Hyperscale Cloud

Technical White Paper

Published: May 2015

Applies to: SQL Server 2016 CTP2, SQL Server 2014, and Microsoft Azure

Summary: With cloud computing comes a new paradigm shift as organizations continue to realize the potential cost benefits of running their database applications in the cloud with greater scale and flexibility. Microsoft SQL Server is built for cloud integration—your organization can easily deploy SQL Server in a private cloud, hybrid cloud, or public cloud, and can use familiar tools for development and management. In the public cloud, you can run SQL Server in a Microsoft Azure Virtual Machine. Microsoft SQL Server in an Azure Virtual Machine provides full feature parity with on-premises SQL Server. The ability to seamlessly cross these logical boundaries and integrate operations with data and services from virtually anywhere is one of the most exciting and impactful features of SQL Server 2016 and Microsoft Azure.

Copyright

Contents

SQL Server evolution

For many years, Microsoft SQL Server has been moving toward easier integration with outside data sources while simplifying management and administration. With the recent development of Microsoft cloud services and the shift toward larger and more varied data sources, the cloud's capabilities have grown at a dramatically increasing rate. Since SQL Server 2012, cloud strategy has been a central focus of development. SQL Server 2014 added significant new cloud features. These features have been enhanced with SQL Server 2016. *Hyperscale cloud* is a term that resonates with SQL Server developers, administrators, and business leaders because of the benefits that hybrid cloud delivers. Analysts have also recognized the dedication of SQL Server to cloud development and its ability to execute on that vision, as illustrated by the Gartner Magic Quadrant for Infrastructure as a Service (Figure 1).

Figure 1: Gartner Magic Quadrant for infrastructure as a service[1]

[1] Gartner, http://aws.amazon.com/resources/gartner-mq-2014-learn-more/, May 2014.

In addition, SQL Server has consistently added innovative features over the last 15 years (Figure 2).

Figure 2: Major SQL Server functionalities added across releases

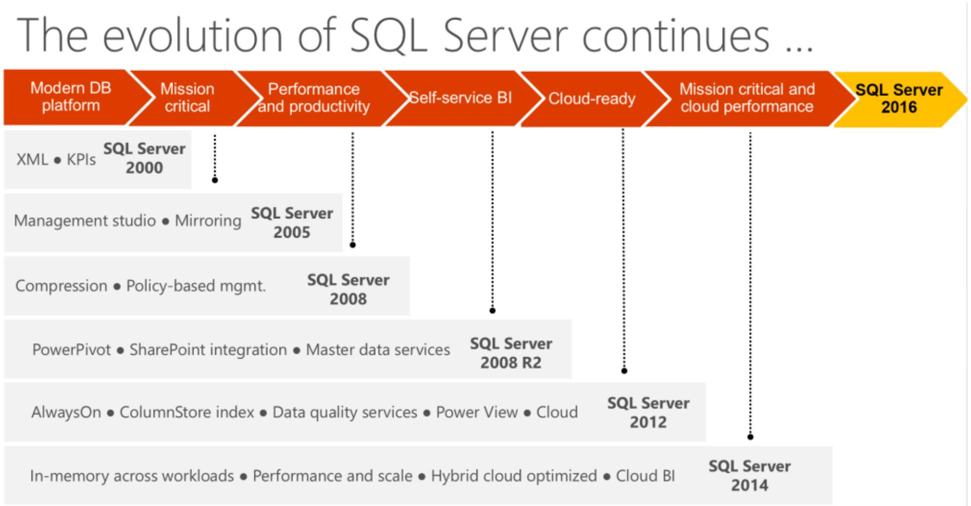

SQL Server 2016 adds to this functionality with:

- Stretch Databases
- Enhanced backups to the cloud
- Enhanced migration to the cloud
- Enhanced availability with cloud replicas
- Enhanced consistency in working with SQL Server on-premises, SQL Server on Azure Virtual Machines, and Azure SQL Database

Hyperscale cloud with SQL Server

SQL Server 2014 and SQL Server 2016 were designed to work in a hybrid environment. 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 data centers. 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.

Hybrid solutions

In almost any application scenario, from the largest public websites to small departmental applications, the vast majority of applications rely on some sort of database management system. The relational database is so ubiquitous that many developers and IT professionals have become desensitized to it—considering it simply part of the stack for a modern application. As organizations look to take advantage of cloud computing, the availability of cloud-enabled database systems will be critical to their success.

This white paper explains the Microsoft vision of relational database management systems in the context of cloud computing. It is a hybrid IT vision that gains leverage from the industry-standard Microsoft SQL

Server technology set and makes it available across the spectrum of deployment approaches that organizations adopt today (Figure 3).

