

Microservices and DevOps

DevOps and Container Technology

Container Orchestration

Henrik Bærbak Christensen

From One to Many

- Container technologies like Docker moves the start-up time of ‘a server’ from hours to seconds
 - Consider installing separate machines for SkyCave
 - Install Ubuntu, Redis, start
 - Install Ubuntu, Web framework, Quote Service, start
 - Install Ubuntu, java, gradle, skycave daemon, start
- However, we still do quite a few manual steps as we focus on each individual container
 - Run Redis container, run quote service, manage network, run daemon
- And worse – it was still tied to **one piece of hardware**
 - Which sort of defeats the whole purpose of scaling...

From One to Many

- Orchestration tools

While the CLI meets the needs of managing one container on one host, it falls short when it comes to managing multiple containers deployed on multiple hosts. To go beyond the management of individual containers, we must turn to orchestration tools.

[Janakiram, 2016]

- That is, a ‘dashboard’ of sets of servers/hardware that allows automatic distribution and scaling of containers...
- *Newman: deployment interface*
 - *One command line that deploys the architecture*

Central Concepts

- Provided Functionality/Tooling
 - Automate all aspects of deployment and monitoring
 - Placement of service on 'convenient' server
 - Deployment and starting services
 - State-state activities: update, roll-back, health monitoring
- *Infrastructure-as-code* Declarative configuration
 - State *what you want*, not *how to do it*
 - The hierarchical 'YAML' format predominant
 - YAML = YAML Ain't Markup Language 😊

Central Concept

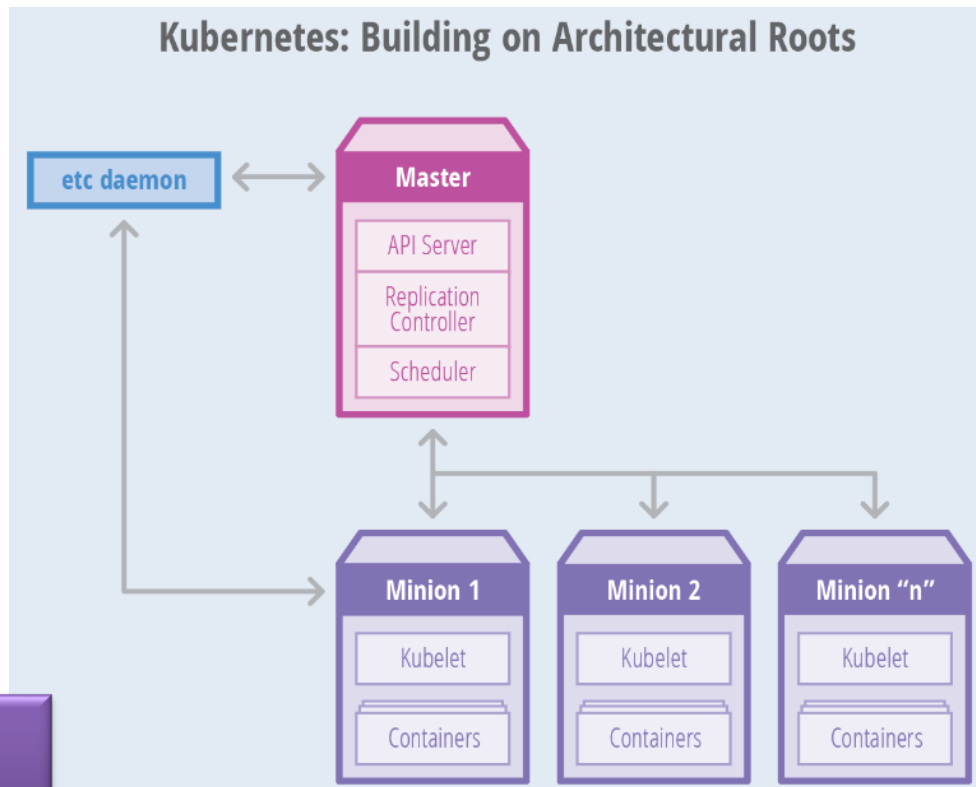
- Rules and Constraints
 - Service placements can be controlled by rules/constraints
 - "put mongodb on a machine with large hard disks"
 - "put master and slave on different hardware"
 - "make three replicas of this service"
- Provisioning
 - Distribute services 'efficiently' across available hardware
- Discovery
 - Services must interact, so some form of DNS is required internally in the cluster to facilitate that

Central Concepts

- Health Monitoring
 - Track and monitor health of services, containers, and hosts
 - Relocate containers from failing hosts
 - Relaunch replacement container if any crashes
 - That is, actively striving to provide the configuration set forth by the *infrastructure-as-code* declarative configuration...

