

# SQL Server 2016

Everything built-in

Csom Gergely – Microsoft

Adat platform szakértő





SQL Server ❤️ Linux

# SQL Server 2016: Everything built-in

Industry leader in Mission Critical OLTP

**built-in**


## Industry leader



Most secure database

**built-in**


## 6 years in a row least vulnerable



Highest performing data warehouse

**built-in**

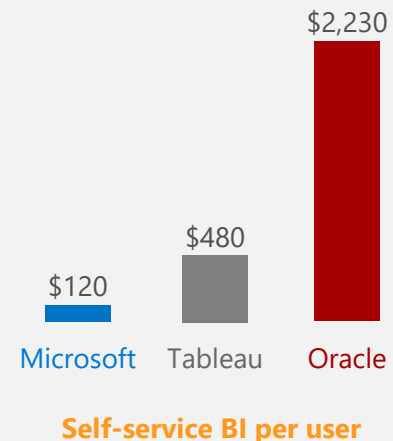
## #1 performance



End-to-end mobile BI on any device

**built-in**


## A fraction of the cost



In-database Advanced Analytics

**built-in**

## R + in-memory



at massive scale

In-memory across all workloads



Consistent experience from on-premises to cloud



The above graphics were 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.

National Institute of Standards and Technology Comprehensive Vulnerability Database update 10/2015

TPC-H 10TB non-clustered results as of 04/06/15, 5/04/15, 4/15/14 and 11/25/13, respectively. [http://www.tpc.org/tpch/results/tpch\\_perf\\_results.asp?resulttype=noncluster](http://www.tpc.org/tpch/results/tpch_perf_results.asp?resulttype=noncluster)

# Microsoft is a leader for...

## Magic Quadrant for Operational Database Management Systems<sup>1</sup> Furthest in vision and ability to execute



## Magic Quadrant for Business Intelligence and Analytics Platforms<sup>2</sup> Furthest in vision; leader 9 years running



## Magic Quadrant for Data Warehouse Database Management Solutions<sup>3</sup> A leader for the fifth consecutive year




[1] \*Gartner "Magic Quadrant for Operational Database Management Systems," by Donald Feinberg , Merv Adrian , Nick Heudecker, Adam Ronthal, October 2015

[2] \*Gartner "Magic Quadrant for Business Intelligence and Analytics Platforms," by Josh Parenteau, Rita L. Sallam, Cindi Howson, Joao Tapadinhas, Kurt Schlegel, Thomas W. Oestreich, February 4, 2016

[3] \*Gartner "Magic Quadrant for Data Warehouse and Data Management Solutions for Analytics," by Roxane Edjlali and Mark Beyer, February 25, 2016

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.

