



*Great Lakes Oracle Conference*

**MAY 15 & 16, 2019**

**CLEVELAND PUBLIC AUDITORIUM, CLEVELAND, OHIO**

**[WWW.NEOOUG.ORG/GLOC](http://WWW.NEOOUG.ORG/GLOC)**

There and back again.

or how to connect Oracle and Big Data.



Gleb Otochkin

2

Senior Cloud Architect

Started to work with data in 1992

Area of expertise:

- Data Integration
- Oracle RAC
- Oracle engineered systems
- Virtualization
- Performance tuning
- Big Data
- Cloud technologies.

[gleb.otochkin@eclipsys.ca](mailto:gleb.otochkin@eclipsys.ca)

@sky\_vst







# AGENDA



- Big Data?
- Business cases.
- Offload the data.
- Back again.
- Cloud.
- QA

# Big Data?



**Big Data?**  
**In Texas we call it just data.**

# Something about data.

Why do we talk about it and where it is coming from.

Stored as a ROW format

|       |       |       |       |
|-------|-------|-------|-------|
| Col A | Col B | Col C |       |
| Col A |       |       |       |
| Col A | Col B |       |       |
| Col A |       |       | Col D |
| Col A |       |       |       |
| Col A | Col B |       |       |

|       |       |       |  |
|-------|-------|-------|--|
| Col A |       |       |  |
| Col A | Col B |       |  |
| Col A |       |       |  |
| Col A |       |       |  |
| Col A | Col B |       |  |
| Col A |       | Col C |  |

# Something about data.

Why do we talk about it and where it is coming from.

## Columnar format

|       |       |       |  |
|-------|-------|-------|--|
| Col A | Col A | Col A |  |
| Col A | Col A |       |  |
| Col A | Col A |       |  |
| Col A | Col A |       |  |
| Col A | Col A |       |  |
| Col A | Col A |       |  |

|       |       |  |  |
|-------|-------|--|--|
| Col B | Col B |  |  |
| Col B | Col B |  |  |
| Col B |       |  |  |
| Col B |       |  |  |
| Col B |       |  |  |
| Col B |       |  |  |



# Something about data.

Why do we talk about it and where it is coming from.


## Unstructured

|       |       |       |       |
|-------|-------|-------|-------|
| Col A | Col B | Col C |       |
| Col A |       |       |       |
| Col A | Col B |       |       |
| Col A |       |       | Col D |
| Col A |       |       |       |
| Col A | Col B |       |       |

|       |        |        |        |       |
|-------|--------|--------|--------|-------|
| Col A |        |        |        |       |
| Col A | Col B  |        |        |       |
| Col A | Col B1 | Col B2 |        |       |
| Col A |        |        |        |       |
| Col A | Col B  |        |        |       |
| Col A |        |        | Col C1 | ColC2 |

# Something about data.

Why do we talk about it and where it is coming from.



Easy for GBs scale tables.

What about TB or PB ?

# Big Data ecosystem<sup>11</sup>

COMPLEX

VOLUME

STRUCTURE

ANALYSIS

PROCESSING

GROWS

DISCOVERY



# Some Big Data tools and terms.


Some terms and tools in the Big Data:

- HDFS - Hadoop Distributed File System.
- Apache Hadoop - framework for distributed storage and processing.
- HBase - non-relational, distributed database.
- Kafka - streaming processing data platform.
- Flume - streaming, aggregation data framework.
- Cassandra - open-source distributed NoSQL database.
- MongoDB - open-source cross-platform document-oriented database.
- Hive - query and analysis data framework on top of Hadoop.
- Avro - data format widely used in BD (JSON+binary)


# Business cases




Operation activity on DB with rest on a Data Lake.




Operation activity on DB with Data in Big Data and BI on a RDBMS.




Operation activity on DB with Data in Big Data and BI on a RDBMS.



Operation and BI activity on DB with main Data body in a Big Data platform.



Operation and BI activity on DB with main Data body in a Big Data platform.



- Reasons:
  - UNstructured data.
  - Growing Data Volume.
  - Data retention policy.
  - Storage cost.
  - New data formats.
  - New API and options.
- Tools and methods:
  - Data replication tools.
    - Oracle GoldenGate.
  - Batch load tools.
    - Sqoop.
  - Data Integration tools.
    - Oracle Data Integrator.
  - Streaming.
    - Kafka, Spark
  - Presenting Big Data to RDBMS:
    - Oracle Big Data SQL.
    - Oracle Big Data Connectors.

# Offload the data.

Real time replication OLTP data to a Big Data platform.




### Oracle Goldengate:



- Proc.
  - Real time streaming.
  - No or minimal impact to the source.
  - Supports most of the data types.
  - Database engine agnostic.
  - Different BD targets.
  - Enterprise support.
- Cons.
  - Licensing fee.
  - Closed source code.


