

REFERENCE ARCHITECTURE

SQL Server Workloads on Nutanix with Pure Storage

Deploying SQL Server databases on Nutanix AHV with
Pure Storage FlashArray as external block storage

Contents

Executive summary	3
Solution overview	3
Solution benefits	4
Technology overview	4
Nutanix Cloud Infrastructure	4
Pure Storage FlashArray	5
Connectivity	5
Technical solution design	5
Compute layer	6
Network layer	6
Storage layer	7
Operating system layer	8
Solution validation	9
HammerDB benchmark suite	9
Environment	10
Workload configuration	11
Results	11
Deployment guidance	12
Pre-deployment checklist	12
Milestone 1: FlashArray integration	12
Milestone 2: SQL Server VM deployment	13
Milestone 3: Operating system configuration	13
Milestone 4: SQL Server installation and configuration	13
Additional resources	14
Conclusion	14



Executive summary

This reference architecture provides validated guidance for deploying Microsoft SQL Server databases on Nutanix Cloud Infrastructure using Pure Storage® FlashArray™ as external block storage. The architecture leverages the integration between FlashArray and Nutanix AHV to deliver virtual machine (VM)-centric storage management, streamlined operations, and enterprise-grade data services for SQL Server workloads.

Organizations consolidating SQL Server environments onto Nutanix require storage that maintains consistent performance as database density increases. FlashArray scales storage performance independently from compute resources, enabling higher consolidation ratios without the latency degradation that limits workload density on alternative storage configurations. Design validation demonstrates that FlashArray delivers higher consolidation density, higher concurrent user capacity, and lower latency than alternative storage approaches on similar platforms.

This reference architecture covers solution design across compute, network, storage, and operating system layers, with specific configuration guidance for SQL Server on Nutanix AHV. The document includes deployment guidance and validation results that demonstrate the ability of FlashArray to support demanding database consolidation workloads.

Solution overview

SQL Server databases run on Nutanix AHV VMs with Pure Storage FlashArray providing external block storage. This architecture separates storage from the Nutanix compute layer, enabling independent scaling of database capacity and throughput while providing access to FlashArray enterprise data services.

FlashArray integrates with Nutanix Prism for unified management of storage provisioning, VM-level snapshots, and operational monitoring. Each SQL Server VM connects to dedicated FlashArray volumes over NVMe/TCP (see Figure 1), enabling granular application of storage policies, quality of service controls, and data protection at the individual database level.

This architecture applies to SQL Server on Nutanix AHV with FlashArray external storage. It does not address Nutanix-native AOS storage, Nutanix Database Service (NDB), or ESXi-based deployments.




FIGURE 1 Solution overview



Solution benefits

By combining Nutanix compute with FlashArray external storage, this architecture provides the following key benefits:

- **Independent storage scaling**—FlashArray storage performance scales independently from Nutanix compute resources. Database workloads can grow in capacity and throughput without requiring additional AHV nodes or oversized VMs.
- **Consolidation without compromise**—Storage latency remains consistent as SQL Server instance density increases. Organizations can consolidate more database workloads per Nutanix node without degrading performance for individual instances.
- **Unified management through Prism**—FlashArray integrates with Nutanix Prism for storage provisioning, snapshot management, and monitoring. Platform teams manage the combined environment from a single interface rather than switching between management consoles.
- **VM-level data services**—Each SQL Server VM maps to dedicated FlashArray volumes, enabling per-VM snapshots, replication policies, and quality of service controls. Data protection and storage policies apply at the database level, not the cluster level.
- **Inline data reduction**—The always-on compression and deduplication of FlashArray reduce the physical storage footprint for SQL Server databases, lowering capacity requirements without performance trade-offs or administrative overhead.

Technology overview

This section describes the components of the solution.

Nutanix Cloud Infrastructure

Nutanix Cloud Infrastructure (NCI) provides the compute and virtualization layer for SQL Server workloads. In this architecture, NCI operates in a compute configuration where Nutanix AHV hosts VMs while Pure Storage FlashArray provides all persistent block storage via NVMe/TCP.

- [Nutanix AHV](#) serves as the hypervisor, providing VM lifecycle management, live migration, high availability, virtual networking, and network security with Nutanix Flow and built-in business continuity/disaster recovery features.
- Key features for database workloads include VM templates for standardized SQL Server deployments, VM startup policies for controlling multi-tier application recovery order, and Acropolis Dynamic Scheduler (ADS) to mitigate resource contention across cluster nodes.
- [Nutanix Prism Central](#) provides unified management for the environment, enabling multi-cluster management, VM provisioning, monitoring, and operational workflows. The FlashArray integration extends Prism Central to manage external storage provisioning and VM-level snapshots directly from the Nutanix management interface, maintaining a single operational experience for platform teams.

