



Unit 8

System storage overview



Unit objectives

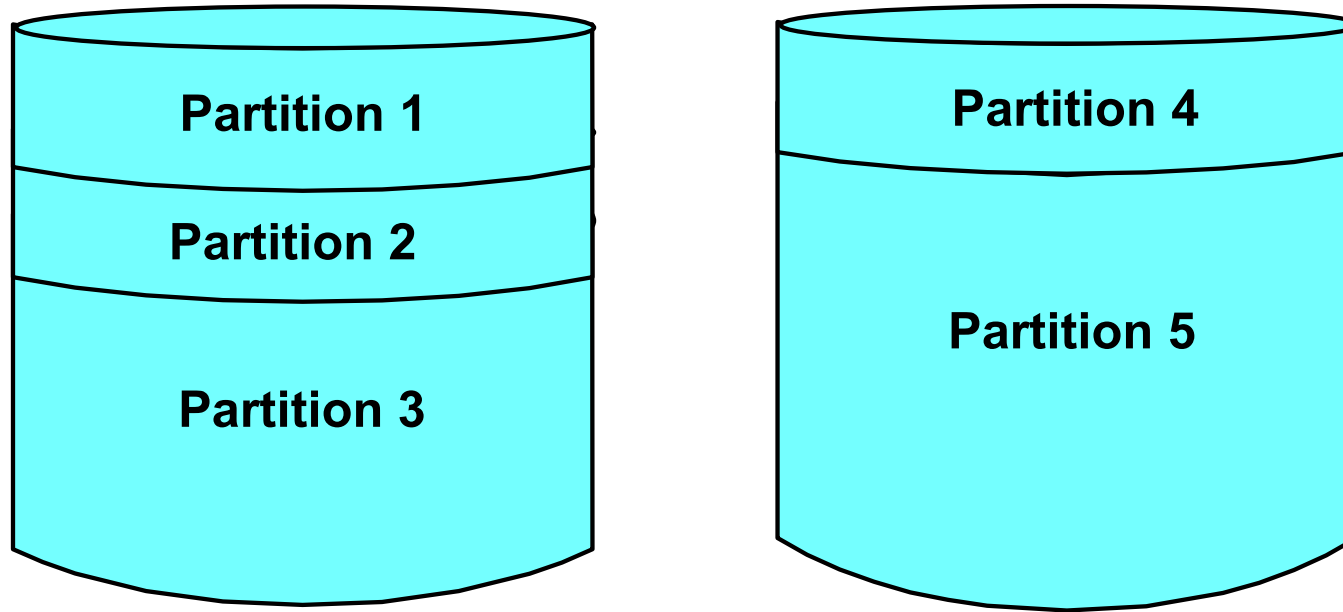
After completing this unit, you should be able to:

- Describe the terminology and concepts associated with:
 - Physical volumes
 - Volume groups
 - Logical volumes
 - Physical partitions
 - Logical partitions
- Describe how file systems and logical volumes are related

Components of AIX storage

- Files
- Directories
- File systems
- Logical storage
- Physical storage
- Logical Volume Manager (LVM)

Traditional UNIX disk storage



PROBLEMS:

- Fixed partitions
- Expanding size of the partition
- Limitation on size of a file system and a file
- Contiguous data requirement
- Time and effort required in planning ahead

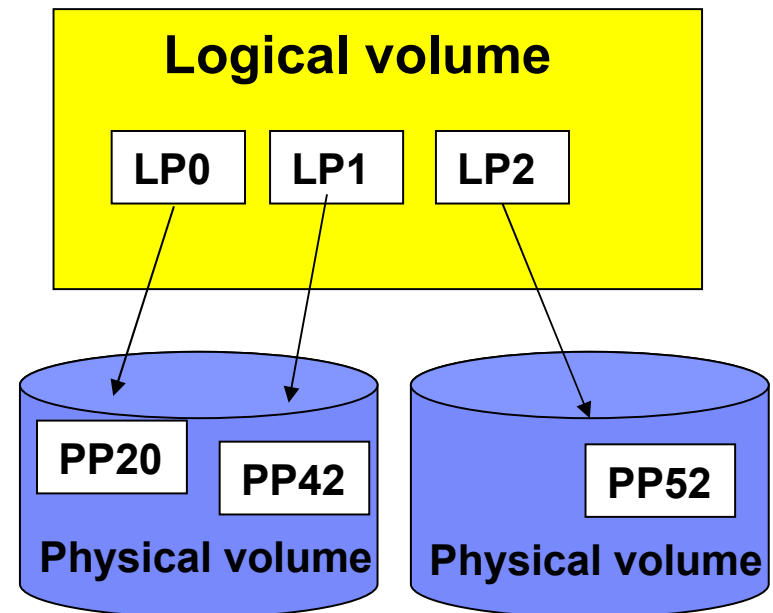
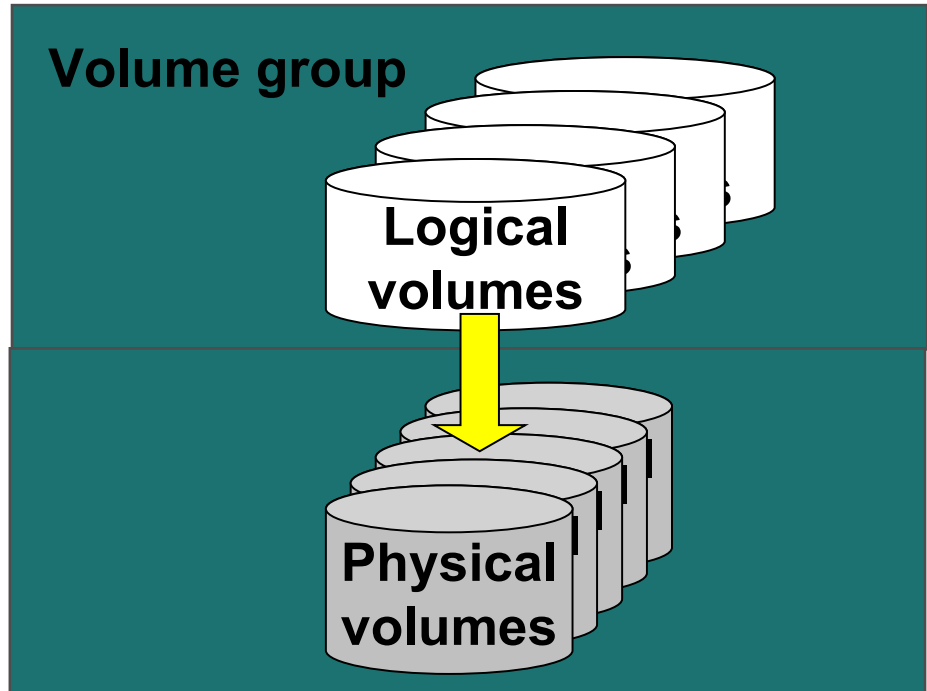
Benefits of the LVM

- Logical volumes solve noncontiguous space problems
- Logical volumes can span disks
- Logical volume sizes can be dynamically increased
- Logical volumes can be mirrored
- Physical volumes are easily added to the system
- Logical volumes can be relocated
- Volume group and logical volume statistics can be collected

