



Microsoft SQL Server

Achieving mission-critical application performance with SQL Server

The volume and complexity of data is increasing everywhere, and the ability to use it is becoming more important for organizations across industries. Digital transactions and experiences are expected to be available across geographies and time zones. Today the data generated by those millions of interactions touches every device, application and process.

IT organizations must balance the impact of this explosion in data against the need for constant uptime, worldwide availability, increasing regulatory compliance, and an evolving digital world where customers expect and demand more.

More than ever, companies need mission-critical performance, stability, security, availability, and the ability to use their data to develop real insights about their business.

This white paper describes how the continuing advancement of Microsoft SQL Server is designed to deliver all of those capabilities at a lower cost and higher return.

Copyright

The information contained in this document represents the current view of Microsoft Corp. on the issues discussed as of the date of publication. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information presented after the date of publication.

This white paper is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS, IMPLIED, OR STATUTORY, AS TO THE INFORMATION IN THIS DOCUMENT.

Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in, or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording or otherwise), or for any purpose, without the express written permission of Microsoft Corp.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights or other intellectual property.

© 2015 Microsoft Corp. All rights reserved.



Contents

- Executive Summary **4**
- SQL Server evolution **5**
- Mission-critical application performance with SQL Server..... **7**
 - Performance..... **7**
 - In-memory online transaction processing..... **8**
 - 30x performance gains..... **8**
 - Reduced database size and increased performance: data and backup compression..... **9**
 - Native JSON support..... **9**
 - Security and compliance..... **9**
 - Secure by default: lowering vulnerability..... **9**
 - Transparent Data Encryption..... **10**
 - Backup Encryption..... **10**
 - New in 2016: Always Encrypted..... **11**
 - New in 2016: Row-Level Security..... **12**
 - Built-in tools for enabling compliance: SQL Server audit tools..... **12**
 - Availability..... **13**
 - AlwaysOn..... **13**
 - Enhanced AlwaysOn..... **14**
 - Scalability..... **15**
 - Columnstore..... **15**
- The cloud and mission-critical systems..... **16**
 - Traditional bare-metal deployment..... **16**
 - Hybrid solutions..... **17**
 - Cloud deployment: public cloud and private cloud..... **18**
 - Private cloud..... **19**
 - Stretch Databases..... **19**
 - Always Encrypted..... **20**
- Comparing the cost of mission-critical capabilities..... **21**
 - Expecting the fundamentals..... **21**
 - Understanding options..... **21**
- Conclusion..... **23**
- More information..... **23**

Executive summary

Web interactions, brick-and-mortar transactions, mobile apps on cellular phones, and intelligent devices on the factory floor: Today there is an explosion of data sources that is driving an exponential boom in the volume of data.

It may seem overwhelming, but the good news is organizations of all types have the opportunity to use that data to drive smarter decisions. Today data is a currency, a compass, and a new natural resource that can be tapped to provide new forms of value.

Companies that recognize this are differentiating themselves, and thriving. A recent IDC study looked at companies that have a data-centric culture versus those without. The results showed a massive “data dividend” for companies adapting to this new data-heavy environment — as much as \$1.6 trillion in new value worldwide.

Hungry to harness those returns, organizations are looking at data in new ways. They want to use information that is both relational and nonrelational, internal and external. They are applying new analytical models on historical data to predict the future. They are finding new insights and sharing them broadly in the organization. And they are doing all of this at speed — near real time in many cases. The result is more productivity, more efficiency, and faster innovation, all of which help increase sales and revenue streams.

These realities of business are driving the next wave of innovation on the Microsoft platform. To support those needs for today and tomorrow, we are investing heavily into several core areas as SQL Server continues to evolve:

- **Performance:** SQL Server’s integrated in-memory toolset goes far beyond isolated features and provides support for improving performance dramatically in a wide range of scenarios.
- **Security and compliance:** As SQL Server progresses, we’re adding new capabilities to protect data both at rest and in motion, with new features like Always Encrypted and Row-Level Security.
- **Availability:** Already known for rock-solid, reliable performance, SQL Server is adding significant new enhancements to AlwaysOn, better load balancing, and new features for flexible and efficient backups.
- **Scalability:** New advancements in compute, storage and networking will provide a direct impact on mission-critical SQL Server workloads.
- **Cloud services:** New tools in SQL Server and Microsoft Azure make it even easier to scale to the cloud; to build patching, backup and disaster recovery solutions; and to access resources wherever they are — on-premises, private cloud or public cloud.

This white paper discusses the new features, and those that are coming soon, that continue to make SQL Server the most cost-effective, flexible and powerful solution for managing data today.

SQL Server evolution

We believe our customers want to utilize the key capabilities that SQL Server and the broader Microsoft Data Platform have to offer, and this is the unique design approach we took back when we first started improving analytics by building in-memory into PowerPivot for millions of rows of data analysis in Excel.

In SQL Server 2012, we expanded our in-memory footprint with the same built-in approach by adding in-memory to Analysis Services so IT could build data models much faster. In SQL Server 2014, we covered the final workload by introducing an in-memory online transaction processing (OLTP) solution — to significantly speed transactional performance. We also enhanced the in-memory columnstore with faster performance and significantly higher data compression so memory utilization can be optimized.

Customers have responded to this evolution by continuing to show confidence in using SQL Server to manage their mission-critical data. Industry analysts have also responded positively. Gartner recently rated SQL Server as having the most complete vision of any operational database management system (Figure 1).

In addition, SQL Server has consistently added groundbreaking functionality over the past 15 years (Figure 2).



The evolution of SQL Server continues ...

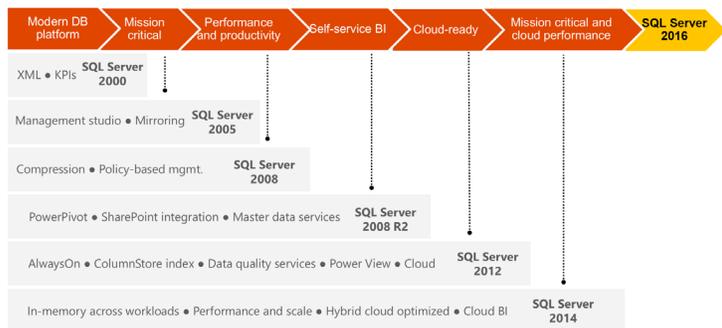


Figure 2: SQL Server functionalities added across releases

in motion, and a world-class high-availability and disaster recovery solution adds new enhancements to AlwaysOn technology.

Organizations will gain deeper insights into all of their data with new capabilities that go beyond business intelligence to perform advanced analytics directly within their database and present rich visualizations for business insights on any device.

You can also gain the benefits of hyper-scale cloud with new hybrid scenarios enabled by new Stretch Database technology that lets you dynamically stretch your warm and cold transactional data to Microsoft Azure in a secured way so your data is always at hand for queries, no matter the size. In addition, SQL Server 2016 delivers a complete database platform for hybrid cloud, enabling you to easily build, deploy and manage solutions that span on-premises and cloud.

SQL Server 2016 continues this tradition of innovation, and in fact may be the biggest leap forward in Microsoft's data platform history, with real-time operational analytics, rich visualizations on mobile devices, built-in advanced analytics, new advanced security technology, and new hybrid cloud scenarios.

SQL Server 2016 delivers breakthrough mission-critical capabilities with in-memory performance and operational analytics built in. Comprehensive security features like new Always Encrypted technology help protect your data at rest and in motion, and a world-class high-availability and disaster recovery solution adds new enhancements to AlwaysOn technology.

¹ Gartner, <http://www.gartner.com/technology/reprints.do?id=1-237UHKQ&ct=141016&st=sb>, October 2014

SQL Server 2016 introduces many new features and enhancements:

- **Enhanced in-memory performance** provides up to 30x faster transactions, more than 100x faster queries than disk-based relational databases and real-time operational analytics
- **New Always Encrypted technology** helps protect your data at rest and in motion, on-premises and in the cloud, with master keys sitting with the application, without application changes
- **Stretch Database technology** keeps more of your customers' historical data at your fingertips by transparently stretching your warm and cold OLTP data to Microsoft Azure in a secure manner without application changes
- **Built-in advanced analytics** provide the scalability and performance benefits of building and running your advanced analytics algorithms directly in the core SQL Server transactional database
- **Business insights through rich visualizations** on mobile devices with native apps for Windows, iOS and Android
- **Simplified management** of relational and nonrelational data by querying both with T-SQL using PolyBase
- **Faster hybrid backups, high-availability and disaster recovery scenarios** to back up and restore your on-premises databases to Microsoft Azure and place your SQL Server AlwaysOn secondaries in Azure

Mission-critical application performance with SQL Server

Today the dramatic shift toward complex, unstructured data types requires organizations to embrace back-end solutions that support all data sources, while also providing high-performance, mission-critical capabilities.

