

Ingesting Data from Oracle Database 12c R2 into a Hadoop System

Building a Data Lake on Linux for Power

Mithun Hr – IBM Systems, India

Wayne T. Martin – IBM Systems, USA

Arun Sar – IBM Systems, USA

Oracle OpenWorld 2017

© 2017 IBM Corporation



Agenda

IBM and Oracle Collaboration

IBM Power Systems

Hortonworks Data Platform

Attunity Replicate

Experiences Ingesting Oracle Data into Hadoop on Power Systems

Conclusion and Acknowledgments



Agenda

IBM and Oracle Collaboration

IBM Power Systems

Hortonworks Data Platform

Attunity Replicate

Experiences Ingesting Oracle Data into Hadoop on Power Systems

Conclusion and Acknowledgments



IBM and Oracle, a shared commitment to Client Value



Sustained Collaboration for Clients

Oracle 30+ years, PeopleSoft 20+ Years, JD Edwards 35+ years, Siebel 15+ years

Mutual Executive Commitment

Dedicated, Executive-led Alliance teams

Senior Executive reviews and functional Management interlocks

More than 100K joint Technology Clients

Hardware and Software support via in-depth certification process

Award-winning Oracle Services Practice

More than 5,500 successful joint services projects, 13,000 skilled Oracle resources worldwide

Vibrant technology Collaboration

Continued joint development delivering Oracle software optimized for IBM hardware, significant skills/resources investment, dedicated International Competency Centers

Cooperative Client Support Process

Dedicated resources and significant program investments

Agenda

IBM and Oracle Collaboration

IBM Power Systems

Hortonworks Data Platform

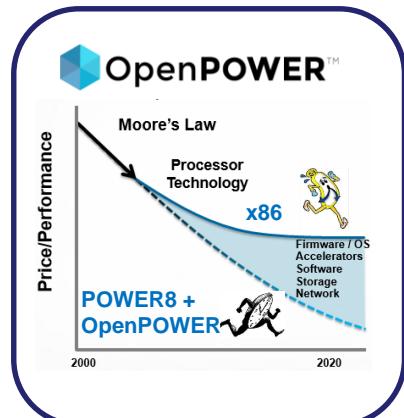
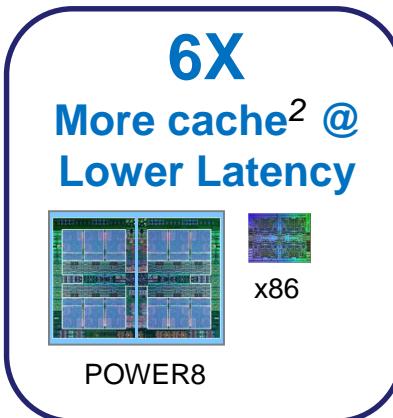
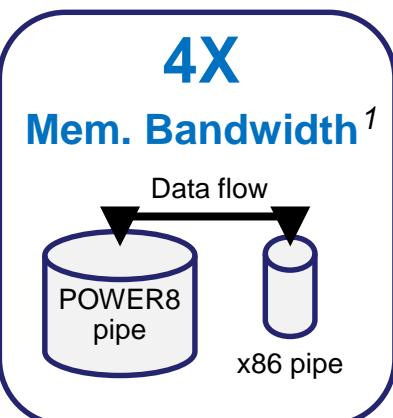
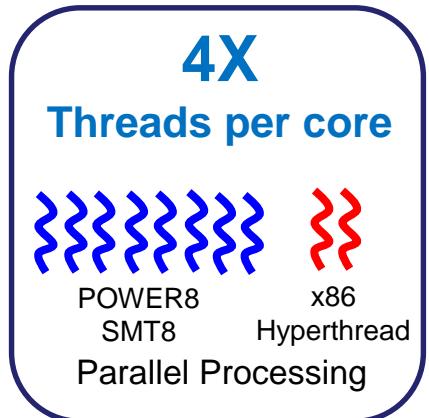
Attunity Replicate

Experiences Ingesting Oracle Data into Hadoop on Power Systems

Conclusion and Acknowledgments



POWER8: Designed for data to deliver breakthrough performance



These design decisions result in best performance for data centric workloads like:
Spark, Hadoop, Database, NoSQL, Big Data Analytics, OLTP

SMT=Simultaneous Multi-Threading
OLTP = On-Line Transaction Processing

1. Up to 4X depending on specific x86 and POWER8 servers being compared
2. Up to 6X more cache comparing Intel e7-8890 servers to 12 core POWER8 servers. See speaker notes for more details

The OpenPOWER Foundation

Market Shifts

- Moore's law no longer satisfies performance gain
- Growing workload demands
- Numerous IT consumption models
- Mature Open software ecosystem



OpenPOWER™

Open Development

open software, open hardware



Performance of POWER architecture

amplified capability

New Open Innovation

- Rich software ecosystem
- Spectrum of power servers
- Multiple hardware options
- Derivative POWER chips

The OpenPOWER Foundation is an open ecosystem, using the POWER Architecture to serve the evolving needs of customers.

