

WELCOME

Backup of OracleVM

Martin Bracher

DOAG – 15. November 2016

BASEL BERN LAUSANNE ZÜRICH DÜSSELDORF FRANKFURT A.M. FREIBURG I.BR. HAMBURG MÜNCHEN STUTTGART WIEN

1

2016 © Trivadis

Backup of OracleVM
November 2016

trivadis
makes IT easier. ■ ■ ■

Our company.

Trivadis is a **market leader in IT consulting, system integration, solution engineering**

and the provision of **IT services** focusing on **ORACLE®** and  **Microsoft** technologies

in Switzerland, Germany, Austria and Denmark. We offer our services in the following strategic business fields:



Trivadis Services takes over the interacting operation of your IT systems.

With over 600 specialists and IT experts in your region.



- 14 Trivadis branches and more than 600 employees
- 200 Service Level Agreements
- Over 4,000 training participants
- Research and development budget: CHF 5.0 million
- Financially self-supporting and sustainably profitable
- Experience from more than 1,900 projects per year at over 800 customers

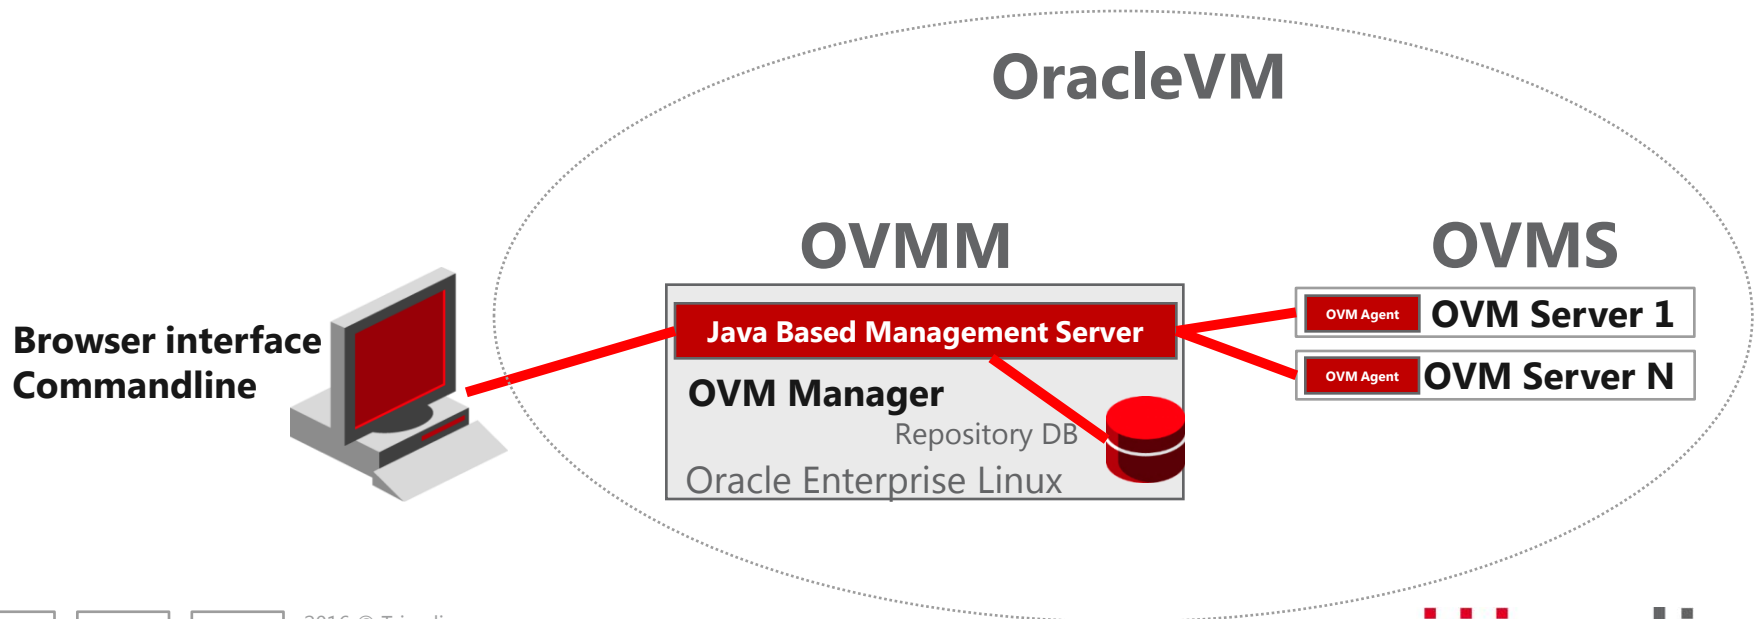
AGENDA – Backup and Recovery of OVM

- 1. Overview of components**
2. Backup
3. Recovery scenarios

Oracle VM infrastructure components

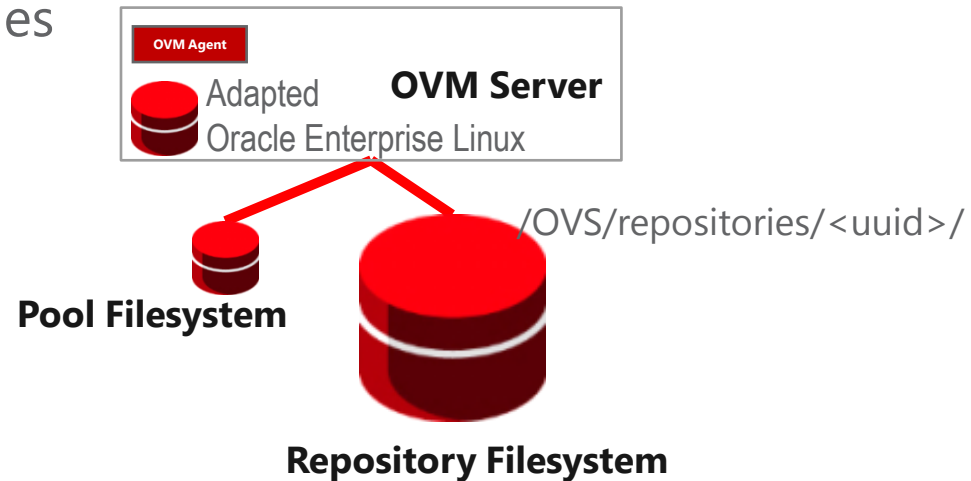
Architecture: Two Components

- 1..n Oracle VM Server (OVMS)
 - Host for the virtual machines
- 1 Oracle VM Manager (OVMM)
 - Management infrastructure for OVS (Repository-DB, Web-frontend, CLI)




Oracle VM infrastructure components: OVMS

- OVMS: Host for the virtual machines
 - Local disk
 - Operating System
 - Pool Filesystem (SAN/iSCSI/NFS)
 - Information shared in a cluster
 - Handling of node membership
 - Repository Filesystem (SAN/iSCSI/NFS)
 - Disk-files of virtual machines
 - ISO-images
 - Definition files of VM's



Oracle VM infrastructure components: OVMM

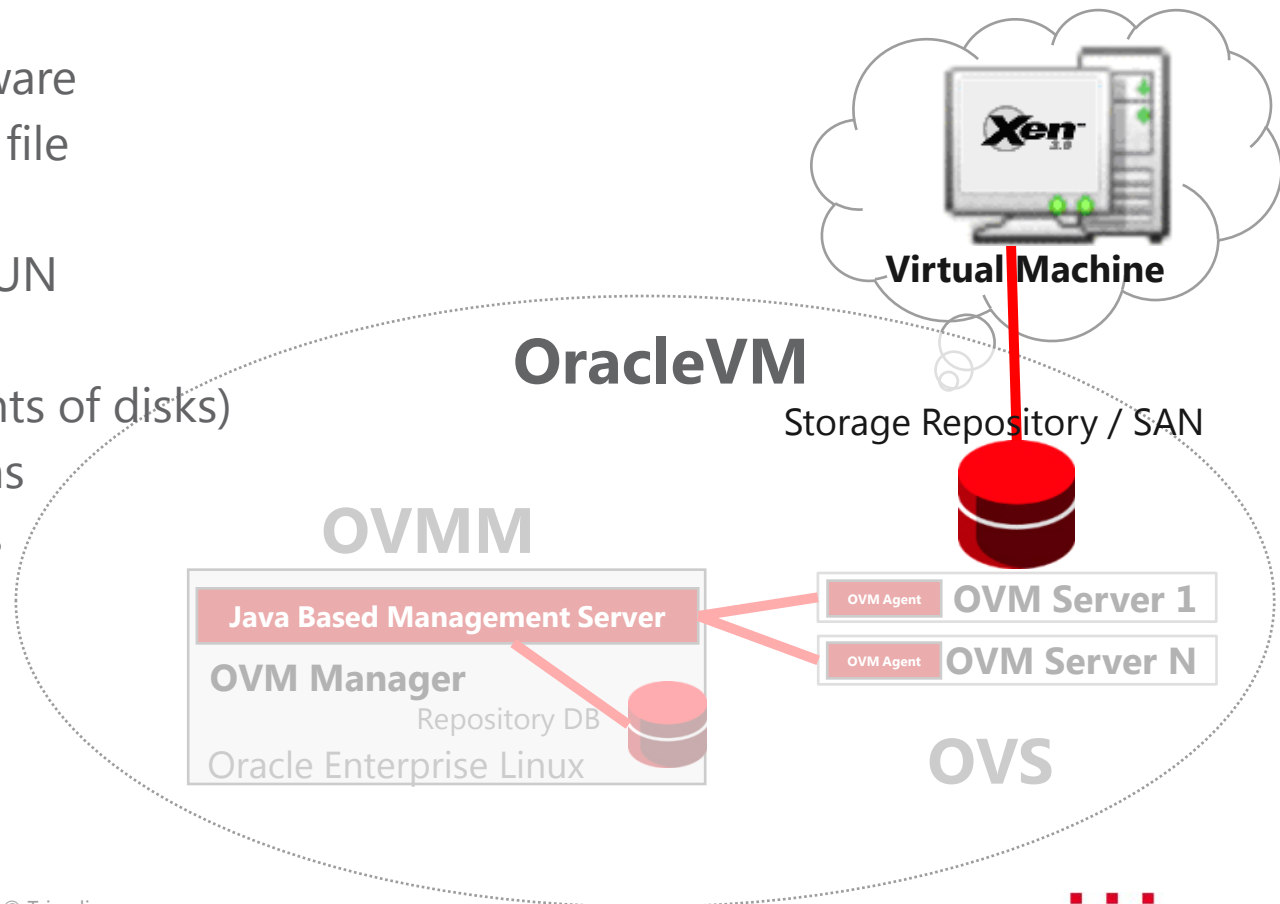
- Oracle VM Manager (OVMM)
 - Oracle/RHEL Enterprise Linux Server
 - Local disks (or SAN)
 - Operating system
 - Software OVMM (Weblogic)
 - Repository database (MySQL)



