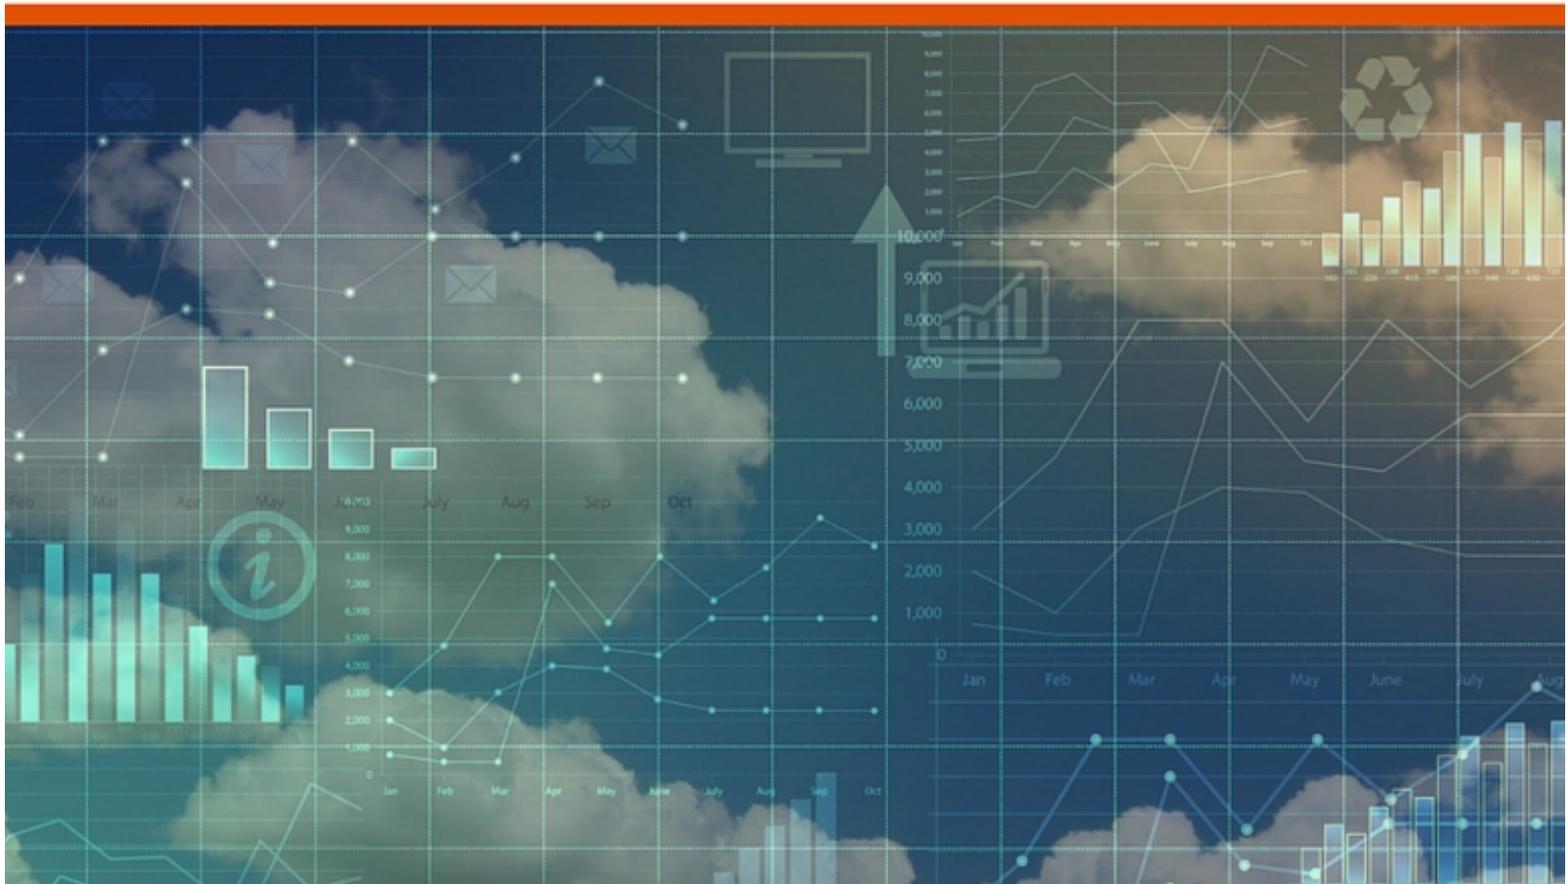


Efficiency in Orchestrating Best-in-Class HPC Cloud Solutions

A V E R E

 **CYCLECOMPUTING**

By Rick Friedman & Scott Jeschonek



1.0 Efficiency in Orchestrating Best-in-Class HPC Cloud Solutions

Cloud solutions offer virtually unlimited capacity plus valuable new capabilities like massively scalable NoSQL and machine learning services for high-performance computing (HPC) applications.

