# GPU, USB, NICs and other physical devices in your containers

ubuntu

Open Source Summit Europe 2017
Prague, Czech Republic

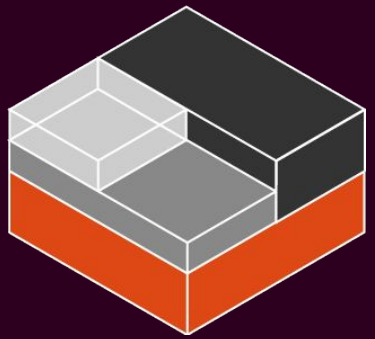

Stéphane Graber
LXD project leader, Canonical Ltd.

stgraber@ubuntu.com @stgraber
https://stgraber.org

# System containers

## What are they?

- The oldest type of containers
  *BSD jails, Linux vServer, Solaris Zones, OpenVZ, LXC and LXD.*
- Behave like standalone systems
  *No need for specialized software or custom images.*
- No virtualization overhead
  *They are containers after all.*

# Device passthrough

Why would you need physical devices?

- Computation
  *CUDA/OpenCV based GPU computation.*
- Very fast networking
  *40/100Gb/s networking including hardware slicing (SR-IOV).*
- Interacting with devices
  *Cell phones, scientific equipments, HSMs, phone cards, …*
- Dedicated block storage
  *Physical disks or partitions.*

# Device passthrough

How is it done?

- Device access is handled by the host kernel
  *So the hardware doesn't need any special capabilities.*
- Device nodes are identified and passed to the container
  *So the workload doesn't need to be container-aware.*
- Devices can be shared very efficiently
  *The same device can be passed to multiple containers, allowing for simultaneous access, so long as the kernel driver allows it.*
- Devices can be attached and detached on the fly
  *They are just files or kernel constructs so can be moved around, added and removed as needed without requiring a reboot of the host or container.*

# Let’s look at LXD

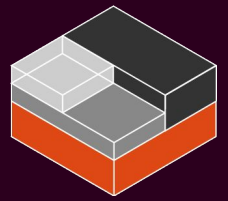

## What it **IS**

- ➔ Simple
  *Clean command line interface, simple REST API and clear terminology.*
- ➔ Fast
  *Optimized storage and image management.*
- ➔ Secure
  *Safe by default. Combines all available kernel security features.*
- ➔ Scalable
  *From a single container on a developer’s laptop to thousands of containers per host in a datacenter.*

# Let’s look at LXD

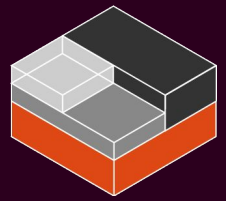

## Device types

- nic
  *A network interface.  Physical device or virtual interface.*
- disk
  *A mounted block device or filesystem path.*
- gpu
  *A physical GPU, available as DRI render nodes and CUDA interfaces.*
- usb
  *A USB device that can be used by libusb.*
- unix-char/unix-block
  *A raw unix char/block device.*

# Demo time!

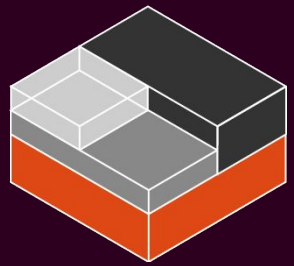

# Physical devices in containers

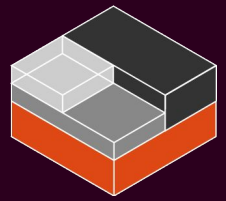

## Let’s recap

- ➔ Just about any device can be attached to a container
  *So long as the host kernel supports the device.*

- ➔ No special hardware or OS configuration required
  *No need for VFIO or virtual devices and no overhead.*

- ➔ Devices can easily be shared with multiple containers
  *The exact same device can be attached to any number of containers and in some cases (like GPUs) can be used by all of them simultaneously.*

# LXD: the container hypervisor

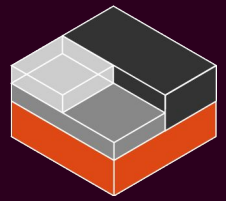

## Daily images

- Alpine (3.3, 3.4, 3.5, 3.6 and edge)
- ArchLinux
- CentOS (6 and 7)
- Debian (wheezy, jessie, stretch, buster and sid)
- Fedora (24, 25 and 26)
- Gentoo
- OpenSUSE (42.2 and 42.3)
- Oracle (6 and 7)
- Plamo (5.x and 6.x)
- Sabayon
- Ubuntu (12.04, 14.04, 16.04 LTS, 17.04, 17.10 and Core 16)

# LXD: the container hypervisor

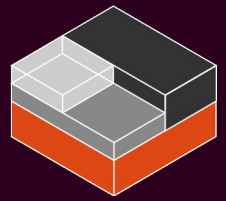

LXD is available on

- Alpine Linux
- Arch Linux
- Debian
- Elementary
- Fedora
- Gentoo
- OpenSUSE
- Solus
- Ubuntu

# LXD: the container hypervisor

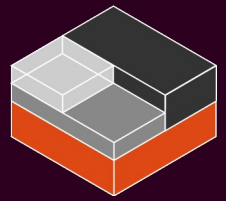

## LXD itself

- ➔ Written in Go
- ➔ Fully translatable
- ➔ API client in Go and python
- ➔ Apache2 licensed
- ➔ No CLA

git clone https://github.com/lxc/lxd

Stéphane Graber
LXD project leader, Canonical Ltd.

stgraber@ubuntu.com @stgraber
https://stgraber.org

https://linuxcontainers.org/lxd
https://github.com/lxc/lxd

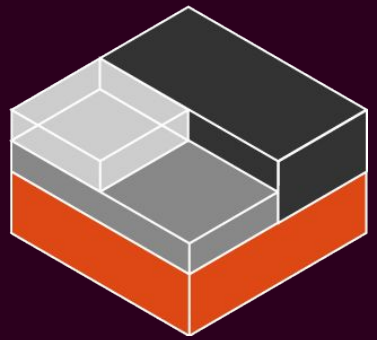

ubuntu

# Questions?

Try LXD at: https://linuxcontainers.org/lxd/try-it
Stickers are available in front!