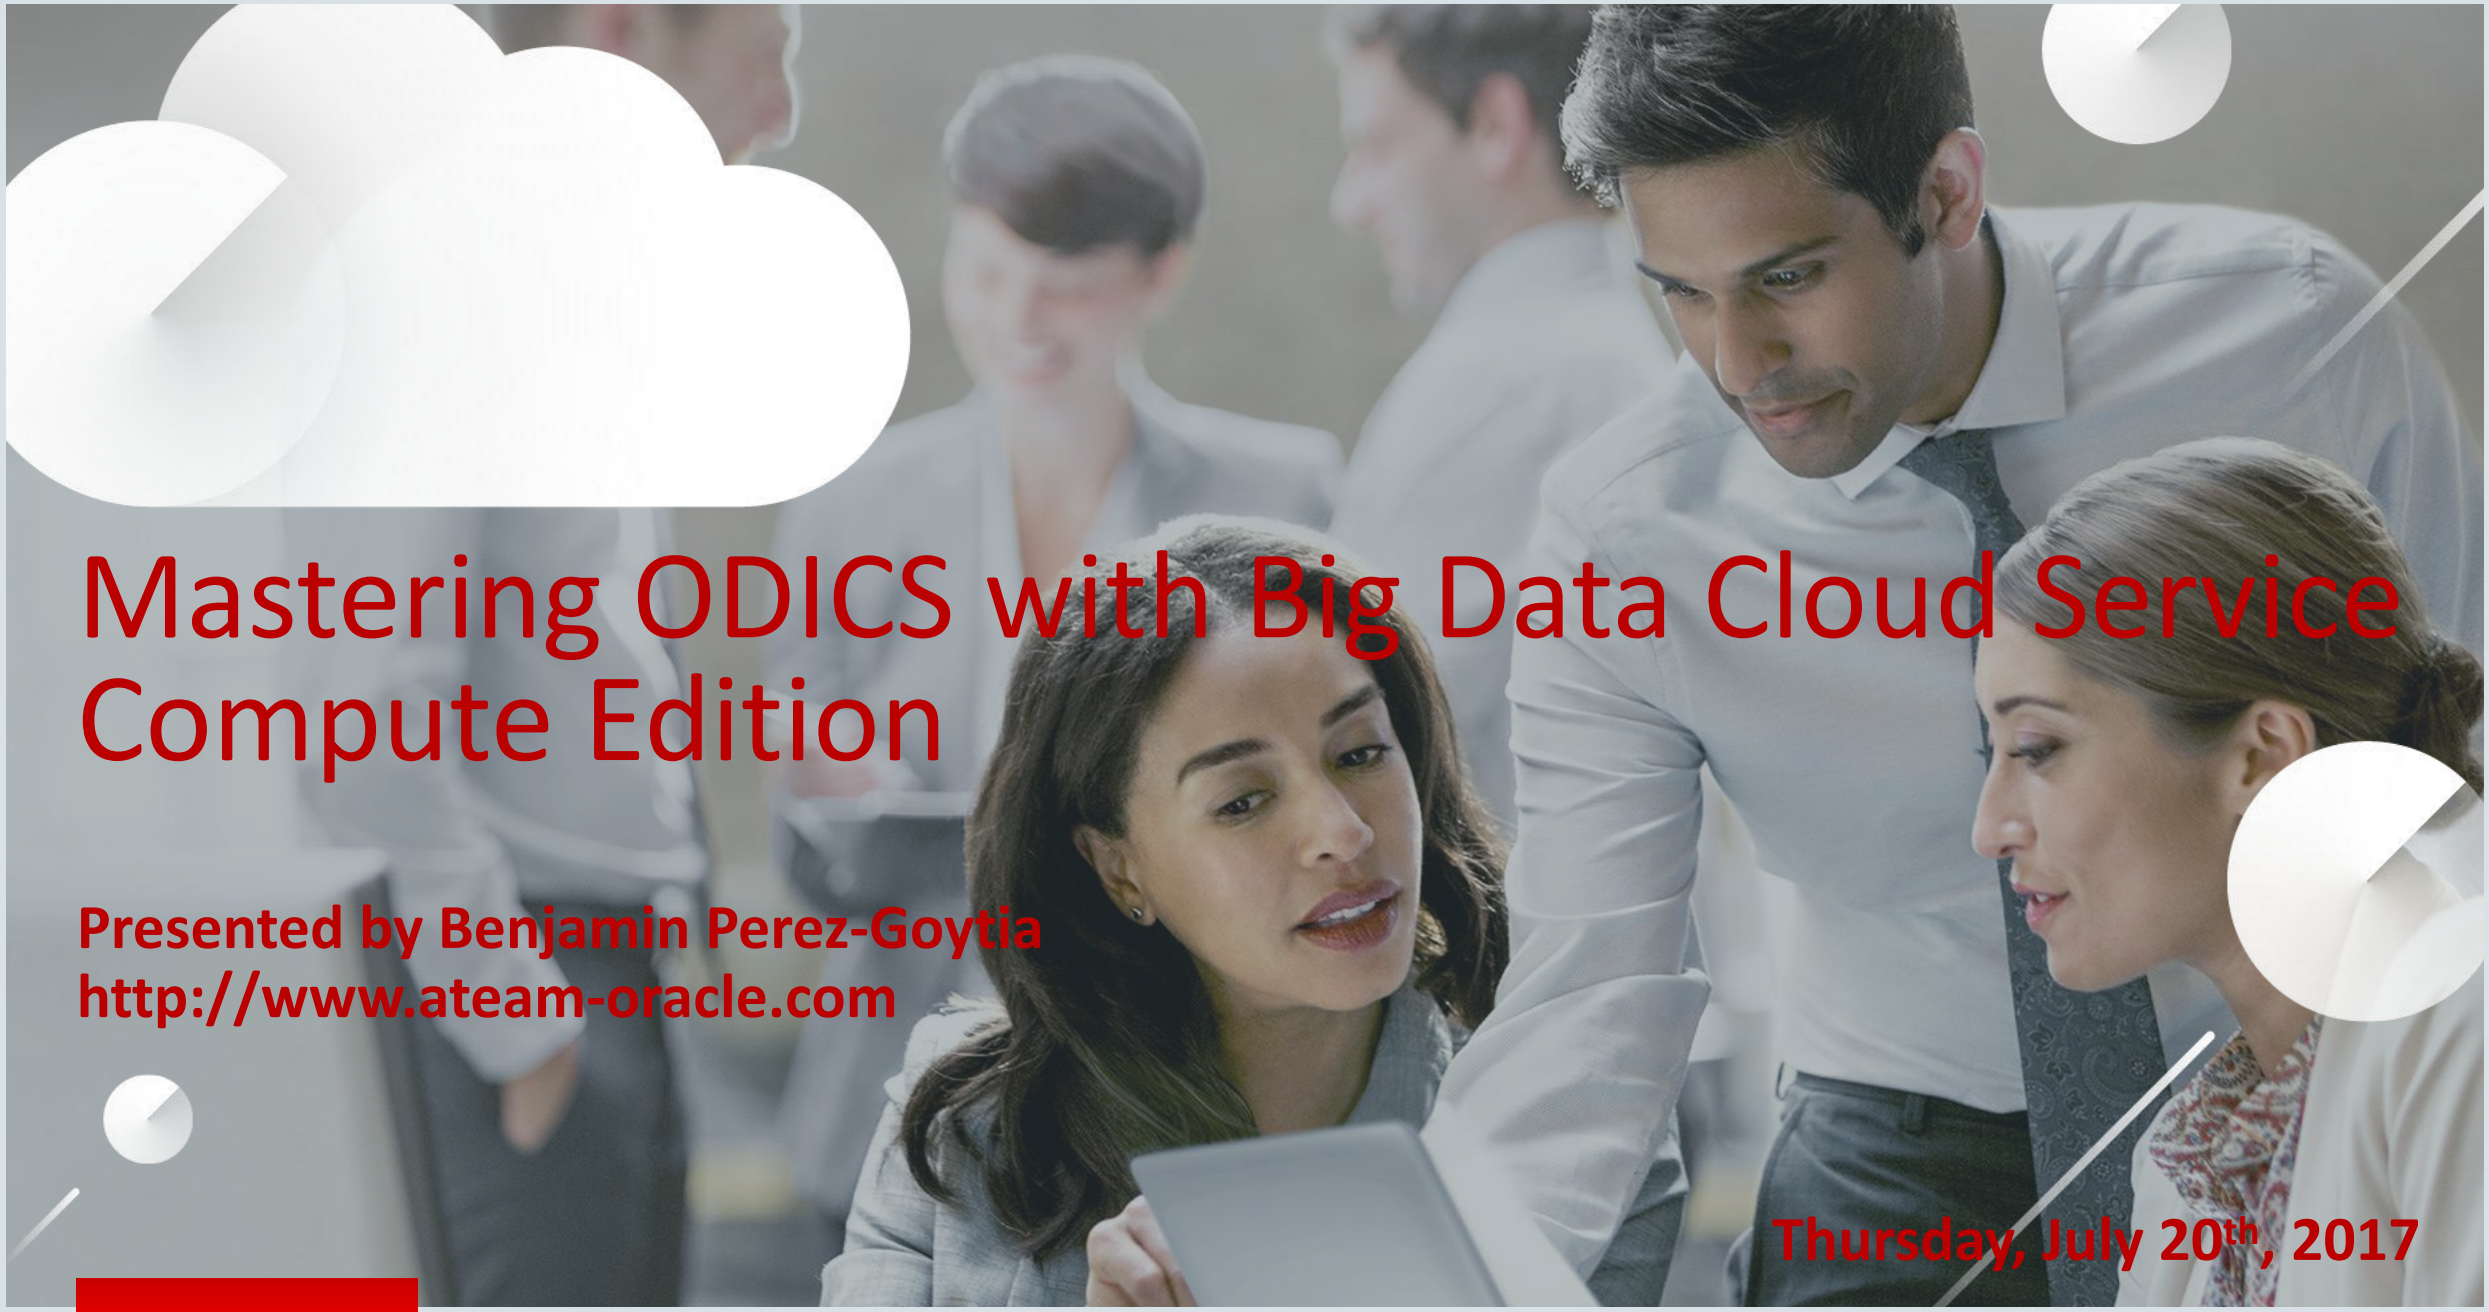

Mastering ODICS with Big Data Cloud Service Compute Edition

Presented by Benjamin Perez-Goytia
http://www.ateam-oracle.com

Thursday, July 20th, 2017

Program Agenda

1. Oracle Data Integrator Cloud Service (ODICS)
2. Big Data Cloud Service Compute Edition (BDCS-CE)
3. ODICS and BDCS-CE Architecture
4. Use Cases
5. Demo

Program Agenda

1. Oracle Data Integrator Cloud Service (ODICS)
2. Big Data Cloud Service Compute Edition (BDCS-CE)
3. ODICS and BDCS-CE Architecture
4. Use Cases
5. Demo

