(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
27 December 2013 (27.12.2013)

WIPO | PCT

(10) International Publication Number
WO 2013/190180 A1

(51) International Patent Classification:
G06F 9/455 (2006.01) *G06F 15/16* (2006.01)
H04L 29/12 (2006.01)

(21) International Application Number:
PCT/FI2013/050664

(22) International Filing Date:
18 June 2013 (18.06.2013)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
20125680 20 June 2012 (20.06.2012) FI

(71) Applicant: **NIXU SOFTWARE OY** [FI/FI]; Keilaranta 15, FI-02150 Espoo (FI).

(72) Inventor: **HOLKKOLA, Juha**; c/o Nixu Software Oy, Keilaranta 15, FI-02150 Espoo (FI).

(74) Agent: **KOLSTER OY AB**; Iso Roobertinkatu 23, P.O. Box 148, FI-00121 Helsinki (FI).

(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, 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, 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).

Declarations under Rule 4.17:

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

— *of inventorship (Rule 4.17(iv))*

[Continued on next page]

(54) Title: METHOD AND APPARATUS FOR IP COMMISSIONING AND DECOM-MISSIONING IN ORCHESTRATED COMPUTING ENVIRONMENTS

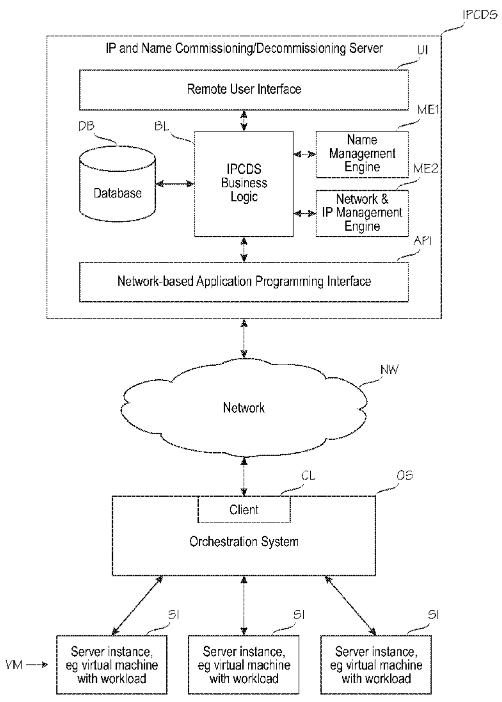

Fig. 1

(57) Abstract: A server computer (IPCDS) for commissioning/decommissioning IP resources to server instances (SI) provisioned using an orchestration solution. The server and client computers constitute a client-server architecture. The server computer comprises a processing system and memory for applications and data, including program code instructing the processing system to implement the following features: a user interface (UI) for remote management, providing access to data (DB) managed by the server computer; a web-based application programming interface (API) supporting service oriented architecture ["SOA"]; a logic to assign and release IP resources to the clients (CL) based on calls via the API. The logic creates unique IP resource(s) for the server instances (SI), which are nodes of networks with overlapping address spaces. The unique IP resources are based on a name of the private network of the respective server instance, and an IP address within that private network.

WO 2013/190180 A1

Published:

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

METHOD AND APPARATUS FOR IP COMMISSIONING AND DECOMMISSIONING IN ORCHESTRATED COMPUTING ENVIRONMENTS

PARENT CASE INFORMATION

[0001] The present application claims priority from Finnish patent application no 20125680, filed 20 June 2012.

FIELD OF THE INVENTION

[0002] The present invention relates to commissioning and decommissioning of IP resources in cloud computing environments.

BACKGROUND OF THE INVENTION

[0003] As has been known for a long time, Internet Protocol v. 4 (IPv4) is rather limited in terms of available address space. To address the issue, standard RFC1918 defines three networks intended for private use, namely 10.0.0.0 (Class A), 172.16.0.0 (Class B) and 192.168.0.0 (Class C). None of these private networks have been routed to the public internet. Large corporations and service providers typically have Class A network (10.0.0.0) address space to expand the address space available to them, while ADSL and cable modems commonly used in homes and small offices distribute IP addresses from private 192.168 networks. Connections to the outside world are provided by utilizing Network Address Translation (NAT) technology, wherein a NAT device located between the public and private networks acts as a bridge. Since several private networks share the same 10.0.0.0 address space, they are overlapping. Overlap has been an insignificant problem as long as these private networks have been run internally, instead of routing them to the public internet.

[0004] Overlapping of private networks becomes a problem in connection with cloud computing, or cloud-based services. For instance IaaS (Infrastructure as a Service) service providers are increasingly deploying multi-tenant computing environments used to offer services concurrently to several business clients, all of whom may use the same 10.0.0.0 address space especially in the context of Virtual Private Cloud and/or other similar technologies. In use-cases such as this, the private networks used by different tenants typically overlap.

[0005] In the following description, “orchestration” is used in its established meaning within Service Oriented Architecture (“SOA”) realm, when discussing automated workflows between data processing and communication systems. Enterprises and service providers use orchestration solutions to align business requests with the applications, data and infrastructure. The said solutions are typically used to define the policies and service levels through automated workflows, provisioning, and change management. With this technology, organizations are able to create an application-aligned infrastructure that can be scaled up, down or sideways based on the needs of each application. Orchestration also provides centralized management of the resource pool, including billing, metering, and chargeback for consumption.

[0006] Allocation of IP addresses, names and other network parameters to the various servers, commonly referred to as workloads, which execute applications in orchestrated environments, has traditionally been accomplished by configuring an IP address in the said servers and by adding the server's name with the corresponding IP address to a domain name server (DNS) manually, or by having such allocation performed dynamically using Dynamic Host Configuration Protocol (DHCP) and dynamic DNS. Since the IP addresses and the names of the physical servers run in traditional orchestrated computing environments have been relatively static, their SOA-based automated workflow management processes have not been extended to integrate with IP and name commissioning mechanisms. As existing orchestration solutions are expanded to cloud-based computing environments, the traditional methods used to manage IP addresses and names described above will create various problems. For instance, as cloud-based computing paradigm requires that new virtual machines are provisioned on demand, the manual IP and name assignment process associated with the prior art methods used for allocation IP resources and names in traditional orchestrated computing environments quickly become a bottleneck as far as the scalability of the entire cloud-based computing environment is concerned. In addition, although the cloud-based on-demand computing paradigm requires that the life-cycle of a virtual server instance can be anywhere from minutes to several years, DHCP servers provide a predefined and fixed lease time for the automatically assigned IP addresses, thereby making it impossible to align the IP lease times with the dynamic nature of the virtual computing environment. Furthermore, the prior art techniques make it impossible to automatically reclaim an IP address when a virtual machine is decommissioned, as even with DHCP, the decommissioning is tied to the pre-defined lease time of the IP addresses that have been issued. The prior art methods thus make it impossible to align the lease time of an IP address with the unique life-cycle of each virtual machine run within the cloud.

[0007] Limitations of DHCP are quickly revealed by attempting to use in connection with cloud computing. One of the reasons for the poor compatibility between DHCP and cloud computing is that DHCP was never designed for cloud computing or web-based integration models. For instance, DHCP operates on OSI Layer 2 (L2). In practice, a client sends a broadcast message to a local-area network (LAN). A DHCP server in that LAN captures the broadcast message, inspects the client's Medium Access Control (MAC) address, which is a unique address of the network interface adapter, and returns an IP address with other network parameters to the MAC address. After that the client configures the network parameters for itself and is able to adopt a TCP/IP connection, which operates on higher OSI layers.

[0008] In practice, the above-described methodology requires that the client and the DHCP server must be interconnected by an L2 connection. In practice, the client and the DHCP server must be connected to the same LAN network. The LAN may be comprised of multiple VLAN networks, but these must be interconnected on L2 layer. In cases where clients have overlapping 10.0.0.0 address spaces, the service provider must isolate them from one another by configuring the overlapping address spaces into distinct LAN networks. As a result, all private networks are isolated from one another, which enables IP traffic within the network on one hand and prevents clients from accessing the networks of other clients.

