

forensic-proof.com

Logical Volume Manager

- Software RAID



Proneer

@pr0neer

proneer@gmail.com

<http://forensic-proof.com>

Outline

1. *Background*
2. *LVM Basics*
3. *LVM Structure*
4. *LVM Advantage & Disadvantage*




Background



Background

Layers to make it easier

- Layers in a typical system (before LVM)

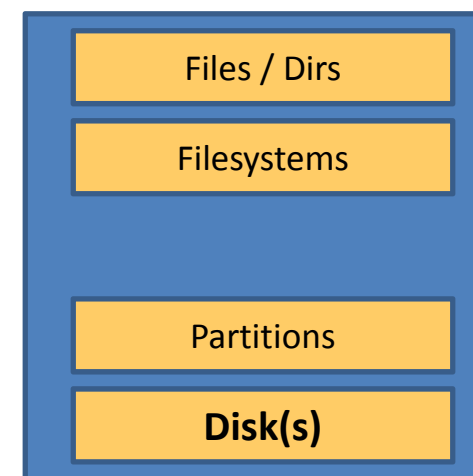


<http://home.techwiz.ca/~peters/presentations/lvm/oclug-lvm.pdf>

Background

What is a Disk?

- Lets look at some terms first, starting from the bottom
- It has various techie things in it:
 - Cylinders, Heads, Sectors, mbr, partitions
 - It's own cpu, cache, firmware etc
- Linux : `/dev/hda`, `/dev/sda`
- Windows : `\\PhysicalDrive0`, `\\PhysicalDrive1`

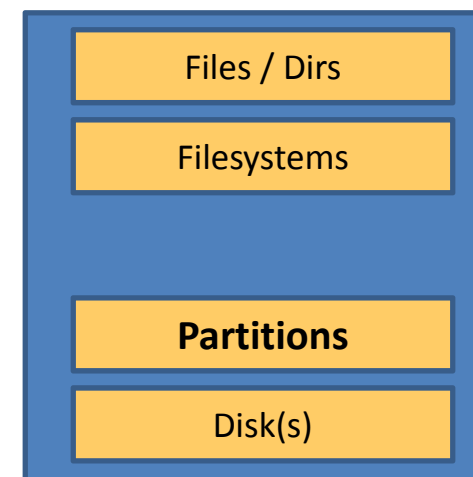


<http://home.techwiz.ca/~peters/presentations/lvm/oclug-lvm.pdf>

Background

Partitions

- *One partition must be a continuous chunk of blocks*
- The original IBM PC from 1981 only had 4 primary partitions
- On a typical linux system you have something like
 - hda1 /
 - hda2 swap
 - hda3 /home
- Or for a dual boot
 - hda1 M\$ Windows C:
 - hda2 swap
 - hda3 /




<http://home.techwiz.ca/~peters/presentations/lvm/oclug-lvm.pdf>

Background

Partition tools

- Partition tools can change the existing partition *without loosing existing data*
 - I never tried this tools
 - *They are all very risky*
- *Parted*
 - A free tool that can handle ext2 and vfat but not NTFS
- *Acronis Disk Director & Partition magic*
 - Commercial tool that can do more including NTFS

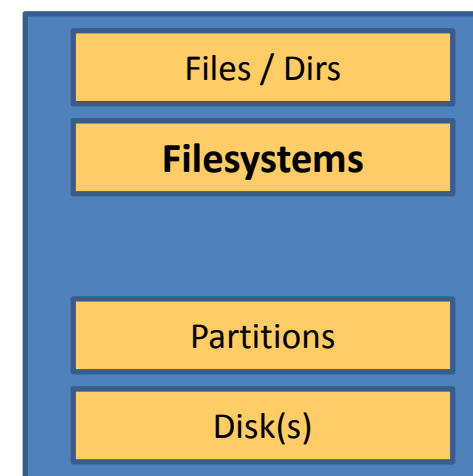


<http://home.techwiz.ca/~peters/presentations/lvm/oclug-lvm.pdf>

Background

Filesystems

- Filesystem are a fancy way to hold together a group of files and directories
- Without LVM *one filesystem = one partition*
- Common File Systems:
 - In the Windows World :
 - FAT12/16/32, exFAT
 - NTFS
 - In Linux
 - Ext2/3/4, Reiserfs, jfs, xfs
 - FAT/NTFS
 - And many many more