# How we develop SQL



Cloud-first but not cloud-only

Use SQL Database to improve core SQL Server features and cadence

Many interesting and compelling on-premises ↔ cloud scenarios

# Only Microsoft delivers On-premises & cloud

## Consistent experience is everything



SQL Server



SQL Server in Azure VM



Azure SQL Database



Analytics Platform System



Azure SQL Data Warehouse

### Common tools

Dev ops tools

Management tools

Identity

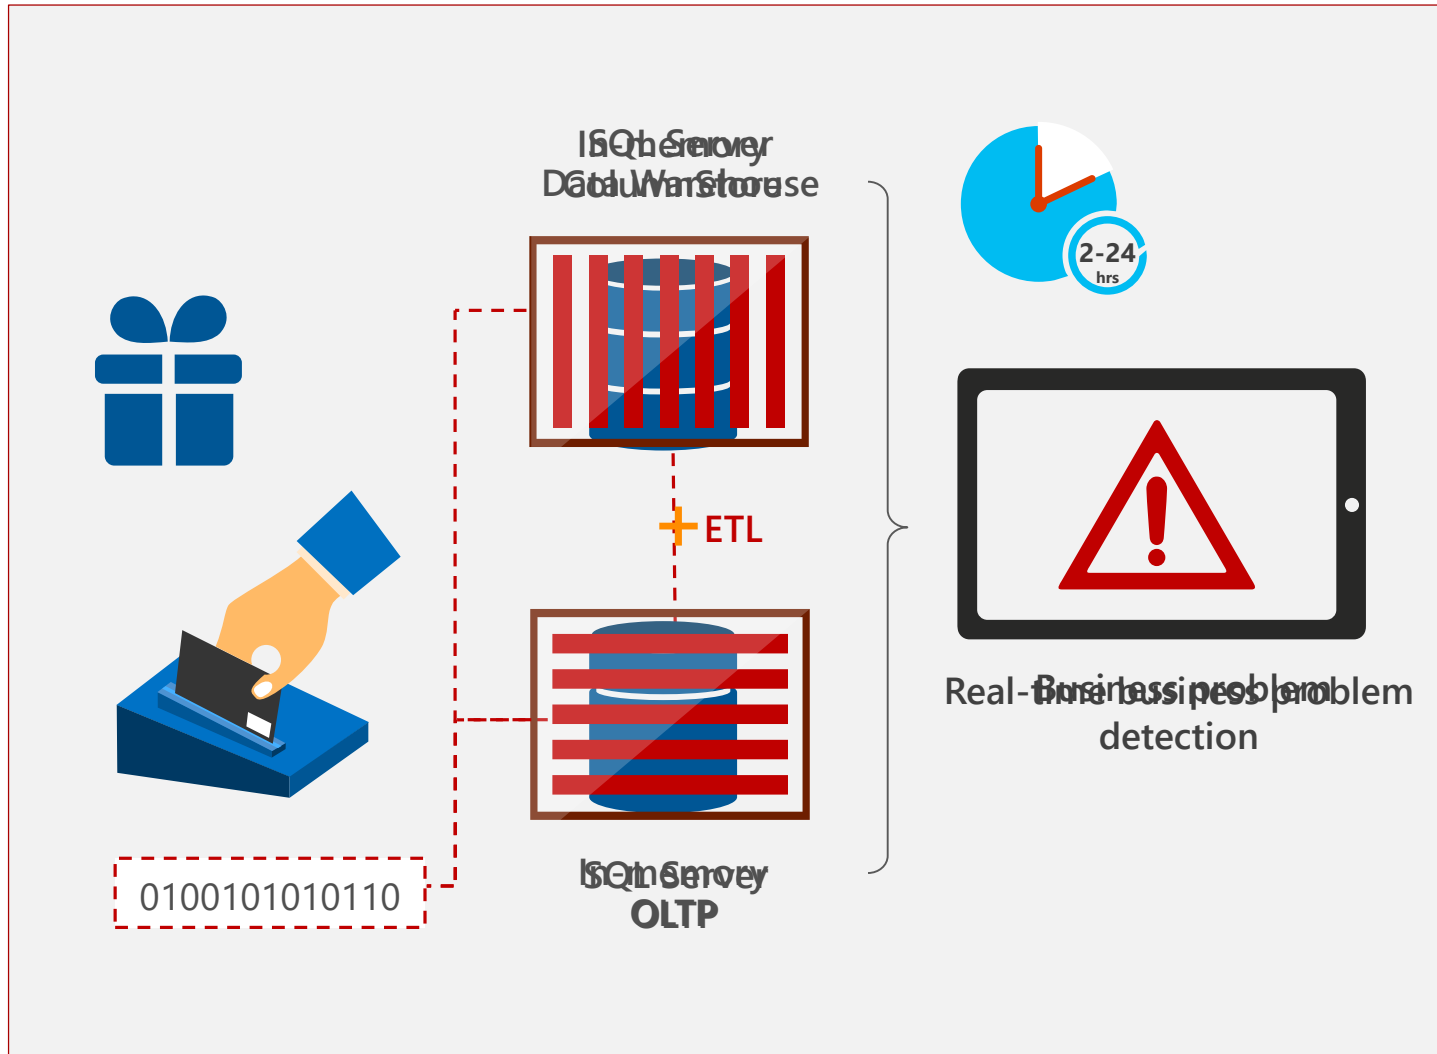
Virtualization

Single-vendor support

- ➔ **Common tools** for development and management
- ➔ **Common T-SQL** surface area
- ➔ **Simple** cloud migration

