

Virtualization at the Edge

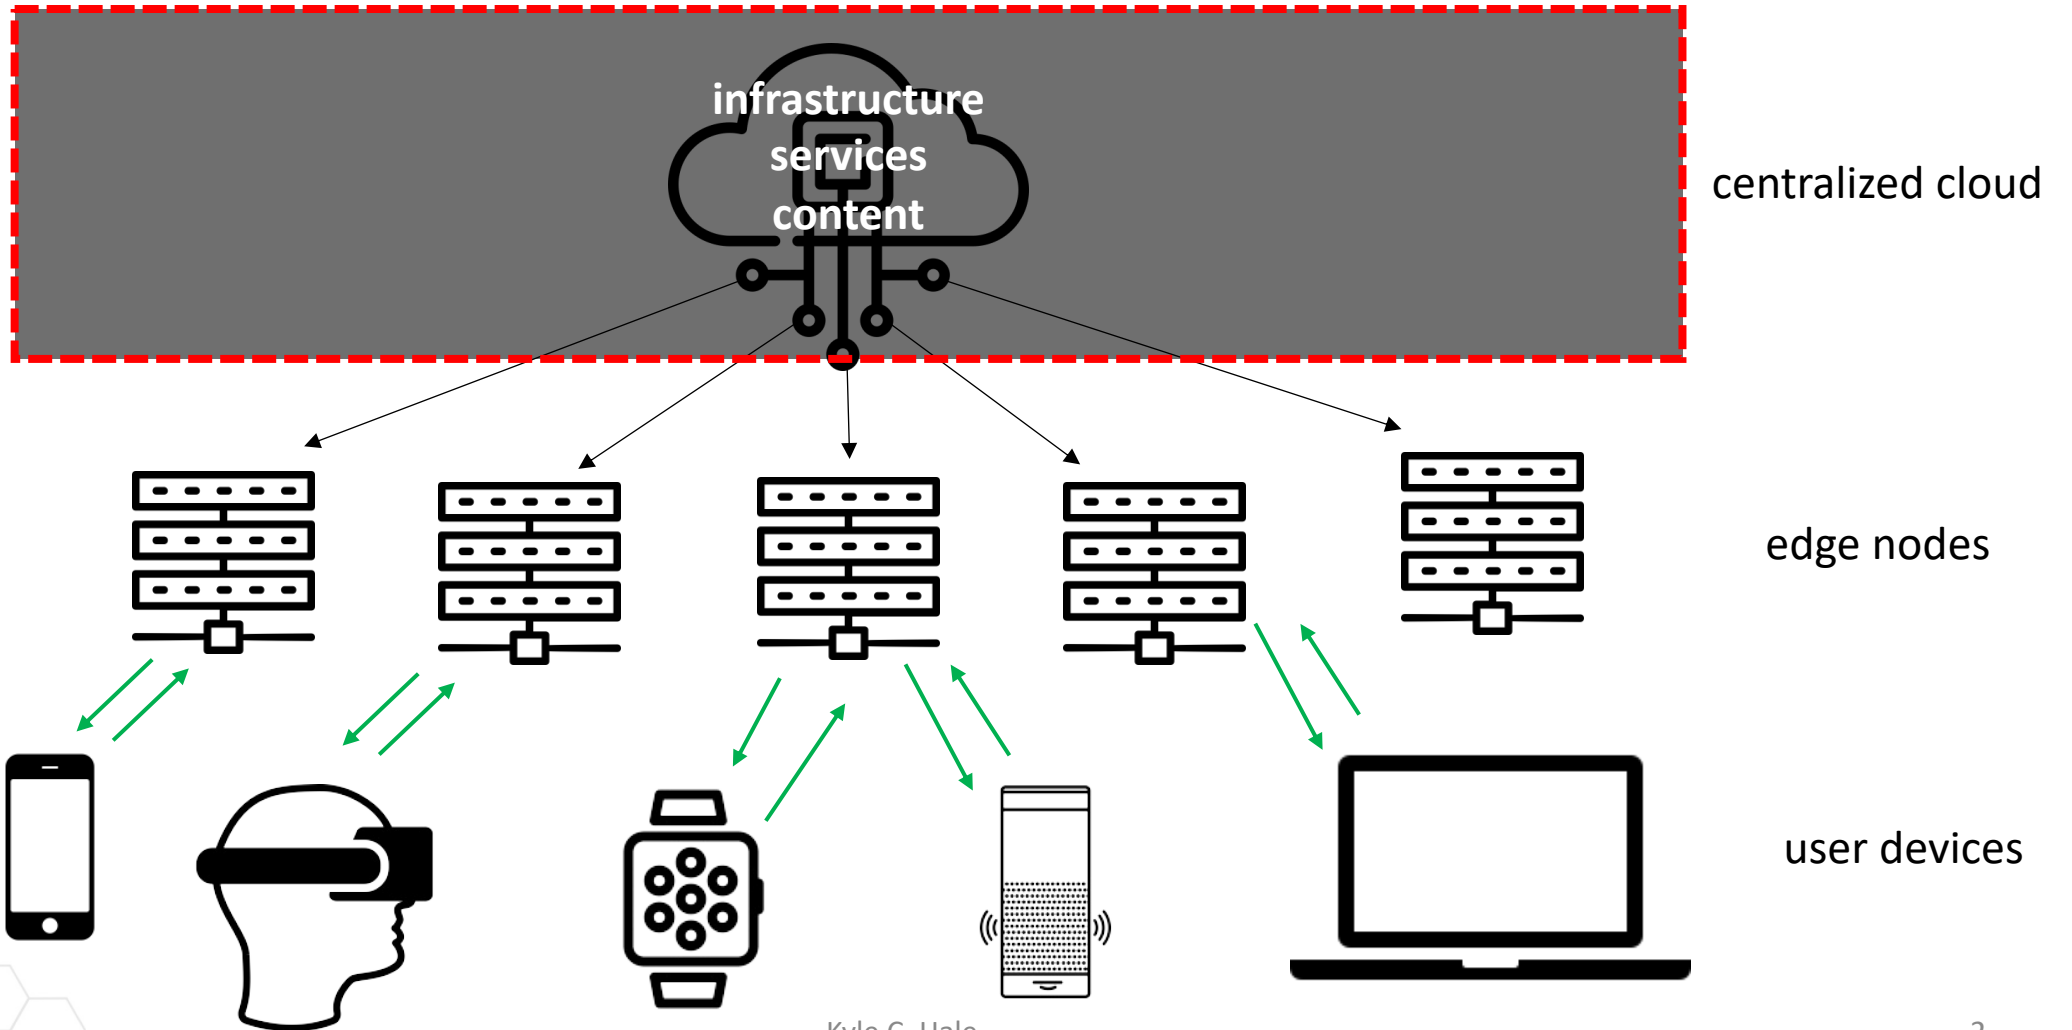
Kyle C. Hale

Laboratory for High-Performance Experimental Systems and Architecture (HEXSA)

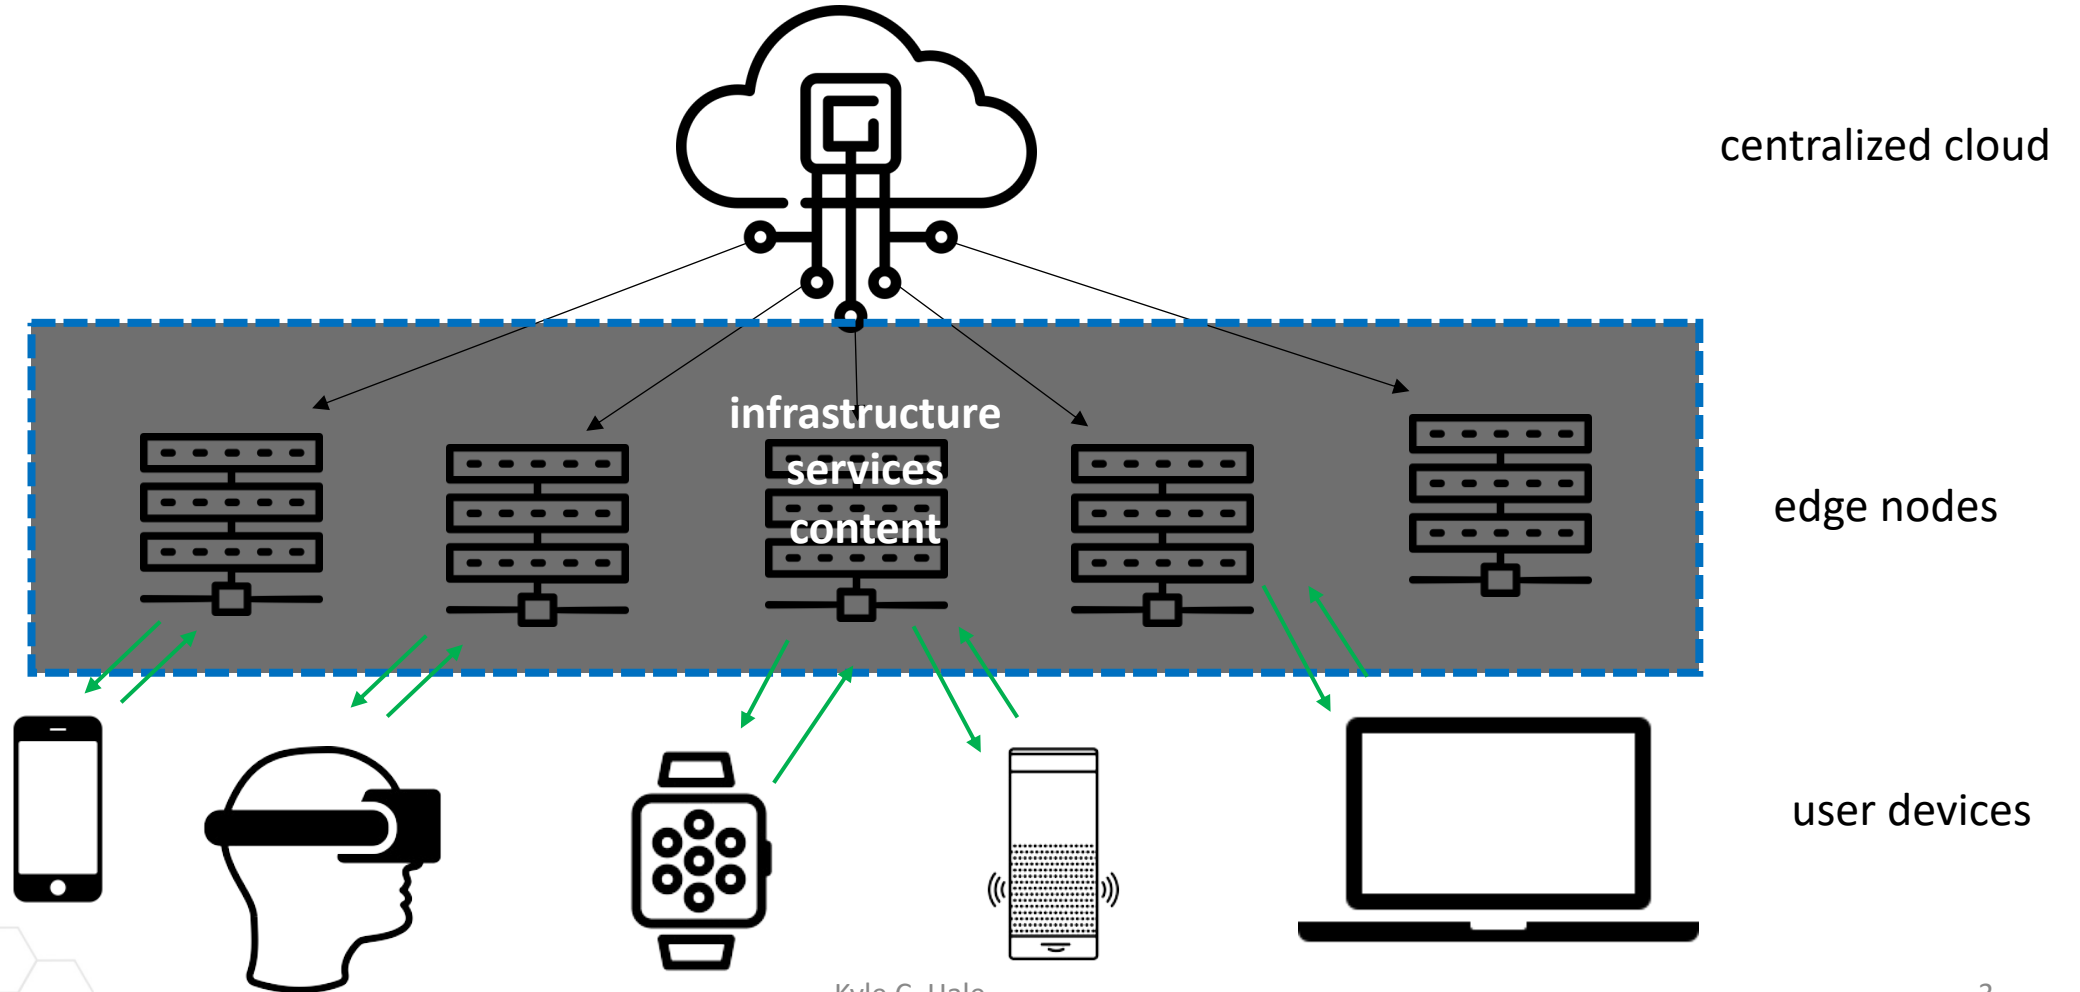
Argonne/IIT Research Seminar - March 30, 2022



