

Unprivileged GPU containers on a LXD cluster

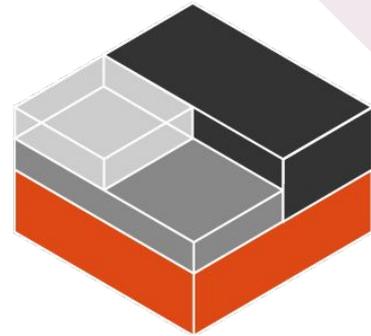
GPU-enabled system containers at scale

Stéphane Graber
LXD project leader

@stgraber
<https://stgraber.org>
stephane.graber@canonical.com

Christian Brauner
LXD maintainer

@brau_ner
<https://brauner.io>
christian.brauner@canonical.com



CANONICAL  ubuntu 

What are system containers?



01

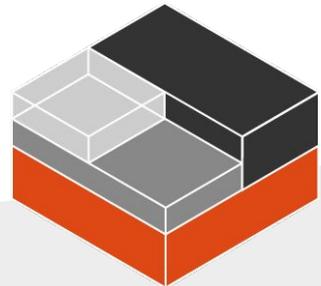
They are the oldest type of containers
BSD jails, Linux vServer, Solaris Zones, OpenVZ, LXC and LXD.

02

They behave like standalone systems
No need for specialized software or custom images.

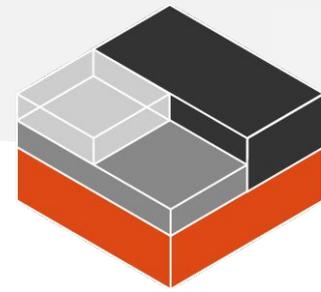
03

No virtualization overhead
They are containers after all.



LXD

System
container
manager



nova-lxd

command line tool

your own client/script ?

LXD REST API

LXD

LXD

LXD

LXD

LXC

LXC

LXC

LXC

Linux kernel

Linux kernel

Linux kernel

Linux kernel

Host A

Host B

Host C

Host ...

What LXD is



01

Simple

Clean command line interface, simple REST API and clear terminology.

02

Fast

Image based, no virtualization, direct hardware access.

03

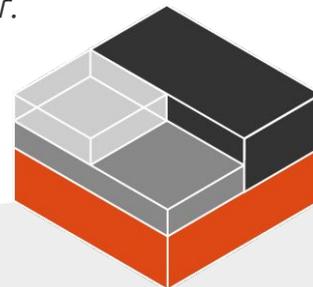
Secure

Safe by default. Combines all available kernel security features.

04

Scalable

From a single container on a laptop to tens of thousands of containers in a cluster.



What LXD isn't



01

Another virtualization technology

LXD offers an experience very similar to a virtual machine.

But it's still containers, with no virtualization overhead and real hardware.

02

A fork of LXC

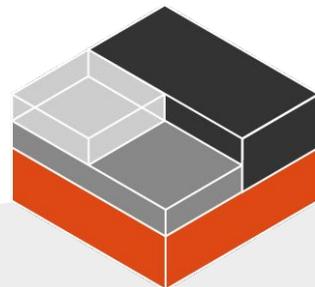
LXD uses LXC's API to manage the containers behind the scene.

03

Another application container manager

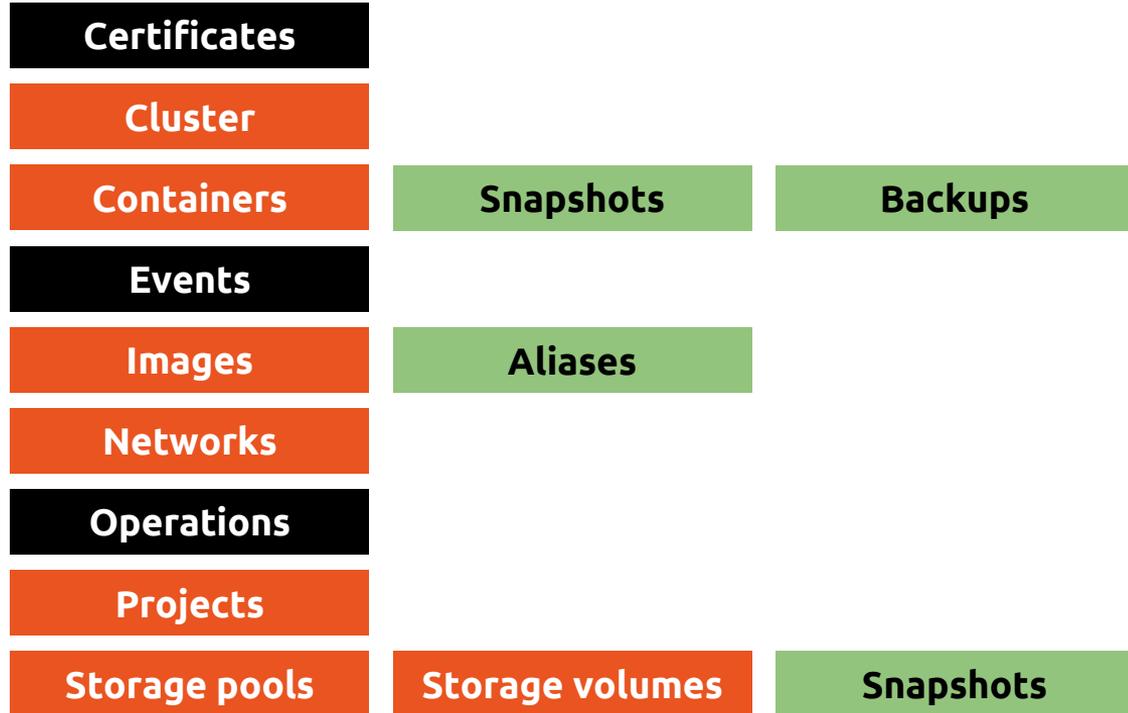
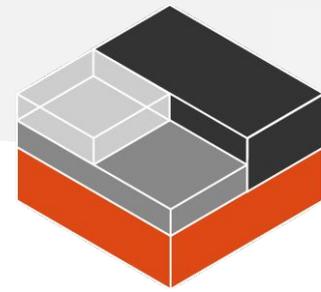
LXD only cares about full system containers.

