

Microsoft technológia pre BigData

L'ubomír Goryl

Solution Professional





Agenda

Súčasný stav

Hadoop = HDInsight

Moderný DWH = APS

Analytika = AzureML, PowerBI

Scenáre pre BigData

Tradiční přístup

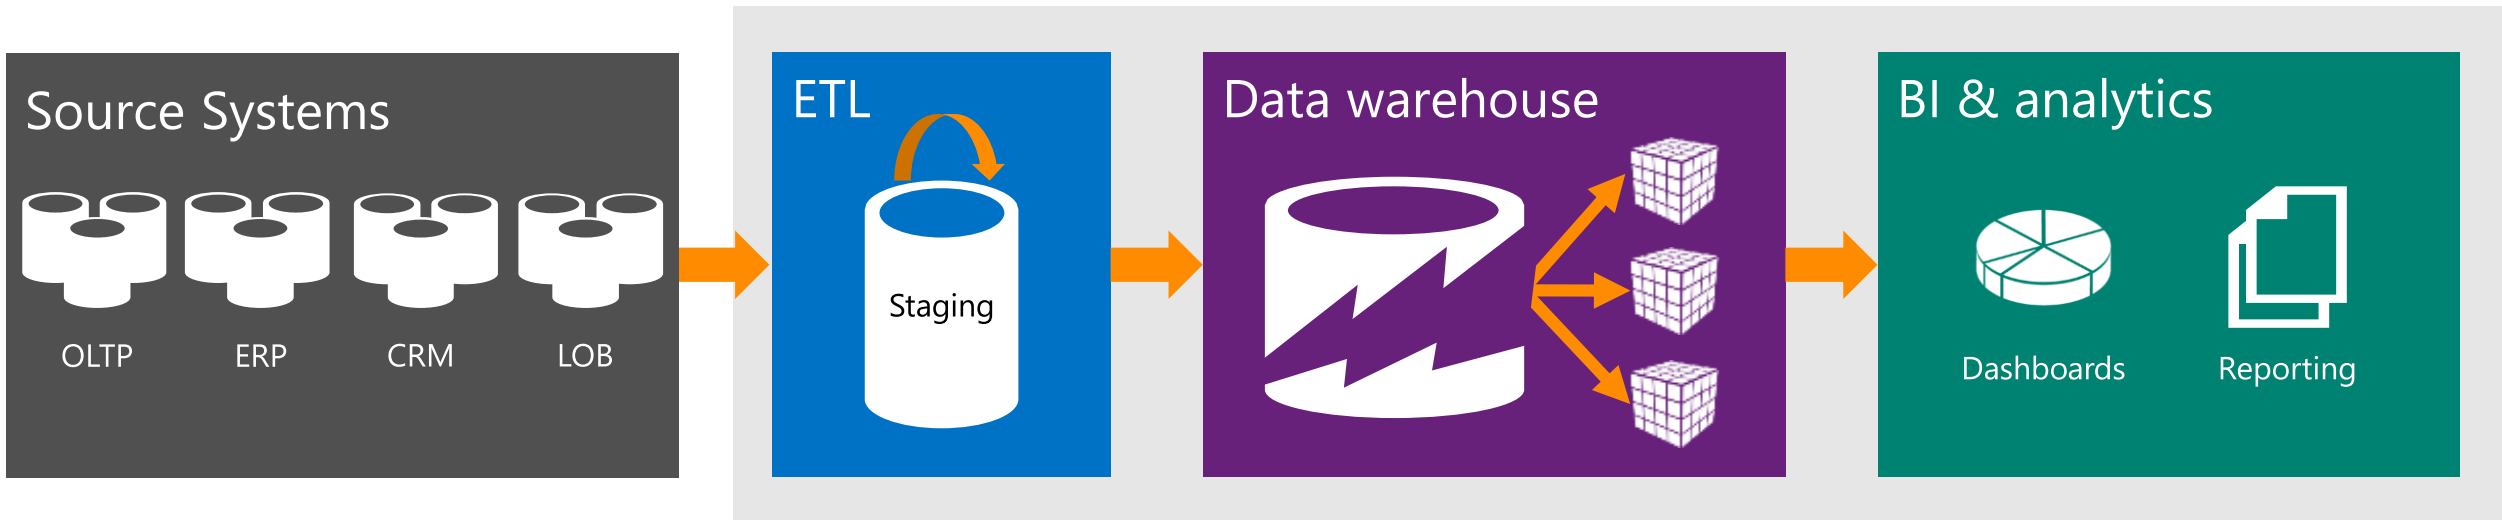
①
Increasing data volumes

1110 1110
1010 1010
1010 1010 1010

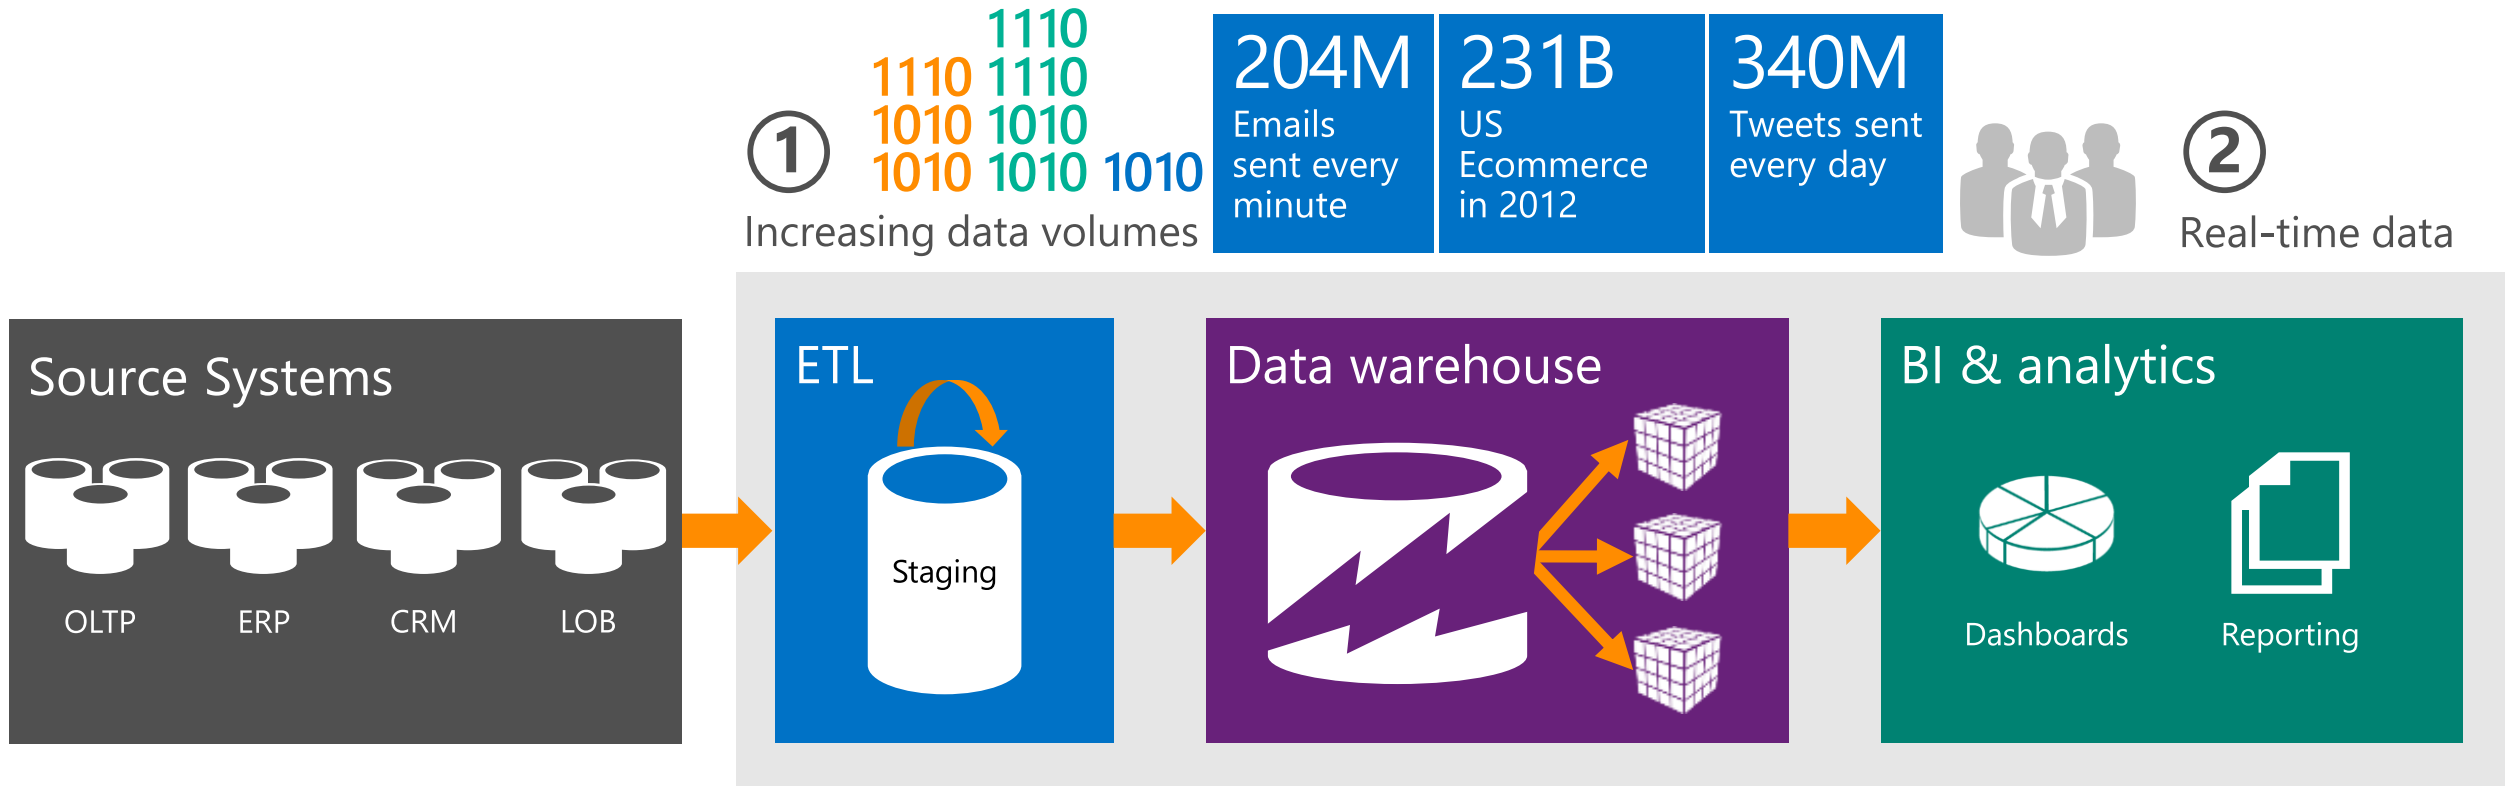
50x
Data
growth
2010-2020

1Trillion
Web pages

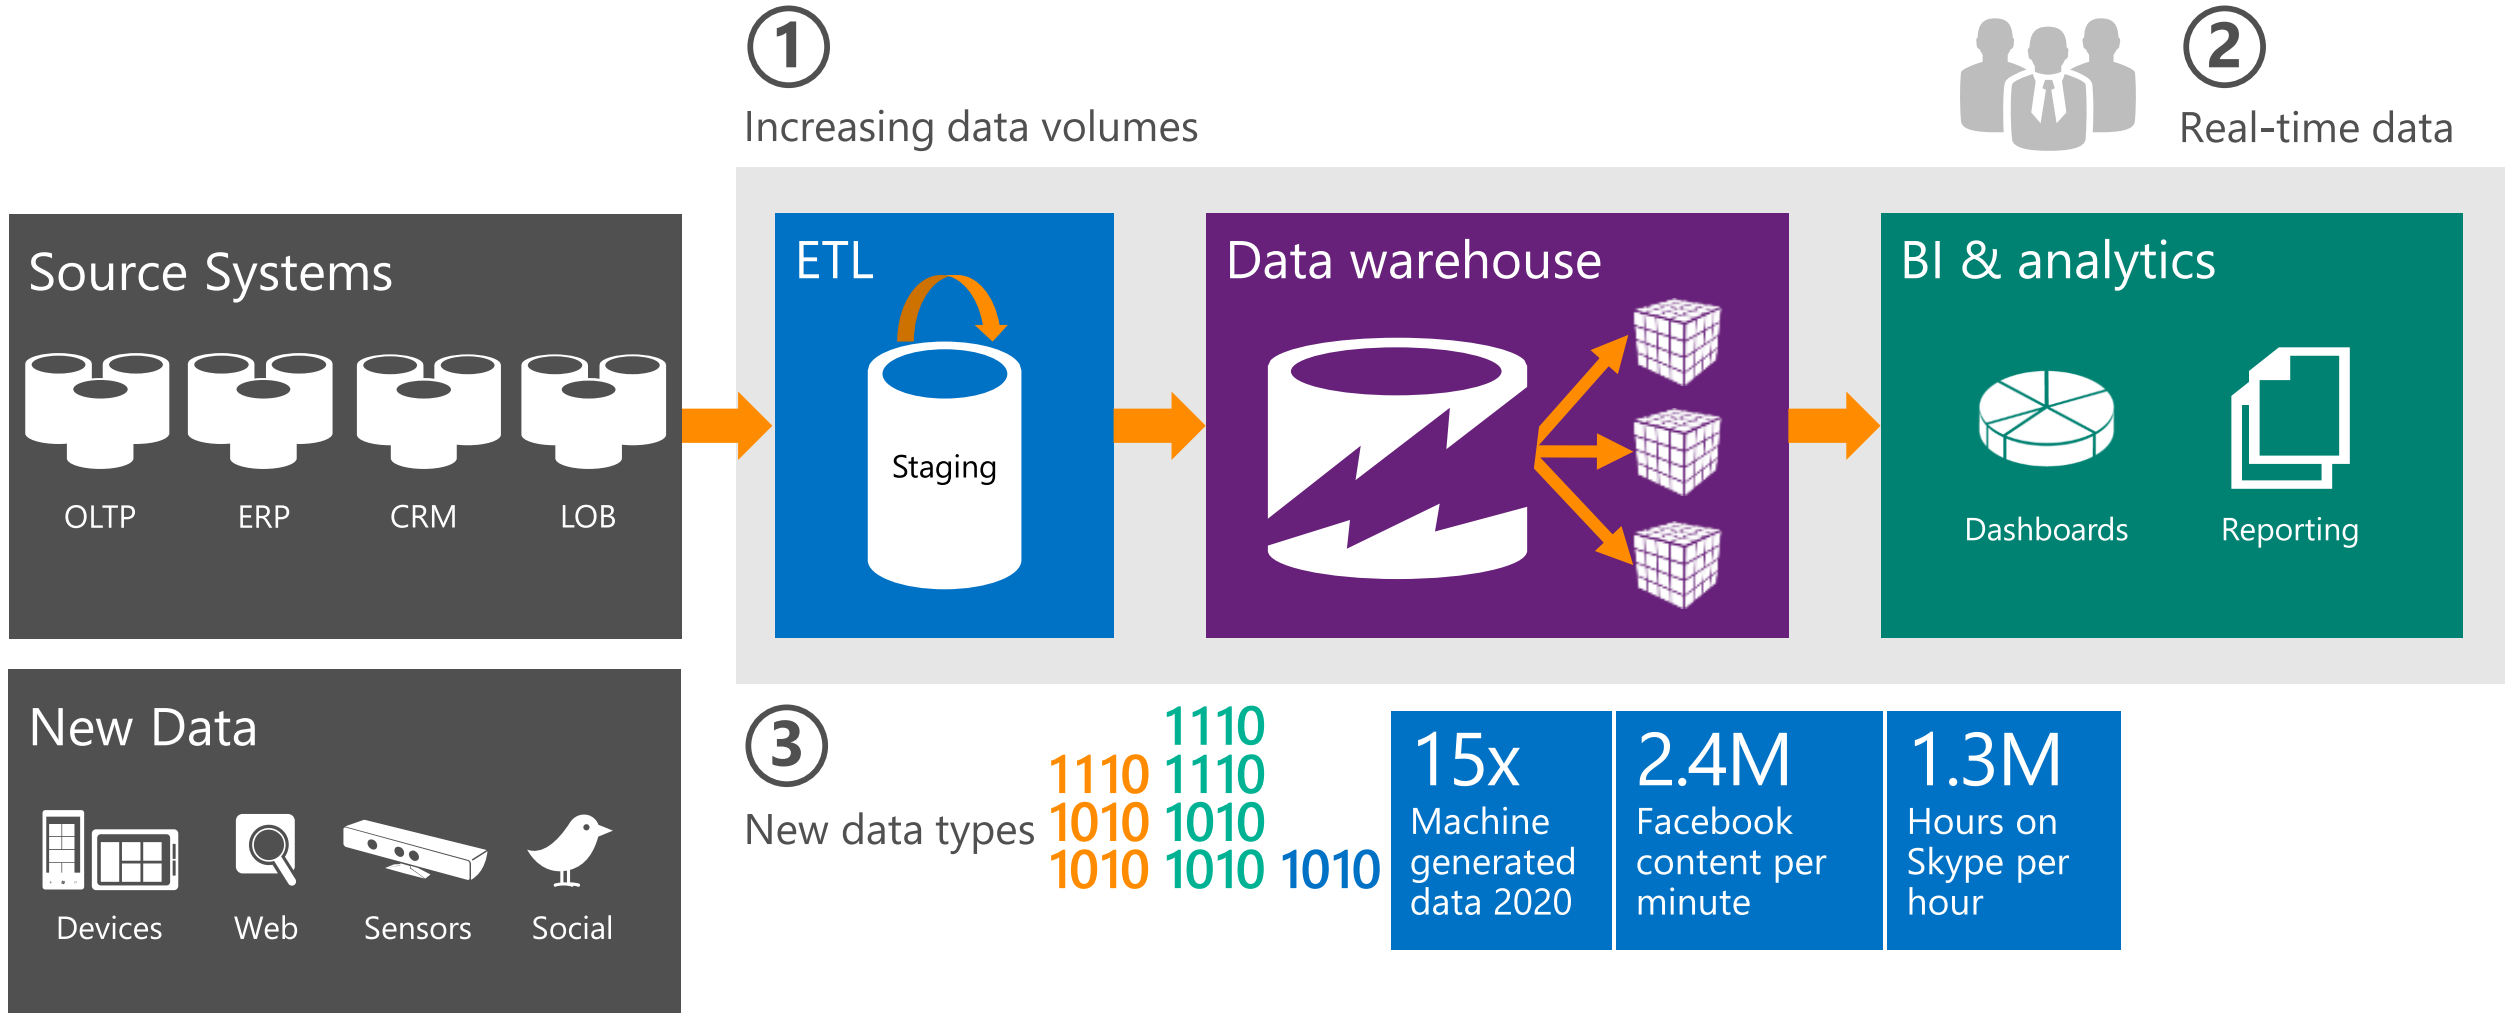
40ZB
Digital
Universe
2020



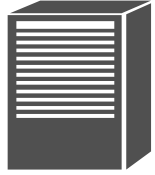
Breaking points of traditional approach



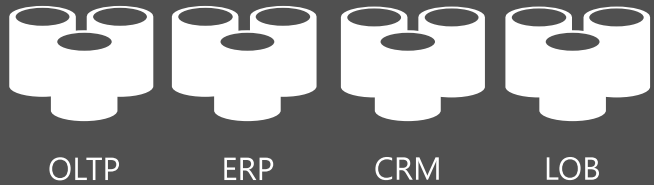
Breaking points of traditional approach



Breaking points of traditional approach



Source Systems



New Data



①

Increasing data volumes



②

Real-time data

ETL

Staging

Data warehouse

BI & analytics

Dashboards

Reporting

③

New data types

④

Cloud-born data

1110 1110
1010 1010
1010 1010 1010

\$100B
spend on
cloud

40%
CRM sold
are SaaS

50%
large orgs
have hybrid
by 2017

HDInsight v Azure (Haddop)

Microsoft Azure | ▼ CREDIT STATUS Subscriptions 🌐 lubo.goryl@outlook.com 👤

hdinsight

You have no HDInsight clusters. Create one to get started!

[CREATE AN HDINSIGHT CLUSTER](#) ➔

NEW

- COMPUTE
- DATA SERVICES**
- APP SERVICES
- NETWORK SERVICES
- STORE PREVIEW

- SQL DATABASE
- STORAGE
- HDINSIGHT**
- RECOVERY SERVICES
- MACHINE LEARNING PREVIEW

- HADOOP**
- HBASE
- CUSTOM CREATE

Quickly create an HDInsight cluster

HDInsight – krok 1/3

NEW HDINSIGHT CLUSTER

Cluster Details

CLUSTER NAME

lghdinsight01

✓

*.azurehdinsight.net

CLUSTER TYPE

Hadoop

▼

HDINSIGHT VERSION ?

default (3.1)

3.1 (HDP 2.1, Hadoop 2.4)

3.0 (HDP 2.0, Hadoop 2.2)

→

234

HDInsight – krok 2/3

NEW HDINSIGHT CLUSTER

Configure Cluster

DATA NODES ?

4

The cluster size affects the cluster price. [Pricing details](#)

REGION/VIRTUAL NETWORK

North Europe

1

←

→

3

4