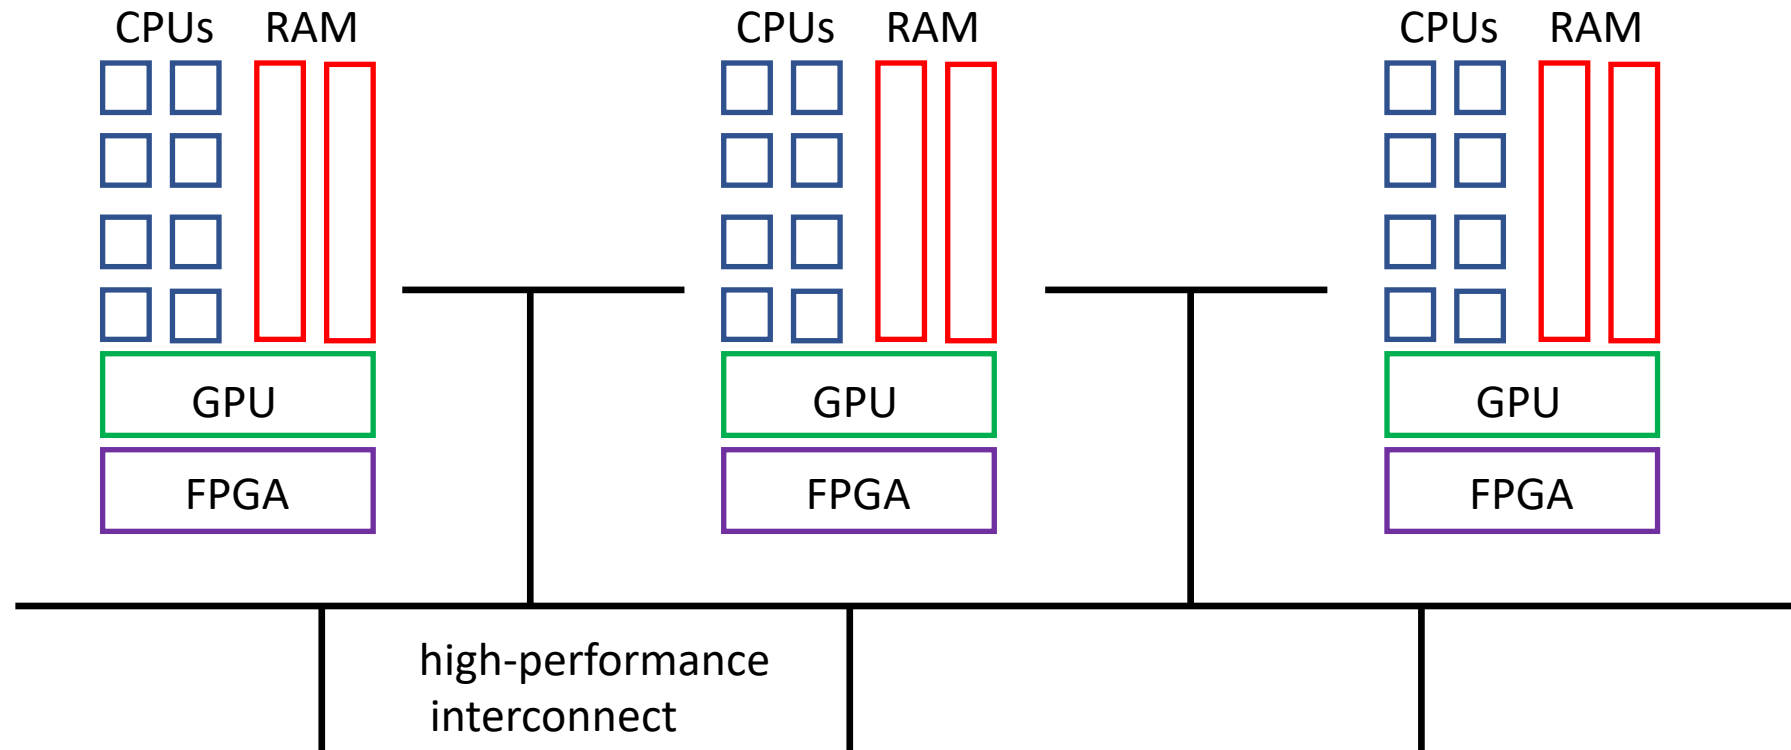
The Cloud is Decentralizing: Services and Infrastructure at the Edge



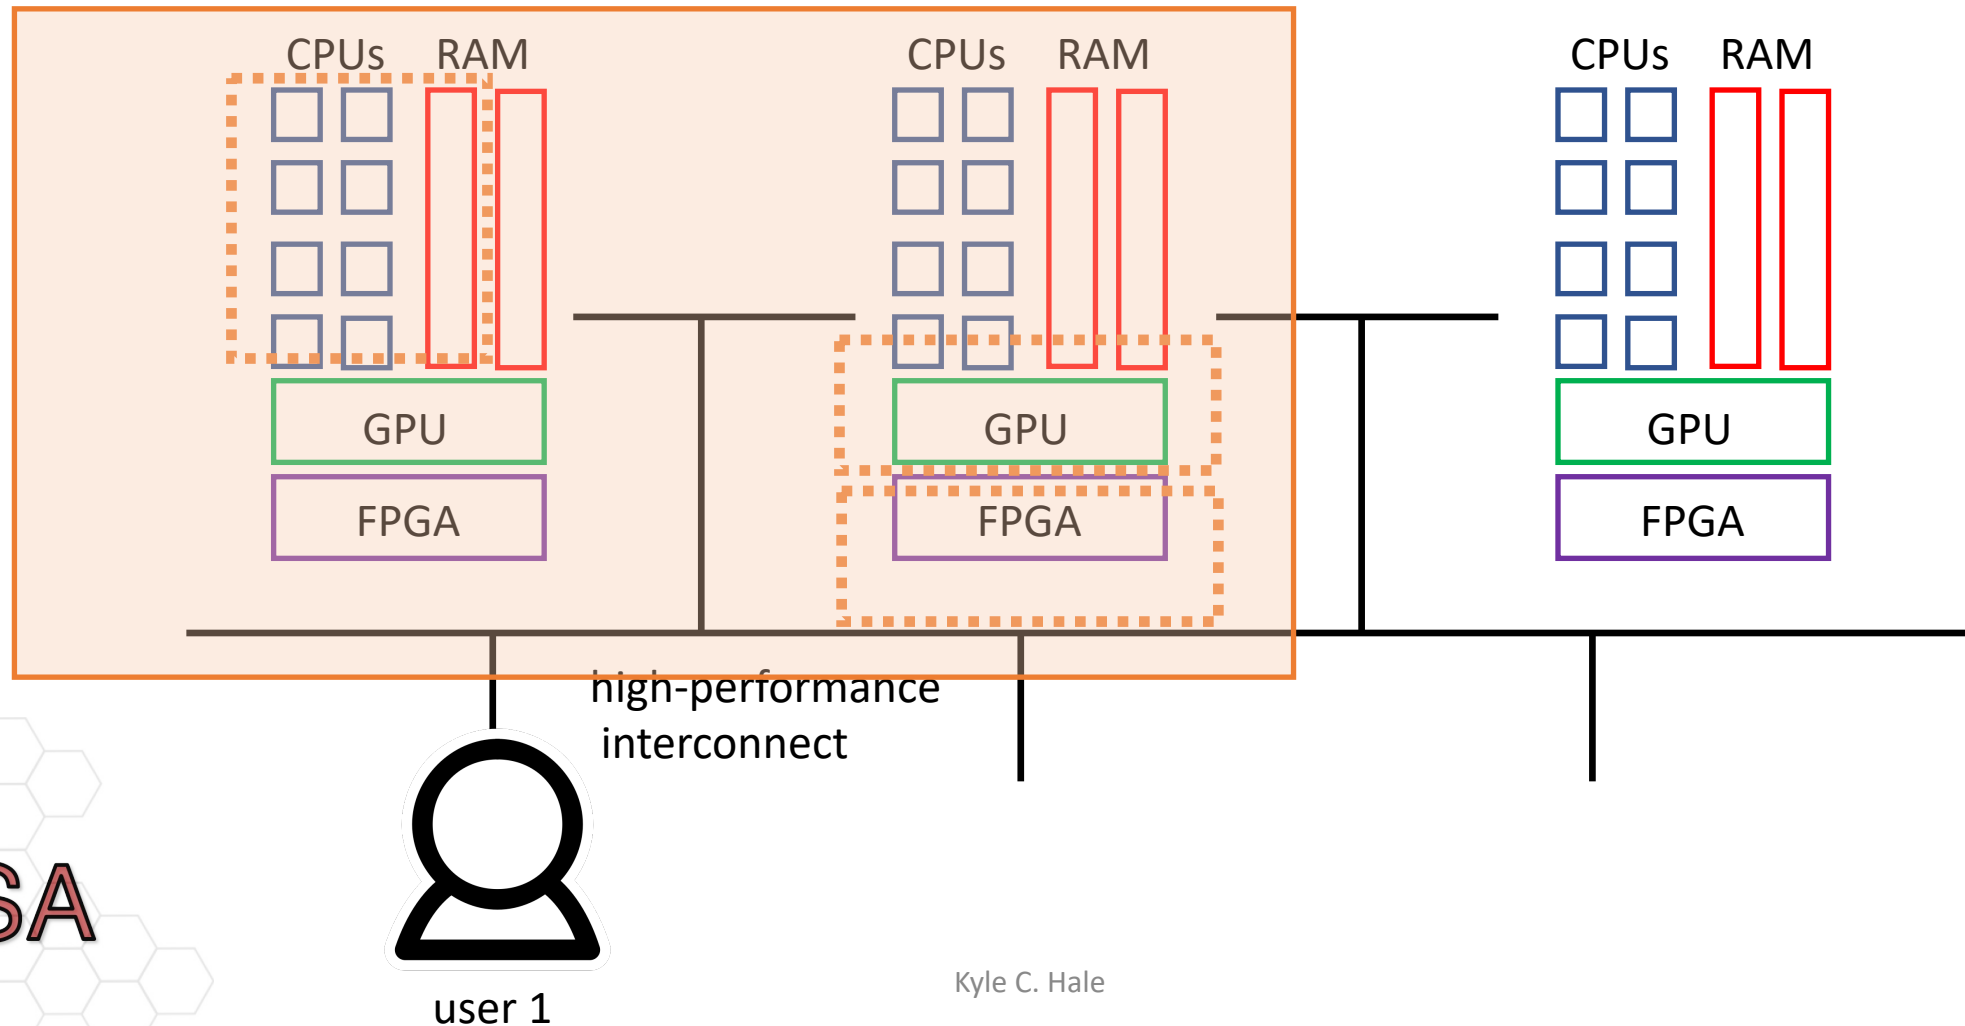
The Cloud is Decentralizing: Services and Infrastructure at the Edge