Oracle VM: components: VM

■ Virtual Machines

- Virtual hardware
 - Definition file
- Disks
 - Physical LUN
 - Diskfile
- Data (Contents of disks)
 - Filesystems
 - Databases



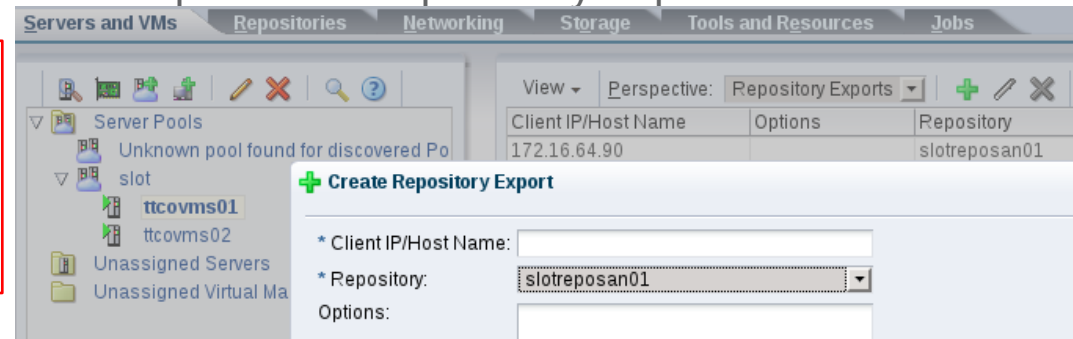
AGENDA

1. Overview of components
- 2. Backup**
3. Recovery scenarios

Backup on OVM Servers

- Not supported to install additional software
 - No backup clients
- Use of default OS tools is OK
 - tar, gzip, cp, scp, ...
- Storage Repositories can be exported by NFS
 - Servers and VMs → servername → Perspective: "Repository Exports"

```
create RepositoryExport
clientHostName=myBackupHost
name="MyExport"
repository=slotreposan01
options="rw, async, no_root_squash"
on Server name=ttcovms01
```




- Put the backup on the exported repository and backup the exported directories via another server (can be a VM)

OVM Server: Operating system

- Create the OS backup with native tools, store it on repo-filesystem

```
tar --one-file-system -zcvf /OVS/repository/<uuid>/root_${HOSTNAME}.tgz /  
tar -zcvf /OVS/repository/<uuid>/boot_${HOSTNAME}.tgz /boot  
dd if=/dev/sda of=/OVS/repository/<uuid>/mbr.raw bs=512 count=1
```

- No crontab on server, run it from remote via ssh (e.g. OVM Manager)
 - Size of such a backup is usually < 1GB
 - Backup it via NFS-Mount or ssh from another server
-
- Alternatively:
 - Do not backup; in case of problem:
 - Delete server from repository
 - Re-install server and add it to repository



OVM Server: Pool Filesystem



Pool Filesystem

- OCFS2 filesystem, acting as a Cluster Registry and Voting Disk
 - On NFS, it is a 12GB file, mounted loopback
- Variant 1: Initially create a backup of the device
 - Stop the cluster
 - SAN/iSCSI: `dd if=<device> bs=1M | gzip - > poolfs.gz`
 - NFS: copy the image file
- Variant 2: Get information required to re-create it (tunefs command on 1 line!)