[0009] A consequence of the facts that, firstly, DHCP operates on L2 and, secondly, overlapping address spaces must be isolated into separate LANs, is that a single DHCP cannot logically reside in multiple LANs separately. In other words, in the case of multiple private networks, each of them must have a dedicated DHCP server.

[0010] Internet Protocol version 6 (IPv6) provides two mechanisms for dynamic IP allocations. These mechanisms are called Stateful autoconfiguration and Stateless autoconfiguration. Neither autoconfiguration mechanism solves the above-identified problems because, firstly, stateful autoconfiguration (DHCPv6) is not really any different from DHCPv4 used in IPv4 environments. This is because each time an IP resource is allocated to a virtual machine, the allocated IP resource obtains a fixed lease value, which basically means that the IP resource shall remain allocated for a predefined period of time, regardless of whether or not the allocated IP resource actually continues to be utilized by the virtual machine.

Within cloud-based environments, this is undesirable, because the IP addresses should be commissioned (issued) whenever a new virtual machine goes live, and decommissioned (released) whenever that virtual machine is removed from the virtualized computing environment.

[0011] On the other hand, stateless autoconfiguration means that a client autonomously obtains an IP address on the basis of router advertisements. As far as SOA architectures and orchestration are concerned, there are two reasons why this scheme may not work. First, in environments where orchestration is used, it is a typical requirement that the IP address is obtained from a network that matches the Virtual Local Area Network (VLAN) in which a virtual machine is intended to run. In other words, the IP address has to be allocated from a specific network corresponding to the VLAN in which the virtual machine is to be deployed, instead of giving the virtual machine an arbitrary IP address that happens to be available (the latter is what stateless autoconfiguration leads to). The second reason why stateless autoconfiguration may not work in this use case is that the environments are typically multi-tenant environments where the administrator has to be able to actively monitor allocation levels of each network and determine what equipments and clients are being run in the networks. In the event that the IP addresses are obtained autonomously by the clients, there is no way to control the IP addresses which a given virtual machine will have obtained, nor will there be any transparency to this process that would allow the administrator to manage these relationships and/or track the allocation of IP allocations.

DISCLOSURE OF THE INVENTION

[0012] An object of the present invention is thus to provide methods, equipments and computer program products so as to alleviate one or more of the problems identified above. The object of the invention is achieved by aspects of the inventions as defined in the attached independent claims. More specifically, the present invention provides a method, equipments and computer program products, which can be used to commission and decommission IP resources , such as IP addresses and names in a way that alleviates at least one of the two problems described above. The dependent claims and the following detailed description and drawings relate to specific embodiments which solve additional problems and/or provide additional benefits.

[0013] Aspects of the invention include a server equipment and client equipment which inter-operate according to principles of client-server architecture. In the present context, the inventive server is called an IP commissioning/decommissioning server apparatus ("IPCD server"). The IPCD server is used to manage and track, in real time, distribution and allocation of networks, the availability status of individual IP addresses within each network, related administrative properties and attributes. In a typical implementation, the IPCD server comprises the following elements:

- a remote user interface, such as a web-based user interface, which can be used to access the data managed in the IP commissioning/decommissioning server;
- a web-based Application Programming Interface (API) that supports Service Oriented Architecture (SOA) plus a logic capable of dynamically assigning and releasing IP resources, such as IP addresses, names, and other network configurations based on calls received via the API;
- another logic configured to associate individual networks with given internal and/or external end-users and, in the event that a network has been fully allocated, the IP commissioning/decommissioning server is configured to automatically provision an IP resource, (eg an IP address) from an alternate network associated with the internal and/or external end-user (client) and/or from a reserve pool of IP resources;
- a client-server architecture that involves a client run in the orchestration solution communicating with the IP commissioning/ decommissioning server over the SOA-based API.

[0014] In a typical implementation, the IPCD client equipment comprises a client logic configured to perform the following tasks:

- requesting an IP resource, such as an IP address and/or a name, from the IP commissioning/decommissioning server when a new virtual machine is commissioned to the cloud computing environment using the orchestration solution; and
- triggering the release of the assigned IP resources, such as the IP address and/or the name of a virtual machine that has been decommissioned.

[0015] Other aspects of the invention include methods for operating the inventive IPCD server and IPCD client, and computer program products whose execution in appropriately equipped server and client computers provide them with features of the IPCD server and IPCD client and thus cause the IPCD server and IPCD client to carry out the inventive method.

[0016] Unlike earlier attempts to couple DHCP with overlapping private networks, the present invention does not provision network parameters on OSI Layer 2 (L2). Instead the calls for requesting and assigning IP resources, communicated over the web-based API, operates on higher layers of the OSI model. This distinguishing feature helps to eliminate the restriction that IP resources must be managed within a single LAN. For instance, a technique for overcoming the single-LAN restriction is that the IPCD server stores information on private networks, it tags those private networks with the LANs, VLANs, Virtual eXtensible Local Networks (VXLAN) and/or other similar LAN-related attributes used by them. In other words, the IPCD server stores information on which LAN(s) are used by which private network(s). Because overlapping address spaces cannot logically used the same VLANs, all managed network are made unique even if the networks per se use identical address spaces. When VLANs are connected to networks and used in provisioning, it is worth noting that address spaces can overlap but VLANs within a data center cannot.

[0017] According to a preferred but optional feature, if a service provider has multiple data centers, the IPCD server may also tags the network with information relating to the data center in which it is being run, in addition LAN-related attributes. This optional method is used to distinguish between overlapping private networks that are connected to identical LANs, VLANs and/or VXLANs in different data centers sharing the same IPCD service.

[0018] The above-identified elements will be described in more detail in the following. The invention involves, firstly, an IP commissioning/decommissioning server apparatus that is used to manage and track, in real time, distribution and allocation of networks, individual IP addresses within each network, related administrative properties and attributes. Unlike a conventional DHCP server which cannot reclaim an issued IP address before the expiry of the pre-configured fixed lease time, the IP commissioning/decommissioning server apparatus according to the present invention is configured to decommission and reclaim unused IP resources, such as IP addresses, as virtual machines are decommissioned. The decommissioning scheme of the IP commissioning/ decommissioning server differs from known IP reconciliation models in the following respects. While some prior art IP reconciliation models permit reclaiming IP addresses that were allocated but are now unused, the prior art IP reconciliation models only permit reclamation of unused IP addresses only after their IP address lease time has expired. Otherwise there would be a risk of duplicate IP addresses, which will cause trouble.

[0019] One of the features of the present invention relates to situations wherein a virtual machine and/or a workload run in the cloud-based computing environment ceases to exist or is migrated to another cloud environment. This is quite common in clouding environments as virtual machines are commissioned and decommissioned on an on-demand basis. In such cases the client running in the orchestration system will notify the IP commissioning system when a workload is about to be decommissioned, or is migrated to another cloud environment. As a result of the notification, the inventive IP commissioning/ decommissioning server will automatically release the resources, such as the IP address, formerly assigned to the client that no longer exists. No comparable mechanism exists in DHCP, since the DHCP specifications do not define a client able to notify the DHCP server about the fact that it will no longer use the IP address it has obtained and that the IP address can be released. The present invention solves this problem by communicating with the orchestration system when the said system is decommissioning or migrating workloads to another environment. Use of client-server architecture provides the benefit that the inventive IPCD server may obtain real-time information concerning whether or not a given IP resource is actually being used for anything. In the prior art, lack of accurate information on which IP resources are being used and which are merely assigned but not used makes it impossible to reclaim unused IP resources until the lease time has expired.

[0020] The invention relates to orchestrated computing environments that utilize SOA-based architectures, and in which the IP allocation should be carried out as part of the commissioning and/or decommissioning process managed by the orchestration system. Since SOA environments are based on open web-based APIs (e.g. SOAP) requiring systems to communicate with each other on the Application Layer (Layer 7 in the OSI model), a traditional DHCP server utilizes the Data Link Layer (Layer 2 in the OSI model) sent on OSI 2 layer, thereby making the two incompatible.