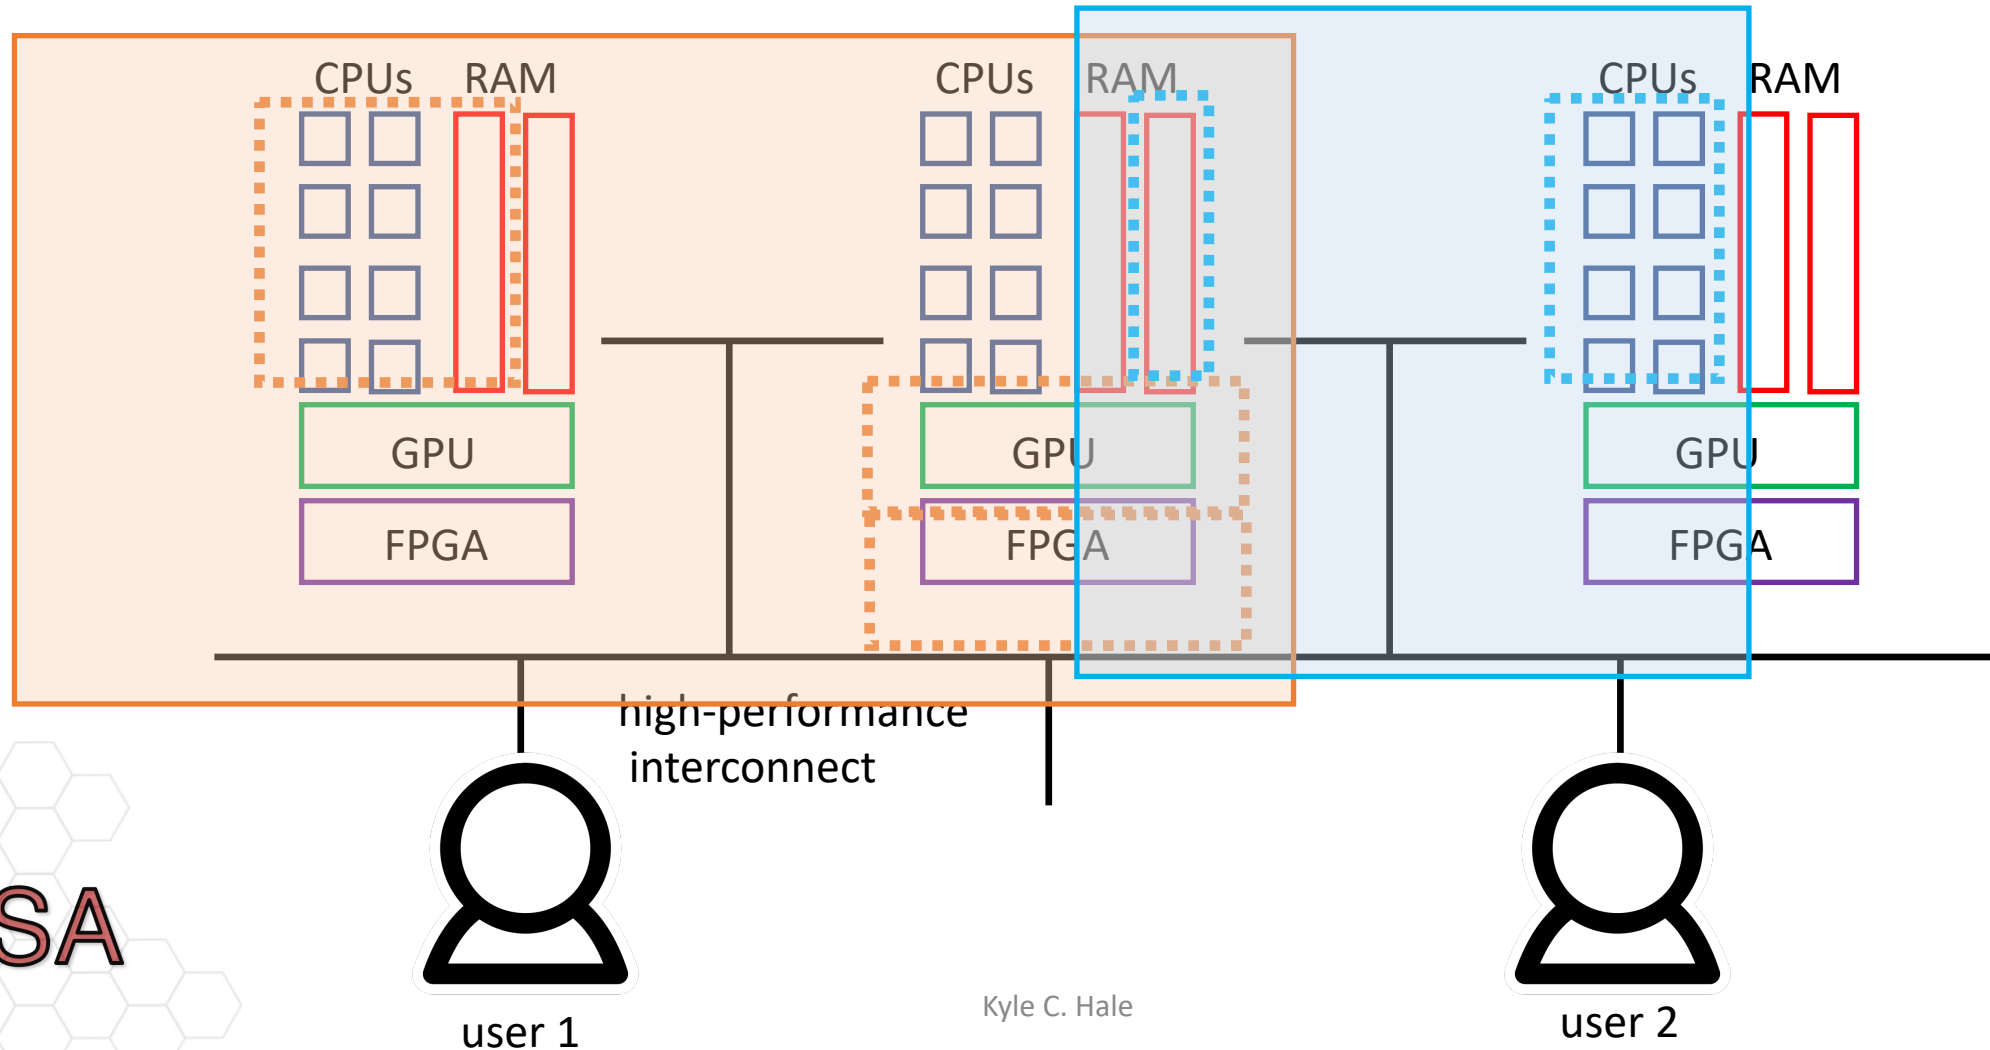
Resources are becoming *disaggregated* in the datacenter



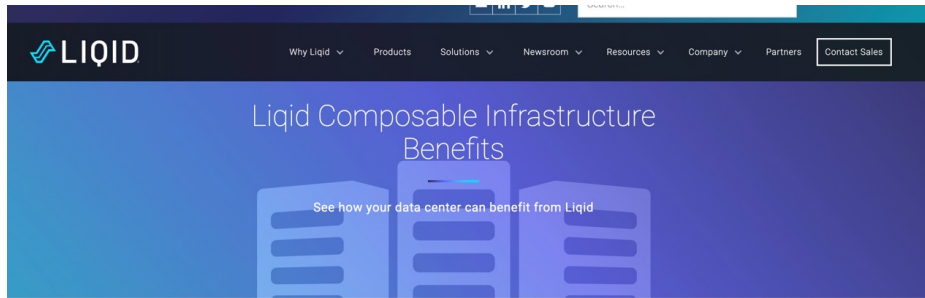

Resources are becoming *disaggregated* in the datacenter



Resources are becoming *disaggregated* in the datacenter



Composable Infrastructure

Accelerate Time-to-Value
Meet Exacting IT Demands in Real-time

- Dynamically configure servers in seconds
- Meet rigid workload requirements

source: <https://www.liqid.com/why-liqid/benefits>

The screenshot shows the Western Digital Store website. The header includes 'Western Digital Store', 'My Cloud', 'ibi', 'Store Login', and a search icon. Navigation links include 'SHOP', 'EXPLORE', and 'SUPPORT'. The main content area features the title 'OpenFlex Composable Infrastructure from Western Digital' with a 'Contact Us' button. Below this is a section titled 'The Future of Data Infrastructure' with a paragraph of text and an image of a server rack.