Pool Filesystem

Repository Filesystem

```
Dev=`mount |grep poolfsmnt|awk '{print $1}'`  
tunefs.ocfs2 -Q "mkfs.ocfs2 -b %B -C %T -J size=4M -N %N -L %V -U %U  
--global-heartbeat --cluster-stack=o2cb  
--cluster-name=`grep cluster /etc/ocfs2/cluster.conf |head -n 1 |awk '{print $3}'`  
--force $Dev\n" $Dev
```

- Backup the contents of the filesystem regularly (with tar/cp/...)

OVM Server: Pool Filesystem - Dependencies

```
mkfs.ocfs2 -b 4096 -C 4096 -J size=4M -N 32 -L OVS_POOL_FILESYSTEM \  
-U 0004fb0000050000e6a66cb83ca0c328 --global-heartbeat \  
--cluster-stack=o2cb \  
--cluster-name=e778eb160b9eab47 \  
--force /dev/mapper/360002ac000000000000000760000d140 ;
```


```
# multipath -ll  
360002ac000000000000000760000d140 dm-2 3PARdata,VV  
size=13G features='1 queue_if_no_path' hwhandler='0' wp=rw ...
```

```
# strings /poolfsmnt/0004fb0000050000e6a66cb83ca0c328/db/server_pool | grep "U "  
U 0004fb0000020000e778eb160b9eab47q
```

```
# df -h |grep poolfsmnt  
/dev/mapper/360002ac000000000000000760000d140          13G  263M   13G   2%  
/poolfsmnt/0004fb0000050000e6a66cb83ca0c328
```

```
# cat /poolfsmnt/0004fb0000050000e6a66cb83ca0c328/.ovspoolfs  
OVS_POOLFS_UID=0004fb0000050000e6a66cb83ca0c328  
OVS_POOLFS_MGR_UID=0004fb000001000079a86253a43b91fb  
OVS_POOLFS_VERSION=3.0  
OVS_POOLFS_POOL_UID=0004fb0000020000e778eb160b9eab47  
OVS_POOLFS_LUN_UID=360002ac000000000000000760000d140
```

OVM Server: Repository Filesystem



Repository Filesystem

- Stores diskfiles and definition files of VM's
- Initially backup the whole filesystem
 - .ovsrepo file and directory structure
- Get ocfs2 information

```
for Dev in `mount |grep /OVS/Repositories/ |awk '{print $1}'`; do
  printf "# " && mount |grep $Dev
  tunefs.ocfs2 \
  -Q "mkfs.ocfs2 -J block64 -b %B -L %V -U %U -T vmstore -N %N $Dev\n" $Dev
done
```

- OCFS2 Can be re-created with this info, then restore the backups on it

```
mkfs.ocfs2 -J block64 -b 4096 -L OVS56e8973883d87 \  
-U 0004fb00000500002dc56e8973883d87 -T vmstore -N 32 /dev/mapper/1IET_00010002
```

OVM Server: Repository Filesystem - Dependencies

```
mkfs.ocfs2 -J block64 -b 4096 \  
-L OVS48c3355bb2934 \  
-U 0004FB0000050000B4C48C3355BB2934 -T vmstore -N 32  
/dev/mapper/360002ac000000000000000740000d140;
```

```
# multipath -ll  
360002ac000000000000000740000d140 dm-1 3PARdata,VV  
size=5.0T features='1 queue_if_no_path' hwhandler='0' wp=rw ...
```

```
/etc/ovs-agent/db # strings repository  
mount_pointq  
U2/OVS/Repositories/0004fb0000030000c43cdd00584279eaq  
...  
fs_locationq  
U-/dev/mapper/360002ac000000000000000740000d140q  
0004fb0000030000c43cdd00584279ea
```

```
# df -h |grep Repo  
/dev/mapper/360002ac000000000000000740000d140 5.0T 3.4T 1.7T 68%  
/OVS/Repositories/0004fb0000030000c43cdd00584279ea
```

```
# cat .ovsrepo  
OVS_REPO_UUID=0004fb0000030000c43cdd00584279ea  
OVS_REPO_VERSION=3.0  
OVS_REPO_MGR_UUID=0004fb000001000079a86253a43b91fb  
OVS_REPO_ALIAS=tr_repo01
```

OVM Server: Repository Filesystem – repeating tasks

- Backup the definition files in these directories
 - VirtualMachines
 - Templates
- Backup ISOs, Assemblies
 - if necessary
- Backup of VirtualDisks directory
 - **Do not** backup files of running VM's
 - Backup of templates and stopped VM's is OK
 - But it is hard to find out which file belongs to a running VM
 - See later for backup concept of virtual disks