250+

Members
And growing



IBM Power Systems is open to the core

OpenPOWER



Open Source Workloads



OpenCAPI

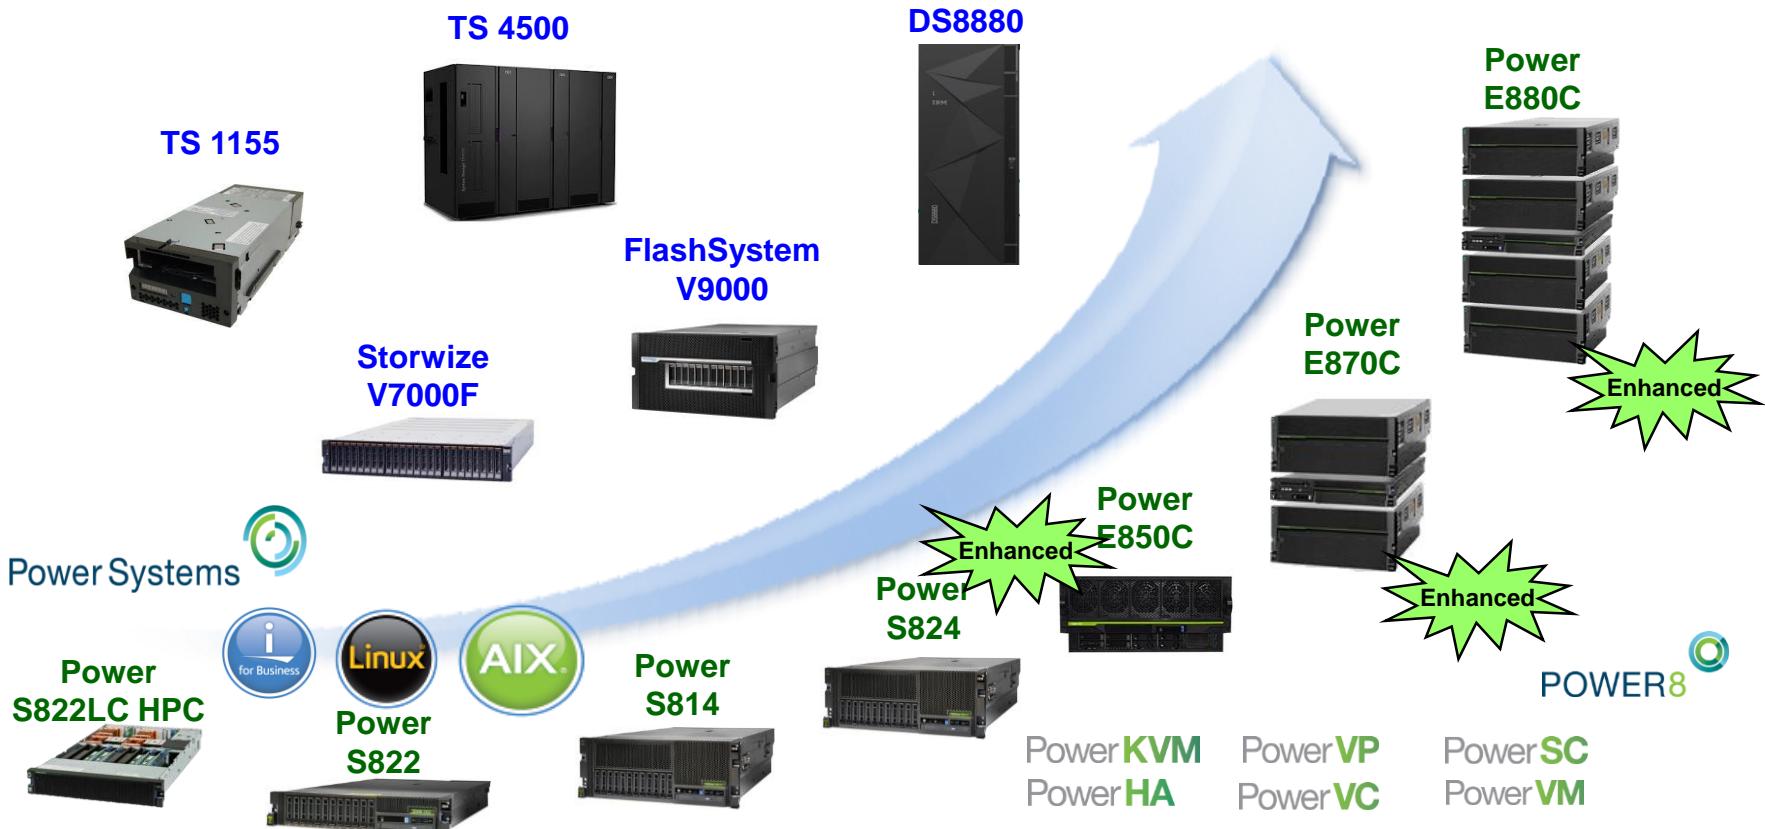


Open Frameworks



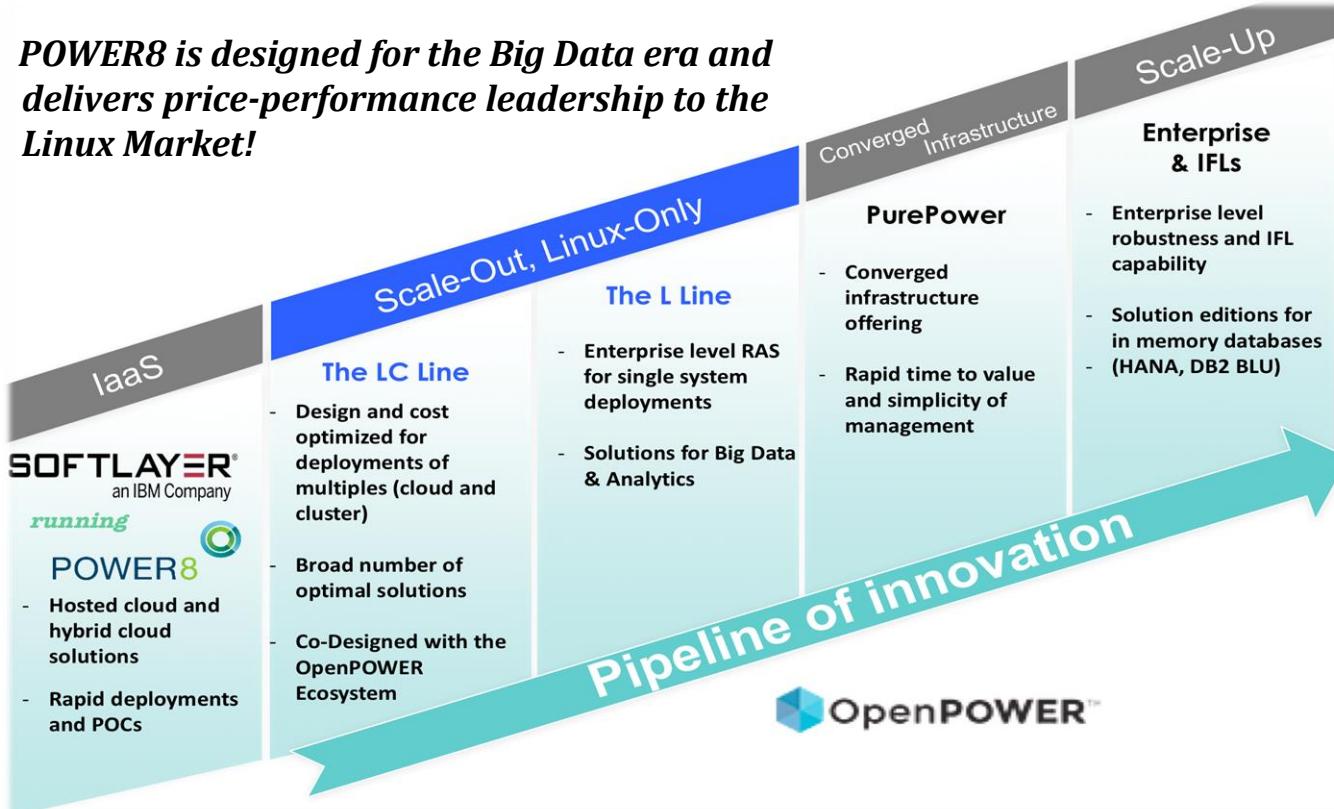
Making Machine Learning and AI more affordable

IBM Systems and Storage leverage Oracle Software



The IBM Power Systems Linux Portfolio

POWER8 is designed for the Big Data era and delivers price-performance leadership to the Linux Market!



Broad Linux portfolio delivers all your Linux deployment needs

IBM Support



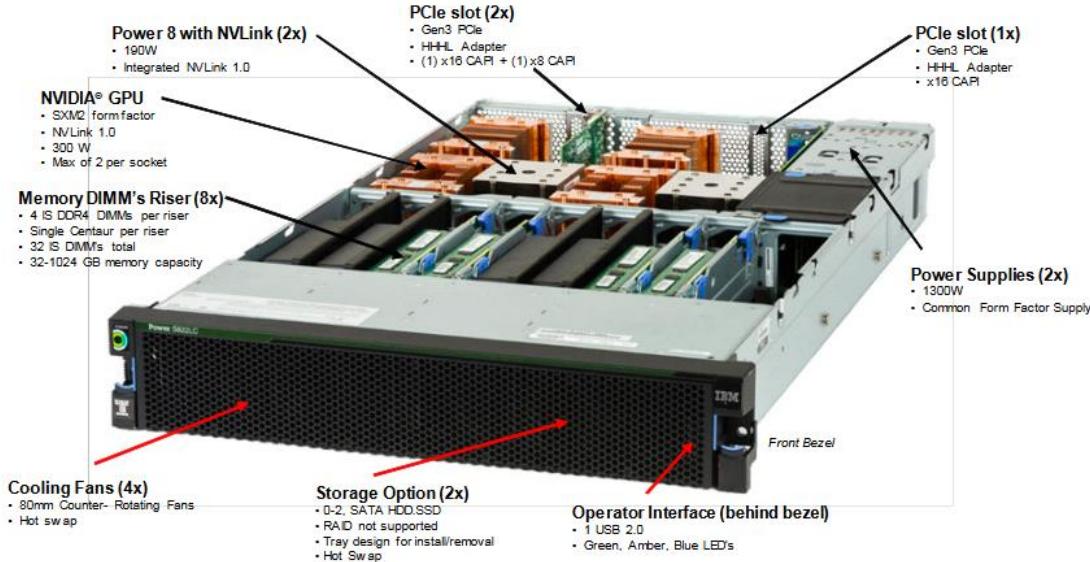
Community / 3rd Party Support



ML/DL Power Server: S822LC for Big Data and HPC

- Included:
- 2 POWER8 with NVLink
- 4 NVIDIA Tesla P100 GPUs
- 256 GB System Memory
- 2 SSD storage devices
- High-speed interconnect (IB or Ethernet, depending on infrastructure)
- Optional:
- Up to 1 TB System Memory
- PCIe attached NVMe storage

2 Socket, 4GPU System with NVLink



Optimized configuration for Big Data workloads

PowerAI Software Distribution

Deep
Learning
Frameworks

Caffe

 Caffe

IBM Caffe

 torch

 TensorFlow™

theano

 Chainer

Supporting
Libraries

DIGITS

OpenBLAS

Distributed
Frameworks

Bazel

NCCL

IBM Power System for HPC, with NVLink

Breakthrough performance for GPU accelerated applications,
including Deep Learning and Machine Learning.



