



Tech Info Library

AIX: Logical Volume Manager Overview (5/96)

Article Created: 31 May 1996

TOPIC -----

This article contains an overview of the Logical Volume Manager in AIX.

DISCUSSION -----

Logical Volume Manager Overview

In a traditional UNIX environment, a hard drive is configured and formatted into distinct sections, called partitions or slices. A disk may be formatted to contain one or more of these partitions, each of which is empty and unassigned until used for a particular purpose. Most partitions later have some kind of file system created with them. These file systems can later be mounted somewhere under the root file system and accessed. Other partitions are deliberately left empty and used for system purposes, such as paging or swapping, or for specific applications such as Oracle, which use the empty space in an application specific way. Under this arrangement, a file can only be as large as the file system that contains it, which in turn can only be as large as the partition that contains it. The partition can never be larger than the single physical disk that contains it.

AIX has a way around these limitations. A traditional disk partition or slice in UNIX is replaced by something called a logical volume. The disk space that makes up a logical volume can span multiple disks, and be made up of blocks of non-contiguous space. Since a logical volume is not hard wired to a fixed amount of physical disk space, it can be dynamically enlarged without affecting the file system contained within it. That file system can then be enlarged later on to take advantage of the additional space added to the volume. Volumes can also be mirrored or striped for reliability or high speed access.

In AIX, an actual hard drive is known as a physical volume. Hard disks are grouped together to form volume groups. A volume group can contain one or more hard drives, and a hard drive can belong to only one volume group. The one thing all hard disks in a volume group have in common is, they all use the same size unit of disk allocation space, which is called a physical partition. A physical partition can be anywhere from 1 to 256 MB in size. Physical partitions are the smallest units of space allocation on a physical disk. AIX comes with one default volume group, known as rootvg. All of the logical volumes necessary to boot and run the system are a part of rootvg. These volumes (and the file systems within them are):

Begin_Table

- /dev/hd4 '/' (root)
- /dev/hd2 /usr
- /dev/hd1 /home
- /dev/hd3 /tmp
- /dev/hd6 Paging space (not a file system)
- /dev/hd8 Journal File System log space (also not a file system)
- /dev/hd5 The boot logical volume (contains kernel, file system and other structures required to boot the system)
- /dev/hd9var /var

End_Table

A system can have many volume groups; all except for the rootvg are user defined and configured.

A logical partition is the same size as a physical partition; each one maps to an actual physical partition somewhere in the volume group. A logical volume consists of one or more logical partitions from within a volume group. Since logical volumes can be increased, it stands to reason that they can be expanded (the size of) one logical partition at a time. The default size of a logical partition in the rootvg (under Apple's version of AIX) is 16 MB. This logical storage method allows for a tremendous amount of storage capacity per system. The system limits are:

Begin_Table

- Volume groups 255
- Physical volumes 32 per volume group
- Physical partitions 1016 per physical volume, up to 256 MB each in size
- Logical volumes 256 per volume group
- Logical partitions 32,512 per logical volume

End_Table

Striping and Mirroring

One of the advantages to a logical storage management system is, striping and mirroring can be easily set up. Striping allows the data blocks that make up a file to be spread over two or more disks. The data blocks used in AIX striping are called stripe units; they can be 4K to 128K bytes in size. These stripe units (contained within logical partitions on different disks) allow parallel access to the data blocks of a file. This can greatly reduce file I/O access time.

Mirroring is when a logical partition maps to more than one physical partition with the same volume group (usually physical partitions on different physical volumes). All the logical partitions with a mirrored logical volume will be mapped to two, or possibly 3 different physical partitions. Mirroring can slow the performance of the logical volume, but it is very useful in cases where data redundancy and availability are critical.

Copyright 1996, Apple Computer, Inc.

Keywords: <None>

=====

This information is from the Apple Technical Information Library.

19960531 08:19:30.00

Tech Info Library Article Number: 19854