HDInsight – krok 3/3

123

NEW HDINSIGHT CLUSTER

Storage Account

STORAGE ACCOUNT

Create New Storage

ACCOUNT NAME

lghadoop

DEFAULT CONTAINER

lghadoop

ADDITIONAL STORAGE ACCOUNTS

0

←

✓

15' Hadoop cluster - running

The screenshot displays the Microsoft Azure portal interface. The top navigation bar includes the 'Microsoft Azure' logo, a 'CREDIT STATUS' button, and user information for 'lubo.goryl@outlook.com'. The left sidebar contains a list of service categories: ALL ITEMS, WEBSITES (0), VIRTUAL MACHINES (1), MOBILE SERVICES (0), CLOUD SERVICES (1), SQL DATABASES (0), STORAGE (1), and HDINSIGHT (1). The main content area is titled 'hdinsight' and features a table with the following data:

NAME	STATUS	CLUSTER TYPE	SUBSCRIPTION NAME	LOCATION	VERSION
lghdinsight01	→ ✓ Running	Hadoop	Windows Azure MSDN - Vi...	North Europe	3.1

Below the table, a dark grey banner displays a success message: 'Creation of cluster lghdinsight01 is complete'. To the right of the message is an 'OK' button with a checkmark icon. A list of steps with green checkmarks indicates the completion of the following tasks:

- ✓ Creating Storage
- ✓ Submitting
- ✓ Accepted
- ✓ Cluster Storage Provisioned
- ✓ Windows Azure VM Configuration
- ✓ Configuring HDInsight
- ✓ Running

The bottom of the interface features a dark grey bar with a '+ NEW' button, a 'DELETE' button (trash icon), a 'MANAGE CLUSTER' button (tag icon), and a cluster count of '1' with a help icon.

Hadoop is a platform with portfolio of projects

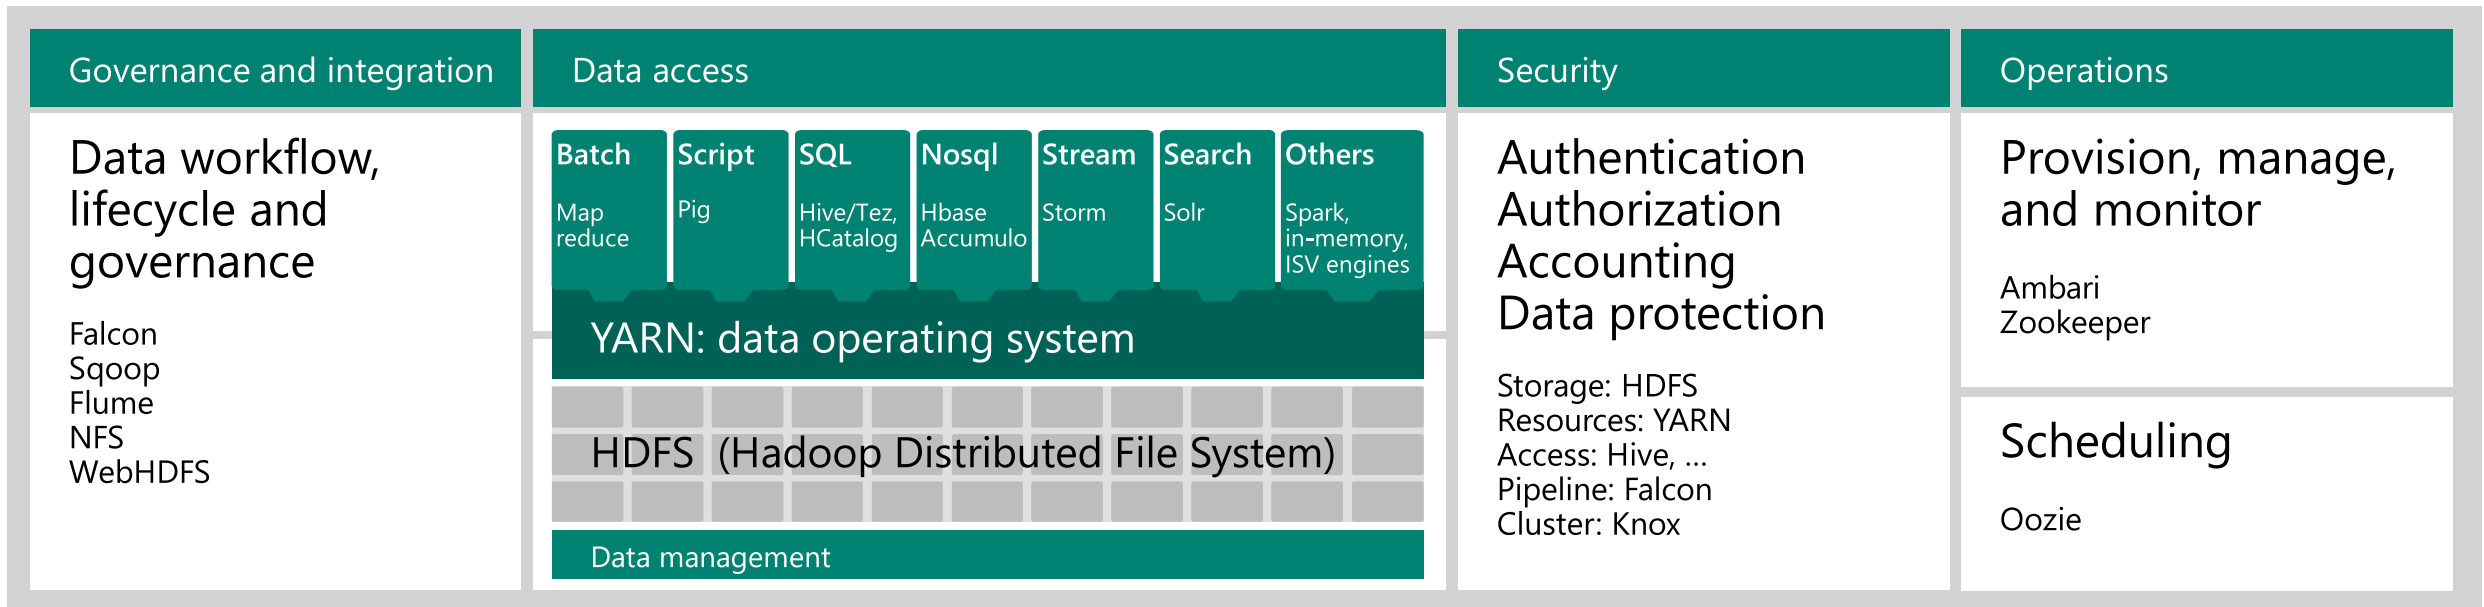
Governed by Apache Software Foundation (ASF)

Comprises core services of MapReduce, HDFS, and YARN

In addition to the core, includes functions across:


Data services which allow you to manipulate and move data (Hive, HBase, Pig, Flume, Sqoop)

Operational services which help manage the cluster (Ambari, Falcon, and Oozie)