Oracle Data Integrator Cloud Service

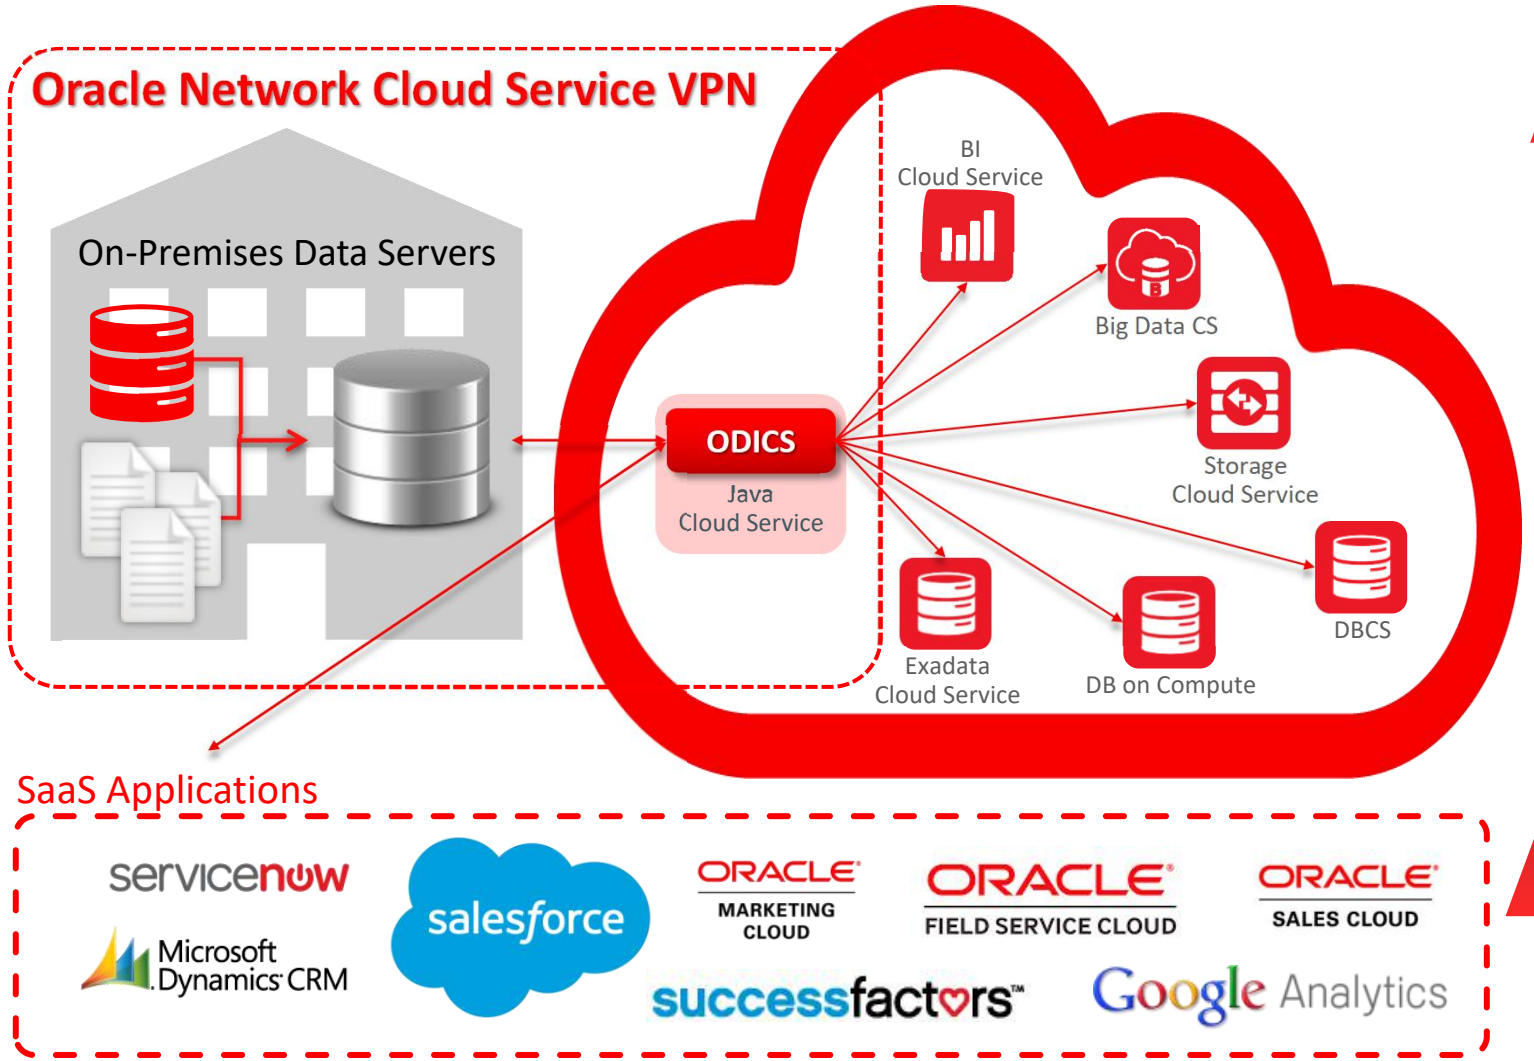

On-Premises to Oracle Cloud

- ODICS can integrate with on-premises resources.
- With Oracle VPN, Oracle Cloud becomes an extension of the customer network.
- Without VPN, on-premises data can be transferred to the Cloud with ODI technologies.

Oracle Public Cloud

- ODICS runs on Java Cloud Service for high availability.
- ODICS natively integrates with other Oracle pubic cloud services.

Software As a Service (SaaS)

- ODICS can integrate with Oracle SaaS and non-Oracle SaaS applications via RESTful services or third party adapters.

