Running Identity Federation Services on Containers and K8s

Muhammad Farhan Sjaugi[1,2] and Irfan Hakim AS[1]
SIFULAN Malaysian Access Federation[1]
School of Data Sciences, Perdana University[2]
farhan@sifulan.my, irfan@sifulan.my

Background

- Since 2018, SIFULAN Malaysian Access Federation (SIFULAN) has been operating in production mode by using Virtualization Machine (VM) to run several identity federation services.
- As the federation grows, SIFULAN plans to offer IdP-as-a-service to the existing and potential future members as an additional service.
- However, the current infrastructure setup has some limitations to support the plan as multi-tenant services were not part of the initial infrastructure design.
- Hence, SIFULAN migrated its federation infrastructure from a VM based to a container-based infrastructure and use Kubernetes (K8s) as the orchestration manager for the containers.

Dissecting Federation Services

- Typical federation core services:
  - Federation Manager – Jagger (a LAMP based application)
  - Metadata signer – xmlsectool (java based application)
  - Metadata repository – basic web repository (apache/nginx)
  - Discovery Services/WAYF – SwitchWAYF (php based application)

- Federation auxiliary services:
  - SAML IdP – Shibboleth/SimpleSAMLphp (java/php based application)
  - Directory Services – OpenLDAP
  - Others (e.g. filesender)

Possible Implementation of Federation Services

- Dedicated baremetal server per application = the best performance, but costly and need to manage many servers
- Dedicated VM per application = cost effective, but still need to manage many VMs and has additional performance overhead
- Dedicated Container per application = cost effective, nearline baremetal performance, but certain limitations when preparing the container apply

Linux Containers

- **LXC** (**Linux Containers**) is an operating-system-level virtualization method for running multiple isolated Linux systems (containers) on a control host using a single Linux kernel.
- Application is contained in a container image along with its runtime libraries.
- Unlike VM which has a dedicated operating system (OS) kernel, the container shares the same OS kernel with the host, hence its performance is nearline to a “baremetal” performance.
- Docker is the most popular containers platform.

Linux Containers (cont.)

