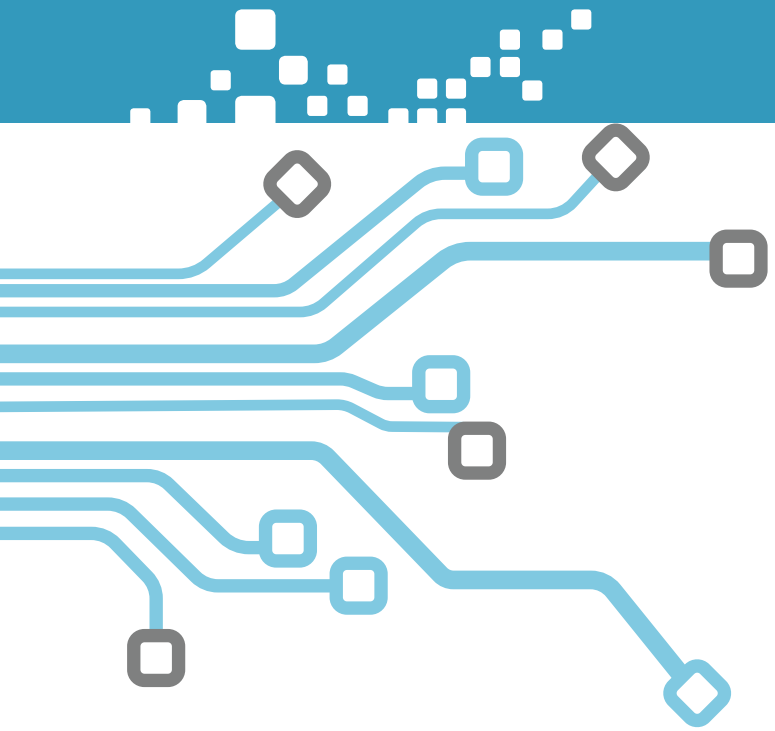

Motivations and IT Roadmap for Cloud HPC

In a competitive market, enterprise companies gain a critical advantage by transforming stagnant IT into a dynamic cloud HPC environment designed to foster innovation and improve product development.

Contents

Introduction

Today's CIOs are under constant pressure to provide the resources and infrastructure that allow their organizations to innovate superior products faster than ever, while maintaining a lean and agile IT structure. Research, design, and development in engineering and science are transforming and expanding at speeds that have never been seen before. Technology is rapidly evolving across industries, and if an organization fails to adapt to this environment, they quickly lose their competitive edge.

As an IT leader, properly allocating resources within the company without wasted expenditures or unmet demands is the true challenge. Leading manufacturers and innovators have realized that relying solely on traditional, rigid, on-premise HPC systems does not allow them to manage the dynamic workload required by their organization to stay ahead of the competition while simultaneously preventing inflated IT costs.

A cloud HPC solution provides the unparalleled agility and unlimited capacity that these organizations need, allowing them to significantly reduce costs, unleash innovation, and accelerate development.

As an IT leader, properly allocating resources within the company without wasted expenditures or unmet demands is the true challenge.

Pain and Limitations with On-Premise HPC Systems

Deployment of on-premise HPC systems has been the standard approach to providing additional computing resources for applications that require IT resources beyond a typical desktop or workstation. There are several critical reasons why exclusively relying on an internal HPC system will significantly hold an organization back:

High Capital Expenditure

Corporate initiatives are constantly requiring that IT organizations 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.

Static Capacity

When organizations build out an internal cluster to meet the "average" or "typical" workloads, they automatically create a situation where peak demand and specialized projects cannot be accommodated. Essential information and results will be delayed for product development teams at the exact time during development cycles when additional simulation and HPC resources are critical for the timely completion of the project.

Unpredictable Resource Planning

Predicting the engineering needs over the life of a cluster can be extremely difficult and tedious. It's often an impossible task for IT teams to complete accurately. Strategic planning over a multi-year period for HPC IT resources requires addressing unknown and unpredictable variance, due to the rapidly changing IT technology landscape and organizational priorities.

Long Procurement Cycles and Hardware Lock In

Organizations require at least six months to procure and install a new HPC cluster, and with the accelerating hardware technology release cycles, this cluster will already be considered dated by the time it can be used by internal teams. With the high cost and time to implement, an organization is locked in with the cluster they've built for an average of three years, while faster hardware becomes available in the marketplace.

There are several critical reasons why exclusively relying on an internal HPC system will significantly hold an organization back.

Cloud HPC is the Solution

There are numerous benefits obtained by migrating or augmenting existing HPC capacity with cloud HPC. Cloud HPC directly results in more efficient productivity and exponentially greater innovation and growth for the organization. With cloud HPC, companies have:

Customized, Scalable HPC

CIOs and IT managers can instantly select the specific hardware needed for their organizations on a scalable basis. Customized hardware is immediately made available without a lengthy procurement process. Companies only access hardware when needed, therefore, resources are never wasted. A truly comprehensive cloud solution offers hardware across different infrastructure providers and global data centers, allowing companies to utilize the configuration best suited for their workload.

Significant Cost Reduction

One of the critical benefits of cloud HPC is the on-demand cost model. Companies only pay for the resources they use, and do not have to make large upfront investments for infrastructure. Since hardware is not on-premise, both capital and operating expenses are drastically reduced.

Consolidated Resources

Cloud HPC providers offer a vital foundation for the simulation process: a consolidated, easy-to-use environment that provides users with all of the resources they'll need within one platform. Users are able to begin running in a shorter time frame with a reduced learning curve. Software, mathematical tools, optimizers, meshing software, and visualization tools are all housed in one location and quickly paired with hardware best suited for the analysis yielding a more efficient and fluid experience for users.

