

Oracle Workloads on Nutanix with Everpure

Deploying Oracle databases
on Nutanix AHV with
Everpure FlashArray as
external block storage

Contents

Executive summary	3
Solution overview	3
Solutions benefits	4
Technology overview	4
Nutanix Cloud Infrastructure	4
Everpure FlashArray	5
Connectivity	5
Technical solution design	5
Compute layer	6
Network layer	7
Storage layer	8
Operating system layer	9
Oracle Database configuration	10
Solution validation	11
HammerDB benchmark suite	11
Environment	12
Workload configuration	13
Results	13
Deployment guidance	14
Pre-deployment checklist	14
Milestone 1: FlashArray integration	14
Milestone 2: Oracle VM deployment	15
Milestone 3: Operating system configuration	15
Milestone 4: Oracle Grid Infrastructure installation	15
Milestone 5: Oracle Database installation and configuration	16
Additional resources	16
Conclusion	16

Executive summary

This reference architecture provides validated guidance for deploying Oracle® databases on the Nutanix Cloud Infrastructure (NCI) platform using Everpure™ FlashArray™ as external block storage. The architecture supports both standalone Oracle Database and Oracle Real Application Clusters (RAC) deployments, leveraging the integration between FlashArray and the Nutanix AHV hypervisor to deliver virtual machine (VM)-centric storage management and streamlined operations for Oracle workloads.

Organizations running Oracle face pressure to reduce infrastructure costs while maintaining the performance that business-critical databases demand. Virtualizing Oracle on traditional platforms often forces trade-offs: either storage that can't sustain production workloads at scale, or RAC shared-storage configurations that block essential operations like snapshots and cloning. FlashArray with Nutanix is designed to address these constraints. Storage performance scales independently from compute resources, enabling higher consolidation ratios without the latency degradation that limits workload density on alternative configurations. For RAC deployments, volume groups backed by FlashArray provide the shared storage required by Oracle Automated Storage Management (Oracle ASM) while supporting the operational flexibility that legacy shared-disk approaches cannot deliver. 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 standalone Oracle Database and Oracle RAC on Nutanix AHV.

Solution overview

Oracle databases run on Nutanix AHV VMs, with Everpure 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 the Nutanix Prism management solution for unified management of storage provisioning, VM-level snapshots, and operational monitoring. For standalone Oracle deployments, each database VM is paired to dedicated FlashArray volumes over NVMe/TCP (see Figure 1), with the hypervisor presenting them to the guest as standard SCSI disks. This enables granular application of storage policies, quality of service controls, and data protection at the individual database level. For Oracle RAC deployments, volume groups backed by FlashArray provide the shared storage required by Oracle ASM, with volume group snapshots enabling crash-consistent, point-in-time copies of clustered database storage independent of individual VM snapshots.

This architecture applies to Oracle Database on Nutanix AHV with FlashArray external storage. It does not address Nutanix-native AOS storage, Nutanix Database Service (NDB), or ESXi-based deployments. Oracle does not provide direct support for Oracle databases running on Nutanix; support for Oracle on Nutanix is provided through Nutanix. See [Nutanix KB-2682](#) for details.

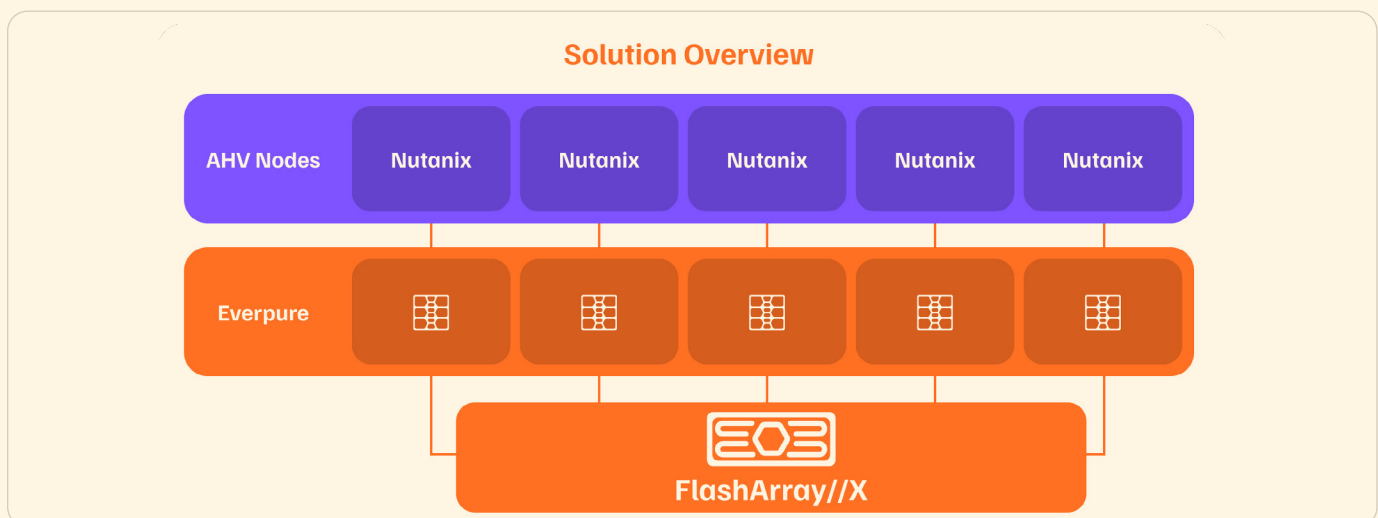


FIGURE 1 Solution overview

Solutions 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. Oracle workloads can grow in capacity and throughput without requiring additional AHV nodes or oversized VMs.
- **Consolidation without compromise**—Storage latency remains consistent as Oracle 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 and monitoring. Platform teams manage the combined environment from a single interface rather than switching between management consoles. Snapshots are managed directly on FlashArray rather than through the Prism integration.
- **Per-VM storage granularity**—Each Oracle virtual disk maps to a dedicated FlashArray volume, enabling independent capacity allocation and performance characteristics for datafiles, redo logs, and archive logs within the same database VM.
- **RAC-ready shared storage**—FlashArray-backed volume groups on Nutanix provide the shared storage Oracle RAC requires, with support for snapshots, cloning, and storage migration. Unlike legacy shared-disk approaches, volume groups enable day-two operations that traditional, multiwriter configurations block. Use Nutanix consistency groups to create crash-consistent snapshots across all vDisks belonging to a database.
- **Space-efficient cloning for test and development**—FlashArray snapshots and clones are metadata operations. Cloning a 10TB production database for test/dev consumes no additional storage until changes are written, enabling database administrators to maintain multiple database copies without capacity overhead.
- **Inline data reduction**—The always-on compression and deduplication of FlashArray reduce the physical storage footprint for Oracle databases, lowering capacity requirements without performance trade-offs or administrative overhead.
- **Enterprise data protection**—FlashArray provides instant snapshots and SafeMode™ protection against ransomware. These capabilities extend to Oracle workloads on Nutanix without additional software or complexity.

Technology overview

This section describes the components of the solution.

Nutanix Cloud Infrastructure

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

- [Nutanix AHV](#) serves as the hypervisor, providing VM lifecycle management, live migration, high availability, and virtual networking with enterprise-grade virtualization capabilities. Key features for database workloads include VM templates for standardized Oracle Database deployments, VM startup policies for controlling multi-tier application recovery order, and Acropolis Dynamic Scheduler (ADS) for automatic workload balancing across cluster nodes.
- [Nutanix Prism](#) provides unified management for the environment. Prism Central enables multi-cluster management, VM provisioning, monitoring, and operational workflows. The FlashArray integration extends Prism 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. Refer to the [Nutanix hardware compatibility list](#) for supported compute platforms.

Note: Nutanix NX-series platforms are not supported for compute deployments with external storage.

For detailed guidance on configuring Nutanix clusters for FlashArray connectivity, refer to the [Nutanix Cloud Platform with FlashArray](#) quick start guide.

Everpure FlashArray

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

- [FlashArray//X™](#) is the performance-optimized member of the FlashArray family, suitable for most Oracle 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 Oracle consolidation or mission-critical deployments requiring maximum throughput.
- [FlashArray//C™](#) is the capacity-optimized member of the FlashArray family, designed for workloads where cost-effective capacity is prioritized over maximum performance. The FlashArray//C series is suitable for Oracle environments with lower I/O intensity, such as development, test, or archive databases.

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 leveraging standard Ethernet infrastructure. NVMe/TCP provides lower latency than iSCSI by eliminating SCSI translation overhead, making it the preferred protocol for performance-intensive Oracle Database workloads.

The Nutanix-Everpure integration presents FlashArray volumes directly to AHV, where they map to individual VM virtual disks. This per-VM granularity enables storage policies, snapshots, and quality of service controls to apply at the database level rather than the cluster level.

For connectivity configuration details, refer to guidance on [connecting Everpure FlashArray as external storage to a Nutanix cluster](#).

Technical solution design

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

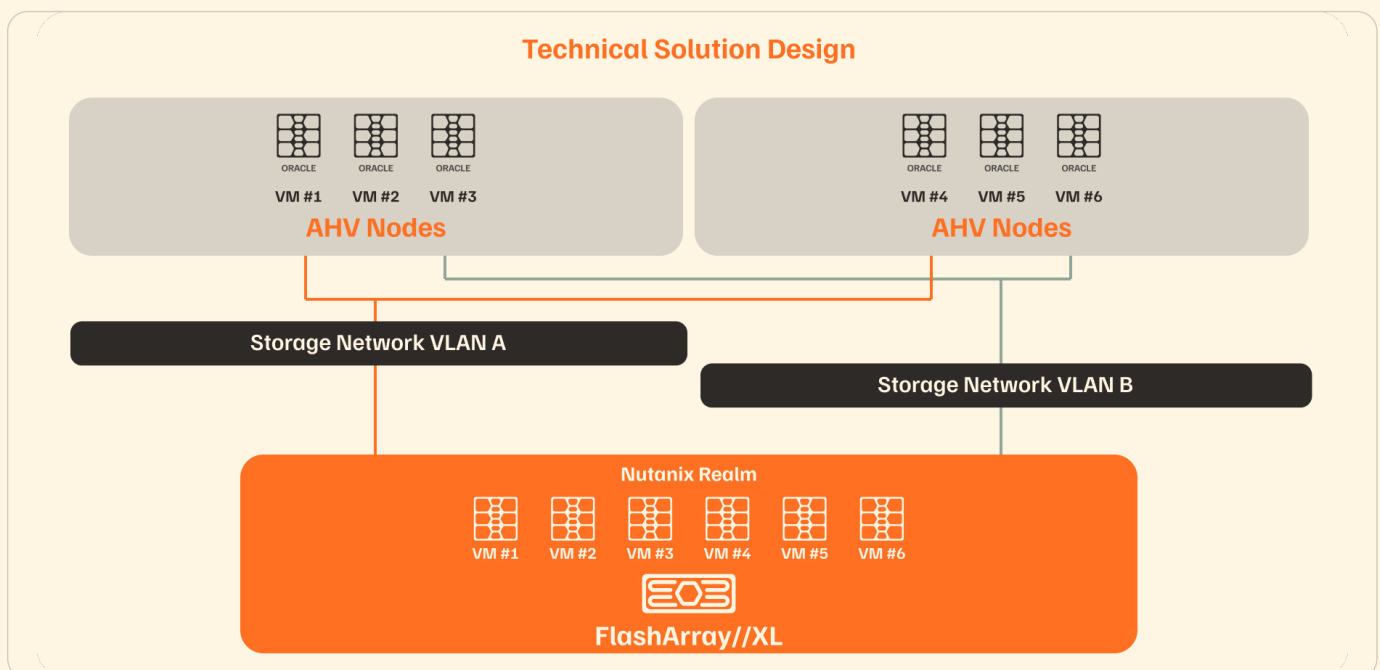


FIGURE 2 Technical solution design

Compute layer

Nutanix AHV nodes provide the compute foundation for Oracle Database VMs. The compute-only configuration dedicates node resources to VM workloads while FlashArray handles all persistent storage I/O.

VM sizing considerations

Oracle VM sizing follows standard virtualization best practices. Memory should be sized to accommodate the system global area (SGA) and program global area (PGA) with headroom for the operating system. vCPU allocation should align with workload requirements and Oracle processor core licensing.

For standalone deployments running multiple Oracle instances, distribute VMs across AHV nodes to balance compute load (ADS automates this placement to mitigate resource contention). For Oracle RAC deployments, configure VM anti-affinity rules to keep RAC nodes on separate physical hosts without pinning them to specific hosts, maintaining high availability while allowing maintenance flexibility.

Note: Nutanix AHV uses soft partitioning from an Oracle licensing perspective. Consult Oracle licensing documentation and your Oracle representative to determine licensing requirements for your deployment.

AHV host considerations

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

Configuration	Recommendation
CPU	High-core-count processors with high clock speed for Oracle Database workloads. For Oracle licensing optimization, different hosts may be preferred depending on core count and clock speed trade-offs.
Memory	Size for VM density targets plus AHV overhead (3.5GB baseline + 2.2% of total system memory per host, plus controller VM allocation).
Local storage	Minimal local NVMe for AHV boot.
Network	Minimum 25GbE for NVMe/TCP storage traffic; 100GbE and above is recommended for high-throughput workloads. For RAC deployments, allocate additional bandwidth for cluster interconnect traffic.

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 four traffic types for Oracle deployments: storage I/O between AHV nodes and FlashArray, VM client traffic for Oracle connections, management traffic for Prism and array administration, and, for RAC deployments, cluster interconnect communication between RAC nodes.

Network segmentation

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

Traffic type	Purpose	Protocol and MTU	Separation
Storage	Oracle datafile, redo, and archive I/O	NVMe/TCP MTU—9000	Physical NICs, dedicated VLANs, and vSwitches recommended
Client	Application connections to Oracle	TCP (1521) MTU—1500	Logical (VLAN)
Management	Prism, FlashArray management, monitoring	HTTPS MTU—1500	Logical (VLAN)
RAC interconnect	Oracle Cache Fusion and cluster heartbeat	UDP MTU—9000	Dedicated VLAN with low latency

TABLE 2 Network segmentation for Oracle traffic types

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-Everpure 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 Oracle I/O requirements. The recommendations in Table 3 apply to the dedicated storage network segment, separate from client, management, and interconnect traffic. For consolidated environments with multiple Oracle instances, aggregate bandwidth requirements across all VMs when sizing network infrastructure.

Workload profile	Minimum bandwidth	Recommended bandwidth
Single Oracle instance	25GbE	100GbE
Consolidated (4–8 instances)	25GbE	100GbE
Oracle RAC (2–4 nodes)	25GbE	100GbE
High-throughput analytics or data warehouse	100GbE	2x 100GbE (bonded)

TABLE 3 Storage network bandwidth recommendations for Oracle workloads

Storage layer

The storage layer uses FlashArray for all Oracle storage. For standalone deployments, the Nutanix-Everpure integration maps each VM virtual disk to a dedicated FlashArray volume. For Oracle RAC deployments, Nutanix volume groups enable shared-disk access across multiple VMs, with storage I/O served by FlashArray. Volume groups provide the shared-storage presentation that Oracle ASM requires while FlashArray delivers the underlying storage performance and data services.

Volume design: Standalone Oracle

The Nutanix-Everpure 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.

Standalone Oracle VMs should use separate virtual disks for different file types (see Table 4). Because the integration maps each virtual disk to its own FlashArray volume, this layout provides granular control over capacity and growth management. This layout assumes one Oracle database per VM. For VMs hosting multiple databases, create additional vDisks as needed to maintain separation between database storage.

VM virtual disk	Purpose
Oracle binaries	Oracle Home and Grid Infrastructure (if applicable)
Datafiles	Dedicated virtual disk for database datafiles
Redo logs	Dedicated virtual disk for online redo logs
Fast recovery area	Dedicated virtual disk for archived redo logs

TABLE 4 Standalone Oracle virtual disk recommendations

Volume design: Oracle RAC

Oracle RAC requires that shared storage be accessible to all cluster nodes simultaneously. Nutanix volume groups enable shared SCSI bus presentation to multiple VMs, with storage I/O served by FlashArray.

For RAC deployments, storage is divided between nonshared and shared virtual disks (see Table 5).

Storage type	Mechanism	Purpose
Nonshared	Per-VM virtual disks via Prism integration	Operating system, Oracle binaries
Shared	Nutanix volume groups	ASM disk groups for DATA, REDO, FRA, and Grid Infrastructure (voting disks, OCR)

TABLE 5 Oracle RAC storage layout

Use ASM external redundancy since FlashArray provides underlying data protection. For Oracle Cluster Ready Services disk groups (voting disks, Oracle Cluster Registry [OCR]), normal or high redundancy may be used depending on site requirements for additional ASM-level protection.

When designing ASM disk groups, use the following layout per database:

- Create separate disk groups for DATA, REDO, and FRA to enable independent management.
- Size individual disks within each disk group consistently for balanced ASM striping.
- Use ASM external redundancy since FlashArray provides underlying data protection.

Snapshot and data protection

FlashArray snapshots provide crash-consistent, point-in-time copies of Oracle volumes. For standalone deployments, coordinate snapshots across datafile and redo log volumes to ensure recoverability.

For application-consistent snapshots, place the database in hot backup mode before triggering storage snapshots. Oracle Recovery Manager (RMAN) can be used to coordinate backup mode with scripted snapshot operations.

For snapshot behavior specific to Nutanix integration, refer to the guidance on [Nutanix snapshot behavior with Everpure FlashArray](#).

Operating system layer

The operating system layer provides the foundation for Oracle Database. Oracle Linux is the recommended operating system for Oracle workloads, though other Oracle-supported Linux distributions including Red Hat Enterprise Linux (RHEL) and SUSE Linux Enterprise Server (SLES) are also supported. The Everpure-Nutanix integration has no operating-system-specific dependencies that would preclude any Oracle-certified Linux distribution. See Table 6 for configuration best practices.

For compute-specific best practices, refer to the [Nutanix Oracle best practice guide](#).

Linux configuration

Configuration	Best practice
File system	XFS or EXT4 with default settings for non-ASM storage (ASM manages raw devices directly)
Mount options	discard, noatime, and nodiratime for file system mounts
I/O scheduler	Blk-mq (none) for NVMe devices
Transparent Huge Pages	Disable
HugePages	Enable and size for SGA
Swappiness	Set vm.swappiness = 1
Oracle preinstall package	Install oracle-database-preinstall for automated kernel parameter configuration

TABLE 6 Linux configuration best practices

Volume structure example (standalone Oracle)

- /u01 Oracle Home and Grid Infrastructure binaries
- /u02/oradata Datafiles
- /u02/redo Online redo logs
- /u03/archive Archived redo logs
- /u03/fra Fast recovery area

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

ASM configuration (Oracle RAC)

For Oracle RAC deployments, Oracle ASM manages all shared storage presented via Nutanix volume groups. Configure ASM disk discovery using either ASMLib or udev rules to ensure consistent device naming across all RAC nodes. See Table 7 for configuration best practices.

Configuration	Best practice
ASM disk discovery	Use ASMLibv3.1 or udev rules (for persistent device naming)
ASM redundancy	External redundancy (FlashArray provides underlying protection)
ASM allocation unit	4MB for OLTP workloads; 8MB or higher for data warehouse
ASM disk group layout	Separate disk groups for DATA, REDO, FRA, and OCR/voting

TABLE 7 ASM configuration best practices for Oracle RAC

For complete Oracle on Linux configuration guidance, refer to the [Oracle Database Installation Guide for Linux](#).

Oracle Database configuration

Oracle Database configuration follows standard best practices regardless of underlying storage (see Table 8 for recommended configuration settings).

Setting	Recommendation
SGA size	Leave 4–6GB minimum for OS (more for large-memory systems)
PGA aggregate target	20% of available memory for OLTP (higher for analytics and sorting workloads)
Database block size	8KB default for OLTP; 16KB or 32KB for data warehouse
Redo log size	Size to achieve 15- to 20-minute log switch intervals under normal load
Processes and sessions	Size for expected connection count plus background processes

TABLE 8 Recommended Oracle Database configuration settings

I/O configuration

Oracle's I/O behavior significantly impacts storage performance. Configure the parameters in Table 9 to optimize I/O for FlashArray storage.

Parameter	Recommendation
filesystemio_options	setall (enables direct I/O and asynchronous I/O; not required for ASM)
disk_asynch_io	TRUE
dbwr_io_slaves	0 (not needed with asynchronous I/O enabled)
db_writer_processes	Scale with CPU count; start with 1 per 4 CPUs for write-heavy workloads

TABLE 9 Oracle I/O configuration parameters

When `filesystemio_options` is set to `setall`, Oracle uses direct I/O (bypassing the operating system buffer cache) and asynchronous I/O (enabling parallel I/O operations). This configuration allows `db_writer_processes` to drive concurrent writes directly to storage. With asynchronous I/O enabled, `dbwr_io_slaves` is unnecessary and should remain at 0.

ASM configuration

Oracle ASM manages all database storage for both standalone and RAC deployments on this architecture (see Table 10).

Setting	Recommendation
ASM disk discovery	Use ASMLib (Oracle Linux) or udev rules (RHEL) for persistent device naming
ASM redundancy	External redundancy (FlashArray provides underlying data protection)
ASM allocation unit	4MB for OLTP workloads; 8MB or higher for data warehouse
ASM disk group layout	Separate disk groups for DATA, REDO, FRA, and OCR/voting

TABLE 10 Oracle ASM configuration parameters

Thin provisioning and ASM rebalancing

FlashArray uses thin provisioning, allocating physical capacity only as data is written. To ensure ASM works efficiently with thin-provisioned storage, do the following:

- **ASMLib v3.1**—Use Oracle ASMLib v3.1 or later to enable thin provisioning support. ASMLib v3.1 provides TRIM/UNMAP functionality that allows ASM to return freed space to FlashArray when data is deleted.
- **Disk group attribute**—Set the THIN_PROVISIONED attribute on ASM disk groups to enable space reclamation. This requires COMPATIBLE.ASM set to 11.2.0.2 or higher.
- **Disk sizing**—Create ASM disks of consistent size within each disk group. ASM distributes data proportionally based on disk size; inconsistent sizing leads to uneven distribution.
- **Rebalancing**—ASM automatically rebalances data when disks are added or removed. Set ASM_POWER_LIMIT to control rebalance speed (1 = minimal impact, 1024 = fastest). For production systems, start with a lower value to limit I/O impact during rebalance operations.
- **Capacity monitoring**—Monitor FlashArray physical capacity consumption rather than ASM reported capacity. ASM reports allocated virtual capacity; FlashArray data reduction may result in significantly lower physical consumption.

Solution validation

This section presents performance validation results for Oracle Database on Nutanix AHV with Everpure FlashArray. The testing quantifies consolidation performance at scale, validating that storage latency remains consistent as multiple Oracle 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 11 and Figure 3 show the consolidation performance testing environment used.

Component	Detail
Compute (Nutanix AHV nodes)	Specification: 7x servers with 2x Intel Xeon Platinum 8260 CPU at 2.40GHz, 24 cores per socket, and 256GB memory Networking: 2x 100GbE ports for storage connectivity (NVMe/TCP), 2x 25GbE ports for cluster management Hypervisor: Nutanix AHV
VMs	Operating system: Oracle Enterprise Linux 8.10 VM properties: 2 CPUs, 16 cores per CPU, 164GB memory vDisks: 4x DATA, 2x REDO, 2x FRA, all using SCSI bus type
Storage	Model: FlashArray//XL170 R5 Storage configuration: 4 data packs—40x 4.56T DirectFlash® Modules Purity OS version: Purity//FA 6.10.3
Oracle Database	Oracle Database 19c (January 2025 release update) TPROC-C database: Single bigfile tablespace, 10,000 warehouses SGA: 96GB PGA: 32GB ASM configuration: DATA (4 disks), REDO (2 disks), FRA (2 disks), all external redundancy Archive Logging was Disabled
HammerDB	HammerDB-Scale (based on HammerDB 5.0) Platform: HammerDB-Scale on Kubernetes Network connectivity: 2x 100GbE to test cluster Location: Separate from Oracle VMs under test

TABLE 11 Test environment configuration for Oracle consolidation performance benchmarking

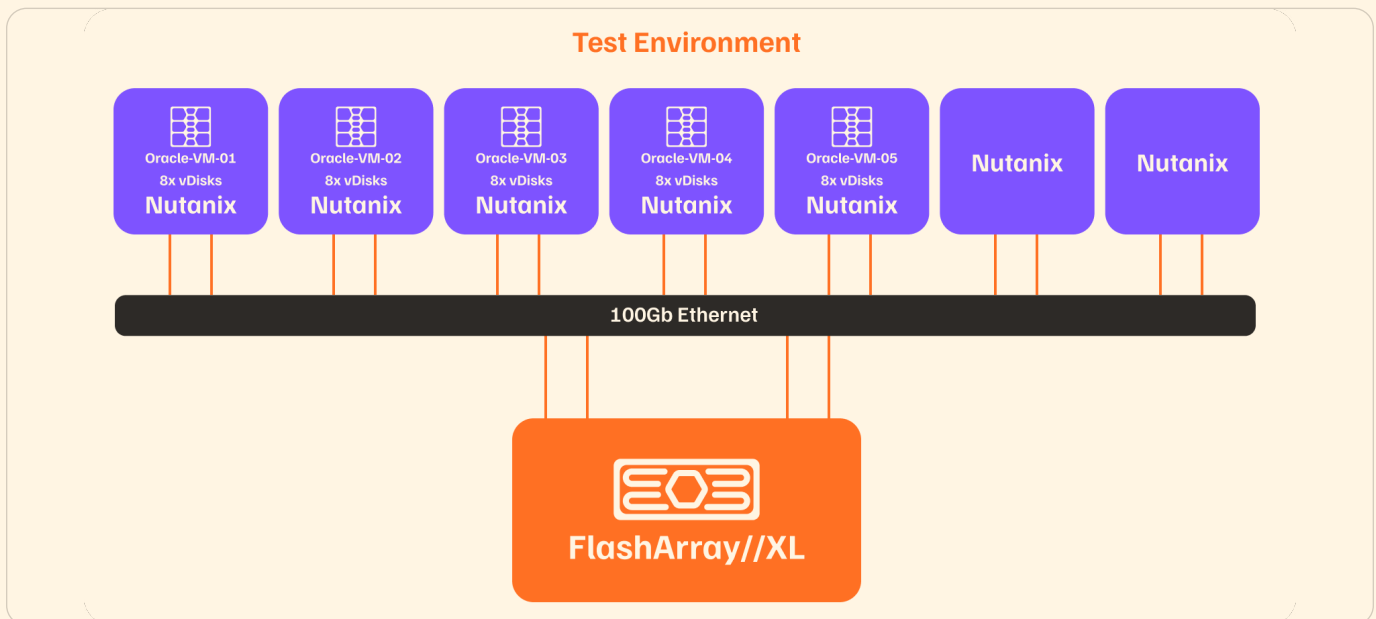


FIGURE 3 Test environment

Workload configuration

To evaluate storage performance under consolidated transactional workloads, five standalone Oracle database instances were deployed on Nutanix AHV, each running as a 32 vCPU, 164GB memory VM with dedicated storage volumes provisioned from FlashArray. Each instance contained a single pluggable database (PDB) initialized with a unique TPROC-C data set. All workloads were executed in parallel to simulate a highly consolidated, high-intensity OLTP environment.

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 standalone Oracle 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.

Results

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

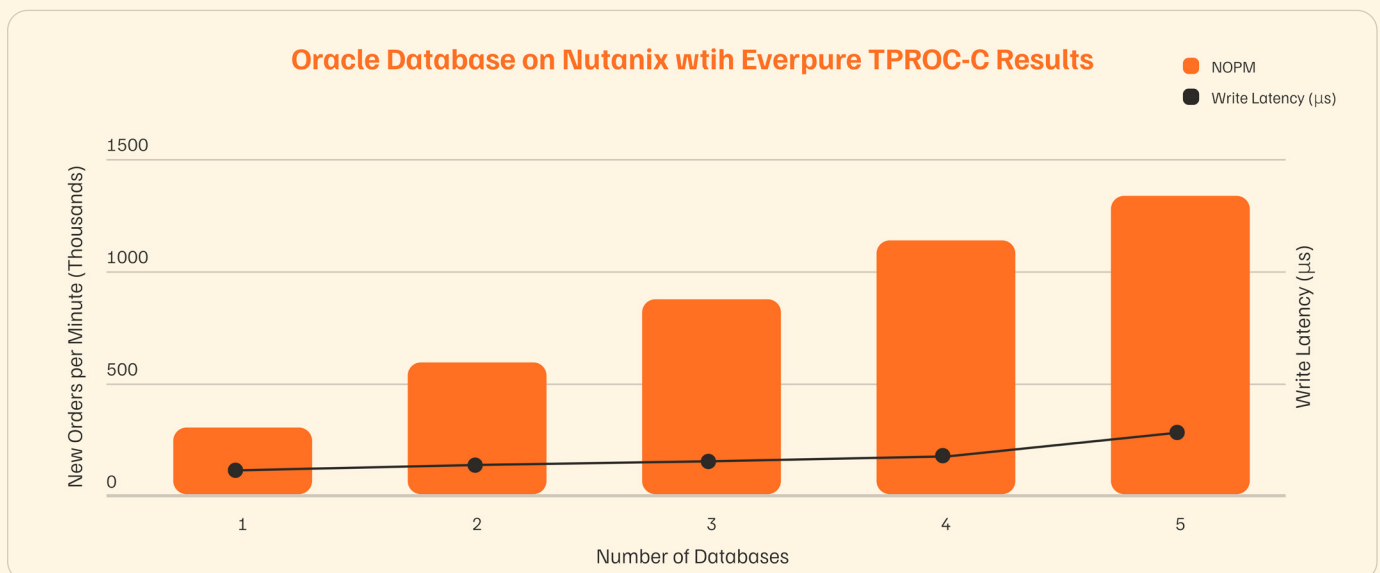


FIGURE 4 Consolidation performance results

Key observations:

- Transaction throughput scaled linearly from one to four standalone instances, adding approximately 280K NOPM per instance.
- Write latency remained flat at 126–161µs through four instances.
- At four instances: 1.15M NOPM with 161µs write latency.
- At five instances, NOPM scaling slowed (193K added vs. ~280K at lower scale points) and write latency nearly doubled to 298µs, indicating storage saturation.

Deployment guidance

This section provides a milestone-based approach for deploying Oracle Database on Nutanix AHV with Everpure 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 Oracle 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 Everpure and NCP](#).

Software requirements

- Nutanix AOS version supporting FlashArray integration (refer to [Nutanix compatibility matrix](#))
- FlashArray Purity version supporting Nutanix integration
- Oracle Enterprise Linux ISO or other supported operating system for Oracle VM deployment
- Oracle Database software and patches
- Oracle Grid Infrastructure software and patches (for ASM and RAC deployments)

Planning inputs

- Oracle 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
- For RAC deployments: cluster interconnect network addressing

Milestone 1: FlashArray integration

Configure the connection between Nutanix Prism and FlashArray.

1. In FlashArray, create a dedicated realm 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 guidance on [connecting Everpure FlashArray as external storage to a Nutanix Cluster](#).

Milestone 2: Oracle VM deployment

Deploy Oracle 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 Oracle storage:
 - System disk (OS and Oracle binaries)
 - ASM disks for DATA disk group
 - ASM disks for REDO disk group
 - ASM disks for FRA disk group
4. Select FlashArray storage container as storage location for ASM disks.
5. Verify virtual disks are provisioned on FlashArray through Prism.

Milestone 3: Operating system configuration

Configure the guest operating system for Oracle Database.

Linux (Oracle Enterprise Linux, Red Hat Enterprise Linux)

1. Install oracle-database-preinstall package (OL) or configure kernel parameters manually (RHEL).
2. Configure ASM disk discovery using ASMLib or udev rules for persistent device naming.
3. Set I/O scheduler to none (blk-mq) for NVMe devices.
4. Disable Transparent Huge Pages.
5. Configure HugePages sized for Oracle SGA.
6. Set `vm.swappiness = 1`.

Windows Server

1. Format non-ASM volumes with NTFS.
2. Set power plan to High Performance.
3. Configure antivirus exclusions for Oracle directories and database file extensions.
4. For ASM on Windows, use Oracle ASM disk stamping for disk management. Disks can be stamped using `asmtool` (command line) or `asmtoolg` (GUI). ASM will identify disks as "Stamped ASM disks" or "Unstamped ASM disks."

Milestone 4: Oracle Grid Infrastructure installation

Install and configure Oracle Grid Infrastructure for ASM.

1. Install Oracle Grid Infrastructure.
2. Create ASM disk groups:
 - DATA disk group with external redundancy
 - REDO disk group with external redundancy
 - FRA disk group with external redundancy
3. Verify ASM disk group creation and availability.

For RAC deployments, install Grid Infrastructure in cluster mode and configure Oracle Clusterware with shared voting disks and OCR on ASM.

Milestone 5: Oracle Database installation and configuration

Install and configure Oracle Database.

1. Install Oracle Database software.
2. Create database using DBCA with storage on ASM disk groups.
3. Configure SGA and PGA memory allocation.
4. Optional if not using Oracle ASM: Set filesystemio_options = setall and disk_asynch_io = TRUE.
5. Configure db_writer_processes based on workload requirements.
6. Verify database files are created on correct ASM disk groups.

Additional resources

- Explore the quick start guide for deploying [Nutanix Cloud Platform with Everpure FlashArray](#).
- Learn how to [configure a Nutanix cluster for FlashArray](#).
- Learn more about [Nutanix snapshot behavior with Everpure FlashArray](#).
- Review the [Oracle Database installation guide for Linux](#).
- Consult the [Oracle Grid Infrastructure installation and upgrade guide for Linux](#) and the [Oracle RAC installation guide for Linux and UNIX](#).

Conclusion

Oracle consolidation projects often stall when storage becomes a bottleneck. Adding database instances increases aggregate I/O demand, latency climbs, and application performance suffers. The architecture validated in this document is designed to address that constraint.

Testing showed FlashArray achieved 1.15M NOPM across four consolidated standalone Oracle instances while write latency held steady at 161µs. The storage scaled with the workload rather than against it, giving administrators clear signals for capacity planning.

The integration between FlashArray and Nutanix AHV helps reduce the operational complexity that typically accompanies external storage. Storage provisioning 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 Oracle RAC deployments, Nutanix volume groups provide shared storage with full support for snapshots and cloning—operations that legacy shared-disk approaches block.

For organizations running Oracle 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 Everpure or Nutanix account team or read more about the [Everpure-Nutanix integration](#).

Visit Our Website

800.379.PURE

