



KubeCon



CloudNativeCon

Europe 2019

**Kubernetes + Encrypted Memory =
Security * Privacy**

Disclaimer



KubeCon



CloudNativeCon

Europe 2019

Please Note:

- IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice and at IBM's sole discretion.
- Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision.
- The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract.
- The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.
- Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.

Who Are We?



KubeCon



CloudNativeCon

Europe 2019



Harshal Patil



harche



Pradipta Banerjee



@pradipta_kr

<https://medium.com/@pradipta.banerjee>

www.cloudgeekz.com

Agenda



KubeCon



CloudNativeCon

Europe 2019

- Securing Data
- Introducing Memory Protection
- Kubernetes Integration

How Do We Secure Data and Code?



KubeCon



CloudNativeCon

Europe 2019

TLS/HTTPS

Data in Transit



Data at Rest



Data in Use ??

- From other software
- Malicious Admins
- Compromised host/hypervisor

What is being done?



KubeCon



CloudNativeCon

Europe 2019

IBM Power

- Secure VM and Protected Execution Facility (PEF)

Intel

- *SGX*
- Total Memory Encryption - TME/MKTME

AMD

- Secure Memory Encryption
- Secure Encrypted Virtualization (SEV)



Create a black box in memory

- Stuffs inside the black box is protected from anything that is outside.

Secure VM and Secure Containers

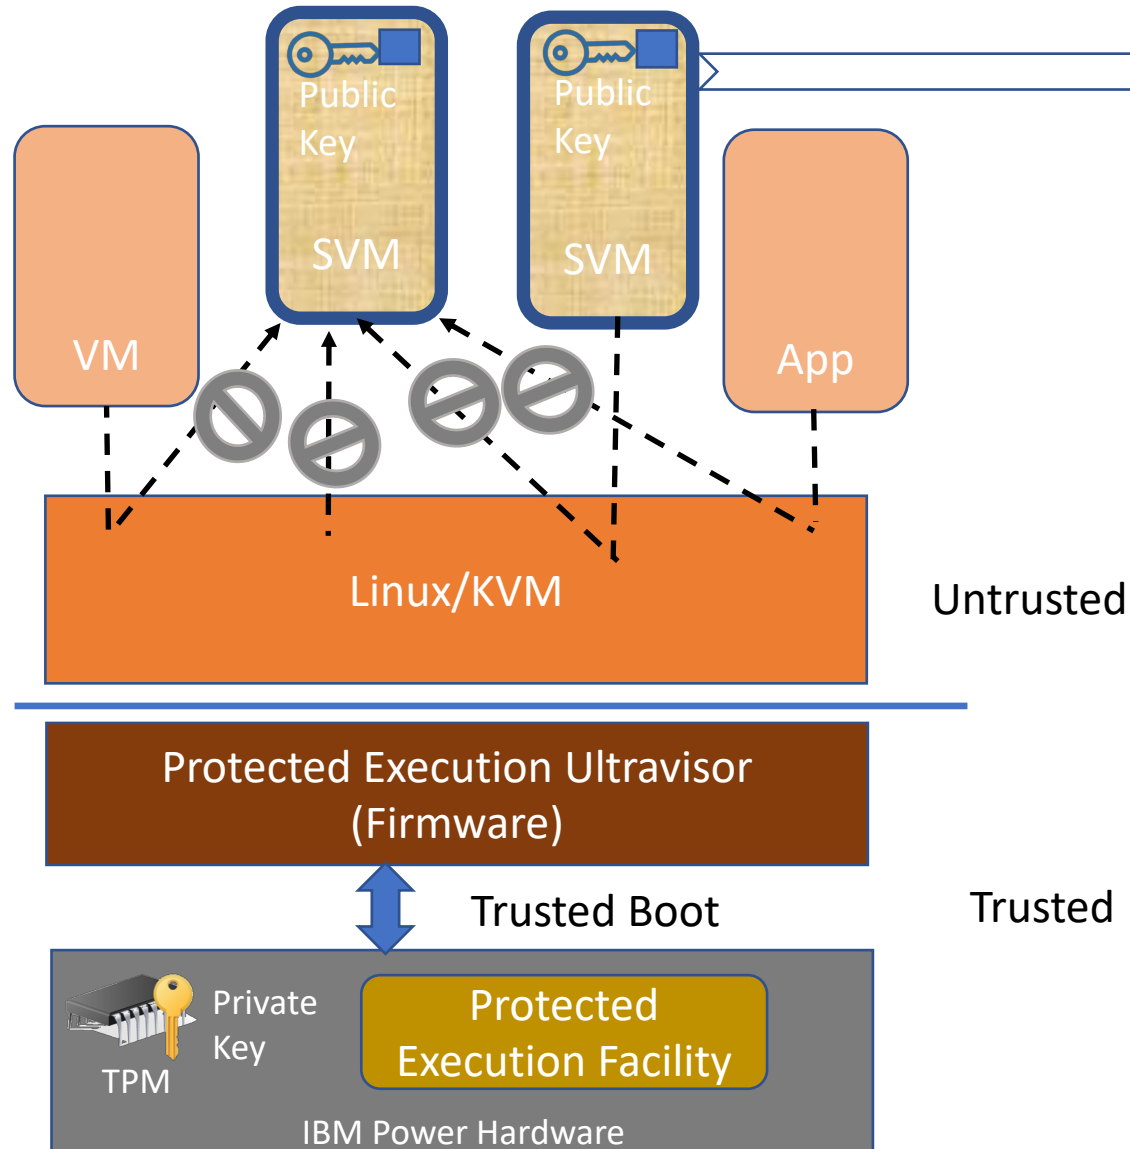


KubeCon



CloudNativeCon

Europe 2019



Secure Containers

- SVM Image = Encrypted RootFS + Lock Boxes + Encrypted Secrets
- Secure Container = SVM Image + Container RootFS
- Encryption Key (for rootfs, secrets) and integrity info put in Lock Box
- Lock Box is wrapped using system public key
- No code changes needed for application or container

Ref: <https://ibm.co/2DOL7LJ>

How can we use it with k8s?



KubeCon



CloudNativeCon

Europe 2019

Leveraging Kata container runtime

- Kata launches Secure containers (SVM + container)
- Aspects related to ephemeral volumes, extraction of container image etc needs to be handled

