REFERENCE ARCHITECTURE

SQL Server Workloads in Azure with Pure Cloud Block Store

A reference architecture for SQL Server databases on Azure VMs with Pure Cloud Block Store™.

Contents

Executive Summary 3

Solution Overview 4

Solution Benefits 4

Technology Overview 5

Microsoft Azure 5

Microsoft SQL Server 6

Pure Cloud Block Store 6

Technical Solution Design 7

Compute Layer 7

Network Layer 7

Storage Layer 9

Operating System Layer 11

Pure Cloud Block Store Solutions for SQL Server 13

Migrating SQL Server Workloads Using Pure Cloud Block Store 13

Intra-cloud Disaster Recovery Using Pure Cloud Block Store 13

Using Volume Snapshots for Development and Test Workflows 14

Design Validation 15

HammerDB Benchmark Suite 15

Workloads and Methodology 15

Deployment Guidance 18

Milestone 1: Deploy Pure Cloud Block Store in Azure 18

Milestone 2: Install and Configure Multipath I/O 18

Milestone 3: Configure and Connect Volumes 19

Milestone 4: Apply Best Practices for Connectivity and SQL Server 19

Conclusion 19

Executive Summary

Pure Storage® Pure Cloud Block Store in Microsoft Azure delivers enterprise-grade storage efficiency, predictable high performance, and advanced data services beyond the capabilities of Azure managed disks, even when deployed on identical virtual machines. Built-in inline data reduction typically shrinks the SQL Server storage footprint by three to five times, significantly reducing Azure storage costs.

With Pure Cloud Block Store, advanced data services such as instant snapshots, near-synchronous replication (ActiveDR®), and synchronous replication (ActiveCluster™) become a cost-effective, consistent way of implementing BC/DR strategies, enabling robust data protection without significant overhead. Pure Cloud Block Store also independently scales storage performance from compute resources, removing the need to upgrade virtual machine size simply to achieve greater throughput or lower latency.

In recent benchmark tests, Pure Cloud Block Store delivered nearly double the analytical query throughput (TPROC-H workload) and over a third higher transactional processing performance (TPROC-C workload) compared to Azure Premium SSD v2, even while consuming substantially less physical storage.

Why it matters:

Disks managed by Azure link higher storage performance to larger virtual machines. This forces organizations to overspend on compute resources, software licensing, and operational complexity to achieve their desired database performance.

Pure Cloud Block Store solves these challenges with:

- **Reduced storage footprint**: Built-in data reduction shrinks data footprints, typically turning five terabytes of raw data into about one terabyte of billable storage.
- **Simplified business continuity/disaster recovery**: Instantaneous snapshots, built-in ActiveDR for near-synchronous replication, and ActiveCluster for synchronous replication provide powerful enterprise-grade data protection without excessive complexity or cost.
- **Performance decoupled from VM size**: Pure Cloud Block Store delivers predictable, high-performance storage independently from Azure virtual machine specifications, enabling flexible scaling of throughput and IOPS without requiring costly VM upgrades.

Solution Overview

Microsoft SQL Server is a robust, enterprise-grade database management platform widely trusted for critical transaction processing, structured data storage, and business intelligence workloads. When deployed in Microsoft Azure Virtual Machines, SQL Server is an essential building block for scalable, resilient, cloud-based business application solutions. Leveraging Pure Cloud Block Store as persistent storage for virtualized SQL Server databases in Azure enables organizations to achieve greater storage efficiency, predictable high performance, and streamlined data protection, all while significantly lowering cloud storage costs.

This reference architecture provides clear guidance for deploying SQL Server databases on Azure Virtual Machines using Pure Cloud Block Store exclusively as the foundation for persistent storage for user databases.

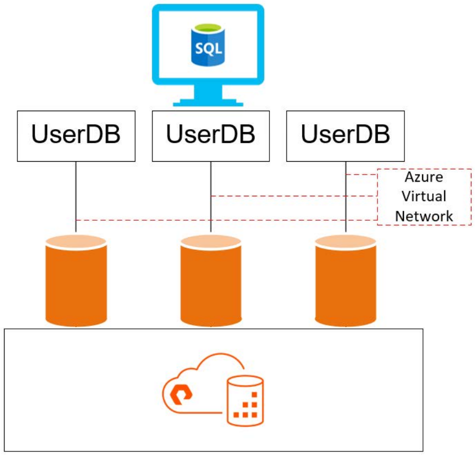

FIGURE 1 Overview of SQL Server deployed on Azure Virtual Machines using Pure Cloud Block Store.

Note: This architecture specifically addresses SQL Server user databases deployed on Azure Virtual Machines. It does not apply to Azure SQL Managed Instance, Azure SQL Database, or Azure SQL Edge.

Solution Benefits

Deploying SQL Server on Azure Virtual Machines with Pure Cloud Block Store delivers several key business and operational advantages:

- **Significant storage cost reduction**: Built-in inline data reduction typically shrinks SQL Server storage footprints by three to five times, dramatically lowering Azure storage expenses and total cost of ownership.
- **Simplified business continuity and disaster recovery**: Pure Cloud Block Store provides enterprise-grade data services such as instantaneous snapshots, near-synchronous replication (ActiveDR), and synchronous replication (ActiveCluster), enabling robust data protection and disaster recovery without additional complexity or high cost.
- **Independent, predictable storage performance**: Pure Cloud Block Store separates storage performance from virtual machine sizing, allowing organizations to flexibly scale IOPS and throughput on the storage array without requiring costly upgrades of virtual machine resources, software licenses, or unnecessary operational overhead.
- **Avoidance of cloud vendor lock-in**: Pure Cloud Block Store runs the same Purity operating environment available in on-premises and other public cloud deployments, giving organizations full flexibility to adopt hybrid or multi-cloud strategies without losing data control or consistency.
- **Operational simplicity through automation**: Pure Cloud Block Store supports comprehensive automation through its REST APIs and common infrastructure-as-code tools, streamlining database provisioning, backup workflows, and operational tasks to reduce manual intervention and operational complexity.