NCI supports industry-standard x86 servers from multiple vendors as well as [Nutanix NX-series platforms](#).

For detailed guidance on configuring Nutanix clusters for FlashArray connectivity, refer to the [Nutanix Cloud Platform with Pure Storage FlashArray Quick Start Guide](#).



Pure Storage FlashArray

FlashArray provides external block storage for SQL Server databases, connecting to Nutanix AHV nodes over NVMe/TCP. Two FlashArray models are applicable to this architecture:

- [FlashArray//X™](#) is the performance-optimized member of the FlashArray family, suitable for most SQL Server deployments. The FlashArray//X series delivers submillisecond latency with effective capacities up to 3.3PB. Models range from FlashArray//X10 for smaller deployments to FlashArray//X90 for larger consolidated environments.
- [FlashArray//XL™](#) provides the highest performance and scale for large SQL Server consolidation or mission-critical deployments requiring maximum throughput.

All FlashArray models share the Purity operating environment, providing consistent data services regardless of model.

Connectivity

FlashArray connects to Nutanix AHV nodes using NVMe/TCP protocol leveraging standard Ethernet infrastructure. NVMe/TCP provides lower latency than iSCSI by eliminating SCSI translation overhead, making it the preferred protocol for performance-intensive SQL Server workloads.

For connectivity configuration details, refer to [Connect Pure Storage FlashArray as External Storage to Nutanix Cluster](#).

Technical solution design

This section provides layered design guidance for deploying SQL Server on Nutanix AHV with Pure Storage FlashArray. The design (see Figure 2) addresses compute, network, storage, and operating system considerations specific to this architecture.

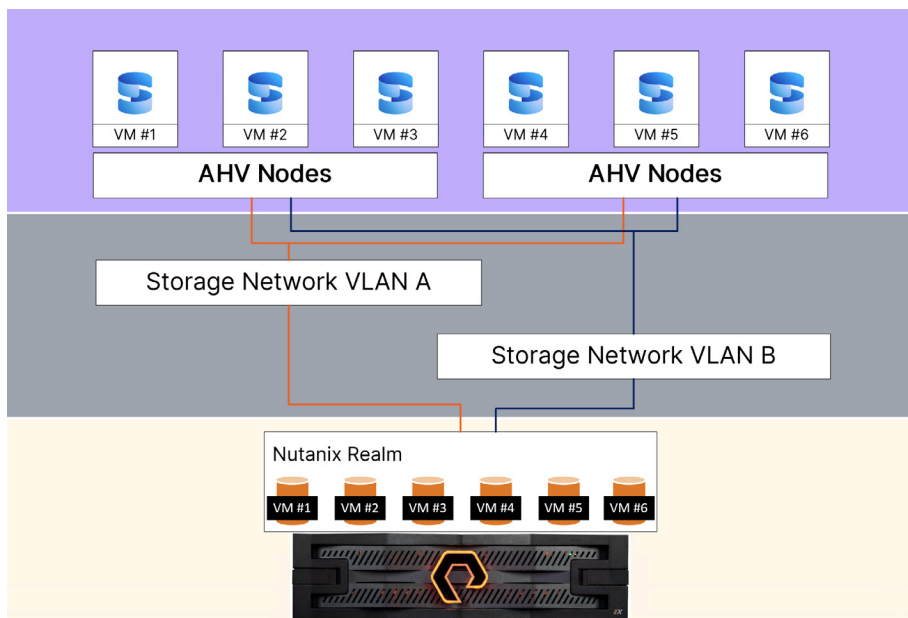


FIGURE 2 Technical solution design



Compute layer

Nutanix AHV nodes provide the compute foundation for SQL Server VMs. FlashArray handles all persistent storage I/O for SQL Server workloads while Nutanix AOS manages cluster operations.

VM sizing considerations

SQL Server VM sizing follows standard virtualization best practices. Memory should be sized to minimize buffer pool misses, with SQL Server maximum memory configured to leave headroom for the operating system. vCPU allocation should align with SQL Server edition core licensing and workload requirements.

