

The 7th Asian Tier Center Forum in 2023

Research on The Container Runtime of The CERN EOS Distributed File System in a Kubernetes Environment

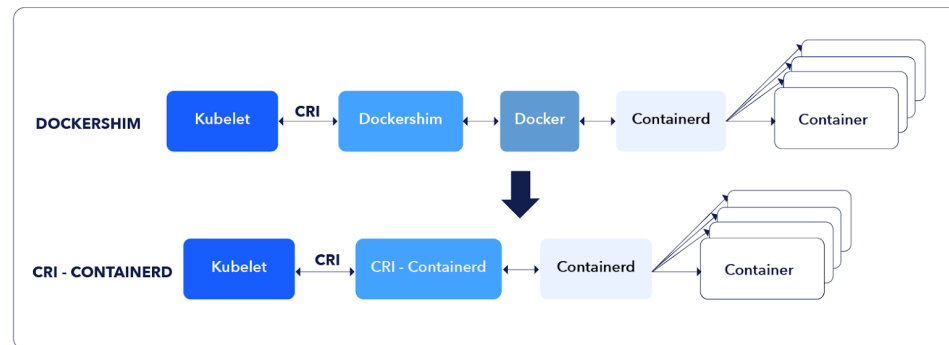
Jung-Bin Kim

Data Computing Laboratory,
Chungbuk National University, South Korea,
jungbin.kim@cern.ch


2023.11.01(Wed)

Container Runtime in K8S

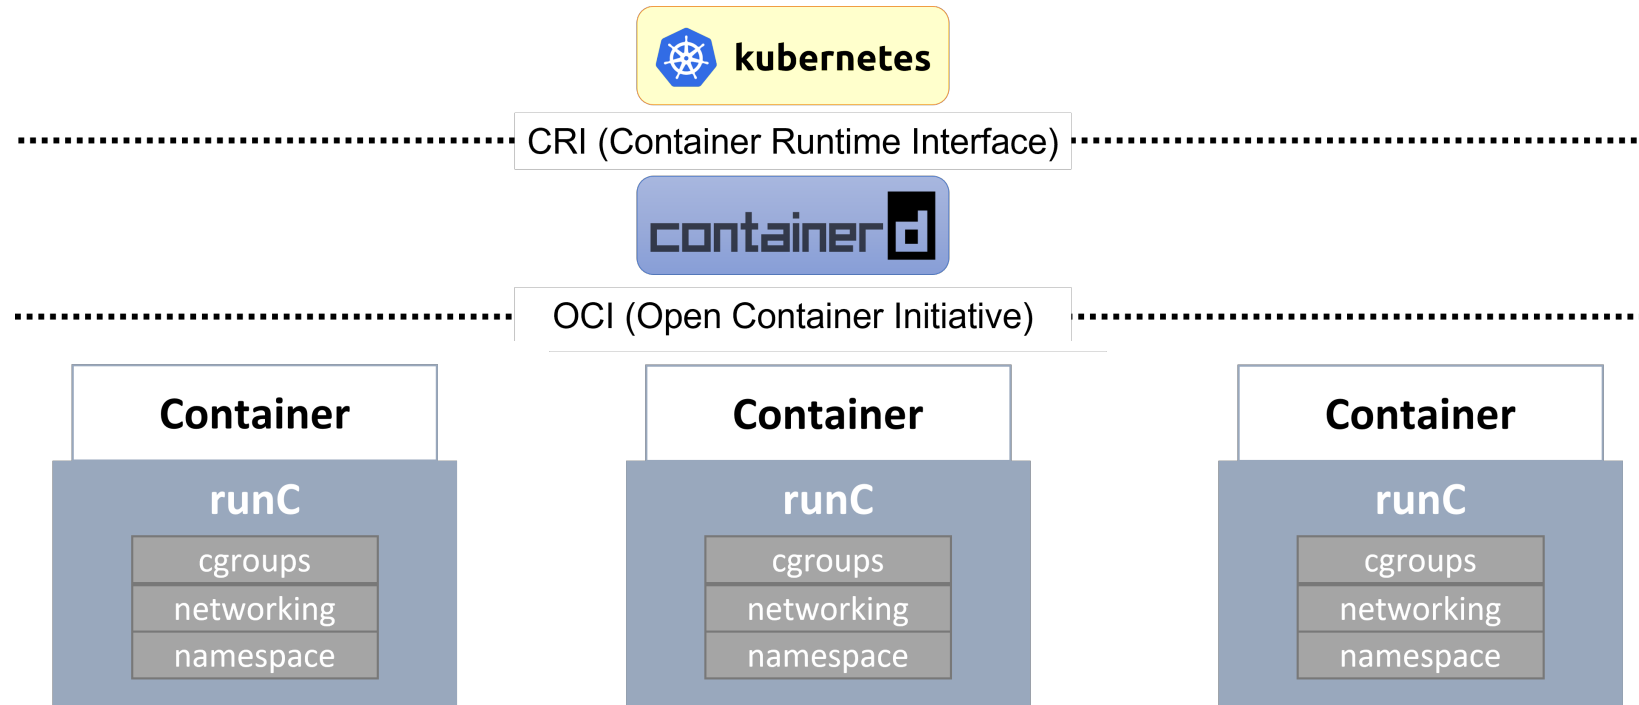
- **The importance of choosing a container runtime in Kubernetes.**
 - In a Kubernetes configuration, the container runtime is one of the key elements responsible for executing and managing applications.
 - The choice of container runtime can impact the overall performance of both applications and the system.



