Ansys

Ansys Cloud: Configured and Optimized for HPC and Ansys Solvers

Date: 09/17/2021

Authors: Neehar Kulkarni Scott Gilmore, Wim Slagter, Ashwini Kumar, Thomas Lejeune, Michael Tooley, Matt Commens, Hardik Shah, Jim DeLap, Bruno Reymond and Navin Budhiraja

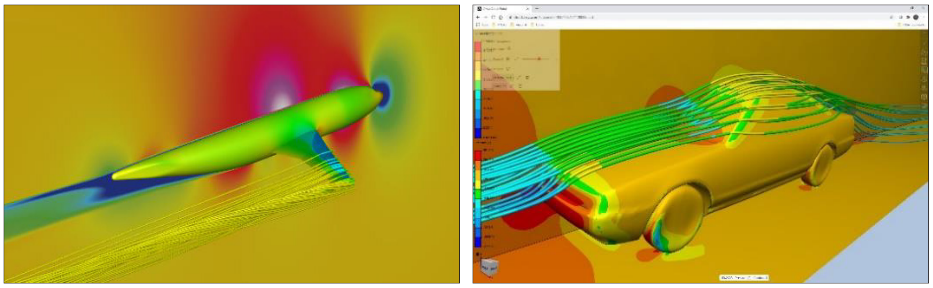

Figure 1: Example of Fluent Solver running on Ansys Cloud

Menu: Click to access the section you want to see.

Ansys Cloud
Cloud & Structures
Cloud & Electronics
Cloud & Fluids

/ MORE CHOICE, MORE FLEXIBILITY, MORE POWER

More and more engineers bump up against the limits of their desktop/workstation computers for engineering simulations. They are looking for ways to conduct more, bigger, and faster simulations. Cloud high-performance computing (HPC) resources enable that in an efficient and economical manner. Cloud computing delivers on-demand access to HPC infrastructure, including the fastest processors, automatic software upgrades, support for a mobile workforce and the ability for enhanced collaboration among geographically distributed teams. While looking to cloud solutions, decision-makers must also ensure that outsourcing IT will not create any cybersecurity and data issues. Additionally, costs can be reduced by only paying for HPC resources that the company uses.

Despite these benefits of cloud computing, most engineers will likely dismiss cloud providers if they must develop their simulation environment on a cloud infrastructure themselves. Matching appropriate cloud HPC resources (processor, memory, storage..) to each simulation application is essential for ensuring that jobs are optimally executed and cloud HPC resources are efficiently used. Building up an enterprise-grade simulation environment with HPC orchestration, storage and remote visualization in the public cloud can easily take many months, and it is beyond many engineers' expertise or training.

This represents a significant task for a busy engineering team to tackle and could even tax small or overworked information technology (IT) departments.

Ansys has addressed these needs by developing and offering Ansys Cloud as a Software-as-a-Service (SaaS). Requiring no additional configurations, the Ansys Cloud provides engineers easy and instant access to on-demand HPC in the Microsoft Azure cloud directly from within Ansys applications.

This paper describes how Ansys Cloud is architected and optimized for Ansys' applications in terms of performance, cost, and usability.

/ CURRENT HARDWARE CONFIGURATIONS

Ansys Cloud provides virtual machine (VM) configurations in partnership with Microsoft Azure across nine global regions. It offers HPC clusters with H, HC, HB, HBv2 and HBv3 VMs, delivering large configurations with high memory and memory bandwidth and excellent price/performance. VMs H16r, H16mr, HB, HBv2, HBv3 and HC are available for batch HPC workflows, whereas H16mr, HB, HBv2, HC, NV6 and NV12sv3 are available to run interactive sessions in-browser, potentially enabling full end-to-end cloud-based simulation workflows. These configurations provide significant power, choice, and flexibility. Additionally, these options enable graphics-intensive, interactive simulation workloads thanks to the use of virtual machines with Nvidia GPUs (NV6, NV12sv3) powering more physics and more applications. By having more cores, more memory and memory bandwidth, engineers can run bigger, more challenging simulations and get higher fidelity insight into how the design will work in the real world.

Infra	Cores per node	Frequency Peak	RAM per Node	Memory Bandwidth	Interconnect	
H16r	16	3.3 GHz	112 GB	80 GB/s	56 Gb/s	Current Configurations
H16mr	16	3.3 GHz	224 GB	80 GB/s	56 Gb/s	
HBv1	60	2.55 GHz	240 GB	263 GB/s	100 Gb/s	Last released
HBv2	120	3.1 GHz	480 GB	350 GB/s	200 Gb/s	
HBv3	120	3.1 GHz	448 GB	350 GB/s	200 Gb/s	
HC	44	3.4 GHz	352 GB	191 GB/s	100 Gb/s	
Nv6	6 cores, M60 GPU	NA	56 GB	NA	In Browser only	
Nv12sv3	12 cores, M60 GPU	NA	112 GB	NA	In Browser only	

Figure 2: Current HW configurations for interactive sessions and batch HPC jobs.

When selecting the number of nodes and number of cores to utilize per-node, it's important to understand the architecture of each configuration and how the number of utilized cores affects the per-core and per-node performance. For example, the HBv2 VM contains 120 cores and has 350 GB/s of memory bandwidth to share between those cores. If we use only 60 out of 120 cores, then the memory bandwidth available per-core will be higher, and the per-core performance will be better. Memory bandwidth is not the only resource shared between cores, but it does have a strong impact on solver performance, particularly for computational fluid dynamics analysis.

For each VM there is an optimal balance between the number of cores utilized, the per-core performance and the per-node performance. The best choice also depends on your objective, for example minimizing solve time or minimizing job cost, and on the solver being used.

Finally, the Non-Uniform Memory Architecture (NUMA) configuration for each VM should be considered for best performance. The HB VM's contain 2 sockets, 15 NUMA domains with 4 cores per NUMA domain and will perform best when exactly 1, 2, 3 or 4 cores are used per NUMA domain. This corresponds to 15, 30, 45 or 60 cores per node; using core counts other than this will create an imbalance between NUMA domains resulting in poor performance. Similarly, the HBv2 VM's should use 30, 60, 90 or 120 cores per node for best performance, since this VM contains 2 sockets, 30 NUMA domains and 4 cores per NUMA domain. HBv3 VM's provide 120 physical cores, but have been pre-configured with 32, 64, 96 or 120 cores exposed, allowing you to simply select the appropriate VM and use all available cores. The H16 and HC VMs are less sensitive to the number of cores utilized and generally give the "expected" performance at any utilized core count.

You will find recommendations below for the number of cores to utilize for each solver type and each VM type.