Secure Containers with Kata

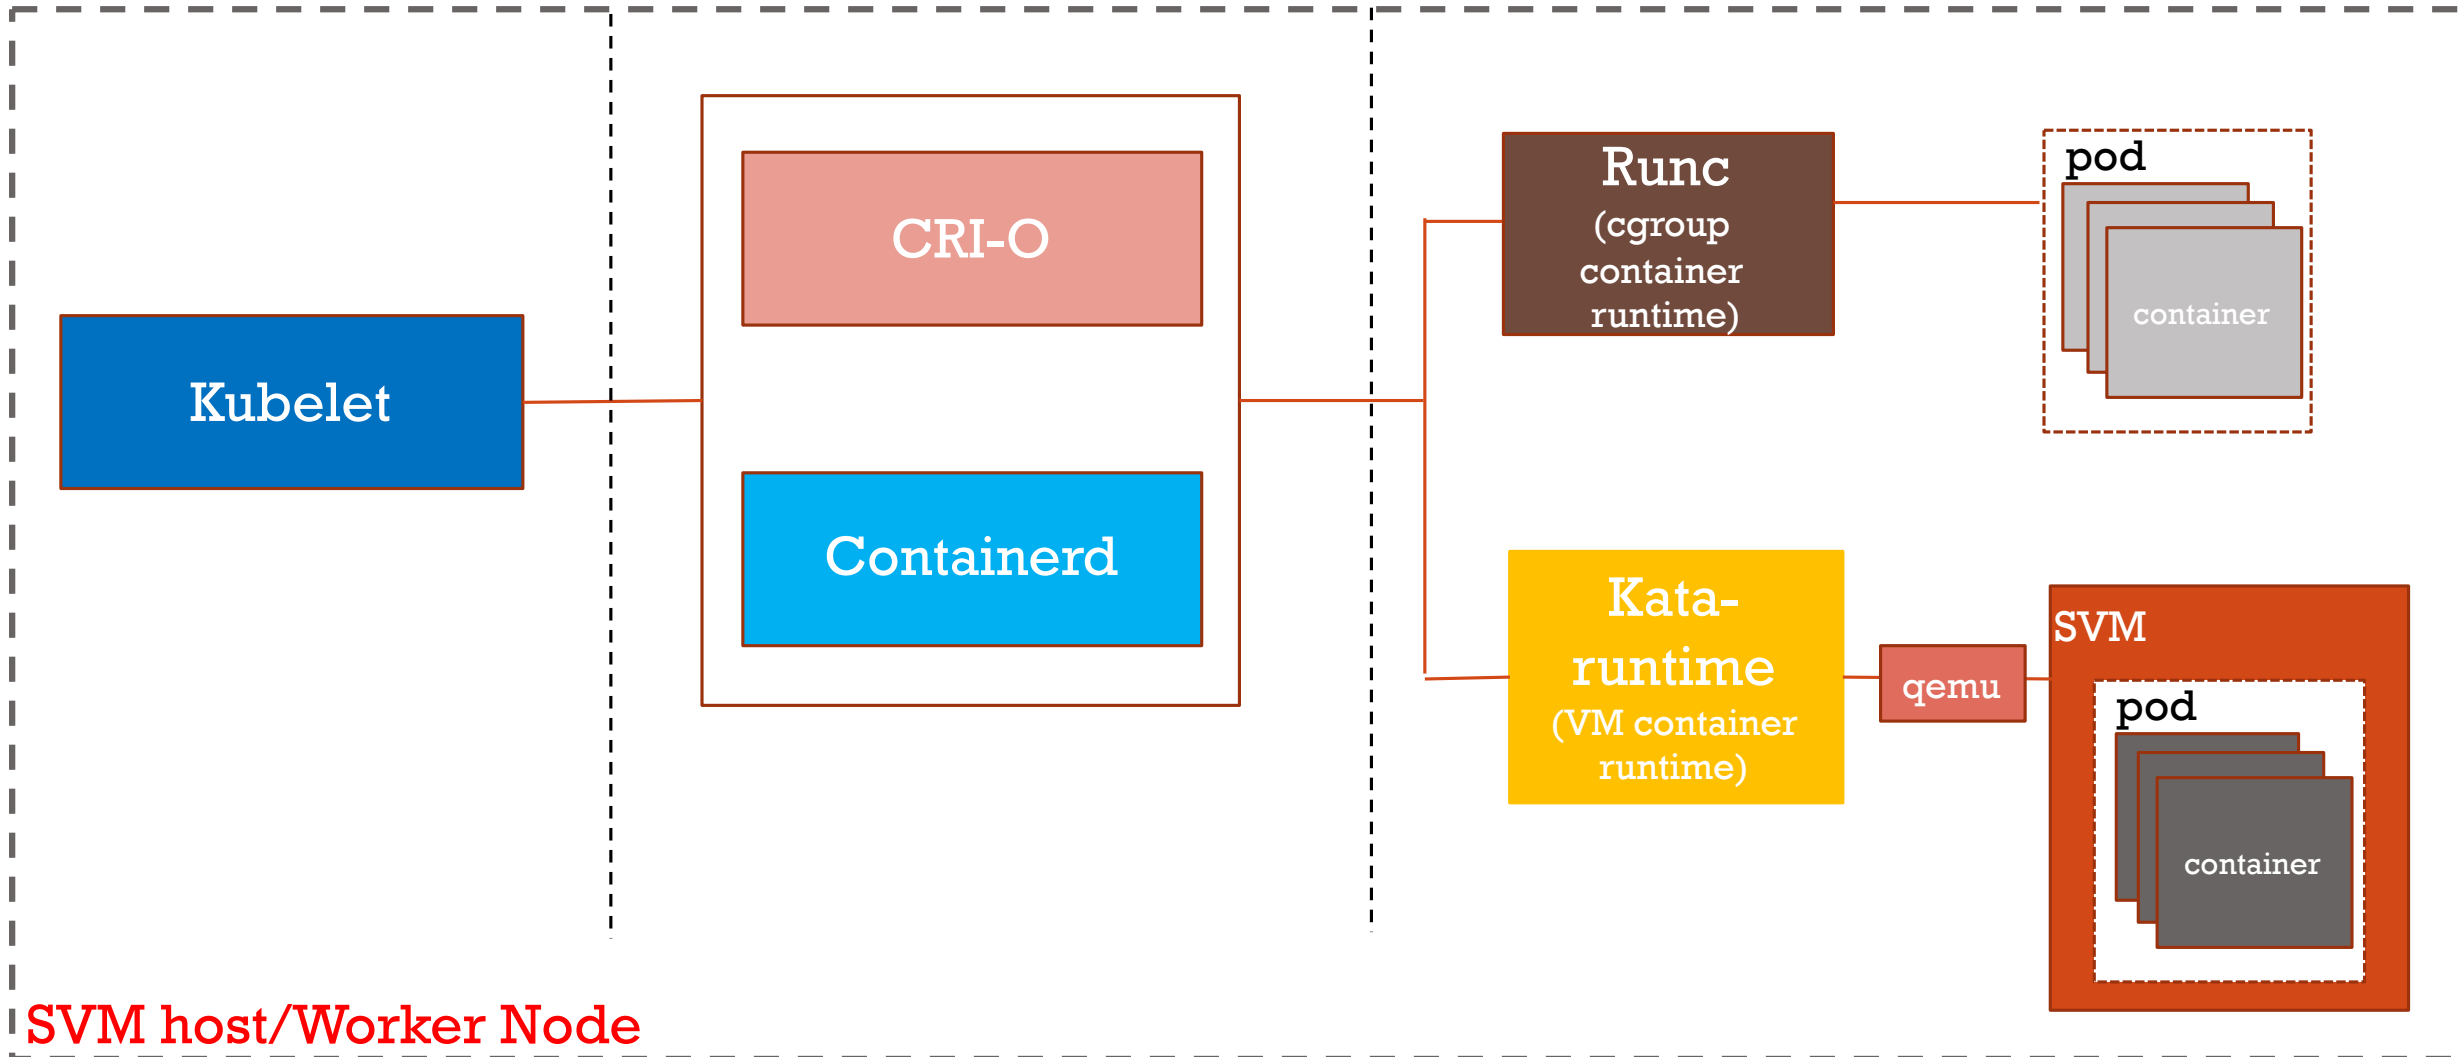


KubeCon



CloudNativeCon

Europe 2019



Kata Containers



KubeCon

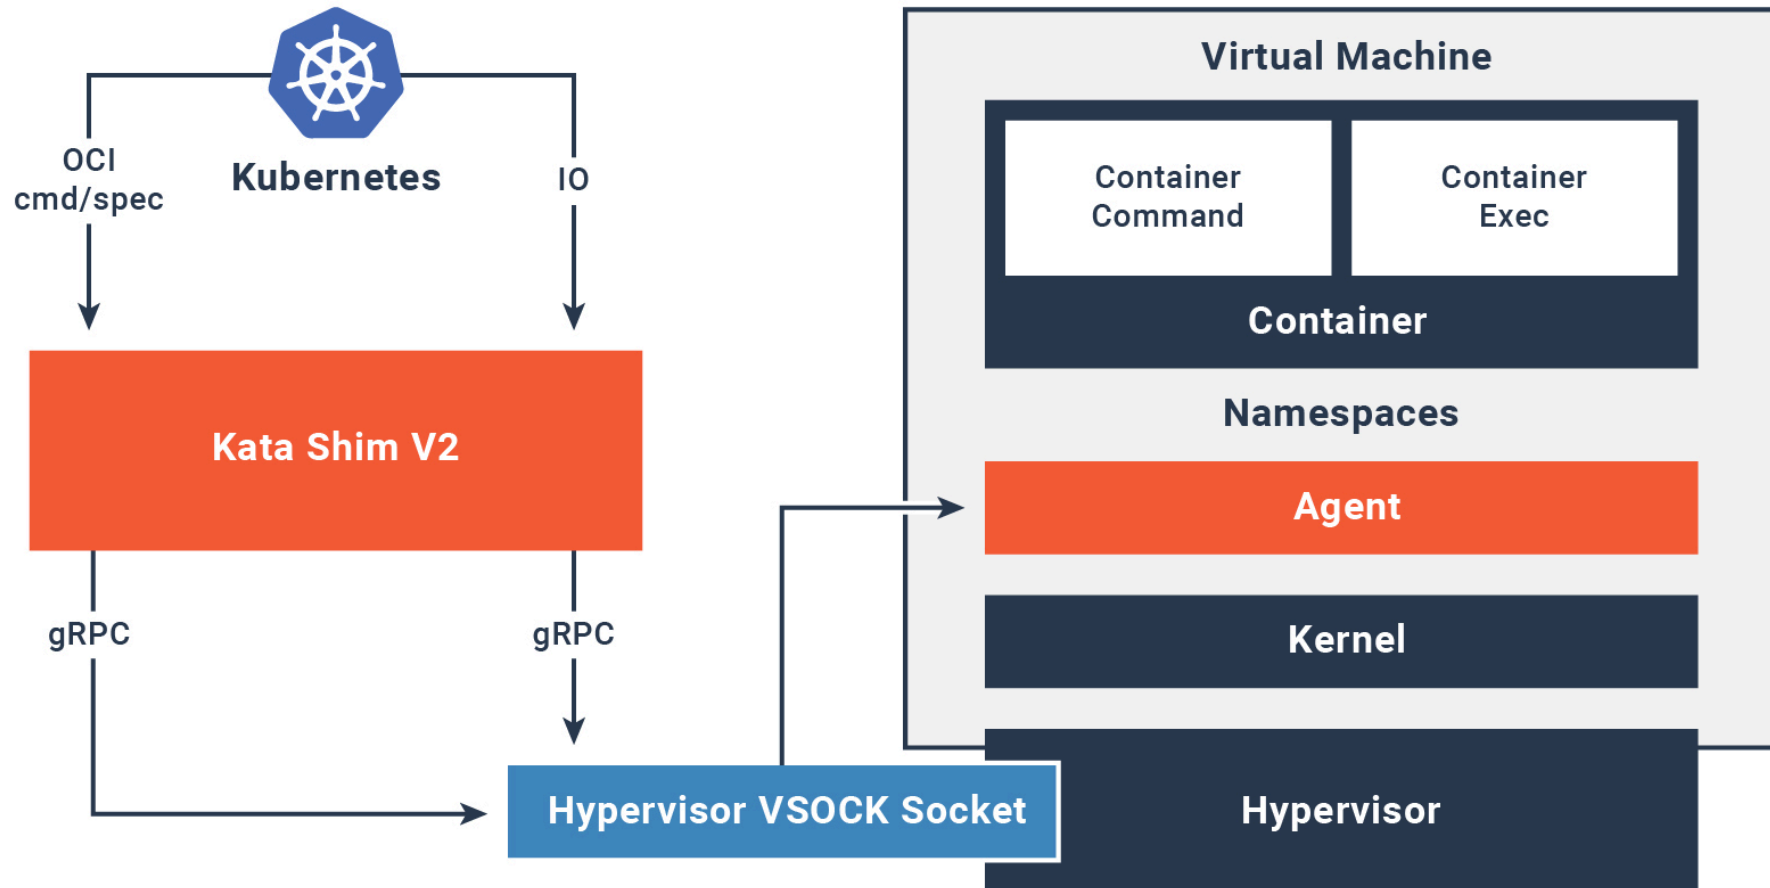


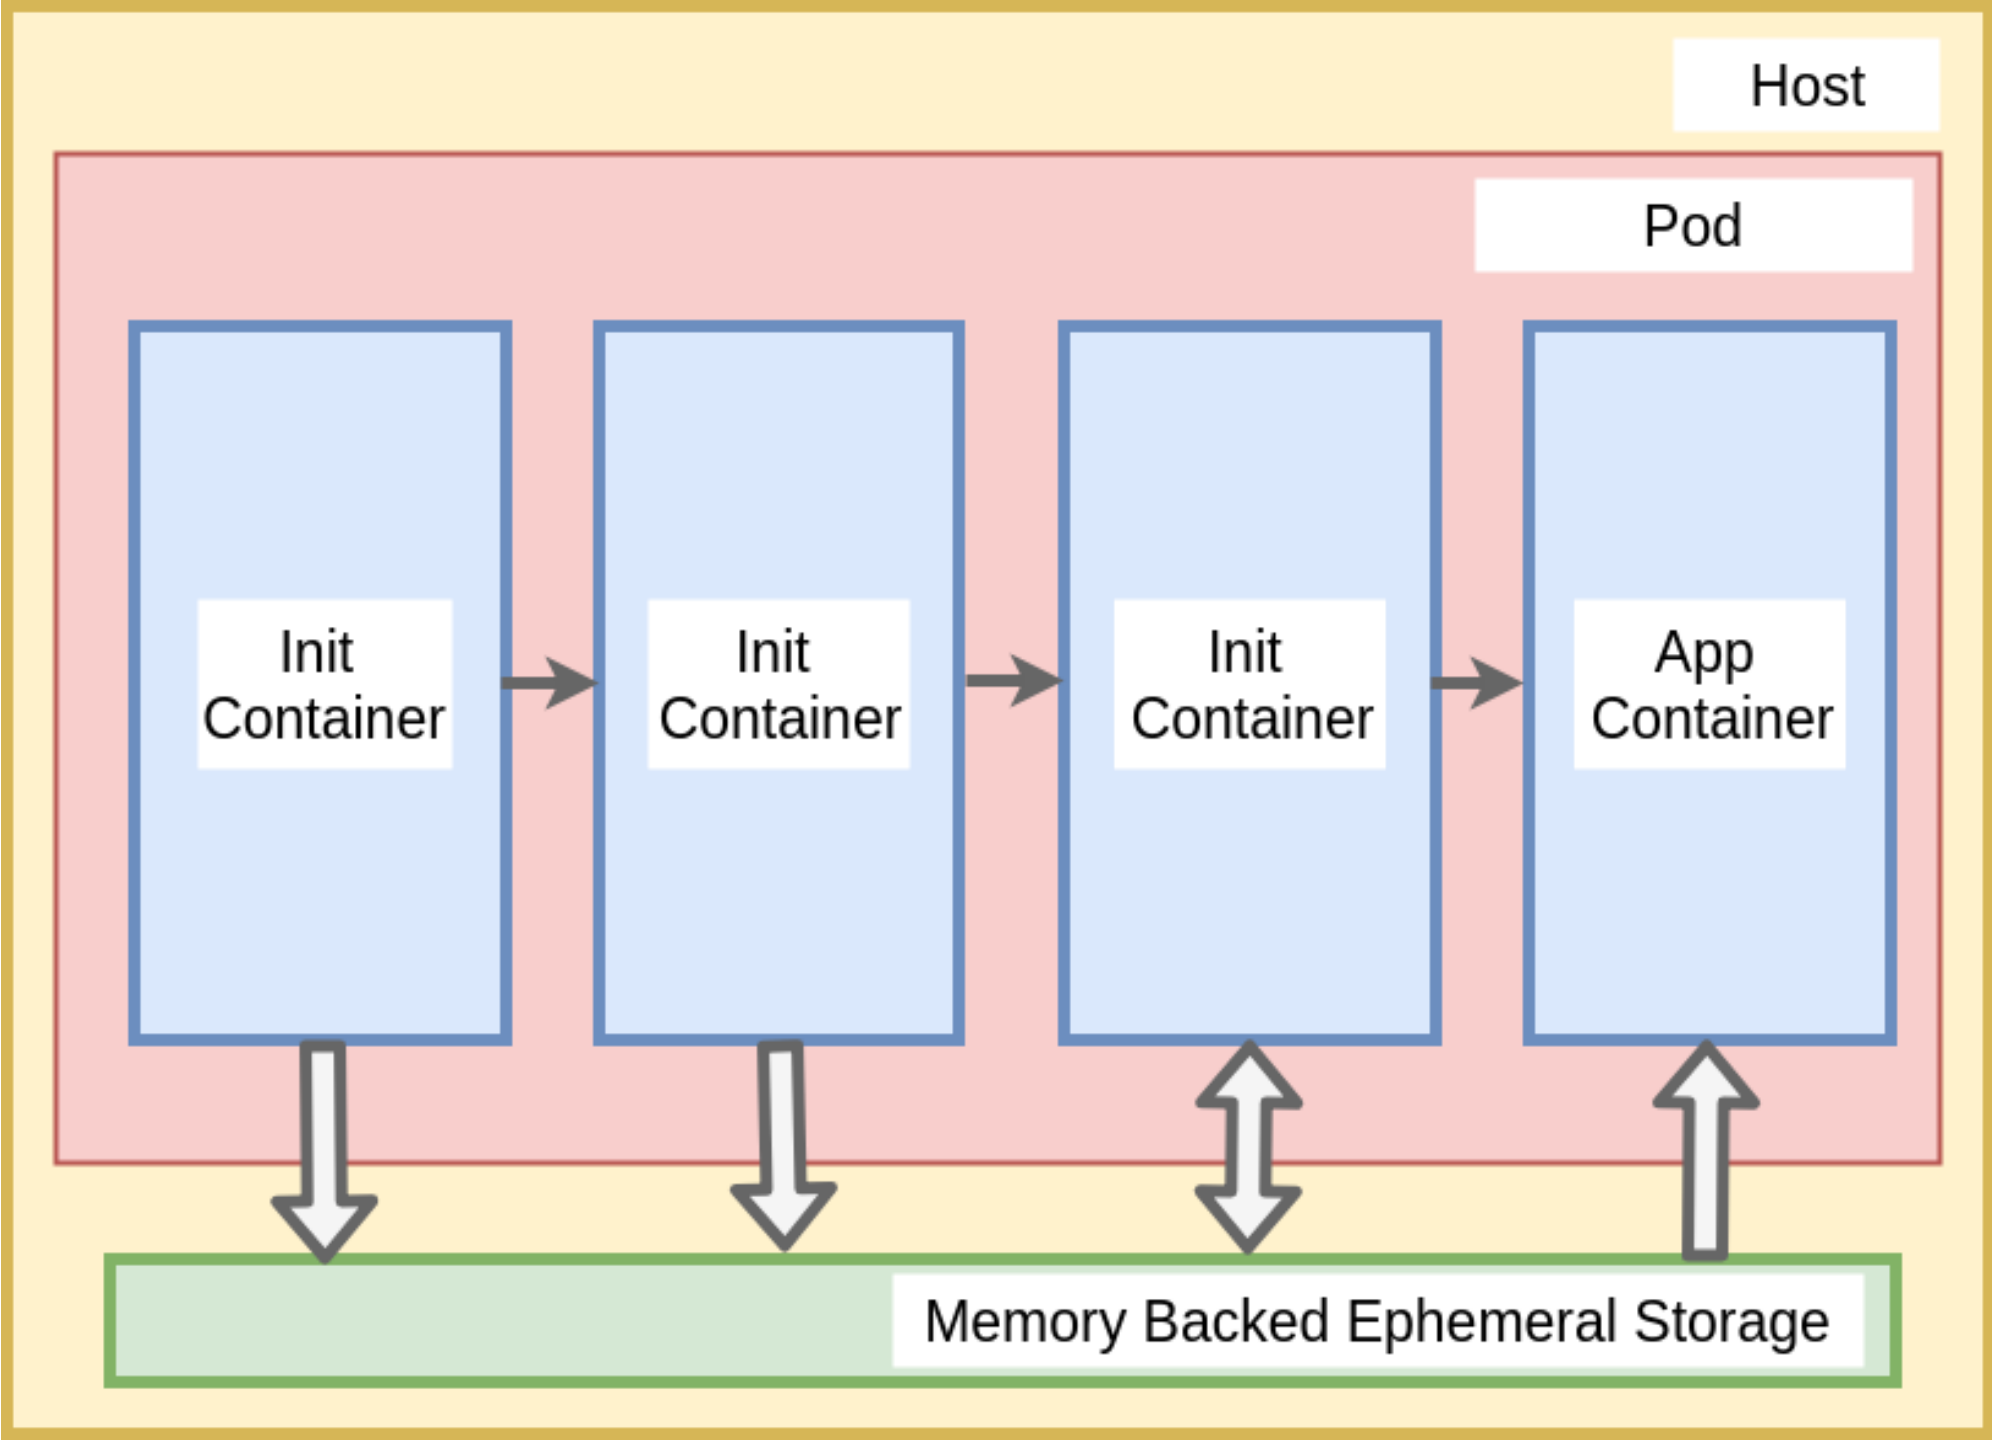
CloudNativeCon

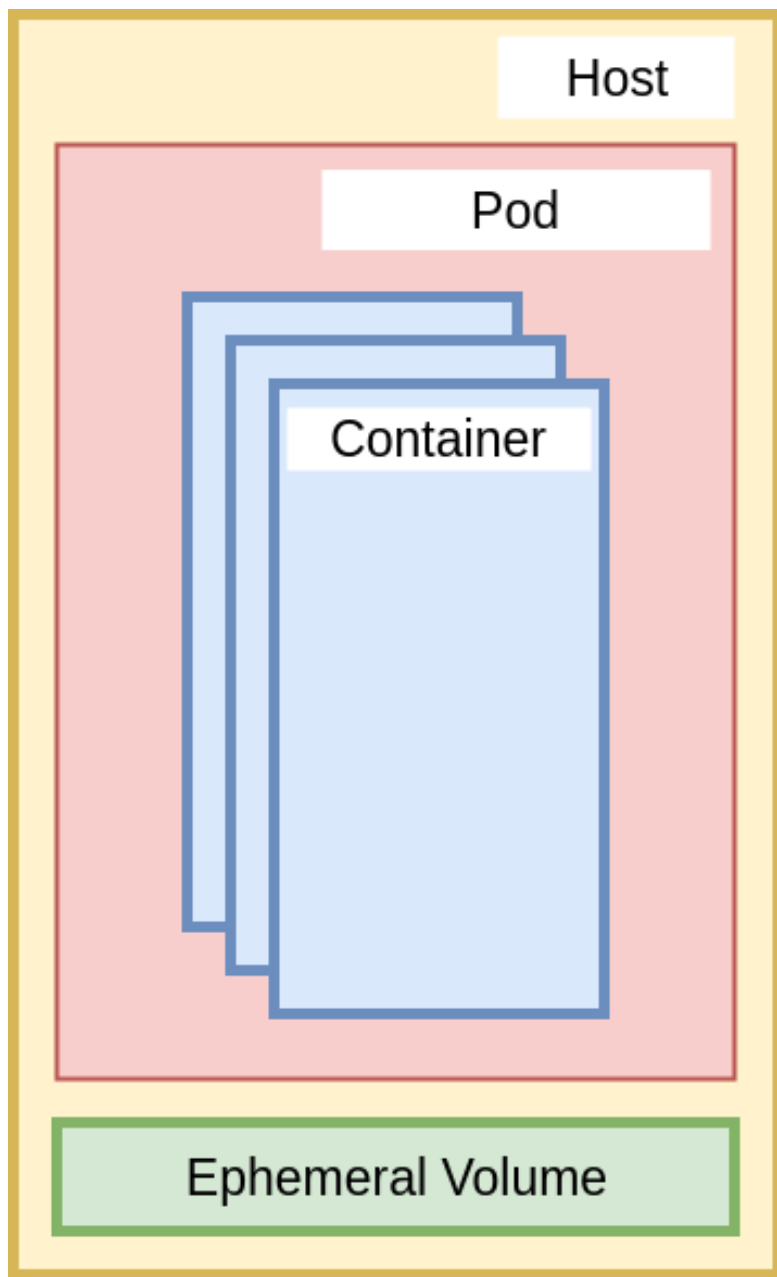
Europe 2019



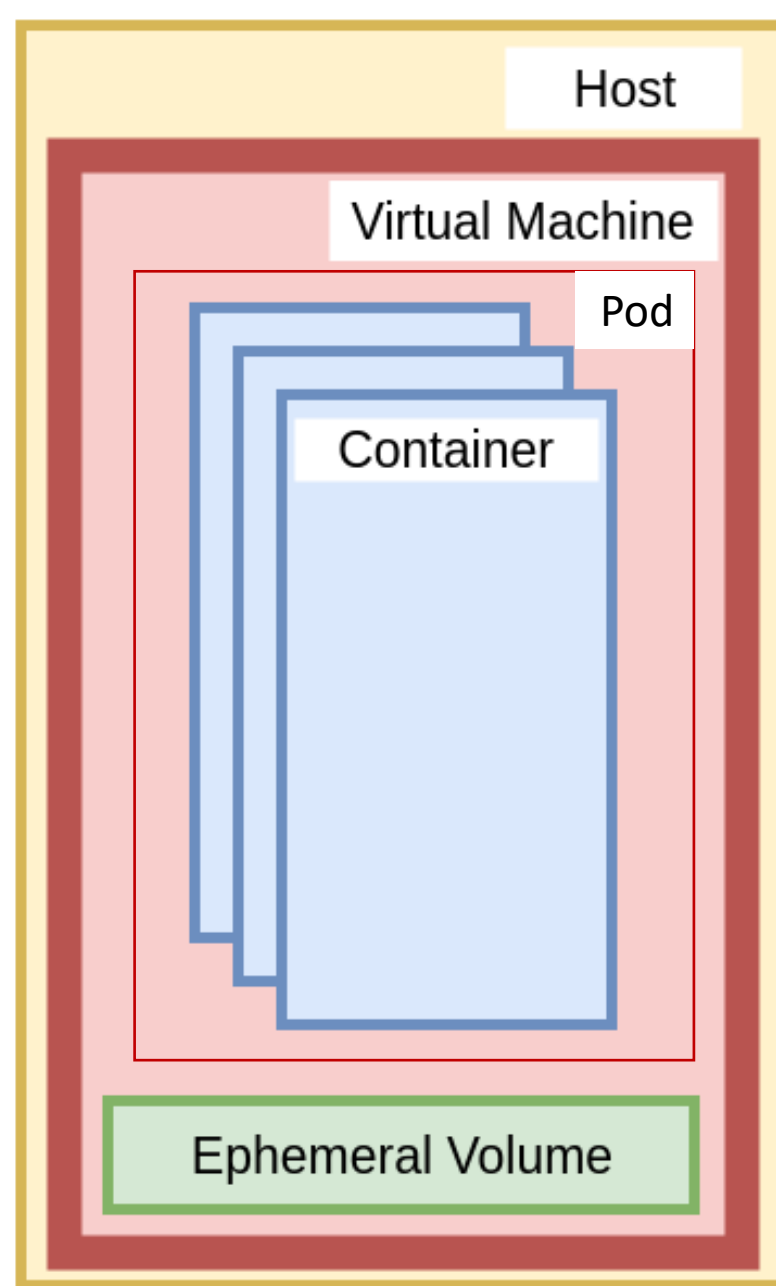
Kata Shim V2







runc




Kata

Host

SVM Protected Memory Pages

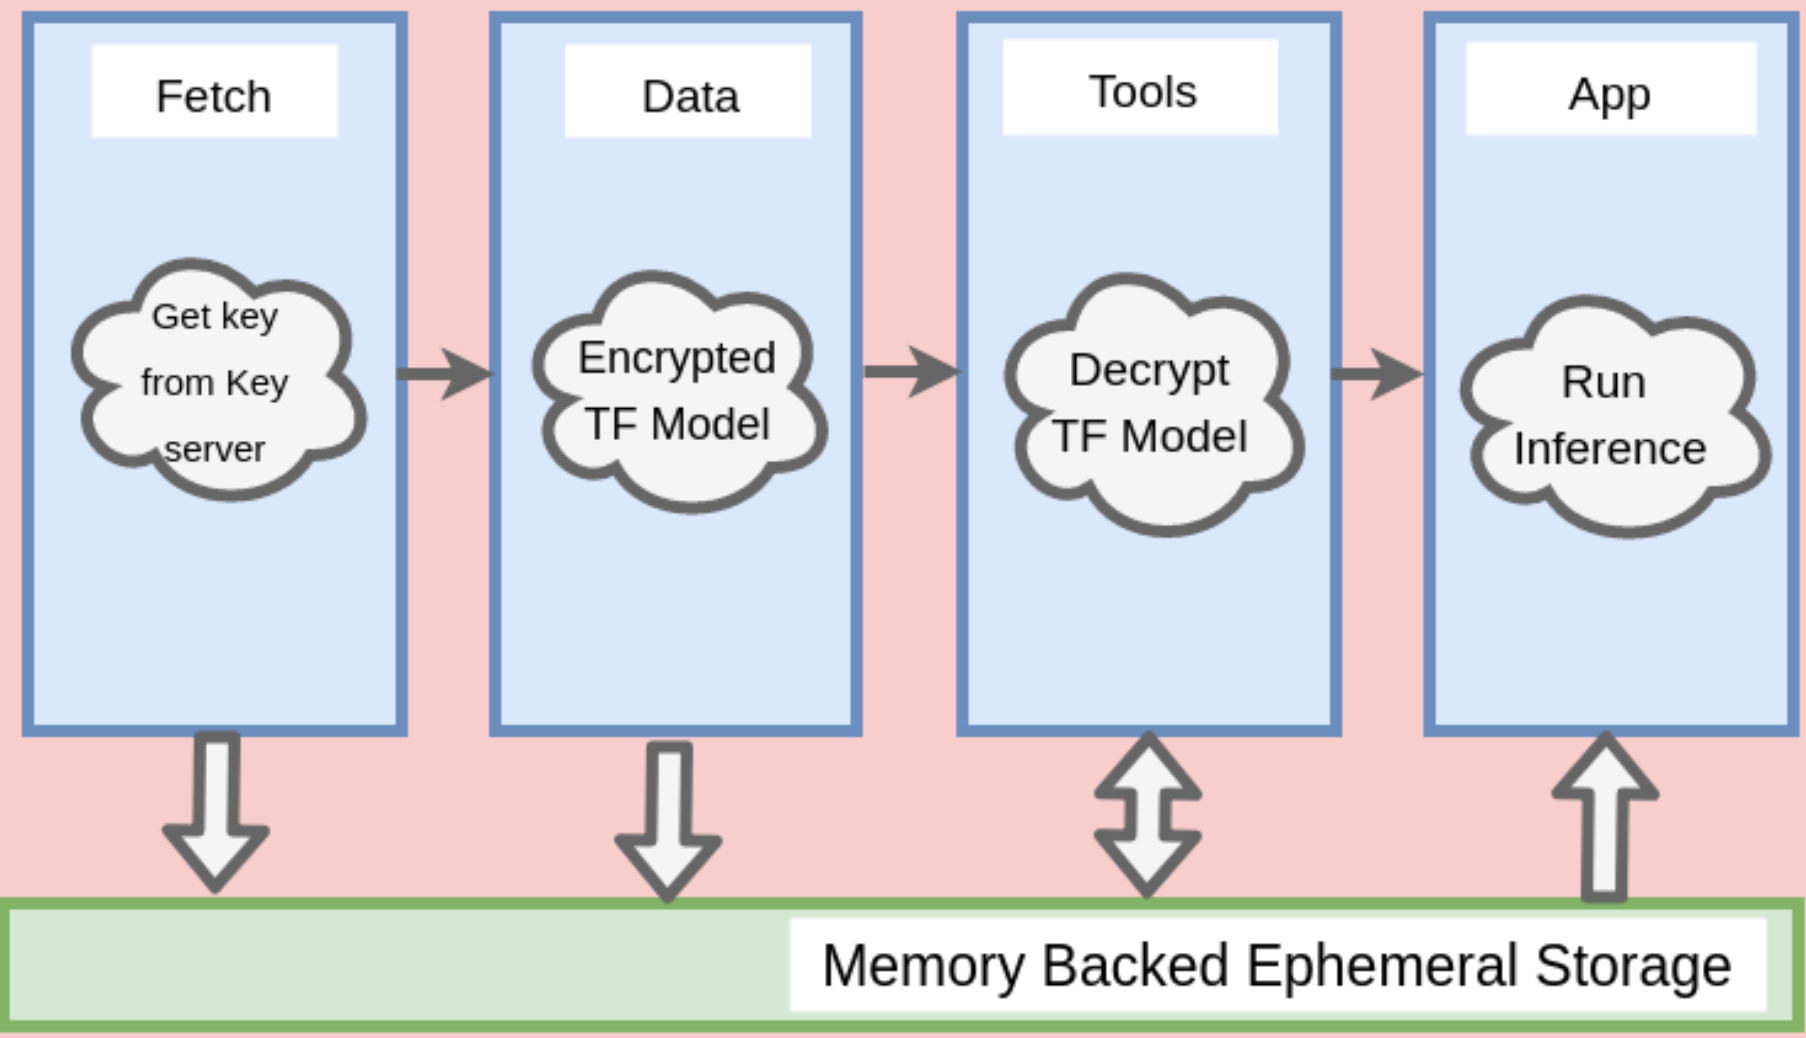
Virtual Machine



Host

SVM Protected Memory Pages

Virtual Machine



Demo




```

last-ports: 8/10
