WHITE PAPER

Ansys Cloud – Configured and Optimized for HPC and Ansys Solvers

Date: 01/26/2021

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

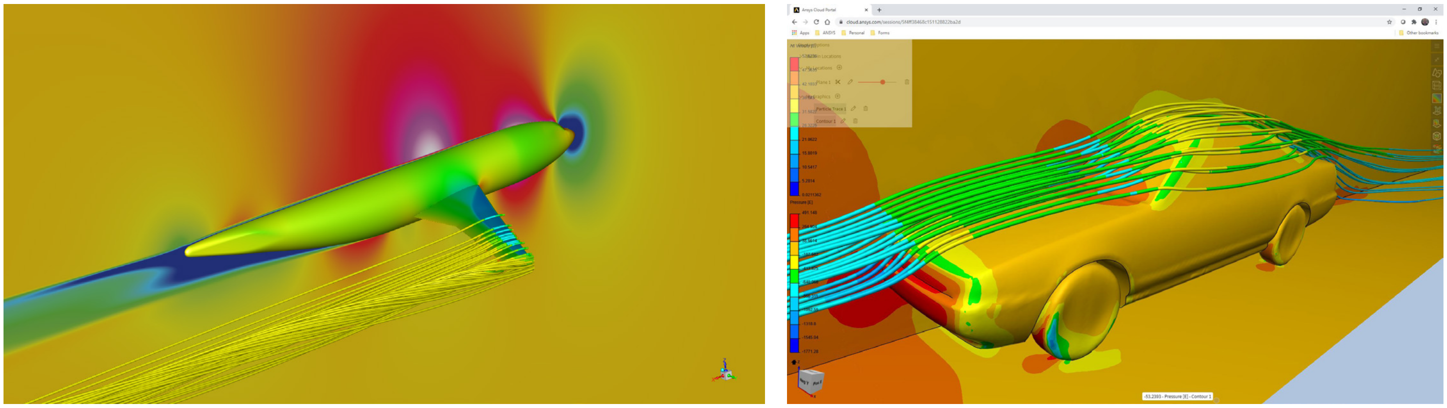

Figure 1: Example of Fluent Solver running on Ansys Cloud

Menu: Click to access the section you want to see

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, etc.) 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, job scheduling, storage and remote visualization in the public cloud can easily take many months, as 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 SaaS cloud service. Requiring no additional configurations, 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 due to the partnership with Microsoft Azure across nine global regions. It has added HPC clusters with HC, HB and HBv2 VMs, delivering larger configurations with more memory and memory bandwidth with better price/performance. VMs H16r, H16mr, HB, HBv2 and HC will be available through HPC, whereas H16mr, HB, HBv2, HC, NV6 and NV12sv3 will be available to run completely in-cloud through an interactive workflow in-browser. These configurations provide a significant boost in power, choice, and flexibility. Additionally, these options enable graphics-intensive, interactive simulation workloads thanks to the use of 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.

VM Type	Cores per node	Frequency Peak	RAM per Node	Memory Bandwidth	Interconnect	
H16r	16	3.3 GHz	112 GB	80 GB/s	56 Gb/s	Existing Configurations
H16mr	16	3.3 GHz	224 GB	80 GB/s	56 Gb/s	Existing Configurations
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 Gbs	Last Released
HC	44	3.4 GHz	352 GB	191 GB/s	100 Gb/s	Last Released
Nv6	6 cores, M60 GPU	2.6 GHz	56 GB	-	-	Last Released
Nv12sv3	12 cores, M60 GPU	2.6 GHz	112 GB	-	-	Last Released

HPC: H16r, H16mr, HBv1, HBv2, HC
Interactive: H16mr, HBv1, HBv2, HC, Nv6, Nv12sv3

Figure 2: Current HW configurations released both for in-browser and HPC

/ Selecting the Cloud Server for Optimum Compute Performance

Typically, your 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 pre-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. In the case that you already have an existing subscription, you can bring your own Azure enrollment (BYOE) and your associated VM's etc. Additionally, we also help select the best combination of number of nodes and number of cores. These recommendations are based on our internal benchmarks, which 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.

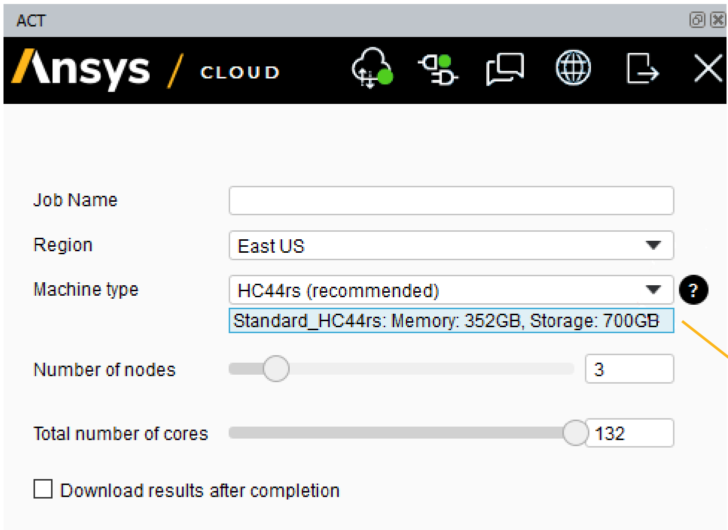

Azure H-Series VM Specifications

	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 HW configuration in Ansys Cloud, you can find VM specifications such as Local Disk space on Azure.

Note that we added 2 new regions (U.S. west and U.S. south central) and our datacenters 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, including data security and GDPR.

Want to **Go back to Menu?**

Ansys Cloud, Configured and Optimized for Mechanical Solver

Figure 4 spotlights solve time for the mechanical direct solver on the HC VM type, supporting typical large and medium models.