/ SELECTING THE CLOUD SERVER FOR OPTIMUM COMPUTE PERFORMANCE

Hardware is configured and optimized for Ansys solvers; however, not all configurations are recommended for every solver. That is why we provide guidance to choose your ideal machine and find the best type of VM. We help you select VMs that work the best with your solvers, and we run many benchmarks that provide best practices that get you up to speed for selecting the best VM. Additionally, we also provide guidance to help select the best combination of number of nodes and number of cores.

These recommendations are based on our internal benchmarks, are detailed below in this document.

We have run hundreds of technical tests to identify the perfect match through all the different parameters and for each physics. This optimization and limitless configuration capabilities enable you to identify your ideal configuration.

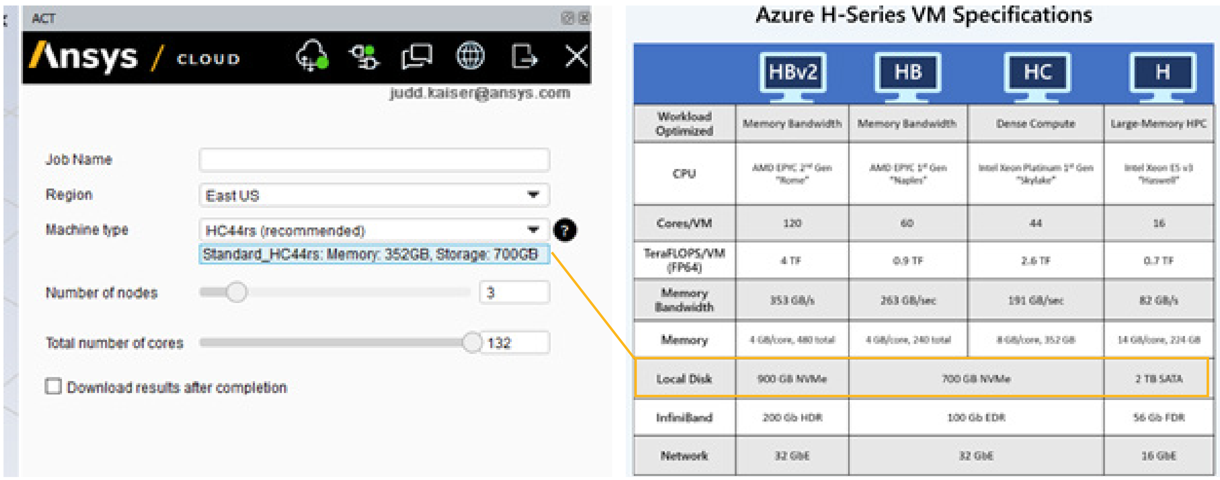

	HBv2	HB	HC	H
Workload Optimized	Memory Bandwidth	Memory Bandwidth	Dense Compute	Large-Memory HPC
CPU	AMD EPYC 2nd Gen "Rome"	AMD EPYC 1st Gen "Naples"	Intel Xeon Platinum 1st Gen "Skylake"	Intel Xeon E5 v3 "Haswell"
Cores/VM	120	60	44	16
TeraFLOPS/VM (FP64)	4 TF	0.9 TF	2.6 TF	0.7 TF
Memory Bandwidth	353 GB/s	263 GB/sec	191 GB/sec	82 GB/s
Memory	4 GB/core, 480 total	4 GB/core, 240 total	8 GB/core, 352 GB	14 GB/core, 224 GB
Local Disk	900 GB NVMe	700 GB NVMe		2 TB SATA
InfiniBand	200 Gb HDR	100 Gb EDR		56 Gb FDR
Network	32 GbE	32 GbE		16 GbE

Figure 3: Select your hardware configuration in Ansys Cloud, you can find VM specifications such as Local Disk space on Azure.

Note that we recently added two new regions (U.S. West and U.S. South central) and our data centers span the world. You can select where you would like to run the job and submit it to the appropriate server. This worldwide coverage provides broader support and availability, and control over data locality supporting GDPR.

/ ANSYS CLOUD, CONFIGURED AND OPTIMIZED FOR MECHANICAL SOLVER

Different VM configurations on Ansys Cloud are optimized for Ansys Mechanical and Ansys LS-Dyna requirements. We evaluated and benchmarked to provide and recommend best possible VM Configuration to our customers focusing on performance in terms of minimum solve time, optimum solver and HPC license utilization, minimum hardware cost and minimum total cost.

New Hardware Introduction

In August 2021 Ansys Cloud introduced new hardware [HBV3] with core-constrained configurations (32, 64, 96, and 120 cores per node) to get maximum performance based on NUMA domains. Leaving some cores idle increases memory bandwidth available for the active cores. Benchmarks were run on all four configurations as well as HC using 44, 33, and 22 cores per node.

Focusing on Ansys Mechanical

- Ansys Mechanical in general shows strong scalability across hardware configurations even at 400+ cores.
- HC is recommended for most situations. HC gives the fastest and most consistent performance. Usually, all 44 cores on the node should be used.
- HBv3 is a close second choice for release 2021R2 and later, especially when considering HW cost.

Ansys Mechanical - Detailed Benchmark Results

Ansys Mechanical [Versions 2020R2 and 2021R2] proves to be scalable even at higher core counts. During our benchmarking, Highly Nonlinear [16 million DOF] Implicit Problem showed scale up even at 400+ core count (Ref: Figure 04).

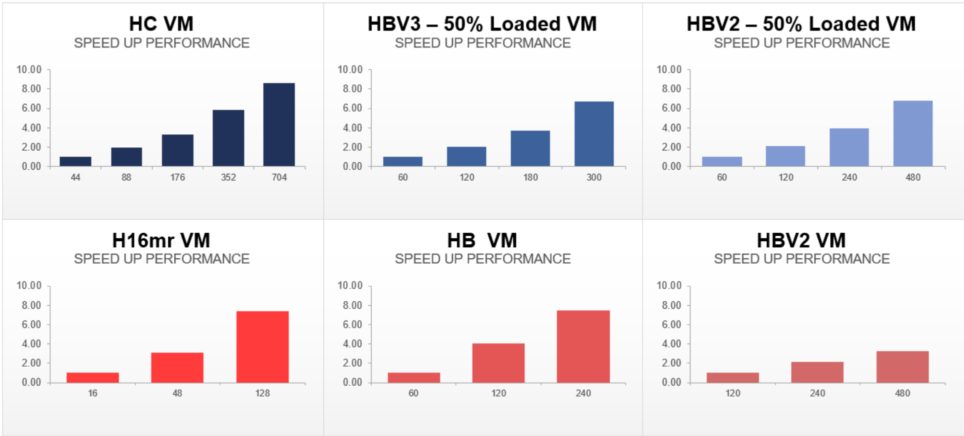

Figure 4 : 2020R2/2021R2 Mechanical Direct Solver shows strong scale up for 16 million DOF Highly Nonlinear Problem