Agenda

IBM and Oracle Collaboration

IBM Power Systems

Hortonworks Data Platform

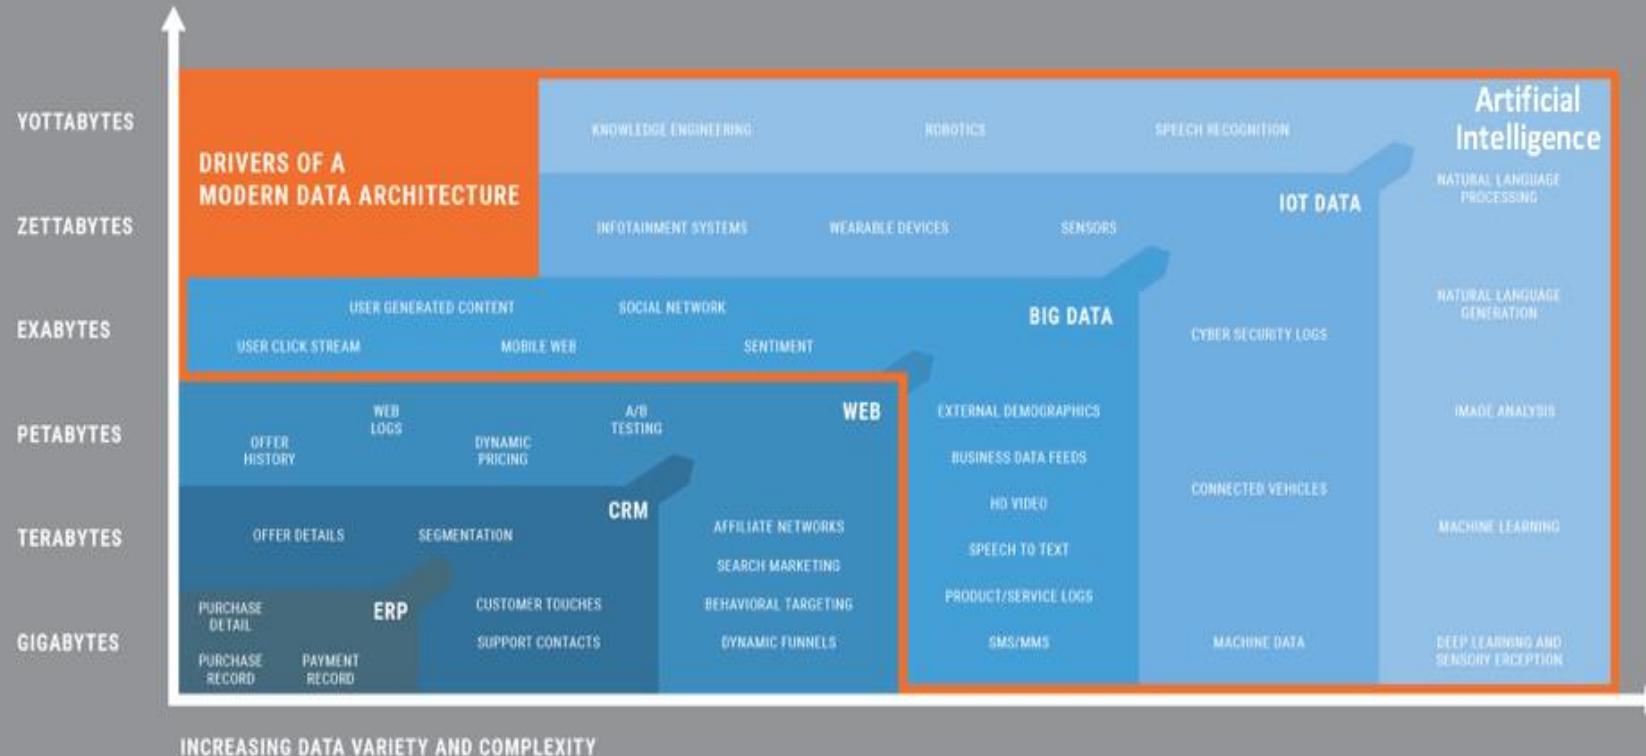
Attunity Replicate

Experiences Ingesting Oracle Data into Hadoop on Power Systems

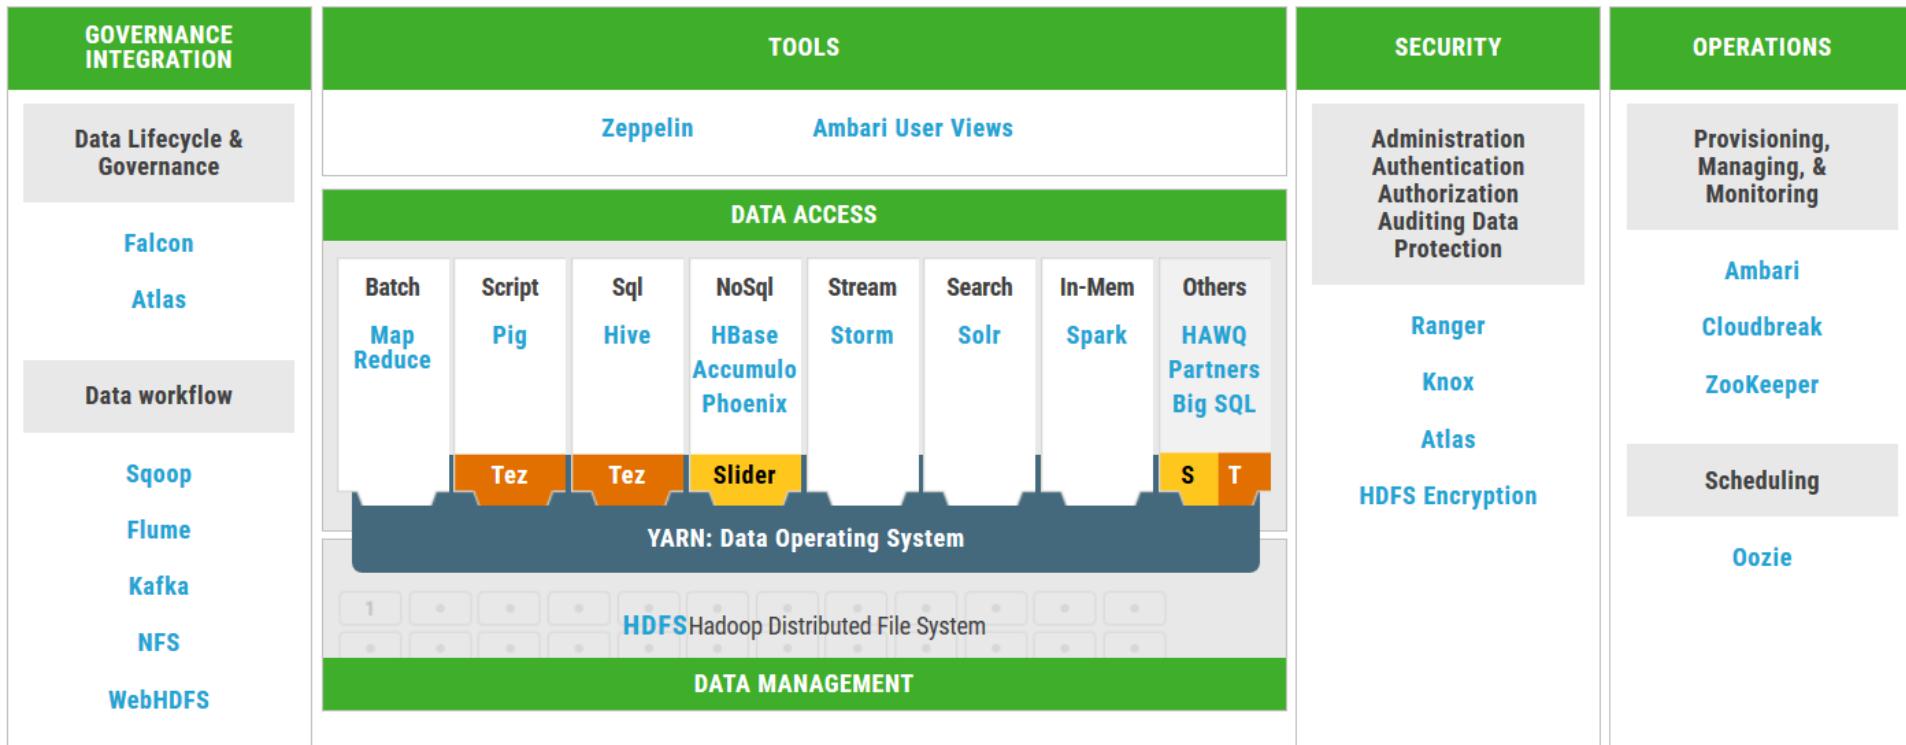
Conclusion and Acknowledgments



The Data Tipping Point that drives Data Lake creation

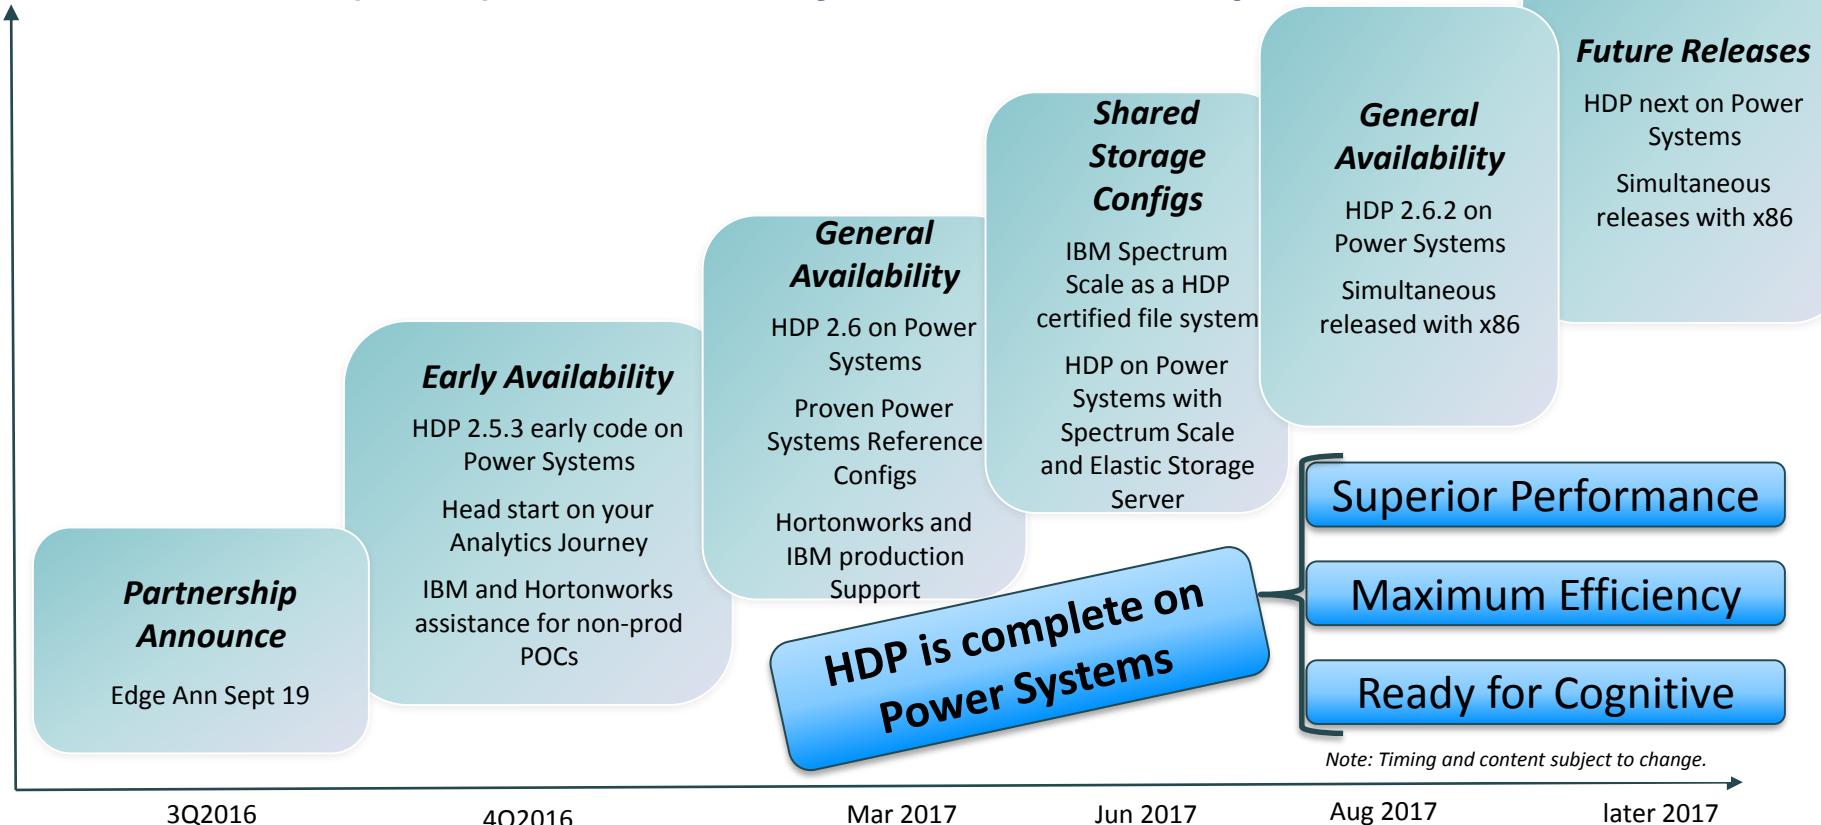


Hortonworks Data Platform, complete on Power Systems



100% Open Source Connected Data Platform

Hortonworks (HDP) on Power Systems Roadmap



IBM and Hortonworks partnered to Deliver:



Data Science Experience (DSX)

Interactive, collaborative, cloud-based environment where data scientists can use multiple tools to activate their insights.

[LEARN MORE](#)



Big SQL

Hybrid SQL engine for Hadoop that can connect to disparate sources using a single database connection or query.

[LEARN MORE](#)



Power Systems

High performance servers built with open technologies and designed for mission-critical applications

[LEARN MORE](#)



Spectrum Scale

Data storage that transforms economics for traditional and next-generation applications with great speed, agility and efficiency.

