Build, Ship, Run Unikernels

Justin Cormack

Justin Cormack

Cambridge based developer at Docker @justincormack

Co-author of Docker in the Trenches: Successful Production Deployment

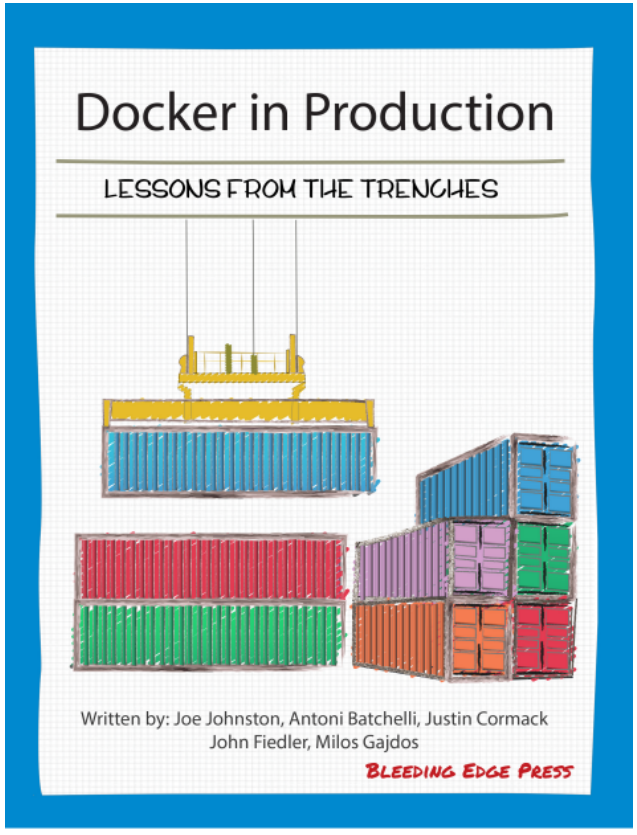

containers


HANJIN
HANJIN
HANJIN
COSCO
COSCO
한진 타이페이
HANJIN TAIPEI
HAMBURG

- “Linux containers are an operating-system-level virtualization environment for running multiple isolated Linux systems on a single Linux control host”
- “Building on top of facilities provided by the Linux kernel, a Docker container, unlike a virtual machine, does not require or include a separate operating system. Instead, it relies on the kernel's functionality and uses resource isolation and separate namespaces to isolate the application's view of the operating system.”

General Purpose Container

20'

Construction	Inside Dimensions			Door Opening		Weights			Capacity	Hapag-Lloyd Serial Number	Foot-note
	Length	Width	Height	Width	Height	Max. Gross	Tare	Max. Payload			
8'6" high	mm ft	mm ft	mm ft	mm ft	mm ft	kg lbs	kg lbs	kg lbs	m^3 cu.ft		
Steel container with corrugated walls and wooden floor	5 895 19'4 1/8"	2 350 7'8 1/2"	2 392 7'10 1/8"	2 340 7'8 1/8"	2 292 7'6 1/4"	30 480 67 200	2 250 4 960	28 230 62 240	33,2 1172	CPSU 100 000 – 108 362 CPSU 108 470 – 182 099 IVLU 955 076 – 957 000 HLXU 200 000 – 212 799 HLXU 212 800 – 239 799 HLXU 300 000 – 310 099	 1) 2) 3) 1) 2) 3)
	5 900 9'4 1/4"	2 352 7'8 5/8"	2 395 7'10 1/4"	2 340 7'8 1/8"	2 292 7'6 1/4"	32 500 71 650	2 370 5 220	30 130 66 430	33,2 1172	HLXU 310 100 – 340 699 HLXU 340 800 – 354 699	1) 2) 3) 1) 2) 3)
and steel floor	5 895 19'4 1/8"	2 350 7'8 1/2"	2 392 7'10 1/8"	2 340 7'8 1/8"	2 292 7'6 1/4"	32 500 71 650	2 570 5 670	29 930 65 980	33,2 1172	HLXU 340 700 – 340 799 CPSU 108 363 – 108 469	1) 2) 3) 5)