Comprehensive IT Management

Managing resources across offices, countries, and continents is a strenuous job. Cloud HPC aggregates all the moving pieces of simulation resources and processes and makes them accessible from a central location where IT leaders gain full visibility and control. IT teams can easily manage hardware and software, allocate budgets, set user limits, track usage, and generate reports across globally-distributed teams, departments, and offices.

Cloud HPC providers offer a vital foundation for the simulation process: a consolidated, easy-to-use environment that provides users with all of the resources they'll need within one platform.

Improved Product Lifecycle

Product design processes are accelerated and improved, and projects are being finished ahead of schedule and under budget. Engineers and scientists are completing more iterations in less time and are able to fully explore the design space with simulation resources optimized for their workloads. Companies are gaining competitive advantages and pushing innovative products to market at speeds never seen before.

Key Roadblocks to Adopting Cloud Solutions

Security is always of utmost concern to every organization. In today's rapidly evolving technological world, staying abreast of security vulnerabilities is becoming increasingly challenging. Creating hardened on-premise systems is expensive, difficult to maintain, and challenging to enforce among many users. The cost and complexity of creating systems to control hardware, software, and data access is increasing while the occurrence of new potential vulnerabilities and security gaps is becoming more frequent.

In most cases, the threat to, or uncertainty of, security is a main hesitation for moving away from an on-premise system. However, a cloud solution can actually tighten administrator control, provide improved data protection permissions, improve security enforcement methods, provide granular visibility, and strengthen internal security policies and tools – ultimately minimizing and mitigating risks an organization is already exposed to. In order to ensure that you create a highly secure cloud environment, there are several factors to evaluate.

Infrastructure

Finding a single cloud infrastructure provider that meets all of your hardware, software, and performance needs while still complying with your security requirements can be difficult. Often, one provider may meet some, but not all of your requirements. Integrating with multiple providers, each with their own terminology, APIs, and security feature sets is a daunting and complicated task. Cloud HPC environments that offer a unified interface to multiple infrastructure providers saves a company a significant amount of work, providing consolidated terminology and a single API to utilize the best of all available cloud infrastructure providers.

Data Management

Cloud solutions can be built from the ground up using the latest security features, without the burden of integrating with the limited feature sets of legacy systems. Encryption can and should be built into the entire lifecycle of data as it flows through, and is stored in, cloud HPC environments. Disaster recovery and backup processes should be implemented across geographic regions.

Integrations

Companies do not have to abandon their existing infrastructure when utilizing cloud resources. Often, it makes economic sense to use a combination of on-premise and cloud resources. Ensure that the cloud provider offers a variety of options to securely enable bi-directional communication between a company's on-premise resources and their cloud-based resources.

Companies do not have to abandon their existing infrastructure when utilizing cloud resources. Often, it makes economic sense to use a combination of on-premise and cloud resources.

Cloud resources have the ability to significantly decrease the HPC cost for an organization over time both in terms of literal investment, as well as the gains in productivity.

Independent Verification

To maintain the highest security standards, cloud HPC providers should submit to periodic external security reviews by independent auditors. The reports generated from these audits can be obtained under NDA from cloud providers. Look for providers whose assurance programs align with company requirements.

Complete Visibility

Cloud HPC environments should provide complete visibility over a company's cloud resources, integrations, users, projects, activity, and costs. High-level visibility for at-a-glance monitoring as well as detailed reporting should be available to ensure a company makes full and efficient use of their resources.

Network Latency

A frequent concern when analyzing a potential cloud migration is the issue of network latency. The main bottleneck with a cloud solution would be network bandwidth between on-premise infrastructure and the cloud. The decision to move to the cloud is predicated on the need for increased scale and a clear limitation with insufficient internal resources to match demand. However, many companies are concerned that the increased efficiency and elasticity gained by moving to the cloud will be minimized by the anticipated delays of file transfer times that are part of this migration. The best way to improve speed is to avoid transferring data when possible. It is critical to choose a cloud provider that offers an end-to-end, comprehensive environment in which users have the ability to remotely visualize files within the cloud interface itself, allowing for pre/post-processing in the cloud and a decreased need for subsequent data transfer.

Limited IT Budget

Limited budgets, combined with the fear of increasing new resources costs, is another common reason for pause as organizations consider the cloud. The reality is just the opposite. Cloud resources have the ability to significantly decrease the HPC cost for an organization over time both in terms of literal investment, as well as the gains in productivity. The agile nature of cloud environments allows for immediate implementation of resources and the ability to quickly respond to unexpected IT demands, yielding increased savings and a decrease in wasted time and resources. Over time, transitioning to a complete cloud solution gives an organization the ability to both retire on-premise hardware and software infrastructure, as well as remove the costs associated with updating, maintaining, and implementing new systems. More significantly, the total cost of ownership (TCO) for a cloud solution has been proven to be dramatically less than that for an on-premise system and will be examined further in this paper.

Unfamiliar Territory

Once an organization realizes that the typical barriers around security and budgets are clearly addressable, the critical barrier that remains is how to build a clear roadmap to adopt cloud solutions for both the short and long term. Education and knowledgeable IT is paramount. Lack of knowledge and understanding of the cloud and how to migrate in an efficient and effective way keeps many organizations from transforming existing systems to a solution that is clearly superior to the status quo. The right knowledge, training, and guidelines will produce an environment that allows your organization to improve and expand on their designs and operations while generating superior business results. Having a properly educated and prepared management and IT team, with a well structured cloud strategy for both the short and long term, is necessary for success in a cloud HPC transition.

The Roadmap: How to Move to the Cloud

Once an organization has recognized that keeping all work on premise is no longer feasible or efficient to stay competitive in today's world, and the decision to migrate to the cloud has been made, the appropriate level of cloud implementation to best suit the organization's needs and workload must be determined. There are four phases of cloud HPC transformation.

