Performance of Ansys Mechanical on Oracle Cloud Infrastructure

October 2022, version 1.0

Public

Disclaimer

This document in any form, software or printed matter, contains proprietary information that is the exclusive property of Oracle. Your access to and use of this confidential material is subject to the terms and conditions of your Oracle software license and service agreement, which has been executed and with which you agree to comply. This document and information contained herein may not be disclosed, copied, reproduced, or distributed to anyone outside Oracle without prior written consent of Oracle. This document is not part of your license agreement, nor can it be incorporated into any contractual agreement with Oracle or its subsidiaries or affiliates.

This document is for informational purposes only and is intended solely to assist you in planning for the implementation and upgrade of the product features described. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described in this document remains at the sole discretion of Oracle. Due to the nature of the product architecture, it may not be possible to safely include all features described in this document without risking significant destabilization of the code.

Revision History

The following revisions have been made to this document.

DATE	REVISION
October 2022	Initial publication

Table of Contents

Purpose

This paper discusses the performance of Ansys Mechanical using Oracle Cloud Infrastructure (OCI). Benchmarks were performed with the latest set of Ansys Mechanical benchmarks on various compute cluster systems. The goal was to determine the best overall configuration to carry out a production simulation workload in a cloud environment that provides performance comparable to leading dedicated on-premises hardware systems.

Background

Manufacturing faces several challenges, from increased competition in the world economy, with an increased burden of regional regulations and market peculiarities, to spontaneous events such as tsunamis, pandemics, global supply chain issues, inflation, and regional conflict. These issues place pressure on manufacturers to accelerate product development to serve a broad, rapidly changing market.

To improve the speed and quality of their product development, manufacturers increasingly rely on virtual product development driven by computer-aided engineering (CAE). As the need for CAE increases, so does the need for manufacturers to harness the high-performance computing (HPC) resources required to improve both product quality and the time-to-market of virtual product design. However, many manufacturers are struggling to maintain their existing IT, let alone improve it to facilitate the increased CAE capabilities that they need. As a result, the growth of HPC in the cloud is now outstripping the traditional in-house, on-premises-based HPC IT used for CAE. Oracle is working with its leading technology partners, such as Ansys, to help meet this growing need for power and accessible HPC CAE resources in the cloud.

Ansys Mechanical is an industry-leading finite element analysis (FEA) solver with structural, thermal, acoustics, transient, and nonlinear capabilities. Implicit FEA problems often place large demands on memory and disk I/O subsystems, particularly for out-of-core solutions, where the problem is too large to fit into the available system RAM. Because of these characteristics, manufacturers have tended to avoid the cloud for resourcing their implicit FEA workloads. However, Oracle has designed its second-generation cloud infrastructure with resources fully capable of delivering outstanding performance for implicit FEA workloads such as Ansys Mechanical. This infrastructure typically matches the latest generation of on-premises hardware built specifically for implicit FEA workloads.

Benchmark Systems

Benchmarks were conducted on OCI within virtual HPC clusters similar to the one shown in Figure 1.

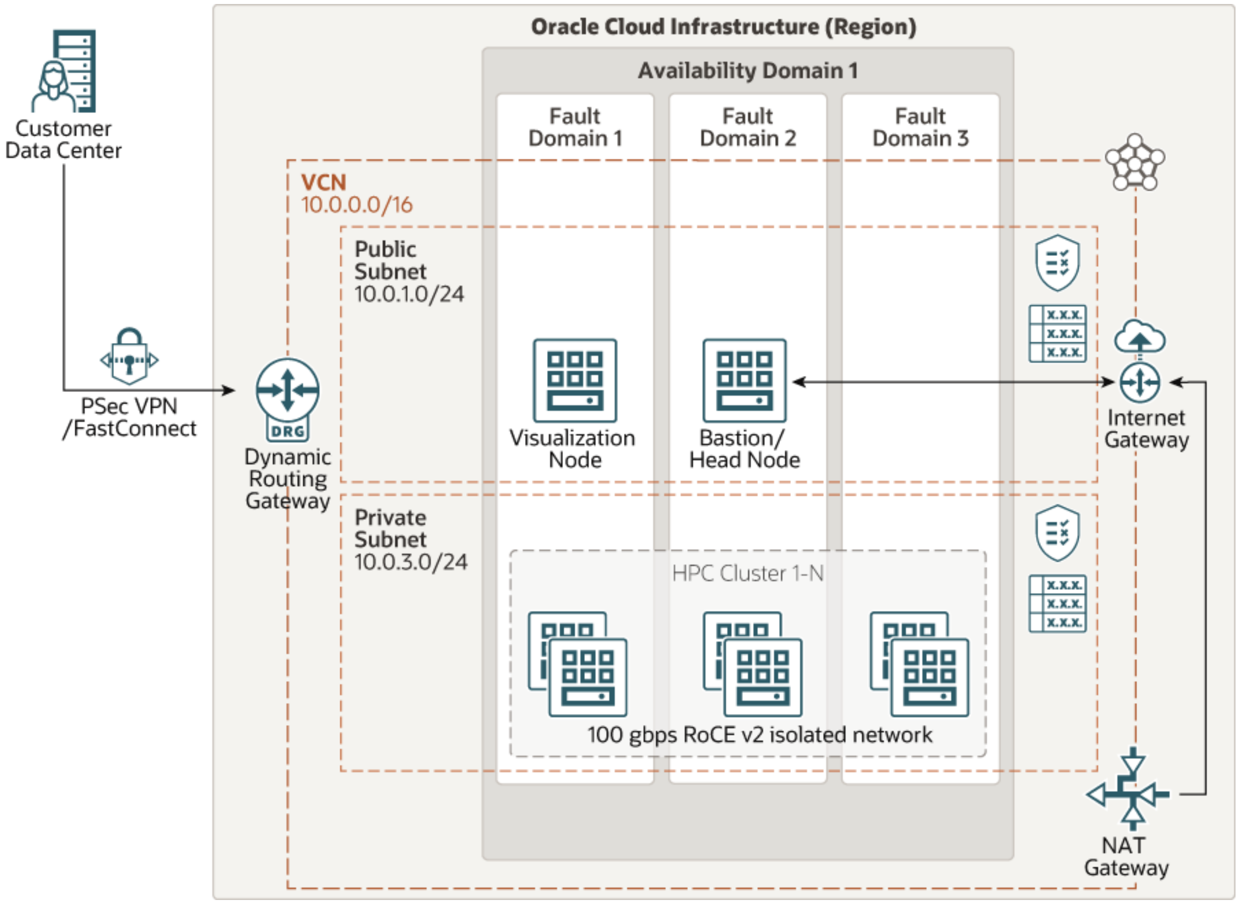

Figure 1: Typical OCI HPC Configuration

For a description of such a cluster with its components, see the Deploy HPC on OCI reference architecture.