Program Agenda

1. Oracle Data Integrator Cloud Service (ODICS)
2. Big Data Cloud Service Compute Edition (BDCS-CE)
3. ODICS and BDCS-CE Architecture
4. Use Cases
5. Demo

Oracle Data Integrator Cloud Service

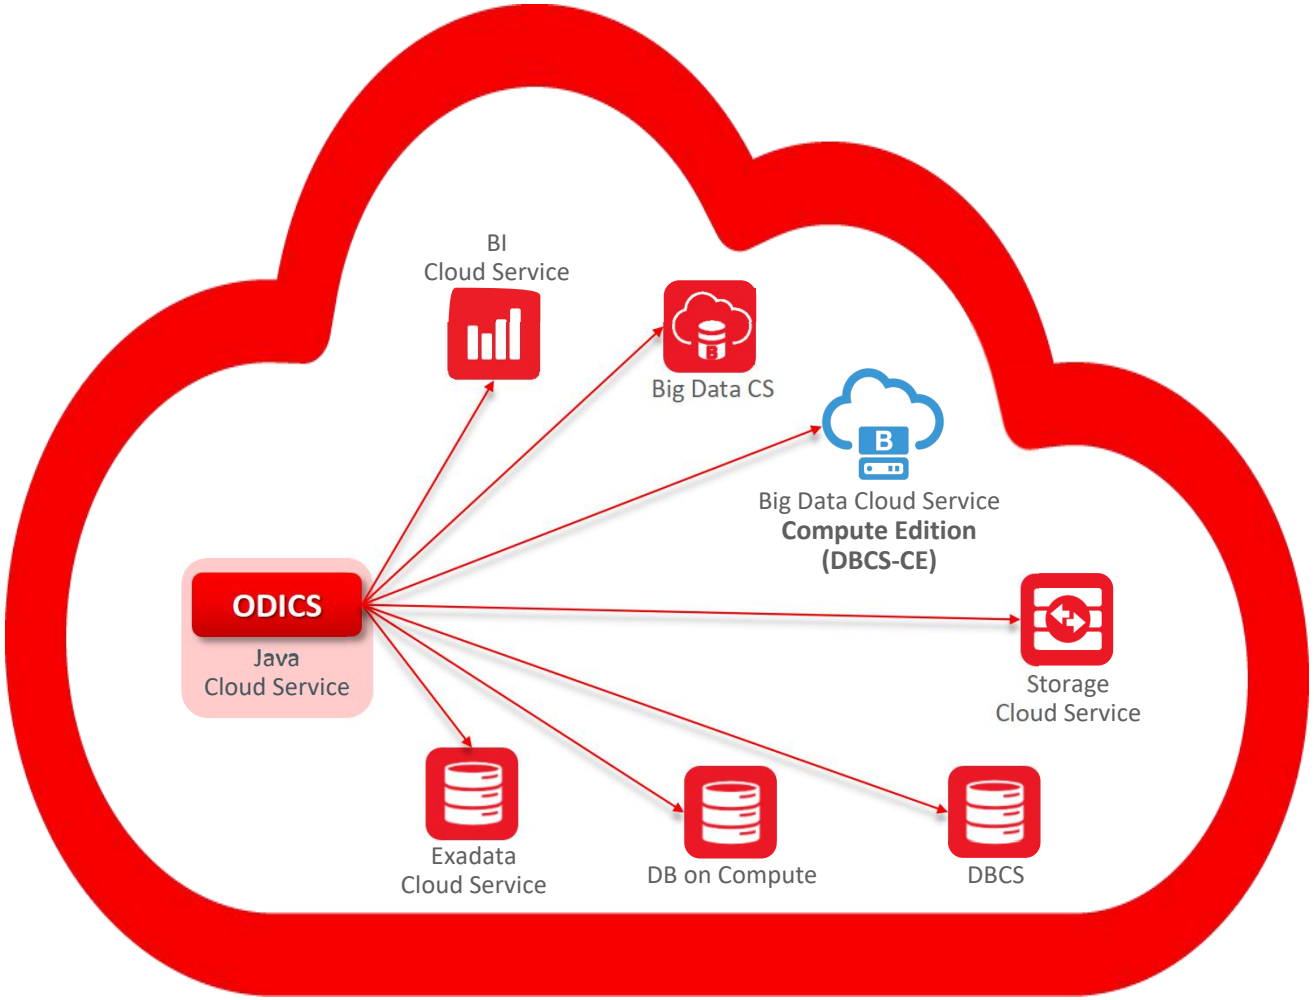

Big Data Cloud Service (BDCS)

- Provisioned and fully configured Hadoop and Spark clusters on demand.
- Dedicated Compute Shapes with Attached Storage.
- Pre-configured Software such as Big Data Connectors, ODI, Big Data Spatial, and Oracle R.
- Big Data SQL integration.

Big Data Cloud Service - Compute Edition (BDCS-CE)

- Hadoop and Spark delivered as a managed, elastic, integrated platform.
- Independent Compute and Storage.
- REST APIs to access all functionality.
- Integration with other Oracle cloud services via association.

Big Data Cloud Service – Overview

Key Features

- Hadoop, Spark, Kafka, Hive, Pig delivered.
  - Cloudera Hadoop Distribution
  - Oracle Big Data Connectors + Oracle R Distribution
  - Oracle Big Data Spatial and Graph
  - Oracle Data Integrator Enterprise Edition
- Dedicated Compute Shapes with Attached Storage
  - Start as small as 3 nodes and grow seamlessly
  - Burst/Shrink Compute when Required
  - Embedded Edge Nodes
  - Full Security
- Platform to Integrate Big Data with Oracle Database
  - Big Data SQL Cloud Service as an Add-On Service

Big Data Cloud Service – Compute Edition Overview

Big Data Cloud Service Compute Edition (BDCS-CE)

Key Features

- Apache Hadoop, Spark, Hive, Pig.
  - Oracle Hadoop Distribution, similar to Hortonworks.
  - Kafka is a separate cloud service: Event Hub Cloud Service.
- Independent Compute and Storage
  - Independently scale compute or storage.
  - Use Storage Cloud Service and Object Store as the data lake.
  - Use Alluxio as the in-memory caching layer for fast data access.
  - Scale up and down BDCS-CE clusters.
- REST APIs
  - REST based API access to all functionality.
- Cloud Service Associations
  - Automatic integration with Oracle, Event Hub, and MySQL cloud services through association.