Phases of Cloud Migration

Phase 0: **No Cloud Resources**
Phase 1: **Cloud for Overflow**
Phase 2: **Hybrid Cloud and On-Premise**
Phase 3: **Full Cloud Migration / On-Premise Replacement**

For the purpose of this paper, the cyclical engineering demand has been simplified in the following phase graphs. True demand will be much more variant.

Prepaid / Reserved cloud HPC has been limited to one year and three year increments for simplicity.

Phase 0: No Cloud

Completely On-Premise

A company has determined that they are content with their current situation, despite underutilized resources, unmet demand, and delayed research and development. At this time, it is essential to perform an assessment of unmet needs and projected expansion in capacity requirements. When evaluating on-premise resource consumption, it is important to factor in not only hardware purchase cost but also maintenance, hardware upgrades, labor, and the percentage of time systems that are not used. You must determine if continuing to maintain and expand on-premise resources is sustainable both from a capital and personnel perspective, and if you are still capable of meeting the variable needs of engineering (Figure A).

If it is determined that on-premise infrastructure alone will not be sufficient, then the next step is moving to the cloud for overflow.

Figure A

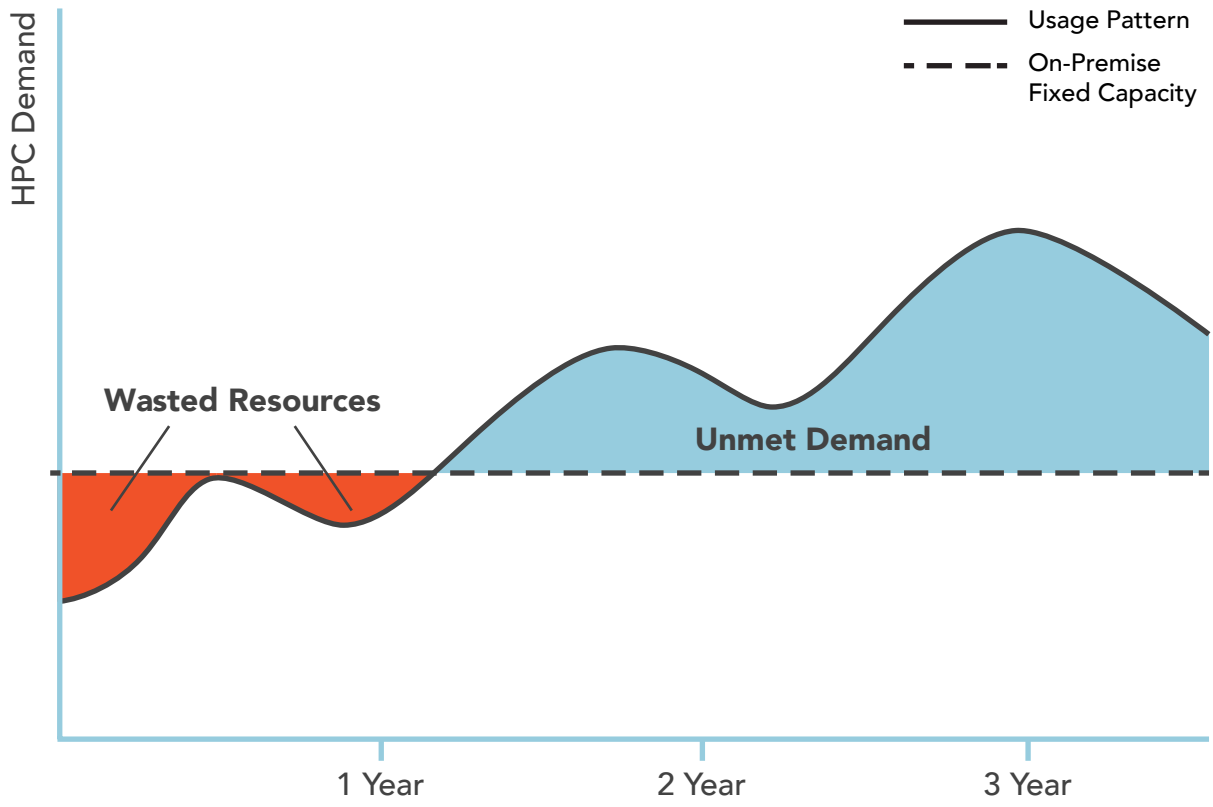

Phase 1: Cloud for Overflow

When the cloud is used for overflow, the objective is to reduce or eliminate queue times, deliver results faster, provide engineers the opportunity to conduct more simulations, and to minimize the capital budget being used to increase on-premise capacity. In this phase of implementation, the on-premise resources and cloud resources are separate with no common interface (Figure B).