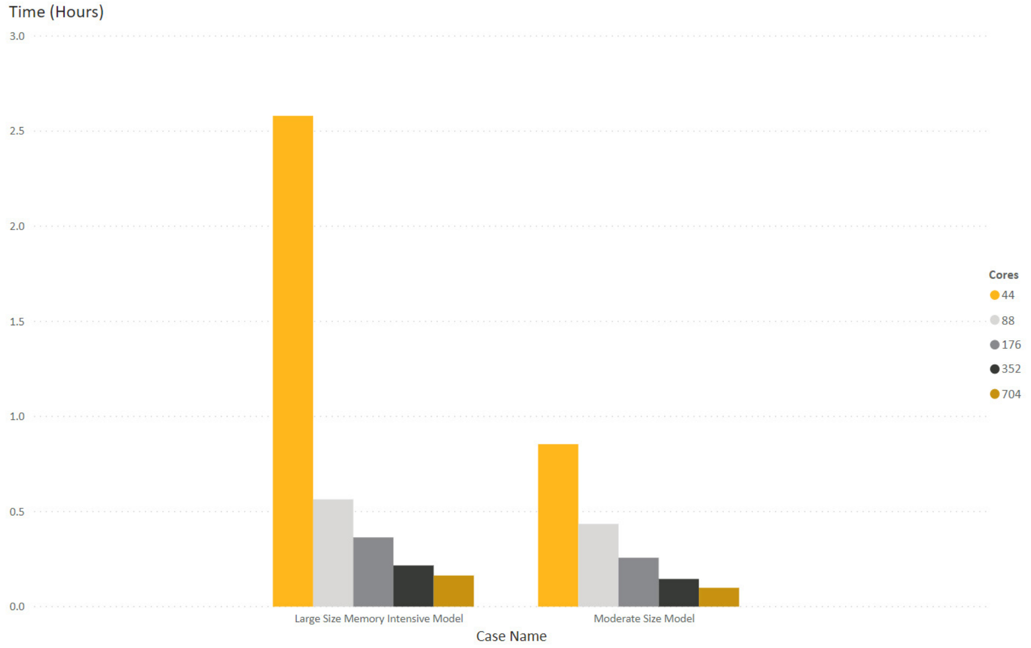

Figure 4: Mechanical Direct Solver solves many times faster on HC, supporting large and moderate sized models.

Figure 5 exhibits similar data for the Mechanical PCG Solver.

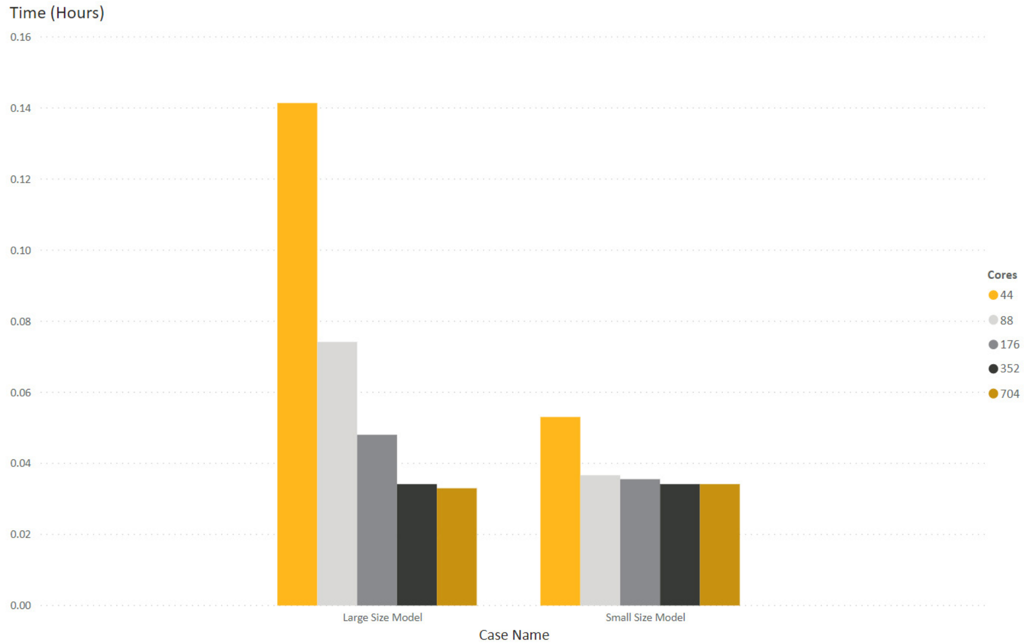

Figure 5: Mechanical PCG Solver solves many times faster on HC, supporting large and small models.

Key takeaways from the above results are as follows:

- **Increasing the number of compute nodes significantly improves performance, especially for memory-intensive problems.**
- **HC hardware shows strong scalability up to 704 cores with the Direct Solver, which is less sensitive to the size and type of analysis.**
- **HC hardware shows good scalability up to 352 cores with the PCG Solver, which is sensitive to the size and type of analysis. For smaller models, we did not observe scalability above 88 cores.**

Figure 6 compares solution times for the HB and HBv2 types at 100% load (all cores) and 50% load (half cores).

