

Unit 11. Working With the Logical Volume Manager

What This Unit Is About

This unit provides information on how to work with logical volumes, physical volumes and volume groups.

What You Should Be Able to Do

After completing this unit, you should be able to:

- Add/Change/Delete Volume Groups
- Add/Change/Delete Logical Volumes
- Add/Change/Delete Physical Volumes
- Describe mirroring
- Describe striping

How You Will Check Your Progress

Accountability:

- Checkpoint questions
- Exercises

References

AIX Version 4/5 System Management Guide: Operating System and Devices
AIX Storage Management

Objectives

After completing this unit, students should be able to:

- Add/Change/Delete Volume Groups
- Add/Change/Delete Logical Volumes
- Add/Change/Delete Physical Volumes
- Describe mirroring
- Describe striping

Figure 11-1. Objectives

Notes:

Logical Volume Manager

Logical Volume Manager

Move cursor to desired item and press enter.

Volume Groups
Logical Volumes
Physical Volumes
Paging Space

F1=Help F2=Refresh F3=Cancel F8=Image
F9=Shell F10=Exit Enter=Do

Figure 11-2. Logical Volume Manager

Notes:

The Logical Volume Manager menu is used to manage many aspects of the system's storage.

- **Volume Groups** - This menu provides facilities to manipulate the volume groups in the system. The "Define a Fixed Disk to the Operating System" menu duplicates some items on this menu.
- **Logical Volumes** - This menu provides facilities to manipulate the logical volumes in the system. Logical volumes which contain journaled file systems, paging space or dump volumes can also be manipulated from their respective menus. However, the facilities on this menu give a much lower level of control over the characteristics of the logical volume. For example, features such as partition allocation policy and mirroring for a logical volume, can only be set using this menu. This menu is also used when a logical volume which will not contain an AIX file system, etc. is being manipulated.
- **Physical Volumes** - This option allows the user to configure the physical volumes (fixed disks) in the system. This menu duplicates options on the "Fixed Disks" menu of "Devices".

- **Paging Space** - This option will allow a user to add, delete, activate and list the paging spaces available.

The Web-based System Manager can also be used to manage the Logical Volume Manager. Use the fastpath **wsm lvm**.

Volume Groups

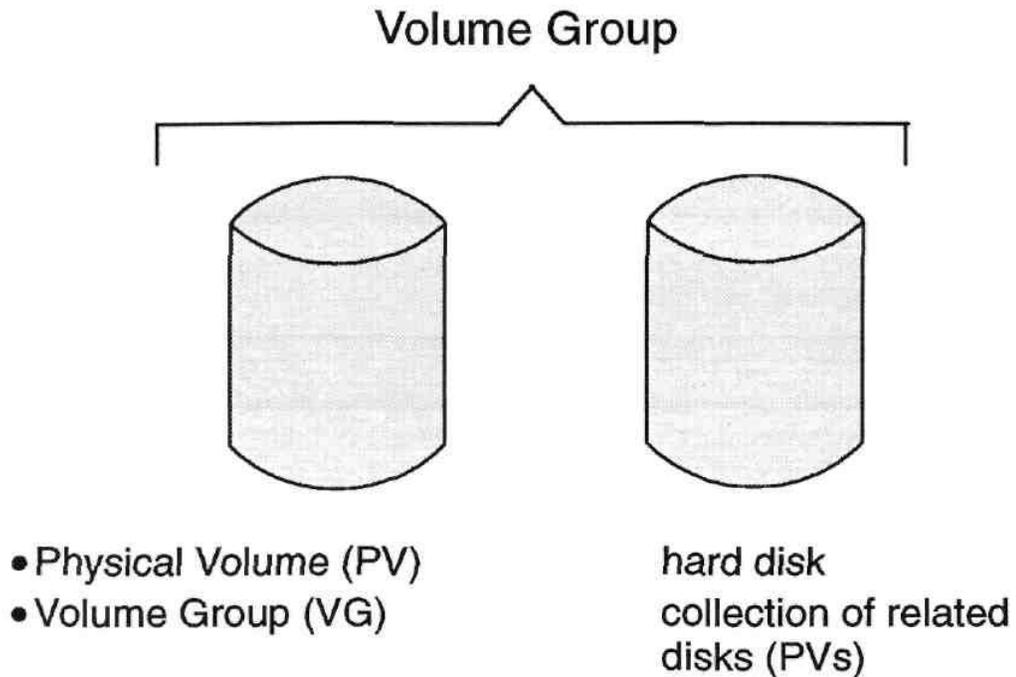


Figure 11-3. Volume Groups

Notes:

Physical Volume - hard disk. There is a limit of up to 32 PVs per volume group.

A **Volume Group** is a collection of related PVs on a processor that:

- are not members of another VG
- share a single physical partition size

When you install your AIX system, one volume group called **rootvg** is automatically created.

There can be a maximum of 255 VGs per system.

A PV that supports removable media should be assigned to a VG containing itself and no other members.

SMIT Volume Groups Menu

smit vg

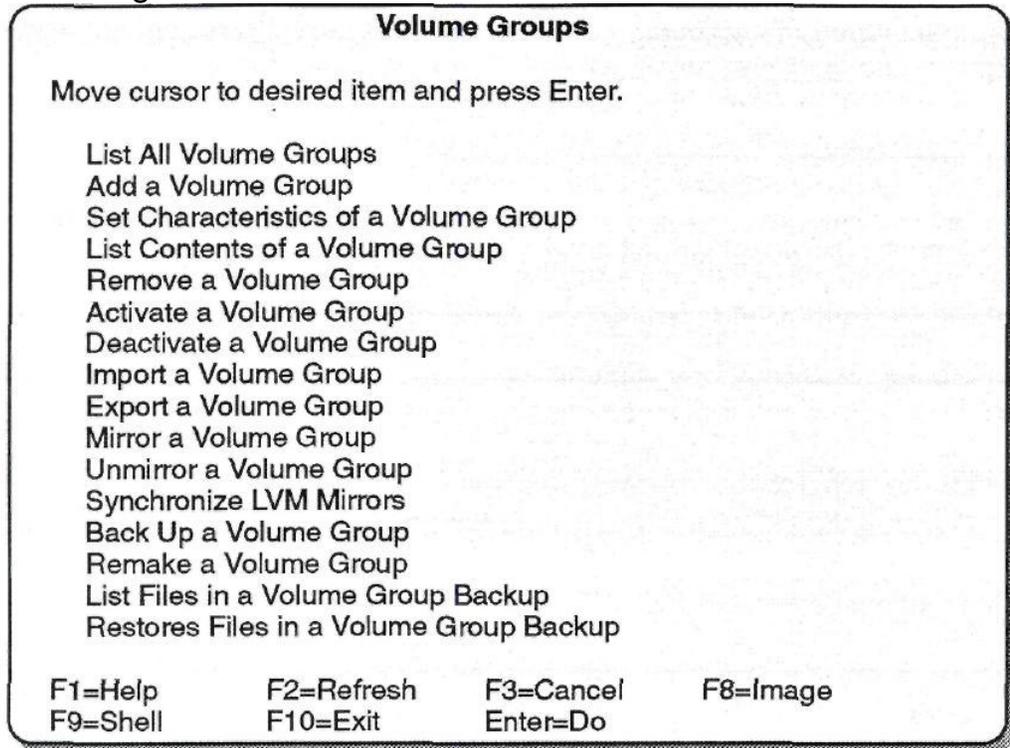


Figure 11-4. SMIT Volume Groups Menu

Notes:

Listing Volume Group Information

List All Volume Groups

```
# lsvg
```

```
rootvg
```

```
payrollvg
```

```
# lsvg -o
```

```
rootvg
```

List Contents of a Volume Group -

Status:

```
# lsvg rootvg
```

Volume GROUP:	rootvg	VG IDENTIFIER:	000008039c386994
VG STATE:	active	PP SIZE:	4MB
VG PERMISSION:	read/write	TOTAL PPs:	515 (2060MB)
MAX LVs:	256	FREE PPs:	281 (1124MB)
LVs:	10	USED PPs:	234 (936MB)
OPEN LVs:	9	QUORUM:	2
TOTAL PVs:	2	VG DESCRIPTORS:	3
STALE PVs:	0	STALE PPs:	0
ACTIVE PVs:	2	AUTO ON:	yes
MAX PPs per PV:	1016	MAX PVs:	32

Figure 11-5. Listing Volume Group Information

Notes:

The **lsvg** command is used to list the volume groups in the system.

It can be used to list the names of all volume groups (default) or only those that are varied on/active (-o).

This command can also be used to list information about the status and content of a particular volume group, for example **lsvg** *Volumegroup*.

The output provides status information about the volume group. The most useful information here is:

State (active or inactive/complete if all PVs are active) PP size
(4MB by default) Total number of PPs Number of free PPs

Instructor Notes:

Purpose —List VG information.

Details — The **lsvg -o** command lists only those volume groups that are *varied on*.

On the SMIT **Volume Group** menu, there is an option, **List Contents of a Volume Group**. This option allows you to view VG information three ways: status, by physical volumes and by logical volumes. We will consider viewing the status first.

