Microsoft

# The SQL Server 2022 Workshop

Microsoft Azure Data


#sqlserver2022

#azuresql

# Goals for this workshop

- Learn in an interactive style **what** is new in SQL Server 2022
- Evaluate **why** you should upgrade to SQL Server 2022
- Learn **how** specific key features work in SQL Server 2022
- Gain **hands-on** experience trying out new features
- Walk away with **resources** you can use to learn more

# The SQL Server 2022 workshop


Module 1
**Introduction to SQL Server 2022**

Module 2
**Connect SQL Server 2022 to Azure**

Module 3
**Accelerate performance with built-in query intelligence**

Module 4
**Power your database with security, scalability, and availability**

Module 5
**Access new sources with data virtualization and object storage**


Module 6
**Enhance your application with new T-SQL capabilities**

Module 7
**Summary and closing**

aka.ms/sql2022workshop
aka.ms/sqlserver2022decks
aka.ms/sqlserver2022series
aka.ms/sqlserver2022demos


Module 1

# Introduction to SQL Server 2022

Overview of SQL Server 2022

Getting started with SQL Server 2022

Knowledge checks and summary

# Introducing SQL Server 2022

# The next step for SQL Server

## SQL Server 2016

Query Store

Polybase

Always encrypted

Row level security

It just runs faster

Std Edition surface area

## SQL Server 2017

SQL Server on Linux

Containers

Adaptive query processing

Automatic Tuning

Graph database

Machine learning services

## SQL Server 2019

Data virtualization

Intelligent query processing

Accelerated database recovery

Data classification

# Microsoft SQL – all edges and clouds

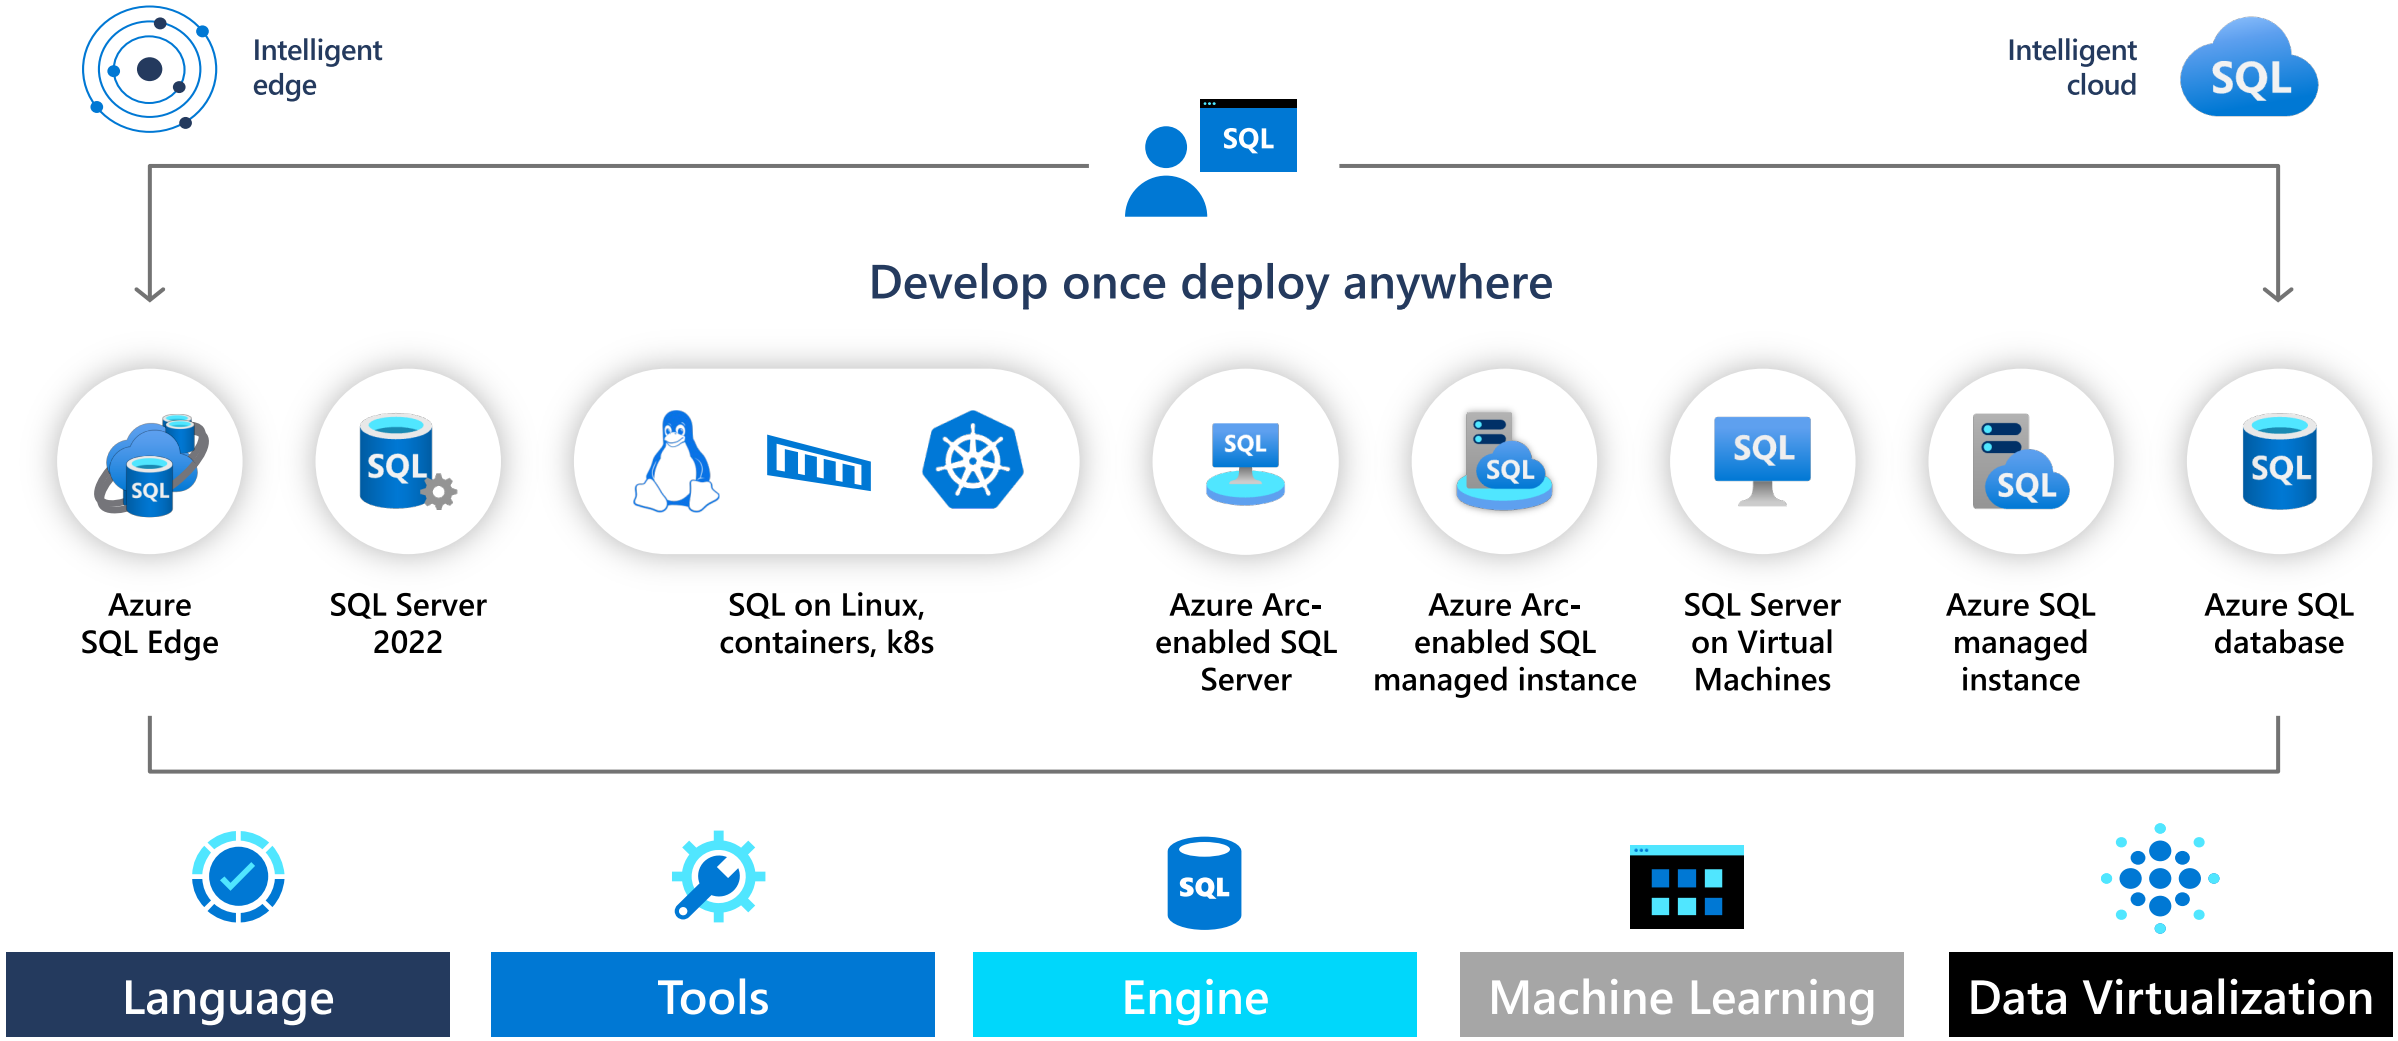

# SQL Server 2022

## Azure-enabled with continued performance and security innovation

### Business continuity through Azure

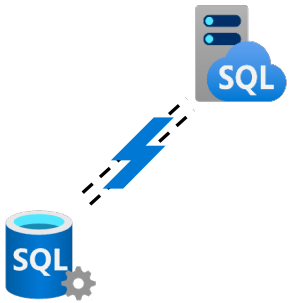


### Seamless analytics over on-prem operational data

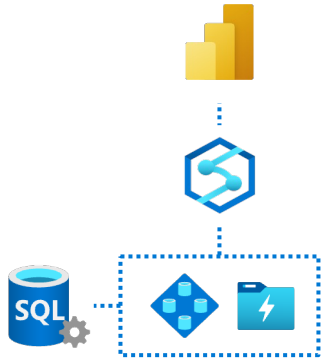


### Visibility over your entire data estate

- Microsoft Purview
- Insights
- Lineage
- Policies

### Most secure over the last 10 years[1]

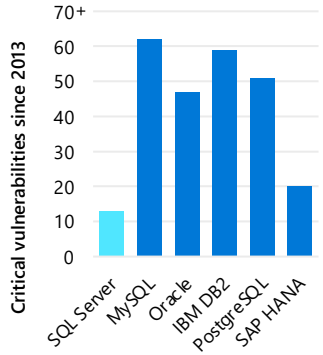


### Industry-leading performance and availability

#1 OLTP performance[2]

#1 Non-Clustered DW performance on 1TB[3], 3TB[4], 10TB[5], 30TB[6], and 100TB[7]

| Hero capabilities | Link feature for Managed Instance | Azure Synapse Link | Microsoft Purview integration | Ledger for SQL Server | Built-in query intelligence |
|---|---|---|---|---|---|

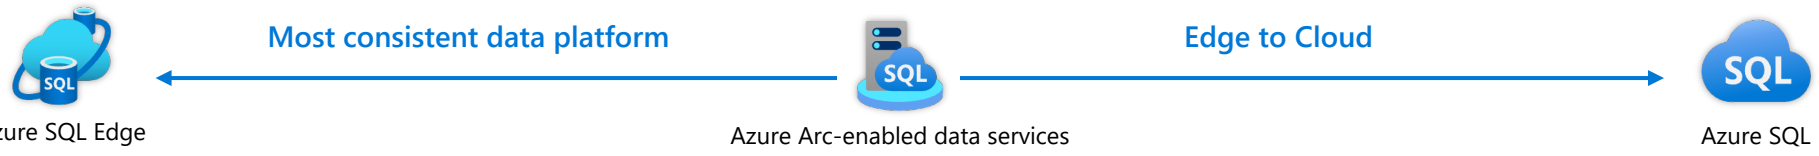


All TPC Claims as of 11/10/2022.

[1]National Institute of Standards and Technology Comprehensive Vulnerability Database [2] http://www.tpc.org/4087; [3] http://www.tpc.org/3386; [4] http://www.tpc.org/3385; [5] https://tpc.org/3383; [6] http://www.tpc.org/3364; [7] https://tpc.org/3382;

# SQL Server 2022 capabilities

BanQsoft

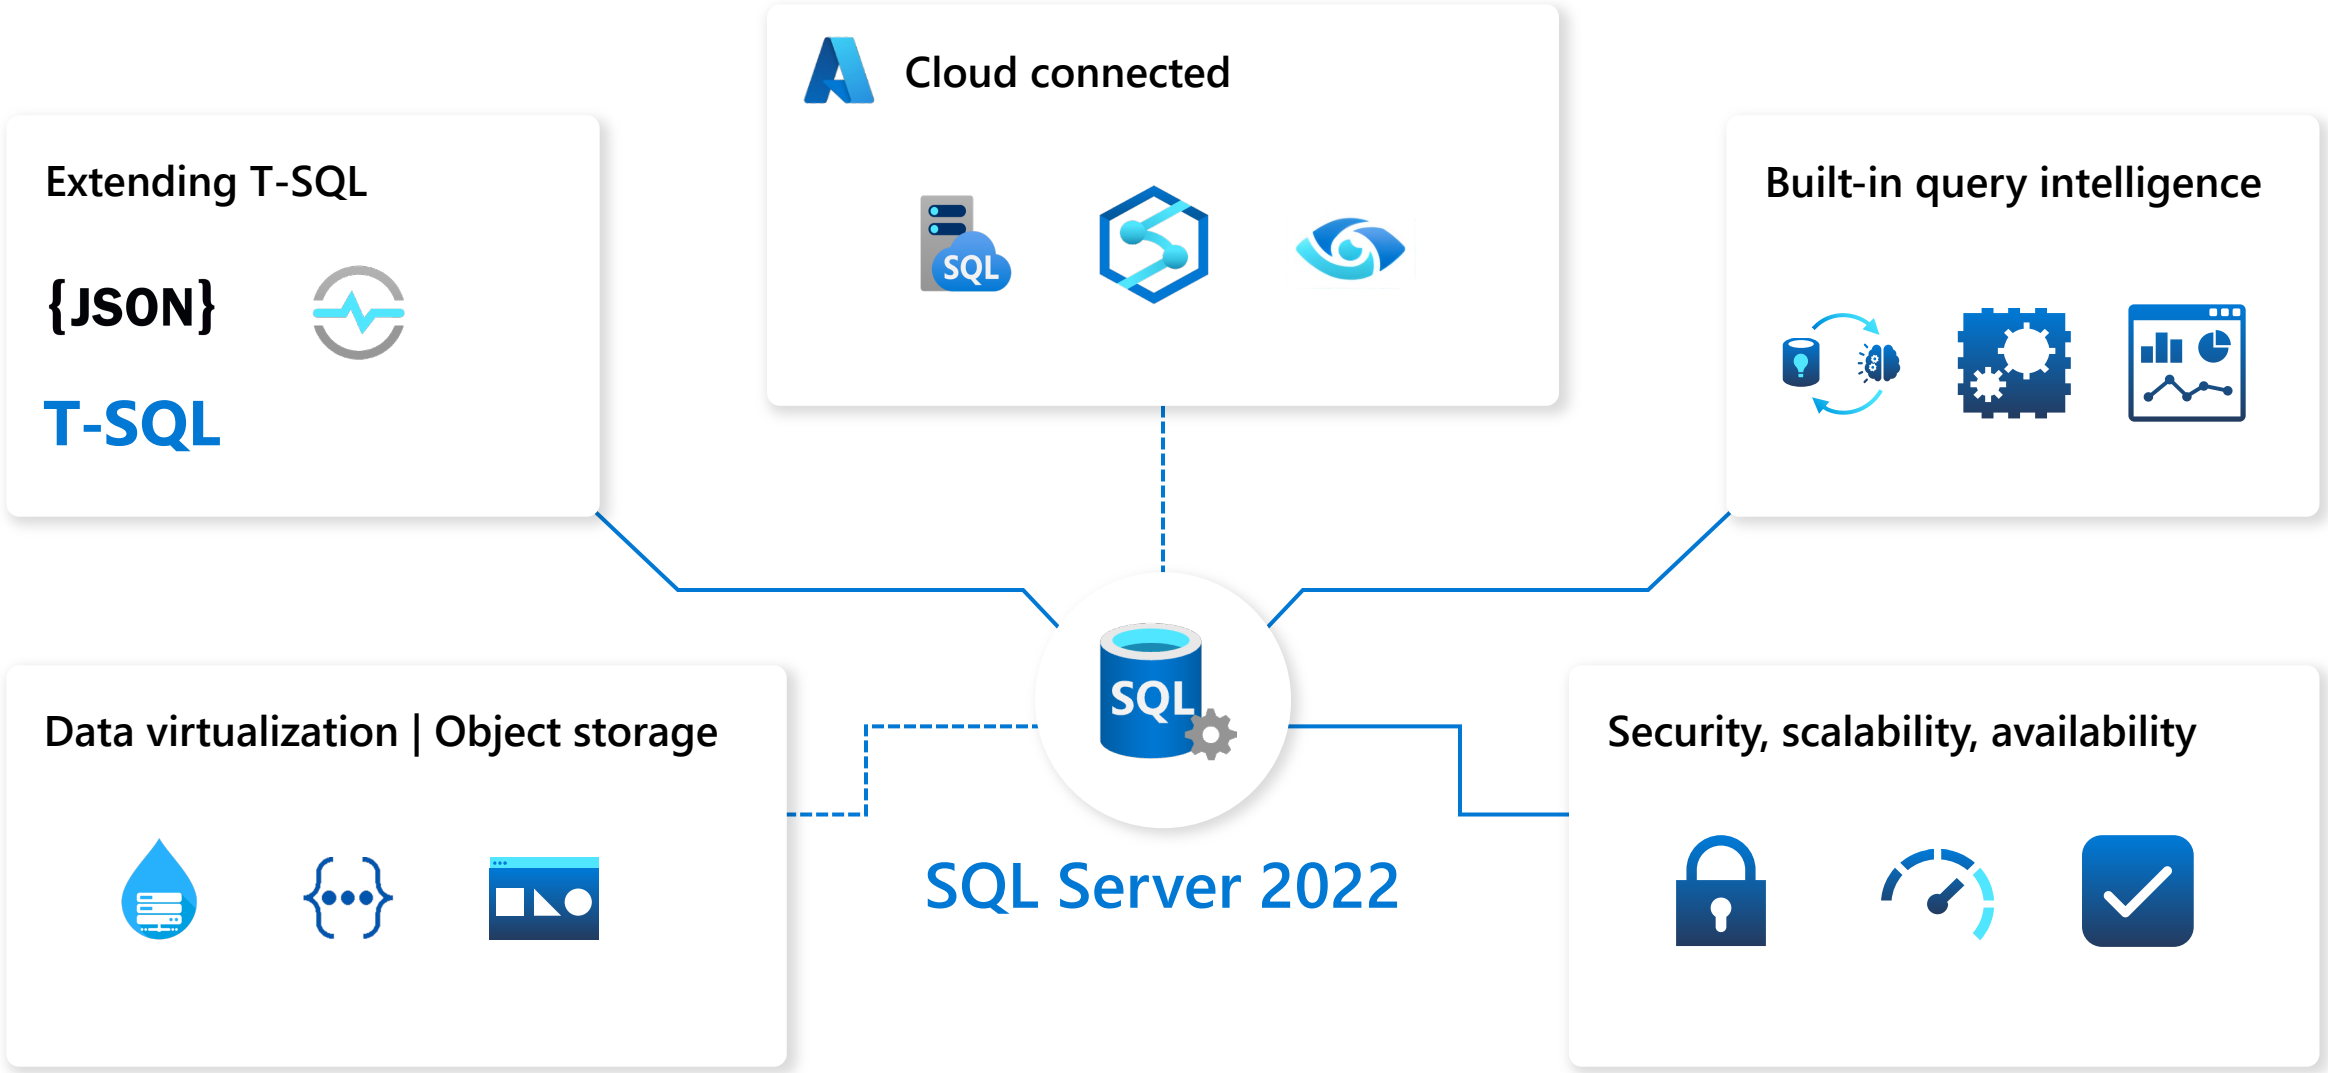

# Deploying SQL Server 2022


SQL Server Instance

Supported OS

Same upgrade options as in previous releases

dbcompat 160

older dbcompat supported

Side by side

Windows GUI or command line

RHEL, Ubuntu, SLES package managers

Pre-installed container images

Azure marketplace images for Virtual Machines

# What's different about SQL Server 2022 setup?

## Removed

R, Python, and Java packages[1]

Polybase Hadoop connectivity with Java

Polybase scale out groups

Machine learning server

Distributed replay[2]

## Deprecated features

Stretch database

## Added

Azure extension for SQL Server

Pay As You go Licensing

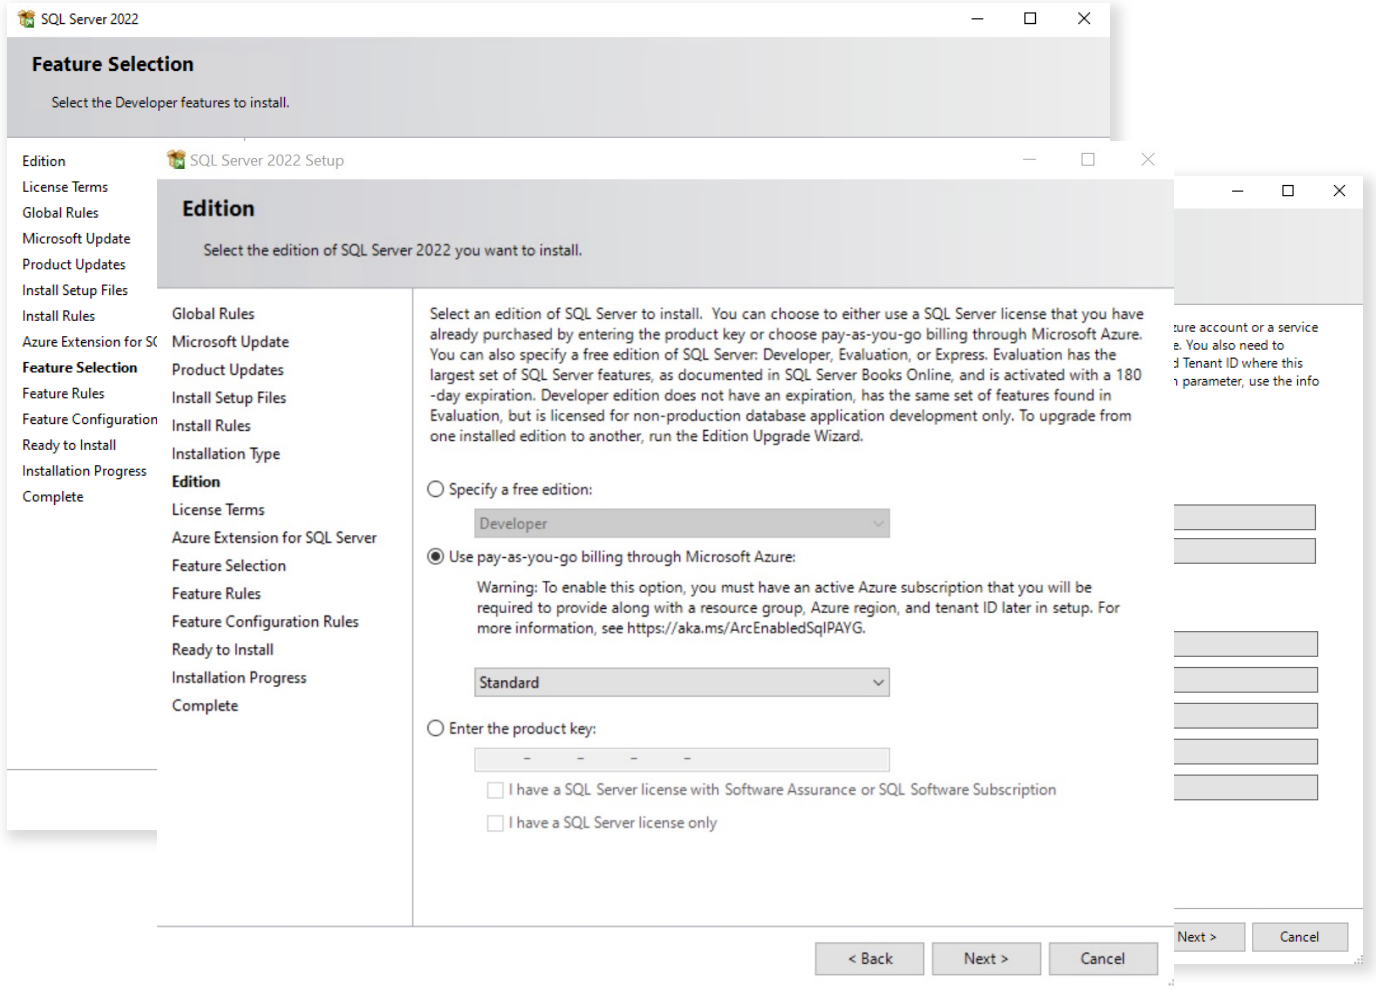


[1] Bring your own runtime packages [2] Will be available for separate download

# Azure extension for SQL Server

## Registers Azure Arc-enabled SQL Server

## Capabilities for SQL Server 2022

AAD authentication

Microsoft Purview access policies

Microsoft Defender (supported by Monitoring Agent extension)

Inventory, Pay as You Go Licensing, BPA, Automated backups

## Extension under the Azure Arc Agent framework

"Mimics" IaaS extension for Azure VM

Uses REST API to communicate with Azure

Supports proxy

Service on Windows and daemon process on Linux

Configure during SQL Server 2022 setup or after with scripts

SQL Server Config Manager to stop or start extension service

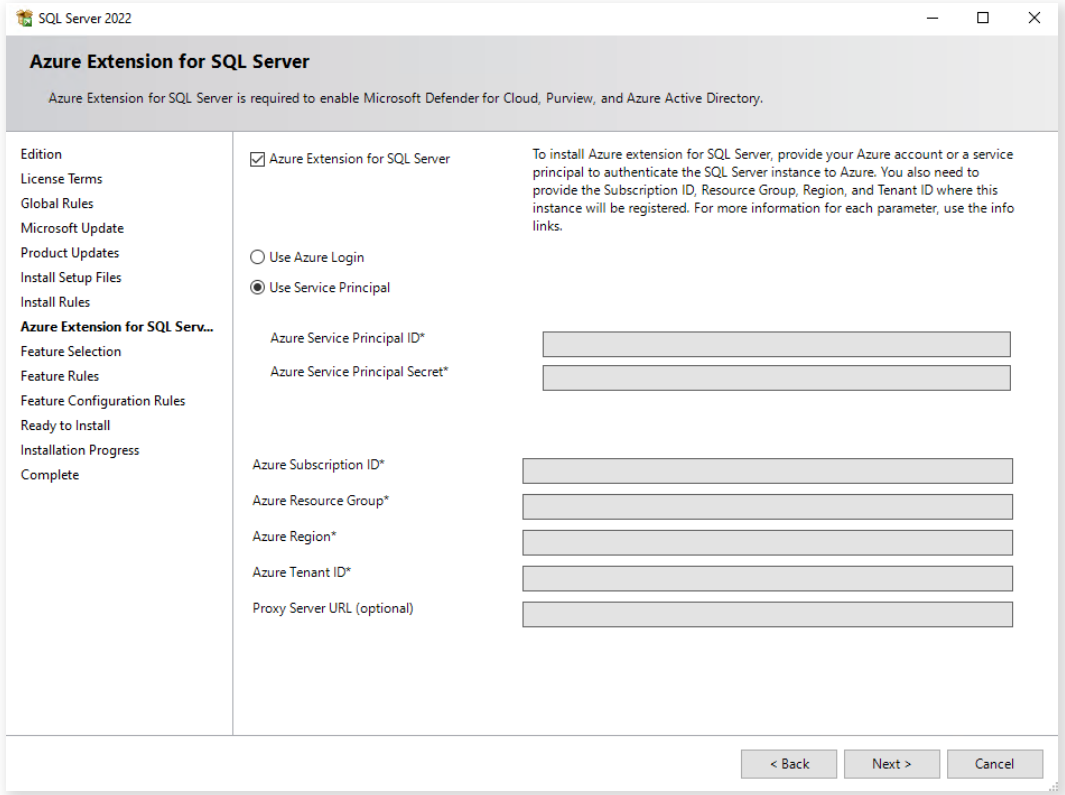

# SQL Server 2022 editions

https://aka.ms/sql2022editions

## Azure-enabled with continued performance and security innovation

**Express**
Free, entry-level database for small web and mobile apps

**Feature highlights**

- Up to 4 cores of CPU
- Up to 1410 MBs of memory
- Microsoft Purview Policies
- Azure AD authentication
- Built-in query intelligence: PSP Optimization, Optimized plan forcing
- Query store on by default for new databases
- Data Lake Virtualization
- Ledger
- Timeseries support

**Standard**
Full featured database with for mid-tier applications and data marts

**Feature highlights**

- Up to 24 cores of CPU
- Up to 128 GBs of memory
- Azure Synapse Link for SQL
- Link feature for Azure SQL Managed Instance (basic availability groups)
- Buffer Pool Parallel Scan
- Backups to S3-compatible object storage

+ **Express features**

**Enterprise**
Mission-critical performance and intelligence for tier 1 databases

**Feature highlights**

- Unlimited cores of CPU
- Unlimited memory
- Azure Synapse Link for SQL (multi-threaded snapshot)
- Link feature for Azure SQL Managed Instance
- Built-in query intelligence: DOP feedback, CE Feedback, Memory Grant Feedback
- Contained Availability Group
- AVX 512 extension for batch mode

+ **Standard features**
+ **Express features**

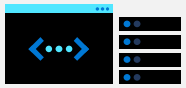

**Developer**
Free to use with all the features of Enterprise Edition specifically for dev/test in non-production environments

Build once and deploy across any SQL Server edition without changing your app

# SQL Server 2022 on Linux, Containers, Kubernetes

## Linux distributions supported

RHEL 8.0 – 8.5

Ubuntu 20.04

SLELS 15

## New DTC transaction management support under WMI

## Features not supported that are new to SQL Server 2022

TLS 1.3

Intel QAT backup compression

## New MCR container images for SQL Server 2022 (2022-latest)

https://aka.mssqllinux
https://aka.ms/sqlcontainers
https://aka.ms/sqlk8s

# SQL Server 2022 on Azure Virtual Machine

## Marketplace images

Windows and Linux

Dev, Web, Std, and EE Editions

## Uses the IaaS Agent Extension

## T-SQL snapshot backups supported with Azure disk snapshots

## Several VM sizes support AVX vector including AVX512

## Portal experience always being updated

## SQL Server 2022 features

AAD authentication

Microsoft Purview policies*

* Planned

https://aka.ms/sqlazurevm

# Knowledge check

**Which of these is one of the major categories for SQL Server 2022 features?**

**A.** Cloud connected

**B.** Azure Arc

**C.** Contained Availability Groups

**D.** Built-in query intelligence

**E.** A and D

**The Azure extension for SQL Server provides which of these capabilities?**

**A.** Synapse Link

**B.** Managed disaster recovery

**C.** AAD authentication

**D.** Data virtualization

**E.** Azure virtual machine

# Summary

SQL Server 2022 is **Azure-enabled** for disaster recovery, analytics, and security

**Built-in query intelligence** reduces query tuning and gains you performance with **no code changes**

SQL Server 2022 has new innovations in **security, scalability, and availability**

