

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
18 March 2010 (18.03.2010)

PCT

(10) International Publication Number  
WO 2010/030919 A2

(51) International Patent Classification:  
G06F 15/16 (2006.01) G06F 9/06 (2006.01)

(21) International Application Number:  
PCT/US2009/056709

(22) International Filing Date:  
11 September 2009 (11.09.2009)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:  
61/097,059 15 September 2008 (15.09.2008) US  
61/122,695 15 December 2008 (15.12.2008) US

(71) Applicant (for all designated States except US):  
PALANTIR TECHNOLOGIES, INC. [US/US]; 100  
Hamilton Avenue, Suite 300, Palo Alto, California 94301  
(US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **AYMELOGLU, Andrew** [US/US]; 432 High Street, Palo Alto, California 94301 (US). **MIYAKE, Nicholas** [US/US]; 460 Channing Ave., Apt. A, Palo Alto, California 94301 (US). **BURR, Brandon** [US/US]; 569 Addison Ave., Apt. 4, Palo Alto, California 94301 (US). **CICERONE, Derek** [US/US]; 145 Hawthorne Ave., Apt. A, Palo Alto, California 94301 (US). **SIMLER, Kevin** [US/US]; 3035 Emerson St., Palo Alto, California 94306 (US). **TAN, Garry** [US/US]; 255 King St., Apt. 518, San Francisco, California 94107 (US).

(74) Agents: REES, Karl T. et al.; 2055 Gateway Place, Suite 550, San Jose, California 95110 (US).

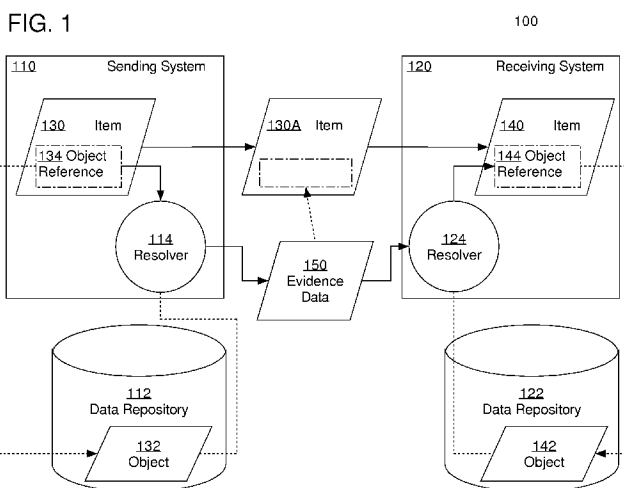
(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, 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, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished upon receipt of that report (Rule 48.2(g))

(54) Title: SHARING OBJECTS THAT RELY ON LOCAL RESOURCES WITH OUTSIDE SERVERS



(57) Abstract: To prevent potential confusion when referencing an object by name, a party may supplementary or alternatively include referential metadata, other than the name by which the object is known to the party, that describes the referenced object. A party receiving information that includes such referential metadata for an object may utilize the referential metadata to identify a local object that matches or may substitute for the object referenced by the sending party. For example, a broker operating a financial system may generate a chart based on a first named object. The broker may wish to send the chart to another broker operating a different financial system. Since the first name object may be known by a different name in the different financial system, the financial system may replace the name of the object with referential metadata prior to sending the chart.

WO 2010/030919 A2

## SHARING OBJECTS THAT RELY ON LOCAL RESOURCES WITH OUTSIDE SERVERS

### INVENTION

**[0001]** Embodiments of the invention described herein relate generally to information sharing, and, more specifically, to techniques for sharing information that relies upon local references across servers for which those references may have different meanings.

### BACKGROUND

**[0002]** The approaches described in this section are approaches that could be pursued, but not necessarily approaches that have been previously conceived or pursued. Therefore, unless otherwise indicated, it should not be assumed that any of the approaches described in this section qualify as prior art merely by virtue of their inclusion in this section.

**[0003]** Collaboration between multiple parties often requires that the parties share various items of information, such as documents, tables, spreadsheets, graphs, and charts. When sharing these items of information, the parties and/or the information sharing mechanisms relied upon by those parties often reference certain objects by name. For example, the parties may need to exchange a document that analyzes the performance of a certain financial instrument. The document may refer to the financial instrument by a name, such as "ALPHAAlpha."

**[0004]** In some situations, a referenced name may not have the same meaning to all parties involved in the communication. For example, a first party may know a particular financial instrument by the name "ALPHAAlpha," while a second party may know the particular financial instrument by the name "109328." When the first party attempts to communicate information that involves a financial instrument named "ALPHAAlpha" to the second party, because the second party knows that financial instrument by the name "109328," the second party will be unable to determine the financial instrument to which the first party is referring.

**[0005]** In fact, in some cases the second party may associate the name "ALPHAAlpha" with a different object than the particular financial instrument. Since the second party will believe that the first party is communicating information that involves a different object than was intended, the second party may entirely misunderstand the information communicated by the first party. Moreover, since the first party cannot be certain that the first party's version of "ALPHAAlpha" is associated with the same data as the second party's version of "ALPHAAlpha," the first party may hesitate to share any information at all, lest the second party be misled as to the meaning of the information.

**[0006]** In some cases, the second party in the above information exchange scenario may lack access to data for a financial instrument that exactly corresponds to the particular financial instrument, but may still have access to an object labeled “ALPH” whose data is close enough to the exchanged information so as to render the object an effective substitute for the particular financial instrument. If the second party were capable of recognizing that an object to which it has access could serve as a substitute for the financial instrument, the second party would be able to understand the information that the first party wishes to share. However, again because of the differences in nomenclature, the second party will be unable to identify the object as a substitute for the financial instrument referenced by the first party.

**[0007]** The inability of a party to correctly identify an object or a substitute for an object referenced by another party can severely cripple an exchange of information. For example, in many cases, parties share items of information that include metadata in place of actual data. This metadata may include the name of one or more objects associated with the actual data, and by which the actual data may be located. The actual data utilized in a document is therefore retrieved dynamically based on the metadata.

**[0008]** Sharing metadata instead of actual data is particularly common when each party may rely upon a separate data source. For example, sharing metadata instead of actual data may be desirable for security or efficiency purposes, in that the parties may be unwilling or unable to transmit actual data. Or, sharing metadata instead of actual data may be useful for the purposes of ensuring that a shared document or graph is always populated with the most recent data available to a party.

**[0009]** However, because each party is relying on a separate data source, there is no way to guarantee that the parties name various objects in a consistent fashion. Nor is there any guarantee that the parties have access to the same type of data. For example, one party may subscribe to data from Bloomberg, while another party may subscribe to data from Reuters. These potential differences between data sources complicate any exchange of documents that rely upon metadata instead of actual data. For example, if a second party is unable to rely upon metadata in a document shared by a first party to correctly identify an object in the second party’s data repository that corresponds to an object that the first party intended to reference, the second party will be unable to understand the information conveyed in that document.

#### BRIEF DESCRIPTION OF THE DRAWINGS

**[0010]** The present invention is illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings and in which like reference numerals refer to similar elements and in which:

- [0011] FIG. 1 illustrates an architecture 100 for sharing information between a sending system and a receiving system;
- [0012] FIG. 2 is a flowchart illustrating a method for sending information from a sending system to a receiving system, wherein the information references non-shared data;
- [0013] FIG. 3 is a flowchart illustrating a method for, at a receiving system, interpreting information received from a sending system, wherein the information references non-shared data from the sending system; and
- [0014] FIG. 4 is a block diagram that illustrates a computer system upon which an embodiment may be implemented.

