



EBOOK

HPC Buyers Guide

Building a modern HPC environment



Major trends impacting HPC decisions

CIOs have become the cornerstone of innovation for product development. They are under pressure to deliver a growing number of resources while keeping their team lean and agile. Failure to meet the ever-increasing demand of R&D teams ultimately risks losing market share. Fixed capital investments risk underutilization in today's uncertainty and rapidly changing business environment.

This tension between building specialized infrastructure, whether that be on-premises or a DIY cloud computing approach, for ever-growing demand and maintaining agility begs these core questions:

How to provide the ideal resources for every situation?

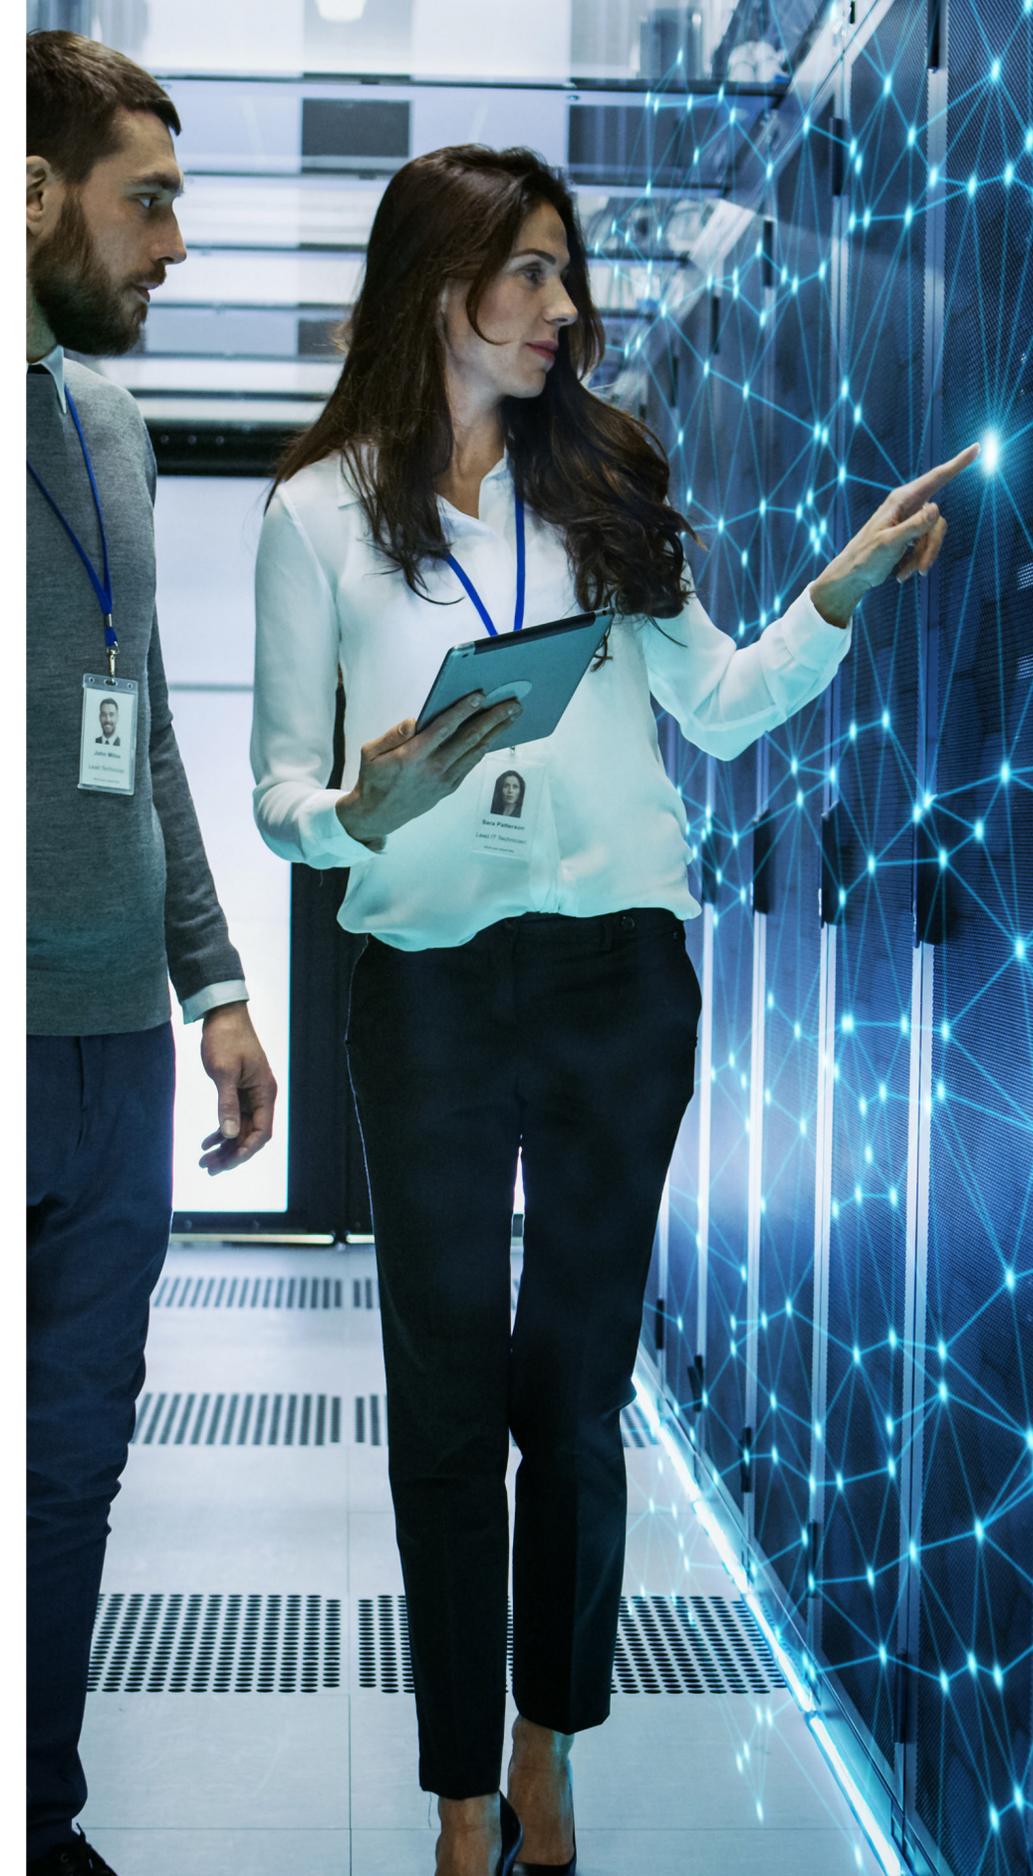
How to optimize overhead needed to maintain HPC infrastructure?