Technology Overview

This reference architecture leverages Azure as the foundational cloud platform for deploying SQL Server workloads in Azure Virtual Machines using Pure Cloud Block Store as the block storage for user databases. Azure serves as a single platform for all components, integrating compute, storage, and networking services to deliver an environment that performs consistently and that is scalable and reliable.

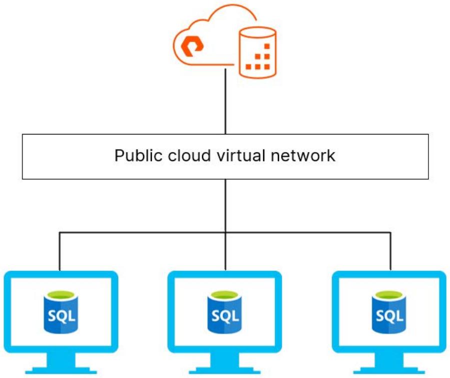

FIGURE 2 Technology components of the solution

Microsoft Azure

Microsoft Azure serves as the foundational platform for deploying SQL Server workloads in this reference architecture. It provides a scalable, secure, and resilient environment that supports modern database requirements with integrated compute, storage, and networking services. Relevant services included in this reference architecture include:

- Azure Virtual Machines, delivering flexible and scalable compute resources capable of handling a wide range of transactional and analytical SQL Server workloads. Azure Virtual Network enables secure communication between resources, while native support for iSCSI and other protocols facilitates efficient connectivity to storage solutions.
- Azure Availability Zones enhance system reliability by ensuring fault tolerance and high availability across geographically separated data centers. Azure services, such as Premium SSD v2 for high-performance storage, further extend the capabilities of the platform, enabling robust disaster recovery and backup strategies.

As a unified cloud environment, Microsoft Azure integrates compute, storage, and networking seamlessly to create a secure, high-performing foundation for mission-critical database deployments.

Microsoft Azure Virtual Machines

Azure Virtual Machines provides the flexible and scalable compute foundation for deploying SQL Server user databases in this architecture. Supporting both Windows Server and assorted Linux platforms, Azure Virtual Machines delivers the performance and reliability required for modern database workloads.

SQL Server workloads can utilize various Azure Virtual Machines families tailored to different requirements. For memory-intensive applications, E-series and M-series virtual machines offer high memory capacity and scalability, while D-series virtual machines provide a balanced option for general-purpose deployments. This flexibility ensures that organizations can select virtual machines aligned with their specific workloads and performance needs.

Microsoft SQL Server

Microsoft SQL Server is a relational database developed by Microsoft. It is designed to store, retrieve, and manage large volumes of data for enterprise applications, and it supports a wide range of workloads, from online transaction processing to complex online analytical processing. SQL Server ensures data integrity, provides robust security features, and supports high availability through features like Always On Availability Groups.

A SQL Server instance represents the installation of the database engine and its associated services on an Azure virtual machine. It is responsible for managing one or more user databases and system databases. A SQL Server database consists of a set of files (data and transaction log) that persist in relationally structured data for use.

Pure Cloud Block Store

Pure Cloud Block Store by Pure Storage is the block storage solution for SQL Server user databases deployed in Azure Virtual Machines. Purpose-built for cloud environments, Pure Cloud Block Store delivers enterprise-grade storage capabilities while ensuring a consistent storage experience across on-premises and Azure environments. This seamless mobility enables organizations to maximize flexibility for their SQL Server workloads.

Pure Cloud Block Store incorporates advanced data reduction technologies, significantly reducing cloud storage costs without compromising performance or scalability. Optimized for Azure, Pure Cloud Block Store integrates with Azure Virtual Machines, to provide low-latency, high-performance storage tailored to the demands of SQL Server deployments. Additionally, Pure Cloud Block Store supports offloading snapshots to Azure Blob Storage, offering a robust disaster recovery solution with rapid restores and long-term retention.

By leveraging Pure Cloud Block Store, organizations can achieve predictable performance, cost efficiency, and simplified management for SQL Server databases in Azure.

Cloud Block Store is available in two versions, //V10 and //V20. Each version provides different capacity and performance capabilities.

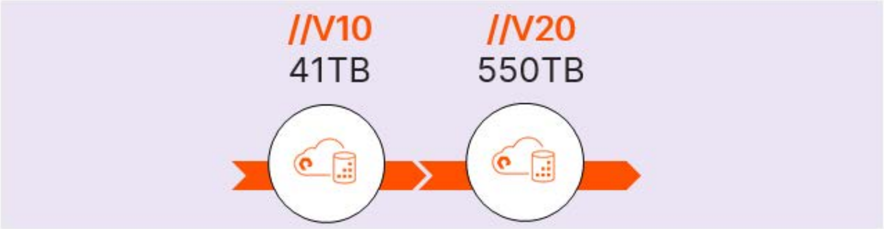

FIGURE 3 The V10 and V20 Pure Cloud Block Store version. The capacity listed for each model is effective with a 4:1 data reduction rate.