Figure 3: Modern IT departments can meet business needs through a combination of on-premises and cloud-hosted delivery

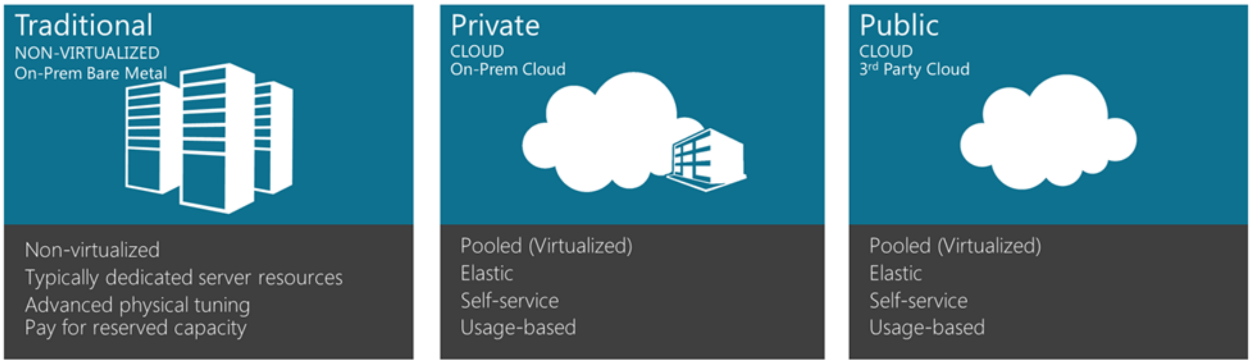

The concept of *hybrid cloud* recognizes that organizations typically have a portfolio of different applications deployed across their business and a breadth of environments with unique requirements. Some applications require detailed and complex hardware configurations that defy deployment into the 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. 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 4).

Figure 4: Each approach to database deployment brings unique benefits and challenges—organizations increasingly are moving workloads to the cloud

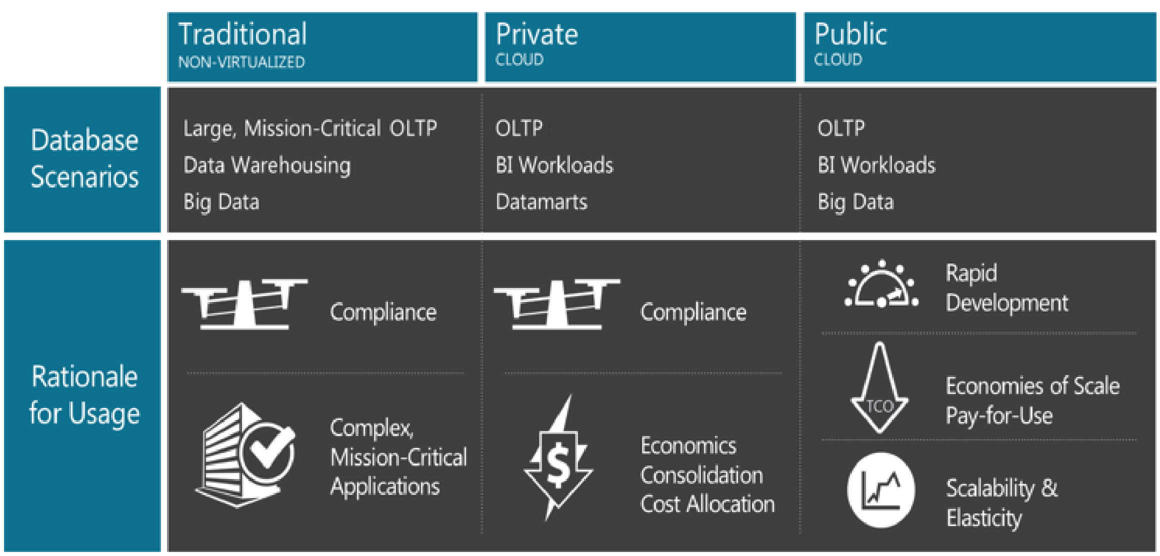

Traditional bare-metal deployment

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."

Non-virtualized, dedicated hardware

For most workloads, virtualization is an ideal approach because it delivers significant benefits in total cost of ownership (TCO). However, in situations where the ability to scale up matters and every bit of extra performance counts, organizations need to run workloads on the metal. 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.

Physical tuning

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 cloud environment.

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 one percent of a 10,000-server cloud does. Elasticity is the most difficult cloud characteristic to achieve in a private data center 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 (such as the ticketing example above) 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.

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.

Usage-based models

Most shared IT environments suffer from the “tragedy of the commons[2]”—if IT capacity is “free” at the margin, there is no incentive for conservation by any one consumer despite this being in the interest of all consumers collectively. Consumers are used to paying on a per-unit basis for other resources such as water, gas, and electricity. The pay-per-use model offered by cloud computing provides incentives to turn off capacity that is not being used.

[2] "The Tragedy of the Commons." *Science* 162 (3859): 1243–1248. 1968. http://www.sciencemag.org/cgi/reprint/162/3859/1243.pdf

Public cloud vendors need to charge for their services, so those environments will always be metered and billed. In private clouds, the situation varies. Implementing a chargeback model is complex, particularly if the business does not have existing accounting systems in place to support it, but there are significant benefits. The goal of pay-per-use in a private cloud environment is to drive user behavior and to ensure that cloud resources are treated as scarce and conserved where possible. Quotas and other resource allocation mechanisms may be more appropriate for some private cloud environments.

One challenge for public and private cloud operators alike is “which meter to use.” What should be counted to determine charges? The metric needs to be well correlated to the actual cost of providing the service but also remain sufficiently simple so that cloud consumers can understand it. It makes little sense to measure “query hours” if cloud consumers do not know how the expected number of required query hours are derived for their particular application.

The charging model is the mechanism by which a cloud provider signals efficiency. In a cloud environment, where layers of virtualization obscure every real resource, organizations should architect their applications to be cost-optimized as a primary consideration.