# Replication to Big Data data.

Replication to HDFS by Oracle GoldenGate.




# Replication to Big Data - capture.


GoldenGate - Source side.



- Oracle 11.2.0.4 and up
- Archivelog mode.
- Supplemental logging:
  - Minimal on DB level.
  - On schema or table level.
- OGG user in database.




- OGG 12.2.0.1 and up.
- BD adapters with replicat.
- Supported BD targets:
  - HDFS
  - Kafka
  - Flume
  - HBase
  - MongoDB (from 12.3)
  - Cassandra (from 12.3)
- Different formats.



- OGG 12.2.0.1 and up.
- BD adapters with replicat.
- DML and DDL.
- Supported BD targets:
  - HDFS
  - Kafka
  - Flume
  - HBase
  - MongoDB (from 12.3)
  - Cassandra (from 12.3)

- With the latest OGG for BD:
  - JDBC , Elasticsearch, Kinesis, Kafka Connect

...



- File Writer Handler.
  - HDFS Event Handler.
  - ORC Event Handler.
  - Oracle Cloud Event Handler.
  - Parquet Event Handler.
  - S3 Event Handler.

- List of target handlers.
  - BigQuery Handler.
  - Cassandra Handler.
  - Elasticsearch Handler.
  - File Writer Handler.
    - HDFS Event Handler.
    - ORC Event Handler.
    - Oracle Cloud Event Handler.
    - Parquet Event Handler.
    - S3 Event Handler.
  - Flume Handler.
  - HBase Event Handler.
  - HDFS Handler.
  - JDBC Handler.
- List of target handlers (continued).
  - Kafka.
  - Kafka connect.
  - Kafka REST Proxy Handler.
  - Kinesis Streams Handler.
  - MongoDB Handler.
  - Oracle NoSQL Handler.
  - Microsoft Azure Data Lake.
  - File Writer Handler.
- Source handler.
  - Oracle GoldenGate Capture for Cassandra.

Data on the source and on the target.


```
orcl> select * from ggtest.test_tab_2;  
ID  RND_STR  USE_DATE  
1   BGBXRKJL 02/13/16 08:34:19  
2   FNMCEPWE 08/17/15 04:50:18
```

- New columns:
  - Operation type.
  - Table name
  - Local and UTC timestamp

```
hive> select * from BDTEST.TEST_TAB_2;  
OK  
I      BDTEST.TEST_TAB_2      2016-10-25 01:09:16.000168      2016-10-24T21:09:21.186000      00000000120000004759      1  
BGBXRKJL      2016-02-13:08:34:19      NULL  
I      BDTEST.TEST_TAB_2      2016-10-25 01:09:16.000168      2016-10-24T21:09:22.827000      00000000120000004921      2  
FNMCEPWE      2015-08-17:04:50:18      NULL
```



## Database



Insert row 1


Insert row 2

Insert row 3

Update row 3

Delete row 2

## HDFS



# Replication to Big Data - supplemental logging.

Primary key for a source table in HBase.

- Having supplemental logging for all columns:
  - Row id as concatenation of values for all columns.
- Adding a primary key and supplemental logging for keys:
  - Row id is primary key.

```
orcl> alter table ggtest.test_tab_2 add constraint pk_test_tab_2 primary key  
(pk_id);
```

```
Table altered.
```

```
orcl> insert into ggtest.test_tab_2 values(9,'PK_TEST',sysdate,null);
```



```
hbase(main):012:0> scan 'BDTEST:TEST_TAB_2'
ROW
7|IJWQRO7T|2013-07-07:08:13:52
7|IJWQRO7T|2013-07-07:08:13:52
7|IJWQRO7T|2013-07-07:08:13:52
8|TEST_INS1|2016-03-29:15:14:37|TEST_ALTER
8|TEST_INS1|2016-03-29:15:14:37|TEST_ALTER
8|TEST_INS1|2016-03-29:15:14:37|TEST_ALTER
8|TEST_INS1|2016-03-29:15:14:37|TEST_ALTER
9
9
9
9
COLUMN+CELL
column=cf:ACC_DATE,timestamp=1459275116849,value=2013-07-07:08:13:52
column=cf:PK_ID,timestamp=1459275116849,value=7
column=cf:RND_STR_1,timestamp=1459275116849,value=IJWQRO7T
column=cf:ACC_DATE,timestamp=1459278884047,value=2016-03-29:15:14:37
column=cf:PK_ID,timestamp=1459278884047,value=8
column=cf:RND_STR_1,timestamp=1459278884047,value=TEST_INS1
column=cf:TEST_COL,timestamp=1459278884047,value=TEST_ALTER
column=cf:ACC_DATE,timestamp=1462473865704,value=2016-05-05:14:44:19
column=cf:PK_ID,timestamp=1462473865704,value=9
column=cf:RND_STR_1,timestamp=1462473865704,value=PK_TEST
column=cf:TEST_COL,timestamp=1462473865704,value=NULL
```

## Some notes about source:

- We are replicating a log of changes.
- Only committed transactions.
- "Passive" commit.
- Different levels of supplemental logging may lead to different captured data (compressed DML).
- DDL only with following DML.