Go through each attribute associated with the VG. Rather than telling the students what each attribute sets, ask the following questions:

- How many disks does rootvg have ?

Two

- What is an open LV and why do you think one of the LVs is closed ?

An open LV is an LV that can be accessed and used. In general, all LVs are open except one, the boot LV, which by default is always closed. This LV is used to boot up the system and once the boot is complete it will not be used until the next boot of the system. For this reason, there is no need to have this LV open. Also, if this LV is corrupted then your system might not boot up. So by closing the LV, there is less chance of corruption occurring.

You might feel that this level of detail is too much at this stage. If so, you may wish to discuss this during the LV topic in this unit. Either way, ensure that you do cover this at some point.

- How much free space is there in this VG ?

281 PPs which makes 1124MB

- How many VGDA's should this system have and is this correct ?

There should be three VGDA's as there are two disks in this system.

The last line in the output of the **lsvg rootvg** command applies to AIX V4.3.1 and later.

Additional Information —

Transition Statement —We'll next view VG information by physical volume and by logical volume.

Listing Volume Group Information

- List Contents of a Volume Group -

Physical Volumes: # `lsvg -p`

rootvg

```
rootvg:
PV_NAME PV_STATE TOTAL PPs FREE PPs FREE DISTRIBUTION
hdisk0 active 159 52 24..00..00..00..28
hdisk1 active 159 78 32..02..00..12..32
```

- Logical Volumes: # `lsvg -l rootvg`

```
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 jfs 1 1 1 open/syncd /var
hd4 jfs 1 1 1 open/syncd /
hd2 jfs 77 77 1 open/syncd /usr
hd3 jfs 3 3 1 open/syncd /tmp
hd1 jfs 11 11 1 open/syncd /home
lv00 jfs 1 2 1 open/syncd /home/john
lv01 jfs 4 4 2 open/syncd /home/fred
```

Figure 11-6. Listing Volume Group Information

Notes:

`lsvg -p` *Volume group*

This gives information about all of the physical volumes within the volume group. The information given is: PV Name, PV State (active or inactive), Total number of PPs, Number of free PPs and how the free space is distributed across the disk (edge..middle..center..middle..edge).

Free distribution is the number of physical partitions allocated within each section of the physical volume: outer edge, outer middle, center, inner middle, inner edge

`lsvg -l` *Volume group*

This gives information about all of the logical volumes within the volume group. The details given are: LV name, Type of LV (for example, file system, paging), number of LPs, number of PPs, number of PVs, LV state and, if the LV contains a journaled file system the mount point is also given.

Instructor Notes:

Purpose —List the contents of a volume group.

Details —

The output on this visual is not meant to match what was on the previous visual. Rather, it is meant to show an example of viewing volume group information.

The visual illustrates how information can be obtained per VG with respect to the PVs and the LVs.

The first command example shows the list of PVs in rootvg, how many partitions are on each disk and how many are free. Also, a distribution chart of the free PPs is given.

The second example shows a list of all the logical volumes in rootvg. The first column gives the name of the LV. Point out that all the system logical volumes are labelled **ndx**. Encourage students to recognize the system LVs. The second column describes the type of data contained on each LV, for example journaled file system (jfs). The next two fields give the LP to PP mappings, and it is from these two fields that you can tell whether any LVs have been mirrored or not.

The PV column states how many disks that particular LV is spanning. The second to last field shows the LV state, i.e. whether it is open or not, and notice that the boot LV is the only one that is closed. The syncd attribute states that the LV is up to date. The final field is only applicable for some LVs and it shows to which directory the file system is mounted.

Additional Information — Ask the students whether they know what the free distribution chart is referring to.

Answer: This coincides with the intra-physical allocation policy and the disk bands. From the output of the command it can be seen that for hdisk0 there are 24 partitions free on the outer edge of the disk, and also 28 partitions free on the inner edge. The center and the middle sections are all used.

Ask the students if there are any mirrored LVs. The answer is yes. **/home/john**

Transition Statement — How can we add a VG to our system?

Adding and Removing Volume Groups

smit mkvg

Add a Volume Group

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

	[Entry Fields]	
Volume GROUP name	[]	+
Physical partition SIZE in megabytes	4	+
* PHYSICAL VOLUME names	[]	+
Activate Volume group AUTOMATICALLY at system restart	yes	
Volume Group MAJOR NUMBER	[]	+#
Create VG Concurrent Capable?	no	+
Auto-varyon in Concurrent Mode?	no	+

F1=Help	F2=Refresh	F3=Cancel	F4=List
F5=Reset	F6=Command	F7=Edit	F8=Image
F9=Shell	F10=Exit	Enter=Do	

• Remove a Volume Group

When the last PV is deleted from a VG, then the VG is removed as well.

Figure 11-7. Adding and Removing Volume Groups

Notes:

The **mkvg** command is used to create a volume group. A new volume group must contain at least one physical volume. The **-y** option is used to indicate the name for the new volume group. If this is not specified, a system generated name is used. The **-s** option is used to specify the physical partition size in MB which must be a power of 2. The default is 4 MB.

The **-n** option means that the volume group is not automatically activated at system startup. This should be done for external disks that may not always be available to the system.

The Volume Group MAJOR NUMBER on the SMIT dialog screen is used by the kernel to access that volume group. This field is most often used for High Availability Network File System (HANFS) and High Availability Cluster Management Program (HACMP) applications.

The two items on the SMIT dialog screen referring to concurrent mode operation have no meaning on systems without HACMP installed. These items are valid on AIX V4.2 and later.

The **Remove a Volume Group** menu item does not have a corresponding high-level command. The correct way to remove a volume group is to use the **Remove a Physical Volume from a Volume Group** option (**reducevg** command). This will remove the volume group when you remove the last physical volume within it.

Instructor Notes:

Purpose —Show students how to add and delete a volume group.

Details — After installation, the system has only one VG containing the operating system and that is the **rootvg**. Other physical volumes that are not in a VG can be grouped together into additional VGs.

Note that physical partition size is a value of 1-1024MB, and must be equal to the power of 2. The default size is 4MB. Sizes of 512MB and 1024MB are valid for AIX V4.3.1 and later.

The menu items dealing with concurrent mode are used only if in the HACMP environment, and are supported in AIX V4.2 and later.

Additional Information —

Transition Statement —Now that we can add and delete a VG, how can we change the characteristics of a VG?

Set Characteristics of a Volume Group

```
# smit vgsc
```

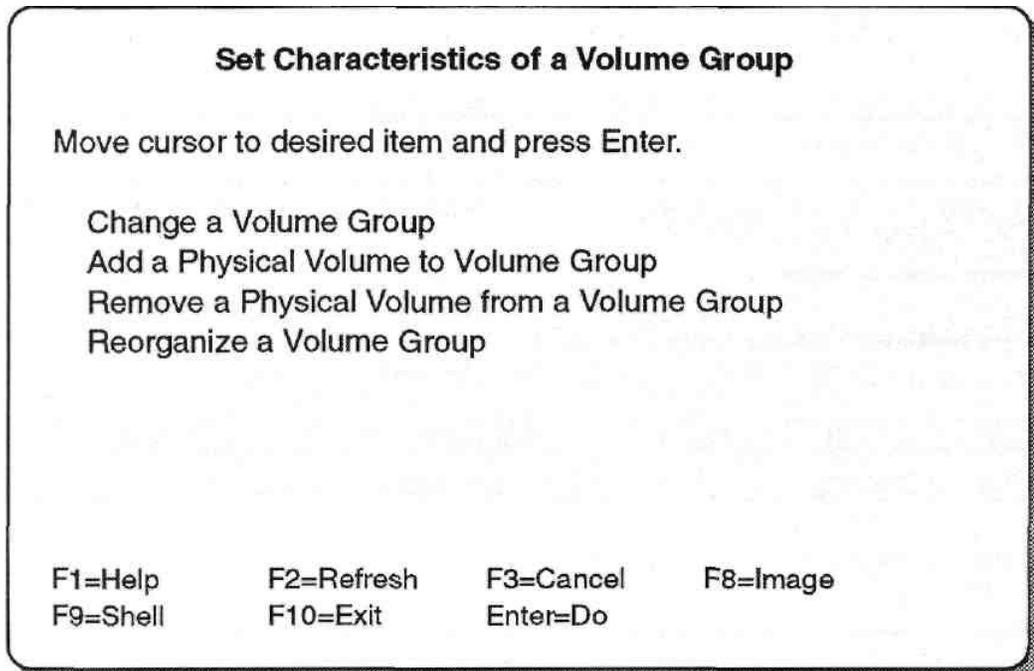


Figure 11-8. Set Characteristics of a Volume Group

Notes:

'Add a Physical Volume to a Volume Group'

extendvg -f Volumegroup hdiskn

The **extendvg** command is used to add a new physical volume to an existing volume group. The fixed disk must be installed in the system or connected to it externally, and must be powered on.

If the existing data on the disk shows that it is part of another volume group, the -f option forces the addition of the disk to the volume group without requesting confirmation. Use this option when adding a disk which has been previously used but which contains no data that is required any further.

