



# Oracle VM Server The Undocumentation

Robert Marz

# Robert Marz

Client

Senior Technical Architect  
with database centric view of the world

its-people

Portfolio Manager Database Technologies  
Blog Editor

DOAG

Active Member Database Community  
in charge of Cloud topics



@RobbieDatabee




blog.its-people.de



Robert.Marz  
@its-people.de



ORACLE  
ACE





# Motivation

## – Why am I here?

### Oracle VM Server

Open Source  
Underlying technology  
of Oracle Cloud  
(one out of  
several Hypervisors)

### Documentation

Describes  
mainly UX and CLI  
Information  
is hard to find



# VM Server: x86-Hypervisors

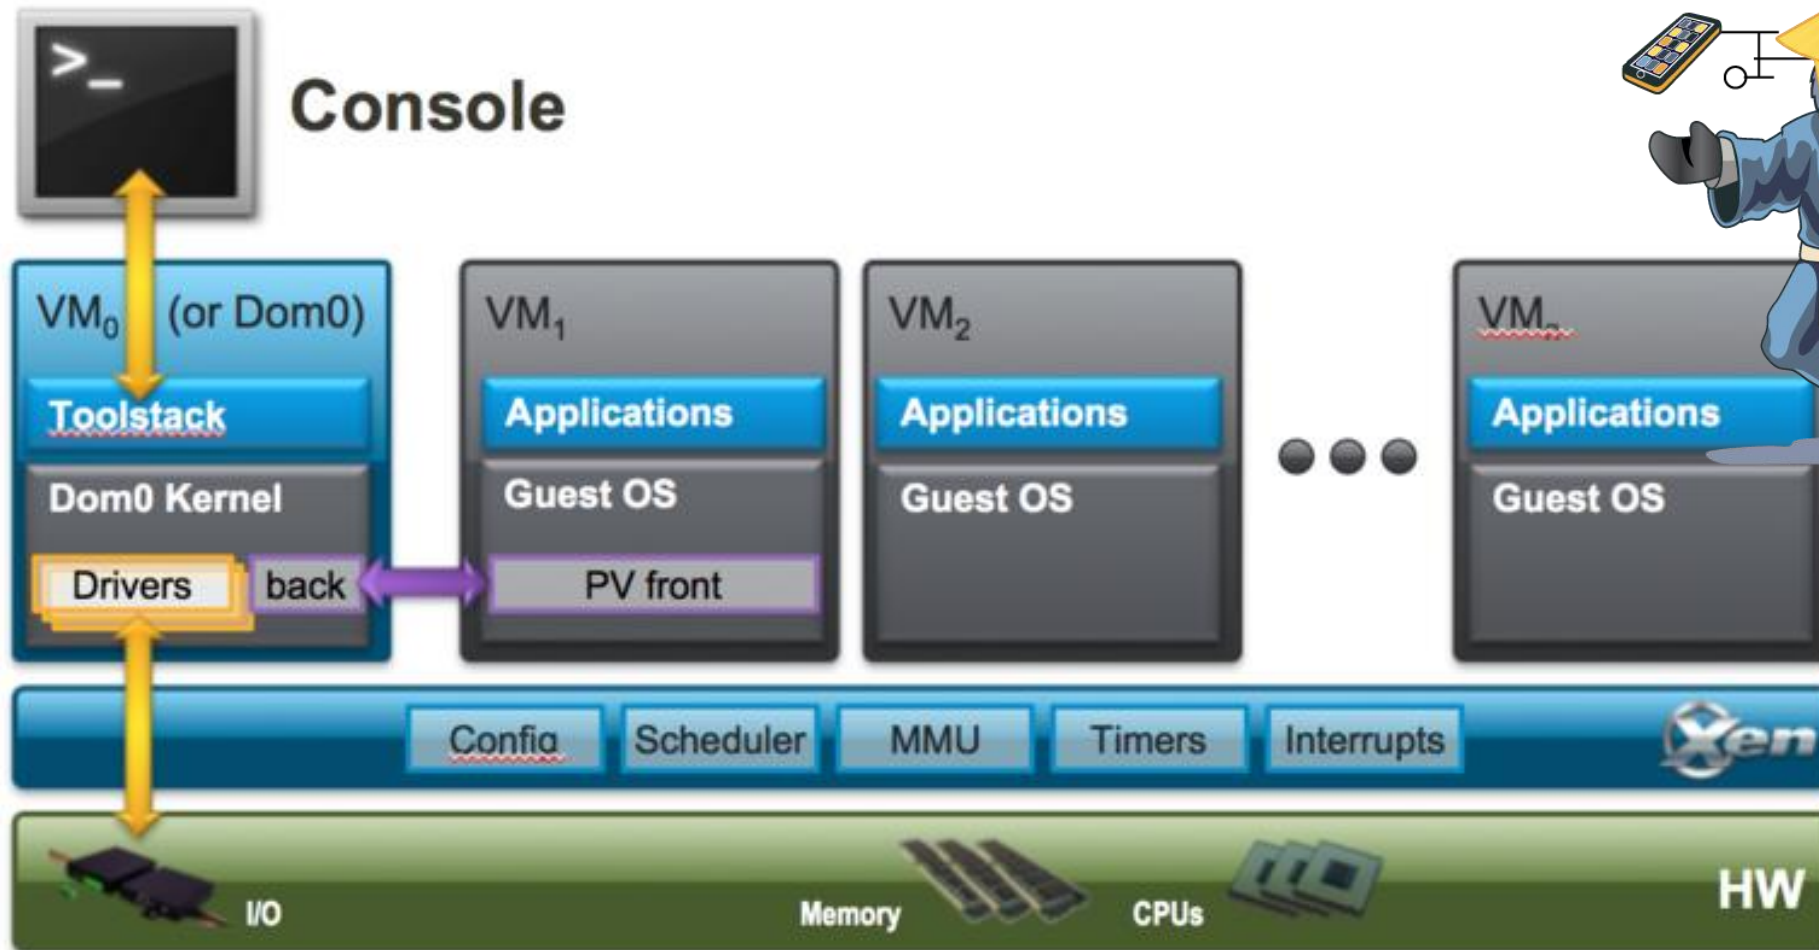


Microsoft  
Hyper-v



vmware®  
vSphere

# Xen: The Big Picture



Source: [https://wiki.xen.org/wiki/File:Xen\\_Arch\\_Diagram.png](https://wiki.xen.org/wiki/File:Xen_Arch_Diagram.png)