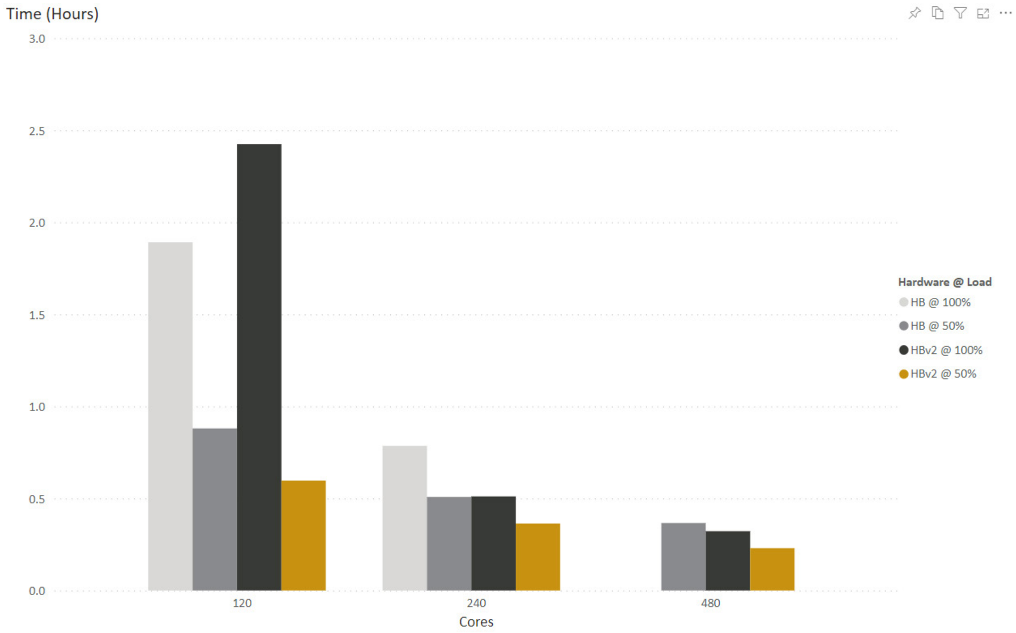

Figure 6: Mechanical Direct Solver solve times on HB and HBv2, supporting large, memory-intensive models

Except for the single-machine configuration, HBv2 outperforms HB at the same total core counts. Interestingly, using two 60-core HBs performs better than a single 120-core HBv2. Furthermore, using four HBs at 50% capacity is more than twice as fast than using two HBs at full capacity, even though the total number of cores is 120 in each case. Similarly, two HBv2s at half capacity outperform a single HBv2 at full capacity. These differences are the result of the increased network and memory bandwidth when running on more machines at half capacity. Results at 240 and 480 total cores show similar patterns.

Figure 7 compares the Direct Solver Solve times on all available VM types.

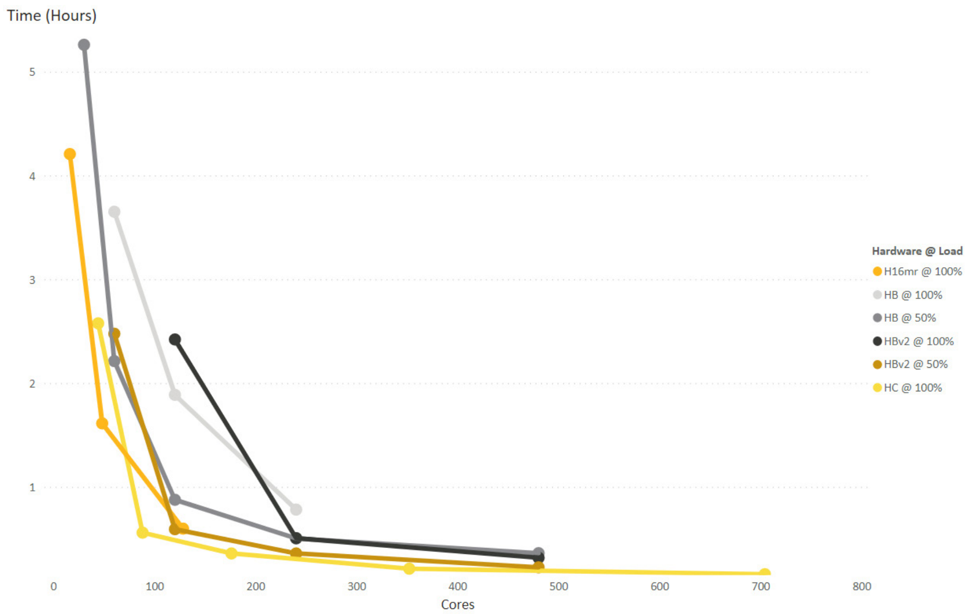

Figure 7: Mechanical Direct Solver time on all available VM types with large, memory-intensive model

This shows that HC generally provides the best performance, followed closely by HBv2 at 50% capacity. Similarly, HB at 50% capacity outperforms HB at 100%, for the same total number of cores.

In summary:

- HC is the recommended type for Mechanical in most situations. HC usually delivers the best performance.
- HBv2 is a good choice for very large models, due to its high RAM, memory bandwidth and interconnect bandwidth, especially when running with half of available cores. With iterative solvers, HBv2 can perform better than HC, especially if you bring your own license (BYOL).
- 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 the Ansys LS-DYNA Solver

Figure 8 shows Ansys LS-DYNA Solve times for a model with 12.5 million elements simulating one second of time.

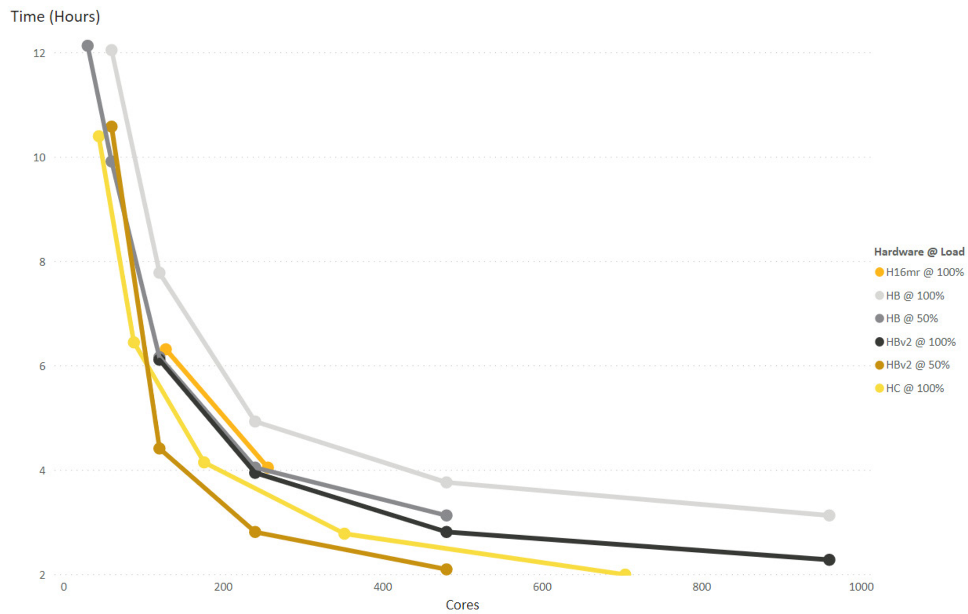

Figure 8: LS-Dyna solver solve time on all available VM types

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.

