

What is the difference between the PAR 1.0 specification and the PAR 2.0 specification?

PAR 1.0 has a number of limitations:

Damaged files cannot be repaired, they must be fully reconstructed instead. This means that a single-byte error in a 10MB file would require the use of one whole PAR file to reconstruct the damaged file.

All of the PAR files are of equal size and contain enough recovery data to reconstruct the largest source file. This means that if you have source files of varied sizes and the smallest one is damaged, then you still need a whole PAR file to reconstruct it. When PAR is used on UseNet, this could mean that you have to download a 10MB PAR file to reconstruct a 3MB data file.

Damaged PAR files are of no use during reconstruction. A single byte error to a PAR file renders all of the recovery data it contains useless.

When used with small numbers of source files, it is very inefficient and you need to create an excessive number of PAR files to achieve a desired level of protection. For this reason, files are normally split into many equal sized pieces and PAR files generated from those pieces.

It cannot handle more than 255 files.

PAR 2.0 either completely removes or significantly reduces these limitations:

Damaged files can be repaired. A single byte error in a 10MB file might only require the use of recovery data from a PAR file that is only 100KB in size.

There is no relationship between the size of the data files and the size of the PAR files. Also, the PAR files will normally be of varied sizes allowing you to pick the size you need appropriate to the amount of damaged data you need to repair.

Damaged PAR files will still be useable. PAR 2.0 can use the undamaged parts of a PAR file.

PAR files can be generated from a single source file without the need to split it. On UseNet, this removes the need to use RAR or any other file splitter. Please note however, that due to the limitations of some newsreaders (which do not permit the download of incomplete files), it is advisable to use RAR to split very large files.

It can handle up to 32768 files.

So how exactly does PAR 2.0 remove all of these limitations?

The limitations of PAR 1.0 are all due to the fact that it operates on "whole" files.

PAR 2.0 operates by "virtually" splitting the files you wish to protect into many smaller "slices" (or blocks) of data. PAR 2.0 then processes these virtual slices in the same way that PAR 1.0 would process whole files. The resulting blocks of recovery data are the same size of these slices, and for convenience, many of them will be placed in a single PAR file.

How does using a PAR 2.0 program differ from a PAR 1.0 program?

A PAR 2.0 program will work equally well with a single 800MB file or with hundreds of files of varied sizes. It is not necessary either to split large files, or to archive large numbers of small files. This makes the files directly useable without the need to rejoin them or unpack them first. This means that it is possible place PAR files on an SVCD to protect a video.

With PAR 1.0, the only option you needed to specify when creating PAR files is how much recovery data you wish to create.

With PAR 2.0, you must also specify how many slices the files will be virtually divided into. Using a larger number of smaller slices allows more accurate detection of errors in files and reduces the amount of recovery data required to achieve a successful repair.

Copyright © 2003 Peter B. Clements. All Rights Reserved.