execute on DBMS_CLOUD to KLM;
grant execute on DBMS_CLOUD_AI to KLM;
```

OpenAI Credentials

```
BEGIN
  DBMS_CLOUD.DROP_CREDENTIAL (
    credential_name => 'OPENAI_CRED'
  );

  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'OPENAI_CRED',
    username => 'openaitest@performing-db.com',
    password => 'youdontwanttoknow'
  );
END;
/
```

Create AI Profile

```
BEGIN
  DBMS_CLOUD_AI.drop_profile(
    profile_name => 'OPENAI_TEST2'
  );
  DBMS_CLOUD_AI.create_profile(
    profile_name => 'OPENAI_TEST2',
    attributes => {"provider": "openai",
                  "model": "gpt-4o",
                  "credential_name": "OPENAI_CRED",
                  "object_list": [{"owner": "KLM", "name": "ARTICLE"},
                                {"owner": "KLM", "name": "CUSTOMER"},
                                {"owner": "KLM", "name": "MAILORDER"},
                                {"owner": "KLM", "name": "MAILORDERPOS"}
                  ]
  });
END;
/
```

Choose AI Profile

```
BEGIN
  DBMS_CLOUD_AI.SET_PROFILE(
    profile_name => 'OPENAI_TEST2'
  );
END;
/
```

```
BEGIN
  DBMS_CLOUD_AI.SET_PROFILE(
    profile_name => 'SOME_TEST1'
  );
END;
/
```

SELECT AI

select AI **/runsql** / what is the customer who ordered the most articles

	Customer Name	Total Articles Ordered
1	Smith	49284

select AI **narrate** what is the customer who ordered the most articles

RESPONSE
1 The customer named Smith has ordered the most articles, with a total of 49,284 articles.

select AI **showsql** what is the customer who ordered the most articles

RESPONSE
1 SELECT C.CU_NAME AS CUSTOMER_NAME, COUNT(MOP.ID) AS TOTAL_ARTICLES_ORDERED FROM

select AI **chat** what is the customer who ordered the most articles

RESPONSE
1 To determine the customer who ordered the most articles, we would need access to

SELECT AI – Test Data Generation

DBMS_CLOUD_AI.GENERATE_SYNTHETIC_DATA()

```
create table synt_customer as select * from customer where 0=1;

begin
  dbms_cloud_ai.generate_synthetic_data(
    profile_name => 'OPENAI_TEST1',
    object_name => 'SYNT_CUSTOMER',
    owner_name => 'KLM',
    record_count => 5,
    user_prompt => 'use German names and addresses'
  );
end;
/
```

jeergebnis 1 x

Alle Zeilen abgerufen:10 in 0,036 Sekunden

ID	CU_NO	CU_NAME	CU_FIRSTNAME	CU_STREET	CU_CITY	CU_POSTCODE	CU_COUNTRY
1	1001	Smith's Electronics	John	123 Elm St	Springfield	12345	USA
2	1002	Green Grocers	Emily	456 Oak St	Shelbyville	54321	USA
3	1003	Blue Skies Travel	Michael	789 Pine St	Capital City	67890	USA
4	1004	Red Apple Bakery	Sarah	321 Maple St	Ogdenville	98765	USA
5	1005	Golden Gate Realty	David	654 Cedar St	North Haverbrook	11223	USA
1	1001	Müller	Hans	Hauptstraße 5	Berlin	10115	Deutschland
2	1002	Schneider	Anna	Bahnhofstraße 12	München	80331	Deutschland
3	1003	Fischer	Peter	Gartenweg 3	Hamburg	20095	Deutschland
4	1004	Weber	Klara	Ringstraße 8	Köln	50667	Deutschland
5	1005	Schulz	Lukas	Bergstraße 10	Frankfurt am Main	60311	Deutschland



SELECT AI - Tuning 😊

select AI chat 'Can you improve this query? select * from customer c, mailorder m;'

"Certainly!

Here is an improved version of the query:

```
SELECT c.*, m.*  
FROM customer c  
JOIN mailorder m  
ON c.customer_id = m.customer_id;
```

This query specifies which columns from the customer and mailorder tables should be selected and uses a JOIN clause to link the two tables based on the customer_id column."

=> Pure ChatGPT

Knows our schema and what we did before!

