



Mission Critical Performance and Scale with SQL Server and Windows Server
Technical White Paper

Published: September 2013 (updated)

Applies to: Microsoft SQL Server 2014 and SQL Server 2012

Summary: As the volume and complexity of data continue to increase, organizations require a new approach to their mission-critical capabilities. In this white paper, we examine how capabilities built into Microsoft SQL Server—including enterprise-grade security, availability, and performance, along with support for a full range of complex data types—define a “new mission critical” that answers what kinds of capabilities organizations need and expect to compete in a dynamic global landscape. We also compare the cost impact of solutions that offer mission-critical functionality that is built into the core database with solutions that offer separate features organizations can add at additional cost.

Copyright

The information contained in this document represents the current view of Microsoft Corporation 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 Corporation.

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.

© 2013 Microsoft Corporation. All rights reserved.

Microsoft, Active Directory, Azure, Excel, SharePoint, SQL Server, Windows, and Windows Server, are trademarks of the Microsoft group of companies.

All other trademarks are property of their respective owners.

Contents

- The new mission critical 5**
- Organizational requirements 5**
- SQL Server 6**
 - Data availability 6**
 - High availability of mission-critical systems: SQL Server AlwaysOn 6
 - Online database operations 7
 - Predictable, efficient, and flexible data backups 7
 - Better together with Windows Server 8
 - Performance and scale 9**
 - In-memory online transaction processing (OLTP) 9
 - In-memory data warehousing (DW): ColumnStore Index..... 10
 - Buffer pool extension..... 11
 - Query-processing enhancements 11
 - Private cloud 11
 - Tier-1 partitioning: Scale to 15,000 partitions 12
 - Scalable real-world application testing: Distributed Replay 12
 - Reduced database size and increased performance: Data and backup compression 12
 - Proactive troubleshooting and diagnostics: Performance Data Collector and Management Studio 12
 - Better together with Windows Server 13
 - Organizational security and compliance 14**
 - Secure by default: Lowering vulnerability 14
 - Built-in tools for enabling compliance: SQL Server audit enhancements 14
 - Restricted access to data at the row level: Label security 14
 - Controlled access to data in motion: Advanced security 15
 - Controlled access to data for administrators: User-Defined Server Roles 15
 - Backup Encryption 15
 - Controlled access to data across business intelligence tools: Microsoft SharePoint and Microsoft Active Directory 15
 - Virtually any data, built-in 15**
 - Embracing unstructured data through support for complex data types..... 15
 - High availability for unstructured data 16

Seamless connection and analysis of Big Data through Hadoop connectors	16
Seamless connection to a variety of platforms through greater interoperability.....	16
Comparing the cost of mission-critical capabilities	17
Expecting the fundamentals.....	17
Understanding options	17
Conclusion	20

The new mission critical

Data is growing everywhere, and the amplification of data affects every imaginable device, application, and process as our world rapidly evolves into digital and virtual transactions and experiences. It is no longer the case that organizations serve customers during standard business hours in single time zone or geography. Today, services are continually available to customers through a range of operational measures, from Internet presence at the very minimum to tracking complex operations globally for the highest efficiencies and customer satisfaction. The notion of a *maintenance window* is no longer an opportunity for IT to maintain and upgrade systems; these necessary activities are now problems. Consumers, whether at work or at home, simply expect continuous service and access to information and experiences on their terms. Mission-critical response is no longer reserved for the few who can afford to pay astronomical costs to purchase and maintain Tier-1 systems.

Additionally, data amplification extends beyond traditional data types. According to Gartner, the total worldwide volume of data is increasing 59 percent per year. Furthermore, Gartner estimates that 70–85 percent of data is unstructured.¹ The dramatic shift from structured to unstructured and complex data types requires organizations to embrace back-end solutions that support complex data types and nontraditional data sources (such as Big Data) with the same level of support for mission-critical capabilities.

IT organizations currently face the need to balance the impact of data amplification and the requirement for global expansion with conservative budgets and ever-tightening compliance policies. They need to accomplish this feat at higher levels of expected uptime and adherence with regulations that are increasingly strict and diverse across geographies. More than ever, organizations need mission-critical operations that are easy to deploy and are balanced with faster time-to-solution.

