ORACLE Platinum Partner

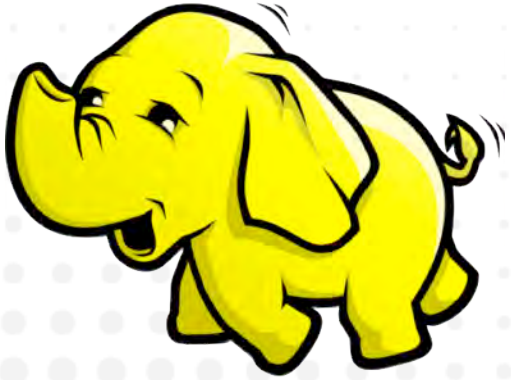

Practical Hadoop by Example

for relational database professioanals

Alex Gorbachev
12-Mar-2013
New York, NY

Pythian
love your data

Alex Gorbachev

ORACLE RACSIG

Sydney ORACLE Meetup

- Chief Technology Officer at Pythian
- Blogger
- OakTable Network member
- Oracle ACE Director
- Founder of BattleAgainstAnyGuess.com
- Founder of Sydney Oracle Meetup
- IOUG Director of Communities
- EVP, Ottawa Oracle User Group

ORACLE ACE Director

OakTable.net

IOUG
independent oracle users group

ORACLE Platinum Partner

Pythian
love your data

Why Companies Trust Pythian

Recognized Leader:

- Global industry-leader in remote database administration services and consulting for Oracle, Oracle Applications, MySQL and SQL Server
- Work with over 150 multinational companies such as Forbes.com, Fox Interactive media, and MDS Inc. to help manage their complex IT deployments

Expertise:

- One of the world's largest concentrations of dedicated, full-time DBA expertise.

Global Reach & Scalability:

- 24/7/365 global remote support for DBA and consulting, systems administration, special projects or emergency response

bridgewater SYSTEMS

TOYOTA

MDS Science advancing health

register.com

ORACLE Platinum Partner

Penn

FOX SPORTS

MySQL Enterprise Authorized Platinum REMOTE DBA SERVICES PARTNER

ORACLE Platinum Partner

Agenda

- What is Big Data?
- What is Hadoop?
- Hadoop use cases
- Moving data in and out of Hadoop
- Avoiding major pitfalls

ORACLE Platinum Partner

What is Big Data?

Pythian
love your data

Doesn't Matter.

We are here to discuss data architecture and use cases.
Not define market segments.

What Does Matter?

Some data types are a bad fit for RDBMS.
Some problems are a bad fit for RDBMS.

We can call them BIG if you want.
Data Warehouses have always been BIG.

Given enough skill and money –
Oracle can do anything.

Lets talk about efficient solutions.

When RDBMS Makes no Sense?

- Storing images and video
- Processing images and video
- Storing and processing other large files
  - PDFs, Excel files
- Processing large blocks of natural language text
  - Blog posts, job ads, product descriptions
- Processing semi-structured data
  - CSV, JSON, XML, log files
  - Sensor data

When RDBMS Makes no Sense?

- Ad-hoc, exploratory analytics
- Integrating data from external sources
- Data cleanup tasks
- Very advanced analytics (machine learning)

New Data Sources

- Blog posts
- Social media
- Images
- Videos
- Logs from web applications
- Sensors

They all have large potential value

But they are awkward fit for traditional data warehouses

Big Problems with Big Data

- It is:
  - Unstructured
  - Unprocessed
  - Un-aggregated
  - Un-filtered
  - Repetitive
  - Low quality
  - And generally messy.

Oh, and there is a lot of it.

Technical Challenges

- Storage capacity
- Storage throughput
- Pipeline throughput
- Processing power
- Parallel processing
- System Integration
- Data Analysis

Scalable storage

Massive Parallel Processing

Ready to use tools

Big Data Solutions

Real-time transactions at very high scale, always available, distributed

- Relaxing ACID rules
  - Atomicity
  - Consistency
  - Isolation
  - Durability

  Example: eventual consistency in Cassandra

ORACLE NOSQL DATABASE

mongoDB

CouchDB relax

Cassandra

riak

redis

Analytics and batch-like workload on very large volume often unstructured

- Massively scalable
  - Throughput oriented
  - Sacrifice efficiency for scale

hadoop

Hadoop is most industry accepted standard / tool

ORACLE Platinum Partner

What is Hadoop?