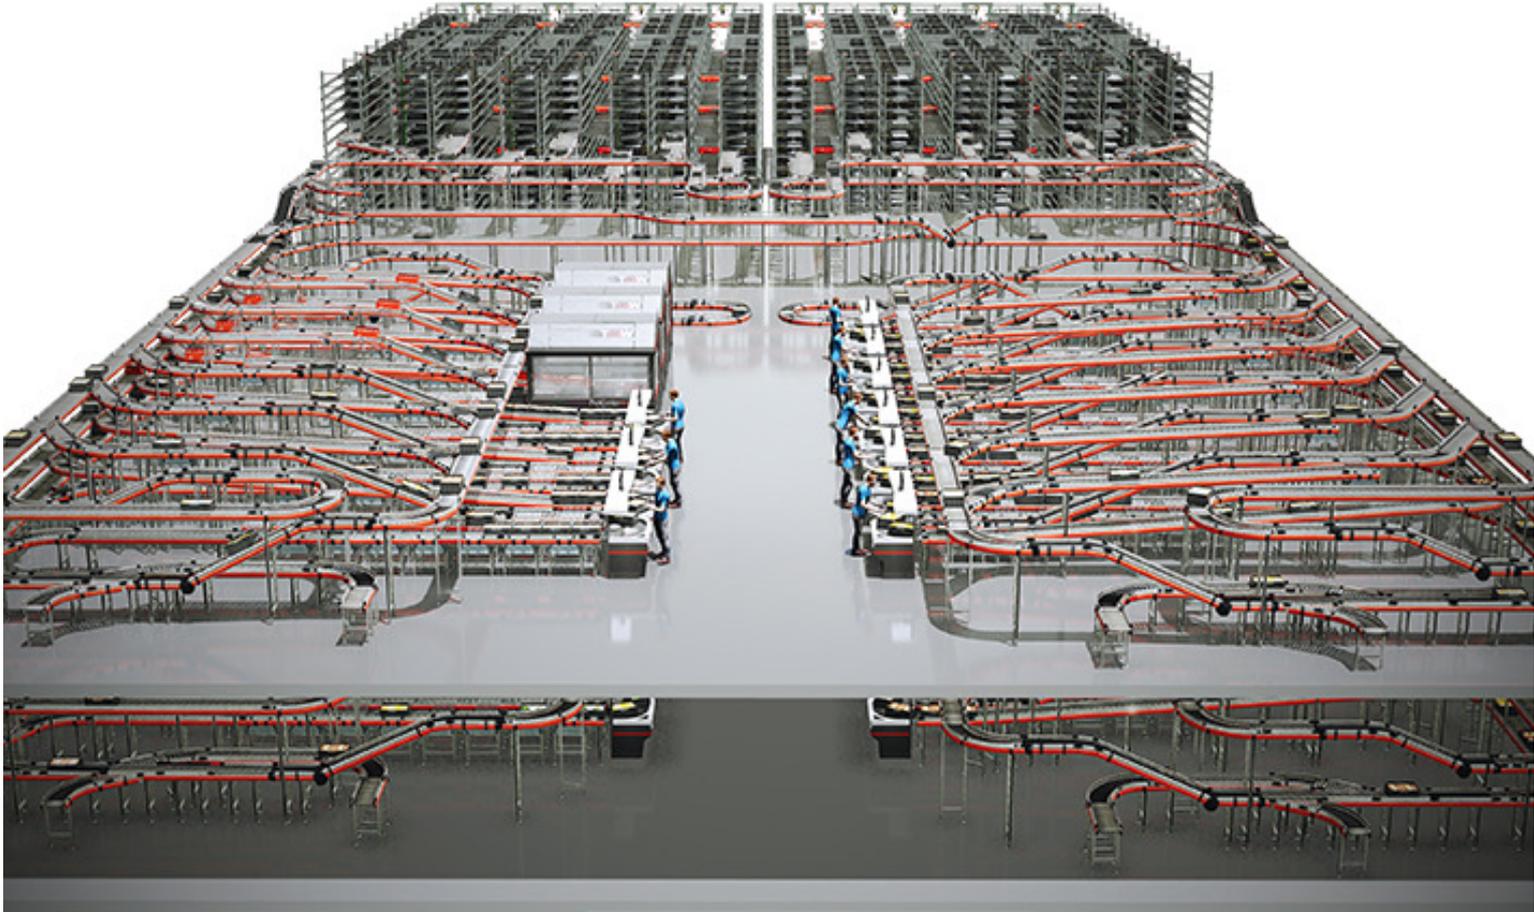
The Real Thing

ORACLE

Machine Sizing for a Logistics Warehouse

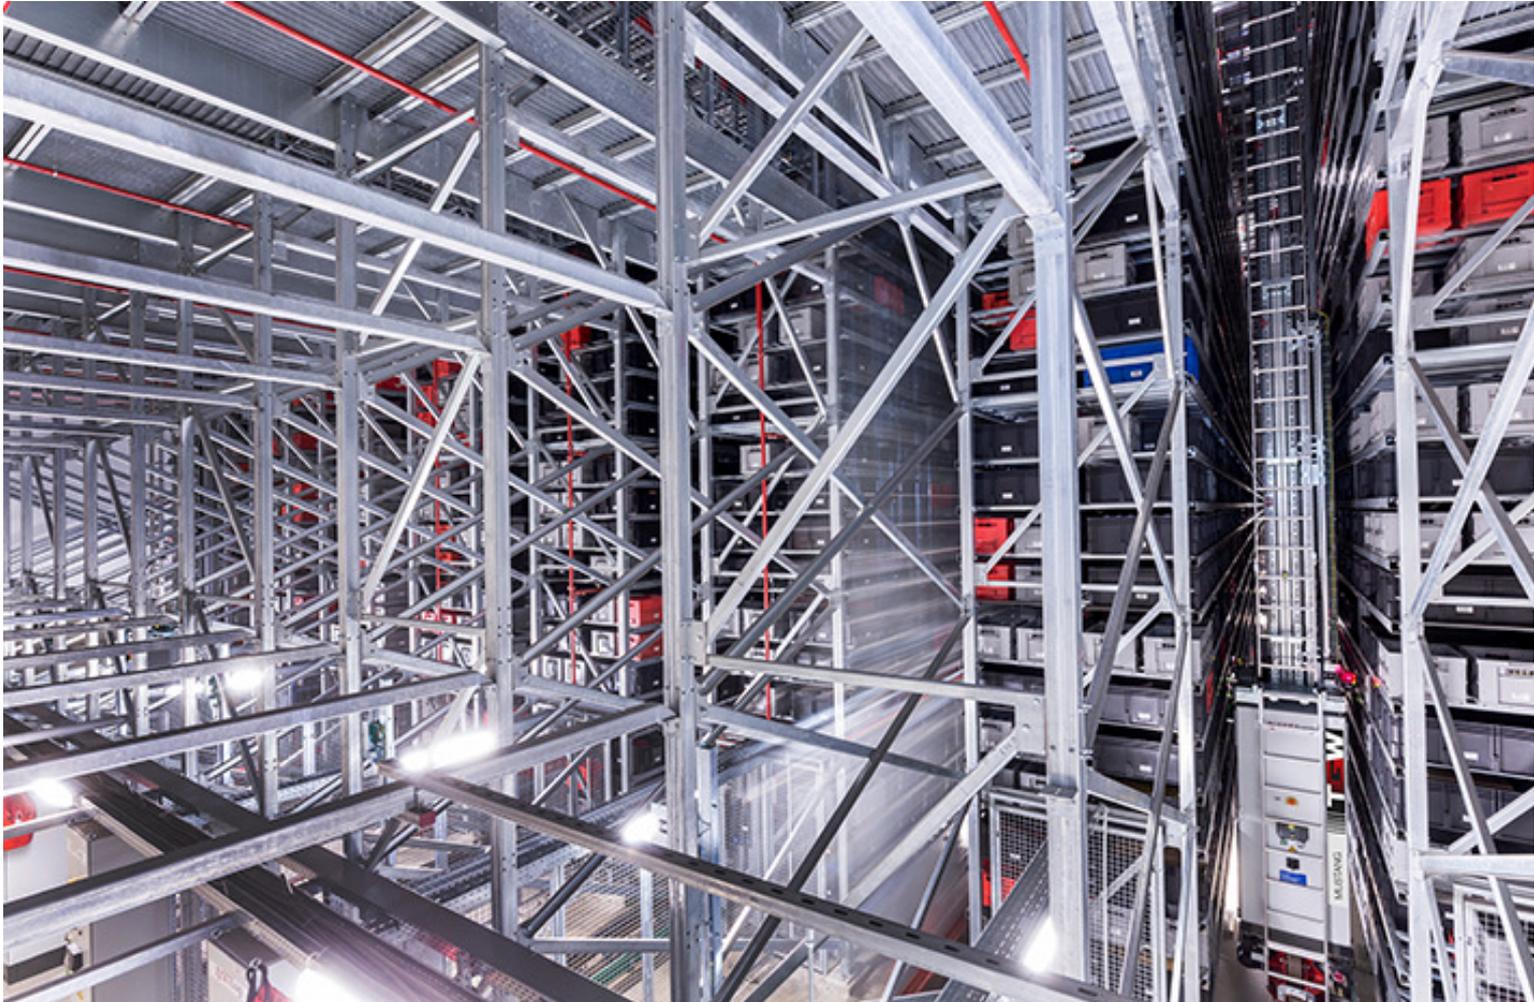
Mission: To Fulfill Orders

Complex Tasks done Fast



TGW Next Generation Fulfillment Center with FlashPick

Storage



TGW Miniload (Mustang)

Fast Storage



TGW Shuttle (Stingray)

Picking



TGW PickCenter One

How much CPU will we need?

(Hardware, Licensing...)

We need to collect data ...

Source 1: V\$SYSMETRIC

METRIC_ID

2057 as SYSMETRIC_HOST_CPU_UTIL_PCT_ABS

2075 as SYSMETRIC_CPU_CENTISEC_PER_SEC

2076 as SYSMETRIC_CPU_CENTISEC_PER_TX

2031 as SYSMETRIC_LOGICAL_READS_COUNT_PER_TX

2005 as SYSMETRIC_PHYSICAL_READS_COUNT_PER_TX

2007 as SYSMETRIC_PHYSICAL_WRITES_COUNT_PER_TX

2003 as SYSMETRIC_USER_TX_PER_SEC

2000 as SYSMETRIC_BUFFER_CACHE_HIT_RATIO_PCT_ABS

(Independent from # of CPUs in host)

(Assumption: All compared Host's CPUs are equally fast)

Source 2: Business

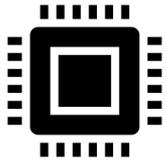
Let's count how many ...

- Orders
- Orders to pick
- Picks
- Mailings
- Goods-Ins
- Boxes moving
- ...