Figure B

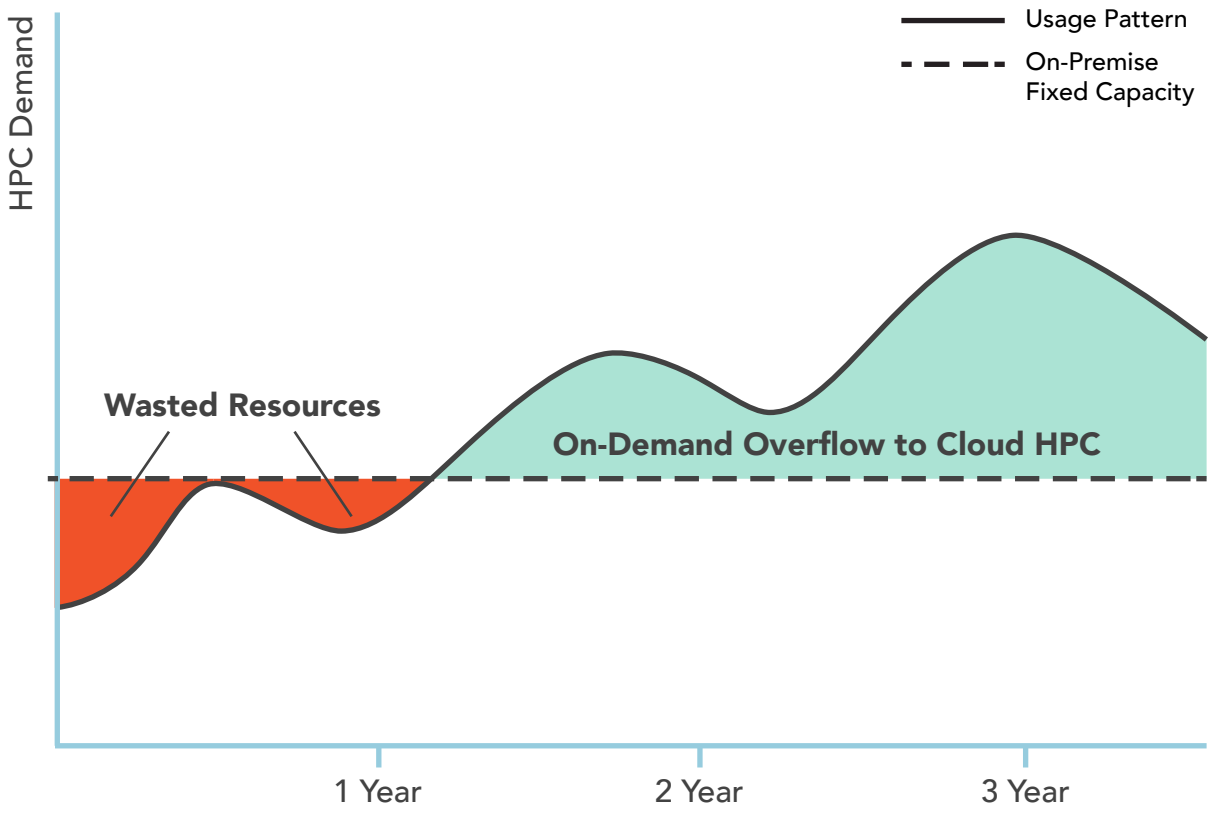

Phase 1 Business Case

Evaluating if cloud for overflow use is justified for your organization requires a comparison of the cost of the on-demand hardware resources with the cost of delayed results or the inability to run all desired simulations during the allotted time. This includes the cost of potential design errors due to the lack of comprehensive analysis, idle engineering time due to long queues, and other subsequent impacts of constrained resources. For most organizations, this is a quick and easy evaluation since the costs of labor and development delays far outweigh the cost of instant cloud resources.

PHASE 1 DEPLOYMENT PLAN

Below are the steps in this phase of cloud integration and the principal teams involved at each step

	Engineers	Management	IT	Cloud Provider
1. Determine how much overflow is needed to begin transitioning to cloud HPC.	●	●		
2. Identify a cloud provider that meets company hardware and software needs.	●	●		●
3. Determine based on company specific firewall and security protocol if IT involvement is required at this phase.	●		●	
4. If required, complete initial security assessment, review and audit.			●	●
5. Determine if engineering teams will be using on-demand, open source or internal licenses.	●			
6. If using internal licenses, establish secure connection between on-premise license server and cloud infrastructure.		●	●	●
7. Determine optimal hardware configuration(s) for each engineering team.	●			●
8. Internal training to understand capabilities, processes and cloud interface.	●			●
9. Set budgets and project / user limitation if necessary.		●		●
10. Identify simulation workflows that are good candidates for the cloud (reduced data transfer required, pleasantly parallel applications such as parameter sweeps, models that benefit from increased core count) and begin using cloud HPC.	●			

Examining the Benefits of Cloud for Overflow

- Yield a shortened product development time and a faster time-to-market with increased access to resources.
- Meet peaks in demand immediately with the option to burst to the cloud on the spot for hardware, software, or both.
- Dramatically reduce capital expenditure that would have gone toward building out internal resources to meet demand peaks.
- Immediately reduce product development schedule risk.

Solving the overflow problem is an easy entry into cloud computing. For some companies this may be sufficient, but for many companies a more cohesive environment between on-premise and cloud resources is needed. At this point, companies can transition to Phase 2 and access existing, on-premise resources and cloud HPC all from one unified platform.

Phase 2: Hybrid Cloud and On-Premise

As cloud overflow usage increases, it becomes more economical for companies to reserve space in the cloud for predictable workloads. While on-demand pricing is valuable for situations that arise unexpectedly, reserved or prepaid capacity is more financially viable as steady-state usage increases. As demand increases, the most cost efficient approach is purchasing block cloud capacity instead of paying on-demand prices or building out internal resources. For example, companies can purchase annual or longer-term customized hardware from a cloud provider to accommodate their changing engineering needs.

As companies invest more in cloud computing, a deeper level of integration between cloud and on-premise resources is required for a smooth and efficient transition. A hybrid cloud / on-premise system, with a unified interface, creates a consolidated environment for both IT and engineering teams. IT managers have full visibility into all utilized resources and can more accurately track usage. Engineering efficiency is drastically improved as engineers have the ability to execute jobs between internal and external hardware without bouncing between separate systems. Since all resources are now available through a common interface, the transition is more fluid and can be accomplished in one step, or gradually over time depending on the needs of the organization. It is critical that a cloud provider can merge a company's existing and cloud resources for a seamless integration (Figure C).