IT organizations need to balance this amplified importance of data with other pressures, such as globalization, conservative budgets, and ever-tightening compliance policies and regulations that are increasingly strict in many areas of the world. And they need to accomplish it all with higher levels of uptime and performance, with layered defenses that provide the utmost security.

For SQL Server 2016, we are continuing to push the envelope on mission-critical performance with new innovations across many components covering performance, security, availability and scalability.

Mission-critical performance

Performance	Security	Availability	Scalability
Operational analytics <ul style="list-style-type: none">• Insights on operational data• Works with in-memory OLTP and disk-based OLTP In-memory OLTP for more applications Query store <p>Monitor and optimize query plans</p> Native JSON <p>Expanded support for un-structured JSON data</p> Temporal database support	Always Encrypted <p>Sensitive data remains encrypted at all times with ability to query</p> Row level security <p>Apply fine-grained access control to table rows</p> Dynamic Data Masking <p>Return masked data to non-privileged users (e.g. credit card numbers)</p> Other enhancements <ul style="list-style-type: none">• Audit success/failure of database operations• TDE support for storage of In-memory OLTP Tables• Enhanced auditing for OLTP with ability to track history of record changes	Enhanced AlwaysOn <ul style="list-style-type: none">• 3 synchronous replicas for auto failover across domains• Round robin load balancing of replicas• Automatic failover based on DB health• DTC for transactional integrity across database instances with AlwaysOn• Support for SSIS with AlwaysOn Enhanced online operations	Enhanced database caching <p>Cache data with automatic, multiple TempDB files per instance in multi-core environments</p> Support for Windows Server v-next

Figure 3: An overview of new features coming in SQL Server 2016

Performance

Data volume is increasing exponentially, and the ability of today's analysis and business intelligence tools to derive insight from that data is more important than ever. Modern database technologies must be able to take advantage of these amplified data streams across faster, parallel processors and great reservoirs of storage in order for businesses to compete.

With SQL Server 2014 and the upcoming SQL Server 2016, performance is enhanced with a number of new technologies, including in-memory, query store, JSON, and temporal support, to name a few.

SQL Server's integrated in-memory toolset goes far beyond isolated features and provides support for improving performance dramatically in a wide range of scenarios. These technologies include in-memory OLTP, primarily for transactional workloads, and in-memory columnstore, primarily for decision support workloads (this last one is discussed in the Deeper Insights white paper).

The new query store feature allows you to monitor query plans to optimize them for particular application scenarios, providing additional performance-tuning opportunities. Native JSON support in the core database engine is another new feature that provides support for working with schema-free data within SQL Server. The new temporal database features allow you to record, audit and query data changes over time.

Read on for a deeper dive into each of these evolving capabilities.

In-memory online transaction processing

SQL Server was originally designed when it could be assumed that main memory was very expensive, so data needed to reside on disk except when it was actually needed for processing. This assumption is no longer valid as memory prices have dropped enormously over the past 30 years. At the same time, multicore servers have become affordable, so that today one can buy a server with 32 cores and 1TB of memory for under \$50K.

Because of this trend to much more available memory and many more cores, the SQL Server team at Microsoft began building a database engine optimized for large main memories and many-core CPUs. This increased performance is obtained by making changes in three main areas compared with the previous disk-based storage for our relational data.

- 1. Data structures.** Completely new data structures have been designed for the rows and indexes of In-Memory OLTP. These structures are designed with multiversioning in mind, and are also designed to be updateable with no locking required.
- 2. No Locking and Latching.** Locking is one of the most common causes of long wait times and slow responses in a multiuser OLTP system. Because of the way the In-Memory OLTP data structures have been designed, no locking is required for any data manipulation operations. In addition, because data is not read in from disk into memory buffers, and the rows are not stored on pages, no page latches are needed for In-Memory OLTP data processing.
- 3. Native compilation.** Even though we talk about “recompiling” for our queries and procedures in SQL Server, the code is not truly compiled. It is translated into a lower level of code that can be interpreted, but it is not compiled into machine code. So execution of normal operations on disk-based tables requires line-by-line interpretation of each operation. But In-Memory OLTP supports natively compiled procedures that can access memory-optimized tables that will be truly compiled and then loaded as DLLs.

These major changes in data storage and access, plus many smaller changes, including much more efficient log writes, allow the incredible performance improvements obtained with SQL Server 2014 In-Memory OLTP.

30x performance gains

In-memory technology for SQL Server dramatically improves the throughput and latency of SQL Server OLTP capabilities. It is designed to meet the requirements of the most demanding transaction processing applications, and Microsoft has worked closely with a number of companies to prove these gains — Dell has achieved a 9x increase in performance. Beth Israel Deaconess Medical Center has cut query times by 75 percent while tripling its data storage from 30 days to 90 days.

We are continuing with our design point of workload optimized in-memory, but now allowing you to gain real-time insights on operational data with the ability to run a columnar index over your in-memory or on disk row store.

What’s unique here is you can gain the speed of In-Memory OLTP and have the ability to gain operation analytics. Even if you are not using in-memory you can gain operational insights. Also you can apply our in-memory technology to more of your applications than ever before with expanded T-SQL surface area support.

For SQL Server 2016, we are still keeping to the workload-optimized approach as customers want to optimize in-memory by workload. When it comes to In-Memory OLTP you will now be able to apply this tuned transaction performance technology to a significantly greater number of applications with expanded T-SQL surface area. In addition to providing up 30x performance gains you will now be able to gain real-time operational insights on your operational data. This data can be in-memory or on disk.

Reduced database size and increased performance: data and backup compression

Many organizations want to increase speed and reliability by putting more data onto specialized disk arrays or a SAN, but often they are prohibited by the cost of these high-end disk resources. Backup and data compression in SQL Server can free up space by dramatically reducing the size of databases. Reduced data size also can increase performance. With additional space, more data can be stored on the SAN. And because storing data on the SAN is more reliable, it also increases availability.

In addition, SQL Server enables data compression for people who use Unicode UCS-2. This capability enables organizations that have global language sets in their data storage to take advantage of data compression and experience the benefits of compression.

New in SQL Server 2016:

Native JSON support

JSON is the storage format used in several NoSQL engines, including Azure DocumentDB. DocumentDB uses Azure Blob storage to store schema-less documents but provides a rich SQL query dialect that allows you to conduct SQL queries over the data contained in the documents.

In addition to DocumentDB, Azure Search also utilizes JSON. Azure Search is a fully managed search solution that allows developers to embed sophisticated search experiences into Web and mobile applications without having to worry about the complexities of full-text search and without having to deploy, maintain or manage any infrastructure.

The combination of SQL Server's new support for JSON with these other Microsoft tools enables many scenarios for moving data back and forth between relational and schema-less storage and the applications that access such data.

We have added native JSON support in the core database engine so now you have support for schema-free data to tackle more diverse data types right in SQL Server. We also added temporal database support so you can record, audit and query data changes over time.

Security and compliance

The security landscape has changed dramatically over the years, but it's as important as ever. Today organizations must contend with an explosion in devices and device types, varying network technologies across the world, and data that resides in a multitude of formats and platforms — including the cloud.

How do you protect data that is constantly in motion? How do you provide the right access to the right people at the right time? As SQL Server continues to evolve, we're adding new capabilities to protect data both at rest and in motion, with new features like Always Encrypted and Row-Level Security. Other new features enhance security in a multitenant environment, with fine-grain access control based on user attributes such as location, role and more.

With new capabilities for 2016, SQL Server is continuing to evolve to provide the kinds of finely tuned, granular control that companies need today.

Secure by default: lowering vulnerability

Microsoft and the SQL Server team take security seriously. More than 10 years ago, Microsoft implemented the Trustworthy Computing initiative. The initiative requires SQL Server engineers to take regular security training and carry that responsibility for security across their job duties, regardless of the group in which they reside. This companywide discipline to protect security and privacy was developed to create software that is secure by design — and to reduce by default the overall risks related to security.