# Xen: Dom0 (The VM Server)

## Dom0

is a privileged Domain

This where you land  
with ssh

guest VMs

Create and Destroy

Manage

runs xend

the Xen Daemon

communicates with hypervisor

xm

Xen management tool

interact with xend

Supplies device and I/O  
services to VMs

backend device drivers

provides domain storage



# Xen: Dom0 – RAM Usage

## Dom0 RAM

RAM is NOT  
overcommittable

$\text{Mem}(\text{Host}) < \text{Mem}(\text{Dom0}) + \text{Mem}(\text{DomUs})$

Default dom0\_mem=

$\text{Mem}(\text{host}) * 0.0205 + 502$

Designed for a few big VMs

Device Drivers

Eat RAM

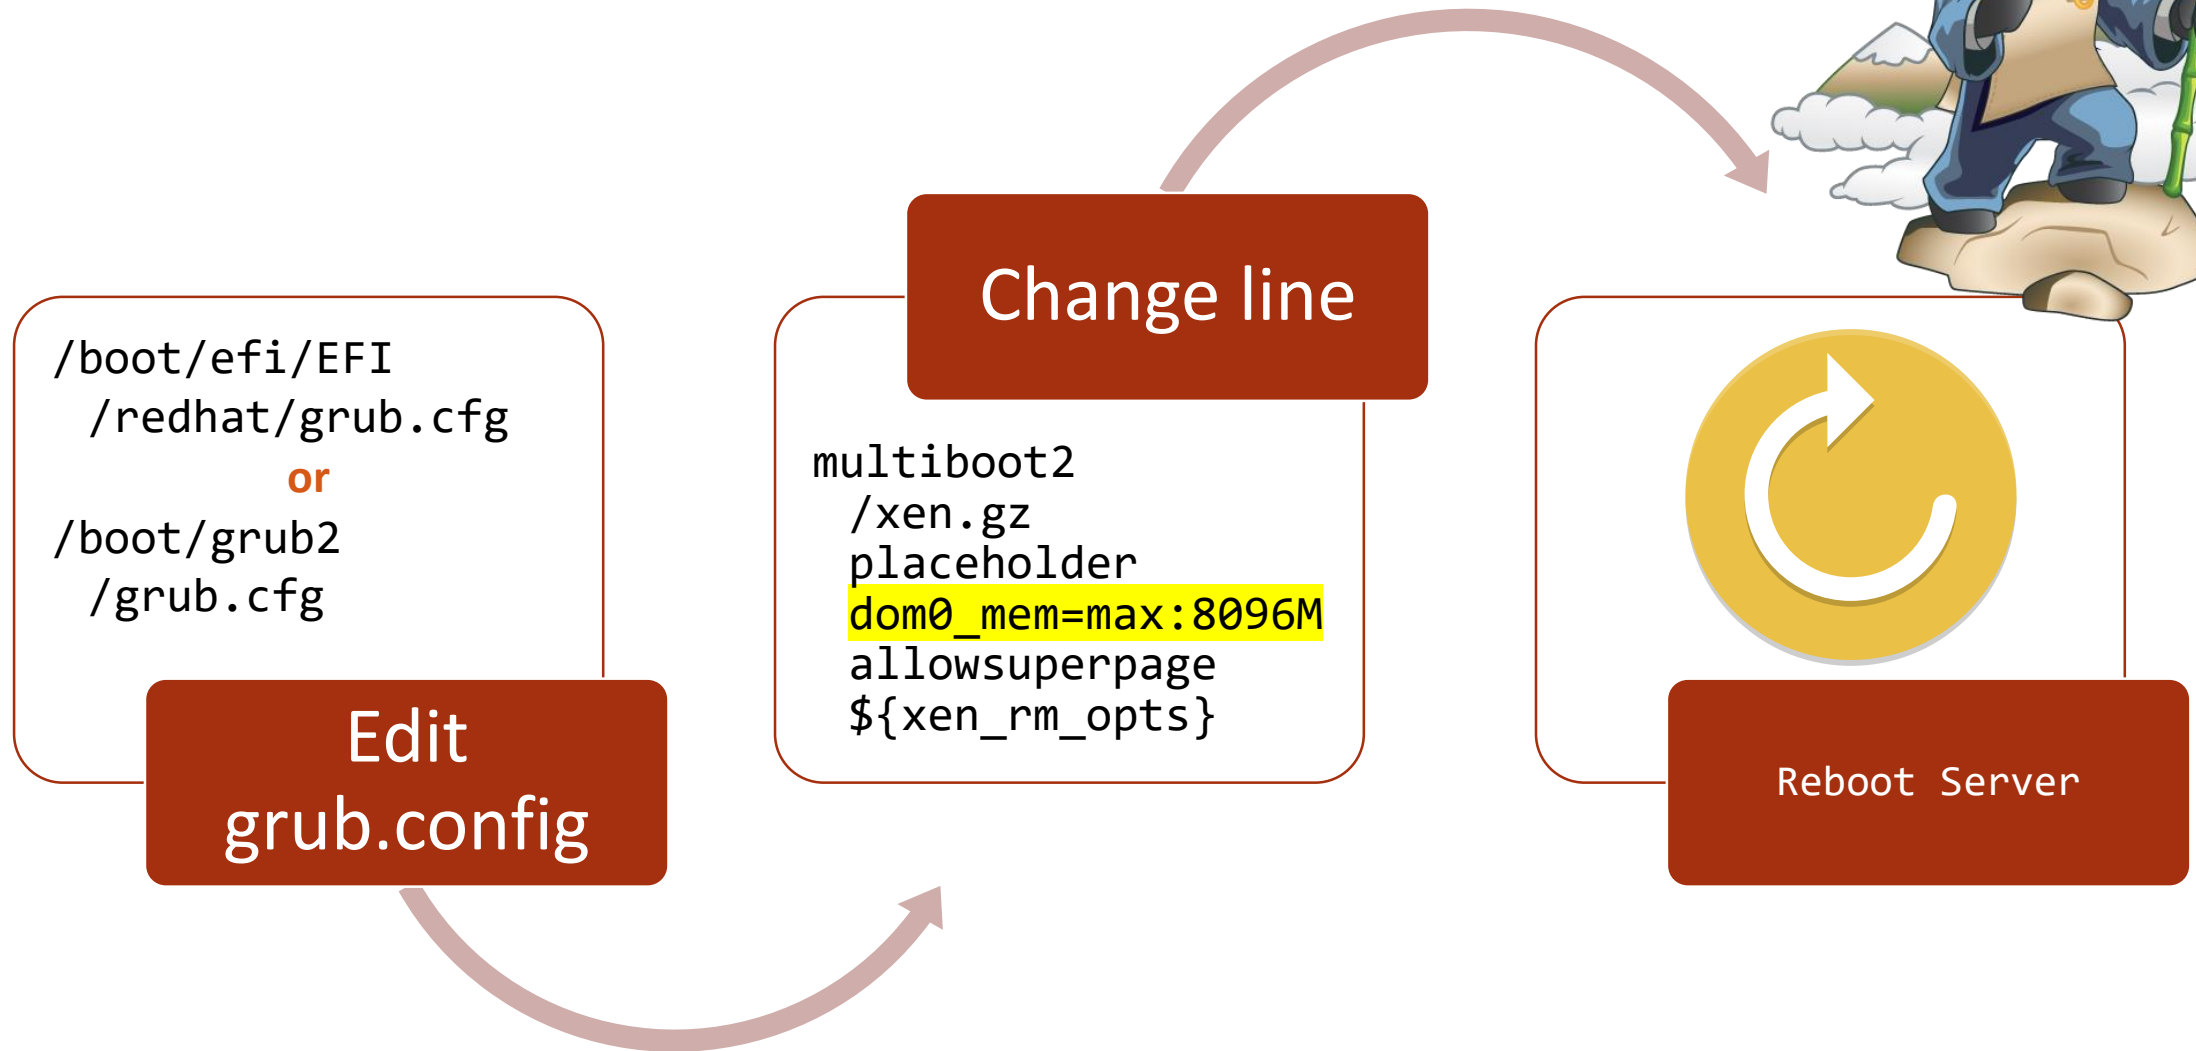
VMs won't start

if Mem(Dom0) is exhausted

weird error messages



# Xen: Dom0 – Adjusting RAM



# Xen management Tool: xm

## xm

- Create, Destroy
- Modify
- List

## OVM agent

- monitors xend Events
- Reports to OVM Manager

xm interacts with the local xend

```
[root@sv091 ~]# xm
Usage: xm <subcommand> [args]

Control, list, and manipulate Xen guest instances.

Common 'xm' commands:

console      Attach to <Domain>'s console.
vncviewer    Attach to <Domain>'s VNC server.
create       Create a domain based on <ConfigFile>.
new          Adds a domain to Xend domain management
delete       Remove a domain from Xend domain management.
destroy      Terminate a domain immediately.
dump-core    Dump core for a specific domain.
help         Display this message.
list         List information about all/some domains.
mem-set      Set the current memory usage for a domain.
migrate      Migrate a domain to another machine.
pause        Pause execution of a domain.
reboot       Reboot a domain.
reset        Reset a domain.
restore      Restore a domain from a saved state.
resume       Resume a Xend managed domain
save         Save a domain state to restore later.
shell        Launch an interactive shell.