Organizational requirements

Our digital age demands a level of response to the demand for mission-critical capabilities that is no longer reserved for the few who can afford costly Tier-1 systems. Organizations require mission-critical capabilities across the following functions:

- **Data availability:** Solutions that are highly available across the globe and that fail over in seconds, deliver reliable backups, and are easy to configure, maintain, and monitor—at a low total cost of ownership (TCO).
- **Performance and scale:** Leading and predictable performance across server activity, including complex queries, data integration, and analysis, and backed by scalable systems to support increasing data.
- **Security:** Clean solutions that include inherently secure database software for reduced risk combined with built-in, easy-to-use tools and controlled access to data that help organizations meet strict compliance policies.
- **Any data, built in:** Support for growing volumes of complex data types and acceptance of nontraditional data sources—with simple support for varied platforms and heterogeneous environments.

¹ Source: Gartner Symposium Presentation, “Information Management Goes ‘Extreme’”: The Biggest Challenges for 21st-Century CIOs, Mark Beyer, October 2011.

SQL Server

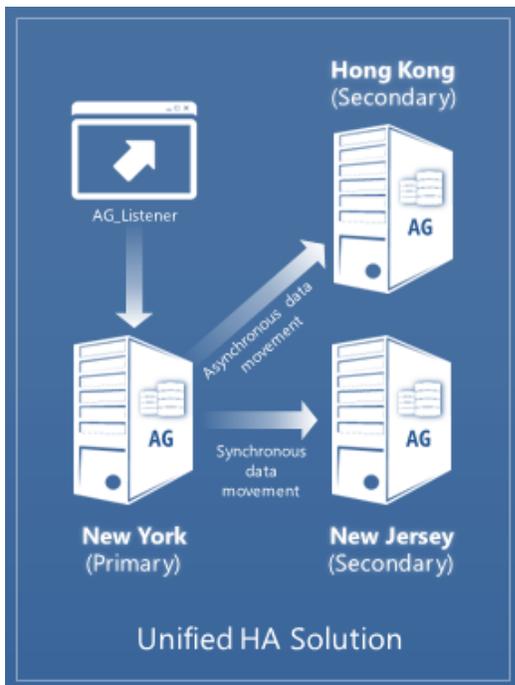
SQL Server sets a new standard for mission-critical platforms by offering organizations the uptime and performance they need at a low TCO, backed by built-in enterprise-level security and support for virtually any type of data.

Data availability

High availability of mission-critical systems: SQL Server AlwaysOn

SQL Server 2014 continues to deliver on its promise of manageability through an incredible user experience with *AlwaysOn*, the enhanced high-availability solution. This integrated high-availability and disaster recovery solution provides redundancy within a data center and across data centers to help enable fast failover of applications during planned and unplanned downtime. AlwaysOn delivers a suite of new capabilities rolled into a single solution.

Figure 1: 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 for backing 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 data centers. 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 Windows Azure Virtual Machine enables organizations to add secondary replicas in a Windows 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.

Online database operations

SQL Server continues to enable organizations to achieve high availability during resource-intensive operations. For example, the ability to rebuild online indexes in a single partition provides partition-level control for users who need continual access to the database. This approach also requires fewer resources (CPU and memory), so it minimizes the impact of rebuilding indexes. Specifically, the ability to manage priorities for locking tables gives organizations greater control over the impact of maintenance operations on running transactions—from table switching to online index rebuild operations—by allowing database administrators to specify whether or not to terminate processes that block their ability to lock tables.

Predictable, efficient, and flexible data backups

Recovery Advisor introduces significant enhancements to user experience through the ways in which database administrators can restore databases by using SQL Server Management Studio. SQL Server offers a variety of backup types, so creating the right recovery sequence for any point in time can get tricky. To help streamline this process, SQL Server Recovery Advisor helps database administrators create a more predictable and optimal restore sequence.

Capabilities include a visual timeline that shows the backup history of the database and the available points in time to which the user can restore the database; algorithms that help streamline the process of identifying the right sets of backup media to get the database back to a specific point in time; and a page restore dialog box in SQL Server Management Studio to do page-level restores of the database.