ORACLE

Big Data Cloud Service – Compute Edition

Big Data Cloud Service Compute Edition (BDCS-CE)

Use Cases

- Streaming Data Analysis
  - High Performance connections with Event Hub Cloud Service.
  - In memory caching layer for fast access to Storage Cloud Service.
  - Spark Streaming.
- Interactive Data Analysis
  - Apache Zeppelin-based notebooks.
  - Ability to import/export notes.
  - Support for Scala, Hive, Spark SQL, Python, and R.
- Batch Jobs – ETL, ELT, Cleansing
  - Job scheduling
  - REST APIs to launch jobs.
  - Client-side CLI

Configuring ODICS for Big Data CS – Compute Edition

Big Data Cloud Service Compute Edition (BDCS-CE)

Configuration Steps:

- Provision ODICS
  - Create an instance of JCS.
  - Choose Software Release and Software Edition.
  - Add the Load Balancer option, and provision JCS.
  - Install ODI and create an ODI repository on DBCS.
  - Create at least two ODI Weblogic Managed Servers for failover.
- Provision BDCS-CE
  - Create an instance of BDCS-CE.
  - Install at least two ODI standalone agents.
  - Configure the agent to access Hadoop paths, libraries and jars.
  - Install additional software (i.e. Sqoop, etc).
  - Configure the ODI Topology for Hadoop, Hive, Spark, and Pig.
- Access Configuration
  - On DBCS, enable access rules to allow both JCS and DBCS-CE to access the ODI repository.
  - On JCS, enable access rules to access BDCS-CE.
  - On BDCS-CE, enable access rules to access the ODI repository on DBCS.

Program Agenda

1. Oracle Data Integrator Cloud Service (ODICS)
2. Big Data Cloud Service Compute Edition (BDCS-CE)
3. ODICS and BDCS-CE Architecture
4. Use Cases
5. Demo

