

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
27 November 2008 (27.11.2008)

PCT

(10) International Publication Number
WO 2008/141852 A1

- (51) International Patent Classification:
G06F 9/46 (2006.01)
- (21) International Application Number:
PCT/EP2008/052937
- (22) International Filing Date: 12 March 2008 (12.03.2008)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
11/751,571 21 May 2007 (21.05.2007) US
- (71) Applicant (for all designated States except US): **INTERNATIONAL BUSINESS MACHINES CORPORATION** [US/US]; New Orchard Road, Armonk, NY 10504 (US).
- (71) Applicant (for MG only): **IBM UNITED KINGDOM LIMITED** [GB/GB]; PO Box 41, North Harbour, Portsmouth, Hampshire PO6 3AU (GB).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): **DESAI, Saurabh** [US/US]; 15808 Pumpkin Ridge Drive, Austin, TX 78717 (US). **DUBEY, Niteesh, Kumar** [IN/US]; 2 Rochambaeu Drive, Apt. A, Yorktown Heights, NY 10598 (US). **JANN, Joefon** [US/US]; 213 Barnes Street, Ossining, NY 10562

(US). **PATTNAIK, Pratap** [US/US]; 213 Barnes Street, Ossining, NY 10562 (US). **SHANKAR, Ravi** [IN/US]; 2600 Gracy Farms Lane, #218, Austin, TX 78758 (US). **VADDAGIRI, Murali** [IN/IN]; 435 Second Cross, OMBR Layout, Bhuvanagiri, Bangalore 560 033 (IN).

(74) Agent: **SEKAR, Anita**; IBM United Kingdom Limited, Intellectual Property Law, Hursley Park, Winchester, Hampshire SO21 2JN (GB).

(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, 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, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, 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,

[Continued on next page]

(54) Title: A FRAMEWORK FOR MANAGING ATTRIBUTES OF OBJECTS

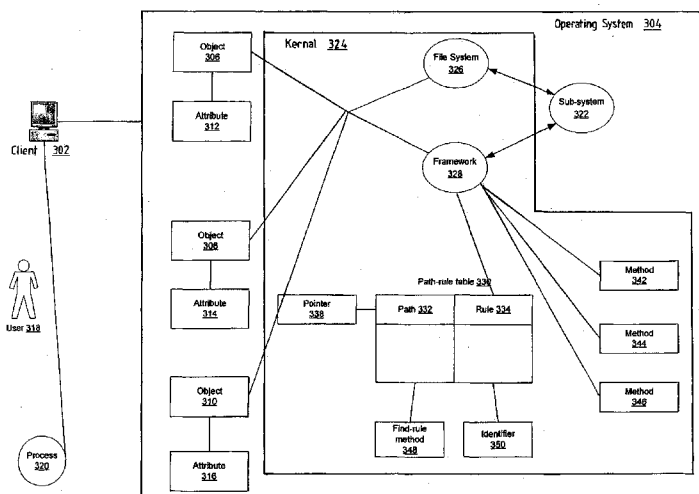


Figure 3

(57) Abstract: A computer implemented method, computer program product, and system for managing objects. Responsive to receiving a find-rule method, and a path-rule table, wherein the path-rule table contains a set of paths, wherein each path references an object, wherein a file system locates the object using the path, and wherein the object has at least one attribute not known to the file system, a path-rule table identifier is created. The path-rule table is associated with the path-rule table identifier to form an associated path-rule table. The find-rule method is associated with the path-rule table identifier to form an associated find-rule method. The path-rule table identifier, the associated path-rule table, and the associated find-rule method are stored. The path-rule table identifier is returned.

WO 2008/141852 A1



FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL,
NO, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG,
CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:
— *with international search report*

A FRAMEWORK FOR MANAGING ATTRIBUTES OF OBJECTS

FIELD OF THE INVENTION

5

The present invention relates generally to data processing systems and in particular to a method, computer program product, and apparatus for managing objects. Still more particularly, the present invention provides a computer implemented method, apparatus, and computer program product for a framework for managing objects.

10

BACKGROUND OF THE INVENTION

An object in a computer is an item which can be accessed, such as, for example, a file, a device, or a command. In a computer, the computer's operating system keeps track of objects. The operating system may control each user's access to specific objects in the computer. For example, a format command in a kernel may be used to format a storage device, erasing all data stored on the storage device. To prevent the loss of important objects, the kernel may prevent specific users from accessing the important objects. Typically, an operating system allows a system administrator to specify an access rule for an object. The access rule specifies the level of access each user has to the object. An access rule is specified using an attribute of the object known to the operating system. Thus, the kernel must know the object's attribute to control the level of access. If the object has an attribute and the kernel is not aware of the attribute, the attribute cannot be used to specify the access rule for the object.

25

DISCLOSURE OF THE INVENTION

The illustrative embodiments described herein provide a computer implemented method, computer program product, and apparatus for managing attributes for objects. A path-rule table containing a set of entries is stored in a location accessible by an operating system. The set of entries in the path-rule table specifies associations between objects and attributes of the objects. The path-rule table is associated with an identifier. When a request

30

to access an object is received from an entity, the type of the object is identified. The identifier is used to identify the path-rule table for the type of object. Access to the object is then selectively provided based on an attribute associated with the object in the path-rule table.

5

BRIEF DESCRIPTION OF THE DRAWINGS

The novel features believed characteristic of the invention are set forth in the appended claims. The invention itself, however, as well as a preferred mode of use, further objectives and advantages thereof, will best be understood by reference to the following detailed description of an illustrative embodiment when read in conjunction with the accompanying drawings, wherein:

10

Figure 1 depicts a pictorial representation of a network of data processing systems in accordance with an illustrative embodiment;

15

Figure 2 is a block diagram of a data processing system in which illustrative embodiments may be implemented;

20

Figure 3 is a block diagram of using a framework in accordance with an illustrative embodiment;

Figure 4 is a block diagram of an exemplary path-rule table in accordance with an illustrative embodiment;

25

Figure 5 is a flowchart for creating a path-rule table in accordance with an illustrative embodiment;

Figure 6 is a flowchart of finding a rule in a path-rule table in accordance with an illustrative embodiment; and

30

Figure 7 is a flowchart of an exemplary find-rule method in accordance with an illustrative embodiment.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

5

With reference now to the figures and in particular with reference to **Figures 1-2**, exemplary diagrams of data processing environments are provided in which illustrative embodiments may be implemented. It should be appreciated that **Figures 1-2** are only exemplary and are not intended to assert or imply any limitation with regard to the environments in which different embodiments may be implemented. Many modifications to the depicted environments may be made.

10

Figure 1 depicts a pictorial representation of a network of data processing systems in which illustrative embodiments may be implemented. Network data processing system **100** is a network of computers in which the illustrative embodiments may be implemented. Network data processing system **100** contains network **102**, which is the medium used to provide communications links between various devices and computers connected together within network data processing system **100**. Network **102** may include connections, such as wire, wireless communication links, or fiber optic cables.

15

20

In the depicted example, server **104** and server **106** connect to network **102** along with storage unit **108**. In addition, clients **110**, **112**, and **114** connect to network **102**. Clients **110**, **112**, and **114** may be, for example, personal computers or network computers. In the depicted example, server **104** provides data, such as boot files, operating system images, and applications to clients **110**, **112**, and **114**. Clients **110**, **112**, and **114** are clients to server **104** in this example. Servers **104** and **106** provide processes for managing object attributes, such as file, device, or command attributes. These processes may manage local objects, or objects distributed across network **102**. Network data processing system **100** may include additional servers, clients, and other devices not shown.

