

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
4 May 2006 (04.05.2006)

PCT

(10) International Publication Number
WO 2006/047009 A1

(51) International Patent Classification:
G06F 17/30 (2006.01)

(21) International Application Number:
PCT/US2005/031430

(22) International Filing Date:
2 September 2005 (02.09.2005)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
10/974,510 26 October 2004 (26.10.2004) US

(71) Applicant (for all designated States except US): **COM-
PUTER ASSOCIATES THINK, INC.** [US/US]; One
Computer Associates Plaza, Islandia, NY 11749 (US).

(72) Inventor; and

(75) Inventor/Applicant (for US only): **PAVAL, Eugen**
[RO/US]; 500 Peconic Street #156B, Ronkonkoma, NY
11779 (US).

(74) Agent: **TOERING, Rick, A.**; Pillsbury Winthrop Shaw
Pittman LLP, P.O. Box 10500, McLean, VA 22102 (US).

(81) Designated States (unless otherwise indicated, for every
kind of national protection available): AE, AG, AL, AM,
AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN,
CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI,

GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE,
KG, KM, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA,
MD, MG, MK, MN, MW, MX, MZ, NA, NG, NI, NO, NZ,
OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL,
SM, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC,
VN, YU, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every
kind of regional protection available): ARIPO (BW, GH,
GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM,
ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM),
European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI,
FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT,
RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA,
GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:


- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))
- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))

Published:

- with international search report
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments

[Continued on next page]

(54) Title: SYSTEM AND METHOD FOR PROVIDING A RELATIONAL APPLICATION DOMAIN MODEL



(57) Abstract: This disclosure provides a system and method for providing a relational application domain model. In one embodiment, the system is operable to or includes software that is operable to receive a structured query language (SQL) query for one or more services by at least a first of a plurality of resources. The software may be further operable to convert the SQL query into an object-oriented request compatible with at least the first resource and execute the object-oriented request at at least the first resource.



For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

SYSTEM AND METHOD FOR PROVIDING A RELATIONAL APPLICATION DOMAIN MODEL

TECHNICAL FIELD

This disclosure generally relates to distributed resourcing and, more
5 specifically, to a system and method for providing a relational application domain
model.

BACKGROUND

Conventional distributed applications typically include one or more
communication protocols between various components physically located on separate
10 machines or logically located in remote segments or environments. Usually there are
two logical pipes of communication: a data pipe that facilitates data exchange between
applications and a metadata pipe communicating type information associated with the
communicated data. The protocols facilitate or manage this exchange of the data and
metadata. At the highest level, a protocol may be considered one form of an API (or
15 Application Programming Interface). Typical APIs are difficult to modify to fit
changing needs or applications and are often difficult to enforce in a distributed team
environment. For example, various APIs may be defined and implemented during a
long period of time by various teams across an enterprise on multiple platforms.
Indeed, many enterprises include or utilize a plurality of different APIs to exchange
20 metadata with various agents in order to traverse a hierarchy of managed resources.

SUMMARY

This disclosure provides a system and method for providing a relational
application domain model. In one embodiment, the system is operable to or includes
software that is operable to receive a structured query language (SQL) query for one or
25 more services by at least a first of a plurality of resources. The system or software
may be further operable to convert the SQL query into an object-oriented request
compatible with at least the first resource and execute the object-oriented request at at
least the first resource. The details of one or more embodiments of the disclosure are
set forth in the accompanying drawings and the description below. Particular features,

objects, and advantages of the disclosure will be apparent from the description and drawings and from the claims.

DESCRIPTION OF DRAWINGS

FIGURE 1 illustrates a environment that provides a relational application domain model in accordance with one embodiment of the present disclosure; and

FIGURE 2 is a flowchart illustrating an example method for processing a client request through a relational application domain model in accordance with one embodiment of the present disclosure