Candidates

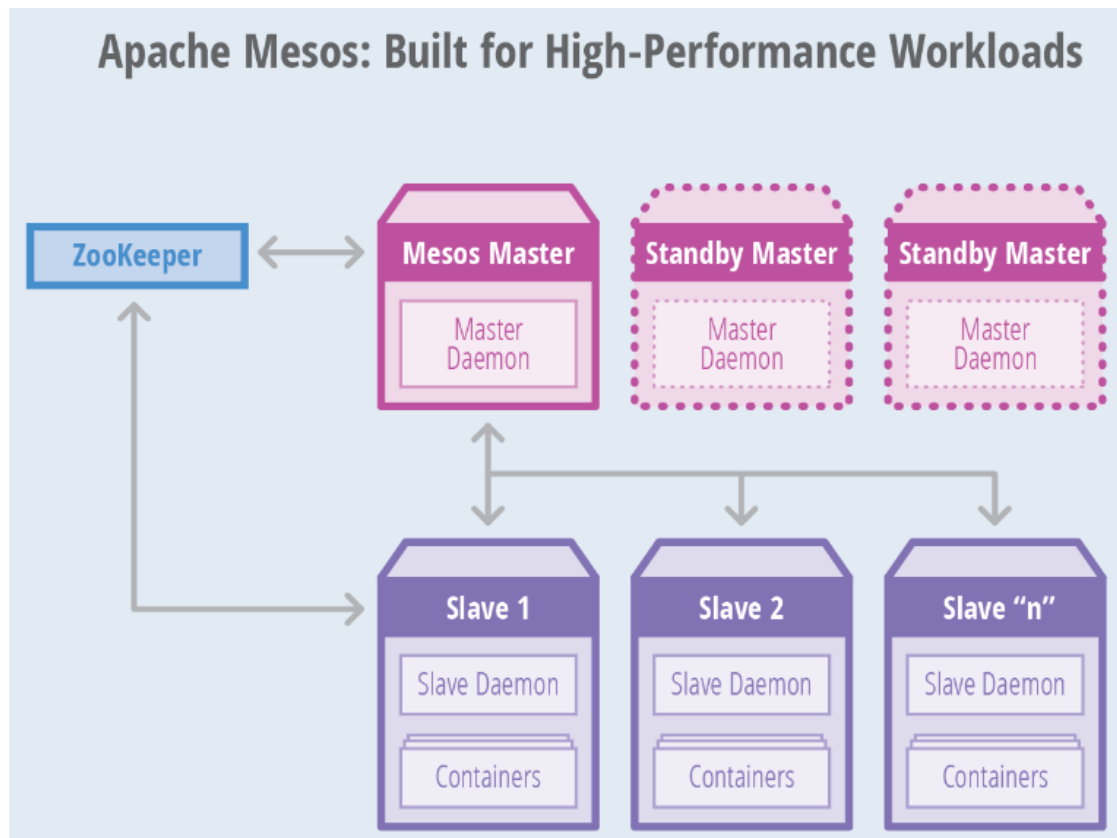
- Kubernetes
 - Google
 - Oldest, proven in war 😊
 - Open source
 - Steep learning curve ☹



Review *'Docker Swarm vs. Kubernetes'* by Pedersen & Gribenco, on <https://baerbak.cs.au.dk/c/ProjectReports/>

- Apache Mesos
 - High availability focus...

