

Harvester

Harvesters are individual machines controlled by a farmer. In a large farming operation, a farmer may be connected to many harvesters.

Harvesters control the actual plot files by retrieving qualities or proofs from disk. The minimum plot size (and by far the most common) is k32, which corresponds to around 100 GiB. With each increment of a k-value, the plot size roughly doubles, so a k33 plot is around 200 GiB, k34 is around 400 GiB, etc.

The difficulty level automatically adjusts every 4608 beacon chain blocks to target one proof of space - across the entire network - for every two signage points. This is the targeted average value - there can also be zero or multiple proofs per signage point. This leads to a difficulty adjustment approximately every 24 hours.

Given a plot, the harvester must perform two tasks to find a valid proof:

1. Fetch the initial quality - this requires around seven random disk seeks, or 70 milliseconds on a slow HDD.
2. (Only performed if the initial quality is sufficiently high) Fetch the full proof - this requires around 64 disk seeks, or 640 milliseconds on a slow HDD.

For most challenges, the quality (step 1) will be very low, so fetching the entire proof (step 2) will not be necessary. A node has 28 seconds to return a proof, so disk I/O will not be a limiting factor, even when proofs are stored on slow HDDs.

Tape drives are too slow for farming. The protocol was designed to support hard disks, but nothing slower. It is possible to use tape for long-term plot storage, only transferring the plots to disks for occasional farming, but this is likely a very rare use case.

Finally, harvesters also maintain a private key for each plot. The blocks are signed with these keys.

Revision #2

Created 5 June 2023 11:30:56 by Admin

Updated 5 June 2023 12:54:11 by Admin