How to mount a VirtualBox VDI image

Posted on January 5, 2011 by Jeff Waugh

Don't believe the hype! It is entirely possible to mount a VirtualBox VDI image, just like a loopback filesystem… all you need are the right tools and know-how. Allow me to illustrate.

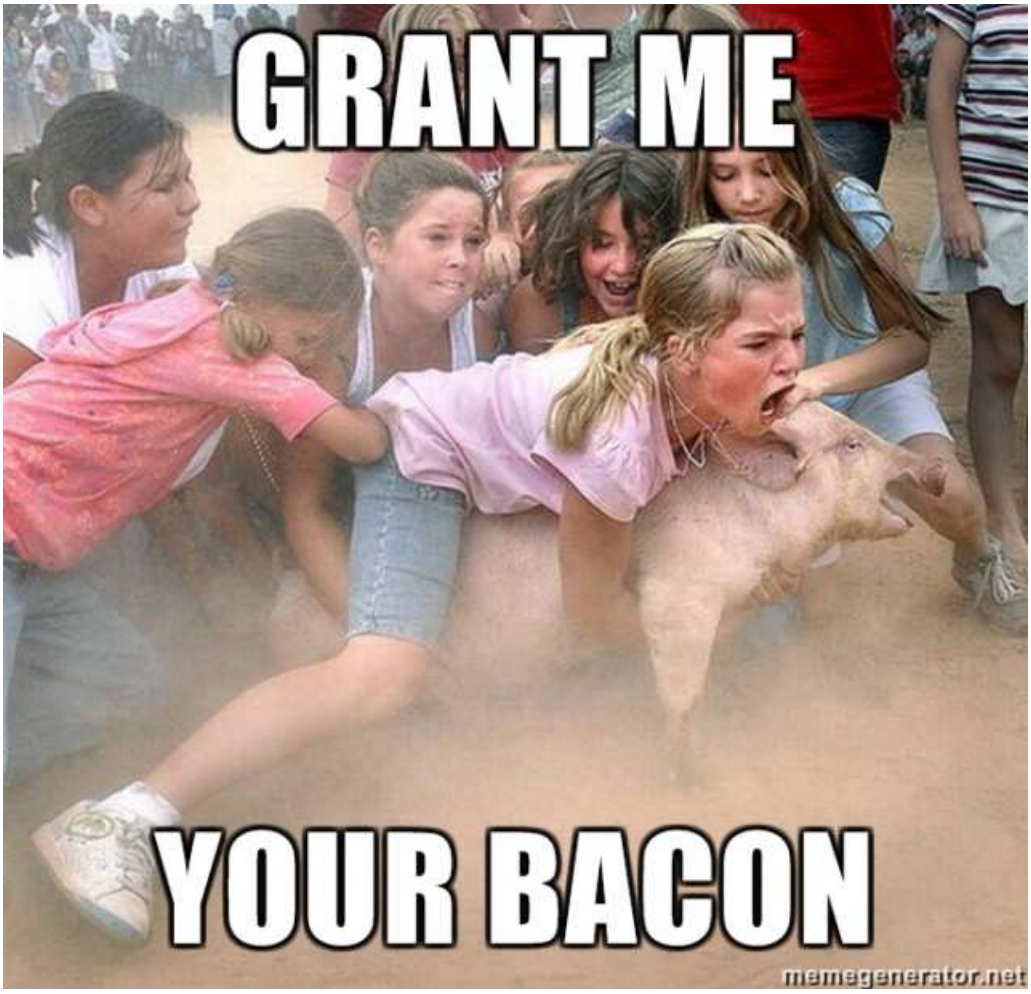

My apologies, that was the wrong illustration. Onward!

Before we start, it should be noted that you don't want to do this while your disk image is already in use. That is to say, if you're running a virtualised host using this image, GTFO.

First, install the QEMU tools. In Ubuntu, you'll find them in the `qemu-kvm` package. Whatever package your distribution ships which contains the `qemu-nbd` binary should be fine.