To that end, according to the National Institute of Standards and Technology (NIST) public security board, SQL Server reportedly has the lowest number of security vulnerabilities across the major database vendors. In addition, SQL Server has been deemed "the most secure database" by the Information Technology Industry Council (ITIC).²

² Information Technology Intelligence Corp. (ITIC), SQL Server Delivers Industry-Leading Security, September 2012

The ability to encrypt data no matter where it resides and while it's in transit enables the kinds of anytime, anywhere access that ensures only the right person with the right credentials can read the data but can do so on any device, virtually anywhere in the world. SQL Server 2016 will provide a number of enhancements to data encryption that ensure that data is safe no matter where it goes, but is still consumable by users in today's dynamic business environment.

For 2016, SQL Server continues to refine the ways in which it can handle and secure sensitive data. These enhancements are designed to help companies remain agile and competitive, while at the same time maintaining compliance with today's more convoluted regulatory landscape.

Transparent Data Encryption

SQL Server's Transparent Data Encryption (TDE) allows organizations to encrypt data when it is stored on a disk, and decrypt it when it is read into memory. TDE uses a database encryption key (DEK), which is stored in the database boot record for availability during recovery. The DEK is a symmetric key secured by using a certificate stored in the master database of the server or an asymmetric key protected by an EKM module.

Simply put, TDE protects data at rest, meaning the data and log files. This enables software developers to encrypt data by using AES and 3DES encryption algorithms, without changing existing applications. Encryption and decryption operations are handled by the database engine in the background. Therefore, organizations do not have to make changes to their applications for SQL Server to secure their data. Because encryption is built into the database engine, it is transparent to applications and users — and it is included in SQL Server Enterprise edition.

In addition, extensible key management works with TDE to store encryption keys outside of the database. With extensible key management, organizations can use a hardware device or a third-party encryption tool to create encryption keys. Storing the keys separately from the encrypted data makes it even harder for unauthorized users to gain access to encrypted data.

For databases protected by TDE, backups of those databases are also encrypted.

Enhanced in SQL Server 2016:

TDE now supports storage of memory-optimized OLTP Tables. This allows for greater security along with the performance enhancements provided by memory-optimization.

Dynamic Data Masking limits exposure to sensitive data by obfuscating it for nonprivileged users. This feature enables you to set up policies at the table and column level that provide multiple masking functions, such as obfuscating the first eight digits and displaying the last four digits of an ID or credit card number. Once the policies have been set up, these masks are applied in queries. You can allow certain privileged logins to see the data unmasked.

Backup Encryption

SQL Server has the ability to encrypt the data while creating a backup. By specifying the encryption algorithm and the encryptor (a certificate or asymmetric key) when creating a backup, you can create an encrypted backup file. On-premises and Windows Azure storage locations are supported. In addition, encryption options can be configured for SQL Server Managed Backup to Windows Azure operations, a new feature introduced in SQL Server 2014. To encrypt during backup, you must specify an encryption algorithm, and an encryptor to secure the encryption key. The following are the supported encryption options:

- Encryption algorithm: The supported encryption algorithms are: AES 128, AES 192, AES 256, and Triple DES.
- Encryptor: A certificate or asymmetric key.

Enhanced in SQL Server 2016:

Backup encryption is now supported with compression.

New in 2016: Always Encrypted

SQL Server 2016 introduces several security innovations. Always Encrypted adds a unique capability to have your data encrypted while at rest and in motion with the ability to query that data while it is encrypted. This is optimal for internal compliance, especially in regulated industries or for handling very sensitive data, and is accomplished with minimal overhead. We are calling this capability Transparent Queryable Encryption.

TDE does not prevent a security administrator or DBA from accessing the encrypted data. The new Always Encrypted feature in SQL Server 2016 protects sensitive data stored in a SQL database from DBAs and other high-privileged yet unauthorized users. Always Encrypted transparently encrypts data in an Always Encrypted-enabled client driver, before the encrypted data is uploaded to the database (or transparently decrypts data prior to returning it to the application). SQL Server guarantees that the data and the corresponding keys are never seen in plain text on the server, yet SQL Server can process queries against the encrypted data.

Always Encrypted is the first data platform solution on the market providing queryable encryption. SQL Server 2016 supports deterministic encryption, which allows equality comparisons on encrypted columns. Equality operations include joins, group by and distinct operators. This will allow encryption of sensitive data such as identification or credit card numbers, which are typically only involved in lookup operations. All other operations will fail gracefully as unsupported when executed on encrypted columns. This means, for example, that an application could run a SELECT statement against a particular credit card number, without any credit card numbers being visible to unauthorized users (even users with SA credentials) (Figure 4).

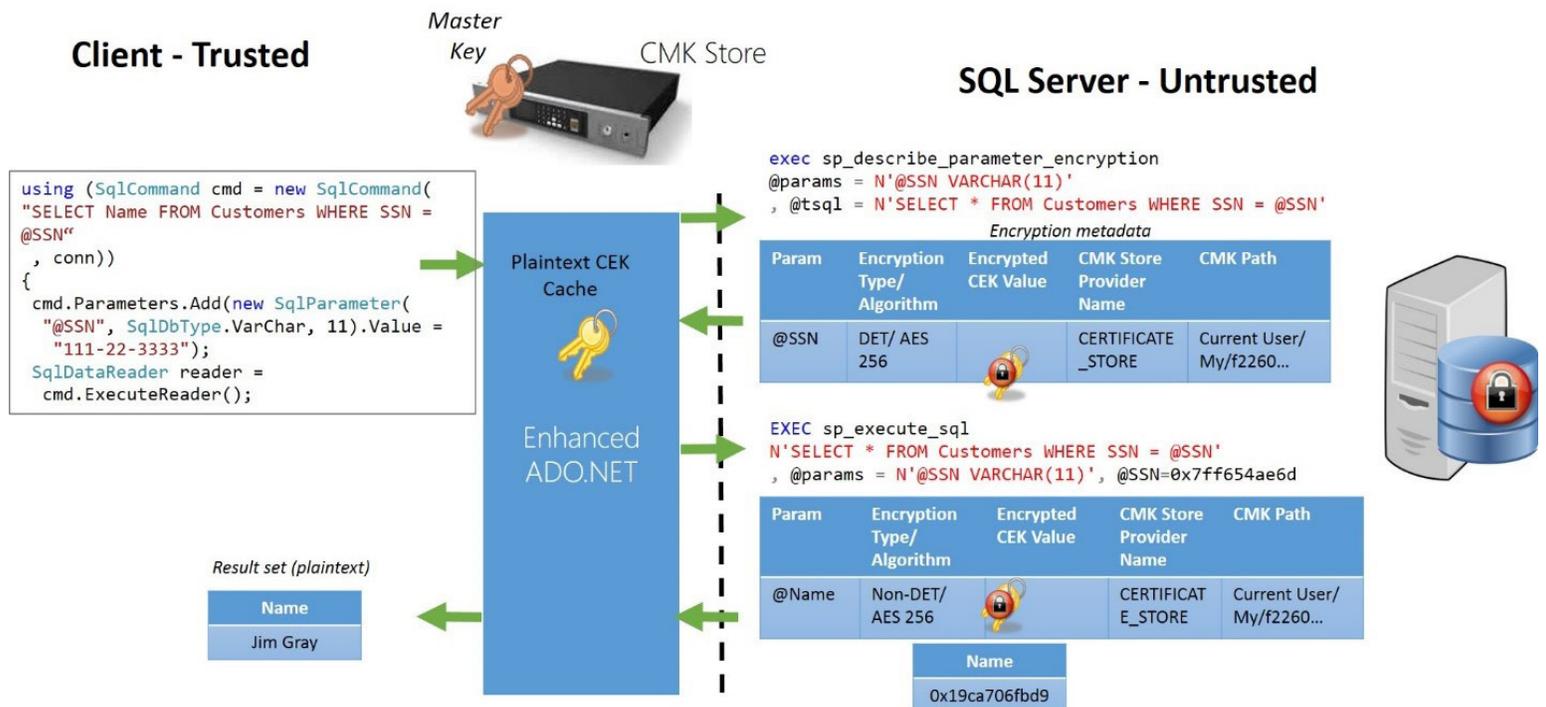


Figure 4: Queryable encryption with Always Encrypted

Using Always Encrypted with client applications requires very few changes either on the client or the server, so the development costs to realize the benefits of this feature are minimal. It does require an Always Encrypted-enabled client driver to sit between the client and the database, though. Whether you are setting up Always Encrypted with a new or existing application, implementation is straightforward using SSMS or SQL Server Data Tools (SSDT) and is supported with tools to make the process as easy as possible.

