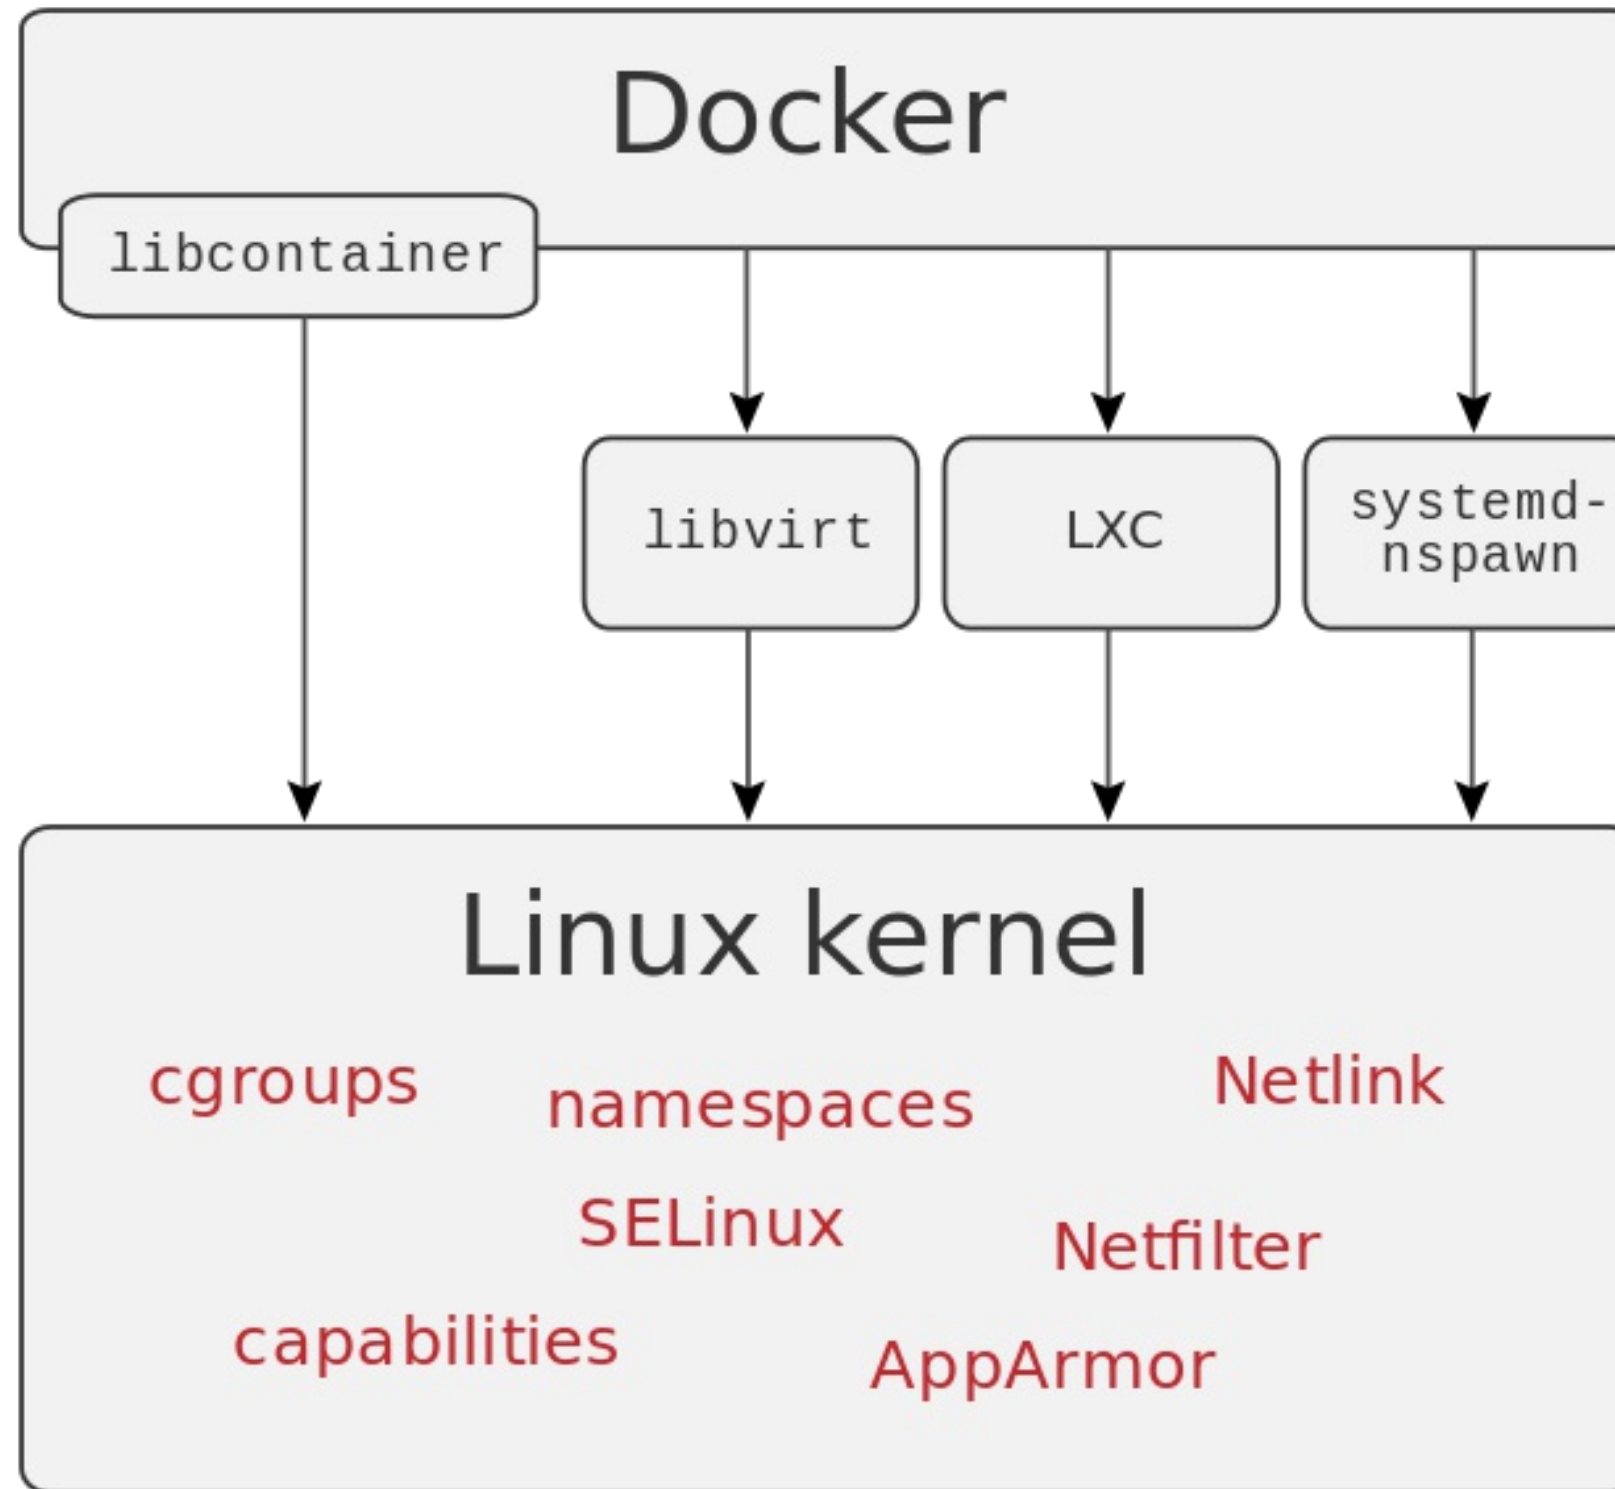


KUBERNETES @ CCIN2P3

A SHIPLoad OF CONTAINERS

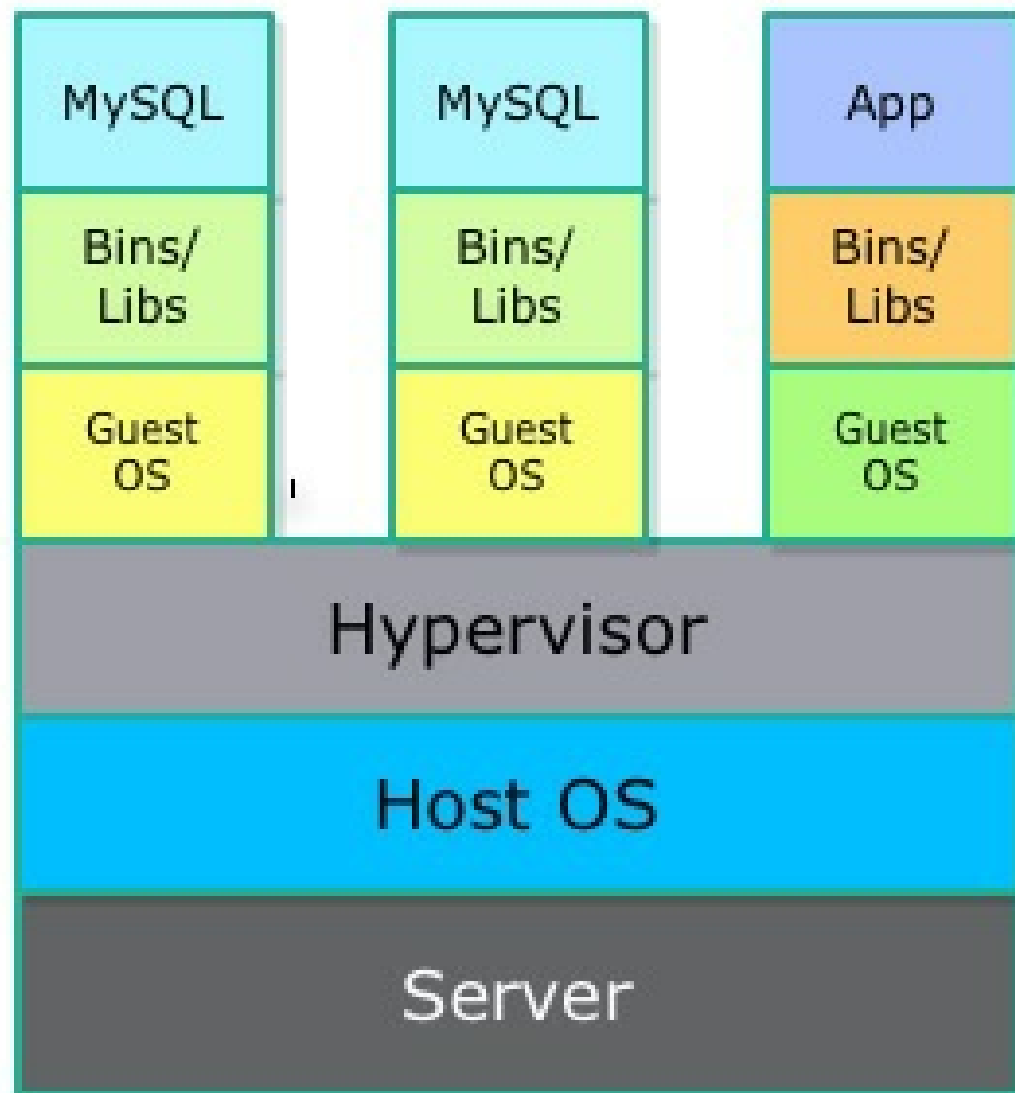
FJPPL - 2018.02.13

- Containers
 - Docker
 - Kubernetes
- Use cases at CCIN2P3
 - Web hosting
 - CI/CD
- Infrastructure
 - Systems and storage
 - Deployment
 - Monitoring

