Apple Business Systems
Technical Notes

Product Technical Support

AWS11 : Configuring UNIX file system inode overhead

Written by : Sandhya Vora & John Sovereign August, 1993

This technical note documents methods for tuning/creating a UNIX file system with lower overheads in terms of inode allocation, which proves to be important in large size file systems by reclaiming disk space.

The *newfs* command makes a Berkeley 4.2 UNIX file system (UFS) on the specified device file. The A/UX version of HD SC Setup invokes *newfs* when creating a UFS.

Each file on the system is represented by an inode, which contains critical information about the file, including ownership, permissions, size, etc. The *newfs* command allocates disk space for inodes when the UFS is created.

By default, *newfs* creates an inode for each 4 kilobytes (kb) of available space (2 kb on A/UX 3.0.1 and previous versions). Since each inode takes up 128 bytes on disk, the space dedicated for inodes (or the inode overhead) amounts approximately 3% of the total filesystem space (6% on A/UX 3.0.1 and previous versions). On a 1 gigabyte drive, this works out to be 32 megabytes (MB) or 250,000 inodes (64 MB or 500,000 inodes on A/UX 3.0.1 and previous versions).

If the average size of the files you use is much greater than 4 kb, for example, large, bit-mapped, color graphics, you may want to reclaim some of the inode overhead. Use the *-i* option to *newfs* to specify a larger value of bytes-per-inode (or the inode density). Figure 1 graphically depicts the inode overhead on differently-sized filesystems given inode densities of 2, 4 and 8 kb.

For example, if all the files you use contain at least 16 kb of data, you would choose the bytes-per-inode to be 16384. You need to manually execute the *newfs* command from the A/UX CommandShell. *Important: This must be done when the filesystem is created. If you already have data on this disk drive, be sure to* ***back-up your data before re-creating the filesystem.*** Assuming the filesystem was previously formatted with HD SC Setup on an external SCSI drive (id 5) on the external motherboard SCSI bus (bus 2), you would enter the following command.

```
    newfs -i 16384 /dev/dsk/c205d0s3 other
```

The filesystem device name is identified as /dev/dsk/c*x* 0*y*d0s*z* where *x* specifies the bus (1-4), *y* is the SCSI ID number (0-6), and *z* is the slice number.

Note: If the data on the filesystem will be predominantly Macintosh files in the AppleDouble format, for example, plain text files, be sure to allocate additional inodes. The AppleDouble format requires two inodes per file, one for the resource fork and one for the data fork.

