CLAIMS

What is claimed is:

1. A method for scheduling and allocating data storage tasks in a plurality of storage nodes, comprising:

determining a plurality of storage tasks for performing a storage request, wherein the plurality of storage tasks are to be assigned to the plurality of storage nodes;

identifying constraints of the plurality of storage nodes;

generating, by a processor, a scheduling problem based at least on states of each of the plurality of storage nodes, the plurality of storage tasks, and the constraints, wherein the scheduling problem is a constraint satisfaction problem;

selecting one or more approaches to solving the scheduling problem based on metadata associated with the storage tasks and the constraints;

generating a scheduling solution based on the one or more approaches;

determining whether the scheduling solution comprises an assignment schedule wherein at least a predetermined threshold percentage of the constraints of the plurality of storage nodes are satisfied by the scheduling solution;

executing, by the processor, the scheduling solution by assigning storage tasks of the plurality of storage tasks to each of the plurality of storage nodes when the scheduling solution comprises the assignment schedule; and

determining another scheduling solution based on the one or more approaches when the scheduling solution does not comprise the assignment schedule.

2. The method of claim 1, wherein generating the scheduling solution is performed using one or more of an integer programming problem, Boolean satisfiability problem or specific scheduling heuristics.

3. The method of claim 1, further comprising optimizing the scheduling solution such that the scheduling solution 1) is optimized for a given objective or objectives and 2) satisfies all of the constraints.

4. The method of claim 3, wherein the given objectives or objectives include one or more of:

minimal power used by a storage node;

no nodes were turned on from stand-by mode; and

specific data-durable distribution was used.

5. The method of claim 1, further comprising:

determining the states of each of the plurality of storage nodes by:

determining which of the plurality of storage nodes are currently online;

determining storage space available in each of the plurality of storage nodes; and

determining a workload of each of the plurality of storage nodes.

6. The method of claim 5, further comprising:

inspecting a data size for each of the plurality of storage tasks; and

determining additional information for each of the plurality of storage tasks.

7. The method of claim 1, wherein the metadata associated with the plurality of storage tasks and the constraints comprises at least one or more of current file location, data size, access and operation (edit, create, scheduled deletion) dates, number of copies, and copy locations.

8. The method of claim 1, wherein generating the scheduling solution comprises using machine learning and heuristics to generate the scheduling solution.

9. The method of claim 1, further comprising:

generating an alternate scheduling solution based on the one or more approaches;

comparing respective efficiencies of the scheduling solution and the alternate scheduling solution;

selecting the alternate scheduling solution responsive to determining that the alternate scheduling solution is comparatively most efficient; and

applying the alternate scheduling solution to the scheduling problem by assigning the plurality of storage tasks according to the alternate scheduling solution.

10. The method of claim 1, wherein the constraints describe prohibited and/or discouraged states of the plurality of storage nodes.

11. The method of claim 10, wherein the constraints define a maximum allowed power, a maximum number of nodes allowed online, or a restriction on which nodes can serve the storage request.

12. The method of claim 1, wherein chunks belonging to a same data file associated with a storage task of the plurality of storage tasks are stored in different storage nodes.

13. A system for scheduling and allocating data storage tasks in a plurality of storage nodes, the system comprising:

a hardware processor configured to:

determine a plurality of storage tasks for performing a storage request, wherein the plurality of storage tasks are to be assigned to the plurality of storage nodes;

identify constraints of the plurality of storage nodes;

generate a scheduling problem based at least on states of each of the plurality of storage nodes, the plurality of storage tasks, and the constraints, wherein the scheduling problem is a constraint satisfaction problem;

select one or more approaches to solving the scheduling problem based on metadata associated with the storage tasks and the constraints;

generate a scheduling solution based on the one or more approaches;

determine whether the scheduling solution comprises an assignment schedule wherein at least a predetermined threshold percentage of the constraints of the plurality of storage nodes are satisfied by the scheduling solution;

executing, by the processor, the scheduling solution by assigning storage tasks of the plurality of storage tasks to each of the plurality of storage nodes when the scheduling solution comprises the assignment schedule; and

determine another scheduling solution based on the one or more approaches when the scheduling solution does not comprise the assignment schedule.

14. The system of claim 13, wherein generating the scheduling solution is performed using one or more of an integer programming problem, Boolean satisfiability problem or specific scheduling heuristics.

15. The system of claim 13, wherein the hardware processor is further configured to optimize the scheduling solution such that the scheduling solution 1) is optimized for a given objective or objectives and 2) satisfies all of the constraints.

16. The system of claim 15, wherein the given objectives or objectives include one or more of:

minimal power used by a storage node;

no nodes were turned on from stand-by mode; and

specific data-durable distribution was used.

17. The system of claim 13, the hardware processor further configured to:

determine the states of each of the plurality of storage nodes by:

determining which of the plurality of storage nodes are currently online;

determining storage space available in each of the plurality of storage nodes; and

determining a workload of each of the plurality of storage nodes.

18. A non-transitory computer-readable medium storing therein computer-executable instructions, the instructions comprising:

determining a plurality of storage tasks for performing a storage request, wherein the plurality of storage tasks are to be assigned to a plurality of storage nodes;

identifying constraints of the plurality of storage nodes;

generating, by a processor, a scheduling problem based at least on states of each of the plurality of storage nodes, the plurality of storage tasks, and the constraints, wherein the scheduling problem is a constraint satisfaction problem;

selecting one or more approaches to solving the scheduling problem based on metadata associated with the storage tasks and the constraints;

generating a scheduling solution based on the one or more approaches;

determining whether the scheduling solution comprises an assignment schedule wherein at least a predetermined threshold percentage of the constraints of the plurality of storage nodes are satisfied by the scheduling solution;

executing, by the processor, the scheduling solution by assigning storage tasks of the plurality of storage tasks to each of the plurality of storage nodes when the scheduling solution comprises the assignment schedule; and

determining another scheduling solution based on the one or more approaches when the scheduling solution does not comprise the assignment schedule.