

High Performance Analytics Platform on Azure With EXASOL



Carsten
Weidmann

- **Intelligent, high performance MPP In-Memory Database for Analytics**
 - MPP, shared-nothing
 - ACID compliant, fully transactional, but focused on analytics
 - intelligent in-memory column store
 - runs on standard x86 hardware
 - standard SQL
 - parallelized in-database analytics with R, Python, Java....
 - extremely scalable, supports high concurrency
 - minimal administration overhead, self-tuning capabilities

The World's Fastest Analytic Database That Just Works!

100% German Engineering

EXASOL



- 15+ years of experience
- Privately funded, growing and profitable German company
- Scientific roots: massively parallel processing
- We work to the highest standards



 made
in
Germany

Let Me Just Prepare Something...

Setting Up EXASOL in Azure

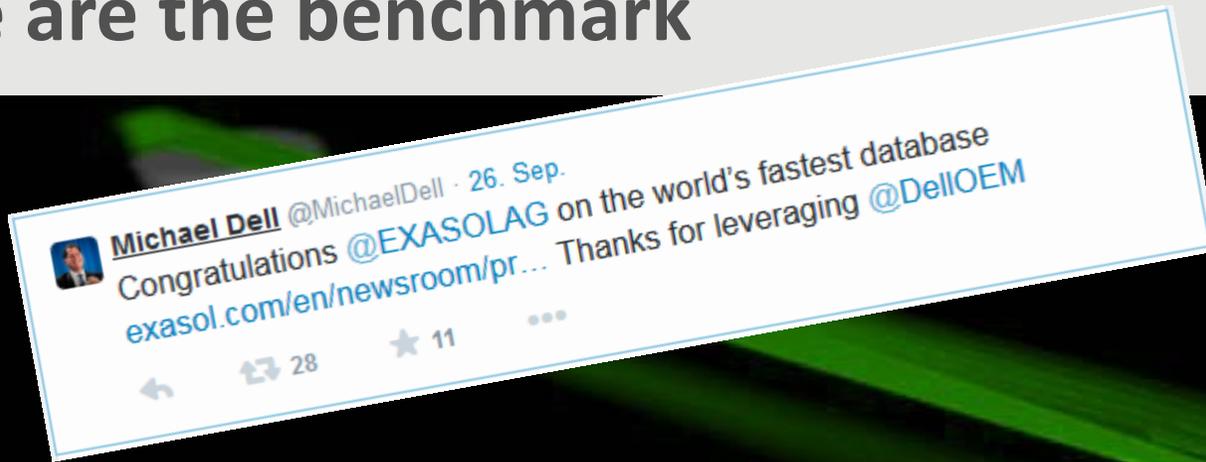


What Gartner says about EXASOL



“EXASOL is a prime example of what Gartner considers to be the future of DBMS”

