MANAGING TCO IN HPC HYBRID CLOUD ENVIRONMENTS

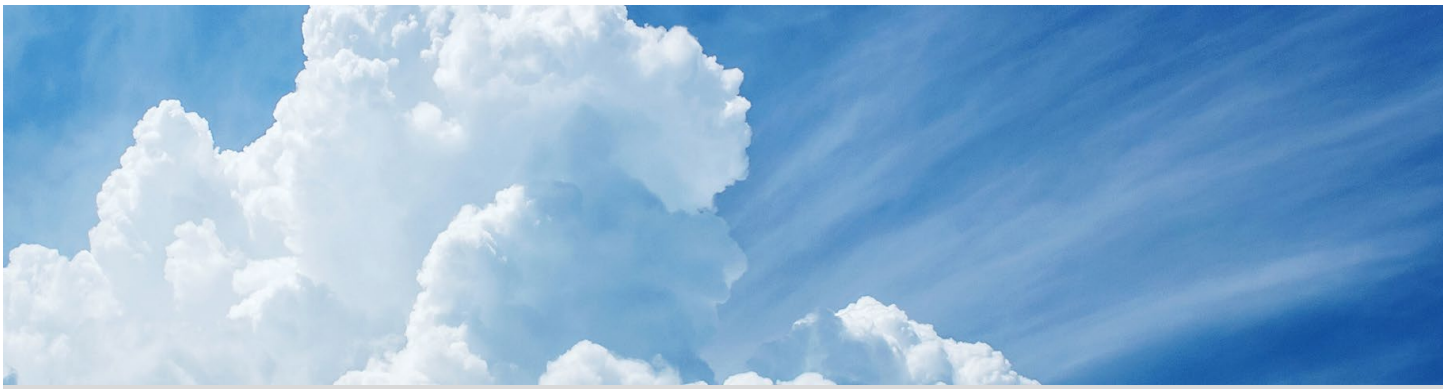

Enabled by improvements in security, new instance types, and fast interconnects, high-performance computing (HPC) users are increasingly shifting workloads to the cloud. With cloud usage increasing, however, managing and containing costs is a growing concern. As organizations become more reliant on cloud, they are also concerned with staying portable and flexible, and avoiding lock-in to a single cloud ecosystem or provider.

In this paper, we make a case for HPC hybrid cloud and explain how operators can manage total cost of ownership (TCO) more effectively. We present a simple TCO model that can help users estimate the cost of hybrid cloud deployments and describe various Altair solutions that can help organizations quickly and cost-effectively implement private and hybrid multi-cloud HPC environments. Using our TCO model, we illustrate how Altair cloud automation and spend management solutions such as Altair® Control™ and Altair® NavOps® can boost productivity and reduce cloud-related expenses while delivering a compelling return on investment (ROI).

HPC in the Cloud

Fueled by the need for fast, convenient access to large-scale computing resources, the use of HPC cloud is accelerating. After years of being stuck at under 10% growth, Hyperion Consulting estimates that HPC cloud revenue will grow at a compound annual growth rate (CAGR) of 17.6% through 2024, making cloud the fastest growing HPC segment.[i] According to Hyperion's 2020 market estimates, cloud revenue is now growing at 2.5x the rate of on-premises server HPC revenue.

Across multiple industries, HPC users see advantages to using cloud computing. Among these advantages are:

- Fast and convenient access to specialized resources such as GPU-capable instances
- The ability to quickly scale infrastructure to accommodate "surge" workloads to meet changing business requirements
- The opportunity to avoid the cost and complexity of managing on-premises infrastructure
- The ability to move to consumption-based pay-per-use models and reduce long-term capital investments
- New artificial intelligence (AI) and data analytics applications that require HPC infrastructure to support the development of new machine learning models

Growing at 17.6% annually, cloud is the fastest-growing HPC segment. HPC cloud revenue is increasing at 2.5x the rate of on-premises HPC server revenue.

New offerings from cloud providers have made cloud more attractive to HPC practitioners. Examples include HPC- and GPU-optimized instances that deliver performance comparable to bare-metal and high-speed interconnects, such as Amazon EFA and Microsoft Azure InfiniBand. Customers can also take advantage of easy-to-deploy, cloud-based parallel file systems and cluster orchestration solutions. As clouds become more capable, traditional barriers to HPC cloud adoption such as performance, security, and data-related challenges are rapidly falling away.

