

(19) World Intellectual Property Organization
International Bureau



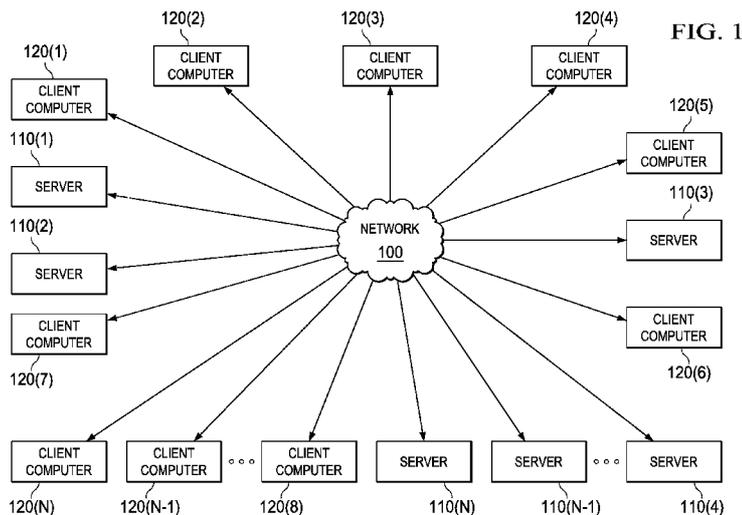
(43) International Publication Date
28 May 2009 (28.05.2009)

PCT

(10) International Publication Number
WO 2009/067707 A2

- (51) **International Patent Classification:**
H04L 12/24 (2006.01) *G06F 17/30* (2006.01)
 - (74) **Agent:** **HITT, David, H.;** Hitt Gaines, PC, P.O. Box 832570, Richardson, TX 75083 (US).
 - (21) **International Application Number:**
PCT/US2008/084429
 - (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, ST, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW
 - (22) **International Filing Date:**
21 November 2008 (21.11.2008)
 - (25) **Filing Language:** English
 - (26) **Publication Language:** English
 - (30) **Priority Data:**
60/989,730 21 November 2007 (21.11.2007) US
 - (71) **Applicant** (*for all designated States except US*): **MO-TIVE, INCORPORATED** [US/US]; 12515 Research Boulevard, Building 5, Austin, TX 78759 (US).
 - (84) **Designated States** (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, 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).
 - (72) **Inventors; and**
 - (75) **Inventors/Applicants** (*for US only*): **PELLEY, Edward, S.** [US/US]; 1904 Nelson Ranch Loop, Cedar Park, TX 78613 (US). **TACKER, Bryan, P.** [US/US]; 15712 Hamden Circle, Austin, TX 78717 (US). **PALKOT, Thomas, A.** [US/US]; 12801 Poquoson Drive, Austin, TX 78727 (US). **SINGH, Amarjit** [US/US]; 8745 Lemens Spice Trail, Austin, TX 78750 (US).
- Published:**
— *without international search report and to be republished upon receipt of that report*

(54) **Title:** NORMALIZATION ENGINE AND METHOD OF REQUESTING A KEY OR PERFORMING AN OPERATION PERTAINING TO AN END POINT



(57) **Abstract:** A normalization engine and a method of requesting a key or performing an operation pertaining to an end point. In one embodiment, the normalization engine includes: (1) a data manager configured to receive a request to execute an operation or obtain a key/value pair pertaining to an end point and determine a normalized version of the key in the key/value pair and (2) a normalization mapping repository configured to contain normalization mechanisms, the data manager further configured to employ the normalized version to obtain one of the normalization mechanisms from the normalization mapping repository, carry out the normalization mechanism and employ a data source corresponding to the end point to retrieve a device-specific value corresponding to the normalized version.

WO 2009/067707 A2

NORMALIZATION ENGINE AND METHOD OF REQUESTING A KEY OR
PERFORMING AN OPERATION PERTAINING TO AN END POINT

CROSS-REFERENCE TO RELATED APPLICATION

[0001] This application claims the benefit of U.S. Provisional Application Serial No. 60/989,730, filed by Dholakia, et al., on November 21, 2007, entitled "Method
5 and System for Remote Device Management," commonly assigned with this application and incorporated herein by reference. This application is also related to the following U.S. patent applications, which are filed on even date herewith, commonly assigned with this
10 application and incorporated herein by reference:

Serial No.	Inventors	Title
[Attorney Docket No. 804144-US-NP]	Dholakia, et. al.	"Service Management System and Method of Operation thereof"
[Attorney Docket No. 804144-US-NP(2)]	Dholakia, et. al.	"System and Method for Identifying and Calling a Function of a Service With Respect to a Subscriber And Service Management System Employing the Same"
[Attorney Docket No. 804144-US-NP(4)]	Pelley, et. al.	"Service Management System and Method of Executing a Policy"
[Attorney Docket No. 804144-US-NP(5)]	Pelley	"System and Method for Generating a Visual Representation of a Service and Service Management System Employing the Same"
[Attorney Docket No. 804144-US-NP(6)]	Pelley, et. al.	"System and Method for Remotely Activating a Service and Service Management System Incorporating the Same"

[Attorney Docket No. 804144-US-NP(7)]	Pelley	"Application and Method for Dynamically Presenting Data Regarding an End Point or a Service and Service Management System Incorporating the Same"
[Attorney Docket No. 804144-US-NP(8)]	Pelley, et. al.	"Service Diagnostic Engine and Method and Service Management System Employing the Same"
[Attorney Docket No. 804144-US-NP(9)]	Pelley	"Self-Service Application for a Service Management System and Method of Operation Thereof"
[Attorney Docket No. 804144-US-NP(10)]	Pelley	"Customer Service Representative Support Application for a Service Management System and Method of Operation Thereof"
[Attorney Docket No. 804144-US-NP(11)]	Pelley, et. al.	"System and Method for Remotely Repairing and Maintaining a Telecommunication Service Using Service Relationships and Service Management System Employing the Same"
[Attorney Docket No. 804144-US-NP(12)]	Pelley, et. al.	"Application and Method for Generating Automated Offers of Service and Service Management System Incorporating the Same"
[Attorney Docket No. 804144-US-NP(13)]	Dholakia, et. al.	"System and Method for Provisioning and Unprovisioning Multiple End Points With Respect to a Subscriber and Service Management System Employing the Same"
[Attorney Docket No. 804144-US-NP(14)]	Dholakia, et. al.	"System and Method for Identifying Functions and Data With Respect to a Service and a Subscriber and Service Management System Employing the Same"

[Attorney Docket No. 804144- US- NP(15)]	Dholakia, et. al.	"System and Method for Invoking a Function of a Service in Response to an Event and Service Management System Employing the Same"
---	----------------------	--

TECHNICAL FIELD

[0002] This application relates to remote management of fixed-line and mobile devices, and, more particularly, to activation, provisioning, support, management and assurance of consumer and business services spanning one or more fixed-line devices and one or more mobile devices .

BACKGROUND

[0003] Network service providers are called upon to support a large variety of networked devices, including devices coupled to home networks [e.g., residential gateways, set-top boxes and voice-over-IP, or VoIP, adapters) and cellular networks (e.gr. smart phones and pocket computers) . Given the proliferation of such devices and the distributed nature of the networks involved, remote management of such devices is highly desirable .

[0004] For example, demand for smart phones and other advanced handsets is growing faster than anticipated as users look for new ways to increase their personal and professional productivity. In 2005, year-over-year growth in the smart phone market exceeded 70%, and the industry experts expect that trend to continue for the next several years. In fact by 2009, it is estimated that smart phones will represent almost 30% of all new handsets sold - up from less than three percent in 2004.

[0005] As smart phones and services for smart phones boom, so do the challenges. Today, the complexity often associated with smart phones is driving customer service costs up and serves as a potential inhibitor as mobile
5 network operators strive to achieve mass-market adoption with these sophisticated devices. In fact, consumers are finding mobile services increasingly confusing and issues around ease-of-use are holding them back from buying and using third generation (3G) handsets and services.

10 [0006] Wireless service providers who sell and support smart phones and their associated data services face the prospect of rising customer support costs due to the complexity associated with these devices and services. In 2007, the support costs for smart phones will surpass
15 that of feature phones. The following are few of the top reasons for this support cost .

- Multiple contacts are made to a helpdesk to solve a single problem.
- 34% of users have never solved a problem with a
20 single contact to the helpdesk.
- Calls last two to three times longer than calls from users of feature phones .
- It is common practice to escalate care from a helpdesk (Tier 1) to expensive technicians (Tier 2 and
25 Tier 3) .
- FMC (Fixed-Mobile Convergence) will add to the support burden. 89% of early adopters are more likely to go to CE vendors for support . Mainstream consumers are three times more likely to look to their service provider
30 for support.

[0007] Similarly, network providers that are coupled to home networks {e.g., Digital Subscriber Link, or DSL,

and cable) find those networks coupled to a variety of customer premises equipment (CPE) within the homes that are gradually becoming more and more sophisticated. Customer issues with such devices are no less taxing upon support staff and support infrastructure.

[0008] The Open Mobile Alliance (OMA) is currently defining a number of standards for managing functionality on mobile devices. These include protocols for device management (OMA-DM), client provisioning (OMA-CP), firmware updates, data synchronization (OMA-DS) and the like. Devices that support at least some of these protocols are becoming prevalent. A support solution that utilizes these protocols and provides a usable console for customer support is the only way network providers and mobile carriers can handle support for the increasing number of devices in the market.

[0009] It is therefore desirable to provide a support solution that allows centralized management and control of remotely networked devices such as smart phones and CPE using protocols established for device management, updates, data synchronization and the like.

SUMMARY

[0010] Various embodiments of a method and system for providing customer support with centralized management and control of mobile phones and customer premises equipment in order to aid users of such equipment with problems related to that equipment. In one embodiment, a user interface driven mechanism is provided to allow customer support representatives to manipulate remote devices in, for example, the following manners: access information about the remote devices and the users thereof, including history of issues with a particular

device, device provisioning, access to diagnostics of a device, ability to upgrade firmware/software of a device, synchronization of data, enablement of security features, remote control of devices, service and application provisioning, defining and following policies related to service management for a variety of devices and resetting devices. Such functionality can be provided, for example, through the use of a device management server that uses a variety of appropriate protocols to communicate with the remote devices.

[0011] Another aspect provides a normalization engine. In one embodiment, the normalization engine includes: (1) a data manager configured to receive a request to execute an operation or obtain a key/value pair pertaining to an end point and determine a normalized version of the key in the key/value pair and (2) a normalization mapping repository configured to contain normalization mechanisms, the data manager further configured to employ the normalized version to obtain one of the normalization mechanisms from the normalization mapping repository, carry out the normalization mechanism and employ a data source corresponding to the end point to retrieve a device-specific value corresponding to the normalized version .

[0012] Yet another aspect provides a method of requesting a key pertaining to an end point. In one embodiment, the method includes: (1) determining a normalized version of the key, (2) employing the normalized version to obtain a normalization mechanism from a normalization mapping repository and (3) carrying out the normalization mechanism, including: (3a) employing an attribute and a role of the end point to

identify a data source and (3b) employing the data source to retrieve a device-specific value corresponding to the normalized version.

BRIEF DESCRIPTION

[0013] Reference is now made to the following descriptions taken in conjunction with the accompanying drawings, in which:

5 [0014] FIG. 1 is a block diagram illustrating a network environment in which commercial transaction processing according to embodiments of the invention may be practiced;

[0015] FIG. 2 is a block diagram illustrating a
10 computer system suitable for implementing embodiments of the invention;

[0016] FIG. 3 is a block diagram illustrating the interconnection of the computer system of FIG. 2 to client and host systems;

15 [0017] FIG. 4 is a diagram illustrating relationships that may exist among a subscriber, a service and various devices and systems;

[0018] FIG. 5 is a diagram of one embodiment of a service description;

20 [0019] FIG. 6 is a diagram illustrating relationships that may exist among management operations, roles, capabilities and attributes;

[0020] FIG. 7 is a high-level block diagram of one embodiment of a service management system;

25 [0021] FIG. 8 is a block diagram of one embodiment of a service normalization block of FIG. 7;

[0022] FIG. 9 is a diagram illustrating examples of different devices having different associated key/value pairs ;

30 [0023] FIG. 10 is a diagram of one embodiment of a device normalization sequence carried out according to the principles of the invention; and

[0024] FIG. 11 is a flow diagram of one embodiment of a method of requesting a value carried out according to the principles of the invention.

DETAILED DESCRIPTION

[0025] The following is intended to provide a detailed description of an example of the invention and should not be taken to be limiting of the invention itself. Rather, 5 any number of variations may fall within the scope of the invention which is defined in the claims following the description. The use of the same reference symbols in different drawings indicates similar or identical items.

