



HPC Cloud On-Demand Data Center

Cloud Systems Management Solution

FEATURES

- From a 'single pane of glass' (UI) deploy and manage clusters in multiple CSPs or multiple regions within a single CSP
- Multi-cloud: dynamically expand your on-premise cluster to any Cloud Service Provider
- Shared Clusters: Collaborate and share clusters across multiple users with only one set of cloud credentials
- Spin up an unlimited number of nodes in the same amount of time as it takes to spin up one – 5 minutes
- Container Support: containers such as Docker, Singularity, and Kubernetes are fully supported, though the use of containers is not required
- Easily switch between and move across clouds
- Cloud multi-node jobs are allocated on optimal resources without administrators micromanaging the process
- Manage homogeneous and heterogeneous clusters
- Comprehensive management across the following environments: Public Cloud, Private Cloud, Corporate Cloud, Containers, Virtualized, and Edge
- Enhanced file management and job output
- Auto-Deploy CI/CD Pipelines: CI/CD is improved by enabling automation at any part of the pipeline and can be enabled quickly to handle a new pipeline with ease
- Provide a single point of control for provisioning and deprovisioning infrastructure resources
- Cost-Effective Automation Testing: ODDC allows developers to test on a large variety of high-performance machines and specific environments
- Schedule and orchestrate both HPC and Enterprise workloads; fully integrated with VMware and OpenStack
- Embedded Application: The solution can be embedded into an existing application interface to run the application in the Cloud
- Frequently used job scripts can be saved to quickly reuse when needed for running the same workflow
- Deliver your HPC applications to any private or public cloud as well as from one Cloud Service Provider to another

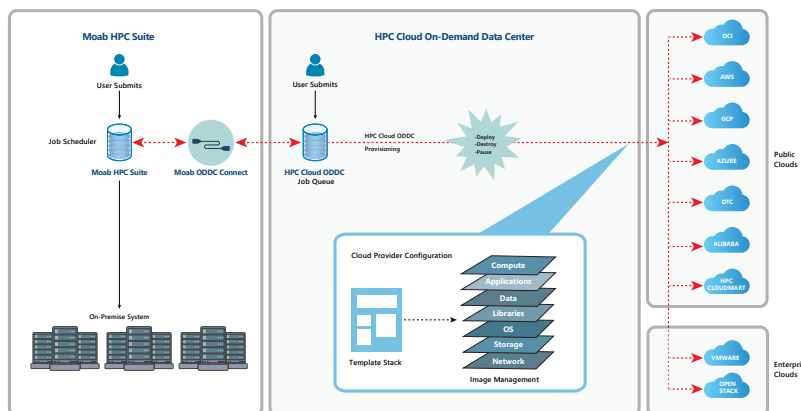
Would you like to spin up an HPC Data Center of any size in the Cloud in under 5 minutes with no up-front cost?

The HPC Cloud On-Demand Data Center (ODDC) solution provides organizations with the ability to rapidly deploy temporary or persistent HPC infrastructure resources, quickly, inexpensively, and on-demand. Access to virtually unlimited HPC compute resources is made available through major Cloud Service Providers, enabling workloads to be run from any geographic location simply by inputting cloud credentials. This enterprise-grade solution automatically deploys and builds clusters in the Cloud, runs applications on those clusters, and then terminates the cloud resources, ensuring that you only pay for what is being used, saving up to 70% of cloud spend. When using the ODDC solution, running workloads in the Cloud costs less than maintaining an on-premise data center, eliminating cost barriers to cloud adoption. Additionally, there will be a power savings of approximately 50%.

ODDC is a scalable cloud systems management solution that empowers organizations to leverage public Cloud Service Provider (CSP) resources without being locked into any single CSP. It provides a simplified way to run HPC workloads in the cloud by serving as an abstraction layer on top of CSP management consoles. Preconfigured Cloud Service Provider access eliminates the complexities of running workloads in the Cloud for users without cloud expertise. OCI, AWS, Google Cloud, Azure, OTC, Alibaba, HPC CloudMart, VMware, and OpenStack are fully integrated and accessible through the user-friendly interface (CLI or GUI) with deployment-ready access. Non-technical end-users can easily access cloud-based HPC resources for hosting compute-intensive workflows, such as simulation models, without needing to navigate their way around any of the Cloud Service Providers.

Stacks and Deployment Images

The HPC Cloud On-Demand Data Center allows users to define stack components, enabling applications to be run in the Cloud. ODDC takes these definitions and automatically builds the deployment image required to execute the workload in the Cloud. The same job script can be used both on-premise and in the Cloud. By employing scripted components to automate manual processes, HPC capabilities can rapidly scale up (and down) without the need for IT intervention. Users can configure stacks, run test jobs, execute custom jobs, run jobs on any major Cloud Service Provider, and view job output from a single interface.