#### DETAILED DESCRIPTION

[0015] In the following description, for the purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be apparent, however, that the present invention may be practiced without these specific details. In other instances, well-known structures and devices are shown in block diagram form in order to avoid unnecessarily obscuring the present invention.

[0016] Embodiments are described herein according to the following outline:

- 1.0. General Overview
- 2.0. Example Architecture for Sharing Information
- 3.0. Functional Description
- 4.0. Implementation Examples
  - 4.1. Evidence Data
  - 4.2. Generating Evidence Data
  - 4.3. Matching
  - 4.4. Miscellaneous
- 5.0. Implementation Mechanism—Hardware Overview
- 6.0. Extensions and Alternatives

#### 1.0. GENERAL OVERVIEW

[0017] In an embodiment, to prevent potential confusion when referencing an object by name, a party may utilize referential metadata, other than the name by which the object is known to the party, to describe the referenced object. This referential metadata is hereinafter referred to as “evidence.” A party receiving information that includes evidence for an object may utilize

that evidence to identify a local object that matches or may substitute for an object referenced by the sending party.

**[0018]** According to an embodiment, a first user at a first system sends an item of information to a second user at a second system. The item of information relies upon data associated with a first object stored at a first data source accessible to the first system. The item of information therefore comprises metadata indicating a name by which the first system references the first object. Prior to sending the item to the second system, the first system generates evidence of the first object. The first system includes the evidence with the item of information when sending the item to the second system. In this manner, the first system provides the second system data by which the second system may identify or verify objects that the second system may substitute for the first object.

**[0019]** According to an embodiment, a second user at a second system receives an item of information from a first user at a first system. The item of information relies upon data from a first object at the first system. The second system also receives evidence for the first object. The second system determines that a second object, stored at a second data source accessible to the second system, matches the evidence. The second system relies upon the second object in place of the first object. In this manner, the second system ensures that the item of information, as seen by the second user, is based upon data that sufficiently corresponds to the data that the first user intended the item of information to be based upon.

**[0020]** According to an embodiment, the evidence describes characteristics that are desirable in or essential to any object that is substituted for the referenced object. According to an embodiment, the evidence uniquely and unambiguously identifies an object. According to an embodiment, the evidence allows one to determine one or more objects that may potentially match the referenced object, along with a confidence level for each potentially matching object. According to an embodiment, the evidence may comprise one or more of the following: data indicating values for one or more attribute fields stored within or associated with the referenced object, wherein at least one value is a value other than the referenced object's name; metadata indicating the referenced object type; and metadata indicating the types and quantities of data stored within or associated with the referenced object.

## 2.0. EXAMPLE ARCHITECTURE FOR SHARING INFORMATION

**[0021]** FIG. 1 illustrates an architecture 100 for sharing information between a sending system 110 and a receiving system 120.

**[0022]** Sending system 110 and receiving system 120 are systems configured for interacting with data in data repositories 112 and 122, respectively. Systems 110 and 120 may comprise

one or more computers operated by one or more users. Systems 110 and 120 may, for example, include a computer upon which executes one or more applications for accessing and analyzing the data in data repositories 112 and 122. Systems 110 and 120 may also include one or more servers for managing data in data repositories 112 and 122, as well as for managing various items of information, such as item 130. In an embodiment, systems 110 and 120 each include a financial analysis server, such as a Palantir Finance server, and a financial analysis client, such as a Palantir Finance client. However, systems 110 and 120 need not necessarily operate the same server and client software.

**[0023]** Data repositories 112 and 122 may be any components of systems 110 or 120 that are capable of storing data, including a database, database server, or file system. Data repositories 112 and 122 store a wide variety of objects, including objects 132 and 142, respectively. For example, Objects 132 and 142 may comprise any structured data, including, but not limited to, structured data related to financial instruments, tables, indexes, and entities. Objects 132 and 142 are said to be associated with the data of which objects 132 and 142 are comprised. Objects 132 and 142 may also be associated with data external to objects 132 and 142. For example, an object may contain one or more fields indicating values by which information associated with objects 132 and 142 are indexed in tables or other data structures of data repositories 112 and 122.

**[0024]** Objects 132 and 142 are identified in their respective systems by reference names 134 and 144, respectively. Reference names 134 and 144 are indicated by metadata associated with or stored in one or more fields of their respective objects. Reference names 134 and 144 uniquely identify their respective objects in their respective systems. In some embodiments, objects 132 and 142 are known within their system only by their respective reference names, while in other systems, objects 132 and 142 may instead be identified by other data fields as well. In some embodiments, objects 132 and 142 are of a common object type, and reference names 134 and 144 are both from a common field associated with or defined for that object type. Reference names 134 and 144 may also be referred to as labels or identifiers.

**[0025]** According to an embodiment, all objects within a system are uniquely identified by their name. According to an embodiment, all objects within a system are uniquely identified by their name and object type. According to an embodiment, all systems involved in the sharing of information utilize a common technique for identifying the name of an object. For example, all systems may expect that an object's name will be specified in a certain field of data stored in or associated with the object.

**[0026]** Sending system 110 shares items of information, such as item 130, with receiving system 120. Item 130 may be any of a wide variety of items that can be shared between multiple parties, including, but not limited to, documents, tables, spreadsheets, graphs, and charts. Item 130 relies, at least in part, on data associated with one or more objects, including object 132, from data repository 112. For example, item 130 may be a graph of data stored within object 132. Rather than storing the data from object 132 directly within item 130, system 110 references object 132 by reference name 134 from within item 130.

**[0027]** Prior to sharing item 130 with receiving system 120, system 110 detects that item 130 references object 132. In response, system 110 utilizes resolver component 114 to generate evidence data 150 for object 132. Evidence data 150 is data that describes characteristics of object 132, including characteristics other than reference name 134. Techniques for generating evidence data 150 are more fully described in subsequent sections.

**[0028]** Resolver 114 may be any software or hardware component of system 110 capable of producing evidence data 150 based on characteristics of object 132. System 110 may feature a single resolver 114 capable of generating evidence data 150 for any object in data repository 112. Alternatively, system 110 may feature a number of different resolvers 114, each adapted to producing evidence data 150 for a particular type of object in data repository 112.

**[0029]** System 110 then sends the evidence data 150 with a copy of item 130, i.e. item copy 130A, to receiving system 120. In some embodiments, system 110 embeds evidence data 150 directly within item copy 130A. In some embodiments, system 110 replaces reference name 134 with evidence data 150. In some embodiments, system 120 stores evidence data 150 separately from item copy 130A. System 120 may then send evidence data 150 to system 120 at the same time as item copy 130A, or at another time via mapping data that identifies evidence data 150 as being associated with reference name 134.

**[0030]** Upon receiving item copy 130A and evidence data 150, receiving system 120 determines that item copy 130A relies upon data associated with a referenced object. Using evidence data 150, system 120 attempts to determine an object in its data repository 122 that matches, or can substitute for, the referenced object. System 120 accomplishes this task by feeding evidence data 150 to a resolver component 124.

