

**(19) AUSTRALIAN PATENT OFFICE**

(54) Title  
Providing secure input to a system with a high-assurance execution environment

(51)<sup>6</sup> International Patent Classification(s)  
**G06F** 1/00 (2006.01) OBMJP **G06F**  
**G06F** 9/46 (2006.01) 15/00  
**G06F** 15/00 (2006.01) 20060101ALI2006010  
**G06F** 21/00 (2006.01) 1BMKR **G06F**  
**G06F** 21/20 (2006.01) 21/00  
G06F 1/00 20060101ALI2005100  
20060101AFI2005122 8BMEP **G06F**  
OBMJP **G06F** 21/20  
9/46 20060101ALI2005122  
20060101ALI2005122 OBMJP

(21) Application No: 2004216665 (22) Application Date: 2004 .09 .30

(30) Priority Data

(31) Number	(32) Date	(33) Country
10/693061	2003 .10 .24	US

(43) Publication Date : 2005 .05 .12  
(43) Publication Journal Date : 2005 .05 .12

(71) Applicant(s)  
Microsoft Corporation

(72) Inventor(s)  
Chew, Christine M.; Yamamoto, Hirofumi; Paff, John E.; Roberts, Paul C.; Avraham, Idan

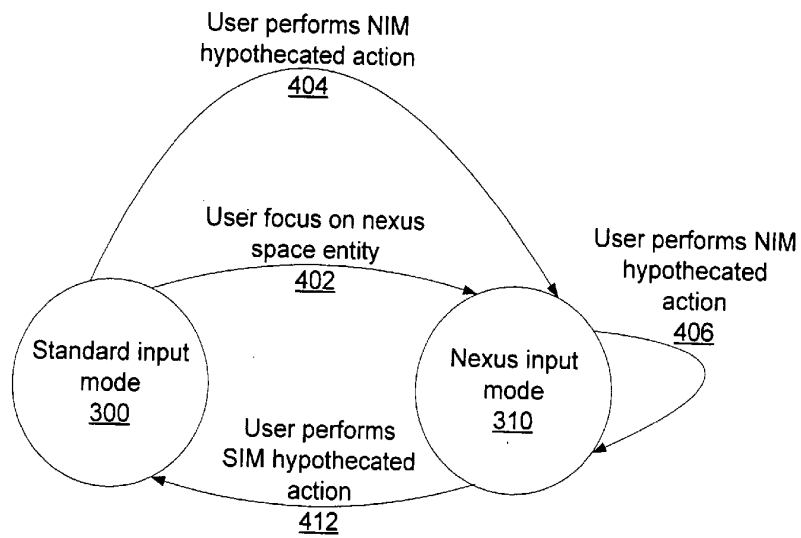
(74) Agent/Attorney  
Davies Collison Cave, 1 Nicholson Street, Melbourne, VIC, 3000

(56) Related Art  
KILPATRICK et al. "Securing The X Window System With SELinux" January 2003  
EPSTEIN et al. "User Interface for a High Assurance Windowing System" TRW Systems  
Division 1993  
US 5822435

2004216665 30 Sep 2004

ABSTRACT OF THE DISCLOSURE

Techniques are disclosed to provide security for user input in which a first, host operating system is used along with a second, high assurance operating system, where the first system provides at least some of the infrastructure for the second system. Two modes are presented. In a first mode, user data is passed to the host operating system. In a second mode, user data is retained in the second operating system for the use of the second operating system or processes running on the second operating system. Transitions between the modes can be accomplished according to hypothesized user actions such as keystroke combinations, or when the user performs an action which indicates a programmatic activation of a process running in the second operating system. Where shadow graphical elements are run by the first operating system to indicate the location of graphical elements from processes running on the second operating system, this programmatic activation may be indicated by programmatic activation of a shadow graphical element.



**FIG. 4**

2004216665 30 Sep 2004

**AUSTRALIA**  
**PATENTS ACT 1990**  
**COMPLETE SPECIFICATION**

NAME OF APPLICANT(S)::

**Microsoft Corporation**

ADDRESS FOR SERVICE:

**DAVIES COLLISON CAVE**  
Patent Attorneys  
1 Nicholson Street, Melbourne, 3000, Australia

INVENTION TITLE:

Providing secure input to a system with a high-assurance execution environment

The following statement is a full description of this invention, including the best method of performing it known to me/us:-

2004216665 02 Feb 2010

**FIELD**

The present invention relates generally to the field of computer security. For example, the invention relates to the use of plural execution environments (e.g., operating systems) on a single computing device, and provides techniques that allow for the integrity of input to be secured in a higher-assurance execution environment from among the plural execution environments.

**BACKGROUND**

In modern computing, many tasks which can be performed on a computer require some level of security. In order to provide a level of security, there are several options. One is to perform all secure applications on a computer which is completely separate from any possibly insecure elements, or to use a virtual machine monitor (VMM) to allow complete separation between two execution environments (e.g. operating systems) running on a single computer system. However, this may be impractical. There may be a need, for cost or convenience reasons, for a secure execution environment to share resources with applications with unassured security, and those applications and those resources may be vulnerable to an attacker. Additionally, where a VMM is used, since a VMM requires full virtualization of the machine and all of its devices (thereby requiring that the VMM provide its own device driver for every possible device), a VMM is not well suited to an open architecture machine in which an almost limitless variety of devices can be added to the machine.

One way to provide the ability to share resources among two execution environments is to provide a computer system in which there is one "main" operating system that controls most processes and devices on a machine, and where a second operating system also exists. This second operating system is a small, limited-purpose operating system alongside the main operating system which performs certain limited tasks. One way to make an operating system "small" or "limited-purpose" is to allow the small operating system to borrow certain infrastructure (e.g., the scheduling facility, the memory manager, the device drivers, etc.) from

the "main" operating system. Since a VMM effectively isolates one operating system from another, this sharing of infrastructure is not practical using a VMM.

[0004] Certain other techniques allow operating systems to exist side-by-side on the same machine without the use of a VMM. One such technique is to have one operating system act as a "host" for the other operating system. (The operating system that the "host" is hosting is sometimes called a "guest.") In this case, the host operating system provides the guest with resources such as memory and processor time. Another such technique is the use of an "exokernel." An exokernel manages certain devices (e.g., the processor and the memory), and also manages certain types of interaction between the operating systems, although an exokernel – unlike a VMM – does not virtualize the entire machine. Even when an exokernel is used, it may be the case that one operating system (e.g., the "main" operating system) provides much of the infrastructure for the other, in which case the main operating system can still be referred to as the "host," and the smaller operating system as the "guest." Both the hosting model and the exokernel model allow useful types of interaction between operating systems that support sharing of infrastructure.

[0005] Thus, these techniques can be used to provide a computer system with at least two execution environments. One of these may be a "high-assurance" operating system, referred to herein as a "nexus." A high-assurance operating system is one that provides a certain level of assurance as to its behavior. For example, a nexus might be employed to work with secret information (e.g., cryptographic keys, etc.) that should not be divulged, by providing a curtailed memory that is guaranteed not to leak information to the world outside of the nexus, and by permitting only certain certified applications to execute under the nexus and to access the curtailed memory.