You can run whatever you want inside a LXD container, including Docker.



LXD

Main components



LXD clustering



01

Built-in clustering support

No external dependencies, all LXD 3.0 or higher installations can be instantly turned into a cluster.

02

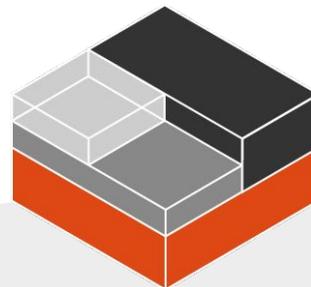
Same API as a single node

Clients that aren't clustering aware just see it as a very large LXD instance.

03

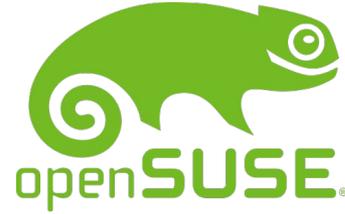
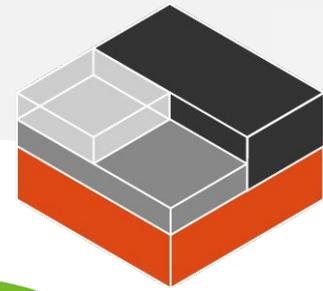
Scales to thousands of containers on dozens of nodes

Uses a built-in distributed database and cross-connections between the nodes to offer a consistent view to clients and load-balance containers.



Wide
selection of
images

Updated daily



gentoo linux™



debian

CentOS



GPUs in LXD containers



01

Support for all GPU vendors

02

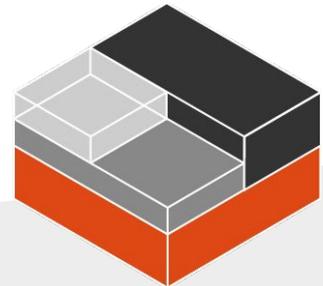
Integration with NVIDIA container (libnvidia-container)

03

Share a GPU with multiple containers

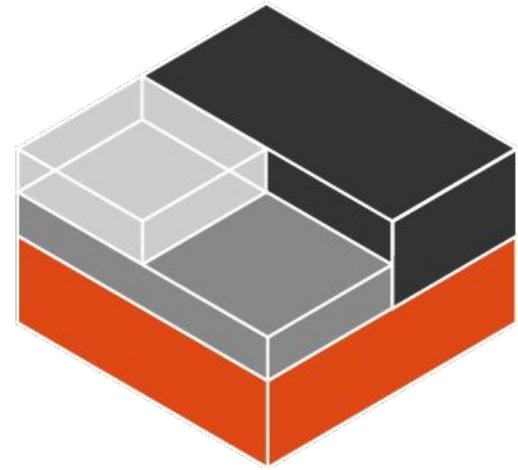
04

Fine grained selection of GPU





Demo time!



Let's recap



01

System containers as alternative to virtual machines

*Very similar workflow to virtual machines or cloud instances.
Without overhead, with direct hardware access and no need for virtualization support.*

02

Large scale management with clustering

*Single entity to manage, highly available and easily scalable.
Combined with CEPH, allows for fault tolerance.*

03

Direct hardware access

No virtualized hardware, directly pass your devices to your containers.

04

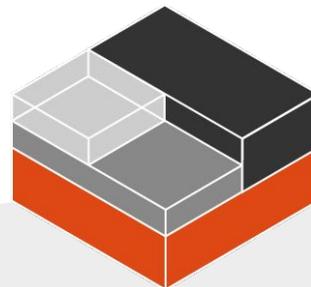
Safe and fast

State of the art container security and isolation.

05

Production ready

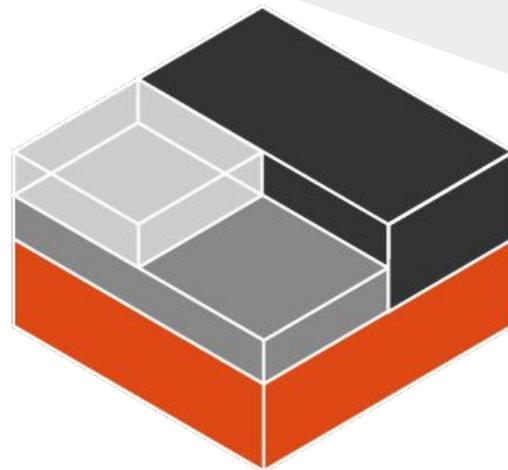
*Long term support releases with 5 years of support.
LXD has been around for over 4 years, LXC for over a decade.*



Questions ?

Website: <https://linuxcontainers.org/lxd>
Code: <https://github.com/lxc/lxd>
Online demo: <https://linuxcontainers.org/lxd/try-it>

We have stickers, come
get them in front!



Stéphane Graber
LXD project leader

@stgraber
<https://stgraber.org>
stephane.graber@canonical.com

Christian Brauner
LXD maintainer

@brau_ner
<https://brauner.io>
christian.brauner@canonical.com