Compliance

Some application scenarios require compliance with specific enterprise or industry-standard policies. These policies typically relate to security, systems management, and legal matters. Policies range from the simple (which antivirus software to install on servers) to the complex (Information Security Management Systems standards such as ISO/IEC 27001).

The more control an organization has over the entire systems stack, the more amenable that system is to compliance with all imaginable policies and requirements. An enterprise policy requiring air gap deployment is unsuitable for deployment in a public cloud, and equally, an isolated private cloud cannot coexist on pooled hardware with a private cloud connected to the Internet. The more onerous the compliance requirements, the more likely they are to require a dedicated environment running within the complete control of the enterprise that deploys them.

The major public cloud vendors have moved quickly to audit and certify their systems against industry-standard frameworks. For many small and midsized organizations, achieving compliance can be too costly. For these smaller organizations, public cloud computing actually presents new opportunities to deploy applications into a certified environment without the cost of implementing those standards within their own data centers.

Some compliance challenges remain insurmountable for public cloud computing. Organizations that require complete jurisdiction over their systems often need to ensure that data is located only within a certain country and that systems are accessible only by their own staff. For these organizations, the use of dedicated systems or private cloud environments remains the only feasible solution.

For more information about the compliance of Azure with various industry standards, see http://www.windowsazure.com/en-us/support/trust-center/compliance/.

SQL Server, for example, has gone through rigorous compliance exercises since SQL Server 2005 to achieve Common Criteria certification. This certification is formally recognized by the governments of 26 countries that have signed the Common Criteria Recognition Arrangement (CCRA) and by as many as 40 more governments on a product-by-product basis. The CCRA is more than just the concise definitions of security functionalities and assurance requirements. It also is a precise evaluation process defined in the Common Evaluation Methodology document. In addition, it is a formal and approved evaluation scheme for each nation performing Common Criteria evaluations, as well as a government certification based on the government working with a private evaluation lab certified in that country. For more information about the compliance of SQL Server, see http://www.microsoft.com/sqlserver.

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:

- 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.

For these users, the only option is to keep all data in the production database. That results in other issues, some of which become increasingly difficult, if not impractical, to solve. Those issues include:

- Maintenance operations become more difficult. Consider the cost of re-indexing a 1-billion row, 1 TB table.
- Restoring a 5 TB database could take an entire day or more, even though much of this data may be cold data that might not need to be brought online quickly.
- Tiered storage solutions (such as SSD + SAS + SATA) can save significant costs compared to physical disks, but often save little on overall TCO because they usually reside on the same SAN infrastructure (rack, trays, interconnects, power, management software) and require the same licenses, maintenance contracts, and administration effort.

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 (Figure 5). It can also help enhance existing applications and facilitate DBA tasks, especially on larger databases.

Figure 5: The Stretch Databases feature migrates selected data to the cloud from on-premises SQL Server

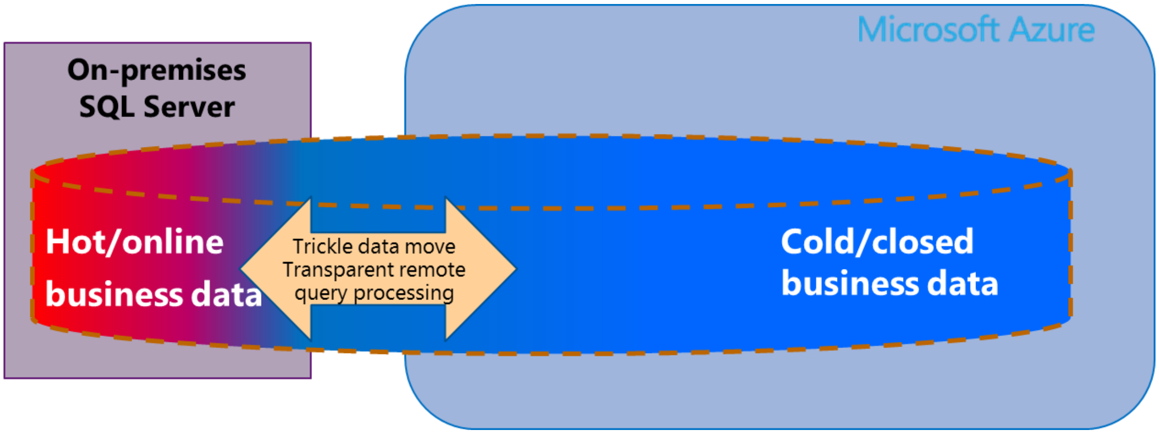

Cloud on your terms

You can choose to stretch specific SQL Server database tables to Azure. You have the option to move the entire table to Azure, or you can set it up to only store the old or cold data there. Moving the entire table would be a good choice if the table only stored archive data. Moving cold data would be a good choice if you had a single table that contained both hot (current) and cold (old) data. Once this is set up, Stretch Databases uses trickle data migration to move the data to Azure. It can also be moved back to the on-premises SQL Server database, so you can achieve transparent bi-directional data movement without any changes to your stored procedures, functions, other application code, or user access control. You are in control of what is stored in the cloud and what is stored on-premises.

Additionally, the data is kept secure at all times during this process. 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.

