




- (51) **International Patent Classification:**
G06F 21/00 (2013.01)
- (21) **International Application Number:**
PCT/US2012/055210
- (22) **International Filing Date:**
13 September 2012 (13.09.2012)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
61/534,220 13 September 2011 (13.09.2011) US
- (71) **Applicant (for all designated States except US):**
PRIVATECORE, INC. [US/US]; 555 Bryant St. #821,
Palo Alto, California 94301 (US).
- (72) **Inventors; and**
- (75) **Inventors/Applicants (for US only):** **HOROVITZ, Oded**
[IL/US]; 1444 Pitnam Avenue, Palo Alto, California 94301
(US). **WEIS, Stephen A.** [US/US]; 280 Fell Street #308,
San Francisco, California 94102 (US). **WALDSPURGER,**
Carl A. [US/US]; 517 Georgia Avenue, Palo Alto, Califor-
nia 94306 (US). **RIIHAN, Sahil** [IN/US]; 103 Chester
Street, Menlo Park, California 94025 (US).
- (74) **Agents:** **PEARCE, Jeffrey** et al.; P.O. Box 51, Redmond,
Washington 98073 (US).
- (81) **Designated States (unless otherwise indicated, for every
kind of national protection available):** AE, AG, AL, AM,
AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY,
BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM,
DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT,
HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP,
KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD,
ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI,
NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU,
RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ,
TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA,
ZM, ZW.
- (84) **Designated States (unless otherwise indicated, for every
kind of regional protection available):** ARIPO (BW, GH,
GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ,
UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ,
TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK,
EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV,
MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM,
TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW,
ML, MR, NE, SN, TD, TG).

[Continued on next page]

(54) **Title:** SOFTWARE CRYPTOPROCESSOR

(57) **Abstract:** Security of information - both code and data - stored in a computer's system memory is provided by an agent (5100) loaded into and at run time resident in a CPU cache (5000). Memory writes from the CPU (1000) are encrypted by the agent before writing and reads into the CPU are decrypted by the agent before they reach the CPU. The cache-resident agent also optionally validates the encrypted information stored in the system memory. Support for I/O devices (6700) and cache protection from unsafe DMA of the cache by devices is also provided.



Published:

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

— before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))

SOFTWARE CRYPTOPROCESSOR

FIELD OF THE INVENTION

[0001] This invention relates to computer security, in particular, to establishing trust in the integrity and privacy of software and hardware components.

BACKGROUND

[0002] Since the beginning of electronic computing, there has been a never-ending concern that the hardware and/or software comprising the computers – to say nothing of user and other data – has been deliberately or inadvertently corrupted. As the complexity of computers increases, so does this concern. At the heart of the problem is trust, or the lack thereof. Software, hypervisors (in virtualized computer systems), operating systems (OS) and applications and the like that run on modern computers, servers, mobile devices, etc., assume, for example, that the hardware on which they are running can be trusted – when software stores information in memory, it expects the information to be the same the next time it is read from memory and not tampered with or leaked to an attacker.

[0003] With an increase of platform modularization, it is possible for a human operator or a hardware component to modify the state of the software out-of-band of its execution. For example, every device that is connected to main memory through an I/O bus (for example, a PCI bus) can communicate with the software by means of writing and reading not only to shared areas but also from the application private state that is stored in the main memory. For these and many other reasons, the problem of trust is more prevalent today than ever before, especially as more and more organizations rely on service providers or manufacturers to run their software and build the systems on which they run.

[0004] An attacker can successfully penetrate a software system just by observing, without modifying, the state of the application. One example of such an attack consists of just observing data stored in memory (for example, credit card numbers, personally identifying information, etc.) without requiring access

credentials such as passwords. Another example is known as a “cold-boot” attack, which exploits the physical property of system memory (in particular, DRAM) that its contents are retained for a short time even when the power is removed. In such attacks, a human can force a reboot of the system, often just by causing a loss of power or reset, and then, while data is still latent, reboot into a small (even USB-based) OS that then can observe or copy the data from the supposedly secure hardware; alternatively, the user can quickly remove the main memory device from one host and mount it on another host in order to read its content. The cold-boot attack is explained, for example, in <http://citp.princeton.edu/memory/>. By observing the state of the software, that is, the contents of memory, the attacker can then obtain secrets and credentials (for example, keys) that later allow the attacker entry through the front door of the software stack.

[0005] Consequently, a whole branch of computer science is dedicated to finding solutions to the problem of trust. Some of the concepts and attempted solutions include the following.

[0006] “Data-at-rest” encryption involves encrypting data at the storage back-end, that is, data that is not, at a given moment, being transmitted between devices or otherwise processed. Although this method provides a relative easy and straightforward way to encrypt substantially the entire non-volatile storage, it does not protect data that is currently being processed or used. Yet another drawback is that any keys used to encrypt the “at-rest” data must be stored and persist as long as the data; keys themselves are subject to theft while in memory, which would defeat the encryption as a whole.

[0007] “Cryptoprocessors” have been in use since the early days of computers. In broad terms, a conventional, modern cryptoprocessor is some form of microprocessor (including smartcards) that includes built-in hardware support for cryptographic operations. In short, existing cryptoprocessor-based security solutions rely on dedicated or specialized hardware and typically also on modifications to the processor circuitry itself. One variant of this concept is “in-line processor RAM encryption”, which is a hardware-based approach to the problem in which the CPU integrates encryption logic into the cache circuitry, and which enables

encryption/decryption whenever cache data is either evicted or filled. Although commodity operating systems can run on some known cryptoprocessors, at present there is no cryptoprocessor implementation, in the sense of full-memory encryption, for an x86 architecture.

[0008] The Trusted Computing Group TCG (<http://www.trustedcomputinggroup.org>) works on a solution that can attest to the software that has started to execute on a processor after a system boot, or in the case of Intel's TXT (Trusted Execution Technology) (<http://download.intel.com/technology/security/downloads/315168.pdf>), whenever a software component decides to execute in a measured environment. Once the system has started to run the measured code, however, the TCG framework no longer provides additional run-time protection.

[0009] "Messaging encryption" is a system concept for secure communication (for example, VPN, IPsec, SSL) that allows hosts or other software entities to communicate privately over public networks. Note that this is a form of encryption for "data-in-motion". Although this method provides a level of security for data in transit between two entities, it is not designed or suitable for protecting data being otherwise processed before or after transit.

[0010] Various researchers have started to experiment with using private storage in the CPU to hide data, in particular the x86 internal state, from the main memory. One such example is TRESOR (http://www.usenix.org/events/sec11/tech/full_papers/muller.pdf), which contains references to "frozen cache". Research in this area has demonstrated the ability to protect some data from reaching RAM, but not the entire software stack. This solution therefore has the weakness that an attacker can easily modify other data that is exposed to the RAM and modify its software component to expose all secrets.

[0011] "Cache management via page coloring" (http://en.Wikipedia.org/wiki/cache_coloring) is a method to split the cache content between multiple applications running on a single x86 processor. This method divides the total set of physical pages into pages that are "known" not to conflict "as

far as cache placement is concerned", assuming that if each application allocates only from one pool of pages having color A, it will never contend with another application that allocates from another pool of pages having color B. In this way, the application using color A cannot cause evictions of cached data associated with application using color B, and vice versa, but it does not provide any guarantee as to which pages in a given pool are currently in the cache or in system memory.

BRIEF DESCRIPTION OF THE DRAWINGS

[0012] Figure 1 illustrates the main hardware and software components of the invention, as well as the main principles of operation.

[0013] Figure 2 illustrates an example of an embodiment in which a security/validation component may be part of system software wholly or partially in a CPU cache.

[0014] Figure 3 illustrates the interaction with system memory by various hardware and software components in typical implementations.