ODICS and BDCS-CE Architecture

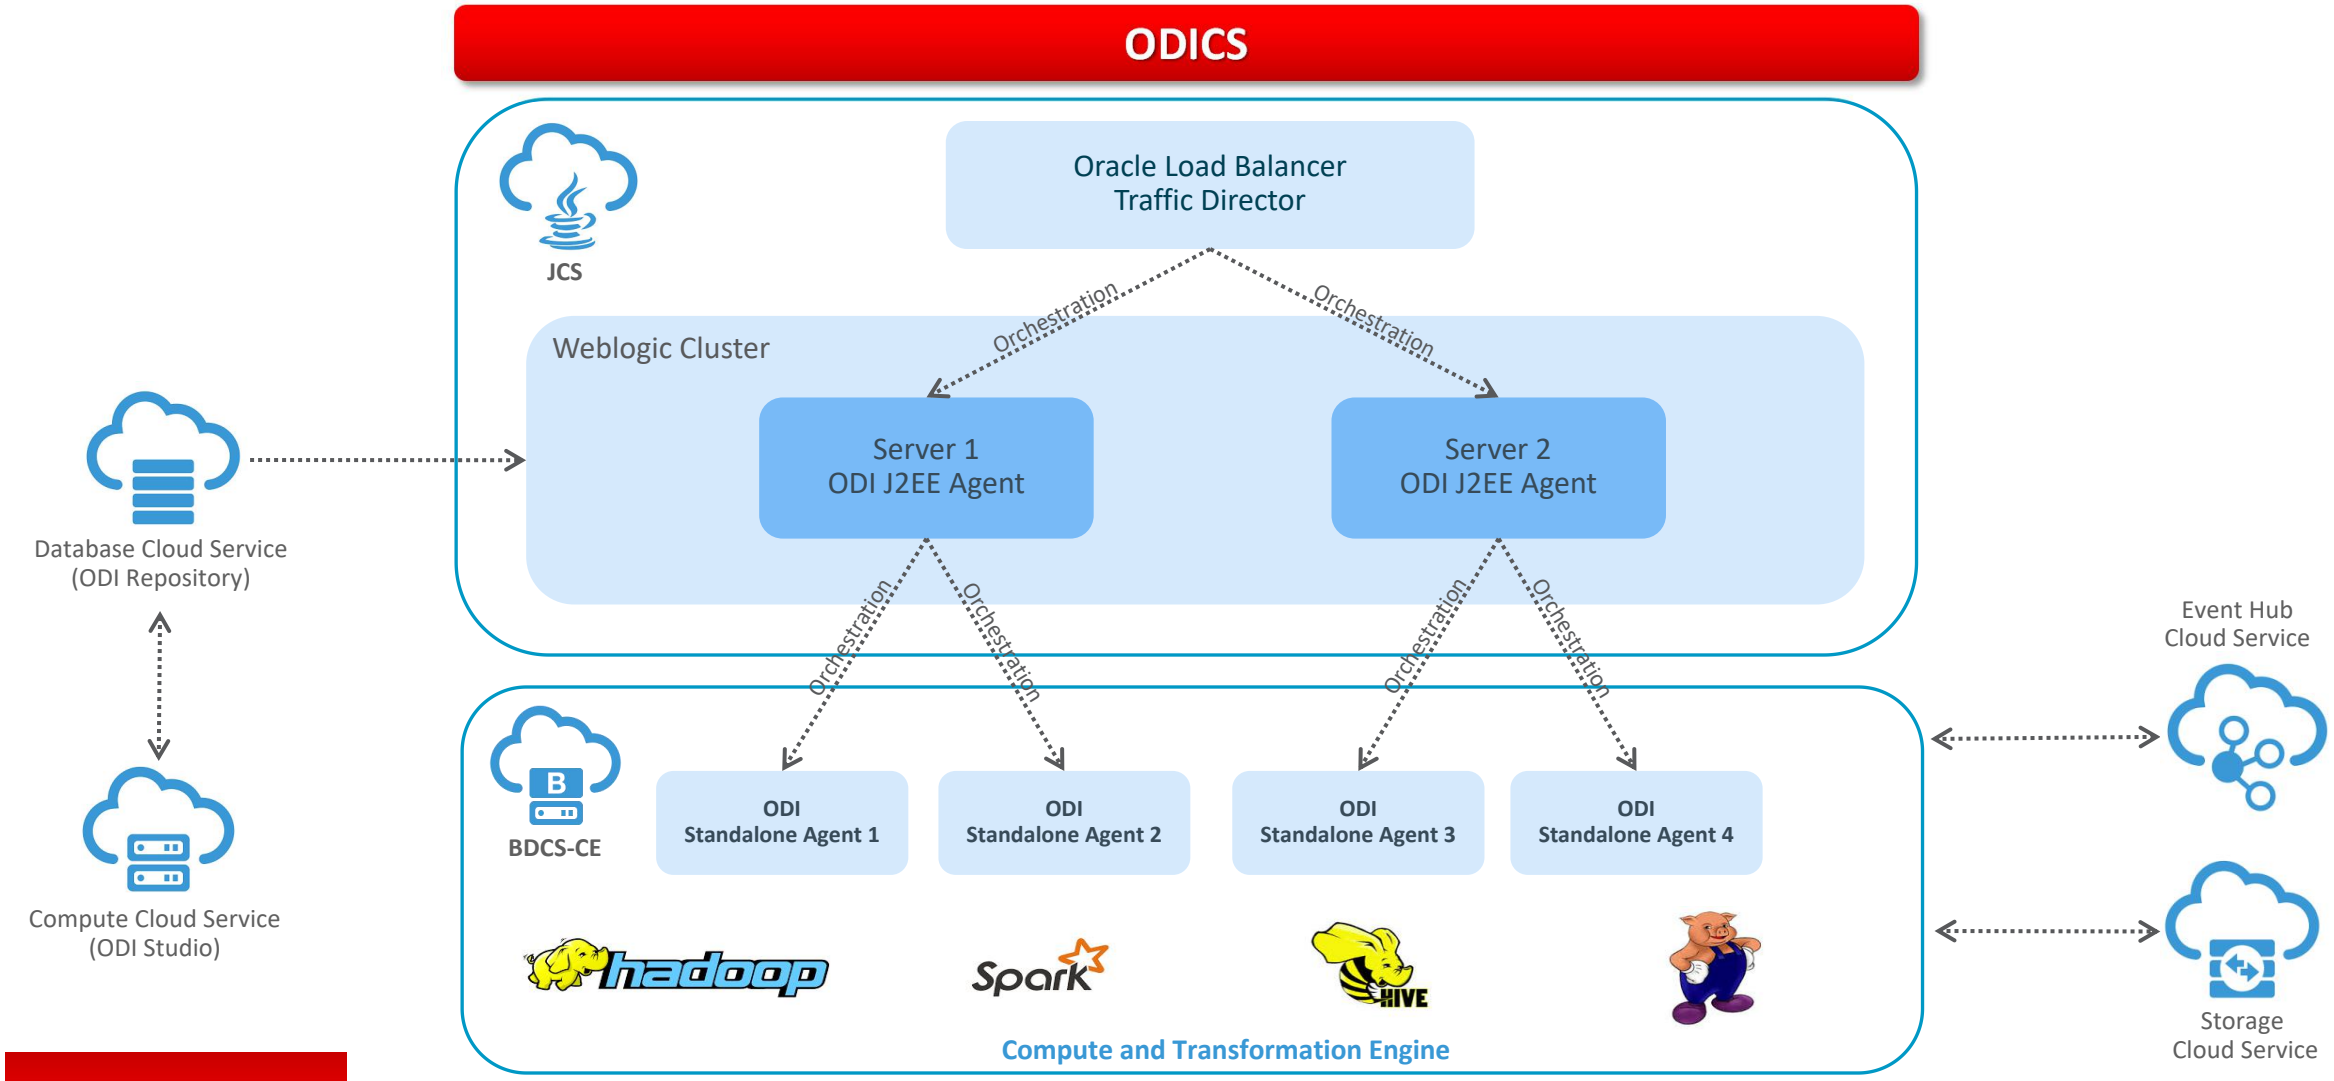

Program Agenda

1. Oracle Data Integrator Cloud Service (ODICS)
2. Big Data Cloud Service Compute Edition (BDCS-CE)
3. ODICS and BDCS-CE Architecture
4. Use Cases
5. Demo

