

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
31 March 2005 (31.03.2005)

PCT

(10) International Publication Number
WO 2005/029326 A1

(51) International Patent Classification⁷: **G06F 9/455**

(21) International Application Number:
PCT/US2004/029611

(22) International Filing Date:
10 September 2004 (10.09.2004)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
10/663,206 15 September 2003 (15.09.2003) US

Hillsboro, OR 97123 (US). **KAGI, Alain** [—/US]; 3624 SE Crystal Springs Boulevard, Portland, OR 97202 (US). **JEYASINGH, Stalinselvaraj** [—/US]; 16893 NW Tucson Street, Beaverton, OR 97006 (US). **ANDERSON, Andrew** [—/US]; 677 SE 68th Avenue, Hillsboro, OR 97123 (US). **UHLIG, Richard** [—/US]; 1564 NE Orenco Station Parkway West, Hillsboro, OR 97124 (US). **COTA-ROBLES, Erik** [—/US]; 11790 SW Lanewood, Portland, OR 97225 (US). **RODGERS, Scott** [—/US]; 452 SW Brookwood Avenue, Hillsboro, OR 97123 (US). **SMITH, Lawrence, III** [—/US]; 14995 NW Northumbria, Beaverton, OR 97006 (US).

(74) Agents: **VINCENT, Lester, J.** et al.; Blakely, Sokoloff, Taylor & Zafman, 12400 Wilshire Boulevard, 7th Floor, Los Angeles, CA 90025 (US).

(71) Applicant (for all designated States except US): **INTEL CORPORATION** [US/US]; 2200 Mission College Boulevard, Santa Clara, CA 95052 (US).

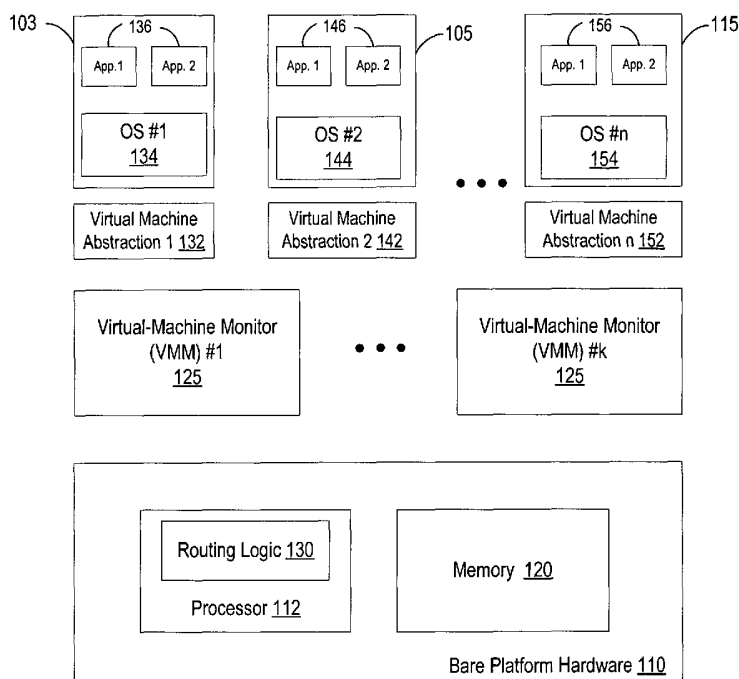
(72) Inventors; and

(75) Inventors/Applicants (for US only): **NEIGER, Gilbert** [—/US]; 2424 NE 11th Avenue, Portland, OR 97212 (US). **BENNETT, Steven** [—/US]; 6469 SE Sigrid Street,

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SY, TJ, TM,

[Continued on next page]

(54) Title: USE OF MULTIPLE VIRTUAL MACHINE MONITORS TO HANDLE PRIVILEGED EVENTS



(57) Abstract: In one embodiment, a method for handling privileged events in a virtual machine environment includes detecting an occurrence of a privileged event, determining which one of multiple virtual machine monitors (VMMs) is to handle the privileged event, and transitioning control to this VMM.

WO 2005/029326 A1



TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— with international search report

— before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

USE OF MULTIPLE VIRTUAL MACHINE MONITORS TO HANDLE PRIVILEGED EVENTS

Field of the Invention

5 [0001] The present invention relates generally to virtual machines, and more specifically to handling privileged events using multiple virtual machine monitors.

Background of the Invention

 [0002] A conventional virtual-machine monitor (VMM) typically runs
10 on a computer and presents to other software the abstraction of one or more virtual machines. Each virtual machine may function as a self-contained platform, running its own "guest operating system" (i.e., an operating system (OS) hosted by the VMM) and other software, collectively referred to as guest software. The guest software expects to operate as if it were running on a
15 dedicated computer rather than a virtual machine. That is, the guest software expects to control various events and have access to hardware resources. The hardware resources may include processor-resident resources (e.g., control registers), resources that reside in memory (e.g., descriptor tables) and resources that reside on the underlying hardware platform (e.g., input-output devices). The
20 events may include internal interrupts, external interrupts, exceptions, platform events (e.g., initialization (INIT) or system management interrupts (SMIs)), execution of certain instructions, etc.

 [0003] In a virtual-machine environment, the VMM should be able to have ultimate control over these events and hardware resources to provide proper

operation of guest software running on the virtual machines and for protection from and between guest software running on the virtual machines. To achieve this, the VMM typically receives control when guest software accesses a protected resource or when other events (such as interrupts or exceptions) occur. For
5 example, when an operation in a virtual machine supported by the VMM causes a system device to generate an interrupt, the currently running virtual machine is interrupted and control of the processor is passed to the VMM. The VMM then receives the interrupt, and handles the interrupt itself or delivers the interrupt to the appropriate virtual machine.

10 Brief Description of the Drawings

[0004] The present invention is illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings and in which like reference numerals refer to similar elements and in which:

[0005] **Figure 1** illustrates one embodiment of a virtual-machine
15 environment, in which the present invention may operate;

[0006] **Figures 2** illustrates a configuration of multiple VMMs in a virtual-machine environment;

[0007] **Figure 3** is a flow diagram of one embodiment of a process for handling privileged events in a virtual-machine environment having multiple
20 VMMs;

[0008] **Figures 4, 6 and 7** illustrate exemplary embodiments of processes of identifying a VMM for handling a privileged event; and

[0009] Figure 5 illustrates an exemplary usage of process 400 in a virtual machine environment having two VMMs.

Description of Embodiments

[0010] A method and apparatus for handling privileged events using multiple virtual machine monitors are described. In the following description, for purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be apparent, however, to one skilled in the art that the present invention can be practiced without these specific details.

[0011] Some portions of the detailed descriptions that follow are presented in terms of algorithms and symbolic representations of operations on data bits within a computer system's registers or memory. These algorithmic descriptions and representations are the means used by those skilled in the data processing arts to most effectively convey the substance of their work to others skilled in the art. An algorithm is here, and generally, conceived to be a self-consistent sequence of operations leading to a desired result. The operations are those requiring physical manipulations of physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared, and otherwise manipulated. It has proven convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, or the like.

[0012] It should be borne in mind, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities. Unless specifically stated otherwise as apparent from the following discussions, it is appreciated that throughout the present invention, discussions utilizing terms such as "processing" or "computing" or "calculating" or "determining" or the like, may refer to the action and processes of a computer system, or similar electronic computing device, that manipulates and transforms data represented as physical (electronic) quantities within the computer system's registers and memories into other data similarly represented as physical quantities within the computer-system memories or registers or other such information storage, transmission or display devices.

[0013] In the following detailed description of the embodiments, reference is made to the accompanying drawings that show, by way of illustration, specific embodiments in which the invention may be practiced. In the drawings, like numerals describe substantially similar components throughout the several views. These embodiments are described in sufficient detail to enable those skilled in the art to practice the invention. Other embodiments may be utilized and structural, logical, and electrical changes may be made without departing from the scope of the present invention. Moreover, it is to be understood that the various embodiments of the invention, although different, are not necessarily mutually exclusive. For example, a particular feature, structure, or characteristic described in one embodiment may be included within other embodiments. The following detailed description is, therefore, not to be taken in

a limiting sense, and the scope of the present invention is defined only by the appended claims, along with the full scope of equivalents to which such claims are entitled.

