

SQL SERVER 2016 ROADMAP

What you need to know to build the right roadmap for your organization.

Why Upgrade to SQL Server 2016?

SQL Server 2016, released in June, 2016, is one of the most significant leaps forward in the history of the platform. So, what new features set it apart from what's happened in the past, and what are the key considerations you should have when looking forward to SQL Server 2016?

Data Intensity's Ukarsh Patel and Phil Harbour, SQL Server experts with more than 25 years combined experience, discussed upgrade road mapping in February 2017 at a Data Intensity webinar.

SQL Server 2016 Runs Much Faster

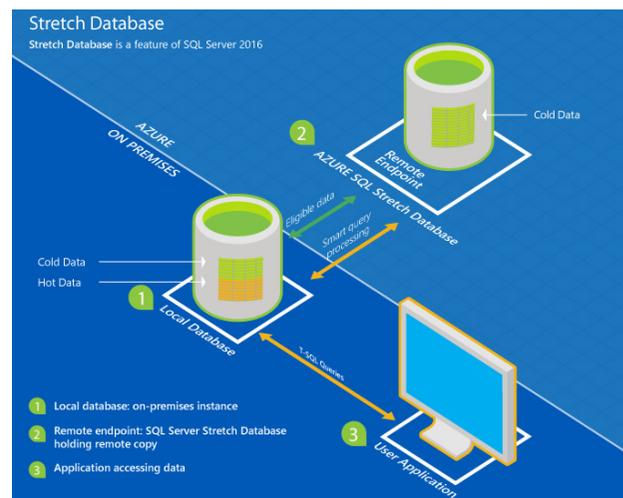
- Database consistency checks run 7x faster, reducing maintenance windows
- TEMPDB issues solved by enabling -T1117 and -T1118 flags automatically, for better scalability and performance
- SQL Server 2016 systems reporting more than eight CPUs per NUMA, now automatically configure Soft NUMA for you on a per-instance basis rather than per computer for double-digit performance gains
- Systems with multiple-NUMA nodes see improved throughput
- Thread scheduling improved
- Log block transportation improves 4x

New Features Are Important and Far-Reaching

STRETCH YOUR DB TO AZURE CLOUD

Until now, security and compliance requirements required you to retain "cold", historical data on expensive storage resources, causing performance issues, widening maintenance windows, and threatening backup and restore SLAs.

Stretch Database solves these problems by putting old data – which averages 70-80% of total storage – in lower cost cloud storage. Stretch an entire table or just the parts that you specify. All this is seamless to the user, who accesses historical data from the Azure cloud automatically. Backups of your local data complete much faster. Recoveries are completed in a hurry.

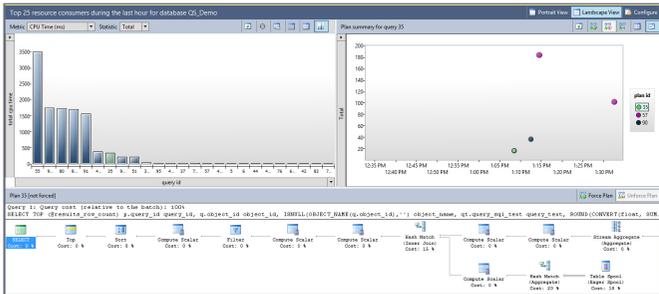


QUERY STORE, A NEW DATABASE PERFORMANCE TROUBLESHOOTING TOOL

Query Store tool is a fast way to put complaints about slow application queries to rest.



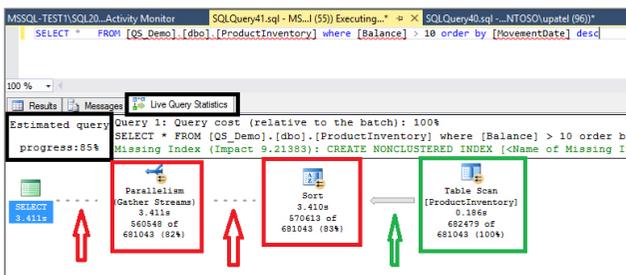
When the DBA enables Query Store using the Alter Database command, it analyzes query plan histories and identifies performance problems that have slowed access to the DBA due to execution plan changes over time.