Construction	Inside Dimensions				Weights			Capacity	Hapag-Lloyd Serial Number	Foot-note
	Length	Width	Height		Max. Gross	Tare	Max. Payload			
			Middle	Side						
8'6" high ISO Size Type Code: 22U6	mm ft	mm ft	mm ft	mm ft	kg lbs	kg lbs	kg lbs	m^3 cu.ft		
	5 886 19'3 3/4"	2 342 7'8 1/8"	2 388 7'10"	2 313 7'7"	30 480 67 200	2 700 5 950	27 780 61 250	32,8 1160	FANU 260 200 – 261 799	4)
Steel container with corrugated walls, wooden floor and removable steel roof	5 886 19'3 3/4"	2 342 7'8 1/8"	2 388 7'10"	2 313 7'7"	30 480 67 200	2 700 5 950	27 780 61 250	32,8 1160	HLXU 365 000 – 365 649	4)
	5 859 19'3 3/4"	2 350 7'8 1/8"	2 390 7'9 1/2"	2 309 7'7 3/4"	32 500 71 650	2 850 6 280	29 650 65 370	32,1 1132	HLXU 365 650 – 365 949	4)

Remarks:
1) 10 lashing rings on each top longitudinal rail; particularly suitable for the transport of hanging garments racks.
2) Provided with passive vents.
3) Provided with extra lashing rings/bars for the transport of liner bags in the corner posts adjacent to the corner castings.
4) For special information please see 20' Hard Top Container.
5) Max Gross 30 480 kg


WUT


COSCO
HYUNDAI
HYUNDAI
HYUNDAI

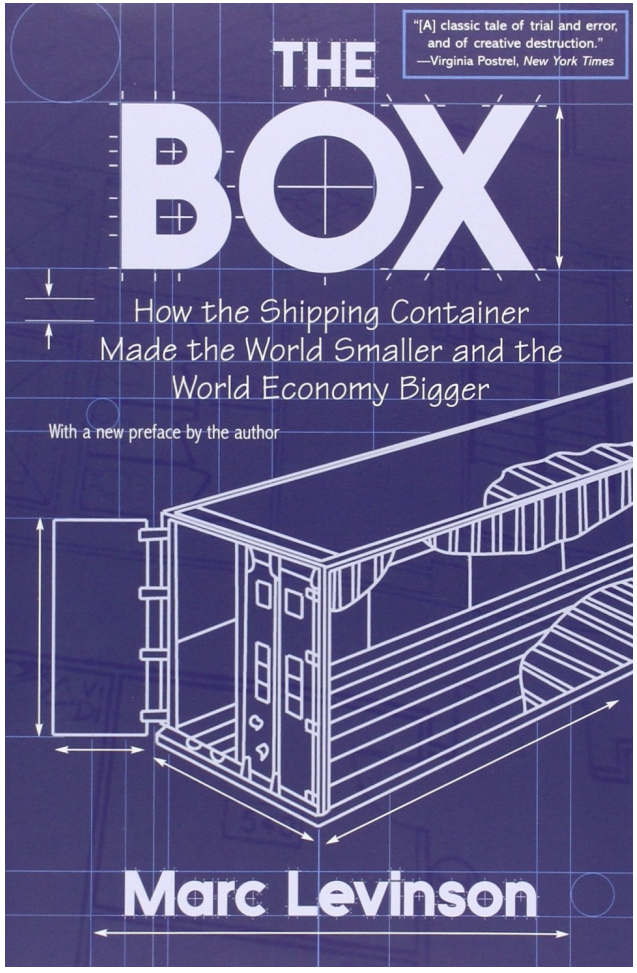
"[A] classic tale of trial and error, and of creative destruction."
—Virginia Postrel, New York Times
THE
BOX
How the Shipping Container Made the World Smaller and the World Economy Bigger
With a new preface by the author
Marc Levinson

“Docker containers wrap up a piece of software in a complete filesystem that contains everything it needs to run: code, runtime, system tools, system libraries – anything you can install on a server. This guarantees that it will always run the same, regardless of the environment it is running in.” https://www.docker.com/what-docker