# Real-time operational analytics

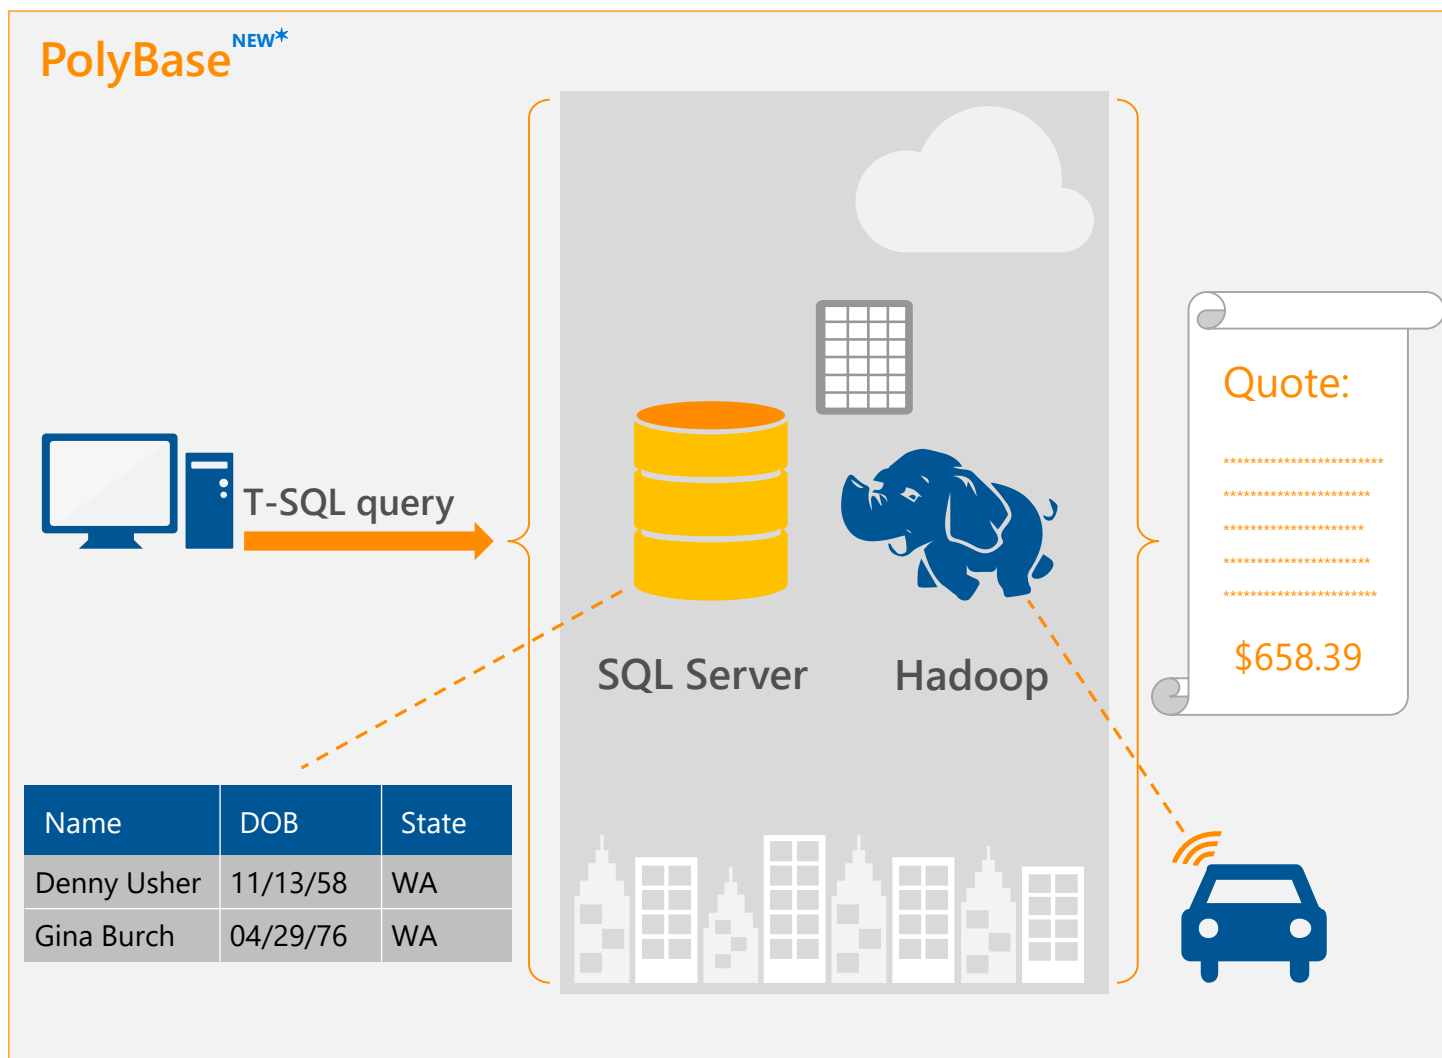
In-memory **built-in**



- ➔ Up to **30x** faster transactions with in-memory OLTP
- ➔ Queries from **minutes to seconds**
- ➔ **Real-time**<sup>NEW\*</sup> operational analytics

# Remove the complexity of big data

## T-SQL over Hadoop



Manage structured & unstructured data

- ➔ **Simple T-SQL** NEW\*  
to query Hadoop data (HDFS)
- ➔ **JSON support** NEW\*



# PolyBase use cases

## Load data

Use Hadoop as an ETL tool to cleanse data before loading to data warehouse with PolyBase

## Interactively query

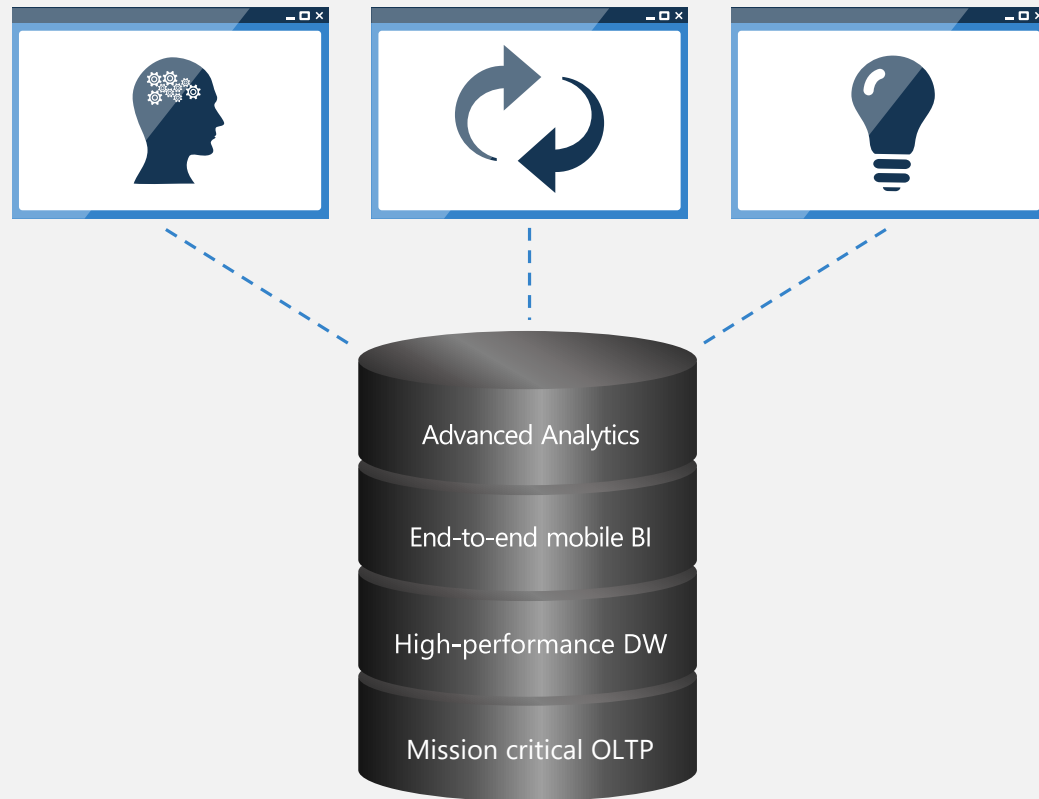
Analyze relational data with semi-structured data using split-based query processing

