

(51) International Patent Classification:
G06F 9/44 (2006.01)(21) International Application Number:
PCT/US2015/019766(22) International Filing Date:
10 March 2015 (10.03.2015)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/951,403 11 March 2014 (11.03.2014) US(71) Applicant: CITRIX SYSTEMS, INC. [US/US]; 851
West Cypress Creek Road, Fort Lauderdale, Florida 33309
(US).(72) Inventors: LACHAMBRE, Sheldon Ferdinand; 9a
Devonshire Square, First Floor, London E2N 4YN (GB).
THORNLEY, David Gregory; 801 Providence Dr., Al-
gonquin, Illinois 60102 (US).(74) Agent: CAPRON, Aaron J.; Finnegan, Henderson, Fara-
bow, Garrett & Dunner, LLP, 901 New York Avenue NW,
Washington, DC, District of Columbia 20001-4413 (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, BN, 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, IR, IS, JP, KE, KG, KN, KP, KR,
KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG,
MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM,
PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC,
SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, 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, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ,
TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU,
TJ, TM), European (AL, 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, RS, SE, SI, SK,
SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ,
GW, KM, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

— as to applicant's entitlement to apply for and be granted a
patent (Rule 4.17(ii))

Published:

— with international search report (Art. 21(3))


(54) Title: COMPUTER-IMPLEMENTED METHODS AND SYSTEMS FOR DETERMINING APPLICATION MATCHING
STATUS

FIG. 9

(57) Abstract: Computer-implemented systems and methods are provided for determining application matching status. In one imple-
mentation, a method is implemented with one or more processors and includes accessing, at a server, a first dependency tree repres-
enting a first application and a second dependency tree, and acquiring one or more values for the first dependency tree and one or
more values for the second dependency tree. The method also includes comparing the one or more values of the first dependency
tree with the one or more values of the second dependency tree. The method further includes determining a matching status between
the first application and an application represented by the second dependency tree based on the comparison, and providing, for dis-
play, an indication of the matching status.

COMPUTER-IMPLEMENTED METHODS AND SYSTEMS FOR DETERMINING APPLICATION MATCHING STATUS

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims the benefit of U.S. Provisional Application No. 61/951,403, filed March 11, 2014, which is herein incorporated by reference in its entirety.

BACKGROUND

[0002] Today's rapidly changing technology requires application users to periodically migrate from one system to another, such from as a system implemented using Windows 7TM platform to a system implemented using Windows 8TM platform, or changes to runtime environment, such as App-VTM. For every migration, users have to validate each application to ensure compatibility with the new system. This process might need to be repeated for every major and/or minor new release of a given system. As users demand new and improved features, the frequency of updating systems increases, thereby also increasing the importance of the migration from one system or version to another.

[0003] Such migration is critical for enterprise users running thousands of applications for any given project on any given system. And each such application can be associated with many different algorithms. That is, verification of applications and their associated algorithms for any given system can be very complex and time consuming. Traditionally, such migration and verification is implemented manually, costing the organization a significant number of man hours (or months or years) and cost.

[0004] While a number of current techniques can be used to perform an accurate analysis to detect compatibility and matching of an application across different systems, the current techniques often involve manual testing of hundreds of applications using diagnostic

tools such as process monitoring and debuggers to discover any defects. A disadvantage for such a process is the specialized knowledge required by a person in-charge of the testing.

[0005] Another drawback with the current techniques is an inefficiency associated with application testing, where each application of a set of identical applications that are residing across different systems is verified for compatibility for the new platform. A reduction in the amount of time it takes to verify compatibility of applications between platforms is useful for applications users, especially when migrating between systems.

[0006] Yet another drawback with the current techniques is an inefficiency associated with verifying that each application of a set of identical applications that are residing across different systems is installed correctly. A reduction in the amount of time it takes to verify installation of applications between platforms is useful for applications users, especially when migrating between systems.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] Reference will now be made to the accompanying drawings showing exemplary embodiments of this disclosure. In the drawings:

[0008] **FIG. 1** is a block diagram of an exemplary network environment, consistent with embodiments of the present disclosure;

[0009] **FIGs. 2A-2B** are block diagrams of an exemplary computing device, consistent with embodiments of the present disclosure;

[0010] **FIG. 3A** is a block diagram of an exemplary process implementing an application matching system, consistent with embodiments of the present disclosure;

[0011] **FIG. 3B** is a block diagram illustrating exemplary application heuristic data points, consistent with embodiments of the present disclosure;

[0012] FIG. 4 is a block diagram illustrating an exemplary application dependency tree, consistent with embodiments of the present disclosure;

[0013] FIG. 5 is a block diagram of an exemplary database of an application matching system, consistent with embodiments of the present disclosure;

[0014] FIG. 6 is a block diagram showing an exemplary embodiment of application matching system performing an analysis of an application, consistent with embodiments of the present disclosure;

[0015] FIG. 7 is a block diagram showing an exemplary embodiment of application matching system receiving a matching status request between two or more applications, consistent with embodiments of the present disclosure;

[0016] FIG. 8 is a block diagram showing an exemplary embodiment of application matching system analyzing dependency trees associated with applications, consistent with embodiments of the present disclosure;

[0017] FIG. 9 is a block diagram showing an exemplary embodiment of application matching system determining and displaying matching status of an application, consistent with embodiments of the present disclosure; and

[0018] FIG. 10 is a flowchart illustrating an exemplary method for determining application matching status, consistent with embodiments of the present disclosure.

DETAILED DESCRIPTION OF DRAWINGS

[0019] Reference will now be made in detail to the exemplary embodiments consistent with the embodiments disclosed herein, the examples of which are illustrated in the accompanying drawings. Wherever possible, the same reference numbers will be used throughout the drawings to refer to the same or like parts.

[0020] The present disclosure generally relates to application matching between applications residing in different systems. Exemplary embodiments relate to computer-implemented methods and systems for machine learning to determine a matching status for applications residing in different systems. Disclosed embodiments provide multiple technological advantages, including computerized application analyses to identify applications compatible with a new or updated platform. Moreover, the disclosed embodiments significantly reduce bandwidth requirements for application testing by transmitting hashes of an application's dependency tree, rather than a fuller description of the application binary and meta data. Thus, the disclosed embodiments enhance the functionality of computer hardware and software by ensuring compatibility while identifying any incompatibilities without the need for extensive manual testing and intervention.

