rescale

# Benchmarking Cloud-HPC Hardware for R&D and Engineering Applications

An IT and HPC Leader's Guide to Optimizing
Their Most Critical Compute-Driven Operations

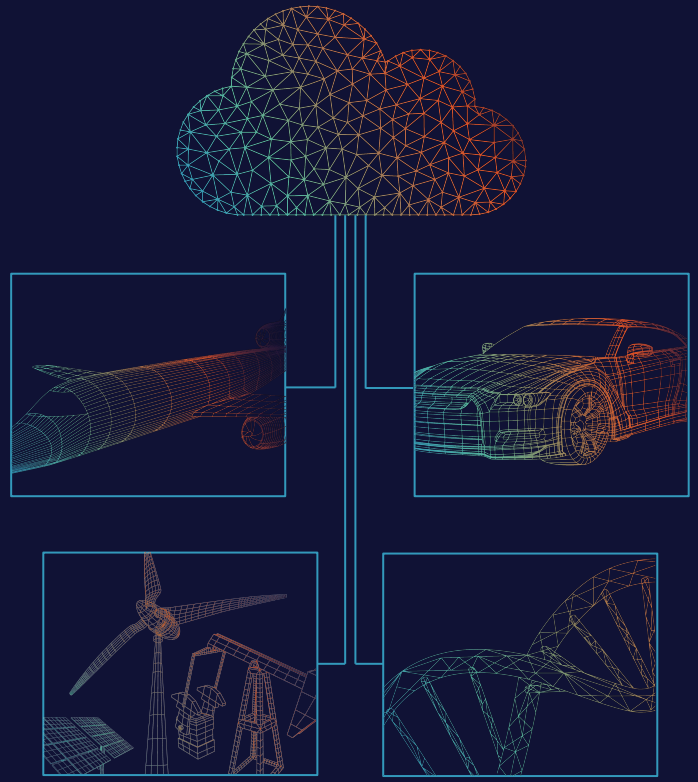

# The case for benchmarking in cloud-HPC modernization

R&D and engineering teams developing new, more advanced products are demanding more high-performance computing (HPC) resources to power their computation-intensive applications. At the same time, there has been an explosion of specialized HPC hardware in the cloud, pushing many organizations to explore new ways to deploy new technologies and techniques for great R&D efficiency and innovation.

While many organizations have a cloud strategy and may have a preferred cloud service provider (CSP), IT and HPC practitioners face a new landscape of challenges and opportunities caused by:

1. Explosion of specialized HPC hardware options in the cloud
2. Broader and more computationally-intensive applications
3. Lack of business visibility and control over HPC efficiency

These conditions make it increasingly complex and important for organizations to optimize their HPC operations to capitalize on their talent and investments in hardware and software.

**In this guide you will learn:** How to use a data-driven framework to capture more value from the latest cloud-HPC technologies. By benchmarking commonly used computer aided engineering (CAE) applications you can identify which hardware and software configurations are best suited for your specific business goals.

## Key considerations for IT & HPC leaders

- ❑ Are you currently using the best cloud architectures for your specific workloads, optimized for cost, speed, or scale?
- ❑ Do you have a strategy to keep up with the latest technologies across all available across multiple providers?
- ❑ What metrics and data sources are you leveraging to maximize the performance and efficiency of your HPC and engineering resources?

**Businesses, on average, can increase workload performance by 30% and reduce overall simulation costs by 20% just by choosing a more optimal hardware available in the cloud.**


**Rescale Customer Impact:** Leading eVTOL manufacturer, Vertical Aerospace, optimized their engineering workloads on Rescale for a **4300% increase in simulation performance and 75% reduction in the cost per job**, enabling faster time-to-market and greater efficiency.

# Cloud options for HPC innovators take off

Are you skeptical about the ROI on multi-cloud, hybrid, and heterogeneous hardware initiatives? Now, more than ever, it pays to pay attention to these trends and determine a modern IT strategy for your business. Today, IT modernization means more than just deploying new technology, it's about having intelligence to know when and how to use them for maximum impact.