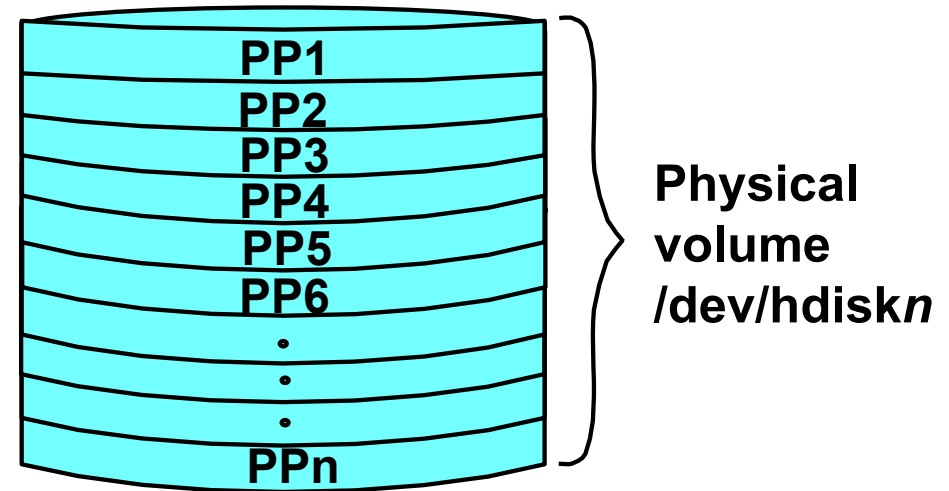
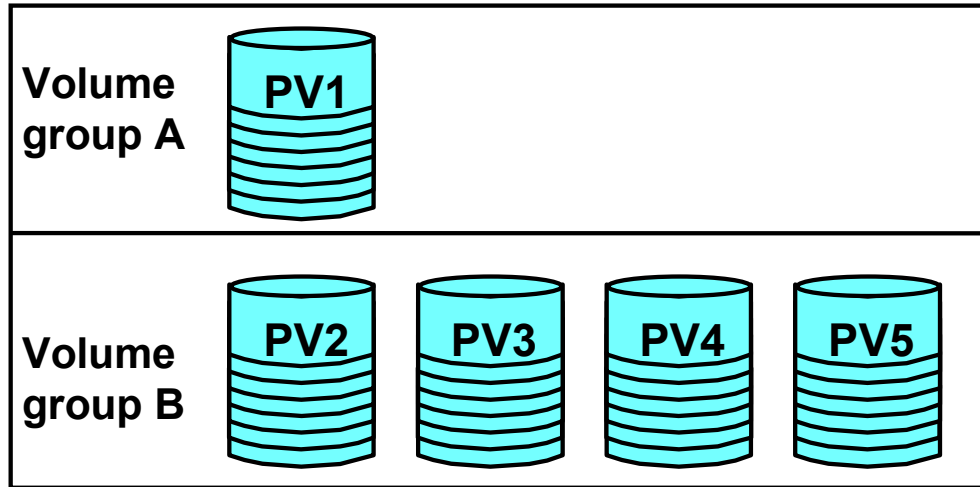
These tasks can be performed dynamically!

Logical Volume Manager components

- Volume group (VG)
- Physical volume (PV)
- Physical partition (PP)
- Logical volume (LV)
- Logical partition (LP)



Physical storage



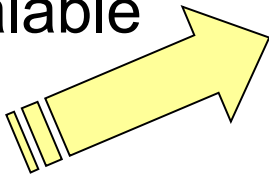
Normal volume groups			Big volume groups	
-t factor	Disks (PVs)	PPs per PV	Disks (PVs)	-t factor
1	32	1016	128	1
2	16	2032	64	2
4	8	4064	32	4
8	4	8128	16	8
16	2	16256	8	16
N/A	N/A	32512	4	32
N/A	N/A	65024	2	64

Volume groups

- Volume group types:

- Normal
- Big
- Scalable

- Limits



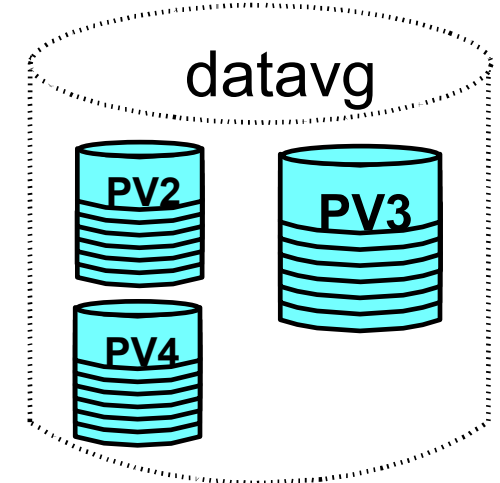
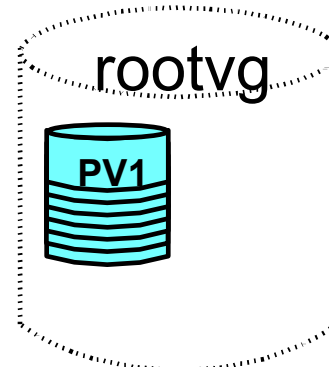
Volume Group Type	Max PVs	Max LVs	Max PPs per VG	Max PP Size
Normal	32	256	32512 (1016 * 32)	1 GB
Big	128	512	130048 (1016 * 128)	1 GB
Scalable	1024	4096	2097152	128 GB

- New physical volumes:

- Add to existing volume groups
- Create new volume group

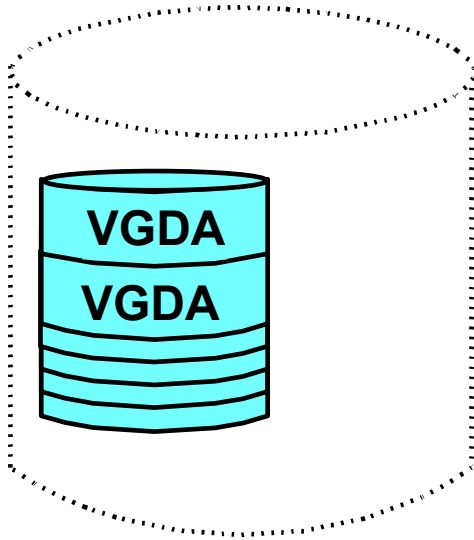
- Why create new volume groups?

- Separate user data from operating system files
- Disaster recovery
- Data portability
- Data integrity and security

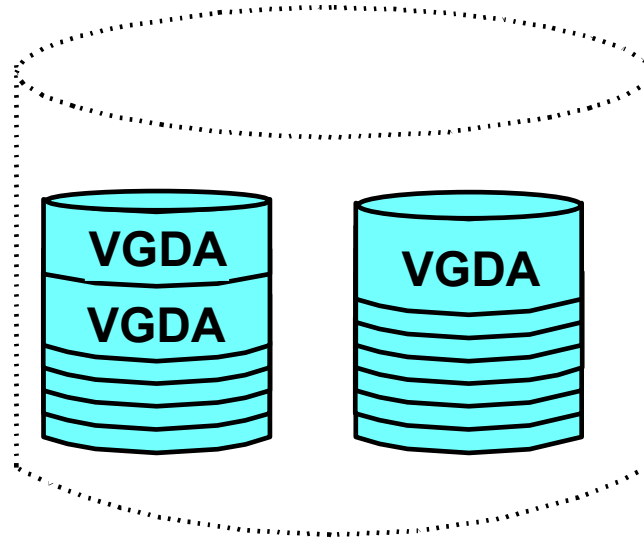


Volume group descriptor area (VGDA)