Ansys Cloud Benchmarks for different VMs with Ansys Mechanical 2020R2 clearly indicate HC as best performing VM from Solve Time, HW cost and Total Cost perspective (Ref: Figure 05). Selected benchmark model demands high memory. HB/HBV2 VM have lower memory per core than HC. By running models at 50% load, we can increase per-core memory allocation significantly which improves performance substantially. This recommendation depends on model size, Memory Requirement, etc.…

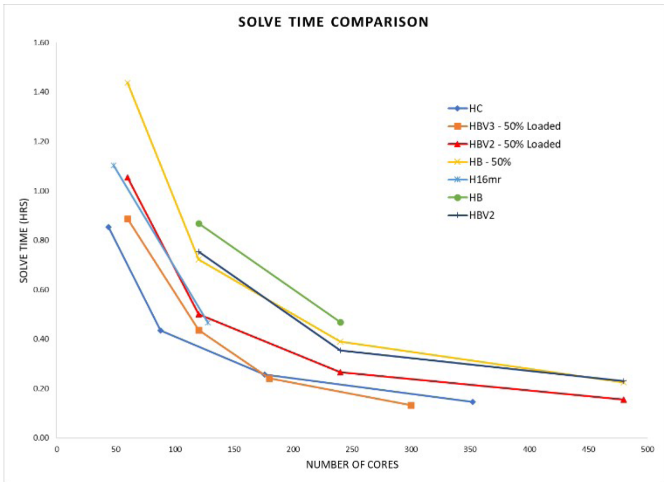

Figure 5: 2020R2 Mechanical Direct Solver performance across different VM Configurations

Ansys Mechanical made solver improvement in release 2021R2 to perform better on AMD processors (HB/HBV2/HBV3) and at the same time HBV3 got introduced to Ansys Cloud. Considering both improvements, when solving with the same memory allocation per core, HC and HBV3 provide similar solver speed for the selected benchmark model which also reduced the cost gap for hardware only and total cost (Ref: Figure 6). This recommendation depends on model size, Memory Requirement, etc.…

The simplest recommendation is to use HC for all MAPDL solves considering Solver Speed, HW cost and Overall Total Cost parameters. In addition to that, HC is less dependent on the number of cores utilized and it also has more disk space per Node/Core. To achieve optimal performance from HBv3 it is important to select the right number of cores. For example, to minimize hardware cost per job, 120 or 96 cores per node are best. To minimize total cost including hardware and software, 32 cores per node is likely the best option. HBv2 tends to perform slightly slower than HBv3, and users must similarly be aware of the number of cores to use to achieve the best performance (30, 60, 90, or 120).

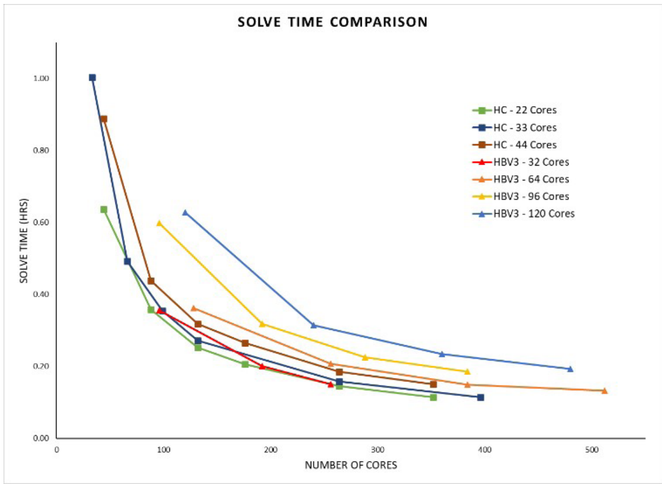

Figure 6.1 : 2021R2 Mechanical Direct Solver performance comparison between HC and HBV3 – Solve Time Focused

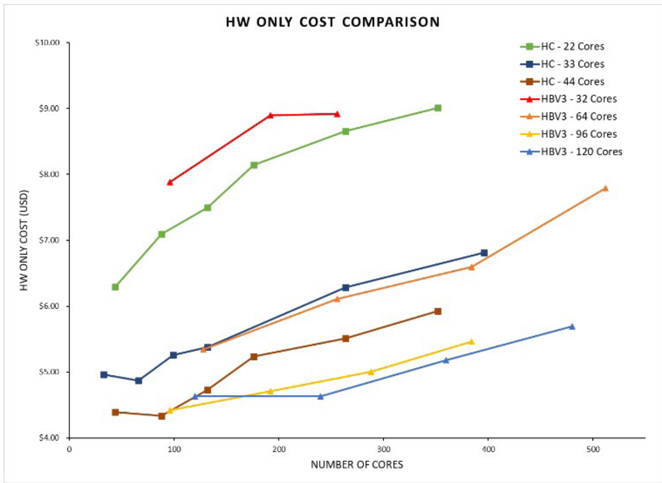

Figure 6.2: 2021R2 Mechanical Direct Solver performance comparison between HC and HBV3 –HW Cost Focused

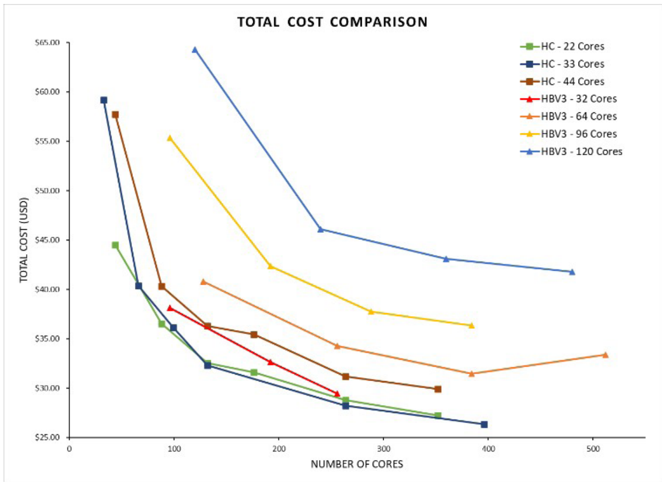

Figure 6.3: 2021R2 Mechanical Direct Solver performance comparison between HC and HBV3 –Total Cost Focused

Figure 7 shows Ansys LS-DYNA solve times for a model with 12.5 million elements simulating one second of time.

HC delivers the fastest performance, but HBv2 at 50% capacity (using half of the available cores) comes in a close second, followed by HBv2 at full capacity. In general, running compute nodes at half capacity improves performance, due to the memory and interconnect bandwidths.