Microsoft


FreeBSD®

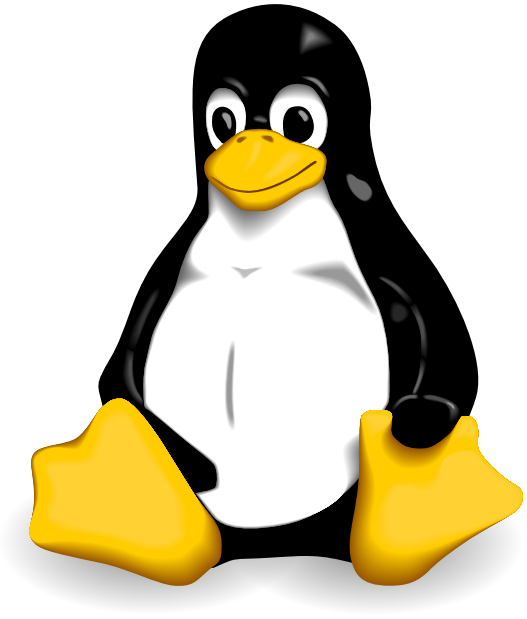

automation, repeatability, reliability

- Repeatable builds
- Ship one artifact
- Do not depend on the runtime environment
- Standard reusable tooling
- Build, test, run pipeline with one container shipped through it

unikernels

- “Unikernels are specialised, single-address-space machine images constructed by using library operating systems.”
- “Unikernels are self contained applications that bundle all their dependencies, and only their dependencies.”
- Containers bundle most dependencies, but rely on the kernel the host is running. Unikernels bundle everything.

Service Model

- Pets are given names like pussinboots.cern.ch
- They are unique, lovingly hand raised and cared for
- When they get ill, you nurse them back to health

- Cattle are given numbers like vm0042.cern.ch
- They are almost identical to other cattle
- When they get ill, you get another one

- Future application architectures should use Cattle but Pets with strong configuration management are viable and still needed

Code you want

to run

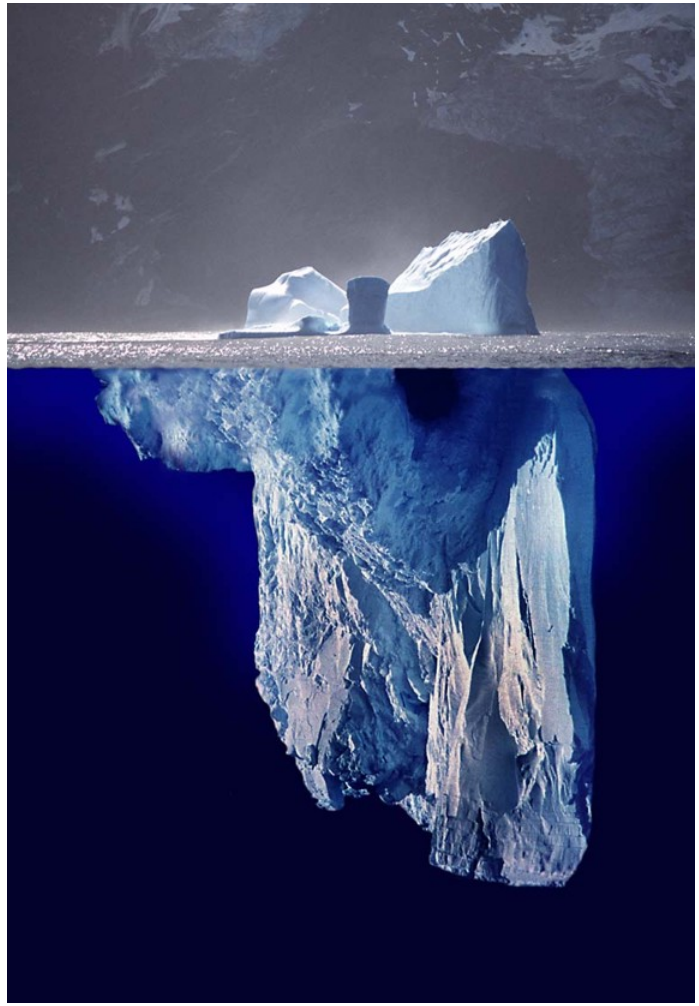

Code your OS

includes

Currently Linux has over 25 million lines of code...