25

30

In the depicted example, network data processing system **100** is the Internet with network **102** representing a worldwide collection of networks and gateways that use the

Transmission Control Protocol/Internet Protocol (TCP/IP) suite of protocols to communicate with one another. At the heart of the Internet is a backbone of high-speed data communication lines between major nodes or host computers, consisting of thousands of commercial, governmental, educational and other computer systems that route data and messages. Of course, network data processing system **100** also may be implemented as a number of different types of networks, such as for example, an intranet, a local area network (LAN), or a wide area network (WAN). **Figure 1** is intended as an example, and not as an architectural limitation for the different illustrative embodiments.

With reference now to **Figure 2**, a block diagram of a data processing system is shown in which illustrative embodiments may be implemented. Data processing system **200** is an example of a computer, such as server **104** or client **110** in **Figure 1**, in which computer usable program code or instructions implementing the processes may be located for the illustrative embodiments.

In the depicted example, data processing system **200** employs a hub architecture including a north bridge and memory controller hub (NB/MCH) **202** and a south bridge and input/output (I/O) controller hub (SB/ICH) **204**. Processing unit **206**, main memory **208**, and graphics processor **210** are coupled to north bridge and memory controller hub **202**. Processing unit **206** may contain one or more processors and even may be implemented using one or more heterogeneous processor systems. Graphics processor **210** may be coupled to the NB/MCH through an accelerated graphics port (AGP), for example.

In the depicted example, local area network (LAN) adapter **212** is coupled to south bridge and I/O controller hub **204** and audio adapter **216**, keyboard and mouse adapter **220**, modem **222**, read only memory (ROM) **224**, universal serial bus (USB) and other ports **232**, and PCI/PCIe devices **234** are coupled to south bridge and I/O controller hub **204** through bus **238**, and hard disk drive (HDD) **226** and CD-ROM **230** are coupled to south bridge and I/O controller hub **204** through bus **240**. PCI/PCIe devices may include, for example, Ethernet adapters, add-in cards, and PC cards for notebook computers. PCI uses a card bus controller, while PCIe does not. ROM **224** may be, for example, a flash binary input/output system (BIOS). Hard disk drive **226** and CD-ROM **230** may use, for example, an integrated drive

electronics (IDE) or serial advanced technology attachment (SATA) interface. A super I/O (SIO) device **236** may be coupled to south bridge and I/O controller hub **204**.

An operating system runs on processing unit **206** and coordinates and provides control of various components within data processing system **200** in **Figure 2**. The operating system may be a commercially available operating system such as Microsoft® Windows® XP (Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both). An object oriented programming system, such as the Java™ programming system, may run in conjunction with the operating system and provides calls to the operating system from Java™ programs or applications executing on data processing system **200**. Java™ and all Java™-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Instructions for the operating system, the object-oriented programming system, and applications or programs are located on storage devices, such as hard disk drive **226**, and may be loaded into main memory **208** for execution by processing unit **206**. The processes of the illustrative embodiments may be performed by processing unit **206** using computer implemented instructions, which may be located in a memory such as, for example, main memory **208**, read only memory **224**, or in one or more peripheral devices.

The hardware in **Figures 1-2** may vary depending on the implementation. Other internal hardware or peripheral devices, such as flash memory, equivalent non-volatile memory, or optical disk drives and the like, may be used in addition to or in place of the hardware depicted in **Figures 1-2**. Also, the processes of the illustrative embodiments may be applied to a multiprocessor data processing system.

In some illustrative examples, data processing system **200** may be a personal digital assistant (PDA), which is generally configured with flash memory to provide non-volatile memory for storing operating system files and/or user-generated data. A bus system may be comprised of one or more buses, such as a system bus, an I/O bus and a PCI bus. Of course the bus system may be implemented using any type of communications fabric or architecture that provides for a transfer of data between different components or devices attached to the

fabric or architecture. A communications unit may include one or more devices used to transmit and receive data, such as a modem or a network adapter. A memory may be, for example, main memory **208** or a cache such as found in north bridge and memory controller hub **202**. A processing unit may include one or more processors or CPUs. The depicted examples in **Figures 1-2** and above-described examples are not meant to imply architectural limitations. For example, data processing system **200** also may be a tablet computer, laptop computer, or telephone device in addition to taking the form of a PDA.

In a computer operating system, the lowest-level portion of the operating system is called a kernel. One part of the kernel is called the file system. The file system is used to locate objects in the computer. An object is any item referenced by a path, such as a file, a command, or a device.

An operating system on a computer may have one or more sub-systems. Each sub-system is a set of related software components for performing a specific purpose. For example, a computer may have a security sub-system which restricts access to important files, commands, and devices. A sub-system uses the kernel to access objects using the path of the object. The kernel associates a set of attributes with each object, such as, for example, the date and time the object was last modified.

A sub-system may associate one or more attributes with a set of objects, and the file system may not be aware of these attributes. The sub-system may have a need to manage the set of objects in the file system using at least one attribute of the set of objects which is not known to the file system. The illustrative embodiments recognize a need for managing a set of objects in the file system independent of the attributes known to the file system.

Figure 3 is a block diagram of using a framework in accordance with an illustrative embodiment. In using a framework for managing objects **300**, client **302** is a computer, such as client **110** in **Figure 1**. An operating system, operating system **304**, runs on client **302**. Operating system **304** contains objects **306**, **308**, and **310**. Operating system **304** may be a UNIX®-based operating system, such as, for example, Advanced Interactive eXecutive (AIX®) from International Business Machines (IBM®). Objects **306-310** are each objects

which may be in a computer, such as, for example, commands, files, and devices. Each object has an attribute. Object **306** has attribute **312**, object **308** has attribute **314**, and object **310** has attribute **316**. User **318**, and process **320** use operating system **304** to access objects **306-310** and perform various actions.

5

Operating system **304** contains various functional components, including sub-system **322**, and kernel **324**. Sub-system **318** is a component of operating system **304**. For example, sub-system **318** may be used to implement an access control mechanism for objects **306-310** to control whether user **318**, and process **320** can access objects **306-310**. Kernel **324** has several components, including file system **326**, and framework **328**. User **318**, process **320**, and sub-system **302** may use file system **326** to access objects **306-310**.

10

For example, assume sub-system **322** uses file system **326** to access objects **306-310**. If file system **326** is aware of attributes **312-316** of objects **306-310**, then sub-system **322** may manage objects **306-310** using attributes **312-316**. For example, sub-system **322** may set attribute **312** to indicate that file system **326** should not allow user **318** to access object **306**. However, if sub-system **322** wishes to associate attributes **312-316** with objects **306-310**, and file system **326** is not aware of attributes **312-316**, then sub-system **322** needs a means for managing the attributes **312-316** of objects **306-310**. In this situation, kernel **324** provides framework **328** for managing attributes **312-316** of objects **306-310**, independent of file system **326**. Framework **328** may be used by more than one file system in operating system **304**.

15

20

Framework **328** allows sub-system **322** to manage attributes **312-316** of objects **306-310**, regardless of whether file system **326** is aware of attributes **312-316**. Sub-system may use framework **328** to manage attributes **312-316** of objects **306-310** even when file system **326** is aware of attributes **312-316**.