## Age out data

Age out data to HDFS and use it as "cold" but queryable storage

# In-database Advanced Analytics

Build intelligent applications with SQL Server R Services



R built-in to SQL Server

- ➔ **R built-in to your T-SQL** <sup>NEW\*</sup>
- ➔ **Real-time operational analytics** <sup>NEW\*</sup> without moving the data
- ➔ **Open Source R with in-memory & massive scale** <sup>NEW\*</sup> - multi-threading and massive parallel processing

# Revolution Analytics product integration



## Community



Revolution R Open



Microsoft R Open



## Commercial



Revolution R Enterprise

SQL Server  
R Services

Windows



Microsoft  
R Server

Red Hat

SUSE

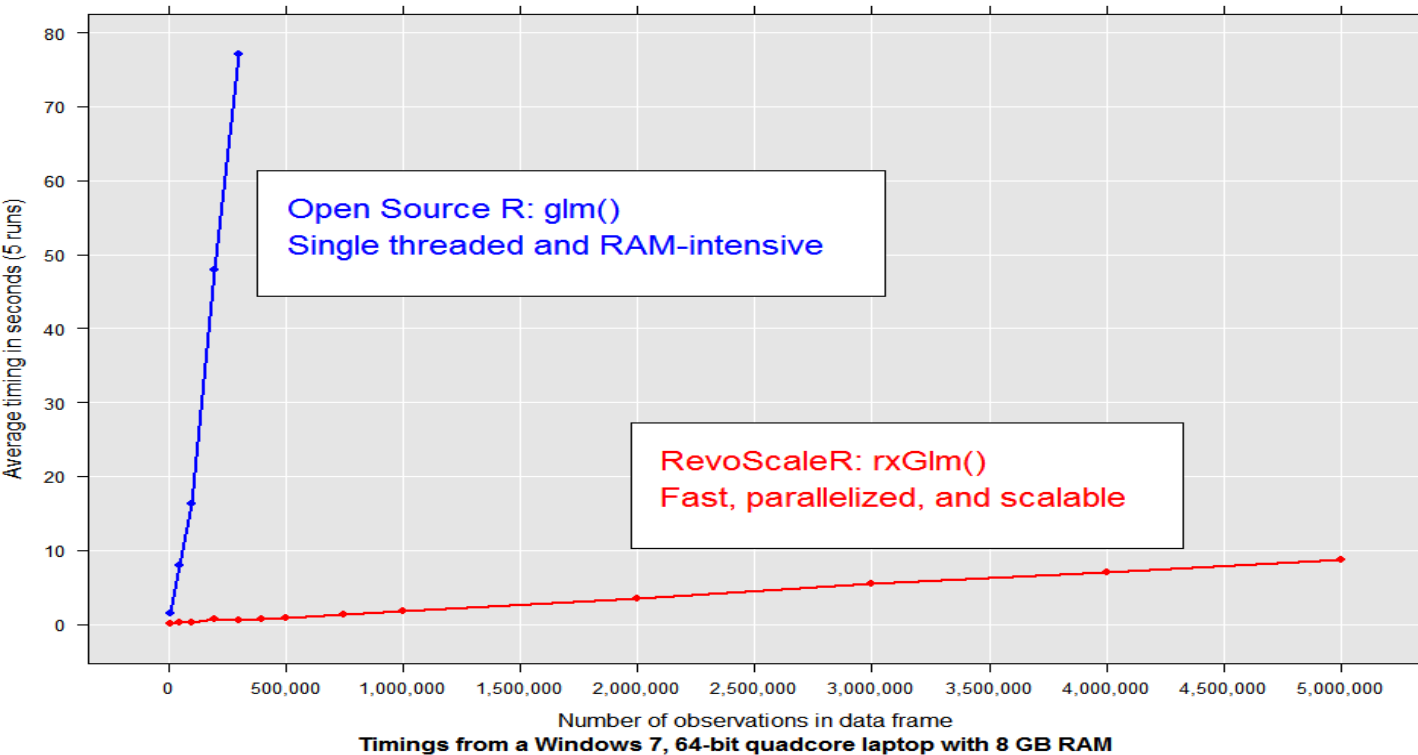
Hadoop

Teradata

# ScaleR - Performance comparison

Microsoft R Server has no data size limits in relation to size of available RAM. When open source R operates on data sets that exceed RAM it will fail. In contrast Microsoft R Server scales linearly well beyond RAM limits and parallel algorithms are much faster.

