

PROCESSOR FEATURE ID RESPONSE FOR VIRTUALIZATION

BACKGROUND

5 [0001] Virtualization technology is employed in a wide variety of contexts. For example, virtualization technology may be employed to abstract physical computing resources from guest partitions, to enable higher utilization of the physical computing resources, to enable portability of guest partitions across physical devices, to protect the physical computing resources from malicious and/or erroneous code running in guest partitions, to protect secrets, to enforce security requirements or policy, and/or the like. In
10 prior virtualization technology, a guest exit (e.g., a transfer of control of a processor from the guest partition to a supervisory partition, such as a hypervisor) may occur in response to certain operations. For example, a guest exit may occur in response to a request for processor feature ID information.

SUMMARY OF THE DISCLOSURE

15 [0002] This Summary is provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used to limit the scope of the claimed subject matter.

[0003] Briefly stated, the disclosed technology is generally directed to virtualization
20 technology. The disclosed technology includes providing processor feature ID information requested by, or from, a virtual machine (VM), virtualized application, Virtualization Based Security (VBS) user mode process, VBS kernel mode process, or other guest partition, by a processor. Such information may be provided based on information provided a priori to the processor, for example, by a supervisory partition,
25 such as a hypervisor. The disclosed technology also includes a supervisory partition, for example, that provides such information to the processor, and includes guest partitions that receive such information.

[0004] Other aspects of and applications for the disclosed technology will be appreciated upon reading and understanding the attached figures and description.

30 BRIEF DESCRIPTION OF THE DRAWINGS

[0005] Non-limiting and non-exhaustive examples of the present disclosure are described with reference to the following drawings. In the drawings, like reference numerals refer to like parts throughout the various figures unless otherwise specified. These drawings are not necessarily drawn to scale.

[0006] For a better understanding of the present disclosure, reference will be made to the following Detailed Description, which is to be read in association with the accompanying drawings, in which:

- 5 [0007] FIG. 1 is a block diagram illustrating a physical view of one example of a suitable computing device according to aspects of the disclosed technology;
- [0008] FIG. 2 is a block diagram illustrating a logical view of the example computing device according to aspects of the disclosed technology;
- [0009] FIG. 3 illustrates an example process according to aspects of the disclosed technology;
- 10 [0010] FIG. 4 illustrates another example process according to additional aspects of the disclosed technology; and
- [0011] FIG. 5 illustrates yet another example process according to other aspects of the disclosed technology.

DETAILED DESCRIPTION

- 15 [0012] The following description provides specific details for a thorough understanding of, and enabling description for, various examples of the technology. One skilled in the art will understand that the technology may be practiced without many of these details. In some instances, well-known structures and functions have not been shown or described in detail to avoid unnecessarily obscuring the description of examples
- 20 of the technology. It is intended that the terminology used in this disclosure be interpreted in its broadest reasonable manner, even though it is being used in conjunction with a detailed description of certain examples of the technology. Although certain terms may be emphasized below, any terminology intended to be interpreted in any restricted manner will be overtly and specifically defined as such in this Detailed Description section.
- 25 Throughout the specification and claims, the following terms take at least the meanings explicitly associated herein, unless the context dictates otherwise. The meanings identified below do not necessarily limit the terms, but merely provide illustrative examples for the terms. For example, each of the terms "based on" and "based upon" is not exclusive, and is equivalent to the term "based, at least in part, on", and includes the
- 30 option of being based on additional factors, some of which may not be described herein. As another example, the term "via" is not exclusive, and is equivalent to the term "via, at least in part", and includes the option of being via additional factors, some of which may not be described herein. The meaning of "in" includes "in" and "on." The phrase "in one embodiment," or "in one example," as used herein does not necessarily refer to the same

embodiment or example, although it may. Use of particular textual numeric designators does not imply the existence of lesser-valued numerical designators. For example, reciting "a widget selected from the group consisting of a third foo and a fourth bar" would not itself imply that there are at least three foo, nor that there are at least four bar, elements.