OVM Manager

- Officially:
 - Backup the configuration file:
`/u01/app/oracle/ovm-manager-3/.config`
- Oracle VM Manager (OVMM) must be installed in an existing Oracle/RedHat Enterprise Linux
 - Create a backup as for every physical server

OVM Manager

- Catalog database
 - OVMM uses MySQL (>3.1)
 - Backup-Tool: /opt/mysql/meb-*/bin/mysqlbackup
 - Automatic backups are created daily in /u01/app/oracle/mysql/dbbackup
 - Have enough space there! ~ 0.5GB/backup is required
 - Will be rotated automatically
 - Should be saved with the filesystem backup
- Manual backups can be created (not automatically rotated)

```
# /u01/app/oracle/ovm-manager-3/ovm_tools/bin/BackupDatabase -w
Enter your OVM Manager username: admin
Enter your OVM Manager password:

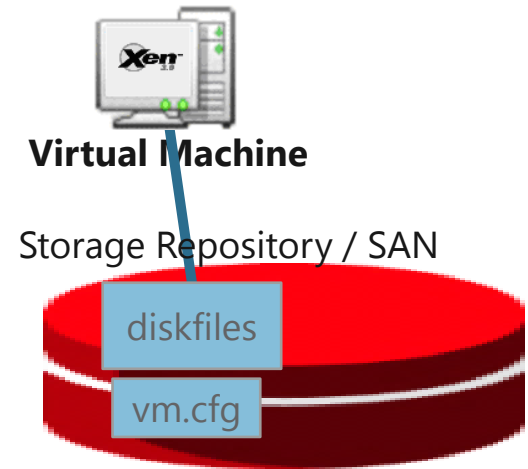
INFO: Backup job starting with destination:
      /u01/app/oracle/mysql/dbbackup/ManualBackup-20161106_123032

      Job Id   = 'Start Backup to: ManualBackup(1478431832638) Uri:
https://localhost:7002/ovm/core/wsapi/rest/Job/1478431832638'
      Job Name = 'Start Backup to: ManualBackup'

INFO: Backup job finished
```

Backup of Virtual Machines

- A VM consists of
 - Storage, (Virtual) disks
 - File or block-device on OVM Server,
 - iSCSI LUN or NFS
 - Virtual Hardware: a configuration file defining the hardware
 - Number of CPU's
 - MAC addr. of network cards
 - Mapping of files/block-devices to virtual disks
 - ...
- 2 different types of backup
 - Backup of the whole VM: → "bare-metal-recovery"
 - Backup of the contents of the (virtual) disks: → Restore of files/filesystems



Backup of Virtual Machines

Backup of the virtual hardware

- `<repo-fs>/VirtualMachines/<vm-uuid>/vm.cfg`
 - Generated by OVM Manager (do not edit manually...)
 - Backup of this file is required to restore the repository filesystem
- Store the create script of the VM
 - If done via GUI → reverse engineering... (show vm name=...)

```
# ssh -p 10000 admin@ovmmanager
OVM> create Vm name=slot011 repository=slotreposan01 domainType=XEN_PVM osType=OL_5
      bootOrder=DISK cpuCount=2 highAvailability=yes memory=2048 on ServerPool name=slot
OVM> create Vnic name=00:21:f6:01:00:0b network=vlanpublic
OVM> add Vnic name=00:21:f6:01:00:0b to Vm name=slot011
OVM> create VmDiskMapping slot=0 storageDevice=slot011_xvda name=xvda on Vm name=slot011
```

- Backup the OVM Manager Repository (as described before)
 - Also contains the definition of the VM's