**Data virtualization and object storage** through standard SQL interfaces

**Extending the T-SQL language** for new developer scenarios

Module 2

# Connect SQL Server 2022 to Azure

SQL Server 2022 Hybrid

Managed disaster recovery with Azure SQL Managed Instance*

Near-real time analytics with Synapse link*

Azure Arc-enabled SQL Server*

Azure Security hybrid capabilities*

Knowledge checks and summary

Demos*

# Has SQL Server been hybrid?

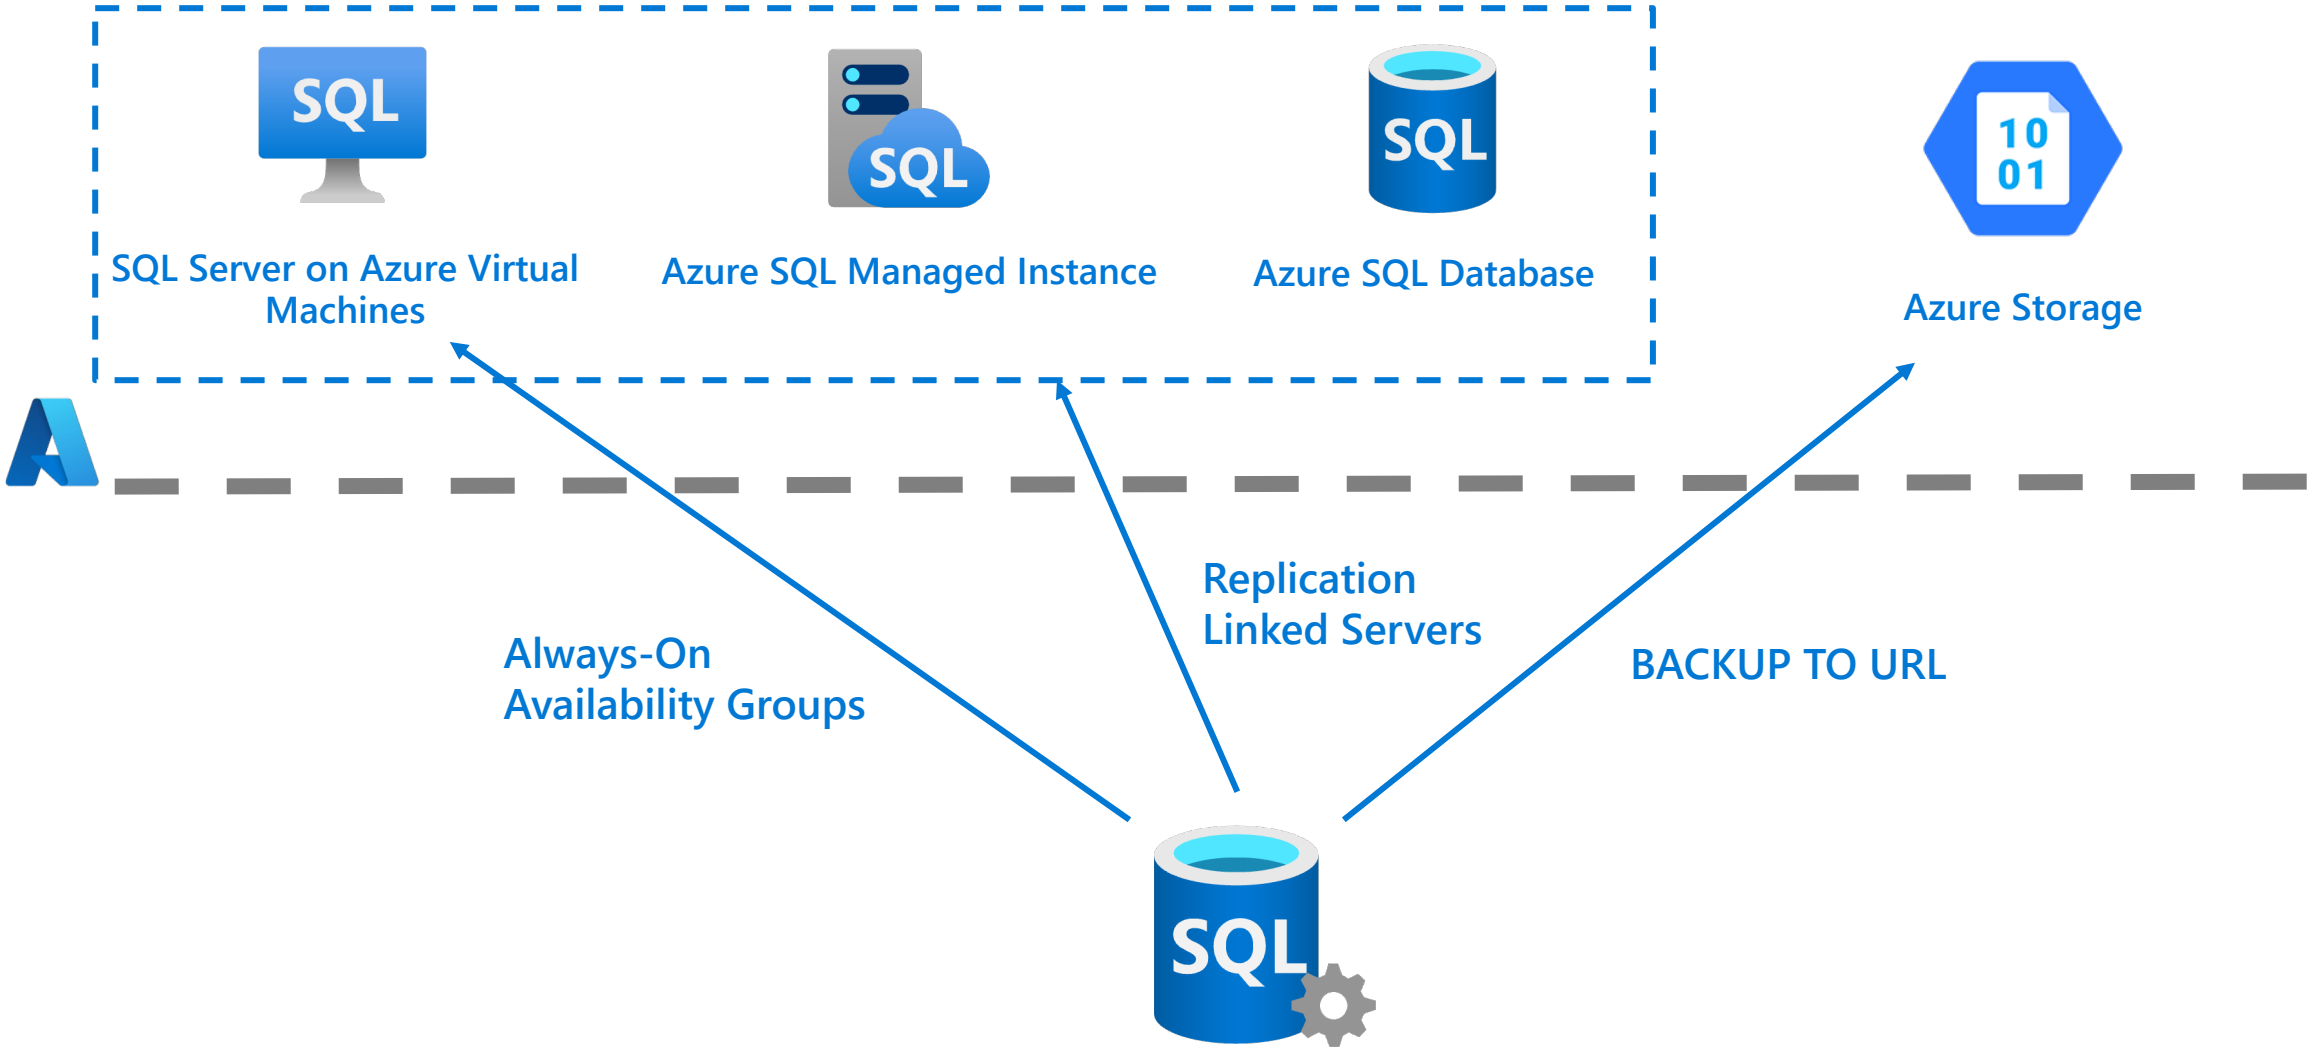

# SQL Server 2022 Cloud connected

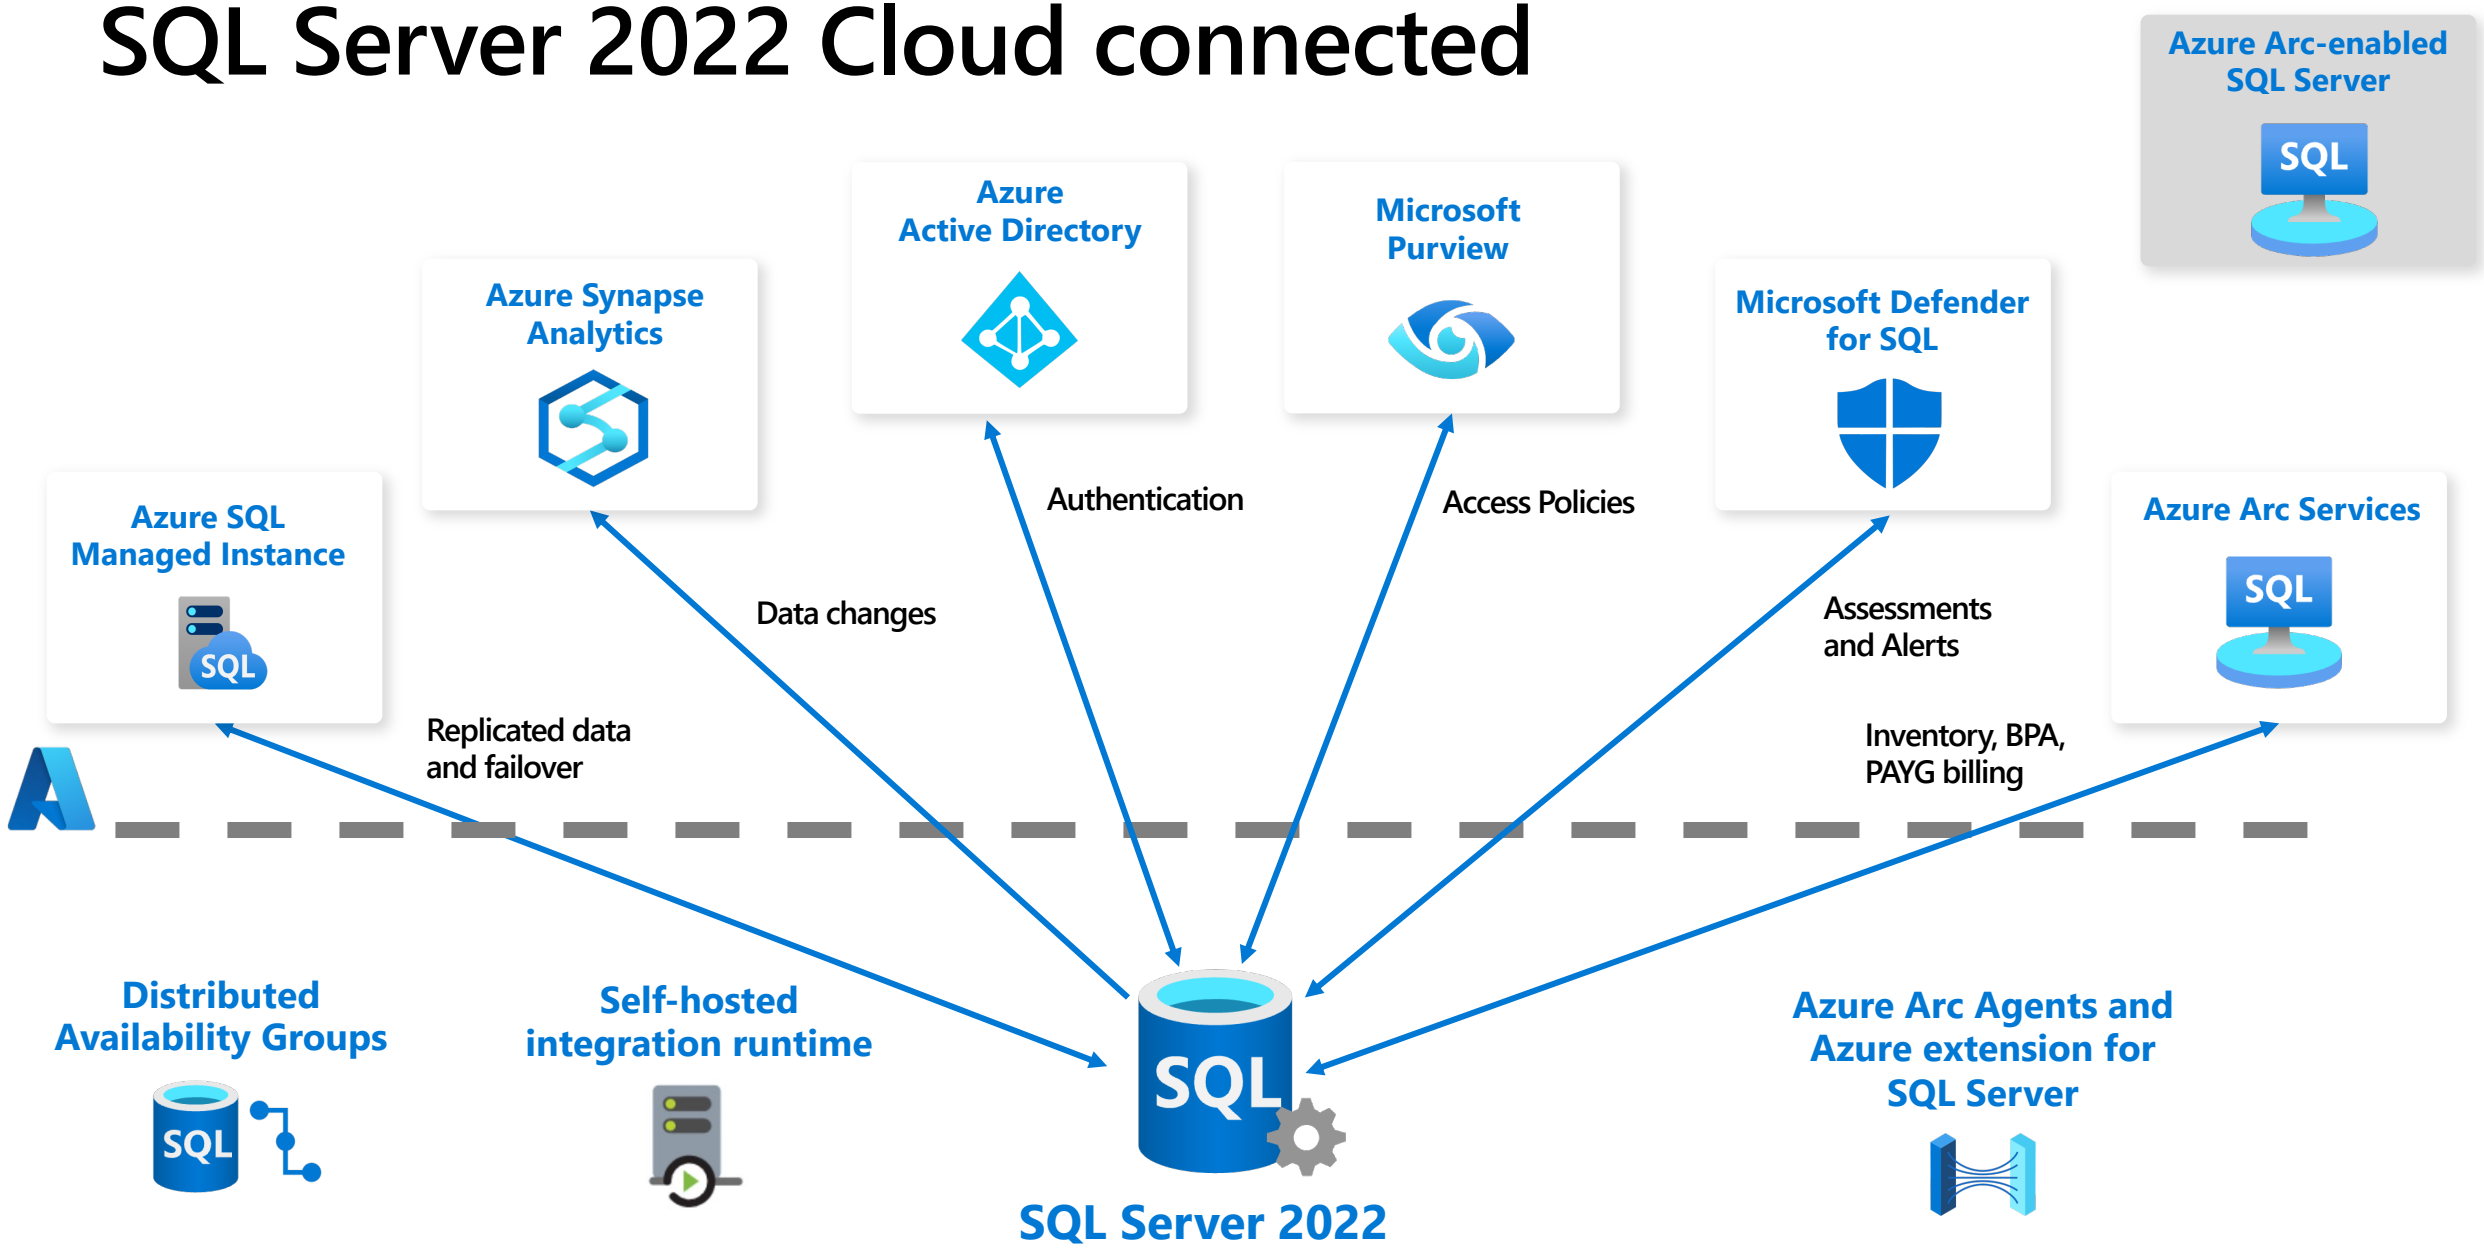

# Business continuity through Azure SQL Managed Instance

## Disaster recovery in the cloud with link feature for Azure SQL Managed Instance

SAXO
BE INVESTED

### Challenge: Difficult to setup and maintain a DR site

- Use PaaS for managed disaster recovery
- Don’t have an AG? We will build one for you with no replica or clustering required
- Built-in distributed availability group (DAG)
- Offline Disaster Recovery using backup/restore
- Online Disaster Recovery in preview

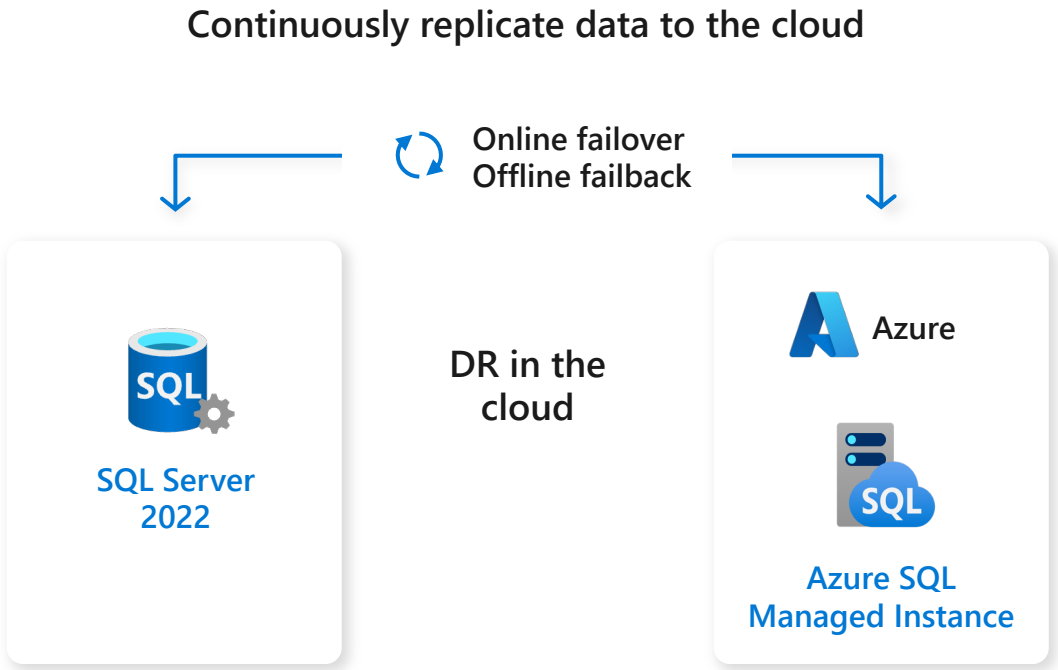


License free passive DR

https:/aka.ms/milink

# SQL Server Distributed Availability Groups (DAG)

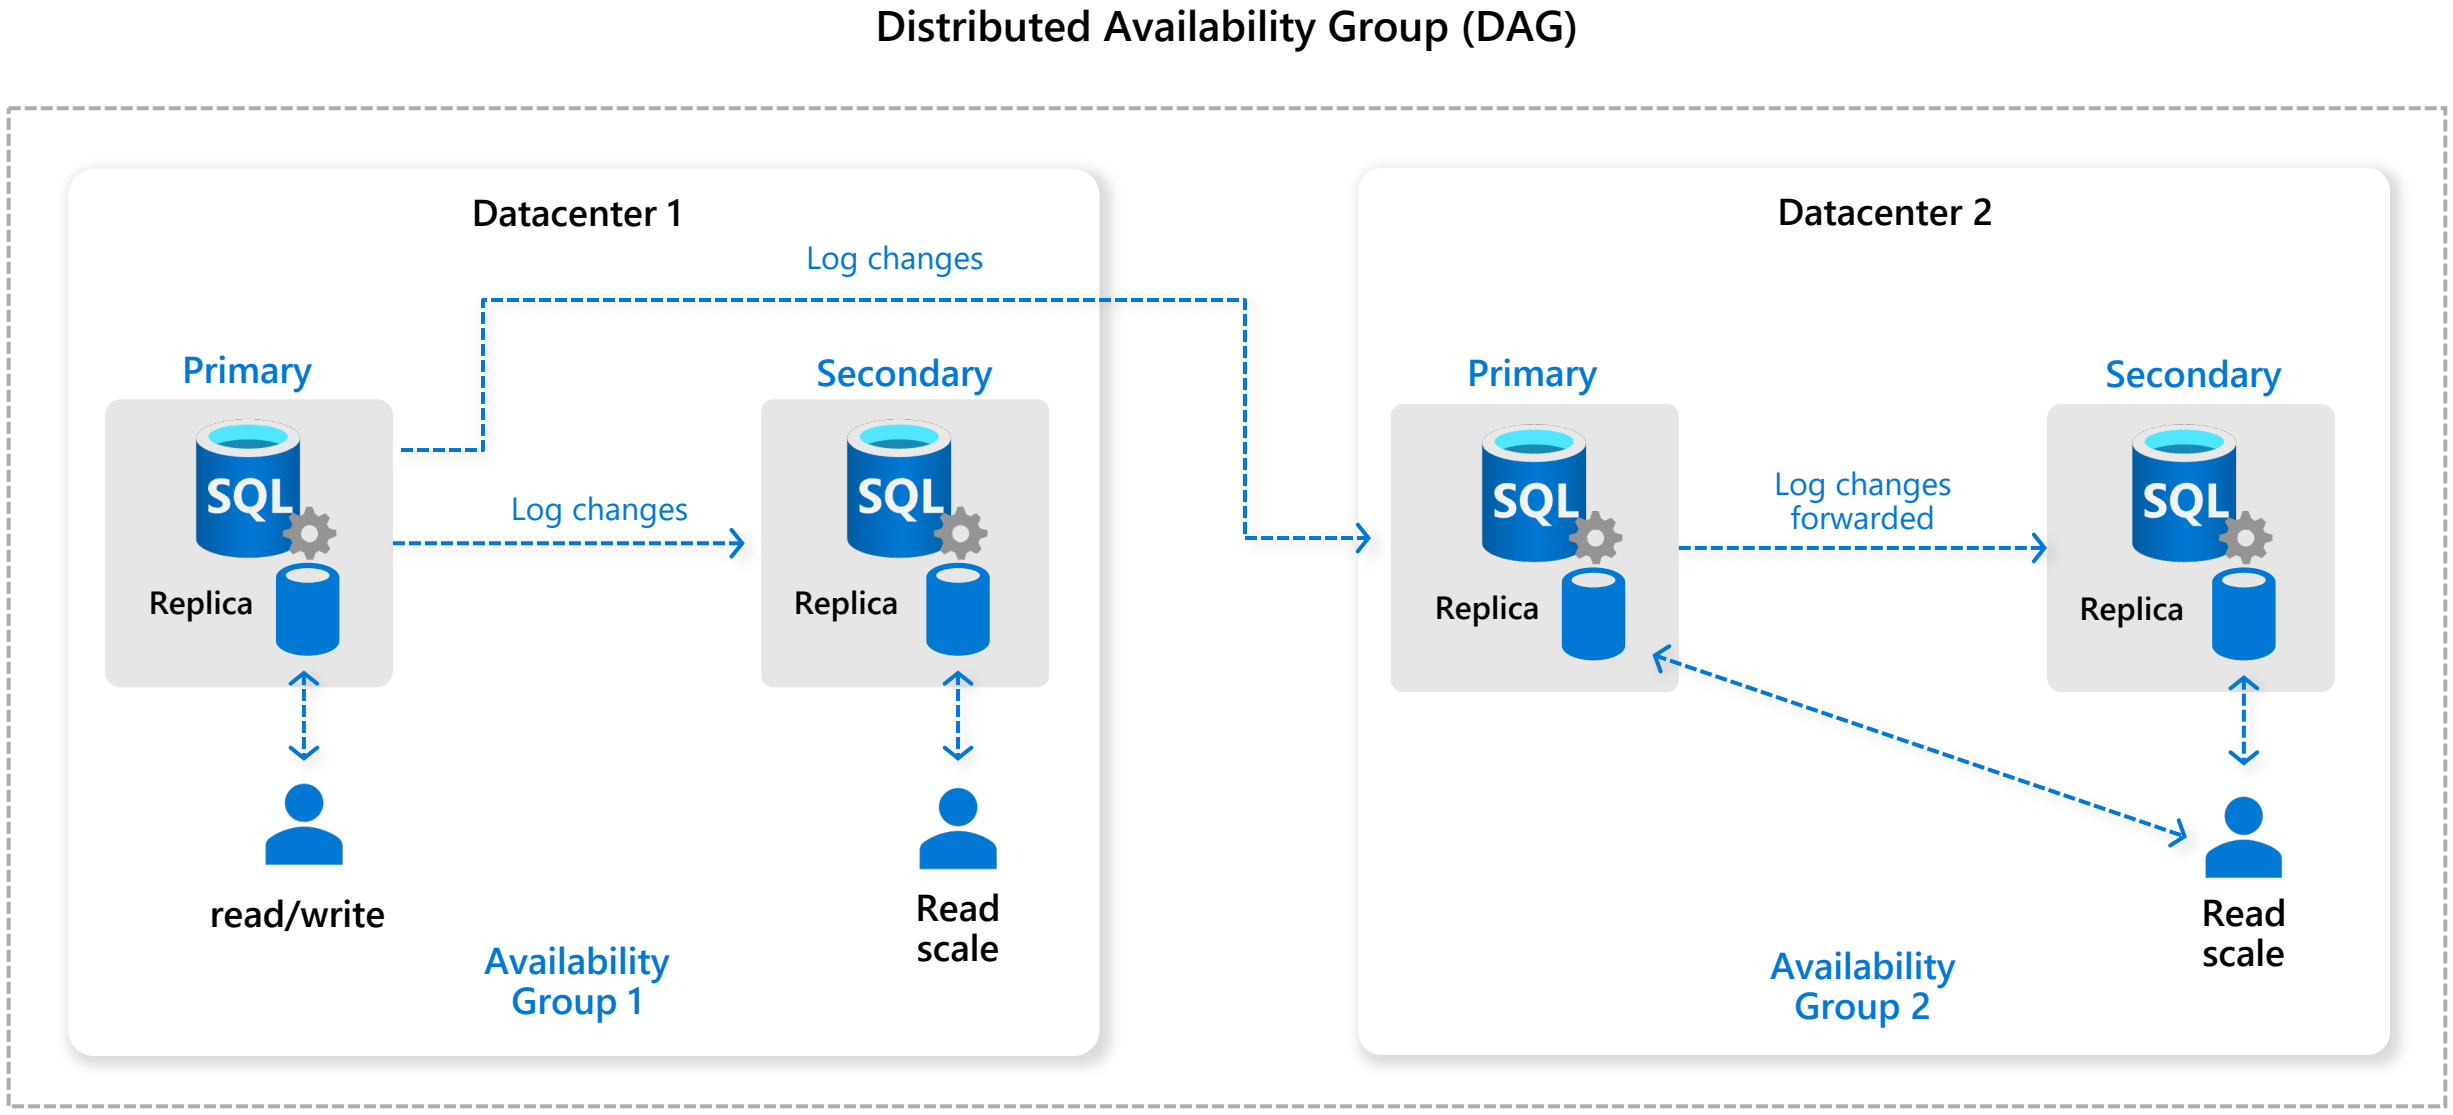

# Link feature for Managed Instance

SQL Server 2016+ → Migration
SQL Server 2022 → DR

Distributed Availability Group (DAG) async 3

Single step with SSMS
using the replicate wizard

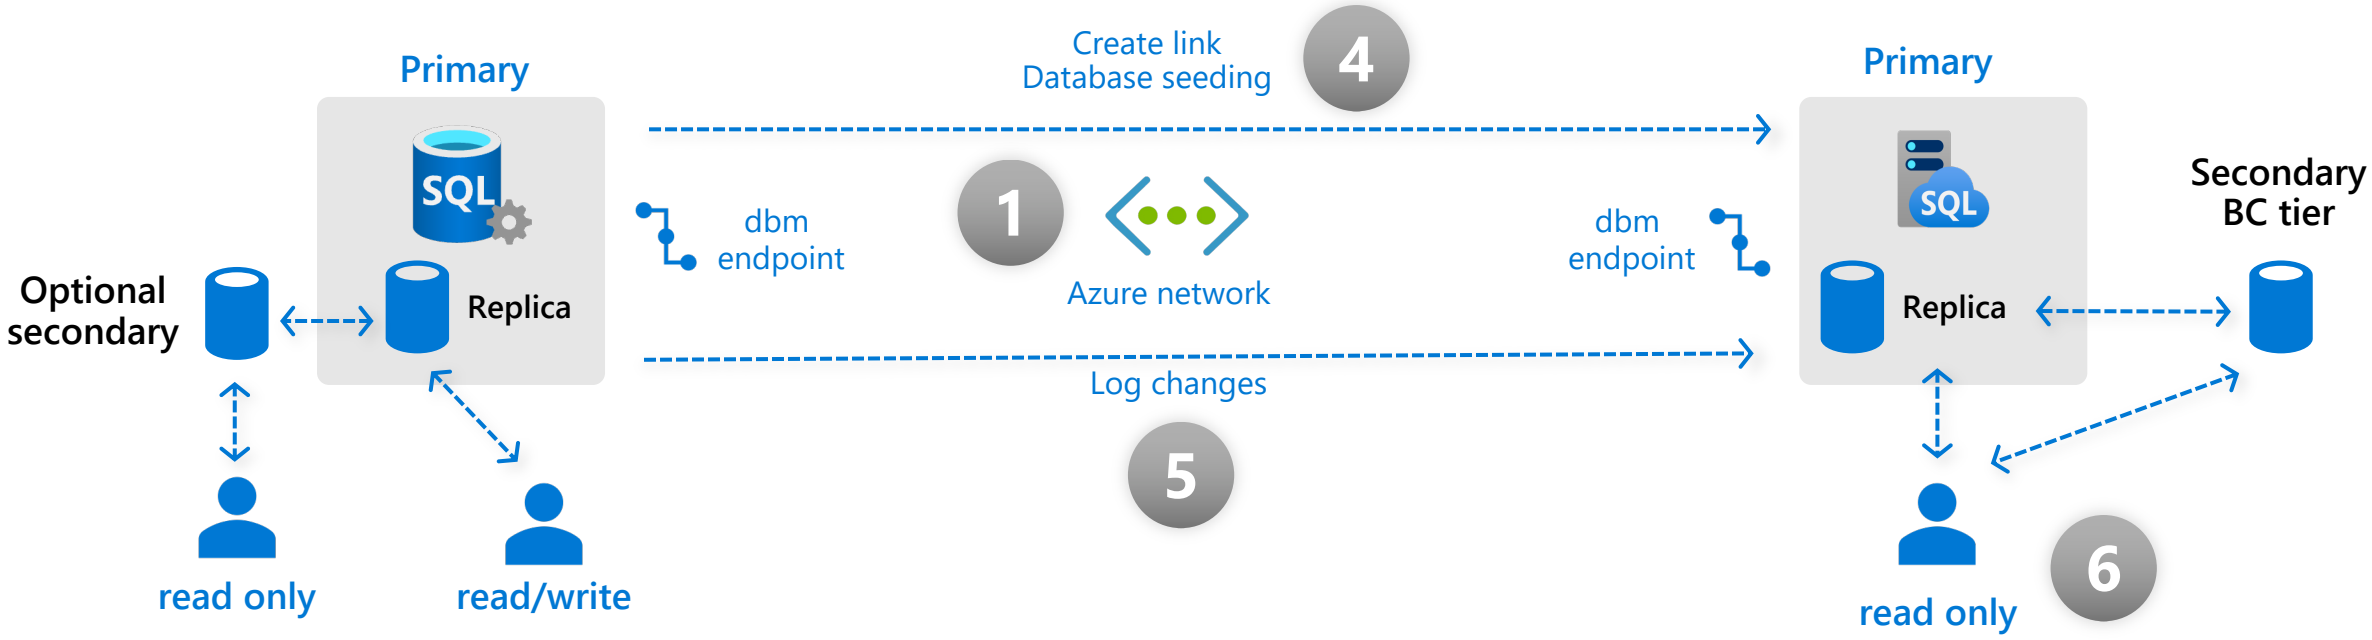

# Online failover to Azure SQL Managed Instance

Single step with SSMS
using the failover wizard

1

## Planned failover

- Stop all writes to database
- Switch mode for MI from async to sync
- Switch DAG to sync
- Verify LSNs match
- Remove link to MI
- Optionally clean up DAG and AG

## Forced failover

- Remove link to MI
- Clean up DAG and AG when SQL Server is available

2 **Database on Azure SQL Managed Instance becomes read/write**

3 **Redirect application to connect to Azure SQL Managed Instance**

4 **Migrate instance level objects to Azure SQL Managed Instance**

# SQL Server database version compatibility

SQL Server 2019
Version 15.X

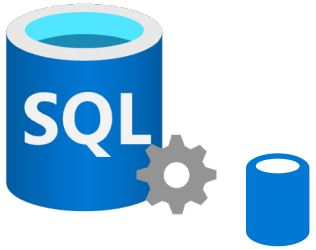


database
version 904

database version != dbcompat

Database 'WideWorldImporters' running the upgrade step from version 904 to version 9XX

Azure SQL
Managed Instance

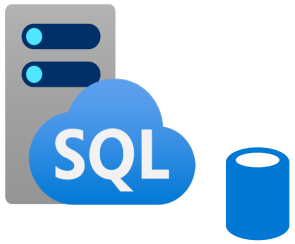


database
version 957

Locked in with
SQL Server 2019 RTM

Doesn't change for CUs

The database was backed up on a server running version 16.00.0312. That version is incompatible with this server, which is running version 15.00.2095. Either restore the database on a server that supports the backup,or use a backup that is compatible with this server.