shutdown     Shutdown a domain.
start        Start a Xend managed domain
suspend      Suspend a Xend managed domain
top          Monitor a host and the domains in real time.
unpause      Unpause a paused domain.
uptime       Print uptime for all/some domains.
usb-add      Add the usb device to FV VM.
usb-del      Delete the usb device to FV VM.
domstate    get the state of a domain
vcpu-set     Set the number of active VCPUs allowed for the
             domain.

<Domain> can either be the Domain Name or Id.
For more help on 'xm' see the xm(1) man page.
For more help on 'xm create' see the xmdomain.cfg(5) man page.

For a complete list of subcommands run 'xm help'.
```

# Xen xm Demo: List

Warning: making manual modifications in the management domain might cause inconsistencies between Oracle VM Manager and the server.

```
[root@sv091 ~]# xm list
```

Name	ID	Mem	VCPUs	State	Time(s)
0004fb00000600001ab4fe3614bd3616	1	16351	4	-b----	2551256.8
0004fb00000600004a27037b0835a489	8	16351	4	-b----	1142964.1
0004fb000006000063c616814464d177	15	16351	4	-b----	419351.5
0004fb000006000075cbdd1c01458b89	2	16351	4	-b----	2518897.3
0004fb0000060000848126d4350e937c	6	16351	4	-b----	2567024.0
0004fb0000060000a31577d2f66b7345	7	16351	4	-b----	2030372.7
0004fb0000060000cd82eeb81e5f78a1	13	16351	4	-b----	435982.3
Domain-0	0	2796	16	r-----	3867699.3


```
[root@sv091 ~]# █
```

# Xen xm Demo: Create / Destroy

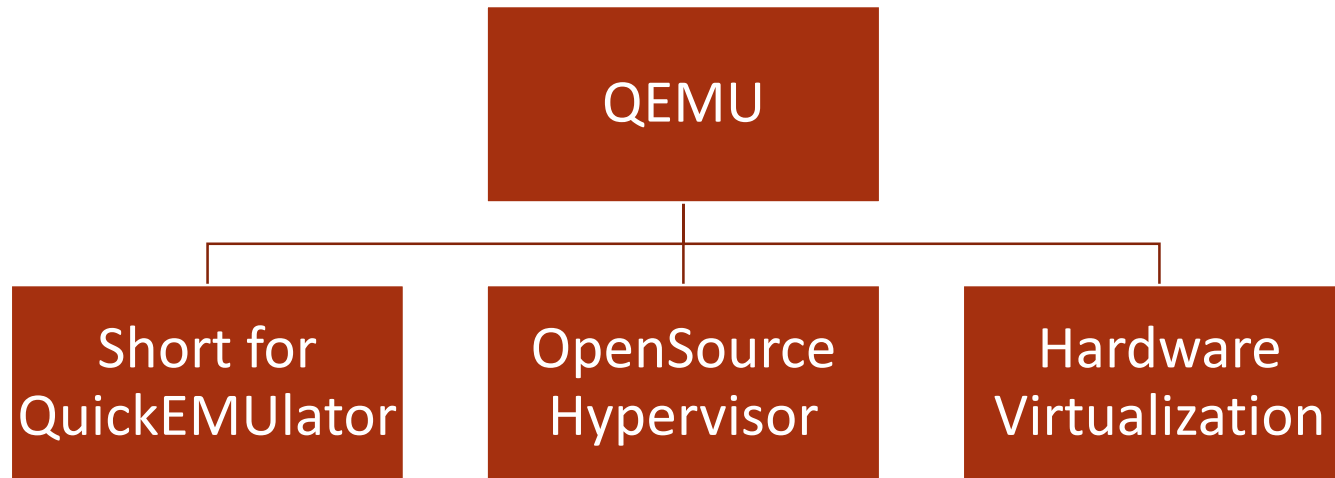
```
[root@sv091 ~]# cd /OVS/Repositories/0004fb000003000024cd350d232b41fc/VirtualMachines/
[root@sv091 VirtualMachines]# xm create 0004fb000006000063c616814464d177/vm.cfg
Using config file "./0004fb000006000063c616814464d177/vm.cfg".
Started domain 0004fb000006000063c616814464d177 (id=19)
[root@sv091 VirtualMachines]#
[root@sv091 VirtualMachines]#
[root@sv091 VirtualMachines]# xm list 19
```

