LXC

- LXC web panel: https://lxc-webpanel.github.io/index.html

Preparation

```
sudo apt-get install bridge-utils
sudo apt-get install lxc lxc-templates
sudo apt-get install cgmanager cgmanager-utils cgroup-bin
sudo lxc-checkconfig
```

Make sure cgroup filesystem is mounted

/etc/fstab

```
cgroup  /sys/fs/cgroup    cgroup    defaults    0    0
```

LXC Files

Priviledged containers

- /var/lib/lxc default container place
- /var/cache/lxc download cache

Unpriviledged containers

```
/home/user
├── .cache
│   ├── lxc
│       ├── download
│       │   └── ubuntu
│       │       ├── precise
│       │       │   └── amd64
│       │       └── xenial
│       │           └── amd64
│       └── run
│           └── lxc
│               └── lock
│                   └── home
├── .config
│   └── lxc
│       └── default.conf
├── .local
```

```
│   └── share
│       ├── lxc
│       │   ├── gitlab
│       │   │   ├── config
│       │   │   ├── rootfs
│       │   │   │   ├── bin
│       │   │   │   ├── boot
│       │   │   │   ├── dev
│       │   │   │   ├── etc
│       │   │   │   ├── ...
│       │   │   └── snaps
│       │   │       └── snap0
│       │   └── apache
│       │       ├── config
│       │       ├── rootfs
│       │       │   ├── bin
│       │       │   ├── boot
│       │       │   ├── dev
│       │       │   ├── etc
│       │       │   ├── ...
│       │       └── snaps
│       │           └── snap0
```

Basic usage

```
lxc-create -n test-container -t ubuntu
lxc-create -n test-container -t ubuntu -B btrfs
lxc-create -n test-container -t download -B btrfs
lxc-destroy -n test-container

lxc-start -n test-container
lxc-start -n test-container --daemon
lxc-stop -n  test-container

lxc-ls --fancy
lxc-info -n test-container

lxc-attach -n test-container
lxc-console -n test-container

lxc-snapshot -n test-container
```

Bind mounts

local/share/lxc/oldgitlab/config

```
lxc.mount.entry = /host/some/folder container/folder none
```

```
bind,create=dir,optional 0 0
```

Templates

Use template “ubuntu” and pass “-r trusty” argument to template:

```
lxc-create -n test-container -t ubuntu --dir/home/LXC/test-container -- -r
trusty -a amd64
```

Every template can show own help:

```
lxc-create -t download --help
```

Pass “List images” parameter to “download” template

```
lxc-create -t download -n test-container -- -l
```

Available templates:

```
/usr/share/lxc/templates/lxc-gentoo
/usr/share/lxc/templates/lxc-centos
/usr/share/lxc/templates/lxc-oracle
/usr/share/lxc/templates/lxc-alpine
/usr/share/lxc/templates/lxc-fedora
/usr/share/lxc/templates/lxc-sshd
/usr/share/lxc/templates/lxc-altlinux
/usr/share/lxc/templates/lxc-opensuse
/usr/share/lxc/templates/lxc-download
/usr/share/lxc/templates/lxc-busybox
/usr/share/lxc/templates/lxc-ubuntu
/usr/share/lxc/templates/lxc-ubuntu-cloud
/usr/share/lxc/templates/lxc-openmandriva
/usr/share/lxc/templates/lxc-cirros
/usr/share/lxc/templates/lxc-plamo
/usr/share/lxc/templates/lxc-archlinux
/usr/share/lxc/templates/lxc-debian
```

Network

Direct bridge

On host: use br0 as main interface:

/etc/network/interfaces

```
auto eth0
iface eth0 inet static
    address 0.0.0.0

auto br0
iface br0 inet dhcp
    bridge_ports eth0
```

Edit container configuration and set net bridge interface:

```
lxc.network.link = br0
```

Unprivileged containers

Create user:

```
sudo useradd -s /bin/bash -c 'gitlab lxc user' -m lxcgitlab
```

```
sudo passwd mylxcusr
```

Find subuids and subgids for created user

```
sudo grep lxc /etc/sub{gid,uid}
```

```
/etc/subgid:lxcgitlab:165536:65536
/etc/subuid:lxcgitlab:165536:65536
```

Allow new user create network interfaces:

/etc/lxc/lxc-usernet

```
lxcgitlab veth br0 10
```

Now login as new user (using ssh or su). Create default lxc configureation in user directory:

```
mkdir -p ~/.config/lxc
cp /etc/lxc/default.conf ~/.config/lxc/default.conf
```

Edit file below and put subuid and subgid:

~/.config/lxc/default.conf

```
...
lxc.network.link = br0
lxc.id_map = u 0 165536 65536
lxc.id_map = g 0 165536 65536
```

```
...
```

Create container:

```
lxc-create -t download -n gitlab
```

Snapshot

```
~$ lxc-snapshot -n gitlab
newgidmap: gid range [165536-165537) -> [331072-331073) not allowed
error mapping child
setgid: Invalid argument
```

lxc-snapshow is calling newgidmap

```
 lxc-snapshot 20160426080144.153 WARN     lxc_confile -
confile.c:config_pivotdir:1877 - lxc.pivotdir is ignored.  It will soon
become an error.
 lxc-snapshot 20160426080144.153 INFO     lxc_confile -
confile.c:config_idmap:1498 - read uid map: type u nsid 0 hostid 165536
range 65536
 lxc-snapshot 20160426080144.153 INFO     lxc_confile -
confile.c:config_idmap:1498 - read uid map: type g nsid 0 hostid 165536
range 65536
 lxc-snapshot 20160426080144.258 WARN     lxc_confile -
confile.c:config_pivotdir:1877 - lxc.pivotdir is ignored.  It will soon
become an error.
 lxc-snapshot 20160426080144.258 INFO     lxc_confile -
confile.c:config_idmap:1498 - read uid map: type u nsid 0 hostid 165536
range 65536
 lxc-snapshot 20160426080144.258 INFO     lxc_confile -
confile.c:config_idmap:1498 - read uid map: type g nsid 0 hostid 165536
range 65536
 lxc-snapshot 20160426080144.377 INFO     lxcbtrfs -
bdev/lxcbtrfs.c:btrfs_snapshot:306 - btrfs: snapshot create ioctl returned 0
 lxc-snapshot 20160426080144.397 WARN     bdev - bdev/bdev.c:bdev_copy:393 -
Failed to update ownership of
/home/lxcgitlab/.local/share/lxc/oldgitlab/snaps/snap2/rootfs
 lxc-snapshot 20160426080144.397 INFO     lxc_container -
lxccontainer.c:copy_file:2622 - Error stat'ing
/home/lxcgitlab/.local/share/lxc/oldgitlab/lxc_rdepends
 lxc-snapshot 20160426080144.398 INFO     lxc_container -
lxccontainer.c:copy_rdepends:2781 - Error copying reverse dependencies
```

Autostart

~/.local/share/lxc/gitlab/config

```
lxc.start.auto = 1
lxc.start.delay = 5
lxc.start.order = 100
lxc.group = onboot
```

```
lxc-autostart  processes  containers with lxc.start.auto set. It lets the
user start, shutdown, kill, restart containers in the right order, waiting
the right time. Supports fil-
tering by lxc.group or just run against all defined containers. It can also
be used by external tools in list mode where no action will be performed and
the list of affected con-
tainers (and if relevant, delays) will be shown.
```

Edit cron

```
crontab -e
```

```
@reboot /usr/bin/lxc-autostart --all
```

Use systemd (**not finished yet**): Enable autostarting systemd for user:

```
sudo loginctl enable-linger lxcgitlab
```

~/.config/systemd/user/lxc-autostart.service

```
...
```

Limit resources

config

```
# 512MB memory limit, 256MB soft limie - system treats it as low mem
condition
lxc.cgroup.memory.limit_in_bytes = 512M
lxc.cgroup.memory.soft_limit_in_bytes = 256M
# total usage memory (swap+ram) limit to 1G
lxc.cgroup.memory.memsw.limit_in_bytes = 1G

# arbitary value which only sets relative priority between containers
lxc.cgroup.cpu.shares = 100

# restrict to use cpu core 0 and 1
lxc.cgroup.cpuset.cpus 0,1
```

```
lxc.cgroup.blkio.weight 500
```