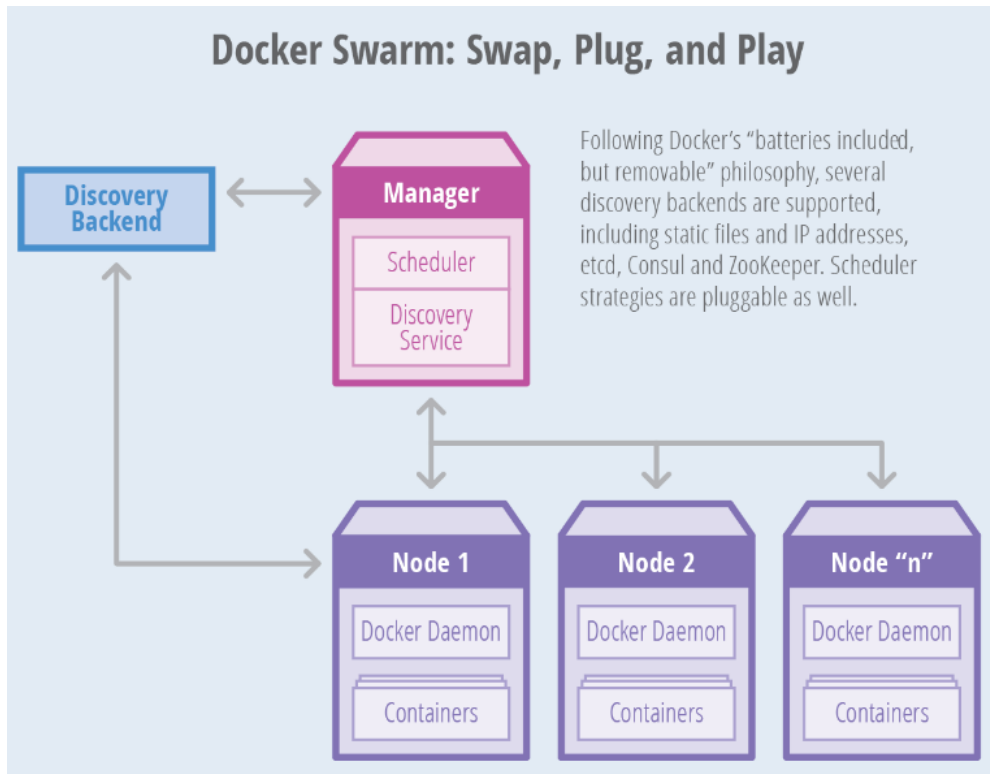
Candidates



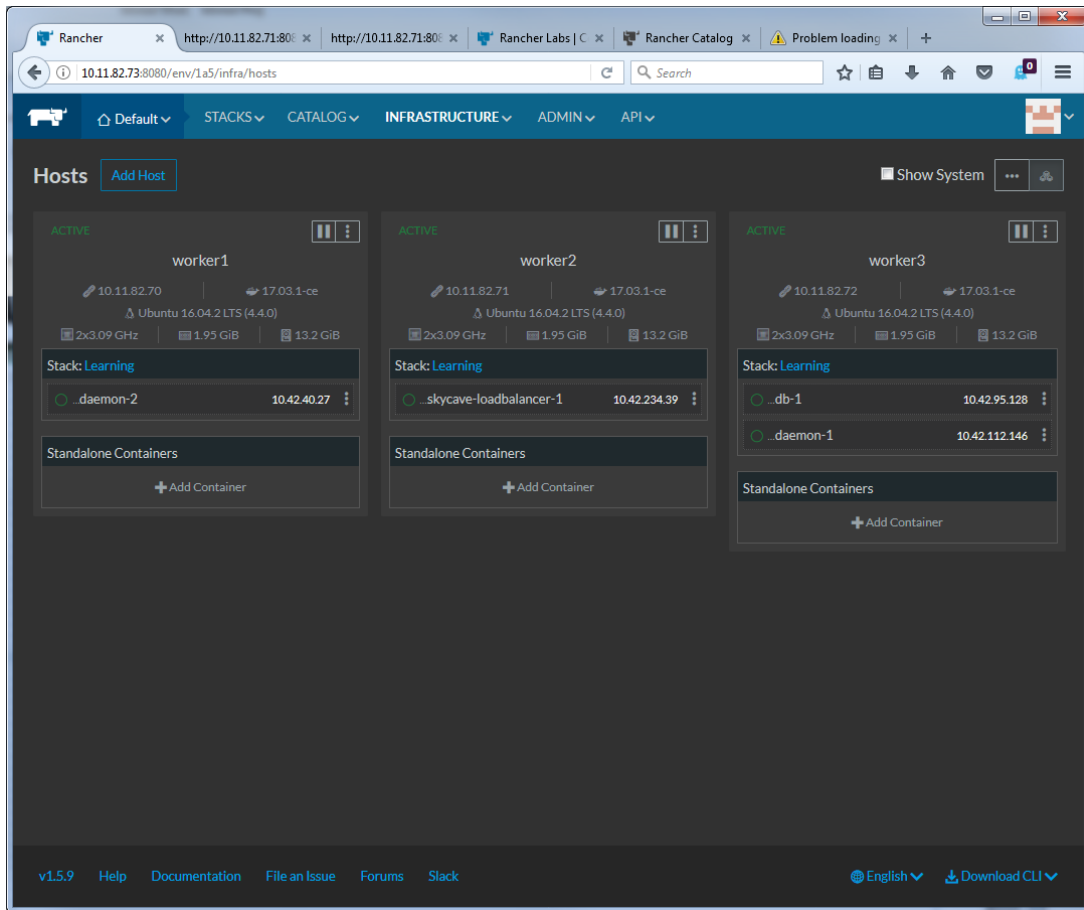
Candidates

- Swarm

- Highly integrated with Docker
- *Infrastructure-as-code*
- Less steep learning curve
 - Claim: Kubernetes compliant
- Caveat:
 - Rumors that it 'does not work' in practice?
 - It does!



- Rancher
 - A docker container having a Web gui
 - ‘click-deploy’ your containers, add build in load balancer
 - And off you go...



Our Choice

- We go for Swarm because
 - Part of the Docker infrastructure
 - Shallower learning curve
 - Very easy to deploy locally on your own machine
 - No hazzle to create a real, costly, cluster
 - Downside
 - Kubernetes seems to be winner that takes all...