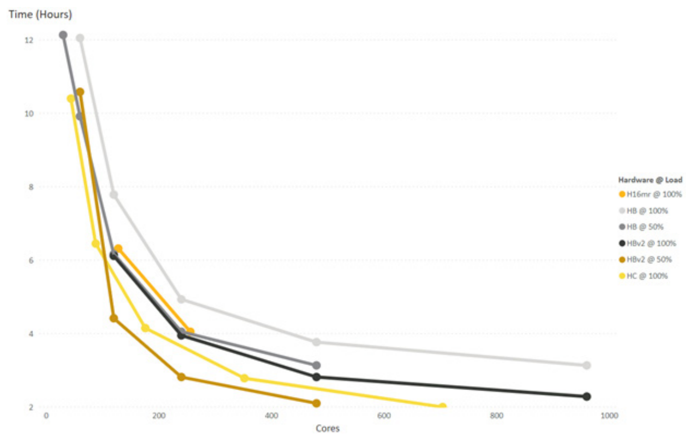

Figure 7: LS-DYNA solver solve time on all available VM types.

Looking more closely at the HBs, Figure 8 further illustrates how running at 50% capacity significantly improves performance. It also shows the performance benefits of the larger, faster HBv2 over the HB at equivalent total core counts. HBv3 improves performance by approximately 15% relative to HBv2.

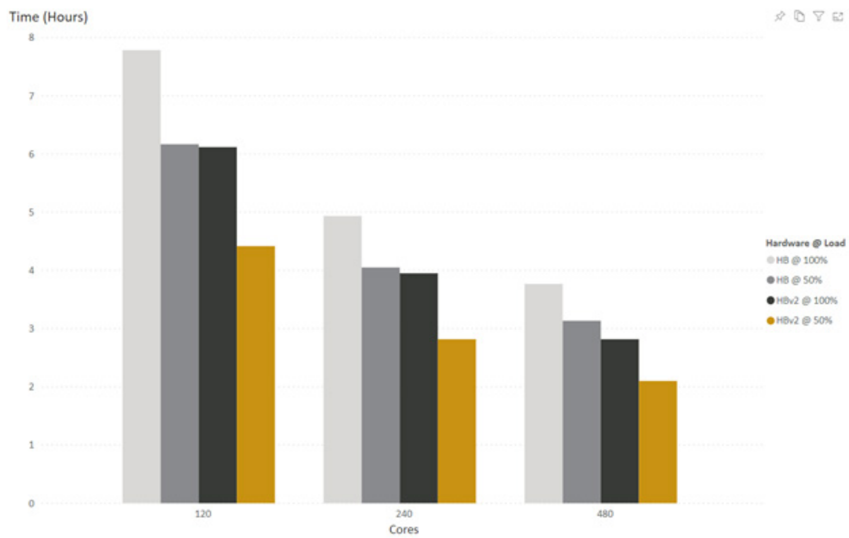

Figure 8: LS-DYNA solver time on all HB and HBv.

In summary:

- HC is the recommended type for LS-DYNA in most situations. It usually delivers the best cost performance.
- HBv2 and HBv3 are also good choices, especially for large models and when running with half of available cores.
- H16mr and HB are suitable for small models and are reasonable choices when the other types are not available.

/ ANSYS CLOUD: CONFIGURED AND OPTIMIZED FOR ELECTRONICS SOLVERS

All Ansys Electronics Desktop products use the same user interface to submit jobs. The following recommendations are based on requirements across all solvers (Ansys HFSS, Ansys Maxwell, Ansys SIwave, Ansys Q3D Extractor and Ansys Icepak). The two factors considered for these recommendations are performance scalability and memory footprint.

Figure 9 compares solution times for two HFSS models on the 16-core H16mr, 44-core HC, and 60-core HB.

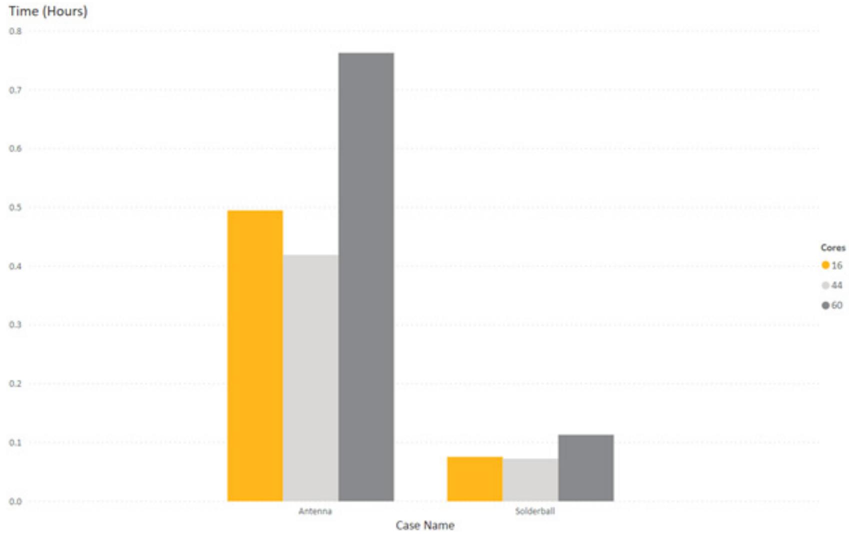

Figure 9: HFSS solution times for two models on H16mr (16 cores), HC (44 cores), and HB (60 cores).

HC shows significant improvements over H16mr. However, HB is slower due to the lack of optimized math libraries for its AMD CPU generation. This will be addressed in the future. HBv2 provides slightly improved performance over HB, and HBv3 is slightly faster than HBv2. For now, HC is the recommended machine type for the electronics solvers.

Following are specific results on HC for two HFSS examples and one Maxwell example.

HFSS SO-DIMM

An eight-layer, eight-module small outline dual in-lina memory module (SO-DIMM) was solved using mixed order elements to address the high degree of geometric complexity by leveraging the distributed memory direct matrix solver. This latter solver technology allows a single matrix solution to be distributed across multiple machines and eliminates the requirement for a single machine's shared memory to be sufficiently large to solve the entire problem. At convergence, the problem size is roughly 7M tetrahedra and 28M matrix unknowns, using 35 GB of RAM per engine. During the frequency sweep, the memory footprint per point is reduced with S-parameter-only matrix solve, which only holds in memory the part of matrix for extracting S-parameters. This allows more frequency points to be solved in parallel and delivers overall faster, and more scalable, frequency sweeps in Ansys Cloud.

Figure 10 shows the timing results on the HC machine type. Performance steadily improves with the addition of more compute nodes.