[0014] Although the below examples may describe embodiments of the present invention in the context of execution units and logic circuits, other
5 embodiments of the present invention can be accomplished by way of software. For example, in some embodiments, the present invention may be provided as a computer program product or software which may include a machine or computer-readable medium having stored thereon instructions which may be
10 used to program a computer (or other electronic devices) to perform a process according to the present invention. In other embodiments, steps of the present invention might be performed by specific hardware components that contain hardwired logic for performing the steps, or by any combination of programmed computer components and custom hardware components.

15 [0015] Thus, a machine-readable medium may include any mechanism for storing or transmitting information in a form readable by a machine (e.g., a computer), but is not limited to, floppy diskettes, optical disks, Compact Disc, Read-Only Memory (CD-ROMs), and magneto-optical disks, Read-Only Memory (ROMs), Random Access Memory (RAM), Erasable Programmable Read-Only
20 Memory (EPROM), Electrically Erasable Programmable Read-Only Memory (EEPROM), magnetic or optical cards, flash memory, a transmission over the Internet, electrical, optical, acoustical or other forms of propagated signals (e.g., carrier waves, infrared signals, digital signals, etc.) or the like.

[0016] Further, a design may go through various stages, from creation to simulation to fabrication. Data representing a design may represent the design in a number of manners. First, as is useful in simulations, the hardware may be represented using a hardware description language or another functional description language. Additionally, a circuit level model with logic and/or transistor gates may be produced at some stages of the design process. Furthermore, most designs, at some stage, reach a level of data representing the physical placement of various devices in the hardware model. In the case where conventional semiconductor fabrication techniques are used, data representing a hardware model may be the data specifying the presence or absence of various features on different mask layers for masks used to produce the integrated circuit. In any representation of the design, the data may be stored in any form of a machine-readable medium. An optical or electrical wave modulated or otherwise generated to transmit such information, a memory, or a magnetic or optical storage such as a disc may be the machine readable medium. Any of these mediums may "carry" or "indicate" the design or software information. When an electrical carrier wave indicating or carrying the code or design is transmitted, to the extent that copying, buffering, or re-transmission of the electrical signal is performed, a new copy is made. Thus, a communication provider or a network provider may make copies of an article (a carrier wave) embodying techniques of the present invention.

[0017] Figure 1 illustrates one embodiment of a virtual-machine environment 100, in which the present invention may operate. In this

embodiment, bare platform hardware 110 comprises a computing platform, which may be capable, for example, of executing a standard operating system (OS) or virtual-machine monitors (VMMs), such as VMMs 125.

[0018] The platform hardware 110 can be of a personal computer (PC),
5 mainframe, handheld device, portable computer, set-top box, or any other computing system. The platform hardware 110 includes at least one processor 112, memory 120 and possibly other platform hardware (e.g. input-output devices), not shown.

[0019] Processor 112 can be any type of processor capable of executing
10 software, such as a microprocessor, digital signal processor, microcontroller, or the like. The processor 112 may include microcode, programmable logic or hardcoded logic for performing the execution of method embodiments of the present invention.

[0020] Memory 120 can be a hard disk, a floppy disk, random access
15 memory (RAM), read only memory (ROM), flash memory, any combination of the above devices, or any other type of machine medium readable by processor 112. Memory 120 may store instructions or data for performing the execution of method embodiments of the present invention.

[0021] Each VMM 125, though typically implemented in software, may
20 emulate and export a bare machine interface to higher level software. Such higher level software may comprise a standard or real-time OS, may be a highly stripped down operating environment with limited operating system functionality, may not include traditional OS facilities, etc. The VMMs 125 may be implemented, for

example, in hardware, software, firmware, or by a combination of various techniques.

[0022] When running, each VMM 125 presents to “guest” software (i.e., software other than that of the VMMs 125) the abstraction of one or more virtual machines (VMs). The VMMs 125 may provide the same or different abstractions to the various guests. The guest software running on each VM may include a guest OS (e.g., a guest OS 134, 144 or 154) and various guest software applications (e.g., applications 136, 146 and 156). Collectively, guest OS and software applications are referred to herein as guest software 103, 105 and 115.

10 [0023] Guest software 103, 105 and 115 expects to access physical resources (e.g., processor registers, memory and I/O devices) within the VMs 132, 142 and 152 on which the guest software is running. The VMMs 125 facilitate access to resources desired by guest software while retaining ultimate control over resources within the platform hardware 110. In addition, the guest software 15 103, 105 and 115 expect to handle various events such as exceptions, interrupts and platform events (e.g., initialization (INIT) and system management interrupts (SMIs)). Some of these events are “privileged” because they cannot be handled by the guest software to ensure proper operation of VMs 132, 142 and 152 and guest software 103, 105 and 115 and for protection from and between guest software 20 103, 105 and 115. The privileged events are handled by the VMMs 125.

[0024] In particular, a specific VMM is identified to handle each privileged event. In one embodiment, a specific VMM is identified using routing logic 130. In some embodiments, the routing logic 130 is implemented as

microcode, programmable logic or hardcoded logic. In other embodiments, the routing logic 130 is implemented as software residing in memory 120. In yet other embodiment, the routing logic 130 is implemented as a combination of hardware and software.

5 **[0025]** Once a specific VMM is identified, it will facilitate functionality desired by guest software while retaining ultimate control over this privileged event. The act of facilitating the functionality for the guest software may include a wide variety of activities on the part of the VMMs 125. These activities of the VMMs 125 should not limit the scope of the present invention.

10 **[0026]** The system 100 may include two or more VMMs 125 executing on the platform hardware 110. In one embodiment, the VMMs 125 run in parallel, and each VMM 125 can receive control from a VM. **Figure 2** illustrates one embodiment of such configuration of the VMMs 125 within the system 100.

[0027] Referring to **Figure 2**, an exemplary parallel configuration of
15 VMMs 210 through 214 is illustrated. According to this configuration, control transfers from a VM to a specific VMM when a privileged event occurs during the operation of a VM. The specific VMM is identified based on certain criteria as will be discussed in more detail below. Once the specific VMM is identified, the information pertaining to the virtualization event is delivered to this VMM, and
20 control is transitioned to it. For example, privileged event 1 occurring during the operation of VM 202 may cause transition from VM 202 to VMM 210, privileged event 2 occurring during the operation of VM 202 may cause transition from VM 202 to VMM 212, privileged event 3 occurring during the operation of VM 206

may cause transition from VM 206 to VMM 212, and privileged event 4 occurring during the operation of VM 206 may cause transition from VM 206 to VMM 214. A corresponding VMM then handles the virtualization event and may transfer control back to a VM from which the control was received. In one embodiment, the transfer of control from a VMM to a VM is achieved by executing a special instruction. The transfer of control from a VMM to a VM is referred to herein as a VM entry and the transfer of control from a VM to a VMM is referred to herein as a VM exit.

[0028] In one embodiment, when a VM exit occurs, control is passed to a VMM at a specific entry point (e.g., an instruction pointer value) delineated in a virtual machine control structure (VMCS) that resides in memory and is maintained by the processor. In another embodiment, control is passed to a VMM after vectoring through a redirection structure (e.g., the interrupt-descriptor table in the processor instruction set architecture (ISA) of the Intel® Pentium® 4 (referred to herein as the IA-32 ISA)). Alternatively, any other mechanism known in the art can be used to transfer control from a VM to a VMM.