# Offline Disaster Recovery with SQL Server 2022

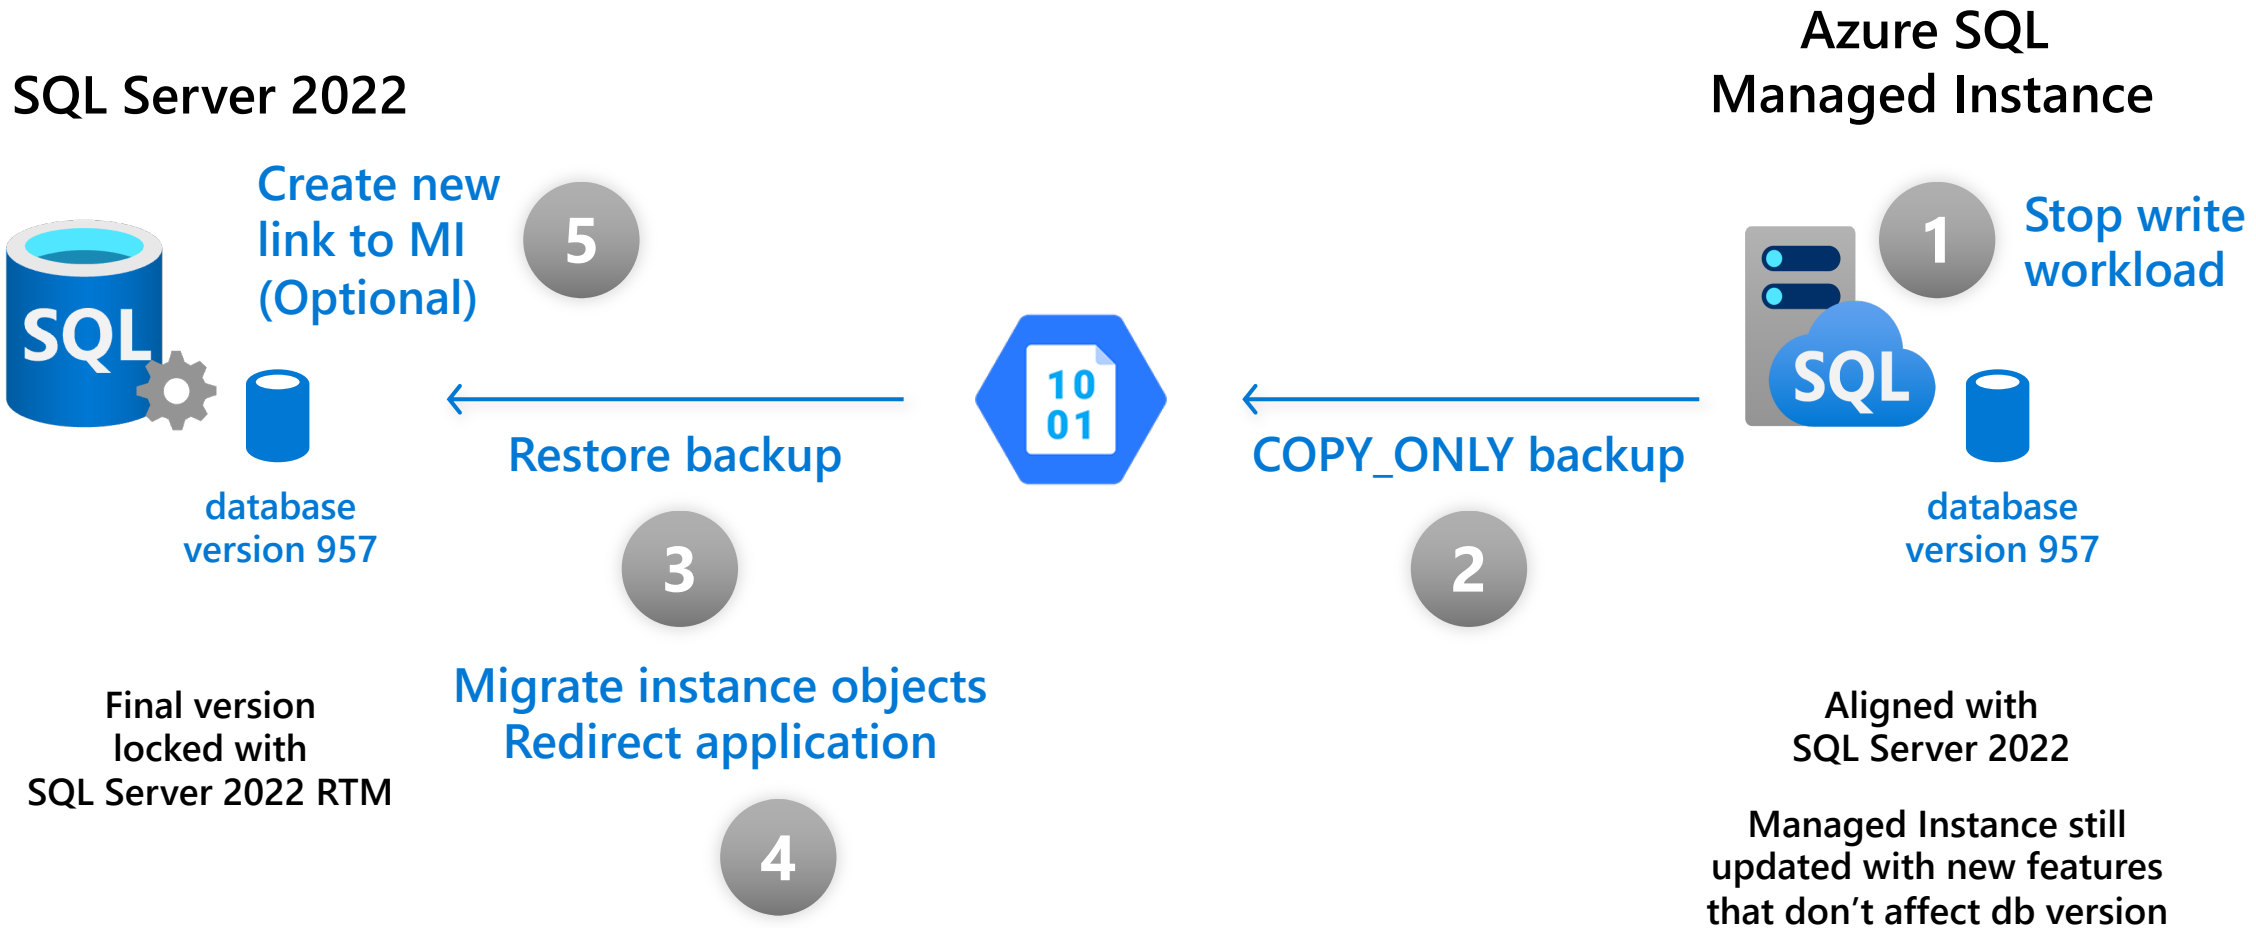

Microsoft

# Demo

## Exercise 2.0

## Link feature for Managed Instance

# Azure Synapse Link for SQL Server

## Seamless analytics over on-prem operational data

### Challenge: ETL expensive, out of date, and affects operational workloads

- Break the wall between operational and analytical stores
- New change feed capability reduces impact on OLTP workloads
- Near real-time latency between SQL Server and Synapse Analytics
- Use SQL pools so harness the full power of a scalable warehouse solution
- Analyze all your data using both Spark and SQL runtimes in Synapse

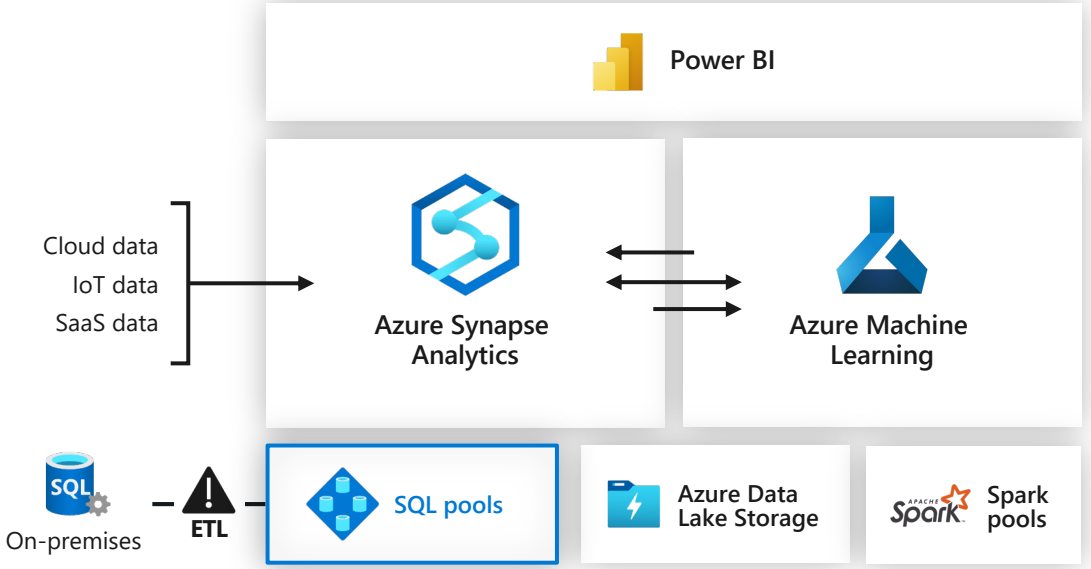


https:/aka.ms/synapselinksql

# Using Synapse Link for SQL Server

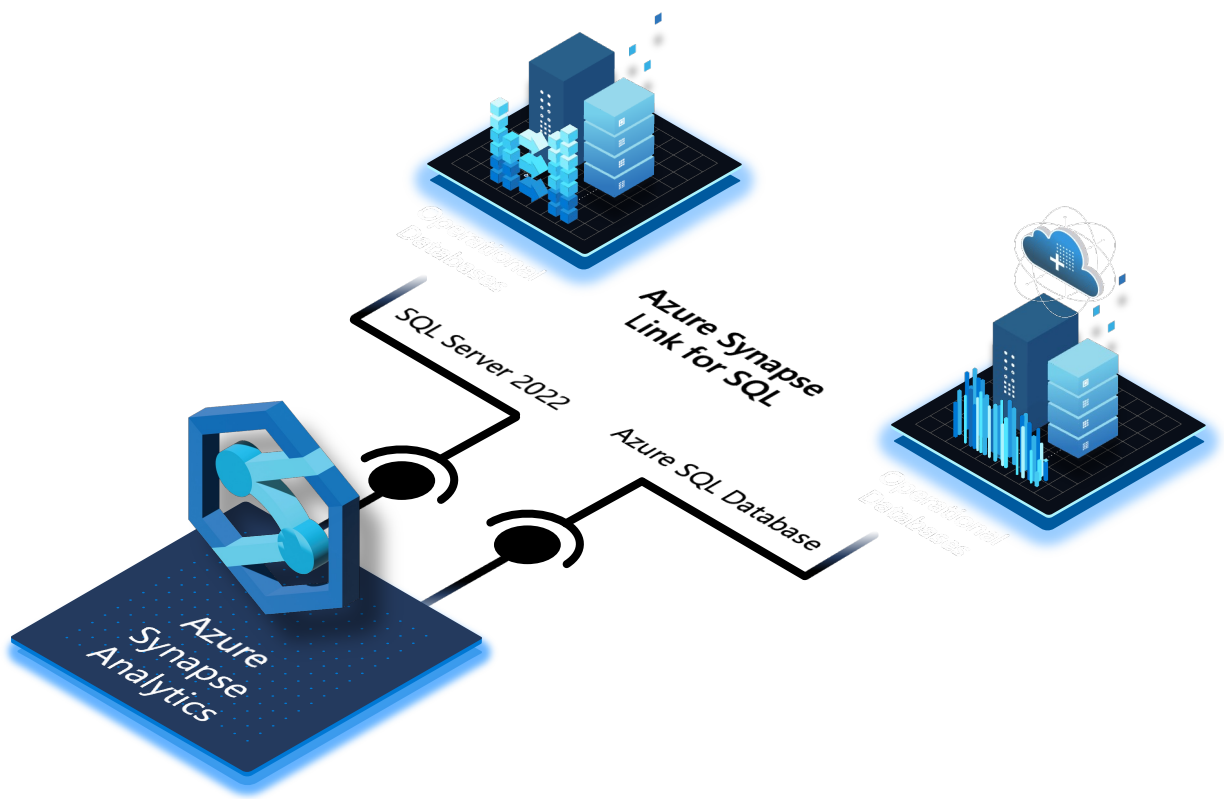


- Synapse Workspace
- Dedicated SQL Pool
- Landing Zone
- Self-hosted Integration Runtime
- SQL Server 2022 database
- Engine enhancements for change feed
- Similar for Azure SQL Database

# Azure Synapse Link for SQL Server

## How it works

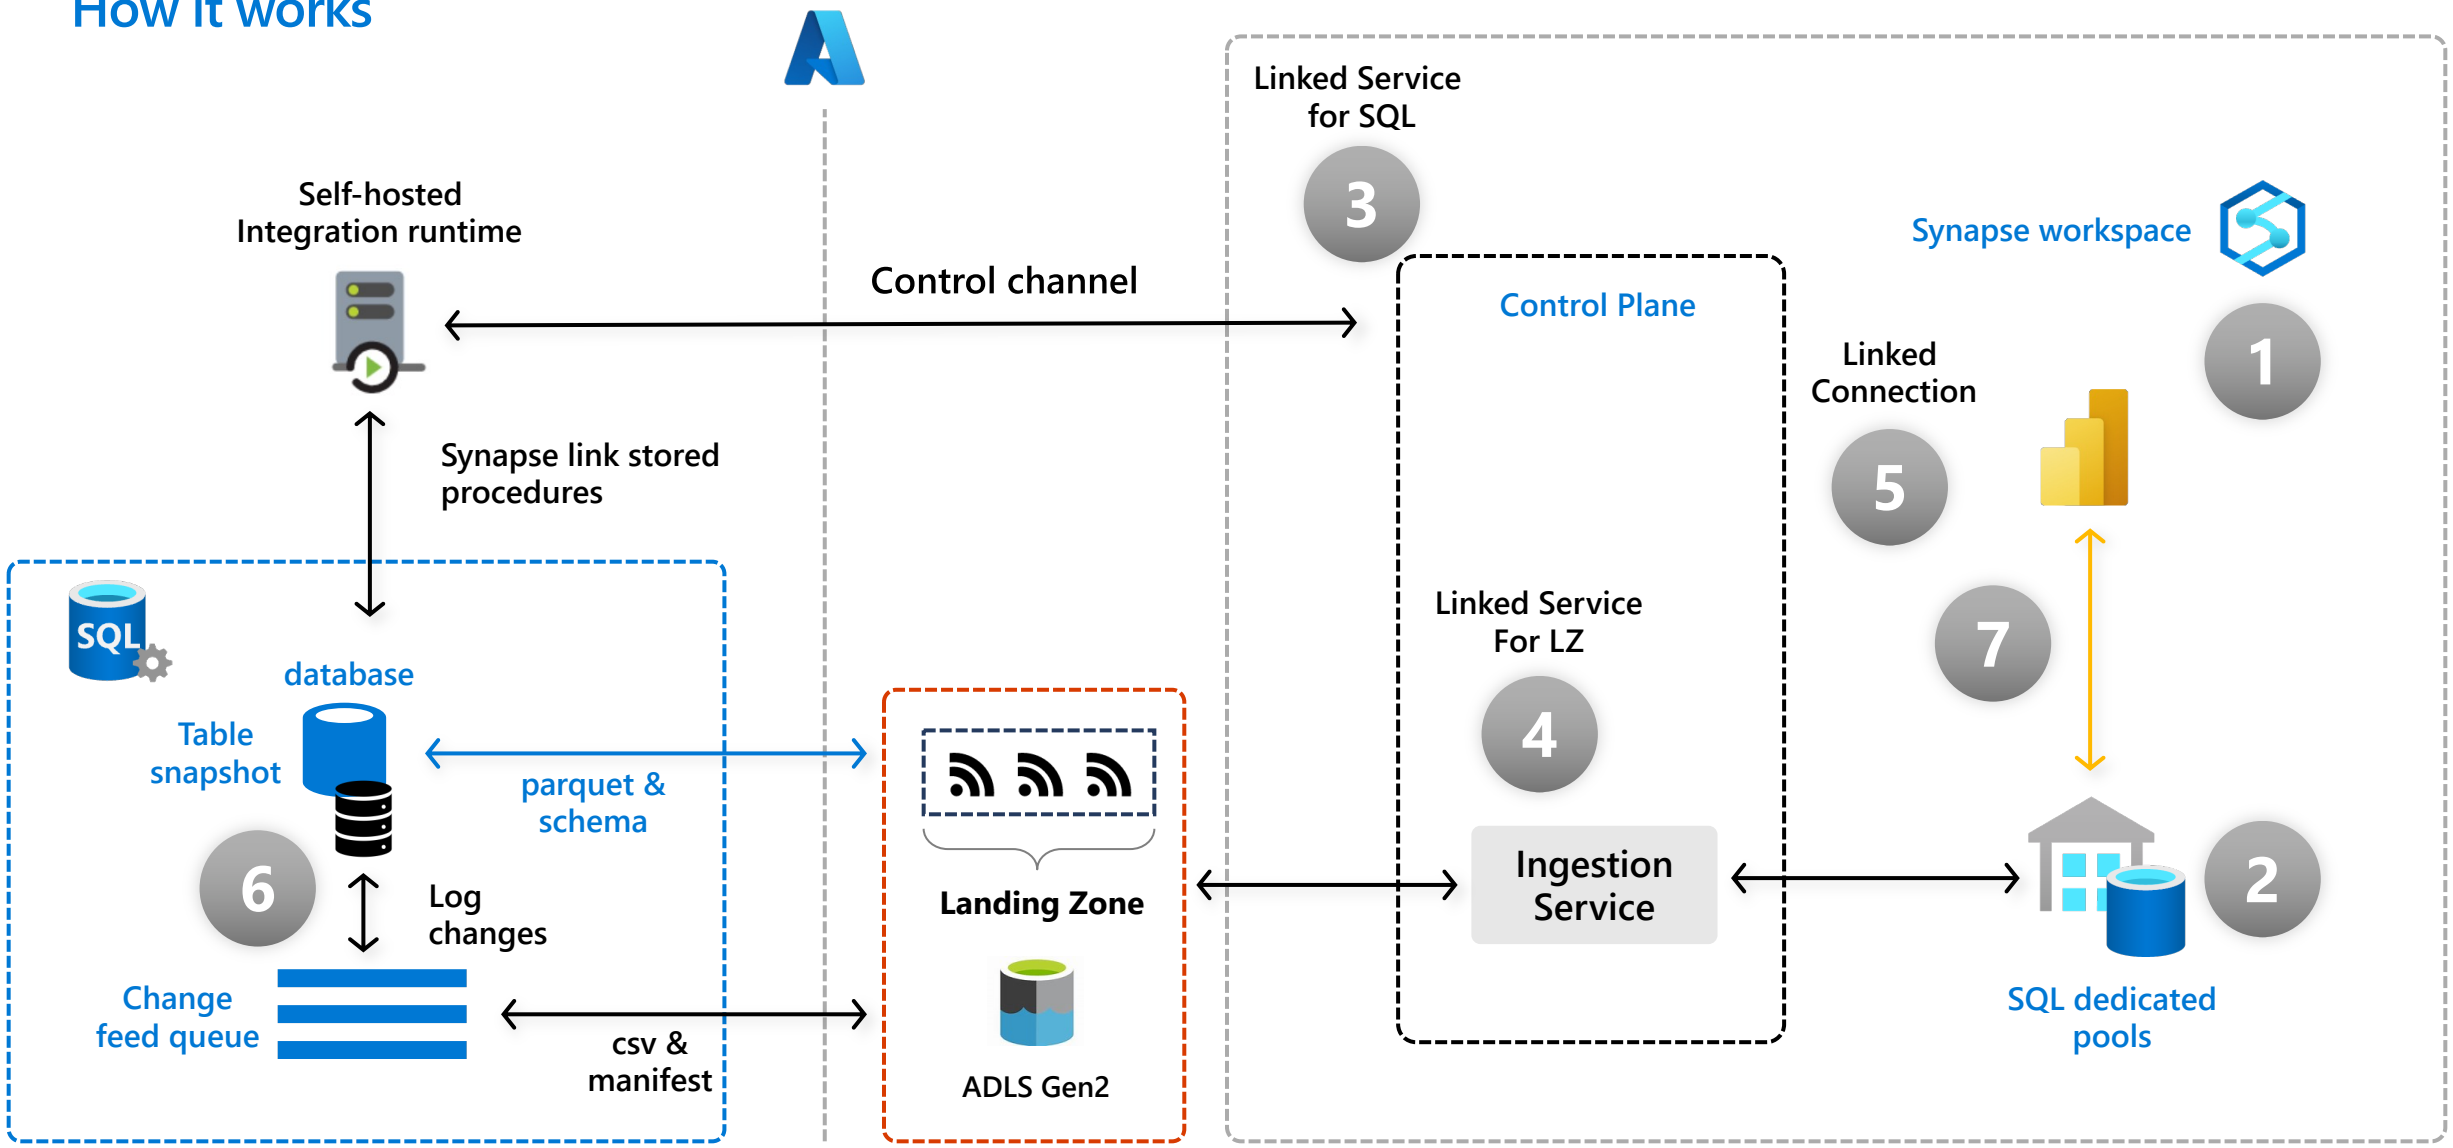

Microsoft

# Demo

## Exercise 2.1
## Synapse link for SQL

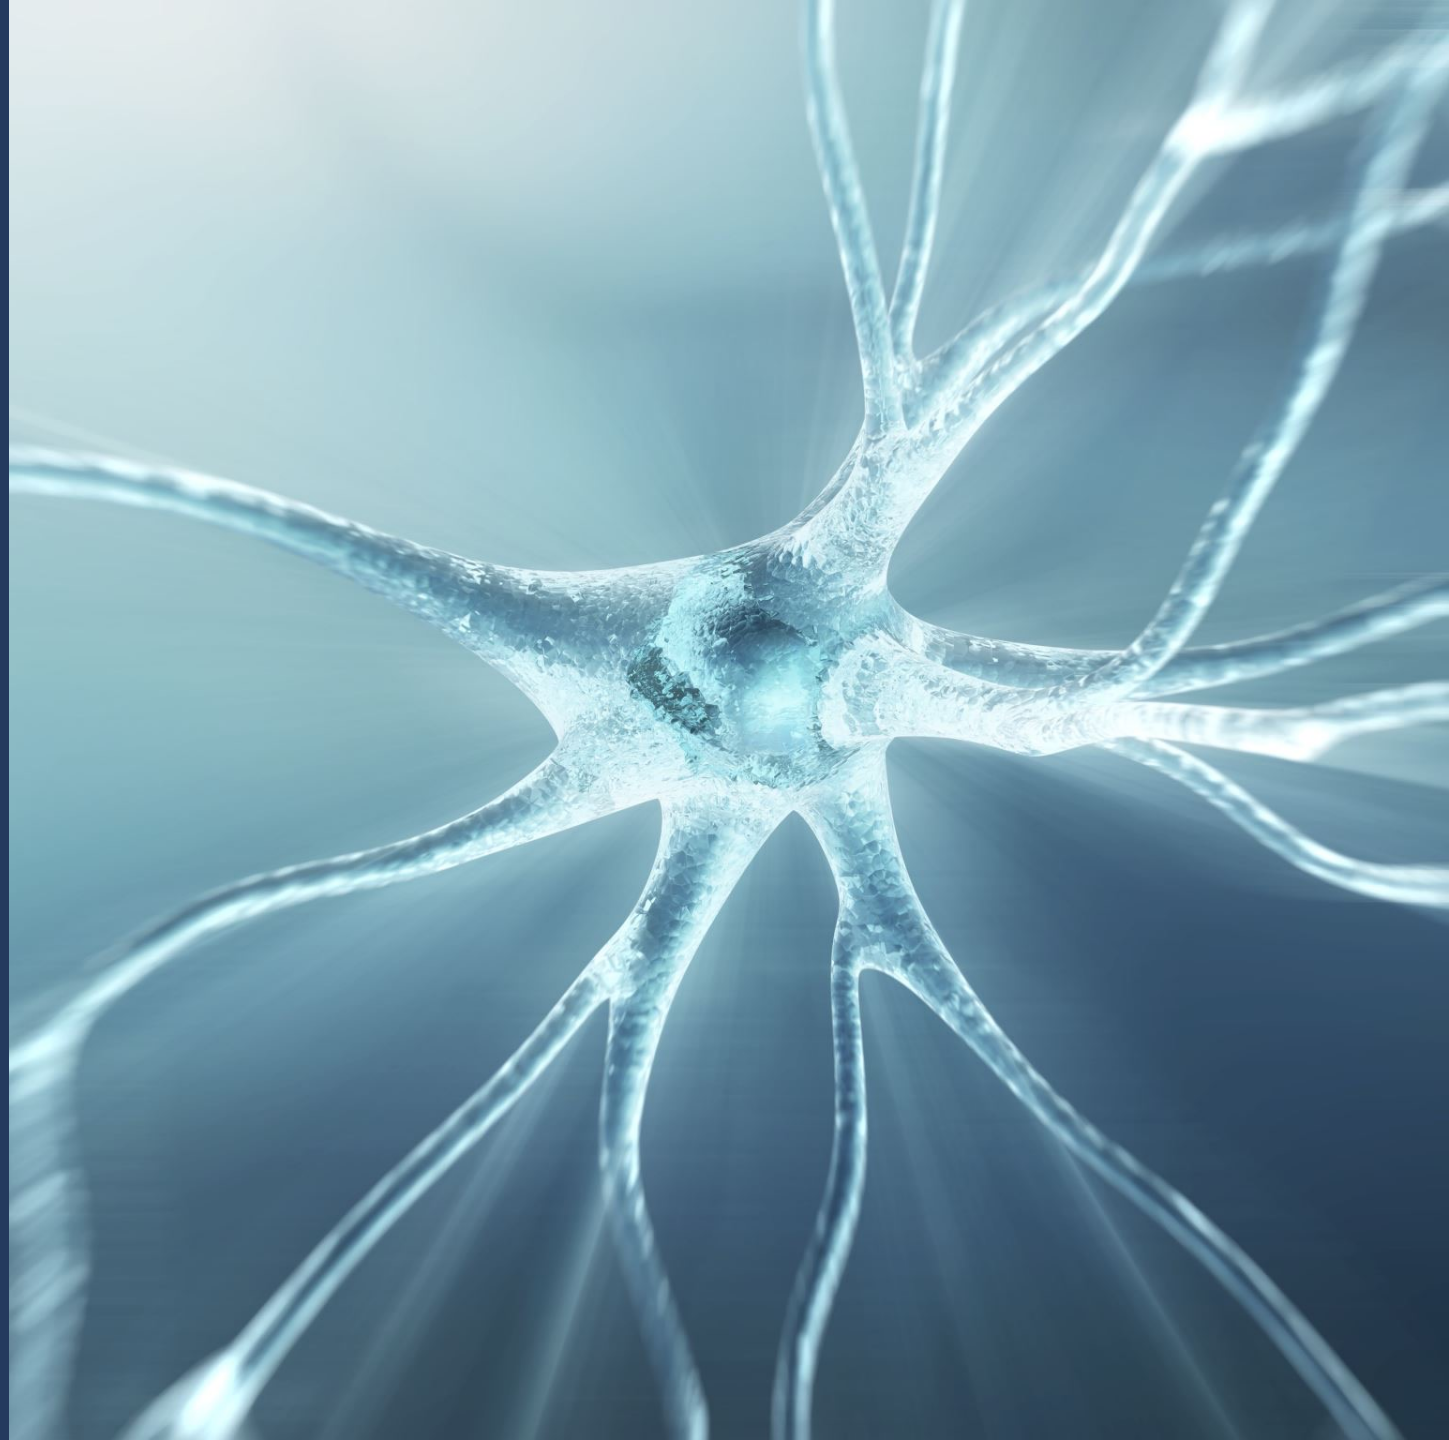

# Azure Arc-enabled SQL Server

**Azure Arc-enabled SQL Server**

Organize, inventory, and monitor
Governance and Security
Use with your existing SQL servers
SQL on Windows or Linux servers


- Pay-as-you-go-licensing
- Free inventory of instances and databases
- SQL Best Practices Assessment
- Automated backups
- Automated security updates
- Microsoft Defender for SQL
- Azure Active Directory*
- Microsoft Purview access policies*