Name	ID	Mem	VCPUs	State	Time(s)
0004fb000006000063c616814464d177	19	16351	1	-----	5.5

```
[root@sv091 VirtualMachines]#
[root@sv091 VirtualMachines]# xm shutdown 19
[root@sv091 VirtualMachines]#
[root@sv091 VirtualMachines]# xm list 19
Error: Domain '19' does not exist.
[root@sv091 VirtualMachines]#
```



# Xen qemu



Xen uses parts of  
QEMU Code Base

QEMU References  
everywhere

Lots of usefull  
Debug Information  
in Logfiles:

```
/var/log/xen/*log
```

```
[root@sv091 xen]# cd /var/log/xen
[root@sv091 xen]# ls -l *log
domain-builder-ng.log
qemu-dm-0004fb00000600001ab4fe3614bd3616.log
qemu-dm-0004fb00000600002a56e34a3ba76342.log
qemu-dm-0004fb00000600004a27037b0835a489.log
qemu-dm-0004fb000006000063c616814464d177.log
qemu-dm-0004fb00000600006cf0850e247a2c6b.log
qemu-dm-0004fb000006000075cbdd1c01458b89.log
qemu-dm-0004fb00000600007b0a258dbc15dedf.log
qemu-dm-0004fb0000060000848126d4350e937c.log
qemu-dm-0004fb000006000097d11afba5e0e2a4.log
qemu-dm-0004fb0000060000a31577d2f66b7345.log
qemu-dm-0004fb0000060000b2020726fa645272.log
qemu-dm-0004fb0000060000cd82eeb81e5f78a1.log
xend-debug.log
xend.log
xen-hotplug.log
[root@sv091 xen]#
```

# Xen Hypervisor: xentop

## Xentop

displays real-time information

Oracle VM Names

Xen system

Xen domains (VMs)

are too long

no chance to identify VM



```
xentop - 01:25:09 Xen 4.4.40VM
7 domains: 7 running, 0 blocked, 0 paused, 0 crashed, 0 dying, 0 shutdown
Mem: 134089040k total, 105872228k used, 28216812k free CPUs: 16 @ 1698MHz
```

NAME	STATE	CPU(sec)	CPU(%)	MEM(k)	MEM(%)	MAXMEM(k)	MAXMEM(%)	VCPUS	NETS	NETTX(k)	NETRX(k)	VBDS	VBD_00	VBD_RD	VBD_WR	VBD_RSECT	VBD_WSECT	SSID
0004fb0000	-----r	2576284	101.9	16744412	12.5	16744448	12.5	4	2	6706307	6246144	6	7	166372565	14152398	20528968189	703751050	0
0004fb0000	-----r	1161531	104.8	16744412	12.5	16744448	12.5	4	2	4388360	2534062	6	4	70259624	5899554	2637507210	149089976	0
0004fb0000	-----r	2536826	104.8	16744412	12.5	16744448	12.5	4	2	5305504	865909	6	2	9549958	6180375	289616933	142755219	0
0004fb0000	-----r	2593169	104.8	16744412	12.5	16744448	12.5	4	2	5208046	947980	6	10	112831934	13913929	2435999873	446094680	0
0004fb0000	-----r	2045759	101.7	16744412	12.5	16744448	12.5	4	2	3456593	3571243	6	1	15131845	10693713	490624383	240997839	0
0004fb0000	-----r	453116	105.0	16744412	12.5	16744448	12.5	4	2	1368957	163406	6	1	7058671	2720478	569534178	96903270	0
Domain-0	-----r	3911056	80.3	2863220	2.1	3538944	2.6	16	0	0	0	0	0	0	0	0	0	0

# Xen Hypervisor: xentop (cont.)

## Xentop Options

--full-name (-f)

Shows full Domain Name

destroys heading alignment

--networks (-n)

vif network data

--vbds (-x)

Block device data

--vcpus (-v)

vCPU data

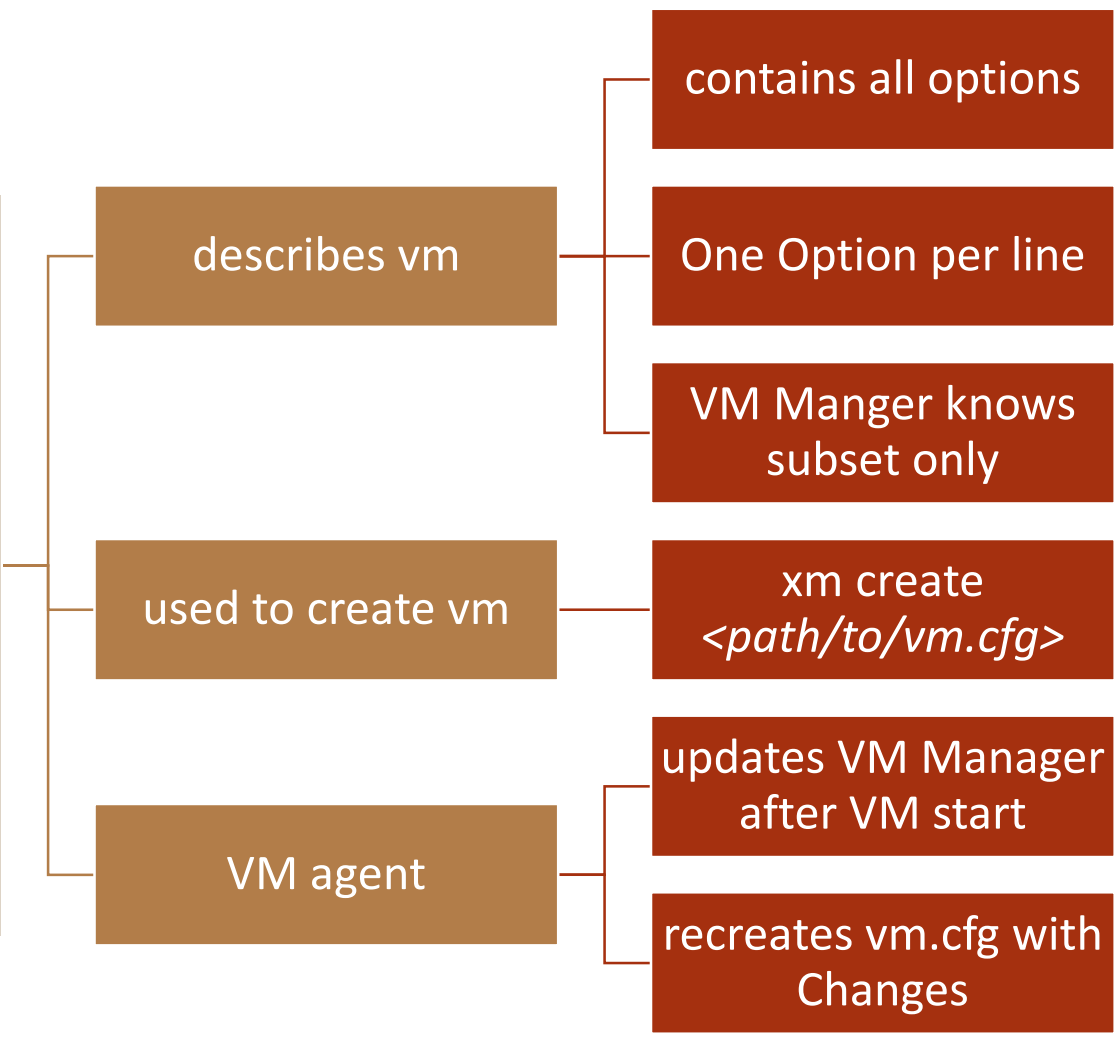


```
xentop - 01:31:28 Xen 4.4.40VM
7 domains: 7 running, 0 blocked, 0 paused, 0 crashed, 0 dying, 0 shutdown
xentop - 01:27:50 Xen 4.4.40VM
7 domains: 7 running, 0 blocked, 0 paused, 0 crashed, 0 dying, 0 shutdown
Mem: 134089040k total, 105872228k used, 28216812k free CPUs: 16 @ 1698MHz
```