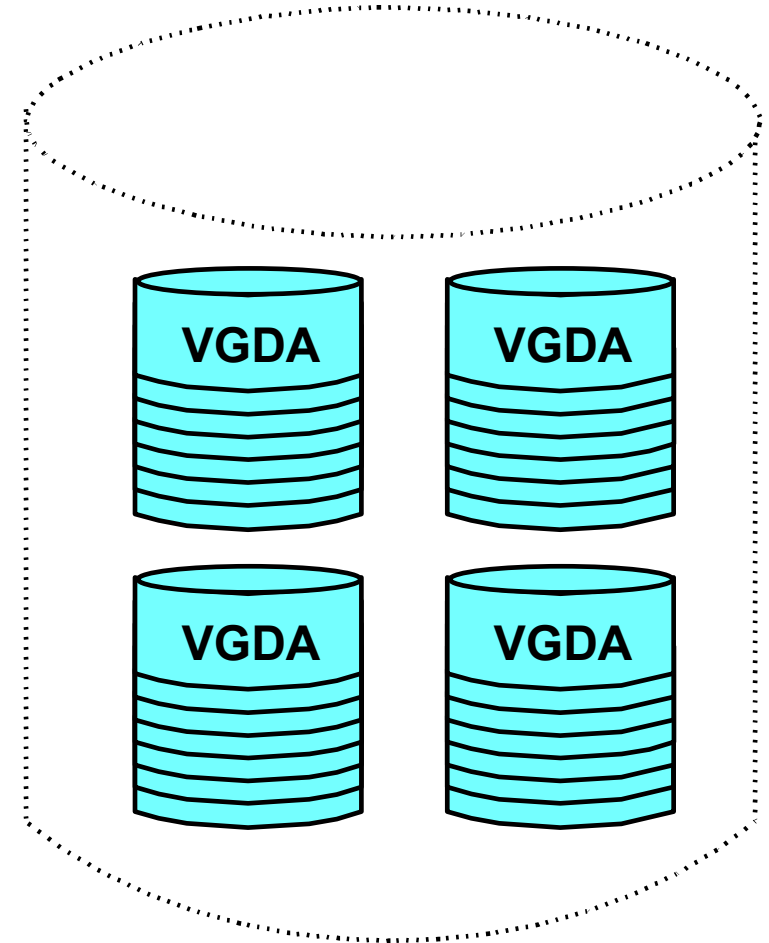
One disk VG



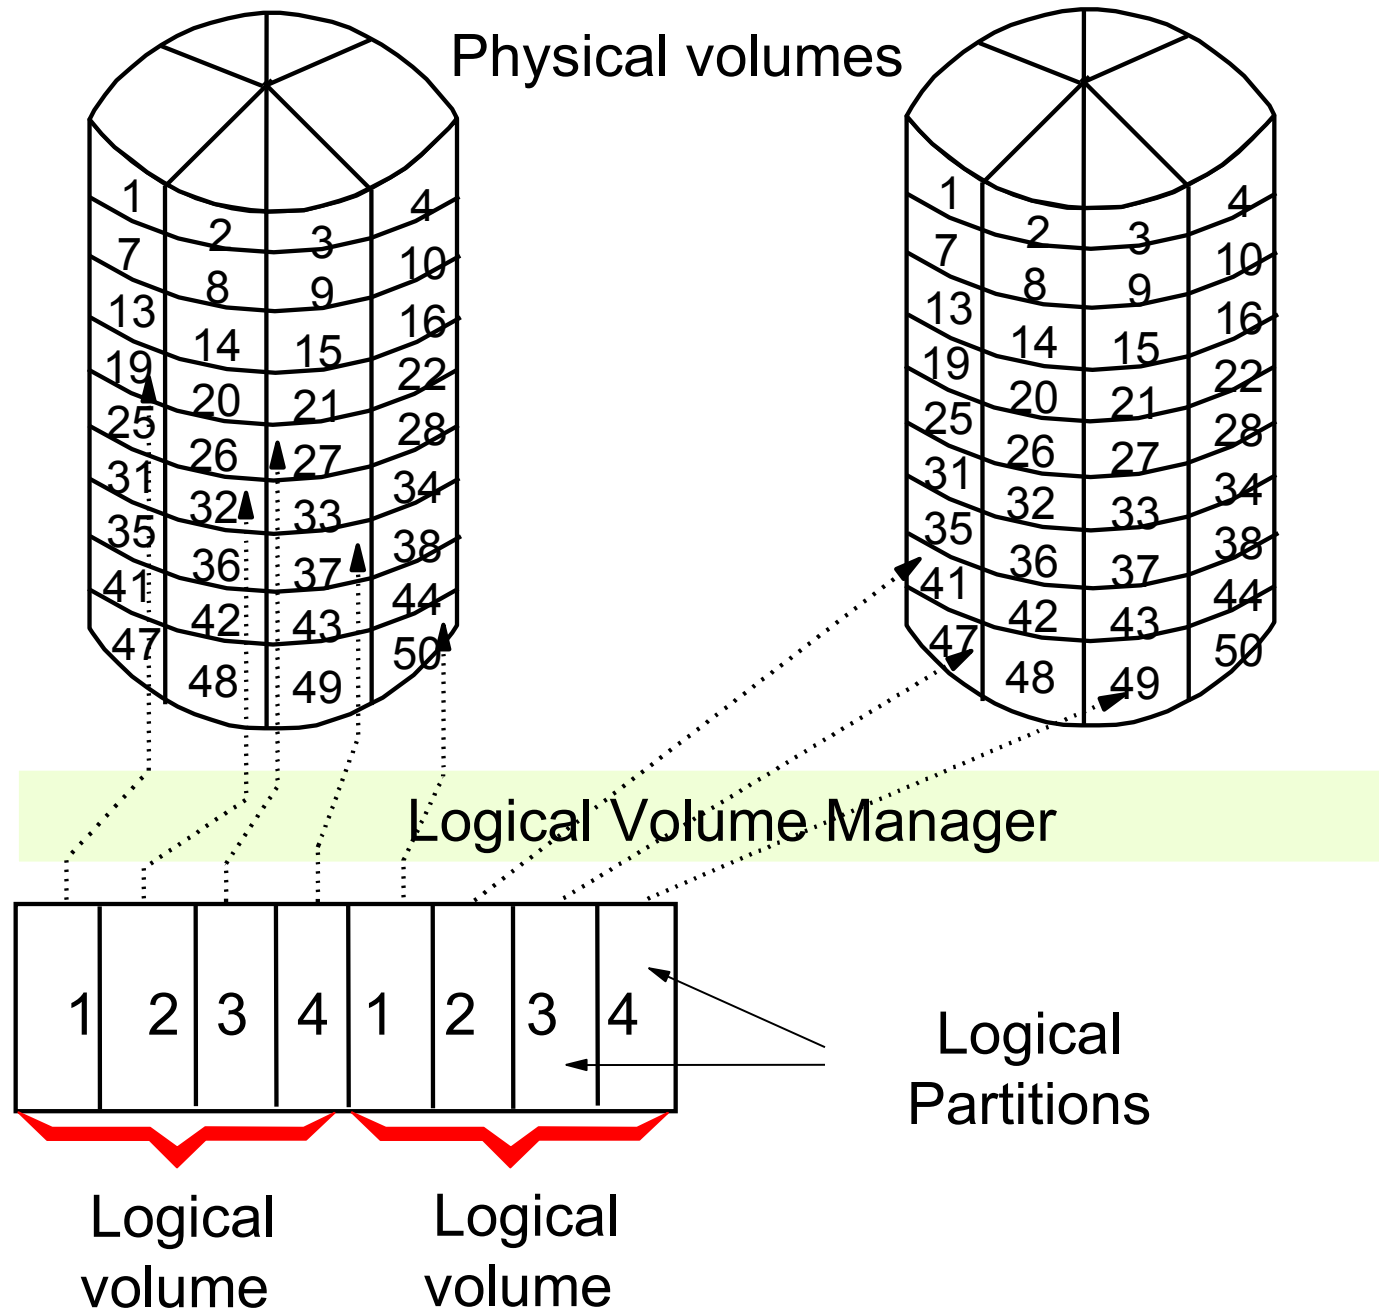
Two disk VG



Three disk or more VG



Logical storage



Uses of logical volumes

- A logical volume may contain one of the following, and only one at a time:
 - Journaled (JFS) or enhanced journaled file system (JFS2)
 - Journal log (**/dev/hd8**)
 - Paging space (**/dev/hd6**)
 - Boot logical volume (**/dev/hd5**)
 - Dump device
 - Nothing (raw logical volume)

- Examples of JFS/JFS2 logical volumes:

/dev/hd1	/home
/dev/hd2	/usr
/dev/hd3	/tmp
/dev/hd4	/
/dev/hd9var	/var
/dev/hd10opt	/opt
/dev/hd11admin	/admin
/dev/lv00	/myfilesystem