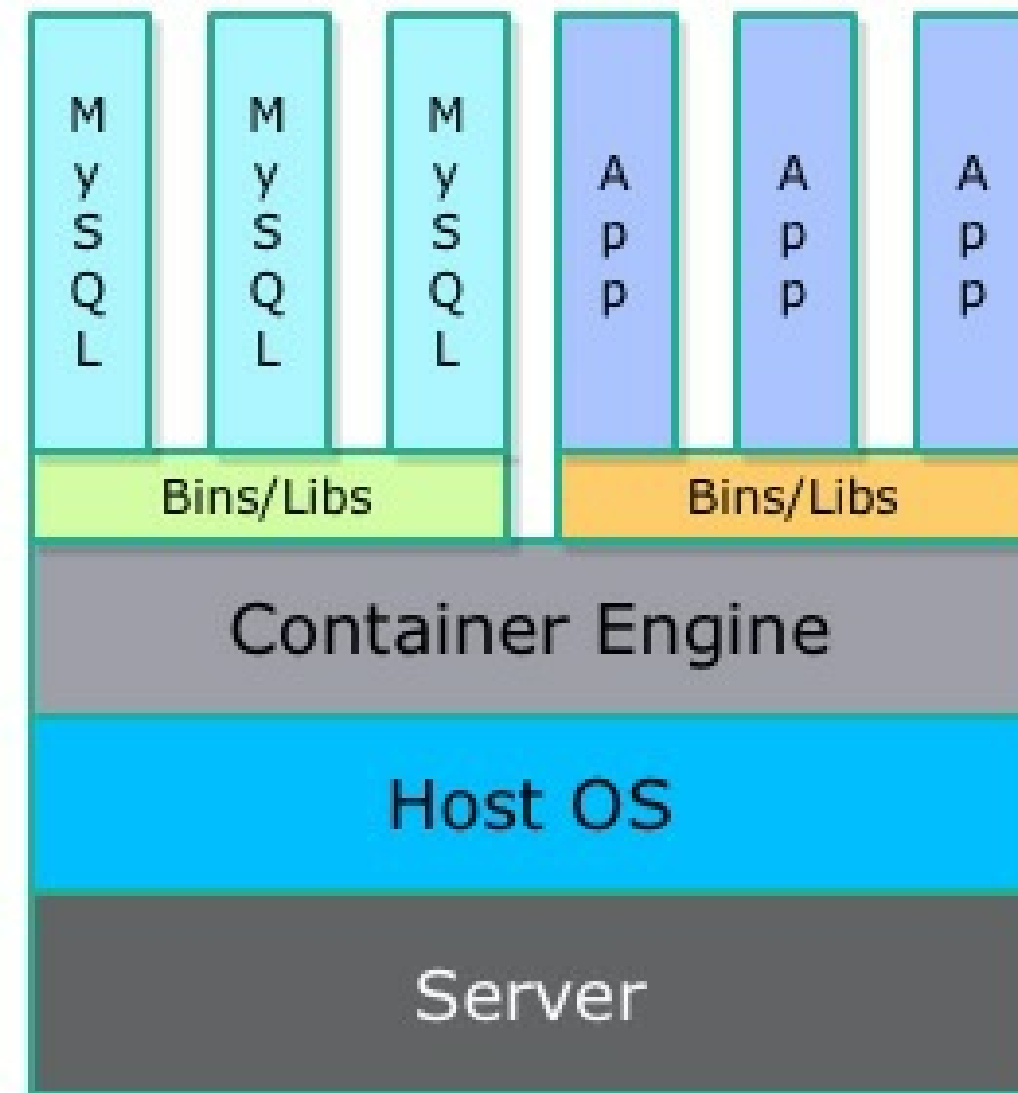


Containers vs. VMs

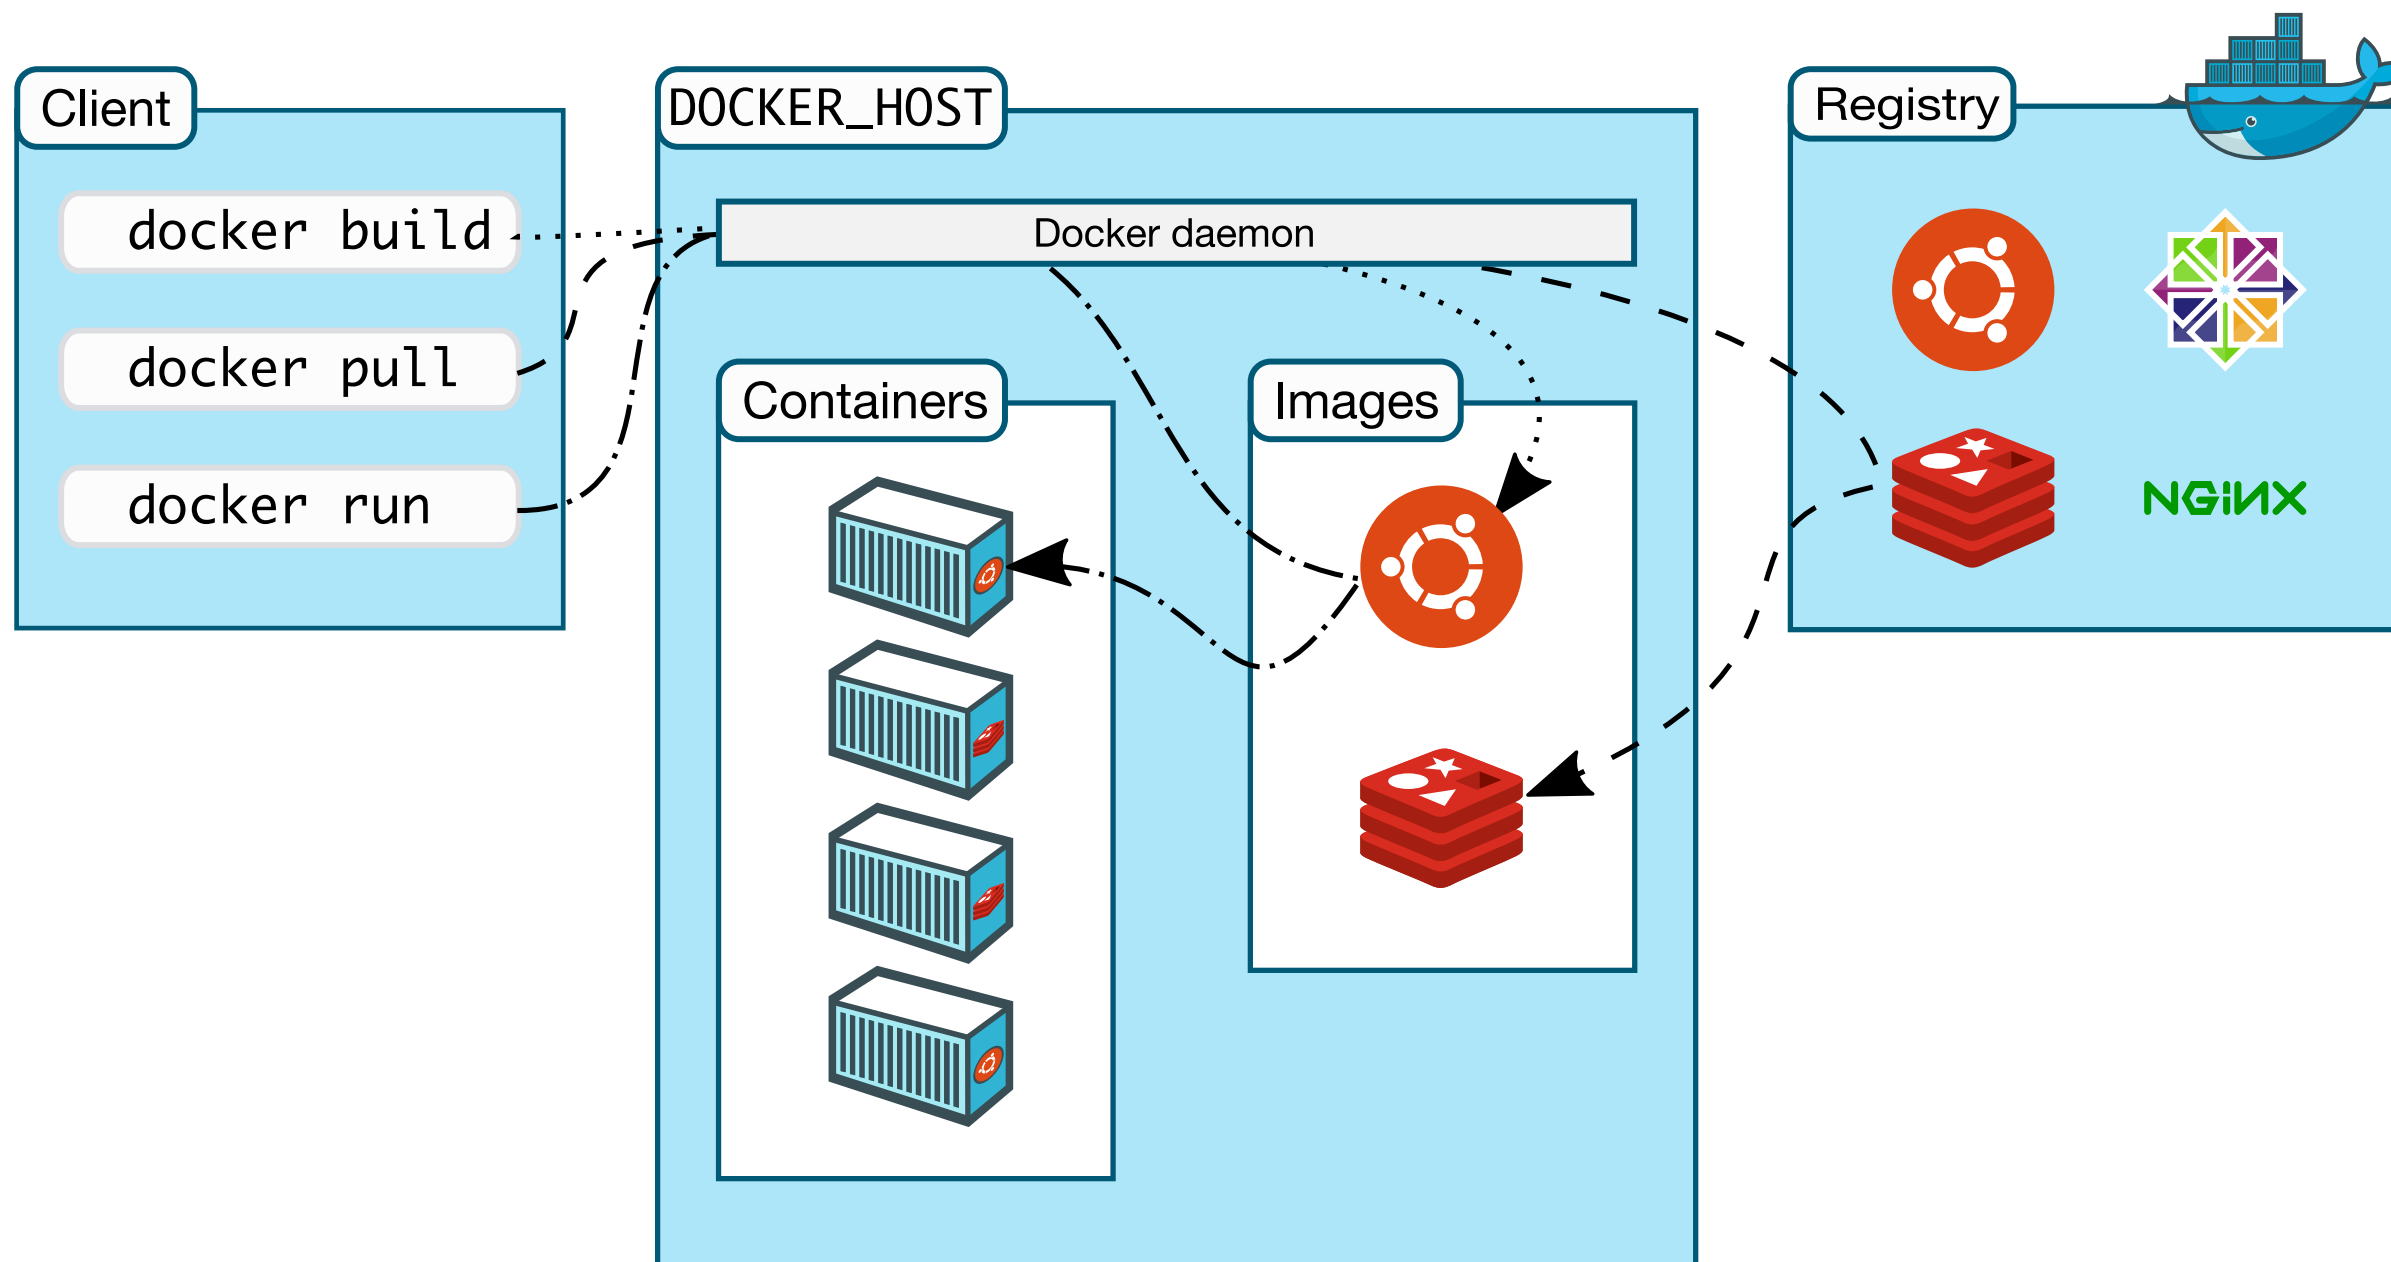
Virtual Machines




Containers



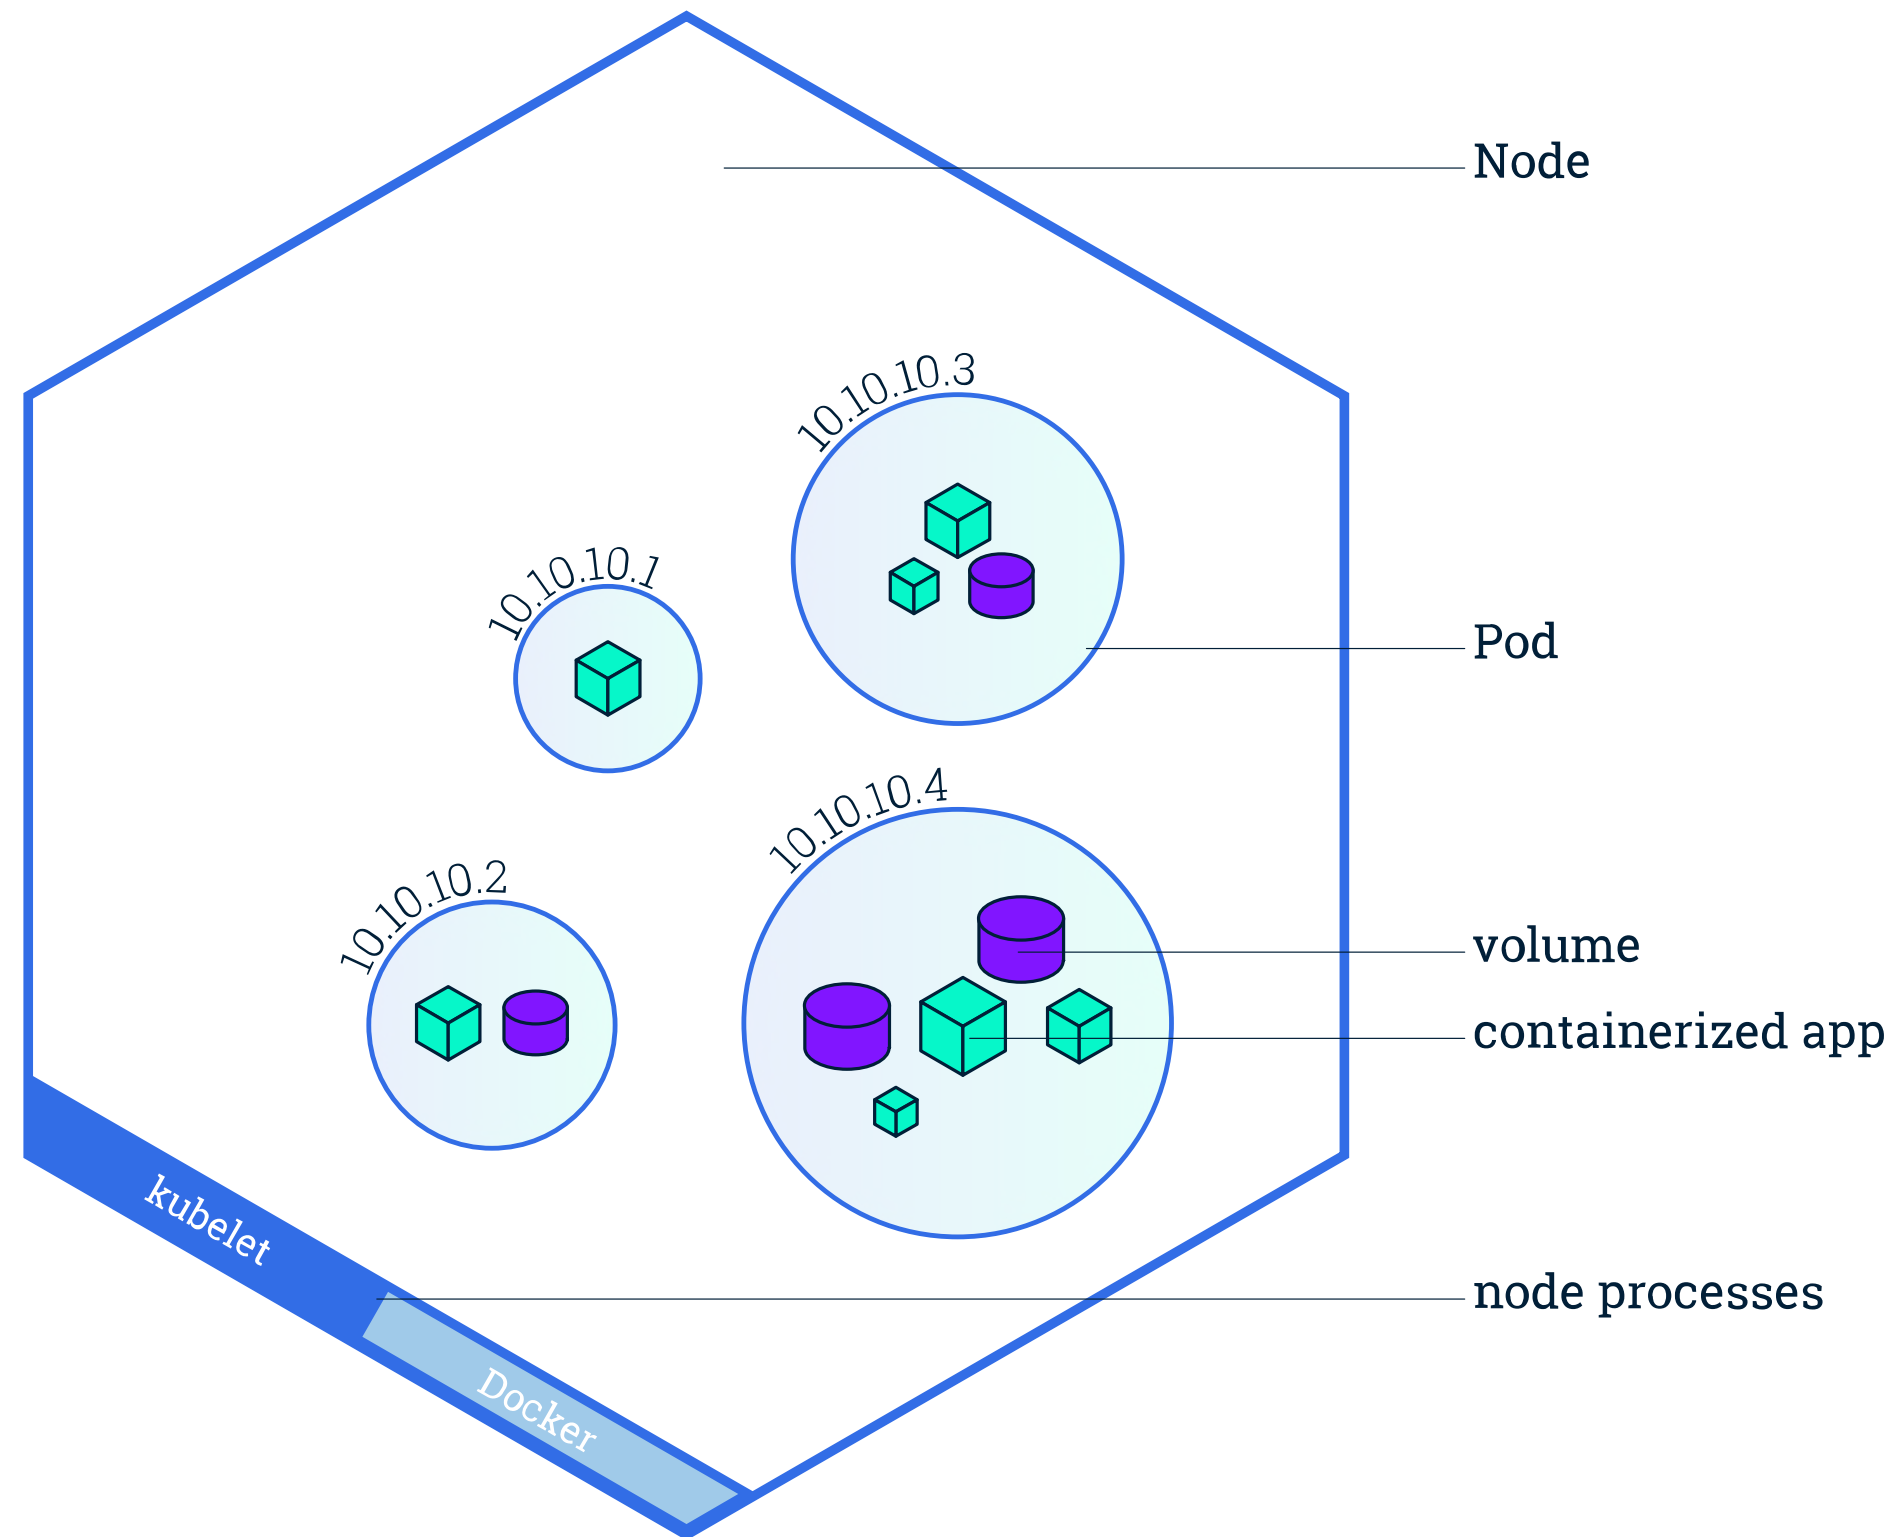
The Docker workflow



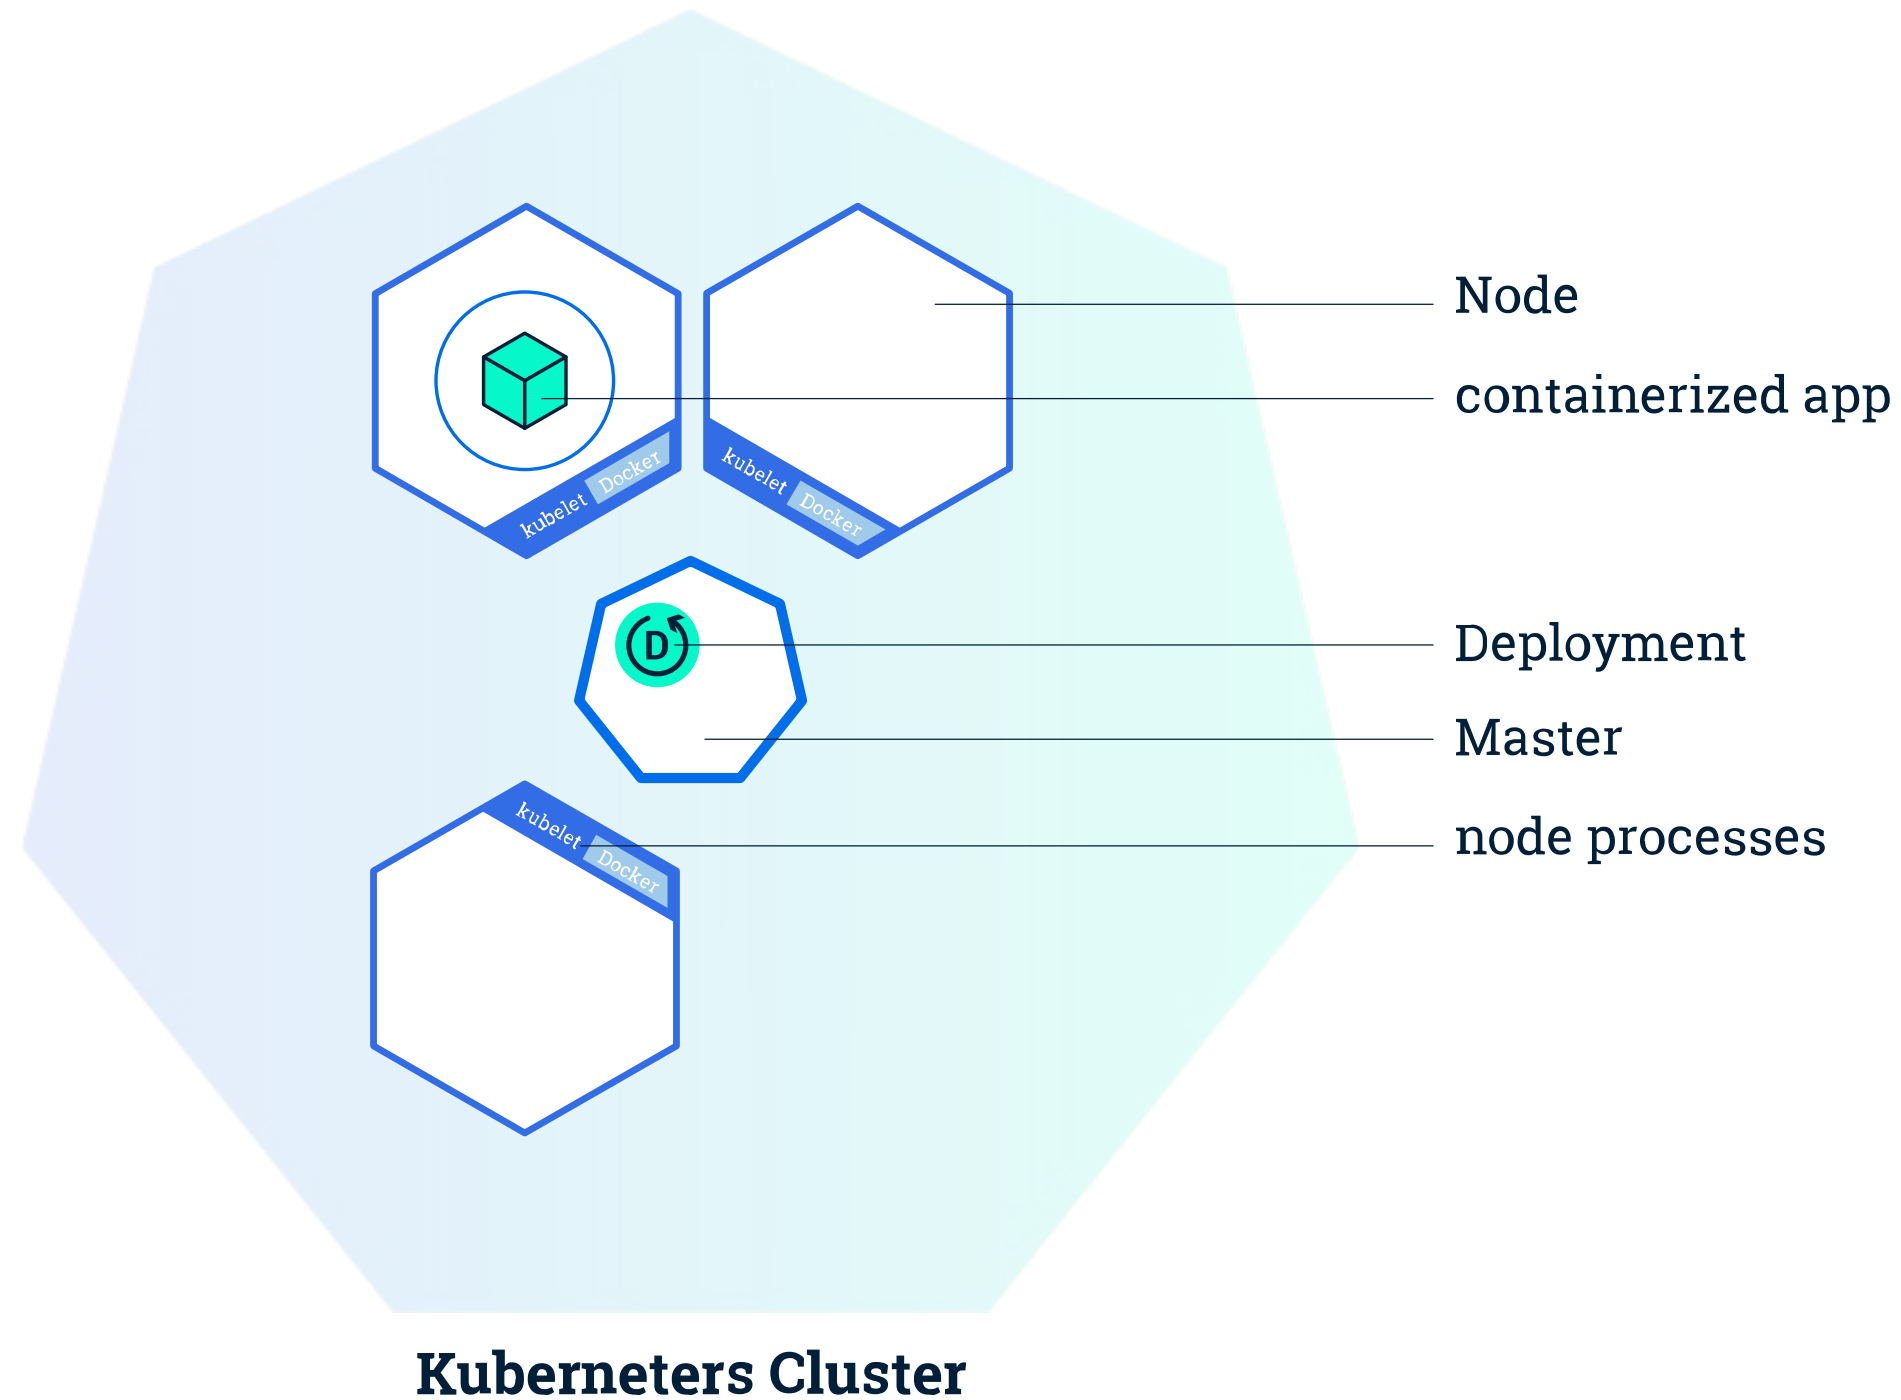
Kubernetes cluster



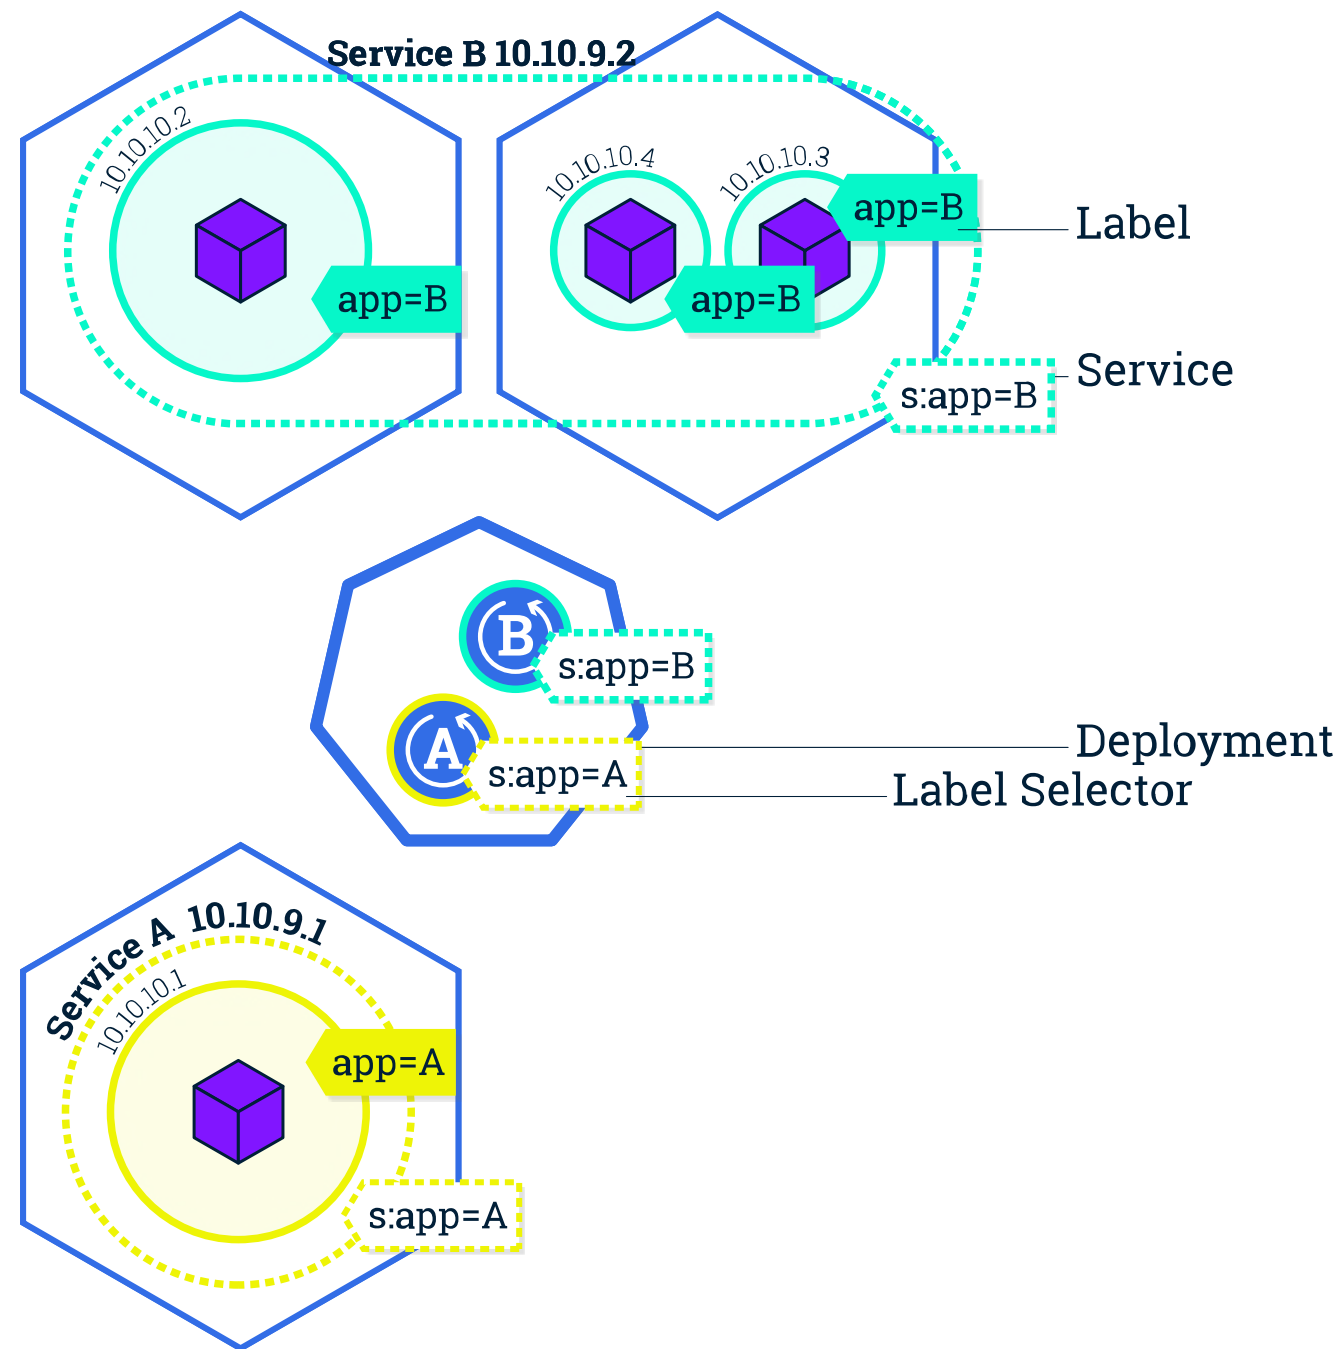
Kubernetes pods




