


Running OpenShift in production "Are we brilliant or stupid?"

Kalle Happonen – SIG-CISS Athens 29.11.2019



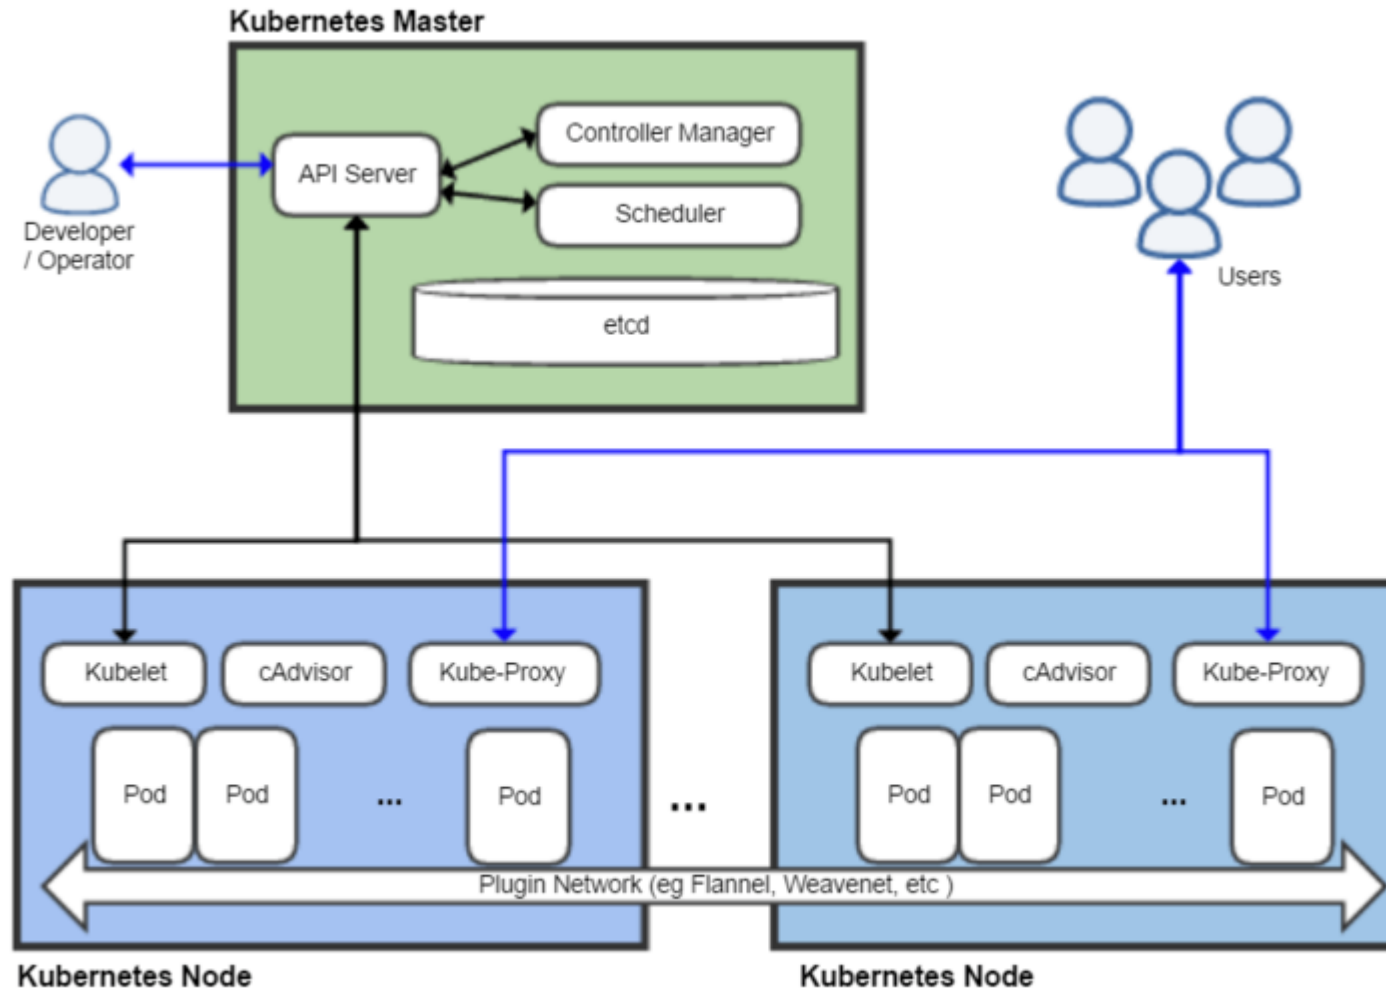
CSC – Finnish research, education, culture and public administration ICT knowledge center

Containers



docker

Kubernetes



OpenShift / OKD

- Kubernetes +
 - Installer
 - Built in container registry
 - Metrics
 - CI/CD workflows
 - ...
 - **!!_**Multitenancy**_!!**
 - Less privileges to pods
 - Non-root users only in containers