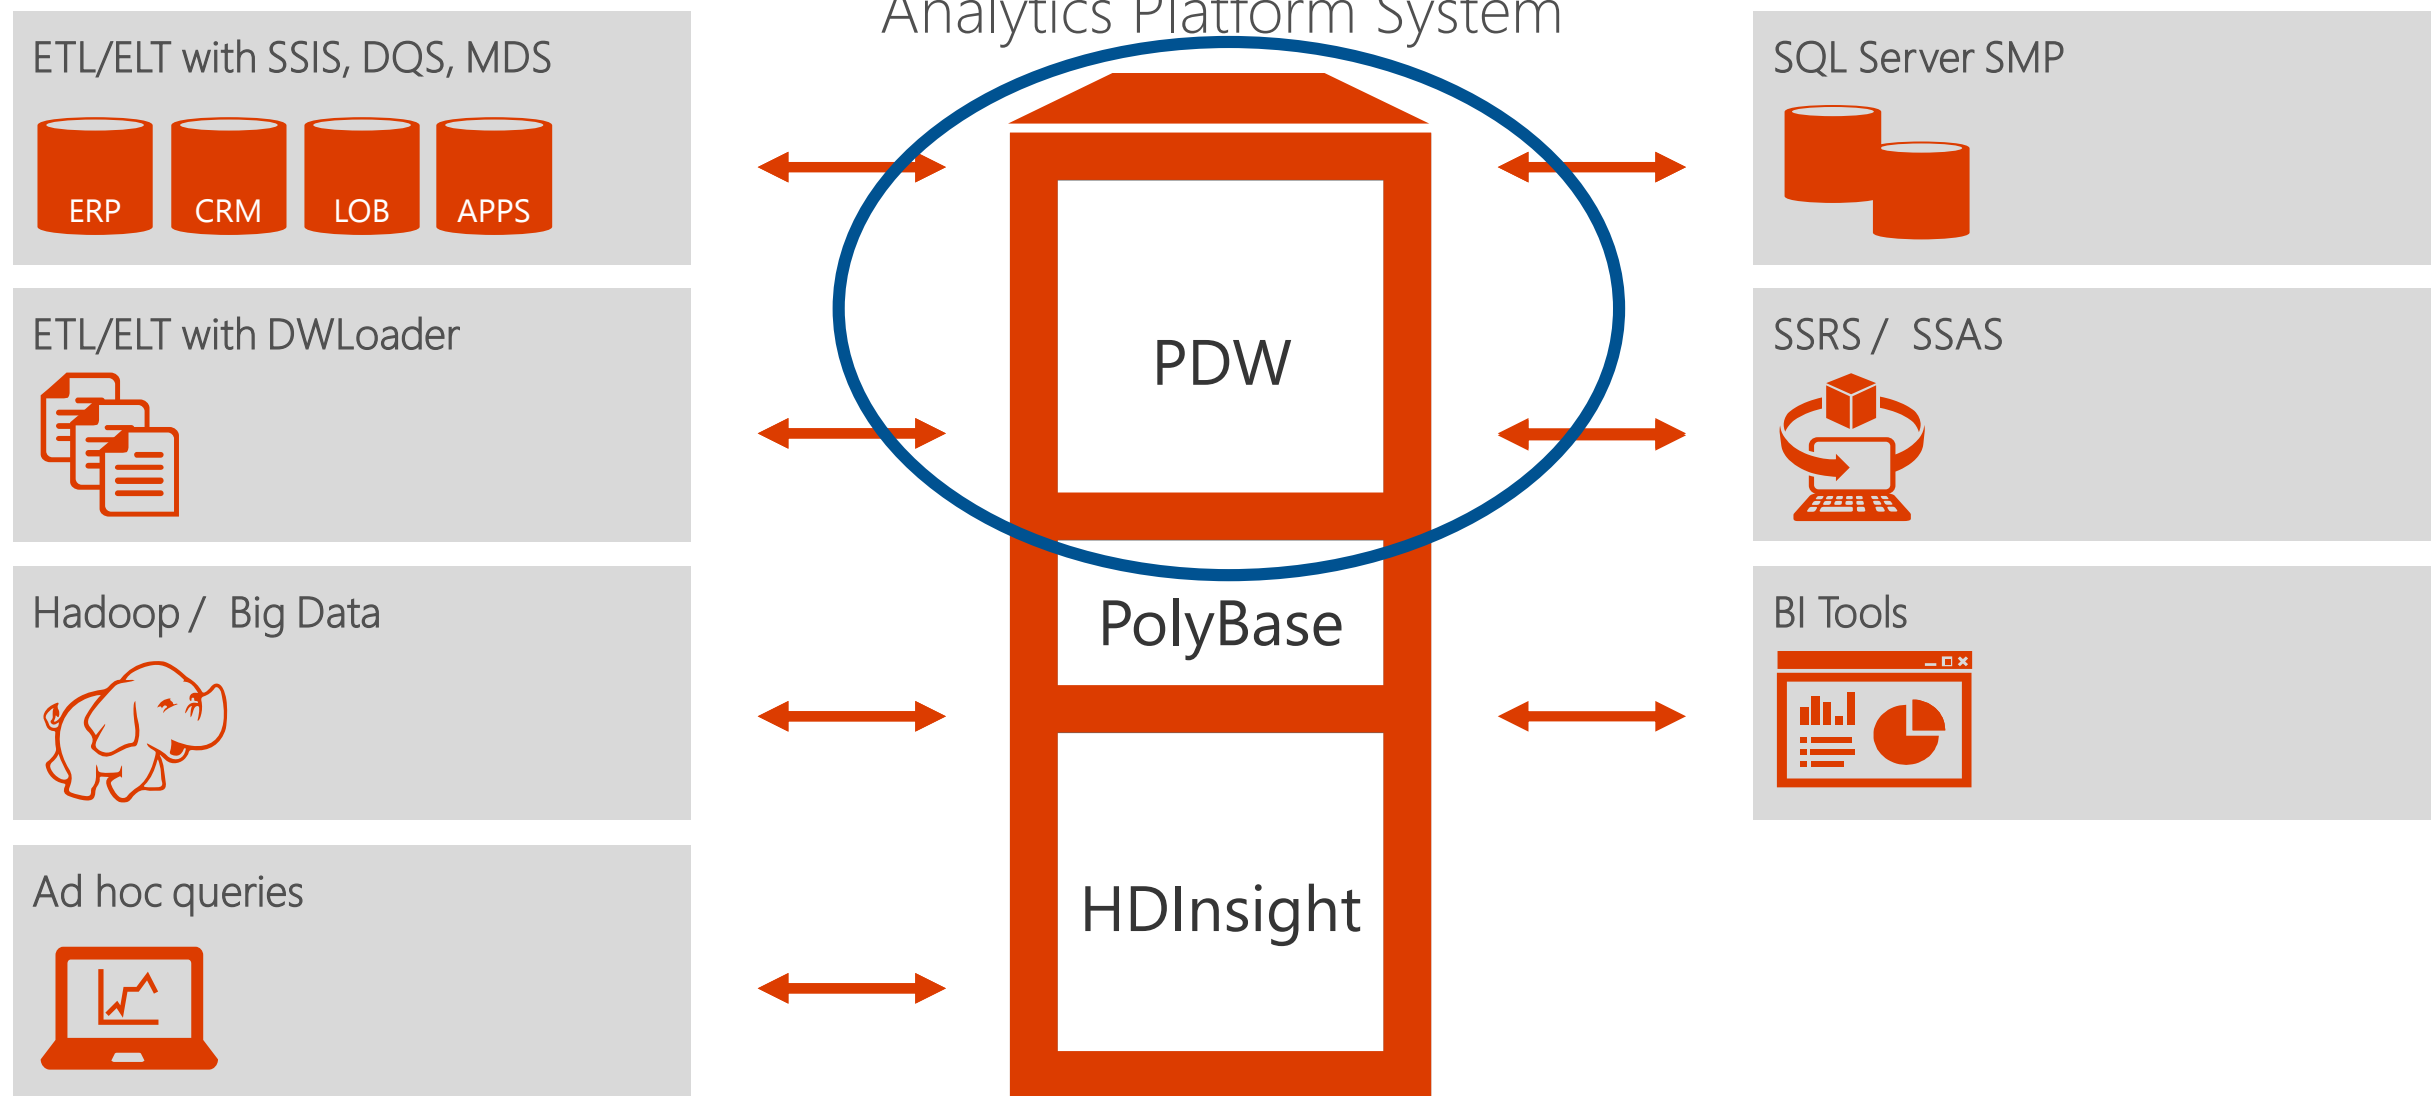
A Hadoop distribution is a package of projects

Tested for consistency across entire package

HDP 2.1 April 2014	2.4.0	0.4.0	0.12.1	0.13.0	0.98.0	4.0.0	1.5.1	0.9.1	0.9.0	4.7.2	0.5.0	1.4.4	1.4.0	1.5.1	4.0.0	3.4.5	.0.4.0
HDP 2.0 October 2013	2.2.0		0.12.0	0.12.0	0.96.1				0.8.0			1.4.4	1.3.0	1.4.4	3.3.2	3.4.5	.0.4.0
HDP 1.3 May 2013	1.1.2		011.0	0.11.0	0.94.6				0.7.0			1.4.3	1.3.1	1.2.5	3.3.2	3.4.5	.0.4.0
	 Hadoop and YARN	Tez	Pig	Hive and HCatalog	HBase	Phoenix	Accumulo	Storm	Mahout	Solr	Falcon	Sqoop	Flume	Ambari	Oozie	Zookeeper	Knox
	Data management	Data access									Governance and integration			Operations		Security	

Microsoft Analytics Platform System

Appliance pre moderný Datawarehouse

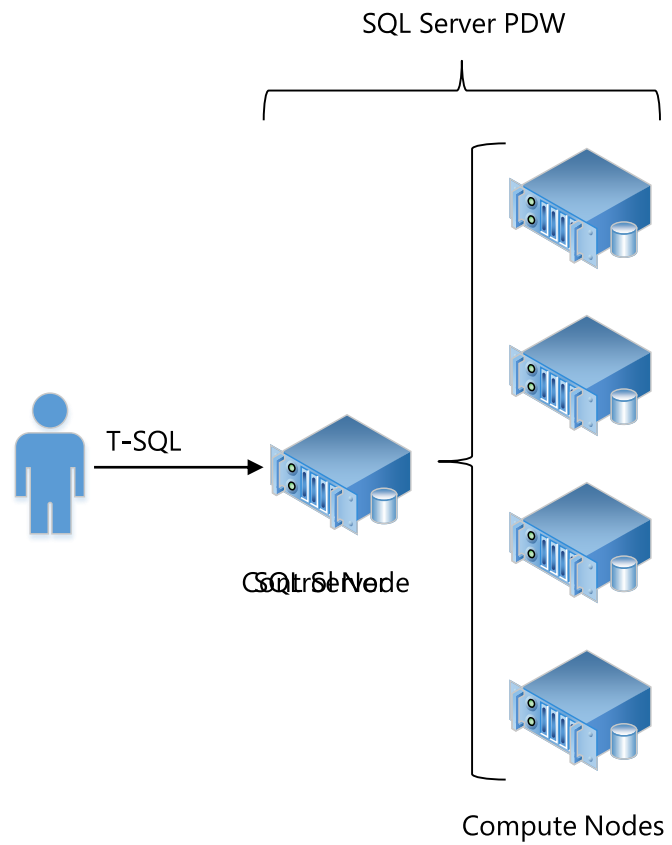


Čo je Parallel Data Warehouse?

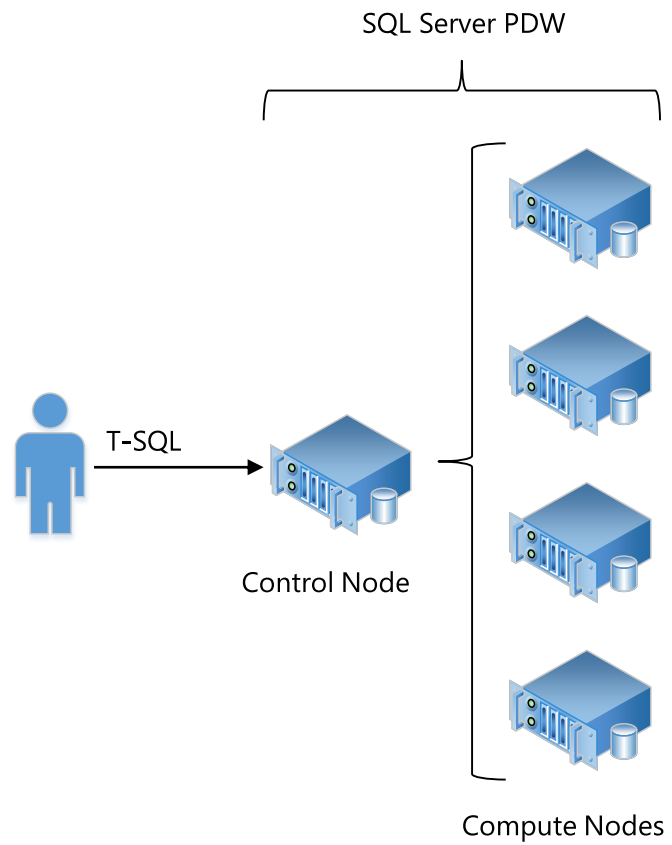
- **Shared-nothing parallel database system**
 - » Massively parallel processing (MPP)
 - » A “Control” server that accepts user queries, generates a plan, and distributes operations in parallel to compute nodes
 - » Multiple “Compute” servers running SQL Server
 - » A “Management” server for administering the system
 - » A “Data Movement Service” that facilitates parallel SQL operations
- **Delivered as an appliance**
 - » Balanced and pre-configured software and industry standard hardware from **HP**
 - » Single Call Support
 - » Fastest Time to Market
 - » Scales from 2 to 56 Nodes