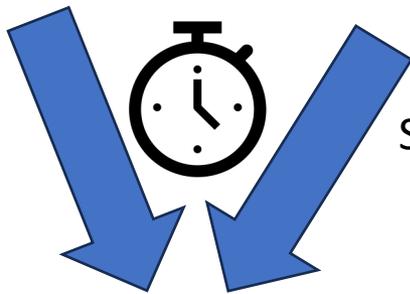
=> It's all in the database anyway

Put it together

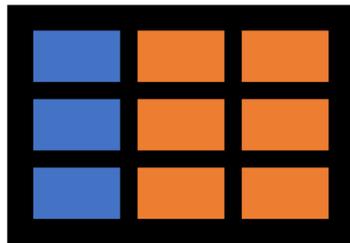
System Load



Business Data



Scheduler Job
5 Minutes
6 Months



LOADCOLLECTORDATA

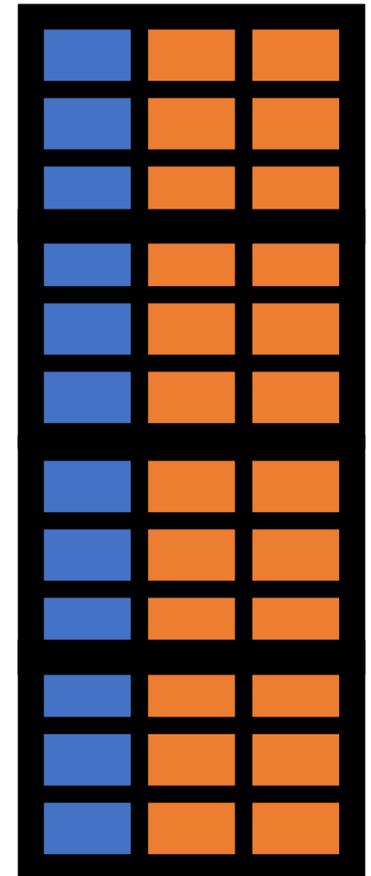
Warehouse 1

Warehouse 2

Warehouse 3

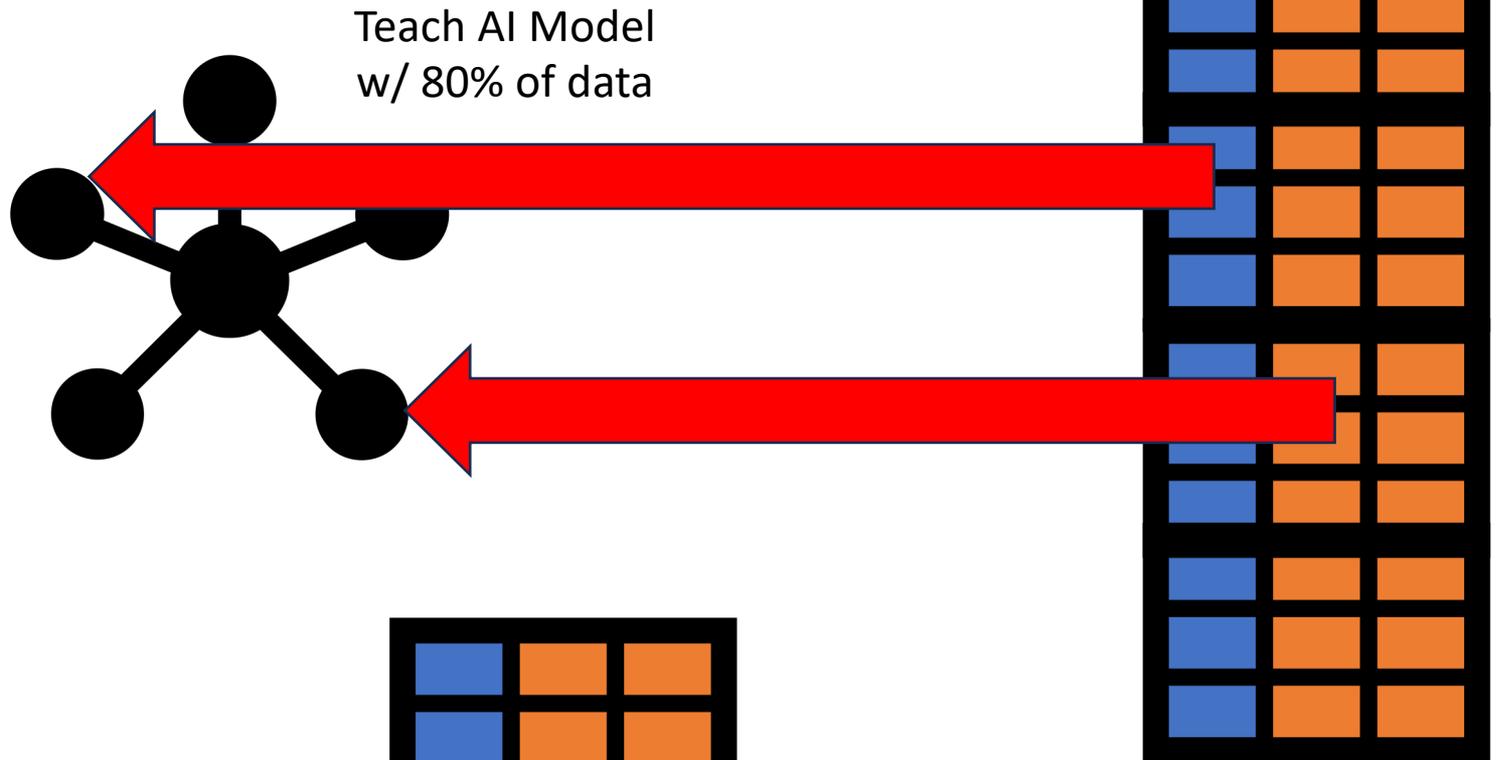
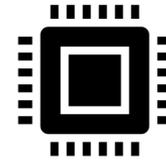
Warehouse 4

...



