



(19) **United States**
(12) **Patent Application Publication**
MOORING et al.

(10) **Pub. No.: US 2016/0156665 A1**
(43) **Pub. Date: Jun. 2, 2016**

(54) **SYSTEMS AND METHODS INVOLVING ASPECTS OF HARDWARE VIRTUALIZATION SUCH AS HYPERVISOR, DETECTION AND INTERCEPTION OF CODE OR INSTRUCTION EXECUTION INCLUDING API CALLS, AND/OR OTHER FEATURES**

(52) **U.S. Cl.**
CPC **H04L 63/20** (2013.01); **G06F 9/45545** (2013.01); **G06F 9/45558** (2013.01); **H04L 63/1416** (2013.01); **G06F 2009/45587** (2013.01); **G06F 2009/45591** (2013.01)

(71) Applicant: **Lynx Software Technologies, Inc.**, San Jose, CA (US)

(57) **ABSTRACT**

(72) Inventors: **Edward T. MOORING**, Santa Clara, CA (US); **Craig HOWARD**, Belmont, CA (US)

Systems, methods, computer readable media and articles of manufacture consistent with innovations herein are directed to computer virtualization, computer security and/or memory access. According to some illustrative implementations, innovations herein may utilize and/or involve a separation kernel hypervisor which may include the use of a guest operating system virtual machine protection domain, a virtualization assistance layer, and/or a instruction execution detection/interception mechanism (which may be proximate in temporal and/or spatial locality to malicious code, but isolated from it). The instruction execution detection/interception mechanism may perform processing, inter alia, for detection and/or notification of, and actions upon by a monitoring guest, code execution by a monitored guest involving predetermined physical memory locations, such as API calls. Such actions may include interception of API calls within the monitored guest and simulation thereof by the monitoring guest or another authorized guest.

(21) Appl. No.: **14/955,018**

(22) Filed: **Nov. 30, 2015**

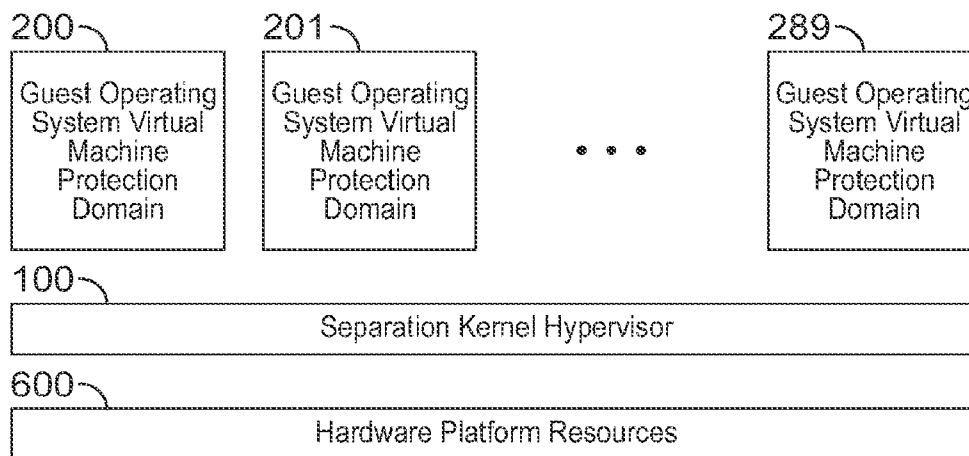
Related U.S. Application Data

(63) Continuation of application No. 14/714,241, filed on May 15, 2015, now Pat. No. 9,203,855.

(60) Provisional application No. 61/993,296, filed on May 15, 2014.

Publication Classification

(51) **Int. Cl.**
H04L 29/06 (2006.01)
G06F 9/455 (2006.01)



