



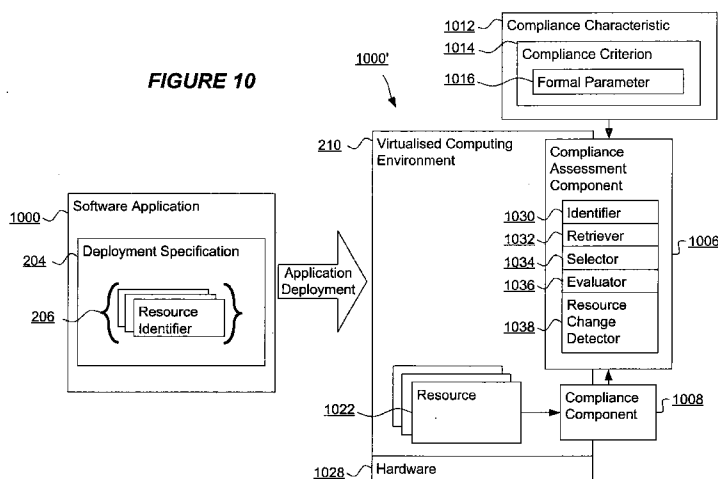
(43) International Publication Date
24 December 2014 (24.12.2014)

- (51) International Patent Classification:
G06F 9/44 (2006.01) G06F 11/30 (2006.01)
- (21) International Application Number:
PCT/GB2014/000231
- (22) International Filing Date:
12 June 2014 (12.06.2014)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
13250068.7 19 June 2013 (19.06.2013) EP
- (71) Applicant: **BRITISH TELECOMMUNICATIONS public limited company** [GB/GB]; 81 Newgate Street, London EC1A 7AJ (GB).
- (72) Inventors: **DIMITRAKOS, Theo**; PP: C5A, 81 Newgate Street, London EC1A 7AJ (GB). **GEORGALAS, Nektarios**; PP: C5A, 81 Newgate Street, London EC1A 7AJ (GB). **EL-MOUSSA, Fadi**; PP: C5A, 81 Newgate Street, London EC1A 7AJ (GB). **PAWAR, Pramod**; PP: C5A, 81 Newgate Street, London EC1A 7AJ (GB). **VAFIADIS, George**; PP: C5A, 81 Newgate Street, London EC1A 7AJ (GB).
- (74) Agent: **ROBERTS, Scott, Anthony**; BT LEGAL, Intellectual Property Department, PP: C5A, BT Centre, 81 Newgate Street, London EC1A 7AJ (GB).

- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, 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, LT, 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, 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).

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

(54) Title: EVALUATING SOFTWARE COMPLIANCE



(57) Abstract: A software compliance assessment apparatus for determining a level of compliance of a software application in execution in a virtualised computing environment, the apparatus comprising: an identifier component operable to identify resources instantiated for execution of the application; a retriever component operable to retrieve a compliance characteristic for the application, the compliance characteristic being retrieved based on the identified resources, and the compliance characteristic having associated a compliance criterion based on a formal parameter; a selector component operable to select a software component for providing an actual parameter corresponding to the formal parameter, the actual parameter being based on data concerning at least one of the resources; an evaluator component operable to evaluate the compliance criterion using the actual parameter; and a detector component operable to detect a change to one or more of the resources, wherein the identifier component, selector component and evaluator component are operable in response to a determination by the detector component that one or more resources is changed, wherein the selector component selects the software component based on an identification of one or more data items that the software component is operable to provide.



Evaluating Software Compliance

Field of the Invention

The present invention relates to software compliance. In particular, it relates to an evaluation of a level of compliance of a software application in execution.

5 Background of the Invention

There is an increasing trend towards the deployment of software applications using software and hardware infrastructures and platforms offered by service providers as services.

Infrastructure as a Service (IaaS) providers offer resources such as hardware or
10 virtualised hardware environments for the deployment of software platforms and applications. IaaS infrastructures can include, inter alia, resources such as hypervisors, storage resources, load balancing resources and network resources. Platform as a Service (PaaS) providers offer platform resources such as, inter alia, operating systems, execution runtime environments, databases, middleware, network services such as web servers and
15 development tools.

Infrastructure and platform services can be implemented so as to abstract any particular deployed application from underlying resources employed. A software application may require specific resources, for example a specific operating system, execution environment, database and web server. The application can be deployed to a platform provided by a
20 platform service provider, the platform having potentially many and numerous alternative resources being selected and configured to satisfy the specific requirements of the application. Further, the platform itself can operate with an infrastructure provided by an infrastructure service provider, certain attributes and resources of which may be at least partly specified for the application. The infrastructure may also have many and numerous
25 alternative resources being selected and configured to satisfy the requirements of the platform and the application. Thus, an application deployment can involve a multiplicity of interconnected resources selected from a potentially greater number of available resources at each of the application, platform and infrastructure level.

A feature of such service-based technologies is that resources can be arranged in a
30 localised or distributed manner. Distributed resources, such as distributed hardware cooperating or managed through physical or logical network arrangements, can provide for the distribution of resources such as, inter alia, distributed execution environments, distributed web servers and distributed database software.

A further feature of such service-based technologies is that resources can be provided in a virtualised manner such as a software implementation of a resource. Thus a hardware device such as a computer system or storage device can be provided as a virtualised device such as a software implementation of a computer system or storage device. Such virtual resources present an abstraction of underlying actual hardware. For example, a computer system resource executing a particular operating system can be provided as a virtual machine (VM) executing with a hypervisor on a hardware device or, potentially, a distributed arrangement of hardware devices. Examples of hypervisor's include native hypervisors that execute in conjunction with specific hardware such as Oracle VM Server for SPARC, Oracle VM Server for x86, the Citrix XenServer, VMware ESX/ESXi, KVM, and Microsoft Hyper-V (Oracle, Oracle VM Server and SPARC are trademarks or registered trademarks of Oracle Corp. in some countries. Citrix and XenServer are trademarks or registered trademarks of Citrix Systems, Inc in some countries. VMware is a trademark or registered trademark of VMware, Inc in some countries. Microsoft and Hyper-V are trademarks or registered trademarks of Microsoft Corp. in some countries.) Additionally, hypervisors can be hosted in existing operating environments, for example BHyVe, VMware Workstation and VirtualBox (VirtualBox is a trademark or registered trademark of Oracle Corp.)

One example of the use of such service-based technologies to deploy software applications is cloud computing. Cloud computing uses hardware and software resources provided as a service over a network, such as the internet. For example, cloud computing service providers can employ IaaS and PaaS to provide cloud computing services for the deployment of software applications. Applications themselves can also be provided as services (known as Software as a Service or SaaS), such as, inter alia, email applications, office applications, social networking applications, virtual desktops, communications applications and games. Thus, applications deployed to cloud computing environments often involve the selection of IaaS and PaaS and potentially SaaS components.

A further feature of such service-based technologies is an abstraction between resource provision and resource consumption such that the deployment of an application with service-based technologies does not require, and indeed preferably does not involve, a complete understanding of the underlying mechanisms and technologies through and with which the resources are provided. Due to the potentially virtualised, distributed and abstracted nature of the services provided, there is reduced transparency of underlying technologies provided to service consumers. This reduced transparency introduces a dependency of a service consumer on the resource service providers with respect to characteristics of the resources. For example, an application requiring a certain standard of information security, security architecture, data governance or resiliency will depend on service providers to commit to

satisfy such requirements and further to actually provide services satisfying the requirements. One way this can be articulated between a service provider and consumer is through a Service Level Agreement (SLA) in which service providers and consumers agree what resources will be provided and what the characteristics of those resources will be.

- 5 While helpful for service consumers, SLAs provide no technical assurance that required characteristics of a particular service level are provided. Indeed the extent of a lack of transparency of a service-based technology will mean that certain resources and characteristics of resources will not be exposed to a service consumer and, accordingly, may not be readily audited by the consumer or an auditor operating on behalf of the consumer.
- 10 For example, a standard of encryption used in communication between deep components in a computing platform or infrastructure, a level of security applied to data stored in a data store, or a level of security applied to computing facilities access may not be exposed or exposable to a service consumer or auditor.

The importance of required characteristics of resources cannot be understated, especially
15 for applications having associated legal or regulatory frameworks or constraints. For example, the location and manner of storage and communication of personal information can require strict control in many territories. Similarly, a level of access control and protection against intrusion can be grounded in legal requirements. It is therefore desirable that an extent or level of compliance with required resource characteristics can be assessed.

20 The Cloud Security Alliance (CSA) has published a set of controls which can be used by cloud computing service consumers in assessing the overall security risk of a cloud provider (Cloud Security Alliance and CSA are trademarks or registered trademarks of the Cloud Security Alliance in some countries). Examples of such controls are listed in a Cloud Controls Matrix (CCM) available at cloudsecurityalliance.org/research/ccm. The controls are mapped
25 to security standards, regulations, and controls frameworks such as: the International Organization for Standardization (ISO) information security standards 27001/27002; the Information Systems Audit and Control Association (ISACA) Control Objectives for Information and Related Technology (COBIT); the Payment Card Industry Data Security Standard (PCI DSS); standards of the National Institute of Standards and Technology (NIST)
30 such as NIST Special Publication 800-53 "Recommended Security Controls for Federal Information Systems and Organizations"; and the North American Electric Reliability Corporation's (NERC) Critical Infrastructure Protection (PIP) (ISO is a trademark or registered trademark of the International Organization for Standardization in some countries. COBIT is a trademark or registered trademark of ISACA and The IT Governance Institute
35 (ITGI) in some countries.) The CCM provides a reference for key compliance characteristics applicable across software applications and service-based technologies. While the CCM is

helpful in assisting cloud computing service providers in identifying desirable characteristics, and the CCM provides a reference for cloud computing service consumers in defining characteristics with which compliance is required, the CCM does not provide for an assessment of an extent or level of compliance with required resource characteristics.

- 5 Manual intervention is required along with service provider transparency to employ the CCM to assess compliance of required resource characteristics.

The CSA further published a Consensus Assessment Initiative (CAI) Questionnaire that provides a set of questions for each control in the CCM that a cloud computing service consumer may ask of a service provider (published at

- 10 cloudsecurityalliance.org/research/cai). The questionnaire provides a series of "yes or no" control assertion questions which can be tailored to suit a service consumer's requirements. While the CAI Questionnaire is helpful to assist service consumers in interrogating service providers, it does not provide for an assessment of an extent or level of compliance with required resource characteristics.

- 15 The CSA has also published a network working group internet draft "CloudAudit – Automated Audit, Assertion, Assessment, and Assurance API (A6)" as Internet Engineering Task Force (IETF) draft "draft-hoff-cloudaudit-00" (Hoff et al, July 2010). CloudAudit provides a namespace and interface that allows cloud computing service providers to make assertions relating to compliance controls at the request of service consumers. The accuracy and
20 appropriateness of the assertions is dependent on the mechanism for making the assertion, and the CloudAudit draft does not contemplate how such assertions are to be founded. Computer Sciences Corporation (CSC) published a précis for a mechanism for requesting and receiving information about compliance controls from cloud computing service providers ("A Précis for the CloudTrust Protocol", CSC, 2010) (CSC is a trademark or registered
25 trademark of Computer Sciences Corporation in some countries.) The CloudTrust protocol (CTP) defines a "question and response protocol" for communication between cloud computing service consumers and cloud computing service providers using the CloudAudit namespace and interface. Requests relate to one of 23 defined "Elements of Transparency" where a subset of the elements relate to "evidence requests" and other elements relate to
30 "policy introduction" requests. For example, one element of transparency can be used to request information relating to a current configuration of a hypervisor. While the CTP provides a mechanism for service consumers to request information from a service provider, it does not provide for an assessment of an extent or level of compliance with resource characteristics that can be relied upon. Cloud computing service providers can choose
35 whether or not to respond to CTP requests, and the response is entirely in the control of the service provider. Further, the CloudAudit interface is fallible. CloudAudit and CTP

repositories may not be secure, private or integrity-guaranteed. The name system of CloudAudit and CTP may be susceptible to attack and servers may not be authenticated. CloudAudit servers may make false assertions or may refer to assertions that do not apply to them.

5 Additionally, service-based technologies such as cloud computing services and deployed applications can have a configuration or architecture that is transient in nature. A feature of service-based technologies is their scalability and “elasticity”. Elasticity refers to the ability of service-based technologies to not only scale up or down as required by a deployed application, but also to transition, move, evolve, add, remove or shift services and resources
10 in accordance with changing needs or requirement of a deployed application or service consumer. In this regard, technologies such as autonomic computing provide self-managing distributed computing resources which adapt to changes in requirements. Such scalability and elasticity of service-based technologies can mean the underlying resources employed to provide services such as IaaS, PaaS or SaaS will change. Accordingly, any change in
15 services and/or resources will require a corresponding review of an extent or level of compliance with required resource characteristics.

