NBD

Network Block Device capability in the Linux kernel

NBD is a capability avaible in the linux kernel by loading a kernel module that lets you connect to a server on the network and use one of its block devices as if it was a local device.

The module is part of the mainline kernel.

The userspace tools are probably available through your distribution

What is a Block Device

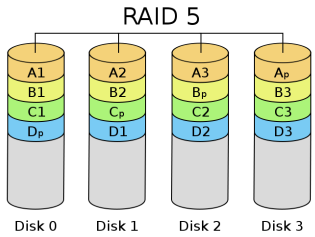

some_file.iso

- It is any kind of random access storage
- It is what we put a filesystem on (usually)

2

Most commonly it is a hard disk

But its also other “disks” like compact disks or flash drives.

RAID is a way to make a bunch of real block devices appear to be one big virtual block device.

Other things like LVM might create virtual block devices

A file isn't really a block device but we can use it as one . Note that the extension, .iso stands for the iso 9660 filesystem. And we usually put fileystems on block devices

We could use the file as a block device with the loopback driver, If we play this iso in a dvd player program it is actaully using it just like a block device.

I mention this because NBD can also serve files as if they were block devices

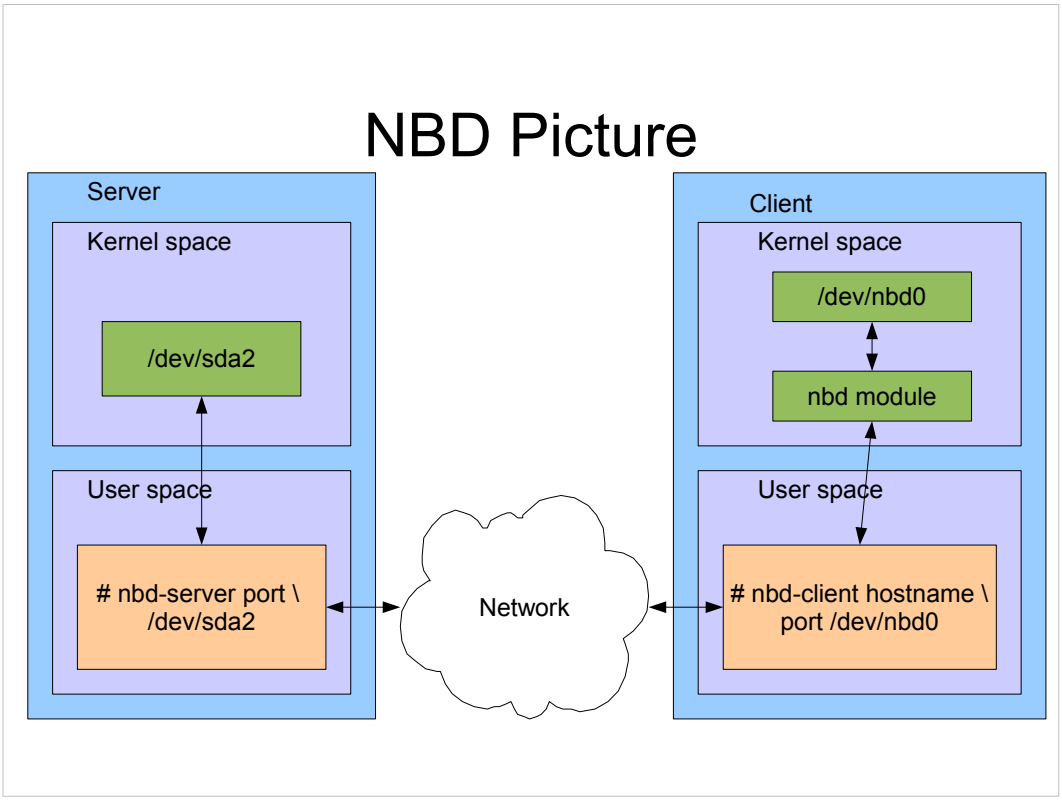

Here's a picture that shows how NBD works.
On the server we only have to run nbd-server. This is what you put on the command line. .....
On the client we have to run nbd-client and here's the command line we use for it.
We have to load the nbd module on the client.
When we access /dev/nbd0 on the client it goes over then network and we are really accessing /dev/sda2.