- 5 References in the singular are made merely for clarity of reading and include plural references unless plural references are specifically excluded. The term "or" is an inclusive "or" operator unless specifically indicated otherwise. For example, the phrases "A or B" means "A, B, or A and B." As used herein, the terms "component" and "system" are intended to encompass hardware, software, or various combinations of hardware and
- 10 software. Thus, for example, a system or component may be a process, a process executing on a computing device, the computing device, or a portion thereof.

[0013] Briefly stated, the disclosed technology is generally directed to virtualization technology. The disclosed technology includes providing processor feature ID information requested by, or from, a virtual machine (VM), virtualized application,

- 15 Virtualization Based Security (VBS) user mode process, VBS kernel mode process, or other guest partition, by a processor. For example, such requests may be generated by code executing on a guest virtual processor within a guest partition. Such information may be provided based on information provided a priori to the processor, for example, by a supervisory partition, such as a hypervisor. The disclosed technology also includes a
- 20 supervisory partition, for example, that provides such information to the processor, and includes guest partitions that receive such information.

[0014] In some examples, the disclosed technology may be employed in virtualization/virtualized systems. For example, the technology may be employed in conjunction with hypervisors, virtual machines, virtualized applications, Virtualization

25 Based Security (VBS) user mode processes, VBS kernel mode processes, and/or the like. For example, the technology may include precomputing or otherwise determining processor feature ID information by a hypervisor or other supervisory partition, and providing the determined processor feature ID information to a processor of a computing device for later use in responding to requests for the processor feature ID information, e.g.,

30 from a guest partition. For example, requests for such information from a guest partition can then be handled by the processor without intervention of the supervisory partition.

[0015] Use of the disclosed technology may be employed to reduce and/or eliminate the frequency of VM exits or other guest exits in which a system exits from execution of guest code in order to execute supervisory partition code. As one example of a guest exit,

in response to a guest's query for processor feature ID information, a supervisory partition process may take control of the processor from the guest partition, set up processor feature ID registers of the processor, then return control of the processor to the guest partition.

[0016] In certain prior virtualization systems, guest exits occur in response to certain operations being executed from a guest partition. For example, a guest exit may occur in response to operation being executed from the guest partition in which a CPU ID or other processor feature ID information is requested. Such exits are typically "expensive" in terms of time, power, and computing performance. For example, guest exits may be associated with use of processing bandwidth for supervisory partition code (e.g., overhead) rather than guest partition code (e.g., the intended workload). Guest exits may also be associated with additional context switching overhead due to over writing of processor data and/or instruction cache locations with supervisory partition code, and/or the like.

Illustrative Physical Computing Device

[0017] FIG. 1 is a diagram illustrating one example of a physical view of computing device 100 in which aspects of the technology may be practiced. Computing device 100 may be virtually any type of general- or specific-purpose computing device. For example, computing device 100 may be a user device such as a desktop computer, a laptop computer, a tablet computer, a display device, a camera, a printer, or a smartphone. Likewise, computing device 100 may also be server device such as an application server computer, a virtual computing host computer, or a file server computer. Computing device 100 may also be an IoT device that connects to a network to receive IoT services. Likewise, computer device 100 may be an example any of the devices illustrated in or referred to in FIGs. 2-5, as discussed in greater detail below. As illustrated in FIG. 1, computing device 100 includes processing circuit 110, operating memory 120, memory controller 130, data storage memory 150, input interface 160, output interface 170, and network adapter 180. Each of these afore-listed components of computing device 100 includes at least one hardware element.

[0018] Computing device 100 includes at least one processing circuit 110 configured to execute instructions, such as instructions for implementing the herein-described workloads, processes, or technology. Processing circuit 110 may include a microprocessor, a microcontroller, a graphics processor, a coprocessor, a field-programmable gate array, a programmable logic device, a signal processor, or any other circuit suitable for processing data. Processing circuit 110 is an example of a core. The

aforementioned instructions, along with other data (e.g., datasets, metadata, operating system instructions, etc.), may be stored in operating memory 120 during run-time of computing device 100. Operating memory 120 may also include any of a variety of data storage devices/components, such as volatile memories, semi-volatile memories, random access memories, static memories, caches, buffers, or other media used to store run-time information. In one example, operating memory 120 does not retain information when computing device 100 is powered off. Rather, computing device 100 may be configured to transfer instructions from a non-volatile data storage component (e.g., data storage component 150) to operating memory 120 as part of a booting or other loading process.