25

To use framework **328**, sub-system **322** first asks framework **328** to create a path-rule table, in this example, path-rule table **330**. Path-rule table **330** contains a set of entries, where a set is one or more entries. Each entry in path-rule table **330** contains at least two fields, path **332**, and rule **334**. Path **332** is an absolute path to an object. Rule **334** is

30

meaningful to sub-system **322**, and is used by sub-system **322** to decide whether to perform an action. For example, sub-system **322** may perform one or more actions such as setting a privilege level, setting an access level, and creating a checksum based on rule **334**.

5 To create a path-rule table, sub-system **322** sends request **336** to framework **328**, requesting framework **328** to create path-rule table **330**. In this example, request 336 is a create request which asks framework **328** to create a path-rule table. In request **336**, sub-system **322** specifies various information, such as the size and the contents of the path-rule table. For example, sub-system **322** may specify the number of entries to be created in the
10 path-rule table, as well as a set of paths, and a set of rules in request **336**. A set of paths is one or more paths, and a set of rules is one or more rules. Each entry in the path-rule table has one path, and one rule. When an object referenced by an absolute path in the path-rule table is accessed, the sub-system may use the rule associated with the absolute path to perform an action.

15 Framework **328** receives request **336** and creates path-rule table **330** using the information contained in request **336**. For example, framework 328 creates a set of entries in path-rule table **330**, and using the information contained in request **336**, adds each path in the set of paths, and each rule in the set of rules, to each entry in path-rule table **330**.

20 Sub-system **322** may optionally specify additional information in request **336**. For example, sub-system **322** may specify pointer **338** in request **336**. Pointer **338** is a pointer to a root directory. The root directory referenced by pointer **338** is used as a reference point for each absolute path in path-rule table **330**. Sub-system **322** may also specify the maximum
25 size of each absolute path in path **332** in the information in request **336**.

 Framework **328** also provides a set of one or more find-rule methods. In this example, framework 328 provides three find-rule methods, methods **342**, **344**, and **346**. For a given path, the find-rule method finds the absolute path in the path-rule table which refers
30 to the same object as the given path refers to. In request **336**, sub-system **322** also specifies the find-rule method which the sub-system wishes to use with the path-rule table. For

example, sub-system **322** may specify in request **336** that framework **328** use method **342** with the path-rule table.

Thus, when framework **328** receives request **336** from sub-system **322**, framework **328** creates path-rule table **330**, and stores the set of paths, and the set of rules in path-rule table **330**. Each entry in path-rule table **330** contains one path, path **332**, and one rule, rule **334**. Framework **328** associates a find-rule method, find-rule method **348**, with path-rule table **330**. Find-rule method **348** identifies the specific method from table **340** chosen by sub-system **322** for use with path-rule table **330**. Framework **328** then returns identifier **350** to sub-system **322**.

Sub-system **322** may use identifier **350** to request framework **328** to perform a variety of operations. For example, sub-system **322** may use identifier **350** to request framework **328** to delete path-rule table **330**, add one or more entries in path-rule table **330**, and change find-rule method **348** to a different method.

Each find-rule method varies in the technique the find-rule method uses to match an object referenced by a given path with the object referenced by a path in path-rule table **330**. For example, when using a path-rule table to manage a level of security for an object, different find-rule methods may be created based on the level of security, and the level of performance desired. Typically, when finding a set of rules for a given path in the path-rule table using a given find-rule method, the greater the level of security, the longer it takes to find the set of rules for the given path using the find-rule method. In this example, the find-rule method balances the need for fast performance with the need for greater security. One find-rule method may provide lower security with fast performance, while another find-rule method may provide may provide higher security with faster performance.

When user **318**, or process **320** accesses an object, such as object **306**, the object is accessed using a path. When the object is accessed using the path, sub-system **322** determines whether the object has an associated rule. Sub-system **322** sends a find request to framework **328**, asking framework **328** to determine if object **306** has an associated rule in path-rule table **330**. Sub-system **322** sends the path for the object accessed, and identifier

350 to framework **328**. For clarity, the path given by sub-system **322** to framework **328** is called the given path.

Upon receiving the find request, framework **328** uses identifier **350** to find the path-rule table, and find-rule method associated with identifier **350**. In this example, framework **328** uses identifier **350** to find path-rule table **330**, and find-rule method **348**. Framework **328** uses find-rule method **348** to determine if the object referenced by the given path is also referenced by a path in the path-rule table. If find-rule method **348** finds the path, such as path **332**, in the path-rule table which refers to the object in the given path, find-rule method returns rule **334** associated with path **332** in the path-rule table.

Thus, rule **334** is returned to sub-system **322** in response to the find request of sub-system **322** to determine whether path of the object accessed has a rule in path-rule-table **330**. Sub-system **322** may then perform an action using rule **334**. For example, if sub-system **322** is a security sub-system, sub-system **322** may use rule **334** to determine whether user **318**, and process **320** have the privileges necessary to access the object.

Figure 4 is a block diagram of an exemplary path-rule table in accordance with an illustrative embodiment. Exemplary path-rule table **400** is an example of a path-rule table, such as path-rule table **326** in **Figure 3**, in a Unix-based operating system, such as Advanced Interactive eXecutive (AIX®). Exemplary path-rule table **400** contains three entries. Each entry has an absolute path, and an access rule. In this example, each path in paths **402**, **404**, and **406** is an absolute path for commands. In another embodiment, each path in the path-rule table may contain an absolute path for a device, and a file. Rule **408** is a rule for determining the access level for path **402**. Similarly, rule **410** is a rule for determining the access level for path **404**. Additionally, rule **412** is a rule for determining the access level for path **406**.

Figure 5 is a flowchart for creating a path-rule table in accordance with an illustrative embodiment. The process in **Figure 5** is executed by a framework, such as framework **328** in **Figure 3**. The process begins when a create request to create a path rule table is received (step **502**). For example, the request to create the path-rule table may be sent by a sub-system, such as sub-system **322** in **Figure 3**. The request contains the number

of entries for the path-rule table, and a set of paths, and a set of rules. Optionally, the request may also contain a pointer to a root directory.

A path-rule table is created with the requested number of entries, and the set of paths, and the set of rules are stored in the path-rule table, resulting in each entry in the path-rule table having a path, and a rule (step **504**). An identifier for the path-rule table is created, and the identifier is associated with the selected find-rule method, and the path-rule table (step **506**). In step **506**, optionally, the pointer to the root directory is stored with the path-rule table. The framework returns the identifier (step **508**). Typically, a sub-system, such as sub-system **322** in **Figure 3** receives the identifier to the path-rule table created by the framework.

Figure 6 is a flowchart of finding a rule in a path-rule table in accordance with an illustrative embodiment. The process in **Figure 6** is executed by a framework, such as framework **328** in **Figure 3**. The process begins when a find request is received containing a path for an object, and an identifier (step **602**). The identifier is an identifier for a path-rule table, such as identifier **350** in **Figure 3**. The identifier is used to find the path-rule table, and the find-rule method associated with the identifier (step **604**). The find-rule method is used to determine if there is a path in the find rule table that references the same object the given path references (step **606**). A determination is made as to whether the find-rule method found a path in the path-rule table (step **608**). If the answer in step **608** is “yes”, because the find-rule method did not find a path, then the rule corresponding to the path in the path-rule table is returned (step **610**), and the process ends. If the answer is “no”, and the find-rule method did not find a path, then a “not found in table” return code is returned (step **612**), and the process ends.