[LEARN MORE](#)

<https://hortonworks.com/partner/ibm/>

Flexibility with HDP on Power Systems

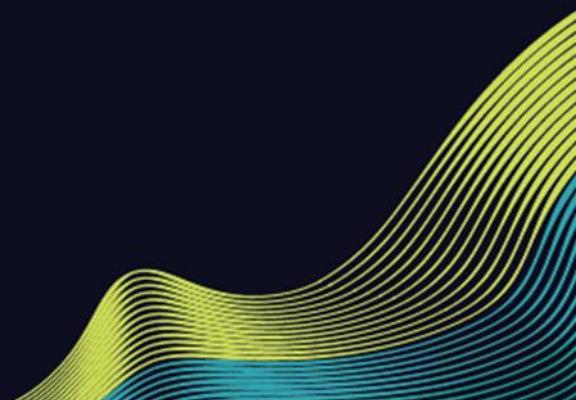
- **Scale Up or Out to Meet Evolving Workloads**
 - Scale up each node by exploiting the memory bandwidth and multi-threading
 - 4X threads per core vs x86 allows you to optimize and drive more workload per node
 - 4X memory bandwidth vs x86 gives you more options as your workloads expand and evolve
- **Unmatched Range of Linux Servers**
 - From 1U, 16-core servers up to 16 socket, 192-core powerhouses with industry leading reliability
 - Virtualization options to host low cost dev environments or rich, multi-tenant private clouds
 - Wide range of OpenPOWER servers offered by OpenPower members for on-prem and the cloud
- **Accelerated Analytics**
 - Add accelerators (flash, GPU, FPGA) with direct access to processor memory through OpenCAPI

Upgrade your Big Data engine. Downsize your Big Data server cost.

Get a 3X price-performance advantage for HDP on
IBM Power Systems vs x86 – Guaranteed.

Act Now - Only Available until Dec 31, 2017!

Visit ibm.biz/powerofhortonworks



Agenda

IBM and Oracle Collaboration

IBM Power Systems

Hortonworks Data Platform

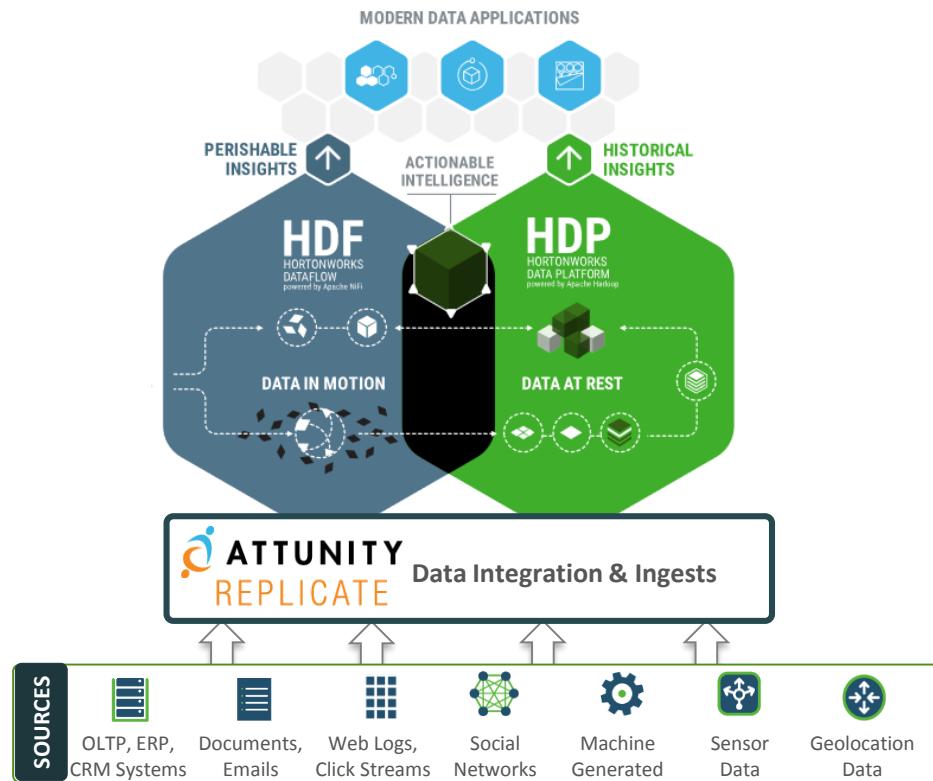
Attunity Replicate

Experiences Ingesting Oracle Data into Hadoop on Power Systems

Conclusion and Acknowledgments



The connected Data Architecture and Attunity



Attunity Replicate for HDP and HDF

Accelerate time-to-insights by delivering solutions faster, with fresher data, from many sources

- Automated data ingest
- Incremental data ingest (CDC)
- Broad support for many sources



Agenda

IBM and Oracle Collaboration

IBM Power Systems

Hortonworks Data Platform

Attunity Replicate

Experiences Ingesting Oracle Data into Hadoop on Power Systems

Conclusion and Acknowledgments



Modernize and expand your enterprise data

Enterprise
Data Warehouse



Off Load

Big Data
Administrators

Augment

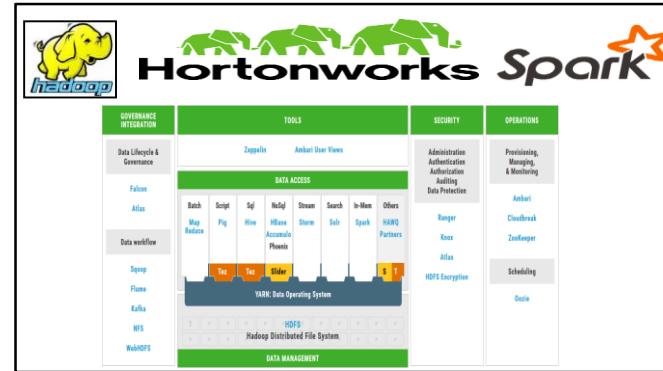


Live
Streams



Data Scientists
+
Marketing Managers

Explore, Discover, Analyze



Running on POWER8:
The Platform Designed for Big Data

4X
threads per
core vs. x86

4X
memory
bandwidth vs.
x86

6X
more cache
vs. x86

Hadoop Use Cases

- Data Ingestion
- Data Discovery
- EDW Optimization
- Predictive Analytics
- Single View
- And more ...



Build Your Data Lake on Linux on Power Systems

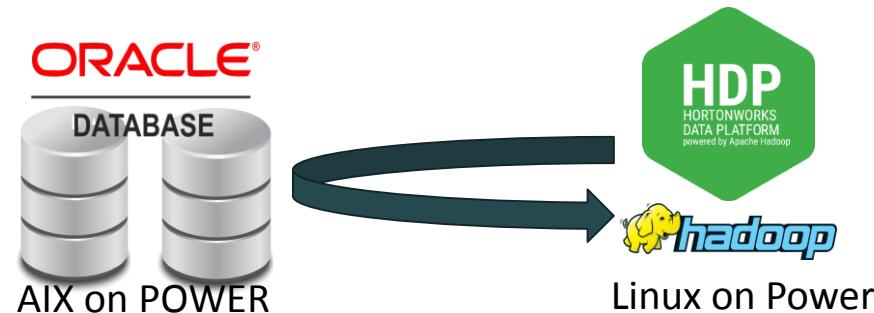
- Bulk Load
 - *Oracle Instant Client*
 - *Sqoop*
 - *Attunity Replicate*
- Change Data Capture
 - *Attunity Replicate*
- Federated Queries
 - *Apache Drill*

The diagram consists of three main components: 1) A large light blue speech bubble containing the text "Extract, Load and Transform data from database sources into Hadoop". 2) A smaller light blue speech bubble containing the text "Leave (some) data in place". 3) A list of data management tasks on the left, with arrows pointing from the text in the blue bubbles to specific items in the list. The "Extract, Load and Transform" bubble points to the "Bulk Load" and "Change Data Capture" items. The "Leave (some) data in place" bubble points to the "Federated Queries" item.

Tools for integrating Oracle and Hadoop on Power Systems are already in place ...

Oracle Instant client for Linux on Power connects Hadoop on Linux on Power to Oracle Databases –

Used to fetch Oracle data from Hadoop or issue a Federated Query from Apache Drill



Download the Oracle client here:

<http://www.oracle.com/technetwork/topics/linux-power-1e-2835260.html>

Tools for integrating Oracle and Hadoop on Power Systems are already in place ...