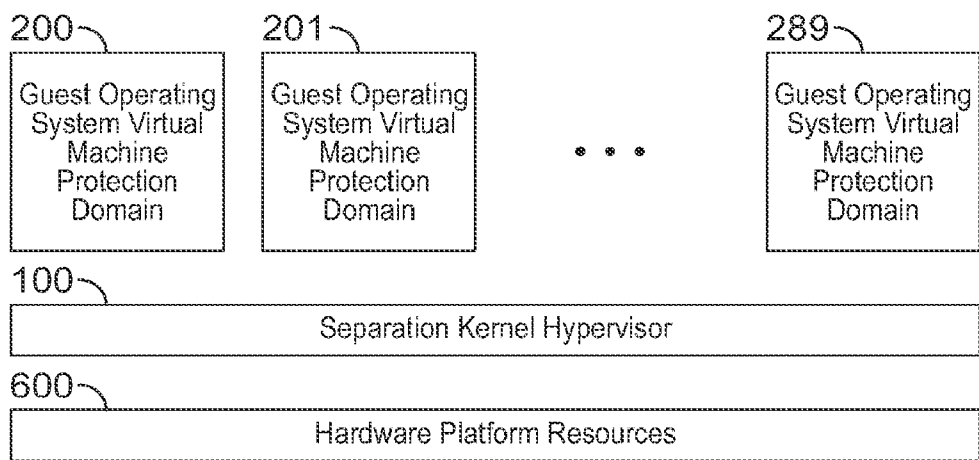


FIG. 1

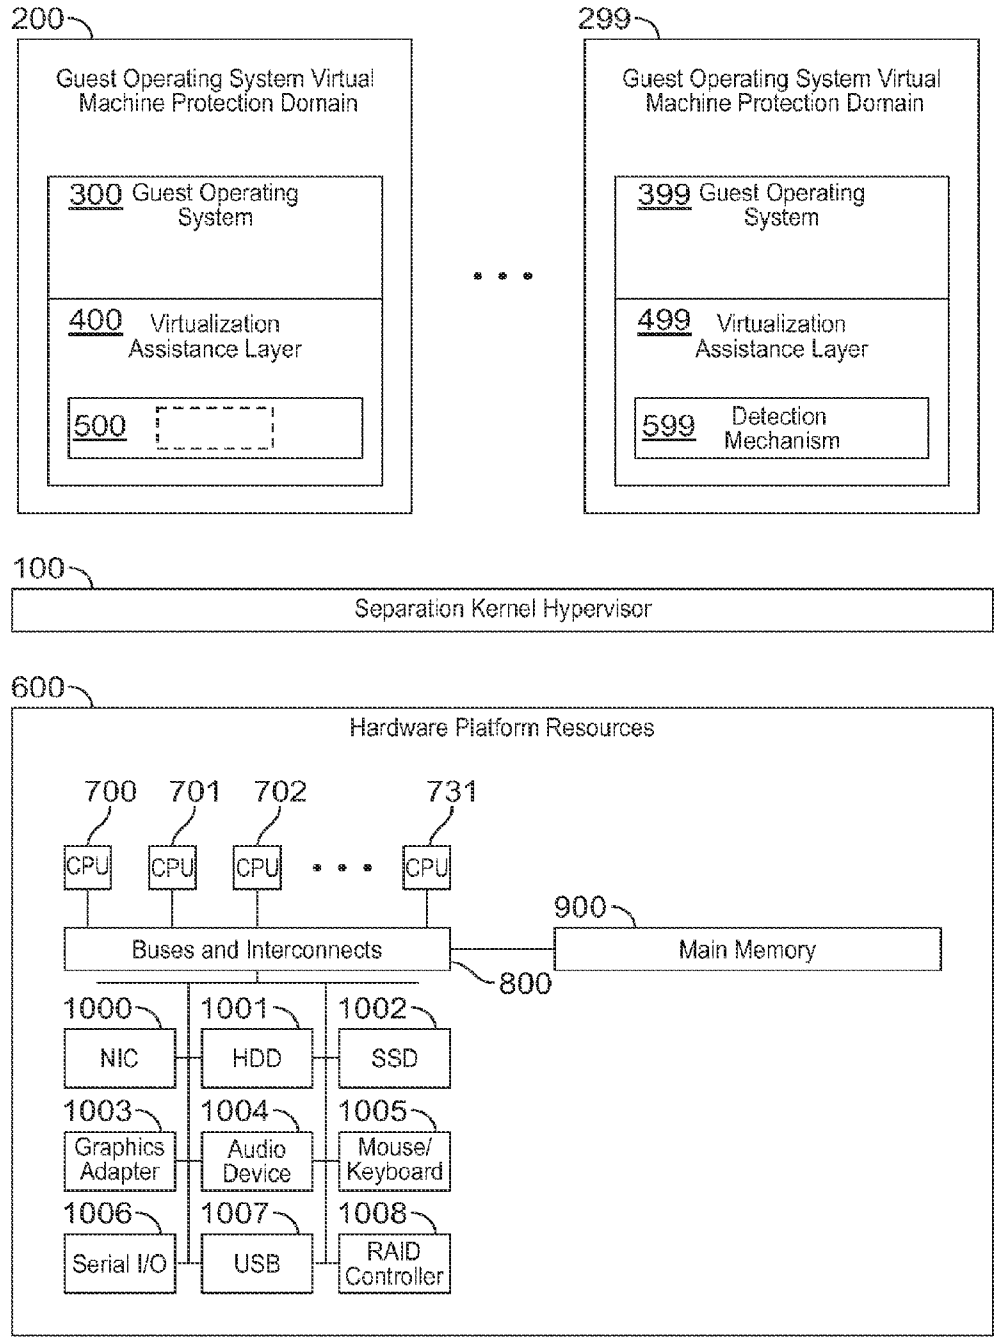


FIG. 2A

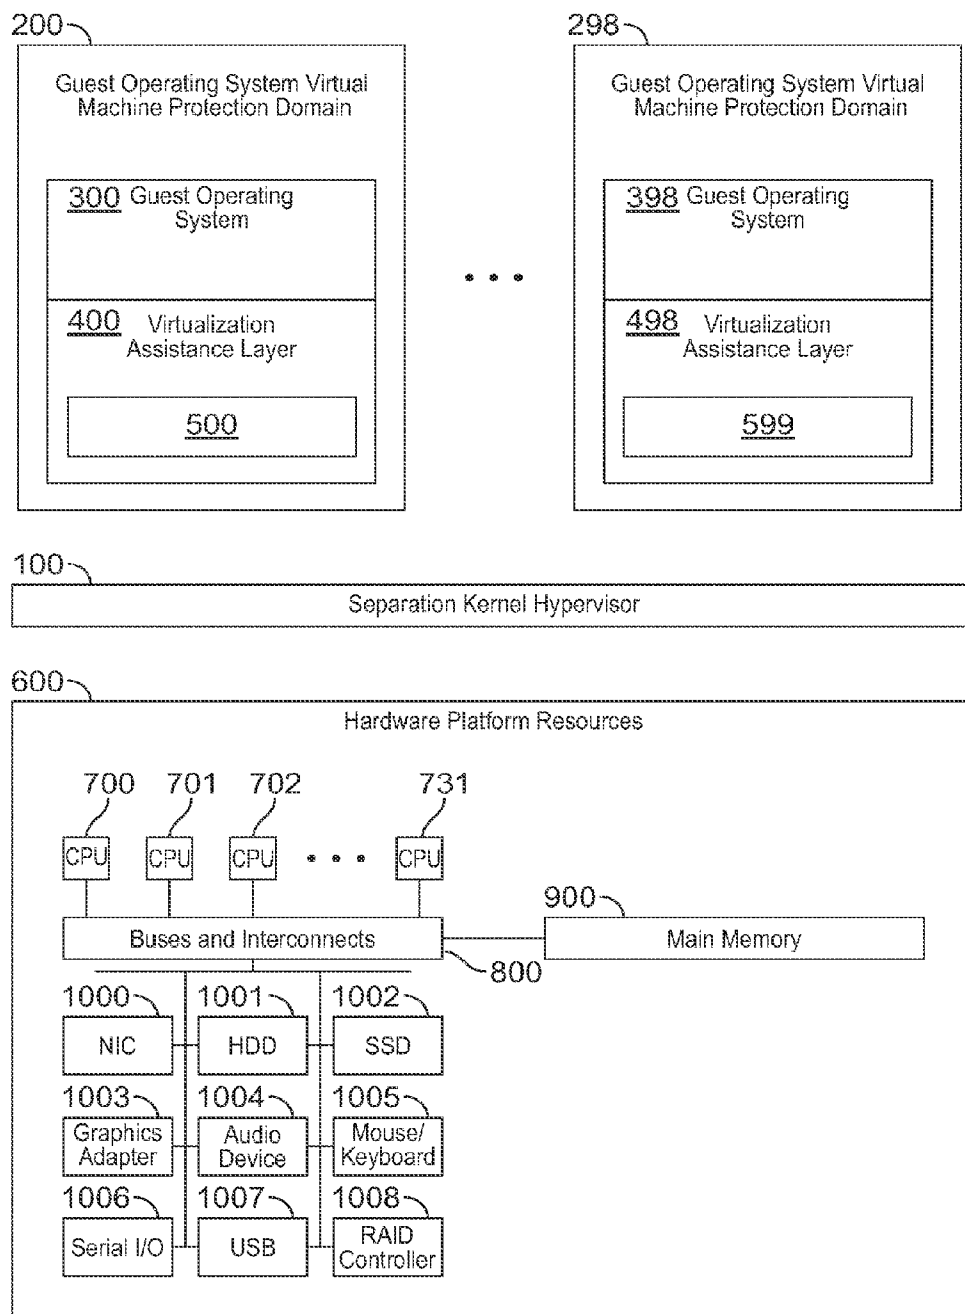


FIG. 2B

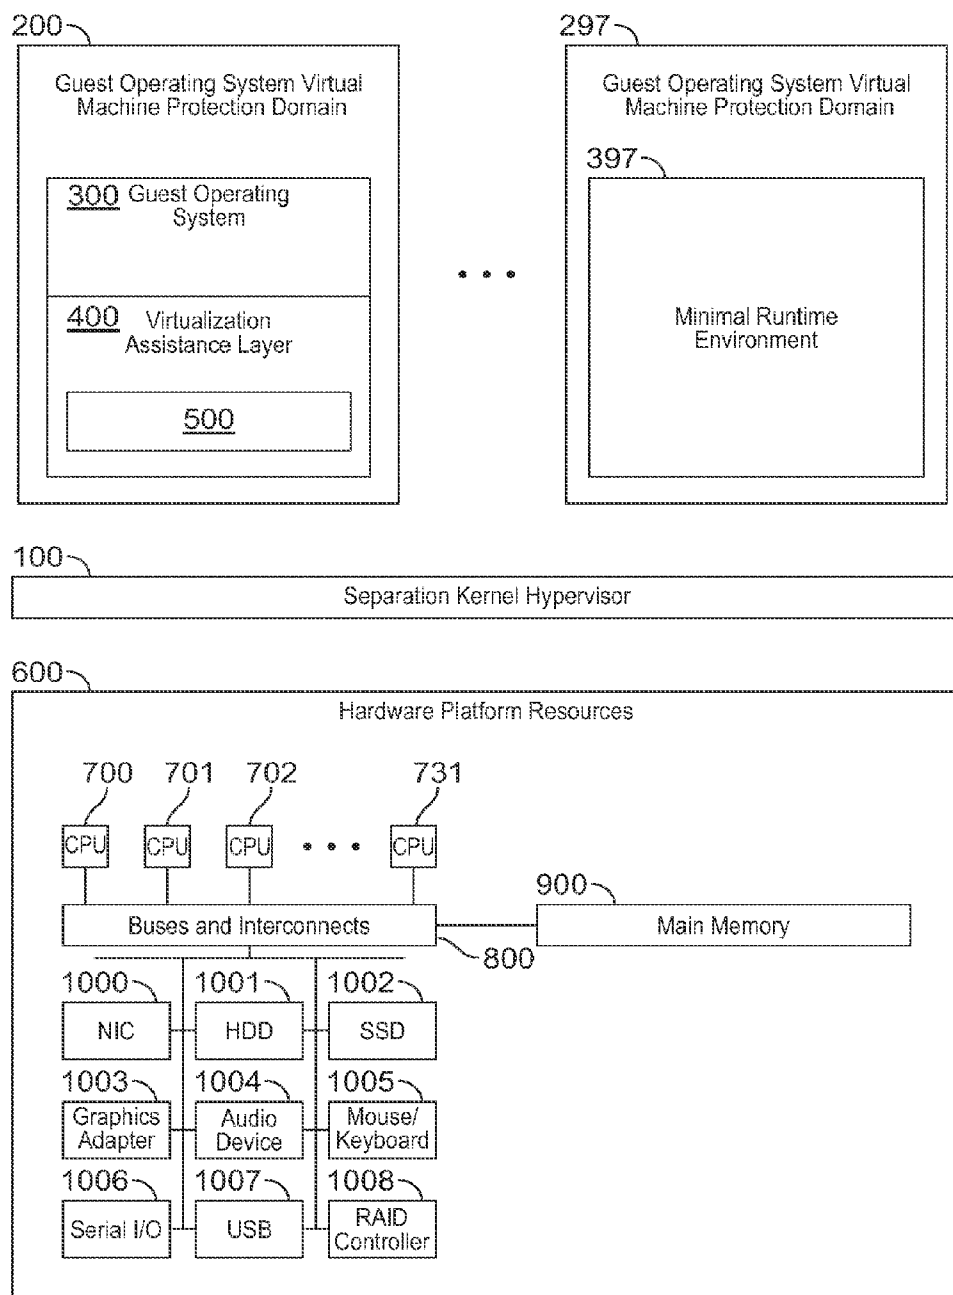


FIG. 2C

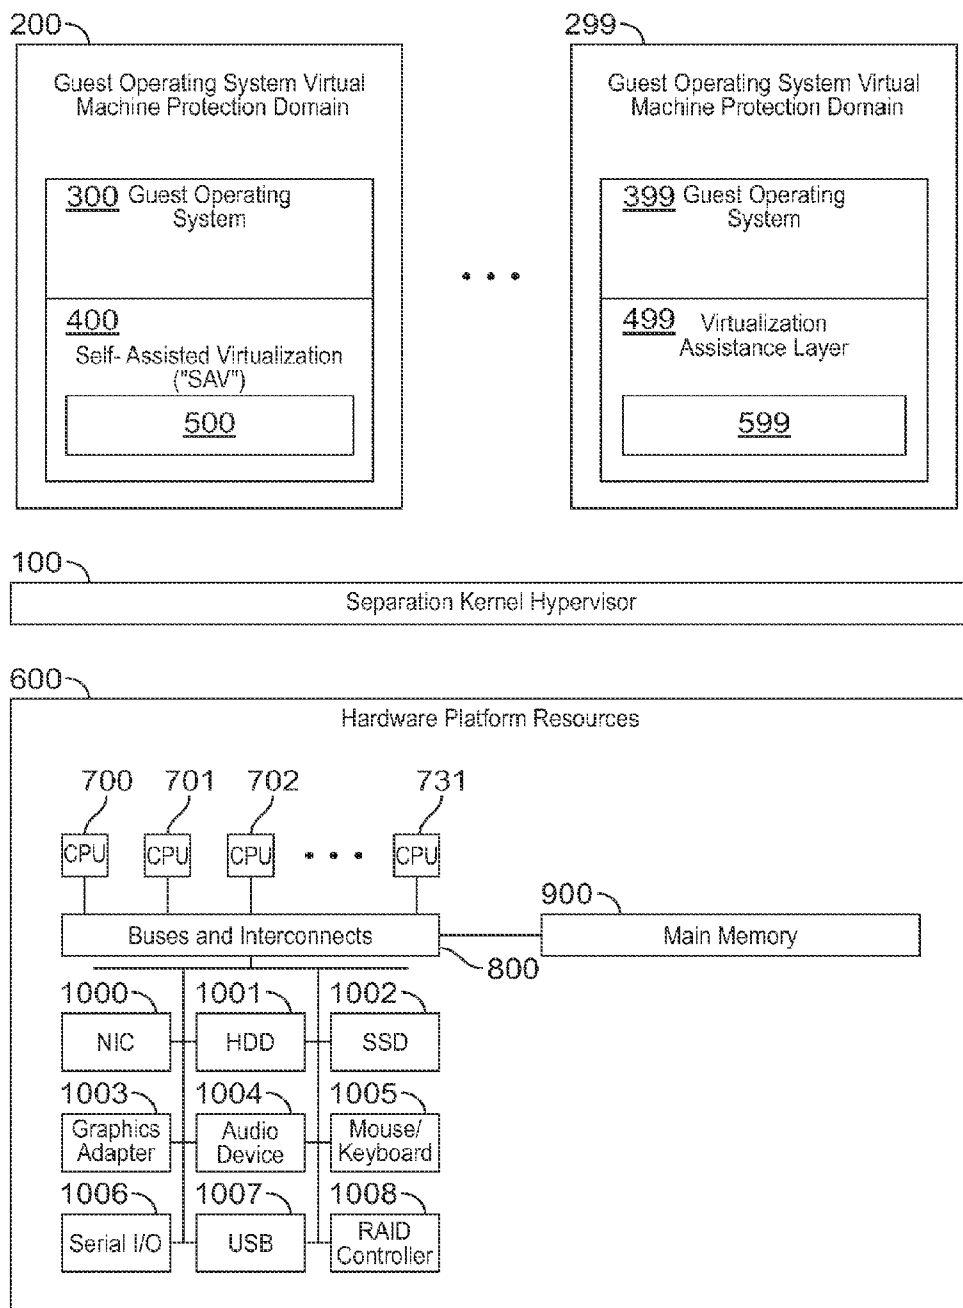


FIG. 2D

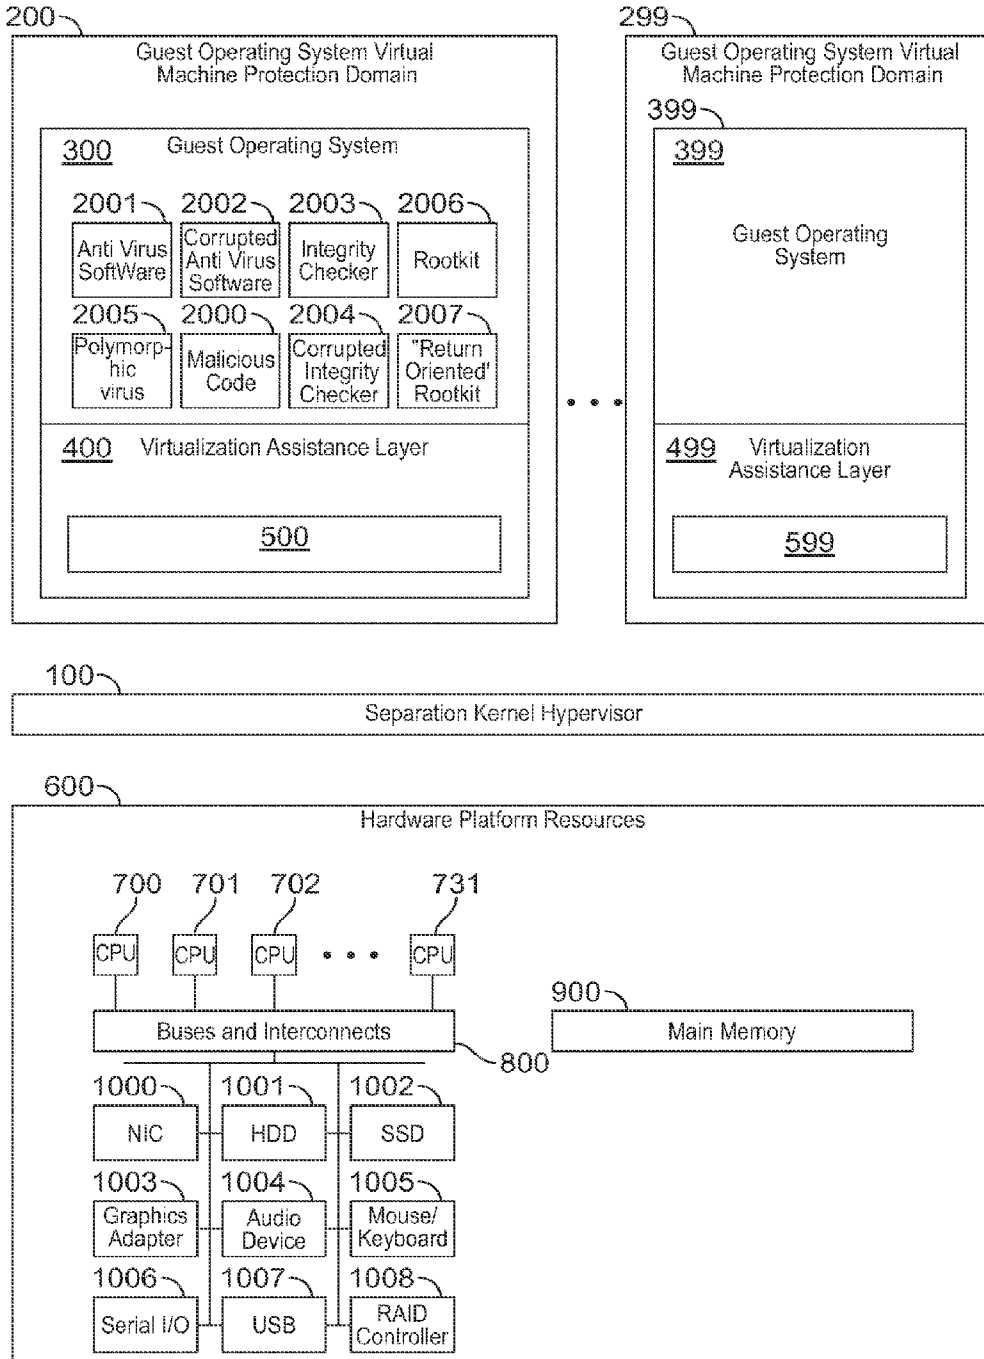


FIG. 3

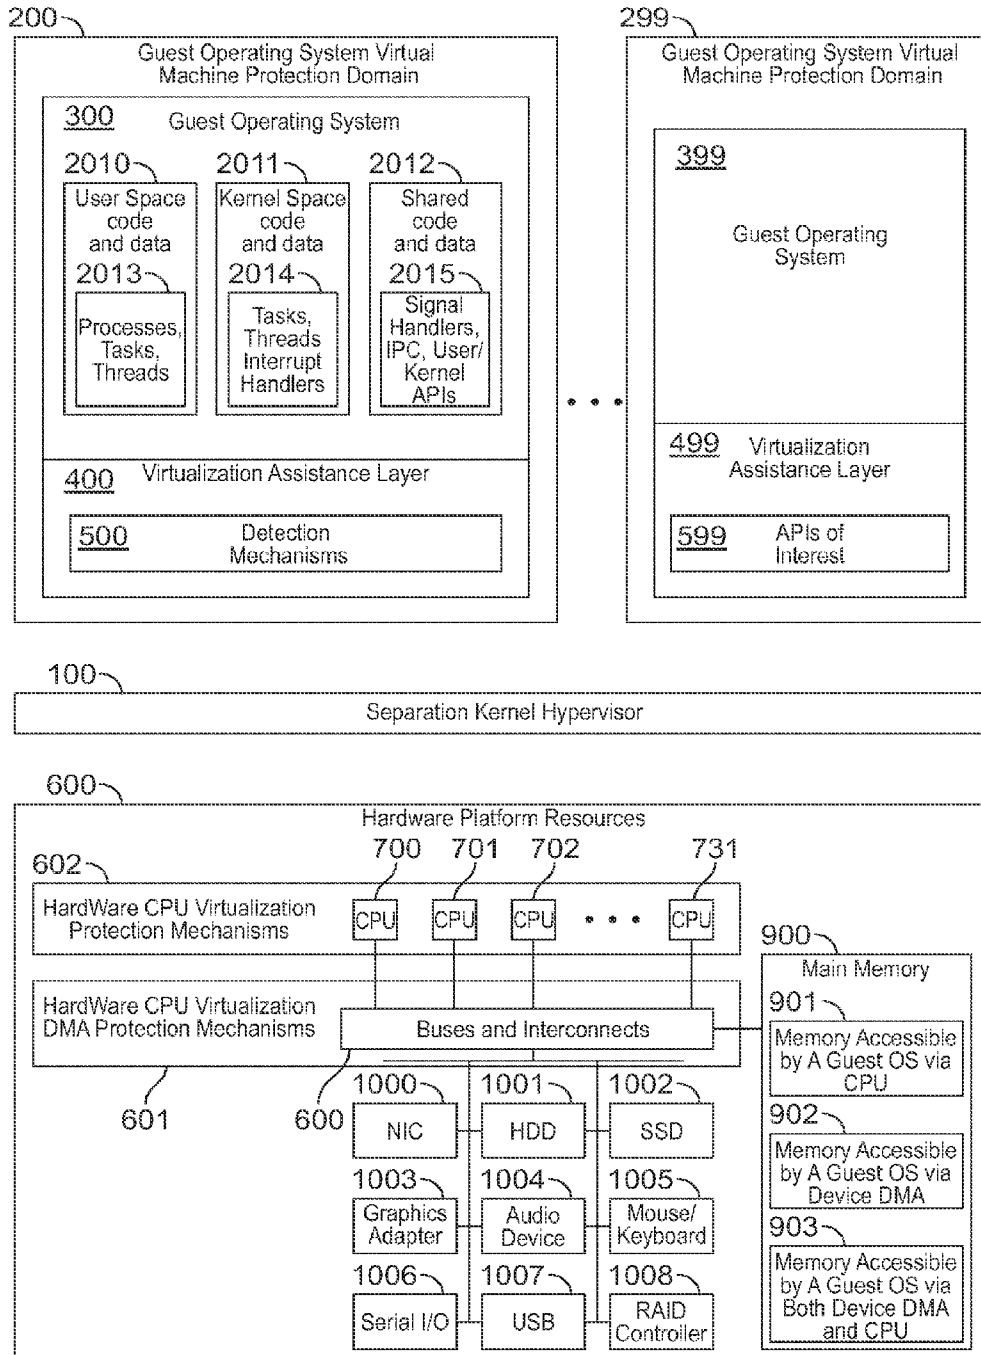


FIG. 4

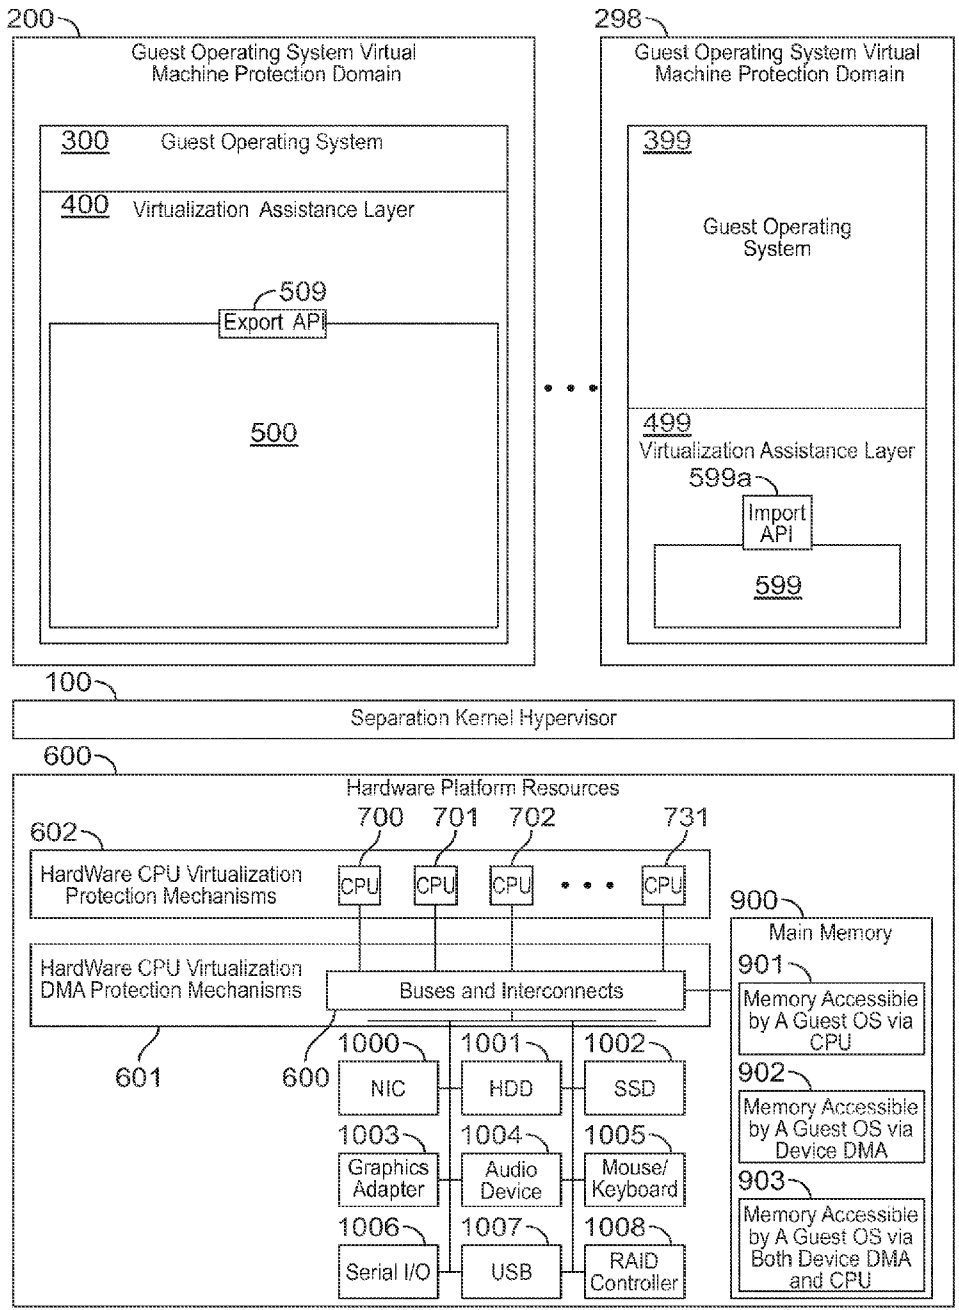


FIG. 5

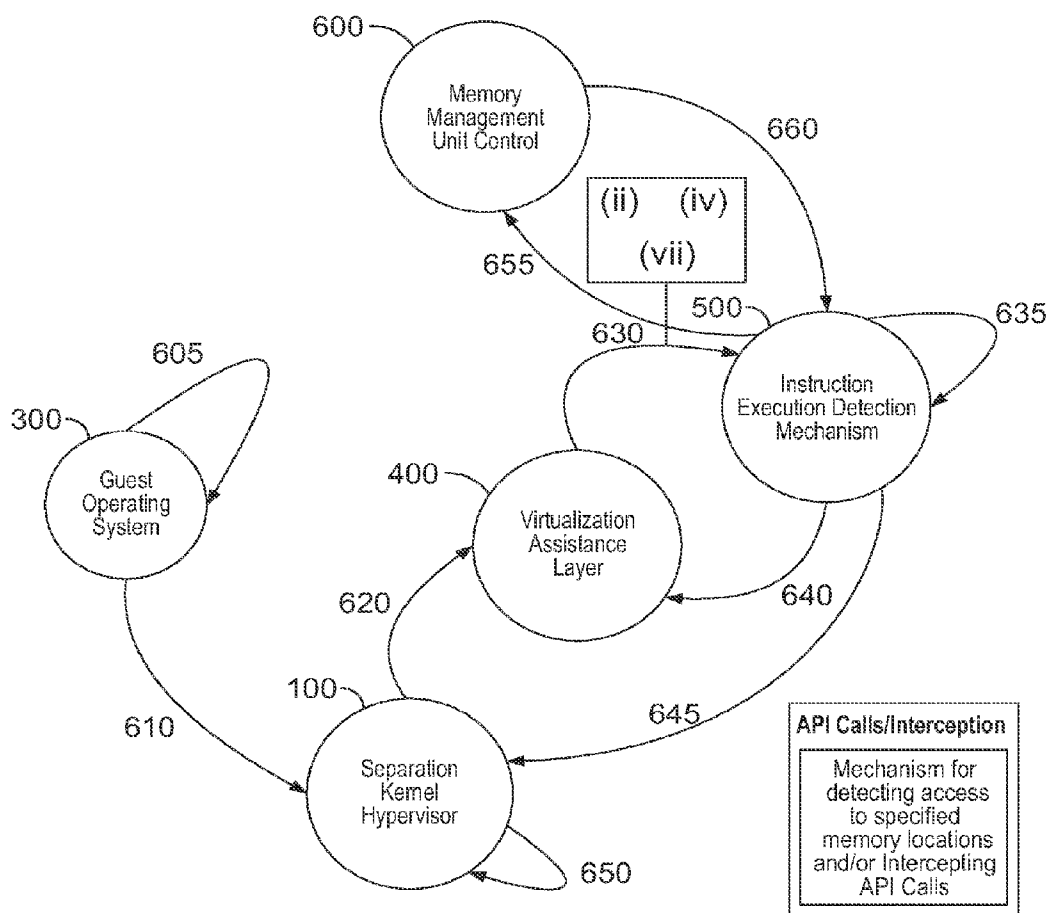


FIG. 6A

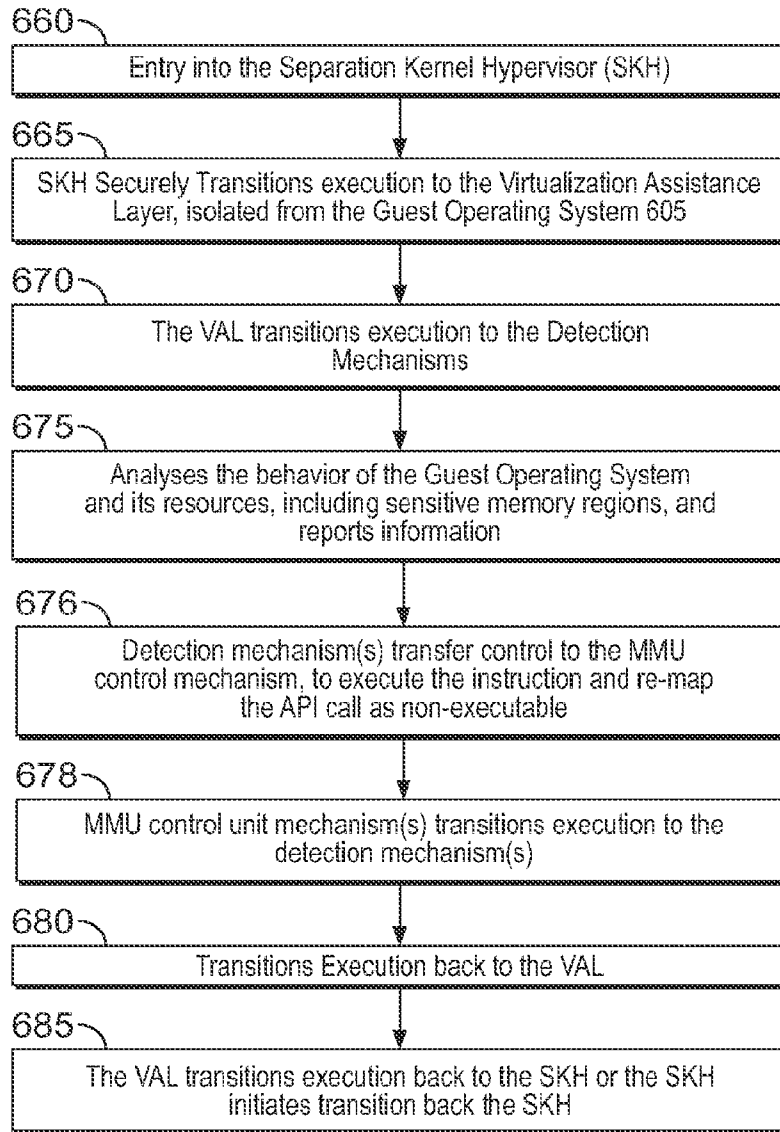


FIG. 6B

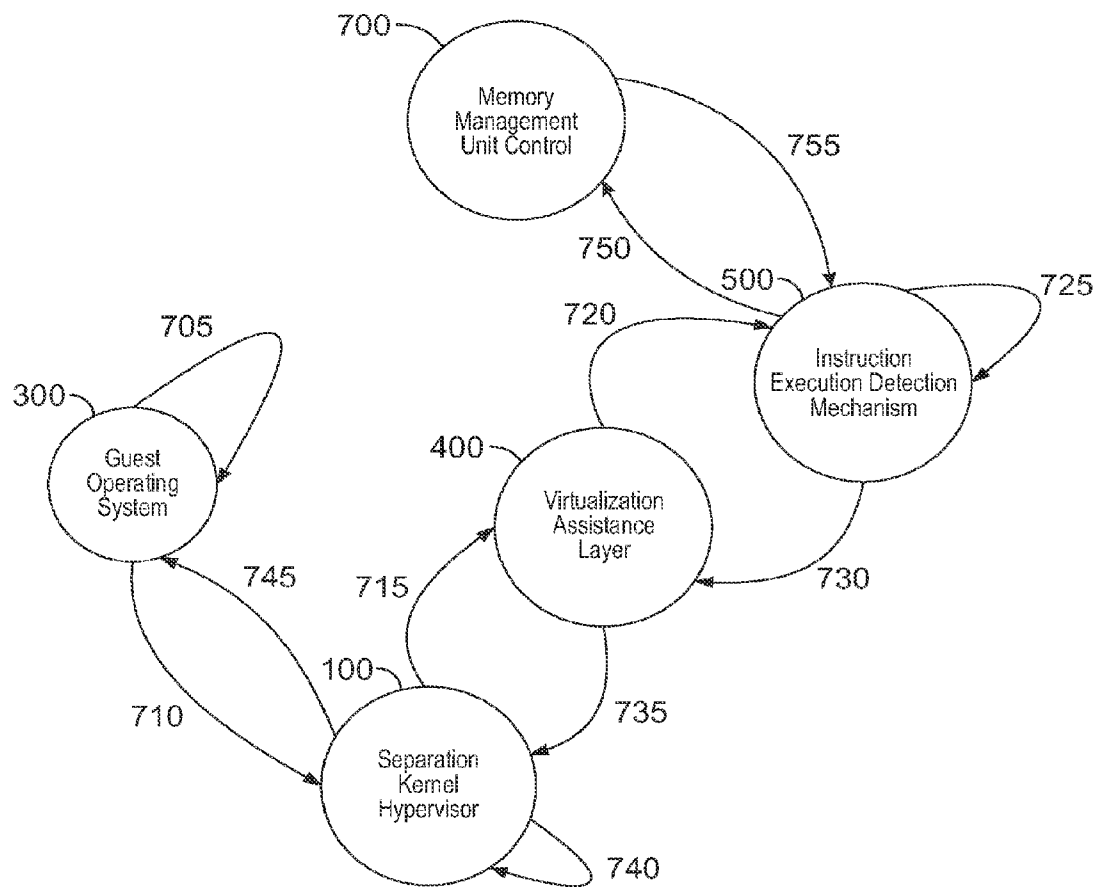


FIG. 7A

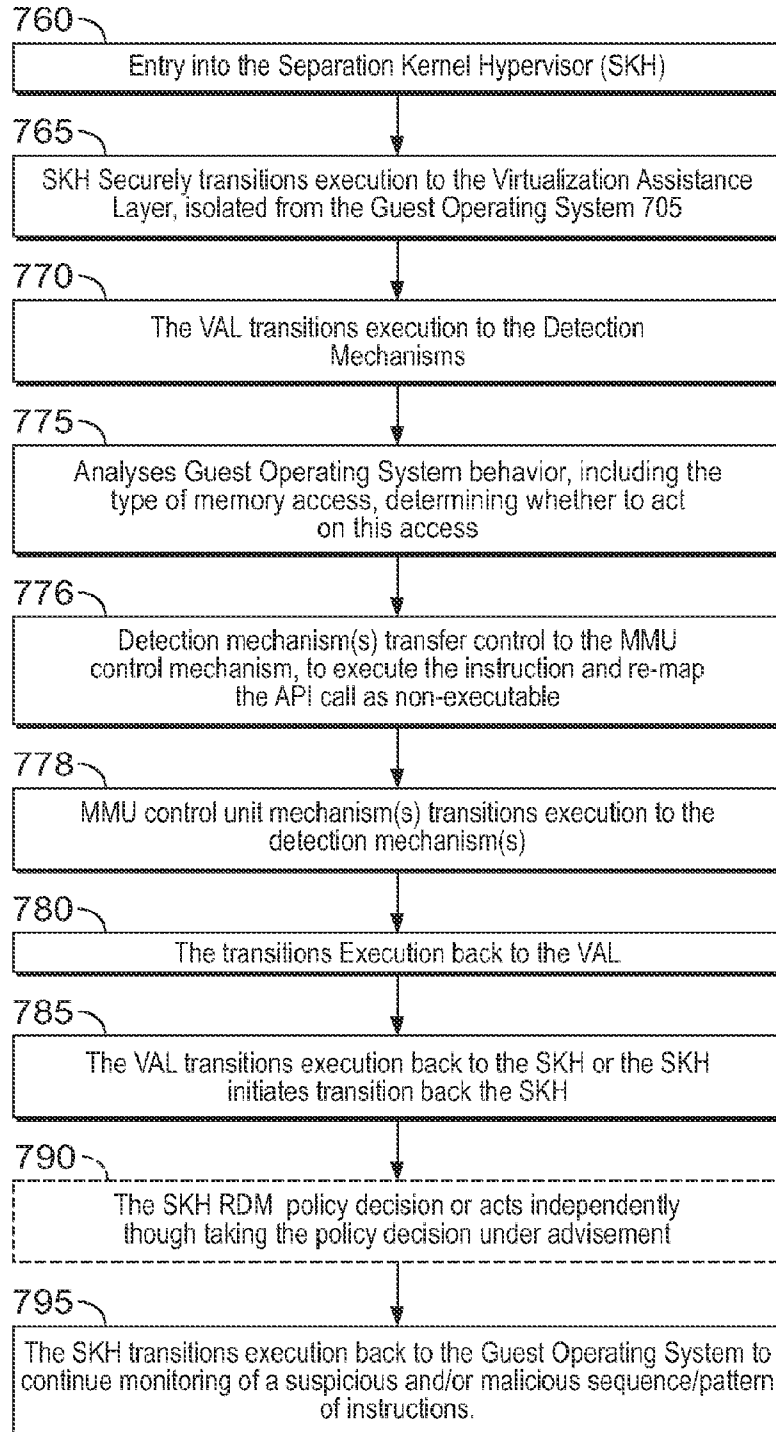


FIG. 7B

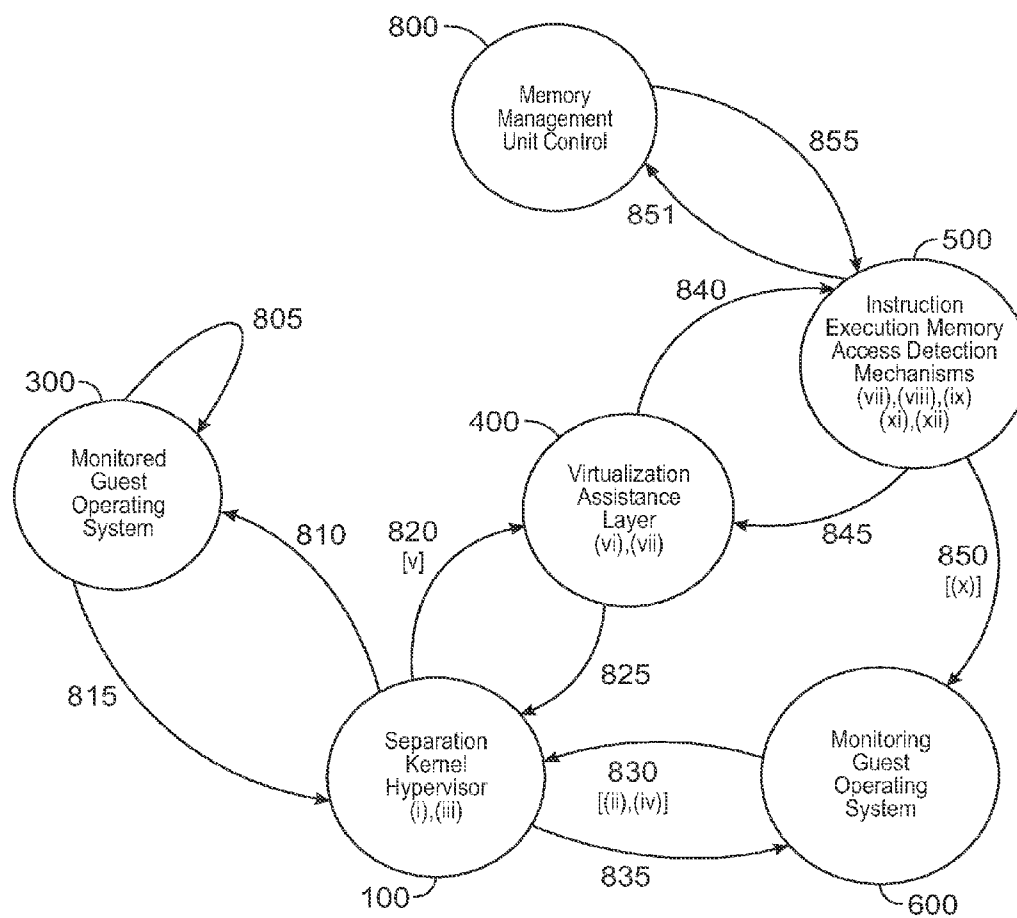


FIG. 8

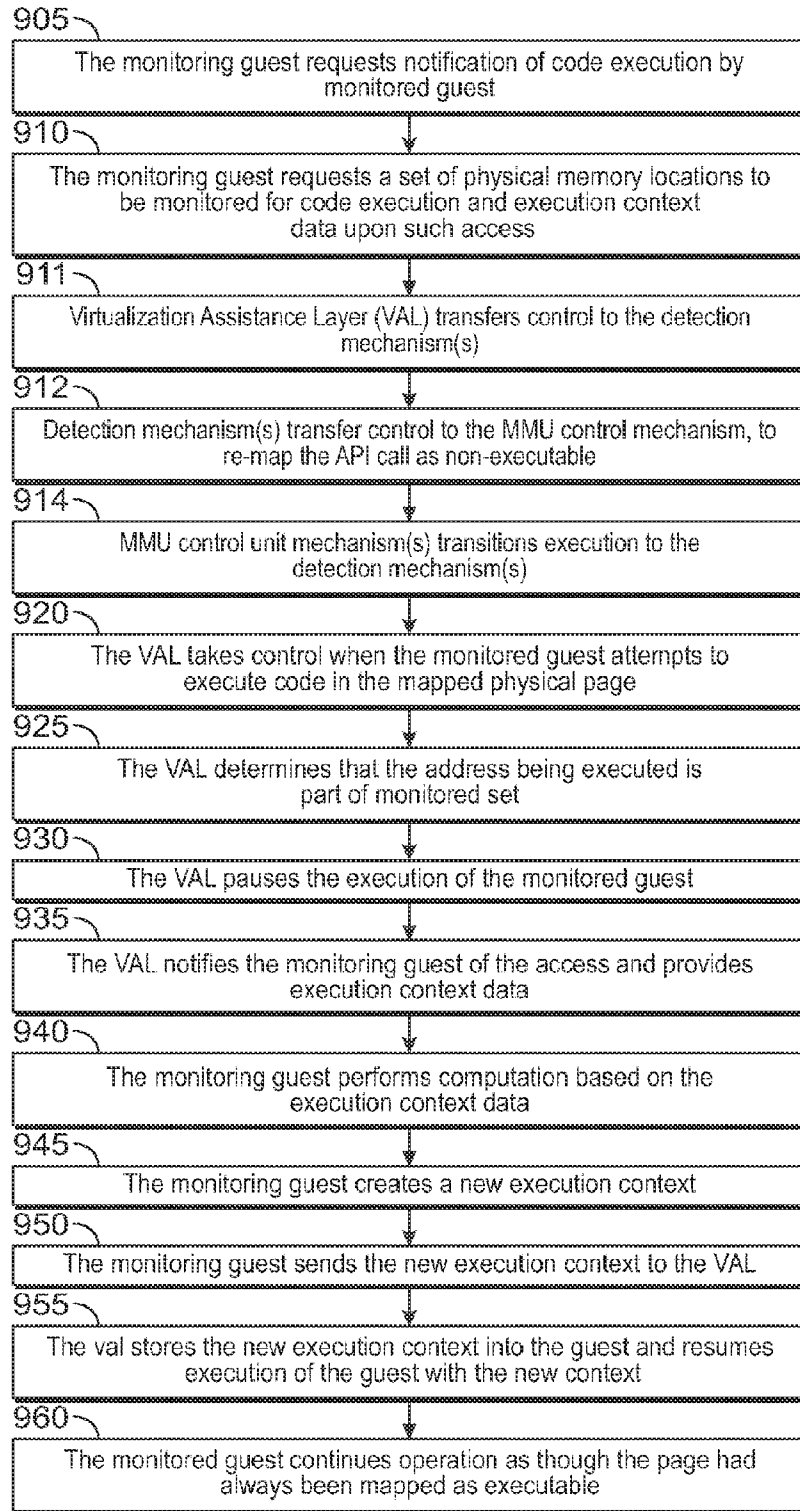


FIG. 9

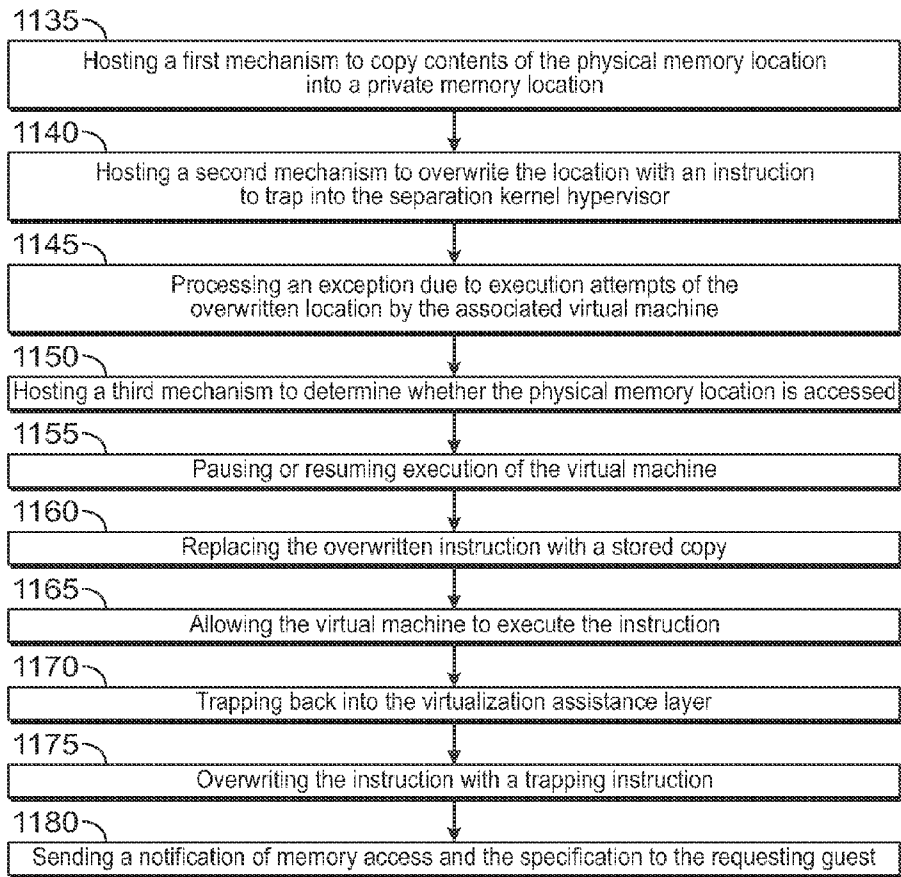


FIG. 11

**SYSTEMS AND METHODS INVOLVING
ASPECTS OF HARDWARE VIRTUALIZATION
SUCH AS HYPERVISOR, DETECTION AND
INTERCEPTION OF CODE OR
INSTRUCTION EXECUTION INCLUDING API
CALLS, AND/OR OTHER FEATURES**

CROSS-REFERENCE TO RELATED
APPLICATION INFORMATION

[0001] This is a continuation of application Ser. No. 14/714,241, filed May 15, 2015, published as US2015/0334126A1, now U.S. Pat. No. 9,203,855, which claims benefit/priority of U.S. provisional patent application No. 61/993,296, filed May 15, 2014, all of which are incorporated herein by reference in entirety.

BACKGROUND

[0002] 1. Field

[0003] Innovations herein pertain to computer software and hardware, computer virtualization, computer security and/or data isolation, and/or the use of a separation kernel hypervisor (and/or hypervisor), such as to detect and/or process information, including notification(s), interception and other processing regarding code/instruction execution by guest software, such as API calls, and which may include or involve guest operating system(s).

[0004] 2. Description of Related Information

[0005] In computer systems with hypervisors supporting a guest operating system, there exist some means to monitor the guest operating system for malicious or errant activity.

[0006] In a virtualized environment, running under control of a hypervisor, a suitably authorized guest may be allowed to monitor the activities of another guest. Among the reasons for such monitoring are debugging and security. However, previous approaches for monitoring other guests may include various drawbacks, such as allowing guests to poll the memory and other information within the monitored guest.

[0007] Due to the constantly evolving nature of malicious code, however, such systems face numerous limitations in their ability to detect and defeat malicious code. One major limitation is the inability of a hypervisor to defend itself against malicious code; e.g., the particular hypervisor may be subverted by malicious code and/or may allow malicious code in a guest operating system to proliferate between a plurality of guest operating systems in the system.

[0008] To solve that issue, the motivation and use of a Separation Kernel Hypervisor is introduced in environments with malicious code. The Separation Kernel Hypervisor, unlike a hypervisor, does not merely support a plurality of Virtual Machines (VMs), but supports more secure, more isolated mechanisms, including systems and mechanisms to monitor and defeat malicious code, where such mechanisms are isolated from the malicious code but are also have high temporal and spatial locality to the malicious code. For example, they are proximate to the malicious code, but incorruptible and unaffected by the malicious code.

[0009] Furthermore the Separation Kernel Hypervisor is designed and constructed from the ground-up, with security and isolation in mind, in order to provide security and certain isolation between a plurality of software entities (and their associated/assigned resources, e.g., devices, memory, etc.); by mechanisms which may include Guest Operating System Virtual Machine Protection Domains (secure entities estab-

lished and maintained by a Separation Kernel Hypervisor to provide isolation in time and space between such entities, and subsets therein, which may include guest operating systems, virtualization assistance layers, and detection mechanisms); where such software entities (and their associated assigned resources, e.g., devices, memory, etc., are themselves isolated and protected from each other by the Separation Kernel Hypervisor, and/or its use of hardware platform virtualization mechanisms.

[0010] Additionally, where some hypervisors may provide mechanisms to communicate between the hypervisor and antivirus software, or monitoring agent, executing within a guest operating system, the hypervisor is not able to prevent corruption of the monitoring agent where the agent is within the same guest operating system; or the guest operating system (or any subset thereof, possibly including the antivirus software, and/or monitoring agent) may be corrupted and/or subverted.