DETAILED DESCRIPTION

[0015] Although different implementations of the invention may operate with more relaxed requirements and assumptions, in broadest terms, this invention provides an "agent" (either software or "firmware") contained within the CPU cache that encrypts/decrypts and/or validates/authenticates data and/or instructions that are passed to or from the processor core(s). As such, the invention provides a cryptoprocessor that is embodied as software – requiring no specific hardware support or modifications – onboard the main processor itself. This cached software module/agent thus forms a software "wall" around the CPU core, which makes available instructions and data coming in to the core in the proper, unencrypted form for processing, but that encrypts all out-going data before the CPU stores it in any external modules, including system memory, device memory, etc. Thanks to this, with at most limited and fabricator-level exceptions (see below), the only software (including the notion of "firmware") or hardware that need be trusted during run-time is the processor and the code that defines the agent. In the context of this

discussion, "trusted" means that the main processor is expected to behave according to its specifications and execute the associated software stack without deviation from publicly available processor specifications. In other words, the main processor itself is trusted to be free of backdoors that might allow attacks on the cached software. According to this aspect, the level of trust of the overall system will therefore be as great as the trust for the main processor itself, together with the cached agent. It will then be possible for end users to load arbitrary applications and process sensitive information knowing that neither human operators nor malicious hardware or software can attack their sensitive data or applications.

[0016] One other potential limitation to the level of trust would stem from the theoretical ability of a malicious actor to examine the state of the CPU at run-time. In general, however, tampering with the internal state of a running CPU is a difficult and expensive operation such that the level of trust that the system described here allows is still greater than for conventional solutions.

[0017] Where the processor (together with the cached agent) is the only trusted component, there is greatly reduced complexity as compared with existing systems, which can be trusted only if one also trusts that each component and human operator is working according to specifications. Of course, human operators cannot be "tested" the same as computer components, which makes trust all the more difficult. Even if one were to trust all components and humans, it would still be necessary to monitor and evaluate whether their performance honors the software stack privacy and integrity. In short, it is much more difficult to trust the integrity of many different links in a chain than it is to trust the integrity of a single link that is known to be strong.

[0018] In one embodiment, existing software solutions (hypervisors, operating systems, and applications) may be run with only a minimal change in the lowermost component of the software stack, such as the operating system (OS) or hypervisor. In particular, in this embodiment, the agent is itself included as a component of the OS/hypervisor, either as a cache-management module, or as a part of this module, which is resident in cache at run-time. This provides support for general-purpose applications and also support for legacy applications. This contrasts with the

"hardware security module" HSM (http://en.wikipedia.org/wiki/hardware_security_module), which is a dedicated external component designed to protect a specific application but that has a limited programming interface that cannot run arbitrary applications.

[0019] In general, this invention focuses on the use of the main processor cache, which is far smaller than the system memory installed in modern systems. Two changes need to be made in the software stack in order to execute in this new architecture. These changes can be made in the OS/hypervisor, whichever runs directly on the main processor. First, in any OS/hypervisor, there is some code that handles interrupts sent by hardware, as well as software faults triggered by applications. This "always present" code generally can fit inside the main processor cache. Second, most modern OSes/hypervisors assume a large main memory, although that assumption is not an absolute requirement - applications are likely to work correctly even if only very limited amounts of physical memory are available, such as with the help of virtual memory; however, without making a change in the OS/hypervisor, performance will suffer.

[0020] According to various aspects, systems and methods are provided to manage the main processor cache so that a software agent always controls which parts of main memory can occupy the cache. Some processors are already configured to enable such granular control and in such cases the configured instructions and procedures can be used to implement this feature. Some main processors lack the granular control over their memory cache, however, and will therefore require careful software algorithms to manipulate the main processor into a state in which a software agent can control which part of the memory occupies the cache at any given time.

[0021] The x86 processor is an example of a processor that lacks granular cache control. The x86 class of processors is well known in the industry and has memory cache controls that are listed, for example, in the manual "Intel 64 and IA-32 Architectures Software Developer's Manual, Volume 3: System Programming Guide" available from Intel Corp. In order to achieve control over the cache content, however, the main processor would need to notify the software stack whenever a

cache line is about to be evicted, which in turn allows the software stack to change the cache line before it is written back to main memory. The x86 lacks such software notification.

[0022] Figure 1 is a simplified illustration of a computer system that could embody different aspects of the invention. At the heart of the system is a main processor CPU 1000, which includes at least one core 1100 that receives and processes instructions using well-known components and circuitry. A CPU will of course include many different components, ranging from internal clock circuitry to an ALU, but since these are well-known, they are not illustrated even though they may be assumed to be part of the CPU 1000.

[0023] System-level software such as an operating system (OS) and/or hypervisor 2000 will typically be included to perform well-known functions. Virtual machines 6200 and applications 6300 are illustrated as running under the supervision of the OS/hypervisor 2000. A hypervisor will typically not be necessary if no virtual machines are included; both options are illustrated in Figure 1 merely for the sake of completeness. Various devices 6700, which can encompass almost any known devices such as storage, network, human interface, chipset, etc., may also be included. Some form of storage 6100 such as a hard disk system will normally be included along with less typically less persistent but faster memory devices such as system memory 7000.

[0024] In Figure 1, the system memory 7000 is shown as a single component MEM, but this is merely for the sake of clarity; in most implementations, the system memory 7000 will comprise different high-speed memory devices that may be included either stand-alone, such as for the main system memory, dedicated and located within different devices, etc. RAM, flash memory, flash-backed RAM, phase-change memory (PCM) and other such technologies may also be encompassed by the general term “system memory 7000” for purposes of this invention. From the perspective of the CPU, the system memory 7000 is therefore some addressable memory space, which does not need to be within one component or contiguous.

[0025] A typical computer system will include various buses 4000 such as an address bus, a data bus, possibly a dedicated I/O bus, etc. In cases where one or more devices is remote, there will generally also be some form of network channel, bus, or point-to-point interconnects such as the Intel QPI (QuickPath Interconnect), which is not separately shown merely for the sake of simplicity.

[0026] An essential component to this invention is the cache 5000, which is part of the CPU 1000. The general structure and properties of a cache are well-understood in the field of computer science and will therefore not be described further here.

[0027] Unique to this invention, however, is an agent 5100, which is a software component that includes an encryption/decryption module 5110 as well as an optional validation module 5120. In one embodiment, the agent 5100 is a software component within the system software 2000 and which resides in the cache 5000 at run time. In some implementations, the system software 2000 may include a cache management module 2100 that also performs various cache-related tasks; in these cases, it is possible for the agent 5100 to be either identical to or a sub-component of such a cache management module 2100. The agent 5100 may be included either as a dedicated component in the system software, or be incorporated into any other appropriate system software component. In another embodiment, the agent 5100 may be an independent component resident in the cache, in which case the cache management module 2100 may not be necessary and the OS/hypervisor may be an unmodified system software layer.

[0028] Different processor architectures will naturally have different cache structures and some have more than one. For the purposes of this invention, if there are options, it is preferable to select a cache that is large enough to hold all the code defining the cache management module 2100/agent 5100 because this will reduce the traffic (similar to memory-to-disk swapping and paging) between the cache and the main memory, and the accompanying increase in encryption/decryption activity this brings. Other caches may be chosen, however, if one is willing to accept that not all components of the agent, plus any other cached software, may fit at once in the cache. In many existing processors, the LLC (last-level cache) will be most

suitable because it is generally the largest. In x86 systems, this will be the L3 cache, which at least partially includes the lower level L1 and L2 caches such that when a lower level cache has a miss, it will read through the next level cache, not directly from memory.

