




# Dispelling the Top 5 Myths of Cloud HPC

Misconceptions and outdated concerns IT and HPC practitioners, managers, and business leaders often believe about implementing cloud high performance computing strategies



# Cloud HPC is going mainstream. Are these outdated misconceptions holding you back?

Beliefs are opinions that the longer they are held, the more cherished they tend to become. The challenge is, as technology advances, a belief can easily morph into a bias. IT decisions that were once correct when first made years or even months ago, as time passes, become assumptions based on myth rather than factual assertions.

For example, a widely held belief in IT today is that high-performance computing (HPC) platforms can't be accessed as a cloud service due to cost-performance shortcomings. The general assumption is that running cloud infrastructure is either not specialized enough or less efficient than utilizing on-premises resources.

In truth, however, cloud computing infrastructure and platforms to access compute resources continue to rapidly evolve and much of the innovation and growth in the past decade has been in cloud-based technologies. Factors influencing HPC applications range from the cost of cloud instances and the level of security provided to assumptions about the number of HPC applications any IT team might need to support have all changed in the last few months. It's clear the time to revisit those HPC decisions has arrived.

With that premise in mind, there are at least five HPC assumptions that are now more myth than actual fact.

- ▶ 1  
HPC is a Niche Application
- ▶ 2  
Cloud is Not Cost-Effective for HPC
- ▶ 3  
Cloud Can't Meet HPC Requirements
- ▶ 4  
Deploying HPC in the Cloud is Difficult
- ▶ 5  
Cloud is Not Secure Enough for HPC

## Myth #1: HPC is a niche application

Not too long ago HPC was considered an IT niche that primarily addressed, for example, use cases involving life sciences applications created and managed by scientists and researchers that required access to supercomputers. However, with increased demand for simulation and modeling along with more reliance on machine and deep learning algorithms that are now almost routinely infused into applications, the need for HPC platforms to create AI models has greatly expanded. The worldwide HPC market already stands at \$46.4 billion, according to a new [report](#) published by Intersect360 Research. That estimate, however, does not include AI infrastructure purchased by cloud service providers.

Any way it is measured, an IT market valued in the tens of billions is hardly a niche industry segment. Almost every vertical industry employs HPC platforms across a broad spectrum of engineering, research and development and product optimization use cases. Cloud computing platforms simply makes HPC more accessible to all types of development initiatives for a few dollars per hour. Clearly, the number of HPC applications that will be running in the cloud has already exploded as more researchers experiment with applications that they previously would not have been able to financially consider.

In fact, the easier it is for researchers to access computational resources the more likely it is that an organization will achieve its business goals. A recent 2022 State of Computational Engineering Report finds that 85% of organizations that either consistently achieve goals on budget (22%) or often achieve goals on budget (63%) enjoy easy access to computational resources.

HPC is also converging with AI and ML technologies which are being integrated into nearly every industry and stage of product development. As organizations invest more heavily in AI-enabled engineering, they will need increased access to the latest GPU and CPU architectures to meet the demand of their engineers and scientists. Being able to operate complex HPC practices efficiently is now a critical lever for new innovation and competitive advantage.

## HPC Market Set to Reach \$50 Billion by 2027

Source: [MarketsandMarkets](#)

From startups to large enterprises, increased investments in HPC resources is tied to increased competitive advantages such as improved R&D productivity and accelerated time to market.

New cloud cost models significantly reduce workload costs up to 72%

Cloud HPC economics now compete with or beat on-premises economics especially when fully-burdened total cost of ownership

## Myth #2: Cloud is not cost-effective for HPC

One of the biggest myths, when it comes to HPC in the cloud, is cost. HPC applications have long run times, so the assumption is that on-demand cloud resources are more expensive over time when compared to the fixed costs of on-premises infrastructure.

However, it is now possible to take advantage of new cost models, or service levels, such as reserved cloud instances/VMs which offer as much as 72% less than typical on-demand pricing. At the same time, the overall performance of cloud platforms continues to improve as well. On average, the cost-performance ratio offered by cloud service providers improves roughly 30% per quarter. That ratio will continue to improve as new classes of processors designed for specific high-performance workload types become available at an increasing rate.

