



Which VMs Work Best in Microsoft Azure® Data Processing and AI Operations?

Prowess Consulting conducted testing in Microsoft Azure to compare the performance of VMs powered by Intel® Xeon® processors and those powered by AMD EPYC™ processors for large data volumes, AI deployment, and high-volume daily workloads.

Executive Summary

As enterprise demand for computing power grows, processor efficiency has become a critical factor in cloud performance and cost control. Traditional CPU benchmarks often focus on artificial peak performance, which might not accurately represent how systems handle real-world, production-scale workloads. To address this gap, this study evaluates processor performance in practical scenarios such as large-scale online transaction processing (OLTP) databases, real-time analytics, and AI-driven decision engines.

This report compares Microsoft® SQL Server® workloads running in Microsoft Azure® across multiple virtual machine (VM) types and processor generations. The study tested Azure E-series VMs powered by 5th Gen and 3rd Gen Intel® Xeon® Scalable processors against E-series VMs powered by 3rd and 4th Generation AMD EPYC™ processors. Additional testing was conducted on FX-series FX48mds v2 and FX4mds v2 VMs, comparing performance between 5th Gen Intel Xeon processors and 2nd Gen Intel Xeon Scalable processors.

The findings show that Azure Eds v6 VMs with 5th Gen Intel Xeon processors deliver measurable gains in compute and AI performance, stronger power efficiency, and integrated security enhancements. These advantages make them a compelling choice for organizations seeking scalable, cost-effective infrastructure that is well-suited for future workloads. By delivering more performance per vCPU and higher workload consolidation, organizations can significantly reduce costs and more effectively maximize the long-term total cost of ownership (TCO) benefits of Intel Xeon processors.

Highlights

Compared to Microsoft Azure® Eads v6 VMs powered by 4th Gen AMD EPYC™ processors, Azure Eds v6 VMs powered by 5th Gen Intel® Xeon® processors provide up to:

21%
higher throughput
(new orders per minute [NOPM])

20–30%
lower total cost of ownership (TCO) for high-volume OLTP workloads

17%
higher throughput per dollar for SQL Server workloads

Benchmarking Azure: Intel vs. AMD in Real-World OLTP Workloads

Enterprises are placing growing emphasis and investment on cloud-based infrastructure to support critical database and analytics workloads, but variations in performance consistency and cost efficiency can restrict their ability to operate smoothly. OLTP systems require high throughput and predictable responsiveness to sustain themselves, making the choice of the right processor architecture essential for organizations that want to balance raw performance, scalability, and TCO in the cloud. Concerns about customer-facing application performance, delivery timelines, and workload balancing can vary across industries. Investing in infrastructure with a high value-to-performance ratio is crucial for long-term scalability.

This study, conducted by Prowess Consulting, compares the performance of SQL Server workloads running in Azure VMs powered by Intel Xeon processors and AMD EPYC processors across multiple generations. Our analysis focused on workloads running on Azure memory-optimized VMs, with special emphasis on OLTP activity to more effectively mirror real-world use cases. To this end, we used HammerDB to simulate real-world, high-volume enterprise database transactions, enabling us to measure performance in categories such as new orders per minute (NOPM), query execution efficiency, and workload scalability.

By grounding our methodology in enterprise-relevant OLTP use cases, we found that processor generation can have a direct impact on both application performance and cost-effectiveness in the cloud. Our findings suggest that organizations running mission-critical database and analytics workloads in Azure can benefit from the performance and efficiency advantages of Azure Eds v6 VMs powered by 5th Gen Intel Xeon processors when compared to prior-generation VMs powered by earlier Intel Xeon processor generations and compared to Eads v6 and Eads v5 VMs powered by current and prior generations of AMD EPYC processors.

How Do We Determine the Best Choice for SQL Server?

The hypothesis guiding this study was that Azure VMs powered by 5th Gen Intel Xeon processors would deliver stronger results than VMs powered by prior-generation Intel Xeon processors and those powered by AMD EPYC processors for SQL Server workloads, particularly across memory-optimized E-series VMs. The evaluation and testing plans focused on OLTP workloads and large-scale data analytics to simulate real-world enterprise database scenarios that rely on consistent throughput, predictable latency, and scalability.