- **A Study of CERN EOS Distributed File System on Container Runtime Performance Analysis in a Kubernetes Environment**
 - The research goal is identifying a suitable container runtime for CERN EOS, which is used for large-scale scientific research.
 - So, we performed a performance evaluation of different container runtimes.

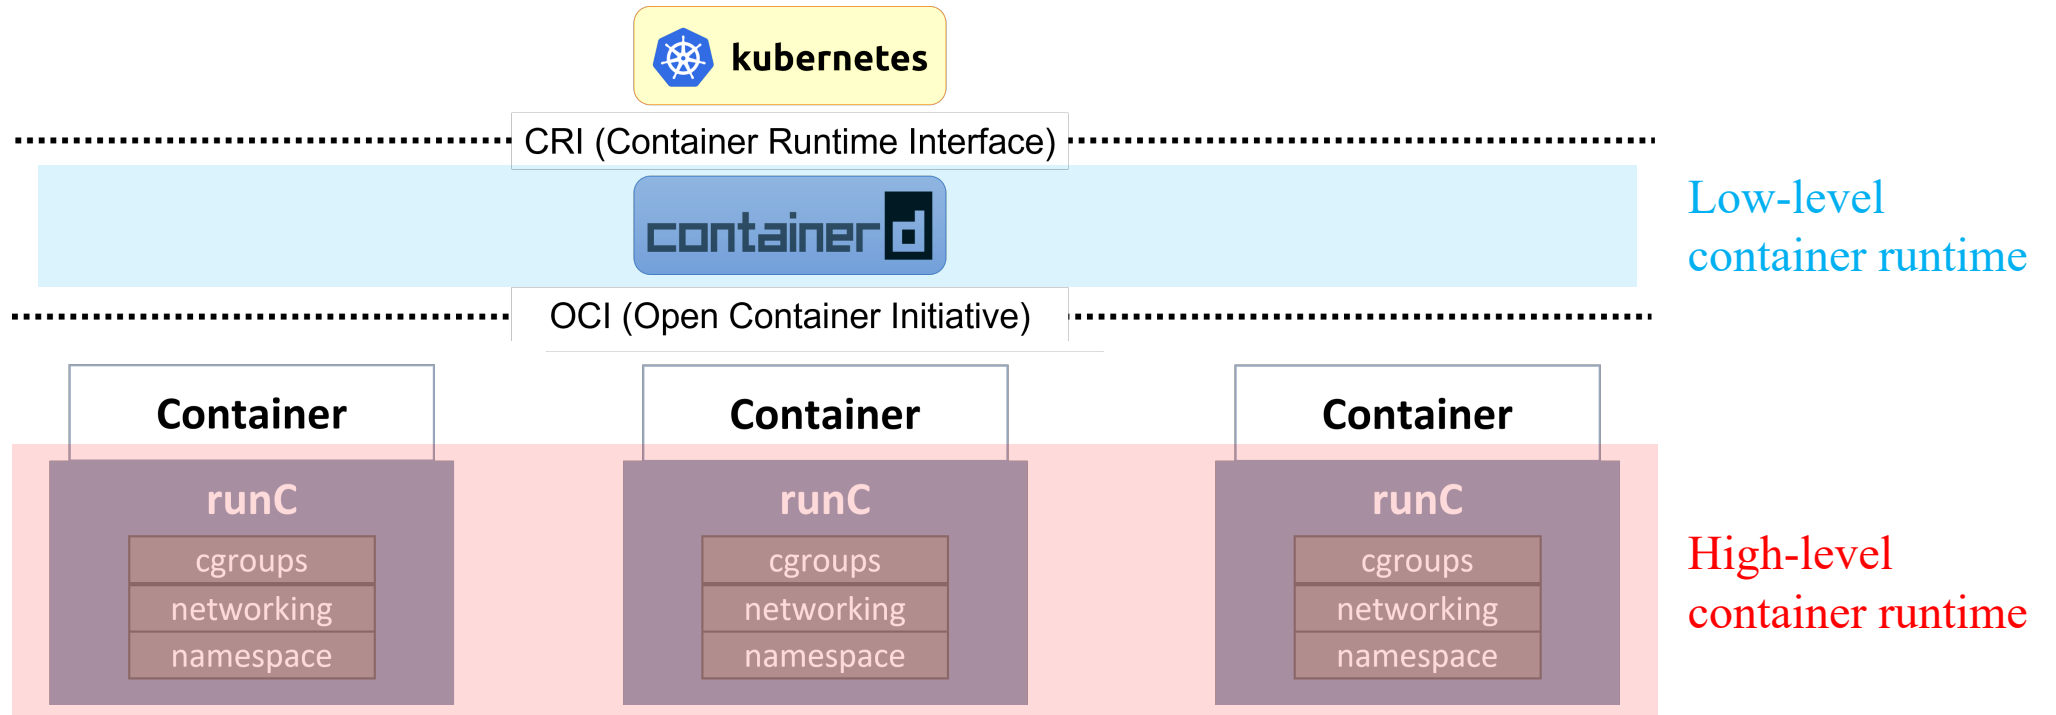


- **The Architecture of Kubernetes** : Kubernetes – **containerd** – **runc**
 - **containerd** is a tool that manages container runtimes.
 - **runc** actually do performs tasks such as cgroups, networking, and namespace isolation.



Container Runtime

- There are 2 type of the container runtime, High- and Low-level container runtime.
- Which type of container runtime do the mentioned belong to?
 - **containerd** - High-level container runtime
 - **runc** - Low-level container runtime



- **Diverse container runtimes**

- As shown in the picture, there are various container runtimes.
- It is expected that each container runtime has distinct characteristics that impact the system differently.