For consolidated environments running multiple SQL Server instances, distribute VMs across AHV nodes to balance compute load (ADS can automate this placement). For consolidated environments running multiple SQL Server instances, distribute VMs across AHV nodes to balance compute load (ADS automates this placement to mitigate resource contention).

AHV host considerations

Table 1 summarizes key AHV host configuration recommendations to ensure optimal performance for SQL Server deployments.

Configuration	Recommendation
CPU	High-core-count processors with high clock speed for SQL Server workloads
Memory	Size for VM density targets plus AHV overhead
Local storage	2 x M.2 boot SSDs in RAID1
Network	Minimum 25GbE for NVMe/TCP storage traffic; 100GbE and above is recommended for high-throughput workloads

TABLE 1 AHV host configuration considerations and recommended settings

For current hardware compatibility and specifications, refer to the [Nutanix hardware platforms](#) documentation.

Network layer

The network layer handles three traffic types: storage I/O between AHV nodes and FlashArray, VM client traffic for SQL Server connections, and management traffic for Prism and array administration.

Network segmentation

Separate network segments isolate traffic types and simplify troubleshooting, as shown in Table 2.

Traffic type	Purpose	Protocol	Separation
Storage	SQL Server data and log I/O	NVMe/TCP	Physical NICs, dedicated VLANs, and vSwitches recommended
Client	Application connections to SQL Server	TCP (1433)	Logical (VLAN)
Management	Prism, FlashArray management, monitoring	HTTPS	Logical (VLAN)

TABLE 2 Network segmentation for SQL Server traffic types



Network resilience

If the storage network becomes isolated on a particular host, the Nutanix–Pure Storage integration automatically reroutes VM I/O traffic over the management network, keeping VMs online and operational. This failover capability is not typically available in traditional three-tier solutions.

NVMe/TCP configuration

NVMe/TCP connectivity between Nutanix AHV and FlashArray requires minimal configuration:

1. Configure virtual networking in Nutanix Prism for storage traffic.
2. Create a separate dedicated vSwitch for storage traffic.
3. Ensure FlashArray has multiple network interface cards (NICs) and virtual local-area networks (VLANs) configured for NVMe/TCP. Nutanix storage interfaces and FlashArray NVMe/TCP ports are on the same VLAN with Layer 2 adjacency.

The Nutanix–Pure Storage integration handles connection management, multipathing, and failover automatically. For detailed configuration steps, refer to [Configure Nutanix Cluster for FlashArray](#).

Storage network bandwidth recommendations

Storage network bandwidth should accommodate peak SQL Server I/O requirements. The recommendations in Table 3 apply to the dedicated storage network segment, separate from client and management traffic. For consolidated environments with multiple SQL Server instances, aggregate bandwidth requirements across all VMs when sizing network infrastructure.

Workload profile	Minimum bandwidth	Recommended bandwidth
Single SQL Server instance	25GbE	50GbE (2×25GbE)
Consolidated (4–8 instances)	50GbE	100GbE
High-throughput analytics	100GbE	2×100GbE

TABLE 3 Storage network bandwidth recommendations for SQL Server workloads

Storage layer

The storage layer uses FlashArray for all SQL Server storage, including user databases and TempDB.

Volume design

The Nutanix–Pure Storage integration creates one FlashArray volume per VM virtual disk automatically within a dedicated realm. Volume provisioning, naming, and mapping are handled by the integration through Prism. Create the realm on FlashArray before configuring the integration.

SQL Server VMs should use separate virtual disks for data files, log files, and TempDB (see Table 4). Because the integration maps each virtual disk to its own FlashArray volume, this layout gives you granular control over protection policies. TempDB can be excluded from replication entirely since it gets recreated on SQL Server startup. Data and log volumes can be grouped in a protection group so that snapshots capture a crash-consistent state across both.



VM virtual disk	Purpose
Data files (.mdf/.ndf)	Dedicated virtual disk per database
Log files (.ldf)	Dedicated virtual disk per database, separate from data
TempDB	Dedicated virtual disk; exclude from replication

TABLE 4 SQL Server virtual disk recommendations

Snapshot and data protection

FlashArray snapshots provide crash-consistent, point-in-time copies of SQL Server volumes. For application-consistent snapshots, use SQL Server T-SQL snapshot backup (SQL Server 2022+).

For snapshot behavior specific to Nutanix integration, refer to [Nutanix Snapshot Behavior with Pure Storage FlashArray](#).

Operating system layer

The operating system layer provides the foundation for SQL Server. Both Windows Server and [Linux](#) are supported (see Tables 5 and 6 for configuration best practices).