[0021] The present IP commissioning/decommissioning server further comprises a remote user interface, such as a web-based user interface, which can be used to access the data managed in the IP commissioning/decommissioning server. The remote user interface provides the ability to track allocation and usage levels of networks in real-time (with regard to stateless autoconfiguration). Other desirable features in the invention is the ability of administrators to transparently monitor these levels in real-time; to support multi-tenant environments, including the possibility to offer internal and/or external end-users a restricted access/viewing rights to the networks used by them; and the ability to manage networks (add, remove, split, merge and/or edit contents) and to be able to associate select attributes (e.g. VLAN associated with a given network; owner of the network; etc.) with the managed networks by tagging the select attributes or properties to them . The reason why the last part is important, is that in order to associate e.g. a given VLAN known by an orchestration solution with a network from which an IP is to be retrieved, there needs to be a common denominator, such as the VLAN tag or name, which is known to the orchestration system and the IP commissioning/decommissioning server alike. And to manage all this, a GUI is preferable for many administrative users.

[0022] The present IP commissioning/decommissioning server further comprises an Application Programming Interface (API) that supports Service Oriented Architecture (SOA) plus a first logic capable of dynamically assigning and releasing IP resources, such as IP addresses and other network configurations based on calls received via the API.

[0023] The IP commissioning/decommissioning server further comprises a second logic configured to associate individual networks with given end-users and, in the event that a network has been fully allocated, the IP commissioning/decommissioning server is configured to automatically provision an IP resource, (e.g. an IP address) from an alternate network associated with the end-user and/or from a reserve pool of IP resources.

[0024] Yet further, the IP commissioning/decommissioning server comprises a client-server architecture that involves a client run in the orchestration solution communicating with the IP commissioning/ decommissioning server over the SOA-based API.

[0025] As regards the client, the client comprises a client logic that requests an IP resource, such as an IP address and/or name, from the IP commissioning/decommissioning server when a new virtual machine is commissioned to the cloud-computing environment by the orchestration solution. The client logic triggers the release of the assigned IP resources, such as the IP address and/or the name of a virtual machine that has been decommissioned.

BRIEF DESCRIPTION OF THE DRAWINGS

[0026] In the following the invention will be described in greater detail by means of specific embodiments with reference to the attached drawings, in which

Figure 1 is a block-level diagram of client-server architecture according to an embodiment of the invention;

Figure 2, which consists of partial drawings 2A and 2B, is a flow chart illustrating operation of an IP commissioning/decommissioning server (IPCD server) according to an embodiment of the invention; and

Figure 3 schematically shows an exemplary block diagram for the various server and/or client computers.

DETAILED DESCRIPTION OF SPECIFIC EMBODIMENTS

[0027] Figure 1 is a block-level diagram of a client-server architecture according to an embodiment of the invention. Reference sign IPCDS denotes the IP commissioning/decommissioning server (IPCD server) according to the present embodiment. Reference sign UI denotes a user interface, which preferably is a remote user interface, such as a web-based user interface. The user interface can be used to access the data managed in the IPCD server IPCDS. Reference sign API denotes an Application Programming Interface that supports Service Oriented Architecture (SOA). Reference sign BL denotes a business logic, which is capable of dynamically assigning and releasing IP resources, such as IP addresses, names, and other network configurations based on calls received via the API. The present embodiment of the IPCD server IPCDS further comprises two management engines ME1 and ME2, which respectively correspond to the first logic and second logic, and which are collectively configured to associate individual networks with given end-users and, in the event that a network has been fully allocated, the IP commissioning/decommissioning server is configured to automatically provision an IP resource, (e.g. an IP address) from an alternate network associated with the internal and/or external end-user and/or from a reserve pool of IP resources.

[0028] Reference sign CL denotes a client computer of the inventive client-server architecture. The client computer CL runs in, or in connection with, an orchestration solution OS communicating with the IPCD server IPCDS over the SOA-based Application Programming Interface API. The orchestration solution OS supports a number of server instances SI. In a typical implementation the server instances SI are virtual machines, each of which has a respective workload.

[0029] Figure 2, which consists of partial drawings 2A and 2B, is a flow chart illustrating operation of an IP commissioning/decommissioning server (IPCD server) according to an embodiment of the invention.

[0030] In step 2-2 the orchestration system sends a request to the IPCD server IPCDS, wherein the request identifies the virtual local-area network VLAN in which a host is to be deployed. In step 2-4 the IPCD server IPCDS examines whether the VLAN is tagged to a network managed by the present IPCD system. If not, the flow proceeds to step 2-46 (see drawing 2B), wherein the IPCD server IPCDS returns an error message stating the cause of the error. If in step 2-4 the outcome is positive, the flow proceeds to step 2-6, wherein the IPCD server IPCDS examines whether issuance of an IP address was part of the request. If yes, the flow proceeds to step 2-14, wherein the IPCD server IPCDS examines if a free IP address is available in the first network associated with the VLAN. If yes, the flow proceeds to step 2-20, wherein the IPCD server IPCDS examines if issuance of a name was part of the request. If it was, the flow proceeds to step 2-30, wherein the IPCD server IPCDS reserves a free IP address and generates a unique name attaching it to a default zone configured to the network associated with the VLAN. The IPCD server IPCDS then returns the unique name and IP address to the orchestration system OS and marks the IP address as used in the network from which it was allocated. If the outcome of step 2-20 was negative, the flow proceeds to step 2-32, wherein the IPCD server IPCDS reserves a free IP address, returns it to the orchestration system OS, and marks the IP address as used in the network from which it was allocated.

[0031] If the outcome of step 2-6 was negative, the flow proceeds to step 2-8, wherein the IPCD server IPCDS examines whether issuance of a name was part of the request. If yes, the flow proceeds to step 2-38, wherein the IPCD server IPCDS returns an automatically generated name attaching it to a default zone configured to the network associated with the VLAN, and returns the unique name to the orchestration system OS. If the outcome of step 2-8 was negative, the flow proceeds to step 2-10, wherein the IPCD server IPCDS examines whether release of an IP address was part of the request. If yes, the flow proceeds to step 2-18, wherein the IPCD server IPCDS examines whether release of a used name was part of the request. If yes, the flow proceeds to step 2-40, wherein the IPCD server IPCDS releases the used name and IP address from the network that matched the associated VLAN, and returns a confirmation to the orchestration system OS. If the outcome of step 2-18 was negative, the flow proceeds to step 2-42, wherein the IPCD server IPCDS releases the used IP from the network matched with the associated VLAN and returns a confirmation to the orchestration system OS.

[0032] In the present example, if the outcome of step 2-10 was negative, the flow proceeds to step 2-12, wherein the IPCD server IPCDS notices that the request did not relate to any of the functions of the IPCD server, and returns an error message. In more ambitious implementations with more functions, the branch from step 2-12 to the right may proceed to further tests and functions.

[0033] If the outcome of step 2-14 was negative, the flow proceeds to step 2-16, wherein the IPCD server IPCDS examines if there are other networks tagged with the same VLAN. If there are, the flow proceeds to step 2-22, wherein the IPCD server IPCDS examines if there are free IP addresses in the other networks. If there are, the IPCD server IPCDS examines in step 2-24 if issuance of a name was part of the request. If yes, the flow proceeds to the above-described step 2-30. If not, the flow proceeds to step 2-34, wherein the IPCD server IPCDS returns the free IP address to the orchestration system OS and marks it as used in the network from which the IP was allocated.

[0034] If the outcome of step 2-16 or 2-22 is negative, the flow proceeds to step 2-36, wherein the IPCD server IPCDS returns an automatically generated name attaching it to a default zone configured to the network associated with the VLAN, and returns the unique name to the orchestration system OS.

