


Case Study: To read the contents of a VHD/LVM block device from a XenServer v6.0.2 /XCP v1.6 (hypervisor)

When we look under the hood of your container we usually will find the following elements:

- VM container
- Virtual Network Interfaces
- Virtual Network
- Physical Network Interface
- Virtual Block devices (Guests HDD)
- Virtual Disk Image
- Storage Repository
- Physical Block Device
- Host

The following picture shows how this comes together:



Copied from : <http://www.crucial.com.au/blog/2011/05/19/xenserver-under-the-hood/>

Create a virtual machine to test this scenario, with an unencrypted LVM.
In this e.g we use VM-name=i-2-2294

Provided below are the commands to find the relevant VM details for VDI.

#xe vm-list

```
uuid ( RO)           : c76a3665-f7de-3918-e1b3-193f7fd45334
  name-label ( RW): i-2-2294-VM
  power-state ( RO): running
```

#xe vm-disk-list power-state=running --multiple

Disk 0 VBD:

```
uuid ( RO)           : 5a6bcb14-fa0b-4243-3035-8c22efc05568
  vm-name-label ( RO): i-2-2294-VM
  userdevice ( RW): 0
```

Disk 0 VDI:

```
uuid ( RO)           : 422facae-3deb-4181-98ed-168418ce5e23
  name-label ( RW): ROOT-2294
  sr-name-label ( RO): ed154d7d-8165-957c-6e18-c9d6f79d3b11
  virtual-size ( RO): 21474836480
```

Few more XenServer commands to display further VM details:

#xe vbd-list uuid=5a6bcb14-fa0b-4243-3035-8c22efc05568

```
uuid ( RO)           : 5a6bcb14-fa0b-4243-3035-8c22efc05568
  vm-uuid ( RO): c76a3665-f7de-3918-e1b3-193f7fd45334
  vm-name-label ( RO): i-2-2294-VM
  vdi-uuid ( RO): 422facae-3deb-4181-98ed-168418ce5e23
  empty ( RO): false
  device ( RO): hda
```

xe vdi-list uuid=422facae-3deb-4181-98ed-168418ce5e23

```
uuid ( RO)           : 422facae-3deb-4181-98ed-168418ce5e23
  name-label ( RW): ROOT-2294
  name-description ( RW):
  sr-uuid ( RO): ed154d7d-8165-957c-6e18-c9d6f79d3b11
  virtual-size ( RO): 21474836480
  sharable ( RO): false
  read-only ( RO): false
```

xe vm-param-list uuid=c76a3665-f7de-3918-e1b3-193f7fd45334

```
uuid ( RO): c76a3665-f7de-3918-e1b3-193f7fd45334
name-label ( RW): i-2-2294-VM
name-description ( RW): Template which allows VM installation from install
media
  user-version ( RW): 1
  is-a-template ( RW): false
  is-a-snapshot ( RO): false
  snapshot-of ( RO): <not in database>
  snapshots ( RO):
  snapshot-time ( RO): 19700101T00:00:00Z
```

```

        snapshot-info ( RO):
            parent ( RO): <not in database>
            children ( RO):
is-control-domain ( RO): false
        power-state ( RO): running
        memory-actual ( RO): 1073688576
        memory-target ( RO): 1073741824
        memory-overhead ( RO): 11534336
        memory-static-max ( RW): 1073741824
memory-dynamic-max ( RW): 1073741824
memory-dynamic-min ( RW): 1073741824
        memory-static-min ( RW): 1073741824
        suspend-VDI-uuid ( RW): <not in database>
        suspend-SR-uuid ( RW): <not in database>
        VCPUs-params (MRW): weight: 44; cap: 0
            VCPUs-max ( RW): 1
        VCPUs-at-startup ( RW): 1
actions-after-shutdown ( RW): Destroy
actions-after-reboot ( RW): Restart
actions-after-crash ( RW): Destroy
        console-uuids (SRO): 731f0627-d0b2-1115-dfdf-8e0847c7aa43
        platform (MRW): timeoffset: 0; nx: true; acpi: 1; apic:
true; pae: true; viridian: true
        allowed-operations (SRO): changing_dynamic_range; migrate_send;
pool_migrate; changing_VCPUs_live; suspend; hard_reboot; hard_shutdown;
clean_reboot; clean_shutdown; pause; checkpoint; snapshot
        current-operations (SRO):
        blocked-operations (MRW):
allowed-VBD-devices (SRO): 1; 2; 4; 5; 6; 7; 8; 9; 10; 11; 12; 13;
14; 15
        allowed-VIF-devices (SRO): 1; 2; 3; 4; 5; 6
        possible-hosts ( RO): f4af9bc3-24c4-4901-9e70-c725e3616c32;
8ccf3858-2369-4a95-bb69-47af99e2ba1e; 12060590-9c64-4a92-939f-095ac30042e8;
530a3863-be9a-4b0b-97bf-37bd7e6c05a8; f843afad-c35d-4c50-94ef-8bd933226564
        HVM-boot-policy ( RW): BIOS order
        HVM-boot-params (MRW): order: dc
HVM-shadow-multiplier ( RW): 1.000
        PV-kernel ( RW):
        PV-ramdisk ( RW):
        PV-args ( RW):
        PV-Legacy-args ( RW):
        PV-bootloader ( RW):
        PV-bootloader-args ( RW):
Last-boot-CPU-flags ( RO): vendor: GenuineIntel; features: 17bee3ff-bfebfbff-
00000001-2c100800
        Last-boot-record ( RO): ('struct' ('uuid' 'c76a3665-f7de-3918-
e1b3-193f7fd45334') ('allowed_operations' ('array')) ('current_operations'
('struct' ('OpaqueRef:8214b5a9-55b9-14e7-0247-b98bd76e022d' 'start_on'))))
('power_state' 'Halted') ('name_label' 'i-2-2294-VM') ('name_description'
'Template which allows VM installation from install media') ('user_version'
'1') ('is_a_template' ('boolean' '0')) ('suspend_VDI' 'OpaqueRef:NULL')
('resident_on' 'OpaqueRef:NULL') ('affinity' 'OpaqueRef:2b155ef5-4f8b-cca9-
c504-cf0520ceecf2') ('memory_overhead' '11534336') ('memory_target'
'1073741824') ('memory_static_max' '1073741824') ('memory_dynamic_max'
'1073741824') ('memory_dynamic_min' '1073741824') ('memory_static_min'
'1073741824') ('VCPUs_params' ('struct' ('weight' '44') ('cap' '0'))))