Bulk load with Sqoop:

- Copy Oracle JDBC driver into sqoop's lib directory

```
# ls /usr/hdp/2.6.0.0-598/sqoop/lib/ojdbc6.jar  
/usr/hdp/2.6.0.0-598/sqoop/lib/ojdbc6.jar
```

- Import data from Oracle database into HDFS

```
# /usr/bin/sqoop import --connect jdbc:oracle:thin:system/system@129.40.76.28:1521/ohdppdb.pbm.ihost.com --username SOE -P  
--table SOE.ADDRESSES  
...  
17/09/08 03:29:20 INFO mapreduce.ImportJobBase: Retrieved 913156293 records.
```

- Data Files in HDFS

```
$ hdfs dfs -ls  
Found 3 items  
drwx-----  - hdfs hdfs      0 2017-09-11 12:12 .Trash  
drwx-----  - hdfs hdfs      0 2017-09-08 03:29 .staging  
drwxr-xr-x  - hdfs hdfs      0 2017-09-08 03:29 SOE.ADDRESSES
```

- Import and Create database SOE in Hive similar to source database schema

```
$ sqoop import --connect jdbc:oracle:thin:system/system@129.40.76.28:1521/ohdppdb.pbm.ihost.com --username SOE -P --table  
SOE.PRODUCT_DESCRIPTIONS --hive-import  
Logging initialized using configuration in jar:file:/usr/hdp/2.6.0.0-598/hive/lib/hive-common-1.2.1000.2.6.0.0-598.jar!/hive-  
log4j.properties  
OK  
Time taken: 3.537 seconds  
Loading data to table soe.product_descriptions  
Table soe.product_descriptions stats: [numFiles=4, numRows=0, totalSize=111808, rawDataSize=0]  
OK  
Time taken: 1.058 seconds
```

Tools for integrating Oracle and Hadoop on Power Systems are already in place ...

Access Oracle data from Spark:

- On Hadoop node running spark, login as spark user and start the spark shell

```
# su - spark
$ cd /usr/hdp/2.6.0.0-598/spark/bin/
$ ./spark-shell
```

- To load the data

```
scala> val x =
  sqlContext.read.format("jdbc").option("driver","oracle.jdbc.driver.OracleDriver").
  option("url","jdbc:oracle:thin:soe/soe@129.40.76.28:1521/ohdppdb.pbm.ihost.com").o
  ption("dbtable", "soe.inventories").load();
x: org.apache.spark.sql.DataFrame = [PRODUCT_ID: decimal(6,0), WAREHOUSE_ID:
  decimal(6,0), QUANTITY_ON_HAND: decimal(8,0)]
scala> x.count()
...
res1: Long = 901478
```

Tools for integrating Oracle and Hadoop on Power Systems are already in place ...

Apache Drill for federated queries:

- Download apache drill from <https://drill.apache.org/>
- Oracle JDBC driver needs to be placed in /root/apache-drill-1.11.0/jars/3rdparty/ojdbc6.jar
- Modify `drill-override.conf` to include nodes in your cluster. Start apache drill on all the nodes

```
drill.exec: {  
  cluster-id: "cluster_id",  
  zk.connect: "zookeeper path"  
  drill.exec.sys.store.provider.local.path="ojdbc6.jar"  
}
```

- Apache drill console could be accessed by <http://hostname:8047/storage>
- SQL query to Oracle and Hive

```
select * from ORCL_DB.SOE.inventories os, NW_HIVE.inventories hv where  
os.warehouse_id=hv.warehouse_id and os.warehouse_id=221 limit 10;
```

Tools for integrating Oracle and Hadoop on Power Systems are already in place ...

Create storage plugin for Oracle and Hive

```
{  
  "type": "jdbc",  
  "driver": "oracle.jdbc.OracleDriver",  
  "url": "jdbc:oracle:thin:soe/soe@hostname:1521/SID",  
  "username": null,  
  "password": null,  
  "enabled": true  
}
```

```
{  
  "type": "hive",  
  "enabled": true,  
  "configProps": {  
    "hive.metastore.uris": "thrift://hostname:9083",  
    "javax.jdo.option.ConnectionURL":  
      "jdbc:mysql://hostname/hive",  
    "hive.metastore.warehouse.dir": "/tmp/drill_hive_wh",  
    "fs.default.name": "hdfs://hostname:8020",  
    "hive.metastore.sasl.enabled": "false"  
  }  
}
```

The screenshot shows the Apache Drill web interface. At the top, a navigation bar includes links for Apache Drill, Query, Profiles, Storage, Metrics, Threads, and Logs. Below the navigation bar, a preview pane displays a sample SQL query: `SELECT * FROM cp.`employee.json` LIMIT 20`. Under the 'Query Type' section, 'SQL' is selected. The 'Query' section contains the following SQL code:

```
select * from ORCL_DB.SOE.inventories os, NW_HIVE.inventories hv where os.warehouse_id=hv.warehouse_id and os.warehouse_id=221 limit 10;
```

A 'Submit' button is located at the bottom left of the query editor.

Change Data Capture for many sources/targets:

Attunity Replicate, Hortonworks HDP, IBM Power Systems

Automated data ingest



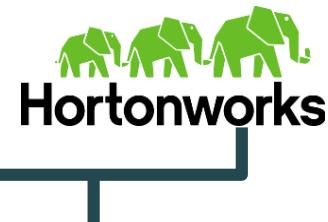
Rapid deployments of HUGE data lakes

Incremental updates with CDC

Continuous data refresh for RELEVANT analytics

Broad support for many enterprise data sources

COMPLETE datasets across DBs, DWs & mainframes



Data Sources Include
DB2, Hadoop, SQL Server, MySQL/MariaDB,
Oracle, PostgreSQL, SAP Sybase ASE,
Teradata Database, ODBC, Amazon RDS for
Aurora, MySQL, Oracle