Backup to Azure

SQL Server enables backup and restore directly to the Windows Azure Blob service. This feature can be used to back up SQL Server databases in an on-premises instance or in an instance of SQL Server running a hosted environment such as Windows Azure Virtual Machine. Backup to the cloud offers benefits such as availability, limitless geographically replicated offsite storage, and ease of transferring data to and from the cloud. Benefits include flexible, reliable, and limitless offsite storage while providing a great mechanism for archiving backups, and virtually no overhead to manage hardware—in addition to cost effectiveness.

Smart Backup

Because it is built on the foundation of backup to Windows Azure, SQL Server Smart Backup provides automatic backup policy to Windows Azure Storage that is context-aware (sensitive to workload and throttling), uses minimum configurations (for settings such as retention period), and is able to manage backups for the entire database instance or particular databases.

Deploy Database to a Windows Azure Virtual Machine

SQL Server provides a wizard for deploying a database to another SQL Server instance running in Windows Azure Virtual Machine. The complexity is fully automated and does not require deep knowledge of Windows Azure. Also, there is no performance overhead for this operation because the major factor in defining deployment time is the size of the database itself.

SQL Server and Windows Azure Storage Integration

SQL Server provides native support for database files stored as Windows Azure blobs. This support is a first step in moving on-premises SQL Server databases to a Windows Azure environment—to deliver an incremental path to a database in the cloud. This feature opens up scenarios such as rapid disaster recovery without data restore, and data encryption in the cloud with encryption keys stored on-premises. Organizations also can move one database at a time without application changes.

Better together with Windows Server

Support for Windows Server Core

SQL Server is supported on Windows Server Core—the Windows Server edition with the smallest footprint. Because Windows Server Core requires less maintenance and fewer OS patches, planned downtime is greatly reduced when you run SQL Server on Windows Server Core. The percentage reduction in patching and OS reboots can be as much as 50 to 60 percent in certain environments, depending on the server roles that are enabled and the type of patches that are applied.²

Support for Windows Server ReFS

SQL Server supports the usage of Windows Server ReFS (Resilient File System) to provide maximum availability, scalability, and integrity. ReFS gives organizations a cost-effective platform that maximizes data availability, scales efficiently to very large data sets across diverse workloads, and guarantees data integrity by providing resilience to corruption (regardless of software or hardware failures).

² Source: “Why Is Server Core Useful,” Microsoft TechNet, <http://technet.microsoft.com/en-us/library/dd184076.aspx>, accessed May 15, 2013.

Faster live migration

Windows Server allows simultaneous migration of as many SQL Server virtual machines as you need, which helps organizations to maintain availability of SQL Server while decreasing planned downtime. Faster live migration also helps organizations to decrease planned downtime by allowing migration of many SQL Server virtual machines (using priority settings) in a clustered environment, and by using as much as 10 GB of network bandwidth.

Live migration for non-clustered virtual machines

Windows Server allows live migration of SQL Server virtual machines in a non-clustered environment both in centrally shared and non-shared virtual machine storage scenarios. This practice helps organizations to reduce the cost and complexity of SQL Server deployments in virtualized environments while maintaining availability during planned downtime.

Cluster-Aware Updating

With Cluster-Aware Updating, organizations can apply updates automatically to the host operating system or other system components in a clustered SQL Server environment while maintaining availability. This approach can significantly help to increase SQL Server availability during the update process in both virtualized and non-virtualized environments.

Dynamic Quorum

Windows Server Failover Clustering Dynamic Quorum enables the SQL Server AlwaysOn cluster to dynamically adjust the number of quorum votes that are required to keep the system running. This adjustment can simplify setup by as much as 80 percent. It also helps increase availability of a SQL Server cluster in failover scenarios in both virtualized and non-virtualized environments—with the ability to recalculate a quorum as needed and still maintain a working cluster.

Performance and scale

In-memory online transaction processing (OLTP)

In-memory technology for SQL Server dramatically improves the throughput and latency of SQL Server online transaction processing (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. In-memory OLTP is designed on the following architectural principles:

- **Optimize for main memory data access:** Storage-optimized engines (such as the current OLTP engine in SQL Server) will retain hot data in a main memory buffer pool based on frequency of access. The data access and modification capabilities, however, are built on the viewpoint that data may be paged in or paged out to disk at any point. With In-memory OLTP, you place tables used in the extreme transaction processing portion of an application into memory-optimized main-memory structures. The remaining application tables, such as reference data details or historical data, are left in traditional storage-optimized structures. This approach lets you optimize hotspots for memory use without having to manage multiple data engines. Main memory structures for In-Memory OLTP eliminate the overhead and indirection of the storage-optimized view while still providing the full atomicity, consistency, isolation, and durability (ACID) properties you expect of a database system.
- **Analyze, Migrate, Report (AMR) Tool:** To identify the appropriate tables to take advantage of In-Memory OLTP, the new AMR (Analyze, Migrate, Report) Tool is integrated into the SQL Server 2014 Management Studio to assist in transition to In-Memory OLTP. It consists of a set of data collectors, leveraging the existing Data Collection framework in SQL Server, as well as a set of Management

Data Warehouse analysis reports. The reports will provide a set of recommended tables or stored procedures for In-memory OLTP to speed the overall performance.

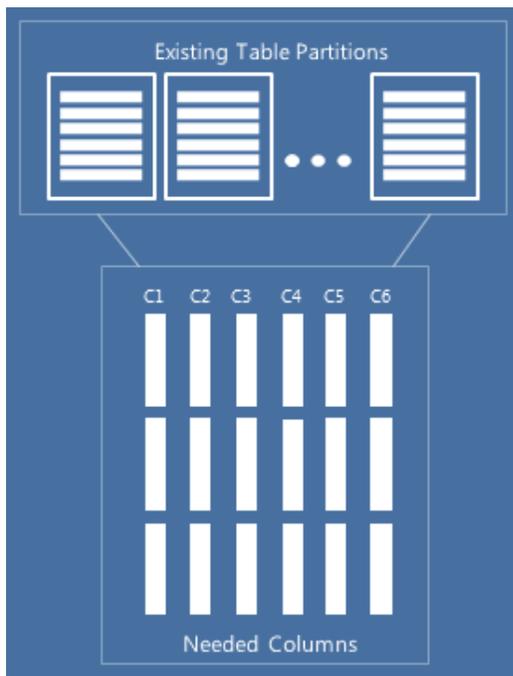
- **Accelerate business-logic processing:** With In-memory OLTP, queries and procedural logic in procedures that are stored in T-SQL are compiled directly into machine code through aggressive optimizations that are applied at compilation time. Consequently, the stored procedure can be executed at the speed of native code.
- **Provide frictionless scale-up:** In-memory OLTP implements a highly scalable concurrency control mechanism and uses a series of lock-free data structures to eliminate traditional locks and latches while guaranteeing the correct transactional semantics that ensure data consistency.
- **Built into SQL Server:** The most impressive thing about In-memory OLTP is that it achieves breakthrough improvement in transactional processing capabilities without requiring a separate data-management product or a new programming model. It is still SQL Server!

In-memory data warehousing (DW): ColumnStore Index

ColumnStore Index introduces in-memory column store technology to the database engine, which makes SQL Server the first of the major general-purpose database systems to have a true column store.

ColumnStore Index combines the VertiPaq technology that was developed in Analysis Services (and is the foundation for PowerPivot) and a new query execution paradigm called *batch processing* to provide significantly greater speed for common data warehouse queries. In test scenarios of star-join and similar queries, organizations have experienced improvements of as much as 100 times faster than in previous versions of SQL Server.

Figure 2: ColumnStore index in-memory columnar storage



ColumnStore Index is built on top of an existing row-level table and provides a view of the data that puts the index on specific columns. It translates the data based on only those required columns and then stores this view, which results in dramatic performance gains. The level of performance gains is dependent on the data and the nature of the query.

Database administrators also can update the Clustered ColumnStore Index online to accommodate real-time data warehouse queries without the need to drop and recreate the index. To save disk space, they can apply a new option called COLUMNSTORE_ARCHIVE for higher compression and storage space savings of as much as 90 percent. Improvements to global batch aggregation also result in better performance and more efficient processing of batch query plans when you use batch mode instead of row mode (which consumes less memory).

Buffer pool extension

SQL Server enables improvement to query performance by allowing the use of non-volatile devices such as solid-state drives (SSDs) to reduce SQL Server memory pressure with no risk of data loss. The configuration is simple and can greatly improve query performance.

Query-processing enhancements

New cardinality estimator

The cardinality estimator improves the querying process and offers the following benefits:

- **Consistent, predictable query performance:** Different operator trees that represent the same relation have the same cardinality estimates.
- **New model for better performance:** Significant changes to the model result in more accurate cardinality estimates and better plan choices.
- **Easier to support:** The querying process separates querying into two steps, *decision-making* and *execution*. It also produces tracing output to facilitate troubleshooting.

