CLAIMS

1. A distributed computing system, comprising:

a plurality of compute nodes, each compute node comprising at least one processor and a memory; and

an interface;

wherein the distributed computing system implements a distributed computing service;

wherein the plurality of compute nodes are configured as a cluster of compute nodes according to a MapReduce distributed computing framework, wherein the cluster is configured to execute a distributed application, wherein the cluster comprises two or more groups, each group comprising a non-overlapping set of compute nodes;

wherein the distributed computing service is configured to:

receive, through the interface from a client of the distributed computing service, input defining an expression that, when evaluated true, represents a trigger condition for performing an automatic scaling operation on a particular group of the two or more groups of the cluster, wherein the input also identifies the particular group to which the automatic scaling operation applies prior to triggering of the automatic scaling operation, and wherein the input also specifies a scaling action to be taken in response to the expression evaluating true, wherein the expression is dependent on values of one or more metrics generated during execution of the distributed application;

collect, during execution of the distributed application, the one or more metrics;

determine, during execution of the distributed application and dependent on the collected metrics, that the expression evaluates true; and

initiate, in response to the determination, performance of the automatic scaling operation on the particular group of the cluster, wherein the automatic scaling operation comprises an operation to add one or more compute nodes to the particular group of the cluster or an operation to remove one or more compute nodes from the particular group of the cluster, wherein at

least one other of the two or more groups is unchanged by performance of the automatic scaling operation.

2. The system of claim 1,

wherein the inputs received through the interface define an automatic scaling policy.

3. The system of claim 1,

wherein the inputs received through the interface define the automatic scaling policy; and

wherein the automatic scaling policy specifies that the scaling action to be taken in response to the expression evaluating true comprises an operation to remove the particular group of the two or more groups of compute nodes from the plurality of compute nodes.

4. The system of claim 1,

wherein the distributed application is configured to emit one or more application-specific metrics that were defined by the client of the distributed computing service; and

wherein the expression is dependent on at least one of the one or more application-specific metrics.

5. The system of claim 1,

wherein the expression is dependent on one or more metrics that are emitted by the cluster or by one or more of the compute nodes by default while operating in the distributed computing system.

6. The system of claim 1,

wherein to collect, during execution of the distributed application, the one or more metrics, the distributed computing service is configured to:

receive one or more metrics from a respective monitoring component on each of two or more of the plurality of compute nodes; and

aggregate the metrics received from the respective monitoring components to generate an aggregate metric for the two or more compute nodes; and

wherein the expression is dependent on the aggregate metric.

7. A method, comprising:

performing, by one or more computers:

creating a cluster of computing resource instances, wherein the cluster comprises two or more instance groups, each comprising a respective one or more computing resource instances;

associating an automatic scaling policy with a particular group of the two or more instance groups, wherein the particular group is identified based at least in part on input received from a client of a distributed computing system that includes the cluster of computing resources, wherein the automatic scaling policy defines a condition that, when met, triggers the performance of an automatic scaling operation on the particular one of the two or more instance groups that changes the respective number of computing resource instances in the particular one of the two or more instance groups without changing the respective number of computing resource instances in at least one other of the two or more instance groups;

after associating the automatic scaling policy with the particular one of the two or more instance groups, detecting, during execution of a distributed application on the cluster, that the trigger condition has been met; and

initiating, in response to said detecting, performance of the automatic scaling operation on the particular one of the two or more instance groups.

8. The method of claim 7,

wherein the trigger condition comprises an expression that, when evaluated true, triggers the performance of the automatic scaling operation on the particular one of the two or more instance groups, and wherein the expression is dependent on one or more metrics generated during execution of the distributed application on the cluster.

9. The method of claim 7,

wherein the trigger condition comprises an expression that, when evaluated true, triggers the performance of the automatic scaling operation on the particular one of the two or more instance groups, and wherein the expression is dependent on a day of the week, a date, a time of day, an elapsed period of time, or an estimated period of time.

10. The method of claim 7, further comprising:

receiving input associating another automatic scaling policy with another one of the two or more instance groups, wherein the other automatic scaling policy defines a second condition that, when met, triggers the performance of a second automatic scaling operation on the other one of the two or more instance groups that changes the number of computing resource instances in the other one of the two or more instance groups;

detecting, during execution of the distributed application on the cluster, that the second trigger condition has been met; and

in response to detecting that the second trigger condition has been met, initiating performance of the second automatic scaling operation on the other one of the two or more instance groups.