[0011] Finally, while some known systems and methods include implementations involving virtualized assistance layers and separation kernel hypervisors to handle various malicious code intrusions, such systems and method possess drawbacks with regard to handling and/or intercepting certain specified attacks, such as those related to API calls.

[0012] Overview of Some Aspects

[0013] Systems, methods, computer readable media and articles of manufacture consistent with innovations herein are directed to computer virtualization, computer security and/or data isolation, and/or the use of a separation kernel hypervisor (and/or hypervisor), such as to detect, process information, provide notification and/or interception features regarding code/instruction execution in specified physical memory location(s) by guest software and which may include or involve guest operating system(s). Information may further be obtained regarding the context of such code/instruction execution, the flow of execution within the guest may be controlled, and the context of the guest may be changed. Here, for example, certain implementations may include a suitably authorized guest running under control of a hypervisor and involving features of being immediately notified of another guest executing code at specified physical memory location(s). Upon access the monitoring guest may be provided with execution context information from the monitored guest. Further, the flow of execution within the guest may be controlled and/or the context of the guest may be changed.

[0014] According to some illustrative implementations, innovations herein may utilize and/or involve a separation kernel hypervisor which may include the use of a guest operating system virtual machine protection domain, a virtualization assistance layer, and/or an instruction (or code) execution detection/interception mechanism (which may be proximate in temporal and/or spatial locality to subject code, but isolated from it), inter alia, for detection, interception etc of code/instruction execution by guest software in specified memory locations. In some implementations, for example, a suitably authorized guest may obtain immediate notification if another guest it is monitoring executes code at specified physical memory location(s). Upon such access, the monitoring guest may be provided with execution context information from the monitored guest. Further, the monitored guest may be paused until the monitoring guest provides a new execution context to the monitored guest, whereupon the monitored guest resumes execution with the new context. Additionally, as indicated herein, the flow of execution within the guest may

be controlled and/or the context of the guest may be changed such that, e.g., API calls within the guest may be intercepted and simulated by the authorized guest.

[0015] It is to be understood that both the foregoing general description and the following detailed description are exemplary and explanatory only and are not restrictive of the inventions, as described. Further features and/or variations may be provided in addition to those set forth herein. For example, the present inventions may be directed to various combinations and subcombinations of the disclosed features and/or combinations and subcombinations of several further features disclosed below in the detailed description.

BRIEF DESCRIPTION OF THE DRAWINGS

[0016] The accompanying drawings, which constitute a part of this specification, illustrate various implementations and features of the present innovations and, together with the description, explain aspects of the inventions herein. In the drawings:

[0017] FIG. 1 is a block diagram illustrating an exemplary system and Separation Kernel Hypervisor architecture consistent with certain aspects related to the innovations herein.

[0018] FIG. 2A is a block diagram illustrating an exemplary system and Separation Kernel Hypervisor architecture consistent with certain aspects related to the innovations herein.

[0019] FIG. 2B is a block diagram illustrating an exemplary system and Separation Kernel Hypervisor architecture consistent with certain aspects related to the innovations herein.

[0020] FIG. 2C is a block diagram illustrating an exemplary system and Separation Kernel Hypervisor architecture consistent with certain aspects related to the innovations herein.

[0021] FIG. 2D is a block diagram illustrating an exemplary system and Separation Kernel Hypervisor architecture consistent with certain aspects related to the innovations herein.

