



Microsoft

SQL Server 2019

SQL Server 2019: Not your Grandpa's SQL Server

Published on September 4, 2019

... when you deploy SQL Server 2019, you are also gaining new capabilities from the amazing work of SQL Server 2012, 2014, 2016, and 2017

iterating and releasing preview builds of the release so you can test, develop, study, and now even deploy in production.

When I first started talking about SQL Server 2019, it was all about the features and how to migrate from legacy SQL Server releases or even other products such as Oracle or PostgreSQL. I realized that in my conversations with customers there were far greater

Resumable Online Index creation has eased the whole process of heavy Index creation, manageable. Otherwise, in older editions, with any pause to the Index creation process, either due to memory constraints or otherwise, the process had to be re-started, making it cumbersome.

... an extension to the existing features of the older SQL server versions, no new installation or environment setup needs to be done

There are many great features of the older SQL server versions, no new installation or environment setup needs to be done. Many feedbacks to the new installation and update include the ease of set up and easy usage of the tool. Overall, it's a reliable and stable product.



Microsoft
SQL Server
2019



Who am I

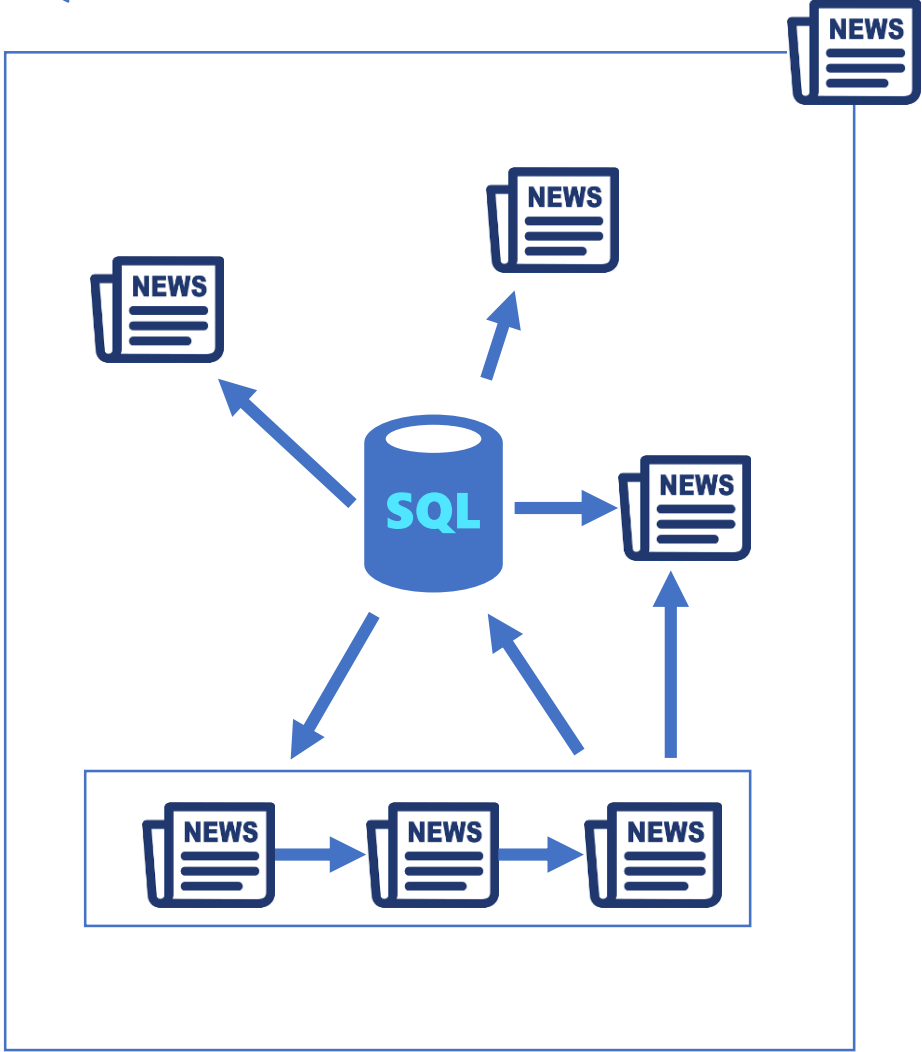
- Steen Dybboe
- Database Developer
- Pragmatic BI
- Plays Golf
- No dogs



SQL Server Nyheder

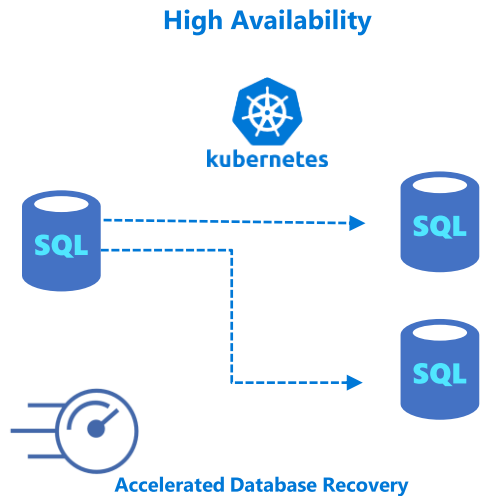
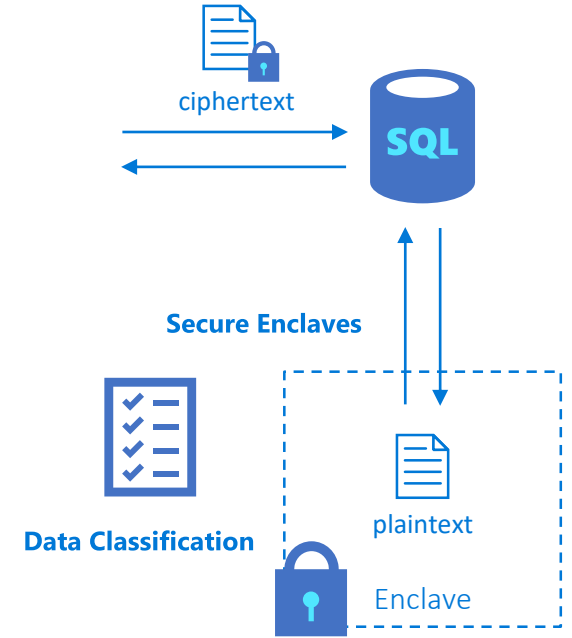
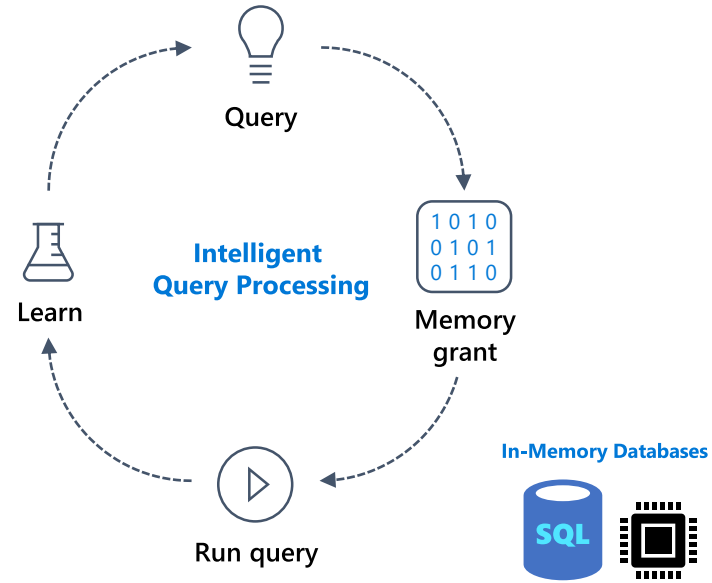
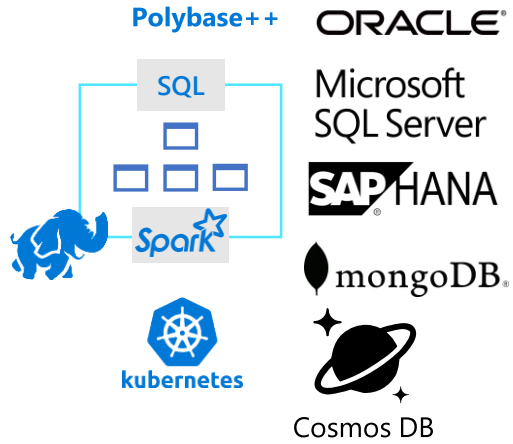


SQL Server 2019

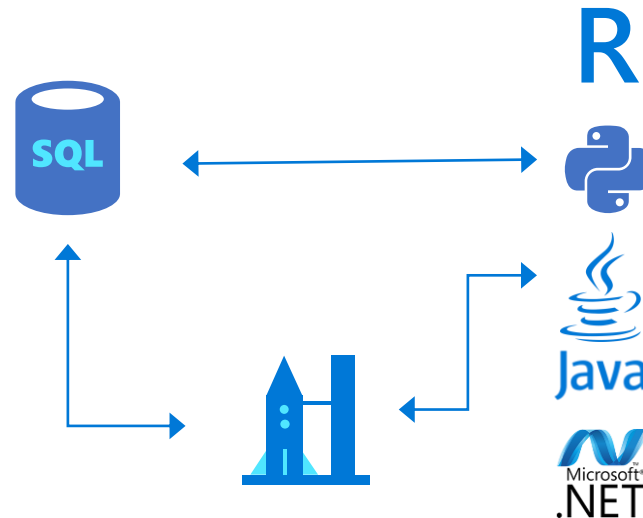


SQL Server 2019

Solving Modern Data Challenges



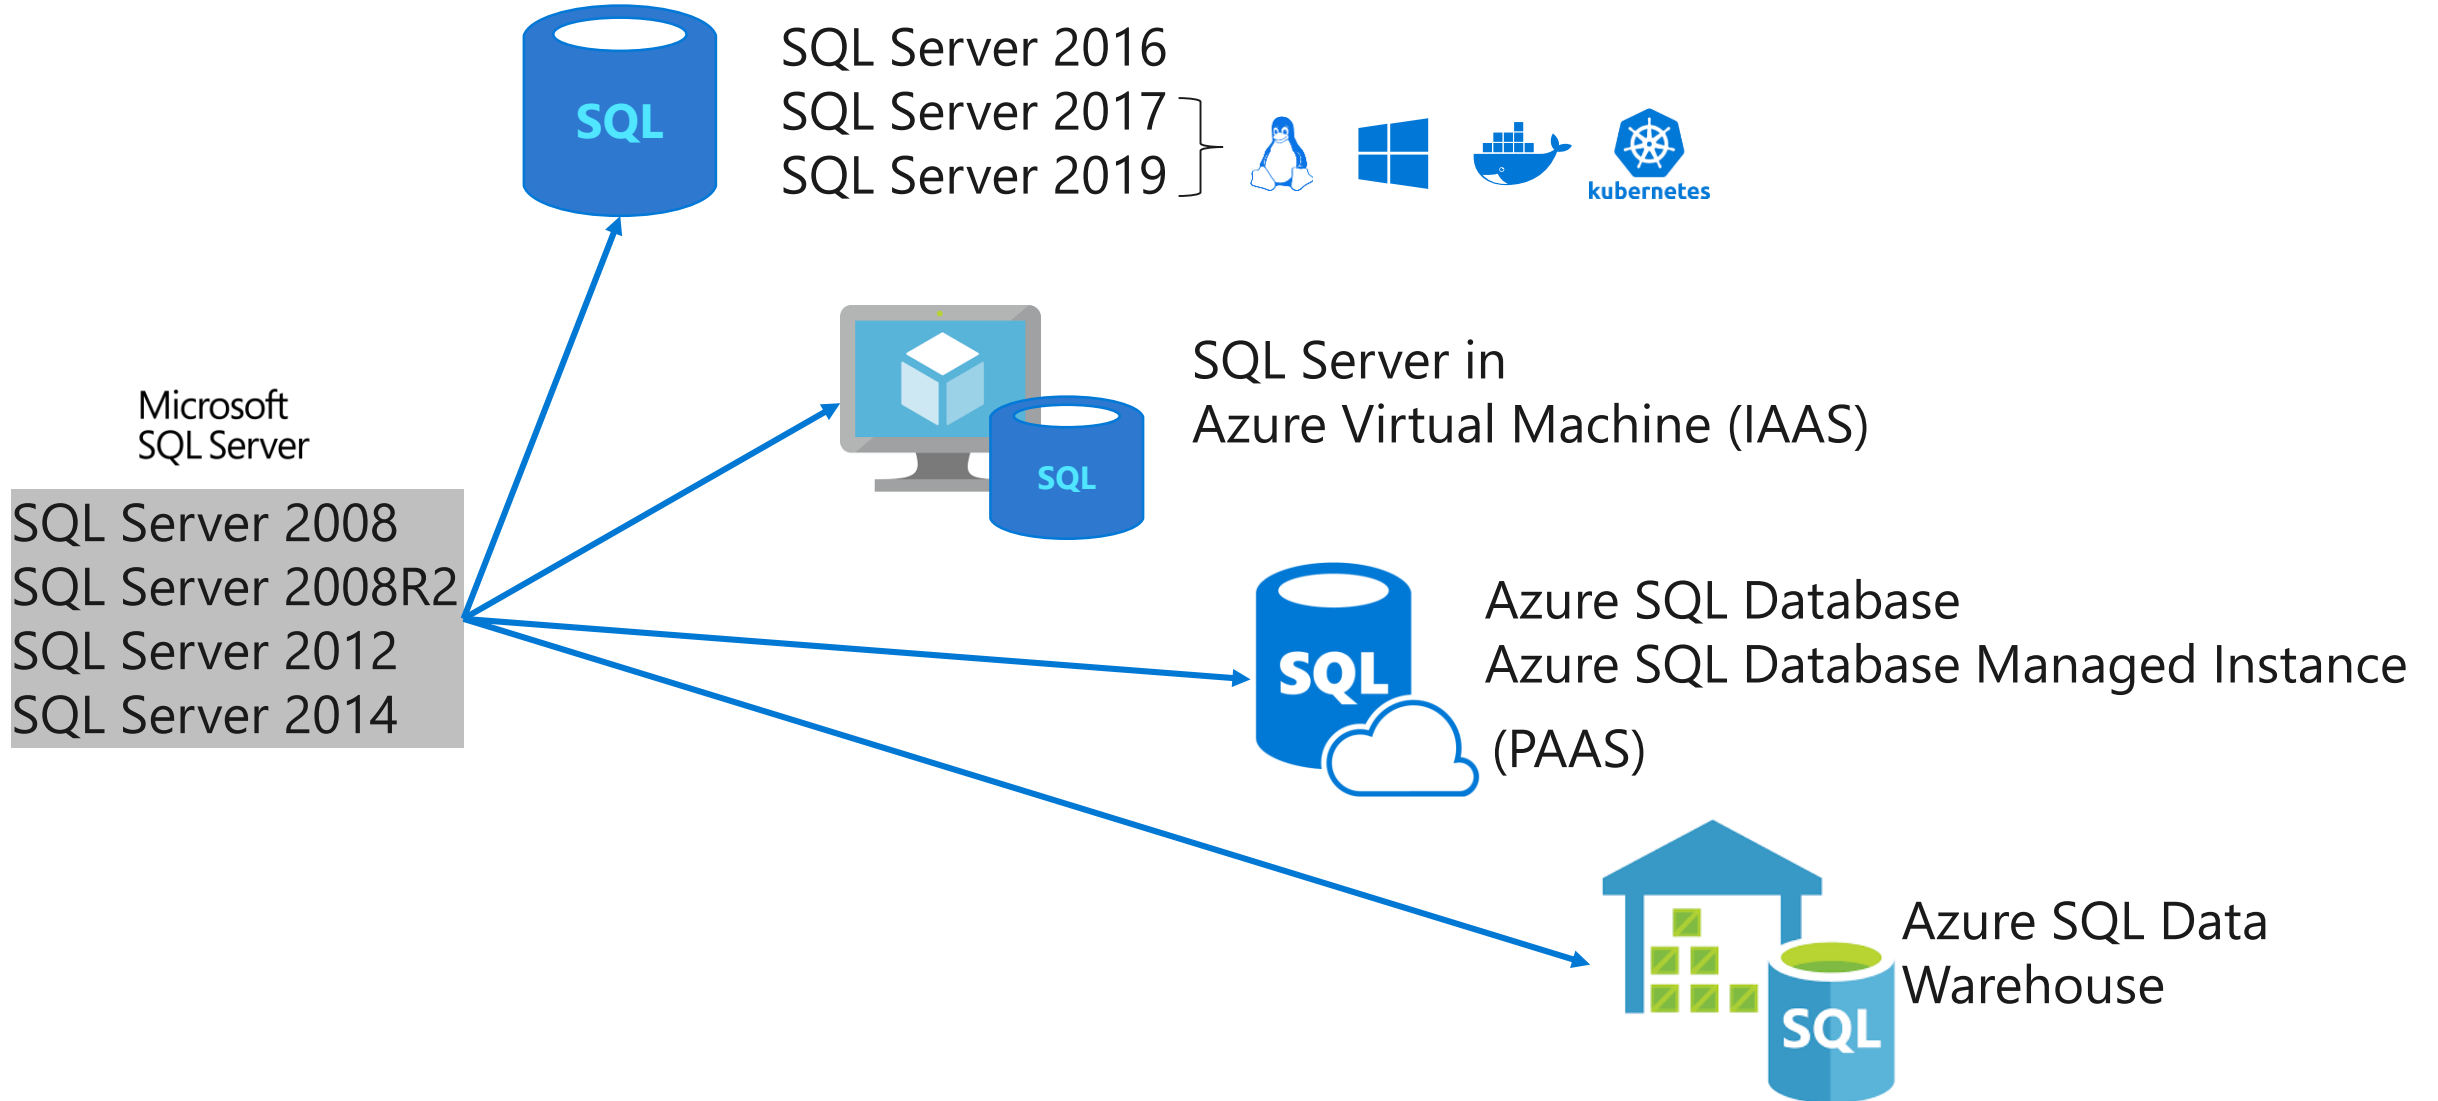
Built-in Machine Learning and Extensibility



Modern Platforms with Compatibility



Azure Data Modernization Choices



Agenda

SQL Server 2019 – Release Candidate 1

- (1) Big Data Clusters
- (2) Data virtualization
- (3) Platform and languages
- (4) Availability
- (5) Performance
- (6) Security
- (7) The Rest

Analysis Services 2019

Reporting Services 2019

Integration Services 2019

Big Data Clusters



Big Data is Just Data



What the heck
is Electronic Mail?

Electronic Mail is a term that's been bandied about data processing circles for years.

Simply put, it means high-speed information transportation.

One of the most advanced methods is terminals talking to one another.

Your mailbox is the terminal on your desk. Punch a key and today's correspondence and messages are displayed instantly.

Need to notify people immediately of a fast-breaking development? Have your messages delivered to their terminal mailboxes electronically, across the hall or around the world.