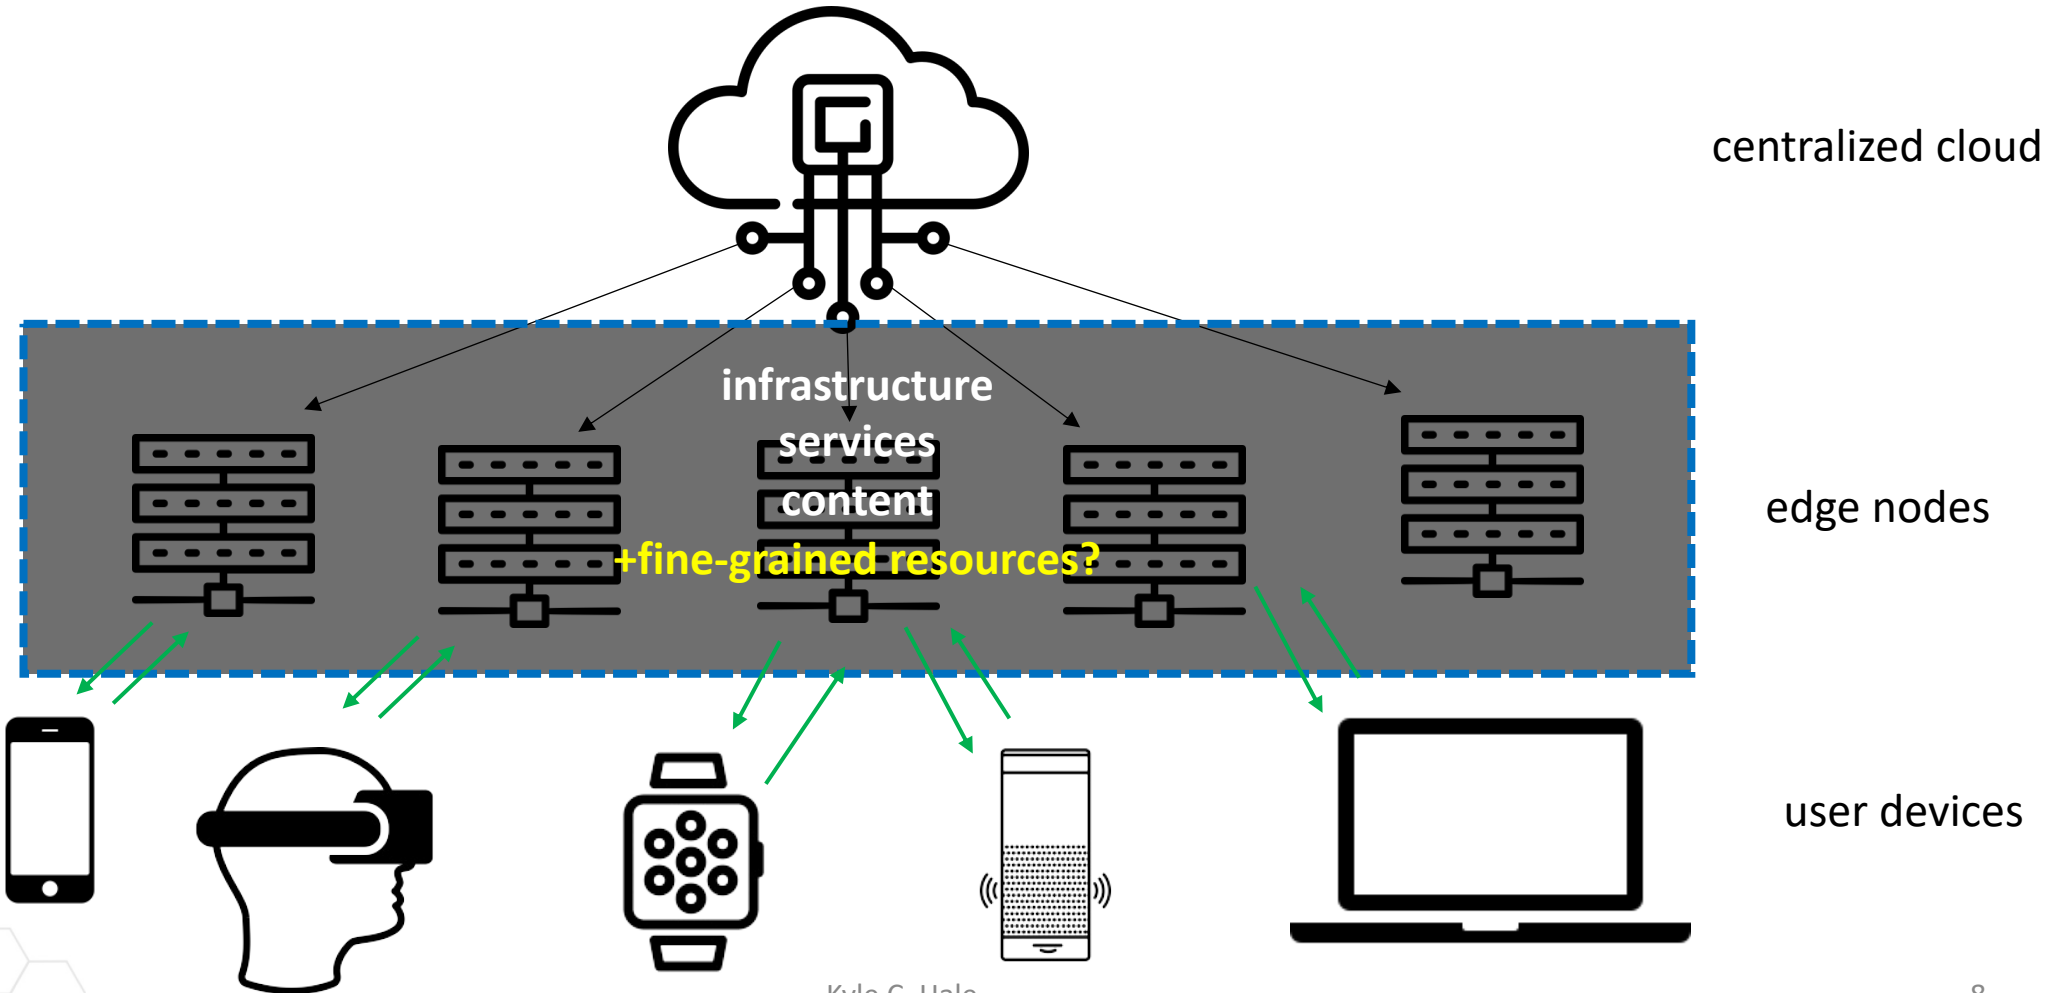
source: <https://www.westerndigital.com/products/data-center-platforms/openflex-composable-infrastructure>

The screenshot shows the HPE GreenLake website. The header includes 'HPE GREENLAKE' and navigation links: Overview, Portfolio, Pricing, Demos and Resources, and Get Started. The main content area features the title 'HPE GREENLAKE FOR COMPOSABLE COMPUTE' and the sub-headline 'Compose workloads at cloud-like speed in a pay-per-use model.' A 'Get pricing ->' button is visible at the bottom.

<https://www.hpe.com/us/en/greenlake/composable-compute.html>



Disaggregation at the Edge



To date, there have been no compelling mass-market applications that require low latencies that cannot be achieved as a web service and that also are too computationally or energy intensive for modern smart phones to run locally. This might be a “chicken or the egg” problem: the lack of cyber foraging infrastructure could potentially be hindering the development of such applications. Later, we discuss one emerging class of applications that could prove to be the compelling application that cyber foraging needs.

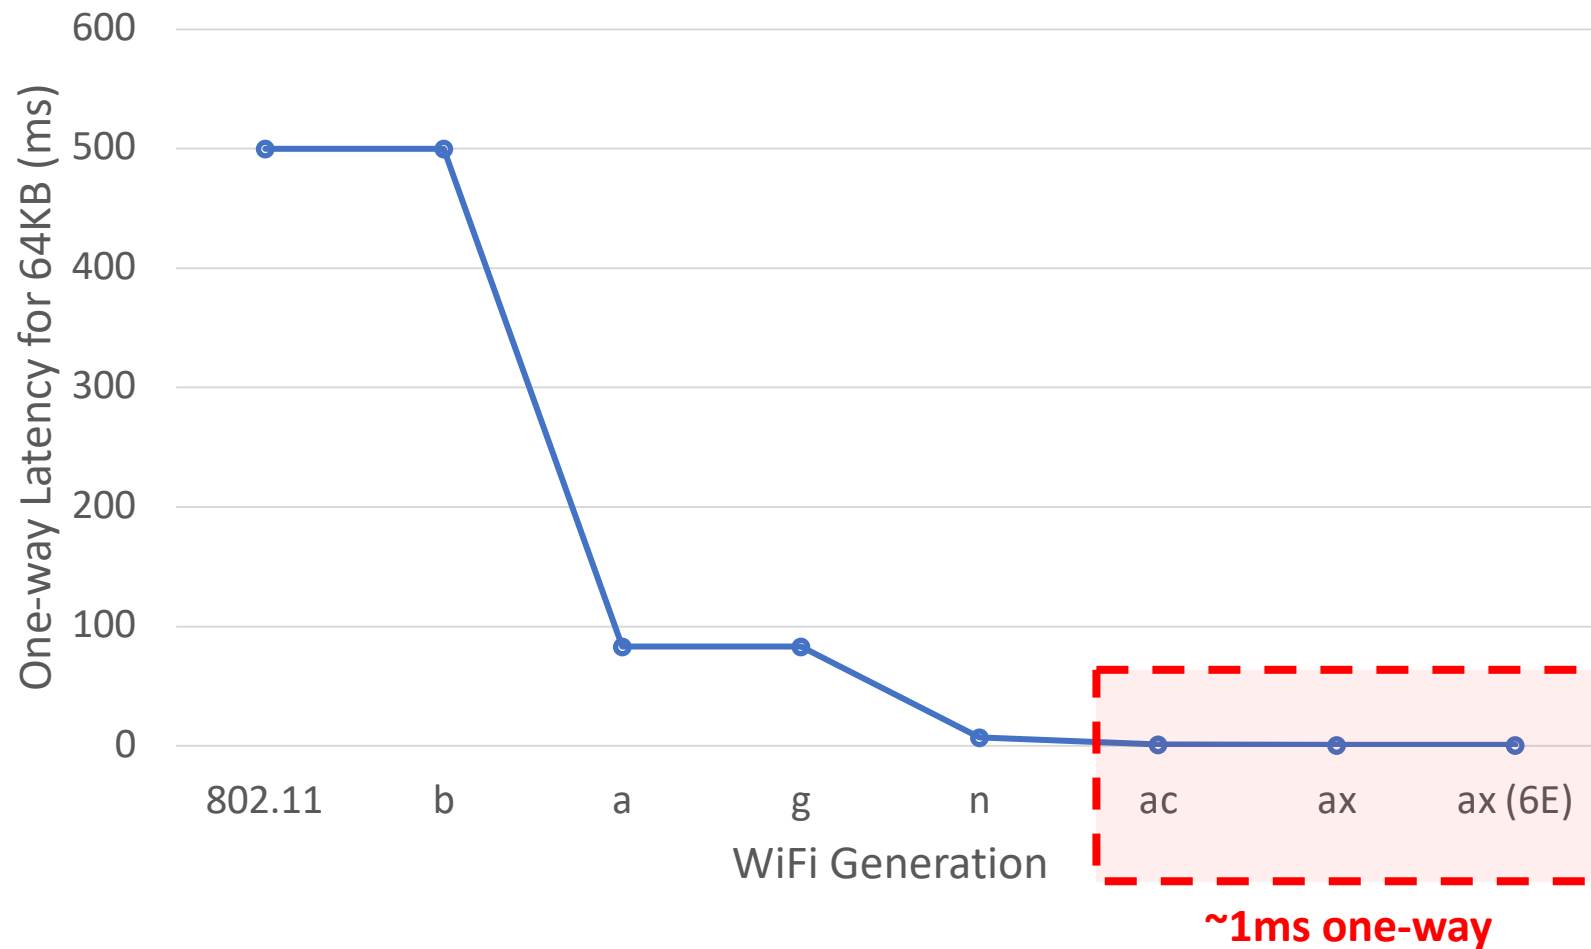
R.K. Balan and J. Flinn, “Cyber Foraging Fifteen Years Later,” IEEE Pervasive Computing, 16(3), July 2017.