[0022] FIG. 3 is a block diagram illustrating an exemplary system and separation kernel Hypervisor architecture consistent with certain aspects related to the innovations herein.

[0023] FIG. 4 is a block diagram illustrating an exemplary system and separation kernel hypervisor architecture consistent with certain aspects related to the innovations herein.

[0024] FIG. 5 is a block diagram illustrating an exemplary system and separation kernel hypervisor architecture consistent with certain aspects related to the innovations herein.

[0025] FIGS. 6A-6B are representative sequence/flow diagrams illustrating exemplary systems, methods and separation kernel hypervisor architecture consistent with certain aspects related to the innovations herein.

[0026] FIGS. 7A-7B are representative sequence/flow diagrams illustrating exemplary systems, methods and separation kernel hypervisor architecture consistent with certain aspects related to the innovations herein.

[0027] FIG. 8 is a representative sequence diagram illustrating exemplary systems, methods, and Separation Kernel Hypervisor architecture consistent with certain aspects related to the innovations herein.

[0028] FIG. 9 is a representative flow diagram illustrating exemplary methodology and Separation Kernel Hypervisor processing consistent with certain aspects related to the innovations herein.

[0029] FIG. 10 is an exemplary state diagram illustrating aspects of memory management unit processing in conjunc-

tion with the hypervisor and VAL, consistent with certain aspects related to the innovations herein.

[0030] FIG. 11 is a representative flow diagram illustrating exemplary methodology and separation kernel hypervisor processing concerning exception-related instructions consistent with certain aspects related to the innovations herein.

DETAILED DESCRIPTION OF ILLUSTRATIVE IMPLEMENTATIONS

[0031] Reference will now be made in detail to the inventions herein, examples of which are illustrated in the accompanying drawings. The implementations set forth in the following description do not represent all implementations consistent with the inventions herein. Instead, they are merely some examples consistent with certain aspects related to the present innovations. Wherever possible, the same reference numbers are used throughout the drawings to refer to the same or like parts.

[0032] To solve one or more of the drawbacks mentioned above and/or other issues, implementations herein may relate to various detection, monitoring, notification(s), interception and/or prevention techniques, systems, and mechanisms, as may be used with a separation kernel hypervisor. Among other things, such systems and methods may include and/or involve the use of the monitoring of the entirety, or suitably configured subset thereof of guest operating system resources including virtualized resources, and/or “physical” or “pass-through” resources. Examples include monitoring of the virtual CPUs, its memory access attempts to execute code involving specified memory such as monitoring and/or intercepting API calls within the guest.

[0033] With regard to certain implementations, in order to perform such advanced monitoring in a manner that maintains suitable performance characteristics in a system that may include a separation kernel hypervisor and a guest operating system, mechanisms such as a separation kernel hypervisor, a guest operating system virtual machine protection domain, virtual machine assistance layer, and/or instruction execution detection/interception mechanisms, may be used to monitor a monitored guest on a corresponding guest operating system.

[0034] Systems and methods are disclosed for detecting and/or notifying executed code by guest software and which may include or involve guest operating system(s). According to some implementations, for example, a suitably authorized guest running under control of a hypervisor may request that it be notified of another guest executing code at a specified physical memory location. Features of real-time notification of, and action(s) regarding obtaining an execution context are provided to the monitoring guest upon access by the monitored guest to executed code at specific physical memory locations. Here, monitoring may also be performed in a timely and expeditious fashion, including by virtue of the monitoring context being proximate (in time and space) to the monitored context. Additionally, isolation may be maintained between the monitor and monitored context. Further, such monitoring may be performed by mechanisms providing a wide and comprehensive set of monitoring techniques and resources under monitoring, inter alia, so as to monitor against threats which are multi-lateral and/or multi-dimensional in nature.