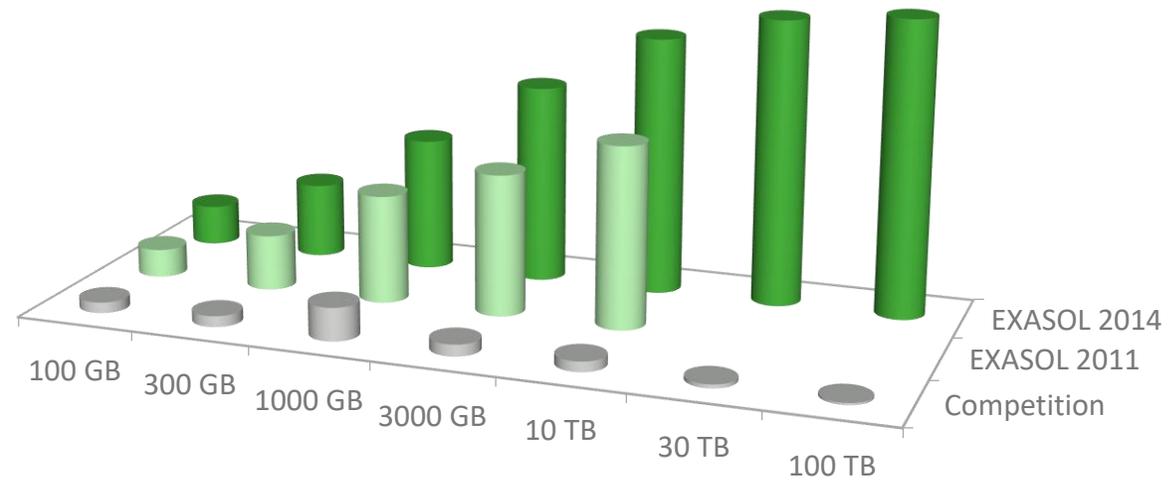
We are the benchmark



- TPC-H is the industry standard benchmark for analytical databases



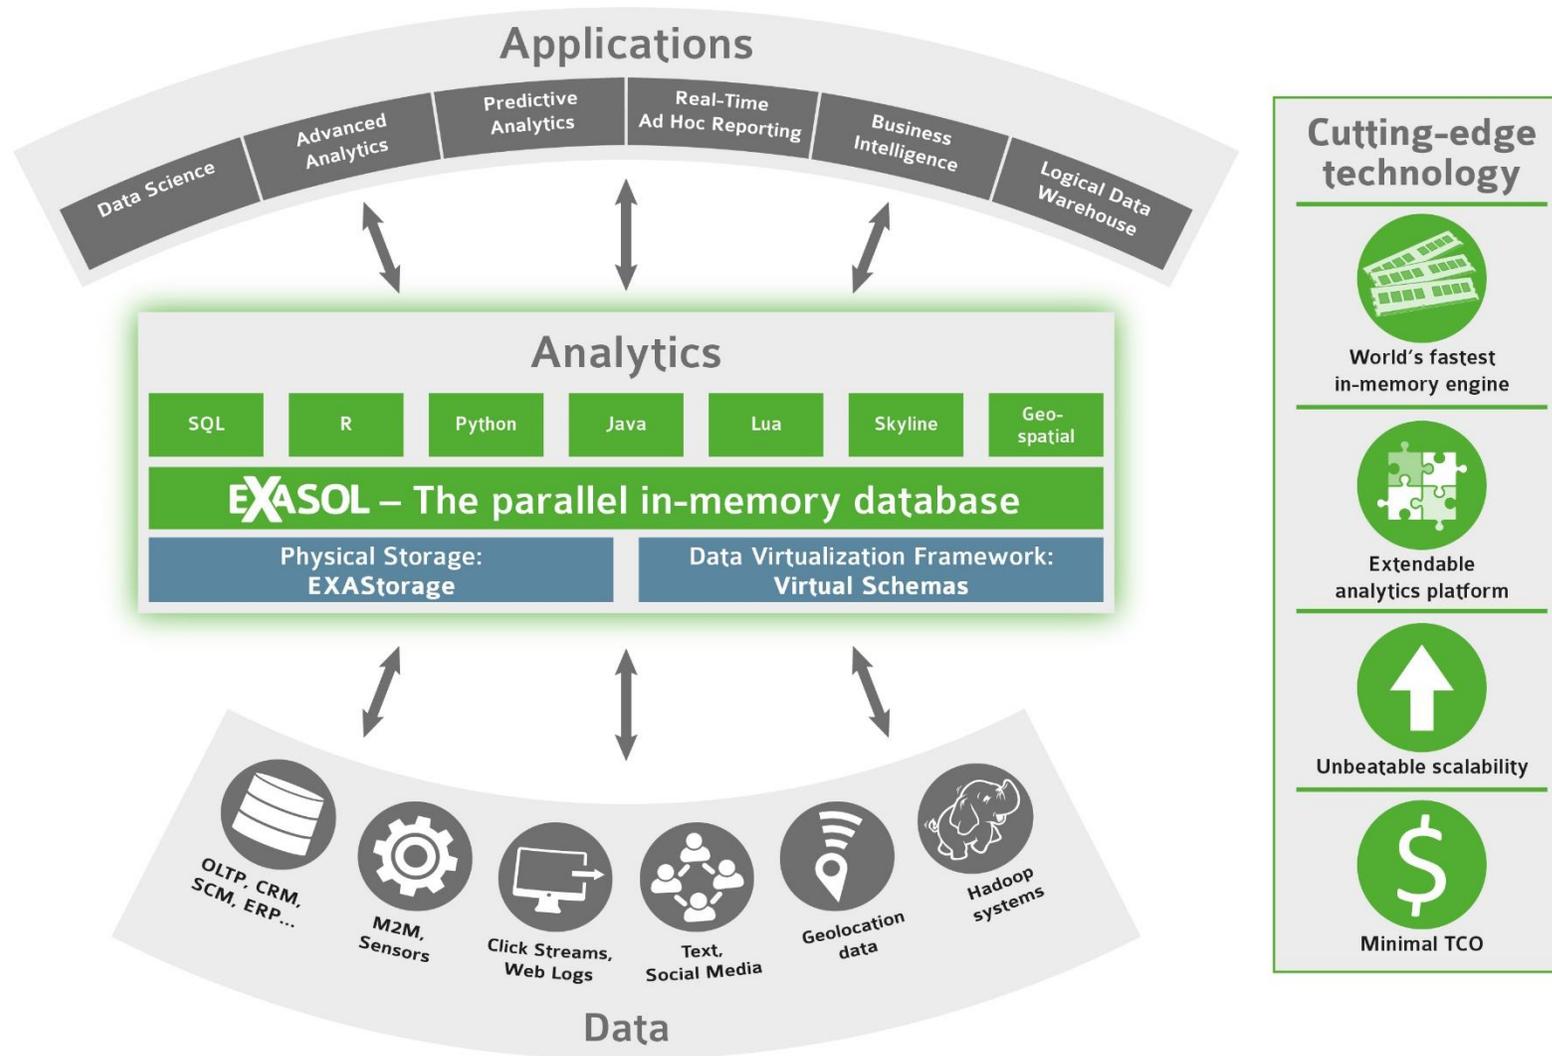
- #1 - dwarfing our followers
- **Best-scalable** solution
- **Best price-performance** ratio