NAME	STATE	CPU(sec)	CPU(%)	MEM(k)	MEM(%)	MAXMEM(k)	MAXMEM(%)	VCPUS	NETS	NETTX(k)	NETRX(k)	VBDS	VBD_00	VBD_RD	VBD_WR	VBD_RSECT	VBD_WSECT	SSID		
0004fb00000600001ab4fe3614bd3616	-----r	2576452	107.2	16744412	12.5	16744448	12.5	16744448	12.5	4	2	6706547	6246264	6	7	166375737	14152671	20529029005	703762135	0
0004fb000006000004a27037b0835a489	-----r	1161702	106.8	16744412	12.5	16744448	12.5	16744448	12.5	4	2	4388565	2534142	6	4	70260141	5899767	2637523706	149094454	0
0004fb0000060000075cbdd1c01458b89	-----r	2536996	126.7	16744412	12.5	16744448	12.5	16744448	12.5	4	2	5305647	865921	6	2	9550296	6180591	289628309	142759505	0
0004fb00000600000848126d4350e937c	-----r	2593335	102.2	16744412	12.5	16744448	12.5	16744448	12.5	4	2	5208174	947980	6	10	112832471	13914117	2436016609	446099042	0
0004fb00000600000a31577d2f66b7345	-----r	2045924	102.8	16744412	12.5	16744448	12.5	16744448	12.5	4	2	3456722	3571243	6	1	15132173	10693875	490634847	241001755	0
0004fb00000600000cd82eeb81e5f78a1	-----r	453286	105.2	16744412	12.5	16744448	12.5	16744448	12.5	4	2	1369085	163406	6	1	7058985	2720665	569544226	96907732	0
Domain-0	-----r	3911141	52.6	2863220	2.1	3538944	2.6	16	0	0	0	0	0	0	0	0	0	0	0	0
Net0 RX: 307652896bytes		3202024pkts																		
0004fb0000	-----r	2593561	101.8	16744412	12.5	16744448	12.5	16744448	12.5	4	2	5208347	947980	6	10	112834542	13914532	2436098449	446106969	0
Net0 RX: 932617850bytes		3291139pkts																		
Net1 RX: 38113951bytes		101076pkts																		
0004fb0000	-----r	2046146	101.8	16744412	12.5	16744448	12.5	16744448	12.5	4	2	3456895	3571243	6	1	15132677	10694137	490650943	241007813	0
Net0 RX: 3638705110bytes		1814499pkts																		
Net1 RX: 18248286bytes		46949pkts																		
0004fb0000	-----r	453516	104.7	16744412	12.5	16744448	12.5	16744448	12.5	4	2	1369258	163406	6	1	7059496	2721014	569560498	96914978	0
Net0 RX: 9600528bytes		51672pkts																		
Net1 RX: 157727241bytes		482514pkts																		
Domain-0	-----r	3911257	54.2	2863220	2.1	3538944	2.6	16	0	0	0	0	0	0	0	0	0	0	0	0

