

Docker Swarm



Alternatives:

- Kubernetes
- Mesosphere
- Apache Mesos
- ...

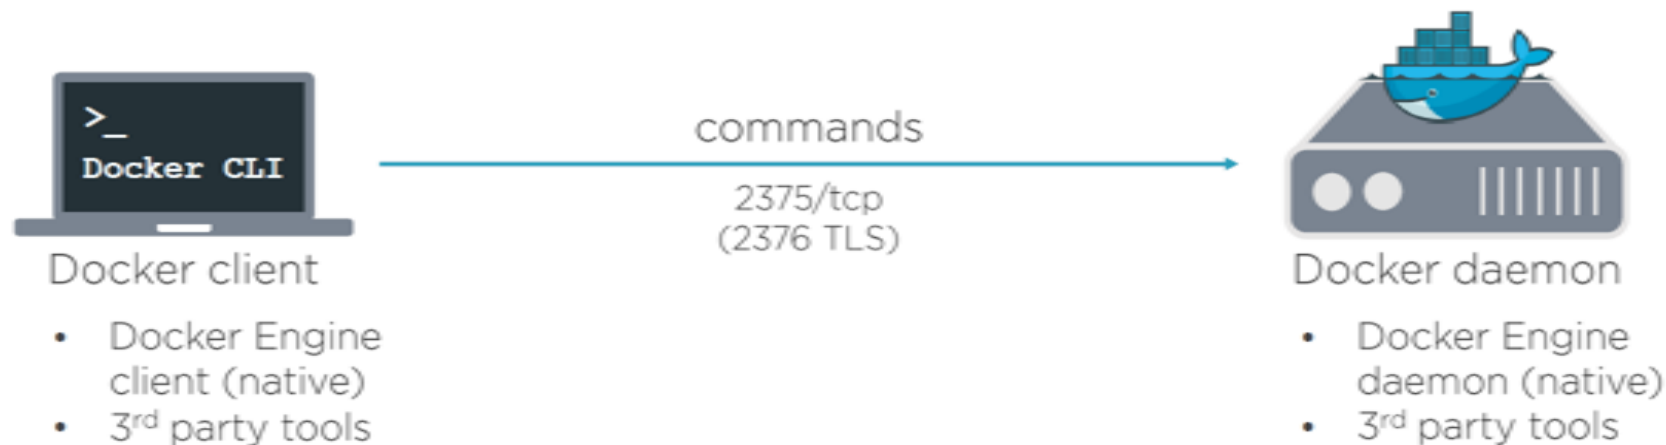


Swarm:


- Native (Docker, Inc.)
- Clusters many Docker hosts into a single large pool
- Simpler management at scale

Basic Docker Architecture:



Client-server





Scaling Challenges




Scaling Challenges




Scaling Challenges



 On prem  Cloud **P** Prod **S** Staging  Fast storage



Scaling Challenges



Scaling Challenges



Architecture & Terminology

Discovery service



Swarm manager



Discovery service

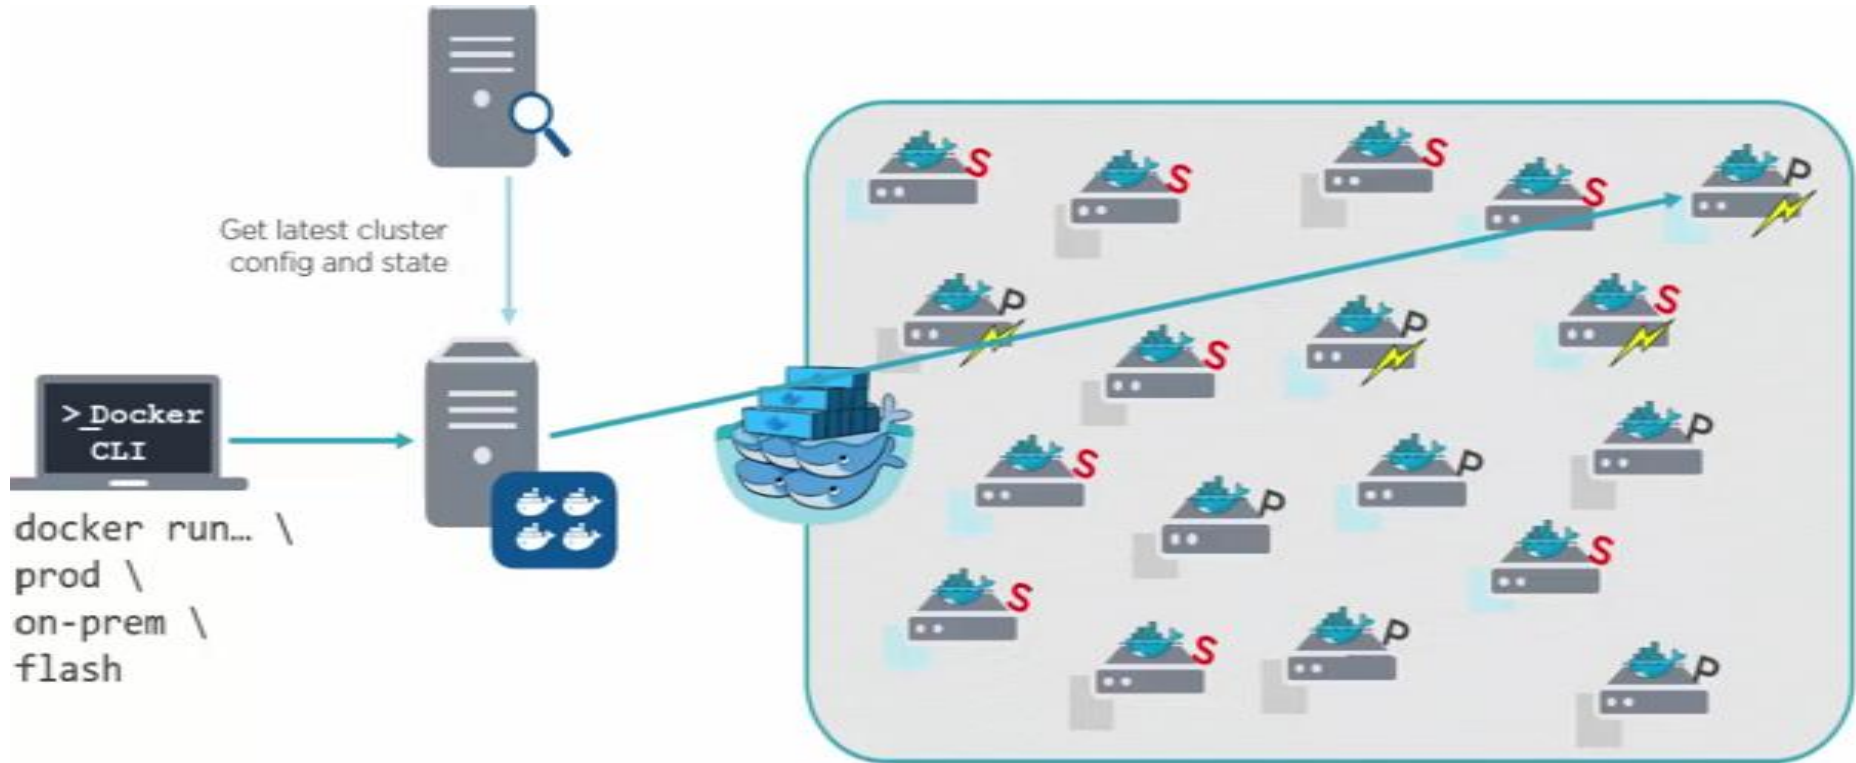


- Key-value store
- Stores cluster state and config
- Needed for the cluster to operate properly