How should it be measured - Utilization? User uptime?

Historically, enterprises have purchased fixed infrastructure for specialized computing, commonly deployed as on-premises, co-located, or shared supercomputing centers.

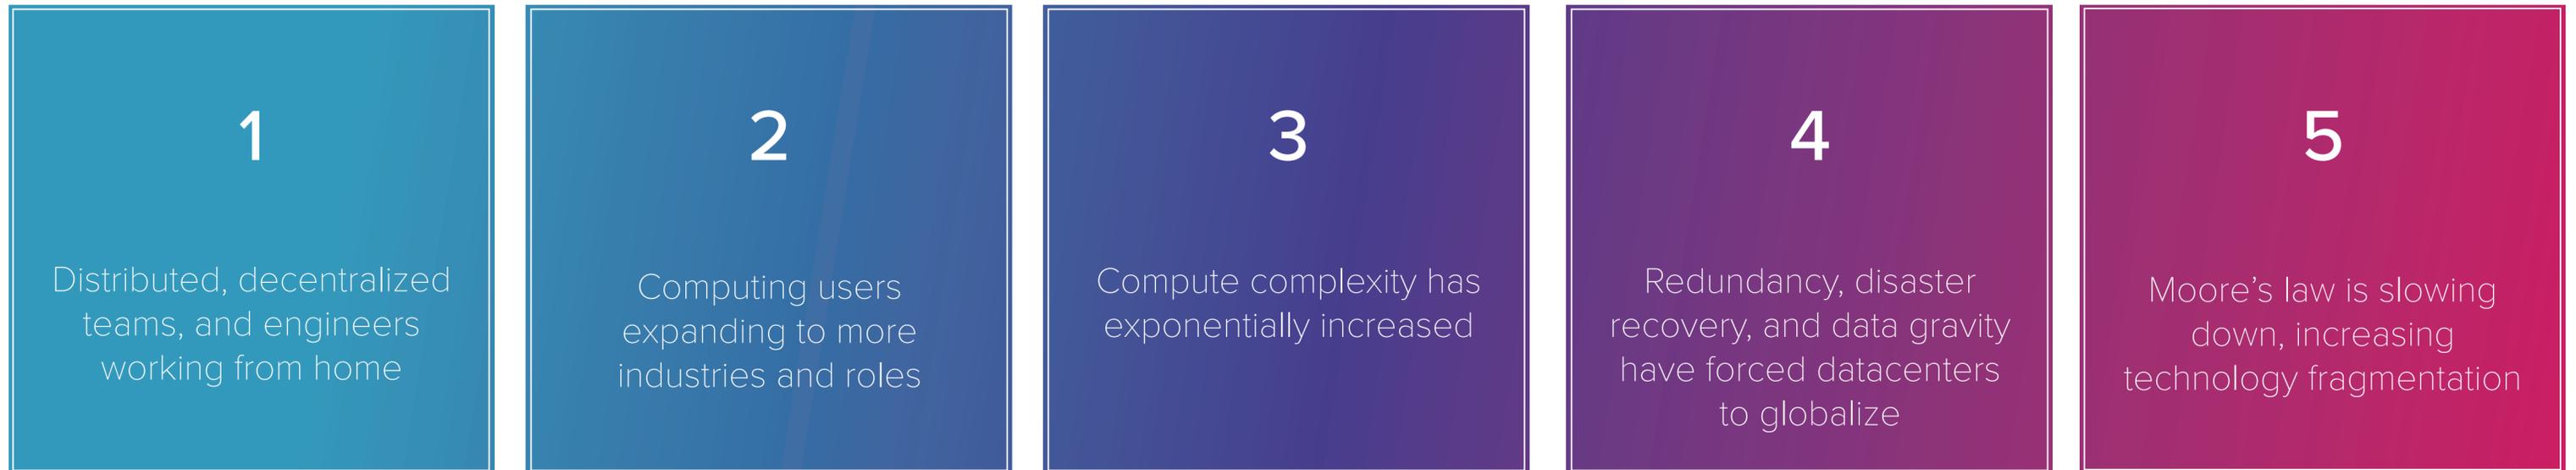
Fixed capacity investments are measured by utilization, a metric well-suited for stable demand. However, workloads that power innovation processes tend to be highly volatile and can drive the demand far beyond capacity. R&D service uptime can be as low as 20%, despite nearly 100% utilization by the fixed infrastructure.

This lack of capacity and elasticity has driven many organizations to migrate to the cloud. However, the lack of skill-sets required to build a customized solution, and the inability to predict business requirements years in advance can lead to significant overspending.



Major trends impacting HPC decisions

We have seen the rapid growth of five key trends since 2010. However, 2020 accelerated this growth beyond what anyone could have predicted.



The above industry trends are independent of the particular solution approach and ultimately force a CIO and CTO in an arena where they must adapt to new solutions. Delay or failure to do so results in an unfortunate parasitic effect on innovation: unsustainable resource limitation.

This guide will ensure your organization understands the pros and cons of the HPC options available and introduce the Rescale managed service solution. Our solution works with both on-premise and cloud clusters, as well as a hybrid solution, and provides a way to get your engineers up and running immediately on a platform specifically optimized for their unique requirements.