Electronic Mail is document distribution that's more timely, accurate and flexible than traditional methods.

There's no mountain of paperwork.

Administrative personnel are more effective.

Managers have access to more up-to-date information. Decision-making is easier.

Tomorrow's automated office will clearly include Electronic Mail. But like the rest of the Office of the Future, it's available at Honeywell today.

For more information call Mr. Laurie Reeves at (800) 925-3992/3 (within the 617 area, call 552-9048). Or write him at Honeywell Office Automation Systems, Three Newton Executive Park Drive, Newton Lower Falls, Massachusetts 02162.

Honeywell

SQL Server 2019 Big Data Cluster

SQL Server Big Data Clusters allow you to deploy scalable clusters of SQL Server, Spark, and HDFS containers running on Kubernetes.

These components are running side by side to enable you to read, write, and process big data from Transact-SQL or Spark, allowing you to easily combine and analyze your high-value relational data with high-volume big data.


Technology	Description
<i>Linux</i>	Operating system used in Containers and Container Orchestration
<i>Docker</i>	Encapsulation level for the SQL Server big data cluster architecture
<i>Kubernetes</i>	Management, control plane for Containers
Microsoft Azure	Cloud environment for services
Azure Kubernetes Service (AKS)	Kubernetes as a Service
<i>Apache HDFS</i>	Scale-out storage subsystem
<i>Apache Knox</i>	The Knox Gateway provides a single access point for all REST interactions, used for security
<i>Apache Livy</i>	Job submission system for Apache Spark
<i>Apache Spark</i>	In-memory large-scale, scale-out data processing architecture used by SQL Server
<i>Python, R, Java, SparkML</i>	<i>ML/AI programming languages used for Machine Learning and AI Model creation</i>
Azure Data Studio	Tooling for SQL Server, HDFS, Kubernetes cluster management, T-SQL, R, Python, and SparkML languages
SQL Server Machine Learning Services	R, Python and Java extensions for SQL Server
Microsoft Data Science Process (TDSP)	Project, Development, Control and Management framework
<i>Monitoring and Management</i>	Dashboards, logs, API's and other constructs to manage and monitor the solution
<i>Security</i>	RBAC, Keys, Secrets, VNETs and Compliance for the solution



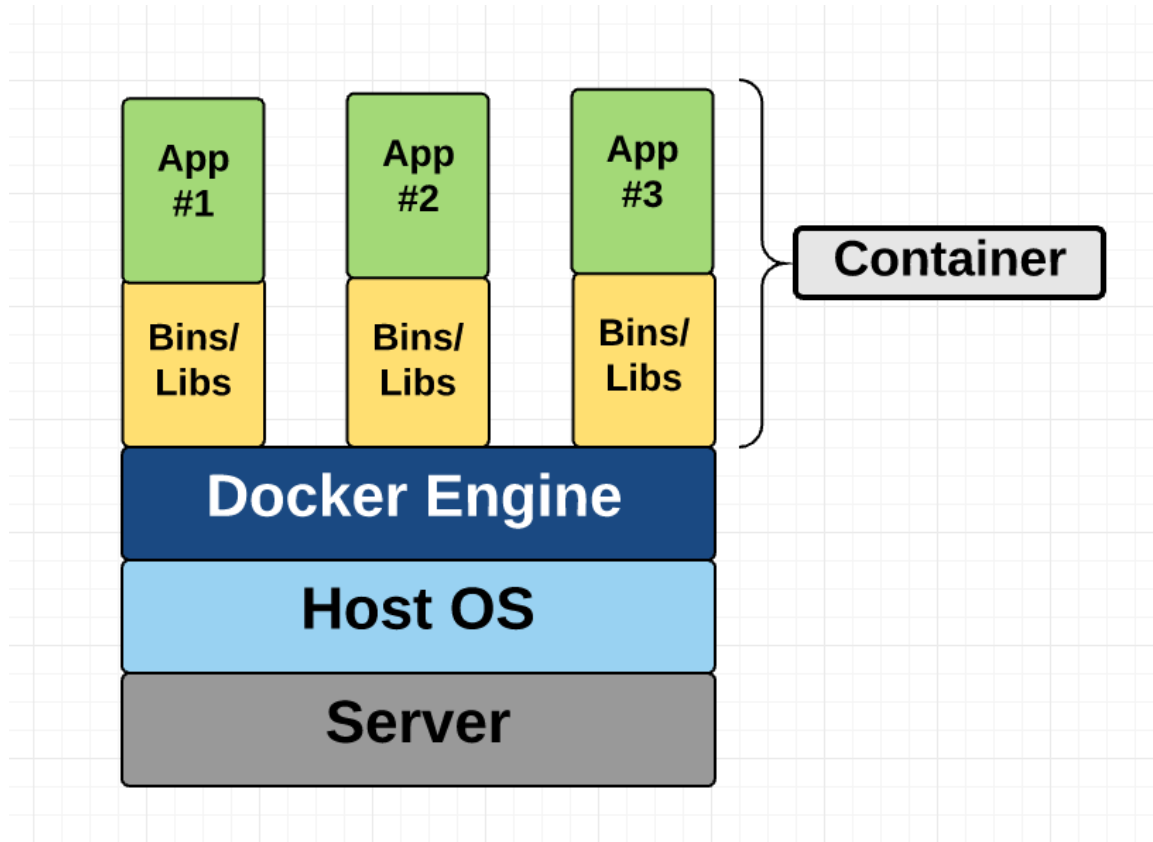
- Docker Container
- Kubernetes
- SQL Server on Linux

Virtual Machines

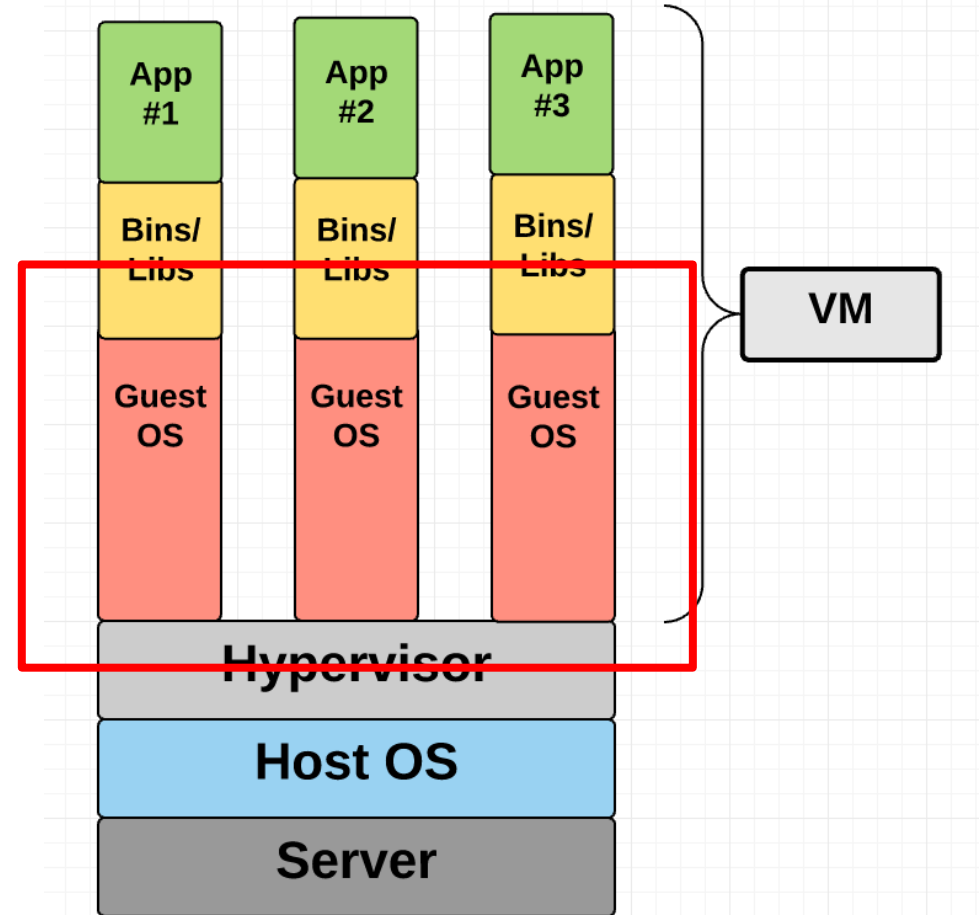
- Emulation of a real computer that executes programs like a real computer.
- VMs run on top of a physical machine using a “*hypervisor*”.
 - Piece of software, firmware, or hardware that VMs run on top of.
- Can run on either a **hosted hypervisor** or a **bare-metal hypervisor**



Container

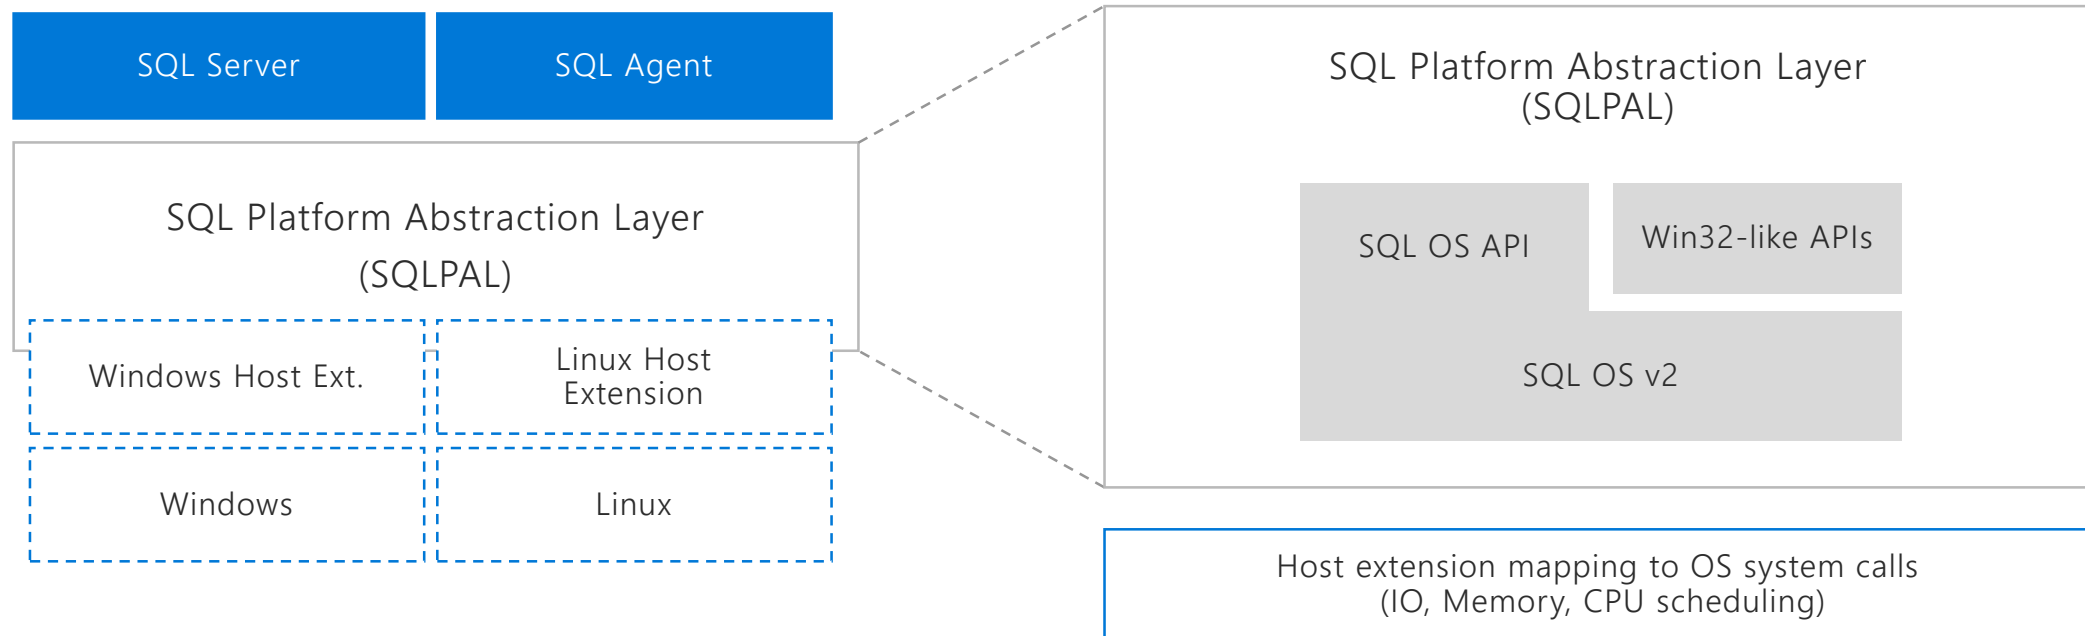


Virtual Machine



SQL SERVER ON LINUX:

SAME AS SQL SERVER ON WINDOWS SQL PLATFORM ABSTRACTION LAYER



SQL Server 2019 and Linux and Containers


New Features for Linux

- Replication
- Change Data Capture
- Distributed transactions
- Machine Learning
- Polybase
- Tempdb files auto-config

Containers

- Microsoft Container Registry
- Red Hat Images
- Non-root Containers
- SQL Server on Windows Containers (Private Preview)

Windows



Containers and Kubernetes



Linux



What's in SQL Server on Linux

		Windows	Linux
Editions	Developer, Express, Web, Standard, Enterprise	●	●
Services	Database Engine, Integration Services, SQL Server Agent Analysis Services, Reporting Services, MDS, DQS	● ●	● ●
Mission critical performance and HADR	Maximum number of cores	Unlimited	Unlimited
	Maximum memory utilized per instance	24 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)	●	●
	SQL Server Replication	●	● NEW
	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, Active Directory Authentication) Advanced security (Transparent Data Encryption)	● ●	● ●
Data warehousing	PolyBase	●	● Planned
	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)	●	●
	Distributed Transactions	●	● NEW
	Machine Learning Services	●	● NEW

Why SQL Server in Containers



Portable

Run anywhere Docker is supported



Lightweight

Reduced disk, CPU, and memory footprint



Consistent

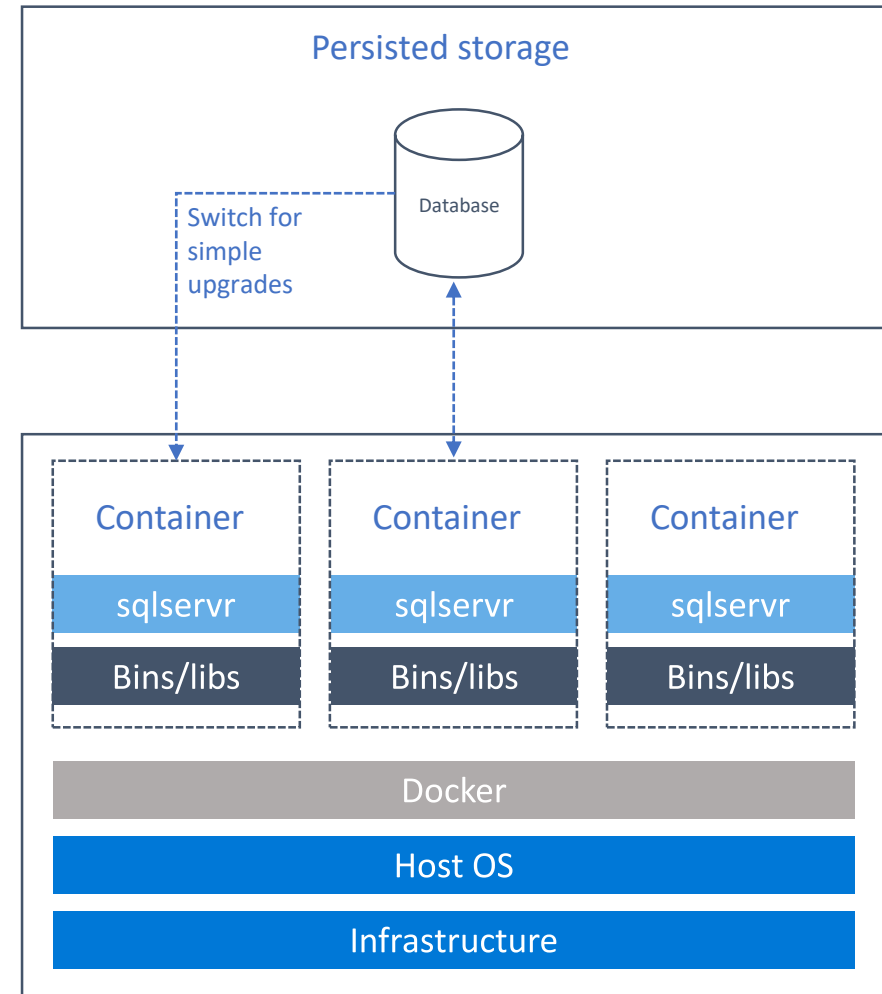
Consistent image of SQL Server, scripts, and tools



Efficient

Faster deployment, reduced patching, and less downtime

Container configuration



SQL Server Images

Public Registries



Docker Hub (<https://hub.docker.com>)

Red Hat Container Catalog (<https://access.redhat.com/containers>)

The Microsoft Container Registry (<mcr.microsoft.com>)

SQL Server 2017

Ubuntu based Images

RTM plus each CU and GDR has separate image

Includes mssql-server and mssql-tools packages

SQL Server 2019 Preview images

Ubuntu and RHEL based images

Includes mssql-server and mssql-tools

SQL Server
Engine pre-
installed

SQL Server
tools pre-
installed

Database and
application
compatibility

Running a SQL Server Container

sudo docker run

```
-e 'ACCEPT_EULA=Y' -e 'MSSQL_SA_PASSWORD=Sql2017isfast'  
-p 1401:1433  
-v sqlvolume:/var/opt/mssql  
--hostname sql1  
--name sql1  
-d  
mcr.microsoft.com/mssql/server:2017-CU8
```

environment variables
(Ex. enable SQLAgent)