Figure C

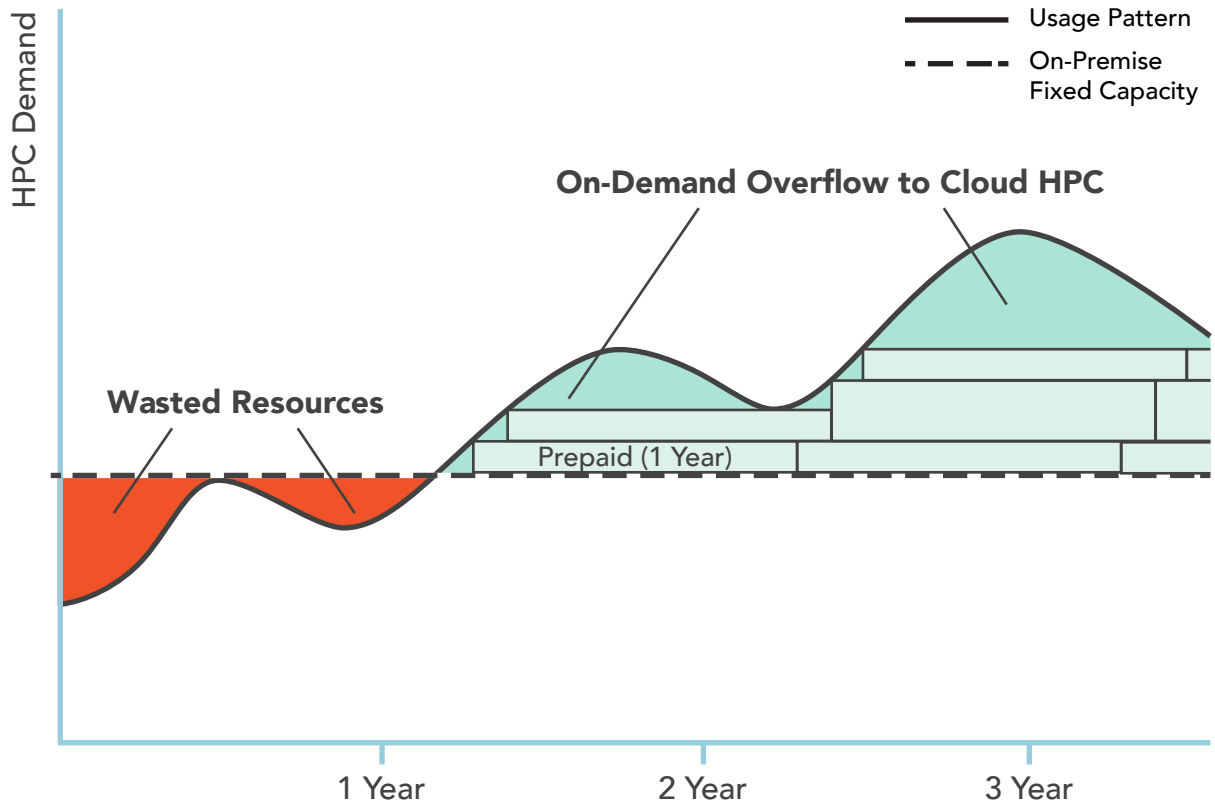

Phase 2 Business Case

Determining if a hybrid environment is the appropriate transition largely hinges on a cost comparison of a company's internal infrastructure and the longer-term savings of cloud HPC. Here, the strategic decision is the opportunity for scalable, customized resource expansion that allows for more competitive engineering growth within a company. While there are often costs associated with a cloud / on-premise integration and security review, these costs are minimal and ensure that a company has limitless turnkey HPC capabilities.

At this point, a company determines how much of its cloud usage is to address steady-state capacity. This level of steady-state cloud usage is not the most economical at a purely on-demand rate. Prepaid computing time is a better path for companies to not only continue to accommodate workload growth, but also reduce longer term capital and operating costs. In addition, migrating steady state capacity has the benefit of shortened procurement times and reduced IT resources. Once companies determine long-term cloud usage, they can move to a hybrid environment.

PHASE 2 DEPLOYMENT PLAN

This is under the assumption that the Phase 1 Deployment Plan has been implemented

	Engineers	Management	IT	Cloud Provider
1. Evaluate historical cloud usage to determine prepaid / reserved capacity and time frame.	●	●	●	
2. Define purchasing rules for overflow and long-term hardware capacity.		●	●	
3. Identify custom hardware configuration best suited for prepaid capacity.	●			●
4. Complete any security reviews and audits before proceeding with hybrid implementation.			●	●
5. Link all company license servers to cloud portal via SSH or VPN.			●	
6. Link on-premise hardware and schedulers as appropriate.			●	
7. Implement process for access to developer API as needed.			●	
8. Set desired security and sign on requirements for users.			●	
9. Administrator defines budgets, controls, permissions, and procedures for departmental and / or access to specific resources.		●	●	
10. Train IT staff on management of platforms through administration portal and train engineers, scientists, and users on platform UI and capabilities.	●	●	●	●

Examining the Benefits of Hybrid Cloud and On-Premise

- Comprehensive enterprise strategy allows for transparency and control for administration and management.
- Leveraging of existing assets avoids underutilization while maximizing design and production capabilities with prepaid cloud infrastructure and overflow resources.
- Improve security, provide a better experience for your engineers, and avoid "shadow IT" by integrating with your existing IT infrastructure.

After a hybrid cloud / on-premise environment is reached, the economic and productivity advantages of cloud HPC are realized. Now, as internal resources become outdated and in need of replacement, companies are able to easily transition to a complete cloud solution for simulation and technical computing workloads.