[0021] **FIG. 1** is a block diagram of an exemplary network environment 100. While exemplary network environment 100 is directed to a virtual network environment, it is appreciated that the network environment can be any type of network that communicates using packets. Network environment 100 can include one or more client devices 102A-F, a public network 104, a private network 110, and a data center 120.

[0022] One or more client devices 102A-F are devices that can acquire remote services from data center 120 through various means. Client devices 102A-F can communicate with data center 120 either directly (e.g., client device 102E) or indirectly through a public network 104 (e.g., client devices 102A-D) or a private network 110 (e.g., client device 102F). When client device 102 communicates through public network 104 or private network 110, a communication link can be established. For example, a communication link can be established by public network 104, thereby providing a client device (e.g. client devices 102A-D) access to data center 120. A communication link can also be established by private network 110, thereby providing a client device (e.g. client

device 102F) access to data center 120. While client devices 102A-F are portrayed as a computer (e.g., client devices 102A, 102E, and 102F), a laptop (e.g., client device 102B), a tablet (e.g., client device 102C), and a mobile smart phone (e.g., client device 102D), it is appreciated that client device 102 could be any type of device (e.g., wearable device) that communicates packets to and from data center 120. For example, client devices 102A-F can be referred to as a computing device, an electronic device, a mobile device, a communications device, and the like.

[0023] Public network 104 and private network 110 can be any type of network such as a wide area network (WAN), a local area network (LAN), or a metropolitan area network (MAN). As an example, a WAN can be the Internet or the World Wide Web, and a LAN can be a corporate Intranet. Public network 104 and private network 110 can be a wired network, a wireless network, or a combination of both.

[0024] Data center 120 can be a central repository, either physical or virtual, for the storage, management, and dissemination of data and information pertaining to a particular public or private entity. Data center 120 can be used to house computer systems and associated components, such as one or physical servers, virtual servers, and storage systems. Data center 120 can include, among other things, one or more servers (e.g., server 122), a desktop delivery controller 124, a virtual desktop 126, applications 128, and a backend system 130. In this disclosure the term system refers to any system capable of hosting software applications. For example, a data center 120 can be a system that is hosting software applications (e.g., applications 128). In some embodiments, data center 120 can include more than one system such that each of those systems can be different from one another. For example, when company A merges (or acquires) company B, each company having their own system that can be different from the other, the combined company (including company A and company B) can have a single system or two different systems

hosting software applications. A system hosting software applications can be referred to as a hosting system.

[0025] Server 122 can be an entity represented by an IP address and can exist as a single entity or a member of a server farm. Server 122 can be a physical server or a virtual server. In some embodiments, server 122 can include a hardware layer, an operating system, and a hypervisor creating or managing one or more virtual machines. Server 122 can provide one or more services to an endpoint. These services can include providing one or more desktops operating on one or more systems and one or more applications 128 to one or more endpoints (e.g., client devices 102A-F). For example, the one or more applications 128 can include WindowsTM- or SAPTM-based applications and computing resources. In some embodiments, server 122 can provide application matching services. For example, application matching between applications residing on one platform and applications residing on another (e.g., Window 7TM and Windows 8TM) can be performed at server 122. Moreover, the matching services can include applications compatibility services and application migration services.

[0026] In some embodiments, server 122 can provide application matching services. For example, application matching between applications residing in different systems (e.g., system A of organization A and system B of organization B) and/or different software platforms or platforms (e.g., Window 7TM and Windows 8TM) can be performed at server 122. Alternatively, application matching can be performed so as to identify any application on disparate systems that matches with a golden reference (an exemplar application). A platform can include, for example, a WindowsTM-based platform (e.g., Windows 8TM), an AppleTM-based platform (Mac OS XTM), or a mobile-based platform such as Apple iOSTM. It will be understood, in some embodiments, that the term platform as used herein can also refer to a hardware platform. It will also be understood, in some embodiments, that a system can

include one or more platforms. For example, system A of organization A can include a plurality of platforms, Window 7TM and Windows 8TM.

[0027] Desktop delivery controller 124 can be a device that enables delivery of services, such as virtual desktops 126 to client devices (e.g., client devices 102A-F). Desktop delivery controller 124 can provide functionality required to manage, maintain, and optimize all virtual desktop communications. In some embodiments, desktop delivery controller 124 can control, manage, maintain, or optimize the provisioning of applications 128. In some embodiments, controlling, managing, maintaining, or optimizing the provisioning of applications 128 can be implemented for one or more systems. That is, desktop delivery controller 124 can enable delivery of application migration and matching services for one or more systems.

[0028] In some embodiments, one or more virtual desktops 126 can provide one or more applications 128. Virtual desktops 126 can include hosted shared desktops allowing multiple users to access a single shared remote-desktop-services desktop, virtual desktop infrastructure desktops allowing each user to have their own virtual machine, streaming disk images, a local virtual machine, individual applications (e.g., one or more applications 128), or a combination thereof. In some embodiments, applications 128 can be associated with one or more systems. That is, one or more virtual desktops 126 can provide one or more applications 128 that can be used to verify application matching and application migration corresponding to a system.

[0029] Backend system 130 can be a single or multiple instances of computer networking hardware, appliances, or servers in a server farm or a bank of servers. Backend system 130 can interface directly or indirectly with server 122. For example, backend system 130 can include MicrosoftTM Active Directory, which can provide a number of network services, including lightweight directory access protocol (LDAP) directory services,

Kerberos-based authentication, domain name system (DNS) based naming and other network information, and synchronization of directory updates amongst several servers. Backend system 130 can also include, among other things, an Oracle backend server, a SQL Server backend, and/or a dynamic host configuration protocol (DHCP). Backend system 130 can provide data, services, or a combination of both to data center 120, which can then provide that information via varying forms to client devices 102 or branch office 140 (not shown).

[0030] As shown in **FIGs. 2A-2B**, each computing device 200 (such as server 122 or a client device interfacing with server 122) includes a central processing unit (CPU) 221 and a main memory 222. CPU 221 can be any logic circuitry that responds to and processes instructions fetched from the main memory 222. CPU 221 can be a single or multiple microprocessors, field-programmable gate arrays (FPGAs), or digital signal processors (DSPs) capable of executing particular sets of instructions stored in a memory (e.g., main memory 222) or cache (e.g., cache 240). CPU 221 can be, for example, one or more general-purpose microprocessors, becoming one or more special-purpose microprocessors during the importing, analyzing, reporting, and/or displaying steps of application matching process as described herein, or can be a reduced instruction set of one or more microprocessors.