Figure 7 is a flowchart of an exemplary find-rule method in accordance with an illustrative embodiment. The process in **Figure 7** is an example of a find-rule method. The flowchart in **Figure 7** is executed by a find-rule method in a framework, such as method **342** in **Figure 3**. The process begins when a find request to find a rule for a path, a path, and an identifier are received (step **702**). The given path may be a relative, or an absolute path. An

absolute path is retrieved from an entry in the path-rule table (step **704**). The path-rule table may be a path-rule table, such as path-rule table **330** in **Figure 3**.

5 A determination is made whether the absolute path from the entry in path-rule table has the same basename as the resolved basename (step **706**). The basename is the name of the object referenced by the path. For example, in path “/usr/doc/text”, “text” is the
10 basename of the path. If the answer in step **706** is “yes”, because the absolute path from the entry in path-rule table has the same basename as the resolved basename, then a determination is made whether the identity of the physical file system vnode of the given
15 path is identical to the identity of the physical file system vnode of the absolute path (step **708**). In a Unix-based operating system, each active file, each current directory, each mounted-on file, text file, and the root has a unique vnode. Thus, a vnode uniquely identifies each object.

20 If the answer in step **708** is “yes”, then the rule corresponding to the path in the path-rule table is retrieved, and returned (step **710**). For example, framework **322** retrieves the rule corresponding to the given path in the path-rule table, and returns the rule to sub-system **322**. If the answer in step **708** is “no”, then a determination is made whether all absolute paths in the path-rule table have been retrieved. If the answer in step **714** is “yes”, because
25 all absolute paths have been retrieved from the path-rule table, then a “path not found” return code is returned (step **714**), and the process ends. If the answer in step **714** is “no”, because all absolute paths in the path-rule table have not been retrieved, then the process returns to step **704** and retrieves an absolute path from an entry in the path-rule table.

30 Thus, the illustrative embodiments described herein provide a computer implemented method, apparatus, and computer program product for accessing an object in a set of objects. Responsive to receiving a path-rule table containing a set of entries, a table identifier is created, wherein the table identifier is used to identify the path-rule table. A rule finder is received, wherein the rule finder is used to identify the rule in the path-rule table for accessing the object.

The flowchart and block diagrams in the figures illustrate the architecture, functionality, and operation of some possible implementations of systems, methods and computer program products according to various embodiments. In this regard, each block in the flowchart or block diagrams may represent a module, segment or portion of code, which comprises one or more executable instructions for implementing the specified logical function(s). It should also be noted that, in some alternative implementations, the functions noted in the block may occur out of the order noted in the figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved.

The invention can take the form of an entirely hardware embodiment, an entirely software embodiment or an embodiment containing both hardware and software elements. In a preferred embodiment, the invention is implemented in software, which includes but is not limited to firmware, resident software, microcode, etc.

Furthermore, the invention can take the form of a computer program product accessible from a computer-usable or computer-readable medium providing program code for use by or in connection with a computer or any instruction execution system. For the purposes of this description, a computer-usable or computer readable medium can be any tangible apparatus that can contain, store, communicate, propagate, or transport the program for use by or in connection with the instruction execution system, apparatus, or device.

The medium can be an electronic, magnetic, optical, electromagnetic, infrared, or semiconductor system (or apparatus or device) or a propagation medium. Examples of a computer-readable medium include a semiconductor or solid state memory, magnetic tape, a removable computer diskette, a random access memory (RAM), a read-only memory (ROM), a rigid magnetic disk and an optical disk. Current examples of optical disks include compact disk – read only memory (CD-ROM), compact disk – read/write (CD-R/W) and DVD.

A data processing system suitable for storing and/or executing program code will include at least one processor coupled directly or indirectly to memory elements through a

system bus. The memory elements can include local memory employed during actual execution of the program code, bulk storage, and cache memories which provide temporary storage of at least some program code in order to reduce the number of times code must be retrieved from bulk storage during execution.

5

Input/output or I/O devices (including but not limited to keyboards, displays, pointing devices, etc.) can be coupled to the system either directly or through intervening I/O controllers.

10

Network adapters may also be coupled to the system to enable the data processing system to become coupled to other data processing systems or remote printers or storage devices through intervening private or public networks. Modems, cable modem and Ethernet cards are just a few of the currently available types of network adapters.

15

The description of the present invention has been presented for purposes of illustration and description, and is not intended to be exhaustive or limited to the invention in the form disclosed. Many modifications and variations will be apparent to those of ordinary skill in the art. The embodiment was chosen and described in order to best explain the principles of the invention, the practical application, and to enable others of ordinary skill in the art to understand the invention for various embodiments with various modifications as are suited to the particular use contemplated.

20

CLAIMS

1. A method for managing attributes for objects, the method comprising:
storing a path-rule table comprising a set of entries associated with a location
5 accessible by an operating system, wherein the set of entries in the path-rule table specifies
associations between objects and attributes of the objects;
associating the path-rule table with an identifier;
responsive to a request from an entity to access an object, identifying a type of the
object;
10 using the identifier to identify the path-rule table for the type of object; and
selectively providing access to the object in accordance with an attribute associated
with the object in the path-rule table.
2. The method of claim 1, wherein storing a path-rule table further comprises:
15 creating the identifier for the path-rule table; and
receiving a rule finder, wherein the rule finder is used to identify the attribute in the
path-rule table for accessing the object.
3. The method of claim 1, wherein selectively providing access to the object in
20 accordance with an attribute associated with the object in the path-rule table further
comprises:
determining a level of access to the object using the path-rule table in accordance
with the entity; and
providing the entity the level of access to the object.
25
4. The method of claim 3, wherein the level of access specifies an authorization or a
privilege.
5. The method of claim 1, wherein a type of object comprises one of a command, a
30 device, or a file.
6. The method of claim 1, wherein a path-rule table is stored for each type of object.

7. The method of claim 1, wherein an entry in path-rule table comprises an absolute path for an object and an associated attribute, wherein the absolute path identifies a location of the object, and wherein the attribute identifies a level of access for the object.

5 8. The method of claim 1, wherein the path-rule table is provided by the operating system to a kernel.

9. A data processing system for managing attributes for objects, the data processing system comprising:

10 a bus;
a storage device connected to the bus, wherein the storage device contains computer usable code;
at least one managed device connected to the bus;
a communications unit connected to the bus; and
15 a processing unit connected to the bus, wherein the processing unit executes the computer usable code to store a path-rule table comprising a set of entries in a location accessible by an operating system, wherein the set of entries in the path-rule table specifies associations between objects and attributes of the objects; associate the path-rule table with an identifier; identify, in response to a request from an entity to access an object, a type of
20 the object; use the identifier to identify the path-rule table for the type of object; and selectively provide access to the object in accordance with an attribute associated with the object in the path-rule table.

10. An apparatus for managing attributes for objects, the apparatus comprising:

25 means for storing a path-rule table comprising a set of entries associated with a location accessible by an operating system, wherein the set of entries in the path-rule table is operable to specify associations between objects and attributes of the objects;
means for associating the path-rule table with an identifier;
means, responsive to a request from an entity to access an object, for identifying a
30 type of the object;
means for using the identifier to identify the path-rule table for the type of object; and

means for selectively providing access to the object in accordance with an attribute associated with the object in the path-rule table.