[0006] In a computer system with two execution environments, one of which is a nexus, it may be desirable for the nexus to be the guest operating system, and a second operating system, not subject to the same level of assurance as to behavior, to be the host operating system. This allows the nexus to be as small as possible. A small nexus allows a higher level of confidence in the assurance provided by the nexus. Therefore operating system functions be run by the host operating system.

2004216665 02 Feb 2010

One such operating system which may be run by the host operating system is a windowing system. When using a windowing system, a user's display will be populated with windows, areas on the screen which display information from an application. An application may have one or more windows. One window of all the windows displayed  
5 may have focus. The focus window may be indicated by a different border around the window, for example.

In traditional windowing systems, when a window has focus, it is generally the object of the user's input. Therefore, if a user types information using a keyboard, in many cases, the keystroke data will be sent by the operating system to the application which  
10 owns the window which has focus. Some keystrokes and other input actions may not be sent to the application which owns the window with focus. For example, in some windowing systems there is a keystroke command which minimizes all windows. Such a command will be handled by the windowing system, and not sent to the application owning the window with focus. The application owning the focused-on window may  
15 receive notification of the minimization of the window; however, the user keystrokes are intended for the windowing system, not the application owning the window having focus, and will not be sent to that application.

When the windowing system is run by the host operating system, rather than by the nexus, it is vulnerable to attack. This causes a problem because, as described, the  
20 windowing system tracks which window will receive user input by tracking which window has focus. Therefore an attack can be mounted on the system by shifting the focus without the user noticing, or even without indications which the user can notice.

An attack may occur, for example, via a program running on the host which switches the focus and captures keystrokes. The user is led to believe that the user's input  
25 will be directed towards a trustworthy entity – some hardware, system software, application, or window which is running in the nexus. However, at some point the focus shifts. The user input instead is being directed to the attacking program, where it is captured and may be used later.

Thus, because the windowing system is under the control of the host operating  
30 system and the user's input is generally directed to the window having focus at the time of the input, the high-assurance nature of the nexus is imperiled. Although embodiments of the invention are described with reference to a windowing system, there may be other

2004216665 02 Feb 2010

systems in which there are two or more levels of assurance, and where a user may choose which entity (e.g. an application) to interact with. In these systems, as well, certain input must be kept secure from lower-assurance applications and entities, and the same problems described above occur.

5 It is desired to address or ameliorate one or more disadvantages or limitations associated with the prior art, or to at least provide a useful alternative.

#### SUMMARY

In accordance with the present invention, there is provided a method for maintaining the security of a secured execution environment on a system comprising said  
10 secured execution environment and a second execution environment both on a single computing device, the method comprising:

accepting at the second execution environment a flow of user input from a trusted input device;

15 initially directing the accepted flow of user input from the second execution environment to the secured execution environment;

determining at the secured execution environment whether said secured execution environment is in a standard input mode;

20 subsequently directing the initially directed flow of user input at the secured execution environment based on the input mode of the secured execution environment including if said secured execution environment is in a standard input mode, transferring at least a first portion of said user input to said second execution environment;

25 determining at the secured execution environment from the initially directed flow of user input at such secured execution environment whether said user input comprises a user NIM indication that said secured execution environment should be in a nexus input mode; and

if said user input comprises said user NIM indication and said secured execution environment is not in said nexus input mode, switching said secured execution environment to said nexus input mode, said user NIM indication being the only way to initiate a transition from said standard input mode to said nexus input mode, there being at  
30 least two ways to transition from said nexus input mode to said standard input mode at least one of which is not a symmetrical counterpart of said user NIM indication.



2004216665 02 Feb 2010

The present invention also provides a method for maintaining the security of a secured execution environment on a system comprising said secured execution environment and a second execution environment, comprising:

- 5 accepting encrypted user input from a trusted input device and passing said encrypted user input from an input stack in the second execution environment to the secured execution environment;
- decrypting said encrypted user input;
- determining whether said secured execution environment is in a standard input mode; and
- 10 if said secured execution environment is in the standard input mode, transferring at least a first portion of the decrypted user input to said input stack in the second execution environment.

The present invention also provides a trusted user interface engine for use in a computer system comprising both a secured execution environment and a second execution environment on a single computing device, said trusted user interface engine comprising:

- 15 an input stack at the second execution environment for accepting a flow of user input from a trusted input device; and
- a trusted input manager at the secured execution environment for receiving the accepted flow of user input from the input stack and determining at the secured execution environment whether said secured execution environment is in a standard input mode; and
- 20 for subsequently directing at least a first portion of said user input to said second execution environment if said secured execution environment is in a standard input mode,
- the trusted input manager for determining at the secured execution environment from the received flow of user input at such secured execution environment whether said
- 25 user input comprises a user NIM indication that said secured execution environment should be in a nexus input mode; and if said user input comprises said user NIM indication and said secured execution environment is not in said nexus input mode, switching said secured execution environment to said nexus input mode, there being at least two ways to transition from said secured execution environment to said standard input mode at least
- 30 one of which is not a symmetrical counterpart of said user NIM indication.

The present invention also provides a computer system comprising a secured execution environment and a second execution environment, said computer system

2004216665 02 Feb 2010

comprising:

an input stack in the second execution environment for accepting encrypted user input from a trusted input device, wherein the encrypted user input is passed from the input stack in the second execution environment to the secured execution environment;

5 a trusted service provider in the secured execution environment, where the encrypted user input is decrypted;

a trusted user interface engine in the secured execution environment having a trusted input manager for determining whether said secured execution environment is in a standard input mode and for directing at least a first portion of the decrypted user input to  
10 said input stack in the second execution environment if said secured execution environment is in the standard input mode.

In embodiments of the present invention, security of a secured execution environment on a system with a secured execution environment and a second execution environment, is maintained using two modes – a standard input mode and a nexus input  
15 mode. Generally, user input from the trusted input device is transferred to the second execution environment only if the secured execution environment is in a standard input mode. User input will be generally directed to processes running in the secured execution environment only when the secured execution environment is in a nexus input mode.

Decryption of user input can be done in the secured execution environment.

20 In at least one embodiment, a transfer from the standard input mode to the nexus input mode occurs on hypothecated user input or upon programmatic activation of a process running in the secure execution environment (e.g. using a mouse pointer or keyboard commands to bring a window to focus, where that window corresponds to a process running in the secure execution environment).

25 In at least one embodiment, transfer from the nexus input mode to the standard input mode occurs on hypothecated user input, or upon exiting the last process running in the secure execution environment which was displaying a graphical user interface.

30 In at least one embodiment, even in standard input mode, user input corresponding to changes to graphical user interface elements corresponding to a process running in the secure execution environment is not directed to the second execution environment.

2004216665 02 Feb 2010

In at least one embodiment security is maintained by maintaining a current state for the secured execution environment (from at least a standard input mode state and a nexus input mode state) and directing the flow of user input according to the current state.