Pythian
love your data

Hadoop Principles

Bring Code to Data

Share Nothing

Hadoop in a Nutshell

hadoop HDFS

Google

Replicated Distributed Big-Data File System

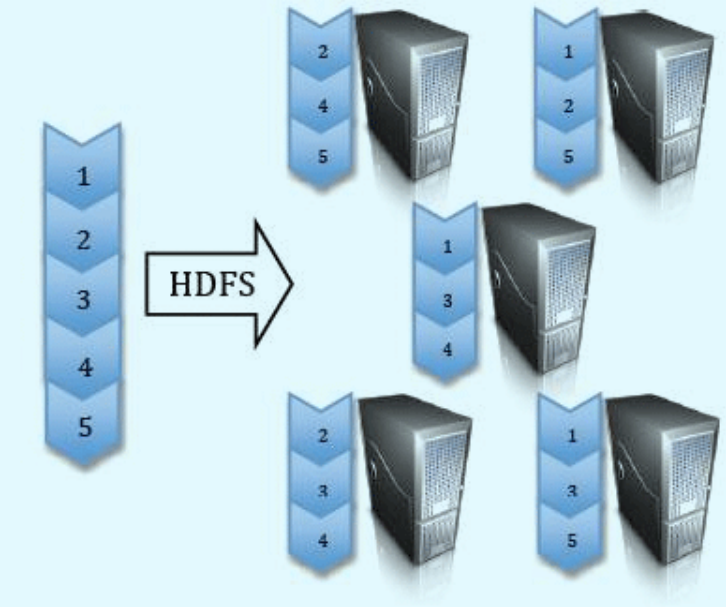

Map-Reduce - framework for writing massively parallel jobs

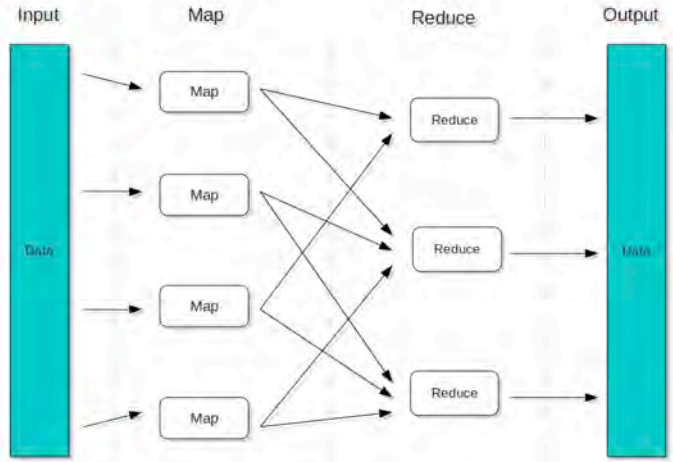

HDFS architecture

simplified view

- Files are split in large blocks
- Each block is replicated on write
- Files can be only created and deleted by one client
  - Uploading new data? => new file
    - Append supported in recent versions
  - Update data? => recreate file
  - No concurrent writes to a file
- Clients transfer blocks directly to & from data nodes
- Data nodes use cheap local disks
- Local reads are efficient