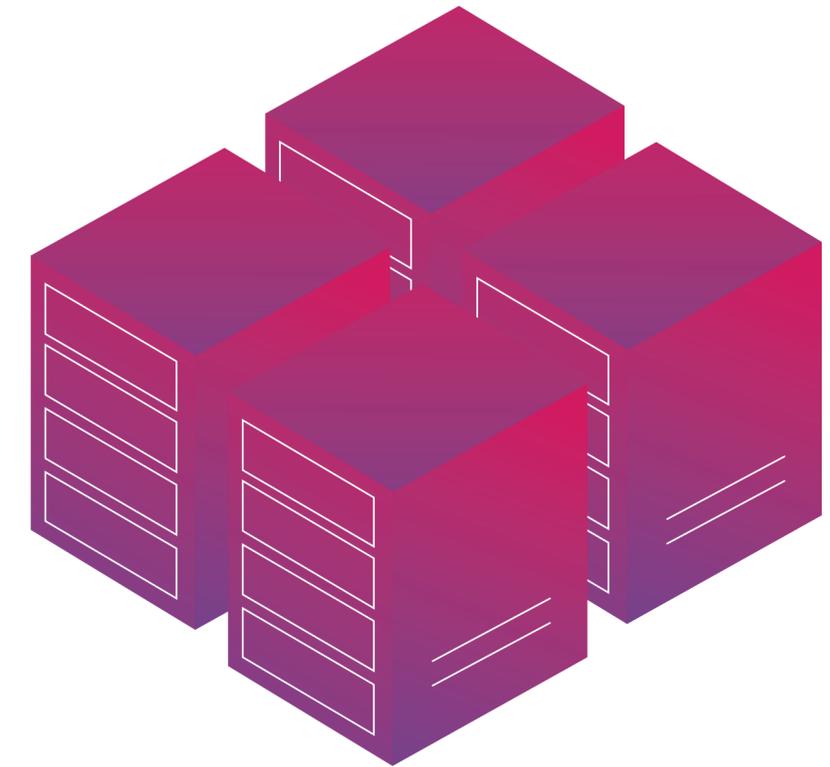
The legacy option: on-premises clusters

The most common reasons CIOs cite in favor of on-premises clusters are security and compliance. In industries such as banking, government, and healthcare, with their strict security and compliance protocols, those CIOs might be right when it comes to certain data. Especially if you already have the expertise in-house to keep on top of those protocols and stay current with security trends and threats.

However, security and compliance are becoming less and less a concern with cloud computing. Cloud platform and software providers are making security a priority and continue to ramp up that investment to the extent that most companies could not compete.

“CIOs need to ensure that their security teams are not holding back cloud initiatives with unsubstantiated cloud security worries,” says Jay Heiser, Vice President Analyst, Gartner. “Exaggerated fears can result in lost opportunity and inappropriate spending.”

Still, some enterprises feel more comfortable keeping certain business functions on-premises until cloud providers can reasonably guarantee security as well as a stamp of approval from relevant regulatory agencies, all while understanding the costs associated with that decision.



Public cloud IaaS workloads will experience 60% fewer security incidents than traditional data centers by 2020. And at least 95% of all security failures in the cloud will be caused by the customer.

- Gartner



Hidden costs of on-premises

Sections

Static Capacity

Unpredictable Resource
Planning

High Capital Expenditure

Hidden costs of on-premises: static capacity

The on-premises HPC market has gradually matured, and with that maturation has come increased specialization of hardware. While specialization offers an initial appeal, it also brings with it a staggering burden to operators. The diversity of machines required leads to more overhead, buried maintenance costs, stability risks, and security exposures.

Additionally, the variety of hardware means more expertise and less utilization. In the long run, teams either suffer tremendous burdens of the cost associated with the latest and greatest type-specific hardware or opt for a middling solution with reduced performance.

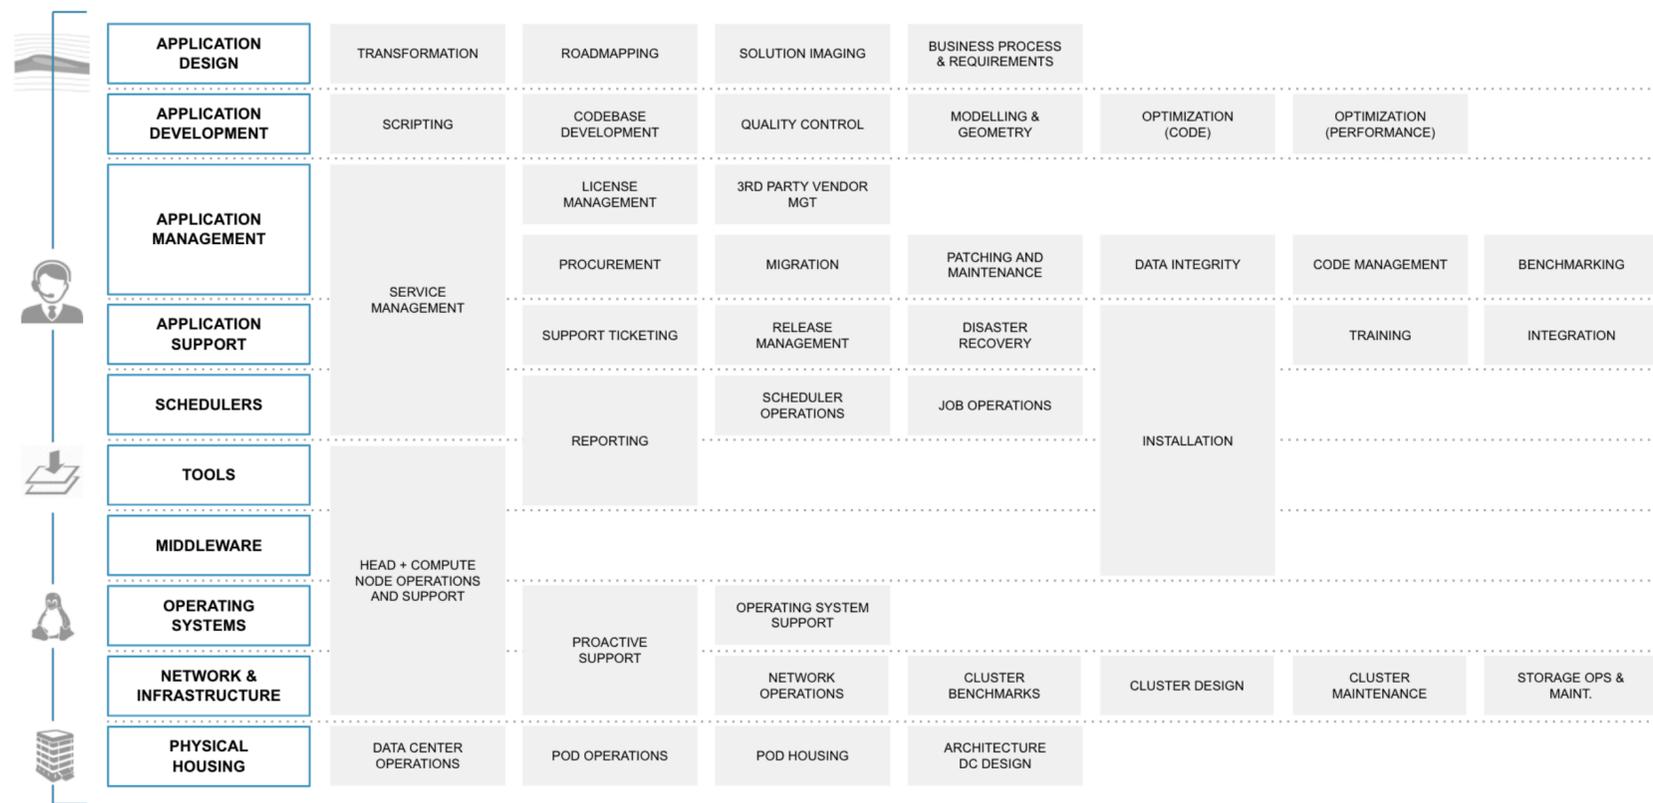
Modern HPC resources are highly optimized, challenging to scale appropriately, and their performance is sensitive to many factors. Unfortunately, product innovation places unpredictable, non-steady demands on the systems: these workloads tend to be highly volatile or “peaky.”

