





Unikernels: Paths to Production & Current Research Trends

Hugo Lefeuvre The University of Manchester

ASPLOS 2022 Unikraft Tutorial, March 1st













~2012: A. Kantee's thesis on Rump kernels

A tad later, Mirage (ASPLOS'13), coined the term "unikernel".

~2012: A. Kantee's thesis on Rump kernels

A tad later, Mirage (ASPLOS'13), coined the term "unikernel".

Now, quite a large body of work

~2012: A. Kantee's thesis on Rump kernels

A tad later, Mirage (ASPLOS'13), coined the term "unikernel".

Now, quite a large body of work

ClickOS (NSDI'24) Mirage (ASPLOS'13)

Specialization / Performance

Unikraft (EuroSys'21)

~2012: A. Kantee's thesis on Rump kernels

A tad later, Mirage (ASPLOS'13), coined the term "unikernel".

Now, quite a large body of work

ClickOS (NSDI'24) Mirage (ASPLOS'13)

Specialization / Performance

Unikraft (EuroSys'21)

Firecracker (NSDI'21) USETL (ApSys)

LightVM (SOSP'17)

Scalability

Solo5 (HotCloud/SoCC)

~2012: A. Kantee's thesis on Rump kernels

A tad later, Mirage (ASPLOS'13), coined the term "unikernel".

Now, quite a large body of work

ClickOS (NSDI'24)		LightVM (SOSP'17)	EbbRT (OSDI'16)
	Mirage (ASPLOS'13)	Firecracker (NSDI'21) USETL (ApSys)	Lupine Linux (EuroSys'20)
Specialization / Performance			
Specialization	n / Performance	Scalability	Compatibility

So, a lot of interest, a lot of nice ideas.

So, a lot of interest, a lot of nice ideas.

And yet, in practice, we don't see unikernels in production (even in the cloud) in 2022.

So, a lot of interest, a lot of nice ideas.

And yet, in practice, we don't see unikernels in production (even in the cloud) in 2022.

Why?

So, a lot of interest, a lot of nice ideas.

And yet, in practice, we don't see unikernels in production (even in the cloud) in 2022.

Why?

Current unikernels are just (mostly academic) research prototypes

Since they are research prototypes, none **managed (or even tried!) to**:

Since they are research prototypes, none managed (or even tried!) to:

- provide good compatibility
- excellent performance
- transparent integration in major deployment workflows

Since they are research prototypes, none managed (or even tried!) to:

- provide good compatibility
- excellent performance
- transparent integration in major deployment workflows
- take security seriously
- meet production-grade testing standards
- good debuggability

Since they are research prototypes, none managed (or even tried!) to:

- provide good compatibility
- excellent performance
- transparent integration in major deployment workflows
- take security seriously
- meet production-grade testing standards
- good debuggability

Now, if we want to see unikernels in production one day, one project must show that it can be done because some of these items are not trivial. **We want this project to be Unikraft.**

Since they are research prototypes, none managed (or even tried!) to:

- provide good compatibility
- excellent performance
- transparent integration in major deployment workflows
- take security seriously
- meet production-grade testing standards
- good debuggability

Now, if we want to see unikernels in production one day, one project must show that it can be done because some of these items are not trivial. **We want this project to be Unikraft.**

Unikraft is not there yet. But it's progressing, and we hope to see it reaching full maturity in the coming year

Since they are research prototypes, none **managed (or even tried!) to**:

- excellent performance
- transparent integration in major —• good debuggability deployment workflows
- meet production-grade testing standards

Now, if we want to see unikernels in production one day, one project must show that it can be done because some of these items are not trivial. We want this project to be Unikraft.

Unikraft is not there yet. But it's progressing, and we hope to see it reaching full maturity in the coming year



Integration and Scalability





• Making Unikraft fit for classical deployment workflows (Kubernetes, etc.)