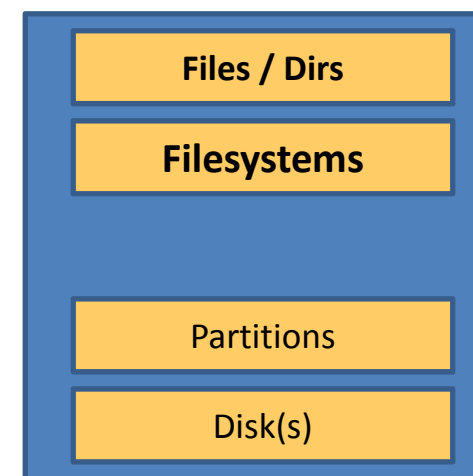


<http://home.techwiz.ca/~peters/presentations/lvm/oclug-lvm.pdf>

Background

Files on a linux (and unix) system

- Within User Level we have
 - Files
 - Directories
 - Filesystems
- Files Hold Data
- Directories hold files and directories
- Filesystems holds directories and files
- Mount points hold Filesystems



<http://home.techwiz.ca/~peters/presentations/lvm/oclug-lvm.pdf>

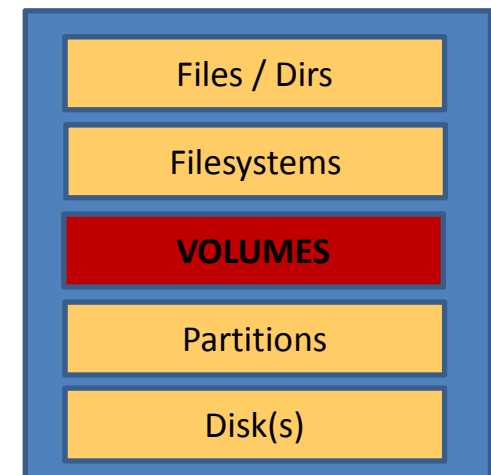
Logical Volume Manager Basics



LVM Basics

Why Volume Management?

- To make your life as System Administrator easier
- To give you a bigger Sandbox to play in
- *One Filesystem = One Volume*
- *One Volume = more than One Partition*
- This has many advantages
 - Can carve out non-continuous filesystems
 - Can add disk partitions together so they look like one BIG




<http://home.techwiz.ca/~peters/presentations/lvm/oclug-lvm.pdf>

LVM Basics

What is LVM?

- Stands for Logical Volume Management
- It manages disk drives and *similar mass-storage devices*, in particular large ones
- *available in Linux* (Arch Linux, Debian, Fedora, Gentoo, MontaVista Linux, openSUSE, SLED, Slackware, SLES and Ubuntu)
- *as well as HP-UX, IBM AIX, Solaris, OS/2 and Windows*
- It is placed between the filesystems and disk partitions