In some cases, changing a chipset family can reduce the compute time required by more than 50%. When organizations build out an internal cluster to meet the “average” or “typical” workloads, they create a situation where peak demand and specialized projects cannot be accommodated with state of the art performance.

This lack of elasticity and the ability to dynamically allocate resources based on workloads have a significant negative impact on a company’s bottom line. Engineers, a high-value resource, sit idle as they await their turn in the queue. This idle time can lead to delays in a vital product launch and missed customer deadlines. Static capacity can also lead to less obvious but certainly not insignificant issues such as lower fidelity simulations and reduced engineering efficiency.

72.8%
of organizations reported pent-up demand was either delayed or cancelled due to lack of resources. Hyperion Research

Hidden costs of on-premises: unpredictable resource planning



Think about how your company has changed over the last five years. Even just this past year. Not only is technology changing rapidly, but so too have many businesses. During the pandemic, many manufacturers found themselves completely retooling their production lines to produce much-needed PPE, and they did so at a pace no one imagined.

Even during “normal” times, hardware and software application resource decisions have to be made years in advance. It would be incredibly cost-prohibitive to design a system based on each departments’ needs, resulting in a one size fits all approach being taken. That approach often leaves engineering lacking the compute power they require, significantly reducing their productivity.

New programs that could make a substantial positive impact on your business bottom line may have to be put on hold for a few years until you can upgrade your hardware to accommodate them. Static infrastructure with long refresh cycles will not accommodate changes in business goals. When engineers do not have access to the latest hardware, simulation costs and software costs rise.

On top of all this, IT organizations often experience a skills gap and lack the in-house expertise required to maintain such a stack.

Hidden costs of on-premises: high capital expenditure

Corporate initiatives always require IT organizations to find ways to cut capital expenditure, while internal teams are simultaneously demanding more capacity. The ultimate goal is maximum elasticity with minimal long-term investments. That elasticity is something that is lacking in on-premise HPC.

Instead, internal teams face long procurement cycles and hardware lock-ins. Due to the high capital investment and long deployment times for new hardware, most organizations are forced into a 3-5 year depreciation cycle. That long cycle means even an optimal cluster might be running on machines that are five years out of date.

A common mitigation strategy employed by on-premises operators is to make partial equipment upgrades of computing equipment. While this increases the average system performance more continuously, it brings the following drawbacks:

1

A constant state of procurement, often taking a substantial percentage of the service life (example: 8 months from RFP to operational, when the equipment's service life can be as short as two years).

2

On-premises technology fragmentation drives compatibility challenges, increased maintenance and support load, and additional system elements to map workloads to hardware.

3

High disposal and end-of-life cycle costs

Cloud migration - DIY cloud

Because of the hidden costs of an on-premises solution, many companies are choosing a DIY Cloud infrastructure. Some organizations are building their infrastructure from scratch, or they are using the cloud to enhance their legacy enterprise systems.

Many of these companies have particular needs that out-of-the-box solutions do not provide. However, to build such a custom infrastructure, organizations must have the expertise in-house, or they must rely on outside consultants to design the solution for them.

Much like an on-premise solution, the DIY solution requires that you can accurately forecast business requirements to some extent but allows you to change certain aspects on the fly when needed. Still, it isn't easy to accurately estimate each department's needs, and therefore, it is not uncommon for companies to overspend.

It's essential to keep in mind that while you don't have the capital expense of hardware and software, it does have costs associated with the time it takes to plan, build, and implement.

80% of IaaS users will overshoot their budgets mostly because they lack necessary process controls to deal with costs in the cloud.
- **Gartner**

Performance and optimization costs

Consider the performance cost of getting locked into one Cloud Service Provider (CSP), limiting your options for optimization. Should you have the option of using multiple providers or simply wish to migrate to a different provider in the future, you must learn the unique ways of setting up virtual machines for each different CSP.

As mentioned previously, while cloud providers invest in security expertise and always design for security enhancements, security is a shared responsibility. Application security will still rest firmly in the hands of the customer requiring in-house skill sets.

Optimizing your system requires continual testing and tweaking to ensure you are getting the biggest bang for your buck, so to speak.

We see the majority of our customers using a hybrid model. Functions, such as those that require strict security protocols, compliance, or even highly specialized hardware, remain on-premise. In contrast, applications that require speed, such as engineering simulations, are moved to the cloud.

If you opt for a DIY cloud solution, be sure to account for the time you will spend planning, executing,

and optimizing the system, as well as training and supporting your users.

1. Requirements gathering

Acquire a set of requirements from each user group and all workflows. A basic framework should be determined in advance to avoid the need to interpret each user group's process.

2. Determine cloud provider

Establish HPC on the public cloud you will be using. For most basic needs, this is not a terribly onerous process. However, many organizations have complex DIY solutions that incorporate several different platforms, each having unique setup processes and requirements.

3. HPC Configuration

Once setup is complete, you must configure your HPC clusters and workflows. While cluster management software helps you configure multiple nodes and resources, each platform has its requirements to manage the cluster and schedule jobs.

4. Application installation

Next, you must install and compile HPC apps and versions. If you are running several applications, this can be an arduous process.

5. UI Configuration

Just as crucial as the cluster's setup and configuration is the setup for the HPC user interface to ensure an effective user experience. Your goal is to simplify complex software configurations and provide consistency across architectures.

6. Training and support

Finally, you must implement a training and support program. This program will be a crucial factor in successful implementation. While we have listed it last as an activity, it should be considered in the early planning stages.

The modern option: platform-managed HPC

Rescale is a managed service provider built with engineering in mind. Our independent software vendor partners provide simulation software to engineers, and we make these programs available on our platform.