Kubernetes deployments



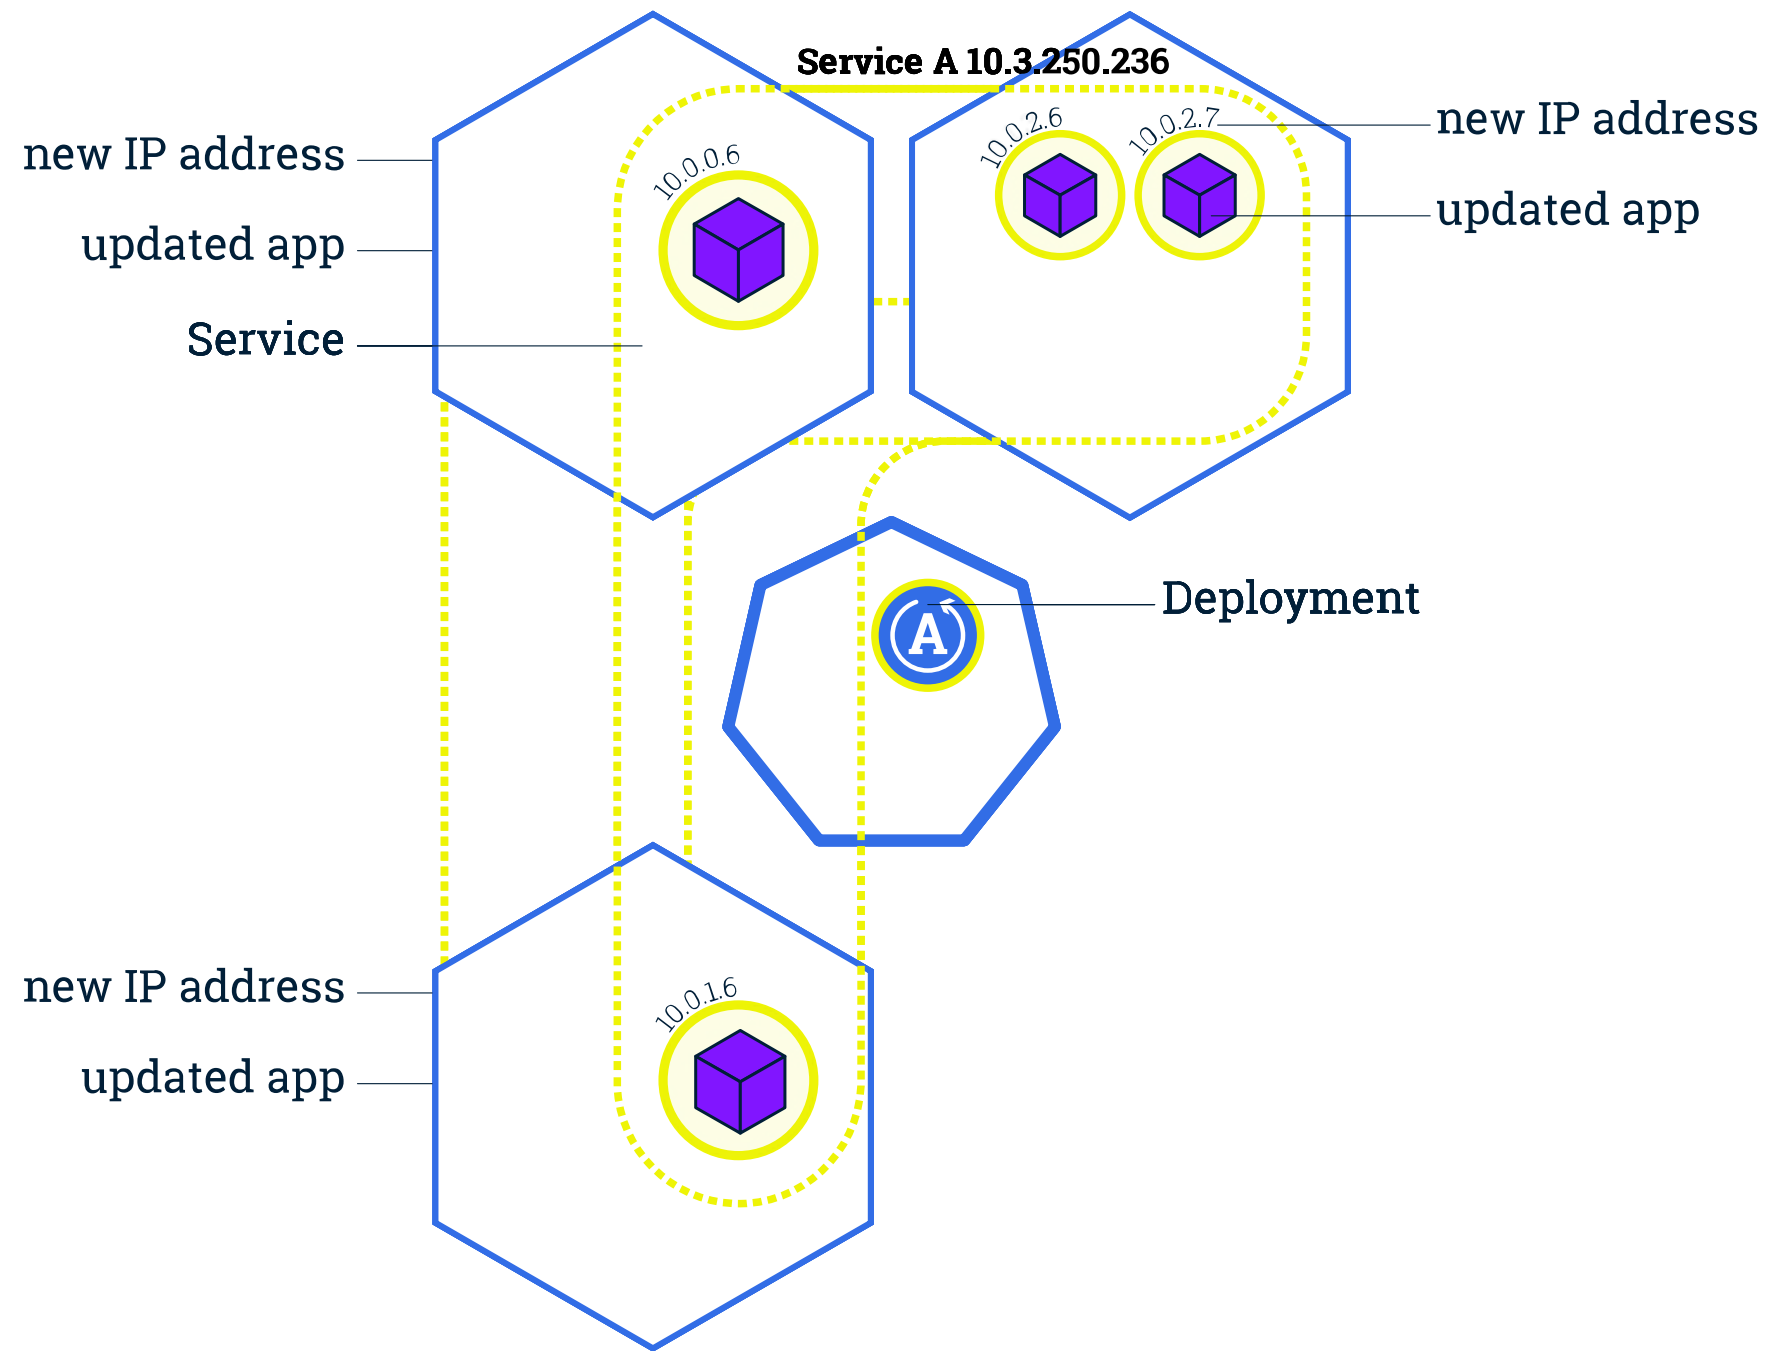
Kubernetes services



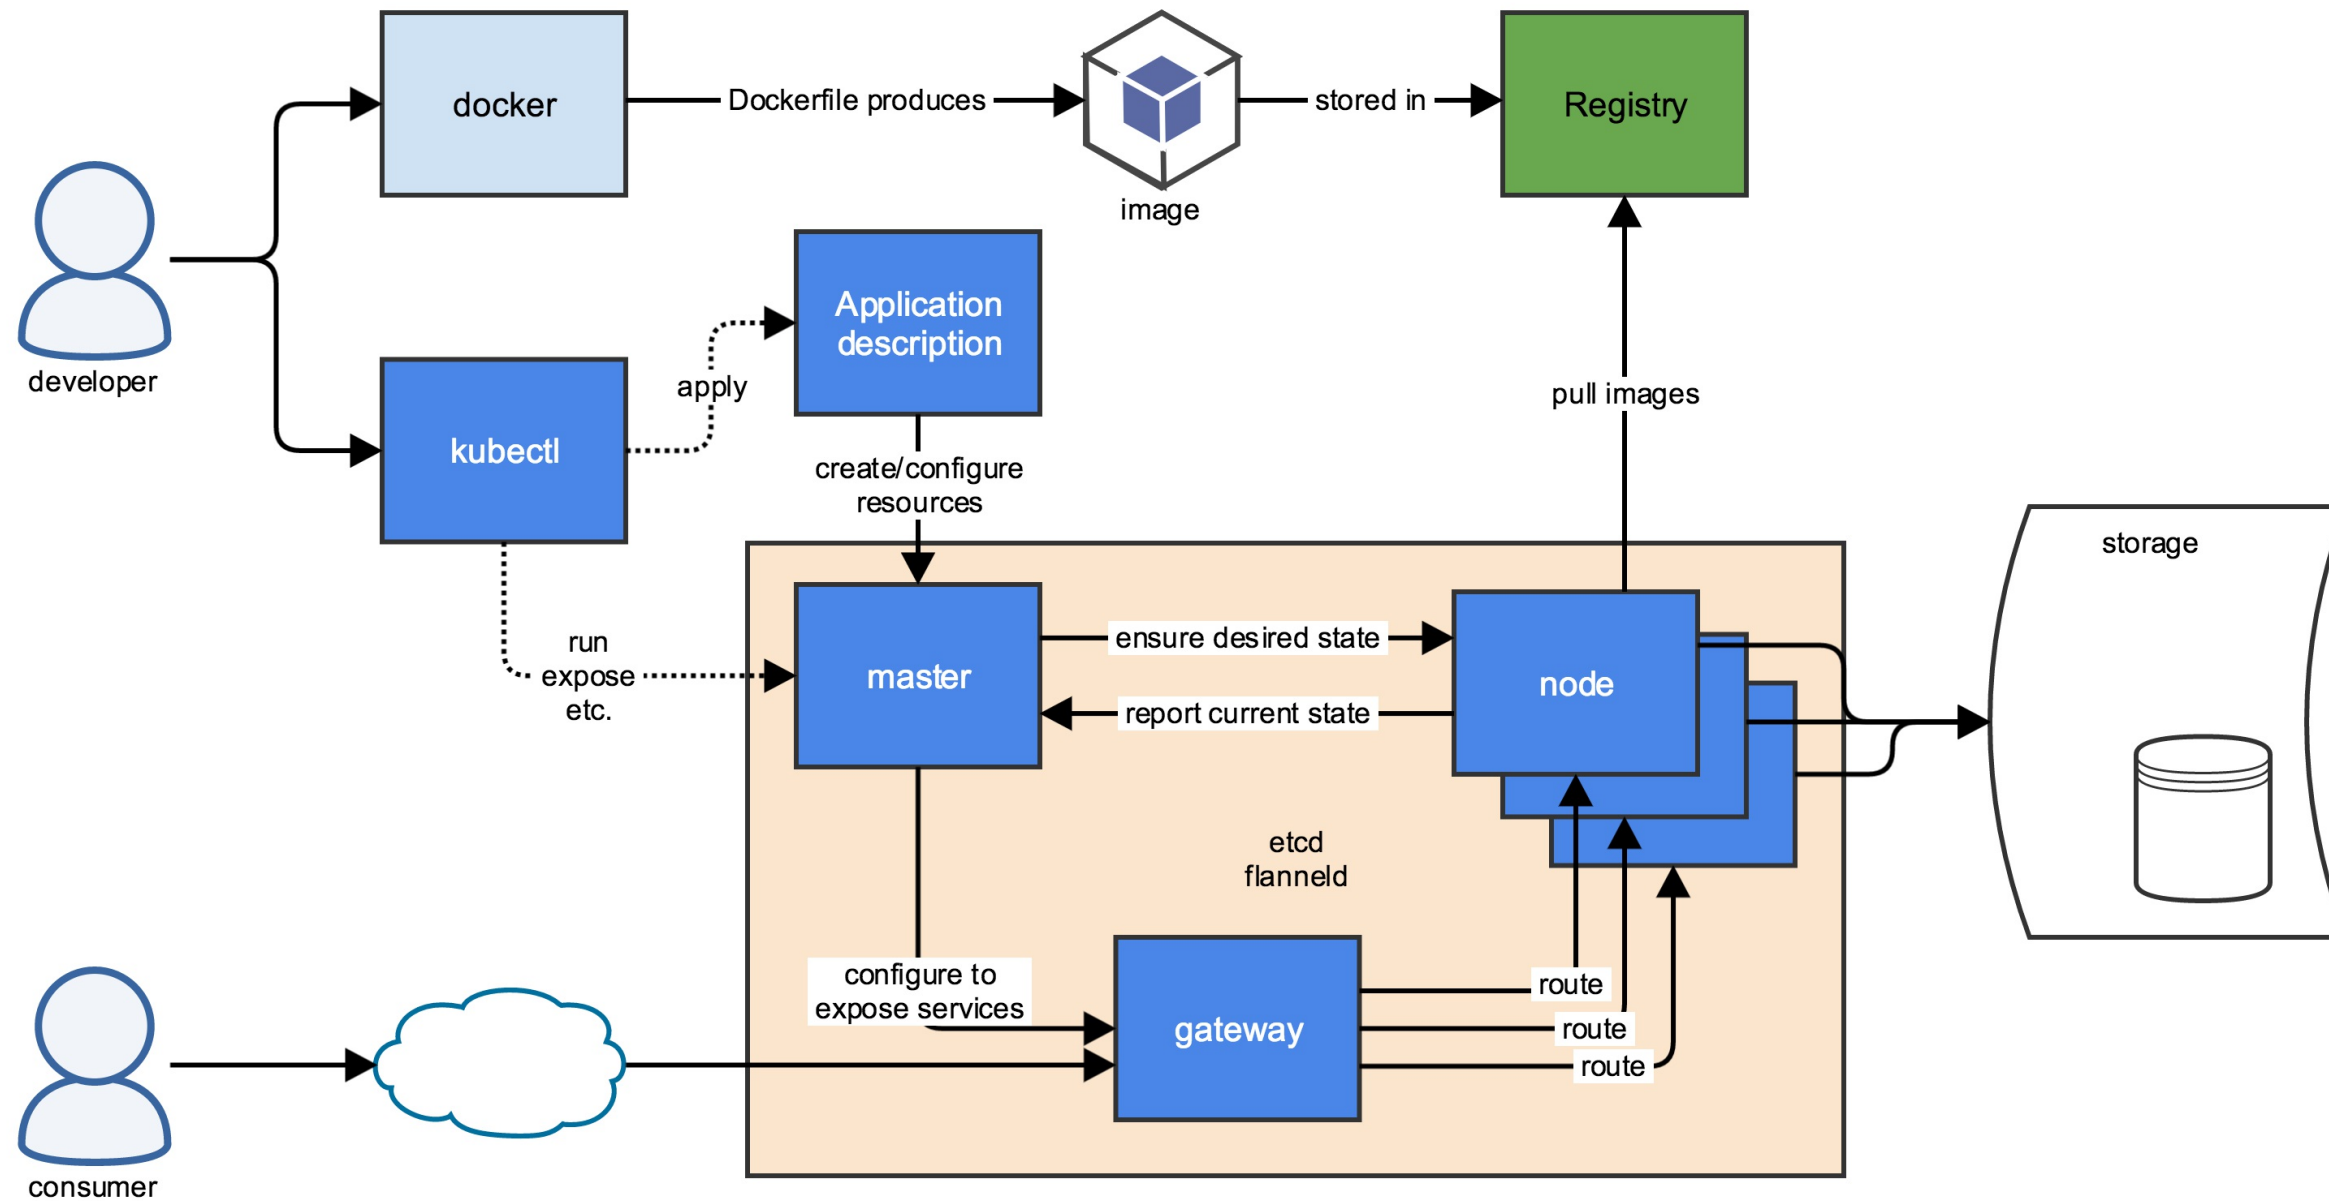
Scaling deployments



Deployment rolling upgrades



Kubernetes architecture overview

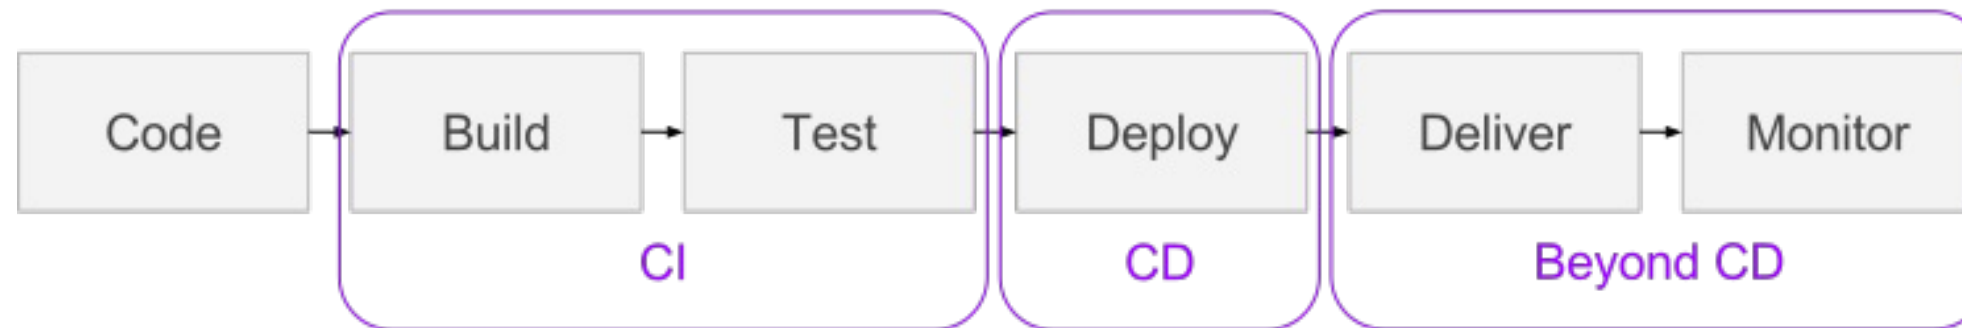


HETEROGEOUS HOSTING OF AROUND 450 WEB SITES AND APPLICATIONS

