

Overview

NODUS Cloud OS™ is an operating system for the cloud. Just like an operating system is an abstraction layer for a hardware platform, NODUS Cloud OS is an abstraction layer for a cloud platform. This highly flexible and intelligent cloud management technology enables seamless access to all compute resources, whether on-premise or in the cloud. Cloud deployment on any of the leading cloud providers becomes easily attainable with NODUS Cloud OS. It simplifies the complexities of running HPC and enterprise workloads and applications in the cloud by including AWS, Google Cloud, Oracle Cloud, Azure, and Open Telekom Cloud within the GUI. Powerful management interfaces (Command line and GUI), with cloud access built in, put you in control. Run jobs on any cloud provider or switch between them. Intelligently manage cloud resources so that they can be used effectively and efficiently.

Expand Your On-premise Resources

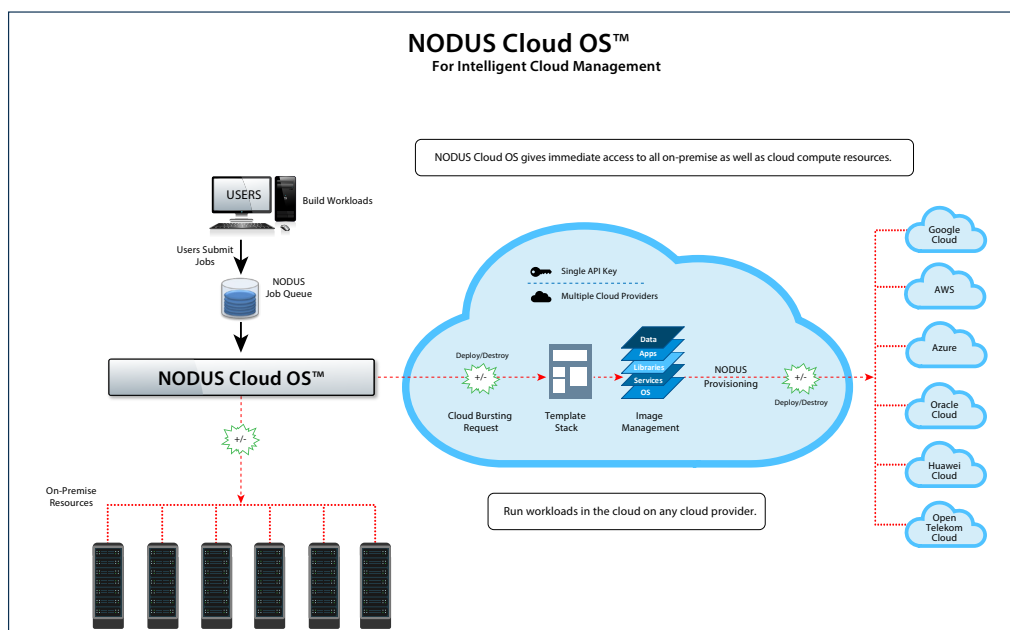
NODUS Cloud OS dynamically expands the capacity of your HPC data center, eliminating the need for an up-front investment in a complex hardware installation. Additional infrastructure is not required.

NODUS Cloud OS supports most popular HPC workload schedulers including the open source scheduler SLURM. Before deployment, there is an option to select SLURM as one of the pre-configured schedulers. NODUS Cloud OS auto builds the SLURM cluster and you do not need to install software.

Make Legacy Applications Portable

NODUS Cloud OS can deliver your applications to the cloud, to on-premise resources or to remote locations. Make applications completely portable between on-premise and cloud infrastructures as well as from one cloud to another. NODUS offers access to computing resources provided through the cloud on an as required basis. NODUS Cloud OS can run without the application owner requiring any knowledge of the cloud itself.

With NODUS Cloud OS a non-technical user can be running applications in the cloud in less than an hour. There is no need for admin/user interaction or cloud expertise. It includes all the necessary tools to facilitate moving workloads and applications to the cloud by “bursting” automatically, based on backlog, or on demand.



NODUS Cloud OS™ Cloud Management System

NODUS Cloud OS™

For Intelligent Cloud Management

Benefits

- Scalability or the immediate availability of resources; instantly launch or scale-up HPC clusters
- Drastically improve the performance of certain workloads by provisioning HPC compute environments and running workloads in the cloud on demand, creating agility and reducing the costs of allocating temporary resources or making additional hardware purchases
- “Burst” the additional workload to an external cloud automatically based on backlog or on demand
- Ease of use for admins
- Seamlessly spin up and spin down on-premise and cloud resources for a hyper-efficient and agile infrastructure strategy
- Run test jobs and run custom jobs on demand
- Can be completely automated to shut down cloud nodes when they are not in use, making it very cost effective
- Run hundreds of simulations with on-demand agility and flexibility
- Increase the capacity of your HPC data center
- Get access to unique specialized resources such as GPUs
- Easy to use, manage, and configure, while integrating seamlessly with on-premise systems; no need to over provision
- Spin up an unlimited number of nodes in the same time as it would take to spin up one
- Offers cost reduction, and unprecedented productivity
- Solve cloud migration challenges
- Move your expenditures from Cap Ex to Op Ex (pay as you go)
- Accelerate time-to-results

Features

- Comprehensive management across the following environments:
 - Virtualized
 - Private Cloud
 - Public Cloud
 - Containers
- Supports Google Cloud, AWS, Oracle Cloud, Azure, and Open Telekom Cloud out of the box from a common provider API
- Automated deployment and release of nodes
- Move across clouds easily and switch between them with the click of your mouse
- Jobs can be run spontaneously; run any workload or application on demand in the cloud
- Run workloads anywhere—in your datacenter, on your laptop, server or in the cloud
- bare metal, VMs, and containers, etc.
- Granular control of resources (whether it is cloud or on-premise)
- Works with any HPC job scheduler or without a workload scheduler
- Cloud Bursting can be automated based on backlog, and can be done in blocks of nodes or by highest priority job
- All required workload resources are automatically deployed as needed, and then retired when the resources are no longer required; instance time-to-live is based on clock or idle time and nodes are dynamically added and removed when not in use
- “Burst” to multiple cloud providers and switch between them
- Supports all public clouds of any size, as well as private cloud infrastructures
- From a single interface the user can configure template stacks (“workflow”), run test jobs, run custom jobs, run jobs on any major cloud provider, and view job output

Corporate Headquarters

1100 5th Avenue South, Suite 201
Naples, FL 34102

Email: info@adaptivecomputing.com
+1 (239) 330-6093

Contact a solutions advisor by phone or email,
or visit our web site today.

www.adaptivecomputing.com