5 11. The apparatus of claim 10, wherein the means for storing a path-rule table further comprises:

means for creating the identifier for the path-rule table; and

means for receiving a rule finder, wherein the rule finder is operable to be used to identify the attribute in the path-rule table for accessing the object.

10 12. The apparatus of claim 10, wherein the means for selectively providing access to the object in accordance with an attribute associated with the object in the path-rule table further comprises:

means for determining a level of access to the object using the path-rule table in accordance with the entity; and

15 means for providing the entity the level of access to the object.

13. The apparatus of claim 12, wherein the level of access is operable to specify an authorization or a privilege.

20 14. The apparatus of claim 10, wherein a type of object comprises one of a command, a device, or a file.

15. The apparatus of claim 10, wherein a path-rule table is operable to be stored for each type of object.

25 16. The apparatus of claim 10, wherein an entry in path-rule table comprises an absolute path for an object and an associated attribute, wherein the absolute path is operable to identify a location of the object, and wherein the attribute is operable to identify a level of access for the object.

30 17. The apparatus of claim 10, wherein the path-rule table is operable to be provided by the operating system to a kernel.

18. A computer program product for managing attributes for objects, the computer program product comprising:

a computer usable medium having computer usable program code tangibly embodied thereon, the computer usable program code comprising:

5 computer usable program code for storing a path-rule table comprising a set of entries in a location accessible by an operating system, wherein the set of entries in the path-rule table specifies associations between objects and attributes of the objects;

computer usable program code for associating the path-rule table with an identifier; and