[0035] According to some implementations, for example, a hypervisor is configured to allow a guest (monitoring guest) to request notifications of code execution by another guest

(monitored guest). The monitoring guest requests that a set of physical memory locations be monitored for code execution, and the execution context data be returned on such access. The virtualization assistance layer (VAL) in the monitored guest maps (e.g., remaps, unmaps) those physical APIs containing those locations as non-executable. This is distinct from the monitored guest's notion of API mappings. When software in the monitored guest attempts to execute code involving an API call, for example, control transitions to the VAL. The VAL determines that the address being executed is part of the set to be monitored. The VAL notifies the monitoring guest of the access and provides the monitoring guest with the execution context data as configured for that access. Various innovative features involving mapping (unmapping, re-mapping) and insertion of exception-causing instruction(s) may be utilized. As such, the monitored guest is allowed to continue operation as though the API has always been mapped executable.

[0036] According to some implementations, for example, a separation kernel hypervisor (SKH) ensures the isolation of multiple guest Operating Systems each in its own virtual machine (VM). The SKH may implement a mechanism whereby a suitably authorized Monitoring Guest sends a list of memory locations to be monitored for another guest. Furthermore, each of the physical memory locations may be associated with a specification for the execution context information to be obtained upon access to the memory location(s). The SKH may send to the other guests the specification for the execution context information associated with the list of memory locations. A Virtualization Assistance Layer of software runs within the same protection domain as the guest Virtual Machine but is not directly accessible by the guest. A Virtualization Assistance Layer implements a virtual motherboard containing a virtual CPU and memory. The VAL and mechanism may process exceptions caused by non-executable API execution attempts by its associated guest virtual machine. The VAL may determine whether the memory address accessed is one of those specified in the list of physical memory locations sent to another guest. The VAL may send a notification of the memory access and associated context information to the requesting guest.

[0037] Systems and methods are disclosed for providing secure information monitoring. According to some implementations, for example, such information monitoring may be provided from a context not able to be bypassed, tampered with or by the context under monitoring. Here, monitoring may also be performed in a timely and expeditious fashion, including by virtue of the monitoring context being proximate (in time and space) to the monitored context. Additionally, isolation may be maintained between the monitor and monitored context. Further, such monitoring may be performed by mechanisms providing a wide and comprehensive set of monitoring techniques and resources under monitoring, inter alia, so as to monitor against threats which are multi-lateral and/or multi-dimensional in nature.

[0038] In one exemplary implementation, there is provided a method of secure domain isolation, whereby an execution context within a virtual machine may monitor another execution context within that virtual machine or another virtual machine, in a manner maintaining security and isolation between such contexts. Innovations herein also relate to provision of these contexts such that neither/none can necessarily corrupt, affect, and/or detect the other.

[0039] Moreover, systems and methods herein may include and/or involve a virtual machine which is augmented to form a more secure virtual representation of the native hardware platform for a particular execution context. And such implementations may also include a virtual representation which is augmented with a wide and deep variety of built-in detection, notification(s), monitoring and/or interception mechanisms, wherein secure isolation between the domains or virtual machines is maintained.

[0040] In general, aspects of the present innovations may include, relate to, and/or involve one or more of the following aspects, features and/or functionality. Systems and methods herein may include or involve a separation kernel hypervisor. According to some implementations, a software entity in hypervisor context that partitions the native hardware platform resources, in time and space, in an isolated and secure fashion may be utilized. Here, for example, embodiments may be configured for partitioning/isolation as between a plurality of guest operating system virtual machine protection domains (e.g., entities in a hypervisor guest context).

[0041] The separation kernel hypervisor may host a plurality of guest operating system virtual machine protection domains and may host a plurality of mechanisms including instruction execution detection/interception mechanisms which may execute within such guest operating system virtual machine protection domains. The instruction execution detection/interception mechanisms may execute in an environment where guest operating systems cannot tamper with, bypass, or corrupt the instruction execution detection/interception mechanisms. The instruction execution detection/interception mechanisms may also execute to increase temporal and spatial locality of the guest operating system's resources. Further, in some implementations, the instruction execution detection/interception mechanisms may execute in a manner that is not interfered with, nor able to be interfered with, nor corrupted by other guest operating system virtual machine protection domains including their corresponding guest operating systems. The instruction execution detection/interception mechanisms include, but are not limited to, performing one or more of the following actions on guest operating systems related to guest code execution at specified memory location(s), such as access to API calls including sensitive memory regions, and/or actions in response thereto such as performing various interception processing.

[0042] Where monitoring may include, but is not limited to, actions pertaining to observation, detection, mitigation, prevention, tracking, modification, reporting upon, of memory access within and/or by a guest operating system and/or by entities configured to perform such monitoring for purposes which may be used to ascertain, and assist in ascertaining, when suspect code, and/or code under general monitoring or instrumented execution/debugging, unit test, regression test, or similar scrutiny, is or may be operating at specified memory location(s); or, therein, hiding and/or concealed, halted, stalled, infinitely looping, making no progress beyond its intended execution, stored and/or present (either operating or not), once-active (e.g., extinct/not present, but having performed suspect and/or malicious action), and otherwise having been or being in a position to adversely and/or maliciously affect the hypervisor guest, or resource under control of the hypervisor guest.

[0043] The term "map" or "mapped" shall broadly mean: setting a memory page with any of the following properties

applied to it (as set and enforced by the hardware MMU via the SKH): mapped (present), executable, readable, writeable.

[0044] The term “unmap” or “unmapped” shall broadly mean: setting a memory page with any of the following properties applied to it (as set and enforced by the hardware MMU via the SKH): unmapped (non-present), non-executable, non-readable, non-writeable.

[0045] FIG. 1 is a block diagram illustrating an exemplary system and separation kernel hypervisor architecture consistent with certain aspects related to the innovations herein. FIG. 1 also shows a separation kernel hypervisor executing on native hardware platform resources, e.g., where the separation kernel hypervisor may support the execution, isolated and partitioned in time and space, between a plurality of guest operating system protection domains. Here, a guest operating system domain may be an entity that is established and maintained by the separation kernel hypervisor in order to provide a secure and isolated execution environment for software. Referring to FIG. 1, a separation kernel hypervisor 100 is shown executing on top of the native hardware platform resources 600. Further, the separation kernel hypervisor 100 supports the execution of a guest operating system virtual machine protection domain 200.

[0046] The separation kernel hypervisor 100 may also support the execution of a plurality of guest operating system virtual machine protection domains, e.g., 200 to 299 in FIG. 1. In some implementations, the separation kernel hypervisor may provide time and space partitioning in a secure and isolated manner for a plurality of guest operating system virtual machine protection domains, e.g., 200 to 299 in FIG. 1. Such features may include rigid guarantees on scheduling resources, execution time, latency requirements, and/or resource access quotas for such domains.

[0047] According to some implementations, in terms of the sequence of establishment, after the native hardware platform resources 600 boot the system, execution is transitioned to the separation kernel hypervisor 100. The separation kernel hypervisor 100 then creates and executes a guest operating system virtual machine protection domain 200, or a plurality of guest operating system virtual machine protection domains, e.g., 200 to 299 in FIG. 1. Some implementations of doing so consonant with the innovations herein are set forth in PCT Application No. PCT/2012/042330, filed 13 Jun. 2012, published as WO2012/177464A1, and U.S. patent application Ser. No. 13/576,155, filed Dec. 12, 2013, published as US2014/0208442 A1, which are incorporated herein by reference in entirety.

[0048] Consistent with aspects of the present implementations, it is within a guest operating system virtual machine protection domain that a guest operating system may execute. Further, it is within a guest operating system virtual machine protection domain that instruction execution detection/interception mechanisms may also execute, e.g., in a fashion isolated from any guest operating system which may also execute within that same guest operating system virtual machine protection domain, or in other guest operating system virtual machine protection domains.

[0049] FIG. 2A is a block diagram illustrating an exemplary system and separation kernel hypervisor architecture consistent with certain aspects related to the innovations herein. FIG. 2A also shows a separation kernel hypervisor executing on native hardware platform resources (where the native platform resources may include a plurality of CPUs, buses and interconnects, main memory, Network Interface

Cards (NIC), Hard Disk Drives (HDD), Solid State Drives (SSD), Graphics Adaptors, Audio Devices, Mouse/Keyboard/Pointing Devices, Serial I/O, USB, and/or Raid Controllers, etc.), where the separation kernel hypervisor may support the execution, isolated and/or partitioning in time and space, between a plurality of guest operating system protection domains. Here, some implementations may involve a guest operating system protection domains which may contain a guest operating system, and/or a virtualization assistance layer (which itself may contain instruction execution detection/interception mechanisms).

[0050] FIG. 2A shows both a guest operating system 300, and a virtualization assistance layer 400 executing within the same guest operating system virtual machine protection domain 200. In some implementations, the virtualization assistance layer 400 may provide the execution environment for the instruction execution detection/interception mechanism(s) 500. Further, the virtualization assistance layer 400 may assist the separation kernel hypervisor in virtualizing portions of the platform resources exported to a given guest operating system (e.g., Virtual CPU/ABI, Virtual chipset ABI, set of virtual devices, set of physical devices, and/or firmware, etc., assigned to a given guest operating system 300 and/or guest virtual machine protection domain 200). Some systems and methods herein utilizing such virtualization assistance layer may include or involve (but are not strictly limited to) a self-assisted virtualization component, e.g., with an illustrative implementation shown in FIG. 2D.

[0051] The guest operating system 300 and the virtualization assistance layer 400, which may include instruction execution detection/interception mechanism(s) 500, are isolated from each other by the separation kernel hypervisor 100. In implementations herein, the guest operating system 300 cannot tamper with, bypass, or corrupt the virtualization assistance layer 400, nor can it tamper with, bypass or corrupt the instruction execution detection/interception mechanisms 500. Since the instruction execution detection/interception mechanisms 500 are isolated from the guest operating system 300, the instruction execution detection/interception mechanisms 500 are able to act on a portion of (or the entirety, depending on policy and configuration) of the guest operating system 300 and its assigned resources in a manner that is (a) is transparent to the guest operating system 300 and (b) not able to be tampered with by the guest operating system 300 or its assigned resources (e.g., errant and/or malicious device DMA originated by devices assigned to the guest operating system 300), and (c) not able to be bypassed by the guest operating system 300. For example, the instruction execution detection/interception mechanisms 500, within the given virtualization assistance layer 400, may read and/or modify portions of the guest operating system 300 and resources to which the Guest Operating System 300 has been granted access (by the Separation Kernel Hypervisor 100), while none of the Guest Operating System 300 nor the resources to which has access may modify any portion of the instruction execution detection/interception mechanisms 500 and/or virtualization assistance layer 400.

[0052] By having a given virtualization assistance layer 400 and a given Guest Operating System 300 within the within the same Guest Virtual Machine Protection Domain 200, isolated from each other by the Separation Kernel Hypervisor 100, various benefits, non-penalties, or mitiga-

tion of penalties, such as the following, may be conferred to the system at large and to the instruction execution detection/interception mechanisms 500.

[0053] Increased Spatial and Temporal Locality of Data

[0054] By being contained within the same Guest Virtual Machine Protection Domain 300, the virtualization assistance layer 200, and/or corresponding private (local) instruction execution detection/interception mechanisms 500 existing in that same Guest Virtual Machine Protection Domain 300, have greater access, such as in time and space, to the resources of the Guest Operating System 300 than would entities in other guest virtual machine protection domains or other Guest Operating Systems; e.g., the subject guest virtual machine protection domain has faster responsiveness and/or has lower latency than if processed in another guest virtual machine protection domain. Though such resources are still accessed in a manner that is ultimately constrained by the separation kernel hypervisor 100, there is less indirection and time/latency consumed in accessing the resources:

[0055] In one illustrative case, the instruction execution detection/interception mechanisms 500 private (local) to a given Guest virtualization assistance layer 200 and its associated Guest Operating System 300 can react faster to code execution physical memory access issues, and not need to wait on actions from another entity in another guest virtual machine protection domain 200 or guest operating system 300 (which may themselves have high latency, be corrupted, unavailable, poorly scheduled, or subject to a lack of determinism and/or resource constraint, or improper policy configuration, etc.).

[0056] Here, for example, if a Guest Operating System 300 was to monitor a Guest Operating System 399 located within another Guest Virtual Machine Protection Domain 107, it would encounter penalties in time and space for accessing that Guest Operating System and its resources; furthermore, there is increased code, data, scheduling, and/or security policy complexity to establish and maintain such a more complex system; such increases in complexity and resources allow for more bugs in the implementation, configuration, and/or security policy establishment and maintenance.

[0057] Scalability and Parallelism

[0058] Each Guest Operating System 300 may have a virtualization assistance layer 200, and instruction execution detection/interception mechanisms 500, that are private (local) to the Guest Virtual Machine Protection Domain 200 that contains both that Guest Operating System 300, the virtualization assistance layer 400, and the instruction execution detection/interception mechanisms.

[0059] Fault Isolation, Low Level of Privilege, Defense in Depth, Locality of Security Policy, and Constraint of Resource Access

[0060] Here, for example, relative to the extremely high level of privilege of the separation kernel hypervisor 100, the virtualization assistance layer 400, the instruction execution detection/interception mechanism 500, and the Guest Operating System 300 within the same Guest Virtual Machine Protection Domain 200 are only able to act on portions of that Guest Virtual Machine Protection Domain 200 (subject to the Separation Kernel Hypervisor 100) and not portions of other Guest Virtual Machine Protection Domains (nor their contained or assigned resources).

[0061] Subject to the isolation guarantees provided by the Separation Kernel Hypervisor 100, the virtualization assistance layer 400 accesses only the resources of the Guest

Operating System 300 within the same Guest Virtual Machine Protection Domain 200 and that virtualization assistance layer 400 is not able to access the resources of other Guest Operating Systems.

[0062] As such, if there is corruption (bugs, programmatic errors, malicious code, code and/or data corruption, or other faults, etc.) within a given Guest Virtual Machine Protection Domain 200 they are isolated to that Guest Virtual Machine Protection Domain 200. They do not affect other Guest Virtual Machine Protection Domains 299 nor do they affect the separation kernel hypervisor 100. This allows the separation kernel hypervisor to act upon (e.g., instantiate, maintain, monitor, create/destroy, suspend, restart, refresh, backup/restore, patch/fix, import/export etc.) a plurality of Guest Virtual Machine Protection Domains 200 and their corresponding virtualization assistance layer 400 and instruction execution detection/interception mechanisms 500 (or even Guest Operating Systems 300) without corruption of the most privileged execution context of the system, the separation kernel hypervisor 100.

[0063] Similarly, the faults that may occur within a virtualization assistance layer 400 or the instruction execution detection/interception mechanisms 500 (e.g., by corruption of software during delivery) are contained to the Guest Virtual Machine Protection Domain 200 and do not corrupt any other Guest Virtual Machine Protection Domain; nor do they corrupt the Separation Kernel Hypervisor 100.

[0064] Furthermore, the faults within a Guest Operating System 300 are contained to that Guest Operating System 300, and do not corrupt either the virtualization assistance layer 400 or the instruction execution detection/interception mechanisms 500.

[0065] FIG. 2B is a block diagram illustrating an exemplary system and separation kernel hypervisor architecture consistent with certain aspects related to the innovations herein. FIG. 2B illustrates a variation of FIG. 2A where a minimal runtime environment 398 executes in place of a (larger/more complex) guest operating system. Here, a minimal runtime environment may be an environment such as a VDS (virtual device server), and/or a LSA (LynxSecure application), etc. The minimal runtime environment 398 can be used for policy enforcement related to activities reported by a virtualization assistance layer and/or instruction execution detection/interception mechanisms; such an environment is also monitored by a virtualization assistance layer and/or instruction execution detection/interception mechanisms private to the guest operating system virtual machine protection domain containing the minimal runtime environment.

[0066] FIG. 2C is a block diagram illustrating an exemplary system and separation kernel hypervisor architecture consistent with certain aspects related to the innovations herein. FIG. 2C illustrates a variation of FIG. 2A and FIG. 2B where a minimal runtime environment executes in place of a (larger/more complex) guest operating system but without a virtualization assistance layer or instruction execution detection/interception mechanisms.

[0067] FIG. 2D is a block diagram illustrating an exemplary system and Separation Kernel Hypervisor architecture consistent with certain aspects related to the innovations herein. FIG. 2D illustrates a variation of FIG. 2 where a self-assisted virtualization (SAV) mechanism is used to implement the virtualization assistance layer.