http://www.tpc.org/downloaded_result_files/tpch_results.xls

The most powerful engine for your analytics

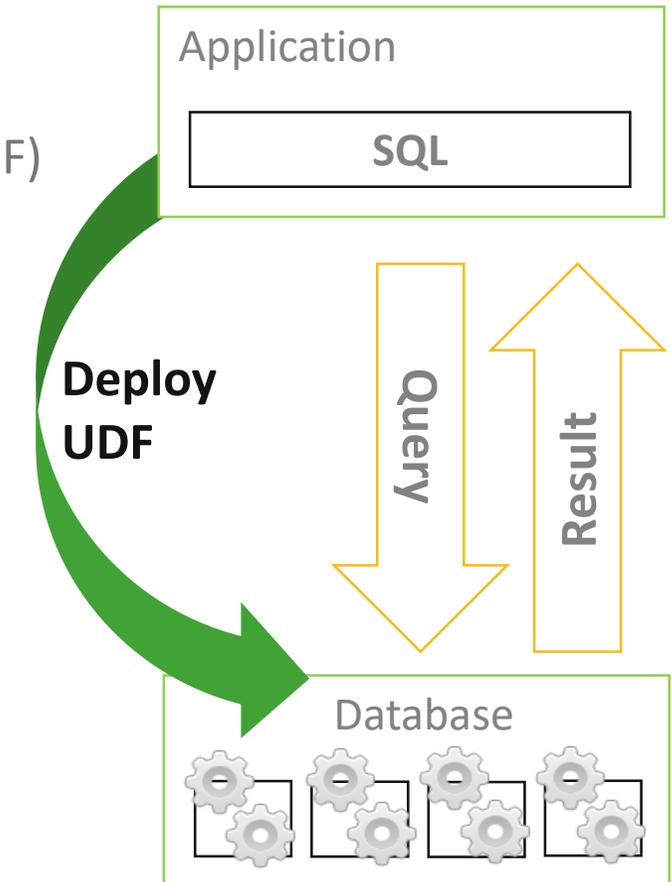
EXASOL



Advanced In-Database Analytics with

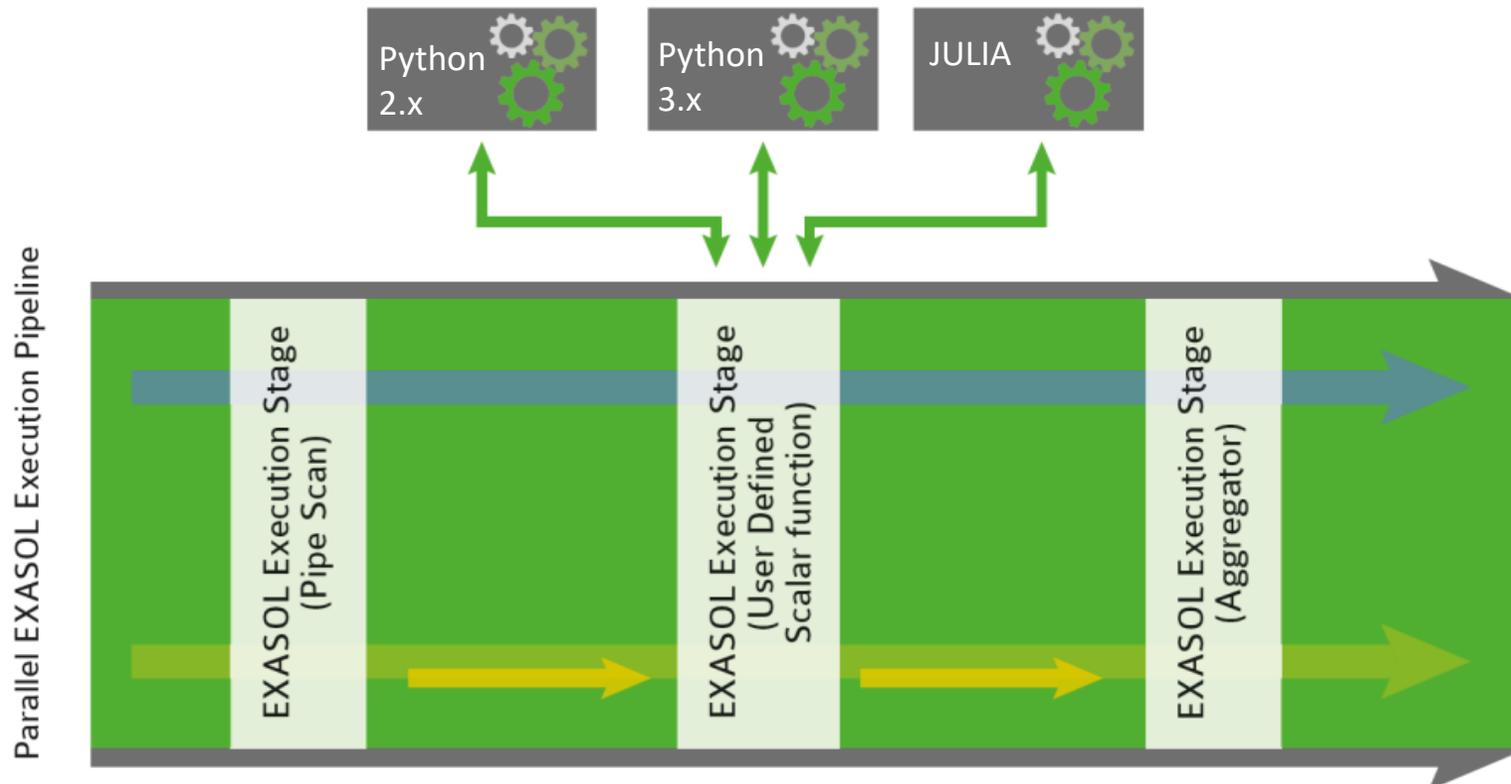


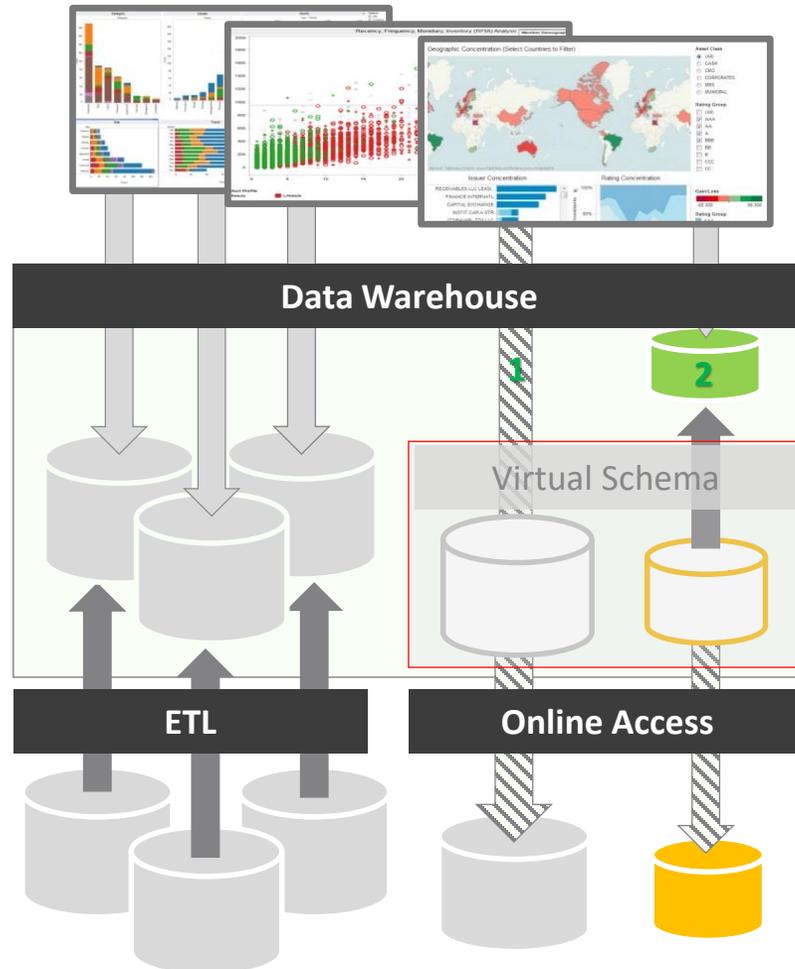
- Bring the algorithms to the data
- Massive parallel execution of R, Java, Python.... scripts
- Embedded into SQL database processing with User Defined Functions (UDF)
- Integrated into standard tools like RStudio
- Seamless integration in existing architectures
- **NEW : Universal/pluggable language support (Scala, Julia, C++ etc)**



V6.0 - Universal/pluggable language support

EXASOL 6.0 offers a framework to integrate any analytical programming language. You are no longer limited to the languages provided by EXASOL out-of-the-box. Just package the programming language of your choice or the language used within your company, deploy it to your EXASOL database and use it for in-database analytics.





