

Home - VMware - VCDX Preparation - - Enterprise Admin - S1 - Storage - O1 - Datastores

# Objective 1.1 - Create and Administer VMFS Datastores Using Advance Techniques

Written by Matthijs van den Berg

Tuesday, 06 January 2009 08:56

## Knowledge

- Describe how to identify iSCSI, Fibre channel, SATA and NFS configurations using CLI commands and log entries

First we take a look at the configured volumes (EXT and VMFS) on the system by using the command `vdf`. This command provides u with a complete list of the configured volumes with their human readable name and the UUID used by VMware to identify the system. Example output:

```
v# vdf Filesystem 1K-blocks Used Available Use% Mounted on
/dev/cciss/c0d0p3 5160496 1822964 3075392 38% /
/dev/cciss/c0d0p1 98747 26659 66989 29% /boot
none 134284 0 134284 0% /dev/shm
/dev/cciss/c0d0p6 2064160 32920 1926388 2% /tmp
/dev/cciss/c0d0p5 2064160 127672 1831636 7% /var/log
/vmfs/devices 103110244 0 103110244 0% /vmfs/devices
/vmfs/volumes/480c68ee-c8f52d00-b97c-001e0bee3026
524025856 464945152 59080704 88% /vmfs/volumes/ATA01
```

To learn more about a specific volume like the adaptor it uses and how large the volume is details can be found by using the command `vmfstools`.

For example:

```
# vmfstools -P -h /vmfs/volumes/480c68ee-c8f52d00-b97c-001e0bee3026
VMFS-3.31 file system spanning 1 partitions.
File system label (if any): DATA01
Mode: public
Capacity 499G, 56G available, file block size 8.0M
UUID: 480c68ee-c8f52d00-b97c-001e0bee3026
Partitions spanned (on "lvm"):
vmhbal:0:11:1
```

The last line show the adaptor in use for this volume.

You can also use the command

```
esxcfg-module -l
```

to look at the modules that are currently loaded. You can find much info regarding the adaptors in use by using the following command:

```
cat /proc/scsi/qla2300/?
```

(this is for the qla2300 fibre channel, find the adaptor in use by a ls of the /proc/scsi directory). The output shows information about the adaptor like BIOS and driver version, card capabilities and the LUNs found by the adaptor.

- Describe the VMFS file system

- Metadata

Metadata is data describing the characteristics of files. VMware uses this metadata for example to describe a snapshot. Also the `vmx` file holding information regarding the Virtual Machines can be considered a metadata file holding information about the disks, snapshot devices, etc.

- Multi-access and locking

VMware uses shared storage between VMware ESX hosts. This means the files on this shared storage are visible to all ESX hosts that are granted access. Because a virtual machine is generally only active in one ESX host a lock file is maintained. VMFS provides on-disk locking to ensure that multiple servers do not power on a virtual machine at the same time and file consistency is maintained. Should a server fail, the on-disk lock for each virtual machine is released so that virtual machines can be restarted on other physical servers.

- Extents

### Menu - VCDX

- Introduction
- More Info
- Enterprise Admin
- Design Exam

### VCDX Cloud

virtual  
server  
vmware network  
host file service  
esxcfg click  
configure machine  
traffic vswitch  
machines storage  
command console  
hosts cluster disk  
client physical  
configuration files port

### Recent Blog Postings

When you are creating a VM in Virtual Infrastructure 3 (ESX 3.x) that needs a specific MAC address VMware allows you to modify this only in a specific range. This range is: 00:50:56:00:00:00 - 00:50:56:3f:ff:ff When creating a address outside this range the GUI will not allow you to do so and .....

[Read more](#)

### Add Banners

VMware volumes can be extended to spread across multiple LUNs. This technique can be used to spread IO's over multiple LUNs increasing performance but is also known to generate some errors. In the recent past it is much used to extend VMFS volumes that were too small, but this need has become less since the introduction of storage vMotion.

#### o Tree structure and files

Virtual machines files are stored in a per VM directory. These files can be found in /vmfs/volumes// by their human readable name as well by their UUID name. VMware has the following table from a blog article about the files used for Virtual Machines.

Extension	File Name	Description
.log	<vmname>.log	This is the file that keeps a log of key VMware Workstation activity. This file can be useful in troubleshooting if you encounter problems. This file is stored in the directory that holds the configuration (.vmx) file of the virtual machine.
	or vmware.log	
.nvram	<vmname>.nvram	This is the file that stores the state of the virtual machine's BIOS.
	or nvram	
.vmdk	<vmname>.vmdk	<p>This is a virtual disk file, which stores the contents of the virtual machine's hard disk drive.</p> <p>A virtual disk is made up of one or more .vmdk files. If you have specified that the virtual disk should be split into 2GB chunks, the number of .vmdk files depends on the size of the virtual disk. As data is added to a virtual disk, the .vmdk files grow in size, to a maximum of 2GB each. (If you specify that all space should be allocated when you create the disk, these files start at the maximum size and do not grow.) Almost all of a .vmdk file's content is the virtual machine's data, with a small portion allotted to virtual machine overhead.</p> <p>If the virtual machine is connected directly to a physical disk, rather than to a virtual disk, the .vmdk file stores information about the partitions the virtual machine is allowed to access.</p> <p>Earlier VMware products used the extension .dsk for virtual disk files.</p>
	<diskname>-<###>.vmdk	This is a redo-log file, created automatically when a virtual machine has one or more snapshots. This file stores changes made to a virtual disk while the virtual machine is running. There may be more than one such file. The ### indicates a unique suffix added automatically by VMware Workstation to avoid duplicate file names.
.vmem	<uuid>.vmem	The virtual machine's paging file, which backs up the guest main memory on the host file system. This file exists only when the virtual machine is running, or if the virtual machine has crashed.
	<snapshot_name_and_number>	Each snapshot of a virtual machine that is powered on has an associated .vmem file, which contains the guest's main memory, saved as part of the snapshot.
.vmsd	<vmname>.vmsd	This is a centralized file for storing information and metadata about snapshots.
.vmsn	<vmname>-Snapshot.vmsn	This is the snapshot state file, which stores the running state of a virtual machine at the time you take that snapshot
	<vmname>- Snapshot<###>.vmsn	This is the file which stores the state of a snapshot
.vmss	<vmname>.vmss	<p>This is the suspended state file, which stores the state of a suspended virtual machine</p> <p>Some earlier VMware products used the extension .std for suspended state files</p>