Your app connects
to port 1401

Map default directory to host
volume for persisted storage

@SERVERNAME


Use with docker
client

Run process in background.
Remove this to *debug*

SQL Server 2017
CU8 image

Welcome to the new world of patch and upgrade

volume



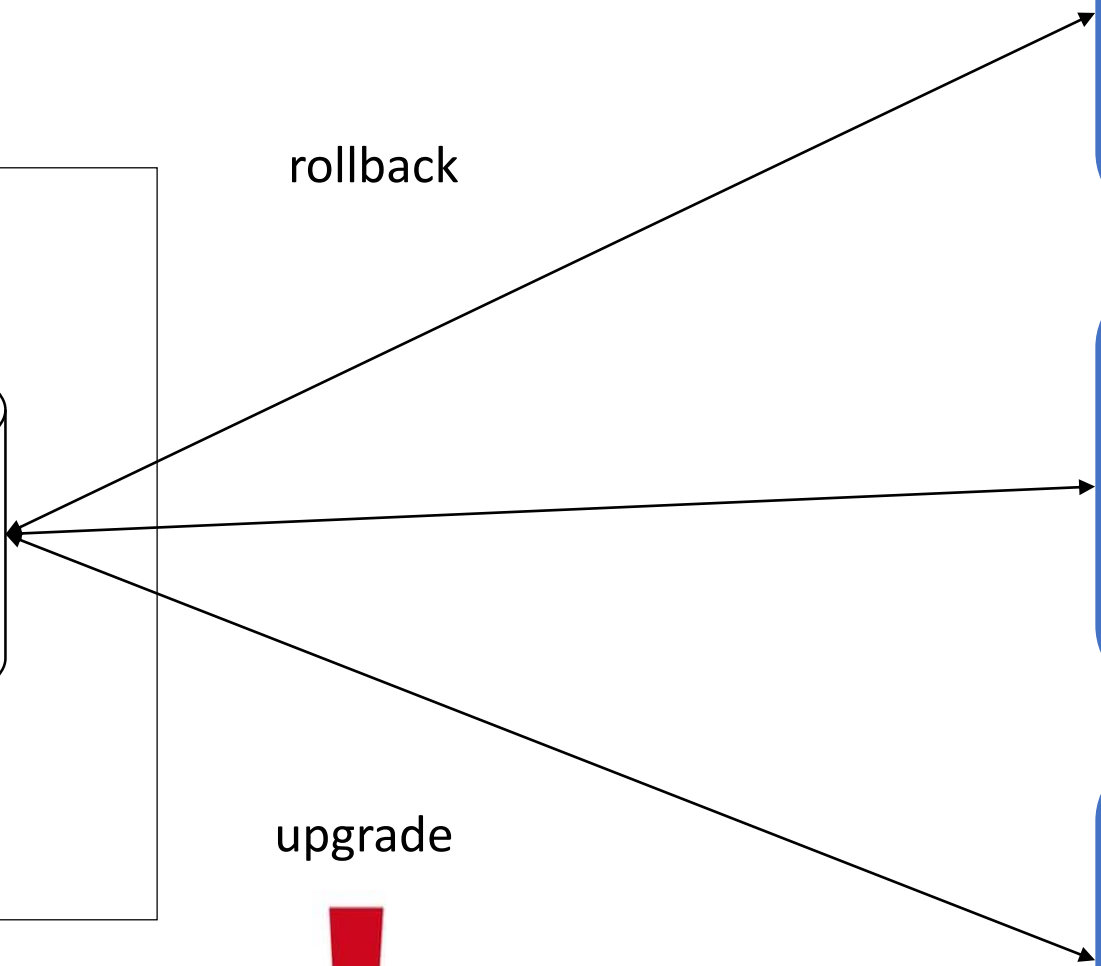
rollback

SQL Server 2017 CU8
Container

SQL Server 2017 Latest
Container

upgrade

SQL Server 2019



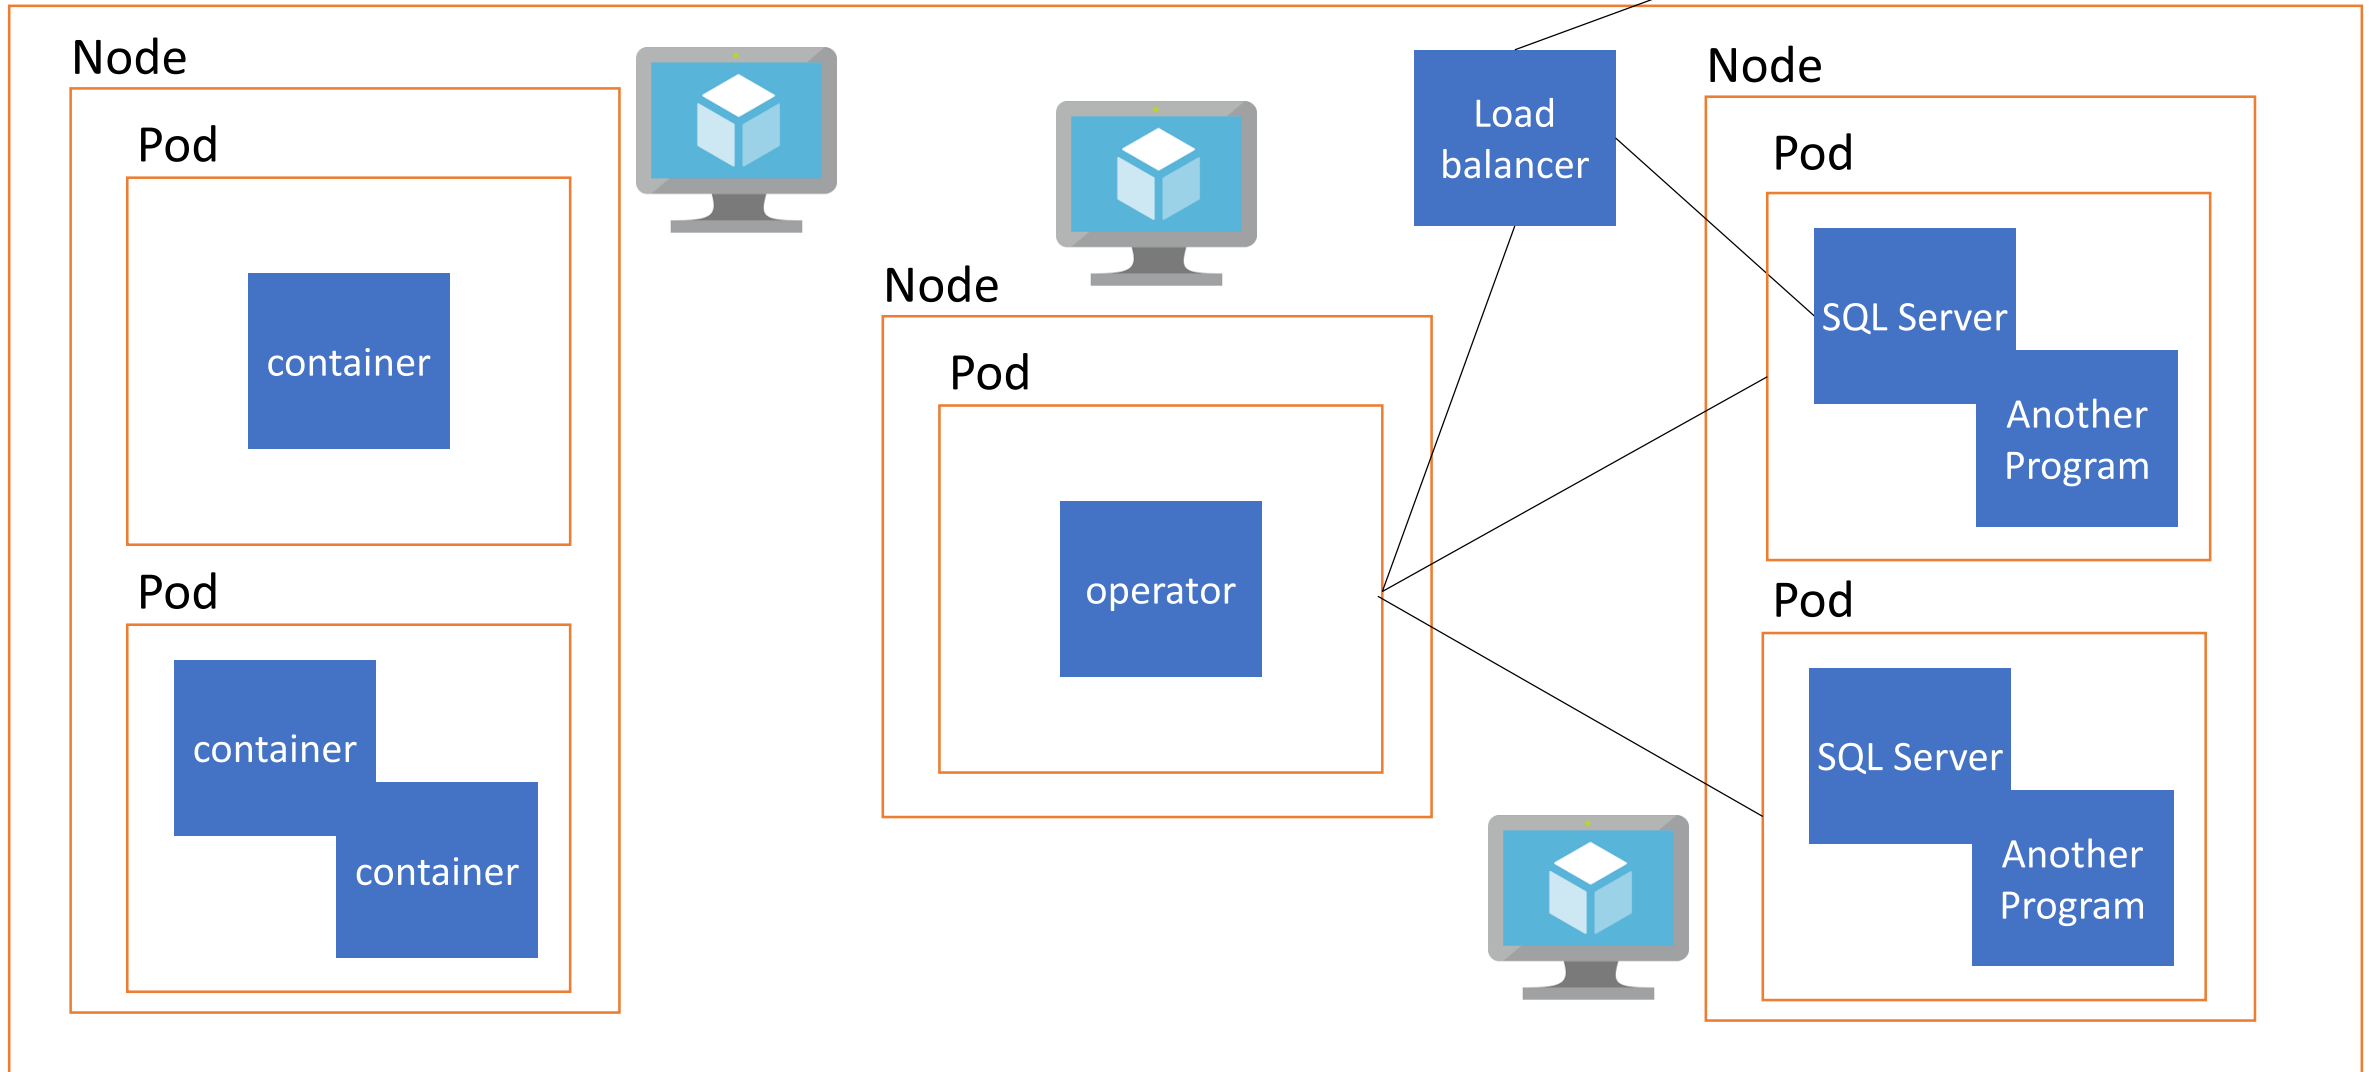


kubernetes

Kubernetes (k8s) in a nutshell



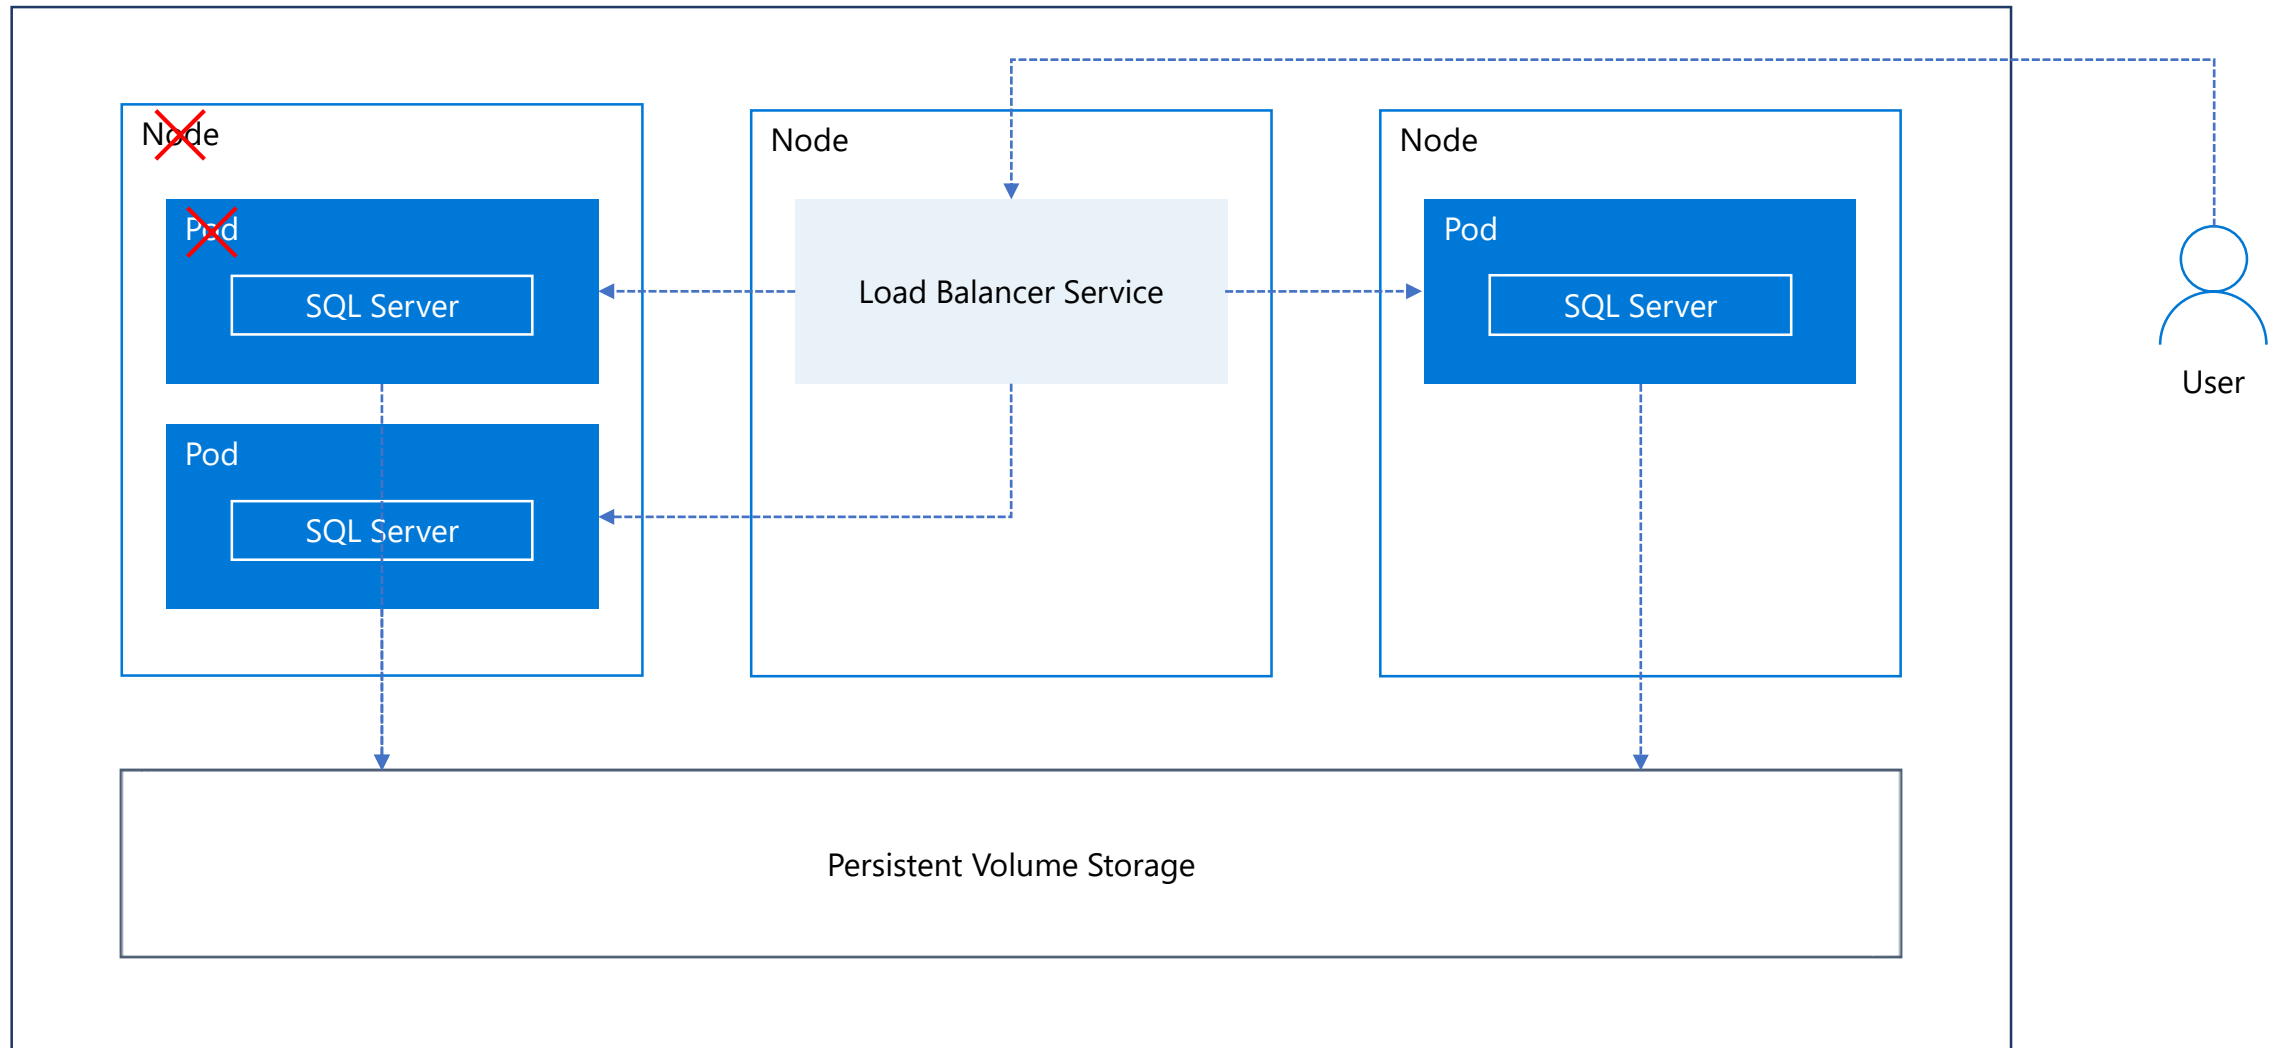
Cluster



Persistent Volume Storage

SQL Server - Shared storage HA in Kubernetes

Built-in HADR orchestration with no clustering required






What *IS*
Big Data Cluster

SQL Server 2019 – Intelligence over all data




Integration of data

Unified access to all your data with performance



Managing all data

Easily and securely manage data (big and small)