Load the *nbd* kernel module. Yes, I'm serious, the *network block device* module! (Note: All of the following commands require superuser permissions, so escalate your privileges in whatever way makes you most comfortable.)

```
modprobe nbd
```

Then run qemu-nbd, which is a user space loopback block device server for QEMU-supported disk images. Basically, it knows all about weird disk image formats, and presents them to the kernel via *nbd*, and ultimately to the rest of the system as if they were a normal disk.

```
qemu-nbd -c /dev/nbd0 <vdi-file>
```

That command will expose the entire image as a block device named `/dev/nbd0`, and the partitions within it as subdevices. For example, the first partition in the image will appear as `/dev/nbd0p1`.

Now you could, for instance, run cfdisk on the block device, but you will most likely want to mount an individual partition.

```
mount /dev/nbd0p1 /mnt
```

Gadzooks! Now you can muck around inside the filesystem to your heart's content. Go ahead and copy stuff in or out, or if you're feeling fruity, have some chrooty: `chroot /mnt`.

When you're done, unmount the filesystem and shut down the qemu-nbd service.

```
umount /mnt
qemu-nbd -d /dev/nbd0
```

This entry was posted in General and tagged Linux, nbd, qemu, Ubuntu, vdi, VirtualBox. Bookmark the permalink.

78 Responses to *How to mount a VirtualBox VDI image*

Chinmay K *says:*
November 30, 1999 at 11:00

Works like a charm! Thanks, it saved me a painful trip to the office just to get one file off a VM!

Reply

unity fan *says:*
January 5, 2011 at 22:31

hehe 😉 lol

Reply

dremon *says:*
January 6, 2011 at 01:19

Thanks a lot. I think there must be a "-c" parameter specified to qemu-nbd (at least on my machine).

Reply

Jeff Waugh *says:*
January 6, 2011 at 02:48

Oops, crucial missing parameter! Thanks for pointing that out, I'll fix the post.

Reply

Hubert Kario *says:*
January 31, 2011 at 01:21

I've noticed that "max_part" parameter for nbd kernel module in my distribution is 0, so for this to work, one has to do "modprobe nbd max_part=16" or similar

Reply

zzeroo *says:*
October 8, 2012 at 05:48

Yeah, my debian wheezy is one of them.
Without Your `modprobe nbd max_part=16` I'm not able to mount some partitions other the first one.

Reply

Paul Goller *says:*
January 10, 2011 at 11:47

I also have to run partprobe (from parted package) before I'm able to access the partitions (tested under Debian Squeeze RC). After that the partitions will accessible through /dev/disk /by-id/*nbd*part[0-9]*.

Sadly the nbd module seems to be very unstable for this kind of operation. After copying ~1.5GB of data from a VDI-File the module dies and the copying process (cp and rsync tested) will produce I/O Errors.

Reply

Jeff Waugh *says:*
January 11, 2011 at 03:49

The versions of qemu and nbd worked well for me on Ubuntu 10.10.

In fact, I was meant to be fetching packaging stuff off the disk image (because the new VirtualBox was wedging the machine) when I found myself doing the packaging and building while chrooted in it! Not a terrible stress test. 😉

Reply

shabeer *says:*
October 21, 2011 at 06:42

You are right Paul!

Me too getting same IO errors are using rsync over 1.5GB

Reply

aperson *says:*
February 4, 2011 at 15:32

I believe this to only be applicable to fixed-size images. Dynamically-sized images cannot be mounted, IIRC.

Reply

Jeff Waugh *says:*
February 4, 2011 at 15:39

Not so! All of my disk images are dynamic. Hooray!

Reply

zzeroo *says:*
October 8, 2012 at 05:53

Full ack, I copy 11GB data to a dynamically allocated .VDI file, at the moment without any trouble.

Reply

Ulli Hochholdinger *says:*
March 3, 2011 at 03:30

Hi,
Many thanks for the Tip.

I use debian squeeze. To get this to work there you have to insert the nbd module with the paramter max_part with appropriate number of partitions to use:

rmmod nbd
modprobe nbd max_part=16

without this Parameter you won't get the partitions. partprobe gives a weird Error:
"Error: Error informing the kernel about modifications to partition /dev/nbd0p1 — Invalid argument. This means Linux won't know about any changes you made to /dev/nbd0p1 until you reboot — so you shouldn't mount it or use it in any way before rebooting."