Characteristics of the software application can be many and varied and can be distributed throughout the application. Additionally, determining a level of compliance of the software application can require information from multiple sources including the software application
20 itself, a service based environment with which the application operates such as a virtualised computing environment, and software components operating external to both the application and the environment.

US published patent application number US 2011/0321033 A1 (Kelkar et al) describes the use of an application blueprint augmented with a deployment model for the provisioning of an
25 application. US 2011/0321033 further describes how compliance policies can be defined in the blueprint/deployment model. The mere definition of policies for an application is not sufficient for identifying or assessing an extent or level of compliance with required resource characteristics of an application. Further, in view of the elasticity of service-based technologies, providing for such an assessment as underlying services and/or resources for
30 a deployed application change or adapt cannot be achieved by defining compliance policies in a blueprint or deployment model for application provisioning.

It would therefore be advantageous to provide a mechanism for determining an extent or level of technical compliance of a service-based technology for the deployment of a software application accounting for the elasticity of the service based technology that and without the
35 aforementioned disadvantages.

Summary of the Invention

In accordance with a first aspect, the present invention accordingly provides a software compliance assessment apparatus for determining a level of compliance of a software application in execution in a virtualised computing environment, the apparatus comprising: an
5 identifier component operable to identify resources instantiated for execution of the application; a retriever component operable to retrieve a compliance characteristic for the application, the compliance characteristic being retrieved based on the identified resources, and the compliance characteristic having associated a compliance criterion based on a formal parameter; a selector component operable to select a software component for
10 providing an actual parameter corresponding to the formal parameter, the actual parameter being based on data concerning at least one of the resources; an evaluator component operable to evaluate the compliance criterion using the actual parameter; and a detector component operable to detect a change to one or more of the resources, wherein the identifier component, selector component and evaluator component are operable in response
15 to a determination by the detector component that one or more resources is changed, wherein the selector component selects the software component based on an identification of one or more data items that the software component is operable to provide.

Embodiments of the invention provide a separation of concerns between a compliance assessment component and a software component selected to provide an actual parameter.
20 Such separation is advantageous where the resources for a deployed application can change at runtime, such as due to deployment of the application using a service based environment such as a cloud computing environment. In particular, the software component is selected to provide the actual parameter such that the selection of an appropriate software component is based on the data requirements for evaluating the criterion. Accordingly, where
25 one or more of the resources change, the selection of a software component can result in a different software component able to provide the actual parameter for the changed application. Thus the separation of concerns between the compliance assessment component and the software component provides for the selection of an appropriate software component based on the data requirements for evaluating the criterion and the resources
30 instantiated for the deployed application. Further, embodiments of the invention provide an adaptable approach to compliance assessment for software applications executing with service based infrastructures where resources can change at runtime, such as in response to platform or infrastructure re-provisioning, or where a platform or infrastructure exhibits characteristics of resource elasticity as is typical in cloud computing environments.
35 Embodiments of the present invention further provide for such compliance assessment

without a need to interrupt or redeploy the software application, or redeploy a compliance architecture.

Preferably, the software component is a first software component and the first software component is operable to select a second software component, the second software component providing at least some of the data concerning the at least one of the resources. In this way the concerns of the software components are separated. Such separation is advantageous when information from multiple information sources is required to determine a level of compliance with a compliance characteristic or to determine an actual parameter. The approach is particularly advantageous in service based software environments such as cloud computing environments where resources employed in the deployment of an application can change, including changes at runtime. The use of software components with separate responsibilities to provide data for the actual parameter allows for the selection of potentially different components at runtime in response to changes to the resources for the deployed application.

Preferably the data concerning the at least one of the resources includes at least one of: a measurement of a characteristic of the resource; an indicator of a state of the resource; and indicator of an occurrence of an event associated with the resource.

Preferably the actual parameter is determined based on a transformation of the data concerning the at least one of the resources.

Preferably the compliance criterion is a first criterion in a set of multiple compliance criteria for the compliance characteristic, each of the multiple compliance criteria being based on a formal parameter, wherein the selector component is further operable to select one or more additional software components for providing an actual parameter corresponding to a formal parameter for each other criterion in the set of multiple compliance criteria, and wherein the evaluator component is further operable to evaluate each of the compliance criteria using a corresponding actual parameter to determine the level of compliance.

Preferably the software component provides an identification of one or more data items that the software component is operable to provide in response to a request for such identification.

Preferably the second software component provides an identification of one or more data items that the second software component is operable to provide, and the selection of the second software component is based on the identification by the second software component.

The present invention accordingly provides, in a second aspect, a method of a compliance assessment component for determining a level of compliance of a software application in execution in a virtualised computing environment, the method comprising: identifying resources instantiated for execution of the application; retrieving a compliance characteristic for the application, the compliance characteristic being retrieved based on the identified resources, and the compliance characteristic having associated a compliance criterion based on a formal parameter; selecting a software component for providing an actual parameter corresponding to the formal parameter, the actual parameter being based on data concerning at least one of the resources; evaluating the compliance criterion using the actual parameter; and in response to a determination that one or more resources is changed, repeating at least the identifying, selecting and evaluating steps, wherein the selection of the software component is based on an identification of one or more data items that the software component is operable to provide.

The present invention accordingly provides, in a third aspect, a computer program element comprising computer program code to, when loaded into a computer system and executed thereon, cause the computer to perform the steps of the method set out above.

Brief Description of the Drawings

A preferred embodiment of the present invention will now be described, by way of example only, with reference to the accompanying drawings, in which:

Figure 1 is a block diagram of a computer system suitable for the operation of embodiments of the present invention;

Figure 2 is a component diagram of a compliance augmentation tool arranged to augment a deployment specification for a software application in accordance with an exemplary embodiment of the present invention;

Figure 3 is a detailed component diagram of the arrangement of Figure 2 in accordance with an exemplary embodiment of the present invention;

Figure 4 is a flowchart of a method of the compliance augmentation tool of Figures 2 and 3 in accordance with an exemplary embodiment of the present invention;

Figure 5 is a component diagram illustrating resource identification and compliance characteristic selection processes in accordance with an exemplary embodiment of the present invention;

Figure 6a is a component diagram illustrating a resource identification process for a software application having a deployment specification in accordance with an exemplary embodiment of the present invention;

Figure 6b illustrates resources identified by the resource identification component of
5 Figure 6a in accordance with an exemplary embodiment of the present invention;

Figure 7 is a component diagram illustrating a compliance characteristic selection process for resources identified by the process of Figures 6a and 6b in accordance with an exemplary embodiment of the present invention;

Figure 8 is a component diagram of a compliance augmentation tool arranged to
10 augment the deployment specification of Figure 6 in accordance with an exemplary embodiment of the present invention;

Figure 9 is a component diagram illustrating an augmented deployment specification in accordance with an exemplary embodiment of the present invention;

Figure 10 is a component diagram illustrating a deployment of a software application
15 with a virtualised computing environment in accordance with a preferred embodiment of the present invention;

Figure 11 is a component diagram of a plurality of compliance components in accordance with a preferred embodiment of the present invention;

Figure 12 is a flowchart of a method of the compliance assessment component of
20 Figure 10 in accordance with a preferred embodiment of the present invention; and

Figure 13 is a schematic illustration of an arrangement for determining a level of compliance of the software application of Figure 10 with a compliance characteristic in accordance with a preferred embodiment of the present invention.

Detailed Description of the Preferred Embodiments

25 Figure 1 is a block diagram of a computer system suitable for the operation of embodiments of the present invention. A central processor unit (CPU) 102 is communicatively connected to a storage 104 and an input/output (I/O) interface 106 via a data bus 108. The storage 104 can be any read/write storage device such as a random access memory (RAM) or a non-volatile storage device. An example of a non-volatile storage
30 device includes a disk or tape storage device. The I/O interface 106 is an interface to devices for the input or output of data, or for both input and output of data. Examples of I/O devices

connectable to I/O interface 106 include a keyboard, a mouse, a display (such as a monitor) and a network connection.

Figure 2 is a component diagram of a compliance augmentation tool 220 arranged to augment a deployment specification 204 for a software application 202 in accordance with an exemplary embodiment of the present invention. The compliance augmentation tool 220 is a software or hardware component operable to: receive one or more compliance characteristics 212 via receiver 222; select one or more software components 218 via selector 224; and modify the deployment specification 204 via modifier 226.

The receiver 222 is a software or hardware component operable to receive one or more compliance characteristics 212. A compliance characteristic is a characteristic of a deployed software application, such as application 202 deployed for execution in a virtualised computing environment 210. Each of the compliance characteristics 212 received by the receiver 222 are used to determine an extent or level of compliance of the software application 202 when deployed as is described in detail below. For example, compliance characteristics 212 can be defined in a Cloud Compliance Matrix (CCM) provided by the Cloud Security Alliance (CSA).

The selector 224 is a software or hardware component operable to select one or more software components as compliance software components 218 for assessing, measuring or determining an extent or level of compliance of the software application 202 when deployed.

The modifier 226 is a software or hardware component operable to modify a deployment specification 204 associated with the software application 202 to identify the selected compliance software components 218 such that, on deployment of the software application 202, the selected compliance software components 218 are operable to determine a level or extent of compliance of the software application 202 with the received compliance characteristics 212. For example, the modifier 226 can modify the deployment specification 204 by amending, supplementing or augmenting the deployment specification 204 such that the selected compliance software components 218 are instantiated at runtime of the deployed application 202.

Thus, in this way the compliance augmentation tool 220 is operable to modify the deployment specification 204 for the application 202 to incorporate compliance software components 218 selected by the selector 224 based on the received compliance characteristics 212. The deployed software application 202 in operation at runtime is thus augmented by the selected compliance software components 218 such that the compliance software components 218 function to retrieve, inter alia, data, metrics, configuration details,

measurements, test results or any other information suitable for determining characteristics of the deployed software application 202. Such information on the characteristics of the application 202 are suitable for determining an extent or level of compliance of the application 202 with the received compliance characteristics 212. The inclusion of the compliance software components 218 as part of the deployment of the application 202 means that the extent or level of compliance can be continually determined irrespective of whether the software application 202 and/or the virtualised computing environment 210 is adapted, redeployed, adjusted or otherwise changed at runtime, such as in response to changes in the operation of the application 202 or in response to service based facilities such as IaaS, PaaS, SaaS or cloud computing facilities. For example, changes to the workload of application 202 may result in a corresponding change to the provisioning of resources for the application by an infrastructure or platform service provider. Such changes can reflect a feature of services provided environments known as 'elasticity' involving the scaling up or down and/or transitioning of resources to accommodate changing needs of the application 202 or changing demands on the resources or services of the service providers. The deployment of the compliance software components with the application 202 allows for the determination of an extent or level of compliance even when such changes are effected to the application or virtualised computing environment 210.

Figure 3 is a more detailed component diagram of the arrangement of Figure 2 in accordance with an exemplary embodiment of the present invention. Many of the features of Figure 3 are identical to those described above with respect to Figure 2 and these will not be repeated here.

The virtualised computing environment 210 is an environment for the deployment of the software application 202. For example, the virtualised computing environment 210 can be provided as a particular operating system executing within a virtual machine with a hypervisor on a hardware device or, potentially, a distributed arrangement of hardware devices. The virtualised computing environment 210 can be provided as a service-based technology such that the environment 210 is delivered as a service for the installation and execution of a software application such as application 202. In an exemplary embodiment, the virtualised environment is provided as part of a cloud computing service provided by a cloud computing service provider such as BT Cloud Compute available from British Telecommunications plc. Additionally or alternatively, the virtualised computing environment 210 can be provided as, or operate with, a service based infrastructure and/or platform such as IaaS and/or PaaS.

Deployment of the software application 202 includes any or all of installing, configuring, arranging and adapting the software application 202 such that the application 202 is executable with the virtualised computing environment 210. For example, a web based software application can be installed to execute with an operating system executing on a virtual machine, the virtual machine being configured to include networking facilities and the virtual machine also having installed thereon a web server having a certain configuration, a database and certain other requirements defined for the application. All such installation and configuration such that the web based software application is executable in the virtualised computing environment 210 is part of the deployment of the application.

The software application 202 has associated a deployment specification 204 suitable for use in deploying the software application 202 with the virtualised computing environment 210. For example, the deployment specification 204 can include a specification of an architecture of the software application 202 and/or an architecture of software components required for the application 202. Additionally or alternatively, the deployment specification 204 can include specifiers or descriptors of application or other software or platform components that are required for the deployment of the application 202.