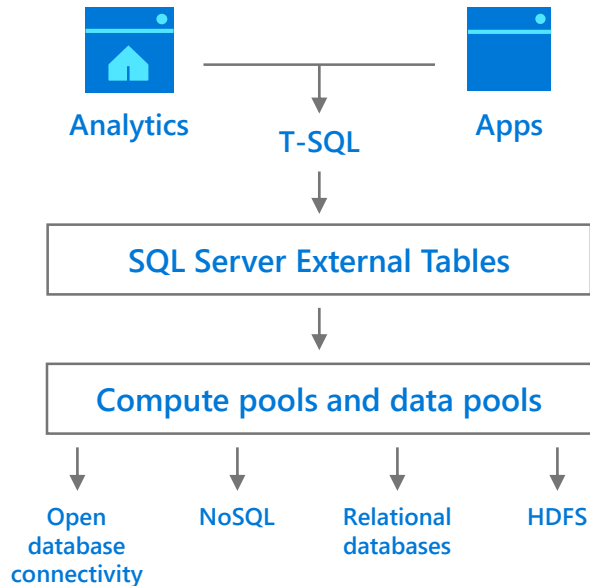


Analyzing all data

Build intelligent apps and AI with all your data

SQL Server 2019 Big Data Cluster Scenarios

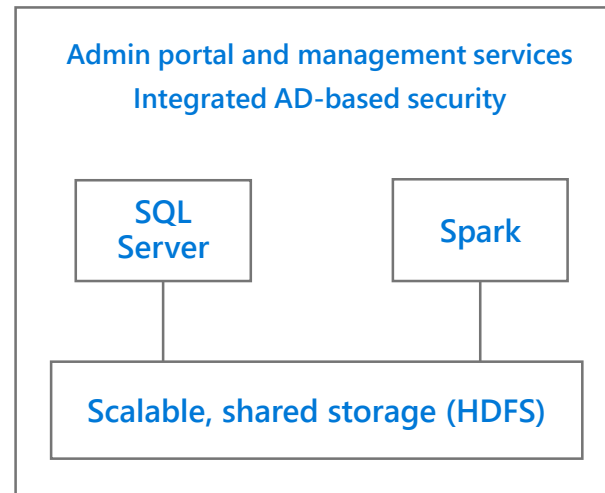
Data virtualization



Combine data from many sources without moving or replicating it

Scale out compute and caching to boost performance

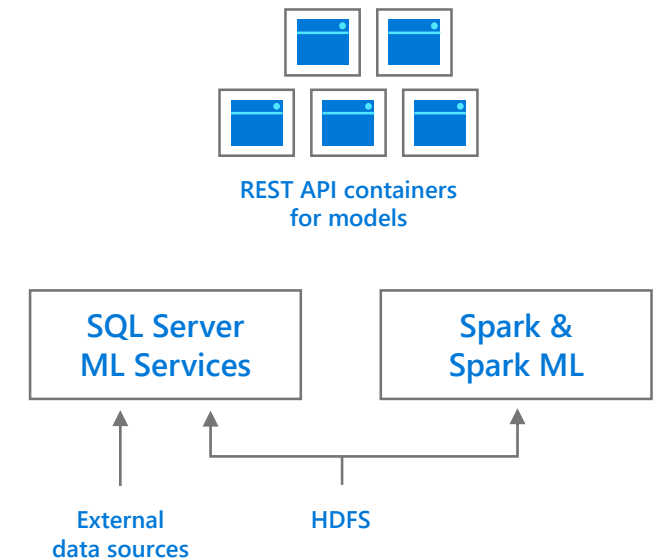
Managed SQL Server, Spark, and data lake



Store high volume data in a data lake and access it easily using either SQL or Spark

Management services, admin portal, and integrated security make it all easy to manage

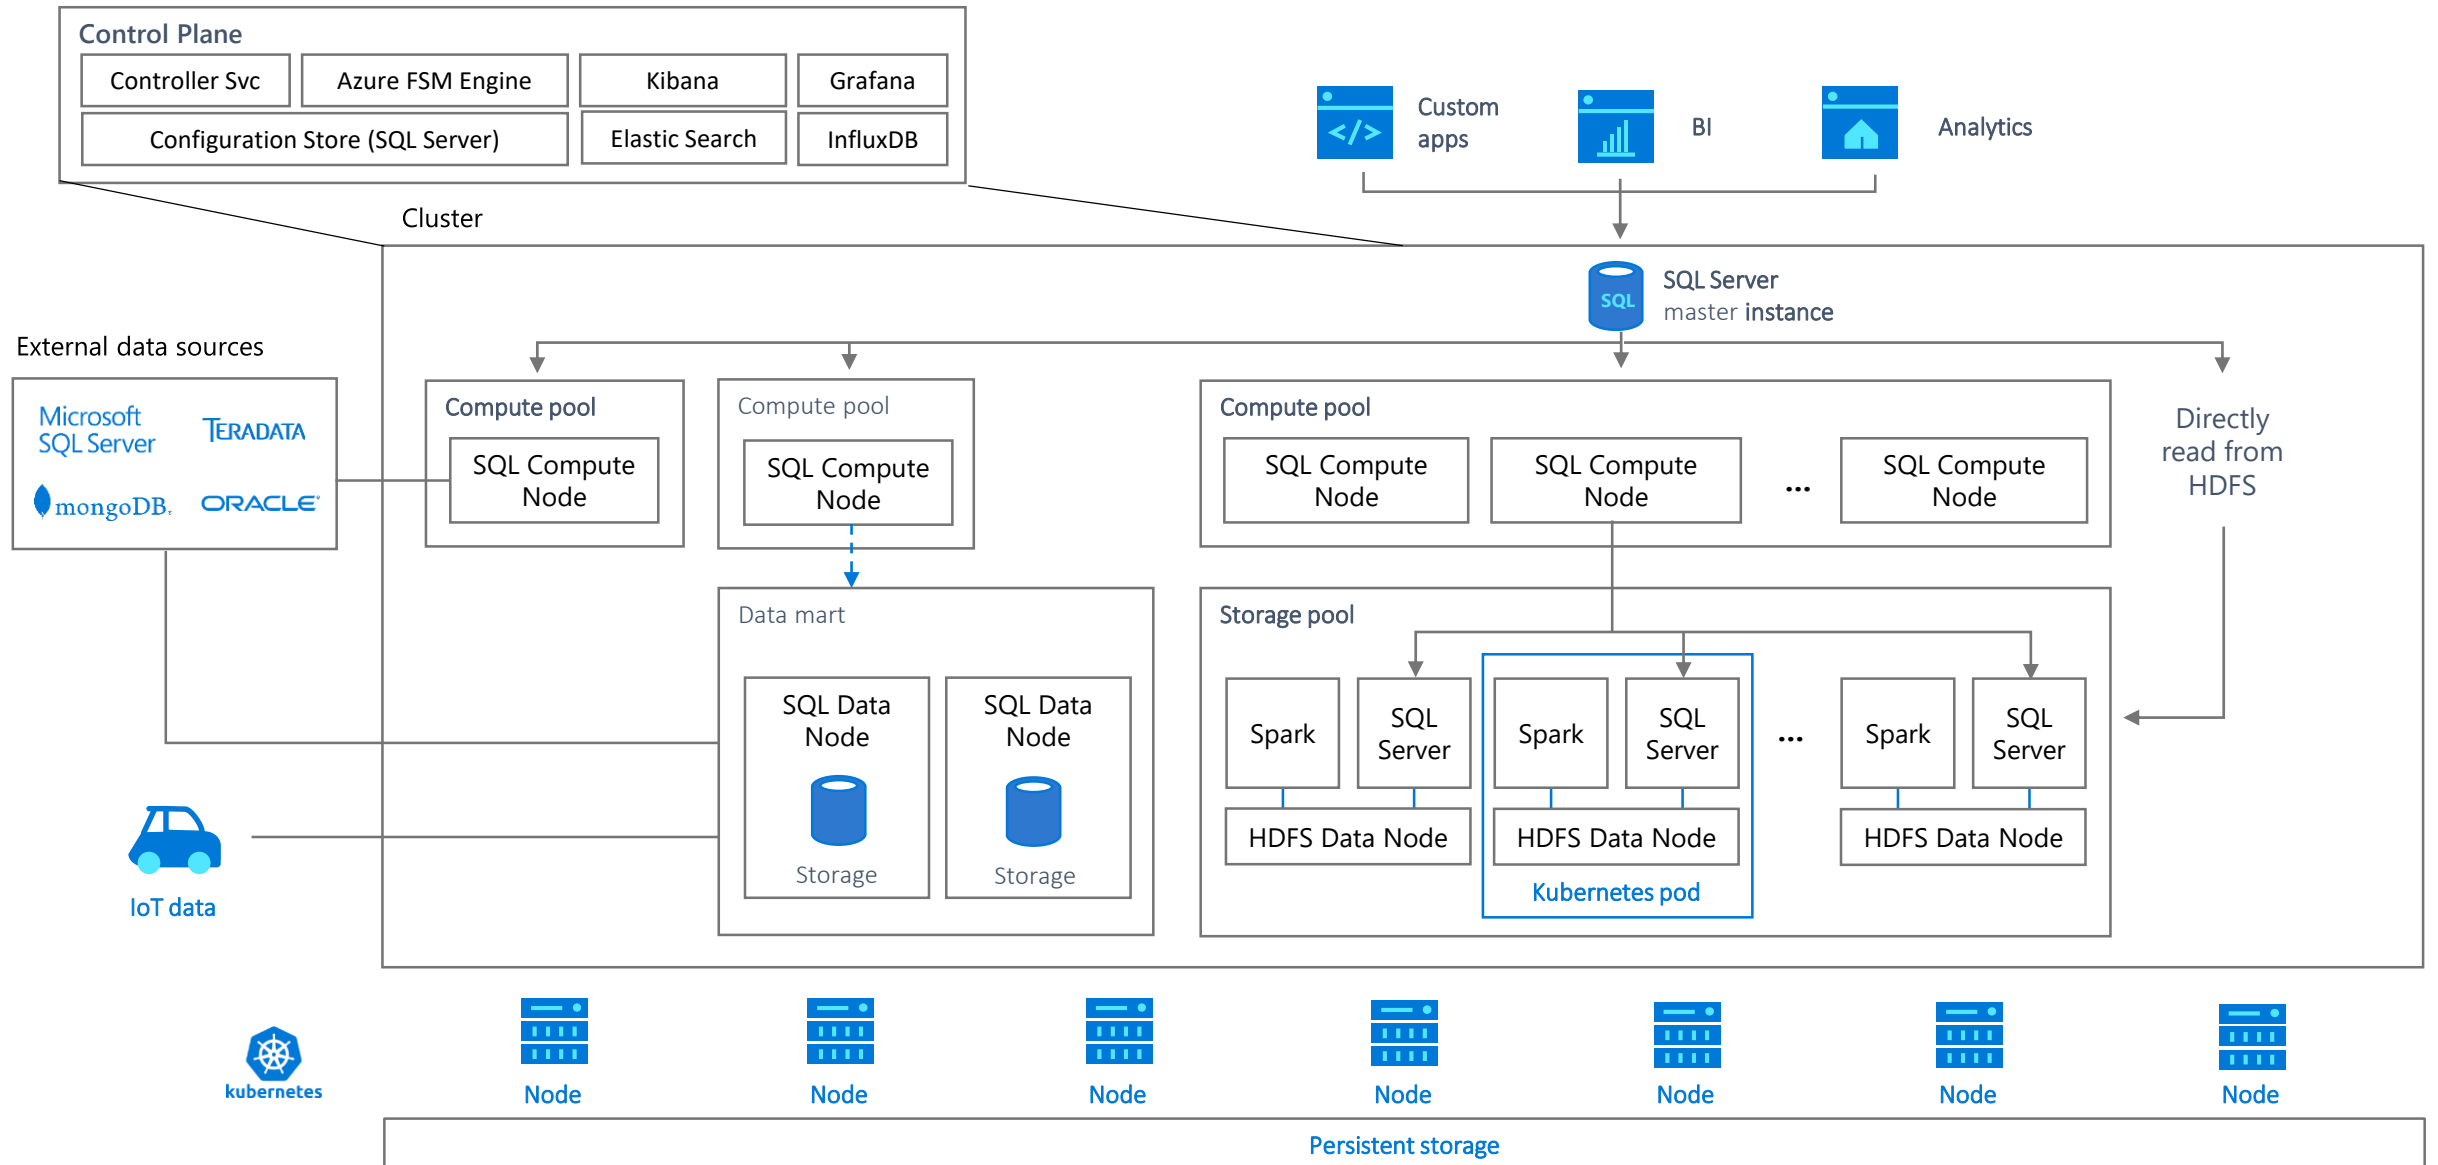
Complete AI platform




Easily feed integrated data from many sources to your model training

Ingest and prep data and then train, store, and operationalize your models all in one system

SQL Server Big Data Cluster Architecture



Data
virtualization



Data
virtualization
with
PolyBase

Data virtualization allows you to have a single querypoint where you run your SQL code

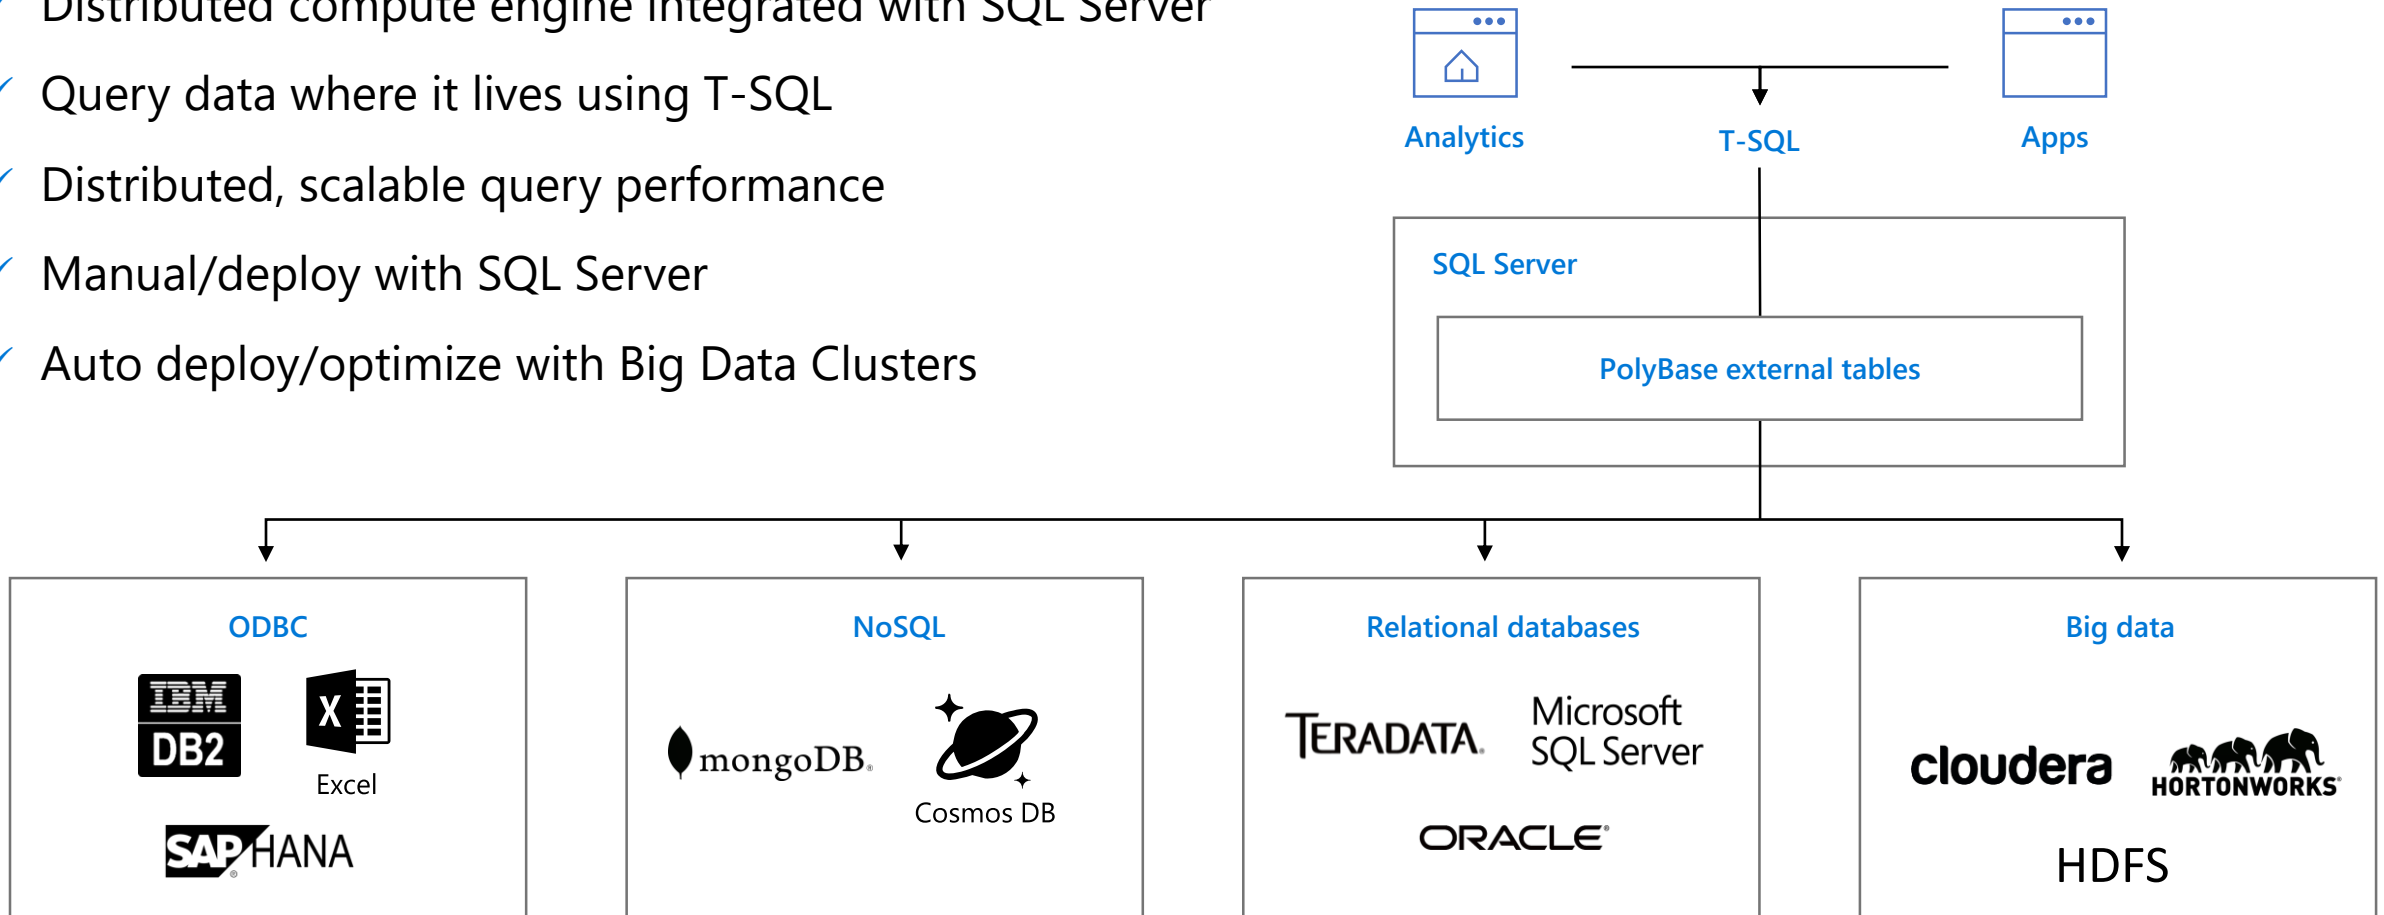
No data movement, just a semantic layer to abstract the complexity of your underlying estate.