* SQL Server 2022 only

Microsoft

# Demo

## Exercise 2.2

## Azure Arc-enabled SQL Server

# Secure your SQL Server 2022 with Azure

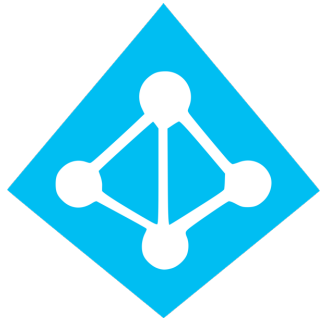

Azure Active Directory Authentication

Microsoft Purview

Microsoft Defender for SQL

# Azure Active Directory for SQL Server

https://aka.ms/aadsqlserver

Challenge: I need an alternative for authentication than SQL and Windows domains

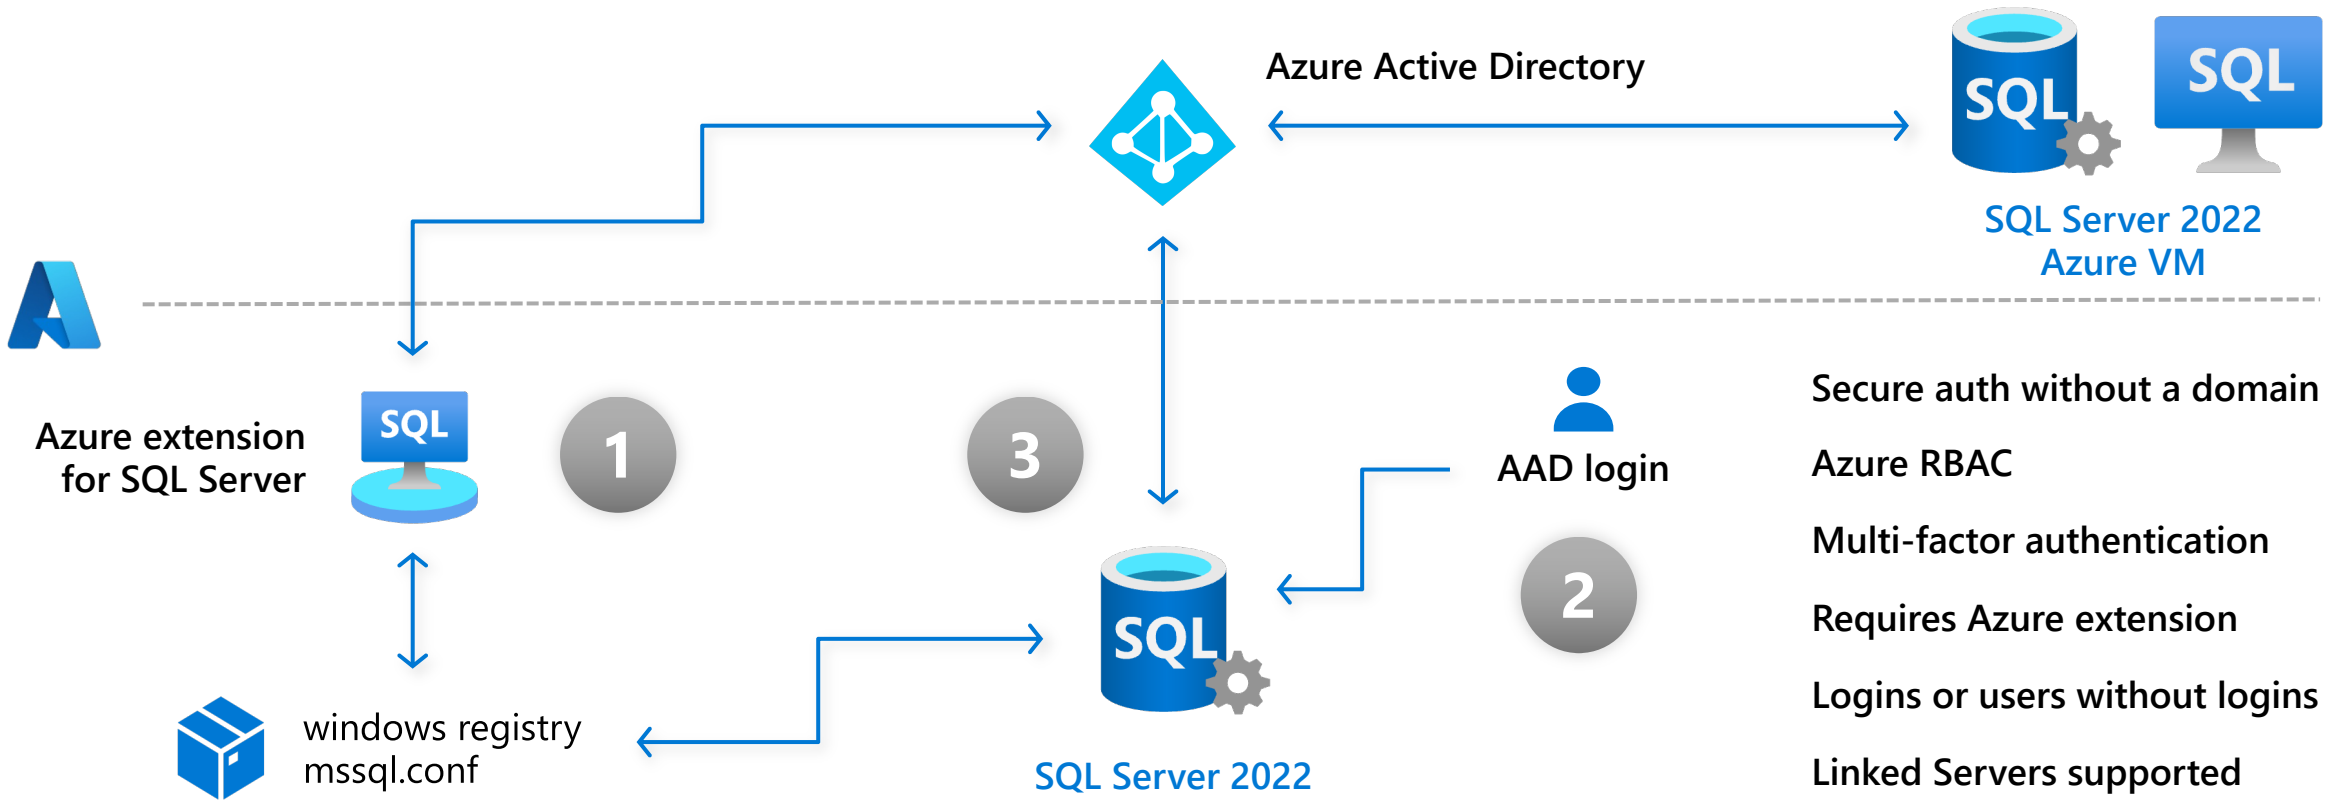


Secure auth without a domain

Azure RBAC

Multi-factor authentication

Requires Azure extension

Logins or users without logins

Linked Servers supported

Microsoft

# Demo

## Exercise 2.3

## Azure Active Directory (AAD)

# Microsoft Purview integration

## Insights, lineage, and governance over your entire data estate

### Challenge: I need central insights and governance across all of SQL

- Automatically scan on-prem SQL Server to capture metadata
- Classify data using built-in and custom classifiers and Microsoft Information Protection sensitivity labels
- Track lineage for usage and dependencies
- Set up and control specific access rights to SQL Server with policies

# Microsoft Purview access policies for SQL Server

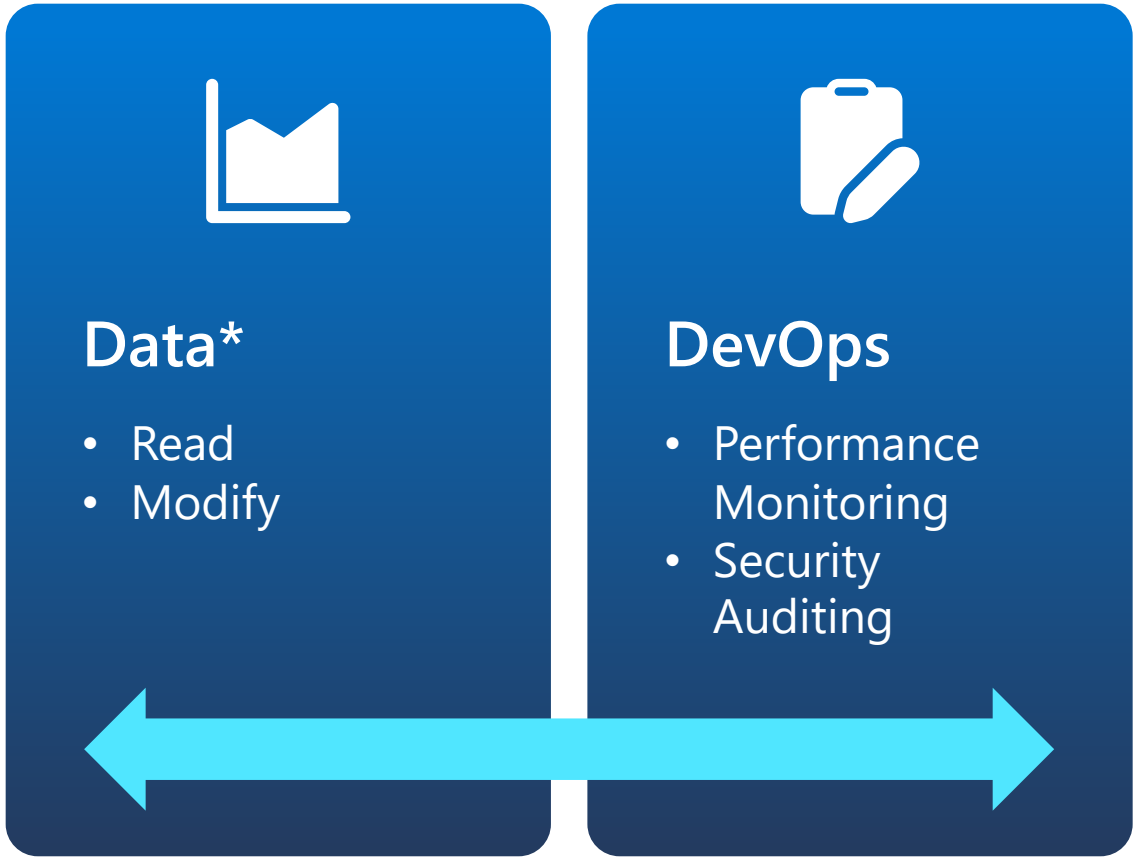


- Publish a policy from Purview
- Works only with AAD accounts
- No explicit login required
- Engine understands policies to grant login and specific permissions
- Engine integrates policies with standard logins and permissions
- Works across multiple SQL instances
- Remove policy removes authentication

**Futures:**

More granular scope

New policies

Self-service access

https://aka.ms/purviewsqlaccesspolicies

*Preview

# Microsoft Purview SQL Server Policy Architecture

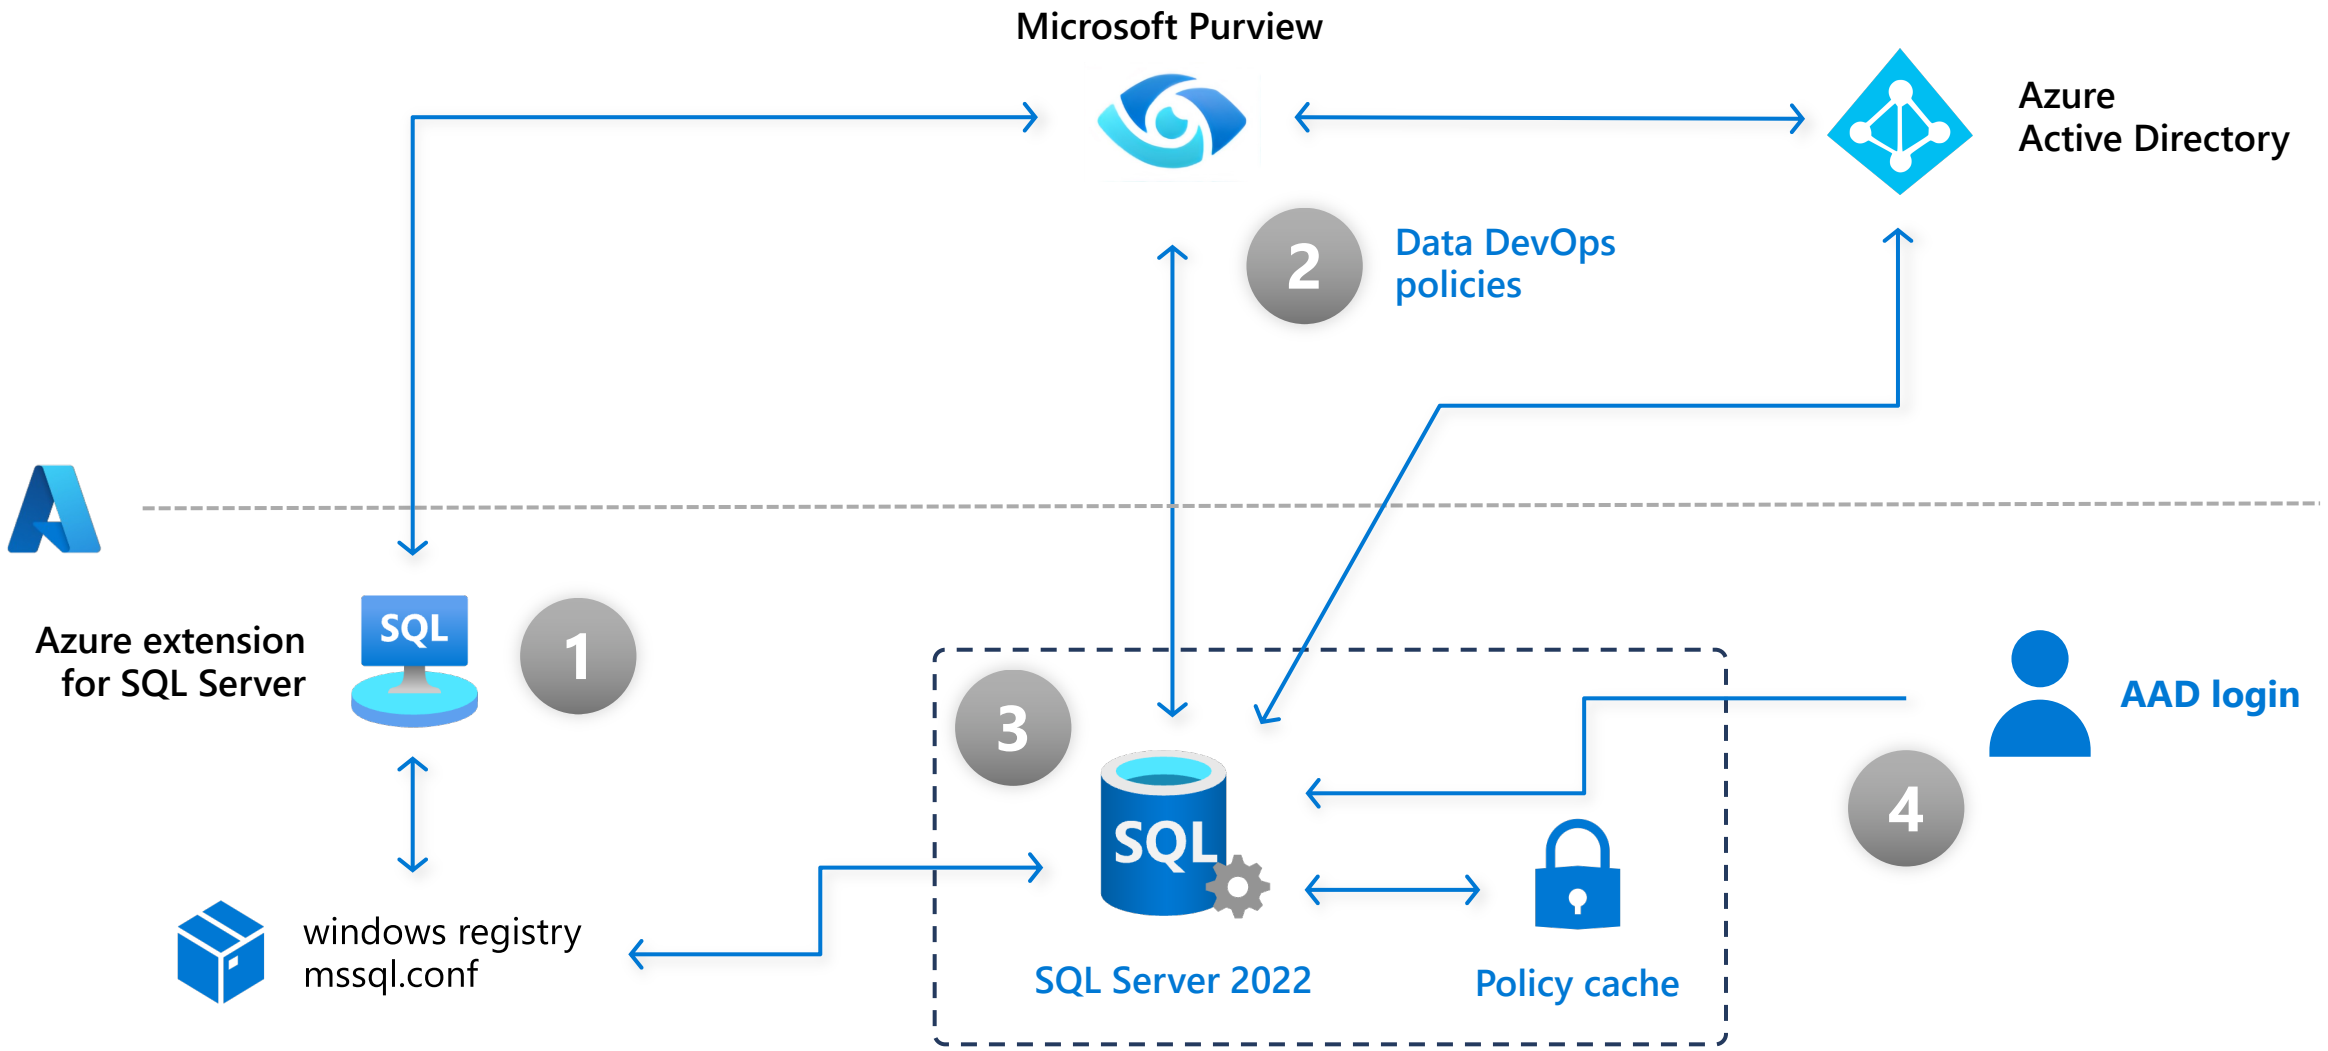

Microsoft

# Demo

## Exercises 2.4

## Microsoft Purview Access Policies

# Microsoft Defender for Cloud—Databases Protection

## Protect SQL workloads through security posture management and allow timely responses to threats

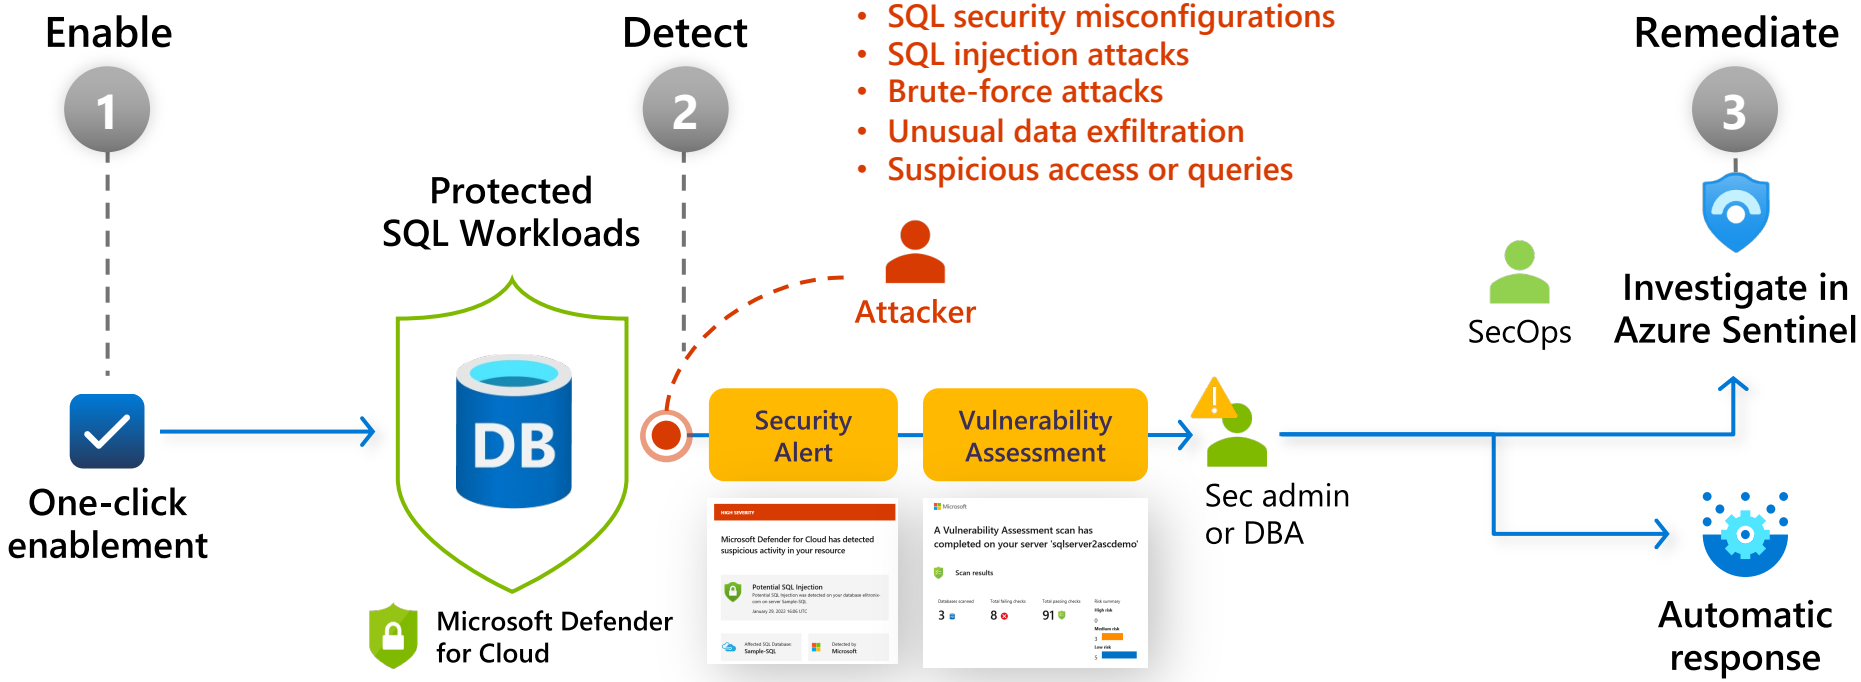


### Cloud native security

1-click enablement of protect different type of SQL workloads (IaaS or PaaS)

### Security posture management

Discover, track, and remediate SQL workloads security misconfigurations

### Advanced threat protection

Detect and response unusual and harmful attempts to breach SQL workloads

### Centralized and integrated

Centralize security across all data assets managed by Azure and built-in integration with Sentinel and Purview

# Knowledge check

**Link for Managed Instance feature is supported on which SQL Server versions?**

**A.** SQL Server 2016

**B.** SQL Server 2019

**C.** SQL Server 2022

**D.** B and C

**E.** A, B, and C

**SQL Server sends changes to Synapse Link to what location?**

**A.** SQL Pools

**B.** Landing Zone

**C.** Azure Data Factory

**D.** Linked Service

**E.** Linked Connection

# Knowledge check

**Azure Arc-enabled SQL Server provides what capabilities?**

**A.** Free inventory of instances and databases

**B.** Managed Database Service

**C.** Synapse Link

**D.** Best Practices Assessment

**E.** A and D

**Microsoft Purview policies provide access to SQL Server by:**

**A.** You create a login and then map the login to a Purview policy

**B.** You create a policy in Purview based on an Azure Active Directory (AAD) account

**C.** You create a role in SQL Server that maps to a Purview policy

**D.** You create a AAD login and add it to a policy role in SQL Server

**E.** Purview scans SQL Server for logins

# Summary

Link SQL Server to Azure SQL Managed Instance for **managed disaster recovery**

Link SQL Server to Synapse for automatic **near real-time analytics**

**Azure Arc-enabled SQL Server** provides hybrid capabilities like inventory, PAYG, and BPA

Authenticate SQL Server logins and users with **Azure Active Directory (AAD)**

Publish **policies for data access** without creating logins or roles with Microsoft Purview

Module 3

# Accelerate performance with built-in query intelligence

Overview

The new Query Store

IQP NextGen foundations

IQP NextGen with dbcompat 140+

IQP NextGen with dbcompat 160

Exercises

Knowledge Checks and Summary

# Query Store and Intelligent Query Processing

## Accelerate query performance and tuning with no code changes

### Challenge: Query tuning is expensive and time consuming

- Query Store on by default for new databases
- Query Store support for read replicas from availability groups*
- Query Store hints to shape plans with no code changes
- New IQP scenarios enabled through better together capabilities

* Currently requires trace flag 12606

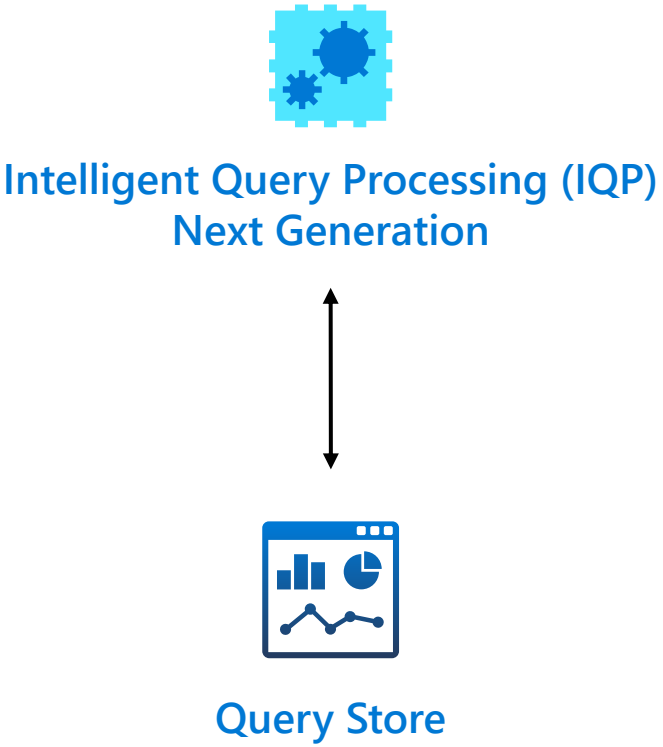

# Query Store and Intelligent Query Processing

## Accelerate query performance and tuning with no code changes

### Challenge: Query tuning is expensive and time consuming

- Query Store on by default for new databases
- Query Store support for read replicas from availability groups*
- Query Store hints to shape plans with no code changes
- New IQP scenarios enabled through better together capabilities

* Currently requires trace flag 12606

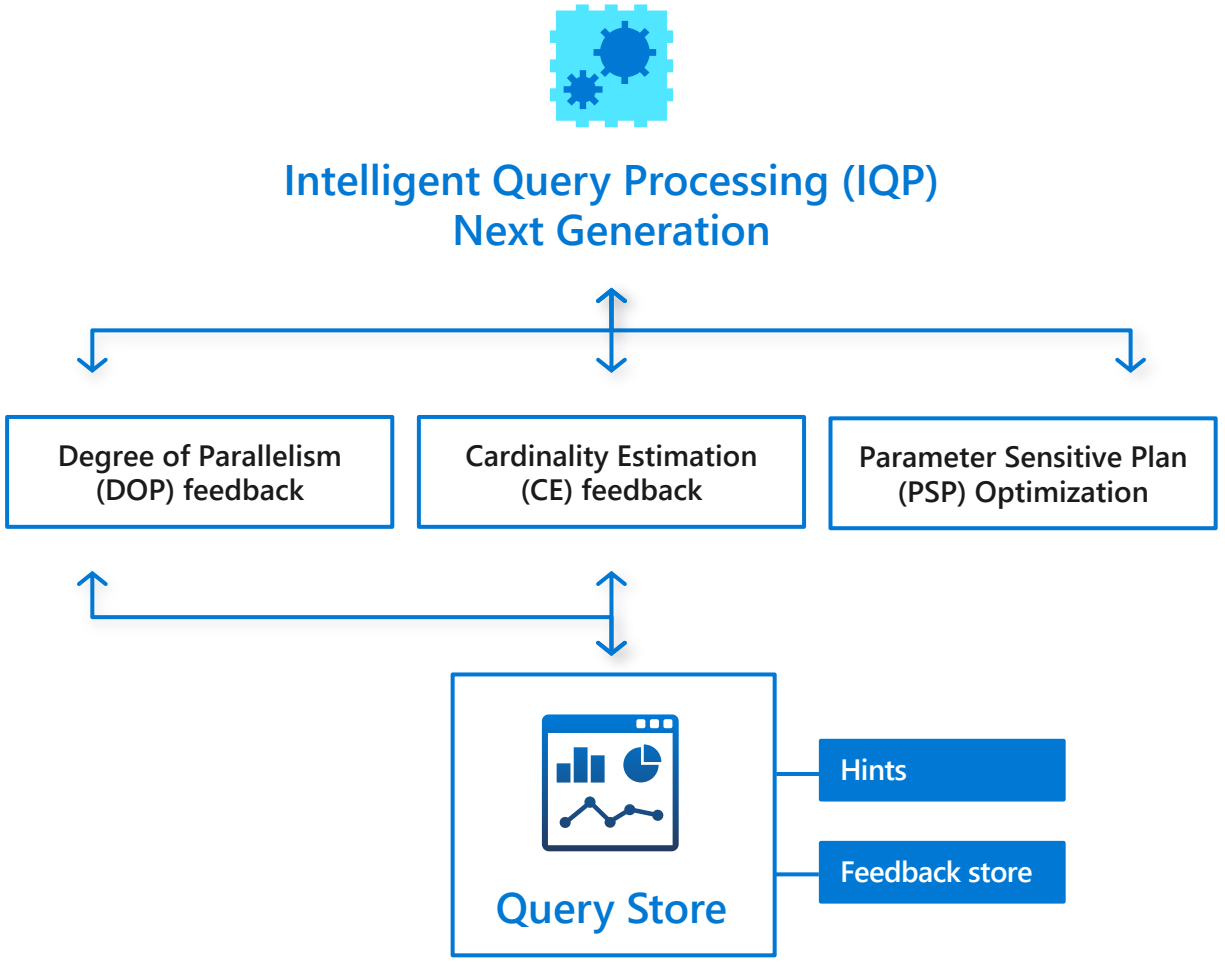

# Intelligent Query Processing (IQP) NextGen

## Gain consistent performance with no code changes

Query Store

### Upgrade to SQL Server 2022

#### Approximate Percentile

- Approximate versions for PERCENTILE_CONT and PERCENTILE_DISC

#### Optimized Plan Forcing

- Persist compile steps for forced plans
- Reduce compilation overhead

dbcompat 140 = IQP Gen1
dbcompat 150 = IQP Gen1+2

### dbcompat 140+

#### Memory Grant Feedback Percentiles

- Smooth out oscillation with percentile-based calculation

#### Memory Grant Feedback Persistence

- Persist feedback for cache eviction

**Key principles**

- "Do no harm"
- Options to disable

Adaptive Joins
Interleaved Execution
Memory Grant Feedback
Table Variable Deferred Compilation
Batch Mode on Rowstore
Scalar UDF inlining
Approximate Count Distinct

### dbcompat 160

#### Parameter Sensitive Plan (PSP) Optimization

- Cache multiple query plans for the same parameterized statement
- Save the world from parameter sniffing?

