

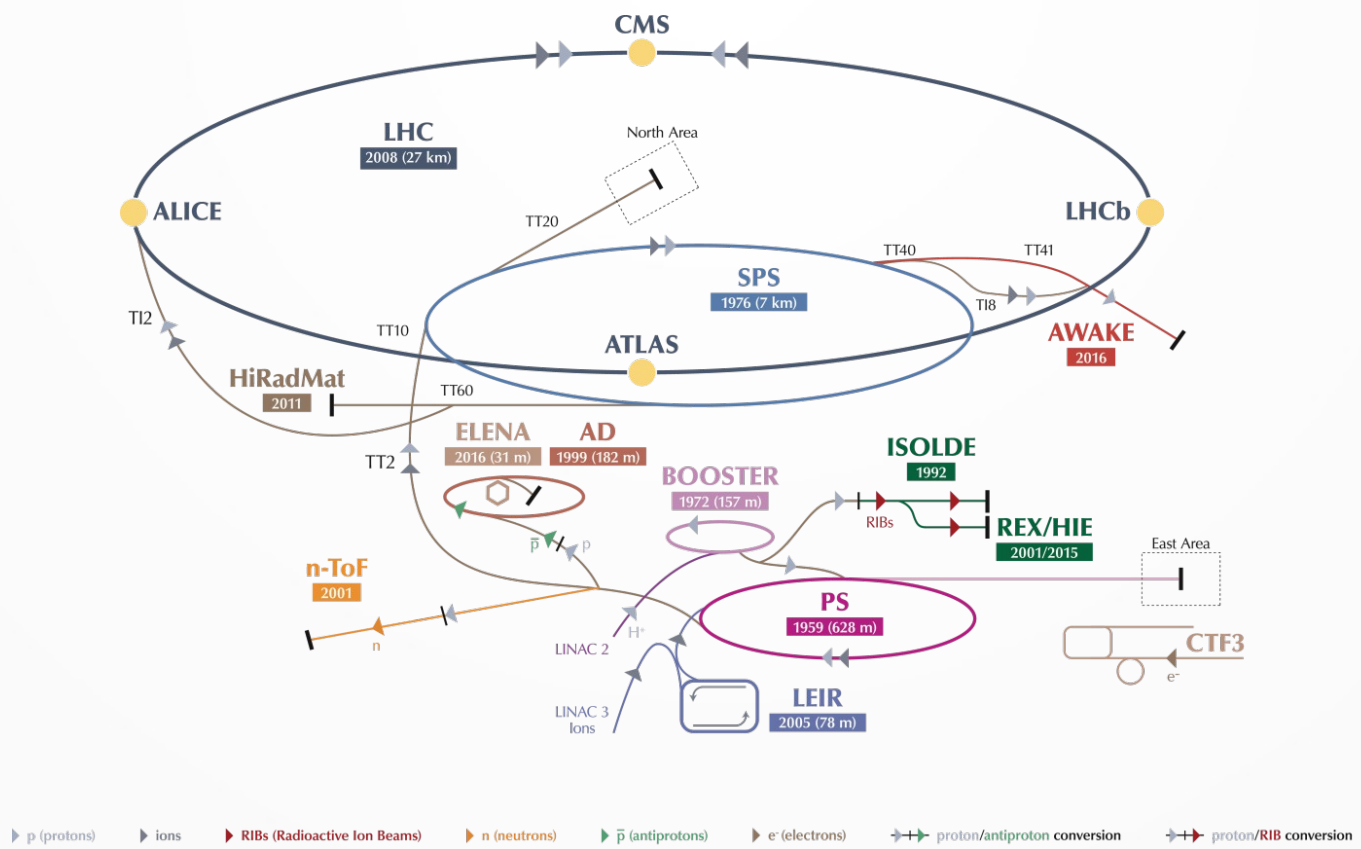
Unikernels – The New Black

Hristo Mohamed on behalf of the LHCb
collaboration
SREcon18 EMEA, Dusseldorf

What's CERN?




CERN Accelerator Complexes



LHC Large Hadron Collider SPS Super Proton Synchrotron PS Proton Synchrotron AD Antiproton Decelerator CTF3 Clic Test Facility
 AWAKE Advanced WAKefield Experiment ISOLDE Isotope Separator OnLine REX/HIE Radioactive EXperiment/High Intensity and Energy ISOLDE
 LEIR Low Energy Ion Ring LINAC LINear ACcelerator n-ToF Neutrons Time Of Flight HiRadMat High-Radiation to Materials

LHCb is a specialized b-physics experiment, designed primarily to measure the parameters of CP violation in the interactions of b-hadrons (heavy particles containing a bottom quark). Such studies can help to explain the Matter-Antimatter asymmetry of the Universe. The detector is also able to perform measurements of production cross sections, exotic hadron spectroscopy, charm physics and electroweak physics in the forward region.