Cheers Ulli

Reply

Pablo Mateos *says:*
October 12, 2012 at 07:37

I've tried with debian squeeze and I did too have to use the modprobe nbd max_part=16, otherwise got the same error.

Thank you Uli !

Reply

vinicius *says:*
March 28, 2014 at 06:49

+1

I also had to load it with max_part, otherwise no partitions.

Thanks Ulli!

Reply

Erik N *says:*
March 16, 2011 at 07:36

Works like a charm on Ubuntu 10.10

Reply

Simon W. *says:*
May 17, 2011 at 08:05

Thanks YOU!!!

you R my life saver! 🙂
Yours tutorial 100% works 🙂
Ubuntu 11.04

Reply

jorge *says:*
June 11, 2011 at 04:00

so, this is if you are running virtualbox on an ubuntu system, how would I go about it if I'm running virtualbox on a winxp machine?
thanx.

Reply

antiplex *says:*
June 24, 2011 at 01:05

finally a solution, many thanks jeff!
i thought i might want to bring this to ubuntuforums as well, so i created a posting there mentioning your blog-entry here: http://ubuntuforums.org/showpost.php?p=10972071&postcount=6

Reply

gmart *says:*
June 24, 2011 at 05:30

Sadly the nbd module is very unstable for this kind of operation. It only mostly works in Ubuntu 9:10, 10:04, 10:10 and now 11:04. Unfortunately, 'mostly' is not good enough for production work. Sooner or later the copy process (also cp and rsync tested) will produce I/O Errors. This has caused me no end of headaches. I have investigated many fixes and work arounds. The I/O error proble is easy to duplicate:

```
root@somehost:~# qemu-nbd --connect=/dev/nbd0 ./gmos/tmp4UatFh.qcow2
root@somehost:~# dd if=/dev/nbd0p1 of=/dev/null bs=1M
dd: reading `/dev/nbd0p1': Input/output error
1175+0 records in
1175+0 records out
1232076800 bytes (1.2 GB) copied, 16.2789 s, 75.7 MB/s
```

Reply

gmart *says:*
June 24, 2011 at 05:59

My latest suspect for the I/O errors was double disk caching. I have found that adding the -n (–nocache) option to the qemu-nbd connect command seems to solve the problem (more testing needed). Thanks to http://superuser.com/users/79064/exot for that tip.

Reply

Tom H *says:*
July 4, 2011 at 02:15

had to use the max_part=16 trick, but got it working on Linux Mint 9. Very cool info, thanks for sharing!

Reply

The pabster *says:*
July 13, 2011 at 01:57

Thank you for this. 🙂

But it doesn't work for me. 🙁

qemu-nbd dies without doing nothing. Neither passing the max_part=16 to the module, nor running partprobe afterwards.

If I run qemu-nbd with the -v (verbose) option, I always get an error: "Failed setting NBD block size". I've googled that, but didn't find anything understandable. Any suggestion?

Reply

someone in the world [of linux ;)] *says:*
August 7, 2011 at 20:52

(Y)(Y)(Y)(Y)(Y)
(Y)(Y)(Y)(Y)
(Y)(Y)(Y)
(Y)(Y)
(Y)

Thank you man,
thank you many times
there should be more people and blogs like you/rs
makes linux better than already is

Reply

someone in the world [of linux ;)] *says:*
August 7, 2011 at 20:53

BTW
works perfectly in ubuntu 11 natty as it is

Reply

Erwin Kooi *says:*
October 4, 2011 at 05:23

In case you get
qemu-nbd: Could not access '/dev/nbd0': No such file or directory

Do this first:
apt-get install nbd-server

Reply

jlettvin *says:*
November 4, 2011 at 22:57

Nice access of vdi from host.
How about the other way, guest accesses passive partition in host?
Host system has 2 partitions Windows 7 and ubuntu 11.10 with Windows hosting.
Normal ubuntu guest running in vbox needs access to host ubuntu partition.
Any suggestions?