A Complicated Economic Picture

For sites running HPC workloads in the cloud, evaluating the economics can be challenging. Unlike their commercial counterparts, most HPC data centers already run at high utilization, so savings related to virtualization are often less significant.

Some HPC operators report that cloud is more expensive than on-prem environments, while others cite cost as a reason to shift workloads to the cloud.

Many HPC users report that operating in the cloud is significantly more expensive than on-premises infrastructure. However, others cite cost-effectiveness as a reason to shift workloads to the cloud.[ii] While this sounds contradictory, both assertions can be true at the same time. For example, an HPC operator may enjoy lower operating costs on-premises, but if they need to run a surge workload or pilot a new application, running in the cloud may be faster and more cost-efficient.

Much of the debate centers around how organizations account for their spending. Views on the relative economics of cloud vs. on-premises deployments can vary for reasons that include:

- How organizations account for sunk investments such as data centers, HVAC, supporting infrastructure, and the cost of labor
- Depreciation policies on capital assets such as servers, switches, and storage
- The nature of HPC workloads, their suitability to running in the cloud, and the types of cloud instances required
- The efficiency of cloud-orchestration solutions and the ease with which on-premises workloads can be shifted to the cloud

Complicating things further, there are multiple models for HPC cloud adoption. Some customers may prefer easy-to-use software as a service (SaaS) solutions such Altair One™. Others may prefer a more infrastructure-centric approach and rely on workload managers such as Altair® PBS Professional® or Altair® Grid Engine® to managing bursting to public cloud resources as needed.

HPC Cloud Can Provide Compelling Value

Western Digital ran a multiphysics simulation that required 20 days to run on-premises on a one-million-vCPU cloud environment in just 8 hours – a 60X improvement.

Despite the complicated economics, running HPC applications in the cloud can provide compelling business value. Western Digital's well-publicized use of the cloud to conduct a large-scale multiphysics simulation is a good example. Working with engineers from Amazon Web Services (AWS) and Univa (acquired by Altair), Western Digital deployed a 1-million-vCPU cluster comprised of over 40,000 spot instances and reduced the runtime for a job comprised of 2.5 million simulation tasks from 20 days to just eight hours.[iii] Western Digital reported that the cluster, powered by Altair Grid Engine, ran the 8-hour simulation at the cost of USD 137,307 (USD 17,164 per hour) – approximately half the cost of running the same workload in-house.[iv]

This level of scale and cost efficiency was possible only because scheduling and the use of cloud resources were highly efficient. Resources were deployed only for the time required. The workload manager, Altair Grid Engine, and cloud orchestration tools kept resources fully utilized, tolerated instances coming and going at runtime, and took advantage of AWS Spot Fleets to minimize cloud-related costs.

Cloud Spend Management Is a Major Challenge

Regardless of where HPC operators come down on the economics of cloud vs. on-premises environments, nearly all agree that managing spending in the cloud is a growing challenge. Gartner estimates that 80% of organizations will overshoot their infrastructure as a service (IaaS) budgets.[v] There are several reasons for this:

- Spending on cloud resources as a percentage of total HPC spending is increasing, making cloud a more visible budget line item.
- Researchers and engineers have an insatiable appetite for computing resources, and unmanaged, some will consume as much resource as they can get.

- As much as 35% of IaaS resources are wasted owing to orphaned instances and data, idle services, improperly sized resource requests, and other factors.[vi]

"If you can't measure it, you can't manage it."

84% of HPC users see value in cloud spend association, but 76% have no automated solution.

Most organizations lack the tools and cost controls needed to monitor and manage cloud spending. In a Univa-sponsored survey conducted by InsideHPC in 2019, 84% of HPC users saw value in cloud spend automation, but 76% had no automated solution to associate cloud spending to workloads – a major gap.[vi]

More Sophisticated Tools Are Required

It is tempting to think that simple cloud bursting solutions will suffice, but real environments are complicated. Decisions about whether to burst and how much cloud resources to allocate can depend on many factors. These include the cost of cloud resources, month-to-date spending against departmental budgets, workload priorities, and data staging considerations. Workload managers generally do not have visibility into all these metrics. Cloud-related costs also need to be measured and tracked even when instances are not bound to clusters and are therefore invisible to the workload manager.