[0068] FIG. 3 is a block diagram illustrating an exemplary system and separation kernel Hypervisor architecture consis-

tent with certain aspects related to the innovations herein. FIG. 3 also shows certain detailed aspects with respect to FIGS. 2A/B, where the guest operating system may attempt to process API calls at specified memory locations that may include a plurality of code and/or data which may constitute execution contexts which may include the following types of software including any/all of which malicious code may attempt to corrupt or utilize: malicious code, anti-virus software, corrupted anti-virus software, integrity checkers, corrupted integrity checkers, rootkits, return oriented rootkits, etc. The inventions herein are not limited to memory access attempts to malicious code and is discussed below as illustrative examples.

[0069] For example, in FIG. 3, if antivirus software 2001 executes within a given guest operating system 300, and such anti-virus software 2001 is itself corrupted, and itself executes malicious code 2002 or fails to prevent the execution of malicious code 2002, the corruption is constrained to the given guest operating system 300, and the corruption may be acted upon (e.g., detected, notified, prevented, mitigated, reported, tracked, modified/patched, suspended, halted, restarted, eradicated, etc.) by the instruction execution detection/interception mechanisms 500 that monitors/acts on code execution in specified memory location(s) such as API calls, and is provided within the same guest virtual machine protection domain 200 as the guest operating system 300.

[0070] With regard to other exemplary implementations, as may be appreciated in connection with FIG. 3, if an integrity checker 2003 (e.g., a “security” component or driver within a guest operating system 300) executes within a given guest operating system 300, and such integrity checker 2003 is itself corrupted into a corrupted integrity checker 2004 (and executes malicious code, or fails to prevent the execution of malicious code), the corruption is constrained to the given guest operating system 300, and the corruption may be acted upon (e.g., detected, notified, prevented, mitigated, reported, tracked, modified/patched, suspended, halted, restarted, eradicated, etc.) by the instruction execution detection/interception mechanisms 500 that monitors/acts on code executed at the specified memory location(s), and is provided within the same guest virtual machine protection domain 200 as the guest operating system 300.

[0071] With regard to another illustration, again with reference to FIG. 3, if a rootkit 2006 executes within the guest operating system 300 (e.g., by having fooled the Integrity Checker 2003 by the nature of the root kit being a return oriented rootkit 2007, which are designed specifically to defeat integrity checkers) the corruption is constrained to the given guest operating system 300, and the corruption may be acted upon (e.g., detected, notified, prevented, mitigated, reported, tracked, modified/patched, suspended, halted, restarted, eradicated, etc.) by the instruction execution detection/interception mechanisms 500 that monitors/acts on code execution in specified memory location(s), and is provided within the same guest virtual machine protection domain 200 as the guest operating system 300.

[0072] In another example, again with respect to FIG. 3, if a polymorphic virus 2005 (an entity designed to defeat integrity checkers, among other things) executes within the guest operating system 300 (e.g., by having fooled the integrity checker 2003, or by having the a corrupted integrity checker 2003) the corruption is constrained to the given guest operating system 300, and the corruption may be acted upon (e.g., detected, notified, prevented, mitigated, reported, tracked,

modified/patched, suspended, halted, restarted, eradicated, etc.) by the instruction execution detection/interception mechanisms 500 that monitors/acts on code execution in specified memory location(s), and is provided within the same guest virtual machine protection domain 200 as the guest operating system 300.

[0073] In general, referring to FIG. 3, if a malicious code 2000 executes within the guest operating system 300 (e.g., by means including, but not limited strictly to bugs, defects, bad patches, code and/or data corruption, failed integrity checkers, poor security policy, root kits, viruses, trojans, polymorphic viruses, and/or other attack vectors and/or sources of instability within the guest operating system 300 etc.), the corruption is constrained to the given guest operating system 300, and the corruption may be acted upon (e.g., detected, notified, prevented, mitigated, reported, tracked, modified/patched, suspended, halted, restarted, eradicated, etc.) by the instruction execution detection/interception mechanisms 500 that monitors/acts on code execution in specified memory location(s), and is provided within the same guest virtual machine protection domain 200 as the guest operating system 300.

[0074] Furthermore, in the examples above and other cases, such corruption of the guest operating system 300, and the resources to which it has access, do not corrupt the instruction execution detection/interception mechanisms 500, the virtualization assistance layer 400, the guest virtual machine protection domain 200, or plurality of other such resources in the system (e.g., other guest virtual machine protection domains 299), or the separation kernel hypervisor 100.

[0075] In some implementations, the instruction execution detection/interception mechanisms 500, in conjunction with the virtualization assistance layer 400, and the separation kernel hypervisor 100, may utilize various methods and mechanisms such as the following, given by way of illustration and example but not limitation, to act with and upon its associated guest operating system 300 the resources assigned to the guest operating system 300, and the systems behavior generated thereto and/or thereby.

[0076] FIG. 4 is a block diagram illustrating an exemplary system and separation kernel hypervisor architecture consistent with certain aspects related to the innovations herein. For example, FIG. 4 illustrates resources that may be assigned to a Guest Operating System 300 consistent with certain aspects related to the innovations herein.

[0077] FIG. 4 shows an illustrative extension of either FIG. 2, and/or FIG. 3, where the guest operating system may have a plurality of code and/or data which may constitute execution contexts which may include the following types of software mechanisms and/or constructs user space code and data that may be associated with an unprivileged mode of CPU code execution (as used herein ‘user space’ being an execution environment of low privilege, versus an execution environment of high privilege, such as kernel space), which may contain processes, tasks, and/or threads, etc.; kernel space code and data, that may be associated with a privileged mode of CPU execution, which may contain tasks, threads, interrupt handlers, drivers, etc.; shared code and data, that may be associated with either privileged and/or unprivileged modes of CPU execution, and which may include signal handlers, Inter Process Communication Mechanisms (IPC), and/or user/kernel mode APIs. It also may include main memory that may be accessed by the CPU, by DMA from devices, or both. It also shows protection mechanisms including hardware

CPU virtualization protection mechanisms, and hardware virtualization DMA protection mechanisms. Instruction execution detection/interception mechanisms **500**, **599** such as API call interception/simulation mechanisms may reside within corresponding Virtualization Assistance Layers **400**, **499**

[0078] Such resources, explained here by way of example, not limitation, may include a subset of (a) hardware platform resources **600**, virtualized hardware platform resources (hardware platform resources **600** subject to further constraint by the separation kernel hypervisor **100**, the hardware CPU virtualization protection mechanisms **602**, and/or the hardware virtualization DMA protection mechanisms **601**), and execution time on a CPU **700** (or a plurality of CPUs, e.g., **700** to **731**) (scheduling time provided by the separation kernel hypervisor **100**), and space (memory **900** provided by the separation kernel hypervisor) within which the guest operating system **300** may instantiate and utilize constructs of the particular guest operating system **300**, such as a privileged (“kernel” space) modes of execution, non-privileged (“user” space) modes of execution, code and data for each such mode of execution (e.g., processes, tasks, threads, interrupt handlers, drivers, signal handlers, inter process communication mechanisms, shared memory, shared APIs between such entities/contexts/modes, etc.

[0079] FIG. **5** is a block diagram illustrating an exemplary system and separation kernel hypervisor architecture consistent with certain aspects related to the innovations herein. FIG. **5** shows an illustrative implementation as may be associated with FIG. **2**, FIG. **3**, and/or FIG. **4**, where the instruction execution detection/interception mechanisms, that may be within the virtualization assistance layer, may include the following monitoring systems and mechanisms: memory monitor, an instruction monitor, etc. FIG. **5** also illustrates import/export mechanism that may be used by a virtualization assistance layer and/or instruction execution detection/interception mechanisms to communicate between themselves and other virtualization assistance layer and/or instruction execution detection/interception mechanisms in other guest operating system virtual machine protection domains (subject to the security policies established, maintained, and enforced by the separation kernel hypervisor), in an isolated, secure, and even monitored fashion.

[0080] FIG. **5** illustrates mechanism and resources that may be used by the instruction execution detection/interception mechanisms **500** to monitor a guest operating system **300**. Such mechanisms and resources may include a memory monitor **501**, and an instruction monitor **502**.

[0081] The virtualization assistance layer **400** and/or the instruction execution detection/interception mechanisms **500** may also use an export API **509** and/or an import API **599** (as may be configured and governed by the separation kernel hypervisor **100**), in order to provide secure communication between a plurality of virtualization assistance layers (e.g., virtualization assistance layers **400** to **499**) and/or a plurality of instruction execution detection/interception mechanisms (e.g., instruction execution detection/interception mechanisms **500** to **599**).

[0082] Innovations set forth herein, as also described in additional detail elsewhere herein via notation to the reference numerals in the description below, reside around various combinations, subcombinations and/or interrelated functionality of the following features or aspects: (i) a separation kernel hypervisor that ensures the isolation of multiple guest Operating Systems each in its own Virtual Machine (VM); (ii)

a separation kernel hypervisor as in (i) that implements a mechanism whereby a suitably authorized guest can send a list of physical memory locations to be watched to another guest; (iii) a separation kernel hypervisor as in (i) that implements a mechanism whereby each of the physical memory locations in (ii) is associated with a specification for what execution context information is to be obtained on access to that location; (iv) a separation kernel hypervisor as in (i) that implements a mechanism whereby the specifications associated with the list of memory locations in (ii) can be sent to the other guest as in (ii); (v) a separation kernel hypervisor as in (i) that implements a mechanism whereby the execution context specified in (iii) can be sent to the other guest as in (ii); (vi) a virtualization assistance layer (VAL) of software that runs within the same protection domain as the guest virtual machine but is not directly accessible by the guest; (vii) a virtualization assistance layer as in (vi) that implements a virtual motherboard containing a virtual CPU and memory; (viii) a VAL as in (vi) that implements a mechanism to map physical memory pages as non-executable; (ix) a VAL as in (vi) that processes exceptions caused by non-executable page execution attempts by its associated guest virtual machine; (x) a VAL as in (vi) that implements a mechanism to determine whether the address accessed is one of those specified in (ii); (xi) a VAL as in (vi) that can send a notification of the memory access and associated context information as in (iii) to the requesting guest; (xii) a VAL as in (vi) that implements a mechanism receive context information as in (iii) from the requesting guest; (xiii) a VAL as in (vi) that can replace the context information in its associated virtual machine; (xiv) a VAL as in (vi) that can pause the execution of its virtual machine; and/or (xv) a VAL as in (vi) that can resume the execution of its virtual machine.

[0083] Systems and mechanisms, and example embodiments, of the instruction execution detection/interception mechanisms **500** may include:

[0084] 1. Monitoring of CPU (and CPU cache based) guest OS memory access (originated from a plurality of resources available to the guest operating system **300** (in FIGS. **3** and **4**), including CPUs and/or caches assigned and/or associated with such), as directed by execution and resources (shown in FIG. **3**) within the guest OS **300**. For memory assigned to the guest OS **300**, such as a subset of the main memory **900** (in FIGS. **2**, **3**, **4**, and **5**) the separation kernel hypervisor **100** may trap access to that memory, and then pass associated data of that trap to the virtualization assistance layer **400**. The virtualization assistance layer **400** may then pass the associated data of that trap to the instruction execution detection/interception mechanisms **500**.

[0085] The virtualization assistance layer **400**, instruction execution detection/interception mechanisms **500**, and/or the separation kernel hypervisor **100** may use feedback mechanisms between themselves to recognize and monitor patterns of guest operating system **300** memory access; not strictly one-off memory access attempts.

[0086] The monitoring of guest operating system **300** memory access includes, but is not limited to, constructs in guest operating system **300** memory (including the resources in the guest operating system **300** in FIGS. **3** and **4**) which may have semantics specific to a particular guest operating system **300** or a set of applications hosted by the guest operating system **300** (possibly including antivirus software).

[0087] The virtualization assistance layer **400**, instruction execution detection/interception mechanisms **500**, and/or the

Separation Kernel Hypervisor **100** may use feedback mechanisms between themselves to recognize and monitor patterns of Guest Operating System **300** DMA access to memory; not strictly one-off access attempts. Illustrative aspects, here, are shown in FIGS. **6A-6B**.

[0088] 2. Monitoring of specific Guest Operating System **300** instruction execution attempts, and/or specific instruction sequence execution attempts.

[0089] For all such attempts by the Guest Operating System **300**, the Separation Kernel Hypervisor **100** (when configured to do so, or via feedback receive from the virtualization assistance layer **400** and/or the instruction execution detection/interception mechanisms **500**) may trap such access attempts, then pass associated data of that trap to the virtualization assistance layer **400** and/or instruction execution detection/interception mechanisms **500**.

[0090] The virtualization assistance layer **400** and/or the instruction execution detection/interception mechanisms **500** can respond to such instruction sequences; including, but not limited to, recognition of a significant fraction of a given sequence, then prevent/block the final instructions of the malicious sequence from execution.

[0091] Illustrative aspects, here, are shown in FIGS. **7A-7B**.

[0092] FIGS. **6A-6B** are representative sequence and flow diagrams illustrating exemplary systems, methods and separation kernel hypervisor architecture consistent with certain aspects related to the innovations herein. FIGS. **6A-6B** relate, inter alia, to behavior relating to the handling of guest operating system attempts to access main memory.

[0093] Turning to the illustrative implementations/aspects of FIG. **6A**, at step **605** a Guest Operating System receives a command for memory access to a specified memory location. Then, at step **610**, the Guest Operating System attempts to execute code in the memory location(s). The memory usage attempt triggers entry into the separation kernel hypervisor. Then, at step **620**, the separation kernel hypervisor securely transitions execution to the virtualization assistance layer; in a manner isolated from the Guest Operating System. Next, in step **630** the virtualization assistance layer transitions execution to the instruction execution detection/interception mechanisms. Step **630** may include and/or involve aspects (ii), (iv) and (viii), including aspect (ii) where the separation kernel hypervisor implements a mechanism whereby a suitably authorized guest can send a list of memory locations to be watched to another guest. A virtualization assistance layer (VAL) of software that runs within the same protection domain as the guest Virtual Machine but is not directly accessible by the guest (aspect vi). The VAL that processes unmapped memory exceptions taken by its associated guest virtual machine (aspect viii). Then, at step **635** the instruction execution detection/interception mechanisms analyze the behavior of the Guest Operating System and its resources and makes a policy decision; in this example, it has been configured to understand the memory locations which are sensitive (e.g., involve API calls), thus decides to disallow, pause or continue the code execution. The instruction execution detection/interception mechanism detects access or processing related to specified memory locations, for example. Then, at step **655**, the instruction execution detection mechanism **500** transfers control to a memory management unit (MMU) control mechanism **600**. This mechanism **600** performs the memory management unit control operations need to execute the instruction(s) and map the appropriate page as non-executable.

Additional details of the MMU functionality, here, are set forth further below in connection with FIG. **10**. Then, at step **660**, the MMU control mechanisms transition execution to the instruction execution detection mechanism. Next, at step **640** the instruction execution detection/interception mechanisms transition execution to the virtualization assistance layer, passing it the policy decision. Then, at step **645** the virtualization assistance layer transitions execution back to the Separation Kernel Hypervisor, or the Separation Kernel Hypervisor transitions execution from the virtualization assistance layer back to the Separation Kernel Hypervisor. Next, at step **650** the Separation Kernel Hypervisor acts on the policy decision generated by the instruction execution detection/interception mechanisms (in one example, it intercepts API calls for simulation by the authorized guest), or the Separation Kernel Hypervisor acts independently of the policy decision, but in a manner that takes the policy decision under advisement (depending on configuration). The SKH may receive, analyze, and/or act upon policy decisions from multiple sources, which may include multiple detection/notification mechanisms; including cases where multiple mechanisms monitor a given Guest OS.

[0094] As explained above in connection with FIG. **6A**, the Guest Operating System accesses a specified memory location. The memory access may be monitored and identified as including API calls requiring interception via the instruction execution detection/interception mechanism to generate a policy decision. The memory access attempt triggers entry into the Separation Kernel Hypervisor.

