



# Tech Info Library

## Apple Workgroup Server: AppleShare Pro & AppleShare 4.0 (5/93)

Revised: 6/18/93  
Security: Everyone

Apple Workgroup Server: AppleShare Pro & AppleShare 4.0 (5/93)

=====

Article Created: 24 May 1993

TOPIC -----

The AWS 95 comes with AppleShare Pro. This is the highest performance version of AppleShare and has been enhanced to get the most out of the added I/O performance features and tuning of A/UX 3.0.1. For the most part anyone familiar with AppleShare 3.0 will find it extremely easy to switch to AppleShare Pro. There are just a few A/UX related differences that an administrator will have to watch out for. After creating AppleShare Pro the team took much of the code (except for the A/UX-only pieces) and created AppleShare 4.0. The AWS 60 and AWS 80 come with AppleShare 4.0. Any System 7 based system with a 68040 processor can also run AppleShare 4.0. The features described below pertain to both AppleShare Pro & AppleShare 4.0 unless otherwise noted.

DISCUSSION -----

Performance Features explained

-----

- Read-ahead  
This is the single most important performance feature of the new AppleShares. When a client makes a request of the server that requires reading, the server will read more than the client requested, or "read ahead". The server makes an assumption of what the next logical section of disk should be and reads it into RAM. In general, the client will ask for that section in its next network request and receive an immediate response because the data is already in RAM. If the client does not request that data or requests something else, then nothing is lost because the server will read the new data into the same RAM space (overwriting the unnecessary bits).
- Write behind  
The most time consuming operation from a user's point of view is to write

to a file server. This is the order of operation for previous versions of AppleShare:

- 1) server receives data block from network and stores it in RAM
- 2) server gives filesystem pointers to data in RAM, requests it be written to disk and waits for reply from filesystem
- 3) filesystem writes data to disk then verifies the write
- 4) filesystem replies to server
- 5) server notifies client of write completion and asks for next data block

In the new AppleShares, the server uses "write behinds". After step 1 above, the server skips to step 5 and informs the client of write completion even though it had not actually been done at the time the server sent the notification. In theory, the writes are lagging behind the messages. Because of the difference in speed between networks and high speed SCSI devices, by the time the user gets the message the write may have actually completed (depending on the filesystem and load).

Since A/UX 3.0.1 has intelligent I/O, whenever multiple writes are pending the filesystem will spool several write requests, organize them in disk order, position the disk head, and then SCSI burst them all onto the disk in one pass. This makes AppleShares Pro the fastest server under heavy load and multiple write conditions.

- File, directory, and icon caching

The new AppleShares also allow for the use of main memory as a level 3 cache. There are a few parameters that must be set by the administrator in the File Server Cache Preferences window. Following are the general guidelines for setting cache sizes.

Set the Number of Files to Cache to 20. If many files will be open at the same time, raise this value to 40. Leave the Cache Size for Each File at its default of 64K.

In determining the Folder Cache Size for a server with random usage patterns (that is, when users are as likely to access one file or folder as they are any other), multiply the number of shared files and folders by 0.2K. A large server (for example, a server with 10,000 folders and files) might allocate 2000K for the Folder Cache Size setting. For a server on which specific folders are used repeatedly and most others are accessed infrequently, multiply the number of shared files and folders by 0.05K (for the same server of 10,000 folders and files, use 500K).

Leave the Icon Cache Size at its default of 256K.

Check to make sure that the Total Cache Memory Size should be less than half the value of the Total Available Memory.

- Number of simultaneous users

AppleShare Pro can support a maximum of 200 simultaneous users; however, the Maximum Number of Connections is set to 100 by default. AppleShare 4.0 can support a maximum of 150 simultaneous users; however, the Maximum

Number of Connections is set to 75 by default.

The administrator can change this value in the File Server Preferences window.

#### Apple II support, but not booting

-----  
Apple II net booting support will not be activated; but, ProDOS will continue to be supported. This will allow Apple II's to continue to log onto AppleShare Pro. For Apple II booting, AppleShare 3.0 will still be available as a product.

#### Full International Support

-----  
All changes that were requested to upgrade AppleShare 3.0 to provide international support have been made to both AppleShare 4.0 and AppleShare Pro .

#### Anti-Piracy in AppleShare 4.0

-----  
Each legitimate AppleShare 4.0 File Server copy will have its own unique serial number assigned at manufacturing time. This serialization will make people more aware of software piracy. This will help honest people stay honest and is not a heavy-handed, fool-proof, copy-protection scheme. It will allow people to demo the software, but will make life annoying for them unless they purchase it. When a duplicate serial number is detected on the network, both servers will display an alert asserting that fact. After several hours the greeting message for clients will show the same alert. The longer the server runs after detection, the more frequently the duplicate detection alert will be displayed to all users. A tool will be produced to allow network administrators and Systems Engineers to discover how many duplicate copies of AppleShare 4.0 are running on an internetwork.

#### 2+ GB disk support in AppleShare Pro

-----  
Previous versions of AppleShare were limited to a 1.999 GB volume size because the Macintosh Finder is unable to handle larger volumes. Since the AWS 95 uses the A/UX filesystem, it can handle disks 2 GB and larger; however, the Finder-based AppleShare clients can not. To use a larger disk on the AWS 95 simply share partitions of less than 2 GB each as separate AppleShare volumes.

#### Mount point considerations in AppleShare Pro

-----  
UNIX filesystems are hierarchical with only one root, the startup volume. If there is more than one physical disk (or partition) all volumes must be connected to that startup volume. These connections are made with "mount points". The administrator chooses a folder on the startup volume for each additional volume to connect to. The additional volume then appears as a folder within the chosen folder (this is the "mount point").

Since AppleShare Pro runs atop the A/UX filesystem, there is one special consideration when sharing. Folders with mount points in them can not be

shared. This includes the A/UX startup volume (which has the Mac startup volume mount point in it). Any number of specific folders on the startup volume can be shared as long as they do not have mount points inside them. If an administrator needs to share the entire contents of a folder with a mount point in it, he/she must open the folder and share everything inside it, then select and share the volume that is connected to that mount point. This can become an administration headache; therefore, it is best for the administrator to choose the location of the mount points so that it is not necessary to share the folders that contain them.

Copyright 1993, Apple Computer, Inc.

Tech Info Library Article Number:12361