command:
  /usr/bin/tensorflow_model_server
  --port=8888
  --model_name=inception
  --model_base_path=/models/inception
state: Running
started: Tue, 18 Sep 2018 18:34:56 -0500
ready: True
restart-count: 0
environment: {}
mounts:
  /models/inception from cache-volume [rw]
  /var/run/secrets/kubernetes.io/serviceaccount from default-token-kgq6 [ro]
conditions:
  Type           Status
  Initialized     True
  Ready           True
  ContainersReady True
  PodScheduled    True
volumes:
  cache-volume:
    type: EmptyDir (a temporary directory that shares a pod's lifetime)
    medium: Memory
  default-token-kgq6:
    type: Secret (a volume populated by a Secret)
    secretName: default-token-kgq6
    optional: false
QoS Class: BestEffort
nodeSelectors: {}
tolerations:
  node.kubernetes.io/not-ready: NoExecute for 300s
  node.kubernetes.io/unreachable: NoExecute for 300s
events:
  Type       Reason      Age    From          Message
  ---
  Normal    Scheduled   15s    default-scheduler    Successfully assigned default/model-serving to 127.0.0.1
  Normal    Pulled     14s    kubelet, 127.0.0.1   Container image "ghansal/tfmodel:latest" already present on machine
  Normal    Created    14s    kubelet, 127.0.0.1   Created container
  Normal    Started    14s    kubelet, 127.0.0.1   Started container
  Normal    Pulled     13s    kubelet, 127.0.0.1   Container image "ghansal/tfutils:latest" already present on machine
  Normal    Created    13s    kubelet, 127.0.0.1   Created container
  Normal    Started    13s    kubelet, 127.0.0.1   Started container
  Normal    Pulled     8s     kubelet, 127.0.0.1   Container image "ghansal/tfutils:latest" already present on machine
  Normal    Created    8s     kubelet, 127.0.0.1   Created container
  Normal    Started    8s     kubelet, 127.0.0.1   Started container
  Normal    Pulling    7s     kubelet, 127.0.0.1   pulling image "ghansal/tferving:latest"
  Normal    Pulled     6s     kubelet, 127.0.0.1   Successfully pulled image "ghansal/tferving:latest"
  Normal    Created    6s     kubelet, 127.0.0.1   Created container
  Normal    Started    5s     kubelet, 127.0.0.1   Started container

```

But...