In some embodiments the virtualised computing environment 210 is provided as, or operates with, a service based infrastructure and/or platform such a cloud computing service, an IaaS service and/or a PaaS service. In such embodiments the deployment specification 204 is further suitable for use in deploying the software application 202 with such services. The deployment specification 204 identifies one or more resources required for the deployment of the application 202 such that the application executes with the virtualised computing environment 210. Resources can include functions, dataflows and/or technologies. Examples of function resources include bespoke functions, procedures, modules or components provided for the software application 202, such as a library containing functions embodying or supporting the application 202 or a class of instantiable objects providing methods and routines of or for the application 202. Examples of dataflow resources include communications between software components such as the invocation of a function, routine or method of a first component by a facility of a second component. A further example of a dataflow resource is a coupling between two or more components such that messages are passed, requests are sent or data is shared between the two components. Such components can be internal to the application 202 when deployed, part of the virtualised computing environment 210 or external to the application and the virtualised computing environment 210. Examples of technology resources include particular software components, applications or facilities to be installed to deploy the application 202. For example, a technology resource can be a database software component from a particular

technology vendor at a particular version, release or level. Further examples of technology resources include intrusion detection or prevention technologies, virus scanning technologies such as antivirus software, web servers, operating systems, middleware and message handling technologies.

5 Thus, resources can include, inter alia, software or hardware components, software packages, modules, applications, services or solutions, networking facilities, protocols, storage facilities including databases, middleware facilities, user interface facilities and connectivity services. The deployment specification 204 may explicitly identify resources such as an explicit identification of a particular database or web server facility. Alternatively
10 or additionally, the deployment specification 204 can be suitable for identifying a resource such that the identification is not explicit but is discernable. For example, an explicit identification of a web server resource by the software application further identifies dataflows between web page repositories, server side script repositories and the web server. Such dataflows are identified by the deployment specification 204 while such identification is not
15 necessarily explicit. In all cases the deployment specification identifies resources as is illustrated schematically in Figure 3 as a set of resource identifiers 206.

In identifying resources for the deployment of the software application 202, the deployment specification 204 can include an architecture specification as a specification of an environment such as the virtualised computing environment 210, potentially including a
20 definition of an architecture of technology components such as software components, software packages, applications, services or solutions required for the deployment of the application. For example, the deployment specification 204 can be at least partly embodied as an architecture, environment, platform, topology or software stack specification. Such specifications can be expressed in formal terms such as through formal specification
25 languages, semantic definitions, models or diagrams. For example, an architecture of the virtualised computing environment 210 can be expressed as a unified modelling language (UML) model such as a physical UML model, a deployment model or a component model such as are described in "UML Distilled – A Brief Guide to the Standard Object Modelling Language" (Martin Fowler and Kendall Scott, 2000). Further, an architecture of the virtualised
30 computing environment 210 can be built using a formal modelling tool such as the tools in the IBM Rational suite of products (IBM and Rational are trademarks or registered trademarks of IBM Corp. in some countries.) The architecture of the virtualised computing environment 210 is expressed in a parseable or machine readable manner such as within a standard document format, data structure, specification language, modelling language or
35 markup language.

For example, a configuration of the virtualised computing environment 210 can be partly or completely expressed in configuration specifications such as XML files. An example of an XML specification of a virtual machine component of a virtualised computing environment 210 is a domain specification parsed by the 'libvirt' tool. The domain specification is a virtual machine specification stored in an XML file which, when processed by the libvirt tool, can be used to create a new virtual machine in a virtualised computing environment 210. The libvirt domain specification can include, inter alia, the specification of: a hypervisor; an operating system; a boot mechanism such as BIOS bootloader; CPU allocation; CPU tuning; memory allocation; CPU topology; power management; clock and time configuration; input/output device configuration; storage device configuration; filesystem configuration; device busses and controllers; network interfaces; input devices; graphic framebuffers; video devices; consoles; and other physical or software characteristics of a virtualised computing environment 210. Thus a libvirt domain specification can comprise at least part of the deployment specification 204 for an application.

Additionally or alternatively, in identifying resources for the deployment of the software application 202 the deployment specification 204 can include a specification or description of the application 202 and how the application should be deployed, such as by identifying constituent parts of the application and defining installation and configuration details. For example, the paper "Solution Deployment Descriptor 3 Specification 1.0" (Organization for the Advancement of Structured Information Standards (OASIS), 2008) describes a standard for a set of XML documents that define deployment metadata about deployment artifacts and the aggregation of deployment artifacts. The solution deployment descriptor (SDD) provides a standard way to encode deployment information. Such XML documents containing SDD information can be parsed to identify resources associated with the deployment specification. Deployment artifacts are package contents that can be processed to create or modify software resources in the deployment environment. These resources collectively make up the software whose deployment is described by the SDD and include items such as executable files and database table definitions. Examples of deployment artifacts are Linux RPM files, Microsoft MSI files, setup.exe, ZIP, and custom installation executable files (see "Solution Deployment Descriptor (SDD), Part 1: An emerging standard for deployment artifacts", McCarthy & Miller, 2008). Thus the SDD is suitable for identifying resources required for the deployment of an application 202 to a virtualised computing environment 210. Such resources can include applications, software components or technologies installed or deployed for the operation of the software application 202, such as database software, middleware, message handling software, security software, intrusion detection or prevention software etc.

Further techniques suitable for the identification of resources for the deployment of software application 202 include, inter alia: functional decomposition; data model definitions; data schema definitions; library linkages; class and/or object models; component introspection such as object introspection; code analysis such as source code analysis; software component models; component interaction models; and data structures such as those identified by, or available from, integrated software development environments. Such techniques can be employed to identify functional and data components within the application 202. For example, application source and/or packaging code including package information, library linkages such as static or dynamic linkages, build scripts, install scripts or similar, can be used to identify resources. Functional components within the application 202 and resources employed by, linked to, or otherwise associated with, the software application 202 such as resources in the virtualised computing environment 210, software components installed with or for the software application 202 or software components constituting the platform, PaaS, IaaS or other aspects of the architecture of the application 202 or environment 210 can be identified. Further, by reference to data schemas, application and library linkages, build and packaging scripts it is possible to identify dataflow resources. All such insights obtainable about the software application and the virtualised computing environment and being suitable for identifying, directly or indirectly, resources for the software application, constitute part of the deployment specification 204.

It will be apparent to those skilled in the art that, while the deployment specification 204 is illustrated in Figure 2 as being comprised within the software application 202, the deployment specification 204 need not be so integrated and alternatively the deployment descriptor can be associated with the software application 202 or generated for the software application 202. Equally, the deployment specification 204 can exist independent of the software application 202 such that the deployment specification 204 specifies how technologies, software, hardware etc. are to be deployed in order to constitute the software application 202.

One or more of the resources identified by the deployment specification 204 are resources about which compliance characteristics 212 of the software application 202 can be assessed on deployment of the software application 202. Which of the compliance characteristics 212 is relevant to the software application 202 is determined based on the resources identified by the deployment specification 204 as is described in detail below with respect to Figure 5. As a characteristic of the software application 202 when deployed, each of the relevant compliance characteristics 212 can relate to characteristics of the software application 202 itself and/or characteristics of the virtualised computing environment 210 with which the software application 202 executes. Yet further, relevant compliance characteristics 212 can

relate to characteristics of software, hardware, functions, features or services employed in deploying the software application 202 such as software installed with the virtualised computing environment 210, and/or software, hardware, functions, features or services external to both the software application 202 and the virtualised computing environment 210, such as software components providing services or functions to the software application 202 or the virtualised computing environment 210. Thus, the compliance characteristic 212 can include characteristics of the application 202, the virtualised computing environment 210, any environment for the deployment of the application 202 such as a cloud computing service, an IaaS service, a PaaS service, or a service or function provided external to any virtualised, Cloud, IaaS or PaaS service but operating in conjunction with such service.

Examples of characteristics of software applications include, inter alia: features, facilities, attributes and services of an application such as: resources used; algorithms employed; protocols supported; versions of features, algorithms, services or protocols supported or used; performance characteristics such as speed, overhead or throughput; a level or standard of security; adherence to one or more defined standards; update or refresh intervals used; level of up-to-datedness of features, facilities, attributes or services; environments, systems, protocols or functions used; particular versions or levels of environments, systems protocols or functions used; hardware or software supported; audit facilities available; data governance technology or services employed; user access controls employed; hardware requirements; languages used; encryption standards used; patch management processes employed; intrusion detection or prevention facilities available; virus-detection, protection and prevention facilities available; financial handling facilities available; diagnostic tools employed; diagnostic services available; legal or regulatory requirements adhered to; policies employed; third-party access controls in place; reliability facilities provided; accessibility features available; stability features employed; database used; database facilities supported; geographic location of hardware or software; particular geographic distribution, or non-distribution, of hardware or software; features of physical equipment security; networks supported; data integrity facilities used or measures available; and any other characteristic conceivably attributable to a software application as will be apparent to those skilled in the art.

Each of the compliance characteristics 212 is defined by a set of compliance criteria 214. The set of compliance criteria 214 for a compliance characteristic 212 is used to determine an extent or level of compliance of the deployed software application 202 with the compliance characteristic 212. Each criterion in the set of compliance criteria 214 concerns a resource identified by the deployment specification 204. For example, a compliance criterion may explicitly relate to a resource identified by one of the resource identifiers in the set of

resource identifiers 206. Alternatively or additionally, a compliance criterion can concern a feature, attribute, characteristic or component associated with a resource. For example, a criterion may relate to, inter alia, a provider of a resource, a counterpart to a resource, a configuration of a resource or a function of the resource.

5 Where multiple compliance criteria 214 define a compliance characteristic 212 then satisfaction of all the compliance criteria 214 is normally required for the deployed software application 202 to be fully compliant. Satisfaction of anything less than all the criteria 214 will normally constitutes non-compliance. In some embodiments a single criterion in the set of compliance criteria 214 is sufficient to define a compliance characteristic. Further, in some
10 embodiments, a more complex set of compliance criteria 214 may be conceived such that satisfaction of a subset of the compliance criteria 214 by a deployed software application 202 is determined to be sufficient to constitute full compliance with the compliance characteristic 212. For example, multiple alternative compliance criteria 214 may be provided, any or all of which are satisfactory alternatives to each other. Yet further, in an alternative embodiment
15 the set of compliance criteria 214 may be comprised of a plurality of subsets of compliance criteria, any or all of which being sufficient to constitute compliance with the compliance characteristic 212. Thus, an extent to which a deployed software application 202 satisfies compliance criteria 214 in the set of compliance criteria is suitable for determining a level of compliance of the software application 202 with the compliance characteristic 212 when the
20 application 202 is deployed. One way to measure a level of compliance for the deployed software application 202 is to evaluate a proportion of all the compliance criteria 214 in the set of compliance criteria 214 that are satisfied and use such proportion as a quantitative measure of a level or extent of compliance. In some embodiments different compliance criteria 214 can have different weights associated such that an evaluation of a quantitative
25 level of compliance includes applying weights, such as multiplicative factors, to certain of the compliance criteria 214 when determining a proportion of all the compliance criteria 214 that are satisfied. In this way it is possible to impart a greater emphasis on certain of the compliance criteria 214 in the set.

The compliance software components 218 belong to a set stored in a compliance
30 component library 216. The library 216 can be a data store, database, single or multiple static or dynamic software libraries, repository, software object library or any other suitable mechanism for storing the compliance software components 218 as will appear to those skilled in the art. Each compliance software component 218 is selectable for one or more compliance characteristics 212 such that a compliance software component 218 is operable
35 to contribute to an assessment of an extent or level of compliance of one or more compliance characteristics 212. Thus, in preferred embodiments, the compliance software components

218 are operable to assess the satisfaction of the one or more compliance criteria 214 associated with one or more compliance characteristics 212.

The compliance software components 218 can be embodied as, inter alia, software routines, agents, modules, functions, methods or objects for determining an extent or level of compliance of the software application 202 with the received compliance characteristics 212. The level of compliance of the software application 202 is assessed, determined or measured when the application 202 is deployed and can include a level of compliance of the application 202 itself by virtue of the functions, operations, behaviours, data items and communications of the software application 202 and additionally the compliance of the execution environment in which the application 202 operates including, inter alia, the virtualised computing environment 210 and any additional internal or external facilities, services or components with which the application 202 operates.