You can query data stored in Oracle, Teradata, MongoDB, HDFS or any other data sources without moving or replicating the data

SQL Server Polybase

“It’s all about Data Virtualization”

- ✓ Distributed compute engine integrated with SQL Server
- ✓ Query data where it lives using T-SQL
- ✓ Distributed, scalable query performance
- ✓ Manual/deploy with SQL Server
- ✓ Auto deploy/optimize with Big Data Clusters





Polybase
as a
Semantic
Layer

**External
table**

Use EXTERNAL table to map common names familiar to your SQL Server

Exclude

Exclude fields or columns from data source that are not needed (Planned)

Views

Use SQL Server views to abstract joins and data source access

UNION

Use UNION to join similar data from data sources and local SQL Server

Polybase and Other Connectors (2019)

Built-in Connectors (use ODBC)

sqlserver

oracle

teradata

mongodb

No client software install required

SQL Server, Azure SQL Database, Azure SQL Data Warehouse

MongoDB or CosmosDB (using MongoDB API)

ODBC Connectors

odbc

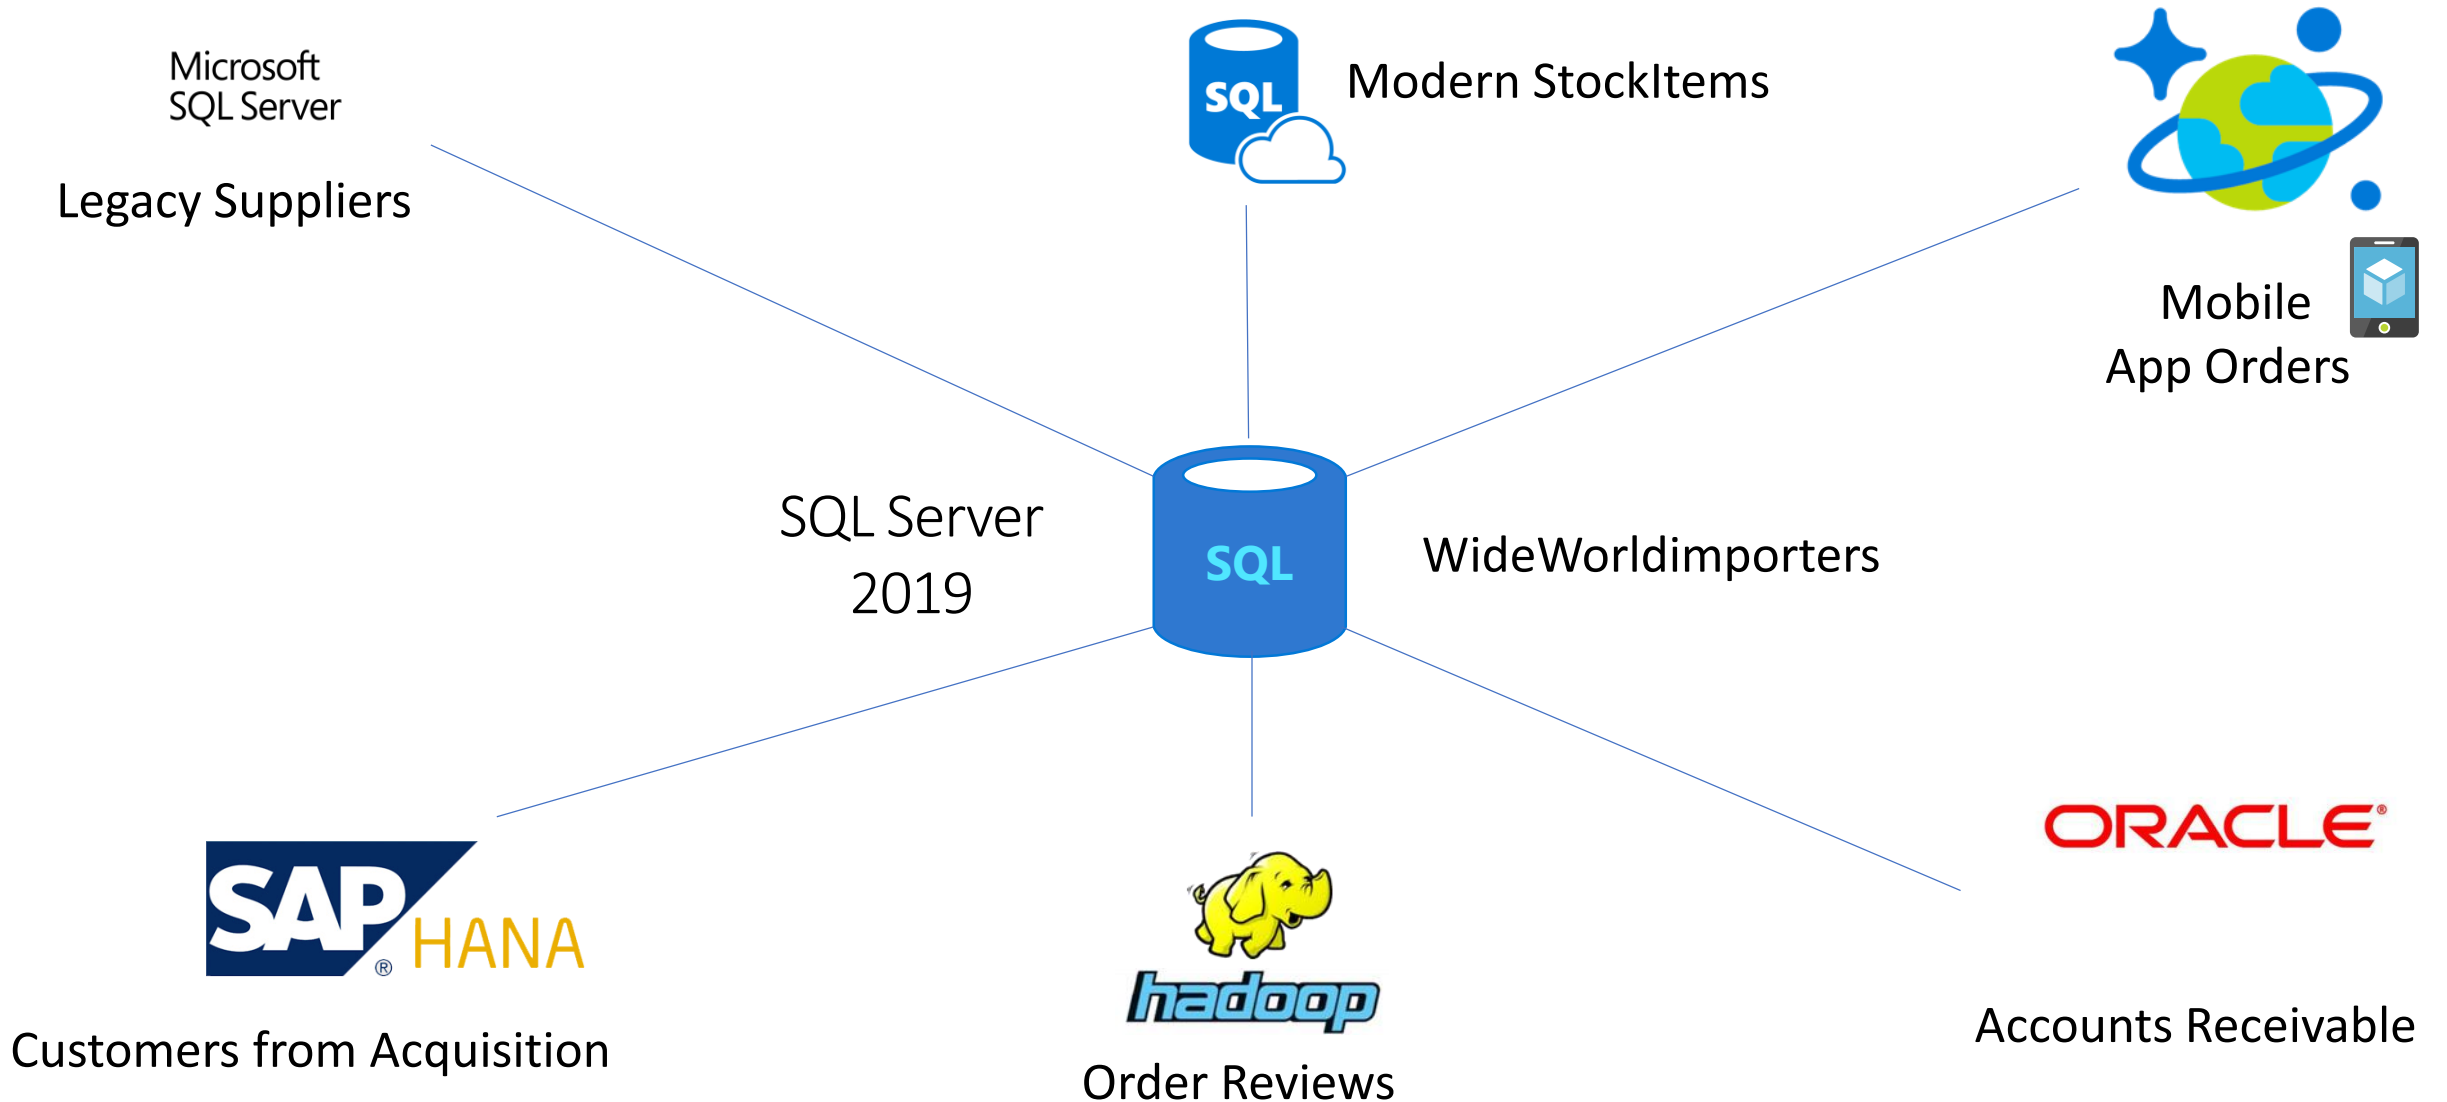
You install the driver

64bit ODBC 3.0+ compliant

Ex. SAP HANA (HDBCODBC Driver)

`\binn\Polybase\ODBC Drivers`

SQL Server 2019: Data Virtualization



platform and
language



Platform and languages



SQL Server Java extension



SQL Graph enhancements



UTF-8 support



Machine Learning Services enhancements

Java Support

SQL Server



Java

Microsoft Extensibility SDK for
Java.

[Azul Zulu Enterprise Distribution](#)

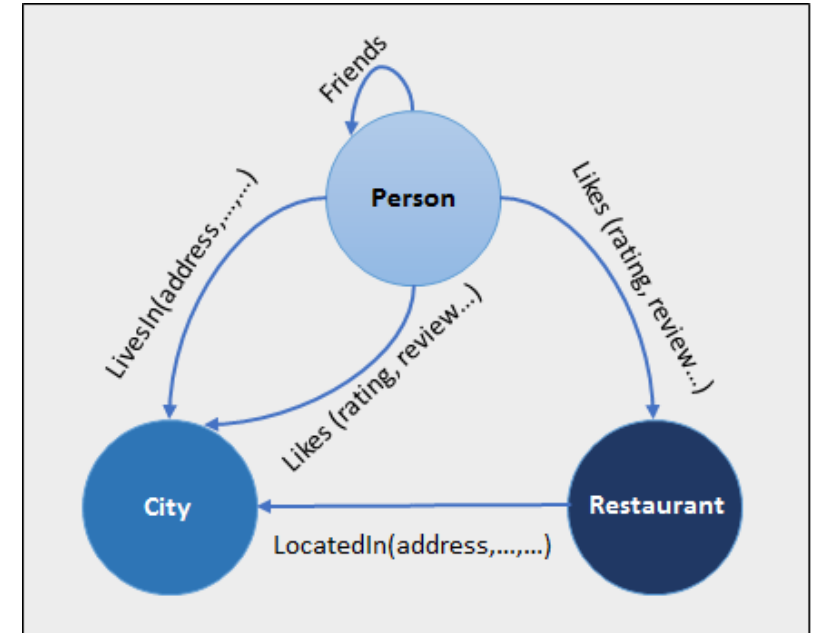
Zulu Open JRE

Full free support

SQL Server Graph Database

SQL Server 2019

- Edge constraint cascade delete actions
- New graph function - SHORTEST_PATH
- Match support in MERGE



```
INSERT INTO Person VALUES (1, 'John')
```

```
INSERT INTO likes VALUES (  
    (SELECT $node_id FROM Person WHERE id = 1)  
    ,(SELECT $node_id FROM Restaurant WHERE id = 1)  
    , 9);
```

```
-- Find Restaurants that John's friends like  
SELECT Restaurant.name  
FROM Person person1, Person person2, likes, friendOf,  
Restaurant  
WHERE MATCH(person1-(friendOf)->person2-(likes)->Restaurant)  
AND person1.name='John';
```

Table	Table Type
Person	NODE
Restaurant	NODE
City	NODE
Likes (rating)	EDGE
friendOf	EDGE
livesIn	EDGE
LocatedIn	EDGE

Utf-8 Support

Can be used to "replace" Unicode

If a UTF-8 enabled collation (`_UTF8`) is used, previously non-Unicode data types (**char** and **varchar**) become Unicode (UTF-8) data types.

```
SELECT Name, Description
FROM fn_helpcollations()
WHERE Name like 'Danish%UTF8';
```

UTF-8 Encoding Table

Bits	Byte 1	Byte 2	Byte 3	Byte 4	Byte 5	Byte 6
7	0xxxxxxx					
11	110xxxxx	10xxxxxx				
16	1110xxxx	10xxxxxx	10xxxxxx			
21	11110xxx	10xxxxxx	10xxxxxx	10xxxxxx		
26	111110xx	10xxxxxx	10xxxxxx	10xxxxxx	10xxxxxx	
31	1111110x	10xxxxxx	10xxxxxx	10xxxxxx	10xxxxxx	10xxxxxx

	Name	Description
1	Danish_Greenlandic_100_CI_AI_SC_UTF8	Danish-Greenlandic-100, case-insensitive
2	Danish_Greenlandic_100_CI_AI_WS_SC_UTF8	Danish-Greenlandic-100, case-insensitive
3	Danish_Greenlandic_100_CI_AI_KS_SC_UTF8	Danish-Greenlandic-100, case-insensitive
4	Danish_Greenlandic_100_CI_AI_KS_WS_SC_UTF8	Danish-Greenlandic-100, case-insensitive
5	Danish_Greenlandic_100_CI_AS_SC_UTF8	Danish-Greenlandic-100, case-insensitive
6	Danish_Greenlandic_100_CI_AS_WS_SC_UTF8	Danish-Greenlandic-100, case-insensitive
7	Danish_Greenlandic_100_CI_AS_KS_SC_UTF8	Danish-Greenlandic-100, case-insensitive
8	Danish_Greenlandic_100_CI_AS_KS_WS_SC_UTF8	Danish-Greenlandic-100, case-insensitive
9	Danish_Greenlandic_100_CS_AI_SC_UTF8	Danish-Greenlandic-100, case-sensitive
10	Danish_Greenlandic_100_CS_AI_WS_SC_UTF8	Danish-Greenlandic-100, case-sensitive
11	Danish_Greenlandic_100_CS_AI_KS_SC_UTF8	Danish-Greenlandic-100, case-sensitive
12	Danish_Greenlandic_100_CS_AI_KS_WS_SC_UTF8	Danish-Greenlandic-100, case-sensitive
13	Danish_Greenlandic_100_CS_AS_SC_UTF8	Danish-Greenlandic-100, case-sensitive
14	Danish_Greenlandic_100_CS_AS_WS_SC_UTF8	Danish-Greenlandic-100, case-sensitive
15	Danish_Greenlandic_100_CS_AS_KS_SC_UTF8	Danish-Greenlandic-100, case-sensitive
16	Danish_Greenlandic_100_CS_AS_KS_WS_SC_UTF8	Danish-Greenlandic-100, case-sensitive