Enhance existing applications

By choosing which data resides where, SQL Server 2016 Stretch Databases can help you enhance existing applications. For example, if you have company policies, SLAs, or TCO goals that may impact how data is archived, Stretch Databases can help meet those goals or comply with those policies. There is also a limited amount of computation that can occur on Azure, and Stretch Databases automatically decides that the compute should happen at the location where it makes the most sense. Further, Stretch Databases will transparently scale out by creating additional shards on Azure as needed.

Empower DBAs

SQL Server 2016 Stretch Databases has several features that facilitate database administration (Figure 6). There is symmetry in manageability functions and tools (security, access control, monitoring, maintenance, policies, and so on) between on-premises and stretched tables. You can still back up and restore as usual, but because the archived or old rows are on Azure, the backup takes less time; this is especially important with very large databases. The same principle applies to index maintenance, which is also less costly in a Stretch Databases scenario. Most importantly, you can get almost infinite database storage without increasing the complexity or TCO of an external archiving solution.

Figure 6: Stretch Databases: Data migration between on-premises and the cloud

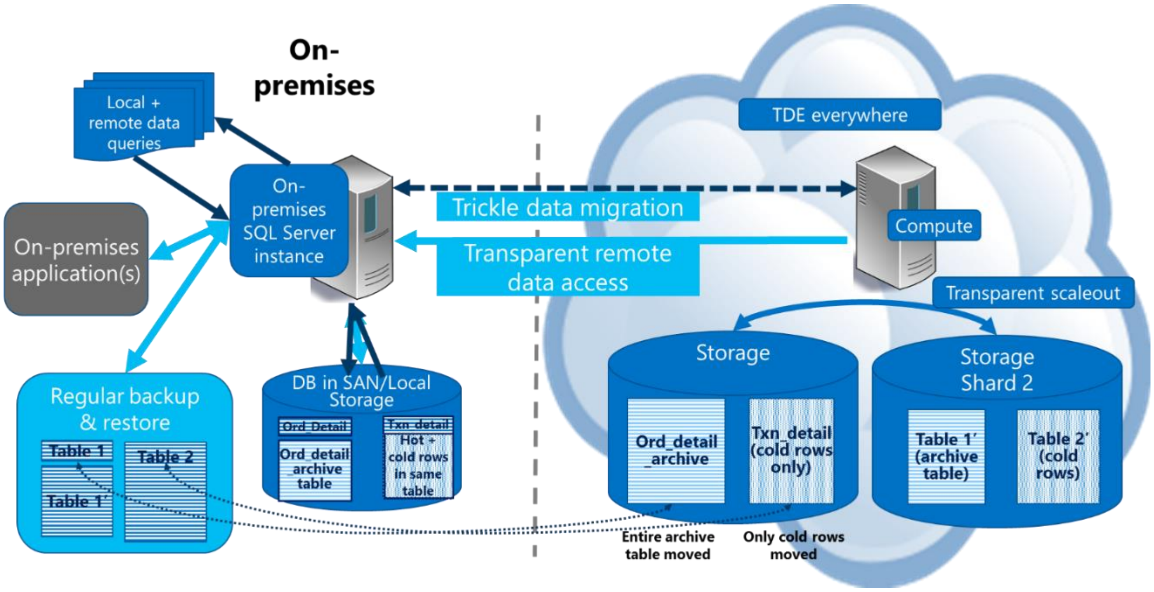

With Stretch Databases in SQL Server 2016, you get a set of extremely elegant solutions to problems common to larger databases. That being said, there are some moderate performance reductions when accessing the cold data as compared to keeping all data on-premises, and any updates or deletes are handled as an administrative function. However, based on the nature of archived/old data, those tradeoffs will generally be worth the lower TCO, reduced complexity, and other benefits that Stretch Databases provides.

Simplicity

In addition to providing better scalability and lower TCO, a hybrid cloud environment can also simplify operations. SQL Server supports several tools that facilitate administrative functions, including backup to Azure, migration of on-premises SQL Server to Azure, and the ability to easily add an Azure node to an AlwaysOn Availability Group in a hybrid environment. SQL Server 2016 adds enhancements to these features. (Figure 7)

Figure 7: SQL Server tools that facilitate administrative functions on the cloud

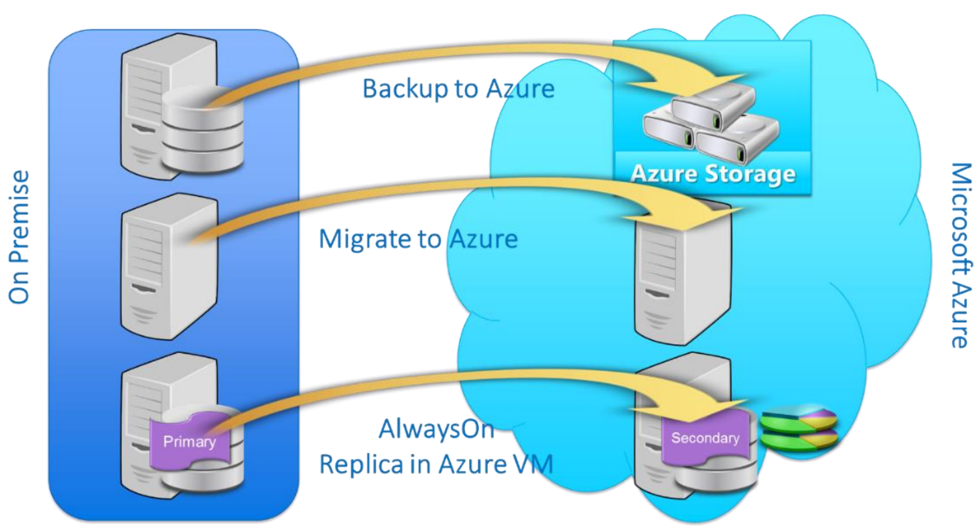

Enhanced in SQL Server 2016:

Enhanced backup to Azure

Your backup strategy can be greatly enhanced with a hybrid cloud approach. SQL Server has several options for backing up to Azure, including managed backup, backup to Azure Block Blobs, and Azure Storage snapshot backup. SQL Server 2016 has made enhancements in each of these backup options.