Integration and Scalability



٠



workflows (Kubernetes, etc.)Making Unikraft just as debuggable as any userland application

Making Unikraft fit for classical deployment

Integration and Scalability





Integration and Scalability



- Making Unikraft fit for classical deployment workflows (Kubernetes, etc.)
- Making Unikraft just as debuggable as any userland application
- Open question: real world high-density: achieving 100s of unikernels per host?



Integration and Scalability

- Making Unikraft fit for classical deployment workflows (Kubernetes, etc.)
- Making Unikraft just as debuggable as any userland application
- Open question: real world high-density: achieving 100s of unikernels per host?
- Matching the security/hardening features of mainstream OSes





Integration and Scalability



- Making Unikraft fit for classical deployment workflows (Kubernetes, etc.)
- Making Unikraft just as debuggable as any userland application
- Open question: real world high-density: achieving 100s of unikernels per host?
- Matching the security/hardening features of mainstream OSes
- Production-grade testing and fuzzing of Unikraft



Integration and Scalability



- Making Unikraft fit for classical deployment workflows (Kubernetes, etc.)
- Making Unikraft just as debuggable as any userland application
- Open question: real world high-density: achieving 100s of unikernels per host?
- Matching the security/hardening features of mainstream OSes
- Production-grade testing and fuzzing of Unikraft
- Specialization for the masses: automatic reasoning about Unikraft configurations



Integration and Scalability



- Making Unikraft fit for classical deployment workflows (Kubernetes, etc.)
- Making Unikraft just as debuggable as any userland application
- Open question: real world high-density: achieving 100s of unikernels per host?
- Matching the security/hardening features of mainstream OSes
- Production-grade testing and fuzzing of Unikraft
- Specialization for the masses: automatic reasoning about Unikraft configurations
- Beyond the single trust domain: compartmentalizing Unikraft?

Integration & Scalability

Integration and Scalability

Why no unikernels in production?

[-] 🔺 😔 yonkeltron 13 days ago | link

It's always been a mystery to me why Unikernels haven't caught on more. Especially with earlier toolkiits like <u>UniK</u> and continuing work such as <u>OSv</u>. Does anyone have production experience or a tale of why they *didn't* pick Unikernels?

[-] 🔺 💥 david_chisnall 13 days ago | link

The problems are not technical, they're economic. If I want to deploy a unikernel in the cloud, I am deploying as IaaS (i.e. a VM) where the unit of accounting is typically pairs of vCPUs and gigabytes of RAM, on an hourly basis. If I have the kind of problem where a unikernel would be a good solution, then I can deploy it as a FaaS system and be billed per CPU second and per RAM MiB second. None of the cloud providers (yet?) have a way of deploying unikernels with FaaS-like pricing and so if you make something small and efficient as a unikernel then it will have a load of unused CPU time and RAM that you're still being charged for. Unikernels only make economic sense if you're deploying your own datacenter.

No good integration.

https://lobste.rs/s/cyyx7a/unikraft_fast_secure_open_source (NOT an official Microsoft comment)

Integration and Scalability

Why no unikernels in production?



https://lobste.rs/s/cyyx7a/unikraft_fast_secure_open_source (NOT an official Microsoft comment) https://thenewstack.io/good-luck-debugging-unikernels-joyents-chief-technology-says/

People need integration of unikernels into orchestration frameworks to truly leverage their benefits.

People need integration of unikernels into orchestration frameworks to truly leverage their benefits.

- Dynamically and quickly provision new services
- Schedule/Reschedule services based on workload (etc.)



People need integration of unikernels into orchestration frameworks to truly leverage their benefits.

- Dynamically and quickly provision new services
- Schedule/Reschedule services based on workload (etc.)
 No need to reinvent the wheel: make unikernels fit in these frameworks



People need integration of unikernels into orchestration frameworks to truly leverage their benefits.

- Dynamically and quickly provision new services
- Schedule/Reschedule services based on workload (etc.)
 No need to reinvent the wheel: make unikernels fit in these frameworks

We are almost there.

- Integration of unikernels in Kubernetes infrastucture
 - OCI-compliant unikernel runtime interface



People need integration of unikernels into orchestration frameworks to truly leverage their benefits.

