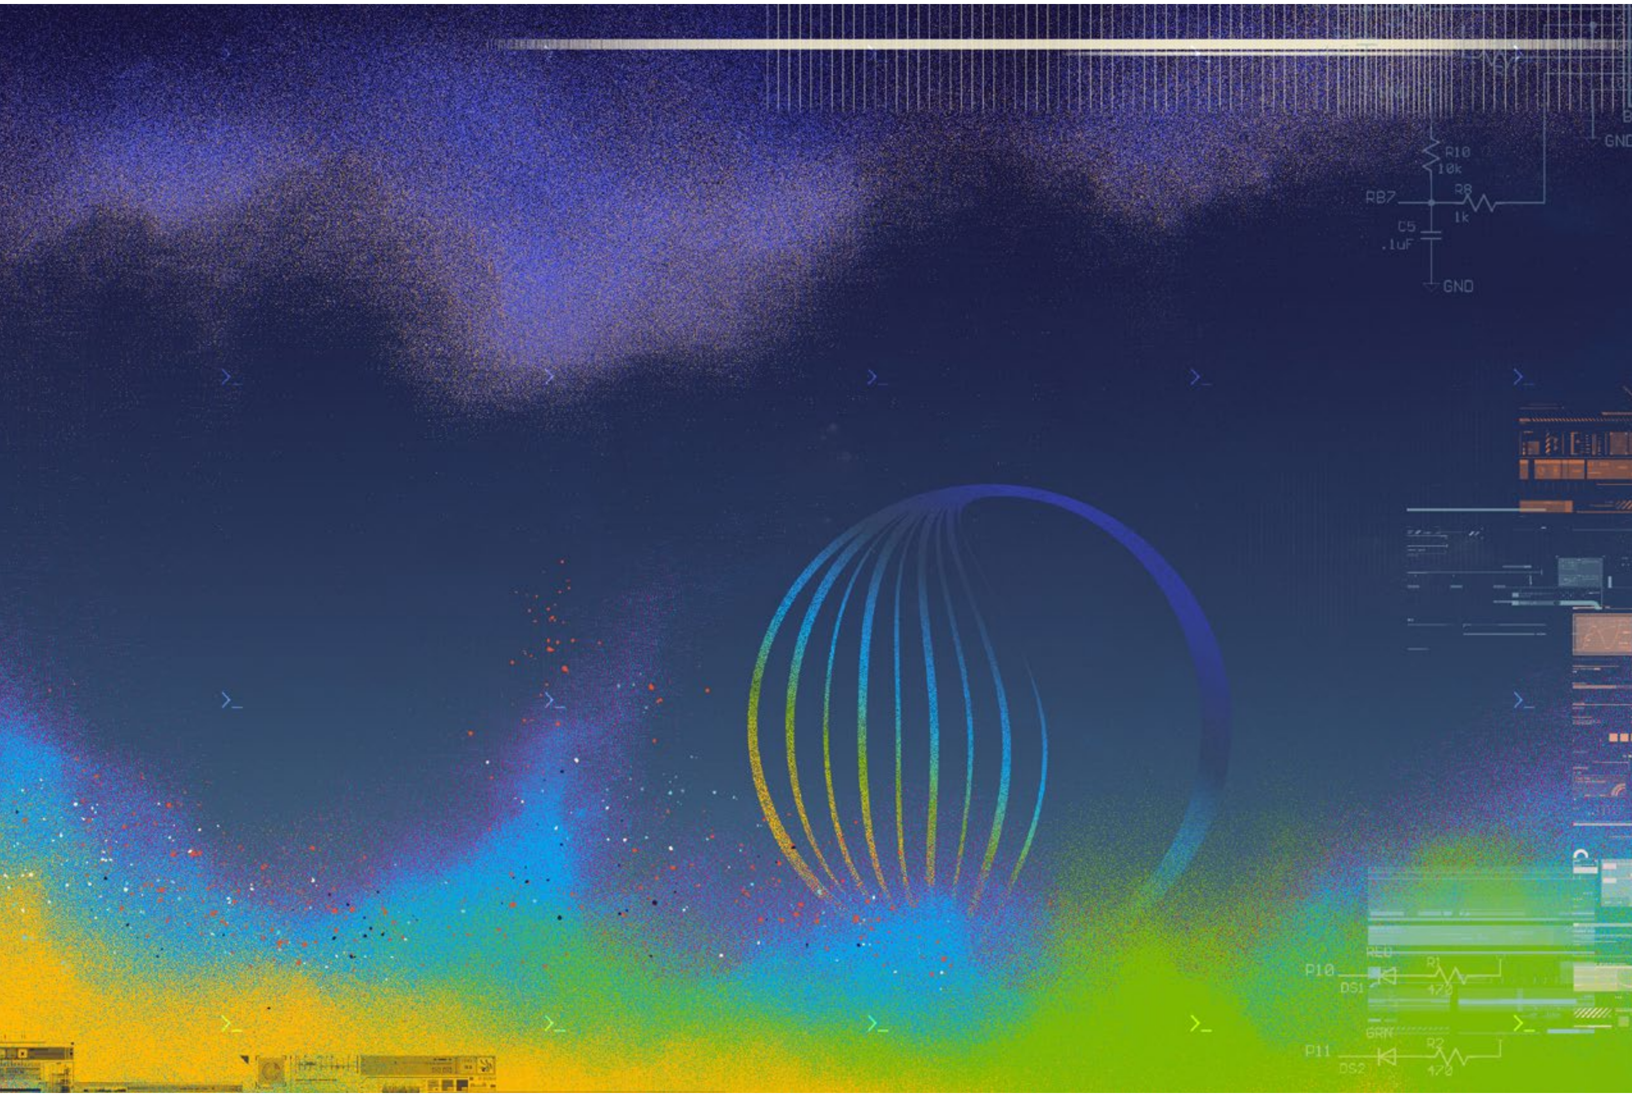

MAIN-STREAM HPC:

A Technical and Economic Evaluation of Azure HBv5 Instances

AUTHORS
Jonathan Fellows
Performance Validation Engineer | Signal65
Russ Fellows
VP, Labs | Signal65

SPONSORED BY
Microsoft

APRIL 2026

Executive Summary

High Performance Computing (HPC) is undergoing a fundamental transformation, transitioning from a specialized tool reserved for academic research into a mainstream engine of industrial innovation. As enterprises across the globe face mounting pressure to accelerate results, ranging from weather modeling to discovery of pharmaceuticals, the demand for high-scale computational throughput has never been higher. In this competitive landscape, the primary constraint is no longer just the cost of hardware, but rather the "Money Value of Time." The ability to reach a solution faster than a competitor or before a critical deadline provides a decisive advantage that often dwarfs the infrastructure investment itself.