'Remove a Physical Volume from a Volume Group'

reducevg [-d] Volumegroup hdiskn

The **reducevg** command is used to remove a physical volume from a volume group. If it is the last physical volume, the volume group will be removed.

Instructor Notes:

Purpose —To show how to add and remove physical volumes in a VG.

Details —This graphic shows four options available for setting the characteristics of a volume group. Discuss the options for adding and removing a physical volume from the volume group. The other two options will be covered on the following pages.

The student notes show a -d option for the **reducevg** command. This option will cause automatic deletion of all logical volumes on the physical volume before removing the physical volume from the volume group. If a logical volume spans multiple physical volumes, the removal of any of those physical volumes may jeopardize the integrity of the entire logical volume. SMIT provides a menu item to indicate if you wish this forced deletion to take place.

Typically, a physical volume may be removed from a volume group if it is damaged and maintenance work needs to be done on it.

Briefly describe the steps that need to be taken to add or remove a physical drive volume from a volume group.

Additional Information —

Transition Statement —Logical volumes can be moved within a volume group to get better compliance with requested allocation policies.

Change a Volume Group

smit chvg_

Change a Volume Group

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

	[Entry Fields]	
* VOLUME GROUP name	[]	
* Activate Volume group AUTOMATICALLY at system restart?	yes	+
* A QUORUM of disks required to keep the Volume group on-line?	yes	+
Convert this VG to Concurrent Capable?	no	+
* Autovaryon VG in Concurrent Mode?	no	+

F1=Help	F2=Refresh	F3=Cancel	F4=List
F5=Reset	F6=Command	F7=Edit	F8=Image
F9=Shell	F10=Exit	Enter=Do	

Figure 11-9. Change a Volume Group

Notes:

The chvg command is used to change the startup characteristics of a volume group. The -a y option sets the volume group to be used at startup. The -a n option resets this characteristic to 'no'.

Instructor Notes:

Purpose — Show how to change the characteristics of a VG.

Details — Once a VG has been set up, there are only a few attributes that can be changed. The third option determines if the VG is automatically varied off (deactivated) after losing its quorum of PVs. Selecting 'no' means that the VG stays active until it loses all of its PVs. However, note that if this option is set to 'no' then you are in danger of having backdated VGDA's. To activate a non-quorum user-defined volume group, all of the physical volumes within the volume group must be assessable or the activation fails.

The menu items referring to concurrent mode are used only with HACMP.

Additional Information —

Transition Statement — Other operations that can be carried out on the VG follow.

Reorganize a Volume Group

```
# smit reorgvg
```

Reorganize a Volume Group

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

	[Entry Fields]	
* VOLUME GROUP name	vg3	+
LOGICAL VOLUME names	[lv04 lv07]	+

F1=Help	F2=Refresh	F3=Cancel	F4=List
F5=Reset	F6=Command	F7=Edit	F8=Image
F9=Shell	F10=Exit	Enter=Do	

Figure 11-10. Reorganize a Volume Group

Notes:

The **reorgvg** command is used to redistribute the physical partitions of the logical volumes of a volume group according to their preferred allocation policies. This should improve disk performance. Preference is given in the order listed on the command line.

```
reorgvg volume group [lvname]
```

```
# reorgvg vg3 lv04 lv07
```

In AIX V4.2 and later, if you enter the **reorgvg** command with the volume group name and no other arguments, the entire volume group is reorganized.

Instructor Notes:

Purpose —Show how logical volumes in a volume group can be reorganized to handle changing logical volume requirements.

Details — **reorgvg** can be used to change the allocation of LVs in a VG based on the allocation policy set on the LV. This command will try to place the specified LV in the new place (or as close to the requested place as possible) disregarding the other LV allocation policies. This command will only be successful if there is space to maneuver.

The volume group must be varied on to use this command.

In AIX V4.1, if you run this command with the volume group name and no other arguments, only the first logical volume in the volume group will be reorganized (i.e. the first one listed by the **lsvg -l volumegroup** command). In AIX V4.2 and later, if you run the command with the volume group name and no other arguments, the entire volume group is reorganized.

When running this command, priority is given to the first logical volume name in the parameter list and lowest priority is given to the last logical volume in the list.

Additional Information —

Transition Statement —VGs can be taken off line and brought on-line again. Let's see how this can be done.

Activate/Deactivate a Volume Group

- **Activate a Volume Group (make it available for use)**

```
varyonvg [ -s ] Volumegroup #
```

```
varyonvg datavg
```

- **Deactivate a Volume Group (make it unavailable for use)**

```
varyoffvg [ -s ] Volumegroup
```

```
# varyoffvg datavg
```

Figure 11-11. Activate/Deactivate a Volume Group

Notes:

The **varyonvg** command is used to activate a volume group that is not activated at system startup (or has been added to the system since startup.) A volume group can be brought up into a system management mode (using the -s option) which prevents users from accessing it but allows the use of administrative commands.

The **varyoffvg** command is used to deactivate a volume group (for example, if it contains an external disk that was going to be moved to another machine.) No logical volumes should be open when this command is issued. To bring a volume group down into system management mode use the -s option. Removing a disk without deactivating the volume group could cause errors and loss of data in the volume group descriptor areas and the logical volumes within that volume group.

Instructor Notes:

Purpose —To show how volume groups, other than the **rootvg**, can be activated and deactivated.

Details —With non-rootvg VGs, a useful operation that can be done is to deactivate the VG, with the option of removing the VG completely from the system (provided the VG is on an external disk pack).

This proves very useful especially in high security environments, where the VG can be kept on external disks, and at night the entire VG can be deactivated and disconnected. The external disk pack can then be physically removed and placed in a safe place where it can be locked away until the next day. In the morning the VG can be activated again.

Before the **varyoffvg** command can be used, the logical volumes within it must first be closed. For example, if the logical volume is a file system, it must be unmounted.

A volume group that has a paging space on it cannot be varied off while the paging space is active.

Additional Information —

Transition Statement —Volume groups can be moved from one system to another. Let's see how this can be done.

Import / Export a Volume Group

```
# smit importvg
```

Import a Volume Group

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

	[Entry Fields]	
VOLUME GROUP name	[]	
* PHYSICAL VOLUME name	[]	+
Volume Group MAJOR NUMBER	[]	+#
Make this VG Concurrent Capable?	no	+
Make default varyon of VG Concurrent?	no	+

F1=Help	F2=Refresh	F3=Cancel	F4=List
F5=Reset	F6=Command	F7=Edit	F8=Image
F9=Shell	F10=Exit	Enter=Do	

Figure 11-12. Import / Export a Volume Group

Notes:

If you have a volume group on one or more removable disks that you want to access on another system, you must first export the volume group from the current system using the **exportvg** command. This removes any information about the volume group from the system. To export a volume group it must be inactive.

To access an exported volume group on a system, it must be imported to the system using the **importvg** command. Do not attempt to import a rootvg. Also, unless instructed to do so by support personnel, never interrupt an LVM command.

Instructor Notes:

Purpose —Show how to import and export a VG.

Details — There may be times when a volume group may need to be moved from one RISC system to another so that the logical volume's and file system's data in the volume group can be accessed directly on the target system. It may be necessary to remove all knowledge of a volume group from the system if the file system and logical volume within it are no longer being accessed.

To remove all knowledge of a volume group from the ODM database, the volume group needs to be exported. This command does not remove any user data in the volume group, but only removes its definition from the ODM database.

Similarly, when a volume group is moved, the target system needs to be made aware of the new volume group. This is done by importing.

Once exported, a volume group can no longer be accessed. `rootvg` cannot be exported because you cannot vary it off.

Once a volume group is changed or created to hold more than 1016 physical partitions per PV (as supported in AIX V4.3.1), it cannot be imported into AIX versions earlier than 4.3.1. Also, volume groups created in AIX V4.3.1 using a PP size of 512MB or 1024MB cannot be imported into previous versions of AIX.

Additional Information — You may import an AIX V3.2 VG into an AIX V4 system and vice versa provided striping has not been applied.

Testing shows that these commands can be used to change the name of a volume group.

Transition Statement —Now that we have seen some of the things that can be done with volume groups, let's look more closely at what can be done with logical volumes within a volume group.

