

## Introduction

Containerization offers a lightweight, portable solution for software packaging and execution. Various container orchestration tools are available, which enable the production of computing clusters.

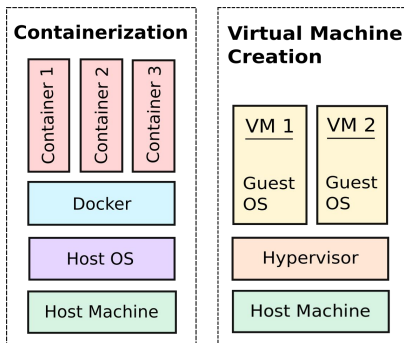
Efforts are ongoing within the ATLAS experiment at CERN to study container applications in particle physics research. In particular, feasibility studies are underway to explore the possibility of using container clusters as ATLAS computing sites.

## Containerization

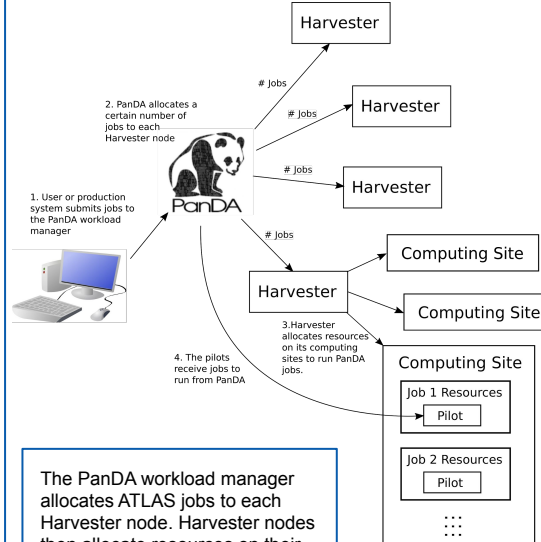
Containers are similar in principle to virtual machines (VMs). The distinction is that containers are much more lightweight than VMs, because rather than virtualizing all the computer hardware, containers only virtualize the operating system. Hardware is shared with the host machine. Because of the minimal virtualization, containers tend to be smaller, more efficient, and faster to boot than VMs.



Docker is a popular tool for containerization.




## ATLAS Workload Management



The PanDA workload manager allocates ATLAS jobs to each Harvester node. Harvester nodes then allocate resources on their computing sites to run the jobs.

## Harvester Integration




Using the k8s API python client, ATLAS developed a method to integrate k8s sites into its workload management. The client receives the number of jobs allocated to the k8s site and creates pods to run the jobs.

## K8s Site Testing

We aim to test and improve the performance of a k8s cluster as a production-level ATLAS computing site. We use terraform to create VMs on Compute Canada's Arbutus cloud, and kubespray to combine the VMs into a k8s cluster. Important features to test and/or improve upon include:

- Site reliability
- Efficient resource utilization
- Site scalability

The cluster currently runs ATLAS jobs reliably with 20 8-core VMs.



## Container Orchestration

Container orchestration tools such as Kubernetes (k8s) make it possible to manage a cluster of containers over multiple computing nodes. The cluster is maintained by one or more "master nodes", which provision and deploy containers on the nodes, monitor cluster health, and handle the addition and removal of cluster nodes. K8s launches pods to create containers in which jobs can be run on cluster nodes. The k8s API enables the creation, management, and destruction of k8s objects such as pods.

## Plans for the Future

We plan to continue adding nodes to the k8s cluster to test its ability to operate at scale. We will soon be adding an additional k8s site on the Arbutus cloud, designed and maintained by UVic Research Computing Services.