Multi-cloud environments also pose a growing challenge for organizations. Despite their best efforts to standardize on a single cloud provider, HPC users may find themselves operating across multiple clouds. A Gartner survey of public cloud users conducted in 2019 found that 81% of respondents were using two or more cloud providers.[vii] Reasons include mergers and acquisitions, external collaborations, and the need to access datasets residing in various public clouds. In other cases, HPC application providers may have an affinity to a particular cloud.

To accommodate this multi-cloud reality, hybrid cloud solutions need to interact with multiple cloud providers, each having different APIs, feature sets, and identity and access management (IAM) frameworks. HPC operators also need to consolidate reporting on resource utilization and spending across multiple cloud providers and track spending by project and cost center.

Organizations can develop in-house solutions or employ plug-ins to burst to a single cloud provider but solving these problems across multiple cloud providers is much more difficult. For these reasons, most HPC cloud operators will see value in tools that provide cloud automation and spend management across multiple clouds.

A Simple Model for Evaluating HPC Cloud TCO

To help organizations evaluate TCO in the cloud, we present a simple TCO model detailed in Appendix A. The model estimates a three-year TCO for an on-premises HPC environment that includes hardware depreciation, power, and cooling costs. The model also allows users to explore the TCO associated with different hybrid cloud configurations based on user-provided inputs.

To keep our model simple, we make two key assumptions:

1. We assume that most HPC operators will adopt cloud gradually. The InsideHPC survey data referenced earlier supports this.[ii] As such, we do not attempt to account for financial variables unlikely to change in the short term, such as data center space and fixed infrastructure.
2. We intentionally omit some secondary cost items, both on-premises and in the cloud. For example, we do not consider on-premises network infrastructure and storage, variable cloud costs such as VPCs, object storage, VPNs, or data egress costs. These costs will vary depending on the customer environment.

Figure 1 illustrates a typical cloud bursting scenario detailed in Appendix A. In a baseline scenario (the red line), we assume 200 dual-socket servers spread across ten on-premises data center racks. The solid blue line represents the variable workload. In this example, average utilization may be in the range of 75%.[viii]

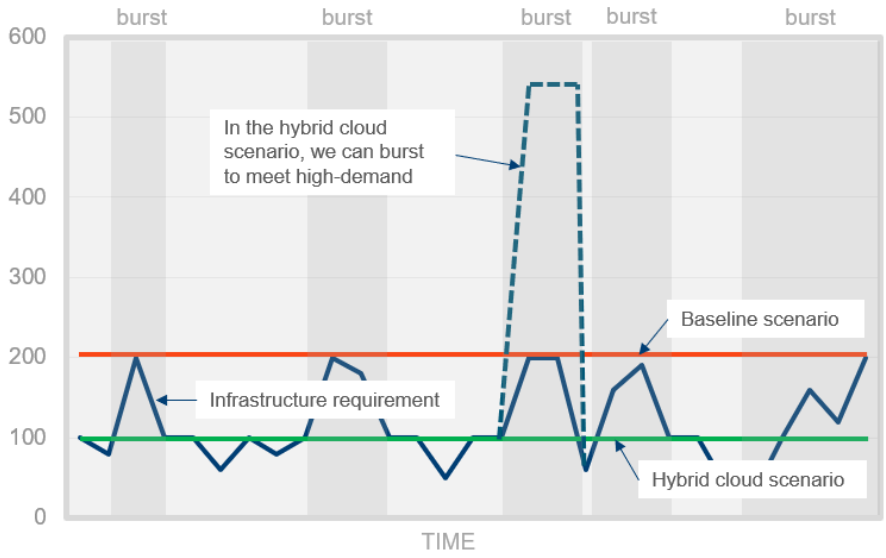

Figure 1 – HPC Hybrid Cloud Bursting vs. a Static On-premises Cluster

In the hybrid cloud scenario (represented by the solid green line), we reduce the on-premises server footprint by 50% to 100 servers. We assume that when capacity beyond 100 servers is required, cloud capacity will be automatically provisioned and bound to the cluster using a cloud automation tool or bursting features built into the workload manager. Cloud resources will only be deployed for the time that they are needed.