In an exemplary embodiment, each of the compliance software components 218 is a functional component for executing with the deployed application 202. In some embodiments a software component 218 can be embodied as a configuration change to the software application 202, such as a selection of a mode of operation of a component of the software application 202. For example, a compliance software component 218 can be embodied as a change of a mode of operation of a function of the application 202 such that the function operates to generate trace output, such as debug or verbose status information. Such a compliance software component 218, possibly in conjunction with other compliance software components 218 or other functionality, can cause the generation of additional data that can be used to determine an extent or level of compliance with a compliance characteristic 212. Further, in some embodiments, a compliance software component 218 can be operable to prepare and/or send messages or invoke functions or engage application programming interfaces (APIs) of components comprised in the application 202, the environment 210 or other components, features or services executing with the application 202. In some embodiments a compliance software component 218 can be operable to test or challenge a feature or service of the application 202, environment 210 or other components. Further examples of the operation of compliance software components 218 include, inter alia, components that analyse, check, inspect, evaluate, scrutinise, probe or review features or services of the application 202, environment 210 or other components.

In an exemplary embodiment, the virtualised computing environment 210 provides interfaces such as application programming interfaces (APIs) accessible to compliance software components 218. The compliance software components 218 can use such interfaces to, inter alia: collect data; request information; retrieve or set configuration; activate

or deactivate functionality; and other features and functionality provided by the interface. Such information and functions provided by the interface are suitable for the compliance software component 218 to contribute to an assessment of an extent or level of compliance of the software application 202 operating with the virtualised computing environment 210.

5 Further, in a preferred embodiment, resources external to both the application 202 and the virtualised computing environment 210 can provide interfaces such as application programming interfaces (APIs) accessible to compliance software components 218. Such components can include, inter alia: third party services or functions; supplementary facilities; external routines; collaborating applications; or other external resources. The compliance
10 software components 218 can use such interfaces in a manner similar to the use of interfaces of the virtualised computing environment 210 described above.

While the receiver 222, selector 224 and modifier 226 are illustrated as being comprised within the compliance augmentation tool 220, it will be appreciated by those skilled in the art that these components may be provided external to the compliance augmentation tool 220
15 such as in association with, in communication with or otherwise operable with the compliance augmentation tool 220.

Figure 4 is a flowchart of a method of the compliance augmentation tool 220 of Figures 2 and 3 in accordance with an exemplary embodiment of the present invention. The method augments a deployment specification 204 of a software application 202 such that an extent
20 or level of compliance of the application 202 with a compliance characteristic 212 can be determined when the application 202 is deployed. Initially, at step 402, a definition of the compliance characteristic 212 is received including a set of one or more compliance criteria 214. The satisfaction of the compliance criteria 214 for the compliance characteristic 212 is suitable for determining a level of compliance of an application 202 with the compliance
25 characteristic when the application is deployed. At step 404 at least one compliance software component 218 is selected from a compliance component library 216. The selection is based on the definition of the compliance characteristic 212 as will be described in detail below with respect to Figure 5. The compliance software component 218 is operable to determine a state of satisfaction of at least a subset of the set of compliance criteria 214 for the
30 compliance characteristic 212. At step 406 the deployment specification 204 for the application 202 is modified to identify the selected compliance software components 218 such that, on deployment of the application 202, the application 202 is operable to determine a level of compliance of the application 202 with the compliance characteristic 212.

Figure 5 is a component diagram illustrating resource identification and compliance
35 characteristic selection processes in accordance with an exemplary embodiment of the

present invention. Many of the features of Figure 5 are identical to those described above with respect to Figures 2 to 4 and these will not be repeated here. The deployment specification 204 for the software application 202 identifies resources required for the deployment of the application. In the embodiment of Figure 5 the resources are identified using an architecture specification 502 and a deployment descriptor 504. A resource identification component 506 is a software or hardware component for establishing resource identifiers 206 based on the deployment specification 204. The resource identification component 506 can be an integral part of the compliance augmentation tool 220 or, alternatively, the resource identification component 506 can be at least partly external to, and operable or associated with, the compliance augmentation tool 220. In an exemplary embodiment the resource identification component 506 includes three further software or hardware components: a function identifier 508; a dataflow identifier 510; and a technology identifier 512. The function identifier 508 is operable to identify function resources required for the deployment of the software application 202, such as function resources previously described. The dataflow identifier 510 is operable to identify dataflow resources required for the deployment of the software application 202, such as dataflow resources previously described. The technology identifier 512 is operable to identify technology resources required for the deployment of the software application 202, such as technology resources previously described. The function identifier 508, dataflow identifier 510 and technology identifier 512 are operable to identify resources 206 based on the architecture specification 502 and the deployment descriptor 504. It will be appreciated that while the resource identification component 506 is illustrated as including the function identifier 508, dataflow identifier 510 and the technology identifier 512, a subset of these components or one or more alternative components suitable for identifying at least a subset of resources required to deploy the software application 202 could alternatively be employed.

In an exemplary embodiment, the resource identifiers 206 of the resources required to deploy the software application 202 are used to identify a set of compliance characteristics 212 for the application 202. A compliance characteristic selector 514 is a software or hardware component for selecting one or more compliance characteristics 212 from a dictionary 516 of compliance characteristics. The selected compliance characteristics 212 are those compliance characteristics 212 with which an extent or level of compliance of the deployed application 202 is measured or assessed. The compliance characteristic selector 514 can be an integral part of the compliance augmentation tool 220 or, alternatively, the compliance characteristic selector 514 can be at least partly external to, and operable or associated with, the compliance augmentation tool 220.

In an exemplary embodiment the dictionary 516 is a repository of references to compliance characteristics, some or all of which can be selected by the compliance characteristic selector 514 for applicability to the application 202. The compliance characteristic selector 514 can include rules for determining which compliance characteristics should be selected from the dictionary 516. In an exemplary embodiment the dictionary 516 provides a mapping between resources and compliance characteristics such that the compliance characteristics 212 can be selected based on the identified resources 206. For example, the dictionary 516 can be a correspondence table relating resources to compliance characteristics. Identified resources 206 can be associated with attributes such as: a resource type; a resource name; a resource version etc. The dictionary 516 can provide a correspondence, mapping, association or other identification of one or more compliance characteristics 212 for certain attributes of a resource. For example, resources of the type "database" can be associated with compliance characteristics 212 relating to database characteristics. Thus, in use, the compliance characteristic selector 514 selects compliance characteristics 212 for receipt by the receiver 222 of the compliance augmentation tool 220. In this way the compliance characteristics 212 with which an extent or level of compliance of the deployed software application 202 is measured can be determined based on the deployment specification 204 for the software application 202.

Some compliance characteristics in the dictionary 516 may be designated as mandatory or broadly applicable such that the compliance characteristics are used for all applications irrespective of the identified resources 206. Further, compliance characteristics 212 can be identified for the software application 202 based on specific operational requirements determined for the application 202. In some embodiments, a set of compliance characteristics 212 can be defined to reflect operational standards required for the deployed application 202 and/or relevant compliance characteristics 212 can be identified in dependence on the particular constitution and/or configuration of the deployed software application 202 as reflected by the identified resources 206. For example, a software application dealing with personal confidential information may be required to comply with legal and regulatory requirements reflected by one or more compliance characteristics. Accordingly, where an assessment of the identified resources 206 indicates that such information is handled by the application, compliance characteristics relating to such regulatory requirements can be selected by the compliance characteristic selector 514.

An exemplary embodiment of the invention will now be described in use. Figure 6a is a component diagram illustrating a resource identification process for a software application 702 having a deployment specification 704 in accordance with an exemplary embodiment of the present invention. The deployment specification 704 provides a specification of resources

required for the deployment of the software application 702 to a virtualised computing environment 710. The deployment specification 704 includes a specification of architectural components, functional components, technologies, data and data flows. The virtualised computing environment 710 is a service based environment such that aspects of the environment 710 are provided as services by service providers. For example, the virtualised computing environment 710 is a cloud computing environment. Physical resources 7102 such as central processing units (CPUs), memory, disk storage and network hardware are available to the virtualised computing environment 710. Software implementations of physical machines are deployed in the virtualised computing environment 710 using the physical resources 7102. Such software implementations of physical machines are known as virtual machines (VMs) and are created and managed by one or more hypervisor software components 7204. Each hypervisor component 7204 includes one or more of each or any of: operating systems (OS); virtual or physical storage devices; virtual or physical networking devices or networks; and virtual or physical access control facilities. It will be apparent to those skilled in the art that other virtual or physical devices, services and facilities can be created, handled and managed by hypervisors 7204. The virtualised computing environment further includes a pool 7106 of applications, technologies and software services including one or more of each or any of: database products (DB); virus detection and prevention products; payment protection services (PPS); integrated development environments (IDS); load balancing facilities (Load); web servers (Web); intrusion prevention or prevention products (IPS); messaging middleware products (MSG); electronic commerce applications (Store); and transaction handling software (TXN). It will be apparent to those skilled in the art that other applications, software, technologies and services can be provided by pool 7106 of the virtualised computing environment 710. The virtualised computing environment 710 can be provided as a distributed environment such that physical resources 7102 are distributed on a number of physical machines communicatively connected by a network.

The deployment specification 704 for the application 702 includes an architecture specification 802 and a deployment descriptor 804. The architecture specification 802 expresses architectural resources required for deployment of the software application 704 to a virtualised computing environment 710. In the exemplary embodiment the architecture specification 802 is expressed in a standard form such as XML (not shown). The architecture specification 802 includes a specification of resources of a virtualised computing environment including: a hypervisor "Hypervisor1"; an operating system "OS-A"; a virtual storage component "Disk_Storage"; a virtual networking resource "Network_Ethernet_TCP"; and a user access control service "User_Access_Control".

The deployment descriptor 804 includes a specification of application resources required to deploy the application 702. In the exemplary embodiment the deployment descriptor 804 is expressed in a standard form such as XML (not shown). The deployment descriptor 804 indicates resources including: a database software component "DB_MYSQL"; a database
5 scheme for the application 702 "DB_Schema"; a financial payment processing facility "Payment_Processor"; a web server "Apache_WebServer"; a web application "Web_Application"; and an intrusion prevention system "IPS".

The deployment specification 704 further includes an indication of dataflows between architectural components, functional components and technologies for the deployed
10 application 702. Such dataflows may be explicitly specified in the deployment specification 704 or may be indicated by, or discernible from, the deployment specification 704. In the exemplary embodiment of Figure 6a the deployment descriptor 804 includes indications of five dataflows: a dataflow 860 between the "DB_MYSQL" resource and the "Payment_Processor" resource; a dataflow 862 between the "DB_MYSQL" resource and the
15 "Web_Application" resource; a dataflow 864 between the "Apache_WebServer" resource and the "Web_Application" resource; a dataflow 866 between the "Web_Application" resource and the "Payment_Processor" resource; and a dataflow 868 between the "DB_MYSQL" resource and the "DB_Schema" resource.

A compliance augmentation tool 720 includes a resource identification component 806 for
20 identifying resources required for deployment of the application 702. The resource identification component 806 is operable to receive the architecture specification 802 and the deployment descriptor 804 and generate a list of identified resources 706. In the exemplary embodiment the resource identification component 806 parses the XML representations of the architecture specification 802 and the deployment descriptor 804 to identify resources
25 including functional resources, technology resources, data resources and data flow resources.

Figure 6b illustrates resources identified by the resource identification component 806 of Figure 6a in accordance with an exemplary embodiment of the present invention. The identified resources 706 are stored in a data structure. For each identified resource, the
30 resource identification component 806 identifies: a category 7062 of the resource, such as "Technology", "Functional", "Data Flow" or "Data"; a type 7064 of the resource; an identifier 7066 of the resource; and where applicable, a version 7068 of the resource. "Hypervisor1" is identified as a technology resource 960 of type "Hypervisor" and has a version of "AcmeVisor 1.1". "OS-A" is identified as a technology resource 962 of type "Operating System" and has a
35 version of "AcmeOS 3.0". "Disk_Storage" is identified as a technology resource 964 of type

“Storage”. “Network_Ethernet_TCP” is identified as a technology resource 966 of type “Network”. “DB_MYSQL” is identified as a technology resource 968 of type “Database” and has a version of “MySql 5.6.11”. “DB_Schema” is identified as a functional resource 970 of type “Bespoke Database Schema”. A data flow resource 868 is identified between

5 “DB_Schema” and “DB_MYSQL”. “Payment_Processor” is identified as a technology resource 974 of type “Payment Processing System” and has a version of “Acme PPS 2.1”. A data flow resource 860 is identified between “Payment_Processor” and “DB_MYSQL”. “Apache_WebServer” is identified as a technology resource of type “Web Server” and has a version of “Apache HTTP Server 2.4.4”. “Web_Application” is identified as a functional

10 resource 980 of type “Bespoke Application”. A data flow resource 864 is identified between “Apache_WebServer” and “Web_Application”. A data flow resource 862 is identified between “DB_MYSQL” and “Web_Application”. A data flow resource 866 is identified between “Payment_Processor” and “Web_Application”. “IPS” is identified as a technology resource 994 of type “Application Software” and has a version of “Intrusion Protect 2.8”. Within

15 “Web_Application” three further resources are identified: resource “HTML_WebPages” is identified as a data resource 982 of type “Data”; resource “PHP_Scripts” is identified as a functional resource 984 of type “Bespoke Application”; and resource “Web_Commerce” is identified as a technology resource 986 of type “Application Software” and having version “OS Commerce 3.0.2”.

