EDR MEDESO

Best Practices

Fluent CloudConnect

Workflow and hardware recommendations

In today's business world, staying competitive is crucial. Cloud computing is a flexible solution to boost efficiency. Cloud computing offers a wide range of benefits, but certain simulation tasks, like geometry preparation, mesh generation, result analysis, and post-processing, work better on dedicated user workstations. So, investing in powerful workstations for these tasks is recommended.

	Local Workstation	Rescale
pre-processing	• Geometry preparation • Meshing • Solver set-up • Validation/Test runs	• Validation/Test runs
solution	• Real time solution monitoring	• Run calculations • Create report files, graphics objects and animations
post-processing	• Post-processing • Solution Analysis • Reporting	• Data storage

Fluent CloudConnect is most effective when the mesh and simulation setup originate from a local workstation. During this setup phase, preliminary test runs are often required. These can be conducted on either the local workstation or the Rescale platform. The user's dedicated workstation should be powerful enough to manage less demanding simulations effectively.

Alternatively, organizations may utilize their proprietary cluster infrastructure in conjunction with cloud computing resources for enhanced computational capabilities. Keep in mind that the computing requirements within the company may evolve over the lifespan of the hardware. Therefore EDRMedeso recommends to phase out all onsite high performance computing environments.

To figure out an optimal local workstation for your analysts, we suggest exploring options on the ANSYS Hardware partners' websites.

Hardware requirements for workstations vary based on simulation complexity and usage. Here are considerations for pre- and post-processing in CFD simulations:

1. Processor (CPU)

- The number of CPU cores significantly impact meshing speed when meshing in parallel with Fluent Meshing.
- When solving with ANSYS Fluent Solver, CPU core scaling is excellent. Additional cores typically decrease the solution time almost linearly.
- In actual usage, minimum of 8 cores is recommended and up to 20 cores in case the workstation is used for solving in addition to pre-processing.
- Latest generation AMD and Intel CPUs are recommended.

2. Memory (RAM)

- Creating the mesh and running simulations will slow down significantly if there is not enough RAM
- A general guideline for a balanced and efficient simulation is to have about

4 GB of RAM for every 1 million cells. However, memory consumption vary between used methods.

- Advised RAM per core for Ansys Fluent is 8 GB. With larger core count system the number can be smaller.
- Opt for high-frequency RAM to optimize memory bandwidth and accelerate data transfer speed.

3. Storage

- CFD simulations generate large datasets that consume a lot of memory. Transient simulations for a large simulation models may consume hundreds of gigabytes of storage. Consider local storage size based on target simulations.
- Solid-state drives (SSDs) are recommended, but Ansys CFD applications are not overbearingly disk read/write intensive so hard-disk drives (HDDs) are acceptable.

4. Other considerations:

- While it is important to have a GPU for post processing purposes, the GPU model itself has limited significance in hardware selection. Even if your solver benefits of GPU acceleration, it is better to do these in the Cloud as the availability to high end GPU accelerators onsite is very limited.
- Ensure your power supply unit (PSU) can handle the power demands of your hardware.
- A fast, reliable network connection is essential for seamless data transfer.

Rescale Cloud Computing Performance Benchmarks

Mixing Tank

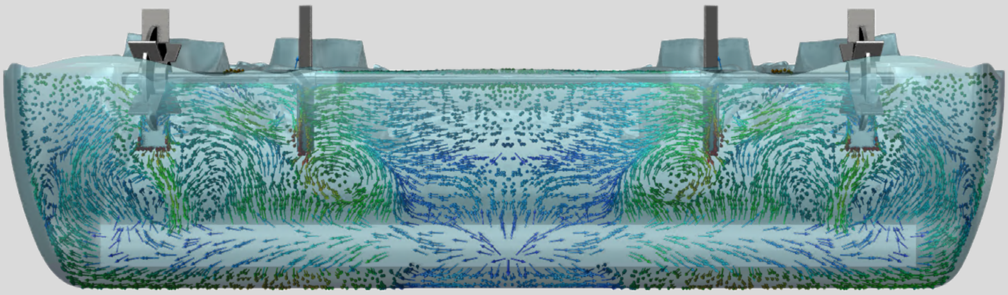

The Case
A steady-state, 3D and two-phase Volume of Fluid (VOF) mixing tank simulation was used to evaluate solution times across different hardwares. The benchmark test involved two personal workstations and five different core configurations in Rescale cloud, which utilized CloudConnect for job submissions to Rescale.

The case has approximately 4.7 million poly-hexcore cells, featuring a symmetry boundary in the middle. Impeller movement was modelled with Moving Reference Frame (MRF). A pressure-based coupled solver with the k-omega viscous model was used to run 2000 iterations. Energy equations was left out.