The Microsoft Azure HBv5 virtual machine series, developed in collaboration with AMD, represents a strategic response to these business imperatives. By integrating High Bandwidth Memory 3 (HBM3) directly onto custom EPYC silicon, the HBv5 effectively removes the "memory wall" that has historically stalled the most demanding scientific and industrial workloads. This report evaluates the performance gains and Total Cost of Ownership (TCO) implications of migrating to the HBv5 generation, demonstrating how technical breakthroughs translate into significant economic and operational value.

Key Highlights

Up to 4.9x faster time-to-solution for memory-bound HPC workloads

Up to 9x greater sustained memory bandwidth with HBM3 on AMD EPYC

Up to 56% lower infrastructure cost per simulation

Azure has addressed these system constraints through a co-designed partnership with AMD, resulting in the custom EPYC 9V64H processor together with significant networking and storage infrastructure that powers the HBv5 series. Unlike traditional high performance computing nodes that rely on external memory modules, the HBv5 integrates HBM3 directly onto the processor package, delivering an order of magnitude increase in sustained memory bandwidth. This allows researchers to leverage cloud native agility with hardware capabilities that rival or exceed dedicated on premises HPC systems.

The following analysis is based on performance tests conducted by Signal65 across five critical high performance computing benchmarks: OpenFOAM, GROMACS, NAMD, WRF, and CP2K. These workloads represent a cross section of the most demanding scientific applications in use today. The data demonstrates that the HBv5 generation provides significant performance gains, in some cases more than tripling the simulation speed compared to the previous Azure HBv4 Genoa-X instances.

Executive Summary of Technical Findings

The transition from the HBv4 generation, which featured the AMD EPYC 9V33X with 3D V Cache, to the HBv5 series represents a generational leap in efficiency. The key findings of Signal65's evaluation include:

- **Unprecedented Memory Throughput:** The HBv5 series delivers approximately 6.7 to 6.9 Terabytes per second of sustained memory bandwidth, representing a ninefold increase over the HBv4 generation.
- **Massive Parallelism:** By providing up to 368 physical cores per virtual machine, the HBv5 enables massive node level parallelism, reducing the reliance on internode communication for large scale simulations.
- **I/O Throughput Enhancement:** Local NVMe performance has increased to 50 Gigabytes per second for reads and 30 Gigabytes per second for writes, facilitating near instant checkpointing and high-speed scratch space for data intensive runs.
- **Application Speedups:** Performance results show median speedup gains of up to 4.9x for OpenFOAM, 3.2x for WRF, and 2.6x for large scale CP2K simulations.
- **Consistently Superior Efficiency:** Across all tested workloads, the reduction in wall clock time provides value that far exceeds the hardware expense by accelerating research cycles and reducing idle labor costs for highly paid engineering teams.

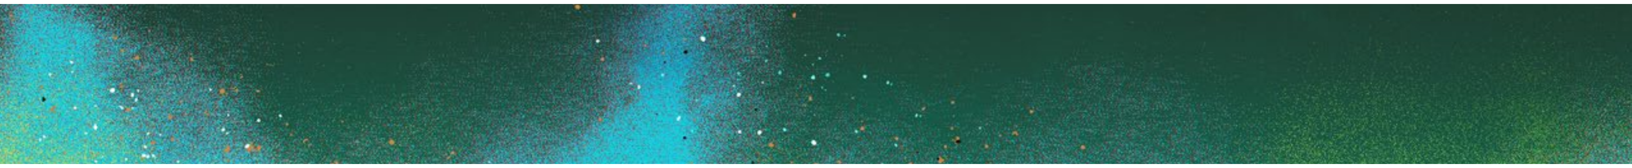

The Strategic Imperative: Beyond Raw Compute

Today, the success of a computational projects is measured by "Time-to-Solution." For decades, the industry focused on peak theoretical FLOPS. But as HPC workloads have become more mainstream, such as those used in computational fluid dynamics (CFD), weather forecasting, genomics and more, the limitation is how fast data can move between the processor and memory.

The Economic Value of Time

The Economic Value of Time is one of the most critical components of a modern HPC strategy. In sectors like Formula 1 racing, strict regulations limit the number of "allocation unit hours" a team can use for simulations. A nearly 5x performance increase from using HBv5 allows an engineering team to complete three high-fidelity design iterations in the same time a competitor completes one, providing a massive edge in aerodynamic optimization. Similarly, in the pharmaceutical industry, reducing the drug discovery timeline by even a single year can save millions of dollars in development costs and secure first-mover market advantages worth billions.

The value of time is most consequential in the context of weather forecasting and disaster response. Utilizing the Weather Research and Forecasting (WRF) model at high resolutions (3km grid spacing or less) is essential for "nowcasting" extreme events like hurricanes or flash floods. Global models often miss these fine-scale phenomena. Because the HBv5 can deliver simulation results up to 3.2x faster than previous generations, emergency management teams can issue evacuation orders hours earlier, directly impacting the ability to save lives and protect property during catastrophic atmospheric events.

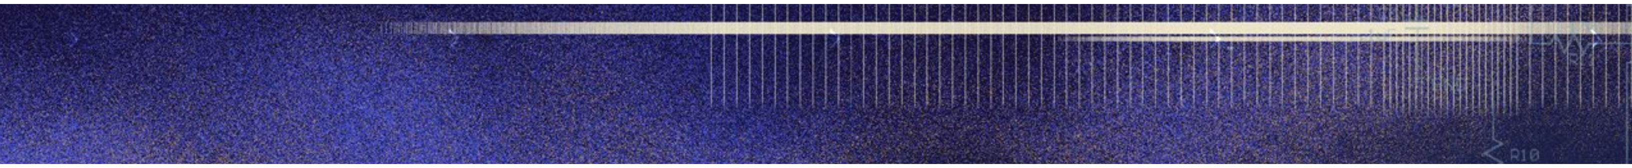

Azure HBv5 Architectural Deep Dive

The move from the previous Azure HBv4 to the HBv5 is not merely an incremental update; it is a foundational redesign of the server node. To appreciate the magnitude of this change, it is necessary to examine the evolution of the custom AMD silicon and the transition from DDR5-based cache strategies to integrated HBM3 subsystems.

