



Microsoft Cloud Data Platform

Het data platform van de toekomst

Sam Straube

Technology Advisor

sam.straube@microsoft.com

Data is the
new electricity



Industry trends push the boundaries

Devices

more than half
of information workers
across 17 countries report
using **3+ devices** for work

Apps

one quarter
of external app
implementation spending
will be on **mobility, cloud,
analytics & social**, by 2016

Big data

40ZB
the size of digital universe
by 2020, of which 90%
will be **unstructured data**

Existing enterprise apps under pressure to keep up


Legacy hardware + software

De-prioritized IT resources

Pressure to go mobile



Big data is driving us to a tipping point




Real-time data creation

Multiple data types


Large data volumes

Hardware and storage economics



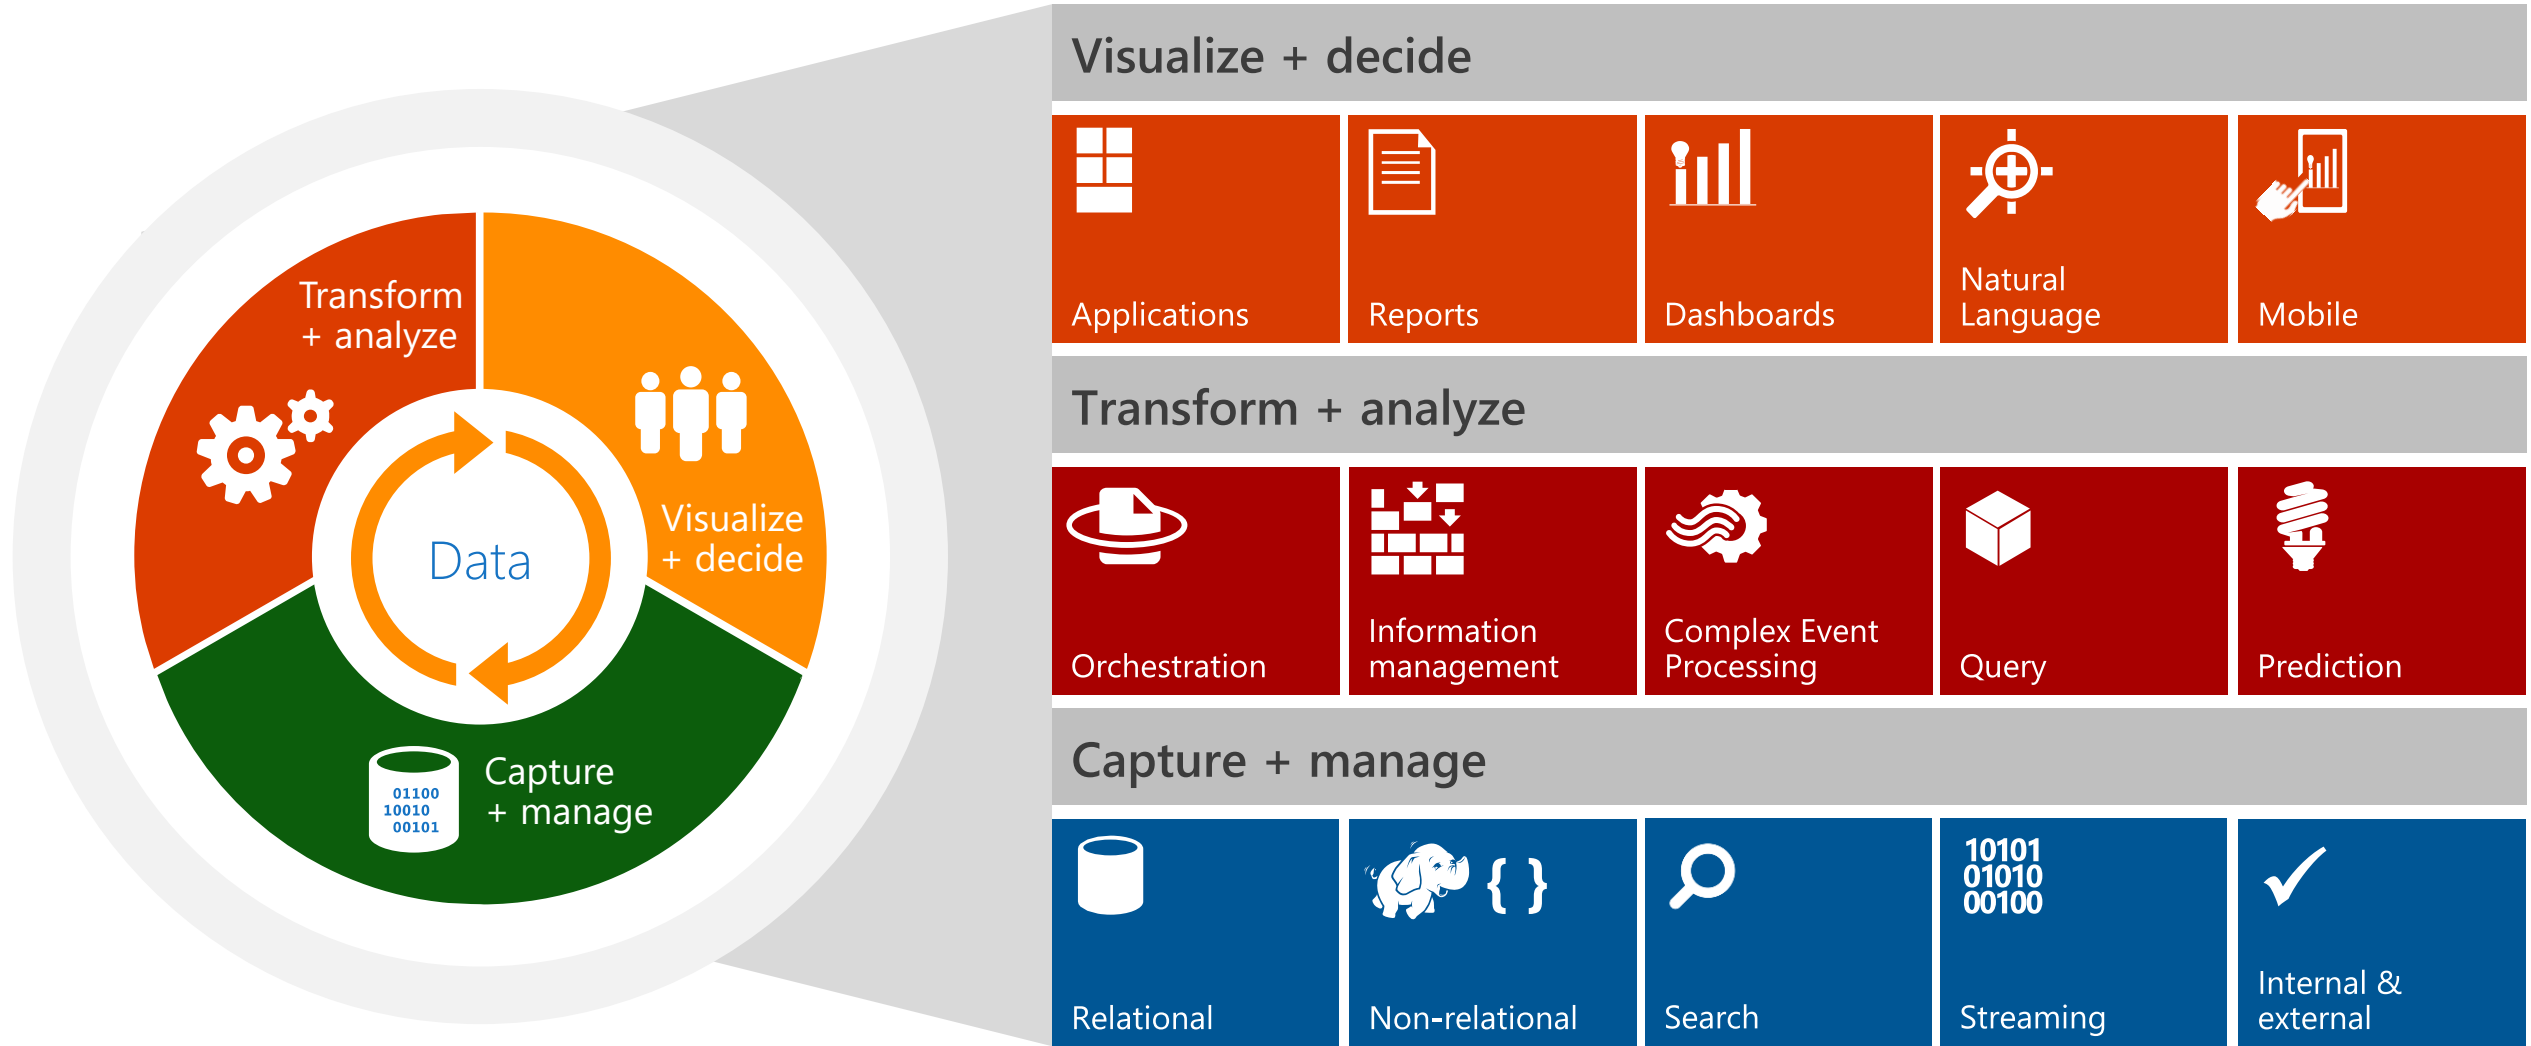