[0029] As is illustrated by the dashed lines, instructions and data submitted to and transmitted back from the core can be made visible to and be intercepted by the agent 5100. According to one aspect of the invention, whenever information (data and/or instructions) is transmitted from the CPU, in particular from the core 1100 or some other internal CPU component, this transmission is intercepted by the agent 5100 and is encrypted by the agent module 5110 before it is returned outside of the logical boundaries of the CPU 1000, in particular, to system memory 7000. Instructions and data inbound to the CPU core or internal components are then decrypted by the agent 5110 before they are submitted for processing. In some implementations, all in-coming and out-going (from the perspective of the CPU) information may be decrypted and encrypted, although in other implementations this encryption and decryption could be carried out only for certain portions of system memory 7000. The system designer may freely choose which encryption and decryption algorithm the module 5110 uses. The designer will then consider such factors as the available size within the cache for the agent and how much performance he is willing to sacrifice to obtain higher levels of cryptologic security. Suitable procedures will be implemented to store the encryption keys, typically within the cache 5000 itself, and preferably only for the duration of key use.

[0030] It would in general be unnecessary and cause unacceptable performance degradation to individually encrypt/decrypt every byte of information passing between the CPU and the main system memory by way of the cached agent 5100. Instead, a larger granularity will generally be preferable. In x86 systems, for example, the virtual memory subsystem usually moves information at the level of pages. In one implementation, the agent 5100 therefore traps on accesses at the granularity of full pages and since most software execution will generate page-size traffic, the agent will typically encrypt and decrypt by page. This operation can be triggered, for example, by detecting the page faults that are generated to instruct a

hypervisor (in virtualized embodiments) to fetch the next portion of system memory that is necessary for making progress in the execution of software.

[0031] In addition to application execution, the agent 5100 might have its own meta-data that may be too big to fit inside the cache. Such meta-data may be stored in encrypted form in main memory and be accessed proactively by the agent, not requiring triggering by a page fault; consequently, such accesses do not need to be at the same page-size granularity.

[0032] It is not necessary for the agent to execute software directly on the CPU; rather, the agent could be configured to interpret the software, for example, unmodified virtual machines and applications. This would allow the agent to proactively detect a need for memory access by the software and allow it to read in such memory in smaller chunks, thereby improving throughput and performance. This is a design choice, which might also include such known techniques as pre-fetching and code scanning. One example of an interpreter that might be used is an x86 instruction set interpreter.

[0033] Although the system software 2000, the virtual machines 6200 and the applications 6300 are often conceptualized as separate entities in system diagrams such as Figure 1, these are of course in reality respective bodies of data and code that are stored as any other body of "1s" and "0s" in system storage 6100 and are loaded either as a whole or in units such as pages into the higher-speed system memory 7000 for execution under the control of the system software 2000. The contents of their code and data in respective memory address locations may therefore be cached by the CPU just as any other in order to improve performance. One other consequence of this, however, is that the agent 5100 can secure even the code that defines such software components just like any other information stored in system memory.

[0034] Even at this point, one can see a difference between this invention and known solutions. For one, except for the always present code necessary to initialize the CPU itself, the illustrated embodiment does not require action on the part of any hardware or software component other than what is part of the main processor 1000.

One other important distinction is that the agent 5100 may be implemented as a software module resident in the cache 5000. In short, the agent 5100 converts what may be a commodity processing system into a cryptoprocessor with no requirement for hardware modification or, at run-time, any trusted hardware or software components external to the CPU.

[0035] Although cache-based encryption of information transmitted from the CPU to system memory can ensure the secrecy of such information, encryption alone does not also ensure its integrity – even encrypted information can be tampered with or even accidentally altered. As Figure 1 also illustrates, the agent 5100 may therefore also include a validation software module 5120, which computes some form of validation data to associate with what the module 5110 encrypts and has stored in system memory. This validation data could be as simple as a checksum, although in most cases it would be more secure to use any known message integrity algorithm to compute hash values, message authentication codes (MACs), or digital signatures associated with each block of information (such as a page) that is encrypted. The validation module 5120 may then later re-compute the validation data for the same memory block, compare the re-computed validation data with the previously computed and stored validation data, and thereby detect any change, which would be indicated by a mis-match. This verification procedure could be carried out whenever the memory block is to be read, as a background process, or on any other schedule the system designer may prefer.

[0036] Encryption does not strictly require validation, and vice versa. Consequently, it would be possible to configure the agent to do validation alone, or to selectively enable both encryption and validation. In other words, the agent may include both the encryption/decryption module 5110, or the validation module 5120 or both (the embodiment illustrated in Figure 1 and the embodiment that, in the most secure implementation, will be the preferred one). As for a validation-only implementation, this may be sufficient in environments where privacy is not required or desired, but data integrity is. Note that even a validation-only implementation would still have the benefit of the highest level notion in this invention – an agent within the trusted cache of a hardware processor.

[0037] Even in implementations that incorporate both modules 5110 and 5120 and corresponding functionalities, it would be possible to make these selectable. For example, one may wish to maintain privacy of data, in which case the encryption module would encrypt that data before it is written to memory, whereas one may not care to incur the performance cost (however slight) of encryption/decryption for other information, such as the already well-known code of a commodity operating system or application. It may be desirable, however, to ensure that even non-encrypted memory blocks (such as pages) have not been tampered with, and in such case validation only may be sufficient. A suitable data structure such as a simple indexed list could then be included in the agent to indicate whether a particular memory block (such as a page) should be or has been encrypted and/or validated.

[0038] Figure 2 illustrates, at run-time, the embodiment in which the agent 5100 is a sub-component of the cache management module 2100 of the OS/hypervisor 2000. In other words, in this embodiment, the OS/hypervisor 2000 is modified to include the agent 5100, either as part of an existing cache management module 2100 (if it has other unrelated functions as well), or as the cache management module itself. The cache management module 2100 will maintain the list of pages, pages ranges, etc., that are allowed access to the OS/hypervisor and actively manage the cache, for example, monitor that the cache never misses via the performance counter and takes action if any error is discovered.

[0039] As Figure 2 illustrates, at run-time, the OS/hypervisor 2000 will generally be loaded at last partially into the cache 5000; this is normal even in conventional systems. Depending on their relative sizes in a given implementation, it is possible that the entire OS/hypervisor 2000 may not fit entirely within the cache 5000. In such case, as Figure 2 illustrates by the shaded upper portion of the OS/hypervisor 2000 extending logically outside the cache 5000, some portion of the OS/hypervisor 2000 may instead be stored until needed in an encrypted system memory region indicated as 9100. As the figure shows, user-level software such as the VMs 6200 and applications 6300 may also be loaded in the cache 5000, at the same time as the system software 2000 and, as needed, portions (shown shaded) of the corresponding code may at least temporarily be moved out to or left in the

encrypted system memory region 9100. This is analogous the manner in which memory pages are frequently swapped between system memory and disk storage in conventional computers. In order to ensure proper functioning of all the features of this invention, however, in particular, the feature that only the CPU and agent need to be trusted, all of the agent 5100 should preferably remain resident in the cache throughout run time.

[0040] One example of a way to manage the cache content for main processors that lack granular cache control is:

1. Determine the cache geometry and associativity, for example, using supported interfaces or by direct evaluation (for example, one known technique is to measure the time to access memory regions of different sizes; a cache hit is much faster than a cache miss);
2. Determine the maximum set of physical pages that can coexist in the cache for a typical cache that operates on physical addresses. For a processor that instead uses virtual addresses to index cache memory (possibly with additional physical address tagging), the agent 5100 may be configured to manage the assignment of virtual addresses to ensure that all mapped memory can co-exist simultaneously in the cache.
3. Maintain the discovered set of physical pages, for example, as the only pages that can be assigned to virtual addresses to be used by the hypervisor, OS or the application. Note, however (as described below), as one alternative, that virtual-to-physical page mappings and/or specific regions of physical memory can be marked as uncacheable, so arbitrary mappings are also possible provided that the only cachable pages are the ones in the set of physical pages that can co-exist (see step 2 above).