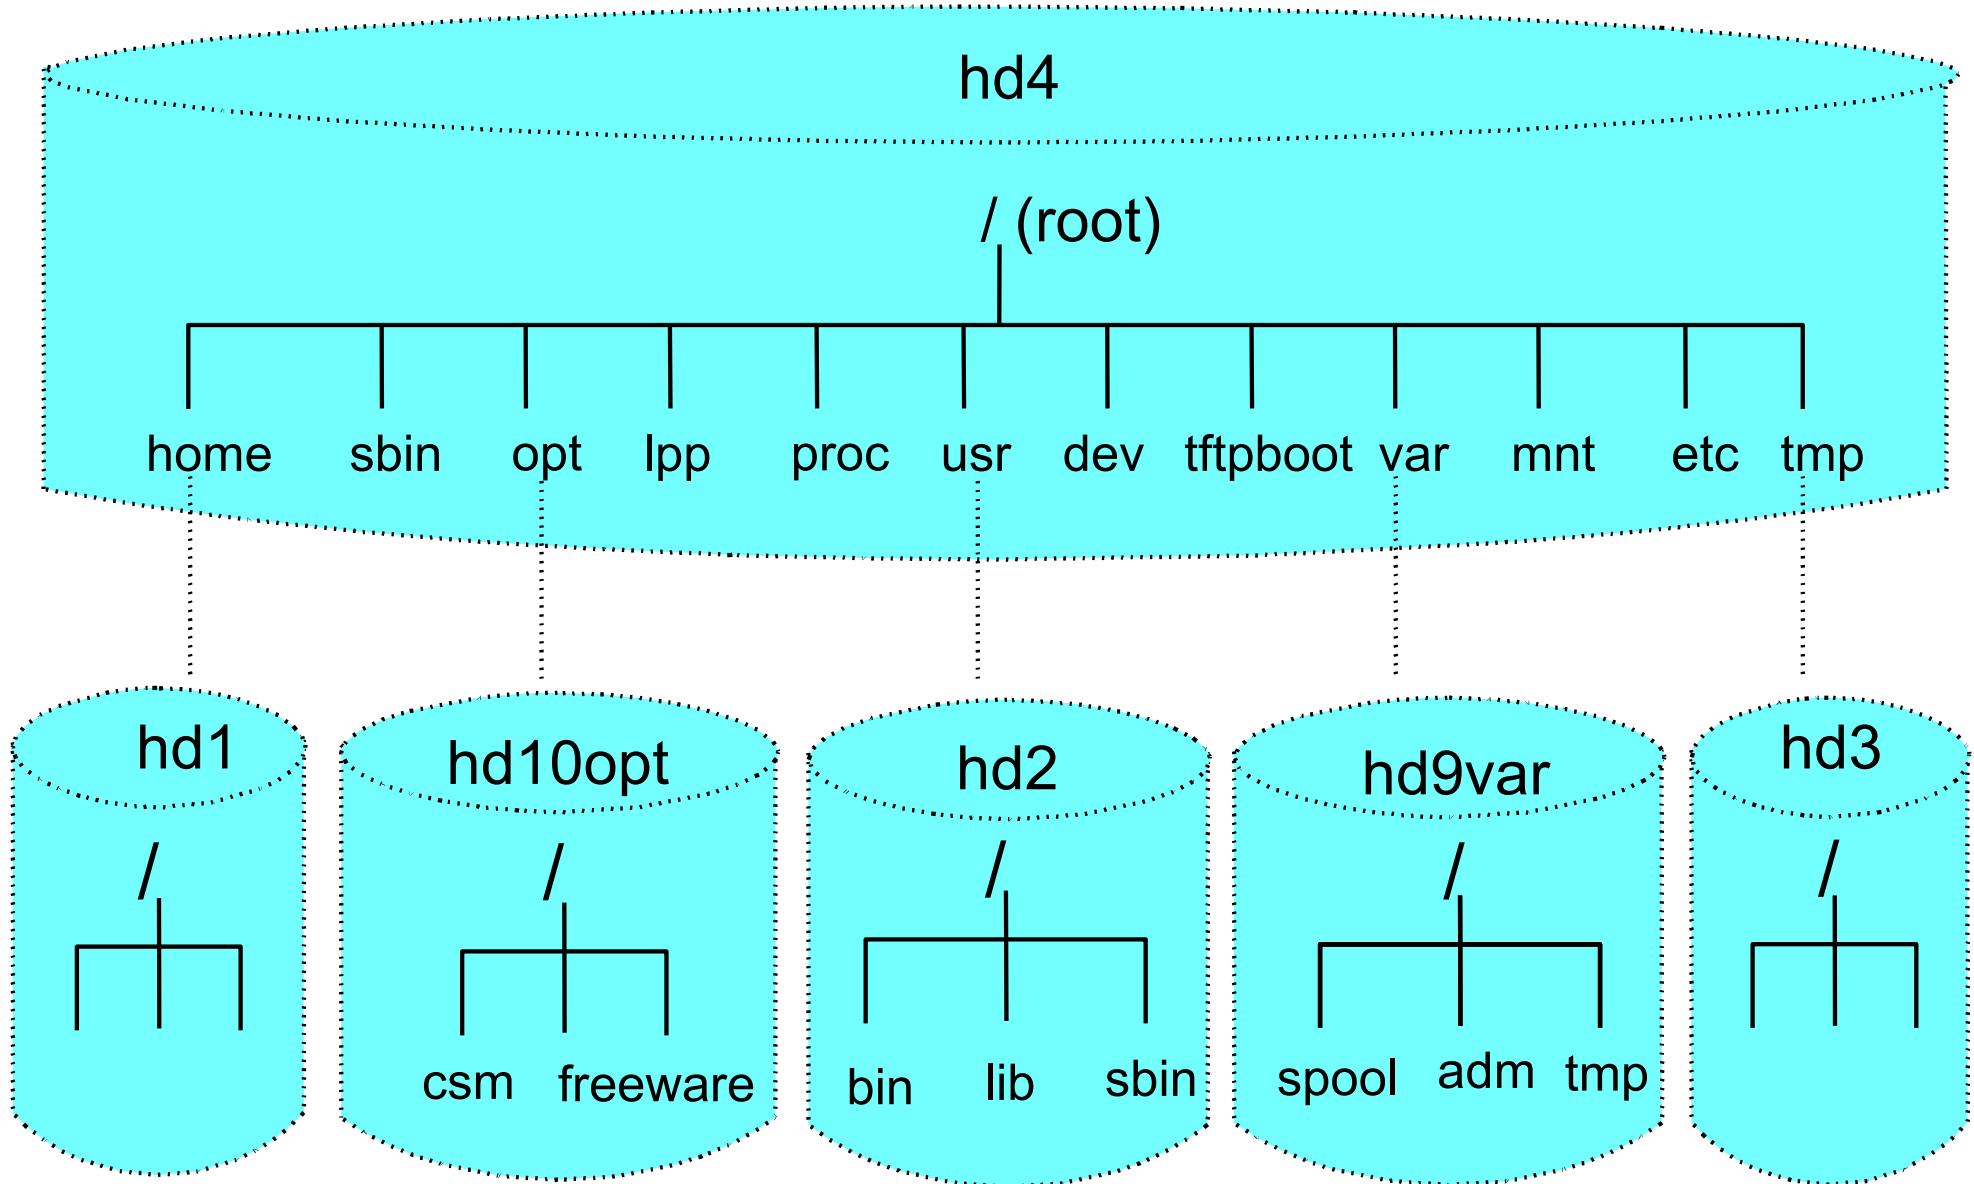
What is a file system?

- A file system is:
 - Method of storing data
 - Hierarchy of directories
- Seven types supported:
 - Journaled File System (JFS)
 - Enhanced Journaled File System (JFS2)
 - CD-ROM File System (CDRFS)
 - DVD-ROM File System (UDFS)
 - Network File System (NFS)
 - Common Internet Filesystem (CIFS)
 - Proc File System (PROCFS)
- Different file systems are connected together via directories to form the view of files users see

Why have multiple file systems?