[0031] The memory includes a tangible and/or non-transitory computer-readable medium, such as a flexible disk, a hard disk, a CD-ROM (compact disk read-only memory), MO (magneto-optical) drive, a DVD-ROM (digital versatile disk read-only memory), a DVD-RAM (digital versatile disk random-access memory), flash drive, flash memory, RAM, one or more caches, one or more registers, or a semiconductor memory. Main memory 222 can be one or more memory chips capable of storing data and allowing any storage location to be directly accessed by CPU 221. Main memory 222 can be any type of random access memory (RAM), or any other available memory chip capable of operating as described herein. In the exemplary embodiment shown in **FIG. 2A**, CPU 221 communicates with main

memory 222 via a system bus 250. Computing device 200 can also include a visual display device 224 and an input/output (I/O) device 230 (e.g., a keyboard, mouse, or pointing device) connected through I/O controller 223, both of which communicate via system bus 250. One of ordinary skill in the art would appreciate that CPU 221 can also communicate with memory 222 and other devices in manners other than through system bus 250, such as through serial communication manners or point-to-point communication manners. Furthermore, I/O device 230 can also provide storage and/or an installation medium for the computing device 200.

[0032] **FIG. 2B** depicts an embodiment of an exemplary computing device 200 in which CPU 221 communicates directly with main memory 222 via a memory port 203. CPU 221 can communicate with a cache 240 via a secondary bus, sometimes referred to as a backside bus. In some other embodiments, CPU 221 can communicate with cache 240 via system bus 250. Cache 240 typically has a faster response time than main memory 222. In some embodiments, such as the embodiment shown in **FIG. 2B**, CPU 221 can communicate directly with I/O device 230 via an I/O port. In further embodiments, I/O device 230 can be a bridge 270 between system bus 250 and an external communication bus, such as a USB bus, an Apple Desktop Bus, an RS-232 serial connection, a SCSI bus, a FireWire bus, a FireWire 800 bus, an Ethernet bus, an AppleTalk bus, a Gigabit Ethernet bus, an Asynchronous Transfer Mode bus, a HIPPI bus, a Super HIPPI bus, a SerialPlus bus, a SCI/LAMP bus, a FibreChannel bus, or a Serial Attached small computer system interface bus.

[0033] As shown in **FIG. 2A**, computing device 200 can support any suitable installation device 216, such as a floppy disk drive for receiving floppy disks such as 3.5-inch, 5.25-inch disks or ZIP disks; a CD-ROM drive; a CD-R/RW drive; a DVD-ROM drive; tape drives of various formats; a USB device; a hard-drive; or any other device suitable for installing software and programs such as any client agent 220, or portion thereof. Computing

device 200 can further comprise a storage device 228, such as one or more hard disk drives or redundant arrays of independent disks, for storing an operating system and other related software, and for storing application software programs such as any program related to client agent 220. Optionally, any of the installation devices 216 could also be used as storage device 228.

[0034] Furthermore, computing device 200 can include a network interface 218 to interface to a LAN, WAN, MAN, or the Internet through a variety of connections including, but not limited to, standard telephone lines, LAN or WAN links (e.g., 802.11, T1, T3, 56 kb, X.25), broadband connections (e.g., ISDN, Frame Relay, ATM), wireless connections, or some combination of any or all of the above. Network interface 218 can comprise a built-in network adapter, network interface card, PCMCIA network card, card bus network adapter, wireless network adapter, USB network adapter, modem or any other device suitable for interfacing computing device 200 to any type of network capable of communication and performing the operations described herein.

[0035] **FIG. 3A** depicts a block diagram of an exemplary process 300A implementing an application matching system. Exemplary process 300A can comprise steps including importing applications 310, analyzing the imported applications 320, and generating reports 330. Applications can be the raw material for the application matching system. When applications are imported, the application matching system can interrogate each application's files, registry entries, and application programming interface (API) usage to expose the application's intrinsic data that can uniquely identify the application.

[0036] Each application can include a plurality of data points, which can be referred to as the application's intrinsic data. In some embodiments, the data points of the application's intrinsic data can be based on heuristics. For example, each application can include over 68,000 data points. An example of these data points is illustrated in **FIG. 3B**.

As shown in **FIG. 3B**, these data points can include information associated with an application's static data 370 and also runtime data 360. Application's static data 370 can include information associated with install logic, install conditions, and/or install actions. Such data can also include information associated with application's files, registry settings, and/or configuration settings. It can further include information associated with API calls, API imports, and/or API exports. Additionally, static data can include information associated with relationships and/or dependencies within an application (e.g., dependency tree of **FIG. 4**). Runtime data 360 of application's intrinsic data can include either runtime internal data regarding performance and resources, or runtime external data regarding dependencies between network, database, etc. These data points can also include an operating system's static data 380. For example, operating system's static data can include information associated with operating system's files, registry settings, and/or configuration settings. It can also include information associated operating system's API calls, API imports, and/or API exports. Additionally, operating system's static data can also include information associated with relationships and/or dependencies of an operating system. These data points can also include information associated with external context 350. For example, external context 350 can include information associated with users and/or machines used by the users. External context 350 can also include information associated with usage and/or utilization of application matching system resources.

[0037] In some exemplary embodiments, an application's intrinsic data can be organized as a dependency tree that represents a relationship between various executable files and API calls. An exemplary dependency tree (e.g., dependency tree 400) is shown in **FIG. 4**. Dependency tree can be created for any application that uses accessible linkage such as, for example, API imports and API exports that can be queried. For example, accessible linkage can include import address table (IAT) and export address table (EAT) in a portable

execution or any other executable file. A dependency tree can be created by traversing through files (e.g., executable files like .exe files and/or .dll files) and/or API calls (e.g., calls associated with IAT and EAT) included in the application. As shown in **FIG. 4**, dependency tree 400 can include an executable file (e.g., File 1.exe 410) at the top-level hierarchy, this could also be a library file like a DLL at the top-level hierarchy. File 1.exe 410 can include import functions (e.g., Imports 415). Imports 415 can include functions, AFunc() 416, BFunc() 417, and GetDate() 418. Functions AFunc() 416 and BFunc() 417 can depend on executable file File 2.dll 430. That is, File 1.exe 410 can directly depend on a second file, File 2.dll 430. Similarly, File 1.exe can also directly depend on another file, Kernel32.dll 460 through function, GetDate() 418. As shown in **FIG. 4**, Kernel32.dll 460 can include an export function (e.g., Exports 465) that comprises the function, GetDate() 466.