SQL Server Extensibility

R, Python, or Java
on the same
computer but not
in SQL process

Extend the T-SQL
language

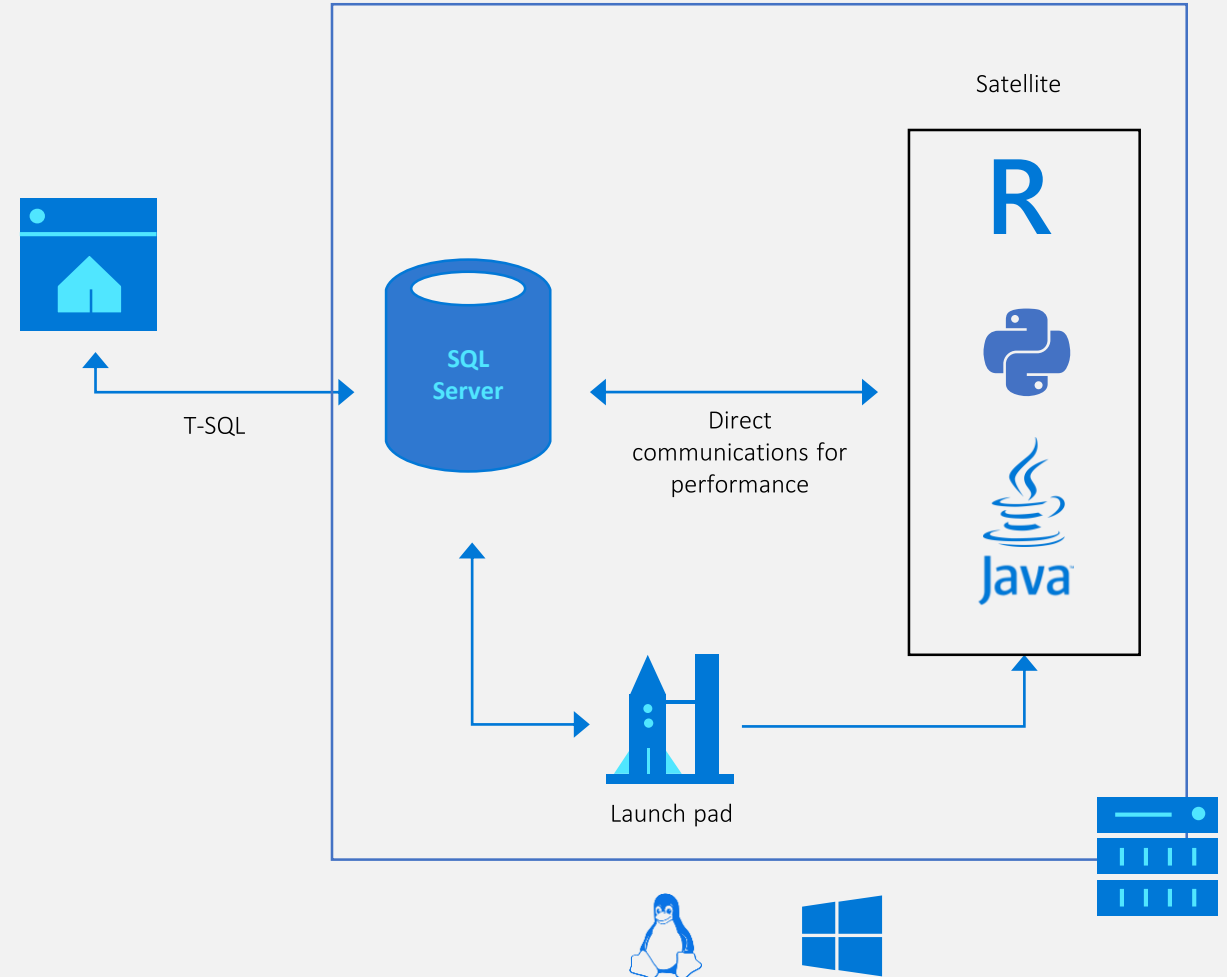
Machine Learning

Compat across
Windows and Linux

Secure and isolated

Resource governor
enabled

Extend T-SQL with R, Python, and Java




SQL Server Language Extensions

- Used for executing external code
- Uses the [Extensibility Framework](#) for executing external code.
 - Code execution is isolated from the core engine processes
 - Launchpad
- CREATE EXTERNAL LANGUAGE
 - Create new language (only Java.... Yet)
- EXECUTE EXTERNAL SCRIPT permissions (SQL Server 2019)
 - external script execution can be granted on specific languages

```
GRANT EXECUTE EXTERNAL SCRIPT ON EXTERNAL LANGUAGE  
::Java TO mylogin;
```


Availability



Availability



**RESUMABLE ONLINE
INDEX CREATION**




**ONLINE CLUSTERED
COLUMNSTORE INDEX
CREATION AND REBUILD**



**COMPATIBILITY
CERTIFICATION**



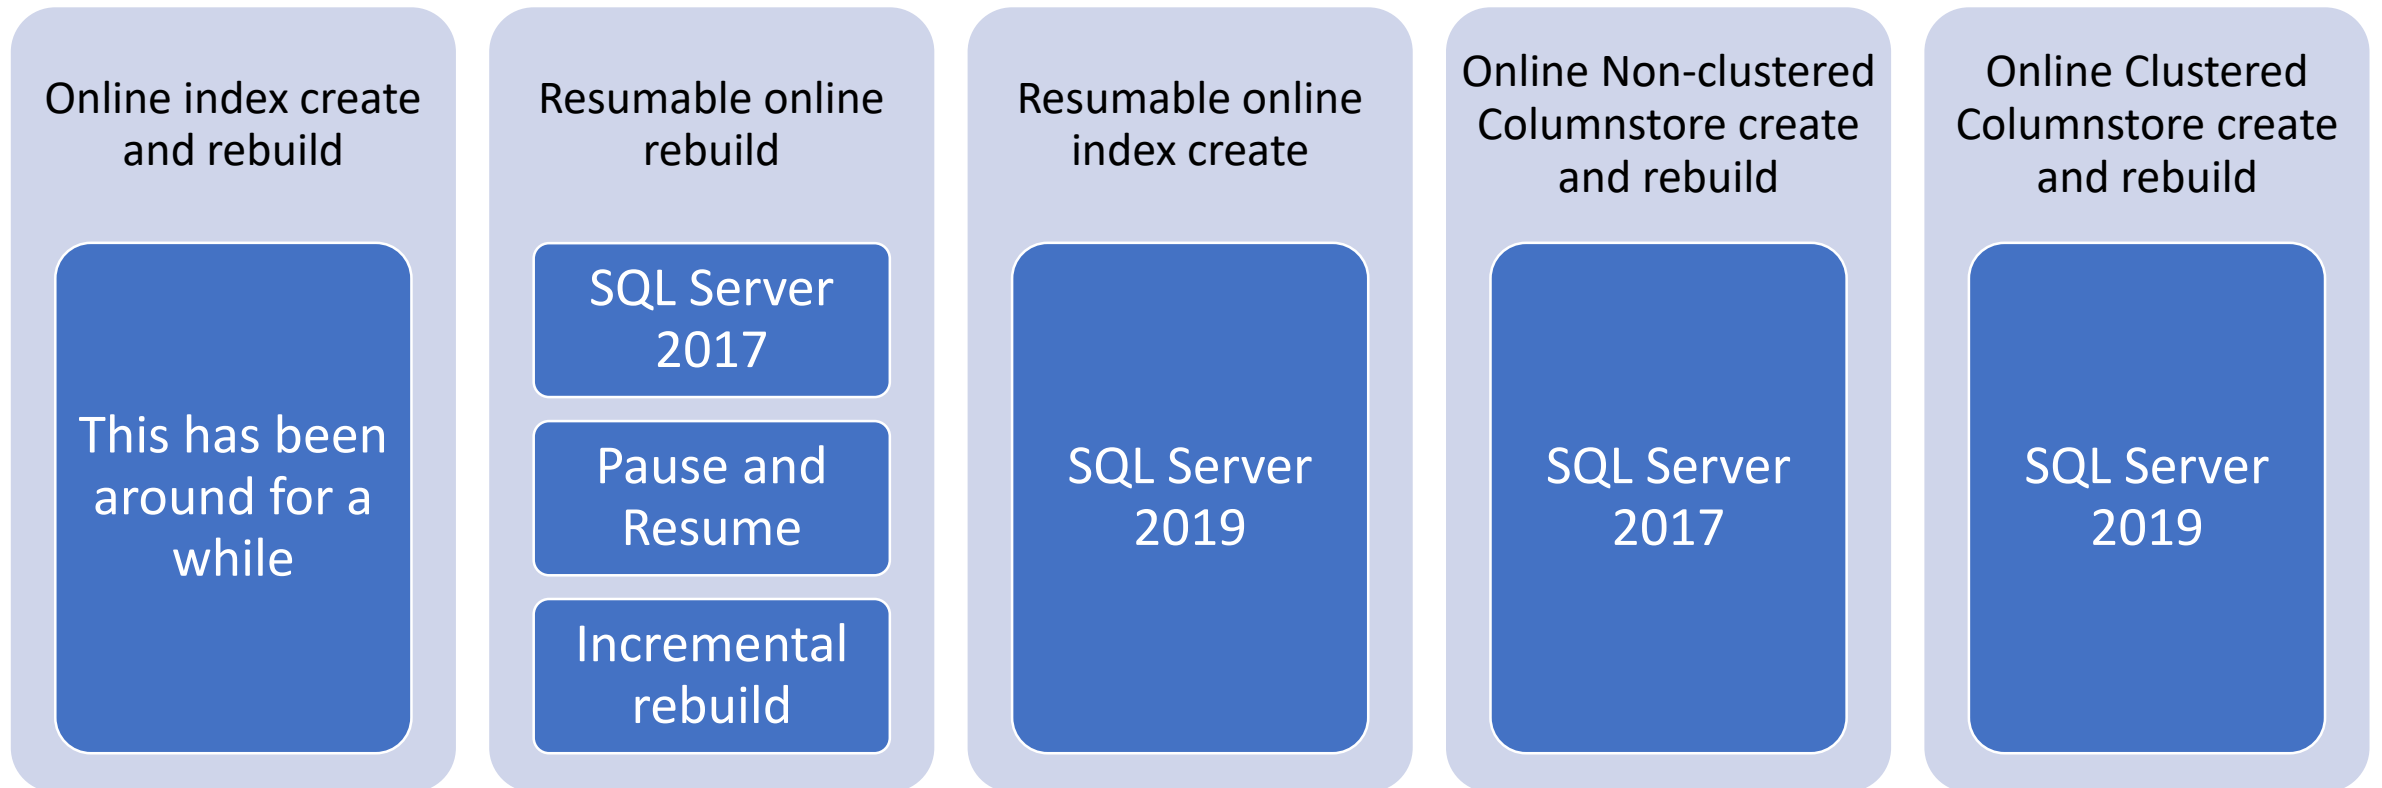
**ACCELERATED
DATABASE RECOVERY**



**AVAILABILITY GROUPS
ON KUBERNETES**

Using Online Indexing in SQL Server 2019

Index maintenance causes concurrency problems and application downtime




Accelerated Database Recovery

Instant Rollback and Database Recovery

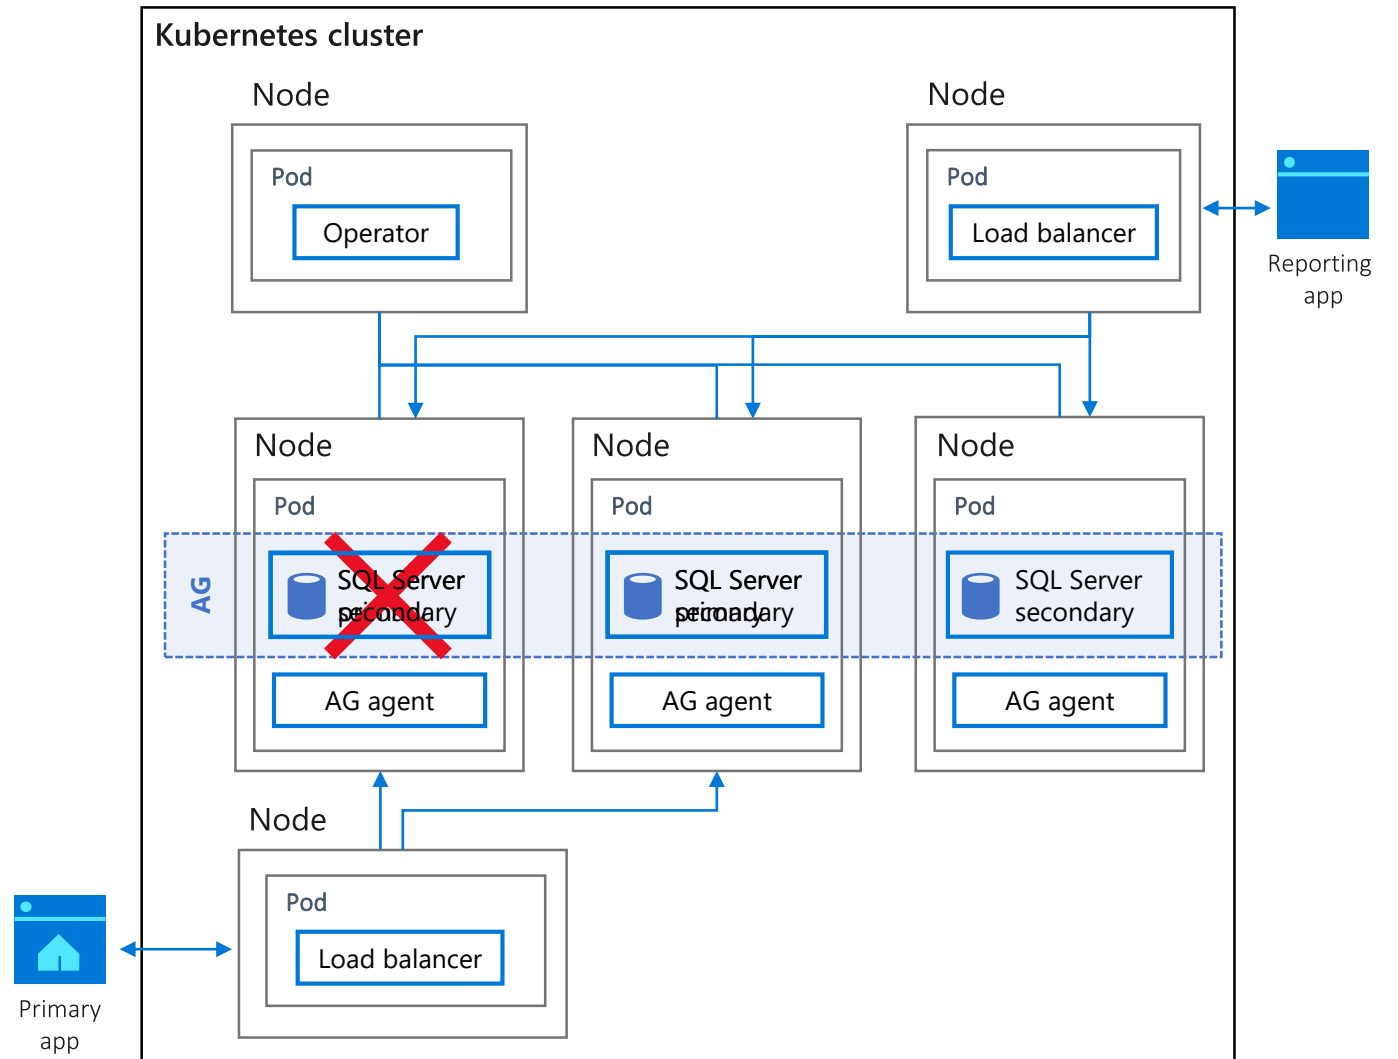
Before ADR
Three phases of database recovery.
(Abnormal query end)

Once we enabled **Accelerated Database Recovery** on a SQL Server Database, it stores the version of all modifications.



It is similar to versioning in **Read Committed Snapshot Isolation** level.

Availability groups on Kubernetes



Database Compatibility Certification

Stop certifying for any given platform (Cloud, on-prem)!

Stop certifying for a named SQL Server version!

- Any certification process should be thought in terms of "*which target database compatibility level am I certifying to?*"
- Updated public documentation: <http://aka.ms/dbcompat>

Performance and Availability



Intelligent
Performance
with SQL
Server

Gain
performance
without
changing the
application

Intelligent Query Processing

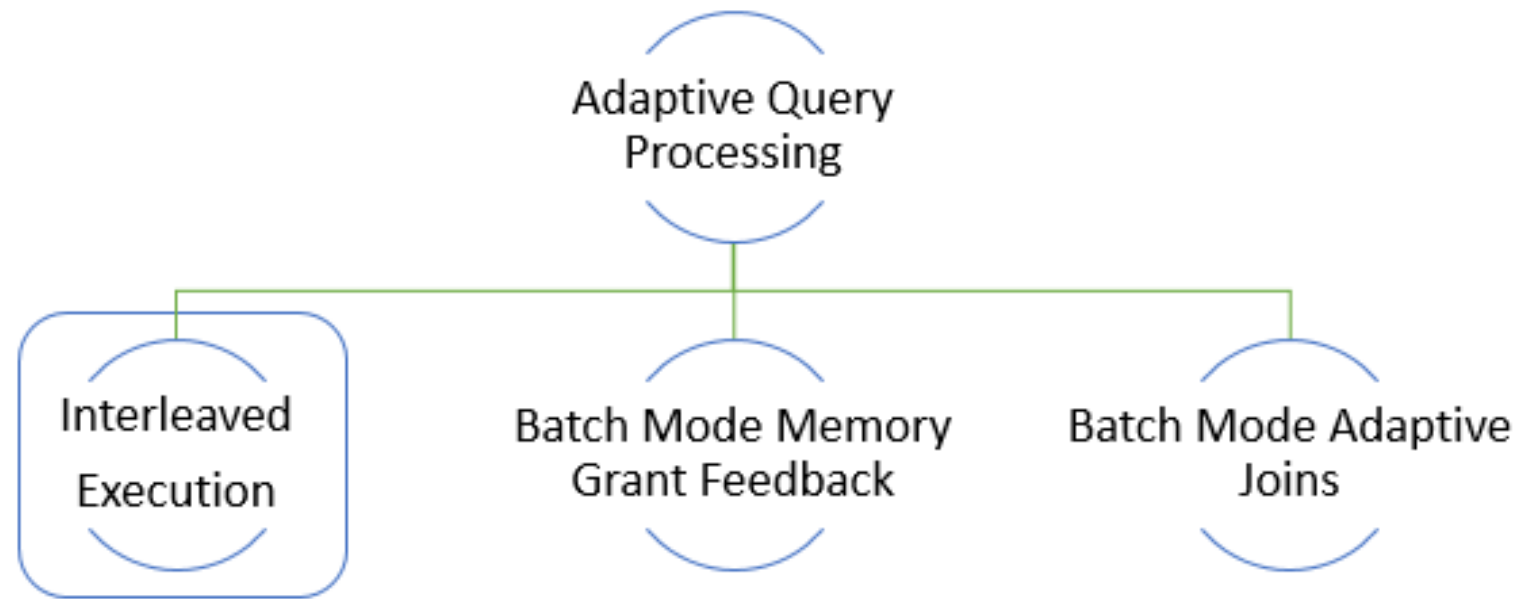
Performance insights anytime and anywhere with
Lightweight Query Profiling