Logical Storage

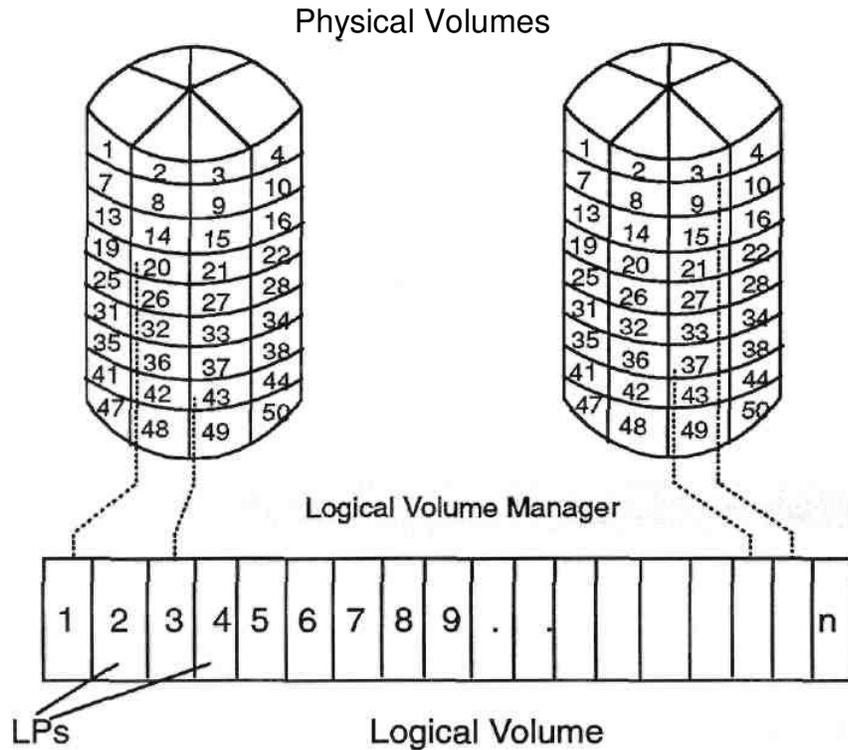


Figure 11-13. Logical Storage

Notes:

Logical Volumes

A logical volume is a set of logical partitions which may span physical volumes (as long as the physical volumes are in the same volume group). A file system sits on top of a logical volume (LV). A logical volume can be dynamically extended. There is a theoretical limit of 255 LVs per VG.

Logical Partitions

Are mapped one-to-one to physical partitions unless there is mirroring.

Instructor Notes:

Purpose — Review logical volume structure as an introduction to more detailed discussion.

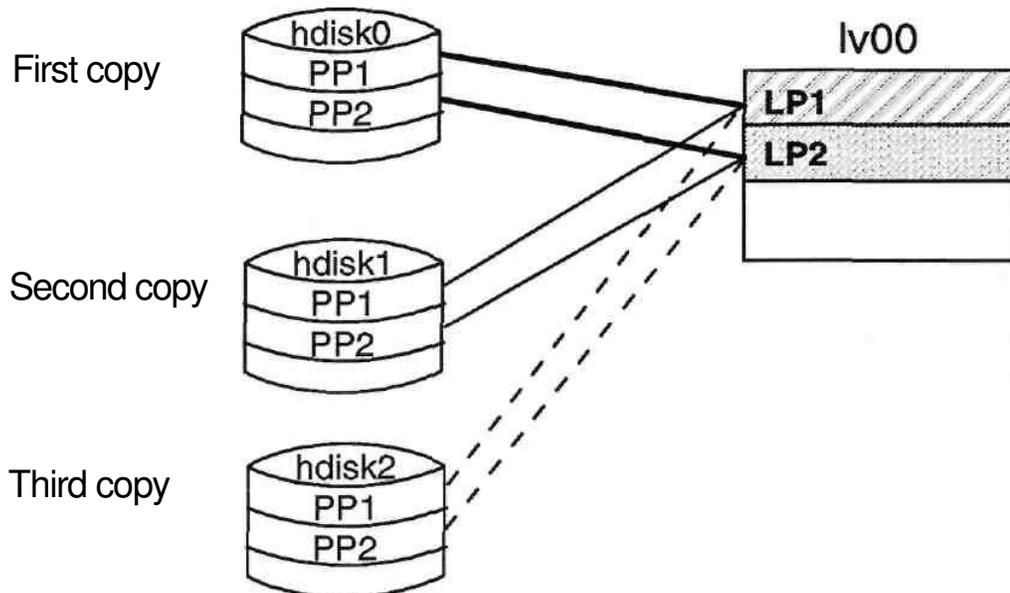
Details — Simply review logical volumes and logical partitions and how they map to physical partitions.

Review the fact that the Logical Volume Manager (LVM) is a device driver that assimilates the physical volumes into a logical view.

Additional Information —

Transition Statement — Because of the logical partition/logical volume structure, you can have more than one copy of a logical volume on a system. This is called mirroring.

Mirroring



- Mirroring is when a logical partition maps to more than one physical partition of the same volume group
- Scheduling Policy:
 - Parallel PPs written simultaneously
 - Sequential PPs written in sequence

Figure 11-14. Mirroring

Notes:

When creating a logical volume you can implement mirroring of the logical partitions of the logical volume. This means that either 2 or 3 copies are kept on disk of the logical partitions so that the data is still intact and available in the event of a disk failure. Normally, each copy must reside on a separate disk, but this restriction can be removed if required. Physical partitions do not need to be contiguous.

Implementing mirroring will inhibit the performance of the logical volume. However, this can be partly overcome by setting the scheduling policy for logical partition copies. The two options are:

- Parallel - Write requests are scheduled for each of the copies simultaneously. The write request returns when the copy that takes the longest time to **update** completes. Control is passed back to the application after the slowest disk has completed its write. This means there is a chance that the data integrity could be lost if a disk failure occurred while the copies were being updated. To overcome this problem, the mirror write consistency option can be set on. When turning on mirroring on an already existing logical volume, the copies have to be synchronized so the new copy contains a perfect image of the existing copy at that point in time. This can be done by using the **-k** option on the **mklvcopy**

command at the time mirroring is turned on or with the **syncvg** command at a later time. Until the copies are synchronized, the new copy will be marked stale.

- Sequential - when data is written to a logical partition, control is returned to the program only after all copies have been updated. Each physical partition is updated in turn making performance slower than with parallel mirroring.

On read operations:

- Parallel - reads the nearest copy to get improved response time.
- Sequential - always attempts to read the primary copy first.

Instructor Notes:

Purpose —Define what mirroring is.

Details — It is the entire LV that is mirrored, and so all the LPs within an LV will have multiple copies. By default, an LV is not mirrored.

To increase availability, mirror the copies onto physical volumes that are attached to separate busses, adapters and power supplies. Upon failure of one of these components, the copies may still be maintained, since they are not attached to the failing device.

A mirrored copy of the data can be removed. This will be shown later in this unit. Similarly, extra copies of the data can be added later on, the maximum being 3 copies. However, in order for the copies to be useful, they have to be synchronized.

Additional Information — If you are configuring your storage subsystem for maximum performance, mirroring should be disabled. If mirroring is required, then the scheduling policy should be set to parallel and allocation policy to require the LVM to place copies on separate physical volumes. In addition, write verification and mirror write consistency should also be set to no. This will prevent an extra disk revolution on every write to read back the data for validity. It also stops the LVM from waiting for all writes to copies to succeed before returning successful completion of a write.

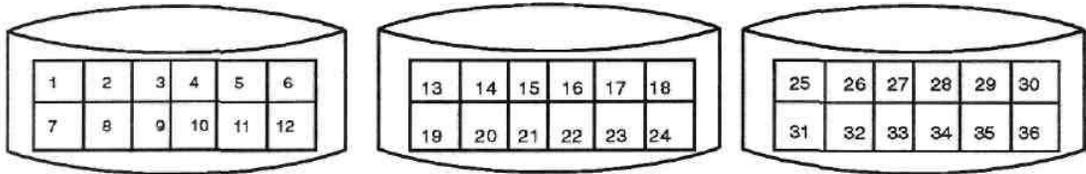
Discussion Items -

Why would you use mirroring ? **Answer:** The main use of mirroring is to enhance availability. If one disk containing the data goes down, and there is a second disk on the system containing exactly the same data, the system may be able to carry on with normal operations as if nothing had happened until the failing disk can be replaced and mirroring resumed.

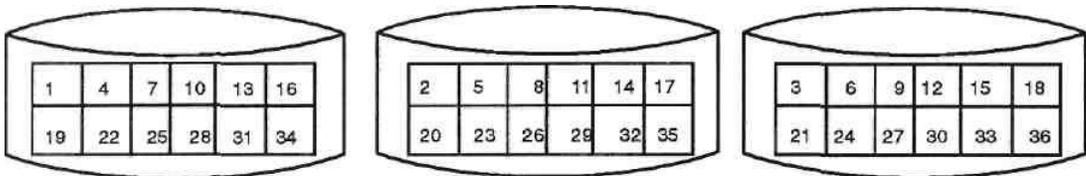
Transition Statement — Another thing that you can do with logical volumes is spread the data across physical volumes using a technique called striping.

Striping

Normal flow of data blocks when a logical volume is spread across physical volumes.



The layout of stripe units when a logical volume is set up to stripe.



- Consecutive stripe units are created on different physical volumes
- Striping increases read/write sequential throughput by evenly distributing partitions among disks
- Stripe unit size is specified at creation time

Figure 11-15. Striping

Notes:

Striping is a technique for spreading the data in a logical volume across several disks such that the I/O capacity of the disk drives can be used in parallel to access data on the logical volume. The primary objective of striping is geared toward enabling very high-performance reading and writing of large sequential files.

In non-striped logical volumes, data is accessed using addresses to data blocks within physical partitions. In a striped logical volume, data is accessed using addresses to stripe units. The size of the stripe unit is specified at creation time. It is specified as a power of 2 in the range of 4K to 128K bytes.