[0038] File 2.dll 430 can include export functions (e.g., Exports 435) and import functions (e.g., Imports 445). As shown in **FIG. 4**, Exports 435 can include functions AFunc() 436, BFunc() 437, and CFunc() 438. Imports 445 can include functions FEFunc() 446 and FFunc() 447, which can depend on another file, File 3.dll 450. That is, File 2.dll 430 can directly depend on a third file File 3.dll 450 and File 1.dll 410 can indirectly (through second file, File 2.dll 430) depend on the third file File 3.dll 450. **FIG. 4** also shows that File 3.dll 450 can include export functions (e.g., Exports 455) that comprise functions FEFunc() 456 and FFunc() 457. As shown in **FIG. 4**, dependency tree 400 can depict relationships and/or linkages between various files and API calls included in an application. While dependency tree 400 shows a dependency tree that includes only three levels of hierarchy (e.g., File 1.exe 410 depending on File 2.dll 430 and further on File 3.dll 450), it is understood that an application can include dependencies between files (e.g., executables files) and/or API calls (e.g., calls associated with import address table and export address table) that can include any number of hierarchies.

[0039] Referring back to **FIG. 3A**, applications imported in importing step 310 can comprise any kind of application including, but not limited to, desktop applications (e.g., MSI applications 302, non-MSI applications 304), or web based applications (for example, web applications 306). MSI applications 302 can include, for example, Microsoft Installer™ or Windows Installer™ based applications. That is, MSI applications 302 can be applications that use a .msi file (Microsoft Installer™ or Windows Installer™) in the process of installation. Non-MSI applications 304 can be applications that use an interface other than Microsoft Installer™ or Windows Installer™ interface for installation, such as Microsoft App-V™ or a custom developed setup.exe file. In some embodiments, the imported applications can include any application or process that includes dependent linkages that, for example, depict relationships and/or linkages between various files and API calls.

[0040] Importing step 310 can be implemented differently between importing desktop applications and web applications. For example, Windows™ desktop applications can be imported using the applications' installation packages. These installation packages can include an .msi file or any other type of installation package used in the industry. In some embodiments, such installation packages can be App-V™ (e.g., .sft or .appv) packages. Web applications can be imported into the application matching system by using one of the two different exemplary approaches described below, or a combination of those approaches. In some embodiments, web applications can be imported by a directed web spider or web crawler to crawl over the run-time HTML pages and capture them for import into the application matching system. In some embodiments, importing can be accomplished by importing the web applications' source files into a database of application matching system.

[0041] **FIG. 5** depicts such an exemplary database, database 500, of an application matching system. In accordance with some embodiments, database 500 can be a structured collection of tables, lists, or other data for design verification setup and runtime execution.

As such, the structure may be organized as a relational database or an object-oriented database. In some embodiments, database 500 can be a hardware system comprising physical computer readable storage media and input and/or output devices configured to receive and provide access to tables, lists, or other data structures. Furthermore, configured as a hardware system, database 500 can include one or more processors and/or displays.

[0042] In some embodiments, database 500 can reside within a server (e.g., server 122) hosting the application matching system. In some embodiments, database 500 can reside on a server (or on a distributed network of servers) remote from a server hosting the application matching system. Applications imported in importing step 510 of the application matching system can reside in database 500. Each such application can be associated with one or more matching algorithms (not shown), which can also reside in database 500. In some embodiments, matching algorithms are created by a user of the application matching system, whereas in other embodiments, such algorithms can already be created and saved in a universal algorithm library residing in database 500.

[0043] Database 500 can also include information associated with operating systems (e.g., operating system intrinsic data) 520, running on computing device 200. For example, operating systems running on computing device can include Windows 7™, Windows Server 2003™, Windows XP™, and Windows Server 2008 R2™. In some embodiments, information associated with operating system (e.g., operating system image or snapshot) can be uploaded by a user of the application matching system, whereas in other embodiments such information can be extracted by an application matching system itself. For example, operating system intrinsic data can include information describing a certificate chain on an operating system image. Using the signing certificate authority on the certificate chain, an algorithm can simply process the certificate chain until it identifies the certificate as either trusted or untrusted. It will be understood that the application matching system can also

allow users to load multiple operating system images concurrently to enable the user to execute algorithms for matching with all such platforms in a single pass. Database 500 can further include information associated with usage and/or runtime data 530. Database 500 can further include information associated with auditing, performance, users, groups, and/or devices 540, as depicted in **FIG. 5**.

[0044] **FIG. 6** depicts an exemplary analysis step 320 of an application matching system. After the applications are imported into the application matching system, those applications can be analyzed to check for matching between applications associated with different systems. An exemplary analysis process can include analyzing an application's intrinsic data comprising information associated with an application's install logic, files, registry settings, configuration settings, or API calls. In some embodiments, an analysis process can also include analyzing any of an operating system image data 630, external data 620, journal entries (not shown), or configuration manager (not shown) data associated with runtime environment. The analysis process can include analyzing an application's intrinsic data to create a dependency tree (e.g., dependency tree 400) of the application residing in a particular system or platform. A typical analysis process involves analyzing imported applications 610 just once. There may be embodiments, however, where imported applications 610 may need to be re-analyzed, for example, based on changes to any attributes of the one or more systems including an operating system of the one or more systems. Furthermore, if there are any changes to any of the other attributes of the one or more systems, imported applications 610 can also be re-analyzed.

[0045] In some embodiments, a matching status of an application can be determined by computing and comparing a digital signature associated with the application that uniquely identifies the application. A digital signature associated with an application can be a specific pattern. For example, an application's digital signature can be a hash value returned by a

hash function. It is well understood that a hash function can be any algorithm that maps data of variable length to data of a fixed length (e.g., MD5 and SHA-2). It is well understood that even a one-bit change in the input data can result in a significant change in the hash value of the input data. Said differently, hash functions resulting in hash values are a good exemplary function that can uniquely identify a variable length of data including binary files. The values returned by a hash function can be called hash values, hash codes, hash sums, checksums, or simply hashes. In some embodiments where an application's digital signature is defined as a hash value returned by a hash function operated on an entire dependency tree of the application, an application matching system can determine a matching status between two applications by comparing each application's hash value. If the comparison results in a match between the hash values of the two applications, the application matching system can identify that the two applications are the same applications that may be residing on two different platforms or systems.