**[0031]** Resolver 124 may be any hardware or software component of system 120 capable of determining an object accessible to system 120 that corresponds to evidence data 150. Resolver 124 accepts, as input, evidence data 150. Resolver 124 then determines that an object 142 in data repository 122 corresponds to the evidence data 150. Resolver 124 then identifies object 142 to system 120 as being a match or substitute for the object referenced in item copy 130A.

**[0032]** System 120 then saves item copy 130A as item 140. In contrast to item 130, which included a reference name 134 referencing object 132, item 140 includes reference name 144 referencing object 142. In embodiments where item copy 130A included reference name 134, system 120 simply changes reference name 132 to reference name 142 when saving item copy 130A as item 140. In embodiments where item copy 130A includes evidence data 150 in lieu of reference name 134, system 120 substitutes reference name 144 for evidence data 150.

**[0033]** Architecture 100 is but one of many environments in which the techniques described herein may be utilized. For example, other environments may feature additional systems and/or data repositories. Still other environments may feature only one system operated by a first party who relies upon objects from a first set of data and a second party who instead relies on objects from a second set of data. Still other environments may feature a single data repository to which a sending party has different access rights than the receiving party.

### 3.0. FUNCTIONAL DESCRIPTION

**[0034]** FIG. 2 is a flowchart illustrating a method for sending information from a sending system to a receiving system, wherein the information references non-shared data.

**[0035]** At step 210, a sending system, such as system 110, receives instructions to send an item, such as item 130, to a receiving system, such as system 120. For example, a user may be working with a financial analysis document at a financial analysis software client or other document viewing application executing within the sending system. Via the client, the user may instruct the financial analysis server to send the financial analysis document to a user at the receiving system.

**[0036]** At step 220, the sending system may detect that the item references one or more objects, such as object 132, within a data repository, such as data repository 112, that may not be shared with or otherwise accessible to the receiving system. For example, the financial analysis server may detect that the financial analysis document contains several graphs and tables that rely upon data associated with financial instrument objects named ALPH and BETA. These associations may be detected, for instance, because the financial analysis document may include metadata referencing the objects by their names. The metadata may be included, for instance, in instructions that cause a document viewing application to dynamically construct the document's graphs and tables from data stored in the sending system's data repository in association with the named objects. In some embodiments, the financial analysis server may assume that all objects referenced in the item are inaccessible to the receiving system, or that the receiving system does not have access to the same underlying data relied upon by the referenced objects. In other embodiments, the financial analysis server may further employ additional steps to determine



whether or not particular referenced objects and/or data repositories are accessible to the receiving system.

**[0037]** At step 230, the sending system generates evidence data, such as evidence data 150, for each of the one or more referenced objects. For example, the financial analysis server may send data for the ALPH and BETA financial instrument objects to an evidence generation component. Using techniques such as those discussed in later sections of this application, the evidence generation component may produce and return evidence data for both ALPH and BETA.

**[0038]** At step 240, the sending system may send a copy of the item and the evidence data to the receiving system. The evidence data may be inserted into the item copy along with or in place of the names of one or more referenced objects. For example, the financial analysis server may send a copy of the financial analysis document to the receiving system, in which copy the financial analysis server has replaced the metadata that identifies the names ALPH and BETA with the generated evidence data. Evidence data may instead accompany the item as separate data. Or, evidence data may be sent at a different time than the item.

**[0039]** In some embodiments, evidence data may be pre-compiled for certain objects. In such embodiments, step 230 precedes the other steps of FIG. 2. In some embodiments, the sending system may already have sent to the receiving system mapping data indicating evidence data for various objects at the sending system. In these embodiments, step 220 may therefore be optional, while the sending of the evidence data aspect of step 240 may precede the other steps of FIG. 2.

**[0040]** FIG. 3 is a flowchart illustrating a method for, at a receiving system, interpreting information received from a sending system, wherein the information references non-shared data from the sending system.

**[0041]** In step 310, the receiving system receives a copy of an item, such as item copy 130A, and evidence data from the sending system. The item copy and evidence data may be received together or separately. For example, a financial analysis server at the receiving system may receive the copy of a financial analysis document from the financial analysis server at the sending system.

**[0042]** In step 320, the receiving system determines that the received item copy relies on data associated with one or more referenced objects that may not be accessible to the receiving system. For example, the financial analysis server at the receiving system may detect that the copy of the financial analysis document contains data that, when interpreted by a document viewer, causes the document viewer to retrieve and graph data associated with objects were

stored at the sending system—namely, the financial instrument objects ALPH and BETA. In some embodiments, the receiving system makes this determination implicitly, based on detecting that the financial analysis document contains or was received with evidence data. In some embodiments, all referenced objects are assumed to be inaccessible.

**[0043]** In step 330, for each of the one or more referenced objects relied upon by the item copy, the receiving system determines an internal object that matches the evidence data for the referenced object. Like object 140, the internal object is accessible to the receiving system via, for instance, a data repository connected to the receiving system, such as data repository 122. For example, the receiving system may have received a set of evidence data for each of two referenced financial instrument objects in the received item copy. The receiving system may send the sets of evidence data to a resolving component for financial instrument objects. The resolving component may in turn determine that two financial instrument objects stored within the receiving system's data repository—named ALPHALPHA and BETA—match the two sets of evidence data, respectively. A match may be determined using any algorithm capable of identifying an object that may be substituted for the object for which the evidence data was generated. Some example algorithms are discussed in subsequent sections.

**[0044]** In step 340, the receiving system may substitute each identified internal object for the corresponding referenced object from the sending system. The manner of substitution is not important, so long as the substitution results in receiving system relying on, for use with the received item, data associated with the internal objects instead of the referenced objects. For example, the financial analysis server at the receiving system may save the copy of the financial analysis document to the receiving system. In the saved copy of the financial analysis document, the financial analysis server may replace all metadata referencing the names ALPH and BETA (or evidence data that had been substituted for the names ALPH and BETA) with metadata referencing ALPHALPHA and BETA. The financial analysis document may subsequently be accessed by a user using a document viewing application at the receiving system. Because of the substitution, instead of failing or graphing incorrect data, the document viewing application will construct graphs and tables within the financial analysis document based on data associated with the ALPHALPHA and BETA objects in the receiving system's own data repository.

**[0045]** FIG. 2 and FIG. 3 have been discussed with respect to the particular example of a financial analysis document shared between two financial analysis servers. However, the item described in FIG. 2 and FIG. 3 may in fact be any item of information, and the sharing parties may be any server, software, or component of any computer-based system. Moreover, while the

given examples have been with respect to objects that represent financial instruments, the objects described in FIG. 2 and FIG. 3 may be objects that represent any type of data or entity.

#### 4.0. IMPLEMENTATION EXAMPLES

##### 4.1. EVIDENCE DATA

**[0046]** Any type of data capable of communicating information that identifies a match or substitute for an object may be utilized as evidence data. Specific examples of evidence data are described below, however the scope of the evidence data is by no means limited to the described examples.

**[0047]** In an embodiment, objects may conform to pre-defined type definitions. For example, some pre-defined object types may include equities, bonds, derivatives, commodities, funds, persons, and entities. Accordingly, evidence data for an object may include data indicating the object type.

