

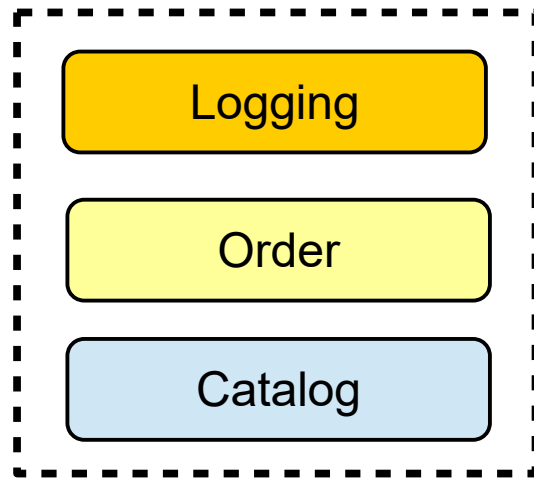
Building a Cloud Infrastructure to Deploy Microservices as MicroVM Toro Guests

www.torokernel.io

Matias Vara Larsen
matiasvara@gmail.com

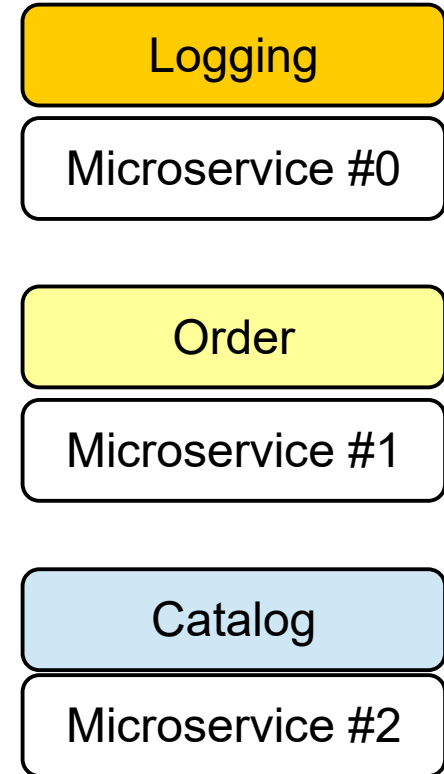
What are microservices?

e.g., Amazon website

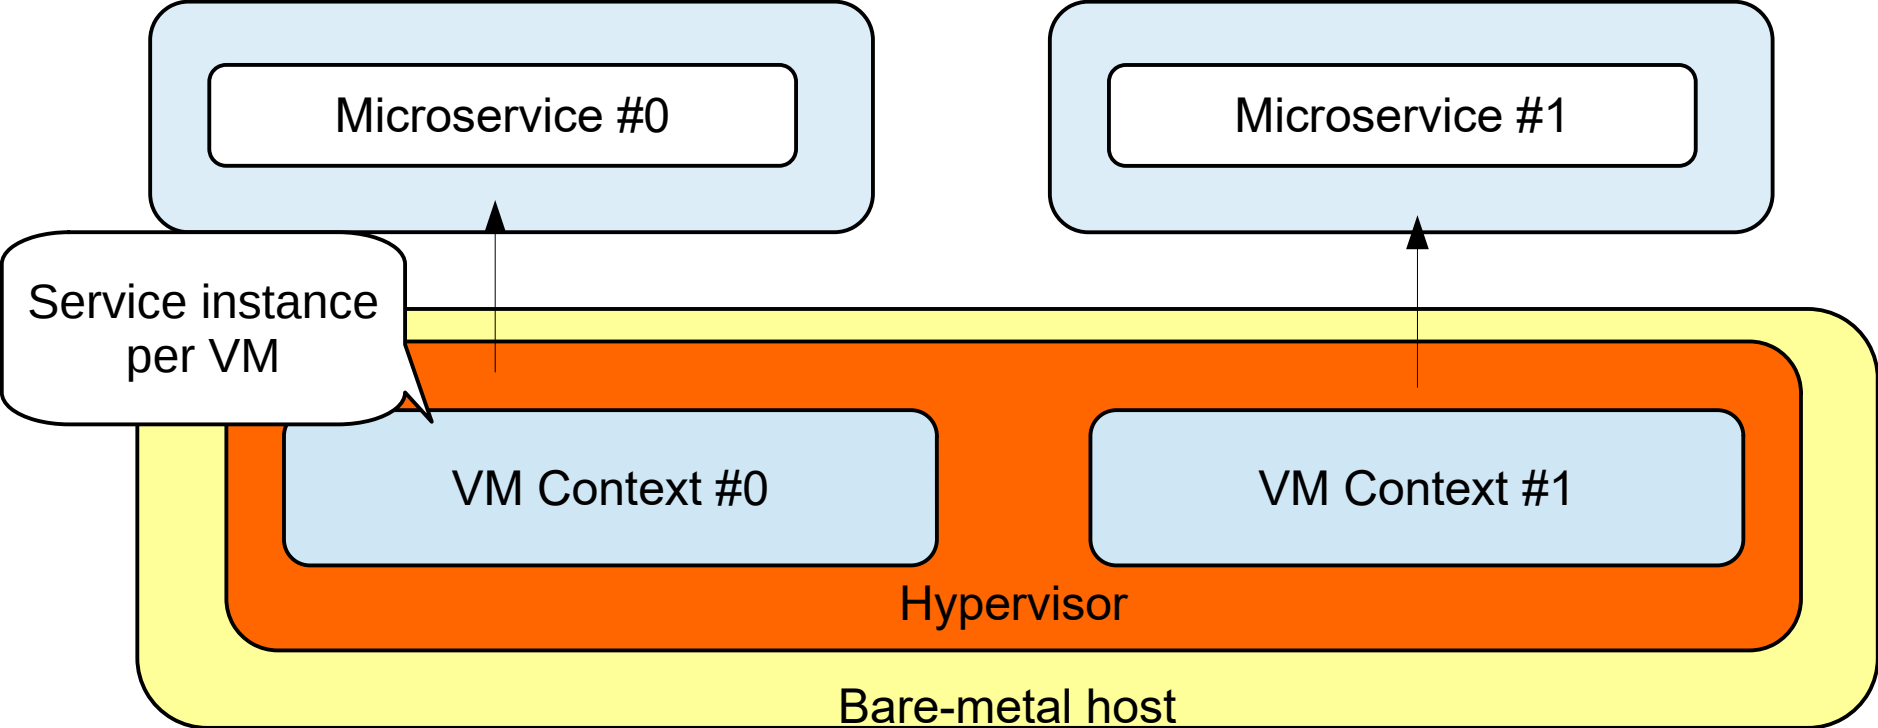


Monolithic Application

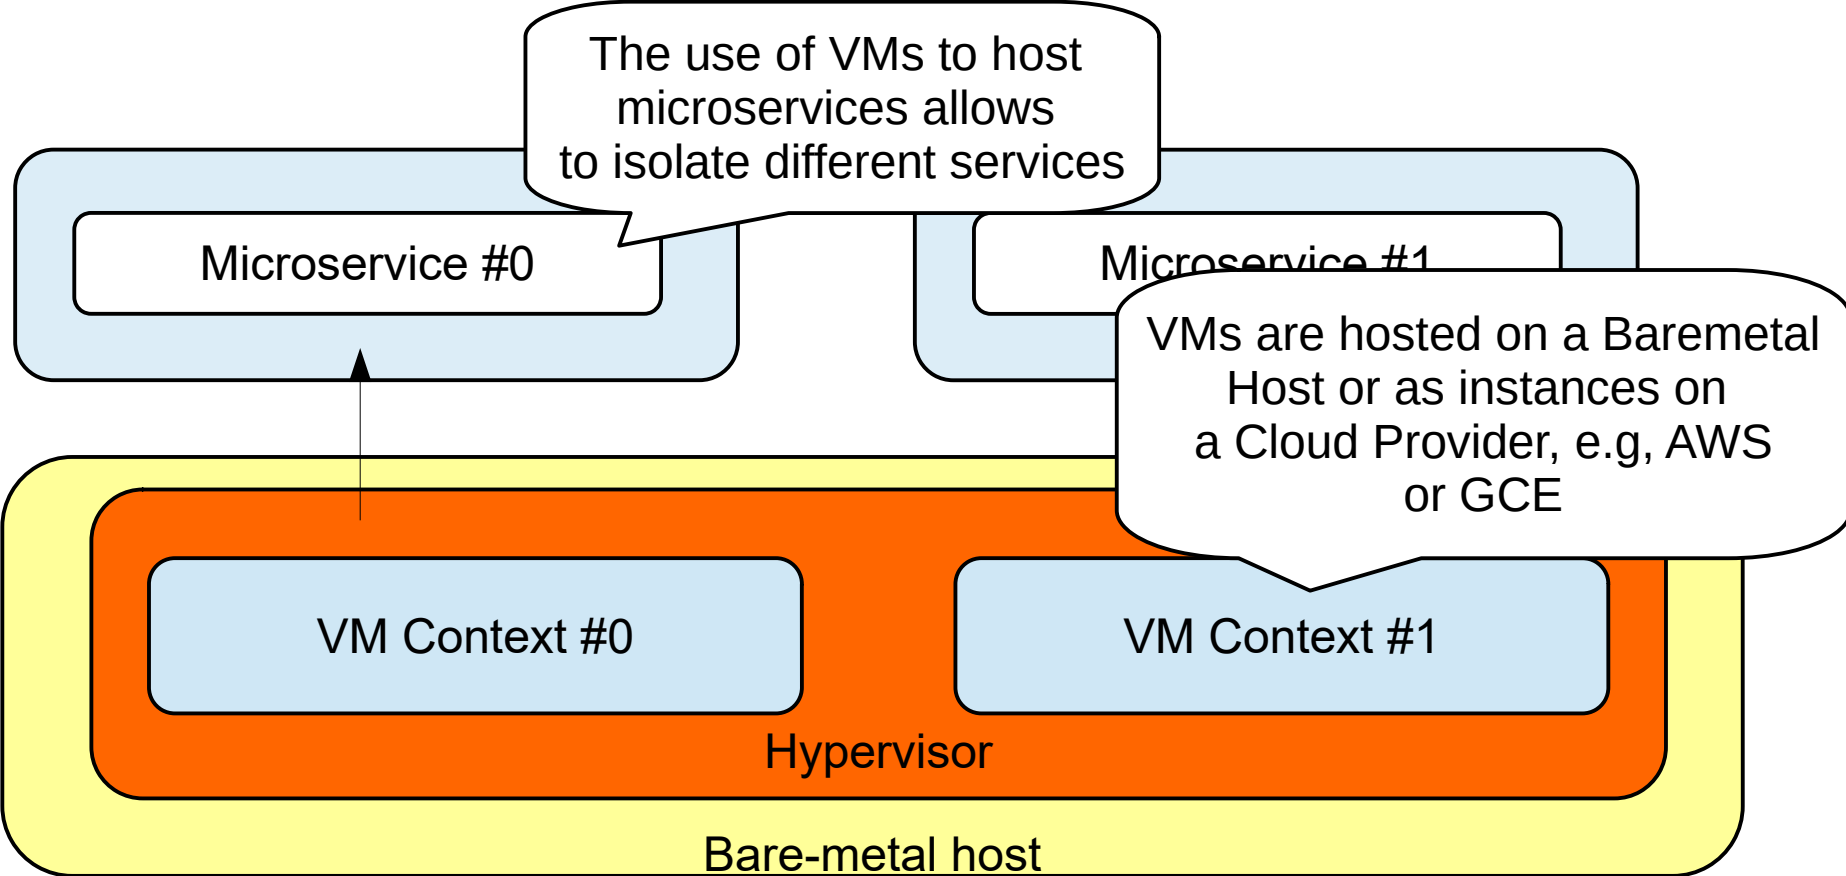
Decomposed Application
into Services



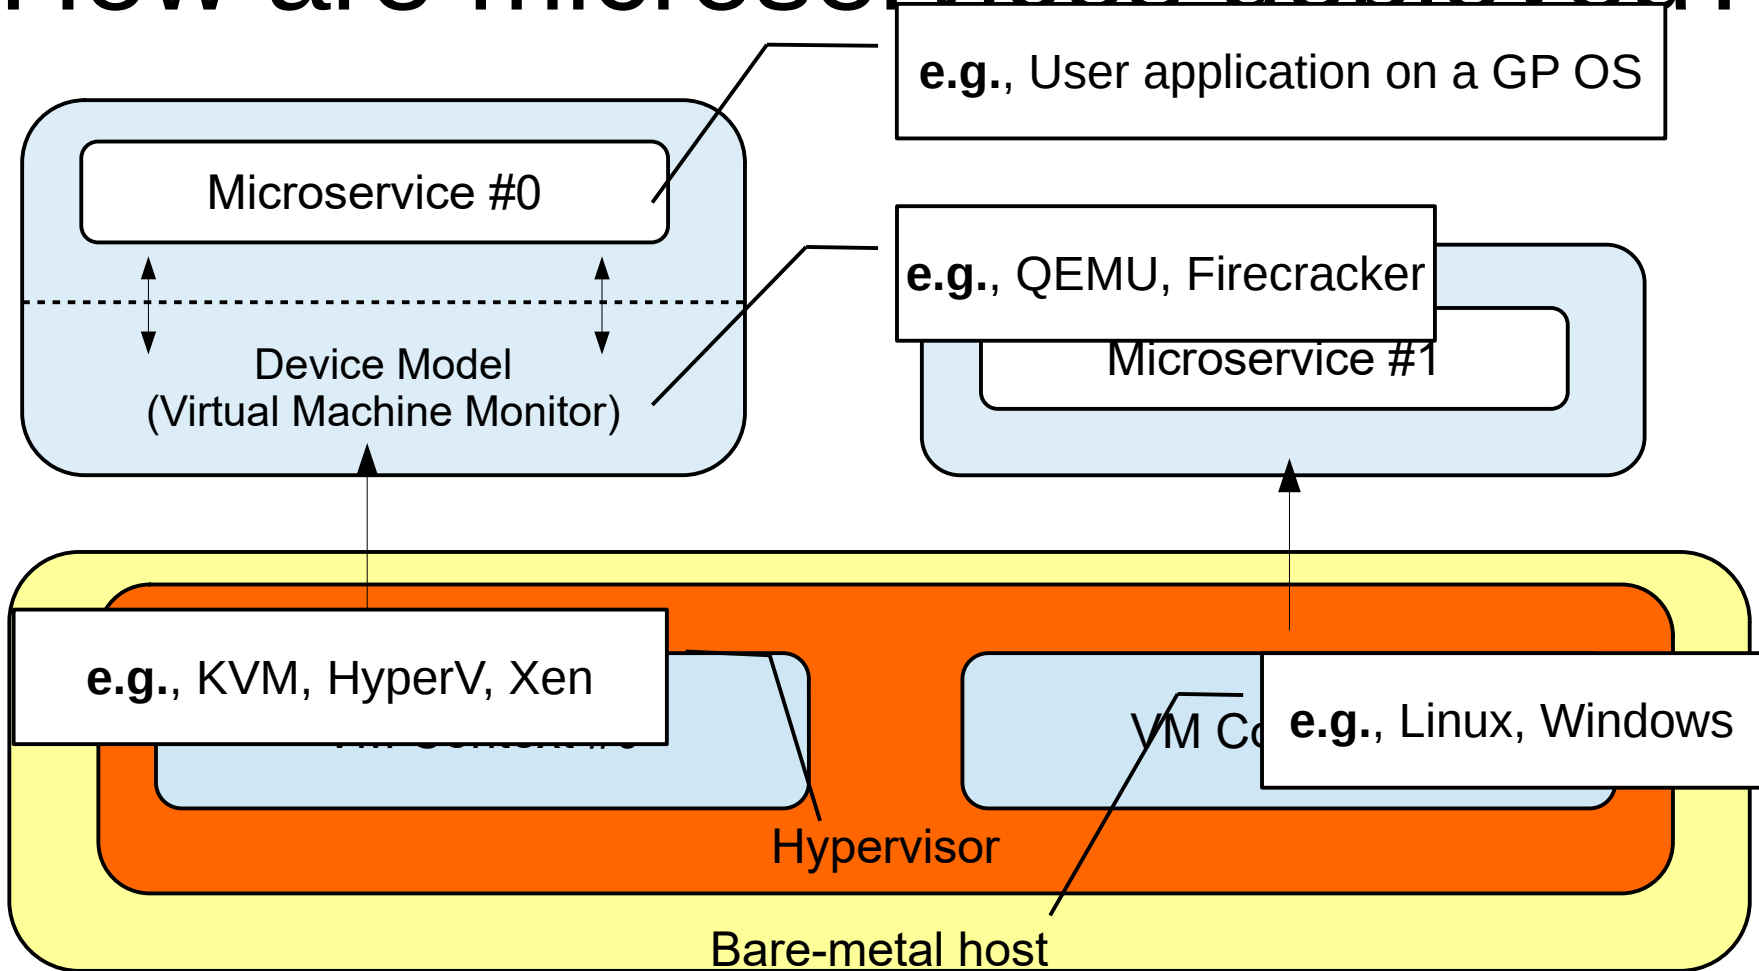
How are microservices deployed?