[0046] If, on the other hand, the comparison does not result in a match between the hash values of the two applications, the application matching system can identify that the two applications do not match. In some embodiments, when the comparison does not result in a match between a top-level hash values of the dependency trees of the two applications (e.g., File 1.exe 410), a hash value can be computed on a second-level hierarchy of the dependency tree (e.g., File 2.dll 430) for further comparison. If the comparison of hash values at the second-level hierarchy of the dependency trees results in a match between the two applications, the application matching system can identify that the two applications include the same files at the second-level hierarchy (e.g., File 2.dll 430) that were used for computing the hash value. If, however, the comparison of the hash values at the second-level hierarchy of the dependency tree does not match, then the process can be repeated to the third-level hierarchy of the dependency tree (e.g., File 3.dll 450), and so on. It will be understood that

the hash values can be computed and later compared for any number of hierarchies of an application's dependency tree. It will be understood that an application's digital signature value can be defined in many ways other than and different from the exemplary hash function described above.

[0047] Imported applications can be associated with matching algorithms that can be used in the analysis process. In some embodiments, each imported application can be associated with one or more matching algorithms. Matching algorithms can be implemented as heuristic algorithms 640. In some embodiments, matching algorithms can relate to information associated with the one or more systems and/or platforms including internal states, dependencies, relationship between one application to another, relationship between an application and an operating system, or external data. Matching algorithms can also relate to operating system image-dependent algorithms, including dependencies on features that are provided by an operating system. When relevant, these algorithms can interrogate the operating system image data that has been loaded into the database (e.g., database 500) of the application matching system. These algorithms can check a variety of operating system image information 630, including: APIs, registry information, file management system for each fixed drive partition, the certificate store, or some matching settings.

[0048] The application matching system can generate reports (e.g., report step 330), as an output of analysis step 320. For example, such reports can contain information about a matching status for two or more applications that have been analyzed by the application matching system to determine whether the two or more applications match. In some embodiments, reports can include a plurality of icons to indicate a matching status for two or more applications. For example, icons represented by red (R), amber (A), and green (G) colors or indications can signify different matching statuses for applications (illustrated in **FIG. 9**). An icon G can indicate that a matching status between two or more applications is

at least above a predetermined percentage. An icon R can indicate that a matching status between two or more applications is below a predetermined percentage. An icon A can indicate either that a matching status between two or more applications is unknown or that a matching determination is unsuccessful. It will be understood the above-listed definitions of icons are merely exemplary and other definitions are possible. An exemplary method describing a process of determining a matching status between two or more applications is described in **FIGs. 7-10** below.

[0049] **FIGs. 7-9** depict block diagrams illustrating how application matching system 701 determines a matching status between two or more applications, consistent with embodiments of the present disclosure. **FIG. 7** depicts an application matching system receiving, at 715 and 725, a request for determining a matching status between a first application (e.g., first application 710) and a second application (e.g., second application 720). In some embodiments, each of first application 710 and second application 720 can be located in a different hosting system. Requests 715 and 725 can be received at application matching system 701. Application matching system 701 can be similar to an application matching system described in **FIG. 3**. For simplicity, a block diagram of application matching system 701 is depicted with only server 122 that can host an application matching system and any one or more client devices 102A-F. Receiving 715 can be associated with first application 710 and receiving 725 can be associated with second application 720. First application 710 and second application 720 can either be an MSI application, a non-MSI application, or a web application. In some embodiments, each of first application 710 and second application 720 can be a new application that can be received from either within the application matching system 701 or external to application matching system 701. Alternatively, each of first application 710 and second application 720 can be an existing application that is currently residing at application matching system 701. In some

embodiments, each of first application 710 and second application 720 can be a newer version of an existing application that is currently residing at application matching system 701. In some embodiments, a received request can include both the first application and the second application. Alternatively, each of first application 710 and second application 720 can be an application whose installer is missing and is currently running (or executing) such that the application's main executable can be identified and further processed to build a dependency tree.

[0050] Server 122 can host a database (e.g., database 500), where a plurality of applications can reside. In some embodiments, each of the plurality of applications residing at server 122 can be organized into a dependency tree (e.g., dependency tree 400) associated with each of the plurality of the applications. Dependency tree can be created for any application that uses accessible linkage that can be queried (such as, for example, API imports and API exports). For example, accessible linkage can include IAT and EAT in a portable execution (or a new executable) file. It is understood that an application can include dependencies between files (e.g., executables files) and/or API calls (e.g., import address table and export address table) that can include any number of hierarchies (e.g., three-levels of hierarchy shown in dependency tree 400).

[0051] **FIG. 8** depicts, at 815 and 825, an application matching system organizing each of first application 710's intrinsic data and second application 720's intrinsic data into two separate dependency trees. The process of organizing an application's intrinsic data into a dependency tree can be the same as the dependency tree described above in **FIG. 4**. As described in **FIG. 4**, a dependency tree can be created by traversing through files (e.g., executable files like .exe and/or .dll) and/or API calls (e.g., calls associated with import address table and export address tables) included in the application. In some embodiments,

first application 710 and/or second application 720 can already be organized into their respective dependency trees before requests 715 and 725 are received.

[0052] **FIG. 8** further depicts that application matching system 701 can access a first dependency tree representing first application 710 and a second dependency tree representing second application 720. As shown by labels 815 and 825 in **FIG. 8**, one or more values (e.g., hash values) can be computed for each application and then the one or more computed values for each application can be compared between the two applications. Application matching system 701, at 815 and 825, can perform a computation of one or more hash values for the first dependency tree and one or more hash values for the second dependency tree. In some embodiments, one or more hash values for a particular dependency tree can represent the entire dependency tree of the particular dependency tree. For example, a hash value can be for a top-level file associated with the dependency tree (e.g., File 1.exe 410 of dependency tree 400). Alternatively, the one or more hash value can be computed for one or more portions of the particular dependency tree (i.e., each portion can be associated with each file and/or API call of the particular dependency tree). For example, one or more hash values can be computed for a one or more files of the dependency tree (e.g., File 2.dll 410 and File 3.dll 450 of dependency tree 400). In some embodiments, the one or more values of the first dependency tree and/or the one or more values of the second dependency tree can be acquired by application matching system 701, wherein the acquired values have been previously computed and stored in a database (e.g., database 500). That is, the one or more values can either be computed within application matching system 701 or external to application matching system 701.