But harnessing that capacity and exploiting predictive analytics and other advanced cloud services is not without challenge. Differentiating between cloud hype and real value, minimizing infrastructure decisions risk, and architecting the best solutions to address both immediate and long-term HPC needs can try the patience—and sanity—of the most informed and practiced IT team.

This white paper reviews common HPC-environment challenges and outlines solutions that can help IT professionals deliver best-in-class HPC cloud solutions—without undue stress and organizational chaos. The paper:

- Identifies current issues—including **data management, data center limitations, user expectations, and technology shifts**—that stress IT teams and existing infrastructure across industries and HPC applications
- Describes the potential **cost savings, operational scale, and new functionality** that cloud solutions can bring to big compute
- Characterizes technical and other barriers to an all-cloud infrastructure and describes how IT teams can leverage a **hybrid cloud for compute power, maximum flexibility, and protection against locked-in scenarios**

- Describes the value of implementing **performant data access and orchestration/management layers for simplicity and scale**
- Summarizes how IT can leverage **Avere Systems and Cycle Computing solutions to derive more value and productivity from the cloud**, using its resources to tackle even more challenging problems

Companies of nearly every size and in all industries can use cloud HPC for unprecedented access to compute resources, scale, and efficiency. Informed decisions early in the process of cloud deployment will ensure HPC IT teams can build cost-effective, manageable infrastructures that deliver required access for now with built-in flexibility for future application and technology opportunities.

2.0 HPC Challenges Driving Cloud Consideration

In the past, computational requirements drove HPC infrastructure design, and IT teams planned technology refreshes years in advance. Today, the pace of change is faster, and compute resources no longer represent the only gating factor for the size and types of workloads that can be supported in an HPC environment. Beyond delivering sufficient compute resources, HPC IT now must address a myriad of other challenges, including:

- **More data.** IT teams are increasingly challenged to deliver storage capacity and responsive data access due to the massive growth of datasets. Data management can be a nightmare at today's scale—IT must track data types and owners, administer long-and short-term data stores, and manage security and privacy for petabytes of capacity and more.

- **Data center limitations.** Data centers are complex and costly to maintain. Power and cooling optimization takes time and resources. Many organizations have already run out of space and can't afford to build new facilities, forcing difficult decisions about what infrastructure and services to keep on premises.
- **Expanding user expectations.** Tasked with supporting an ever-broadening base of users with widely varying requirements, IT teams must manage as a homogeneous environment a complex set of operating systems, compilers, libraries, and applications. At the same time, users hope for shorter job queues and faster time-to-results.
- **Rapid technology shifts.** HPC IT teams want to adopt best practices and fully leverage technology advances. But evaluating, testing, and validating new approaches and systems overly tax limited technical resources.

Over the past 15 years, the only option was to build out datacenters to handle all workloads, and to expand the data centers as demand grew. But this growth has also highlighted the inefficiencies of on-premises infrastructure—including resource access and allocation—at larger scales.

Cloud HPC has the potential to resolve each of these challenges, giving IT the capability to more quickly and cost-effectively allocate needed infrastructure to HPC users. On-demand access to virtually unlimited resources means researchers, scientists, engineers, and other users no longer have to restrict project size or compete for on-premises infrastructure.

But to derive maximum benefits from cloud deployments, IT strategists must take care to leverage cloud services as part of an overall computation approach and to consider requirements for data management and orchestration to simplify provisioning and automate workflows in the cloud.

3.0 Cloud Benefits: Savings, Scale, and Functionality

Cloud computing offers compelling scale and financial benefits. Economies of scale make cloud infrastructure attractively priced, and additional savings can be realized from reduced maintenance and management costs, increased productivity, and built-in disaster recovery. Virtually unlimited amounts of parallel computing resources are available without capital investment. IT can maintain operational flexibility during scale-out jobs while the cloud provider deals with the many challenges of scaling the physical infrastructure. Users benefit from nearly instant access to exactly the quantity and types of resources required, all delivered in a pay-as-you-go model that ensures no wasted capacity. In contrast to a fixed data center – whose value is driven by the customer’s ability to utilize as close to 100% of resources -- cloud infrastructure exactly matches resources to usage patterns.