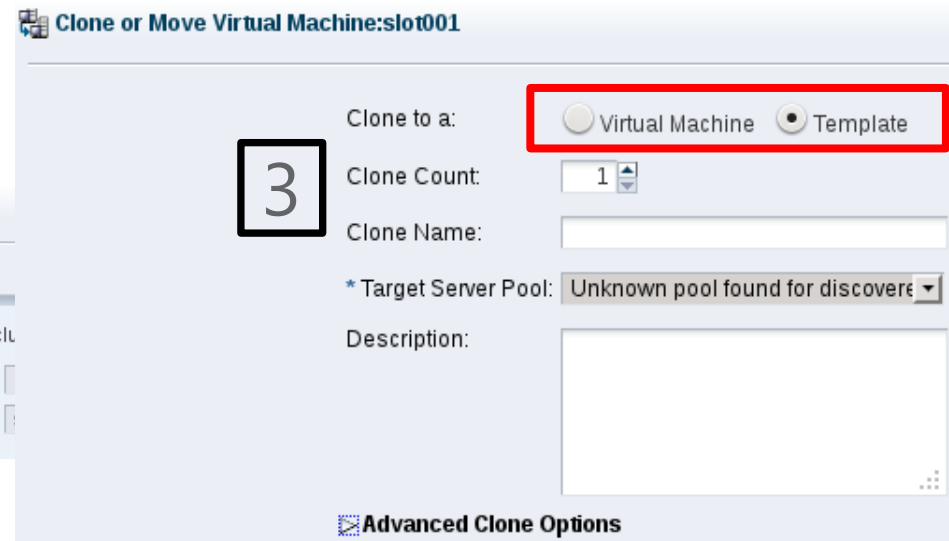
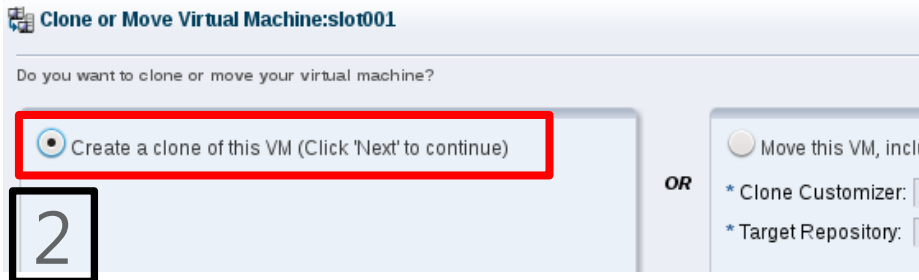
Backup of Virtual Machines

Backup of the virtual disks

- For bare-metal-recovery
- If virtual disk is a LUN on OVM server
 - With OS tool "dd" (requires stop of VM to be consistent)
 - Maybe your storage box offers a snapshot mechanism
- If virtual disk is a file on the storage repository
 - Save the whole file
 - If possible, create a **snapshot** and save the snapshot afterwards
 - A snapshot is NOT a backup, it is still on the same physical media!
 - Repository on SAN/iSCSI uses OCFS2, which offers snapshot functionality
 - If snapshots are not possible, stop the VM and backup the file
 - For Oracle it is recommended to set "alter database begin/end backup"

How to create Snapshots: Clone

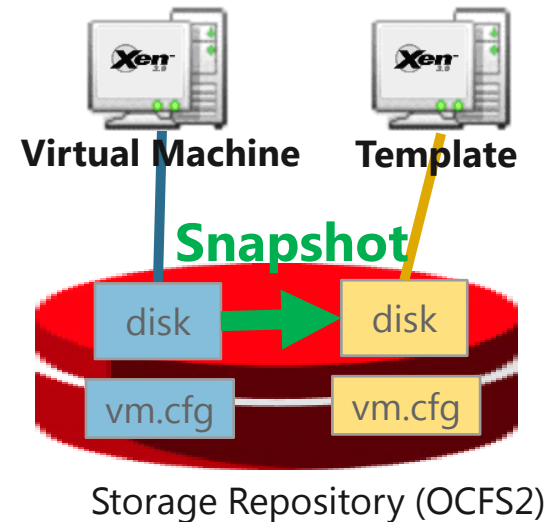
- What is a Snapshot?
 - Saved state of a VM
- There are no snapshots like in VMware or VirtualBox
 - Take a snapshot and revert to it at a later time
- But: you can clone a VM
 - Create a template (recommended)
 - Create a new VM



```
clone Vm name=vm1 destType=VmTemplate destName=clone1 serverPool=myServerpool
```

How to create Snapshots: Clone

- What cloning will do:
 - Copy files (if possible, a snapshot copy is created (on ocfs2))
 - Create new vm.cfg file
 - Register it as a template or new VM
- How to backup this template
 - Get the disk-names of the template
 - Backup it via NFS (repository-export)
 - Also save the mapping of this file (which diskslot it used; stored in vm.cfg)
 - → after the backup the template can be deleted
- How to restore
 - Get the name and slot-number of the lost diskfile
 - Restore the corresponding backed up diskfile of the template (rename it)



Backup of Virtual Machines: Files


Backup of the files / filesystems

- Files/Filesystems are entities inside the VM
- On this level, the backup concept is identical to a physical host
 - Install a backup client inside the VM
 - Store the backup on disk / nfs-share
 - Store the backup via network on tape
- Often backup clients are licensed per hostname
 - Every VM needs a license
 - But there is a workaround. See next slide...