What's Changed?

- Virtualization technology has improved significantly
- Infrastructure provisioning has become more sophisticated (NB serverless research)
- Composable infrastructure
- Hardware design more democratic
- AR/VR/XR is here






...also, wireless latency continues to drop





Coalescent Computing

Ephemeral Single-System Image at the Edge



Coalescent Computing

Ephemeral Single-System Image at the Edge


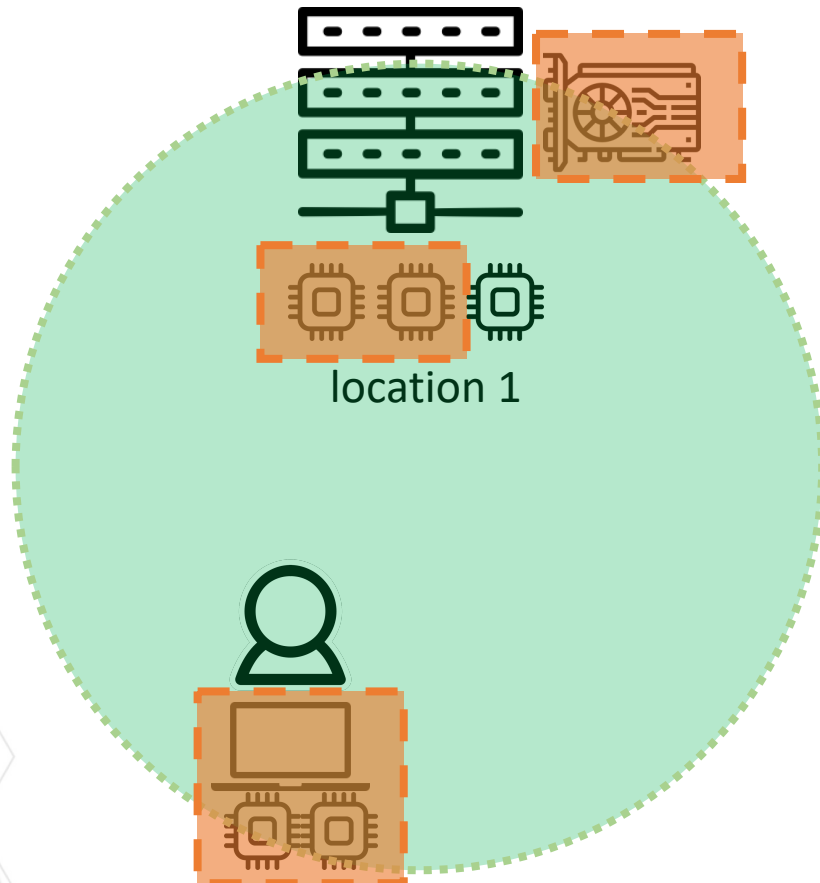


User nears physical proximity of edge system

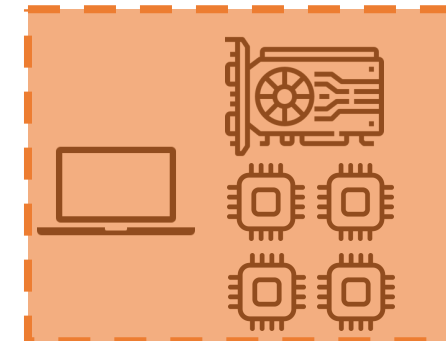


Coalescent Computing

Ephemeral Single-System Image at the Edge






Resources coalesced into one logical system



Coalescent Computing

Ephemeral Single-System Image at the Edge




**User leaves environment,
resources relinquished**


Kyle C. Hale

Coalescent Computing

Ephemeral Single-System Image at the Edge





User approaches another edge system



Coalescent Computing

Ephemeral Single-System Image at the Edge

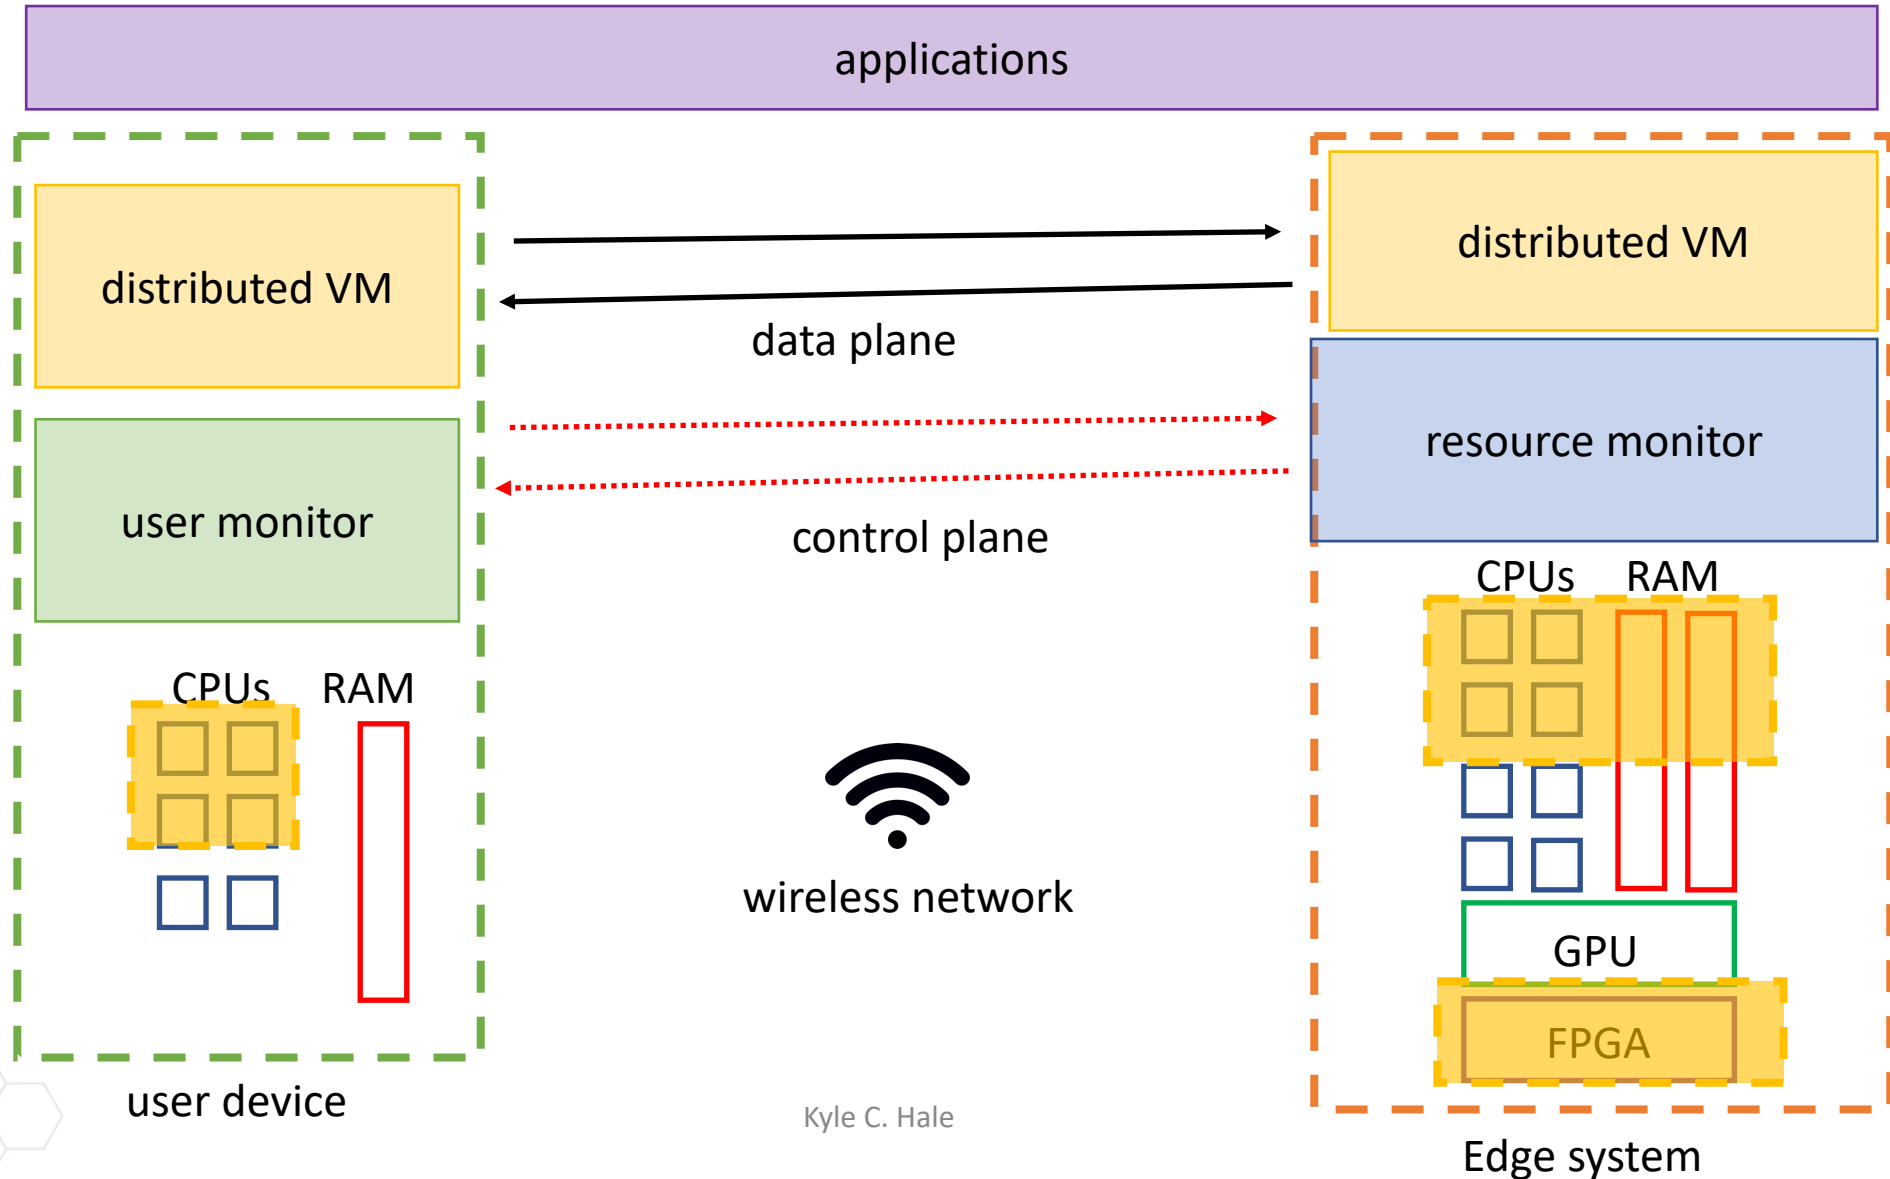


Resources coalesced again, subject to performance, policy constraints



[Hale, Coalescent Computing, APSys '21]

Coalescent System Software



Current Work

- Adapting DSM-based approaches to the edge (e.g., GiantVM)
- Building a prototype co-designed hypervisor/OS for CC
- Applying PL techniques for coalescent offloading policies (collab. with Stefan Muller)



Virtines

Isolating Functions at the Hardware Limit (to appear in EuroSys '22)



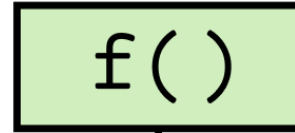
Developer-Friendly, Fast, Function Isolation

- Function isolation
 - Web browsers – sandboxing [CVE-2009-2555, CVE-2009-2935, CVE-2017-2505...]
 - Serverless/FaaS – containers, vms
 - DB *UDF* - high level languages
- Low latency Startup, Short Lived Runtime
 - Spawn and manage many functions w/o significant impact
- Easy programming interface



What might function isolation look like?