- Dynamically and quickly provision new services
- Schedule/Reschedule services based on workload (etc.)
 No need to reinvent the wheel: make unikernels fit in these frameworks

We are almost there.

- Integration of unikernels in Kubernetes infrastucture
 - OCI-compliant unikernel runtime interface
- "A Unikernel in OCI Clothing": make unikernels look and feel like containers

A. Jung @ CNCF'21 https://www.youtube.com/watch?v=cV-xawN9_cg



People need integration of unikernels into orchestration frameworks to truly leverage their benefits.

- Dynamically and quickly provision new services
- Schedule/Reschedule services based on workload (etc.)
 No need to reinvent the wheel: make unikernels fit in these frameworks

We are almost there.

- Integration of unikernels in Kubernetes infrastucture
 - OCI-compliant unikernel runtime interface
- "A Unikernel in OCI Clothing": make unikernels look and feel like containers
- More progress to make on the FaaS side?

A. Jung @ CNCF'21 <u>https://www.youtube.com/watch?v=cV-xawN9_cg</u>



In the Works: Debugging



Vast engineering effort towards seamless introspection and debugging

S. Kuenzer, M. Rittinghaus @ FOSDEM'22 https://fosdem.org/2022/schedule/event/skuenzer/

In the Works: Debugging



Vast engineering effort towards seamless introspection and debugging

- Production-grade monitoring with Prometheus:
 - Monitor unikernels like any general-purpose VM
 - Setup alarms when values pass thresholds, etc.



S. Kuenzer, M. Rittinghaus @ FOSDEM'22 https://fosdem.org/2022/schedule/event/skuenzer/

In the Works: Debugging



Vast engineering effort towards seamless introspection and debugging

- Production-grade monitoring with Prometheus:
 - Monitor unikernels like any general-purpose VM
 - Setup alarms when values pass thresholds, etc.
- OS-level native GDB debugger support:
 - Debug unikernels like any userland application
 - Support for threads, OS-specific constructs (asserts, kernel crashes, etc.)
- Uniform debugging experience all platforms



S. Kuenzer, M. Rittinghaus @ FOSDEM'22 https://fosdem.org/2022/schedule/event/skuenzer/
Vast engineering effort towards seamless introspection and debugging

- Production-grade monitoring with Prometheus:
 - Monitor unikernels like any general-purpose VM
 - Setup alarms when values pass thresholds, etc.
- OS-level native GDB debugger support:
 - Debug unikernels like any userland application
 - Support for threads, OS-specific constructs (asserts, kernel crashes, etc.)
- Uniform debugging experience all platforms





In the Works: Debugging

As unikernels are evolving towards production, more challenges are likely to arise

We (as a community) claimed **massive packing** of such VMs on a single host

As unikernels are evolving towards production, more challenges are likely to arise

We (as a community) claimed massive packing of such VMs on a single host



As unikernels are evolving towards production, more challenges are likely to arise

We (as a community) claimed **massive packing** of such VMs on a single host

... but are we ready to see 50 Nginx instances on a single host? Let alone 100s?



As unikernels are evolving towards production, more challenges are likely to arise

We (as a community) claimed **massive packing** of such VMs on a single host

... but are we ready to see 50 Nginx instances on a single host? Let alone 100s?

Pretty much all unikernel papers evaluate systems with 1 CPU = 1vCPU static pinning...



As unikernels are evolving towards production, more challenges are likely to arise

We (as a community) claimed **massive packing** of such VMs on a single host

... but are we ready to see 50 Nginx instances on a single host? Let alone 100s?

Pretty much all unikernel papers evaluate systems with 1 CPU = 1vCPU static pinning...

With such density: how do things look on the networking side? Have hypervisors really been thought for this kind of usage?



Security & Stability

Security and Stability

Why no unikernels in production?

11 Conclusion

Because of security.