In OCI, it's possible to quickly create a variety of clusters, which can lead to some confusion over which are best-suited for various workloads. For HPC performance, the following components are key and are the focus of this paper:

1. Compute nodes
2. Network interconnect
3. File storage

These components aren't completely independent within OCI. For example, only select compute node shapes can interconnect with the highest-performing 100-Gbps RDMA network interconnect. For descriptions of the various compute node shapes available in OCI, see Compute Shapes in the documentation.

Table 1 lists the compute node shapes tested in this paper. These nodes were equipped with Intel-based Xeon and AMD EPYC processors from various generations, and with varying amount of memory.

Table 1: Benchmark Node Description

SHAPE	TYPE	PROCESSOR	CORES	MEMORY	STORAGE	NETWORK
BM.Standard2 (BM.S2)	Bare metal	Intel 8167M	26x2	768 GB	Block	2x25 Gbps
BM.HPC2	Bare metal	Intel 6154	18x2	384 GB	6.4 TB NVMe	1x100 Gbps RDMA
BM.Standard3.64 (BM.S3)	Bare metal	Intel 8358	32x2	1024 GB	Block	2x50 Gbps
BM.Optimized3 (BM.O3)	Bare metal	Intel 6354	18x2	512 GB	3.8 TB NVMe	1x100 Gbps RDMA
BM.Standard.E4.128 (BM.E4)	Bare metal	AMD 7J13	64x2	2048 GB	Block	2x50 Gbps
VM.Standard3.Flex (VM.S3F)	Virtual machine	Intel 8358	32	512 GB	Block	1x32 Gbps
BM.GPU4.8	Bare metal	AMD 7542 NVIDIA A100	64 8	2048 GB 40 GB	27.2 TB NVMe	8x200 Gbps RDMA

In OCI, both virtual machine (VM) and bare metal node shapes are available for HPC. Although bare metal shapes typically work best for HPC, VM shapes are suitable for some workloads because of their potential for increased flexibility and ease of use. With bare metal shapes, the user has all the node's memory and CPU cores at their disposal, and they're charged accordingly. The VM shapes provide some degree of variable sizing, depending on how many processor cores are requested, which allows for flexible pricing. For the VM.Standard3.Flex test in this paper, 512 GB of memory was requested, which is the minimal allowed for the use of 32 cores per system.

All node shapes contained some networked block storage for the boot image. Three shapes—BM.Optimized3, BM.HPC2, and BM.GPU4.8—also contained additional local NVMe storage. By default, Ansys Mechanical places temporary files in the execution (or working) directory. For the nodes with local NVMe storage available, Ansys Mechanical was benchmarked using the NVMe-based local file system as the working directory, which allowed the bulk of the simulation I/O to benefit from the high performance provided by direct-attached NVMe. For the other shapes, the working directory was placed in the networked block storage. The amount of local storage used for each benchmark varied. All attempts were made to use the in-core equation solver to minimize the amount of I/O traffic and storage needed.

Additionally, those three shapes, optimized for HPC workloads, were interconnected with NVIDIA ConnectX-based high-speed RDMA networks, which were used for the multi-node benchmarks. All of these benchmarks were carried out on systems using Oracle Linux, typically version 7.9, with version 8 used for the BM.Standard3.64, BMStandard.E4.128, and VM.Standard3.Flex shapes.

Although overall performance is usually critical for HPC, it's subject to cost constraints. As a result, the total cost of a particular simulation is of comparable importance. The total simulation cost includes hardware cost, software cost, support cost, and various human cost such as the simulation engineer's time. Indicative pricing for the various OCI shapes is located on the Compute Pricing page.

The compute cost often dominates the overall cloud hardware cost for HPC, and the hardware cost often relates directly to the number of cores used. Software licensing costs for premier applications such as Ansys Mechanical can be complex and variable, but they're often associated with how many compute cores are used for each job. Although there's no simple formula to directly calculate the total cost that a customer would incur for a simulation with a particular cloud-based system, it's reasonable to assume that the overall optimal cost can be obtained with systems that have fewer total cores. This suggests that a superior performance-per-core metric would lower both the hardware and application software costs per simulation. This assumption is examined further in this paper.

Single-Node Results

All benchmarks in this paper were obtained by using version 2022R1 of Ansys Mechanical for the benchmark sets V22Benchmark-Workstation and V22Benchmark-Cluster obtained from Ansys. A description of these benchmarks is provided in Tables 2 and 3.

Table 2: V22Benchmark-Workstation Descriptions

DATASET	SOLVER	SIZE
V22direct-1	Sparse solver, transient, nonlinear	10 MDOFs
V22direct-2	Block Lanczos eigensolver with 50 modes	3.4 MDOFs
V22direct-3	Sparse solver, transient, nonlinear	16.1 MDOFs
V22iter-1	PCG solver, static, linear	23.3 MDOFs
V22iter-2	PCG solver, static, linear	19.2 MDOFs
V22iter-3	PCG Lanczos eigensolver with 10 modes	25 MDOFs

Table 3: V22Benchmark-Cluster Descriptions

DATASET	SOLVER	SIZE
V22direct-4	Sparse solver, nonsymmetric, transient, nonlinear, thermos-electric coupled field	4 MDOFS
V22direct-5	Sparse solver, transient, nonlinear	15 MDOFs
V22direct-6	Sparse solver, transient, linear regression, 1 frequency	11 MDOFs
V22iter-4	JCP solver, static, linear	30 MDOFs
V22iter-5	PCG solver, static, linear	63 MDOFs
V22iter-6	PCG solver, static, linear	125 MDOFs

The benchmarks were created in two groups, Workstation (1,2,3) and Cluster (4,5,6), based on problem size, with the Cluster benchmarks being larger than the corresponding Workstation benchmarks. Each group contained benchmarks using both iterative and direct solvers. Benchmarks were run on systems in a dedicated fashion with only one job running on a system (or a cluster of systems) at a time. All results were reported in terms of overall job performance based on the Ansys solver core rating. This rating is defined as the number of solver jobs that can be run in a 24-hour period. Results with a higher performance rating run quicker than comparable jobs with a slower performance rating, making higher better. These performance ratings directly correspond to both the elapsed time to run a single job and the number of jobs that could be run in a day in an inverse manner, and would be used in predicting expected cluster performance and capacity.