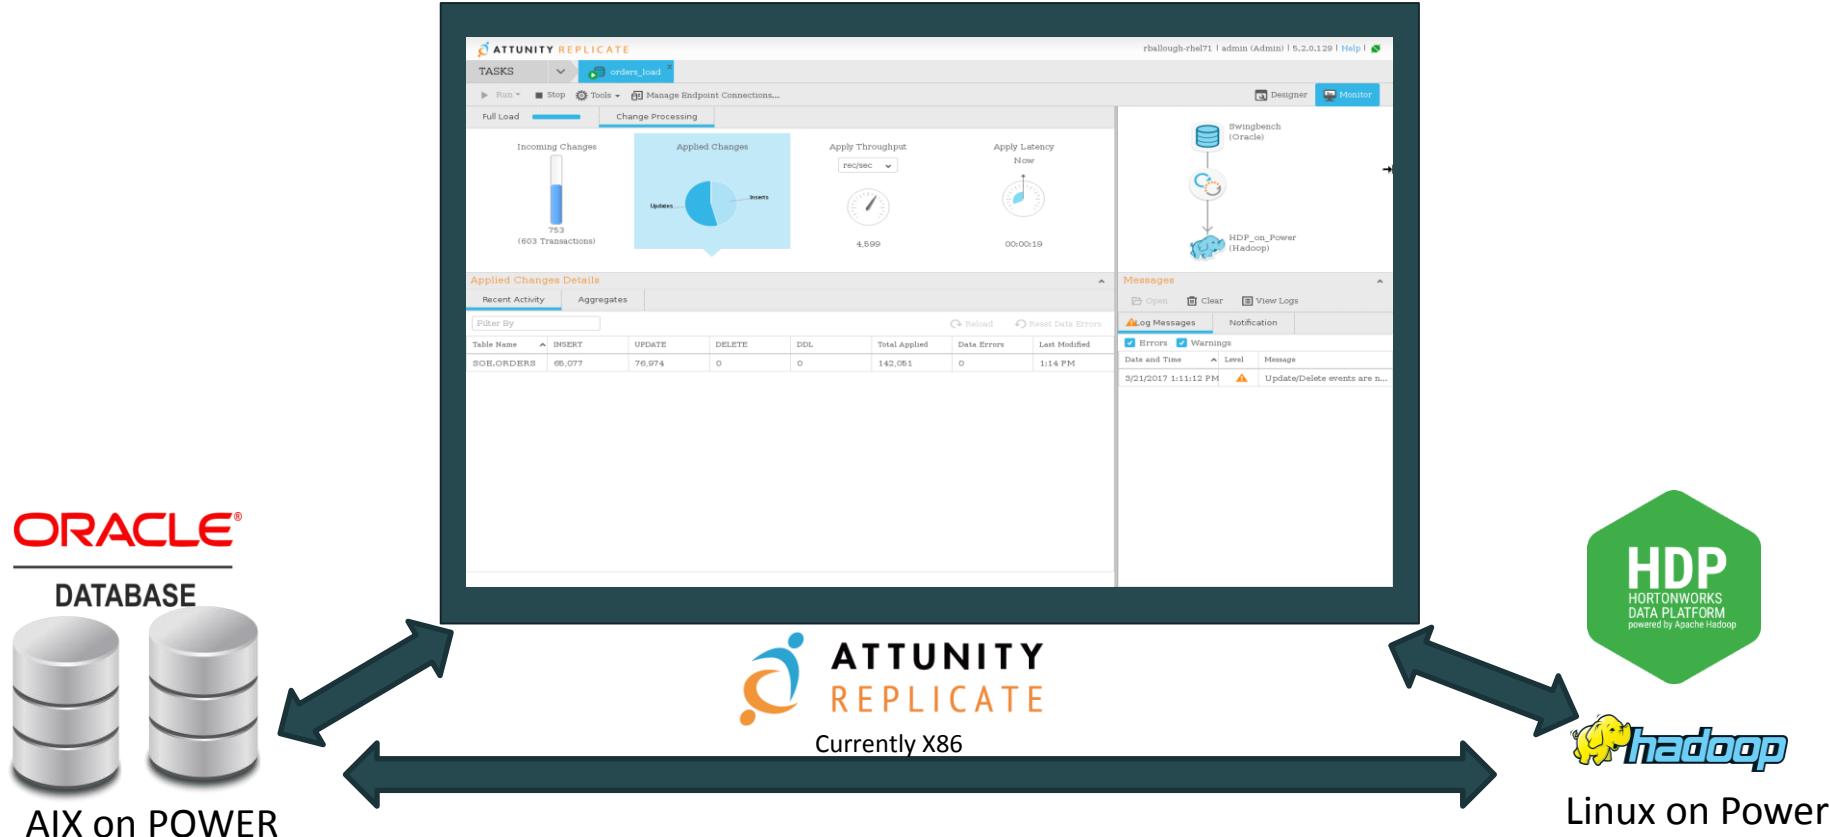
Running on POWER8:
The Platform Designed for Big Data

4X
threads per core

4X
memory bandwidth

6X
more cache

Experiences: Bring Oracle data into Hive using Attunity Replicate



Configuration Details

Oracle Database

- *Version 12cR2*
- *AIX 7.2 on POWER8*
- *ARCHIVELOG Mode is On*
- *Supplemental Logging is enabled*
- *Basic encryption, OLTP and HCC supported*
- *Required privileges to be granted be V\$xxx*

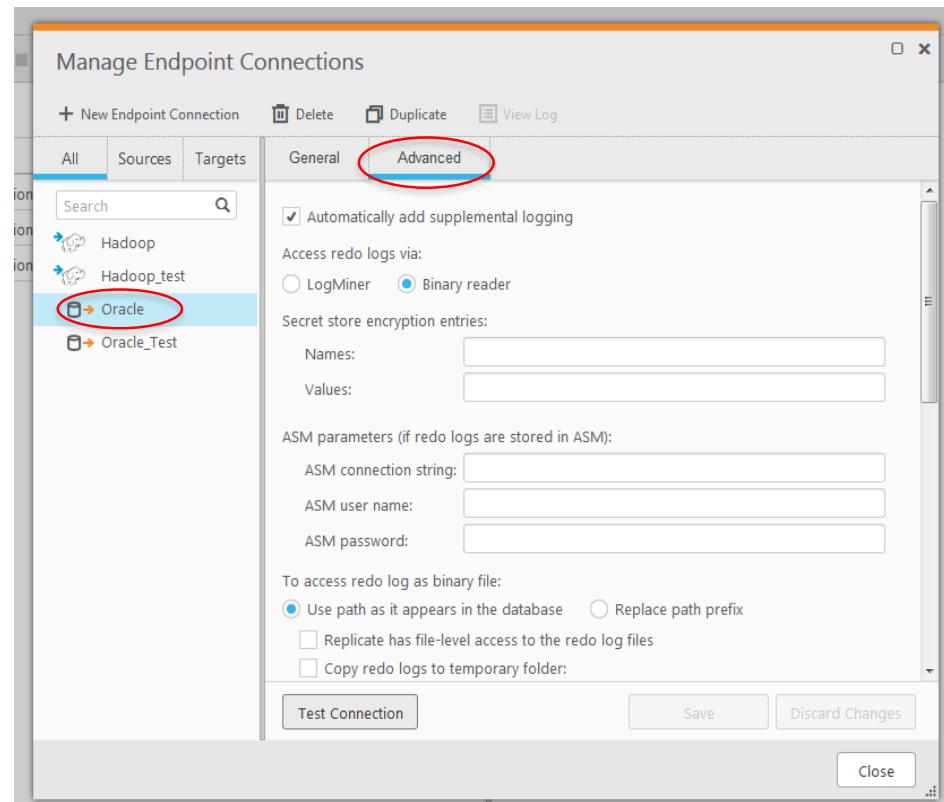
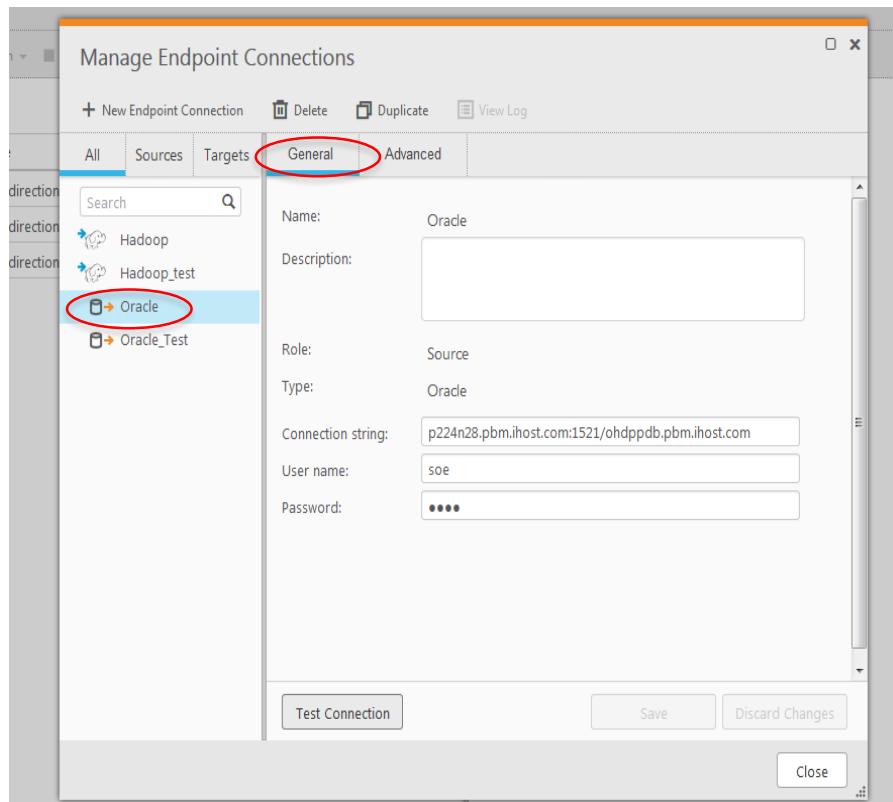
Attunity Replicate

- *Version 5.5.0.143*
- *RHEL 7.2*
- *Oracle Instant Client for Linux (x86-64)*
- *Minimum of 8 GB RAM, and 320 GB of disk space*
- *Mozilla Firefox Version 38 and above,*
- *Minimum Network 1 Gb interface*

Horotnworks HDP

- *Version 2.6.1*
- *RHEL 7.2 on POWER8*
- *Hadoop WebHDFS must be accessible*
- *Hadoop WebHCat service must be running*
- *Hadoop target settings must have access to Hcatalog*
- *Access Hive using ODBC 2.2 and above*

Attunity Replicate: Manage Endpoint – Oracle as source



Attunity Replicate: Manage Endpoint – Hadoop as target

The image displays six screenshots of the Attunity Replicate 'Manage Endpoint Connections' interface, illustrating the configuration of a Hadoop target endpoint. The interface is divided into 'General' and 'Advanced' tabs, with 'General' being the active tab in all shown configurations.