Pure Cloud Block Store deployments are powered by Purity for FlashArray™. Purity delivers rich, enterprise level data services that ensure data is stored in the most secure and efficient way while providing additional functionality to extend storage capabilities.

For more technical details, refer to the Pure Cloud Block Store for Azure documentation.

Technical Solution Design

This technical solution design is structured around a layered approach to deliver an optimized SQL Server deployment in Azure. Through the coupling of Azure Virtual Machines to Pure Cloud Block Store, this solution addresses performance, scalability, and reliability across distinct compute, network, storage, operating system, and database layers.

Compute Layer

Azure Virtual Machines serve as the compute layer, hosting SQL Server instances on either Windows Server or supported Linux platforms. Azure virtual machines provide scalable compute resources to meet the demands of transactional and analytical workloads.

Key considerations for the compute layer include:

- **Performance optimization**: Ensure virtual machines are configured with high-clock-speed processors and sufficient memory to minimize reliance on disk input/output.
- **Scalability**: Resources should align with workload demands, scaling virtual machines as required to accommodate growth and performance requirements.

Virtual Machine Families for SQL Server

Azure Virtual Machines supports a variety of families optimized for different workload types:

- **E-series**: Memory-optimized virtual machines offering up to 672GiB of RAM, ideal for large-scale transactional and analytical workloads.
- **M-series**: High-memory virtual machines designed for extreme-scale and mission-critical SQL Server deployments.
- **D-series**: General-purpose virtual machines that balance compute, memory, and storage performance, suitable for smaller or moderately demanding workloads.

Note: Minimum and recommended configurations depend on workload requirements, with more processor cores and higher clock speeds recommended for large-scale transactional and analytical workloads. The Microsoft Azure virtual machine best practice checklist for SQL Server provides a detailed and up to date overview of the optimized virtual machines for performance workloads

Network Layer

The network layer ensures secure, low-latency communication between Azure Virtual Machines, Pure Cloud Block Store, and client connections. Azure Virtual Network provides the foundation for connectivity, with subnets segmented for specific roles to optimize traffic flow and maintain isolation.

Network Connection between Pure Cloud Block Store and Azure Virtual Machines

Pure Cloud Block Store communicates with Azure Virtual Machines using iSCSI over TCP/IP. This setup requires proper configuration of iSCSI initiators on Azure Virtual Machines to establish secure connections with Pure Cloud Block Store volumes.

To minimize latency and enhance security:

- Deploy Pure Cloud Block Store and Azure Virtual Machines with storage protocol (iSCSI for Windows, NVM Express over TCP for Linux) connectivity in the same Azure Virtual Network.
- Configure dedicated subnets for SQL Server client connections, storage communication, and management traffic.
- Configure the Azure virtual machine's network interface with accelerated networking.

Key Network Design Principles

Effective network design ensures optimal performance, reliability, and security. The following principles guide the configuration:

Azure Virtual Network:

- Host SQL Server virtual machines and Pure Cloud Block Store within a dedicated Azure Virtual Network to ensure isolated and secure communication.
- Define subnets for specific roles:
  - **System subnet**: For management and control traffic
  - **iSCSI subnet**: For storage communication between SQL Server databases and Pure Cloud Block Store
  - **Application subnet**: For client application traffic

Network latency and proximity:

- Deploy SQL Server virtual machines and Pure Cloud Block Store in the same Azure region and availability zone to minimize latency and avoid unnecessary network hops.
- Enable Azure Accelerated Networking to improve throughput and reduce network jitter on SQL Server virtual machines.

Bandwidth recommendations:

- Configure virtual network interface controllers and Azure Virtual Network to meet the performance demands of SQL Server workloads.
- Use high-throughput network interface controllers supported by specific Azure Virtual Machines families for data-intensive SQL Server operations.

Protocols and Traffic Configuration

Efficient communication between SQL Server, Pure Cloud Block Store, and client applications require optimized traffic flow and protocol usage:

- **Storage traffic**: Use iSCSI over TCP/IP for storage communication. For Linux-based SQL Server deployments, consider NVM Express over TCP for lower latency.
- **Application traffic**: Configure Azure Virtual Network to support SQL Server listeners for Always On Availability Groups or Failover Cluster Instances. Assign static IP addresses within appropriate subnets to support seamless client traffic routing, dynamic failover, and high availability for SQL Server applications.

Security and Isolation

To safeguard communication between SQL Server and Pure Cloud Block Store, this architecture leverages Azure-native security features:

- **Network security groups**: Apply network security groups to Pure Cloud Block Store and Azure Virtual Machine subnets to restrict traffic to required ports (for example, port 3260 for iSCSI) and authorized source/destination IPs.
- **Azure Firewall**: Optionally deploy Azure Firewall for centralized control and monitoring of network traffic between SQL Server and Pure Cloud Block Store.

For detailed guidance, refer to Azure network security best practices and documentation on network security groups, Azure Firewall, and Accelerated Networking.

Storage Layer

The storage layer in this reference architecture integrates Azure managed disks, local ephemeral disks, and Pure Cloud Block Store to optimize SQL Server performance, scalability, and reliability. Azure managed disks are utilized primarily for system databases, ensuring high IOPS and low latency for critical SQL Server system operations. TempDB is deployed on local ephemeral disks to leverage their superior performance characteristics, benefiting workloads that require rapid temporary data operations.

