

ABSTRACT

Methods, systems, and computer program products are provided for dynamic data allocation to nodes in distributed systems. The method may include storing data blocks at nodes; receiving, from the nodes, node health data associated with the nodes; and for a data block, performing a dynamic data allocation including: receiving, from a node on which that data block was stored, age data associated with an age of that data block; receiving, from each client that accessed that data block at the node, network latency data associated with a network latency of that client for accessing that data block; generating, based on (i) the node health data, (ii) the age data, and (iii) the network latency data, a node assignment for that data block; and in response to determining that the node assignment is different than a current node, moving that data block from the current node to the node assignment.

< Figure 1>