DETAILED DESCRIPTION

FIGURE 1 illustrates a distributed application environment 100 for at least a portion of enterprise or data processing environment in accordance with one embodiment of the present disclosure. At a high level, environment 100 represents a relational application domain model and includes or is communicably coupled with some or all of server 102, one or more clients 104, and a plurality of data resources 106. For example, distributed application environment 100 may include server 102 that is operable to receive a Structured Query Language (SQL) query 150 from a user of one of the clients 104, convert query 150 to an object-oriented request associated with at least one of the data resources 106 referencing or storing data 160 requested by query 150, and execute the object-oriented request on the appropriate resource 106. Therefore, distributed application environment 100 may provide substantially uniform access to metadata exposed by agents associated with data resources 106. Put another way, server 102 may provide a SQL interface or other Application Programming Interface (API) for applications executing on clients 104 to resources 106. As used herein, "SQL" describes or includes any of the plurality of versions of the SQL relational database query and manipulation language such as, for example, SEQUEL, ANSI SQL, any other proprietary or public variant of SQL, or other suitable or generic database query language (XML). Distributed application environment 100 may be a distributed client/server system that allows users of clients 104 to submit queries 150 for execution on any of the plurality of external resources 106. But environment 100 may be a standalone computing environment or any other suitable environment, such as an administrator accessing data stored on server 102, without departing from the

scope of this disclosure. For example, resources 106 may be distributed across a plurality of virtual or logical partitions resident on server 102. The term “dynamically,” as used herein, generally means that certain processing is determined, at least in part, at run-time based on one or more variables. The term “automatically,”
5 as used herein, generally means that the appropriate processing is substantially performed by at least part of distributed application environment 100. It should be understood that “automatically” further contemplates any suitable administrator or other user interaction with environment 100 without departing from the scope of this disclosure.

10 Returning to the illustrated embodiment, server 102 includes memory 120 and processor 125 and comprises an electronic computing device operable to receive, transmit, process and store data associated with environment 100. For example, server 102 may be any computer or processing device such as a mainframe, a blade server, general-purpose personal computer (PC), Macintosh, workstation, Unix-based
15 computer, or any other suitable device. Generally, FIGURE 1 provides merely one example of computers that may be used with the disclosure. In other words, the present disclosure contemplates computers other than general purpose computers as well as computers without conventional operating systems. As used in this document, the term “computer” is intended to encompass a personal computer, workstation,
20 network computer, or any other suitable processing device. For example, although FIGURE 1 illustrates one server 102 that may be used with the disclosure, environment 100 can be implemented using computers other than servers, as well as a server pool. Server 102 may be adapted to execute any operating system 110 including z/OS, Linux-Intel or Linux/390, UNIX, Windows Server, or any other
25 suitable operating system. According to one embodiment, server 102 may also include or be communicably coupled with a web server and/or an SMTP server.

Memory 120 may include any memory or database module and may take the form of volatile or non-volatile memory including, without limitation, magnetic media, optical media, random access memory (RAM), read-only memory (ROM), removable
30 media, or any other suitable local or remote memory component. In this embodiment, illustrated memory 120 includes transformation agents 140 and transformation policies

145, but it may also include any other appropriate data such as a job history, a security or audit log, print or other reporting files, HTML files or templates, and others.

Each transformation agent 140 is any hardware, software, firmware, or other component associated with at least one data resource 106 and operable to convert
5 query 150 to an object-oriented form compatible with the particular one or more resources 106. Each agent may be further operable to call or execute Business Logic Layer (BLL) or Data Access Layer (DAL) logic as appropriate. For example, agent 140 may comprise an API implemented through a dynamic linked library (DLL), a daemon, an object, or other such software module. Each transformation agent 140
10 may be written or described in any appropriate computer language including C, C++, C#, Java, J#, Visual Basic, assembler, Perl, any suitable version of 4GL, and others or any combination thereof. It will be understood that while each transformation agent 140 is illustrated in FIGURE 1 as a single multi-tasked module, the features and functionality performed by this engine may be performed by multiple modules such as,
15 for example, a plurality of objects within a particular agent 140. Further, agents 140 may be collectively stored in one file or library or distributed across a plurality of files or libraries without departing from the scope of the disclosure. While illustrated as residing on server 102, each agent 140 may be located or executed on the respective data resource 106. In certain embodiments, transformation agent 140 may convert
20 query 150 to an object-oriented framework based one or more transformation policies 145. As used herein, "each" means one or all of a particular subset as appropriate.