Much to the contrary of grandiose security claims often made by unikernel developers, the evidence thus far indicates that unikernels are decidedly *not* secure. [Bue] Having examined two major unikernels, Rumprun and IncludeOS, a worrying trend is already apparent: unikernels often lack even the most basic security features, especially with regard to memory corruption. ASLR, consistent W^X policy, and stack, heap, and standard library hardening are generally either missing, improperly implemented, or intentionally disabled. This would be bad enough in a full, general-purpose operating system, but it is made even worse in unikernels, where application and kernel code run together and share an address space. An attacker who gains code execution in the application can immediately go on to invoke kernel-level functionality, make hypercalls, perform raw packet I/O, and so on. This makes unikernels a particular liability when running alongside other types of hosts, as they can be used as pivot points from which to attack their neighbors with even more potency than would be possible on a full-OS VM or container (at least without privilege escalation).

Given how low the bar has been set, there are numerous ways in which the currently abysmal state of unikernel security could improve. Aside from the protections we tested for – i.e. those typically found in modern, full-featured operating systems – there are several hypervisor-specific features that can be taken advantage of in order to improve unikernel security. For instance, many privileged operations, e.g. page table management, packet I/O, etc. can be performed via requests to the hypervisor rather than directly by the guest itself through emulated devices; such functionality is akin to syscalls or ioctls in a full OS.

Nonetheless, as it stands, unikemels remain an unsuitable and unappealing choice for production use, and will likely remain so until their security measures are at least brought in line with those of modern, full-featured operating systems.

https://www.nccgroup.com/globalassets/our-research/us/whitepapers/2019/ncc_group-assessing_unikernel_security.pdf

Security and Stability

11 Conclusion

Because of security.



Much to the contrary of grandiose security claims often made by unikernel developers, the evidence thus far indicates that unikernels are decidedly *not* secure. [Bue] Having examined two major unikernels, Rumprun and IncludeOS, a worrying trend is already apparent: unikernels often lack even the most basic security features, especially with regard to memory corruption. ASLR, consistent W^X policy, and stack, heap, and standard library hardening are generally either missing, improperly implemented, or intentionally disabled. This would be bad enough in a full, general-purpose operating system, but it is made even worse in uniker-



together and share an address space. An attacker who gains iately go on to invoke kernel-level functionality, make hypermakes unikernels a particular liability when running alongside rot points from which to attack their neighbors with even more /M or container (at least without privilege escalation).

are numerous ways in which the currently abysmal state of 1 the protections we tested for – i.e. those typically found in ere are several hypervisor-specific features that can be taken ecurity. For instance, many privileged operations, e.g. page erformed via requests to the hypervisor rather than directly by

the guest itsen through emulated devices; such functionality is akin to syscalls or loctls in a full OS.



Why no unikernels in production?

Nonetheless, as it stands, unikemels remain an unsuitable and unappealing choice for production use, and will likely remain so until their security measures are at least brought in line with those of modern, full-featured operating systems.

https://www.nccgroup.com/globalassets/our-research/us/whitepapers/2019/ncc_group-assessing_unikernel_security.pdf https://thenewstack.io/unikernels-will-create-security-problems-solve/

Write or Execute



Write or Execute



Write or Execute



Write or Execute

... and stack protection, KASan, etc.





Write or Execute

... and stack protection, KASan, etc.