[0026] Introduction

10 [0027] Described herein are various embodiments of a management system that allows users to create, define and maintain services by defining the roles of its constituent devices and systems. Certain of the embodiments have the ability to map a given set of 15 devices and systems into roles. The roles may then be used to select key/value pairs, alerts and management function from each device. Certain of the embodiments allow relationships to be specified between roles and between other services. Using the roles and 20 relationships as a lens upon the service's constituent devices, service-wide key/value pairs, alerts and management functions may be created.

[0028] In various embodiments to be described and illustrated herein, a method, apparatus and process are 25 disclosed that allow for activation, provisioning, support (by call center, functions or self), management (by call center, functions or self) and assurance of consumer and business services spanning one or more fixed-line devices and one or more mobile devices, such 30 as PCs, AAA servers, email servers, web servers and devices of every kind. Before describing the embodiments, an example computing and network environment

within which the embodiments may operate will be described.

[0029] An Example Computing and Network Environment

[0030] FIG. 1 is a block diagram illustrating a
5 network environment in which a system according to the invention may be practiced. As is illustrated in FIG. 1, network 100, such as a private wide area network (WAN) or the Internet, includes a number of networked servers HO(I)-(N) that are accessible by client computers
10 120 (1) - (N) .

[0031] Communication between the client computers 120(I)-(N) and the servers HO(I)-(N) typically occurs over a publicly accessible network, such as a public switched telephone network (PSTN) , a DSL connection, a
15 cable modem connection or large bandwidth trunks (e.g., communications channels providing T1 or OC3 service) . The client computers 120(I)-(N) access the servers HO(I)-(N) through, for example, a service provider. This might be, for example, an Internet Service Provider
20 (ISP) such as America On-Line™, Prodigy™, CompuServe™ or the like. Access is typically had by executing application specific software (e.g., network connection software and a browser) on the given one of the client computers 120 (1) - (N) .

25 [0032] One or more of the client computers 120 (1) - (N) and/or one or more of the servers HO(I)-(N) may be, for example, a computer system of any appropriate design, in general, including a mainframe, a mini-computer or a personal computer system. Such a computer system
30 typically includes a system unit having a system processor and associated volatile and non-volatile memory, one or more display monitors and keyboards, one

or more diskette drives, one or more fixed disk storage devices and one or more printers . These computer systems are typically information handling systems which are designed to provide computing power to one or more users, 5 either locally or remotely. Such a computer system may also include one or a plurality of I/O devices (i.e., peripheral devices) which are coupled to the system processor and which perform specialized functions. Examples of I/O devices include modems, sound and video 10 devices and specialized communication devices. Mass storage devices such as hard disks, CD-ROM drives and magneto-optical drives may also be provided, either as an integrated or peripheral device . One such example computer system, discussed in terms of the client 15 computers 120(I)-(N), is shown in detail in FIG. 2 .

[0033] It will be noted that the variable identifier "N" is used in several instances in FIG. 1 to more simply designate the final element (e.gr., the servers HO(I)-(N) and the client computers 120(I)-(N)) of a series of 20 related or similar elements (e.g., servers and client computers) . The repeated use of such variable identifiers is not meant to imply a correlation between the sizes of such series of elements, although such correlation may exist. The use of such variable 25 identifiers does not require that each series of elements has the same number of elements as another series delimited by the same variable identifier. Rather, in each instance of use, the variable identified by "N" may hold the same or a different value than other instances 30 of the same variable identifier.

[0034] FIG. 2 depicts a block diagram of a computer system 210 suitable for implementing the invention and

example of one or more of the client computers 120(1)-(N). A computer system 210 includes a bus 212 which interconnects major subsystems of the computer system 210 such as a central processor 214, a system memory 216
5 (typically random-access memory, or RAM, but which may also include read-only memory, or ROM, flash RAM, or the like), an input/output controller 218, an external audio device such as a speaker system 220 via an audio output interface 222, an external device such as a display
10 screen 224 via a display adapter 226, serial ports 228, 230, a keyboard 232 (interfaced with a keyboard controller 233), a storage interface 234, a floppy disk drive 236 operative to receive a floppy disk 238 and a CD-ROM drive 240 operative to receive a CD-ROM 242. Also
15 included are a mouse 246 (or other point-and-click device, coupled to the bus 212 via the serial port 228), a modem 247 (coupled to the bus 212 via the serial port 230) and a network interface 248 (coupled directly to the bus 212).

20 [0035] The bus 212 allows data communication between a central processor 214 and a system memory 216, which may include RAM, ROM or flash memory, as previously noted. The RAM is generally the main memory into which the operating system and application programs are loaded and
25 typically affords at least 16 megabytes of memory space. The ROM or flash memory may contain, among other code, the Basic Input-Output system (BIOS) which controls basic hardware operation such as the interaction with peripheral components. Applications resident with the
30 computer system 210 are generally stored on and accessed via a computer readable medium, such as a hard disk drive {e.g., fixed disk 244}, an optical drive {e.g., CD-ROM

drive 240), a floppy disk unit 236 or other storage medium. Additionally, applications may be in the form of electronic signals modulated in accordance with the application and data communication technology when
5 accessed via a network modem 247 or an interface 248.

[0036] The storage interface 234, as with the other storage interfaces of the computer system 210, may connect to a standard computer readable medium for storage and/or retrieval of information, such as a fixed
10 disk drive 244. The fixed disk drive 244 may be a part of computer system 210 or may be separate and accessed through other interface systems. Many other devices can be connected, such as a mouse 246 connected to the bus 212 via serial port 228, a modem 247 connected to the bus
15 212 via serial port 230 and a network interface 248 connected directly to the bus 212. The modem 247 may provide a direct connection to a remote server via a telephone link or to the Internet via an internet service provider (ISP). The network interface 248 may provide a
20 direct connection to a remote server via a direct network link to the Internet via a POP (point of presence). The network interface 248 may provide such connection using wireless techniques, including digital cellular telephone connection, Cellular Digital Packet Data (CDPD)
25 connection, digital satellite data connection or the like.

[0037] Many other devices or subsystems (not shown) may be connected in a similar manner (e.g., bar code readers, document scanners, digital cameras and so on).

30 [0038] Conversely, it is not necessary for all of the devices shown in FIG. 2 to be present to practice the invention. The devices and subsystems may be

interconnected in different ways from that shown in FIG. 2. The operation of a computer system such as that shown in FIG. 2 is readily known in the art and is not discussed in detail in this application. Code to implement the invention may be stored in computer-readable storage media such as one or more of the system memory 216, the fixed disk 244, the CD-ROM 242 or the floppy disk 238. Additionally, the computer system 210 may be any kind of computing device and so includes personal data assistants (PDAs), network appliance, X-window terminal or other such computing device. The operating system provided on the computer system 210 may be MS-DOS®, MS-Windows®, OS/2®, UNIX®, Linux® or other known operating system. The computer system 210 may also support a number of Internet access tools, including, for example, a Hypertext Transfer Protocol (HTTP)-compliant web browser having a JavaScript interpreter, such as Netscape Navigator® 3.0, Microsoft Explorer® 3.0 and the like.

[0039] The foregoing described embodiment wherein the different components are contained within different other components (e.g., the various elements shown as components of the computer system 210). It is to be understood that such depicted architectures are merely examples, and that in fact many other architectures can be implemented which achieve the same functionality. In an abstract, but still definite sense, any arrangement of components to achieve the same functionality is effectively "associated" such that the desired functionality is achieved. Hence, any two components herein combined to achieve a particular functionality can be seen as "associated with" each other such that the

desired functionality is achieved, irrespective of architectures or intermediate components. Likewise, any two components so associated can also be viewed as being "operably connected," or "operably coupled," to each other to achieve the desired functionality.

[0040] FIG. 3 is a block diagram depicting a network 300 in which the computer system 210 is coupled to an internetwork 310, which is coupled, in turn, to client systems 320, 330, as well as a server 340. An internetwork 310 (e.g., the Internet or a wide-area network, or WAN) is also capable of coupling the client systems 320, 330 and the server 340 to one another. With reference to the computer system 210, the modem 247, the network interface 248 or some other method can be used to provide connectivity from the computer system 210 to the internetwork 310. The computer system 210, the client system 320 and the client system 330 are able to access information on the server 340 using, for example, a web browser (not shown). Such a web browser allows the computer system 210, as well as the client systems 320, 330, to access data on the server 340 representing the pages of a website hosted on the server 340. Protocols for exchanging data via the Internet are well known to those skilled in the art. Although FIG. 3 depicts the use of the Internet for exchanging data, the invention is not limited to the Internet or any particular network-based environment.

[0041] Referring to FIGs. 1, 2 and 3, a browser running on the computer system 210 employs a TCP/IP connection to pass a request to the server 340, which can run an HTTP "service" (e.g., under the WINDOWS® operating system) or a "daemon" (e.g., under the UNIX® operating

system) , for example. Such a request can be processed, for example, by contacting an HTTP server employing a protocol that can be used to communicate between the HTTP server and the client computer. The HTTP server then
5 responds to the protocol, typically by sending a "web page" formatted as an HTML file. The browser interprets the HTML file and may form a visual representation of the same using local resources {e.g., fonts and colors} .

[0042] Example Embodiments of a Service Management
10 System

[0043] The functions referred to herein may be modules or portions of modules (e.g., software, firmware or hardware modules) . For example, although the described embodiment includes software modules and/or includes
15 manually entered user commands, the various example modules may be application specific hardware modules. The software modules discussed herein may include script, batch or other executable files, or combinations and/or portions of such files. The software modules may include
20 a computer program or subroutines thereof encoded on computer-readable media.

[0044] Additionally, those skilled in the art will recognize that the boundaries between modules are merely illustrative and alternative embodiments may merge
25 modules or impose an alternative decomposition of functionality of modules. For example, the modules discussed herein may be decomposed into sub-modules to be executed as multiple computer processes and, optionally, on multiple computers. Moreover, alternative embodiments
30 may combine multiple instances of a particular module or sub-module. Furthermore, those skilled in the art will recognize that the functions described in example

embodiment are for illustration only. Operations may be combined or the functionality of the functions may be distributed in additional functions in accordance with the invention.

5 [0045] Alternatively, such actions may be embodied in the structure of circuitry that implements such functionality, such as the micro-code of a complex instruction set computer (CISC), firmware programmed into programmable or erasable/programmable devices, the
10 configuration of a field-programmable gate array (FPGA), the design of a gate array or full-custom application-specific integrated circuit (ASIC), or the like.

[0046] Each of the blocks of the flow diagram may be executed by a module (e.g., a software module) or a
15 portion of a module or a computer system user using, for example, a computer system such as the computer system 210. Thus, the above described method, the functions thereof and modules therefore may be executed on a computer system configured to execute the functions of
20 the method and/or may be executed from computer-readable media. The method may be embodied in a machine-readable and/or computer-readable medium for configuring a computer system to execute the method. Thus, the software modules may be stored within and/or transmitted
25 to a computer system memory to configure the computer system to perform the functions of the module.

[0047] Such a computer system normally processes information according to a program (a list of internally stored instructions such as a particular application
30 program and/or an operating system) and produces resultant output information via I/O devices. A computer process typically includes an executing (running) program

or portion of a program, current program values and state information and the resources used by the operating system to manage the execution of the process. A parent process may spawn other, child processes to help perform
5 the overall functionality of the parent process. Because the parent process specifically spawns the child processes to perform a portion of the overall functionality of the parent process, the functions performed by child processes (and grandchild processes,
10 etc.) may sometimes be described as being performed by the parent process.

[0048] Such a computer system typically includes multiple computer processes executing "concurrently." Often, a computer system includes a single processing
15 unit which is capable of supporting many active processes alternately. Although multiple processes may appear to be executing concurrently, at any given point in time only one process is actually executed by the single processing unit. By rapidly changing the process
20 executing, a computer system gives the appearance of concurrent process execution. The ability of a computer system to multiplex the computer system's resources among multiple processes in various stages of execution is called multitasking. Systems with multiple processing
25 units, which by definition can support true concurrent processing, are called multiprocessing systems. Active processes are often referred to as executing concurrently when such processes are executed in a multitasking and/or a multiprocessing environment.