.vmtm	<vmname>.vmtm	This is the configuration file containing team data.
.vmx	<vmname>.vmx	This is the primary configuration file, which stores settings chosen in the New Virtual Machine Wizard or virtual machine settings editor. If you created the virtual machine under an earlier version of VMware Workstation on a Linux host, this file may have a .cfg extension
vmxf	<vmname>.vmxf	This is a supplemental configuration file for virtual machines that are in a team. Note that the .vmxf file remains if a virtual machine is removed from the team.

There can be other files in the directory, some of which are present only while a virtual machine is running.

- **Applicability to clustered environment**

The file system is also ready to deliver volumes to more than one virtual machine. This functionality can be used if a cluster consisting of two virtual machines is created and a shared disk (for example for a windows cluster a quorum disk) is needed.

- **Journaling**

Journaling is a technique used to prevent file system corruption. Changes to the file system are written to a journal log before they are committed to disk. VMFS3 also used distributed journaling to enhance this protection provided by journaling and recover VMs faster and more reliable.

- **Explain the process used to align VMFS partitions**

Aligning partitions can increase performance for the virtual machines. If the VMFS and the virtual machines disk are properly aligned the number of chunks on a LUN that need to be read can be decreased.

To check that your existing partitions are aligned, issue the command:

```
fdisk -lu /dev/sd*
```

The output is similar to:

Device	boot	Start	End	Blocks	Id	System
/dev/sdj	1	128	167766794	83883333+	fb	Unknown

Aligned partitions start at **128**. If the Start value is **63** (the default), the partition is not aligned. If you create a VMFS from the GUI the partition will automatically be aligned to the LUN. If you create it manually you will need to align the partition. IP based storage does not need to be aligned.

*More information regarding aligning can be found in this VMware whitepaper:*

[http://www.vmware.com/pdf/esx3\\_partition\\_align.pdf](http://www.vmware.com/pdf/esx3_partition_align.pdf) .

- **Explain the use cases for round-robin load balancing**

Round Robin load balancing can be used to load balance storage traffic over multiple adaptors. If for example a Fibre channel based storage array is used with an active active controller and two HBA's in a ESX servers each LUN will be seen four times. The best practice would be to load balance the LUNs over all available HBAs and storage adaptors spreading the load and optimizing performance whilst maintaining high availability in case of hardware failure.

Round Robin can be used to configure the multi pathing to spread the load automatically. You can configure when to switch, what target to use and what HBA to use. Basic configuration is done via the Virtual Center Client, advanced configuration via the CLI.

## Skills and Abilities

- **Perform advanced multi-pathing configuration**

- **Configure multi-pathing policy**

- **Configure round-robin behavior using command-line tools**

Round robin load balancing for HBA's can be configure via the CLI using the esxcfg-mpath command. This command is also used to configure static multi-pathing. The options regarding round robin load balancing for this command are:

**--lun**

Points out the LUN you would like to change the setting for. Use the vml LUN number with you can find with

```
esxcfg-mpath -l -v
```

**--policy rr**

The multi-pathing policy. rr Stands for round robin. Other options are mru, fixed and custom.

- **Manage active and inactive paths**

Active and inactive paths can be managed from the Virtual Infrastructure Client and the CLI. In the VIC you can click a ESX server, Configuration, Storage Adaptors, right click a LUN and select Manage paths. You can set the policy (fixes,

MRU or RR) and the preferred path can be set.

- **Verify SAN LUN accessibility**

GUI:

Browse the LUN from the GUI and see if you find VM's on the LUN

CLI:

Go to the virtual machine on the LUN. The VM can be found in /vmfs/volumes/. If you can find VM files the LUN is accessible.

- **Configure and use NPIV HBAs**

NPIV is a functionality to have multiple World Wide Names (WWNs) over a single adaptor. By using this technology you can provide Virtual machines with their own WWN and assigning LUNs directly to this VM.

NPIV needs to be configured in the entire SAN to be operational. You can check if the HBA supports NPIV by checking what module you use for your storage in the /proc/scsi directory. QLogic HBAs are listed as qla2300. Emulex HBAs are listed as lpfc. For the qla2300 module the command is:

```
cat /proc/scsi/qla2300/? |grep NPIV
```

- **Manage VMFS file systems using command-line tools**

The command vmkfstools is used for VMFS volumes and Virtual Machines disks to creation, upgrade, maintenance and delete partitions. There are various options

- **Configure NFS datastores using command-line tools**

The esxcfg-nas command is used to list, mount and dismount NFS exports for the VMkernel. In the first example we list the NFS datastores which the VMkernel has mounted.

To mount a NFS share the following command can be used:

```
esxcfg-nas -a -o -s
```

- **Configure iSCSI hardware and software initiators using command-line tools**

TODO

## Tools

- VI client
- CLI
  - esxcfg-vmhbadevs
  - vdf
  - fdisk
  - vmkfstools

Next >