

JULY 2026

How Leading Organizations Optimize for Database Workloads

Simon Robinson, Principal Analyst, and Monya Keane, Senior Research Analyst

Abstract: Modern database workloads require fundamentally different infrastructure, especially for artificial intelligence (AI). The Everpure Platform addresses those requirements through unified data management, policy-driven automation, ultra-low latency, immutable ransomware protection, and zero-planned-downtime architecture, enabling organizations to operationalize AI while maintaining existing database SLAs.

Risk, complexity, and cost threaten to derail enterprise AI ambitions

Databases and database-driven applications support mission-critical applications, from transaction processing and advanced analytics to web-scale, cloud-native apps. As organizations incorporate AI into operations, databases play a pivotal role.

However, IT teams need to deploy and operationalize technologies within existing environments. This creates challenges due to a mismatch between modern database demands and legacy infrastructure operating models. Workload shifts (such as AI vector search in databases) expose the limitations and operational risks of legacy data infrastructure. Legacy models cannot safely support the modernization enterprises need. Specific challenges include:

- **Compounding complexity.** IT leaders struggle with complexity, and AI workloads compound it. According to our research, IT has become more complex for six in ten organizations than it was two years earlier.¹ One pain point is storage, as 62% of respondents said storage complexity slows IT operations and digital initiatives.² Legacy storage environments are fragmented, with different systems for different protocols and storage tiers (e.g., file, block, object, primary, backup), all managed separately. That adds a huge “tax” on TCO. Many database environments rely on fragmented infrastructure stacks for transactional workloads, analytics, backup, disaster recovery, and other operations. Not surprisingly, 63% of IT leaders said storage silos and data sprawl are significant barriers to AI success. Additionally, a combined 62% either agreed (41%) or strongly agreed (21%) that managing the infrastructure for AI is currently more difficult than building the AI models themselves.³

Omdia research finds that most organizations routinely have multiple storage silos deployed, including separate solutions for SAN, NAS, object data, as well as distinct offerings for backup, DR, and archiving—each of which often has its own management tools and frameworks. This adds up to a huge operational burden as we move into the AI era. It will place a premium on a holistic understanding of the storage infrastructure and rapid data access across all organizational data. Such silos are going to create substantial challenges.

¹ Source: Enterprise Strategy Group (now Omdia) Research Report, [2025 Technology Spending Intentions Survey](#), December 2024.

² Source: Omdia Research Report, [The Storage Modernization Imperative](#), April 2026. All research references and charts in this Showcase are from this report unless otherwise noted.

³ Source: Omdia Research Report, [The State of AI Inference Strategies: Optimizing Deployment, Performance, and Impact](#), April 2026.

- **Spiraling risk.** AI offers huge potential but introduces risk. Emerging AI workloads such as vector search place huge capacity and performance demands on infrastructure, raising the risk of failing to meet application SLAs. Either infrastructure can't keep pace with AI demands, or resources must be diverted to AI workloads, impacting other applications. Research found 87% of IT leaders said their adoption of AI will drive or is driving substantial data growth,⁴ and performance was the most commonly cited storage-related challenge as AI workloads move from proof of concept to production. Additionally, performance of data services, such as slow database restores, was reported to be a top-two storage challenge.⁵ In another survey, managing downtime emerged as a key concern, with nearly one in five respondents identifying it as one of their most significant challenges.⁶

Omdia research highlights that IT leaders continue to struggle with delivering key aspects related to data and application availability and recovery, including dealing with planned and unplanned downtime. Recent research found that ensuring the availability of mission-and business-critical data is a significant, top four challenge for block storage (e.g., SAN) environments. In particular, the performance of data services—such as snapshots and replication, essential tools in meeting downtime SLAs—is consistently highlighted as a top-four challenge for primary block storage environments, with more than a third of respondents (37%) highlighting this as a significant concern.⁷

- **Cost challenges.** Organizations are under pressure to invest aggressively in new technologies, but IT budgets are stretched due to combating cyber threats, paying for virtualization licensing, and supporting AI initiatives that may require expensive GPU infrastructure, with 65% of IT leaders saying storage performance bottlenecks GPU utilization. Omdia research shows implementation-related cost is a major concern with AI adoption.⁸ IT organizations struggle to onboard AI technologies without creating additional infrastructure or data silos that compound complexity and drive up costs.

AI marks a key decision point for IT infrastructure leaders

AI vector search capabilities are increasingly being embedded directly into mainstream database management systems (DBMSs) such as Oracle (with Oracle 26ai) and SQL Server (with the release of SQL Server 2025). Those vendors are embedding capabilities such as vector search, retrieval-augmented generation (RAG), and improved hybrid cloud capabilities into their offerings. An AI vector database is where the “embeddings”—numerical representations of raw text, images, and video that capture the data's meaning—are created and stored. This enables similarity and semantic searches to take place to support user or agent queries and prompts.

AI vector databases present challenges to IT teams charged with operationalizing those workloads. According to Omdia research, 60% of IT leaders said AI integration was already making their data environment more complex,⁹ and they don't want to stand up separate infrastructure stacks to manage AI workflows such as vector search databases. New workloads need to be integrated into their existing infrastructure and management frameworks.

⁴ Source: Enterprise Strategy Group (now Omdia) Research Report, [The Critical Role of Storage in Building an Enterprise AI Infrastructure](#), September 2025.

⁵ Source: Enterprise Strategy Group (now Omdia) Complete Survey Results, [Data Readiness for Impactful Generative AI](#), April 2025.

⁶ Source: Omdia Research Report, [Modernizing IT Operations in the Agentic AI Era](#), May 2026.

⁷ Source: Omdia Complete Survey Results, [IT Modernization Report: Cloud, Applications, and Infrastructure Amid the AI Revolution](#), June 2026.

⁸ Source: Omdia Research Report, [Modernizing IT Operations in the Agentic AI Era](#), May 2026.

⁹ Ibid.

New versions of databases are relatively frequent, and IT organizations typically have established processes to ensure upgrades complete with minimal disruption. However, organizations contemplating migrating or upgrading to Oracle 26ai and/or Microsoft SQL Server 2025 must make a high-stakes decision because those modernization choices have a direct influence over operational simplicity, downtime exposure, and the ability to safely adopt next-generation workloads such as AI.

New database capabilities can place extra strain on infrastructure that is often nearing its breaking point. The performance characteristics and infrastructure requirements of a vector database can be very different from those of traditional relational databases, yet multiple OLTP, JSON, graph, and vector searching may all be running on the same instance. Vector databases typically require lots of RAM to load large index structures. When text is converted into embeddings, the underlying data will expand significantly (without significant compression), throwing traditional storage capacity assumptions out the window.

Additionally, AI workloads' I/O patterns can differ from traditional database workloads. AI is becoming increasingly multi-modal and is being applied to rich media such as video and images as well as text. At the storage layer, that drives mixed I/O patterns rather than predictable block reads; consistent performance under variable loads matters more than peak IOPS. It also requires supporting file systems and object storage in addition to block storage. The net effect of this workload diversity, especially when it runs alongside existing demanding database workloads, places immense strain on the underlying infrastructure and the IT teams responsible for maintaining uptime, performance, and cost SLAs.

Time for a new approach

Given the challenges, it's no surprise many organizations are looking to upgrade their storage infrastructure to handle today's diverse and increasingly AI-driven workloads. When asked which storage capabilities would be critical in deciding whether to retain or change their storage infrastructure in their next planned refresh, the responses from IT decision makers were tied to (see **Figure 1**):

- **Performance.** The top requirement that IT leaders are prioritizing in their next storage refresh is higher levels of performance overall. This is logical given the demands of diverse workloads, especially in the database realm. Here, organizations need performance to be consistent and predictable, providing ultra-low latency, high throughput, and instant data copies to keep databases running at peak speed. Although high performance has always been table stakes for running database workloads, the stakes are raised in the AI era since storage bottlenecks introduce latencies that can leave CPUs and GPUs underutilized and degrade user experience.
- **Simplified management and operations/automation.** Storage complexity is nothing new, but IT managers desire a different approach that replaces a predominantly manually driven storage management process with something simplified, intelligent, and automated. Database workload automation, including replacing manual database storage tasks with automated presets, can eliminate errors, reduce the risk of outages, improve consistency, and enforce storage SLAs across a range of functions (e.g., RPO, ransomware protection, DR readiness, and audit compliance across a dispersed environment). This is important given the need for greater security capabilities within the storage environment.

- **Unified platform across the hybrid cloud.** For too long, storage has existed in silos, each dedicated to a group of applications and managed separately. At a small scale, this is workable, but the proliferation of silos across a distributed environment (i.e., on-premises and in the public cloud) creates a management and efficiency nightmare. IT leaders wish to replace it with a single, unified platform that provides teams with a consistent way to deploy, manage, and monitor all database workloads across all data types.
- **Non-disruptive upgrades and predictable economics.** With storage migrations being a major pain point, IT leaders want to escape that disruptive storage upgrade cycle (which often results in significant amounts of planned downtime), replacing it with something simpler and less costly. This is particularly important given that many organizations are planning significant database upgrades to Oracle 26ai/SQL Server 2025.

Figure 1. Key factors in decisions to change or retain storage infrastructure

What will be the most significant factors in your organization’s decision to retain or change its current storage infrastructure in the next refresh or purchase? (Percent of respondents, n=193, three responses accepted)



Source: Omdia

Database environments place heavy demands on storage infrastructure. When AI workloads, real-time transactions, and large data volumes hit the same traditional storage system simultaneously, it can't keep up, resulting in slow performance, limited scale, higher outage risk, and frustration for IT teams and end users. It's no wonder that IT leaders want differentiated solutions and a new approach, such as that offered by Everpure.

The Everpure Platform is ideally suited to the evolving database landscape

Everpure's approach is fundamentally different from that of other storage vendors, many of which are bolting on functionality and capabilities as the market evolves—a sub-optimal architectural approach that comes with scaling challenges, additional complexity, and greater downtime risk. By contrast, [the Everpure Platform](#) powers an organization's Enterprise Data Cloud, delivering cloud-scale storage and data management to cut operational effort, lower risk and cost, and ensure consistent outcomes everywhere. It's an approach suited to the demanding requirements of modern database applications, offering:

- Continuous availability.
- Nondisruptive changes and migrations.
- Simplified operations across on-premises, cloud, and hybrid environments.

Removing risk from database and storage upgrades with the Everpure Platform

The Everpure Platform unifies data, automates control, and reduces risk across the database estate. Block, file, and object storage, plus the data protection environment, come together in a single data layer so control stays centralized. Policies are defined once and applied everywhere, with management spanning on-premises, public cloud, and hybrid environments. Core capabilities such as protection, security, and performance are built in, not bolted on later.

Consequently, organizations can reduce operational effort, cost, and risk, while making it easier for IT teams to support what the business needs today and future-proofing for what comes next.

With its approach, Everpure accommodates organizations migrating to [Oracle 26ai](#) and [Microsoft SQL Server 2025](#), regardless of their starting point. Following are real-world benefits that arise as a result of using the Everpure Platform.

Overcoming compounding complexity – performance and an efficient footprint

Running or migrating production databases is highly sensitive to latency, and it requires consistently low latency. [FlashArray](#) meets this need, along with providing six-nines availability. NVMe-based FlashArray delivers consistently low latency under mixed-workload pressure, not just at peak, so GPUs and CPUs spend less time waiting and more time executing. This enables continuous database storage performance tuning without increasing operational complexity.

With AI vector search, OLTP transactions and analytics run concurrently on the same instance, while Everpure's unified storage pool prevents resource contention from degrading production workloads. Unlike

legacy tiered architectures that force tradeoffs between performance and capacity, FlashArray maintains throughput consistency across workloads, eliminating the storage bottleneck. According to Everpure, these storage efficiency technologies reduce data stored for vector embedding by up to 50%. Organizations can also consolidate backups and snapshots on [FlashBlade](#), which offers scalable throughput and fast restores.

Overcoming compounding complexity – operational simplicity

Organizations can replace database silos with a simple data management experience across all workloads (on-premises, in the cloud, or in a hybrid manner). The solution provides seamless data mobility, including snapshot replication, while maintaining consistency, integrity, and availability.

Everpure Fusion enables IT teams to easily scale and manage their global Oracle and Microsoft storage with policy-driven automation and intelligent workload placement. Critically, Everpure Fusion enforces database storage SLAs as policy at the point of deployment, not as a remediation step after something goes wrong. Recovery point objective (RPO) targets, ransomware protection via SafeMode snapshots, disaster recovery (DR) readiness, and audit compliance requirements are all defined once and applied automatically across every workload and location. This shifts the operational model from firefighting to proactive assurance, removing the manual configuration steps that introduce risk during database upgrades, migrations, and new workload onboarding. Utilizing Everpure Fusion, storage operates as a pooled resource, not a collection of individual systems, supporting scale, density, and efficiency without increasing costs. This approach supports database administrator best practices by enforcing consistent policies, reducing manual intervention, and improving operational consistency across Oracle Database and Microsoft SQL Server environments.

Versatile Everpure snapshots provide database volume snapshot capabilities, enabling teams to restore original volumes, overwrite the data, create copies on new volumes, replicate asynchronously to another array, or offload to object storage for long-term retention. FlashBlade also delivers S3-compatible object storage for cloud-based offerings such as DBMS_CLOUD and external table workloads under the same data management plan as the primary storage.

Overcoming spiraling risk – data protection and an end to planned downtime

Database data is protected with always-on encryption, always-on logging, and SafeMode—Everpure's immutable snapshot technology that prevents ransomware and insider threats from deleting or encrypting recovery points. Unlike backup copies that can be targeted and destroyed, immutable SafeMode snapshots cannot be modified or removed, even by a storage administrator, providing a guaranteed clean recovery state for mission-critical databases. This is backed by a contractual 99.9999% availability SLA and a zero-planned-downtime guarantee. Together, these capabilities strengthen database disaster recovery, while the contractual 99.9999% availability SLA and zero-planned-downtime guarantee reinforce business continuity.

When combined with the Evergreen architecture's non-disruptive controller and capacity upgrades, storage maintenance becomes permanently decoupled from database operations, eliminating the planned outage windows that historically forced DBAs to coordinate costly downtime events around upgrade cycles. Instead, controllers upgrade in place, and capacity expands without data migration.

Overcoming cost challenges, now and in the future

The Everpure Platform offers enterprise-grade data services with no additional licensing (replication, inline deduplication/compression) or complexity to drive efficiency, lower energy consumption, and save costs. It replaces “best effort” promises with guaranteed performance, availability, and efficiency. Flexible, as a service consumption options enable IT teams to consume a platform without owning or operating hardware, resulting in a cloud-like experience with capacity scaling on demand and with predictable costs.

Conclusion

The database modernization wave driven by Oracle 26ai and SQL Server 2025 represents far more than a routine technology refresh; it’s a strategic inflection point that will determine which organizations can operationalize AI at scale and which will remain mired in complexity, risk, and escalating costs.

The convergence of AI workloads with mission-critical database operations exposes a fundamental truth. Legacy storage architectures were not designed for the multi-modal, unpredictable I/O patterns and explosive data growth that AI demands. Organizations attempting to bolt AI capabilities onto fragmented infrastructure stacks are discovering that incremental approaches fail when workload characteristics change fundamentally.

The research is unambiguous. When 63% of IT leaders identify storage silos as barriers to AI success, and 65% report that storage bottlenecks are undermining GPU utilization, we’re witnessing infrastructure becoming a limiting factor in AI adoption. This is unacceptable, given the competitive imperative AI represents.

What separates leaders from laggards in the transition is architectural philosophy. Organizations that consolidate around unified platforms designed for workload diversity, policy-driven automation, and non-disruptive operation will operationalize AI while maintaining existing SLAs. Those clinging to siloed, manually intensive infrastructure models will face a compounding tax on every AI initiative: longer deployment cycles, higher operational overhead, greater outage risk, and, ultimately, competitive disadvantage.

The Everpure Platform exemplifies the architectural approach this moment demands. It is purpose-built for heterogeneous workloads, operationally simplified through intelligent automation, and economically predictable through consumption flexibility.

It boils down to a choice between perpetuating technical debt that will constrain AI ambitions or establishing a foundation that turns infrastructure from bottleneck to enabler. In an era where AI capabilities are beginning to impact competitive position, that choice will echo across organizations for years to come.

For more information, explore [Everpure’s database storage solutions](#).

Copyright notice and disclaimer

The Omdia research, data, and information referenced herein (the “Omdia Materials”) are the copyrighted property of TechTarget, Inc. and its subsidiaries or affiliates (together “Informa TechTarget”) or its third-party data providers and represent data, research, opinions, or viewpoints published by Informa TechTarget and are not representations of fact.

The Omdia Materials reflect information and opinions from the original publication date and not from the date of this document. The information and opinions expressed in the Omdia Materials are subject to change without notice, and Informa TechTarget does not have any duty or responsibility to update the Omdia Materials or this publication as a result.

Omdia Materials are delivered on an “as-is” and “as-available” basis. No representation or warranty, express or implied, is made as to the fairness, accuracy, completeness, or correctness of the information, opinions, and conclusions contained in Omdia Materials.

To the maximum extent permitted by law, Informa TechTarget and its affiliates, officers, directors, employees, agents, and third-party data providers disclaim any liability (including, without limitation, any liability arising from fault or negligence) as to the accuracy or completeness or use of the Omdia Materials. Informa TechTarget will not, under any circumstance whatsoever, be liable for any trading, investment, commercial, or other decisions based on or made in reliance of the Omdia Materials.

Get in touch: www.omdia.com askananalyst@omdia.com