# Xen Hypervisor: vm.cfg

vm.cfg

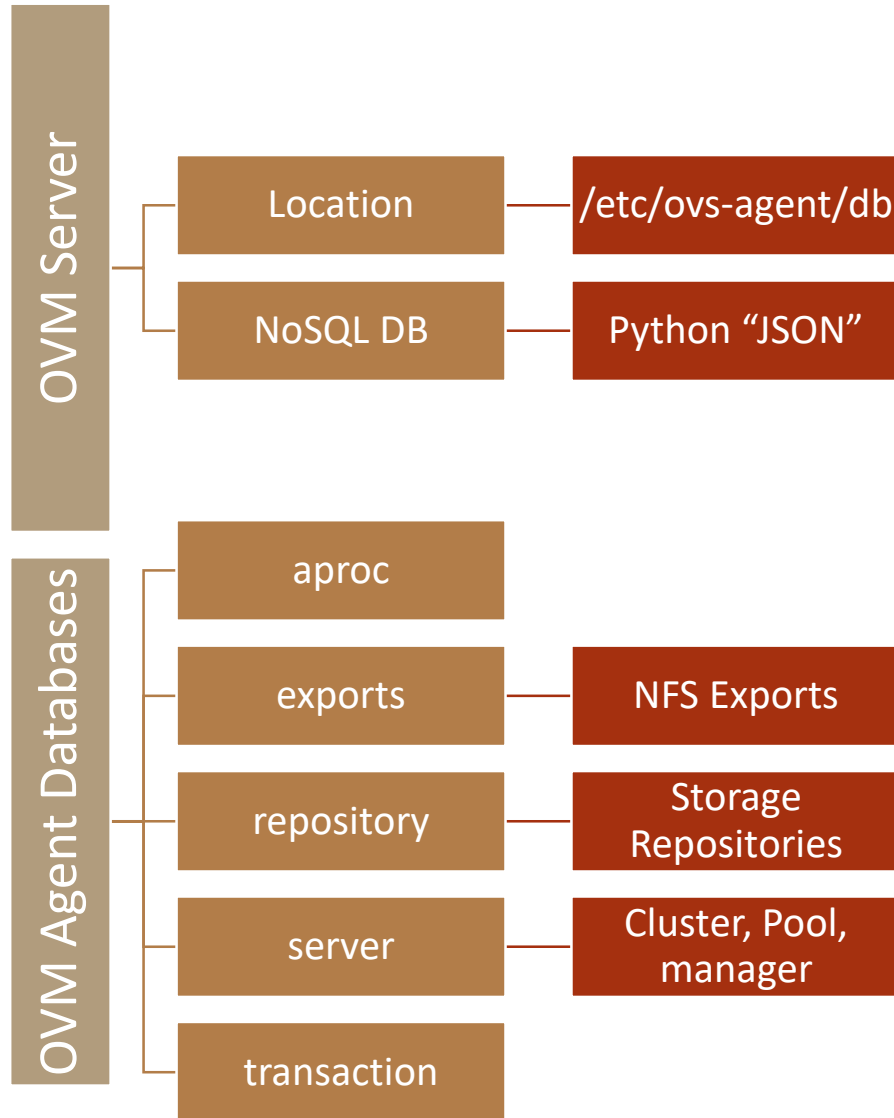


```

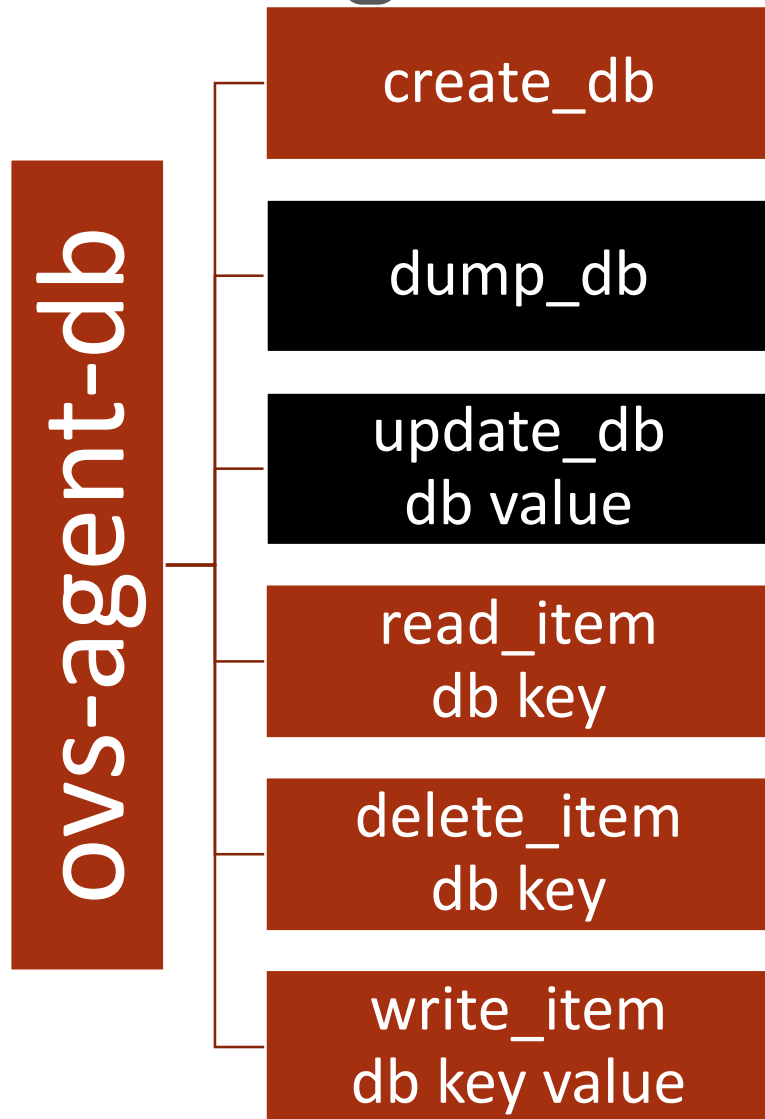
vif = ['mac=00:21:f6:2a:00:68,bridge=172.24.19.0', '
      mac=00:21:f6:2a:01:68,bridge=100bc68c90']
OVM_simple_name = 'sv104'
vnclisten = '127.0.0.1'
serial = 'pty'
disk = ['file:/OVS/Repositories/0004fb000003000024cd350d232b41fc/VirtualDisks/
0004fb0000120000b4aa1d5fea4c543a.img,xvda,w', 'file:/OVS/Repositories/
0004fb000003000024cd350d232b41fc/VirtualDisks/
0004fb0000120000a0af3357d6cafd05.img,xvdb,w', 'file:/OVS/Repositories/
0004fb000003000024cd350d232b41fc/VirtualDisks/
0004fb00001200005428824c138a9ebf.img,xvdc,w', 'file:/OVS/Repositories/
0004fb000003000024cd350d232b41fc/VirtualDisks/
0004fb0000120000b171b415380cb844.img,xvdd,w', 'file:/OVS/Repositories/
0004fb000003000024cd350d232b41fc/VirtualDisks/
0004fb000012000097776037e15d0175.img,xvde,w', 'file:/OVS/Repositories/
0004fb000003000024cd350d232b41fc/VirtualDisks/
0004fb0000120000538b8ba47b973b25.img,xvdf,w']
vncunused = 1
uuid = '0004fb00-0006-0000-63c6-16814464d177'
on_reboot = 'restart'
boot = 'cd'
name = '0004fb000006000063c616814464d177'
cpu_weight = 27500
memory = 16348
cpu_cap = 0
maxvcpus = 4
OVM_high_availability = False
vnc = 1
OVM_description = ''
on_poweroff = 'destroy'
on_crash = 'restart'
guest_os_type = 'linux'
usbdevice = 'tablet'
builder = 'hvm'
vcpus = 4
keymap = 'de'
OVM_os_type = 'Oracle Linux 7'
OVM_cpu_compat_group = ''
OVM_domain_type = 'xen_hvm_pv'
  
```



# The OVM Agent Databases



# OVM Agent Database Manipulation



Use python syntax to specify the "value" parameters:

- \* None: "None"
- \* string: "'some string'"
- \* number: "1234.5678", "0x1234"
- \* boolean: "True", "False"
- \* list: "['foo', 1234]"
- \* tuple: "('foo', 1234)"
- \* dict: "{'foo': 'bar', 1: 2}"