## and BD destination:

- JSON, Text, Avro (different types) and XML formats.
- Hive support through text format.
- Trimming out leading or trailing whitespaces.
- Truncates only as DML.

What can be used as an alternative.


- Kafka using JDBC to pull data.
- Custom CDC solution built using JDBC and Java applications.
- Custom solutions based on triggers and AQ.
- Quest Shareplex.
  - Replication to Kafka.
  - Supports most of data types.
  - <https://www.quest.com/landing/shareplex/>
- Attunity replicate.
  - <https://www.attunity.com/products/replicate/platform-support>

# And what about Sqoop?



Batch offloading.


# Batch offloading from Oracle to Big Data.



- Sqoop one of the oldest.
- Works in both ways.
- Gluent is more than a just batch tool.
- BD SQL Copy to Hadoop.
  - Using Oracle DataPump format




Generate MapReduce Jobs.  
Connects through JDBC.  
Good for batch processing.



# Back again

Using data from Big Data in Oracle (and Oracle data in Big Data)





- Datasource for Apache Hadoop.
  - Hive access to Oracle Database tables.
  - Predicate pushdown and partition pruning.
  - Directly convert SQL data.
  - Parallel access to data.
  - Kerberos, SSL, Oracle Wallet.
  - Write data back to Oracle.
- SQL Connector for HDFS
  - SQL access to data in Hadoop.
  - Partition-aware access of Hive.
  - Supports parallel query and load.
  - Kerberos security.
  - Can use Oracle Data Pump files in HDFS.

- Oracle Loader for Hadoop.
  - Partition-aware load.
  - Parallel load.
  - Supports Parquet, JSON, Text ...
  - Offload data type conversion to Hadoop.
  - Kerberos security.
- Oracle XQuery for Hadoop.
  - Integration with Oozie workflows, Cloudera search and XML for Hive.
  - XQueries execute where the data is located.
  - Works with HDFS, Hive, or Oracle NoSQL Database.
  - Parallel XML parsing.
  - Can use Oracle Loader for Hadoop to write back to Oracle.


# From Big Data to Oracle.


| Tool                          | Filtering | Data Conversion | Query | Offload query | Parallel |
|-------------------------------|-----------|-----------------|-------|---------------|----------|
| Sqoop                         | Oracle    | Oracle          |       |               | Yes      |
| ODBC gateway                  | Hadoop    | Oracle          | Yes   | Yes           |          |
| Oracle loader for HDFS        | Oracle    | Hadoop          |       |               | Yes      |
| Oracle SQL Connector for HDFS | Oracle    | Oracle          | Yes   |               | Yes      |
| BD SQL                        | Hadoop    | Hadoop          | Yes   | Yes           | Yes      |
| ODI                           | KM        | KM              | Yes   |               | Yes      |
| Gluent                        | Hadoop    | Hadoop          | Yes   | Yes           | Yes      |












- Logical mapping is separated from physical implementation.

Have a big monitor. It helps a lot!




- Different physical mapping for the same logical map
- Different knowledge modules.

The same advice - big monitor. It helps a lot!

- ELT:
  - Extract.
  - Load.
  - Transform.
- Logical and physical models separation.
- Knowledge modules:
  - RKM - reverse engineering.
  - LKM - load.
  - IKM - integration.
- Knowledge modules for most platforms.
- Uses filters, mappings, joins and constraints.
- Batch or event(stream) oriented integration.

Architecture.



**Exadata technology for Big Data**

- Works with:
  - Apache Hive.
  - HDFS
  - Oracle NoSQL Database.
  - Apache HBase
- Fully supports SQL syntax.
- Predicate push down.
- Smart scan:
- Storage indexes.
- Bloom filters.

```
orcl> select cluster_id,database_name, owner, table_name from all_hive_tables where
database_name='bdtest';
```

| CLUSTER_ID  | DATABASE_NAME | OWNER  | TABLE_NAME |
|-------------|---------------|--------|------------|
| bigdatalite | bdtest        | oracle | test_tab_1 |
| bigdatalite | bdtest        | oracle | test_tab_2 |