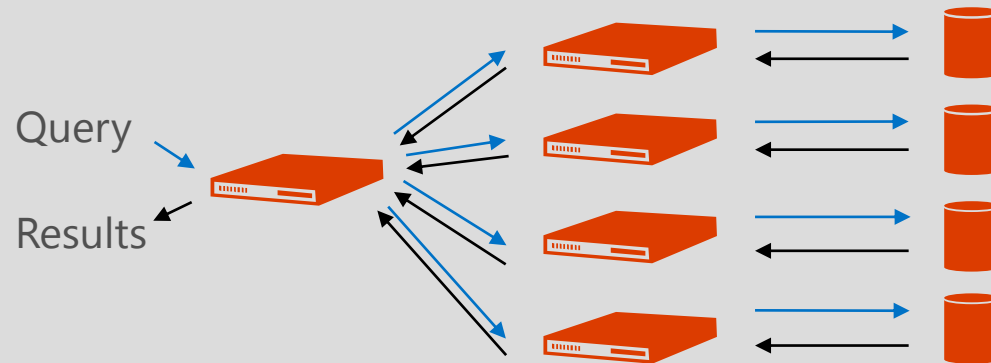
SQL Server Parallel Data Warehouse



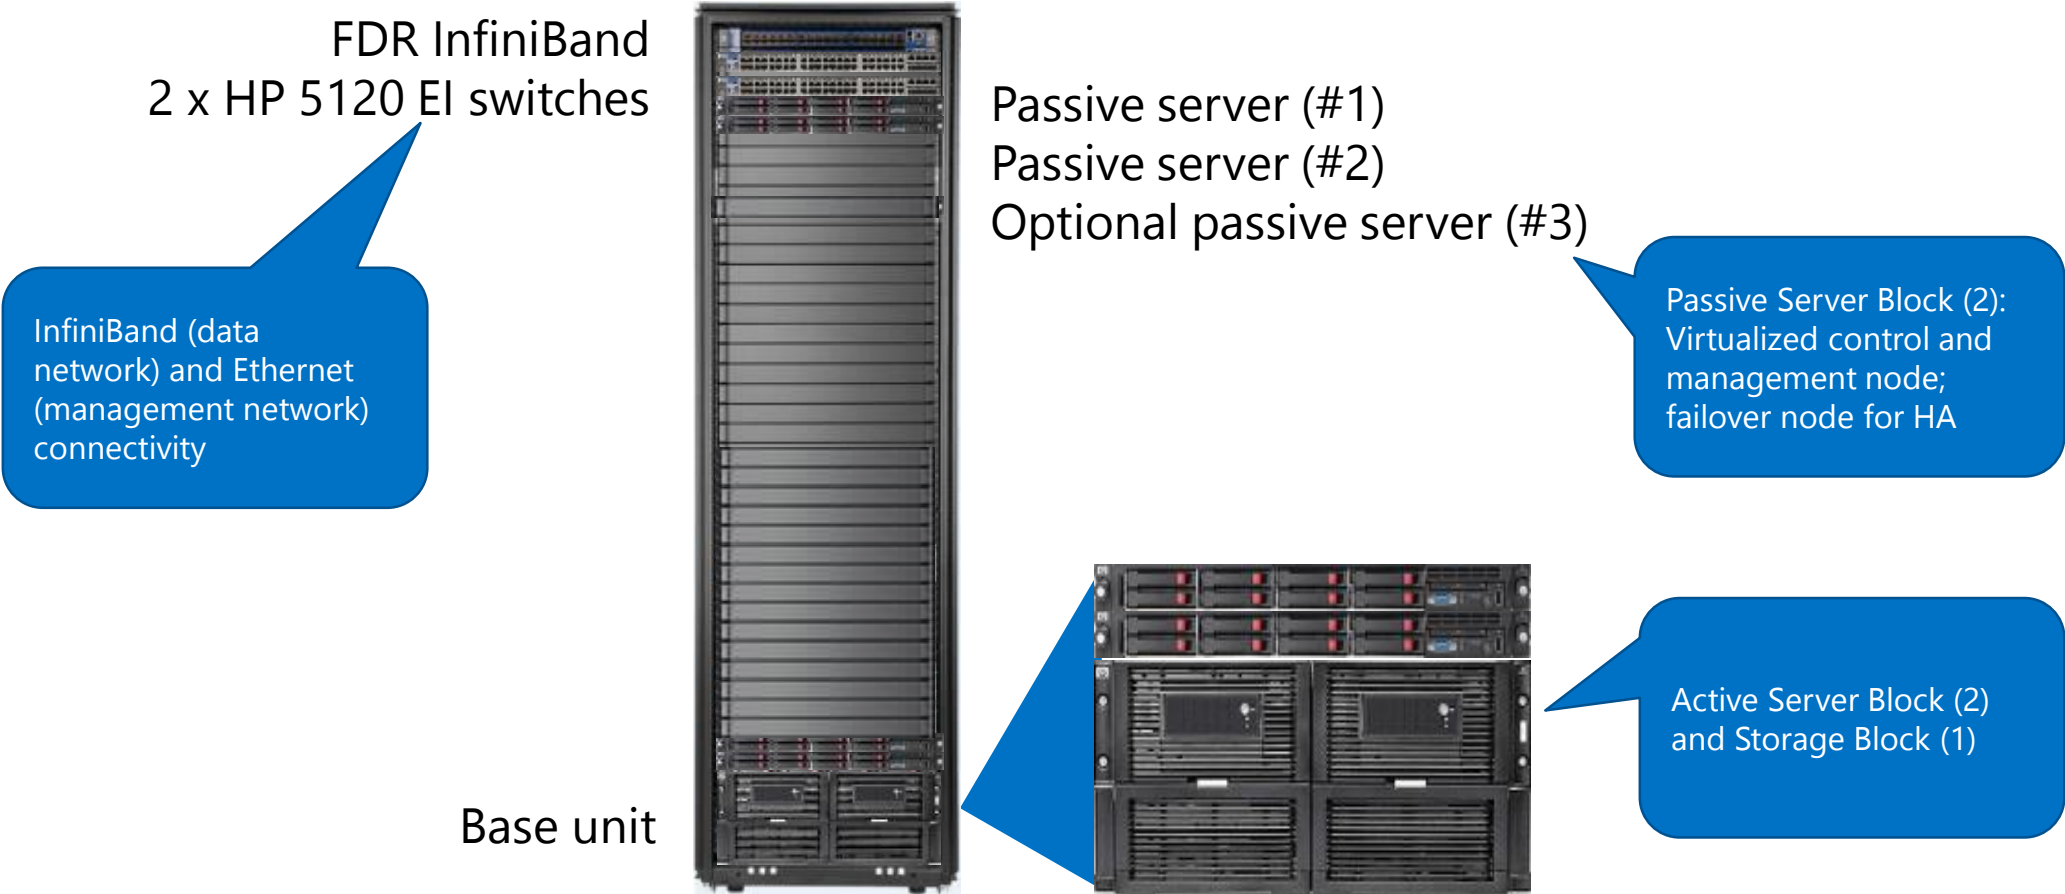
SQL Server Parallel Data Warehouse



Parallel query execution

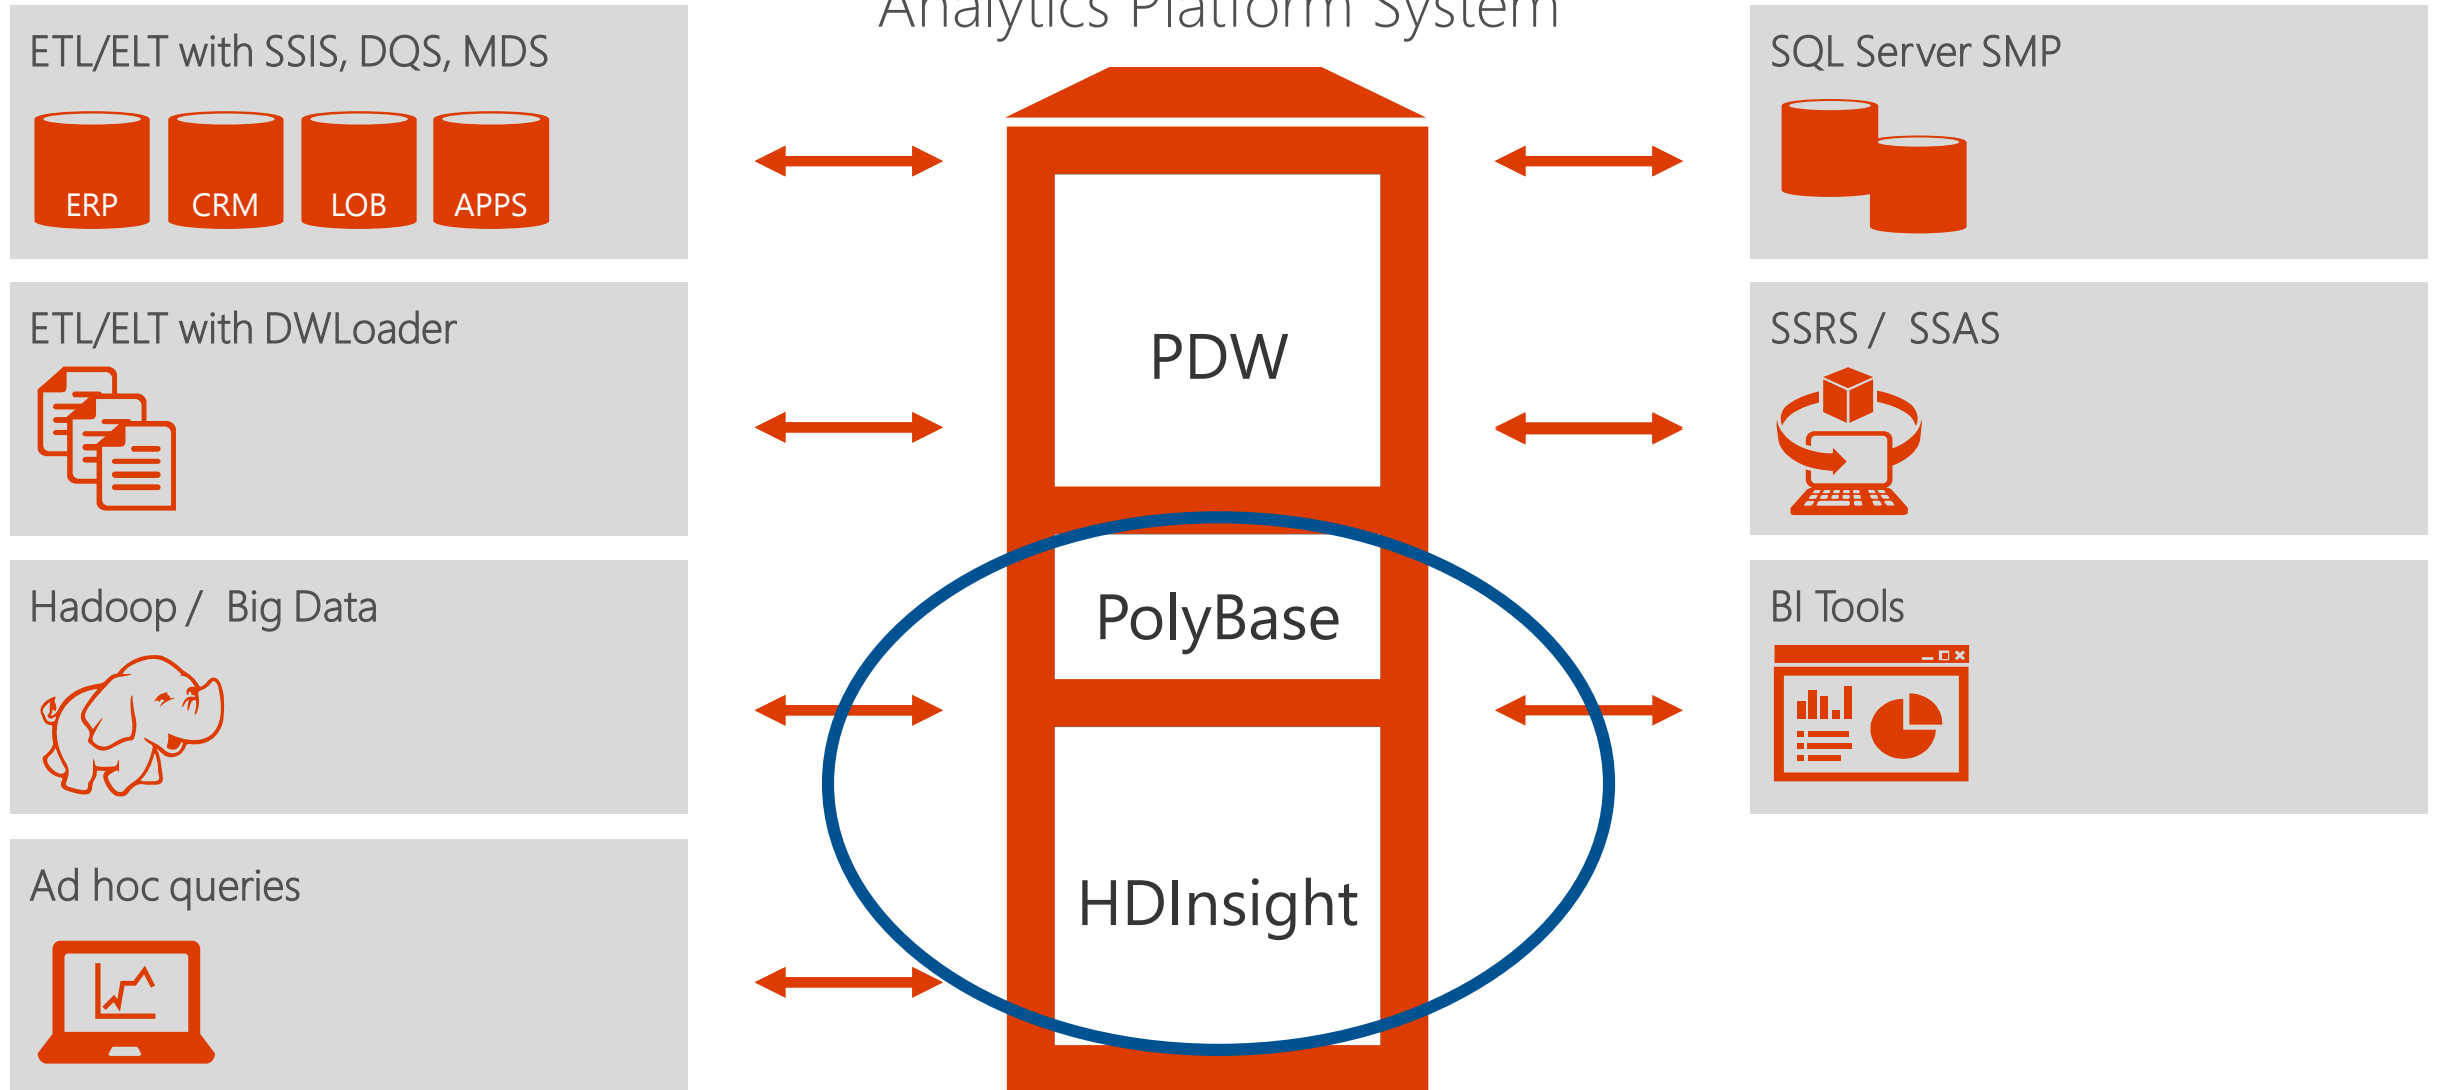


HP ConvergedSystem 300 for Microsoft Analytics Platform Base unit



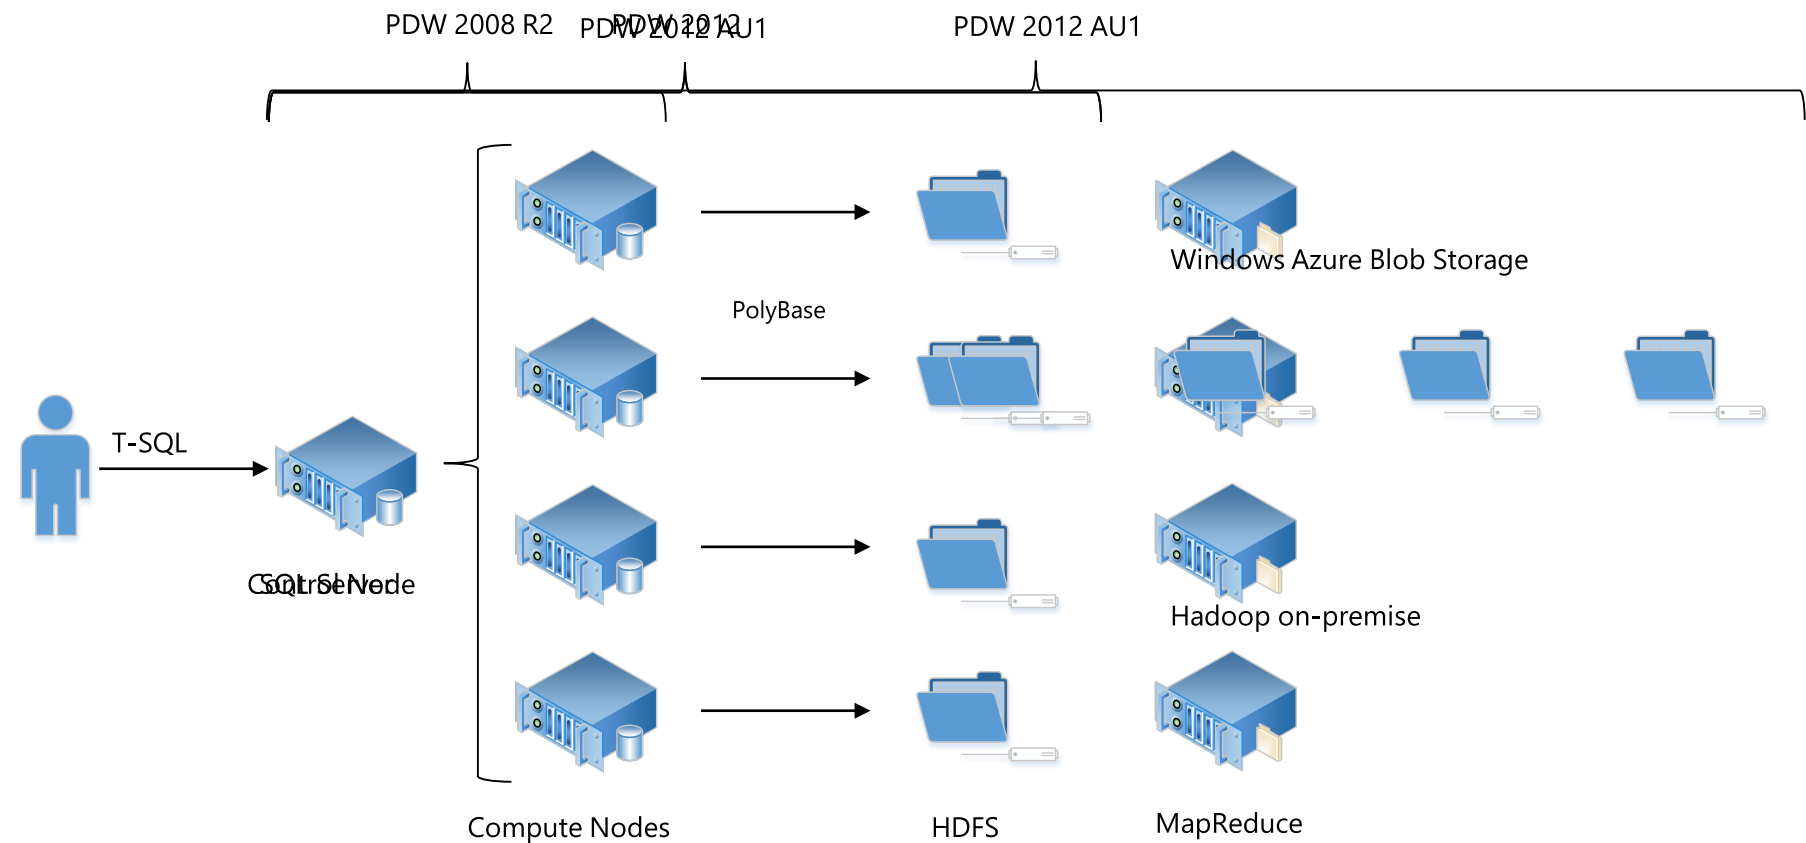
Microsoft Analytics Platform System

Appliance pre moderný datawarehouse



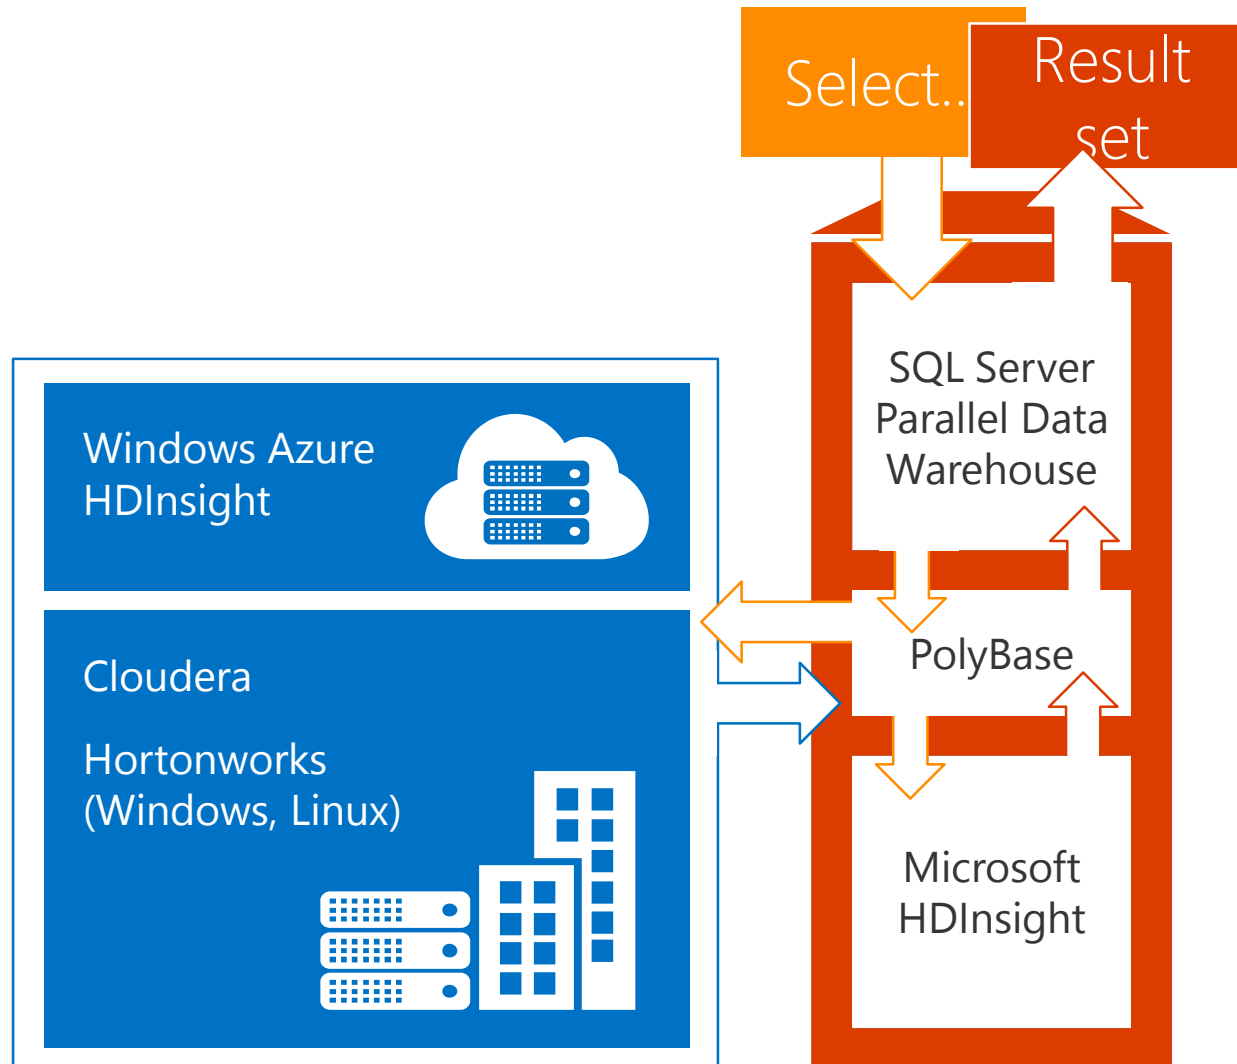
SQL Server Parallel Data Warehouse – Appliance Update 1

Extending the distributed Data Warehouse further



Query Hadoop data with T-SQL using PolyBase

Bringing the worlds of big data and the data warehouse together for users and IT