User expectations

Mobility

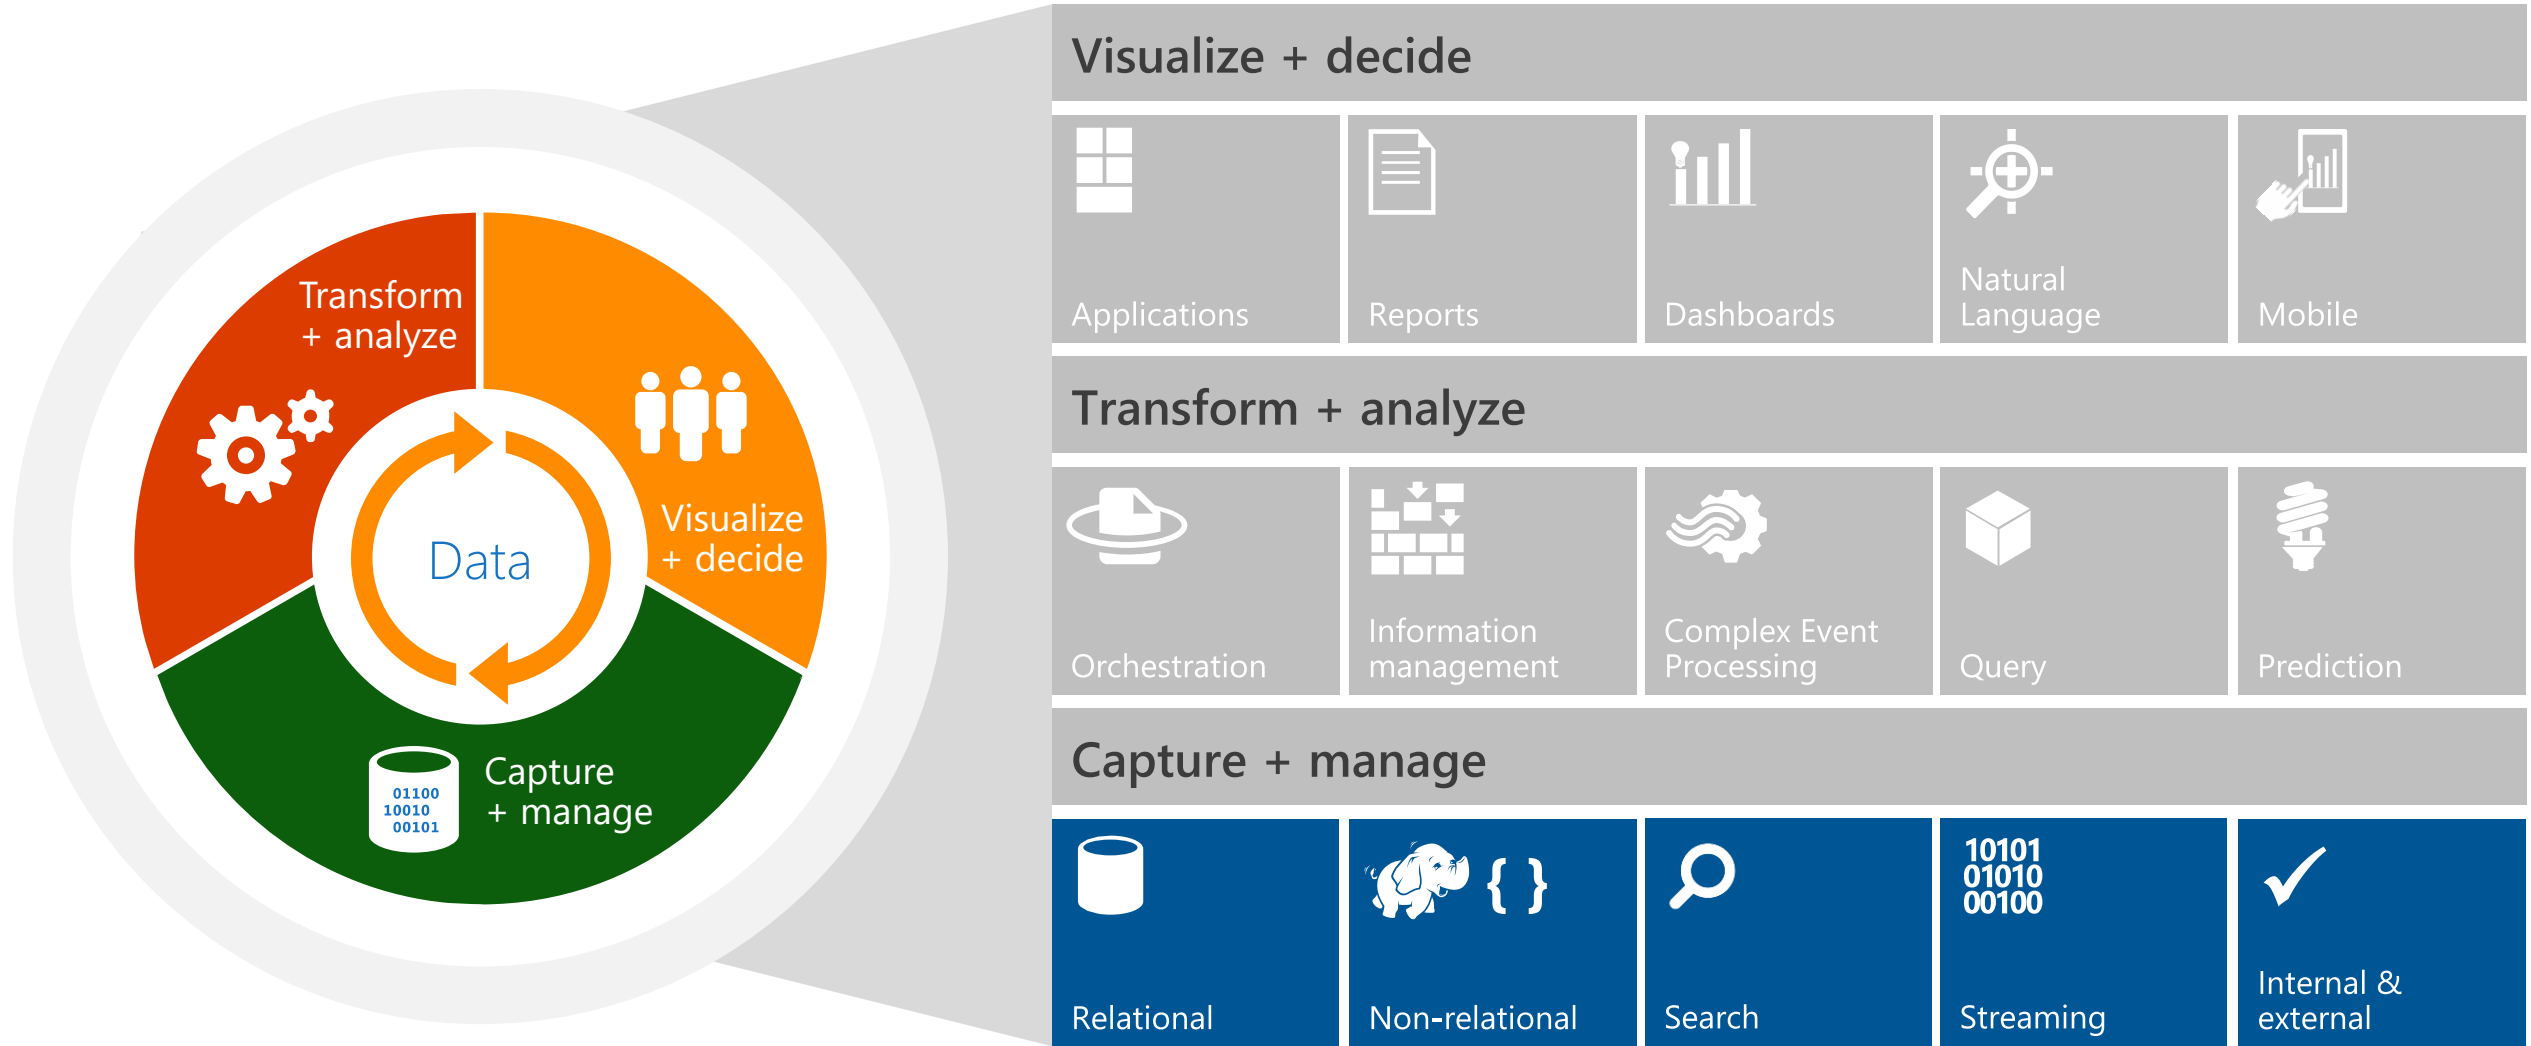




The Microsoft data platform capabilities

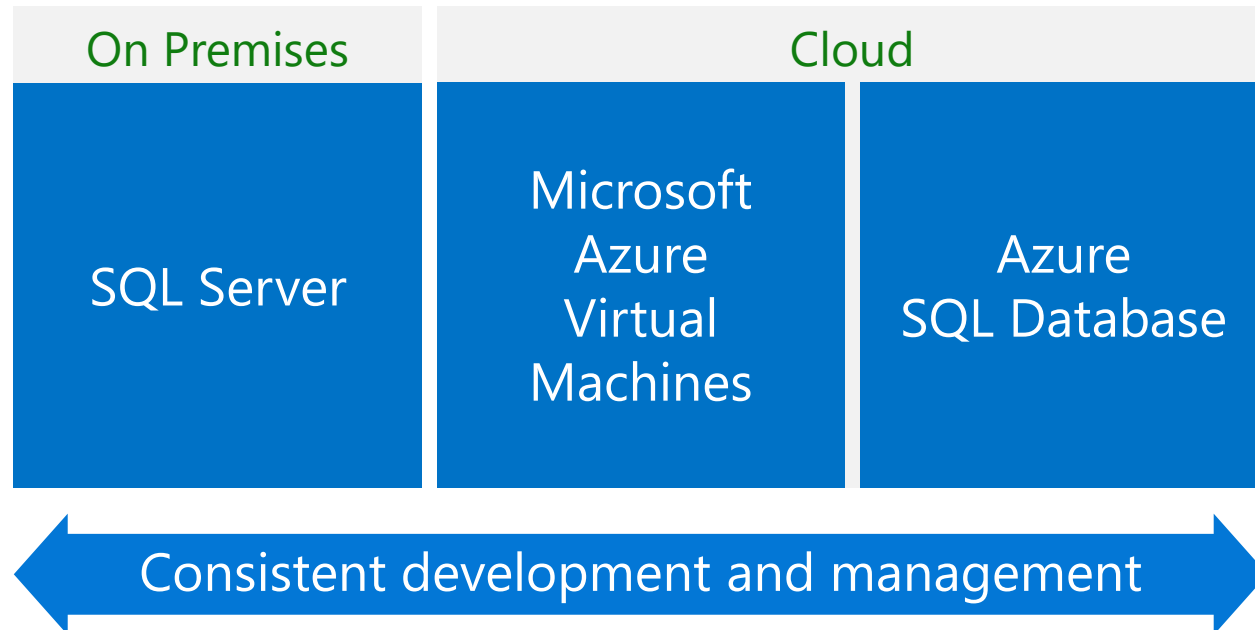


The Microsoft data platform capabilities



The data platform continuum

One consistent platform
with common tools.



