

(12) UK Patent Application (19) GB (11) 2519547 (13) A

(43) Date of A Publication

29.04.2015

(21) Application No: 1318778.6

(22) Date of Filing: 24.10.2013

(71) Applicant(s):
Eaton Industries France SAS
Immeuble Le Viséo -Bâtiment A Innovallé,
110 Rue Blaise Pascal, Montbonnot St. Martin 38334,
France (including Overseas Departments and Territories)

(72) Inventor(s):
Dominique Louis Lallement
Aurélien Louis Bégou
Emilien Benoit Kia

(74) Agent and/or Address for Service:
Eaton Industries GmbH
Eaton Patentabt./Morges, Hein-Moeller-Strasse 7-11,
Bonn 53115, Germany

(51) INT CL:
G06F 9/50 (2006.01)

(56) Documents Cited:
US 20130108259 A1 US 20110125894 A1
US 20110029673 A1 US 20080244595 A1
US 20080123559 A1 US 20040073673 A1

(58) Field of Search:
INT CL G06F
Other: EPODOC, WPI

(54) Title of the Invention: **Method of controlling a data centre architecture equipment**
Abstract Title: **Method of controlling a data centre architecture**

(57) A method of controlling a data centre architecture, with the data centre 10 comprising a plurality of devices 2 to 8 of different types, and a plurality of software applications 9 being executed on resources provided by the devices, comprises: acquiring information on links 11, 12, P between the devices, between the software applications, or between the devices and the software applications; processing this information to generate or update a model of the data centre architecture, where relations between devices, between software applications, or between devices and software applications are established; and implementing command sequences for controlling the devices or software applications, the sequences being set up on the basis of the generated data centre model. Acquired information may relate to network routing information. The model may include a virtualized architecture of the data centre, involving virtual machines V1, V2, V3. Software applications may include a power management tool, or infrastructure management application, or virtualization management application. Relations may be distinguished into compulsory relations and redundant relations. The method is intended to facilitate automatically discovering dynamic changes in the data centre architecture.

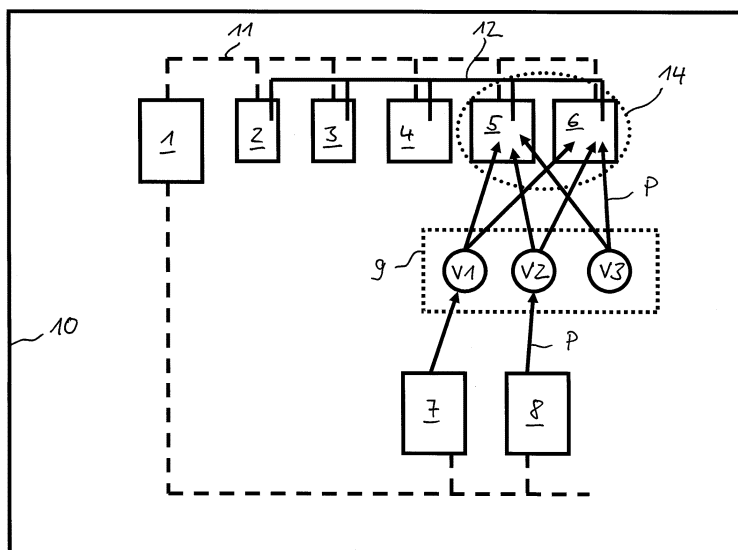


FIG.

