

Microsoft Tech Summit

Build your skills with the latest in
cloud technologies



Microsoft Tech Summit

28 February – 1 March 2018, Trafo Baden



SQL Server 2017 on Linux, Windows and containers

Marko Hotti

Sr. Technical Product Manager / SQL Server
Microsoft Corporation

marko.hotti@microsoft.com



EVOLUTION OF SQL SERVER


Businesses are embracing choice



Heterogeneous environments



Multiple data types



Different development languages



On-premises, cloud, and hybrid environments

Microsoft is delivering on choice



HDInsight on Linux



R Server on Linux



Linux in Azure



SQL Server drivers and connectivity




Visual Studio Code extension for SQL Server

The world is demanding SQL Server on Linux

20K+

applications for private preview



36%

enterprise DB market runs on Linux

FLEXIBLE, RELIABLE DATA MANAGEMENT SQL Server on the platform of your choice


Support for RedHat Enterprise Linux (RHEL),
Ubuntu, and SUSE Enterprise Linux (SLES)

Linux and Windows Docker containers


Windows Server / Windows 10

Package-based installation: Yum Install, Apt-Get,
and Zypper


Windows



Linux



Linux/Windows container

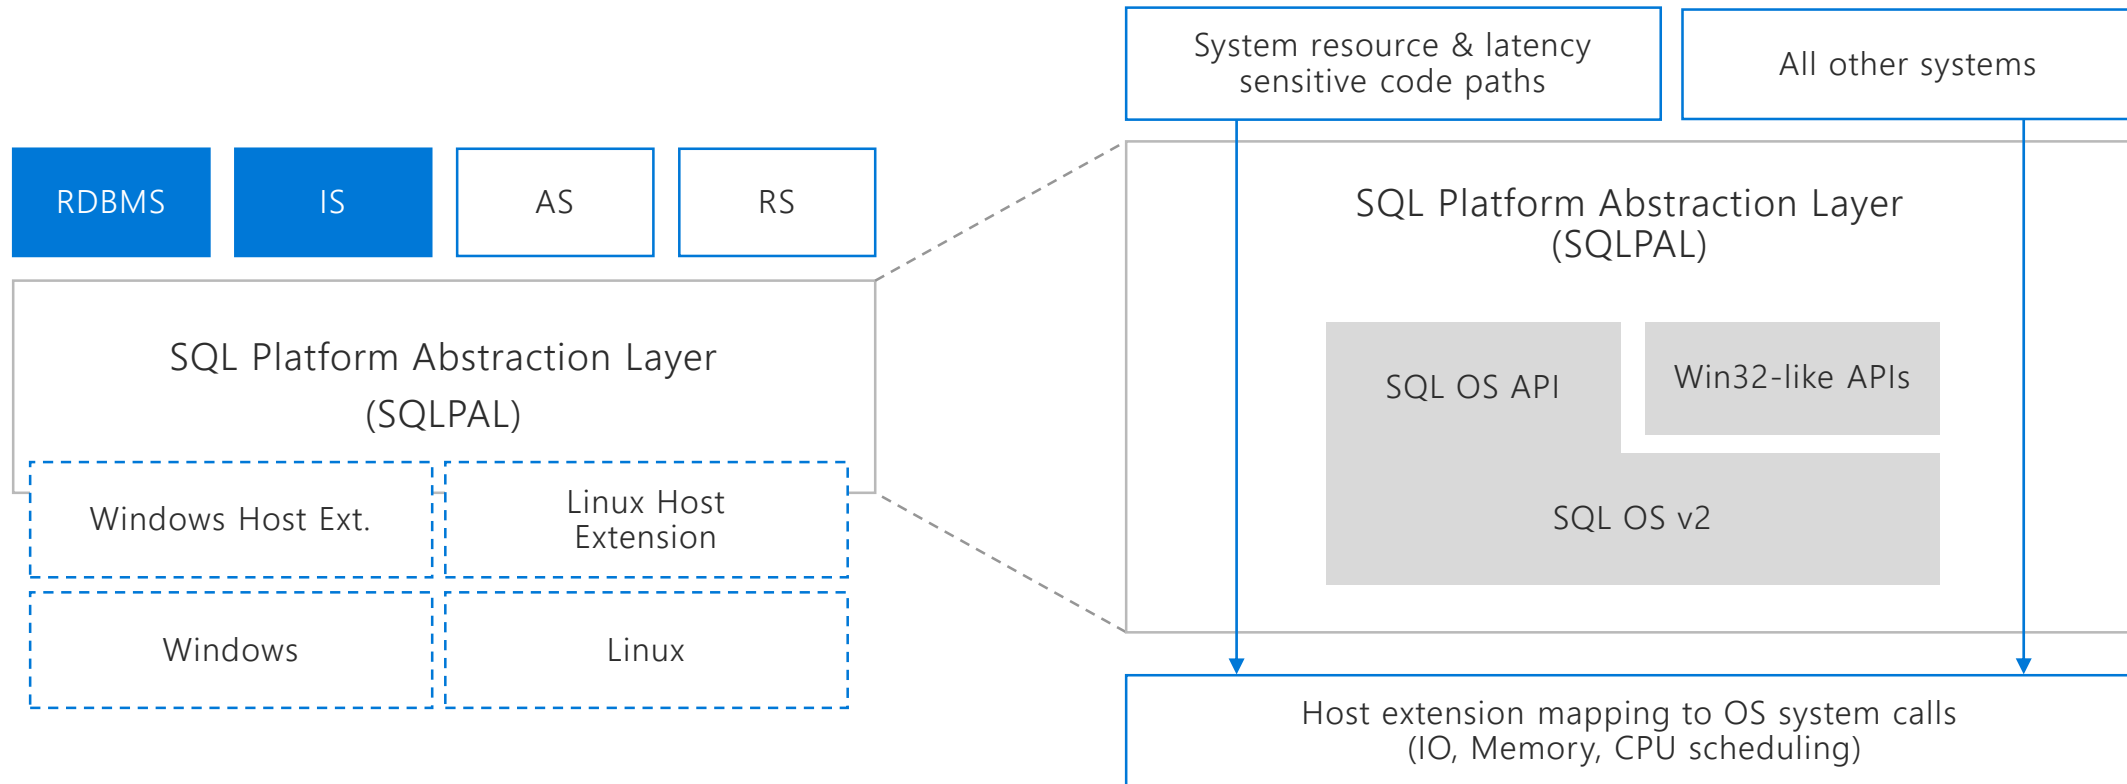


WHAT'S IN SQL SERVER ON LINUX

		Windows	Linux
Editions	Developer, Express, Web, Standard, Enterprise	•	•
Services	Database Engine, Integration Services	•	•
	ML Services, Analysis Services, Reporting Services, MDS, DQS	•	
Mission critical performance	Maximum number of cores	Unlimited	Unlimited
	Maximum memory utilized per instance	12 TB	12 TB
	Maximum database size	524 PB	524 PB
	Basic OLTP (Basic In-Memory OLTP, Basic operational analytics)	•	•
	Advanced OLTP (Advanced In-Memory OLTP, Advanced operational analytics, adaptive query processing)	•	•
	Basic high availability (2-node single database failover, non-readable secondary)	•	•
	Advanced HA (Always On - multi-node, multi-db failover, readable secondaries)	•	•
Security	Basic security (Basic auditing, Row-level security, Data masking, Always Encrypted)	•	•
	Advanced security (Transparent Data Encryption)	•	•
