

IHME Boosts Productivity

Avoids Cost by Deploying Docker and HPC Workloads on Kubernetes

A leading health research center at University of Washington pioneers mixed container and HPC workloads for their global health registry with Navops Command and Univa Grid Engine



The Institute for Health Metrics and Evaluation (IHME) is an independent population health research center at the University of Washington. With a global reputation for rigorous measurement and analysis, the IHME seeks to understand the world's most challenging health problems and evaluate strategies to address them.

IMHE's Global Health Data Exchange (GHDx) provides a comprehensive catalog of surveys, census data and vital statistics, and in keeping with their core principles, IMHE shares their datasets and findings broadly in a spirit of open collaboration.

Running the infrastructure behind a data science effort of this scale is a formidable task. To support their nearly 300 employees and thousands of collaborating researchers, IHME operates a variety of big data and analytic tools. At the heart of their analytic environment is a shared 500 node high-performance computing (HPC) cluster with 13 petabytes of storage powered by Univa Grid Engine.

The 20,000 core cluster supports a variety of simulation and modeling applications developed using tools including R and Stata.

THE CHALLENGE

To keep pace with evolving applications, IHME needed a modern, shared infrastructure

IHME use virtualization, orchestration tools and sophisticated workload scheduling to ensure that researchers have access to needed hardware and software to run their analytic models.

In recent years, the nature of workloads has been evolving. With the increased popularity of Docker and microservices-based architectures, researchers have increasingly seen value in packaging application logic in Docker containers.

With portable, ready-to-run containers, scientists can evolve models faster, share them more easily, and move seamlessly from development to production without fear of minor environmental differences breaking functionality.

CHALLENGE

- Evolving research workloads
- Need for more infrastructure
- Multiple complex environments

SOLUTION

- Consolidated Kubernetes cluster
- Navops Command
- Univa Grid Engine

RESULT

- Reduced infrastructure cost
- Improved capacity
- Simplified administration

With a Global reputation for rigorous measurement and analysis, the IHME seeks to understand the world's most challenging health problems, and evaluate strategies to address them.

Tyrone Grandison, Chief Information Officer for the IHME, quickly saw the challenge: IHME had the choice of deploying a separate cluster for containerized workloads or finding a better way to manage all workloads on their existing infrastructure.

THE SOLUTION

A single, shared cluster supporting multiple workload types

For containerized workloads, IHME selected Rancher, a Docker container management platform that uses Kubernetes to manage containerized applications. While Kubernetes excels at managing scaled-out container applications, it lacked critical features needed by scientists running existing HPC workloads.

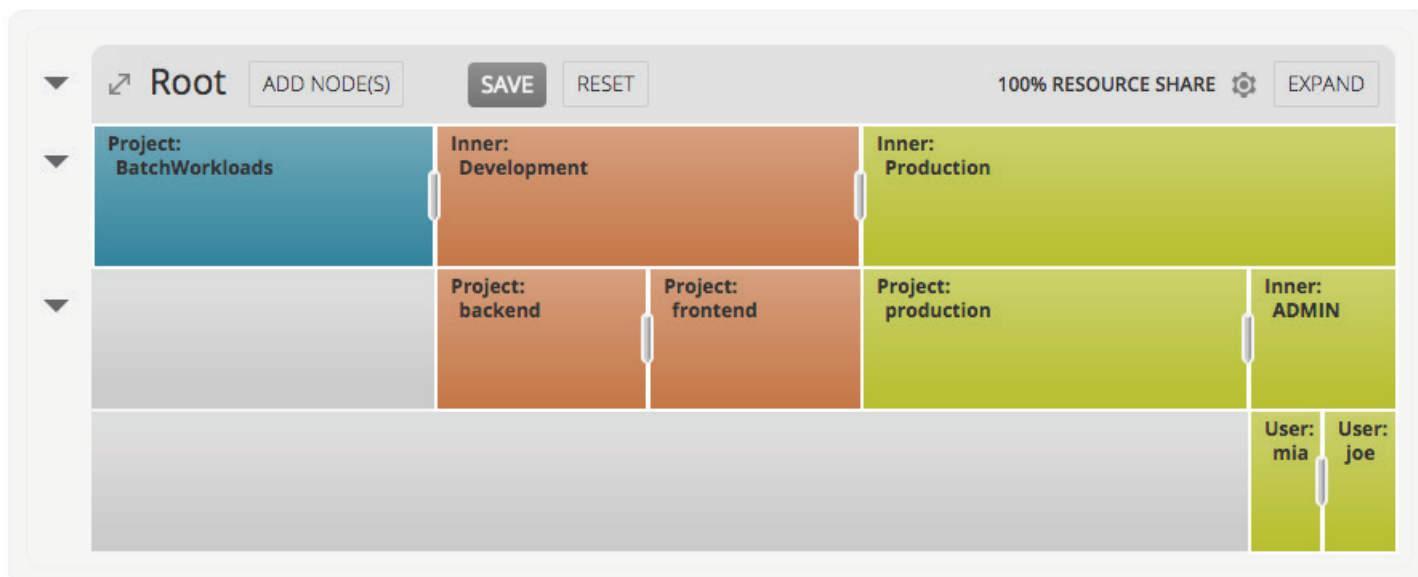
Reliably managing HPC workloads depended on critical Grid Engine features like array jobs, short-job scheduling, fairshare scheduling, and resource based preemption policies.