**[0048]** Evidence data may also include data indicating values associated with the object, such as the value of a specific field of data inside of the object, or values associated with specific object metadata. For example, an object of the type "equity" may include fields labeled ISIN, EXCH, TICKER, and COUNTRY. The evidence data generated for the object may include values for one or more of these fields. The values may be stored within the evidence data, for example, as a name value pair such as "ISIN=112233445566". Values may also be patterns (e.g. regular expressions) instead of actual values stored within the object. The included values may be of any data type, including textual, numerical, graphical, or audiovisual.

**[0049]** Evidence data may also include data describing the structure of the data associated with the object. For example, the evidence data may indicate a count of the fields, tables, columns, rows, and/or other elements associated with object. As another example, the evidence data may indicate the names, types, and/or sizes of some or all of the data elements associated with the object. As another example, evidence data may include data indicating specific dates or time ranges covered by the data associated with the object. As another example, evidence data may indicate a frequency or granularity of the data associated with the object.

**[0050]** Evidence data may be in any format capable of communicating the above information, including XML, plain text, code-based, and binary data formats. The evidence data may or may not include the evidenced object's name.

**[0051]** For example, the following XML data may represent evidence for an S&P 500 object. The evidence data includes the object model type, name, and symbol property value ('SPX').

```
<object class=" StockMarketIndex">
```

```

<name>S&P 500 INDEX</name>
<modelId>af11e33b-f54f-4a57-aec9-176e107c7232</modelId>
<modelType> StockMarketIndex</modelType>
<propertyTypes>
  <values>
    <PropertyType>
<canonicalName>com.palantir.finance.SYMBOL</canonicalName>
    </PropertyType>
    <PropertyValue>
      <value class="string">SPX</value>
      <timeRange/>
    </PropertyValue>
  </values>
</propertyTypes>
</object>

```

#### 4.2. GENERATING EVIDENCE DATA

**[0052]** According to an embodiment, a system may generate evidence data based on evidence generation rules that dictate which of the above described elements should be included in the evidence data. Some or all of the rules may be hard-coded within one or more evidence generating components, such as resolver 114, of the system. For example, the system may comprise evidence generating components for each type of object the system encounters, or the system may have a universal evidence generating component. Some or all of the rules may instead be user-configurable via, for instance, an XML file, script, or user interface. For example, a system administrator may create evidence generation scripts for each object encountered by the system.

**[0053]** When a system or a component thereof generates evidence data for an object, the system may follow rules that cause the system to include within the evidence data criteria that uniquely identifies the object. For example, an object may be of a pre-defined object type. The system may have some knowledge of the nature of various fields of data associated with that data type. Based on this knowledge, the system may be able to determine that no two objects of the same type should have the same value or combination of values for one or more fields common to that object type. Accordingly, the system may generate evidence data that includes criteria indicating the values for those one or more fields.

**[0054]** For example, an object of the type “equity” may include an ISIN number. The system may recognize that no two equities should have the same ISIN number. Accordingly, the system may use the ISIN number and the object type as criteria in the evidence data. As another

example, an equity may lack an ISIN number but include values for EXCH and TICKER fields. The system may be able to determine that no other equity should have the same values for those two fields. Thus, the values for those two fields may be used as criteria in the evidence data.

**[0055]** The system may also follow rules that cause the system to generate evidence data that includes substitutability criteria. The substitutability criteria indicate characteristics that a substitute for the evidenced object should or may possess. The system may create substitutability criteria instead of or in addition to data that uniquely identifies an object. For example, the system may include substitutability criteria in evidence data when the system is unable to generate data that uniquely identifies the object. As another example, the system may insert substitutability criteria in evidence data to provide another system with information about potential substitutes for an object, should the other system be unable to locate an exact match for the object. As another example, substitutability criteria may also be useful for verifying that a matching object actually contains the data that the sending system expects the matching object to contain.

**[0056]** As an example of substitutability criteria, the system may generate evidence data for an equity object that includes values for the TICKER and COUNTRY fields. Although the system cannot guarantee that TICKER and COUNTRY will uniquely identify an object, the system may nonetheless include these fields in the evidence data, since these fields may still effectively identify a match or at least a substitute for the evidenced object. As another example, the object for which the system is generating evidence may contain data indicating a price for each week of the year 2007. The system may generate evidence data for the object that includes criteria for a substitute object, such as a requirement that the substitute object should likewise have price data for each week of the year 2007.

**[0057]** The system may further generate evidence data that indicates which substitutability criteria are required of the substitute object as opposed to which characteristics are desirable of the substitute object. For example, an object may be required to have data for each week of the year 2007, while it may only be desirable that the object include data for each week of 2006.

**[0058]** When generating evidence data for an object referenced in a shared item, a system may also follow rules that consider how the shared item uses the data associated with the object. For example, when generating evidence data, a system may further narrow data indicating time ranges covered by the object to just those time ranges utilized by the shared item. Thus, even though the object contains data for an entire year, if the shared item only utilizes data for a particular month of that year, the criteria in the evidence data may indicate that a substitute for the object need only contain data for that particular month. As another example, if an object

contains data indicating both opening and closing prices, but the shared item only utilizes closing prices, the criteria in the evidence data may indicate that a substitute for the object need only contain data for the closing prices.

**[0059]** The system may further or alternatively adhere to other rules having other objectives when generating evidence data.

#### 4.3. MATCHING

**[0060]** As with the event-generation process, a system may also rely on resolution rules for determining objects that match or may substitute for an object evidenced by evidence data.

This process of determining a match or substitute may be known as “resolving” the evidence data to an object. For example, a system may feature one or more resolver components, such as resolver 124, that resolve evidence data to objects based on a set of resolution rules. As with the evidence generation rules, some or all of the resolution rules may be hard-coded or user-configurable. According to an embodiment, both resolution and event generation are performed by the same system components. Thus, event generation rules and resolution rules may be intermingled or the same.

**[0061]** When a system attempts to resolve evidence data to an object, the system may follow rules that cause it to search for objects having associated data that matches unique identifying criteria stored within the evidence data. The system may instead or additionally follow rules that cause it to search for objects whose data matches substitutability criteria stored within the evidence data. In both cases, matching may be performed based on any suitable technique, including textual, phonetic, or pattern-based comparisons. In some embodiments, the search for matching objects is performed without any consideration of an object’s name. In other embodiments, the object’s name may also be considered.

**[0062]** Certain criteria specified in the evidence data may be considered “required” criteria, based on a designation in the evidence data or on a resolution rule. The system may eliminate any object that does not possess a required criterion as a match. Thus, objects that match uniquely identifying criteria may nonetheless be screened according to whether or not the objects also match required substitutability criteria. For example, if the system contains an object whose unique identifier matches that described in the evidence data, but whose data does not cover a required date range specified in the criteria, the system may determine that the object does not match the evidence data. Or, if the unique identifier is a strong match, such as is typically the case with ISIN values, the system would inform a user that a match was found, but data for the required date range does not exist.

**[0063]** According to an embodiment, multiple potential matches may be identified during the matching process. In such embodiments, the system may request that a user identify the object to which the evidence data should be resolved. The system may, for instance, prompt the user via a user interface to identify which object in a list of potential matching objects the system should rely upon for the shared item. The system may then use the identified object as the match for the evidence data.