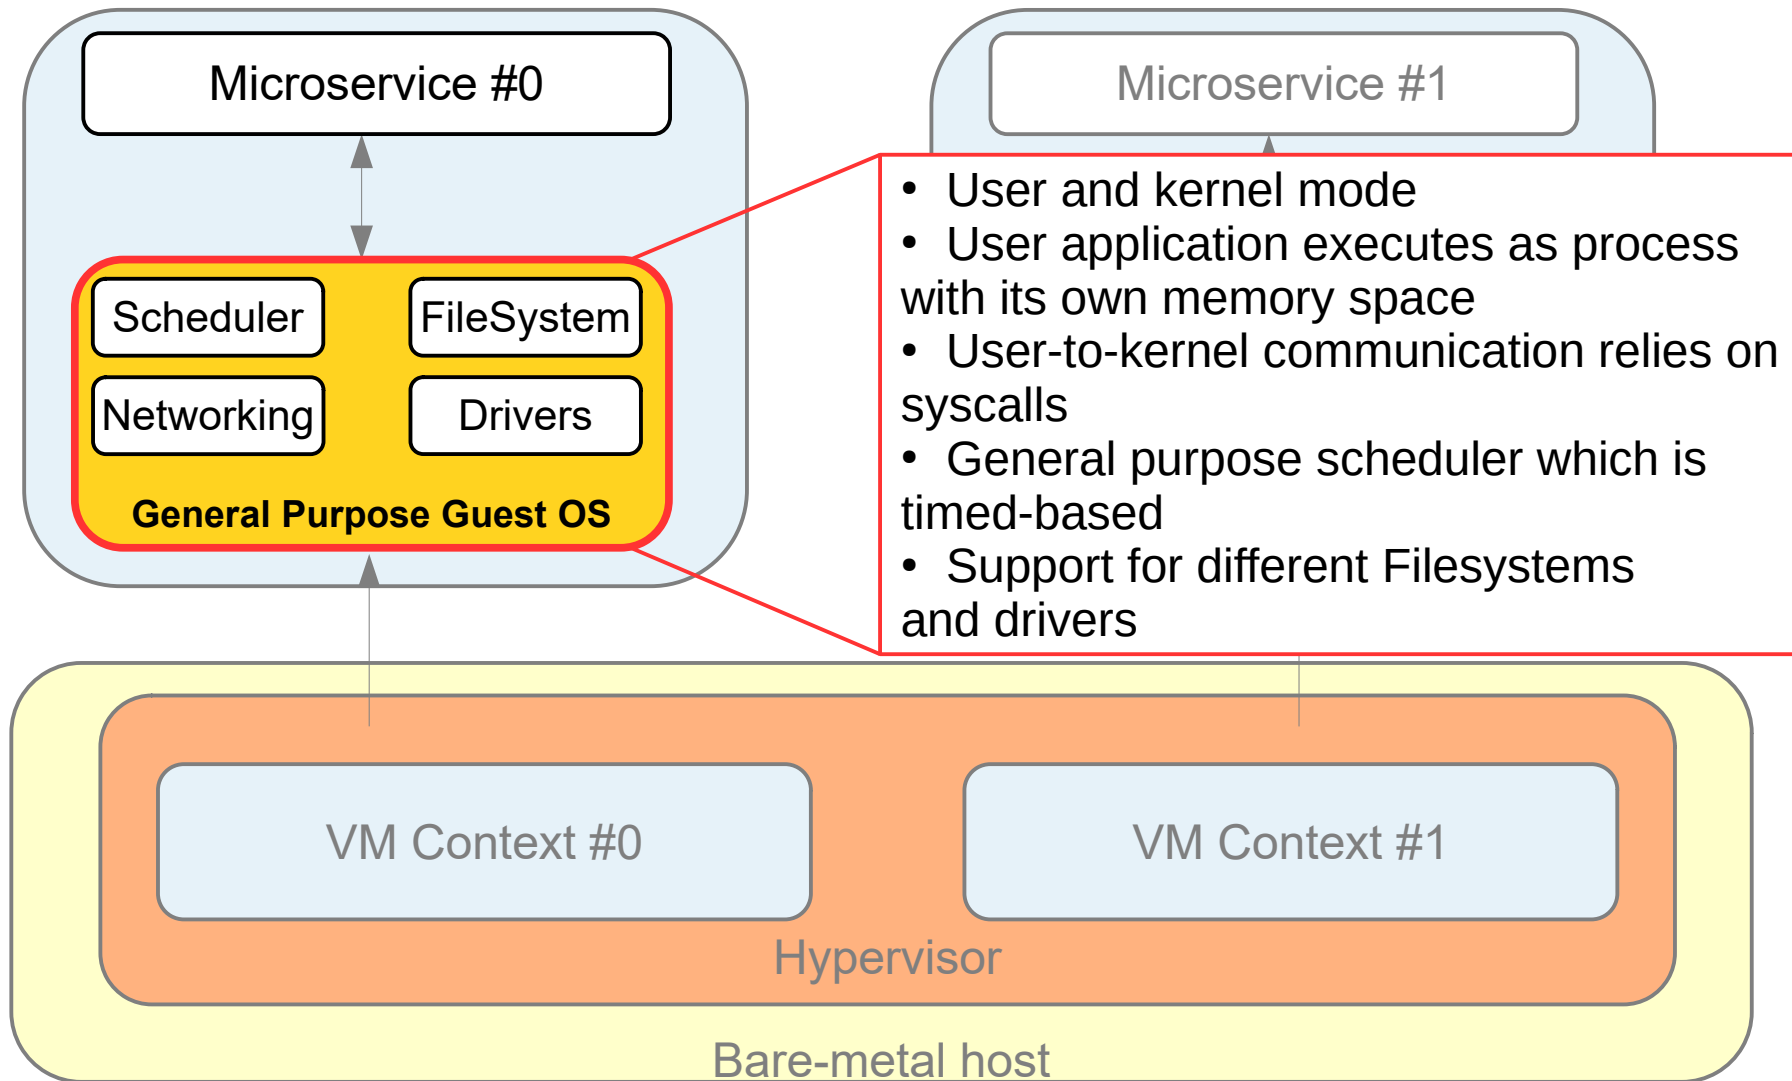


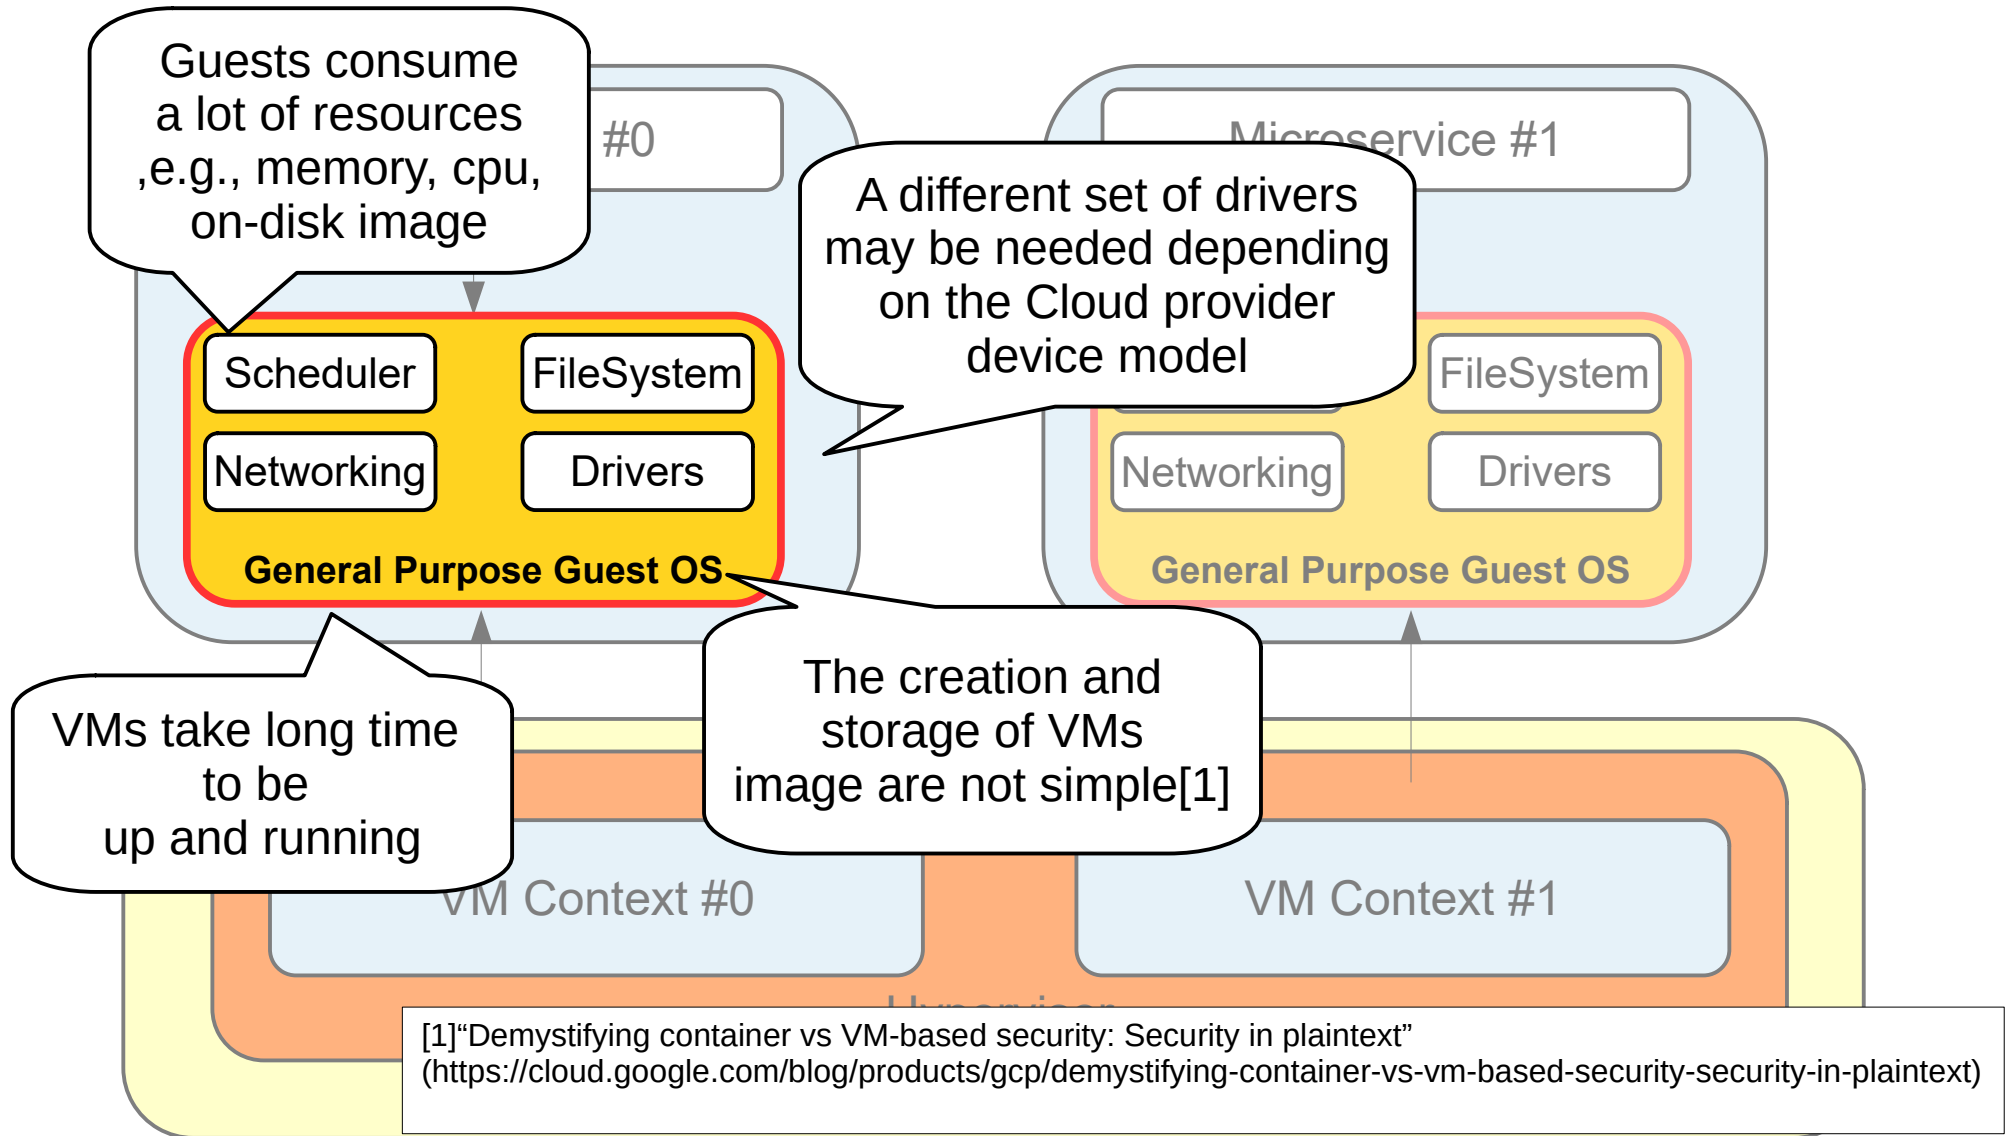
How are microservices deployed?