Data warehousing	PolyBase	•	
	Basic data warehousing/data marts (Basic In-Memory ColumnStore, Partitioning, Compression)	•	•
	Advanced data warehousing (Advanced In-Memory ColumnStore)	•	•
	Advanced data integration (Fuzzy grouping and look ups)	•	•
Tools	Windows ecosystem: Full-fidelity Management & Dev Tool (SSMS & SSDT), command line tools	•	•
	Linux/OSX/Windows ecosystem: Dev tools (VS Code), DB Admin GUI tool, command line tools	•	•
Developer	Programmability (T-SQL, CLR, Data Types, JSON, Graph)	•	•
	Windows Filesystem Integration - FileTable	•	
BI & Advanced Analytics	Corporate Business Intelligence (Multi-dimensional models, Basic tabular model)	•	
	Machine Learning Services (R and Python integration)	•	
Hybrid cloud	Stretch Database	•	

SQL SERVER ON LINUX:

SAME AS SQL SERVER ON WINDOWS SQL PLATFORM ABSTRACTION LAYER



LINUX-NATIVE USER EXPERIENCE

Standard installation process



Package-based installation using yum for Fedora-based distributions, apt-get for Debian-based distributions, and zypper for SLES
Existing package update/upgrade processes for SQL Server upgrade

Cross-platform tools



SQL Server command-line tools (sqlcmd, bcp) available for Linux (and soon on macOS)
Existing Windows tools such as SQL Server Management Studio (SSMS), SQL Server Data Tools (SSDT), and PowerShell module (sqlps) to manage SQL Server on Linux from Windows
Cross Platform open source tools such SQL Server Ops Studio, mssql-cli
Visual Studio Code extension for SQL Server on macOS, Linux, or Windows

Familiar experience



SQL Server service runs natively using systemd
Linux file paths are supported in T-SQL statements and scripts (defining/changing the path, database backup files)
Popular Linux high-availability solutions like Pacemaker and Corosync

ANYWHERE AVAILABILITY FOR SQL SERVER


Any platform, any cloud

Run the world's leading data platform on any environment

Works on Windows, Linux, and Docker

Runs on any cloud including Azure, 3rd party hosting providers

Use on-premises licenses in the cloud with License Mobility, a Software Assurance benefit



