

An internal view on Unicorn Orchestrator

Unicorn is a novel framework that helps SMEs and START-UPS to adopt cloud computing for the deployment of services. Developers can focus on developing their services using the suggested development paradigm and let Unicorn Orchestrator handle elasticity, security, resilience and performance across heterogeneous clouds, such as public or private enterprise clouds, or even hybrid clouds. In this blog, we present how Unicorn provides multi-cloud application orchestration and though dynamic allocation of cloud resources in order to help administrators minimise operating expenses, without sacrificing the QoS of their application. Unicorn's policy-aware, multi-cloud orchestrator combines resource management, containerized services deployment over microVMs, and constant policy enforcement for dynamically managing cloud resources.

Unicorn orchestration capabilities are based on resource abstraction and interoperability, thus allowing heterogeneous cloud management and the support for inter-cloud connectivity allow the deployment of applications even in aggregated and hybrid cloud scenarios. In other words, Unicorn policy-aware orchestrator allows the deployment and management of microservice based applications on top of multi-cloud execution environments. In specific, Unicorn's Orchestration mechanism offers:

- Support for multi-cloud deployment of micro-services on OpenStack, AWS and Google Cloud.
- Support for inter-cloud connectivity of micro-service components for the creation of a service mesh in a multi-cloud environment.
- Continuous orchestration and automatic optimization of portable and dynamic cloud services running on virtual instances or micro-execution containers.
- Policy-aware orchestrator that combines resource management, containerized services deployment and management over microVMs, and constant policy enforcement for managing optimized elasticity.
- A configurable elasticity mechanism that allow apps to elastically (de-)allocate resources based on user-defined scaling policies and real-time monitoring and analytics.

This orchestrator is the main outcome of Work Package 3 (WP3) and it is composed of a *Multi-Cloud Resource Manager*, the *Unicorn Cloud Orchestrator*, a *Policy Engine* and finally an *Optimization Manager*.

- *The Multi-Cloud Resource Manager* is responsible for on-boarding new IaaS resources and continuously reporting their availability.
- *The Cloud Orchestrator* is responsible for the creation of the deployment artefacts that are needed for each service defined in the Unicorn compose file and orchestrate the deployment and management of them on IaaS resources. Unicorn relies on Docker containerization technologies in order to bundle stateless services and create the appropriate scalable-by-design microservices.
- *The Policy Engine* is responsible for the enforcement of the policies defined by the developer of a Unicorn microservice.
- *The Optimization Manager* is responsible for suggesting the optimal deployment configuration or re-configuration based on predefined placement constraints.