30 [0049] The software modules described herein may be received by such a computer system, for example, from computer readable media. The computer readable media may

be permanently, removably or remotely coupled to the computer system. The computer readable media may non-exclusively include, for example, any number of the following: magnetic storage media including disk and tape
5 storage media, optical storage media such as compact disk media (e.g., CD-ROM, CD-R, etc.) and digital video disk storage media, nonvolatile memory storage memory including semiconductor-based memory units such as flash memory, EEPROM, EPROM, ROM or application-specific
10 integrated circuits (ASICs), volatile storage media including registers, buffers or caches, main memory, RAM and the like, and data transmission media including computer network, point-to-point telecommunication and carrier wave transmission media. In a UNIX-based
15 embodiment, the software modules may be embodied in a file which may be a device, a terminal, a local or remote file, a socket, a network connection, a signal, or other expedient of communication or state change. Other new and various types of computer-readable media may be used
20 to store and/or transmit the software modules discussed herein.

[0050] Before describing various embodiments of management systems constructed according to the principles of the invention, some use cases or
25 interactions will be described that provide a framework for understanding the management systems. Various of the embodiments of the management systems are directed to addressing the following categories of use cases or interactions: service activation, service management,
30 service interruption and resumption and service offering. Service activation refers to all the use cases that involve creating (provisioning) and deleting

(unprovisioning) a new service instance, or "subscription." Service management refers to day-to-day management tasks involving a given service or subscription. Service interruption and resumption may be
5 thought of as special types of service management that involve the loss and restoration of service. Service offering refers to new services that may be offered to a subscriber .

[0051] The categories of use cases that set forth
10 above may be illustrated with reference to FIG. 4. FIG. 4 is a diagram illustrating relationships that may exist among a subscriber 410, a service 420 and various devices and systems 430. A service provider, such as a cellular phone company, an Internet service provider, a cable
15 television company or a combination of these, offers one or more services to subscribers that involve devices and systems such as cell phones, set-top boxes, routers, cell towers, email servers, DSLAMs, landline phones and other mobile and customer premises equipment and network
20 infrastructure. In the context of FIG. 4, the subscriber 410 takes out a subscription 440 to a service 420 offered by a service provider. The subscription contains state and other information that is both relative to the subscriber 410 and the service 420. The subscription
25 calls for one or more associations 450 to be made between the subscriber 410 and various devices and systems 430. The subscriber 410 may own or lease one or more devices 430. The subscriber 410 may also be associated with one or more systems 430. Once the associations 450 are made,
30 the devices and systems 430 assume roles 460 in delivering the service 420 to the subscriber 410. The

roles describe to the service management system how the device should be managed.

[0052] FIG. 4 may be employed to illustrate two example use cases: activating a device for a subscriber
5 and managing and provisioning a subscription.

[0053] To activate a device (one of the devices and systems 430) for a subscriber 410, the following steps may be taken. First, for a given device 410, the subscriber 410 associated with that device 430 is found.
10 This is done by employing the associations 450. Once the associated subscriber 410 has been identified, a corresponding subscription 440 may then be used to determine the service or services 420 that should be provisioned on the device 430. For each service 420 that
15 needs to be activated on the device 430, two alternative actions may be taken. Based upon the role the device 430 plays with respect to the service 420, settings on the device 430 may be set to provision the device. Alternatively or additionally, based on the roles of
20 other devices and systems 430 relative to the service 420, settings on the other devices and systems may be set to provision them for the new device's presence.

[0054] Managing and provisioning a subscription involves either adding a new service to a subscriber or
25 managing an existing service. To manage and provision a subscription 440 for a subscriber 410, the following steps may be taken. First, a service 420 is added to a subscriber 410, or an existing service 420 is modified. Devices and systems 430 associated with the subscriber
30 are collected. Then, each device associated with the service 420 is mapped into the different roles 460 in the targeted service. This reveals what actions should be

taken with respect to each device or system to provision the service 420.

[0055] Another use case that is a variant on the one above is a bulk change to an existing service for all subscribers. In this use case, existing subscriptions are retrieved to obtain a list of subscribers. Then, a list of devices and systems 430 is assembled for each subscriber. Using roles, changes are then applied to the devices and systems 430 to enable the bulk change.

10 [0056] Having described various use cases, one manner in which interaction may occur among roles, devices and systems, and the service level management interfaces will now be described. FIG. 5 is a diagram of one embodiment of a service description 500. FIG. 5 shows that the service description 500 includes service alerts 505, functions 510 and key/value pairs 515. Roles are associated with the service alerts 505, the functions 510 and the key/value pairs 515. The roles, designated Role A 520, Role B 525 and Role C 530 are associated as shown with the service alerts 505, the functions 510 and the key/value pairs 515 as various arrows show. Meta data 535 is also associated with the key/value pairs 515. Devices or systems 540, 545 are associated with the roles 520, 525, 535 as various arrows show.

25 [0057] FIG. 5 illustrates, in the context of the service description 500, how a set of roles (e.g., the roles 520, 525, 535) can be mapped onto a set of devices and systems (e.g., the devices or systems 540, 545). The roles define the functions, alerts and functions that are of interest from each device or system. In the illustrated embodiment, the meta data 535 contains both service-wide and subscriber/subscription data that is

specific to the service description. Items like level of service and current state of activation would be a part of the meta data. The alerts, functions, and key/value pairs 505, 510, 515 together constitute services 5 supported by the devices and systems 540, 545 exposed through the roles 520, 525, 530. The service description 500 is configured to contain arbitrary, named relationships that can be used to affect a service. The service description 500 may also be configured to contain 10 one or more references to other service with which it has relationships or upon which it has dependencies. Accordingly, a further service description 550 is associated with the meta data 535 as an arrow shows. Similar to the relationships between roles, relationships 15 between services are exposed to the logic in the service description and to external logic associated with the service description.

[0058] To make a role useful, the role is matched with a device. FIG. 6 is a diagram illustrating relationships 20 that may exist among management functions (i.e., values, alerts and functions 505, 510, 515), roles (e.g., 460), capabilities and attributes 610 and devices and systems (e.g., 430). In the embodiment of FIG. 6, the mechanism for doing this is twofold. First, a role may be matched 25 to a device or a system based on the known attributes of the device or system. Second, a role may be matched to a device or a system based upon the known capabilities of the device or system.

[0059] Device attributes are known aspects of the 30 device, e.g., the type, serial number, MAC address, manufacture date, make, model, service tags, device ID or operating system of the device or system. Other

attributes may include the firmware version, hardware version, embedded device, locale (language), and physical location. Device attributes, in the simplest form, could be a list of known key/value pairs associated to the
5 device.

[0060] Capabilities are similar to device attributes. In this case, instead of a list of key/value pairs, these are a list of values (without the key) of known capabilities about the device, e.g., generic email
10 client, Microsoft Outlook® email client, phone, router or IPTV device. Other examples include network attached storage, media server, media renderer, camera, MMS client, SMS client, wireless access provider, wireless access client, printer, GPS, vibrate, Bluetooth, USB, Wi-
15 Fi, clock, browser, QVGA, flight mode, caller ID, touchscreen, or fax.

[0061] Both the capabilities and the attributes can be provided to or retrieved from an external system, deduced or derived from known attributes and capabilities,
20 queried directly from the device or system or a combination of these. For example, prior knowledge may exist that any device that has a serial number from a given manufacture starting with the letter W- has built-in Wi-Pi capabilities, or that Windows® Mobile phone
25 supports OMA-DM.

[0062] It can be determined whether a given device or system matches a role by matching its attributes and capabilities (derived, discovered, or known) against the required attributes and capabilities of a given role.
30 Each role defines a set of key/value pairs, alerts, and functions that are relevant for devices of that role in the service description.

[0063] It should be noted that roles do not imply device type, model or brand. Direct mappings between devices and roles may occur in practice, but the mappings are flexible such that they may change as device
5 attributes or capabilities change. For example, newer devices may support more roles than do older devices. One example of a role is a phone capable of functioning as an email client may play an "EmailClient" role in a service description associated with an email service.
10 Other roles in an email service may include "SMTPServer," "POPServer" and "IMAPServer." Roles in a data connectivity service may include "Host," "Router," "Wireless Access Point," "Head End," "Border Gateway" and "Authentication, Authorization, Account Server."

15 [0064] FIG. 7 is a high-level block diagram of one embodiment of a service management system. The service management system includes a service normalization block
705 that will be described in greater detail in conjunction with FIG. 8 below. The service normalization
20 block interacts with an optimal settings block 710, a device (and/or system) normalization block 715, a diagnostic engine block 720 and a content repository block 725. A segmentation block 730 and a knowledge base
735 interact with the optimal settings block 710, the
25 device normalization block 715, the diagnostic engine block 720 and the content repository block 725 as shown.

[0065] The service normalization block 705 employs an application programming interface (API), allowing it to exchange information with an interactive voice response
30 (IVR) system 745, a console 750 for a customer service representative (CSR), a Self-Service Management (SSM)

application module 755 and other applications 760, 765 as may be found advantageous in a particular environment .

[0066] The optimal settings block 710 is a repository of predefined known good values used for the purposes of
5 comparing key/value pairs to determine diagnostic and state information. The key/value pairs are also used during provisioning to set up a system or a device. The optimal settings block 710 may be regarded as a configuration repository that contains meta data about
10 the configuration so that applications and other systems can look up known good values for the purpose of configuring (provisioning), diagnostic, and repair. The illustrated embodiment of the optimal settings block 710 is configured to define optimal values for any given
15 key/value pair based on the context of the device, subscriber, customer, or any other segmentation scheme that could be used to define different values for the same attribute (key) . These values may be used by both the script engine 830 and directly by the service management
20 engine 805 to determine whether or not a given key/value pair is "optimal." Optimal values may fall into three categories: (1) values predefined in the context of the service as being correct, (2) values defined based by a call to an extrinsic system or by subscriber input as
25 being correct and (3) absolute values (which is often built into the logic of the script or service description logic and not stored externally) . An example of a predefined optimal value is a POP server. The subscriber is aware of its identity, and it is the same for all
30 subscribers. An example of an absolute value is "connectivity = good." An example of a value defined by a subscriber as being correct is a password, something

that the subscriber chooses and is not defined before the subscriber chooses it.

[0067] The device normalization block 715 is configured to map normalized key/value pairs to device-specific or system-specific key value pairs. Mapping may
5 be performed by transformation, executing a script, or undertaking any other appropriate normalization mechanism.

[0068] The diagnostic engine 720 is configured to
10 contain diagnostic rules and cause diagnostic rules to be executed in order to identify, characterize and present potential solutions to problems that may exist with devices or systems. The content repository is configured to provide a channel-independent mechanism for
15 associating bearer (e.g., IVR voice flows, self-service portal web content and customer service articles) to diagnose a problem.

[0069] The segmentation block 730 and knowledge base 735 likewise employ a data sources abstraction layer 770,
20 allowing it to communicate with systems 775, 790 and provisioning servers and device managers 780, 785.

[0070] Different subscribers subscribe to different levels of service and live in different locations and under different circumstances. The segmentation block
25 730 is configured to enable other portions of the service management system to tailor responses to a subscriber based on his level of service, location and/or circumstance .

[0071] The knowledge base 735 is configured to contain
30 articles associated with known device, system and/or service problems. When the diagnostic engine 720 identifies a problem area or a specified problem, it may

provide articles from the knowledge based 735 to an application to allow the application to provide the article m turn to a subscriber or other user for informational purposes.

5 [0072] The data sources abstraction layer 770 is configured to operate as a protocol implementation and adaptation layer, allowing generic logic to interact with specific devices and systems without having to employ a device-specific or system-specific protocol.

10 [0073] The systems 775, 790 are typically added by a particular service provider and interact with the service management system as needed. The provisioning servers and device managers 780, 785 support various devices 795 intended for use by subscribers. In the illustrated
15 embodiment, the provisioning servers and device managers 780, 785 are management systems that manage large groups of devices that typically share the same protocol {e.g., a mobile device manager that manages 10 million phones using the OMA-DM protocol). The devices 795 are just
20 CPE, such as phones and routers.

[0074] FIG. 8 is a block diagram of one embodiment of the service normalization block 705. The API 740 provides a mechanism by which the service normalization block 705 may be called from within and without the
25 service management system. By means of the API 740, subscribers may have services added (provisioned), removed (unprovisioned), modified or otherwise managed. Devices of subscribers may also be managed m the context of a particular service. Management access to key/value
30 pairs of constituent devices and systems may also be dynamically determined and managed based on their roles m providing particular services.

[0075] A service management engine 805 enables a service provider to implement and manage services by means of the service normalization block 705 according to the various use cases described above. The illustrated
5 embodiment of the service management engine 805 functions in two primary ways. First, the service management engine 805 manages functions defined by service descriptions. Second, the service management engine 805 provides a dynamic view of a given service.

10 [0076] The management of functions allows service descriptions to define named functions that can be called with contextual data derived from analyzing constituent devices and systems m their associated roles of a service .