Looking more closely at the HBs, Figure 9 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.

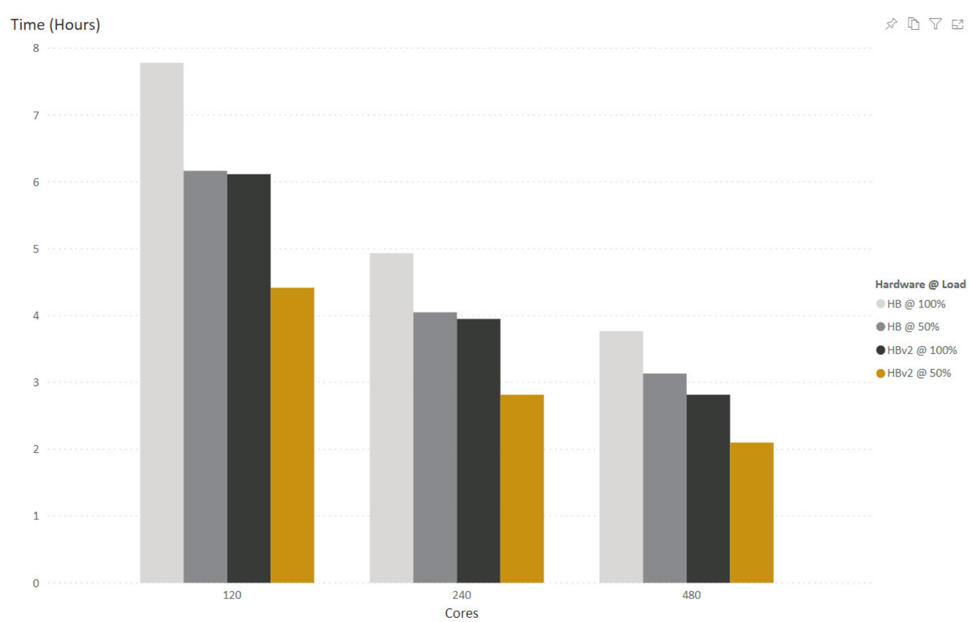

Figure 9: LS-Dyna solver time on all HB and HBv2

In summary:

- HC is the recommended type for LS-DYNA in most situations. It usually delivers the best cost performance.
- HBv2 is also a good choice, 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.

Want to **Go back to Menu?**

/ Ansys Cloud, Configured and Optimized for Electronics Solvers

All Ansys Electronics Desktop products use the same UI 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 10 compares solution times for two HFSS models on the 16-core H16mr, 44-core HC, and 60-core HB.

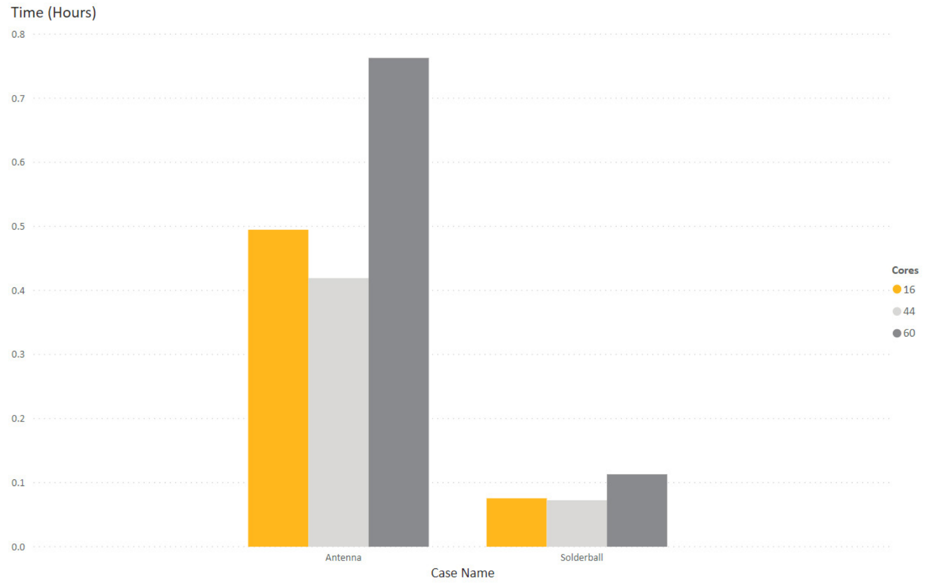

Figure 10: HFSS solution times for two models on H16mr (16 cores), HC (44 cores), and HB (60 cores)

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

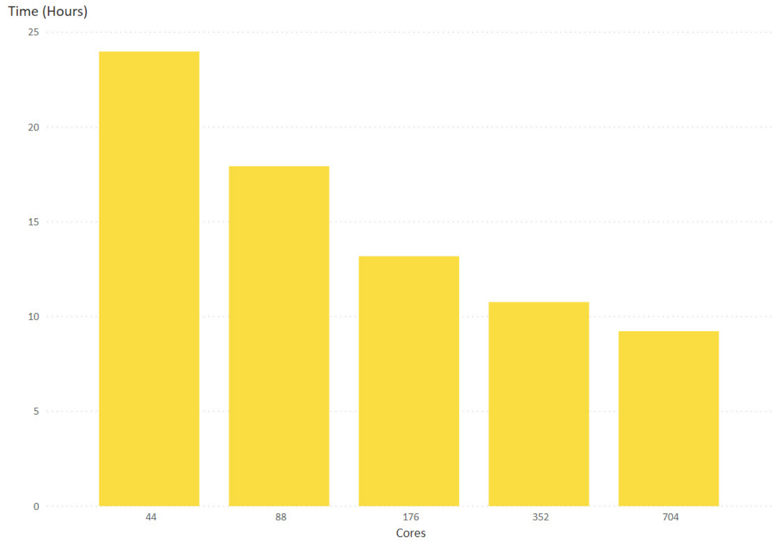

Figure 11: HFSS SO-DIMM on HC

HFSS Anechoic Chamber

This is a simulation of an automobile in an 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 12.