The limitations are as follows:

- A striped logical volume cannot be mirrored
- The number of physical partitions allocated to a striped logical volume must be able to be evenly distributed among the disks
- At least 2 physical volumes are required
- Use as many adapters as possible
- Create on a volume group dedicated to striped logical volumes

Instructor Notes:

Purpose —Define what striping is.

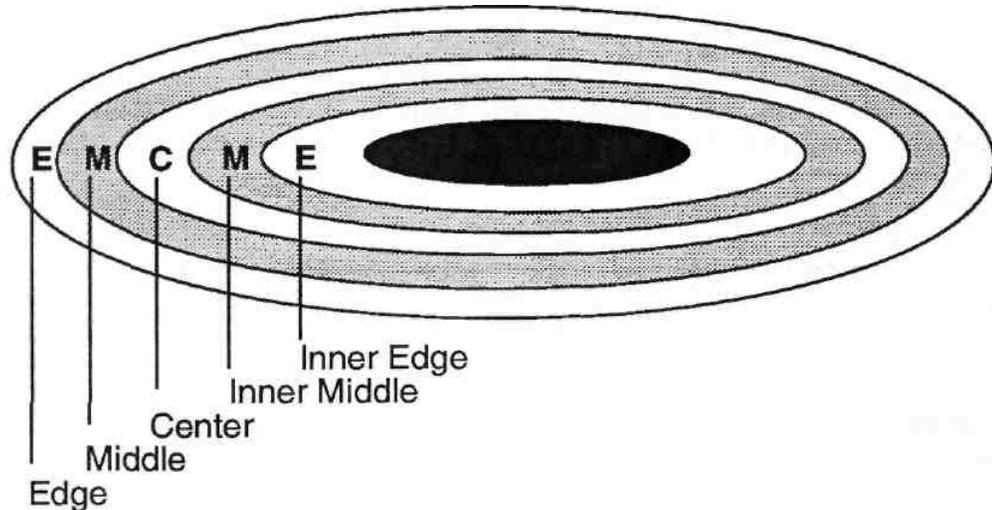
Details — When we add an LV there are a number of attributes that can be set, one of them being whether the LV will be striped or not. The default is "Not Striped". If striping is desired, the stripe size may be chosen from the pop-up list. The smallest stripe size is one block (4K). Care must be taken when creating an LV with striping as certain LV parameters are not compatible with striping. In particular, a striped LV cannot be mirrored and the number of PPs allocated must be a multiple of the number of disk drives used.

Additional Information — Data in a striped logical volume is no longer accessed using data block addresses. The LVM tracks which blocks on which physical drives actually hold the data. If the data being accessed is on more than one physical volume, the appropriate number of simultaneous disk I/O operations will be scheduled for all drives.

Transition Statement —You can also specify where on the disk, or on which disk(s) the LV is to be placed. Let's look at this next.

Logical Volume Policies

- Intra-physical volume allocation policy:



- Inter-physical volume allocation policy:
 - Maximum number of PVs to use
 - Range of PVs to use

Figure 11-16. Logical Volume Policies

Notes:

When creating/changing a logical volume you can set the way the Logical Volume Manager decides on which physical partitions to allocate to the logical volume. This will affect the performance of the logical volume.

The **intra-physical** volume allocation policy indicates where on the physical volume partitions will be allocated to the logical volume. The choices are: center, middle, edge, inner edge, and inner middle. Partitions in the center of the disk will have the shortest disk seek time and hence faster access time than those in the middle and the edge. Usually, the logical volumes that will be accessed most frequently will be allocated to the center, those that are rarely accessed to the edge and other logical volumes will be set to the middle (which is the default value.)

The **inter-physical** volume allocation policy indicates how many physical volumes can be used to contain the physical partitions of the logical volume. The maximum number of physical volumes that can be used by the logical volume can be specified (this is normally set to the number of physical volumes in the volume group). The range of volumes used can be: minimum (only allocate partitions on one physical volume, or as many as there are copies) or maximum, (allocate partitions across all physical volumes up to the maximum number of physical volumes.)

Instructor Notes:

Purpose —Define the LV allocation policies.

Details —By default you do not have to worry where an LV is placed on disk and on which disk it is placed. The LVM will handle all the policies for you. However, if you have a specific requirement you can specify this through SMIT. Remember that the intra-allocation policy can be changed later on too, after you have set up the LV. However, this operation will require a little bit of work because the VG will have to be reorganized.

Additional Information — Ask the students why, in their opinion, certain LVs are good candidates to be placed either at the center or on the edge.

Answer: The average time for the disk head to get to the center, rather than to the edge, is far less, so any critical LVs that are frequently used can be placed in the center, and any LVs that are not used as frequently can be placed on the edges. You could take this discussion one step further and ask students if they know of any LVs which could be placed at the edge or the center.

Logical volumes on 200MB, 540MB and 1GB disks that contain large sequential files should be at the edge because sequential performance is better there (there are more blocks per track at the edge than farther in).

For inter-physical policy, choose 'minimum' on nonmirrored logical volumes to provide the greatest availability (access to data in case of hardware failure). With this option, you reduce the risk of losing data because of a disk failure. Each additional physical volume used for a single physical copy increases that risk.

The 'maximum' setting spreads the physical partitions of the logical volume as evenly as possible over as many physical volumes as possible. This is a performance-oriented option because spreading the physical partitions over several disks tends to decrease the average access time for the logical volume.

Transition Statement —Let's look at SMIT's LVM screens.

SMIT Logical Volumes Menu

smit lv

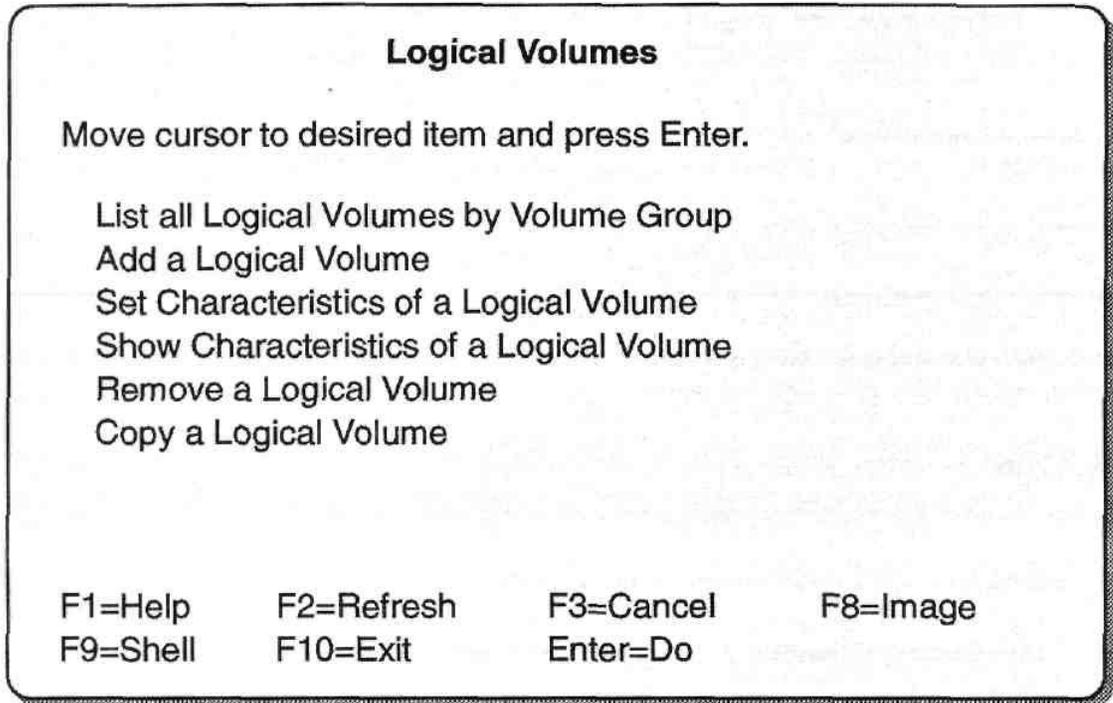


Figure 11-17. SMIT Logical Volumes Menu

Notes:

Listing Logical Volume Information

- List all Logical Volumes by Volume Group: #

lsvg -o | lsvg -i -l

LVNAME	TYPE	LPs	PPs	PVs	LV STATE	MOUNT POINT
hd6	paging	8	8	1	open/syncd	N/A
hd5	boot	1	1	1	closed/sync	N/A
hd8	jfslog	1	1	1	open/syncd	N/A
hd9var	jfs	1	1	1	open/syncd	/var

- Show Characteristics of a Logical Volume:

Status: # **lslv lv02**

LOGICAL VOLUME: lv02

LV IDENTIFIER: 000008039c7e37c4.5
 VG STATE: active/complete
 TYPE: MAX
 LPs: COPIES: jfs
 LPs: 128
 STALE PPs: 1
 INTER-POLICY: 10
 INTRA-POLICY: 0
 MOUNT POINT: minimum
 middle

VOLUME GROUP: course
 PERMISSION: read/write
 LV STATE: opened/syncd
 WRITE VERIFY: off
 PP SIZE: 4 megabyte(s)
 SCHED POLICY: parallel
 PPs: 10
 BB POLICY: relocatable
 RELOCATABLE: yes
 UPPER BOUND: 32
 LABEL: /home/malcolm

MIRROR WRITE CONSISTENCY: on EACH
 LP COPY ON A SEPARATE PV?: yes

Figure 11-18. Listing Logical Volume Information

Notes:

List all **Logical Volumes by Volume Group** uses **lsvg -o** to find out the active volume groups and then **lsvg -il** to list the logical volumes within them. The **-i** option of **lsvg** reads the list of volume groups from standard input.