While much skepticism of cloud-based HPC is tied to limited options in its early days, things have changed dramatically over the past few years. Cloud has largely transformed the way enterprises procure and manage business-critical infrastructure. And in just the past decade, the public cloud has seen considerable growth in HPC (e.g. simulation and modeling), which has urged public cloud providers to meet the demand for more specialized architectures from hardware providers like Intel, AMD, and Nvidia.

With more businesses bringing HPC into the cloud with the rest of their enterprise applications, IT and HPC leaders are looking for opportunities to take advantage of the pace of innovation in the latest x86, GPU, ARM, RISC-5, and FPGA chip types. Unlike other enterprise applications, HPC applications often benefit from using a variety of multi-cloud architectures, especially when data movement is less important. Because of this HPC practitioners try to stay up-to-date on the latest hardware and software versions to tune and optimize their applications, but the rate of change and complexity often stifle any meaningful business outcome.

Today, engineers and scientists can flexibly select cloud-based hardware for each specific workload they run. However organizations now face a paradox of choice, and making the best decisions is challenging without the intelligence and tools to take action on benchmarking data across the multitude of applications and hardware configurations.

# Application variety and complexity on the rise

In 2020 organizations the majority of organizations using HPC were using 3 or more digital R&D and engineering applications.* Making informed decisions based on their organizations' specific criteria is a key reason why IT, HPC, and engineering managers choose Rescale to manage their growing HPC portfolio complexity. Figure 1 below shows how customers can use this data to determine which CSP has the best hardware across a variety of workloads and CSPs.

**Figure 1** - Average Cloud Provider Performance Across Applications

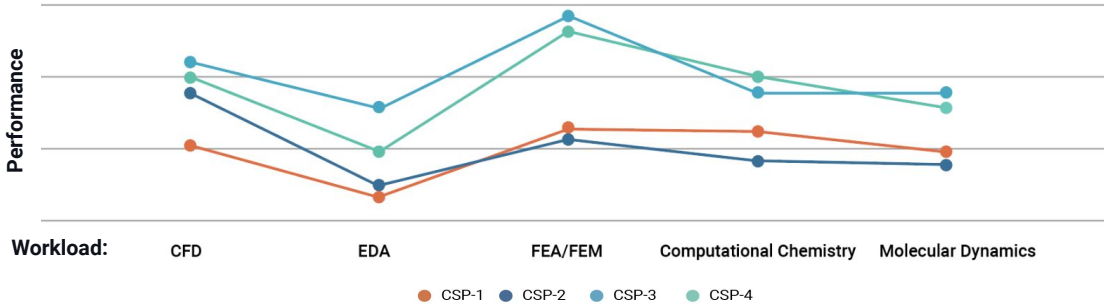


Across industries that commonly utilize HPC technologies, the cost of CAE activities can be significant with licensing and compute expenditures in the millions and increasing each year. On average, organizations can reduce their overall simulation costs by 20% and increase their application performance by 30% by selecting the most optimal hardware available from public cloud providers. As CSP offerings get updated quarterly, continuously optimizing workloads can produce significant results and cost savings.

***Source:** State of Cloud HPC Report, www.bigcompute.org/resources/2021-state-of-cloud-hpc-report

**Rescale Customer Impact:**
Innovative technology company, Arrival, **optimized various R&D applications for up to 400% performance increase across** CFD and FEA to pioneer sustainable automotive manufacturing.

# Unlocking performance intelligence

High-impact HPC decisions are difficult to make without access to up-to-date benchmarking data based on your specific needs and all the possible configurations available. To make this analysis accessible and useful, Rescale compiles performance intelligence to make optimization decisions fast and programmatic.

In Figure 2, we can see how different each offering is, not only across CSPs, but across the hardware types of each provider. In this example we use a popular CAE application, STAR-CCM+, to demonstrate a recent benchmarking of nearly 20 different suitable cloud-HPC architectures. This STAR-CCM+ workload is commonly used in automotive organizations and is a useful application for comparing infrastructure performance.

From this example we can see that certain cloud hardware has better scalability, meaning if organizations can make the right selections they can improve their workload efficiency, time to solution, and overall higher return on investment.