Transformation policies 145 include any parameters, variables, mappings, algorithms, instructions, rules or other policies for dynamically converting from SQL to an object-oriented framework. Moreover, transformation policies 145 may also
25 define or implement various relationships between transformation agents 140 or other processes associated with data resources 106. For example, transformation policies 145 may include an enterprise-wide or a default mapping policy 145, thereby providing a uniform map from SQL to the appropriate resource 106. In another example, transformation policies 145 may store a plurality of individual API
30 transformation policies 145, each of which may be associated with an agent 140 or resource 106 and allow the particular agent to customize query 150 for the particular resource 105. In certain embodiments, transformation policies 145 allow server 102 to

quickly adapt to changing resources 106 and to provide a standard or substantially uniform interface to clients 104. In one embodiment, transformation policies 145 may comprise one or more tables stored in a relational database described in terms of SQL statements or scripts. In another embodiment, transformation policies 145 may store or define various data structures as text files, eXtensible Markup Language (XML) documents, Virtual Storage Access Method (VSAM) files, flat files, Btrieve files, comma-separated-value (CSV) files, internal variables, or one or more libraries. While illustrated separately, in a further embodiment, one or more transformation policies 145 may be rules or variables internal to certain transformation agents 140. In short, transformation policies 145 may comprise one variable, table, or file or a plurality of variables, tables, or files stored on one computer or across a plurality of computers in any appropriate format. Moreover, all or a portion of transformation policies 145 may be local or remote to server 102 without departing from the scope of this disclosure and store any type of appropriate data.

Server 102 also includes processor 125. Processor 125 executes instructions and manipulates data to perform the operations of server 102 such as, for example, a central processing unit (CPU), a blade, an application specific integrated circuit (ASIC), or a field-programmable gate array (FPGA). Although FIGURE 1 illustrates a single processor 125 in server 102, multiple processors 125 may be used according to particular needs and reference to processor 125 is meant to include multiple processors 125 where applicable. In the illustrated embodiment, processor 125 executes parsing engine 130, which performs at least a portion of the analysis or other processing of incoming SQL queries 150.

Parsing engine 130 could include any hardware, software, firmware, or combination thereof operable to receive queries 150 that are compatible with SQL, dynamically convert queries 150 to an object-oriented framework associated with one or more data resources 106, and automatically execute converted queries 150 on the resources 106 to gather requested data 160. For example, parsing engine 130 may receive a query 150 that is in SQL and output a copy of query 150 as one or more POJOs (Plain Old Java Objects). Parsing engine 130 may be further operable to aggregate the various data 160 received from a plurality of data resources 106. In certain embodiments, parsing engine 130 may include, reference, or execute one or

more transformation agents 140 associated with resources 106. In other words, parsing engine 130 and transformation agents 140 may be used interchangeably as appropriate without departing from the scope of the disclosure. Parsing engine 130 may be written or described in any appropriate computer language including C, C++,
5 C#, Java, J#, Visual Basic, assembler, Perl, any suitable version of 4GL, and others or any combination thereof. It will be understood that while parsing engine 130 is illustrated in FIGURE 1 as a single multi-tasked module, the features and functionality performed by this engine may be performed by multiple modules such as, for example, one or more agents 140. For example, parsing engine 130 may facilitate
10 interoperability among various transformation agents 140 and transparency to client 104. Further, while illustrated as internal to server 102, one or more processes associated with parsing engine 130 may be stored, referenced, or executed remotely. Moreover, parsing engine 130 may be a child or sub-module of another software module (not illustrated) without departing from the scope of this disclosure. In one
15 embodiment, parsing engine 130 may be referenced by or communicably coupled with applications executing on client 104.

Server 102 may also include interface 117 for communicating with other computer systems, such as client 104, over network 112 in a client-server or other distributed environment. In certain embodiments, server 102 receives queries
20 from internal or external senders through interface 117 for storage in memory 120 and/or processing by processor 125. Generally, interface 117 comprises logic encoded in software and/or hardware in a suitable combination and operable to communicate with network 112. More specifically, interface 117 may comprise software supporting one or more communications protocols associated with communications network 112
25 or hardware operable to communicate physical signals.

Network 112 facilitates wireless or wireline communication between computer server 102 and any other local or remote computer, such as clients 104. Indeed, while illustrated as two networks, 112a and 112b respectively, network 112 may be a continuous network without departing from the scope of this disclosure, so long as at
30 least portion of network 112 may facilitate communications between senders and recipients of queries 150 and results 160. In other words, network 112 encompasses any internal or external network, networks, sub-network, or combination thereof

operable to facilitate communications between various computing components in environment 100. Network 112 may communicate, for example, Internet Protocol (IP) packets, Frame Relay frames, Asynchronous Transfer Mode (ATM) cells, voice, video, data, and other suitable information between network addresses. Network 112
5 may include one or more local area networks (LANs), radio access networks (RANs), metropolitan area networks (MANs), wide area networks (WANs), all or a portion of the global computer network known as the Internet, and/or any other communication system or systems at one or more locations.