Reply

Jeff Waugh *says:*
November 5, 2011 at 03:57

You can use a raw disk from VirtualBox: http://www.virtualbox.org/manual/ch09.html#rawdisk

Reply

M. Belbut *says:*
June 12, 2012 at 06:50

If you are running a linux-based guest OS, you may consider doing the reverse: publish a nbd 'share' from your host, and accessing it from your guest. For that, instead of using qemu-nbd (which is specific for handling disk images on files), you should use nbd-server.

Reply

Anyeos *says:*
November 26, 2011 at 13:09

Mine worked one time but next it say:

fdisk /dev/nbd0

Unable to read /dev/nbd0

Wtf? nbd0 is unreadable? why?

I just tried with nbd1 and it worked but again it freezes last when I try to -d it.
Now I have two qemu-nbd freezed in my terminals. So? What I must wait for?
killalled it all (it hangs with a message).

SOoo!!! I cannot use this method guys!!

I need something to read my VDI image and for now I cannot get nothing to work.

Thanks but no thanks... :S

Reply

AJ *says:*
July 25, 2012 at 12:40

should be ndb0 (zero)

Reply

Anyeos *says:*
November 26, 2011 at 15:25

Sorry guys my qemu-nbd is using 100% cpu when I give to it the -d parameter. It does not disconnect anything just use 100% CPU.

What happen with that? How to update to the lastest version? I'm using Ubuntu 10.10 and don't want to update to 11.

I searched in google but there are no ppa related to backports of qemu for Ubuntu 10.10. I'm using custom kernel manually built (2.6.38-rc3).

I see a possible solution there are some bug report about qemu-nbd hangued. It is related with a kernel 2.6.38 update. I guess that what I'm using (rc3) have the bug. I will recompile a kernel but 2.6.39.

Bye

Reply

henri *says:*
January 11, 2012 at 03:12

On my slackware box (Q6600, kernel 3.2), partprobe -s /dev/nbd0 could not notify the creation of the new device /dev/nbd0p1 (ERROR in modification transmitting to the kernel)!! although fdisk /dev/nbd0 shows correctly partition 1

Reply

henri *says:*
January 11, 2012 at 03:39

I just not read the Ulli Hochholdinger message... SORRY

In my Slackware 13.37, the max_part parameter is also mandatory!

Reply

Olaf L *says:*
January 18, 2012 at 23:18

Works flawlessly on Mint-11 (Ubuntu-11.04) as well – when the correction (max_part=16 on modprobe; and qemu-nbd -c as stated above) are applied. Thank you for a great tip.

Reply

roberto *says:*
January 21, 2012 at 20:29

On suse 11.4 the tip:

rmmod nbd
modprobe nbd max_part=16

make the difference
It works like a charm thanks ;)) very useful

Reply

Rabin *says:*
January 26, 2012 at 01:21

This is the only solution for Dynamic virtual disks!!!. You better emphasize this. All other solution are deploying loop back concept which is lacking the awareness to partition table and dynamic drives.
R.

Reply

Antonio *says:*
January 26, 2012 at 03:09

Thanks a lot for the useful post. I test with two vdi images: ext4 and ntfs, in read and write. My VDI disks are dynamically-sized images.
It Works ^_^

Bye,
Antonio

Reply

Dast *says:*
January 28, 2012 at 13:14

Awesome! Worked beautifully on Ubuntu 11.04. Thanks!

Reply

Pat *says:*
February 7, 2012 at 16:09

Woohoo!!! I searched google all night to figure out how to mount a dynamic vdi image, and this worked like a charm. Most other methods required finding the data offset number, or didnt work right with dynamic images (for example, after I finally figured out the correct offset number, I kept getting a “Failed to read last sector” error for an ntfs filesystem trying to mount it with a /dev/loop0).
I did have to use the “max_part” option as listed in the above post.