Why do you need Rescale's platform-managed HPC? Optimization through automation.

Sections

Optimization of In-house
IT Team

Application Optimization

Resource Optimization

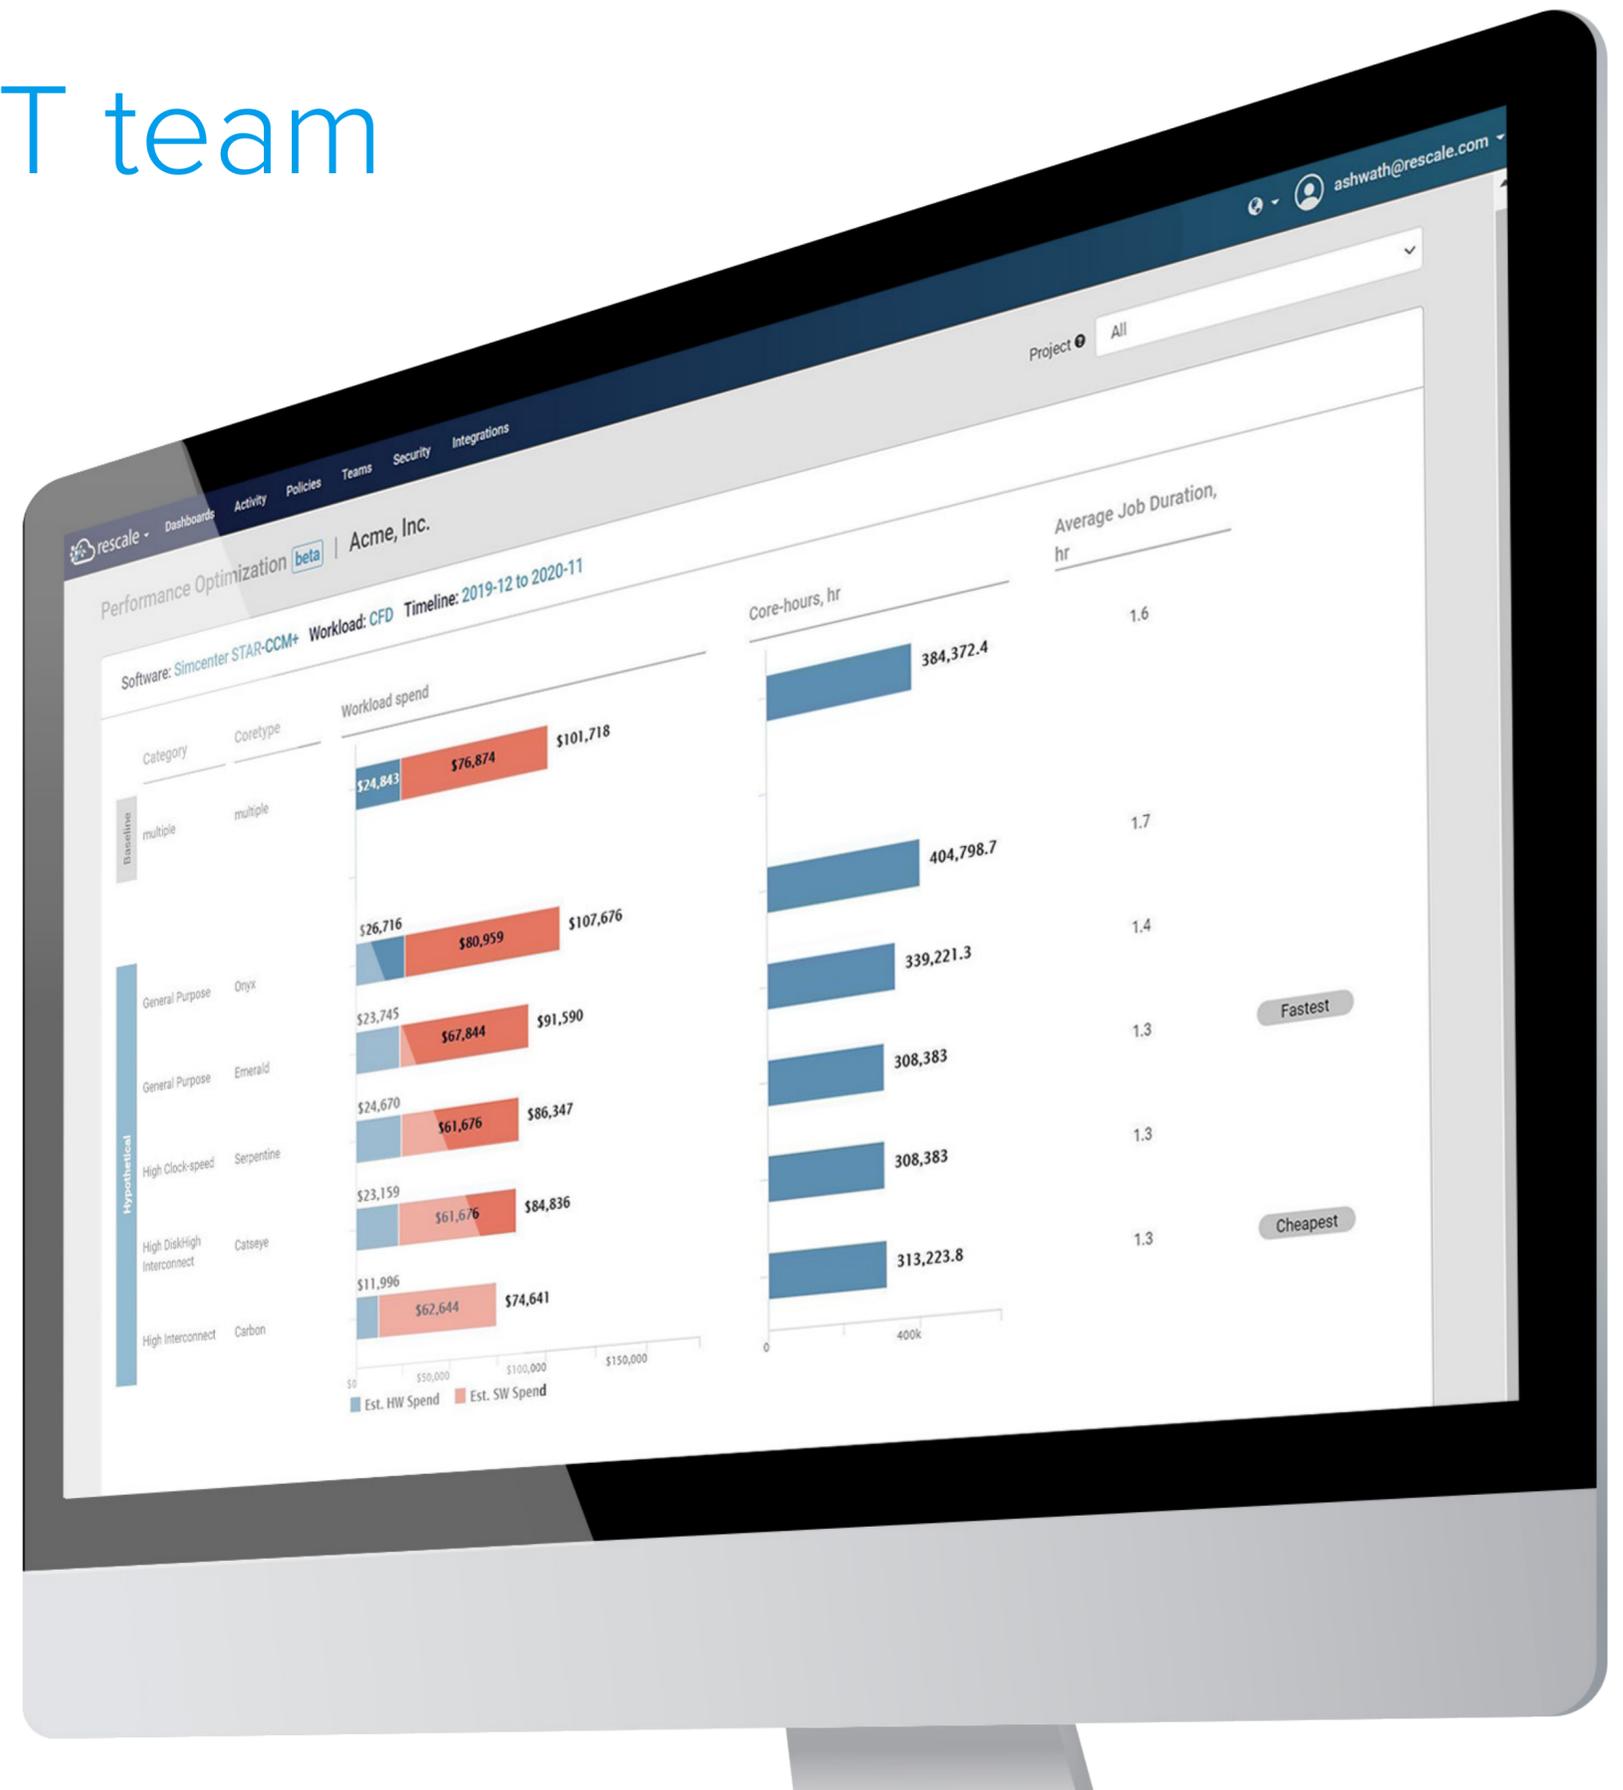
Cost Optimization

Summary Comparison
Table

Optimization of in-house IT team

Most on-premises deployments don't include a user-friendly interface for working on engineering problems. With a managed HPC platform, the centralization of resources can dramatically reduce the IT burden in onboarding new users, documenting processes, and maintaining command line job submissions. By combining all the software and hardware pairings in one location, we can improve simulation start and finish time. This results in higher efficiency from engineering teams.

Managing resources across offices, countries, and continents is a tricky job. Platform infrastructure models can easily aggregate simulation usage metadata and make it accessible from a central location where IT leaders gain full visibility and control. IT teams can easily manage hardware and software, allocate and enforce budgets, set user limits, track usage, and generate reports across globally-distributed teams, departments, and offices.





Application optimization

Rescale supports over 370 enterprise HPC simulation applications. Our simple installation and tuning process ensures each application is optimized for your underlying infrastructure. Not only is installation a quick process, but you will also have your engineers up and running immediately.