Single T-SQL query model for PDW and Hadoop with rich features of T-SQL including joins without ETL

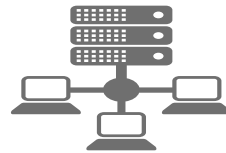
Leverages the power of MPP to enhance query execution performance

Supports Windows Azure HDInsight to enable new hybrid cloud scenarios

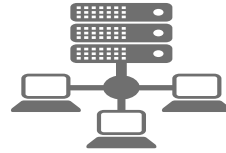
Query non-Microsoft Hadoop distributions such as Hortonworks and Cloudera

Access Hadoop on different cluster (cloud or on premise)

Hortonworks Data
Platform (HDP) on
Windows or Linux

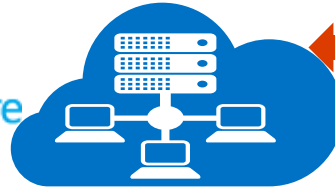


← Polybase →

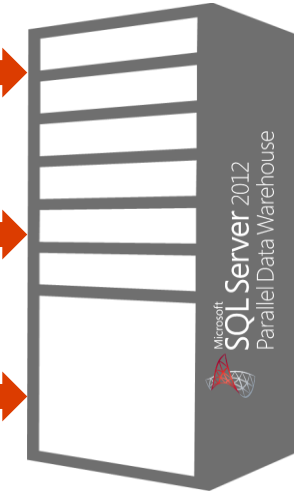


← Polybase →

HDInsight on Azure
(rebranded HDP)