200k rows

Teach the Machine



Teach AI Model
w/ 80% of data

160k rows

“TRAINING_DATA”

Blue	Orange	Orange
Blue	Orange	Orange
Blue	Orange	Orange

Leave 40k rows
= 20% unknown to the model

Create the Model

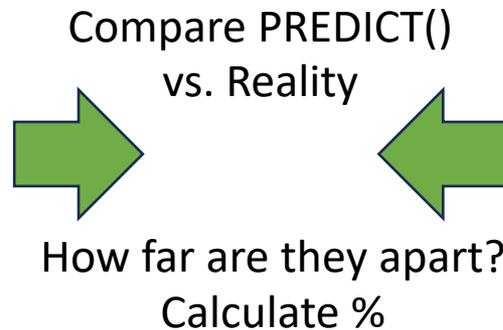
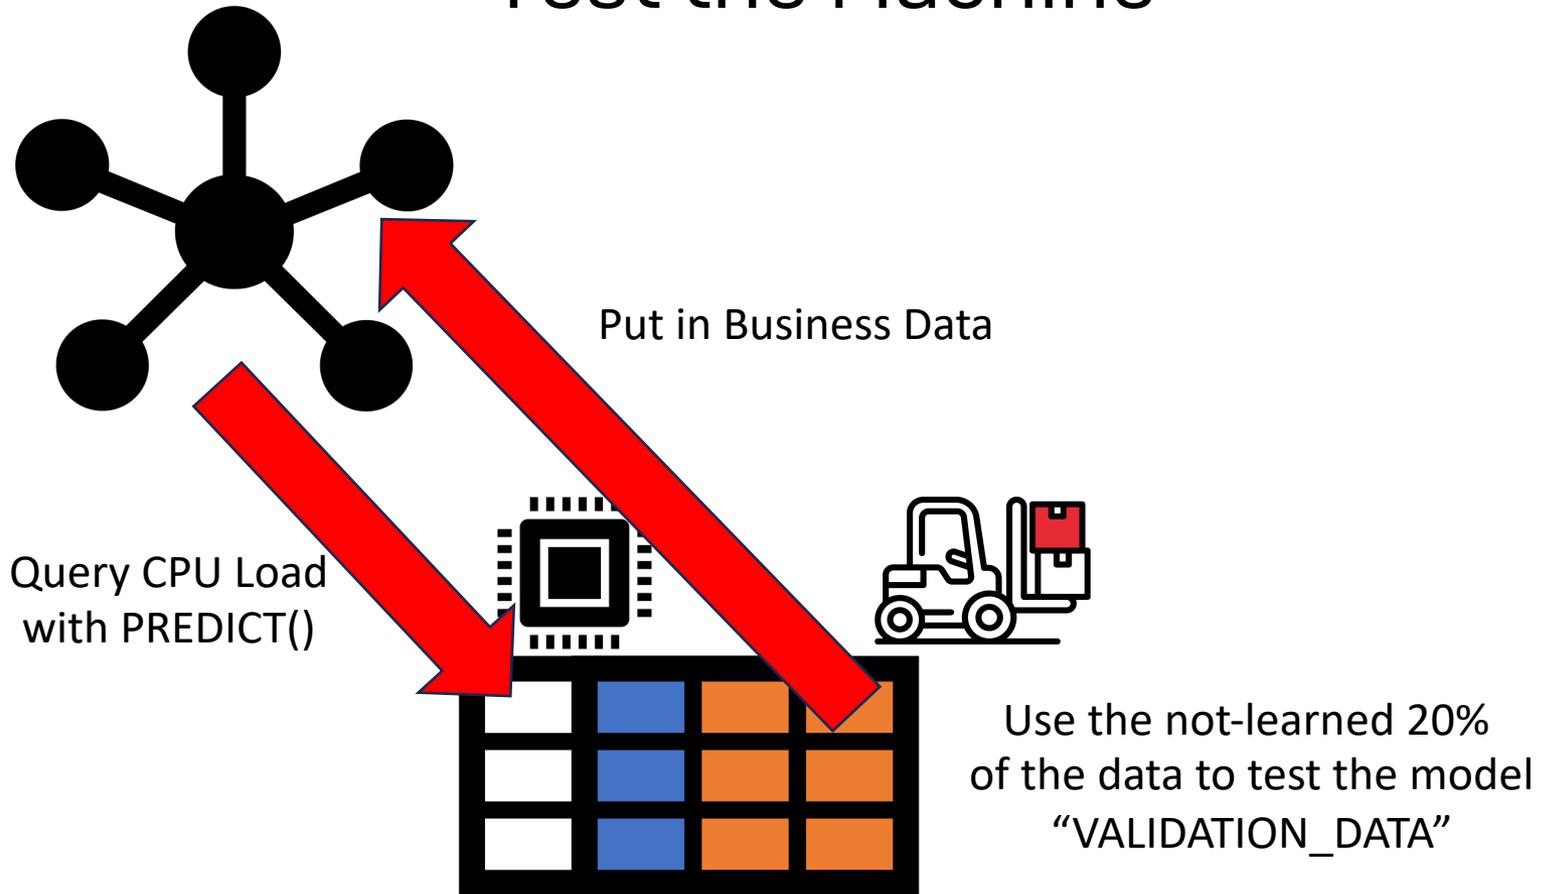
```
INSERT INTO MODEL_SETTINGS (setting_name, setting_value) VALUES
(dbms_data_mining.algo_name, dbms_data_mining.algo_support_vector_machines);
<...>
(dbms_data_mining.prep_auto, dbms_data_mining.prep_auto_on); -- Auto Preparation
<...>
(dbms_data_mining.svms_kernel_function, 'SVMS_GAUSSIAN'); -- Kernel type, default LINEAR
<...>
(dbms_data_mining.SVMS_EPSILON, '0.001'); -- Adjust epsilon in loss function, default 0.01