Phase 3: Complete On-Premise Replacement / Cloud Integration

At this point the advantages and savings gained through both prepaid and on-demand cloud resources, and the hindrances and limitations of static and expiring on-premise HPC, justify the adoption of a complete cloud solution. By retiring internal infrastructure and shifting to a fully cloud model, a company can mimic an on-premise cluster with a long term cloud investment without the procurement expenses, implementation time, operating costs, and exhaustive IT resources, while decreasing capital investment (Figure D). For some companies the financial model of cloud HPC also allows them to transfer their simulation expenses on to their customers or their internal supply chain.

In addition to a significant cost reduction, productivity and engineering efficiency is improved. Management and procurement teams eliminate the lengthy periods of hardware acquisition while engineering teams have fully customized resources available at the exact time they're needed. IT teams no longer have to worry about maintaining complex infrastructure systems or dealing with persistent problems such as down hardware or software package installations. With cloud HPC, companies and organizations are able to achieve full cost optimization while keeping pace with dynamic and expanding engineering workloads.

Figure D

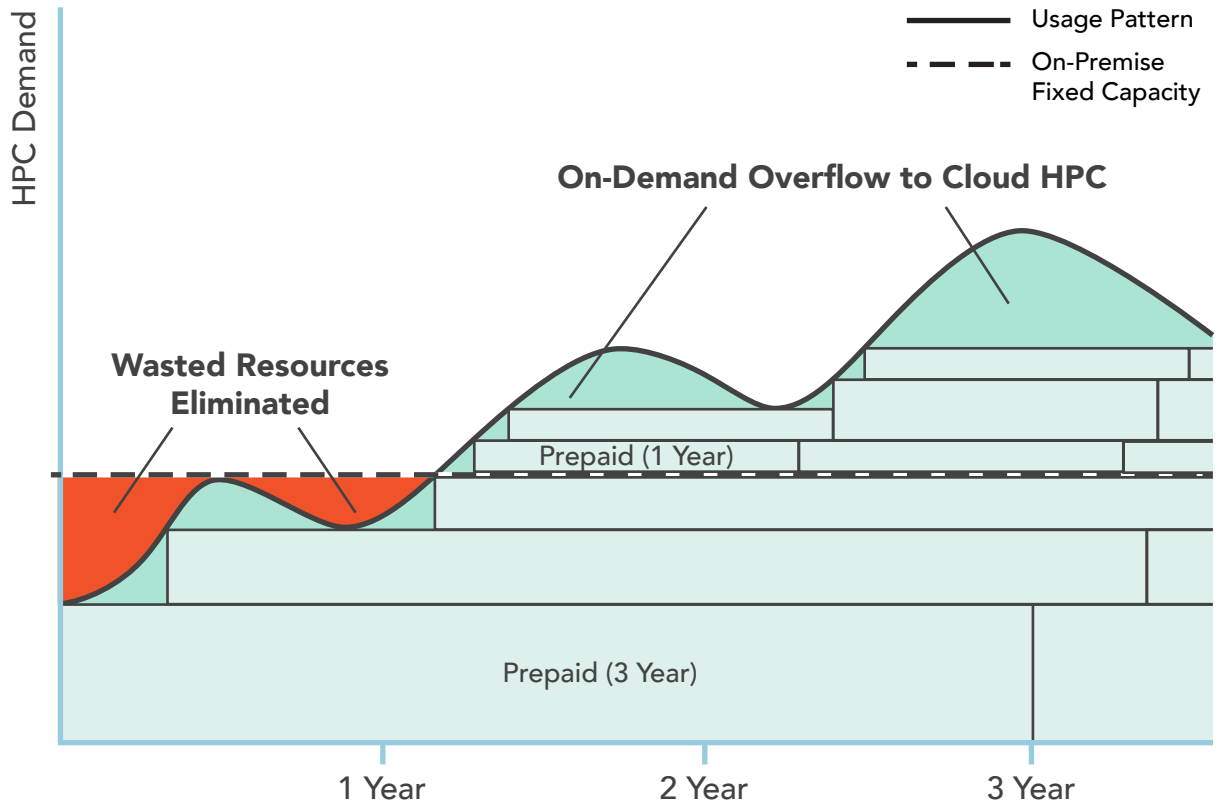

Phase 3 Business Case

Justification for using the cloud for future expansion of fixed capacity is a straightforward comparison of fully burdened internal costs for on-premise capacity and the cost for prepaid cloud capacity. There are a couple of key variables associated with this analysis:

- **Complete internal HPC cost** - Many companies only factor in the cost of the physical hardware when comparing internal expenses to cloud HPC. In addition to the actual hardware costs, companies must also evaluate the cost of electricity, personnel, and facilities. While some assume that these costs are fixed, it's important to factor in these expenses for a true evaluation since no costs are truly fixed.

- **Cloud HPC cost** - This cost should be evaluated for the same time period that a company would keep an internal HPC system. If a company maintains hardware for three years, then the same time frame should be used to evaluate cloud costs.

- **Utilization rates** - Realistic utilization rates for on-premise resources need to be applied in the cost comparison evaluation. For on-premise, 100% utilization is not feasible due to software upgrades, down systems, and other maintenance issues.

Even at an unrealistic 100% utilization, a cloud HPC system over a three-year period is more cost-effective than an internal system. Coupled with the combined advantages of productivity, management, security, and scalability, a complete cloud HPC environment is the best solution for organizations looking to expand their engineering and science work-loads at a competitive rate.

Example of Cloud Integration vs HPC Costs

The graph shown below compares the full expenses of on-premise infrastructure to cloud HPC (Figure E). Cost is shown per core-hour. For on-premise, this analysis is assuming 100% utilization. For cloud computing, there is only one component, the per-hour charge for computing. Specifically, a prepaid price for a 3-year commitment was chosen to closely align with the average on-premise ownership period. For on-premise systems, the cost is broken into three components:

- **Equipment cost** - A 3-year period of ownership is the time period.
- **Electricity cost** - This cost will vary by locally negotiated electricity rates.
- **Personnel and facilities** - This includes the cost of direct and indirect personnel associated with the HPC system, facilities costs, and other miscellaneous operating expenses.