Quickly identify top resources consumers by asking Query Store to check measures like CPU, duration, execution count, logical reads and writes, and more. The DBA can then force SQL Server to use an older, higher-performance execution plan to return performance to earlier levels.

LIVE QUERY STATISTICS FOR REAL-TIME INSIGHT AND DEBUGGING ON THE FLY

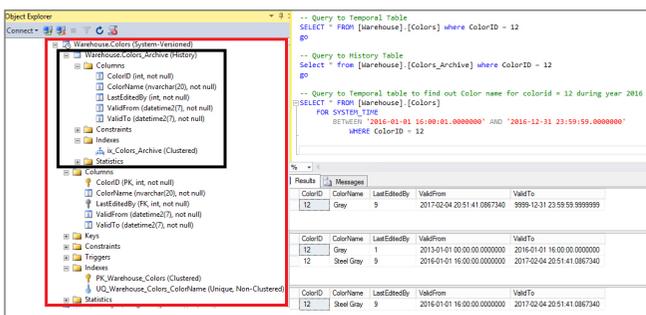
A new button next to the Actual Execution Plan called Include Live Query Statistics helps debug slow queries as they are running so you can accelerate completion. In the single-session example shown below...



... the green box indicates that the first part of the query has been completed. The middle red box shows that sorting is 83% complete, while total progress on the query (left hand red box) is 85% complete.

TEMPORAL TABLE FOR ANALYZING DATA CHANGES OVER TIME

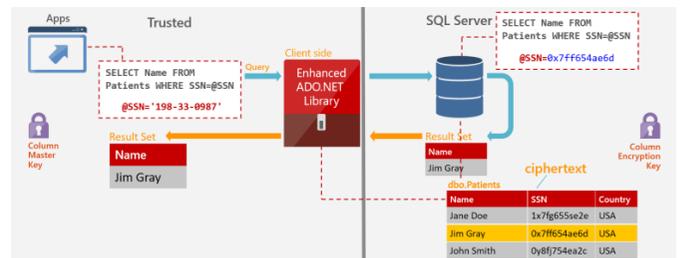
This feature captures each data update and deletion. The resulting history table allows you to carry out point-in-time analyses for data auditing, calculating business trends, investigating errors, data recovery reconstructing the state of data at any point, and more.



Stretch the history table to Azure to save storage and maintenance costs. Details on this feature are found at the 19:05 minute mark in the archive for the webinar presentation SQL Server 2016 Roadmap.

NEW FEATURE AUTOMATICALLY ENCRYPTS AND DECRYPTS DATA

The Always Encrypted function encrypts at the application layer using ADO.NET drivers within the application before passing data over the network. It is supported over both SQL Server and Azure SQL databases. Encryption keys are never revealed to the database; only data owners have access. It frees you to use third-party data administration while reducing security clearance overhead for your DBA staff.

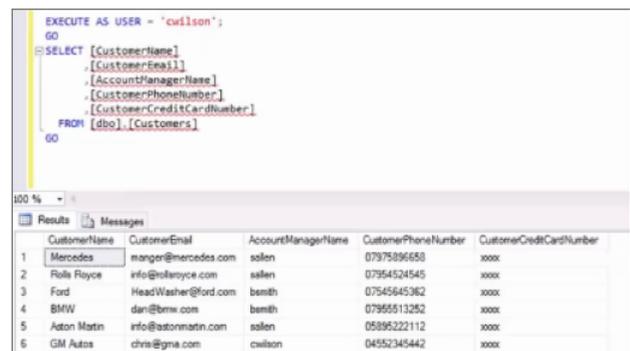


ROW LEVEL SECURITY: NEW APPROACH TO MASKING SENSITIVE DATA

Control access to data stored in the SQL Server database with the new Row Level Security (RLS) feature, implemented in the database. Filters can be applied so, for example, each account manager can see his or her customers' data, but no other's. Blocks can be applied to block write operations that violate a filter setting.