For clarity and to help others, the following worked for me on openSuse 11.4 – 64bit host, mounting an ntfs filesystem from a dynamic vdi created from a virtualbox windows7 guest:

1) unload network block devicei kernel module, in case it is loaded without the max_part option (if you have already linked any nbd devices to vdi files, remove them with “qemu-nbd -d /dev/nbd0” or similar)

```
rmmod nbd
```

2) reload nbd with max_part option

```
modprobe nbd max_part=16
```

3) link your nbd to your vdi file

```
qemu-nbd -c /dev/nbd0 VDIFILE.vdi
```

4) scan for partitions, this should create the new device files in dev

```
partprobe
```

5) list your device files to make sure they were created

```
ls -l /dev/nbd0*
```

6) to view the sizes of any partions found in the vdi, you can run fdisk and print the partition table

```
fdisk /dev/nbd0
p
```

7) finaly, mount the desired partition (I used '-r' for readonly to be safe)

```
mount -t ntfs -r /dev/nbd0p2 /MNTPOINT
```

Reply

js200300953 *says:*
February 15, 2012 at 17:03

Thanks a lot.

Reply

alby1729 *says:*
February 20, 2012 at 20:22

YOU ARE GREAT!

Reply

Marcus Downing *says:*
March 15, 2012 at 21:01

Ulli, that hint with rmmod saved my skin. Thank you!

Reply

M. Belbut *says:*
April 7, 2012 at 06:29

Hi,

I tried your technique and it works for me. But I wanted to save the pain of having to go to the terminal and eventually create mountpoins for my partitions, so I was hoping 'partprobe' to allow me to handle the nbdXpY volumes through nautilus...
Should it work?
When I run partprobe, I get this error:

```
miguel@cdrsp-laptop-miguel:~$ sudo rmmod nbd
miguel@cdrsp-laptop-miguel:~$ sudo modprobe nbd max_part=16
miguel@cdrsp-laptop-miguel:~$ sudo qemu-nbd -c /dev/nbd0 '/home/miguel/VirtualBox
VMs/Tests/NewHardDisk1.vdi'
```

```
miguel@cdrsp-laptop-miguel:~$ sudo partprobe -s/dev/sda: msdos partitions 1 2 3
Warning: Unable to open /dev/sr0 read-write (Read-only file system). /dev/sr0 has
been opened read-only.
Warning: Unable to open /dev/sr0 read-write (Read-only file system). /dev/sr0 has
been opened read-only.
Error: Can't have a partition outside the disk!
/dev/nbd0: msdos partitions 1
miguel@cdrsp-laptop-miguel:~$
```

Any hint?

Reply

Jeff Waugh *says:*
April 11, 2012 at 03:55

Firstly, I don't understand why you're using `-s` or specifying `/dev/sda`... but I'm not entirely sure that `partprobe` is going to do what you want anyway. 🙂

Reply

romain dartigues *says:*
May 2, 2012 at 23:39

My Ubuntu didn't want to create the `/dev/nbd0p1` device, and `partprobe` failed with "Error: Error informing the kernel about modifications to partition /dev/nbd0p1 — Invalid argument. This means Linux won't know about any changes you made to /dev/nbd0p1 until you reboot — so you shouldn't mount it or use it in any way before rebooting.".

So i used another way:

```
$ fdisk -lu /dev/nbd0
Device Boot Start End Blocks Id System
/dev/nbd0p1 * 2048 16775167 8386560 83 Linux

$ mount -o loop,offset=1048576 /dev/nbd0 /mnt # sector * size (2048*512)
```

That was about it. Thank's.

Reply

Manuel H.P. *says:*
August 22, 2012 at 03:41

Thanks a lot Jeff,

Your excellent blog has allowed the recovery of a full VDI disk containing critical working documents of one colleague of mine.

All the best from the Pyrenees,

Manuel

Reply

dobbs *says:*
August 24, 2012 at 05:02

Sweet, just used qemu-nbd and gdisks to move the backup table of a resized VDI with GUID/GPT partition table.