<http://home.techwiz.ca/~peters/presentations/lvm/oclug-lvm.pdf>

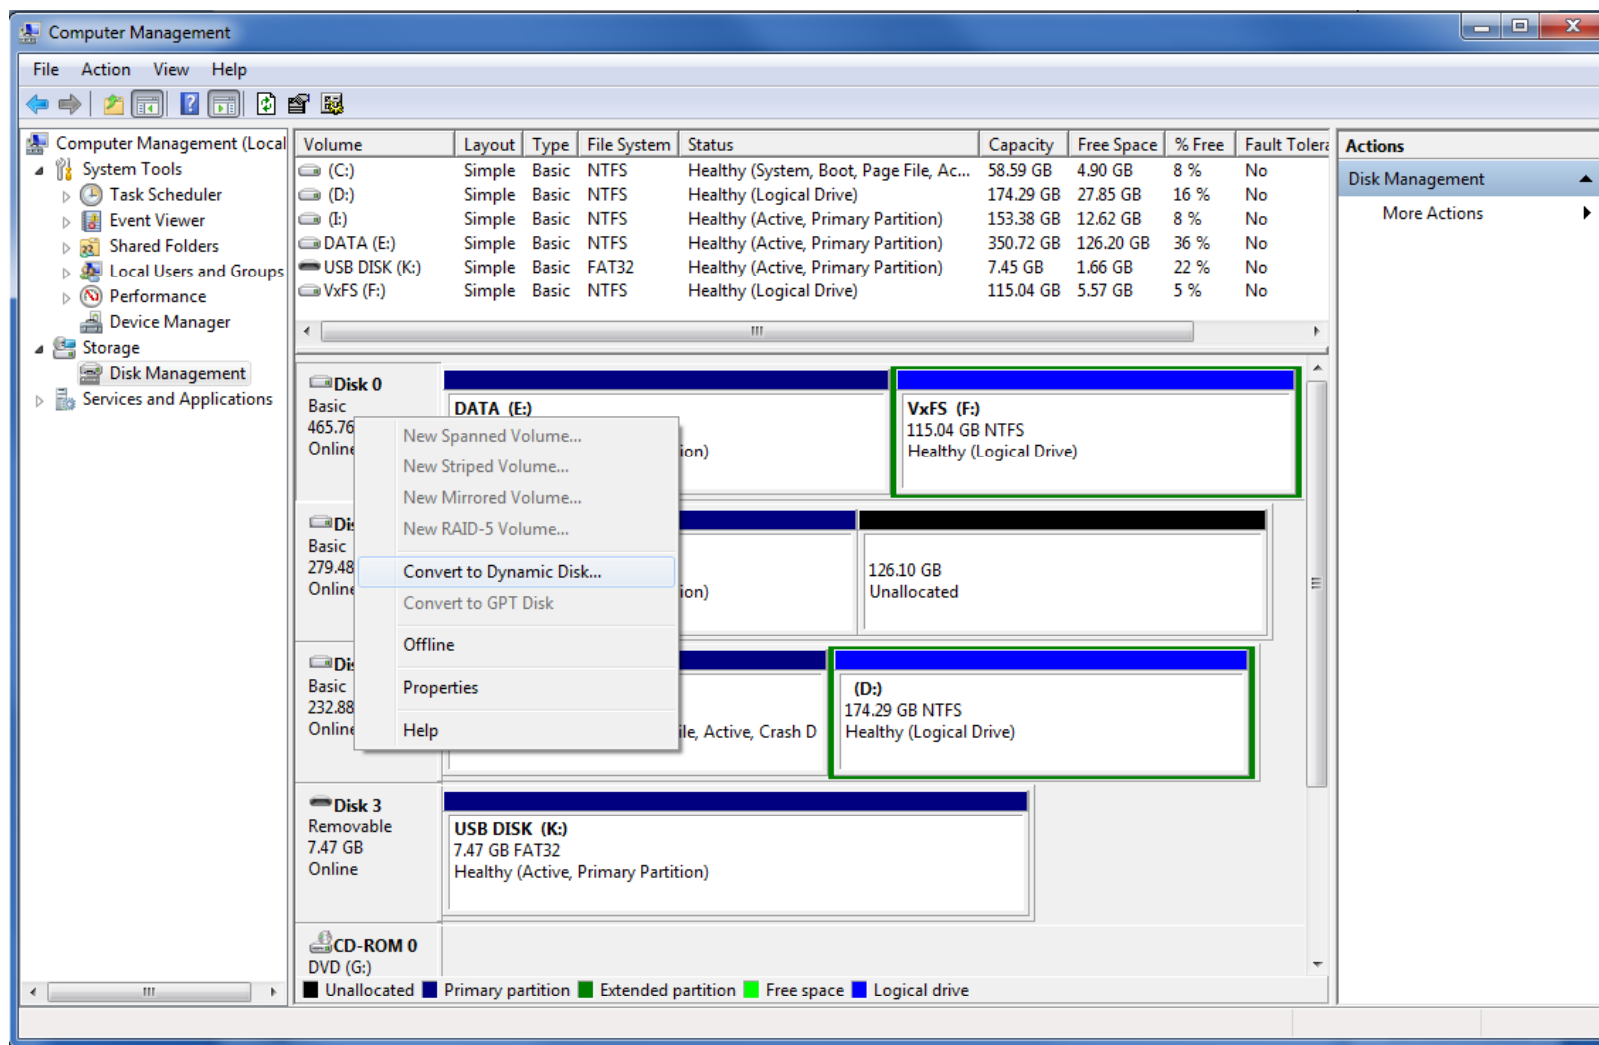
LVM Basics

Implementations each OS

Vendor	Introduced in	Volume manager	Allocate anywhere	Snapshots	RAID 0	RAID 1	RAID 5	RAID 10
IBM	AIX 3.0	Logical Volume Manager	YES	NO	YES	YES	NO	YES
Hewlett-Packard	HP-UX 9.0	HP Logical Volume Manager	YES	YES	YES	YES	NO	YES
	FreeBSD	Vinum Volume Manager	YES	NO	YES	YES	YES	
	NetBSD	Logical Volume Manager	YES	NO	YES	YES	NO	NO
	Linux 2.2	Logical Volume Manager	YES	YES	YES	YES	NO	
	Linux 2.4	Enterprise Volume Management System	YES	YES	YES	YES	YES	
	Linux 2.6	Logical Volume Manager	YES	YES	YES	YES	YES	
Silicon Graphics	IRIX or Linux	XVM Volume Manager	YES	YES	YES	YES	YES	
Sun Microsystems	SunOS	Solaris Volume Manager	NO	NO	YES	YES	YES	YES
Sun Microsystems	Solaris 10	ZFS	YES	YES	YES	YES	YES	YES
Veritas	Cross-OS	Veritas Volume Manager	YES	YES	YES	YES	YES	YES
Microsoft	Later NT-based	Logical Disk Manager	YES	YES	YES	YES	YES	

LVM Basics

Windows Dynamic Disk – *required 8MB area*




Logical Volume Manager Structure



LVM Structure

What is LVM? cont.