The SMIT option **Show Characteristics of a Logical Volume** uses the **lslv lvname** to show status information about the selected logical volume.

Instructor Notes:

Purpose —List the attributes of a logical drive.

Details — Do not spend too much time going through all the attributes. The student should by now be familiar with this type of output.

The following attributes are specific to LVs so spend a few minutes going through them: Mirror Write Consistency, Bad Block Policy, Upper Bound (maximum number of disks that this LV can span) and Write Verify. (For further information on these attributes see the Web-based documentation.)

Mirror Write Consistency

Ensures data consistency among mirrored copies of a logical volume. This setting enhances availability by making sure that good mirrored copies are available in case of hardware failures.

Write Verify

Specifies whether to verify all writes to the logical volume with a follow-up read.

Bad Block

Indicates whether the LVM should try to relocate a bad block if one is encountered.

Additional Information —You could pose the following questions:

- If mirroring is set up, which write policy would be used?

Parallel

- Is this LV mirrored?

No - as the number of copies is one

Transition Statement —Just like the **lspv** command, the **lslv** command can be used to view the logical volume by status, physical volume and logical volume. This page showed the logical volume status. Let's look at the other two options on the next page.

Listing Logical Volume Information

- Show Characteristics of a Logical Volume:

-Physical Volume map:

```
# lslv -l lv00
```

```
lv00:/home/john
PV    COPIES    IN BAND    DISTRIBUTION
hdisk0 010:000:000 70%        000:000:007:003:000
```

-Logical Partition map:

```
# lslv -m lv00
```

```
lv00:/home/john
LP    PP1    PV1    PP2    PV2    PP3    PV3
1     0134   hdisk0
2     0135   hdisk0
3     0136   hdisk0
```

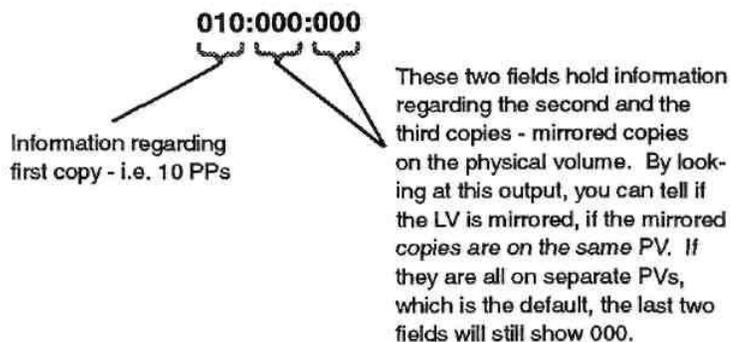
Figure 11-19. Listing Logical Volume Information

Notes:

lslv -l lvname

This gives information about the distribution of a particular logical volume's logical partitions for each physical volume. The information includes the number of logical partitions on the disk and its copies, if any, on that disk; the percentage of physical partitions which match the intra-physical volume allocation policy; the distribution of physical partitions on the physical volume (outer edge, outer middle, center, inner middle, inner edge).

Copies can be interpreted as:



The **IN BAND** attribute shows the percentage of the physical partitions on the physical volume that belong to the logical volume and were allocated within the region specified by the intra-allocation policy.

distribution

There is a relationship between the numbers 000:000:007:003:000 and 010:000:000 whereby the 007:003 numbers indicate the distribution of the 010. The interpretation is: "of the 10 PP's, 7 PP's are located in the center and 3 PP's in the inner-middle of the disk, respectively.

lslv -m lvname

This gives a map of which physical volumes contain which physical partitions for the logical partitions of the logical volume. Three columns are given, one for each copy of a logical partition.

Add a Logical Volume

smit mklv

Add a Logical Volume

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

<p>[TOP]</p> <p>Logical Volume NAME</p> <p>* Volume GROUP name</p> <p>* Number of LOGICAL PARTITIONS</p> <p>PHYSICAL VOLUME names</p> <p>Logical Volume TYPE</p> <p>POSITION on physical volume</p> <p>RANGE of physical volumes</p> <p>MAXIMUM NUMBER of PHYSICAL VOLUMES to use for allocation</p> <p>Number of COPIES of each logical partition</p> <p>Mirror Write Consistency</p> <p>Allocate each logical partition copy on a SEPARATE physical volume?</p> <p>[MORE... 10]</p>	<p>[Entry Fields]</p> <p>[]</p> <p>rootvg</p> <p>[] #</p> <p>[] +</p> <p>[]</p> <p>middle +</p> <p>minimum +</p> <p>[] #</p> <p>1 +</p> <p>yes +</p> <p>yes +</p>
---	--

F1=Help	F2=Refresh	F3=Cancel	F4=List
F5=Reset	F6=Command	F7=Edit	F8=Image
F9=Shell	F10=Exit	Enter=Do	

Figure 11-20. Add a Logical Volume

Notes:

The mklv command creates a logical volume. The name of the logical volume can be specified or alternatively a system generated name will be used. The volume group the logical volume will belong to and size (in logical partitions) must be specified. Other characteristics that can be set are the allocation policy, copies (mirroring), scheduling policy and striping.