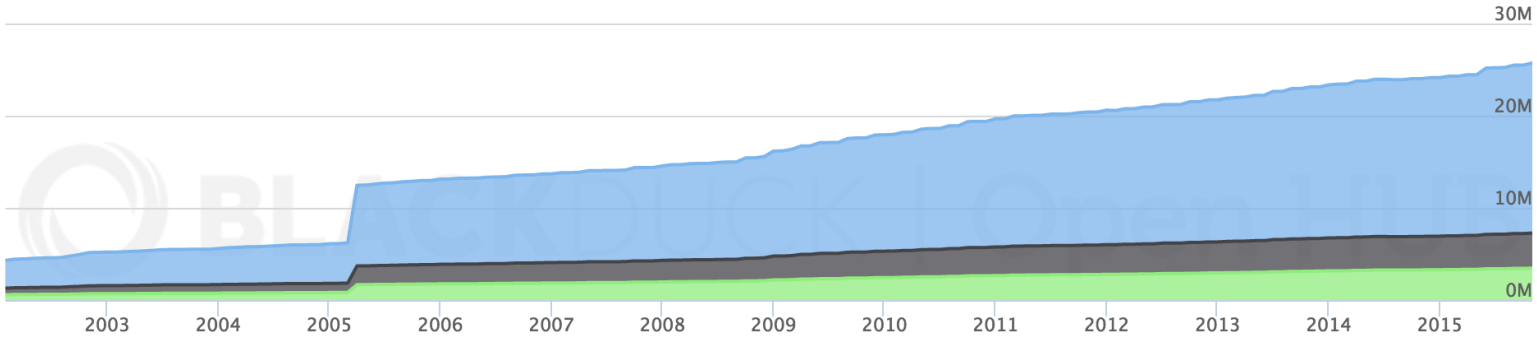

... and Windows has 50 million.

Code you want to run

Code your OS includes

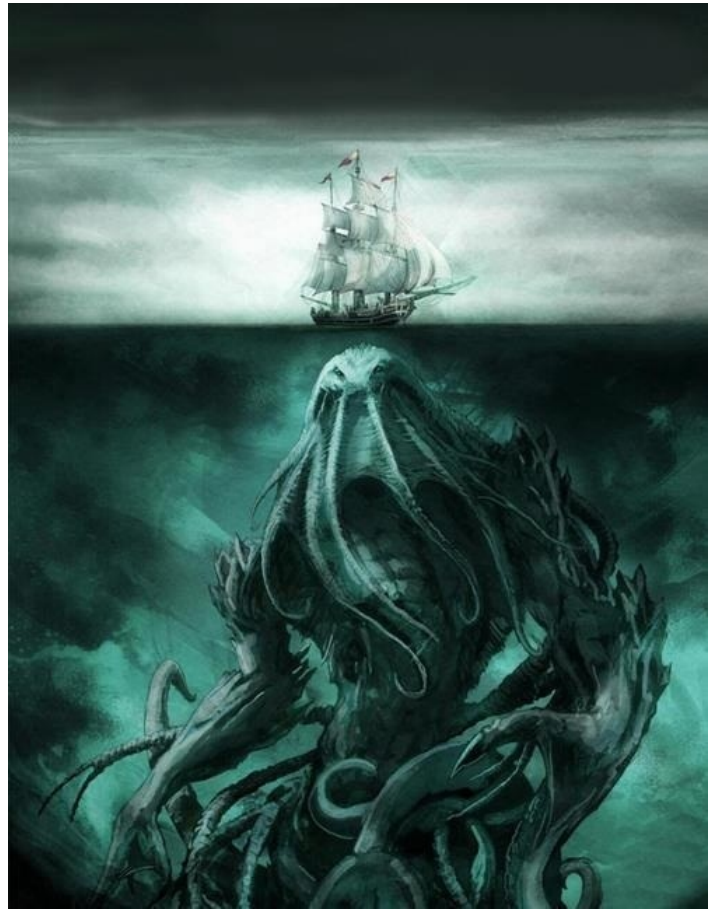

Unikernels

- First cut down on the amount of operating system dependencies lurking under your code.
- Just link exactly what you need as libraries, eg tcp, filesystems, etc
- Then they can also make that code less scary.

A security hardened container

- No large OS attack surface
- Just what you need, no extra shell or other executables, so small attack surface
- Can run inside virtual machine for sandboxing
- Language guarantees, like type safety and memory safety
- Can use additional sandboxing techniques: ASLR, NaCl etc
- Whole system hardening
- Ideal for embedded systems

Making systems programming less scary

- Systems programming is unusually difficult compared to other forms of programming
- OS development and design are the pinnacle or programming achievement, and the highest calling for any programmer
- Systems programmers are inherently superior to other kinds of programmers
- A competent systems programmer will naturally be gifted in all other forms of programming

I find these assumptions laughable. – Jay Osako

- Systems programming has a huge amount of technical debt
- Operating systems are huge, not very modular.
- Rebuilding with modern agile development is not as hard as people pretend.
- Easier when not working inside a kernel in C.