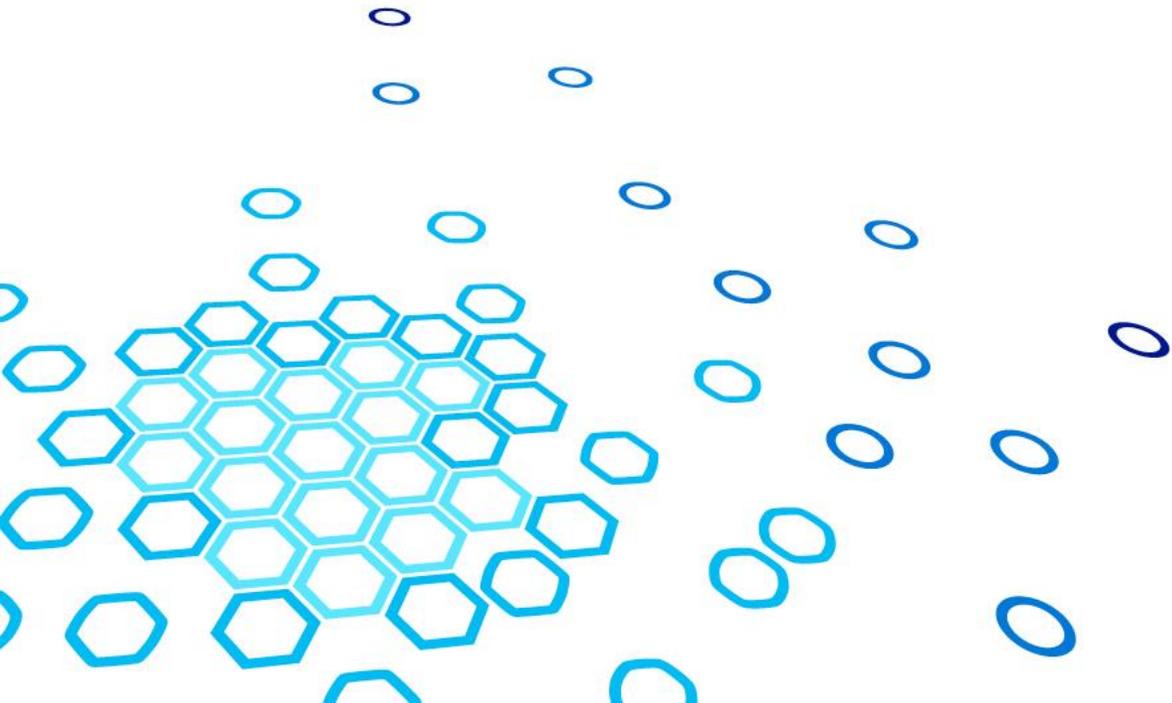
Technical details

- ✓ Intelligent predicate pushdown based on connected data source capabilities
- ✓ Support for **joins** between:
 - ✓ Virtually-connected and physical tables
 - ✓ Multiple virtually-connected tables

Advantages

- ✓ Agile access to most recent information
- ✓ No/reduced redundancy
- ✓ Fewer ETL-jobs
- ✓ No waste of disk space

So, who's using it?