The table below displays the results. Total time refers to the duration from clicking calculate to until when the results are obtained. For Rescale jobs, the total time encompasses the time it takes to upload files to Rescale, start and stop the cluster, and download files to your local machine.

The time taken for sending jobs to Rescale and receiving results back on the local workstation was consistent across all test cases. The process of writing, preparing, and uploading files (case, data, journal and command) to Rescale, create the job and start the calculation took 1 minute. Downloading the case and data files and reading them took 3 minutes. The workstation employed for this task was a Lenovo laptop. The Rescale Command Line Interface (CLI) was used and the network upload/download speed was 20 MB/s.

In an additional test, the same case was simulated using a poly-hexcore mesh containing 18.3 million cells. Identical settings and geometry were maintained. This time 1000 iterations were solved.

Results

4.7 MILLION CELLS

	Core Count	Processor	System memory	Run time	Total time
Workstation - Lenovo P15v	6	Intel i7-10850H @ 2.70GHz	32.0 GB	local solve	24 h 59 min
Workstation - Fujitsu Celsius R970	20	Intel Xeon Gold 6148 @ 2.40GHz	192 GB	local solve	5 h 48 min
Rescale - Luna	24	2nd Gen. Intel Xeon @ 3.0 GHz	96 GB	4 h 35 min	4 h 44 min
Rescale - Emerald	36	Intel Xeon Platinum P-8124 @ 3.0 GHz	144 GB	3 h 26 min	3h 37 min
Rescale - Hematite	64	AMD EPYC 7V73X @ 2.2 GHz	448 GB	1 h 27 min	1 h 41 min
Rescale - Amber V2	120	AMD EPYC 7742 @ 2.5 GHz	480 GB	1 h 37 min	1h 51 min
Rescale - Rozenite	120	AMD EPYC 7V73X @ 2.2 GHz	448 GB	1 h 15 min	1h 29 min

The usual cost for Rescale hardware was about $14 with priority prices. The cheapest option was only 53 % of the most expensive. Hematite was the least expensive coretype. The running time difference between Hematite and the coretypes with a large number of cores was not significant.

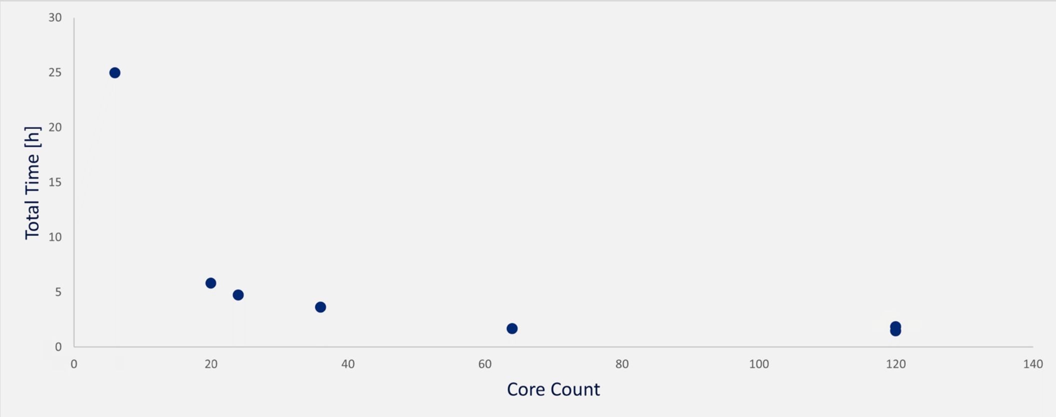

18.3 MILLION CELLS

	Core Count	Processor	System memory	Run time	Total time
Workstation - Fujitsu Celsius R970	20	Intel Xeon Gold 6148 @ 2.40GHz	192 GB	local solve	13 h 59 min
Rescale - Hematite	64	AMD EPYC 7V73X @ 2.2 GHz	448 GB	3h 32 min	3h 55 min
Rescale - Amber V2	120	AMD EPYC 7742 @ 2.5 GHz	480 GB	3h 36 min	4h 2 min
Rescale - Rozenite	120	AMD EPYC 7V73X @ 2.2 GHz	448 GB	3h 3 min	3h 27 min

In the scenario with 18.3 million cells, the average cost was $25. Hematite continued to be the most budget-friendly coretype. The Hematite CPU is a bit newer than the one in the Amber V2 and the test case includes the multiphase model, likely impacting scalability negatively. As a result, it appears that the limit for efficient solution time is around 64 cores per node, and adding more cores beyond this point does not seem to decrease the solution time significantly.