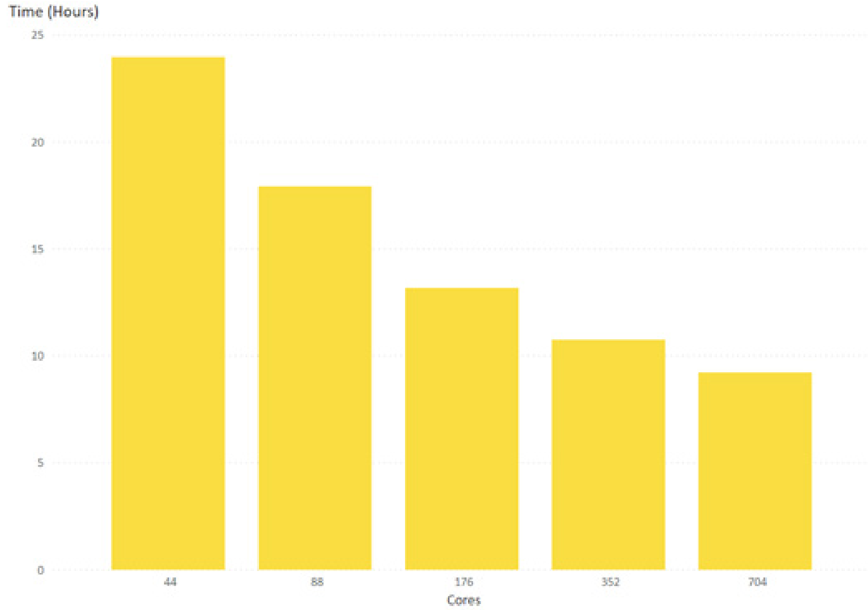

Figure 10: HFSS SO-DIMM on HC.

HFSS Anechoic Chamber

This is a simulation of an automobile in an electromagnetic interference (EMI) anechoic test chamber studied at 400 MHz, where the orientation of the vehicle is parameterized with the variations solved in parallel. The solution uses domain decomposition and mixed order elements to address the electrically large nature of this design. The typical problem size at convergence is 2.5M tetrahedra and 25M matrix unknowns.

This type of parametric analysis exhibits near linear scalability on HC, as shown in Figure 11.

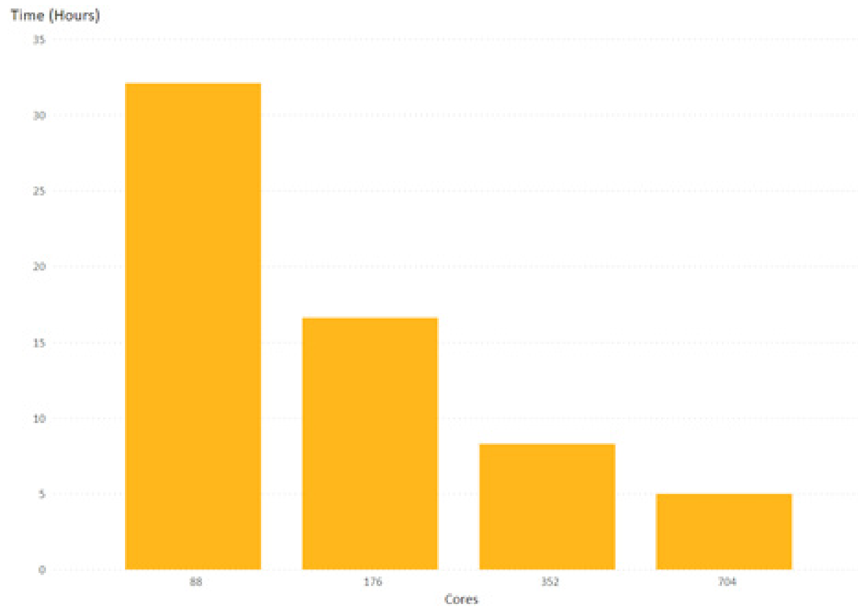

Figure 11: HFSS anechoic test chamber on HC.

Maxwell Induction Motor

This is a steady-state transient simulation for an 8-kW induction motor. The resulting mesh contains 469,192 tetrahedra. It was solved for 400-time steps with HPC TDM technology.

This analysis exhibits excellent scalability up through 352 cores on HC (eight compute nodes), as shown in Figure 12. By employing a larger time-domain, further compute time savings can be achieved with more cores.

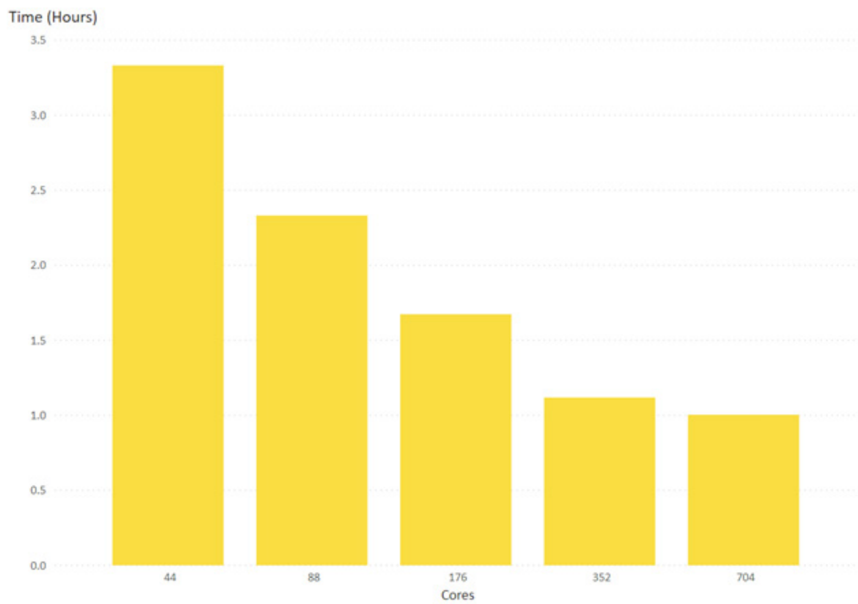

Figure 12: Maxwell on HC

/ ANSYS CLOUD, CONFIGURED AND OPTIMIZED FOR FLUIDS SOLVERS

We ran numerous benchmarks to test and compare different hardware configurations across a range of use cases. Results varied depending on use case size, complexity and physical models used, but showed consistent trends. The goal of these benchmarks was to identify the preferred hardware configuration for different scenarios with fluids solvers running on Ansys Cloud.

As discussed earlier, the number of cores to utilize per node is a key consideration. First, we will establish recommended cores-per-node for each VM.

Figure 13 shows various performance metrics for a single HB node with 15, 30, 45 and 60 cores utilized. As noted earlier, these core counts align with the NUMA architecture of HB VM and we do not recommend any other core counts. This data set is from Fluent 2021 R1 running the 14M cell aircraft benchmark, but the trends are consistent across all tested benchmarks.

Figure 13: Example HB performance metrics for Fluent.