Windows Server configuration

Configuration	Best practice
File system	NTFS with 64KB allocation unit size for SQL Server volumes
Power plan	High Performance
Antivirus exclusions	Exclude .mdf, .ndf, .ldf, and .bak file extensions and SQL Server directories
Instant file initialization	Enable for SQL Server service account

TABLE 5 Windows Server configuration best practices

Volume structure example (Windows)

- D:\DATA\ Data files mount point
- L:\LOGS\ Log files mount point
- T:\TEMPDB\ TempDB data files mount point
- S:\TEMPDB-LOG\ TempDB log files mount point

Place each database's files in dedicated subdirectories under these mount points. Avoid placing SQL Server database files on the operating system volume.



Linux configuration

Configuration	Best practice
File system	XFS or EXT4 with default settings
Mount options	noatime, nodiratime
I/O scheduler	Blk-mq (none) for NVMe devices
Transparent Huge Pages	Disable

TABLE 6 Linux configuration best practices

For complete SQL Server on Linux configuration guidance, refer to [Microsoft SQL Server on Linux performance best practices](#).

SQL Server configuration

SQL Server configuration follows standard best practices regardless of underlying storage (see Table 7 for recommended configuration settings).

Setting	Recommendation
Max server memory	Leave 4–6GB minimum for OS (more for large-memory systems)
Max degree of parallelism	Start with number of physical cores per NUMA node
Cost threshold for parallelism	50 (adjust based on workload)
TempDB data files	One file per logical processor, up to eight files

TABLE 7 Recommended SQL Server configuration settings

Solution validation

This section presents performance validation results for SQL Server on Nutanix AHV with Pure Storage FlashArray. The testing quantifies consolidation performance at scale, validating that storage latency remains consistent as multiple SQL Server instances run concurrently under demanding workload conditions.

Note: Real-world results vary with schema design, data volume, and concurrency. Customers should run a proof of concept that mirrors their own workload mix.

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) databases. It provides standardized benchmarks such as TPROC-C, widely recognized for evaluating database system performance and scalability.

Validation was conducted using [HammerDB-Scale](#), an orchestration framework based on HammerDB 5.0 that enables consistent, repeatable benchmark execution across multiple database instances.



Environment

Table 8 and Figure 3 show the consolidation performance testing environment used.

Component	Detail
Compute (Nutanix AHV nodes)	Specification: 7 x servers with 2 x Intel Xeon Platinum 8260 CPU (24 cores per socket) at 2.40GHz and 256GB memory Networking: 2 x 100GbE ports for storage connectivity (NVMe/TCP), 2 x 25GbE ports for cluster management Hypervisor: Nutanix AHV
VMs	Operating system: Windows Server 2025 Datacenter VM properties: 2 CPUs, 16 cores per CPU, 144GB memory vDisks: SQL data volumes = 8 x 512GB, log volume = 768GB, all using SCSI bus type Sizing considerations: AHV VMs use default settings with SCSI bus type for all virtual disks; no additional VM-level tuning is required
Storage	Model: FlashArray//XL170 R5 Storage configuration: 4 data packs—40 x 4.56T DirectFlash® Modules Purity OS version: Purity//FA 6.10.3
Microsoft SQL Server	SQL Server 2025 Enterprise Developer Edition TPROC-C database: 8 data files (128GB 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
HammerDB	HammerDB-Scale (based on HammerDB 5.0)

TABLE 8 Test environment configuration for SQL Server consolidation performance benchmarking