As the variety and scale of HPC workloads and hardware optionality grow, benchmarking in this way becomes more impactful to the business. Because this data is typically not readily available, organizations might work with a CSP, but, as expected, these benchmarks will be CSP-specific and thus limited in scope.

**Figure 2** - Performance Across Various CSP Hardware Options

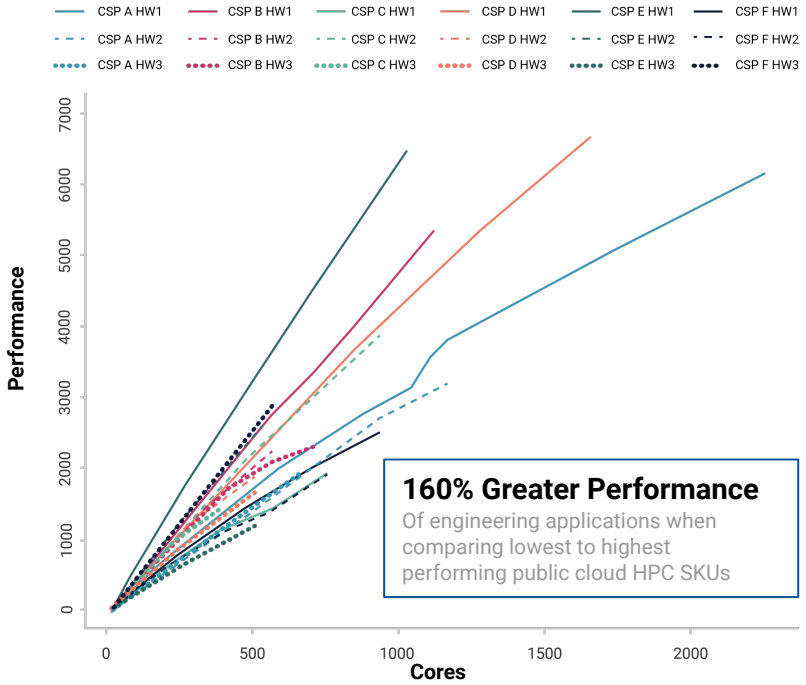


**Workload:** STAR-CCM+ (Lemans - 104 million cells)

**Figure 3** - Cost-Performance Improvement of CSP Hardware Options Over Time

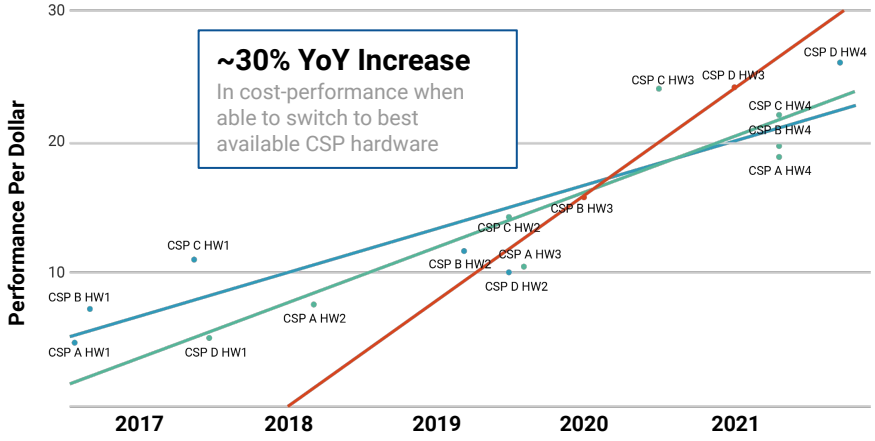


**Workload:** LS-Dyna Car2Car Crash Test

## Rescale Customer Impact:

Nissan optimized R&D workloads on Rescale to achieve **50% Cost Reduction and 18% optimization** of engineering applications and productivity.