# OVM Agent Database: server

```
{'cluster_state': 'DLM_Ready',
'clustered': True,
'fs_stat_uuid_list': ['0004fb000005000004b84b442df68cb2'],
'is_master': False,
'manager_event_url': 'https://172.24.19.254:7002/ovm/core/wsapi/rest/internal/Server/01:c7:3f:76:8e',
'manager_ip': '172.24.19.254',
'manager_statistic_url': 'https://172.24.19.254:7002/ovm/core/wsapi/rest/internal/Server/01:c7:3f:7',
'manager_uuid': '0004fb0000010000081397618a5f4bb8',
'node_number': 3,
'pool_alias': 'I.NPP 42 ServerPool sv080',
'pool_master_ip': '172.24.19.95',
'pool_member_ip_list': ['172.24.19.91',
                        '172.24.19.94',
                        '172.24.19.92',
                        '172.24.19.95',
                        '172.24.19.96'],
'pool_uuid': '0004fb000002000031fb31cdfa8a3db5',
'poolfs_nfsbase_uuid': '',
'poolfs_target': '/dev/mapper/36e843b66f5d2554d35f7d4504d9adbdd',
'poolfs_type': 'lun',
'poolfs_uuid': '0004fb0000050000e456be54b953205c',
'registered_hostname': 'sv095',
'registered_ip': '172.24.19.95',
'roles': set(['utility', 'xen'])}
```

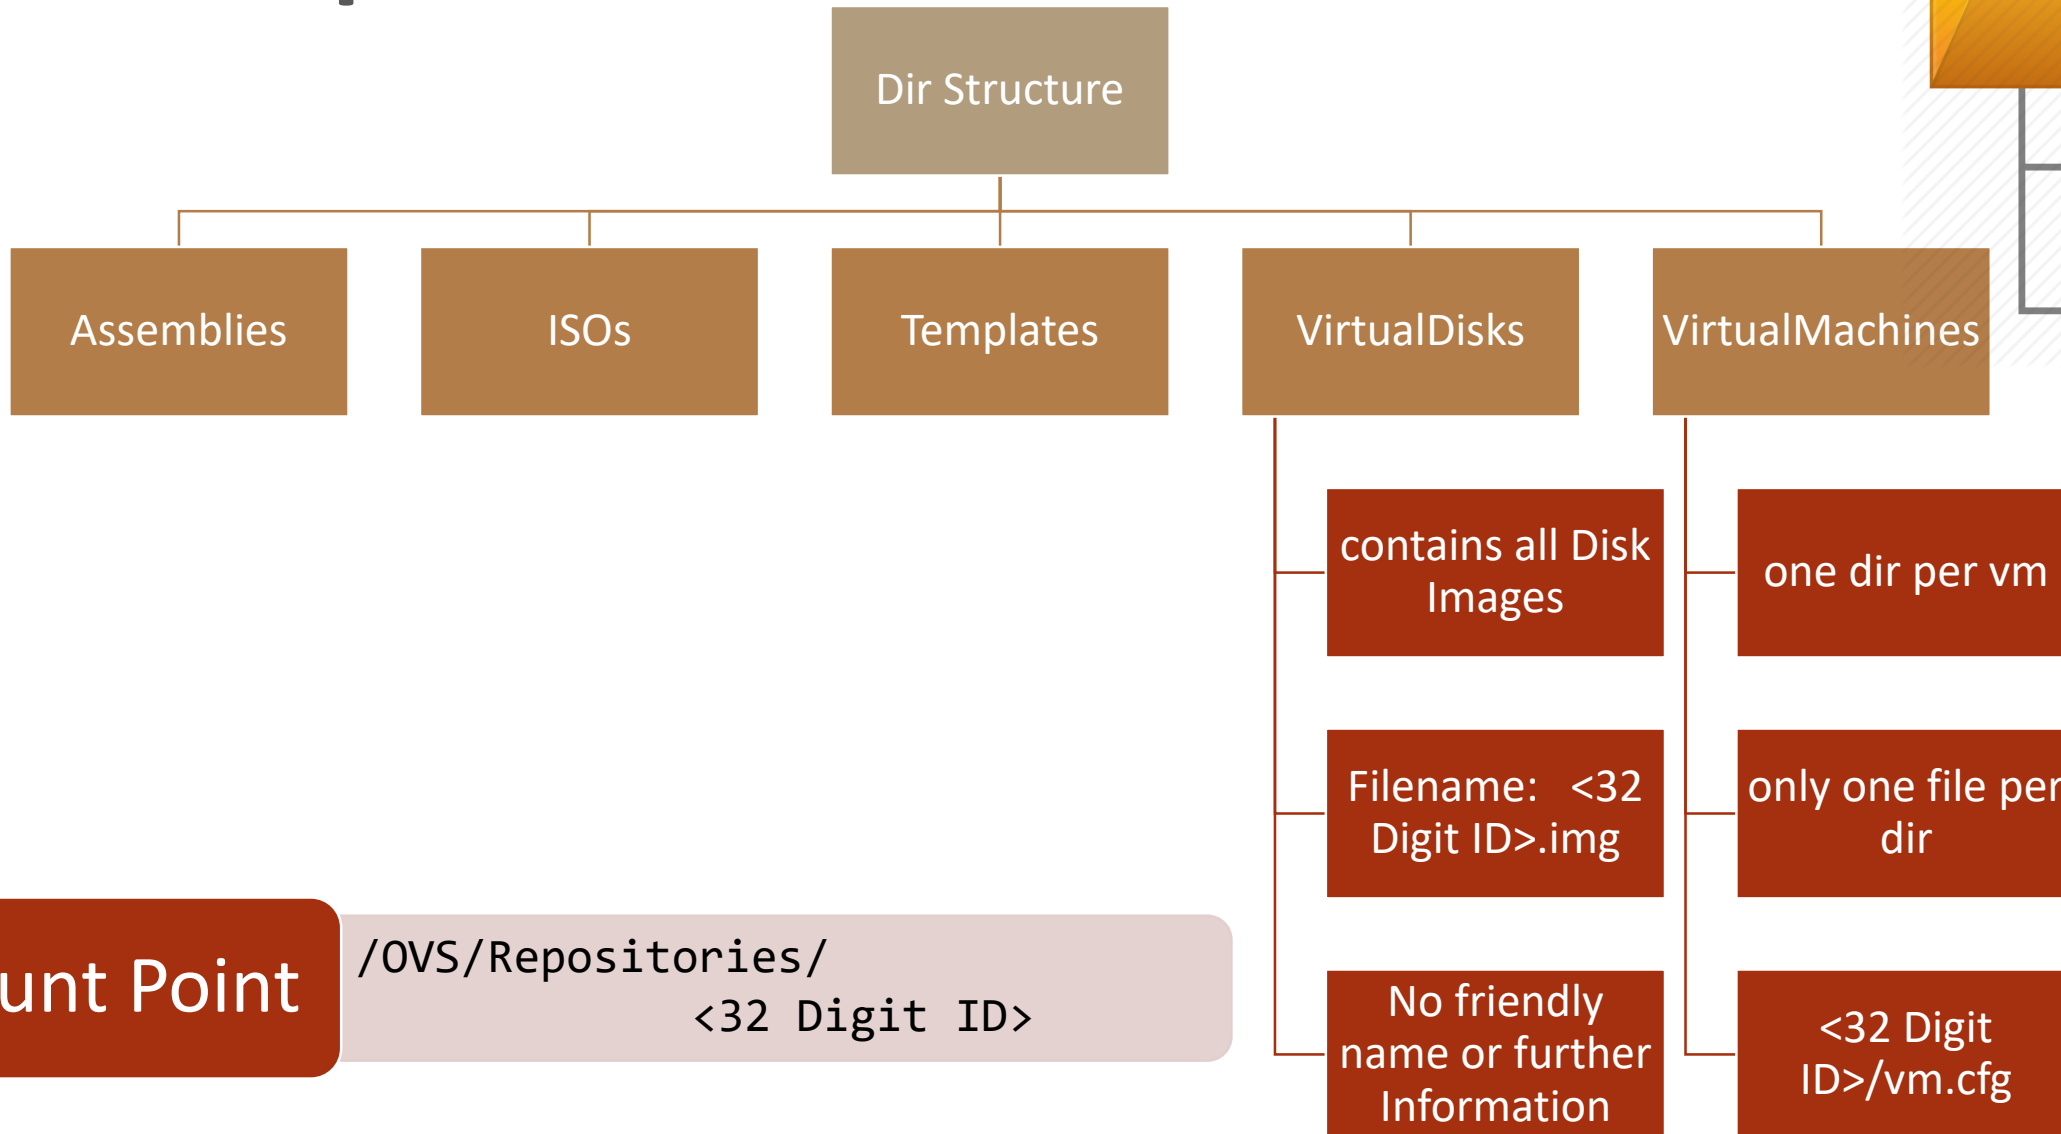
```
[root@sv095 db]# ovs-agent-db dump_db server
```