Managed backup

Managed backup provides the capability to easily manage and automate SQL Server backups to the Microsoft Azure Blob Storage. You can manage the whole instance or particular databases with an easy interface that can be accessed directly in the SQL Server Management Studio Object Explorer in the Management node. It provides you with off-site, geo-redundant backups with the ability to control the retention period.

SQL Server 2016 provides several enhancements to this tool. For example, you now have more granular control of the backup schedule for both full and log backups. Also, backups are made locally and then uploaded to the cloud, allowing for faster recovery and resiliency to transient network issues. You can also now run managed backups on system databases. Additionally, SQL Server 2016 now supports the simple recovery mode.

Backup to Azure with Block Blobs

Backup to Azure Block Blobs allows you to manage backups to Azure Blob Storage with fine-grained control over the process. This service, which was called *backup to URL* in previous versions, has been enhanced with SQL Server 2016. Now, the process has been incorporated into the database engine and the type of blob storage used has been altered from 3-copy page blobs to significantly less expensive parity-based block blobs. This provides the following benefits:

- Cost savings on storage.
- Almost all backup and restore features now available, including mirrors and stripes.
- Significantly improved restore performance.
- More granular control over Azure Storage.

Backup to Azure with File Snapshot

SQL Server 2014 introduced Data Files in Microsoft Azure. This enables native support for SQL Server database files stored as Azure blobs. It allows you to create a database in SQL Server running on-premises or in a virtual machine on Azure with a dedicated storage location for your data in Azure Blob Storage. This provides an alternative storage location for your database backup files by allowing you to restore from or to Azure Storage.

SQL Server 2016 builds on this capability with Backup with File Snapshot. Backup with File Snapshot provides the fastest and cheapest method for creating backups and running restores. Unlike Backup to Azure with Block Blobs, data is not actually moved. Instead, when SQL Server database files are directly stored in Azure Storage, the snapshot of those files is created. You only need to run a full backup once to establish the backup chain. Snapshot backups minimize the use of SQL Server resources to accomplish the backup. This is especially useful for moderate to very large databases, where the impact of backups can be significant.

Enhanced in SQL Server 2016:

Easy migration of on-premises SQL Server

SQL Server 2016 enhances the process of migrating SQL Server from on-premises to Azure Virtual Machines. In previous versions, when you moved SQL Server to Azure Virtual Machines, only the schema and data were migrated. In SQL Server 2016, you can migrate system objects (like logins, jobs, and certificates) and SQL Server settings (like trace flags, default language, and memory settings) (Figure 8). Additionally, the migration wizard has been enhanced to provide recommendations for image size and virtual machine size. The migration is literally as simple as point-and-click.

Figure 8: SQL Server 2016 enhanced migration

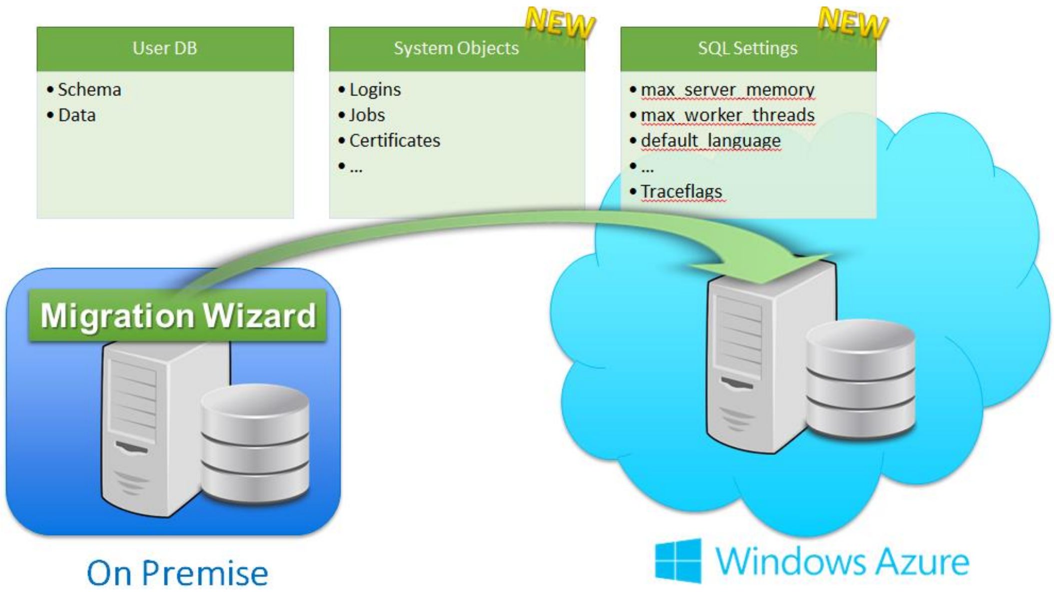

Enhanced in SQL Server 2016:

Simplified Add Azure Replica Wizard

Within the context of an AlwaysOn Availability Group, the primary role and secondary role of availability replicas are typically interchangeable in a process known as *failover*. An Availability Group fails over at the availability replica level. That is, an Availability Group fails over to one of its secondary replicas (the current failover target). An Availability Group needs a group listener to provide connectivity to the database.