Why Azure

We chose Azure as the test environment for this study because it is one of the world's largest cloud providers, making it representative of real-world customer usage, thereby ensuring that our findings are directly applicable to enterprises running SQL Server workloads in the cloud. The Azure platform's ubiquity makes it an ideal choice for testing the performance and scalability of 5th Gen Intel Xeon processors versus older-generation Intel Xeon processors and 3rd and 4th Gen AMD EPYC processors in real-world enterprise workloads.

Why E-Series

We selected E-series VMs in Azure for this study because these VMs are ideal for memory-intensive workloads such as large relational database servers, data analytics engines, and in-memory caches. These VMs provide a balance of a high memory-to-CPU ratio, large storage bandwidth, and predictable performance, making them ideal for enterprise production scenarios.

By using E-series VMs, we were able to simulate real-world workloads like financial transaction processing, e-commerce order management, logistics planning, and warehouse inventory management, making these VMs ideal for achieving realistic testing results. This approach helped us ensure that the testing environment closely mirrored what customers and businesses often experience in production environments, allowing performance metrics, scaling behavior, and cost-to-performance evaluations to be meaningful, actionable, and accurate.

Why FX-Series

We included FX-series VMs in this study because many enterprises use FX-series VMs to handle real-time workloads, such as analytics, and a new FX-series VM based on 5th Gen Intel Xeon processors was just released. These VMs work best for workloads that require high single-threaded performance, low latency, and a high memory-to-CPU ratio, such as real-time analytics, high-frequency trading, and operational decision engines.

These new VMs benefit workloads that require a high CPU clock speed and a high memory-to-CPU ratio, in addition to workloads with high per-core performance needs and applications that require high single-core performance. FX-series VMs have a higher cost-to-performance ratio compared to memory-optimized VMs, but they provide better performance in scenarios where responsiveness is more critical than raw throughput.

Usage of HammerDB 4.12

HammerDB 4.12 is an open-source, cross-platform benchmarking tool that implements a derived version of the TPC-C® benchmark, known as TPROC-C. This workload simulates OLTP scenarios, making it ideal for assessing the performance of relational database management systems (RDBMSs) under high-volume transactional loads. For more information on HammerDB, refer to [Appendix A](#).

Intel vs. AMD: Comparisons and Performance

Before reviewing the test results, it's important to understand the metrics used to evaluate performance. Each metric represents a different dimension of database and infrastructure efficiency, making each one relevant to our results and conclusions.

- Transactions per minute (TPM):** TPM measures the total number of database transactions processed every minute. In this study, a "transaction" includes the full set of operations that make up an OLTP event, such as order creation, payment, shipping, and inventory updates.
 - Why it matters:** A higher TPM indicates greater workload throughput for industries like e-commerce, banking, and logistics. A higher TPM means processing more orders, payments, or shipments per unit of time.
- Latency:** Latency represents the average response time per transaction, measured in milliseconds (ms). This reflects how quickly the system can complete a transaction from initiation to completion.
 - Why it matters:** Low latency is critical for real-time applications, such as fraud detection, financial trading, and point-of-sale (POS) systems. Even a small increase in latency can have an exponential impact on user experience and downstream systems.
- Cost per transaction:** Cost-per-transaction measures financial efficiency, calculated by dividing the Azure VM operational cost by the number of transactions processed within a given timeframe.
 - Why it matters:** Cost per transaction demonstrates cloud cost efficiency, enabling organizations to compare infrastructure choices on a per-workload basis, which is especially important for workloads with per-core licensing fees (like SQL Server), where higher per-core performance directly reduces licensing overhead.
- Scaling efficiency:** This metric tracks how effectively a system leverages additional compute resources, measured by the performance increase per added vCPU or per-core licensing value.
 - Why it matters:** In cloud environments, scaling is rarely linear. Without strong scaling efficiency, adding more cores might deliver diminishing performance returns, driving up costs unnecessarily.

E-Series VMs: Driving OLTP Efficiency and Scalability

As enterprises migrate mission-critical workloads to the cloud, performance differences between processor families can directly impact cost efficiency, scalability, and customer experience. Database and analytics workloads, particularly OLTP systems, require consistent throughput, low latency, and predictable scalability to ensure smooth operations. Variability in performance can lead to slower transaction processing, increased infrastructure costs, and reduced responsiveness for customer-facing applications.

The goal of this test is to determine which processor family offers the best performance and business value for high-volume OLTP workloads using 5th and 3rd Gen Intel Xeon Scalable processors and 4th and 3rd Gen AMD EPYC processors. By focusing on workloads representative of financial services, e-commerce, logistics, and data analytics, the testing provides a practical view into how processor choice influences both technical performance and TCO.