With new applications, you select the columns to be encrypted and encryption settings, set up the content master key (CMK) and the content encryption key (CEK) using the key setup tool, and identify any impacts on the schema or application queries using the schema analysis tool (SSDT only).

With existing applications, the setup additionally requires encrypting the (previously) plain text data in the selected columns. This can be accomplished in two ways:

4. Creating new, encrypted columns and copying the data from the unencrypted columns, then swapping the old columns for the new in the schema and recreating any dependencies from the old to the new. This process is facilitated by the encryption tool.
5. If you are migrating the database to a new target server, you can use the encryption tool in conjunction with Import/Export to migrate the data into an encrypted column.

New in 2016: Row-Level Security

SQL Server is in the clear leadership position when it comes to security for your mission-critical applications. We are also delivering what you have been asking for some time, Row-Level Security. Row-Level Security provides fine-grained access control over rows in a table based upon conditions you set up.

Using RLS, you can store data for different customers, departments or tenants in the same table, while restricting access to rows based on a query's execution context. RLS works transparently at query time, with no application changes required. It uses a centralized security logic that resides inside the database and is schema-bound to the table it protects, providing greater security. Implementing RLS in the database can greatly reduce client application maintenance and complexity.

SQL Server is in the clear leadership position when it comes to security for your mission-critical applications.

In RLS, you first create a predicate function that encapsulates your access logic. This is simply a user-defined, inline table-valued function that implements the security logic. It can be as simple or complex as you need it to be and it can reference as many tables as needed. The predicate function filters the rows that any given user can access.

For example, imagine a function allows hospital staff to access rows in a patient table only where there is a match between the staff member's assigned hospital wings and the dates that they were assigned to each wing. RLS will allow the hospital to create a security policy that binds the search function to one or more tables. Once bound to the table, all access to the table is routed through the security policy. So a staff member who queries patients would only see those patients who were in her wing during the time she was assigned to that wing.

Built-in tools for enabling compliance: SQL Server audit tools

Database auditing is built into SQL Server to make auditing easy because it is continually available, and to help organizations audit database activities, including database reads, with minimal impact to performance. As compliance policies get increasingly tighter, organizations can use built-in tools such as the following:

- **SQL Server Audit (all editions):** Enables organizations to extend the benefits of SQL Server Audit from Enterprise edition to all SQL Server editions. This extensibility allows for more thorough auditing practices across SQL Server databases, and it enables standardization, better performance and richer features.
- **User-Defined Audit:** Allows the middle-tier application to write custom events into the audit log, which enables more flexibility to store audit information.
- **Audit Filtering:** Provides greater flexibility to filter unwanted events in an audit log.
- **Audit Resilience:** Gives the ability to recover auditing data from temporary file and network issues to help ensure that audit logs are not lost during failover.

Availability

Today's world of business spans geographies and time zones. In this 24/7 environment, there are no longer any acceptable windows for downtime. Simply put, companies need their apps up and running. SQL Server continues to refine its availability features to provide mission-critical uptime, fast failover, improved manageability, and better use of hardware resources.

Already known for rock-solid, reliable performance, SQL Server is getting new features in the coming months that will make it more dependable than ever. AlwaysOn continues to get better and more powerful with every release, and SQL Server 2016 adds significant new enhancements to AlwaysOn. SQL Server 2016 also adds better load balancing, resolves compatibility issues with Distributed Transaction Coordinator (DTC) and SSIS, and provides new features for flexible and efficient backups.

AlwaysOn

Introduced in SQL Server 2012, AlwaysOn maximizes the availability of a set of user databases for an enterprise. SQL Server 2014 expanded on those capabilities by delivering an integrated and enhanced high-availability and disaster recovery solution that provides redundancy within and across datacenters to help enable fast failover of applications during planned and unplanned downtime. AlwaysOn delivers a suite of capabilities rolled into a single solution.

We continue to work to make AlwaysOn better and more powerful with every release, and SQL Server 2016 will deliver even more from this powerful feature. The next version of SQL Server will support up to three synchronous replicas, as well as readable asynchronous replicas. These replicas can exist in different domains and be either on-premises or on Azure virtual machines. We have also added in better load balancing of replicas using a round-robin methodology, and the Distributed Transaction Coordinator allows you to synchronize transactions across different databases, especially beneficial to customers with IBM, Oracle and SQL, for example, who want to replicate transactions across their distributed systems.

For 2016, we are enhancing the high-availability and scalability capabilities of Standard Edition by adding AlwaysOn basic for two-node failover and removing core and memory restrictions. However, all of the in-memory, security and enhanced HA capabilities are only included in Enterprise Edition.

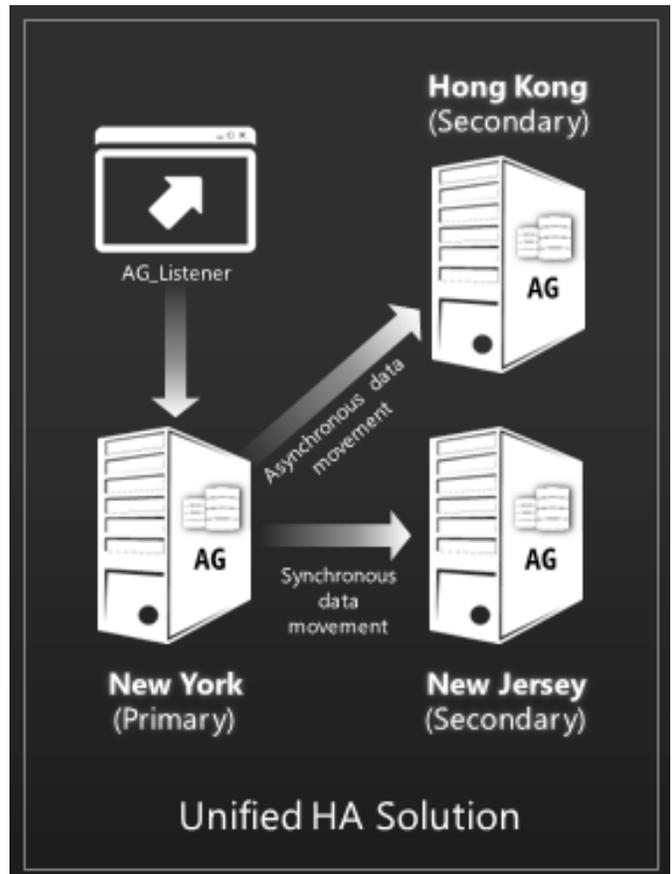


Figure 5: Unified high-availability solution

SQL Server AlwaysOn Availability Groups is a high-availability and disaster recovery solution that provides an enterprise-level alternative to database mirroring. Availability Groups are an integrated set of options that include automatic and manual failover of a group of databases, support for as many as eight secondary replicas ("secondaries"), faster failover for applications, and automatic page repair. Each availability group is a container for a discrete set of user databases known as availability databases that fail over together. An availability group can have many possible failover targets (secondary replicas). Moreover, organizations can easily configure secondary replicas to support read-only access to secondary databases and back up secondary databases. The addition of Availability Groups removes the requirement of shared disk storage such as storage area network (SAN) or network-attached storage (NAS) for deployment of a Failover Cluster Instance.

SQL Server AlwaysOn Failover Cluster Instances enhance SQL Server Failover Clustering and support multisite clustering across subnets, which helps enable failover of SQL Server instances across datacenters. Faster and more predictable failover of instances is another key benefit that helps ensure faster recovery of applications. By supporting Windows Server Cluster Shared Volumes, AlwaysOn further improves use and management of SAN storage through increased resilience of storage failover and avoidance of the drive-letter limitation in SAN.

SQL Server AlwaysOn Multiple, Active Secondaries enables use of as many as eight secondary instances for running report queries (many times faster than replication) and backup operations, even in the presence of network failures, which helps in repurposing idle hardware and improving resource utility. It also helps to dramatically improve performance for both primary and secondary workloads because they are no longer competing for resources.