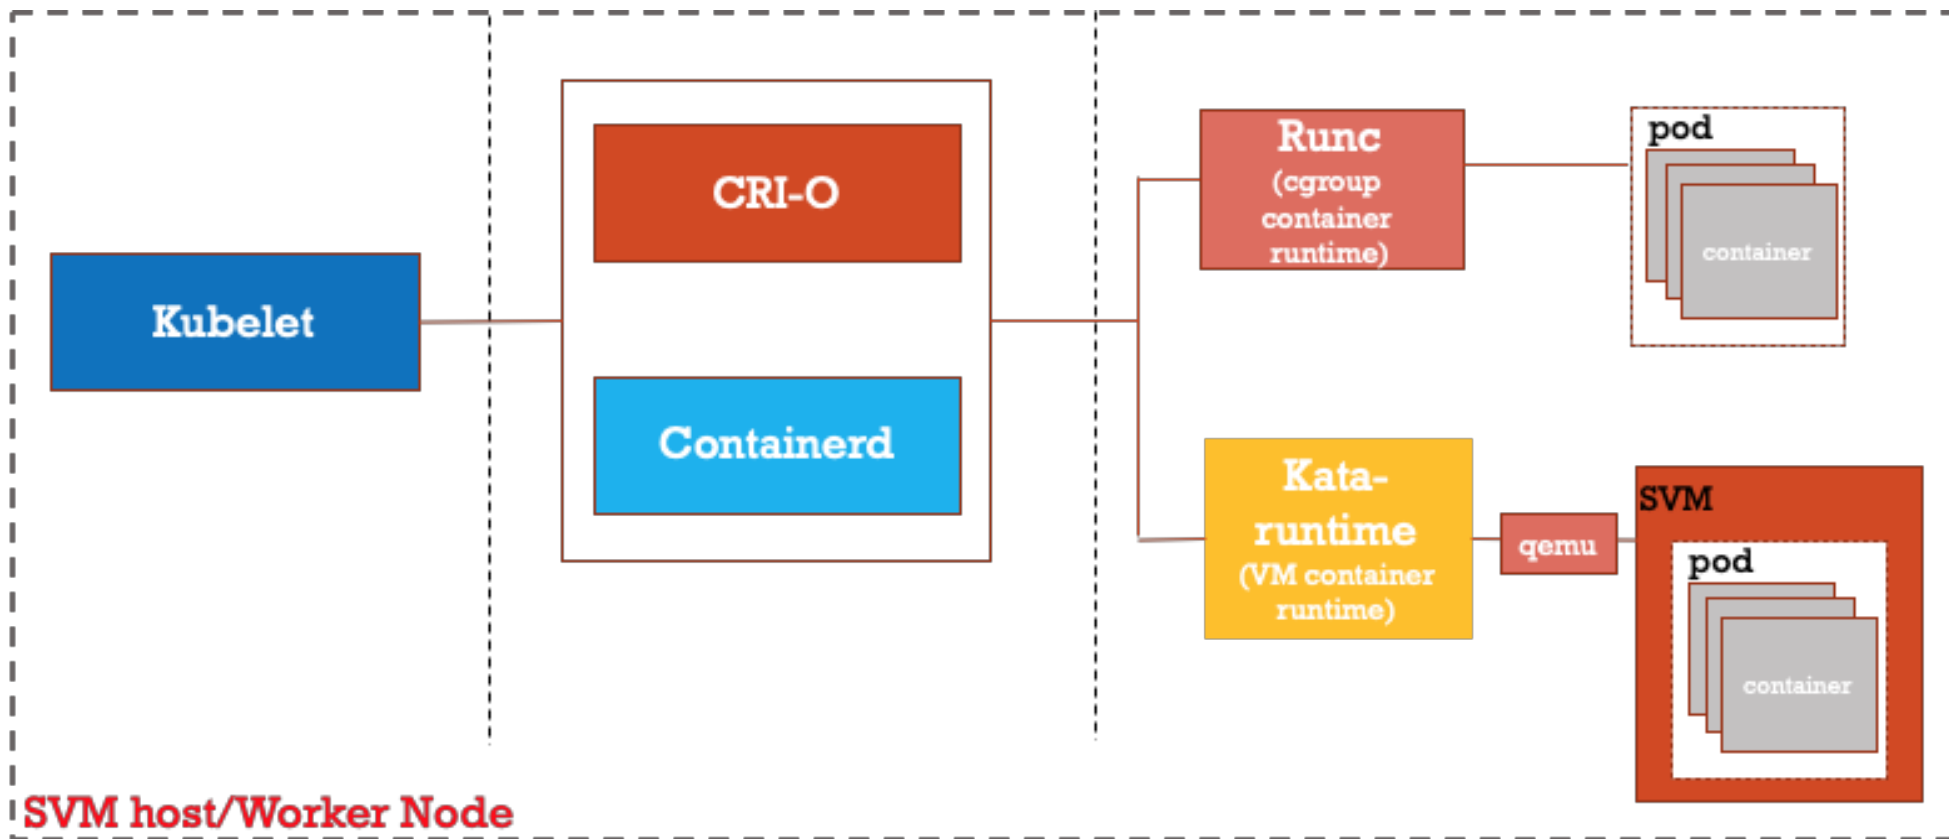


KubeCon



CloudNativeCon

Europe 2019



Registry is not encrypted

Images are extracted on the host

Work in Progress



KubeCon



CloudNativeCon

Europe 2019

Encrypt Container images

- Ongoing effort to bring encryption to container images
- Presented in DockerCon 2019 - <https://bit.ly/2LQhq3v>
- KEP with Kubernetes community to add support for Encrypted Container Images
- Join us in Kubecon Shanghai 2019 where we will talk in detail

Enable the OCI runtime to pull Images

- Directly inside the confines of the SVM

Kata support for EmptyDir type volumes of k8s

- <https://github.com/kata-containers/runtime/issues/61>

Blog post

- <https://mawacake.blogspot.com/2018/09/trust-tensorflow-and-cloud.html>

Kubernetes KEP for Image Encryption

- <https://github.com/kubernetes/community/issues/2970>

IBM Power Protected Computing

- <https://developer.ibm.com/articles/l-support-protected-computing/>

Acknowledgements



KubeCon



CloudNativeCon

Europe 2019


Community

- Containerd
- Kata
- Kubernetes
- Linux
- Qemu

Teams

- IBM Cognitive Systems
- IBM Linux Technology Center
- IBM Research

Many others...




