LVM / VxVM

Vocabulaire

LVM Term	VxVM Term
Both LVM and VxVM enable online disk storage management. They both build virtual devices, called volumes, on physical disks. Volumes are not limited by the underlying physical disks, and can include other virtual objects such as mirrors. Volumes are accessed through the HP-UX file system, a database, or other applications in the same manner as physical disks would be accessed.	
Physical Volume	VxVM Disk
An LVM physical volume and a VxVM disk are conceptually the same. A physical disk is the basic storage device (media) where the data is ultimately stored. You can access the data on a physical disk by using a device name (devname) to locate the disk. In LVM, a disk that has been initialized by LVM becomes known as a physical volume. A VxVM disk is one that is placed under the Volume Manager control and is added to a disk group. VxVM can place a disk under its control without adding it to a disk group. The VxVM Storage Administrator shows these disks as “free space pool”.	
Logical Volume	Volume
An LVM logical volume and a VxVM volume are conceptually the same. Both are virtual disk devices that appear to applications, databases, and file systems like physical disk devices, but do not have the physical limitations of physical disk devices. Due to its virtual nature, a volume (LVM or VxVM) is not restricted to a particular disk or a specific area of a disk. An LVM volume is composed of fixed length extents. LVM volumes can be mirrored or striped. VxVM volumes consist of one or more plexes/mirrors holding a copy of the data in the volume which in turn are made up of subdisks with arbitrary length. The configuration of a volume can be changed by using the VxVM user interfaces. See the VERITAS Volume Manager Administrator's Guide for more information. VxVM volumes can be one of four types: mirrored, RAID-5, striped, or concatenated.	
Volume Group	Disk Group
LVM volume groups are conceptually similar to VxVM disk groups. An LVM volume group is the collective identity of a set of physical volumes, which provide disk storage for the logical volumes. A VxVM disk group is a collection of VxVM disks that share a common configuration. A configuration is a set of records with detailed information about related VxVM objects, their attributes, and their associations. In addition, both LVM and VxVM have the following characteristics: Volumes can be mapped to multiple VxVM disks or LVM physical volumes. VxVM disks must reside in only one disk group, and LVM physical volumes must reside in one volume group.	
Physical Extent	Subdisk
User data is contained in physical extents in LVM and subdisks in VxVM. The LVM physical extents are of a fixed length. LVM allocates space in terms of physical extents which is a set of physical disk blocks on a physical volume. The extent size for all physical volumes within a volume group must be the same, and is usually 4 MB. VxVM allocates disk space in term of subdisks which is a set of physical disk blocks representing a specific portion of a VxVM disk and is of arbitrary size.	
LVM metadata	Private Region
LVM metadata and the Private Region are similar conceptually. In LVM, metadata is stored in a reserved area in the disk. In VxVM, the private region of a disk contains various on-disk structures that are used by the Volume Manager for various internal purposes. Private regions can also contain copies of a disk group's configuration, and copies of the disk group's kernel log.	
Unused Physical Extent	Free Space
LVM contains unused physical extents that are not part of a logical volume, but are part of the volume group. Similarly, free space is an area of a disk under VxVM that is not allocated to any subdisk or reserved for use by any other Volume Manager object.	
Mirrors	Mirrors (Plexes)