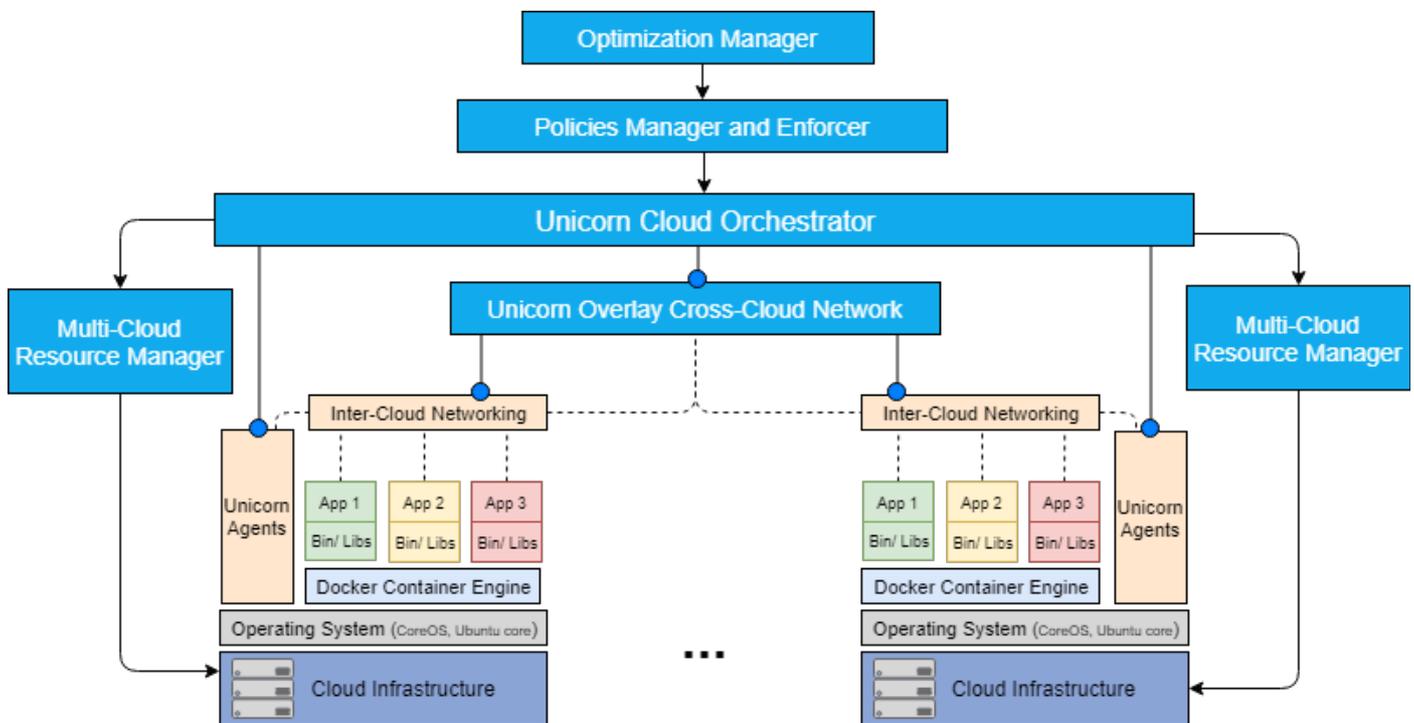


Figure 1: High Level view of Unicorn Orchestrator

Infrastructure in computing refers to resources which can be virtual or physical and include computing, storage and networking resources. In the scope of Unicorn, we consider the usage of virtualized resources that are offered for lease by the infrastructure provider and can be used by the service developers. Moreover, as the focus of the project is towards the deployment of scalable services towards a multi-cloud environment, we need to invest on technologies that can enable the easy, fast and secure deployment over virtualized infrastructure. To achieve this, Unicorn is suggesting a combination of containerized applications paradigm and usage of virtualized multi-cloud infrastructure, and as depicted in the figure below, we enrich this combination with a set of agents that are used by the orchestrator.

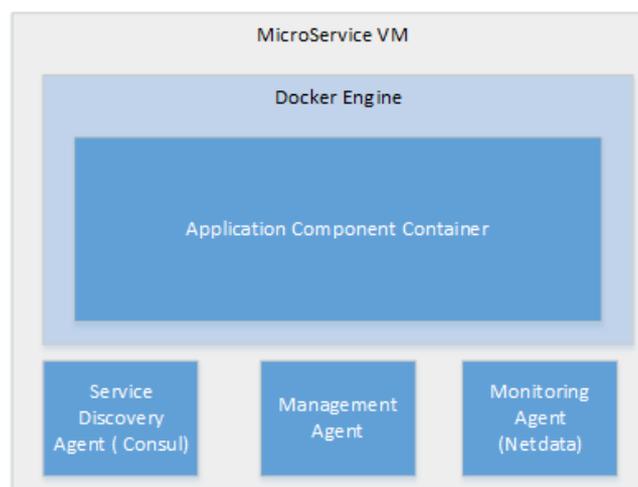


Figure 2: Representation of a Microservice VM as part of Unicorn

With Unicorn Orchestrator, multi-cloud deployment and orchestration capabilities are provided in order to support the vision of Unicorn for containerized execution environments that addresses data portability and interoperability issues. In addition, cloud applications should be highly scalable and adaptable during runtime. Therefore, it should be possible to retrieve real-time measurements of resource consumption, and part of our work done in the project is to

provide the means to realize cloud application monitoring, auto-scaling and management in multi-cloud execution environments. The way we achieve this in Unicorn is depicted in the figure below.