Accelerating I/O performance with persistent
memory

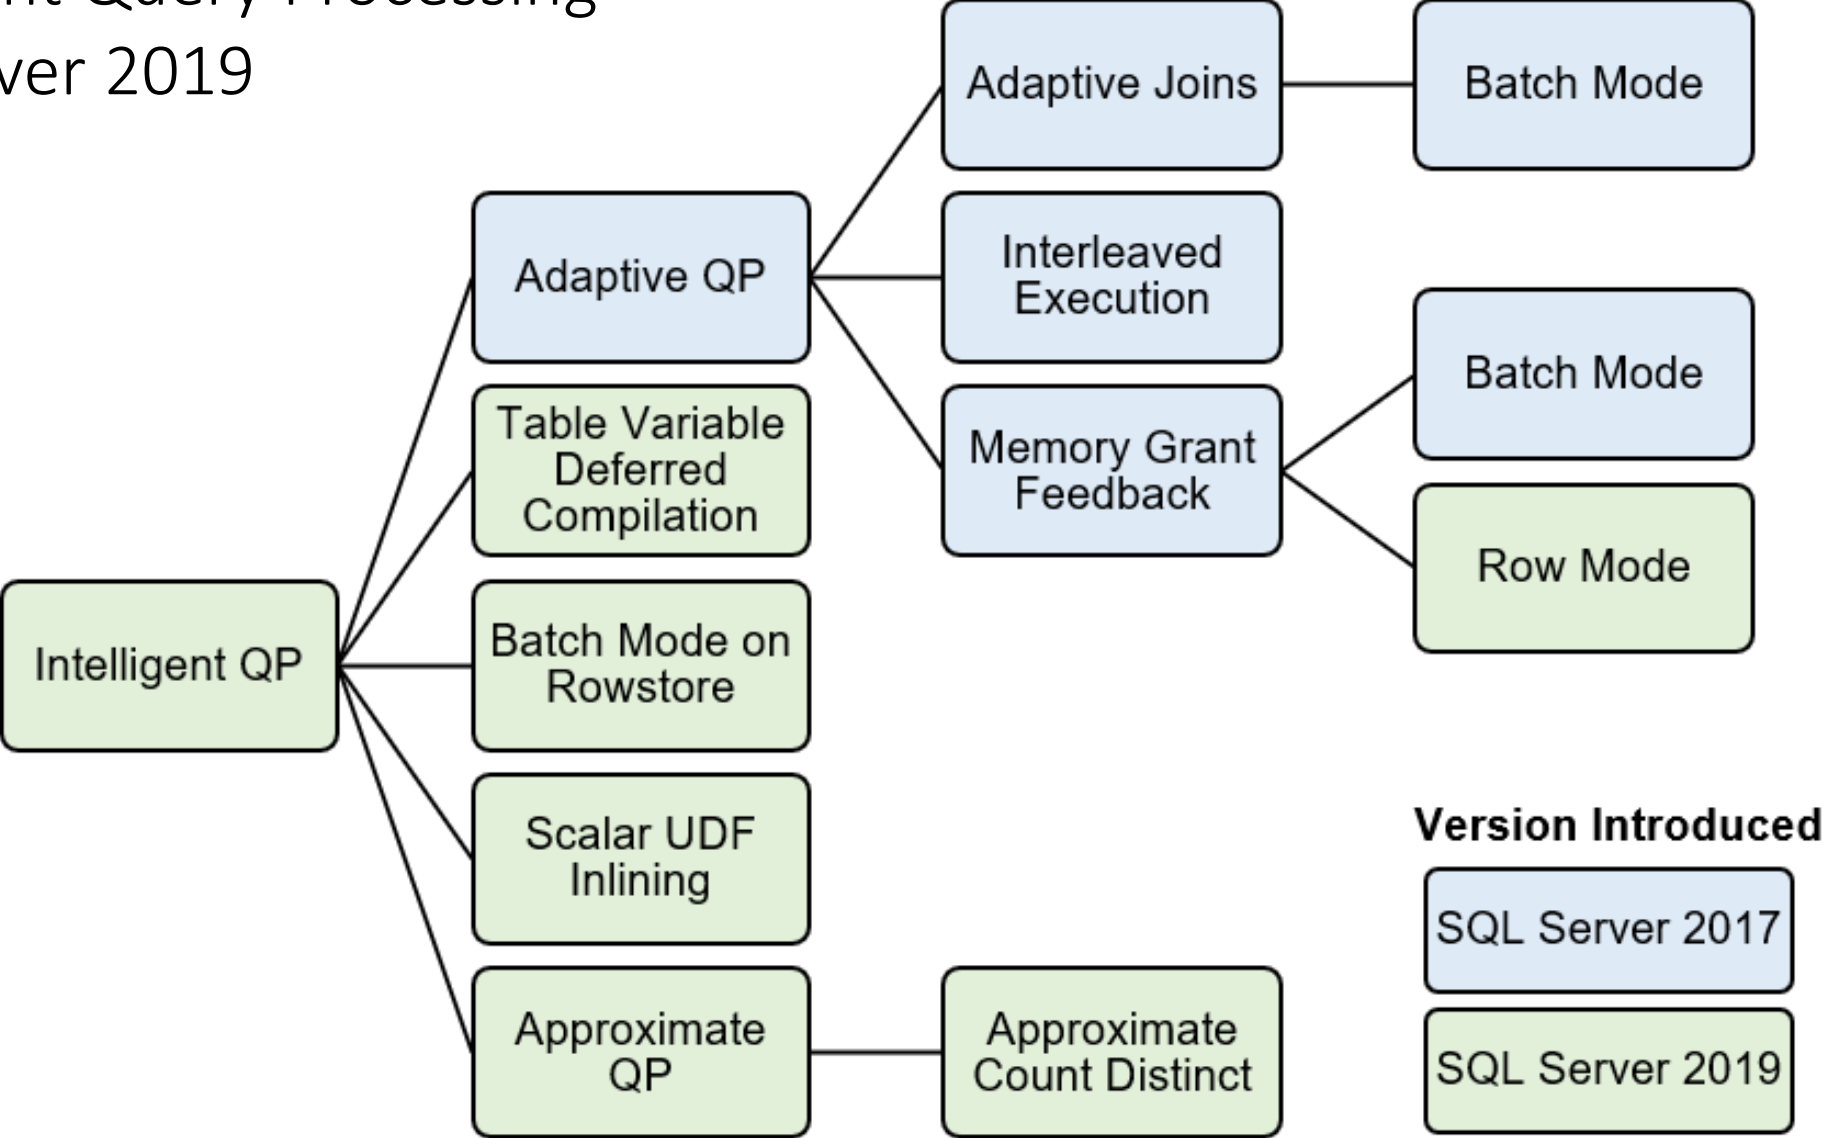
Hybrid Buffer Pool

Memory-Optimized TempDB Metadata

SQL Server
2017



Intelligent Query Processing SQL Server 2019





Intelligent Query Processing

Adaptive Joins (Batch Mode)	2017 (140)	Adaptive joins dynamically select a join type during runtime based on actual input rows.
Approximate Count Distinct	2019 - CTP2.0	Provide approximate COUNT DISTINCT for big data scenarios with the benefit of high performance and a low memory footprint.
Batch Mode on Rowstore	2019 – CTP2.0 (150)	Provide batch mode for CPU-bound relational DW workloads without requiring columnstore indexes.
Interleaved Execution	2017 (140)	Use the actual cardinality of the multi-statement table valued function (MSTVF) encountered on first compilation instead of a fixed guess.
Memory Grant Feedback (Batch Mode)	2017 (140)	If a batch mode query has operations that spill to disk, add more memory for consecutive executions. If a query wastes > 50% of the memory allocated to it, reduce the memory grant side for consecutive executions.
Memory Grant Feedback (Row Mode)	2019- CTP2.0 (150)	If a row mode query has operations that spill to disk, add more memory for consecutive executions. If a query wastes > 50% of the memory allocated to it, reduce the memory grant side for consecutive executions.
Scalar UDF Inlining	2019- CTP2.1 (150)	Scalar UDFs are transformed into equivalent relational expressions that are "inlined" into the calling query, often resulting in significant performance gains.
Table Variable Deferred Compilation	2019- CTP2.0 (150)	Use the actual cardinality of the table variable encountered on first compilation instead of a fixed guess.

Inlineable Scalar UDFs requirements

A scalar T-SQL UDF can be inline if all of the following conditions are true:

- The UDF is written using the following constructs:
 - `DECLARE` , `SET` : Variable declaration and assignments.
 - `SELECT` : SQL query with single/multiple variable assignments¹.
 - `IF / ELSE` : Branching with arbitrary levels of nesting.
 - `RETURN` : Single or multiple return statements.
 - `UDF` : Nested/recursive function calls².
 - Others: Relational operations such as `EXISTS` , `ISNULL` .
- The UDF does not invoke any intrinsic function that is either time-dependent (such as `GETDATE()`) or has side effects³ (such as `NEWSEQUENTIALID()`).
- The UDF uses the `EXECUTE AS CALLER` clause (the default behavior if the `EXECUTE AS` clause is not specified).
- The UDF does not reference table variables or table-valued parameters.
- The query invoking a scalar UDF does not reference a scalar UDF call in its `GROUP BY` clause.
- The UDF is not natively compiled (interop is supported).
- The UDF is not used in a computed column or a check constraint definition.
- The UDF does not reference user-defined types.
- There are no signatures added to the UDF.
- The UDF is not a partition function.

Lightweight Query Profiling

I want to see the details of a query plan at the operator level for any **active** executing query

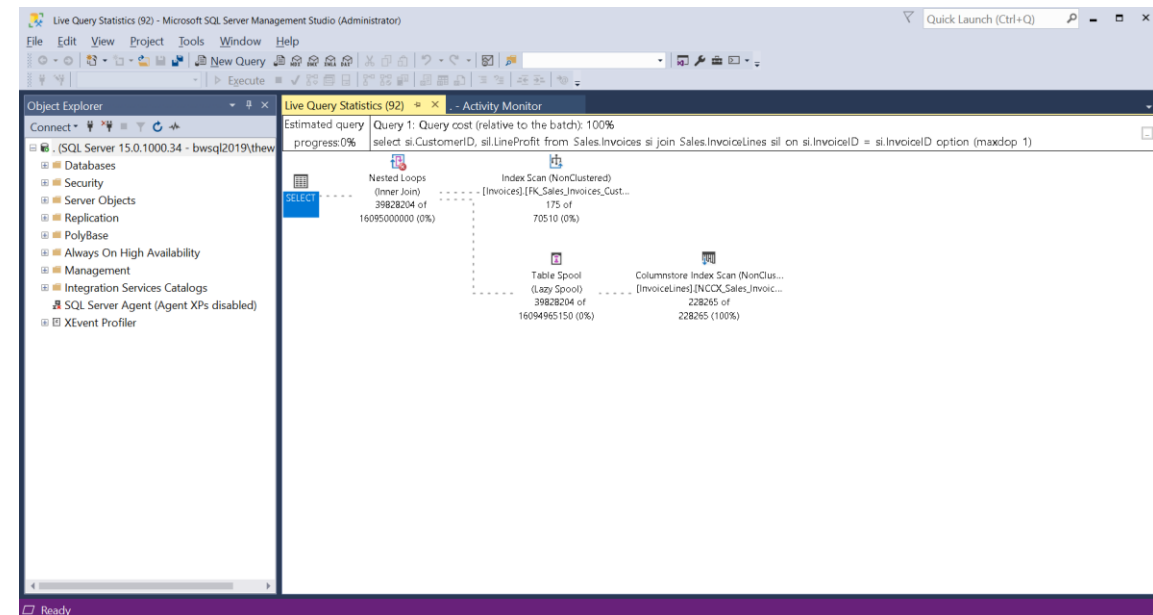
SQL Server 2016 SP1/2017

- Trace flag 7412
- XEvent query_thread_profile

SQL Server 2019

- On by default. No knobs required

dm_exec_query_profiles
dm_exec_query_statistics_xml



Security



Security

Confidential Computing



Always Encrypted with secure enclaves



Data Classification and auditing built-in




Manage certificates easier with SQL
Configuration Manager



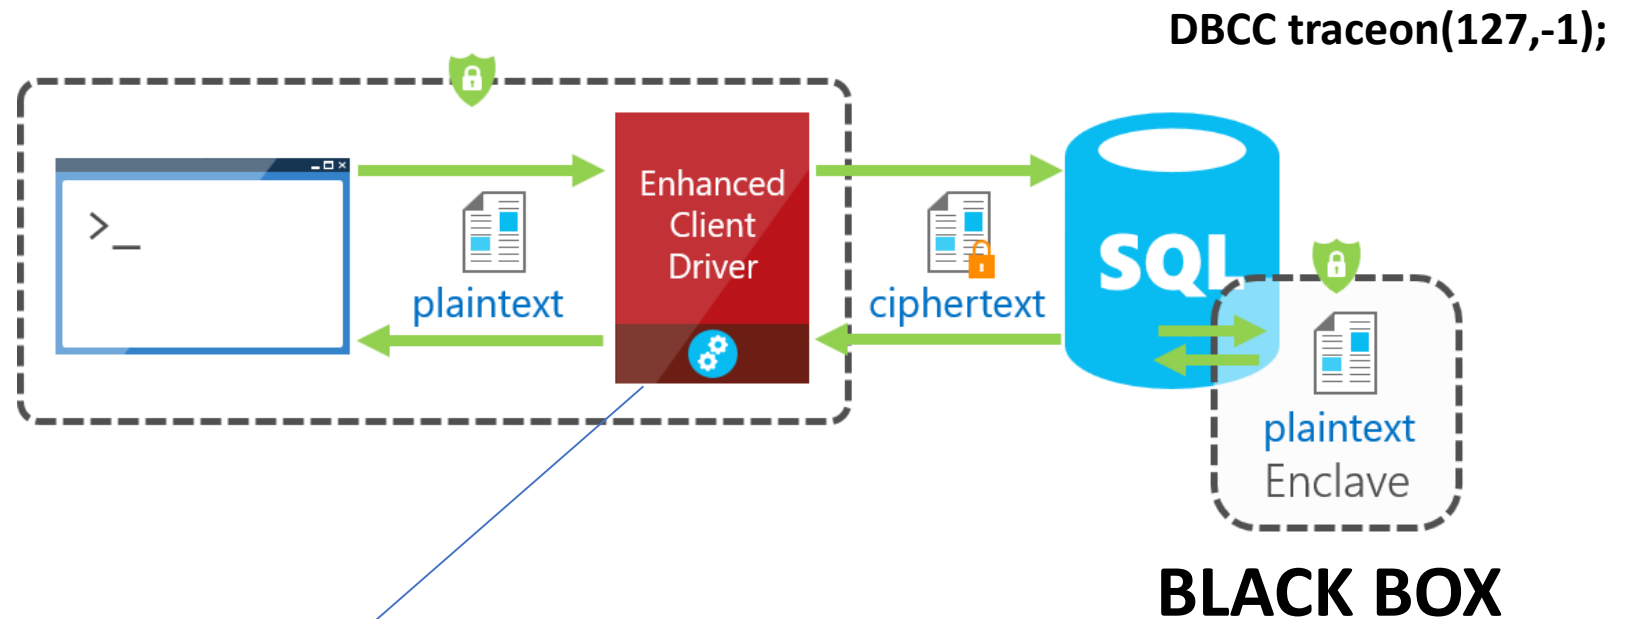
Static Data Masking

Always Encrypted (SQL 2016)

- Transparent to applications
- Protects Confidentiality from
 - malware
 - high-privileged *unauthorized* users
- Data on the server is Encrypted – done on the client side
- Functionality is restricted
 - Only comparison (if encryption is deterministic)



Always Encrypted with Secure Enclaves

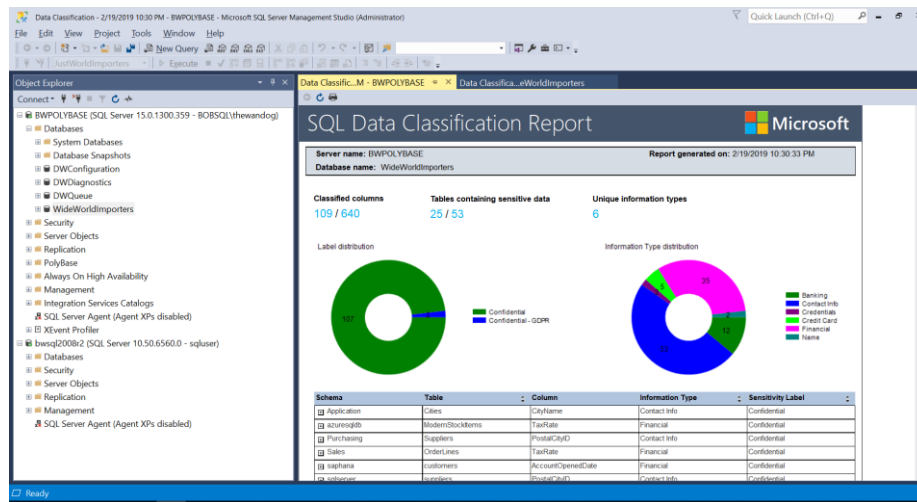


sends the column encryption keys required for the operations to the secure enclave over a secure channel

Data Classification and Auditing

Classify data in SQL Server and audit access to the classified data

SQL Server 2017



SQL Server 2019 and Azure SQL Database



ADD SENSITIVITY CLASSIFICATION TO
dbo.sales.price, dbo.sales.discount
WITH (LABEL='Highly Confidential',
INFORMATION_TYPE='Financial')

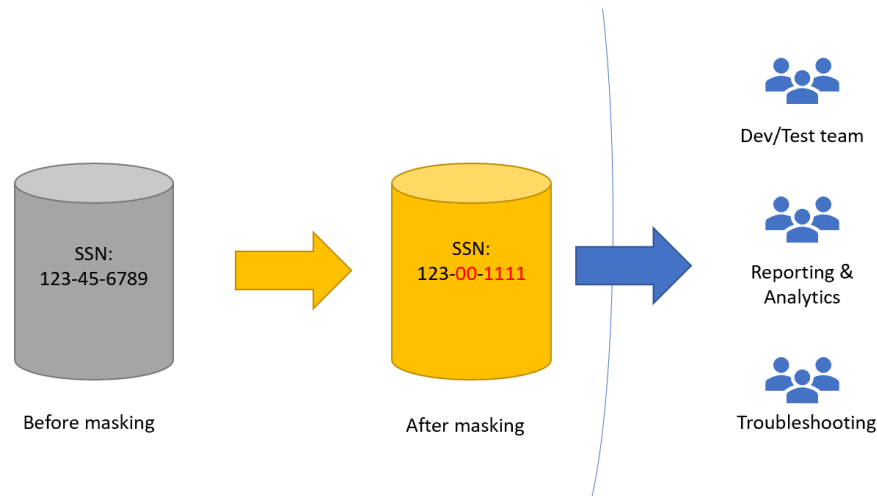
SQL Server Auditing

data_sensitivity_information

Who, what, and when accessed my classified data?