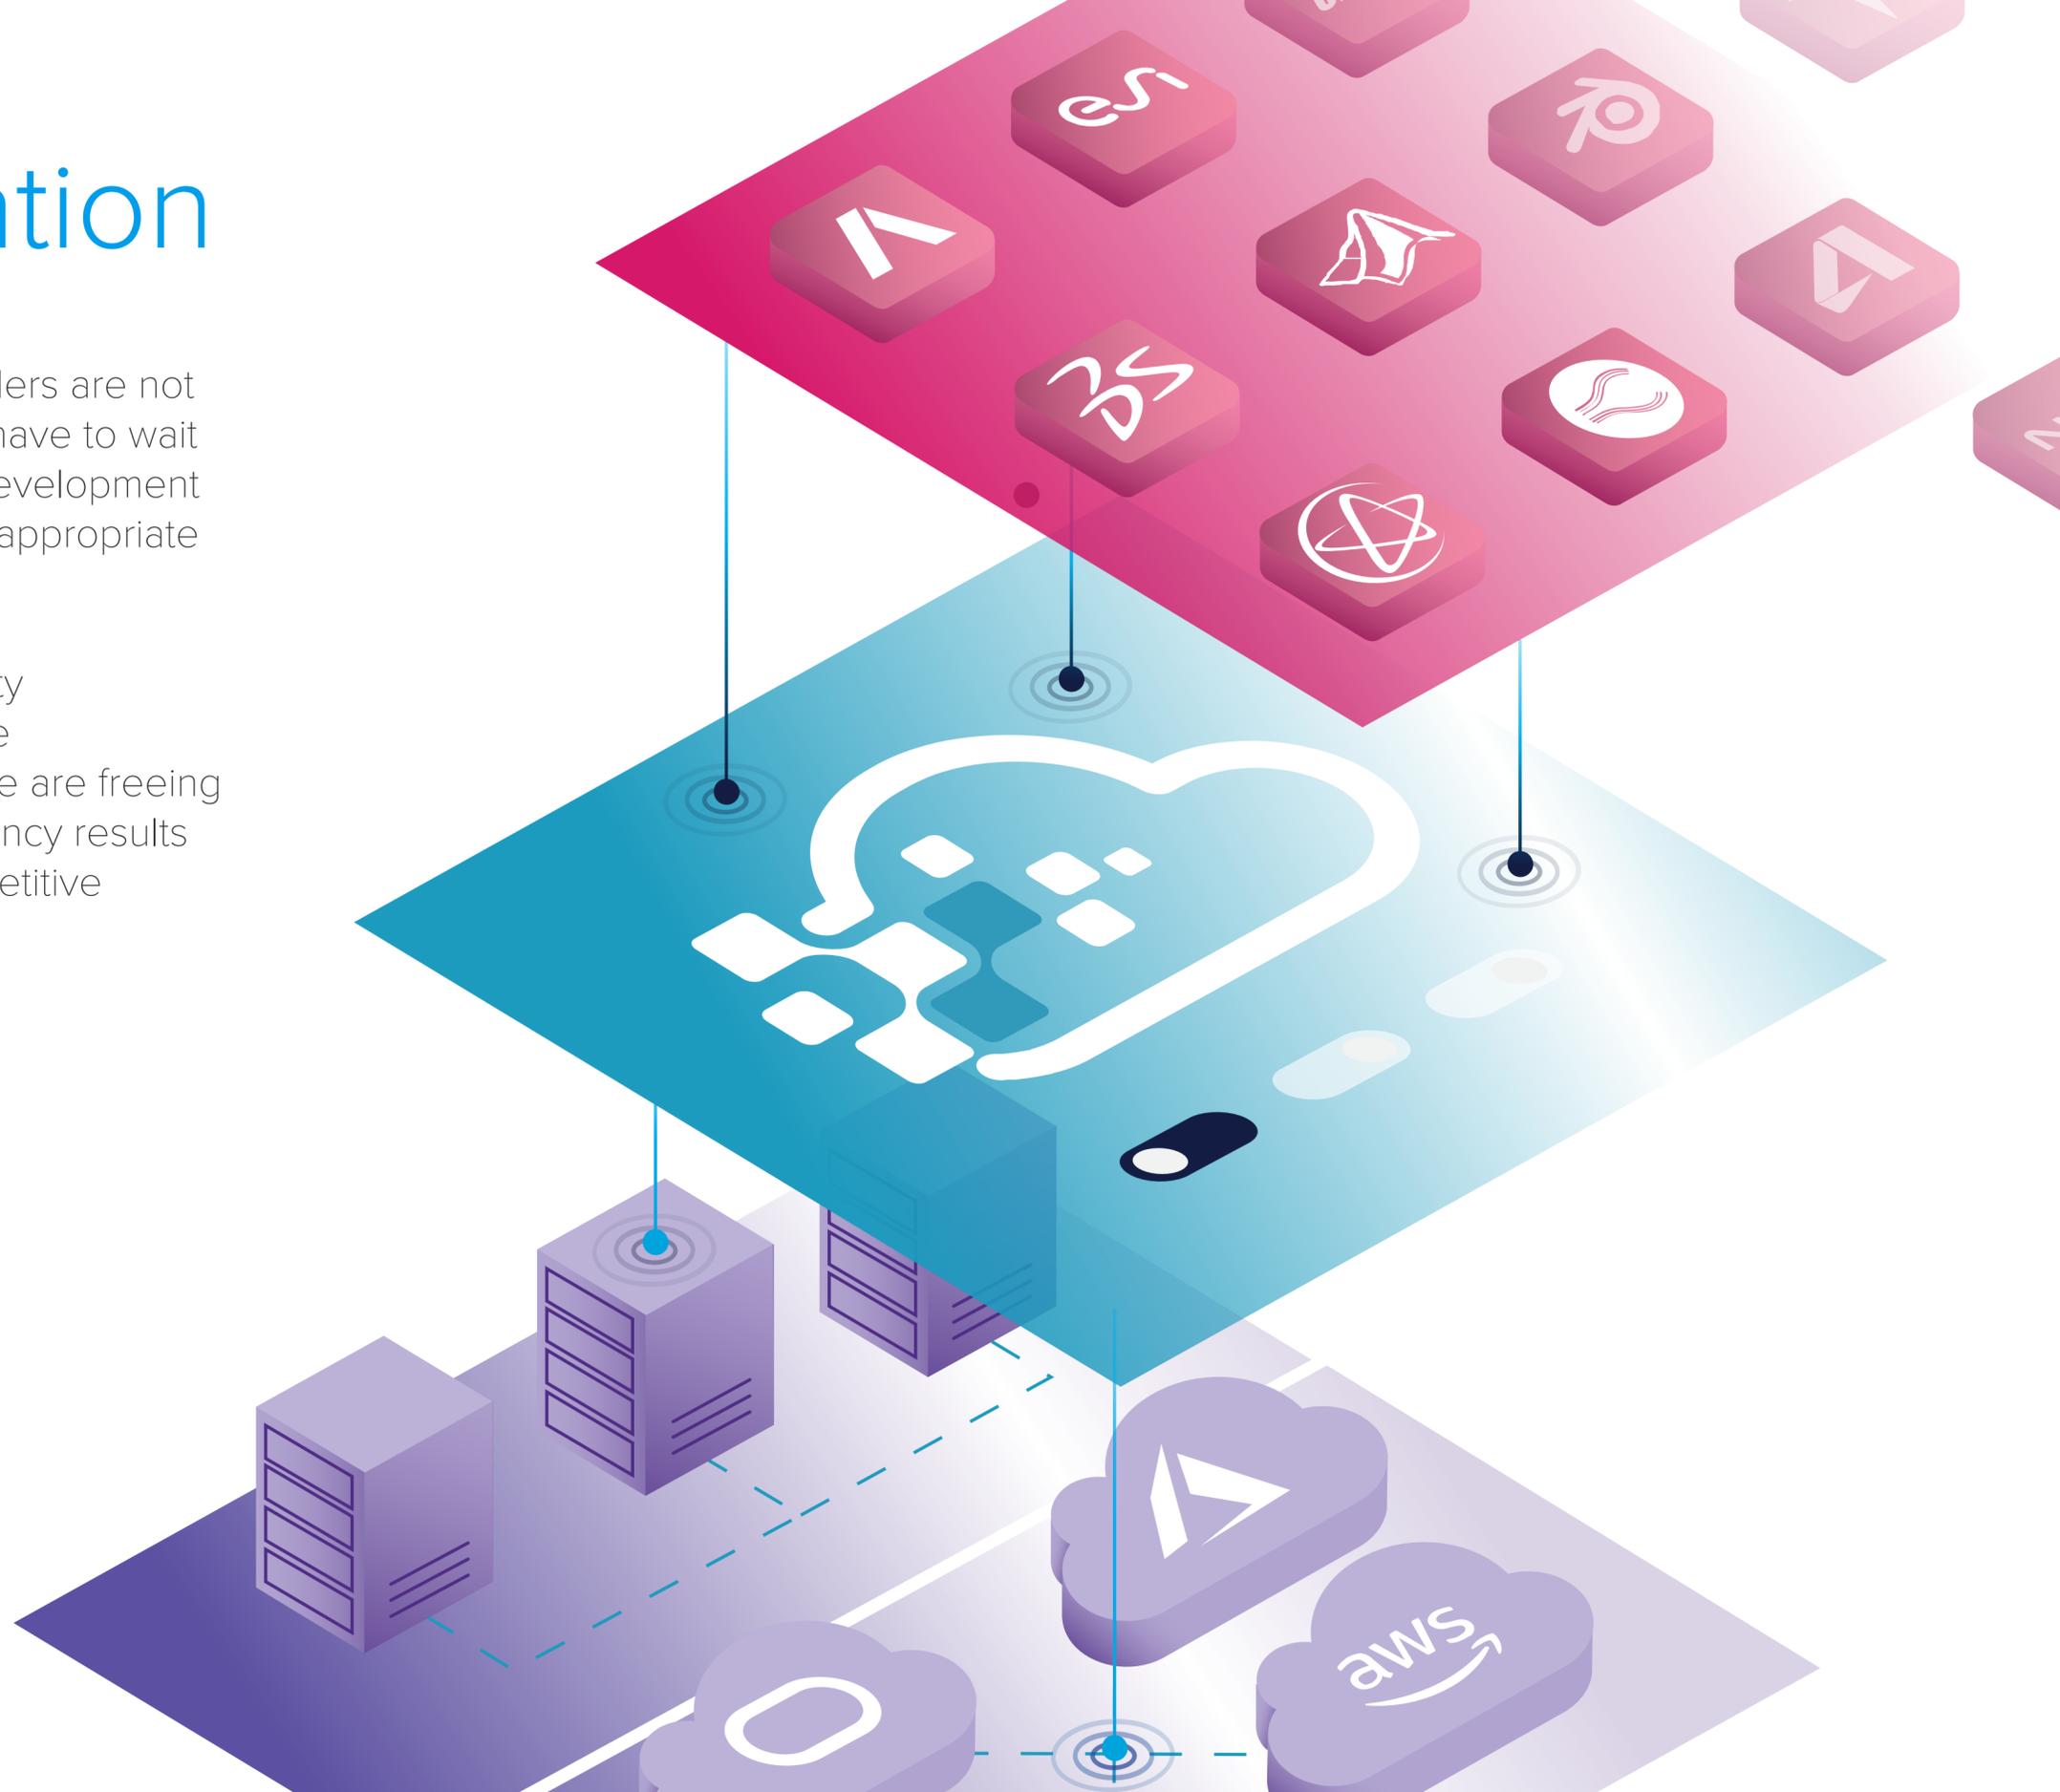
Each application running on Rescale provides the same, familiar software experience offered by your Independent Software Vendor (ISV).