15 [0077] The provision of a dynamic view of a given service enables a service description to associate key/value pairs (data) with different roles and dynamically to gather data from the devices and systems so that data can be presented without the need for
20 intrinsic knowledge of the data that is being collected. For example, a service-view dashboard capable of creating a map of devices with their associated data of interest (each categorized by their role m the service) may employ dynamic views of services. In the illustrated
25 embodiment, the data itself is self-describing and typically presented m tabular form.

[0078] In an alternative embodiment, the service management engine 805 is also capable of providing a view of a given service, m which case the application does
30 have prior intrinsic knowledge regarding the data that is being collected.

[0079] The first step in managing the service is to collect a list of the devices and systems that are associated to the subscriber of the service. A device repository 835 serves this purpose. In the illustrated
5 embodiment, the device repository is external to the service normalization block 705

[0080] The service management engine 805 employs a capabilities repository to obtain an expanded view on the capabilities of a device so that it may map it into a
10 role. Often the only pieces of information obtained from extrinsic systems are the unique identifier of the device, *e.g.*, its make and model. A device may be viewed as a list of attributes (*e.g.*, further key/value pairs). From those attributes, the capabilities of the device may
15 be extrapolated. Extrapolation may involve expanding the known attributes of a device by deriving new attributes of a device based on the original attributes. For example, once the make and model of a system or device is obtained by means of a query, built-in rules may then be
20 able to determine whether or not it has Wi-Fi capability.

[0081] A service description repository contains service descriptions. As described above, a service description includes at least some of: functions, key/value pairs, alerts, roles (along with their
25 associated key/value pairs, alerts and actions) and relationships. Functions, at the service description level, can be actions exposed by a device, a script that can be executed, or a process (a series of scripts executed in a state engine).

30 [0082] A script engine 830 is configured to execute service-level functions. As described previously, a service-level function could be a script, a process or

action (a series of scripts) or any other type of computer program. In the illustrated embodiment, the service management engine 805 retrieves a named script from a service description based on an event or request
5 from a consumer of the service and passes it, along with a set of parameters, to the script engine 805 for execution. In the illustrated embodiment, the set of parameters includes: references to the constituent devices (categorized by role) and the service
10 description. The script, once started, has access to the optimal values, the devices and systems (abstracted by device normalization or directly) and the service management system in general .

[0083] The service normalization block 705 has access
15 to a capabilities repository 820. The capabilities repository 820 is configured to derive new attributes and capabilities through rules based on existing, known attributes. For example, it is known that Windows® Mobile cell phones have Internet browsers.

[0084] The service normalization block 705 employs the
20 device normalization engine 715 configured to create an abstraction that provides a normal view of extrinsic device and systems. This allows the service description to be defined generically for devices and systems of the
25 same class without having to include logic and cases for each device. For example, if one device is managed by OMA-DM and has an email client while another device that also has an email client is managed by Digital Subscriber Line (DSL) Forum standard TR069, both devices will have a
30 Simple Mail Transfer Protocol (SMTP) server. However, the manner in which values are obtained can differ by

protocol (OMA-DM versus TR069) and by key (the name for the value)

[0085] At least some of the systems and devices 775, 795 have the capability to generate alerts or events
5 The alert/event engine 815 is configured to receive these alerts or events and apply them against each service description to determine whether the alert applies to that service description. If a particular alert or event does apply to a particular service, the service
10 management system is configured to obtain a corresponding action, script or process from the service description that may be executed to respond to the alert or event.

[0086] Systems and Methods for Data Normalization

[0087] Data that characterizes one device rarely
15 matches the data that characterizes another device, though many similarities may exist. Data normalization, referred to above, is the abstracting away of differences that exist within a group of heterogeneous device types such that a common management mechanism may be applied to
20 allow solutions and applications to span the devices. Data retrieved from a device comes in the form of key/value pairs. The key/value pairs that characterize one device do not necessarily match the key/value pairs that characterize another device. In one embodiment,
25 data normalization occurs with respect to multiple types of data: keys, values, alerts, operations and capabilities across device types, bearers and/or protocols that allow solutions and applications to handle or treat devices in a consistent manner for the purpose
30 of service management.

[0088] FIG. 9 is a diagram illustrating examples of different devices having different associated key/value

pairs. In the example of FIG 9, email accounts on Phone A 905 only have one username and password for both receiving and sending. However, email accounts on Phone B 910 have separate receiving and sending usernames and passwords. The key/value pairs returned for the Phone A 905 and the Phone B 910 are different, because they reflect the different numbers of usernames and passwords. Data normalization carried out according to the principles of the invention makes the Phone A 905 and the Phone B 910 appear the same to an application. Through data normalization, the subject involves email user names and password, but the actual implementation as between the Phone A 905 and the Phone B 910 differs substantially .

15 [0089] Operation Normalization

[0090] An operation is a complex action that changes the state of a device beyond the getting and setting of data. An operation may be a factory reset, a reboot, a remote lock, a remote wipe, or a connect/disconnect of a Point-to-Point Protocol over Ethernet (PPPoE) session. Operations themselves may be triggered by the setting of a key/value pair, or there may be other means of triggering an action on a given end point.

[0091] Operation normalization allows an application to execute an operation with respect to a device without needing specific knowledge about the device. The application does not need to know how the operation is carried out, what keys need to be modified or what new values need to be assigned to carry out the operation.

30 [0092] Protocol Normalization

[0093] A protocol is a communication vehicle used to exchange data (key/value pairs) with an end point. In

addition to defining things such as the reading and writing of data, protocols can also define triggers and alerts on specific data. Typically, an end point manager manages end points. End point managers support at least
5 one management protocol so that communication with the end point can take place. Normalization of the protocol allows the consuming application to manage the end point without knowing which end point manager or protocol is facilitating the communication. Normalization happens
10 independent of how the data is obtained.

[0094] It is, in general, more advantageous to normalize as close to the end point as possible so that as many layers can deal with the generic as opposed to the specific. However, there are cases where
15 normalization can occur at higher layers and be beneficial in interoperating with different end point managers .

[0095] Challenges for Normalization

[0096] Normalization faces certain challenges. Any
20 given key/value pair that describe the same thing can be different in both the key and the value. Further, often complex relationships and meta-rules differ between end points key/value pairs. Normalization may be carried out as a normalization sequence in a normalization layer that
25 exists between an application and the physical systems and devices that are being normalized.

[0097] To normalize a given key/value pair fully, one embodiment of a normalization layer is configured to normalize the key, the value and any meta-rules and
30 relationships that the key/value pair has with other key/value pairs. Differences in keys may be as simple as using a different name to represent the same value. In

simple scenarios, a one-to-one mapping file is sufficient to normalize the key. However, other, more complex scenarios exist in which different numbers of keys represent the same concept. Often values also differ.

5 In scenarios where the relationships between keys are more complex than a simple key translation, one embodiment of a normalization layer is configured to employ complex logic to manage the translation.

[0098] Values may be implemented differently by using
10 different units of measurement, different data types, or different conceptual ways of modeling concepts. Sometimes the differences can be managed by a simple mapping or type conversion, but other times the mapping requires complex logic to reflect values separately.

15 Sometimes a clean or clear translation is unavailable. One embodiment of a normalization layer is configured to employ logic to choose a best answer.

[0099] In addition to the above-described challenges, one embodiment of a normalization layer is configured to
20 filter a group of data based on the values of the key/value pairs. An example is found in email for a Windows Mobile® device. When retrieving the email accounts for a Windows Mobile® device, the Multimedia Messaging Service (MMS) account is included in the list.

25 The MMS email account should be excluded from the returned accounts based on the name. One embodiment of a normalization layer is configured to remove this account and not force the application to do additional filtering. To accomplish this, the embodiment of the normalization
30 layer is configured to understand that the key/value pairs returned from a device have additional concepts and information associated with them beyond what the device

description documents (e.g., Device Description Framework, or DDF, documents, or Management Information Bases, or MIBs) and the key/value pairs themselves convey. More specifically, the embodiment of the
5 normalization layer is configured to understand that for a Windows Mobile® device the key/value pairs will sometimes be representing email accounts and sometimes MMS accounts.

[00100] Another challenge occurs where the keys for the
10 same key/value pair differ among end points. As an example, the email display name is found in tree node SERVICENAME on Windows Mobile® and in NAME on the E61. This problem is one of the more simple problems that is encountered in end point normalization and can be
15 typically handled via a simple mapping file.

[00101] Further, not all devices support all the same key/value pairs, either literally or conceptually. An example of this is in Windows Mobile®, where no key/value pair represents a sending username. In such scenarios,
20 one embodiment of a normalization layer is configured to describe the supported key/value pairs of the specific end point and make this available to the solution. It is the application's responsibility to query the normalization layer to determine which key/value pairs on
25 a given end point are available. One embodiment of a normalization layer is configured to consider the total set of known key/value pairs. Further the contracts defined in the normalization layer should be well documented to applications so that ultimate consumers
30 (professional services) are aware of the requirements when building, extending and modifying applications.

[00102] End points will often differ in the way that specific values in the same conceptual key/value pair are implemented. These values can be different in the valid values (e.g., ranges and enumerations) and the types.

5 Windows Mobile® returns 1 for true. Type conversion can typically be handled through simple mapping files, but valid values are often more complex. One embodiment of a normalization layer is configured to allow applications to query to determine what values are available for a

10 given type. This is often complex and should be done dynamically based on the end point which is being managed.

[00103] The relationships between the allowed values of a given key/value pair and the values of other key/value

15 pairs are often complex. In these scenarios, the allowed values for a given key/value pair are determined based on the values of other key/value pairs in the end point . One embodiment of a normalization layer is configured logically to determine the value of a key/value pair

20 during a set operation using values of other key/value pairs. Windows Mobile® WiPi authentication values have to be logically determined.

[00104] Often key/value pairs set functions have an implied order of operation. One embodiment of a

25 normalization layer is configured to break a set operation into dependent logical operations based upon other key/value pairs. On a Windows Mobile® device, the SSID of the WiFi connection is part of the tree node. In order to edit the SSID, a new tree node should be created

30 and the old one deleted. These operations may or may not be identifiable in the device description documents (e.g., DDFs and MIBs) . As an example, one device may

require a reboot to make use of the key/value pairs while another may make use of them immediately

[00105] Key/value pairs often have complex relationships beyond mapping values based on each other and logical orders of operations. Some end points key/value pairs should be set together in atomic operations. As an example, Windows Mobile® requires security mode be set when encryption mode is set or the operation will fail. One embodiment of a normalization layer is configured to allow an application to depend on the normalization layer to get the required key/value pairs for a set operation independent of what was in the applications initial query. To continue the Windows Mobile® Security mode/encryption mode example, if the application asks for the security mode, one embodiment of a normalization layer is configured to query the encryption mode so that, if applications modify and set the security mode, the normalization layer is also able to set the encryption mode.

[00106] End points often have key/value pairs that reference other key/value pairs on the same end point. For example, MMS contains a reference to the NAP profile that is needed to complete the data to be presented to the CSR. At the very highest level, the common end point normalization engine (or normalization Engine) defines a set of logical key/value pairs and exposes them to the consuming technology through an API that allows the consumer to get, set or query meta data about the normalized key/value pairs. One embodiment of a normalization layer is configured to present a second API that can be customized to integrate into existing systems .

[00107] FIG. 10 is a diagram of one embodiment of a device normalization sequence carried out according to the principles of the invention. The normalization engine 715 of FIG. 7 is illustrated as containing a data manager 1005, a normalization mapping repository 1010 and a plurality of normalization mechanisms: mechanism A 1015, mechanism B 1020 and mechanism C 1025. Mechanism A 1015 could be, for example, the script engine 830 of FIG. 8. Mechanism B 1020 could be, for example, a transformation engine. However, those skilled in the pertinent art should understand that these and other normalization mechanisms may be employed in alternative embodiments. Other and further conventional and later-developed normalization mechanisms fall within the scope of the invention. The normalization engine 715 is also illustrated as containing a plurality of data sources 1030, 1035, 1040. The data source 1030 could be, for example a Mobile Device Manager (MDM) data source. The data source 1035 could be a Windows Mobile Tunneling Server (WMTS) data source 1035. The data source 1040 could be a Motive Common Client Infrastructure (MCCI) data source or a Home Device Manager (HDM) data source. Other and further conventional and later-developed data sources fall within the scope of the invention. The data sources correspond to a plurality of end point managers 1045, 1050, 1055. The end point managers 1045, 1050, 1055 correspond to different types of end points (i.e., devices and systems of various conventional or later-developed make and model).