Client 104 is any local or remote computing device operable to receive requests
10 from the user via a GUI 116, a CLI (Command Line Interface), or other user interface. At a high level, each client 104 includes at least GUI 116 and comprises an electronic computing device operable to receive, transmit, process and store any appropriate data associated with environment 100. It will be understood that there may be any number of clients 104 communicably coupled to server 102. For example, illustrated clients
15 104 include one local client 104 and two clients external to the illustrated portion of enterprise 102. Further, "client 104," "recipient," and "user" may be used interchangeably as appropriate without departing from the scope of this disclosure. Moreover, for ease of illustration, each client 104 is described in terms of being used by one user. But this disclosure contemplates that many users may use one computer
20 or that one user may use multiple computers to submit or review jobs 150 via GUI 116. As used in this disclosure, client 104 is intended to encompass a personal computer, touch screen terminal, workstation, network computer, kiosk, wireless data port, wireless or wireline phone, personal data assistant (PDA), one or more processors within these or other devices, or any other suitable processing device. For example,
25 client 104 may comprise a computer that includes an input device, such as a keypad, touch screen, mouse, or other device that can accept information, and an output device that conveys information associated with the operation of server 102 or clients 104, including digital data, visual information, or GUI 116. Both the input device and output device may include fixed or removable storage media such as a magnetic
30 computer disk, CD-ROM, or other suitable media to both receive input from and provide output to users of clients 104 through the display, namely GUI 116.

GUI 116 comprises a graphical user interface operable to allow the user of client 104 to interface with at least a portion of environment 100 for any suitable purpose. Generally, GUI 116 provides the user of client 104 with an efficient and user-friendly presentation of data provided by or communicated within environment 5 100. For example, GUI 116 may be a front-end of an application executing on client 104 that is operable to submit SQL queries 150 to one or more data resources 106 through parsing engine 130. GUI 116 may comprise a plurality of customizable frames or views having interactive fields, pull-down lists, and buttons operated by the user. In one embodiment, GUI 116 presents information associated with queries 150 and buttons and receives commands from the user of client 104 via one of the input 10 devices. Moreover, it should be understood that the term graphical user interface may be used in the singular or in the plural to describe one or more graphical user interfaces and each of the displays of a particular graphical user interface. Therefore, GUI 116 contemplates any graphical user interface, such as a generic web browser or touch 15 screen, that processes information in environment 100 and efficiently presents the results to the user. Server 102 can accept data from client 104 via the web browser (*e.g.*, Microsoft Internet Explorer or Netscape Navigator) and return the appropriate HTML or XML responses using network 112. For example, server 102 may receive a job submission or data query 150 from client 104 using the web browser and then 20 execute the parsed query in the appropriate one or more operating resources 106.

Environment 100 includes, references, or is communicably coupled with a plurality of data resources 106. Data resource 106 typically comprises a distributed application, data repository, web server, job queue, media service, file system agent, or any other target machine or process operable to store, retrieve, generate, or otherwise 25 identify and collect data based on a received query 150. For example, resource 106 may include a plurality of services operable to communicate requested data 160 to data consuming services on client 104 based on queries 150 transformed by parsing engine 130. As illustrated by second resource 106b, resources 106 may also include an enterprise agent 140 residing on the target machine. It will be understood that the 30 target machine or process may be remote or on-site; further, the target machine may represent a separate process or repository residing on server 102 without departing from the scope of this disclosure. In other words, data resource 106 is any module or

component that receives queries 150 and retrieves data 160, normally (but not always) residing on the target machine, based on received query 150. In certain embodiments, some or all data resources 106 are each associated with an API or other object-oriented framework. Of course, two or more resources 106 may share the same API or may be
5 compatible with identical frameworks without departing from the scope of the disclosure. Moreover, each data resource 106 may include one server or other computer or may be distributed across a plurality of computers. In certain embodiments, retrieved or communicated data 160 is in tabular format. In other embodiments, data 160 may be in XML format. Regardless of the particular format,
10 data 160 is often operable to be normalized or self-describing. Once the data has been gathered, data resource 106 may communicate data 160 to server 102 or server 102 may retrieve data 160 from data resource 106, as appropriate.