- Ephemeral state by default (call stack)
- `return` destroys the context



`f()`

Macro Goals

Function
Granularity





Low Latency
Startup

Easy to integrate



Virtines: Virtual Subroutines

- **Hardware-virtualized** isolated functions
- **Microsecond** level boot times
- Paravirtualization
- General purpose



The *lower bounds* of virtualization



It's not the hardware that's expensive...

- What is the latency of VM creation?
 - HW/SW state

SW VM Allocation is
expensive

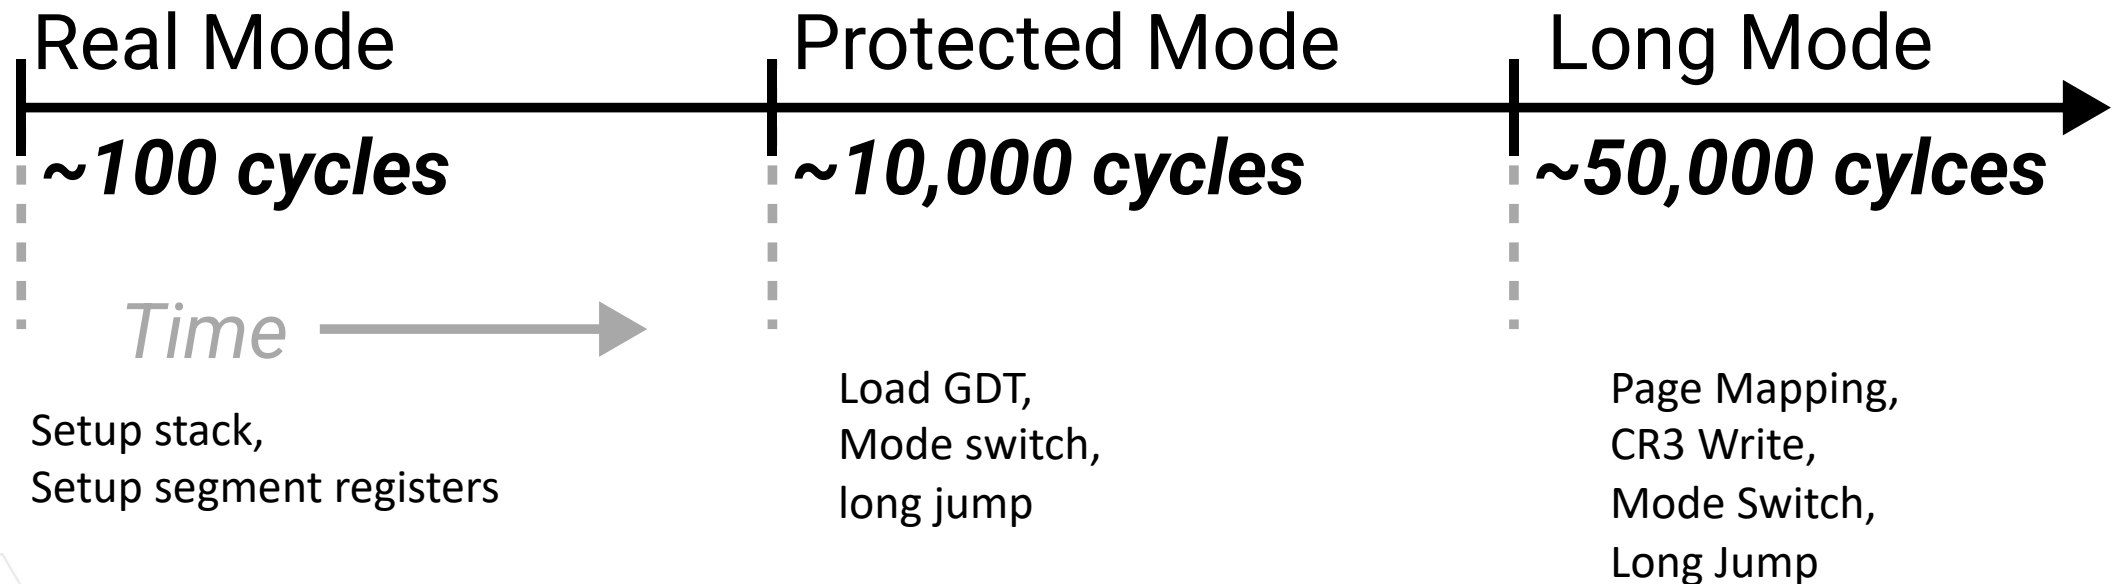
VM Interaction is *cheap*

KVM pthread vmrun



(on AMD EPYC 7281)

Bootstrapping woes

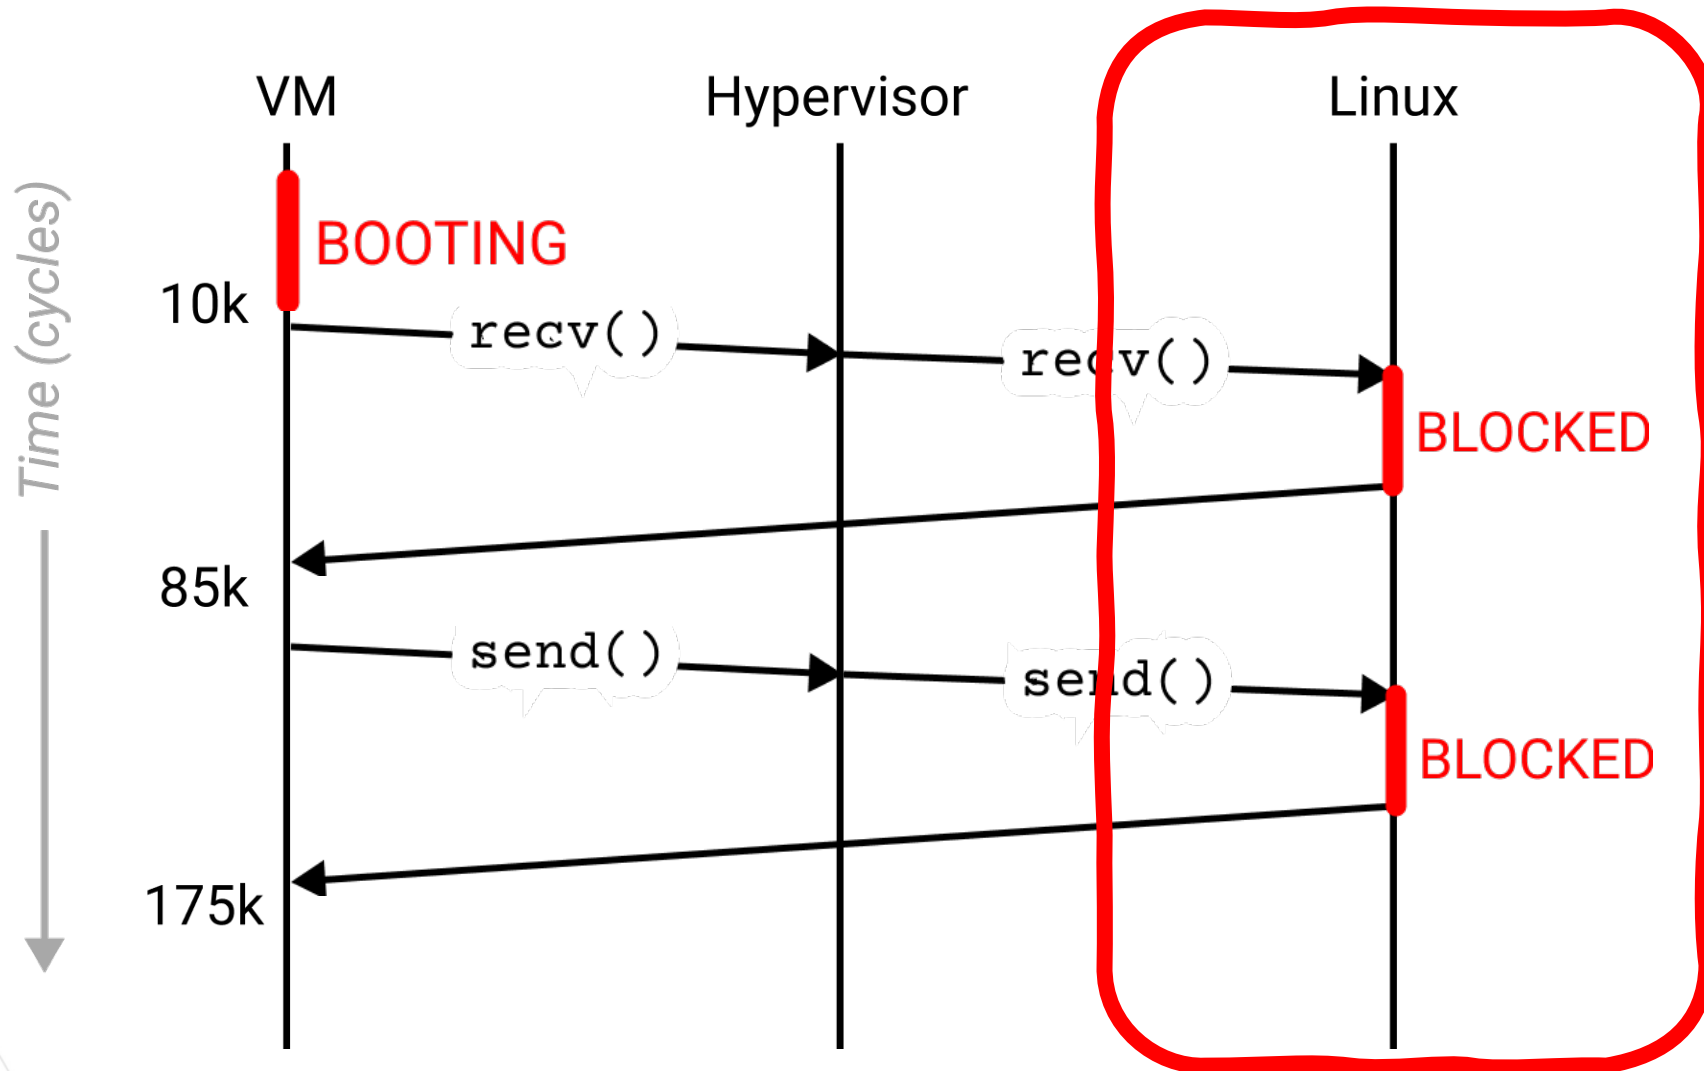


Traditional I/O is *very* expensive

- Try to make a VM feel like real hardware
 - Requires large device drivers
 - Lots of VM Exits for single ops (Expensive!)
- Paravirtualization
 - Codesigned, VM aware of the Hypervisor
 - I/O via ***hypercalls***



