

What's New in SQL Server 2016?

In this article I will be discussing about the SQL Server 2016 new features in brief.

In-place Upgrade options for SQL Server 2016

- SQL Server 2008 SP4
- SQL Server 2008 R2 SP3
- SQL Server 2012 SP2
- SQL Server 2014 RTM
- SQL Server 2014 SP1

Supported Operating Systems for SQL Server 2016

- Windows 8
- Windows Server 2012
- Windows 8.1
- Windows Server 2012 R2
- Windows 10
- Windows Server 2016

Memory, CPU, Licensing for SQL Server 2016

- Same editions as SQL2014
- Same Limitations
- Same Licensing

List of New Features in SQL Server 2016

Database Engine

Always Encrypted

- With Always Encrypted, SQL Server can perform operations on encrypted data, and best of all the encryption key resides with the application inside the customer's trusted environment and not on the server.
- Always Encrypted secures customer data so DBAs do not have access to plain text data.
- Encryption and decryption of data happens transparently at the driver level minimizing changes that have to be made to existing applications.

Backup to Microsoft Azure

- SQL Server backup to URL using Microsoft Azure Blob storage services now supports using block blobs instead of page blobs. Block blobs have a size limitation of 200GB. However, by striping your backup across multiple files, the maximum backup size is 12.8 TB.

Columnstore Indexes

- Kind of like each column getting its own index.
- In SQL Server 2012:
 - Nonclustered only
 - Made the table read-only when created
- In SQL Server 2014:
 - Clustered columnstore introduced, writeable
 - Massive compression (60-90%)
 - Nonclustered still read-only
- Nonclustered columnstore indexes now updateable, can be filtered
- Clustered columnstore indexes can get regular b-tree indexes atop them
- That improves concurrency with row-level locking
- Better performance for:
 - Aggregates (MIN, MAX, SUM, COUNT, AVG)
 - String predicates (myfieldname LIKE '%foo')
- Better concurrency (support of Snapshot Isolation and Read Committed Snapshot Isolation)
- Readable on AlwaysOn Availability Group secondary replicas
- Better index reorganize results (removes deleted rows, less memory pressure)

Dynamic Data Masking

- Dynamic data masking limits sensitive data exposure by masking it to non-privileged users.
- Dynamic data masking helps prevent unauthorized access to sensitive data by enabling customers to designate how much of the sensitive data to reveal with minimal impact on the application layer.
- It's a policy-based security feature that hides the sensitive data in the result set of a query over designated database fields, while the data in the database is not changed

High Availability Enhancements

- Adds Distributed Transaction Coordinator (DTC) support
- Faster throughput
- Load balancing for readable secondaries
- Database-level health detection, failover
- Available in Standard Edition, but:
 - Only one replica
 - Only one database per Availability Group (you can set up multiple Availability Groups per server though)
 - Secondaries aren't readable, can't do backups

In-Memory Online Transaction Processing (OLTP)

- Improvements to In-Memory OLTP enable scaling to larger databases and higher throughput in order to support bigger workloads. In addition, a number of limitations concerning tables and stored procedures have been removed to make it easier to migrate your applications to and leverage the benefits of In-Memory OLTP.
- Support for performing ALTER operations for memory-optimized tables and natively compiled stored procedures.
- Use MARS (Multiple Active Result Set) connections to access memory-optimized tables and natively compiled stored procedures
- This allows leveraging In-Memory OLTP in existing applications that rely on MARS.
- Support for natively compiled, scalar user-defined functions.
- Complete support for collations
- Enhancements to transaction performance analysis reports and Storage improvements
- Support for nesting of natively compiled stored procedures, and other increases in the Transact-SQL surface area.
- Support for subqueries in natively compiled stored procedures.
- PowerShell cmdlet for evaluating the migration fitness of multiple objects in a SQL Server database.

Format Query Results as JSON

- Use the FOR JSON clause, for example, to delegate the formatting of JSON output from your client applications to SQL Server.

- Here's a sample query that uses the FOR JSON clause.

Query

```
Transact-SQL
SELECT name, surname
FROM emp
FOR JSON AUTO
```

Result

```
JSON
[
  { "name": "John" },
  { "name": "Jane", "surname": "Doe" }
]
```

- Parsing and storing native JSON data (coming)

Live Query Statistics

- SQL Server Management Studio provides the ability to view the live execution plan of an active query.
- This live query plan provides real-time insights into the query execution process as the controls flow from one query plan operator to another.
- The live query plan displays the overall query progress and operator-level runtime execution statistics such as the number of rows produced, elapsed time, operator progress, etc.
- Because this data is available in real time without needing to wait for the query to complete, these execution statistics are extremely useful for debugging query performance issues.