GB 2519547 A

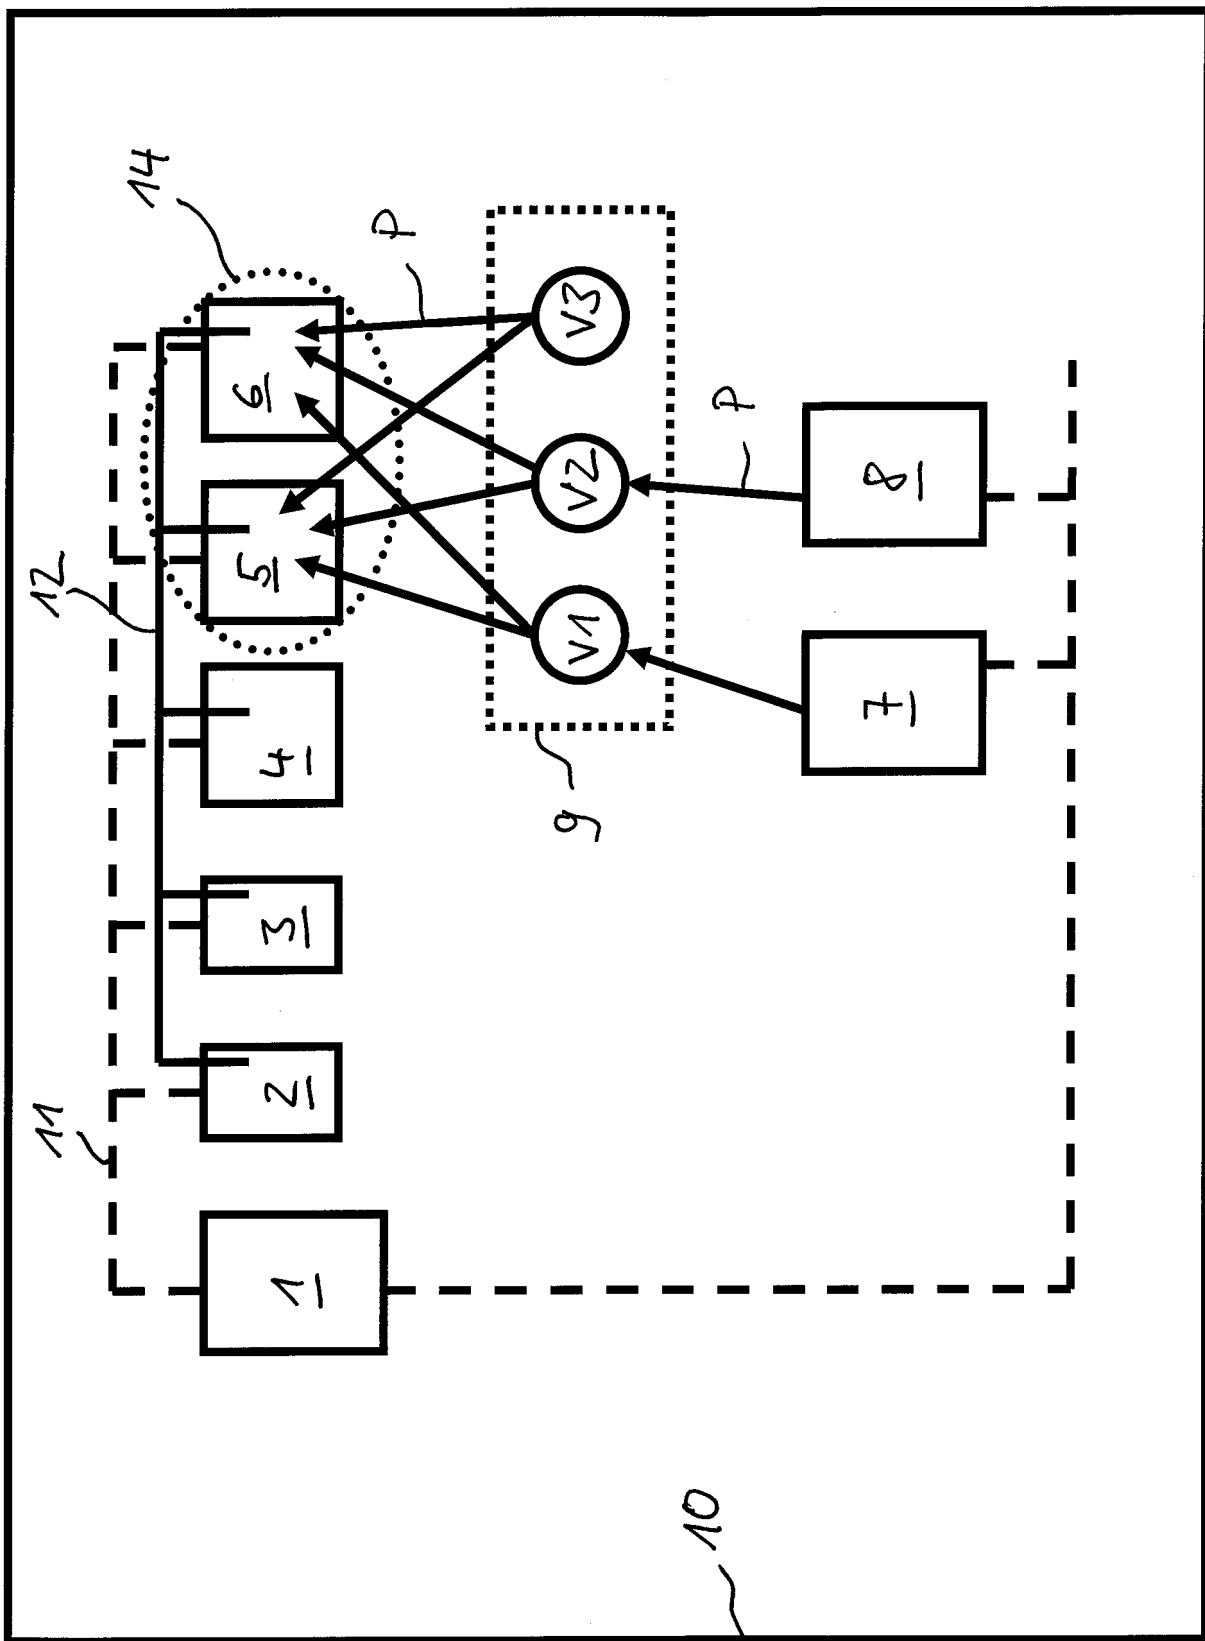


FIG.

5 **Method of controlling a data centre architecture equipment**

The present invention relates to a method of controlling a data centre architecture and to a data centre comprising a plurality of devices of different types, a plurality of software applications being executed on resources provided by the devices, wherein
10 a computer device, which is temporarily or permanently connected to the data centre is adapted to execute an application for controlling the data centre architecture.

Data centres are widely used to house various types of electrical equipment devices, including computer systems and the physical infrastructure needed to support such
15 systems, such as power supplies, including uninterruptible power supplies and backup power supplies, environmental systems, like air conditioning, fire suppression, etc., physical data centre security, and other monitoring devices. Companies that depend on the proper and efficient operation of their data centres use various tools to monitor and operate the physical infrastructure, including multiple monitoring
20 systems that are coordinated to provide centralized collection and reporting of critical infrastructure events. A method and apparatus for data centre automation is disclosed in WO 2011/031459 A2.

With the large number of mutually connected information technology equipment
25 devices, a problem arises if any action like implementing an individual setting or reconfiguration or individually controlling on an application or on a device is executed, as it is difficult to anticipate the consequences on other related applications or application parts or equipment devices. Even more difficult to anticipate are the consequences on other related applications or application parts or equipment
30 devices, if any action like a global setting or a set of settings or a set of reconfigurations or control commands is executed on an application or on a device. A lack of awareness about dependencies between different parts or devices in a data centre does not permit to define the proper set of related devices and related applications in order to optimize a command sequence. User defined static or manually defined