[0029] A privileged event may occur during the operation of a VMM. Examples of privileged events that may occur during the operation of a VMM may include, for example, system management interrupts (SMIs), INIT, non-maskable interrupts (NMIs), hardware interrupts, interprocessor interrupts (IPIs), etc. In one embodiment, if a privileged event occurs during the operation of a VMM, routing logic is employed to identify a VMM designated to handle this privileged event. If the designated VMM is not the VMM that is currently

operating, the information pertaining to the privileged event is delivered to the designated VMM, and control is transitioned to it. For example, privileged event 5 occurring during the operation of VMM 210 may cause transition from VMM 210 to VMM 212. VMM 212 then handles the privileged event and may transfer control back to VMM 210. Hence, in interactions between VMM 210 and VMM 212, VMM 210 plays a role of a VM. Accordingly, in one embodiment, a VM exit and VM entry mechanism similar to the one described above with respect to the transfer of control between a VM and a VMM is used to transfer control between the VMMs.

10 **[0030]** **Figure 3** is a flow diagram of one embodiment of a process 300 for handling privileged events in a virtual-machine environment having multiple VMMs. The process may be performed by processing logic that may comprise hardware (e.g., circuitry, dedicated logic, programmable logic, microcode, etc.), software (such as run on a general purpose computer system or a dedicated 15 machine), or a combination of both.

[0031] Referring to **Figure 3**, process 300 begins with processing logic detecting the occurrence of a privileged event (processing block 302). A privileged event is an event that is not to be handled by the currently running software. Such privileged events may include exceptions, interrupts, platform 20 events, execution of a "privileged" instruction (an instruction whose execution causes a VM exit), etc. A privileged event may occur during the operation of guest software or during the operation of a VMM.

[0032] At processing block 304, processing logic determines which one of the multiple VMMs is to handle the privileged event. The determination may be based on various factors such as characteristics of the privileged event, current values of designated fields modifiable by the VMMs, operational parameters of the VMMs, operational parameters of guest software, etc. In one example, the determination is based on the type of the privileged event as will be discussed in greater detail below in conjunction with **Figure 4**. In another example, the determination is based on current values of control fields configured by one of the multiple VMMs as will be discussed in more detail below in conjunction with **Figure 6**. In yet another example, the determination is based on the evaluation of load and/or usage characteristics of the VMMs, as will be discussed in greater detail below in conjunction with **Figure 7**. In still another example, the determination is based on the combination of the above factors.

[0033] Once the processing logic determines which VMM is to handle the privileged event, it delivers information pertaining to the privileged event to the designated VMM and transitions control to this VMM (processing block 306). The VMM may then handle the privileged event itself or route it to guest software for handling.

[0034] **Figures 4, 6 and 7** illustrate exemplary embodiments of processes of identifying a VMM for handling a privileged event. The processes may be performed by processing logic that may comprise hardware (e.g., circuitry, dedicated logic, programmable logic, microcode, etc.), software (such as run on a

general purpose computer system or a dedicated machine), or a combination of both.

[0035] Referring to **Figure 4**, process 400 begins with processing logic detecting the occurrence of a privileged event during the operation of guest software or a VMM (processing block 402). At processing block 404, processing logic identifies the type of the privileged event. Next, processing logic determines which VMM is designated to handle privileged events of this type (processing block 406). In one embodiment, each event type is statically mapped to a particular VMM (e.g., using hardcoded logic). In another embodiment, the determination is dynamic, as described below with regard to **Figure 6**. In yet another embodiment, a combination of statically determined VMMs and dynamically determined VMMs is used. That is, some event types may be statically mapped to corresponding VMMs while other event types may require some additional processor operation for the determination.

[0036] If the currently-operating entity is not the designated VMM (decision box 408), processing logic transitions control to the designated VMM (processing block 410). In one embodiment, the transition to the designated VMM may be accomplished by generating a VM exit. Alternatively, any other mechanism known in the art may be used to transition control to the designated VMM.

[0037] If the currently-operating entity is the designated VMM, the event is delivered to the currently running VMM (processing block 412). In one embodiment, the delivery of the event to the VMM is performed by synthesizing a

VM exit from the VMM to itself. In another embodiment, the event is delivered as it would be delivered in a non-virtual machine environment (e.g., by vectoring through an interrupt descriptor table or causing a transition to system management mode). It should be noted that a variety of other mechanisms
5 known in the art may be used to deliver the event to the currently running VMM.

[0038] **Figure 5** illustrates an exemplary usage of process 400 in a virtual machine environment having two VMMs, according to one embodiment of the present invention. One of the two VMMs (e.g., VMM 508) is exclusively designated to handle certain platform events (e.g., system management interrupts
10 (SMIs), IPIs, non-maskable interrupts, etc.). VMM 508 is referred to as the Platform VMM (PVMM). In some embodiments, the PVMM is designated to handle only SMIs. In other embodiments, PVMMs handle additional event types, as well as SMIs. In yet other embodiments, multiple PVMMs are used, each handling different platform events.

15 [0039] When a privileged event occurs, routing logic 510 determines the type of the event. If the event is a platform event to be handled by the PVMM 508, routing logic 510 routes it to VMM 508. Otherwise, the event is routed to VMM 506. As illustrated in **Figure 5**, routing logic 510 may route events that occur during the operation of guest software or during the operation of a VMM.

20 [0040] The routing of an event to a VMM may differ depending on what entity was running when the event occurs. If guest software was running, the transition to the VMM selected by the routing logic 510 is performed, in one embodiment, via a VM exit. If the VMM selected by the routing logic was

running when the event occurred, the event may be delivered to the VMM through a variety of mechanisms, as described above with regard to **Figure 4**.

[0041] The use of a second VMM designated exclusively to handle certain platform events eliminates product dependency between hardware vendors and OS vendors. That is, it allows platform hardware vendors to develop the code for the second VMM (the PVMM) independently from OS vendors. Similarly, OS vendors can develop the code for the first VMM independently from the hardware vendors. In addition, the use of two VMMs performing different functionality enhances system security and limits the exposure of the code that needs to be validated for security.

[0042] Referring to **Figure 6**, process 600 begins with processing logic detecting the occurrence of a privileged event during the operation of guest software or a VMM (processing block 602). At processing block 604, processing logic receives information that identifies the privileged event. Next, processing logic accesses a resource (e.g., an in-memory data structure, a processor register, memory in the platform chipset, a register in an input-output device, etc.) that controls the selection of a VMM for handling privileged events (processing block 606) and reads the current value of a resource field associated with the identifier of the privileged event (processing block 608).

[0043] In one embodiment, the identifier of the privileged event is the type of the privileged event, and processing logic uses the type of the privileged event to identify a resource field value associated with this type of privileged event.

[0044] In another embodiment, the identifying information identifies a specific input-output address associated with the privileged event. In this embodiment, processing logic analyzes the input-output access associated with the privileged event, determines what input-output address is being accessed, and then uses this input-output address to find a resource field value associated with an input-output address range to which the determined input-output address belongs. In one embodiment, the values of the resource fields are set during the initialization and cannot be modified. In another embodiment, the values of the resource fields can be dynamically configured by a VMM (e.g., a main VMM) based on such factors as security requirements, the VMM's knowledge of the platform, etc. This dynamic configuration of resource field values allows for VMM functionality to be partitioned, possibly improving security, system performance, development methodologies, etc.

[0045] Further, processing logic determines which VMM is designated to handle the privileged event based on the current value of the corresponding field (processing block 610) and, if the currently-operating entity is not the designated VMM (decision box 612), transitions control to the designated VMM (processing block 614). In one embodiment, the transition to the designated VMM may be accomplished by generating a VM exit. Alternatively, the transition to the designated VMM may be accomplished through any other mechanism in the art.

[0046] If the designated VMM is the currently-operating entity, processing logic delivers the event to that VMM (processing block 616), as described above in conjunction with **Figure 4** and **Figure 5**.

[0047] Referring to **Figure 7**, process 700 begins with processing logic detecting the occurrence of a privileged event during the operation of guest software or a VMM (processing block 702). At processing block 704, processing logic identifies resource usage/load parameters of each VMM in the system.