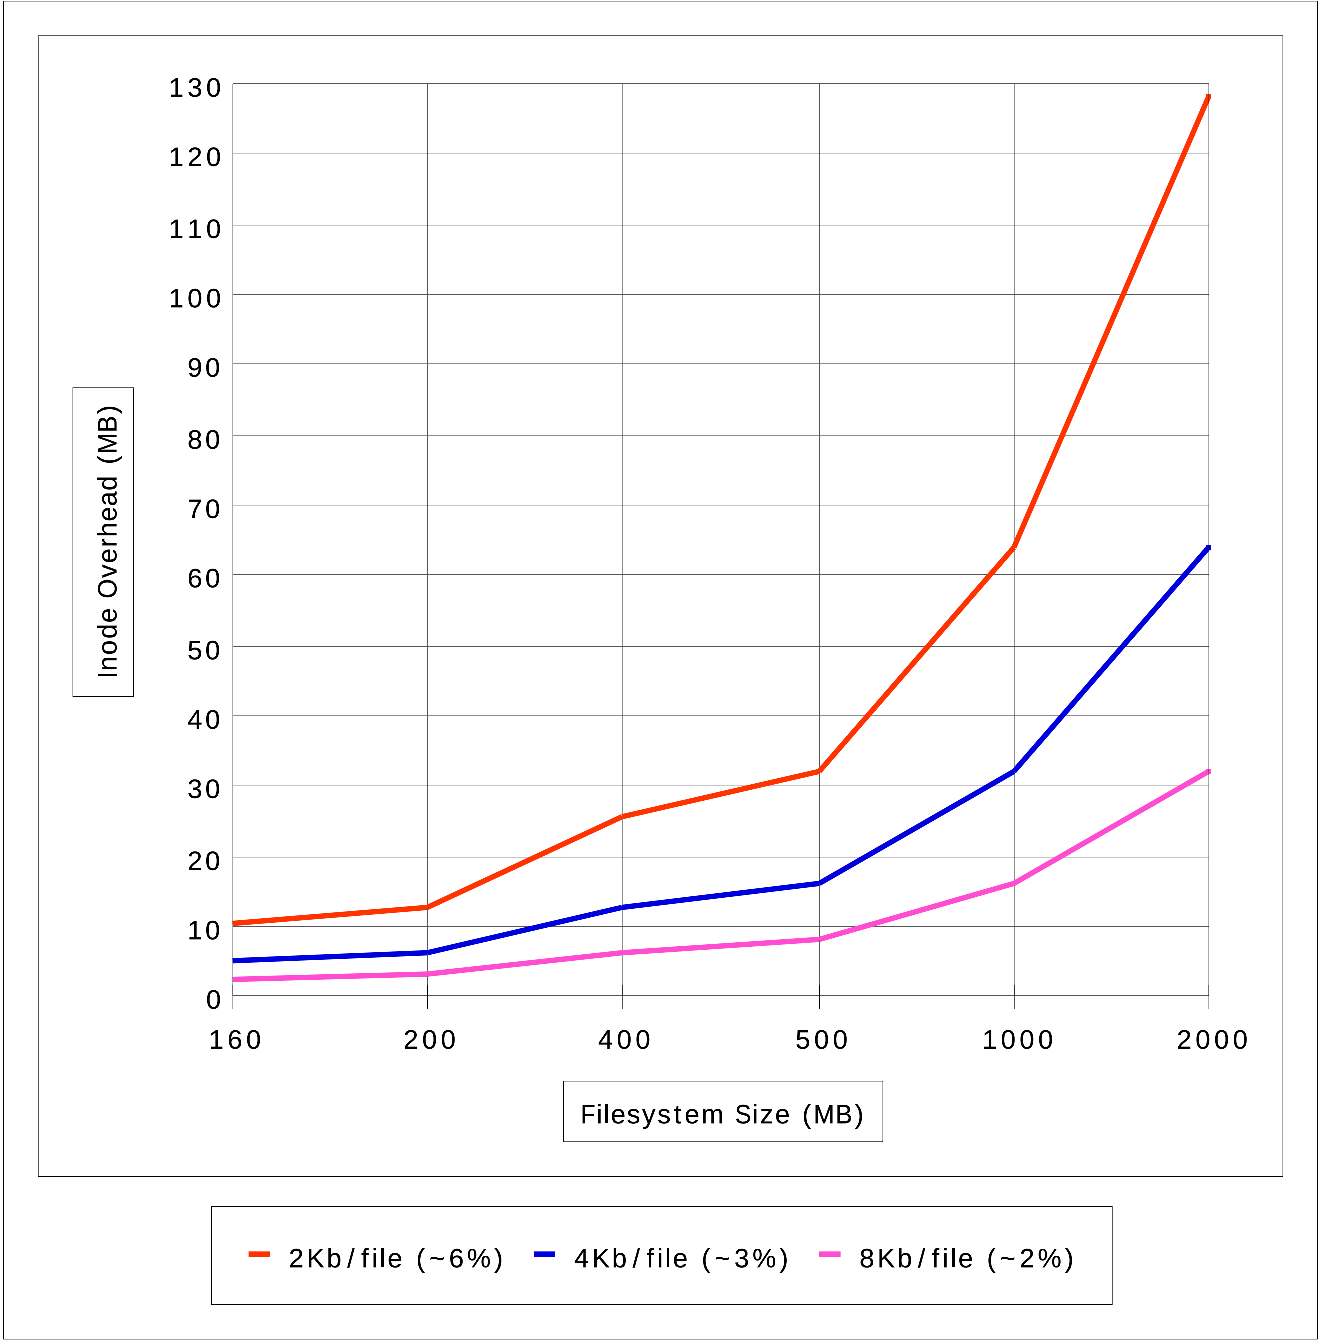

Figure 1. Inode Overhead (MB) vs. Filesystem Size (MB)