#### Cardinality Estimation (CE) feedback

- Feedback loop for CE model choices and query execution
- Auto use query hints

#### Degree of Parallelism (DOP) feedback

- Feedback loop for DOP for a specific query
- Adjust query DOP without recompilation

Optimizer fixes

# Query Store enhancements

- Query Store on by default for new databases
- Query Store support for read replicas from availability groups*
- Query Store hints to shape plans with no code changes
- Query store is the perfect “store” for IQP

* Currently requires trace flag 12606

# Query Store for read replicas*

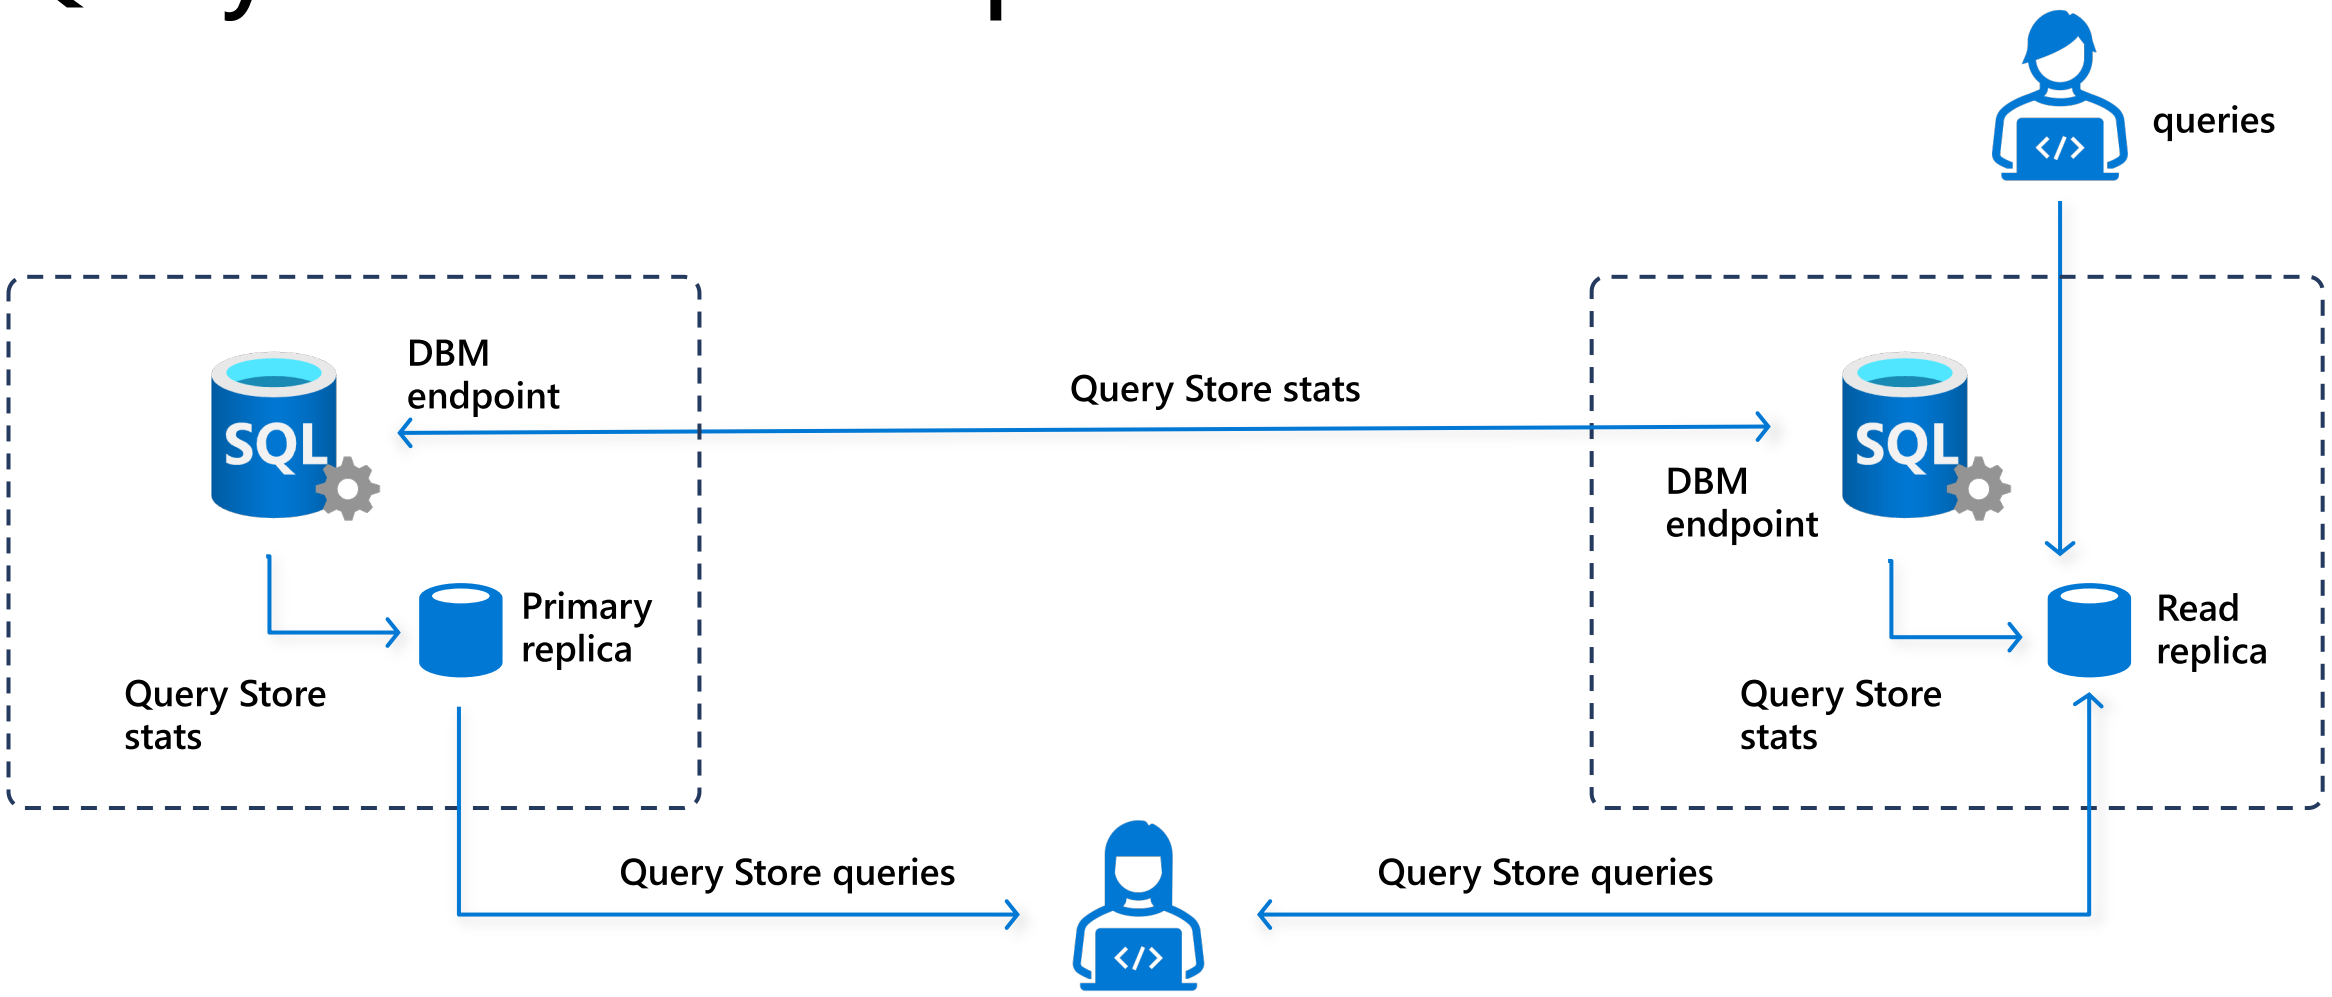


* Currently requires trace flag 12606

# Intelligent Query Processing (IQP) NextGen

## Gain consistent performance with no code changes

Query Store

### Upgrade to SQL Server 2022

#### Approximate Percentile

- Approximate versions for PERCENTILE_CONT and PERCENTILE_DISC

#### Optimized Plan Forcing

- Persist compile steps for forced plans
- Reduce compilation overhead

# Intelligent Query Processing (IQP) NextGen

## Gain consistent performance with no code changes

Query Store

### Upgrade to SQL Server 2022

#### Approximate Percentile

- Approximate versions for PERCENTILE_CONT and PERCENTILE_DISC

#### Optimized Plan Forcing

- Persist compile steps for forced plans
- Reduce compilation overhead

### dbcompat 140+

#### Memory Grant Feedback Percentiles

- Smooth out oscillation with percentile-based calculation

#### Memory Grant Feedback Persistence

- Persist feedback for cache eviction

# What is a memory grant?

```
SELECT customer_id, customer_type, customer_date
FROM customers
WHERE customer_date < GETDATE()
ORDER BY customer_date, customer_type
```

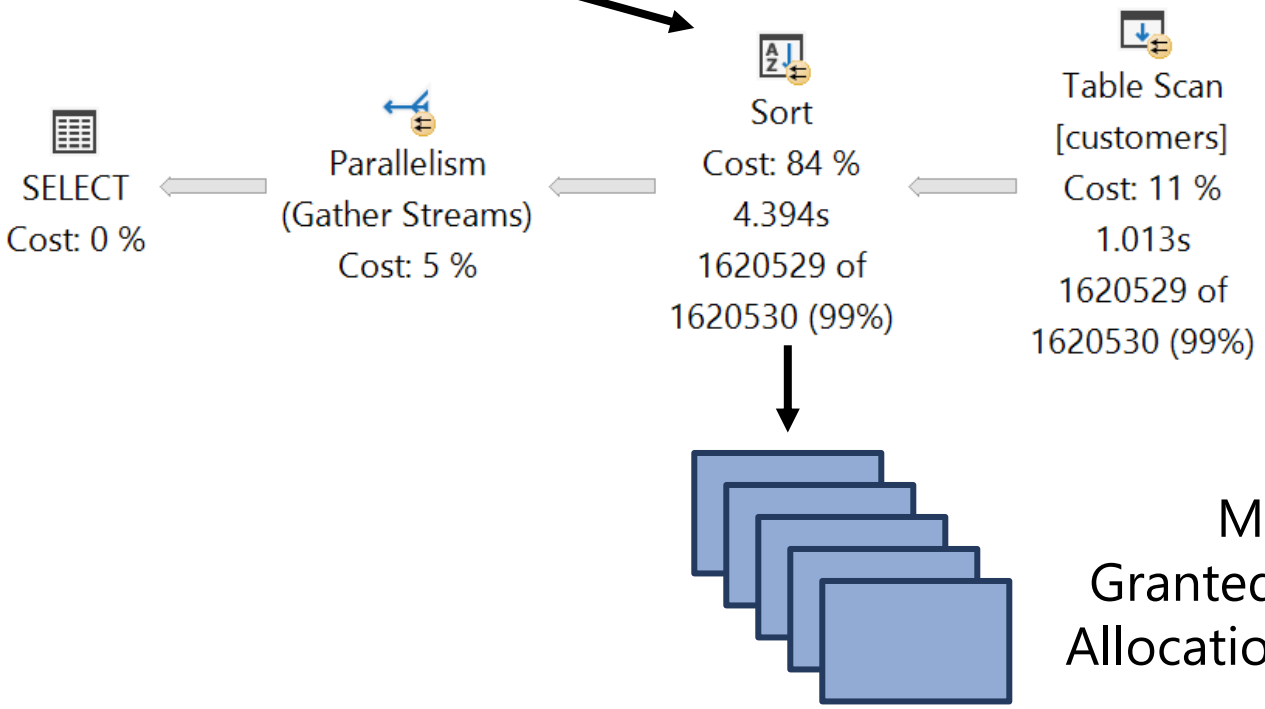


Memory Grant
Granted before execution
Allocation during execution

1,000,000, rows x
500 bytes per row

= **~500MB**

# When can memory grants become a problem?

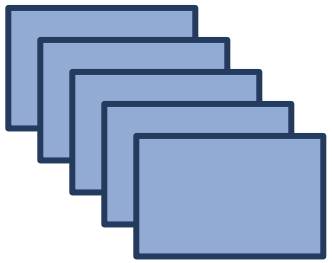

Memory Grant
Granted before execution

1,000,000, rows x
500 bytes per row

= ~**500MB**

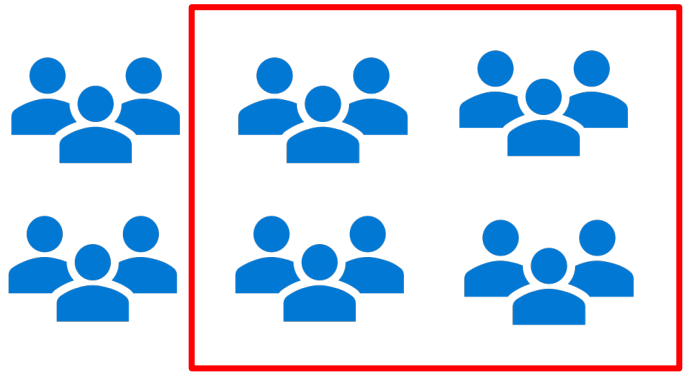

**RESOURCE_SEMAPHORE**

---

Memory Grant
1MB

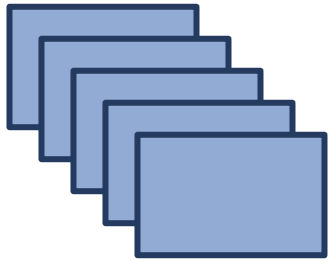

Actual Memory Needed
500MB

tempdb

**IO_COMPLETION**

# What is Memory Grant Feedback (MGF)?

## Memory Grant

Memory allocated at start of query execution for operators like hashes and sorts

Estimates could cause memory grant to be wrong

Too small = tempdb spill

Too big = RESOURCE_SEMAPHORE waits

Workarounds include hints and resource governor

## SQL Server 2017 and 2019

First batch mode only and then row store

Observe actual grant after execution

Adjust next execution if grant was incorrect

# MGF Percentiles and Persistence

## MGF Percentiles

Solves problem if memory grant feedback is constantly moving the grant back and forth

We would disable before with volatility

Now use a trend over time (favor avoiding spills)

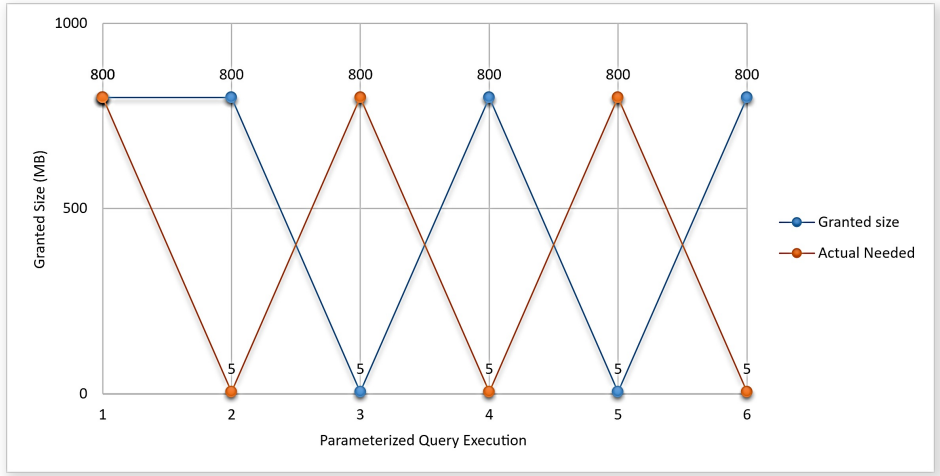


## MGF Persistence

Store feedback in the Query Store to survive plan cache eviction and restarts

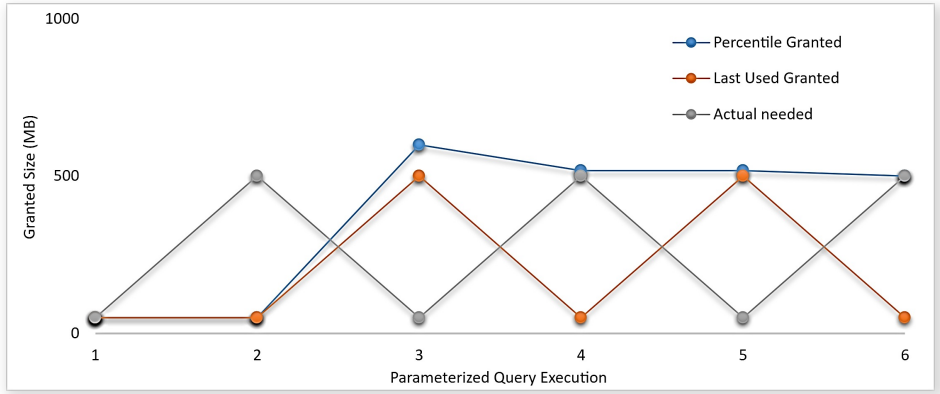

# Intelligent Query Processing (IQP) NextGen

## Gain consistent performance with no code changes

Query Store

### Upgrade to SQL Server 2022

#### Approximate Percentile

- Approximate versions for PERCENTILE_CONT and PERCENTILE_DISC

#### Optimized Plan Forcing

- Persist compile steps for forced plans
- Reduce compilation overhead

### dbcompat 140+

#### Memory Grant Feedback Percentiles

- Smooth out oscillation with percentile-based calculation

#### Memory Grant Feedback Persistence

- Persist feedback for cache eviction

### dbcompat 160

#### Parameter Sensitive Plan (PSP) Optimization

- Cache multiple query plans for the same parameterized statement
- Save the world from parameter sniffing?

#### Cardinality Estimation (CE) feedback

- Feedback loop for CE model choices and query execution
- Auto use query hints

#### Degree of Parallelism (DOP) feedback

- Feedback loop for DOP for a specific query
- Adjust query DOP without recompilation

# Parameter Sensitive Plan Optimization

Before




With PSP optimization

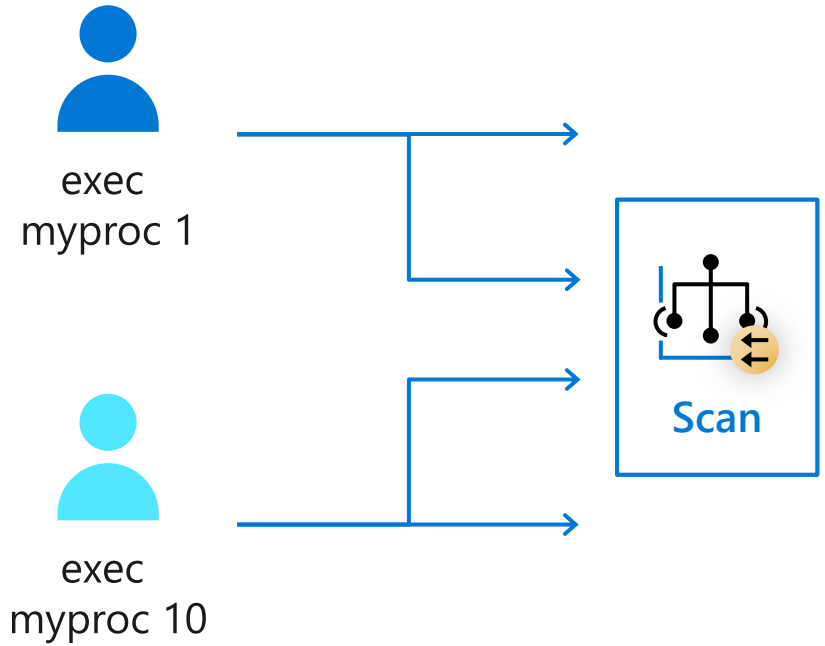

# How does PSP optimization work?

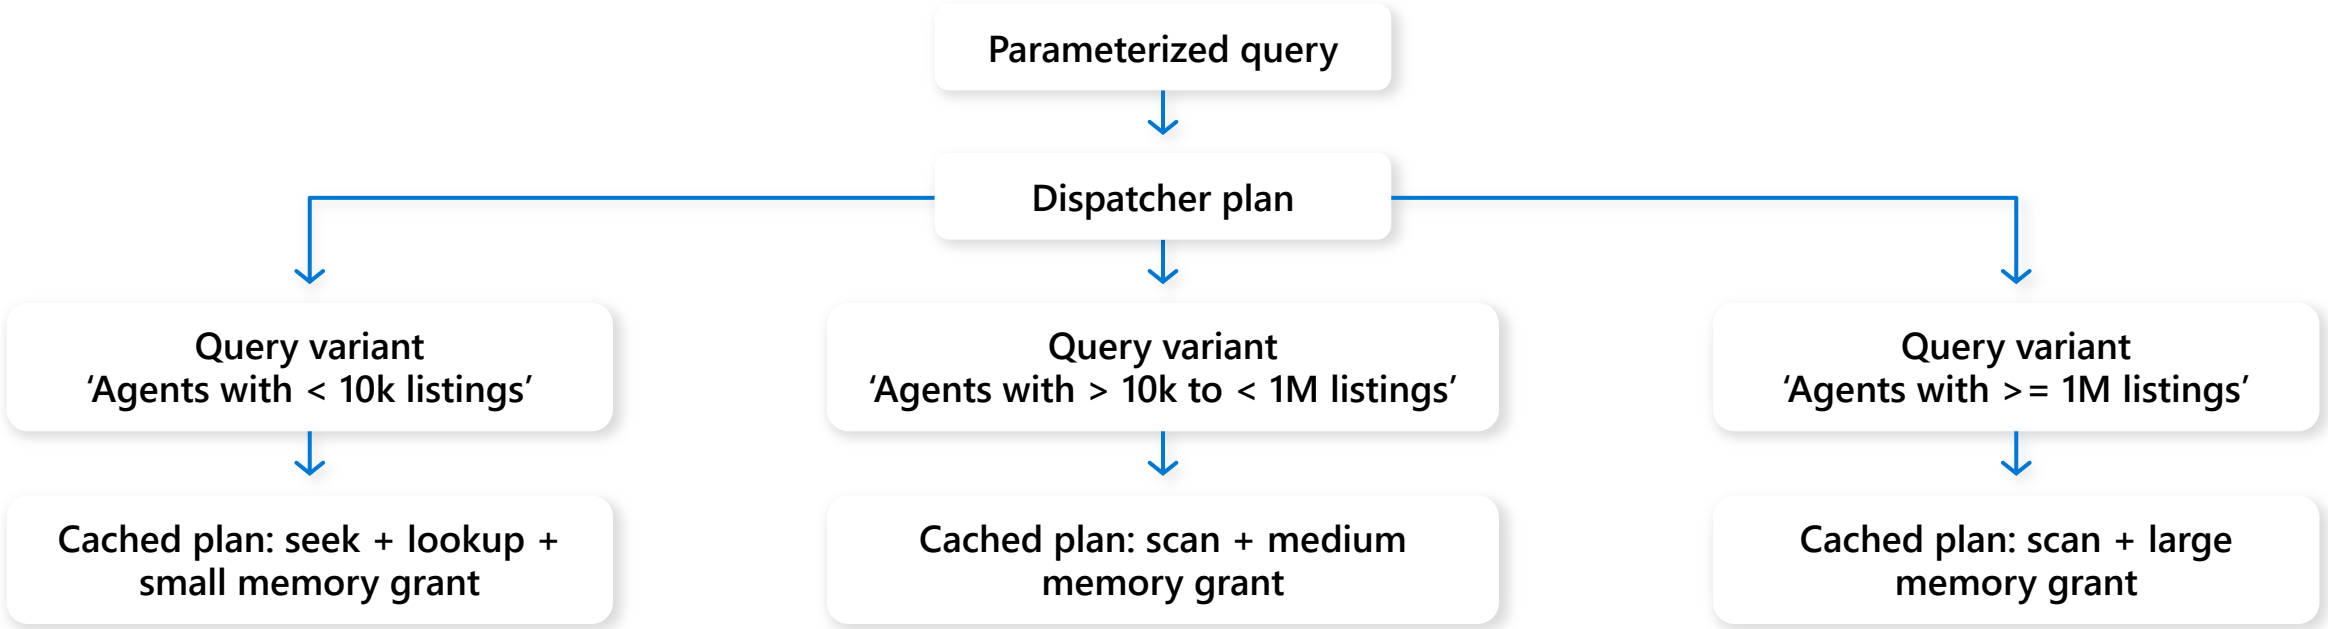


- Not one plan per user
- Dispatcher plan and variant plans
- We have buckets to represent *ranges* of rows matching parameter values
- Dependent on correct statistics
- Variants have different query text but same query hash

# Is Parameter Sniffing gone?

- ❑ Only = operator supported
- ❑ Finite limit of variants
- ❑ Conservative ratio for min/max cardinality
- ❑ You could get a plan not optimal for all parameter values *within* a variant
- ❑ Multiple parameters are supported
- ❑ New Query Store DMVs required to examine stored procs
- ❑ Issues using DMVs for plan cache
- ❑ Won't use if you use conflicting hints, … (`psp_skipped_reason_enum`)

# The problem of the Cardinality Estimation (CE) model

**What cardinality to use for query optimization is not always easy and obvious**

## CE Versions

- Legacy CE Model < dbcompat 120
- New CE Model >= dbcompat 120

## Why the change?

- Assumptions from the original model we thought no longer correct and matched modern workloads
- Uh…this didn't quite work out like we thought
- Customers upgraded to SQL Server 2014 and used dbcompat 120 and…problems.

## Workarounds

- Don't change dbcompat (ouch)
- Trace flags
- ALTER DATABASE SCOPED CONFIGURATION SET LEGACY_CARDINALITY_ESTIMATION
- Query hints

# An example CE model problem

```
SELECT AddressLine1, City, PostalCode
FROM Person.Address
WHERE StateProvinceID = 79
AND City = 'Redmond';
```

## Is there a correlation between Redmond and StateProvinceID 79?

1. **Full independence**
   Legacy model
   Hint or TF
2. **Partial correlation**
   New model
   Hint or TF
3. **Full correlation**
   Either with a hint or TF

# Cardinality Estimation (CE) Feedback

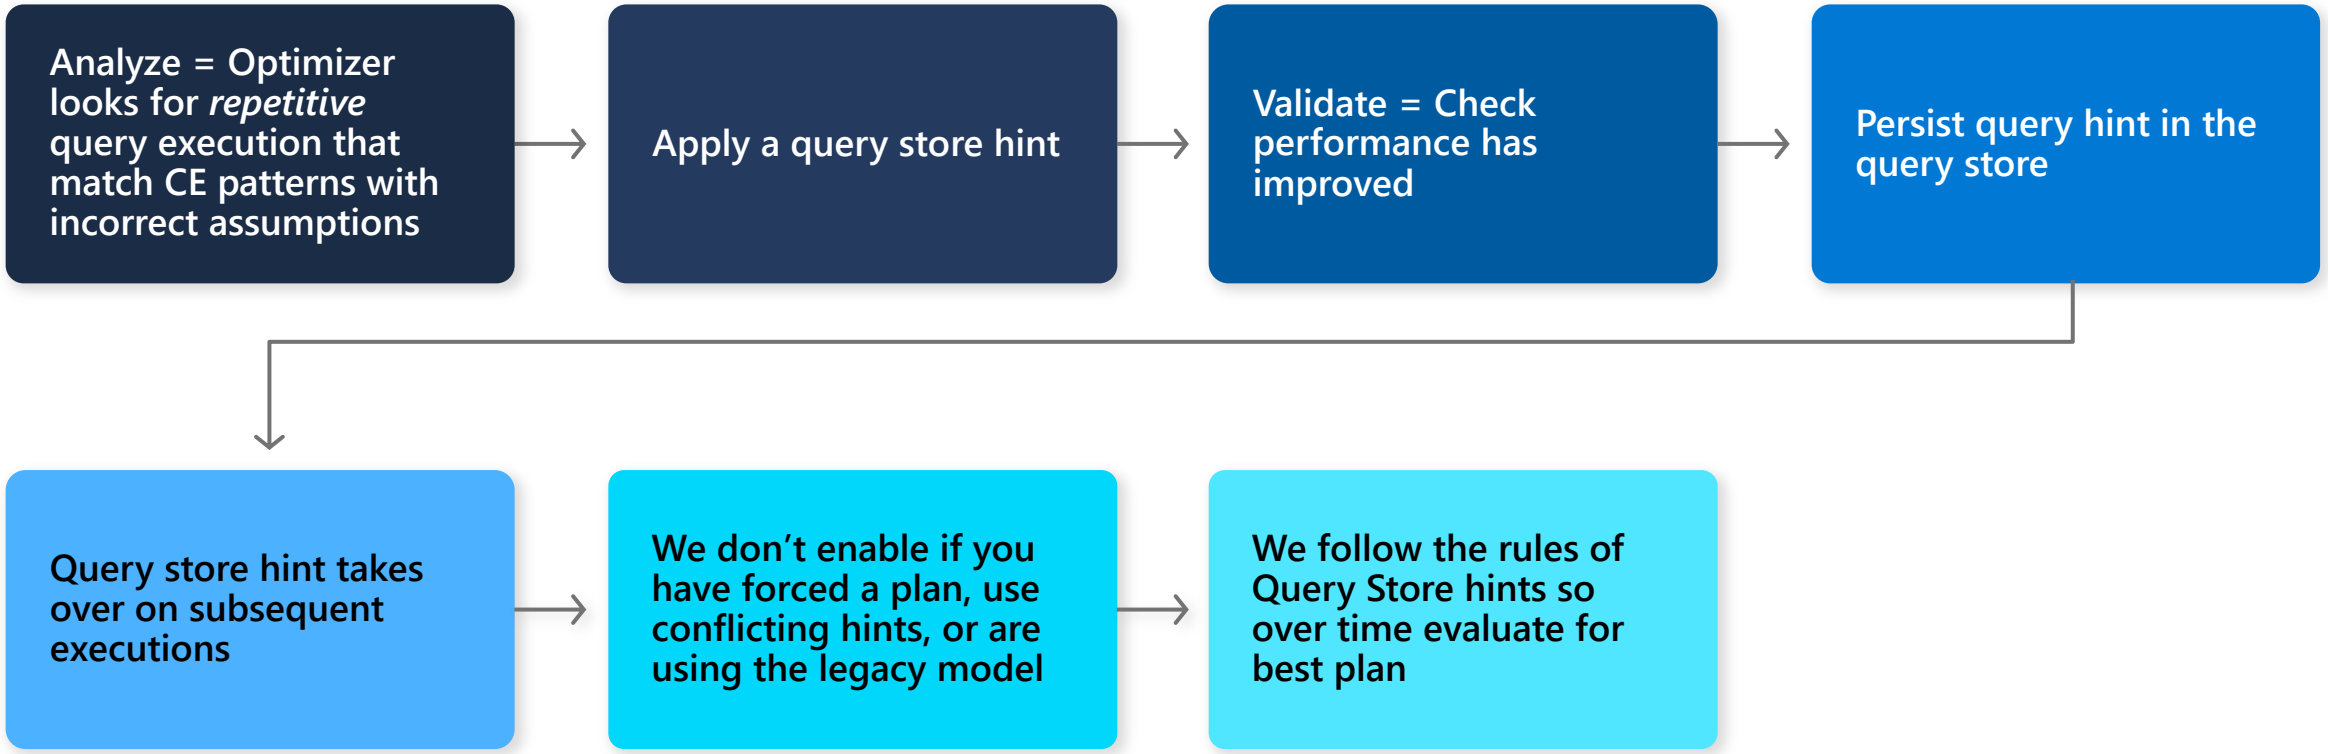

# The pain of MAXDOP

Setting MAXDOP just *right* is a pain

Is parallelism bad?