[00108] The illustrated embodiment of the device normalization sequence begins in a start step 1, in which an application sends a request (calls) the normalization

engine 715 to get a list of normalized key/value pairs or execute an operation. In one embodiment, the request is asynchronous. In an alternative embodiment, the request is synchronous. In one embodiment, the call contains a list of requested keys or the operation to execute (in a normalized format). In another embodiment, the call also contains a list of attributes associated with the endpoint (including its make and manufacture).

[00109] The illustrated embodiment continues in a step 2 m which the data manager 1005 handles the request. If the operation is asynchronous, the data manager 1005 manages the jobs needed to fulfill the request and generate the response. In the illustrated embodiment, the data manager uses the normalization mapping repository 1010 to resolve the request. The normalization mapping repository 1010 returns a normalization {e.g., an execution script or a transformation} that the data manager 1005 can then employ to normalize the request and gather data response thereto. In the illustrated embodiment, the normalization mapping repository 1010 uses attributes to determine the normalization to return to the data manager 1005. The normalization mapping repository 1010 of FIG. 10 is structured as hierarchically related meta data points that can inherit definitions to reach a given normalization. The data manager 1005 evaluates the normalization returned to resolve the normalized response needed for key/value pairs or operation (s) requested.

[00110] The illustrated embodiment continues in a step 3 m which the normalization executes in the appropriate normalization mechanism 1015, 1020, 1025. In the illustrated embodiment, the normalization mechanism 1015,

1020, 1025 evaluates what is required to retrieve the needed data. More specifically, the normalization mechanism 1015, 1020, 1025 has implicit knowledge of the end point key/value pairs and how to map them to the
5 normalized key/value pairs for data retrieval or operation execution. The normalization mechanism 1015, 1020, 1025 also handles any complex conversions needed to fulfill the request.

[00111] The illustrated embodiment continues in a step
10 4 in which the normalization mechanism 1015, 1020, 1025 makes a call to the relevant data source or sources 1030, 1035, 1040 to communicate with endpoint managers 1045, 1050, 1055 or other backoffice systems (not shown). In the illustrated embodiment, the data sources 1030, 1035,
15 1040 are implemented with a common API (not shown) to allow them to be exchanged with little or no changes to above layers. The API provides mechanisms for reading and writing end point key/value pairs. The data sources 1030, 1035, 1040 normalize the protocol used for
20 communication either by employing the end point managers 1045, 1050, 1055 or implementing the details of the protocol needed.

[00112] The illustrated embodiment continues in a step
5 in which results are collected, a response is generated and the sequence is complete. The data manager 1005 may
25 trigger an event to notify the consuming application that the results are available if the application made the request asynchronous.

[00113] FIG. 11 is a flow diagram of one embodiment of
30 a method of requesting a key carried out according to the principles of the invention. The method begins in a start step 1105 when an application requests one or more

keys. In a step 1110 normalized versions of the one or more keys, along with end point attributes and roles, are employed to obtain a normalization mechanism (i.e., execution script, transformation or other normalization) pertaining to the requested key or keys from the normalization mapping repository. In a step 1115, the normalization mechanism is evaluated to determine how it should be carried out. For purposes of this example, the normalization mechanism is assumed to be a script. In a step 1120, the normalization mechanism is executed in the script engine. In a step 1125, the normalization mechanism requests one or more device-specific values using the data source corresponding to the device to which the key or keys correspond. In a step 1130, the values are received and converted or otherwise manipulated to a normalized data type. In a step 1135, the requested key/value pairs are returned in normalized form. In a step 1140, a response is compiled containing the requested key/value pairs. In a step 1145, the response is returned to the application.

[00114] Those skilled in the art to which this application relates will appreciate that other and further additions, deletions, substitutions and modifications may be made to the described embodiments.

WHAT IS CLAIMED IS:

1. A normalization engine, comprising:
 - a data manager configured to receive a request to execute an operation or obtain a key/value pair pertaining to an end point and determine a normalized version of said key in said key/value pair; and
 - a normalization mapping repository configured to contain normalization mechanisms, said data manager further configured to employ said normalized version to obtain one of said a normalization mechanisms from said normalization mapping repository, carry out said normalization mechanism and employ a data source corresponding to said end point to retrieve a device-specific value corresponding to said normalized version.
2. The normalization engine as recited in Claim 1 wherein said data manager is further configured to employ an attribute and a role of said end point to identify said data source .
3. The normalization engine as recited in Claim 2 wherein said role is based on at least some of descriptions, capabilities and attributes of said end point .
4. The normalization engine as recited in Claim 1 wherein said normalization engine is further configured to employ one of a script engine, a transformation engine and another normalization mechanism to carry out said normalization mechanism.
5. The normalization engine as recited in Claim 1 wherein said normalization mechanism is based upon meta data provided in a query to said normalization mapping repository .

6. The service management system as recited in Claim 5 wherein said data manager evaluates said meta data against hierarchical configuration data.

7. The service management system as recited in Claim 1 wherein said normalization mapping repository has a data structure such that device types can inherit normalization mechanisms from one another.

8. The service management system as recited in Claim 1 wherein said data manager is further configured to normalize said key, a value and meta-rules and relationships that said key/value pair have with other key/value pairs .

9. The service management system as recited in Claim 1 wherein said data manager is further configured to allow an application to query to determine what values are available for a given type of key.

10. The service management system as recited in Claim 1 wherein said data manager is further configured logically to determine a value of a key/value pair during a set operation using values of other key/value pairs.

11. A method of requesting a key pertaining to an end point, comprising:

determining a normalized version of said key;

employing said normalized version to obtain a normalization mechanism from a normalization mapping repository; and

carrying out said normalization mechanism, including :

employing an attribute and a role of said end point to identify a data source, and

employing said data source to retrieve a device-specific value corresponding to said normalized version.

12. The method as recited in Claim 11 further comprising converting said device-specific value to a normalized data type.

13. The method as recited m Claim 11 further
5 comprising returning said device-specific value to an application m normalized form.

14. The method as recited m Claim 11 wherein said normalization mechanism is selected from the group consisting of:

10 an execution script, and
a transformation.

15 15. The method as recited m Claim 11 further comprising evaluating said normalization mechanism is evaluated to determine how said normalization mechanism should be carried out.

16. The method as recited m Claim 11 wherein said carrying out comprises executing said normalization mechanism m a script engine.

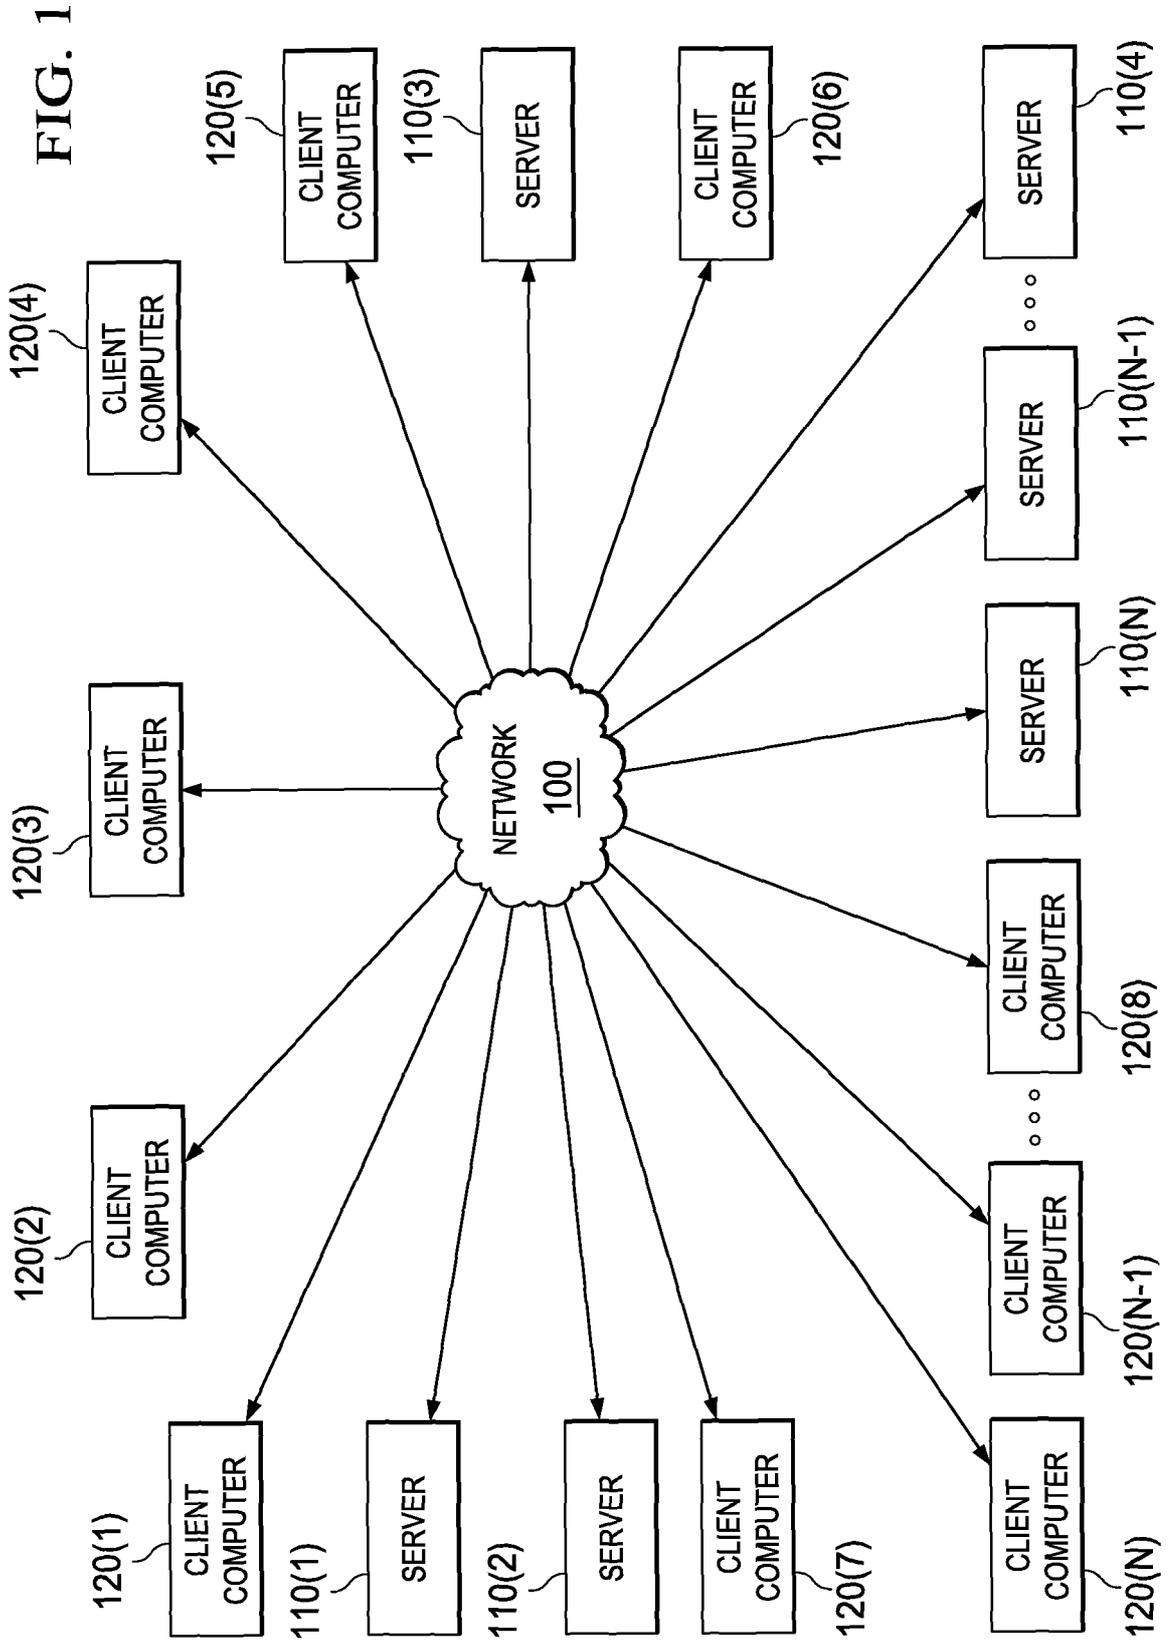
17. The method as recited m Claim 11 further
20 comprising basing said normalization mechanism upon meta data provided m a query to said normalization mapping repository.

18. The method as recited m Claim 11 further
25 comprising evaluating said meta data against hierarchical configuration data.