- The appearance of New Terms
- *Physical Volumes PVs*
 - collects all disk partitions
- *Volume Group VGs*
 - creates one big virtual disk
- *Logical Volumes LVs*
 - From the VG you can then create filesystem within LVs
- *Extents*
 - Physical & Logical Extents



<http://home.techwiz.ca/~peters/presentations/lvm/oclug-lvm.pdf>

LVM Structure

LVM Structure

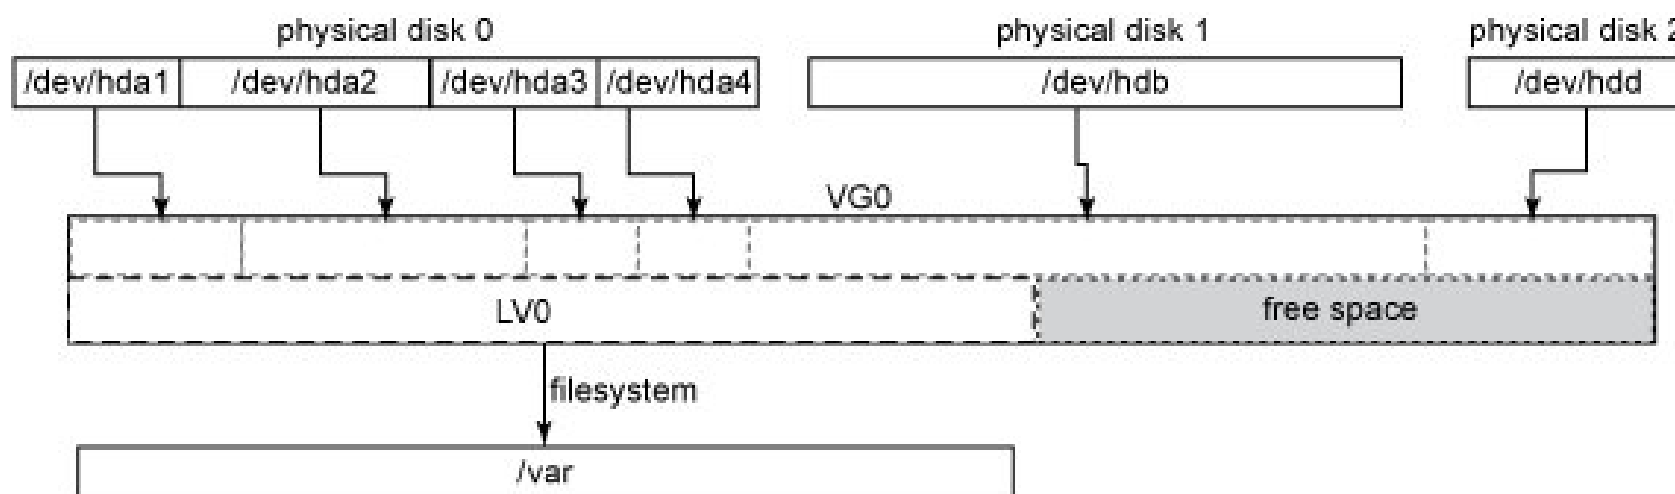


linuxconfig.org

http://www.linuxconfig.org/Linux_lvm_-_Logical_Volume_Manager

LVM Structure


LVM Structure cont.