HTTP server using hypercalls



Wasp: an implementation of Virtines

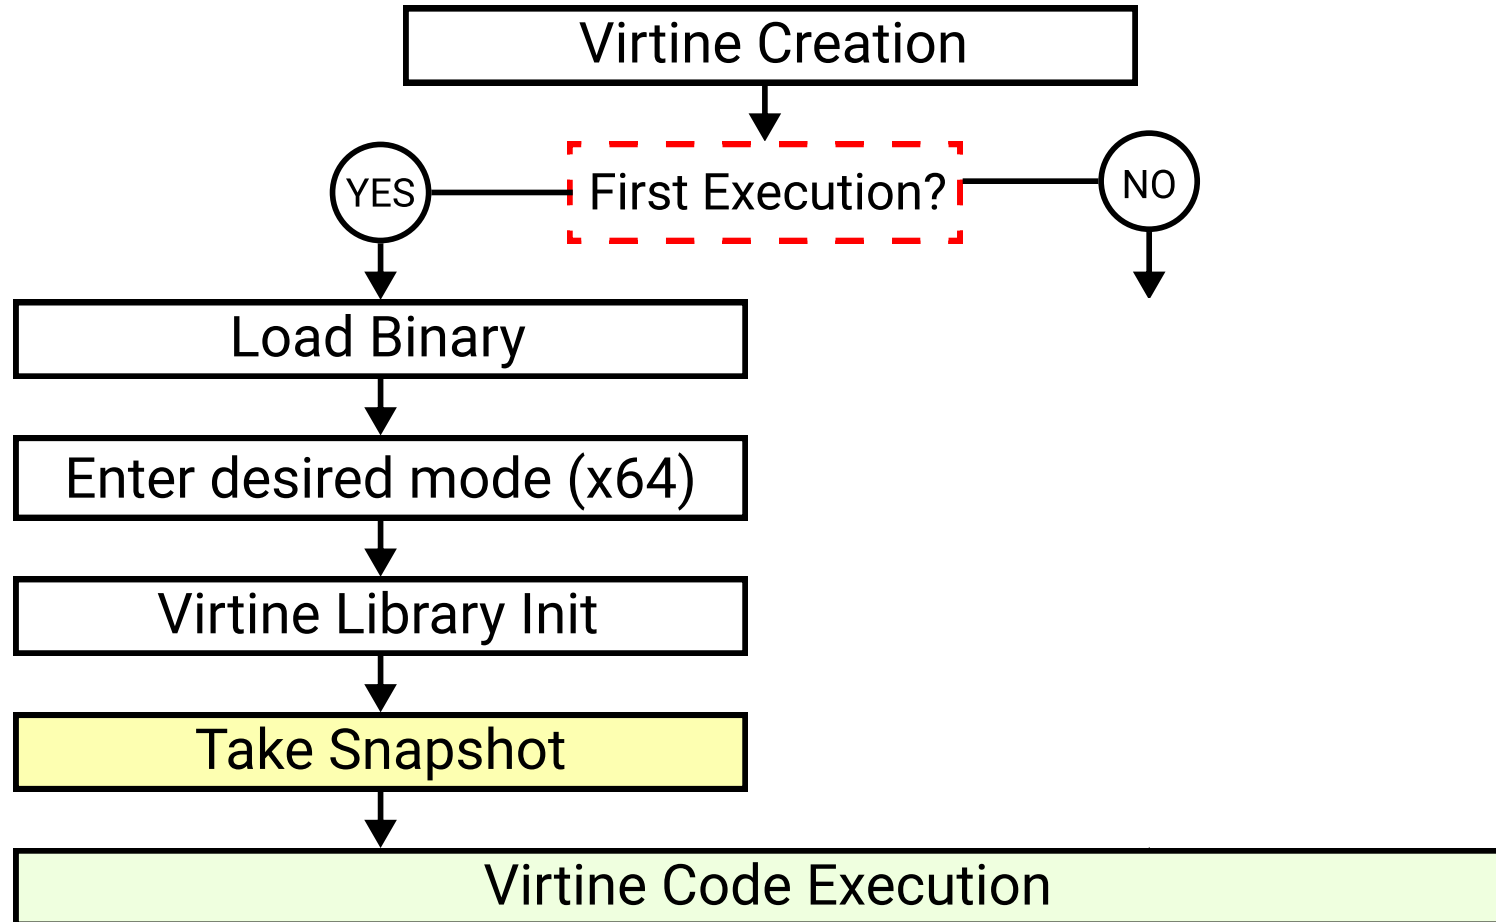


Wasp

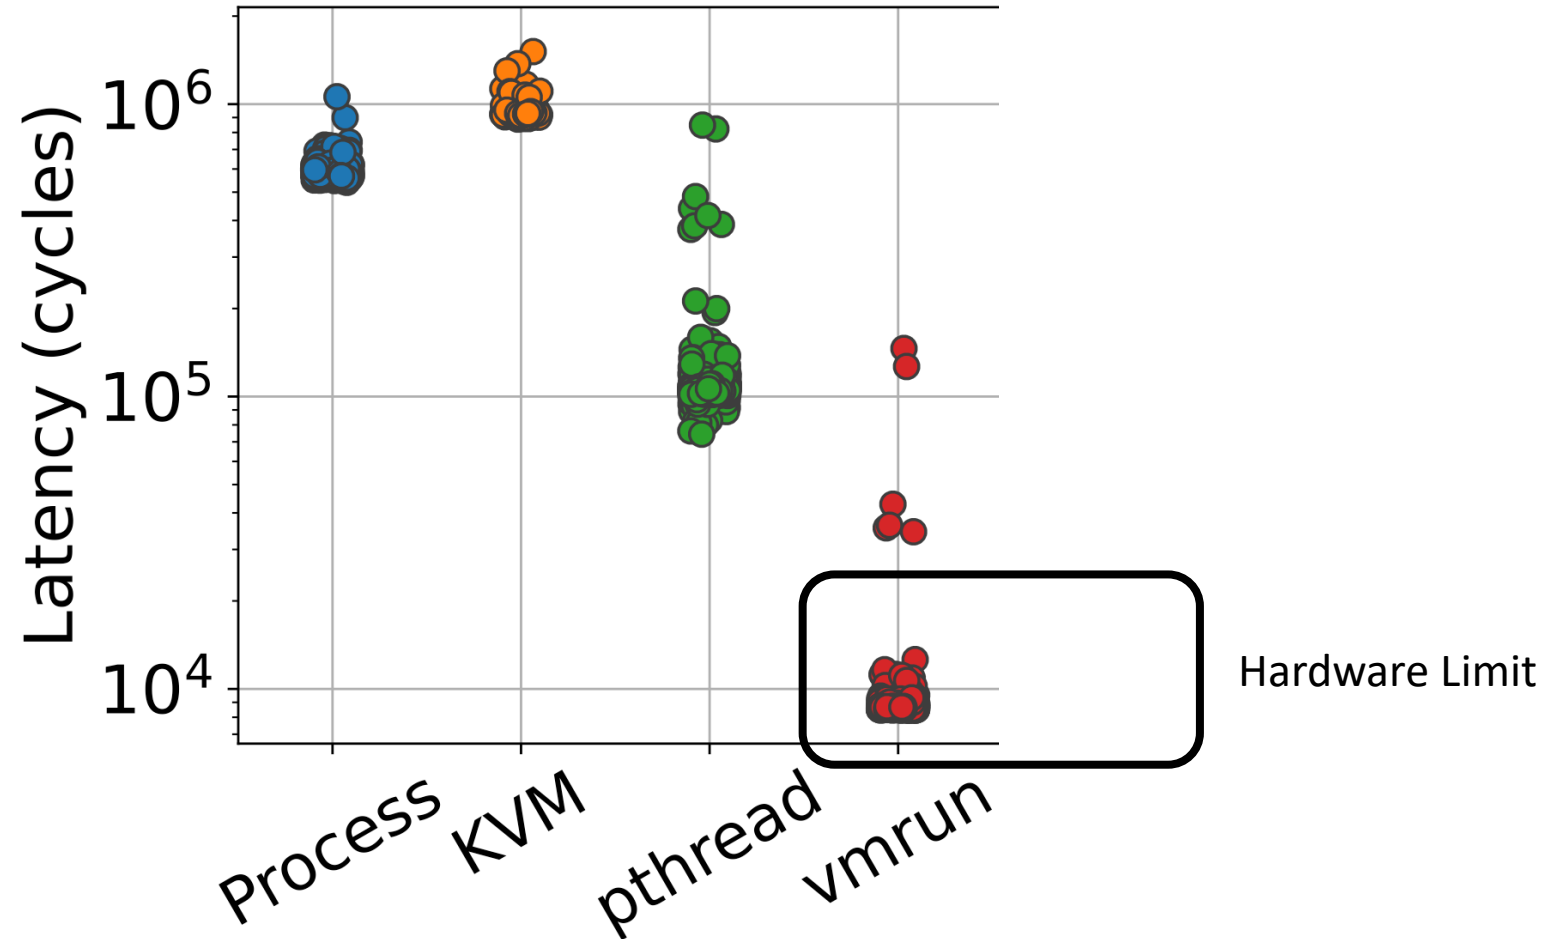
- A micro-hypervisor library
- Abstracts hardware specific interfaces
- Lean
- ***Heavily optimized***



Allocating virtines with Wasp




Wasp is close to the hardware limit



C Language Extension

```
void db_run_udf(db_udf_t *code) {  
    // ...  
}
```



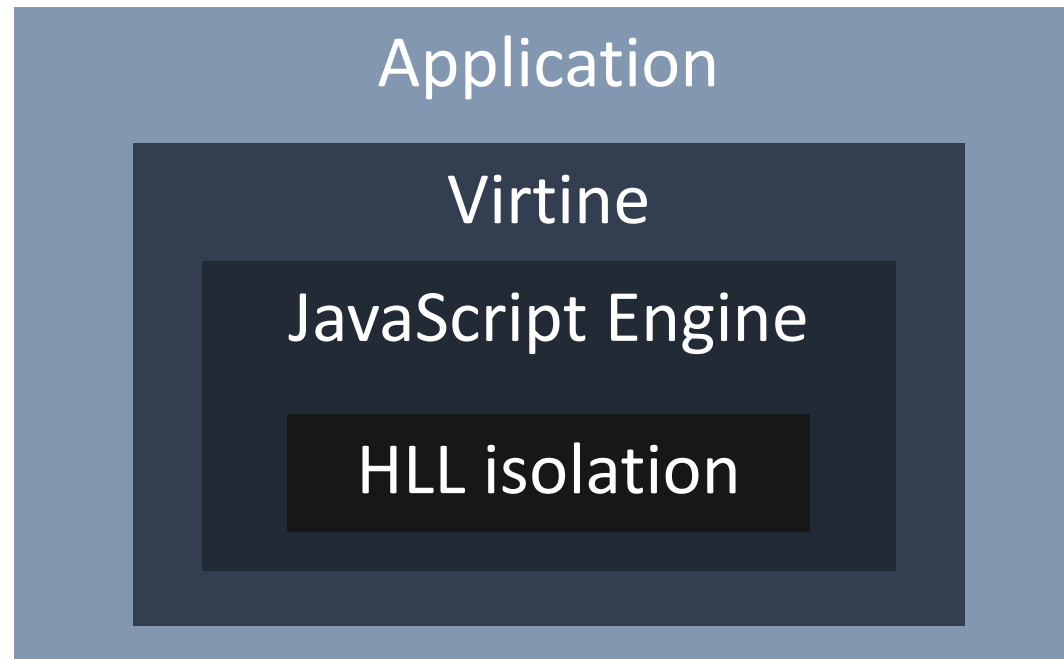
```
virtine void db_run_udf(db_udf_t *code) {  
    // ...  
}
```

