Software Product Description

Acorn Fragmentation Avoider, V1.0

Description: VMS disks tend to become fragmented rapidly once placed in use, causing slower file access than might be possible, and making more work for defragmenting utilities later. This happens mainly for two reasons: First, whenever a file is deleted or truncated in VMS, the blocks are placed in an "extent cache" which is then used first in allocating new space. Unfortunately, this space tends to be in tiny, scattered chunks and guarantees that newly allocated space is badly fragmented. Second, many utilities write long files, particularly journal files, and use small extents (~5 blocks is common) even though the files may become thousands of blocks long. Because these files are held open most of the time, they not only fragment, but are difficult to defragment because defragmentation of open files is far more difficult and intrusive than of files which are closed. As these files pepper a disk with little bits of themselves, they make defragmentation of the rest less effective.

The Fragmentation Avoider consists of an intercept driver and a control utility which reduce the rate of VMS disk fragmentation in two ways:

1. Files are extended in "contiguous best try" fashion, preventing the accumulation of many tiny extents in extent cache which would otherwise cause new files to be allocated in small bits.

2. Files are extended by a fraction of their current size when this is larger than the default, subject to minimum and maximum extents and to a constraint of never forcing the extent to be more than 1/8 of volume free space. Thus files tend to be extended less often and in larger pieces than by default which leads to their being far less fragmented than otherwise.

These actions are configurable on a per volume basis, and FragmentAvoid can be enabled or disabled separately on any disk on a system. Configuration options include:

* Selection of the fraction to extend a file by (default is 25% of current size, but can be tailored).
* Selection of maximum number of blocks to extend by, even for giant files
* Selection of minimum number of blocks to extend by
* Selection of how often to force contiguous best try extension. This is normally done every time, but can be done every Nth time, where N can range from 1 to 1000. (Note that requests for contiguous extension are not touched.)
* Selection of whether to modify extension amounts where the application has selected using volume default extension only, or whether to do this every time. (Using the extension modification every time overcomes problems with utilities which write very long files with hardcoded extensions of only a few blocks, leading to badly fragmented files. Occasionally, though, this leads to the need to manually truncate files via the $SET FILE/TRUNCATE command. If size is altered only for users of volume default allocation, files never need to be truncated, but some files will fragment the disk.)

Fragmentation Avoider works on all VMS disks and is a "set it and forget it" utility. Users will notice it in that long files will tend to be written faster (some measurements show up to a 30% speedup) due to less frequent use of the (expensive) extend operation. They will also notice slower volume fragmentation if they monitor it, and will generally see fewer extents in any file if they use the DUMP/HEADER command to examine this. The effect is particularly dramatic on system disks with files like ACCOUNTNG.DAT or the OPCOM logfiles, which are generally hard to defragment.

A menu-driven setup script makes setup of FA simple to do, even without training.

Processors supported: All VMS processors, running VMS 5.5 and above on VAX, or VMS 1.5 on AXP. AXP VMS 6.1 support is in the works.

Requirements: Sufficient memory to hold one multiunit driver in pool plus its control blocks. No other software is required.