The transition from HBv4 to HBv5 marks a pivot from scaling core counts to breaking the "memory wall" through an HBM3-integrated "supercomputer-on-a-chip" architecture. By moving memory directly onto the AMD EPYC 9V64H package, the HBv5 achieves a ninefold bandwidth increase that keeps its 368 cores fully saturated during data-intensive scientific simulations. This architectural shift delivers performance gains across a range of HPC applications, fundamentally redefining the economics of time-to-solution for modern HPC workloads.

The AMD EPYC 9V64H: A Custom HPC Core

The heart of the Azure HBv5 is the AMD EPYC 9V64H, a custom-designed processor co-developed by Microsoft and AMD specifically for the Azure cloud environment. This processor is architecturally distinct from the standard EPYC product line. Industry analysis suggests that the 9V64H is a derivative of the AMD Instinct MI300 series, specifically a CPU-only variant (often referred to as the MI300C) that replaces traditional GPU chiplets with additional "Zen 4" CPU dies.

The EPYC 9V64H is built on a 5nm process node and utilizes the "Zen 4" microarchitecture. While Zen 5 processors have begun to enter the broader market, the Zen 4 design in the 9V64H is optimized for the thermal and power delivery constraints of HBM3 integration. The processor features a base clock of 3.5 GHz, approximately 1 GHz higher than standard high-core-count EPYC processors, with boost frequencies reaching 4.0 GHz across all cores. This high base frequency is critical for HPC workloads where maintaining consistent per-thread performance is essential for parallel scaling efficiency.

Feature	Azure HBv4 (Genoa-X)	Azure HBv5 (Custom AMD HBM3)
Processor Model	AMD EPYC 9V33X	Custom AMD EPYC 9V64H
Microarchitecture	Zen 4 with 3D V-Cache	Zen 4 (HBM-Optimized)
Max Physical Cores	176 cores	368 cores
Memory Subsystem	12-Channel DDR5	Integrated HBM3
Peak Memory Bandwidth	~780-800 GB/s	~6.7 to 6.9TB/s
Memory Capacity	768 GB	432 GB
Interconnect	400 Gb/s NDR InfiniBand	800 Gb/s NDR InfiniBand

Figure 1: *Azure HBv Architecture Comparison*

Breaking the Memory Wall with HBM3

The defining characteristic of the HBv5 architecture is its use of High Bandwidth Memory 3 (HBM3). In the previous HBv4 generation, Microsoft utilized AMD's 3D V-Cache to provide 1.1 GB of L3 cache per VM. While this was highly effective for workloads that fit within that cache window, performance degraded sharply once the dataset required access to the underlying DDR5 memory channels.

The HBv5 bypasses the narrow DIMM channels entirely. HBM3 stacks are placed directly onto the CPU package on a shared silicon interposer. This physical proximity minimizes data travel distance and enables a massive parallel interface. Each HBM3 stack has a 1024-bit wide interface, compared to the 64-bit pipe of a single DDR5 channel. With four custom processors per HBv5 node, the aggregate memory bandwidth reaches 6.7 to 6.9 TB/s, representing a nearly 9x improvement over the HBv4. This massive increase in sustained throughput allows the 368 Zen 4 cores to process data as fast as the ALU can perform operations, effectively dismantling the memory wall for most solvers.

Compute Topology and NUMA Mapping

The HBv5 server features a sophisticated topology designed to maximize data locality. The server is configured with 16 NUMA domains (4 per physical socket), where each domain contains 23 physical cores in the flagship configuration. The underlying hardware consists of 48 Core Chiplet Dies (CCDs), with each CCD containing eight cores and access to a 32 MB L3 cache.

A critical architectural decision in the HBv5 is the reservation of cores for the hypervisor. To prevent the "noisy neighbor" effect and ensure that background virtualization tasks do not interfere with the simulation, Azure reserves 16 physical cores per server. This leaves 368 cores available for the guest virtual machine, ensuring that time-critical MPI ranks are never pre-empted by host-level activities.

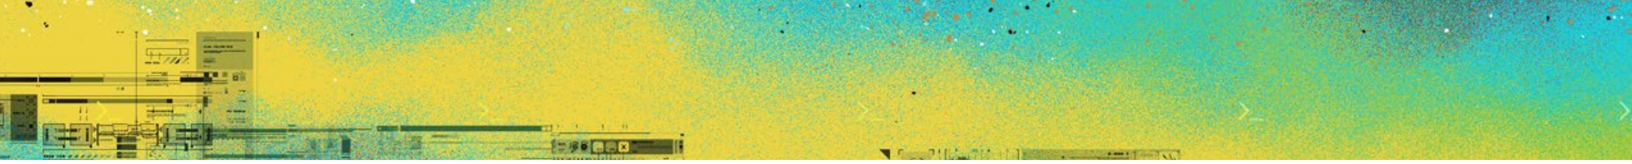

Azure HBv5 System Architecture

While the memory subsystem and custom silicon provide the raw compute power, the scalability and reliability of the HBv5 are driven by major advancements in the supporting infrastructure, specifically in networking offload and I/O throughput.

For distributed HPC simulations, internode communication is often the primary bottleneck. The HBv5 features a significant upgrade to the backend network, utilizing 800 Gb/s NVIDIA Quantum-2 InfiniBand. This is achieved through four 200 Gb/s ConnectX-7 network interface cards (NICs) per virtual machine, with each NIC directly connected to one of the four physical CPUs in the node.

The InfiniBand implementation in the HBv5 utilizes SR-IOV pass-through, allowing the guest VM to interact directly with the hardware and achieve latency as low as 1.25 microseconds. Furthermore, the ConnectX-7 adapters feature onboard processors that support the hardware acceleration of MPI collectives. Operations such as MPI_Allreduce and MPI_Barrier, which are critical for synchronizing thousands of cores across a cluster, are offloaded to the NIC, freeing up the AMD EPYC cores to focus on actual simulation work.

Another advantage of Azure HBv5 instance performance comes from Azure Boost. By offloading some network processing, Boost can help accelerate access to remote, networked storage used for running distributed HPC workloads. By offloading these tasks to a dedicated Data Processing Unit (DPU) or FPGA, Azure Boost eliminates the performance jitter that can plagues shared cloud environments.

For the HPC user, this means that the 368 available cores provide near-deterministic performance. Furthermore, Azure Boost enables the HBv5 to support industry-leading storage performance. Each node includes 15 TB of local NVMe SSD storage, which achieves read bandwidth of up to 50 GB/s and write bandwidth of 30 GB/s. For simulations that generate massive amounts of data for checkpointing or post-processing, such as high-resolution CFD meshes, this storage throughput reduces I/O wait times from hours to minutes.