As the number of cores utilized reduces from 60 to 45 to 30 and then to 15, we observe a significant increase in per-core performance which offsets some of the drop from using fewer cores. Utilizing 30 cores achieves 82% of the performance compared to 60 cores and using 45 cores achieves 94% of the performance. If we only consider hardware costs, then 60 cores-per-node provides the lowest cost per job, but when we place value on license usage then lower core utilization per node is clearly the better choice. The license usage costs in this data set are based on Ansys Elastic Currency (AEC) costs, but the same conclusions will be reached whenever license usage value is high compared to hardware costs.

When using AEC, 15 cores-per-node is recommended with HB VMs up to 240 cores total, and 30 cores-per-node is recommended when between 240 and 480 cores total are used. Note that the above chart shows 15 cores-per-node has the highest Total Job Cost on a single node, but on an equal core-count basis it will have the lowest Total Job Cost (i.e. 2 nodes with 15 cores-per-node will have a lower Total Job Cost than 1 node with 30 cores-per-node).

When using bring-your-own-licensing (BYOL), 45 cores-per-node with HB is recommended for the best performance/cost ratio while 30 cores-per-node will give better per-core performance at a higher hardware job cost. 60 cores-per-node is not recommended with either AEC or BYOL

Figure 14 shows the same metrics for a single HBv2 VM. The trends are like HB, but with a larger increase in per-core performance as we reduced the number of utilized cores, resulting in 90 cores outperforming 120 cores on a per-node basis.

Figure 14: Example HBv2 performance metrics for Fluent.

When using HBv2 VM's we recommend 60 or 90 cores-per-node with Bring-Your-Own-Licensing (BYOL. Using 90 cores-per-node gives lower hardware costs per job, while 60 cores-per-node provides better per-core performance. 30 cores-per-node can be used when the best per-core performance is needed at a higher hardware cost, and this is also the recommended option with AEC. 120 cores-per-node is not recommended in any situation.

Figure 15 shows the performance metrics for a single HC VM for the same test case. We still see per-core performance gains when fewer cores are used, but the gains are less than we see with HB and HBv2 VM's. The lowest hardware cost per job and the best performance/cost ratio with AEC are seen when all 44 cores are used, and this is the recommended core utilization in most situations. When using AEC, a reduction in total job cost is observed when using 34 cores-per-node with a reduction in per-node performance of approximately 7%.

Figure 15: Example HC performance metrics for Fluent.

Lastly, for HBv3 VM's we examined the four pre-configured sizes, with 32, 64, 96 and 120 cores available. Each of these VM's contain the same hardware but have been pre-configured to expose different numbers of cores, avoiding the need to understand the NUMA architecture and under-utilize cores. These VM's have been optimized to perform well when using all available cores.

Figure 16: example HBv3 performance metrics for Fluent

Figure 16 shows the trends for HBv3 follow those seen for HBv2, but with slightly better performance. The 64 and 96-core VM are the best options with BYOL, with the 64-core VM providing better per-core performance and the 96-core VM providing lower hardware costs per job. The 32-core VM provides the best per-core performance and is recommended with AEC. 120 cores-per-node is not recommended in any situation.

The above best practices for cores-per-node have been established using performance data for a single node. We believe these practices will remain valid when using multiple nodes, since inter-node communication is not a bottleneck when using high performance interconnects, as provided with Microsoft Azure. While we have not performed exhaustive benchmarks at different cores-per-node with multiple nodes, we have spot-checked several data points and they align with the single-node best practices.

With the cores-per-node best practices established, we can now compare benchmark results on the different VM types. Results from four benchmark cases are presented below, with mesh sizes from 14 million to 140 million cells, both steady and transient cases and a range of physical models. We have selected ten different configurations to compare, corresponding to the best practices for each VM type, and have measured the solver performance, the performance/hardware-configuration-cost ratio, and the performance/total job cost ratio (with the total job cost based on hardware job cost plus elastic licensing cost). These measurements help you select the best configuration to meet your goals, whether that is to minimize job cost, maximize performance or achieve the best performance/cost ratio.

Figure 17 shows the Fluent solver performance for the four benchmark cases. The Aircraft Wing is a 14M cell steady-state external aerodynamics benchmark. The Exhaust System is a 33M cell steady-state internal flow benchmark. The Combustor is a 71M cell transient case, with spray injection and DPM tracking. Lastly, the F1 car is a 140M cell steady-state case solving flow around a Formula 1 car. The charts plot the Solver Rating (higher is better) on the y-axis and the number of cores used on the x-axis. Fives tested configurations are:

- HC with 44 cores-per-node (blue squares)
- HB with 15, 30 and 45 cores-per-node (magenta, gold and light blue circles)
- HBv2 with 30, 60 and 90 cores-per-node (brown, orange and grey triangles)
- HBv3 32-core, 64-core and 96-core VM (green, navy blue and red diamonds)