[0035] Figure 3 schematically shows an exemplary block diagram for the various information processing elements. The data processing architecture shown in Figure 3, generally denoted by reference numeral 3-100, can be used to implement the servers and clients of the present invention. A configuration appropriate for a server is shown; the configuration of a client computer may be simpler. The two major functional blocks of the are a processing system 3-100 and a storage system 3-190. The processing system 3-100 comprises one or more central processing units CP1 ... CPn, generally denoted by reference numeral 3-110. The processing units may be native or virtual processing units. Embodiments comprising multiple processing units 3-110 are preferably provided with a load balancing unit 3-115 that balances processing load among the multiple processing units 3-110. The multiple processing units 3-110 may be implemented as separate processor components or as physical processor cores or virtual processors within a single component case. The processing system 3-100 further comprises a network interface 3-120 for communicating with various data networks, which are generally denoted by reference sign DN. The data networks DN may include local-area networks, such as an Ethernet network, and/or wide-area networks, such as the internet. Reference numeral 3-125 denotes a mobile network interface, through which the processing system 3-100 may communicate with various access networks AN, which in turn serve the mobile terminals used by end users or clients. A configuration supporting multiple different networks enables the processing system 3-100 to support multiple types of clients, such as land-based terminals 3-200 and mobile terminals 3-210.

[0036] The processing system 3-100 of the present embodiment may also comprise a local user interface 3-140. Depending on implementation, the user interface 3-140 may comprise local input-output circuitry for a local user interface, such as a keyboard, mouse and display (not shown). Alternatively or additionally, management of the processing system 3-100 may be implemented remotely, by utilizing the network interface 3-120 and any internet-enabled terminal that provides a user interface. The nature of the user interface depends on which kind of computer is used to implement the processing system 3-100. If the processing system 3-100 is a dedicated computer, it may not need a local user interface, and the processing system 3-100 may be managed remotely, such as from a web browser over the internet, for example. Such remote management may be accomplished via the same network interface 3-120 that the computer utilizes for traffic between itself and the client terminals.

[0037] The processing system 3-100 also comprises memory 3-150 for storing program instructions, operating parameters and variables. Reference numeral 3-160 denotes a program suite for the processing system 3-100.

[0038] The processing system 3-100 also comprises circuitry for various clocks, interrupts and the like, and these are generally depicted by reference numeral 3-130. The processing system 3-100 further comprises a storage interface 3-145 to the storage system 3-190. When the processing system 3-100 is switched off, the storage system 3-190 may store the software that implements the processing functions, and on power-up, the software is read into semiconductor memory 3-150. The storage system 3-190 also retains operating and variables over power-off periods. In large-volume implementations, that is, implementations wherein a single processing system 3-100 serves a large number of clients via respective mobile terminals MT, the storage system 3-190 may be used to store the dynamic dialog matrices associated with the clients and mobile terminals MT. The various elements 3-110 through 3-150 intercommunicate via a bus 3-105, which carries address signals, data signals and control signals, as is well known to those skilled in the art.

[0039] It will be apparent to a person skilled in the art that, as the technology advances, the inventive concept can be implemented in various ways. The invention and its embodiments are not limited to the examples described above but may vary within the scope of the claims.

CLAIMS

1. A server computer (IPCDS, 3-100) for commissioning/decommissioning Internet Protocol ["IP"] resources to a plurality of server instances provisioned using an orchestration solution, wherein the server computer and each of several clients constitute a client-server architecture, wherein the server computer comprises:

- a processing system (3-110) comprising at least one processing unit (CP1 ... CPn;
- memory (MEM, 3-150) for storing applications and data;

wherein the memory comprises program code instructions (3-160) for instructing the processing system (3-110) to implement the following features:

- a user interface (UI) for remote management of the server computer, wherein the user interface is configured to provide access to data (DB) managed by the server computer;
- a web-based application programming interface (API) capable of supporting service oriented architecture ["SOA"];
- a first logic configured to dynamically assign and release IP resources to each of a number of the clients (CL) based on one or more calls received from a respective client via the application programming interface;
- wherein the first logic comprises program code instructions for creating at least one unique IP resource for each of two or more server instances (SI), wherein the two or more server instances are nodes of networks with overlapping address spaces;
- wherein the at least one unique IP resource is based on a combination of a name of the private network of the respective one of the two or more server instances, and an IP address within that private network.

2. The server computer according to claim 1, wherein the at least one unique IP resource is based on a combination of a name of the private network of the respective one of the two or more computers, and an IP address within that private network.

3. The server computer according to claim 1 or 2, further comprising a second logic configured to associate individual networks with given internal and/or external end-users wherein, in case one or more resources of a network have been fully allocated, the second logic is configured to automatically provision an IP re-

source from an alternate network, which is associated with the internal and/or external end-user, and/or from a reserve pool of IP resources.

4. The server computer according to claim 1, 2 or 3, wherein the processing system comprises multiple processing units and a load-balancing unit for distributing processing load among the multiple processing units.

5. A client computer for requesting commissioning/decommissioning Internet Protocol ["IP"] resources, the client computer comprising:

- at least one processing unit;
- memory for storing applications and data;
- wherein the memory comprises program code instructions for instructing the at least one processing unit to implement a client logic configured to perform the following tasks:
- requesting an IP resource from an IP commissioning/decommissioning server when a new virtual machine is commissioned to the cloud computing environment by the orchestration solution ; and
- triggering release of IP resources assigned to a virtual machine in response to the virtual machine having been decommissioned.

6. A method for operating a server computer for commissioning/decommissioning Internet Protocol ["IP"] resources to a plurality of clients provisioned using an orchestration solution, wherein the server computer and each of the clients constitute a client-server architecture;
the method comprising:

- providing the server computer with an application programming interface ["API"] capable of supporting service oriented architecture ["SOA"];
- configuring a remote interface to provide access to data managed by the server computer;
- remotely managing the server computer via the remote user interface;
- dynamically assigning and releasing IP resources to each of several clients based on one or more calls received from a respective client via the application programming interface;
- creating at least one unique IP resource for each of two or more server instances, wherein the two or more server instances are nodes of networks with overlapping address spaces;

- wherein the at least one unique IP resource is based on a combination of a name of the private network of the respective one of the two or more server instances, and an IP address within that private network.

7. A method for operating a client computer for requesting commissioning/decommissioning Internet Protocol ["IP"] resources, the method comprising:

- requesting an IP resource from an IP commissioning/decommissioning server when a new virtual machine is commissioned to the cloud computing environment by the orchestration solution ; and
- triggering release of IP resources assigned to a virtual machine in response to the virtual machine having been decommissioned.

8. A tangible program carrier comprising program code instructions, wherein the program code instructions, when executed on a server computer as defined in claim 1, cause the server computer to carry out the method of claim 6.

9. A tangible program carrier comprising program code instructions, wherein the program code instructions, when executed on a client computer as defined in claim 5, cause the client computer to carry out the method of claim 7.