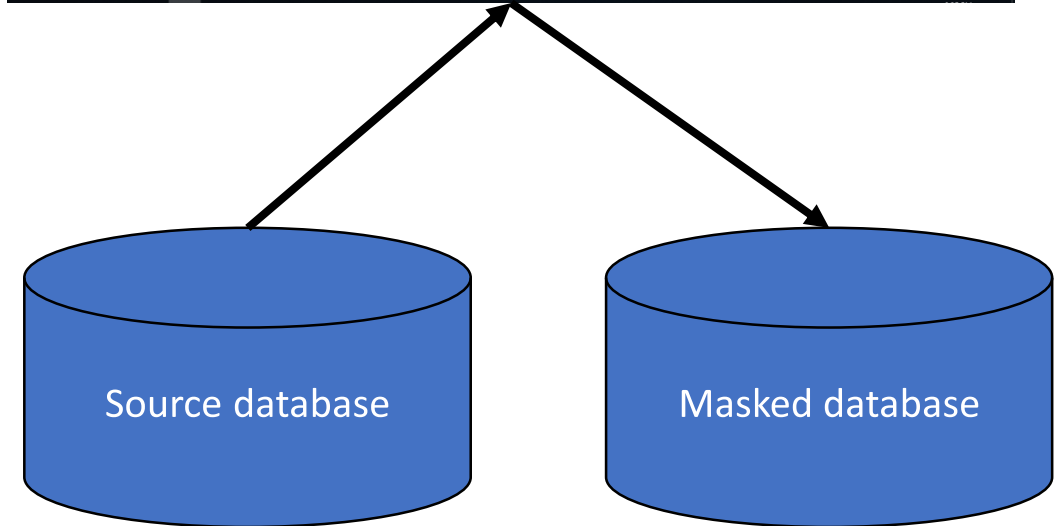

Static Data Masking

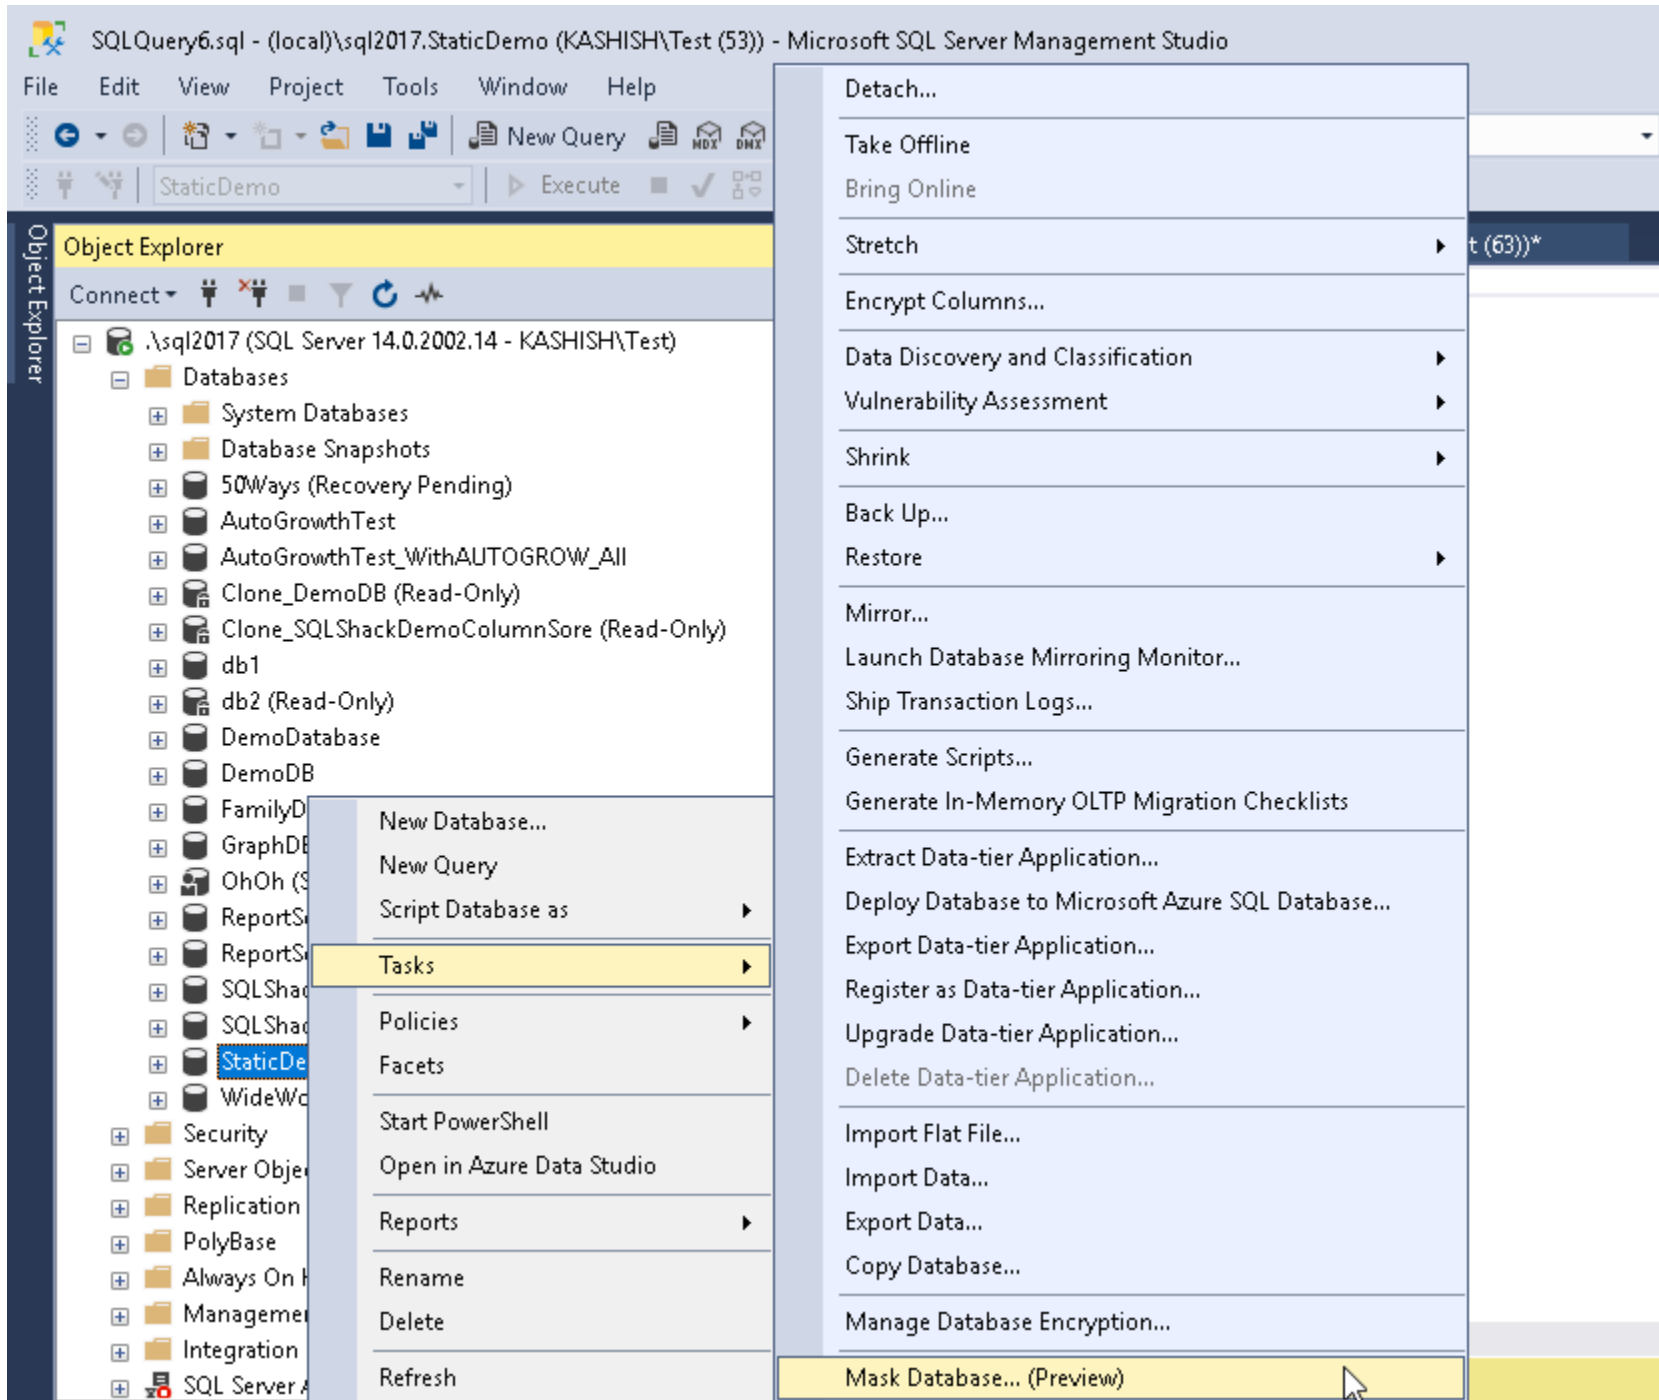
provide a copy of production data with certain data masked



- Choose masking functions
- Save choices to configuration file
- Some limitations
- Supports Azure SQL Database
- Update stats and use TDE

SSMS 18.0

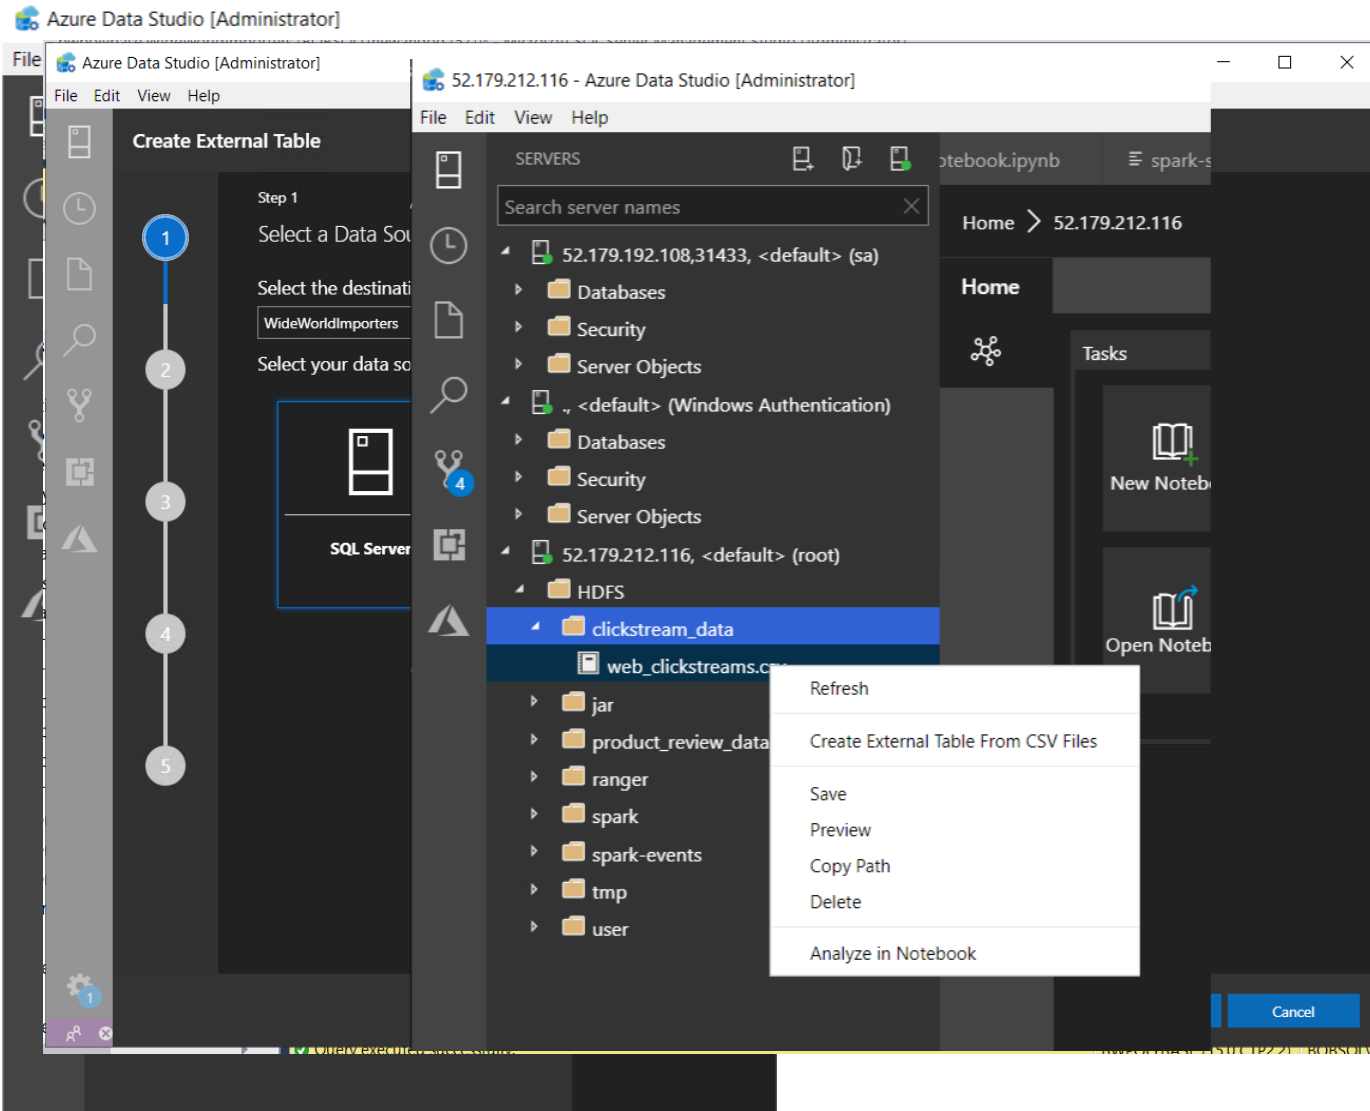




Other additions



Azure Data Studio and Data Virtualization





Using [PROSE](#) for intelligent import and schema detection

Code generation in Notebooks

Import wizards

External Table detection for HDFS



That's All ???



Transact-SQL Additions

```
SELECT APPROX_COUNT_DISTINCT ( id )  
FROM dbo.PoolEntrance
```



Microsoft®
SQL Server®
Analysis Services

SQL Server 2019 Analysis Services RC1

- [Calculation Groups](#)
 - Calculation groups for calculation reusability in complex models ([CTP 2.3](#))
 - Dynamic measure formatting with calculation groups ([CTP 3.0](#))
 - Custom ordering of calculation items in calculation groups
- [Query interleaving](#) with short query bias for high-concurrency workloads
- [Online attach](#) for optimized synchronization of read-only replicas
- [Many-to-many relationships](#) can help avoid unnecessary “snowflake” models ([CTP 2.4](#))
- [Governance](#)
 - Governance settings to protect server memory from runaway queries ([CTP 2.4](#))
 - Governance setting to control Power BI cache refreshes

Calculation Groups

- [Calculation groups](#) address the issue of proliferation of measures in complex BI models often caused by common calculations like ***time-intelligence***.
- SSAS models are reused throughout large organizations, so they tend to grow in scale and complexity.
- Calculation groups work with **explicit** DAX measures, e.g. **Sales** could be an explicit measure already created in the model.

Lets see the Magic

The screenshot displays a Power BI report interface. On the left, a table shows sales data for 2012 and 2013, categorized by month. The right side of the interface features the 'Visualizations' and 'Fields' panes. The 'Visualizations' pane shows a table icon selected, and the 'Fields' pane shows a search bar and a list of data sources including 'DimDate' and 'Time Intelligence'. The 'Time Calculation' checkbox is visible in the 'Time Intelligence' section.

CalendarYear	Sales
2012	
January	\$495,364
February	\$506,994
March	\$373,483
April	\$400,336
May	\$358,878
June	\$555,160
July	\$444,558
August	\$523,917
September	\$486,177
October	\$535,159
November	\$537,956
December	\$624,502
2013	
January	\$857,690
February	\$771,349
March	\$1,049,907
April	\$1,046,023
May	\$1,284,593
June	\$1,643,178
July	\$1,371,676
August	\$1,551,066
September	\$1,447,496
October	\$1,673,293
November	\$1,780,920
December	\$1,874,360

Visualizations

Rows

- Calendar
- CalendarYear
- EnglishMonthName

Columns

Add data fields here

Values

- Sales

Fields

Search

- Averages
- Currency Conversion
- DimCurrency
- DimCustomer
- DimDate
- DimProduct
- FactCurrencyRate
- FactInternetSales
- Time Intelligence
 - Time Calculation

Filters

Visual level filters

- CalendarYear is (All)
- EnglishMonthName is (All)
- Sales is (All)
- Time Calculation is Current, MTD, YTD, ...

Page level filters

Calculation Groups Custom Ordering

Custom order

Time Calculation

- Current
- MTD
- QTD
- YTD
- PY
- PY MTD
- PY QTD
- PY YTD
- YOY
- YOY%

Alphabetical order (less intuitive)


Time Calculation

- Current
- MTD
- PY
- PY MTD
- PY QTD
- PY YTD
- QTD
- YOY
- YOY%
- YTD


Query interleaving

Allows system configuration for improved user experiences in **high-concurrency [scenarios](#)**

FIFO



Query interleaving with short query bias



Online attach

Can be used for synchronization of read-only replicas in [on-premises query scale-out](#) environments (read only).

A typical usage pattern could be as follows:

1. DB1 (version 1) is already attached on read-only server B.
2. DB1 (version 2) is processed on the write server A.
3. DB1 (version 2) is detached and placed on a location accessible to server B
4. The <Attach> command with **AllowOverwrite=True** is executed on server B with the new location of DB1 (version 2).



Microsoft
SQL Server
Reporting Services

SQL Server Reporting Services

- Support for hosting your database catalog with Azure SQL Managed Instance
 - host your database catalog used for SQL Server Reporting Services
- A new version of Microsoft Report Builder updated for use with SSRS 2019
- Accessibility
 - AltText (alternative text) in support of WCAG and 508 compliance for your report elements
 - [PDF/UA](#) support for PDF documents (coming soon)
- Rollup of several security updates and bug fixes previously released in Power BI Report Server



Microsoft®
SQL Server®
Integration Services

SQL Server Integration Services 2019



SSIS is now an
[Extension](#) in Visual
Studio

Like SSRS and
SSAS became
in VS2017



Azure Enabling
takes up a lot of
space

Lift and shift
packages to
Azure



New Power Query [Source](#)

Summing Up

Top 10 reasons to choose SQL Server 2019

Bring the industry-leading performance and security of SQL Server to your choice of language, platform, and data—structured and unstructured

1. Harness the power of big data



Big data clusters with scalable compute and storage composed of SQL Server, Spark, and HDFS. Cache data in scale-out data marts.

2. Bring AI to your workloads



A complete AI platform to train and operationalize models in SQL Server ML Services or Spark ML using Azure Data Studio notebooks.

3. Eliminate the need for data movement



Data virtualization allows queries across relational and non-relational data without movement or replication.

4. Explore and interact with visual data



Visual data exploration and interactive analysis using SQL Server BI tools and Power BI Report Server.

5. Run real-time analytics on operational data



In-memory technologies for analytics on operational data using HTAP. Higher concurrency and scale through persistent memory.

6. Automatically tune SQL Server



Intelligent Query Processing improves scaling of queries and Automatic Plan Correction resolves performance problems.

7. Reduce database maintenance and increase business uptime



Greater uptime with more online indexing operations. Now run Always On availability groups on containers using Kubernetes.

8. Boost security and protect data in use



SQL Server enables layers of security including protection of computations in Always Encrypted secure enclaves.

9. Track compliance with sophisticated resources



Data Discovery & Classification labeling for GDPR and Vulnerability Assessment tool to track compliance.

10. Optimize with choice and flexibility



Support for your choice of Windows, Linux, and containers. Run Java code on SQL Server and store and analyze graph data.



Microsoft®
SQL server 2008



Pragmatic BI

Pragmatic BI Aps

Åhusene 4,9,3

8000 Aarhus C

Telefon +45 42 606 202

E-mail steen@pragbi.dk

Internet www.pragbi.dk

CVR nr. 36974133