More than 120 Customers in 16 Countries



Zalando: Rising star of e-commerce



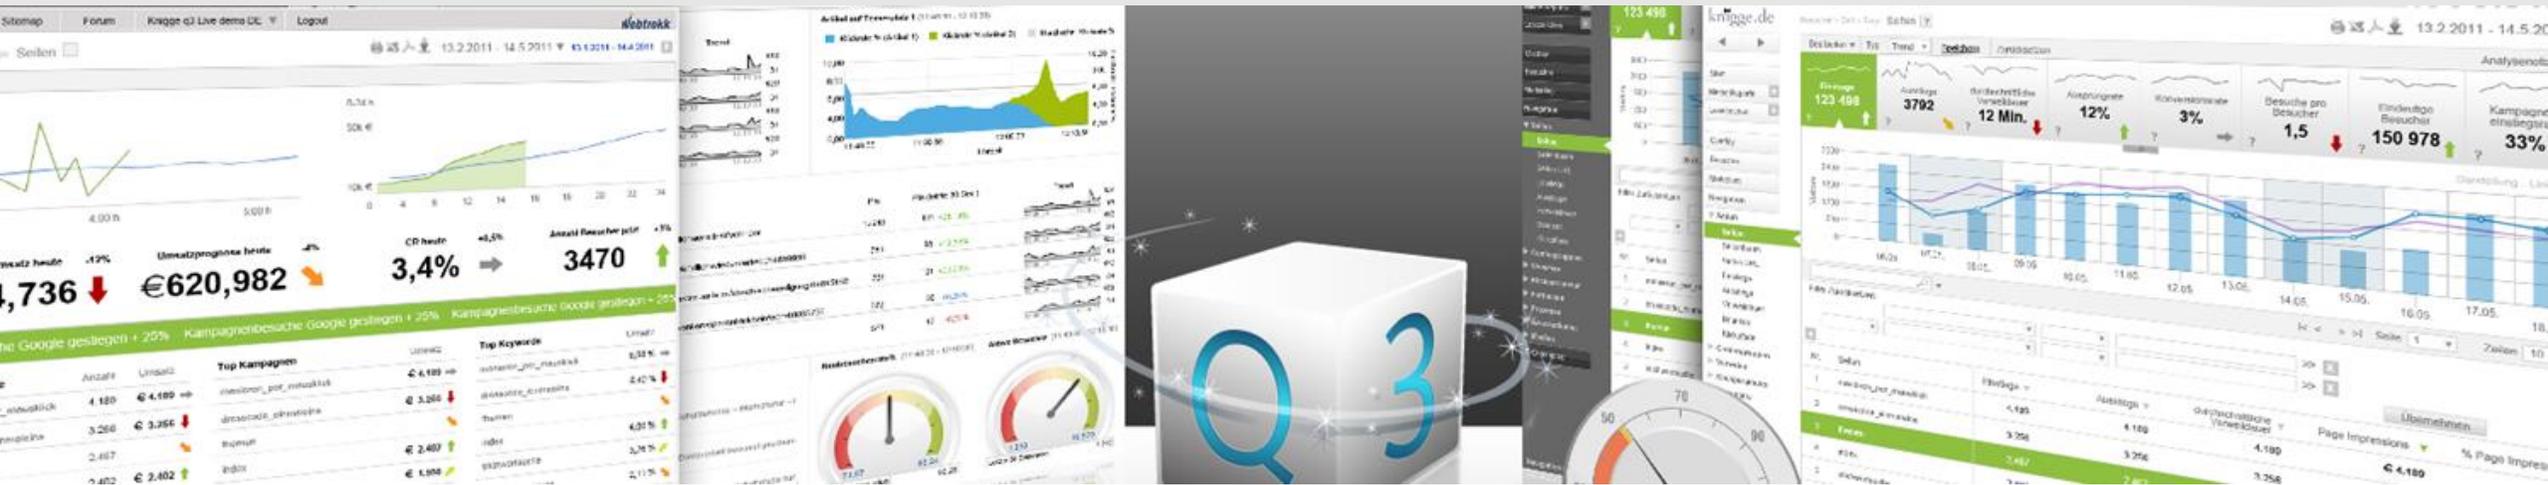
- Online fashion retailer with **14m+ customers across Europe**
- **150,000 products** available online
- EXASOL complements DWH to enable **fast analytics and reporting**
- Database **optimizes stock availability**, returns process and targeted marketing
- EXASOL performs **shopping cart analyses** and powers **product recommendations**



EXASOL database: 15TB

Webtrekk: tracking >30,000 domains

EXASOL



- High-end digital intelligence provider
- EXASOL in **business-critical, 24x7** operation
- Ad hoc **segmentation** analysis in **real-time**
- Continuous near real-time loading stream

Webtrekk
Drive Your Profit.