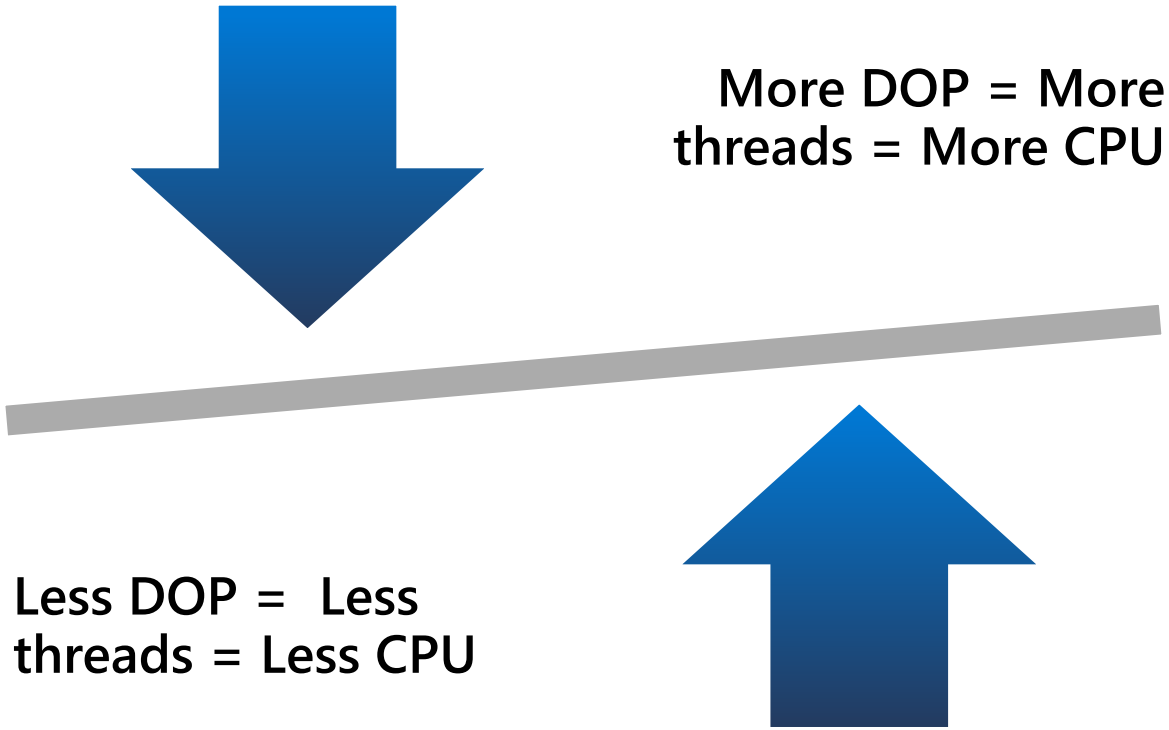


**Parallel efficiency**
The lowest DOP to achieve the same query duration

# DOP feedback architecture

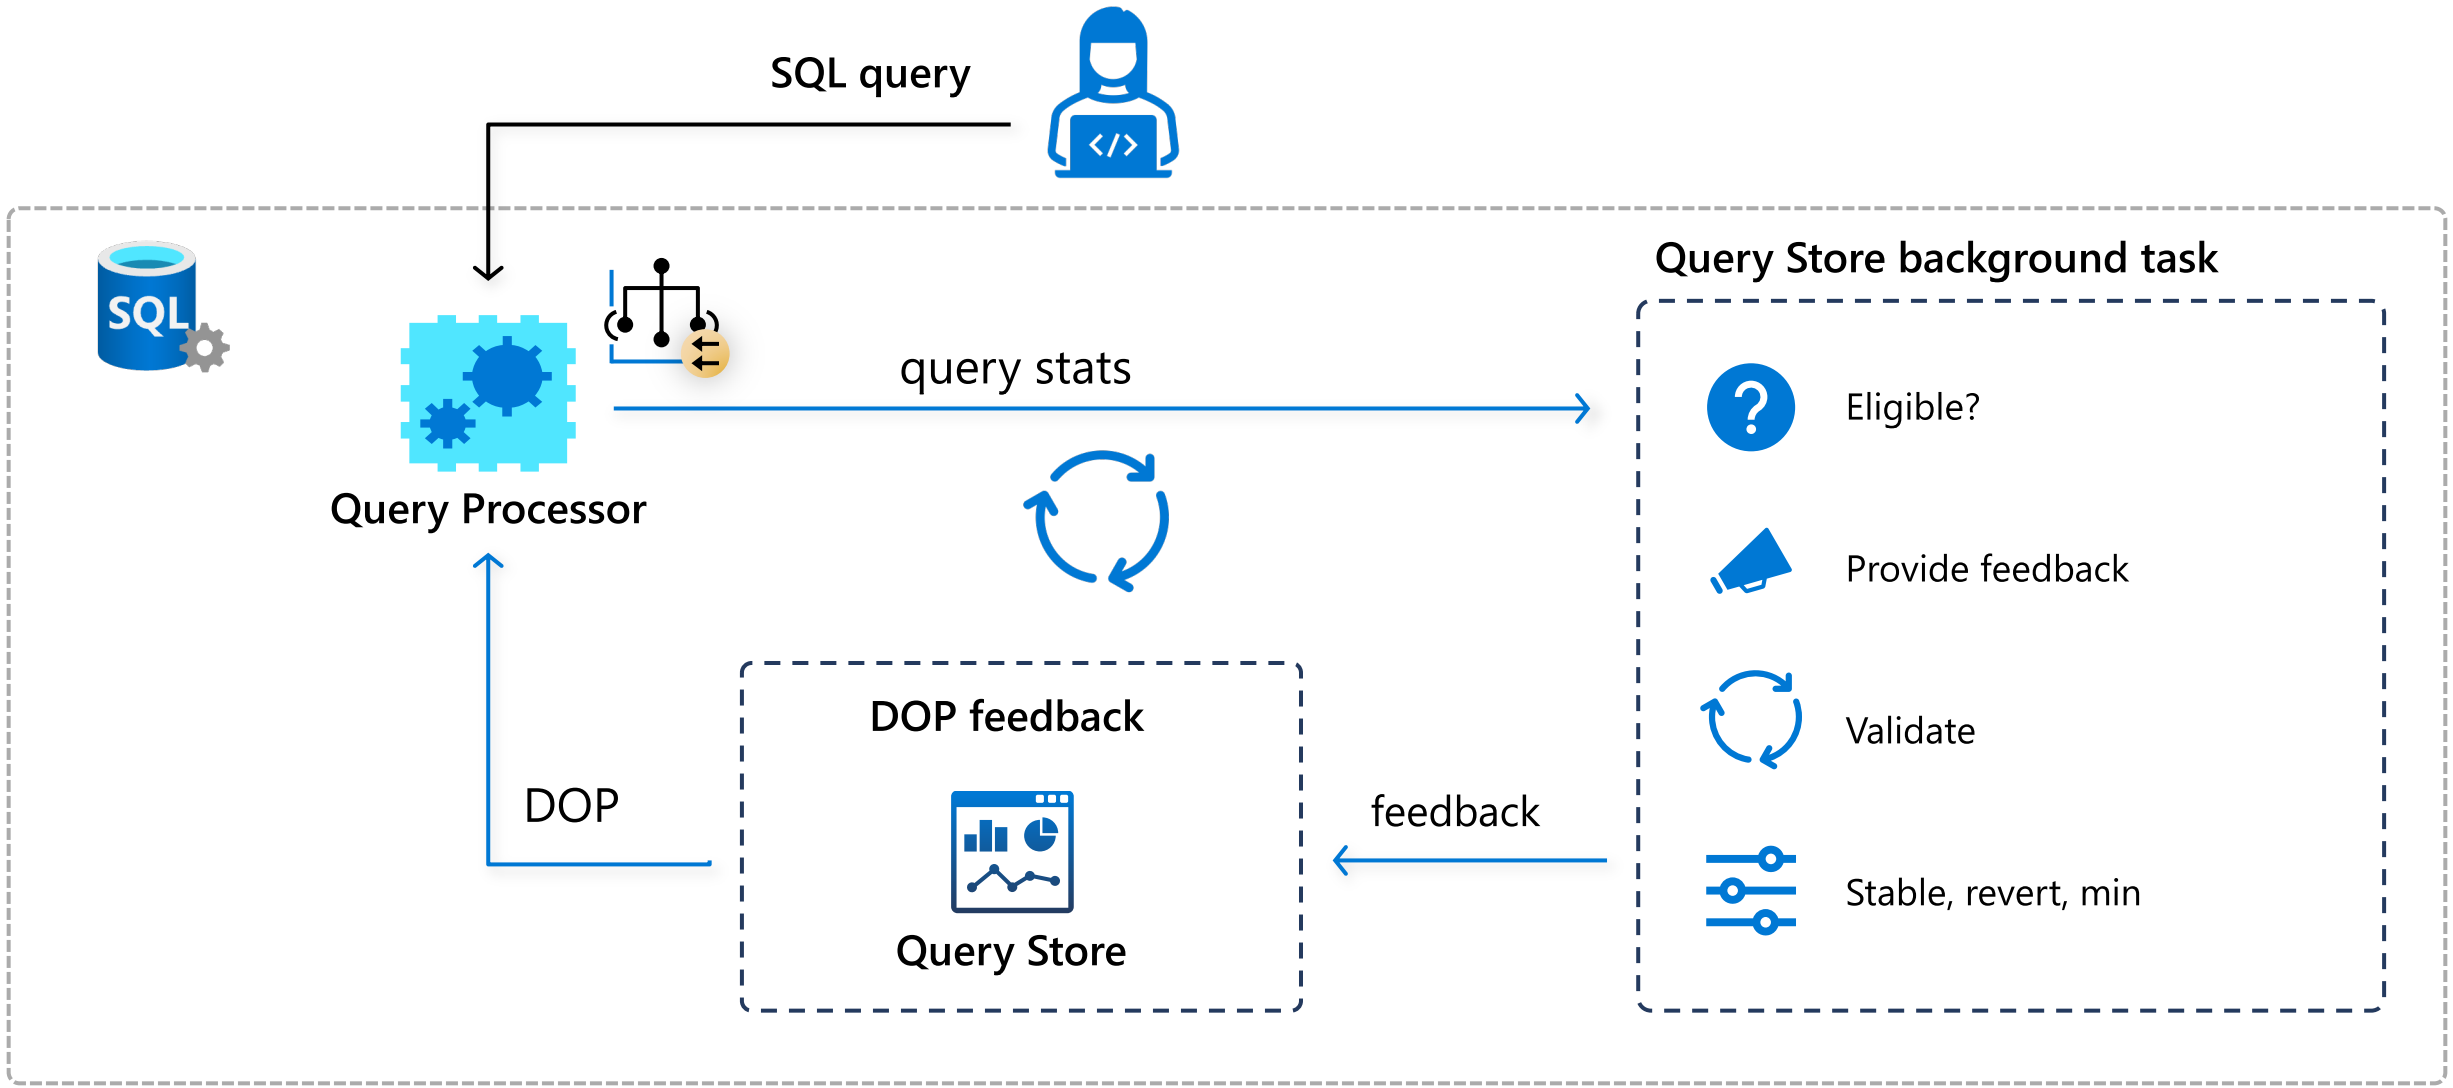

# DOP feedback details

- Eligible queries must have a longer duration and repetitive
- It's all about lowering overall CPU usage of the query
- Less CPU for DOP query = more CPU for everyone else
- Start from MAXDOP precedence setting and *lower (never higher)*
- Factor out common waits = Focus on CPU
- Analyze again at recompile

# “IQP is a symphony” – Kate Smith

PSPO creates a variant plan

CE feedback creates a query store hint

MGF feedback adjusts the memory grant

DOP feedback adjusts DOP

Microsoft

# Exercises

## Exercise 3.1
## Parameter Sensitive Plan (PSP) optimization

## Bonus: Exercise 3.0
## Persisted Memory Grant Feedback

## Advanced: Exercise 3.2
## DOP feedback

# Knowledge check

**Query Store is now on by default for SQL Server 2022 under which circumstances:**

**A.** Create a new database with SSMS in SQL Server 2022

**B.** Use CREATE DATABASE in SQL Server 2022

**C.** Restore a database backup from SQL Server 2019

**D.** A and B

**E.** Restore a version compatible backup from Azure SQL Managed Instance

**Parameter Sensitive Plan optimization can help which scenario:**

**A.** Data is skewed and you use ad-hoc queries with your application

**B.** Scenarios where more than one plan in cache would help with stored procedure performance

**C.** Stored procedure performance using parameters in filters with >= operators

**D.** All of the above

# Knowledge check

## CE Feedback works by:

**A.** Applies a query hint in plan cache

**B.** Disables the legacy cardinality estimation model

**C.** Triggers a recompile to see if the query gets a different plan

**D.** Dynamically change between legacy and new CE model

**E.** Persist a query store hint for queries that improve in performance

## Parallel efficiency is defined as:

**A.** Dynamically calculating the best setting for sp_configure

**B.** The lowest DOP setting for a query to achieve similar duration across executions while lowering CPU

**C.** A MAXDOP setting for the database that lowers the overall CPU for the system

**D.** B and C

# Summary

The **new Query store** is on by default for new databases, provides hints, works with replicas, and is a store for new query intelligence

Upgrade to SQL Server 2022 with **any dbcompat** and get approx. percentile and optimized plan forcing

Upgrade and use **dbcompat 140+** to enable memory grant percentiles and feedback

Upgrade and use **dbcompat 160** to light up parameter sensitive plan optimization, ce feedback, and dop feedback

Don’t forget all the great innovations from **IQP Gen1 and 2** that come with dbcompat 140 and 150+

Module 4

# Power your database with security, scalability, and availability

Overview

Security enhancements

Scalability enhancements

Availability enhancements

Purvi's List

Exercises

Knowledge Checks and Summary

# Industry-proven database engine

## Security, scalability, and availability

Entain

### Secure

Ledger for SQL Server

Always encrypted enhancements

Strict connected encryption (TDS 8.0)

### Scalable

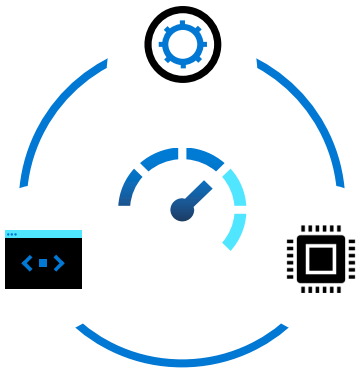

Buffer pool parallel scan

“Hands-free” tempdb

Auto async update stats concurrency

### Available

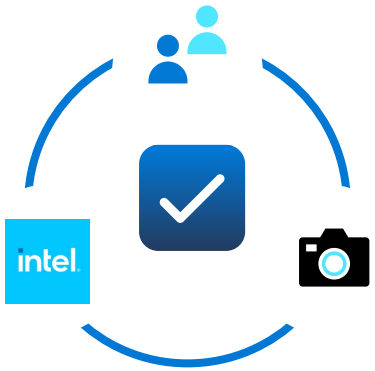


Contained Availability Groups

Cross-platform SNAPSHOT backups

Intel QuickAssist™ Technology backup compression

# Security enhancements

- Ledger for SQL Server
- Always encrypted enhancements
- Strict connected encryption (TDS 8.0 and TLS 1.3)
- PFX certificate support
- New granular fixed server roles
- Enhancements for Dynamic Data Masking

# Ledger for SQL Server

## Tamper-evidence track record of data over time

Challenge: I want the power of blockchain in a centralized system like SQL Server

- Use a cryptographically hashed ledger detect tampering by malicious actors
- Built into SQL Server with T-SQL
- Establish digital trust in a centralized system using blockchain technology.
- Attest to other parties that data integrity has not been compromised
- Automatic digest storage

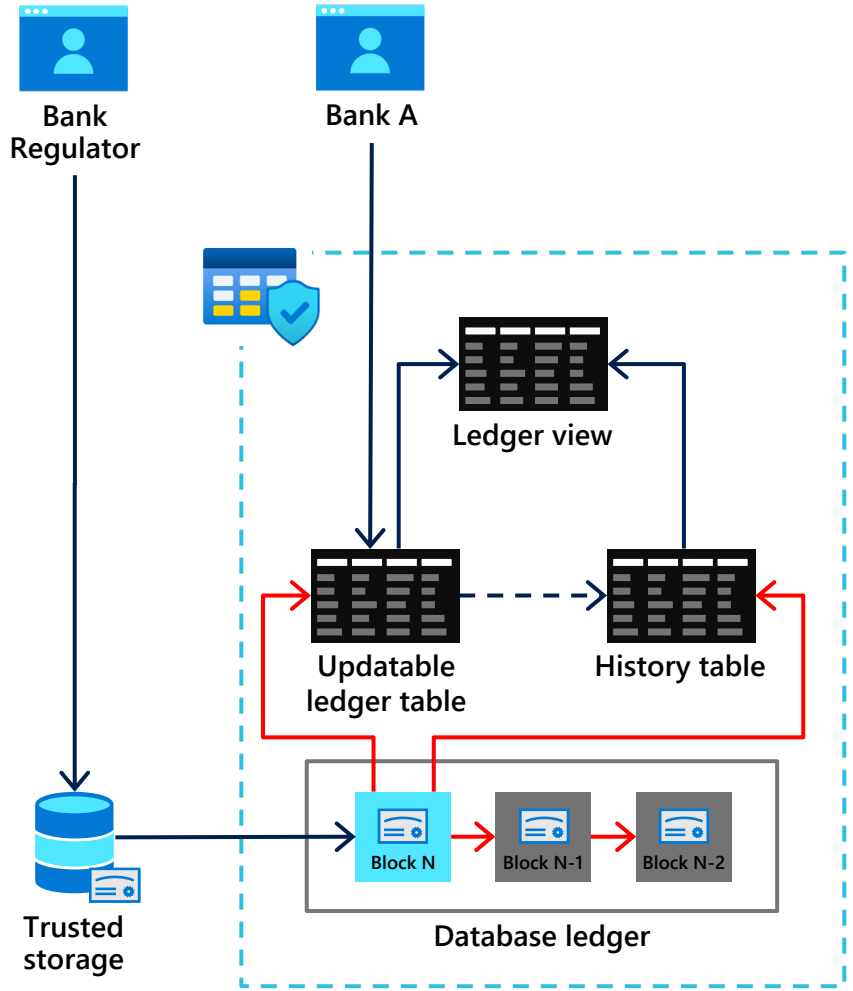

# How to use Ledger for SQL Server

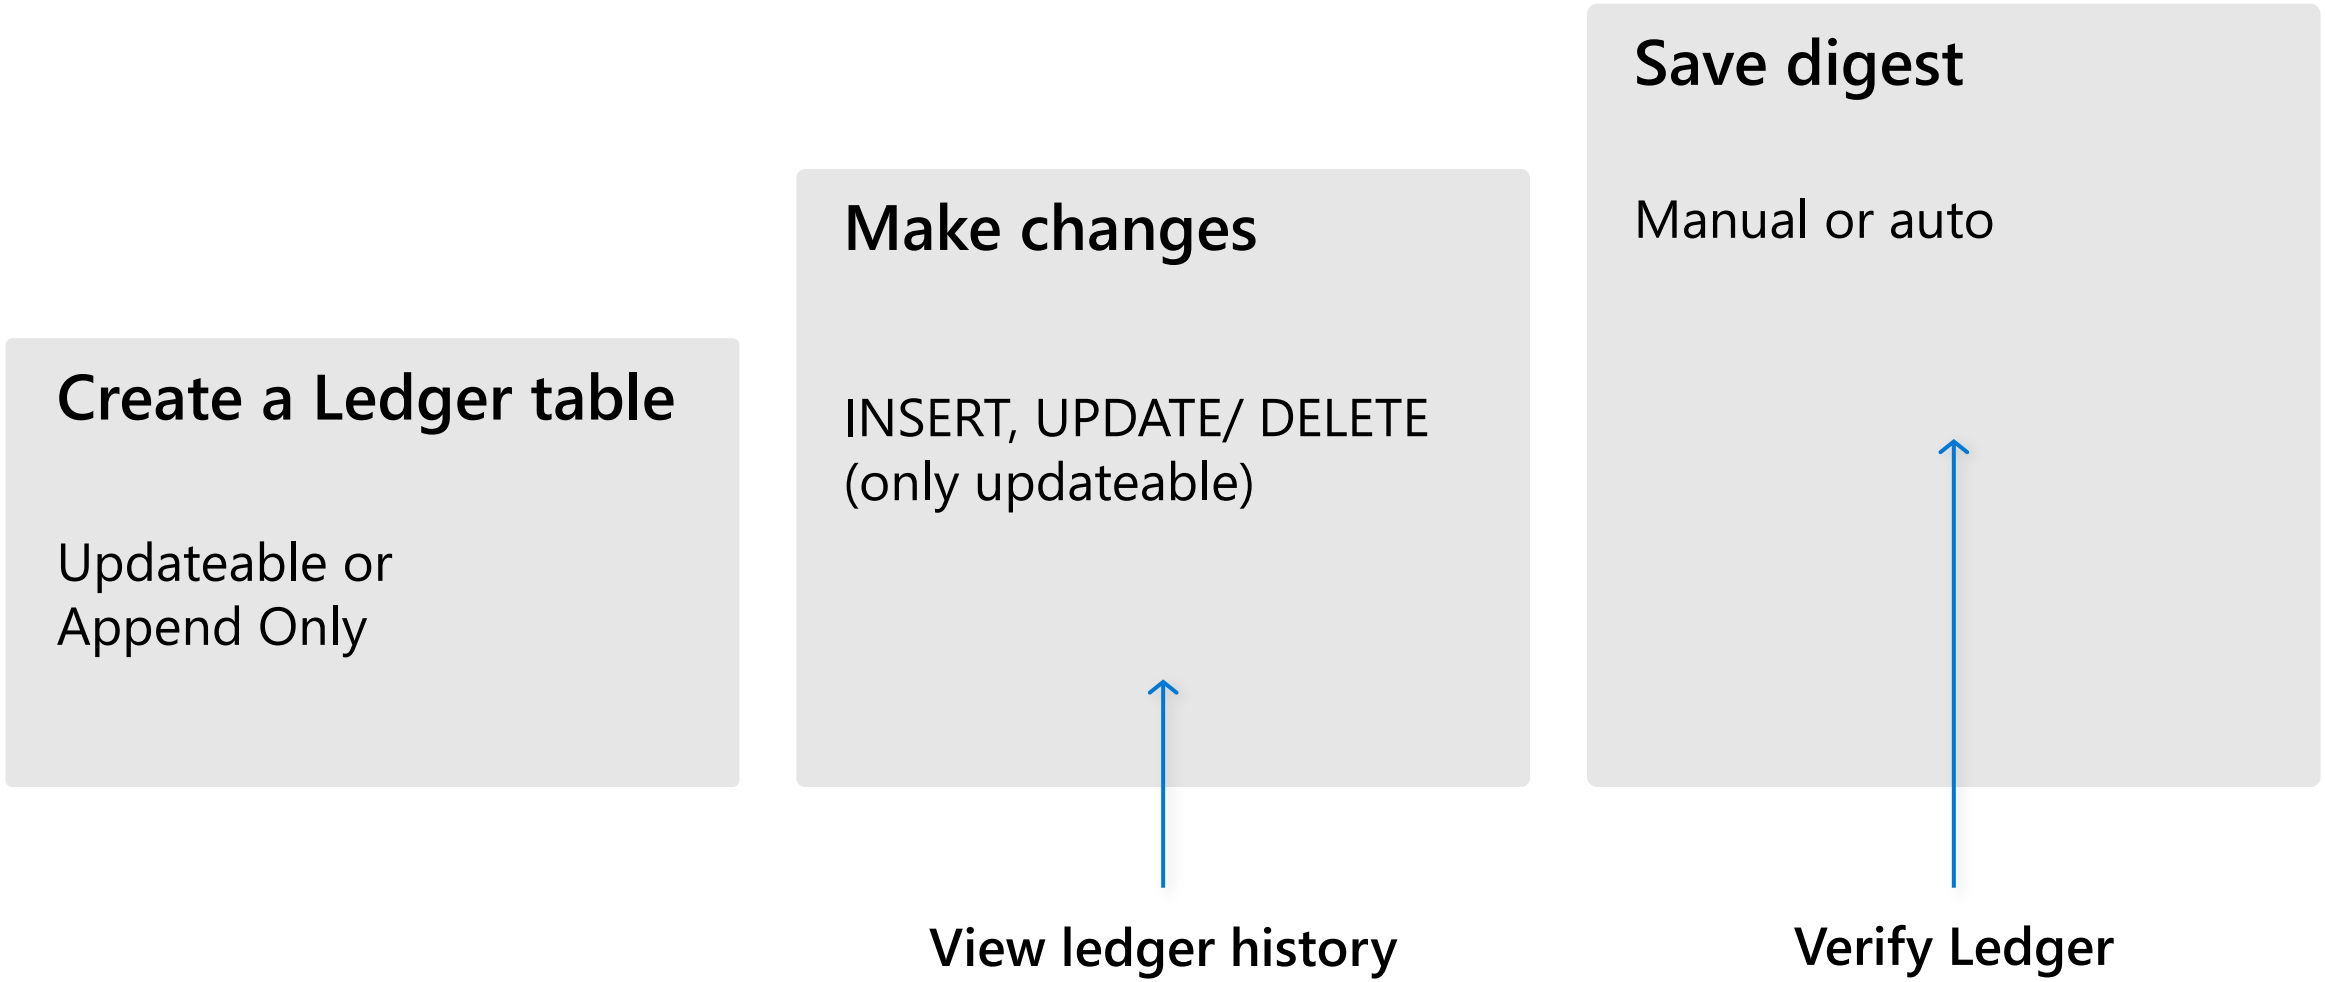

# Ledger FAQ

## How is this different from a temporal table?

- Built-in transaction auditing
- Append-only
- Database Ledger and digest for tamper evidence

## How is this different than SQL Server Audit?

- Transaction history and audit built into database
- Digest verification

## Can I ALTER a ledger table?

- You cannot ALTER a table to “turn off” ledger or disable system versioning
- You can alter the schema but there are some limitations

## Can I drop a ledger table?

- Yes, but a history of the dropped table and ledger is kept

# Ledger FAQ

## How often do I need to save the digest?

- As frequent as you need to ensure the ledger is tamper proof

## Does Ledger require more space?

- Updateable requires similar extra space as temporal
- **Plus** database ledger requires some minimal extra space for hashes and blockchain
- Digests small and separate but you may need a long history
- You can’t archive or truncate ledger tables and database ledger

## Any perf impact?

- Append only should see minimal impact
- Updateable would have similar impact as temporal

# Scalability enhancements

- Buffer pool parallel scan
- “Hands-free” tempdb
- Auto async update stats concurrency
- Auto drop statistics
- XML compression
- Advanced Vector Extension (AVX) 512 extension to improve batchmode operations.
- Improved columnstore segmentation
- Ordered clustered columnstore index

# “Hands-free” tempdb

Tempdb lazy log flush
Tempdb != FUA

**Pre SQL Server 2019**

- 1 file = PFS, GAM, SGAM contention
- Add multiple files
- Trace flags 1117 and 1118
- SQL 2016 setup auto adds multiple files
- Trace flags not required by SQL Server 2016

**SQL Server 2019**

- PFS concurrency
- Autogrow and uniform default for tempdb
- Now system table pages become hotspot
- Tempdb metadata optimization ON
- SGAM and GAM contention remain

**SQL Server 2022**

- SGAM and GAM concurrency
- Latch contention gone

# Availability enhancements

- Contained Availability Groups
- Cross-platform SNAPSHOT backups
- Intel QuickAssist™ backup compression
- ADR enhancements
- In-memory OLTP memory enhancements
- Shrink database concurrency
- Parallel redo enhancements
- DAG enhancements
- Resumable ADD CONSTRAINT
- CREATE INDEX concurrency
- Last valid restore time in backupset msdb table
- LWW wins auto conflict peer-to-peer replication

# Contained Availability Groups

## Synchronize instance level objects with data

Challenge: I have to manually synchronize instance level objects with Availability Groups

- Create an Availability Group and declare it be CONTAINED
- We add a contained master and msdb database to the AG
- Connect with listener and contained master and msdb are used
- Connect directly and default master and model are used.
- SQL Server agent enhanced to understand contained msdb

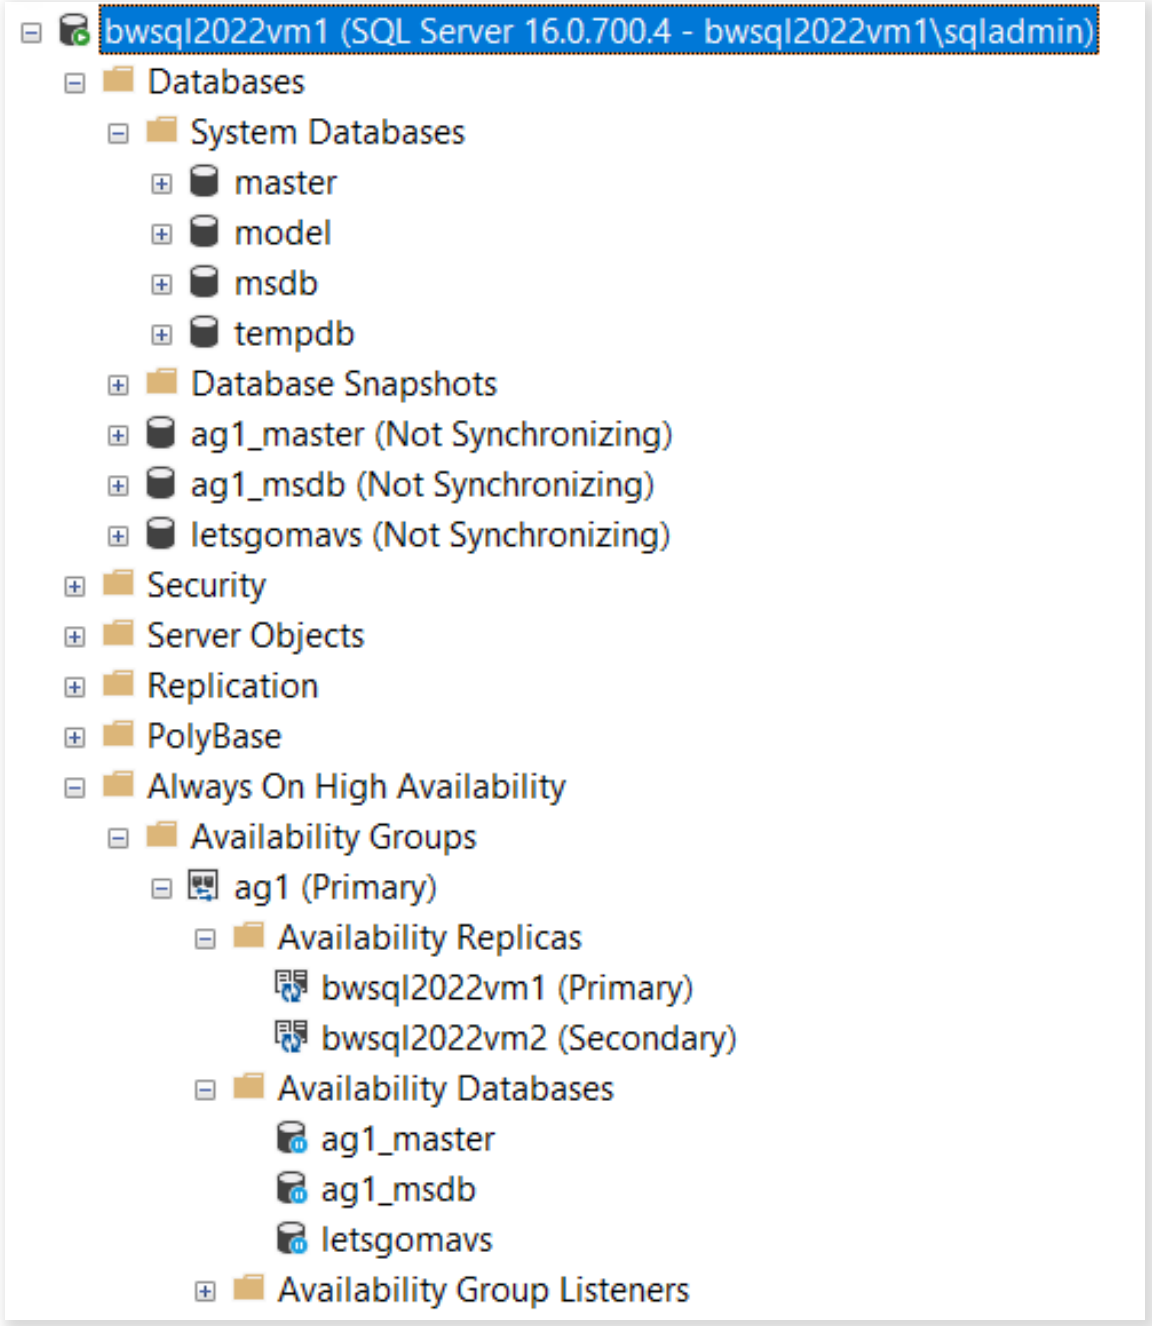

# How Contained AGs work

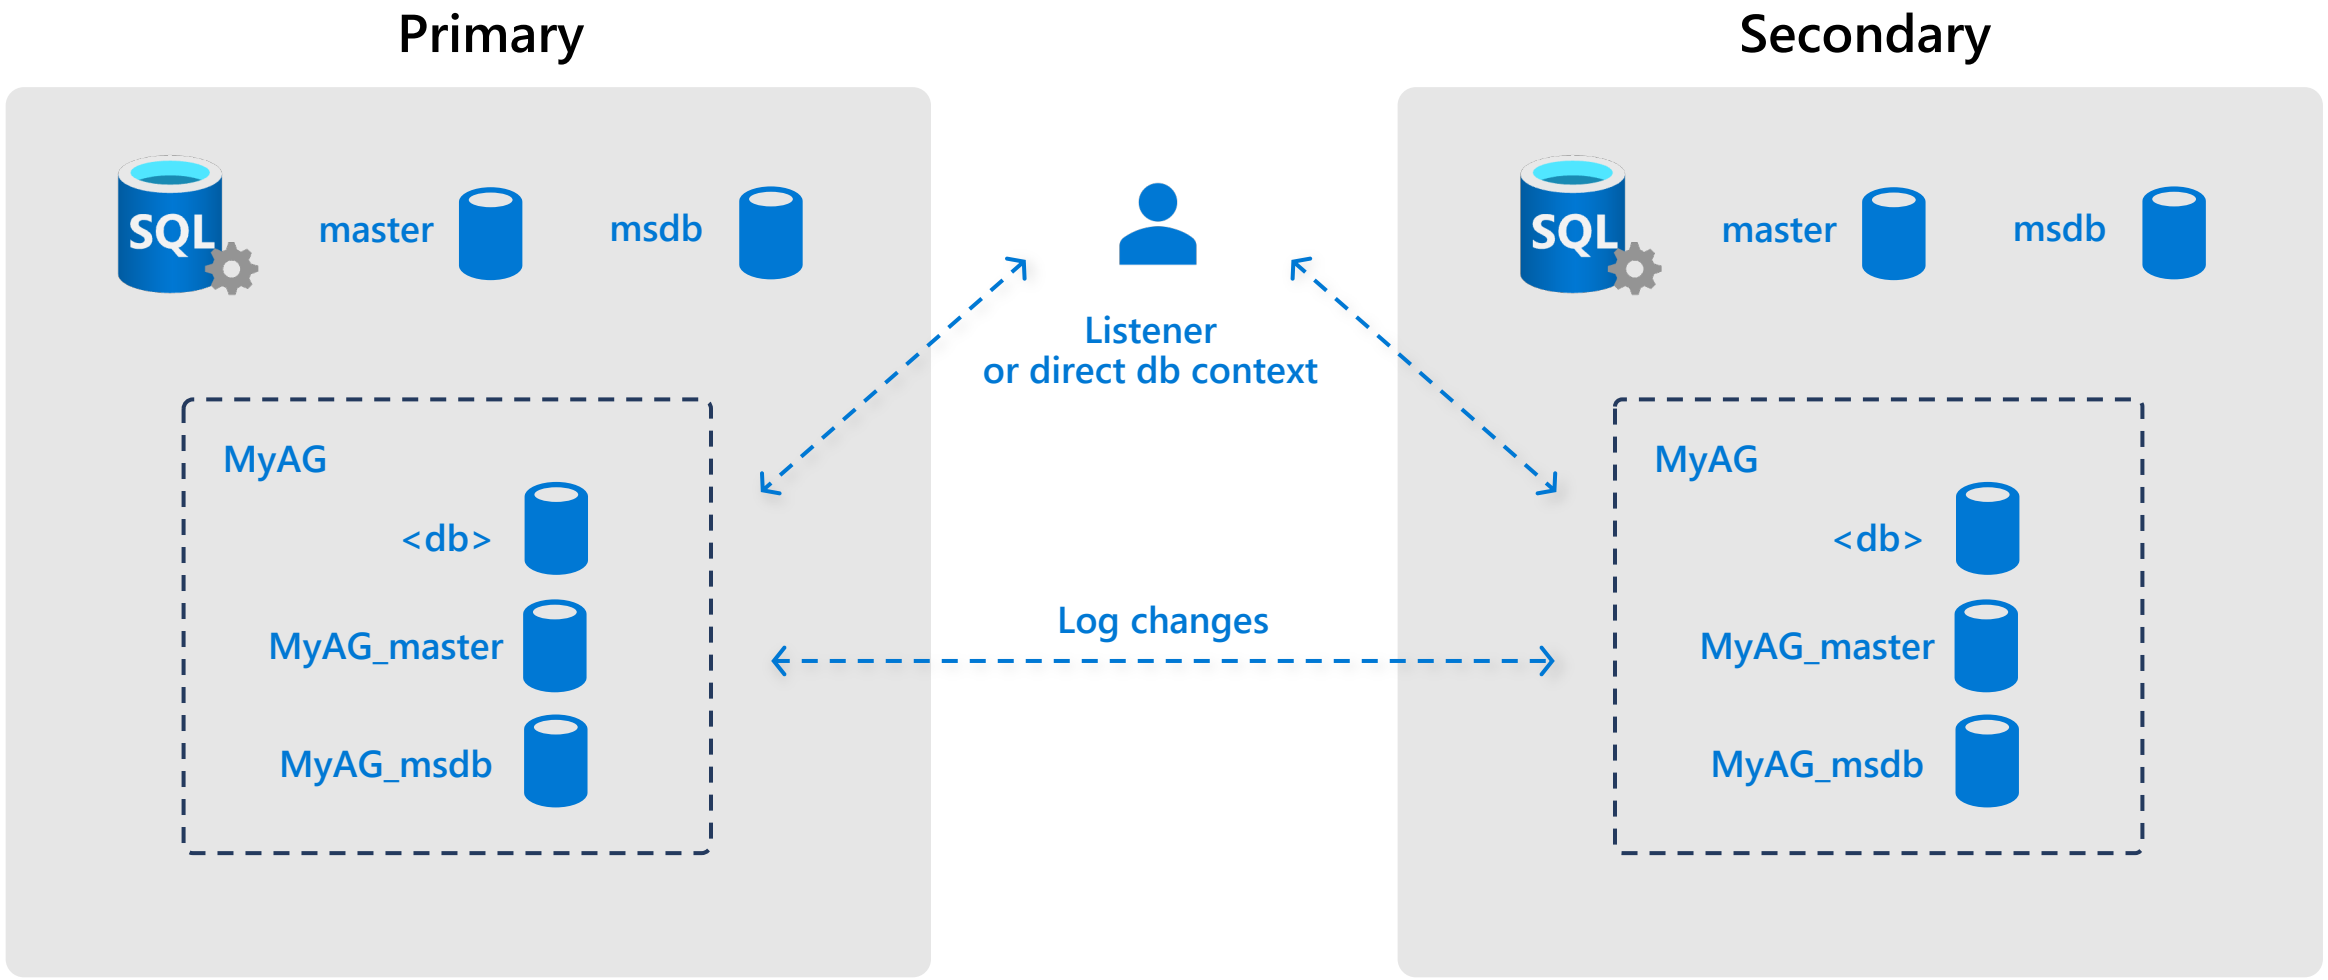

Microsoft

# Demo

## Exercise 4.2
## Contained Availability Groups

https://aka.ms/sql2022mechanicsupdate

# Cross-Platform Snapshot Backups

## Before

VDI Program

VSS and SQL Writer

## Now

T-SQL suspend I/O

Perform your snapshot

T-SQL resume I/O

- Why? Backup TB in seconds or milliseconds? Constant performance regardless of size
- Why would I halt I/O vs *normal* BACKUP? **ANSWER**: Restore faster!
- Use storage that supports snapshots—Tip: Azure Storage does by default
- Is this a SQL backup? **ANSWER**: Not a separate backup stream but you can restore it
- Works on Windows and Linux
- Users most likely see WRITELOG waits on suspend

**Tip**: Use undocumented and unsupported trace flag 3661 to test it

# SQL Server 2022—Hardware acceleration & offloading

## Current customer demands

Stability for critical workloads

Consistent performance without oversizing

Remove scalability blockers

Make the most of current hardware

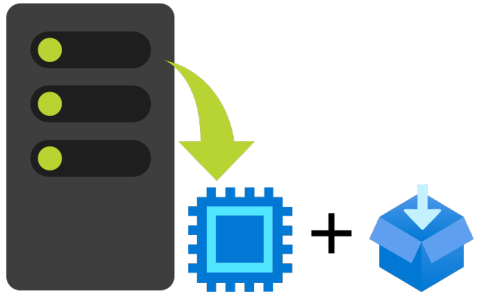

## Intel QuickAssist Technology (QAT)

Intel QuickAssist Technology (QAT) provides a compression solution that integrates hardware acceleration of compute-intensive workloads on Intel platforms.