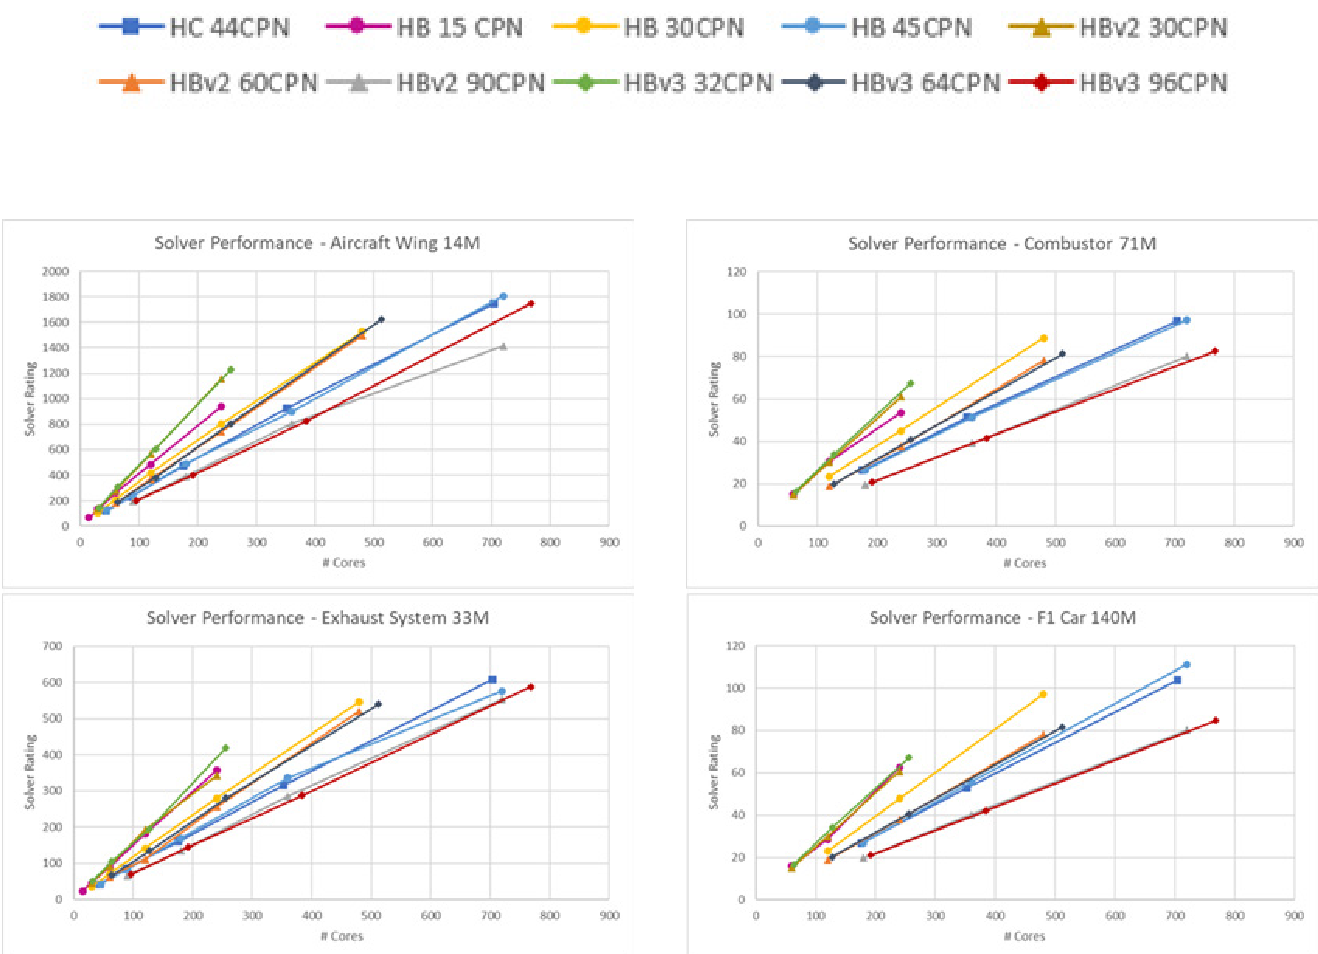

Figure 18: Fluent performance/hardware-cost ratio for four benchmark cases.

Note that a horizontal line in the figures above indicates linear scaling as the number on nodes used increases, whereas lines that slope up/down indicate super/sub-linear scaling as more nodes are used. The HB VM with 45 cores-per-node provides the best performance/hardware-config-cost ratio, followed by HB with 30 cores-per-node. The HBv3 96-core VM provides a good alternative, particularly at low mesh cell counts per-core where the super-linear scaling behavior of this VM is most pronounced. This can be clearly seen on the Aircraft 14M benchmark with the upwards sloping line for the HBv3 96-core VM.

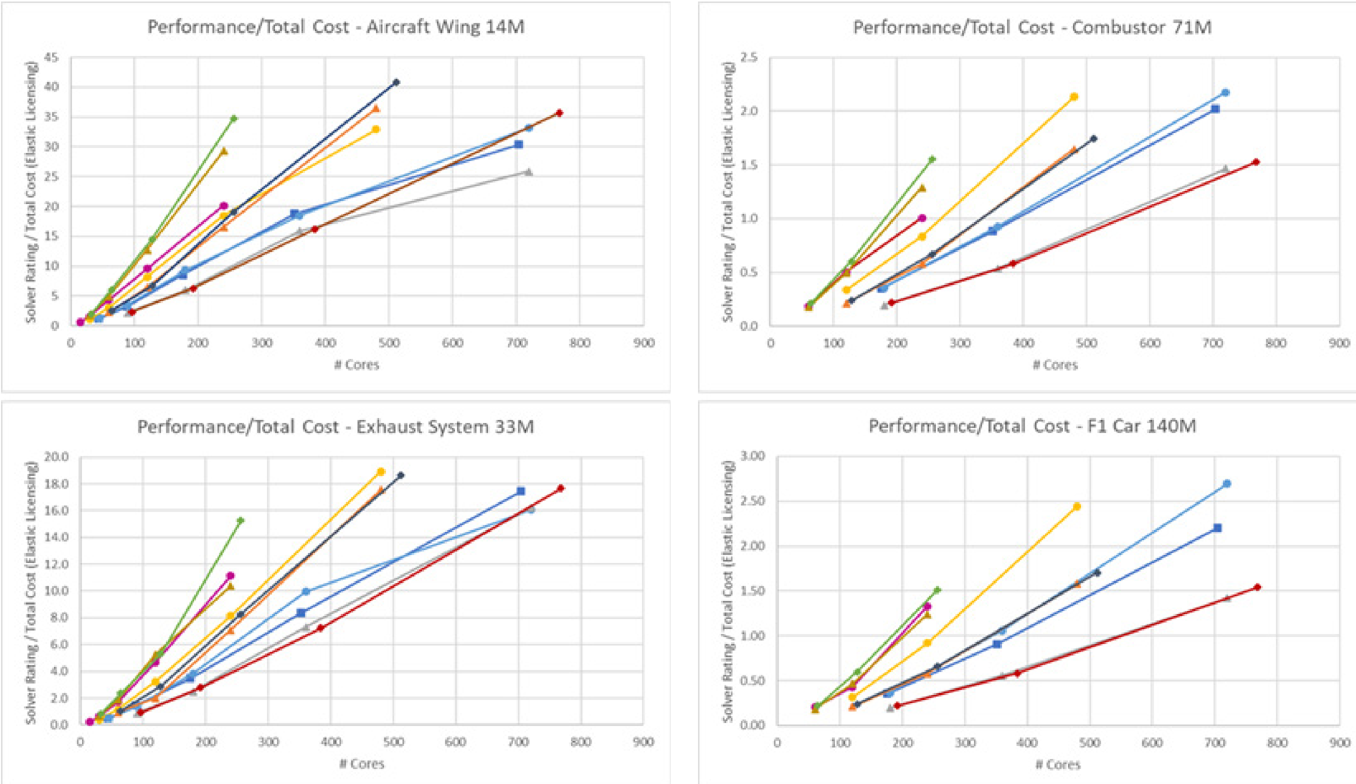

Figure 19: Fluent performance/total-cost ratio for four benchmark cases.

Figure 19 shows the performance to total cost ratio, where the total cost is the sum of the hardware cost and the elastic licensing cost. This data is a good guide when using elastic licensing. Here we see that the HBv3 32-core VM provides the best performance/total cost ratio, followed closely by HBv2 with 30 cores-per-node and HB with 15 core-per-node. Since hardware costs are lower than licensing costs, investing more in hardware to get better per-core performance can results in lower total costs. Note that in all cases we see improved performance/total costs ratios at higher cores counts. This is driven by the HPC licensing cost structure, which provides lower per-core costs as the number of cores increases.