For example, let's estimate it costs \$.03/core-hour to run an HPC application in an on-premises environment. This estimate doesn't include the cost of the overhead of managing on-premises environments so, in reality, it's much more likely that the fully loaded cost is closer to \$.05/core-hour to run an HPC application in an on-premises environment. A recent 2021 [State of Cloud Computing Report](#) finds that on a pure price-performance basis on-premises, home-grown HPC environments are less and less cost-effective in the wake of recent advances made by cloud service providers who offer some VMs running at \$.02-.03/core-hour. Better still, a cloud service enables organizations to instantly take advantage of burst capacity when required without having to install additional infrastructure. In many cases, that level of flexibility alone is priceless.

Finally, accounting rules have changed. Organizations are no longer required to treat these cloud platforms as an operational expense. Organizations can take advantage of capital budgets to consume reserved instances of cloud infrastructure or employ a mix of both financing models as they best see fit. And because cloud hardware selections get better over time, versus five years of the same on-prem hardware, many IT teams are shifting to cloud to maximize the performance and efficiency of other major investments like their talent and costly software licensing.



## Myth #3: The cloud can't meet HPC requirements

Early on the only types of processors available via the cloud were variants of x86 platforms. Today there is now a wide range of classes of processors capable of driving HPC applications. In addition to x86 processors, there are now GPUs from NVIDIA and floating-point gate arrays from Xilinx that are readily accessible to anyone from anywhere. It's clear that every type of processor that can be found in an on-premises IT platform is now available in the cloud.

The truth is the cloud enables organizations to accelerate the pace of HPC innovation in a way that is simply not possible in on-premises IT environments. HPC applications benefit from being able to tap multiple types of processors to optimize performance. Engineers and scientists can flexibly select cloud-based hardware for each specific workload they run using analytics tools that identify what workloads should be deployed based on real benchmark data.

In fact, a recent Benchmarking Public Cloud-HPC Hardware [report](#) finds 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.

Not surprisingly, the number of HPC applications being deployed in the cloud has been increasing at double-digit rates. A recent report published by Intersect360 finds more than 70% of organizations make use of the cloud, with more than half (53%) consistently opting to deploy HPC applications in the cloud. Similarly, a [report](#) published by Forrester finds 99% of respondents indicated their organizations are open to shifting more HPC workloads to the cloud. The respondents highlighted several factors that will enable this shift, including an expansion of HPC services in cloud environments, the ability to access workload automation management solutions across their organizations' toolset, the availability of turnkey hybrid and multi-cloud platforms, and greater flexibility to choose hardware and tools that fit their organizations' specific use.

**99% of  
decision  
makers are  
open to  
shifting HPC  
to the Cloud**

Source: [Forrester](#)

Top benefits of cloud cited include:  
"Availability of more HPC Services",  
"Workload automation",  
and "Turnkey Hybrid and Multicloud".

## 95% reduction in deployment time reported with cloud HPC

Source: NOV

Simplified resource, application, and budget management enables global teams to quickly deploy enterprise-grade computing capabilities.

## Myth #4: Deploying HPC in the cloud is difficult

Historically, many HPC applications have been challenging to deploy because they depended on custom-built and maintained infrastructure. Each application required a specific environment and dedicated tooling that needed to be planned, procured, and then learned by end users and administrators.

Modern HPC applications today make use of HPC-as-a-service (HPCaaS) and automation frameworks to simplify application and infrastructure deployments. IT and engineering tasks that once required weeks, even months, can now be completed in a few days using pre-configured compute environments, utilities, and workflows. Cloud enablement solutions like Rescale work with customers to automate connections between hybrid and multicloud resources as well as complex workflows. With thousands of automations already deployed for various customers, other organizations can realize immediate benefit with Rescale's turnkey solution.

Previously, IT organizations would have needed to hire a team of specialists to integrate the disparate components that make up an HPC platform, for example: schedulers, license servers, specialized chip architectures, file system, operating system, compilers, and libraries.

Of course, the most expensive element of an IT solution is talent, which, in the HPC industry, is becoming increasingly difficult to find. HPCaaS platforms reduce the need for significant personnel overhead by pre-packaging components in a way that makes it easier for end users of HPC applications to programmatically invoke HPC services. There's still a need for some level of HPC expertise but now that there are platforms that abstract away the underlying complexity of the HPC environment there are also going to be more developers than ever building HPC applications.

In effect, the cloud and cloud enablement solutions make it possible to shift more time to realizing the benefits of an HPC system instead of setting up and managing it.