DYNAMIC DATA MASKING KEEPS SENSITIVE DATA FROM NON-PRIVILEGED USERS

This built-in feature thwarts non-privileged users from getting to sensitive data, and it's supported on both SQL Server 16 and the Azure SQL Database. It takes advantage of centralized logic to mask data whenever it's queried. It can be applied without affecting database operations and does not require application code changes. In the example below, customer credit card numbers are masked:



ALWAYSON AVAILABILITY GROUPS NOW AVAILABLE WITH STANDARD EDITION

Originally introduced in 2012 as an Enterprise Edition feature, AlwaysON Availability Groups have several feature enhancements:

- **Basic Availability Groups.** Bring a limited, but welcome level of high availability to Standard Edition Customers, with a failover environment for a single database, and the ability to operate across on-premises and Microsoft Azure.
- **Failover Policy enhancements.** Earlier versions only looked at the health of the instance itself; the SQL Server 2016 version looks after the health of the database too.
- **Distributed Transaction Coordinator support.** You'll need to be running Windows Server 2016 or patch Windows Server 2012 to run this. Best implemented during a maintenance window or downtime.
- **Up to three synchronous replicas as automatic failover targets.** Adds a third failover target and choose which one to fail over to. Improves availability while removing the need to manually manage failovers.
- **Improved synchronization throughput performance.** A near-10X improvement one that was required for SQL Server to keep up with newer, high-speed hardware. You will also experience faster database recovery times in case a failover.
- Load balancing across readable secondary replicas is not a true load balancing function (because one cannot be certain that each readable secondary will have the same amount of work to do), but it still delivers an overall performance improvement.

TEMPDB ENHANCEMENT WORTH NOTING

Database files can hinder performance if they're not configured properly, and SQL Server 16 TempDB enhancement makes it easy to create files based on the number of logical processors. This helps them to grow equally, up to a newly increased default size of 64MB.

How to Upgrade to SQL Server 2016

The following instances are upgradable to the latest version of SQL Server: 2008, 2008 R2, 2012 and 2014. Running SQL Server 2005? Migrate first to one of the above versions. A few upgrade scenarios are not supported:

- Cross-platform upgrades. SQL Server 2016 is 64bit only.
- WOW
- Can't upgrade from previous Evaluation Edition version.

STEP-BY-STEP APPROACH

Follow this step-by-step roadmap for your SQL Server 2016 upgrade:

- Check with applications vendors to confirm their application is supported on the new platform
- Run Data Migration Assistant (DMA)
- Perform a current system configuration health check
- Gather 24-hour performance statistics to create a benchmark
- Carry out your upgrade side by side, not in place
- Run the upgrade first in a test and development environment with full regression testing performed by the application team, and have them sign off on it.
- Build the environment in the production infrastructure and carry out production dry runs to check that everything – security, jobs and necessary file configuration – is behaving as expected
- Perform the migration itself, using HA methods or backup strategies to minimize downtime
- Complete a post-upgrade review and benchmark the new system to ensure performance that is equal to, if not superior to the previous version.

Learn more

For more information on this subject, you're welcome to contact us by email at Info@dataintensity.com, or follow us on **Twitter: @DataIntensity**

About Data Intensity

Data Intensity is the leading independent provider of managed and cloud consulting services for enterprise databases, applications, business intelligence solutions and analytics. The company combines best-of-category technology, world-class services, a flexible business model, and deep-rooted expertise gained from hundreds of successful deployments. Our mission is to cost-effectively support the full scope of a customer's enterprise data lifecycle. [Contact us](#) for more information. Copyright © 2017 Data Intensity, LLC. All Rights Reserved.