**GLM 'Gamma' Simulation Timings**  
Independent Variables: 2 factors (100 and 20 levels) and one continuous



File Name	Compressed File Size (MB)	No. Rows	Open Source R (secs)	Revolution R (secs)
Tiny	0.3	1,235	0.00	0.05
V. Small	0.4	12,353	0.21	0.05
Small	1.3	123,534	0.03	0.03
Medium	10.7	1,235,349	1.94	0.08
Large	104.5	12,353,496	60.69	0.42
Big (full)	12,960.0	123,534,969	Memory!	4.89
V.Big	25,919.7	247,069,938	Memory!	9.49
Huge	51,840.2	494,139,876	Memory!	18.92

- US flight data for 20 years
- Linear Regression on Arrival Delay
- Run on 4 core laptop, 16GB RAM and 500GB SSD

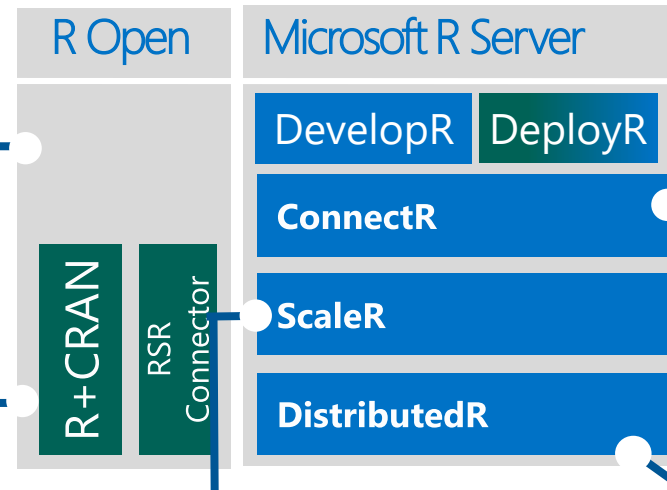
# The Microsoft R Server Platform

## R+CRAN

- Open source R interpreter
  - R 3.2.3
- Freely-available huge range of R algorithms
- Algorithms callable by RevoR
- Embeddable in R scripts
- 100% Compatible with existing R scripts, functions and packages

## RevoR

- Performance enhanced R interpreter
- Based on open source R
- Adds high-performance math library to speed up linear algebra functions



## ConnectR

- High-speed & direct connectors

### Available for:

- High-performance XDF
- SAS, SPSS, delimited & fixed format text data files
- Hadoop HDFS (text & XDF)
- Teradata Database & Aster
- EDWs and ADWs
- ODBC





## ScaleR

- Ready-to-Use high-performance big data big analytics
- Fully-parallelized analytics
- Data prep & data distillation
- Descriptive statistics & statistical tests
- Range of predictive functions
- User tools for distributing customized R algorithms across nodes
- Wide data sets supported – thousands of variables

## DistributedR

- Distributed computing framework
- Delivers cross-platform portability

# SQL Server R Services Samples

 This repository  [Pull requests](#) [Issues](#) [Gist](#)   


[Microsoft](#) / **SQL-Server-R-Services-Samples** [Watch](#) 29 [Star](#) 52 [Fork](#) 45

[Code](#) [Issues](#) 0 [Pull requests](#) 0 [Wiki](#) [Pulse](#) [Graphs](#)

Advanced analytics samples and templates using SQL Server R Services

[39 commits](#) [1 branch](#) [0 releases](#) [5 contributors](#)

[Branch: master](#) [New pull request](#) [Create new file](#) [Upload files](#) [Find file](#) [Clone or download](#)

 **hlums** committed on **GitHub** Update README.md Latest commit e8fc09a 5 days ago

<a href="#">Churn</a>	fixed bug in predictRx	12 days ago
<a href="#">EnergyDemandForecasting</a>	Update README.md	5 days ago
<a href="#">FraudDetection</a>	Adding sql script to create table for online Scoring	3 months ago
<a href="#">PerfTuning</a>	Added MSDN link for the Perf Tuning documentation	14 days ago
<a href="#">PredictiveMaintenance</a>	Invoke stored procedure with Powershell instead of SQL script in a file	3 months ago
<a href="#">RetailForecasting</a>	Implementations of Retail Forecasting machine learning template with ...	a month ago
<a href="#">README.md</a>	Update README.md	5 days ago

# End-to-end mobile BI on any device

## Lightning fast queries & reports

- ➔ **In-memory** **built-in**
- ➔ **Reduce time to insight** with direct query
- ➔ **Powerful modeling** support for 50+ new DAX functions