**[0064]** According to an embodiment, the system follows rules that cause it to calculate a level of confidence that an object in the system matches the evidence data. The system may calculate a level of confidence for all objects in the system. Alternatively, the system may calculate a level of confidence for just those objects that match required criteria, such as those objects that are of the same object type as the evidenced object. The level of confidence, which may be based on any decimal or enumerated scale, indicates the degree to which an object matches the evidence data. For example, the rules may assign weights to each criterion in the evidence data, including both uniquely identifying and substitutability criteria, whether required or optional. Using the weights and the matched criteria, the system may compute a level of confidence. In some embodiments, the system may further rely upon an object's name in determining the level of confidence.

**[0065]** In some embodiments, the rules may require that an object have a level of confidence above some threshold value in order for that object to be considered a match or a potential match. In some embodiments, the system may also report the level of confidence to a user when requesting that the user select a match. In some embodiments, where multiple potential matches are found, the system may automatically select the potential match having the highest level of confidence, assuming that match is above a threshold level. If the potential match with the highest level of confidence is not above the threshold level, the system may prompt the user to identify a match.

**[0066]** In some cases, the system may be unable to determine a match. In such cases, the system is said to have failed to resolve the evidence data for the object. The system may take any of a number of reactions to such a failure, including deleting the shared document, removing any elements of the shared document that rely on the unresolved object, using dummy data for the unresolved object, automatically requesting the unresolved object from the system sharing the item, warning the user of the error and requesting further user guidance, and/or attempting to render the document without correcting the failure.

**[0067]** The system may further or alternatively adhere to other resolution rules and other matching techniques.

#### 4.4. MISCELLANEOUS

**[0068]** In an embodiment, when a user sends an item to another user who may not have access to a data repository upon which the item relies, the sending or receiving system may communicate a warning message to the user advising the user that the recipient may not be able to correctly view the data in the item. The receiving system may further warn the recipient that the receiving system will render the item using data from the recipient's own data repository instead of the sender's data repository, and thus the data may not be the same as was intended by the sender.

**[0069]** According to an embodiment, items within a system store references to objects within the system. The system automatically replaces the references with evidence when the item is sent outside of the system. According to another embodiment, the system also replaces references with evidence when an item is sent to a user of the system who lacks access permissions or connectivity to a relied upon data repository. According to another embodiment, the system automatically generates and stores evidence for all items, regardless of whether those items are intended to remain within the system.

**[0070]** In some embodiments, evidence data may be used to verify referenced objects even when the referenced object has a name that matches the name of an object found in the receiving system's repository. For example, if the sending system sends an item that references an object named BETA, the receiving system may utilize evidence data to verify that its object named BETA is the same or substantively similar to the object named BETA at the sending system.

#### 5.0. IMPLEMENTATION MECHANISM—HARDWARE OVERVIEW

**[0071]** According to one embodiment, the techniques described herein are implemented by one or more special-purpose computing devices. The special-purpose computing devices may be hard-wired to perform the techniques, or may include digital electronic devices such as one or more application-specific integrated circuits (ASICs) or field programmable gate arrays (FPGAs) that are persistently programmed to perform the techniques, or may include one or more general purpose hardware processors programmed to perform the techniques pursuant to program instructions in firmware, memory, other storage, or a combination. Such special-purpose computing devices may also combine custom hard-wired logic, ASICs, or FPGAs with custom programming to accomplish the techniques. The special-purpose computing devices may be desktop computer systems, portable computer systems, handheld devices, networking devices or any other device that incorporates hard-wired and/or program logic to implement the techniques.



**[0072]** For example, FIG. 4 is a block diagram that illustrates a computer system 400 upon which an embodiment of the invention may be implemented. Computer system 400 includes a bus 402 or other communication mechanism for communicating information, and a hardware processor 404 coupled with bus 402 for processing information. Hardware processor 404 may be, for example, a general purpose microprocessor.

**[0073]** Computer system 400 also includes a main memory 406, such as a random access memory (RAM) or other dynamic storage device, coupled to bus 402 for storing information and instructions to be executed by processor 404. Main memory 406 also may be used for storing temporary variables or other intermediate information during execution of instructions to be executed by processor 404. Such instructions, when stored in storage media accessible to processor 404, render computer system 400 into a special-purpose machine that is customized to perform the operations specified in the instructions.

**[0074]** Computer system 400 further includes a read only memory (ROM) 408 or other static storage device coupled to bus 402 for storing static information and instructions for processor 404. A storage device 410, such as a magnetic disk or optical disk, is provided and coupled to bus 402 for storing information and instructions.

**[0075]** Computer system 400 may be coupled via bus 402 to a display 412, such as a cathode ray tube (CRT), for displaying information to a computer user. An input device 414, including alphanumeric and other keys, is coupled to bus 402 for communicating information and command selections to processor 404. Another type of user input device is cursor control 416, such as a mouse, a trackball, or cursor direction keys for communicating direction information and command selections to processor 404 and for controlling cursor movement on display 412. This input device typically has two degrees of freedom in two axes, a first axis (e.g., x) and a second axis (e.g., y), that allows the device to specify positions in a plane.

**[0076]** Computer system 400 may implement the techniques described herein using customized hard-wired logic, one or more ASICs or FPGAs, firmware and/or program logic which in combination with the computer system causes or programs computer system 400 to be a special-purpose machine. According to one embodiment, the techniques herein are performed by computer system 400 in response to processor 404 executing one or more sequences of one or more instructions contained in main memory 406. Such instructions may be read into main memory 406 from another storage medium, such as storage device 410. Execution of the sequences of instructions contained in main memory 406 causes processor 404 to perform the process steps described herein. In alternative embodiments, hard-wired circuitry may be used in place of or in combination with software instructions.

**[0077]** The term “storage media” as used herein refers to any media that store data and/or instructions that cause a machine to operation in a specific fashion. Such storage media may comprise non-volatile media and/or volatile media. Non-volatile media includes, for example, optical or magnetic disks, such as storage device 410. Volatile media includes dynamic memory, such as main memory 406. Common forms of storage media include, for example, a floppy disk, a flexible disk, hard disk, solid state drive, magnetic tape, or any other magnetic data storage medium, a CD-ROM, any other optical data storage medium, any physical medium with patterns of holes, a RAM, a PROM, and EPROM, a FLASH-EPROM, NVRAM, any other memory chip or cartridge.

**[0078]** Storage media is distinct from but may be used in conjunction with transmission media. Transmission media participates in transferring information between storage media. For example, transmission media includes coaxial cables, copper wire and fiber optics, including the wires that comprise bus 402. Transmission media can also take the form of acoustic or light waves, such as those generated during radio-wave and infra-red data communications.

**[0079]** Various forms of media may be involved in carrying one or more sequences of one or more instructions to processor 404 for execution. For example, the instructions may initially be carried on a magnetic disk or solid state drive of a remote computer. The remote computer can load the instructions into its dynamic memory and send the instructions over a telephone line using a modem. A modem local to computer system 400 can receive the data on the telephone line and use an infra-red transmitter to convert the data to an infra-red signal. An infra-red detector can receive the data carried in the infra-red signal and appropriate circuitry can place the data on bus 402. Bus 402 carries the data to main memory 406, from which processor 404 retrieves and executes the instructions. The instructions received by main memory 406 may optionally be stored on storage device 410 either before or after execution by processor 404.