In one aspect of operation, server 102 receives a SQL query 150 from one of the communicably coupled clients 104. Parsing engine 130 parses query 150 into one
15 or fields, parameters, or other elements to identify the requested data or other services and determines the location of the request services. For example, parsing engine 130 may identify and collect one or more fields requested by SQL query 150. Moreover, parsing engine 130 may identify a requested sort order, filter, limit, or other properties or parameters of query 150. In this example, these fields may be processed as an array
20 of strings. Once query 150 is parsed into the appropriate fields or other elements, parsing engine 130 identifies the one or more data resources 106, based on the determined location, and the appropriate transformation agents 140 that are compatible with the identified resources 106. In certain embodiments, parsing engine may load or instantiate the identified transformation agents 140, thereby possibly allowing quicker
25 conversion of SQL query 150 into an object-oriented request. Parsing engine 130 may also identify, load, or otherwise process the appropriate transformation policies 145 based on the identified agents 140. Next, parsing engine 130 maps, converts, or otherwise transforms SQL query 150 into one or more object-oriented requests in one or more object-oriented frameworks, each framework compatible with at least one of
30 the identified resources 106. Parsing engine then uses these object-oriented requests to obtain the requested data or other services from resources 106. For example, parsing 130 may use an API associated with the one or more resources 106 to execute the

object-oriented requests. In certain embodiments, at least a subset of the data resources 106 return data 160 in an object-oriented format operable to generate tabular results represented by tuples or name-value pairs. For example, a first data resource 106 may communicate results 160 as a cache of data objects that are XML-enabled.

5 These tabular results 160 may implement any or all of tables, views, relations, constraints, and other database tabular characteristics. Moreover, tabular results 160 may provide built-in data filtering, sorting, computation, transformation, aggregation, and other processes without departing from the scope of this disclosure.

FIGURE 2 is a flowchart illustrating example method 200 for processing a
10 client request or other SQL query 150 through a relational application domain model through parsing engine 130 in accordance with one embodiment of the present disclosure. Generally, method 200 describes one example technique for parsing engine 130 to convert standardized queries 150 to object-oriented queries or commands compatible with one or more data resources 106. It will be understood that
15 method 200 is for illustration purposes only and these or similar techniques may be performed at any appropriate time, including concurrently, individually, or in combination. The following descriptions will focus on the operation of parsing engine 130 in performing this method. But, server 102 contemplates using any appropriate combination and arrangement of logical elements implementing some or all of the
20 described functionality.

Method 200 begins at step 202, where server 102 receives a SQL query 150 from client 104, often via network 112. Next, parsing engine 130 identifies one or more data resources 106 associated with the received SQL query 150 at step 204. For example, query 150 may request data that resides on a first data resource 106a and a
25 second data resource 106b. In this example, the requested data may reside on multiple resources 106 because first data resource 106 is a production instance and second data resource 106 is an archive instance, first resource 106 includes one requested table and second resource 106 includes another requested table, or because of any other reason. Returning to the illustration, parsing engine 130 selects the first of the identified
30 resources 106 at step 206. Once selected, parsing engine 130 identifies, locates, or otherwise loads transformation agent 140 associated with the particular resource 106 at step 208. For example, if the selected first resource 106 is a mainframe DBMS

environment, then parsing engine 130 may load the particular transformation agent 140 that is compatible with API to the mainframe DBMS environment. In another example, if the selected first resource 106a is a job scheduler, then parsing engine 130 may load a transformation agent 140 operable to generate objects that may execute on
5 or in the job scheduling resource 106a. Next, parsing engine 130 parses the received query 150 into one or more distinct elements at step 210. As described above, these elements may be data fields, sort orders, filters, or any other data characteristic or parameter of query 150.

Once query 150 is suitably parsed, parsing engine 130 identifies the first of the
10 various elements at step 212. Once identified, parsing engine 130 maps the identified element to an appropriate class definition at step 214. This class definition may include a class, inheritance, method, or property definition without departing from the scope of the disclosure. In certain embodiments, this mapping may be based on one or more transformation policies 145. For example, parsing engine 130 may load one or
15 more policies 145 based on the identified data resources 106 (or agent 140) and locate a particular mapping entry using the identified query element. At decisional step 216, parsing engine 130 determines if there are remaining unprocessed elements in parsed query 150. If there are, and parsing engine 130 identifies the next element at step 218 and execution returns to step 214.