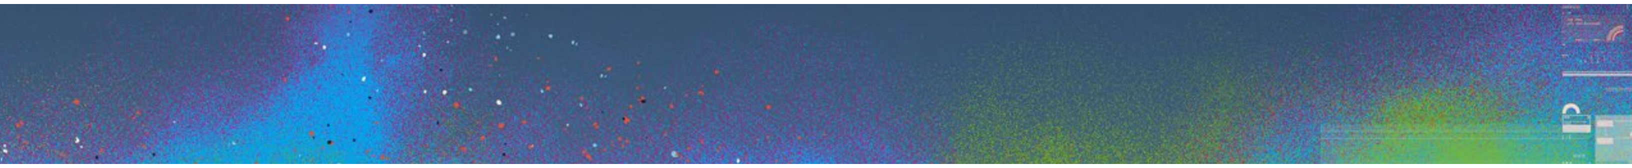

Total Cost of Ownership (TCO) Analysis

In the high-performance computing domain, the value of an infrastructure investment is rarely determined by the hourly rate of the instance alone. Instead, organizations evaluate the Total Cost of Ownership through the lens of "Time-to-Solution" and "Cost-per-Job". The HBv5, while possessing a higher hourly rate than the HBv4, provides significant economic advantages when the accelerated throughput is factored into the modeling.

Infrastructure Cost Modeling

As of March 2026, there are significant cost differences between the older HBv4 and HBv5 instances. The Azure HBv4 (Standard_HB176rs_v4) is priced at $8.64 per hour in the US South Central region, while the HBv5 (Standard_HB368rs_v5) is priced at $23.76 per hour in the same region. This results in an infrastructure cost ratio of approximately 2.75x higher for HBv5 instances.

However, the economic analysis shifts when evaluating the efficiency of a single simulation. For a job that is bound by memory bandwidth, the performance speedup on HBv5 often far exceeds the 3.3x cost increase. For example, in OpenFOAM simulations, the HBv5 is 4.9x faster than the HBv4.

The cost-per-job can be calculated using the formula: $[\ Cost_{Job} = Rate_{Hourly}\ x\ Time_{Solution}\ ]$

For an OpenFOAM simulation that takes 10 hours on an HBv4 (costing $72.00), the same simulation on an HBv5 would take approximately 2.04 hours, resulting in a 56% reduction in infrastructure costs per simulation, despite the higher hourly rate of the virtual machine.

In their research and testing Microsoft presents similar performance differences and provides some price / performance metrics, including for commercially licensed products. When adding in the licensing costs, the reduced run-times of the HBv5 can often further reduce license costs, increasing the advantages of HBv5 systems.[1]

The HBv5's efficiency is particularly impactful for organizations running proprietary software stacks such as Ansys Fluent or Simcenter STAR-CCM+, which are typically licensed on a per-core basis. Because the HBv5 completes complex simulations in a fraction of the time, it reduces the total "core-hours" consumed per job, effectively lowering the cost of commercial software licenses while simultaneously accelerating the time-to-solution. By optimizing the synergy between high-bandwidth hardware and core-based licensing models, the HBv5 provides a superior Total Cost of Ownership for mission-critical engineering and research workflows, allowing teams to maximize their output without a proportional increase in software spend.

Software Licensing Efficiency: The Force Multiplier

For many enterprise HPC users, the cost of software licenses from vendors like Ansys, Siemens, or Dassault Systemes is significantly higher than the cost of the hardware. These licenses are often structured per core. For this reason, HBv4 and HBv5 instances may run with lower CPU core counts, but with all the memory and other resources, to help lower software license costs. This allows organizations to maximize the "Performance per License Dollar," ensuring that every licensed core is operating at its maximum possible throughput.

Because one HBv5 node can perform the work of nearly five HBv4 nodes in memory-bound applications, a company can complete the same volume of work with a much smaller license footprint. For example, achieving a specific level of throughput in ANSYS might require licensing 880 cores on the HBv4 generation (across 5 nodes). On the HBv5, that same throughput can be achieved with only 368 cores on a single node, representing a reduction in the required license count of approximately 58%.

---

[1] Microsoft Azure: HBv5 vs. HBv4 Performance (Retrieved, 20, Apr. 2026).

The Economic Value of Time

The most significant, albeit qualitative, component of TCO is the Economic Value of Time. In competitive industries like Formula 1 racing, pharmaceutical drug discovery, or energy exploration, the ability to finish a simulation faster is not just about saving money—it is about gaining a market advantage.

If an engineering team can run four design iterations in a day on an HBv5 cluster compared to only one iteration on an HBv4 cluster, the research cycle is shortened by 75%. This acceleration allows for more robust design exploration, leading to higher-performing products and a faster path to commercialization. In the context of pharmaceutical research, reducing the design cycle by even a few weeks can be worth millions of dollars in first-mover advantage, far outweighing the premium for high-bandwidth memory hardware.

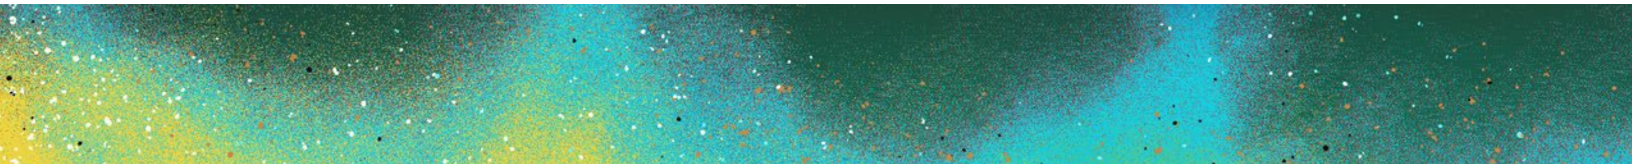

Performance Results and Application Benchmarks

To validate the architectural claims of the HBv5, a series of benchmark tests were conducted across five critical HPC domains. These tests compared the Standard_HB176rs_v4 flagship to the new Standard_HB368rs_v5.

Computational Fluid Dynamics: OpenFOAM Analysis

OpenFOAM is an open-source CFD suite used extensively for automotive and aerospace aerodynamic optimization. CFD is inherently bandwidth-sensitive because it involves iterative sparse matrix-vector multiplications across massive meshes.