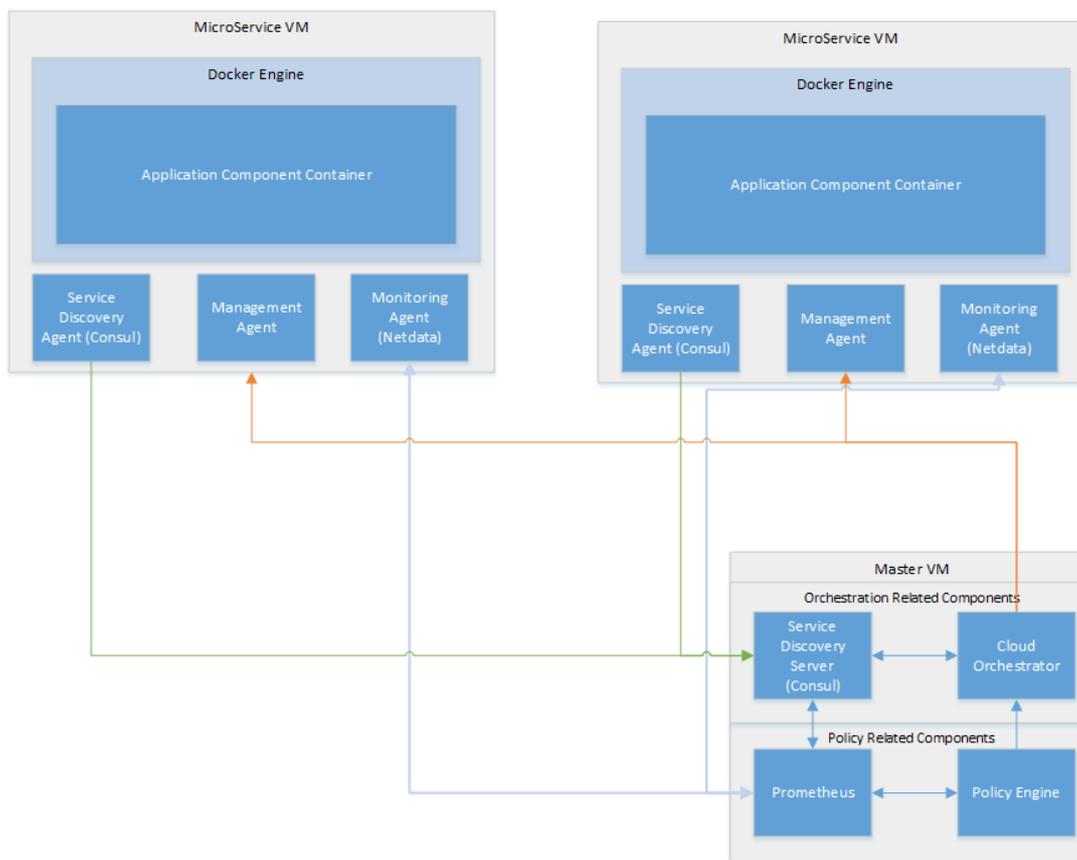


Figure 3:Representation of detailed architecture of Policy Engine (Master Node and micro-service node)

The *Cloud Orchestrator* and the *Multi-Cloud Resource Manager* are actually used to orchestrate the deployment and management of deployment artefacts on multiple IaaS resources. The deployment artefacts are described in detail as they are needed for each service defined in the Unicorn compose file of each application. When an application is deployed, the service discovery agent of each microservice VM is registering itself to the Service Discovery Server and the Cloud Orchestrator is then using dedicated agents to talk and manage each VM.

The *Policy Engine* and the *Optimization Manager* allow the constant scaling and the optimal deployment configuration of the microservice based on predefined placement constraints and the facts provided by the monitoring mechanism of Unicorn. This monitoring information is retrieved using Prometheus monitoring system that collects the data from the monitoring agents.



UNICORN

UNICORN project has received funding from the European Union’s Horizon 2020 Research and innovation programme under grant agreement No 731846.

Copyright 2016 © All Rights Reserved