CERN Accelerator Complexes



Unikernels – take one

Unikernels are specialized, single-address-space machine images constructed by using library operating systems.

Going back in time – wayyy back



Going back in time – more recent times

IBM Virtualization initial steps ~ 1960

Multics came to life in 1969

And then came the disruptive technology of the 70s – UNIX

Linux in 1991

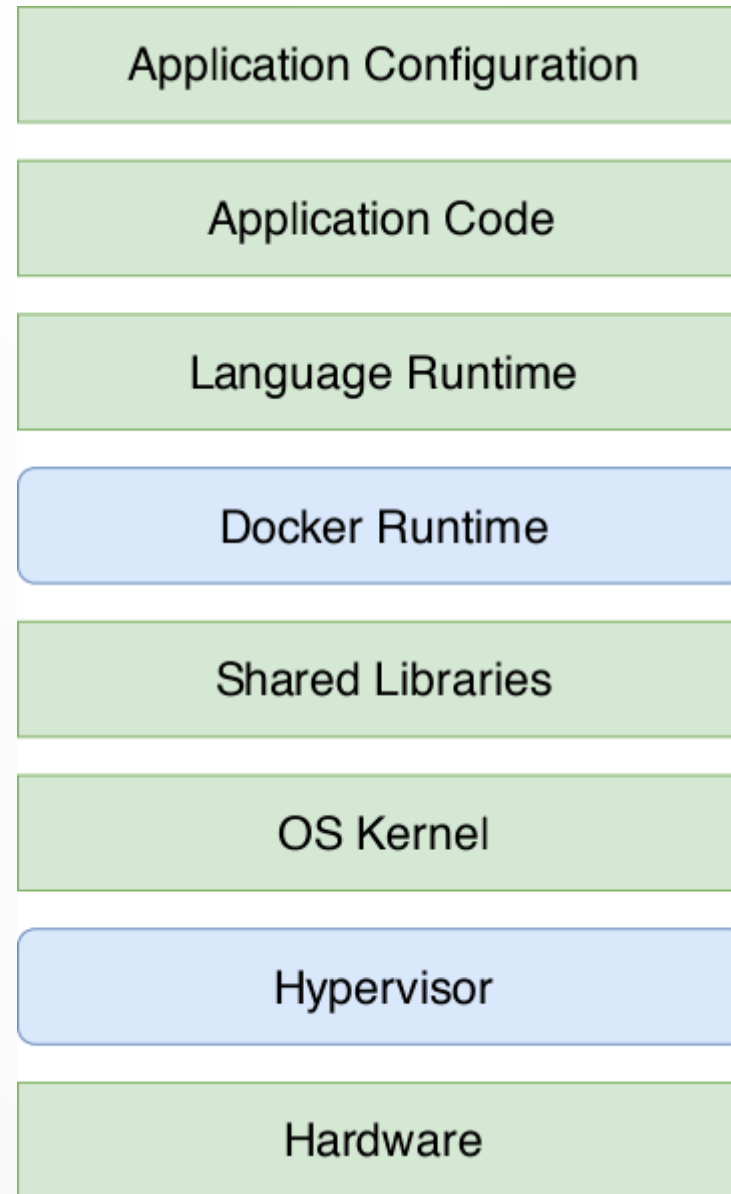
XEN came to life in 2003

KVM merged into mainline Linux kernel in 2007


LXC in 2008

Docker in 2013

How software is run



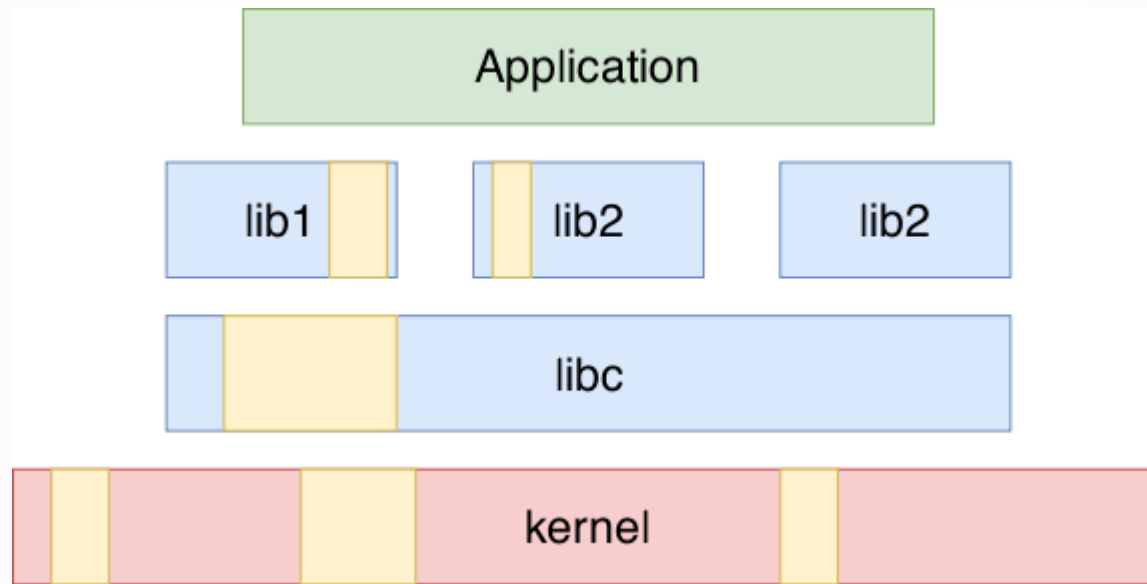
Average Application

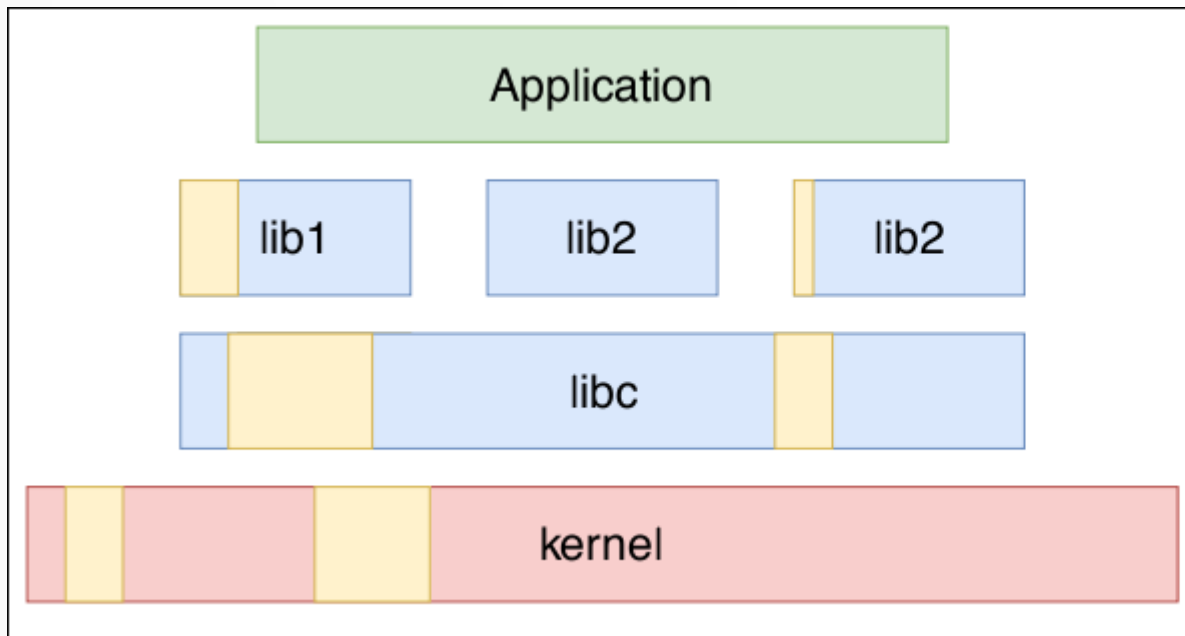


Unikernels – take two


Unikernels are specialized, single-address-space machine images constructed by using library operating systems.

Average – but unikernel Average!





Unikernel
Mystery
Machine



Unikernel

Hypervisor

Hardware

So, what do we gain by all this?

- Based on library OS, contain only needed components
- Actual single process & address space
- No virtual memory/context switching/different modes of execution
- Less code => less attack surface
- Completely immutable
- Small footprints & low boot times
- No characteristics of time-sharing Oses – permission checks, protection from other users, etc

Two camps – POSIX-complaint and purist

- Rumprun
- OsV
not just run-time, but
complete OS
compatibility
- MirageOS - oCaml
- IncludeOS - C++
- HalVM - Haskell
- LING – Erlang
- RuntimeJS (died :())

Is anything actually runnable right now?

Absolutely! Little demo time

So what can I run on a unikernel?

- Stateless services
- Honey Pots
- TOR nodes
- Network devices
- Anything highly specialized

Where is the catch?

- Unikernels are hard :(
Unikraft
- Yes debugging is nowhere near normal time shared OS, but work is being done
uniprof: Xen Domain Stack Profiler by Florian Schmidt
xenctx
gdb
- Logs forward data to somewhere – syslog protocol is basically a string in correct format

Questions?