Backup of Virtual Machines: Files

Backup of the files / filesystems via snapshot

- Create a snapshot of the diskfiles
- NFS-Export of Repository-filesystem to backup-server
- Loopback-Mount of snapshot-file and backup contents




```
mount ttcovms01.ttc.trivadis.com:/OVS/Repositories/0004fb000003000078778888fdcdb96e /mnt/  
cd /mnt/snapshots/  
kpartx -a slot009.20121121-102606.slot009_xvda  
kpartx -l slot009.20121121-102606.slot009_xvda  
mount /dev/mapper/loop0p2 /backups/slot009  
cd /backups/slot009  
tar zcvf ../slot009.20121121-102606.slot009_xvda.tgz *  
cd  
umount /backups/slot009  
kpartx -d slot009.20121121-102606.slot009_xvda
```

Backup of Virtual Machines: Files

Backup of the files / filesystems via snapshot

- If the backup server is a VM itself
 - You can add the cloned devices to the VM
 - No Network / NFS overhead, same performance as original



```
OVM> create VmDiskMapping slot=1 storageDevice=vm1_xvda name=vm1_xvda on Vm name=bckvm
```

```
bckvm# mount /dev/xvdb1 /backup
```

```
bckvm# ## run backup-script for /backup/
```

Backup of Virtual Machines: Database


- No difference to a physical server
 - Create **RMAN** backups to disk or tape
 - For tape: install media management software in VM → license costs
 - Configure Data Guard
- Save license costs of the tape-module: Backup of a snapshot
 - Similar to the filesystem-backup described above
 - Like "Split Mirror"
 - Put the database in the backup mode, then create a snapshot
 - Mount the snapshot file on a server with media manager rman module
 - Start the database in MOUNT mode (NOT open!)
 - Create an rman backup to tape

AGENDA

1. Overview of components
2. Backup
- 3. Recovery scenarios**

Loss of OVM Server

- Loss of local disk / OS
 - Variant 1
 - Replace and partition the disk and restore from a Backup
 - Variant 2
 - In a cluster environment, servers are replaceable
 - Data is on repository filesystem
 - Remove server from OVM Manager configuration
 - Replace the disk
 - Re-install the server
 - Re-add the server to the configuration




Loss of OVM Server

- Loss of other hardware components (e.g. Network)
 - Every server has an **UUID** derived from hardware, e.g. network MAC address
 - If OVMM does not re-configure automatically
 - Get UUID in ovmmmanager
 - Choose server → perspective "Info" → Advanced
 - Add this UUID on server in /etc/ovs-agent/agent.ini

```
[server]
fakeuuid = 35:38:33:39:31:34:43:5a:32:31:32:38:30:32:4a:59
```


- If it does not work, it is often easier to use variant 2 above

Loss of Pool-Filesystem



- Server is running, but filesystem not mounted, can't be used for VM's
 - Server: <server> does not have a virtual machine role
 - Restore the filesystem
 - On NFS, restore the 12 GB disk-file
 - Restore the file from the initial backup
 - Loopback mount the file and restore the latest filesystem backup
 - On SAN/iSCSI it is a OCFS2 filesystem
 - Restore the initial backup with "dd" -or-
 - re-create the filesystem
- ```
mkfs.ocfs2 -b 4K -C 4K -J size=4M -N 32 -L OVS_POOL_FILESYSTEM \
-U 0004fb00000500005c00fb0704a361b4 --global-heartbeat --cluster-stack=o2cb \
--cluster-name=cc55c8771248b2a0 --force /dev/mapper/33000000100000001
```
- Mount the filesystem temporarily (service o2cb must be started before)
  - Restore the latest filesystem backup
- Reboot the OVM servers

# Loss of Repository Filesystem



- Present the volume

Important!

- The restore or re-creation of the filesystem only works if the WWID of the LUN has not changed!
  - It is stored in binary configuration files on the server
  - The device is only mounted if it has the old WWID!

```
multipath -ll
360002ac0000000000000000740000d140 dm-1 3PARdata,VV
size=5.0T features='1 queue_if_no_path' hwhandler='0' wp=rw ...
```

```
/etc/ovs-agent/db # strings repository
...
fs_locationq
U-/dev/mapper/360002ac0000000000000000740000d140q
```



# Loss of Repository Filesystem

- Re-create the filesystem

```
mkfs.ocfs2 -J block64 -b 4096 -L OVS56e8973883d87 \
-U 0004fb00000500002dc56e8973883d87 -T vmstore -N 32 /dev/mapper/1IET_00010002
```

- Restore the structural filesystem backup
  - The backup without the diskfiles
- Restore the static files: vm.cfg, ISO-files, ...
- Restore the backup of the VM diskimages
  - Move the snapshot files to the original location
  - Same procedure as restoring a snapshot

# Loss of OVM Manager

- If OVM Manager is lost/corrupt, OVM servers are not directly affected
  - Running VM's will continue running without manager
- OVM is based on XEN virtualization
  - You can start / stop / relocate VMs manually with XEN commands

```
cd /OVS/Repositories/*/VirtualMachines
grep <vm-name> */vm.cfg #which file is the configuration file of <vm-name>
xm create <uuid>/vm.cfg
```