# OVM Agent Database: repository

```
{'0004fb000003000024cd350d232b41fc': {'alias': u'nas-088',  
    'filesystem': 'ocfs2',  
    'fs_location': '/dev/mapper/36e843b67fd0b542d4551d417dd8bdad8',  
    'manager_uuid': u'0004fb0000010000081397618a5f4bb8',  
    'mount_point': '/OVS/Repositories/0004fb000003000024cd350d232b41fc',  
    'version': u'3.0'}},  
'0004fb00000300003fbb494d1ddbff7d': {'alias': 'sv095-repository',  
    'filesystem': 'ocfs2',  
    'fs_location': '/dev/mapper/3600508b1001cde2abd0e04515e3e2b34',  
    'manager_uuid': '0004fb0000010000081397618a5f4bb8',  
    'mount_point': '/OVS/Repositories/0004fb00000300003fbb494d1ddbff7d',  
    'version': u'3.0'}}}
```

```
[root@sv095 db]# ovs-agent-db dump_db repository
```

# File Repositories: Structure



**Mount Point** /OVS/Repositories/  
<32 Digit ID>

# File Repositories: Metadata

## Repository Metadata

Repository- Root	.ovsrepo	Repo-ID Manager-ID friendly name
	.ovsmeta	JSON-File Virtual Disks Info New with OVM 3.4
ovs-agent- db	Repositories	Update repo & .ovsrepo together



# File Repository Metadata: .ovsrepo

```
[root@sv095 0004fb000003000024cd350d232b41fc]# cat .ovsrepo
OVS_REPO_UUID=0004fb000003000024cd350d232b41fc
OVS_REPO_VERSION=3.0
OVS_REPO_MGR_UUID=0004fb0000010000081397618a5f4bb8
OVS_REPO_ALIAS=nas088-repository
```

# File Repository Metadata: .ovsmeta

```
{ "0004fb00001200009c02221b285ad76d.img": {
  "OBJECT_TYPE": "VIRTUAL_DISK",
  "SIMPLE_NAME": "template-oracle-12.1-2018KW06"},
"sv254-ovmm-01.img": {
  "OBJECT_TYPE": "VIRTUAL_DISK",
  "SIMPLE_NAME": "sv254-ovmm-01.img"},
"0004fb0000120000538b8ba47b973b25.img": {
  "OBJECT_TYPE": "VIRTUAL_DISK",
  "SIMPLE_NAME": "sv104-stredax-db-2018KW20"},
"0004fb00001200003a0a9885c3bf95a1.img": {
  "OBJECT_TYPE": "VIRTUAL_DISK",
  "SIMPLE_NAME": "sv100-oracle-12.1-2017KW51"},
"0004fb000012000096097eddef0e4270.img": {
  "OBJECT_TYPE": "VIRTUAL_DISK",
  "SIMPLE_NAME": "sv103-system-2017KW51"},
"0004fb0000120000de7cc7b37bdde7cd.img": {
  "OBJECT_TYPE": "VIRTUAL_DISK",
```

# Enlarging a Repository on a shared LUN

## Step 1:

- physically extend your LUN

## Step 2:

- reread partitiontable
  - on every server
  - **multipath -ll /dev/mapper/36e843b67fd0b542d4551d417dd8bdad8**
  - or every multipath blockdev:
    - **blockdev -rereadpt /dev/sdc**

## Step 3:

- refresh multipathd
  - on every server
  - **multipathd -k"resize map 36e843b67fd0b542d4551d417dd8bdad8"**

## Step 4:

- extend OCFS2 Filesystem
  - on one server only
  - **tunefs.ocfs -S /dev/mapper/36e843b67fd0b542d4551d417dd8bdad8**

# Enlarging a Repository on a shared LUN V2

## Step 1:

- physically extend your LUN

## Step 2 (since VM Manager 3.1):

- Open VM Manager
- In the „Storage“-Tab Refresh LUN

## Step 3:

- In the „Repository“-Tab
- Refresh Repository

# Conclusion

- Configuring the VM Server directly gives great power
- There is always something new to learn
- Not every Operation is supported

**PLEASE**

**DO  
TRY THIS  
AT HOME**

**On a Testing Machine**



# Herzlichen Dank für Ihre Aufmerksamkeit

we make the difference  
www.its-people.de

## Questions ?



**its-people GmbH**

Frankfurt  
Hamburg  
Köln  
München

Tel. 069 2475 2100  
Tel. 040 2360 8808  
Tel. 0221 1602 5204  
Tel. 089 5484 2401

**its-people ERP Beratungsgesellschaft mbH**

Frankfurt

Tel. 069 2475 1980

[www.its-people.de](http://www.its-people.de) [info@its-people.de](mailto:info@its-people.de)