Reply

Irregular Joe *says:*
August 26, 2012 at 05:59

Thanks for a great tutorial, just something to add for those that use LVM within their virtual machines: Once you mount the virtual drive, you can issue a 'pvscan' to check for any LVM physical volumes, then you may need to 'vgchange -ay vgname' to activate any LVM volumes. Once that has been done, you can do a normal mount of the logical volume.
The only potential problem may come if the volume group name on the virtual disk matches your host machine's group name. I resolved this by mounting a live-cd with my virtual machine and renaming the group with 'vgrename oldvgname newvgname'. You will also have to change the /etc/fstab and /boot/grub/grub.cfg to reflect the newvgname.
Hope this helps someone! Great info!

Reply

Derek *says:*
August 26, 2012 at 21:53

Couldn't get this working at first – the partitions didn't show up as devices. Before I read all of the comments and saw the 'partprobe' method, I'd discovered a Partition option to qemu-nbd, so I simply mounted the one partition that I wanted as nbd0. Thanks for the pointers!

Reply

Menner *says:*
September 2, 2012 at 06:16

Thank you very much for this post – it took only a few minutes to resize the root partition of the virtual machine!
First, I increased the VDI image size via
$ VBoxManage modifyhd archlinux.vdi –resize 2500
There's no VirtualBox GUI for that.
After that I used the nbd trick from this post to mount the VDI in the host.
I only needed to install qemu-utils; the kernel module was already there.
Then I could resize the partition via gparted GUI:
$ sudo gparted /dev/nbd0

Reply

Calmarius *says:*
October 8, 2012 at 06:20

I get I/O errors too even if -nocache is on.

Reply

kendall *says:*
February 15, 2013 at 04:54

works like a charm everytime on nearly any distro!

Reply

calamari *says:*
March 18, 2013 at 14:25

Your method works great for VMDK images too (even the expandable ones). Thanks!

Reply

arashium *says:*
May 5, 2013 at 21:16

Failed setting NBD block size

make sure you run qemu-nbd as *root*

Reply

pek *says:*
September 30, 2013 at 23:16

```
qemu-nbd -P 1 -c /dev/nbd0 /data/virtual-box/android.vdi
```

Work for me.

Reply

molecular *says:*
October 7, 2013 at 19:59

I had some problems with the partition ("/dev/nbd0p1") not showing up.

using: qemu-nbd -P 1 -c /dev/nbd0

worked. The partition #1 was then mapped directly to the existing node /dev/nbd0

thanks!

Reply

Roel de Cock *says:*
October 21, 2013 at 19:00

Damn, why didn't I come up with this solution myself! It even mounts a Windows-7 partition.

I was afraid for a moment that, after the untimely demise of vdfuse, I'd have to use libguestfs (which is less of a 'lib' and more of a 'massive dependency takeover of your PC').

But I still have to figure out how to make it work with an LVM volume.

Reply

Parad *says:*
January 8, 2014 at 14:06

OK, edge case here.
It's possibly a bad habit.
I sometimes create virtualHD's with no_partitions.
The entire HD is the file system.
So far I've not found a way to mount such a vdi.
fdisk see the disk once looped back to /mnt/nbd0
Though mounting cannot see a valid file system (type).
Any clues?

Thanks for a well simplified page.

M

Reply

David Strauss *says:*

January 18, 2014 at 08:34

> First, install the QEMU tools. In Ubuntu, you'll find them in the qemu-kvm package.
Whatever package your distribution ships which contains the qemu-nbd binary should be fine.

For Fedora systems, the package is qemu-img.

Reply

WJCarpenter *says:*

June 13, 2014 at 06:55

On Ubuntu 12.04, I found it in qemu-utils. You get that if you install qemu-kvm. If you don't need the other stuff, then the qemu-utils has a smaller footprint.

Reply

wyatt *says:*

February 5, 2014 at 13:01

thanks! this is great for moving files to and from my freeDOS VM!

Reply

LS *says:*

February 14, 2014 at 13:16

This guide works like a charm except for one specific issue that I ran into.

In my particular scenario, I want to mount a vdi disk used by my Win Server 2008 R2 guest to access some program files. After I mount the vdi and access the filesystem according to this guide, I can see *most* files on the disk. I can, however, not see the files that belong to some programs that I installed on the guest myself. I can not even see the directories used by these programs, and yet they are not hidden or anything of that sort.