Reflection

Using Rescale delivers significant advantages. Even with similar core counts the Rescale run outpaced the performance of the local Workstation. This is attributed to Rescale's provision of the latest technology hardware components, while on-premise hardware tends to become outdated. The impact of using newer technology on simulation times can also be observed within Rescale. For instance, the CPU used in the Amber V2 was launched in 2019, while the CPU in the Hematite, launched in 2021, represents more recent technology.

The up-front costs of the tested workstations are 2500 € and 8000 €, but these number do not consider the additional costs for maintenance and upkeep. Furthermore, on-premise hardware leads to employee downtime while awaiting results. The cost of employees' work hours is considerably higher than the hardware cost on Rescale.

During simulation runs on workstations, engaging in other activities is either impossible or very slow. Cloud computing, on the other hand, frees up the simulation engineer's personal workstation for additional tasks, such as pre- or post-processing for another case. However, monitoring the job that is running in Rescale is straightforward in real-time.

In our test case, finding time to execute benchmark tests on local workstations was challenging. Particularly, simulations with six cores that took more than one night to solve had to be scheduled for weekends. This delay wouldn't work for real-life situations needing quick results Rescale enable faster simulations and handles multiple simulations at once when necessary.

The 18.3 million cells simulation couldn't have been possible on the 6-core laptop. The laptop's memory was insufficient even for efficient mesh creation. However, on the 20-core workstation, meshing and case setup were seamless. Sending the job to Rescale was also straightforward. Using Rescale's cloud computing opens up new possibilities to simulate processes that were impossible before due to hardware limits.

Selecting the Ideal Core Type for Your Job

Rescale offers users the flexibility to select from a range of core types to power their cloud simulations. However, for newcomers, determining the most appropriate core type for their specific requirements can often be a challenging task. For guidance, recommended hardware types for Computational Fluid Dynamics (CFD) simulations are outlined in the table below and different factors affecting hardware needs are discussed later on.

Follow these five steps to find the right core type for you

1. Determine the typical parallel scalability limit using following equation
   $$cores = \frac{cells}{40000}$$

2. Refer to the table below and identify the coretype with the closest match to the required number of cores. Note that Emerald offer configurations with 1, 2, 4, 8 and 18 cores per node and Luna 24 cores per node. Choose the coretype with the most cores per node instead of using multiple smaller nodes.

3. Run your job using the selected coretype.

4. Check if the hardware features of the selected coretype met your needs. Learn how to monitor Rescale hardware performance and costs for detailed instructions.

5. For repeated simulations, run benchmark tests to find the most suitable core type for your needs.

Typical Coretypes for CFD simulations

Coretype	Processor	Cores/node	Clock speed	Memory/node	Storage/node
Emerald	Intel Xeon Platinum P-8124 (Skylake)	36	3.0 GHz	144.00 GB	1296 GiB
Carbon	Intel Xeon Platinum 8168 (Skylake)	44	2.7 GHz	352.00 GB	700 GiB
Luna	2nd Generation Intel Xeon Scalable Processors (Cascade Lake)	48	3.0 GHz	192.00 GB	1728 GiB
Hematite	AMD EPYC 7V73X (Milan-X)	64	2.2 GHz	448.00 GB	1920 GiB
Amber V2	AMD EPYC 7742 (Rome)	120	2.5 GHz	480.00 GB	960 GiB
Rozenite	AMD EPYC 7V73X (Milan-X)	120	2.2 GHz	448.00 GB	1920 GiB

Other considerations

Choosing the right coretype for cloud simulations is important for achieving optimal performance and cost-effectiveness. It’s not always about opting for the cheapest core hours; faster processors deliver results more quickly. This means you’ll need to purchase fewer hardware hours, and more importantly, it saves your employees’ time.

ANSYS Fluent exhibits excellent parallel scalability, meaning that more cores generally accelerate your simulations up to a certain limit. This limit typically falls in the range of 30,000 to 50,000 cells per core. However, the mesh size is just one of several factors influencing parallel scalability and solution time.

Consider the following factors when choosing the appropriate core type for your simulations:

- Solver Type and Methods: The simulation's complexity varies depending on the chosen numerical behavior.
- Physics Models and Boundary Conditions: More intricate physics require higher computational power.
- Steady-state or Transient: Transient simulations often require more iterations due to time-dependent calculations and the solution time might be much longer.
- Mesh Adaption: The mesh size may differ significantly during the simulation.
- UDFs, Calculation activities or other computationally heavy additions
- Mesh Complexity: Alongside mesh size, the mesh type also impacts computational demands.
- Memory Capacity: Ensure your chosen coretype provides sufficient RAM to avoid memory-related bottlenecks.
- Consider GPU: ANSYS Fluent native GPU solver support some of the simulation models. If applicable, consider GPU solver and coretypes with GPU options for enhanced performance.
- On-Demand Priority or On-Demand Economy: On-Demand Economy offers a lower unit price but run time and costs may increase if there are multiple stops during the calculation.