```

```

('VCPUs_max' '1') ('VCPUs_at_startup' '1') ('actions_after_shutdown'
'destroy') ('actions_after_reboot' 'restart') ('actions_after_crash'
'destroy') ('consoles' ('array')) ('VIFs' ('array' 'Opaqu
eRef:91acfb6e-cfd1-addb-436f-1d2324c096a6')) ('VBDs' ('array'
'OpaqueRef:4aa80198-3577-7958-89bc-2ecbcb574787' 'OpaqueRef:ef157ec6-2db7-
c832-38c2-24efce5afa69')) ('crash_dumps' ('array')) ('VTPMs' ('array'))
('PV_bootloader' '') ('PV_kernel' '') ('PV_ramdisk' '') ('PV_args' '')
('PV_bootloader_args' '') ('PV_legacy_args' '') ('HVM_boot_policy' 'BIOS
order') ('HVM_boot_params' ('struct' ('order' 'dc'))) ('HVM_shadow_multiplier'
('double' '1')) ('platform' ('struct' ('nx' 'true') ('acpi' '1') ('apic'
'true') ('pae' 'true') ('viridian' 'true'))) ('PCI_bus' '') ('other_config'
('struct' ('base_template_name' 'Other install media') ('mac_seed' '54d837fa-
2bd2-90f9-87e7-a623077e6402') ('install-methods' 'cdrom'))) ('domid' '-1')
('domarch' '') ('last_boot_CPU_flags' ('struct')) ('is_control_domain'
('boolean' '0')) ('metrics' 'OpaqueRef:dab92c67-5a45-293b-77df-1b2a20edf06a')
('guest_metrics' 'OpaqueRef:NULL') ('last_booted_record' '')
('recommendations' '<restrictions><restriction field="me mory-static-max\
max="\137438953472\ " /><restriction field="vcpus-max\ " max="\16\
"/><restriction property="number-of-vbds\ " max="\7\ " /><restriction
property="number-of-vifs\ " max="\7\ " /></restrictions>') ('xenstore_data'
('struct')) ('ha_always_run' ('boolean' '0')) ('ha_restart_priority' '')
('is_a_snapshot' ('boolean' '0')) ('snapshot_of' 'OpaqueRef:NULL')
('snapshots' ('array')) ('snapshot_time' ('dateTime.iso8601'
'19700101T00:00:00Z')) ('transportable_snapshot_id' '') ('blobs' ('struct'))
('tags' ('array')) ('blocked_operations' ('struct')) ('snapshot_info'
('struct')) ('snapshot_metadata' '') ('parent' 'OpaqueRef:NULL') ('children'
('array')) ('bios_strings' ('struct')) ('protection_policy' 'OpaqueRef:NULL')
('is_snapshot_from_vmpp' ('boolean' '0')) ('appliance' 'OpaqueRef:NULL')
('start_delay' '0') ('shutdown_delay' '0') ('order' '0') ('VGPIs' ('array'))
('attached_PCIs' ('array')) ('suspend_SR' 'OpaqueRef:NULL') ('version' '0'))'
resident-on ( RO): f843afad-c35d-4c50-94ef-8bd933226564
affinity ( RW): f843afad-c35d-4c50-94ef-8bd933226564
other-config (MRW): vgpu_pci; ; base_template_name: Other
install media; mac_seed: 54d837fa-2bd2-90f9-87e7-a623077e6402; install-
methods: cdrom

dom-id ( RO): 68
recommendations ( RO): <restrictions><restriction
field="memory-static-max" max="137438953472" /><restriction field="vcpus-max"
max="16" /><restriction property="number-of-vbds" max="7" /><restriction
property="number-of-vifs" max="7" /></restrictions>
xenstore-data (MRW): vm-data:
ha-always-run ( RW) [DEPRECATED]: false
ha-restart-priority ( RW):
blobs ( RO):
start-time ( RO): 20140228T15:18:41Z
install-time ( RO): 20140228T15:19:27Z
VCPUs-number ( RO): 1
VCPUs-utilisation (MRO): 0: 0.001
os-version (MRO): name: Ubuntu 12.04.3 LTS; uname: 3.8.0-29-generic; distro:
ubuntu; major: 12; minor: 04
PV-drivers-version (MRO): major: 6; minor: 1; micro: 0; build: 61809
PV-drivers-up-to-date ( RO): truevm-vif-list
memory (MRO):
disks (MRO):
networks (MRO): 0/ip: 146.169.44.139; 0/ipv6/0: fe80::417:b4ff:fe00:3f
other (MRO): platform-feature-multiprocessor-suspend: 1; feature-balloon: 1