[0041] Another example of a procedure to manage the cache content is:

1. Determine the cache geometry and associativity, for example, using known, supported interfaces or by direct evaluation
2. Determine the maximum set of physical pages that can coexist in the cache while noting that the cache operates on physical addresses. As in the method described above, for a processor that instead uses virtual addresses to index cache

memory (possibly with additional physical address tagging), the agent 5100 may be configured to manage the assignment of virtual addresses to ensure that all mapped memory can co-exist simultaneously in the cache.

3. Maintain the determined set of physical pages, for example, as the only pages that can be assigned to virtual addresses to be used by the hypervisor, OS or the application. Note, as above, the alternative that virtual-to-physical page mappings and/or specific regions of physical memory can be marked as uncacheable.

4. To read any memory outside the set of pages mentioned above, the OS or hypervisor may use any supported processor methods to read memory into a CPU register while bypassing the cache. Once the copied memory is in a processor register, it is safe to copy the register content into one of the physical pages that occupy the cache.

5. To write any memory outside the set of pages mentioned above, use reverse logic relative to step 4.

6. To protect the CPU cache from access via DMA, configure the IOMMU (provided in many modern processor architectures, such as the x86 architectures) or some other CPU-specific method to block DMA to pages that occupy the cache.

[0042] Various methods are known to allow the cache to be bypassed when accessing a particular address or range of addresses. Examples of such methods that are available for, among other architectures, the x86, include: Memory Type Range Registers (MTTRs), which can command the x86 processor not to cache certain memory areas; the PAT extension to the MMU, which controls caching of individual pages; and special non-temporal load and store instructions, such as MOVNTDQ and MOVNTDQA.

[0043] In most existing x86 systems, a set of physical pages that can coexist in the CPU cache is any contiguous set of pages that starts on a multiple of the cache size. For example, if the cache size is 8 MB, a first set of pages are the pages that occupy the system address space from [0-8) MB, a second set is the pages that occupy the system address space from [8-16) MB, and so on. This feature can be exploited in various ways to test whether a given set of physical pages can coexist, assuming a software stack that can fit within the tested set of pages. This can be

accomplished in various ways, for example, by determining whether cache evictions have occurred using either a processor cache performance counter or a hardware bus analyzer (that is, a “memory protocol analyzer”). These methods comprise, more specifically, the following:

[0044] Using the processor cache performance counter.

1. Invalidate the cached content.
2. Read every byte (or other memory unit) in the set of pages to be tested.
3. Record the value of the processor cache performance counter (number of cache misses).
4. Repeat step 2, that is, read every byte (or other memory unit) again.
5. Again read the processor cache performance counter (number of cache misses).
6. Compare the results. If the number of cache misses is the same as in step 3, then all pages coexisted in the cache, since this will show that re-reading the whole range did not cause any miss.

[0045] Using a hardware bus analyzer:

1. Read every byte (or other memory unit) in the set of pages to be tested.
2. Using a bus analyzer, start measuring any access to main memory by the main processor.
3. Repeat step 1.
4. Analyze results of the bus analyzer; in particular, search for any reads from the addresses in the set of pages. If there is no read in that range then all the pages coexisted in the cache.

[0046] There are other methods for determining if cache misses occur, which may be used instead of the method outlined above. Just one example is to use a device capable of performing non-cache-coherent reads of memory (bypassing the cache).

[0047] Another alternate method involves using the x86 INVD instruction to invalidate the cache contents, and then inspecting the associated system memory

pages to see if any lines were written back. In particular, the cache region may be filled with an initial pattern and then flushed to memory via WBINVD. The region is then filled with a distinct test pattern, and the cache is invalidated via the INVD instruction. The memory backing the cache is then examined and the information about which lines were evicted may be used to determine the sets of pages that conflict in the cache, in order to find the non-conflicting set. Just one example of a way to determine this would be to observe that any cache lines that do not contain the original pattern must have been evicted due to conflicts when writing the test pattern.

[0048] Inspection of the cache as just described can be carried out at different times and by different entities, depending on a given implementation choice. In one embodiment, for example, cache inspection could be done offline and manually for each processor configuration. In another embodiment, a loadable Linux kernel module may run the conflict test. In other embodiments, cache inspection could be done at boot-time when the cache management/agent software is first loaded. In yet other embodiments, such checks could be performed periodically at run-time by the secure OS/hypervisor, as an extra layer of protection to detect any unexpected cache evictions that may have occurred (and allowing them to be scrubbed quickly to reduce the exposure time of sensitive data, as is explained further below).

[0049] Yet another approach is to time memory accesses (for example, using the x86 RDTSC "read timestamp counter" instruction) to detect misses, for example, as an alternative to using the cache-miss performance counter.

[0050] Note that the more recent generation of Intel x86 processors (starting with those code-named "Sandy Bridge") use so-called "complex cache addressing", where the hardware hashes a memory address to determine where it will reside in the cache. In other words, page coloring in such systems may no longer work. In such systems, the method for computing a set of non-conflicting pages will typically require the generate-and-test procedures outlined above.

[0051] In many common systems, a modern hypervisor or OS can be given as a constraint a set of physical pages that are allowed to be used, with the expectation

that the hypervisor or OS will respect that constraint and reference only the given pages for execution. One additional step will normally be required, however, to prevent hardware devices from reading the state of the CPU cache. By default, it will normally be possible for devices to issue DMA (direct memory access) requests to read from any system address. If that memory happens to be cached by a DMA-cache-coherent CPU, such as in an x86 system, the CPU will allow access to such memory, which could defeat the trust module and expose the software state.

[0052] An IOMMU (I/O memory management unit) 1300 is a device configurable by software that can restrict hardware devices from accessing system addresses, such that a device cannot read or write to memory that hasn't been explicitly allocated (mapped) for it. Many commodity processors include the IOMMU as part of the CPU circuitry itself; in other cases, it will be a separate component, in which case known measures will need to be taken to ensure its security. In addition to an IOMMU, some CPUs have other technologies that can provide DMA protection, for example, the Intel x86 TXT technology. One method to provide DMA protection is therefore to use the IOMMU, although it is possible to implement various embodiments of this invention using other CPU-specific alternatives. To mitigate access to the CPU cache by devices, an IOMMU may therefore be configured to prevent such access. This is normally done by configuring the IOMMU so that each page in the set of pages occupying the cache will have no mapping, or a mapping to an invalid address range, such as an address that is not mapped by any device or memory in the system address space. In normal operation, the software stack will not require any of the devices to read from the memory that occupies the cache, so normal operation can be expected. If a malicious hardware device attempts to read the software state, however, this will be prevented by the IOMMU. One exception to this are some pages that occupy the cache and are knowingly selected by the software stack for I/O communication with devices. For such pages, the IOMMU will be configured to map the pages so that they are allowed to be accessed by hardware devices. One reason for having such pages is that some old hardware devices may be able to address only small amounts of the address space; for example, ISA devices can address only the first 16 MB of the system address space.

Communicating with such devices might therefore need to happen using memory that is within the cache.

[0053] Even when allowing access to certain pages inside the cache for I/O purposes, however, the pages that are allowed access for devices are only those that are used at that point in time for I/O communication. Once I/O communication is completed, the IOMMU protection can be established once again and the page can be used for other purposes, such as for running other software code. Note that this is not a security risk, as the agent is aware of any page being used for I/O and can treat it accordingly.

[0054] Since the agent 5100 contains executable code loaded in the CPU cache 5000, and this cache is empty at the time the CPU is first powered on or after a reset, the issue arises as to how to load the agent in a secure manner. Each processor architecture will have different ways to address this issue that skilled designers will know how to use, but for the purpose of illustration, assume that the invention is to be implemented in a system based on the current x86 architecture. As mentioned above, Intel at present provides the Trusted Execution Technology (TXT), which includes a Trusted Platform Module (TPM) chip as a hardware extension of the x86-type CPU. TXT is a known method for authenticating software launch on a system with a TPM. According to some proposals, the TPM may be incorporated into the CPU itself in the future. Figure 1 shows these as TXT 3000 and TPM 3100.