Use Case: Transform SQL Data into BDCS-CE and load it into BICS

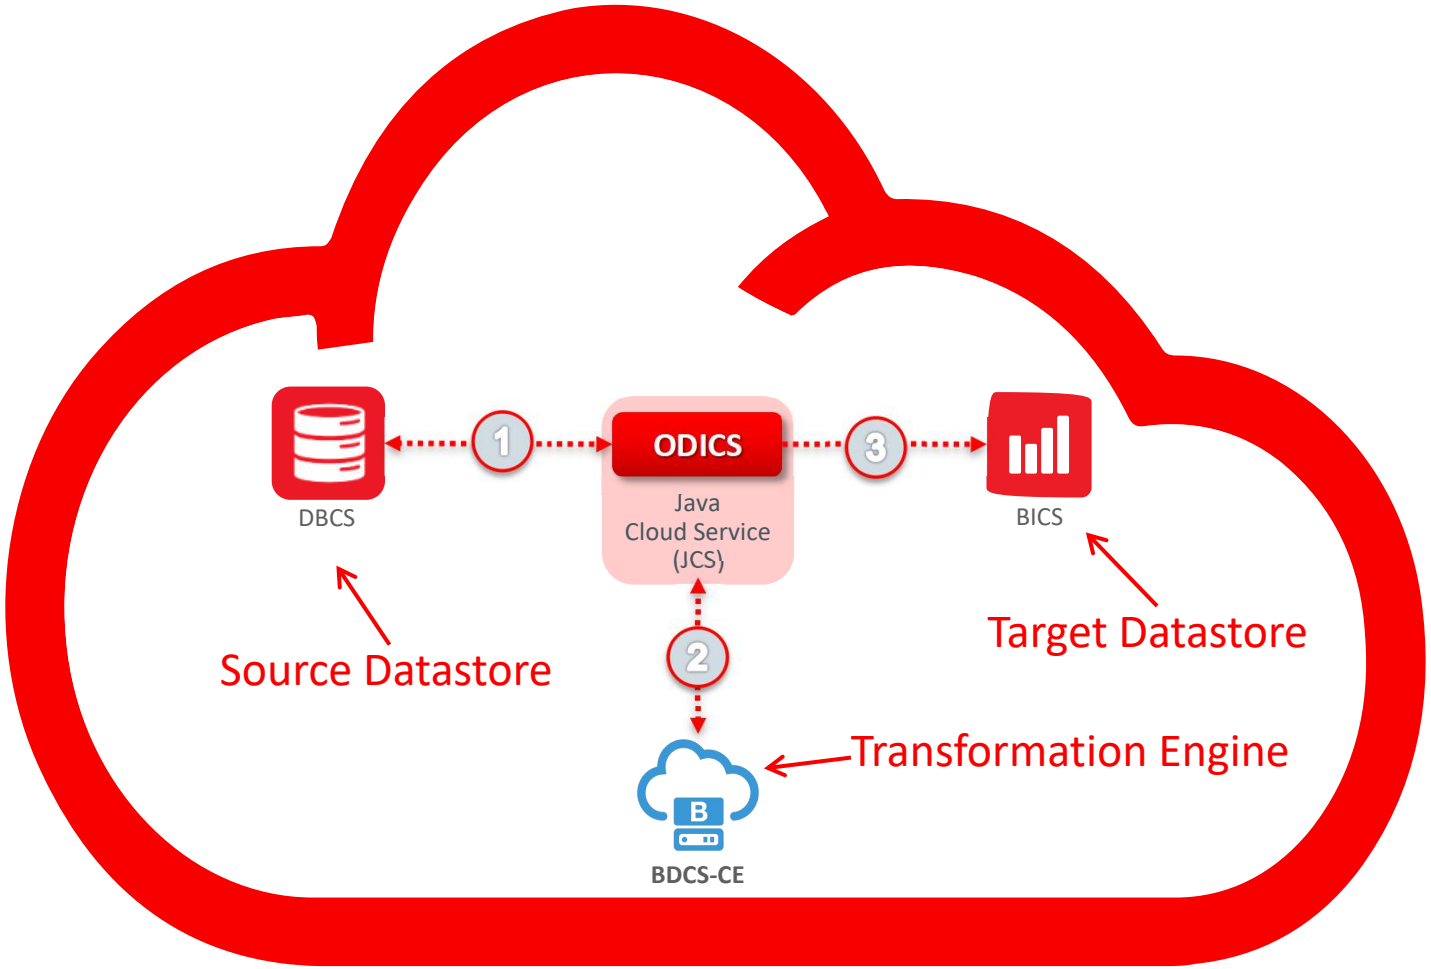

1. ODICS extracts data from DBCS.
2. Data is transformed on BDCS-CE.
3. Transformed Data is loaded from BDCS-CE into BICS.