10 computer usable program code for identifying, in response to a request from an entity to access an object, a type of the object;

computer usable program code for using the identifier to identify the path-rule table for the type of object; and

15 computer usable program code for selectively providing access to the object in accordance with an attribute associated with the object in the path-rule table.

19. The computer program product of claim 18, wherein the computer usable program code for storing a path-rule table further comprises:

computer usable program code for creating the identifier for the path-rule table; and

20 computer usable program code for receiving a rule finder, wherein the rule finder is used to identify the attribute in the path-rule table for accessing the object.

20. The computer program product of claim 18, wherein the computer usable program code for selectively providing access to the object in accordance with an attribute associated with the object in the path-rule table further comprises:

25 computer usable program code for determining a level of access to the object using the path-rule table in accordance with the entity; and

computer usable program code for providing the entity the level of access to the object.

30

21. The computer program product of claim 20, wherein the level of access specifies an authorization or a privilege.

22. The computer program product of claim 18, wherein a type of object is one of a command, a device, or a file.

5 23. The computer program product of claim 18, wherein a path-rule table is stored for each type of object.

10 24. The computer program product of claim 18, wherein an entry in path-rule table comprises an absolute path for an object and an associated attribute, wherein the absolute path identifies a location of the object, and wherein the attribute identifies a level of access for the object.

25. The computer program product of claim 18, wherein the path-rule table is provided by the operating system to a kernel.

15

Figure 1


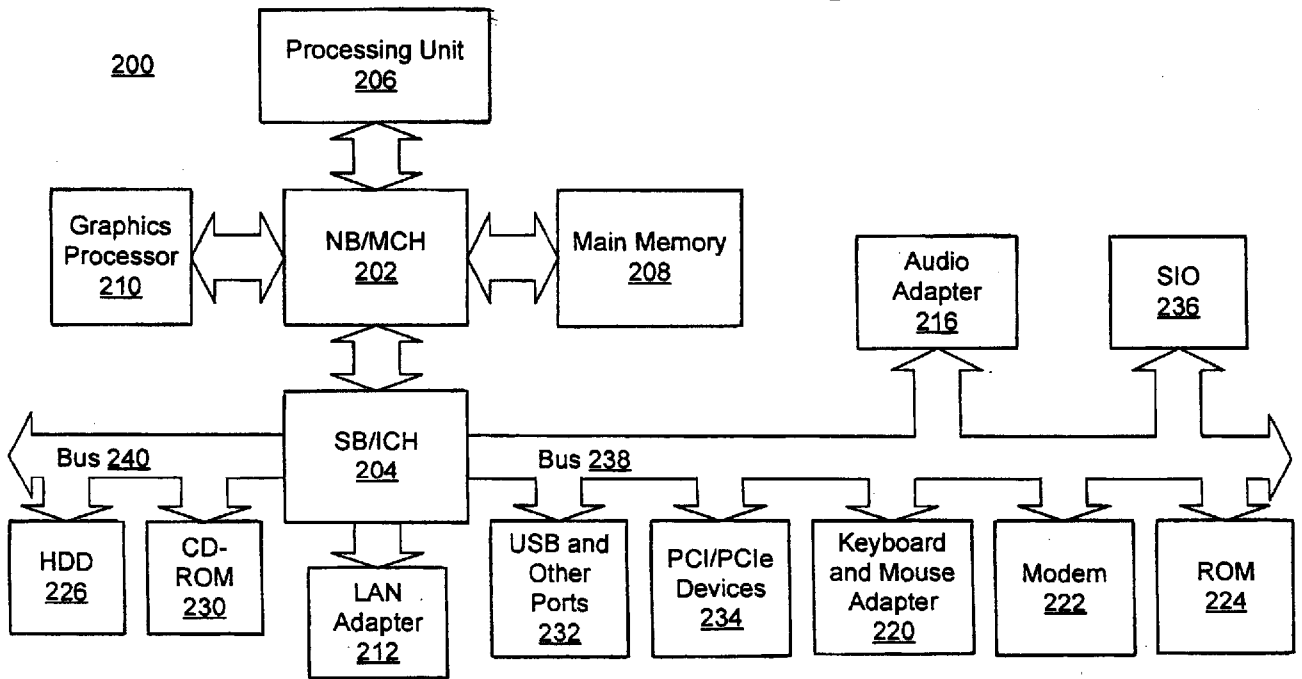
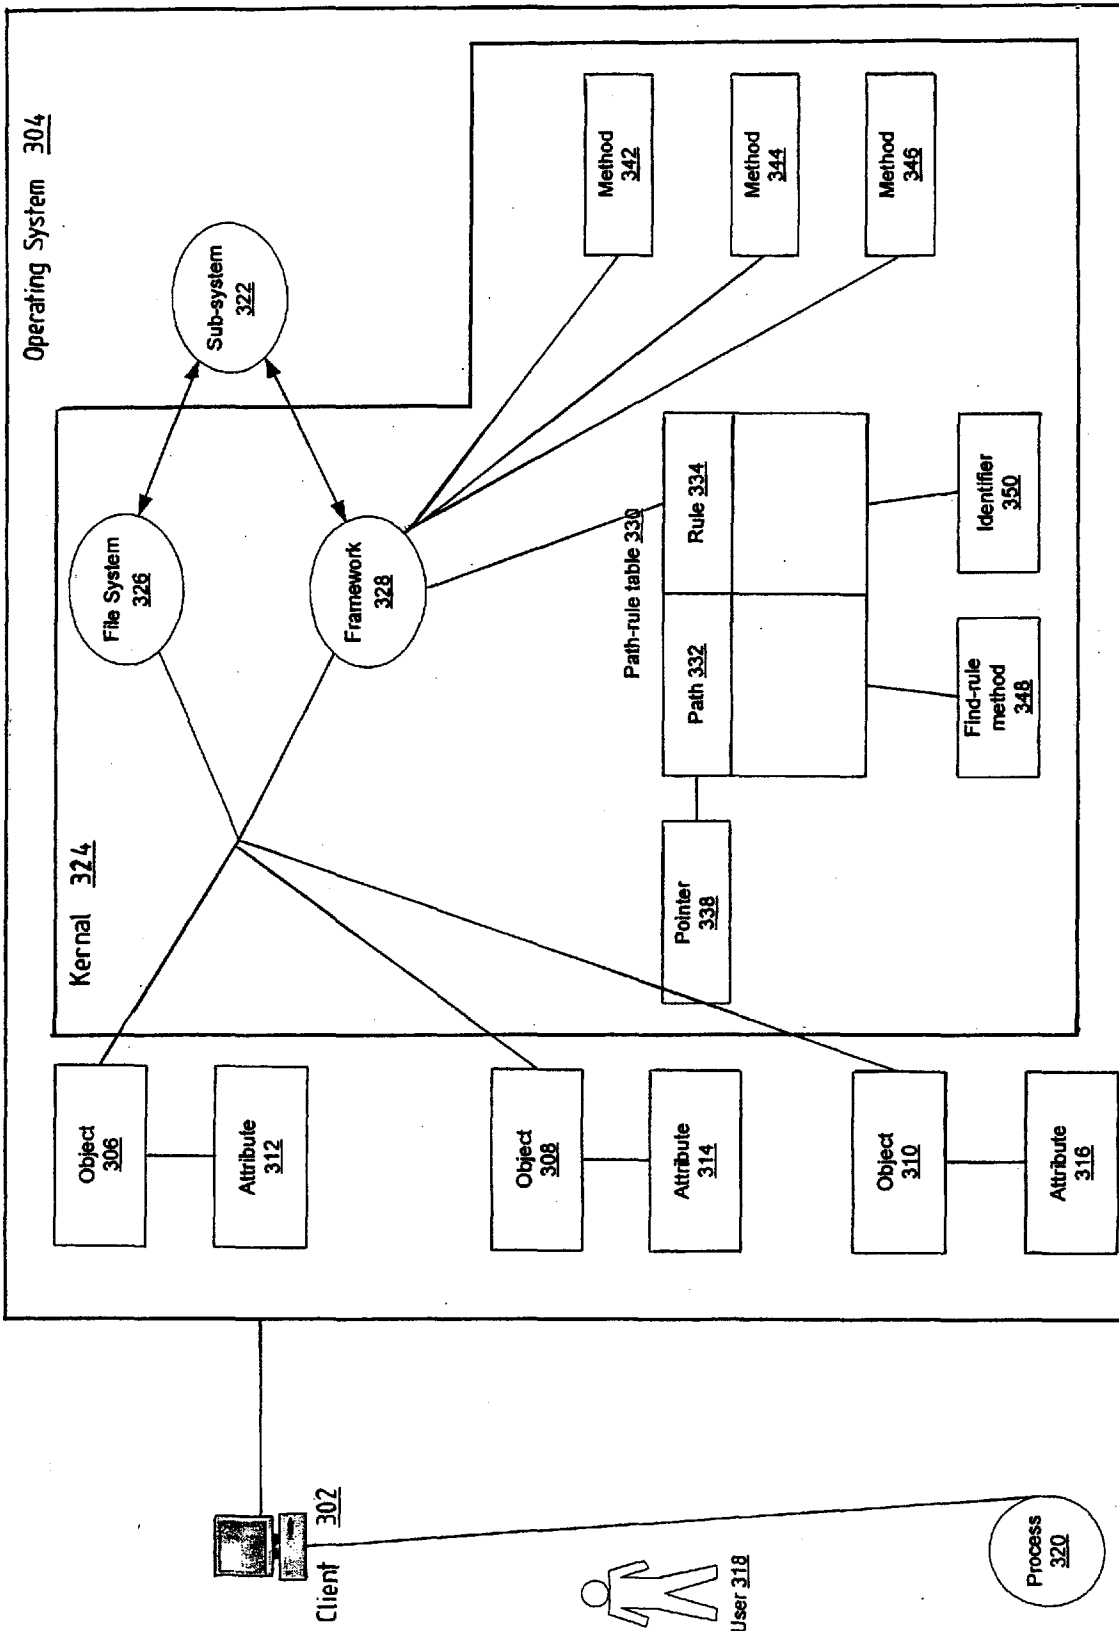