SQL Server AlwaysOn Availability Groups Listener enables faster failover in client connections for AlwaysOn in scenarios that employ multiple subnets. Now, client applications can achieve failover across multiple subnets (as many as 64) almost as fast as they can achieve failover within a single subnet. Meanwhile, the ability to set the connection from within applications to read-only (instead of read and write) empowers organizations to control the type of workloads that run on their high-availability servers, so they can more efficiently manage their resources.

SQL Server AlwaysOn to Azure Virtual Machine enables organizations to add secondary replicas in an Azure Virtual Machine through the Add Azure Replica Wizard. They can then use this replica for disaster recovery, reporting and backup operations. This configuration can lower capital expenses by eliminating the need to purchase additional hardware for AlwaysOn secondaries.

Enhanced in SQL Server 2016:

Enhanced AlwaysOn

SQL Server 2016 makes major improvements in AlwaysOn in the areas of scalability and manageability. For scalability, SQL Server 2016 adds in load balancing of readable secondaries. This allows you to define one or more groups of readable secondaries to load balance. Connections are assigned round-robin to members of the group.

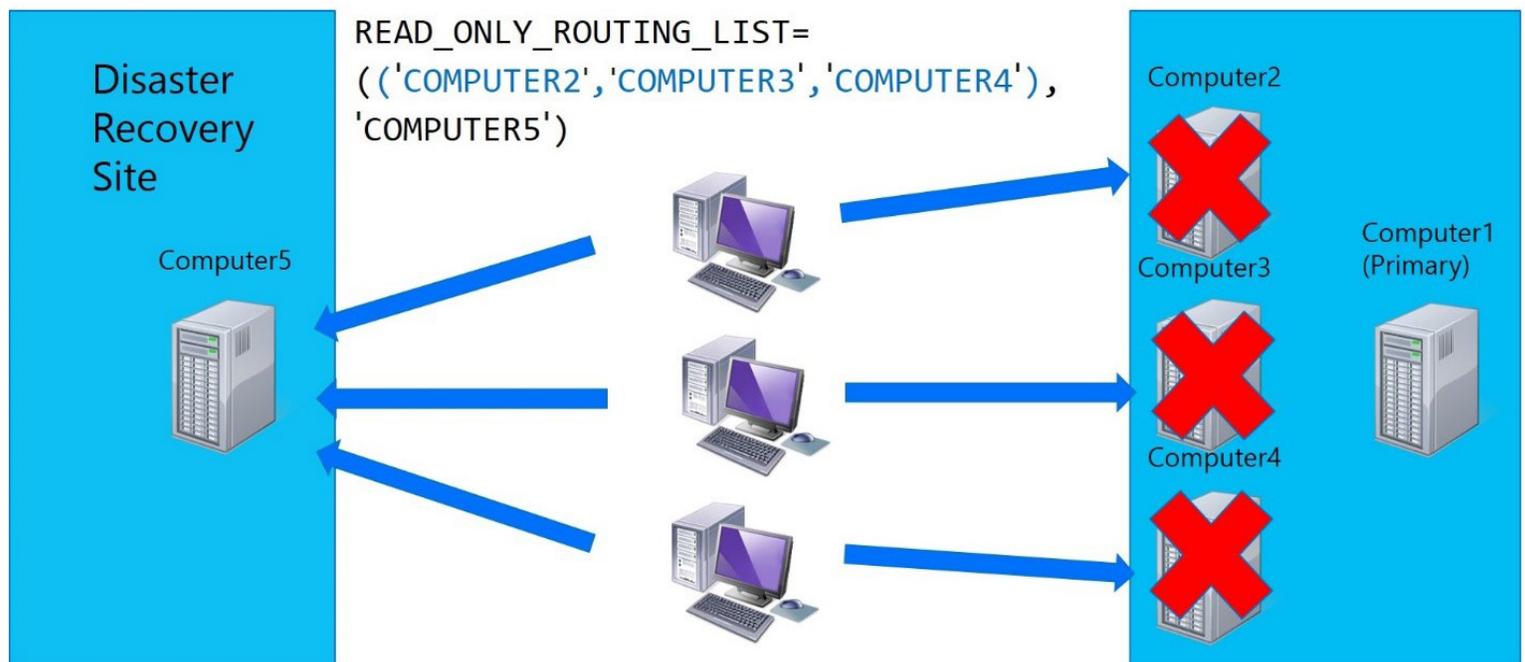


Figure 6: AlwaysOn — readable secondary load balancing

This version also increases the number of auto-failover targets from two to three. Finally, log transport performance has been improved. With In-Memory OLTP and other technologies pushing database speeds, the high-availability pipeline becomes more critical. SQL Server 2016 offers improvements in both throughput and latency.

Manageability has also been improved in several areas with SQL Server 2016, including support for Distributed Transaction Coordinator (DTC) — enrolled transactions for Availability Group (AG) databases. DTC resources are tied to the database instead of the instance so, on failover, the DTC sees the same resource on the new primary, and transaction outcomes can be resolved.

Another manageability improvement is database-level health monitoring. In SQL Server 2014, AG health only monitors the health of the instance. A database can be offline or corrupt, but as long as the instance itself is healthy, SQL Server won't trigger a failover. SQL Server 2016 allows you to optionally change the health monitoring to also consider the health of the databases in the AG.

Group managed service accounts (GMSA) are another manageability improvement. GMSAs are domain-level accounts that are automatically managed. These are similar to the default service account, but with a domain scope. This enables setting permissions for network resources across the AG instances without requiring user accounts.

Scalability

Whether it's running analytics on huge datasets or running a retail site on Black Friday, today's business environment demands scalability. The interaction between SQL Server and Windows Server is an area that can lead to large improvements in scalability, and SQL Server can take advantage of the full feature set in Windows Server for a number of areas.

From advancements in compute, to storage and networking, all have a direct impact on mission-critical SQL Server workloads. With SQL Server and Windows Server, physical processing now scales up to 640 logical processors, and virtual machines scale up to 64 logical processors. SQL Server also utilizes storage spaces and network virtualization to optimize resources in new ways. A number of other tools and techniques can help boost scalability as well.

SQL Server also has several enhancements that can make setting up a private cloud much easier. In addition, SQL Server is integrated into Microsoft's Cloud Platform System, discussed below, which provides new possibilities for scaling a private cloud very quickly.

Whether it's running analytics on huge datasets or running a retail site on Black Friday, today's business environment demands scalability.

Columnstore

The in-memory columnstore index is the standard for storing and querying large data warehousing fact tables. It uses column-based data storage and query processing to achieve up to 10x query performance gains in your data warehouse over traditional row-oriented storage, and up to 10x data compression over the uncompressed data size. This very high level of data compression can reduce your data warehouse storage cost significantly. Plus, for analytics they offer an order of magnitude better performance than a btree index. They are the preferred data storage format for data warehousing and analytics workloads.

Starting with SQL Server 2016 Community Technology Preview 2 (CTP2), you can use columnstore indexes for real-time analytics on your operational workload.

Reasons why columnstore indexes are so fast:

- Columns store values from the same domain and commonly have similar values, which results in high compression rates. This minimizes or eliminates IO bottleneck in your system while reducing the memory footprint significantly.
- High compression rates improve query performance by using a smaller in-memory footprint. In turn, query performance can improve because SQL Server can perform more query and data operations in-memory.
- Batch execution improves query performance, typically 2 to 4x, by processing multiple rows together.
- Queries often select only a few columns from a table, which reduces total I/O from the physical media.

The cloud and mission-critical systems

There is no question the cloud is here to stay, and will only become more central to the way organizations operate in the future. As companies look to move more data to the cloud, SQL Server is evolving with new capabilities to facilitate the move, and keep data safe, no matter where it resided.

New tools in SQL Server and Microsoft Azure make it even easier to scale to the cloud; to build patching, backup and disaster recovery solutions; and to access resources wherever they are — on-premises, private cloud or public cloud.

These tools provide an easy on-ramp to the cloud for on-premises SQL Server databases, enabling customers to use their existing skills to take advantage of Microsoft global datacenters. This is achieved by an architecture that was designed to facilitate hybrid solutions, while also maintaining simplicity in common tasks and providing a set of consistent tools that work in a similar manner whether operating on-premises or in the cloud.

Traditional bare-metal deployment

Despite massive improvements in virtualization technology in the past 10 years, the fact remains that there is still a significant performance penalty for virtualizing certain workloads. Large, complex, and mission-critical online transactional processing systems (OLTP) remain the province of massive, dedicated servers that have the operating system and database platform installed directly “on the metal.”

In situations where the ability to scale up matters and every bit of extra performance counts, organizations need to retain this capability. These situations usually involve organizations that need to extract the most performance possible out of, in many cases, some of the largest server machines money can buy. As a consequence, applications typically have specific server hardware dedicated to their operation.



A key benefit of running significant dedicated hardware resources is that there are many opportunities for advanced physical tuning. The most significant area for a database deployment such as SQL Server is the physical configuration of the storage subsystem. The ability to perform physical tuning is something organizations lose when they move to a purely cloud environment.

Hybrid solutions

The concept of hybrid cloud recognizes that organizations typically have a portfolio of different applications deployed across their business and a breadth of environments with unique requirements. Some applications require detailed and complex hardware configurations that defy deployment into the type of commoditized, “one-size-fits-all” environment offered by cloud computing. Equally, there are workloads in many businesses that are extremely compelling for massive-scale public clouds — it can be economically infeasible to allocate sufficient levels of hardware for applications that experience massive peaks and troughs in demand.

Thus, as organizations look to take advantage of cloud computing, the availability of cloud-enabled database systems will be critical to their success. SQL Server allows the hybrid IT vision to gain leverage and makes it available across the spectrum of deployment approaches that organizations adopt today (Figure 7).

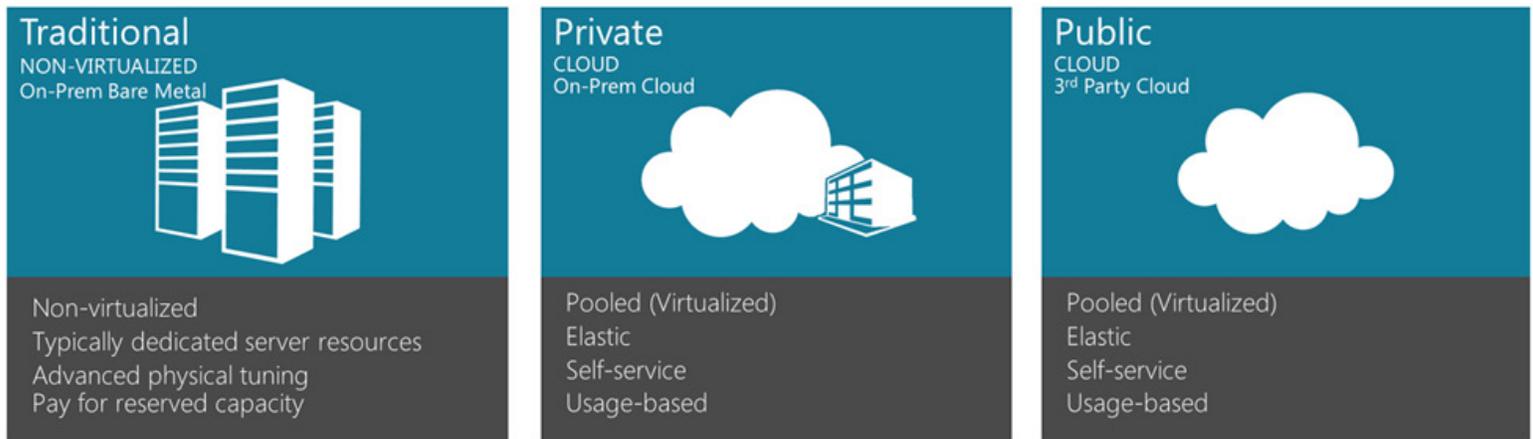


Figure 7: Modern IT departments can meet business needs through a combination of on-premises and cloud-hosted delivery

The Microsoft goal for hybrid cloud is to offer organizations breadth of choice in how and where they run their applications, while at the same ensuring they can use a common set of server products, tools and expertise across a portfolio of solutions (Figure 8).

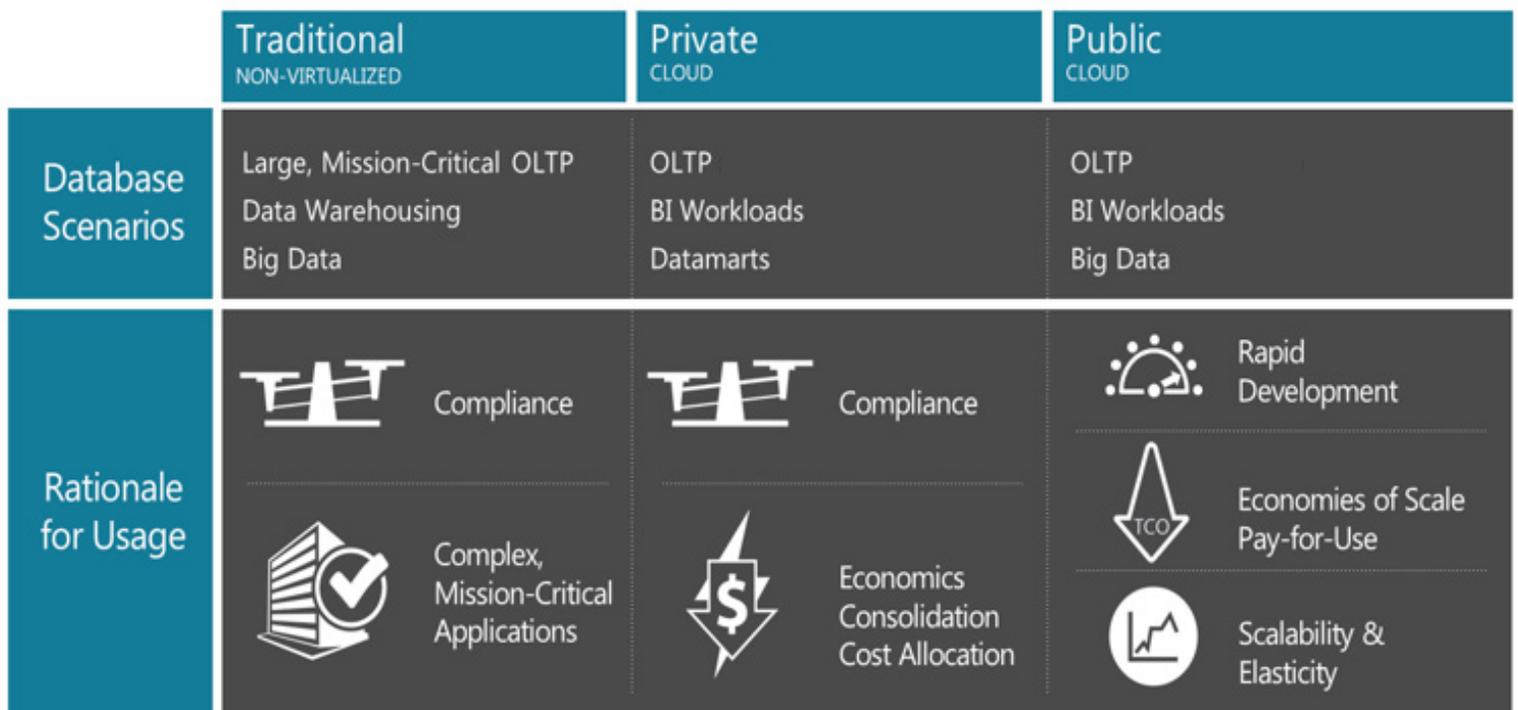


Figure 8: Each approach to database deployment brings unique benefits and challenges — organizations increasingly are moving workloads to the cloud

Cloud deployment: public cloud and private cloud

While a private cloud may have all the characteristics of a public cloud, that does not necessarily mean it must have all of those characteristics. For example, many private clouds do not implement a full chargeback accounting mechanism. Nevertheless, as organizations mature their private cloud strategy, the service and service levels offered by private clouds begin to align more closely with those offered by public cloud providers.

Pooled and virtualized resources

Server virtualization underpins both private and public cloud environments. However, a cloud-based approach to computing requires more than just the mere virtualization of workloads. Many on-premises virtualization environments have specific application targets. Though the environment is virtualized, applications must run on specific, dedicated server hosts. In some cases, this is by technical necessity; in others, because a particular department “owns” that node. A cloud environment is predicated on the pooling of hardware resources, and while virtualization is a key to pooling capacity, it is not enough in and of itself.

Pooling is the mechanism by which resources are aggregated and then made available as a homogenous pool of capacity capable of running any workload. Workloads that run in a pooled cloud environment are agnostic as to the physical hardware on which they are actually deployed.

Because of the advanced physical tuning required, the Tier-1 workloads discussed above are a pooling anti-pattern. For example, a SQL Server workload that requires a particular approach to physical tuning and certain hard drive spindle layouts could be virtualized, but does not lend itself to the use of pooled resources because it has unique resource demands that are unlikely to be demanded by other applications. Put those specific spindle configurations into a pool, and chances are nobody else will want to use them.

Elasticity

Elasticity refers to the ability of the cloud to respond to peaks and troughs in demand. Many business processes are seasonal in nature. For example, during the annual haymaking process, most farmers bring in outside contractors with the necessary machinery to make hay because it is uneconomical to have the requisite large tractors and hay balers lying idle for most of the year. Information technology workloads also are highly seasonal, yet the machinery deployed to support them is typically purchased in sufficient capacity to meet the peak load and “stored in the shed” for the remaining time.

A typical example of a seasonal workload is the sale of tickets for sporting and cultural events. When a large event goes on sale, the number of customers seeking tickets can, in many cases, outstrip supply. Historically, customers would camp all night outside the ticketing office to obtain their tickets. In the online world, this natural queuing mechanism breaks down, and instead prospective event-goers swarm the virtual ticketing office, often overloading it.

Because cloud resources are both generic and pooled, it is easy to justify having spare capacity. Cloud providers, whether public or private, typically try to have a portion of their capacity available at all times to deal with peaks. Public clouds are at a distinct advantage here. Because public clouds operate at massive scale, with thousands of customers accessing their pooled resources, they are able to maintain significantly more absolute headroom than a smaller private cloud: One percent of a 100-server cloud does not permit much of a spike in load, whereas 1 percent of a 10,000-server cloud does. Elasticity is the most difficult cloud characteristic to achieve in a private datacenter because it requires an organization to have capacity lying idle. However, avoiding idleness is usually a key justification for cloud-based deployment in the first place.

Some workloads are not feasible in a private cloud environment. A good test of a cloud’s caliber is to ask the question, “How many times more capacity does the cloud have deployed than my expected elastic demand?” Capacity should be measured in orders of magnitude and not just mere multiples. If you expect to need tens of servers on a burst basis, then look for a cloud that has at least thousands of nodes.

Private cloud

Self-service

Self-service in cloud computing addresses two complementary goals. First, it helps to further drive down the costs of providing the service by reducing or eliminating the labor typically required to provision resources. Second, if done well, it is a measure that benefits users by providing self-service capability.

Cloud consumers are empowered to directly access resources. There is no complicated approval process and no need to wait for the request to become a business priority for IT administrators.

A cloud environment gives users delegated rights to provision resources on demand from the pool. It ensures that users' workloads cannot interfere with other workloads and that users may only provision resources up to the capacity level to which they are entitled (or in the case of a public cloud, the limit their credit limit extends to). Self-service drives business agility by allowing organizations to try new things and reach new markets quickly. Whether in a private cloud inside the enterprise or out in Azure, applications can be taken from development to production much more quickly than through other deployment approaches.

With SQL Server 2016 Stretch Databases, you can stretch an on-premises table to Microsoft Azure transparently for near-infinite capacity with low TCO storage.

New in SQL Server 2016:

Stretch Databases

Data is continuously growing at a high rate, and users generally want to retain all data — including closed business (archive/cold) data — for many possible reasons, such as these:

- Regulatory compliance (for example, taxes)
- Audit (for example, fraud investigation)
- Planning (for example, comparing past results)
- Nature of business (for example, retailer transaction details history)
- Inability to determine with certainty what can be safely deleted (for example, what might a government agency or major institutional investor ask for?)

Traditional archiving solutions typically require third-party software and a completely different data store and application to access. Some archiving solutions depend on backups or offline storage. These may be acceptable for some environments, but many enterprises want their archive stored where the data was born. They also want it accessible using the same application, as needed, without having to wait for the data to be restored or brought online.

With SQL Server 2016 Stretch Databases, you can stretch an on-premises table to Microsoft Azure transparently for near-infinite capacity with low TCO storage. Applications continue to work without code changes; existing database administrator (DBA) skills and processes remain relevant; and developers can continue using current tools and APIs. With Stretch Databases, you do not have the complexity or costs associated with managing external archiving and hardware. Stretch Databases is a unique feature that solves these common problems.

Stretch Databases can help you selectively begin migrating data to the cloud as needed. It can also help enhance existing applications and facilitate DBA tasks, especially on larger databases.

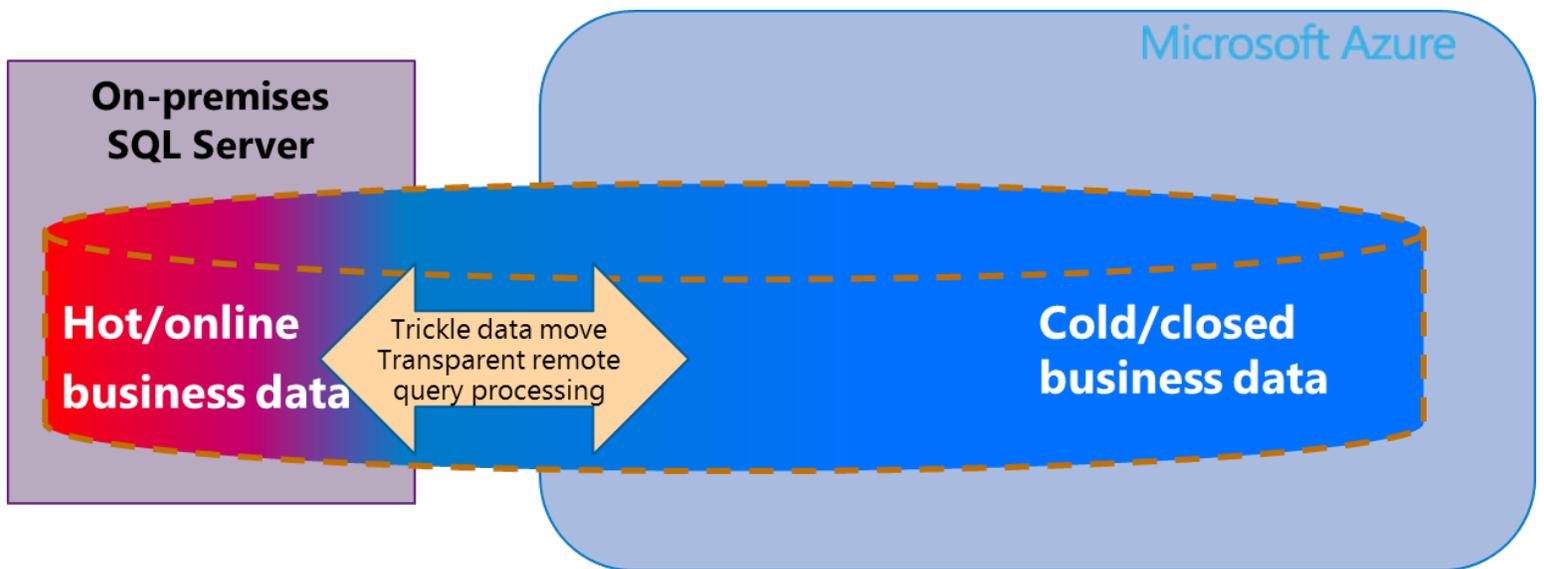


Figure 9: The Stretch Databases feature migrates selected data to the cloud from on-premises SQL Server

Always Encrypted

Always Encrypted adds a unique capability to have your data encrypted while at rest and in motion with the ability to query that data while it is encrypted. This is optimal for internal compliance, especially in regulated industries or for handling very sensitive data, and is accomplished with minimal overhead.

The new Always Encrypted feature in SQL Server 2016 protects sensitive data stored in a SQL database and in the cloud. Always Encrypted transparently encrypts data in an Always Encrypted-enabled client driver, before the encrypted data is uploaded to the database (or transparently decrypts data prior to returning it to the application). SQL Server guarantees that the data and the corresponding keys are never seen in plain text on the server, yet SQL Server can process queries against the encrypted data.

Always Encrypted is the first data platform solution on the market providing queryable encryption. SQL Server 2016 supports deterministic encryption, which allows equality comparisons on encrypted columns. This will allow encryption of sensitive data such as identification or credit card numbers, which are typically only involved in lookup operations. Transparent Data Encryption (TDE) is applied at rest, in transit, and in use by selective encryption of individual columns in a table with keys that are never given to the database system or cloud service provider.