scripts or sequences are error-prone, because there is no automatic discovery of dynamic changes in the data centre architecture.

5 It is an objective of the present invention to provide a data centre architecture control that is capable of automatically discovering dynamic changes in the data centre architecture.

10 The objective is achieved by a method of controlling a data centre architecture according to the independent claim. The objective is further achieved by a data centre, wherein a computer device, which is temporarily or permanently connected to the data centre, is adapted to execute an application for controlling the data centre architecture according to the method according to the invention. Preferred embodiments are subject to the independent claims.

15 The method of controlling a data centre architecture according to the invention is applied in a data centre comprising a plurality of devices of different types, in particular information technology equipment, like servers, network devices and also electrical and mechanicals equipment, like UPS, PDU, cabinets with related environment devices, like thermal conditioning devices. The devices are connected
20 by one or more networks. A plurality of software applications is executed on resources provided by the devices, in particular by the servers. The method comprises the steps of

- a) acquiring information on links between the devices, between the software applications and/or between the devices and the software applications;
- 25 b) processing the information acquired in step a) to generate and/or update a model of the data centre architecture, wherein relations between devices, between software applications and/or between devices and software applications are established;
- c) implementing command sequences for controlling the devices and/or software
30 applications, the sequences being set up on the basis of the data centre model generated in step b).