[0055] One method to arrange secure loading of the agent 5100 is as follows:

- 0) Using a trusted system, compute a hash value of the code to be loaded as the agent 5100 and store this securely.
- 1) Load agent loader code in the system main memory as with other applications. At this stage, the agent code might be visible in main memory, and therefore might come under malicious attack.
- 2) Following Intel-specified and therefore known operation of the TXT technology, execute the agent loader code in an MLE (Measured Launch Environment) to load the agent 5100 into the cache 5000 and record the agent hash value in the TPM 3100. The MLE will provide protection from DMA for all the agent

loader code. At this stage, there are no secrets loaded in memory, only the agent loader code.

3) Rehash the cached agent code and, via a remote system, communicate with the TPM (through a preferably dedicated network agent) and compare this rehashed value with the previously stored agent hash value. If these two hash values match, then the cached agent will be authenticated, that is, it will be known that at run-time it is running unmodified.

4) Private information (storage keys, VM images, etc.) may now be communicated to the host running the agent 5100 since the agent, now securely loaded, will ensure that all execution happens only within the CPU cache. Note that it is possible for a CPU to generate encryption keys securely that reside in the cache with no need for communication external to the CPU itself.

[0056] Figure 3 illustrates the interaction with system memory by various hardware and software components in typical implementations, and also illustrates how the invention is able to support not only operations with system memory, but also I/O devices that may include memory of their own. Figure 3 also illustrates the distinction between safe/private system memory address regions (9000) and unsafe, potentially malicious system memory address regions (9001).

[0057] As for safe/private memory 9000, in general, any memory currently in the cache 5000 may be considered safe to access – it is trusted. Information in memory that has been encrypted by the agent, although not trustworthy in the sense of cached information, is nonetheless safe by virtue of the encryption and further by validation. Software in the stack may consider such address regions as safe in the sense that they view the system memory that they address to be physical addresses, but in reality are within the cache 5000 and can therefore be trusted. In other words, the software stack thinks that it is addressing system memory as usual, but is in reality addressing the cache 5000.

[0058] Note that encryption by the module 5110 (Figure 1) provides confidentiality of the information stored in these memory regions (including such information as the state of the virtual machines 6200 and applications 6300 and any portions of the OS/hypervisor 2000 that cannot fit in the cache 5000); validation by

the module 5120 (Figure 1) provides data integrity; but actual availability is not guaranteed. Thus, a malicious entity could still tamper with code or data, but this tampering would be detected as explained above: tampering would cause a mismatch in the hash or other validation data for the affected block. Note that even without the validation module 5120 and its related operations, tampering with executable code stored in encrypted form would almost certainly lead to detectable runtime errors or an outright failure of the affected code to make forward progress.

[0059] In Figure 3, the software stack 2000, 6200, 6300 is shown as being outside of the cache. This is merely for sake of illustration. As explained above, at run-time most and preferably all of the OS/hypervisor 2000, in particular, the agent 5100, will be resident in the trusted cache 5000, as well as any other parts of the stack (such as VMs 6200 and/or applications 6300) that the OS/hypervisor 2000 chooses to load into the cache as well for conventional performance reasons.

[0060] Unencrypted memory regions 9001 cannot be trusted and are not safe from attack. In general, any memory not currently in the cache is considered unsafe and must be encrypted by the agent to prevent privacy and integrity attacks. The un-trusted memory space 9001 includes normal system memory as used by a conventional system. No special attributes need to be associated with a particular address in such a system memory space, which can therefore be used by executing code (although preferably not code itself), storage or device I/O. As explained above, the un-trusted memory space may also be used to store executable code that is encrypted/validated; this code may be paged in and executed after re-validation and decryption by the agent 5100.

[0061] Unencrypted system memory may not be the only mappable memory in the overall system. Depending on what they are and how they are built, devices 6700 may, for example, be able to access their own dedicated, typically onboard device memory 6710 (such as the memory on a video card). Note that, even in conventional systems, it is rare to map device memory 6710 for any purpose other than I/O with devices; in particular, code is not typically executed from device memory. As the CPU cache 5000 provides access to system addresses, some of the addresses might not be backed by system memory at all; instead, they might be

backed by the device memory 6710, for example, a video card with 512MB, which is not part of the system memory, but in theory can be used as such.

[0062] Devices may also be enabled to access an untrusted system memory address space 9510 via a DMA I/O channel 9600. In a trusted setup, however, DMA protection 9201 (for example, provided by IOMMU 1300 – see Figure 1) is configured to prevent access to the system address ranges that occupy the CPU cache. This is symbolized by the arrow from the I/O channel 9600 “bouncing off of” the DMA protection module 9210. This IOMMU protection mechanism is described in greater detail above.

[0063] The encrypted storage region 9610 shown in Figure 3 is simply undedicated memory. For a system that does not trust the system memory in general, it is beneficial to include a region such as 9610 as an untrusted storage area to store an encrypted RAM disk, as encrypted swap storage, etc. This can improve performance compared to slow storage devices, and can compensate for the limited processor cache size.

[0064] A system memory area that backs the CPU cache is indicated as 9500. As one example, this area might contain zeros or other patterns (for example, randomly generated bits, in which case the seed should be stored for subsequent recomputation and validation) that indicate and validate that no memory is leaked back from the cache to system memory, such as by an accidental write-back. In other words, if, upon later inspection by the agent, the system memory region 9500 does not contain the same pattern, then the system will know there has been some cache leakage. Some pattern other than all zeroes has the advantage that a reset won't be confused with a correct state. Note that this backing memory region 9500 is still not to be trusted, as it is neither in the cache nor in a region of encryption, but it may still be safely used by the agent as described.

[0065] The region of non-conflicting memory 9500 may also be used during the initial loading of the agent 5100. One example of a procedure that uses the cache-backing memory region 9500 includes the following steps:

1. T0: The image of the OS/hypervisor 2000 kernel image is read into the memory region 9500. At this point, all information in region 9500 may be in plain text, that is, none of the information is yet safe.
2. T1: The OS/hypervisor kernel initializes the agent 5100, (which will then manage the caching property of physical memory to limit the kernel to only the set of non- conflicting pages), configure the IOMMU, and enter into the proper run-time state.
3. T2: The agent 5100 fills the memory range 9500 backed by the cache with the eviction-detection-fill-pattern.
4. T3-Tn: The system is run from the cache only. Note that the memory region 9500 backing the cache is not an encrypted version of the running state; rather, it is “stale” memory containing a fill pattern. Periodically, the agent 5100 may examine the memory region 9500 to check that this fill pattern is intact, and, thereby, to detect if any cache lines have leaked due to eviction.

[0066] It is well known that caches operate on aligned, contiguous chunks of data, referred to as cache lines. For example, 64 bytes is a common cache line size in modern x86 processors, and movement of data between the cache and system memory occurs at cache-line granularity. The processor cache is typically organized in a set-associative manner; each cache index can store up to N cache lines in an N-way set-associative cache. For modern x86 processors, N is commonly 8 or 16. The processor typically employs a deterministic mapping from memory addresses to cache indexes, such that a large number of lines in system memory will be mapped to the same cache index, although the cache may only hold up to N of them at a time. When more than N memory addresses that map to the same cache index are accessed, they cannot all fit in the cache at the same time, resulting in cache evictions. Modern processors commonly employ a least-recently-used (LRU) policy (or some approximation) to determine which already-cached line to evict in order to make space for a newly-accessed line from system memory. When a modified (also known as “dirty”) line is evicted, it is written back to system memory. In general, processors do not provide any way for software to interpose on cache evictions or writebacks.

[0067] This invention does not rely on the existence of special hardware support or processor modes for locking down portions of the cache contents, as required by the “bus encryption” method described in Xi Chen, Robert P. Dick, and Alok Choudhary. “Operating System Controlled Processor-Memory Bus Encryption”, in Proceedings of Design, Automation and Test in Europe, (DATE 2008), Munich, Germany, March 2008, or the CARMA system (cited below).