Intel QuickAssist Technology (QAT) with SQL Server 2022 off-loads data compression, enabling faster OLTP operations while reducing the impact on the system's transaction throughput.

**Database Backup Improvements for SQL Server 2022 with Intel QuickAssist Technology (QAT):**

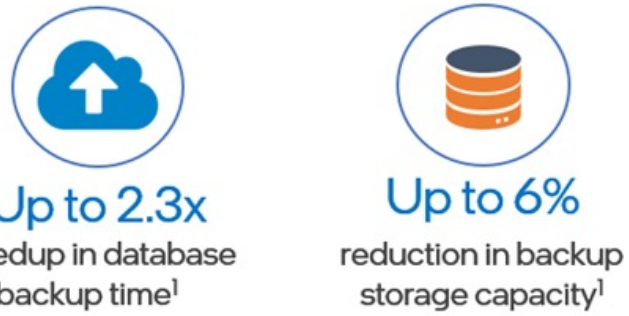


[1] **SOFTWARE BASELINE:** Test by Intel as of 10/26/2021. 1-node, 2x Intel® Xeon® Platinum 8358 Processor on Wilson City (ref. platform), 32 cores with 1024 GB (16 slots/ 64GB/ 3200[3200]) total DDR4 memory, ucode 0xD000323, HT on, Turbo on, Windows Server 2019,10.0.17763.1999, [Hyper-V VM], 1x Intel X550, SQL Server 2022 CTP1.1 private test build, 525GB Database, Database Backup with Compression. **WITH INTEL® QUICKASSIST ADAPTER 8970 TECHNOLOGY:** Test by Intel as of 10/26/2021. 1-node, 2x Intel® Xeon® Platinum 8358 Processor on Wilson City (ref. platform), 32 cores with 1024 GB (16 slots/ 64GB/ 3200[3200]) total DDR4 memory, ucode 0xD000323, HT on, Turbo on, with Intel® QuickAssist Adapter 8970, 1.70.16.4 driver version, Windows Server 2019,10.0.17763.1999, [Hyper-V VM], 1x Intel X550, SQL Server 2022 CTP1.1 private test build, 525GB Database, Database Backup with Compression. **WITH INTEL® Intelligent Storage Acceleration Library:** Test by Intel as of 10/26/2021. 1-node, 2x Intel® Xeon® Platinum 8358 Processor on Wilson City (ref. platform), 32 cores with 1024 GB (16 slots/ 64GB/ 3200[3200]) total DDR4 memory, ucode 0xD000323, HT on, Turbo on, with Intel® Intelligent Storage Acceleration Library v2.30.0.0, Windows Server 2019,10.0.17763.1999, [Hyper-V VM], 1x Intel X550, SQL Server 2022 CTP1.1 private test build, 525GB Database, Database Backup with Compression. Results may vary.

# Intel QAT Details

## How to use

- Install Intel drivers
- Hardware or software based
- Server configuration option
- BACKUP WITH COMPRESSION option

## Preview observations

- Idle systems don’t benefit from QAT in hardware mode
- Hardware mode shines with multiple backup files and heavy CPU workloads
- QAT in software mode will usually outperform the default XPRESS

# Purvi's list

- Reduce Buffer Pool I/O promotions
- Enhanced core spinlock algorithms
- Improved Virtual Log File (VLF) algorithms
- Instant File Initialization (IFI) *behavior* for Log file growth even with TDE enabled (does not require special privilege)

Microsoft

# Exercises

## Exercise 4.0
## Ledger for SQL Server

## Bonus: Exercise 4.1
## “Hands-free” tempdb

# Knowledge check

**Independent verification for Ledger for SQL Server is done by:**

**A.** The database ledger through system tables

**B.** The Ledger history view

**C.** Database digests

**D.** Blockchain data in SQL Server tables

**SQL Server 2022 now eliminates what type of bottleneck for tempdb:**

**A.** PAGEIOLATCH

**B.** PAGELATCH

**C.** Spinlocks

**D.** A and B

**E.** Row locks on system pages

# Knowledge check

**A Contained Availability Group replicates which objects:**

**A.** Linked Servers

**B.** SQL Server Agent jobs

**C.** Logins

**D.** All of the above

**What are benefits of using Intel QuickAssist with SQL Server 2022?**

**A.** Improved backup compression

**B.** Cross-platform backups on SQL Server for Linux

**C.** Offload CPU resources for SQL query workloads

**D.** A and C

**E.** All of the above

# Summary

SQL Server 2022 is an industry proven engine with enhancements in **security, scalability, and availability**

**Ledger for SQL Server** provides a tamper evidence record using blockchain technology

Latch page concurrency improvements made management of **tempdb "hands-free"**

**Enhancements in availability** include contained AG, cross-platform snapshot backups, and Intel QuickAssist backup compression

There are over **25+ major engine enhancements** in SQL Server 2022

Module 5

# Access new sources with data virtualization and object storage

Data Virtualization with Polybase

Data Virtualization with SQL Server 2022

Backup and Restore with S3 Object Storage

Exercises

Knowledge Checks and Summary

# Remove data silos with Polybase—SQL Server 2019

Challenge: Integrating data from numerous data silos is complex with traditional ETL methods

## Solution

**Access from multiple sources without moving data**

- Integrate data with data virtualization
- T-SQL language with external tables
- Distributed/scalable performance with Polybase