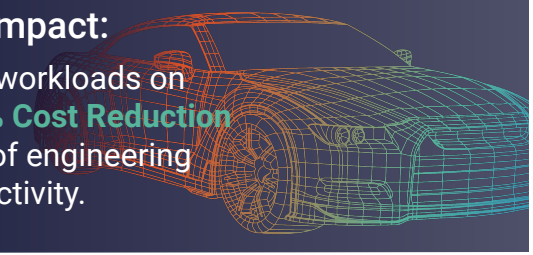

# Realizing Cost-Performance Gains

Benchmarking hardware available at a moment in time (Figure 2) can be valuable for decisions made at that time, but Figure 3 (shown left) illustrates the importance of continuous benchmarking. Here we use another popular automotive application, LS-Dyna, but to compare cost-performance over a period of time (2017-2021).

Among the 3 cloud providers plotted, we can see varying cost efficiencies for each hardware with an overall trend of increased cost-efficiency. This steady improvement of cloud infrastructure is great for end customers; Even with a single cloud posture organizations can expect to get more compute performance at lower costs.

But when organizations can take advantage of multiple CSP offerings, they can realize significant gains when they compete. We can note how the pace of new hardware releases increases and the CSP leader in cost-performance changes frequently.

We expect these year-on-year hardware improvements and variation of best cost-performance to continue, but as seen in this chart, organizations must have both multi-cloud intelligence and flexibility to realize these gains when the hardware becomes available.

# A framework for better business decisions

The typical HPC practice incurs a variety of costs, from compute infrastructure to sizeable software licensing and even unforeseen costs from failed jobs and overruns. As R&D and engineering teams grow their usage of compute-intensive HPC applications, the modern IT & HPC teams need to start thinking about managing all these costs holistically.

Rescale orchestrates hundreds of HPC software and hardware and works directly with customers to implement and optimize complex workflows. Because of this, we have built the platform to take into account the different sources of cost and inefficiency across the stack - then we optimize them. We call this concept “Full-Stack Optimization” and it’s made possible by continuously benchmarking multiple data points across various application and compute architectures.

The result of all this data is a framework (Figure 4), comprised of 3 main indices that track performance, maturity, and value. This framework can be used by HPC managers and practitioners to easily assess the best-fit hardware for their needs, whether their goal is to minimize cost, maximize performance, or ensure successful scaling of massive projects.

This model can inform business decisions and set policies to maximize project spend and ensure new users are given the most optimal resources. Organizations are using Rescale to automate these decisions to avoid cost overruns or job failures caused by selecting the suboptimal hardware, and administrators can also set failover rules if capacity for certain architectures are temporarily unavailable.