LVM Term	VxVM Term
Both LVM and VxVM support mirrors. Mirrors can be used to store multiple copies of a volume's data on separate disks. In LVM, you can create mirrors using the MirrorDisk/UX product. Mirrors allow duplicate copies of the extents to be kept on separate physical volumes. MirrorDisk/UX supports up to 3 copies of the data. A VxVM mirror consists of plexes. Each plex is a copy of the volume. A plex consists of one or more subdisks located on one or more disks. VxVM volumes can have up to 32 mirrors (where each plex is a copy of data). Mirroring features are available with an additional license.	
Export	Deport
In LVM, exporting removes volume group information from /etc/lvmtab. The volume group must have already been deactivated. Similarly in VxVM, deport makes a disk group inaccessible by the system.	
Import	Import
In LVM, import adds a volume group to the system and the volume group information to /etc/lvmtab but does not make the volumes accessible. The volume group must be activated by the vgchange -a y command in order to make volumes accessible. In VxVM, import imports a disk group and makes the diskgroup accessible by the system.	
Bad Block Pool	No similar term
In LVM, the bad block pool provides for the transparent detection of bad disk sectors, and the relocation of data from bad to good disk sectors. The bad block reallocation feature does not exist in VxVM because the vectoring of bad blocks is now done by most hardware.	
/etc/{lvmtab,lvmtab_p}	No similar term
/etc/lvmtab is a file that contains information about volume groups that are accessible by a system.	

Quelques commandes

Description	LVM	VxVM
Create a disk	**pvcreate** /dev/rdsk/disk_name	**vxdiskadd** device_name Option 1 in the vxdiskadm menu adds a disk and initializes it
Create a volume group	**vgcreate** /dev/vol_grp /dev/dsk/disk_name	**vxdg init** disk_group disk_name
Add a new disk to the existing volume group	**vgextend** /dev/vol_grp /dev/rdsk/disk_name	**vxdg -g** disk_group **adddisk disk=**devicename
Extend a logical volume or increase space allocated to a logical volume	**lvextend -l** 50 /dev/vol_grp/lvol_name (-l indicates the number of logical extents in the logical volume) **fsadm resize** /dev/vol_grp/lvol_name new_size (to extend FS)	**vxresize -g** disk_group **-F vxfs** vol_name length **vxassist growto** vol_name new_length **vxassist growby** vol_name length_change Grow the file system after growing the volumes
Reduce a logical volume	**lvreduce -L** to_size /dev/vol_grp/lvol_name (-L indicates the number of megabytes.)	**vxresize -g** disk_group **-F vxfs** vol_name to_length **vxassist -b shrinkby** vol_name length **vxassist -b shrinkto** vol_name newlength shrink the file system before reducing the volume
Import and activate a volume group	**vgimport -v** /dev/vol_grp /dev/dsk/disk_name **vgchange -a y** /dev/vol_grp	**vxdg -tfC -n** newname **import** disk_group

Description	LVM	VxVM
Export and deactivate a volume group, and its associated logical volumes	**vgchange -a n** vol_group **vgexport** /dev/vol_group	**vxdg deport** disk_group
Back up volume group configuration information	**vgcfgbackup -f** /pathname/filename vol_grp	**dgcfgbackup -f** /pathname/filename vol_grp
Restore volume group configuration to a particular physical volume	**vgrestore -n** /dev/vol_grp /dev/rdsk/disk_name	**dgcfgrestore -n** /dev/vol_grp /dev/rdsk/disk_name
Increase or decrease secondary swap space. Enlarge an existing swap logical volume, or add a new swap logical volume.	lvextend to increase swap space. lvreduce to decrease swap space	vxassist make swapvol2 size
Remove a volume group	This destroys a volume group by removing its last disk and removing it from /etc/lvmtab. **vgremove** /dev/vol_grp. This is preceded by **lvremove** and **vgreduce** down to the last disk.	**vxdg deport** disk_group **vxdg init** disk_group
Extend a volume group by adding disks to the volume group	**vgextend** /dev/vol_grp/ /dev/dsk/disk_name	**vxdiskadd** disk_name (option 1 in the vxdiskadm main menu)
Reduce a volume group by reducing the number of disks	**vgreduce** /dev/vol_grp /dev/dsk/disk_name	**vxdg -g** disk_group **-k rmdisk** disk_name
Mirroring a disk	1. **pvcreate** /dev/rdsk/second_disk 2. **vgextend** /dev/vol_grp /dev/dsk/second_disk 3. **lvextend -m** no_of_mirrors /dev/vol_grp/lvol_name /dev/dsk/second_disk	**vxmirror -g** disk_group **-d** yes or no disk_name [new_disk_name] **vxmirror -d** yes disk_name (Option 6 in the vxdiskadm menu)
Mirroring a root disk	**pvcreate -B** /dev/rdsk/second_disk **mkboot -l** /dev/dsk/second_disk **vgextend** /dev/vol_grp /dev/dsk/second_disk **lvextend -m** no_of_mirrors /dev/vol_grp/root_lvol /dev/dsk/second_disk **lvlnboot -r** /dev/vol_grp/lvol_name	**vxrootmir** [-v][-t tasktag] disk_access_name or disk_media_name
Create a logical volume	**lvcreate -L** vol_size /dev/vol_grp	**vxassist make** vol_name length **vxassist make** vol_name length **layout=**mirror, stripe, raid5
Display information about logical volumes	**lvdisplay** /dev/vol_grp/vol_name	**vxprint -vt** **vxprint -ht** vol_name
Display information about volume groups	**vgdisplay -v** /dev/vol_grp	**vxdisk list** **vxdg list** **vxdisk list** disk_group **vxprint -g** disk_group