19. The method as recited m Claim 11 wherein normalization mapping repository has a data structure such that device types can inherit normalization mechanisms from one another.

30 20. The method as recited m Claim 11 further comprising normalizing said key, a value and meta-rules

and relationships that said key/value pair have with other key/value pairs.



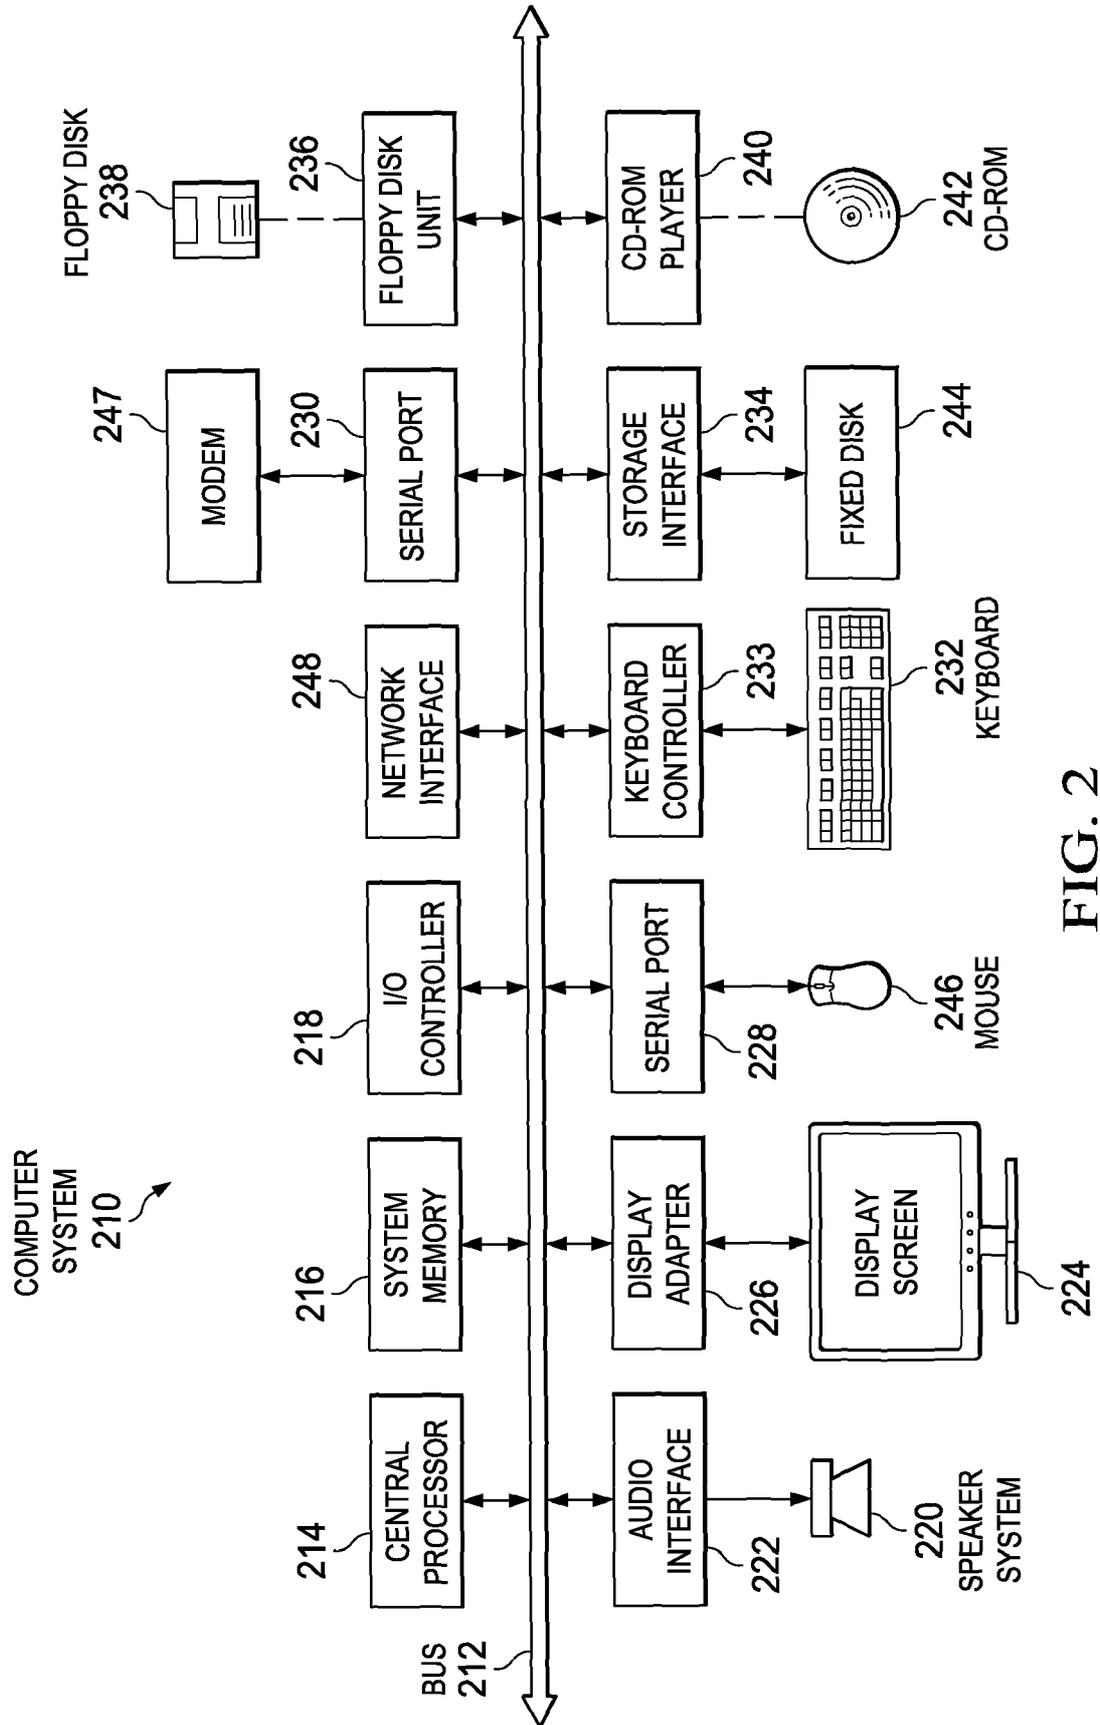
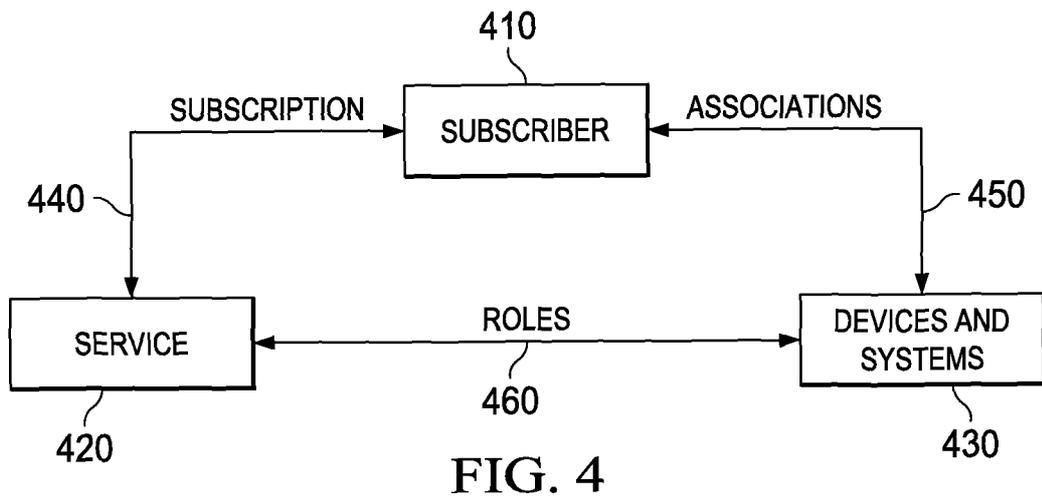
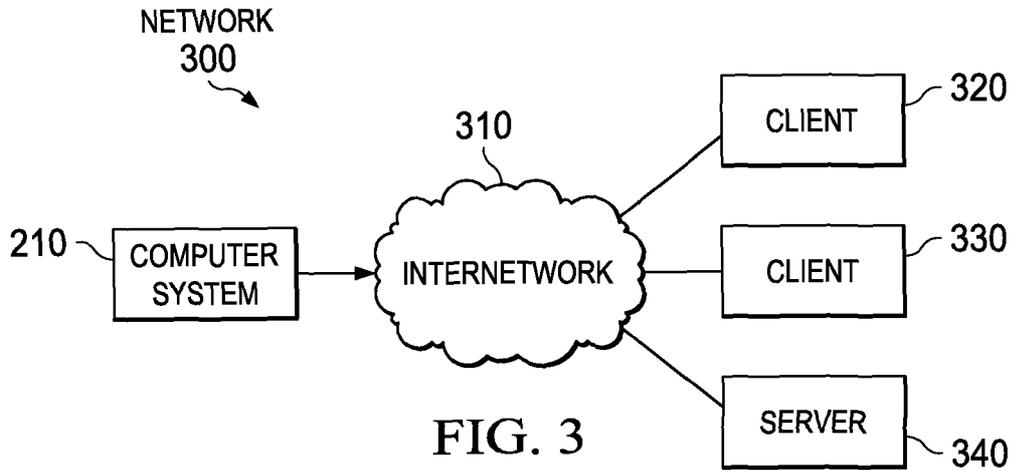