SQL Server 2017



Azure



BENEFITS OF CONTAINER PLATFORMS FOR SQL SERVER

Faster time to deployment

Packaging dependencies with SQL Server

CI/CD process automation - treat SQL Server like a dependency


Deploy exactly what you dev/test

Cross-platform – develop on Windows/Linux/macOS, deploy on Linux


Microservices architectures

Higher density, lower disk space requirements


APPLICATION DEPLOYMENT PATTERNS USING CONTAINERS



Centralized SQL Server



Docker Compose



Monolithic app or
microservice

DEMO

Install SQL Server on Linux in under 60 seconds

HA AND CLUSTERING FOR MICROSOFT SQL SERVER ON RED HAT ENTERPRISE LINUX

Simple HADR

VM failure

Resilience against guest & OS level failures

Planned & unplanned events

Minimum downtime for patching and upgrades

Minutes RTO

Backup/restore

Protection against accidental or malicious data corruption

DR protection

Minutes to hours RTO

Standard HADR

Failover cluster

Instance level protection

Automatic failure detection & failover

Seconds to minutes RTO

Resilience against OS and SQL Server failures

Basic availability groups

AG with 2 replicas

Log shipping

Warm standbys for DR

Mission-critical HADR

Availability groups

Database level protection

Seconds RTO

No data loss

Recover from unplanned outage

No downtime for planned maintenance

Offload read/backup workload to active secondaries

Failover to geographically distributed secondary site

MISSION CRITICAL AVAILABILITY ON ANY PLATFORM

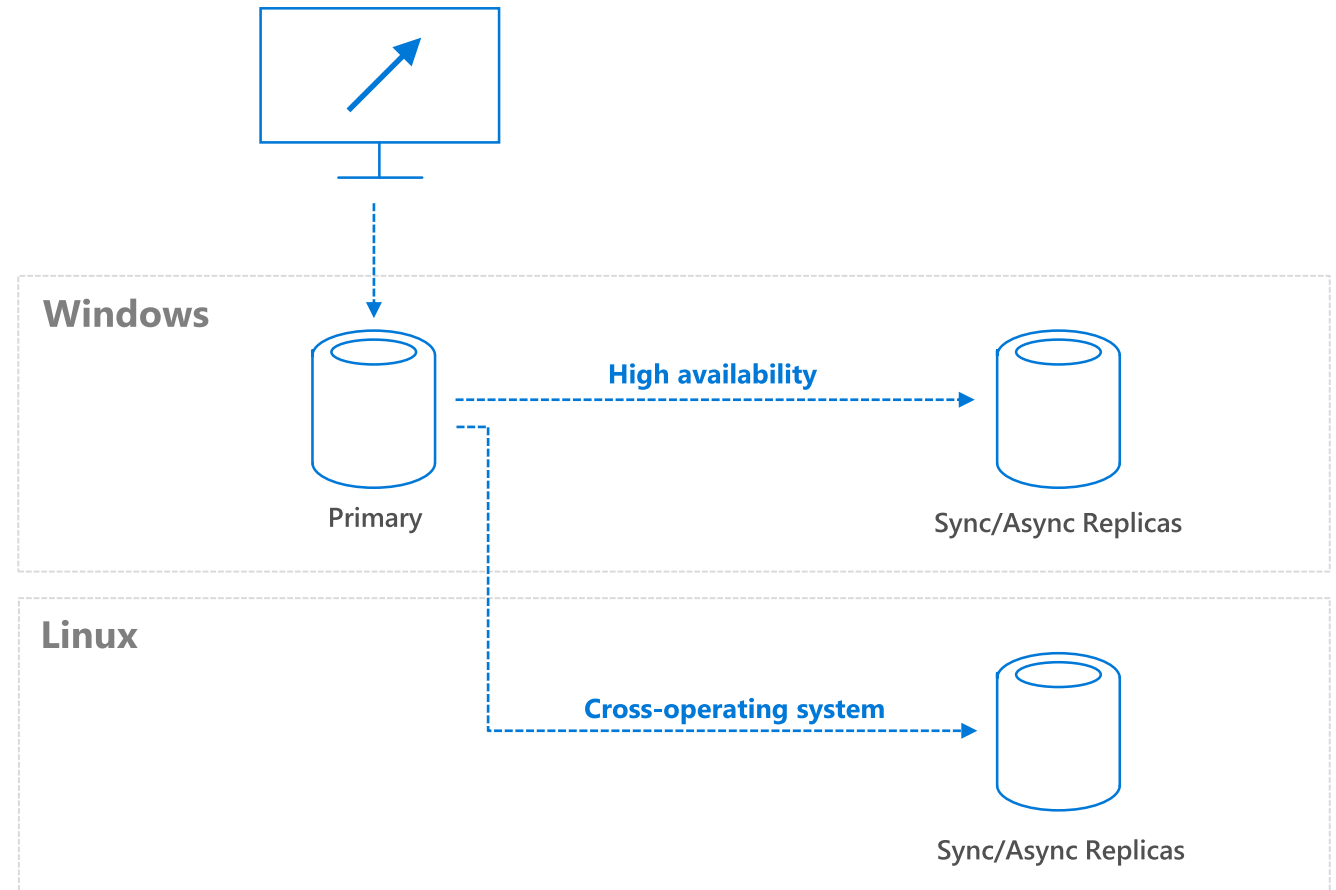
**Always On cross-platform
capabilities**

HA and DR for Linux and Windows

Support for clusterless Availability Groups

Ultimate HA with OS-level redundancy
and low-downtime migration

Load balancing of readable secondaries




DEMO

SQL Server 2017 Always On Availability Groups on RHEL


INTEGRATING RED HAT WITH MICROSOFT ACTIVE DIRECTORY

Integration Options

Direct integration



Indirect integration



INDUSTRY-LEADING PERFORMANCE WITH SQL SERVER 2017


Real-time operational analytics

Bring analytics to your operational data

Improve transactional performance with In-Memory OLTP

Speed analytics and reduce storage needs with In-Memory ColumnStore

LOB compression in ColumnStore index



INVESTORS GET INSIGHTS 15X FASTER

REAL TIME ANALYTICS AT SCALE NOW ON LINUX

CHALLENGE

FinTech company burdened by slow query speeds and time-consuming maintenance sought the familiarity of Linux OS