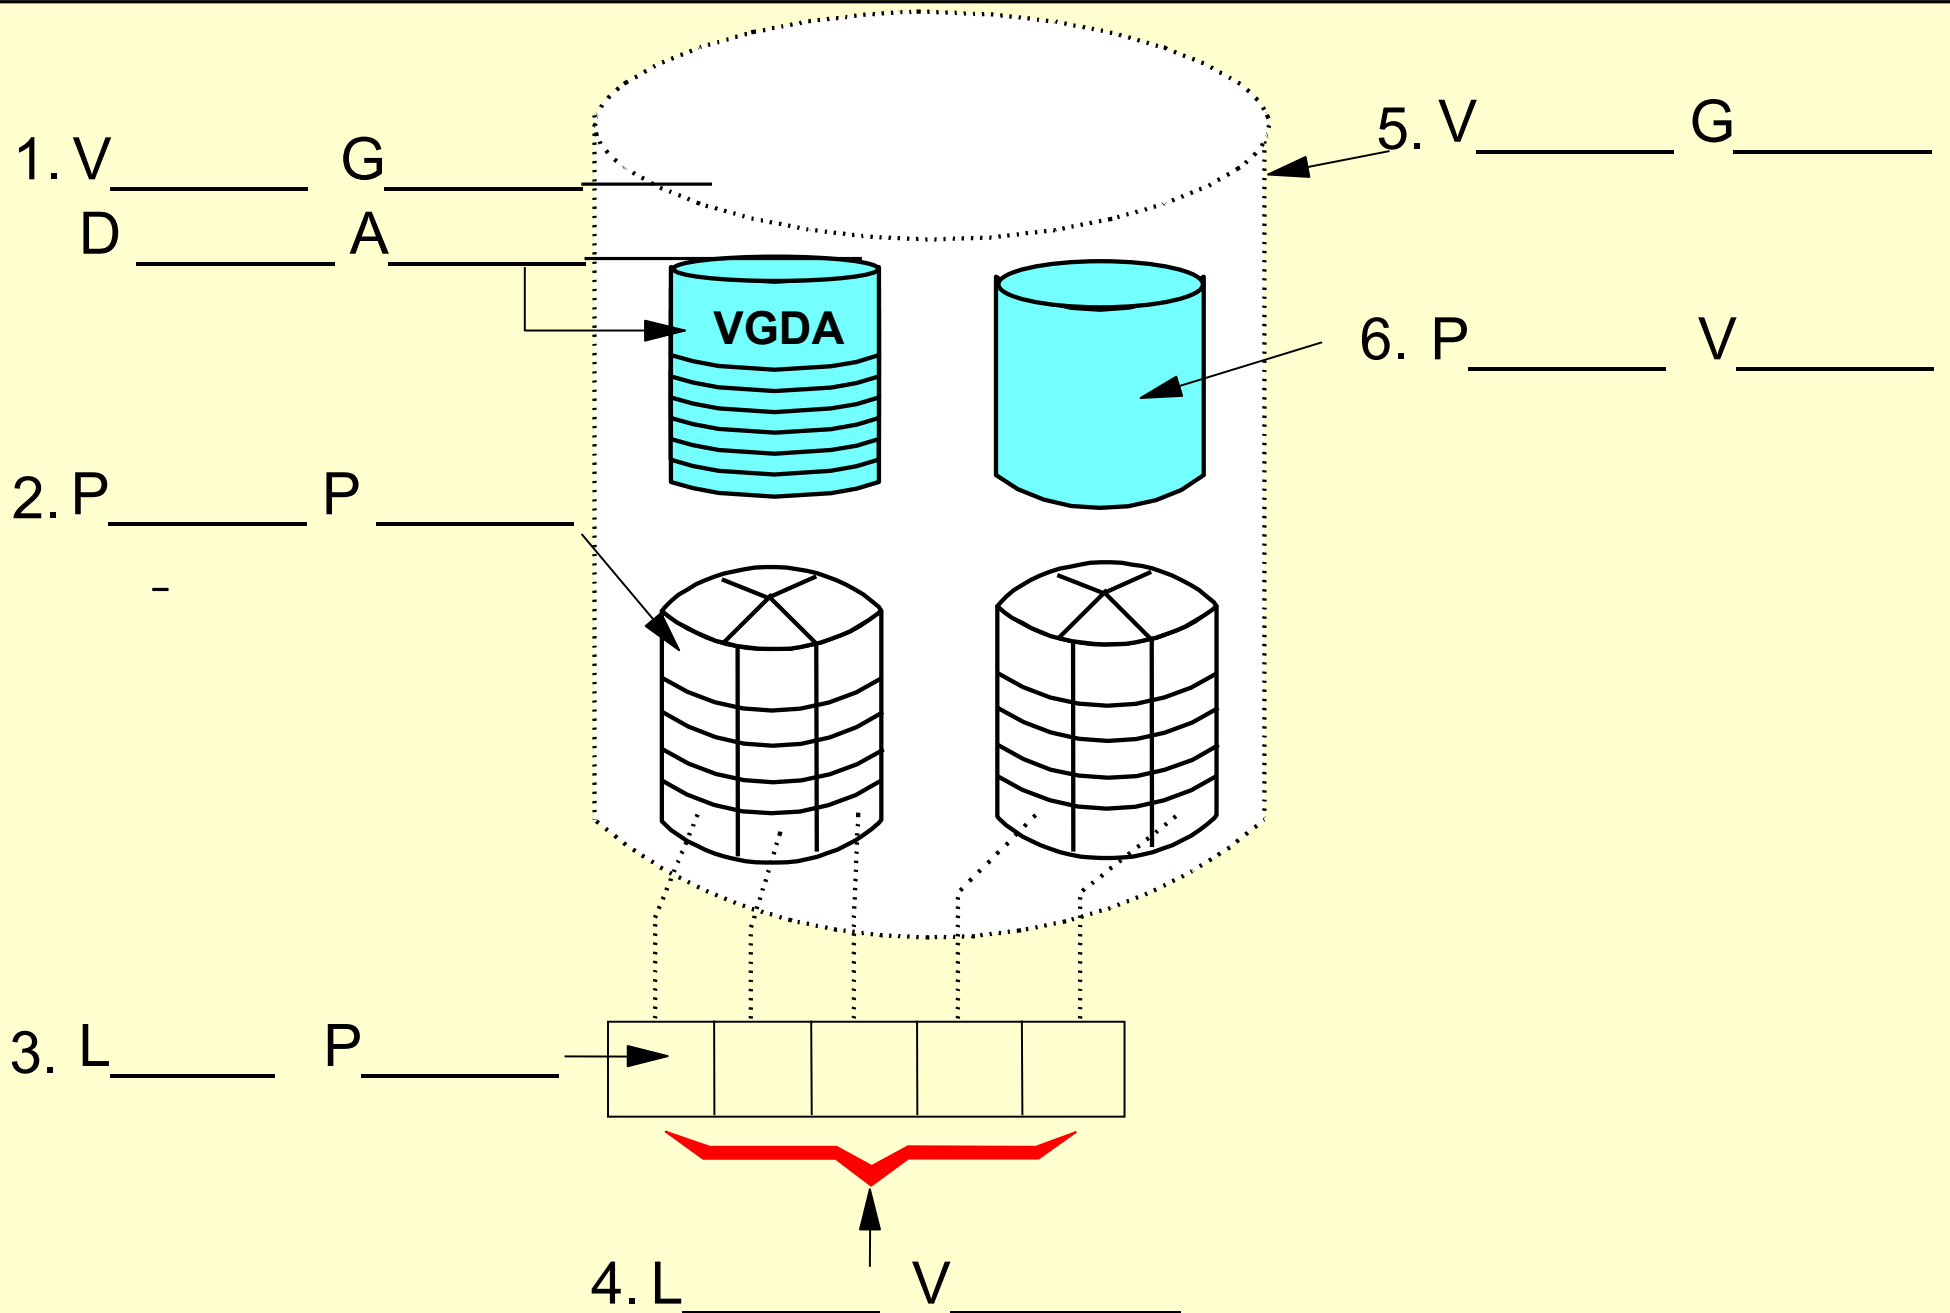
- Can strategically place it on disk for improved performance
- Some tasks are performed more efficiently on a file system than on each directory within the file system, for example, back up, move, secure an entire file system
- Can limit disk usage of users by file system (quotas)
- Maintain integrity of the entire file system structure, for example, if one file system is corrupted, the others are not affected
- Special security situations
- Organize data and programs into groups for ease of file management and better performance