Although Ansys Mechanical offers SMP, DMP, and hybrid parallel methods, our experience has been that most simulations perform better using the DMP method, even on a single node. All the benchmarks in this paper were performed using the DMP method, assigning one MPI domain to each available core.

Figure 2d shows the single-node results obtained on the OCI shapes listed in Table 1 for the direct solver datasets. Figure 2i shows the single-node results for the corresponding iterative solver datasets.

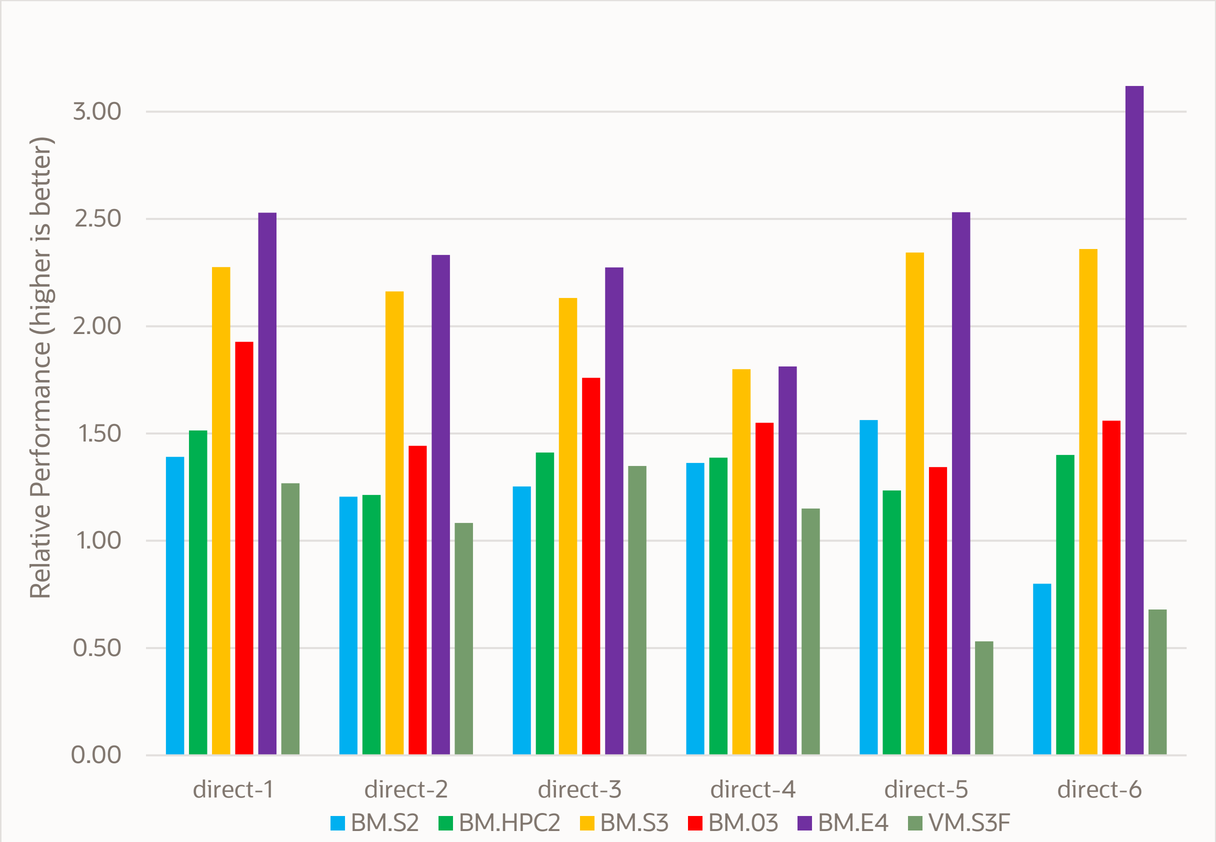

Figure 2d: Ansys Mechanical Single-Node Direct Solver Benchmarks

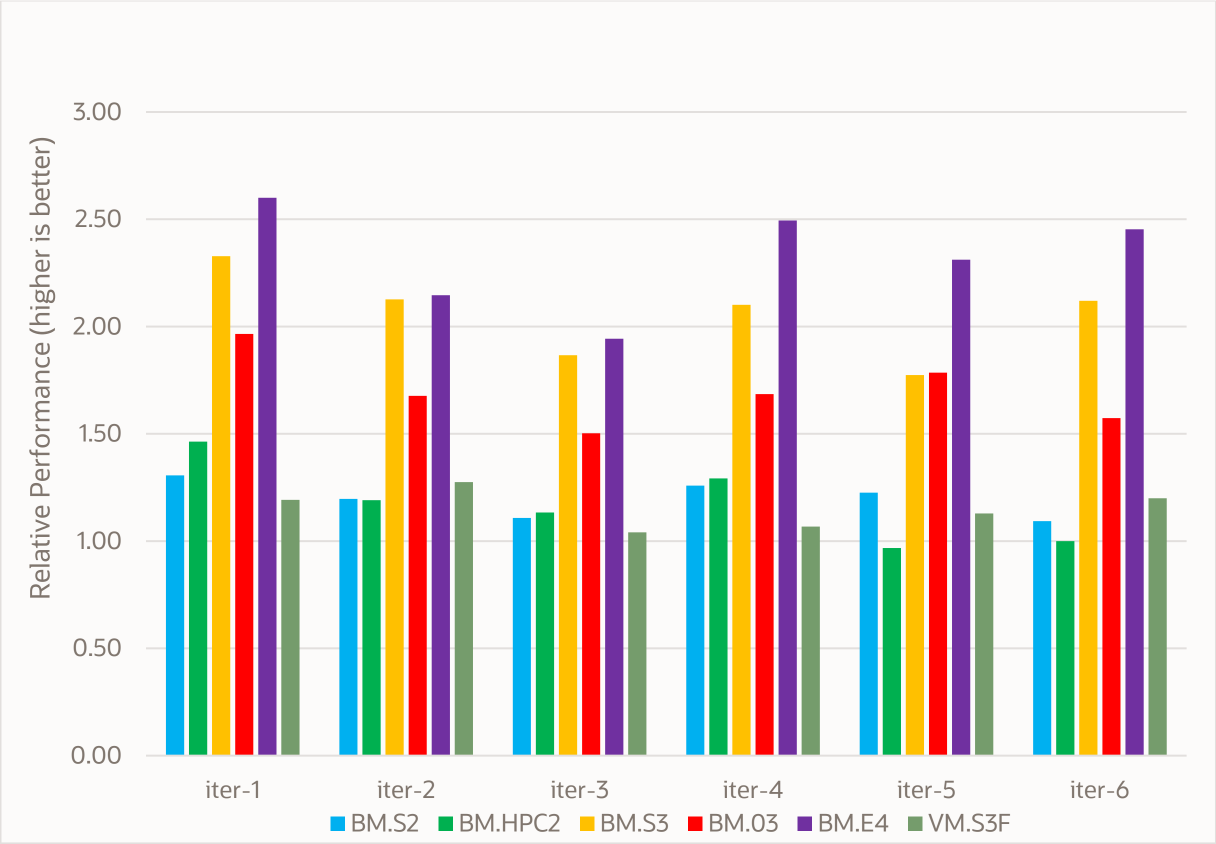

Figure 2i: Ansys Mechanical Single-Node Iterative Solver Benchmarks

The performance reference was based on the results obtained by using the 32-core results of the BM.Standard2 shape, which is similar to many customers' current HPC production systems with dual 16-core Intel Xeon Gold 61xx series processors. This reference should be helpful for customers who want to migrate off their current systems.

The results shown in this paper offer the most insight when they're viewed as trends, not as individual performance numbers. These results were obtained on modest-sized clusters typically running one job on the cluster at a time. As such, they were highly repeatable. For benchmarks such as this, a normal deviation is typically from 1% to 3% from run to run. Using a system with modern processors in turbo mode induces a performance distribution in this range alone. Typically, customers see greater variation in their production workloads when they encounter shared resource constraints for multiple concurrent jobs, typically involving shared storage and occasionally shared network access. Best practices for CAE HPC involve methods to minimize the shared constraints, such as using local file storage rather than shared file storage. With the use of such practices, the performance observed here should be attainable in a large production scenario in OCI.