```

BEGIN

```
  DBMS_DATA_MINING.CREATE_MODEL(
    model_name           => 'klm_cpu_usage_model',
    mining_function      => DBMS_DATA_MINING.REGRESSION,
    data_table_name     => 'TRAINING_DATA', -- here is our input
    case_id_column_name => 'ROW_IDF', -- no PK available
    target_column_name  => 'SYSMETRIC_CPU_CENTISEC_PER_SEC',
    settings_table_name => 'model_settings');
```

END;

Test the Machine



Test the Model

SELECT

SYSMETRIC_CPU_CENTISEC_PER_SEC,

PREDICTION(klm_cpu_usage_model USING *) as PRED_SYSMETRIC_CPU_CENTISEC_PER_SEC

FROM v_validation_data vvd;

	SYSMETRIC_CPU_CENTISEC_PER_SEC	PRED_SYSMETRIC_CPU_CENTISEC_PER_SEC
41	105	116,54189646425778
42	35	71,7081008026
43	116	103,87113877716968
44	121	91,374779967629095
45	166	72,954513240133366
46	116	78,84580940523955
47	131	75,009826459854594
48	110	52,462489207008943
49	138	83,284919265141895
50	117	109,47182580082976
51	61	59,8170250358
52	55	59,804159170250358
53	56	59,804159170250358
54	87	65,0139601157
55	109	60,699176044346712
56	142	85,579637590962193
57	170	86,651310908673238
58	117	107,47143166218066
59	111	112,1195927101
60	130	73,602164329988325

Bad – 100% off

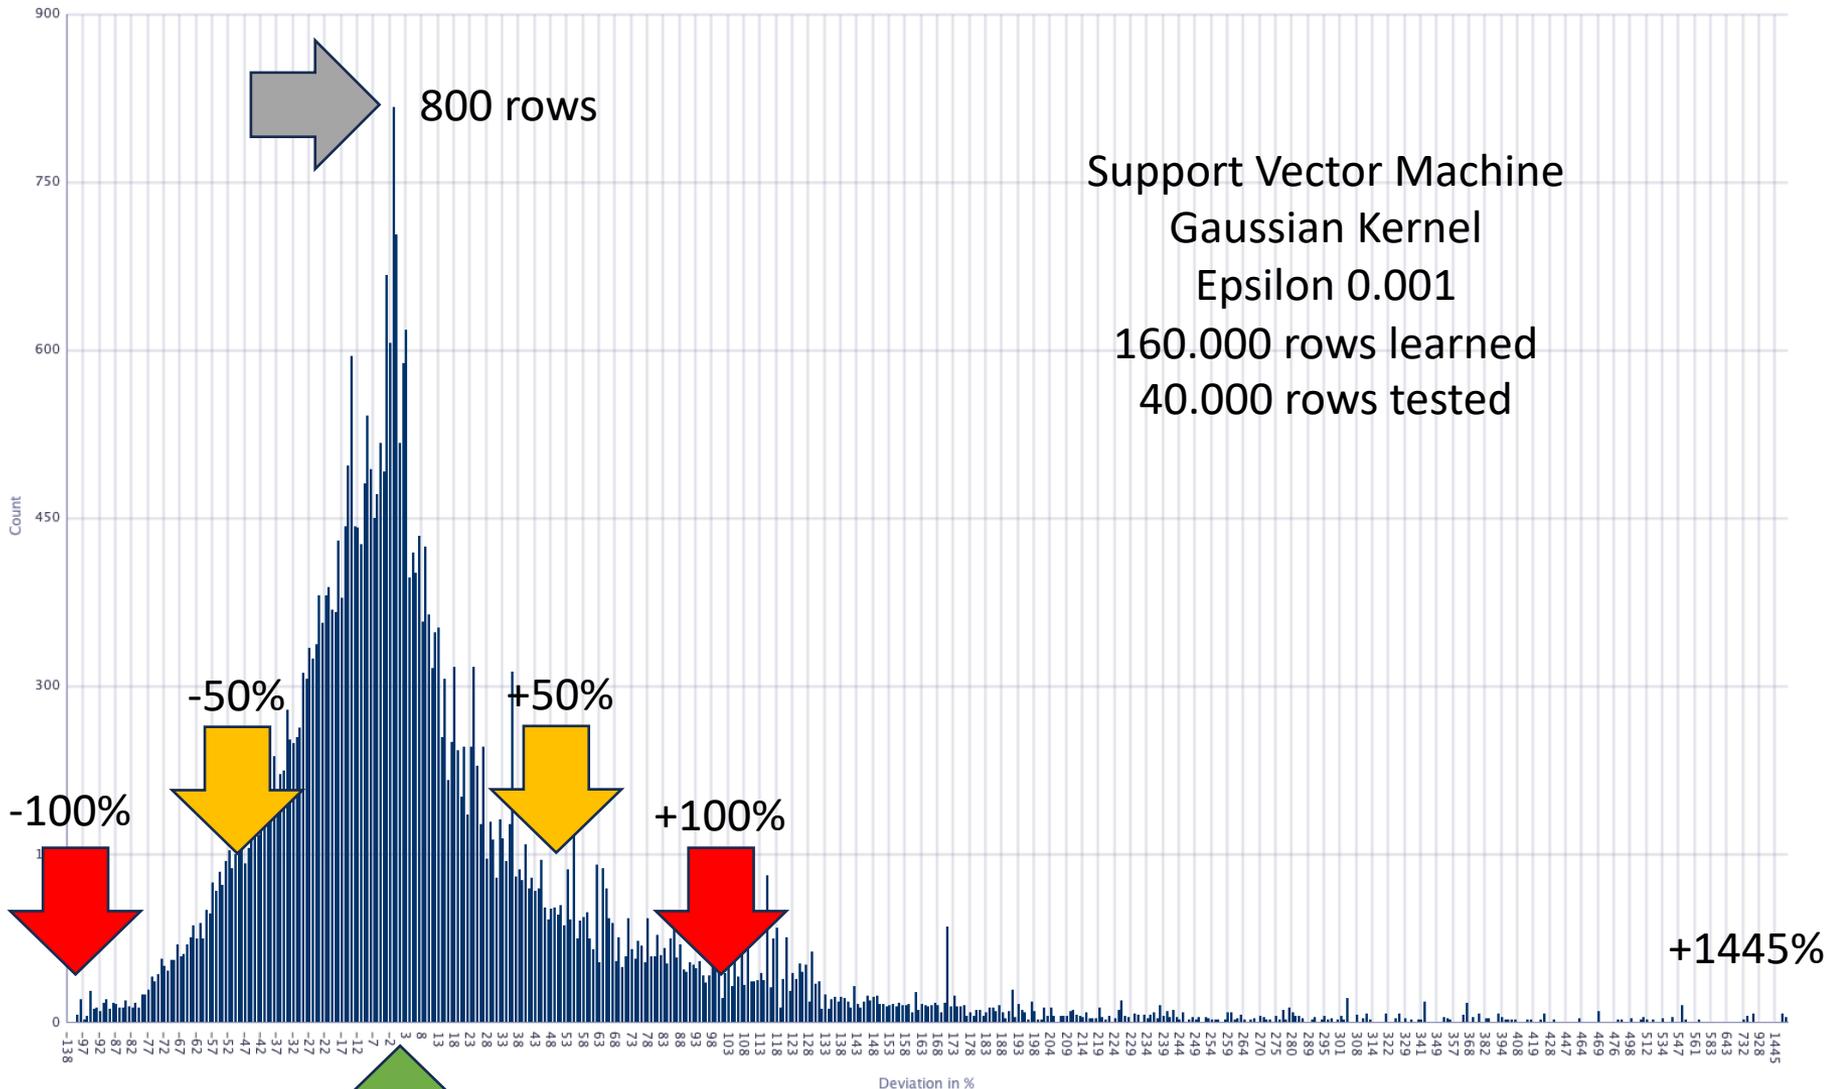
Quite good

Well, it depends 😊

Quite good @MartinKlierDBA
www.performing-databases.com

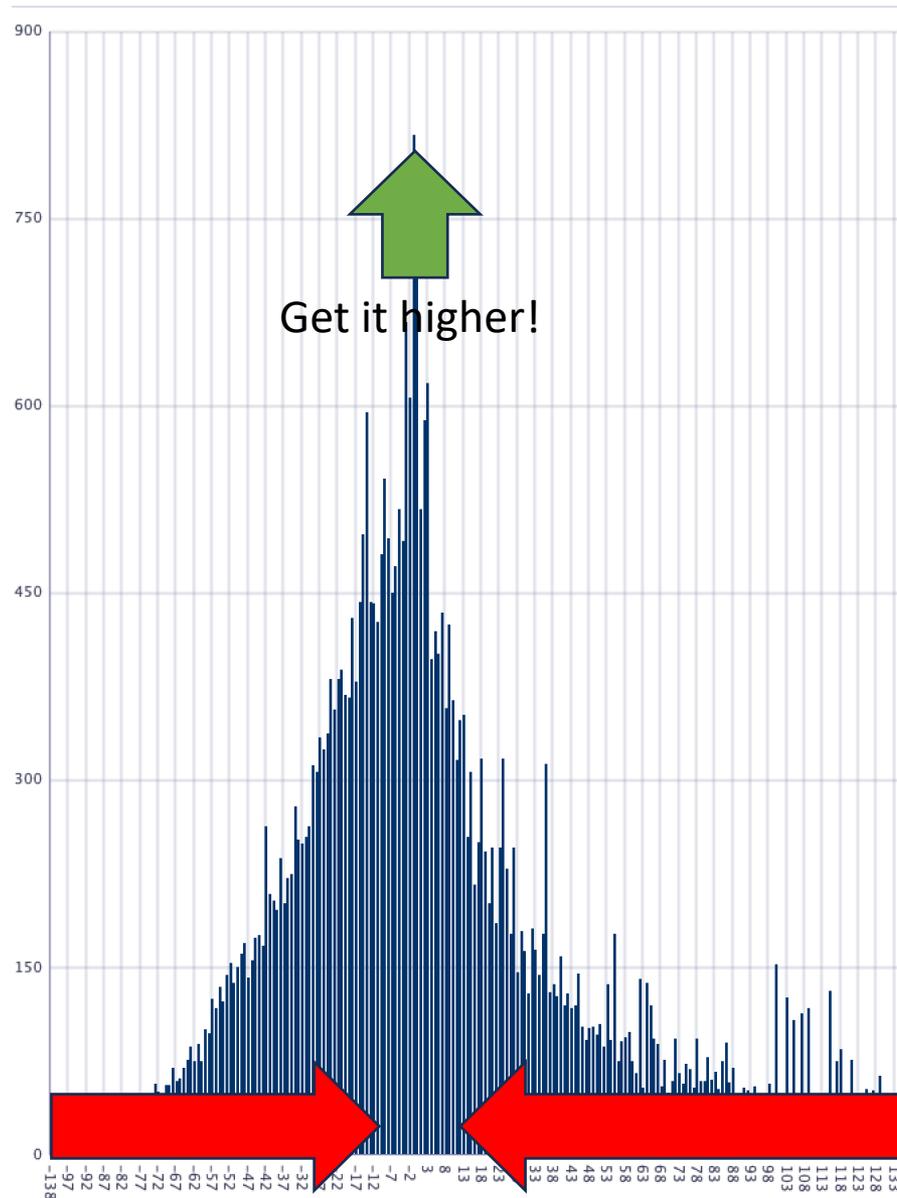