A more detailed breakdown of the on-premise costs vs cloud costs is show in the Appendix.

Figure E

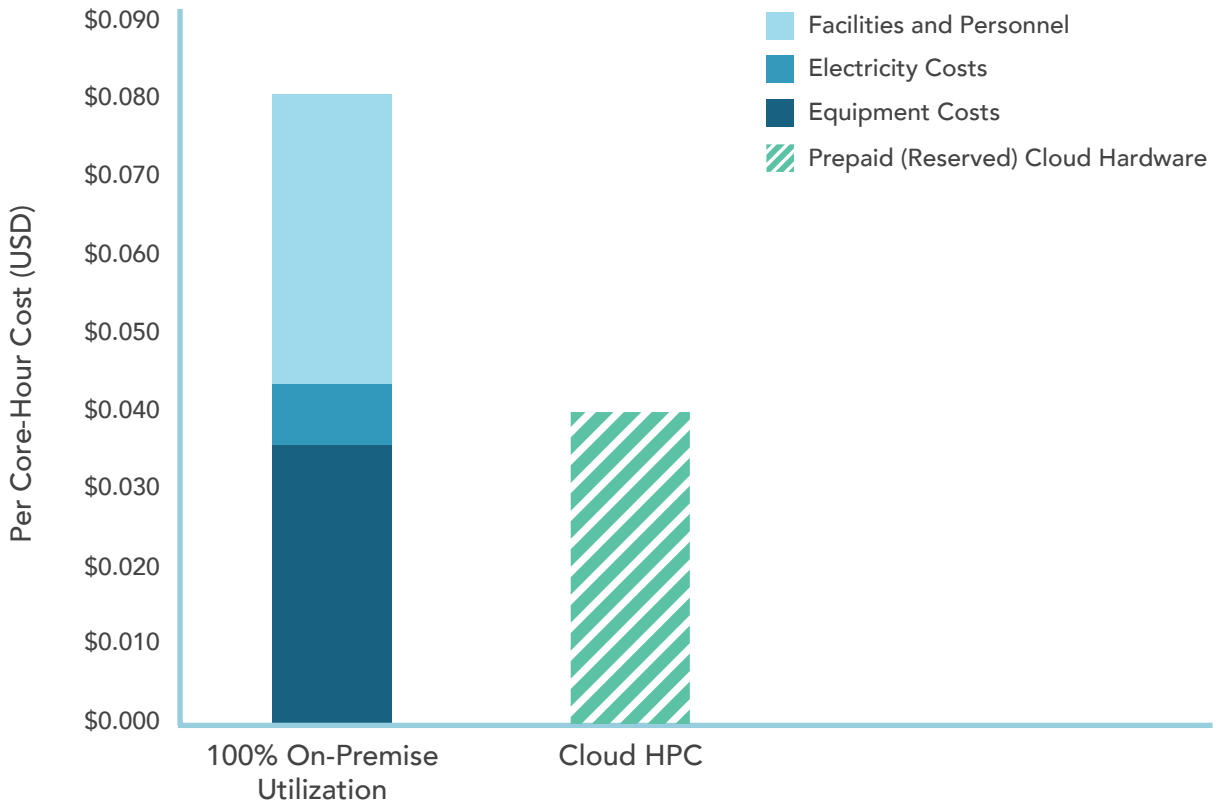

*Please note that power costs here are included in Facilities charge, not electricity.

PHASE 3 DEPLOYMENT PLAN

This is under the assumption that both the Phase 1 and Phase 2 Deployment Plans have been implemented.

	Engineers	Management	IT	Cloud Provider
1. Set up decommission plan for existing data centers and HPC clusters.		●	●	
2. Using historical cloud data, identify long-term, prepaid plan for cloud HPC investment to seamlessly replace internal infrastructure, maximize agility, and achieve the optimal ROI.		●	●	●
3. Develop and implement algorithmic rules for hardware and license purchasing.		●	●	
4. Avoid hardware-lock by identifying cloud providers who offer a range of HPC infrastructure. This maximizes the possibility of leveraging the latest, state-of-the-art hardware.		●		●
5. Deploy simulation platform throughout entire organization via SSO integration.			●	

Examining the Benefits of Complete On-Premise / Cloud Integration

- Provide a fully agile, instantly scalable, and consistency best-in-class HPC system to engineering teams.
- Maintain cost-efficient infrastructure that has minimal capital expenses, eliminated operating expenses, and reduced IT requirements.
- Employ improved management and tracking capabilities by utilizing built-in reporting features from the cloud provider.
- Take advantage of additional on-demand resources including simulation software, optimizers, and mathematical tools.

Companies today who are migrating to cloud HPC solutions have the competitive edge at a time when everything needs to be instant and customizable to meet both internal and external demands.

- Accelerate production and development by removing procurement and implementation latency, ensuring engineers are never using outdated or obsolete hardware.
- Productivity is dramatically improved as job submission queues and underutilized resources are eliminated.

After a hybrid cloud / on-premise environment is reached, the economic and productivity advantages of cloud HPC are realized. Now, as internal resources become outdated and in need of replacement, companies are able to easily transition to a complete cloud solution for simulation and technical computing workloads.

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 been struggling to keep up and to keep their internal hardware and software systems relevant and at sufficient levels necessary to meet their engineering demands. Those who do not take note of the new technologies and transforming environment we live in and do not adapt appropriately will fall behind. Maintaining a strictly or predominantly on-premise solution while leading edge competitors are taking advantage of the limitless scale and ease of deployment afforded by the cloud is not only uneconomical, it is no longer feasible for businesses to sustain.