http://sunoano.name/ws/public_xhtml/lvm.html

LVM Structure

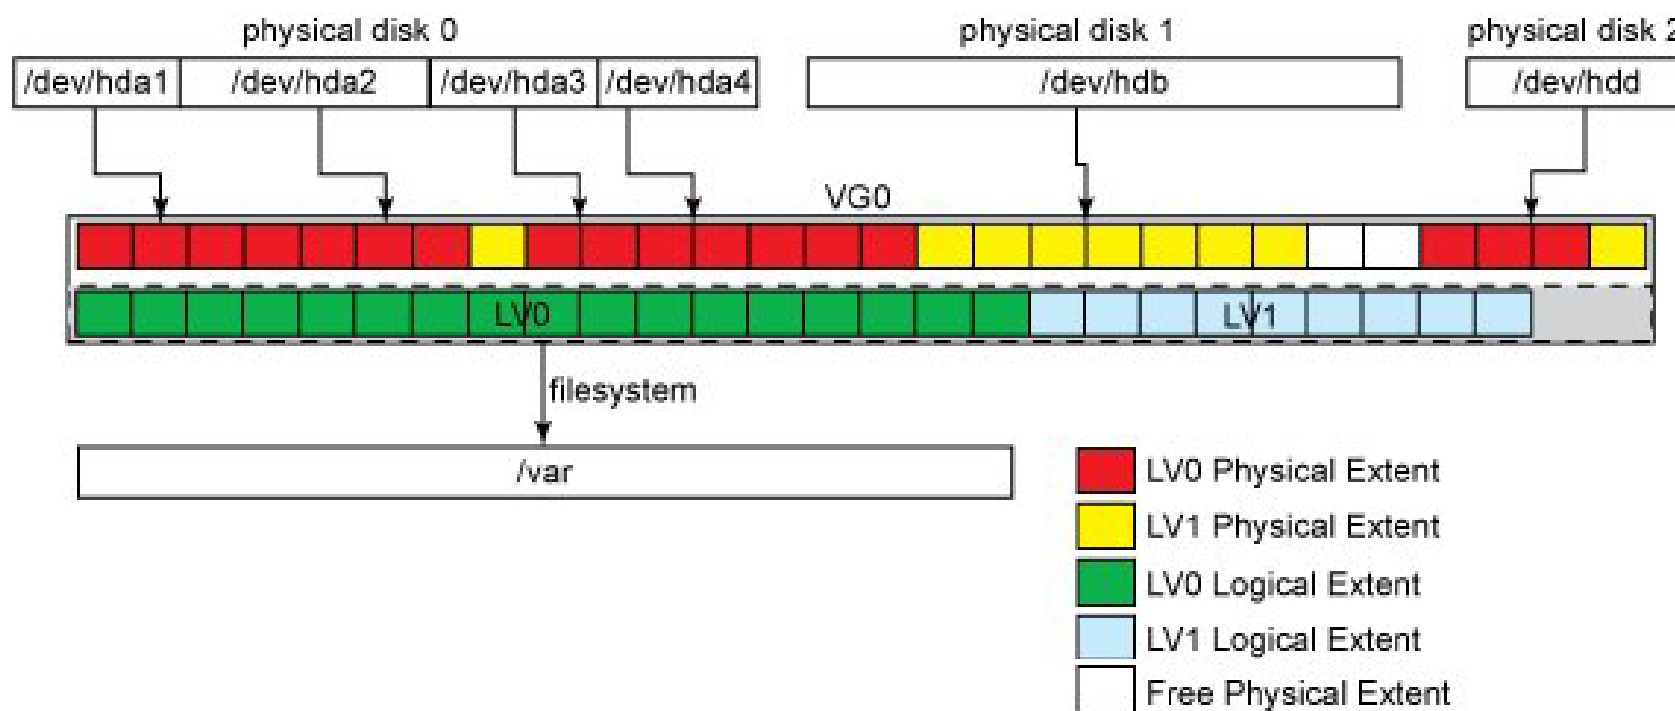
LVM Structure cont.



http://sunoano.name/ws/public_xhtml/lvm.html

LVM Structure


LVM Structure cont.



http://sunoano.name/ws/public_xhtml/lvm.html

LVM Advantage & Disadvantage

Lowest Level of LVM



- VGDA (Volume Group Descriptor Area)
- VGSA (Volume Group Status Area)
- LVCB (Logical Volume Control Block)



Logical Volume Manager Advantage & Disadvantage



LVM Advantage & Disadvantage

Advantages



/boot, /dev, /etc, /lib ...



LVM Advantage & Disadvantage

Advantages

- *Resizing*
 - Resize volume groups online by absorbing new physical volumes(PV) or ejecting existing ones
 - Resize logical volumes (LV) online by concatenating extents onto them or truncating extents from them
- *Snapshot*
 - Create read-only snapshots of logical volumes (LVM1)
 - Create read-write snapshots of logical volumes (LVM2)



LVM Advantage & Disadvantage

Advantages

- *Similar to RAID 0*
 - Stripe whole or parts of logical volumes across multiple PVs
- *Similar to RAID 1*
 - Mirror whole or parts of logical volumes
- *Move online*
 - Move online logical volumes between PVs



LVM Advantage & Disadvantage

Disadvantages

- The volume managers *can complicate the boot process* and *make disaster recovery difficult*
- Logical volumes can *suffer from external fragmentation* when the *underlying storage devices do not allocate their PEs contiguously*



Question & Answer

