CLAIMS

1. A system, comprising:

a plurality of computing devices configured to implement:

a current cluster having a plurality of nodes storing cluster data, wherein each node comprises a respective at least one storage device that stores a respective portion of the cluster data, wherein the current cluster receives access requests for the cluster data at a network endpoint for the current cluster; and

a cluster control interface configured to:

receive a cluster scaling request for the current cluster, wherein said cluster scaling request indicates a change in a number or type of nodes in the current cluster;

in response to receiving the cluster scaling request:

create a new cluster having a plurality of nodes as indicated in the cluster scaling request;

initiate a copy of the cluster data from the plurality of nodes of the current cluster to the plurality of nodes in the new cluster, wherein after completion of the copy of a respective portion of the cluster data from one of the nodes of the current cluster and before completion of the copy of another respective portion of the cluster data from another one of the nodes of the current cluster, the current cluster continues to respond to all read requests for the cluster including a read request directed to the respective portion that has already been copied to the new cluster; and

subsequent to completion of the copy of the cluster data to the plurality of nodes in the new cluster:

move the network endpoint for the current cluster to the new cluster, wherein after the network endpoint is moved, access requests directed to the network endpoint are sent to the new cluster; and

disable the current cluster.

2. The system of claim 1, wherein in response to receiving the cluster scaling request, the cluster control interface is further configured to set the current cluster to operate in read-only mode.

3. The system of claim 1, wherein each node of the current cluster stores one or more data slices of the cluster data, and wherein the copy of the cluster data from the current cluster to the plurality of nodes in the new cluster comprises:

for each data slice in the nodes of the current cluster, assign the data slice to one or more data slices in the plurality of nodes in the new cluster according to an assignment scheme; and

transfer the cluster data from each data slice in the nodes of the current cluster to the assigned one or more data slices in the plurality of nodes in the new cluster.

4. The system of any one of claims 1 to 3, wherein the plurality of computing devices are part of a larger collection of computing devices implementing a network-based cluster hosting service, and wherein each of the nodes of the current cluster of nodes and the plurality of nodes in the new cluster are virtual compute nodes hosted by the network-based cluster hosting service.

5. A method, comprising:

performing, by a plurality of computing devices:

detecting a cluster scaling event for a current cluster having a plurality of nodes storing cluster data, wherein each node comprises a respective at least one storage device that stores a respective portion of the cluster data, wherein the current cluster receives access requests for the cluster data at a network endpoint for the current cluster;

determining a change in the nodes of the current cluster based on the cluster scaling event;

creating a new cluster having a plurality of nodes as indicated by the change;

initiating a copy of the cluster data from the plurality of nodes of the current cluster to the plurality of nodes in the new cluster, wherein after completion of the copy of a respective portion of the cluster data from one

of the nodes of the current cluster and before completion of the copy of another respective portion of the cluster data from another one of the nodes of the current cluster, the current cluster continues to respond to all read requests directed to the network endpoint for the current cluster including a read request directed to the respective portion that has already been copied to the new cluster; and

moving the network endpoint for the current cluster to the new cluster when the copy of the cluster data to the new cluster is complete, wherein after the network endpoint is moved, access requests directed to the network endpoint are sent to the new cluster.

6. The method of claim 5, wherein the change is to a different number or type of nodes in the current cluster.

7. The method of claim 6, wherein the change in type of node is a different storage capacity for the node or a different computing capacity for the node.

8. The method of claim 5, wherein the plurality of computing devices are part of a larger collection of computing devices implementing a network-based cluster hosting service, wherein at least one or more of the collection of computing devices implement a cluster control interface for the network-based cluster hosting service, and wherein the method further comprises:

performing, by the cluster control interface for the network-based cluster hosting service on the at least one or more computing devices:

wherein said detecting the cluster scaling event for the current cluster having a plurality of nodes storing cluster data comprises receiving a cluster scaling request for the current cluster, wherein said cluster scaling request indicates the plurality of nodes different in number or in type from the current cluster;

wherein said determining a change in the nodes of the current cluster based on the cluster scaling event comprises determining the change the plurality of nodes different in number or in type from the current cluster as indicated by the cluster scaling request; and

in response to receiving the cluster scaling request, directing said creating, said initiating, and said moving to be performed.

9. The method of claim 8, further comprising in response to receiving the cluster scaling request, directing, by the control interface, the current cluster to operate in read-only mode.

10. The method of claim 8, wherein each of the nodes of the current cluster of nodes and the plurality of nodes in the new cluster are virtual compute nodes hosted by the network-based cluster hosting service.

11. The method of claim 5, wherein each node comprises one or more data slices for storing cluster data, and wherein the copy comprises sending the cluster data from each data slice in the nodes of the current cluster to one or more data slices in the plurality of nodes in the new cluster according to an assignment scheme.

12. The method of claim 5, wherein each node comprises one or more data slices for storing cluster data, and wherein the method further comprises distributing the cluster data on the data slices on the nodes of the new cluster based on a distribution scheme different from another distribution scheme used to distribute the cluster data on the data slices on the nodes of the current cluster.

13. The method of claim 5, wherein a user sends the access requests before and after the movement of the network endpoint to a same network address.

14. The method of any one of claims 5 to 13, further comprising:

wherein the current cluster stores replicated cluster data in addition to the cluster data;

wherein said copy of the cluster data from the current cluster to the plurality of nodes in the new cluster is performed without sending the replicated cluster data from the current cluster to the new cluster;

in response to the completion of the copy, replicating the cluster data among the plurality of nodes in the new cluster.

15. A non-transitory, computer-readable storage medium, storing program instructions that when executed by a plurality of computing devices cause the plurality of computing devices to implement:

detecting a cluster scaling event for a current cluster having a plurality of nodes storing cluster data, wherein each node comprises a respective at least one storage device that stores a respective portion of the cluster data, and wherein the current cluster receives access requests for the cluster data at a network endpoint for the current cluster;

in response to detecting the cluster scaling event:

directing the current cluster to operate in read-only mode;

determining a change in the nodes of the current cluster based on the cluster scaling event;

creating a new cluster having a plurality of nodes as indicated by the change;

initiating a copy of the cluster data from the plurality of nodes of the current cluster to the plurality of nodes in the new cluster, wherein after completion of the copy of a respective portion of the cluster data from one of the nodes of the current cluster and before completion of the copy of another respective portion of the cluster data from another one of the nodes of the current cluster, the current cluster continues to respond to all read requests directed to the network endpoint for the current cluster including a read request directed to the respective portion that is copied to the new cluster; and

subsequent to completion of the copy of the cluster data to the plurality of nodes in the new cluster, moving the network endpoint for the current cluster to the new cluster, wherein after the network endpoint is moved, access requests directed to the network endpoint are sent to the new cluster.