**[0080]** Computer system 400 also includes a communication interface 418 coupled to bus 402. Communication interface 418 provides a two-way data communication coupling to a network link 420 that is connected to a local network 422. For example, communication interface 418 may be an integrated services digital network (ISDN) card, cable modem, satellite modem, or a modem to provide a data communication connection to a corresponding type of telephone line. As another example, communication interface 418 may be a local area network (LAN) card to provide a data communication connection to a compatible LAN. Wireless links may also be implemented. In any such implementation, communication interface 418 sends and receives electrical, electromagnetic or optical signals that carry digital data streams representing various types of information.

[0081] Network link 420 typically provides data communication through one or more networks to other data devices. For example, network link 420 may provide a connection through local network 422 to a host computer 424 or to data equipment operated by an Internet Service Provider (ISP) 426. ISP 426 in turn provides data communication services through the world wide packet data communication network now commonly referred to as the "Internet" 428. Local network 422 and Internet 428 both use electrical, electromagnetic or optical signals that carry digital data streams. The signals through the various networks and the signals on network link 420 and through communication interface 418, which carry the digital data to and from computer system 400, are example forms of transmission media.

[0082] Computer system 400 can send messages and receive data, including program code, through the network(s), network link 420 and communication interface 418. In the Internet example, a server 430 might transmit a requested code for an application program through Internet 428, ISP 426, local network 422 and communication interface 418.

[0083] The received code may be executed by processor 404 as it is received, and/or stored in storage device 410, or other non-volatile storage for later execution.

#### 6.0. EXTENSIONS AND ALTERNATIVES

[0084] In the foregoing specification, embodiments of the invention have been described with reference to numerous specific details that may vary from implementation to implementation. Thus, the sole and exclusive indicator of what is the invention, and is intended by the applicants to be the invention, is the set of claims that issue from this application, in the specific form in which such claims issue, including any subsequent correction. Any definitions expressly set forth herein for terms contained in such claims shall govern the meaning of such terms as used in the claims. Hence, no limitation, element, property, feature, advantage or attribute that is not expressly recited in a claim should limit the scope of such claim in any way. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.

## CLAIMS

What is claimed is:

1. A method for sharing information, the method comprising, at a first system:  
receiving an instruction to send a copy of a first item to a second system;  
wherein the first item comprises a reference to a first object;  
wherein the first object is accessible to the first system;  
wherein the first object is inaccessible to the second system;  
generating first data that describes the first object;  
wherein the first data is not identical to the reference;  
sending, to the second system, the copy of the first item and the first data;  
wherein the method is performed by one or more computing devices.
2. The method of Claim 1, wherein the copy of the first item includes the first data.
3. The method of Claim 1, further comprising generating the copy of the first item, wherein generating the copy of the first item comprises replacing the reference with the first data.
4. The method of Claim 1, further comprising, in response to receiving the instruction, determining that the first item comprises the reference, wherein generating the first data occurs in response to determining that the first item comprises the reference.
5. The method of Claim 1, wherein the first data uniquely identifies the first object.
6. The method of Claim 1, wherein:  
generating the first data comprises, based at least on the first object, determining one or more criteria for identifying another object that may substitute for the first object;  
and  
the first data comprises the one or more criteria.
7. The method of Claim 6, wherein determining one or more criteria for identifying another object that may substitute for the first object is further based upon a manner in which the first item relies upon data in the first object.
8. The method of Claim 1, wherein the first data includes data indicating at least one of the type of the first object or a plurality of characteristics of data associated with the first object.

9. The method of Claim 1, wherein the first object represents a financial instrument, and wherein the first data is data describing characteristics of data stored for the financial instrument at a data repository accessible to the first system.
10. The method of Claim 1, wherein generating the first data occurs in response to determining that the first object is inaccessible to a recipient at the second system.
11. The method of Claim 1, wherein the first system relies upon a first data repository in which the first object is stored, and wherein the second system does not have access to the first data repository.
12. A method for sharing information from a first system with a second system, the method comprising, at the second system:
- receiving (a) a first item that relies upon a first object accessible to the first system and (b) first data describing the first object;
  - wherein the first object has a first name;
  - wherein the first data is not identical to the first name;
- in response to determining that the first item relies upon the first object, using the first data to identify a second object accessible to the second system that may substitute for the first object; and
- generating and storing a copy of the first item, wherein the copy of the first item is identical to the first item except that that the copy of the first item relies upon the second object instead of the first object;
- wherein the method is performed by one or more computing devices.
13. The method of Claim 12, wherein:
- the first item references the first object using the first name;
  - generating the copy of the first item comprises replacing references to the first name with references to a second name;
  - the second name is a name by which the second system identifies the second object; and
  - the second name is different from the first name.
14. The method of Claim 12, wherein:
- the first item references the first object using the first data;
  - generating the copy of the first item comprises replacing the first data with a name by which the second system identifies the second object.

15. The method of Claim 12, wherein the first system relies upon a first data repository in which the first object is stored, wherein the first data repository is inaccessible to the second system.
16. The method of Claim 12, wherein the first data uniquely identifies the second object.
17. The method of Claim 12, wherein the first data identifies one or more mandatory criteria of an object that may substitute for the first object, wherein the second object satisfies the mandatory criteria.
18. The method of Claim 12, wherein:
  - the first data identifies one or more criteria of an object that may substitute for the first object;
  - the second system has access to a plurality of objects that may substitute for the first object; and
  - the second object is an object in the plurality of objects that best satisfies the one or more criteria.
19. The method of Claim 18, wherein identifying the second object comprises:
  - determining an object type for the first object; and
  - determining one or more resolution rules for the first object, wherein the one or more resolution rules identify weights for each of the one or more criteria;
  - based on the one or more criteria and the identified weights, determining which object in the plurality of objects best satisfies the one or more criteria.
20. The method of Claim 12, wherein the first data includes data indicating at least one of the type of the first object or a plurality of characteristics of data associated with the first object.
21. The method of Claim 12, wherein the first object and the second object both represent financial instruments, wherein the first data is data describing characteristics of data stored for the financial instrument at a data repository at the first system.
22. The method of Claim 12, wherein identifying the second object comprises:
  - determining a plurality of objects that may substitute for the first object;
  - prompting a user to identify which of the plurality of objects should be substituted for the first object.

23. One or more storage media storing instructions for sharing information, which instructions, when executed by one or more processors at a first system, cause:
- receiving an instruction to send a copy of a first item to a second system;
    - wherein the first item comprises a reference to a first object;
    - wherein the first object is accessible to the first system;
    - wherein the first object is inaccessible to the second system;
  - generating first data that describes the first object;
    - wherein the first data is not identical to the reference;
  - sending, to the second system, the copy of the first item and the first data.
24. The one or more storage media of Claim 22, wherein the copy of the first item includes the first data.
25. The one or more storage media of Claim 22, wherein the instructions, when executed by the one or more processors, further cause generating the copy of the first item, wherein generating the copy of the first item comprises replacing the reference with the first data.
26. The one or more storage media of Claim 22, wherein the instructions, when executed by the one or more processors, further cause, in response to receiving the instruction, determining that the first item comprises the reference, wherein generating the first data occurs in response to determining that the first item comprises the reference.
27. The one or more storage media of Claim 22, wherein the first data uniquely identifies the first object.
28. The one or more storage media of Claim 22, wherein:
- generating the first data comprises, based at least on the first object, determining one or more criteria for identifying another object that may substitute for the first object;
  - and
  - the first data comprises the one or more criteria.
29. The one or more storage media of Claim 27, wherein determining one or more criteria for identifying another object that may substitute for the first object is further based upon a manner in which the first item relies upon data in the first object.