← Polybase →

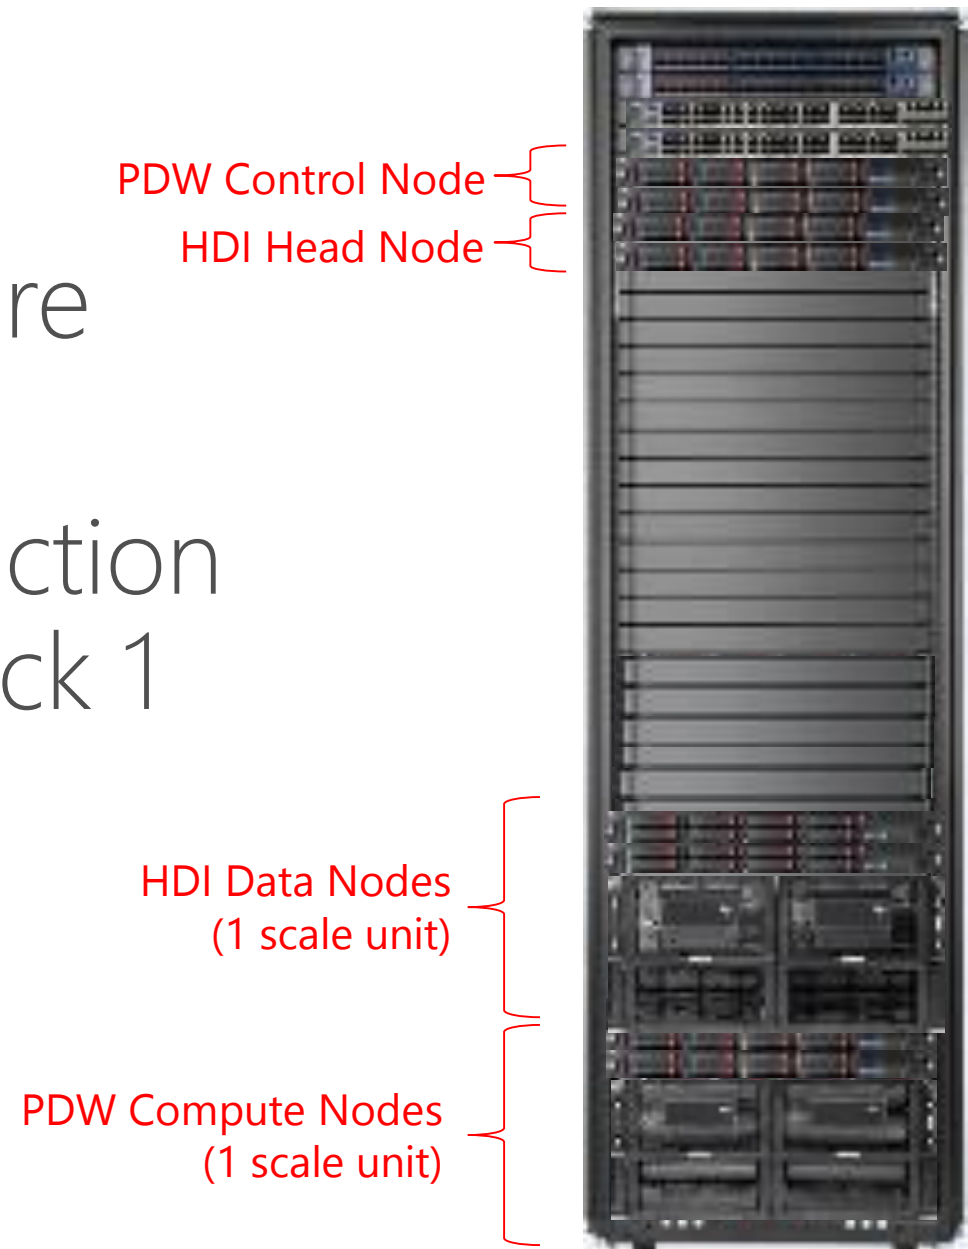


HDInsight is the Microsoft branded Hortonworks Data Platform

- **We made it work on Windows**
- **We brought SystemCenter support to it**

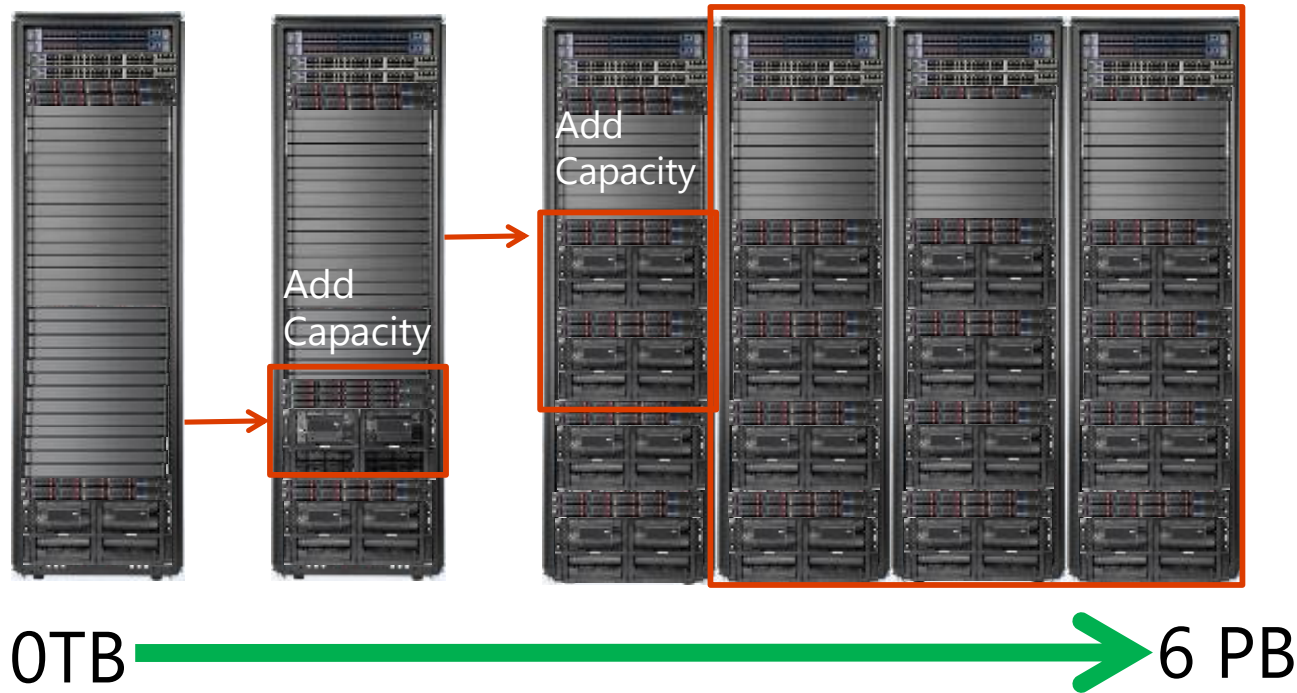
Hardware topology overview

- Uses the same PDW hardware and topology
- The key difference is introduction of 2 additional servers on rack 1 for the HDI Head Node
 - 1 active server and 1 fail over server.



Seamlessly add capacity

Scale from a Quarter Rack with 2 Compute Nodes up to 56 Compute Nodes!



Smallest (0TB) To Largest (6PB)

- Start small with a few Terabyte warehouse
- From 2 compute nodes to 56 compute nodes
- 1 quarter rack up to 7 full racks
- Add capacity up to 6 Petabytes

Start Small
And Grow



Largest
Warehouse

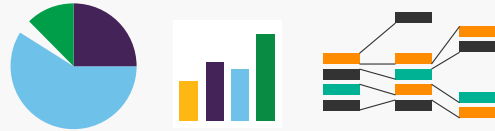
PB

Minimal
Downtime



Bringing Hadoop to a billion people

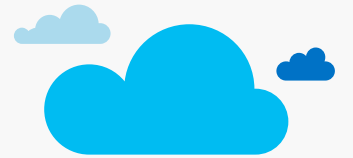
Excel as the BI tool
for everyone



1 Billion
Microsoft Office users

- Connect to HDInsight
- Analyze
- Visualize

Power BI for collaboration
& new experiences



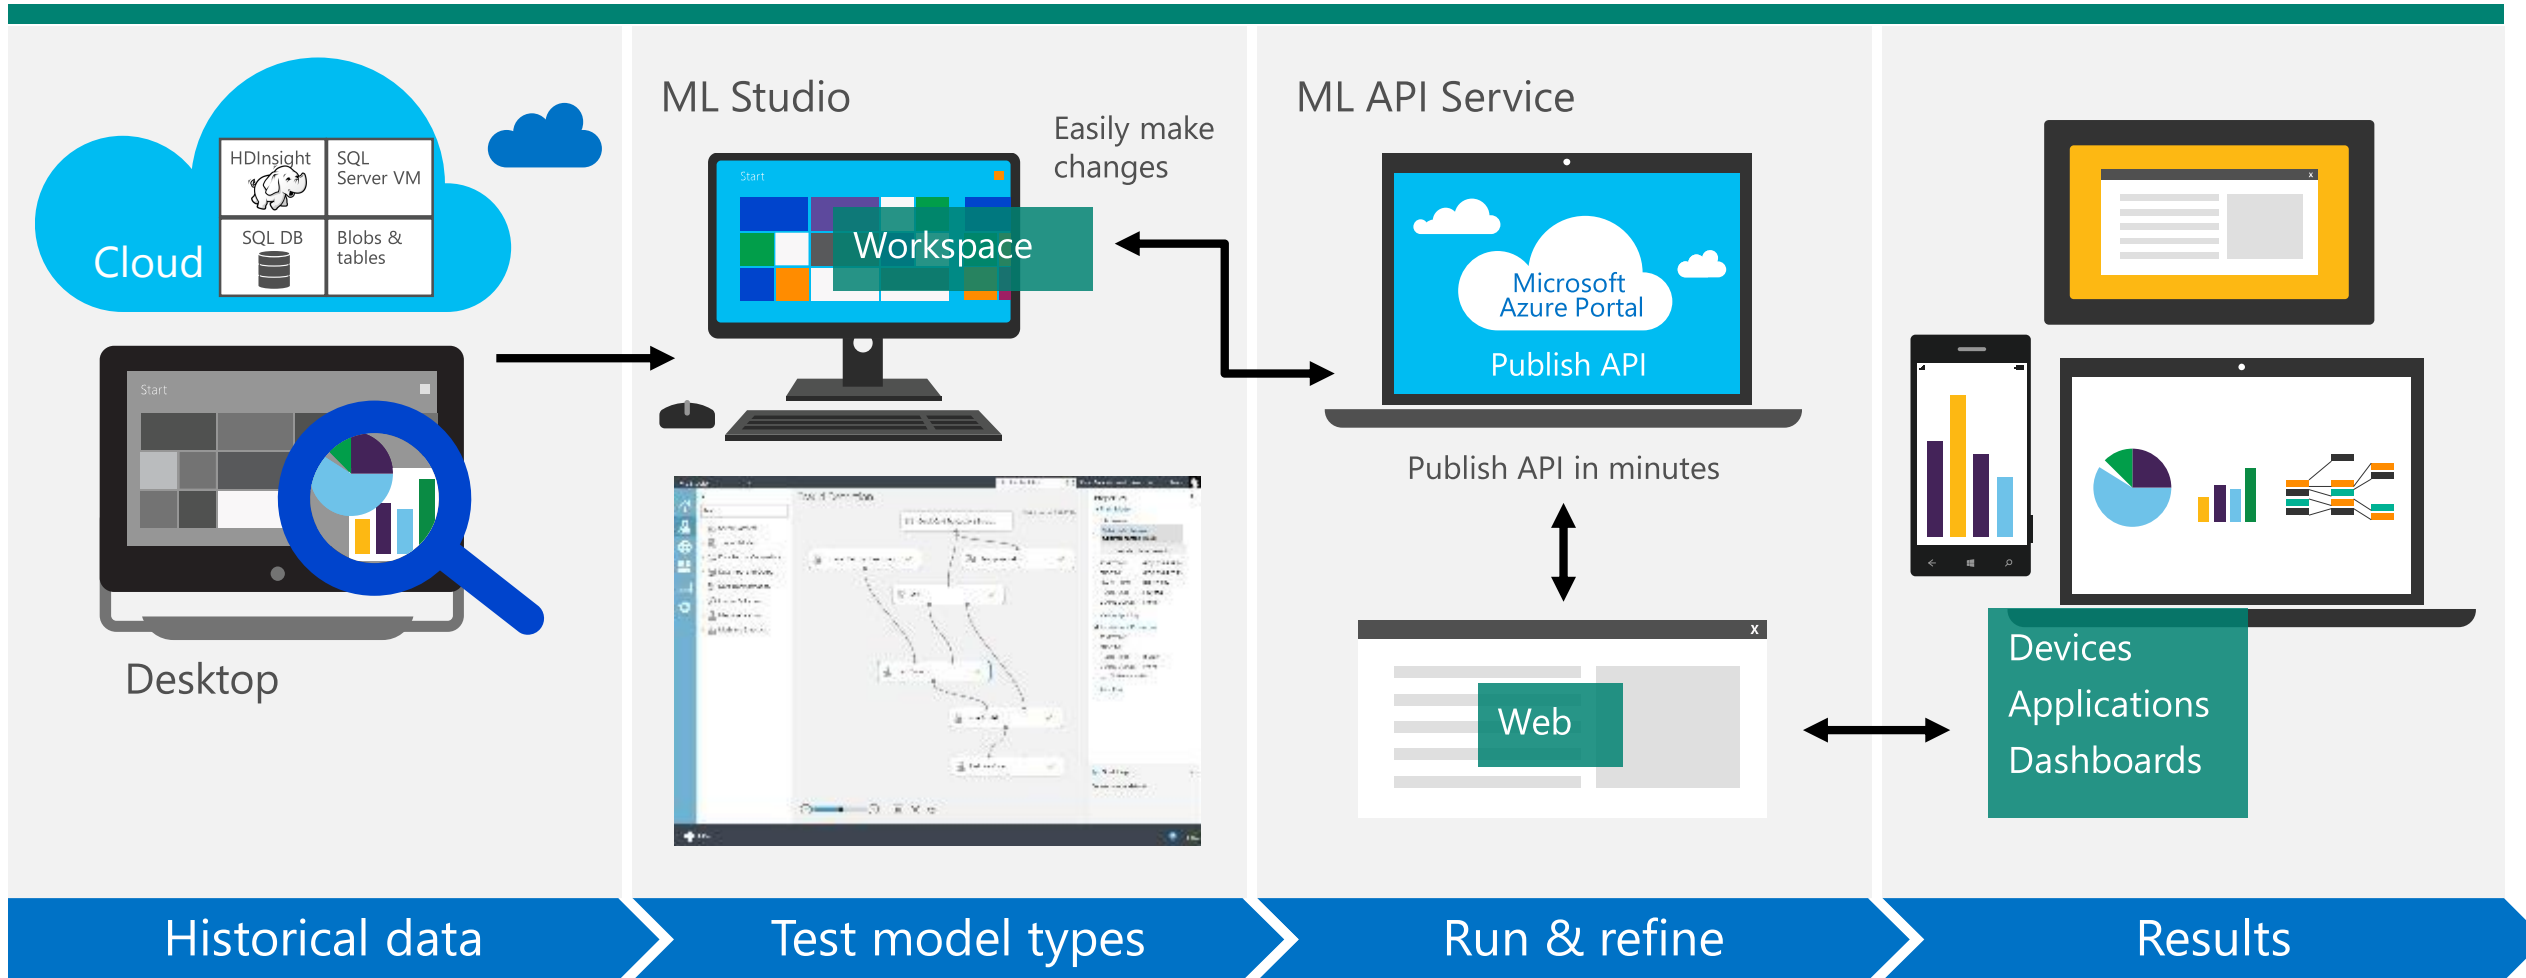
Office 365 is our fastest-growing
commercial product ever