5 [0019] Operating memory 120 may include 4th generation double data rate (DDR4) memory, 3rd generation double data rate (DDR3) memory, other dynamic random access memory (DRAM), High Bandwidth Memory (HBM), Hybrid Memory Cube memory, 3D-stacked memory, static random access memory (SRAM), or other memory, and such memory may comprise one or more memory circuits integrated onto a DIMM, SIMM, SODIMM, or other packaging. Such operating memory modules or devices may be organized according to channels, ranks, and banks. For example, operating memory devices may be coupled to processing circuit 110 via memory controller 130 in channels. One example of computing device 100 may include one or two DIMMs per channel, with one or two ranks per channel. Operating memory within a rank may operate with a shared clock, and shared address and command bus. Also, an operating memory device may be organized into several banks where a bank can be thought of as an array addressed by row and column. Based on such an organization of operating memory, physical addresses within the operating memory may be referred to by a tuple of channel, rank, bank, row, and column.

15 [0020] Despite the above-discussion, operating memory 120 specifically does not include or encompass communications media, any communications medium, or any signals per se.

[0021] Memory controller 130 is configured to interface processing circuit 110 to operating memory 120. For example, memory controller 130 may be configured to interface commands, addresses, and data between operating memory 120 and processing circuit 110. Memory controller 130 may also be configured to abstract or otherwise manage certain aspects of memory management from or for processing circuit 110. Although memory controller 130 is illustrated as single memory controller separate from processing circuit 110, in other examples, multiple memory controllers may be employed,

memory controller(s) may be integrated with operating memory 120, or the like. Further, memory controller(s) may be integrated into processing circuit 110. These and other variations are possible.

5 [0022] In computing device 100, data storage memory 150, input interface 160, output interface 170, and network adapter 180 are interfaced to processing circuit 110 by bus 140. Although, FIG. 1 illustrates bus 140 as a single passive bus, other configurations, such as a collection of buses, a collection of point to point links, an input/output controller, a bridge, other interface circuitry, or any collection thereof may also be suitably employed for interfacing data storage memory 150, input interface 160, output interface 170, or network
10 adapter 180 to processing circuit 110.

[0023] In computing device 100, data storage memory 150 is employed for long-term non-volatile data storage. Data storage memory 150 may include any of a variety of non-volatile data storage devices/components, such as non-volatile memories, disks, disk drives, hard drives, solid-state drives, or any other media that can be used for the non-
15 volatile storage of information. However, data storage memory 150 specifically does not include or encompass communications media, any communications medium, or any signals per se. In contrast to operating memory 120, data storage memory 150 is employed by computing device 100 for non-volatile long-term data storage, instead of for run-time data storage.

20 [0024] Also, computing device 100 may include or be coupled to any type of processor-readable media such as processor-readable storage media (e.g., operating memory 120 and data storage memory 150) and communication media (e.g., communication signals and radio waves). While the term processor-readable storage media includes operating memory 120 and data storage memory 150, the term "processor-readable storage media," throughout the specification and the claims whether used in the
25 singular or the plural, is defined herein so that the term "processor-readable storage media" specifically excludes and does not encompass communications media, any communications medium, or any signals per se. However, the term "processor-readable storage media" does encompass processor cache, Random Access Memory (RAM),
30 register memory, and/or the like.