In our financial model in Appendix A, we model the same workload (the solid blue line) in both on-premises and hybrid cloud scenarios. However, it is worth noting that in the hybrid cloud scenario, we can burst to arbitrarily large amounts of resources to support periods of high demand. The blue dotted-line spike shows the added flexibility afforded by hybrid cloud. Assuming our application and workload manager can scale to use the additional capacity, we can choose to spend more money on cloud capacity to increase productivity.

Figure 2 shows an annual TCO comparison for server infrastructure before and after implementing a hybrid cloud. This comparison reflects our cost assumptions detailed in Appendix A and assumes the same workload in both scenarios. In this example, we expect to save money in the hybrid cloud scenario. Savings occur because we are paying for cloud resources only when they are needed while realizing higher utilization of our on-premises environment.

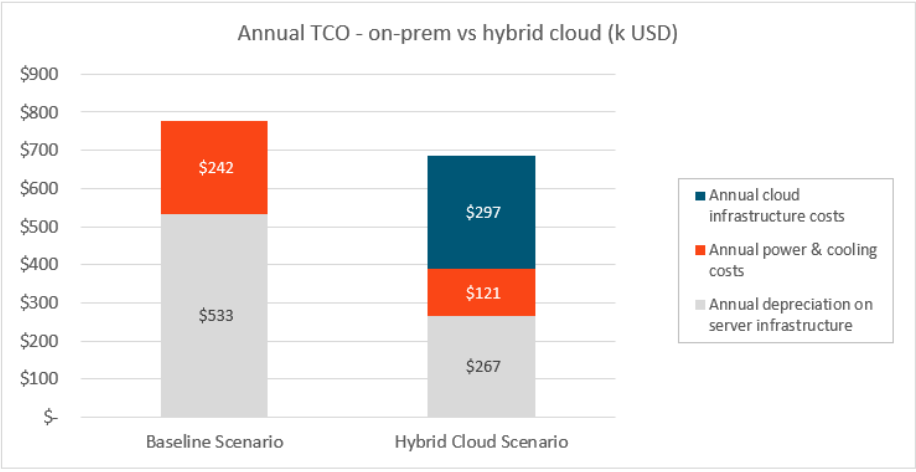

Figure 2 – Comparing Annual TCO of On-premises vs. Hybrid Cloud

In other scenarios, hybrid cloud costs may be higher than on-premises costs, which is perfectly acceptable to many organizations. Hybrid cloud costs will be higher if HPC operators decide that improved productivity warrants additional cloud spending.

A variety of factors can influence costs in the cloud. Examples include:

- Embarrassingly parallel jobs manifest on a cluster as discrete workloads. These workloads are frequently preemptible and re-runnable and can take advantage of lower-cost spot instances or spot fleets.
- MPI workloads may require high-performance interconnects such as Amazon EFA or Azure InfiniBand along with parallel file systems. These workloads will require more expensive infrastructure, both on-premises and in the cloud. MPI workloads generally do not tolerate instances being pre-empted so these workloads will require more expensive on-demand or reserved instances.
- Perhaps most importantly, we have assumed that bursting to the cloud is essentially "frictionless" and is transparent to applications and users. This level of automation is difficult to achieve in practice, especially in multi-cloud environments. If cloud bursting requires administrator intervention or manual steps on the part of users, the hybrid cloud scenario becomes relatively less efficient and more costly.

Automation and Cloud Spend Management Are Critical to Reducing Cost

Cloud resources need to be provisioned quickly and automatically based on flexible policies in a manner that is transparent to HPC users and applications. It is simply not practical to have a cloud or cluster administrator in the middle of decisions to scale cloud resources up or down.

Where cloud resources are concerned, time is money. Cloud instances need to be provisioned quickly and accurately with the correct software images. To avoid instances sitting idle, they need to be made available to the workload manager immediately. They must have access to needed datasets, applications, containers, and software licenses.