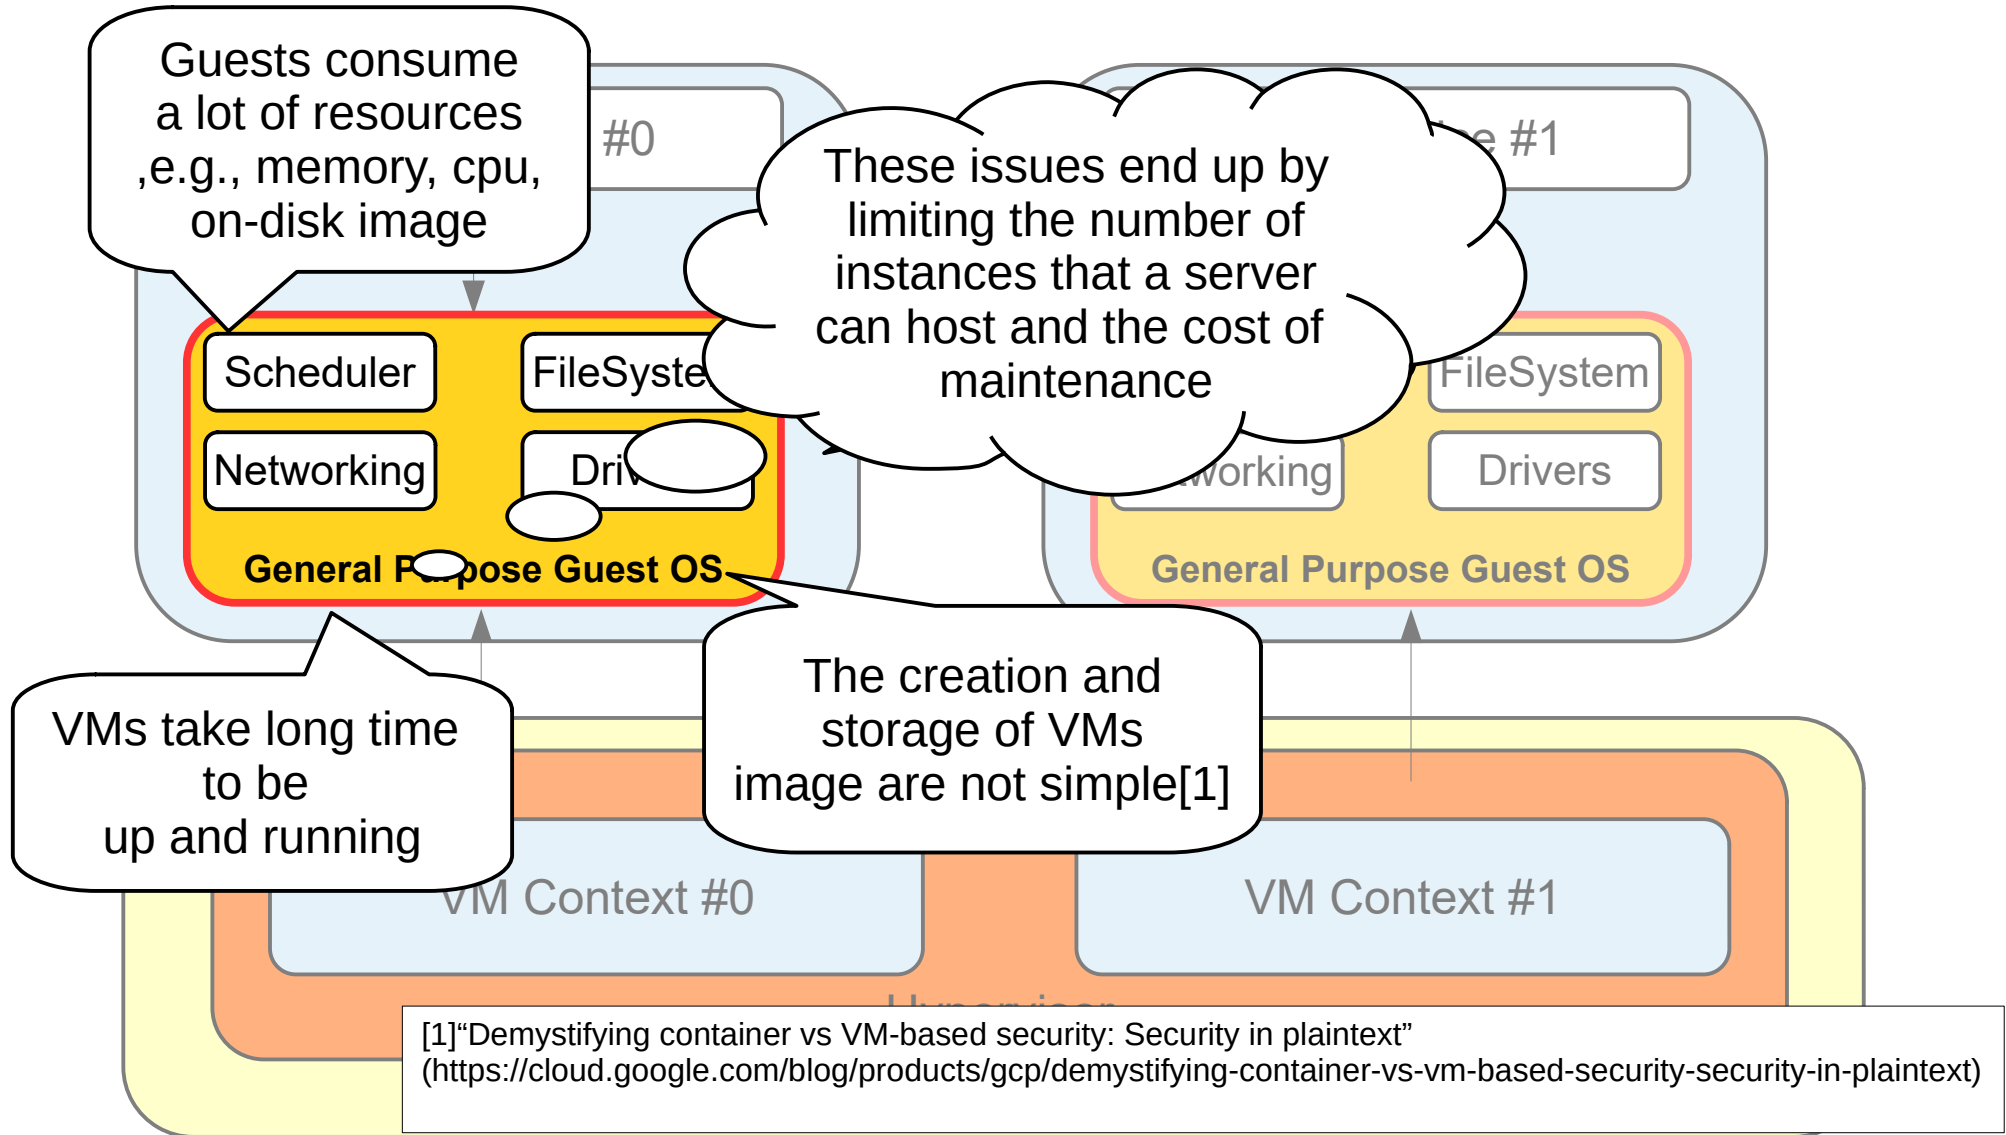


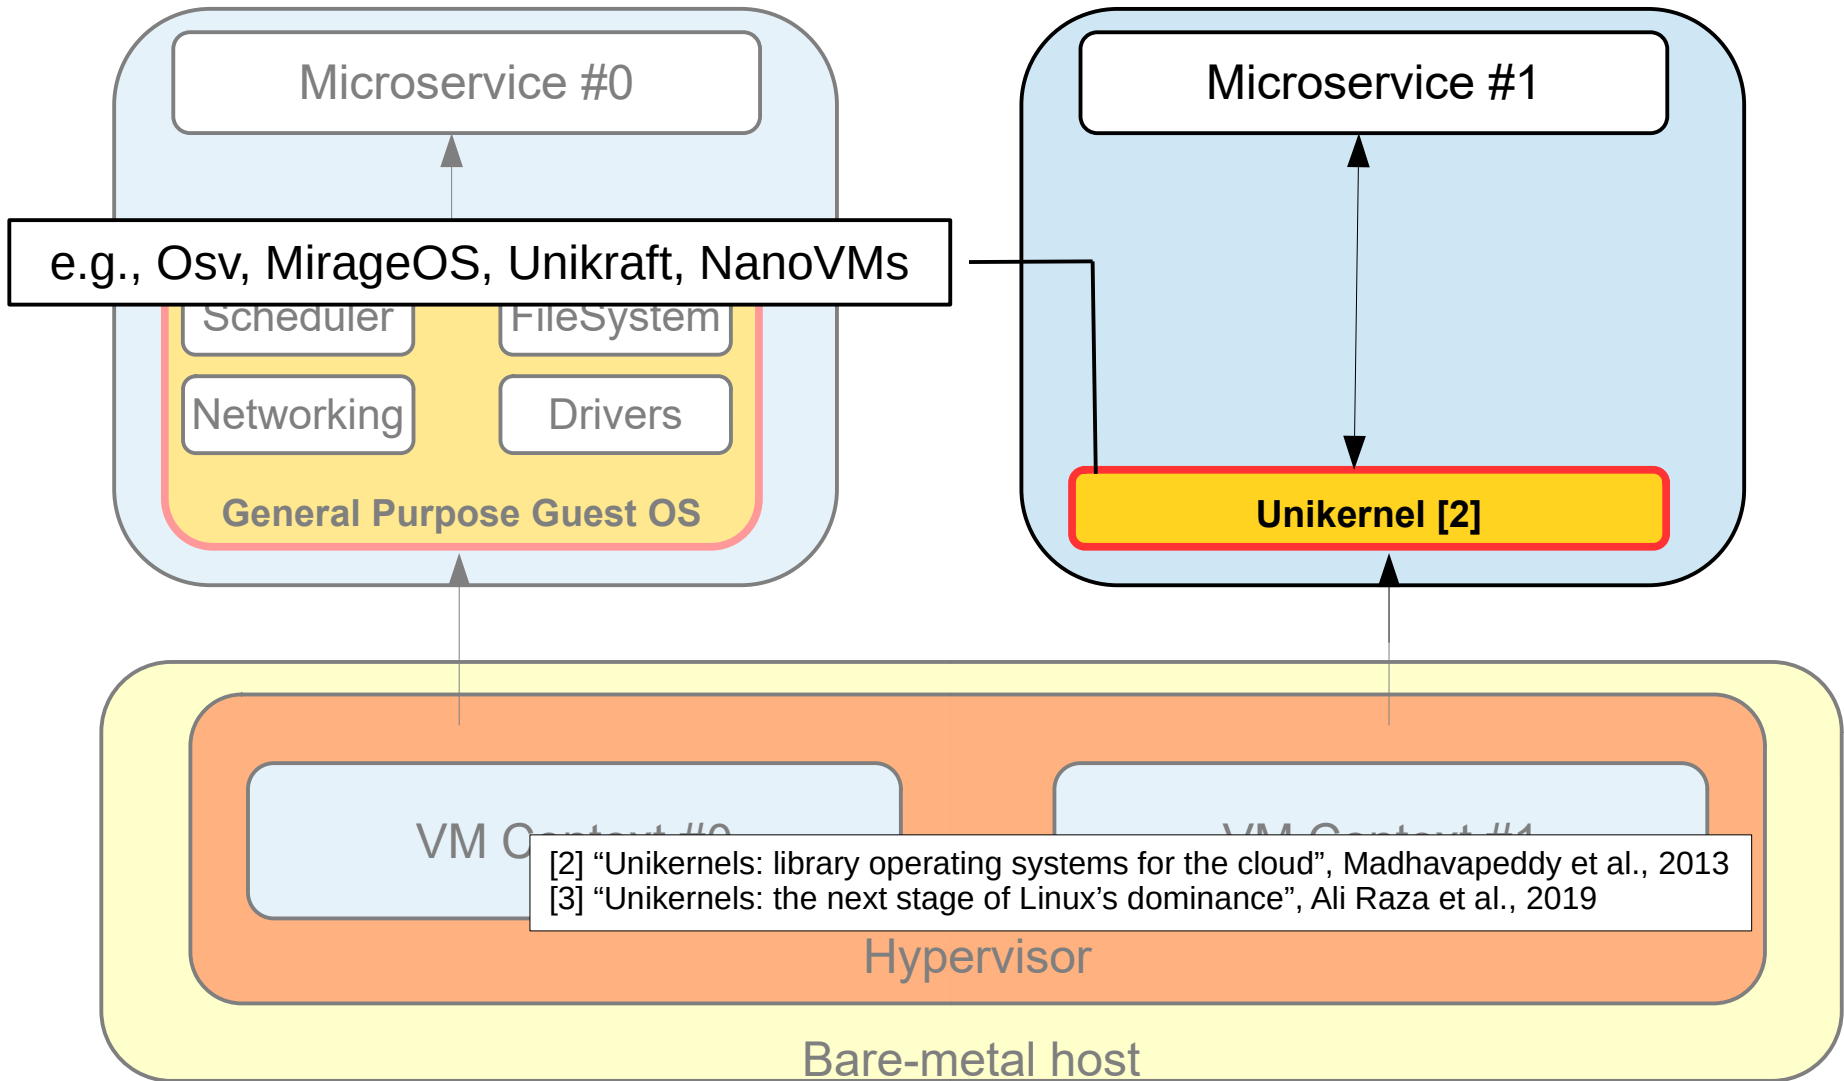
How are microservices deployed?



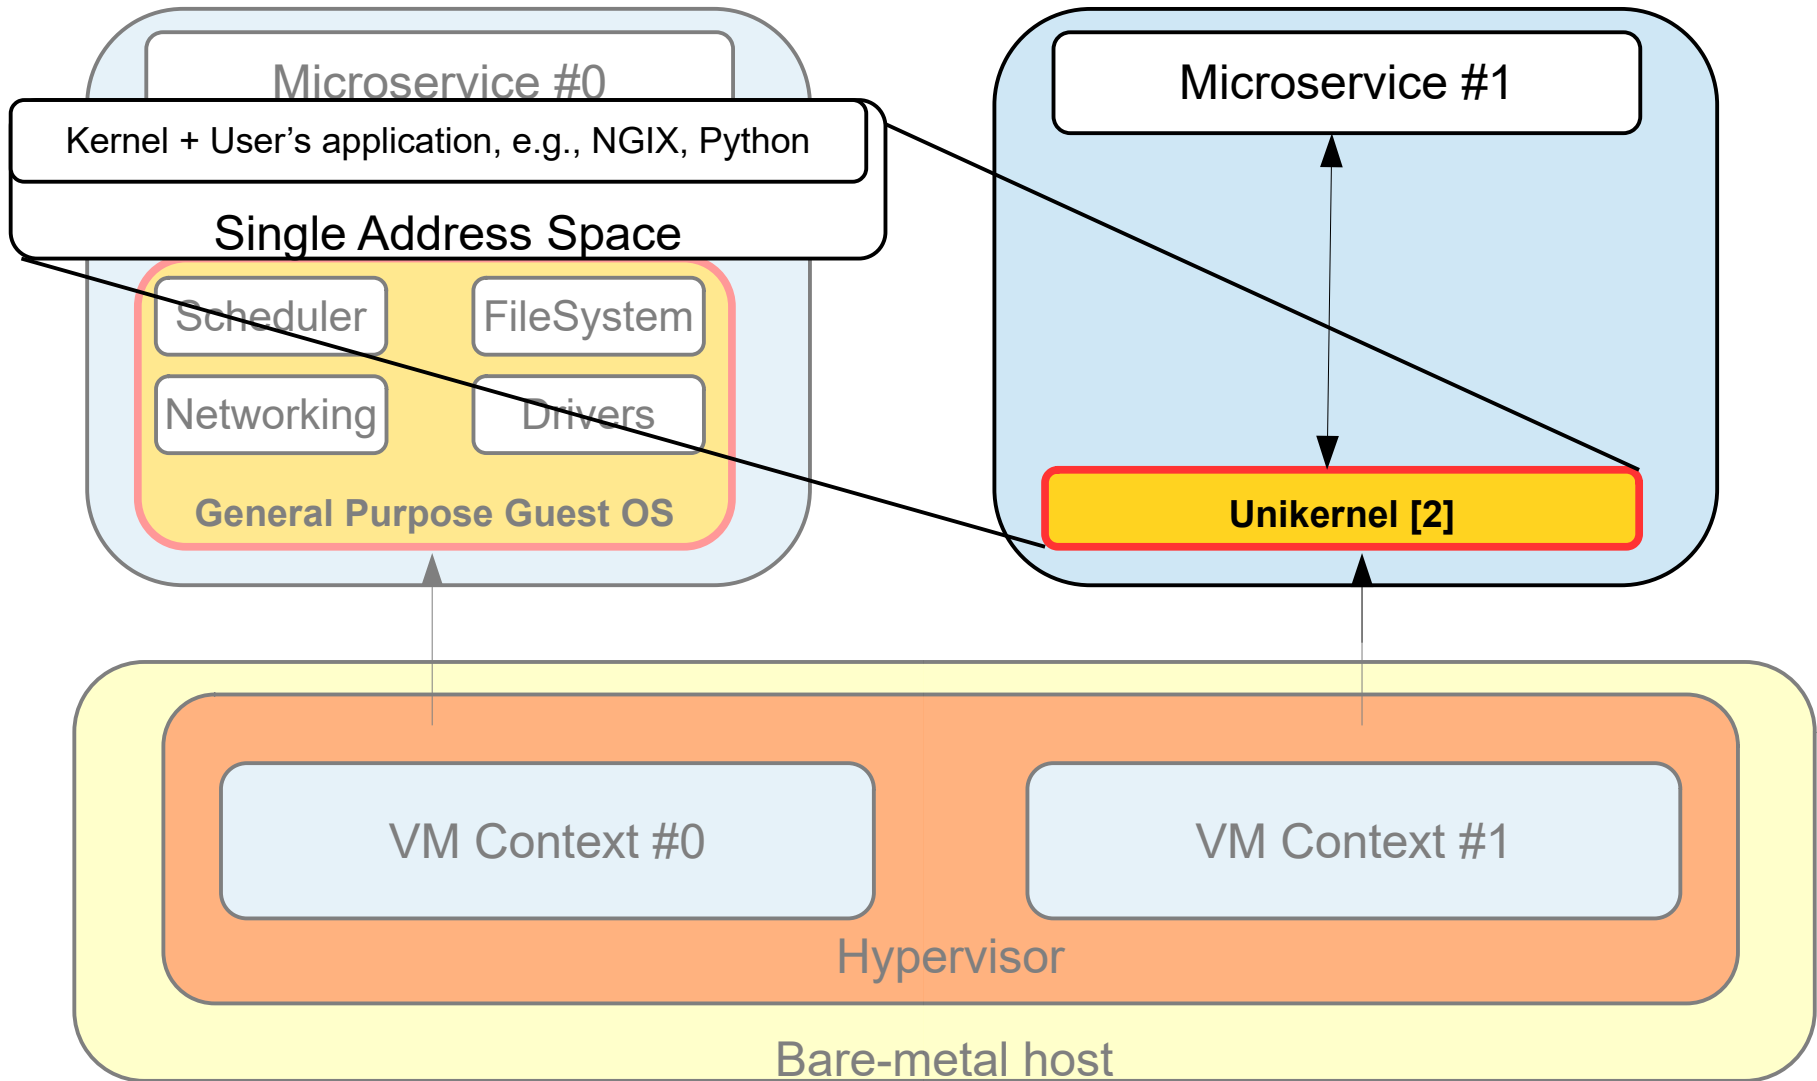


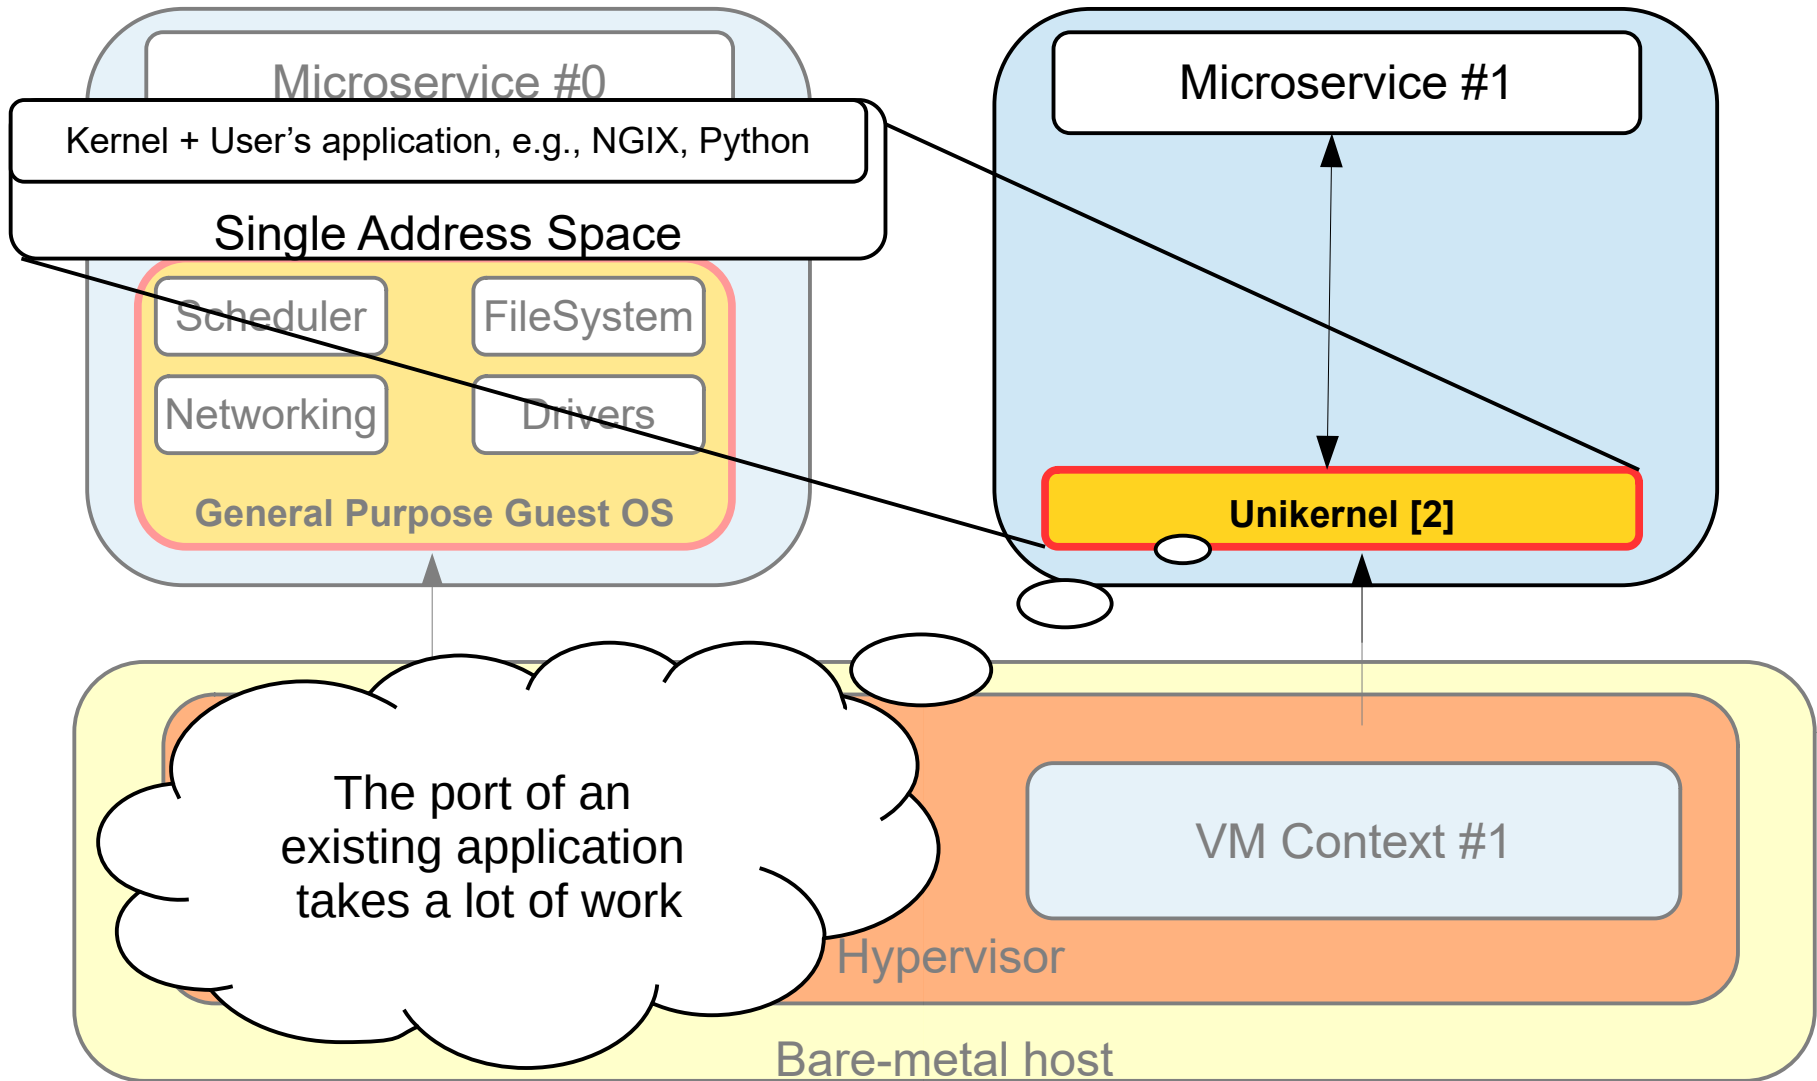


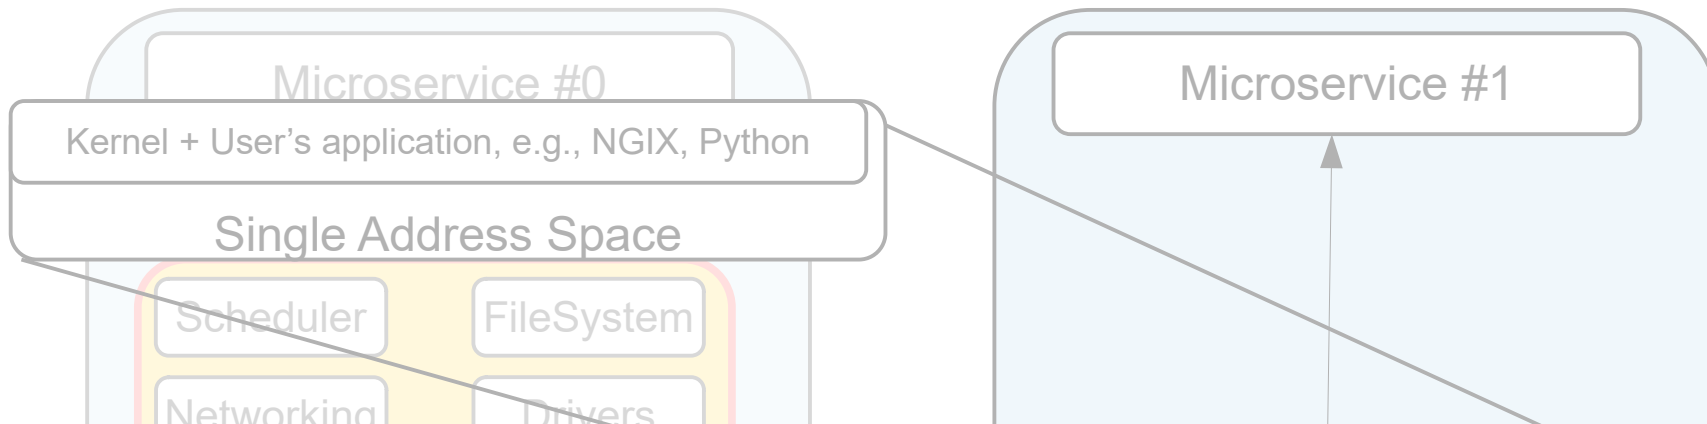





[2] "Unikernels: library operating systems for the cloud", Madhavapeddy et al., 2013
[3] "Unikernels: the next stage of Linux's dominance", Ali Raza et al., 2019