The relative performance among the various shapes remained constant throughout the entire range of datasets. As a result, the choice of the optimal shape for a particular simulation type shouldn't be wildly variable for different classes of simulations. For all the datasets except one, the following bare metal shapes significantly outperformed the others:

- BM.Standard3.64 and BM.Optimized3 (Intel 3rd Gen Intel Xeon Scalable processor)
- BM.Standard.E4.128 (128-core-based AMD EPYC 7003 processor)

The exception was the V22direct-5 dataset, in which the 52-core-based BM.Standard2 shape slightly outperformed the 36-core-based BM.Optimized3 shape.

The only VM shape tested, VM.Standard3.Flex, significantly underperformed compared to its two 3rd Gen bare metal contemporary shapes (BM.Standard3.64 and BM.Optimized3). Because it's similar to previous-generation bare metal shapes, it might be suitable for occasional use, but it's not ideally suited for a heavy production workload.

Although the single-node performance of the 64-core-based BM.Standard3.64 shape and the 128-core-based BM.Standard.E4.128 shape significantly surpassed that of the 36-core-based BM.Optimized3 shape, they might be less cost-effective to use in a typical cloud environment, where pricing is normally closely aligned with the number of cores used. As noted previously, the performance per core is an effective indicator of the total overall cost for the simulation, including both hardware and software cost.

Figures 3d and 3i contain the single-node performance data for the various shapes normalized for performance per core, which is a better indicator of the total cost to conduct a simulation.