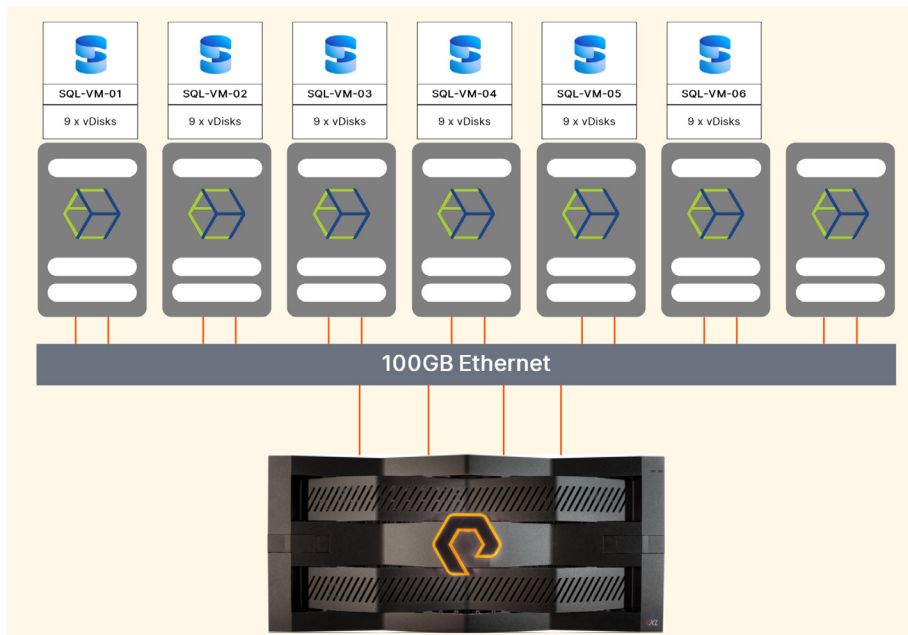


FIGURE 3 Test environment



Workload configuration

To evaluate storage performance under consolidated transactional workloads, multiple SQL Server database instances were deployed on Nutanix AHV, each running as a 44 vCPU, 144GB memory VM with dedicated storage volumes provisioned from FlashArray. Each database was initialized with a unique TPROC-C data set. All workloads were executed in parallel to simulate a highly consolidated, high-intensity OLTP environment. Instances are limited to 64GB of memory to showcase disk pressure and memory optimization.

The primary performance metrics were transactions per minute (TPM) and new orders per minute (NOPM), along with storage read and write latency. Testing scaled from one to five concurrent SQL Server instances to identify the linear scaling range and saturation point.

The following configuration was used for the transaction processing workload:

- Number of warehouses: 10,000
- Virtual users: 200 per instance
- User delay (ms): 500
- Repeat delay (ms): 500
- Total transactions per user: 10,000,000
- Checkpoint when complete
- Minutes of ramp-up time: 5
- Minutes for test duration: 10
- Use all warehouses
- Results averaged across multiple runs with outliers removed

The 200 virtual users per instance represents a demanding consolidation scenario, [significantly higher than workload intensities used in many published benchmarks for similar platforms](#).

Results

Testing scaled from a single SQL Server instance to five concurrent instances, measuring TPM, NOPM, IOPS, and storage latency at each scale point (see Figure 4).

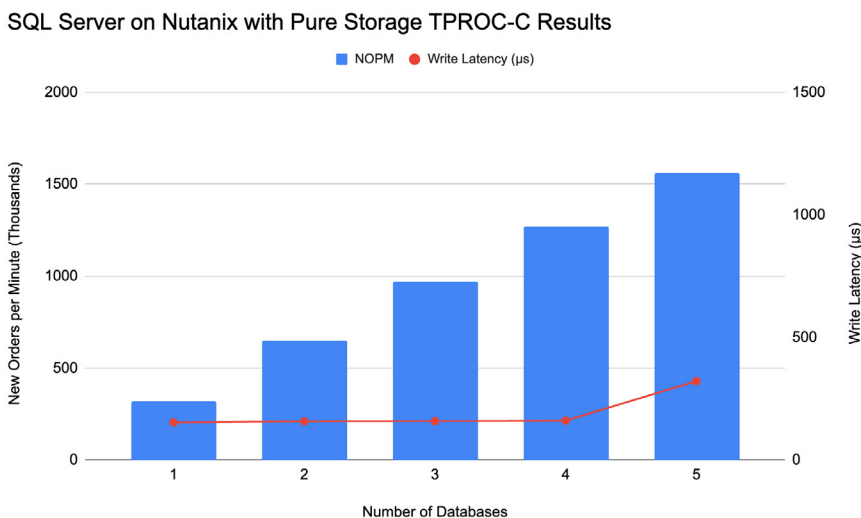


FIGURE 4 Consolidation performance results



Key observations:

- Transaction throughput scaled linearly from one to five instances, adding ~750K TPM per instance.
- Write latency remained under 500µs with no regression through five instances.
- At five instances: 1.5M NOPM (3.6M TPM), 755K combined IOPS, and 322µs write latency.

Deployment guidance

This section provides a milestone-based approach for deploying SQL Server on Nutanix AHV with Pure Storage FlashArray.

Pre-deployment checklist

Before beginning deployment, verify the following prerequisites are in place.

