


# Scalable and Lightweight Cloud-Native Application Sandboxing

Anastassios Nanos, Charalampos Mainas,  
Georgios Ntoutsos, and Ilias Lagomatis

Nubis PC



ML SysOps



# About us



Diverse team of researchers & engineers, working on facilitating and optimizing application execution in Cloud and Edge environments.

Focus areas:

- Low-level systems software
- Hardware acceleration
- Serverless computing

Involved in Research & Innovation Actions and commercial projects



# About us

- Hardware acceleration abstractions
- Cloud-native IoT
- Container runtimes



NUBIS




nubificus

Anastassios (Tassos) Nanos  
Systems Researcher  
[ananos@nubis-pc.eu](mailto:ananos@nubis-pc.eu)



Ilias Lagomatis  
Systems Software Engineer  
[ilago@nubificus.co.uk](mailto:ilago@nubificus.co.uk)



Apostolos Giannousas  
Software Engineer  
[agian@nubificus.co.uk](mailto:agian@nubificus.co.uk)



Georgios Ntoutsos  
Software & Systems Engineer  
[gntouts@nubis-pc.eu](mailto:gntouts@nubis-pc.eu)



Charalampos (Babis) Mainas  
Research Engineer  
[cmainas@nubificus.co.uk](mailto:cmainas@nubificus.co.uk)



Kostis Papazafeiropoulos  
Systems Researcher  
[papazof@nubis-pc.eu](mailto:papazof@nubis-pc.eu)



Panos Mavrikos  
Software Engineer / K8s  
[pmavrikos@nubificus.co.uk](mailto:pmavrikos@nubificus.co.uk)




Maria-Rafaella Gkeka  
Research Engineer  
[mgkeka@nubificus.co.uk](mailto:mgkeka@nubificus.co.uk)

# Overview



- Challenges
- Lightweight applications
  - containers
  - unikernels
    - urunc/bunny (bima/pun)
- Hardware acceleration\*
  - Cloud-native integration
  - Roadmap & Plan

# Cloud-Edge-IoT continuum



# Cloud-Edge- IoT continuum



Diverse requirements at each stage of execution:

- Cloud:
  - Vast resources
  - Mostly homogeneous
- Challenges:
  - data security & privacy
  - multi-cloud management
  - interoperability & flexibility

# Cloud-Edge- IoT continuum



Diverse requirements at each stage of execution:

- Edge:
  - Lots of different devices available for the Edge
    - How to deliver applications ?
    - How to manage multi-tenancy ?
    - How to use & expose hardware accelerators ?

# Cloud-Edge- IoT continuum



Diverse requirements at each stage of execution:

- IoT:
  - Even more types of devices available
    - All with their own proprietary SDK
    - No OS – deploy applications OTA (requires manual/user intervention)




# Application Deployment & Execution



We want to execute a single application in a "containerized" environment

1. How do we sandbox this container so that it doesn't affect the rest of the infrastructure?
2. Do we need the whole systems stack (OS, all libraries, loaders) ?
  - o For instance in a serverless example: Knative




# Application Deployment & Execution




# Application Deployment & Execution




Isolation  
Boundaries



# Application Deployment & Execution




# Application Deployment & Execution



# Introducing unikernels



- A unikernel is:
  - specialized
  - single address space
  - constructed using a LibOS
- In other words:
  - Tailored for a single application
  - No separation between kernel and user space
  - Contains only what is needed



# Containers vs Unikernels



- Lightweight -> **More lightweight**
- Fast spawn time -> **Even faster**
- Portable -> **Similar**
- Scalable -> **Similar**
- Somewhat isolated -> **Truly isolated**

poor Cloud-native integration

# urunc: unikernel containers



nubificus/urunc

- CRI-compatible runtime written in Go
- Treats unikernels as processes -- directly manages applications
- Spawns unikernel VMs with generic hypervisors
- Extensible, easy to add support for more unikernel frameworks & hypervisors
- Hides complexity of unikernel framework-specific hypervisor and command line options



# urunc: unikernel OCI images



nubificus/bima

To facilitate packaging, we build a specialized image builder.

- **bima** uses a dockerfile-like syntax to create OCI images:

```
1 FROM scratch
2
3 COPY test-redis.hvt /unikernel/test-redis.hvt
4 COPY redis.conf /conf/redis.conf
5
6 LABEL com.urunc.unikernel.binary=/unikernel/test-redis.hvt
7 LABEL "com.urunc.unikernel.cmdline"='redis-server /data/conf/redis.conf'
8 LABEL "com.urunc.unikernel.unikernelType"="rumprun"
9 LABEL "com.urunc.unikernel.hypervisor"="qemu"
```

- Sample **bima** invocation:

```
$ bima build -t image:tag .
```

# Lightweight sandboxes for Serverless Functions

deploy a simple HTTP header echo  
(httpprepy) service

Point to:


<https://hellocontainer.hipeac.nbfc.io>

<https://hellors.hipeac.nbfc.io/>

<https://helloclh.hipeac.nbfc.io/>

<https://helloqemu.hipeac.nbfc.io/>

<https://hellofc.hipeac.nbfc.io/>




# Lightweight sandboxes for Serverless Functions

deploy a simple HTTP header echo  
(httpprepy) service

Point to:

<https://hellouruncfc.hipeac.nbfc.io>


<https://hellouruncqemu.hipeac.nbfc.io>



# The vAccel Framework



## High-level Architecture




# The vAccel Framework

## Image Classify Example

- operation:  
`vaccel_image_classify`
- Multiple plugins
- Chosen based on specific criteria




```
export VACCEL_BACKENDS=libtorch-gpu.so:libtorch-cpu.so:libtf-gpu.so:libjetson-inference.so  
./Application
```



# The vAccel Framework

## Remote Execution:


- Transfer inputs
- Forward operation
- Receive results



# The vAccel Framework

Cloud-native integration

- kata-containers++
- Integrate agent into the runtime



# Roadmap & plan



nubificus/urunc



nubificus/bima



nubificus/bunny



blog.cloudkernels.net

- Enhance unikernel support
- Extend urunc support to lightweight apps (kernel+init)
  - Talk @ FOSDEM2025: [Less overhead, strong isolation: Running containers in minimal specialized Linux VMs](#)
- Extend bima to bunny:
  - Build tool to produce lightweight app packages [WiP]
- Introduce WASM sandboxing
  - Talk @ FOSDEM2025: [WASM meets unikernels: Secure and Efficient Cloud-Native Deployments](#)
- Breakdown analysis of time spent in each phase of execution
- Explore confidential unikernels\*




This work is partially funded through Horizon Europe Research and Innovation actions, MLSysOps (GA: 101092912), DESIRE6G (GA: 101096466), and EMPYREAN (GA: 101136024)



**6GSNS**



# Thanks!



Contact info:

- Anastassios Nanos, [ananos@nubificus.co.uk](mailto:ananos@nubificus.co.uk)
- urunc team, [urunc@nubificus.co.uk](mailto:urunc@nubificus.co.uk)

Source code:

- GitHub org: <https://github.com/nubificus>
- urunc:  [nubificus/urunc](https://github.com/nubificus/urunc)
- bima:  [nubificus/bima](https://github.com/nubificus/bima)
- kata:  [kata-containers/kata-containers](https://github.com/kata-containers/kata-containers)