Earlier statistics invalidation

Improvements to trigger the automatic update of statistics enable better query results, because statistics invalidation occurs earlier in SQL Server 2014. Faster and more frequent refreshes of statistics are possible because the invalidation threshold has been set to 20 percent of a single partition.

Parallel SELECT INTO

Data-loading is significantly faster because data insertion into a table can occur in parallel through the SELECT INTO operation.

Private cloud

Resource Governor enhancements

Resource Governor enables organizations to further ensure consistent performance for concurrent and mixed workloads across different SQL Server applications and within private clouds. Database administrators can define which workloads can take what percentage of performance on any given CPU, memory, and I/O resource. Resource Governor in SQL Server brings performance scale with the maximum number of resource pools to 64; use of minimum and maximum capacity settings in the CPU, memory, and input/output operations per second (IOPS); and affinity of resource pools with CPU schedulers and Non-Uniform Memory Access (NUMA) nodes. Governance of I/O resources allows administrators to control physical I/O for users by adding a setting for maximum and minimum IOPS per volume to Resource Governor resource pools.

Sysprep for SQL Server

SQL Server supports the preparation of virtual machine templates through SQL Server Sysprep. Administrators can prepare images with the desired features and then deploy them later in private and public cloud environments. SQL Server Sysprep supports SQL Server Database Engine, SQL Server Reporting Services, SQL Server Analysis Services, SQL Server Integration Services, and shared features. With the addition of cluster support, SQL Server Sysprep can be used in a wider variety of image-preparation scenarios.

Tier-1 partitioning: Scale to 15,000 partitions

SQL Server supports as many as 15,000 table partitions. This support enables large *sliding window* scenarios, which means that applications such as SAP, that take tens of thousands of snapshots of data in daily or hourly partitions, can significantly extend the length of time data is held before it's "pushed out" to allow for new data to enter—generally making it easier to manage these large amounts of data. This capability also helps administrators streamline maintenance of large sets of data within file groups that need data switched in and out according to the needs of the data warehouse.

Scalable real-world application testing: Distributed Replay

Organizations need a way to apply real-world application loads to their applications within test environments. Previously, they could use SQL Server Profiler, which only allowed simulation of a workload from a single computer. This limitation made it difficult to test large-scale workload simulation.

Distributed Replay helps organizations to simplify application testing and minimize errors with application changes, configuration changes, and upgrades. This multithreaded replay utility enables simulation to test production workload scenarios after upgrades or configuration changes—ultimately leading to protected performance during changes. Additionally, integration with SQL Server Upgrade Assistant can help organizations assess the impact of future SQL Server upgrades.

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.

Additionally, 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.

Proactive troubleshooting and diagnostics: Performance Data Collector and Management Studio

Organizations want to proactively manage the health of their systems and quality of queries across their environments to ensure the best possible performance. SQL Server delivers a suite of diagnostics and tuning tools built-in and at no extra cost. Performance Data Collector allows administrators to view SQL Server diagnostics from performance counters, dynamic management views, SQL Trace, and other sources for baseline and historical comparisons. They can view performance data with built-in reports on topics such as server activity, disk usage, and query activity. Additionally, SQL Server Profiler can capture server events for real-time diagnosis, and it can correlate traces with performance counters to analyze events and diagnose issues. Dynamic Management Views and functions that relay server state

information help IT administrators to monitor the health of server instances, diagnose problems, and tune performance. The Database Engine Tuning Advisor helps administrators select and create an optimal set of indexes, indexed views, and partitions without requiring an expert understanding of the structure of the database or the internal workings of SQL Server. Simply select the databases to tune, and Database Engine Tuning Advisor generates indexing and partitioning recommendations.

Better together with Windows Server

High-capacity virtual machines

High-capacity virtual processors and memory in Windows Server enable organizations to deploy mission-critical workloads by using SQL Server in a virtualized environment. A SQL Server virtual machine can use as many as 64 virtual processors and 1 terabyte of memory. In addition, support for as many as 640 logical processors and four terabytes of memory enables deployment of mission-critical SQL Server workloads in a non-virtualized environment.