**Figure 4** - Rescale’s Data-Driven Framework for Intelligent HPC Optimization

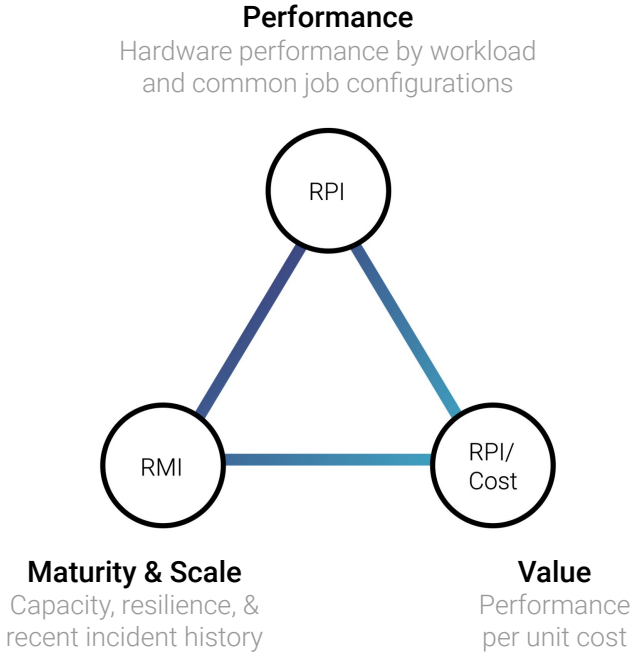

**Figure 5** - Holistic Benchmarking of Hardware Options for a Given Workload

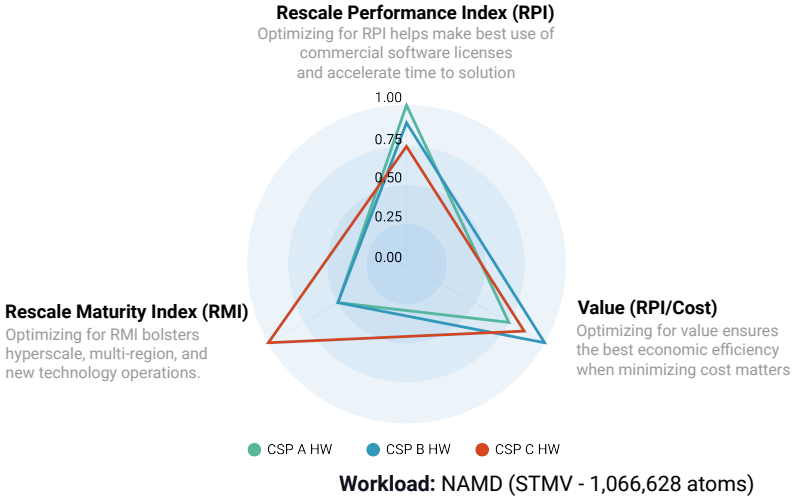


**Workload:** NAMD (STMV - 1,066,628 atoms)

## Rescale Customer Impact:

Biotechnology company, AzothBio, **accelerated their AI-driven drug candidate discovery process by over 2x** leading to faster decision making for which candidates to develop.

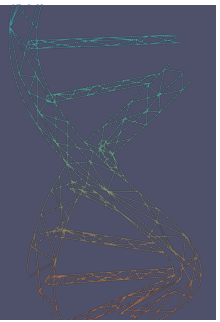

# Putting benchmark intelligence to work

Figure 5 shows how to put all this benchmarking data to use for better business outcomes. This example workload, molecular dynamics application commonly seen in life science applications, we compare 3 different hardware options with different performance, maturity, and value profiles.

For many organizations running commercial software, an important consideration is maximizing the utilization of available licenses and therefore you would want to select hardware that scores highly in performance. Alternatively, for open-source applications without licensing you may want to optimize for value if cost savings is more important than time to solution.

So, in this particular example with NAMD which is available without paid licensing, the blue triangle in Figure 5 would indicate that Hardware 2 has the best value and still relatively strong performance. As organizations begin to scale up their compute requirements or test newly released hardware the maturity score becomes more important to ensure that a given hardware option has sufficient capacity and stability to successfully complete jobs.

The Rescale platform integrates this intelligence to enable companies to deploy their resources quickly and efficiently. As new technologies become available, engineers and scientists can take advantage of them sooner while IT and HPC managers can minimize risk to the business.


www.rescale.com

# Exploring the new possibilities of *Intelligent Computing* on Rescale

With Rescale's benchmarking intelligence for data-driven decision making, you can implement new strategies from automation to optimization and beyond!

**Get in touch with us to discover how your HPC practice could:**

- ❑ Give users on-demand access to the software and hardware they need without queuing delays
- ❑ Take advantage of the best cloud infrastructure as soon as it becomes available, regardless of the cloud provider
- ❑ Ensure that each user, project, or application was automatically configured to run on the most optimal architecture
- ❑ Continuously update and optimize your software and hardware stack so that your engineers and scientists can continue to innovate
- ❑ Adapt to meet end-users resource demands as they fluctuate without having to worry about underutilization
- ❑ Provide complete visibility of compute efficiency and alerts when better resources become available

rescale

Intelligent Computing for Digital R&D

## About Rescale

Rescale helps organizations accelerate science and engineering breakthroughs by eliminating computing complexity. From supersonic jets to personalized medicine, industry leaders bring new product innovations to market with unprecedented agility, speed, and efficiency with the Rescale Platform - an intelligent full-stack automation HPC solution for digital R&D in the cloud.

IT leaders use Rescale to deliver high-performance computing-as-a-service to their organization by harnessing the power of automation on a hybrid cloud control plane with security, architecture, and financial controls. Learn more at www.rescale.com or follow the links below.

Get in touch with our team of experts

Benchmark Your Workloads with our Team

View additional resources