Pure Cloud Block Store, itself built on top of managed disk and serves as the primary storage solution for user database data and log files. It delivers data reduction, efficient resource allocation, and consistent performance tailored specifically for SQL Server user workloads.

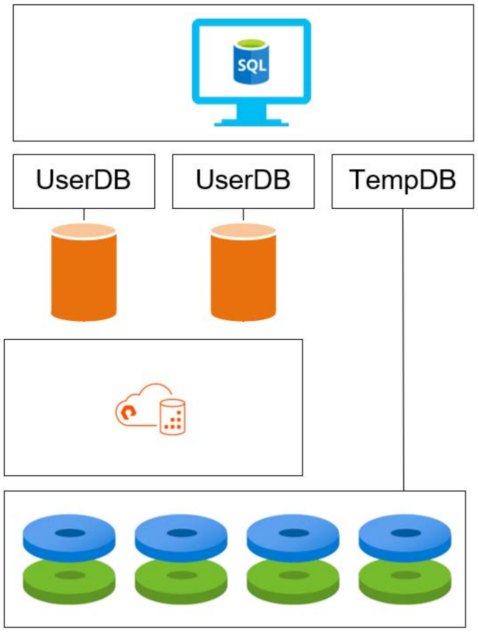

FIGURE 4 Storage layer overview diagram of the SQL Server in Azure Virtual Machines with Pure Cloud Block Store solution.

Pure Cloud Block Store

Pure Cloud Block Store is the primary block storage solution for SQL Server user databases, providing enterprise-grade capabilities that optimize performance and reduce storage costs. This layer ensures consistent performance and seamless integration with Azure services.

Storage considerations include:

- **Storage placement**: Deploy SQL Server data files (.mdf/.ndf) and log files (.ldf) on separate Pure Cloud Block Store volumes to optimize input/output throughput and reduce contention.
- **Snapshot management**: Use crash consistent volume snapshots in Pure Cloud Block Store for fast, space-efficient backups and rapid recovery. Consider offloading these snapshots to Azure Blob Storage for cost-effective, long-term retention. Application consistency can be achieved through the use of T-SQL database snapshots.

The following table outlines design recommendations for Pure Cloud Block Store.

Configuration Category	Design Recommendation
Volume configuration	• Configure Pure Cloud Block Store to host the volumes for SQL Server user databases, such as data files (.mdf and .ndf) and log files (.ldf). • As a design best practice, dedicate a group of storage volumes exclusively to a single database instance. Avoid placing multiple databases on the same set of volumes to maintain optimal performance, simplify management, and enhance reliability.
High availability, disaster recovery and data protection	• Use Pure Cloud Block Store space-efficient volume snapshots for backups, quick restores, and disaster-recovery scenarios. Volume snapshots should complement an existing data protection strategy that utilizes a T-SQL or ISV-driven backup solution. • Implement ActiveDR between Pure Cloud Block Store or FlashArray deployments for near-synchronous replication to a secondary availability zone, ensuring robust disaster recovery with no impact on primary zone performance. • Implement synchronous active/active replication with ActiveCluster to provide high availability for SQL Server, offering zero or near-zero recovery point objectives and recovery time objectives.
Proximity and latency	• Deploy Pure Cloud Block Store in the same Azure region and availability zone as the SQL Server virtual machine to minimize latency and reduce network hops.

TABLE 1 Pure Cloud Block Store design recommendations.

For more information see the Pure Cloud Block Store for Azure Overview.

Azure Managed Disks

Azure managed disks complement Pure Cloud Block Store by hosting SQL Server system databases and TempDB, critical for database operations and performance.

Key design considerations:

- **TempDB Configuration**: Provision TempDB on Azure Premium SSD v2 or Ultra SSD managed disks to efficiently handle intensive input/output workloads. This aligns with Microsoft's best practices, leveraging the disks' native performance capabilities and ensuring optimal performance under heavy TempDB utilization.

Operating System Layer

The operating system layer serves as the foundation for SQL Server operations, providing the necessary environment to support performance, scalability, and stability. It encompasses key optimizations for file systems, resource management, and input/output handling to maximize the efficiency of SQL Server workloads.

This architecture supports Windows Server and Linux operating systems in line with the requirements for SQL Server, both of which are fully compatible with SQL Server. These operating systems can be deployed on Azure Virtual Machines using templates provided by Azure, which streamlines the setup process.

For details on supported operating systems for SQL Server, refer to the SQL Server hardware and software requirements. Azure templates for deploying these operating systems are available in the Azure Marketplace.

Best Practices for Deployments on Windows Server

Windows Server best practices include the actions detailed in Table 2.