🖟 unik	craft / unikraft (Public)			
Code ⊙ Issues 93 1: Pull requests 51 ♀ Discussions ⊙ Actions ⊞ Projects 1 ⑦ Security Pointer authentication				
Virtual Memory API (x86_64,		La unikraft / unikraft Public	ASLR	
11 Ope	marcrittinghaus wants to merge 11 commits into un:	Code ① Issues 93 1 Pull requests 51	f 📮 unikraft / unikraft (Public)	Q Notification
ୟ ୦୦	onversation 7 - Commits 11 🖓 Checks	arch/arm/arm64: Introduce pointer authenticat	Code O Issues 93 1 Pull requests 51 O Discussions O Actions E Projects 1	③ Security
	marcrittinghaus commented on Nov 22, 2021	Copen michpappas wants to merge 1 commit into unikraft:staging from michpappas:arm64_	build: Option to compile as PIE #239	
	 Read the contribution guidelines regarding submittin Tested your changes against relevant architectures and the submitting of the	Q Conversation 7 - Commits 1 E→ Checks 0 E→ Files changed 5	13 Open danield20 wants to merge 1 commit into unikraft:staging from danield20:ddinca/build-as-PIE	
	 Ran the checkpatch.pl on your commit series before Updated relevant documentation. 	michpappas commented on Dec 12, 2021	Q Conversation 7 ↔ Commits 1 🖓 Checks 0 🗄 Files changed 3	
	Base target	Prerequisite checklist	danield20 commented on Jun 25, 2021	(Member) ···
	 Architecture(s): x86_64 Platform(s): kvm Application(s): N/A 	 Read the contribution guidelines regarding submitting new changes to the project; Tested your changes against relevant architectures and platforms; Ran the checkpatch.pl on your commit series before opening this PR; Undated relevant documentation 	This patch adds the option to compile the unikernel as a position-independent executable so we can have ASLR. If the unikernel is compiled as PIE, it cannot run on it's own. A bootloader will be needed that will come in a future PR.	
	Additional configuration The PR introduces a configuration option under Platfor memory.	Base target • Architecture(s): arm64	Signed-off-by: Daniel Dinca dincadaniel97@gmail.com	
	CONFIG_PAGING=y	Platform(s): kvm		
	This will automatically disable all platforms currently not ukfallocbuddy , the physical memory allocator libraries	Application(s): All		
		Additional configuration		
		This PR introduces CONFIG_ARM64_FEAT_PAUTH to enable Pointer Authentication support.		

To go mainstream, unikernels need not only hardening that's in line with mainstream OSes, but also **production-grade testing**

To go mainstream, unikernels need not only hardening that's in line with mainstream OSes, but also **production-grade testing**

Production-grade testing =

- Continuous testing (CI/CD, test suite)
- Destructive testing (Fuzzing)

To go mainstream, unikernels need not only hardening that's in line with mainstream OSes, but also **production-grade testing**

Production-grade testing =

- Continuous testing (CI/CD, test suite)
- Destructive testing (Fuzzing)

Significant efforts on continuous testing:



To go mainstream, unikernels need not only hardening that's in line with mainstream OSes, but also **production-grade testing**

Production-grade testing =

- Continuous testing (CI/CD, test suite)
- Destructive testing (Fuzzing)

Significant efforts on continuous testing:

- CI/CD pipeline tests patches systematically (Concourse)
- Application-level tests but also kernel unit-tests (uktest)

A. Jung @ FOSDEM'22

https://fosdem.org/2022/schedule/event/massive_unikernel_matrices_with_unikraft_concourse_and_more/



To go mainstream, unikernels need not only hardening that's in line with mainstream OSes, but also **production-grade testing**

Ongoing efforts on fuzzing as well

To go mainstream, unikernels need not only hardening that's in line with mainstream OSes, but also **production-grade testing**

Ongoing efforts on fuzzing as well

• Not entirely trivial, as most OS fuzzers are tailored for Linux

To go mainstream, unikernels need not only hardening that's in line with mainstream OSes, but also **production-grade testing**

Ongoing efforts on fuzzing as well

• Not entirely trivial, as most OS fuzzers are tailored for Linux

Not "just" a matter of porting Syzkaller to Unikraft

To go mainstream, unikernels need not only hardening that's in line with mainstream OSes, but also **production-grade testing**

Ongoing efforts on fuzzing as well

- Not entirely trivial, as most OS fuzzers are tailored for Linux
 - Coverage measurement (no Kcov, porting to gcov not without changes)

Not "just" a matter of porting Syzkaller to Unikraft

• Not every system call is fully implemented

To go mainstream, unikernels need not only hardening that's in line with mainstream OSes, but also **production-grade testing**

Ongoing efforts on fuzzing as well

- Not entirely trivial, as most OS fuzzers are tailored for Linux
 - Coverage measurement (no Kcov, porting to gcov not without changes)

Not "just" a matter of porting Syzkaller to Unikraft

- Not every system call is fully implemented
- How does unikernel fuzzing impact the architecture of fuzzers?
- How to design a fuzzer that's ready to "plug and play" in any POSIX OS?