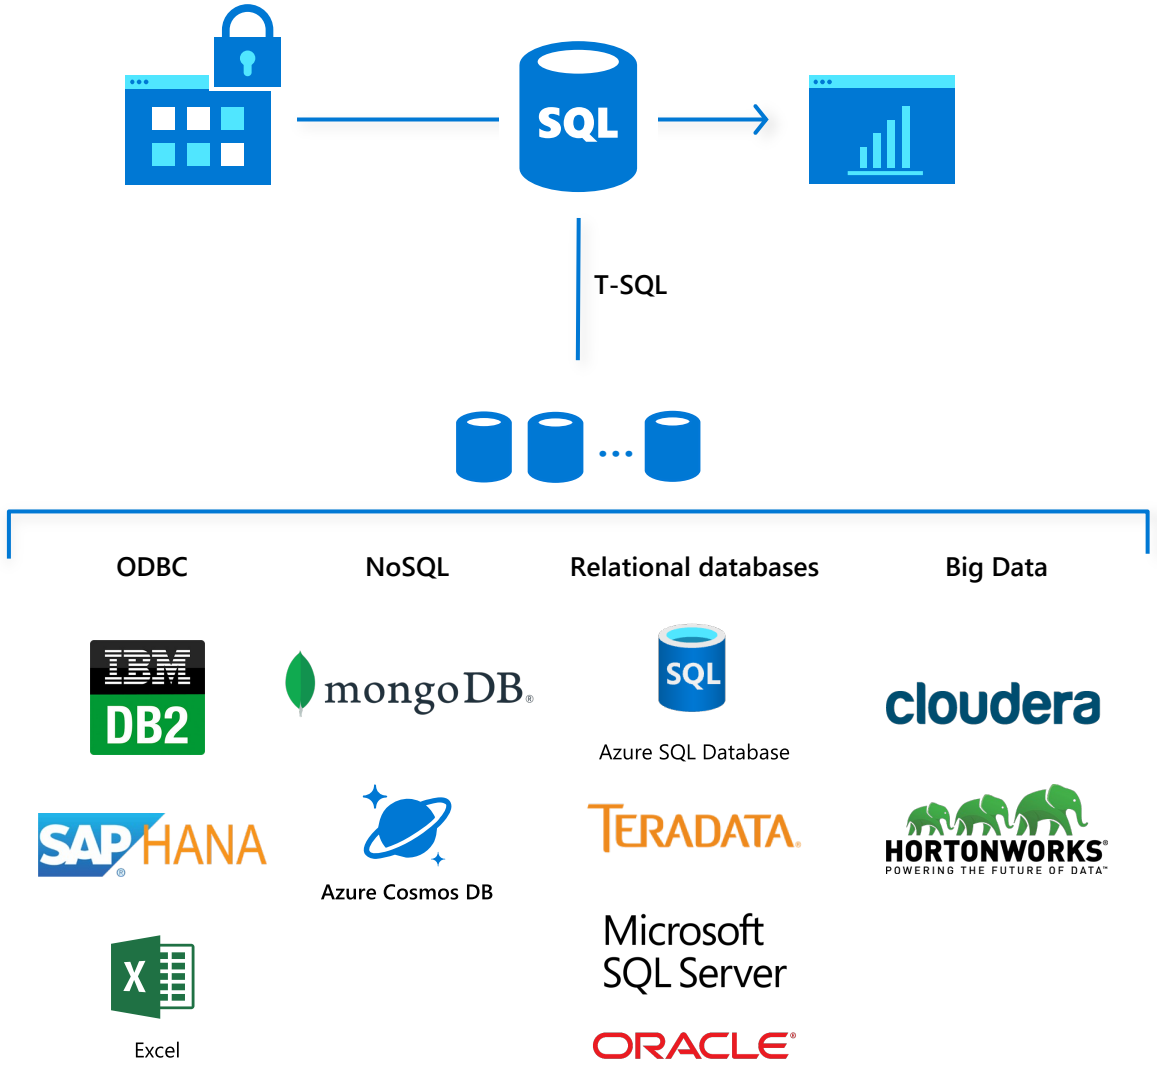

# Data virtualization and object storage

## Directly access any data on object storage

**Challenge:** I need to access data on modern object storage systems through SQL

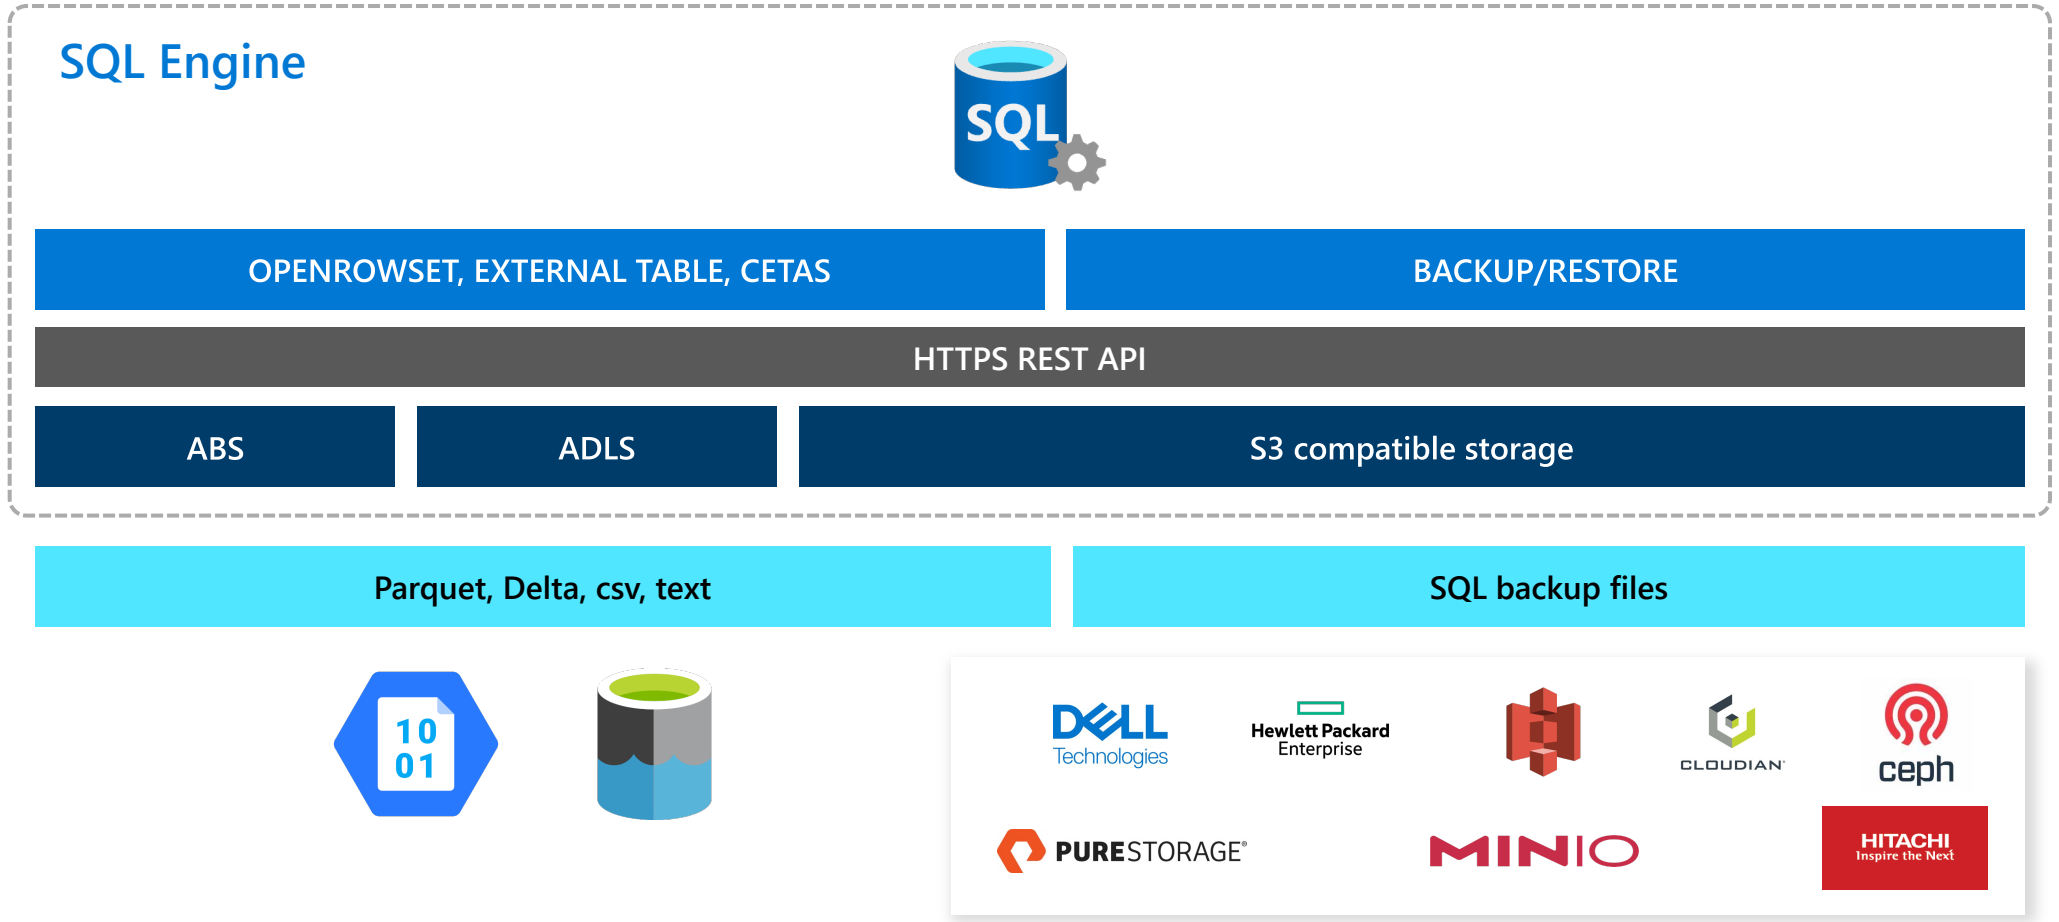

# Why S3, parquet, and delta?

S3 created by Amazon but…open protocol

Popular object storage protocol

Inexpensive, durable, unlimited

Parquet is an open-source file format

Schema built-in

Columnar format for efficient compression and retrieval


Delta lake is an open-source storage architecture with delta tables

Parquet files +

Basic transaction log support

# Data virtualization in SQL Server 2022

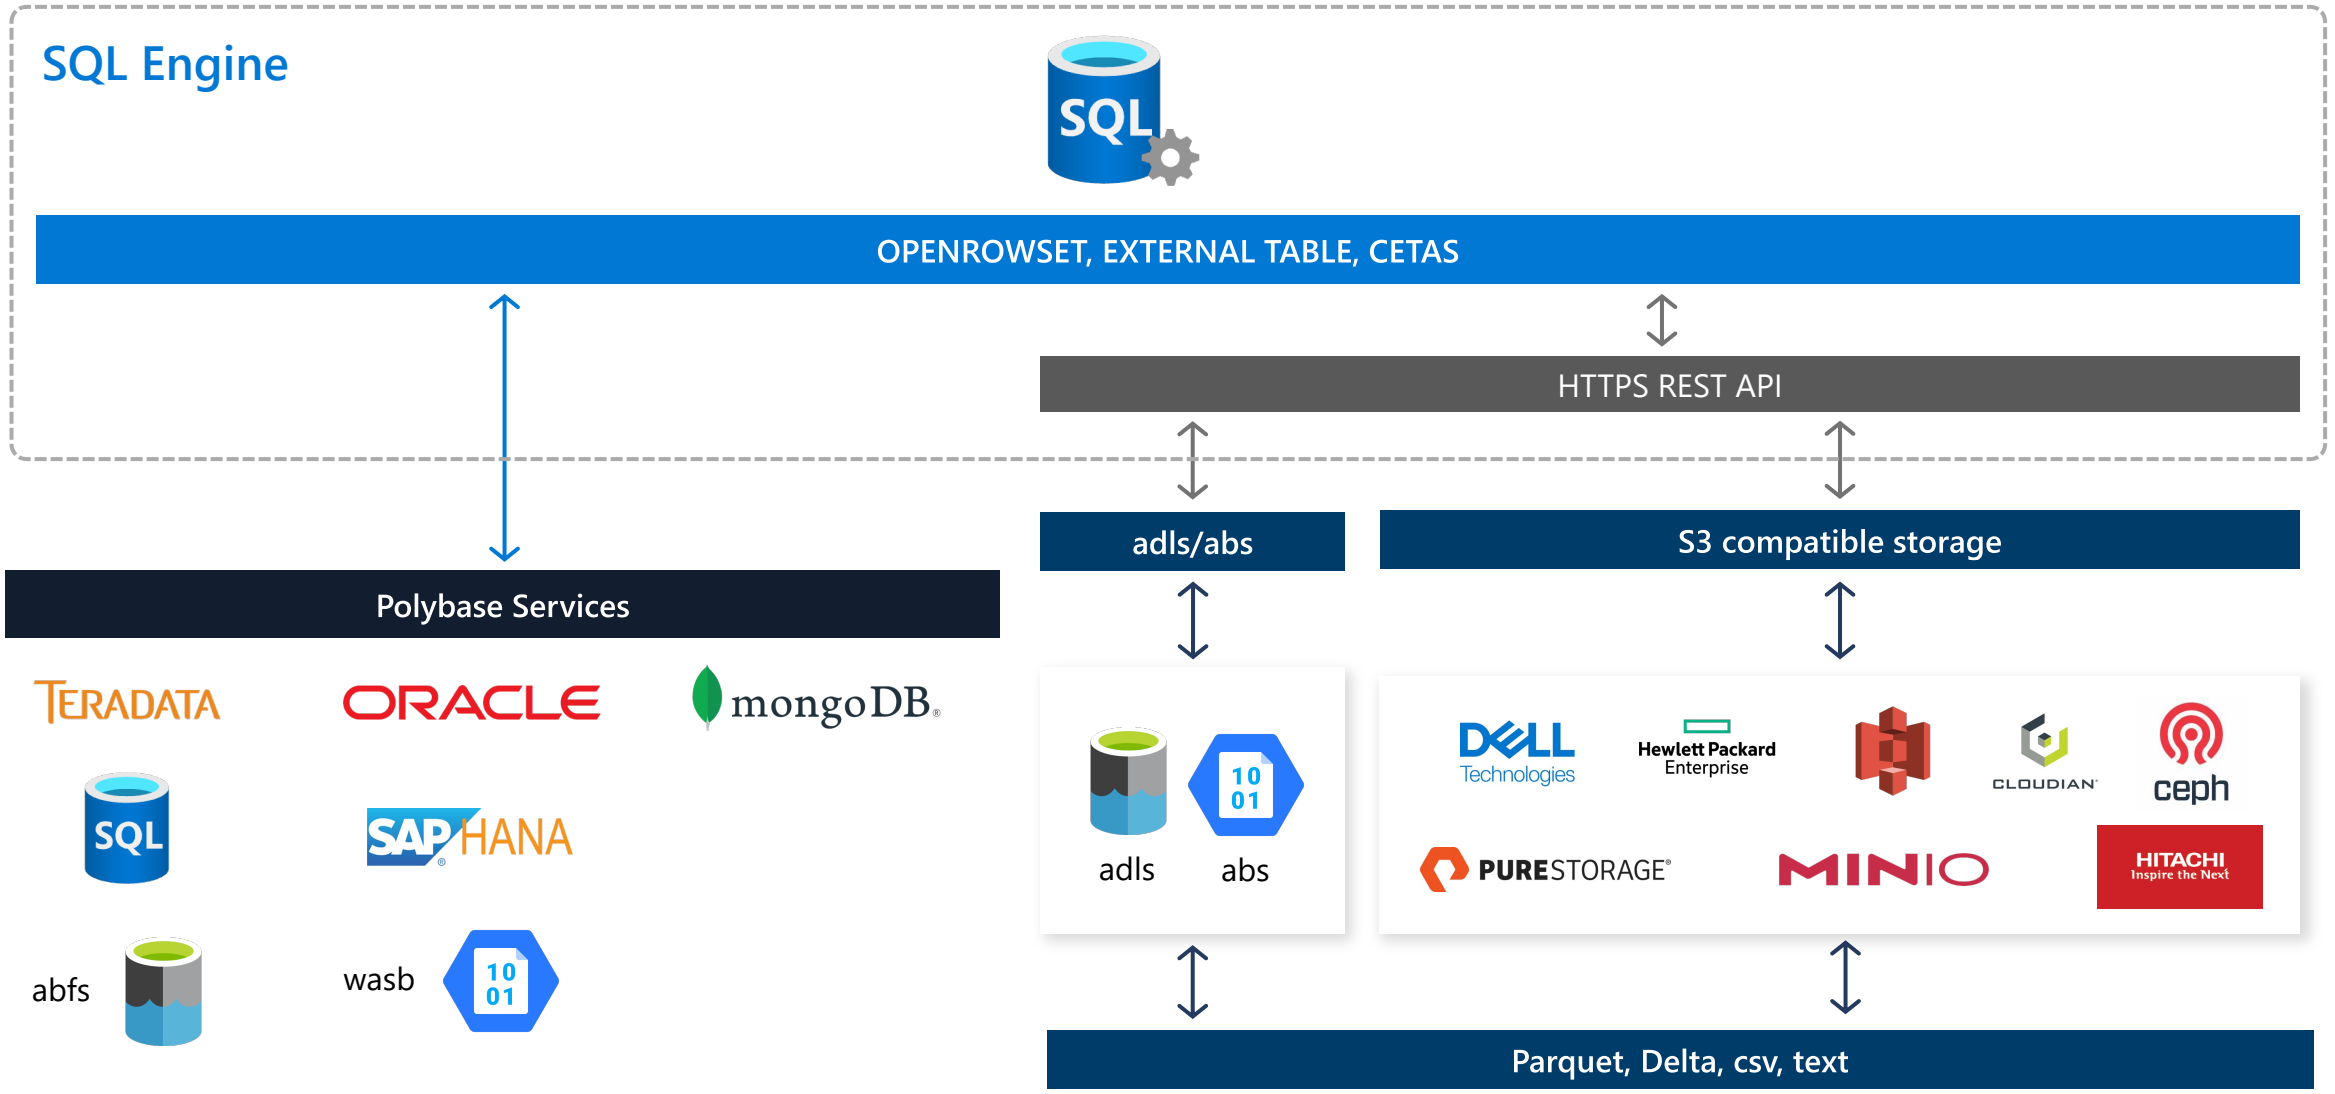

# Using REST API based data virtualization

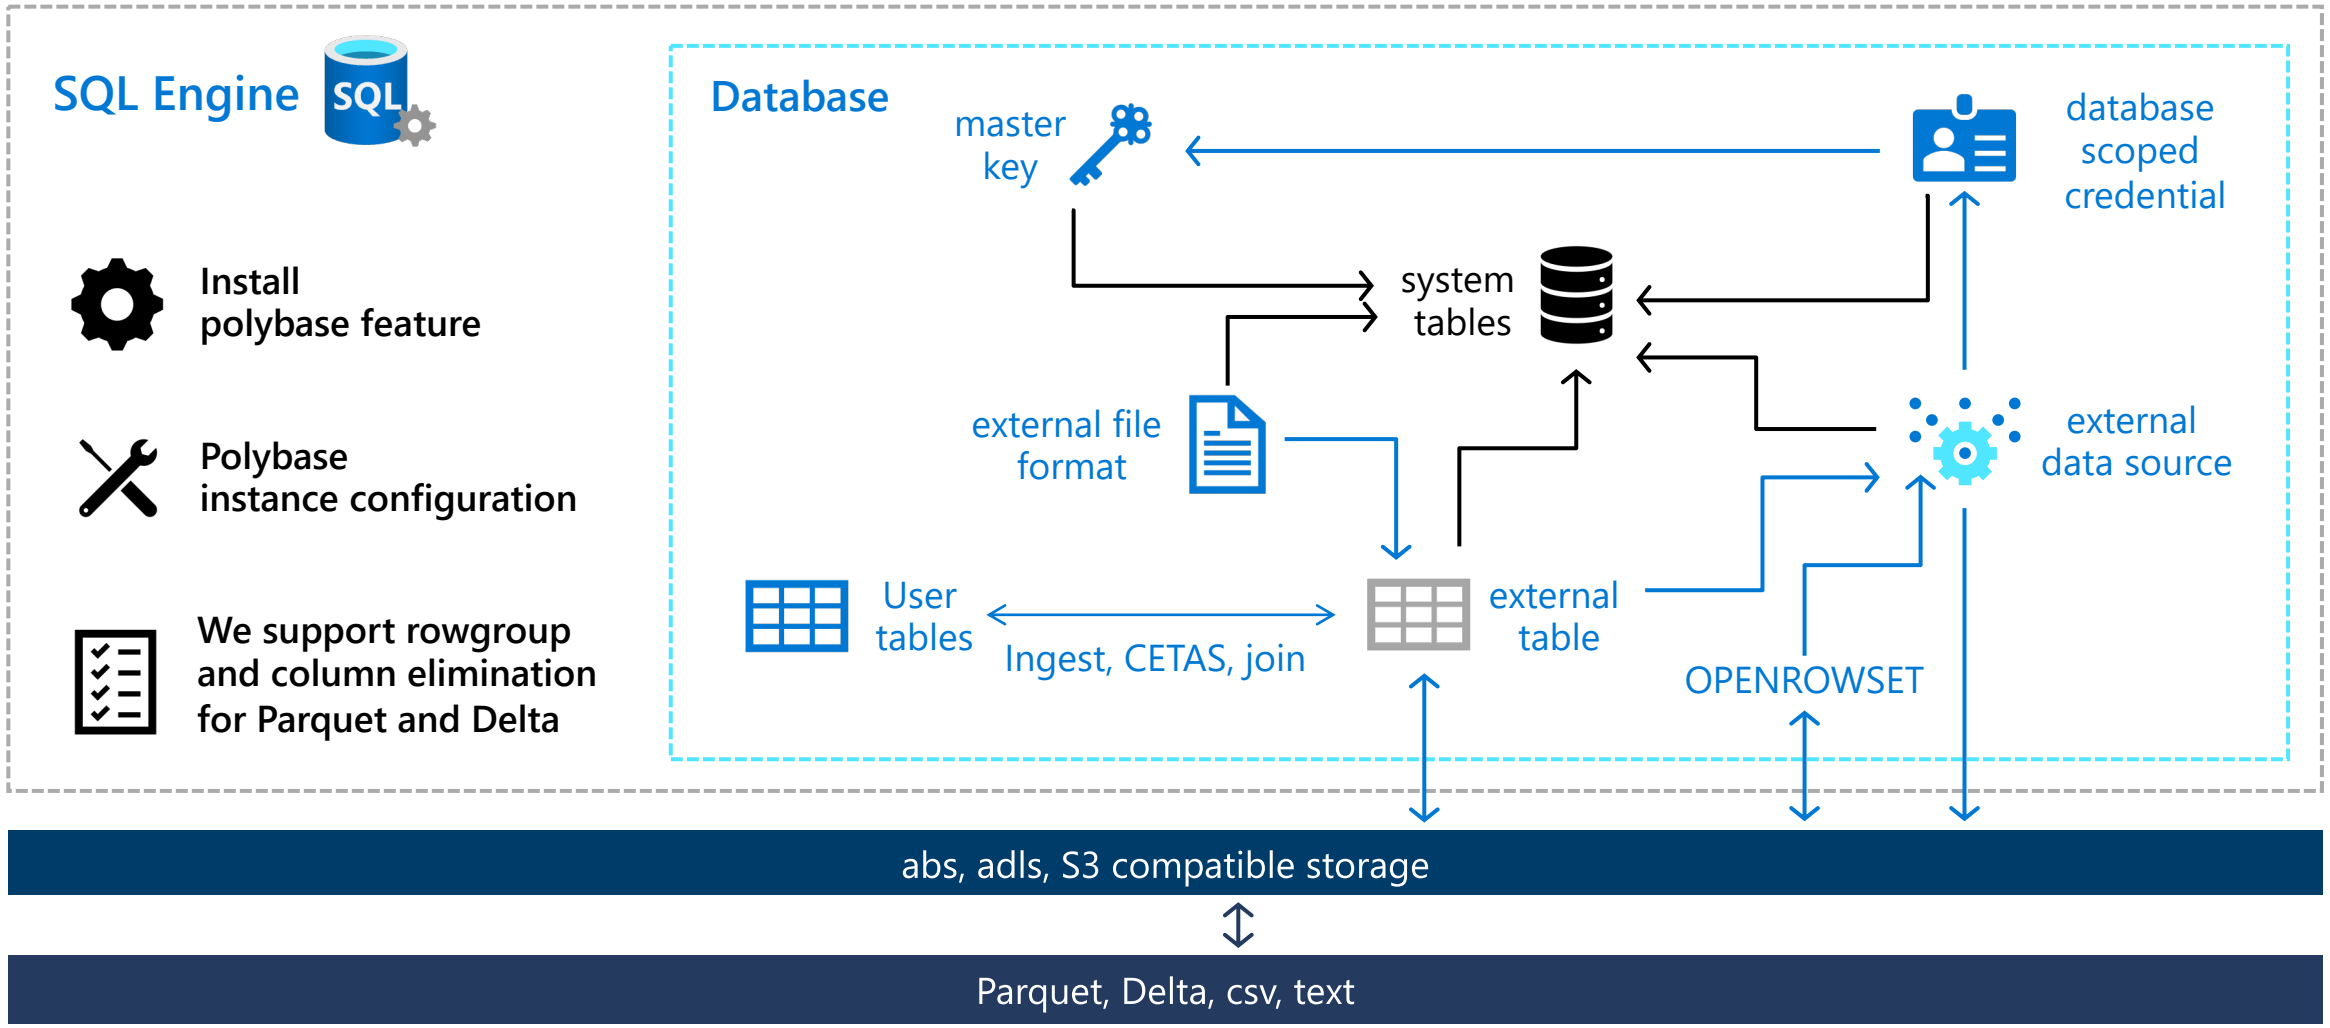

# Archiving cold data to a lake - CETAS

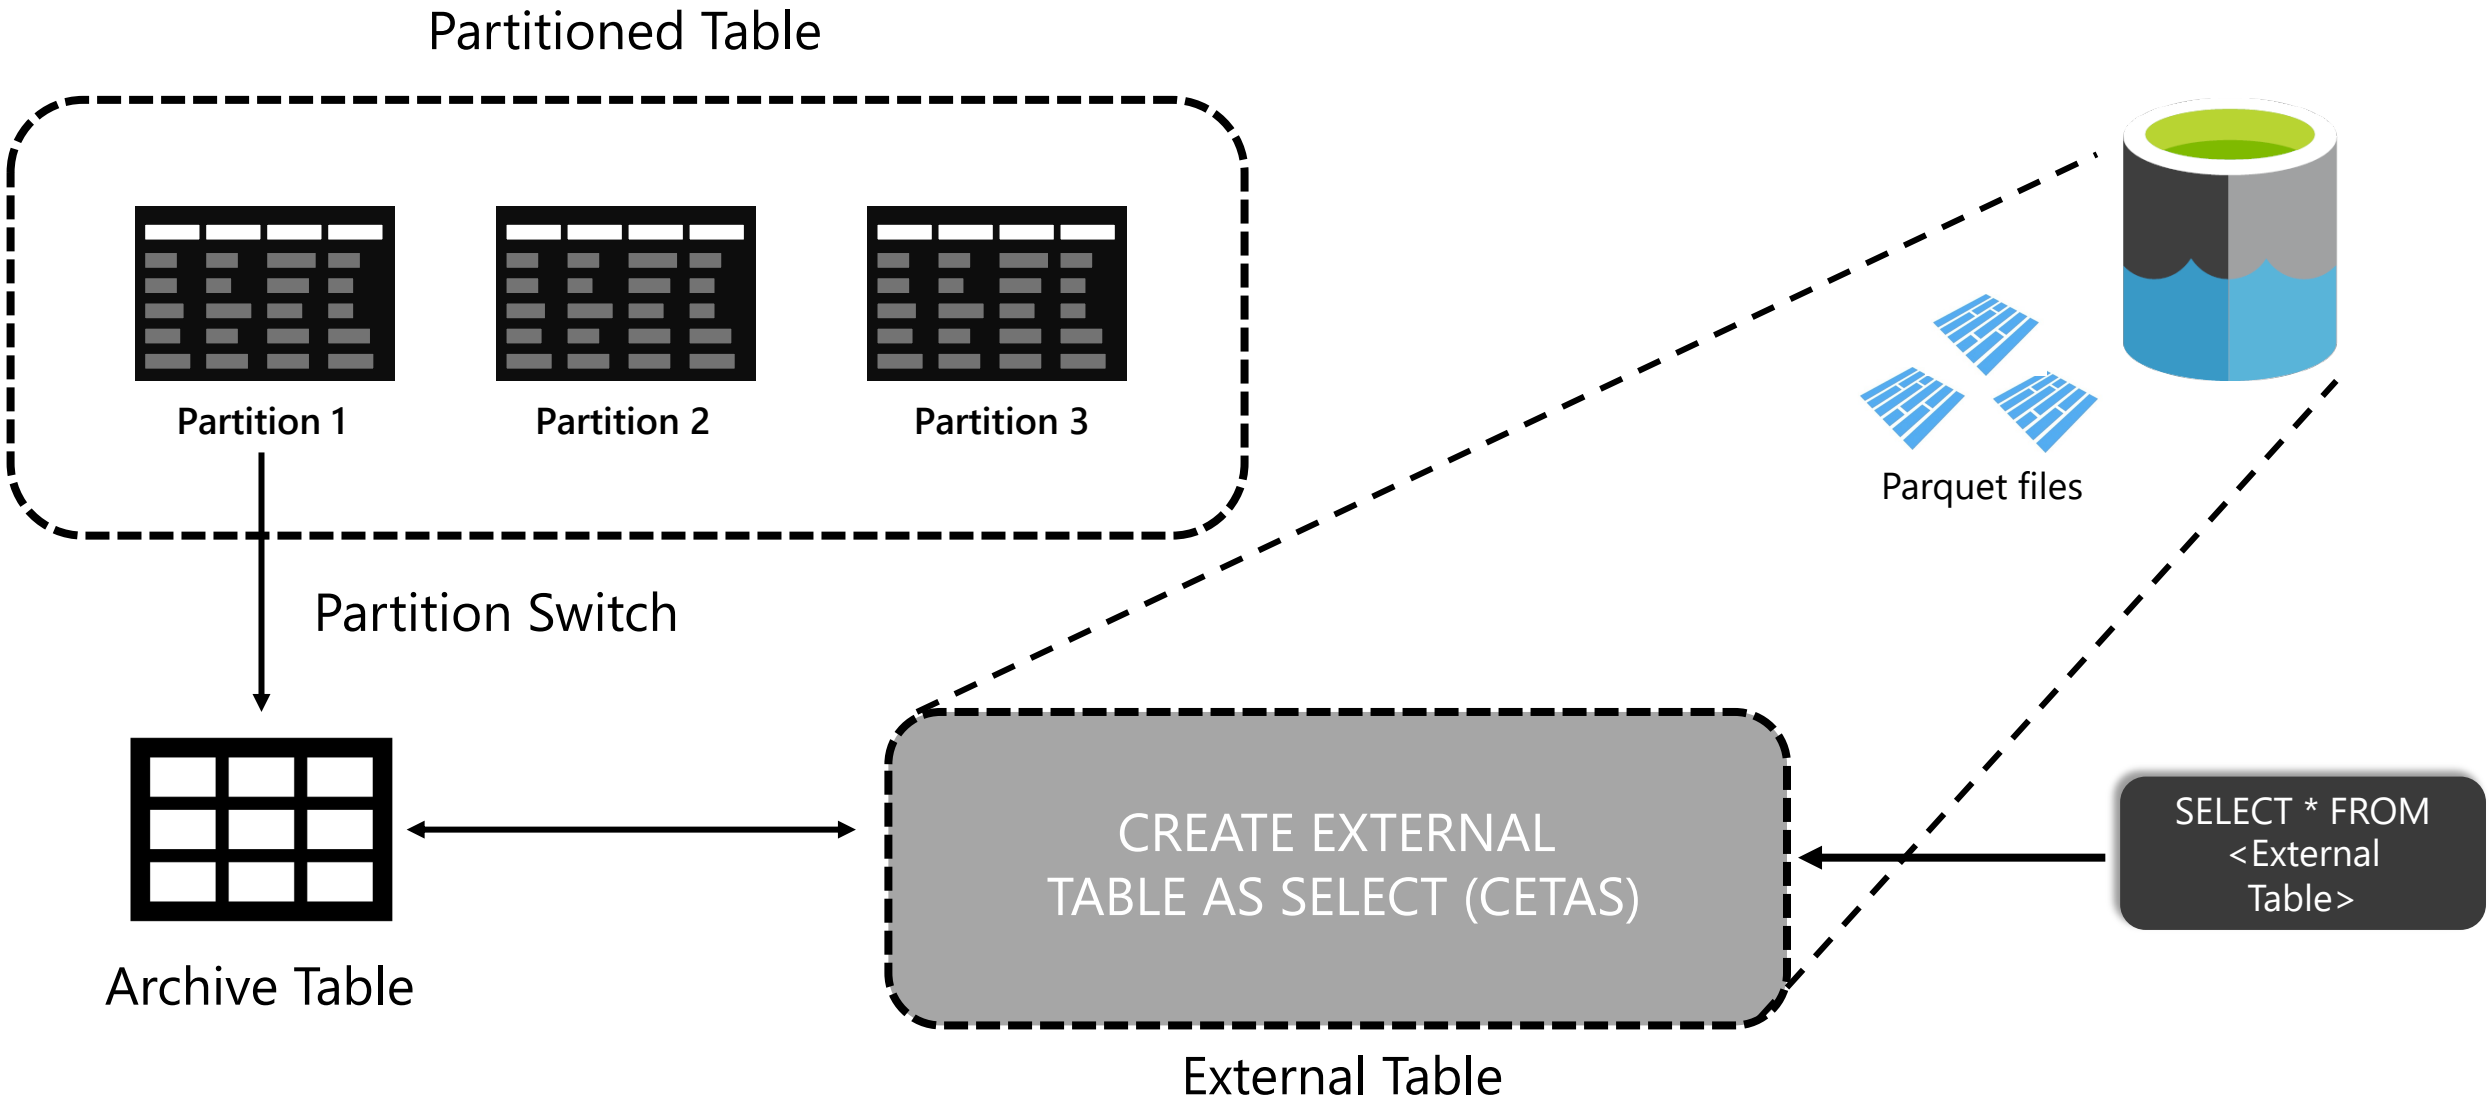

# Integrating SQL with Microsoft Fabric

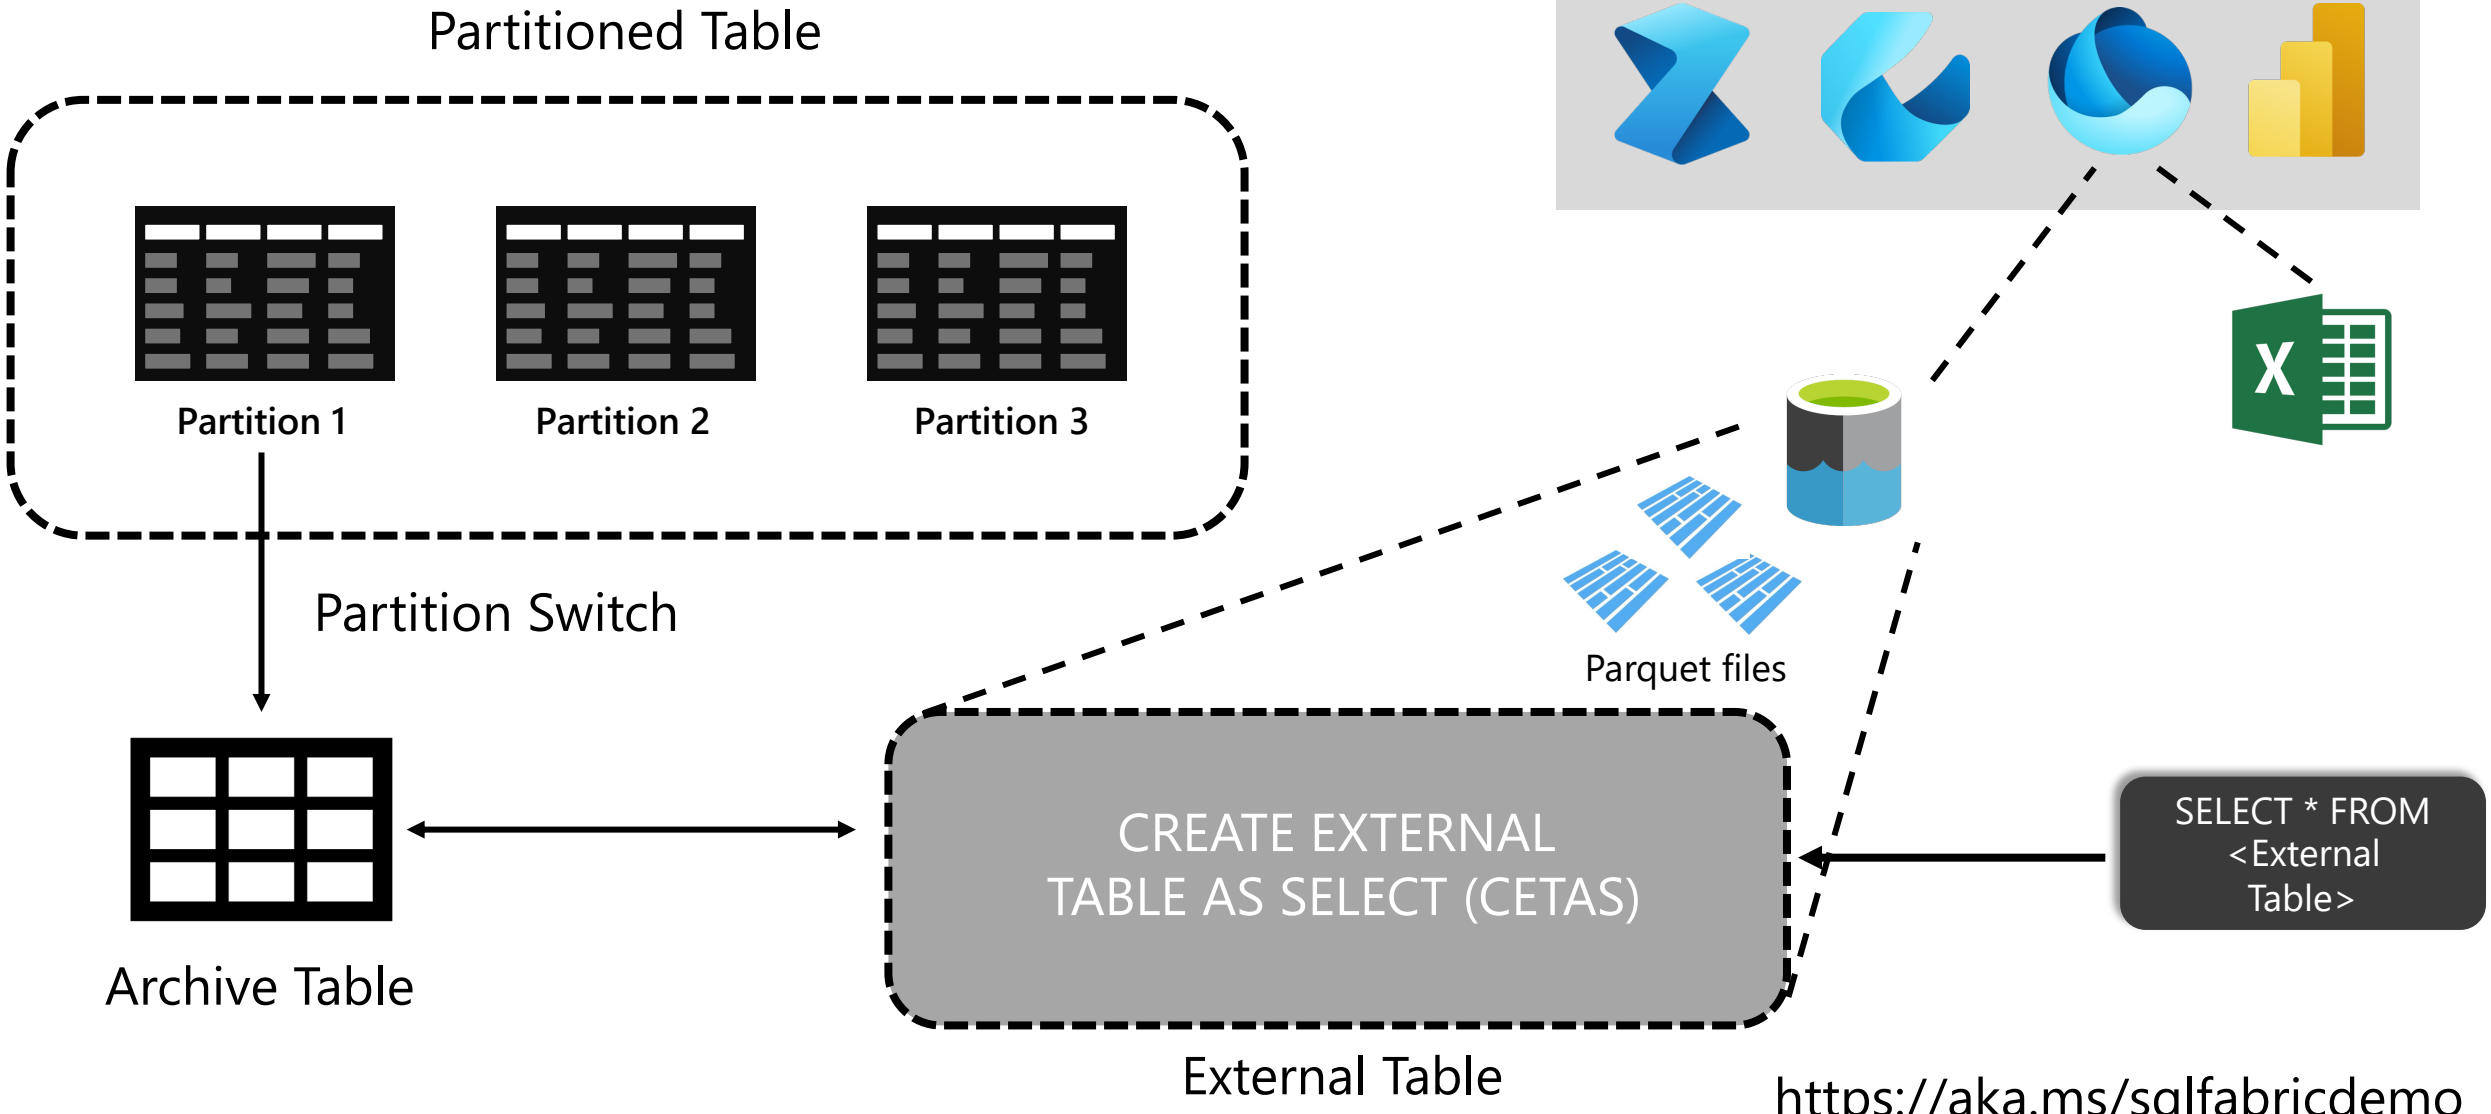


https://aka.ms/sqlfabricdemo

# Using BACKUP/RESTORE with S3 compatible object storage

```
BACKUP DATABASE <db>
TO URL = 's3:<bucket>/<file>'

RESTORE DATABASE <db>
FROM URL = 's3:<bucket>/<file>'
```

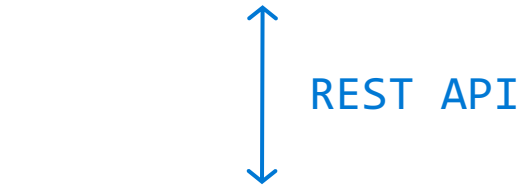



Can I migrate from
AWS RDS to SQL Server 2022?

Microsoft

# Exercises

**Exercise 5.0**
Try out OPENROWSET
**Exercise 5.1**
Notebooks with parquet and S3
**Exercise 5.2**
Notebooks with delta and S3

**Bonus: 5.3**
Notebooks with backup/restore and S3
**Advanced: 5.4**
Using minio server with S3

# Knowledge check

**SQL Server 2022 can now access external data sources through:**

**A.** Using REST API to S3 object providers

**B.** Export data to parquet files on REST API data sources from a SELECT

**C.** Access parquet files on ODBC data sources

**D.** A and B

**You can restore a database from AWS RDS to SQL Server 2022:**

**A.** True

**B.** False

# Summary

**New REST API data sources** include Azure Blob Storage, Azure Data Lake Storage Gen2, and S3 compatible object store providers

SQL Server 2022 now natively recognizes **parquet and delta table** file formats

SQL Server and **backup and restore** native backups to and from **S3** compatible object store providers


Module 6

# T-SQL

# Enhance your application with new T-SQL capabilities

Overview

Exercises

Knowledge Checks and Summary

# Language enhancements for developers

## Use your T-SQL skills to do more than ever

Challenge: I want to get new features without learning a new language

{JSON}

ISJSON
JSON_PATH_EXISTS
JSON_OBJECT
JSON_ARRAY

T-SQL

GREATEST
LEAST
STRING_SPLIT
TRIM functions
DATETRUNC
IS [NOT] DISTINCT FROM
WINDOW clause
LEFT_SHIFT
RIGHT_SHIFT
BIT_COUNT
GET_BIT
SET_BIT

DATE_BUCKET
GENERATE_SERIES
FIRST_VALUE
LAST_VALUE

Microsoft

# Exercises
# You pick and choose

**Exercise 6.0**
JSON functions

**Exercise 6.1**
Core T-SQL

**Exercise 6.2**
Time Series

**Exercise 6.3**
Bit manipulation functions

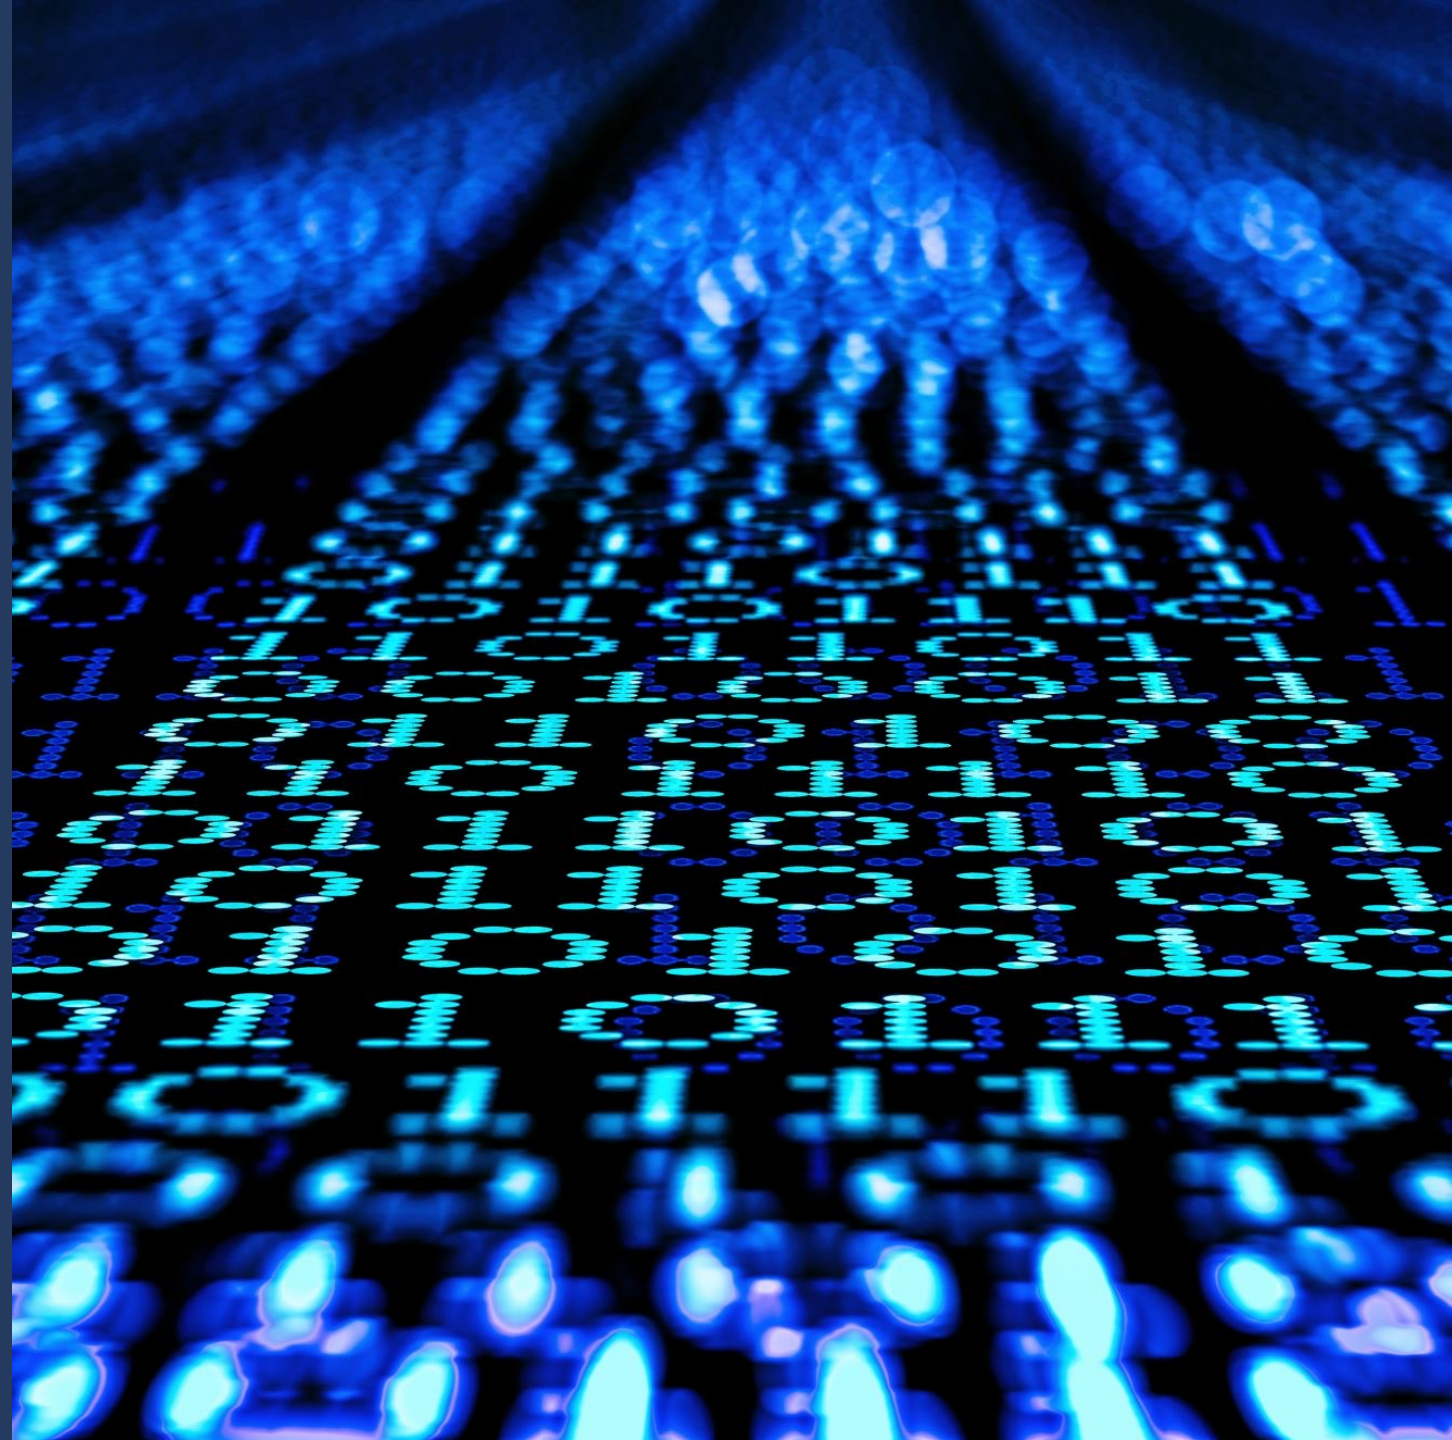

# Knowledge check

**Which T-SQL function can produce JSON formatted data?**

**A.** JSON_OBJECT

**B.** JSON_ARRAY

**C.** JSON_TEXT

**D.** A and B

**E.** All of the above

**Which important enhancement was added to the STRING_SPLIT function in SQL Server 2022?**

**A.** separator

**B.** multi-value expression

**C.** ordinal

**D.** New string types

# Summary

The **most enhancements** in T-SQL than many previous releases

Enhancements to process **JSON**

Core engine T-SQL enhancements for **ANSI compatibility** and **industry standards**

New T-SQL functions to process **time series data**

Module 7

# Summary and closing

Summary of SQL Server 2022

Resources

Group Discussion and Final Q&A

# Summary

SQL Server 2022 is **Azure-enabled** for disaster recovery, analytics, and security

**Built-in query intelligence** reduces query tuning and gains you performance with **no code changes**

SQL Server 2022 has new innovations in **security, scalability, and availability**

**Data virtualization and object storage** through standard SQL interfaces

**Extending the T-SQL language** for new developer scenarios

# Learn more

## Learn more about SQL Server 2022
aka.ms/sqlserver2022

## Download SQL Server 2022
aka.ms/getsqlserver2022

## Get our decks
aka.ms/sqlserver2022decks

## What's new for SQL Server 2022
aka.ms/sqlserver2022docs

## Read the book


aka.ms/sql2022book

## Try out the workshop
aka.ms/sql2022workshop

## Try our demos
aka.ms/sqlserver2022demos

## Watch Microsoft Mechanics
aka.ms/sqlserver2022mechanicsupdate

## Don't miss us on Data Exposed


aka.ms/sqlserver2022series

## Read the SQL Server 2022 Blog Series
aka.ms/sqlserver2022blogs

## Take the Microsoft Learning Path
aka.ms/learnsqlserver2022

# Thank you!

# Applications, Tools, and T-SQL







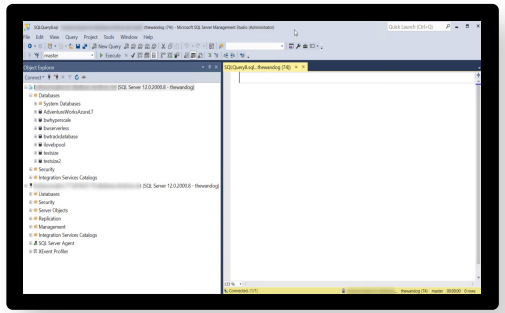

SQL Server Management Studio (SSMS)

Windows Only
New 19.X version

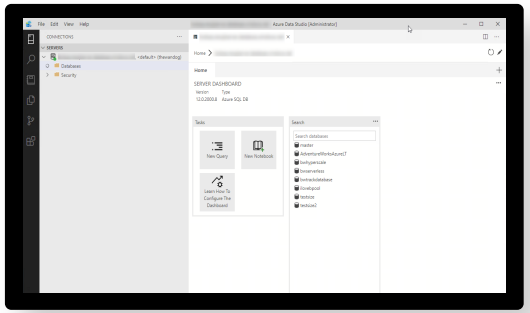

Azure Data Studio



Powershell

sqlcmd

bcp

Tabular Data Stream (TDS) 8.0 and T-SQL

# SSRS and SSAS in SQL Server 2022

## SSAS

- Power BI support for SQL Server 2022 Analysis Services tabular models as a DirectQuery connection for composite models
- Improved MDX query performance
- Improved resource governance
- Query interleaving now includes short query bias with fast cancellation

## SSRS

- Enhanced Windows Narrator support for the new Windows OS (Operating Systems) and Windows Server
- Security enhancements
- Browser performance improvements with Angular
- Accessibility bug fixes
- Support for SQL Server 2022 instances report server catalog
- Reliability updates

# What is a hybrid data platform?

A product or service offered both on-premises and the cloud with consistency

A product on-premises connected to the cloud to enhance data capabilities

# Notes about Synapse Link

Catalog views, DMVs, and Extended Events for monitoring and debugging

Datatypes not supported include image, text, xml, timestamp, UDT, geometry, geography

Can't be used with database that use these: temporal, memory optimized, CDC, columnstore, graph

Use scheduled mode to spend less on ingestion

Create separate linked connections for same databases to separate ingestion

Only committed transactions sent to Landing Zone (LZ)

"Think" replication. Transaction commit not affected but log can't be truncated until committed in LZ

Landing Zone is cleaned up periodically for files no longer needed

aka.ms/synapselinksql

Microsoft