AIX & Journaled File System: Logical Volume Limitation (12/96)

Article Created: 18 December 1996

TOPIC ---

This article provides an understanding of journaled file system (JFS) size limitations.

DISCUSSION ---

The maximum size for a JFS is defined when the file system is created. When you create a JFS, there are two significant issues to consider, maximum journaled file system size and journaled file system log size.

Maximum Journaled File System Size

The maximum JFS size is a function of the following three variables, which are set when the file system is created:

- Logical partition size of the volume group
- Fragment size
- Number of i-nodes

The Logical Volume Manager (LVM) limits a volume group to 32 physical volumes and restricts a physical volume to 1016 logical partitions. This means that the maximum file system size as a function of volume group geometry is as follows:

  32 x 1016 x PartitionSize

The fragment size limits the JFS size due to addressability constraints. Each fragment within a file system must be addressable by a 28-bit value. This means that the maximum file system size as a function of fragment size is as follows:

  228 x FragmentSize

Journaled file systems are also restricted to 224 i-nodes. This limitation yields the following maximum file system size:

  number of bytes per i-node (NBPI) x 224

The following table summarizes the restrictions on journaled file system size:

Begin_Table

NBPI Ratio (bytes)	Fragment Size (Megabytes)	Partition Size (Gigabytes)	Maximum File System Size
512	512, 1024, 2048, 4096	2	8
1024	512, 1024, 2048, 4096	2	16
2048	512, 1024, 2048, 4096	2	32
4096	512, 1024, 2048, 4096	4	64

End_Table

Journaled File System Log Size Issues

Another size-related issue is the size of the JFS log. Changes to JFS on disk control structures are recorded in a separate log logical volume.

In most instances, multiple journaled file systems use a common log configured to be 4MB in size. For example, after initial installation all file systems within the root volume group use logical volume hd8 as a common JFS log. The default logical volume partition size is 4MB, and the default log size is one partition, therefore, the root volume group normally contains a 4MB JFS log. When file systems exceed 2GB or when the total amount of file system space using a single log exceeds 2GB, the default log size may not be sufficient. In either case, the log sizes should be scaled upward.

NOTE: The log should contain 4MB of record space for each 2GB of file system space.

• The default log size may be implicitly increased by creating a volume group with a larger than 4MB partition size.

• On a newly created volume group, the crfs command creates the default log for the volume group. The size of this default log can be specified by using the -l <log_partitions> option of the crfs command.

• JFS logs may be created on an individual basis by using the mklv -t jfslog command followed by the logform command. To associate the new log with a file system, use the -a log=<log_name> option of either the crfs or the chfs command.

• The size of an existing log may be increased only if all file systems using the log are unmounted. After unmounting all necessary file systems, use the extendlv command to increase the size of the log logical volume followed by the logform command. Failure to unmount all file systems using the target log may result in file-system corruption and data loss.

Keywords: ksts

==========

19961219 08:36:56.00
Tech Info Library Article Number: 20749