The Add Azure Replica Wizard adds a replica of your databases to Azure Blob Storage. In SQL Server 2016, the group listener is created and configured within the wizard. This allows clients to connect seamlessly to the Azure replica after failover, as soon as the wizard completes its setup and without additional complex steps (Figure 9).

Figure 9: Validating the group listener for an Availability Group within the Add Azure Replica Wizard

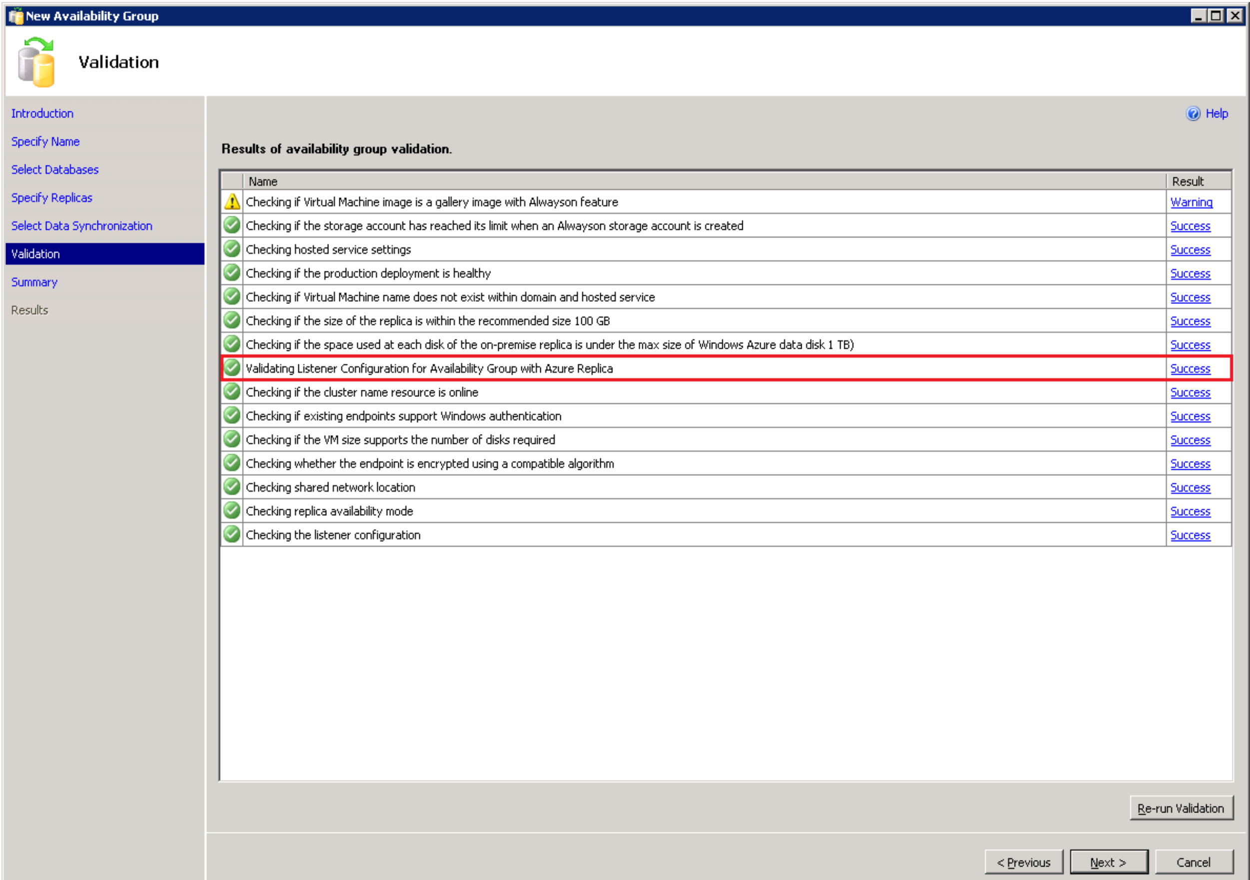

Consistency

Common development, management, and identity tools

The foundations of the common Microsoft development, management, and identity tools are Windows Server and Microsoft Azure. These are complemented by the full breadth of Microsoft technology solutions, including Microsoft SQL Server, Microsoft System Center, Microsoft Active Directory, and Microsoft Visual Studio. Together, these technologies provide one consistent platform for infrastructure, applications, and data that can span your data center, service provider data centers, and the Microsoft public cloud.

This consistent platform removes complexity so enterprises can focus on identifying business needs and determining how applications can address them. This vision of consistency can help enterprises be flexible in thinking about the best ways to deploy applications across the three clouds. The main components of this platform are integrated virtualization, a complete data platform, unified management and DevOps, a flexible development paradigm, and a common identity (Figure 10).