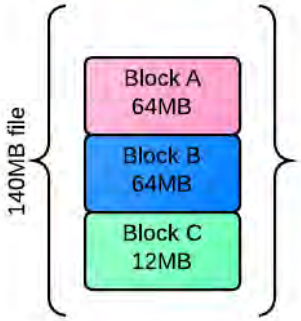

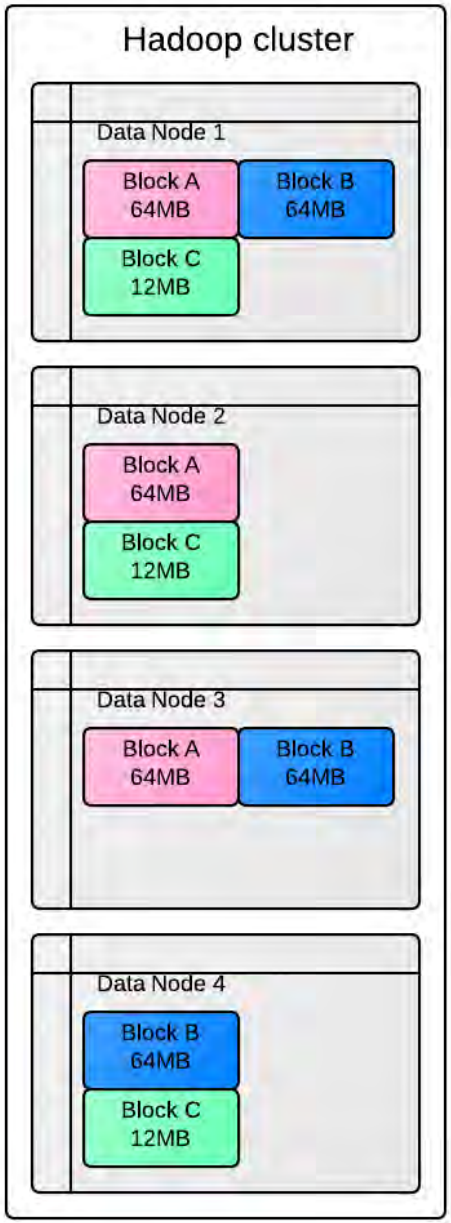

HDFS design principles

Map Reduce example *histogram calculation*

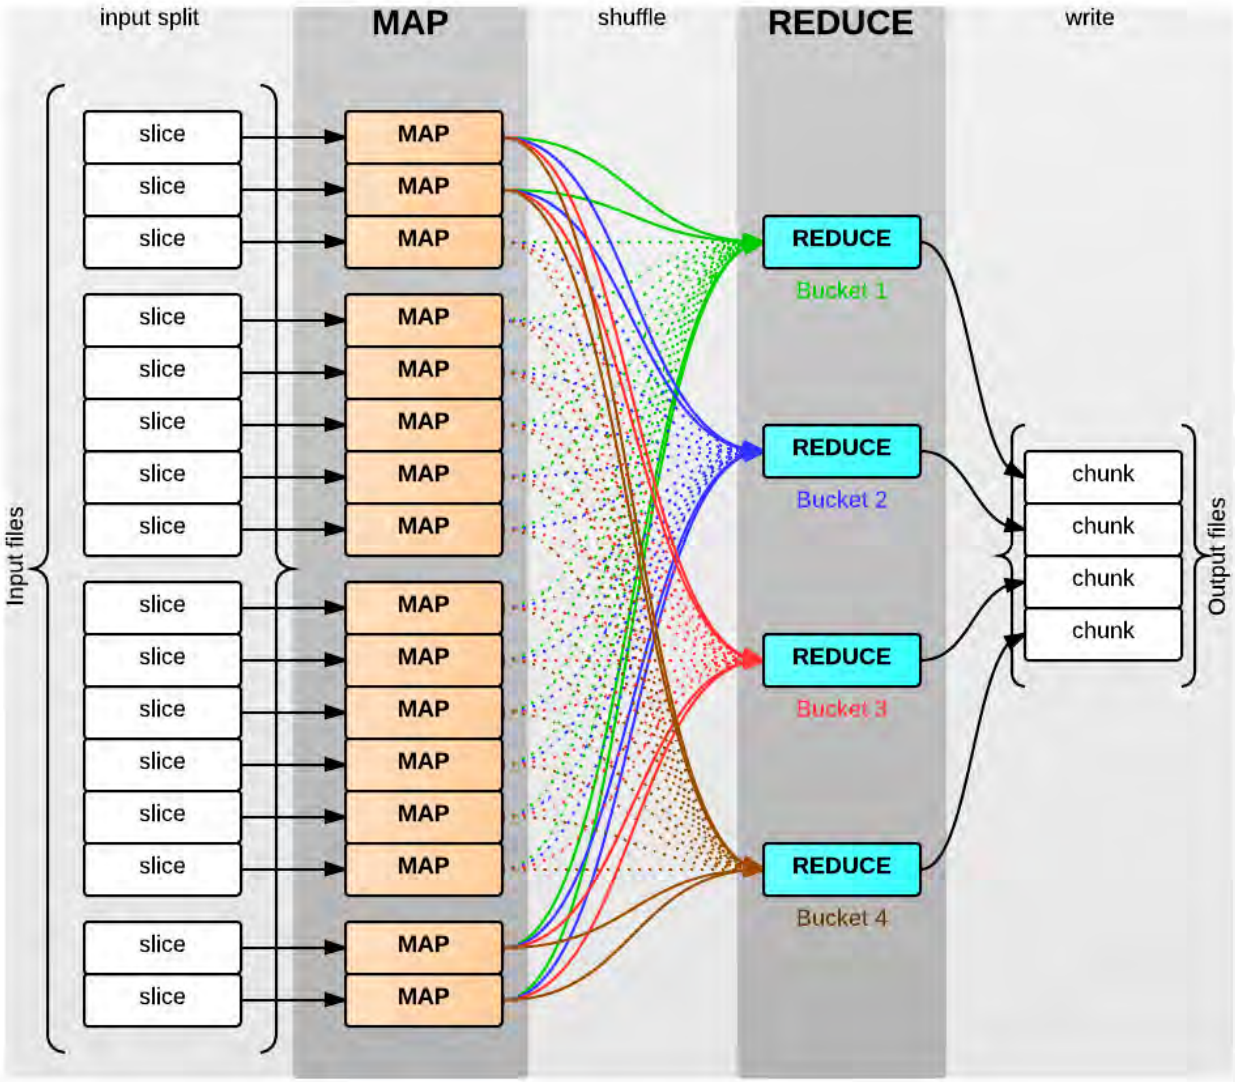

Map Reduce pros & cons

Advantages

- Very simple
- Flexible
- Highly scalable
  - Good fit for HDFS – mappers read locally
- Fault tolerant

Pitfalls

- Low efficiency
  - Lots of intermediate data
  - Lots of network traffic on shuffle
- Complex manipulation requires pipeline of multiple jobs
- No high-level language
- Only mappers leverage local reads on HDFS

Main components of Hadoop ecosystem

- Hive – HiveQL is SQL like query language
  - Generates MapReduce jobs
- Pig – data sets manipulation language (like create your own query execution plan)
  - Generates MapReduce jobs
- Zookeeper – distributed cluster manager
- Oozie – workflow scheduler services
- Sqoop – transfer data between Hadoop and relational

Non-MR processing on Hadoop

- HBase – columnar-oriented key-value store (NoSQL)
- SQL without Map Reduce
  - Impala (Cloudera)
  - Drill (MapR)
  - Phoenix (Salesforce.com)
  - Hadapt (commercial)
- Shark – Spark in-memory analytics on Hadoop

Hadoop Benefits

- Reliable solution based on unreliable hardware
- Designed for large files
- **Load data first, structure later**
- Designed to maximize throughput of large scans
- Designed to leverage parallelism
- Designed to scale
- Flexible development platform
- **Solution Ecosystem**

Hadoop Limitations

- Hadoop is scalable but not fast
- Some assembly required
  - Batteries not included
- Instrumentation not included either
- DIY mindset (remember MySQL?)