Links in the sense of the invention are any kind of information about an attribution of a device to another device, or of a device to an application, or of an application to

another application, independent from a quality of the attribution. The generated model reflects the acquired information on the links, supplemented, if applicable, by additional information about the link, in the sense of an information on a quality of the respective link, which is expressed by the expression relation, which is used after the
5 acquired information on a link has been processed in step b).

It is an advantage of the present invention that the relations between the different devices and applications are automatically discovered and updated. It is avoided to set the same information many times, as the existing information is reused by
10 applying the model for executing any actions on the data centre. The information available from the devices and software applications is advantageously acquired automatically, which information provides, in general, an identification of each device and application, but also information on the links between devices and/or applica-
15 tions. Any sequence of commands to execute an action on the data centre may advantageously be implemented considering the model with the relations, which allows to exactly anticipate the consequences incurred on related devices and/or applications of the data centre.

According to a preferred embodiment, a plurality of different protocols is applied to
20 acquire the information on links according to step a). In particular, the plurality of different protocols includes protocols related to different abstraction layers, with respect to an OSI model. Advantageously, links will be discovered on any layer of communication perused in the data centre, from a hardware layer to an application layer.

25 According to a further preferred embodiment, acquiring the links according to step a) implies an identification of the devices on a basis of characteristic data related to each device. The identification of a device, in particular the detection of a type of device, provides important information for the analysis of the links of the device in
30 order to provide a qualified statement on the relation of the device to other devices and/or applications.

The plurality of devices will generally include at least one network device, like a router etc. and acquiring the links according to step a) thus preferably implies

acquiring network routing information from each of the present network devices, which, for example are stored as routing tables containing information on a plurality of links established via the respective network device.

- 5 Where the plurality of software applications includes a power management tool, acquiring the links according to step a) preferably implies acquiring manually defined or automatically discovered power links between the devices, which are handled by the power management tool.
- 10 Where the plurality of software applications includes an infrastructure management application, acquiring the links according to step a) preferably implies acquiring information about locations of the devices, power cable connections and/or network cable connections managed by the infrastructure application.
- 15 According to a furthermore preferred embodiment, the data centre architecture model generated according to step b) includes a virtualized architecture of the data centre, wherein the devices include a plurality of physical servers, wherein each server comprises one or more server resources allocable to one or more virtual machines on each server and wherein each virtual machine handles operations for a different
- 20 one of the plurality of software applications.

The plurality of software applications will generally include at least one virtualization management application to control a resource allocation of the one or more resources to the one or more virtual machines running on each server and wherein

25 acquiring the links according to step a) preferably implies acquiring a list of the virtual machines with the related resources.

Furthermore preferred, acquiring the links according to step a) implies acquiring a definition and/or composition of at least one set of resources being logically linked to

30 a cluster by the virtualization management application.

According to a furthermore preferred embodiment, the relations between devices, between software applications and/or between devices and software applications established in step b) are distinguished into compulsory relations and redundant

relations. The relations reflected by the model may on the one hand imply that a set of devices are necessary to accomplish and support one application, which may be expressed as logical “AND” relations. The relations may on the other hand imply that parts of the devices are redundant and that a subset of such devices is sufficient to accomplish and support one particular application, which may be expressed by logical “OR” relations. Examples of an “OR” relation may, for example, be found in server clusters and in parallel configuration of UPS devices.

Preferably, the steps a), b) and c) of the method according to the invention are repeated regularly, in particular, to initially generate the model of the data centre and to subsequently update the model.

Subject to the present information is further a data centre comprising a plurality of devices of different types, a plurality of software applications being executed on resources provided by the devices, wherein a computer device temporarily or permanently connected to the data centre is adapted to execute an application for controlling the data centre architecture using a method according to the invention as described here above.

The invention is now further described with respect to an example of an application of the method according to the invention. The description refers both to the method and the data centre according to the application. It is only exemplary and does not limit the scope of the invention.

The only Figure schematically illustrates an embodiment of a data centre according to the present invention.