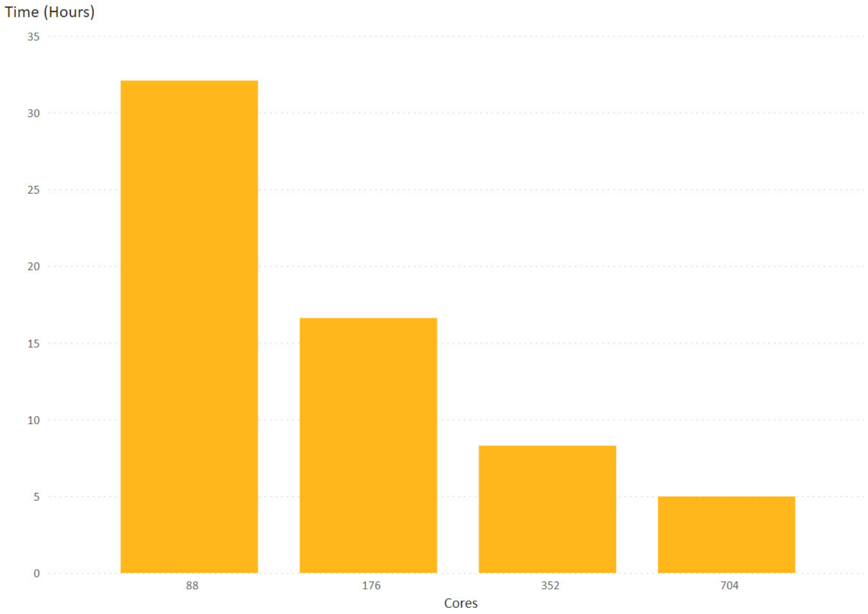

Figure 12: HFSS anechoic test chamber on HC

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 (8 compute nodes), as shown in Figure 13. By employing a larger time-domain, further compute time savings can be achieved with more cores.

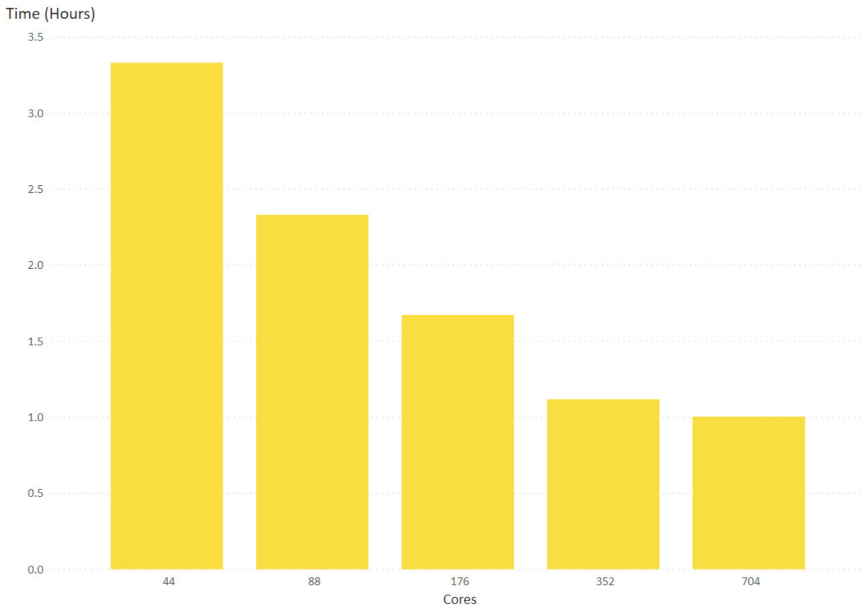

Figure 13: Maxwell on HC

Want to **Go back to Menu?**

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. The goal of these benchmarks was to identify the preferred hardware configuration for different scenarios with fluids solvers running on Ansys Cloud.

Results from three different use 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. The benchmarks compare different hardware configurations and the number of cores used, and they measure the time taken to complete the solution. Note that solver start up and shut down times are included in the reported times.

The first use case is a 71 million cell transient combustion simulation which includes spray injection with DPM tracking. Figure 14 shows that HC performs very well and is the preferred hardware configuration. HBv2 also performs well when utilizing 50% of the available cores; for example, HBv2 @50% on 480 cores is using 8 nodes which have 960 cores available, but only 480 cores are used. HBv2 performs significantly better on a per-core basis when utilizing 50% of the cores because fluids solver performance is highly dependent on the available memory bandwidth. By utilizing half of the available cores, the memory bandwidth per core is doubled, resulting in much better per-core performance. To be clear, the total solve time for HBv2 @50% load is not faster than HBv2 @100% load, but there is not much difference and running at 50% load requires less HPC licensing and is therefore the recommended choice. Note that running at core utilizations other than 50% does not result in the same per-core performance gains; for example, running at 70% core utilization results in similar per-core performance to running at 100% core utilization. While HC and HB hardware configurations will also provide a per-core performance boost by running at 50% load, the improvement is smaller than on HBv2 because the baseline per-core memory bandwidth is higher on HC and HB. We recommend running HC and HB at (or close to) 100% load, since the per-core performance boost by running at 50% load is not sufficient to offset the additional hardware cost.