Description	LVM	VxVM
Display information about physical volumes	**pvdisplay** /dev/dsk/disk_name	**vxinfo** or **vxprint**
Remove a logical volume	**lvremove** /dev/vol_grp/lvol_name	**vxedit rm** vol_name
Set up alternate links to a physical volume	If a disk has two controllers, you can make one primary and the other an alternate link **vgcreate** /dev/vol_grp /dev/dsk/disk_name /dev/dsk/disk_name_2 To remove the link **vgreduce** /dev/vol_grp /dev/dsk/disk_name	The Multipathing disk devices in the Volume Manager represent virtual devices with one or more physical access paths to a particular physical disk. Dynamic Multipathing provides reliability of disk access by dynamically switching to another physical path in the event of failure of a path. The DMP feature in VxVM sets up links automatically. It is not required to set up links separately. Refer to the manual page vxdmp(7) and the VERITAS Volume Manager Administrator's Guide for more information on DMP.
Create a mirrored logical volume	**lvcreate -l** num_log_extents **-m** 1 **-n** mirr_lv /dev/vol_grp	**vxassist make** vol_name length **layout=mirror**
Remove a mirrored logical volume	**lvreduce -m 0** /dev/vol_grp/mirr_lv **lvremove** /dev/vol_grp/mirr_lv	**vxplex -o rm dis** plex_name **vxedit -rf rm** vol_name
Increase the number of mirror copies	**lvextend -m 2** /dev/vol_grp/lvol_name	**vxassist mirror** vol_name
Split a logical mirror volume	**lvsplit -s backup** /dev/vol_grp/lvol_name	
Combine two logical volumes back into a mirrored logical volume	**lvmerge** /dev/vol_grp/split_vol_name /dev/vol_grp/lvol_name (split_vol_name= active logical volume)	
Snapshot a volume and create a new volume		**vxassist snapshot** vol_name new_vol_name
Returns the snapshot plex to the original volume from which it was snapped		**vxassist snapback** new_vol_name
Move a mirrored logical volume from one disk to another	**pvmove -n** /dev/vol_grp/lvol_name /dev/dsk/disk_name /dev/dsk/disk_name2	**vxplex mv** orig_plex new_plex
Synchronize a mirrored logical volume	**lvsync** /dev/vol_grp/lvol_name	**vxvol** resync
Synchronize a volume group	**vgsync** /dev/vol_grp	**vxrecover -s** vol_name
Start a volume	**lvchange -a y** /dev/vol_grp/lvol_name	**vxvol start** vol_name
Stop a volume	**lvchange -a n** /dev/vol_grp/lvol_name	**vxvol stop** vol_name
Make a disk available as a hot spare	**pvchange -z y** /dev/dsk/disk_name	**vxedit set spare=on** disk_name

That's all Folks !