20 Figure 7 is a component diagram illustrating a compliance characteristic selection process for resources 706 identified by the process of Figure 6a in accordance with the exemplary embodiment of the present invention. The compliance augmentation tool 720 further includes a compliance characteristic selector 814 for selecting compliance characteristics from a compliance characteristic dictionary 816 based on the identified resources 706. The

25 compliance characteristic dictionary 816 in the exemplary embodiment of Figure 7 includes a mapping 8162 for mapping resources onto one or more compliance characteristics. The mapping 8162 can include mapping resources based on their category 7062, type 7064, version 7068 or combinations thereof. The mapping can include wildcards such that resources of any type, category or version can be mapped to compliance characteristics. The

30 mapping 8162 maps resources (by type, category or version) to identifiers of compliance characteristics.

In the exemplary embodiment of Figure 7 the compliance characteristic selector 814 selects a payment processor compliance characteristic 712 based on the compliance characteristic dictionary 816. The payment processor compliance characteristic 712 includes

35 a set of compliance criteria 714 for assessing an extent or level of compliance of the software application 702. Criterion 7142 is satisfied if a user access control level is set to “HIGH” and

so concerns the "User_Access_Control" resource of the architecture specification 704. Criterion 7144 is satisfied if a set of most up-to-date rules are associated with an intrusion prevention system. Thus, criterion 7144 concerns the "IPS" resource of the deployment descriptor 804. Criterion 7146 is satisfied if financial cardholder information, such as credit
5 cardholder information, is not stored. Thus, criterion 7146 concerns the "Payment_Processor" resource of the deployment descriptor 804.

Figure 8 is a component diagram of a compliance augmentation tool 720 arranged to augment the deployment specification 704 of Figure 6a in accordance with the exemplary embodiment of the present invention. A compliance component library 716 includes
10 compliance components 7182, 7184 and 7186 corresponding to each of the compliance criteria 7142, 7144 and 7146. It will be appreciated that while the compliance components 7182, 7184 and 7186 are depicted as separate components, the compliance components 7182, 7184 and 7186 could equally be combined into a single compliance component associated with the compliance characteristic 712. Thus, in the exemplary embodiment there
15 is a one-to-one mapping between compliance criteria and compliance components. Equally, there could be a one-to-one mapping between compliance characteristics and compliance components.

Compliance component 7182 is a user access control evidence collector. Component 7182 is operable to collect evidence by way of information pertaining to the user access
20 control facilities of the deployed software application 702. For example, at runtime of the deployed application 702, component 7182 can interface with an API of the "User_Access_Control" resource implemented in the virtualised computing environment 710, such as a user access facility provided by a selected hypervisor 7204 for the application 702. The user access control evidence collector component 7182 can then be operable to request
25 configuration information from the "User_Access_Control" resource. The component 7182 is further operable to evaluate the compliance criterion 7142 to determine if the user access control level for the deployed application 702 is "HIGH". Thus, at runtime, the user access control evidence collector 7182 is operable to determine the satisfaction or otherwise of the deployed application 702 with compliance criterion 7142.

30 Compliance component 7184 is an intrusion prevention system version checker. Component 7184 is operable to collect information on a version of intrusion detection rules employed by the "IPS" resource of the deployed application 702. For example, at runtime of the deployed application 702, component 7184 can interface with an API of the "IPS" resource implemented in the virtualised computing environment 710, such as an intrusion
35 prevention system provided from the applications pool 7106 of the virtualised computing

environment 710. The component 7184 is further operable to communicate with a provider of the intrusion prevention system, such as an external entity developing rules for the particular deployed intrusion prevention system, to identify a most recent version of the rules for the applicable intrusion prevention system. In this way the intrusion prevention system rule version checker 7184 is operable to evaluate the compliance criterion 7144 to determine if the intrusion prevention system rules correspond to a latest version of rules available from the intrusion prevention system service provider. Thus, at runtime, the intrusion prevention system rule version checker 7184 is operable to determine the satisfaction or otherwise of the deployed application 702 with compliance criterion 7144.

10 Compliance component 7186 is a cardholder data storage evidence collector. Component 7186 is operable to collect information on the storage of data in the "DB_MYSQL" resource by the "Web_Application" and "Payment_Processor" resources of the deployed application 702. For example, at runtime of the deployed application 702, component 7186 can monitor: data stored in the "DB_MYSQL" resource; data exchanged in the data flow between the
15 "Payment_Processor" and the "Web_Application" resources; and data exchanged in the data flow between the "Payment_Processor" resource and the "DB_MYSQL" resource. In this way the cardholder data storage evidence collector 7186 is operable to evaluate the compliance criterion 7146 to determine if cardholder data is stored by the application 702. Thus, at runtime, the intrusion prevention system rule version checker 7184 is operable to determine
20 the satisfaction or otherwise of the deployed application 702 with compliance criterion 7146.

In combination the compliance components 7182, 7184 and 7186 are operable to evaluate an extent or level of compliance of the software application 702 with the compliance characteristic 712 at a runtime of the application 702. The extent or level of compliance can be quantitative such that each of the compliance criteria 7142, 7144 and 7146 constitute a
25 third of a total compliance measure.

The compliance augmentation tool 720 is operable to augment the deployment specification 704 for the software application 702 based on the payment processor compliance characteristic 712 identified by the process depicted in Figure 7. The receiver 722 initially receives the payment processor compliance characteristic 712. Subsequently the selector 724 selects the compliance components 7182, 7184 and 7186 corresponding to the compliance characteristic 712 from the compliance component library 716. The modifier 726 is then operable to augment the deployment specification 704 for the application 702 by modifying the deployment specification 704 such that the deployment specification 704 includes a specification for the deployment of the compliance components 7182, 7184 and
35 7186.

Figure 9 is a component diagram illustrating an augmented deployment specification 704' in accordance with the exemplary embodiment of the present invention. As can be seen in Figure 9, the deployment descriptor 804 of the deployment specification 704 has been modified to produce a modified deployment descriptor 804'. The modified deployment descriptor 804' includes references to: the user access control evidence collector 7182 as new "User_access_control_evidence_collector" resource; the intrusion prevention system rule version checker 7184 as new "IPS_rule_version_checker" resource; and the cardholder data storage evidence collector 7186 as new "Cardholder_data_storage_evidence_collector" resource.

10 Additionally, the modified deployment descriptor 804' includes new data flows resulting from, or supporting, the operation of the compliance components 7182, 7184 and 7186. Data flow 906 between the "User_access_control_evidence_collector" resource and the "User_Access_Control" resource in the architecture specification 802' reflects communication between the user access control evidence collector 7182 and the user access control facility
15 provided by the hypervisor 7204 in the virtualised computing environment 710.

The new data flow 902 between the "IPS_rule_version_checker" resource and the "IPS" resource reflects communication between the intrusion prevention system version checker 7184 and the intrusion prevention system provided in the virtualised computing environment 710.

20 Further, the new data flows 904, 908 and 910 reflect communication between the cardholder data storage evidence collector 7186 and various other resources including: the "DB_MYSQL" database; the data flow between the "Web_Application" resource and the "Payment_Processor" resource; and the data flow between the "Payment_Processor" resource and the "DB_MYSQL" database. The new data flows 904, 908 and 910 thus reflect
25 evidence collection of the cardholder data storage evidence collector 7186.

Thus, in this way the compliance augmentation tool 720 is operable to modify the deployment specification 704 for the application 702 to incorporate compliance software components 7182, 7184 and 7186 based on the received compliance characteristic 712. The deployed software application 702 in operation at runtime is thus augmented by the selected
30 compliance software components 7182, 7184 and 7186 such that the compliance software components 7182, 7184 and 7186 function to evaluate or assess an extent or level of compliance of the software application 702 with the payment processor compliance characteristic 712. The inclusion of the compliance software components 7182, 7184 and 7186 as part of the deployment of the application 702 means that the extent or level of
35 compliance can be continually determined irrespective of whether the software application

702 and/or the virtualised computing environment 710 is adapted, redeployed, adjusted or otherwise changed at runtime, such as in response to changes in the operation of the application 702 or in response to service based facilities such as IaaS, PaaS, SaaS or cloud computing facilities.

5 Figure 10 is a component diagram illustrating a deployment of a software application 1000 with a virtualised computing environment 210 in accordance with a preferred embodiment of the present invention. The software application 1000 of Figure 10 includes a deployment specification 204 identifying resources 206 required for deployment of the application 1000 to the virtualised computing environment 210.

10 The deployed software application 1000' includes one or more resources 1022 operating with the virtualised computing environment 210. The deployed application 1000' has associated a compliance assessment component 1006 or compliance assessor. The compliance assessment component 1006 is a software or hardware component operable to determine a level of compliance of the deployed application 1000' based on at least a
15 compliance criterion 1014 and a compliance software component 1008. The compliance assessment component 1006 is executed, instantiated or otherwise deployed in conjunction with the deployed application 1000'. One way to deploy the compliance assessment component 1006 is to include an identifier of the component 1006 with the deployment specification 204 so as to cause the deployment of the compliance assessment component
20 1006 along with the application 1000. Alternatively the compliance assessment component 1006 can be predefined, predeployed, preinstalled or configurably installed, such as in association with a component of the virtualised computing environment 210 such as a hypervisor 1026 or operating system.

The compliance assessment component 1006 includes: an identifier 1030; a retriever
25 1032; a selector 1034; an evaluator 1036; and a resource change detector 1038. The identifier 1030 is a software or hardware component operable to identify resources 1022 instantiated for execution of the deployed application 1000'. For example, the identifier 1030 may receive the deployment specification 204 indicating the resources instantiated for the application 1000. Alternatively or additionally the identifier 1030 can monitor the deployed
30 application 1000', the virtualised computing environment 210 or the resources 1022 themselves to identify the resources 1022. In one embodiment, the identifier 1030 receives an indication of the resources deployed for the application 1000' from a component associated with the virtualised computing environment 210 such as a hypervisor component. In an alternative embodiment, the identifier 1030 interfaces with a component of the

virtualised computing environment 210 to identify the resources 1022 via in interface such as an API.

The retriever 1032 is a software or hardware component for retrieving a compliance characteristic 1012 for the deployed application 1000'. The retrieval of the compliance characteristic 1012 is based on the resources 1022 identified by the identifier 1030. In one embodiment, the retrieval of the compliance characteristic 1012 is pre-specified or predetermined on or before deployment of the application 1000'. For example, the compliance characteristic 1012 can be specified in a configuration of the deployed application 1000'. In an alternative embodiment the compliance characteristic 1012 is retrieved by the retriever 1032 using a compliance characteristic selector 512, 814 such as is described with respect to Figures 5 and 7.

The compliance characteristic 1012 has associated a compliance criterion 1014 being based on a formal parameter 1016. The formal parameter 1016 is a parameter required for an evaluation of the compliance criterion 1014. A data item, argument, or variable supplied to evaluate the compliance criterion 1014, such data item constituting the formal parameter 1016, is known as an actual parameter. A single compliance characteristic 1012 is illustrated in Figure 10, the compliance characteristic 1012 having a single compliance criterion 1014 with a single formal parameter 1016. This representation is chosen for simplicity though it will be appreciated that alternative embodiments can include any number of similar or disparate compliance characteristics each having potentially numerous compliance criteria, each criterion being based on potentially numerous formal parameters. Further, each criterion can specify dependencies between formal parameters. All such compliance characteristics 1012 are retrievable by the retriever 1032.

The selector 1034 is a software or hardware component for selecting a compliance software component 1008 for providing an actual parameter corresponding to the formal parameter 1016. The actual parameter can include or be based on, inter alia: data relating to, about or from one or more resources 1022; data concerning a state of one or more resources 1022; data indicating an occurrence of an event associated with one or more resources 1022; data including a measurement of a characteristic of one or more resources; or a transformation of data associated with one or more resources 1022. The compliance component 1008 contributes to a determination of a level or extent of compliance of the software application 1000 with a compliance characteristic 1012 by providing the actual parameter. The compliance component 1008 is executed, instantiated or otherwise deployed in conjunction with the deployed application 1000'. One way to deploy the compliance component 1008 is to include an identifier of the component 1008 with the deployment

specification 204 so as to cause the deployment of the compliance component 1008 along with the application 1000. Thus, in one embodiment, the deployment specification 204 is augmented by the inclusion of a compliance software component identifier, such as in accordance with one of the methods described hereinbefore with respect to Figures 2 to 9.