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 5/4/2015



TPC-H 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

Drive Real-Time Business with Real-Time Insights

Speed insights across all data, with SQL Server on-premises or in Windows Azure

PowerBI for Office 365

Better business decision, faster



Firewall




In-Memory OLTP

Up to 30x faster



In-Memory DW

Over 100x faster

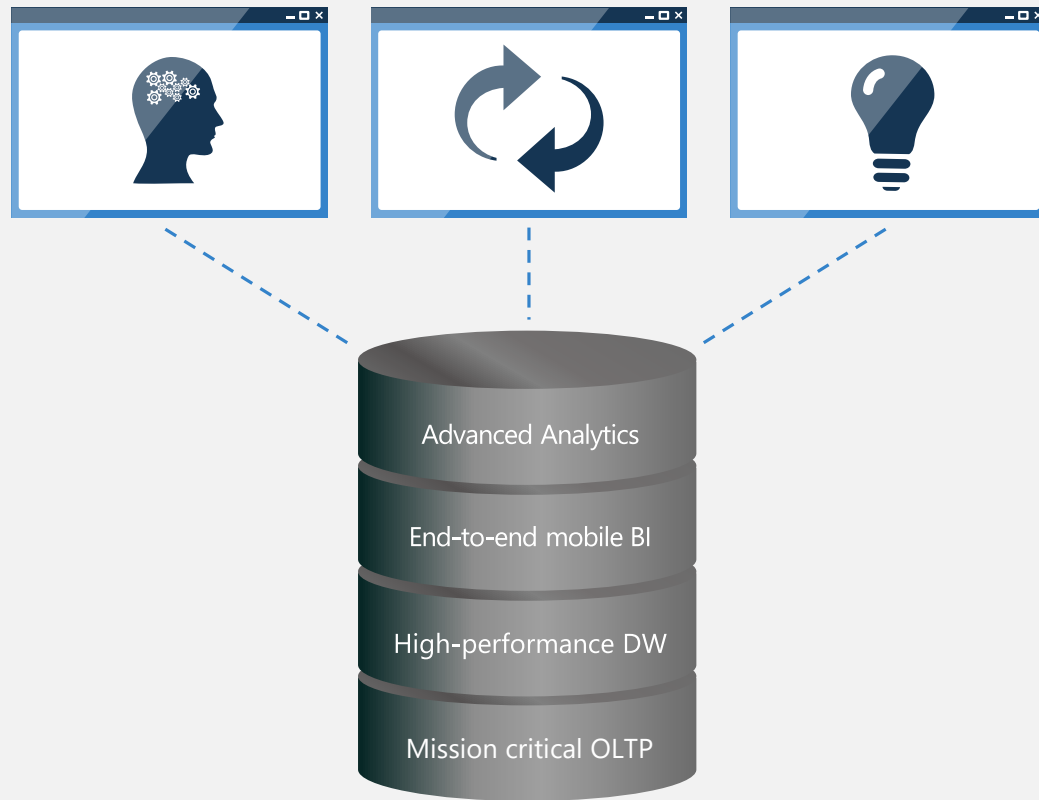


In-Memory Analytics

Billions of rows in seconds

Advanced Analytics

Bring Advanced Analytics to your data



Enterprise

- R built-in to your T-SQL
- Enhanced R APIs with full parallelism and no memory limits for scale/performance
- Built-in In-memory Advanced Analytics
- Advanced tabular model
- Direct query
- Advanced data mining
- SSDT in Visual Studio

Standard

- Connectivity to Microsoft "R" Open
- R APIs with serial execution and memory limitations (2 cores & memory of the host machine)

Enterprise includes all Standard features.

Query Store

Your flight data recorder for
your database



Introducing Query Store

Your flight data recorder for your database

Have You Ever...?

- ...had your system down/slowed down and everyone waiting for you to magically fix the problem ASAP?
- ...upgraded an application to the latest SQL Server version and had an issue with a plan change slowing your application down?
- ...had a problem with your SQL Database and been unable to determine what was going wrong?

With Query Store, ...

- I **CAN** get full history of query execution
- I **CAN** quickly pinpoint the most expensive queries
- I **CAN** get all queries that regressed
- I **CAN** easily force better plan from history with a single line of T-SQL
- I **CAN** safely do server restart or upgrade

What gets captured by Query Store?

Compile time stats	Query text
	Semantic-affecting settings
	Containing object: SP, TVF, trigger
	Parametrization type
	Compilation, binding, optimization stats
	Query plan + initial & last compile / exec times
Runtime stats (aggregated on an <u>interval</u>)	Count of executions, first / last execution time
	AVG, MIN, MAX, STDEV for <i>{metrics}</i> <ul style="list-style-type: none">• Duration• CPU time• Logical IO reads / writes• Physical IO reads• DOP• Memory grants• Number of rows



Microsoft Azure SQL Database

What is SQL Database?

A relational **database-as-a-service**, fully managed by Microsoft.

For cloud-designed apps when **near-zero administration** and **enterprise-grade** capabilities are key.

Perfect for organizations looking to dramatically **increase the DB:IT ratio**.