- Share
- Ask
- Access

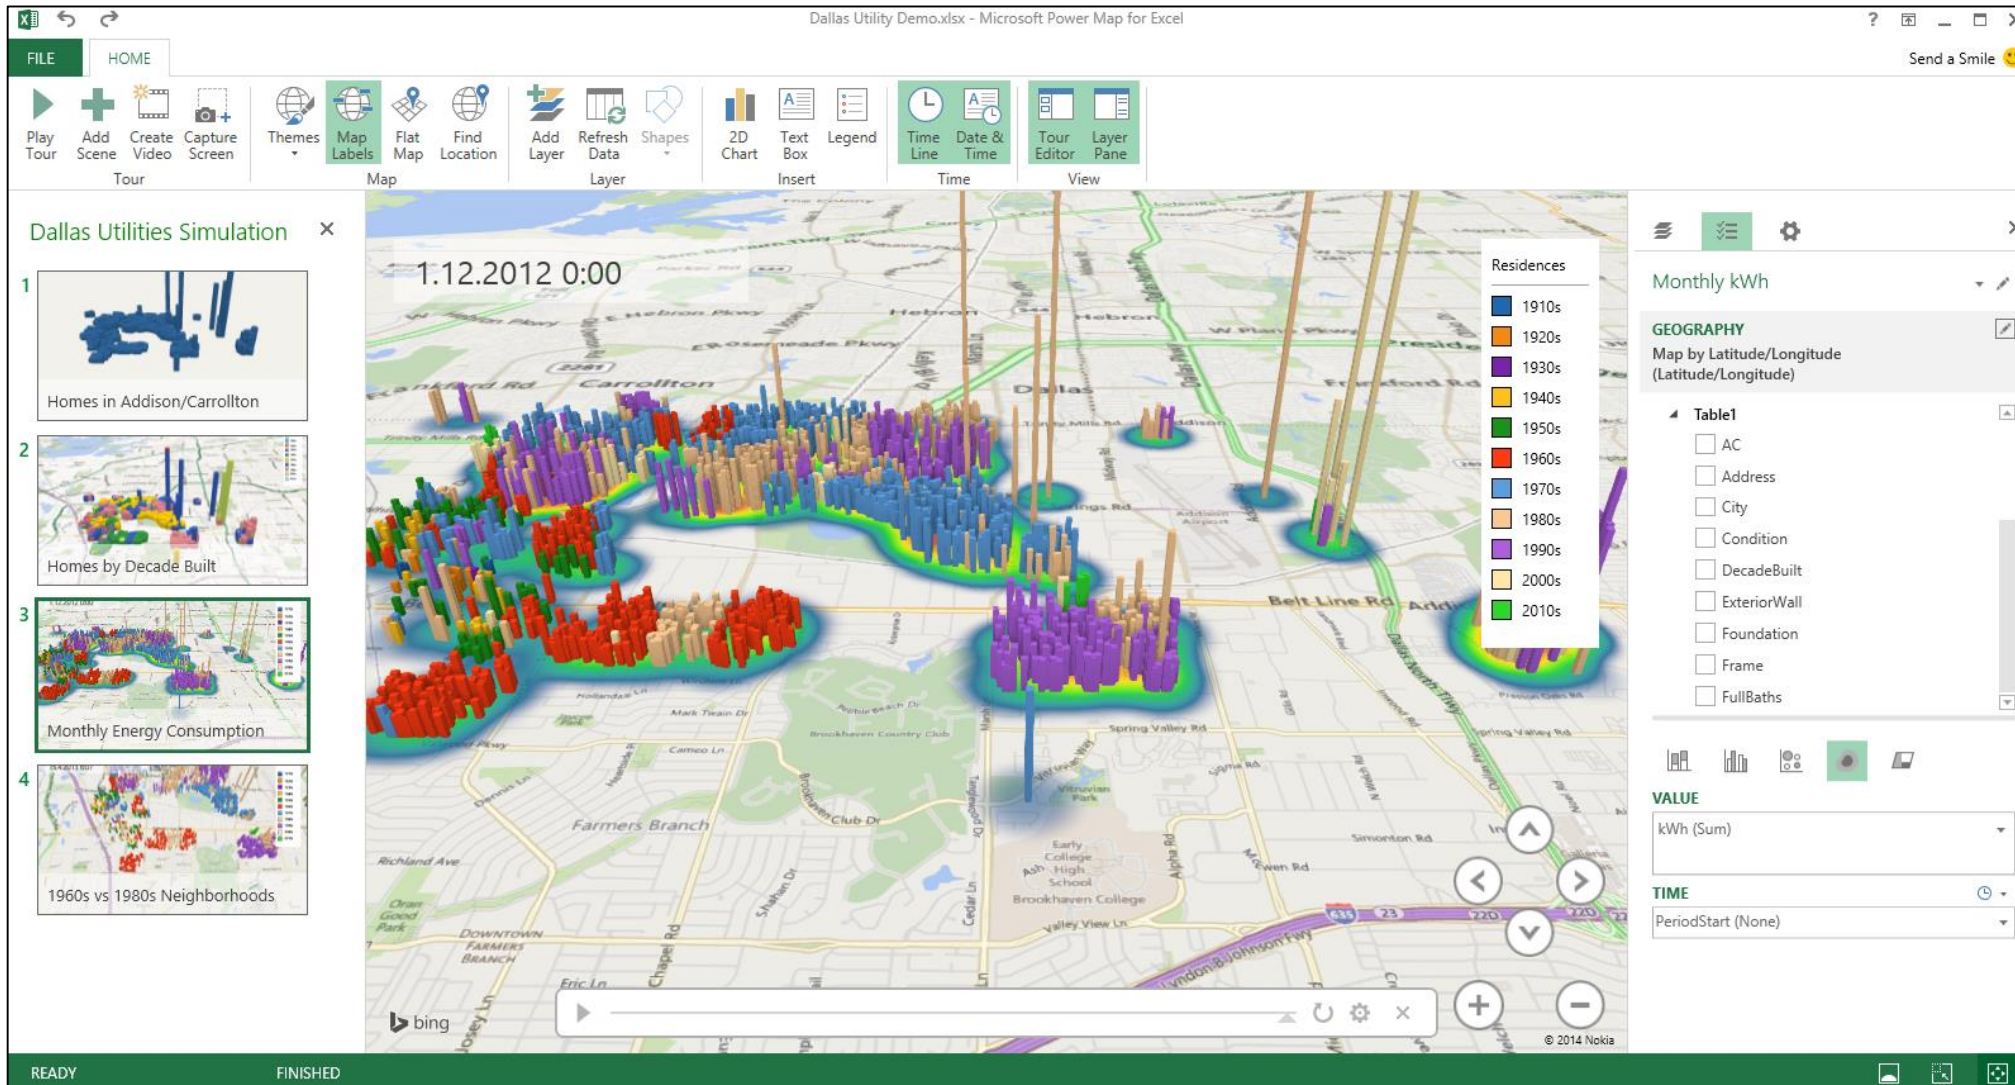
Scalable, manageable, trusted

Making advanced analytics accessible to Hadoop

Microsoft Azure Machine Learning

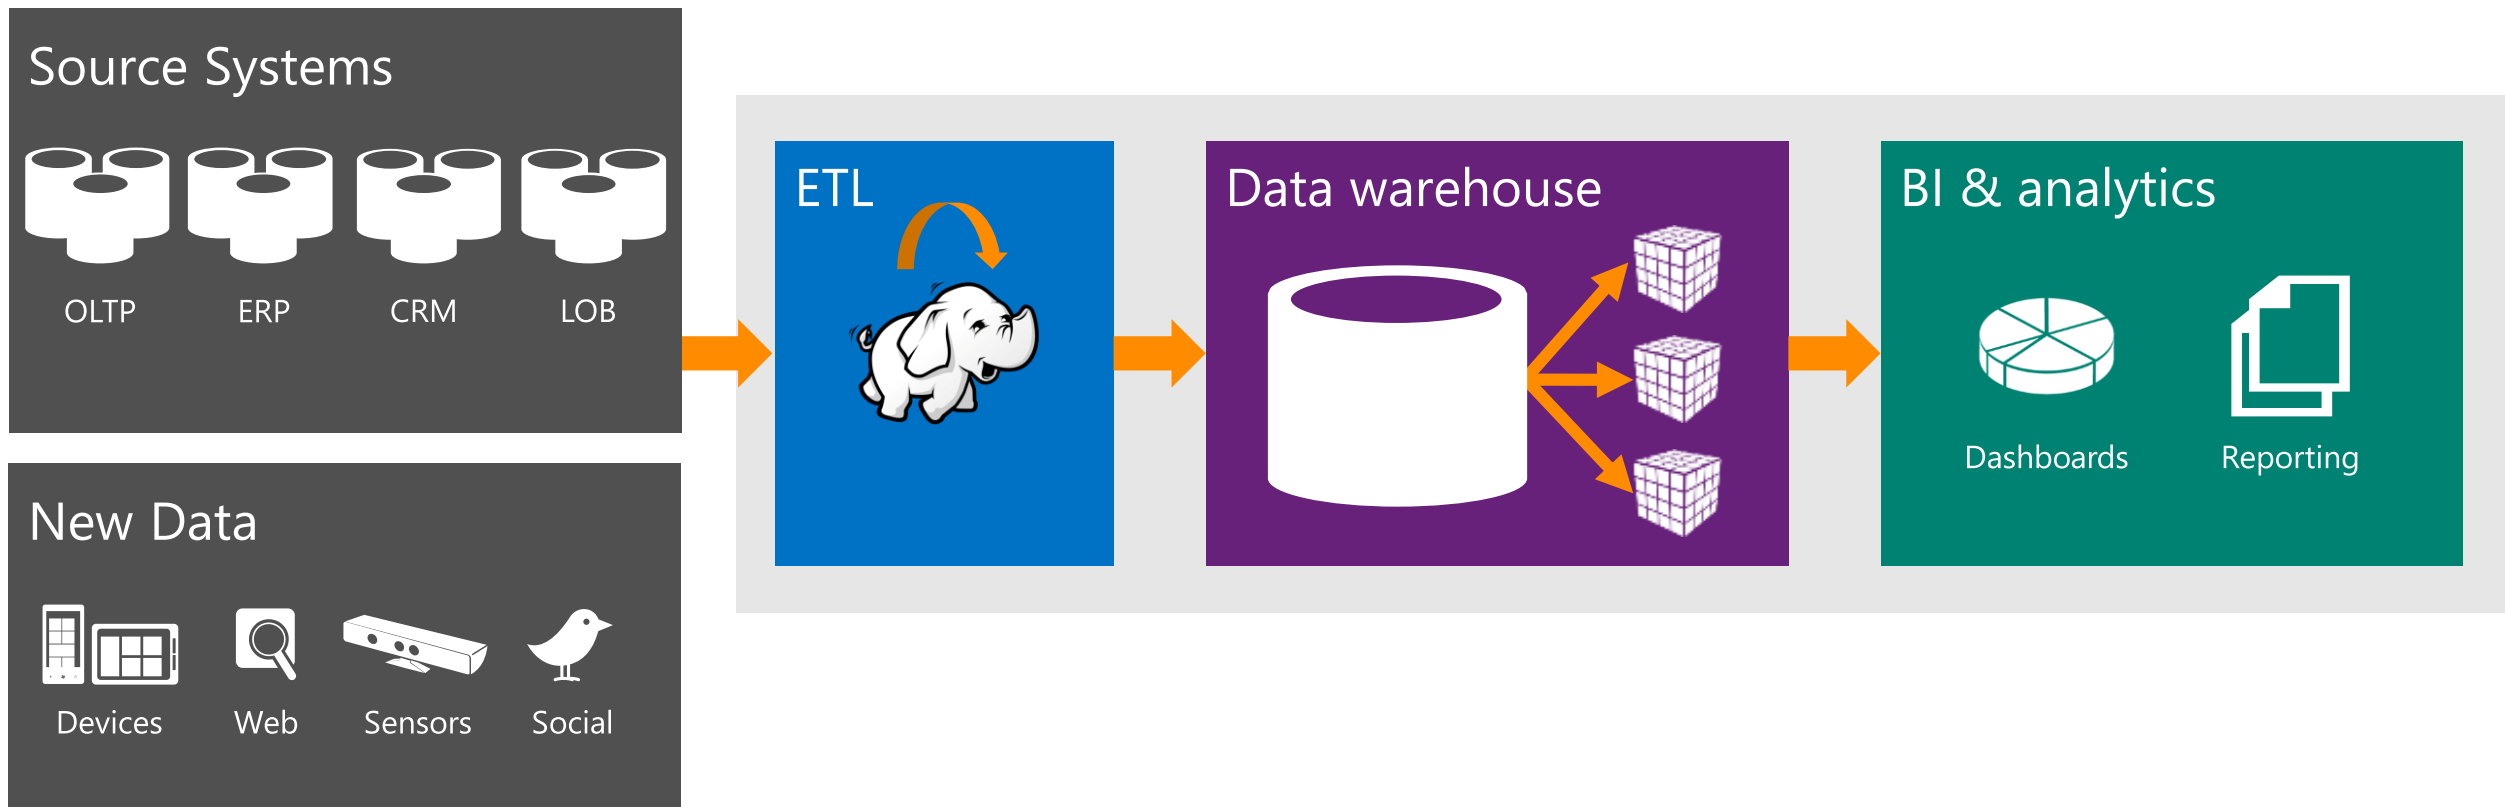


PowerBI, Excel PowerMap, PowerQuery, ...



Hadoop Scenario 1: pre-process ETL

Shift the pre-processing of ETL in staging data warehouse to Hadoop
Shifts high cost data warehousing to lower cost Hadoop clusters

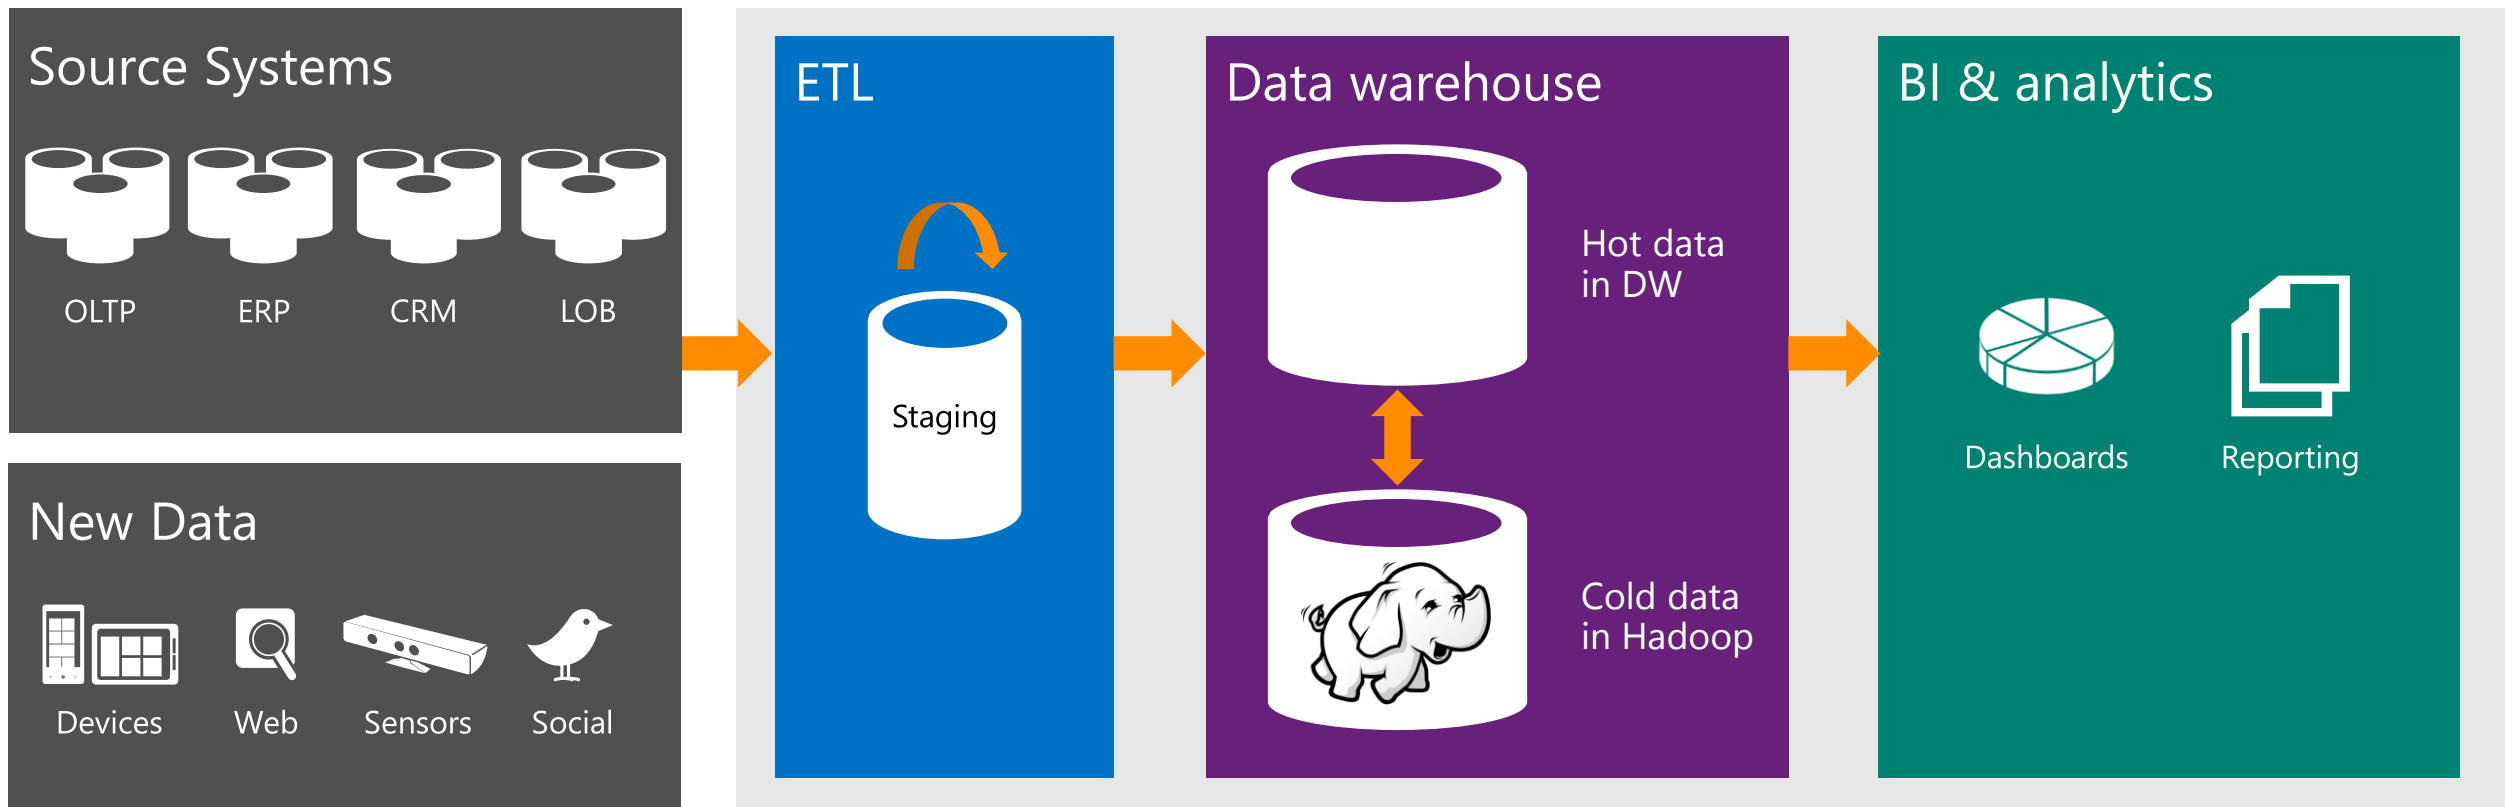


Hadoop Scenario 2: hot and cold storage

Offloading large volume of historical data into cold storage with Hadoop

Keep data warehouse for hot data to allow BI and analytics

When data from cold storage is needed, it can be moved back into the warehouse

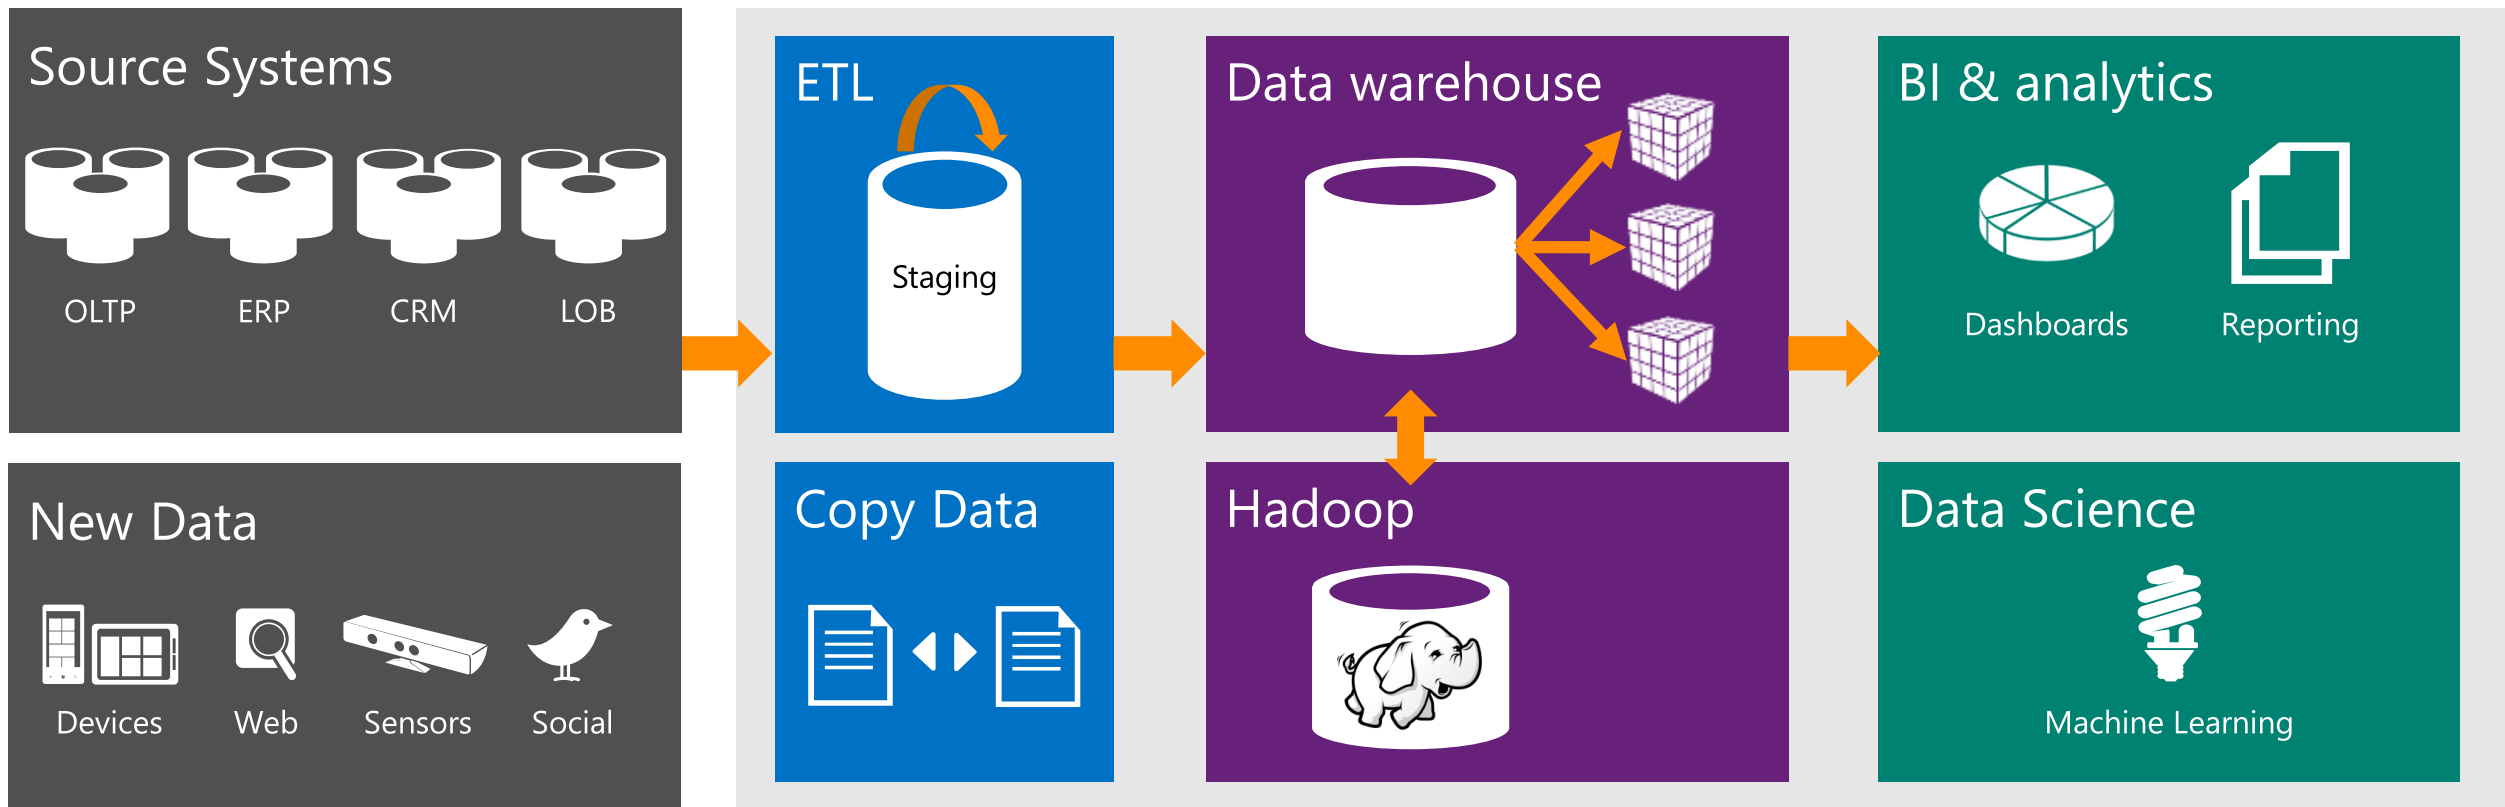


Hadoop Scenario 3: true data discovery

Keep data warehouse for operational BI and analytics

Allow data scientists to gain new discoveries on raw data (no format or structure)

Operationalize discoveries back into the warehouse



Industry Use Cases of Hadoop

Financial services

- New account risk screens
- Fraud prevention
- Trading risk
- Maximize deposit spread
- Insurance underwriting
- Accelerate loan processing



Retail

- 360° view of the customer
- Analyze brand sentiment
- Localized, personalized promotions
- Website optimization
- Optimal store layout



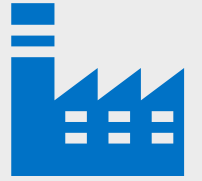
Telecom

- Call detail records (CDRs)
- Infrastructure investment
- Next product to buy (NPTB)
- Real-time bandwidth allocation
- New product development



Manufacturing

- Supplier consolidation
- Supply chain and logistics
- Assembly line quality assurance
- Proactive maintenance
- Crowd source quality assurance



Healthcare

- Genomic data for medical trials
- Monitor patient vitals
- Reduce re-admittance rates
- Store medical research data
- Recruit cohorts for pharmaceutical trials



Utilities, oil and gas

- Smart meter stream analysis
- Slow oil well decline curves
- Optimize lease bidding
- Compliance reporting
- Proactive equipment repair
- Seismic image processing



Public sector

- Analyze public sentiment
- Protect critical networks
- Prevent fraud and waste
- Crowd source reporting for repairs to infrastructure
- Fulfill open records requests





Get started today!

- For more information visit: <http://azure.microsoft.com/en-us/services/hdinsight/>



© 2014 Microsoft Corporation. All rights reserved. Microsoft, Windows, Windows Vista and other product names are or may be registered trademarks and/or trademarks in the U.S. and/or other countries. The information herein is for informational purposes only and represents the current view of Microsoft Corporation as of the date of this presentation. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information provided after the date of this presentation. MICROSOFT MAKES NO WARRANTIES, EXPRESS, IMPLIED OR STATUTORY, AS TO THE INFORMATION IN THIS PRESENTATION.