Companies today who are migrating to cloud HPC solutions have the competitive edge at a time when everything needs to be instant and customizable to meet both internal and external demands. Due to the cloud's dynamic nature, companies can transform their IT infrastructure at a rate that fits their needs. By adopting a cloud solution and retiring their legacy, on-premise infrastructure, companies are significantly reducing capital and operating expenditures, dramatically increasing productivity, and developing next-generation innovative products at a pace that surpasses their competition.

Appendix

Cloud HPC vs. On-Premise HPC Cluster Total Cost of Ownership (TCO)

Number of HPC Servers (Node):	100	Adjust this to model
Number of Servers per Rack:	10	power density
Number of Ports per Server:	2	
Number of Switches per Rack:	2	ToR switches, HA pair
Peak Power per Server:	1800	W
Average Power per Server:	1400	W
Metered Power Cost per KWH:	$ 0.100	
Number of Cores/HPC Server:	16	
Total Cores:	1600	

Depreciation in Months/Rate:	36	Change to 36 for standard
Risk Adjusted Rate:	8%	3-year, 12 for accelerated
Estimated Enterprise Discounts:	40%	1-year discount rate for capital budgeting

Server Ports per Rack:	20	
Peak Power per Rack:	18	KW
Number of Racks:	10	38U ded. servers
Number of Switches:	20	
Design Peak Power:	180	KW
Total Metered Power:	140	KW
Estimated Power Cost:	**$ 10,080.00**	Month

Market / Industry Benchmarks

	Server, dual E5-2670 with 64GB RAM, 640GB SSD, support
Dell R730xd	$ 12,337
HP ProLiant DL380p	$ 12,343
Avg. Industry Price (As of 3/4/2015)	$ 12,340

Cost of On-Premise HPC:

Fixed Assets	**Price Ea**	**Qty**	**Total**	
Server, dual E5-2670 with 64GB RAM, 640GB SSD, support	$ 7,404	100	$ 740,400	Active servers in the rack
Server Maintenance at 15%/server/year over 3 year period	$ 1,111	100	$ 333,180	
Provision for Spare Servers @5% of total	$ 7,000	5	$ 35,000	Cold spares
ToR switches, 10G 48 port, 3-year support	$ 5,000	20	$ 100,000	Active pairs supporting 10-12 servers each
Aggregation Switch	$ 18,000	2	$ 36,000	HA pairs
Racks for servers, CPI TeraFrame or equivalent	$ 3,000	10	$ 30,000	With chimney cooling
PDUs, dual 208V per rack, 3VN3G60 12.6KW or equiv	$ 540	20	$ 10,800	Supporting mgmt & monitoring
10G cabling in-rack	$ 180	200	$ 36,000	Direct attach SFP cabling, fiber
Rack and Stack 1-time deployment, per-rack	$ 1,000	10	$ 10,000	Estimated
Software Cost	$ 5,000	1	$ 5,000	Estimated
Subtotal			**$ 1,331,380**	
Total Fixed Assets			**$ 1,331,380**	
Monthly expense of above fully depreciated			**$ 41, 721**	Monthly
OPEX - Monthly				
Data Center Metered Power	$ 0.100		$ 10,080	
[1] Data Center base cost per peak KW	$ 150	180	$ 27,000	Multi-year contract
Assumes multi-tenant wholesale facility inclusive of cooling, backup power, building security and walled or caged space.				
Operations Mgmt. Non-Staff Cost per Server (nagios or equiv)	$ 50	100	$ 5,000	Per month cost
Operations and IT Management Staffing cost	$ 10,000	1	$ 10,000	One FTE @120K burdened
Subtotal (Monthly)			**$ 52,080**	Monthly
Total monthly OPEX+CAPEX			**$ 93,801**	Monthly
Total spending over term			**$ 3,376,821.97**	

Adjustable costs based on individual company

Appendix cont.

Comparison of On-Premise Versus Cloud

	Per Node	Per Core	Actual Monthly Cost Per Core	
Monthly Cost for On-Premise HPC:	**$ 938.01**	$ 58.63	$ 58.63	
Hourly cost at 100% utilization:	**$ 1.30**	$ 0.08	$ 58.63	
Hourly cost at 80% utilization:	**$ 1.63**	$ 0.10	$ 58.63	
Hourly cost at 60% utilization:	**$ 2.17**	$ 0.14	$ 58.63	
Monthly Cloud Cost:				**Estimated Savings**
Hourly Cloud cost 100% utilization at Prepaid rate:	**$ 0.64**	$ 0.04	$ 28.80	**51%**
Hourly Cloud cost 80% utilization at Prepaid rate:	**$ 0.51**	$ 0.04	$ 28.80	**61%**

Note 1: Cost for data center is assumed to be based on peak power requirements and covers cooling, floor space, N+1 generators, battery backup, building and space preparation, security. Cost as typical for US based data centers.

On-Premise Cost ($/core-hour)

Equipment Only	$ 0.04
Equipment + Electricity	$ 0.04
Fully Burdened Cost*	$ 0.08

Example Cloud cost per core hour	
3-year "Prepaid"	$ 0.04
1-year "Prepaid"	$ 0.08

*(at 100% utilization)

Authors

Sarah Dietz, Business Development Manager, Rescale
Ilea Graedel, Senior Account Executive, Rescale
Ryan Kaneshiro, Chief Architect, Rescale
Adam McKenzie, Chief Technology Officer, Rescale

For more information, please contact Rescale at:
944 Market Street, Suite 300
San Francisco, CA 94102
info@rescale.com
+1 855 737 2253