Figure 10: One consistent platform

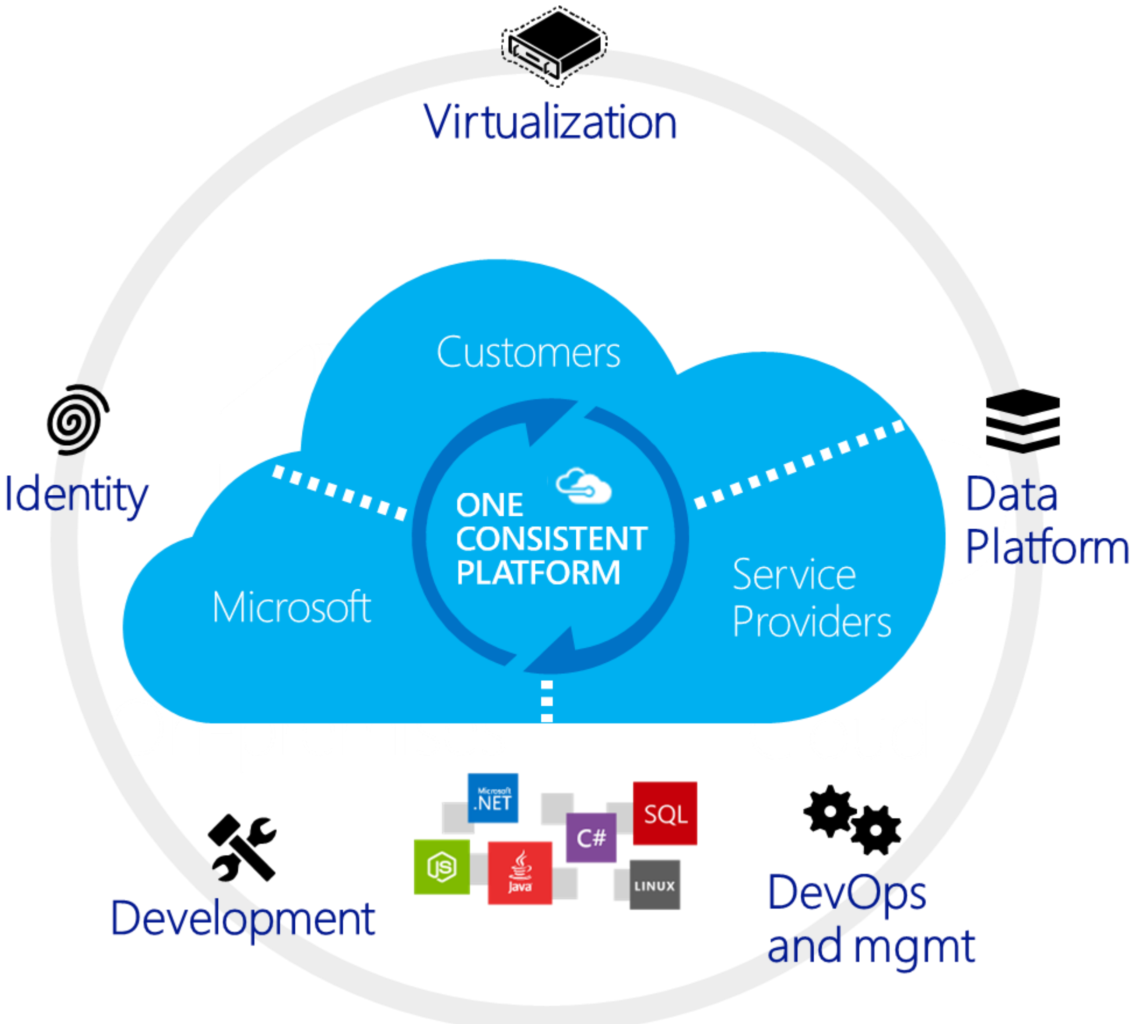

Integrated virtualization

The integrated and portable virtualization built into Windows Server allows your team to virtualize not just servers, but also the network, storage, and applications across clouds. With integrated virtualization, enterprises can generate compute capability for software-defined networking and virtualization of storage, regardless of where the virtual machines reside.

Complete data platform

A complete data platform powered by SQL Server and Microsoft Azure allows you to manage petabytes of data, power mission-critical applications, and give business users BI solutions with a range of tools—from Microsoft Excel to Hadoop.

Unified management

With the power of unified management, enterprises can employ tools like System Center to administer all the automation scripts used to manage and pool resources in a data center. For example, you can provision and move virtual machines, objects, or data from the data center through a service provider or to Azure for bottomless storage, redundancy, and disaster recovery.

Flexible development

Flexible development allows your organization's developers to use their choice of tools, languages (Microsoft or open source), and open standards to quickly build applications; connect them with other apps and data; and then deploy on-premises, in the cloud, or in a hybrid model. Visual Studio and Microsoft Team Foundation Server enable application lifecycle management—from the idea through the deployment of an app.

Common identity

The Microsoft vision for a consistent hybrid cloud platform gives enterprises the ability to manage and connect all their applications seamlessly—regardless of where those applications are hosted. Active Directory and Microsoft Azure Active Directory provide a powerful base for single identity across clouds to securely extend applications to people and their devices. For example, an enterprise can leverage its existing Active Directory group policies with Active Directory Federation Services and Azure Active Directory so that group policies extend to the cloud automatically.