FIG. 2



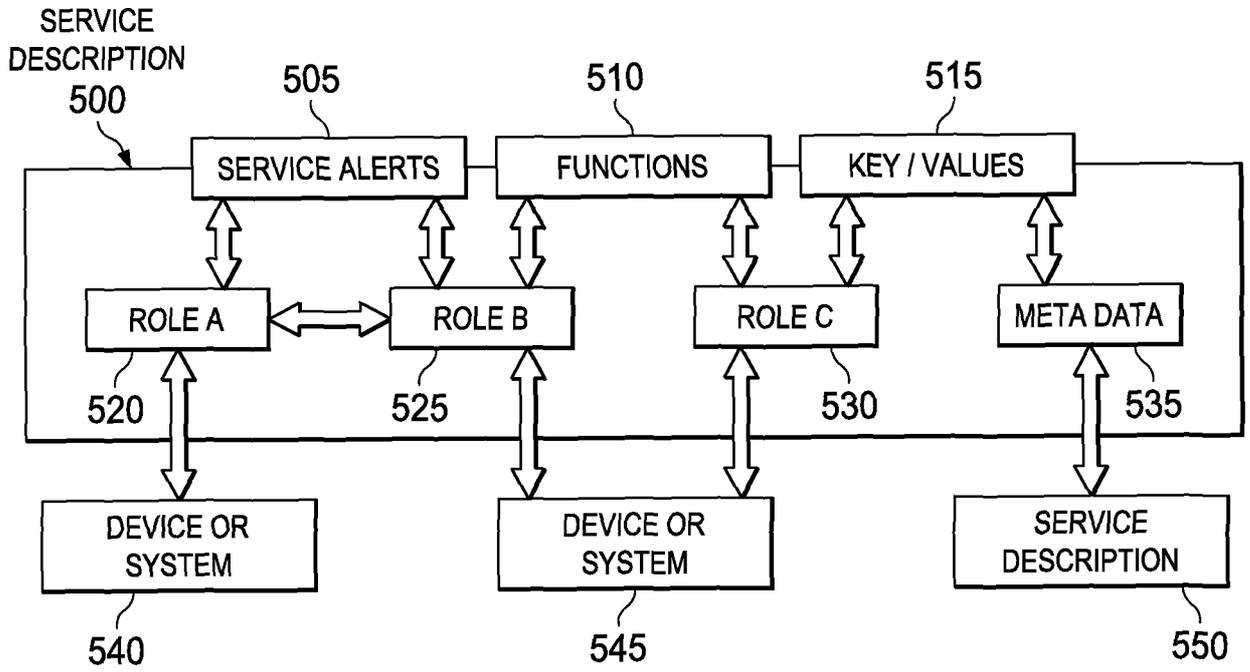


FIG. 5

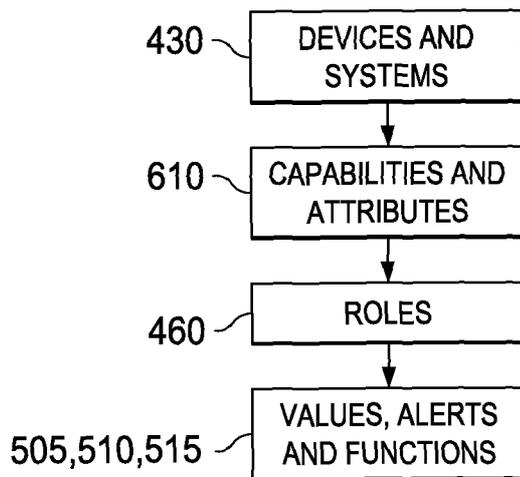


FIG. 6

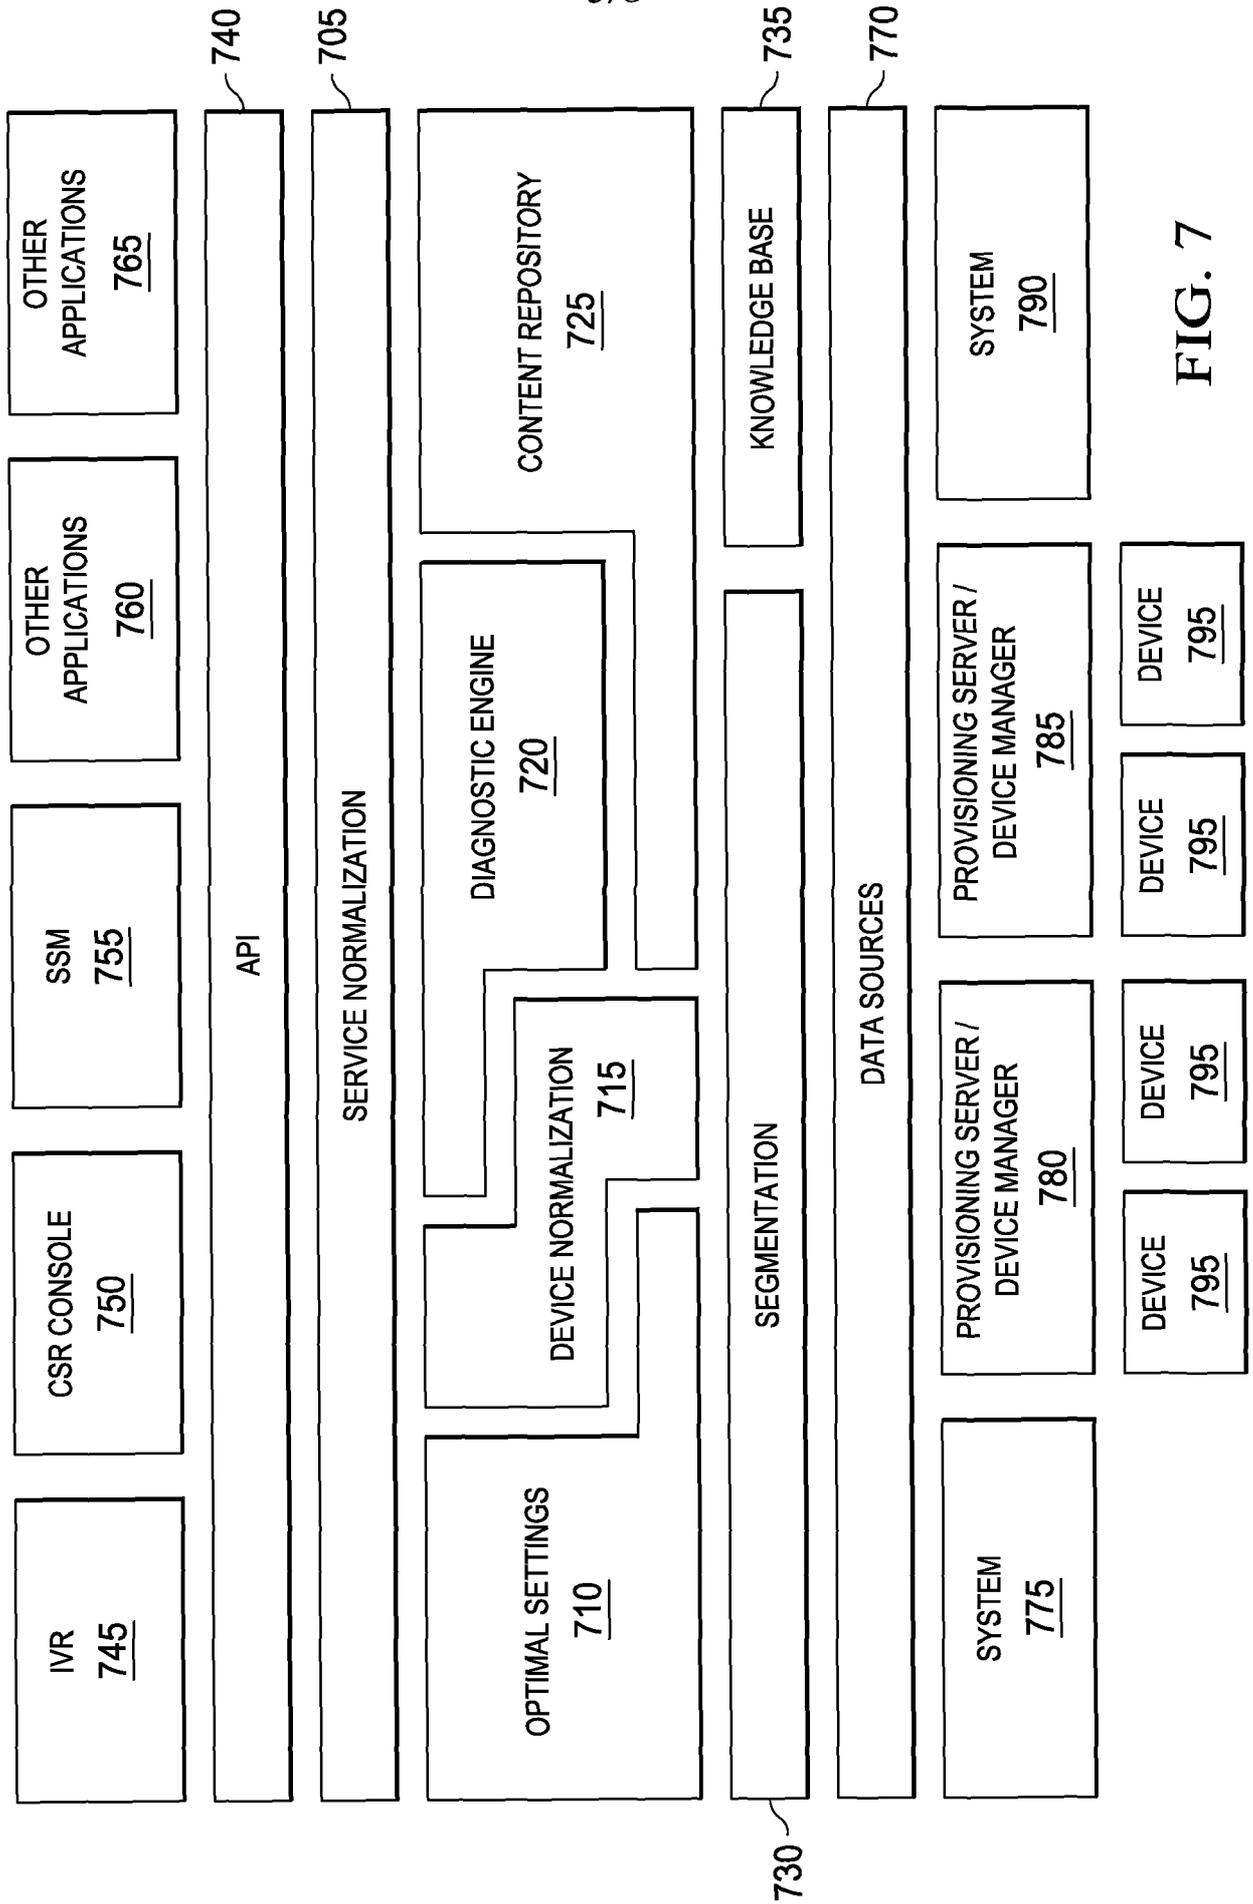
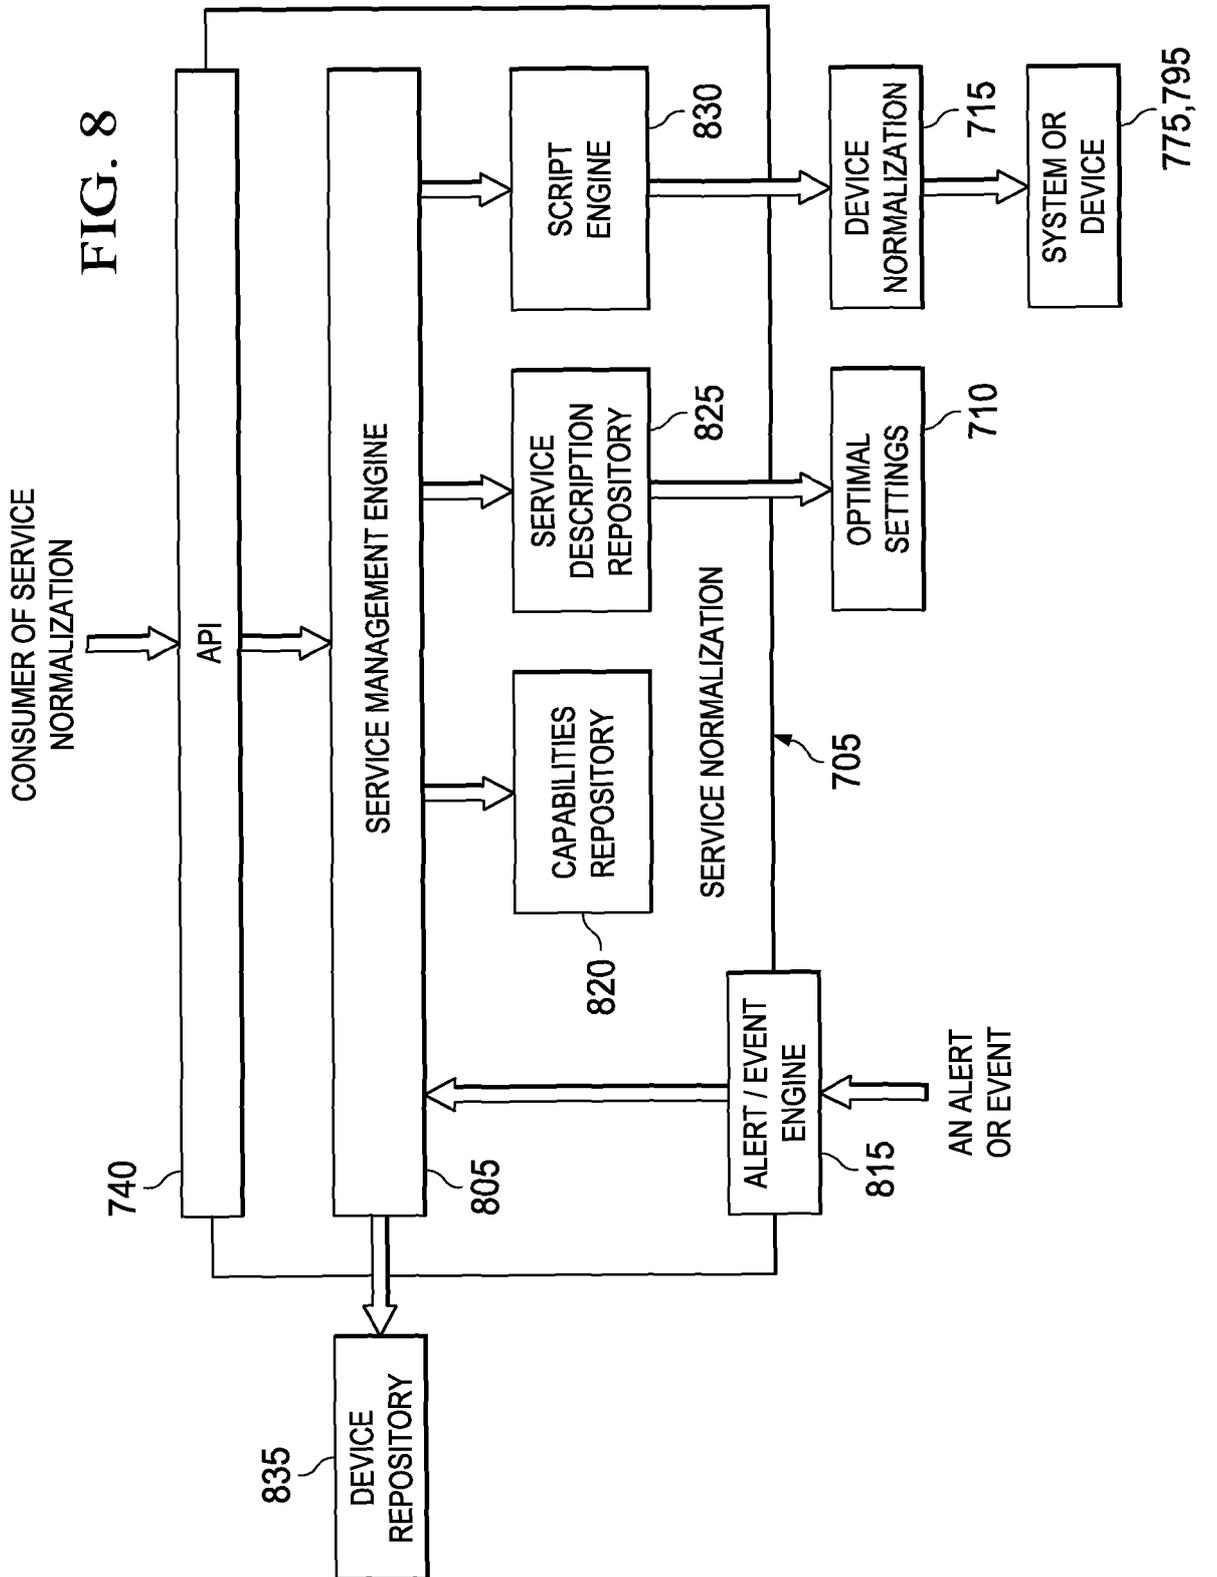