- Afterwards you have time to fix your OVM Manager


# Loss of OVM Manager

- Loss of the database

- Stop services on the manager and restore the database

```
/u01/app/oracle/ovm-manager-3/ovm_tools/bin/RestoreDatabase.sh <backup>
```

- Afterwards, restart ovmm components
- OVMM will re-synchronize an older backup with information on OVM servers



- Loss of the database and no backup ☹️

- Re-install OVMM with the same uuid

- can be found in .config file or /etc/sysconfig/ovmm
- `./runInstaller --uuid 0004EC000000100001C87C2AEE23B7`

- Re-Discover existing OVM server

- OVMM will reverse-engineer the configuration from the server
- Information not on server is lost → better to have a backup...
  - e.g. name of virtual disks




# Loss of OVM Manager

## ■ Loss of OVMM

- Restore the OS backup and database backup and re-start OVMM
- If no OS backup and no database backup
  - Re-install OVMM with the same uuid
    - can be found on OVM server /OVS/Repositories/\*/.ovsrepo
    - `./runInstaller --uuid 0004EC000000100001C87C2AEE23B7`
  - Discover existing OVM server
    - OVMM will reverse-engineer the configuration from the server
    - Information not on server is lost → better to have a backup...
  - Re-scan repository (will load info from vm.cfg files to repository)
    - If a repository is not visible, manually start a VM using this repo and re-start discovery

**restore all**  
**- or -**  
**reinstall all**



# Desaster scenario: Loss of all components

- Loss of "all" (except VM backups)
  - Fresh install of servers and OVMM
  - Restore of diskfiles to /OVS/Repositories/<uuid>/VirtualDisks
  - Re-scan repository
  - Re-create the VMs with a script and map the files to it

```
refresh repository name=slotreposan01
create Vm name=slot011 repository=slotreposan01 domainType=XEN_PVM osType=OL_5
 bootOrder=DISK cpuCount=2 highAvailability=yes memory=2048 on ServerPool name=slot
create Vnic name=00:21:f6:01:00:0b network=vlanpublic
add Vnic name=00:21:f6:01:00:0b to Vm name=slot011
create VmDiskMapping slot=0 storageDevice=slot011_xvda name=xvda on Vm name=slot011
```

# Loss of a VM

- Corrupted diskfile (logically/physically)
  - Restore the latest snapshot file and move it to the original location
  - Maybe it has to be restored from tape via NFS to repository-filessystem
- Logical loss (files deleted inside VM)
  - Present the snapshot file as a new disk-device and mount it read-only
  - Copy the lost/corrupted files back

```
host> reflink snapshots/vm021_xvdc_20130313 VirtualDisks/
OVM> refresh repository name=slotreposan01
OVM> create VmDiskMapping slot=50 storageDevice=slot021_xvdc_20130313
 name=xvdc_20130313 on vm name=slot021
slot021> mount /dev/xvday1 /mnt
slot021> cp -p /mnt/... /...
slot021> umount /mnt
OVM> delete VmDiskMapping id=<id_from_create_VmDiskMapping> #checks if umounted
OVM> delete VirtualDisk id=slot021_xvdc_20130313
```

# Loss of a Database in a VM

- If database is lost or point-in-time recovery required
  - Rman usually does a restore of a full-backup, then recovers the backup
  - Do not restore the database (the most time-consuming job)
  - Replace the diskdevice(s) of the database files with a snapshot
    - Usually DO NOT replace the diskdevice with the archivelogs (required for recovery)

```
slot021> umount /dev/xvdc1
OVM> show vm name=slot021
 VmDiskMapping 3 = 0004fb000013000049b9f5f2e1ceab18 [xvdc]
OVM> show VmDiskMapping id=0004fb000013000049b9f5f2e1ceab18
 Slot = 2
OVM> delete VmDiskMapping id=0004fb000013000049b9f5f2e1ceab18 #checks if unmounted
host> mv slot021_xvdc ../snapshots/slot021_xvdc.old
host> reflink ../snapshots/slot021_xvdc.20130313 slot021_xvdc
OVM> create VmDiskMapping slot=2 storageDevice=slot021_xvdc name=slot021_xvdc
 on vm name=slot021
slot021> mount /dev/xvdc1
```

- Recover the database as usual

# THANK YOU.

Trivadis AG

Martin Bracher

Europa-Strasse 5  
8152 Glattbrugg

Tel. +41 31 928 09 60 / +41 44 808 70 20  
Fax +41 44 808 70 21

info@trivadis.com  
www.trivadis.com

BASEL    BERN    LAUSANNE    ZÜRICH    DÜSSELDORF    FRANKFURT A.M.    FREIBURG I.BR.    HAMBURG    MÜNCHEN    STUTTGART    WIEN