

Navops Launch

Migrate HPC Workloads to the Cloud



OPEN SOURCE PATH TO THE CLOUD

Navops Launch is based on the open source project Tortuga, a general purpose cluster and cloud management framework with applicability to a broad set of applications including High Performance Computing, Big Data frameworks, Kubernetes and scale-out machine learning / deep learning environments. Tortuga automates the deployment of these clusters in local on-premise, cloud-based and hybrid-cloud configurations through repeatable templates.

Navops Launch is the commercial solution for organizations experiencing increased volumes of high priority workloads who need the confidence of consistent, repeatable template-based cloud provisioning.

Automatically configure Univa® Grid Engine® clusters in leading cloud services to harness public and private cloud resources to meet increasing workload demands, while delivering increased operational efficiencies.

Admins or users can select optimum configurations and cloud providers to execute unique workloads. Out-of-the box adapters enable easy use of all leading public and private cloud providers.

KEY FEATURES AND CAPABILITIES

Automate the Use of Cloud

Navops Launch provisions and manages both virtual and bare-metal environments and includes cloud-specific adapters for AWS, Google Cloud, Microsoft Azure, OpenStack and Oracle Cloud Infrastructure with full support for bring-your-own image (BYOI). Multiple cloud providers can be combined for additional capacity and capability to achieve improved business continuity, performance and cost efficiencies. Cloud-based compute nodes are automatically configured to specification by leveraging Puppet for configuration management. Automated orchestration through Puppet also ensures configurations remain in sync as the cluster changes and/or needs evolve.

Cloud On-Demand

Scaling up or down of cloud-based infrastructures are dynamically managed according to workload through use of a flexible policy engine that captures the specifics of cloud use cases for Univa Grid Engine. The built-in policy engine allows users to dynamically create, scale and teardown cloud-based infrastructure in response to changing workload demand. Management, monitoring and accounting of cloud resources is the same as for local servers.

Navops Launch Highlights

- Fully automated cloud deployment and scaling
- Lower total cost of ownership
- Accelerated time-to-results with improved quality
- Enhanced capacity and capability on demand
- Optimal flexibility to ensure maximal reuse
- Experience-based best practices for successful deployments