Configuration Category	Best Practices
Folder structure	Both virtual and physical disks can use either drive letters or mount points. Make sure to provision specific locations for each database file function in alignment with the following examples: • **SQL Server binaries:** Install SQL Server on the operating system drive. • **Data files:** Do not place database and backup files on the same disk where the operating system and binaries exist. Do not put SQL Server data and log files on the root of a drive or a mount point. Do not create data files on the root of a file system in Windows Server. Use naming conventions for drive letters or mount points that make sense for that environment. The following are two examples of naming conventions: • Use a folder on the D drive, such as **D:\DATA**. All the data files can be in a single virtual disk, or each database can have a virtual disk provisioned for it (recommended) and then mounted under this structure. • The data files in “UserDB” would be placed in a virtual disk mounted **D:\DATA\UserDB**. • **Transaction log files:** All transaction log files can be located on a single virtual disk, or each database can have a virtual disk provisioned for it (recommended) and then mounted under this structure. The following is an example of a naming convention: • Create a folder named LOGS on the L drive, **L:\LOGS**. • The log files in “UserDB” would be placed in a virtual disk mounted at **L:\LOGS\UserB**. • **TempDB files:** Provision a single virtual disk to contain the data and log files for TempDB.
File system and disk formatting	Format all SQL Server volumes with the NTFS file system and a minimum of 64KB allocation units *before* placing any data on the disks.
Block alignment	Check if the block alignment is configured properly. See the Windows Server File Systems FAQ for FlashArray knowledge article for more information.
iSCSI Configuration	Configure iSCSI with a minimum of four sessions per volume for optimal throughput, reliability, and redundancy. This configuration helps maximize network bandwidth utilization and overall storage performance.
Antivirus and threat-detection systems	The minimum recommendation is to configure antivirus and threat detection, so they exclude files with .mdf, .ndf, .ldf, and .bak extensions. Other folders and files might need to be excluded as well. For more information, see Configure antivirus software to work with SQL Server.
Page file size	Enable instant file initialization as per Database instant file initialization.

TABLE 2 Azure Virtual Machines Windows Server best practices.

Best Practices for Deployments on Linux

Linux is only supported as a platform from SQL Server 2017 onwards. Best practices for Linux virtual machines include those detailed in Table 3.

Configuration Category	Best Practices
Disk storage-protocol selection	While this reference architecture is validated for Windows Server deployments using iSCSI, NVM Express over TCP offers a compelling option for Linux environments, where its benefits align with workload requirements. NVM Express over TCP is recommended for Linux virtual machines to take advantage of: • **Lower latency**: Provides faster data access, improving performance for transactional and analytical workloads • **Higher throughput**: Supports workloads with high input/output demands, ensuring consistent performance under heavy load
Folder structure	The default folder for all SQL Server data files is **/var/opt/mssql**. Default folders for data, transaction log, and TempDB files are created during installation: • **Data**: **/var/opt/mssql/data** • **Log**: **/var/opt/mssql/log** • **TempDB**: **/var/opt/mssql/tempdb** TempDB files should be migrated to a separate virtual or physical disk to achieve the best performance. The virtual or physical disk can be mounted at the default TempDB location after the files have been migrated to it. Create separate folders for data and transaction log files for any additional user databases.
File system and mount options	Format virtual disks with the XFS or EXT4 file systems and mount them with the "**rw,attr2,noatime**" attributes.
Open file limitations	Set a soft limit of 16,000 and a hard limit of 32,727.
Forced unit access input/ output subsystem capability	For Linux operating systems that support this capability, configure forced unit access using the commands **`control.writethrough =1`** and **`control.alternewritethrough = 0`.**
Disk settings	• **Input/output scheduler:** Use the none scheduler for all virtual disks that contain SQL Server data files. For more information, see the How to use the Noop or None IO Schedulers solution page. • **Disk readahead**: Set the readahead setting to **4,096**. • **vm.swappiness:** Set vm.swappiness to 1 for high-performance systems. • **vm.dirty_***: Set **vm.**dirty_ratio to 80 and vm.dirty_background_ratio to 3. • **kernel.sched_***:Set kernel.sched_min_granularity_ns to 15,000,000 and kernel.sched_wakeup_granularity_ns to 2,000,000.
SQL Server	The Microsoft documentation details out the best practices for SQL Server in Azure virtual machines.

TABLE 3 Azure Virtual Machines Linux best practices.

Pure Cloud Block Store Solutions for SQL Server

This section outlines specific technical solutions using Pure Cloud Block Store to address key SQL Server workload requirements on Azure Virtual Machines. The solutions demonstrate practical implementations for common use cases, including database migration, disaster recovery, and database provisioning for development and testing environments. Each use case highlights how Pure Cloud Block Store simplifies infrastructure management, accelerates data mobility, and enhances operational efficiency in Azure-based SQL Server deployments.

Migrating SQL Server Workloads Using Pure Cloud Block Store

Overview

Pure Cloud Block Store simplifies SQL Server migration to Azure Virtual Machines through efficient snapshot-based database copies. Snapshots move only data-reduced blocks, greatly accelerating database migration and minimizing network traffic.

Key Benefits

- **Rapid migration**: Quickly provision test/dev environments from snapshots without duplicating full data copies.
- **Reduced infrastructure:** Minimize storage needs by leveraging Pure Storage data reduction technology.
- **Faster data transfer**: Efficient snapshot transfers reduce the amount of data sent over the network, speeding up migration workflows

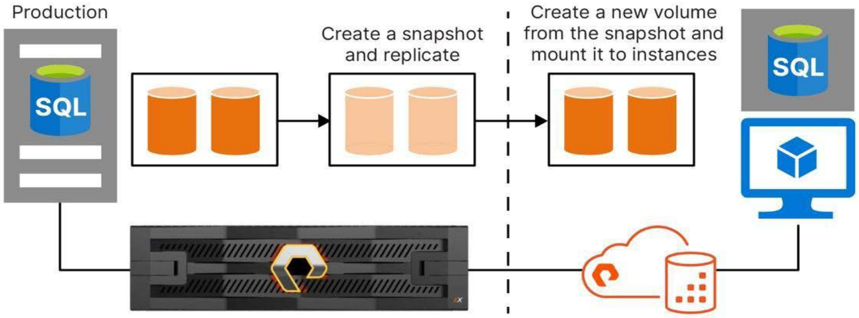
Production
SQL
Create a snapshot and replicate
Create a new volume from the snapshot and mount it to instances
SQL

FIGURE 5 Overview of the process of using Pure Cloud Block Store to migrate SQL Server workloads to Azure Virtual Machines.