IMPACT

Query times cut from **30 to <2 seconds** by switching from PostgreSQL to SQL Server

Management time **reduced 90%** and Linux platform simplified dev operations


▲dv01

Before



Queries timing out
with PostgreSQL

SQL Server on Linux



- ✓ In-database analytics
- ✓ MPP scale
- ✓ R and Python support
- ✓ Faster query processing

FLEXIBILITY

STRENGTHENED CYBERSECURITY BY OPTIMIZING DATA

Seamless integration in a heterogenous environment

Challenge

- Cybersecurity company faced sluggish performance on its mission-critical data management platform
- Delayed threat reporting due to slow analytical queries and log-collection process

Impact

- 3X faster query performance and 50% reduction in storage costs
- Enhanced log analysis and anomaly-detection system with in-database analytics and machine learning



"We gained an easy-to-use database that can accommodate all of our data and run new analytical queries quickly."

Peter Kleinert, Solution Architect, Binary Confidence

PETABYTE-SCALE DATA WAREHOUSING

The world’s first enterprise-class
diskless database

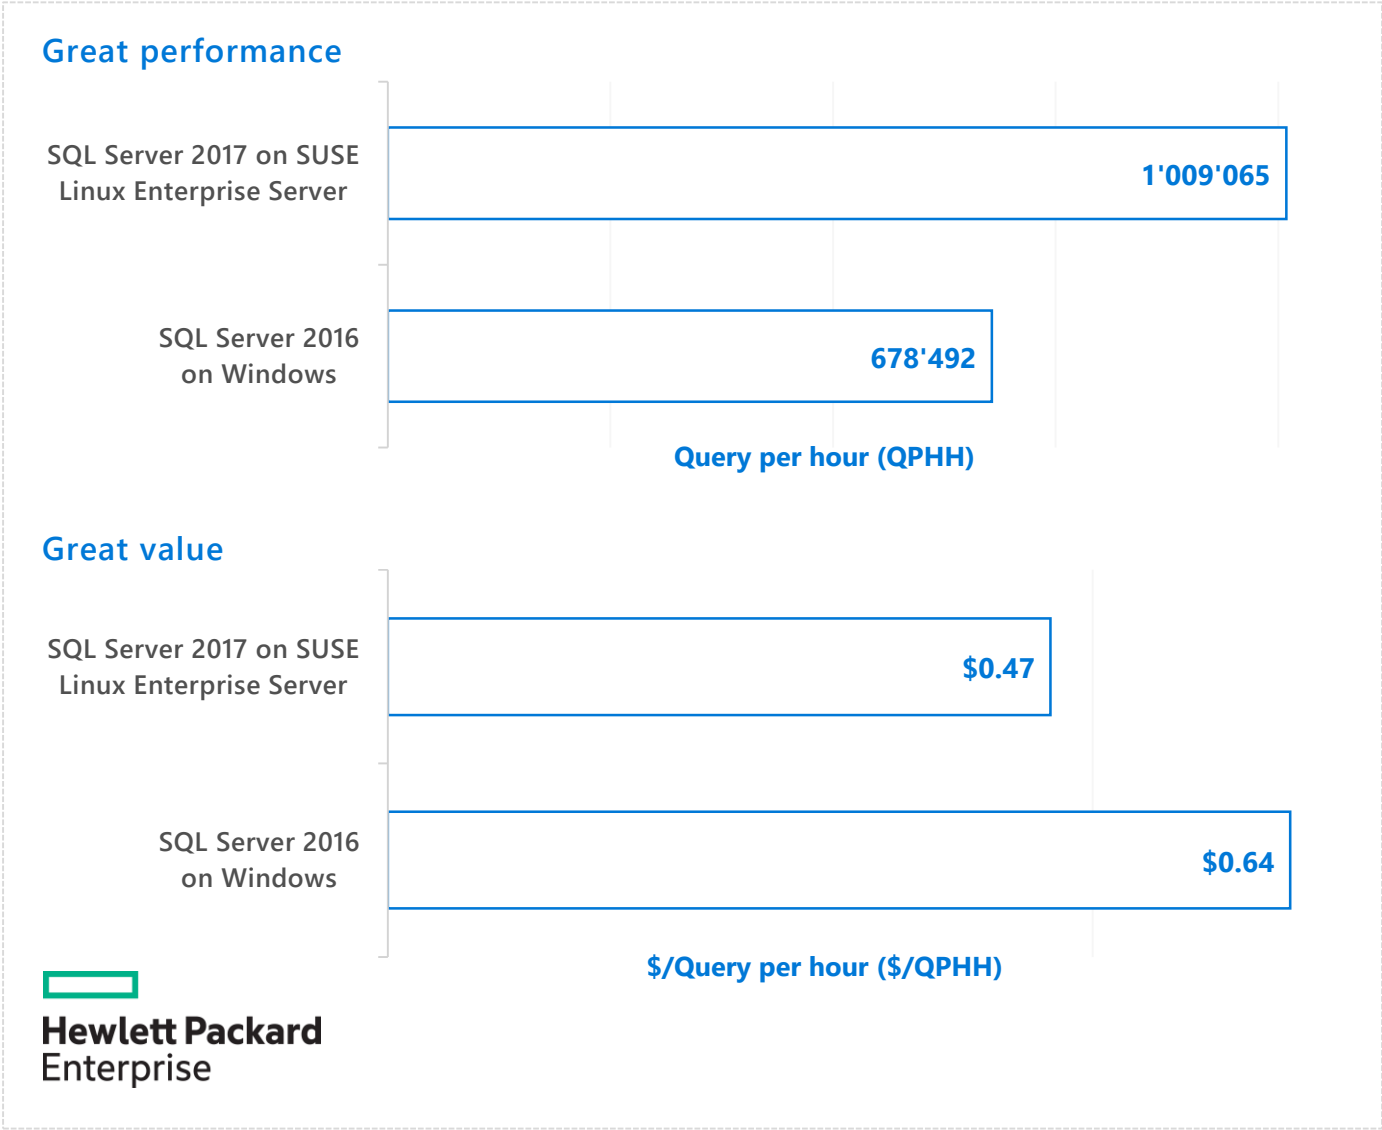
Available on Windows and Linux

#1 price/performance in 1TB, 10TB and 30 TB TPC-H non-clustered benchmark

Reference architectures from major hardware partners

> Industry-leading performance

1TB TPC-H non-clustered world record benchmark