- Smart scan:
- Storage indexes.
- Bloom filters.

```
PROCEDURE CREATE_EXTDDL_FOR_HIVE
```

| Argument Name   | Type           | In/Out | Default? |
|-----------------|----------------|--------|----------|
| CLUSTER_ID      | VARCHAR2       | IN     |          |
| DB_NAME         | VARCHAR2       | IN     |          |
| HIVE_TABLE_NAME | VARCHAR2       | IN     |          |
| HIVE_PARTITION  | PL/SQL BOOLEAN | IN     |          |
| TABLE_NAME      | VARCHAR2       | IN     |          |
| PERFORM_DDL     | PL/SQL BOOLEAN | IN     | DEFAULT  |
| TEXT_OF_DDL     | CLOB           | OUT    |          |

```
CREATE TABLE test_tab_2 ( tran_flag VARCHAR2(4000), tab_name VARCHAR2(4000), .....)  
ORGANIZATION EXTERNAL (TYPE ORACLE_HIVE DEFAULT DIRECTORY DEFAULT_DIR ACCESS PARAMETERS  
( com.oracle.bigdata.cluster=bigdatalite  
com.oracle.bigdata.tablename=bdtest.test_tab_2) ) PARALLEL 2 REJECT LIMIT UNLIMITED
```

| TRAN_FLAG | ID | RND_STR  | USE_DATE            |
|-----------|----|----------|---------------------|
| I         | 1  | BGBXRKJL | 2016-02-13:08:34:19 |
| I         | 2  | FNMCEPWE | 2015-08-17:04:50:18 |
| I         | 1  | BGBXRKJL | 2016-02-13:08:34:19 |

- The hive table can be queried in Oracle now.



## PLAN\_TABLE\_OUTPUT

Plan hash value: 3629305071

| Id | Operation                          | Name       | Rows | Bytes | Cost (%CPU) | Time     | TQ    | IN-OUT | PQ Distrib |
|----|------------------------------------|------------|------|-------|-------------|----------|-------|--------|------------|
| 0  | SELECT STATEMENT                   |            | 1    | 8008  | 2 (0)       | 00:00:01 |       |        |            |
| 1  | PX COORDINATOR                     |            |      |       |             |          |       |        |            |
| 2  | PX SEND QC (RANDOM)                | :TQ10000   | 1    | 8008  | 2 (0)       | 00:00:01 | Q1,00 | P->S   | QC (RAND)  |
| 3  | PX BLOCK ITERATOR                  |            | 1    | 8008  | 2 (0)       | 00:00:01 | Q1,00 | PCWC   |            |
| 4  | EXTERNAL TABLE ACCESS STORAGE FULL | TEST_TAB_2 | 1    | 8008  | 2 (0)       | 00:00:01 | Q1,00 | PCWP   |            |

## Note

- Degree of Parallelism is 2 because of table property

- Parallel FULL table access to the external table.

- Access drivers :
  - ORACLE\_HDFS
  - ORACLE\_HIVE
- Smart Scan for HDFS:
  - Parallel processing of the Hadoop.
  - Reduces data movement.
  - Returns much smaller result sets.
- Storage Indexes:
  - Built automatically.
  - Up to 32 columns.
- Predicate Push Down.
  - Partitioned Hive tables are pruned.
  - Apache Parquet and ORC files.
  - CTAS to external table.
  - Oracle NoSQL Database.
  - Apache HBase.

# Some other tools.

What can be used as an alternative.

- Gluent Data Platform.
  - Can offload data to Hadoop.
  - Allows query data of any data source in Hadoop.
  - Offload engine .
  - Supports Cloudera, Hortonworks or any Hadoop with Impala or or Hive SQL engine installed.

<https://gluent.com/>

Cloud.


- Easy setup.
- Integrated with other Oracle Cloud services.
- Secure.
- Highly customizable and flexible.
- Supports multiple types of targets.

ORACLE<sup>®</sup> Cloud

Platform > Integration > GoldenGate




GoldenGate



Data transformation, integration, replication, analysis, and governance..

- Connect to data sources and prepare, transform, replicate, analyze, govern, and monitor data
- Analyze data streams.
- Set up policies based on metrics to receive notifications.
- Manage all your data sources from a single platform.
- On-premises to Cloud.
- Cloud to Cloud.
- Cloud to On-premises.
- On-premises to On-premises.

# Oracle Data Integration Platform Cloud







A conceptual image featuring a hand pointing towards a city skyline. The image is heavily layered with a blue and teal color palette, creating a sense of depth and focus. The text 'QA' is prominently displayed in the upper left quadrant.

QA

# THANK YOU

Email: [gleb.otochkin@eclipsys.ca](mailto:gleb.otochkin@eclipsys.ca)

<https://eclipsys.ca/blog/>

@sky\_vst