Rahti

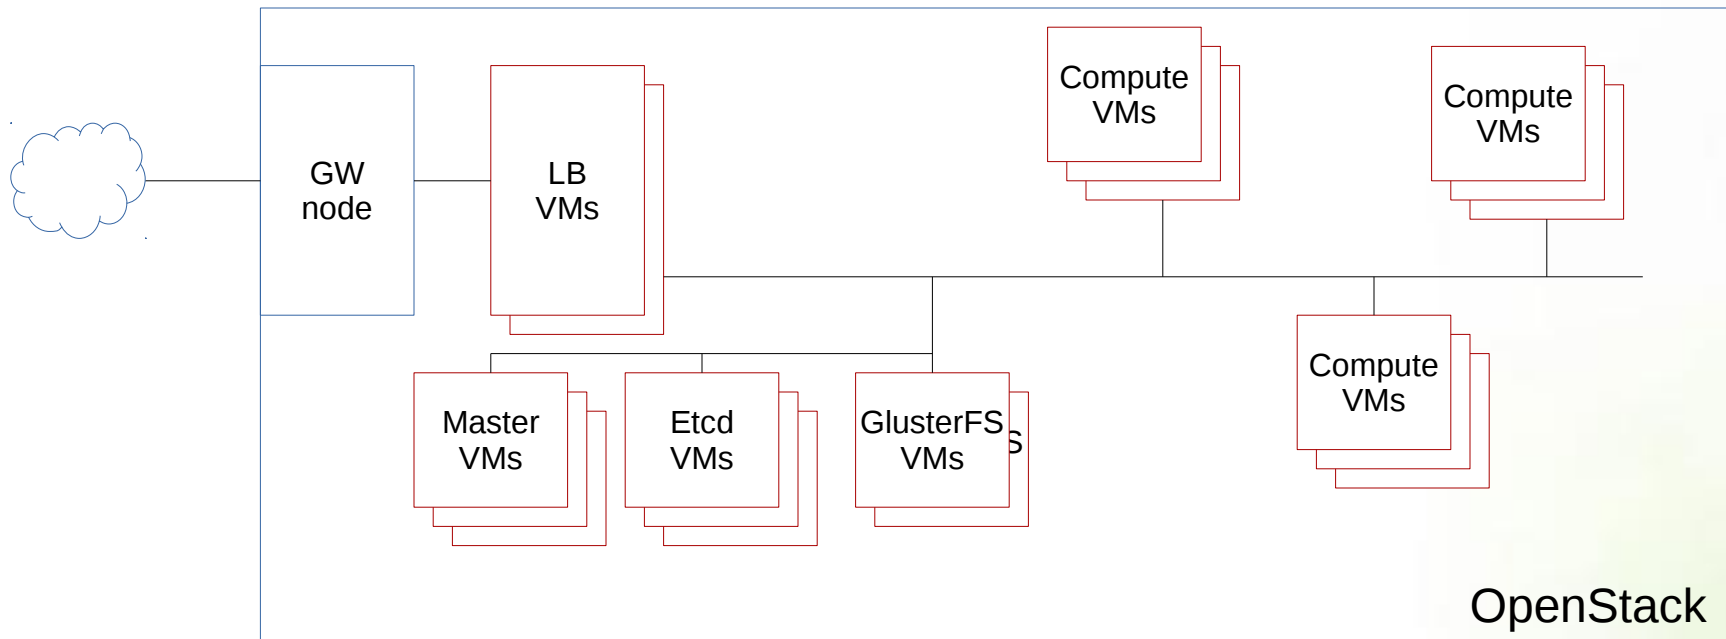
- Started developing ~2017
- Closed beta 2018
- Open Beta July 2019

- Running CSC production services

- Currently on 16 servers
 - 80 vcores / server
 - 384 GB memory/server
 - 5 TB SSD
 - 25 GbE nics

Architecture

- Runs on CSC's cPouta OpenStack IaaS



Storage

- GlusterFS on Cinder
 - For read-write-many use cases
 - Replicated design had problems with stability
 - Current non-replicated okish, but unmaintainable in the long run
 - Problems specifically with postgres
- Cinder
 - For read-write-once use cases
 - Better distributed failure tolerance
 - 26 volumes / VM limit (being fixed)
 - Cinder API calls if VM goes down
 - API load

Network

- VXLAN over OpenStack tenant nets (VLANs)
 - Working fine.
 - Performance not native.
- Ingress
 - Ingress bottleneck in current design (reworking) – single IP for access
- Egress
 - Dedicated egress IPs possible, with work, using floating IPs
- Future
 - How to grow the network stack?

Compute

- Pod limit per VM
 - Current kubernetes limit 256 pods per server
 - Our servers may be way too large for many use cases
- GPU
 - Non-privileged pods running GPUs.
 - Had to write custom SELinux module

Monitoring

- Prometheus resource consumption
 - Our system starts to be large
 - This can be seen in Prometheus

Registry

- No real issues
- Recently moved registry backed to object store

Integration

- “Project” concept
 - Group of users that gets resources and services
 - Doesn’t map directly to OpenShift easily
 - Solution: Tag namespaces with project information for accounting
 - Automatic quota management per project?

Upgrades

- OpenShift 4 upgrade
 - New deployment tools
 - Need to move over all workloads
 - What do we do with storage?
 - A lot of big questions
- Bonus: RedHat and IBM and OpenShift
 - What does the future hold?



Click icon to add
picture



<https://www.facebook.com/CSCfi>



<https://twitter.com/CSCfi>



<https://www.youtube.com/c/CSCfi>



<https://www.linkedin.com/company/csc---it-center-for-science>