Swarm manager



- Cluster admin
- Accepts Docker commands
- Executes commands on the cluster



Swarm Discovery Service




Supports a pluggable backend



Swarm Discovery Service



Supports a pluggable backend



Filtering & Scheduling

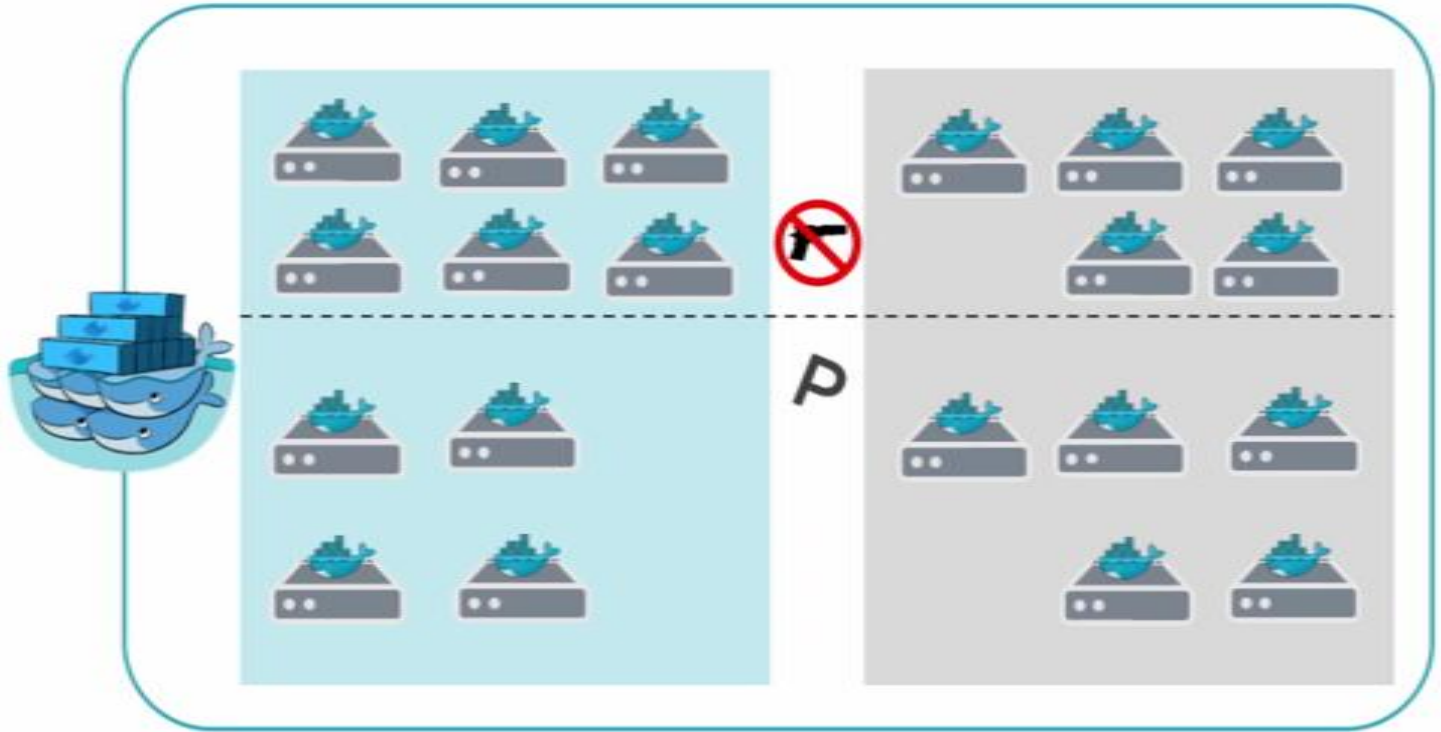
Filtering & Scheduling





Stage 1

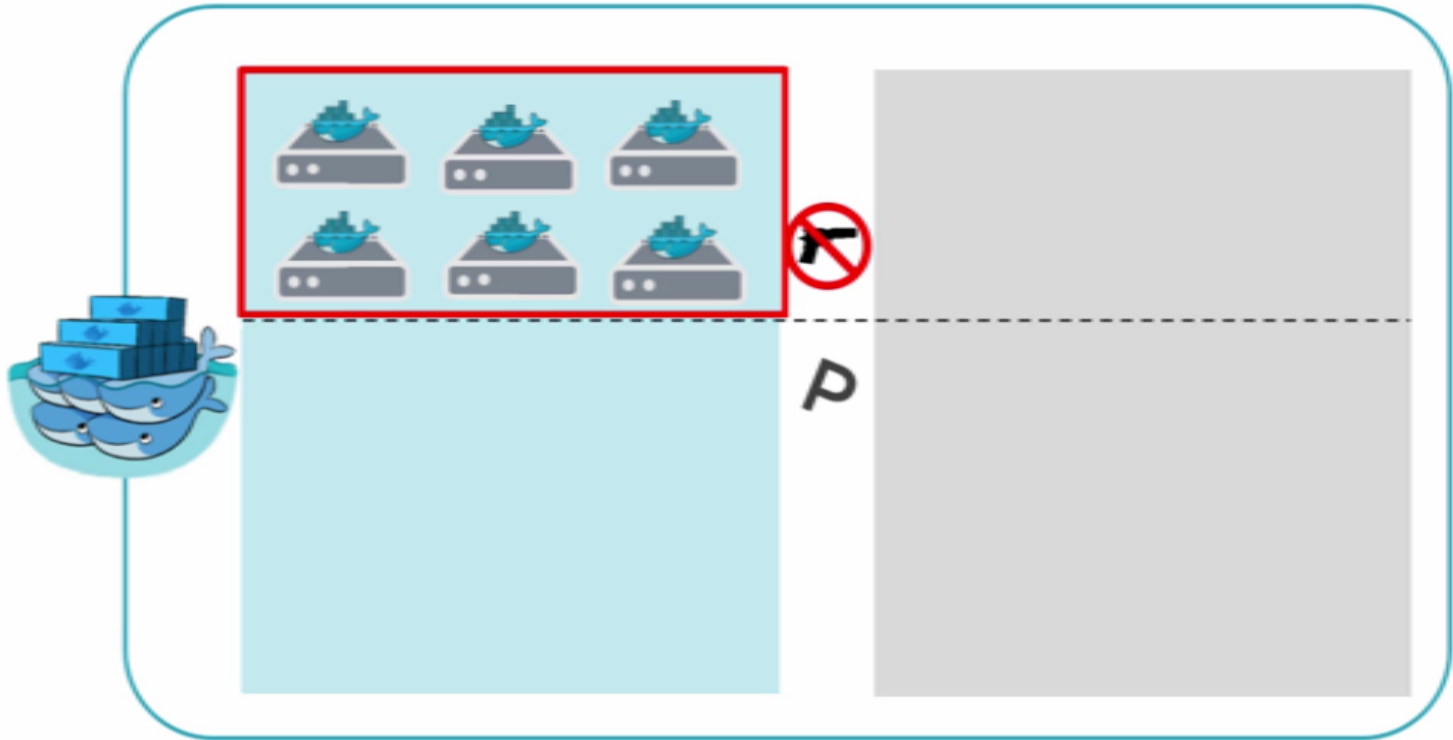


Stage 2









Filtering

Affinity

Run new containers on the same nodes as existing containers or images

Constraint

Standard: Data returned from **docker info** command

Custom: Labels assigned to Engine daemons

Resource

Run new containers on nodes with particular resources free



Scheduling

Random

Spread

Binpack

Scheduling Strategies


Binpack



Stopped containers are considered

Starts on the smallest node
(reserves larger nodes for large containers)

Additional nodes only used if you
specify CPU or RAM requirements



Spread



Stopped containers are considered

Attempts to balance containers evenly
(unless CPU or RAM constraints override)