[0025] Computing device 100 also includes input interface 160, which may be configured to enable computing device 100 to receive input from users or from other devices. In addition, computing device 100 includes output interface 170, which may be configured to provide output from computing device 100. In one example, output

interface 170 includes a frame buffer, graphics processor, graphics processor or accelerator, and is configured to render displays for presentation on a separate visual display device (such as a monitor, projector, virtual computing client computer, etc.). In another example, output interface 170 includes a visual display device and is configured to render and present displays for viewing. In yet another example, input interface 160 and/or output interface 170 may include a universal asynchronous receiver/transmitter ("UART"), a Serial Peripheral Interface ("SPI"), Inter-Integrated Circuit ("IIC"), a General-purpose input/output (GPIO), and/or the like. Moreover, input interface 160 and/or output interface 170 may include or be interfaced to any number or type of peripherals.

[0026] In the illustrated example, computing device 100 is configured to communicate with other computing devices or entities via network adapter 180. Network adapter 180 may include a wired network adapter, e.g., an Ethernet adapter, a Token Ring adapter, or a Digital Subscriber Line (DSL) adapter. Network adapter 180 may also include a wireless network adapter, for example, a Wi-Fi adapter, a Bluetooth adapter, a ZigBee adapter, a Long Term Evolution (LTE) adapter, or a 5G adapter.

[0027] Although computing device 100 is illustrated with certain components configured in a particular arrangement, these components and arrangement are merely one example of a computing device in which the technology may be employed. In other examples, data storage memory 150, input interface 160, output interface 170, or network adapter 180 may be directly coupled to processing circuit 110, or be coupled to processing circuit 110 via an input/output controller, a bridge, or other interface circuitry. Other variations of the technology are possible.

[0028] Some examples of computing device 100 include at least one memory (e.g., operating memory 120) adapted to store run-time data and at least one processor (e.g., processing unit 110) that is adapted to execute processor-executable code that, in response to execution, enables computing device 100 to perform actions.

Illustrative Logical Computing Device

[0029] FIG. 2 is a diagram illustrating one example of a logical view of computing device 200 in which aspects of the technology may be practiced. Computing device 200 may be an example of computing device 100 of FIG. 1. In the illustration of FIG 2, the logical components of computing device 200 include guest partitions 211-213, supervisory partition 230, physical resources 241-243, and processor feature ID 250.

[0030] Physical resources 241-243 may include any variety of physical components,

such as processor components, input/output (I/O) components, and/or other components or devices. For example, physical resources 241-243 may include any suitable combination of physical components, such as those discussed in conjunction with FIG. 1. Although illustrated as part of computing device 200, one or more of physical resources 241-243 (e.g., one or more data storage memories) can be implemented external to computing device 200. Various components or modules running on computing device 200, including supervisory partition 230, may access functionality(-ies) provided via physical resources 241-243 directly and/or indirectly via other components or modules.

[0031] Supervisory partition 230 may generate any number of guest partitions, e.g., guest partitions 211-213. Each of guest partitions 211-213 may be a VM, a virtualized application, a VBS execution environment, user mode process, and/or the like. For example, guest partition 211 is illustrated as a VM with operating system (OS) 221 and application 222, guest partition 212 is illustrated as a virtualized application 223, and guest partition 224 is illustrated as having process 224 executing therefrom.

[0032] Each of guest partitions 211-213 is a logical unit of isolation from which an operating system and/or other software executes. Each of guest partitions 211-213 may also include a guest virtual processor. The software executing within each of guest partitions 211-213 is isolated from the software executing in each of the other of guest partitions. E.g., the software executing in each of guest partitions 211-213 is not able to access, and need not be aware of, the software executing in each of the other guest partitions. Physical resources 241-243 are virtualized to guest partitions 211-213, and access to physical resources 241-243 is managed by supervisory partition 230.

[0033] As illustrated, computing device 100 includes supervisory partition 230. Supervisory partition 230, may include hypervisor such as a virtual machine monitor that manages access to the functionality provided by physical resources 241-243. In another example, supervisory partition 230 is kernel or kernel mode process of an OS, such as an OS that employs VBS.