```

```

Live ( RO): true
guest-metrics-last-updated ( RO): 20140228T15:18:57Z
cooperative ( RO) [DEPRECATED]: true
protection-policy ( RW): <not in database>
is-snapshot-from-vmpp ( RO): false
tags (SRW):
appliance ( RW): <not in database>
start-delay ( RW): 0
shutdown-delay ( RW): 0
order ( RW): 0
version ( RO): 0

```

```

# xe vm-vif-list uuid=c76a3665-f7de-3918-e1b3-193f7fd45334
uuid ( RO) : 172dc8fe-d8d8-f8b7-a7ff-7cd366fdf53c
vm-name-label ( RO): i-2-2294-VM
device ( RO): 0
MAC ( RO): 06:17:b4:00:00:3f
network-uuid ( RO): 832ba4a4-6457-f4c6-7d80-e3d5da6876ad
network-name-label ( RO): Bond 2+3

```

To locate the actual disk at the OS level: search with the VDI name format → **VG_XenStorage\${sr-uuid}VHD\${vdi-uuid}**

Search for the actual block device on the hypervisor using the last bit of the VDI uuid:

```

#ls -l /dev/mapper/ | grep 168418ce5e23
brw-rw---- 1 root disk 252, 23 Feb 28 16:09 VG_XenStorage--ed154d7d--8165--957c--6e18--c9d6f79d3b11-VHD--422facae--3deb--4181--98ed--168418ce5e23

```

```

# lvs | grep 422facae-3deb-4181-98ed-168418ce5e23
VHD-422facae-3deb-4181-98ed-168418ce5e23 VG_XenStorage-ed154d7d-8165-957c-6e18-c9d6f79d3b11 -wi-ao 20.05G

```

Can read the data using (not in order):

```

#strings /dev/mapper/VG_XenStorage--ed154d7d--8165--957c--6e18--c9d6f79d3b11-VHD--422facae--3deb--4181--98ed--168418ce5e23

```

Check/Confirm the different partitions available inside the running VM:

```

/dev/xvda1 -> /boot -> ext3 filesystem
/dev/xvda5 -> lvm for root and swap vols -> ext4 filesystem

```

=====

```

guest@demo-test:~$ sudo fdisk -l

```

```

Disk /dev/xvda: 21.5 GB, 21474836480 bytes
255 heads, 63 sectors/track, 2610 cylinders, total 41943040 sectors Units =
sectors of 1 * 512 = 512 bytes Sector size (logical/physical): 512 bytes / 512
bytes I/O size (minimum/optimal): 512 bytes / 512 bytes Disk identifier:
0x00098c7c

```

Device	Boot	Start	End	Blocks	Id	System
/dev/xvda1		2048	499711	248832	83	Linux
/dev/xvda2		501758	41940991	20719617	5	Extended
/dev/xvda5		501760	41940991	20719616	8e	Linux LVM

```

Disk /dev/mapper/systemvg-root: 14.8 GB, 14759755776 bytes