20 Once there no more elements at decisional step 216, parsing engine 130 instantiates one or more objects based on the mapped class definitions in step 220. At step 222, parsing engine 130 executes the instantiated object on the selected resource 106. It will be understood that execution of the particular object may include executing object using processor 125, thereby requesting information from resource
25 106, communicating the instantiated object to the particular resource 106 via a resident agent 140, creating and executing the object on resource 106, or any other appropriate execution. Next, parsing engine 130 receives, requests, or otherwise obtains certain data 160 from the selected resource 106 at step 224. Parsing engine 130 then determines if there are more resources 106 associated with received SQL query 150 at
30 decisional step 226. If there are more resources 106, then parsing engine 130 selects the next resource 106 at step 228 and execution returns to step 208.

Once the SQL query 150 has been suitably processed, then parsing engine 130 may collate or otherwise aggregate the received data 160 at step 230. For example, parsing engine 130 may cache data 160 as it is received at step 224. In another example, parsing engine 130 may generate a temporary table for storing data 160 upon receiving SQL query 150 and add rows to the table as resources 106 communicate the requested data 160. In certain embodiments, parsing engine 130 then converts data 160, whether collated or not, to a tabular format at step 232. For example, this may allow the result of queries 150 to be normalized and expressed as a single well-defined data type based on a relational concept of a table. In other words, received data 160 may be self descriptive and may support a number well-defined operations including filtering (vertically and horizontally), sorting, and iterating. Moreover, in certain embodiments, received data 160 may be immutable or uniform across at least a subset of managed resources or resources 106. Once data 160 has been appropriately processed, server 102 may communicate data 160 to client 104 at step 234. It will be understood that this communication may involve any other suitable processing including generation of web pages and others. Moreover, it will be understood that requesting client 104 and the receiving client 104 may or may not be the same user or computer.

The preceding flowchart and accompanying description illustrate exemplary method 200. In short, environment 100 contemplates using any suitable technique for performing these and other tasks. Accordingly, many of the steps in this flowchart may take place simultaneously and/or in different orders than as shown. Moreover, environment 100 may use methods with additional steps, fewer steps, and/or different steps, so long as the methods remain appropriate.

Although this disclosure has been described in terms of certain embodiments and generally associated methods, alterations and permutations of these embodiments and methods will be apparent to those skilled in the art. For example, parsing engine 130 and the example plurality of agents 140 may collectively be considered a super-API operable to accommodate a plurality of applications, data resources 106, and clients 104 and provide a single representation of metadata associated with data 160. Accordingly, the above description of example embodiments does not define or

constrain this disclosure. Other changes, substitutions, and alterations are also possible without departing from the spirit and scope of this disclosure.

WHAT IS CLAIMED IS:

1. Software for providing a relational application domain model operable to:
receive a structured query language (SQL) query for one or more services by at
5 least a first of a plurality of resources;
convert the SQL query into an object-oriented request compatible with at least
the first resource; and
execute the object-oriented request at at least the first resource.
- 10 2. The software of Claim 1, wherein the software operable to execute the
object-oriented request at at least the first resource comprises software operable to:
communicate the object-oriented request to at least the first resource; and
request execution of the object-oriented request by the first resource.
- 15 3. The software of Claim 1 wherein the software operable to convert the
SQL query into an object-oriented request compatible with at least the first resource
comprises software operable to:
map the SQL query to one or more classes in an object-oriented framework
compatible with at least the first resource; and
20 instantiate one or more objects based on the one or more mapped classes.
4. The software of Claim 3, further operable to parse the SQL query into a
plurality of fields and wherein the SQL query is mapped to one or more classes based
on the parsed fields.
- 25 5. The software of Claim 1, further operable to:
receive a result from at least the first resource;
convert the result to tabular format; and
communicate the tabular result to a client.

30

6. The software of Claim 1, wherein at least a subset of the plurality of resources comprise external resources.

7. The software of Claim 1, further operable to:
5 identify that the structured query language (SQL) query is for one or more services by a second of the plurality of resources;
convert the SQL query into a second object-oriented request compatible with the second resource; and
execute the second object-oriented request at the second resource.