Remove a Logical Volume

```
# smit rmlv
```

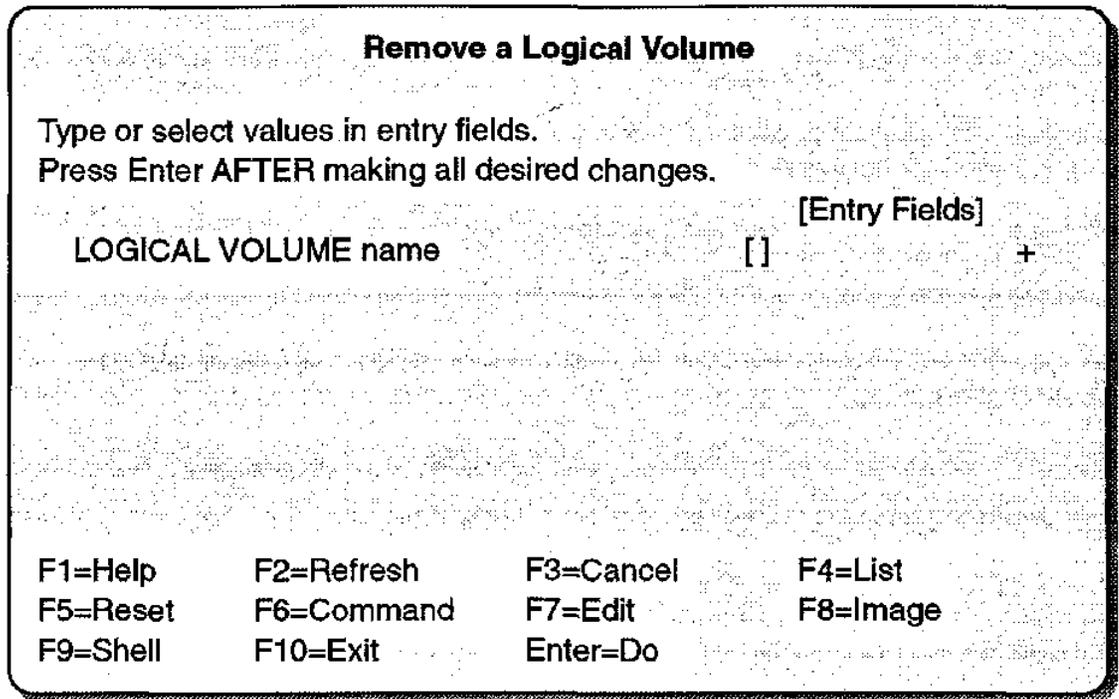


Figure 11-21. Remove a Logical Volume

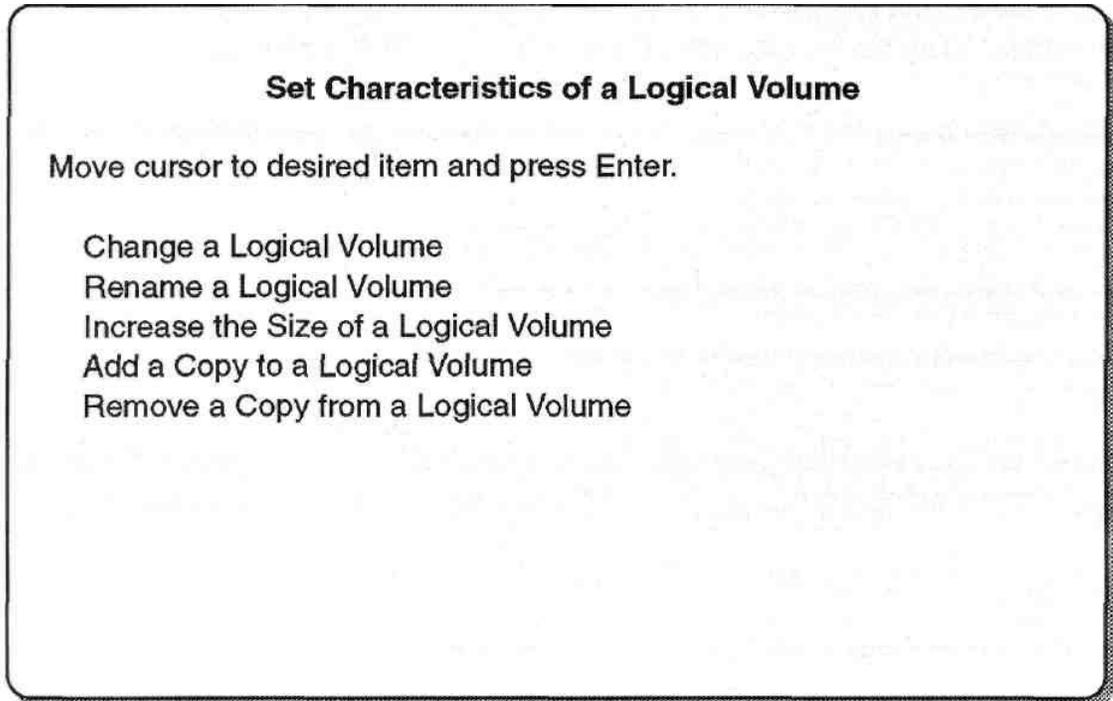
Notes:

The **rmlv** command removes a logical volume. The **-f** option will prevent the command from prompting for confirmation.

Do not use **rmlv** to remove journaled file systems or paging space volumes. These high-level structures have information relating to them saved in the ODM database and in files such as the **/etc/filesystems** file. This information will not be removed by the **rmlv** command. You should use the appropriate command for that type of data structure.

Set Characteristics of a Logical Volume

smit lvsc



F1=Help
F9=Shell

F2=Refresh
F10=Exit

F3=Cancel
Enter=Do

F8=Image

Figure 11-22. Set Characteristics of a Logical Volume

Notes:

The **chlv** command is used to change the characteristics of a logical volume. Characteristics that can be changed are the allocation and scheduling policies and the permissions. (When a logical volume is created it always has read/write permission, but this can be changed to read- only later.)

You can change the name of a logical volume using the **chlv** command with the **-n** option. No other **chlv** options can be specified if **-n** is used.

The size of a logical volume may be increased at any time, assuming that there is sufficient space in the volume group. To do this the **extendlv** command is used. You can set the allocation policies for the new partitions to different values than used by the original logical volume.

The size of a logical volume may not be decreased automatically. To make a logical volume smaller, back it up, delete it, create a new logical volume of the desired size and restore the data.

Add / Remove a Logical Volume Copy

smit mklvcopy

Add Copies to a Logical Volume

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

	[Entry Fields]	
* LOGICAL VOLUME name	lv00	
* NEW TOTAL number of logical partition copies	1	+
PHYSICAL VOLUME names	[]	+
POSITION on physical volume	middle	+
RANGE of physical volumes	minimum	+
MAXIMUM NUMBER of PHYSICAL VOLUMES to use for allocation	[32]	#
Allocate each logical partition copy on a SEPARATE physical volume?	yes	+
File containing ALLOCATION MAP	[]	
SYNCHRONIZE the data in the new logical partition copies?	no	+

F1=Help	F2=Refresh	F3=Cancel	F4=List
F5=Reset	F6=Command	F7=Edit	F8=Image
F9=Shell	F10=Exit	Enter=Do	

Figure 11-23. Add/Remove a Logical Volume Copy

Notes:

The **mklvcopy** command is used to add copies (mirroring) to a logical volume that has none or to increase the copies from two or three. Specify the logical volume to change the desired total number of copies. This will only succeed if there are enough physical partitions to satisfy the requirements on the physical volumes that are specified to be used (i.e., if all copies are to be on different physical volumes).

Also, in order for the copies to match, the logical volume has to be synchronized using the **syncvg** command. This can be done with the -k option when the copy is originally started. It can be done later using the **syncvg** command.

The **rmlvcopy** command is used to reduce the total number of copies for a logical volume. Specify the desired total number (for example, two if you are reducing the number of copies from three to two).

Once a logical volume has been created, striping cannot be imposed or removed.

Physical Volumes

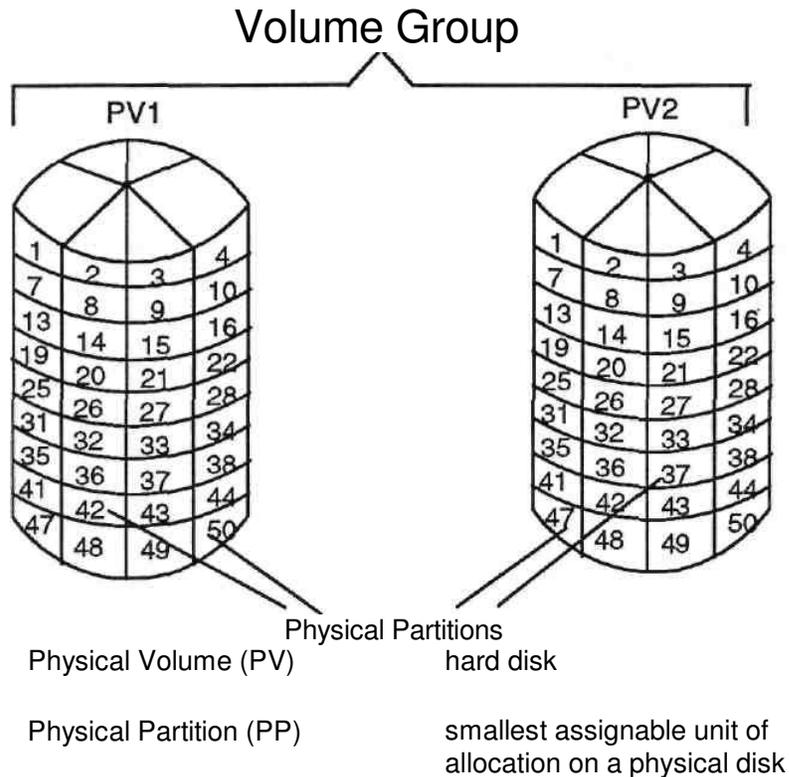


Figure 11-24. Physical Volumes

Notes:

A Physical Partition is a fixed size, contiguous set of bytes on a physical volume (PV). Physical partitions (PP) must be the same size across an entire VG. However, there may be multiple VGs on a single system, each having a different PP size. A PP can be 1 - 256 megabytes. AIX V4.3.1 also supports physical partition sizes of 512MB and 1024MB.

Prior to AIX V4.3.1 the use of 4MB or smaller physical partitions were unsupported or disks larger than 4GB due to the 1016 PPs per PV limit.

AIX V4.3.1 provides support for more than 1016 physical partitions per physical volume. This support provides for multiples of 1016 PPs per PV. Using more than 1016 PPs per PV in this manner will reduce the total number of disks that can exist in the volume group by the same fraction as indicated in the following table:

# of PPs	Max # of PVs
1016	32
2032	16
4064	8
8128	4
16256	2

SMIT Physical Volumes Menu

smit pv

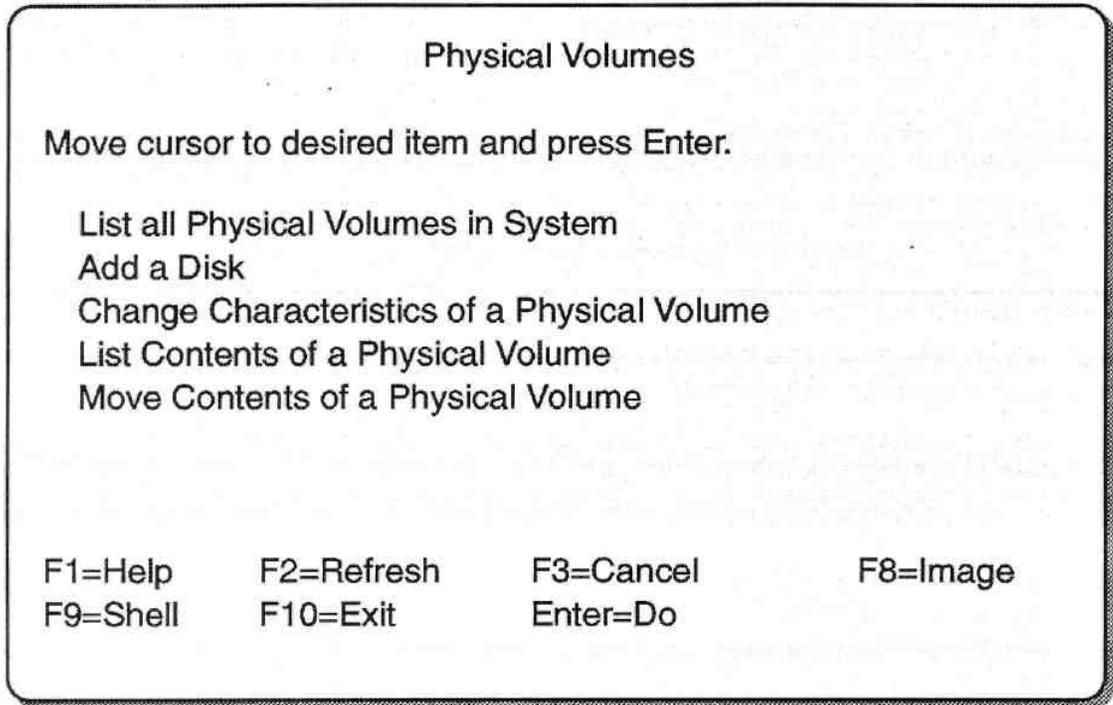


Figure 11-25. SMIT Physical Volumes Menu

Notes:

Listing Physical Volume Information

- List all Physical Volumes in System:

```
#lspv
      hdisk0    0000803eb227d386      rootvg
      hdisk1    00000060e2c69d6d      rootvg
```

List Contents of a Physical Volume: -

Status

```
# lspv hdiskO
      PHYSICAL VOLUME:  hdiskO          VOLUME GROUP:  rootvg
      PV IDENTIFIER:    0000803eb227d386
      VG IDENTIFIER:    00008039C386994
      PV STATE:         active
      STALE PARTITIONS: 0
      PP SIZE:          4 megabyte(s)    ALLOCATABLE:   yes
      TOTAL PPs:        95 (380 MBs)     LOGICAL VOLUMES: 6
      FREE PPs:         3 (12 megabytes)  VG DESCRIPTORS: 2
      USED PPs:         92 (368 megabytes)
      FREE DISTRIBUTION: 00..03..00..00..00
      USED DISTRIBUTION: 19..16..19..19..19
```

Figure 11-26. Listing Physical Volume Information

Notes:

List All Physical Volumes in System actually uses the undocumented command **getlvodm -C** to list the physical volumes in the system.

The **lspv** command with no parameters can be used to list the physical volume name, PV identifier and volume group for all physical volumes in the system.

The **lspv pvname** command gives status information about the physical volume. The most useful information here is: state (active or inactive), number of PP copies that are stale (are not up to date with other copies), total number of PPs, number of free PPs and distribution of free space on the PV.

Listing Physical Volume Information

- List Contents of a Physical Volume:

Logical Volumes

lspv -l hdisk0

hdisk0:

LV NAME	LPs	PPs	DISTRIBUTION	MOUNT POINT
hd1	12	12	00..00..00..12..00	/home
hd3	3	3	00..03..00..00..00	/tmp
hd2	29	29	00..00..17..12..00	/usr
hd4	13	13	00..00..13..00..00	/
hd8	1	1	00..00..01..00..00	N/A
hd6	8	8	00..00..00..08..00	N/A
hd5	1	1	01..00..00..00..00	N/A

Figure 11-27. Listing Physical Volume Information

Notes:

lspv -l pvname lists all the logical volumes on a physical volume including number of logical partitions, physical partitions and distributions on the disk.

Listing Physical Volume Information

- List Contents of a Physical Volume:

Physical Partition Map

lspv -p hdisk0

hdiskO:

PP RANGE	STATE	REGION	LV NAME	TYPE	MOUNT POINT
1-1	used	outer edge	hd5	boot	N/A
2-31	used	outer edge	hd2	jfs	/usr
32-32	free	outer edge			
33-40	used	outer middle	hd6	paging	N/A
41-64	free	outer middle			
65-65	used	center	hd8	jfslog	N/A
66-66	used	center	hd4	jfs	/
67-73	used	center	hd2	jfs	/usr
74-74	used	center	hd9var	jfs	/var
75-76	used	center	hd3	jfs	/tmp
77-77	used	center	hd1	jfs	/home
78-84	used	center	hd2	jfs	/usr
85-92	used	center	paging00	paging	N/A
93-95	used	center	hd2	jfs	/usr
96-127	used	inner middle	hd2	jfs	/usr
128-159	used	inner edge	hd2	jfs	/usr

Figure 11-28. Listing Physical Volume Information

Notes:

lspv -p pvname lists all the logical volumes on a disk and the physical partitions to which its logical partitions are mapped. It is listed in physical partition order and shows what partitions are free and which are used, as well as the location; i.e., outer edge, outer middle, center, etc.

Add / Move Contents Of Physical Volumes

- **Add a disk:**

The device can either be added through SMIT or it can be configured through configuration manager when the system boots up.

- **Move the contents of a Physical Volume:**

```
migratepv [ -l lvname ] sourcePV targetPV.. #
```

```
migratepv -l lv02 hdisk0 hdisk6
```

Figure 11-29. Add/Change Physical Volumes

Notes:

To add a physical volume to the system, the option performed is the **Add a Disk** option from the **Fixed Disks** menu under **Devices**. This adds the disk and assigns it an hdisk number. Once the disk has been added, it needs to be added to a volume group so that it can be used. See the **Volume Groups** or **Define a Fixed Disk to the Operating System** menus.

The alternative method is to power down the system, connect the new disk to the system, power up the system, and in so doing **cfgmgr** will be invoked, which will pick up the new device (if it is a detectable device).

In AIX V4.3.1 and later, if you wish to add a disk that violates the 1016 PP/PV limitation to a pre-existing volume group, first convert the volume group so that it can hold multiples of 1016 partitions per disk. This is done using the **chvg -t (factor)** command, where *factor* is a value between 1 and 16. Thus, the maximum number of physical partitions per physical volume for this volume group changes to factor X 1016.

The **migratepv** command can be used to move all partitions (or partitions from a selected LV) from one physical volume to one or more other physical volumes in a volume group. This would be used if the physical volume is about to be taken out of service and transferred to a different machine or to balance disk usage. Note: Striped logical volumes cannot be migrated using the **migratepv** command.

Instructor Notes:

Purpose —To discuss how to add a disk to the system, and move the contents of it from one disk to the other.

Details —Note that whatever way you try to add a disk drive to the system, if you do not have the correct software installed for that type of device, the system will prompt you to install the software first from installation media before continuing. Remind the students that this is true for any device.

If the PP/PV limitation is changed for the volume group, it cannot be imported into AIX V4.3.0 or earlier versions.

The **migratepv** command enables LVs to be moved from one PV to the next. The example shows how lv02 can be moved from hdisk0 to hdisk6 provided there is room to do so.

Additional Information —The **migratepv** command cannot migrate striped logical volumes. To move data from one physical volume to another, use the **chlv** command to copy the data and then use the **rmlv** command to remove the old copy.

Transition Statement —We have covered the major things that can be done with the Logical Volume Manager in AIX Version 4. Let's see how to document the storage setup.

Documenting the Disk Storage Setup

- List of the disks on the system (PVID and VG)

lspv

- List the VGs

lsvg

- List what LVs are contained in each VG

lsvg -l *vgname*

- List the LVs on each disk

lspv -l *pvname*

Figure 11-30. Documenting the Disk Storage Setup

Notes:

Instructor Notes:

Purpose — Document storage setup.

Details — Encourage the students to have this information readily available in case they have a problem with their system or in the very worst case, a totally crashed system.

You could also mention having a copy of **/etc/filesystems**, although this file has not yet been discussed.

Transition Statement — Before we summarize, let's consider a few questions.

Unit 11 Checkpoint

1. True or false? An LV can span more than one physical volume.
2. True or false? An LV can span more than one volume group.
3. True or false? The contents of a PV can be divided between two VGs.
4. True or false? If mirroring LVs, it is not necessary to perform a backup.
5. True or false? SMIT can be used to easily increase or decrease the size of a logical volume.
6. True or false? Striping is done at a logical partition level.

Figure 11-31. Unit 11 - Checkpoint

Notes:

T F 1. An LV can span more than one physical volume.

— **Correct Answer**

1. True

T F 2. An LV can span more than one volume group.

— **Correct Answer**

2. False

T F 3. The contents of a PV can be divided between two VG's.

— **Correct Answer**

3. False

T F 4. If mirroring LVs, it is not necessary to perform a backup.

— **Correct Answer**

4. False

You still need to back up to external media.

T F 5. SMIT can be used to easily increase or decrease the size of a logical volume.

— **Correct Answer**

5. False

SMIT can only be used to increase a file system. Decreasing one requires backing up the file system, removing it, recreating it, and then restoring.

T F 6. Striping is done at a logical partition level

— **Correct Answer**

6. False

It is done at a stripe unit level.

Unit Summary

- SMIT or high level commands can be used to Add/Change/Delete volume groups, physical volumes and logical volumes
- Mirroring is a way to have two or three copies of a logical volume for high availability requirements
- Disk striping is used to provide high performance in large, sequentially accessed file systems

Figure 11-32. Unit Summary

Notes: