

Big Data and Power BI

Darwin Schweitzer

Contact Info

Darwin Schweitzer


Email: darsch@microsoft.com

Github: [datasnowman](https://github.com/datasnowman)

Twitter: [@DataSnowman](https://twitter.com/DataSnowman)

LinkedIn: [Cortana Intelligence Solutions](#)

Big data and analytics has huge potential



90%

of business leaders expect big data to dramatically change how they do business, *having a landmark affect similar to the development of the Internet.* Source: Accenture, 2016

5X

improvement in decision making by using advanced analytics

Source: Bain & Co., 2016

6-8%

productivity gains that result from investment in data analytics. *"This is a higher rate of return than other recent technologies have yielded, surpassing even the computer investment cycle in the 1980s"*

Source: McKinsey & Co., 2016

#devdaysasia2017

#azure

#azureiot

#Microsoft



NTRA @NTRA · May 20

What a finish. #Preakness2017 #CloudComputing pic.twitter.com/Rp6QR5OtAZ



2



96



143

#cloudcomputing

#bigdata

@MSPowerBI

@MSAdvAnalytics

#iot

Capture

Curate

Consume

Data Sources

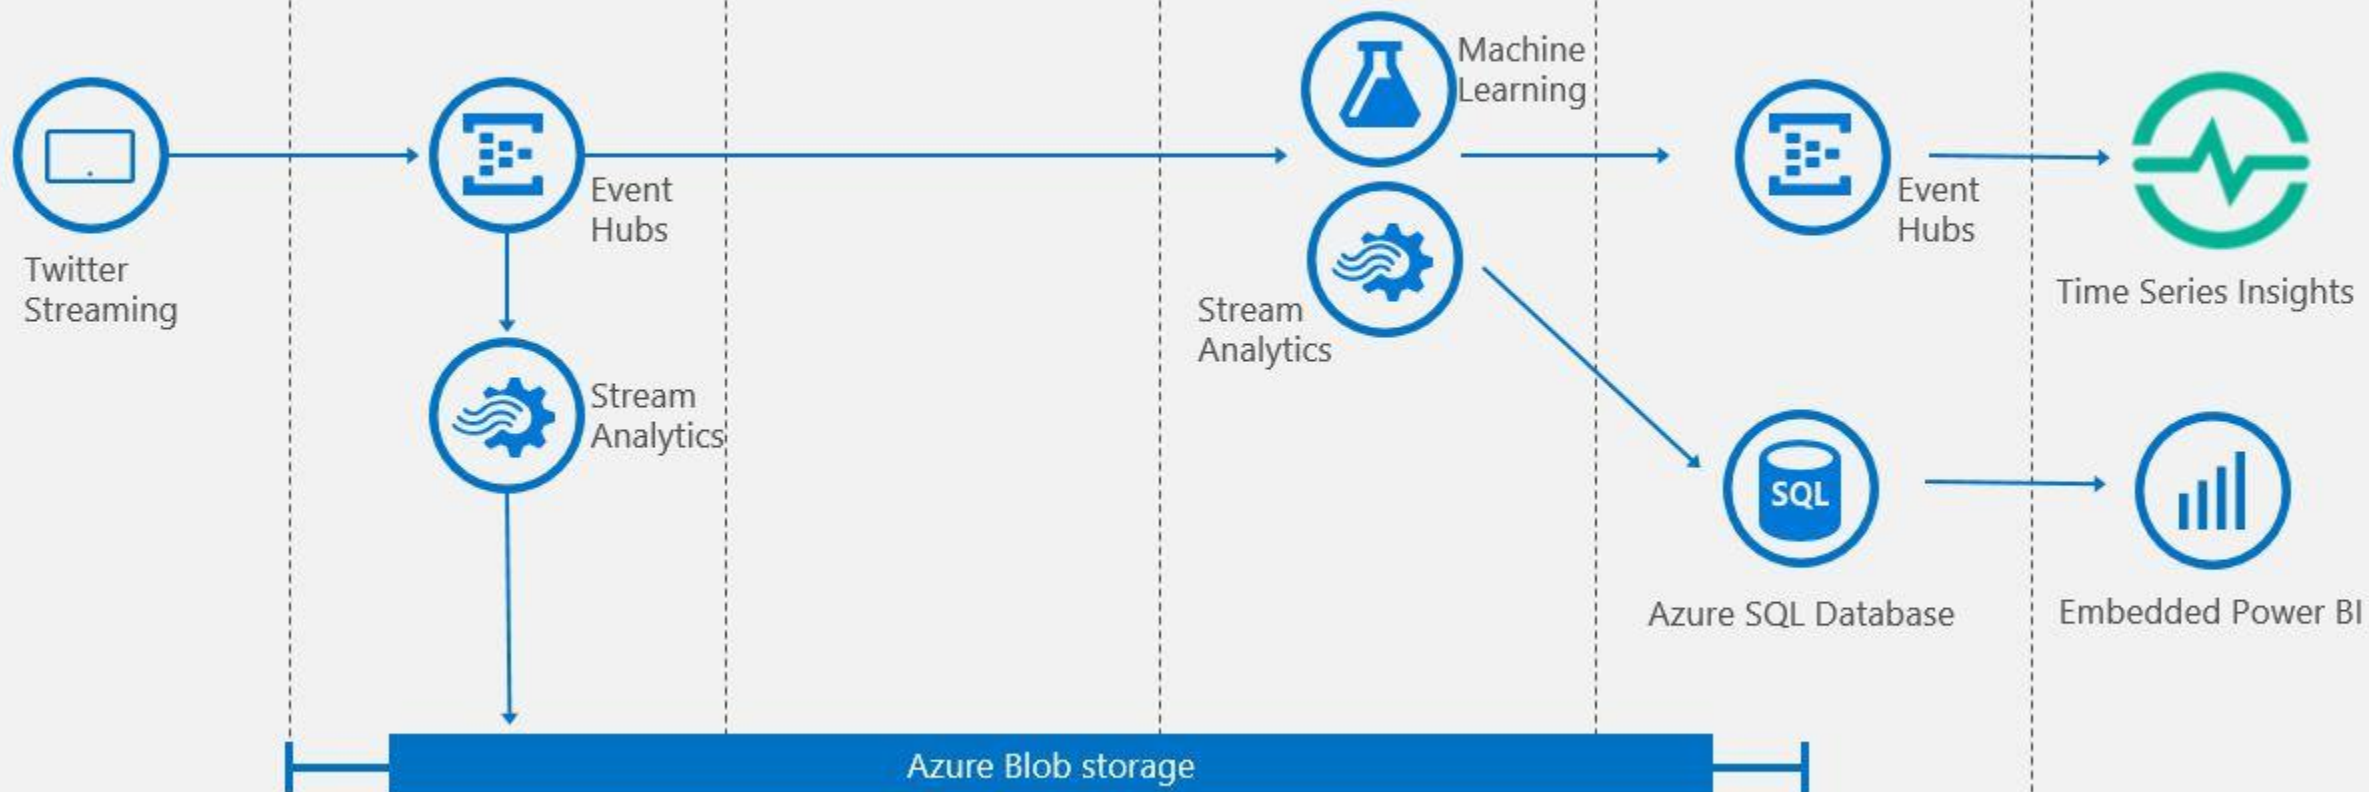
▶ Ingest

▶ Prepare

▶ Analyze

▶ Publish

▶ Visualize



<https://gallery.cortanaintelligence.com/Tutorial/Time-Series-Insights-using-Twitter>






Time Series Insights using Twitter

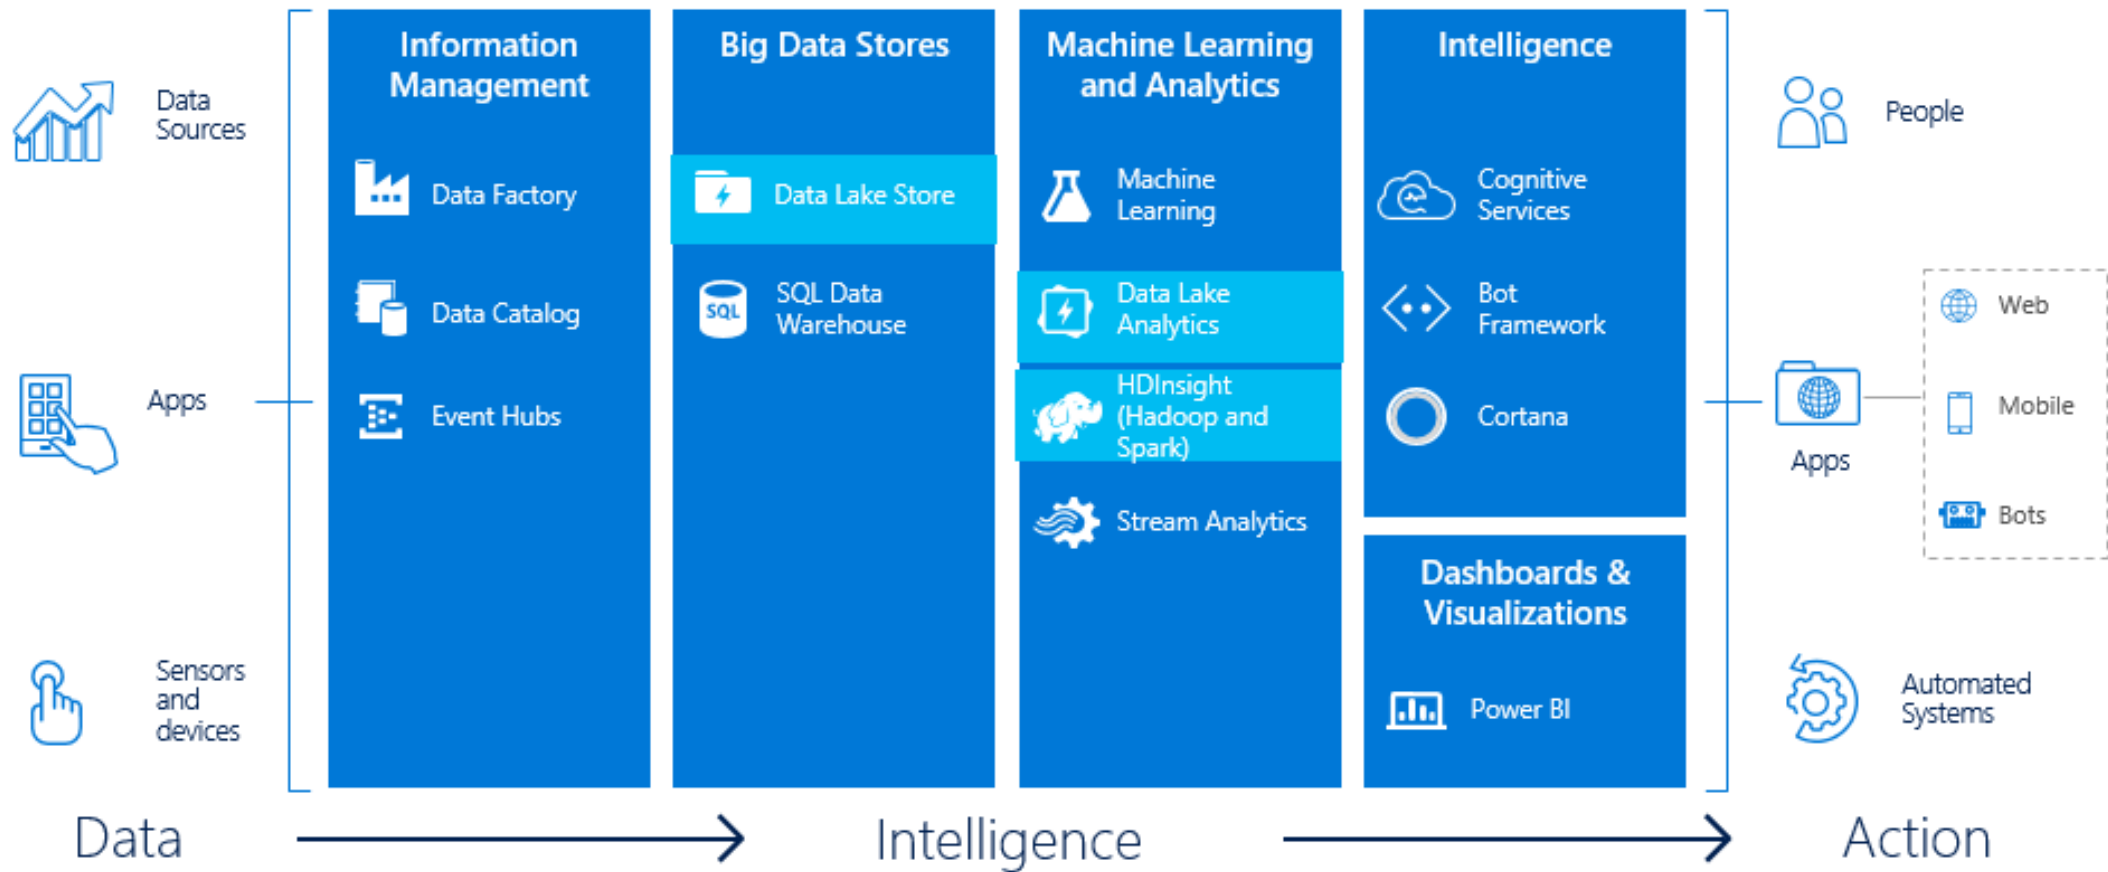
Demo

<https://gallery.cortanaintelligence.com/Tutorial/Time-Series-Insights-using-Twitter>

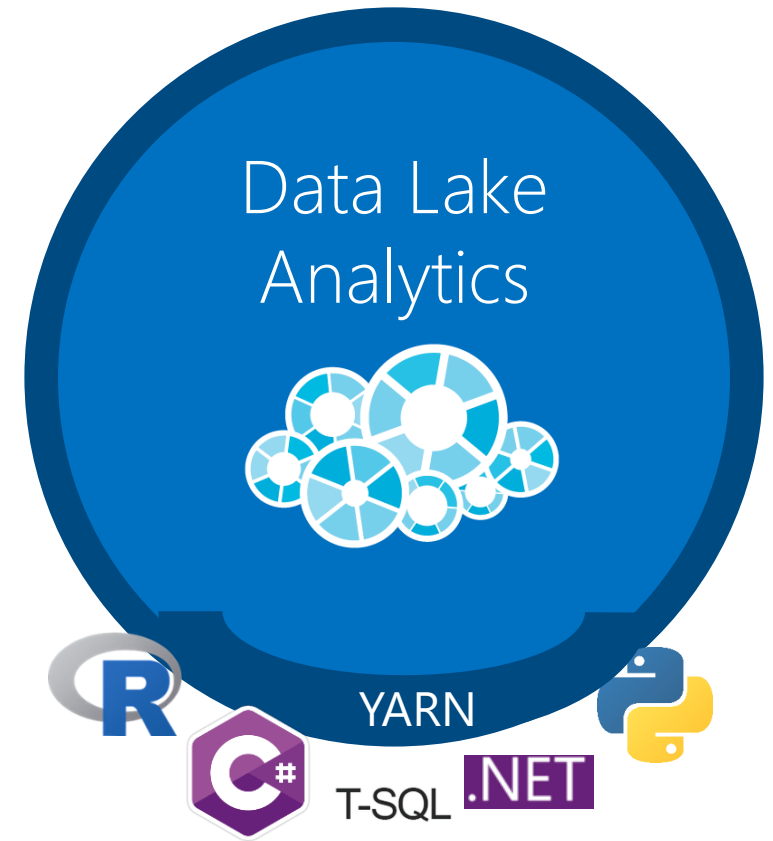
Big Data is driving transformative changes

	Traditional	Big Data
 Data characteristics	Relational data with highly modeled schema	All data with schema agility
 Costs	Specialized HW	Commodity HW
 Culture	Operational reporting Focus on rear-view analysis	Experimentation leading to intelligent action With machine learning, graph, a/b testing


Big Data as a cornerstone of Cortana Intelligence



Microsoft Big Data



Why Hadoop in the Cloud?



No HW costs



Unlimited scale

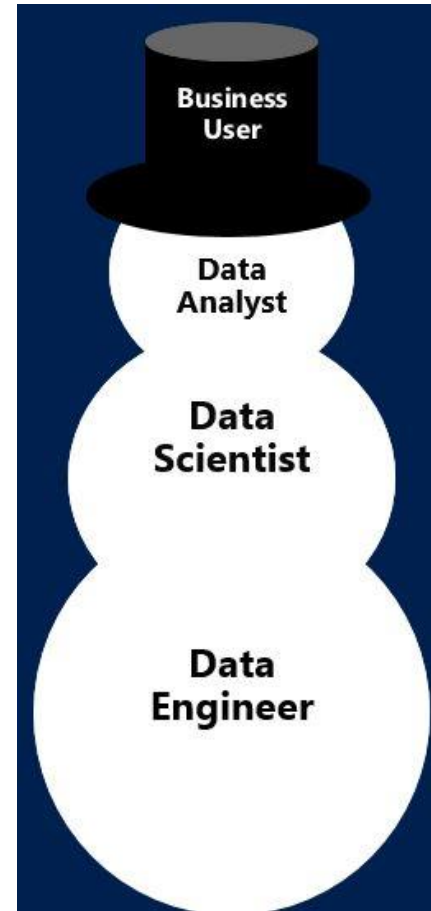


Pay what
you need



Deployed
in minutes

Something for all roles




Lower total cost of ownership



- No hardware
- Hadoop support included with Azure support
- Pay only for what you use
- Independently scale storage and compute
- No need to hire specialized operations team
- 63% lower total cost of ownership than on-premises*

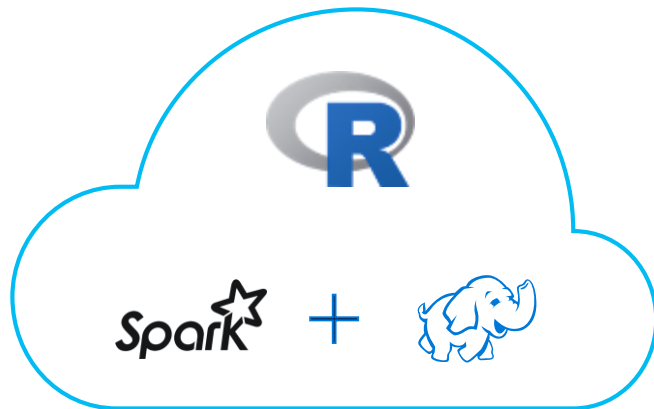
*IDC study "The Business Value and TCO Advantage of Apache Hadoop in the Cloud with Microsoft Azure HDInsight"

Easy for business analysts with interactive reports over big data



- Interactive BI with big data
- Spark 2.0 integration
- Interactive Hive with LLAP - keeps data compressed running in-memory 25x faster
- ODBC driver to use Power BI or third party tools (Tableau, SAP, Qlik, etc.)

Easy for data scientists with familiar R language

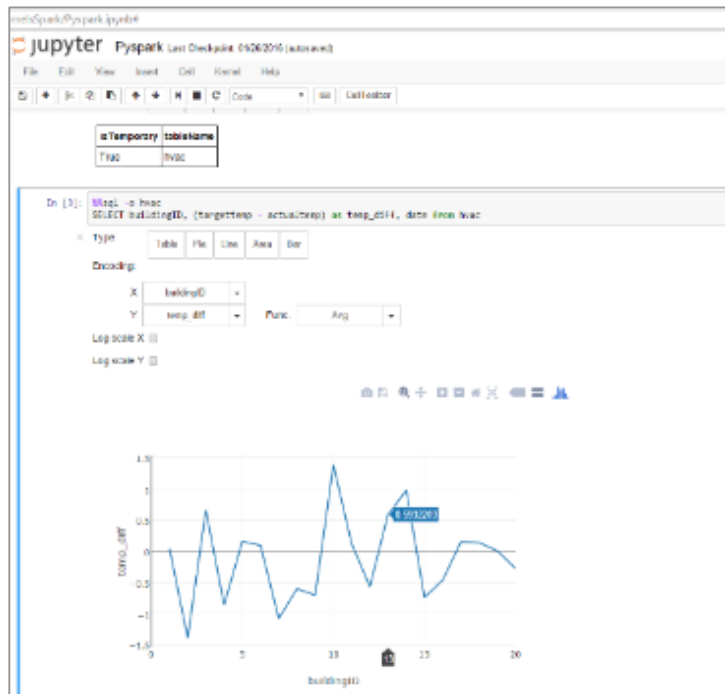


*Applies to HDInsight only

R Server for HDInsight

- Largest portable R parallel analytics library
- Terabyte-scale machine learning—1,000x larger than in open source R
- Up to 100x faster performance using Spark and optimized vector/math libraries
- Enterprise-grade security and support

Easy notebook experience for data engineers



- Most popular notebooks, Jupyter and Zeppelin out-of-the-box
- Combine code, statistical equations and visualizations
- Worked w/ Jupyter community to enhance kernel to allow Spark execution through REST endpoint

Debug and Optimize your Big Data programs with ease

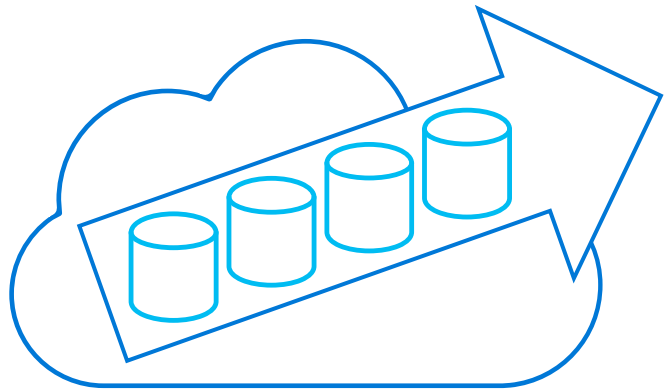
The screenshot displays the HiveServer2 IDE interface. On the left, a 'Summary' panel shows job details: State (FINISHED), Progress (100%), Name (HIVE-6358730e-c491-4ef2-b76), Cluster Id (1445304274202), Elapsed Time (00:01:41.0570000), Start Time (11/1/2015 4:52:04 AM), End Time (11/1/2015 4:53:45 AM), Application Type (TEZ), and User Name (hdp). The main area shows a 'Task Execution Detail' graph with a tree structure of Map and Reduce tasks. Below the graph is a Hive query editor with the following code:

```
1 use tpods_text_2;  
2 set hive.execution.engine=tez;  
3 select i_item_id  
4 , i_item_desc  
5 , s_state  
6 , count(ss_quantity) as store_sales_quantitycount
```

At the bottom, a 'HiveServer2 Output' window shows logs, including the message: 'INFO : Tez session hasn't been created yet. Opening session' and 'INFO : Status: Running (Executing on YARN cluster with App id application_1445304274202_0026)'. The output also shows progress for Map and Reduce tasks.

- Deep integration with IDEs for developer productivity: Visual Studio, Eclipse, & IntelliJ
- Integrated with Hive, Pig, Storm, and Spark
- Visually see execution of Hive jobs ran by the Tez execution engine
- Full Intellisense

Easy for administrators to spin up quickly



- Deploy big data projects in minutes
- No hardware to install, tune, configure or deploy
- No infrastructure or software to manage
- Scale to tens to thousands of machines instantly

Recognized by top analysts




Forrester Wave for Big Data Hadoop Cloud

- Named industry leader by Forrester with the most comprehensive, scalable, and integrated platforms*
- Recognized for its cloud-first strategy that is paying off*

*The Forrester Wave™: Big Data Hadoop Cloud Solutions, Q2 2016.

ISV Integration




Integration with leading productivity applications

Spin up Hadoop and Spark clusters pre-integrated and pre-tuned with ISV applications out-of-the-box

Runs on the HDInsight clusters; does not require separate VMs

Fast and easy way to spin up applications



Azure Data Lake Store

A No limits Data Lake that
powers Big Data Analytics



Petabyte size files and **Trillions of objects**

Scalable throughput for **massively parallel analytics**

HDFS for the cloud

Always encrypted, role-based security & auditing

Enterprise-grade support

Azure Data Lake Analytics

A No limits Analytics Job
Service to power intelligent
action



Start in seconds, scale instantly, pay per job

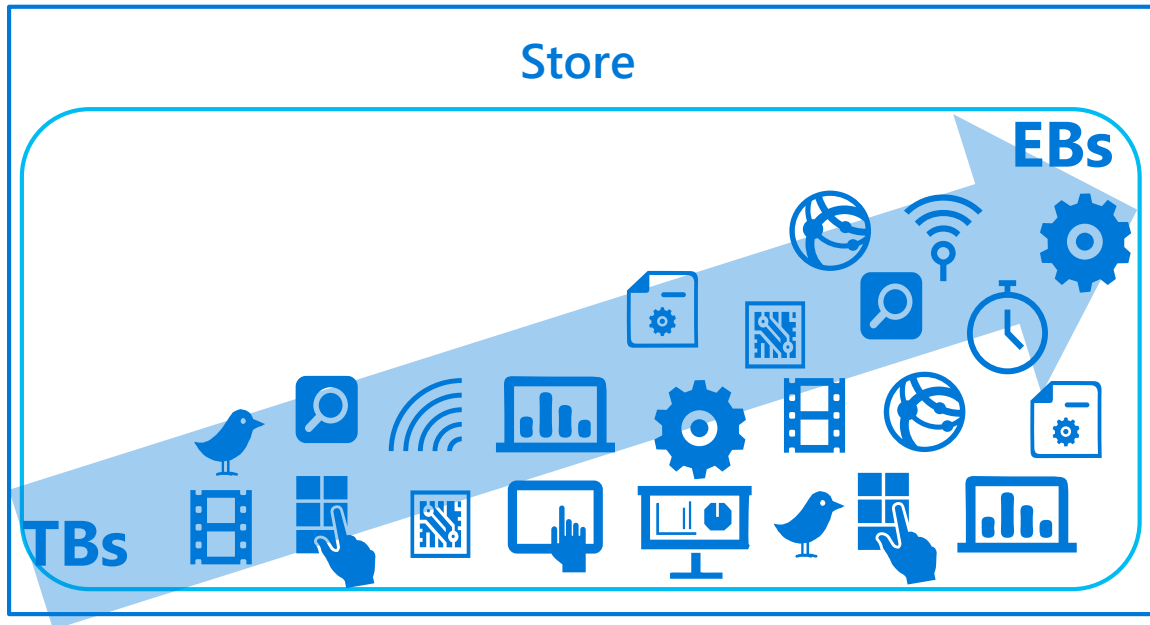
Develop massively parallel programs **with simplicity**

Debug and **optimize** your big data programs **with ease**

Virtualize your analytics

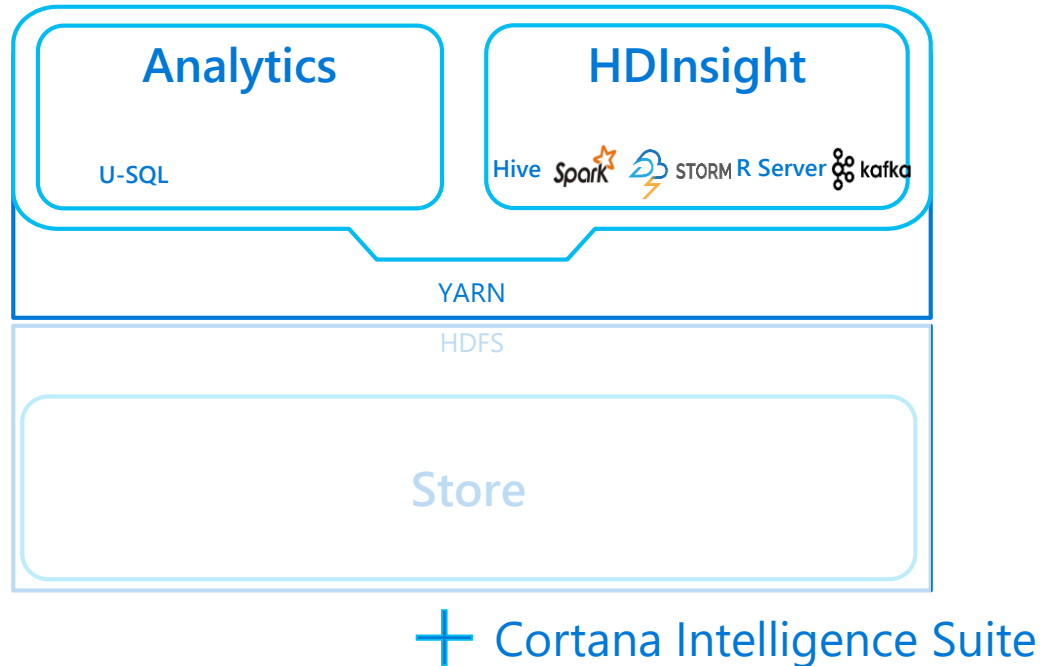
Enterprise-grade security, auditing and **support**

Petabyte size files and Trillions of objects




- Store data in it's native format
- PB sized files, **200x** larger than anyone else
- Scalable throughput for massively parallel analytics
- No need to redesign application or reparation data at higher scale

Any type of analytics



- Batch, interactive, streaming, machine learning
- Allows for exploratory analytics over data
- Analyze with Hadoop and Microsoft solutions

Start in seconds, Scale instantly, Pay per job



- Process big data jobs in **30 seconds**
- No infrastructure to worry about (no servers, no VMs, no clusters)
- Instantly scale analytic units up or down (processing power)
- Architected for cloud scale and performance
- Frees you up to focus only on your business logic

Develop massively parallel programs with simplicity

```
@t = EXTRACT date string
      , time string
      , author string
      , tweet string
FROM "/Input/MyTwitterHistory.csv"
USING Extractors.Csv();

@res = SELECT author AS author
        , COUNT(*) AS tweetcount
FROM @t
GROUP BY author;


OUTPUT @res TO "/Output/MyTwitterAnalysis.csv"
ORDER BY tweetcount DESC
USING Outputters.Csv();
```

- U-SQL: a simple and powerful language that's familiar and easily extensible
- Unifies the declarative nature of SQL with expressive power of C#
- Leverage existing libraries in .NET languages, R and Python
- Massively parallelize code on diverse workloads (ETL, ML, image tagging, facial detection)

Azure SQL Data Warehouse

A relational **data warehouse-as-a-service**, fully managed by Microsoft.
Industries first **elastic** cloud data warehouse with proven SQL Server capabilities.
Support your **smallest to your largest** data storage needs.

Elastic scale & performance

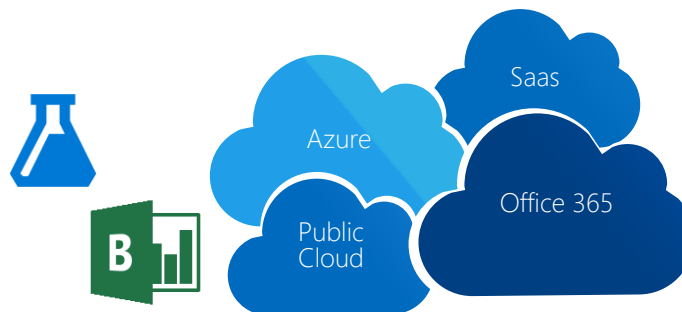


- Scales to petabytes of data
- Massively Parallel Processing
- Instant-on compute scales in seconds
- Query Relational / Non-Relational

Powered by the Cloud

Get started in minutes

Integrated with Azure ML, PowerBI & ADF



Market Leading Price & Performance



Simple billing compute & storage

Pay for what you need, when you need it
with dynamic pause

A fully managed Platform-as-a-Service

- Azure cloud data warehouse service
- Elastic scale
- Separate storage and compute
- Use existing tools and skills
- Deploy and use in minutes!

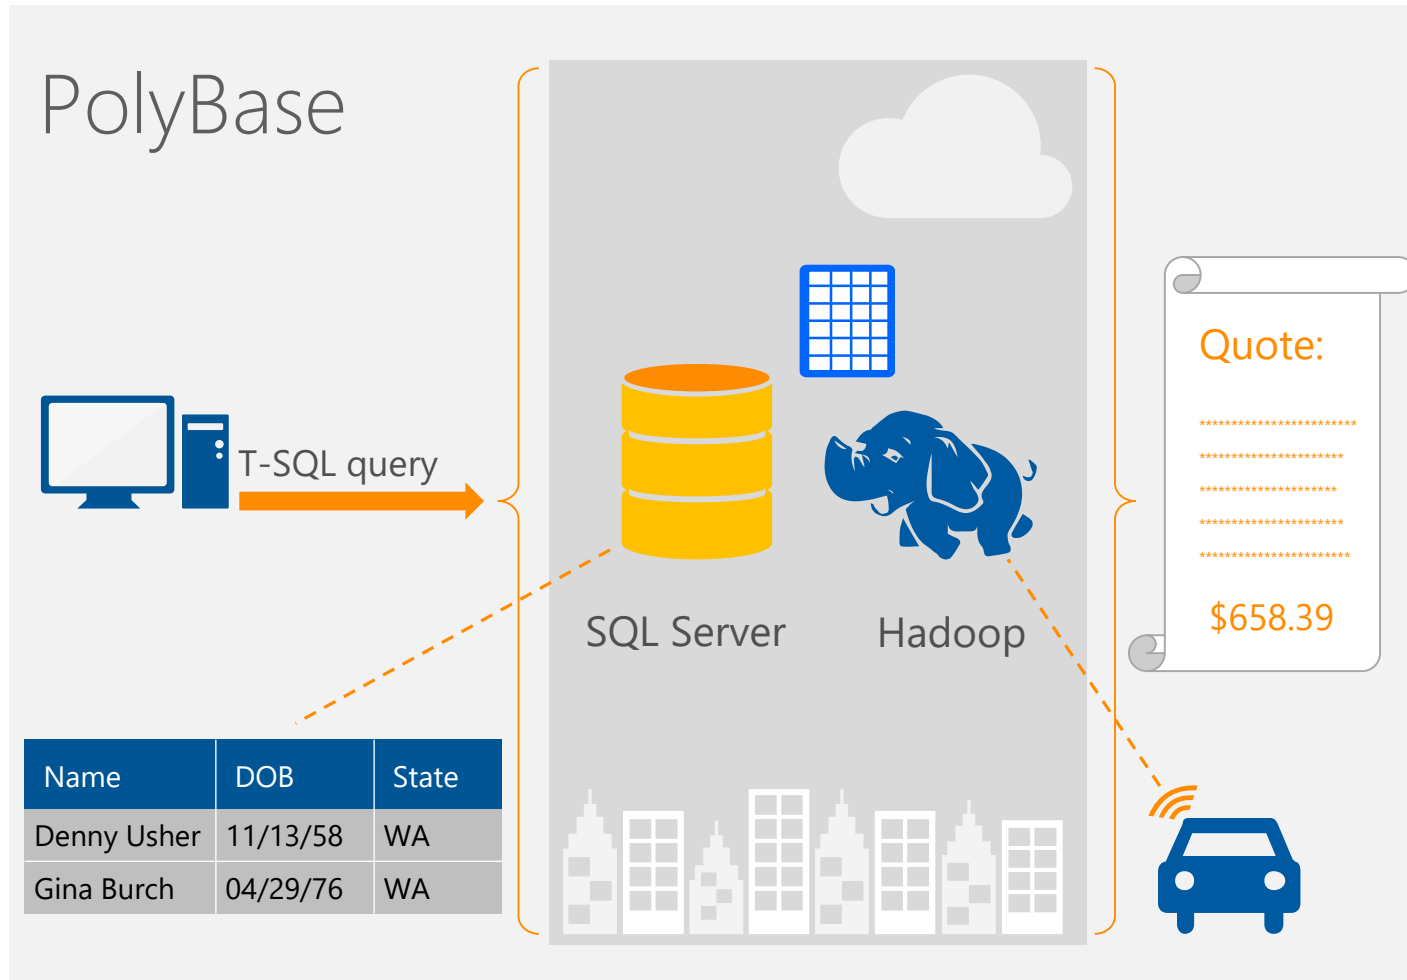
The screenshot displays the Azure portal interface for a ContosoRetailDW SQL Data Warehouse. The interface is divided into several sections:

- Essentials:** Displays key information about the resource group and server.

Property	Value
Resource group	jjjwestusrg
Location	West US
Subscription name	ElasticScaleDev_657854
Server name	jjjwestus.database.windows.net
Status	Online
Connection strings	Show database connection strings
- Monitoring:** Shows a line chart for Query Activity over time, with a summary of 48 successes and 1 failure.
- Settings:** Provides a list of configuration options such as Troubleshoot, Audit logs, Resource health, Locks, and Auditing & Threat detection.
- Auditing & Threat detection:** Shows the status of auditing (ON) and threat detection (OFF), along with options to view server auditing settings and send alerts to email addresses.

Remove the complexity of big data

T-SQL over Hadoop



Leverage existing T-SQL skills and BI tools to gain insights from different data stores

Key Benefits

- Reduce development costs performing query in structured and non-structured data using T-SQL
- Combine SQL with Hadoop enabling database for modern applications
- Simplified architecture with built-in hybrid integration
- Leverage existing T-SQL skills and BI tools to gain insights from different data stores



Data Warehousing and Modern BI on Azure

Demo

<https://gallery.cortanaintelligence.com/Tutorial/Data-Warehousing-and-Modern-BI-on-Azure>

Capture

Curate

Consume

Data Sources

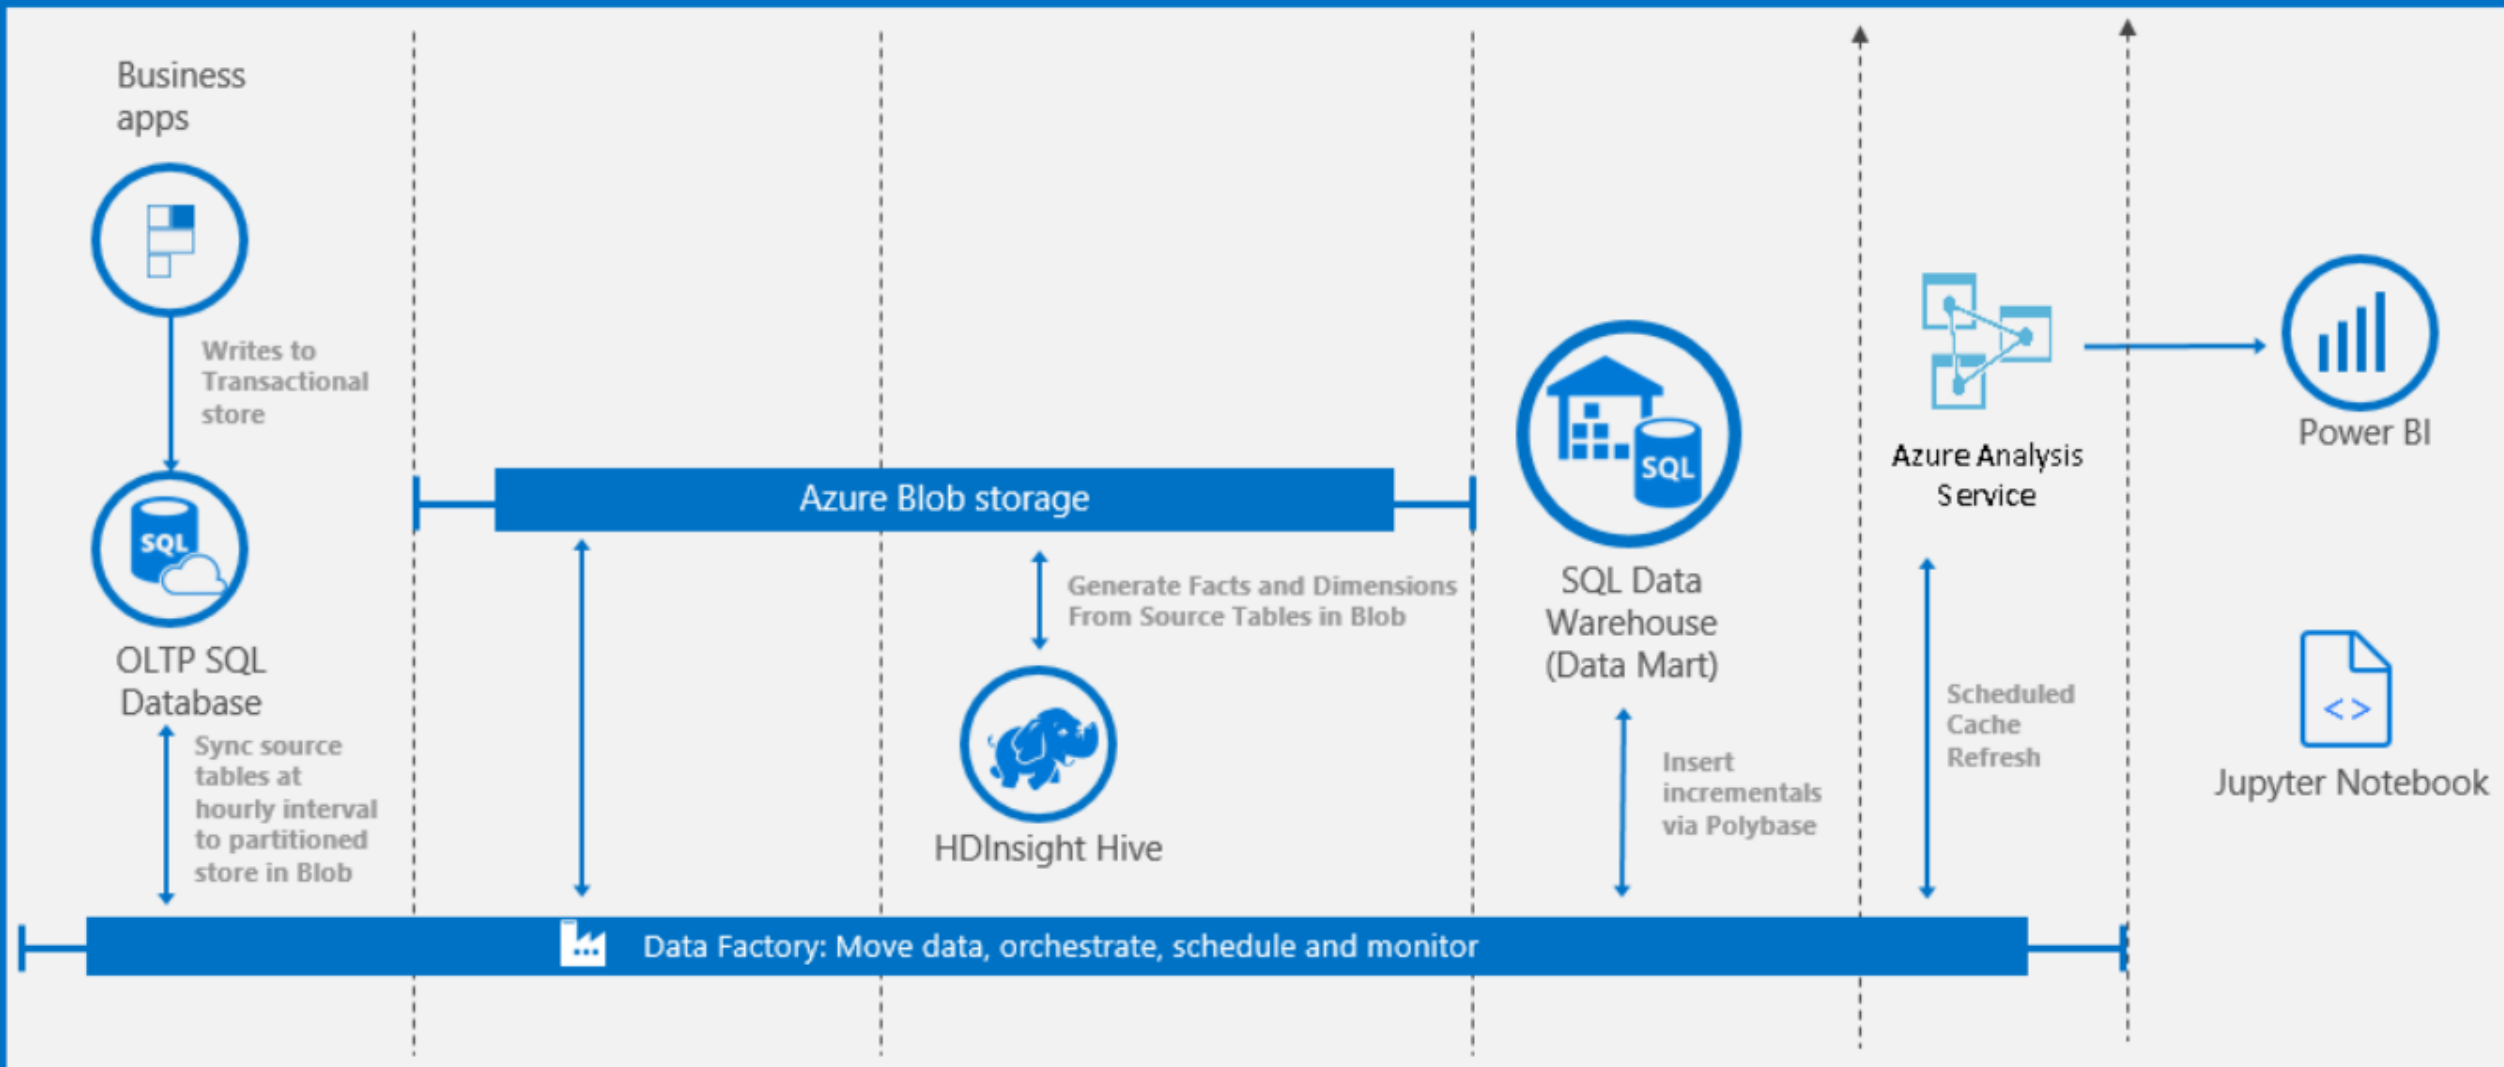
▶ Ingest

▶ Sanitize

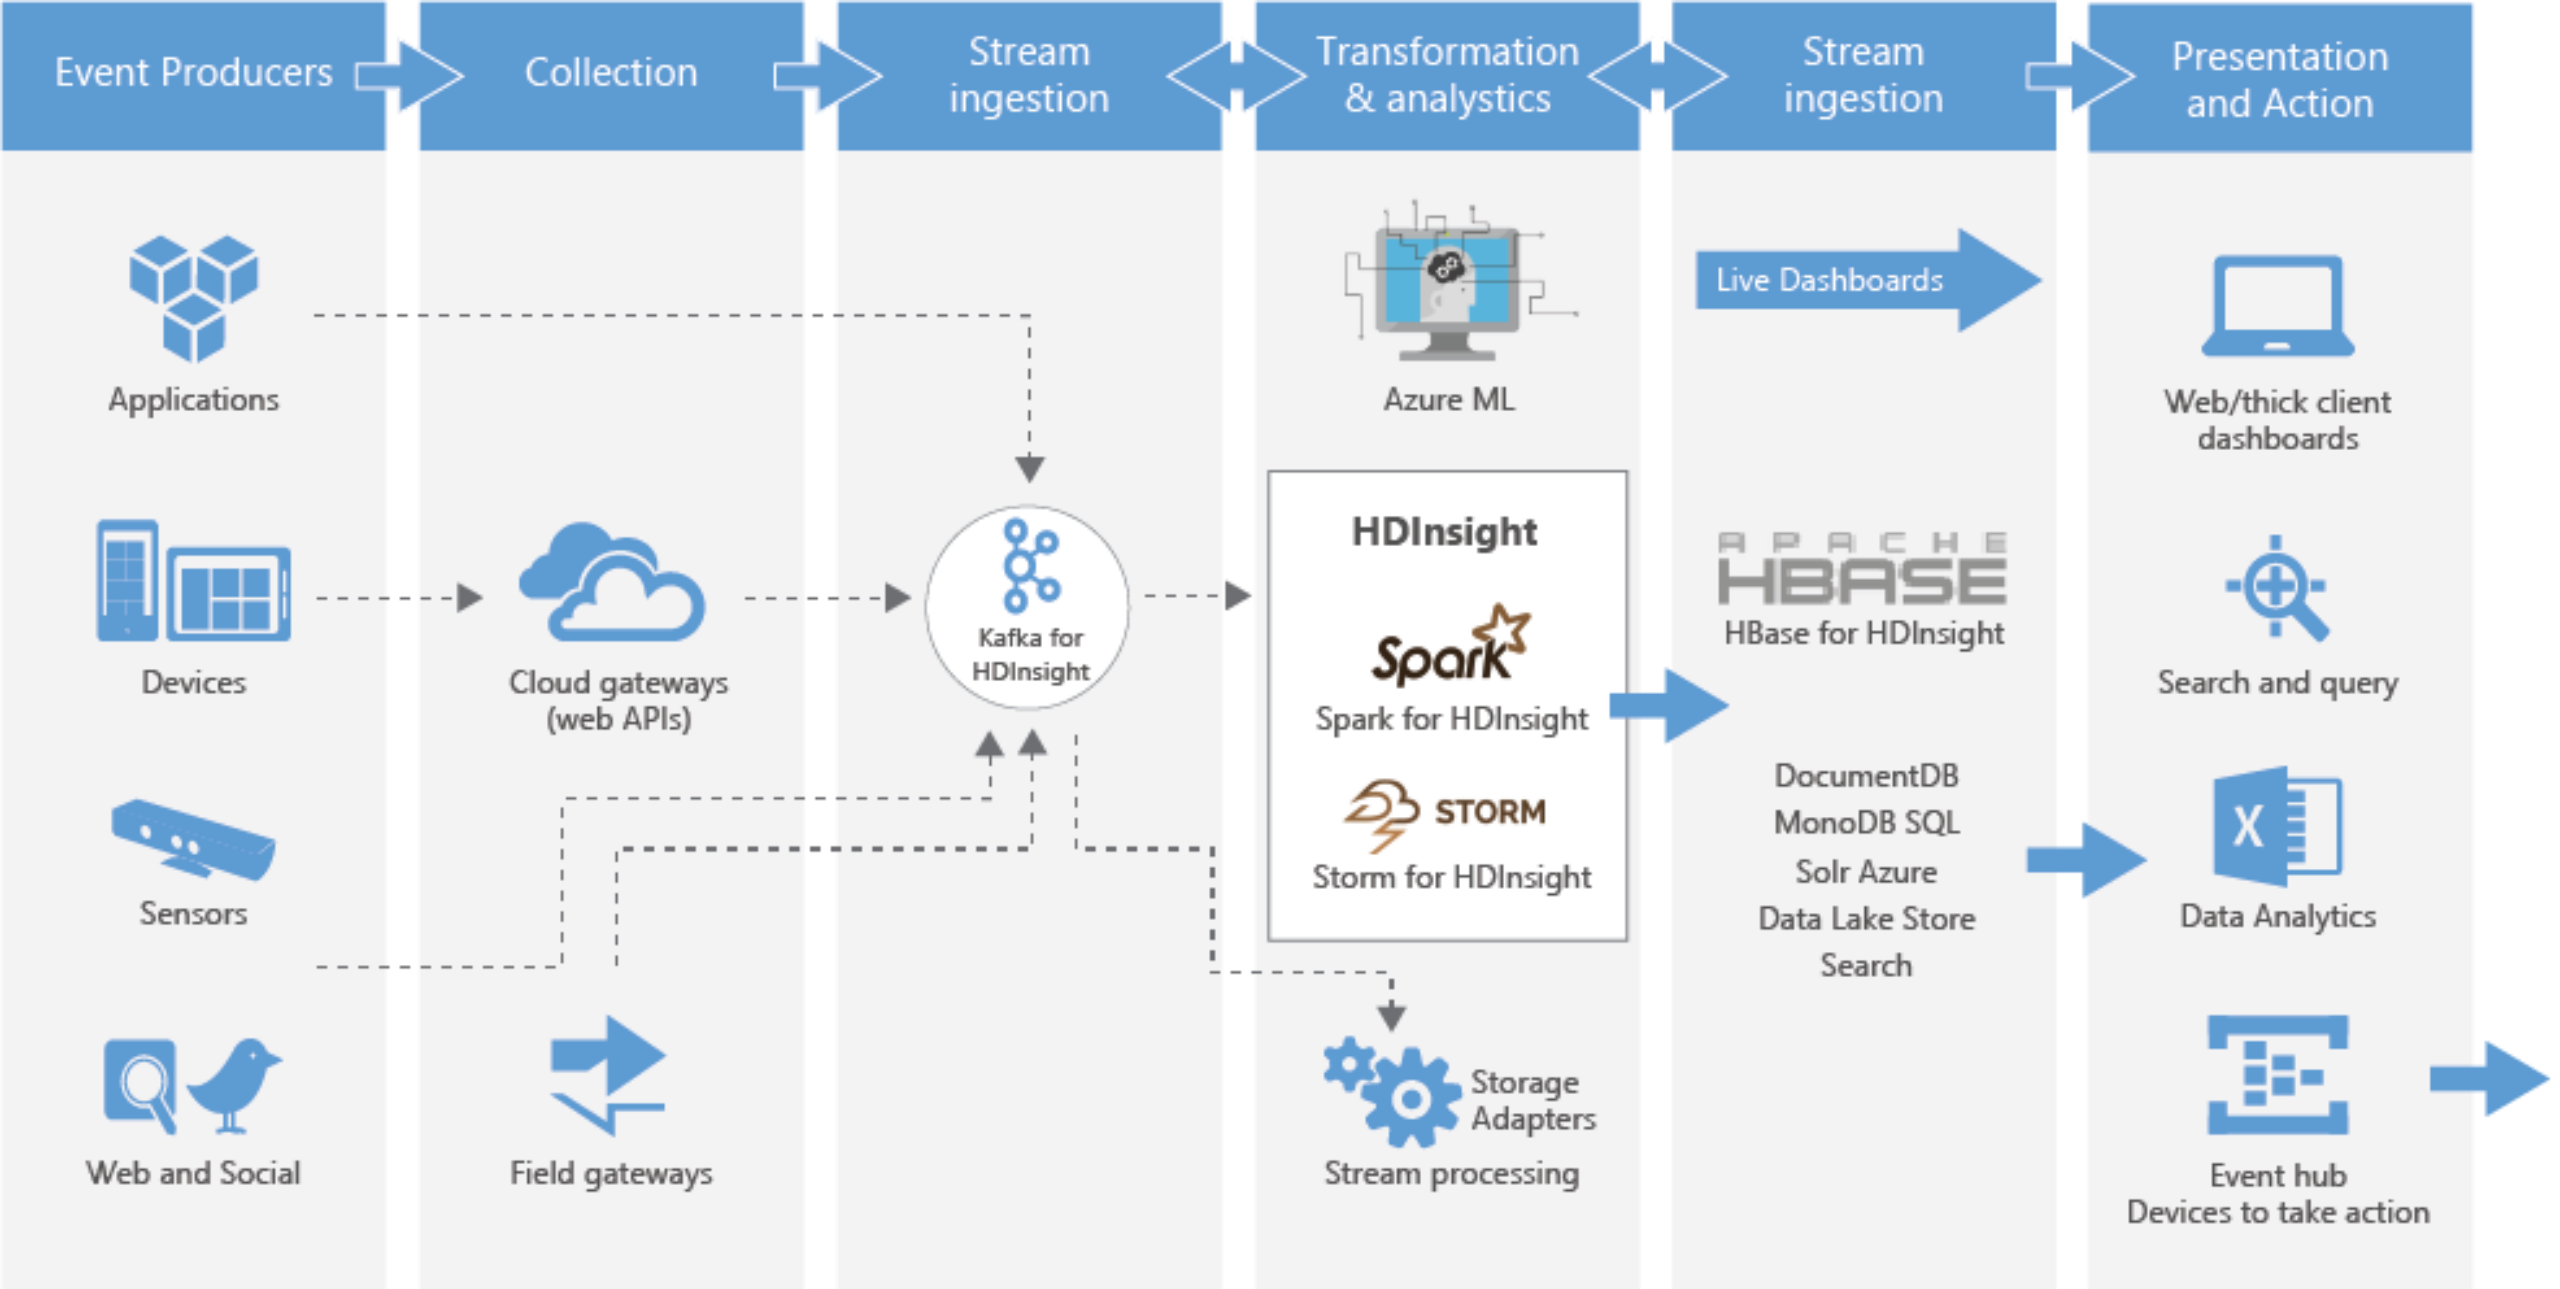
▶ Warehousing

▶ Caching

▶ Visualize



<https://gallery.cortanaintelligence.com/Tutorial/Data-Warehousing-and-Modern-BI-on-Azure>



Power BI: experience your data

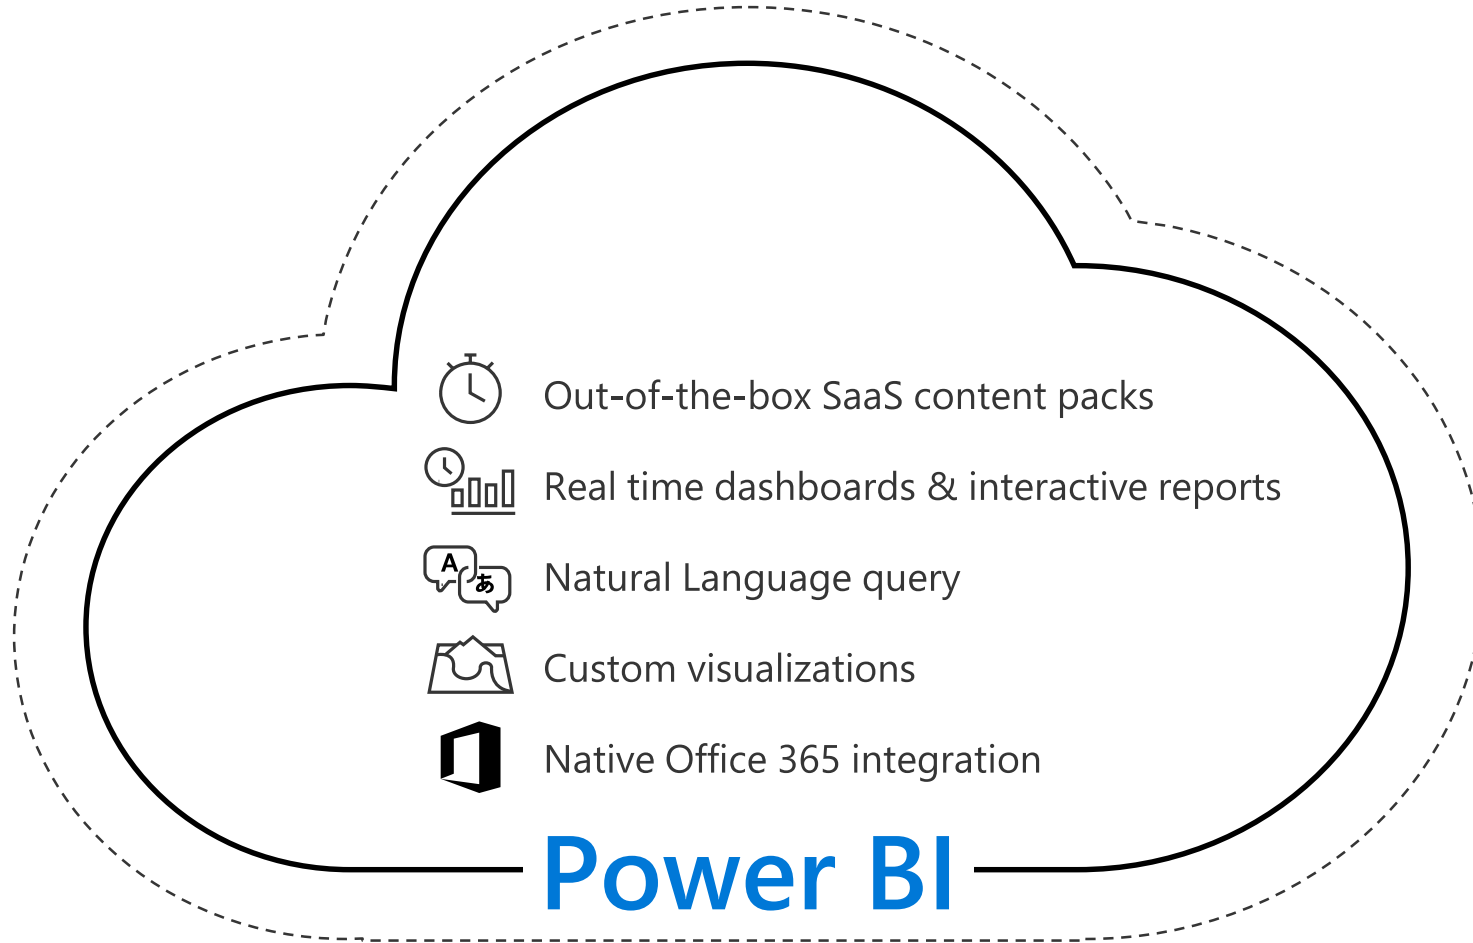
Any data, any way, anywhere



Cloud data

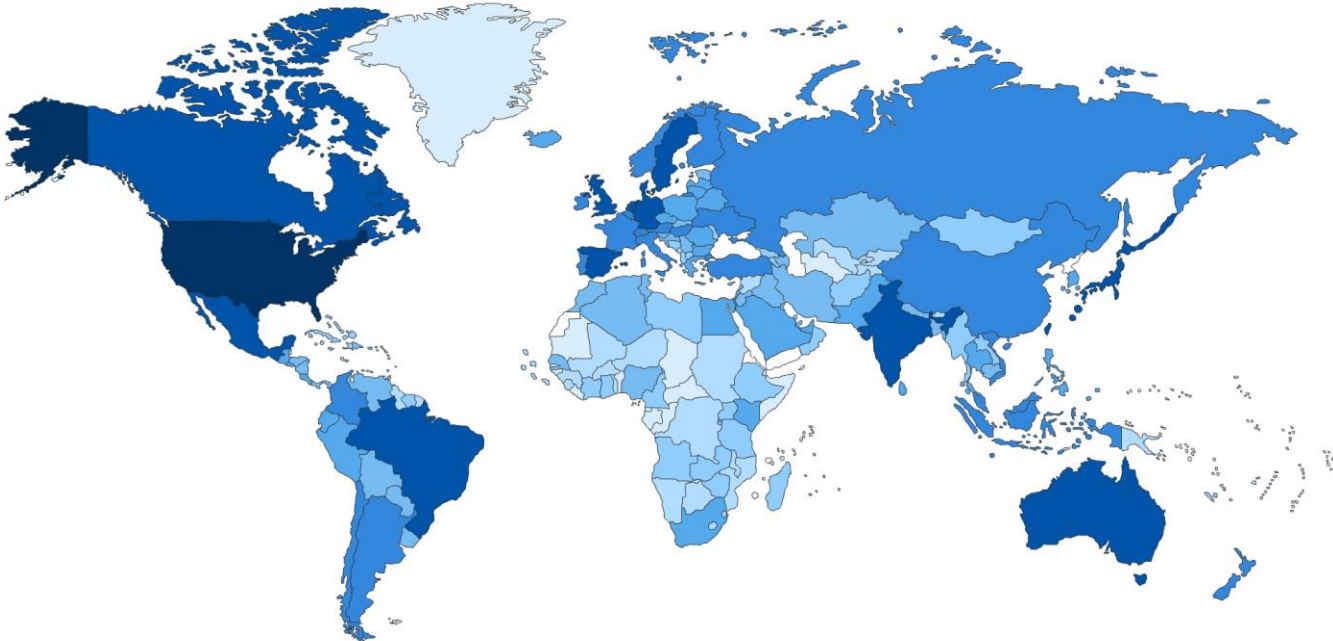


On-premises data



- Web
- Mobile
- Excel
- Embedded
- Cortana
- "....."

Worldwide reach



18,000

Cities globally access the service daily

235

Countries & territories

43

Languages

Global coverage


19 datacenters worldwide

National Clouds

US Government
China
Germany



EU Model Clauses



China GB 18030
RESILIENCE PLUGFEST
JUNE 19-23, 2017

Shaded areas show countries with Power BI usage. The darker the color the greater the number of users



MAY 2017

Editor's choice

The screenshot shows a PCMag.com article for Microsoft Power BI. At the top left is the PCMag logo with 'PCMAG.COM' and 'EDITORS' CHOICE' below it. The article title is 'Microsoft Power BI' with an 'EXCELLENT' editor rating (5 red dots) and '0 COMMENTS' as of 'MAY 26, 2017'. Below the title is a share count of '218 SHARES' and social media icons for Facebook, Twitter, LinkedIn, Pinterest, and Reddit. A navigation bar contains 'PROFILE', '0 COMMENTS', 'SPECS', and 'COMPARE'. Two buttons are visible: a red 'Visit Site' button and a white 'Compare Prices' button with a red border. The main heading 'Power BI' is accompanied by a bar chart icon. The article content is organized into three columns: 'PROS', 'CONS', and 'BOTTOM LINE'. The 'PROS' column lists data source connectors and ease of use. The 'CONS' column notes the refresh cycle on the free version. The 'BOTTOM LINE' column describes the tool as powerful and intuitive.



February 2017

A Leader in the Magic Quadrant for BI & Analytics*


*Gartner "Magic Quadrant for Business Intelligence and Analytics Platforms," by Rita L. Sallam, Cindi Howson, Carlie J. Idoine, Thomas W. Oestreich, James Laurence Richardson, and Joao Tapadinhas February 16, 2017

This graphic was published by Gartner, Inc. as part of a larger research document and should be evaluated in the context of the entire document. The Gartner document is available upon request from Microsoft. Gartner does not endorse any vendor, product or service depicted in its research publications, and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner research publications consist of the opinions of Gartner's research organization and should not be construed as statements of fact. Gartner disclaims all warranties, expressed or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose.



As of February 2017

Power BI Premium



Dedicated resources in the cloud

Cortana Intelligence Solutions

Improve business processes and increase competitiveness

Intelligent

Interactive dashboards that help highlight actionable insights from data

Proven solution architectures that enable transformation of data into insights

Flexible

Deploy and operationalize a complex big data architecture with ease.

Scale and customize the solution to your needs

Trusted

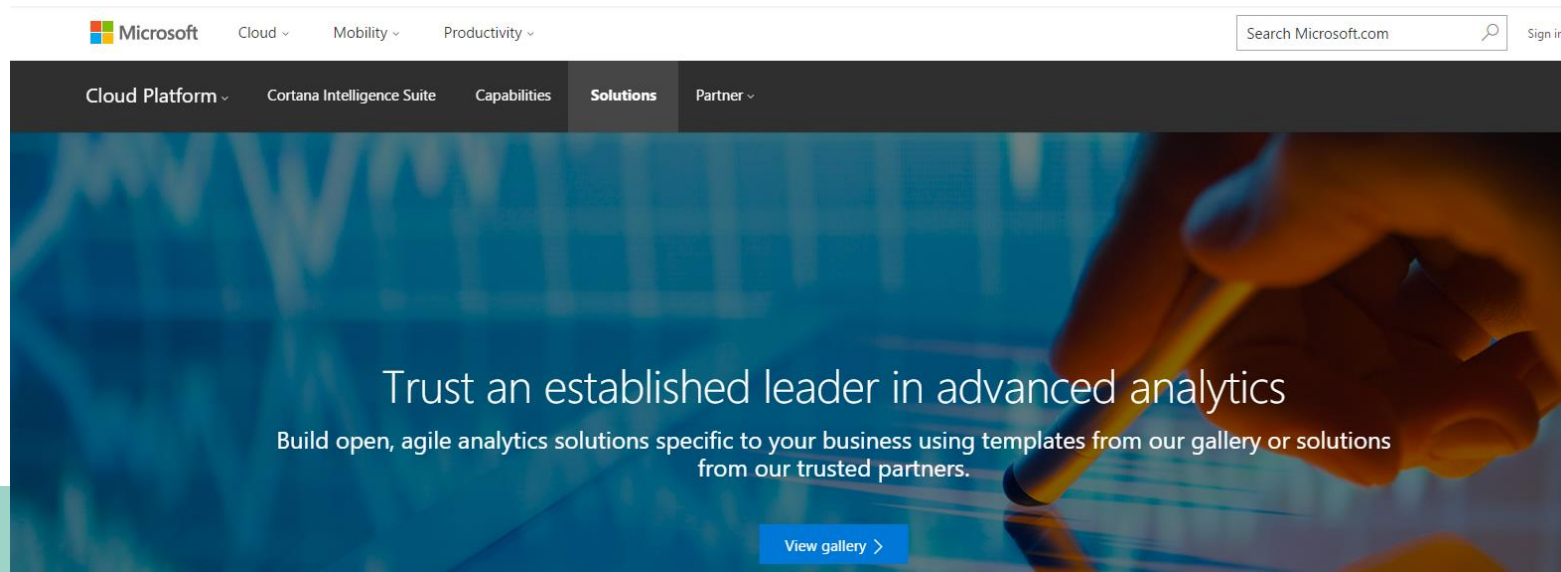
Qualified partners to help with proof-of-concepts and operationalization

Rich technical guides to build an enterprise-grade solution



Cortana Intelligence Solutions

- Solution can be discovered from Microsoft.com
<https://www.microsoft.com/en-us/cloud-platform/cortana-intelligence-suite-solutions>
- Solutions and Tutorials from the Cortana Intelligence Gallery
<https://gallery.cortanaintelligence.com/solutions>
- From aka.ms short urls <https://aka.ms/cisolutions> and <https://aka.ms/citutorials>
- AppSource
<https://appsource.microsoft.com/en-us/marketplace/apps?product=cortana-intelligence&page=1>



Solutions Implementation Partners

Implementation Partners we are training on these solutions



Thank You!

Questions?