How much does it cost?

$300K DIY on SuperMicro

- 100 data nodes
- 2 name nodes
- 3 racks
- 800 Sandy Bridge CPU cores
- 6.4 TB RAM
- 600 x 2TB disks
- 1.2 PB of raw disk capacity
- 400 TB usable (triple mirror)
- Open-source s/w

ORACLE Platinum Partner

Hadoop Use Cases

Pythian
love your data

Use Cases for Big Data

- Top-line contributions
  - Analyze customer behavior
    - Optimize ad placements
    - Customized promotions and etc
  - Recommendation systems
    - Netflix, Pandora, Amazon
  - Improve connection with your customers
    - Know your customers – patterns and responses
- Bottom-line contributors
  - Cheap archives storage
  - ETL layer – transformation engine, data cleansing

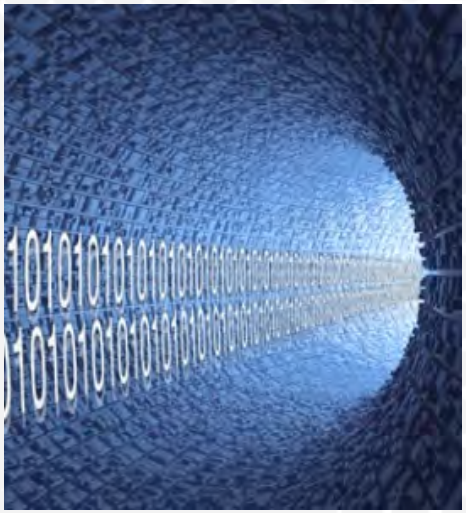

Typical *Initial* Use-Cases for Hadoop in modern Enterprise IT

- Transformation engine (part of ETL)
  - Scales easily
  - Inexpensive processing capacity
  - Any data source and destination

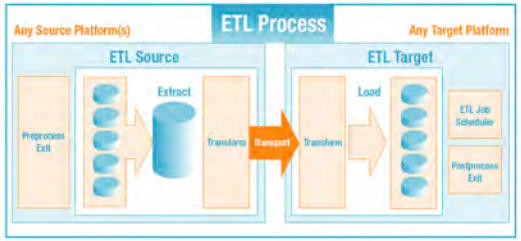

- Data Landfill
  - Stop throwing away any data
  - Don't know how to use data today? Maybe tomorrow you will
  - Hadoop is very inexpensive but very reliable

Advanced: Data Science Platform

- Data warehouse is good when questions are known, data domain and structure is defined
- Hadoop is great for seeking new meaning of data, new types of insights
- Unique information parsing and interpretation
- Huge variety of data sources and domains

- When new insights are found and new structure defined, Hadoop often takes place of ETL engine
- Newly structured information is then loaded to more traditional data-warehouses (still today)

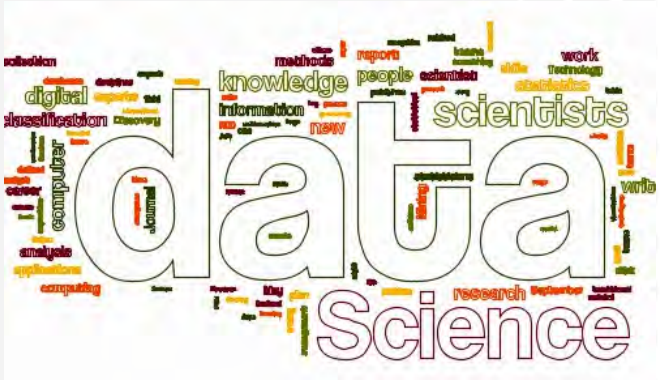

Pythian Internal Hadoop Use

- OCR of screen video capture from Pythian privileged access surveillance system
  - Input raw frames from video capture
  - Map-Reduce job runs OCR on frames and produces text
  - Map-Reduce job identifies text changes from frame to frame and produces text stream with timestamp when it was on the screen
  - Other Map-Reduce jobs mine text (and keystrokes) for insights
    - Credit Cart patterns
    - Sensitive commands (like DROP TABLE)
    - Root access
    - Unusual activity patterns
  - Merge with monitoring and documentation systems