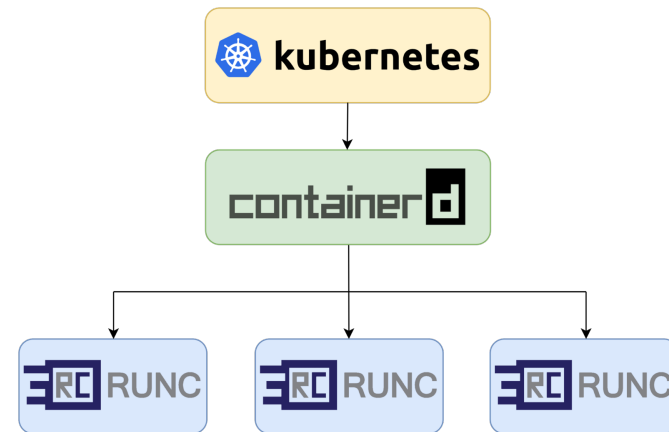
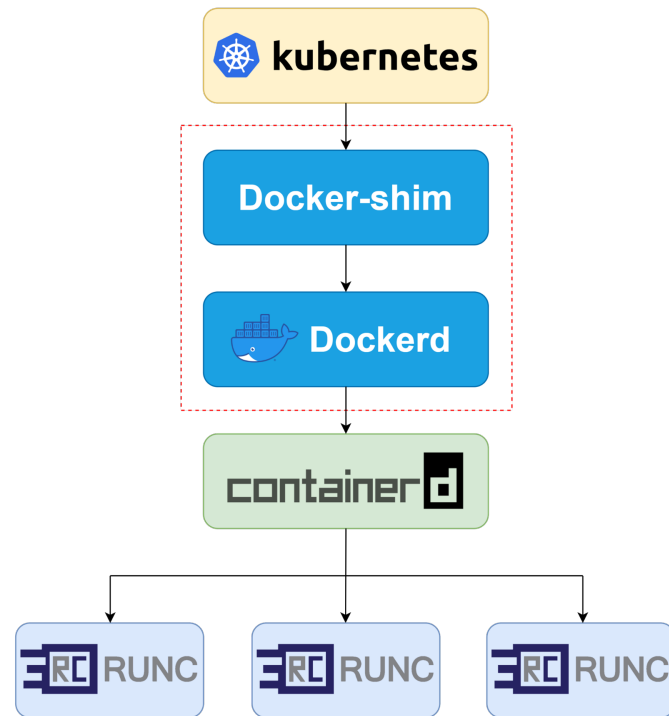


High-Level Container Runtime


Low-Level Container Runtime



- **Problems with Docker as a container runtime**
 - Compared to using `containerd` directly as the container runtime, `Docker-Engine` goes through an additional `Docker` intermediate layer.



- **Experimental environment**
- The experiments were conducted on a typical Kubernetes cluster prior to evaluating the container runtime for CERN EOS.
 - Because of measuring how performance varies as the choice of container runtime.



Server Specification	
Chassis	Power Edge R440
OS	CentOS 7
CPU	Intel(R) Xeon(R) Silver 4208 CPU @ 2.10GHz
RAM	128GB
Disk	960 GB SSD (Boot), 2 TB HDD (Data)