#### 5 **BRIEF DESCRIPTION OF THE DRAWINGS**

Preferred embodiments of the present invention are hereinafter further described, by way of example only, with reference to the accompanying drawings, wherein:

FIG. 1 is a block diagram of an exemplary computing environment in which embodiments of the invention may be implemented;

10 FIG. 2 is a block diagram of two exemplary execution environments that maintain some interaction with each other and some separation from each other;

FIG. 3 is a state diagram for input modes according to at least one embodiment of the invention;

15 FIG. 4 is a state diagram for input modes with additional transition information, according to at least one embodiment of the invention.

FIG. 5 is a block diagram of two exemplary execution environments that maintain some interaction with each other and some separation from each other according to at least one embodiment of the invention and

20 FIG. 6 is a flow diagram of a method for maintaining the security of a secured execution environment according to at least one embodiment of the invention.

#### **DETAILED DESCRIPTION OF THE INVENTION**

##### Overview

When two execution environments, such as operating systems, run side-by-side on a single machine, it must be determined how user input should be accessed by the  
 25 operating systems. Moreover, one of the operating systems may need to be protected from accessing user input which is directed at the second operating system. Embodiments of the present invention provide techniques that allow user input directed at a high assurance entity on a nexus to be protected from possible discovery by host operating system entities.

30 Reference to "aspects" of the invention hereinafter is intended to refer to embodiments or examples of the invention and to not be limiting.

2004216665 02 Feb 2010

Exemplary Computing Arrangement

FIG. 1 shows an exemplary computing environment in which aspects of the invention may be implemented. The computing system environment 100 is only one  
5 example of a suitable computing environment.

Embodiments of the invention are operational with numerous other general purpose or special purpose computing system environments or configurations. Examples of well known computing systems, environments, and/or configurations that may be suitable for  
10 use include, but are not limited to, personal computers, server computers, hand-held or laptop devices, multiprocessor systems, microprocessor-based systems, set top boxes, programmable consumer electronics, network PCs, minicomputers, mainframe computers, embedded systems, distributed computing environments that include any of the above systems or devices, and the like.

Embodiments of the invention may be described in the general context of  
15 computer-executable instructions, such as program modules, being executed by a computer. Generally, program modules include routines, programs, objects, components, data structures, etc. that perform particular tasks or implement particular abstract data types. Embodiments of the invention may also be practiced in distributed computing environments where tasks are performed by remote processing devices that are linked  
20 through a communications network or other data transmission medium. In a distributed computing environment, program modules and other data may be located in both local and remote computer storage media including memory storage devices.

With reference to FIG. 1, an exemplary system for implementing embodiments of the invention include a general purpose computing device in the form of a computer 110.  
25 Components of computer 110 may include, but are not limited to, a processing unit 120, a system memory 130, and a system bus 121 that couples various system components including the system memory to the processing unit 120. The processing unit 120 may represent multiple logical processing units such as those supported on a multi-threaded processor. The system bus 121 may be any of several types of bus structures including a  
30 memory bus or memory controller, a peripheral bus, and a local bus using any of a variety of bus architectures. By way of example, and not

2004216665 30 Sep 2004

limitation, such architectures include Industry Standard Architecture (ISA) bus, Micro Channel Architecture (MCA) bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA) local bus, and Peripheral Component Interconnect (PCI) bus (also known as Mezzanine bus). The system bus 121 may also be implemented as a point-to-point connection, switching fabric, or the like, among the communicating devices.

[0032] Computer 110 typically includes a variety of computer readable media. Computer readable media can be any available media that can be accessed by computer 110 and includes both volatile and nonvolatile media, removable and non-removable media. By way of example, and not limitation, computer readable media may comprise computer storage media and communication media. Computer storage media includes both 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, CDROM, digital versatile disks (DVD) or other optical disk storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired information and which can be accessed by computer 110. Communication media typically embodies computer readable instructions, data structures, program modules or other data in a modulated data signal such as a carrier wave or other transport mechanism and includes any information delivery media. The term "modulated data signal" means a signal that has one or more of its characteristics set or changed in such a manner as to encode information in the signal. By way of example, and not limitation, communication media includes wired media such as a wired network or direct-wired connection, and wireless media such as acoustic, RF, infrared and other wireless media. Combinations of any of the above should also be included within the scope of computer readable media.

[0033] The system memory 130 includes computer storage media in the form of volatile and/or nonvolatile memory such as read only memory (ROM) 131 and random access memory (RAM) 132. A basic input/output system 133 (BIOS), containing the basic routines that help to transfer information between elements within computer 110, such as during start-up, is typically stored in ROM 131. RAM 132 typically contains data and/or program modules that are

immediately accessible to and/or presently being operated on by processing unit 120. By way of example, and not limitation, FIG. 1 illustrates operating system 134, application programs 135, other program modules 136, and program data 137.

[0034] The computer 110 may also include other removable/non-removable, volatile/nonvolatile computer storage media. By way of example only, FIG. 1 illustrates a hard disk drive 140 that reads from or writes to non-removable, nonvolatile magnetic media, a magnetic disk drive 151 that reads from or writes to a removable, nonvolatile magnetic disk 152, and an optical disk drive 155 that reads from or writes to a removable, nonvolatile optical disk 156, such as a CD ROM or other optical media. Other removable/non-removable, volatile/nonvolatile computer storage media that can be used in the exemplary operating environment include, but are not limited to, magnetic tape cassettes, flash memory cards, digital versatile disks, digital video tape, solid state RAM, solid state ROM, and the like. The hard disk drive 141 is typically connected to the system bus 121 through a non-removable memory interface such as interface 140, and magnetic disk drive 151 and optical disk drive 155 are typically connected to the system bus 121 by a removable memory interface, such as interface 150.

[0035] The drives and their associated computer storage media discussed above and illustrated in FIG. 1, provide storage of computer readable instructions, data structures, program modules and other data for the computer 110. In FIG. 1, for example, hard disk drive 141 is illustrated as storing operating system 144, application programs 145, other program modules 146, and program data 147. Note that these components can either be the same as or different from operating system 134, application programs 135, other program modules 136, and program data 137. Operating system 144, application programs 145, other program modules 146, and program data 147 are given different numbers here to illustrate that, at a minimum, they are different copies. A user may enter commands and information into the computer 20 through input devices such as a keyboard 162 and pointing device 161, commonly referred to as a mouse, trackball or touch pad. Other input devices (not shown) may include a microphone, joystick, game pad, satellite dish, scanner, or the like. These and other input devices are often connected to the processing unit 120 through a user input interface 160 that is coupled to the system bus, but may be connected by other interface and bus structures, such as a parallel port, game port or

2004216665 02 Feb 2010

a universal serial bus (USB). A monitor 191 or other type of display device is also connected to the system bus 121 via an interface, such as a video interface 190. In addition to the monitor, computers may also include other peripheral output devices such as speakers 197 and printer 196, which may be connected through an output peripheral interface 190.