Technical approach to integrating the customer's on-premise assets with those within the Cloud Service Providers

BENEFITS

- Easy to use, manage, and configure
- Instantly launch or scale up HPC infrastructure
- Prevent Cloud Service Provider vendor lock-in
- Eliminate the costs and lead time of making additional hardware purchases
- Automated infrastructure provisioning: Automatically deploy and build clusters in the cloud, then use automation to take unused capacity off-line saving cloud usage costs
- Highly flexible and customizable
- Get true scalability and elasticity
- Fire off jobs from your laptop or create and manage a cluster right from your iPad...On-demand
- Access advanced computing power, and gain virtually unlimited capacity in the Cloud
- Extend your on-premise resources to the Cloud to meet peak demand, increase capacity, or eliminate backlog
- Solve cloud migration challenges
- Intelligently manage cloud resources so that they can be used cost-effectively and efficiently
- Balance workloads between on-premise and cloud infrastructures; Deliver Hybrid IT by spreading your tech infrastructure across different Cloud Service Providers and on-premise infrastructure
- Flexible pricing and licensing models
- Increase productivity and accelerate time to results by taking advantage of automation
- Users without technical knowledge can create temporary or persistent HPC clusters quickly on any Cloud Service Provider
- Run HPC workloads in the Cloud or increase on-premise compute capacity immediately without an up-front investment on an as-needed basis
- Improve management by providing controls for one-off projects with contractors

Move Your Electricity Consumption to the Cloud

Move your electricity consumption from expensive and scarce locations to abundant locations.

Use Case: No electricity is available for increasing compute capacity at a large Air Force base that includes an airfield.

Issue: When a large Air Force base that includes an airfield needed to upgrade the computing capacity to increase security due to security violations, the base learned there was no additional electricity available on the local power grid.

Solution: The issue was solved by utilizing Adaptive Computing's On-Demand Data Center solution. This involved bursting to multiple cloud data centers located offsite based on location and time of day. This allowed for additional computing capacity to be accessed, taking advantage of the maximum available electricity for compute at the lowest possible cost.

Cloud-based Disaster Recovery

If you experience a network or power outage, system failure, natural disaster, or any other incident resulting in disruption or loss of your physical data center, or if your Cloud Service Provider experiences an outage, the HPC Cloud On-Demand Data Center offers a cost-effective disaster recovery plan to resume normal operations. This cloud-based disaster recovery option provides an affordable alternative to duplicating your on-premise systems in an off-site location. Additionally, significant business costs associated with implementing and maintaining a second data center can be avoided. By eliminating your current disaster recovery plan expense, you can allocate the savings to your cloud usage costs.

With ODDC's robust disaster recovery capability, you can quickly restore affected data, applications, and servers to another infrastructure environment running within one or several Cloud Service Providers. Scripts and images previously built within ODDC can replicate physical infrastructure and workload components to virtual or bare metal servers and other cloud resources, offering failover in the event of a disaster. Utilizing ODDC will streamline initial cloud setup and preparation in the event of a disaster. You can spin up and maintain key applications in a matter of minutes.

Application Deployment and Portability

Adaptive Computing's On-Demand Data Center solution is 'application agnostic' enabling migration of any application to cloud-based resources or portability from one cloud to another. Application portability with multiple cloud providers enables organizations to leverage the scalability, flexibility, and cost efficiency of the cloud while maintaining compatibility with existing HPC workflows and applications.

Limited Compute Capacity or No On-Premises Data Center

When users do not have access to a physical data center but still require access to High-Performance Computing (HPC) infrastructure, they can leverage resources on all four major CSPs with the On-Demand Data Center. The HPC Cloud On-Demand Data Center, using Adaptive Computing's patented 'cloud bursting' technology delivers scalable capacity that can be ramped up or down, depending on demand, to balance meeting end-user needs while controlling costs. It also provides access to specialized processing resources inherent to cloud computing such as large memory, storage, fast networking, and GPUs.

ODDC acts as an operating system for the Cloud. Just like an operating system is an abstraction layer for a hardware platform, ODDC is an abstraction layer for cloud platforms. This simple, yet powerful solution enables seamless access to all compute resources, whether on-premise or in the Cloud.

With the HPC Cloud On-Demand Data Center your organization can access specialized HPC capabilities, when you need them, on-demand, from anywhere in the world, and deployed within 5 minutes.