Read the performance brief at hpe.com/servers/benchmarks. Microsoft and Windows are U.S. registered trademarks of Microsoft Corporation. Red Hat, Red Hat Enterprise Linux, and the Shadowman logo are registered trademarks of Red Hat, Inc. Linux is a registered trademark of Linus Torvalds. Intel and Xeon are trademarks of Intel Corporation in the U.S. and other countries. TPC and TPC-H are trademarks of the Transaction Processing Performance Council. TPC-H results show the HPE ProLiant DL380 Gen10 with a result of 1,009,065 QphH @ 1000GB and \$0.47/QphH USD with system availability as of 02-28-18 (results published 11-17-2017; see www.tpc.org/3331); the HPE ProLiant DL380 Gen9 with a result of 678,492 QphH @1000GB and \$0.64/QphH @ 1000GB with system availability as of 07-31-2016 (results published 03-24-2016; see tpc.org/3320). The TPC believes that comparisons of TPC-H results published with different scale factors are misleading and discourages such comparisons. Please see tpc.org for up-to-date information. Competitive claims valid as of 04-19-2017.

BUSINESS INSIGHTS IN 1/5TH THE TIME

Mission critical data warehousing

CHALLENGE

High costs, limited performance, and slow reporting pushed UAE health insurer to consider enterprise DW migration

IMPACT

Queries **64% faster** than Sybase, **98% faster** than Oracle, and **13% faster** than SAP HANA