5 The inclusion of a compliance software component identifier in the deployment specification 204 is such that, on deployment of the software application 1000, compliance software component 1008 is deployed. Alternatively the compliance component 1008 can be predefined, predeployed, preinstalled or configurably installed, such as in association with a component of the virtualised computing environment 210 such as a hypervisor or operating
10 system.

Preferably the compliance component 1008 is one of a set of compliance components executable or executing in association with the virtualised computing environment 210. The selector 1034 is arranged so as to select a compliance component 1008 such that the compliance component 1008 is operable to provide an actual parameter for the compliance
15 criterion 1014. Thus the selector 1034 is operable to select a compliance component 1008 that is operable to access, obtain, retrieve or receive such data on which the actual parameter is based.

Notably, it is not a prerequisite for instantiation of the compliance component 1008 that the component is identified in the deployment specification 204. Rather, the compliance
20 component 1008 can be deployed by default, by design, as a consequence of the deployment of a resource identified for the application 1000 or otherwise automatically. In one embodiment, the compliance assessment component 1006 is deployed along with the application 1000 and the compliance assessment component 1006 is operable to cause the deployment of the compliance component 1008.

25 Most preferably the compliance assessment component 1006 and the compliance software component 1008 execute with the deployed application 1000' in a trusted mode of operation such that the compliance assessment component 1006 and the compliance software component 1008 have trusted access to aspects of the deployed application 1000'. Such aspects can include: configuration information; interfaces; technologies; configuration
30 information and data flows. Examples of interfaces include logical or software interfaces such as APIs of any or all of the resources 1022 instantiated for the deployed application 1000' or any other component operable with, or as part of, the deployed application 1000'. Examples of technologies include technical components such as software components provided by software suppliers or service providers and providing functions or services such that the
35 compliance component 1008 can request or retrieve information or functions of the

components. Examples include components, or providers of components, for intrusion prevention, virus detection, middleware or databases. Typically such technologies are uniquely identifiable such as by a version of the technology.

Compliance software component 1008 enjoys a sufficient level of trust that it is able to
5 retrieve, obtain, receive or access information or functionality of resources in order to provide the actual parameter. Thus, where the compliance characteristic 1012 relates only to a single resource required for the deployment of the application 1000, then trusted access to the single resource may be sufficient. However, it will be apparent to those skilled in the art that trusted access to resources other than a resource to which the compliance characteristic
10 1012 explicitly relates may be required to provide the actual parameter.

The evaluator 1036 is a software or hardware component operable to evaluate the compliance criterion 1014 using the actual parameter supplied by the compliance component 1008. Such evaluation is suitable for contributing to a determination of a level or extent of compliance of the deployed application 1000' with the compliance characteristic 1012.

15 The resource change detector 1038 is a software or hardware component operable to detect a change to the resources 1022 instantiated for the deployed application 1000'. Changes to resources can arise numerous including, inter alia: changes to the configuration of a resource by another resource, component or an operator; changes to the configuration of the virtualised computing environment 210; upgrades to a resource; failure of
20 a resource; addition of a new resource; changes to the software application 1000; redeployment of the software application 1000; and reprovisioning of a service based environment provided for the deployed application 1000'. Such reprovisioning is common with cloud computing services, IaaS, PaaS and SaaS environments and can arise in response to a change in the resource requirements of the deployed application 1000' at
25 runtime. For example, the resource demands of the deployed application 1000' can vary based on usage of the application 1000' or throughput of the application 1000'. For example, software applications providing web-based services receiving and reacting to requests received over a network can see a rate of receipt of requests fluctuate over time. Accordingly, a cloud computing service provider may change the resource provisions
30 allocated to such an application in response to fluctuations of resource requirements resulting from such fluctuations in requests. This contributes to the elasticity of such service based environments. The resource change detector can detect changes to the resource instantiated for the application 1000' in numerous ways including, inter alia: the obtaining and monitoring of profiles of resources such as process monitoring; hardware resource
35 monitoring; resource consumption; and configuration settings monitoring. Further, changes to

resources can be flagged by the virtualised computing environment 210 or other service based environment such as via an indicator, notification, message or otherwise to indicate a resource change. In one embodiment, the resource change detector 1038 is operable in conjunction with the identifier 1030 to identify a change in resources 1022 instantiated for the
5 deployed application 1000'.

A single compliance component 1008 is illustrated in Figure 10 for simplicity. It will be appreciated that alternative embodiments can employ multiple and potentially disparate compliance components. Multiple compliance components can be employed such that compliance component 1008 or compliance assessment component 1006 further select
10 other compliance components to obtain information required to supply the actual parameters. Thus, compliance components can be organised in a network, hierarchy, or other suitable arrangement such that information required to evaluate the compliance criterion 1014 can be obtained.

While the identifier 1030, retriever 1032, selector 1034, evaluator 1036 and resource
15 change detector 1038 are illustrated as being comprised with the compliance assessment component 1006 it will be apparent to those skilled in the art that any or all of these components could be alternatively provided as a separate component, or part of a separate component, external to and operable in association with the compliance assessment component 1006. Further, while the compliance assessment component 1006 is illustrated
20 as being partly comprised within the virtualised computing environment 210 it will be appreciated by those skilled in the art that the compliance assessment component 1006 could equally be implemented entirely within the virtualised computing environment 210; or alternatively the compliance assessment component 1006 could be implemented external to the virtualised computing environment 210 and associated with the deployed application
25 1000' such as being operable in communication with the deployed application 1000' via software components, a software interface, a network or any suitable communication means.

Figure 11 is a component diagram of a plurality of compliance components in accordance with a preferred embodiment of the present invention. In the arrangement of Figure 11 compliance component 1008 selects further compliance components 1102 and 1104.
30 Compliance component 1104 further selects compliance component 1106. The additional compliance components 1102 to 1106 can be instantiated as a result of augmentation of the deployment descriptor 204 for the application 1000. Alternatively, the compliance component 1102 to 1106 can be instantiated dynamically at runtime, automatically in association with any of the resources of the deployed application 1000', or in response to instantiation
35 requests by the compliance assessment component 1006 or other instantiated compliance

components. The compliance component 1008 of Figure 11 selects compliance components 1102 and 1104 to provide data to it, each supplying data constituting at least some of the data required to provide an actual parameter corresponding to the formal parameter 1016. Alternatively, compliance components 1102 and 1104 could be selected by the compliance assessment component 1006. An exploded view of an exemplary embodiment of compliance component 1008 is also illustrated in Figure 11. The compliance component 1008 includes: an identification 10082 of data provided by the compliance component 1008; an identification 10086 of data required by the compliance component 1008; and logic 10084 of the compliance component 1008. The identification 10082 of data provided by the compliance component 1008 is an identification of data that the compliance component 1008 can provide as an output, such as an output to another compliance component or to the compliance assessment component 1006. The identification 10082 can be, inter alia, an advertisement, a publication, a statement or a configuration setting indicating what type, class or category of data the compliance component 1008 is operable to provide. The indication 10086 of data required by the compliance component 1008 is an identification of data that the compliance component 1008 requires in order to generate the data provided by the compliance component 1008. The required data can be obtained from other compliance components, such as components 1102 and 1104 in Figure 11. Thus identification 10086 identifies prerequisite data for the compliance component 1008. Logic 10084 can include functionality and operations performed by the compliance component 1008 including, inter alia: accessing, retrieving or receiving data from resources of the deployed application 1000'; interface operations for cooperating with resources over an API; measurement logic for measuring characteristics of resources; modification or transformation logic to modify or transform data; logic to combine, fuse or integrate data or information; and logic suitable for identifying patterns, themes or characteristics from data or information. Such data or information can include data received from a resource, data received from another compliance component or data resulting from a measurement operation.

This arrangement of the compliance component 1008 is replicated across all compliance components to provide for the interoperation and cooperation of components in obtaining actual parameters required to evaluate the compliance criterion 1014. The selection of the compliance component 1008 by the compliance assessment component 1006 is based on the formal parameter 1016 such that the compliance component 1008 includes an identification 10082 of data it provides that is suitable for constituting an actual parameter corresponding to the formal parameter 1016.

In the exemplary embodiment, the identifications 10082 and 10086 for the compliance component 1008 and the formal parameter 1016 are specified using a common format

and/or namespace such that data provided by and required by compliance components can be compared with the formal parameter 1016. In this way it is possible for the compliance assessment component 1006 to select one or more appropriate compliance components to provide data required to evaluate the compliance criterion 1014. Further, it is possible for
5 each compliance component to select further compliance components to provide any required prerequisite data. The common format and/or namespace can be organised in a hierarchy or network such that prerequisite data requirements can be discerned from the namespace.

While the compliance software components 1102 to 1106 are described as software
10 components it will be appreciated by those skilled in the art that any or all of compliance component 1102 to 1106 could be implemented in software, hardware, firmware or combinations of any of software, hardware and firmware. For example, each of the compliance software components 1102 to 1106 can be implemented as a hardware component such as an evaluator component operable to perform the function of a
15 compliance software component.

Figure 12 is a flowchart of a method of the compliance assessment component 1006 in accordance with a preferred embodiment of the present invention. At step 1202 the identifier 1030 identifies resources 1022 instantiated for execution of the application 1000'. Such an identification of resources 1022 can be determined based on, inter alia: configuration
20 information for the virtualised computing environment 210; processes and services executing in the virtualised computing environment 210 identified using a process monitoring tool, a process and/or service registry and the like; referring to software components operable to interrogate resources for the application 1000'; accessing resource information via an API of one or more resources 1022; and other techniques as will be apparent to those skilled in the
25 art. At step 1204 the retriever 1032 retrieves a compliance characteristic 1012 for the application. The retrieval 1204 is based on the resources identified at step 1202. Compliance characteristics can be associated with resources 1022 such as by way of a compliance characteristic dictionary 816 as is illustrated in Figure 7. Alternatively, associations between resources and compliance characteristics can be more complex such as: rule-based
30 associations depending on multiple resources; associations based on attributes or characteristics of resources such as configurations, settings and or arrangements of resources; associations based on versions of resources; and other associations as will be apparent to those skilled in the art. The retrieved compliance characteristic 1012 has associated the compliance criterion 1014 based on the formal parameter 1016.
35 Subsequently, at step 1206, the selector 1034 selects a compliance software component 1008 to provide an actual parameter corresponding to the formal parameter 1016. The actual

parameter is based on data concerning at least one of the resources 1022 such that the compliance criterion 1014 can be evaluated. The selection of the compliance component 1008 is based on an identification, by the compliance component 1008, of one or more data items 10082 that the compliance component 1008 is operable to provide. At step 1208 the
5 evaluator 1036 evaluates the compliance criterion 1014 using the actual parameter. The evaluation contributes to a determination of a level of compliance of the deployed application 1000'. At step 1210 the resource change detector 1038 determines if one or more resources 1022 instantiated for the software application 1000' is changed. Where a resource 1022 is changed, the method returns to step 1202 to repeat the method steps 1202, 1204, 1206 and
10 1208. In one embodiment, step 1204 is not repeated following a positive determination at step 1210 and the compliance characteristic 1012 from a previous iteration of the method is retained.

Thus embodiments of the present invention provide a separation of concerns between a compliance assessment component 1006 and a compliance software component 1008. Such
15 separation is advantageous where the resources for the deployed application 1000' can change at runtime, such as due to deployment of the application 1000' using a service based environment such as a cloud computing environment. In particular, the software component 1008 is selected to provide the actual parameter such that the selection of an appropriate software component is based on the data requirements for evaluating the compliance
20 criterion 1014. Accordingly, where one or more of the resources 1022 changes, the selection of a software component can result in a different software component able to provide the actual parameter for the changed application. Thus the separation of concerns between the compliance assessment component 1006 and the software component 1008 provides for the selection of an appropriate software component based on the data requirements for
25 evaluating the criterion 1014 and the resources 1022 instantiated for the deployed application 1000'.

Embodiments of the invention thus provide an adaptable approach to compliance assessment for software applications executing with service based infrastructures where resources can change at runtime, such as in response to platform or infrastructure
30 reprovisioning, or where a platform or infrastructure exhibits characteristics of resource elasticity as is typical in cloud computing environments. Embodiments of the present invention further provide for such compliance assessment without a need to interrupt or redeploy the software application, or redeploy a compliance architecture.