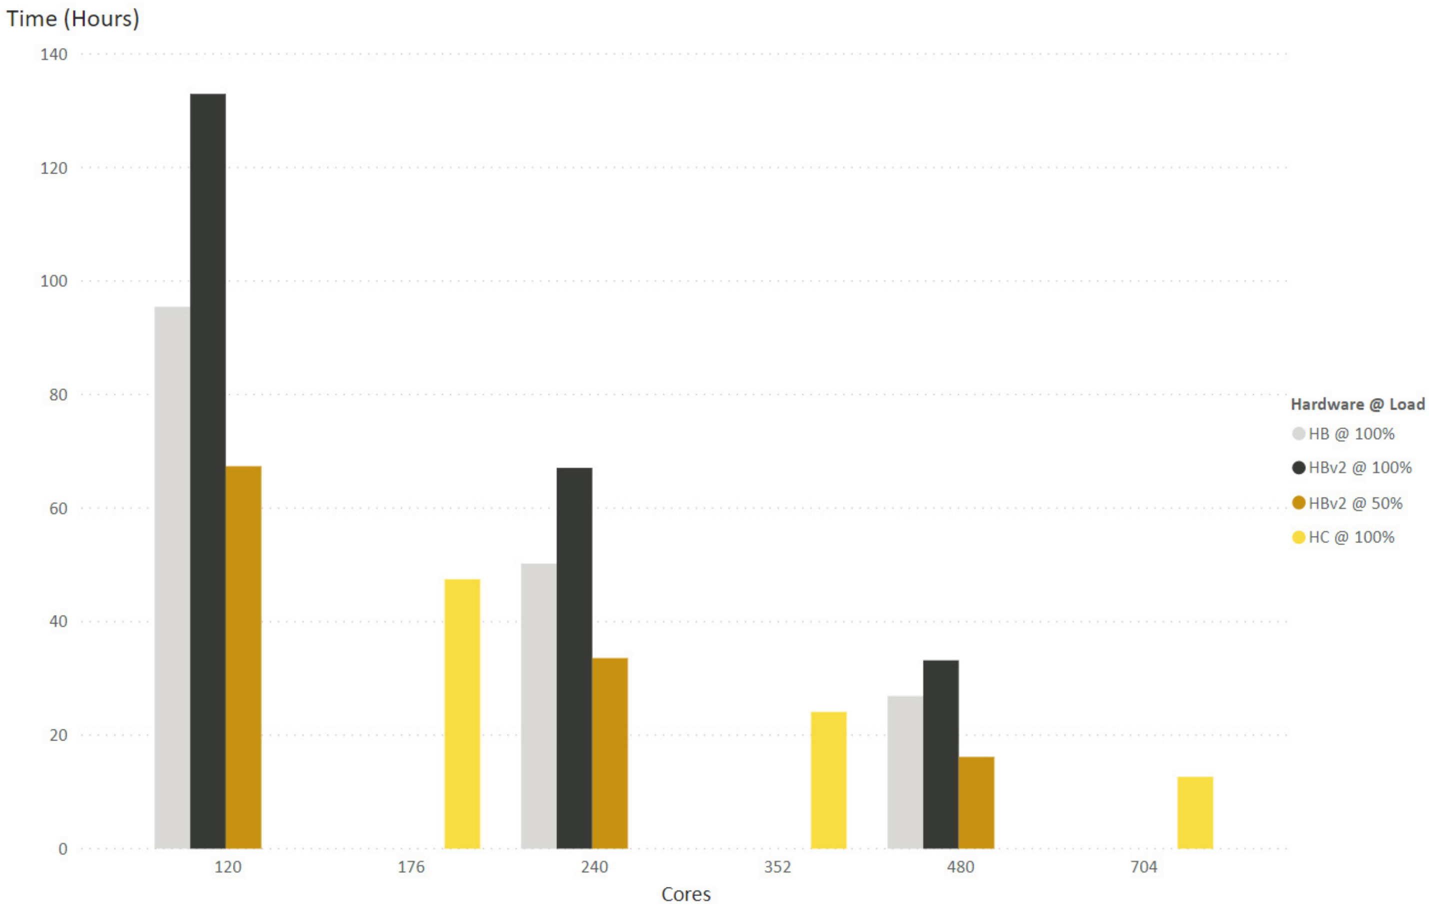

Figure 14: Fluent combustion use case transient 71M, time per VM type

The second use case on the graph below shows a 140 million cell external aerodynamics case for a Formula 1 race car. It displays the same hardware configuration trends as the combustor case, with HC delivering the best performance, closely followed by HBv2 @50% load. The x-axis lists the number of nodes and cores used, for example HBv2*8 480 is using 8 nodes and 480 cores (50% load) whereas HBv2*8 960 is using 8 nodes and 960 cores (100% load). HC and HB are always at 100% load.