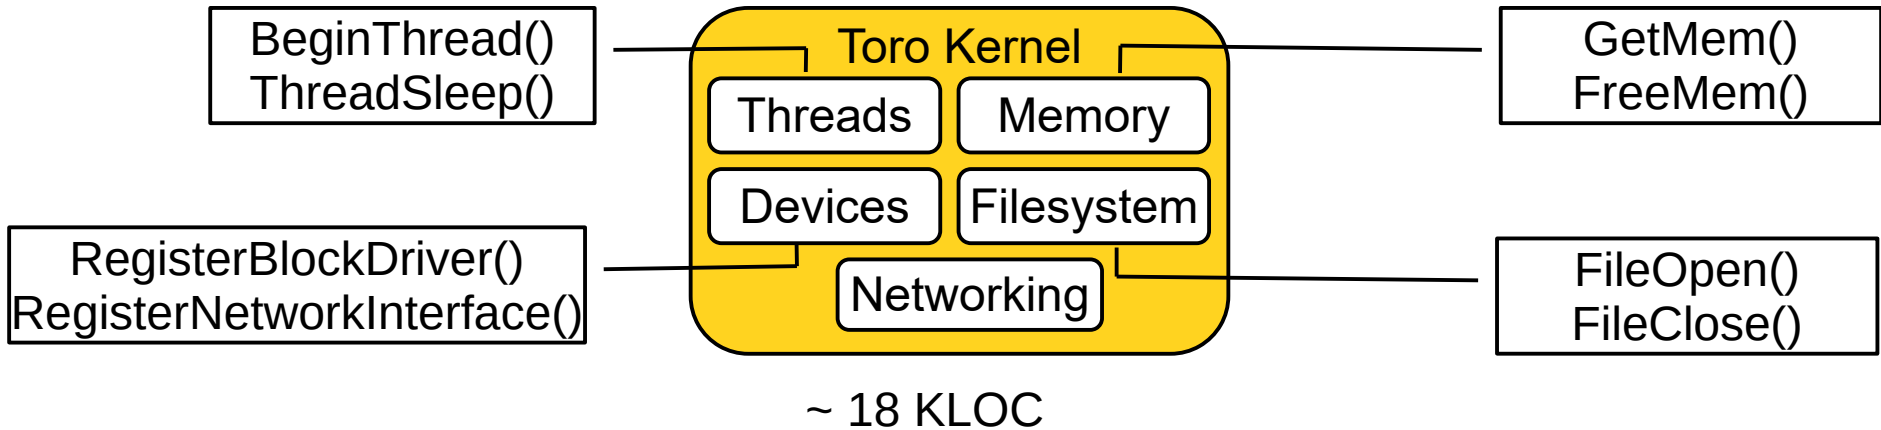




Toro is an application-oriented unikernel that allows **microservices** to run efficiently in **VMs** thus leveraging the strong isolation VMs provide.



Application-oriented Kernel

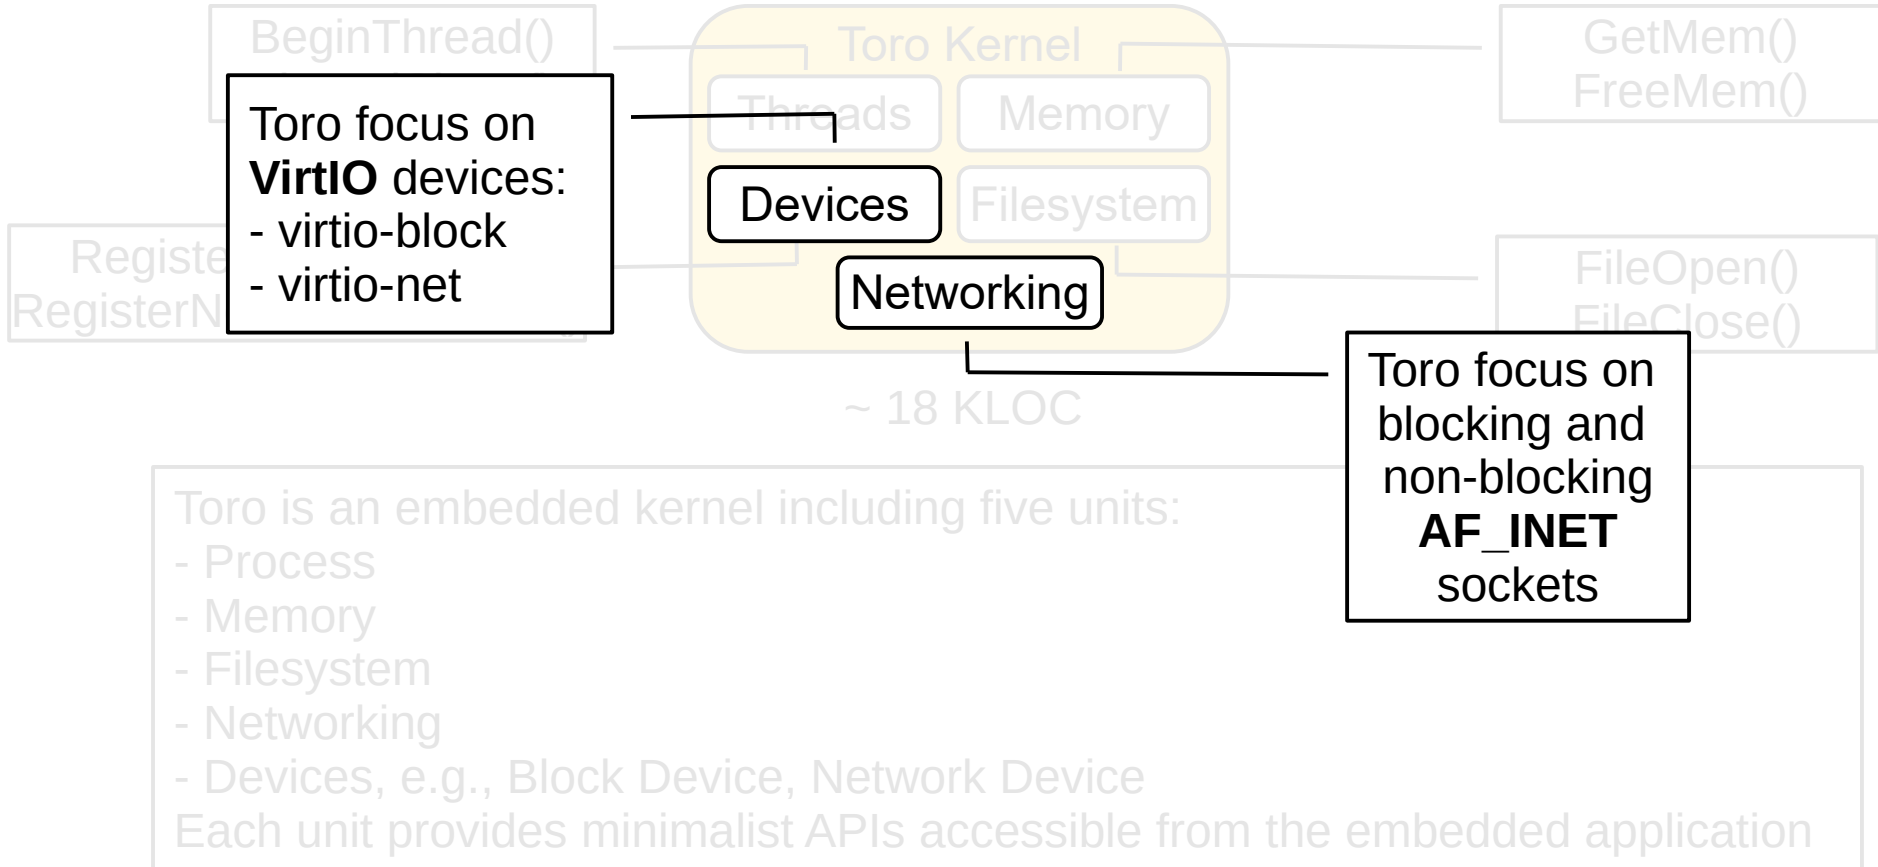


Toro is an embedded kernel including five units:

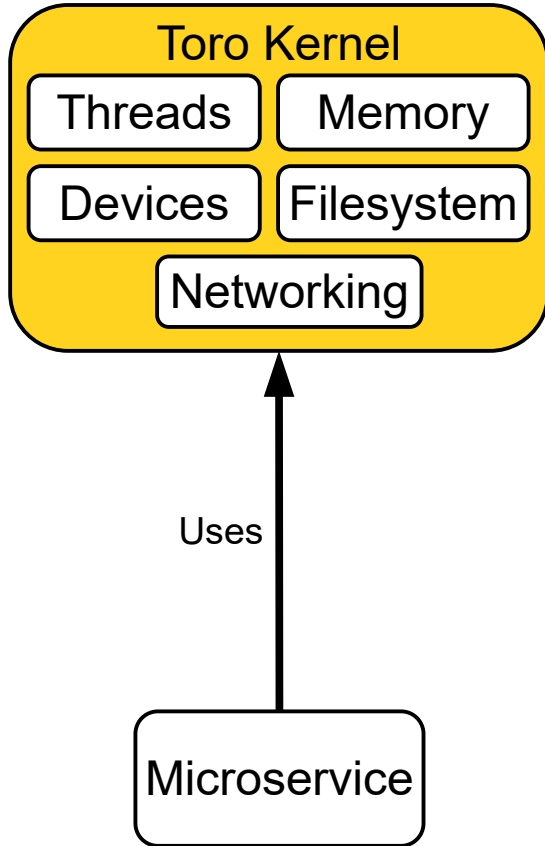
- Process
- Memory
- Filesystem
- Networking
- Devices, e.g., Block Device, Network Device

Each unit provides minimalist APIs accessible from the embedded application

Application-oriented Kernel

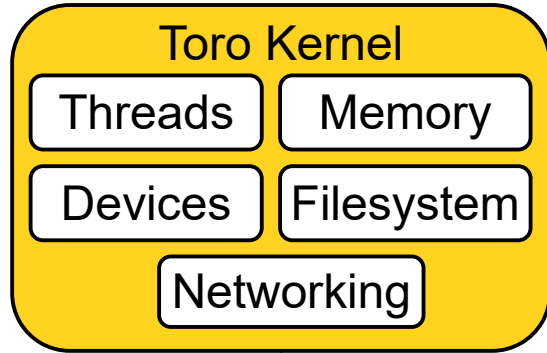


Application-oriented Kernel



- User application and kernel units are compiled in a single binary
- The application includes only the component required

Application-oriented Kernel




Uses

Microservice

- User application and kernel units are compiled in a single binary
- The com

```
program WebServerAppliance;  
uses  
    Memory,  
    Filesystem,  
    Threads,  
    Networking,  
    Fat,  
    Virtio-blk,  
    Virtio-net;  
Begin  
//  
// Your Code Goes Here  
//  
End.
```


Application-oriented Kernel



Uses

Microservice

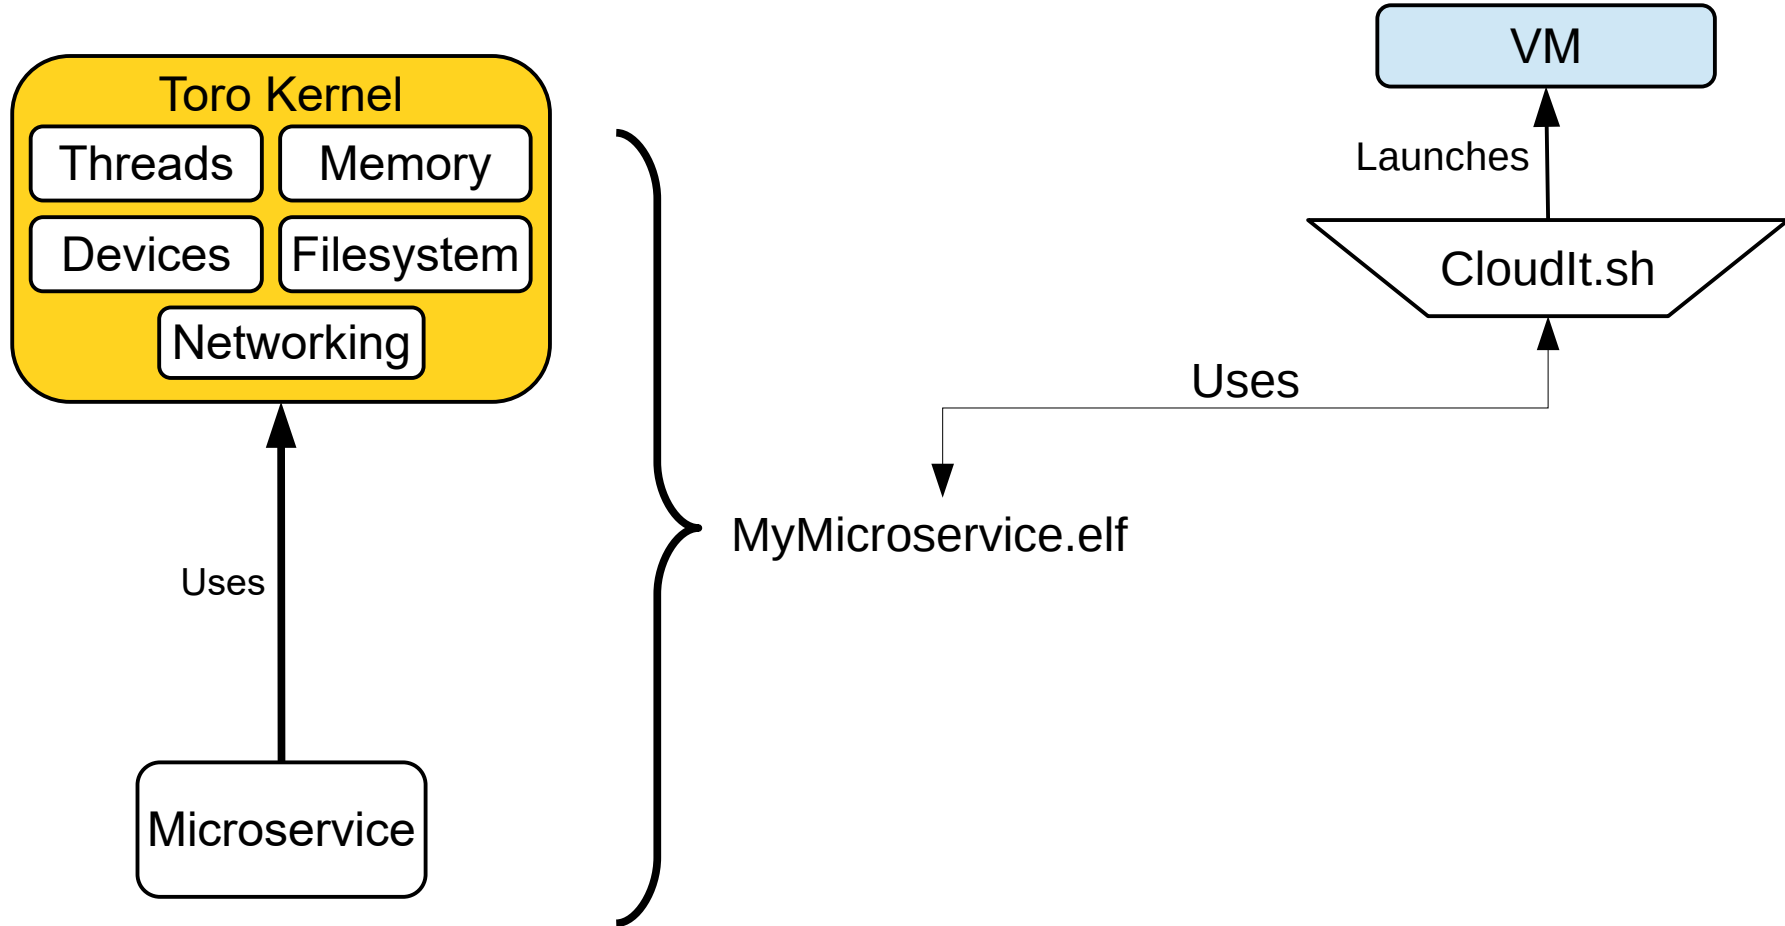
Kernel cannot be modified during its life