New simpler, more secure stacks in high level languages

- Static typing
- Memory safety
- Use of formal methods
- zero-cost abstractions
- Test driven development
- Fuzz testing

Rust

- zero-cost abstractions
- guaranteed memory safety
- threads without data races
- type inference
- minimal runtime

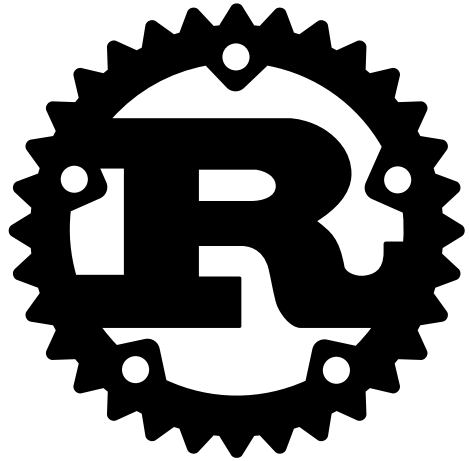

Go

- From the Plan 9 operating system heritage
- memory safety
- strong distributed programming libraries

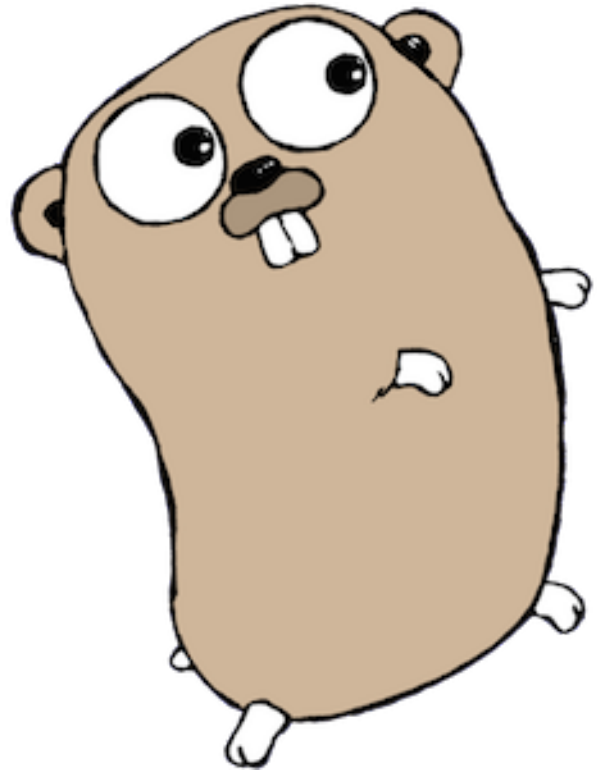

OCaml

- Functional language
- Full network stack implemented from TCP to SSL
- Memory and type safe

- Haskell
- C++
- LuaJIT
- Elixir
- JavaScript
- Swift

Examples of unikernels

- Mirage OCaml
- IncludeOS C++
- HalVM Haskell
- Ling Erlang
- runtime.js JavaScript
- ClickOS C++
- Rumprun C

how to get there?

- hack on some systems code
- implement protocols
- apply modern tools, processes, languages, methods
- have fun

- Reduce dependency on OS
- Don't shell out to command line
- Write portable code
- Just ship applications
- Do not try to introspect your environment

Zvi
@nivertech

Container with Ubuntu
Container with Alpine Linux
Linux ABI-compatible fat Unikernel
slim Unikernel

1:18 PM - 28 Feb 2016

6 8

Build, Ship, Run

Unikernels are still at the stage that Linux containers were three years ago before Docker

- Few users
- Hard to build
- Hard to ship
- Hard to run

Clearly this needs to be fixed for widespread use...

Unikernels are being used in production

- Specialist use cases
- Classified
- Networking devices
- Easrly adopters

Clearly this needs to be fixed for widespread use...

Unikernel.org

- Common community to share tooling, code and tests
- Working on ways to reuse existing code across languages
- Working on standard configuration and other layers
- Take the learnings from Mirage and apply more broadly.

Integrating unikernels into Docker

- Build: Dockerized toolchains
- Ship: Artifacts on Docker Hub
- Run: Same commands to run unikernels as containers

Questions?

- @justincormack
- justin.cormack@docker.com