

CLAIMS

What is claimed is:

1. ~~A method of task scheduling, comprising:~~

~~analyzing network resources required for reading and writing tasks across a cluster to obtain statuses of required network resources while executing the reading and writing tasks; and~~

~~scheduling tasks based on the obtained statuses of the required network resources.~~

2. ~~The task scheduling method according to claim 1, wherein the scheduling the task according to the usage information of the network resources needed for reading and writing comprises:~~

~~scheduling, if network resources needed for reading are more than network resources needed for writing, the task to a target cluster where the read dependent data is located.~~

3. ~~The task scheduling method according to claim 1, wherein the analyzing network resources needed for a task to perform cross-cluster reading and writing comprises:~~

~~obtaining, according to a history record, an amount of input data and an amount of output data for each task; and~~

~~calculating an input-output ratio for indicating a proportion of the network resources needed for reading and writing for each task, wherein the input-output ratio is equal to a ratio of the amount of input data to the amount of output data.~~

4. ~~The task scheduling method according to claim 3, wherein the scheduling the task according to the usage information of the network resources needed for reading and writing comprises:~~

~~determining whether the task meets a preset selection criteria, wherein the selection criteria comprises that the input-output ratio is greater than a preset first threshold, wherein the first threshold is greater than 1; and~~

~~scheduling, if the task meets the selection criteria, the task to the target cluster where the dependent data read by the task locates.~~

~~5. The task scheduling method according to claim 4, after the determining whether the task meets a preset selection criteria, the method further comprises:~~

~~obtaining a task identifier for a task meeting the selection criteria; and
generating scheduling information for recording the task identifier.~~

~~6. The task scheduling method according to claim 5, wherein the scheduling, if the task meets the selection criteria, the task to the target cluster where the dependent data read by the task locates comprises:~~

~~obtaining, when a to-be-scheduled task is received, a task identifier for the to-be-scheduled task; and
scheduling the to-be-scheduled task to a target cluster where dependent data for the to-be-scheduled task locates if the task identifier of the to-be-scheduled task matches the task identifier in the scheduling information.~~

~~7. The task scheduling method according to claim 5 or 6, wherein the obtaining a task identifier comprises:~~

~~determining whether a type of the task is SQL;
performing hash processing on the task to obtain a hash digest if the task type is SQL, and using the hash digest as the task identifier; and
if the task type is not SQL, using a number of the task as the task identifier.~~

~~8. The task scheduling method according to claim 4, wherein the selection criteria further comprises that the amount of output data is smaller than a second threshold;~~

~~and/or occupied cluster resources are less than a preset quota, wherein the occupied cluster resources comprise at least one of operating overhead, operating frequency, and cluster load.~~

~~9. The task scheduling method according to any one of claims 1 to 6, wherein the network resources comprise at least one of a network bandwidth and a network bandwidth-delay product.~~

~~10. A task scheduling apparatus, comprising:~~

~~an analysis module, configured to analyze network resources needed for a task to perform cross-cluster reading and writing to obtain usage information of the network resources needed for the task to perform the cross-cluster reading and writing; and~~

~~a scheduling module, configured to schedule the task according to the usage information of the network resources needed for reading and writing.~~

~~11. The task scheduling apparatus according to claim 10, wherein~~

~~the scheduling module is specifically configured to schedule, if network resources needed for reading are more than network resources needed for writing, the task to a target cluster where the read dependent data is located.~~

~~12. The task scheduling apparatus according to claim 10, wherein the analysis module comprises:~~

~~an obtaining unit, configured to obtain, according to a history record, an amount of input data and an amount of output data for each task; and~~

~~a calculation unit, configured to calculate an input-output ratio for indicating a proportion of the network resources needed for reading and writing for each task, wherein the input-output ratio is equal to a ratio of the amount of input data to the amount of output data.~~

~~13. The task scheduling apparatus according to claim 12, wherein the scheduling module comprises:~~

~~a determining unit, configured to determine whether the task meets a preset selection criteria, wherein the selection criteria comprises that the input-output ratio is greater than a preset first threshold, wherein the first threshold is greater than 1; and~~

~~a scheduling unit, configured to schedule, if the task meets the selection criteria, the task to the target cluster where the dependent data read by the task locates.~~

~~14. The task scheduling apparatus according to claim 13, wherein the scheduling module further comprises:~~

~~an identification unit configured to obtain a task identifier for a task meeting the selection criteria; and~~

~~a generation unit configured to generate scheduling information for recording the task identifier.~~

~~15. The task scheduling apparatus according to claim 14, wherein~~

~~the identification unit is further configured to obtain, when a to-be-scheduled task is received, a task identifier for the to-be-scheduled task; and~~

~~the scheduling unit is specifically configured to obtain the task identifier for the to-be-scheduled task obtained by the identification unit when receiving the to-be-scheduled task; and schedule the to-be-scheduled task to a target cluster where dependent data for the to-be-scheduled~~

~~task locates if the task identifier of the to-be-scheduled task matches the task identifier in the scheduling information.~~

~~16. The task scheduling apparatus according to claim 14 or 15, wherein the identification unit comprises:~~

~~a determining subunit, configured to determine whether a type of the task is SQL;~~

~~a hash subunit, configured to perform hash processing on the task to obtain a hash digest if the task type is SQL, and use the hash digest as the task identifier; and~~

~~a numbering subunit, configured to use a number of the task as the task identifier when the task type is not SQL.~~