MyMicroservice.elf

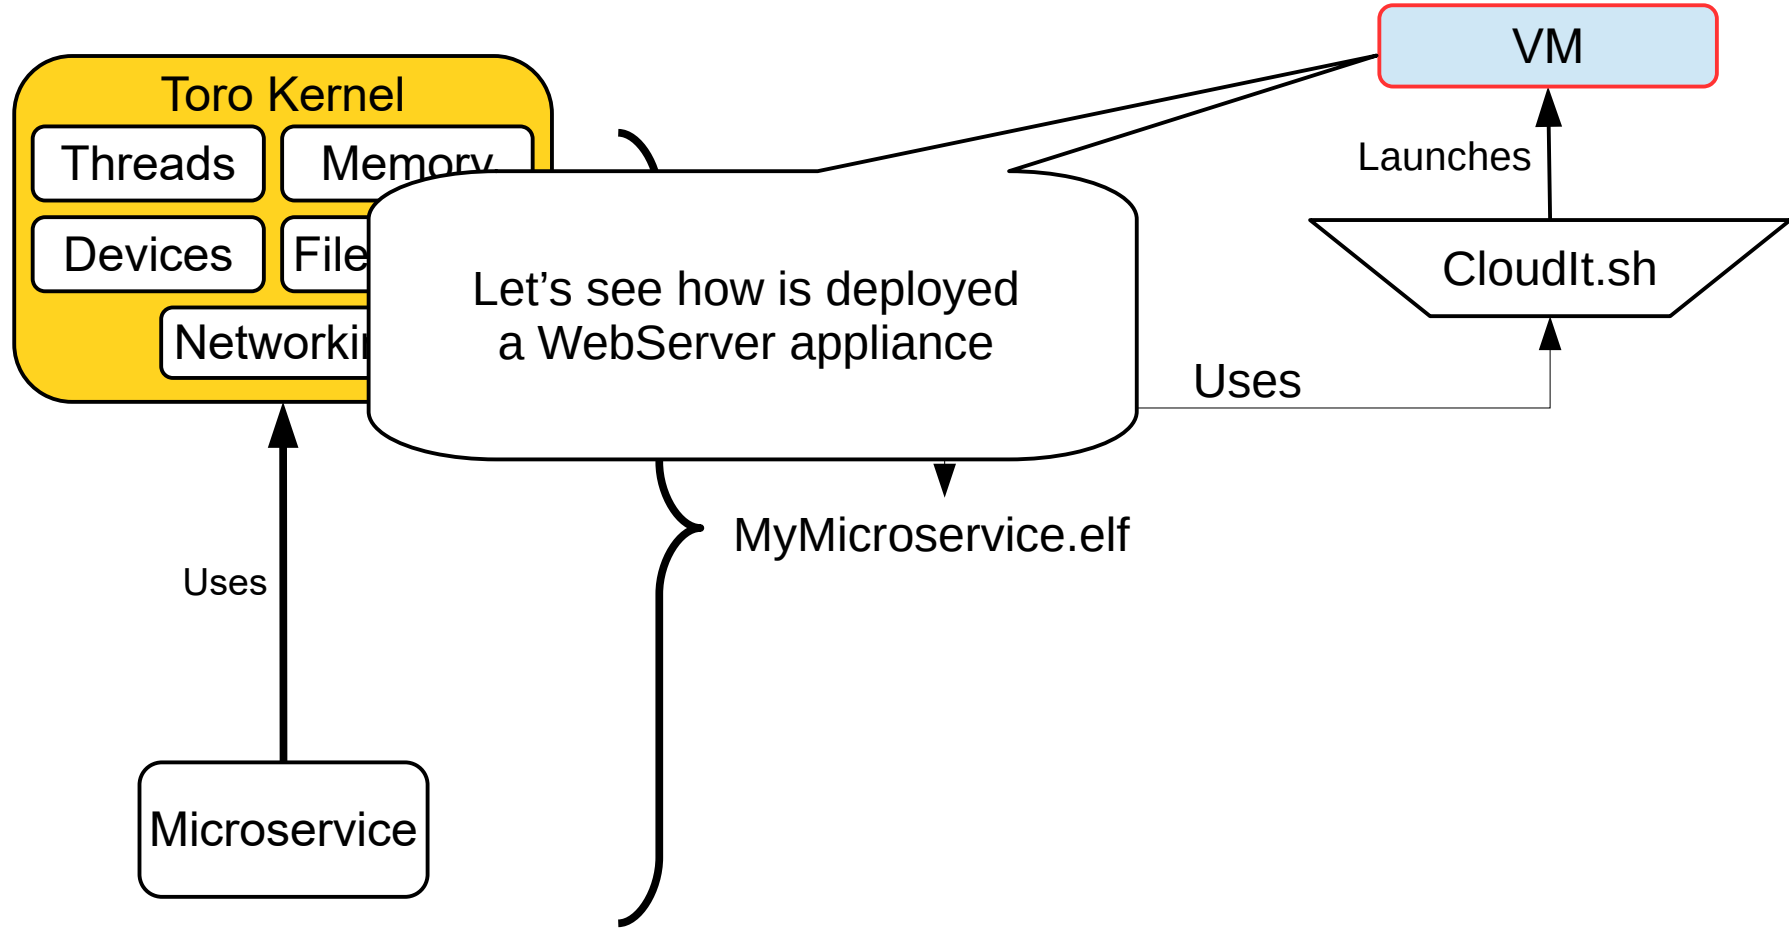
[1]"Immutable Infrastructure", Stella

The generated binary is **Immutable**[1], i.e., the generated image can be used across different hypervisors without the need to recompile it.

Application-oriented Kernel



Application-oriented Kernel

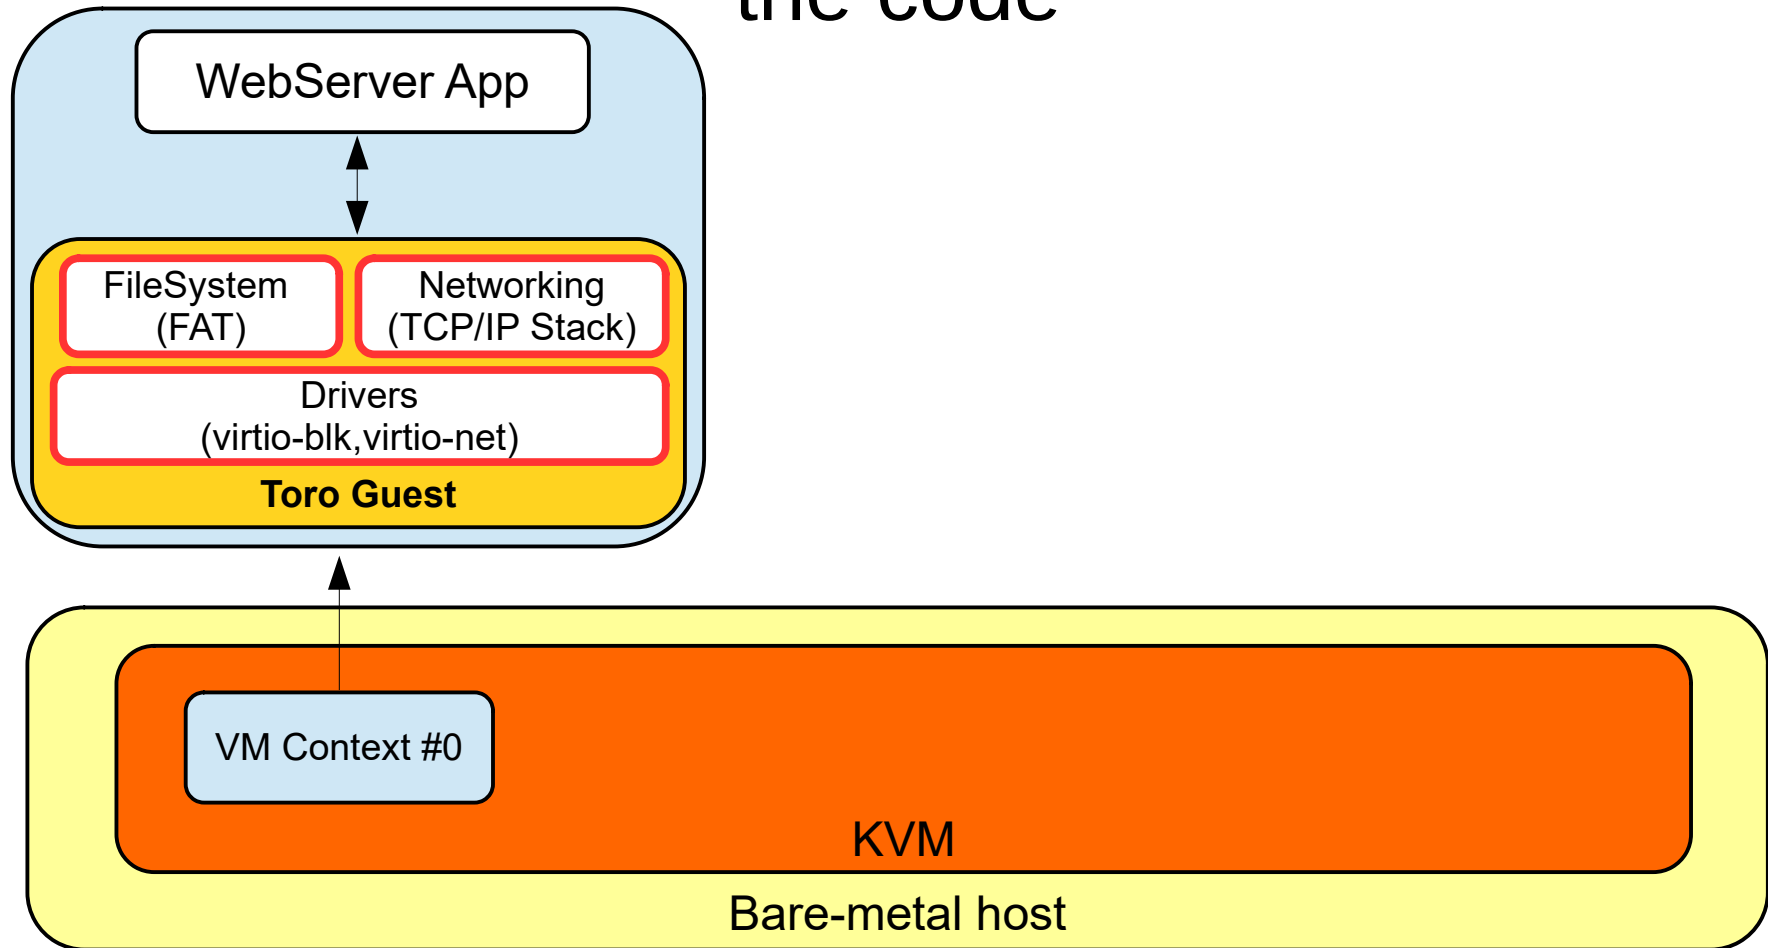


The WebServer Appliance

- Simple microservice that serves files by using the HTTP protocol
 - Find it at <https://github.com/torokernel/torokernel> among other examples
 - This appliance is used to host Toro's website (<http://www.torokernel.io> and click on "View on Toro")

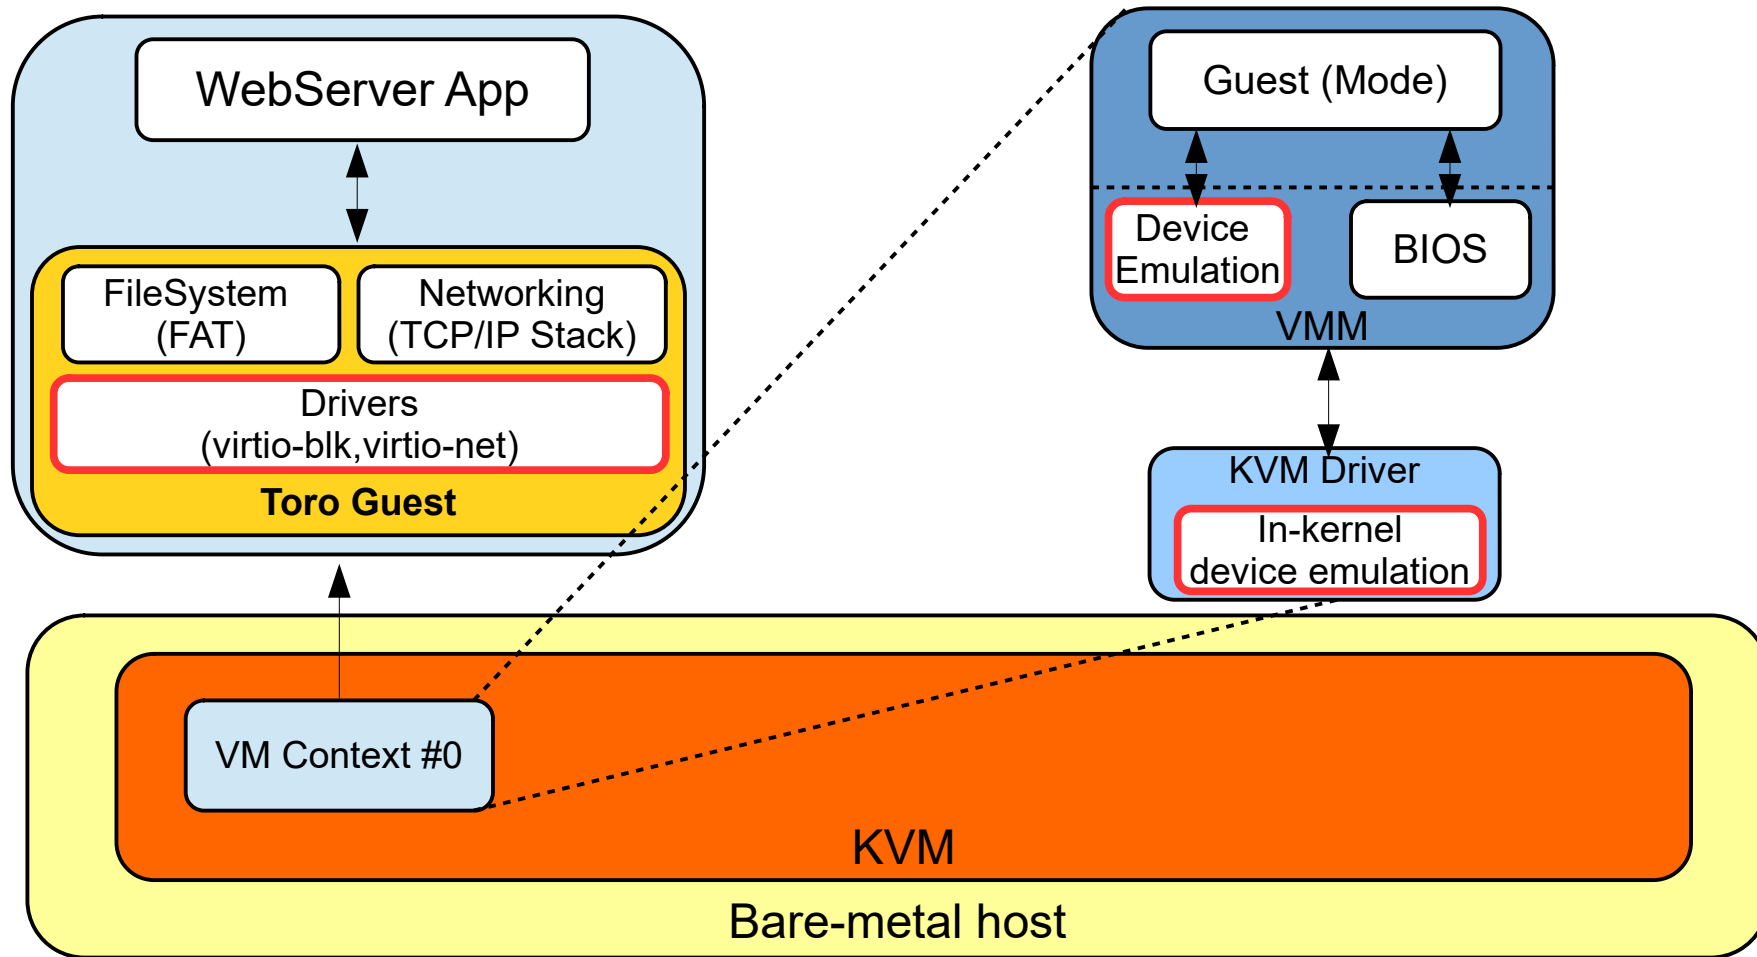
How the appliance is setup?

the code



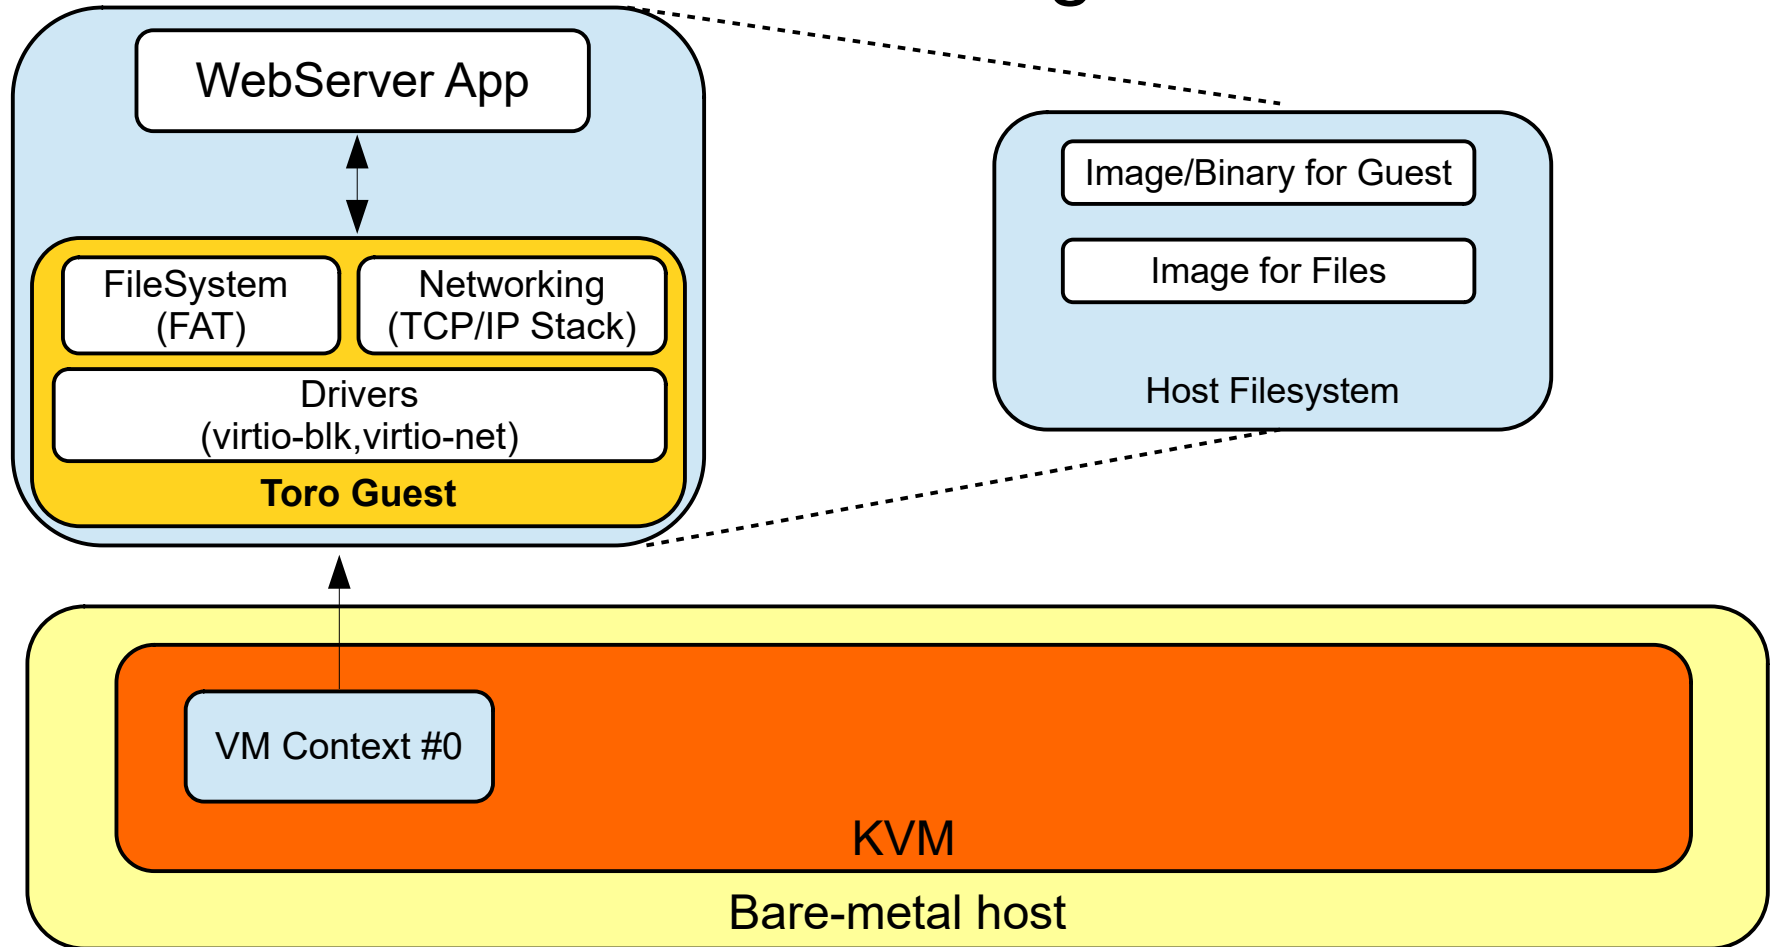
How the appliance is setup?

