

Kubernetes security guidelines automated validation

Paweł Wieczorek

11.06.2019

Intro

- While various underclouds are supported, Kubernetes is widely used for container orchestration
- Clusters are deployed in multiple ways:
 - Rancher (Casablanca default)
 - RKE (Dublin+ default)
 - KRD (with Kubespray)
- Are they sufficiently secured by default?
- Is there a way to quickly locate potential security issues?

Problem

- Broad attack surface if cluster is misconfigured – several internal services: API server, scheduler, controller manager, etc.
- Available mitigation solutions involve careful inspection:
 - [Kubernetes](#)
 - [CNCF](#)
 - [Rancher](#)
- Deployment defaults need adjustments to provide appropriate balance (ease of use shall not compromise security)
- Support for repeated inspection and monitoring might prove itself useful

Ongoing efforts

- Utility for security guidelines validation development (ONAP-focused, yet applicable to other clusters as well)
- Extensibility in mind – support for various guideline providers
- Least performance cost possible – not to introduce avoidable overhead

Solution design

- Collect data from available sources
- Process gathered information
- Point out weaknesses (and provide reason)
- (Suggest remediation)

Proposal

- In-depth resilience testing
- Migration to Dublin+
- Get feedback
- Adjust and adapt

Thank you

- That's all