Fig. 1

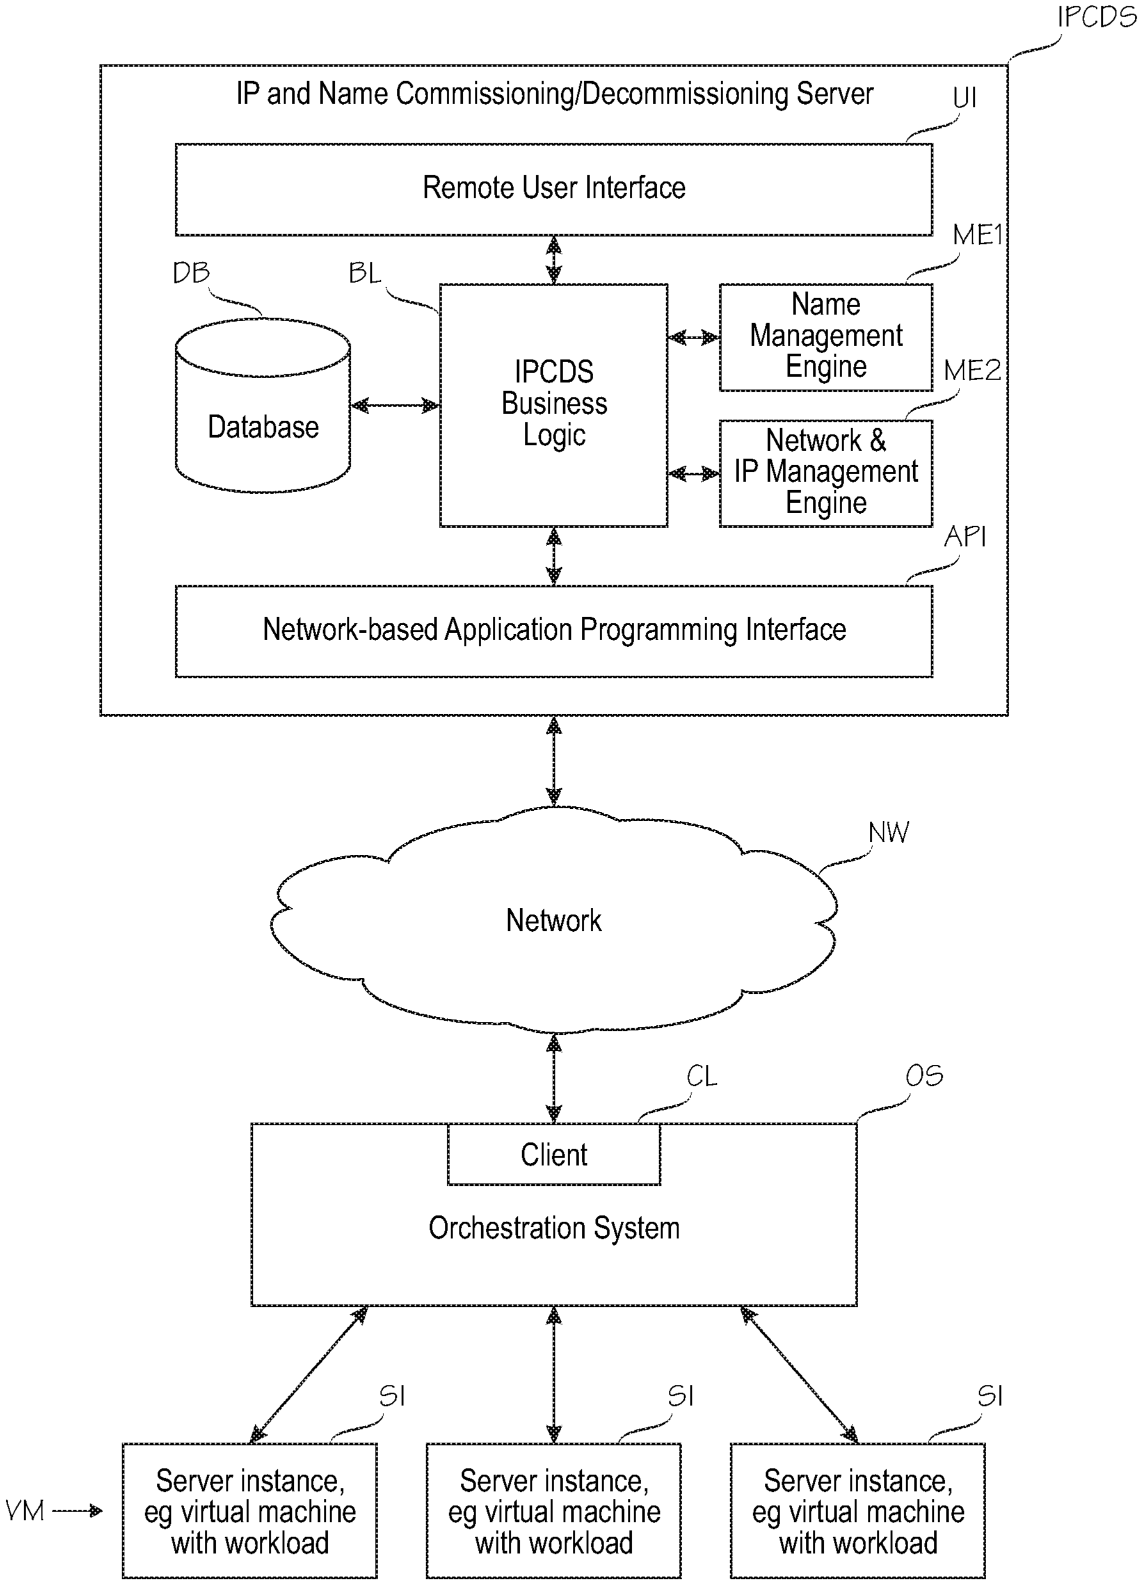
IPCDS
IP and Name Commissioning/Decommissioning Server
UI
Remote User Interface
ME1
DB
BL
Name Management Engine
IPCDS Business Logic
Database
ME2
Network & IP Management Engine
API
Network-based Application Programming Interface
NW
Network
CL
OS
Client
Orchestration System
SI
SI
SI
VM
Server instance, eg virtual machine with workload
Server instance, eg virtual machine with workload
Server instance, eg virtual machine with workload