With regard to the Figure, the steps of the method according to the invention will be illustrated with respect to a data centre 10. The data centre 10 comprises exemplary devices 1-8 and exemplary network connections 11, 12, a software application 9 being executed on a plurality of virtual machines V1, V2, V3, the virtual machines using resources of the devices 5, 6 and receiving information from devices 7, 8, which is illustrated by arrows, of which only two are denoted with the reference P.

The first step of the method according to the invention (step a) is, to acquire information about the links between the devices, between the software applications and/or between the devices and the software applications, preferably through
5 different communication protocols. A power equipment device 1 is powering the devices 2-8 via a power connection 11, depicted as a broken line. A network device 3 is connected to devices 2, 4, 5, 6 by a network communication connection 12, which will usually use a different communication protocol than the power connection 11, for example. Using different communication protocols is thus advantageous to detect
10 and acquire as many links of the data centre 10 as possible.

The device 2, for example, could be a computer device 2, which is temporarily or permanently connected to the data centre 10, where the method of the invention is executed in the form of a software application. The second step (step b) of the
15 method is, to process the information acquired in step a) to generate and/or update a model of the architecture of the data centre 10, wherein relations between devices, between software applications and/or between devices and software applications are established, i.e. to compute the information about the links in order to create and update the model and to manage the model. The functionality of step b) is to
20 recognize the relations between the devices and application by logically analyzing the acquired links. The relations of the exemplary data centre 10 are, for example:

- The power equipment device 1 is powering devices 2, 3, 4, 5 and 6.
- The network equipment device 3 provides a network link 12 to devices 4, 5, and 6, which have been identified as servers.
- 25 - The application 9 is composed of virtual machines V1, V2, V3 running in servers 5 and 6.
- The virtual machines V1 and V2 are using devices 7 and 8, which have been identified as data servers 7, 8.
- The servers 5,6 are part of a common server cluster 14.

30 The model is constituted on the basis of such relations. A part of the relations imply that a set of devices are necessary to accomplish and support one application in the sense of a compulsory relation or a logical "AND" relation. Another part of the relations also imply that some parts of the set of devices are redundant and that a subset of such devices is sufficient to accomplish and support one particular

application in the sense of logical “OR” relations. An example of a redundant relation can be found in the server cluster 14 in the exemplary data centre 10 or in parallel configuration of UPS devices (uninterruptible power supply, not depicted).

- 5 The third step of the method according to the invention (step c) is, to use the model in order to realize appropriate sequences, i.e. to implement command sequences for controlling the devices 1-8 and/or software applications 9, the sequences being set up on the basis of the data centre model generated in step b). There are many ways to advantageously use the relations of the model. For example, if it is desired to
- 10 execute a complete shutdown sequence, a shutdown command sequence will be defined, based on the relations. Referring to the exemplary data centre 10 and the relations described above, for example, as the following sequence will be executed:
- Shutdown virtual machines V1 and V2.
 - Shutdown data servers 7 8, unless they are used by other virtual machines.
 - 15 - Shutdown servers 5 and 6, unless they are used by other virtual machines.
 - Shutdown network device 3, unless it provides network to other devices than 7, 8, 5 and 6.
 - Shutdown power device 1, unless it provides power to other equipment's than 7, 8, 5 and 6

20

The three steps a), b) and c) of the method according to the invention are preferably executed regularly in order to update the relations of the model and the command sequences. The method of the invention will advantageously support data centre operators or automatic management systems to determine the appropriate set of

25 settings and the appropriate sequence of control commands that permit safe and coherent behaviour of the applications and devices.

As an example of setting, the operator or automatic setting system may need to determine how much time of power backup provided by UPS is necessary for each

30 server or for other IT devices, according to their relations with a certain application.

As an example of control commands, the operator or automatic protection system needs to define the proper command sequence to stop or restart certain applications and devices.

As an example of reconfiguration, the operator or automatic power management system needs to reconfigure the distribution of parts of the application, i.e. a set or subset of virtual machines, through different servers in order to optimize an energy consumption or a reliability of the application.