High-density virtual machine cluster

When it is used with Windows Server, SQL Server can achieve greater cluster density for deployment in virtualized environments, and allows as many as 8,000 SQL Server virtual machines per cluster.

High-scale cluster

Windows Server 2012 enables cluster scalability by supporting SQL Server clusters that have as many as 64 nodes—which is four times the number supported by the previous version of Windows Server. This increased capacity provides a range of benefits that include enhanced scalability, improved configuration and management, and ease of maintenance for large SQL Server clusters in virtualized and non-virtualized environments.

Network quality of service

Quality of service (QoS) in Windows Server enables administrators to enforce network bandwidth into a network adapter for SQL Server services such as virtual machines, storage, Live Migration, and Cluster Shared Volumes. This capability can help organizations to lower capital and operating expenses by converging network traffic onto a single network adapter.

Network Interface Card teaming

By using Windows Server Network Interface Card (NIC) teaming to configure multiple NICs for load-balancing, administrators can increase bandwidth for SQL Server network traffic. In addition, configuring multiple NICs for hardware failover helps organizations to maintain SQL Server availability in virtualized and non-virtualized environments.

Support for Server Message Block

Running SQL Server on Windows Server enables SQL Server to take advantage of Server Message Block (SMB) file server features. SQL Server can store data files in remote shared folders that use SMB Direct and SMB Multichannel on industry-standard network adapters. Ultimately, this capability can result in significant storage benefits that include reduced cost, higher availability, and increased performance for SQL Server deployments in virtualized and non-virtualized environments.

Support for Fibre Channel in Virtual Machines

SQL Server virtual machines can connect directly to Fibre Channel to support N_Port ID Virtualization (NPIV), virtual SAN, and multipath IO (MPIO) to ensure continuous connectivity. This improvement helps

to increase storage capacity, storage compatibility, and overall performance for SQL Server deployments in virtualized environments.

Storage Pools

Organizations have an opportunity to lower costs by using industry-standard storage in SQL Server deployments for non-virtualized environments—and in some cases, even replace expensive SAN solutions. Storage Pools can improve the flexibility of SQL Server storage with resilient storage (mirroring and parity) and multitenancy isolation (ACLs).

Organizational security and compliance

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 company-wide 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).³

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

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 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.

Restricted access to data at the row level: Label security

Many industries and government agencies require the ability to restrict access to data at the row and cell levels, to ensure that only authorized people have access to the data. With SQL Server, organizations can enact this detailed approach to security by using the Microsoft toolkit SQL Server Label Security at no additional cost. Organizations can use this toolkit to support critical scenarios or niches within their database environments—without having to purchase costly add-ons. The toolkit, which is developed and maintained by Microsoft engineers, is available on <http://sqlserverlst.codeplex.com>.

³ Source: Information Technology Intelligence Corp. (ITIC), SQL Server Delivers Industry-Leading Security, September 2012.

Controlled access to data in motion: Advanced security

SQL Server allows organizations to encrypt data when it is stored on a disk and decrypt it when it is read into memory. 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. Encryption also protects database backups by automatically encrypting the backups.

Additionally, extensible key management works with transparent data encryption 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.

Controlled access to data for administrators: User-Defined Server Roles

User-Defined Server Roles increases flexibility and manageability, and it facilitates compliance through clearer separation of duties. It allows creation of server roles to suit different organizations that separate administrative duties according to roles. Roles also can be nested to allow more flexibility in mapping to hierarchical structures in organizations. User-Defined Server Roles also helps prevent organizations from using sysadmin for database administration. For example, a special database administration role can be created for common database administration without the ability to access user data.

Backup Encryption

Encryption is a straightforward way to increase security of the data. Encryption of backups is the way to increase security of backups stored separate from the instance, and probably in the different environment. Separate encryption settings for backup allow configuring backup encryption differently from database encryption and benefit from the both. In case database is not encrypted backup compression might be effectively used together with encryption improving as security, same as storage and transfer costs.

Controlled access to data across business intelligence tools: Microsoft SharePoint and Microsoft Active Directory

As Microsoft continues to deliver business intelligence tools that are used by a broader set of users, security concerns also increase because of broader implications if security is compromised. SQL Server helps organizations secure end-user data analytics with built-in IT controls, including new SharePoint and Active Directory security models for end-user reports that are published and shared in SharePoint. Enhanced security models provide control at row and column levels.