Monitor Rescale hardware performance and costs

Hardware performance

Keep an eye on two important performance factors when using cloud hardware resources for simulations. Ensure that the average CPU consumption is close to 100%, and make sure there is available free memory in the cluster. You can check the cluster status both during the simulation and immediately afterwards in the job status page, example shown in the image below.

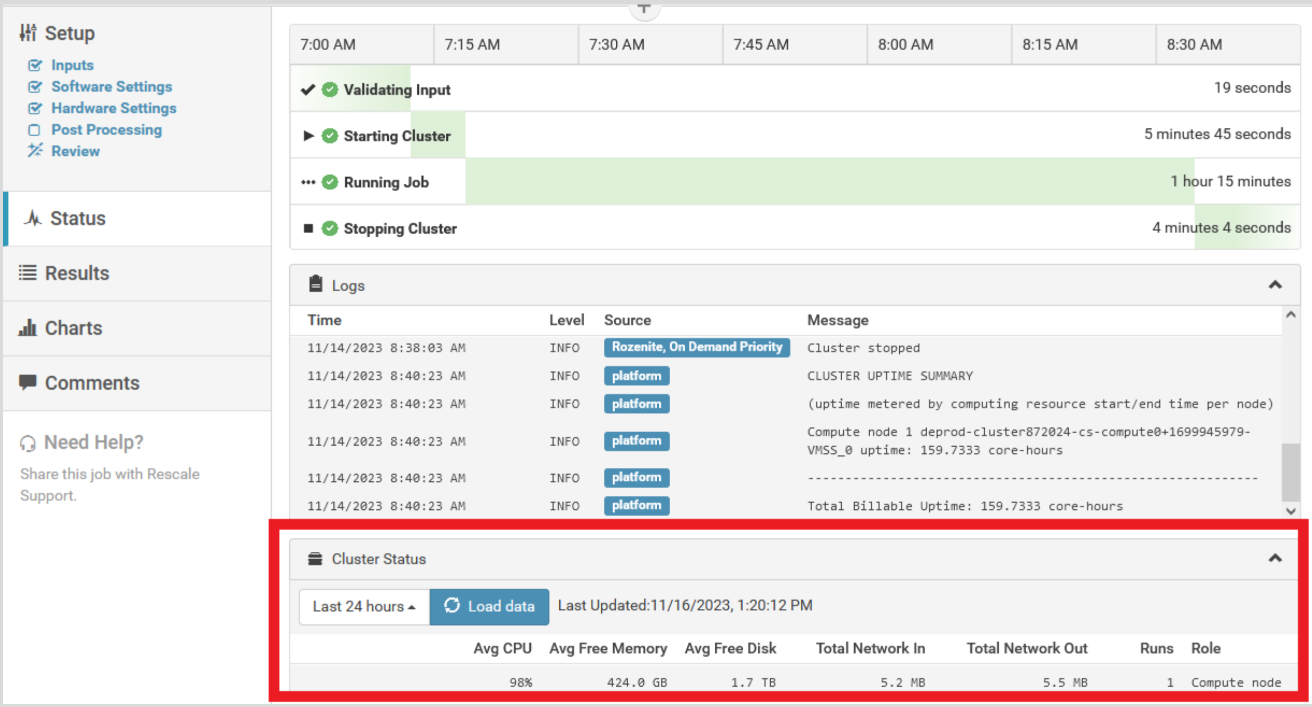

Another option for monitoring hardware consumption during or after solving is to use automatically generated hardware consumption images. These images provide information on CPU consumption, RAM usage, as well as available memory and storage space. You can find these images in the file list, named as follows: all_metrics.png, cpu.png, disk_used.png and mem_used.png.

If the CPU utilization averages 100%, as shown in the image below, it means that all CPU units are being effectively utilized, and parallel scaling is well-executed. In such cases, testing with more cores can be considered. However, if the average line is below 100%, the job does not scale further to that core count, and the cluster core count is not optimally utilized. Subsequent calculations should be performed with fewer cores.

When the memory utilization is near 100%, there is a risk that the RAM may not handle the memory requirements. It's crucial to ensure that RAM does not exceed its capacity during the solve. Solution might crash if there is not enough RAM. If the maximum memory utilization is near 90%, it's advisable to consider a core type with larger memory for your next similar job. You can calculate what was the required RAM from the memory utilization or memory usage images.