[0034] Computing device 200 also includes one or more processor feature ID 250. For example, processor feature ID 250 may represent a physical hardware ID register (or set of registers) of a physical processor, such as a register (or set of registers) containing x86 CPU ID leaf information, ID registers for an Advanced Reduced Instruction Set Computing Machine (ARM) processor, and/or the like. Processor feature ID 250 may also represent a feature supported by a processor, a feature set supported by the processor, a physical characteristic of the processor, and/or the like. For example, processor feature ID

250 may represent a processor frequency, a supported physical address width, a clock multiplier, a power setting, availability of an instruction, a stepping number, a serial number, and/or the like.

Illustrative Processes

5 **[0035]** For clarity, the processes described herein are described in terms of operations performed in particular sequences by particular devices or components of a system. However, it is noted that other processes are not limited to the stated sequences, devices, or components. For example, certain acts may be performed in different sequences, in parallel, omitted, or may be supplemented by additional acts or features, whether or not
10 such sequences, parallelisms, acts, or features are described herein. Likewise, any of the technology described in this disclosure may be incorporated into the described processes or other processes, whether or not that technology is specifically described in conjunction with a process. The disclosed processes may also be performed on or by other devices, components, or systems, whether or not such devices, components, or systems are
15 described herein. These processes may also be embodied in a variety of ways. For example, they may be embodied on an article of manufacture, e.g., as processor-readable instructions stored in a processor-readable storage medium or be performed as a computer-implemented process. As an alternate example, these processes may be encoded as processor-executable instructions and transmitted via a communications medium.

20 **[0036]** FIG. 3 illustrates example process 300 which is illustrated from the viewpoint of a processor of a computing device, e.g., processing circuit 110 of FIG. 1 or physical resource 241, 242, or 243 of FIG. 2. Process 300 starts at 381 where a request for a processor feature ID is received. For example, this request may be received by a physical processor, and this request may have originated from a guest partition. For example, this
25 request may be generated by code executing on a guest virtual processor within a guest partition.

30 **[0037]** From 381, processing flows to 382 where, for example, the processor looks up the processor feature ID. In one example, the processor looks up the processor feature ID from a processor feature ID register, from a memory structure accessible to the processor, from a processor feature ID lookup table, and/or the like. This look up may also be based on information that was provided a priori to the processor by a hypervisor or other supervisory partition, e.g., as further discussed in conjunction with process 400 of FIG. 4.

35 **[0038]** Processing then flows to 383 where the processor feature ID is provided, e.g., without a guest exit. This may include the processor providing the requested processor

feature ID to the guest partition, such as to a requesting process from the guest partition. The provided processor feature ID may be based on the information that was provided a priori to the processor. Following 383, processing returns to other operations.

5 [0039] FIG. 4 illustrates example process 400 which is illustrated from the viewpoint of a supervisory partition, e.g., a hypervisor. For example, process 400 is illustrated from the viewpoint of supervisory partition 230 of FIG 2.

[0040] Process 400 starts at 481 where a value for a processor feature ID is obtained. This value may be obtained by the supervisory partition in conjunction with, or in response to, a setup operation for at least one of a guest partition, guest virtual processor, 10 or guest virtual processor virtual trust level. For example, such setup operation may include at least one of a generation, instantiation, or startup of the guest partition, guest virtual processor, or guest virtual processor virtual trust level. The value may be obtained, for example, by reading a corresponding hardware register value of a physical processor of the computing device. As one example, the obtained processor feature ID value may be a 15 value assigned by a manufacturer of the physical processor (e.g., an on-chip value), and/or be representative and or indicative of a feature/feature set supported by the processor.

[0041] Processing optionally flows to 482 where the processor feature ID that is to be stored is calculated or otherwise determined. For example, 482 may include determining whether the obtained processor function ID value is to be stored without change, or 20 whether a different value is to be stored. For example, a different value may be stored if the supervisory partition is to "offer" a different set of processor features to the guest partition than natively supported by the processor, e.g., for enhanced guest partition portability, performance, and/or other reasons. At 482, multiple values for an obtained processor function ID value may also be determined, e.g., each associated with and/or 25 specified for a different guest partition and/or guest virtual processor and/or trust level contexts. In some examples, different processor function ID values may be employed on a per virtual processor context or virtual processor's virtual trust level context. Examples of a per virtual processor context or virtual processor's virtual trust level context are the Virtual Machine Control Structure (VMCS), Virtual Machine Control Block (VMCB), a 30 set of system registers containing guest context, or other virtualization instruction-set-architecture-specific collection of guest context.