FIG. 7



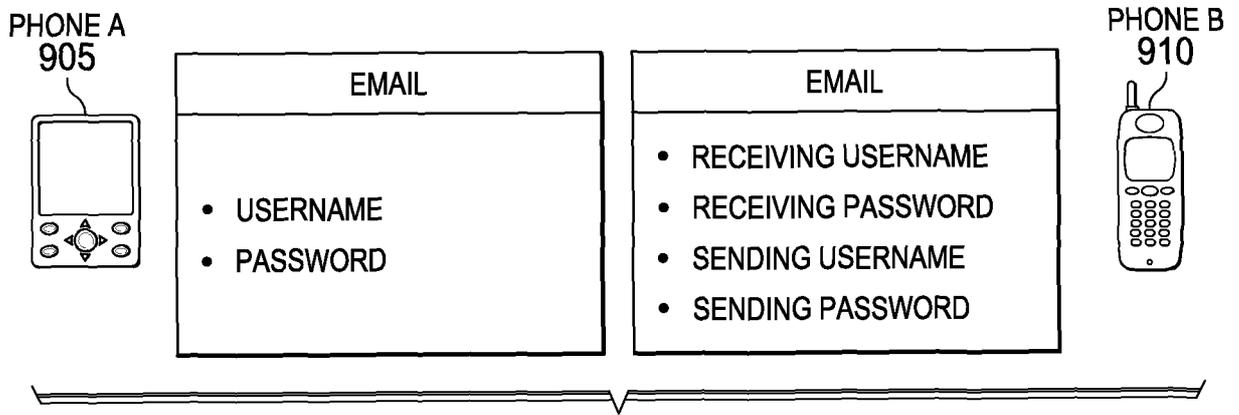


FIG. 9

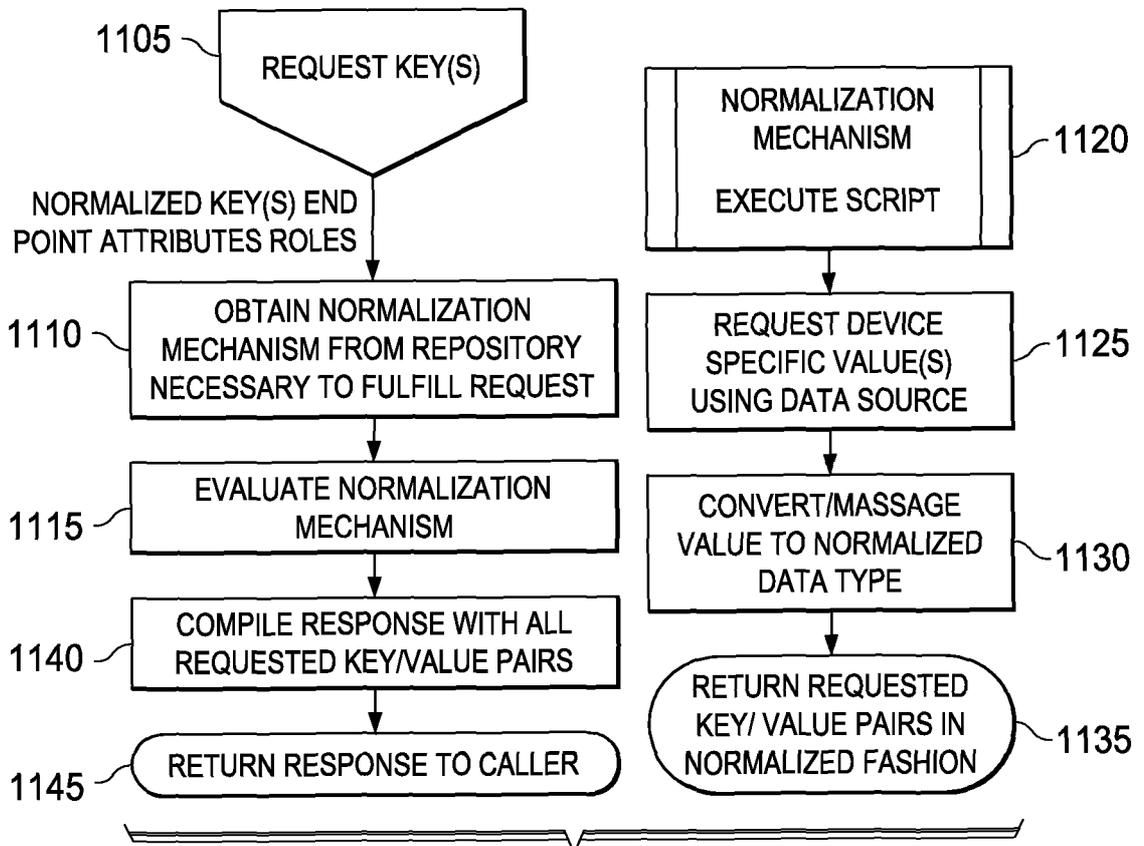


FIG. 11

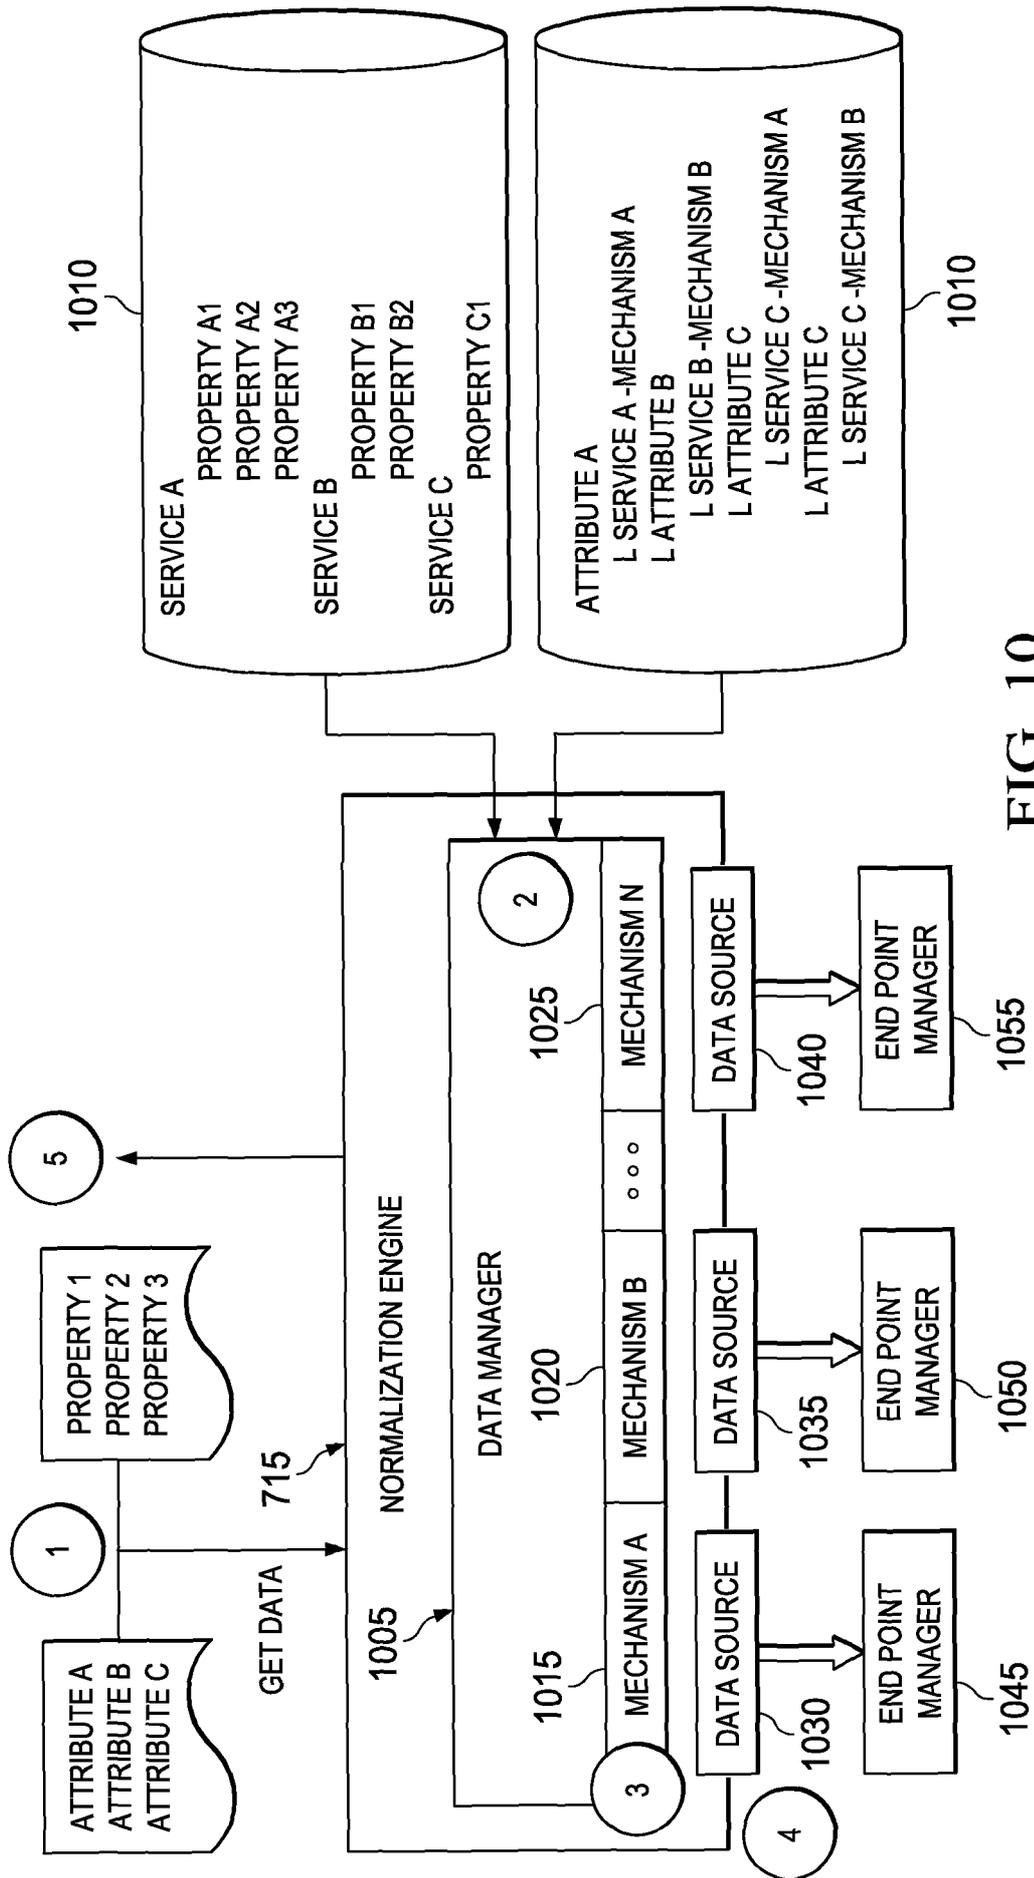


FIG. 10