Rescale supports over 370 enterprise HPC simulation applications

Resource optimization

There are many cloud resource providers, but all providers are not equal. With our HPC platform, engineers will no longer have to wait in the queue to run their job, which takes up valuable development time. Instead, Rescale sends the simulation to the most appropriate service optimized to run the data the fastest.

This access to specialized hardware with elastic capacity accelerates the product design process. R&D teams are empowered to complete more iterations in less time. We are freeing them up to explore the design space fully. All this efficiency results in more competitive products in the market and a competitive advantage for the organization.



Cost optimization

Rescale's modern HPC experience and intelligence compute platform eliminates job silos. Instead of looking at one particular job, we are looking holistically at your entire organization. We understand that teams use resources differently, so we provide budget controls to ensure meaningful, high-impact jobs receive higher priority than low-impact jobs. That allocation gives you peace of mind that you are optimizing the cost of HPC.

While a DIY cloud solution can take months or years to get just right, a Rescale deployment is as simple as signing up for an account. We can have your engineers up and running in a matter of hours, not months or years.

Summary comparison table

There are five main Enterprise Alternatives available for IT organizations to serve HPC needs to its users:

Given the comparison between cloud HPC and on-premise deployments, organizations are searching for the best ways to build their HPC practices. These five strategies are the most common:

On-premises The historical standard found in most company's datacenters

Private Cloud An extension to on-premises for overflow computing with a high degree of control.

Hybrid Cloud Bursting An on-premises system configured with a scheduler software to reroute jobs to one or more cloud providers. Includes public Cloud/ In-house development and integration of all HPC workloads around a single cloud provider.

DIY Cloud A build-your-own, custom approach to cloud migration. The DIY configuration may include several different cloud providers and requires expertise and time to develop a solid migration and testing strategy.

Managed Multi-cloud Platform A purpose-built platform to leverage all the scalability, flexibility of architectures, administrative controls and customization while abstracting away maintenance, security and compliance.

	On-Premises	Cloud DIY	Platform-HPC	Responsibility
Compliance	●	●	●	Customer
Security operations	●	●	●	Shared
User security	●	●	●	Shared
Logging, monitoring	●	●	●	Provider
File security and integrity	●	●	●	Provider
Redundancy and backups	●	●	●	Provider
Application security	●	●	●	Customer
OS security	●	●	●	Provider
Network security	●	●	●	Provider
Physical security	●	●	●	Provider

Conclusion

Technology today is evolving at an unprecedented rate. The cloud too has matured and is constantly evolving. What was once seen as a new, exciting, and unproven option when it came to computing has now shown itself to be a mature, highly secure, and extremely flexible solution at a time when organizations across industries are looking for new answers to long-standing IT challenges.

Companies at every level have struggled to manage limited resources and maintain internal hardware and software systems, while meeting the engineering productivity demands.

Maintaining a fixed capacity, on-premises solution is not a viable strategy in an environment where aggressive competitors leverage limitless scale, ease of deployment, and economical systems.

Enterprises today who take the leap and migrate to platform-managed, cloud-enabled HPC solutions will benefit from a competitive advantage in productivity and greater ROI of their innovation investments. Flexible, dynamic of cloud-enabled systems offer companies near-instant capability to adjust their IT infrastructure at a rate that fits their users' needs.

By adopting a platform solution and retiring the fixed capacity, on-premises infrastructure model, companies significantly can significantly reduce capital expenditures, dramatically increase productivity, and develop next-generation innovative products at a pace that surpasses their competition.

Enterprises today who take the leap and migrate to platform-managed, cloud-enabled HPC solutions will benefit from a competitive advantage in productivity and greater ROI of their innovation investments