The computer 110 may operate in a networked environment using logical connections to one or more remote computers, such as a remote computer 180. The remote computer 180 may be a personal computer, a server, a router, a network PC, a peer device or other common network node, and typically includes many or all of the elements described above relative to the computer 110, although only a memory storage device 181 has been illustrated in FIG. 1. The logical connections depicted in FIG. 1 include a local area network (LAN) 171 and a wide area network (WAN) 173, but may also include other networks. Such networking environments are commonplace in offices, enterprise-wide computer networks, intranets and the Internet.

When used in a LAN networking environment, the computer 110 is connected to the LAN 171 through a network interface or adapter 170. When used in a WAN networking environment, the computer 110 typically includes a modem 172 or other means for establishing communications over the WAN 173, such as the Internet. The modem 172, which may be internal or external, may be connected to the system bus 121 via the user input interface 160, or other appropriate mechanism. In a networked environment, program modules depicted relative to the computer 110, or portions thereof, may be stored in the remote memory storage device. By way of example, and not limitation, FIG. 1 illustrates remote application programs 185 as residing on memory device 181. It will be appreciated that the network connections shown are exemplary and other means of establishing a communications link between the computers may be used.

#### Plural Computing Environments on a Single Machine

For two operating systems executing side-by-side on a single computing device, embodiments of the present invention can be used to provide at least some level of separation between the two operating systems, while still providing for at least some level of interaction between the two operating systems.

FIG. 2 shows a system in which two operating systems 134(1) and 134(2) execute on a single computer 110. Some type of logical separation 202 exists between operating

2004216665 02 Feb 2010

systems 134(1) and 134(2), such that a certain amount of interaction 204 is permitted between operating systems 134(1) and 134(2), while still allowing at least one of the operating systems to be protected against events that originate in the other operating system. In the example of FIG. 2, operating system 134(1) is a host operating system, and  
5 operating system 134(2) is a guest operating system, such as a "nexus" as described above. When operating system 134(2) is a nexus, it is desirable to construct separation 202 such that operating system 134(2) can interact with operating system 134(1) in order to borrow operating system 134(1)'s infrastructure, while still allowing operating system 134(2) to protect itself from actions (either malicious or innocent) that arise at operating system  
10 134(1) and might cause operating system 134(2) to behave in a manner contrary to its behavioral specifications.

The separation 202 between operating systems 134(1) and 134(2) may, optionally, be enforced with the aid of a security monitor. A security monitor is a component external to both operating systems 134(1) and 134(2), which provides some security services that  
15 may be used to protect operating system 134(2) from operating system 134(1). For example, a security monitor may control access to certain hardware, may manage the use of memory (to give operating system 134(2) exclusive use of some portions of memory), or may facilitate the communication of data from operating system 134(1) to operating system 134(2) in a secure way. It should be noted that the use of a security monitor  
20 represents one model of how operating system 134(2) can be protected from operating system 134(1). As another example, operating system 134(2) could include all of the functionality necessary to protect itself from operating system 134(1).

It should be noted that FIG. 2 shows operating system 134(1) as a "host" and operating system 134(2) as a "guest." In general, this characterization refers to the fact  
25 that, in these examples, operating system 134(1) provides certain operating system infrastructure that is used by both operating systems 134(1) and 134(2) (e.g., device drivers, scheduling, etc.), and operating system 134(2) is a "guest" in the sense that it preferably lacks this infrastructure but rather uses the infrastructure of operating system 134(1). However, it should be noted that the



parameters of what makes an operating system a "host" or a "guest" are flexible. Moreover, it should be noted that traditional concepts of "host" and "guest" operating systems presume that the host needs to protect itself from actions of the guest. In the example of FIGS. 2, however, guest operating system 134(2) is presumed to be a high-assurance operating system that needs to protect itself from host operating system 134(1). In the examples that follow, we shall generally refer to operating system 134(1) as the "host" and operating system 134(2) as the "guest" or "nexus" for the purpose of distinguishing between them. It should be appreciated that the techniques described herein can be applied to the interaction of any two or more operating systems running on the same machine (or even on the same set of connected machines).

#### User Input For Plural Computing Environments on a Single Machine

[0042] When a user interacts with programs on a computer system containing a high-assurance operating system, the user does so by means of a user input device, such as mouse 161 or keyboard 162 (from Figure 1). As discussed above, allowing the windowing system running on host operating system 134(1) control the destination of the stream of input events may allow an attack using a compromised host operating system or application. Thus the destination of the stream of input events is controlled by a trusted user interface engine running on the nexus 134(2). Two modes are established in which the computer 110 can be operating. These two modes are standard input mode (SIM) and nexus input mode (NIM).

[0043] Details will be described below, however, generally, nexus input mode is used when user input is to be directed towards a nexus-mode process, and standard input mode is used when user input is not to be directed towards a nexus-mode process. Thus, these modes are used to provide a level of assurance that the user input will not be misdirected to an entity connected with host operating system 134(1) when it should be directed toward an entity connected with nexus 134(2). Avoiding this misdirection allows increased security for nexus-side data and processes. Because the use of the SIM and NIM modes militate against attacks on the nexus 134(2), control over these modes must be protected. Therefore, the modes are controlled by the nexus 134(2) or, where a security monitor is present in the computer system 110, it may be controlled by the security monitor.

[0044] In standard input mode, generally, user input events are unencrypted in the nexus and then passed by the nexus to a standard-mode input pipeline. This standard-mode input pipeline is the input pipeline which the host operating system 134(1) draws user input events from for its own use and for the use of applications running under the host operating system 134(1). Certain user input events, however, trigger a switch from the standard input mode to the nexus input mode. Therefore, before input events are passed to the standard-mode input pipeline, they are examined by a process running on the nexus 134(2) in order to determine whether they are events which should trigger a switch to the nexus input mode.

[0045] In addition to events which trigger a switch to the nexus input mode, some mouse or other user input events may be retained and acted on by the nexus rather than passed to the host side, even in SIM. For example, where a windowing system is being used, in standard input mode only a window owned by a non-nexus entity can have focus. In one embodiment, where a windowing system is used, the mouse is used to control the movement of a cursor over the display and the nexus controls the tracking and display movement of the mouse over windows controlled by a nexus-side application or other entity in both NIM and SIM. In such a case, where input events are mouse movements, in SIM input events which are passed to the host side from the nexus do not include mouse movements over areas with a nexus user interface element.

[0046] In the nexus input mode, user input events are processed within nexus space – that is, within the portion of computer system 110 which contains the nexus 134(2) and is separated from other portions of the computer system 110 by logical separation 202. During this mode, the standard-mode input pipeline receives no user input events. Input events are passed to the destination entities in the nexus space, or are used by the nexus 134(2) for operating system functions and also to determine whether a switch to standard input mode should occur.

[0047] In one embodiment, there are some exceptions which, even in NIM, are passed to the host 134(1). For example, commands regarding the indicator lights on a keyboard, volume control commands, and screen contrast and brightness commands may be passed to the host 134(1) in one embodiment. In an alternate embodiment, however, these commands are simply ignored when the computer system is in NIM.