Infrastructure requirements

- Nutanix cluster running AHV with Prism Central deployed
- FlashArray with sufficient capacity for planned SQL Server workloads and projected growth
- NVMe/TCP storage network configured (dedicated VLANs recommended for storage and management traffic)
- Dedicated storage VLANs and vSwitches configured for NVMe/TCP traffic

For further information on networking best practices, refer to the [Nutanix documentation on network topology with Pure Storage and NCP](#).

Software requirements

- Nutanix AOS version supporting FlashArray integration (refer to [Nutanix compatibility matrix](#))
- FlashArray Purity version supporting Nutanix integration
- Windows Server or Linux ISO for SQL Server VM deployment
- SQL Server installation media and license keys

Planning inputs

- SQL Server instance count and sizing requirements
- Database capacity and growth projections
- RPO/RTO requirements for data protection
- Network IP addressing for storage, client, and management traffic

Milestone 1: FlashArray integration

Configure the connection between Nutanix Prism and FlashArray.

1. In FlashArray, create a dedicated realm and pod for Nutanix AHV volumes.
2. Configure NVMe/TCP network interfaces on FlashArray.
3. In Nutanix Prism, navigate to storage configuration and add FlashArray as external storage.
4. Provide FlashArray management credentials and verify connectivity.
5. Confirm FlashArray appears as available storage in Prism.

For detailed steps, refer to [Connect Pure Storage FlashArray as External Storage to Nutanix Cluster](#) and the [Nutanix deployment guide for Pure Storage](#).



Milestone 2: SQL Server VM deployment

Deploy SQL Server VMs on Nutanix AHV.

1. Create VM from template or new installation.
2. Configure vCPU and memory per sizing guidelines.
3. Add virtual disks for SQL Server storage:
 - System disk (OS and SQL Server binaries)
 - Separate volume for system databases
 - Data disk (user database data files)
 - Log disk (user database log files)
 - TempDB disk (TempDB files)
4. Select FlashArray storage container as storage location for data, log, and TempDB disks.
5. Verify virtual disks are provisioned on FlashArray through Prism.

Milestone 3: Operating system configuration

Configure the guest operating system for SQL Server.

Windows Server

1. Format data volumes with NTFS and 64KB allocation unit size.
2. Create mount points or drive letters for data, log, and TempDB.
3. Set power plan to High Performance.
4. Configure antivirus exclusions for SQL Server directories and file extensions.
5. Enable instant file initialization for SQL Server service account.

Linux

1. Format data volumes with XFS.
2. Configure mount options (noatime, nodiratime).
3. Set I/O scheduler to none for NVMe devices.
4. Disable Transparent Huge Pages.

Milestone 4: SQL Server installation and configuration

Install and configure SQL Server.

1. Install SQL Server.
2. Configure data file locations to use virtual disks backed by FlashArray.
3. Set maximum server memory (leave 4–6GB minimum for OS).
4. Configure TempDB with appropriate file count and location.
5. Set maximum degree of parallelism based on workload requirements.
6. Verify database files are created on correct volumes.



Additional resources

- [Nutanix Cloud Platform with Pure Storage FlashArray Quick Start Guide](#)
- [Configure Nutanix Cluster for FlashArray](#)
- [Nutanix Snapshot Behavior with Pure Storage FlashArray](#)
- [Microsoft SQL Server on Linux performance best practices](#)

Conclusion

SQL Server consolidation projects often stall when storage becomes the bottleneck. Adding database instances increases aggregate I/O demand, latency climbs, and application performance suffers. The architecture validated in this document addresses that constraint directly.

FlashArray delivered 3.6M TPM (1.5M NOPM) across five consolidated SQL Server instances while write latency held steady at 322µs. The storage scaled with the workload rather than against it, giving administrators clear signals for capacity planning.

The integration between FlashArray and Nutanix AHV eliminates the operational complexity that typically accompanies external storage. Storage provisioning, snapshots, and monitoring flow through Prism alongside VM management. If the storage network becomes isolated on a host, VM I/O automatically reroutes over the management network, maintaining availability without manual intervention. For high-performance VMs, volume group load balancing across multiple nodes and network segments can further boost throughput. Each SQL Server instance maps to dedicated FlashArray volumes, enabling granular data protection without cluster-wide policies that protect everything or nothing.

For organizations running SQL Server on Nutanix today, or planning to, FlashArray provides the storage performance headroom to consolidate aggressively without compromising the applications that depend on those databases.

To explore this architecture further, contact your Pure Storage or Nutanix account team, or visit purestorage.com/nutanix.