- Default-deny access to host services
- Custom LLVM module pass to compile and manage virtines

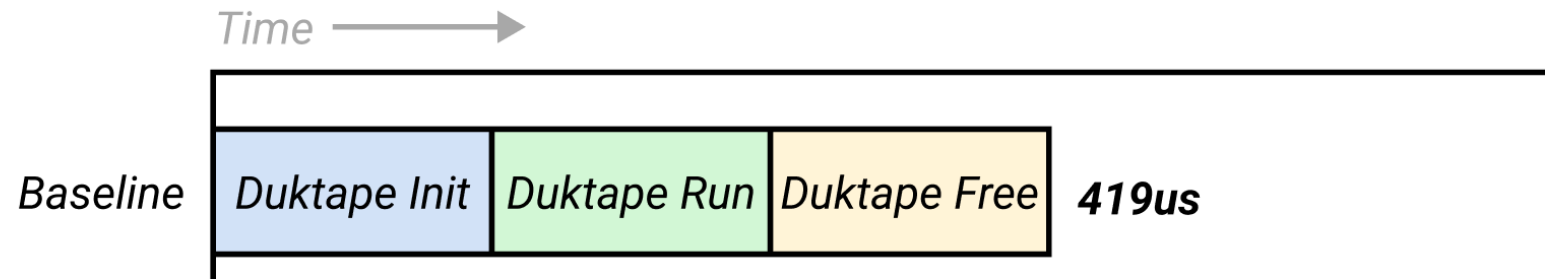
HEXSA

Duktape JavaScript Engine

- **Duktape** JavaScript Engine: <https://duktape.org/>
 - Embeddable, Portable



Do our latency optimizations work?



Effectiveness of language extensions

```
void handl  
    // ...  
}
```

```
virtine v  
    // ...  
}
```

Baseline: **50us**, $\sigma \approx 10\text{us}$

Virtines: **80us**, $\sigma \approx 30\text{us}$

```
(int fd) {
```

How easy is it to integrate?

- ~20 lines of code changed
 - Mostly glue logic
- Significantly slower... But we expect that
- 21kb

```
virtine void do_vpAES_cbc_encrypt(  
    struct virtine_aes_state *state,  
    const AES_KEY key,  
    int encrypt) {  
    // ...  
}
```

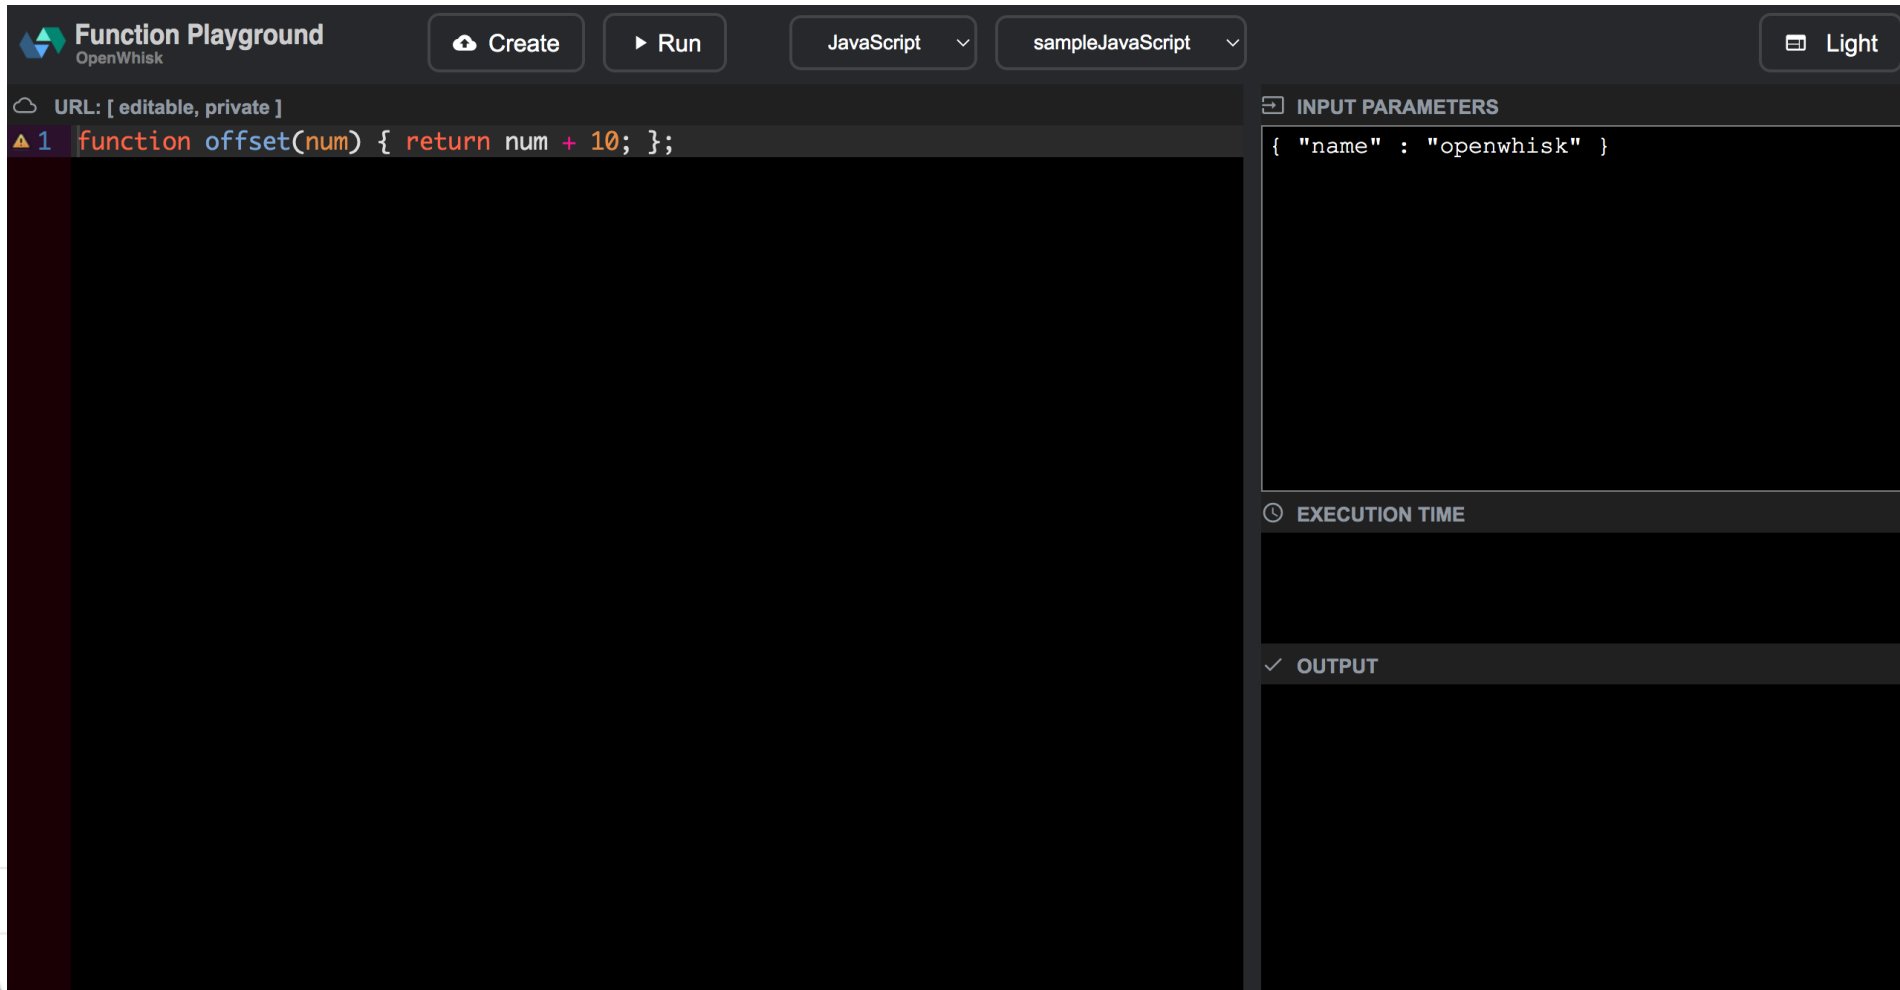


Serverless Vitrines

- OpenWhisk/AWS Lambda are good examples of modern serverless platforms
- Weak isolation between function instances!
- We developed our own using vitrines (based on OW)




Serverless Platform Interface



The screenshot displays the OpenWhisk Function Playground interface. At the top, there is a header with the OpenWhisk logo, the text "Function Playground", and buttons for "Create", "Run", "JavaScript", "sampleJavaScript", and "Light". Below the header, a URL field is labeled "URL: [editable, private]". The main area is a code editor containing a single line of JavaScript code: `1 function offset(num) { return num + 10; };`. To the right of the code editor, there are three panels: "INPUT PARAMETERS" showing a JSON object `{ "name" : "openwhisk" }`, "EXECUTION TIME" which is currently empty, and "OUTPUT" which is also empty.

HEXSA

We can use virtines as drop-in replacement for containers!





Isolating Functions at the Hardware Limit with Virtines

Nicholas C. Wanninger*
ncw@u.northwestern.edu
Northwestern University
Evanston, Illinois, USA

Joshua J. Bowden
jbowden@hawk.iit.edu
Illinois Institute of Technology
Chicago, Illinois, USA

Kirtankumar Shetty
kshetty11@hawk.iit.edu
Illinois Institute of Technology
Chicago, Illinois, USA

Ayush Garg
agarg22@hawk.iit.edu
Illinois Institute of Technology
Chicago, Illinois, USA

Kyle C. Hale
khale@cs.iit.edu
Illinois Institute of Technology
Chicago, Illinois, USA


Abstract

An important class of applications, including programs that leverage third-party libraries, programs that use user-defined functions in databases, and serverless applications, benefit from isolating the execution of untrusted code at the granularity of individual functions or function invocations. However, existing isolation mechanisms were not designed for this use case; rather, they have been adapted to it. We introduce

Keywords: virtines, virtualization, isolation

ACM Reference Format:

Nicholas C. Wanninger, Joshua J. Bowden, Kirtankumar Shetty, Ayush Garg, and Kyle C. Hale. 2022. Isolating Functions at the Hardware Limit with Virtines. In *Seventeenth European Conference on Computer Systems (EuroSys '22)*, April 5–8, 2022, RENNES, France. ACM, New York, NY, USA, 19 pages. <https://doi.org/10.1145/3492224.3510552>



Thanks!

- HExSA Lab: hexsa.halek.co
- Thanks to MD Ali, Conghao Liu, Brian Tauro, Stefan Muller
- Virtines code available: github.com/virtines
- **Thanks to Nick Wanninger (virtines lead author)**
 - Email: ncw@u.northwestern.edu
 - Website: <https://nickw.io>
- Also Josh Bowden, Kirtan Shetty, Ayush Garg