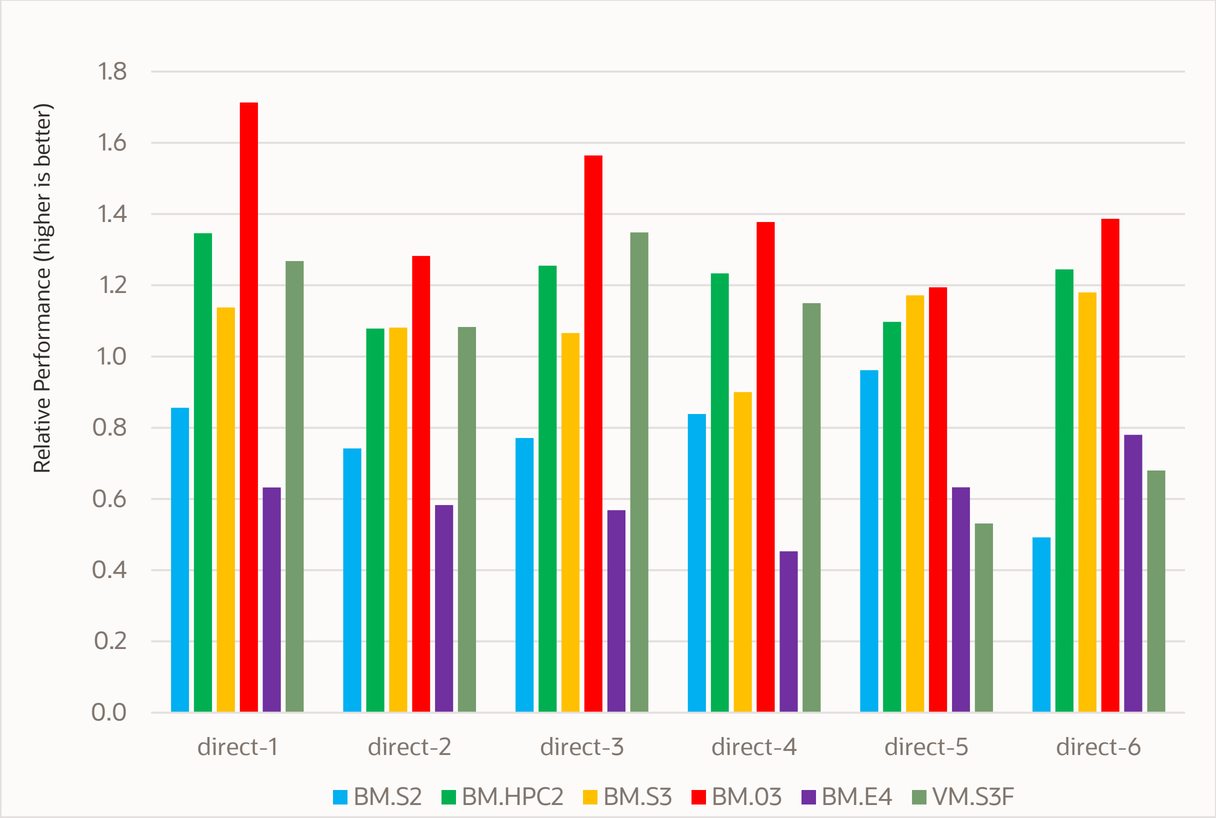

Figure 3d: Ansys Mechanical Single-Node Direct Solver Performance per Core

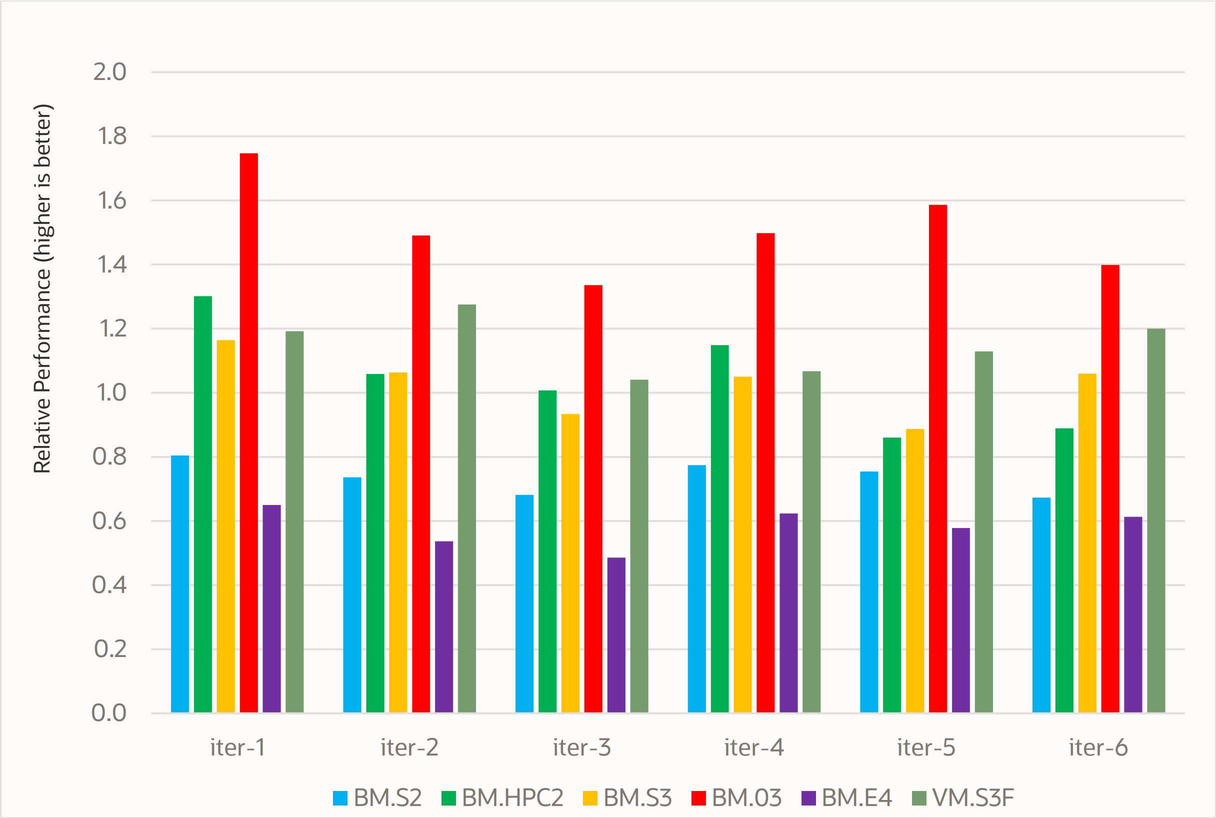

Figure 3i: Ansys Mechanical Single-Node Iterative Solver Performance per Core

The most noticeable difference between Figures 2 and 3 is that the price-performance of the BM.Optimized3 shape is significantly better than that of the BM.Standard3.64 and BM.Standard.E4.128 shapes. When considering both overall performance and performance per core, the BM.Optimized3 appears to be well suited for most Ansys Mechanical workloads.

Multi-Node Results

Many larger Ansys Mechanical simulations benefit from running on a cluster of nodes. Our testing indicated that using the DMP method for multi-node runs was preferred over the hybrid-parallel method. All our multi-node results were obtained using the DMP method, with one domain for each available core. Our multi-node benchmarks were performed on clusters composed of either BM.Standard3.64 or BM.Optimized3 nodes.

As noted in Table 1, the cluster of BM.Standard3.64 nodes used an Ethernet-based interconnect, whereas the cluster of BM.Optimized3 nodes used a state-of-the-art, RDMA-capable 100-GbE NVIDIA ConnectX6 network. Having such a network is a big advantage for pushing the limits of HPC parallel scalability. Because of the network differences, we obtained results only up to two-node jobs for the BM.Standard3.64 cluster, but we obtained results for jobs run across up to eight nodes of the BM.Optimized3 cluster. These results are shown in Figures 4d and 4i.