Reports are **5x faster** and built-in security **saved \$250k** in expenses



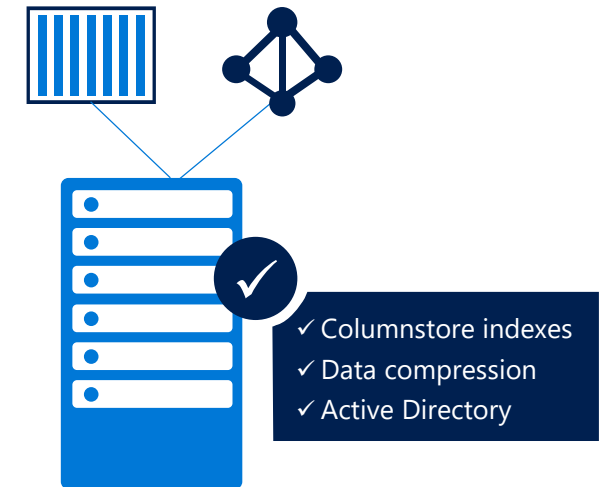
Before



Sybase DW

Monthly reports
in 48 hours

SQL Server DW



SQL Server DW

Monthly reports
in 10 hours

DEMO

Clustered Columnstore Index (CCI) on RHEL

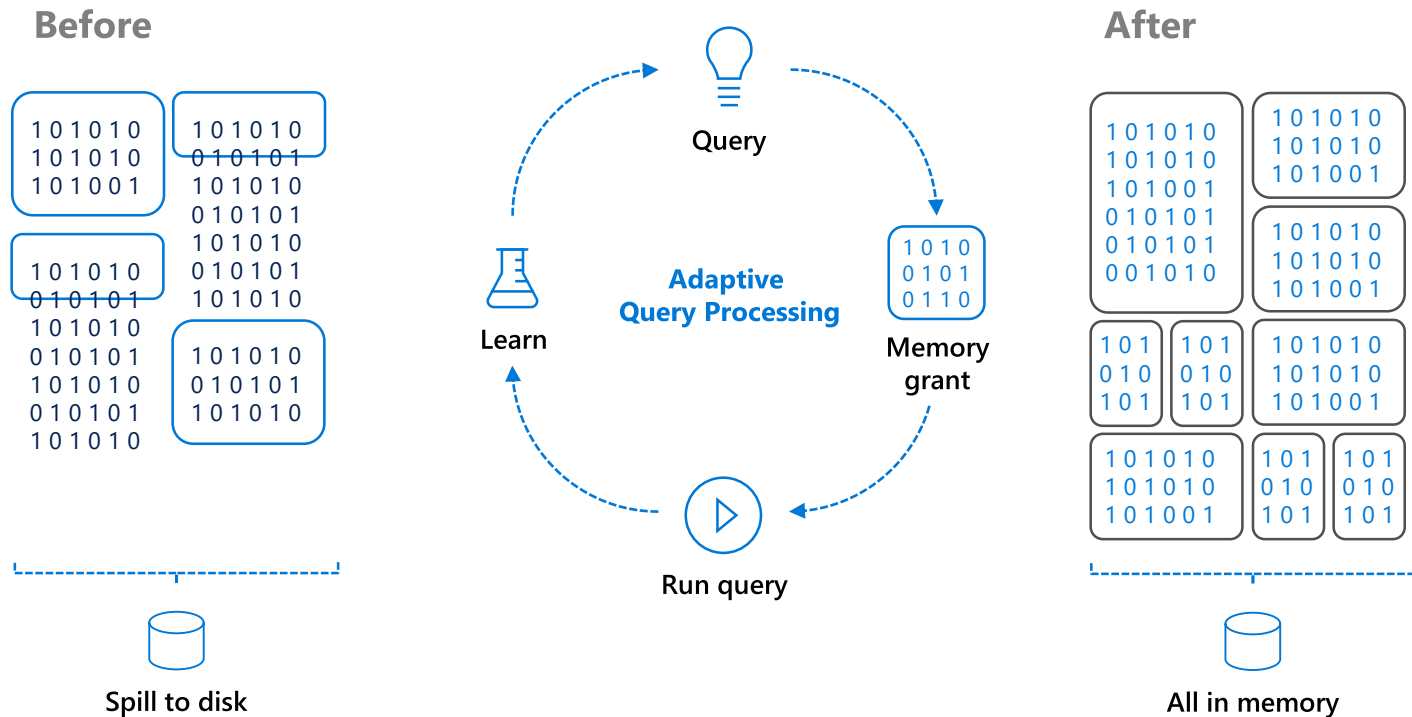
OPTIMIZED QUERY PROCESSING

Improved efficiency with Adaptive Query Processing

Optimize memory grants for repeatable queries to avoid over or under allocating

Adjust data join strategy for small or large tables to speed joins

Batch mode for memory grant feedback and adaptive joins




AUTOMATICALLY FIX PROBLEMS WITHOUT TUNING

Better performance with Automatic Plan Correction

Continuous performance plan monitoring and analysis

Detect problematic plans

Automatically fix performance problems caused by SQL plan choice regressions



DEMO

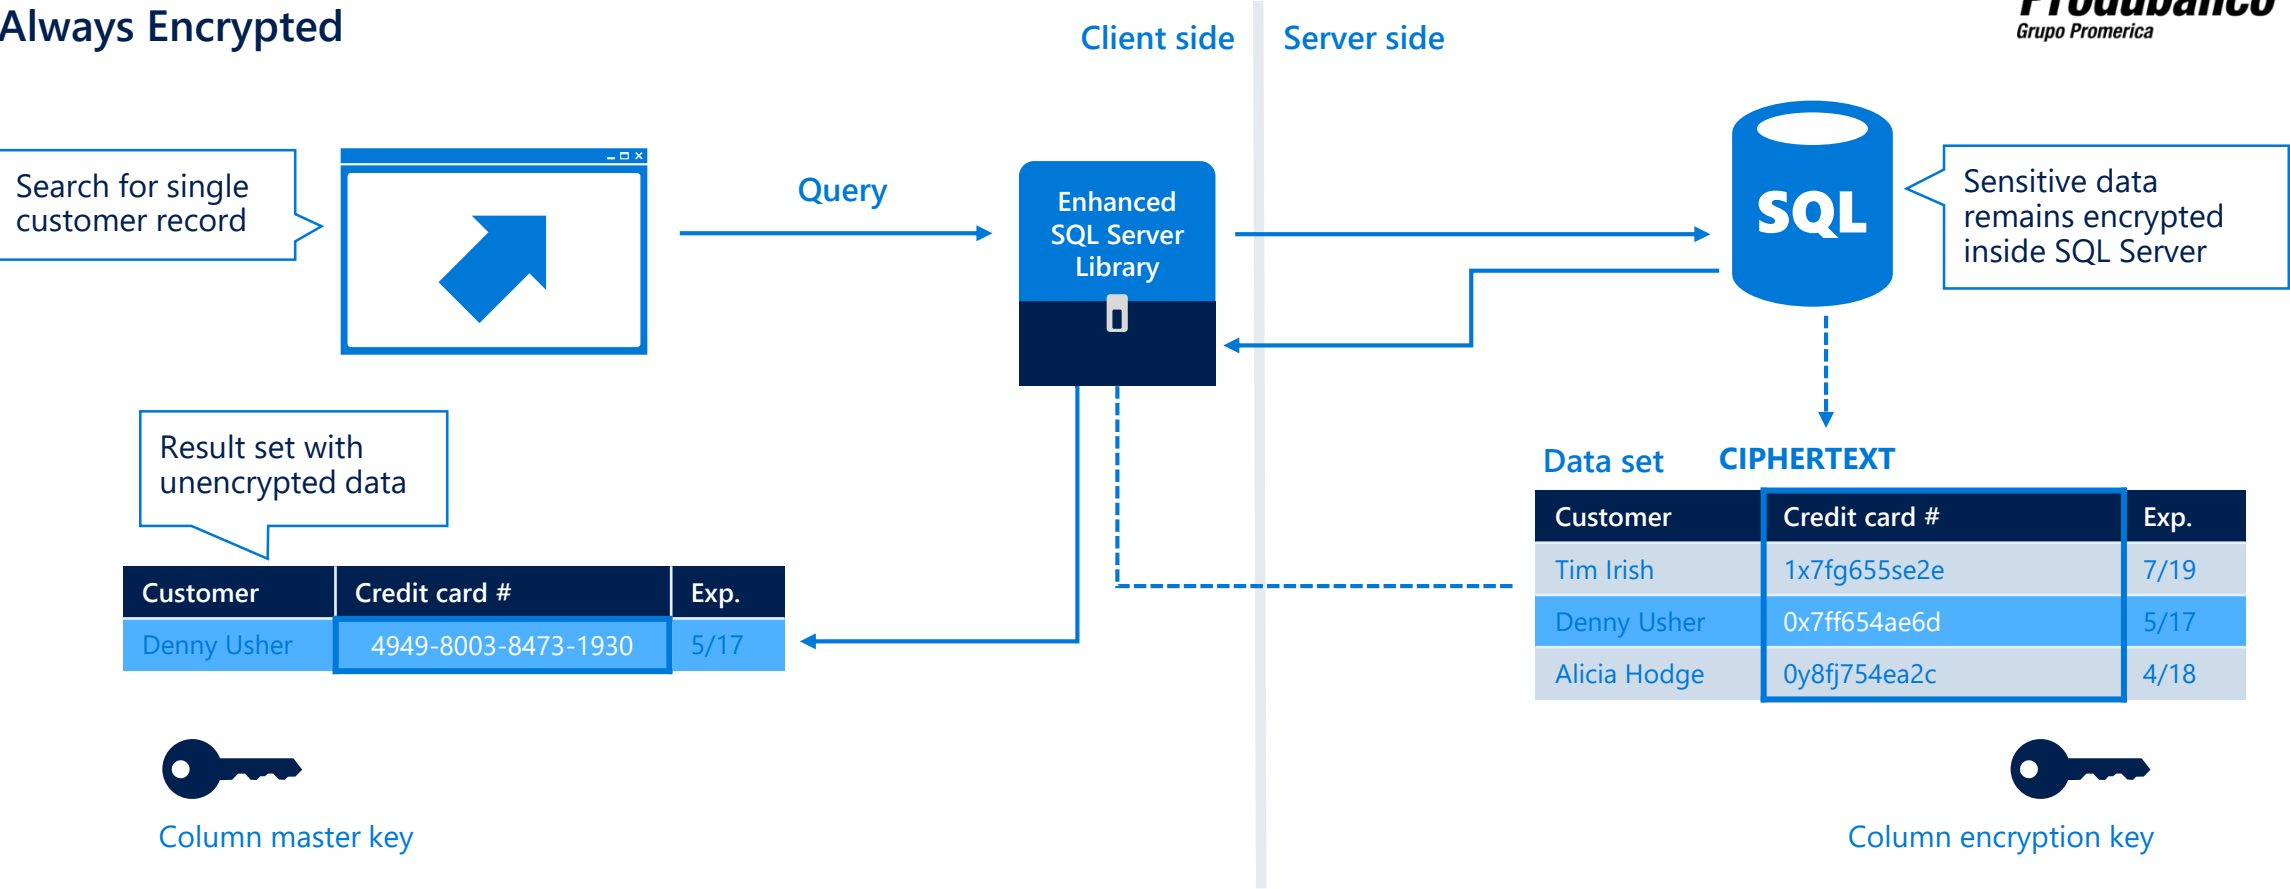
Auto tune performance demo



ENCRYPT DATA AT REST AND IN MOTION



Always Encrypted



GENERATE INSIGHTS ACROSS DIVERSE DATA