Intra-cloud Disaster Recovery Using Pure Cloud Block Store

Overview

Pure Cloud Block Storesimplifies intra-cloud disaster recovery for SQL Server workloads on Azure Virtual Machines through efficient storage-based replication technologies. By using snapshot-based asynchronous replication and ActiveDR continuous replication, Pure Cloud Block Store delivers low recovery point objectives (RPO), reduced downtime, and improved operational resilience across Azure availability zones or regions.

Key Benefits

- **Rapid recovery:** ActiveDR supports near-zero RPO with automated failover.
- **Efficient replication:** Snapshot-based asynchronous replication reduces data transfer, enabling RPOs as low as five minutes.
- **Cost-optimized:** Data deduplication and compression minimize storage costs and network traffic.
- **Enhanced protection:** Immutable SafeMode™ Snapshots safeguard against ransomware and data corruption.

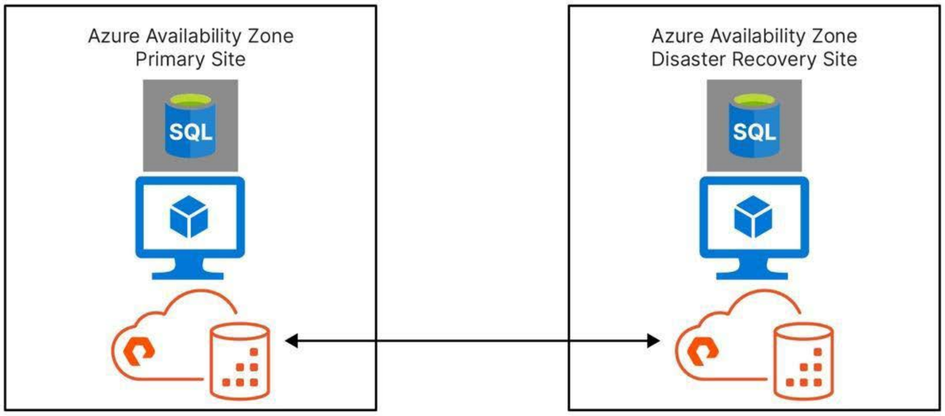

FIGURE 6 Intra-cloud disaster recovery for SQL Server on Azure.

Using Volume Snapshots for Development and Test Workflows

Overview

Pure Cloud Block Store (Pure Cloud Block Store) snapshots enable rapid, efficient creation of SQL Server database copies for development, testing, or analytics on Azure Virtual Machines. Snapshot-based cloning reduces provisioning time and avoids performance impacts on production workloads.

Key Benefits

- **Faster dev/test cycles:** Instantly create point-in-time database copies without performing full backups.
- **Reduced storage and costs**: Data deduplication and compression minimize Azure storage usage and data-transfer fees.
- **No production disruption**: Offload resource-intensive operations, such as maintenance, upgrades, and analytics, to isolated, non-production environments.
- **Simple and automated**: Manage snapshots through cloud-native tools from Pure Storage, APIs, or automation frameworks.

Snapshots can be crash-consistent or application-consistent, depending on the recovery needs, ensuring flexible and reliable database provisioning.

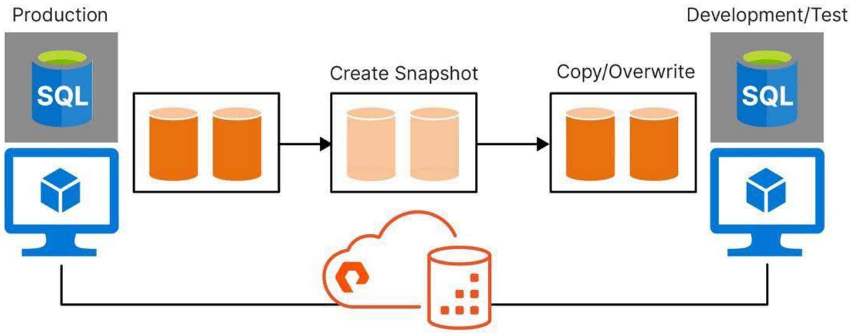

FIGURE 7 Cloned SQL Server database for dev/test.

Design Validation

This design validation evaluates the performance benefits of deploying SQL Server workloads on Azure Virtual Machines using Pure Cloud Block Store (V20 with Premium disks) compared to Azure Premium SSD v2 Managed Disks. The testing methodology simulates realistic transactional and analytical database workloads to characterize storage performance and throughput.

HammerDB Benchmark Suite

HammerDB is an industry-standard, open-source database benchmarking tool designed to simulate realistic database workloads, measuring the performance of transactional (OLTP) and analytical (OLAP) databases. It provides standardized benchmarks such as TPC-C and TPC-H, widely recognized for evaluating database system performance and scalability.