Reports in **minutes** not days

**PB scale DW** in SQL Server

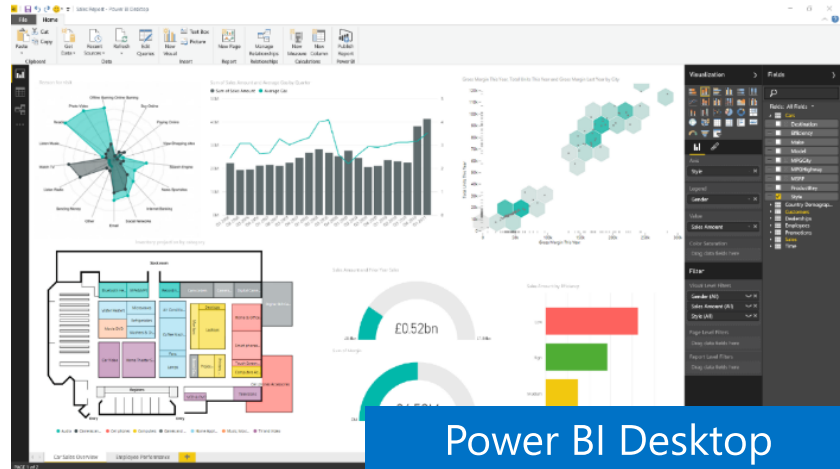
## On all mobile devices

- ➔ **Mobile BI** **built-in** <sup>NEW\*</sup>
- ➔ **Online & offline** access <sup>NEW\*</sup>
- ➔ **Rich visualizations** using Power BI or enhanced Reporting Services <sup>NEW\*</sup>