	SQL Server in a VM	Azure SQL Database
Best for...	<i>Existing applications which requires full box product functionality.</i>	<i>Applications that need elastic scale and/or reduced overhead.</i>
Resources	<i>Customer has ecosystem of IT resources for support and maintenance.</i>	<i>Customer does not want to add additional IT resources for support and maintenance.</i>
TCO benefits	<i>Removing CAPEX.</i>	<i>Avoiding CAPEX and reducing OPEX.</i>
Scalability	<i>Scale up to 20,000 IOPS</i>	<i>Scale out to thousands of DBs, process TBs of OLTP data</i>

SQL Database – ready for business-class apps



Increased from 99.9% to **99.99%** uptime SLA



New service design point enables scale up of resources, delivering **predictable throughput & performance**



Point-in-time-restore, geo-restore, and standard and active geo-replication **protect against human & environmental-initiated events**



Azure certifications: ISO, HIPAA BAA, EU Model Clause
Auditing on SQL Database



Hourly billing & broad set of price points

Elastic databases



Elastic databases at a glance



Elastic database pools
and elastic database
pricing model (preview)



Elastic database tools—
client library and split-
merge service (GA)




Elastic database job
(preview)



Elastic database
queries and transactions
(coming soon)

Concepts



An elastic database pool is a collection of DTUs and storage that is used by multiple databases

Elastic database jobs allow you to perform tasks across databases in the pool

Scenarios include

- Performing administrative tasks, such as deploy new schema
- Update reference data; for example, product information common across all databases
- Rebuild indexes to improve query performance

Performance



Database Throughput Unit – DTU

Bounding box


Monitoring database workload utilization within bounding box

Represents the relative power (resources) assigned to the database

Blended measure of CPU, memory, and read-write rates


Compare the power across performance levels

Simplifies talking about performance, think IOPS vs. %



Performance predictability

Web / Business

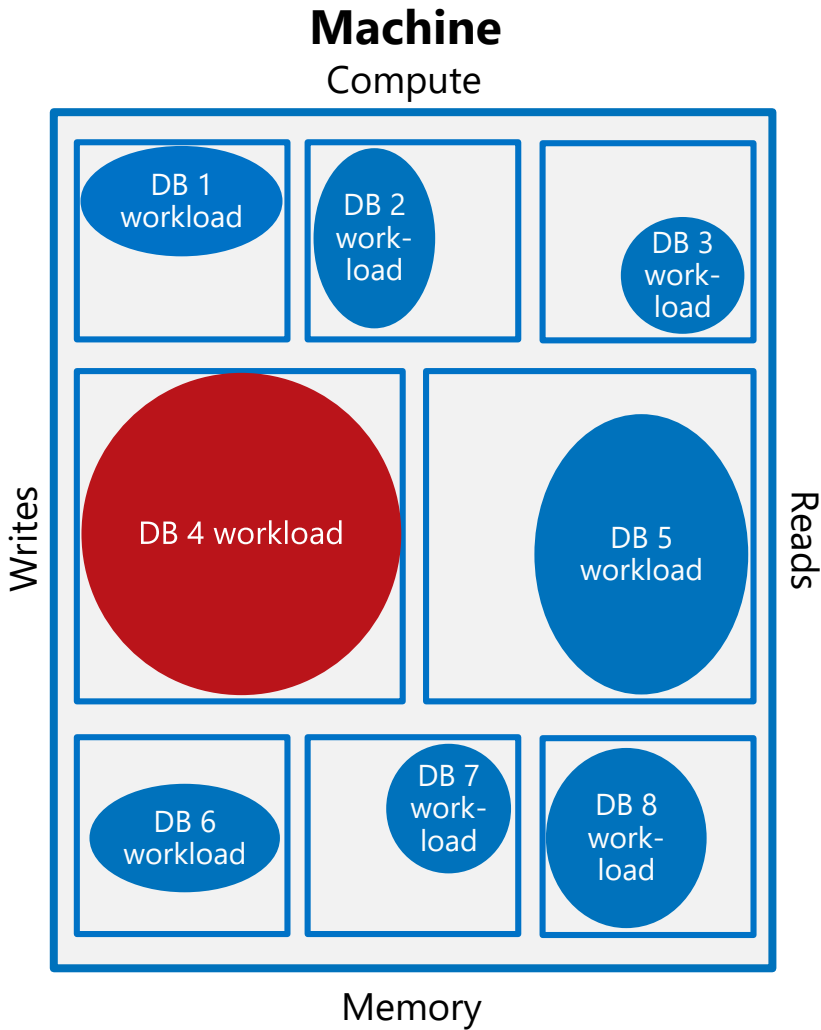


Noisy neighbor!

Introduction of
bounding boxes
to eliminate the
noisy neighbor problem




Basic / Standard / Premium



Migrate to Azure SQL Database



Migrating to Azure SQL Database (v12)



Migrate an on-premises SQL Server database to Azure SQL Database (v12)