```

255 heads, 63 sectors/track, 1794 cylinders, total 28827648 sectors Units = sectors of 1 * 512 = 512 bytes Sector size (logical/physical): 512 bytes / 512 bytes I/O size (minimum/optimal): 512 bytes / 512 bytes Disk identifier: 0x00000000

Disk /dev/mapper/systemvg-root doesn't contain a valid partition table

Disk /dev/mapper/systemvg-swap_1: 6434 MB, 6434062336 bytes
255 heads, 63 sectors/track, 782 cylinders, total 12566528 sectors Units = sectors of 1 * 512 = 512 bytes Sector size (logical/physical): 512 bytes / 512 bytes I/O size (minimum/optimal): 512 bytes / 512 bytes Disk identifier: 0x00000000

Disk /dev/mapper/systemvg-swap_1 doesn't contain a valid partition table

=====

Using kpartx, we create device mappings:

```
# fdisk -l $(ls /dev/sm/backend/ed154d7d-8165-957c-6e18-c9d6f79d3b11/422facae-3deb-4181-98ed-168418ce5e23)
```

```
# kpartx -a $(ls /dev/sm/backend/ed154d7d-8165-957c-6e18-c9d6f79d3b11/422facae-3deb-4181-98ed-168418ce5e23)
```

Newly created volumes:

```
# ls /dev/mapper/  
422facae-3deb-4181-98ed-168418ce5e23p1  
422facae-3deb-4181-98ed-168418ce5e23p5
```

All three volumes are visible:

```
# ls /dev/mapper/ | grep 168418ce5e23  
422facae-3deb-4181-98ed-168418ce5e23p1  
422facae-3deb-4181-98ed-168418ce5e23p5  
VG_XenStorage--ed154d7d--8165--957c--6e18--c9d6f79d3b11-VHD--422facae--3deb--4181--98ed--168418ce5e23
```

Mount for /boot (ext3) partition works:

```
# mount /dev/mapper/422facae-3deb-4181-98ed-168418ce5e23p1 /mnt  
# cd /mnt/  
# ls  
abi-3.8.0-29-generic config-3.8.0-29-generic grub initrd.img-3.8.0-29-generic  
lost+found memtest86+.bin memtest86+_multiboot.bin System.map-3.8.0-29-generic  
vmlinuz-3.8.0-29-generic
```

Found that the hypervisor version we are currently using doesn't support ext4 as there are no modules available.

```
# mount -t ext4 /dev/mapper/422facae-3deb-4181-98ed-168418ce5e23p5 /mnt
mount: wrong fs type, bad option, bad superblock on /dev/mapper/422facae-3deb-
4181-98ed-168418ce5e23p5,
       missing codepage or other error
       In some cases useful info is found in syslog - try
       dmesg | tail or so
```

VM cannot be stopped unless the newly created device/partition mappings are deleted.

I found that even after running "kpartx -d", the mappings still existed, used "dmsetup remove" instead.

```
# kpartx -d /dev/mapper/422facae-3deb-4181-98ed-168418ce5e23p5
# kpartx -d /dev/mapper/422facae-3deb-4181-98ed-168418ce5e23p1
```

Partitions still exist after running "kpartx -d"

```
# ls -l /dev/mapper/ | grep 168418ce5e23
brw----- 1 root root 252, 24 Feb 28 17:16 422facae-3deb-4181-98ed-
168418ce5e23p1
brw----- 1 root root 252, 25 Feb 28 17:16 422facae-3deb-4181-98ed-
168418ce5e23p5
brw-rw---- 1 root disk 252, 23 Feb 28 17:31 VG_XenStorage--ed154d7d--8165--
957c--6e18--c9d6f79d3b11-VHD--422facae--3deb--4181--98ed--168418ce5e23
```

After running "dmsetup remove", I was able to stop the VM.

```
# dmsetup remove 422facae-3deb-4181-98ed-168418ce5e23p1
```

```
# ls -l /dev/mapper/ | grep 168418ce5e23
brw----- 1 root root 252, 25 Feb 28 17:16 422facae-3deb-4181-98ed-
168418ce5e23p5
brw-rw---- 1 root disk 252, 23 Feb 28 17:31 VG_XenStorage--ed154d7d--8165--
957c--6e18--c9d6f79d3b11-VHD--422facae--3deb--4181--98ed--168418ce5e23
```

```
# dmsetup remove 422facae-3deb-4181-98ed-168418ce5e23p5
```

After stopping the VM no device mappings were found with the matching VDI.

```
# ls -l /dev/mapper/ | grep 168418ce5e23
```