5 Next, processing logic evaluates the resource usage and/or load parameters of the VMMs in the context of load balancing (processing block 706) and determines which VMM is the best candidate for handling the privileged event based on the above load balancing evaluation (processing block 708). Further, if the currently operating entity is not the VMM that was identified as the best candidate
10 (decision box 712), processing logic transitions control to the identified VMM (processing block 714). In one embodiment, the transition to the identified VMM may be accomplished by generating a VM exit. Alternatively, the transition to the identified VMM may be accomplished through any other mechanism in the art.

[0048] If the designated VMM is the currently-operating entity,
15 processing logic delivers the event to that VMM (processing block 716), as described above in conjunction with **Figures 4-6**.

[0049] Thus, a method and apparatus for handling privileged events using multiple VMMs have been described. It is to be understood that the above description is intended to be illustrative, and not restrictive. Many other
20 embodiments will be apparent to those of skill in the art upon reading and understanding the above description. The scope of the invention should, therefore, be determined with reference to the appended claims, along with the full scope of equivalents to which such claims are entitled.

CLAIMS

What is claimed is:

1. A method comprising:
 - detecting an occurrence of a privileged event;
 - 5 determining which one of a plurality of virtual machine monitors (VMMs)
is to handle the privileged event; and
 - transitioning control to said one of the plurality of VMMs.
2. The method of claim 1 wherein the plurality of VMMs includes a main
10 VMM and one or more parallel VMMs.
3. The method of claim 1 wherein determining which one of a plurality of
VMMs is to handle the privileged event comprises:
 - determining a type of the privileged event; and
 - 15 identifying one of the plurality of VMMs that is designated to handle
privileged events of the determined type.
4. The method of claim 1 wherein determining which one of a plurality of
VMMs is to handle the privileged event comprises:
 - 20 accessing a field associated with the privileged event in a resource; and
 - identifying one of the plurality of VMMs that is designated to handle the
privileged event based on a value of the field.

5. The method of claim 4 wherein the field associated with the privileged event is a field associated with a type of the privileged event.

6. The method of claim 4 wherein the field associated with the privileged event is a field associated with an input-output address range to which an input-output address of the privileged event belongs.

7. The method of claim 4 wherein the value of the field associated with the privileged event is either predetermined or dynamically configurable.

10

8. The method of claim 1 wherein determining which one of a plurality of VMMS is to handle the privileged event comprises:

evaluating resource usage parameters of the plurality of VMMS; and

identifying one of the plurality of VMMS that is designated to handle the

15 privileged event based on evaluation of the resource usage parameters.

9. The method of claim 4 wherein the resource resides in any one of a memory, a processor, a chipset, and an input-output device.

20 10. The method of claim 1 wherein the privileged event represents any one of an interrupt, an exception, an execution of a privileged instruction, and a platform event.

11. The method of claim 1 wherein the privileged event occurs during an operation of guest software.

12. The method of claim 1 wherein the privileged event occurs during an
5 operation of one of the plurality of VMMs.

13. A system comprising:
a plurality of virtual machine monitors (VMMs); and
routing logic to detect an occurrence of a privileged event, to determine
10 which one of the plurality of VMMs is to handle the privileged event, and to
transition control to said one of the plurality of VMMs.

14. The system of claim 13 wherein the plurality of VMMs includes a main
VMM and one or more parallel VMMs.

15

15. The system of claim 13 wherein the routing logic is to determine which one
of the plurality of VMMs is to handle the privileged event by determining a type
of the privileged event, and identifying one of the plurality of VMMs that is
designated to handle privileged events of the determined type.

20

16. The system of claim 13 wherein the routing logic is to determine which one
of the plurality of VMMs is to handle the privileged event by evaluating resource
usage parameters of the plurality of VMMs, and identifying one of the plurality of

VMMs that is designated to handle the privileged event based on evaluation of the resource usage parameters.

17. The system of claim 13 wherein the privileged event represents any one of
5 an interrupt, an exception, an execution of a privileged instruction, and a platform event.

18. The system of claim 13 wherein the privileged event occurs during an operation of guest software.

10

19. The system of claim 13 wherein the privileged event occurs during an operation of one of the plurality of VMMs.

20. A system comprising:

15 a memory having stored therein guest software and a plurality of virtual machine monitors (VMMs); and

a processor, coupled to the memory, to execute the guest software, to detect an occurrence of a privileged event, to determine which one of the plurality of VMMs is to handle the privileged event, and to transition control to said one of
20 the plurality of VMMs.

21. The system of claim 20 wherein the plurality of VMMs includes a main VMM and one or more parallel VMMs.

22. The system of claim 20 wherein the processor is to determine which one of the plurality of VMMs is to handle the privileged event by determining a type of the privileged event, and identifying one of the plurality of VMMs that is
5 designated to handle privileged events of the determined type.

23. The system of claim 20 wherein the processor is to determine which one of the plurality of VMMs is to handle the privileged event by evaluating resource usage parameters of the plurality of VMMs, and identifying one of the plurality of
10 VMMs that is designated to handle the privileged event based on evaluation of the resource usage parameters.

24. The system of claim 20 wherein the privileged event represents any one of an interrupt, an exception, an execution of a privileged instruction, and a platform
15 event.

25. The system of claim 20 wherein the privileged event occurs during operation of any one of guest software and one of the plurality of VMMs.

20 26. A machine-readable medium containing instructions which, when executed by a processing system, cause the processing system to perform a method, the method comprising:

detecting an occurrence of a privileged event;

determining which one of a plurality of virtual machine monitors (VMMs) is to handle the privileged event; and

transitioning control to said one of the plurality of VMMs.

5 27. The machine-readable medium of claim 26 wherein the plurality of VMMs includes a main VMM and one or more parallel VMMs.

28. The machine-readable medium of claim 26 wherein the privileged event represents any one of an interrupt, an exception, an execution of a privileged
10 instruction, and a platform event.

29. The machine-readable medium of claim 26 wherein the privileged event occurs during operation of any one of guest software and one of the plurality of VMMs.