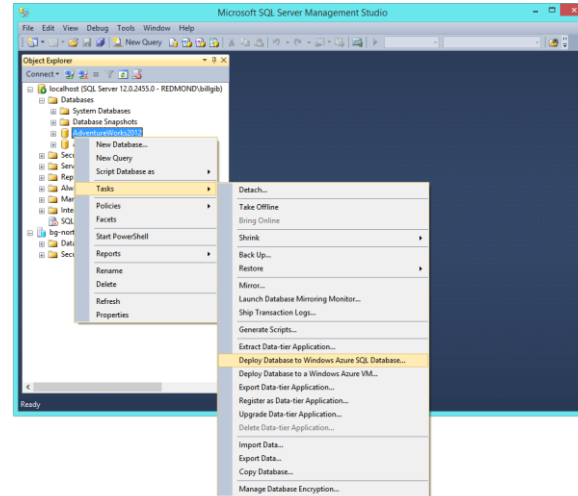
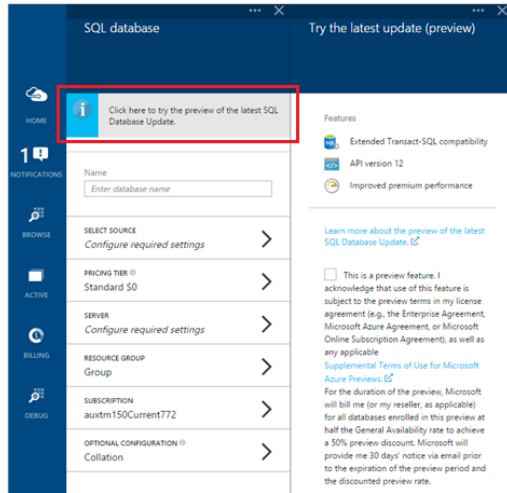
Simplify migration:

- Straightforward migration for the majority of databases

- Requires few or no changes to schema

- No re-engineering of applications**

Migration tools

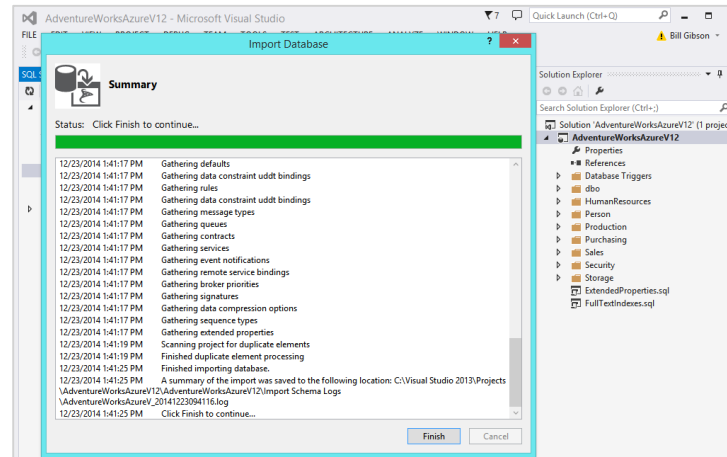
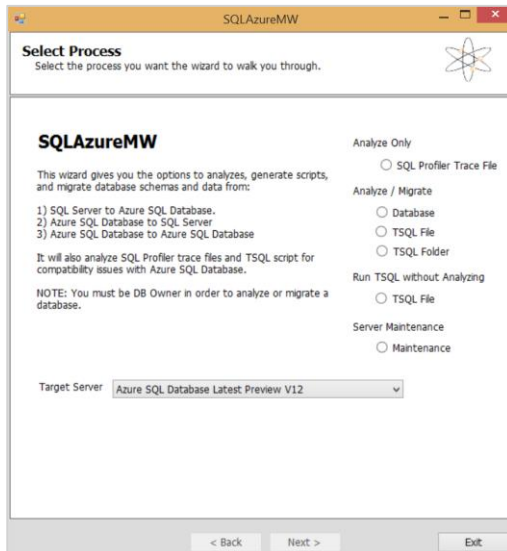


Preview Azure Management Portal

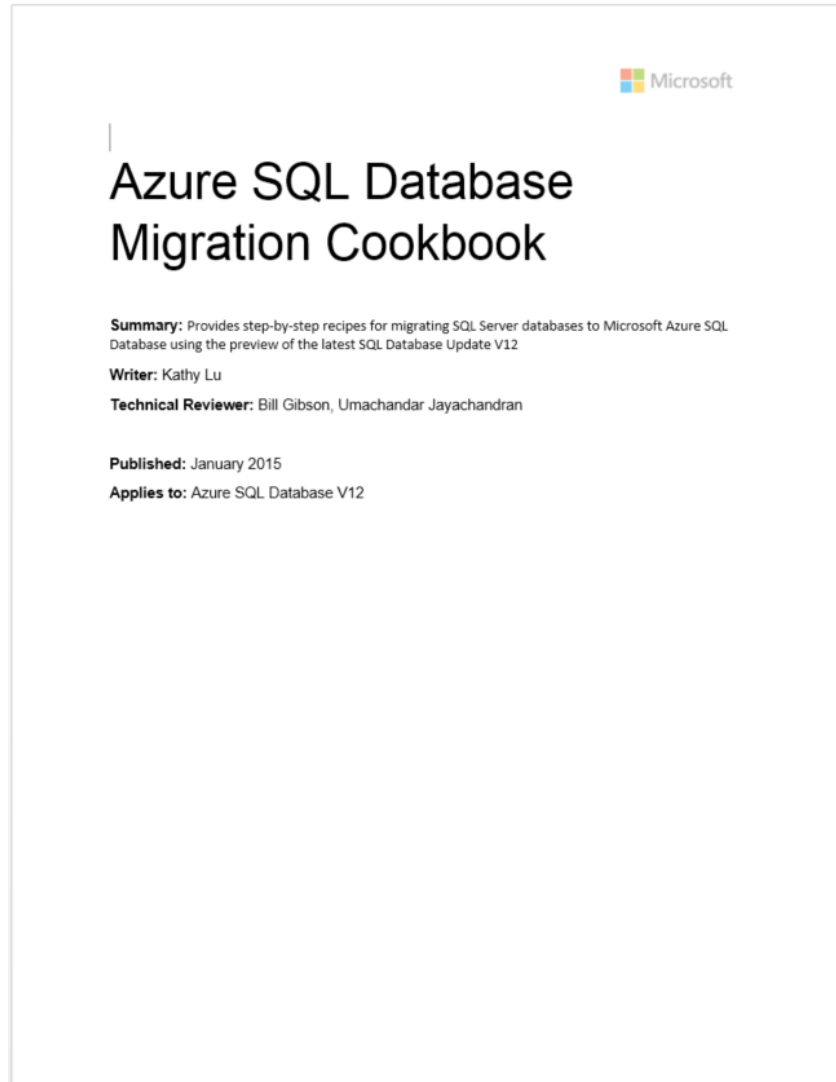
SQL Server Management Studio (SSMS)