Figure 2





300

Figure 3

Figure 4

400

| Absolute Path | Access Rule | Rule <u>408</u> |
|-----------------|---|-----------------|
| /usr/sbin/mount | accessauths=aix.fs.manage.mount innateprivs=PV_DAC_W,PV_FS_MOUNT,PV_SU_UID secflags=FSF_EPS | Rule <u>408</u> |
| /usr/bin/chrole | accessauths=aix.security.role.change innateprivs=PV_AU_ADD,PV_AU_PROC,PV_SU_UID secflags=FSF_EPS accessauths=aix.fs.object.owner | Rule <u>410</u> |
| /usr/bin/chown | innateprivs=PV_DAC_X,PV_DAC_O,PV_FS_CHOWN secflags=FSF_EPS | Rule <u>412</u> |

Path 402

Path 404

Path 406

4/6

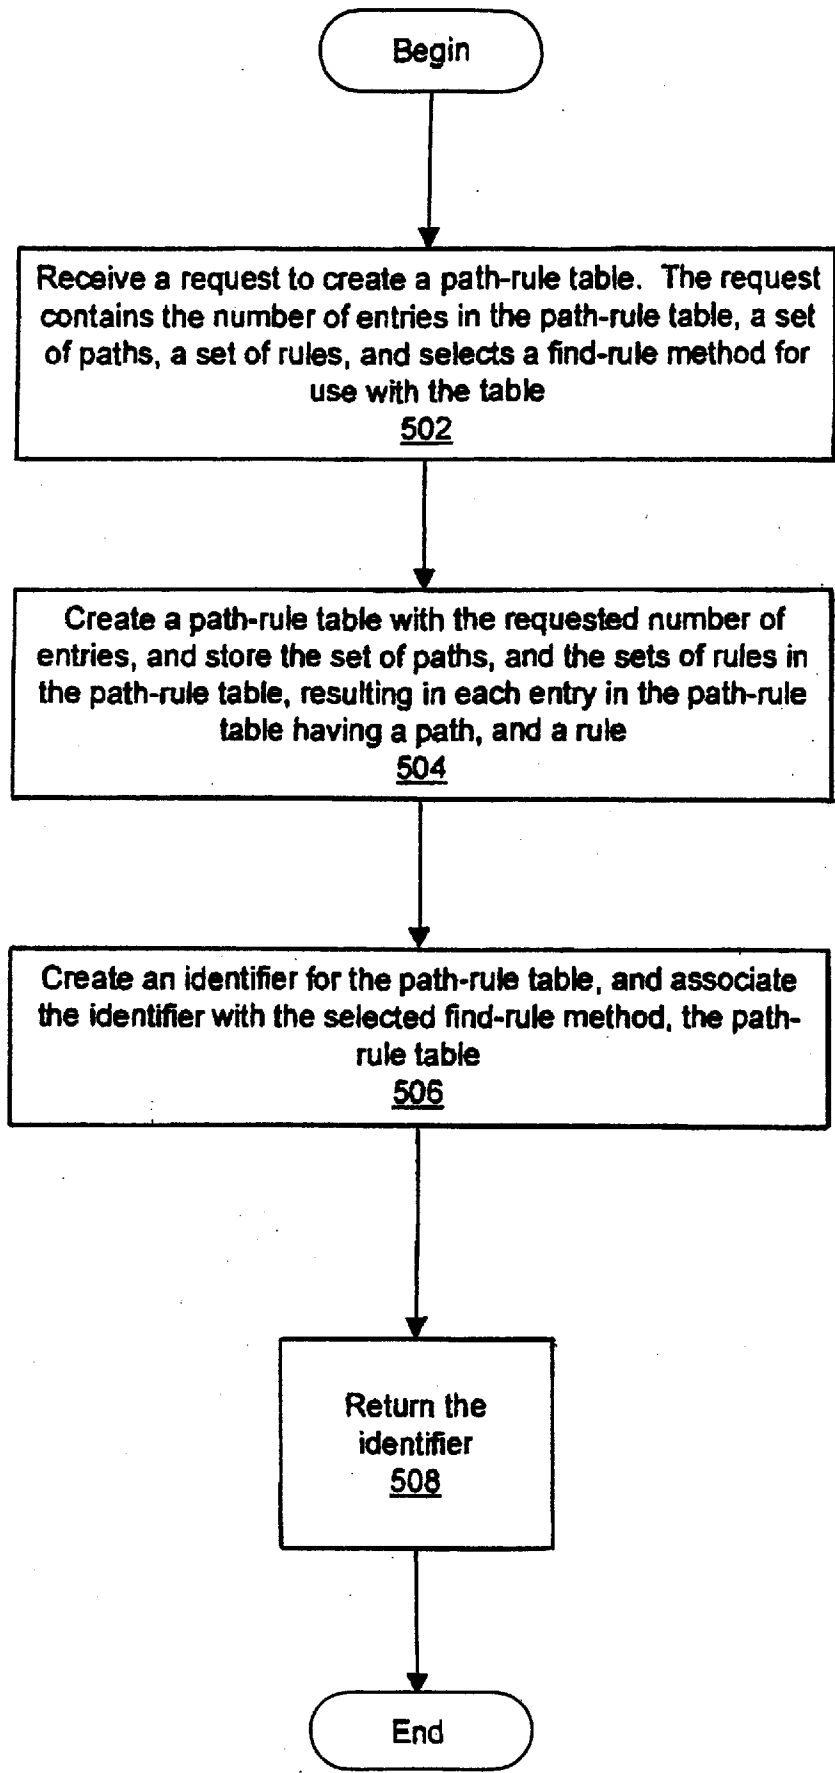


Figure 5

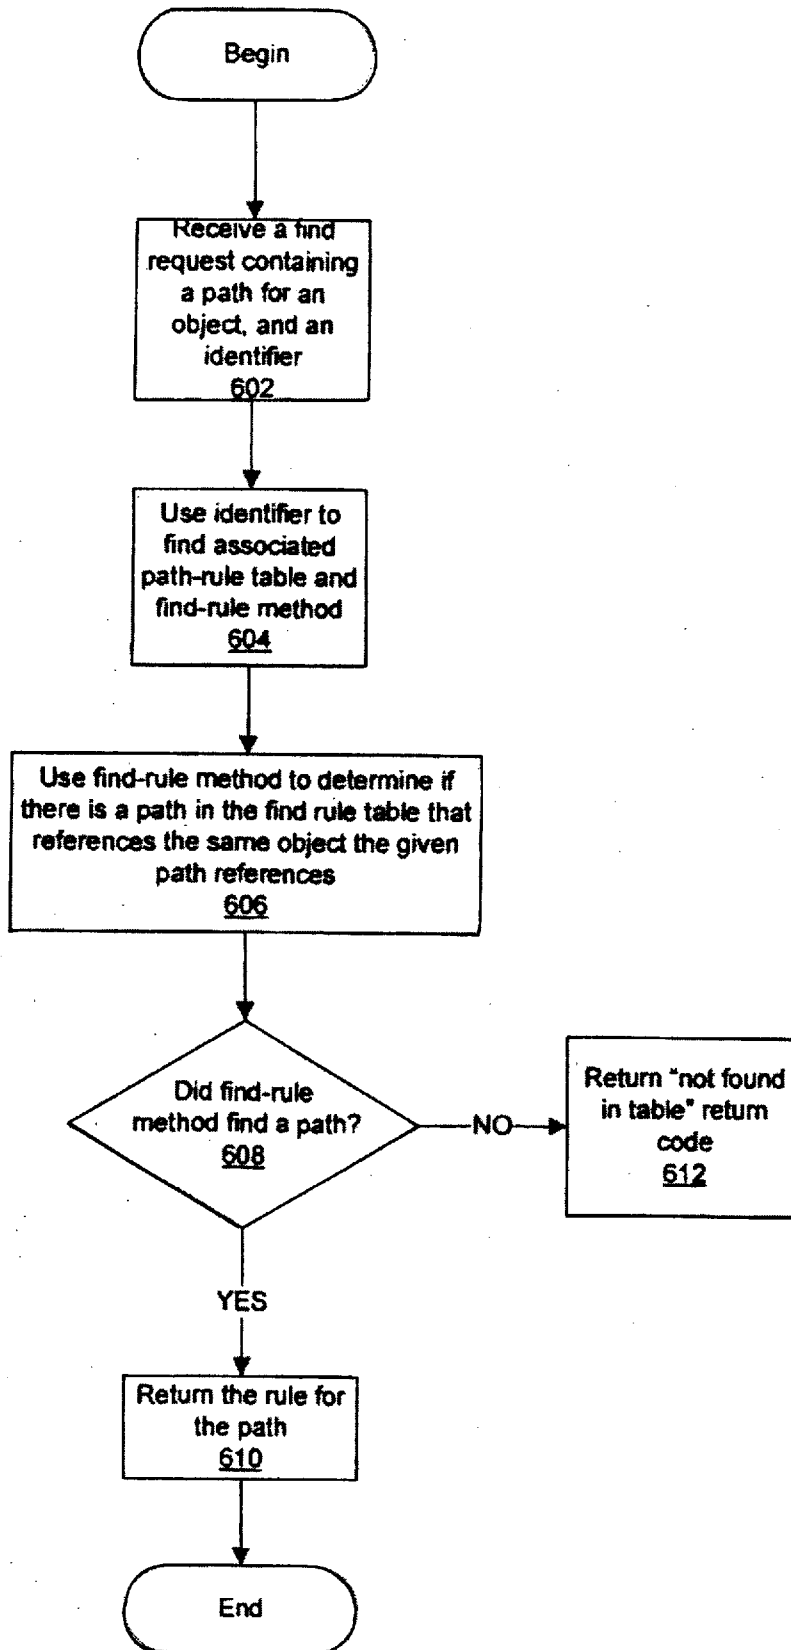
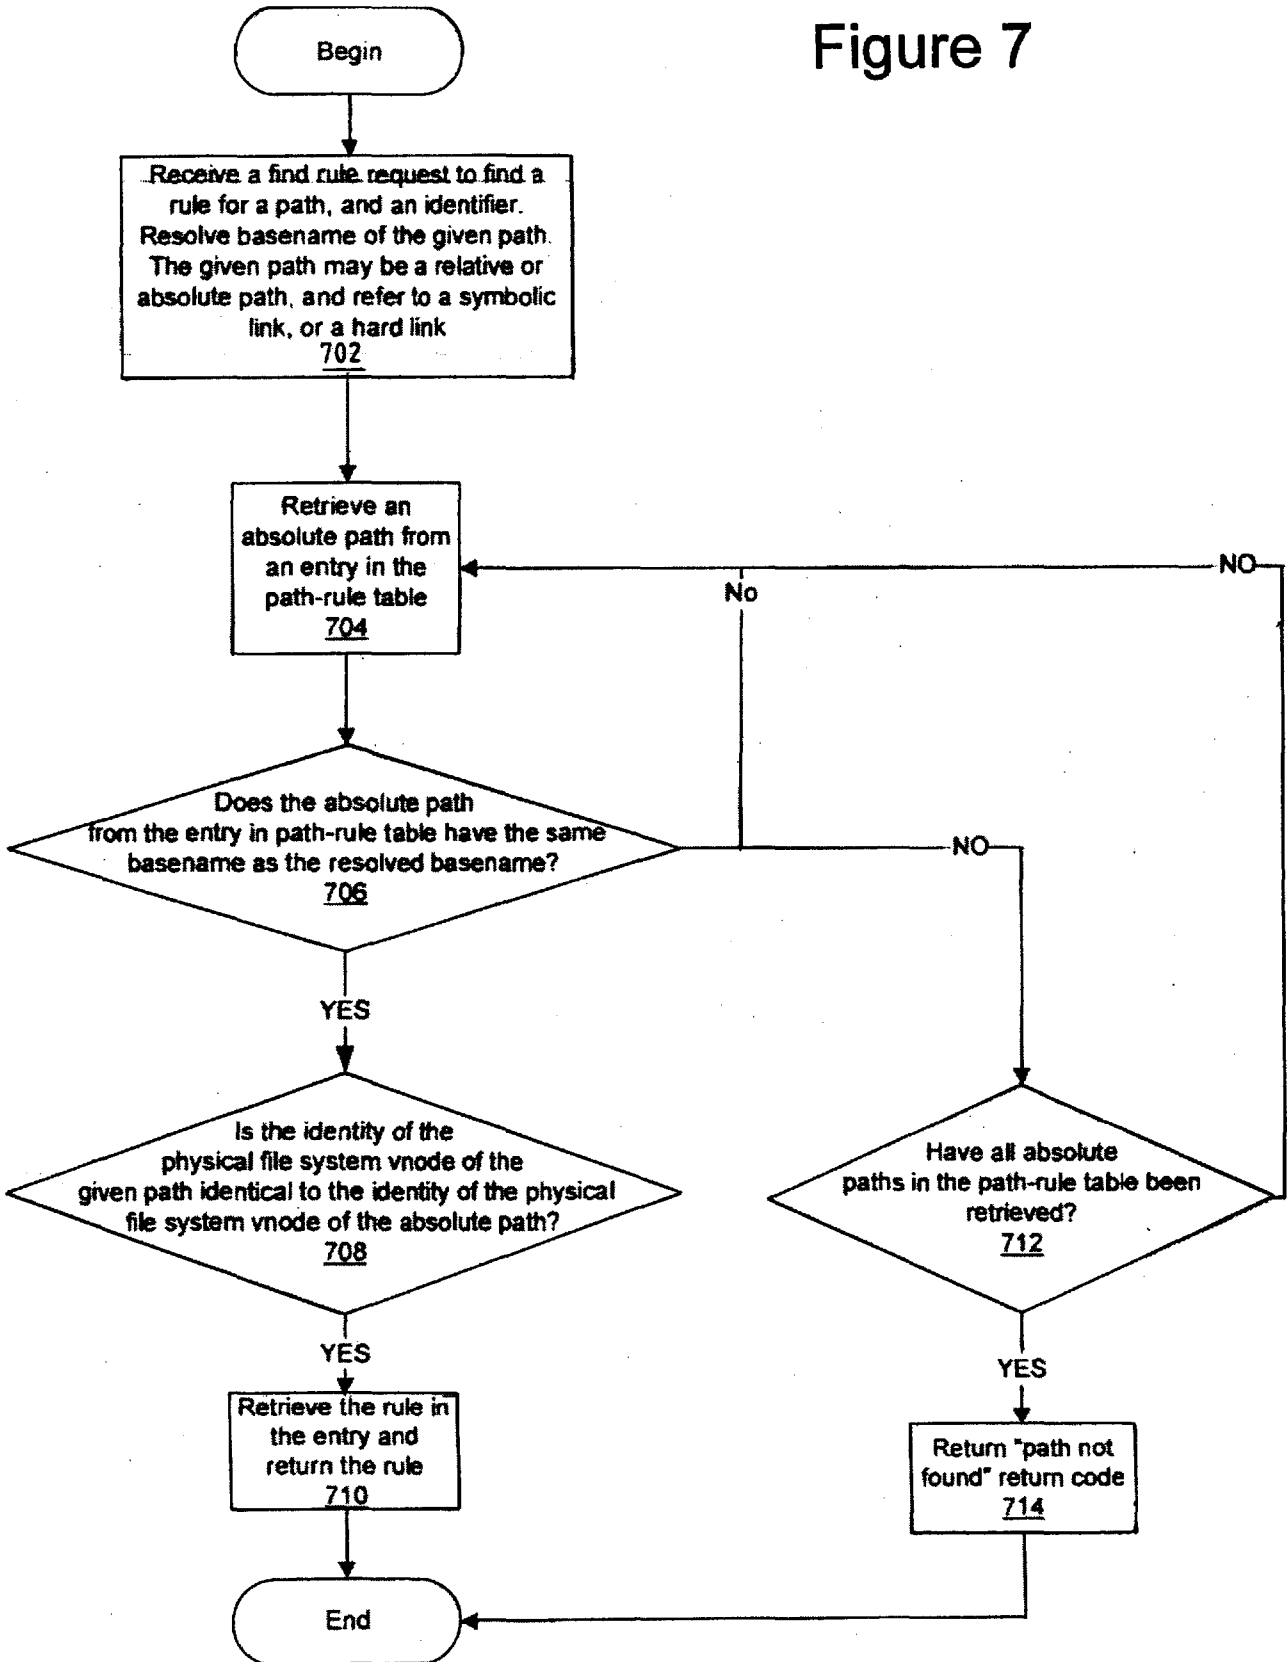


Figure 6

Figure 7



INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2008/052937

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F9/46

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 H04L

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

C. DOCUMENTS CONSIDERED TO BE RELEVANT

| Category* | Citation of document, with indication, where appropriate, of the relevant passages | Relevant to claim No. |
|-----------|--|-----------------------|
| X | <p>MARK G. SOBELL: "A Practical Guide to Red Hat Linux - Chapter The Linux Filesystem > ACLs: Access Control Lists" [Online] 27 June 2006 (2006-06-27), PRENTICE HALL, XP002486070 Retrieved from the Internet: URL: http://proquest.safaribooksonline.com/0132280272/ch06lev1sec7 [retrieved on 2008-06-27] the whole document</p> | 1-25 |
| A | <p>US 4 525 780 A (BRATT RICHARD G [US] ET AL) 25 June 1985 (1985-06-25) the whole document</p> | |

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.
- *Z* document member of the same patent family

Date of the actual completion of the international search

Date of mailing of the international search report

27 June 2008

25/07/2008

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

Mühlenbrock, Martin

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/EP2008/052937

| Patent document cited in search report | Publication date | Patent family member(s) | Publication date |
|--|------------------|-------------------------|------------------|
| US 4525780 | A | 25-06-1985 | NONE |