Adaptive Computing's Moab HPC Suite is included with the HPC Cloud On-Demand Data Center

Moab HPC Suite consists of several high-performance computing products, including the patented Moab HPC Workload Manager. This system serves as a job scheduling and management tool for use on clusters, supercomputers, and grids by global HPC customers. Highly optimized and configurable, Moab can support a wide array of scheduling and fairness policies, dynamic priorities, and extensive reservations. It is widely used across the HPC landscape to optimize the manageability and efficiency of systems ranging in complexity from clusters of a few processors to multi-teraFLOPS supercomputers

Only Moab provides predictive scheduling and supports the breadth of HPC, Big Data, AI, ML, etc., both on-premise and in the Cloud. The HPC Cloud On-Demand Data Center solution is fully integrated with Moab to automate HPC Cloud infrastructure deployment, monitoring, scaling, and shutdown of cloud resources, enabling true 'hybrid cloud'. The combination of the Moab HPC Suite solution and the HPC Cloud On-Demand Data Center solution provides monitoring and management of all on-premise and cloud HPC resources. Furthermore, the HPC Cloud On-Demand Data Center solution automatically scales cloud resources up and down to meet demand while controlling costs. Additionally, clusters can be automatically resized up or down without bringing down the cluster while deployed.

Top Concerns with HPC in the Cloud

The top concerns with HPC in the Cloud are performance, data, security, and cost. The HPC Cloud On-Demand Data Center addresses all four of these concerns. To achieve the same or better performance in the Cloud as you would have on-premise, the application must be matched to the appropriate bare-metal compute infrastructure, and a 'Data Anywhere Strategy' must be properly applied. When these two elements are correctly addressed, there are no issues with performance.

Performance and Composable Infrastructure: With the HPC Cloud On-Demand Data Center, users can select infrastructure resources in the Cloud on a case-by-case basis to meet specific workload requirements. As long as they are offered as standalone components within the Cloud Service Provider, custom infrastructure elements such as CPUs, GPUs, size of memory, storage, and network type can be selected individually. Choosing these items separately allows for unlimited infrastructure configuration options, ideal for matching cloud resources to the specific requirements of particular workloads. Most major CSPs offer a broad range of compute instances for customization of infrastructure to match the specific workload requirements. All infrastructure elements can be selected through the ODDC interface before running your applications. Studies have shown that by matching the instance type to the workload requirements, there is a 65% reduction in workload time-to-complete using ODDC.

InfiniBand connectivity is supported between all nodes to minimize latency. Fast interconnection network configurations include High Bandwidth, Low Latency Cluster Networking for the highest level of performance for HPC workloads in the Cloud. Accelerate visualization of complex data models and structures using specialized resources such as NVIDIA GPUs. Deploy and scale up heterogeneous clusters and specialized infrastructure when you need it from any geographic location or device connected to the internet.

Data Anywhere Strategy: ODDC utilizes a strategy called "Data Anywhere". This means your input and/or output data can be located anywhere; either on-premise, in one or more Cloud Service Providers, or in a structured or unstructured file system. It does not matter where your data or file systems reside. Your data can be easily managed by using the ODDC scripting language. The 'Data Anywhere' function pulls data from external sources or includes it in the deployment stack.

Scalable and independent storage is available through the Cloud Service Providers. Gain accessibility to high-speed cloud storage and fast parallel file system performance. Share data between compute nodes via NFS, BeeGFS, and other supported file systems. VPN protocols are common for data transfer, however, a leased line or a high-speed data connection can be used to increase data transfer throughput.

Data Staging: Storing and accessing the large amounts of data that some HPC applications generate can become a challenge. However, when using Moab with the HPC Cloud On-Demand Data Center, several data staging use cases are supported such as staging data to and from a shared file system, staging data to or from node storage on a single compute node, and staging data from on-premise to cloud storage and back.

Security: All major Cloud Service Providers have security measures designed to protect cloud-based infrastructure, applications, and data. These measures ensure user and device authentication, data and resource access control, and data privacy protection. When the HPC Cloud On-Demand Data Center deploys your infrastructure resources in the Cloud, all CSP security protections such as visibility and compliance, compute-based security, network protections, and identity security apply. Within the HPC Cloud On-Demand Data Center administrators can set up user accounts allowing for access control and cloud cost control.

Cost Control: Avoid an up-front investment in an expensive hardware purchase and installation by using ODDC instead of expanding your on-premise data center. There is no need for additional floor space, air conditioning, power, or personnel.

The most cost-effective option to access HPC in the Cloud without an up-front investment is to use Adaptive Computing's HPC CloudMart (powered by OCI). This is an all-inclusive pay-as-you-go offering with the Moab HPC Suite and HPC Cloud On-Demand Data Center running in OCI. You will never pay an up-front cost or be required to make a spending commitment as is usually the case with the Tier 1 CSPs. There are also additional discounts based on particular infrastructure requirements.