Figure 13 is a schematic illustration of an arrangement for determining a level of
35 compliance of the software application 1000' with a compliance characteristic 1312 in

accordance with a preferred embodiment of the present invention. The compliance characteristic 1312 includes two compliance criteria 1314a and 1314b being expressed in simplified form for ease of understanding. Compliance criterion 1314a is based on a formal parameter "a" 1316a. Compliance criterion 1314b is based on a formal parameter "b" 1316b.

5 A compliance assessment component 1306 is operable to determine a level of compliance of a software application 1000' with the compliance characteristic 1312. In the embodiment of Figure 13 the compliance assessment component 1306 achieves this determination by selecting compliance software components 1308a and 1308b as "criterion tester" components operable to evaluate the compliance criteria 1314a and 1314b
10 respectively. In an alternative embodiment the compliance assessment component 1306 is operable to test the criteria 1314a and 1314b itself, based on data provided by other compliance software components.

Compliance components 1308a and 1308b advertise their ability to provide "criteria satisfaction indicators" as output data items. Compliance component 1308a includes an
15 identification of required data indicating that the component 1308a requires actual parameter data corresponding to parameter "a" 1316a. Compliance component 1308b includes an identification of required data indicating that the component 1308b requires actual parameter data corresponding to parameters "b" 1316b and "c" 1316c. Compliance component 1308a achieves its purpose by selecting a further compliance component 1308c, a "data
20 transformer" compliance component. Component 1308c advertises its ability to provide actual parameter data corresponding to parameter "a" 1316a. Component 1308c further indicates its dependency on data indicated as "raw data (a)". To satisfy this dependency, component 1308c selects compliance component 1308e, a "data collector" compliance component. Component 1308e advertises its ability to provide data as "raw data (a)". Data
25 collector component 1308e is operable to interface with one or more resources in the deployed application 1000' to access the raw data. For example, data collector 1308e can access a resource using an API for the resource, or by intervening in a data flow, or any other suitable access mechanism.

Compliance component 1308b achieves its purpose by obtaining actual parameter data
30 corresponding to parameter "b" 1316b by selecting compliance component 1308f, an "event detector" compliance component. Component 1308f advertises its ability to provide actual parameter data corresponding to parameter "b" 1316b. Event detector component 1308f is operable to interface with one or more resources in the deployed application 1000' to detect events, generating actual parameter data corresponding to parameter "b" 1316b.

Compliance component 1308b further achieves its purpose by obtaining actual parameter data corresponding to parameter "c" 1316c by selecting compliance component 1308d, a "data transformer" compliance component. Component 1308d advertises its ability to provide actual parameter data corresponding to parameter "c" 1316c. Component 1308d
5 further indicates its dependency on data indicated as "raw data (c)". To satisfy this dependency, component 1308d selects compliance component 1308g, a "data collector" compliance component. Component 1308e advertises its ability to provide data as "raw data (c)". Data collector component 1308g is operable to interface with one or more resources in the deployed application 1000' to access the raw data, such as is described above with
10 respect to component 1308e.

Thus, each compliance component 1308a to 1308d can provide further information by supplementing, adapting, processing, verifying or reacting to the data from downstream components. In this way it is possible to separate the concerns of the compliance components 1308a to 1308g. Such separation is advantageous when information from
15 multiple information sources is required to determine a level or extent of compliance with a compliance characteristic 1312. For example, different compliance software components can enjoy different privileges in relation to a deployed application such that one compliance software component may have trusted access to resources that another compliance software component does not have. Further, complex deployed applications can have associated
20 many and varied compliance characteristics, each having potentially many and varied compliance criteria. Such criteria can relate to numerous and differing resources required for application deployment, with the differing resources having associated information in a multiplicity of forms. Where there are overlaps in information requirements to assess a level or extent of compliance with multiple compliance characteristics it is advantageous to
25 centralise data gathering for a resource such that any duplication in the retrieving or obtaining of data for assessing compliance criteria is reduced. Further, it is advantageous to distribute responsibility for information collection between compliance software components which can specialise in, dedicate to, relate to or associate with particular resources, data formats, information types, information gathering methods or other variable attributes for a
30 deployed application. Such distribution reduces a degree of coupling in the compliance determination methods and systems and further provides for a granular approach to information gathering.

The approach to determining a level of compliance described with reference to the preferred embodiments is particularly advantageous in service based software environments
35 such as cloud computing environments. The elasticity of such service based technologies can result in adaptations or modifications to the resources employed in and for a deployed

application, including changes in real-time at runtime. Elasticity can also result in the supplementing of resources with additional resources or the replacement of resources with alternative or new resources. Such changes to the resources for a deployed application require repeat assessment of compliance characteristics to ensure a determination of an extent or level of compliance accurately reflects a current configuration of the application. This is particularly important where a particular minimum level of compliance is required for continuing operation of the deployed application such as, for example, to ensure a requisite level of security is provided. The selection of compliance components by a compliance assessment component and/or other compliance components can be undertaken dynamically at runtime. Accordingly, compliance components can change along with the resources for a deployed application.

Selection of, and communication between, compliance components such as components 1308a to 1308g can be achieved using any suitable mechanism known in the art including inter alia: a directory system; a publish-subscribe infrastructure; a request-response protocol; and a message passing scheme such as a brokered messaging infrastructure. In one example, the identifications of data provided by each compliance component can be stored in a directory accessible to other compliance components and/or the compliance assessment component such that when a compliance component is required for a particular data type, parameter or data item, identification of a suitable compliance component can be achieved by reference to the directory.

In a second example, a compliance component can advertise an identification of data it is capable of providing by publishing messages over a publish-subscribe infrastructure such that subscribing components, such as other compliance components or a compliance assessment components, are able to receive such publications by subscribing to receive such publications, such as by subscribing on a topic basis. A topic scheme can be devised, as is known in the art, whereby publications on a particular topic are related. One approach to implementing such a topic scheme uses an identification of a type of data from a global namespace of data types, such as an identification of a formal parameter, such that compliance components requiring data of that type can subscribe to publications on that topic.

In a third example, compliance components can communicate with each other directly or via a compliance assessment component using a predefined protocol such as a request-response protocol. Such a protocol can include a definition of messages for requesting an identification of data provided by a compliance component and requesting data itself. Using such a protocol, compliance components can form a compliance component network having

one of any number of potential topologies including, inter alia, hierarchical, star, tree, mesh or combinations thereof.

In a fourth example, compliance components can communicate with each other via a message passing scheme such as a brokered messaging infrastructure. Message broker
5 components are suitable for communicating messages between entities in connected networks of entities and can further adapt or translate messages where communicating components have different formats, styles or needs. Such messages can be used to communicate information about compliance components such as indications of data provided by components. Further, messages can be used to request and receive data from
10 components.

Thus, Figure 13 illustrates how the compliance components are operable to interoperate to provide potentially multiple layers of data abstraction and granularity, for example ranging from raw data to evidence about compliance criterion satisfaction; and/or multiple data collection or transformation components that enable, for example, the fusion, aggregation,
15 measurement, determination or derivation of data and/or evidence of compliance requirement satisfaction.

The identified extent or level of compliance is suitable for affecting the operation of the application when deployed and the configuration of a virtualised computing environment. For example, applications having a level of compliance that does not meet a predetermined
20 threshold can be precluded from all or some operations. Thus, in some embodiments the invention provides an access control function to allow or preclude access to the application or resources for a deployed application in response to an assessment of a level or extent of compliance. Further, in some embodiments the invention provides a compliance enforcement function where compliance requirements defining technical requirements of an application
25 are imposed on the application automatically at runtime of the application based on an assessment of a level or extent of compliance of the application according to embodiments of the present invention. Yet further, embodiments of the present invention can be operable to provide safety, security, reliability and/or stability features of an application by assessing a level or extent of compliance of the application with technical compliance requirements for
30 assuring a predefined level of safety, security, reliability and/or stability and indicating such level to inform a determination of future operation and/or to inform a compliance enforcement process. Thus applications that are safety critical, security critical or high-reliability critical can be monitored and affected using the approaches described with respect to embodiments of the present invention.

Insofar as embodiments of the invention described are implementable, at least in part, using a software-controlled programmable processing device, such as a microprocessor, digital signal processor or other processing device, data processing apparatus or system, it will be appreciated that a computer program for configuring a programmable device, apparatus or system to implement the foregoing described methods is envisaged as an aspect of the present invention. The computer program may be embodied as source code or undergo compilation for implementation on a processing device, apparatus or system or may be embodied as object code, for example.

Suitably, the computer program is stored on a carrier medium in machine or device readable form, for example in solid-state memory, magnetic memory such as disk or tape, optically or magneto-optically readable memory such as compact disk or digital versatile disk etc., and the processing device utilises the program or a part thereof to configure it for operation. The computer program may be supplied from a remote source embodied in a communications medium such as an electronic signal, radio frequency carrier wave or optical carrier wave. Such carrier media are also envisaged as aspects of the present invention.

It will be understood by those skilled in the art that, although the present invention has been described in relation to the above described example embodiments, the invention is not limited thereto and that there are many possible variations and modifications which fall within the scope of the invention.

The scope of the present invention includes any novel features or combination of features disclosed herein. The applicant hereby gives notice that new claims may be formulated to such features or combination of features during prosecution of this application or of any such further applications derived therefrom. In particular, with reference to the appended claims, features from dependent claims may be combined with those of the independent claims and features from respective independent claims may be combined in any appropriate manner and not merely in the specific combinations enumerated in the claims.

CLAIMS

1. A software compliance assessment apparatus for determining a level of compliance of a software application in execution in a virtualised computing environment, the apparatus
5 comprising:
 - an identifier adapted to identify resources instantiated for execution of the application;
 - a retriever adapted to retrieve a compliance characteristic for the application, the compliance characteristic being retrieved based on the identified resources, and the compliance characteristic having associated a compliance criterion based on a formal
10 parameter;
 - a selector adapted to select a software component for providing an actual parameter corresponding to the formal parameter, the actual parameter being based on data concerning at least one of the resources;
 - an evaluator adapted to evaluate the compliance criterion using the actual parameter;
 - 15 and
 - a detector adapted to detect a change to one or more of the resources, wherein the identifier, selector and evaluator are operable in response to a determination by the detector that one or more resources is changed,
wherein the selector selects the software component based on an identification of
20 one or more data items that the software component is operable to provide.
2. The apparatus of claim 1 wherein the software component is a first software component and the first software component is operable to select a second software component, the second software component providing at least some of the data concerning
25 the at least one of the resources.
3. The apparatus of any preceding claim wherein the data concerning the at least one of the resources includes at least one of: a measurement of a characteristic of the resource; an indicator of a state of the resource; and indicator of an occurrence of an event associated
30 with the resource.
4. The apparatus of any preceding claim wherein the actual parameter is determined based on a transformation of the data concerning the at least one of the resources.
- 35 5. The apparatus of any preceding claim wherein the compliance criterion is a first criterion in a set of multiple compliance criteria for the compliance characteristic, each of the multiple compliance criteria being based on a formal parameter,

wherein the selector is further adapted to select one or more additional software components for providing an actual parameter corresponding to a formal parameter for each other criterion in the set of multiple compliance criteria, and

5 wherein the evaluator is further adapted to evaluate each of the compliance criteria using a corresponding actual parameter to determine the level of compliance.

6. The apparatus of claim 1 wherein the software component provides an identification of one or more data items that the software component is operable to provide in response to a request for such identification.

10

7. The apparatus of claim 2 wherein the second software component provides an identification of one or more data items that the second software component is operable to provide, and the selection of the second software component is based on the identification by the second software component.

15

8. A method of a compliance assessment component for determining a level of compliance of a software application in execution in a virtualised computing environment, the method comprising:

20 identifying resources instantiated for execution of the application;
retrieving a compliance characteristic for the application, the compliance characteristic being retrieved based on the identified resources, and the compliance characteristic having associated a compliance criterion based on a formal parameter;

25 selecting a software component for providing an actual parameter corresponding to the formal parameter, the actual parameter being based on data concerning at least one of the resources;

evaluating the compliance criterion using the actual parameter; and
in response to a determination that one or more resources is changed, repeating at least the identifying, selecting and evaluating steps,

30 wherein the selection of the software component is based on an identification of one or more data items that the software component is operable to provide.

9. The method of claim 8 wherein the software component is a first software component and the first software component is operable to select a second software component, the second software component providing at least some of the data concerning the at least one
35 of the resources.

10. The method of any of claims 8 or 9 wherein the data concerning the at least one of the resources includes at least one of: a measurement of a characteristic of the resource; an indicator of a state of the resource; and an indicator of an occurrence of an event associated with the resource.

5

11. The method of any of claims 8 to 10 wherein the actual parameter is determined based on a transformation of the data concerning the at least one of the resources.

