

Taking advantage of the asymmetry of offline compression

 devblogs.microsoft.com/oldnewthing/20180626-00

June 26, 2018



Raymond Chen

Online compression is a difficult balancing act. Time you spend trying to get good compression is time the user could have been using to do something useful. For the same reason, the decompression algorithm needs to be relatively fast as well. And since files support random access, you need to be able to compress and decompress an arbitrary location in the file, which typically means that the file is broken down into chunks which are compressed and decompressed independently.

Windows comes in a large compressed image, and starting in Windows 8.1, systems can be configured so that the “files” on disk are really pointers into the compressed image which Windows can decompress on the fly as if they were real files. Windows 10 expands this technique to all systems if an assessment determines that the system would benefit from it.

Offline compression with online decompression means that you can spend a lot of time in the compression step, since users never have to sit around waiting for the compression to take place. The compression can take place in a lab on a machine with a large amount of memory and CPU resources available, in order to eke out as much compression goodness as possible; only the compressed results need to be delivered to the user. On the other hand, decompression is still done in real-time, so the compression algorithm needs to be one that can still decompress relatively quickly.



Raymond Chen

Follow