Virtually any data, built-in

Embracing unstructured data through support for complex data types

With the dramatic shift from structured to complex data types, the new mission critical requires support for complex data types to be built-in, not added on at extra cost and effort. SQL Server supports a growing number of types and volumes of complex data with FILESTREAM, Remote Blob Storage, and Spatial support—enhancements on top of the already robust and built-in foundation that extend beyond relational capabilities. This support allows organizations to build rich and innovative applications without paying extra premiums.

SQL Server FileTable builds on FILESTREAM to bring Win32 namespace support and application compatibility to the file data stored in SQL Server. Countless applications maintain their data in two

“worlds:” *unstructured* (documents, media files, and other unstructured data in file servers) and *structured* (related, structured metadata in relational systems). FileTable lowers the barrier to entry for organizations who have files on servers that currently run Win32 applications, while reducing the effort caused by maintaining two disparate systems and keeping them in sync.

High availability for unstructured data

With SQL Server, complex data types are handled with the same attention as common data types. Organizations can use FILESTREAM to store and manage complex data in a variety of ways, as if it were part of the database. Additionally, with SQL Server, organizations can enjoy the high-availability benefits of AlwaysOn for complex data managed through FILESTREAM—even when they take advantage of Remote Blob Storage and FileTable.

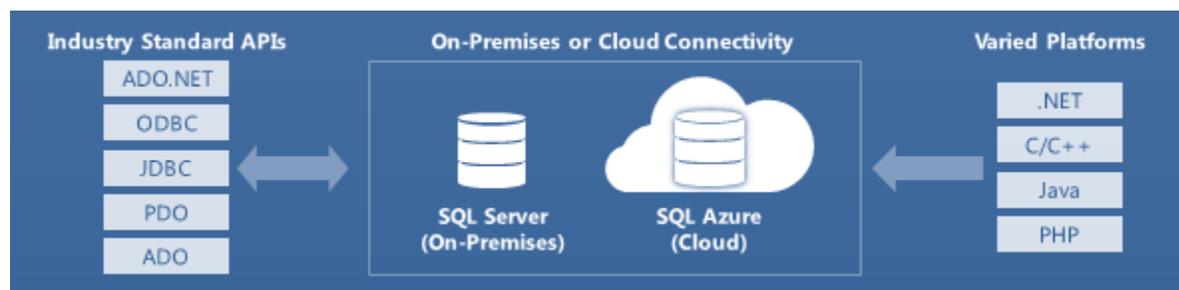
Seamless connection and analysis of Big Data through Hadoop connectors

Hadoop connectors for SQL Server and SQL Server Parallel Data Warehouse are available for download to organizations that have licenses for SQL Server and Parallel Data Warehouse. These connectors enable bidirectional data movement across SQL Server and Hadoop, so users can work effectively with both structured and unstructured data. Additionally, people can use the leading business intelligence platform from Microsoft to perform analysis on Hadoop data sets. Users can gain access to and create mashups of Hadoop data sets by using familiar productivity tools such as Microsoft Excel and award-winning business intelligence clients such as PowerPivot and Power View to perform analysis in an immersive and interactive way.

Seamless connection to a variety of platforms through greater interoperability

Realistically, many organizations own and manage heterogeneous infrastructures. They need interoperable tools and systems to maximize their existing investments and evolve on their own terms. SQL Server helps organizations extend heterogeneous environments by connecting to SQL Server and Windows Azure SQL Database applications, and provides support for additional industry-standard APIs across a variety of platforms. By using SQL Server, organizations can maximize existing investments while they modernize legacy mission-critical workloads on the SQL Server data platform.

Figure 3: Interoperability across platforms and standards



Interoperability with SQL Server includes the following:

- **Microsoft Driver for PHP for SQL Server:** Designed to enable reliable, scalable integration with SQL Server for PHP applications that are deployed on the Windows platform.
- **Connectivity for Java:** Provides secure and highly available connectivity from Java applications to SQL Server for enterprise-level organizations.
- **Microsoft JDBC Driver for Linux and UNIX:** Provides connectivity to Linux and UNIX, which allows organizations that run applications and workloads on legacy platforms to more easily migrate to SQL Server with minimal effort.