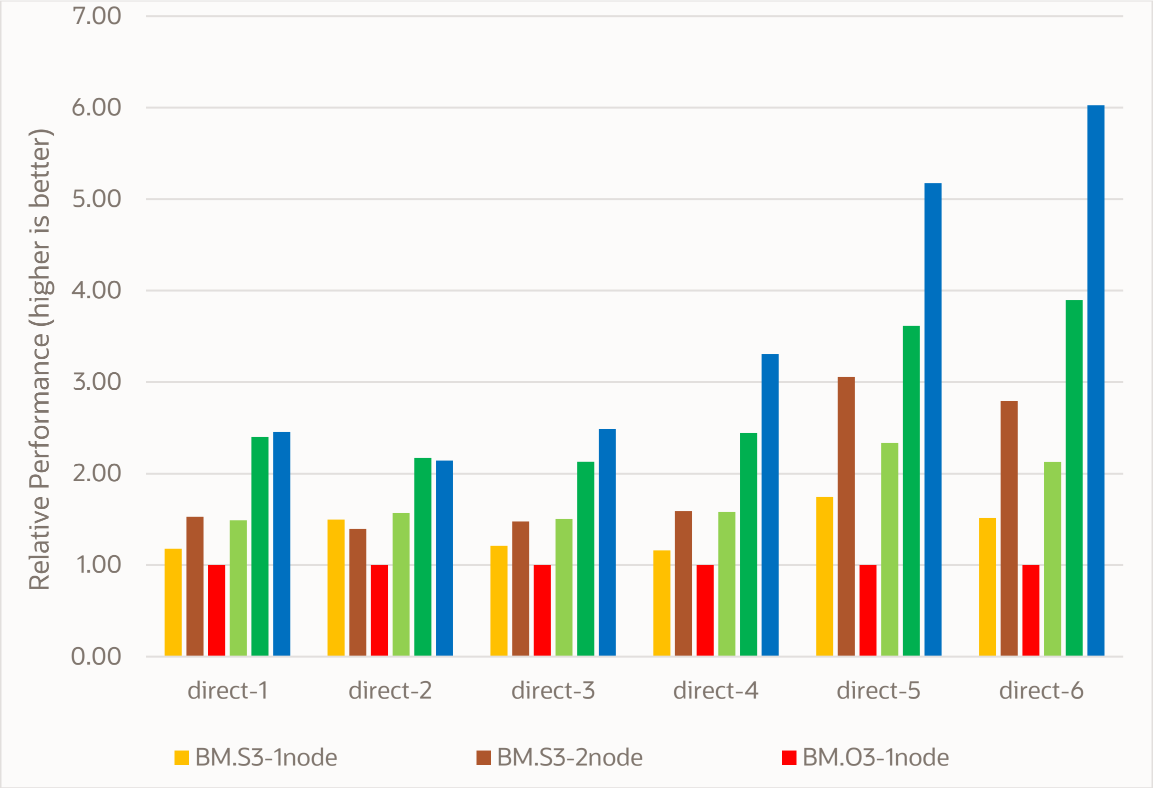

Figure 4d: Ansys Mechanical Multi-Node Direct Solver Benchmarks

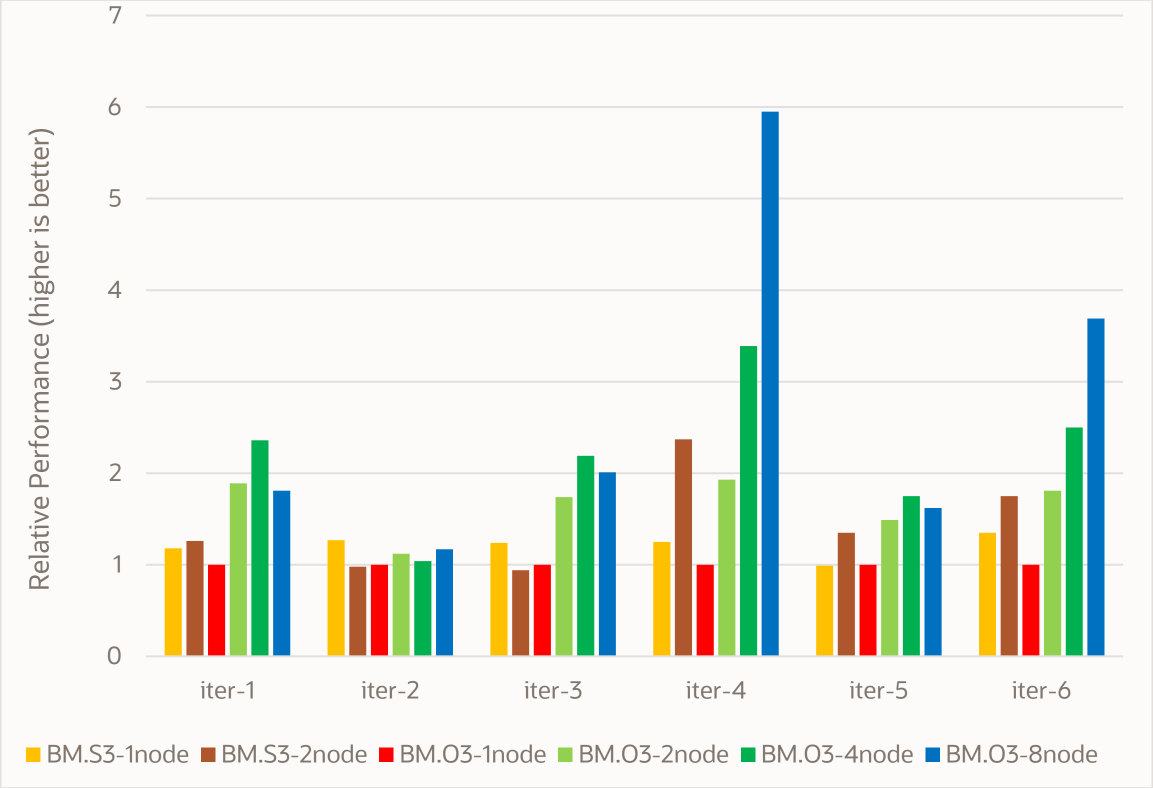

Figure 4i: Ansys Mechanical Multi-Node Iterative Solver Benchmarks

Again, the results are displayed in terms of relative performance per job (higher is better). Here, we chose the reference of 1.0 to be the single-node BM.Optimized3 results. Using this as a reference makes it easier to see the multi-node parallel speedup of the BM.Optimized3 cluster.

Several observations can be made from these results:

- As expected, the parallel performance scalability was generally better for the larger Cluster problems (4,5,6) than the smaller Workstation problems (1,2,3).
- The two-node parallel speedup of the BM.Standard3.64 results was marginally better than the corresponding single node results because of the use of the slower Ethernet-based network.
- Although the BM.Optimized3 nodes have only 36 cores compared to the 64 cores of the BM.Standard3 node, because of the faster underlying cluster network, the performance obtained at two nodes is similar, delivering superior price-performance.
- Apart from the V22iter-5 dataset, all the larger Cluster datasets show good parallel scalability up to eight nodes.

GPU Results

Various Ansys Mechanical simulations can take advantage of GPU acceleration with GPUs such as NVIDIA. Benchmarking tests were carried out with the V22 benchmark datasets on a single node of an OCI BM.GPU4.8 shape, which is equipped with eight NVIDIA A100-40G GPUs with an SXM2 GPU interconnect. Figure 5 shows the single-node performance results obtained with the 12 datasets when using either 8 or 64 CPU cores (c) and 0 or 8 GPUs (g).