- Legacy web cluster
 - Bare metal shared hosting
 - Hard to maintain and update (shared libraries and software)
 - Resource sharing (noisy neighbors)
- Kubernetes based web cluster
 - Each stack component is containerized and evolves independently
 - Dedicated resources for each application (quotas)
 - Environment isolation for more security (at both container and system levels)
 - Horizontal scaling and automatic provisioning

CONTINUOUS DEPLOYMENT OF OUR APPLICATIONS ON KUBERNETES

- Using Gitlab CI
- Easing development and deployment processes (autodevops)
- Several versions of the same software stack can coexist



WORKING TOWARDS MICRO SERVICES

- Each stack component is independent
- A Kubernetes pod should host one and only one component
- Data storage and data processing are separated

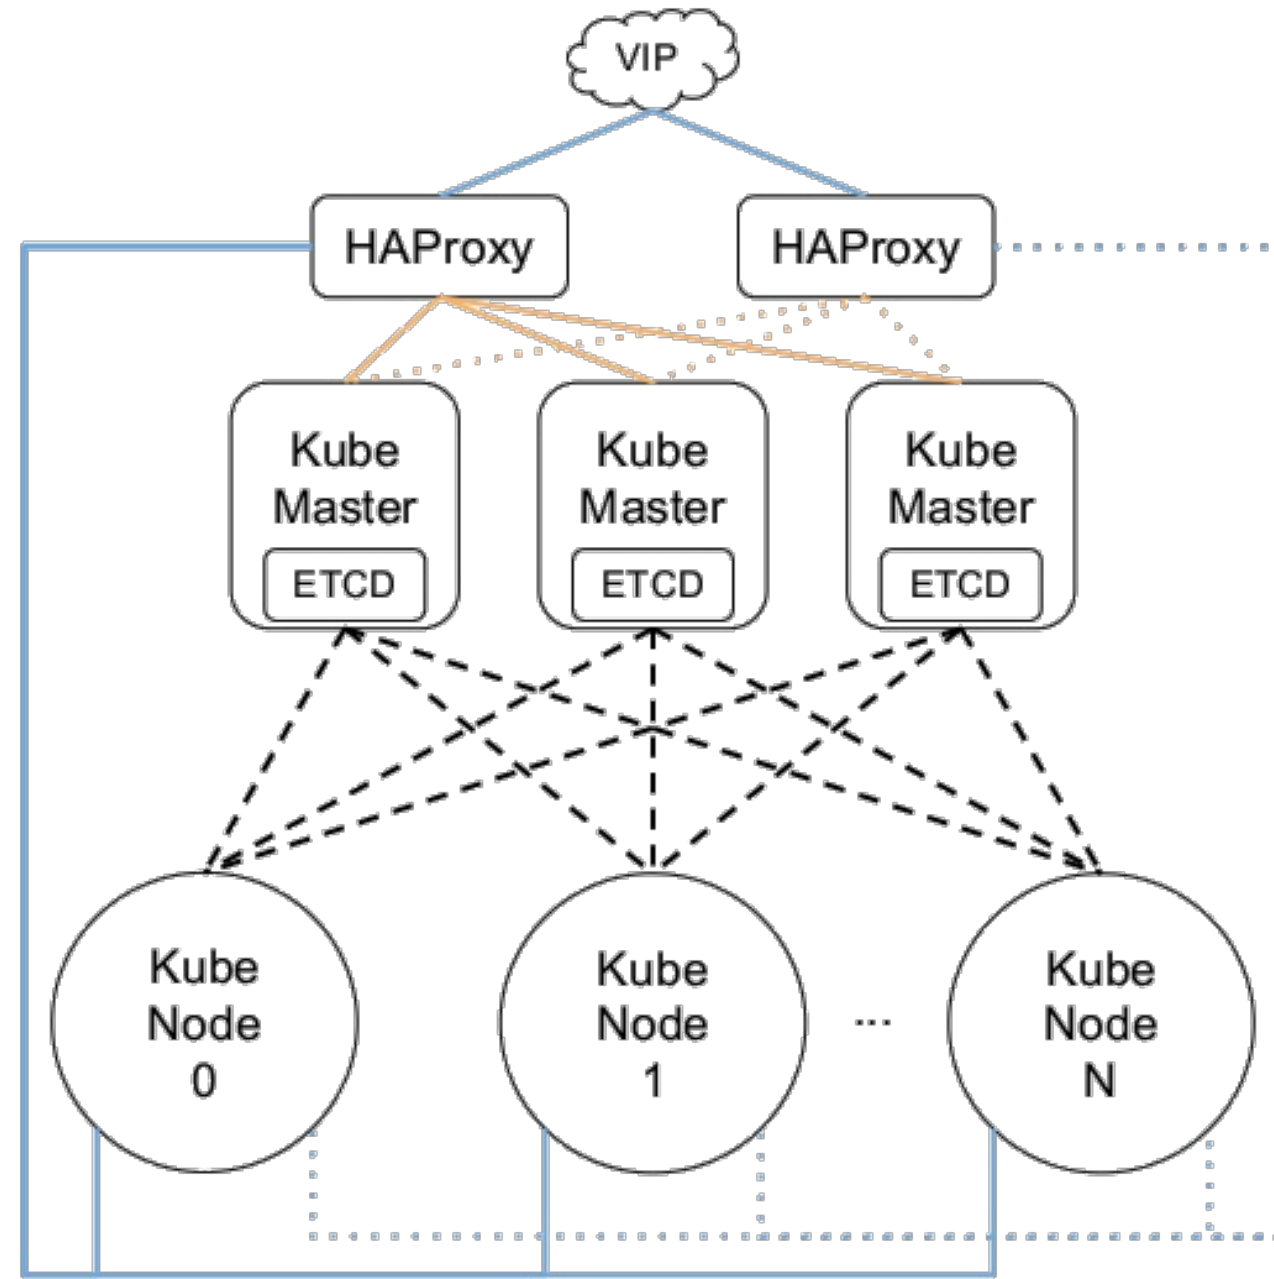
Kubernetes infrastructure



kubernetes



docker



- Dedicated hardware resources for the Kubernetes cluster
 - Bare metal hosting (masters and workers)
 - Cloud hosting on Openstack @ CCIN2P3 (offloading workers)



- Using existing CCIN2P3 storage systems
 - NFS/Isilon (legacy)
 - Ceph (CephFS / CephRBD)



- Fully integrated into CCIN2P3 operations workflows
 - Foreman (host system installation)
 - Puppet (configuration)



- Automated application deployments
 - Docker compose -> Deployments
 - Kubernetes Helm Charts
 - Ansible + Tower / Rundeck (delegation)



- Using CCIN2P3 Coloss

- Syslog-ng + collectd
- Kibana and Grafana dashboards



- Nagios for alerting

Nagios®

- Advanced application dashboards when available

- Prometheus
- Piwik (websites)