Standard file systems in AIX

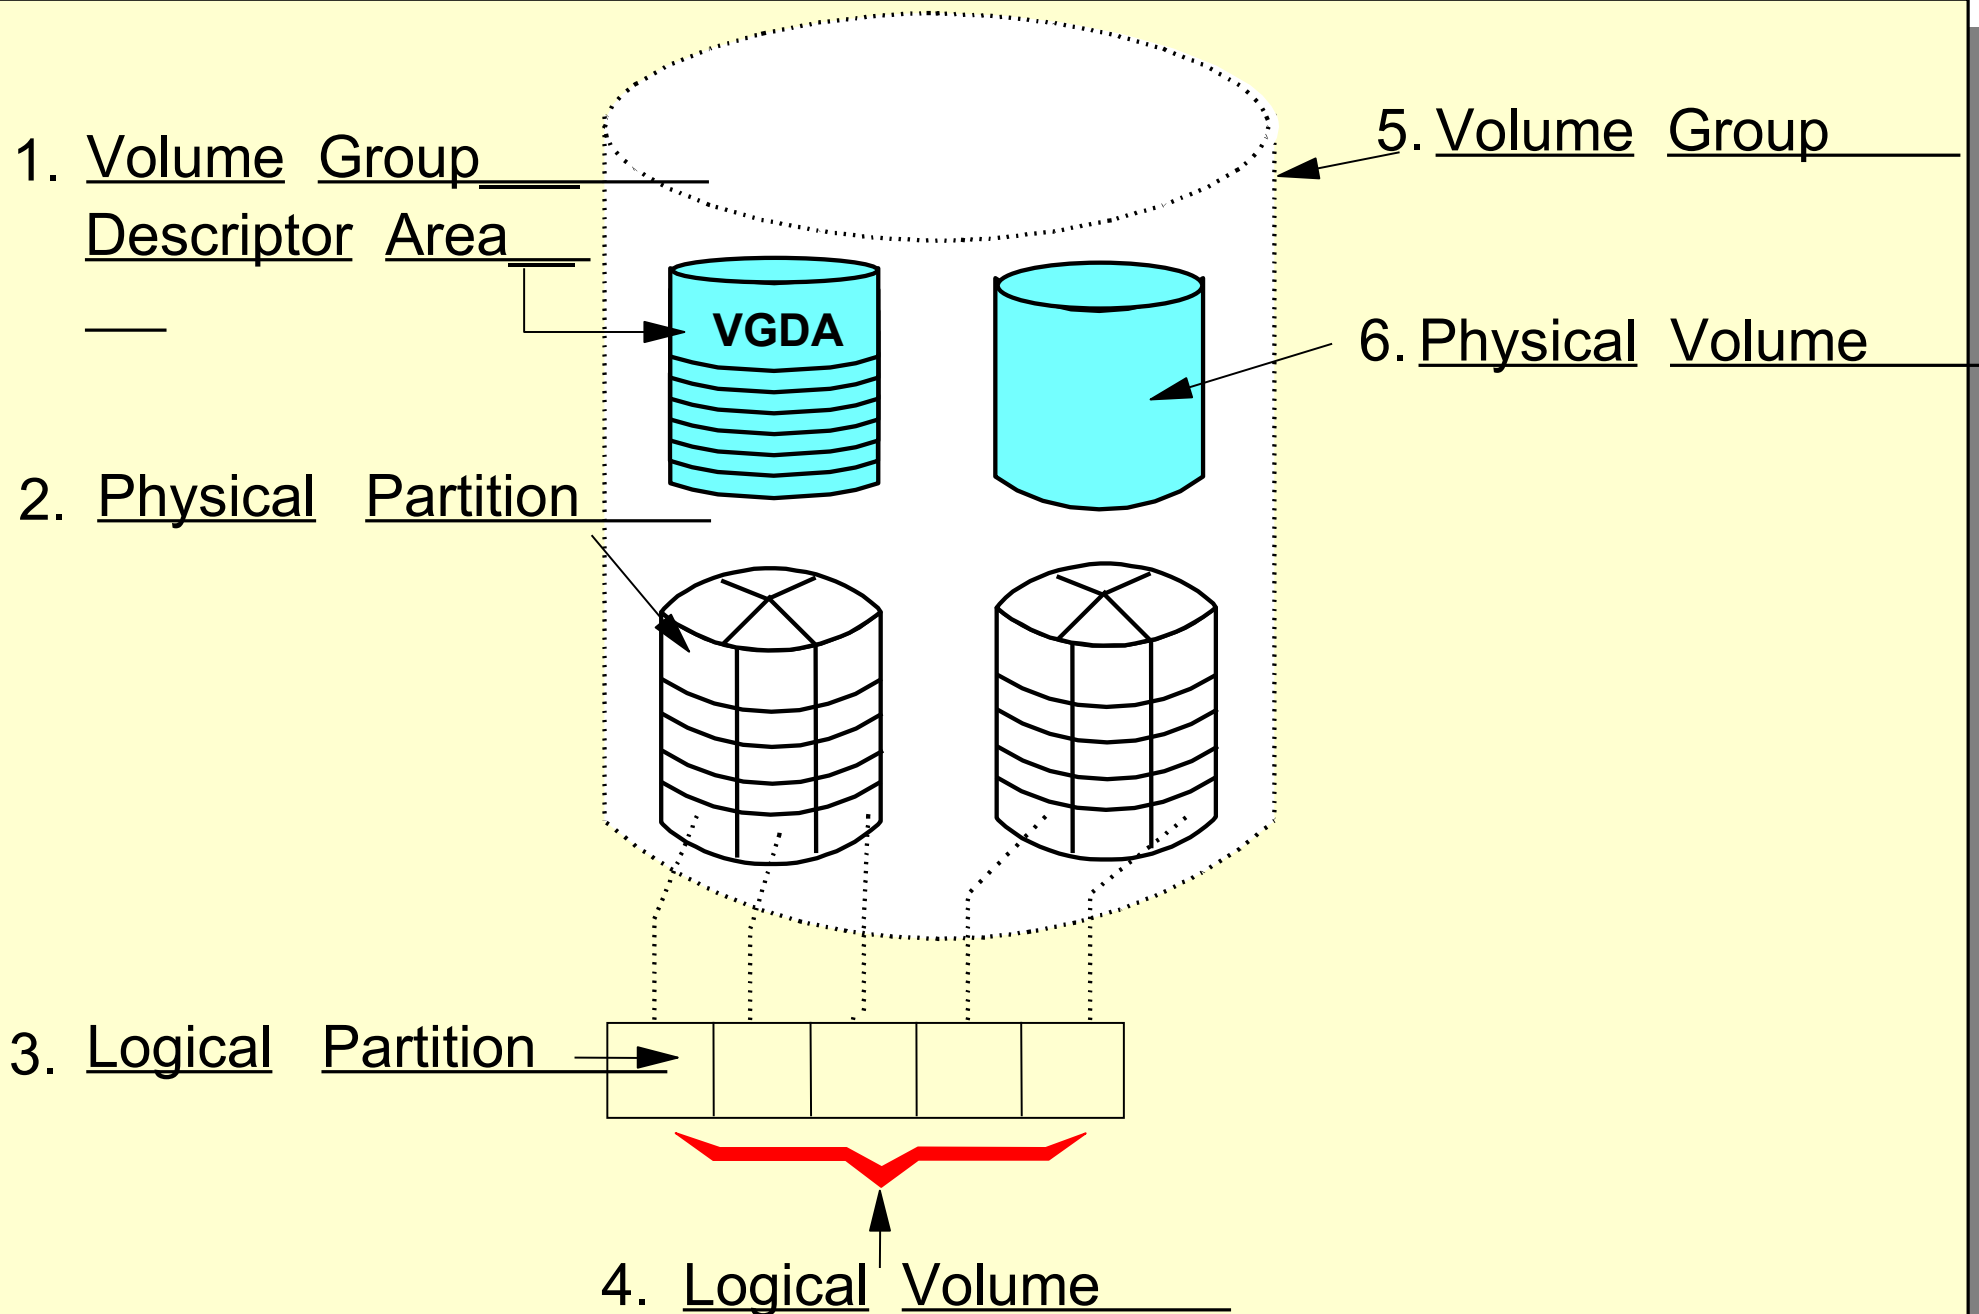


NOTE: The drawing depicts logical not physical volumes

Let's review



Let's review solution



/etc/filesystems

/:

```
dev      = /dev/hd4
vol      = root
mount    = automatic
check    = false
vfs      = jfs2
log      = /dev/hd8
type     = bootfs
```

/home:

```
dev      = /dev/hd1
vol      = /home
mount    = true
check    = true
vfs      = jfs2
log      = /dev/hd8
```

/home/team01:

```
dev      = /dev/fs1v00
vfs      = jfs2
log      = /dev/log1v00
mount    = true
options  = rw
account  = false
```

Mount

- **mount** is the glue that logically connects file systems to the directory hierarchy
- File systems are associated with devices represented by special files in **/dev** (the logical volume)
- When a file system is mounted, the logical volume and its contents are connected to a directory in the hierarchical tree structure