These tests compare raw throughput but also aim to understand how different processor architectures scale under real-world conditions, how efficiently they handle concurrency and high transaction volumes, and how these factors translate into business value.

We measured performance using HammerDB v4.12, simulating high-volume OLTP workloads with TPROC-C, HammerDB’s derived TPC-C benchmark. We focused on NOPM and TPM as key throughput indicators, testing across 32-, 48-, and 64-vCPU Azure E-series VM configurations. We scaled virtual users in proportion to warehouse counts to identify optimal performance ratios, and all VMs used premium storage v2 for data/log volumes and local storage for TempDB. The environment was configured as follows:

Environment Configuration

- **Platform:** Azure
- **VM image:** Windows Server® 2022 Datacenter: Azure Edition x64 Gen2
- **Database:** SQL Server 2022 Evaluation Edition
- **VM architecture:** x64
- **Security type:** Standard
- **Storage:**
 - Local storage for TempDB
 - Premium SSD v2 for data and log volumes

HammerDB Configuration

- **HammerDB version:** 4.12
- **vCPU configurations tested:** 32, 48, and 64 vCPUs
- **Key settings:**
 - **Trust Server Certificate:** True
 - **Encrypt Connection:** True
 - **Ramp-Up Time:** 2 minutes
 - **Run Time:** 5 minutes
 - **Total Transactions per User:** 10,000,000
 - **Time Driver Script:** Enabled
- **Additional test:** Unique testing performed with varying warehouse counts for the 32-vCPU VM only to determine optimal tuning.

Table 1. Processor families tested for the E-series VM comparison

Vendor	Generation	Microsoft Azure® VMs	Processor Model
Intel	5th Gen Intel® Xeon® processors	E64ds v6, E48ds v6, and E32ds v6	Intel Xeon Platinum 8573C processors
Intel	3rd Gen Intel Xeon Scalable processors	E64ds v5, E48ds v5, and E32ds v5	Intel Xeon Platinum 8370C processors
AMD	4th Gen AMD EPYC™ processors	E64ads v6, E48ads v6, and E32ads v6	AMD EPYC 9004 processors
AMD	3rd Gen AMD EPYC processors	E64ads v5, E48ads v5, and E32ads v5	AMD EPYC 7763 processors

Azure Eds v6 VMs, powered by 5th Gen Intel Xeon processors, delivered enhanced performance compared to Eds v5 VMs, which are powered by 3rd Gen AMD EPYC processors for transactional workloads. This performance advantage was driven by several factors, including superior central processing unit (CPU) cache, faster memory bandwidth, NVMe Express® (NVMe®) storage accelerated by Azure Boost, and high-throughput networking.

These improvements make Eds v6 VMs significantly more efficient for transactional workloads such as OLTP. For OLTP workloads, Intel Xeon processor–based Azure VMs could provide IT teams with confidence that their infrastructure can handle growth without costly overprovisioning or a loss in performance. Here’s a quick summary of the results:

- **Generational improvements:** Azure VMs powered by 5th Gen Intel Xeon processors show 20–22% improvement over those powered by 3rd Gen Intel Xeon Scalable processors.
- **Scaling efficiency:** Intel Xeon processor–based VMs maintain near-linear performance as vCPUs increased.
- **Small database (32 vCPUs):** Current-generation Intel Xeon processors show a consistent 18–22% lead in throughput, making them more suitable for departmental databases and lightweight analytics.
- **Large workloads (64 vCPUs):** 5th Gen Intel Xeon processor–powered VMs deliver up to 9% higher performance than those powered by 4th Gen AMD EPYC processors, and up to 50% higher performance than Azure VMs powered by 3rd Gen AMD EPYC processors, showing Intel’s advantage in scaling high-volume OLTP workloads.

See Figure 1 for test results.