[0048] Thus, the input mode of the computer system 110 can be described by the state diagram shown in Figure 3. Figure 3 shows a state diagram for input modes according to one embodiment of the invention. The standard input mode state 300 represents the standard input mode, in which events are passed to the standard-mode input pipeline, after being examined to ensure they are not events which trigger a transition to the nexus input mode state 310. The nexus input mode state 310 represents the nexus input mode, in which the standard input mode receives no user input events.

#### Transitioning Between Input Modes

[0049] The SIM to NIM transition 302 occurs, in one embodiment, when a user activates a nexus user interface element. In the windowing context, this may occur when the user selects or clicks on a window which is controlled by an entity in nexus space. This indicates that the user's focus has moved to that window. In one embodiment, for each user interface element controlled by an entity in nexus space, there is a corresponding "shadow" which is maintained on the host side. When this shadow user interface element is selected (e.g. by clicking on the title bar of the shadow user interface element) a transition is made to NIM and the focus is switched to the corresponding user interface element owned by a nexus-side entity.

[0050] The SIM to NIM transition may also occur when a hypothecated user interaction occurs. A hypothecated user interaction is a user interaction which, in the context of the computer system 110, will always result in a specific consequence. Therefore, there may be a hypothecated user interaction for which the consequence is that the computer system 110 is in the nexus input mode.

[0051] These two ways of transitioning from SIM state 300 to NIM state 310 according to one embodiment are more fully described in Figure 4. Figure 4 shows a state diagram for input modes with more specific transition information, according to one embodiment of the invention. Figure 4 shows that when the user focuses on a nexus space entity 402, the transition is made from SIM to NIM. This is one kind of hypothecated action. Additionally, when the user performs any other NIM hypothecated action 404, that is, an action which is intended to always result in the computer being in NIM state, the transition is made from SIM to NIM. It may also be the case that a user can perform a NIM hypothecated action while in NIM state 310. If so, as

shown by NIM hypothesized action 406, the state remains in NIM state 310. A NIM hypothesized action may be, for example, a set or combination of keystrokes. For example, where a user presses the "control" key, the "alt" key, and the "tab" key, all at the same time, this series of keystrokes provides the guarantee that the computer system 110 will immediately switch to NIM state 310.

[0052] In one embodiment, where such a transition occurs as a result of a combination of keystrokes, the host side will receive mock-up user input data which resolves any keystroke information which it had received. Therefore, using the previous <CTRL><ALT><TAB> example, if the host side had received information that the control key and the alt key had been pressed, when the tab key is pressed, the transition to NIM state 310 occurs, and the host side receives mock-up user input which indicates that the control key and the alt key had been released. Thus, the host side input tracking is flushed or reset when the transition to NIM is made. Similarly, when a transition to SIM is made, the host side receives information about the state of the input devices, e.g. which keys are currently depressed on the keyboard, so that it can respond properly to further changes which arrive from the input devices.

[0053] With reference again to Figure 3, a NIM to SIM transaction 312 is also provided for. In one embodiment, as shown in Figure 4, in order to provide high assurance for nexus-space entities, this transition only occurs when the user specifically requests this switch, by performing a SIM hypothesized action 412. It can be seen that there are two ways to transition from NIM state 310 to SIM state 312, but only one way to transition back. In one embodiment, no transition exists from NIM state 310 to SIM state 312 which is "symmetrical" the transition 402 from SIM state 312 to NIM state 310 which occurs when the user focuses on a nexus space entity. More generally, with reference again to Figure 3, in one embodiment, transitions from SIM to NIM 310 do not occur in exactly symmetrical ways to the ways in which transitions from NIM to SIM 312 occur. This is to provide increased security by assuring that a user of the computer system 110 must cause any transition to standard input mode knowingly and not accidentally.

[0054] In other embodiments, there are alternate ways to switch from SIM to NIM. For example, if a user mode process executing in the nexus (a "nexus agent") reaches a point at which it requires user input, that causes a transition from SIM to NIM. This allows some

additional security. For example, if a nexus agent is displaying confidential information on the monitor 191 (from Fig. 1) it may require a password, at intervals, in order to maintain the confidential information on the monitor. If the nexus agent can force a transition to NIM and request this data, when it does not receive the password, it can remove the confidential information from the monitor 191 and perform any other actions that may be required. Another event which may trigger a transition to NIM mode is the receipt of a power management event. Such events may include powering up, powering down, switching to another form of power (battery, etc.), hibernation, etc.

[0055] In other embodiments, there are alternate ways to switch from NIM to SIM. For example, if a windowing system is being used, and the last window owned by a nexus agent is closed by the user, this times causes a transition from NIM to SIM. Because NIM is used only to allow input to nexus agents, in this embodiment, once all windows associated with any nexus agents are closed, there is no need for the computer system 110 to be in NIM, and it switches to SIM.

#### Path of Trusted Input In a Plural Computing Environment

[0056] In one embodiment, some or all of the user input devices such as mouse 161 or keyboard 162 (from Figure 1) may support encryption. Such encrypted devices allow the input device and the nexus to establish a secure connection and can thus prevent a hardware attack. In one embodiment, the NIM and SIM modes apply only to input from an encrypted device. Other input devices may not be encrypted, however, input from unencrypted user input devices is only accepted for the host system, and not for the nexus or any entities running on the nexus.

[0057] Figure 5 is a block diagram of two exemplary execution environments that maintain some interaction with each other and some separation from each other according to one embodiment of the invention. As shown in Figure 5, a trusted input device 500 sends encrypted user input data to computer system 110. The encrypted data arrives in the input stack 505 in the host-side operating system 134(1). Because it is encrypted, it can not be used by the host-side operating system 134(1) but is passed to the nexus 134(2). In Figure 5, both the host side (to the left of logical separation 202) and the nexus side (to the right of logical separation 202) are divided further by into user level 510 and kernel level 520. These correspond to distinctions

between user level and kernel level execution in operating systems. When the nexus 134(2) receives the encrypted user input data, it is passed to the input trusted service provider (TSP) 525, where the encrypted user input data is decrypted.

[0058] Decrypted data is then passed to the trusted user interface engine (TUE) 530, specifically to the trusted input manager 540. Trusted input manager 540, in conjunction with the trusted window manager (TWinMgr) 545 tracks the state of the input mode. Trusted input manager 540 and TWinMgr 545 also track whether a transition event is indicated by the user input data and changes the user input mode accordingly.

[0059] In the embodiment described above where shadow user interface elements (e.g. shadow windows) on the host side are created for each nexus-side user interface element, when a focus event is received on a shadow user interface element, notification of the event is received by TWinMgr 545, and this causes a switch to NIM. The TWinMgr 545 functions as a window manager for nexus-side user interface elements. Thus, a nexus-side window which requests focus will signal the TWinMgr 545 and cause a switch to NIM.

[0060] If the input mode is SIM, most user input data (as described above) is restored to input stack 505 for use on the host side, as shown by arrow C 590. One exception is data concerning mouse movements over nexus-side user interface elements. The display of nexus-side user interface elements is handled by the trusted output manager (TOM) 550 and so mouse movements over nexus-side user interface elements are passed to TOM 550, as shown by arrow B 580. Other exceptions include trigger events, and events as described above, which may be passed elsewhere in the nexus side.