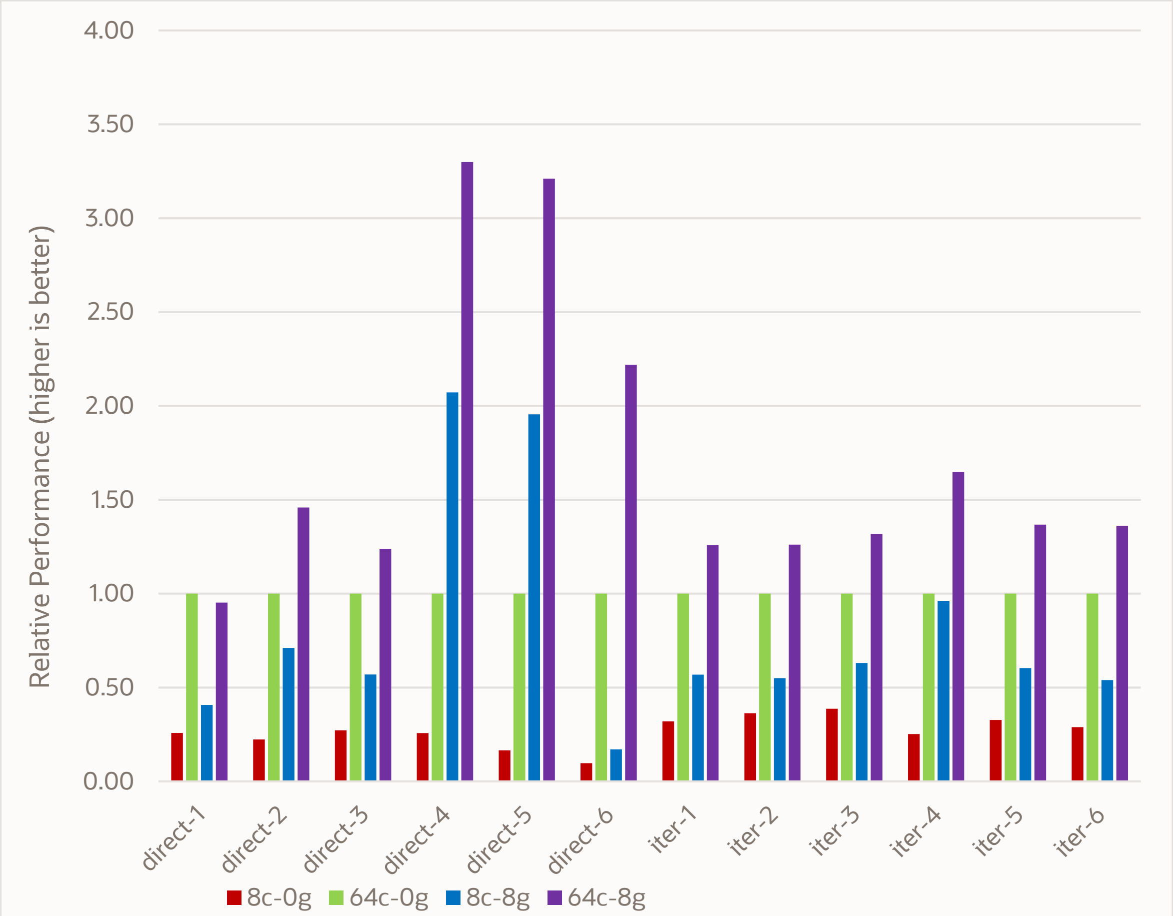

Figure 5: Ansys Mechanical Single-Node GPU Benchmarks, Various Numbers of CPUs and GPUs

The performance reference (1.0) for this chart was the 64-core CPU-only result (64c-0g). For these cases, using all 64 cores created a significant performance improvement over using 8 CPU cores, both with and without GPUs. For all but the direct-4 and direct-5 datasets, the benefit of increasing the CPU core count from 8 to 64 exceeded the benefit of adding 8 GPUs. Except for the direct-1 dataset, the best overall performance was obtained by using all CPUs and GPUs. A sizable improvement in performance with GPUs was obtained when increasing the CPU cores from 8 to 64. This result indicates that for many workloads, a significant portion of the solver workload doesn't appear to benefit from GPU acceleration, and care must be taken to account for the non-GPU related work within the simulation.

Figure 6 shows the performance gains when using varying numbers of GPUs with 8 CPU cores.

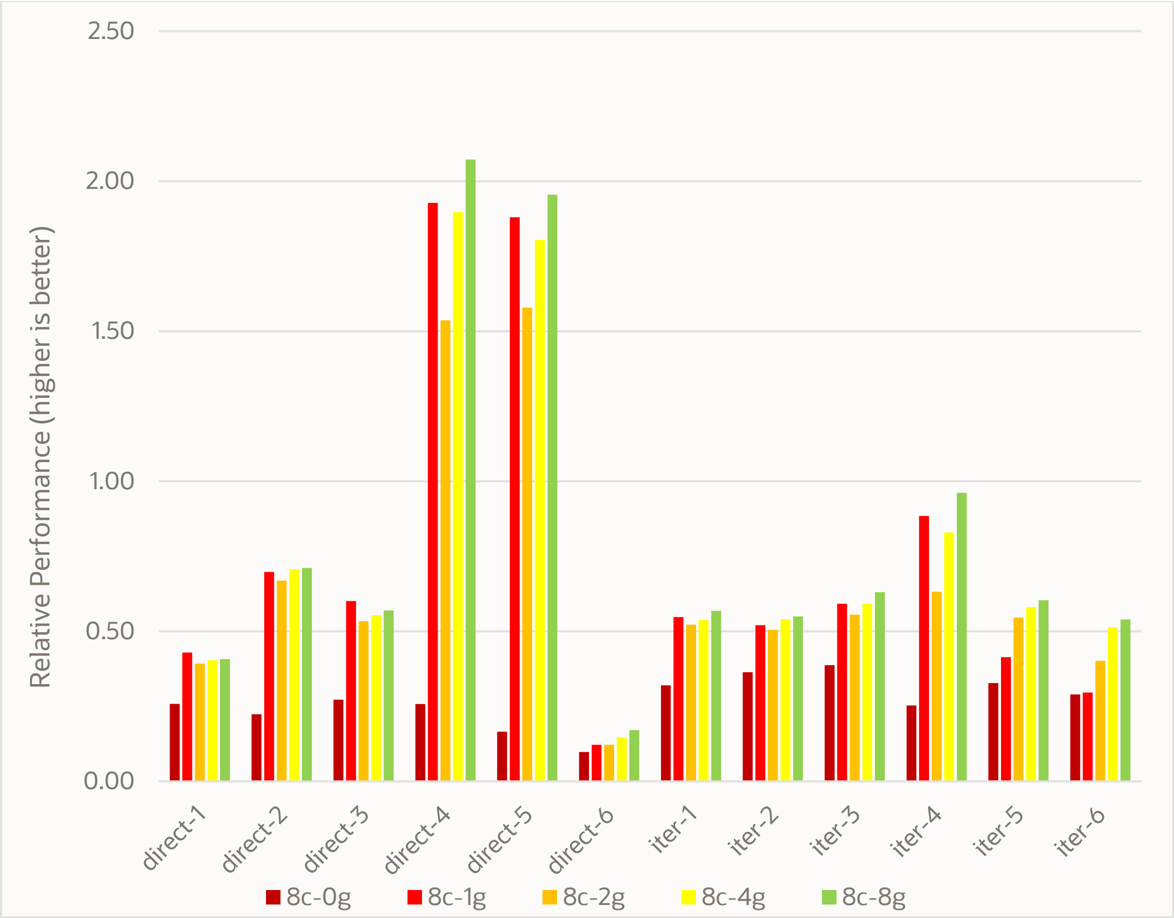

Figure 6: Ansys Mechanical Single-Node GPU Benchmarks, Various Numbers of GPUs with 8 CPUs