ORACLE Platinum Partner

Hadoop in the Data Warehouse

Use Cases and Customer Stories

Pythian

love your data

ETL for Unstructured Data

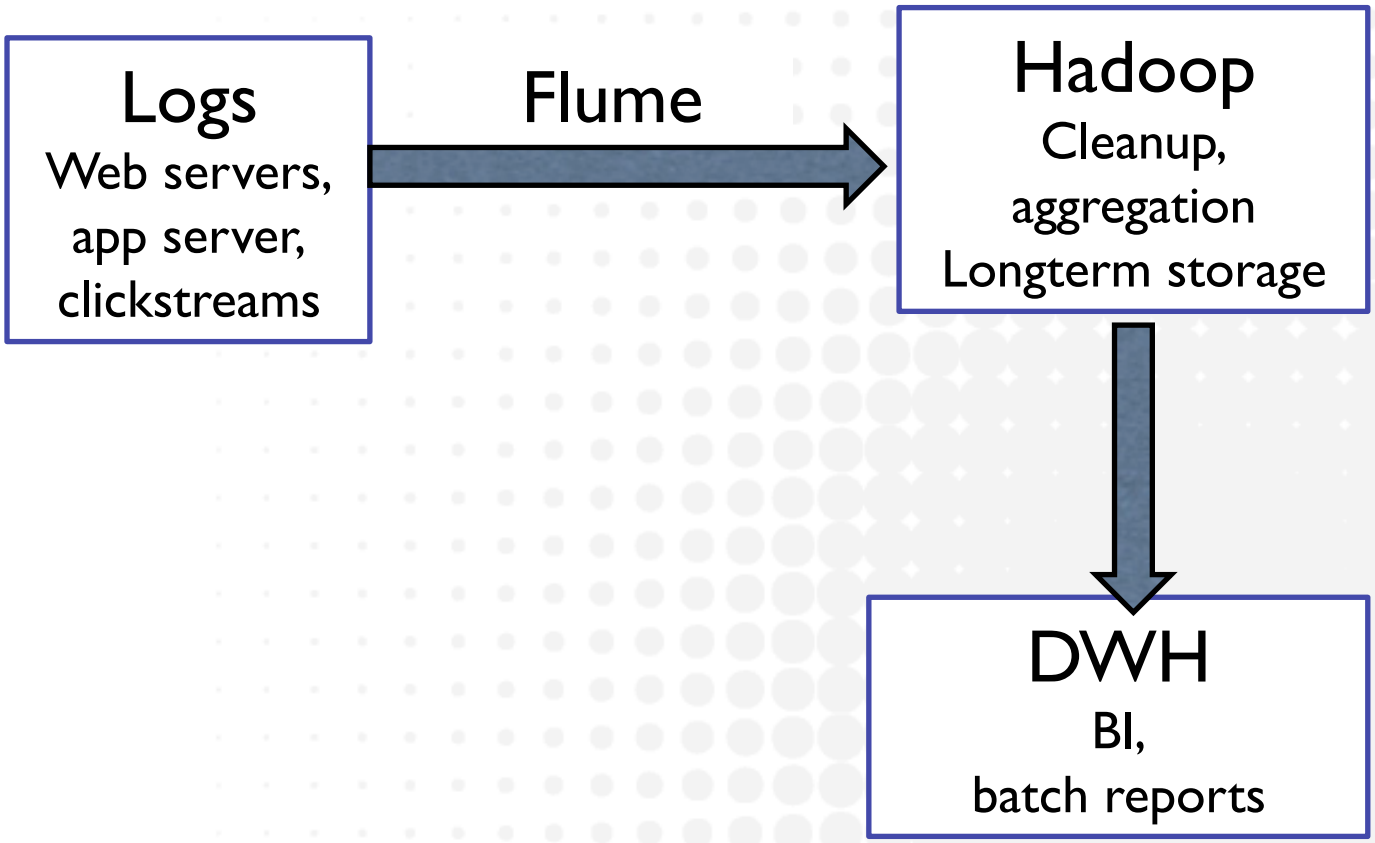

ETL for Structured Data

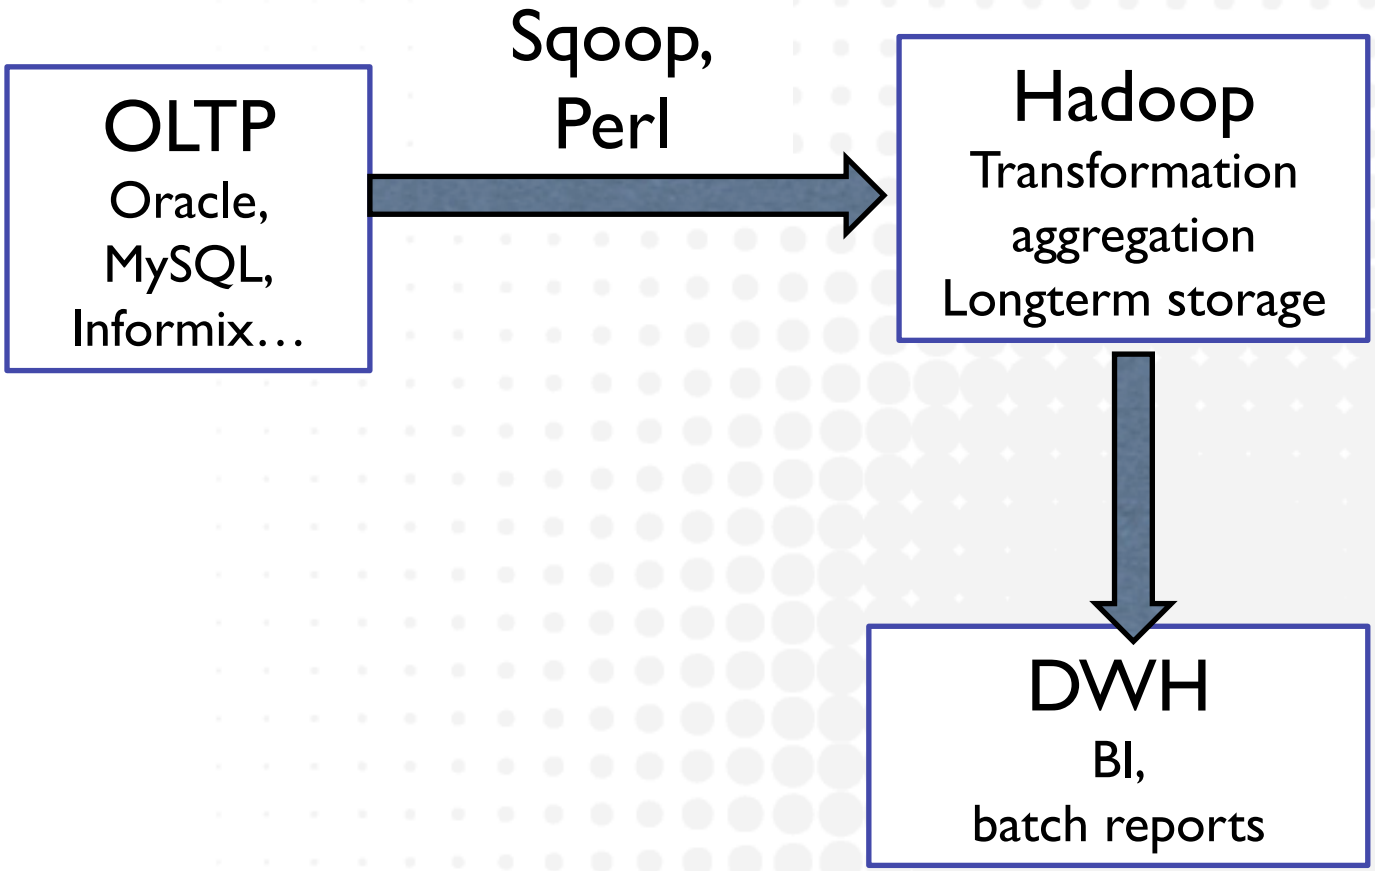

Bring the World into Your Datacenter

Rare Historical Report

Find Needle in Haystack

Hadoop for Oracle DBAs?

- alert.log repository
- listener.log repository
- Statspack/AWR/ASH repository
- trace repository
- DB Audit repository
- Web logs repository
- SAR repository
- SQL and execution plans repository
- Database jobs execution logs