Success? Deviation in %

Prediction Deviation in % (SVM)



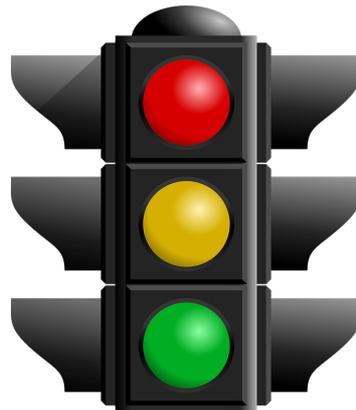
Support Vector Machine
Gaussian Kernel
Epsilon 0.001
160.000 rows learned
40.000 rows tested

Tune it!



What's next?

Traffic Light-Style Chaos Prevention?



Make your life easier with AI

by:

- Offloading Work you don't want do do anyway
- Using it as cheap labor
- Making it recognize patterns you won't see.

=> Just use it!

THANK YOU!

Meet & Greet

martin.klier@performing-db.com

www.performing-databases.com

Many national // international events

OUGN

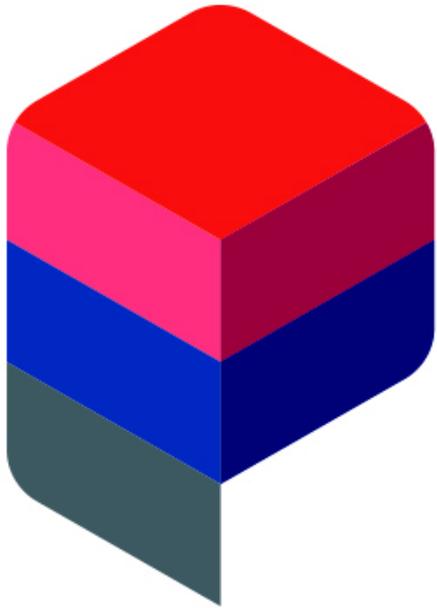
ora2know

DOAG

Meetups

UTOUG

...



**performing
databases**