11. The method of claim 7,

wherein the automatic scaling operation comprises an operation to add capacity to the particular one of the two or more instance groups.

12. The method of claim 7,

wherein the automatic scaling operation comprises an operation to remove capacity from the particular one of the two or more instance groups.

13. The method of claim 12,

wherein the method further comprises:

determining which of the one or more of the computing resource instances to remove from the particular one of the two or more instance groups; and

removing the determined one or more of the computing resource instances from the one of the two or more instance groups; and

wherein said determining is dependent on one or more of:

determining that one of the computing resource instances in the particular one of the two or more instance groups stores data that would be lost if the computing resource were removed;

determining that removal of one of the computing resource instances in the particular one of the two or more instance groups would result in a replication requirement or quorum requirement not being met;

determining that one of the computing resource nodes in the particular one of the two or more instance groups has been decommissioned;

determining that one of the computing resources nodes in the particular one of the two or more instance groups is currently executing a task on behalf of the distributed application; or

determining progress of a task that is currently executing on one of the computing resource instances in the particular one of the two or more instance groups.

14. The method of claim 7,

wherein the automatic scaling policy further defines an amount by which the automatic scaling operation changes the capacity of the particular one of the two or more instance groups or a percentage by which the automatic scaling operation changes the capacity of the particular one of the two or more instance groups.

15. The method of claim 7,

wherein each one of the two or more instance groups comprises computing resource instances of a respective different type or computing resource instances having a respective different role in the execution of the distributed application on the cluster.

16. The method of claim 7,

wherein said detecting is performed by an external service implemented on computing resources outside of the cluster of computing resource instances; and

wherein said initiating is performed in response to receiving an indication from the external service that the trigger condition has been met.

17. The method of claim 7, wherein said creating the cluster comprises configuring a collection of computing resource instances that includes the one or more computing resource instances in each of the two or more instance groups as a cluster of compute nodes according to a MapReduce distributed computing framework.

18. The method of claim 7, wherein the cluster of computing resource instances comprises one or more virtualized computing resource instances or virtualized storage resource instances.

19. A non-transitory computer-accessible storage medium storing program instructions that when executed on one or more computers cause the one or more computers to implement a distributed computing service;

wherein the distributed computing service comprises:

a cluster of virtualized computing resource instances configured to execute a distributed application;

an interface through which one or more clients interact with the service; and

an auto-scaling rules engine;

wherein the distributed computing service is configured to:

receive, through the interface from a client of the distributed computing service, input defining an automatic scaling policy, wherein the input comprises information defining an expression that, when evaluated true, represents a trigger condition for performing an automatic scaling operation, information specifying a scaling action to be taken in response to the expression evaluating true, and input identifying a subset of the virtualized computing resource instances of the cluster to which the automatic scaling policy applies; and

wherein the auto-scaling rules engine is configured to:

after the subset of the virtualized computing resource instances of the cluster is identified, determine, during execution of the distributed application and dependent on one or more metrics generated during the execution, that the expression evaluates true; and

initiate, in response to the determination, performance of the automatic scaling operation, wherein the automatic scaling operation comprises an operation to add one or more instances to the subset of the virtualized computing resource instances of the cluster to which the automatic scaling policy applies or an operation to remove one or more instances from the subset of the virtualized computing resource instances of the cluster to which the automatic scaling policy applies, wherein the virtualized computing resource instances of the cluster and outside of the subset are unaffected by performance of the automatic scaling policy.

20. The non-transitory computer-accessible storage medium of claim 19, wherein the expression is dependent on one or more of: a value of one of the one or more metrics generated during the execution of the distributed application, a minimum or maximum threshold specified for one of the metrics generated during the execution of the distributed application, a length of time that a minimum or maximum threshold for one of the metrics generated during the execution of the distributed application is violated, a day of the week, a date, a time of day, an elapsed period of time, an estimated period of time, a resource utilization metric, a cost metric, an estimated time to complete execution of a task on behalf of the distributed application, or a number of pending tasks to be performed on behalf of the distributed application.

21. The non-transitory computer-accessible storage medium of claim 19, wherein the expression is dependent on one or more of:

a metric that is emitted by the application, by the cluster, or by one or more of the virtualized computing resources instances by default while operating in the distributed computing system; or

an application-specific metric that was defined by the client of the distributed computing service and that is emitted by the distributed application during its execution.

22. The non-transitory computer-accessible storage medium of claim 19, wherein the input defining the automatic scaling policy conforms to an application programming interface (API) that is defined for providing input to the auto-scaling rules engine.