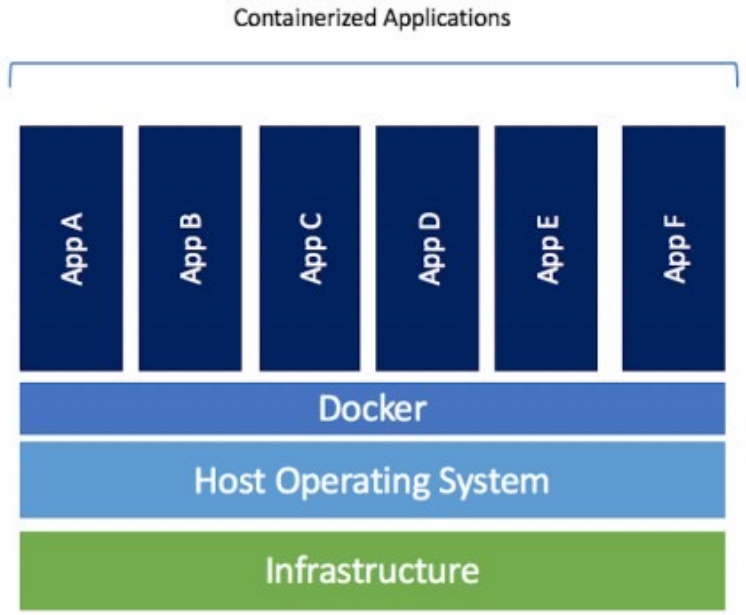

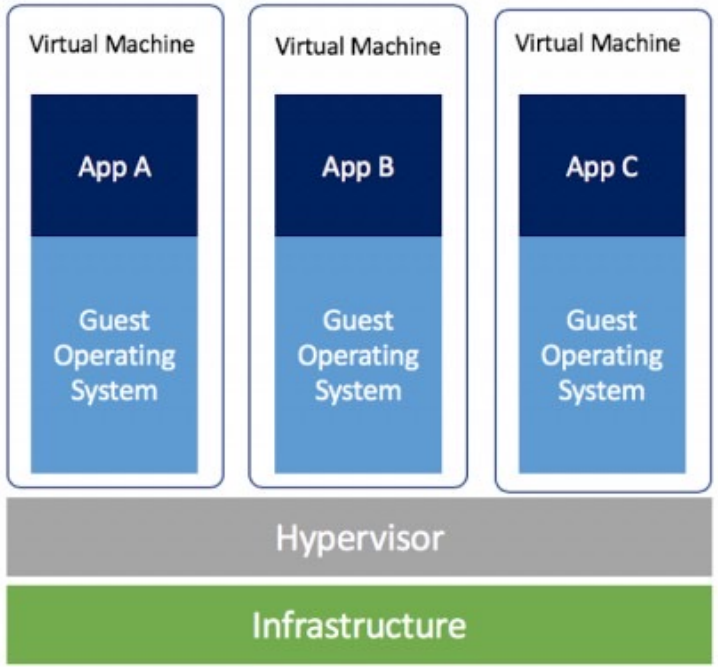

Linux Containers Operational Challenges

- Even though container is good, however the runtime APIs are well suited to managing individual containers.
- It would become a big challenge when it comes to managing applications that might comprise hundreds of containers spread across multiple hosts, like running multiple shibboleth IdP or LDAP for IdP-as-a-service services.
- Containers need to be managed and connected to the outside world for tasks such as scheduling, load balancing, and distribution.
- Therefore, container orchestration system like Kubernetes comes to rescue.

Kubernetes

- Kubernetes (K8s), developed by Google, is an open-source container-orchestration system for automating computer application deployment, scaling, and management.
- It serves as an infrastructure framework for system developer to develop and run their containerized applications.
- K8s has several great features such as load-balance, failover, high-availability, partitioning (namespaces), key-value database (configmap, secrets), and among others.
- K8s is horizontally scaled, i.e. the more worker node is added to the cluster, the bigger your K8s system capacity.