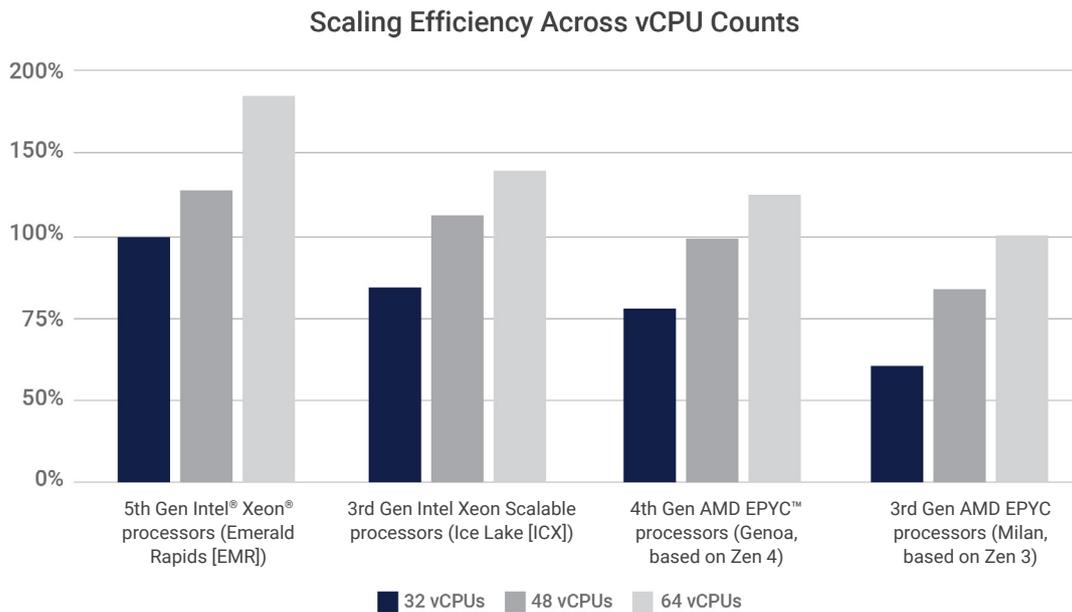


Figure 1. Comparison of vCPU performance between Microsoft Azure® VMs powered by Intel® Xeon® processors and those powered by AMD EPYC™ processors

Note: Percentages are normalized relative to 5th Gen Intel Xeon processors at 32 vCPUs.

FX-Series VMs: Maximizing Real-Time Performance for High-Stakes Workloads

FX-series Azure VMs are engineered for low-latency, real-time workloads that require high per-core performance and rapid transaction responsiveness. Unlike E-series VMs, which prioritize memory-intensive OLTP and analytics, FX-series VMs target scenarios where single-threaded performance and minimal response times drive business outcomes.

We measured performance using HammerDB v4.12, simulating high-frequency, low-latency OLTP workloads with TPROC-C, HammerDB’s derived TPC-C benchmark. We focused on NOPM and TPM as primary performance indicators, testing across 4- and 48-vCPU FX-series Azure VM configurations. We scaled virtual users in proportion to warehouse counts to determine optimal performance ratios for latency-sensitive workloads, and all VMs used premium storage for data and log volumes and local storage for TempDB. The environment was configured as follows:

Environment Configuration

- **Platform:** Azure
- **VM image:** Windows Server 2022 Datacenter: Azure Edition – x64 Gen2
- **Database:** SQL Server 2022 Evaluation Edition
- **VM architecture:** x64
- **Security type:** Standard
- **Storage:**
 - Local storage for TempDB
 - Premium SSD v2 for data and log volumes

HammerDB Configuration

- **HammerDB version:** 4.12
- **vCPU configurations tested:** 4 and 48 vCPUs
- **Key settings:**
 - **Trust Server Certificate:** True
 - **Encrypt Connection:** True
 - **Ramp-Up Time:** 2 minutes
 - **Run Time:** 5 minutes
 - **Total Transactions per User:** 10,000,000
 - **Time Driver Script:** Enabled

Table 2. Processor families tested for FX-series VM comparison

Generation	Microsoft Azure® VMs	Processor Model
5th Gen Intel® Xeon® processors	FX48mds v2 and FX4mds v2	Intel Xeon Platinum 8573C processors
2nd Gen Intel Xeon Scalable processors	FX48mds v1 and FX4mds v1	Intel Xeon Gold 6246R processors

FX v2 VMs, powered by 5th Gen Intel Xeon processors, consistently outperformed the previous FX v1 generation, delivering faster response times and higher per-core throughput. Here’s a quick summary of the results:

- **Performance gains:** FX v2 VMs achieved 30–35% higher NOPM and TPM compared to FX v1 VMs, showing the clear advantage of the latest Intel® architecture for high-frequency workloads.
- **Per-core efficiency:** Higher throughput per core directly reduces per-core SQL Server licensing costs, translating to measurable cost savings for enterprises with heavily licensed database environments.
- **Predictable responsiveness:** FX v2 VMs maintain low-latency performance as virtual users and transaction loads increase, providing IT teams confidence in handling bursty, time-sensitive workloads without overprovisioning.

See Figure 2 for more results.

FX-Series Performance Scaling Across vCPU Counts

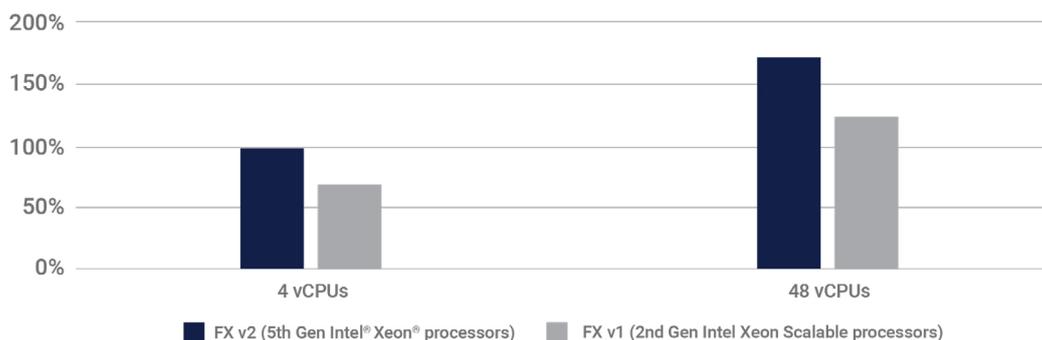


Figure 2. Comparison of vCPU performance between FX-series Microsoft Azure® VMs powered by 5th Gen Intel® Xeon® processors and those powered by 2nd Gen Intel Xeon Scalable processors

Note: Percentages are normalized relative to 5th Gen Intel® Xeon® processor-based Microsoft Azure® VMs at 4 vCPUs.

Which Configuration Is Best for Businesses?

For the tests involved in this study, Prowess Consulting assessed the performance of SQL Server workloads running on Azure VMs powered by different generations of Intel Xeon processors and AMD EPYC processors. Our assessment focused on OLTP workloads using HammerDB to simulate the high-volume enterprise database activity common in industries such as finance, e-commerce, logistics, and analytics. We measured transaction throughput (in NOPM), execution time, and overall workload efficiency across generations to determine both raw performance and price-performance considerations for enterprises seeking to optimize their database workloads in the cloud.

The benchmark tests demonstrated clear, measurable advantages for VMs powered by 5th Gen Intel Xeon processors across both E-series and FX-series Azure VMs, with consistent performance leadership in nearly every configuration tested. Extensive benchmarking in SQL Server demonstrated that Azure VMs powered by 5th Gen Intel Xeon processors deliver superior performance across a wide range of OLTP, big data processing, and data analytics scenarios. These results represent enterprise-scale use cases, such as financial services, supply chain and logistics, e-commerce platforms, and large-scale analytics pipelines.

Cost-to-Performance Analysis

Modern enterprises are under increasing pressure to maximize performance while minimizing TCO. By evaluating processor performance using workload-specific benchmarking, we identified opportunities to achieve better efficiency, reduce licensing costs, and improve customer experiences.

Similar Performance at Lower TCO

This comparison demonstrates how fewer Intel processor cores can deliver comparable performance to a larger 3rd Gen AMD EPYC processors deployment, enabling organizations to do more with less.

By consolidating workloads on fewer virtual cores, enterprises can reduce per-core licensing costs, refresh costs, and simplify overall infrastructure, thereby driving meaningful TCO savings without sacrificing performance.

Key takeaway: With half the vCPUs, the Azure VM powered by current-generation Intel Xeon processors achieves slightly higher performance, creating opportunities for license cost consolidation and operational savings.

Table 3. Comparison between processors with different vCPU counts

vCPU Group	Platform	Normalized NOPM
32 vCPU	5th Gen Intel® Xeon® processor–based E32ds v6	1.04
64 vCPU	3rd Gen AMD EPYC™ processor–based E64ads v5	1.00

Better Performance at Similar TCO

When matched core-for-core, Intel processor–based Azure VMs provide significantly higher throughput, enabling faster transaction processing for OLTP workloads. This improvement translates into quicker insights, better customer experiences, and more responsive OLTP applications, all while maintaining a similar TCO.

Key takeaway: The Intel Xeon processor–based Azure VM delivers up to 21% higher performance at a comparable TCO, ensuring better price/performance ratios for mission-critical OLTP applications.

Table 4. Comparison between processors at the same vCPU counts

vCPU Group	Platform	Normalized NOPM
32 vCPU	4th Gen AMD EPYC™ processor–based E32ads v6	1.00
32 vCPU	5th Gen Intel® Xeon® processor–based E32ds v6	1.21

Transactions per Dollar

While raw throughput is essential, cost efficiency ultimately determines the business value of any Azure deployment. A critical financial metric can be used to evaluate performance:

- **Transactions per dollar** measures how many transactions can be completed for every dollar spent on computing.

This inverse metric shows return on investment (ROI) per dollar spent, making it easy to identify which VM family delivers the most transactions for the least cost.

Example Scenarios

By evaluating processor performance based on cost-to-performance, enterprises can balance throughput performance and operational costs, getting a clearer picture of their TCO and allowing them to prioritize when to allocate budget for upgrades or adjustments based on customer needs.

Consider two 64-vCPU Azure VMs running identical SQL Server workloads, as shown in Table 5. A 64-vCPU E64ds v6 VM powered by a 5th Gen Intel Xeon processor processes 4,680,129 transactions per hour at \$8/hour, giving 585,016 transactions per dollar, whereas an E64ads v6 VM powered by a 4th Gen AMD EPYC processor processes 4,364,818 transactions per hour at the same \$8/hour, giving 545,602 transactions per dollar, which is more than 6% less performance per dollar.

Table 5. Two 64-vCPU Microsoft Azure® VMs running Microsoft® SQL Server® workloads

VM Type	Normalized Performance	Normalized TCO Efficiency
E64ds v6 (5th Gen Intel® Xeon® processor)	(Intel) 1.07	1.07
E64ads v6 (4th Gen AMD EPYC™ processor)	(Intel) 1.00	1.00

E-Series VMs (Memory-Optimized)

- Intel Xeon processor–powered Azure VMs deliver faster, more predictable scaling for high-volume workloads, minimizing the need for additional compute resources.
- At higher core counts, Intel processor–powered VMs sustain performance growth, while AMD processor–powered VMs experience diminishing returns, creating potential inefficiencies in infrastructure spending.
- Generation-over-generation improvements in Intel Xeon processors enhance both throughput and cost-effectiveness, providing a reliable foundation for enterprise resource planning (ERP), logistics, and analytics workloads.

FX-Series VMs (High-Performance, Low-Latency)

- FX v2 VMs powered by 5th Gen Intel Xeon processors enable faster response times and higher per-core efficiency compared to FX v1 VMs, directly reducing SQL Server licensing costs for latency-sensitive applications.
- The improved per-core performance of 5th Gen Intel Xeon processors supports critical workloads such as real-time analytics, fraud detection, and financial trading, delivering tangible operational advantages.

Budget-Friendly Configurations

For budget-conscious or optimization-focused deployments, VMs powered by Intel Xeon processors demonstrated better throughput per core, requiring fewer VMs and cores to meet performance targets. This can translate into lower Azure infrastructure costs and reduced SQL Server per-core licensing expenses, delivering 20–30% TCO savings in high-volume workloads.

In latency-sensitive workloads, FX v2 VMs powered by Intel Xeon processors achieved 30–35% performance gains over FX v1 VMs powered by prior-generation Intel Xeon processors, making them ideal for industries like financial trading, fraud detection, and online gaming, where single-thread responsiveness is critical.

Which Is Best for Enterprise Workloads?

Across all categories tested, the VMs powered by 5th Gen Intel Xeon processors were the clear leaders, offering the best performance per core, the highest throughput per VM, and TCO benefits for OLTP, analytics, and real-time workloads in Azure. Additionally, all Azure VMs based on 5th Gen Intel Xeon processors include the Intel® Total Memory Encryption feature for enhanced security, further adding to the benefits of the Intel platform.

Table 6. Key findings

VM Type	Performance Trend	Cost Impact
E-series with 5th Gen Intel® Xeon® processors	Lowest cost per transaction at 48 and 64 vCPUs	Ideal for high-volume OLTP workloads like ERP and supply chain management (SCM)
E-series with 4th Gen AMD EPYC™ processors	Competitive in midrange, but 20–30% lower throughput than Intel processor–based VM at high core counts	Higher per-core licensing costs and more VMs needed to match the throughput of VMs with Intel Xeon processors
FX-series v2 with Intel Xeon processors	Highest performance per core, 30–35% faster than FX v1 VMs	Premium Microsoft Azure® costs justified for real-time, latency-sensitive workloads like financial trading or fraud detection

AI Acceleration Built-In

Native AI acceleration enables enterprises to handle end-to-end data pipelines, from ingestion and indexing to analytics and machine learning (ML). Intel® Advanced Memory Extensions (Intel® AMX) is the built-in AI accelerator on all Azure VMs based on 5th Gen Intel Xeon processors. This is particularly valuable for predictive analytics, fraud detection, and real-time personalization workloads.

Azure-Optimized VM Families

Several Azure VM families are built on 5th Gen Intel Xeon processors, including:

- General-purpose VMs: Dls v6, Dlds v6, Ds v6, and Dds v6-series
- Compute-optimized VMs: FXv2
- Memory-optimized VMs: Es v6 and Eds v6-series
- Storage-optimized VMs: Ls v4

These VMs deliver the flexibility to match workload requirements, from web front-end interfaces to high-memory OLTP systems.

Broad Workload Applicability

5th Gen Intel Xeon processors perform exceptionally well for:

- E-commerce systems
- Web applications and APIs
- Desktop virtualization solutions
- Big data and analytics pipelines

Conclusion

Prowess Consulting’s testing and analysis found that 5th Gen Intel Xeon processors set the standard for Azure VM performance and cost-efficiency. Whether the goal is maximizing throughput in OLTP workloads, scaling data warehouses, or optimizing real-time applications, for E-series and FX-series VMs, VMs powered by 5th Gen Intel Xeon processors consistently deliver better results than those powered by competing and previous-generation solutions.

For businesses looking to save money, Azure VMs powered by 5th Gen Intel Xeon processors could result in 20–30% savings in total cloud costs while simultaneously improving performance metrics such as TPM and latency. These savings compound at scale, especially for enterprises processing hundreds of millions to billions of transactions monthly.

By deploying Azure VMs powered by 5th Gen Intel Xeon processors, businesses position themselves for:

- High efficiency through optimal throughput per VM
- Future-ready growth with near-linear scalability up to 64 vCPUs
- Real-time responsiveness for mission-critical workloads like fraud detection and high-frequency trading
- Built-in AI capabilities to accelerate next-generation data-driven initiatives

To learn about Azure workloads with 5th Gen Intel Xeon processors, [explore Intel on Microsoft Azure](#).

Appendix A

What Is TPROC-C in Hammer DB?

TPROC-C is the HammerDB version of the TPC-C benchmark, which is the industry standard for OLTP performance testing. While it retains the core elements of TPC-C, TPROC-C is more flexible and easier to use for internal testing and comparative analysis. The workload consists of five transaction types:

- **New-order (45%):** Simulates the entry of new customer orders
- **Payment (43%):** Processes customer payments
- **Delivery (4%):** Delivers orders to customers
- **Order-status (4%):** Checks the status of customer orders
- **Stock-level (4%):** Monitors inventory levels

These transactions mimic the operations of a wholesale supplier, reflecting the complexities and demands of real-world business environments.

Why HammerDB?

We selected HammerDB for this study due to its ability to simulate realistic enterprise workloads using the TPROC-C benchmark. HammerDB is widely recognized for measuring database performance, focusing on key metrics such as:

- **Transaction throughput:** The number of transactions processed per minute
- **Latency:** The time taken to process each transaction
- **Scalability:** The system's ability to handle increasing loads

These metrics are crucial for assessing a database system's performance under high-volume OLTP and analytical workloads, offering insights into its efficiency and capacity.

In this study, HammerDB functioned as a testbed for benchmarking SQL Server performance on Azure E-series and FX-series VMs. Our team configured various TPROC-C workloads with varying warehouse sizes to provide a balanced test environment that reflects typical enterprise workloads. We collected performance metrics such as TPM and response times and analyzed those metrics to assess the impact of different hardware configurations on database performance.



Legal Notices and Disclaimers

The analysis in this document was done by Prowess Consulting and commissioned by Intel.

Results have been simulated and are provided for informational purposes only. Any difference in system hardware or software design or configuration may affect actual performance.

Prowess Consulting and the Prowess logo are trademarks of Prowess Consulting, LLC.

Copyright © 2025 Prowess Consulting, LLC. All rights reserved.

Other trademarks are the property of their respective owners.