30. The one or more storage media of Claim 22, wherein the first data includes data indicating at least one of the type of the first object or a plurality of characteristics of data associated with the first object.
31. The one or more storage media of Claim 22, wherein the first object represents a financial instrument, and wherein the first data is data describing characteristics of data stored for the financial instrument at a data repository accessible to the first system.
32. The one or more storage media of Claim 22, wherein generating the first data occurs in response to determining that the first object is inaccessible to a recipient at the second system.
33. The one or more storage media of Claim 22, wherein the first system relies upon a first data repository in which the first object is stored, and wherein the second system does not have access to the first data repository.
34. One or more storage media storing instructions for sharing information from a first system with a second system, which instructions, when executed by one or more processors at the second system, cause:
- receiving (a) a first item that relies upon a first object accessible to the first system and
  - (b) first data describing the first object;
  - wherein the first object has a first name;
  - wherein the first data is not identical to the first name;
- in response to determining that the first item relies upon the first object, using the first data to identify a second object accessible to the second system that may substitute for the first object; and
- generating and storing a copy of the first item, wherein the copy of the first item is identical to the first item except that that the copy of the first item relies upon the second object instead of the first object.
35. The one or more storage media of Claim 33, wherein:
- the first item references the first object using the first name;
  - generating the copy of the first item comprises replacing references to the first name with references to a second name;
  - the second name is a name by which the second system identifies the second object; and
  - the second name is different from the first name.
36. The one or more storage media of Claim 33, wherein:



the first item references the first object using the first data;  
generating the copy of the first item comprises replacing the first data with a name by  
which the second system identifies the second object.

37. The one or more storage media of Claim 33, wherein the first system relies upon a first data repository in which the first object is stored, wherein the first data repository is inaccessible to the second system.

38. The one or more storage media of Claim 33, wherein the first data uniquely identifies the second object.

39. The one or more storage media of Claim 33, wherein the first data identifies one or more mandatory criteria of an object that may substitute for the first object, wherein the second object satisfies the mandatory criteria.

40. The one or more storage media of Claim 33, wherein:  
the first data identifies one or more criteria of an object that may substitute for the first object;  
the second system has access to a plurality of objects that may substitute for the first object; and  
the second object is an object in the plurality of objects that best satisfies the one or more criteria.

41. The one or more storage media of Claim 39, wherein identifying the second object comprises:  
determining an object type for the first object; and  
determining one or more resolution rules for the first object, wherein the one or more resolution rules identify weights for each of the one or more criteria;  
based on the one or more criteria and the identified weights, determining which object in the plurality of objects best satisfies the one or more criteria.

42. The one or more storage media of Claim 33, wherein the first data includes data indicating at least one of the type of the first object or a plurality of characteristics of data associated with the first object.

43. The one or more storage media of Claim 33, wherein the first object and the second object both represent financial instruments, wherein the first data is data describing characteristics of data stored for the financial instrument at a data repository at the first system.

44. The one or more storage media of Claim 33, wherein identifying the second object comprises:

determining a plurality of objects that may substitute for the first object;

prompting a user to identify which of the plurality of objects should be substituted for the first object.