SQL Server and Azure

The Microsoft Data Platform leverages SQL Server technology and makes it available across physical on-premises machines, private clouds, private clouds hosted by a third party, and the public cloud. This enables you to meet unique and diverse business needs through a combination of on-premises and cloud-hosted deployments, while using the same set of server products, development tools, and expertise across these environments.

As seen in Figure 11, each offering can be characterized by the level of administration you have over the infrastructure (on the X axis) and by the degree of cost efficiency achieved by database-level consolidation and automation (on the Y axis).

Figure 11: SQL Server technology: from on-premises to the public cloud

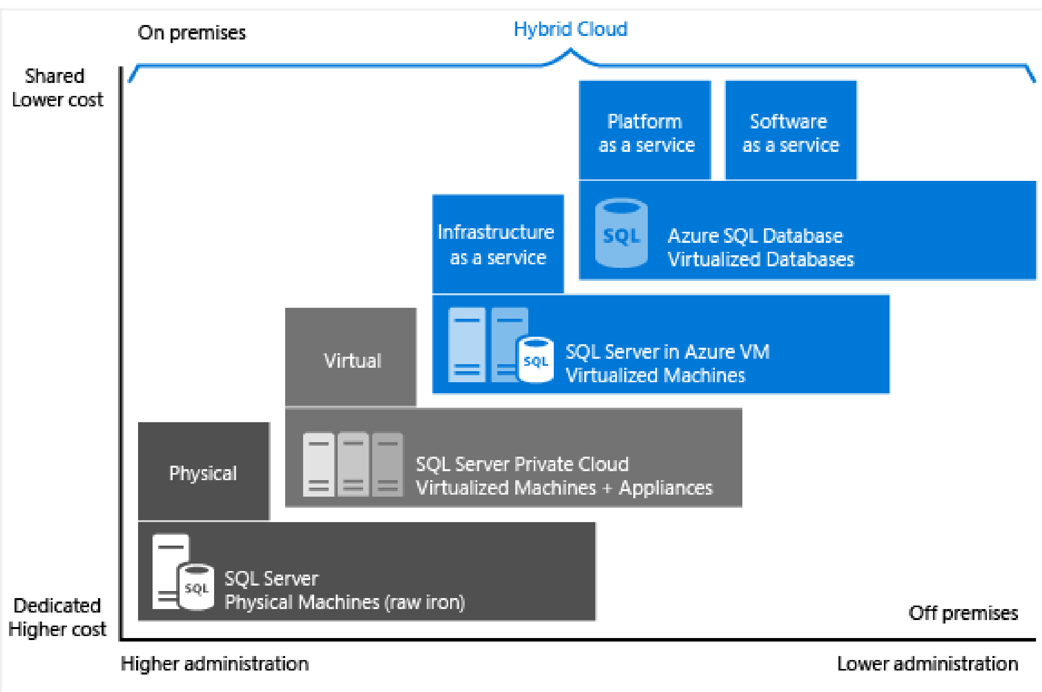

When designing an application, you have four basic options for hosting the SQL Server portion:

- SQL Server on non-virtualized physical machines.
- SQL Server in on-premises virtualized machines (private cloud).
- SQL Server in Azure Virtual Machines (public cloud).
- Azure SQL Database (public cloud).

New and Enhanced in SQL Server 2016:

SQL Server 2016 creates new functionality and enhances existing services to facilitate moving to a hybrid cloud. Scalability, availability, security, identity, backup and restore, and replication all have enhancements for working with SQL Server in Azure Virtual Machines or other Azure services.

In addition, a plethora of new data sources available (many residing natively in the cloud) means that integration, processing, and analytics are also all impacted by enhancements in SQL Server 2016. Analysis Services, Reporting Services, and Integration Services all have additional features that make working with cloud data and hybrid solutions easier. SQL Server 2016 enhances the consistency of your development and administration efforts by making it much easier to work within a hybrid cloud environment.

Conclusion

SQL Server has a long history of moving toward easier integration with outside data sources while simplifying management and administration. SQL Server 2016, specifically, has been designed to facilitate working in a hybrid cloud environment, where data and services can reside in many locations. The Stretch Databases feature in SQL Server 2016 illustrates the scalability that this kind of interactivity can provide. Administrative functions, such as backups, migrations, and availability, are all enhanced and made simpler with features in SQL Server 2016. SQL Server and Microsoft Azure work better together because the Microsoft hybrid cloud concept drives consistency in tools and processes, making it easier to realize the benefits of hyperscale cloud.

More information

The following websites offer more information about topics discussed in this white paper:

- http://www.microsoft.com/en-us/server-cloud/products/sql-server-2016/: SQL Server 2016 website
- http://technet.microsoft.com/en-us/sqlserver/: SQL Server TechCenter
- http://msdn.microsoft.com/en-us/sqlserver/: SQL Server DevCenter
- http://azure.microsoft.com/en-us/: Microsoft Azure

Feedback

Did this paper help you? Please give us your feedback by telling us on a scale of 1 (poor) to 5 (excellent) how you rate this paper and why you have given it this rating. More specifically:

- Are you rating it highly because of relevant examples, helpful screen shots, clear writing, or another reason?
- Are you rating it poorly because of examples that don't apply to your concerns, fuzzy screen shots, or unclear writing?

This feedback will help us improve the quality of white papers we release.

Please send your feedback to: mailto:sqlfback@microsoft.com