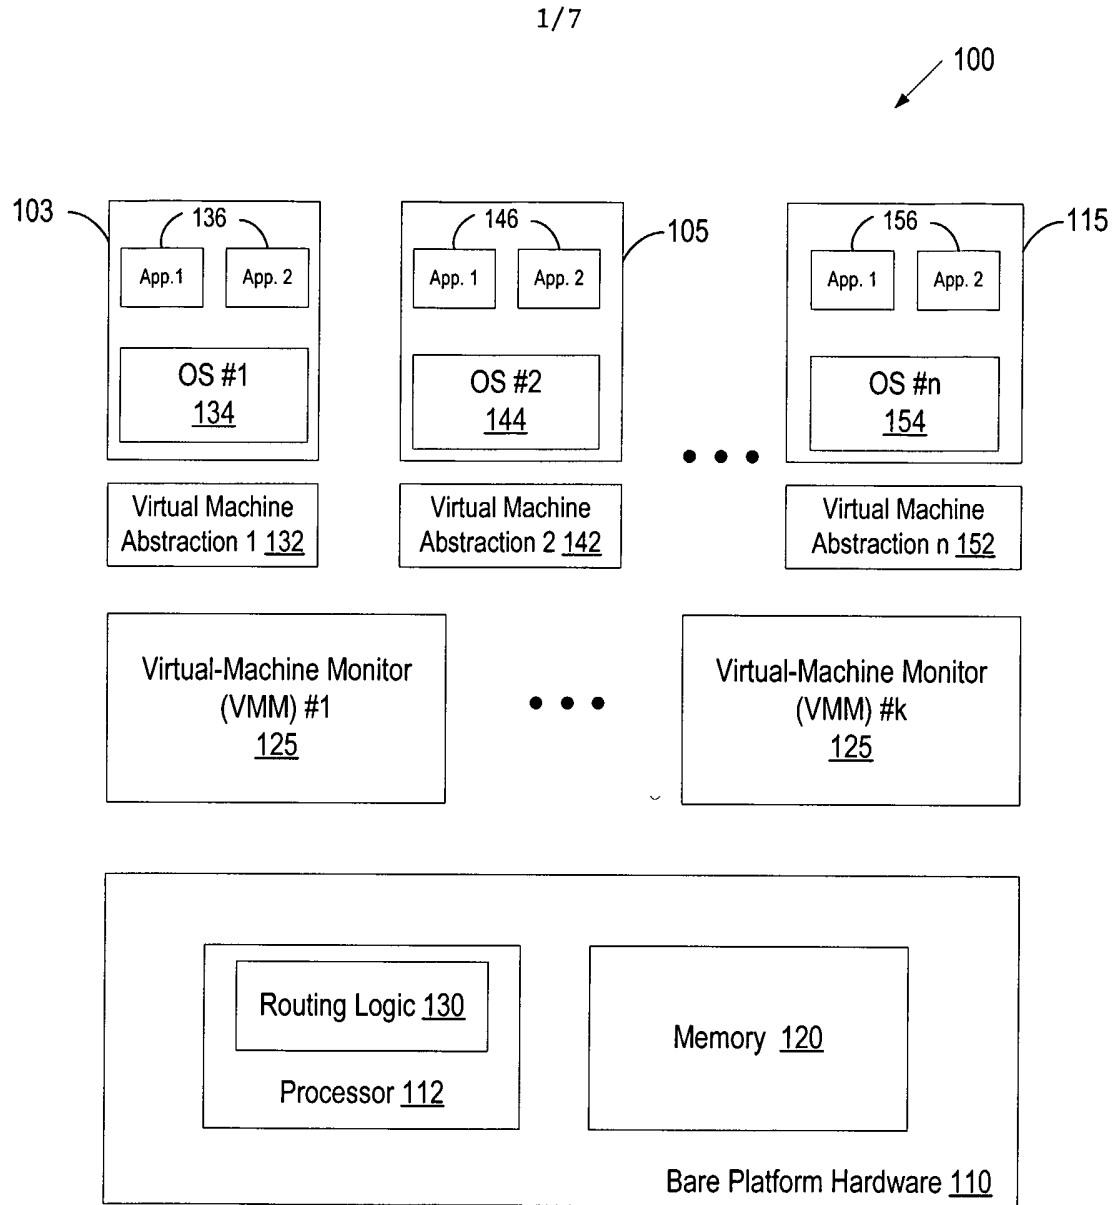


FIG. 1

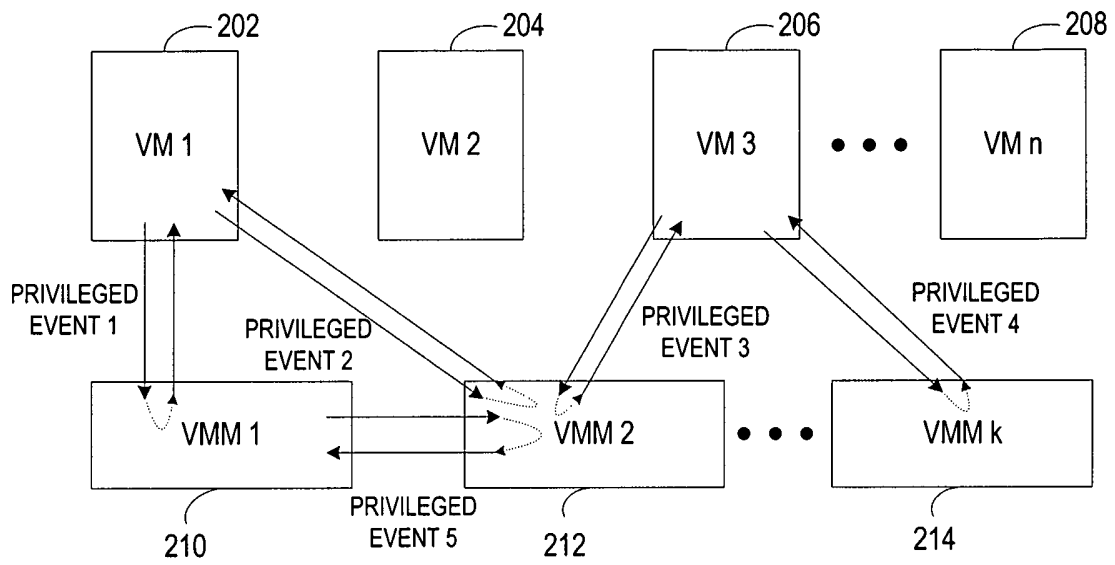


FIG. 2

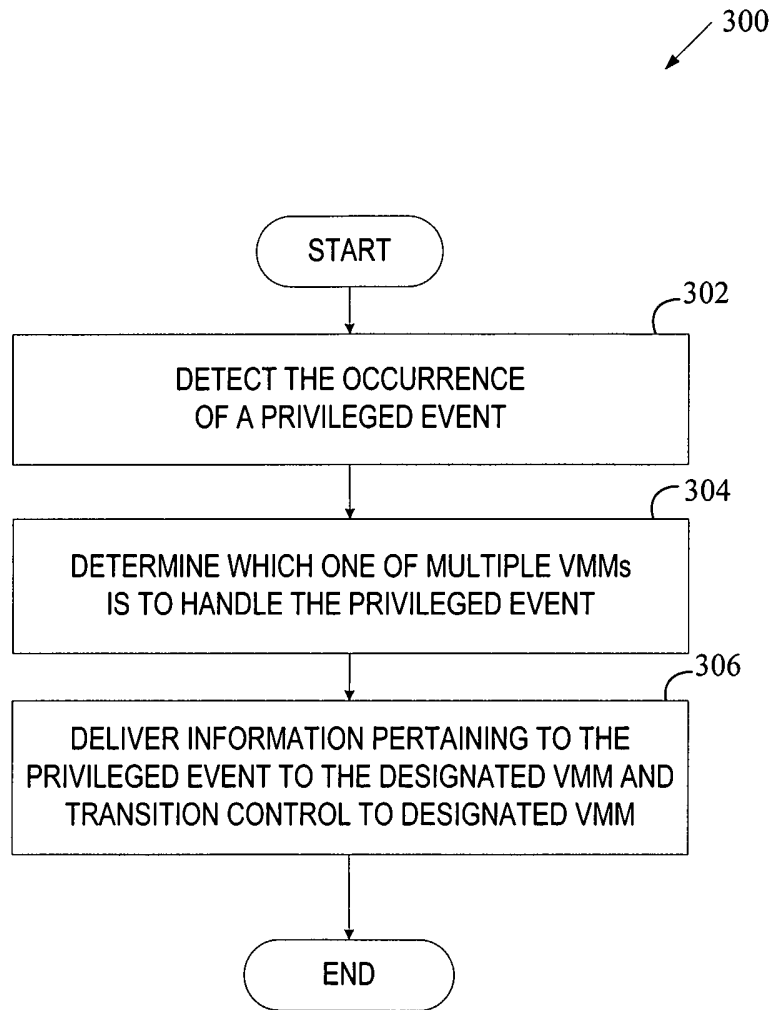


FIG. 3

4/7

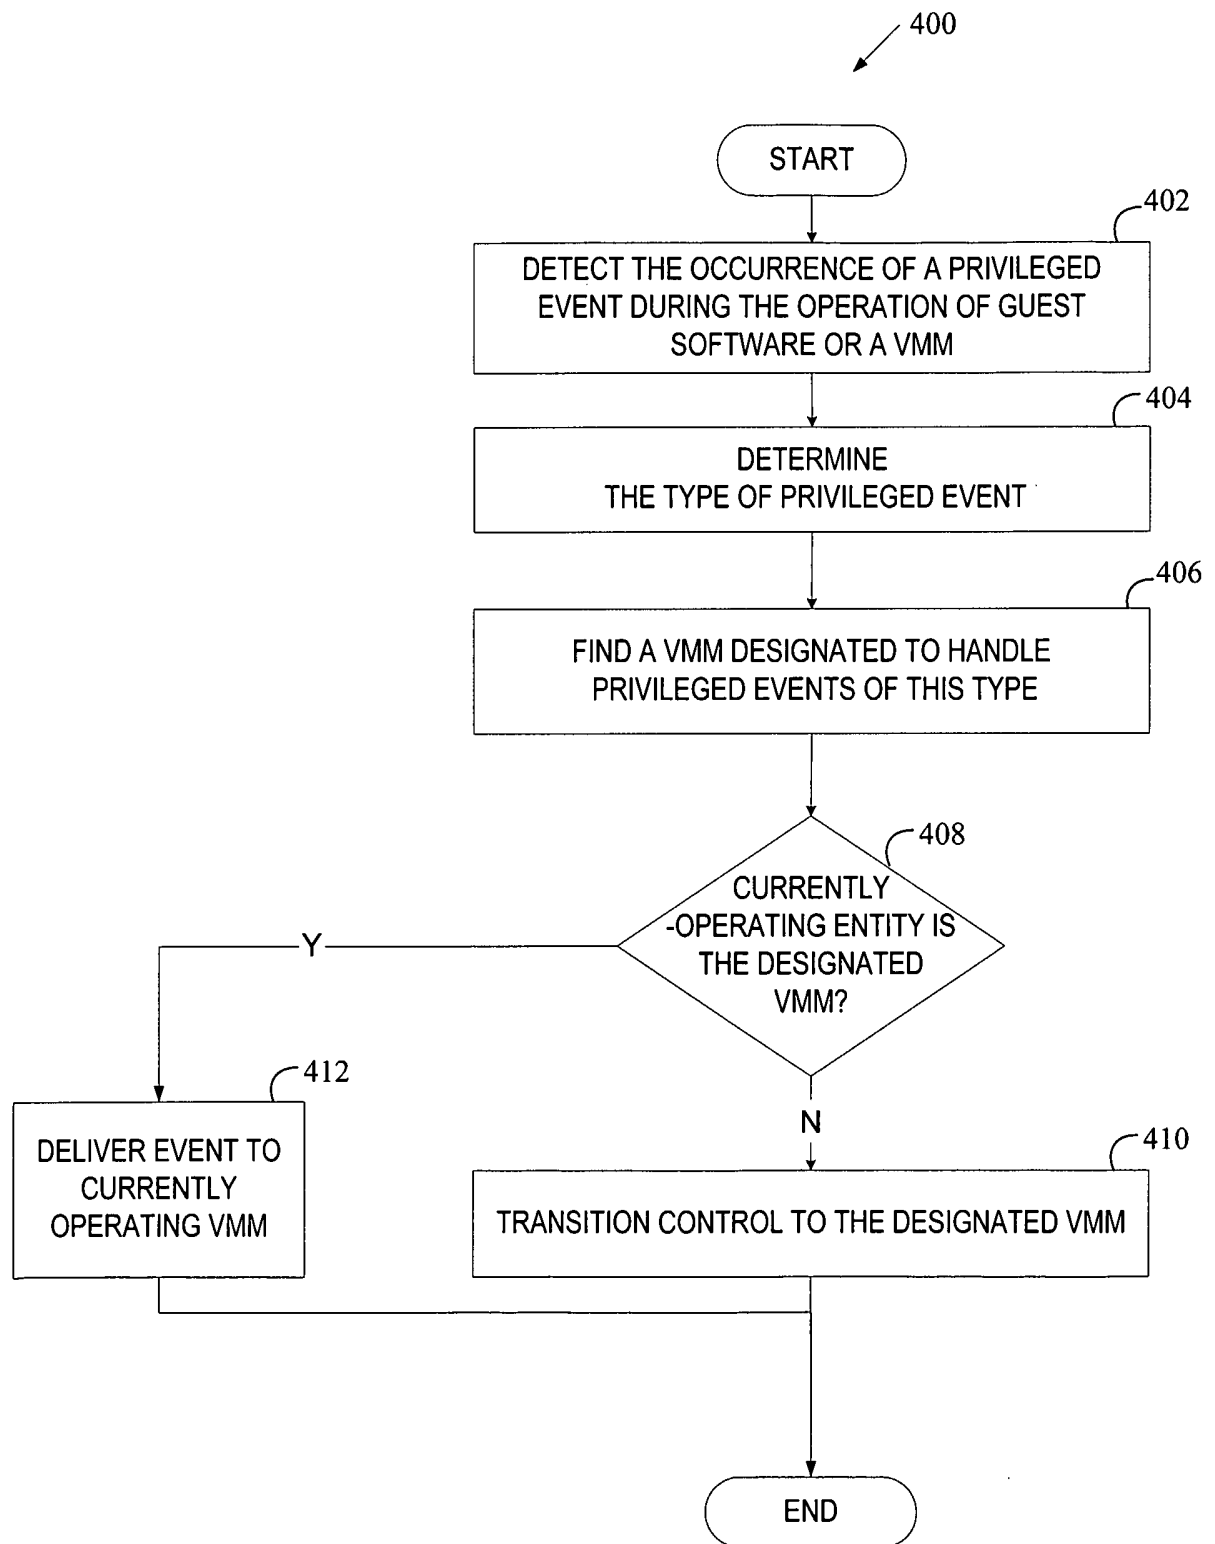


FIG. 4

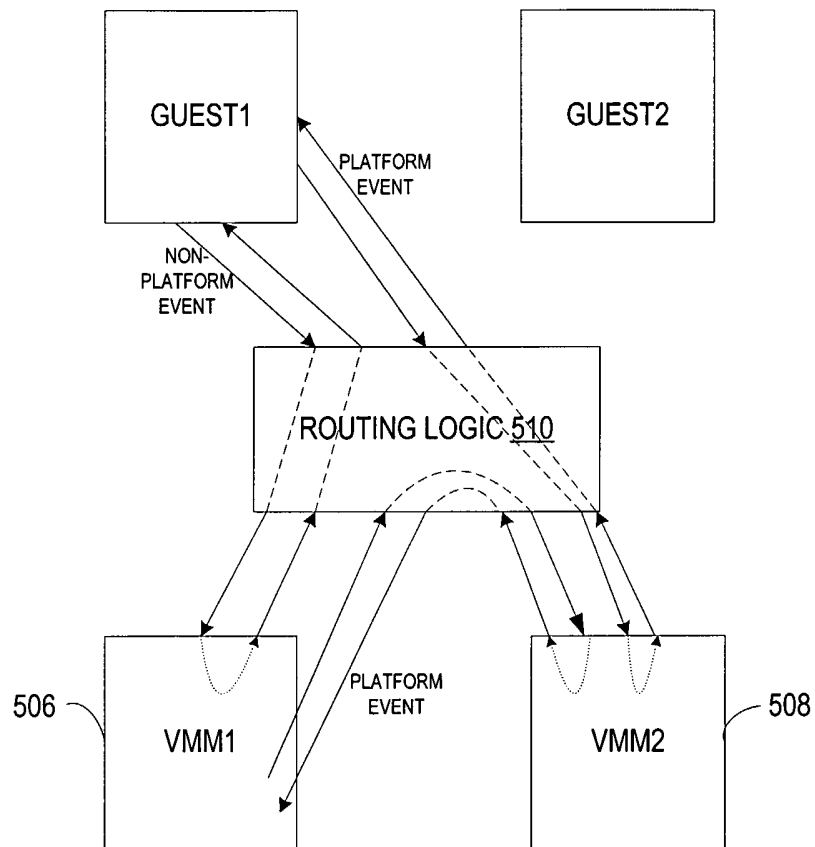


FIG. 5

6/7

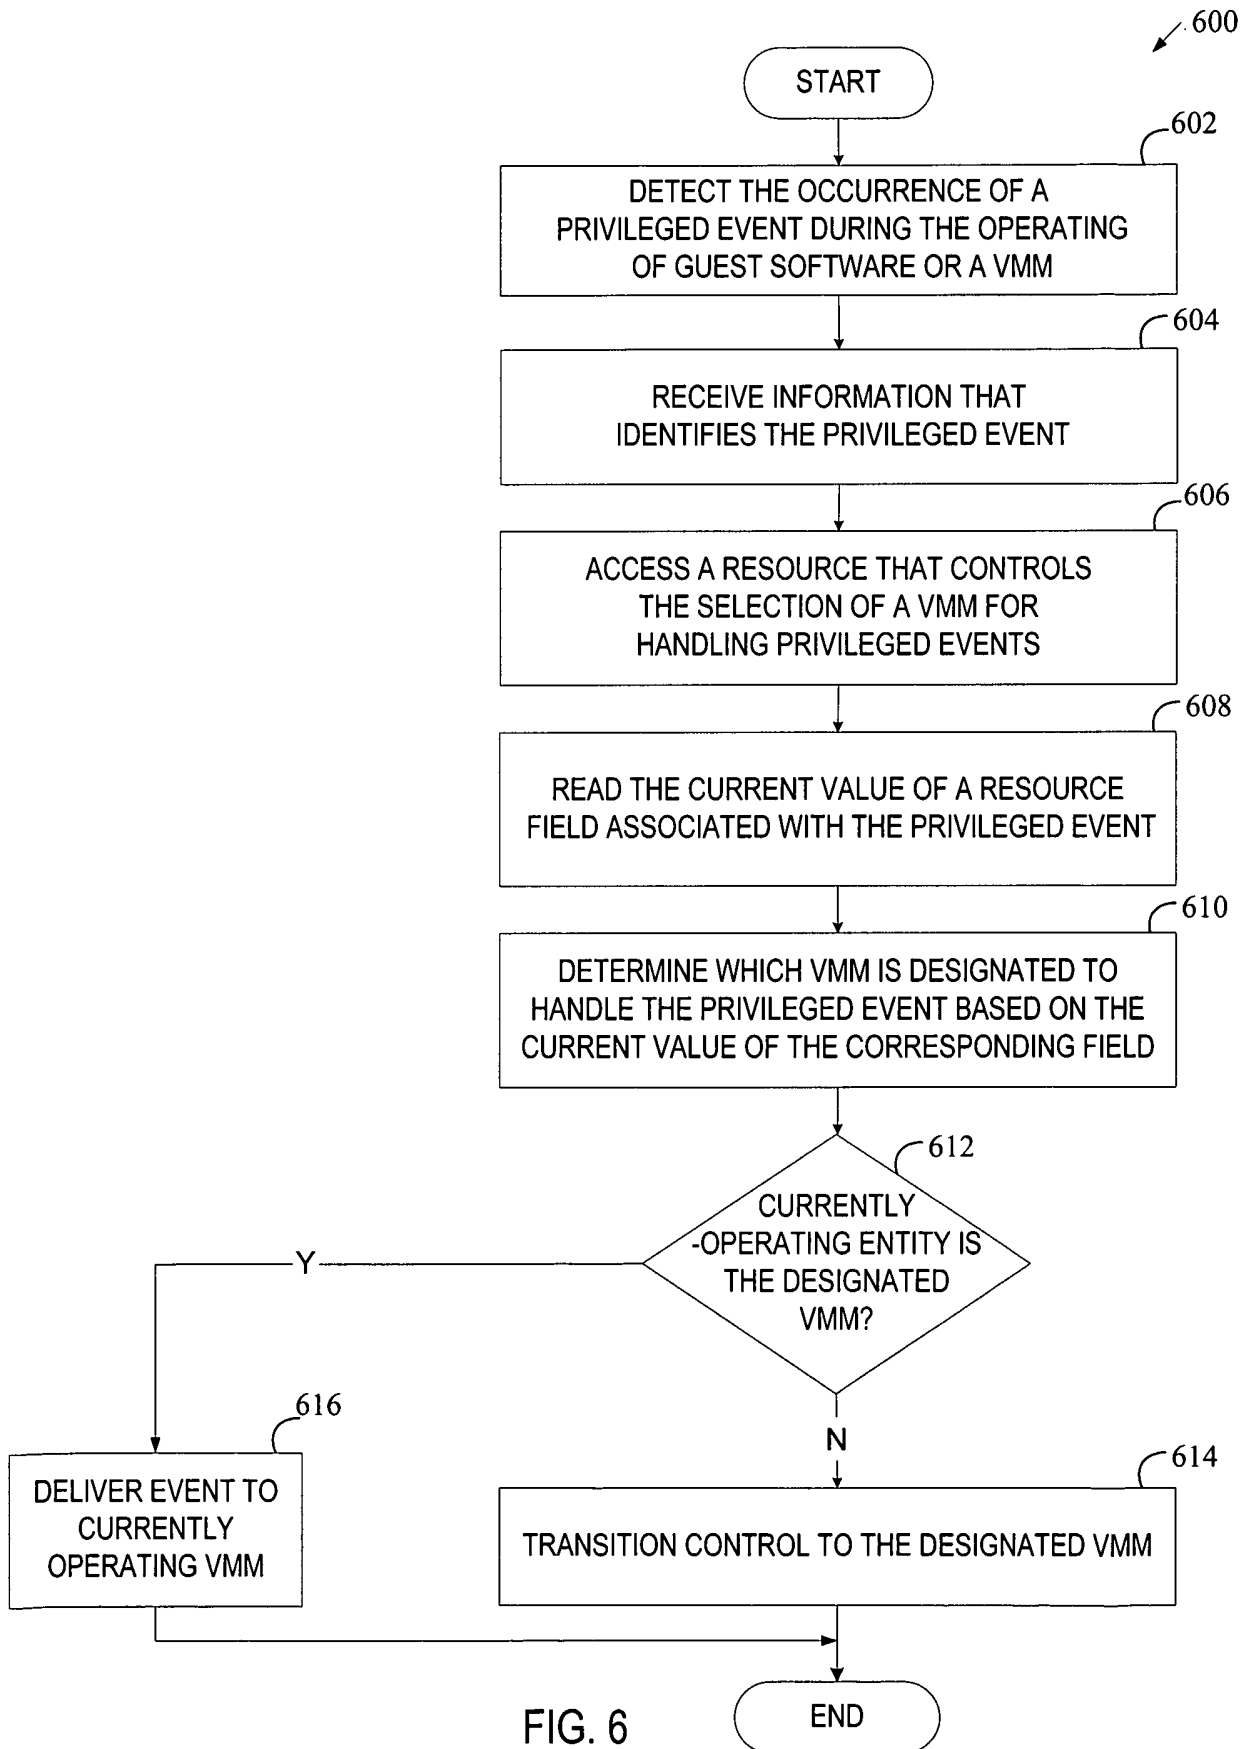


FIG. 6

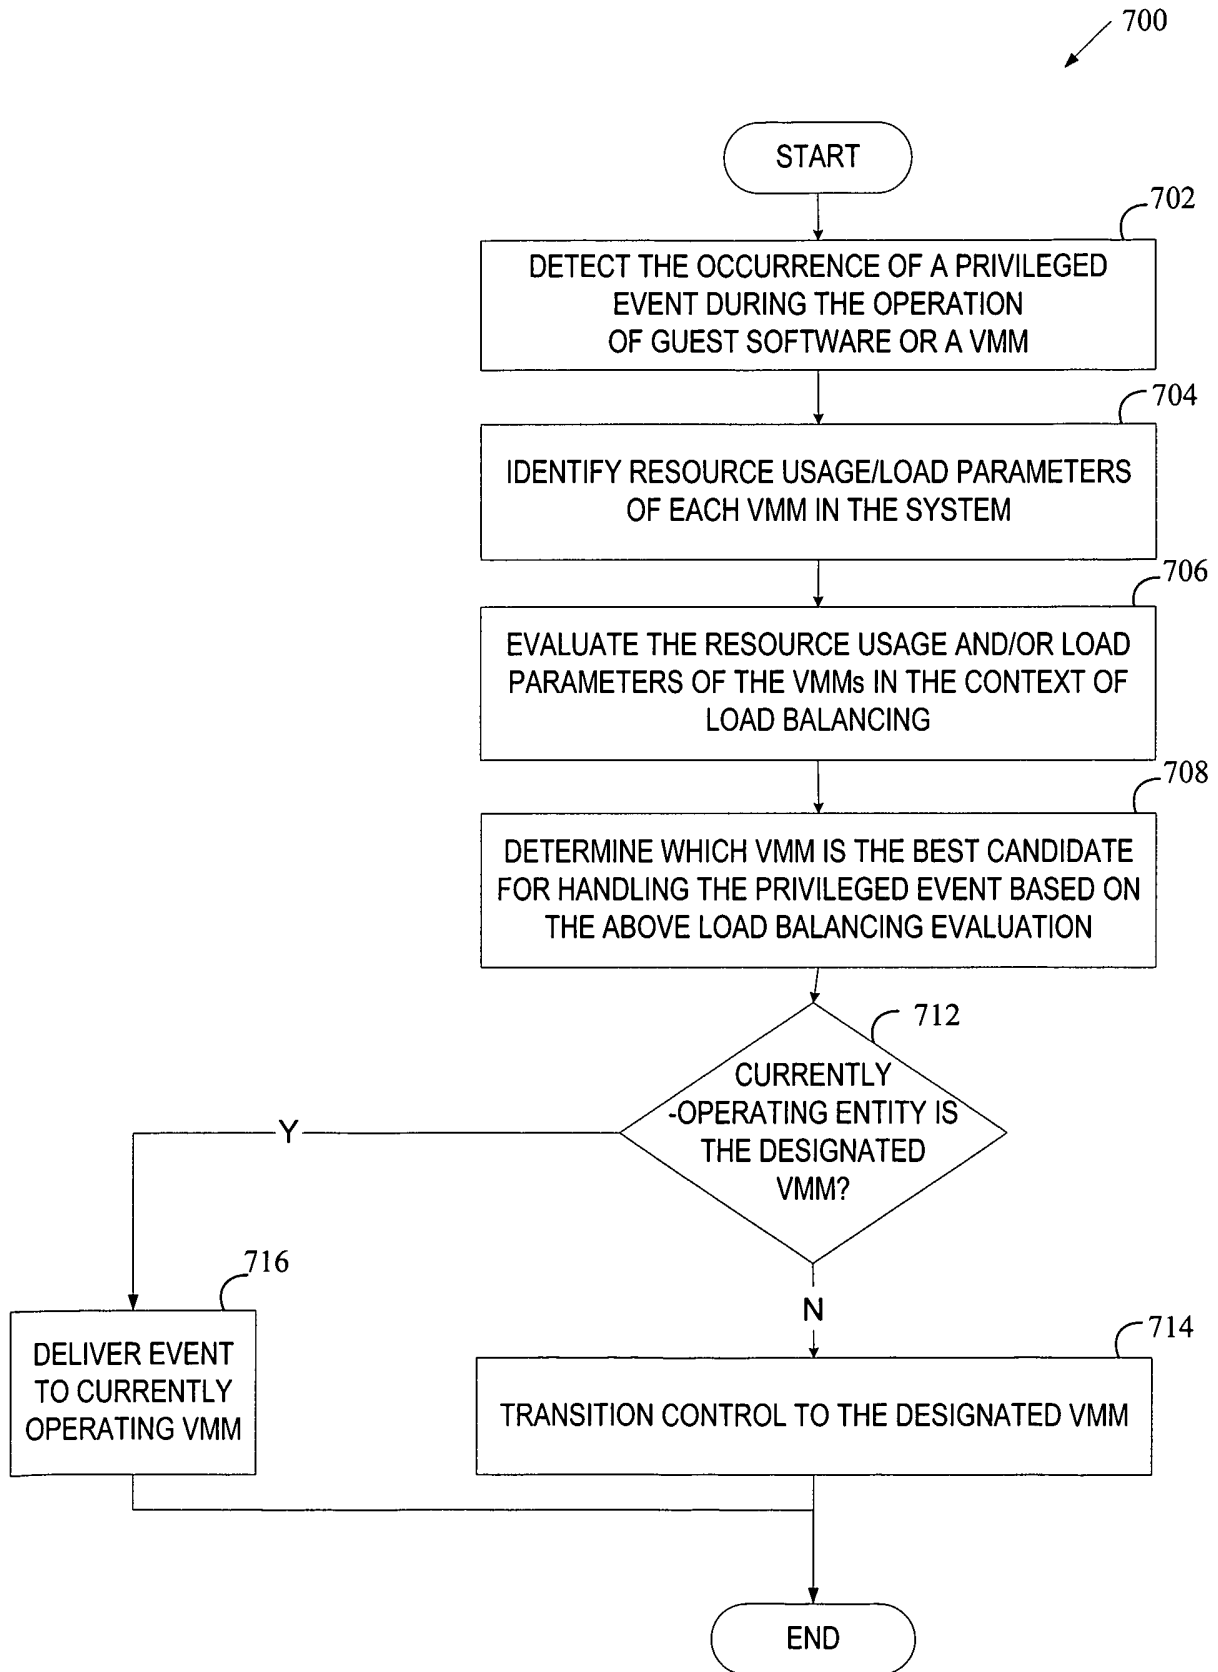


FIG. 7

INTERNATIONAL SEARCH REPORT

International Application No
PCT/US2004/029611

A. CLASSIFICATION OF SUBJECT MATTER IPC 7 G06F9/455		
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 7 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 practical, search terms used) EPO-Internal, INSPEC		
C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 4 843 541 A (BEAN ET AL) 27 June 1989 (1989-06-27) column 1, line 25 - column 2, line 10 column 6, line 56 - column 7, line 9 column 9, line 25 - column 10, line 50 -----	1-29