Am I missing something obvious? I will be very grateful for your help.

Reply

Chathura *says:*

February 17, 2014 at 21:01

Thank you for infinite times and once more!!! This article was a LIFE SAVER!!!! Saved my time and Work!!!

Reply

Nelson Zamith *says:*

March 2, 2014 at 01:05

It saved me like a few people here, the time for a trip in the office 🙂
Really, thank you 🙂

Reply

ernemir *says:*
March 7, 2014 at 10:37

Works great for me. Thanks you!!!

Reply

Stuart Rothrock *says:*
April 23, 2014 at 02:54

Thanks so much. Works on Fedora as written. Appreciated!

Reply

Stuart Rothrock *says:*
April 23, 2014 at 02:57

Followup – Had to use kpartx -a /dev/nbd0 before mounting partitions. kpartx makes partitions visible via /dev/mapper directory.

* moderator may combine this and my previous post if desired.

Reply

Sven *says:*
April 26, 2014 at 18:20

It seems that to see the partitions in the image, one must load the nbd module with the max_part option, i.e., modprobe nbd max_part=16

Reply

Thomas P *says:*
August 6, 2015 at 00:19

Hear, Hear!
I too was suffering from being unable to write data to the actual image. To fdisk and parted, the mounted image appeared just as it should. But once you try to write data to the image, you hot a 3.5GB ceiling. Then, when you check the image file size, it is still just a few megs. My first clue something was very wrong was a simple dd from /dev/zero to the image filled 3.5GB in 2 seconds. Impossible write speeds for spinning rust but not for memory. So I was writing into a buffer and not the disk. did a search on qemu-nbd and found another site that showed the module option to set max_part=63. IF you run 'modinfo nbd', you will see this bit of info: "parm: max_part:number of partitions per device (default: 0) (int)"
Odd how the default is 0. My advice would e to set that to 8 or more and you should be good. Why you see 16 or 63 is beyond me, but they work.

Reply

Paul *says:*
May 1, 2014 at 05:49

Hi,
thanks a lot, work perfectly for me even with dynamically allocated VDI disk (thin

provisioning).

Just a note:
on my debian squeeze x86 I have to install qemu-utils package which contains qemu-nbd not found in the qemu-kvm package.

Thanks again
paul

Reply

Gannet *says:*
September 5, 2014 at 06:06

I am an XP refugee running Mint 17 with Cinnamon. Been trying to find out how to get into the .vdi files for ages. That worked a dream for me. Thank you. To get qemu-nbd I had to install qemu-utils using the software manager.

Reply

John *says:*
January 22, 2015 at 02:44

Thanks, Jeff! Four years after your post you're still helping folks (me!) figure this out. What worked for me on Fedora 21 was the following:

```
sudo modprobe nbd max_part=16
sudo qemu-nbd -c /dev/nbd0 VM_Win7x32.vdi
sudo mount /dev/nbd0p2 /mnt/vmfile
sudo umount /dev/nbd0p2
```

Reply

Muteezy *says:*
September 29, 2015 at 03:39

...aaand now it's my turn to jubilate! I couldn't make a copy of the original vdi file because of some errors so I bit the bullet and mounted it anyway using your short method. It did a short repair after mounting then BOOM!

I had access to all files. Gracias mucho!

Reply

su *says:*
February 17, 2015 at 09:11

Problem with partitions I solved by kpartx:
kpart -a /dev/nbd0
mount /dev/mapper/nbd0p1 /mnt

umount /mnt
kpart -d /dev/nbd0

Reply

John McLean *says:*
December 23, 2015 at 21:30

Worked a charm on Ubuntu 14.04.3 LTS! I've now managed to backup all the data from a virtual android phone.

Reply

Eric Duncan *says:*
February 3, 2016 at 08:49

@su: THANK YOU!

Had to use xpartx as well.

Reply

Avi Deitcher *says:*
April 7, 2016 at 18:14

This was extremely helpful… so I made a script to make it work more easily. See https://github.com/deitch/vdimount

Probably would be great if it did things like interactive, or dealing with LUKS filesystems, but it is a good start

Reply

Be the signal