SQL Azure Migration Wizard (SAMW)

SQL Server Data Tools in Visual Studio



Migration Cookbook



Migrate an on-premises SQL Server database to Azure SQL Database.

The Migration Cookbook describes various approaches you can use to migrate an on-premises SQL Server database to the latest Azure SQL Database Update.

Download:

<http://aka.ms/azuresqlmigration>

Can it support terabytes of data?

With SQL Database Elastic Scale technology, scale out to 10s of terabytes



SQL Database
scale up limits

DTU (throughput) currently up to 800 DTU or ~735 tx/second*
Max DB size (up to 500GB)



Customer
dimensions to
consider

Pure max data size
Active portion of total data
Amount of transactional workload the app will generate
Largest amount of data that needs to live in the same
transactional space (i.e. database)



Scale out
options

100 S2 for 5000 DTUs (max db size 25TB)
25 P2 for 5000 DTUs (max db size 12.5TB)

* Based on Azure SQL Database Benchmark estimation and specific OLTP workload configuration

Azure SQL VMs vs SQL Database

Scenarios

Public Cloud



SQL Server in Azure Virtual Machine



Azure SQL Database

KEY SCENARIOS



Move Existing
Tier 2 & 3
Apps



Develop
& Test



Backup
Database to
the Cloud



Extend
On-Prem Apps
to Azure VM



Develop New
Cloud-Designed Apps



Extend On-Prem Apps
to Azure
SQL Database



Use both **SQL Server in Azure VM** and **SQL Database** in a single app.

Appendix

SQL Server 2016 Enterprise

Enterprise	OLTP Performance	Security	Data Warehousing	Business Intelligence	Advanced Analytics	Hybrid Cloud
	<ul style="list-style-type: none"> • OS max cores and memory • Enhanced in-memory OLTP performance • Operational analytics • Enhanced AlwaysOn with no domain join (WS 2016) • QueryStore • Temporal 	<ul style="list-style-type: none"> • Always Encrypted • Row-level security • Dynamic data masking • Enhanced separation of duties • Enhanced SQL Server auditing • Transparent data encryption 	<ul style="list-style-type: none"> • Enhanced in-memory ColumnStore • PolyBase in scale-out configuration (head and compute nodes) • Deployment rights for APS • Distributed query processing • Support for JSON 	<ul style="list-style-type: none"> • End-to-end mobile BI on all major platforms • Enhanced direct query • In-memory analytics • Advanced data mining • Advanced tabular • Web portal experience (all reports in 1 place) • Modernized reports • Pin report to Power BI • Enhanced multi-dimensional models 	<ul style="list-style-type: none"> • In database Advanced Analytics • R integration with massive parallel processing for performance and scale • Works with in-memory technology • Run in database or standalone • Connectivity to R Open 	<ul style="list-style-type: none"> • Stretch Database • Enhanced backup to Azure • Enhanced HA and DR with Azure – ease of use, no domain join (WS2016) • SSIS integration with Azure Data Factory and Azure SQL Data Warehouse

Enterprise includes all Standard features.

Subject to NDA

SQL Server 2016 Standard

	OLTP Performance	Security	Data Warehousing	Business Intelligence	Advanced Analytics	Hybrid Cloud
Standard	<ul style="list-style-type: none">• Disk-based OLTP• 24 cores max and 128 GB max memory• 2-node single database failover (non-readable secondary)• Temporal	<ul style="list-style-type: none">• Row-level security• Dynamic data masking• Basic auditing• Separation of duties	<ul style="list-style-type: none">• PolyBase (compute node only)• Support for JSON	<ul style="list-style-type: none">• Basic tabular (16GB memory per instance)• Modernized reports• Pin report to Power BI• Enhanced multi-dimensional models	<ul style="list-style-type: none">• Single-threaded for RRE• Connectivity to R Open	<ul style="list-style-type: none">• Stretch Database• Backup to Azure

Azure SQL Database

SQL Server 2016 and SQL Database v12: 130

Starting in Mid-June the default Compatibility Level will be 130 **for newly created databases**