[0068] Uncontrolled cache evictions leak data from trusted cache memory to untrusted system memory, violating privacy. Also, a subsequent access to an evicted line will load data from untrusted system memory into the trusted cache, enabling an attacker to violate integrity. As a result, effective cache management is essential for preserving the privacy and integrity of both code and data. One key challenge is therefore preventing evictions from cache conflicts, a problem that is not discussed in related work. This involves carefully selecting the portions of memory that are allowed to be cached, and marking the rest of memory uncacheable. In general, the processor memory management unit (MMU) 1200 (Figure 1) permits software to control, map, and protect memory at the granularity of pages, which are typically much larger than individual cache lines. For example, the smallest page size is 4096 bytes for x86 processors.

[0069] One method for addressing this issue is: 1) Find a non- conflicting set of pages; then 2) use only these pages for running the current software, since only these pages will have a caching policy that allows the CPU to cache them; remaining pages in the system will typically have a caching policy that prevents caching. Since the set of pages that need to be found as non-conflicting are the only pages ever to be mapped into the cache, even when data is read from the rest of the untrusted memory, it is still read in such a way that does not place them into the cache and that generally will not cause eviction of memory that is currently in the cache.

[0070] The set of non-conflicting pages depends on the particular function used by the processor to map memory addresses to cache indexes, which in turn depends on details of the processor implementation. In some cases, once the cache organization (size and associativity) is determined, the mapping can be computed and controlled via page coloring. In other cases, such as the most recent

generations of Intel x86 processors (starting with “Sandy Bridge”), the mapping is an opaque, complex hash function that is not disclosed by the processor vendor. As a result, an experimental approach is required to find a suitable set of non-conflicting pages, which in turn demands a reliable procedure for detecting conflicts. PrivateCore has developed several eviction-detection techniques, including methods that bypass the cache to perform non-coherent reads from system memory, as well as methods that involve monitoring hardware performance counters. The further ability to detect evictions at runtime, during production system operation, is also important, in order to catch and scrub any unexpected leaks quickly and reduce exposure time.

[0071] At this point, those skilled in computer security at the system level will appreciate some of the advantages of the various aspects of this invention relative to other solutions proposed in the literature. One clear advantage is that the software-only approach of the invention does not require any special hardware support, and may even be used with commodity hardware, such as systems based on modern Intel x86 processors. If both encryption/decryption and validation features are implemented (not required) then the disclosed system is designed to guarantee both privacy (confidentiality) and integrity for all code and data in the system.

[0072] The trusted computing base (TCB) of the implementation of the invention described above does not need to extend beyond the processor, and may operate according to a threat model that explicitly assumes that all other hardware is untrusted and potentially malicious. However, for secure launch and attestation, the TCB should preferably also encompass the components required to establish a root of trust, such as Intel Trusted Execution Technology (TXT), which may include off-processor Trusted Platform Module (TPM) hardware in some implementations.

[0073] In particular, system memory itself may be considered untrusted, and its contents may be assumed to be vulnerable to being read or written by an attacker, violating privacy or integrity, respectively. Similarly, all other hardware devices and interconnects outside the processor are untrusted, and potentially under the control of an attacker. For example, a compromised I/O device, such as a common network interface card (NIC) capable of performing DMA, may be exploited

to violate privacy or integrity by maliciously inspecting or corrupting the contents of code and data in main memory or caches.

[0074] In this invention, the only memory that needs to be trusted is the processor cache, which is a physically-integrated part of the processor itself; any other memory that is used may be encrypted or, in the case of device memory or possible DMA, controlled using IOMMU. This contrasts with the Cryptkeeper system described in Peter A. H. Peterson, "Cryptkeeper: Improving Security with Encrypted RAM", in IEEE International Conference on Technologies for Homeland Security (HST 2010), November 2010. Cryptkeeper divides main-memory into separate cleartext and encrypted portions, and does not attempt to manage cache contents; all RAM is cacheable. As a result, Cryptkeeper fails to preserve privacy for its cleartext memory region, making an explicit choice to allow some memory contents to be leaked, in order to improve performance. Note that even if Cryptkeeper reduced the size of its cleartext region to fit in the cache, without additional modifications, accesses to encrypted regions would still induce cache evictions and therefore leak cleartext data back to memory. Cryptkeeper also does not address integrity, and thus does not provide protection against attacks that modify or corrupt code and data in either the cleartext or encrypted portions of RAM.

[0075] Similarly, the bus-encryption scheme proposed in Chen, et al., cited above, assumes that an attacker is "unable to tamper with the contents of memory" and unable to "modify the kernel", which is assumed to reside in read-only non-volatile memory off-processor, as is the case for some embedded systems.

[0076] Merely by way of comparison with an approach proposed publicly only after the priority date of this application, the CARMA system described in Amit Vasudevan, Jonathan M. McCune, James Newsome, Adrian Perrig, and Leendert van Doorn. "CARMA: A Hardware Tamper-Resistant Isolated Execution Environment on Commodity x86 Platforms", in Proceedings of the ACM Symposium on Information, Computer and Communications Security (ASIACCS 2012), Seoul, Korea, May 2012, relies on cache memory exclusively, so that system memory is never accessed and need not be present. The authors state explicitly that "execution will be entirely within cache and not interact with untrusted RAM", and that they were

able to successfully remove all DRAM in a prototype implementation. This choice constrains the types of workloads that can be supported by CARMA significantly; in particular, it makes CARMA essentially useless for large, general-purpose VMs, complicated applications, or most unmodified commercial software, which are almost all too big to fit entirely within the cache. Moreover, the only mechanism CARMA includes for any notion of support for I/O devices is via an in/out command to passive devices, initiated by the CPU itself. This is a severe restriction as compared with the present invention, which enables even devices to initiate memory writes, through shared memory, in order to share data with the CPU.

[0077] The TRESOR system (mentioned also above) described in Tilo Müller, Felix C. Freiling, and Andreas Dewald. “TRESOR Runs Encryption Securely Outside RAM”, in Proceedings of the 20th USENIX Security Symposium, San Francisco, California, August 2011, is designed to address the more limited problem of storing disk encryption keys securely, so that physical attacks on RAM cannot compromise encrypted disk contents. TRESOR stores encryption keys in privileged machine-specific processor registers (x86 debug registers), and implements the AES encryption algorithm without storing any sensitive information in RAM. In contrast to the architecture of the invention described in this application, which protects privacy and, optionally, integrity for all code and data in the system, TRESOR protects only encryption keys. The approach of using privileged processor registers to store information securely does not scale beyond a very small amount of state; TRESOR is, for example, able to store only 256 bits securely (in four 64-bit debug registers), just enough to hold a single AES-256 key. This (mis)use of the x86 debug register mechanism also precludes its intended use for hardware breakpoints and watchpoints. Furthermore, without privacy and integrity protection for the whole code and data, an attacker can manipulate the TRESOR implementation, and through that alteration, gain access to the keys stored in CPU debug registers.

CLAIMS

1. A system for securing information in a computer CHARACTERIZED BY:
 - a central processing unit (CPU) (1000) that includes at least one core (1100) that receives and executes instructions and processes data;
 - a cache (5000) within the CPU;
 - system memory that has a system memory address space;
 - a software stack (2000, 6200, 6300) that runs on the CPU and is embodied at least partially as instructions and data stored in the system memory; and
 - an agent (5100) running within the cache and comprising an encryption/decryption module (5110) that encrypts information, which may comprise code and/or data, before it is written to the system memory by the CPU and decrypts the information read from system memory for, and before, processing by the CPU core as the instructions and data.
2. A system as in claim 1, further CHARACTERIZED IN THAT:
 - the agent (5100) encrypts the information in blocks; and
 - the agent includes a validation module (5120) that associates with each encrypted, stored information block a validation value.
3. A system as in claim 2, further CHARACTERIZED IN THAT the validation module (5120) subsequently re-computes the validation value for each information block and compares it with the corresponding, previously associated validation value, whereby the validation module (5120) can detect any changes to the information block.
4. A system as in claim 1, further CHARACTERIZED IN THAT the software stack executes on a first portion of the system memory address space, which backs the cache (5000), such that any read or write requests addressed to the first portion by a entity in the software stack are fulfilled by the cache without accessing system memory.