Reference numerals

1	Power supply device
2	Computer device
3	Network device
4-6	Server
7, 8	Data server
9	Software application
10	Data centre
11	Power supply line
12	Communication network
14	Cluster
V1, V2, V3	Virtual machines
P	Arrow

Claims

1. Method of controlling a data centre architecture, the data centre (10) comprising a plurality of devices (2-8) of different types, a plurality of software applications (9) being executed on resources provided by the devices, the method comprising the steps of
 - a) acquiring information on links between the devices, between the software applications and/or between the devices and the software applications;
 - b) processing the information acquired in step a) to generate and/or update a model of the data centre architecture, wherein relations between devices, between software applications and/or between devices and software applications are established;
 - c) implementing command sequences for controlling the devices and/or software applications, the sequences being set up on the basis of the data centre model generated in step b).
2. Method according to claim 1, wherein a plurality of different protocols is applied to acquire the information on links according to step a).
3. Method according to claim 2, wherein the plurality of different protocols includes protocols related to different abstraction layers.
4. Method according to one of the preceding claims, wherein acquiring the links according to step a) implies an identification of the devices on a basis of characteristic data related to each device.
5. Method according to one of the preceding claims, wherein the plurality of devices includes at least one network device (3) and wherein acquiring the links according to step a) implies acquiring network routing information from each network device.

6. Method according to one of the preceding claims, wherein the plurality of software applications includes a power management tool and wherein acquiring the links according to step a) implies acquiring manually defined or automatically discovered power links between the devices.
7. Method according to one of the preceding claims, wherein the plurality of software applications includes an infrastructure management application and wherein acquiring the links according to step a) implies acquiring information about locations of the devices, power cable connections and/or network cable connections.
8. Method according to one of the preceding claims, wherein the data centre architecture model generated according to step b) includes a virtualized architecture of the data centre, wherein the devices (2-8) include a plurality of physical servers (4-6), wherein each server comprises one or more server resources allocable to one or more virtual machines (V1, V2, V3) on each server, wherein each virtual machine handles operations for a different one of the plurality of software applications.
9. Method according to claim 8, wherein the plurality of software applications includes at least one virtualization management application to control a resource allocation of the one or more resources to the one or more virtual machines running on each server and wherein acquiring the links according to step a) implies acquiring a list of the virtual machines with the related resources.

10. Method according to claim 9, wherein acquiring the links according to step a) implies acquiring a definition and/or composition of at least one set of resources (5, 6) being logically linked to a cluster (14) by the virtualization management application.
11. Method according to one of the preceding claims, wherein the relations between devices, between software applications and/or between devices and software applications established in step b) are distinguished into compulsory relations and redundant relations.
12. Method according to one of the preceding claims, wherein steps a), b) and c) are repeated regularly.
13. Data centre (10) comprising a plurality of devices (2-8) of different types, a plurality of software applications (9) being executed on resources provided by the devices, wherein a computer device (2) temporarily or permanently connected to the data centre is adapted to execute an application for controlling the data centre architecture using a method according to one of the preceding claims.



Application No: GB1318778.6

Examiner: Dr Christopher Batty

Claims searched: 1 to 13

Date of search: 26 March 2014

Patents Act 1977: Search Report under Section 17

Documents considered to be relevant:

Category	Relevant to claims	Identity of document and passage or figure of particular relevance
X	1 to 13	US2008/0244595 A1 (EILAM et al.) See paragraphs [0049] to [0052] for example.
X	1 to 13	US2008/0123559 A1 (HAVIV et al.) See paragraphs [0073] to [0076], and Figure 3, for example.
X	1 to 13	US2011/0125894 A1 (ANDERSON et al.) See paragraphs [0040], [0109] and [0110] for example.
X	1 to 13	US2011/0029673 A1 (JAISINGHANI) See claim 10 for example.
X	1 to 13	US2004/0073673 A1 (SANTOS et al.) See claim 1 for example.
X	1 to 13	US2013/0108259 A1 (SRINIVAS et al.) See claim 1 for example.

Categories:

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art.
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.

Field of Search:

Search of GB, EP, WO & US patent documents classified in the following areas of the UKC^X :

Worldwide search of patent documents classified in the following areas of the IPC

G06F

The following online and other databases have been used in the preparation of this search report

EPODOC, WPI



International Classification:

Subclass	Subgroup	Valid From
G06F	0009/50	01/01/2006