Sharing block devices

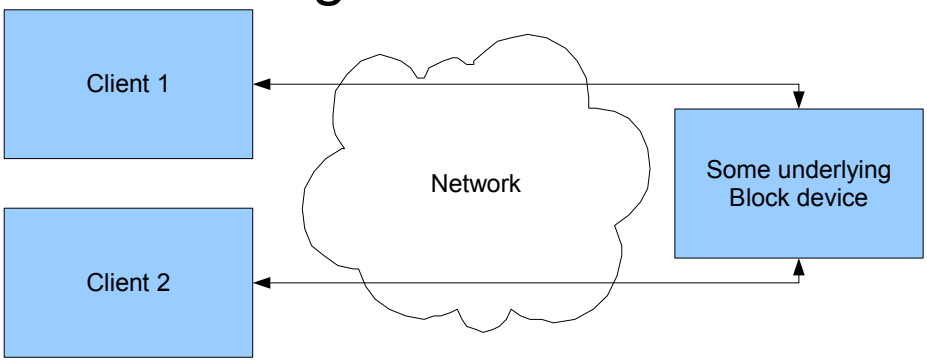

- The filesystem needs to support this or device needs to be read only.... or Bad Things will happen.

4

- Most of the time your computer is not expecting a block device to spontaneously change, which is what would happen if two clients were writing to the same filesystem.
- One way we could get aroung that is to do everything is read only
- You can have read/write if you use a cluster aware filesystem like gfs on the shared device
- GFS handles coordination between the clients so they don't get out of sync.
- There is also a cluster version of LVM that would let all the clients use and manage logical volumes on the device.

Access control / encryption

- Roll your own
  - SSH
  - VPN
  - Trusted Network
  - Inetd Style Wrappers
  - LUKS or similar disk level encryption

- Most likely you are going to be using this on a trusted network anyway
- You could encrypt the connection through vpn or ssh or whatever
- Roll your own solution with inetd or similar
- Or you could just decrypt only on the client with LUKS or similar

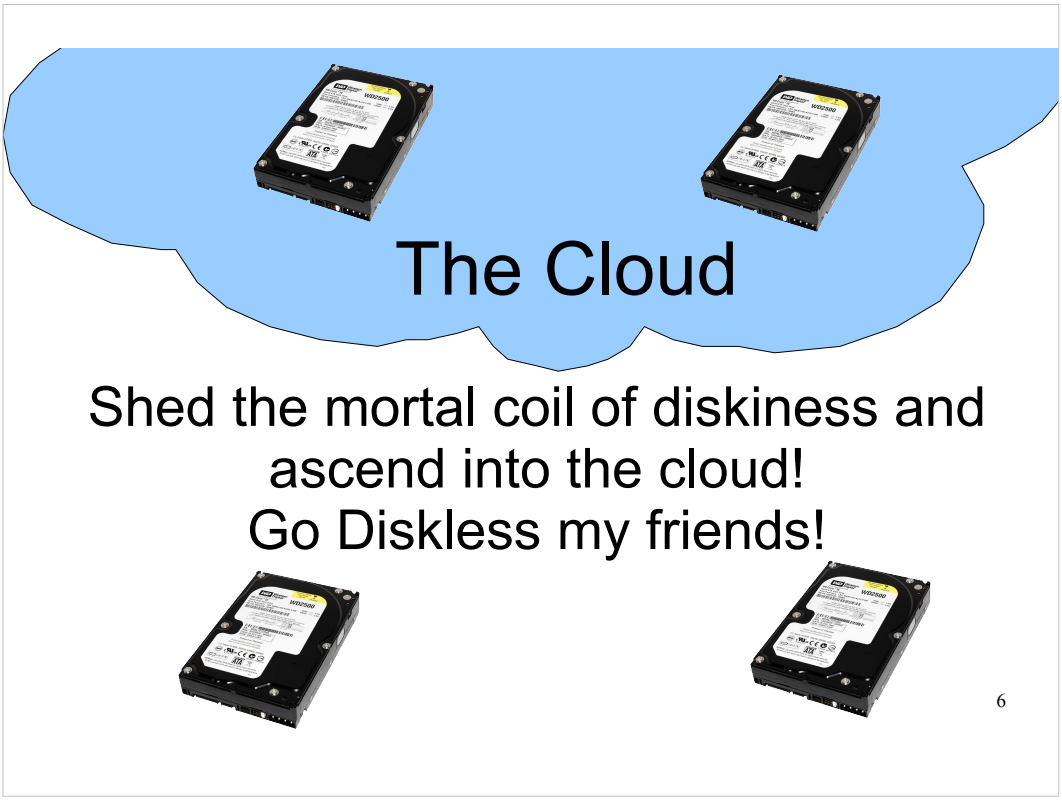

- Yes I have done this. I backed up my disk image over the network and then made a boot usb key that would use this image over the network.
- Put nbd-client and the nbd module into an initramfs image
- Inside early userspace, setup network and connect to nbd server.
- Mount root filesystem and continue boot as usual
- Make sure you don't mess up your network configuration or game over!
- Make sure that your startup/shutdown scripts don't mess with the network

Thin Client Idea

- Have a copy of your OS on a central server in an LVM logical volume
- Run nbd under inetd with a wrapper script to create a new snapshot for each connection.
- Have clients PXE boot and mount their root filesytems from the central server
- ????
- Profit