Comparing the cost of mission-critical capabilities

In a technology landscape where organizations expect vendors to intuitively know what “mission critical” means and to provide easy and cost-effective solutions, Microsoft answers these expectations with enterprise-class tools and abilities that are built into database technology — without the need to purchase costly add-ons. SQL Server delivers the mission-critical capabilities required by organizations to compete in a dynamic digital world. The features discussed in this paper are all included in SQL Server Enterprise Edition and don’t require costly options to deliver a complete and modern database solution.

Expecting the fundamentals

Many legacy vendors deliver limited functionality in their premium editions and only provide mission-critical requirements through additional options or feature packs. Examples of additional options are security features, high-availability, performance and spatial capabilities. Organizations have evolved — and it’s no longer optional for vendors to ensure enterprise-grade security features, availability, performance, or support for complex data types, to name a few.

Just like home buyers expect a roof, windows and doors to be included in their purchase, an organization can expect an enterprise-class database to include built-in availability, performance and security features. Figure 11 shows the difference between two similar database solutions, Microsoft SQL Server and Oracle Database, with the budget impact of adding options to arrive at a similar end state.

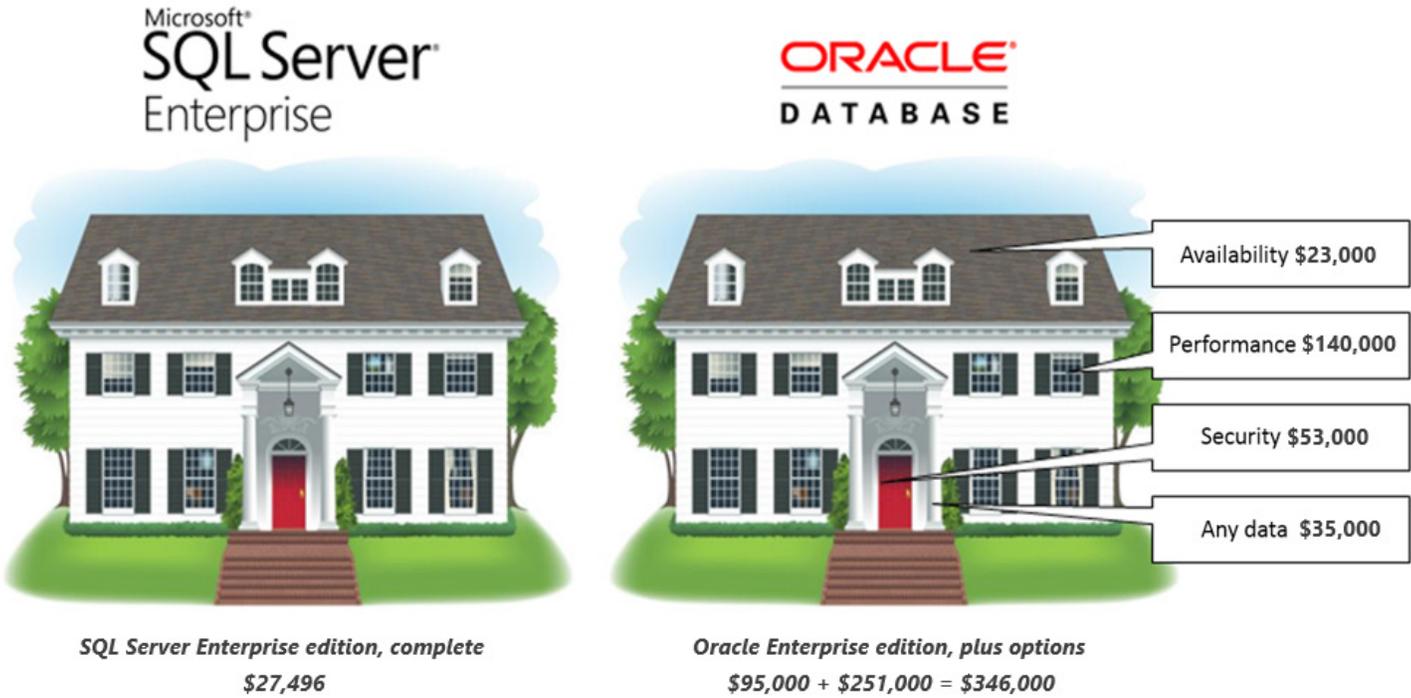


Figure 10: SQL Server and Oracle Database compared

Understanding options

Table 1 highlights the options required across the major database management system (DBMS) vendors to meet mission-critical needs in modern organizations. What used to be optional is more often required by organizations to meet the new standard in mission-critical operations. It is easy to see how achieving mission-critical readiness by adding options or feature packs can dramatically change the total cost of a database solution.

	Microsoft SQL Server	Oracle Database (all options not shown)
Enterprise Edition base license (includes 1-year support)	\$27,496	\$95,000
Data availability	Included	\$11,500 (Active Data Guard) \$23,000 Total
Performance and scale	Included	\$11,500 (Advanced Compression) \$23,000 (Database In-Memory) \$11,500 (Partitioning) \$11,500 (Real Application Testing) \$7,500 (Diagnostics Pack) \$5,000 (Tuning Pack) \$140,000 Total
Enterprise security	Included Free Download	\$15,000 (Advanced Security) \$11,500 (Label Security) \$53,000 Total
Any data, built-in	Included	\$17,500 (Spatial & Graph) \$35,000 Total
Total cost	\$27,496	\$346,000

Table 1: Comparison of mission-critical solutions from Microsoft and Oracle

Note Microsoft prices are based on estimated retail price. All Microsoft and Oracle prices are per-processor (based on a quad core Intel Xeon processor) database pricing for purchases within the United States and are in U.S. dollars. Pricing is based on information available on vendor websites. Oracle prices are based on the Oracle Technology Global Price List, April 9, 2015.

In addition to its mission-critical database functionality, SQL Server Enterprise includes a range of capabilities for data integration, data management, data warehouse, data cleansing, and end-to-end business intelligence. According to the Gartner Magic Quadrant for Data Warehouse and Data Management Solutions for Analytics, Microsoft is positioned as a leader in bringing enterprise data warehouse platform to companies by product innovation such as in-memory columnstore technology. In addition, according to the Gartner Magic Quadrant for Business Intelligence and Analytics Platforms, Microsoft is positioned as a leader in helping organizations to enable broad end-user insight and productivity — balanced with IT oversight through managed self-service business intelligence tools that work both as standalones and within Microsoft SharePoint. SQL Server delivers access to these industry-leading business intelligence capabilities, without requiring costly add-ons. With SQL Server, business intelligence tools are built into the base Enterprise license and are also available in the new Business Intelligence edition. Organizations also can increase cost-savings through built-in data integration, management and cleansing tools. These tools enable data quality managers to easily cleanse and manage data through SQL Server Integration Services, Master Data Management and Data Quality Services. Similar business intelligence and data management tools with other vendors can add up to hundreds of thousands of dollars in additional costs.

³ Gartner, Magic Quadrant for Data Warehouse and Data Management Solutions for Analytics, February 2015

⁴ Gartner, Magic Quadrant for Business Intelligence and Analytics Platforms, February 2015

Conclusion

SQL Server delivers a new standard in enabling mission-critical operations — with true enterprise-class availability, performance and security features built into the solution. Integrated high-availability solutions enable faster failover and more reliable backups — and they are easier to configure, maintain and monitor, which helps organizations reduce the total cost of ownership (TCO). SQL Server also delivers mission-critical performance and scale, with predictable performance across server activities including complex queries, data integration and analysis. Because SQL Server is designed to security standards, it has minimal total surface area and database software that is inherently more secure. Enhanced security, combined with built-in, easy-to-use tools and controlled data access, helps organizations meet strict compliance policies. SQL Server supports complex data types and nontraditional data sources, and it handles them with the same attention — so organizations experience seamless support for a variety of platforms and heterogeneous environments. Finally, SQL Server delivers mission-critical capabilities at low TCO — with full enterprise capabilities that are built into the solution, not provided as costly add-ons. Ultimately, organizations can rely on a comprehensive, integrated solution that helps to contain costs and manage compliance requirements while meeting the demands of the evolving digital world.

More information

For more information about topics discussed in this white paper, see the SQL Server website at <http://www.microsoft.com/en-us/server-cloud/products/sql-server-2016>.