Container Runtime Version	
containerd	1.6.19
cri-o	1.26.1
Docker Engine	23.0.3


** To use docker-shim, we used version 1.23.0 of Kubernetes.*

Analytics - typical Kubernetes cluster



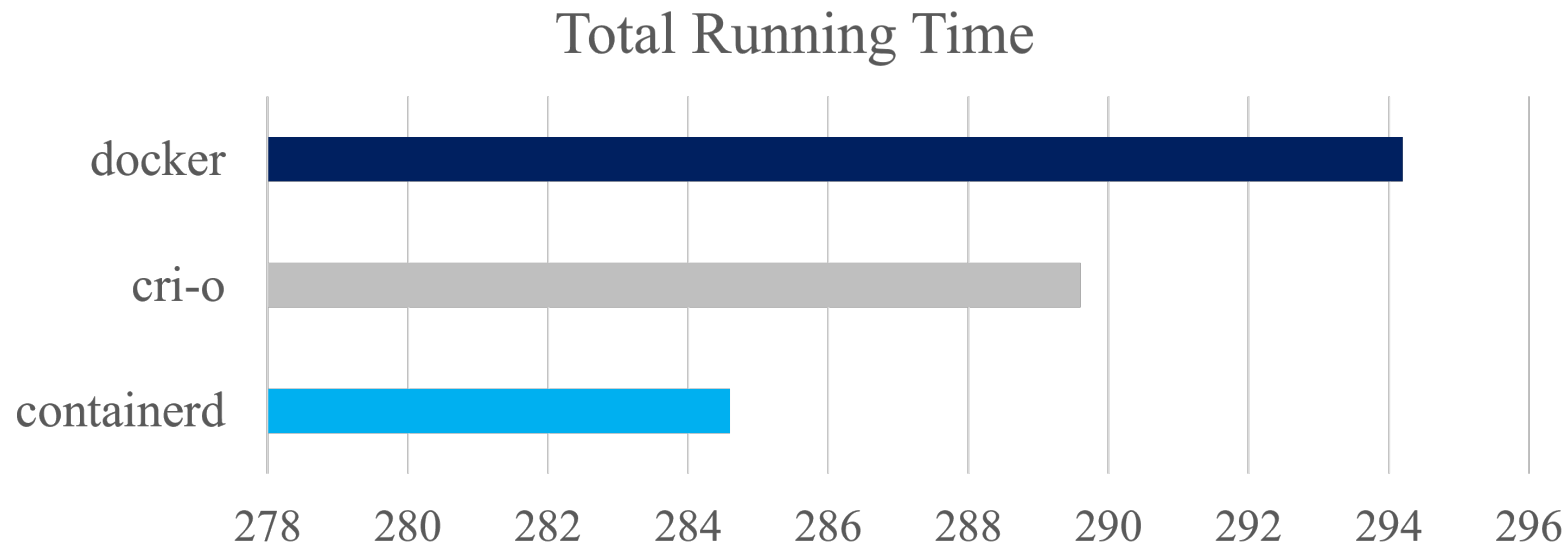
- **Evaluation of Container Runtimes – Container Runtime**

- `containerd` is generally stable and high performance.
- Docker-Engine is similar to `containerd`.
- CRI-O performs poorly compared to the other two container runtimes.




** Higher is good performed*

- **Evaluation of Container Runtimes – Total Running Time**
 - **containerd's** is 284 seconds and has the fastest running time
 - **Docker-Engine's** is 294 seconds and has the lowest running time
 - **CRI-O's** is 289 seconds and has the middle running time