10

8. The software of Claim 7 comprises:
a first agent operable to convert the SQL query into the object-oriented request compatible with at least the first resource; and
a second agent operable to convert the SQL query into the second object-
15 oriented request compatible with the second resource.

9. The software of Claim 1, the SQL query comprising a eXtensible Markup Language (XML) query.

20 10. The software of Claim 1, the SQL query comprising one of the following:
a select statement; or
an execute statement.

11. A method for providing a relational application domain model comprising:

receiving a structured query language (SQL) query for one or more services by at least a first of a plurality of resources;

5 converting the SQL query into an object-oriented request compatible with at least the first resource; and

executing the object-oriented request at at least the first resource.

12. The method of Claim 11, wherein executing the object-oriented request at at least the first resource comprises:

communicating the object-oriented request to at least the first resource; and requesting execution of the object-oriented request by the first resource.

13. The method of Claim 11 wherein converting the SQL query into an object-oriented request compatible with at least the first resource comprises:

15 mapping the SQL query to one or more classes in an object-oriented framework compatible with at least the first resource; and

instantiating one or more objects based on the one or more mapped classes.

20 14. The method of Claim 13, further comprising parsing the SQL query into a plurality of fields and wherein the SQL query is mapped to one or more classes based on the parsed fields.

15. The method of Claim 11, further comprising:

25 receiving a result from at least the first resource;

converting the result to tabular format; and

communicating the tabular result to a client.

16. The method of Claim 11, wherein at least a subset of the plurality of resources comprise external resources.

17. The method of Claim 11, further comprising:
5 identifying that the structured query language (SQL) query is for one or more services by a second of the plurality of resources;
converting the SQL query into a second object-oriented request compatible with the second resource; and
executing the second object-oriented request at the second resource.

10

18. The method of Claim 11, the SQL query comprising a eXtensible Markup Language (XML) query.

19. The method of Claim 11, the SQL query comprising one of the
15 following:
a select statement; or
an execute statement.

20. A system for providing a relational application domain model
20 comprising:
memory storing at least a portion of a data map; and
one or more processors operable to:
receive a structured query language (SQL) query for one or more services by at least a first of a plurality of resources;
25 convert the SQL query into an object-oriented request compatible with at least the first resource based, at least in part, on the data map; and
execute the object-oriented request at at least the first resource.

21. The system of Claim 20, wherein the one or more processors operable to execute the object-oriented request at at least the first resource comprises one or more processors operable to:

5 communicate the object-oriented request to at least the first resource; and
request execution of the object-oriented request by the first resource.

22. The system of Claim 20 wherein the one or more processors operable to convert the SQL query into an object-oriented request compatible with at least the first
10 resource comprises one or more processors operable to:

map the SQL query to one or more classes in an object-oriented framework compatible with at least the first resource; and

instantiate one or more objects based on the one or more mapped classes.

15 23. The system of Claim 22, the one or more processors further operable to parse the SQL query into a plurality of fields and wherein the SQL query is mapped to one or more classes based on the parsed fields.

24. The system of Claim 20, the one or more processors further operable to:
20 receive a result from at least the first resource;
convert the result to tabular format; and
communicate the tabular result to a client.

25. The system of Claim 20, wherein at least a subset of the plurality of
25 resources comprise external resources.

26. The system of Claim 20, the one or more processors further operable to:
identify that the structured query language (SQL) query is for one or more services by a second of the plurality of resources;
30 convert the SQL query into a second object-oriented request compatible with the second resource; and
execute the second object-oriented request at the second resource.

27. The system of Claim 26, wherein the one or more processors execute:
a first agent operable to convert the SQL query into the object-oriented request
compatible with at least the first resource; and

5 a second agent operable to convert the SQL query into the second object-
oriented request compatible with the second resource.

28. The system of Claim 20, the SQL query comprising a eXtensible
Markup Language (XML) query.

10

29. The system of Claim 20, the SQL query comprising one of the
following:

a select statement; or
an execute statement.

15

30. A system for providing a relational application domain model
comprises:


means for receiving a structured query language (SQL) query for one or more
services by at least a first of a plurality of resources;

20 means for converting the SQL query into an object-oriented request compatible
with at least the first resource; and

means for executing the object-oriented request at at least the first resource.