ORACLE Platinum Partner

Connecting the (big) Dots

Pythian
love your data

Sqoop

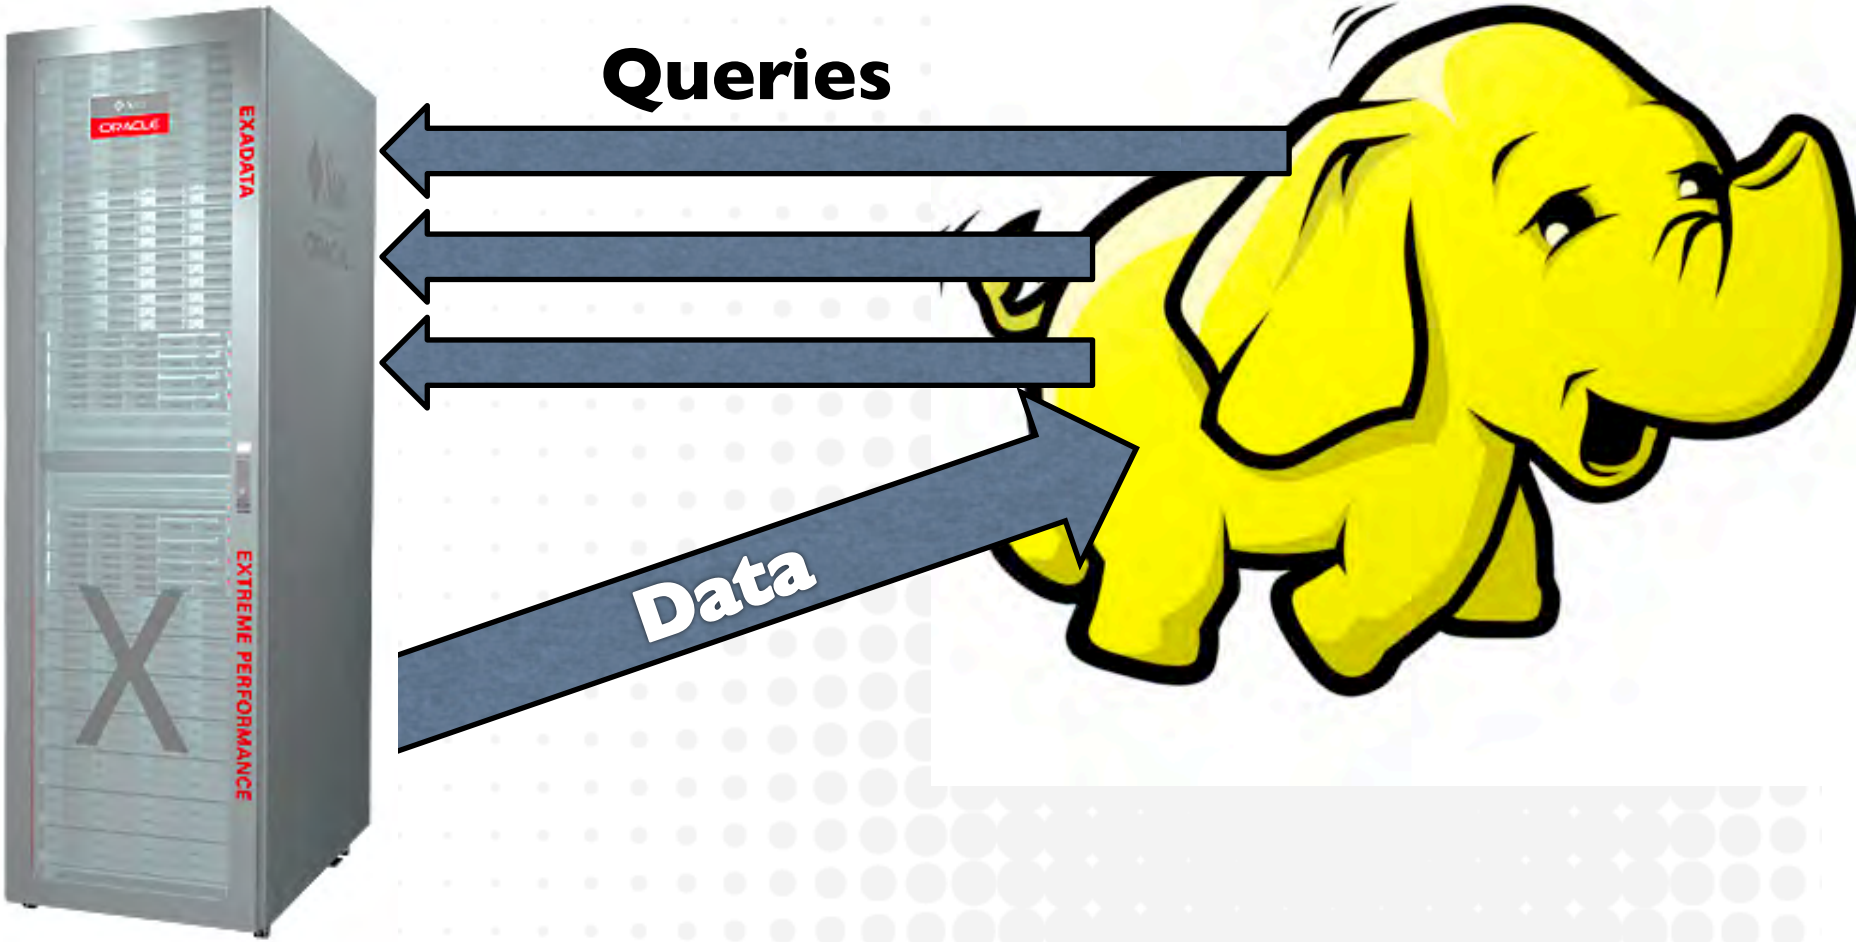

Sqoop is Flexible Import

- Select <columns> from <table> where <condition>
- Or <write your own query>
- Split column
- Parallel
- Incremental
- File formats

Sqoop Import Examples

- Sqoop import --connect jdbc:oracle:thin:@//dbserver:1521/masterdb
  --username hr --table emp
  --where “start_date > ’01-01-2012’”

- Sqoop import jdbc:oracle:thin:@//dbserver:1521/masterdb
  --username myuser
  --table shops --split-by shop_id
  --num-mappers 16

Must be indexed or partitioned to avoid 16 full table scans

Less Flexible Export

- 100 row batch inserts
- Commit every 100 batches
- Parallel export
- Merge vs. Insert

Example:

```
sqoop export
--connect jdbc:mysql://db.example.com/foo
--table bar
--export-dir /results/bar_data
```