Fig. 2A

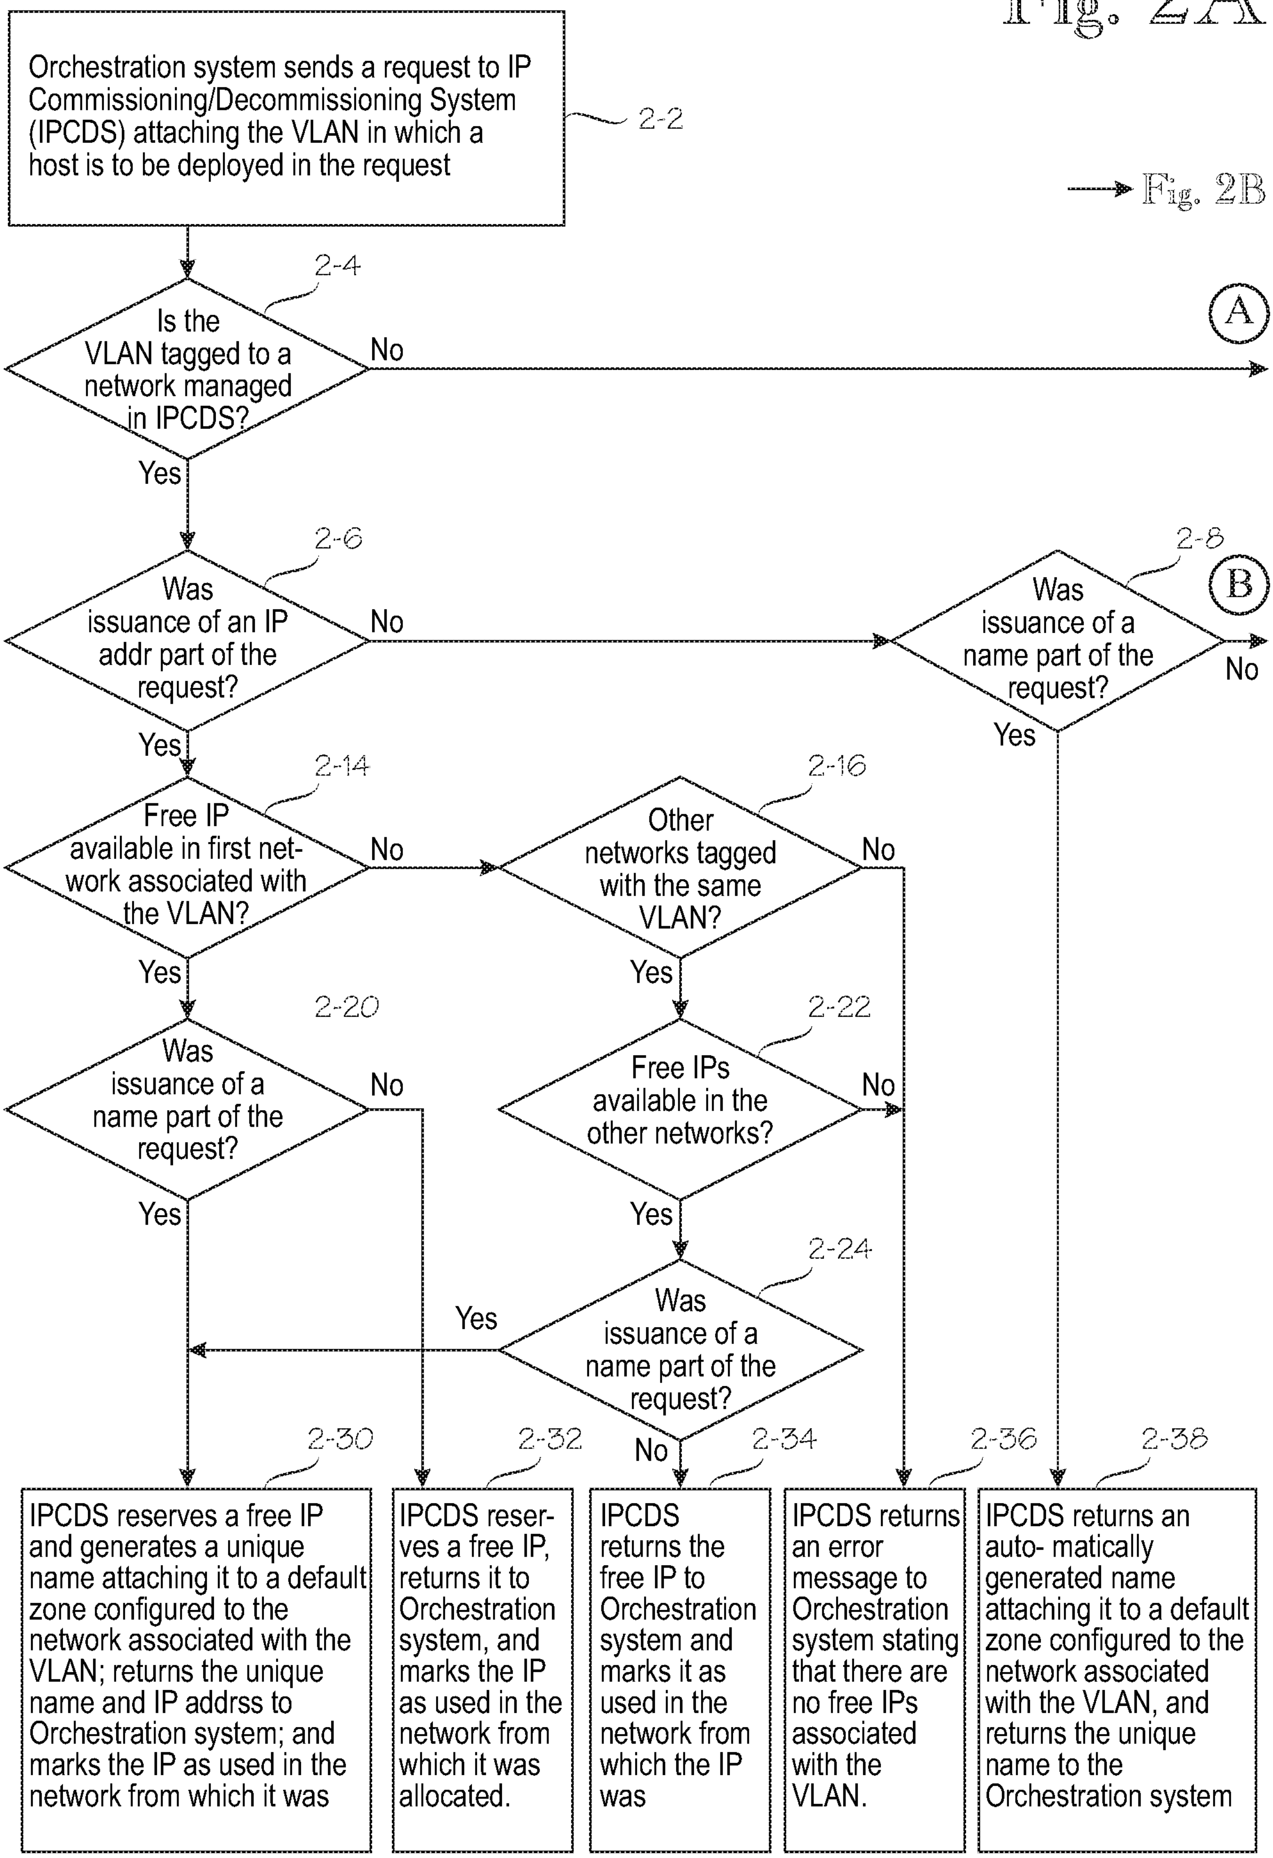

Orchestration system sends a request to IP Commissioning/Decommissioning System (IPCDS) attaching the VLAN in which a host is to be deployed in the request
2-2
→ Fig. 2B
2-4
Is the VLAN tagged to a network managed in IPCDS?
No
A
Yes
2-6
Was issuance of an IP addr part of the request?
No
2-8
Was issuance of a name part of the request?
B
No
Yes
Yes
2-14
Free IP available in first net-work associated with the VLAN?
No
2-16
Other networks tagged with the same VLAN?
No
Yes
Yes
2-20
Was issuance of a name part of the request?
No
2-22
Free IPs available in the other networks?
No
Yes
Yes
2-24
Was issuance of a name part of the request?
Yes
No
2-30
IPCDS reserves a free IP and generates a unique name attaching it to a default zone configured to the network associated with the VLAN; returns the unique name and IP addrss to Orchestration system; and marks the IP as used in the network from which it was
2-32
IPCDS reser-ves a free IP, returns it to Orchestration system, and marks the IP as used in the network from which it was allocated.
2-34
IPCDS returns the free IP to Orchestration system and marks it as used in the network from which the IP was
2-36
IPCDS returns an error message to Orchestration system stating that there are no free IPs associated with the VLAN.
2-38
IPCDS returns an auto- matically generated name attaching it to a default zone configured to the network associated with the VLAN, and returns the unique name to the Orchestration system