[0061] If the input mode is NIM, the data will be passed from the trusted input manager to the correct destination nexus-mode process. This is illustrated by arrow A 570, which shows user input data being sent to nexus agent 575. Other data may be passed to other entities on the nexus side. As described above, in one embodiment some data, for example, volume control data, may be passed to the host side.

[0062] Figure 6 is a flow diagram of a method for maintaining the security of a secured execution environment according to one embodiment of the invention. As shown in Figure 6, in step 600, user input is accepted from a trusted input device. This may be direct, or take an indirect path, for example, the path shown in Figure 5. In step 610, a determination is made

2004216665 02 Feb 2010

whether the secured execution environment is in a standard input mode. If it is, step 620, at least a first portion of the user input is transferred to the second execution environment.

Conclusion

5 It is noted that the foregoing examples have been provided merely for the purpose of explanation of the embodiments of the present invention. The invention is not intended to be limited to the particulars disclosed herein; rather, the invention extends to all functionally equivalent structures, methods and uses, such as are within the scope of the appended claims. Those skilled in the art, having the benefit of the teachings of this  
10 specification, may effect numerous modifications thereto and changes may be made without departing from the scope and spirit of the invention.

Throughout this specification and the claims which follow, unless the context requires otherwise, the word "comprise", and variations such as "comprises" and "comprising", will be understood to imply the inclusion of a stated integer or step or group  
15 of integers or steps but not the exclusion of any other integer or step or group of integers or steps.

The reference to any prior art in this specification is not, and should not be taken as, an acknowledgement or any form of suggestion that the prior art forms part of the common general knowledge in Australia.

2004216665 02 Feb 2010

THE CLAIMS DEFINING THE INVENTION ARE AS FOLLOWS:

1. A method for maintaining the security of a secured execution environment on a system comprising said secured execution environment and a second execution environment both on a single computing device, the method comprising:
- 5 accepting at the second execution environment a flow of user input from a trusted input device;
- initially directing the accepted flow of user input from the second execution environment to the secured execution environment;
- 10 determining at the secured execution environment whether said secured execution environment is in a standard input mode;
- subsequently directing the initially directed flow of user input at the secured execution environment based on the input mode of the secured execution environment including if said secured execution environment is in a standard input mode, transferring at
- 15 least a first portion of said user input to said second execution environment;
- determining at the secured execution environment from the initially directed flow of user input at such secured execution environment whether said user input comprises a user NIM indication that said secured execution environment should be in a nexus input mode; and
- 20 if said user input comprises said user NIM indication and said secured execution environment is not in said nexus input mode, switching said secured execution environment to said nexus input mode, said user NIM indication being the only way to initiate a transition from said standard input mode to said nexus input mode, there being at least two ways to transition from said nexus input mode to said standard input mode at
- 25 least one of which is not a symmetrical counterpart of said user NIM indication.
2. The method of claim 1, further comprising:
- decrypting said user input.
- 30 3. The method of one of claims 1 or 2, further comprising:
- if said secured execution environment is in a nexus input mode, determining a specific process running in said secured execution environment to which said user input is



2004216665 02 Feb 2010

directed; and  
directing said user input to said specific process.

4. A method for maintaining the security of a secured execution environment  
5 on a system comprising said secured execution environment and a second execution  
environment, comprising:  
accepting encrypted user input from a trusted input device and passing said  
encrypted user input from an input stack in the second execution environment to the  
secured execution environment;  
10 decrypting said encrypted user input;  
determining whether said secured execution environment is in a standard input  
mode; and  
if said secured execution environment is in the standard input mode, transferring at  
least a first portion of the decrypted user input to said input stack in the second execution  
15 environment.

5. The method of claim 4, further comprising:  
if said secured execution environment is in a nexus input mode (NIM), wherein the  
nexus input mode is a high assurance input mode, determining a specific process running  
20 in said secured execution environment to which the decrypted user input is directed; and  
directing the decrypted user input to said specific process.

6. The method of claim 4, further comprising:  
determining whether the decrypted user input comprises a user nexus input mode  
25 (NIM) indication that said secured execution environment should be in a nexus input  
mode; and  
if the decrypted user input comprises said user nexus input mode indication and  
said secured execution environment is not in said nexus input mode, switching said  
secured execution environment to said nexus input mode.

30  
7. The method of any one of claims 1 to 3 and 6, where said user NIM  
indication comprises a combination of keystrokes on a keyboard.

2004216665 02 Feb 2010

8. The method of any one of claims 1 to 3 and 6, where said user NIM indication comprises a programmatic activation of a process running in said secured execution environment.

5

9. The method of claim 8, where said programmatic activation of a first process running in said secured execution environment comprises selecting a graphical user interface element corresponding to said process.

10

10. The method of claim 9, where said graphical user interface element is a shadow graphical user interface element displayed using a second process, where said process is running on said second execution environment, and where said shadow graphical user interface element corresponds to a secured graphical user interface element displayed by said first process.

15

11. The method of claim 9, where said graphical user interface element is displayed using a second process, where said second process is running on said second execution environment, and where said graphical user interface element corresponds to a secured graphical user interface element displayed by said first process.

20

12. The method of any one of claims 1 to 11, further comprising:  
determining whether said user input comprises a user standard input mode (SIM) indication that said secured execution environment should be in said standard input mode;  
and

25

if said user input comprises said user SIM indication and said secured execution environment is not in said standard input mode, switching said secured execution environment to said standard input mode.

30

13. The method of claim 12, where said user SIM indication comprises a combination of keystrokes on a keyboard.

2004216665 02 Feb 2010

14. The method of claim 12, where said user SIM indication comprises an action which results in a display with no graphical user interface element which corresponds to a process running on said secured execution environment.

5 15. The method of claim 12, where said user standard input mode indication comprises an action which results in closing the last graphical user interface element which corresponds to a process running on said secured execution environment.

10 16. The method of any one of claims 1 to 15, where if said secured execution environment is in a standard input mode, and a second portion of said user input corresponds to changes to a graphical user interface element displayed by a process running on said secured execution environment, said changes to said graphical user interface element are performed within said secured execution environment.

15 17. The method of claim 16, where said changes to a graphical user interface element displayed by a process running on said secured execution environment comprise the movement of a mouse cursor over a graphical user interface element displayed by a process running on said secured execution environment.

20 18. The method of any one of claims 1 to 17, further comprising:  
switching said secured execution environment to a nexus input mode if a power management change is detected.

25 19. The method of any one of claims 1 to 18, further comprising:  
maintaining a current state for said secured execution environment selected from among a group of possible states comprising: a standard input mode state and a nexus input mode state; and  
directing a flow of user input according to said current state.

30 20. The method of claim 19, further comprising:  
limiting a transfer of said user input to said second execution environment when said current state is said nexus input mode state.

2004216665 02 Feb 2010

21. A method for maintaining the security of a secured execution environment substantially as hereinbefore described, with reference to the accompanying drawings.

5 22. A computer-readable medium containing computer executable instructions to maintain the security of a secured execution environment on a system comprising said secured execution environment and a second execution environment, the computer-executable instructions to perform the method of any one of claims 1 to 21.