** Lower is good performed*

- **The summary about Evaluation of Container Runtimes**



Bandwidth and IOPS

containerd > Docker-Engine > CRI-0

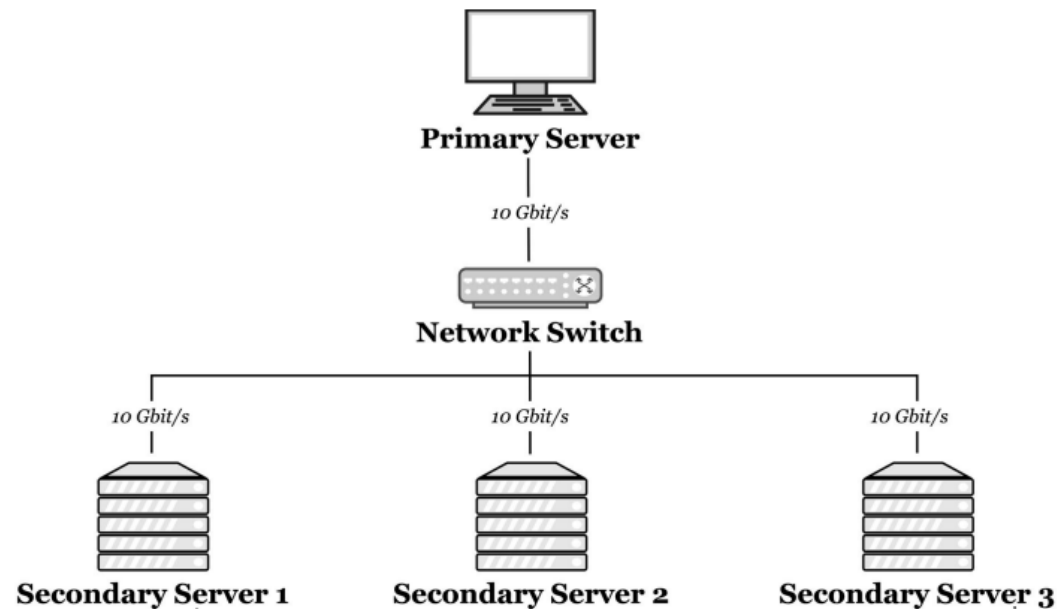
Total Running Time

containerd < CRI-0 < Docker-Engine

- Total running time is not always proportional to bandwidth and IOPS.
 - Docker-Engine has higher bandwidth than CRI-0, but higher total execution time.
- Docker-Engine and containerd are similar in bandwidth and IOPS, but has relatively large differences in total running time.
 - This is likely due to the overhead of Docker-Engine actually using containerd.


- **Experimental environment for CERN EOS K8S Cluster**

- To evaluate the container runtime for CERN EOS, experiments were designed with a different environment from the previous ones.

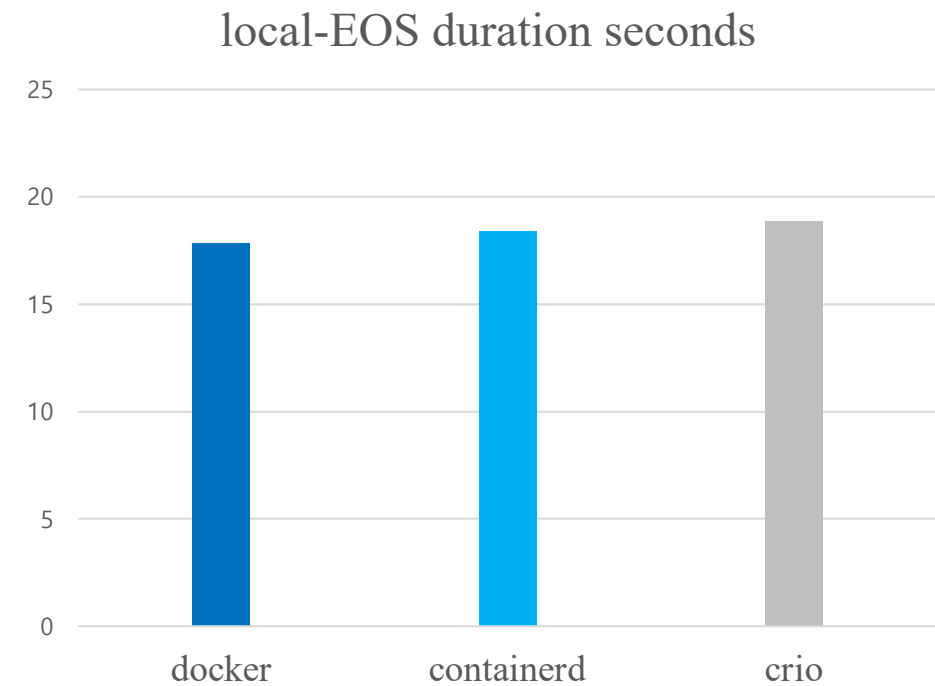


Server Specification	
OS	CentOS 7
CPU	Intel(R) Xeon(R) CPU E5-2680 2.70GHz
RAM	96GB
Disk	600GB HDD(Boot), 2.7 TB RAID Volume * 2
Network	10Gbit/s