Another end of the "testing" topic: how do you determine **how good a configuration really is**?

Another end of the "testing" topic: how do you determine **how good a configuration really is**?

Unikernels pitch specialization:

- Your best Nginx configuration is not your best SQLite configuration
- ...and probably not your best Redis configuration either $oldsymbol{arphi}$

Another end of the "testing" topic: how do you determine **how good a configuration really is**?

Unikernels pitch specialization:

- Your best Nginx configuration is not your best SQLite configuration
- ...and probably not your best Redis configuration either $oldsymbol{arphi}$

The number of possible configurations: astronomical scale

Another end of the "testing" topic: how do you determine **how good a configuration really is**?

Unikernels pitch specialization:

- Your best Nginx configuration is not your best SQLite configuration
- ...and probably not your best Redis configuration either $oldsymbol{arphi}$

The number of possible configurations: astronomical scale

- Small subset of configuration options (Nginx) ~ 10¹³
- How do you explore this? Can you use optimization algorithms? ML?

Traditional understanding of unikernels:

Unikernel = one single trust domain (kernel + application)

(Used to) make sense.

Traditional understanding of unikernels:

Unikernel = one single trust domain (kernel + application)

(Used to) make sense.

Certain applications are large, with heterogeneous components: trust, safety, properties, requirements...

Traditional understanding of unikernels:

Unikernel = one single trust domain (kernel + application)

(Used to) make sense.

Certain applications are large, with heterogeneous components: trust, safety, properties, requirements...

And at the same time we see (re-)appearing a lot of lightweight isolation mechanisms (protection keys, HW capabilities, SFI, etc.)

There is an opportunity to use these mechanisms to make unikernels even safer without yielding their benefits!

There is an opportunity to use these mechanisms to make unikernels even safer without yielding their benefits!



This is what initially motivated our work FlexOS: can we reconcile unikernels/libOSes with isolation to obtain a new OS model that offers not only specialization **towards performance, but also towards safety**?

H. Lefeuvre et al. @ ASPLOS'22, come to our talk Thursday morning! (also @ FOSDEM'22 <u>https://fosdem.org/2022/schedule/event/tee_flexos/</u>)

There is an opportunity to use these mechanisms to make unikernels even safer without yielding their benefits!



This is what initially motivated our work FlexOS: can we reconcile unikernels/libOSes with isolation to obtain a new OS model that offers not only specialization **towards performance**, but also towards safety?

Other groups explored this direction: CubicleOS (also ASPLOS, 2021). Explore intra-unikernel isolation with Intel MPK.

H. Lefeuvre et al. @ ASPLOS'22, come to our talk Thursday morning! (also @ FOSDEM'22 <u>https://fosdem.org/2022/schedule/event/tee_flexos/</u>)
How about Compatibility?

Historical unikernels: hand-written for an application (ClickOS).

Historical unikernels: hand-written for an application (ClickOS)

Historical unikernels: hand-written for an application (ClickOS)

But we're not in the 2010s anymore, Unikraft aims at Linux/POSIX compliance.

Historical unikernels: hand-written for an application (Clickos, patibility)

But we're not in the 2010s anymore, Unikraft aims at Linux/POSIX compliance.

Growing number of supported system calls: now 170+ As a point of comparison, Graphene (Gramine) also supports about ~170

Historical unikernels: hand-written for an application (ClickOS)

But we're not in the 2010s anymore, Unikraft aims at Linux/POSIX compliance.

Growing number of supported system calls: now 170+ As a point of comparison, Graphene (Gramine) also supports about ~170

To be clear: Unikraft is neither fully Linux compatible, nor fully POSIX compliant!