12. The method of any of claims 8 to 11 wherein the compliance criterion is a first
10 criterion in a set of multiple compliance criteria for the compliance characteristic, each of the multiple compliance criteria being based on a formal parameter,
wherein the selecting step further comprises selecting one or more additional software components for providing an actual parameter corresponding to a formal parameter for each other criterion in the set of multiple compliance criteria, and
15 wherein the evaluating step further comprises evaluating each of the compliance criteria using a corresponding actual parameter to determine the level of compliance.

13. The method of claim 8 wherein the software component provides an identification of one or more data items that the software component is operable to provide in response to a
20 request for such identification.

14. The method of claim 9 wherein the second software component provides an identification of one or more data items that the second software component is operable to provide, and the selection of the second software component is based on the identification by
25 the second software component.

15. A computer program element comprising computer program code to, when loaded into a computer system and executed thereon, cause the computer to perform the steps of a method as claimed in any of claims 8 to 14.

30

FIGURE 1




FIGURE 2




FIGURE 3




FIGURE 5

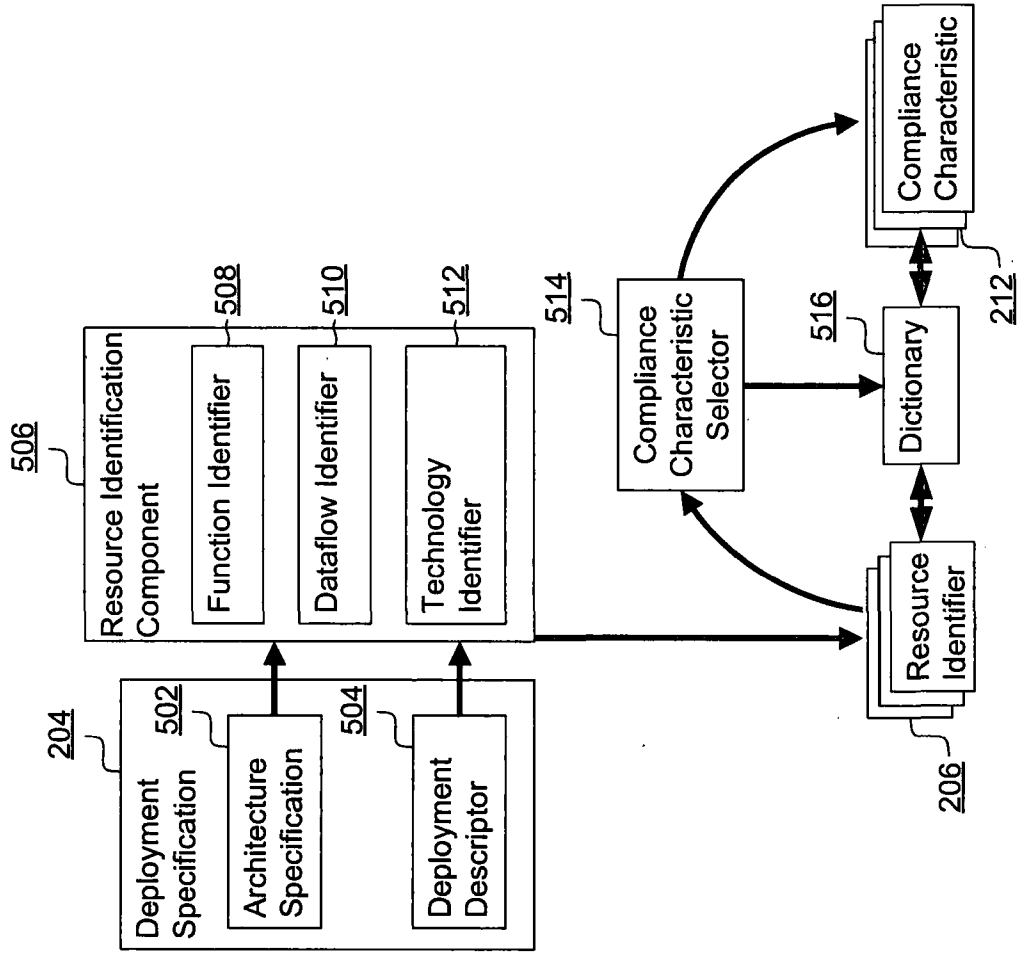


FIGURE 4





FIGURE 6a



7062	7064	7066	7068
Resource	Type	Identifier	Version
Technology	Hypervisor	Hypervisor1	Acme Visor 1.1 <u>960</u>
Technology	Operating System	OS-A	AcmeOS 3.0 <u>962</u>
Technology	Storage	Disk_Storage	<u>964</u>
Technology	Network	Network_Ethernet_TC P	<u>966</u>
Technology	Database	DB_MYSQL	MySql 5.6.11 <u>968</u>
Functional	Bespoke Database Schema	DB_Schema	<u>970</u>
Data Flow	Data Flow	DB_Schema ↔ DB_MYSQL	<u>868</u>
Technology	Payment Processing System	Payment_Processor	Acme PPS 2.1 <u>974</u>
Data Flow	Data Flow	Payment_Processor ↔DB_MYSQL	<u>860</u>
Technology	Web Server	Apache_WebServer	Apache HTTP Server 2.4.4 <u>978</u>
Functional	Bespoke Application	Web_Application	<u>980</u>
Data	Data	HTML_WebPages	<u>982</u>
Functional	Bespoke Application	PHP_Scripts	<u>984</u>
Technology	Application Software	Web_Commerce	OS Commerce 3.0.2 <u>986</u>
Data Flow	Data Flow	Apache_WebServer ↔Web_Application	<u>864</u>
Data Flow	Data Flow	DB_MYSQL ↔Web_Application	<u>862</u>
Data Flow	Data Flow	Payment_Processor ↔Web_Application	<u>866</u>
Technology	Application Software	IPS	Intrusion Protect 2.8 <u>994</u>

FIGURE 6b

FIGURE 7






FIGURE 8








FIGURE 10

FIGURE 12




FIGURE 11


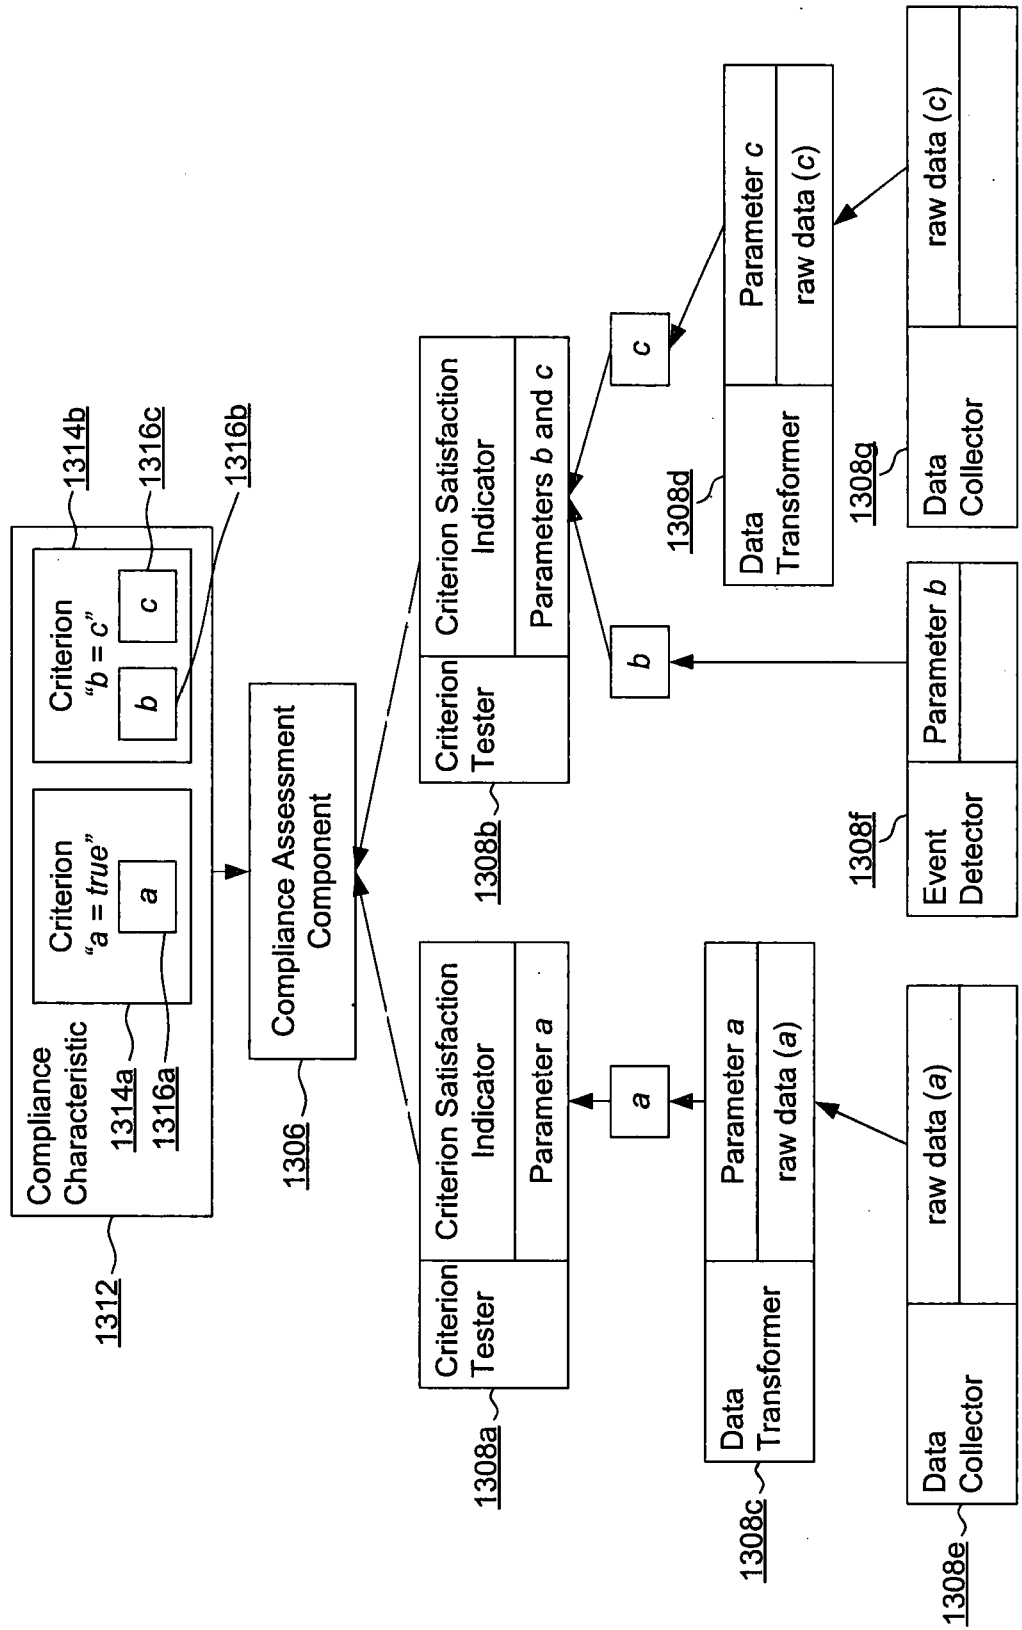


FIGURE 13



INTERNATIONAL SEARCH REPORT

International application No
PCT/GB2014/000231

A. CLASSIFICATION OF SUBJECT MATTER

INV. G06F9/44 G06F11/30
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, WPI Data, COMPENDEX, INSPEC, IBM-TDB

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2013/132933 A1 (RAJARAM SADAGOPAN [IN] ET AL) 23 May 2013 (2013-05-23) paragraph [0001] - paragraph [0005] paragraph [0039] paragraph [0020]	1-15
X	----- US 2008/222631 A1 (BHATIA RITU [US] ET AL) 11 September 2008 (2008-09-11) paragraph [0001] - paragraph [0050]	1-15
X	----- EP 2 568 383 A1 (ACCENTURE GLOBAL SERVICES LTD [IE]) 13 March 2013 (2013-03-13) paragraph [0001] - paragraph [0006] paragraph [0029] -----	1-15

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

21 July 2014

Date of mailing of the international search report

29/07/2014

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

Lo Turco, Salvatore

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No PCT/GB2014/000231
--

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
US 2013132933	A1	23-05-2013	NONE	

US 2008222631	A1	11-09-2008	NONE	

EP 2568383	A1	13-03-2013	CN 103002005 A	27-03-2013
			EP 2568383 A1	13-03-2013
			US 2013060933 A1	07-03-2013
			US 2014075005 A1	13-03-2014
