

Spice Up Your Workloads With Kata Containers

Ricardo Aravena

branch.io

[@raravena80](https://twitter.com/raravena80)



Who Am I?



Work @ Branch Metrics


Cloud Ops

Kata Containers contributor

Slack: rico



Containers @ Branch



Outline

Kata

- What
- Why

Workloads

- API Services
- Databases
- Credentials Store
- Big Data

Kubernetes


- Containerd
- CRI-O
- Docker
- Serverless

Future


- Hotplug
- Hypervisors
- GPUs
- Blockchain




Traditional Containers



Speed vs Security




Kata Containers




How does it work?

<http://bit.ly/kata-containers-2-pager>



Traditional Containers
Prone to exploits



Kata Containers
Each container or container pod isolated in its own lightweight VM

Requirements

Machine	• Bare Metal • Nested Virtualization
OS	• Linux
Public Cloud	• GCP - Nested • Azure - Nested • AWS - i3.metal instances
Private Cloud	• Openstack • Bare-metal & Nested Virt Providers
Platform	• Kubernetes - CRIO/Containerd • Docker



Installation

Requirements

- `$ kata-runtime kata-check`

Kubernetes

- Kata Deploy @egernst
- github.com/egernst/kata-deploy

Docker

- Kata Installation guide



Kata Workloads

APIs or Microservices

Credentials Store

Databases

Kubernetes

Big Data & Analytics



Workloads

APIs or Microservices



Microservices

Backend for web apps

- React.js 
- Angular 
- Vue.js 

MVCs


- Rails (Ruby) 
- Django (Python) **django** 
- Play! Framework (Java, Scala) 
- Laravel (PHP) 

High performance APIs


- gRPC or Rest with GoKit 
- go-micro 
- gin-gonic 



Microservices



Microservices



Workloads

Databases



Databases

What?	• Relational • NoSQL • Containerized
Why?	• Data Protection • Compliance • Memory Management
Advantages	• Easy to set up • Upgrades • Better resource utilization



Databases

Relational DBs

- MySQL 
- PostgreSQL 
- CockroachDB 

NoSQL DBs



- Redis 
- CouchDB 
- MongoDB 

Cluster KV DBs



- Cassandra 
- Scylla 
- Hbase 



Relational DBs



DB Cluster



cassandra

OR



SCYLLA.



Workloads

Credentials Store



Credentials Store

Why?	• Centralized credentials • Microservices auth • Compliance
Encryption	• In transit • At rest
Solutions	• Vault (Hashicorp) • Knox (Pinterest) • SPIFEE/SPIRE • KeyWhiz (Square) • Confidant (Lyft)




Vault

What?	• Manage secrets • Audit logs • Microservice Auth
Architecture	• Redundant • Active/Standby • Client Server
Features	• Multiple Storage Options • API Keys • Employee credentials



Credentials Store



Demo



Demo

Vault cluster in Kata

Create Credentials

**Retrieve
Credentials**

Security in VM



Demo

Vault with Consul

<https://github.com/raravena80/oss-vaultsa>

Vault Operator

<https://github.com/kubevault/operator>



Workloads

Kubernetes



Kubernetes

How?

- CRIO 
- Containerd 

Requirements


- Bare Metal
- Nested Virtualization
- QEMU

Microservices


- Standalone
- Service Mesh
 - Istio 
 - Conduit 



Kubernetes



Kubernetes



Kubernetes CRIO

```
apiVersion: v1
kind: Pod
metadata:
  name: nginx
  annotations:
    io.kubernetes.cri-o.TrustedSandbox: "false"
  labels:
    env: test
spec:
  containers:
    - name: nginx
      image: nginx
      imagePullPolicy: IfNotPresent
  nodeSelector:
    kata-runtime: "true"
```



Kubernetes Containerd

```
apiVersion: v1
kind: Pod
metadata:
  name: nginx
  annotations:
    io.kubernetes.cri.untrusted-workload: "true"
  labels:
    env: test
spec:
  containers:
    - name: nginx
      image: nginx
      imagePullPolicy: IfNotPresent
  nodeSelector:
    kata-runtime: "true"
```



Demo



Demo

K8s Cluster

**Run Stateful App
with Containerd**

**Run Trusted With Untrusted
Workloads**

Add Sample Data



Workloads

Big Data & Analytics



Big Data & Analytics

Why?	• Containerized workloads • Tie to Kubernetes
Security	• Data In transit • Data at rest
Applications	• Spark • Hadoop • Flink
Workloads	• Streaming • Batch



Experimental



Spark & Flink Streaming


What?	• Real-time processing • Large amounts of data
Workloads	• Ingest from Kafka • Process and change data
Storage	• Object Storage • Columnar Storage (Druid/Vertica)



Experimental



Spark & Flink Streaming



Apache Kafka


What?	• Streaming Platform • Real-time Processing • Geo-replication
Architecture	• Distributed Cluster / Data Partition • Producers • Consumers
Use Cases	• Website Activity Traffic • Events Tracking • Metrics



Experimental



Big Data & Analytics



Publish

Subscribe



Apache Druid


What?	• High Performance • Distributed Data Store • Columnar Based
Architecture	• Different Node Types • Data Indexing in Hadoop or Local • Fast Real Time Queries
Features	• Different Storage Options • Data Compression • Different Data Formats (parquet, avro, etc)



Experimental



Big Data & Analytics



Future



Kata Containers

Hypervisors	• Hyper-V • VMware • Xen
Public Cloud	• AWS Nested Virtualization • More Bare Metal Offerings • GKE?, AKS?, ACS?
Features	• VM Hotplug CPUs, Memory, Network • Faster Shim Implementation



Workloads

NFV	• Multus • https://github.com/intel/multus-cni
Edge Computing	• 5G • IoT
AI/HPC	• Tensorflow / Kubeflow • GPUs • ML model training



Resources

Kata Containers

- <https://katacontainers.io>

Microservices

- <http://microservices.io/>

Hashicorp

- <https://github.com/hashicorp>

CNCF

- <https://github.com/cncf/landscape>

Kafka

- <https://kafka.apache.org/>

Druid

- <http://druid.io>



Thank You!