[0095] Turning to FIG. **6B**, such system or process may initiate upon entry into the SKH, at **660**. Then, at **665**, the Separation Kernel Hypervisor securely transitions execution to the Visualization Assistance Layer; in a manner isolated from the Guest Operating System. Next, at **670**, the Visualization Assistance Layer transitions execution to the instruction execution detection/interception mechanisms. The instruction execution detection/interception mechanisms may then analyze, at **675**, the behavior of the Guest Operating System and its resources and makes a policy decision; for example, it may be configured to understand the memory locations which are sensitive (e.g. handle the API calls for interception), thus decides to deny, pause or continue the memory processing/access attempt. At **676**, the detection mechanism(s) may transfer control to a memory management unit (MMU) control mechanism, to execute the instruction(s) and re-map the appropriate page as non-executable. Additional details of the MMU functionality, here, are set forth further below in connection with FIG. **10**. Then, at **678**, the MMU control mechanism(s) may transition execution to the detection mechanism. Once the policy decision(s) have been made, the instruction execution detection/interception mechanisms transition execution to the virtualization assistance layer, at **680**, passing it the policy decision. Then, at **685**, the virtualization assistance layer transitions execution back to the separation kernel hypervisor, or the separation kernel hypervisor transitions execution from the virtualization assistance layer back to the separation kernel hypervisor. Finally, at **690**, the separation kernel hypervisor acts on the policy decision generated by the instruction execution detection/interception mechanisms (in this example, it may intercept/simulate the API call, and/or deny processing with respect to the API call, although it may also allow or pause the memory access), or the Separation Kernel Hypervisor acts independently of the policy decision, but in a manner that

takes the policy decision under advisement (depending on configuration). Further, the SKH may receive, analyze, and/or act upon policy decisions from multiple sources, which may include multiple mechanisms; inducing cases where multiple mechanisms monitor a given Guest OS.

[0096] FIGS. 7A-7B are representative sequence/flow diagrams illustrating exemplary systems, methods and Separation Kernel Hypervisor architecture consistent with certain aspects related to the innovations herein. FIGS. 7A-7B relate, inter alia, to behavior relating to an attempt to access specified API calls such as by the handling of guest operating system instruction sequences (e.g., execution attempts of a repeated pattern/series of MOV, RET, or MOV IRET instruction on an Intel IA32e architecture; such patterns of which may constitute code of “return oriented” attacks/rootkits). Here, in such illustrative cases, memory access within the guest operating system will attempt to corrupt and/or subvert antivirus software and/or software integrity checkers within the guest operating system via a “return oriented” attack (attacks constructed to evade integrity checkers); and the instruction execution detection/interception mechanisms detects/prevents the attack.

[0097] Turning to the illustrative implementations/aspects of FIG. 7A, at step 705, a Guest Operating System receives a command for memory access to a specified memory location. Then at step 710 an attempt to process an API call such as a specific sequence and/or pattern of CPU instructions is performed, that either triggers transition into the SKH for (2a) every instruction in the sequence and/or pattern (a single stepping behavior), or (2b) for a number of instructions of size greater than one of the sequence and/or pattern (multiple stepping). The (2a) or (2b) behavior is based on system configuration. Next, at step 715 the Separation Kernel Hypervisor securely transitions execution to the virtualization assistance layer; in a manner isolated from the Guest Operating System. Then, at step 720 the virtualization assistance layer transitions execution to the instruction execution detection/interception mechanisms. Next, at step 725 the instruction execution detection/interception mechanisms analyzes the behavior of the Guest Operating System and its resources and makes a policy decision. Then, at step 750, the instruction execution detection/interception mechanism 500 transfers control to a memory management unit (MMU) control mechanism 700. This mechanism 700 performs the memory management unit control operations need to execute the instruction(s) and map the appropriate page as non-executable. Additional details of the MMU functionality, here, are set forth further below in connection with FIG. 10. Then, at step 755, the MMU control mechanisms transition execution to the instruction execution detection mechanism. Then, the instruction execution detection/interception mechanisms transition execution to the virtualization assistance layer, at 730, passing it the policy decision. Next, at 735 the virtualization assistance layer transitions execution back to the Separation Kernel Hypervisor, or the Separation Kernel Hypervisor transitions execution from the virtualization assistance layer back to the Separation Kernel Hypervisor. Then, in step 740 the Separation Kernel Hypervisor acts on the policy decision generated by the instruction execution detection/interception mechanisms (in this example it suspends the Guest OS, preventing the Guest OS from accessing the memory and executing the “Return Oriented” attack; a type of attack that thwarts code integrity checkers in the Guest OS), or the Separation Kernel Hypervisor acts independently

of the policy decision, but in a manner that takes the policy decision under advisement (depending on configuration). The SKH may receive, analyze, and/or act upon policy decisions from multiple sources, which may include multiple mechanisms; including cases where multiple mechanisms monitor a given Guest OS. Finally, in step 745, in order to continue to recognize sequences and/or patterns of instructions, execution may cycle a multiple times between steps 705 through 740.

[0098] As explained above in connection with FIG. 7A, the guest operating system attempts specific memory access relating to an API call. Here, for example, the API call may involve a specified memory location. The attempt triggers entry into the Separation Kernel Hypervisor.

[0099] Turning to FIG. 7B, such illustrative system or process may initiate upon entry into the SKH, at 760. Then, at 765, the Separation Kernel Hypervisor securely transitions execution to the Visualization Assistance Layer; in a manner isolated from the Guest Operating System. Next, at 770, the Visualization Assistance Layer transitions execution to the instruction execution detection/interception mechanisms. The instruction execution detection/interception mechanisms may then analyze, at 775, the behavior of the Guest Operating System and its resources and makes a policy decision; in this example it recognizes the Guest Operating System instruction sequence and/or pattern as an attempt to process an API call, and the policy decision is to made to deny further (and/or future) execution of the sequence and/or pattern, preventing the Guest Operating System from providing the API call to the monitored guest. At 776, the detection mechanism(s) may transfer control to a memory management unit (MMU) control mechanism, to execute the instruction(s) and map the appropriate page as non-executable. Additional details of the MMU functionality, here, are set forth further below in connection with FIG. 10. Then, at 778, the MMU control mechanism(s) may transition execution back to the detection mechanism. Once the policy decision(s) have been made, the instruction execution detection/interception mechanisms transition execution to the virtualization assistance layer, at 780, passing it the policy decision. Then, at 785, the virtualization assistance layer transitions execution back to the Separation Kernel Hypervisor, or the separation kernel hypervisor transitions execution from the virtualization assistance layer back to the separation kernel hypervisor. Optionally, at step 790, the separation kernel hypervisor acts on the policy decision generated by the instruction execution detection/interception mechanisms (in this example it denies processing of the API call), or the Separation Kernel Hypervisor acts independently of the policy decision, but in a manner that takes the policy decision under advisement (depending on configuration). Further, the SKH may receive, analyze, and/or act upon policy decisions from multiple sources, which may include multiple mechanisms; inducing cases where multiple mechanisms monitor a given Guest OS. In a final step 795, in order to recognize sequences and/or patterns of instructions (and/or further monitor an existing monitored sequence and/or pattern of instructions), execution may cycle a multiple times between steps 760 through 790.

[0100] FIGS. 8 and 9 are representative sequence/flow diagrams illustrating exemplary systems, methods, and separation kernel hypervisor architecture consistent with certain aspects related to the innovations herein. FIGS. 8 and 9 relate, inter alia, to the guest operating system attempting to access a specified memory location where the detection mechanisms

monitors, detects, and notifies the access and determines an action in response to the detected access.

[0101] Turning to the illustrative implementations/aspects of FIG. 8, at step **805**, a Monitored Guest Operating System **300** attempts to access a memory location. Then, at step **815**, the request to access the memory location is sent to the SKH. The separation kernel hypervisor **100** that ensures the isolation of multiple guest Operating Systems each in its own virtual machine (VM) (step i). Another Monitored Guest Operating System **600** allows a suitably authorized Monitoring Guest **600** to send **830** a list of physical memory locations to be monitored for another guest **300** (step ii). Furthermore, each of the physical memory locations may be associated with a specification for the execution context information to be obtained upon access to the memory location(s) (step iii). A response from the SKH **100** is provided to the Monitored Guest Operating System **600** at step **600**. Next, at step **820** the separation kernel hypervisor securely transitions execution to the virtualization assistance layer **400** in a manner isolated from the Guest Operating System. Step iv sends the other guests the specification associated with the list of memory locations. Step v includes the SKH **100** that implements a mechanism where the execution context information specified is sent to the other guests. A virtualization assistance layer **400** of software runs within the same protection domain as the guest virtual machine but is not directly accessible by the guest (step vi). Step (vii) includes a virtualization assistance layer **400** that implements a virtual motherboard containing a virtual CPU and memory. Step (viii) of the VAL **400** implements a mechanism to map physical memory pages as non-executable. Step (ix) of the VAL **400** processes exceptions caused by non-executable page execution attempts by its associated guest virtual machine. Step (x) of the VAL **400** determines whether the memory address accessed is one of those specified in feature (ii) in the list of physical memory locations sent to another guest. Step (xi) of the VAL **400** sends a notification of the memory access and associated context information as in step (iii) to the requesting guest. Step (xii) of the VAL **400** receives context information as in (iii) from the requesting guest. Step (xii) of the VAL **400** replaces the context information in its associated virtual machine. Step (xiv) of the VAL **400** pauses the execution of its virtual machine. Step (xv) of the VAL **400** resumes the execution of its virtual machine

[0102] Then, at step **840**, the virtualization assistance layer transitions execution to the instruction execution detection/interception mechanisms **500**. Next, the instruction execution detection/interception mechanisms analyze the behavior of the guest operating system and its resources and makes a policy decision. Here, for example, at **851**, the instruction execution detection mechanism **500** may transfer control to a memory management unit control mechanism **800**. This mechanism **800** may perform memory management unit (MMU) control operations needed to execute the instruction (s) and map the appropriate page as non-executable. Additional details of the MMU functionality, here, are set forth further below in connection with FIG. 10. After this, at step **855**, the MMU control mechanisms transition execution to the instruction execution detection mechanism. Then, the instruction execution detection/interception mechanisms transition execution to the virtualization assistance layer, at **845**, passing to it the policy decision. Next, at **825** the virtualization assistance layer transitions execution back to the separation kernel hypervisor, or the separation kernel hyper-

visor transitions execution from the virtualization assistance layer back to the Separation Kernel Hypervisor. At step **810**, the SKH **100** transitions execution to the monitored guest operating system **300** based on the policy decision. At step **825**, the separation kernel hypervisor acts on the policy decision generated by the instruction execution detection/interception mechanisms, or the separation kernel hypervisor acts independently of the policy decision, but in a manner that takes the policy decision under advisement (depending on configuration). The SKH may receive, analyze, and/or act upon policy decisions from multiple sources, which may include multiple code execution mechanisms; including cases where multiple mechanisms monitor a given Guest OS. Then, in order to continue to recognize sequences and/or patterns of memory access, execution may cycle a multiple times between steps **805** through **850**.

[0103] As explained above in connection with FIG. 8, the guest operating system executes code at a specified memory location, obtains and replaces execution context information from the monitored guest. The attempt triggers entry into the separation kernel hypervisor for monitoring, detection, notification and/or interception.

[0104] Turning to FIG. 9, such illustrative system or process begins at step **905** where a hypervisor is configured to allow a guest (monitoring guest) to request notifications of code execution by another guest (monitored guest). The monitoring guest requests that a set of physical memory locations be monitored for code execution, and the execution context data be returned on such access, at step **910**. At **912**, the detection mechanism(s) may transfer control to a memory management unit (MMU) control mechanism, to execute the instruction and re-map the appropriate page or call as non-executable. Additional details of the MMU functionality, here, are set forth further below in connection with FIG. 10. Then, at **914**, the MMU control mechanism(s) may transition execution to the detection mechanism. The VAL in the monitored guest maps the physical pages containing those locations as non-executable, at step **915**. This is distinct from the monitored guest's notion of page mappings. At step **920**, when software in the monitored guest attempts to execute code in such a page, control transitions to the VAL. The VAL determines that the address being executed is part of the set to be monitored, at step **925**. The VAL pauses the execution of the monitored guest, at step **930**. The VAL notifies the monitoring guest of the access and provides the monitoring guest with the execution context data as configured for that access, at step **935**. The monitoring guest then performs computation based on the execution context data, at step **940**. The monitoring guest creates a new execution context, at step **945**, which may be the same as the original execution context. This new execution context is sent to the VAL, at step **950**. The VAL stores the new execution context into the guest and resumes execution of the guest with the new context at step **955**. The monitored guest is allowed to continue operation at step **960** as though the page has always been mapped executable.

[0105] FIG. 10 is an exemplary state diagram illustrating aspects of memory management unit processing in conjunction with the hypervisor and VAL, consistent with certain aspects related to the innovations herein. In FIG. 10, control is passed to the Memory Management Unit (MMU) Control **1019** via any of the following control paths **1005** including step **655** (from FIG. 6A), step **755** (from FIG. 7A), and step **851** (from FIG. 8). Step **1015** transitions control from the

Memory Management Control Unit **1010** to the detection mechanisms **1020** to make a policy decision regarding the page of memory the GuestOS had attempted to access. The detection mechanisms **1020** execute a policy decision to either deny or allow the GuestOS to access the memory. In step **1025**, the detection mechanisms **1020** execute the decision to allow the GuestOS access to the memory.

[**0106**] The detection mechanisms may transition execution to the VAL **1035** with a request that the page of memory the GuestOS had attempted to access be remapped (mapped as accessible) to the GuestOS at step **1030**.

[**0107**] The VAL may then transition execution to the SKH with a request that the page of memory the GuestOS had attempted to access be remapped (mapped as accessible) to the GuestOS at step **1040**. The SKH executes a policy decision at step **1045** to allow or deny the request that the page of memory the GuestOS had attempted to access be remapped (mapped as accessible) to the GuestOS. In an exemplary embodiment, the SKH allows the request to map the memory page as accessible to the GuestOS.

[**0108**] The SKH may transition execution back to the VAL at step **1050** with a message that the memory page that the GuestOS had attempted to access has been remapped (mapped as accessible) to the GuestOS. The VAL transitions execution back to the detection mechanisms **1020** at step **1055** with a message that the memory page that the GuestOS had attempted to access has been remapped (mapped as accessible) to the GuestOS.

[**0109**] At step **1060**, the detection mechanisms **1020** execute a policy decision to either allow or deny the GuestOS to complete the execution of the command/instruction that the GuestOS had attempted which had triggered the GuestOS access attempt to the memory page.

[**0110**] At step **1065**, the detection mechanisms **1020** determine to allow the GuestOS to complete execution of the command/instruction that the GuestOS had attempted which had triggered the GuestOS access attempt to the memory page. The detection mechanisms then transition execution to the VAL **1035**.

[**0111**] At step **1070**, the VAL **1035** then transitions execution to the SKH with a request to allow the GuestOS to complete execution command/instruction that the GuestOS had attempted which had triggered the GuestOS access attempt to the memory page. The SKH executes a policy decision at step **1072** to allow or deny the GuestOS to complete execution of the command/instruction that the GuestOS had attempted which had triggered the GuestOS access attempt to the memory page. In this example, the SKH allows the GuestOS to complete the execution of that command/instruction. At step **1074**, the SKH securely transition execution to the GuestOS. At step **1076**, the GuestOS completes execution of the command/instruction that the GuestOS had attempted which triggered the GuestOS access attempt to the memory page. At step **1078**, the protection mechanisms provided by the SKH trigger a transition back to the SKH immediately after completion of the GuestOS command/instruction.

[**0112**] At step **1080**, the SKH transitions execution back to the VAL **1035**, with a message that the GuestOS has completed execution of the command/instruction that the GuestOS had attempted which had triggered the GuestOS access attempt to the memory page. At step **1082**, the VAL **1035** transitions execution to the detection mechanisms **1020** with a message that the GuestOS has completed execution of the

command/instruction that the GuestOS had attempted which had triggered the GuestOS access attempt to the memory page. At step **1084**, the detection mechanisms **1020** determine whether to map the memory page as nonexecutable again. At step **1086**, the detection mechanisms **1020** make a transition back to the VAL via any of the control paths including step **600** (from FIG. 6A), step **750** (from FIG. 7A), and step **855** (from FIG. 8).

[**0113**] FIG. 11 is a representative flow diagram illustrating exemplary methodology and separation kernel hypervisor processing concerning exception-related instructions consistent with certain aspects related to the innovations herein. Innovative processing, here, may occur in the context of hardware platform resources partitioned via a separation kernel hypervisor into a plurality of guest operating system virtual machine protection domains. With regard to list(s) of memory locations of an authorized guest provides to another guest, each physical memory location is associated with a respective specification of execution context information upon access to the each of the plurality of physical memory locations. A message of the specification may be transmitted to the requesting guest, and execution context information may be provided to such other guest. As set forth elsewhere herein, a virtualization assistance layer (VAL) may be provided including a virtual representation of the hardware platform in each of the guest operating system virtual machine protection domains such that the VAL is not directly accessible by the authorized guest.