Workloads and Methodology

This design validation utilizes the following common environment components and software:

Component	Detail
Azure Virtual Machine	**Specification**: Standard E48as v5 (48 vCPUs, 384 GiB RAM, 71,000 IOPS limit) **Operating system**: Windows Server 2025 Datacenter **Networking**: Private IP with accelerated networking enabled **Azure managed disks**: • TempDB – Ultra disk LRS – Max IOPs=3000, Max throughput (MBps)=200 • SQL_data01 – Premium SSDv2 LRS – Max IOPs=20000, Max throughput (MBps)=250 • SQL_data02– Premium SSDv2 LRS – Max IOPs=20000, Max throughput (MBps)=250 • SQL_data03– Premium SSDv2 LRS – Max IOPs=20000, Max throughput (MBps)=250 • SQL_data04– Premium SSDv2 LRS – Max IOPs=20000, Max throughput (MBps)=250 • SQL_log– Premium SSDv2 LRS – Max IOPs=20000, Max throughput (MBps)=250
Pure Cloud Block Store	V20 with Premium SSD (CBS-V20MP2R2) with Purity 6.8.5. **Volumes provisioned**: • CBS_SQL_data01 • CBS_SQL_data02 • CBS_SQL_data03 • CBS_SQL_data04 • CBS_SQL_log **iSCSI sessions:** 4 (2 per controller) **Multiple connected sessions policy**: Least Queue Depth (LQD)
Microsoft SQL Server	SQL Server 2022 Developer Edition **TPCROC-C database**: 4 data files (128GB each) and 1 log file, each within its own volume. **TPCROC-H database:** 4 data files (256GB each) and 1 log file, each within its own volume. **Max Degree of parallelism:** 8 **Cost threshold for parallelism:** 50 **Memory limits (TPROC-C):** 64GB **Memory limits (TPROC-H):** 100GB
HammerDB	Version 4.12

Note: Actual performance outcomes may vary depending on database structure, workload characteristics, and Azure infrastructure specifics. It's advisable for organizations to conduct environment-specific testing to validate these benefits.

Transaction Processing Workload (TPROC-C)

This workload simulated a high-intensity transaction processing environment typical for OLTP databases. The comparative metric used was new orders per minute (NOPM). The schema was built and then backed up, being restored to the relevant storage type being evaluated.

The following configuration was used for the transaction processing workload:

- **Number of warehouses**: 5000
- **Virtual users**: 200
- **User delay** (ms): 500
- **Repeat delay** (ms): 500
- **Iterations**: 1
- **TPROC-C driver script**: Timed Driver Script
- **Total transactions per user**: 10000000
- **Checkpoint when complete**
- **Minutes of ramp-up time**: 5
- **Minutes for** test duration: 30
- **Warehouses**: Use all warehouses
- **Profile**: Time profile

Results and Analysis

In the TPC-C transactional processing benchmark, Pure Pure Cloud Block Store achieved a transactional throughput of 223,214 new orders per minute (NOPM), outperforming Azure Premium SSD v2 managed disks, which reached 159,732 NOPM. This represents approximately 40% performance improvement using Pure Cloud Block Store.

Notably, the Azure managed disks' performance was limited by the Azure VM’s maximum IOPS constraint (71,000 IOPS), indicating that Pure Cloud Block Store's performance scalability extends beyond the inherent limitations of the VM's disk infrastructure. These results highlight the ability of Pure Cloud Block Store to independently scale storage performance, providing significant transactional performance benefits without necessitating more powerful or costly virtual machine configurations.

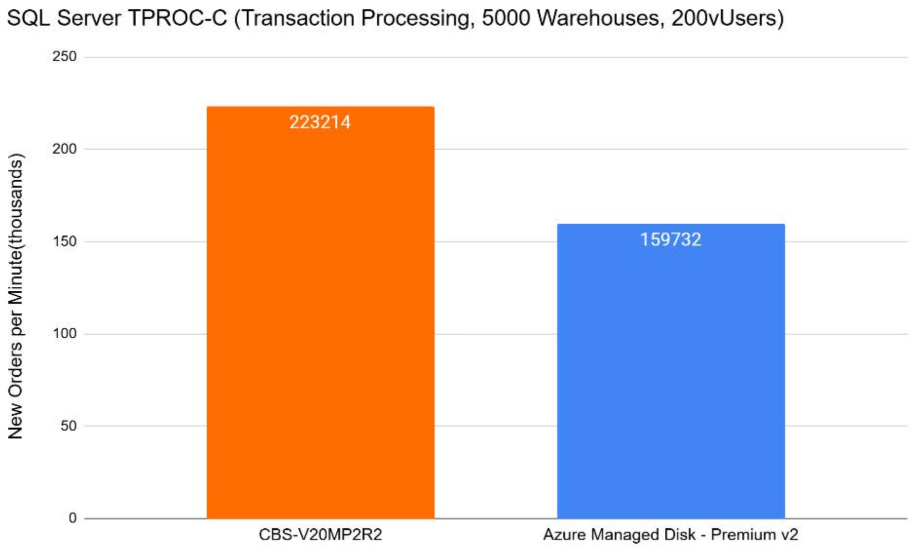

FIGURE 8 TPROC-C performance results comparing Pure Cloud Block Store and managed disk.

Decision Support Workload (TPROC-H)

This workload simulated a complex analytical and decision-support environment, representative of OLAP workloads common in data warehousing scenarios. The comparative metric used was queries per minute (QPM). The database schema was generated and loaded at a defined scale factor, backed up, and subsequently restored onto the evaluated storage solutions before executing the benchmark queries.

The following configuration was used for the analytical processing workload:

- **Scale factor**: 1000
- **MAXDOP**: 8
- **Build option:** Clustered ColumnStore
- **Virtual users**: 1
- **Total query sets per user** : 1
- **Iterations**: One full run of the query set
- **Execution profile**: Sequential execution of queries

In the TPC-H analytical processing benchmark, Pure Pure Cloud Block Store achieved a throughput of 1.45 QPM, significantly surpassing Azure Premium SSD v2 managed disks, which delivered 0.75 QPM. This result represents approximately a 93% improvement in analytical query performance when using Pure Cloud Block Store.