Query Store

- Query store is a new feature in that provides DBAs with insight on query plan choice and performance.
- It simplifies performance troubleshooting by enabling you to quickly find performance differences caused by changes in query plans.
- The feature automatically captures a history of queries, plans, and runtime statistics, and retains these for your review.
- It separates data by time windows, allowing you to see database usage patterns and understand when query plan changes happened on the server.
- The query store presents information by using a Management Studio dialog box, and lets you force the query to one of the selected query plans.

Row-Level Security

- Row level security introduces predicate based access control.
- It features a flexible, centralized, predicate-based evaluation that can take into consideration metadata (such as labels) or any other criteria the administrator determines as appropriate.
- The predicate is used as a criterion to determine whether or not the user has the appropriate access to the data based on user attributes.
- Label based access control can be implemented by using predicate based access control.

Stretch Database

- Stretch Database is a new feature that leverages resources in Windows Azure to store and query archival data.
- Stretch Database automatically archives eligible rows from Stretch-enabled tables and uses computational resources in Azure to offload queries over the archived rows.

Temporal Tables

- A temporal table is a new type of table that provides correct information about stored facts at any point in time.
- Each temporal table consists of two tables actually, one for the current data and one for the historical data.
- The system automatically ensures that when the data changes in the table with the current data the previous values are stored in the historical table.
- Querying constructs are provided to hide this complexity from users.

Integration Services

- AlwaysOn support
- Power Query integration
- Incremental package deployment
- Project Upgrade
- OData 4.0 support
- Hadoop File System (HDFS) support
- JSON support
- Azure Data Factory integration (hybrid)
 - Execute on-premises SSIS packages from Azure Data Factory
 - Azure Data Factory data flow task
 - Azure storage connector
 - Azure commandlets
- New Oracle/Teradata connector (4.0)
- Incremental deployment options
- Custom logging levels
- SSIS package templates to reuse ETL code

Reporting Services

- New look and feel and possibility to apply themes and branding using CSS
- New visualizations, chart types like tree maps and sun bursts
- Improved flexible parameter panel with support for:
 - Autocomplete
 - Search
 - Hierarchical tree display
- Power Query as a data source
- Pin on-premises SSRS reports to Power BI Dashboards
- Report Builder supports High DPI scaling and devices
- Runs in all modern browsers on both desktops as tablets (any device)
- Enhancements to Reporting Services Subscription features
 - Enable and disable subscriptions
 - Subscription description

- Change subscription owner
- Shared credential for file share subscriptions

Analysis Services

Analysis Services - Tabular

- Enhanced modelling capabilities in the semantic layer
- Many-to-many relationships
- BI Directional cross filtering. This means you can not only filter on the 1 side of a 1 to many relationships in your tabular model, but also on the many side. For example, two connected tables, Sales → Product:
 - Product: product, product category
 - Sales: sales date, connection to product table
 - Now select products sold filtering on sales date (many side) while also filtering on product category (1 side). This is not possible in today's version of SSAS tabular.
- Time intelligence
 - Date/time columns are automatically converted to rich date/time tables starting from the column's MIN date till the MAX date found
- New DAX functions
 - A lot of new functions that at the moment require quite complex formulas like present time, date difference, percentile, product, geomean, median, etc.
- Performance improvements
 - For end users
 - ✓ Query engine optimized
 - For developers
 - ✓ Metadata operations; modelling related operations are much faster
 - For data processing
 - ✓ Parallel partition processing
- Expose on-premises tabular models in the cloud (hybrid) → Power BI feature.

Analysis Services – Dimensional

- Netezza as a Data Source (Netezza Data Warehouse | IBM - NDM Technologies)
- Performance improvements
 - Unnatural hierarchies

- Distinct counts
- Other performance improvements in areas where multidimensional is not performant at the moment
- DBCC (DataBase Check Consistency) support. Checks the logical and physical integrity of objects in the specified database.
- Expose on-premises multidimensional cubes in the cloud with Power BI (hybrid)

Master Data Services

- Improved performance for large models
- Row-level compression per entity
- Improved user interface
- Configurable retention settings
- Enhanced security possibilities for read, write, delete and create operations and support for multiple system administrators with specific permissions
- Excel Add-in is 15 times faster and is updated to support bulk entity based staging operation

Tooling

Visual Studio 2015

- Database and BI project types merged into one Visual Studio
- SSIS designer supports previous versions of SQL Server
- Report Builder supports previous versions of SQL Server
- New scripting language for tabular models. Currently tabular models are wrapped into multidimensional constructs and when you deploy it will be reverse engineered to the tabular model. The new native language for tabular will be easy to understand, modify and deploy.

More information

- Public CTP available now
 - <https://www.microsoft.com/en-us/evalcenter/evaluate-sql-server-2016>
- Documentation
 - <https://msdn.microsoft.com/en-us/library/bb500435%28v=sql.130%29.aspx>