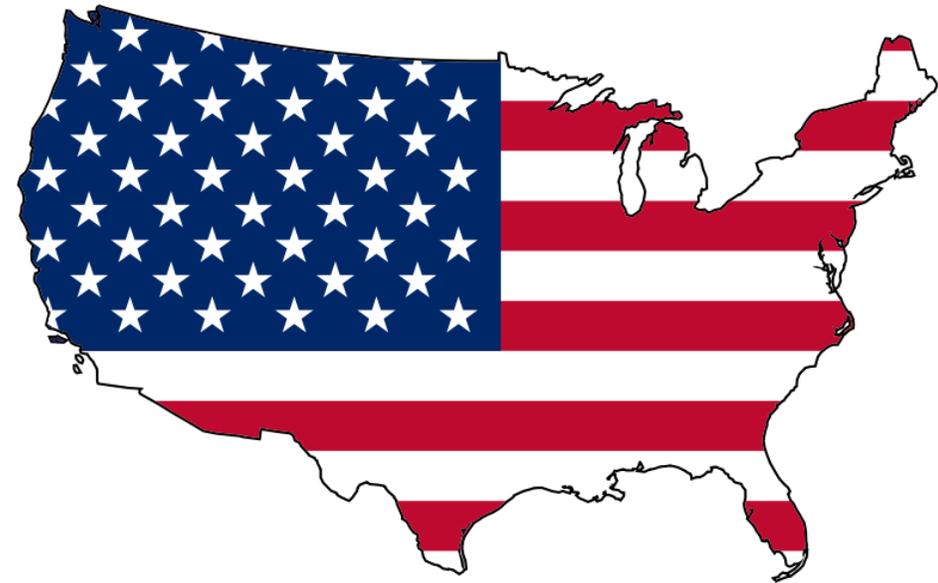
EXASOL database: ~300 TB

King has more players than the entire US



500 m

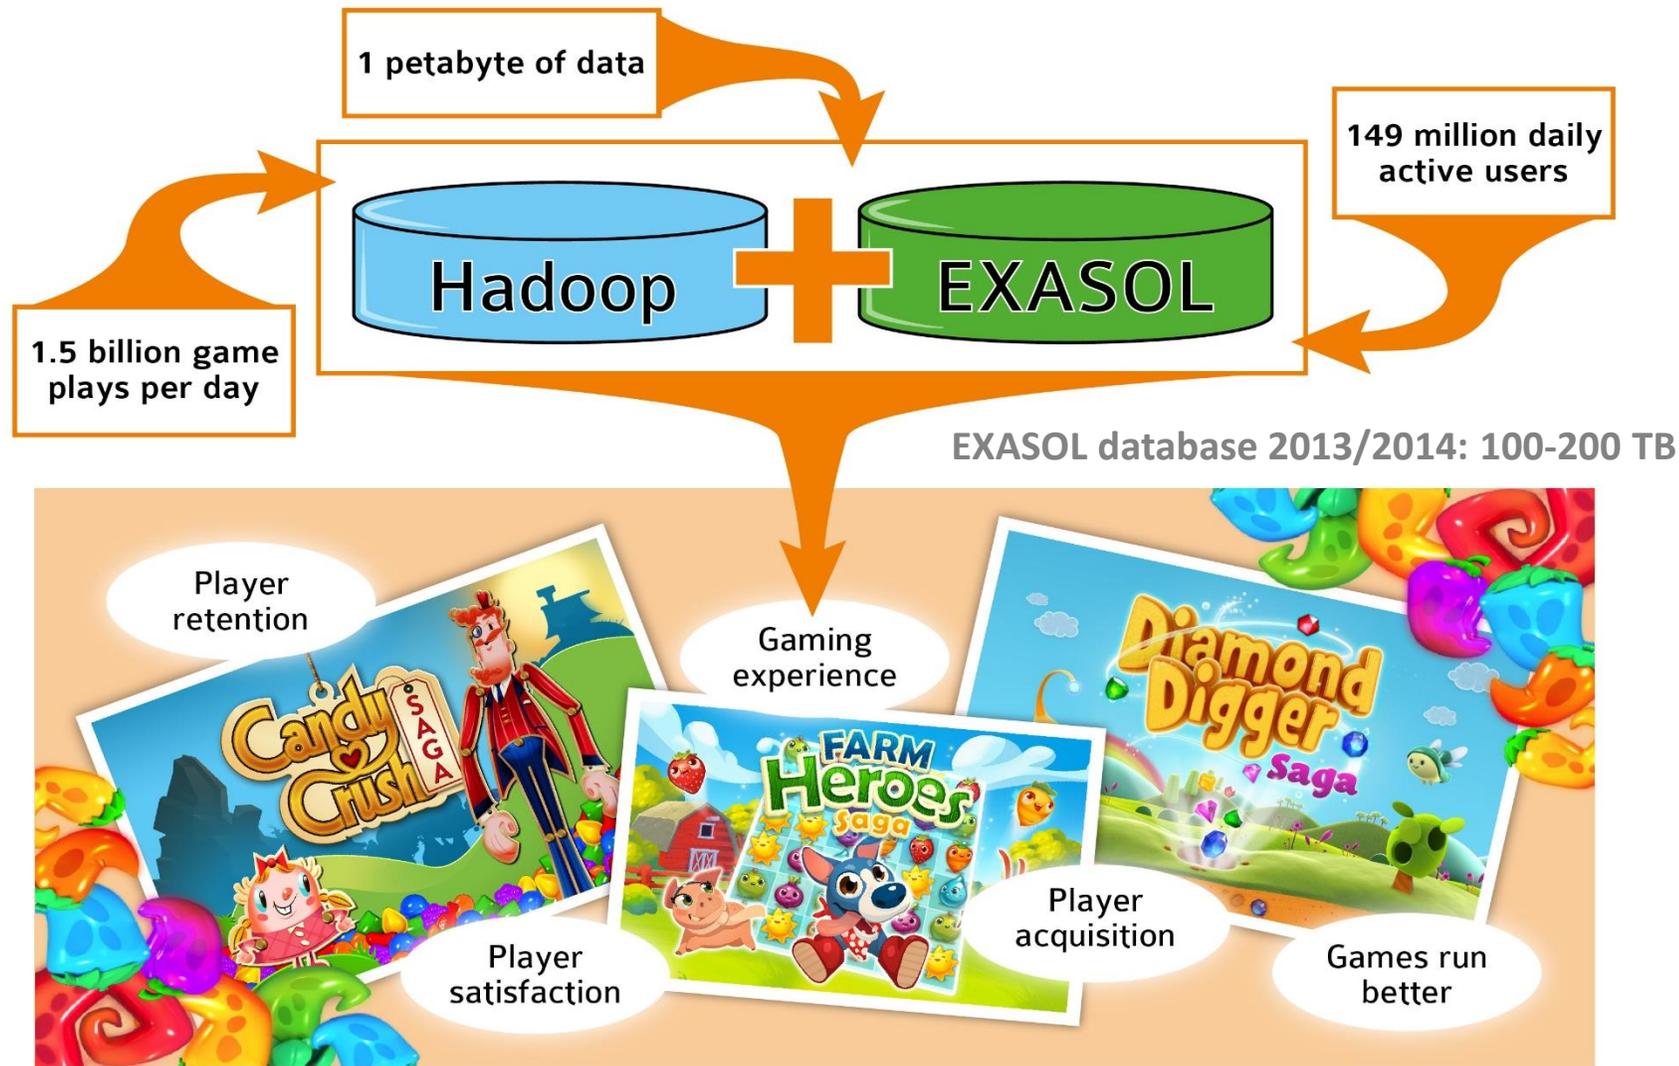
monthly active users in 2015



320 m

Putting data & customers at the heart of its business

2013



King: leading interactive entertainment company



Candy Crush maker King to move to EXASOL system four times the size of its existing 200TB set-up