5. A system as in claim 4, further CHARACTERIZED IN THAT the software stack also addresses a second, encrypted portion (9000) of system memory address space.
6. A system as in claim 5, further CHARACTERIZED BY:
 - system software (2000) that, at run-time, is at least substantially loaded within and executes in the cache (5000);
 - in which:
 - the agent comprises a portion of the system software and is, during run-time, always resident in the cache.
7. A system as in claim 1, further CHARACTERIZED BY:
 - at least one I/O device (6700) that includes device memory (6710); and
 - an I/O memory management unit (1300) configured to allow the I/O device to address and access the cache only for predetermined I/O operations under supervision of the agent (5100).
8. A system as in claim 1, further CHARACTERIZED BY:
 - at least one I/O device (6700) that is configured for direct memory access via an I/O channel (9600); and
 - an I/O memory management unit (1300) configured to block DMA to the cache.
9. A system as in claim 1, further CHARACTERIZED BY a trusted platform module (TPM) component (3100) provided to securely load the agent (5100) into the cache (5000) before execution of other software in the software stack.
10. A system as in claim 1, further CHARACTERIZED IN THAT the software stack includes a hypervisor and at least one virtual machine (6200).
11. A system as in claim 1, further CHARACTERIZED IN THAT the software stack includes an interpreter provided for executing unmodified commodity software.

12. A system as in claim 1, further CHARACTERIZED BY:

a cache-backing memory region (9500) that is a sub-set of the system memory and that is provided to contain a data pattern;

in which:

the agent is provided to inspect the contents of the cache-backing memory region to determine any change in the data pattern, indicating eviction of at least a part of the cache contents into the cache-backing memory region.

13. A method for securing information in a computer that has

a central processing unit (CPU) (1000) that includes at least one core (1100) that receives and executes instructions and processes data;

a cache (5000) within the CPU;

system memory that has a system memory address space; and

a software stack (2000, 6200, 6300) that runs on the CPU and is embodied at least partially as instructions and data stored in the system memory;

CHARACTERIZED BY:

loading an agent (5100) into the cache;

using the agent, encrypting information, which may comprise code and/or data, before it is written to the system memory by the CPU and decrypting the information read from the system memory for, and before, processing by the CPU core as the instructions and data.

14. A method as in claim 13, further CHARACTERIZED BY:

encrypting the information in memory blocks; and

computing and associating validation data with each encrypted memory block.

15. A method as in claim 14, further CHARACTERIZED IN THAT the memory blocks are memory pages.

16. A method as in claim 14, further CHARACTERIZED BY: re-computing the validation data for each information block and comparing it with the corresponding,

previously associated validation data, thereby detecting changes to the information block.

17. A method as in claim 13, further CHARACTERIZED BY:

executing the software stack on a first portion of the system memory address space, which backs the cache (5000);

and fulfilling in the cache any read or write requests addressed to the first portion by an entity in the software stack, without accessing system memory.

18. A method as in claim 17, further CHARACTERIZED BY allowing the software stack also to address a second portion (9000) of the system memory address space, which contains the encrypted information.

19. A method as in claim 13, further CHARACTERIZED IN THAT the computer includes at least one I/O device (6700) that includes device memory (6710), further comprising allowing the I/O device to address and access the cache only for predetermined I/O operations.

20. A method as in claim 13, further CHARACTERIZED IN THAT the computer includes at least one I/O device (6700) that includes device memory (6710), further comprising blocking direct memory access (DMA) to the cache by the I/O device.

21. A method as in claim 13, further CHARACTERIZED BY securely loading the agent (5100) into the cache (5000) before execution of other software in the software stack.

22. A method as in claim 13, further CHARACTERIZED IN THAT the system memory address space is organized into memory blocks, further comprising:

determining a non-cache-conflicting-set of the memory blocks; and

marking only the system memory associated with the non-cache-conflicting as cacheable, thereby preventing cache conflicts by system memory.

23. A method as in claim 13, further CHARACTERIZED BY executing unmodified software in the software stack using interpretation.

24. A method as in claim 13, further CHARACTERIZED BY:

storing a data pattern in a cache-backing memory region (9500); and
inspecting the contents of the cache-backing memory region to determine any change in the data pattern, indicating eviction of at least a part of the cache contents into the cache-backing memory region.

25. In a computer that has a central processing unit (CPU) (1000) that includes at least one core (1100) that receives and executes instructions and processes data; a cache (5000) within the CPU; system memory that has a system memory address space; and a software stack (2000, 6200, 6300) that runs on the CPU and is embodied at least partially as instructions and data stored in the system memory;

a method for preventing cache conflicts by system memory pages
CHARACTERIZED BY:

determining a non-cache-conflicting-set of the memory pages; and
marking only the system memory associated with the non-cache-conflicting as cacheable, thereby preventing cache conflicts by system memory.

26. A method as in claim 25, in which the computer includes a processor cache performance counter and each memory page comprises a set of memory units, the method further CHARACTERIZED IN THAT:

invalidating contents of the cache;
reading in every unit in a set of pages to be tested;
record the value of the processor cache performance counter to determine a number of cache misses;
again reading in every unit in the set of pages to be tested;
again reading the processor cache performance counter; and
determining that the tested memory pages are non-cache-conflicting if the results of the two reads of the performance counter are equal.

27. A method as in claim 25, in each memory page comprises a set of memory units, the method further CHARACTERIZED BY:

reading in every unit in a set of pages to be tested;
measuring any access to the system memory by the CPU;

again reading in every unit in the set of pages to be tested;
searching for any reads from system memory addresses in the tested set of pages; and

if no read is found, determining that the tested memory pages are non-cache-conflicting and may coexist in the cache.

28. A method as in claim 25, further CHARACTERIZED BY:

invalidating the cache contents; and
inspecting a test set of the system memory pages to determine if any cache lines were written back to system memory.

29. A method as in claim 25, further CHARACTERIZED BY:

filling the cache with an initial data pattern;
flushing the cache to a cache-backing region of the system memory;
filling the cache with a test data pattern;
invalidating the cache;
examining the cache-backing region of the system memory to determine which, if any, cache lines were evicted by identifying any cache lines that do not contain the initial data pattern.

30. A method as in claim 25, in which the computer has an x86 architecture, further CHARACTERIZED BY flushing the cache using a WBINVD instruction and invalidating the cache using an INVD instruction.