Use Case: Transform SQL data in BDCS-CE and load it into BICS

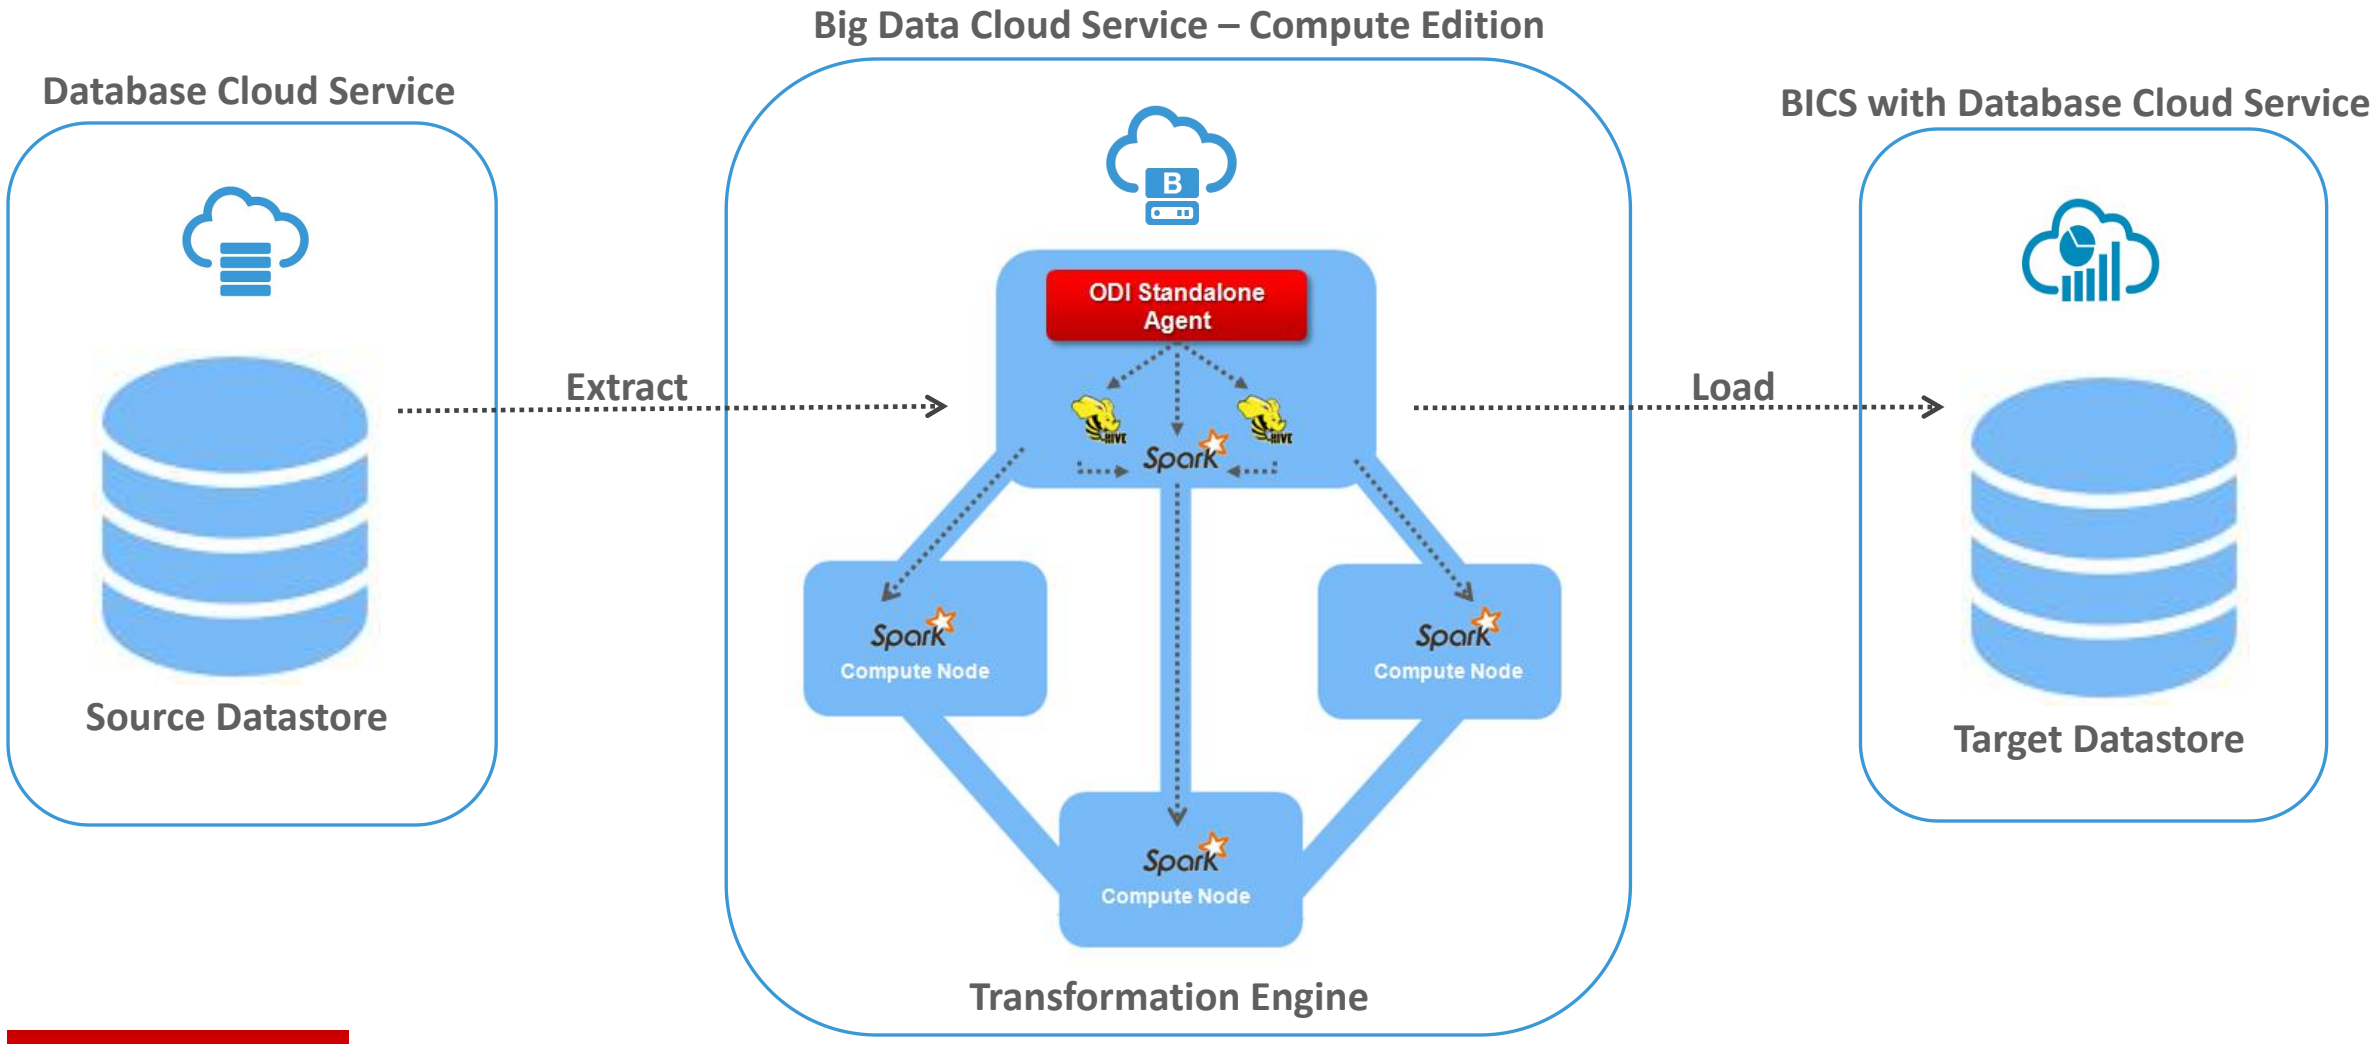

Program Agenda

1. Oracle Data Integrator Cloud Service (ODICS)
2. Big Data Cloud Service Compute Edition (BDCS-CE)
3. ODICS and BDCS-CE Architecture
4. Use Cases
5. Demo

Find More Information

Oracle Big Data Cloud Services

Oracle Big Data Cloud Service Compute Edition	**http://cloud.oracle.com/hadoop**
Oracle Data Integrator Cloud Service	**https://cloud.oracle.com/data-integrator**


ORACLE®