Comparing the cost of mission-critical capabilities

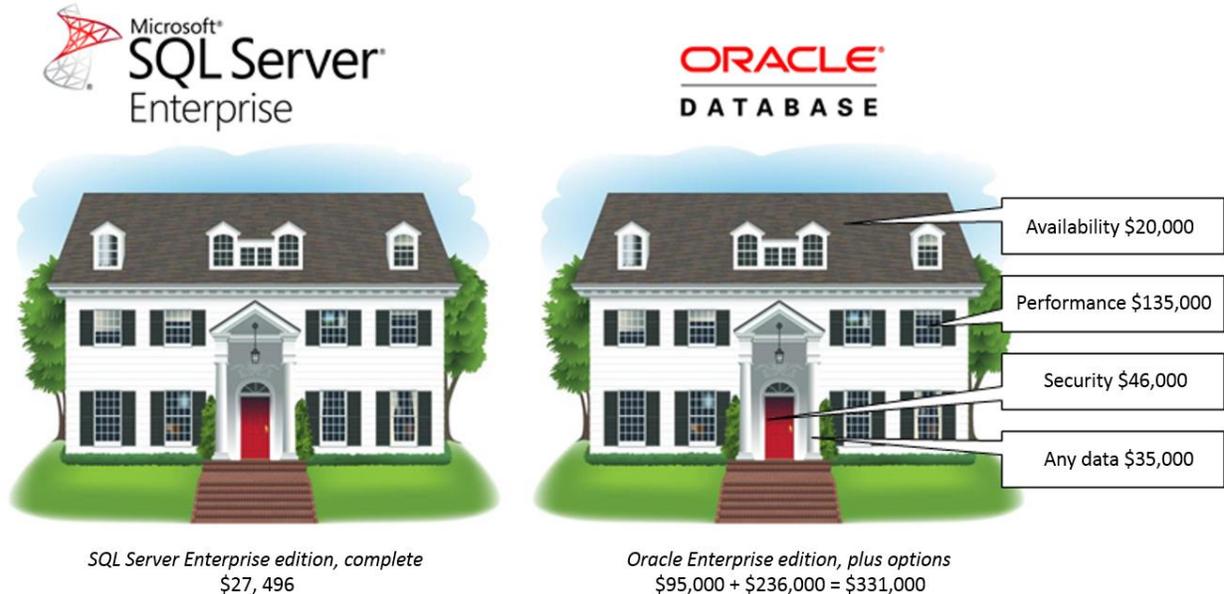
In a technology landscape where organizations expect vendors to intuitively know what “mission critical” means and to provide easy and cost-effective solutions, Microsoft answers these expectations with enterprise-class tools and abilities that are built into database technology—without the need to purchase 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 they 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 4 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 4: 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.

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

	Microsoft SQL Server	Oracle Database (all options not shown)
Enterprise Edition base license (includes 1-year support)	\$27,496	\$95,000
Data availability	Included	\$10,000 (Active Data Guard) \$20,000 Total
Performance and scale	Included	\$11,500 (Advanced Compression) \$23,000 (In-Memory Cache) \$11,500 (Partitioning) \$11,500 (Real Application Testing) \$5,000 (Diagnostics Pack) \$5,000 (Tuning Pack) \$135,000 Total
Enterprise security	Included Free Download	\$11,500 (Advanced Security) \$11,500 (Label Security) \$46,000 Total
Any data, built-in	Included	\$17,500 (Spatial & Graph) \$35,000 Total
Total cost	\$27,496	\$331,000

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 United States dollars. Pricing is based on information available on vendor websites. Oracle prices are based on the Oracle Technology Global Price List, March 15, 2013. IBM prices may vary.

In addition to its mission-critical database functionality, SQL Server Enterprise includes a range of capabilities for data integration, data management, data cleansing, and end-to-end business intelligence. According to the Gartner Magic Quadrant for Business Intelligence and Analytics Platforms,⁴ 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 standalone and within Microsoft SharePoint. SQL Server delivers access to these industry-leading business intelligence capabilities without requiring costly add-ons. With SQL Server, the 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.

⁴ Source: Gartner, Magic Quadrant for Business Intelligence and Analytics Platforms, February 5, 2013.

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 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* non-traditional 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.

For more information, see the SQL Server website at <http://www.microsoft.com/sqlserver/>.