Fig. 2B

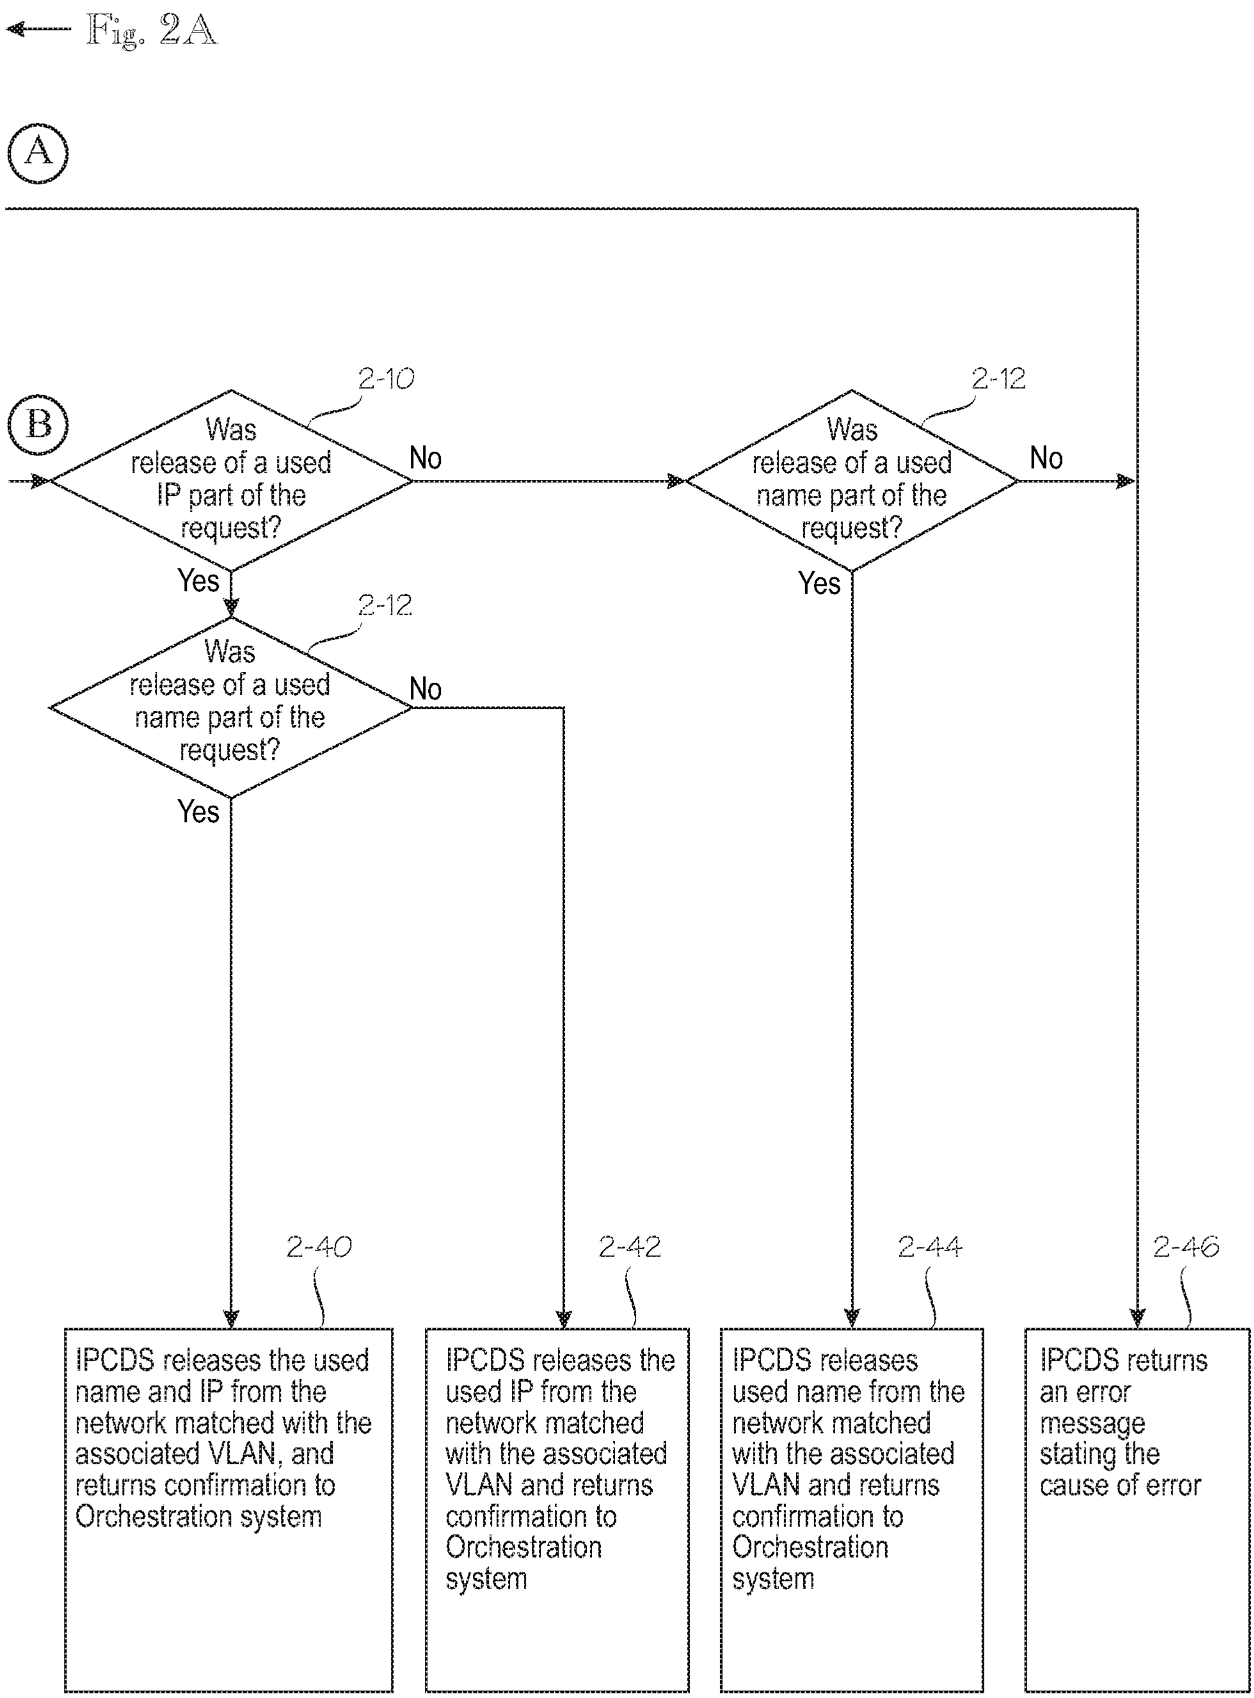
Fig. 2A
A
B
2-10
Was release of a used IP part of the request?
No
Yes
2-12
Was release of a used name part of the request?
No
Yes
2-12
Was release of a used name part of the request?
No
Yes
2-40
IPCDS releases the used name and IP from the network matched with the associated VLAN, and returns confirmation to Orchestration system
2-42
IPCDS releases the used IP from the network matched with the associated VLAN and returns confirmation to Orchestration system
2-44
IPCDS releases used name from the network matched with the associated VLAN and returns confirmation to Orchestration system
2-46
IPCDS returns an error message stating the cause of error