Cloud infrastructure also gives HPC users access to advanced functionality such as Google Cloud Platform’s (GCP) Cloud Bigtable NoSQL Big Data database service, and predictive analytics like Amazon Web Services (AWS) Amazon Machine Learning service and the Microsoft Azure Machine Learning suite. APIs already exist for cloud-based image analysis, speech recognition, translation, and much more. These powerful services, combined with expanded cloud HPC adoption across companies of all sizes and industries, including financial services, life sciences, and manufacturing, suggest that cloud has moved well past the peak of the hype curve. Cloud solutions are already delivering real productivity and value to businesses.

Across the organization, cloud HPC offers several advantages over traditional infrastructure. Consider some of the key benefits to users, IT staff, and the business.

Scientists, engineers, and other end users will be able to:

- Experience zero queue times and access capacity in minutes
- Scale compute to the size of problems, not vice versa
- Quickly try/support new computational approaches, simulations and software

System architects can:

- Dynamically adjust workloads to the lowest-cost/impact provider
- Focus on computational excellence, not hardware management
- Efficiently support a wide range of user types

The organization can more easily:

- Match spending to actual consumption
- Increase responsiveness to business dynamics
- Grow its user base without hitting hardware limitations

4.0 The Best of Both: Hybrid Cloud

The availability of so much functionality and economy in cloud HPC begs the question—why aren't more organizations moving their HPC clusters to the cloud? Reasons range from technical issues to financial constraints, corporate policies, and internal politics. On the financial side, leases and depreciation cycles may make it impractical to abandon existing physical assets in the short term. Software infrastructure can be complex and costly to refactor, and software licensing arrangements may deter the movement of some applications. Annual budget cycles and allocation of operational versus capital dollars may not immediately accommodate the purchase of cloud services.

Another issue is network latency and the ability to procure sufficient bandwidth across all geographies. Current data placement and regulatory requirements for data locality can also restrict change. Business continuity is always a concern—no organization can afford to disrupt ongoing operations to change out infrastructure. Corporate management teams – who drive budgets and priorities - have varying levels of cloud awareness, commitment to infrastructure change, and urgency of transition strategies. Existing vendor relationships may make both IT and executive-level staff resistant to change or hesitant to introduce new providers.

The list of obstacles can be daunting for IT. For many organizations, a hybrid cloud strategy offers a non-disruptive path to cloud. A hybrid approach integrates both on-premises and cloud infrastructure, giving organizations time and flexibility to address obstacles while taking immediate advantage of valuable cloud resources. As shown in Figure 1, a hybrid cloud gives IT the ability to run workloads and store data where it makes the most sense for both technical and business needs.

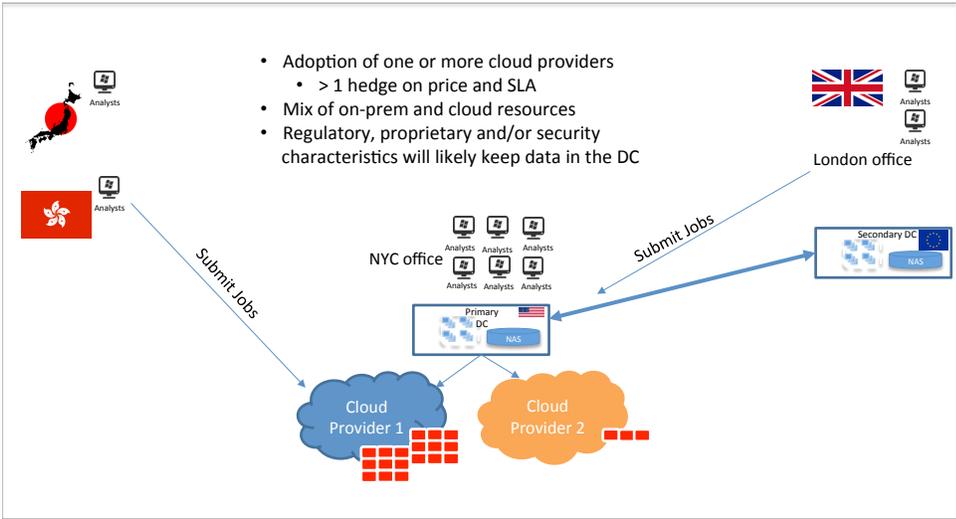


Figure 1: Near Future, Hybrid Cloud

For example, organizations might choose to:

- Keep proprietary or regulated data in their data center
- Run recurring jobs in the cloud while reserving specialized, cutting-edge on-premises hardware/software resources for competitive advantage
- Create multiple isolated cloud environments to give user teams greater flexibility while simplifying IT administration of software versions
- Leverage multiple cloud providers to optimally meet SLA and price objectives across varied workloads

Businesses have already deployed hybrid clouds for exactly these and many more scenarios that are delivering measurable benefits. For most of these organizations, realizing substantial benefits of flexibility and functionality required first overcoming two major challenges: getting application data to cloud compute nodes and orchestrating on-demand clusters or grids of compute nodes in the cloud.

5.0 Data Access/Caching Layer for Speed, Security, and Flexibility

Implementing a data access layer in the compute environment mitigates the time, complexity, and interface compatibility issues associated with moving data from a data center to the cloud or even among cloud providers. Illustrated in Figure 2, the data access or caching layer allows cloud-based applications to run against storage that resides on premises.

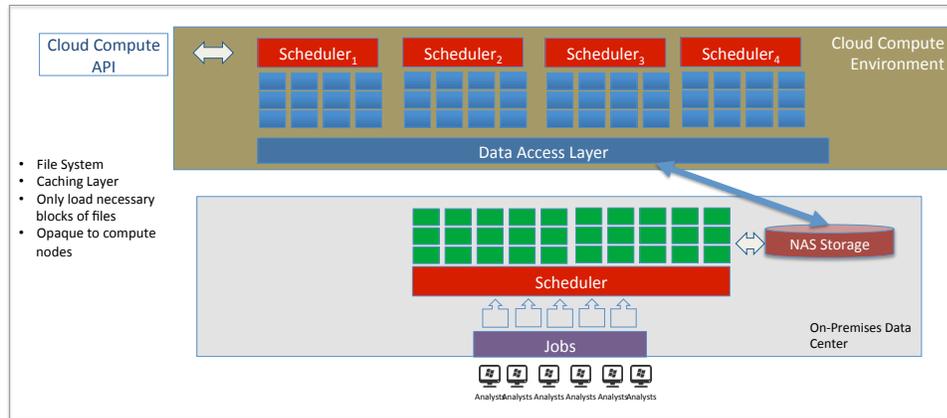


Figure 2: Data Access Layer

This structure offers three key benefits:

- **Speed.** Caching accelerates cloud compute performance by placing most of the required data in RAM deployed close to the compute nodes, helping to avoid ingest latencies and reducing transit latency after the first read. Illustrating the potential benefit of caching in HPC applications, a recent Cloudera study found that in a typical Hadoop job, a frequently read file could be accessed as many as 500 or more times.¹
- **Elimination of security issues.** Using a data access layer reduces security objections by allowing companies to keep data on premises with only transient data stored in the cloud and for only the time that clusters are operating on that data.
- **Flexibility.** The data access layer simplifies access, providing a single mount point to cloud compute regardless of where data is physically stored.

HPC environments can scale across tens of thousands of compute cores while retaining data on premises by leveraging Avere vFXT

¹ Source: <http://blog.cloudera.com/blog/2012/09/what-do-real-life-hadoop-workloads-look-like>

Edge filers. These filers deliver low-latency access by automatically caching active data, storing it in RAM and on SSDs provisioned alongside compute instances in the cloud. Avere provides both read-ahead and write-behind caching. Handling all read, write, and metadata operations near the compute instances, an Edge filer cluster ensures HPC applications run at required levels of performance. Similarly, these virtual appliances serve as a storage gateway to cloud object storage.

6.0 Orchestration and Management Layer for Cloud Workflow Control

The other critical element required to take practical advantage of cloud infrastructure for HPC applications is an orchestration and management layer. Running HPC workflows in the cloud requires a complex process of provisioning the cloud infrastructure, orchestrating workflow execution and job-queue management, automating data placement, monitoring, optimizing, managing security, reporting, auditing, and much more. The orchestration layer eliminates much of the complexity associated with standing up thousands or tens of thousands of compute cores, making cloud infrastructure a practical solution for many more HPC environments and workflows.

