What Is Claimed Is:

1. A computing system including one or more processors and one or more memories configured to perform operations for querying multiple data sources, comprising:

a plurality of server nodes, wherein each server node is configured to communicate with multiple data sources to access data therefrom; and

a second server node in communication with the one or more server nodes, wherein the second server node is configured to:

receive a query instance;

process the query instance to generate a plurality of relational query instances based upon, at least in part, the query instance, with each relational query instance of the plurality of relational query instances to be processed by the plurality of server nodes, wherein the plurality of relational query instances are distributed among the plurality of server nodes to extract portions of data from the multiple data sources in communication therewith corresponding to the plurality of relational query instances;

receive extracted data of the portions of data from the multiple data sources from each of the plurality of server nodes queried; and

aggregate the extracted data of the portions of data from the multiple data sources to respond to the query instance.

2. The computing system according to claim 1, wherein the plurality of server nodes are arranged in a hierarchical structure, and wherein the plurality of relational query instances are distributed among the plurality of server nodes directionally from upper server nodes to lower server nodes.

3. The computing system according to claim 1, wherein the second server node and the plurality of server nodes are arranged across at least two or more different enterprise networks.

4. The computing system according to claim 1, wherein the second server node is configured to receive the query instance further comprises the second server node is configured to receive a natural language query request related to obtaining data stored in the multiple data sources; and wherein the second server node is configured to process the query instance further comprises the second server node is configured to parse the natural language query request to generate the plurality of relational query instances.

5. The computing system according to claim 4, wherein the second server node is configured to parse the natural language query request further comprises the second server node is configured to utilize metadata information.

6. The computing system according to claim 5, wherein the second server node is further configured to generate at least one of an Abstract Semantic Graph (ASG) and an Abstract Semantic Tree (AST) for the natural language query request, to receive the metadata information.

7. The computing system according to claim 1, wherein the second server node is further configured to process the extracted data to generate a response based on the query instance.

8. The computing system according to claim 1, wherein the at least one server node of the one or more server nodes is further configured to optimize the plurality

of relational query instances to distribute among the multiple data sources and at least another server node of the plurality of server nodes.

9. A computer-implemented method for querying multiple data sources, the computer-implemented method comprising:

receiving a query instance;

processing the query instance to generate a plurality of relational query instances based upon, at least in part, the query instance, with each relational query instance of the plurality of relational query instances to be processed by a plurality of server nodes;

distributing the plurality of relational query instances among the plurality of server nodes to extract portions of data from the multiple data sources in communication therewith corresponding to the plurality of relational query instances;

receiving extracted data of the portions of data from the multiple data sources from each of the plurality of server nodes queried; and

aggregating the extracted data of the portions of data from the multiple data sources to respond to the query instance.

10. The computer-implemented method according to claim 9 wherein the distributing the plurality of relational query instances further comprises distributing the plurality of relational query instances among the plurality of server nodes directionally from upper server nodes to lower server nodes.

11. The computer-implemented method according to claim 9 wherein the distributing the plurality of relational query instances further comprises distributing the plurality of relational query instances among the plurality of server nodes across

at least two or more different enterprise networks.

12. The computer-implemented method according to claim 9 wherein the receiving the query instance further comprises receiving a natural language query request related to obtaining data stored in the multiple data sources; and wherein the processing the query instance further comprises parsing the natural language query request to generate the plurality of relational query instances.

13. The computer-implemented method according to claim 12 wherein the parsing the natural language query request further comprises utilizing metadata information.

14. The computer-implemented method according to claim 13 further comprising generating at least one of an Abstract Semantic Graph (ASG) and an Abstract Semantic Tree (AST) for the natural language query request, to receive the metadata information.

15. The computer-implemented method according to claim 9 further comprising processing the extracted data to generate a response based on the query instance.

16. The computer-implemented method according to claim 9 further comprising optimizing the plurality of relational query instances to distribute among the multiple data sources and at least one of another server node of the plurality of server nodes.

17. A computer program product residing on a computer readable storage medium having a plurality of instructions stored thereon which, when executed across

one or more processors, causes at least a portion of the one or more processors to perform operations according to any of claims 1-18.