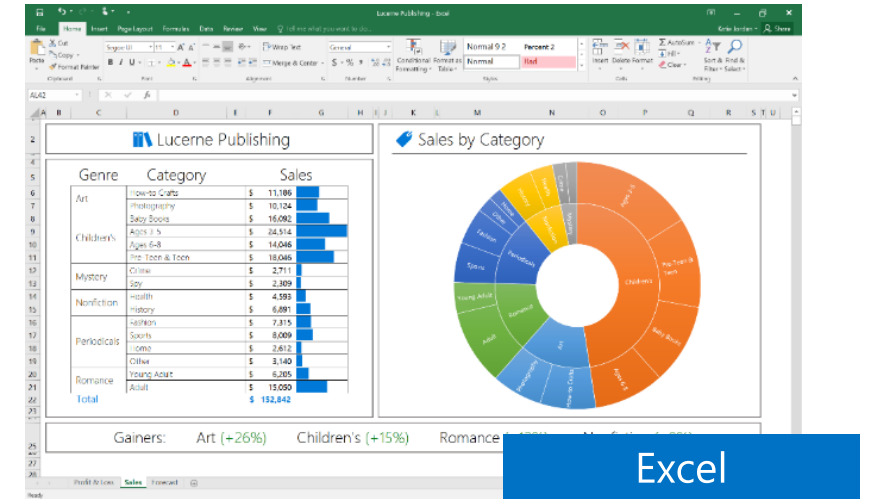


# Vision for Microsoft BI Reporting

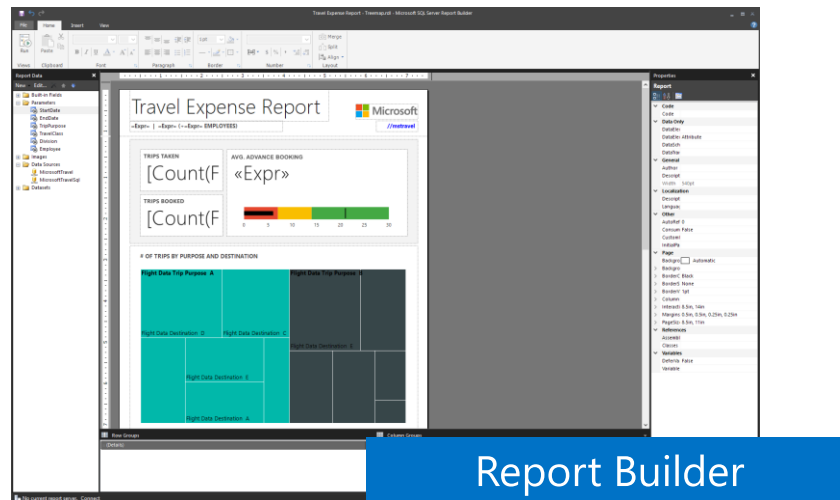
## Interactive Reports



## Spreadsheets



## Paginated Reports



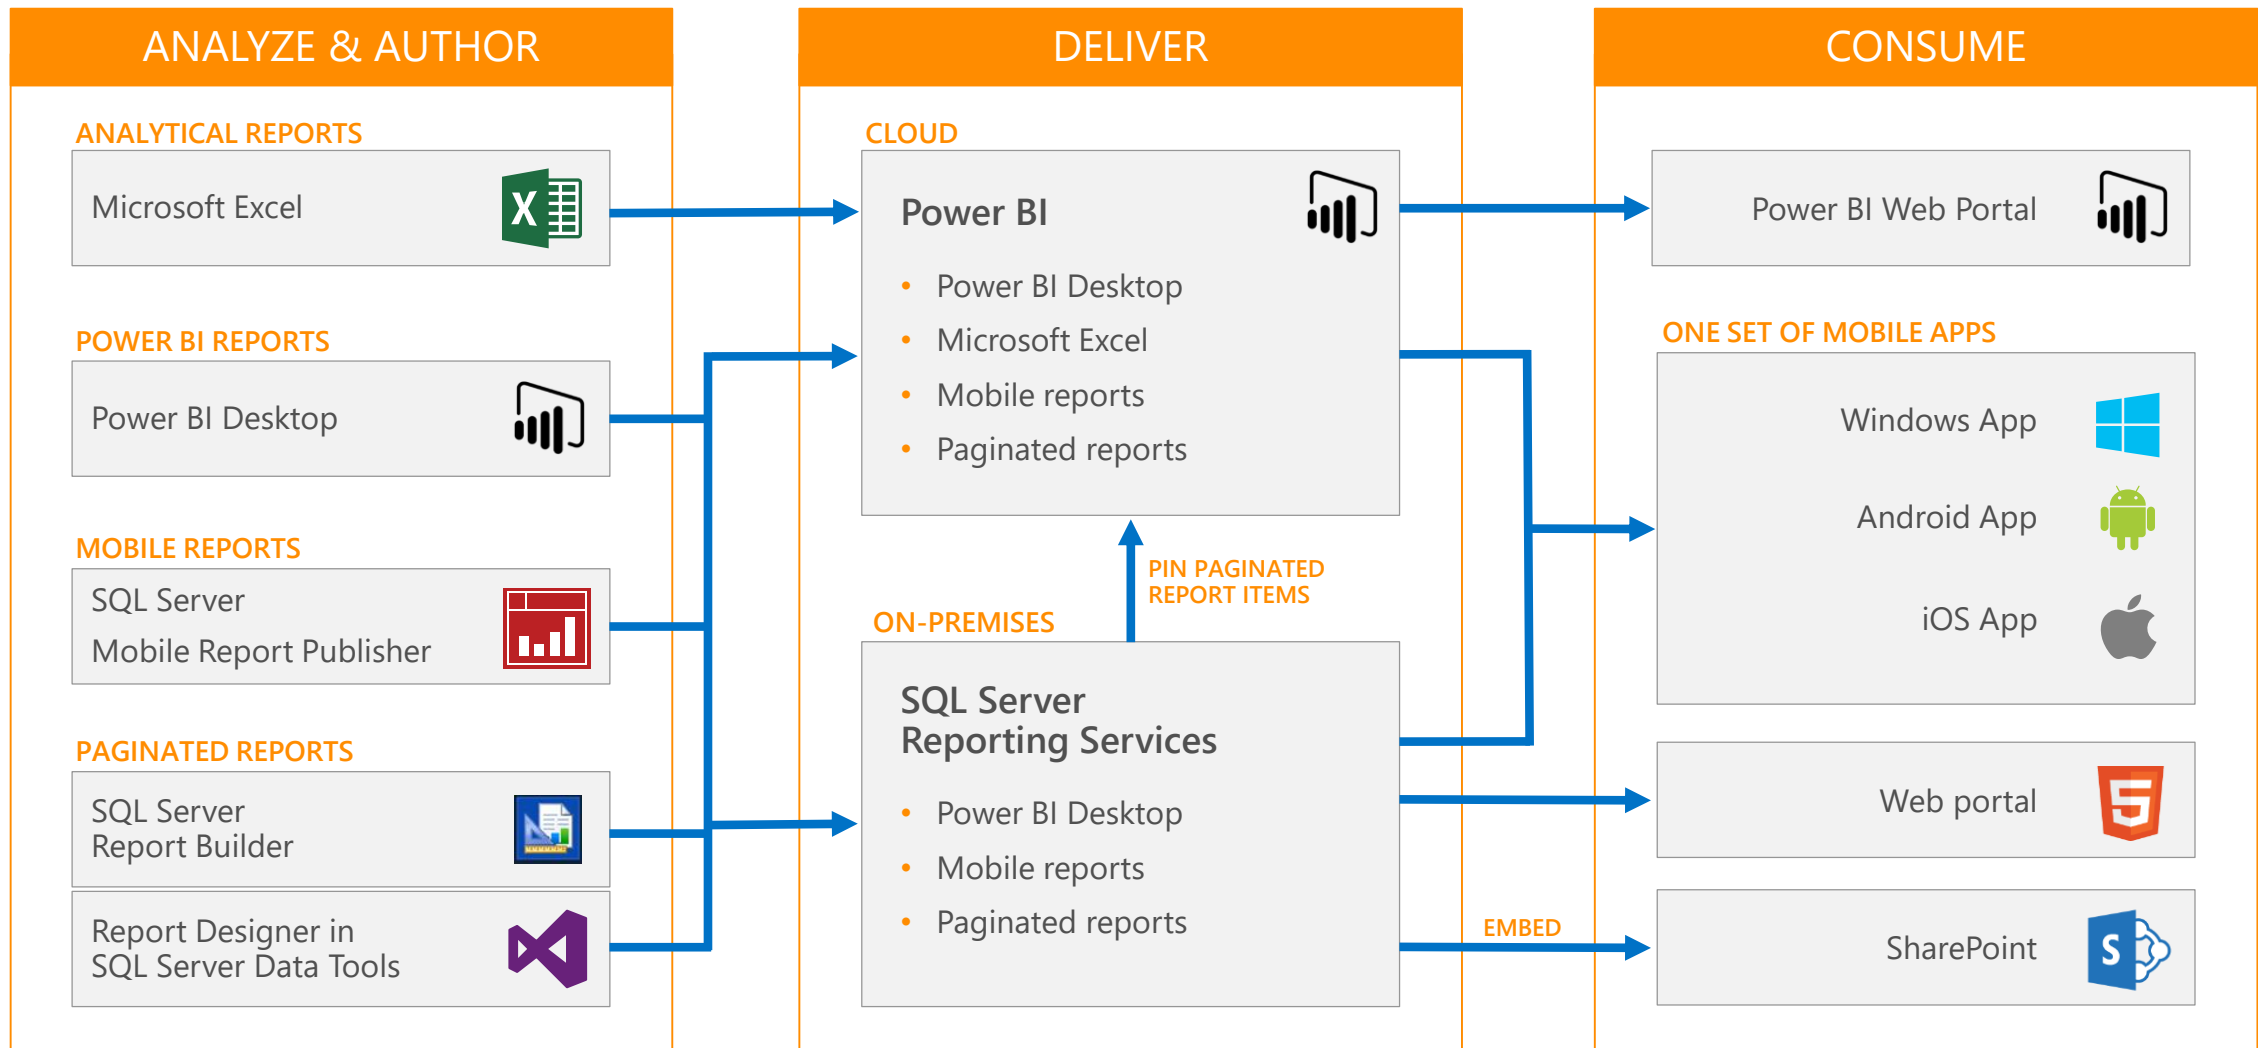
## Mobile Reports





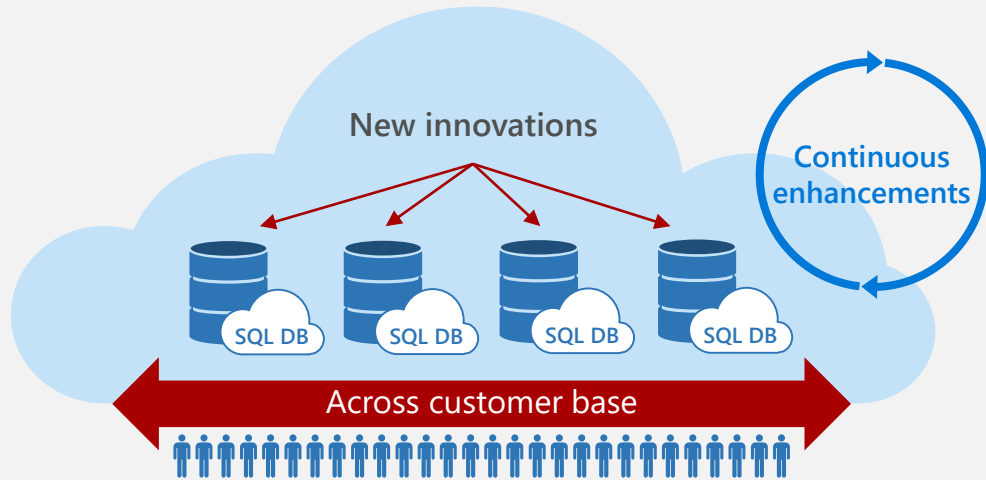
# Future vision for SQL Server BI and Power BI

## Beyond SQL Server 2016



# Continuous Innovation

## Cloud-First Approach



Speed



Proven



Agility



Feedback

## SQL Server 2016



**The best**  
SQL Server  
release in history

# Tovább lépés

- Free Ebooks:
  - Introducing Microsoft Power BI: [https://blogs.msdn.microsoft.com/microsoft\\_press/2016/06/16/free-ebook-introducing-microsoft-power-bi/](https://blogs.msdn.microsoft.com/microsoft_press/2016/06/16/free-ebook-introducing-microsoft-power-bi/)
  - Introducing Microsoft SQL Server 2016: <https://info.microsoft.com/Introducing-SQL-Server-2016-eBook.html>
- SQL Server Data Driven: Technical Deep Dive: [https://www.youtube.com/playlist?list=PL8nfc9haGeb6T3HaGhQWvBz1AqS9d6Zv\\_](https://www.youtube.com/playlist?list=PL8nfc9haGeb6T3HaGhQWvBz1AqS9d6Zv_)
- EDX – Free ondemand courses:
  - Microsoft kurzusok: <https://www.edx.org/school/microsoft>