Kubernetes Architecture

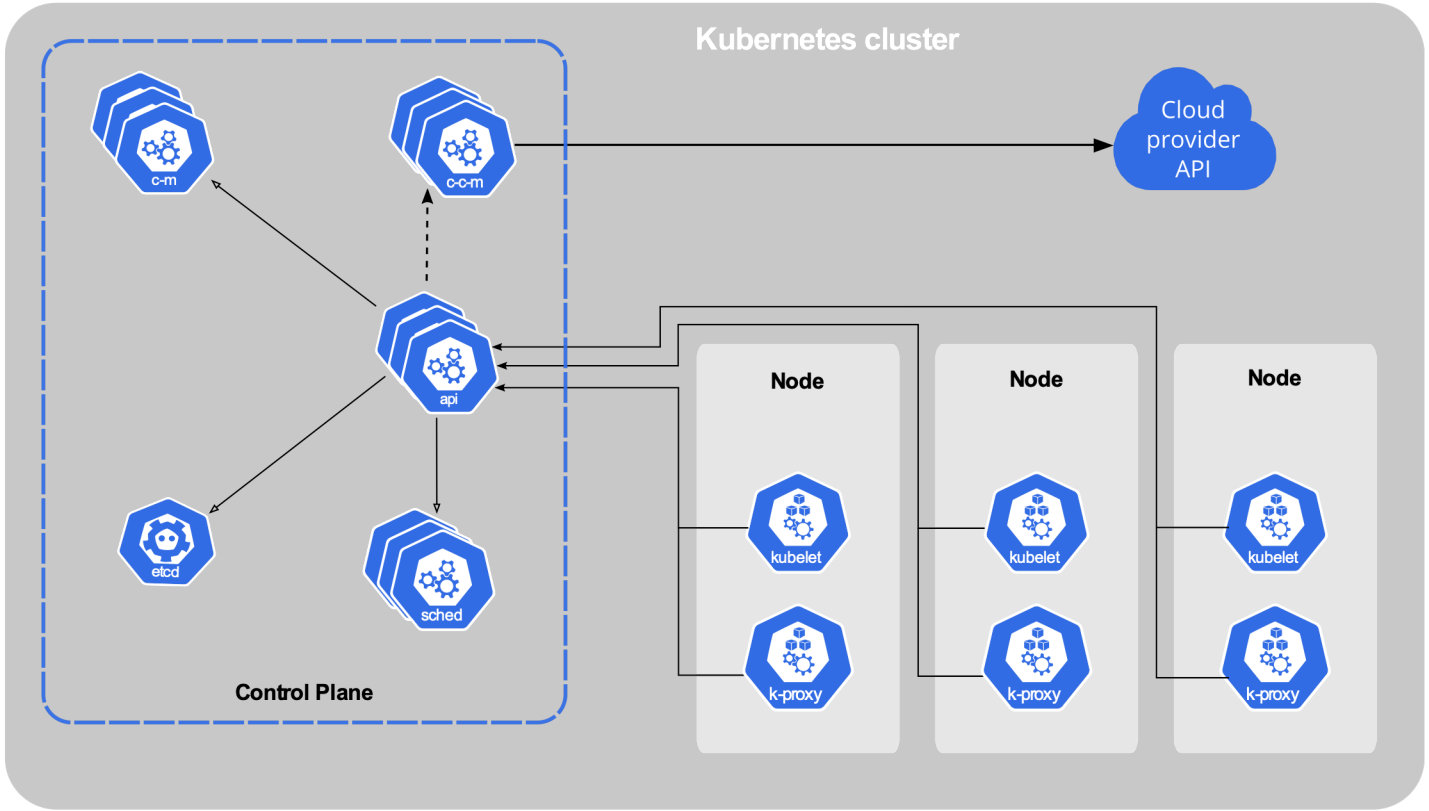

API server

Cloud controller manager ***(optional)***

Controller manager

etcd (persistence store)

kubelet

kube-proxy

Scheduler

Control plane

Node

Migrating Federation Services from VM to Containers: SIFULAN Experience

- Before:
  - Federation manager, metadata repository, and discovery service/WAYF were running on a single VM and partitioned by using Apache's VirtualHost configuration.
  - Metadata signer was running on a separate VM which only have one way access to the internet and a security token attached (Nitrokey HSM)
  - Other Aux services are running on another VM.

Migrating Federation Services from VM to Containers: SIFULAN Experience (cont.)

- After:
  - We use Rancher Kubernetes Engine (RKE) and Docker
  - Federation manager, metadata repository, and discovery service/WAYF are running on a dedicated container.
  - This separation allows each service to scale-out according to their need without affecting other service.
  - Metadata signer is running on a dedicated container. However, the digital certificates are stored as a Secrets object instead inside the HSM.
  - Other Aux services are running on its own container.

Migrating Federation Services from VM to Containers: SIFULAN Experience (cont.)

- After:
  - Each organization who would like to subscribe IdP as a service, can have their own Namespace, and each service (e.g. Shibboleth IdP, OpenLDAP) run individually in a container.
  - Each container image was created as generic as possible so that it can be reusable by other federation/interested party
  - All specific configurations are stored as ConfigMAP object.

Migrating Federation Services from VM to Containers: SIFULAN Experience (cont.)