X	GOLDBERG R P: "Survey of virtual machine research" COMPUTER USA, vol. 7, no. 6, June 1974 (1974-06), pages 34-45, XP009043730 ISSN: 0018-9162 page 37, right-hand column, line 7 - line 23 page 38, left-hand column - right-hand column; figure 3 ----- -/--	1,2,13, 14,20, 21,26,27
<input checked="" type="checkbox"/> Further documents are listed in the continuation of box C.		
<input checked="" type="checkbox"/> Patent family members are listed in 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 document 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 <p style="text-align: center;">11 February 2005</p>	Date of mailing of the international search report <p style="text-align: center;">24/02/2005</p>	
Name and mailing address of the ISA European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Tx. 31 651 epo nl, Fax: (+31-70) 340-3016	Authorized officer <p style="text-align: center;">Ross, C</p>	

INTERNATIONAL SEARCH REPORT

International Application No
PCT/US2004/029611

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT		
Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>US 6 496 847 B1 (BUGNION EDOUARD ET AL) 17 December 2002 (2002-12-17) column 16, line 5 - column 17, line 34; figure 6</p> <p style="text-align: center;">-----</p>	1-29
A	<p>JOHN SCOTT ROBIN ET AL: "Analysis of the Pentium's Ability to Support a Secure Virtual Machine Monitor" PROCEEDINGS OF THE USENIX SECURITY SYMPOSIUM, XX, XX, 14 August 2000 (2000-08-14), pages 1-17, XP002247347 Section 2.1 "Virtual Machine Monitors Types" up to and including section 2.4 "Type II VMM"</p> <p style="text-align: center;">-----</p>	1-29

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/US2004/029611

Patent document cited in search report	A	Publication date	Patent family member(s)	Publication date
US 4843541	A	27-06-1989	AU 606187 B2	31-01-1991
			AU 2001488 A	02-02-1989
			BR 8803742 A	14-02-1989
			CA 1305799 C	28-07-1992
			DE 3850181 D1	21-07-1994
			DE 3850181 T2	12-01-1995
			EP 0301275 A2	01-02-1989
			JP 1037636 A	08-02-1989
			JP 1945910 C	23-06-1995
			JP 6073108 B	14-09-1994
US 6496847	B1	17-12-2002	US 6704925 B1	09-03-2004
			US 6795966 B1	21-09-2004
			US 6397242 B1	28-05-2002
			US 6785886 B1	31-08-2004