Ideally, the workload manager should take advantage of cloud resources as they are provisioned rather than waiting for all cloud instances in an expanded cluster to become available. High-throughput scheduling is also critical, especially in large clusters, to keep cloud resources fully utilized to avoid idle time and unnecessary spending in the cloud.

Finally, to make informed decisions around cloud bursting and manage spending against budgets, the cloud management software needs visibility to cost-related metrics. These metrics include budgets by department or cost center, month-to-date spending against budgets, and the cost of various types of cloud instances.

Altair Improves Productivity and Reduces Costs in the Cloud

Altair provides a variety of cloud-friendly solutions to help HPC users across multiple industries take advantage of hybrid cloud resources while containing costs.

Turnkey HPC in the cloud – For users of Altair simulation, data analytics, IoT, and even third-party software, Altair offers fully managed physical or virtual Altair Unlimited™ appliances deployable on-premises or in your choice of clouds. With this simplified approach, users can simply run their workloads and let Altair Unlimited manage scaling so users can focus on their work rather than managing cloud resources.

Altair One – For users that need access to a broad set of HPC applications along with collaborative access to simulation and data analytics in the cloud, Altair offers Altair One. With Altair One, users can run cloud versions of their favorite software tools and manage on-premises and scalable cloud resources through a single pane of glass.

Cloud automation is critical to enabling HPC hybrid cloud. Cloud resources need to be provisioned quickly and automatically based on flexible policies in a manner that is transparent to HPC users and applications.

Altair PBS Professional, Altair Control, and Altair® Access™ – For users that want to retain more control over their environments, an intuitive cloud bursting graphical user interface (GUI) is built into PBS Professional. HPC customers can also take advantage of cloud bursting facilities in Control, a control center for managing, optimizing, and forecasting HPC resources in the cloud. With these solutions, HPC operators can easily burst to their preferred cloud provider, including Oracle Cloud Infrastructure, Google Cloud Platform, Microsoft Azure, and Amazon Web Services (AWS). Access provides a simple, powerful interface for remote visualization, job submission, and monitoring across remote clusters and clouds.

Altair NavOps – New to Altair's HPC cloud product portfolio is NavOps, a cloud automation and spend management platform for migrating compute-intensive HPC workloads to the cloud. NavOps is designed to work with Altair Grid Engine. It provides organizations with insights into spending against budgets and end-to-end visibility into HPC cloud resources and applications.

Unlike tools that provision cloud resources based on simple policies, NavOps is application, resource, and budget-aware. Pre-packaged or custom automation applets help organizations automate decisions based on up-to-date cloud and workload-related metrics to scale and manage multi-cloud infrastructure, workloads, and data. By combining sophisticated automation with cloud spend management, organizations can boost efficiency, reduce cloud costs, and improve time-to-results, ultimately improving revenue and profitability.

Rapid Scaling – For users of Altair's high-throughput scheduling tool, Altair® Accelerator™, the Rapid Scaling solution provides efficient, cost-conscious scheduling in the cloud. Rapid Scaling is robust enough for even the most demanding semiconductor and EDA workloads, offering sub-millisecond latency for best performance and a scalable, small-footprint, high-capacity architecture capable of scheduling millions of jobs per day.

Designed specifically to optimize workloads for the cloud, Rapid Scaling accounts for speed and license availability to scale cloud resources back down to zero when they're not needed, bringing cloud costs closer than ever to exact demand.

Additional Savings Opportunities in the Cloud

Cloud automation and spend management can help HPC administrators realize additional savings opportunities and further reduce cloud spending. As some examples:

- Automatically detecting and shutting down idle instances and services
- Optimizing the use of lower-cost spot and spot fleet instance types
- Right-sizing cloud instance requests for improved efficiency
- Reducing data storage costs with automated data movement
- Avoiding overshooting IaaS budgets with policy-based automation

By leveraging cloud automation and spend management features in NavOps and other Altair tools, cloud spending can by reduced by up to 35%.

In Table 1, we build on our TCO example and show how cloud automation and spend management tools can further reduce costs in the cloud. This example is based on customer experience deploying NavOps. Actual savings will vary by customer depending on application workloads and infrastructure requirements. PBS Professional customers can realize similar benefits using Control.