31. Software for providing a relational application domain model operable to:
- receive a structured query language (SQL) query for one or more services by at least a first and a second of a plurality of resources, at least a subset of the plurality of resources comprising external resources;
 - map the SQL query to one or more first classes in a first object-oriented framework compatible with the first resource;
 - map the SQL query to one or more second classes in a second object-oriented framework compatible with the second resource;
 - instantiate a first object-oriented request based on the one or more mapped first classes;
 - instantiate a second object-oriented request based on the one or more mapped second classes;
 - execute the first object-oriented request at the first resource;
 - receive a first result from the first resource;
 - execute the second object-oriented request at the second resource;
 - receive a second result from the second resource;
 - convert the first and second results to tabular format; and
 - communicate the tabular result to a client.

1/2



2/2

FIG. 2



INTERNATIONAL SEARCH REPORT

International application No
PCT/US2005/031430

A. CLASSIFICATION OF SUBJECT MATTER G06F17/30		
According to International Patent Classification (IPC) or to both national classification and IPC		
B. FIELDS SEARCHED Minimum documentation searched (classification system followed by classification symbols) G06F		
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched		
Electronic data base consulted during the international search (name of data base and, where practical, search terms used) EPO-Internal, WPI Data, IBM-TDB, INSPEC, COMPENDEX, PAJ		
C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 02/35395 A (ENTIGEN CORPORATION; VLAHOS, HARRY; KASOW, CLAY, M) 2 May 2002 (2002-05-02) paragraphs '0021! - '0054! paragraphs '0069!, '0070! paragraphs '0081! - '0089! paragraphs '0121! - '0143! -----	1-31
X	US 2002/023091 A1 (SILBERBERG DAVID P ET AL) 21 February 2002 (2002-02-21) paragraphs '0018!, '0019! paragraphs '0044!, '0045! paragraphs '0051! - '0057! paragraphs '0081! - '0088! ----- -/--	1-31
☒ Further documents are listed in the continuation of Box C. ☒ See patent family annex.		
* Special categories of cited documents : 'A' document defining the general state of the art which is not considered to be of particular relevance 'E' earlier document but published on or after the international filing date 'L' document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified) 'O' document referring to an oral disclosure, use, exhibition or other means 'P' document published prior to the international filing date but later than the priority date claimed 'T' later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention 'X' document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone 'Y' document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art. '&' document member of the same patent family		
Date of the actual completion of the international search		Date of mailing of the international search report
23 January 2006		08/03/2006
Name and mailing address of the ISA/ European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Tx. 31 651 epo nl, Fax: (+31-70) 340-3016		Authorized officer Díaz de Lezana, C

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2005/031430

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	LIU L ET AL: "An adaptive approach to query mediation across heterogeneous information sources" COOPERATIVE INFORMATION SYSTEMS, 1996. PROCEEDINGS., FIRST IFCIS INTERNATIONAL CONFERENCE ON BRUSSELS, BELGIUM 19-21 JUNE 1996, LOS ALAMITOS, CA, USA, IEEE COMPUT. SOC P, US, 19 June 1996 (1996-06-19), pages 144-156, XP010200757 ISBN: 0-8186-7505-5 the whole document	1-31
X	----- NING WANG ET AL: "Versatile: a scalable CORBA-based system for integrating distributed data" INTELLIGENT PROCESSING SYSTEMS, 1997. ICIPS '97. 1997 IEEE INTERNATIONAL CONFERENCE ON BEIJING, CHINA 28-31 OCT. 1997, NEW YORK, NY, USA, IEEE, US, vol. 2, 28 October 1997 (1997-10-28), pages 1589-1593, XP010276314 ISBN: 0-7803-4253-4	1,11,20, 30
A	the whole document	2-10, 12-19, 21-29,31
A	----- ANONYMOUS: "Java Management Extensions Instrumentation and Agent Specification, v1.0" INTERNET ARTICLE, 'Online! 27 February 2004 (2004-02-27), pages 1-164, XP002363702 Retrieved from the Internet: URL:http://www.jcp.org/aboutJava/community process/first/jsr003/jmx_instr_agent.pdf> 'retrieved on 2006-01-16! the whole document	1-31

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2005/031430

Patent document cited in search report		Publication date		Patent family member(s)	Publication date
WO 0235395	A	02-05-2002	AU	2873902 A	06-05-2002
US 2002023091	A1	21-02-2002	NONE		