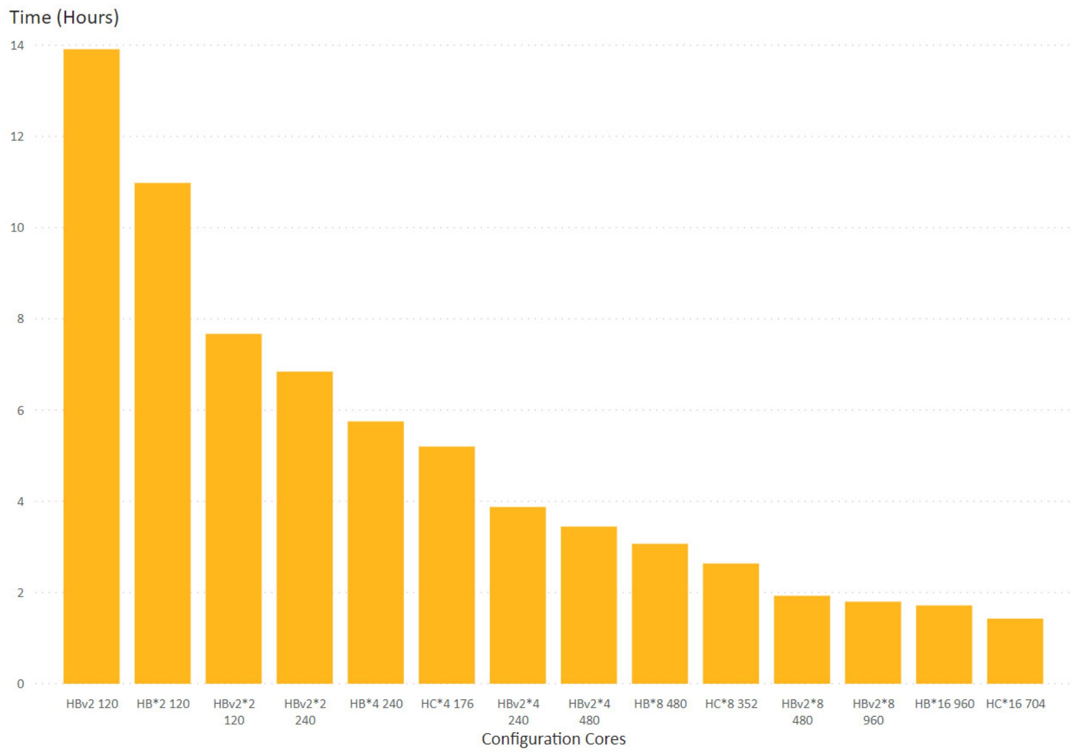

Figure 15: Fluent F1 140M, Time

The last use case is a smaller mesh consisting of 14 million cells, with external flow around an aircraft. Due to the smaller mesh size, we see some drop in scalability at the highest core counts, for example HB*16 960 cores is slightly slower than HB*16 480 cores, which would not be the case on larger mesh sizes. Overall, the trends follow those identified for the combustor and F1 cases.

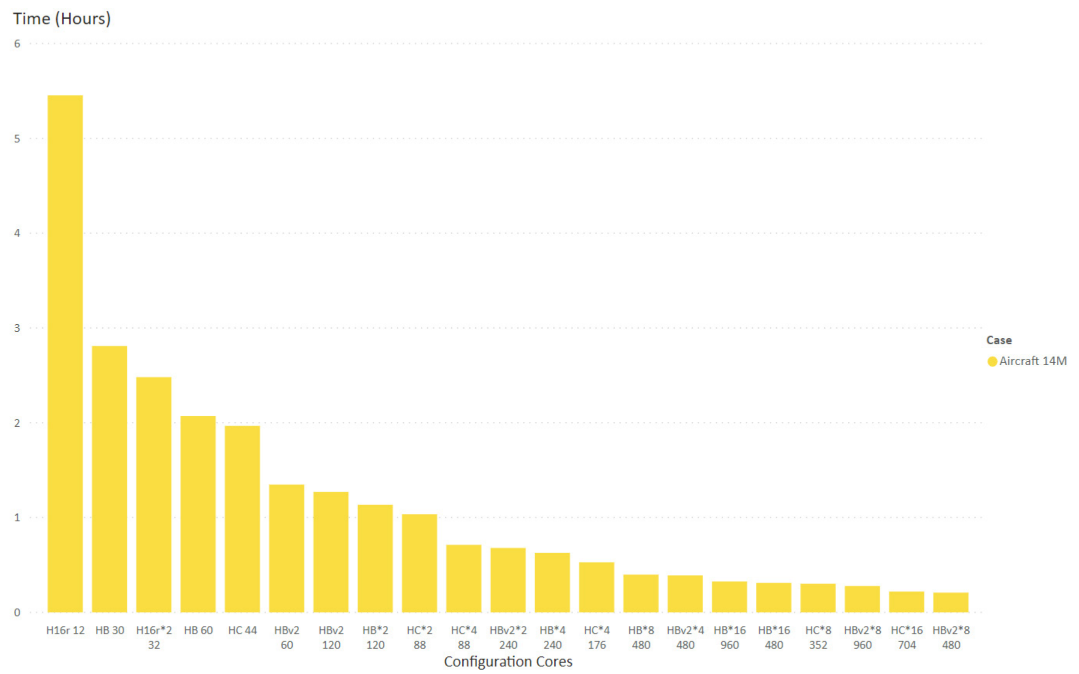

Figure 16: Fluent Aircraft 14M, Time

In summary, for Fluids, we can recommend HC when using elastic licensing or for best performance, with standard licensing. HBv2 is a good alternative to HC when using elastic or BYOL, delivering good performance when using half of the available cores. H16 is suitable for small models and when using one or two HPC Packs. HB provides the lowest hardware cost per job for medium and large models (beyond what can be run on H16), but longer solve times mean this configuration is not recommended with elastic licensing. Note that when using elastic licensing, the hardware cost is less than the licensing cost, so it is generally cheaper to solve on higher core counts utilizing the best performing hardware configuration, which minimizes the solve time.

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 will be good for small models and offer a good price range. HB is good for small models and offers a good price range especially with BYOL. HBv2 is very good for large models when performance matters and is 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 supports from 16 cores up to 120 cores in a cloud-based workstation, available in minutes.