Limiting runtime:

```
lxc-cgroup -n test-container cpu.shares 100
```

~/.local/share/lxc/gitlab/config

```
```

ulimit change for unpriv container

Inside container, this command fails:

```
ulimit -n 65536
```

Debug

```
lxc-start -n test-container
lxc-start: start.c: lxc_init: 402 failed loading seccomp policy
lxc-start: start.c: __lxc_start: 1086 failed to initialize the container
lxc-start: lxc_start.c: main: 341 The container failed to start.
lxc-start: lxc_start.c: main: 345 Additional information can be obtained by
setting the --logfile and --logpriority options.
```

```
lxc-start -n test-container -l DEBUG -o debug.log
```

```
lxc-start 1460629578.157 INFO     lxc_start_ui - lxc_start.c:main:264 -
using rcfile /var/lib/lxc/test-container/config
lxc-start 1460629578.158 WARN     lxc_log - log.c:lxc_log_init:316 -
lxc_log_init called with log already initialized
lxc-start 1460629578.159 WARN     lxc_cgmanager - cgmanager.c:cgm_get:985 -
do_cgm_get exited with error
lxc-start 1460629578.159 INFO     lxc_lsm - lsm/lsm.c:lsm_init:48 - LSM
security driver AppArmor
lxc-start 1460629578.159 ERROR    lxc_start - start.c:lxc_init:402 - failed
loading seccomp policy
lxc-start 1460629578.159 ERROR    lxc_start - start.c:__lxc_start:1086 -
failed to initialize the container
lxc-start 1460629578.159 ERROR    lxc_start_ui - lxc_start.c:main:341 - The
container failed to start.
lxc-start 1460629578.159 ERROR    lxc_start_ui - lxc_start.c:main:345 -
Additional information can be obtained by setting the --logfile and --
logpriority options.
```

Debug levels: FATAL ALERT CRIT ERROR WARN NOTICE INFO DEBUG TRACE

Configure debug levels in config file:

```
lxc.logfile
lxc.loglevel
```

Errors

Failed to load config for gitlab

Error after system upgrade. LXC has been updated from 2.0.1 to v 3.0.1

```
$ lxc-info gitlab
Failed to load config for gitlab
Failure to retrieve information on /home/lxcgitlab/.local/share/lxc:gitlab
```

SOLUTION:

```
cd /home/lxcgitlab/.local/share/lxc/gitlab
lxc-update-config -c config
```

Failed to mount cgroup

```
Failed to mount cgroup at /sys/fs/cgroup/systemd: Permission denied
```

Ubuntu 14.04 has LXC 1.0.7 which doesn't support running systemd inside the container.

You can install the LXC 1.1.4 backport available in trusty-backports which should fix that (enabled backports in /etc/apt/sources.list, then apt-get update, then apt-get -t trusty-backports install lxc) or use the stable LXC PPA at ppa:ubuntu-lxc/stable

failed to attach 'veth'...

Start container in foreground mode `lxc-start -n container -F`

```
lxc-start: conf.c: instantiate_veth: 2594 failed to attach 'veth7LY5W6' to
the bridge 'lxcbr0': Operation not permitted
lxc-start: conf.c: lxc_create_network: 2871 failed to create netdev
lxc-start: start.c: lxc_spawn: 1066 failed to create the network
lxc-start: start.c: __lxc_start: 1329 failed to spawn 'gitlab'
```

Start with debug logging:

```
...
lxc-start 20160418064521.427 ERROR    lxc_conf -
conf.c:instantiate_veth:2594 - failed to attach 'vethSIJAS1' to the bridge
```

```
'lxcbr0': Operation not permitted
lxc-start 20160418064521.456 ERROR    lxc_conf -
conf.c:lxc_create_network:2871 - failed to create netdev
lxc-start 20160418064521.456 ERROR    lxc_start - start.c:lxc_spawn:1066 -
failed to create the network
lxc-start 20160418064521.456 ERROR    lxc_start - start.c:__lxc_start:1329 -
failed to spawn 'gitlab'
...
```

From some reasons `lxcbr0` doesn't exists. Check if `lxc-net.service` works correctly:

```
journalctl -u lxc-net.service
```

```
systemd[1]: Starting LXC network bridge setup...
lxc-net[1280]: dnsmasq: failed to create listening socket for 10.0.3.1:
Address already in use
lxc-net[1280]: Failed to setup lxc-net.
grinnux2 systemd[1]: Started LXC network bridge setup.
```

Dnsmasq starts to spawn own DNS server on port :53 when on host system `bind` daemon is running. Dnsmasq wants to bind only to IP on lxcbr0 interface, so check if other process is listening on port :53

```
lsof -ni :53
```

Probably `bind` daemon is listening on all interfaces. To change this, edit

/etc/bind/named.conf.options

```
listen-on { 127.0.0.1; 192.168.0.231; };
listen-on-v6 { none; };
```

```
systemctl restart bind9
systemctl restart lxc-net
```

umount: /dev/zero: block devices are not permitted on filesystem

During shutdown

```
umount: /dev/zero: block devices are not permitted on filesystem
umount: /dev/urandom: block devices are not permitted on filesystem
umount: /dev/tty: block devices are not permitted on filesystem
```

> Ah - this is happening because the shutdown process is trying to do a force umount. We don't allow those (using seccom) because if the fs is a bind mount from a fuse or nfs, it'll disconnect the original mount.
>
> You can test this yourself by logging in and doing

> umount -f /dev/urandom
>
> versus
>
> umount /dev/urandom