

CLAIMS

1. A computer-implemented method comprising:
 - maintaining a plurality of tiers of storage devices and one or more tiers of caches by a storage system for storing blockchain data, the plurality of tiers of storage devices comprising at least a higher-tier storage device and a lower-tier storage device;
 - determining that a blockchain data object in a data log file stored in a lower-tier storage device is an active data object, wherein the blockchain data object is block data, transaction data, or state data; and
 - writing the data object into a cache of the one or more tiers of caches.
2. The method of claim 1, wherein the higher-tier storage device has a performance characteristic that is superior to the performance characteristic of the lower-tier storage device, wherein the performance characteristic comprises one or more of an access speed, an access bandwidth, or an access latency.
3. The method of claim 1, wherein the one or more tiers of cache comprises a memory cache and a disk cache.
4. The method of claim 1, wherein the data object comprises transaction data, block data, or state data.
5. The method of claim 1, wherein determining that a data object in a data log file stored in a lower-tier storage device is an active data object comprises:
 - determining a number of times that the data object has been accessed within a predetermined time window;
 - determining whether the number of times reaches or exceeds a predetermined threshold; and
 - in response to determining that the number of times reaches or exceeds the predetermined threshold, determining that the data object is an active data object.
6. The method of claim 1, further comprising:

generating an index indicating a storage location of the data object in the cache; and
storing the index into a delta index map in a memory of the storage system.

7. The method of claim 1, further comprising:
receiving a read request of the data object;
identifying the index corresponding to the data object in the delta index map in the memory;
determining the storage location of the data object in the cache according to the index; and
retrieving the data object from the cache based on the storage location of the data object.
8. The method of claim 1, wherein the data log file is an append-only file.
9. An apparatus for data processing and storage, the apparatus comprising a plurality of modules for performing the computer-implemented method of any one of claims 1 to 8.
10. A system for data processing and storage, comprising:
one or more processors; and
one or more computer-readable memories coupled to the one or more processors and having instructions stored thereon that are executable by the one or more processors to perform the computer-implemented method of any one of claims 1 to 8.