The benchmark utilized simplefoam v2406 with a Large Motorbike mesh (v1912) containing 100 million cells. Data I/O was moved to the local ephemeral NVMe disk (/mnt) to ensure that storage latency did not skew the compute results.

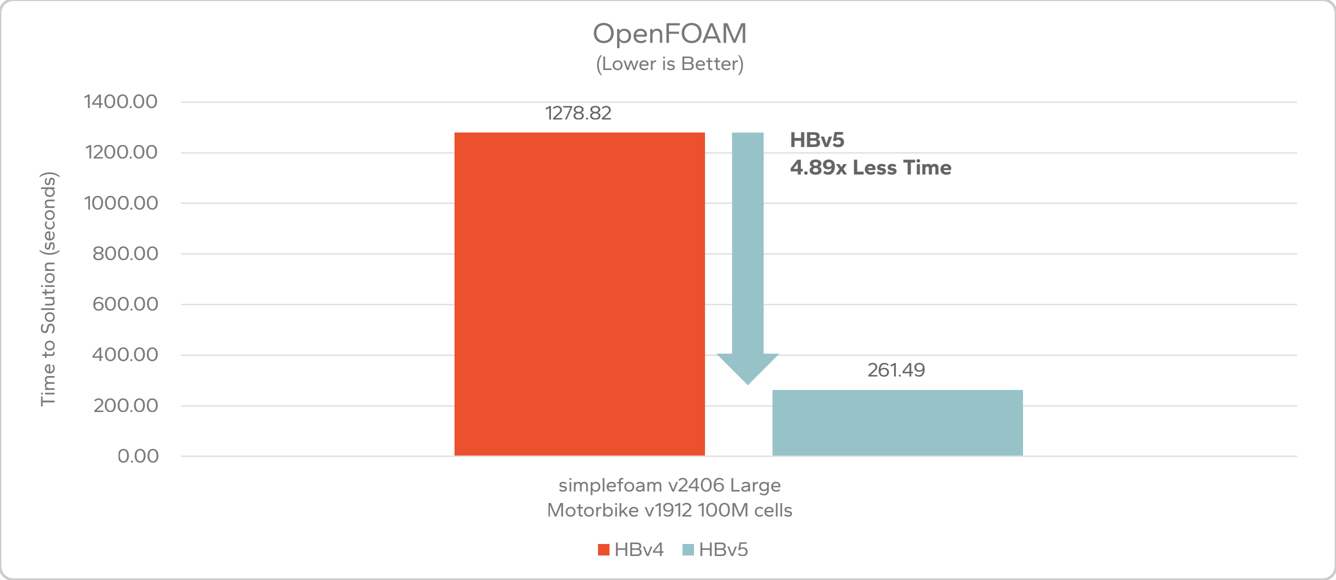

Figure 2: *OpenFOAM Performance*

The 4.9x speedup on HBv5 is a significant result. It indicates that the simulation is scaling almost linearly with the increase in core count and bandwidth. On the HBv4, the 176 cores were likely experiencing "data starvation" as they competed for the DDR5 channels. The high memory bandwidth in the HBv5 provides sufficient headroom to keep all 368 cores operational at maximum capacity, even with a massive 100 million cell dataset.

Molecular Dynamics: GROMACS and NAMD

Molecular dynamics simulations are critical for understanding ligand binding and protein movement in drug discovery. These codes often benefit from both large L3 caches and raw memory throughput.

GROMACS Results

GROMACS was evaluated using two molecular systems: benchPEP and benchRIB.

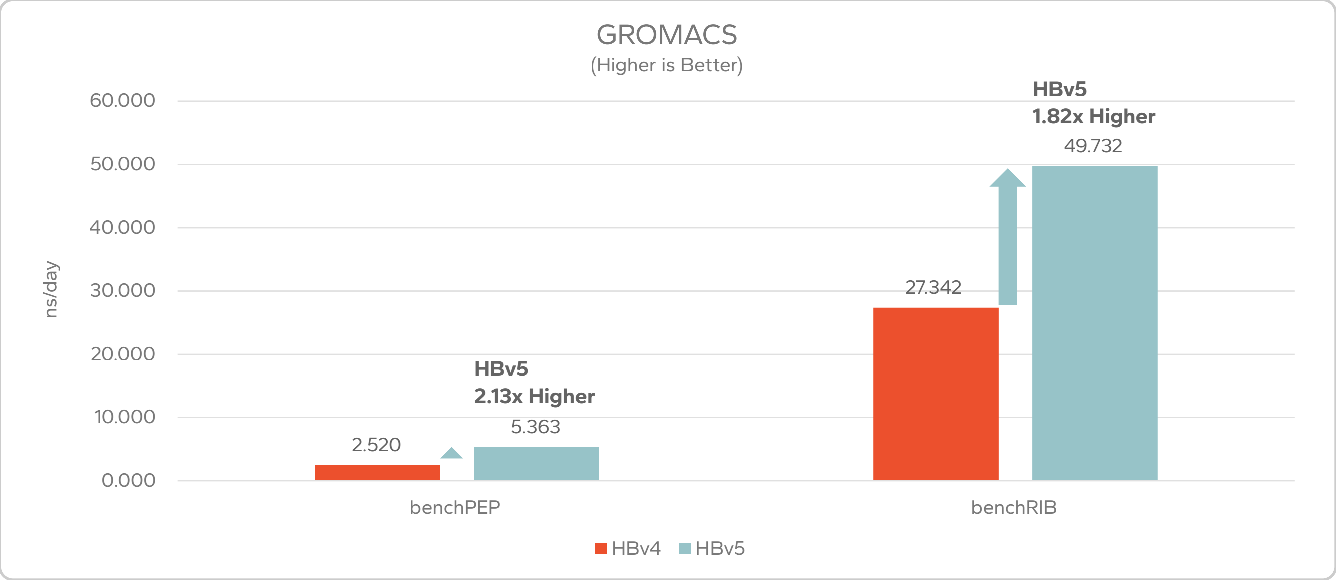

Figure 3: *GROMACS Performance*

The results show a stark contrast between the two cases. The benchPEP system saw a 2.13x improvement, benefiting significantly from the core density and bandwidth of the HBv5. However, the benchRIB system saw only a 1.82x gain. This suggests that benchRIB is more sensitive to cache latency and compute-bound operations rather than streaming bandwidth. The massive L3 cache of the Genoa-X (HBv4) provides a high level of efficiency for smaller, cache-resident datasets, narrowing the gap with the HBM3-powered HBv5.

NAMD Results

NAMD performance was tested using version 3.0.2 binaries.

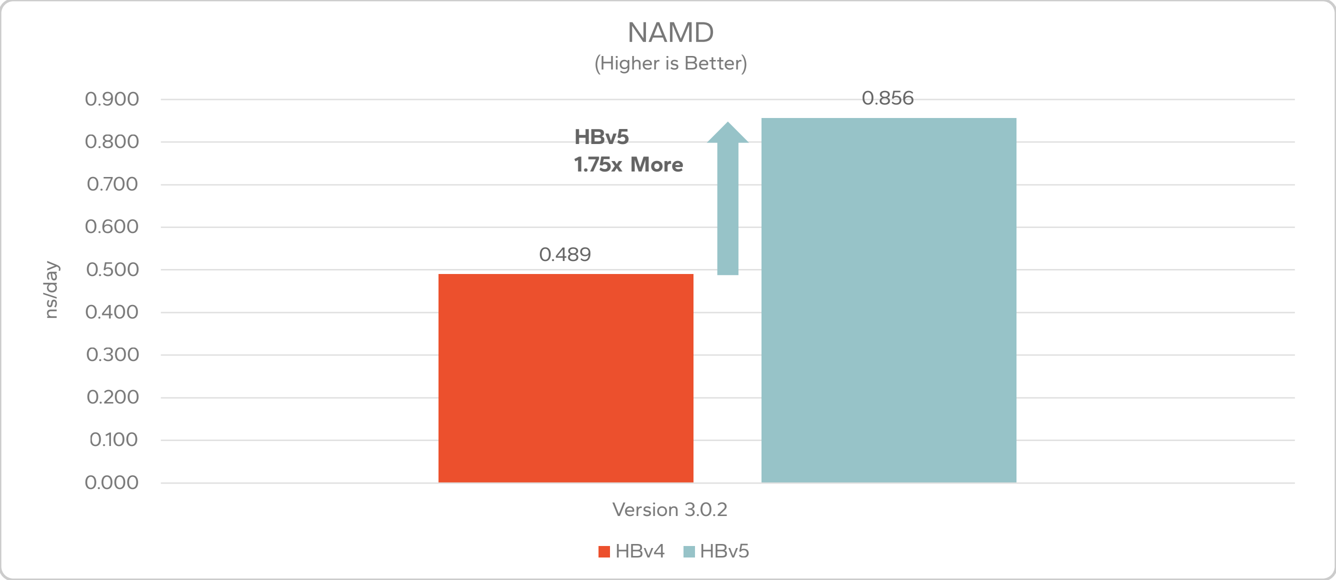

Figure 4: *NAMD Performance*

The speedup factor of 1.75x to 1.82x confirms that the HBv5 delivers substantial value for life sciences. This is likely due to the complex 16 NUMA domain topology. For molecular dynamics codes that are highly sensitive to inter-socket latency, precise process pinning is required to ensure that MPI ranks stay localized to their respective HBM3 pools.

Weather Research and Forecasting: WRF v4.2.2

Weather modeling is a classic "bandwidth-bound" problem. High-resolution forecasting requires capturing fine-scale phenomena like thunderstorms and fog, which global models (at 10km resolution and larger) often miss. These simulations track moisture, air movement, and heat across a kilometer-scale grid.

The benchmark utilized the CONUS NAM 2.5KM 2019 dataset.

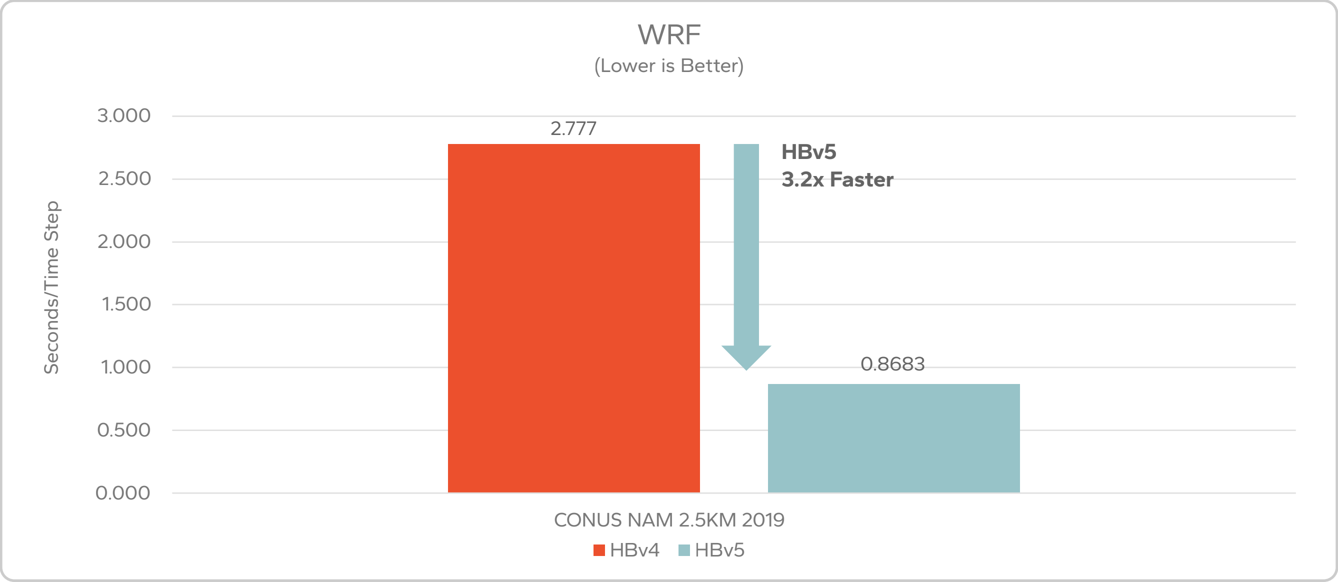

Figure 5: *WRF Performance*

The 3.2x speedup for WRF can have profound life-or-death consequences. For meteorological centers, this generational leap allows them to run higher-resolution ensembles without increasing the time-to-forecast. This is critical for localized forecasting in topographical regions where isolated weather events can develop in minutes.

Quantum Chemistry: CP2K v2024.1

CP2K evaluates atomistic systems using density functional theory (DFT) and molecular mechanics. The tests included Quick-Step and Linear-Scaling solvers for H2O systems of increasing size.

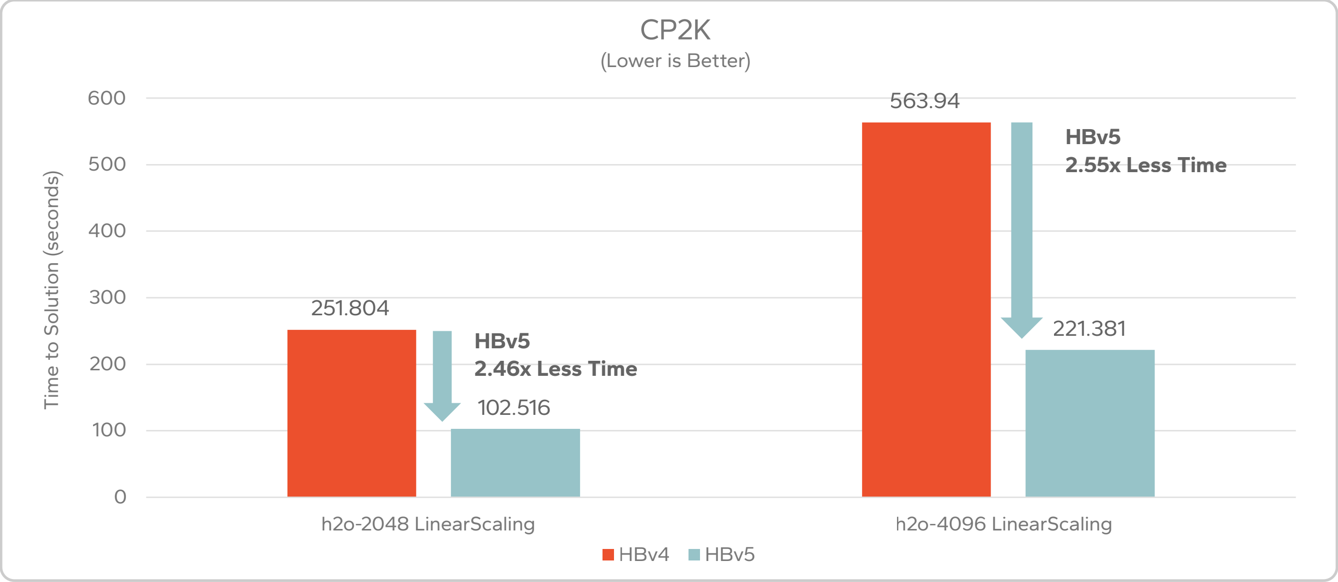

Figure 6: *CP2K Performance*

The CP2K data clearly illustrates the "Memory Wall" phenomenon. At the smallest scale (h2o-256), the HBv5 is only 7% faster than the HBv4, as the Genoa-X processor's massive L3 cache is large enough to hold the entire active dataset, neutralizing the bandwidth advantage of the HBM3.

However, as the system grows to h2o-4096, the data exceeds the cache capacity and must be fetched from main RAM. At this point, the DDR5 bandwidth in the HBv4 becomes a throttling factor, while the HBM3 in the HBv5 provides a consistent, high-speed throughput pipe, leading to a 2.55x speedup. It is also noted that for the h2o-8192 case, both systems failed with an Out-of-Memory error, reminding researchers that HBM3, while fast, currently lacks the massive capacity of DDR5 DIMM-based systems.

Conclusion: The Strategic Value of HBM in the Cloud

The Azure HBv5 represents the democratization of supercomputer-class performance in a public cloud environment. By co-designing a custom EPYC processor with integrated HBM3, Microsoft and AMD have effectively removed the "memory wall" that constrained previous HPC generations. This transition from DDR5 to HBM3 allows the 368 physical cores to operate at their full potential, expanding the limits for memory-bound high-performance computing.

The performance data is conclusive: the HBv5 is not merely an incremental update but a foundational shift that enables time-to-solution improvements ranging from 2x to 5x across the most demanding scientific applications in use today.

From an economic perspective, the HBv5 offers a compelling ROI. While the hourly rate reflects the premium nature of the hardware, the consolidation of workloads onto fewer, more powerful nodes results in a lower TCO per simulation. In sectors where time is the most valuable resource, such as disaster response, aerodynamic development, and pharmaceutical discovery, these technical gains translate into concrete business value. This is realized through shorter drug discovery timelines, faster simulation iterations, and earlier weather warnings.

As AI and simulation workloads continue to grow in complexity, the HBv5 provides the scalability and deterministic reliability required to support the next wave of scientific discovery. Organizations should review their production deployment procedures and engage in proof-of-concept testing to validate these gains within their specific workflows. The era of HPC performance in the public cloud has arrived, and early adopters stand to establish a leadership position that will compound over time.

Appendix

Implementation Best Practices and Optimization Strategies

To realize the gains identified in the benchmarks, administrators and researchers must apply specific tuning strategies that align the software environment with the HBv5 architecture.

Process Pinning and Topology Awareness

The presence of 16 NUMA domains and 48 CCDs necessitate precise MPI rank placement. For optimal performance, users should utilize a symmetric rank distribution per CCD.

Commonly recommended configurations include using 48, 96, 144, 192, 288, or 336 ranks per VM. Because four cores per CCD are required to fully saturate the memory bandwidth, configurations with 144 cores or fewer may not achieve the headline memory throughput figures of 6.9 TB/s. Organizations should also use the HPC compute tuned profile to optimize kernel scheduling and power management for HPC duty cycles.

Compilation Flags and Accuracy

During the evaluation, testers compared the impact of aggressive mathematical optimizations. For WRF and OpenFOAM, Microsoft often utilizes the -Ofast flag to maximize calculation speed. However, empirical testing showed that for OpenFOAM on HBv5, switching from standard gcc to aocc (the AMD Optimizing C/C++ Compiler) and adding the -Ofast flag resulted in only a 7-second improvement.

Given the minimal gain, the use of -Ofast is generally discouraged for weather modeling and safety-critical CFD, as it can reduce numerical accuracy. In meteorological applications, maintaining the physical consistency of the Runge-Kutta schemes used for transport equations is more important than a fractional speedup.

OS and Software Stack Recommendations

The recommended operating system for HBv5 is AlmaLinux 8.10 or Ubuntu 24.04. Testers noted that attempts to use the spack package manager for OpenFOAM were sometimes unsuccessful, suggesting that manual compilation or pre-configured Azure HPC images are the most reliable path for initial deployment. For MPI, the use of HPC-X is strongly recommended, as it is pre-configured to utilize the ConnectX-7's hardware offload features for collective operations.

Detailed Specifications and Performance Tables

This appendix contains the exhaustive data and specifications that were summarized in the main report. These details are intended for systems architects and HPC engineers tasked with cluster design and workload mapping.

Complete HBv5 Size Configurations

HBv5 instances include the following features:

- 6.9 TB/s of memory bandwidth (STREAM Triad) across 400-450 GB of RAM (HBM3)
- Up to 9 GB of memory per core (customer configurable)
- Up to 352 AMD EPYC "Zen4" CPU cores, 4 GHz peak frequencies (customer configurable)
- 2X total Infinity Fabric bandwidth among CPUs as any AMD EPYC™ server platform to date
- SMT disabled, single-tenant only design (1 VM per server)
- 800 Gb/s of NVIDIA Quantum-2 InfiniBand, balanced as 200 Gb/s per CPU SoC
- Azure VMSS Flex to scale MPI workloads to hundreds of thousands of HBM-powered CPU cores
- 160 Gbps of Azure Accelerated Networking via 2nd generation Azure Boost NIC
- 14 TB local NVMe SSD delivering up to 50 GB/s read and 30 GB/s write bandwidth

Some advantages of the HBv5 instances compared to HBv4 include:

- 2.1x greater CPU cores (368 / 176) with each core a newer generation and higher frequency
- 5.6x greater memory bandwidth (6.7 TB/s / 1.2 TB/s = 5.58)
- 2x higher networking connectivity (800 Gb/s InfiniBand vs. 400 Gb InfiniBand)
- 4x higher local storage performance (50 GB/s / 30 GB/s r/w) vs. (12 GB/s / 7 GB/s r/w)

Azure provides a range of constrained-core sizes within the HBv5 series. It is critical to note that while the core count decreases, all other shared assets—including HBM3 capacity, memory bandwidth, L3 cache, and InfiniBand throughput—remain constant.

VM Size	Cores	NUMA Domains	HBM3 Capacity	Memory Bandwidth
Standard_HB368rs_v5	368	16	432 GB	6.7 – 6.9 TB/s
Standard_HB368-336rs_v5	336	16	432 GB	6.7 - 6.9 TB/s
Standard_HB368-288rs_v5	288	16	432 GB	6.7 - 6.9TB/s
Standard_HB368-240rs_v5	240	16	432 GB	6.7 - 6.9 TB/s
Standard_HB368-192rs_v5	192	16	432 GB	6.7 - 6.9 TB/s
Standard_HB368-144rs_v5	144	16	432 GB	6.7 - 6.9 TB/s
Standard_HB368-96rs_v5	96	16	432 GB	6.7 - 6.9 TB/s
Standard_HB368-48rs_v5	48	16	432 GB	6.7 - 6.9 TB/s

Figure 7: *Azure HBv5 Configuration Options*

Pricing Details

Microsoft Azure Online Pricing: https://azure.microsoft.com/en-us/pricing/details/virtual-machines/linux/#pricing

- Region: US South Central
    - HBv4 – HB176-96rs v4: Pay as you go: $6,307.20 / month = $8.64 / hour
    - HBv5 – HB268rs v5: Pay as you go: $17,344.80 / month = $23.76 / hour
- Cost Ratio: $17,344,80 / $6,307.20 = 2.29x cost differential

Detailed Raw Benchmark Results (Full Dataset)

OpenFOAM simplefoam v2406

Large Motorbike v1912 100M cells (Seconds)

- HBv4: 979, 977, 981 (Median: 979)
- HBv5: 293, 293, 294 (Median: 293)
- Note: I/O moved to ephemeral disk.

GROMACS benchPEP

(ns/day)

- **HBv4:** 2.522, 2.415, 2.52 (Median: 2.52)
- **HBv5:** 5.373, 5.363, 5.345 (Median: 5.363)
- Settings: -nsteps 10000.

GROMACS benchRIB

(ns/day)

- **HBv4:** 27.762, 27.214, 27.342 (Median: 27.342)
- **HBv5:** 38.068, 37.117, 37.353 (Median: 37.353)

NAMD v3.0.2 (Binary)

(ns/day)

- **HBv4:** 0.489, 0.487, 0.49 (Median: 0.489)
- **HBv5:** 0.853, 0.899, 0.856 (Median: 0.856)

NAMD v2.15a2 (Binary)

(ns/day)

- **HBv4: 0**.486, 0.484, 0.482 (Median: 0.484)
- **HBv5:** 0.782, 0.883, 0.902 (Median: 0.883)

WRF v4.2.2 (CONUS NAM 2.5KM 2019)

(Seconds per time step)

- **HBv4:** 1.4443, 1.4391, 1.4402 (Median: 1.4402)
- **HBv5:** 0.5712, 0.5638, 0.5691 (Median: 0.5638)

CP2K h2o-4096 LinearScaling and QuickStep

(Seconds)

- **HBv4:** 563.94, 563.906, 564.368 (Median: 563.94)
- **HBv5:** 220.574, 221.381, 223.205 (Median: 221.381)

Test Case (Molecular System Size)	HBv4 Median (sec)	HBv5 Median (sec)	Speedup Gain
h2o-256 QuickStep	128.18	119.749	1.07x
h2o-512 QuickStep	348.435	316.519	1.10x
h2o-1024 QuickStep	1209.896	870.321	1.39x
h2o-1024 LinearScaling	81.175	42.283	1.92x
h2o-2048 LinearScaling	251.804	102.516	2.46x
h2o-4096 LinearScaling	563.94	221.381	2.55x

Figure 8: *CP2K Extended Performance*

Important Information About this Report

AUTHOR

Jonathan Fellows
Performance Validation Engineer | Signal65

Russ Fellows
VP, Labs | Signal65

PUBLISHER

Ryan Shrout
President, GM | Signal65

INQUIRIES

Contact us if you would like to discuss this report and Signal65 will respond promptly.

CITATIONS

This paper can be cited by accredited press and analysts, but must be cited in-context, displaying author's name, author's title, and "Signal65." Non-press and non-analysts must receive prior written permission by Signal65 for any citations.

LICENSING

DISCLOSURES

Signal65 provides research, analysis, advising, and consulting to many high-tech companies, including those mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

ABOUT SIGNAL65

Signal65 is an independent research, analysis, and advisory firm, focused on digital innovation and market-disrupting technologies and trends. Every day our analysts, researchers, and advisors help business leaders from around the world anticipate tectonic shifts in their industries and leverage disruptive innovation to either gain or maintain a competitive advantage in their markets.

signal65

signal65

CONTACT INFORMATION

Signal65 I signal65.com