device model



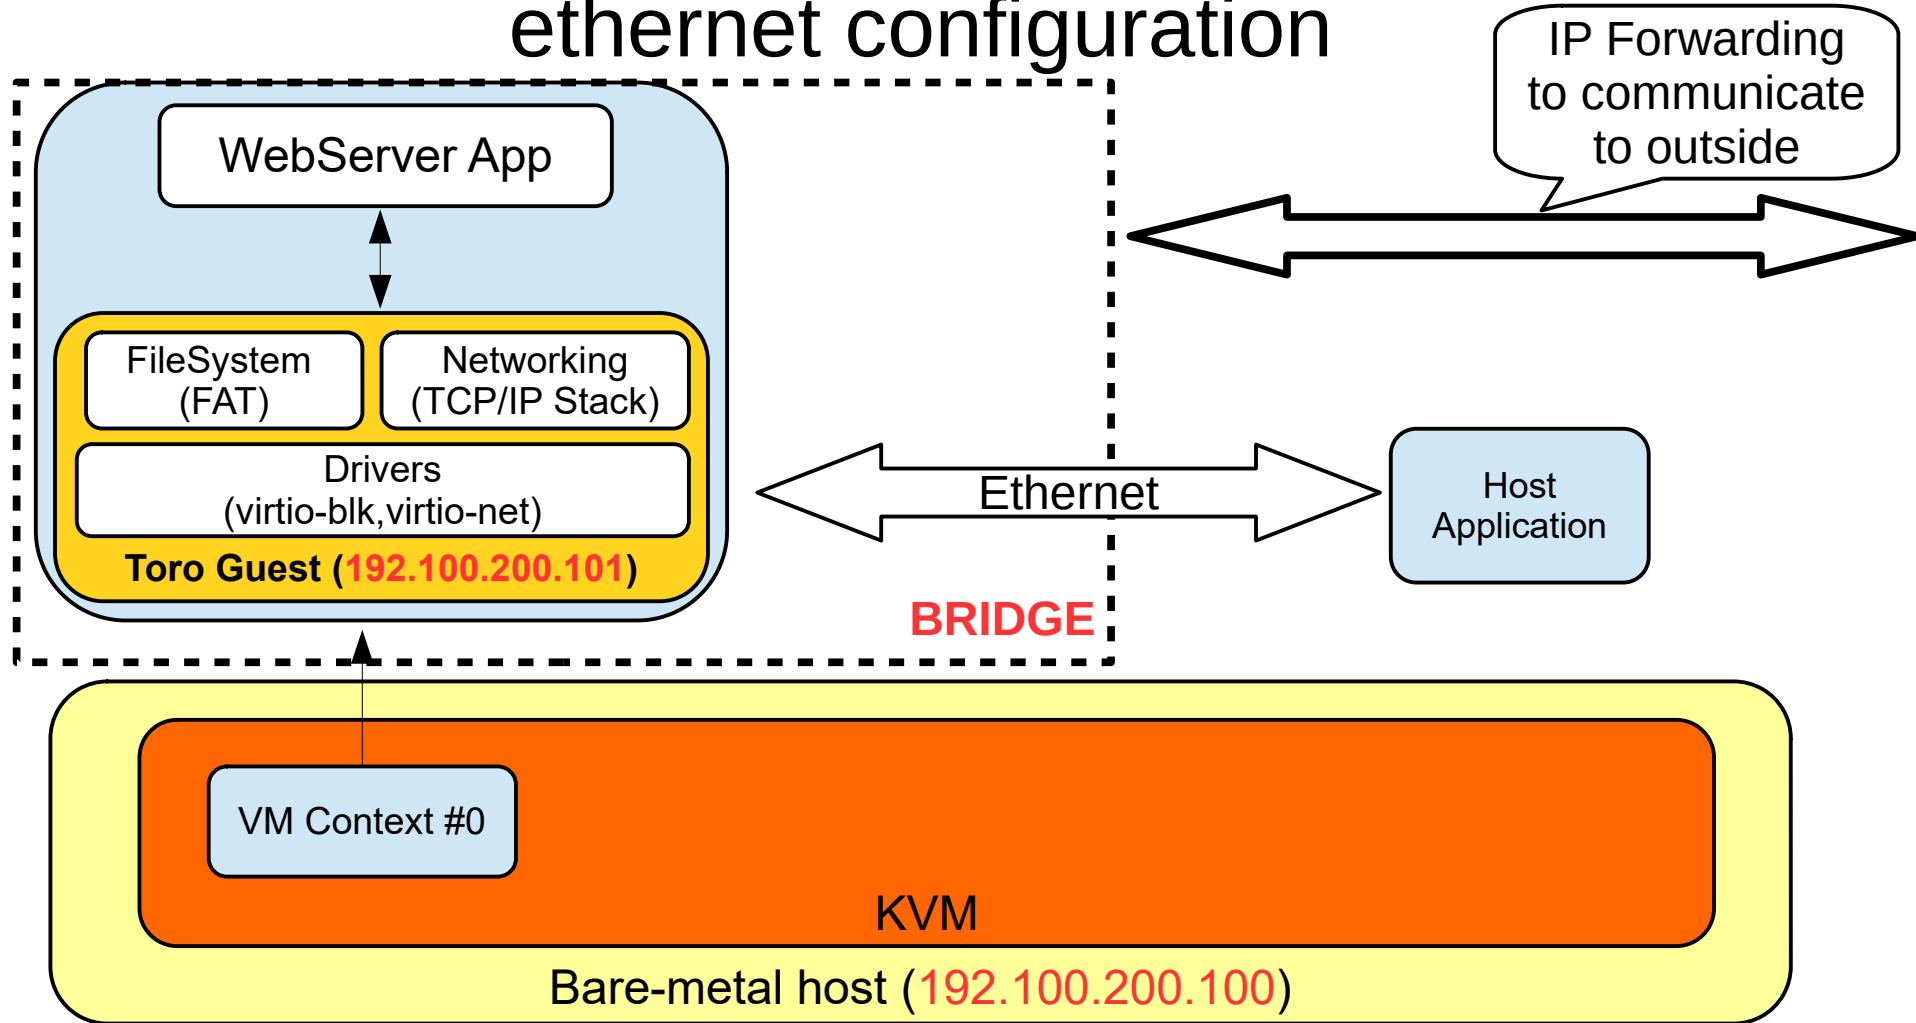
How the appliance is setup?

on-disk images



How the appliance is setup?

ethernet configuration



The Static WebServer drawbacks

- Disk images consume memory and on-disk space, e.g., each guest has its own image
- Disk images have to be distributed in all the nodes
- The use of a TCP/IP stack requires configurations, e.g., bridge, an IP per guest, guest drivers, devices
- The use of more devices increases the attack surface
- Sharing of files between guests and host is hard
- Relying on a specific FS is not good for immutable images

The Static WebServer drawbacks

- Disk images consume memory and on-disk space, e.g., each guest has its own image
- Disk images have to be distributed in all the nodes
- The use of a TCP/IP stack requires configurations, e.g., bridge, an IP per guest, guest drivers, devices
- The use of more devices increases the attack surface
- Sharing of files between guests and host is hard
- Relying on a specific FS is not good for im



Can we do it better?!

The Static WebServer drawbacks


- Disk images consume memory and on-disk space, e.g., each guest has its own image
- Disk images have to be distributed in all the nodes

We propose to use of **virtio-fs** for filesystem, **virtio-vsocket** for networking and **microvm as QEMU** machine to simplify toro unikernel's code, reduce attack surface and ease appliance configuration. Also, we propose to use **CephFS** to provide a distributed FS among VMs.


- Sharing of files between guests and host is hard
- Relying on a specific FS is not good for im

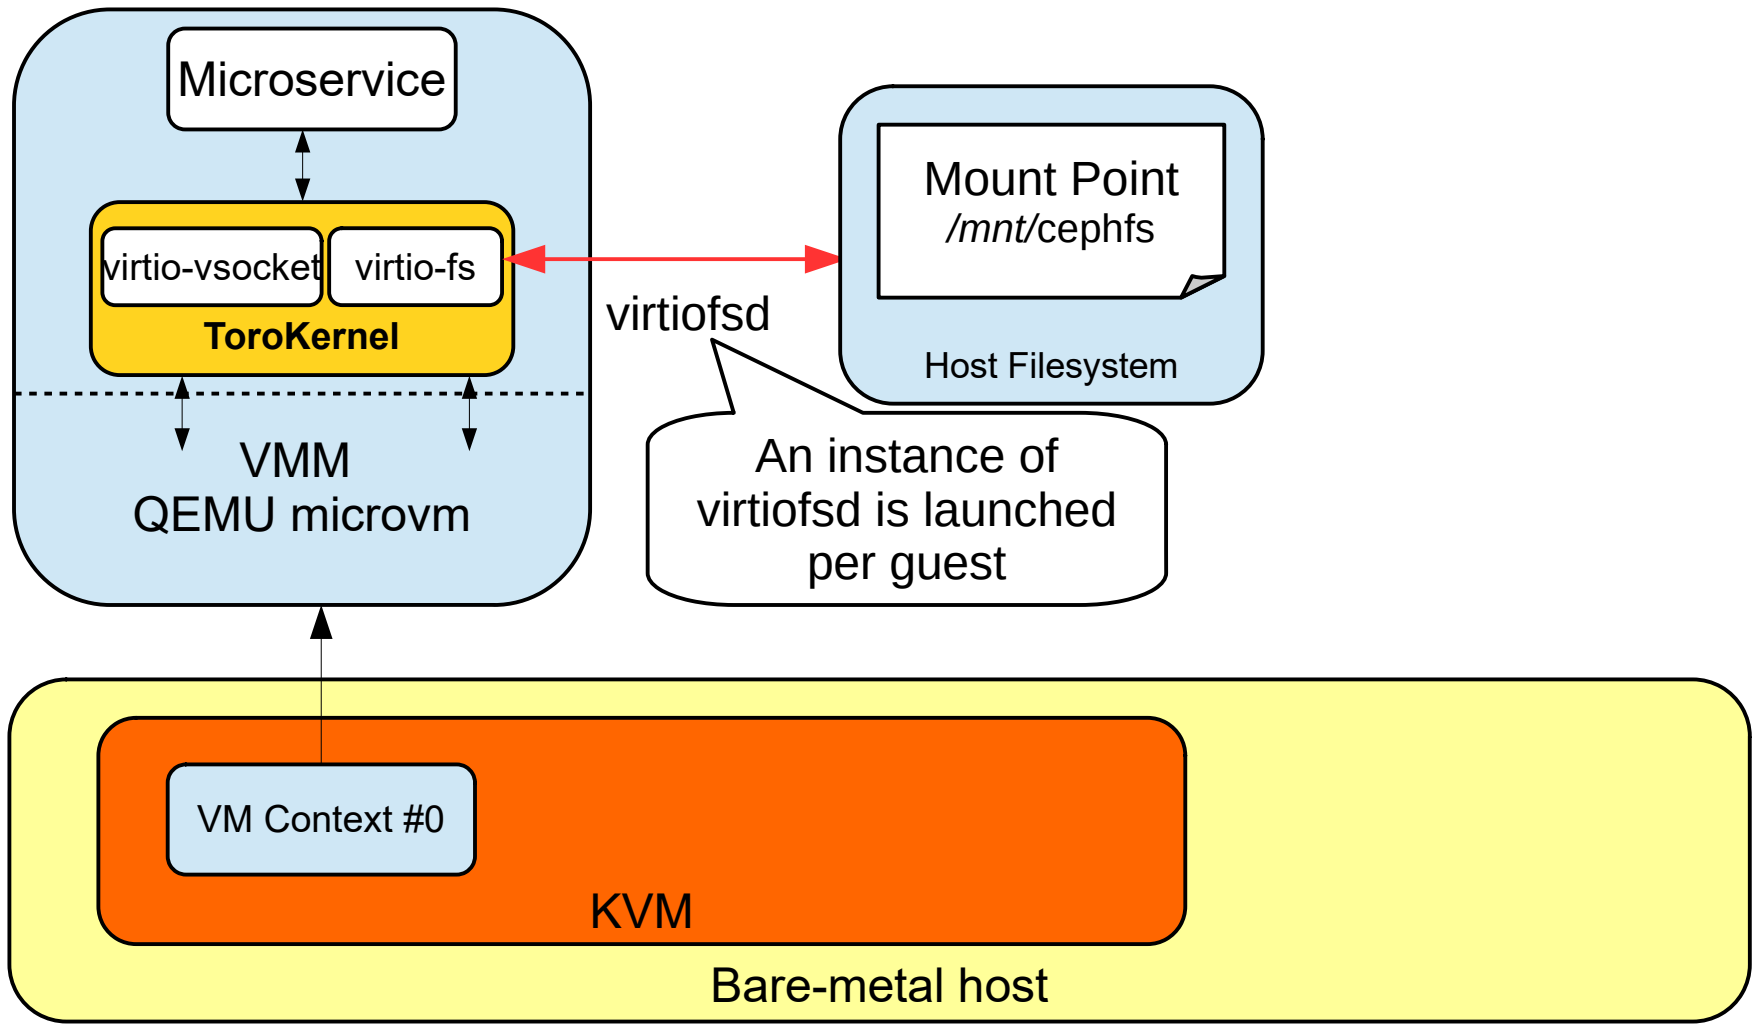


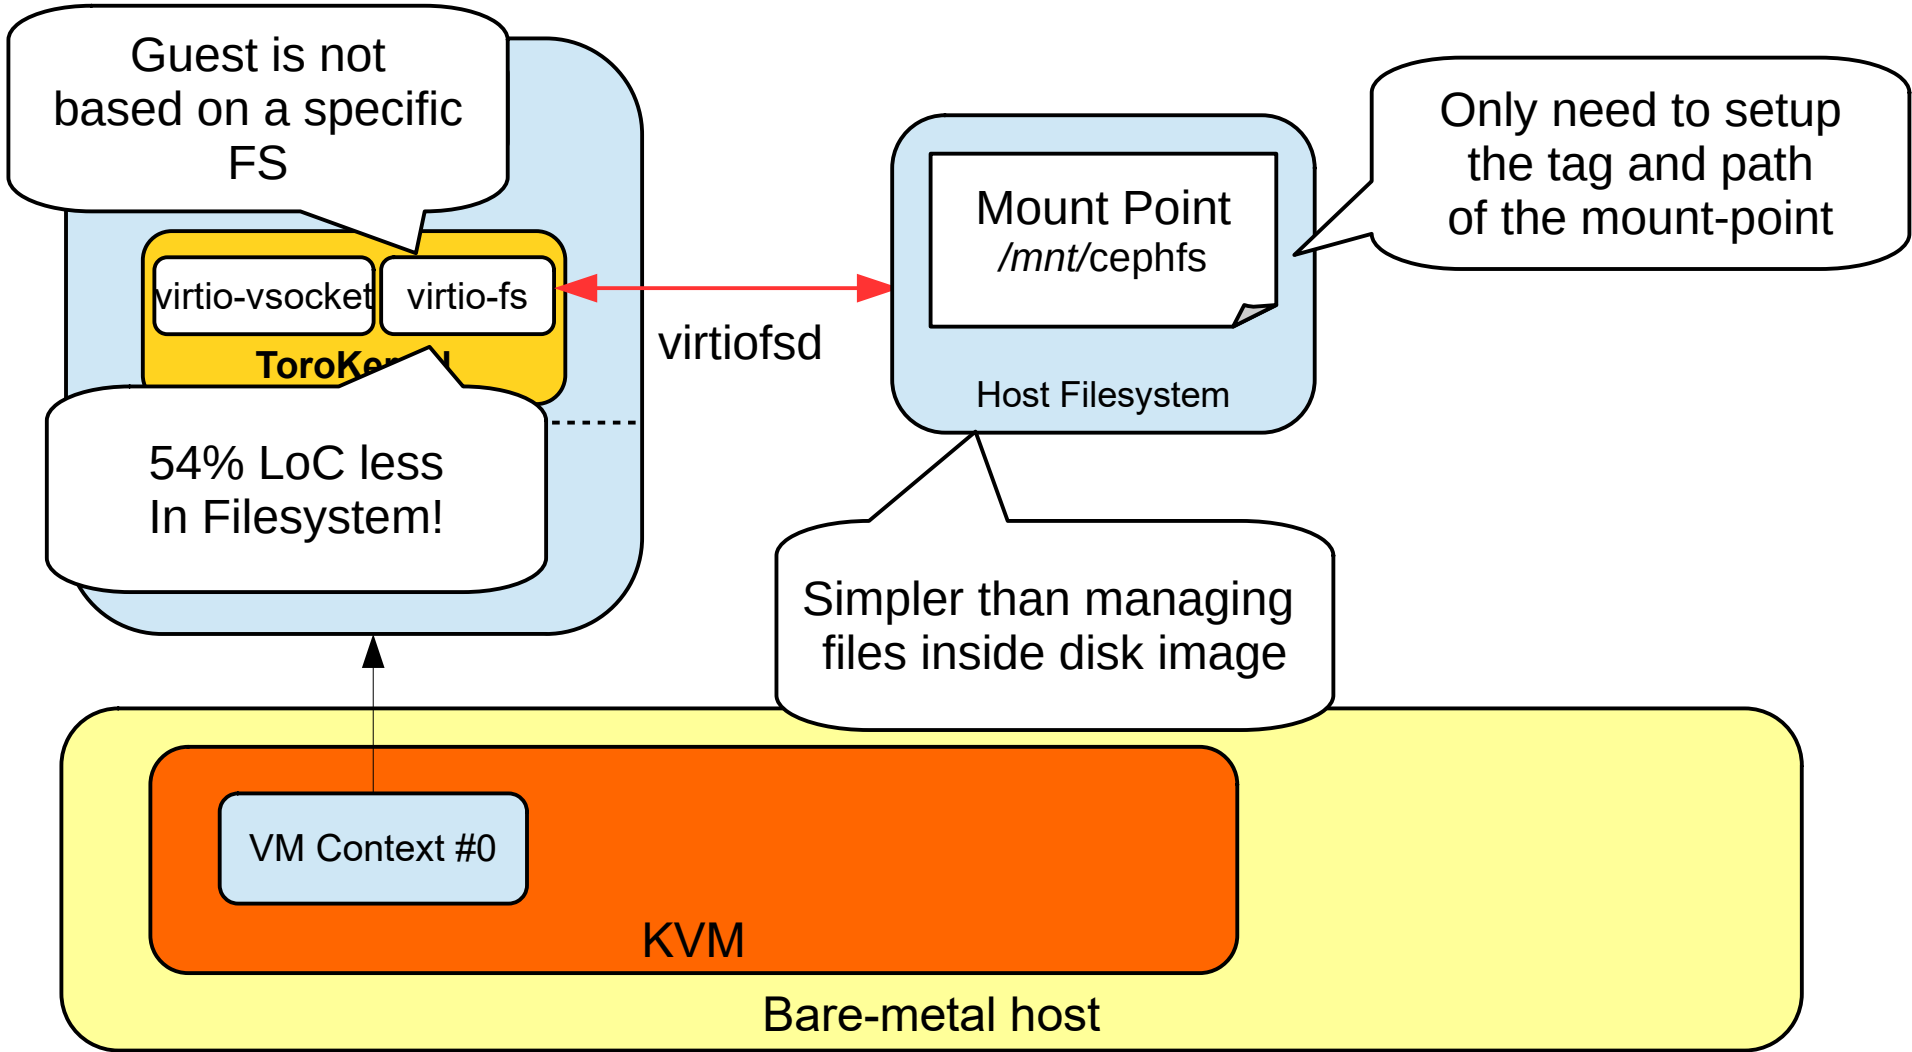
Can we do it better?!