The results demonstrate Pure Cloud Block Store's capability to handle intensive analytical workloads. Organizations leveraging Pure Cloud Block Store for SQL Server analytical workloads can therefore expect substantially improved query responsiveness and reduced processing times compared to similar managed-disk configurations.

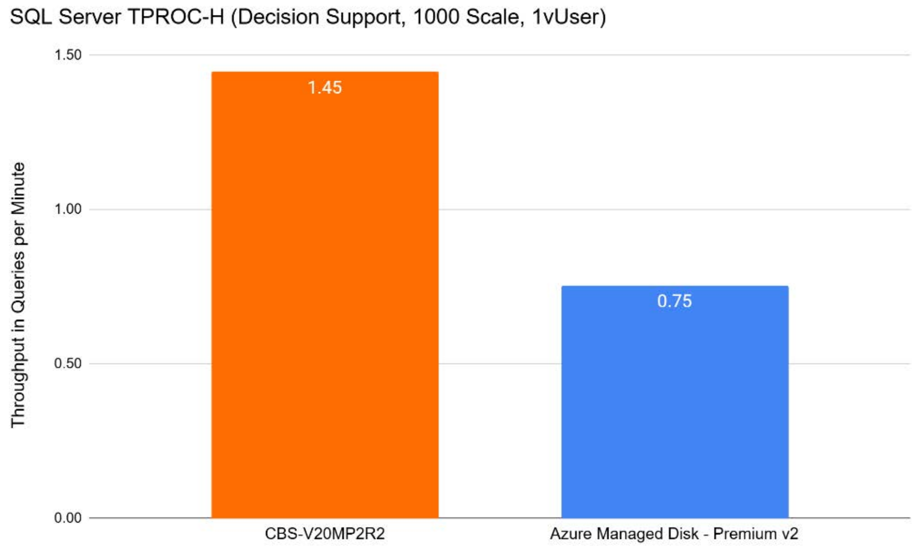

FIGURE 9 TPROC-H performance results comparing Pure Cloud Block Store and managed disk

Deployment Guidance

This section outlines milestone-based deployment guidance for running Microsoft SQL Server workloads on Azure Virtual Machines with Pure Cloud Block Store (Pure Cloud Block Store). Each milestone includes key objectives and links to detailed Pure Storage documentation for step-by-step execution.

Milestone 1: Deploy Pure Cloud Block Store in Azure

Objective: Provision Pure Cloud Block Store using the Azure Marketplace.

1. Deploy Pure Cloud Block Store from the Azure Marketplace in the same region and availability zone as the target Azure VM.
2. Ensure iSCSI endpoints and Pure Cloud Block Store health are validated after deployment.
3. Consult the Pure Cloud Block Store Deployment and Configuration Guide for Azure for more detailed configuration and prerequisite information that should be configured during or before the deployment.
   - The deployment could also be done using a PowerShell or Azure command-line interface script for automation.

Milestone 2: Install and Configure Multipath I/O

Objective: Enable reliable, redundant storage connectivity on Windows Server.

4. Install the Multipath-IO Windows feature.
5. Configure Device Specific Module (DSM) settings for Pure Storage arrays.
6. Enable least queue depth (LQD) load balancing.
7. Reference link: Installing Multipath-IO
8. Reference link: Configuring Multipath-IO for Windows Server

Milestone 3:
Configure and Connect Volumes

Objective: Connect volumes from Pure Cloud Block Store and prepare them for SQL Server data and log placement.

9. Connect Pure Cloud Block Store volumes to the Windows Server host using iSCSI.
10. Format volumes with NTFS using 64KB allocation units.
11. Use separate volumes for data and log files.
12. Reference link: Working with Volumes on a Windows Server Host

Milestone 4:
Apply Best Practices for Connectivity and SQL Server

Objective: Optimize connectivity and SQL Server performance using validated configuration patterns.

13. Configure four or more iSCSI sessions per volume for redundancy and throughput.
14. Follow Pure Storage recommendations for volume layout, TempDB placement, and snapshot usage.
15. Apply Windows and SQL Server tuning parameters for I/O and file initialization.
16. Reference link: SCSI Best Practices for Windows Server and FlashArray
17. Reference link: SQL Server Best Practices for Pure Cloud Block Store in Azure

Conclusion

This reference architecture outlines a comprehensive and highly efficient framework for deploying SQL Server in Azure with Pure Cloud Block Store. By leveraging Pure Cloud Block Store as the foundation for persistent storage, organizations can achieve unparalleled performance, scalability, and resilience in their SQL Server environments.

The integration of services native to Azure, in combination with Pure Cloud Block Store advanced capabilities, ensures seamless communication, optimized storage utilization, and reduced operational complexity. With features like ActiveDR for disaster recovery, efficient data reduction, and robust automation, Pure Cloud Block Store provides a future-proof solution for modern database management.

This architecture not only supports high-demand transactional and analytical workloads, it also minimizes costs while maximizing performance and availability. By implementing the strategies detailed in this paper, businesses can confidently deploy SQL Server in Azure, driving innovation and operational excellence in a rapidly evolving digital landscape.

For more information, visit the Pure Storage product page:

Pure Cloud Block Store