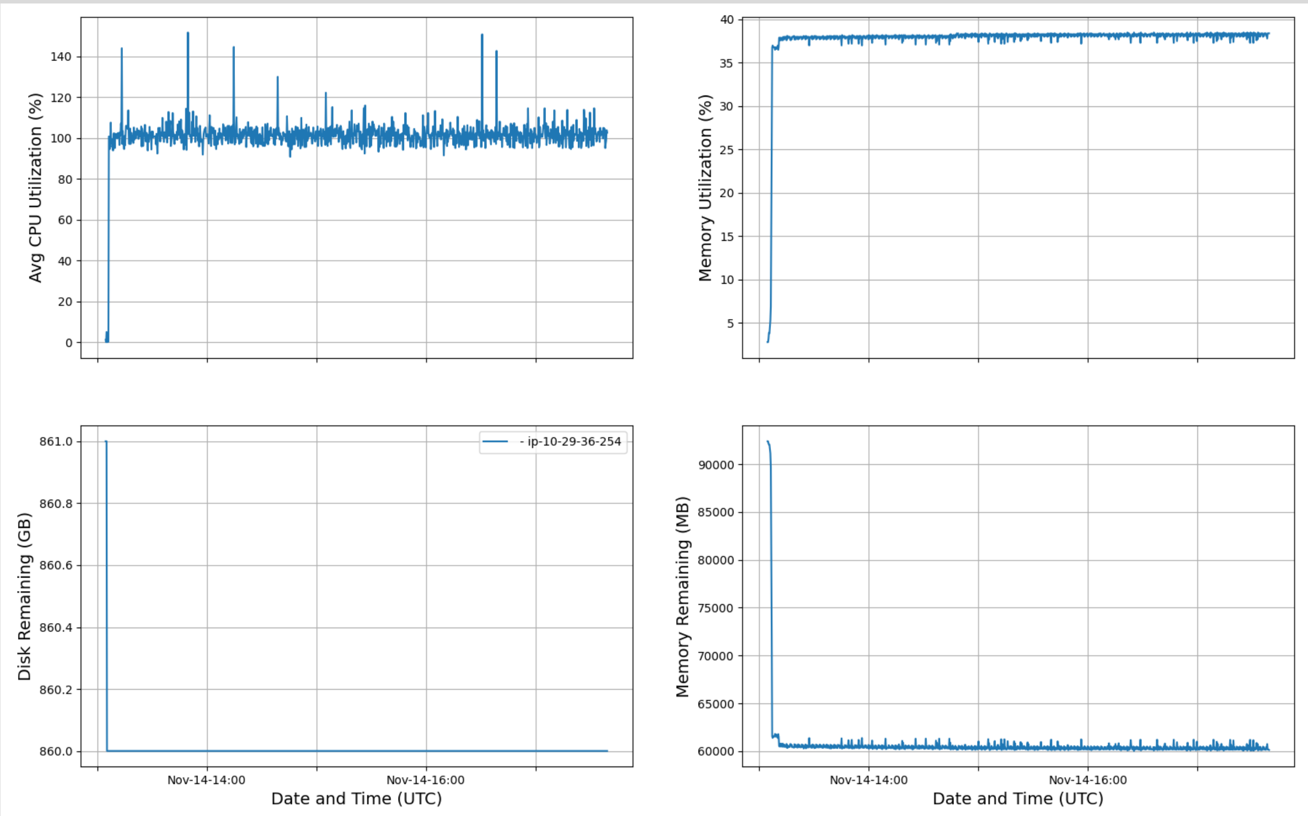

Charges

You can keep track of your cloud computing expenses and usage during and after a simulation by going to **User Profile** and selecting **Charges/Usage**. On this page, you'll find the total consumption for each month. You can also download a CSV file that separates the costs for each job within a month. Remember, each job has two types of costs: hardware and platform. To find the total cost for a job, just add up these two values.

Modify boundary conditions during calculation with Remote Visualization Client

The Ansys Fluent Remote Visualization Client is a tool that lets you monitor and make changes to various aspects of your simulation while it's running. You can edit things like boundary conditions, cell zone conditions, solution methods, create animations and perform many other tasks during the calculation. With this client, you have the ability to send any Text User Interface (TUI) command to the cluster. This means you can control and adjust the simulation without stopping it. The TUI commands take effect right away, even if the job is currently running.

Check out the video below for more details on how to use the Ansys Fluent Remote Visualization Client:

Click here to watch the video