To solve this, and allow HPC workloads to work seamlessly with Kubernetes, IHME implemented Univa's Navops Command, a solution that plugs into Rancher, extending Kubernetes with advanced workload placement and policy management.

By running Univa Grid Engine in Kubernetes pods under control of Navops Command, IHME realized new efficiencies. They retained all the functionality they relied on for HPC workloads and gained the ability to easily share infrastructure among container and non-container workloads.

“By using Navops Command and Univa Grid Engine in our Kubernetes environment we have the best of both worlds. We can easily share resources between container and non-container workloads and even create hybrid workflows, paving the way to further enhance our application environment at our own pace.”

– Tyrone Grandison – IHME CIO



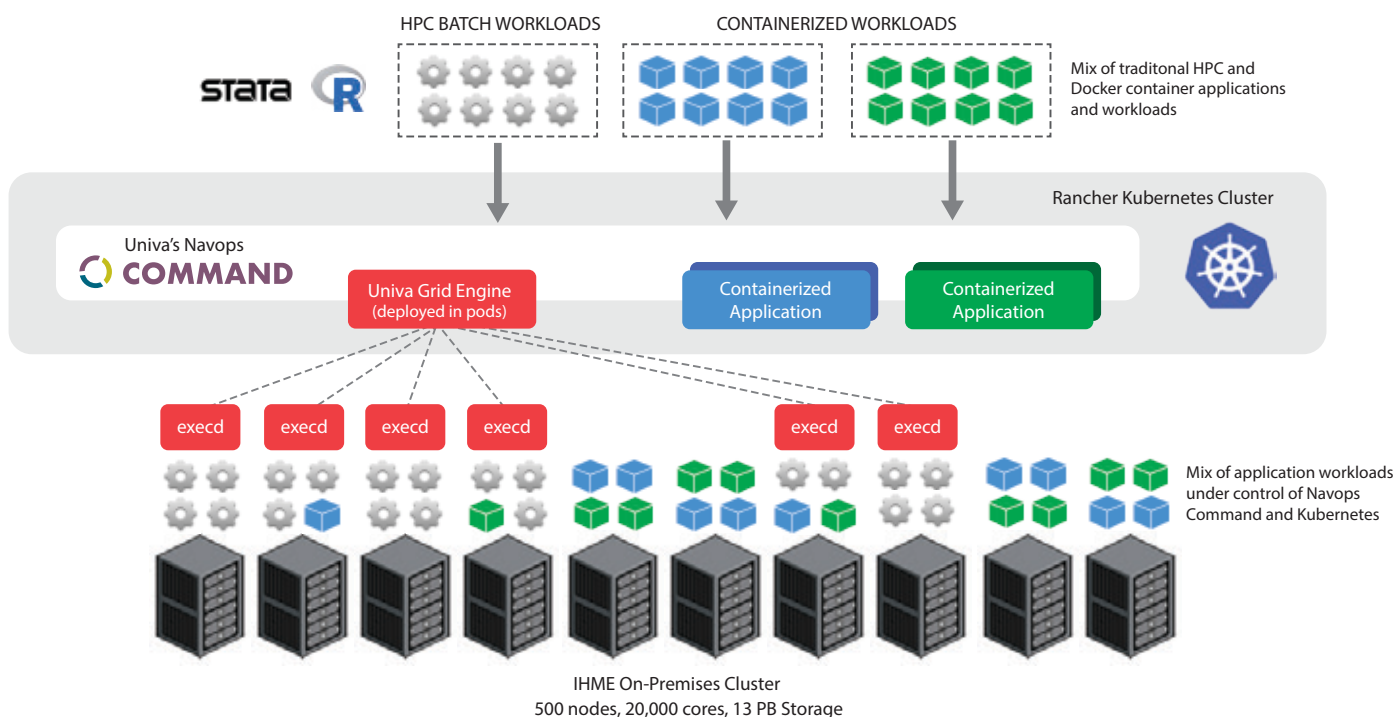
Cluster administrators can adjust resource allocations between container and HPC workloads using an intuitive web interface

BUSINESS RESULTS

With a consolidated environment, IHME avoided significant infrastructure cost, simplified operations and paved the way to continued innovation

By deploying Univa's Navops Command for their Kubernetes environment, IHME realized multiple business benefits:

- **Reduced infrastructure costs** by allowing traditional and microservices workloads to flex dynamically, avoiding the need for separate clusters for batch and containerized workloads
- **Increased flexibility** by providing an easy migration path for applications that cannot be readily containerized
- **Simplified administration and improved efficiencies** by supporting multiple workloads across a single, shared environment



About Navops

Navops is a suite of products that enables enterprises to take full advantage of Kubernetes and provides the ability to quickly and efficiently run containers at scale. Navops utilizes workload placement and advanced policy management across on-premise, cloud, and hybrid infrastructures. With Navops, companies can automate microservices applications and efficiently respond to end user demand. For more information, please visit www.navops.io or follow Twitter [@Navops](https://twitter.com/Navops)



navops.io

Univa Corporation 2300 North Barrington Road, Suite 400, Hoffman Estates, IL, 60195 USA | Tel: +1.647.478.5901