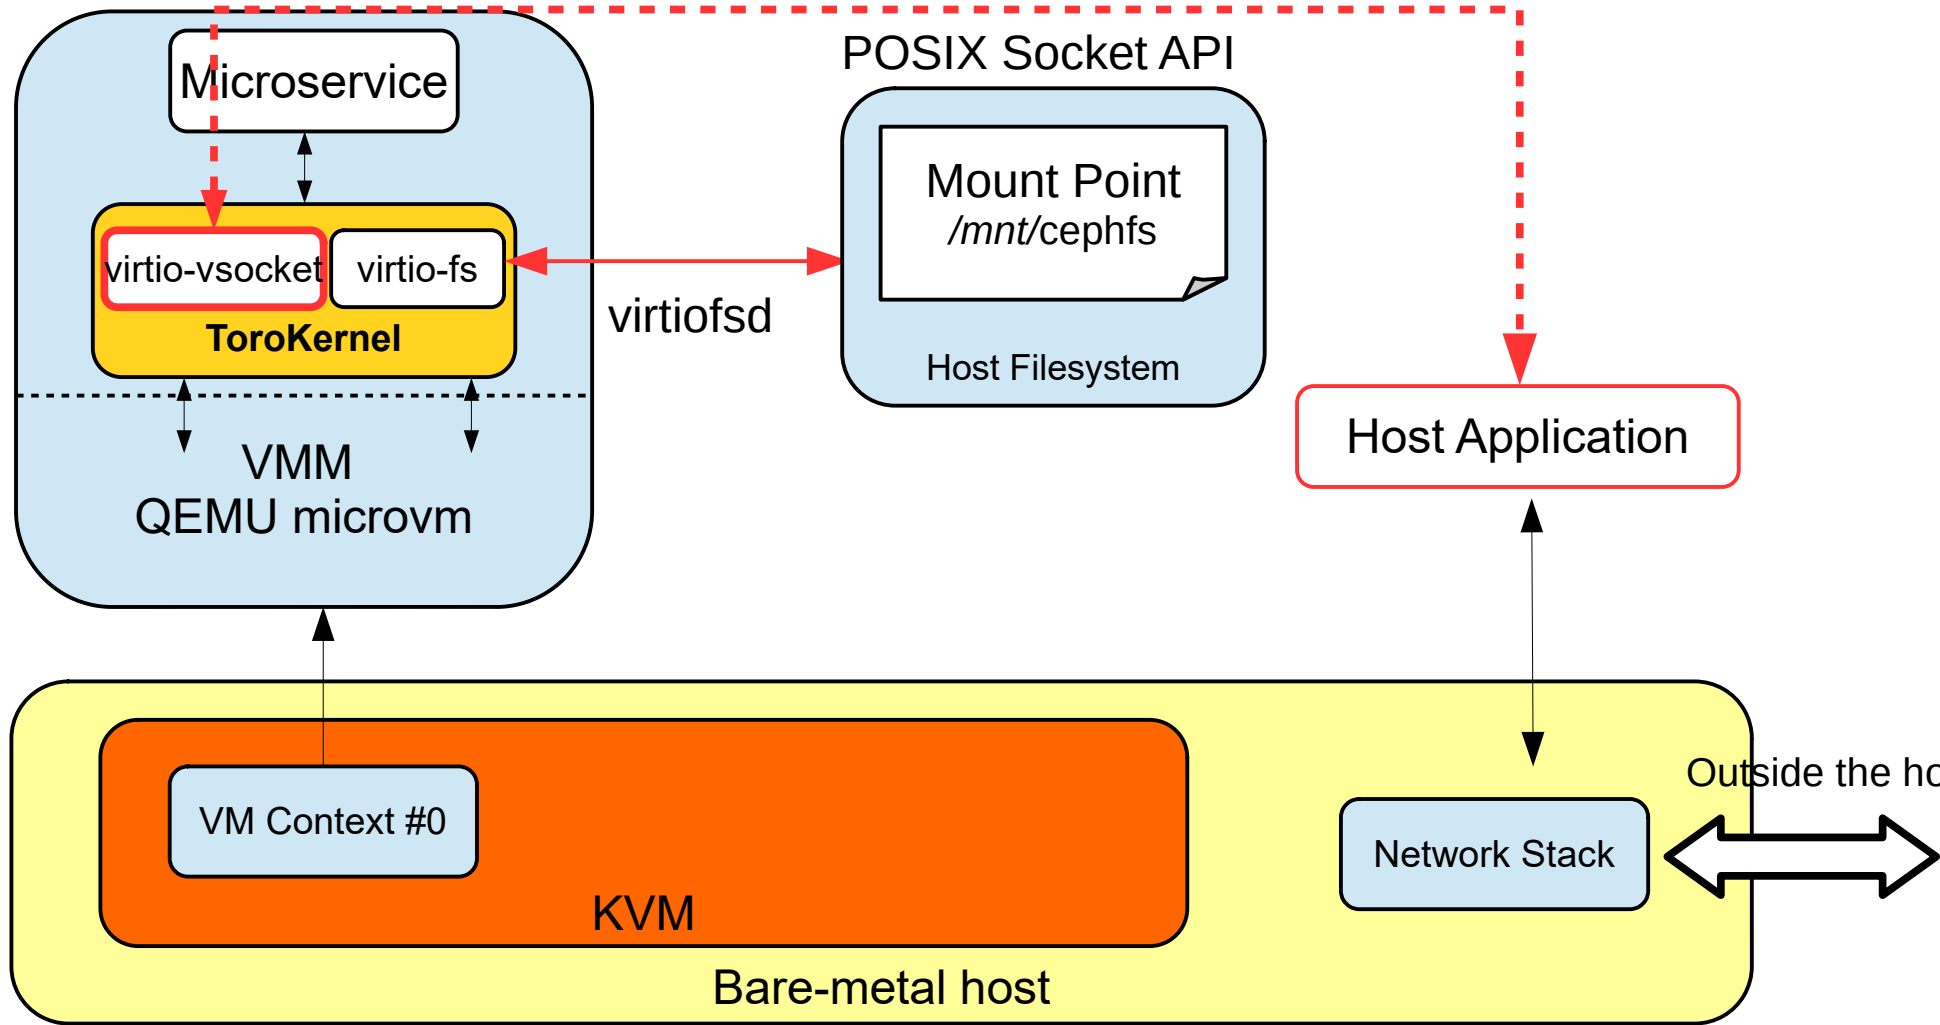


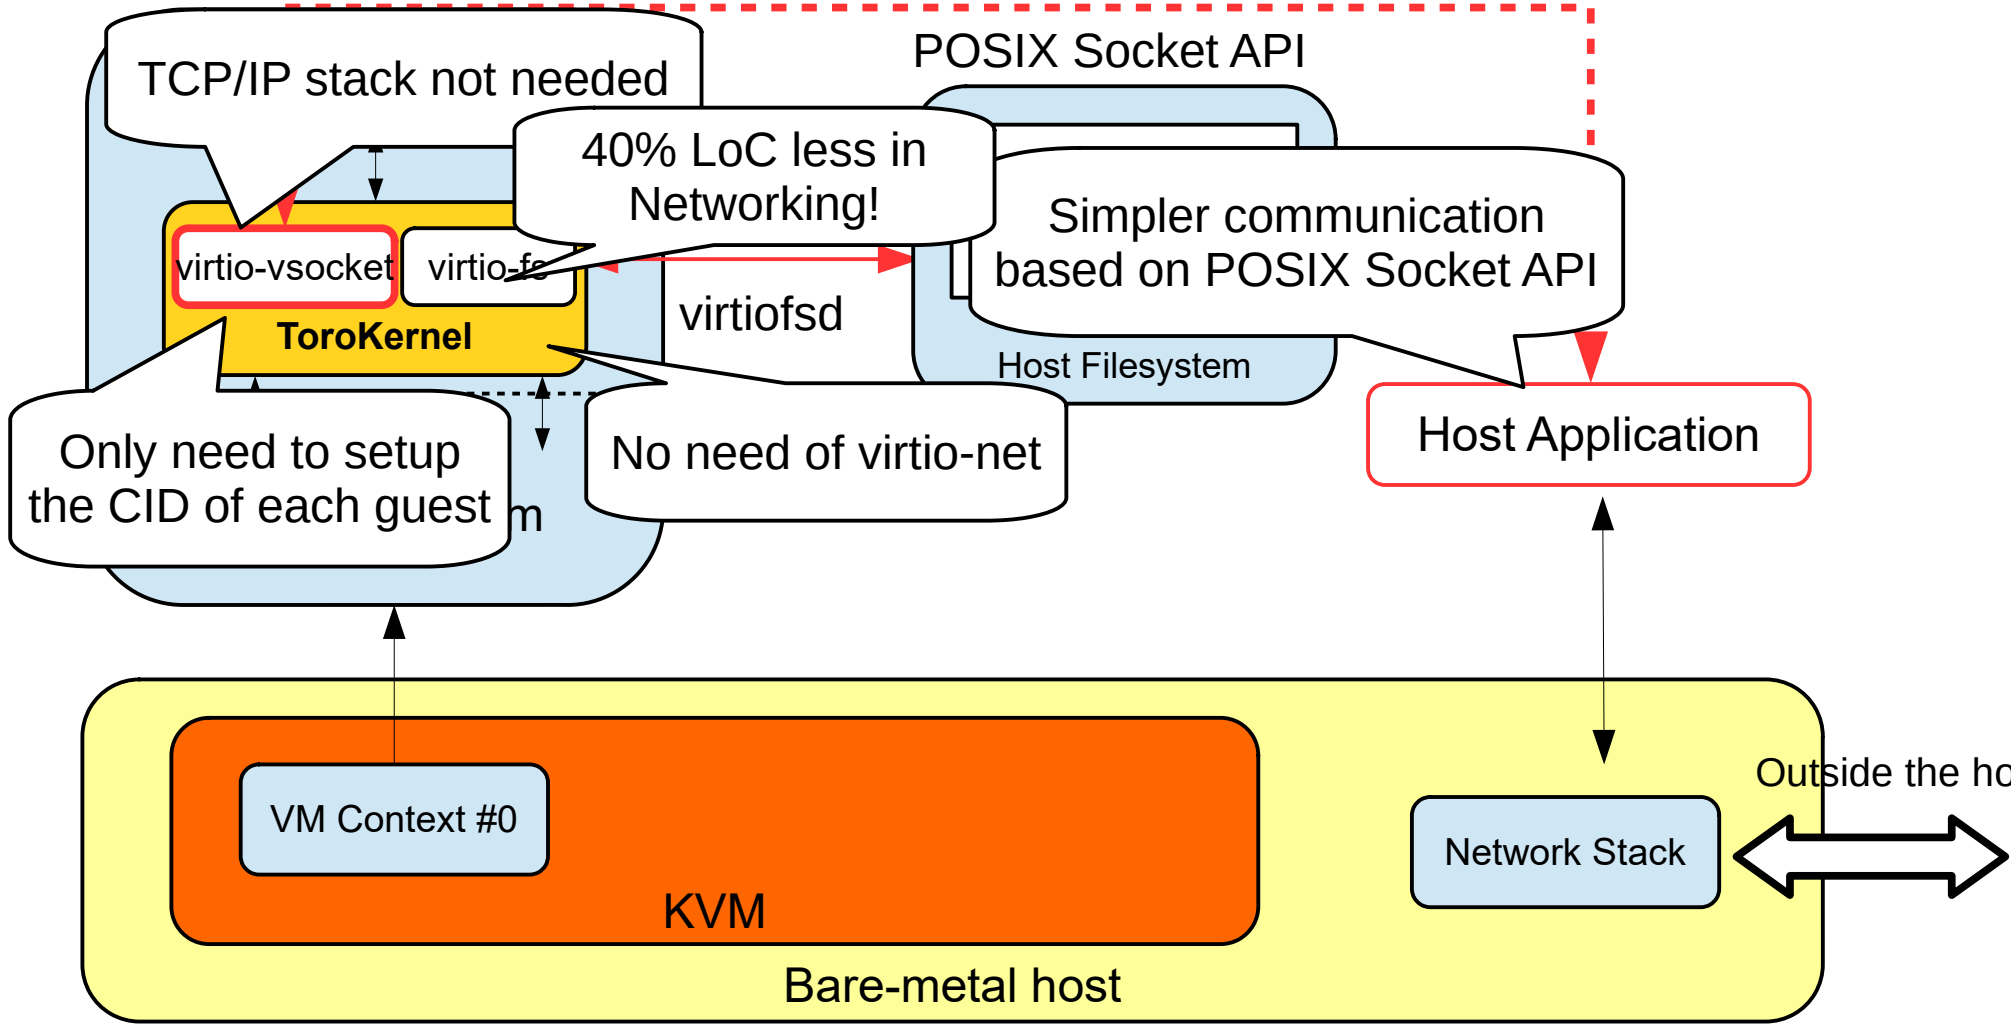
- **Virtio-fs** and **virtio-vsocket** are virtio-devices that are in QEMU since 5.0
- **Microvm** is a minimalist QEMU machine which provides a simplified device-model based on virtio. In QEMU since 5.x