Automatically deprovisioning cloud resources when no longer needed reduces public cloud usage hours by 50-70%. Teams can automatically deploy and build clusters in the Cloud, automatically run applications on those clusters, and then terminate the cloud resources on a daily, weekly, or even hourly basis. (Automated deployment and release of nodes) By automatically shutting down nodes when not in use or using the ODDC destroy cluster command, you can be assured there are no orphan artifacts left running within any Cloud Service Provider. Use scripted components to automate and reduce manual processes and use automation in ODDC to take unused capacity offline saving cloud usage costs.

Patented Cloud Bursting Technology Achieves True 'Hybrid Cloud'

The HPC Cloud On-Demand Data Center solution emphasizes 'cloud bursting' capabilities. Organizations can achieve a true 'hybrid cloud' experience and expand their on-premise resources by 'bursting' their workload backlog to the Cloud, eliminating long wait times in job queues and providing a better end-user experience.

The HPC Cloud On-Demand Data Center solution automates the expansion of on-premise clusters and dynamically utilizes the scalability of the Cloud by also 'bursting' workload backlog to an external public cloud when resources are not sufficient to accommodate peaks in demand. All required workload resources are automatically deployed as needed. When the workload backlog has completed, the cloud resources are automatically deprovisioned from the cloud provider. Cloud Bursting can be set up to deploy applications dynamically or on-demand for true scalability and elasticity. Be assured there are no remaining instances left running within the Cloud Service Provider, racking up unexpected usage charges.

Manage resource expansion and contraction of bursty workloads utilizing additional resources from private/public clouds or other data centers to integrate new physical or cloud-based assets.

The Pure Cloud Solution

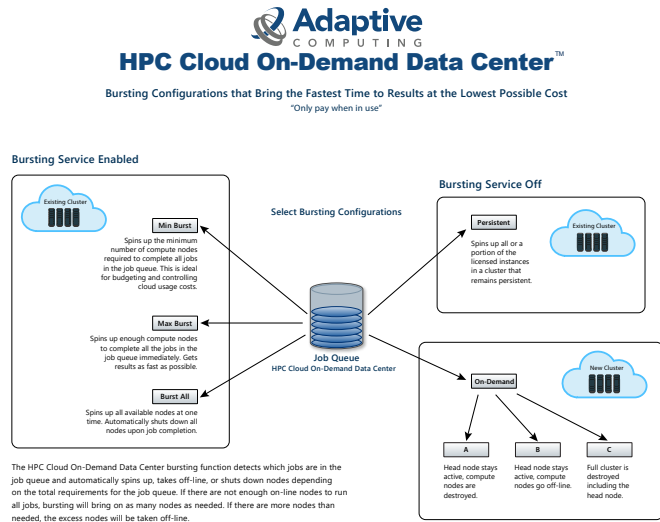
Computation-intensive analysis and simulation can now be performed from any geographic location on a pay-as-you-go basis. When the HPC Cloud On-Demand Data Center solution is used in a cloud-only implementation, there is not any requirement for on-site infrastructure. The user can run applications via a web interface using a laptop or any device connecting to the internet.

No Job Scheduler is Required: The Pure Cloud Installation does not require an HPC scheduler since Adaptive's own open source scheduler Maui is already included with ODDC. However, the HPC Cloud On-Demand Data Center pure cloud installation supports all other HPC job schedulers such as Moab, SLURM, PBS Pro, and LSF.

Use Cases For HPC Pure Cloud: HPC is becoming relevant to a broad range of mainstream businesses that can now access advanced computing without building and maintaining their own HPC Data Centers. Any organization or individual can now run multiple application types and compute-intensive workflows in the Cloud such as AI, Machine Learning, Data Analytics, CAE, EDA, VFX Rendering, Reservoir Simulations, etc. While your models and simulations may run on one Cloud Service Provider initially, the HPC Cloud On-Demand Data Center solution is easily expandable to a multi-cloud environment.

Adaptive Computing Patented Cloud Bursting Technology

U.S. Patent No. 11,134,013 and U.S. Patent No. 11,372,689



Cloud Bursting configurations that bring the fastest time to results at the lowest possible cost

About Adaptive Computing

Adaptive Computing is a global software company headquartered in Naples, Florida, USA and has provided advanced applications and tools to the High-Performance Computing industry for over two decades with hundreds of deployments on the world's largest computing installations. Adaptive Computing products and services are used by organizations of all sizes across a broad range of industries such as High-Tech Manufacturing, Aerospace Engineering, Defense, Universities and Research Labs, Life Sciences, Oil and Gas Exploration, Financial Services, and Data Analytics. Some of the world's largest clusters, grids, and data centers use Adaptive's Moab HPC Suite and Cloud Solutions to maximize performance and value, simplify management, and create a competitive advantage.