Seriously Playful

Home About

Business Intelligence Developer

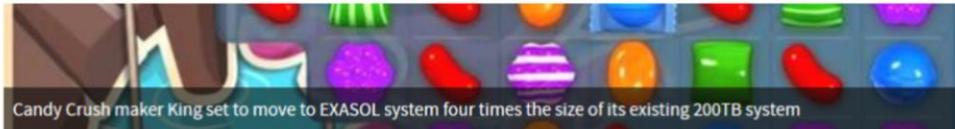
2015

London | Performance, Data & Strategy | Ref: 0587

The Data Analytics organisation at King – split across Barcelona, London and Stockholm – is responsible for processing all the data we receive each day and to use it to drive all of King's data and analytics needs.

With a Hadoop cluster counting 8 Petabytes of data, 16 Terabytes of memory and over 4000 CPU cores and 300 nodes, an Exasol cluster of 64 nodes, 32 Terabytes of RAM and a trillion events generated every day, data at King is in the heart of the business.

As a Business Intelligence Developer, you will help to bring awesome data insights to the business through your data science collaboration with them and your colleagues in our ever-expanding team.



King has outgrown its existing EXASOL system as it 'never throws any data away'

Candy Crush maker King is to move away from its existing EXASOL system to a new one that is quadruple the size, in order to handle increasing data volumes.

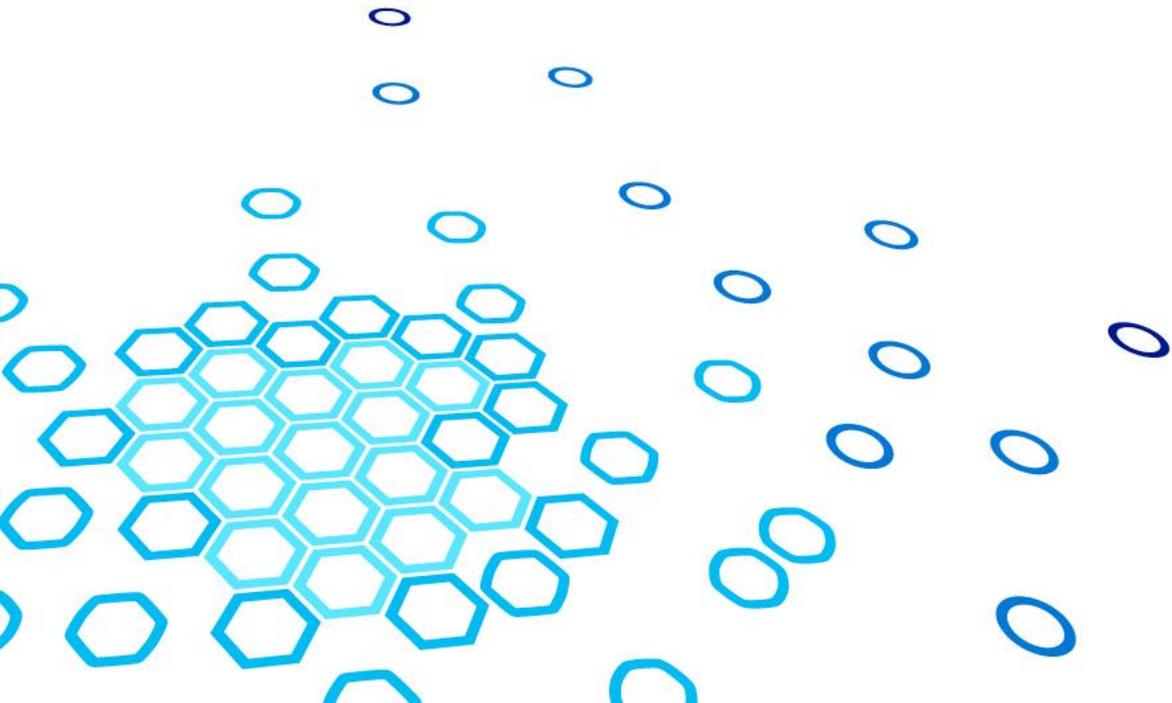
EXASOL database 2016: 64 Nodes & ~0,8 PB

<http://de.slideshare.net/huguk/king-hug-uk>



Let's Have a Look...

EXASOL on Azure



Thank You Very Much!

© 2016 Microsoft Corporation. All rights reserved. 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.