All Fluent benchmarks presented here used version 2021 R1 with Intel MPI 2018, which is the default MPI method for batch solves with Fluent on Ansys Cloud. Process pinning has been used in all cases and is strongly recommended when under-utilizing cores. For example, process pinning is needed with HB 45-cores-per-node and HBv2 90 cores-per-node for best performance, but it is not necessary when using all available cores on the HBv3 VMs. Refer to the Ansys Cloud user guide to set process pinning for Fluent.

If you are using Fluent 2020 R2 or earlier, then process pinning using I_MPI_PIN should not be used. Fluent's default affinity (pinning) settings will do a good job in most cases, but HBv2 90 cores-per-node and HB 45 cores-per-node configurations should be avoided with version 2020 R2 because the default pinning will not set process affinity correctly.

In summary, for Fluids, we can recommend HB 15 or 30-cores-per-node, HBv2 using 30 or 60 cores-per-node, and the HBv3 32-core or 64-core VM when using Ansys Elastic Currency. The lower cores-per-node options are suitable when 256 cores or less are needed, while the higher cores-per-node options allow utilizing up to 512 cores. With standard licensing (BYOL), HB 30-cores-per-node provides excellent performance at a higher hardware cost, and HB 45 cores-per-node provides the best performance/hardware cost ratio. If HB VM's are not available in your region, then we recommend the HBv3 96-core VM or HBv2 with 90 cores-per-node when performance/hardware cost ratio or low job cost are the priority, and the HBv3 64-core VM or HBv2 with 60 cores-per-node for higher per-core performance. Process pinning should be used when underutilizing cores (see the Ansys Cloud users guide) and as a general guideline solving on more cores results in shorter run times and lower job costs, assuming solver scalability is maintained.

/ ANSYS CLOUD, CONFIGURED AND OPTIMIZED FOR ANSYS SOLVERS

Overall, we can observe from these benchmarks that the different hardware configurations can be configured and optimized for Ansys solvers. H16 is good for small models and offers a good price range. HB is good for small models and offers a good price range, especially with BYOL. HBv2 and HBv3 are very good for large models when performance matters and are best when used with half of the available cores. The HC machine type delivers very good price and performance for most situations, making it the recommended VM type for structures, electronics, and fluids solvers.

Do not forget that Ansys Cloud allows you to run your own benchmarks and that we are only giving recommendations based on our tests. You are free to run your own benchmarks and identify the best matching combination for your needs.

/ ANSYS CLOUD PROVIDES AN INTERACTIVE EXPERIENCE ACCESSIBLE VIA A BROWSER

Through Virtual Desktop Interface support, users can also benefit from an end-to-end, cloud-based simulation workflow. The Ansys Cloud in-browser virtual desktop interface (VDI) supports from 16 cores up to 120 cores in a cloud-based workstation and is available in minutes.

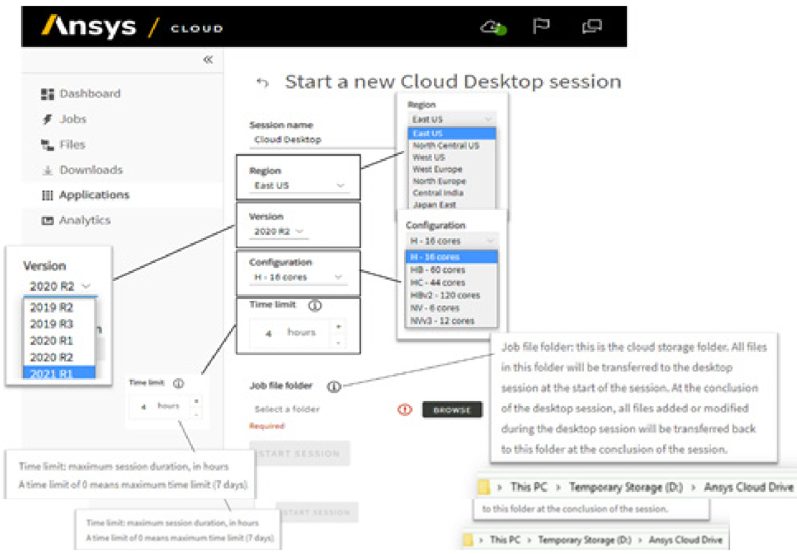

Figure 19 : Start a new Cloud Desktop Session.

The performance is optimized to ensure a reduced latency. You can pre- or post-process in the cloud, or complete a full workstation solve.

Three things are released in this In-Browser:

- More configurations: Access more hardware options for more power and more flexibility.
- In-browser integration: This integration solves the firewall issue of previous RDP VDI solutions.
- Broader product testing coverage.

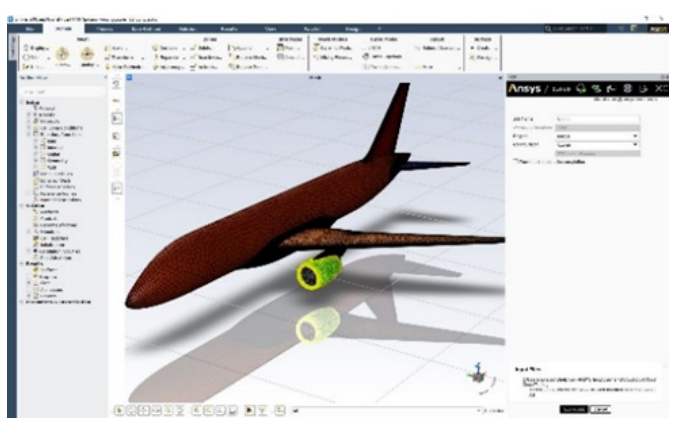

Figure 20: HW config directly through the VDI.

Let us take Ansys Structures as an example. During our benchmarks we have seen that HC outperforms especially nongraphic intensive models and HC is the recommended hardware. HBV2 is good especially for workflows like Design of Experiments (DoE) where the solve time is sensitive to the number of cores. When it comes to graphic usage, Nv6 with GPU is recommended. Graphics process unit (GPU) also helps reduce the solve time.

The combination of a powerful hardware, enhanced VDI and the BYOL model makes Ansys Cloud Platform powerful and provides a strong option to Ansys customers. Customers can meet almost all complex simulation workflows which they currently meet with a combination of workstation and on-premises hardware.

Want to try Ansys Cloud to form your own opinion? Try it for free for one month and experience the power of HPC. Ansys Cloud is HPC as easy as it should be: https://www.ansys.com/products/platform/ansys-cloud/free-trial

ANSYS, Inc.
Southpointe
2600 Ansys Drive
Canonsburg, PA 15317
U.S.A.
724.746.3304
ansysinfo@ansys.com