Estimated 3-year TCO for 200-node On-premises Cluster (Baseline)		$2,327,000
Estimated 3-year TCO for hybrid cloud scenario		**$2,055,000**
Total cloud spending over 3 years in hybrid cloud scenario (above)		**$891,000**
Additional savings using NavOps for cloud automation and spend management based on 3-year total cloud spending of $891K		
1) Automatically detect and shut down idle instances	~5%	$44,500
2) Optimize use of lower-cost Spot and Spot Fleet instances	~10%	$89,100
3) Right-size cloud instance requests for improved cost efficiency	~15%	$133,650
4) Reduce data storage costs with automated data movement	~5%	$44,500
Total savings opportunity on 3-year clous costs with NavOps	~35%	**$311,750**
3-year cloud costs after implementation of NavOps		**$579,250**
Estimated 3-year TCO with the addition of NavOps		**$1,743,250**

Table 1 – Reducing cloud spending with cloud automation and spend management

By leveraging automation and spend management capabilities in NavOps, cloud spending in our hybrid cloud TCO model can be reduced by an estimated $312K over three years. These savings are illustrated in Figure 3.

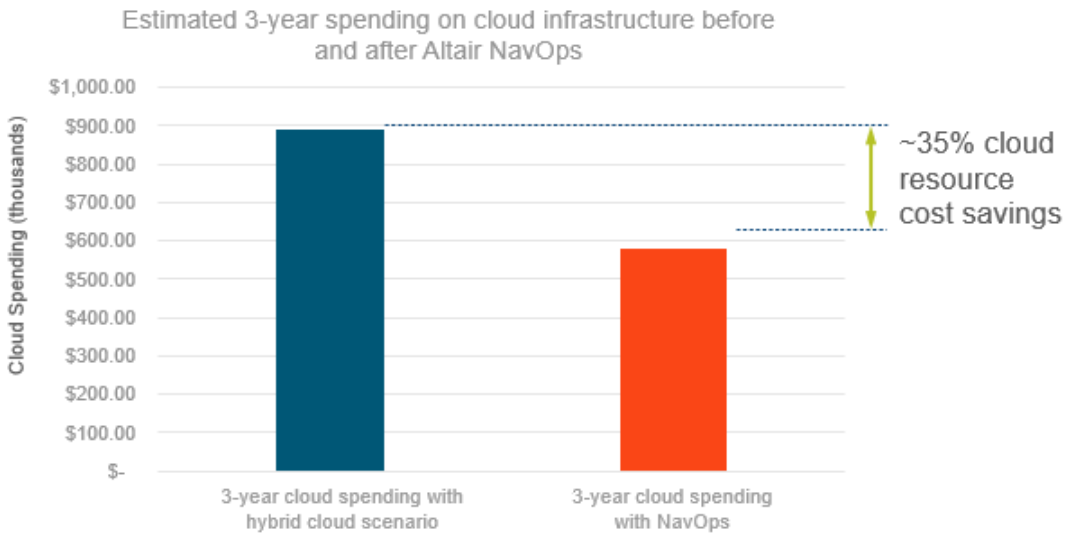

Figure 3 – Savings with NavOps Cloud Automation and Spend Management

A Compelling ROI for HPC Cloud Management Software

The opportunity to reduce cost in HPC hybrid cloud environments is significant. An investment in cloud automation and spend management software can:

- Increase the efficiency of the hybrid cloud environment
- Improve overall productivity and flexibility
- Simplify support for multi-cloud environments
- Provide a compelling ROI based on cloud-related cost savings alone

A sample ROI calculation is provided in Appendix B based on experience deploying NavOps with Altair Grid Engine in a customer's hybrid cloud environment. In one scenario, an investment of $180,000 over three years generated cost savings and efficiency gains of $424,250, an ROI of 136%, more than double the three-year investment.