- The good old fork() problem
- Not all system calls are fully implemented

But **do you really need to be fully compatible** to be useful?

But **do you really need to be fully compatible** to be useful?

Many applications that use unsupported features tend to be bad candidates for unikernelization anyways (system admin tools, heavily multiprocess apps that cannot use threads).

But **do you really need to be fully compatible** to be useful?

Many applications that use unsupported features tend to be bad candidates for unikernelization anyways (system admin tools, heavily multiprocess apps that cannot use threads).

For the rest, partial compatibility is just fine if porting is a reasonable task.

And in 2022, it is.

H. Lefeuvre et al. @ USENIX ;login <u>https://www.usenix.org/publications/loginonline/unikraft-and-coming-age-unikernels</u>

But **do you really need to be fully compatible** to be useful?

Many applications that use unsupported features tend to be bad candidates for unikernelization anyways (system admin tools, heavily multiprocess apps that cannot use threads).

For the rest, partial compatibility is just fine if porting is a reasonable task.

And in 2022, it is.

Given the benefits of Unikraft, a week of porting is a minor annoyance. All you need is a good application test-suite (but you have one, right? 🙂)

H. Lefeuvre et al. @ USENIX ;login <u>https://www.usenix.org/publications/loginonline/unikraft-and-coming-age-unikernels</u>

Conclusion

In the last decade we've seen a lot of unikernels come and go, but **none that managed to**:

In the last decade we've seen a lot of unikernels come and go, but **none that managed to**:

- provide good compatibility
- excellent performance
- transparent integration in major deployment workflows

In the last decade we've seen a lot of unikernels come and go, but **none that managed to**:

- provide good compatibility
- excellent performance
- transparent integration in major deployment workflows
- take security seriously
- meet production-grade testing standards
- good debuggability

In the last decade we've seen a lot of unikernels come and go, but **none that managed to**:

- provide good compatibility
- excellent performance
- transparent integration in major deployment workflows
- take security seriously
- meet production-grade testing standards
- good debuggability

And that's fine, because they were just **research projects**

In the last decade we've seen a lot of unikernels come and go, but **none that managed to**:

- provide good compatibility
- excellent performance
- transparent integration in major deployment workflows
- take security seriously
- meet production-grade testing standards
- good debuggability

And that's fine, because they were just **research projects**

Now, if we want to see unikernels in production one day, one project must manage it. We want this project to be Unikraft.

In the last decade we've seen a lot of unikernels come and go, but **none that managed to**:

- provide good compatibility
- excellent performance
- transparent integration in major deployment workflows
- take security seriously
- meet production-grade testing standards
- good debuggability

And that's fine, because they were just **research projects**

Now, if we want to see unikernels in production one day, one project must manage it. We want this project to be Unikraft.

Unikraft is not there yet. But it's progressing, and we hope to see it reaching full maturity in the coming year

Over time, doing all this engineering proved fruitful on the research side

Over time, doing all this engineering proved fruitful on the research side

Unikraft turns out to be an **excellent substrate for top-tier publications**

It's open, small, clean, modular, fast to experiment with, such that undergrads and grad students quickly get to understand it

Over time, doing all this engineering proved fruitful on the research side

Unikraft turns out to be an **excellent substrate for top-tier publications**

It's open, small, clean, modular, fast to experiment with, such that undergrads and grad students quickly get to understand it

Probably one of the good examples where starting clean-slate pays out in the long run

Over time, doing all this engineering proved fruitful on the research side

Unikraft turns out to be an **excellent substrate for top-tier publications**

It's open, small, clean, modular, fast to experiment with, such that undergrads and grad students quickly get to understand it

Probably one of the good examples where starting clean-slate pays out in the long run

What will the broader systems community build with Unikraft?







Pushing Unikernels to Production!

Unikraft Community: <u>https://unikraft.org/</u> Unikraft Cloud: <u>https://unikraft.io/</u> Code: <u>https://github.com/unikraft</u>