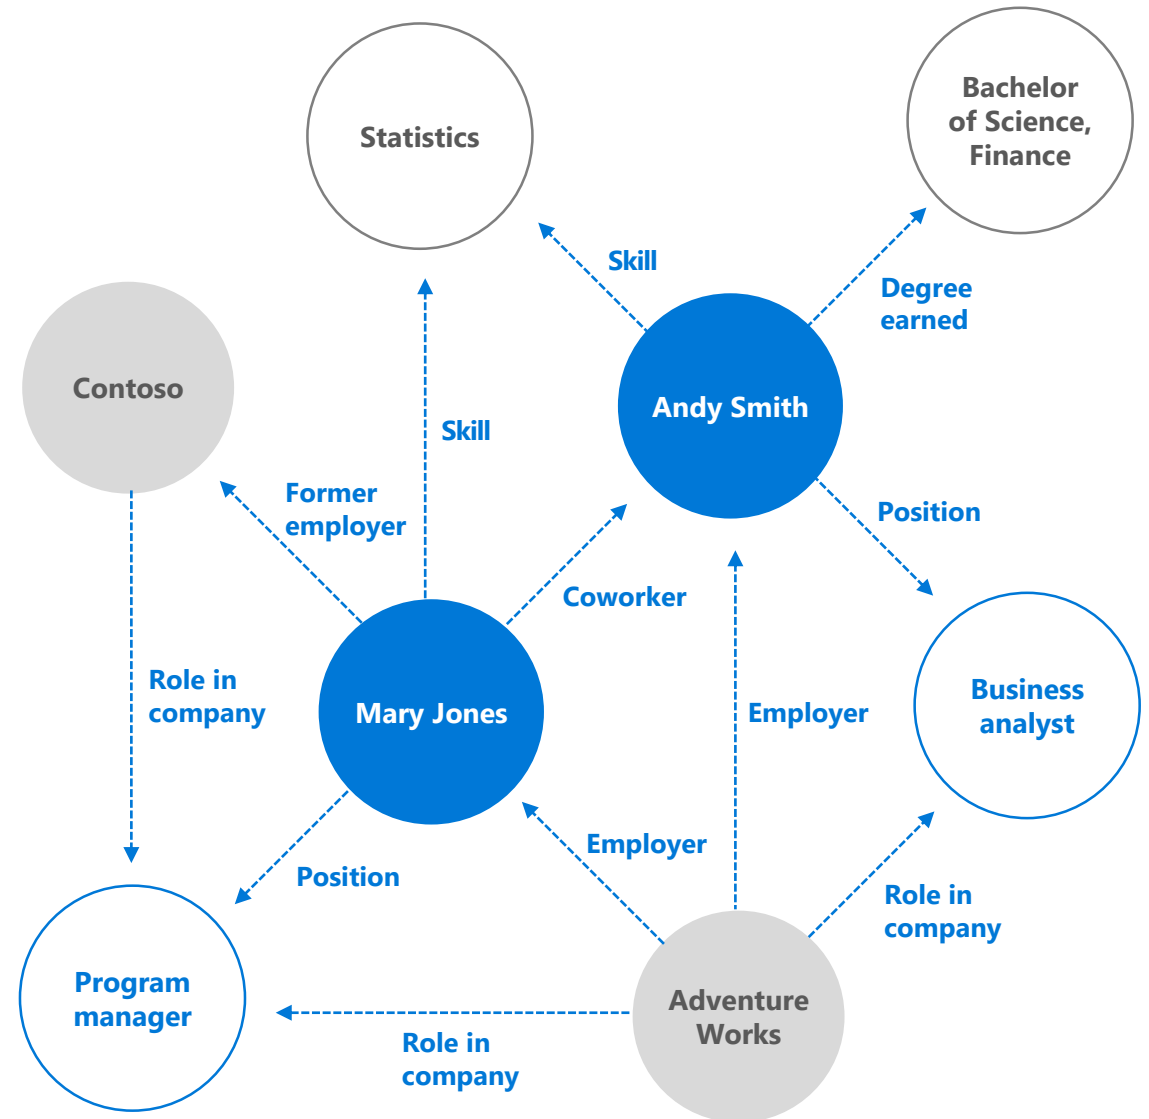
New relationships uncovered with
graph data support

Bring graph data support to your relational
data store

Analyze interconnected data and generate
deeper insights

T-SQL extensions to support pattern matching and
graph traversals

Professional networking app



Machine Learning Services new features

- Python support
 - Python and R scripts are now supported
 - Revoscalepy—Pythonic equivalent of RevoScaleR—parallel algorithms for data processing with a rich API
- MicrosoftML
 - Package of machine learning algorithms and transforms (with Python bindings), as well as pretrained models for image extraction or sentiment analysis

DEMO

Sentiment Analysis using MicrosoftML Pre-Trained Algorithms on
SQL Server 2017 Machine Learning Services

WHAT'S NEXT FOR SQL SERVER ON LINUX

SQL SERVER vNEXT INVESTMENT THEMES

Reason over any data, anywhere

Deeper SQL Server integration with Big Data
Connectors for Data Virtualization across platforms (PolyBase v2)
SQL Graph enhancements

Choice of language and platform

Additional container scenarios
More supported application platforms
Additional external script language extensibility

Industry leading performance and security

Continuing enhancements to industry-leading performance
More intelligent database with additional Adaptive Query Processing features
Additional support for GDPR compliance

Only commercial database with AI built-in

Machine Learning Services on Linux and in containers

Continued improvements to SQL Server on Linux

Merge and transactional replication
Local logins
Bash scripting jobs in agent



Please Complete your Session Evaluations

Get your cool IoT Dev Kit!

Fill out your feedback form and turn it in before you leave.



The background features a dark blue gradient with a glowing horizon line. A large, semi-transparent globe is centered, composed of a network of white dots connected by thin white lines, creating a mesh-like structure. The text "Thank you" is overlaid in the center in a large, white, sans-serif font.

Thank you