[0042] Processing flows to 483 where the processor feature ID of 482 is provided to the processor. For example, the supervisory partition may provide the processor feature ID determined at 482 to the processor for later use by the processor in responding to

requests from guest partition software for processor feature ID information. At 483, the processor may also receive this processor feature ID information from the supervisory partition.

5 [0043] Processing then flows to 484. At 484, the provided processor feature ID is stored. For example, 484 may include actions by the supervisory process that cause the processor to store the provided processor feature ID and/or actions undertaken by the processor to perform the storage. For example, this storage may include storing the processor feature ID in a processor feature ID register, in a processor feature ID lookup table in memory, and/or the like. In yet another example, the supervisory partition code
10 may store the processor feature ID through a processor interface instruction, such as an ID_REGISTER_WRITE instruction.

[0044] In yet another example, the processor feature ID may be written to a memory structure accessible to the processor, and the location of that memory structure programmed to a processor register. In this and other examples, multiple memory structures may be
15 employed, e.g., to enable use of different processor function ID values for different guest partitions, for different VMCSs, virtual processor contexts, trust level contexts, and/or the like. In use, context switching between the different memory structures could be performed by either the processor or by supervisory partition code 230, e.g., in response to control of the processor switching from one guest partition, guest virtual processor, or
20 guest virtual processor virtual trust level to another.

[0045] Optionally, 484 may also include storing conditional expressions to the processor that direct the processor to the appropriate value/memory structure of multiple values/memory structures for a given request from a guest partition. For example, an expression may be stored via an instruction such as an
25 ID_REGISTER_WRITE(VP_CONTEXT, REGISTER_NAME, CONDITIONAL, VALUE) instruction.

[0046] As discussed above, this storage may also enable subsequent use, by the processor, of the processor feature ID to respond to requests from a guest partition without causing a guest exit. Process 400 may be repeated for additional processor function IDs,
30 or multiple processor function IDs may be obtained, determined, provided, and stored in a single iteration of process 400. Following 484, processing returns to other operations.

[0047] FIG. 5 illustrates an example process 500 which is illustrated from the viewpoint of a guest partition, e.g., guest partition 211, 212, or 213 of FIG 2. Process 500 starts at 581 where a determination is made that a processor feature ID is to be obtained.

For example, this determination may be made by or on a guest partition in response to a request for processor feature ID information from an application or other process executing from the guest partition. As another example, this determination may represent a request by a virtual machine that the processor feature ID is to be obtained for use by the
5 virtual machine.

[0048] Processing then flows to 582 where a request for the processor feature ID is sent from the guest partition. This request may be sent to the processor. In response to the request of 582, at 583 the processor feature ID may be received from the processor. As discussed above, the received processor feature ID may be in accordance with information
10 that was provided a priori to the processor by the supervisory partition. From 583, processing may flow to 584 where the received processor feature ID is provided to the requestor, e.g., the application or other process on the guest partition. Following 584, processing returns to other operations.

Conclusion

[0049] While the above Detailed Description describes certain examples of the technology, and describes the best mode contemplated, no matter how detailed the above appears in text, the technology can be practiced in many ways. Details may vary in implementation, while still being encompassed by the technology described herein. As noted above, particular terminology used when describing certain features or aspects of the
20 technology should not be taken to imply that the terminology is being redefined herein to be restricted to any specific characteristics, features, or aspects with which that terminology is associated. In general, the terms used in the following claims should not be construed to limit the technology to the specific examples disclosed herein, unless the Detailed Description explicitly defines such terms. Accordingly, the actual scope of the
25 technology encompasses not only the disclosed examples, but also all equivalent ways of practicing or implementing the technology.