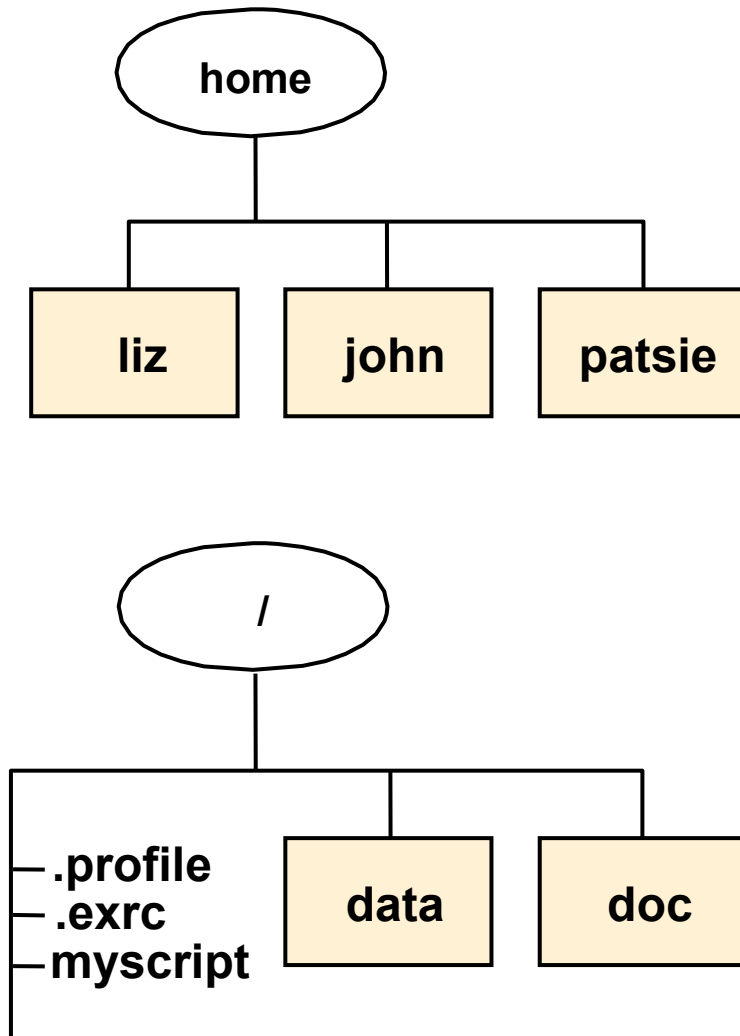
```
# mount /dev/lv00 /home/patsie
```

What to
mount

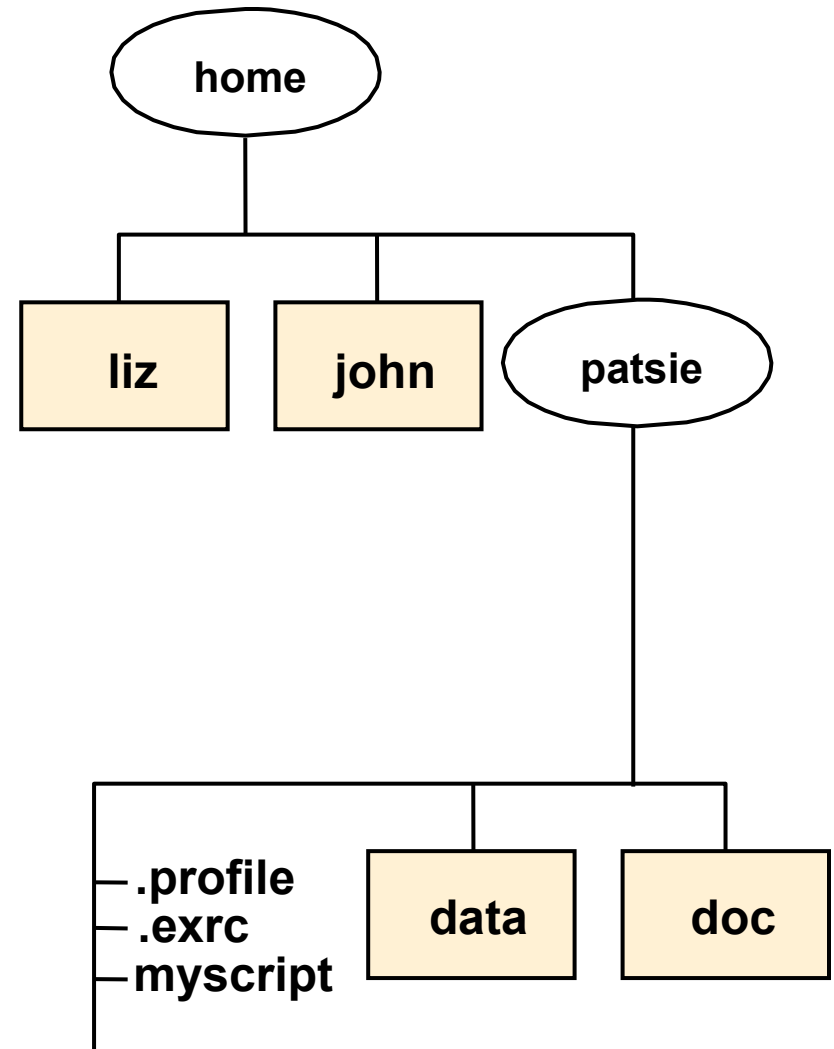
Where to
mount it

Mounting over an empty directory

Before...

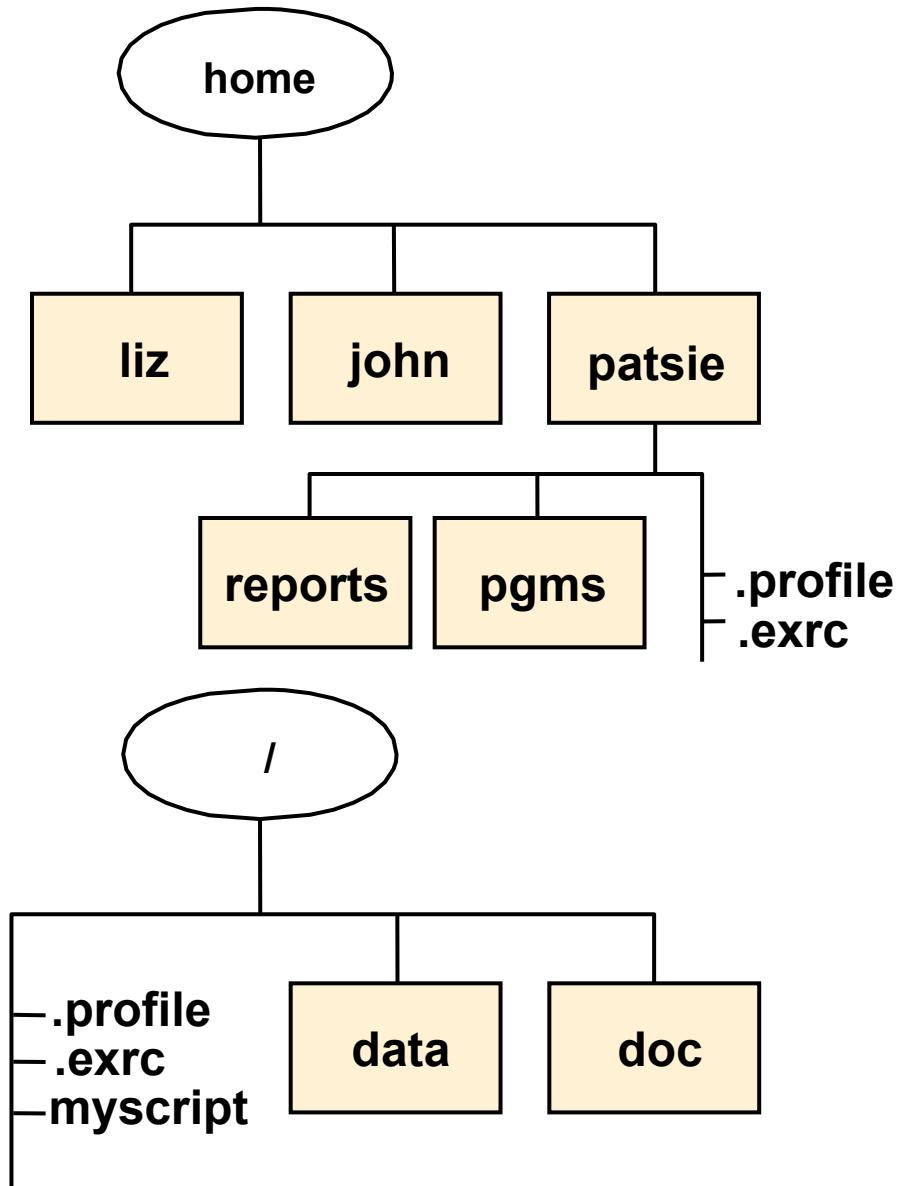


After...