10 23. A trusted user interface engine for use in a computer system comprising both a secured execution environment and a second execution environment on a single computing device, said trusted user interface engine comprising:

an input stack at the second execution environment for accepting a flow of user input from a trusted input device; and

15 a trusted input manager at the secured execution environment for receiving the accepted flow of user input from the input stack and determining at the secured execution environment whether said secured execution environment is in a standard input mode; and for subsequently directing at least a first portion of said user input to said second execution environment if said secured execution environment is in a standard input mode,

20 the trusted input manager for determining at the secured execution environment from the received flow of user input at such secured execution environment whether said user input comprises a user NIM indication that said secured execution environment should be in a nexus input mode; and if said user input comprises said user NIM indication and said secured execution environment is not in said nexus input mode, switching said  
25 secured execution environment to said nexus input mode, there being at least two ways to transition from said secured execution environment to said standard input mode at least one of which is not a symmetrical counterpart of said user NIM indication.

24. The trusted user interface engine of claim 23, where said trusted input  
30 manager, if said secured execution environment is in a nexus input mode, determines a specific process running in said secured execution environment to which said user input is directed; and directs said user input to said specific process.

2004216665 02 Feb 2010

25. A computer system comprising a secured execution environment and a second execution environment, said computer system comprising:

- an input stack in the second execution environment for accepting encrypted user
- 5 input from a trusted input device, wherein the encrypted user input is passed from the input stack in the second execution environment to the secured execution environment;
- a trusted service provider in the secured execution environment, where the encrypted user input is decrypted;
- a trusted user interface engine in the secured execution environment having a
- 10 trusted input manager for determining whether said secured execution environment is in a standard input mode and for directing at least a first portion of the decrypted user input to said input stack in the second execution environment if said secured execution environment is in the standard input mode.

15 26. The computer system of claim 25, where said trusted input manager, if said secured execution environment is in a nexus input mode (NIM), wherein the nexus input mode is a high assurance input mode, determines a specific process running in said secured execution environment to which the decrypted user input is directed; and directs the decrypted user input to said specific process.

20 27. The computer system of claim 25, where said trusted input manager determines whether the decrypted user input comprises a user nexus input mode (NIM) indication that said secured execution environment should be in a nexus input mode; and if the decrypted user input comprises said user nexus input mode indication and said secured

25 execution environment is not in said nexus input mode, switching said secured execution environment to said nexus input mode.

28. The trusted user interface engine of any one of claims 23 to 24, and 27, where said user NIM indication comprises a combination of keystrokes on a keyboard.

30

2004216665 02 Feb 2010

29. The trusted user interface engine of any one of claims 23 to 24, and 27, where said user NIM indication comprises a programmatic activation of a process running in said secured execution environment.

5 30. The trusted user interface engine of claim 29, where said programmatic activation of a first process running in said secured execution environment comprises selecting a graphical user interface element corresponding to said process.

10 31. The trusted user interface engine of claim 30, where said graphical user interface element is a shadow graphical user interface element displayed using a second process, where said process is running on said second execution environment, and where said shadow graphical user interface element corresponds to a secured graphical user interface element displayed by said first process.

15 32. The trusted user interface engine of claim 30, where said graphical user interface element is displayed using a second process, where said second process is running on said second execution environment, and where said graphical user interface element corresponds to a secured graphical user interface element displayed by said first process.

20 33. The trusted user interface engine of any one of claims 23 to 27, where said trusted input manager determines whether said user input comprises a user (standard input mode) SIM indication that said secured execution environment should be in said standard input mode; and if said user input comprises said user SIM indication and said secured execution environment is not in said standard input mode, switches said secured execution environment to said standard input mode.

25 34. The trusted user interface engine of claim 33, where said user SIM indication comprises a combination of keystrokes on a keyboard.

30

2004216665 02 Feb 2010

35. The trusted user interface engine of claim 33, where said user SIM indication comprises an action which results in a display with no graphical user interface element which corresponds to a process running on said secured execution environment.

5 36. The trusted user interface engine of claim 33, where said user standard input mode indication comprises an action which results in closing the last graphical user interface element which corresponds to a process running on said secured execution environment.

10 37. The trusted user interface engine of any one of claims 23 to 27, where if said secured execution environment is in a standard input mode, and a second portion of said user input corresponds to changes to a graphical user interface element displayed by a process running on said secured execution environment, said changes to said graphical user interface element are performed within said secured execution environment.

15 38. The trusted user interface engine of claim 37, where said changes to a graphical user interface element displayed by a process running on said secured execution environment comprise the movement of a mouse cursor over a graphical user interface element displayed by a process running on said secured execution environment.

20 39. The trusted user interface engine of claim 23 to 27, where said trusted input manager switches said secured execution environment to a nexus input mode if a power management change is detected.

25 40. A computer system substantially as hereinbefore described, with reference to the drawings and/or examples.

41. A trusted user interface engine substantially as hereinbefore described, with reference to the drawings and/or examples.

Computing Environment 100

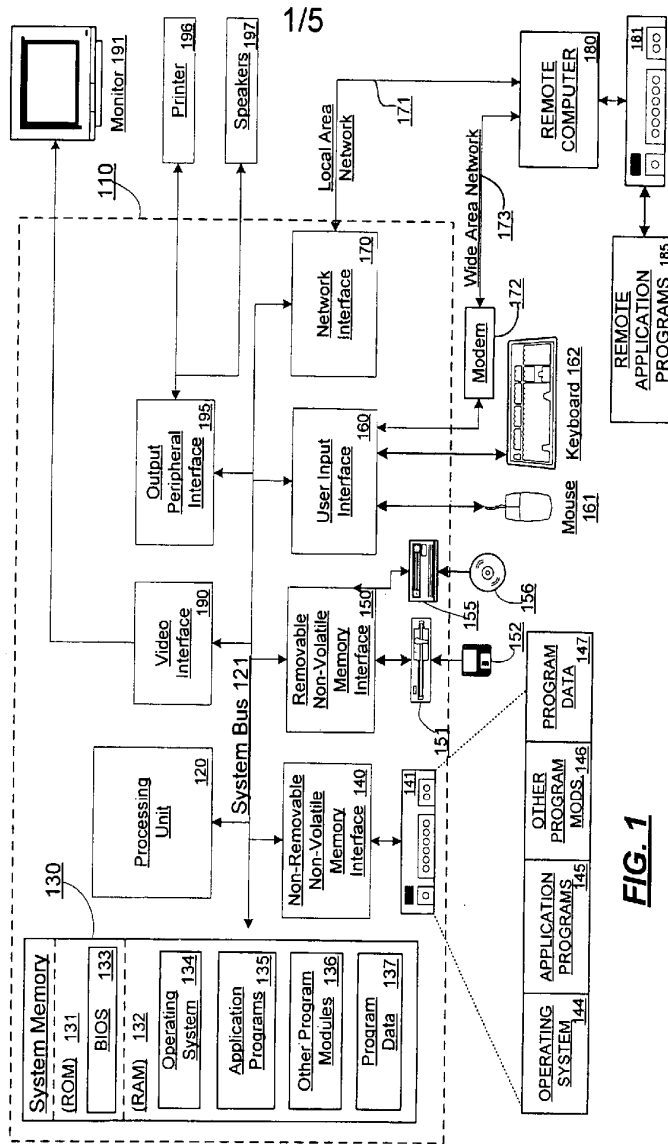
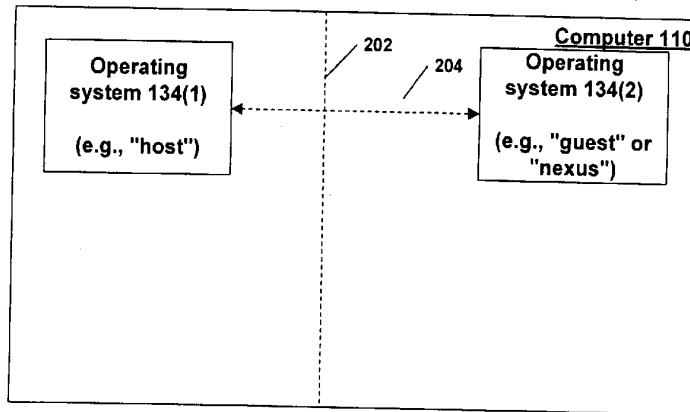
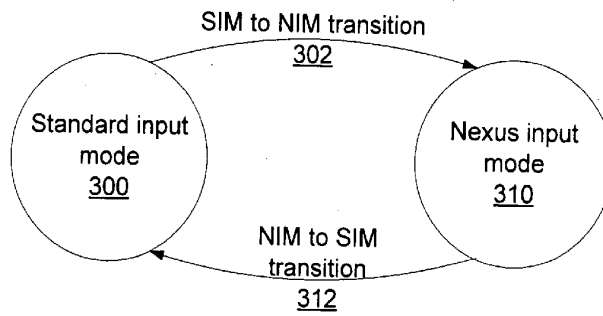