    - Data Science and Machine Learning Essentials:
    - Introduction to Data Analysis using Excel
    - Analyzing and Visualizing Data with Power BI
    - Implementing Real-Time Analytics with Hadoop in Azure HDInsight
    - Introduction to R Programming:
    - Introduction to Python for Data Science
    - Implementing Predictive Analytics with Hadoop in Azure HDInsight
    - Processing Big Data with Azure HDInsight
- R Tools for Visual Studio: <https://www.visualstudio.com/en-us/features/rtvs-vs.aspx>
- Free R Client, which can be used to connect a desktop R environment to R Server environments: <http://aka.ms/rclient>
- Samples
  - SQL Server 2016 R Services Tutorials: <https://msdn.microsoft.com/en-us/library/mt591993.aspx>
  - R, Microsoft R Open és Microsoft R Server samples: <https://github.com/Microsoft/RTVS-docs/tree/master/examples>
  - MSSQL R Services Samples: <https://github.com/Microsoft/SQL-Server-R-Services-Samples>
  - Using Microsoft R Server on a Single Machine for Experiments With 600M Taxi Rides: <https://blogs.technet.microsoft.com/machinelearning/2016/06/16/using-microsoft-r-server-on-a-single-machine-for-experiments-with-600m-taxi-rides/>