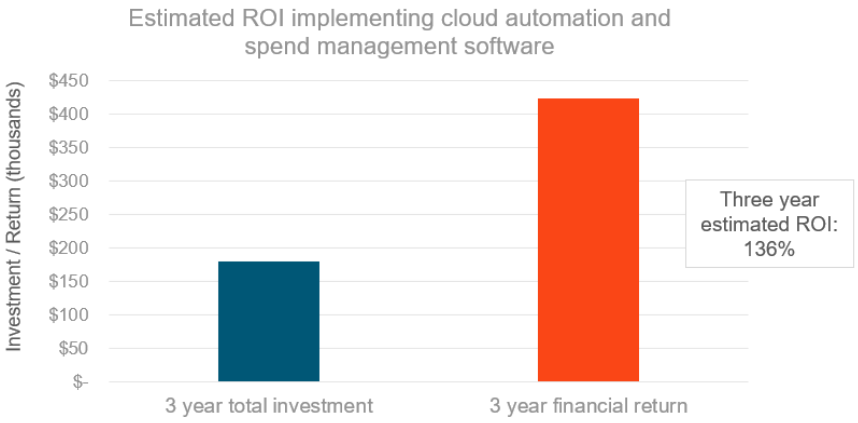

Figure 4 – ROI Implementing NavOps – Details in Appendix B

Summary

For HPC operators, there are compelling benefits to leveraging cloud infrastructure. Organizations can take advantage of the latest state-of-the-art hardware, accommodate spikes in workloads, and avoid the cost, complexity, and risk associated with managing equipment on-premises.

Cloud resources are expensive, however. As organizations increase their reliance on the cloud, they find themselves at significant risk of overshooting IT budgets. This is especially true without automated tools to help them track and manage cloud spending. HPC operators need multi-cloud management solutions to keep pace with evolving requirements, use cloud resources efficiently, improve productivity, and track and control spending against departmental budgets.

Altair provides a variety of cloud-friendly solutions to help HPC operators across multiple industries take advantage of hybrid cloud resources while containing costs. NavOps provides automation and cloud spend management for Altair Grid Engine users. Control provides similar capabilities for PBS Professional users.

With Altair HPC cloud management solutions, customers can:

- Simplify the implementation of HPC hybrid clouds
- Improve overall productivity and efficiency
- Operate seamlessly across multiple clouds
- Significantly reduce cloud spending

Contact us to learn more about Altair cloud solutions or visit altair.com/cloud.

Appendix A: On-premises vs. Hybrid Cloud Financial Model

On-premises server infrastructure

Capital cost per HPC server	$ 8,000.00	USD
Depreciation period	3	years
Cost per server per year	$ 2.67	K USD
Number of servers	200	
Annual cost of servers	$ 533	K USD

Power & Cooling estimates

Power draw per server	800	W
Servers per rack	20	
Number of racks	10	
KW per rack	16	kW
Production load	80	%
Total Kilowatts (compute)	160	kW
PUE (power use efficiency)	1.8	
Total power incl cooling	288	kW
Adjusted power based on load	230	kW
Annual power consumption	2,018,304	kWh per year
Cost per kWh	$ 0.12	$/kWh
Total annual power costs	$ 242	K USD per year

Cloud cost assumptions

Cloud instance type	c5ad.16xlarge	(US East Virginia)
on Demand hourly cost	$ 0.688	per hr.
Estimated instance hours / month	36,000	
Estimates instance hours / year	432,000	

Baseline on-premises scenario (based on 200 server environment)

Annual depreciation cost for server infrastructure	$ 533	K USD
Annual power and cooling costs	$ 242	K USD
Total annual TCO	$ 776	K USD
Total 3 year TCO:	**$ 2,327**	K USD

Hybrid cloud scenario (100 servers on premises)

(assumes that 100 cloud servers are required 14.4 hours per day)

Annual depreciation cost for server infrastructure	$ 267	K USD
Annual power and cooling costs	$ 121	K USD
Annual on-demand cloud costs	$ 297	K USD
Total annual TCO	$ 685	K USD
Total 3 year TCO:	**$ 2,055**	K USD

Notes about the financial model:

- Shaded boxes represent inputs to the model. All other cells are calculated based on input cost assumptions.
- Pricing estimates are based on the online Dell build & price tool - https://www.dell.com/en-us/work/shop/cty/pdp/spd/poweredge-r6525/pe_r6525_13783_vi_vp?configurationid=90be80f3-aeea-4cef-95f6-fde0e0f91589
- Dell model R6525 1U 2P server, 2 x AMD EPYC 2nd Gen 7F52 processor, 10*64GB DIMM modules @ 3200 MT/s = 640GB, 2x SATA 480GB SSDs, 2 x SATA 800GB SSDs – list price is roughly $17K per system configured. Cost per server reflects a significant discount reflecting the street price of whitebox server alternatives.
- Power & Cooling efficiency assumptions:
  - Assume a PUE of 1.8;
  - total power consumption calculated using the methodology described at https://www.racksolutions.com/news/blog/server-rack-power-consumption-calculator/;
  - 200 server environment - with dual-socket servers using 800W power supplies (based on specs of the reference server described above), and
  - assume average 80% load for power calculation purposes.
- Assume that the on-premises environment can be shrunk from 200 servers to 100 servers with cloud infrastructure picking up the remaining workload.
- Assume that there are no immediate savings in data center space or infrastructure since these costs are fixed – savings are related to depreciation on server infrastructure, power, and cooling only. The costs of racks and switches are neglected in this model, but if counted, would reflect additional savings. The cost of host adapters is included in server cost estimates.
- Cloud-related cost assumptions:
  - Assume we need 32 physical cores (64 vCPUs) of AMD EPYC 2 2nd gen processors with SSD storage (3.3 GHz) - https://aws.amazon.com/ec2/instance-types/ - The c5ad.16xlarge model at an on-demand price of $0.68.8 per hour is a close (but not perfect) fit reflecting the same capabilities as the on-premises servers - 32 physical cores on 75F2 requires 64 vCPUs for equivalent compute capacity.

Appendix B: Detailed ROI Calculation

NavOps 3-year ROI Calculation	Per Year	3 Years
NavOps server subscription and support	$40,000	$120,000
Initial installation, training, and start-up costs (first year only)	$15,000	$15,000
Variable usage costs	$15,000	$45,000
Total 3-year investment		**$180,000**
Assumptions: • $297K in annual cloud spending • 35% annual reduction in cloud spending with NavOps • 25% efficiency gain for a single HPC cloud admin (burden cost $150K/year)		
Savings in cloud spending (35% of cloud costs)	$103,950	$311,850
Savings in HPC administrator time (25% savings)	$37,500	$112,500
Return on investment over 3 years		**$424,350**
3-year ROI (total savings - total investment) / (total investment)		**~136%**

[i] November 18th, 2020 – Hyperion Research: COVID-19 Changing the Global 2020 Market Landscape - https://www.hpcwire.com/2020/11/18/hyperion-research-covid-19-changing-the-global-2020-hpc-market-landscape/

[ii] Based on Univa / InsideHPC market survey and customer feedback – Cloud Adoption for HPC: Trends and Opportunities - https://insidehpc.com/white-paper/cloud-adoption-for-hpc-trends-and-opportunities/

[iii] Univa Powers Million-core Cluster on AWS for Western Digital - https://insidehpc.com/2019/02/univa-powers-million-core-cluster-on-aws-for-western-digital/

[iv] Western Digital HDD Simulation at Cloud Scale – 2.5 Million HPC Tasks, 40K EC2 Spot Instances - https://aws.amazon.com/blogs/aws/western-digital-hdd-simulation-at-cloud-scale-2-5-million-hpc-tasks-40k-ec2-spot-instances/

[v] The Real Source of Cloud OverSpend? The Shift from CapEx to OpEx - https://www.networkcomputing.com/cloud-infrastructure/real-source-cloud-overspend-shift-capex-opex

[vi] Gartner May 24th, 2019 – Are you Ready for Multicloud and Intercloud Data Management? – https://www.gartner.com/en/documents/3923929/are-you-ready-for-multicloud-and-intercloud-data-managem

[vii] Gartner – Why Organizations Choose a Multicloud Strategy - https://www.gartner.com/smarterwithgartner/why-organizations-choose-a-multicloud-strategy/

[viii] This number will vary depending on how utilization is defined. Note that some HPC centers claim average utilization close to 100%, but these centers tend to measure utilization based on “job slots” rather than the actual percentage of CPU and memory resources consumed.