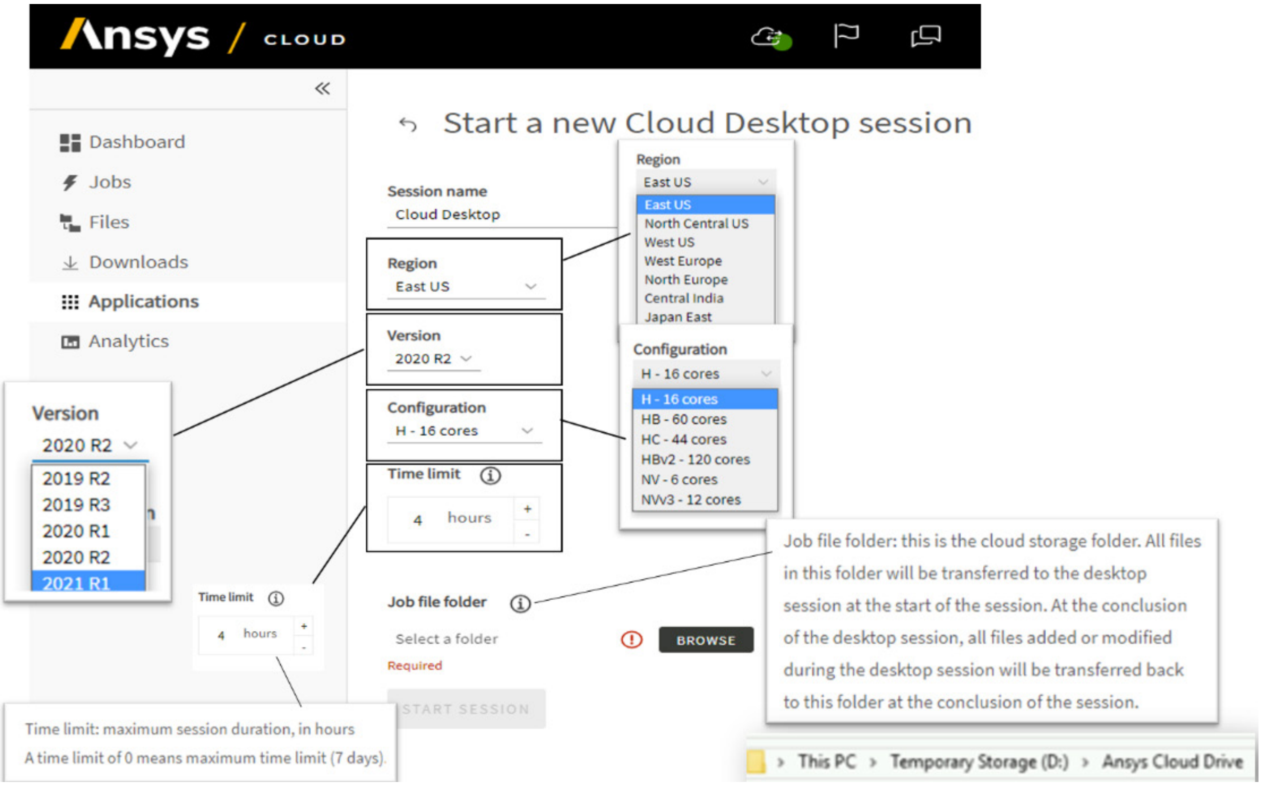

Figure 17: Start a new Cloud Desktop Session

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.

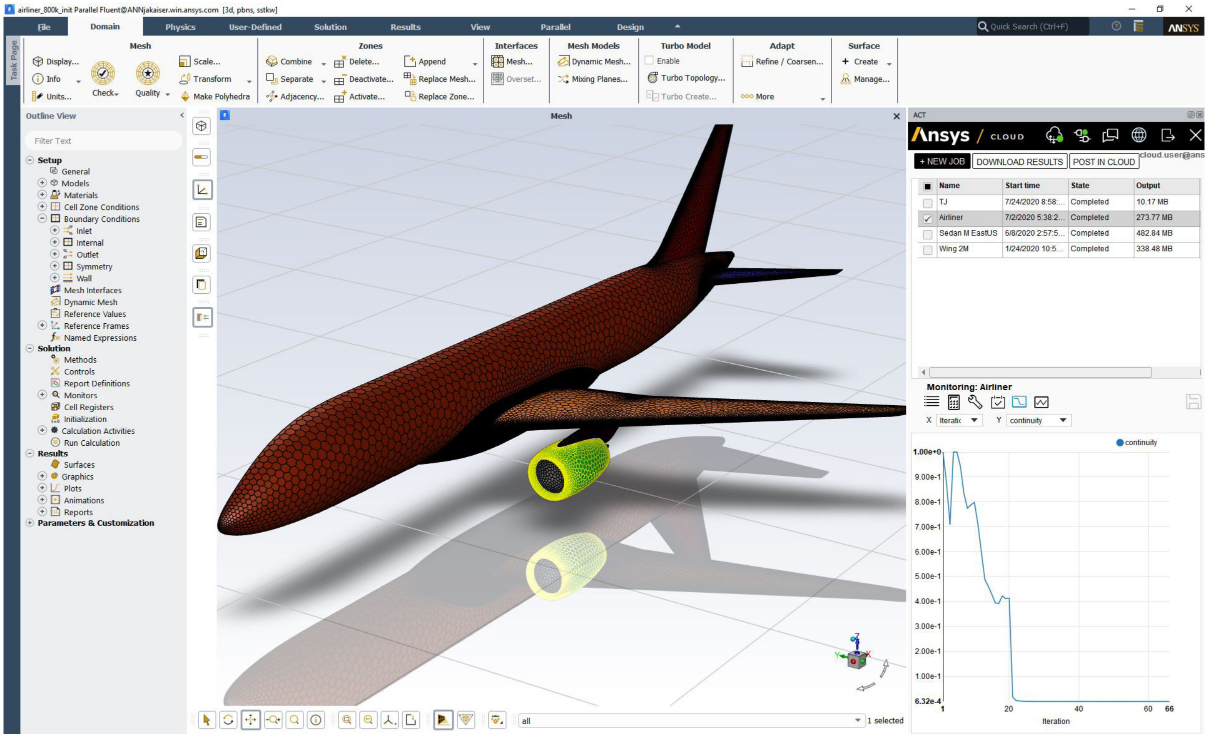

Figure 18: HW config directly through the VDI.

For example, 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 DOE where the solve time is sensitive to the number of cores. When it comes to graphic usage, Nv6 with GPU is recommended. GPU also helps reduce the solve time.

The combination of Multiple 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

If you've ever seen a rocket launch, flown on an airplane, driven a car, used a computer, touched a mobile device, crossed a bridge or put on wearable technology, chances are you've used a product where Ansys software played a critical role in its creation. Ansys is the global leader in engineering simulation. We help the world's most innovative companies deliver radically better products to their customers. By offering the best and broadest portfolio of engineering simulation software, we help them solve the most complex design challenges and engineer products limited only by imagination.

Visit www.ansys.com for more information.