

Import Xen virtuals to KVM/Xen

Warning: Preliminary. These are my notes I'm getting ready to try.

This is a work in progress, as I've not been able to find anything on the 'net about it. Everything I've seen assumes your Xen virtuals are being managed by libvirt already, so it is a matter of using some built in utilities.

The way I propose to do it is to manually generate a libvirt configuration file, modified to use values from the original Xen machine. We'll use the virt-install command with the -import and -print-xml options.

Lets try to convert the following Xen configuration file

```
builder='hvm'
memory=4096
vcpus=4      # change this to the number of virtual cpu's it should have
name='server.example.com'# change this to the name displayed on xm list

# change the MAC address to be unique.
# change the vifname so it shows up well on ifconfig output
vif= [
    'mac=00:16:3f:97:db:f3,bridge=br_dmz,vifname=server0',
]
disk=[
    # this points to the "drive" we created
    'phy:/dev/disk/by-path/ip-10.22.209.32:3260-iscsi-
iqn.2019-09.com.nas:server.disk0-lun-0,hda,w',
    # this points to the installation media
    # note, you may (should) remove the disk image after the install is
done
    'file:/mnt/devuan_beowulf_3.1.1_amd64_netinstall.iso,ioemu:hdc:cdrom,r',
]

# NOTE: you must change this to C after the initial installation or it will
boot right back into the installer
boot='cd' # this makes it boot from the installation media. Change it to 'c'
after install
vnc=1      # allow us to use vnc to connect during install (and after)
vncdisplay=4
usbdevice='tablet'
localtime=0
serial='pty' # allow us to connect from xm console
xen_platform_pci=1 # allows us to use the pv drivers for Linux hvm
uuid='44f33f93-f606-4513-8dc0-7e1eac977025'
```

To do this, will run virt-install with the -print-xml flag, which will simply show the configuration. In this case, I'm writing it to a file so it can be edited, then used to test, then define the virtual.

```
virt-install \
  --print-xml \
  --hvm \
  --connect qemu:///system \
  --name **CHANGEME** \
  --memory **CHANGEME** \
  --vcpus **CHANGEME** \
  --disk path=**CHANGEME**,bus=virtio,target=sda \
  --graphics vnc,port=**CHANGEME** \
  --noautoconsole \
  --os-variant debian9 \
  --metadata
uuid=**CHANGEME**,name=**CHANGEME**t,title=**CHANGEME**,description='**CHANGEME**' \
  --boot hd,cdrom,menu=on \
  --network bridge=**CHANGEME**,mac=**CHANGEME**,model=virtio \
  --network bridge=**CHANGEME**,mac=**CHANGEME**,model=virtio \
  > test.xml
```

The items marked **CHANGEME** in the above command are all populated from a Xen configuration file. **NOTE:** it is very important to place the boot drive first in the list, before any other drives or cdrom's

Once you have this done, edit the resulting file with any changes you may need. I change the CPU as per [Migrating with virsh](#) so I can migrate easily.

Now, shut down the virtual on its Xen hypervisor and test it with the following command:

```
virsh create test.xml --console
```

. This will start the new definition, putting you in the console so you can watch it boot. **Be very sure to shut down the virtual in its origin first.**

If all works well, shut down the virtual, then make the definition permanent with

```
virsh define test.xml
```

You are done.

From:
<https://kb.unixservertech.com/> - **Unix Server Tech Knowledge Base**

Permanent link:
<https://kb.unixservertech.com/unix/virtualization/virtlib/importxen?rev=1653020011>

Last update: **2022/05/20 04:13**