## Myth #5: Cloud is not secure enough for HPC

The last HPC myth in the cloud that needs to be dispelled pertains to security and compliance. It has been long assumed organizations cannot run sensitive workloads on public cloud infrastructure because of compliance and security standards. Today cloud service providers routinely achieve compliance with a wide range of standards and requirements, including:

- **FedRAMP:** The Federal Risk and Authorization Management Program (FedRAMP) is a government-wide program that provides a standardized approach to security assessment, authorization, and continuous monitoring for cloud products and services. FedRAMP empowers agencies to use modern cloud technologies, with emphasis on security and protection of federal information as part of an effort to accelerate the adoption of secure cloud platforms.
- **ITAR:** The International Traffic in Arms Regulations (ITAR) are a set of U.S. government regulations administered by the U.S. Department of State's Directorate of Defense Trade Controls (DDTC). These regulations control the export of defense articles and defense services, including the placing of restrictions on the handling of software and technical data controlled on the United States Munitions List (USML).
- **SOC-2:** A System and Organization Controls audit conducted by a third party assesses the information security and privacy compliance with standards established by the American Institute of Certified Public Accountants (AICPA). These reports verify that an organization has effective security, processing integrity, privacy, and availability of information.
- **ISO 27001:** ISO/IEC 27001:2013 (ISO 27001) is an international standard that helps organizations manage the security of their information assets. It provides a management framework for implementing an information security management system to ensure the confidentiality, integrity, and availability of all corporate data. It is the only internationally recognized certifiable information security standard.

Many protected and sensitive workloads are already running in the cloud. In fact, it may be more likely to experience a breach in an on-premises IT environment than a cloud platform with leading security and compliance protocols. The number of government agencies around the world that make use of cloud platforms to process and secure data using private cloud instances running on shared infrastructure has dramatically increased in the last three years. It is important to evaluate that both the infrastructure and the services utilized and built on top also have the necessary security from end-to-end.

**67% of (US) federal agencies are shifting mission-critical workloads to the cloud**

Source: [FedScope](#)

Many IT leaders now believe a FedRAMP-authorized cloud solution can provide greater security over on-prem deployments



85% of organizations will embrace cloud-first by 2025

Source: [Gartner](#)

Digital strategies will require the use of cloud-native architectures and technologies.

## It's time to modernize HPC with the best in cloud

It is always constantly evolving. Where to run HPC applications is not the only IT decision that organizations have revisited as cloud computing environments mature. The first cloud HPC deployments essentially amounted to little more than deploying the same type of enterprise applications IT teams ran in an on-premises IT environment in the cloud. The new and current era of cloud computing enables new use cases on top of making it more efficient to build, deploy, share, update, and manage applications more efficiently than ever.

HPC in the cloud benefits organizations by eliminating on-premises infrastructure constraints in a way that makes it easier to pay only for the capacity used no matter how much is required. They enable organizations to innovate without constraint, improve flexibility and deliver faster results.

Naturally, there may still be a handful of HPC applications that don't lend themselves to the cloud but as time marches on, the number of those applications continues to grow smaller with each passing day.

### Triggers That Often Initiate Cloud HPC Transformations

If you are not sure if the timing or conditions are right to explore a cloud HPC, the following are common triggers that prompt organizations to kick off a cloud transformation:

- On-premises data center or workstations are due for replacement
- Growth in capacity demands or expanding engineering teams
- Widening workload diversity across custom, commercial, and open source apps
- Cloud-First directives for business resilience e.g. business continuity, collaboration
- Competitive pressure to innovate with new products, processes, and tools





## High Performance Computing Built for the Cloud



Digital  
Engineering



Workload  
Optimization



Intelligent  
Automation



Security &  
Compliance

### About Rescale

Rescale is high performance computing built for the cloud to empower engineers while giving IT security and control. From supersonic jets to personalized medicine, industry leaders are bringing new product innovations to market with unprecedented speed and efficiency with Rescale, a cloud platform delivering intelligent full-stack automation and performance optimization. IT leaders use Rescale to deliver HPC-as-a-Service with a secure control plane to deliver any application, on any architecture, at any scale on their cloud of choice.

Learn more about high performance computing built for the cloud:



Get in touch  
with our team of  
experts



Start your free  
trial today



View additional  
resources