FIG. 1

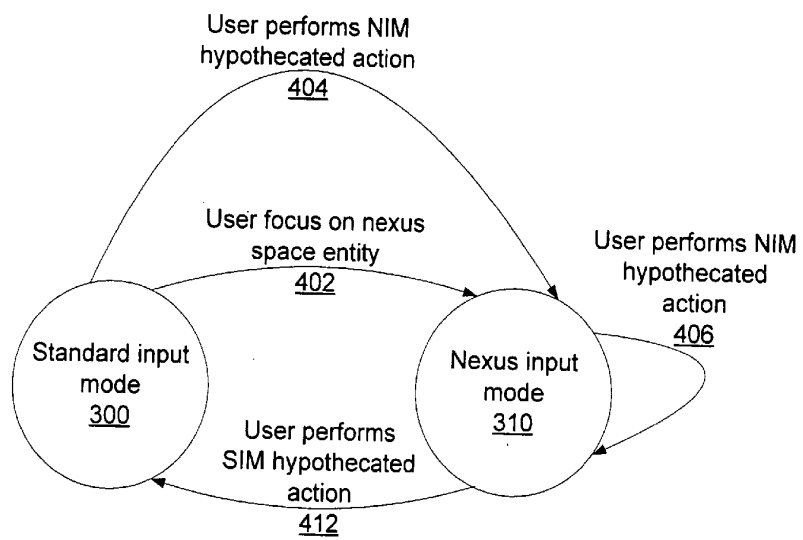




**FIG. 2**



**FIG. 3**



**FIG. 4**

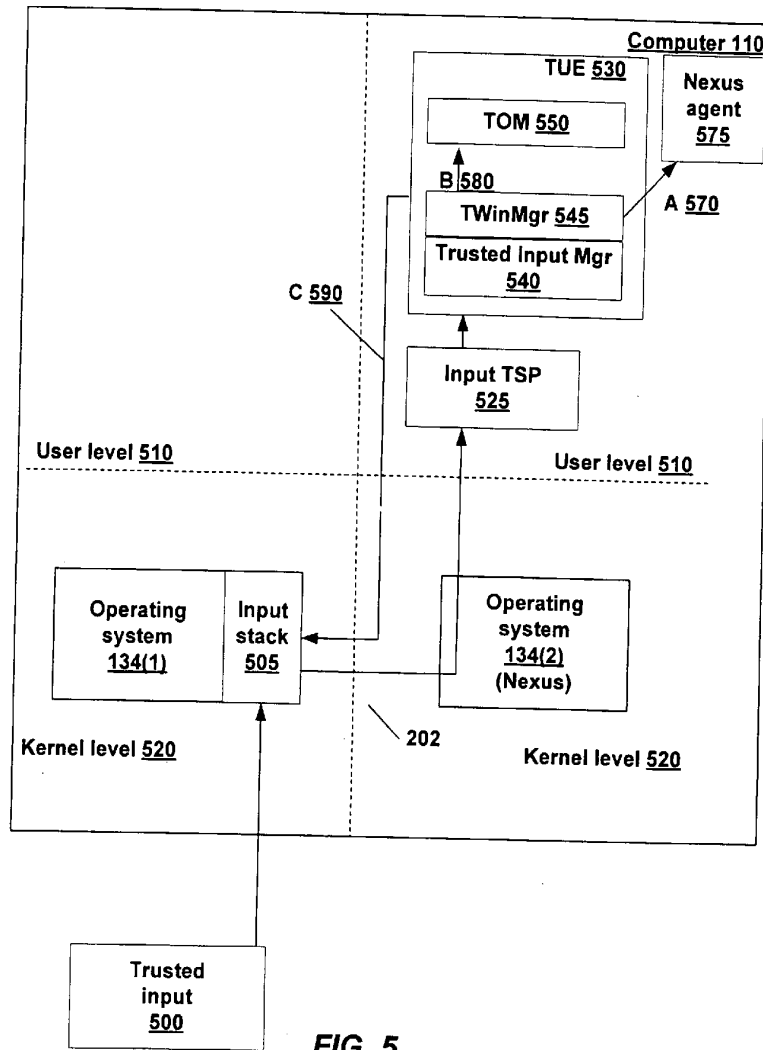
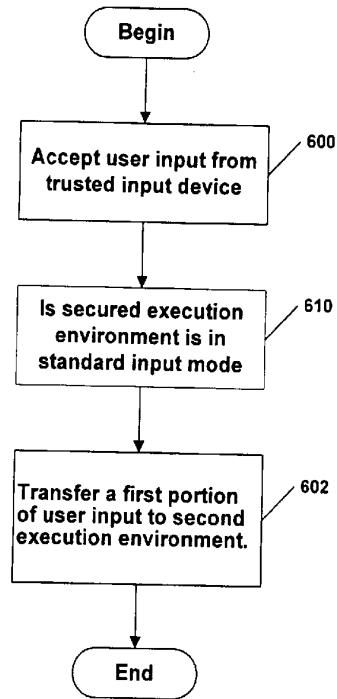


FIG. 5

5/5



**FIG. 6**