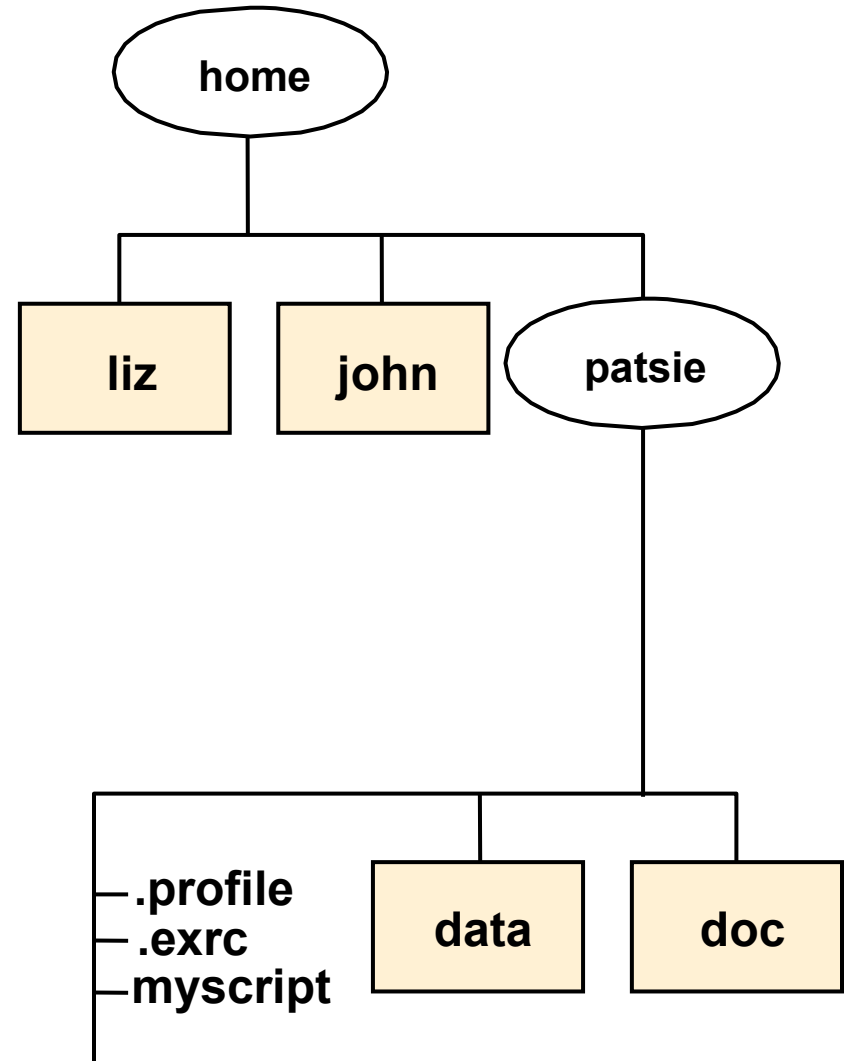


Mounting over files

Before...



After...



Listing file systems

lsfs

Name	Nodename	Mount Pt	VFS	Size	Options	Auto	Accounting
/dev/hd4	--	/	jfs2	294912	--	yes	no
/dev/hd1	--	/home	jfs2	32768	--	yes	no
/dev/hd2	--	/usr	jfs2	3309568	--	yes	no
/dev/hd9var	--	/var	jfs2	65536	--	yes	no
/dev/hd3	--	/tmp	jfs2	131072	--	yes	no
/dev/cd0	--	/infocd	cdrfs	--	ro	yes	no
/dev/lv00	--	/home/john	jfs2	32768	rw	yes	no
/proc	--	/proc	procfs	--	--	yes	no
/dev/hd10opt	--	/opt	jfs2	65536	--	yes	no
/dev/hd11admin	--	/admin	jfs2	262144	--	yes	no

Listing logical volume information

List all logical volumes for a volume group

```
# lsvg -l rootvg
```

LVNAME	TYPE	LPs	PPs	PVs	LV STATE	MOUNT POINT
hd6	paging	32	32	1	open/syncd	N/A
hd5	boot	2	2	1	closed/syncd	N/A
hd8	jfslog	1	1	1	open/syncd	N/A
hd4	jfs2	9	9	1	open/syncd	/
hd2	jfs2	101	101	1	open/syncd	/usr
hd9var	jfs2	2	2	1	open/syncd	/var
hd3	jfs2	4	4	1	open/syncd	/tmp
hd1	jfs2	1	1	1	open/syncd	/home
hd10opt	jfs2	5	5	1	open/syncd	/opt
hd11admin	jfs2	8	8	1	open/syncd	/admin

Checkpoint (1 of 3)

- How many different physical partition (PP) sizes can be set within a single VG? _____
 - By default, how big are PPs?

5. How many volume groups (VGs) can a physical volume (PV) belong to?
- a) Depends on what you specify through SMIT
 - b) Only one
 - c) As many VGs as exist on the system
6. True or False? All VGDA information on your system is identical, regardless of how many volume groups (VGs) exist.

Checkpoint solutions (1 of 3)

- How many different physical partition (PP) sizes can be set within a single VG? One
- By default, how big are PPs? Traditionally 4 MB, but LVM chooses an optimal size based on the #PPs/PV and the size of largest PV in the VG.
- How many volume groups (VGs) can a physical volume (PV) belong to?
 - Depends on what you specify through SMIT
 - ☒ **Only one**
 - c) As many VGs as exist on the system
- True or False? All VGDA information on your system is identical, regardless of how many volume groups (VGs) exist. False. All VGDA's within a VG are the same.

Checkpoint (2 of 3)

Use the following output to answer the questions below:

```
# lsfs
Name          Nodename  Mount Pt  VFS      Size    Options  Auto    Accounting
/dev/hd4      --        /         jfs2     294912  --       yes     no
/dev/hd1      --        /home     jfs2     32768   --       yes     no
/dev/hd2      --        /usr      jfs2     3309568 --       yes     no
/dev/hd9var   --        /var      jfs2     65536   --       yes     no
/dev/hd3      --        /tmp      jfs2     131072  --       yes     no
/dev/hd10opt  --        /opt      jfs2     163840  --       yes     no
/dev/cd0      --        /infocd   cdrfs    --      ro       yes     no
/dev/lv00     --        /home/john jfs2     32768   rw       yes     no
/dev/hd11admin--  /admin    jfs2     262144  --       yes     no
```

- With which logical volume is the **/home** file system associated? _____
- What type of file systems are being displayed? _____
- What is the mount point for the file system located on the **/dev/lv00** logical volume? _____
- Which are the system supplied logical volumes and their associated file systems? _____
- Which file system is used primarily to hold user data and home directories? _____

Checkpoint solutions (2 of 3)

Use the following output to answer the questions below:

```
# lsfs
Name          Nodename  Mount Pt  VFS      Size  Options  Auto  Accounting
/dev/hd4      --        /         jfs2     294912 --        yes   no
/dev/hd1      --        /home     jfs2     32768  --        yes   no
/dev/hd2      --        /usr      jfs2     3309568 --        yes   no
/dev/hd9var   --        /var      jfs2     65536  --        yes   no
/dev/hd3      --        /tmp      jfs2     131072 --        yes   no
/dev/hd10opt  --        /opt      jfs2     163840 --        yes   no
/dev/cd0      --        /infocd   cdrfs                ro    yes   no
/dev/lv00     --        /home/john jfs2     32768  rw       yes   no
/dev/hd11admin--        /admin    jfs2     262144 --        yes   no
```

- With which logical volume is the **/home** file system associated? [/dev/hd1](#)
- What type of file systems are being displayed? [Journalized file systems \(JFS\), enhanced journaled file systems \(JFS2\), and CD-ROM \(CDRFS\)](#)
- What is the mount point for the file system located on the **/dev/lv00** logical volume? [/home/john](#)
- Which are the system supplied logical volumes and their associated file systems? [/dev/hd4 \(/\), /dev/hd1 \(/home\), /dev/hd2 \(/usr\), /dev/hd9var \(/var\), /dev/hd3 \(/tmp\), /dev/hd10opt \(/opt\)](#)
- Which file system is used primarily to hold user data and home directories? [/home](#)

Checkpoint (3 of 3)

Use the following output to answer the question below:

```
# lsvg -l rootvg
```

LVNAME	TYPE	LPs	PPs	PVs	LV State	MOUNT POINT
hd6	paging	8	8	1	open/syncd	N/A
hd5	boot	1	1	1	closed/syncd	N/A
hd8	jfslog	1	1	1	open/syncd	N/A
hd9var	jfs2	1	1	1	open/syncd	/var
hd3	jfs2	2	2	1	open/syncd	/tmp
lv00	jfs2	1	1	1	closed/syncd	/home/john

15. Which of the logical volumes above are examples of logical volumes with journaled file systems on them?

Checkpoint solutions (3 of 3)

Use the following output to answer the question below:

```
# lsvg -l rootvg
```

LVNAME	TYPE	LPs	PPs	PVs	LV State	MOUNT POINT
hd6	paging	8	8	1	open/syncd	N/A
hd5	boot	1	1	1	closed/syncd	N/A
hd8	jfslog	1	1	1	open/syncd	N/A
hd9var	jfs2	1	1	1	open/syncd	/var
hd3	jfs2	2	2	1	open/syncd	/tmp
lv00	jfs2	1	1	1	closed/syncd	/home/john

15. Which of the logical volumes above are examples of logical volumes with journaled file systems on them?

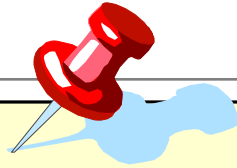
[hd9var, hd3, and lv00](#)

Exercise 8: LVM commands



- Display information about your file systems
- List the LVM information on your system

Unit summary



- The LVM is organized as follows:
 - A volume group consists of one or more physical volumes
 - Each physical volume is divided into physical partitions
 - A logical volume is made up of logical partitions
 - Logical partitions are mapped to physical partitions
- Logical volumes are used to contain:
 - JFS or JFS2 file systems
 - Journal log
 - Paging space
 - Dump space
 - Boot logical volume
 - Raw space
- The most common use of logical volumes is to contain JFS or JFS2 file systems