[0053] It will be understood that a number of hash values that can be computed for each dependency tree can be based on a total number of files (e.g., executable files) included in the dependency tree. For example, if a dependency tree includes 4 files as in dependency

tree 400 (File 1.exe 410; File 2.dll 430; File 3.dll 450; and Kernel32.dll 460), 4 different hash values can be computed for dependency tree 400. That is, a hash value can be computed for each of the files that are included in the dependency tree. Accordingly, in some embodiments, a plurality of hash values can be computed for each dependency tree and therefore for each application. In some embodiments, a hash value for a particular file can include all files that it may further depend on. For example, a hash value for File 2.dll 430 can also include File 3.dll 450. Alternatively, a hash value for a particular file can only include the particular file and might not include any file that it may further depend on. For example, a hash value for File 2.dll 430 need not include File 3.dll 450.

[0054] In some embodiments, a hash value can be computed for all hash values that have been computed for all files included in dependency tree 400 (e.g., 4 files listed above). In some embodiments, this hash value can be the application's fingerprint that uniquely identifies the application. An advantage of using a fingerprint hash value can be that when a fingerprint hash value of two applications match, there is no need to compare hash values for each file at a lower level of hierarchy of an application.

[0055] Application matching system 701 can perform a comparison between the one or more values of the first dependency tree and the one or more values of the second dependency tree. In some embodiments, the comparison can include bit-by-bit comparison of the values. For example, an MD5 hash function results in a 128-bit hash value, whereas an SHA-2 hash function can result in a hash value ranging from 224 bits to 512 bits. In embodiments using MD5 hash function, the comparison step can include comparing each bit of the 128-bit hash value between the one or more hash values of the first dependency tree and the one or more hash values of the second dependency tree. The bit-by-bit comparison can result in either a match or no match between any pair of hash values. Application matching system 701 can perform a determination of a matching status between the first

application and the second application based on the comparison of the hash values. If the bit-by-bit comparison, for example, results in a match between the one or more hash values of the first dependency tree and the one or more hash values of the second dependency tree, a matching status signifying a perfect match (the hash values are identical between the first dependency tree and the second dependency tree) between the first application and the second application can be reported. If, on the other hand, the bit-by-bit comparison results in a no match (e.g., difference of at least one-bit) between the first hash value and the second hash value, a matching status signifying no match between the first application and the second application can be reported. Application matching system 701 can report a matching status in a report (e.g., Report 330) depicted at 905 of **FIG. 9**.

[0056] **FIG. 9** depicts an application matching system displaying, at 905, a matching status between the first and second applications. In some embodiments, the matching status can be displayed on a graphical user interface. The displayed matching status can include displaying at least one status icon for each pair of applications (e.g., first application 710 and second application 720). For example, a displayed status icon can include one of the following icons: red (R), amber (A), and green (G), which can represent different statuses of matching status between first application 710 and second application 720, or portions thereof. In some embodiments, icon G can indicate that a matching status between two or more applications is at least above a predetermined percentage. An icon R can indicate that a matching status between two or more applications is below a predetermined percentage. An icon A can indicate either that a matching status between two or more applications is unknown or that a matching determination is unsuccessful. It will be understood the above-listed definitions of icons are merely exemplary and other definitions are possible. In some embodiments, a matching status between the first application 710 and the second application 720 can be displayed on a graphical user interface on one or more client devices 102A-F. In

some embodiments, a matching status between the first application 710 and the second application 720 can be displayed on a graphical user interface on a display associated with server 122 hosting application matching system 701.

[0057] **FIG. 10** illustrates a flowchart that shows an exemplary method for determining a matching status between two applications, consistent with embodiments of the present disclosure. Referring to **FIG. 10**, it will be appreciated that the illustrated procedure can be altered to change the order of the steps, delete steps, or further include additional steps.

[0058] In step 1020, an application matching system (e.g., application matching system 701) can access a first dependency tree representing a first application (e.g., first application 710) and a second dependency tree representing a second application (e.g., second application 720). In step 1030, the application matching system can acquire one or more values for the first dependency tree and one or more values for the second dependency tree, wherein the acquired values have been previously computed and stored in a database (e.g., database 500). That is, the one or more values can either be computed within the application matching system or external to the application matching system. In some embodiments, the application matching system can perform a computation of one or more hash values for the first dependency tree and one or more hash values for the second dependency tree. In some embodiments, as described in **FIG. 8**, one or more hash values for a particular dependency tree can represent the entire dependency tree of the particular dependency tree or one or more portions of the particular dependency tree (i.e., each portion can be associated with each file and/or API call of the particular dependency tree).

[0059] In step 1040, the application matching system can perform a comparison of the computed first hash value (or a first set of hash values) and the computed second hash value (or a second set of hash values). In some embodiments, the comparison can include

bit-by-bit comparison of the values. The bit-by-bit comparison can result in either a match or no match between any pair of hash values.

[0060] In step 1050, the application matching system can determine a matching status between the first application and the second application based on the comparison of the hash values. If the bit-by-bit comparison, for example, results in a match between the one or more hash values of the first dependency tree and the one or more hash values of the second dependency tree, a matching status signifying a perfect match (the hash values are identical between the first dependency tree and the second dependency tree) between the first application and the second application can be reported. If, on the other hand, the bit-by-bit comparison results in a no match (e.g., difference of at least one-bit) between the first hash value and the second hash value, a matching status signifying no match between the first application and the second application can be reported.

[0061] In step 1060, the application matching system can provide, for display a report (e.g., report 330 as depicted in **FIG. 9**) indicating the matching status between the first application and the second application.

[0062] The flowchart of **FIG. 10** can further be described with an exemplary scenario where a matching status between two applications needs to be determined. In this exemplary scenario, the application matching system determines a matching status between a first application, Microsoft Word 2007TM, and a second application, unknown application. A dependency tree exists for each of the first and second applications. In this example, the first dependency tree (representing Microsoft Word 2007TM application) includes a top-level file name, File 1.exe, whereas the second dependency tree (representing an unknown application) includes a top-level file name, SomeOtherName.exe. After the application matching system accesses the first and second applications (step 1020), in step 1030, the application matching system can compute a first hash value for File 1.exe and a second hash value for

SomeOtherName.exe. In step 1040, the application matching system can compare the first hash value with the second hash value, bit-by-bit. If the application matching system determines that the first hash value matches with the second hash value, the application matching system can generate a report showing that there is a perfect match between Microsoft Word 2007TM and the unknown application. Given the perfect match of hash values, even though the second application is unknown, it may be safe to identify the unknown application is also a Microsoft Word 2007TM application. In some embodiments, the second application can reside in a system that is different from a system where the first application resides.