Illustrated in Figure 3, this element orchestrates workloads from the user to the cloud with the objective of making cloud HPC as easy to use as on-premises infrastructure. The orchestration/management layer that receives workload tasks from users, spins up cloud compute nodes and optimally distributes the workload among them, and elegantly shuts down processes upon job completion.

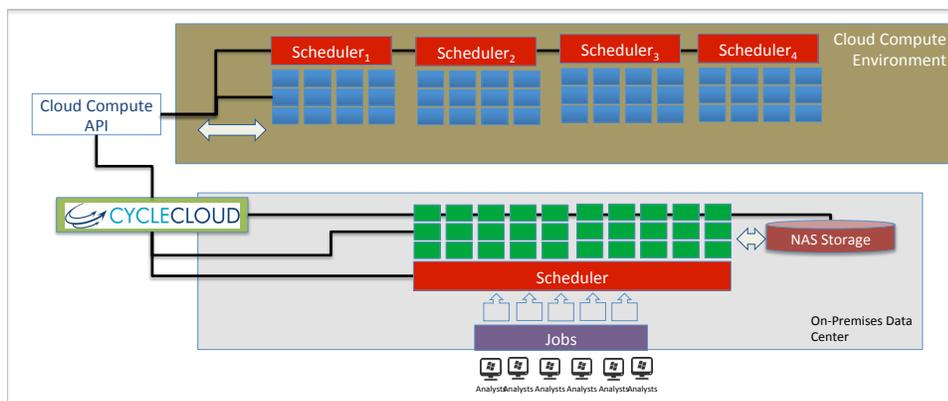


Figure 3: Orchestration and Management Layer

Users benefit from the ability to quickly start and stop tasks, simple workflows, zero queue time, and automatic scaling. System administrators can provide instant access to additional resources, easily link internal systems with resources in multiple clouds, and use reliable tools to enable applications with special requirements. The organization overall benefits from secure and consistent access to cloud resources, the ability to track usage and costs, and flexibility to change and use multiple providers to best meet technical and business requirements.

Multi-cloud workflow solutions like Cycle Computing's CycleCloud solution manages all key workflow elements, including cluster configuration, provisioning, monitoring, and optimization. The elements are handled programmatically to ensure that any options selected can be leveraged across a full range of providers, like AWS, GCP, and Microsoft Azure. Keeping track of the entire workflow, leveraging the feedback loop through monitoring, and providing this functionality across multiple workloads simultaneously makes it possible to expand computational capabilities, rapidly and without major modifications.

7.0 Summary

The cloud undeniably offers unprecedented infrastructure resources and scale, as well as advanced functionality to address pressing needs for capacity, access, new technologies, and simpler data management. Employing hybrid cloud architecture can help IT departments more easily address both technical and business challenges. The availability of cloud-enabling data management and orchestration/management solutions from Avere Systems and Cycle Computing can dramatically reduce the cost, level of effort, and stress associated with delivering best-in-class HPC cloud solutions. These solutions help extend the value of cloud to more HPC environments while protecting flexibility for future needs as well as the sanity of IT teams.

About Avere Systems

Avere is radically changing the economics of data storage. Avere's hybrid cloud solutions give companies—for the first time—the ability to end the rising cost and complexity of data storage and compute via the freedom to store and access files anywhere in the cloud or on premises, without sacrificing the performance, availability, or security of enterprise data. Based in Pittsburgh, Avere is led by veterans and thought leaders in the data storage industry and is backed by investors Lightspeed Venture Partners, Menlo Ventures, Norwest Venture Partners, Tenaya Capital, and Western Digital Capital. For more information, visit www.averesystems.com

Author Scott Jeschonek is Director, Cloud Services at Avere Systems.

About Cycle Computing

Cycle Computing is the leader in Big Compute software to manage simulation, analytics, and Big Data workloads. Cycle turns the Cloud into the innovation engine of your organization by providing simple, managed access to Big Compute. CycleCloud is the enterprise software solution for managing multiple users, running multiple applications, across multiple clouds, enabling users to never wait for compute and solve problems at any scale. Since 2005, Cycle Computing software has empowered customers in Big 5 Life Insurance, Big 10 Pharma, Big 5 Hedge Funds, F500 manufacturing, startups, and Government agencies, to leverage hundreds of millions of hours of cloud based computation annually to accelerate innovation.. For more information, visit www.cyclecomputing.com

Author Rick Friedman leads the Solutions and Marketing group at Cycle Computing.

All trademarks, trade names and service marks referenced herein belong to their respective companies.