- **Evaluation of Container Runtimes – Uploading files from the local system to EOS**
 - The bandwidth - Docker-Engine > containerd > CRI-O
 - The execution time - Docker-Engine < containerd < CRI-O

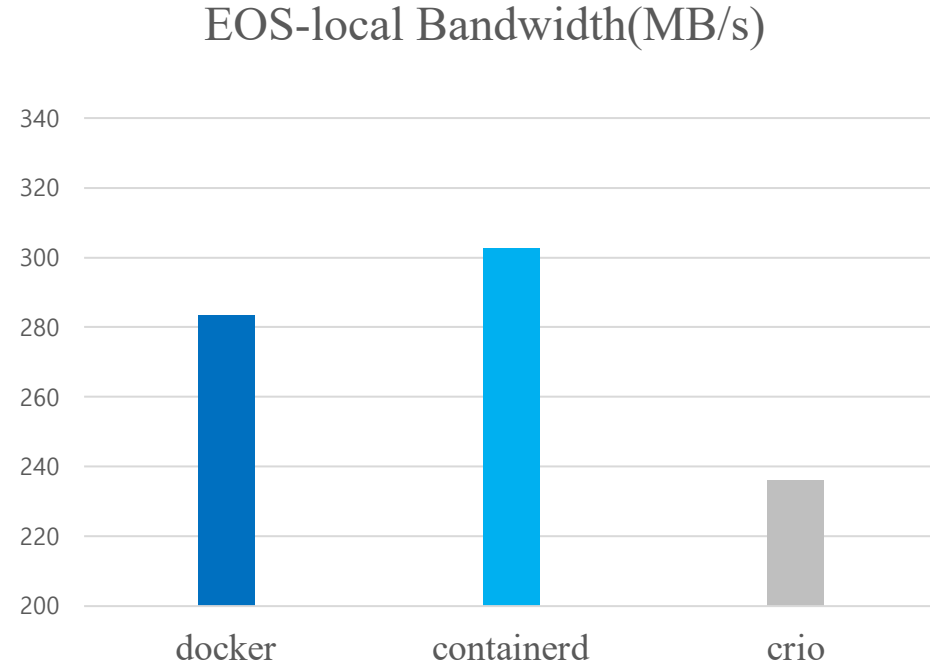


** Higher is better*

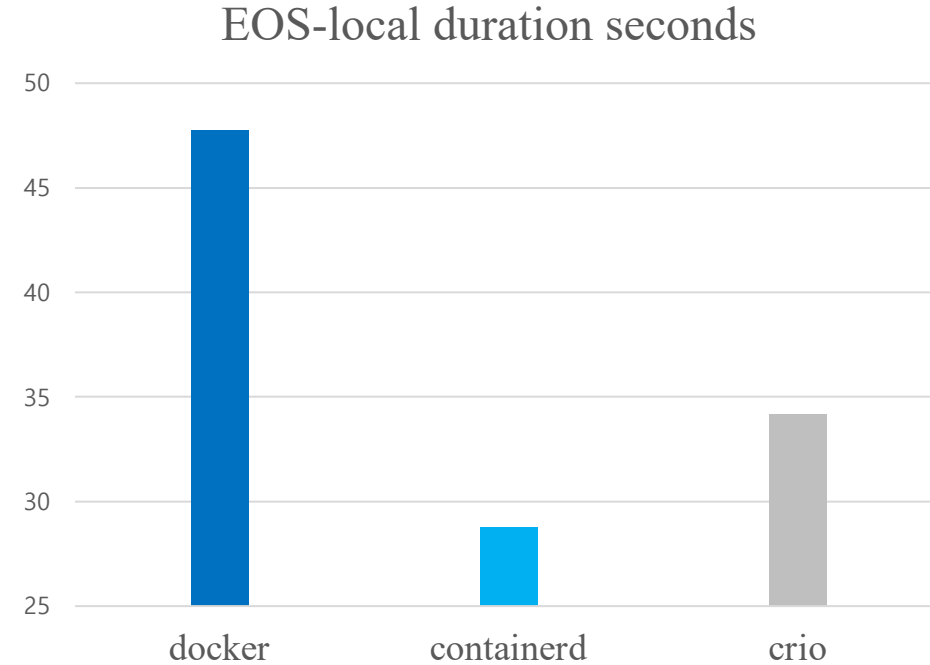


*Lower is better **

- **Evaluation of Container Runtimes – Downloading files from EOS to the local system**
 - The bandwidth - `containerd` > Docker-Engine > CRI-O
 - The execution time - `containerd` < CRI-O < Docker-Engine




** Higher is better*



*Lower is better **

- **The summary about Evaluation of Container Runtimes**

- It is concluded that integrating CERN EOS into Kubernetes using **containerd** is the most stable option.



Average Bandwidth

containerd > Docker-Engine > CRI-0

Duration Time

containerd < CRI-0 < Docker-Engine

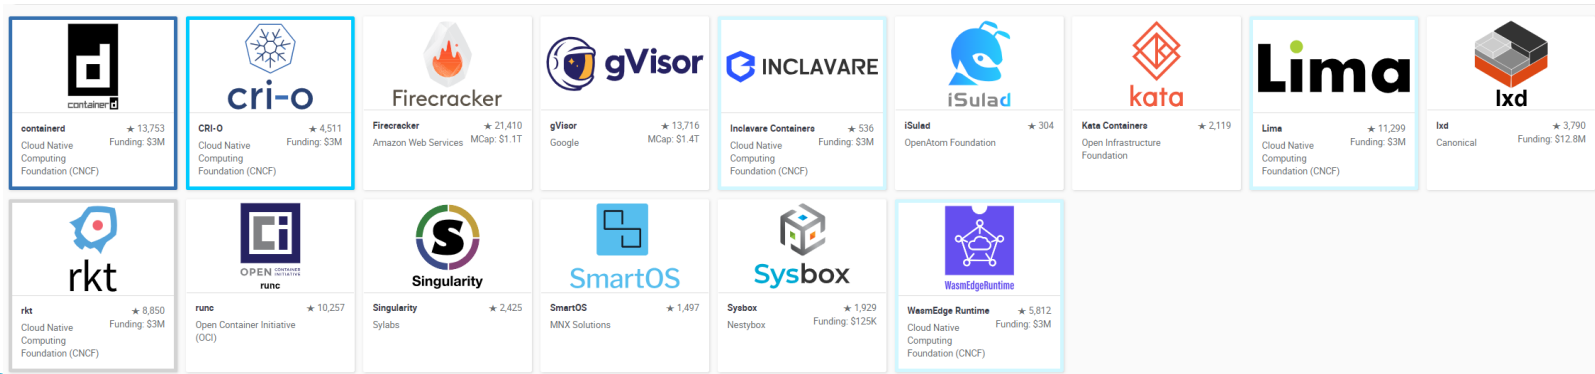
- In the scenario of uploading files from the local system to EOS,
 - **Docker-Engine** unexpectedly shows the best performance with marginal differences.
- In the scenario of downloading files from EOS to the local system,
 - **Docker-Engine** shows a notably unstable behavior and persists even in subsequent evaluations

Conclusion



- **What shows the result of this research?**
 - The **performance of various metrics changes** depending on the container runtime in a Kubernetes environment.
 - It demonstrates the need to **choose the appropriate container runtime** for the situation.
 - It provides performance evaluations to **help users choose the appropriate container runtime**.
- **What is future plans?**
 - we will do performance analysis with **more rigorous evaluation using precise criteria**.
 - We will **analyze more container runtimes** with various benchmark tools to provide users with practical performance evaluations.

[Runtime – Container Runtime, CNCF Cloud Native Interactive Landscape]



Thank you!

Jung-Bin Kim

jungbin.kim@cern.ch

2023.11.01(Wed)