[0063] Continuing with the above exemplary scenario, if the application matching system determines that there is no match between the first hash value and the second hash value, where the hash values are computed for the top-level files of the respective applications (File 1.exe and SomeOtherName.exe), steps 1030 through 1050 can be repeated by traversing down the hierarchy level of dependency trees of the applications. For example, a hash value can be computed for files at a second-level hierarchy of the dependency tree (e.g., File 2.dll 430 of dependency tree 400), for further comparison. If the comparison of hash values at the second-level hierarchy of the dependency trees results in a match between the two applications, the application matching system can identify that the two applications include the same files at the second-level hierarchy (e.g., File 2.dll 430 of dependency tree 400) and below that were used for computing the hash value. That is, while the second application might not be identical to Microsoft Word 2007TM, a second-level hierarchy match of hash values indicates that the second application shares some common files as that of Microsoft Word 2007TM application. If, however, the comparison of the hash values at the second-level hierarchy of the dependency tree does not match, then the process can be repeated to the third-level hierarchy of the dependency tree (e.g., File 3.dll 450 of

dependency tree 400). It will be understood that the hash values can be computed (step 1030) and later compared (step 1040) for any number of hierarchies of an application's dependency tree.

[0064] In some embodiments, a match between two applications can be defined when a comparison between the hash values is above a predetermined threshold. For example, the predetermined threshold can be any number equal to or less than one hundred percent. If the predetermined threshold is equal to one hundred percent, such a match can be referred to as a perfect match. Alternatively, if the predetermined threshold is less than one hundred percent, such a match can be referred to as a partial match. A partial match scenario can be applicable in cases where a hash value is computed for each file of the dependency tree and matching between two applications is determined based on whether a percentage of the number of hash values that are same between the two applications is at least equal to the predetermined threshold. For example, if each of Microsoft Word 2007TM and the unknown application have 4 files (similar to dependency tree 400) included in their respective dependency trees and let us also assume that a predetermined threshold is set as seventy percent. If a comparison of hash values for each of the four files of the two applications results in a match of three of the four files (i.e., seventy five percent), the application matching system will determine that portions of the unknown application matches portions of the Microsoft Word 2007TM application.

[0065] In practice, while a predetermined threshold of less than one hundred percent signifies that the two applications are not identical, a predetermined threshold of seventy five percent, for example, signifies that the two applications share many of the common files and can be closely related. For example, if three of the four files match between Microsoft Word 2007TM and the unknown application, it can signify that the unknown application can be a Microsoft WordTM application that is different from 2007 version of the application (e.g.,

Microsoft Word 2010TM). Another example can signify a situation where the unknown application can still be a Microsoft Word 2007TM but one of the files was not properly installed. Accordingly, a predetermined threshold of less than one hundred percent can result in useful insights in identifying closely related applications, which can help in making application migration from one system to another more efficient. In some embodiments, a match between applications can be based on hash values of both files as well as API calls. Alternatively, a match between applications can be defined such that it is based on hash values of either files or API calls but not both.

[0066] In some embodiments, the application matching system can determine a matching status between two or more applications based on an exemplar application as a reference. For example, when an organization is migrating its applications from one platform to another, applications residing in different systems of the organization need to be verified for compatibility with the new (to-be-migrated) platform, as described in U.S. Patent Application No. 14/094,697, titled “Methods and Systems for machine learning to discover application compatibility,” filed on December 2, 2013, the entirety of which is incorporated herein expressly. The application matching process described in **FIG. 10** (and depicted in **FIGs. 7-9**) can be utilized to make the application migration process more efficient. In some embodiments, the application matching system can receive an exemplar application (e.g., Microsoft Word 2007TM) whose dependency tree and associated hash value (or a set of hash values) can be first computed. The application matching system can operate on applications involved in the migration process such that the application matching system can identify all applications that match the exemplar application by computing (step 1030) and comparing (step 1040) hash values associated with each of the applications involved in the migration process.

[0067] After the application matching system determines (step 1050) a matching status for each application as compared with the exemplar application such that it identifies all applications that match the exemplar application. In such a scenario, an application compatibility system checking for compatibility of applications with a particular platform could only check the exemplar application for compatibility with the new platform. That is, after the application compatibility system verifies that the exemplar application is compatible with the new platform, the application compatibility system can automatically mark that all applications matching the exemplar application are also compatible with the new platform without having to check a compatibility status for any of those applications that match the exemplar application. As will be appreciated, the above-described scenario improves an efficiency of application compatibility with platform by reducing a processing iterations and time by utilizing the application matching process of **FIG. 10**.

[0068] In some embodiments, the above-described process of improving efficiency of an application compatibility system can also be extended to verify whether applications are correctly installed on various platforms. In some embodiments, determining a matching status between applications can be performed over a low bandwidth network connection. In such embodiments, the hash value can be a small string of data (e.g., 32-bit hash value) that can uniquely represent an application and by transmitting only hash values that are relatively short in size, a matching determination can be performed over a low bandwidth network connection. That is, application matching system can perform application matching either on one electronic device (e.g., server 122) or distributed over a network, either a high bandwidth network or a low bandwidth network.

[0069] The methods disclosed herein may be implemented as a computer program product, i.e., a computer program tangibly embodied in an information carrier, e.g., in a machine readable storage device, for execution by, or to control the operation of, data

processing apparatus, e.g., a programmable processor, a computer, or multiple computers. A computer program can be written in any form of programming language, including compiled or interpreted languages, and it can be deployed in any form, including as a standalone program or as a module, component, subroutine, or other unit suitable for use in a computing environment. A computer program can be deployed to be executed on one computer or on multiple computers at one site or distributed across multiple sites and interconnected by a communication network.

[0070] A portion or all of the methods disclosed herein may also be implemented by an application specific integrated circuit (ASIC), a field-programmable gate array (FPGA), a complex programmable logic device (CPLD), a printed circuit board (PCB), a digital signal processor (DSP), a combination of programmable logic components and programmable interconnects, a single central processing unit (CPU) chip, a CPU chip combined on a motherboard, a general purpose computer, or any other combination of devices or modules capable of performing media format substitution disclosed herein.

[0071] In the preceding specification, the systems and methods have been described with reference to specific exemplary embodiments. It will, however, be evident that various modifications and changes may be made without departing from the broader spirit and scope of the disclosed embodiments as set forth in the claims that follow. The specification and drawings are accordingly to be regarded as illustrative rather than restrictive. Other embodiments may be apparent to those skilled in the art from consideration of the specification and practice of the embodiments disclosed herein.