FUSE-DFS

- Mount HDFS on Oracle server:
  - sudo yum install hadoop-0.20-fuse
  - hadoop-fuse-dfs dfs://<name_node_hostname>:<namenode_port> <mount_point>
- Use external tables to load data into Oracle
- File Formats may vary
- All ETL best practices apply

Oracle Loader for Hadoop

- Load data from Hadoop into Oracle
- Map-Reduce job inside Hadoop
- Converts data types, partitions and sorts
- Direct path loads
- Reduces CPU utilization on database

- NEW:
  - Support for Avro
  - Support for compression codecs

Oracle Direct Connector to HDFS

- Create external tables of files in HDFS
- `PREPROCESSOR HDFS_BIN_PATH:hdfs_stream`
- All the features of External Tables
- Tested (by Oracle) as 5 times faster (GB/s) than FUSE-DFS

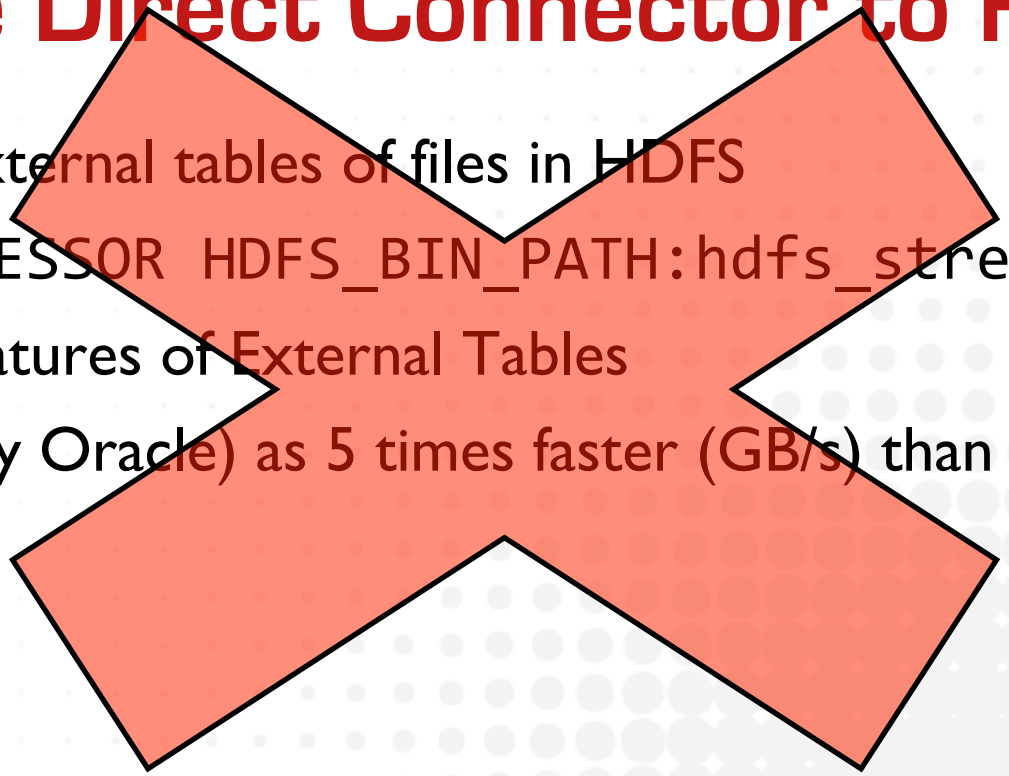

Oracle SQL Connector for HDFS

- Map-Reduce Java program
- Creates an external table
- Can use Hive Metastore for schema
- Optimized for parallel queries
- Supports Avro and compression

ORACLE Platinum Partner

How not to Fail

Pythian
love your data

Data That Belong in RDBMS

Prepare for Migration

Use Hadoop Efficiently

- Understand your bottlenecks:
  - CPU, storage or network?
- Reduce use of temporary data:
  - All data is over the network
  - Written to disk in triplicate.
- Eliminate unbalanced workloads
- Offload work to RDBMS
- Fine-tune optimization with Map-Reduce

Your Data
is NOT
as BIG
as you think

Getting started

- Pick a business problem
- Acquire data
- Get the tools: Hadoop, R, Hive, Pig, Tableau
- Get platform: can start cheap
- Analyze data
  - Need Data Analysts a.k.a. Data Scientists

- Pick an operational problem
  - Data store
  - ETL
- Get the tools: Hadoop, Sqoop, Hive, Pig, Oracle Connectors
- Get platform: Ops suitable
- Operational team

Continue Your Education

COLLABORATE13

TECHNOLOGY AND APPLICATIONS FORUM
FOR THE ORACLE COMMUNITY

Yeah, it's kind of a big deal

Learn More about IOUG's 20th Anniversary Plans

www.collaborate13.ioug.org

Thank you & Q&A

To contact us...

sales@pythian.com

1-866-PYTHIAN

To follow us...

http://www.pythian.com/news/

http://www.facebook.com/pages/The-Pythian-Group/

http://twitter.com/pythian

http://www.linkedin.com/company/pythian