- Screenshot 1 (Top Left):** Shows the 'General' tab for a 'Hadoop' endpoint. It includes fields for 'Name' (Hadoop), 'Description', 'Role' (Target), 'Type' (Hadoop), and 'Security' settings (User name and password: 'hive', Password: '****').
- Screenshot 2 (Top Middle):** Shows the 'General' tab for a 'Hadoop' endpoint, but with 'Role' set to 'Target' and 'Type' set to 'Hadoop'. It includes 'Security' settings and 'HDFS' and 'Hive Access' options.
- Screenshot 3 (Top Right):** Shows the 'Advanced' tab for a 'Hadoop' endpoint. It includes 'File Format' and 'File Attributes' settings, such as 'Hadoop file system block size' (64 MB) and 'Maximum file size(KB)' (1000000).
- Screenshot 4 (Bottom Left):** Shows the 'General' tab for a 'Hadoop' endpoint, with 'Role' set to 'Target' and 'Type' set to 'Hadoop'. It includes 'Security' settings and 'HDFS' and 'Hive Access' options.
- Screenshot 5 (Bottom Middle):** Shows the 'Advanced' tab for a 'Hadoop' endpoint. It includes 'File Format' settings (Target storage format: 'Text', Use default SerDe: 'V001'), 'File Attributes' (Field Delimiter: '\001', Null value: '\000'), and 'Change Processing' (Escape character: '\000').
- Screenshot 6 (Bottom Right):** Shows the 'Advanced' tab for a 'Hadoop' endpoint. It includes 'File Format' and 'File Attributes' settings, and 'Change Processing' settings (Consider state idle when no changes have been processed for (sec): 3, In idle state, apply/store changes when: 'File size reaches(KB)' (50000) or 'Elapsed time reaches' (50000)).

Attunity Replicate: Task (Full Load)

ATTUNITY REPLICATE

x224n32 | admin (Admin) | 5.5.0.143 | Help |

TESTS **TestTask**

Run Stop Tools Manage Endpoint Connections...

Designer Monitor

Full Load Change Processing

Total Completion: 99% 04:49:52

Tables [Select All]

Completed	Loading	Queued	Error
0	5	6	0

Throughput: 841,463

Throughput Details

rec/sec

Source Target

Messages

Open Clear View Logs

Log Messages Notification

Errors Warnings

Date and Time	Level	Message

Diagram: Oracle (Oracle) connected to a circular icon (Replicate) which is connected to Hadoop (Hadoop).

Attunity Replicate: Task (Incremental Change)

ATTUNITY REPLICATE

x224n32 | admin (Admin) | 5.5.0.143 | Help |

TASKS TestTask1

Run Stop Tools Manage Endpoint Connections...

Full Load Change Processing

Designer Monitor

Incoming Changes
1,152,013 (362,378 Transactions)

Applied Changes
Updates Inserts

Apply Throughput
rec/sec

Apply Latency
Now

Oracle (Oracle) **Hadoop (Hadoop)**

Apply Throughput Details

rec/sec

26,000
24,000
22,000
20,000
18,000
16,000
14,000
12,000
10,000
8,000
6,000
4,000
2,000
0

Time

Source Target

Messages

Open Clear View Logs

Log Messages Errors Warnings Notification

Date and Time Level Message

Agenda

IBM and Oracle Collaboration

IBM Power Systems

Hortonworks Data Platform

Attunity Replicate

Experiences Ingesting Oracle Data into Hadoop on Power Systems

Conclusion and Acknowledgments



Live Demonstration Presented @ IBM Booth, Systems Pedestal

Only Available until Dec 31, 2017!

Hortonworks HDP 3X POWER8 Price-Performance Guarantee

IBM Power Systems guarantees the Power S822LC for Big Data system built with POWER8 delivers at least a 3X price-performance advantage vs. x86 based results when running a customer application/workload with Tez/Hive LLAP on Hortonworks HDP under the conditions noted below. A Worker Node is a server carrying out the HDP query functions, with one Worker Node per server.

3X price-performance means that the customer's documented throughput performance on the cluster of S822LC for Big Data Worker Nodes divided by the price of the cluster of Worker Nodes will be at least 3 times higher than the customer's documented throughput performance on the cluster of x86 based Worker Nodes divided by the price of the cluster of x86 Worker Nodes.

EX: If queries per second on the cluster of S822LC Worker Nodes are 30,000 and 10,000 on the cluster of x86 based Worker Nodes, while the price of the S822LC Worker Node cluster is \$10,000, and the price of the x86 based Worker Node cluster is \$10,000, then the Throughput Performance Per Price would be exactly 3 times higher and the guarantee would be met."

The IBM Power S822LC for Big Data servers (22-core/2.89 GHz) used as Worker Nodes must be purchased from IBM or an authorized IBM Business Partner prior to September 30, 2017. The guarantee period is valid for three (3) months from the date of purchase. The x86-based Worker Nodes must be comparably configured branded servers from Cisco, Dell, HP, or Lenovo and the client is responsible for all Hortonworks licenses.

3X throughput performance per price means that the customer's documented throughput performance on the cluster of Power S822LC for BD Worker Nodes based on either queries, operations or transactions per second divided by the price of the cluster of Worker Nodes will be at least 3 times higher than the customer's same documented throughput performance on the cluster of x86 Worker Nodes divided by the price of said cluster of x86 Worker Nodes.

Remediation: IBM will provide additional performance optimization and tuning services consistent with IBM Best Practices, at no charge. If unable to reach the guaranteed level of price-performance, IBM will provide additional equally configured Worker Nodes to those already purchased to reach the guaranteed level of price-performance.

Notes:

1. Client's Power S822LC for BD Worker Nodes and the x86 Worker Nodes must be running at similar utilization rates of at least 50% or higher, using the same software stack as described in Note #4, and which are configured similarly.
2. Client's Power S822LC for BD performance cannot be constrained by I/O subsystem. Specifically, the I/O subsystem on the Power S822LC for BD Worker Node must achieve greater than or equal I/O bandwidth and operations per second than the x86 Worker Node.
3. Client's Power S822LC for BD Worker Node's physical memory must be the same or greater than the physical memory on the x86 Worker Node.
4. Applicable software stack is Tez/Hive LLAP on HDP 2.6 or later for both the Power S822LC and x86-based Worker Nodes.
5. Client is responsible for demonstrating comparable real-world representative workload between the Power S822LC for BD Worker Node and the x86 Worker Node through the use of the IBM provided tools and comparable tools on x86 systems.
6. 3X guarantee is based on a list price for x86 servers from Dell, Cisco, HP or Lenovo based on E5-2600 v4 or earlier processor technology and the IBM S822LC for Big Data.

Acknowledgments

The authors graciously thank:

Mithun Hr for his full support and execution of the runs that supported this presentation.