WHAT IS CLAIMED IS:

1. A system comprising:
 - a database configured to include dependency tree information associated with one or more applications; and
 - a server having one or more processors configured to:
 - access a first dependency tree representing a first application and a second dependency tree;
 - acquire one or more values for the first dependency tree and one or more values for the second dependency tree, wherein one or more values for either the first or second dependency tree represent at least one or more portions of the respective dependency tree;
 - compare the one or more values of the first dependency tree with the one or more values of the second dependency tree;
 - determine a matching status between the first application and an application represented by the second dependency tree based on the comparison; and
 - provide, for display, an indication of the matching status.
2. The system of claim 1, wherein the second dependency tree represents a second application.
3. The system of claim 1 or 2,
 - wherein the one or more values for the first dependency tree include a first value that represents the first dependency tree and the one or more values for the second dependency tree include a second value that represents the second dependency tree, and
 - wherein the comparison between the one or more values of the first dependency tree with the one or more values of the second dependency tree includes a comparison between the first value and the second value.
4. The system of any one of claims 1-3, wherein the one or more values for the first dependency tree include a first value that represents a portion of the first dependency tree and the one or more values for the second dependency tree include a second value that represents a portion of the second dependency tree, and

wherein the comparison between the one or more values of the first dependency tree with the one or more values of the second dependency tree includes a comparison between the first value and the second value.

5. The system of any one of claims 1-4, wherein the one or more portions of the respective dependency tree include at least one of: one or more executable files, one or more library files, and one or more application programming interface (API) calls of the respective application.

6. The system of any one of claims 1-5, wherein the determination of the matching status is based on whether the comparison results in a percentage of matches between the one or more values of the first dependency tree and the one or more values of the second dependency tree that is at least equal to a predetermined threshold.

7. The system of any one of claims 1-6, wherein the application is associated with a hosting system that is different from the first application.

8. The system of any one of claims 1-7, wherein each of the first application and the application is either a new application or a newer version of an existing application.

9. A computer-implemented method comprising:
accessing a first dependency tree representing a first application and a second dependency tree;
acquiring, at a server, one or more values for the first dependency tree and one or more values for the second dependency tree, wherein one or more values for either the first or second dependency tree represent at least one or more portions of the respective dependency tree;
comparing the one or more values of the first dependency tree with the one or more values of the second dependency tree;
determining a matching status between the first application and an application represented by the second dependency tree based on the comparison; and
providing, for display, an indication of the matching status.

10. The method of claim 9, wherein the second dependency tree represents a second application.
11. The method of claim 9 or 10,
wherein the one or more values for the first dependency tree include a first value that represents the first dependency tree and the one or more values for the second dependency tree include a second value that represents the second dependency tree, and
wherein comparing the one or more values of the first dependency tree with the one or more values of the second dependency tree includes comparing the first value with the second value.
12. The method of any one of claims 9-11,
wherein the one or more values for the first dependency tree include a first value that represents a portion of the first dependency tree and the one or more values for the second dependency tree include a second value that represents a portion of the second dependency tree, and
wherein comparing the one or more values of the first dependency tree with the one or more values of the second dependency tree includes comparing the first value with the second value.
13. The method of any one of claims 9-12, wherein the one or more portions of the respective dependency tree include at least one of: one or more executable files, one or more library files, and one or API calls of the respective application.
14. The method of any one of claims 9-13, wherein determining the matching status is based on whether the comparison results in a percentage of matches between the one or more values of the first dependency tree and the one or more values of the second dependency tree that is at least equal to a predetermined threshold.
15. The method of any one of claims 9-14, wherein the application is associated with a hosting system that is different from the first application.

16. The method of any one of claims 9-15, wherein each of the first application and the application is either a new application or a newer version of an existing application.

17. A non-transitory computer-readable medium storing a set of instructions that are executable by one or more processors of one or more servers to cause the one or more servers to perform a method comprising:

accessing a first dependency tree representing a first application and a second dependency tree;

acquiring, at a server, one or more values for the first dependency tree and one or more values for the second dependency tree, wherein one or more values for either the first or second dependency tree are associated with at least one or more portions of the respective dependency tree;

comparing the one or more values of the first dependency tree with the one or more values of the second dependency tree;

determining a matching status between the first application and an application represented by the second dependency tree based on the comparison; and

providing, for display, an indication of the matching status.

18. The computer-readable medium of claim 17, wherein the second dependency tree represents a second application.

19. The computer-readable medium of claim 17 or 18, wherein the one or more values for the first dependency tree include a first value that represents the first dependency tree and the one or more values for the second dependency tree include a second value that represents the second dependency tree, and

wherein comparing the one or more values of the first dependency tree with the one or more values of the second dependency tree includes comparing the first value with the second value.

20. The computer-readable medium of any one of claims 17-19, wherein the one or more values for the first dependency tree include a first value that represents a portion of the first dependency tree and the one or more values for the second

dependency tree include a second value that represents a portion of the second dependency tree, and

wherein comparing the one or more values of the first dependency tree with the one or more values of the second dependency tree includes comparing the first value with the second value.

21. The computer-readable medium of any one of claims 17-20, wherein the one or more portions of the respective dependency tree include at least one of: one or more executable files, one or more library files, and one or more API calls of the respective application.

22. The computer-readable medium of any one of claims 17-21, wherein the determination of the matching status is based on whether the comparison results in a percentage of matches between the one or more values of the first dependency tree and the one or more values of the second dependency tree that is at least equal to a predetermined threshold.

23. The computer-readable medium of any one of claims 17-22, wherein the application is associated with a hosting system that is different from the first application.

24. The computer-readable medium of any one of claims 17-23, wherein each of the first application and the application is either a new application or a newer version of an existing application.




FIG. 1




FIG. 2A




FIG. 2B

300A




FIG. 3A

300B




FIG. 3B




FIG. 4




FIG. 5




FIG. 6




FIG. 7




FIG. 8





FIG. 9

**FIG. 10**

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2015/019766A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F9/44
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, 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	Glenn O ' Donnell: "Smarts Application Discovery Manager 5.0: Accelerating Server/Data Center Consolidations, Application Migrations, and CMDB Projects", 1 January 2007 (2007-01-01), pages 1-22, XP055189108, Retrieved from the Internet: URL: http://www.glennodonnell.com/presentations/2007%20EMC%20World%20-%20ADM.pdf [retrieved on 2015-05-13] pages 5,11 - page 13 -----	1-24



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 application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

13 May 2015

Date of mailing of the international search report

02/06/2015

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Müller, Tobias