FIG. 1




FIG. 2


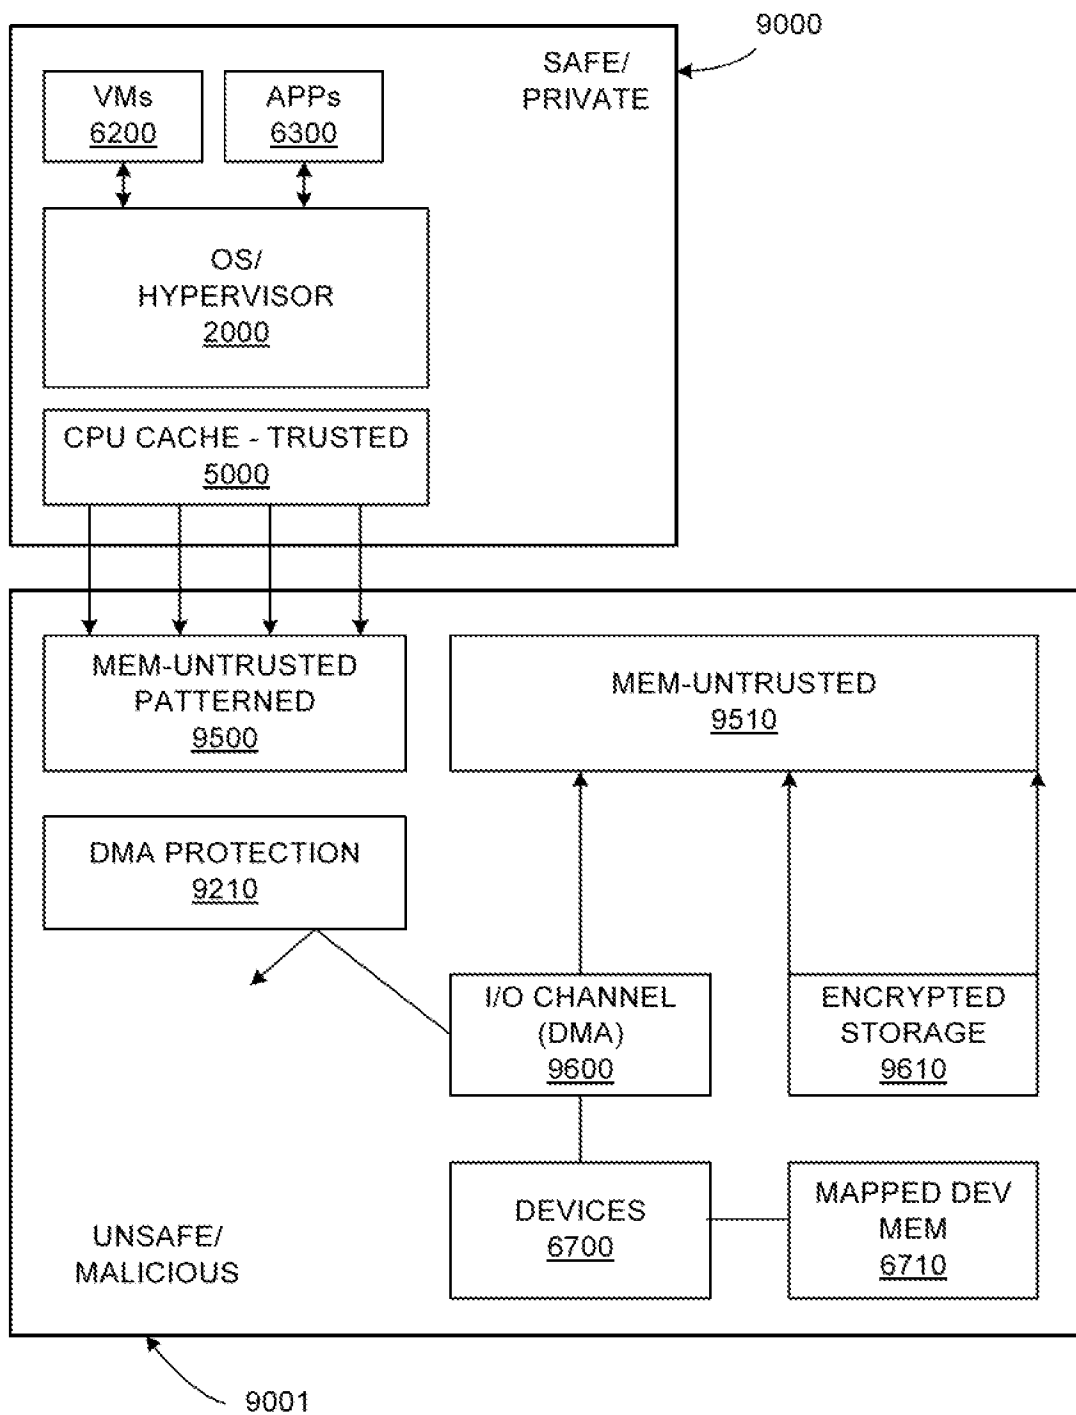


FIG. 3



INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 12/55210

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - G06F 21/00 (2012.01)

USPC - 713/182

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC (8) - G06F 21/00 (2012.01)

USPC - 713/182

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
USPC - 713/182 or 713/189 or 713/190 (See Keywords Below)

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

Patbase; Google Scholar; PubWEST(PGPB, USPT, EPAB, JPAB);

Search Terms: CPU, processor, cache, core, processing unit, internal memory, external memory, system memory, encrypt, decrypt, hash, value, compare, secure, trusted, data, instruction, crypto, hypervisor, virtual machine, agent, block, page, validate, verify, DMA

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2010/0281273 A1 (LEE et al.), 04 November 2010 (04.11.2010), entire document, especially Abstract; para [0016], [0035]-[0037], [0038], [0043], [0045], [0073]-[0075], [0104], [0122], [0139]	1-24
A	US 2011/0167278 A1 (GOTO et al.), 07 July 2011 (07.07.2011), entire document	1-24
A	US 2008/0109660 A1 (MITRA), 08 May 2008 (08.05.2008), entire document	1-24

☐ Further documents are listed in the continuation of Box C.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

09 January 2013 (09.01.2013)

Date of mailing of the international search report

25 JAN 2013

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
P.O. Box 1450, Alexandria, Virginia 22313-1450

Facsimile No. 571-273-3201

Authorized officer:

Lee W. Young

PCT Helpdesk: 571-272-4300
PCT OSP: 571-272-7774

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 12/55210

Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:

2. ☐ Claims Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:

3. ☐ Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:
-see extra sheet-

1. ☐ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying additional fees, this Authority did not invite payment of additional fees.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
4. ☒ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:
1-24

Remark on Protest

- ☐ The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
- ☐ The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- ☐ No protest accompanied the payment of additional search fees.

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 12/55210

Continuation in Box III: Observations where unity of invention is lacking

Invention Groups:

Group I: claims 1-24

Group II: claims 25-30

Reasons for Lack of Unity of Invention:

This application contains the following inventions or groups of inventions which are not so linked as to form a single general inventive concept under PCT Rule 13.1. In order for all inventions to be searched, the appropriate additional search fees must be paid.

Group I: claims 1-24, drawn to a system/method for securing information in a computer comprising an encryption/decryption module that encrypts information.

Group II: claims 25-30, drawn to a method comprising preventing cache conflicts by system memory pages.

The inventions listed as Groups I - II do not relate to a single general inventive concept under PCT Rule 13.1 because, under PCT Rule 13.2, they lack the same or corresponding special technical features for the following reasons:

The special technical feature of the Group I invention is a system/method for securing information in a computer comprising an encryption/decryption module that encrypts information. The special technical feature of the Group II invention is a method comprising preventing cache conflicts by system memory pages.

Groups I and II share the technical feature of a central processing unit (CPU) (1000) that includes at least one core (1100) that receives and executes instructions and processes data; a cache (5000) within the CPU; system memory that has a system memory address space; and a software stack (2000, 6200, 6300) that runs on the CPU and is embodied at least partially as instructions and data stored in the system memory.

However, this shared technical feature does not represent a contribution over the prior art US 2010/0281273 A1 (Lee et al.) (04 November 2010), which discloses a processor is provided having a processor core, a cache memory, a plurality of registers [...] the processor to execute the at least one on-chip instruction [...] the processor encrypts data written to, and decrypts data read from (abstract, para [0016]); "machine memory space" denotes the actual physical memory 22 available in the hardware 16 (para [0037]); and virtualized software stacks, (para [0037]); all software running on the CPU of the present invention is part of a module [...] untrusted parts of the software stack are in a generic untrusted module called module zero (para [0081]). Thus, Groups I - II do not share a same or corresponding special technical feature that would provide a unifying contribution over the prior art. None of these special technical features are common to the other groups. Therefore, unity of invention is lacking.