Fig. 3

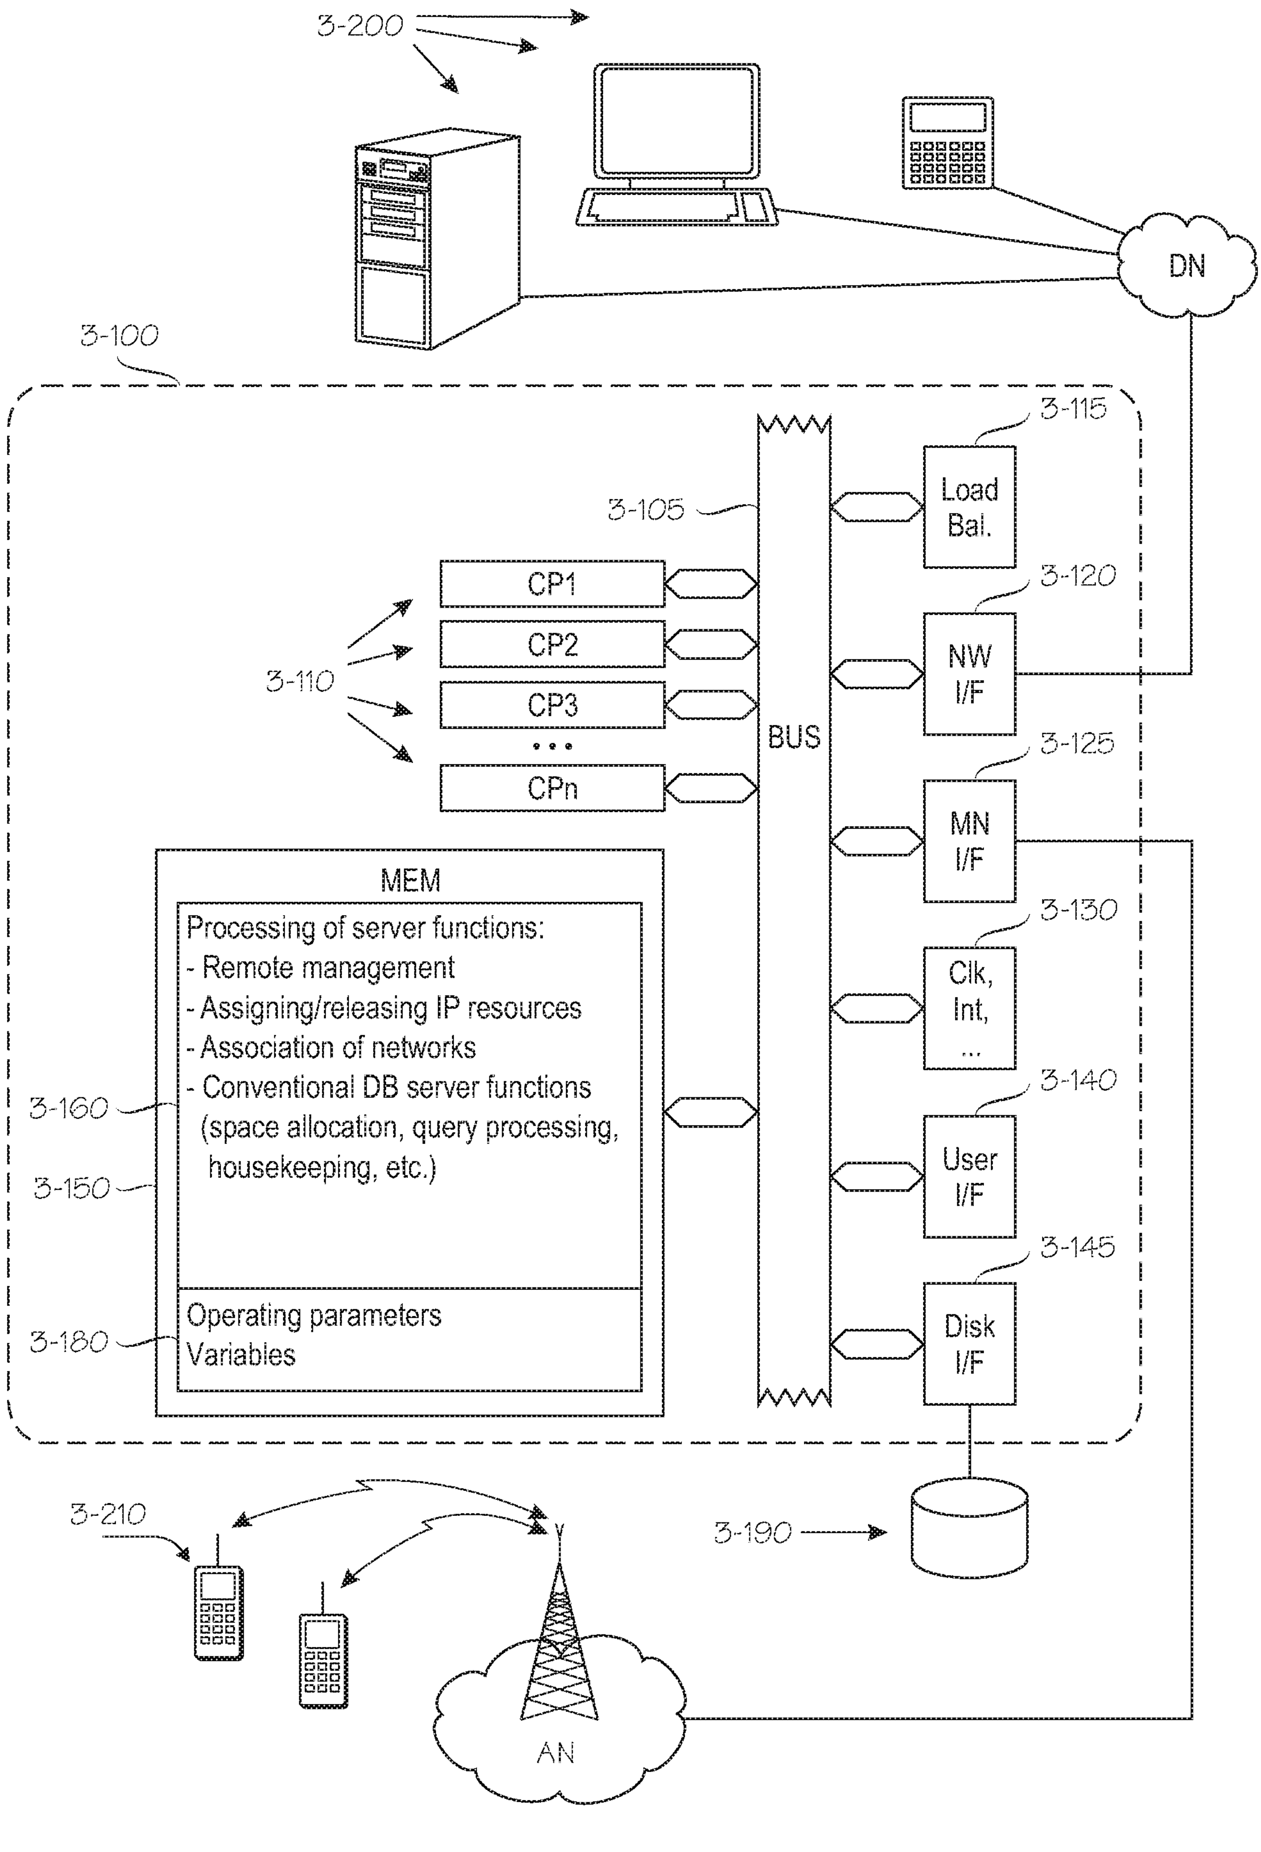
3-200
DN
3-100
3-115
Load Bal.
3-105
CP1
CP2
CP3
CPn
3-110
3-120
NW I/F
BUS
3-125
MN I/F
MEM
Processing of server functions:
- Remote management
- Assigning/releasing IP resources
- Association of networks
- Conventional DB server functions (space allocation, query processing, housekeeping, etc.)
3-130
Clk, Int, ...
3-160
3-140
User I/F
3-150
3-145
Disk I/F
Operating parameters
Variables
3-180
3-210
3-190
AN

INTERNATIONAL SEARCH REPORT

International application No.
PCT/FI2013/050664

A. CLASSIFICATION OF SUBJECT MATTER

See extra sheet

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)

IPC: H04L, H04W, G06F

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

FI, SE, NO, DK

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EPO-Internal, WPI, XPESP, XPESP2, XPRD, XPIOP, XPOAC

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2012131156 A1 (BRANDT MARK S [US] et al.) 24 May 2012 (24.05.2012) Abstract, Figures 5 and 9 and paragraphs [0031]-[0033], [0038], [0069], [0076], [0081], [0082], [0092] and [0097].	1-9
A	EP 2216718 A1 (NOVELL INC [US]) 11 August 2010 (11.08.2010) The whole document, especially abstract and Figure 1.	1-9
A	US 2002124066 A1 (CHANG CHING-JYE [US] et al.) 05 September 2002 (05.09.2002) The whole document.	1-9
A	US 2009259740 A1 (HEPBURN ANDREW NEIL CAMERON [GB]) 15 October 2009 (15.10.2009) The whole document.	1-9

☒ 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	Date of mailing of the international search report
23 September 2013 (23.09.2013)	30 September 2013 (30.09.2013)

Name and mailing address of the ISA/FI
National Board of Patents and Registration of Finland
P.O. Box 1160, FI-00101 HELSINKI, Finland
Facsimile No. +358 9 6939 5328

Authorized officer
Mikko Pitkänen
Telephone No. +358 9 6939 500

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	EP 1251657 A2 (SYSTEM MAN ARTS INC A DELAWARE [US]) 23 October 2002 (23.10.2002) The whole document.	1-9
A	US 6948003 B1 (NEWMAN PETER [US] et al.) 20 September 2005 (20.09.2005) The whole document.	1-9

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.
PCT/FI2013/050664

Patent document cited in search report	Publication date	Patent family members(s)	Publication date
US 2012131156 A1	24/05/2012	AU 2011332008 A1	20/06/2013
		CA 2818838 A1	31/05/2012
		WO 2012071382 A1	31/05/2012
EP 2216718 A1	11/08/2010	US 2010205304 A1	12/08/2010
US 2002124066 A1	05/09/2002	None	
US 2009259740 A1	15/10/2009	US 7908353 B2	15/03/2011
EP 1251657 A2	23/10/2002	US 2002165982 A1	07/11/2002
		US 7383357 B2	03/06/2008
US 6948003 B1	20/09/2005	US RE43051 E1	27/12/2011

CLASSIFICATION OF SUBJECT MATTER

Int.Cl.
G06F 9/455 (2006.01)
H04L 29/12 (2006.01)
G06F 15/16 (2006.01)