KubeCon



CloudNativeCon

Europe 2019

Thank You



KubeCon



CloudNativeCon

Europe 2019

Back up

Prepare the Images

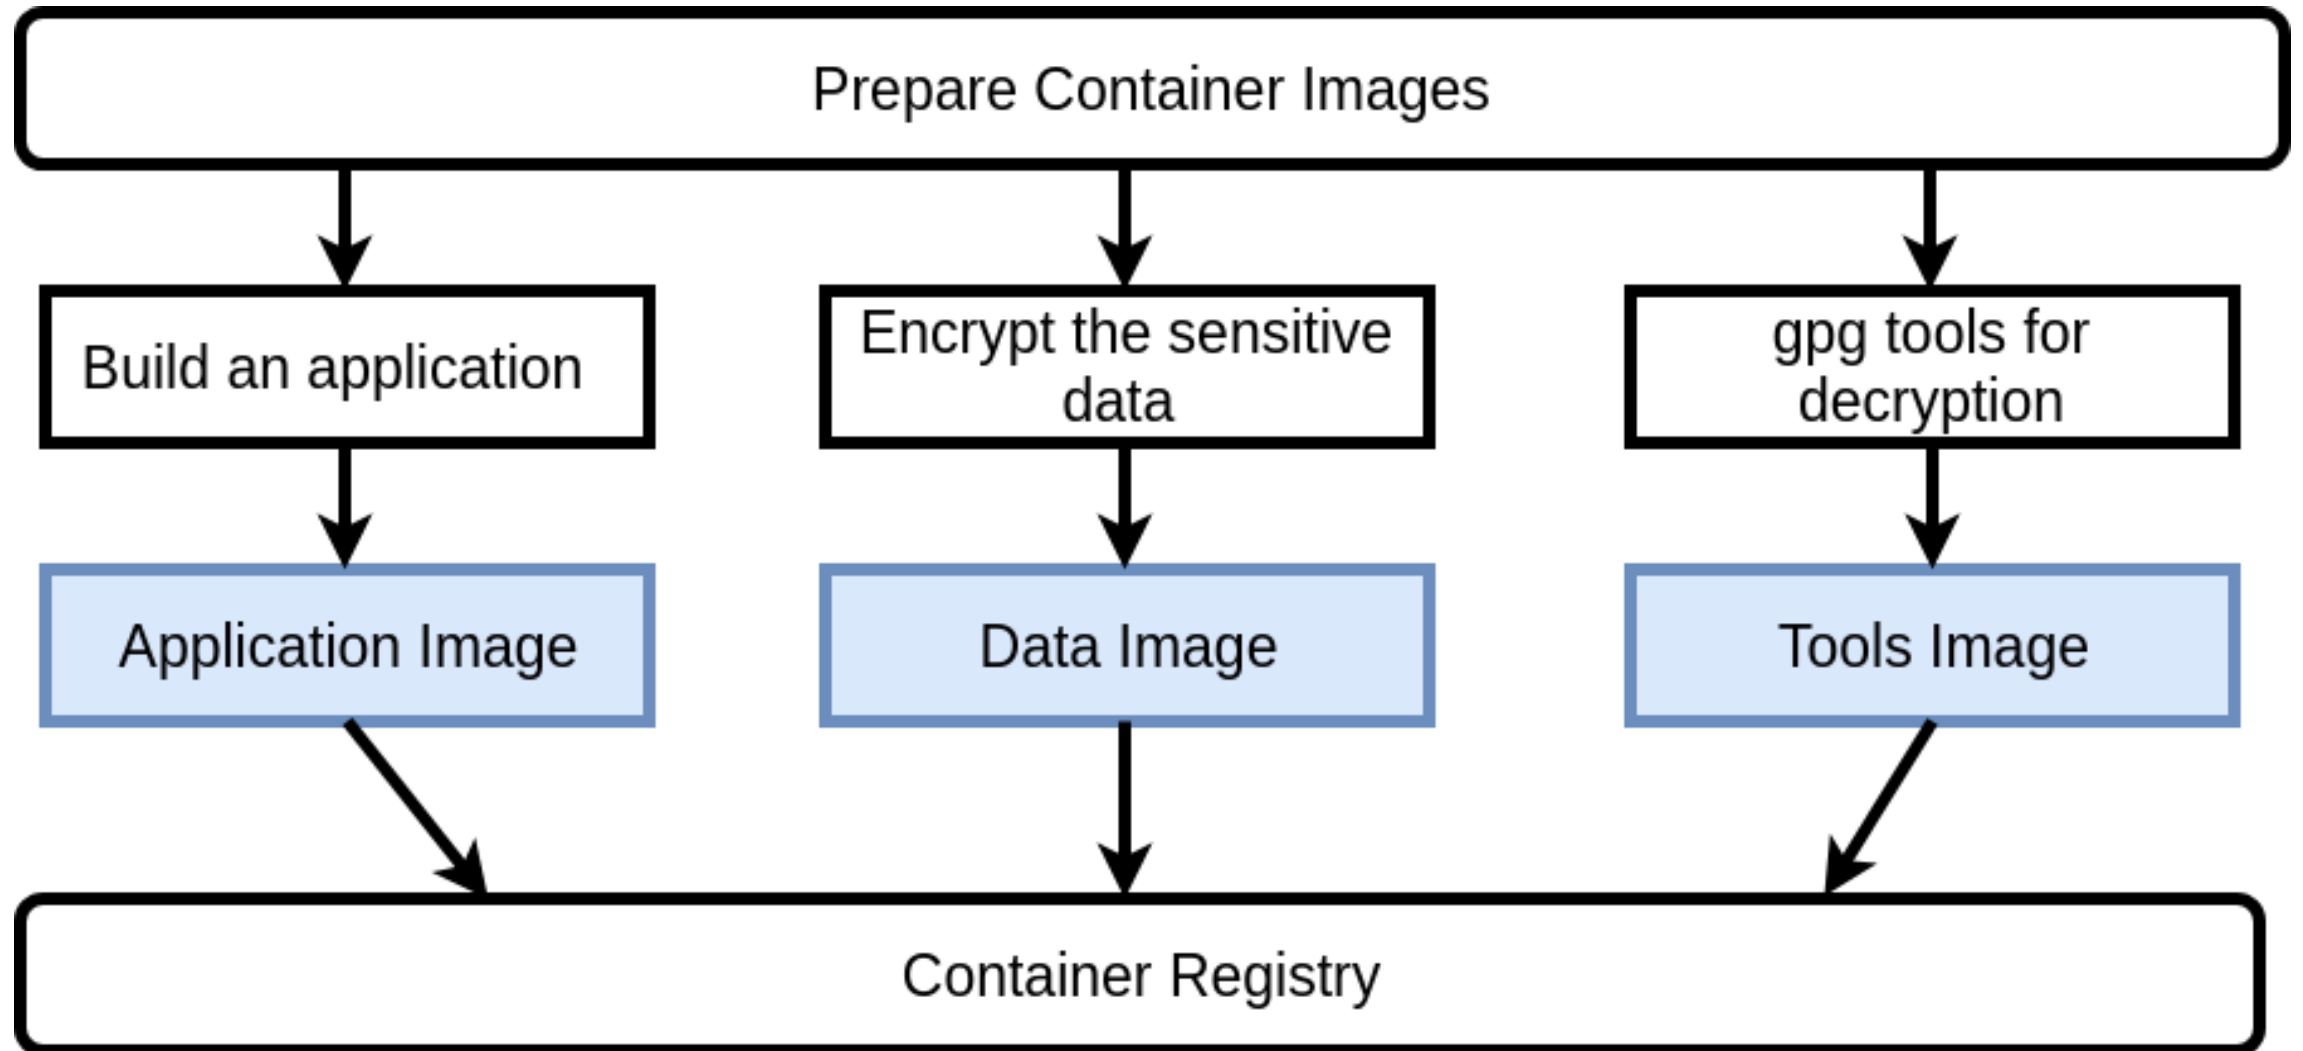


KubeCon



CloudNativeCon

Europe 2019



Root User in the Cloud Systems

- Malicious root can snoop on all containers

System Vulnerabilities Can Lead to Privilege Escalation

- In Multi-tenant environment this could lead to snooping on unauthorized containers
- RunC Vulnerability (CVE-2019-5736)
- Dirty COW(CVE-2016-5195)

Conflict of Interest

- What if your Cloud Provider is also your competitor