[**0114**] Referring to FIG. 11, the virtualization assistance layer and detection mechanism(s) may process exceptions caused by non-executable API execution attempts made by guest virtual machine(s) associated with the VAL. In processing and/or inserting such exception-causing instructions, various mechanisms may be employed. For example, a first mechanism may be hosted, at **1135**, to copy contents of the physical memory location into a private memory location. A second mechanism may then be hosted, at **1140**, to overwrite the location with an instruction to trap into the separation kernel hypervisor. At **1145**, exceptions may then processed as a function of non-executable API execution attempts directed to the overwritten location by the associated virtual machine. A third mechanism may also be hosted, at **1150**, to determine whether the physical memory location is accessed. If access, processing or other specified conditions are detected, execution of the virtual machine may be paused or resumed, at **1155**. The overwritten instruction may then be replaced with a stored copy, at **1160**. The virtual machine is then allowed to execute the instruction, at **1165**, and trap back into the virtualization assistance layer at **1170**. The instruction may then be overwritten with a trapping instruction at **1175**. Finally, a notification of memory access and the specification may be sent to the requesting guest, at **1180**.

[**0115**] At a high level, as may apply to the above examples, the actions taken on monitored activity may include policy based actions taken by, and/or coordinated between, the Separation Kernel Hypervisor **100**, virtualization assistance layer **400**, and/or instruction execution detection/interception mechanisms **500**. Such actions may include and/or involve, though are not limited to any of the following: (1) preventing the monitored activity; (2) allowing the monitored activity; (3) allowing the monitored activity, with instrumentation, and/or partial blocking. It may be that certain sub-sets of the activity are permissible (by configuration policy), and that a portion of the activity may be allowed and a portion blocked

and/or substituted with a harmless surrogate; such as insertion of no-ops in malicious code to render malicious code inert. This may include run-time patching of CPU state of a guest operating system 300, and/or any resources of the guest operating system 300; (4) reporting on the monitored activity, possibly exporting reports to other software in the system, or on remote systems; and/or (5) replay of the monitored activity.

[0116] With regard to (5), immediately above, in separation kernel hypervisor 100 configurations supporting rewind of the guest operating system 300 state, the state of the guest operating system 300 can be rewound and the monitored activity can be replayed and re-monitored (to a degree); e.g., if the instruction execution detection/interception mechanisms 500 requires more systems resources, and/or to map more context of the guest operating system 300, the instruction execution detection/interception mechanisms 500 may request a rewind, request more resources, then request the replay of the monitored activity; so that the instruction execution detection/interception mechanisms 500 may perform analysis of the monitored activity with the advantage of more resources. Systems and methods of monitoring activity, as may be utilized by the separation kernel hypervisor 100, virtualization assistance layer 400, and/or instruction execution detection/interception mechanisms 500; for activities which may include guest operating system 300 activities, and/or separation kernel hypervisor 100, virtualization assistance layer 400, and/or instruction execution detection/interception mechanisms 500 activities (such as feedback between such components), including those activities which may cause transition to the separation kernel hypervisor 100, virtualization assistance layer 400, and/or instruction execution detection/interception mechanisms 500 include (but are not limited to): synchronous mechanisms, bound to a specific instruction stream and/or sequence within a processor, CPU, or platform device and/or ABI, certain elements of which can be used to trap and/or transition to/from the hypervisor. For example, instructions which induce trapping. Such events may be generated by the Separation Kernel Hypervisor 100, virtualization assistance layer 400, and/or instruction execution detection/interception mechanisms 500.

[0117] The innovations and mechanisms herein may also provide or enable means by which software and/or guest operating system vulnerabilities, including improper use of CPU interfaces, specifications, and/or ABIs may be detected and/or prevented; including cases where software vendors have implemented emulation and/or virtualization mechanisms improperly.

[0118] Implementations and Other Nuances

[0119] The innovations herein may be implemented via one or more components, systems, servers, appliances, other sub-components, or distributed between such elements. When implemented as a system, such system may comprise, inter alia, components such as software modules, general-purpose CPU, RAM, etc. found in general-purpose computers, and/or FPGAs and/or ASICs found in more specialized computing devices. In implementations where the innovations reside on a server, such a server may comprise components such as CPU, RAM, etc. found in general-purpose computers.

[0120] Additionally, the innovations herein may be achieved via implementations with disparate or entirely different software, hardware and/or firmware components, beyond that set forth above. With regard to such other components (e.g., software, processing components, etc.) and/or

computer-readable media associated with or embodying the present inventions, for example, aspects of the innovations herein may be implemented consistent with numerous general purpose or special purpose computing systems or configurations. Various exemplary computing systems, environments, and/or configurations that may be suitable for use with the innovations herein may include, but are not limited to: software or other components within or embodied on personal computers, appliances, servers or server computing devices such as routing/connectivity components, hand-held or laptop devices, multiprocessor systems, microprocessor-based systems, set top boxes, consumer electronic devices, network PCs, other existing computer platforms, distributed computing environments that include one or more of the above systems or devices, etc.

[0121] In some instances, aspects of the innovations herein may be achieved via logic and/or logic instructions including program modules, executed in association with such components or circuitry, for example. In general, program modules may include routines, programs, objects, components, data structures, etc. that perform particular tasks or implement particular instructions herein. The inventions may also be practiced in the context of distributed circuit settings where circuitry is connected via communication buses, circuitry or links. In distributed settings, control/instructions may occur from both local and remote computer storage media including memory storage devices.

[0122] Innovative software, circuitry and components herein may also include and/or utilize one or more type of computer readable media. Computer readable media can be any available media that is resident on, associable with, or can be accessed by such circuits and/or computing components. By way of example, and not limitation, computer readable media may comprise computer storage media and other non-transitory media. Computer storage media includes volatile and nonvolatile, removable and non-removable media implemented in any method or technology for storage of information such as computer readable instructions, data structures, program modules or other data. Computer storage media includes, but is not limited to, RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical storage, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired information and can accessed by computing component. Other non-transitory media may comprise computer readable instructions, data structures, program modules or other data embodying the functionality herein, in various non-transitory formats. Combinations of the any of the above are also included within the scope of computer readable media.

[0123] In the present description, the terms component, module, device, etc. may refer to any type of logical or functional circuits, blocks and/or processes that may be implemented in a variety of ways. For example, the functions of various circuits and/or blocks can be combined with one another into any other number of modules. Each module may even be implemented as a software program stored on a tangible memory (e.g., random access memory, read only memory, CD-ROM memory, hard disk drive, etc.) to be read by a central processing unit to implement the functions of the innovations herein. Or, the modules can comprise programming instructions transmitted to a general purpose computer, to processing/graphics hardware, and the like. Also, the modules can be implemented as hardware logic circuitry imple-

menting the functions encompassed by the innovations herein. Finally, the modules can be implemented using special purpose instructions (SIMD instructions), field programmable logic arrays or any mix thereof which provides the desired level performance and cost.

[0124] As disclosed herein, features consistent with the present inventions may be implemented via computer-hardware, software and/or firmware. For example, the systems and methods disclosed herein may be embodied in various forms including, for example, a data processor, such as a computer that also includes a database, digital electronic circuitry, firmware, software, or in combinations of them. Further, while some of the disclosed implementations describe specific hardware components, systems and methods consistent with the innovations herein may be implemented with any combination of hardware, software and/or firmware. Moreover, the above-noted features and other aspects and principles of the innovations herein may be implemented in various environments. Such environments and related applications may be specially constructed for performing the various routines, processes and/or operations according to the invention or they may include a general-purpose computer or computing platform selectively activated or reconfigured by code to provide the necessary functionality. The processes disclosed herein are not inherently related to any particular computer, network, architecture, environment, or other apparatus, and may be implemented by a suitable combination of hardware, software, and/or firmware. For example, various general-purpose machines may be used with programs written in accordance with teachings of the invention, or it may be more convenient to construct a specialized apparatus or system to perform the required methods and techniques.

[0125] Aspects of the method and system described herein, such as the logic, may also be implemented as functionality programmed into any of a variety of circuitry, including programmable logic devices (“PLDs”), such as field programmable gate arrays (“FPGAs”), programmable array logic (“PAL”) devices, electrically programmable logic and memory devices and standard cell-based devices, as well as application specific integrated circuits. Some other possibilities for implementing aspects include: memory devices, microcontrollers with memory (such as EEPROM), embedded microprocessors, firmware, software, etc. Furthermore, aspects may be embodied in microprocessors having software-based circuit emulation, discrete logic (sequential and combinatorial), custom devices, fuzzy (neural) logic, quantum devices, and hybrids of any of the above device types. The underlying device technologies may be provided in a variety of component types, e.g., metal-oxide semiconductor field-effect transistor (“MOSFET”) technologies like complementary metal-oxide semiconductor (“CMOS”), bipolar technologies like emitter-coupled logic (“ECL”), polymer technologies (e.g., Silicon-conjugated polymer and metal-conjugated polymer-metal structures), mixed analog and digital, and so on.

[0126] It should also be noted that the various logic and/or functions disclosed herein may be enabled using any number of combinations of hardware, firmware, and/or as data and/or instructions embodied in various machine-readable or computer-readable media, in terms of their behavioral, register transfer, logic component, and/or other characteristics. Computer-readable media in which such formatted data and/or instructions may be embodied include, but are not limited to, non-volatile storage media in various forms (e.g., optical,

magnetic or semiconductor storage media), though do not include transitory media such as carrier waves.

[0127] Unless the context clearly requires otherwise, throughout the description, the words “comprise,” “comprising,” and the like are to be construed in an inclusive sense as opposed to an exclusive or exhaustive sense; that is to say, in a sense of “including, but not limited to.” Words using the singular or plural number also include the plural or singular number respectively. Additionally, the words “herein,” “hereunder,” “above,” “below,” and words of similar import refer to this application as a whole and not to any particular portions of this application. When the word “or” is used in reference to a list of two or more items, that word covers all of the following interpretations of the word: any of the items in the list, all of the items in the list and any combination of the items in the list.

[0128] Although certain presently preferred implementations of the inventions have been specifically described herein, it will be apparent to those skilled in the art to which the inventions pertain that variations and modifications of the various implementations shown and described herein may be made without departing from the spirit and scope of the inventions. Accordingly, it is intended that the inventions be limited only to the extent required by the applicable rules of law.

1. A method for processing information securely, the method comprising:

- partitioning hardware platform resources via a separation kernel hypervisor into a plurality of guest operating system virtual machine protection domains;
- providing a dedicated virtualization assistance layer (VAL) including a virtual representation of the hardware platform that is a virtual machine in each of the guest operating system virtual machine protection domains such that the dedicated VAL security processing is not performed in the separation kernel hypervisor;
- processing the virtual machine via another guest;
- hosting at least one detection mechanism that executes within the virtual hardware platform in each of the plurality of guest operating system virtual machine protection domains via the separation kernel hypervisor;
- upon detection of suspect behavior, securely transitioning execution to the detection mechanism within the VAL in a manner isolated from the guest operating system;
- securely determining, via the detection mechanism, a policy decision regarding the suspect behavior; and
- transitioning execution back to the separation kernel hypervisor to continue processing regarding enforcement of or taking action in connection with the policy decision.

2. The method of claim **1** further comprising one or more of:

- providing a list of memory locations of an authorized guest to another guest;
- associating each of a plurality of physical memory locations with a respective specification of execution context information upon access to the each of the plurality of physical memory locations;
- providing a message of the specification to the another guest;
- providing the execution context information to the another guest; and
- providing a virtualization assistance layer (VAL) including a virtual representation of the hardware platform in each

of the guest operating system virtual machine protection domains such that the VAL is not directly accessible by the authorized guest.

3. The method of claim **1** wherein the step of securely transitioning execution to the detection mechanism within the VAL in a manner isolated from the guest operating system includes:

performing processing to transition control of the suspect behavior to a memory management control unit for handling; and

transitioning handling of the suspect behavior to an instruction execution detection mechanism to process and handle back to.

4. The method of claim **1** wherein the securely determining, via the detection mechanism, the policy decision regarding the suspect behavior comprises:

analyzing the behavior of the guest operating system and its resources; and

processing unmapped page exceptions taken by the guest operating system based on the unmapped pages.

5.-6. (canceled)

7. The method of claim **1**, wherein the step of securely transitioning execution to the detection mechanism within the VAL in a manner isolated from the guest operating system includes:

performing processing to transition control of the suspect behavior to a memory management control unit for handling; and

transitioning handling of the suspect behavior to an instruction execution detection mechanism to process and handle back to.

8. The method of claim **1**, further comprising one or more of:

hosting an unmapping mechanism to unmap the page; and unmapping, via the unmapping mechanism, the page.

9. The method of claim **1**, further comprising one or more of:

hosting a remapping mechanism to remap unmapped pages;

transitioning execution to the remapping mechanism;

remapping, via the remapping mechanism, the unmapped pages as accessible; and

transitioning, via the remapping mechanism, execution to the detection mechanism.

10. The method of claim **1**, further comprising one or more of:

hosting an unmapping mechanism to unmap specified pages on demand from the another guest;

receiving, via the unmapping mechanism, the demand to unmap the specified pages; and

unmapping, via the unmapping mechanism, the specified pages.

11. The method of claim **1**, further comprising one or more of:

hosting a remapping mechanism to remap unmapped pages;

upon securely determining, via the detection mechanism, the policy decision, transitioning execution to the remapping mechanism;

remapping, via the remapping mechanism, the unmapped pages as accessible; and

transitioning, via the remapping mechanism; execution to the detection mechanism.

12. The method of claim **1**, further comprising: implementing a separation kernel hypervisor that ensures isolation of multiple guest operating systems, each guest operating system in its own virtual machine.

13. The method of claim **1**, further comprising: implementing a separation kernel hypervisor that implements a mechanism wherein a suitably authorized guest can send a list of memory locations and/or physical memory locations to be watched to another guest.

14.-29. (canceled)

30. A method for processing information securely, the method comprising:

partitioning hardware platform resources via a separation kernel hypervisor into a plurality of guest operating system virtual machine protection domains;

isolating the domains in time and space from each other;

providing a list of memory locations of an authorized guest to another guest;

associating each of a plurality of physical memory locations with a respective specification of execution context information upon access to the each of the plurality of physical memory locations;

providing a message of the specification to the another guest;

providing the execution context information to the another guest; and

providing a virtualization assistance layer (VAL) including a virtual representation of the hardware platform in each of the guest operating system virtual machine protection domains such that the VAL is not directly accessible by the authorized guest.

31. The method of claim **30** wherein the virtual representation of the hardware platform is a virtual machine comprising a virtual motherboard including a virtual CPU and memory by the VAL.

32. The method of claim **31** the method further comprising: hosting a mechanism to map physical memory pages as non-executable;

processing exceptions to non-executable page execution attempts by the associated virtual machine;

hosting another mechanism to determine whether the physical memory locations are accessed;

replacing the context information in the virtual machine;

pausing or resuming execution of the virtual machine; and

sending a notification of memory access and the specification to a requesting guest.

33. The method of claim **30** wherein the step of securely transitioning execution to the detection mechanism within the VAL in a manner isolated from the guest operating system includes:

performing processing to transition control of the suspect behavior to a memory management control unit for handling; and

transitioning handling of the suspect behavior to an instruction execution detection mechanism to process and handle back.

34. The method of claim **31** the method further comprising: hosting a mechanism to copy the contents of the physical memory address into a private memory location;

hosting a mechanism to overwrite the address with an instruction that will trap into the separation kernel hypervisor;

processing exceptions due to execution attempts of the overwritten address by the associated virtual machine;

hosting another mechanism to determine whether the physical memory locations are accessed;
pausing or resuming execution of the virtual machine;
replacing the over written instruction with the stored copy;
allowing the virtual machine to execute the original instruction;
trapping back into the virtualization assistance layer;
overwriting the original instruction with the trapping instruction; and
sending a notification of memory access and the specification to a requesting guest.

35. The method of claim **31** further comprising:
hosting a first mechanism to copy contents of the physical memory location into a private memory location.

36. The method of claim **35** further comprising:
hosting a second mechanism to overwrite the location with an instruction to trap into the separation kernel hypervisor; and

processing an exception due to execution attempts of the overwritten location by the associated virtual machine.

37. The method of claim **36** further comprising:
hosting a third mechanism to determine whether the physical memory location is accessed.

38. The method of claim **37** further comprising:
pausing or resuming execution of the virtual machine;
replacing the overwritten instruction with a stored copy;
and

allowing the virtual machine to execute the instruction.

39.-64. (canceled)

* * * * *