100

FIG. 1

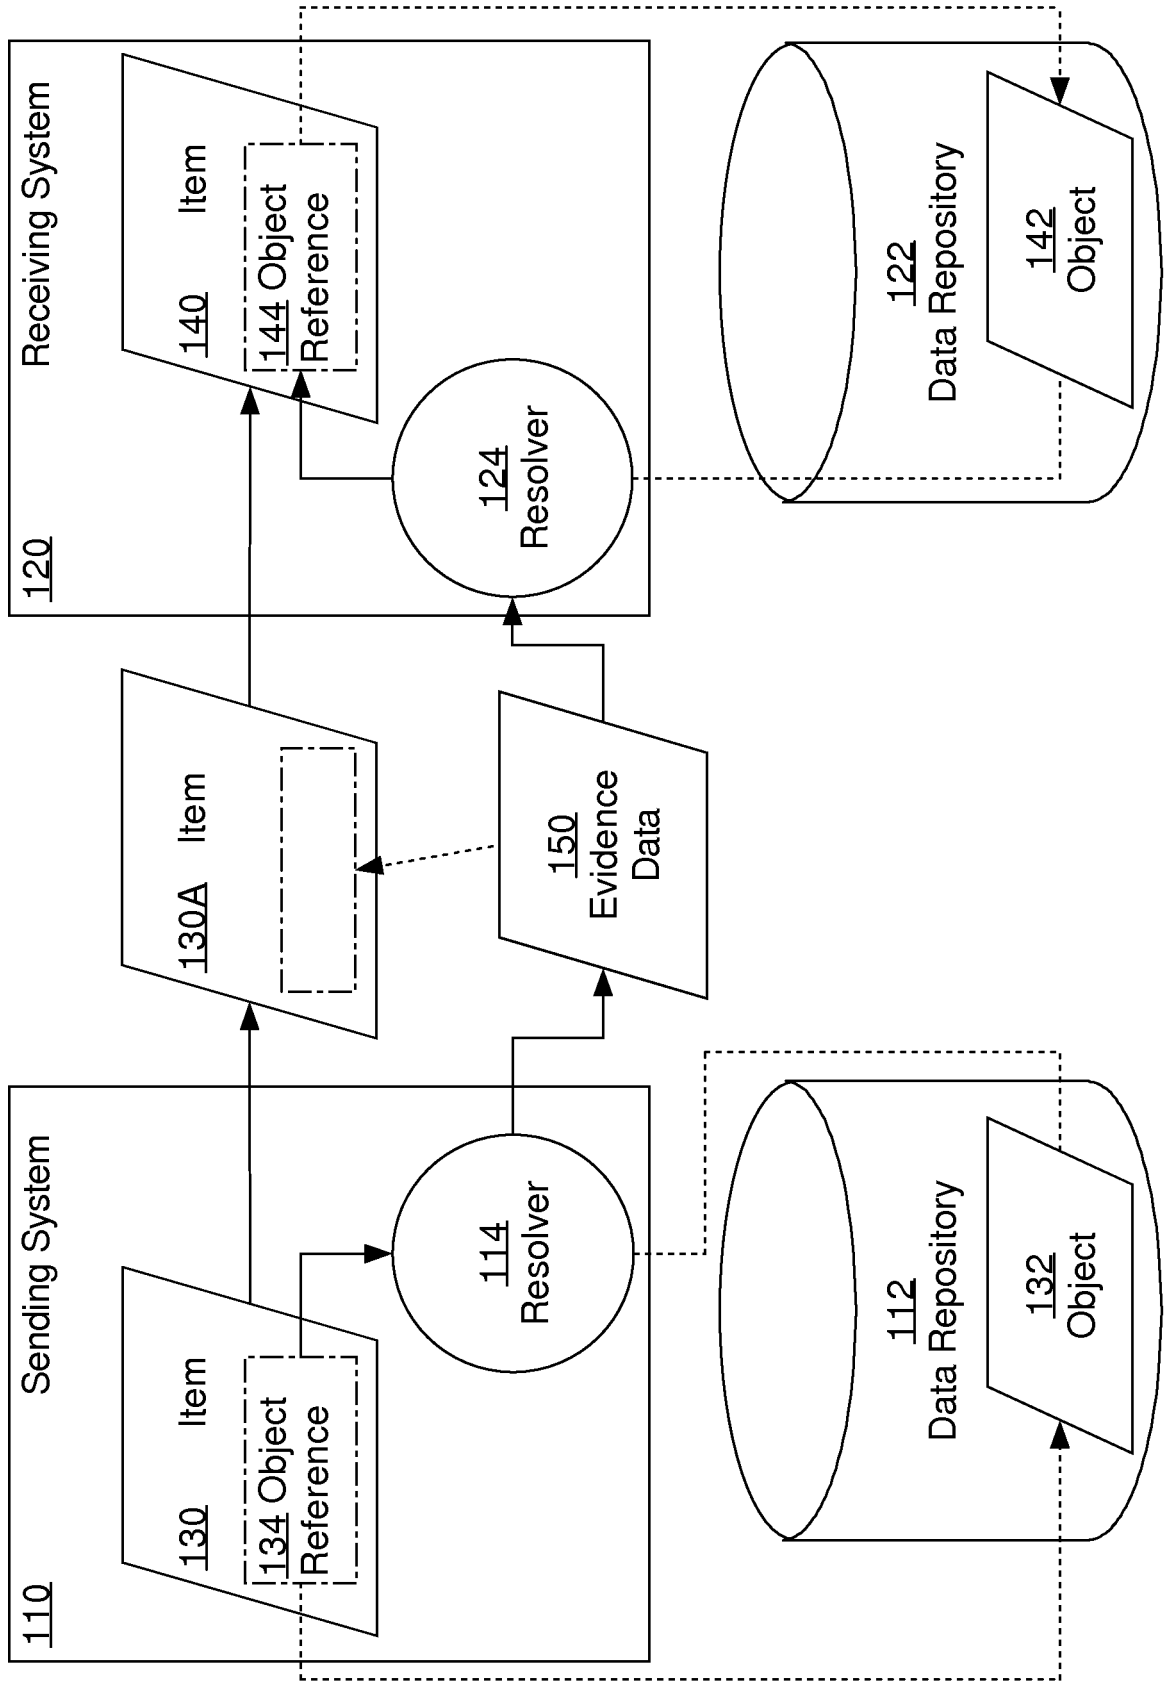


FIG. 2

200

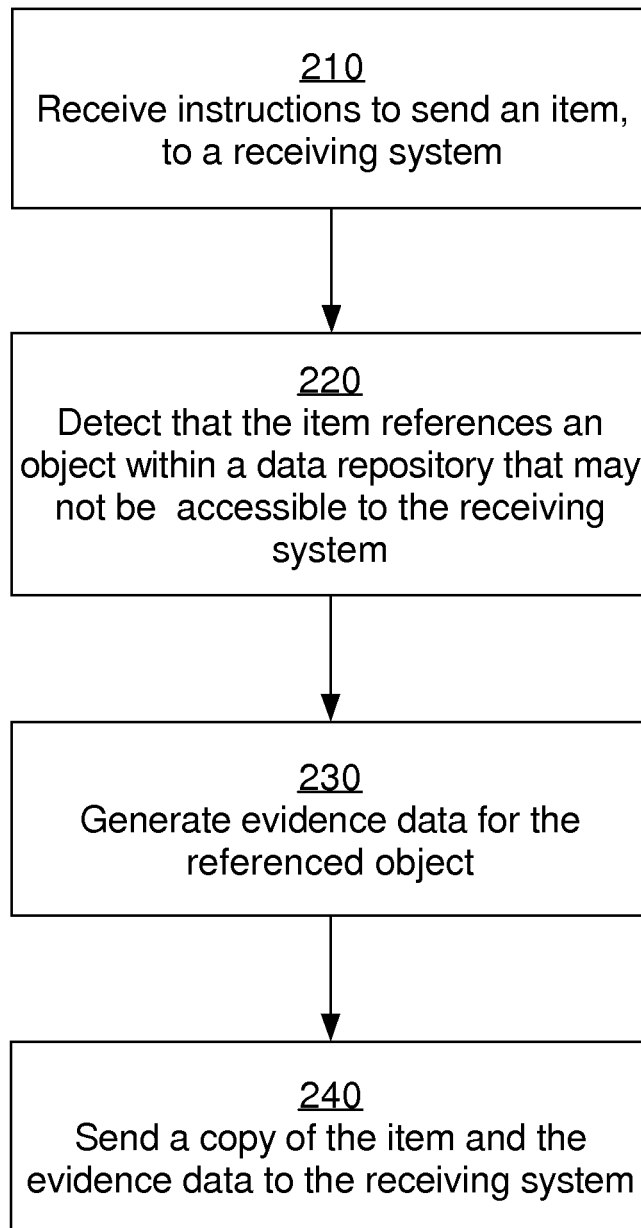


FIG. 3

300

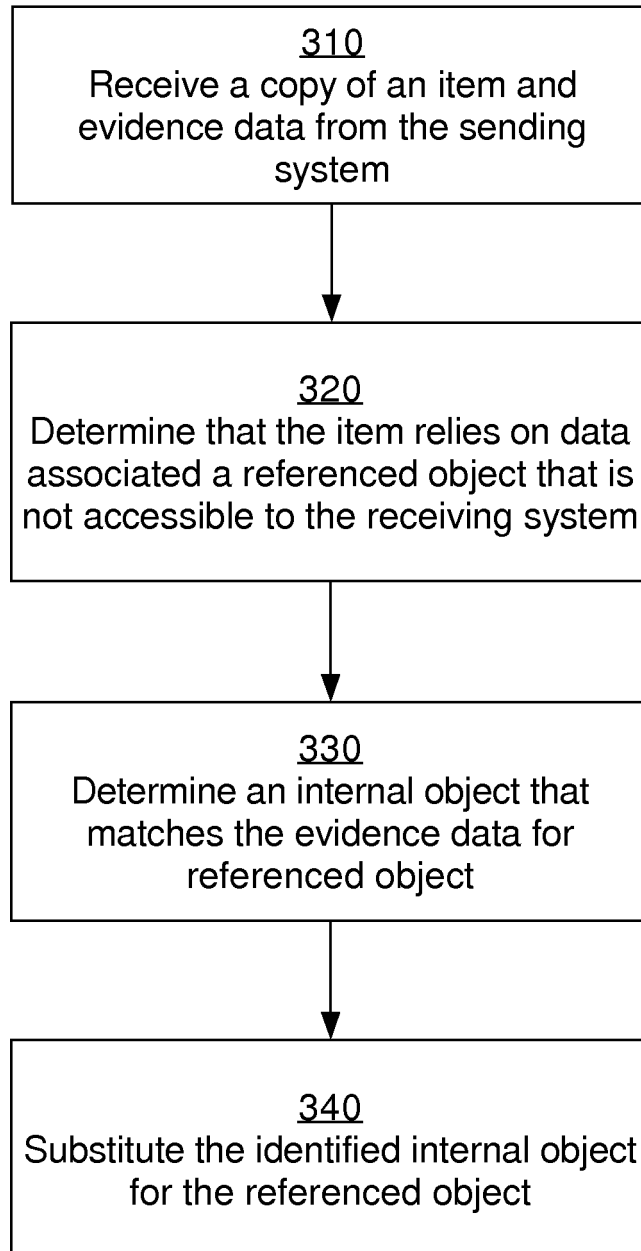


FIG. 4