The performance reference of 1.0 was taken from the 64c-0g CPU-only results shown in Figure 5. These results indicate that the benefit of using more than one GPU was typically better but not consistently significant. A large number of potential CPU-count/GPU-count variations could be tested; here the attempt was made to simply cover the "edge" conditions.

The cost of using HPC either in the cloud or on-premises is constantly changing and situationally dependent, so determining the actual comparative cost of different shapes isn't practical. HPC users should understand the general trends when looking for specific cost optimization but are forced to do the math themselves for their unique situations. For selected workloads, the use of GPUs might be warranted either for overall time-to-solution or total cost-to-solution. However, the best overall performance is likely obtained with a multi-node simulation using the BM.Optimized3 shape.

I/O Performance Impact

One of the advantages of performing HPC in the cloud is the option to create a variety of clusters, which can be modified for particular workloads as needed.

Where possible, we used direct-attached NVMe (local) storage as the working directory to minimize the I/O impact on performance. For node shapes with no direct-attached local storage available, we tested both network-attached block volume storage and storage using the OCI File Storage service (FSS). In OCI, the performance of the network-attached block volume storage can be adjusted. We tested the two lowest levels (which are the most cost-effective performance levels), low-cost (BV-lc) and balanced (BV-bal).

We carried out benchmarks for the three large direct solver datasets on a cluster of four BM.Optimized3 nodes (144 cores total). The results are shown in Figure 6.

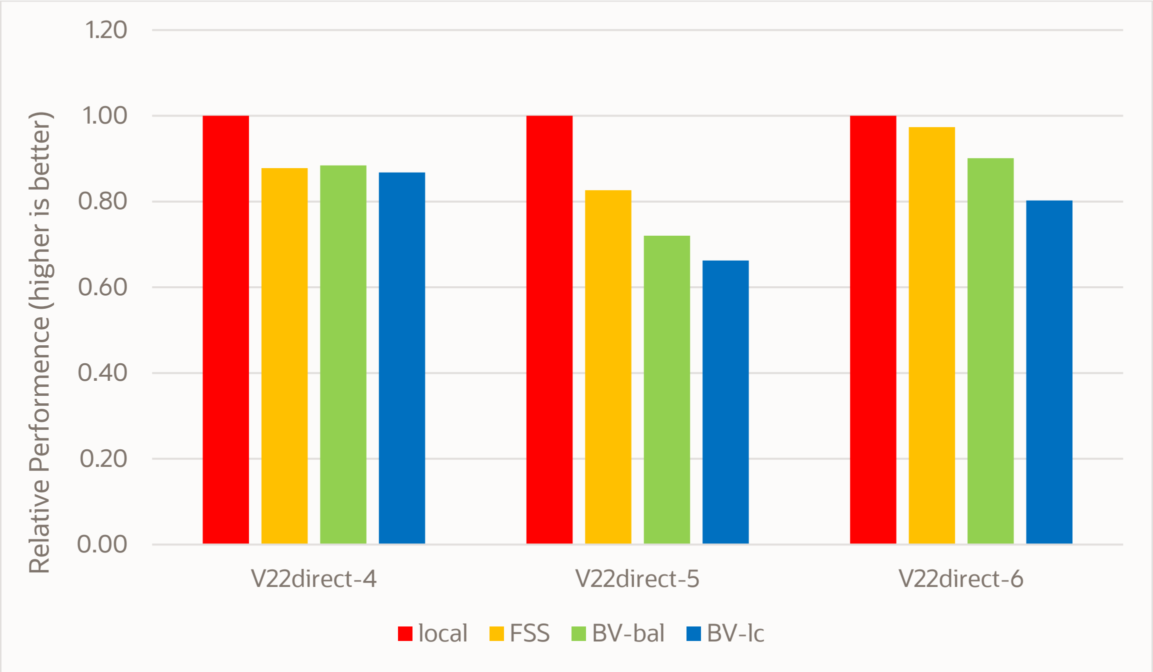

Figure 6: Ansys Mechanical Performance Impact from I/O

Here, the performance using the various NFS-attached storage options is compared to the performance using the direct-attached storage with a reference of 1.0. The amount of overall performance degradation varies but is typically 10% to 20%.

We didn't test the block volume storage option with higher levels of performance because it's likely that customers would chose to perform Ansys Mechanical workloads using the BM.Optimized3 nodes with the direct-attached storage for scratch I/O. However, we could envision a situation in which customers want to perform very large out-of-core simulations for which the local NVMe storage isn't large enough. In such cases, they could mount a block volume of sufficient size to conduct such simulations, and choosing the highest level of I/O performance possible would be cost-effective in terms of limiting overall runtime. We don't think that this situation would be very common, though, and it would be best to test the various options for these extreme cases to maximize performance and minimize cost. Because such cases are unique, there are no simple guidelines to follow.

Conclusion

This paper discusses the various options available for conducting Ansys Mechanical simulations in OCI. The data presented here can help customers determine the best approach to their simulation workloads. For most customers, a cluster system that consists of multiple BM.Optimized3 compute nodes would likely be the optimal overall solution. Such a system compares favorably to the latest generation on-premises offerings, provides the best price-performance for a variety of simulations, and provides the best overall multi-node performance for larger problems.

From a matter of practicality, the best solution might not always be the fastest or the most cost-effective. The Ansys Mechanical workload might be part of a larger overall simulation workload for which it's best to have a cloud infrastructure set up for many different overall requirements and not just one class of Ansys Mechanical simulations. The data in this paper can help to elucidate the trade-offs involved. We encourage Ansys Mechanical users to test their own workloads at OCI using a free 30-day trial.