~~17. The task scheduling apparatus according to claim 13, wherein the selection criteria further comprises that the amount of output data is smaller than a second threshold;~~

~~and/or occupied cluster resources are less than a preset quota, wherein the occupied cluster resources comprise at least one of operating overhead, operating frequency, and cluster load.~~

~~18. The task scheduling apparatus according to any one of claims 10 to 15, wherein the network resources comprise at least one of a network bandwidth and a network bandwidth-delay product.~~

CLAIMS

What is claimed is:

1. A method comprising:

obtaining, by a task management device of a distributed system, network resources needed to perform cross-cluster reading and writing, the network resources associated with a task and comprising an amount of input data and an amount of output data for the task based on a history record of past executions of the task;

analyzing, by the task management device of the distributed system, the network resources by calculating an input-output ratio representing a proportion of the network resources needed for reading and writing for the task, the input-output ratio equal to a ratio of the amount of input data to the amount of output data; and

scheduling, by the task management device, the task according to the network resources.

2. The method of claim 1, the scheduling the task comprising scheduling the task to a target cluster where read dependent data is located if the network resources needed for reading are more than network resources needed for writing.

3. The method of claim 1, the scheduling the task further comprising:

determining whether the task meets a preset selection criteria, the preset selection criteria comprising that the input-output ratio is greater than a preset first threshold, the preset first threshold greater than one; and

scheduling the task to a target cluster where the dependent data read by the task locates.

4. The method of claim 3, further comprising:

obtaining, by the task management device, a task identifier for a task meeting the selection criteria; and

generating, by the task management device, scheduling information that includes the task identifier.

5. The method of claim 4, the scheduling the task to a target cluster where the dependent data

read by the task locates comprising:

obtaining the task identifier for the task when the task is received; and
scheduling the task to a target cluster where dependent data for the task locates if the task identifier of the task matches the task identifier in the scheduling information.

6. The method of claim 4, the obtaining a task identifier comprising:

determining whether a type of the task is Structured Query Language (SQL);
performing hash processing on the task to obtain a hash digest if the task type is SQL, and
using the hash digest as the task identifier; and
using a number of the task as the task identifier if the task type is not SQL.

7. The method of claim 3, the selection criteria comprising one or more of:

an amount of output data being smaller than a second threshold; and
occupied cluster resources being less than a preset quota, the occupied cluster resources comprising at least one of operating overhead, operating frequency, and cluster load.

8. The method of claim 1, the network resources comprising at least one of a network bandwidth and a network bandwidth-delay product.

9. The method of claim 1, further comprising determining, by the task management device, that the task is a cross-cluster task by determining, based on the history record, if the input or output data of the task is located in a cluster different from a current cluster of the task.

10. An apparatus comprising:

a processor; and
a storage medium for tangibly storing thereon program logic for execution by the processor,
the stored program logic comprising:
logic, executed by the processor, for obtaining network resources needed to perform cross-cluster reading and writing, the network resources associated with a task and comprising an
amount of input data and an amount of output data for the task based on a history record of past
executions of the task;

logic, executed by the processor, for analyzing the network resources calculating an input-output ratio representing a proportion of the network resources needed for reading and writing for the task, the input-output ratio equal to a ratio of the amount of input data to the amount of output data; and

logic, executed by the processor, for scheduling the task according to the network resources.

11. The apparatus of claim 10, the logic for scheduling the task comprising logic, executed by the processor, for scheduling the task to a target cluster where read dependent data is located if the network resources needed for reading are more than network resources needed for writing.

12. The apparatus of claim 10, the logic for scheduling the task further comprising:

logic, executed by the processor, for determining whether the task meets a preset selection criteria, the preset selection criteria comprising that the input- output ratio is greater than a preset first threshold, the preset first threshold greater than one; and

logic, executed by the processor, for scheduling the task to a target cluster where the dependent data read by the task locates.

13. The apparatus of claim 12, the stored program logic further comprising:

logic, executed by the processor, for obtaining a task identifier for a task meeting the selection criteria; and

logic, executed by the processor, for generating scheduling information that includes the task identifier.

14. The apparatus of claim 13, the logic for scheduling the task to a target cluster where the dependent data read by the task locates comprising:

logic, executed by the processor, for obtaining the task identifier for the task when the task is received; and

logic, executed by the processor, for scheduling the task to a target cluster where dependent data for the task locates if the task identifier of the task matches the task identifier in the scheduling information.

15. The apparatus of claim 13, the logic for obtaining a task identifier comprising:
logic, executed by the processor, for determining whether a type of the task is Structured
Query Language (SQL);
logic, executed by the processor, for performing hash processing on the task to obtain a hash
digest if the task type is SQL, and using the hash digest as the task identifier; and
logic, executed by the processor, for using a number of the task as the task identifier if the task
type is not SQL.

16. The apparatus of claim 12, the selection criteria comprising one or more of:
an amount of output data being smaller than a second threshold; and
occupied cluster resources being less than a preset quota, the occupied cluster resources
comprising at least one of operating overhead, operating frequency, and cluster load.

17. The apparatus of claim 10, the network resources comprising at least one of a network
bandwidth and a network bandwidth-delay product.

18. The apparatus of claim 10, the stored program logic further comprising logic, executed by the
processor, for determining that the task is a cross-cluster task by determining, based on the
history record, if the input or output data of the task is located in a cluster different from a
current cluster of the task.