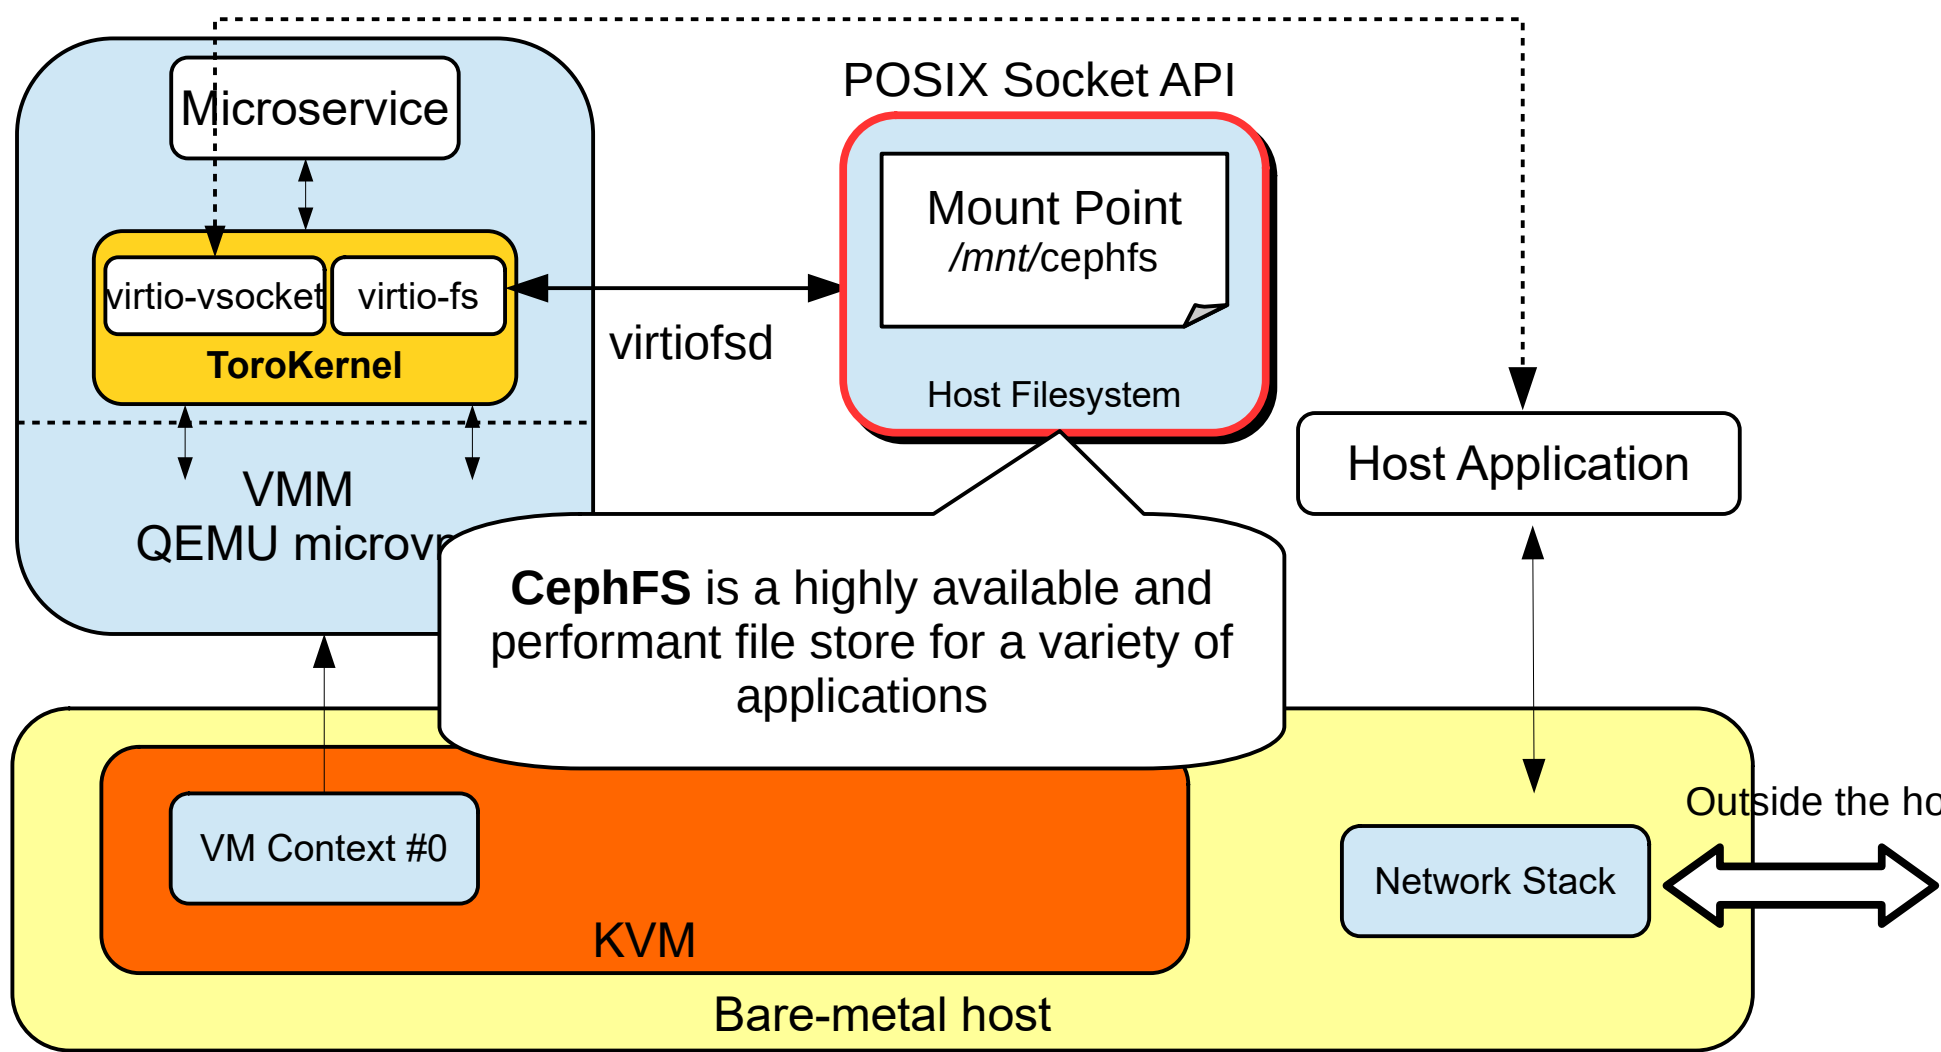




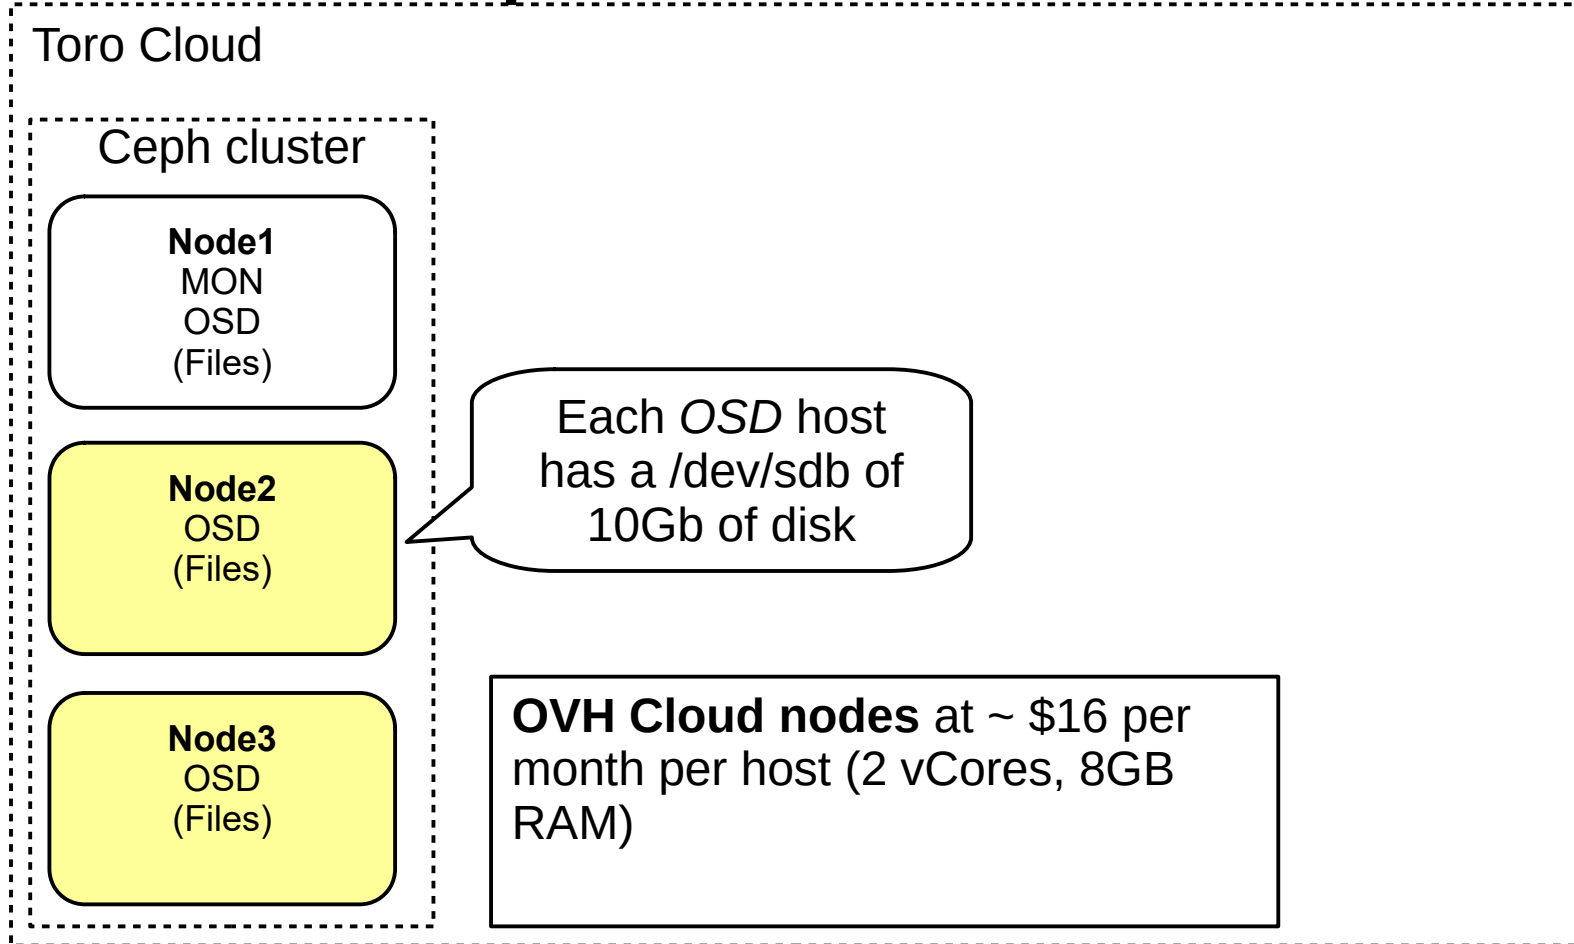




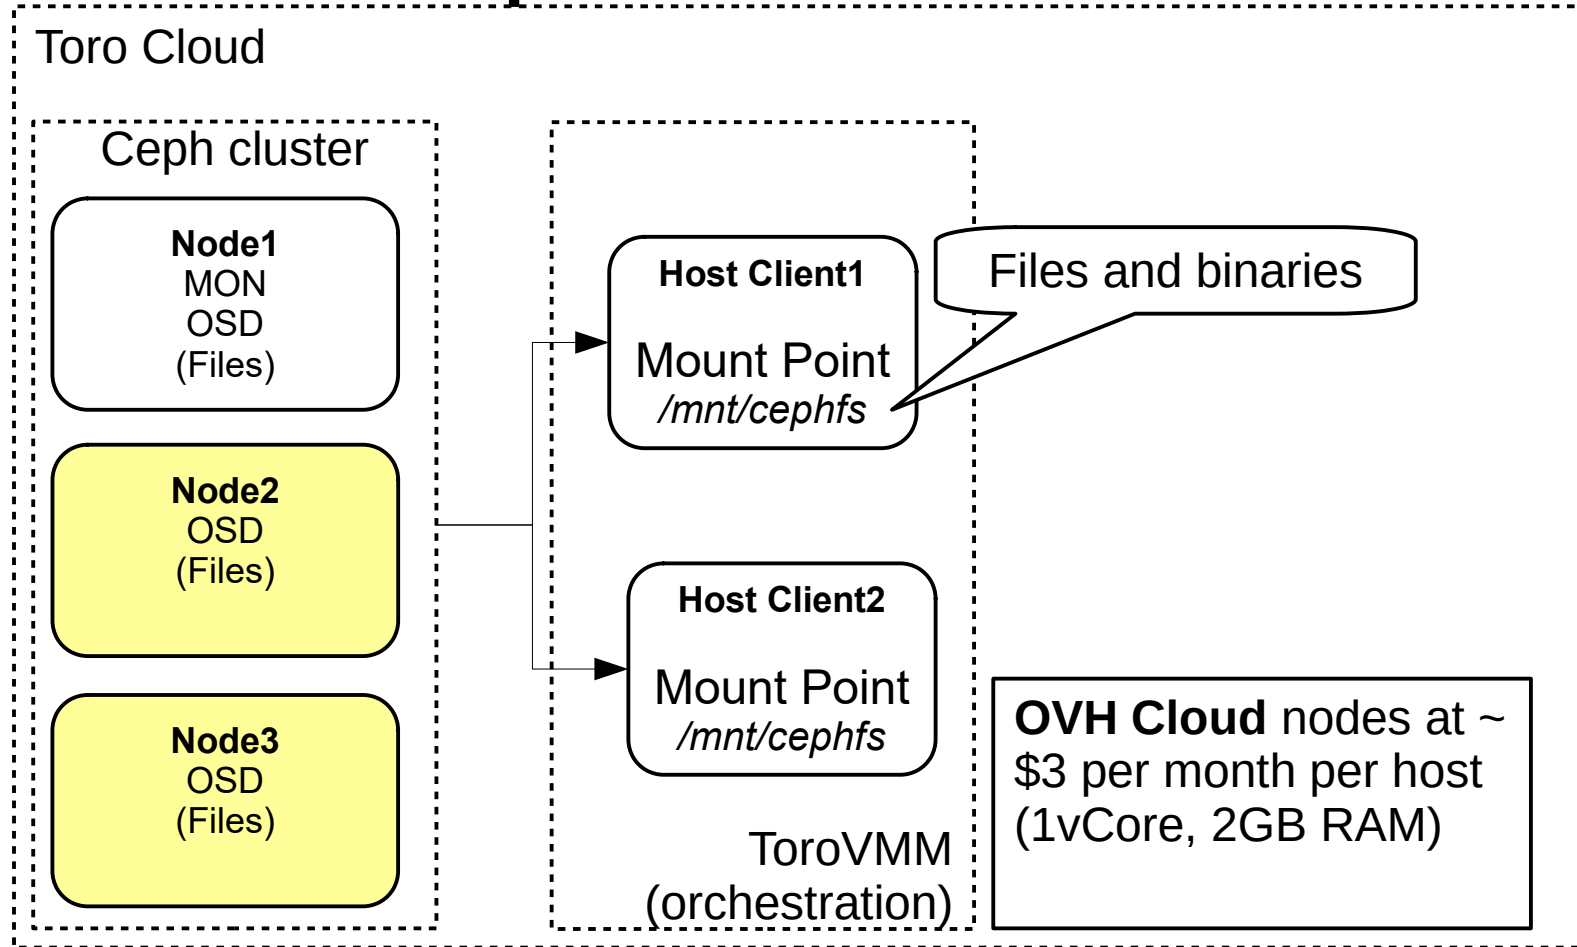




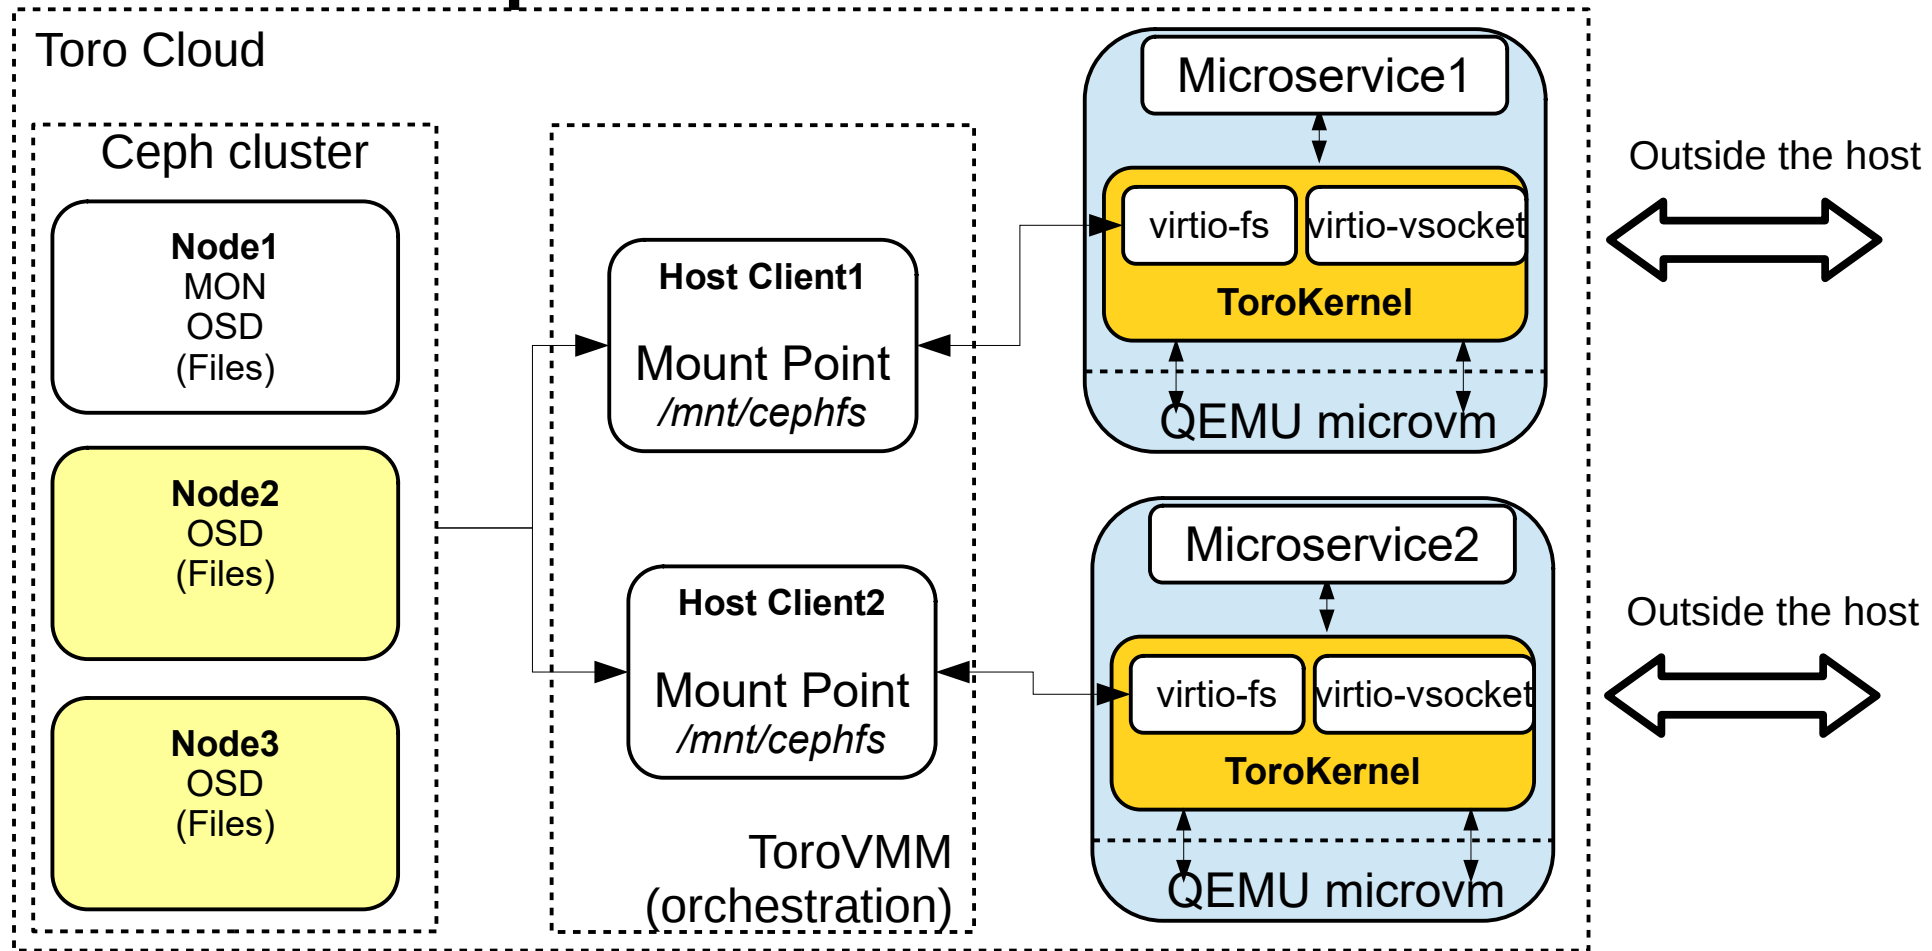
Ceph 3-node cluster



Ceph 3-node cluster



Ceph 3-node cluster



Results

- Binary Size: 235 Kb that includes kernel and user's application
- Time to rebuild the microservice: ~ 500ms
- Boot cycle: ~ 80ms
 - \$echo "Hello World" is ~2.6 ms
- CPU Usage: 90% at high and 10% sleep
- Memory footprint per VM: 2.9% (~ 60Mb) or 35 VMs per hosts
 - QEMU compiled with all enabled
- Price: 58 euros/month ~ 0.85 euros/month per VM
- See <https://github.com/torokernel/torocloudscripts>

Results

- Binary Size: 235 Kb that includes kernel and user's application
- Time to rebuild the microservice: ~ 500ms
- Boot cycle:
 - \$echo 'It is all talk until code runs.' - Ward Cunningham
- CPU Usage: 90% at high and 10% sleep
- Memory footprint per VM: 2.9% (~ 60Mb) or 35 VMs per hosts
 - QEMU compiled with all the configuration
- Price: 58 euros/month ~ 0.85 euros/month per VM
- See <https://github.com/torokernel/torocloudscripts>

Challenges

- Support live-migration which is not currently supported by microvm machine
- Improve bottleneck at vsocket forwarding
- Improve overall performance by using zero copy in virtio-fs and virtio-vsocket
- Improve evaluation by comparing with unikernels/containers/gpos

QA