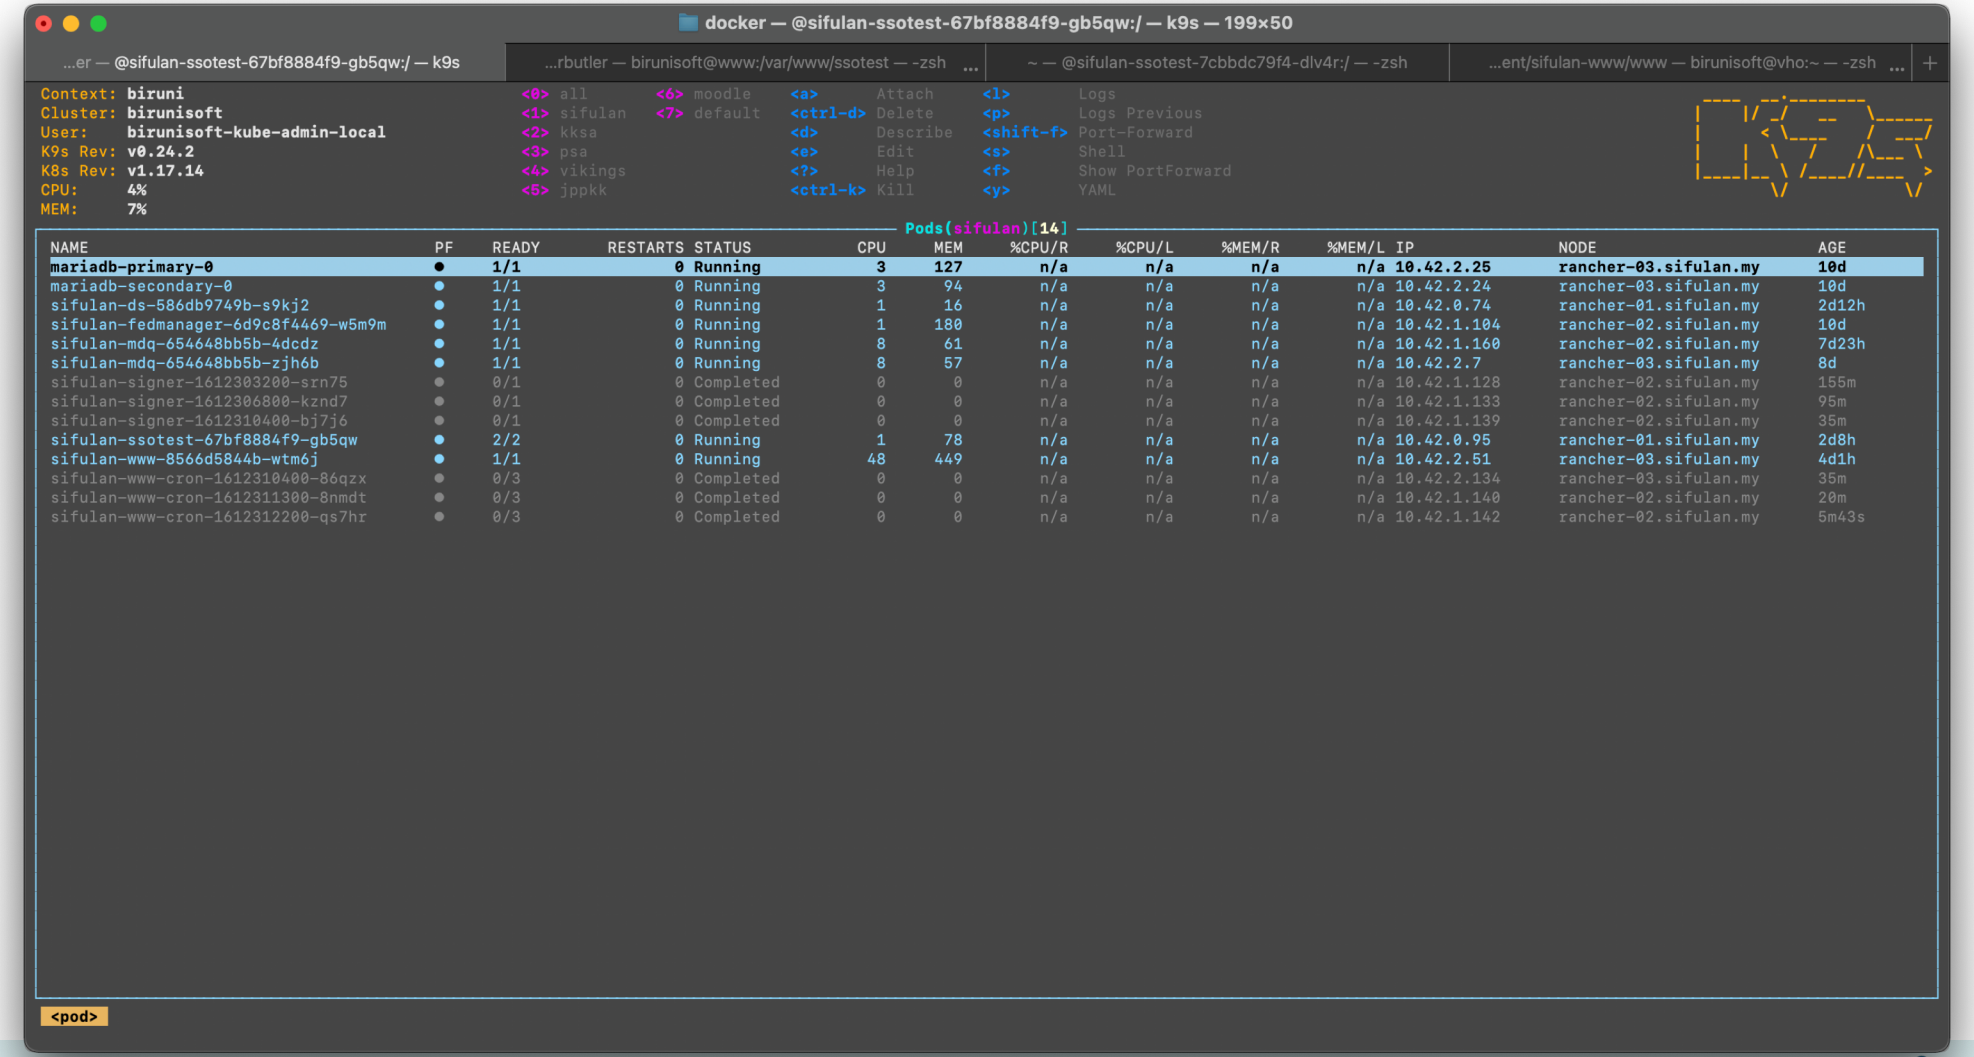

SIFULAN
MALAYSIAN ACCESS FEDERATION

Possible IDMS-as-a-Service Implementation with K8s

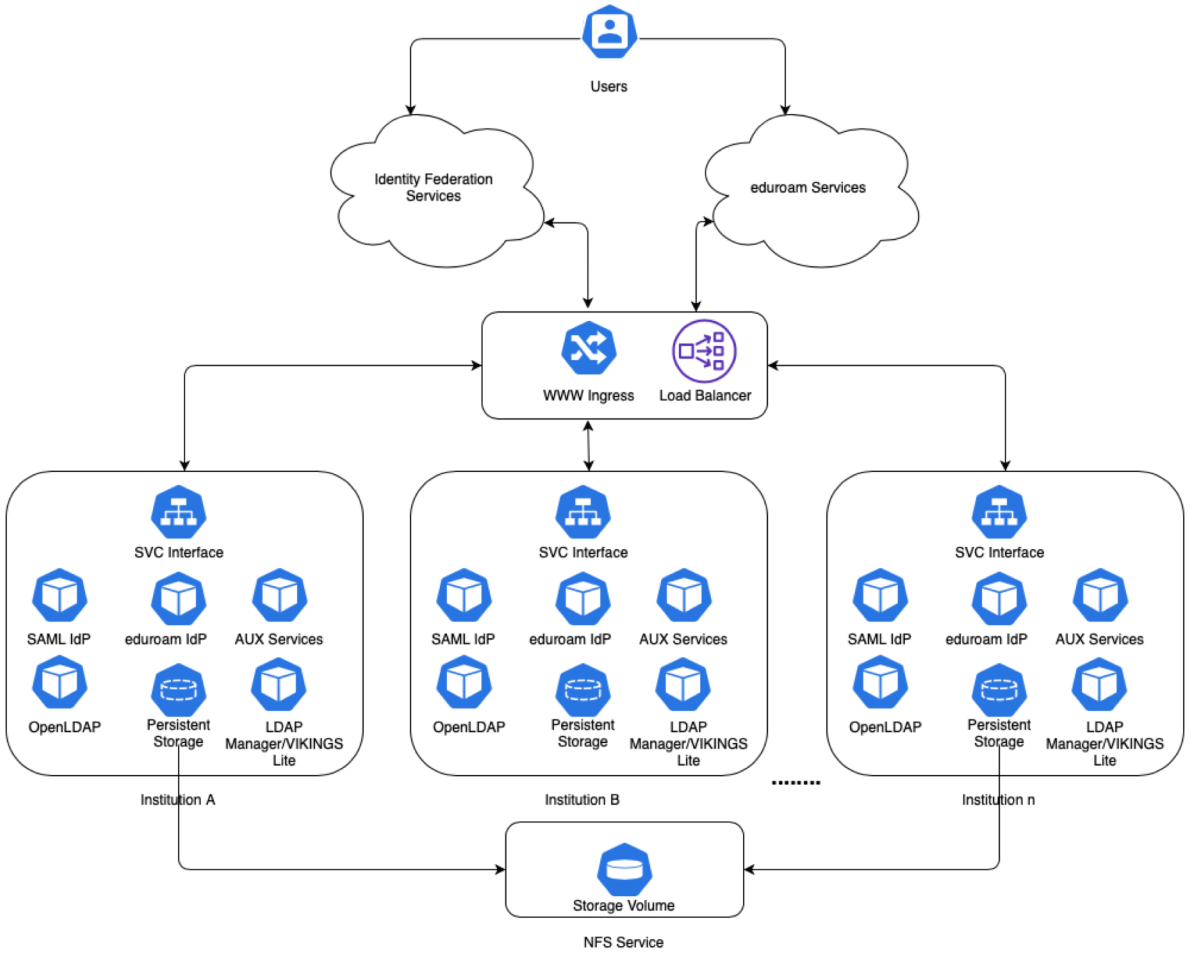


thank you
danke
謝謝
ngiyabonga
شكراً جزيلاً
рахмат
teşekkür ederim
спасибо
gracias
dank je
tapadh leat
Баярлалаа
mersi
barka
welalin
tack
spas
faafetai lava
vinaka
blagodaram
misaotra
matondo
paldies
grazzi
mahalo
хвала
asante
manana
obrigada
tenki
kiitos
dankie
спасибі
dhanyavad
nanni
nandri
mauruuru
hvala
gràcie
bayarlalaa
kösz̈önöm
akun
dankon
aciù
enkosi
bedankt
dziękuję
djere dieuf
tau
mochchakkeram
дякую
mamnun
murakoze
chokrane
chnorakaloutioun
gratias ago
gràcies
sulpáy
go raibh maith agat
sobodi
dèkuji
sagolun
sukriya
kop khun krap
taiku
obrigado
mesi
didi madloba
kam sah hamnida
najis tuke
ありがとう
grazie
arigatô
takk
dakujem
trugarez
terima kasih
rahmat
rahmet
tanemirt
diolch
dhanyavadagalu
shukriya
merce
мерси
তোমাকে ধন্যবাদ
감사합니다
xiexie
ευχαριστώ
merci