CERN'S JOURNEY TOWARDS A KUBERNETES-BASED ACCELERATOR CONTROL SYSTEM

Thomas Oulevey

Controls Software & Services Group, Beams Department, CERN

Goals and Drivers

- **Run all Controls software in an orchestrated, containerized environment before end of LHC long shutdown 3**
  - Optimize infrastructure resources (cost & energy)
  - Improve time to recovery & business continuity
  - Streamline DevSecOps/GitOps practices both internally and with industry
  - Standardize deployment environments across several data centres
- **In turn**
  - Become more agile in hardware & software management
  - Reduce costs associated with diverse practices (infrastructure, development, operations, administration)
  - Facilitate onboarding & mobility of people working across various sub-systems

Background Context

- The Controls Datacenter hosts critical Accelerator services
  - 25 racks
  - 500 bare-metal servers
  - **Isolated network**
- Major upgrades constrained to accelerator schedules
- Diverse in-house DevSecOps practices
- Underused CPU and memory resources

Controls Datacenter

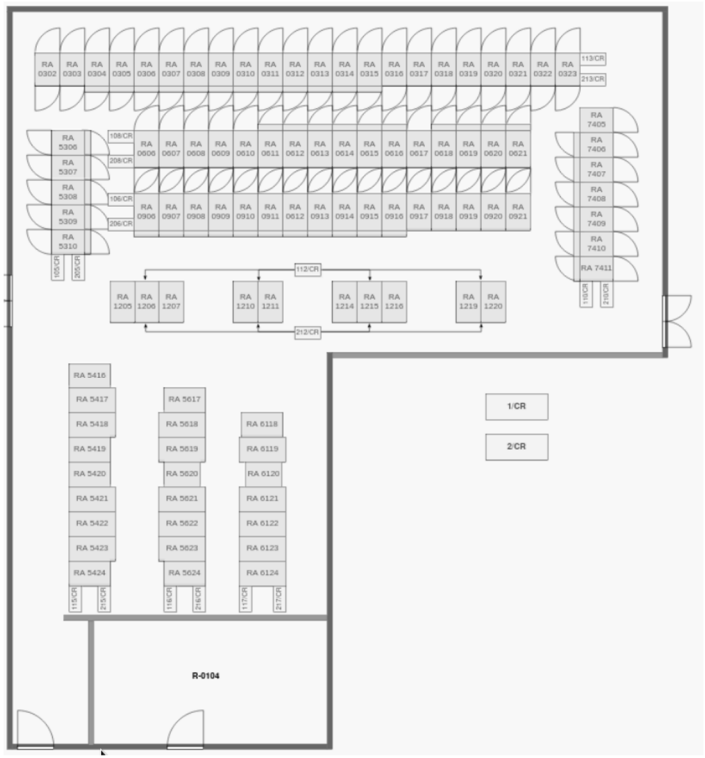

Why Kubernetes as an Engineering Platform ?

- **Define common workflows with our partners**
- **Optimize resources usage (cost and agility)**
  - Power efficiency
  - Better hardware lifecycle control
- **Aim toward transparent service updates and migration**
  - Both for the hardware and the Operating System
- **Industry standard**
  - Containers as a packaging primitive
  - Pre-packaged applications template (commercial and open source) such as Helm charts
- **Experience from CERN IT department with slightly different approach**
  - Kubernetes as a Service approach
  - but infrastructure can be rationalized
  - and a common deployment tool for services across CERN is desirable

★ **Select pilot services and deploy them to gain experience**

Journey's Timeline

First steps with containerisation
First steps with orchestration
Pilot services on Kubernetes

Road toward production
Controls Datacentre infrastructure refresh

Running of the first Production Accelerator Controls Services

Transition of all Remaining Accelerator Controls Services

ICALEPCS '19
Container in Controls workshop

Journey's Timeline

First steps with containerisation
First steps with orchestration
Pilot services on Kubernetes

Road toward production
Controls Datacentre infrastructure refresh

Running of the first Production Accelerator Controls Services

Transition of all Remaining Accelerator Controls Services

Pilot Services Overview

More details in the paper TH2AO03

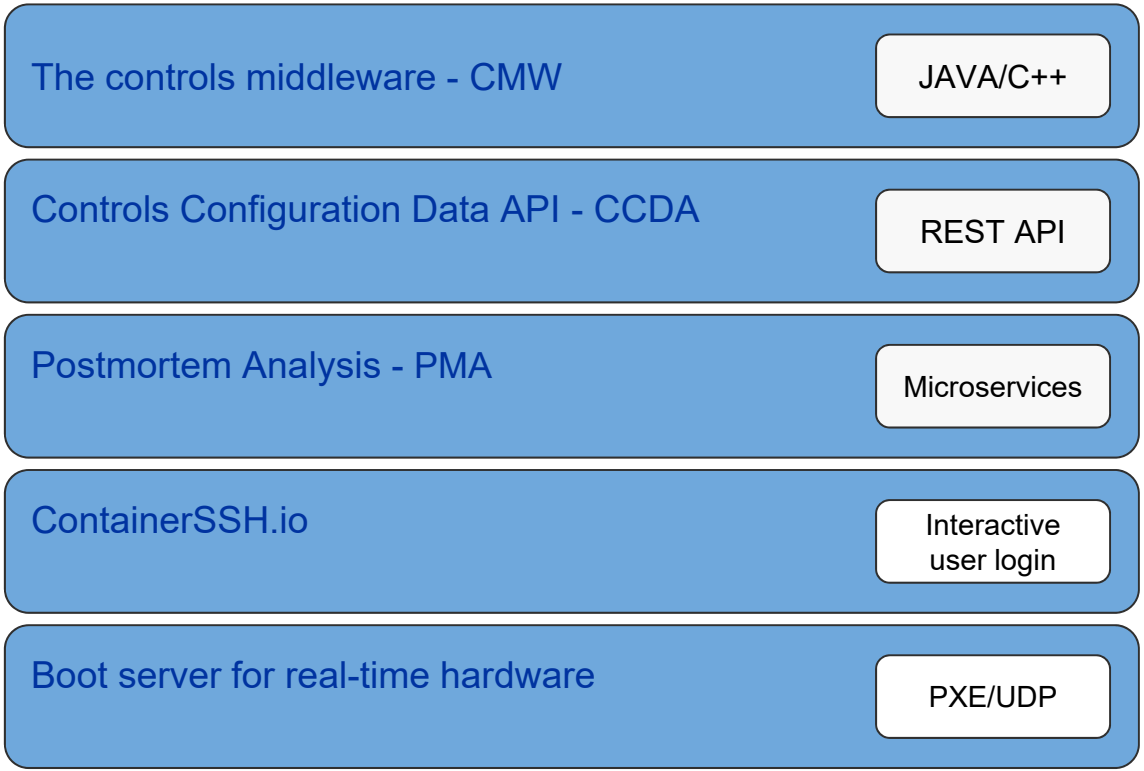

Next talk: Bartek Urbaniec (TH2AO04)

Pilot Phase: key takeaways

Iterative approach with dedicated time for users and frequent meetings

- Breaking silos between teams
- Support, solutions & workarounds discussed weekly to build common cross-team knowledge

Challenges to work on an isolated network

- Containers distribution and curation is critical (Harbor with replication)
- Security scanning and auditing (Trivy or Clair available with Harbor)
- TLS certificates with an external provider (Letsencrypt)

Build applications with orchestration in mind

- Utilise Microservice architecture to enable horizontal scaling or better agility
- Rationalise redundancy mechanisms at the network/kubernetes level
- Provide common configuration and dependencies within the kubernetes ecosystem (e.g. Oracle Database settings as configmap)

From Pilot to Production: infrastructure overview

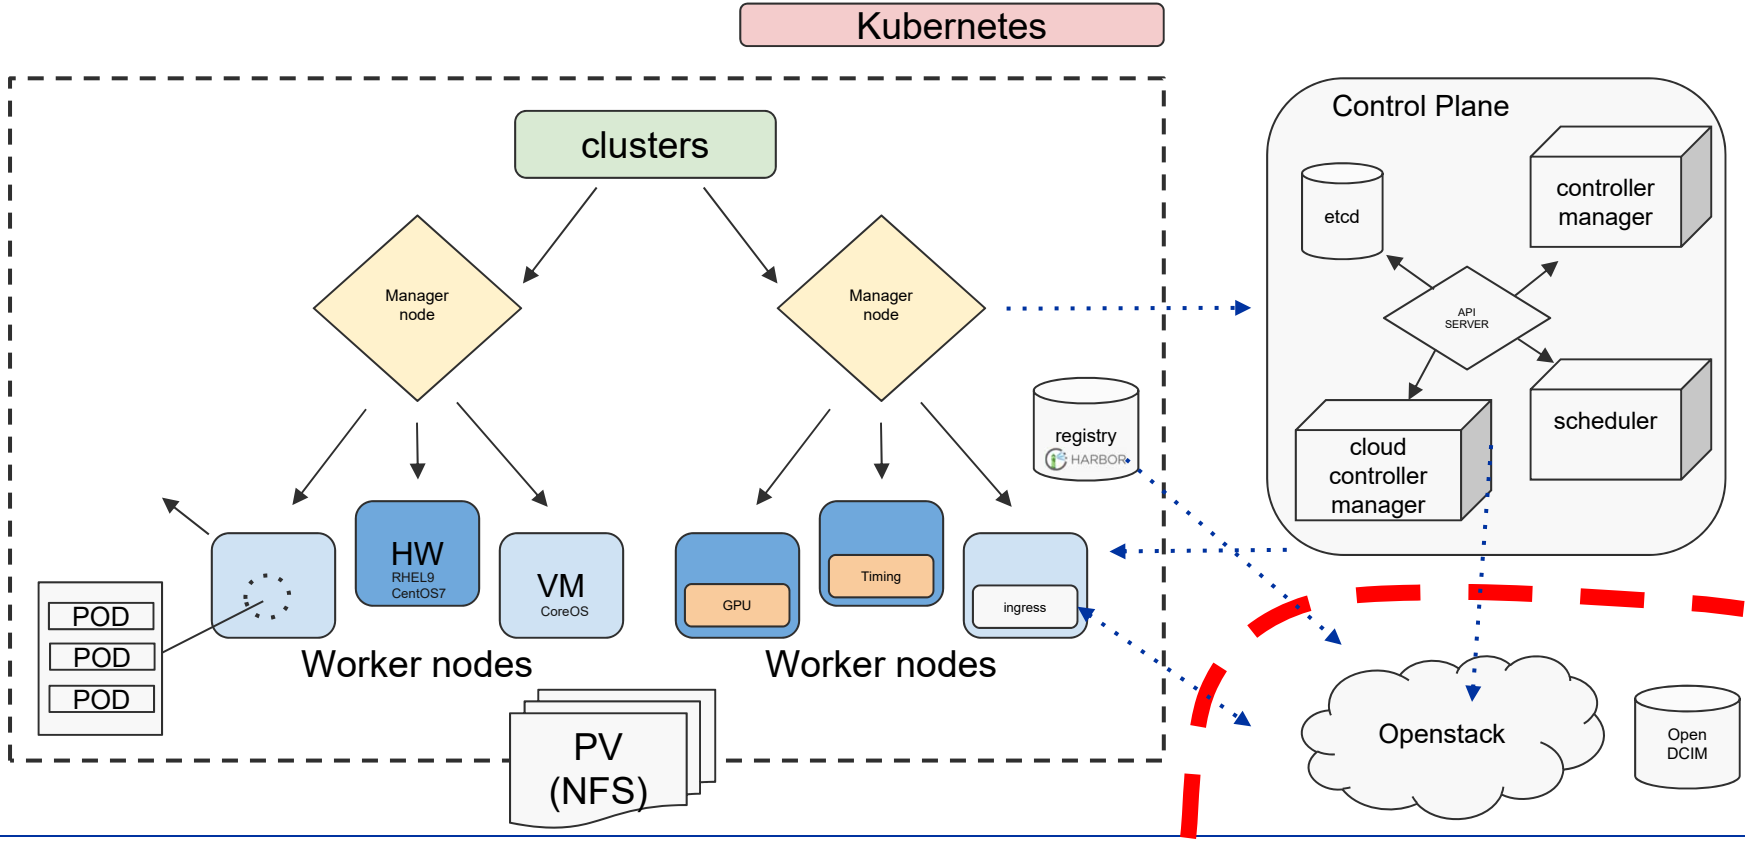

From Pilot to Production: key areas of attention

Reinforcing our infrastructure

- Secret management (e.g: Vault)
- NFS storage for persistent volumes (Longhorn as an alternative)
- Modern monitoring and tracing (e.g: Prometheus)
- Modern network (e.g : vxlan, evpn) and CNI (e.g: cilium)
- Topology constraints (e.g: OpenDCIM integration)

+ Improve security & maintainability

Unique opportunity for a new DevSecOps methodology

- Gitlab and ArgoCD to deploy centrally applications on dedicated clusters
  The control plane and Harbor registry are also deployed with ArgoCD
- Container security scanning
- Software Bill of Materials (SBOM) integration with CERN tools
  MO4BCO03, B. Copy, "Protecting Your Controls Infrastructure Supply Chain"

Summary

Our journey towards Kubernetes is well underway

- **Successful pilot with representative services**
- **Vital experience gained and lessons learned**
  - software architecture ; embrace orchestration primitives
  - datacenter architecture ; modernise network, storage and improve hardware lifecycle
- **Strategic decisions have been taken**
  - Organisation wide commitment for transition to production
- **Many challenges ahead**
  - Revise infrastructure administration and procurement
  - Adaptation of all Controls services to be Kubernetes ready
  - Need to promote and rationalize DevSecOps best practices
  - Improve users' tooling for auditing, debugging and troubleshooting
- ★ **The next step on the journey has already started**
  - Preparing a production ready platform for the first half of 2024