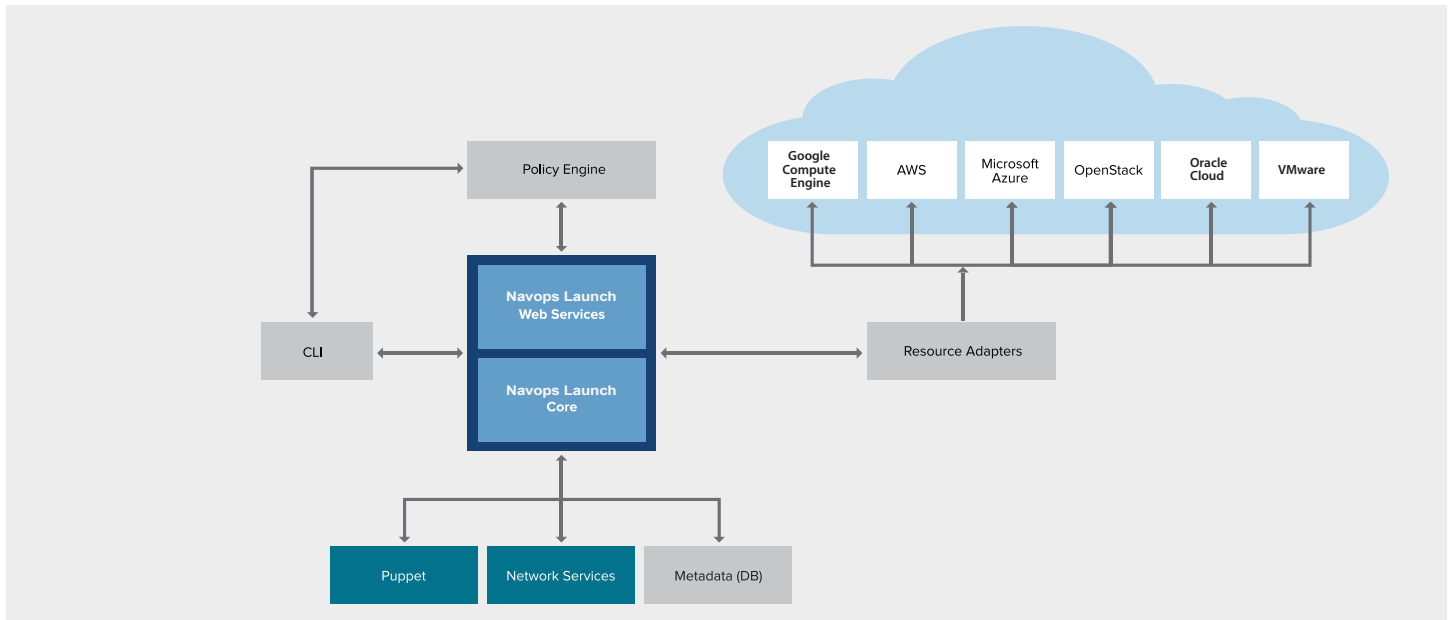
Supported Platforms

- Amazon Elastic Compute Cloud (EC2)
- Google Compute Engine
- Microsoft Azure
- Oracle Cloud
- OpenStack
- VMware vSphere
- On-premise (physical) nodes

Compatible Services

- Name resolution - in-house, Navops Launch, cloud provider
- Identity management - NIS, LDAP, Microsoft Active Directory
- Networking - direct uplinks, VPNs, Amazon VPN/DirectConnect
- Data/storage - cloud provider (e.g., Amazon EFS or S3, Azure Storage) or third-party (e.g., Aspera, Avere, Ceph or Gluster)
- Containerization - Docker
- Big Data - Apache Hadoop, Apache Spark, GPU-based Deep Learning

Navops Launch Component Architecture



Reusable, Interoperable Components

Configuration specifics for hardware, software and services need only be detailed once through template-based profiles that can be easily modified for reuse - includes reuse of customized images, instances and VMs.

Compatible Services

From network services for name resolution and identity management, to use-case specific solutions for managing networks, data and security, Navops Launch's inherent flexibility ensures use of existing compatible services - whether they are provided in-house or by cloud providers.

Complete Visibility

Regardless of cloud use case, or where workloads actually execute, integration with Univa Grid's Engine ensures that all workloads are monitored and reported.

CAPACITY-ON-DEMAND USE CASE

Challenges

- Migrate HPC workloads to the cloud
- Substantially reduce and better manage overall costs and expenditures

Solution

- From basic to highly complex, use-case nuances are addressed with solutions from Univa and its partners

- Employ Univa Grid Engine via Navops Launch to fully automate deployment and scaling using customized software images on cloud-based nodes of choice (including GPUs)

Results

- Numerous 'cloud bursting' scenarios have been deployed in production across multiple enterprise market segments
- Sensibly manage overall expenditures by deferring, reducing, or eliminating on-premise Capex via cloud-based Opex
- Reduce time-to-market and improve quality by intelligently leveraging the cloud during peak usage periods

About Univa

Univa is the leading independent provider of software-defined computing infrastructure and workload orchestration solutions. Univa's intelligent cluster management software increases efficiency while accelerating enterprise migration to hybrid clouds. Millions of compute cores are currently managed by Univa products in industries such as life sciences, manufacturing, oil and gas, transportation and financial services. We help hundreds of companies to manage thousands of applications and run billions of tasks every day. Univa is headquartered in Chicago, with offices in Toronto and Munich. For more information, please visit www.univa.com.



Univa Corporation 2300 North Barrington Road, Suite 400, Hoffman Estates, IL, 60195 USA
Tel: +1.647.478.5901 www.univa.com