



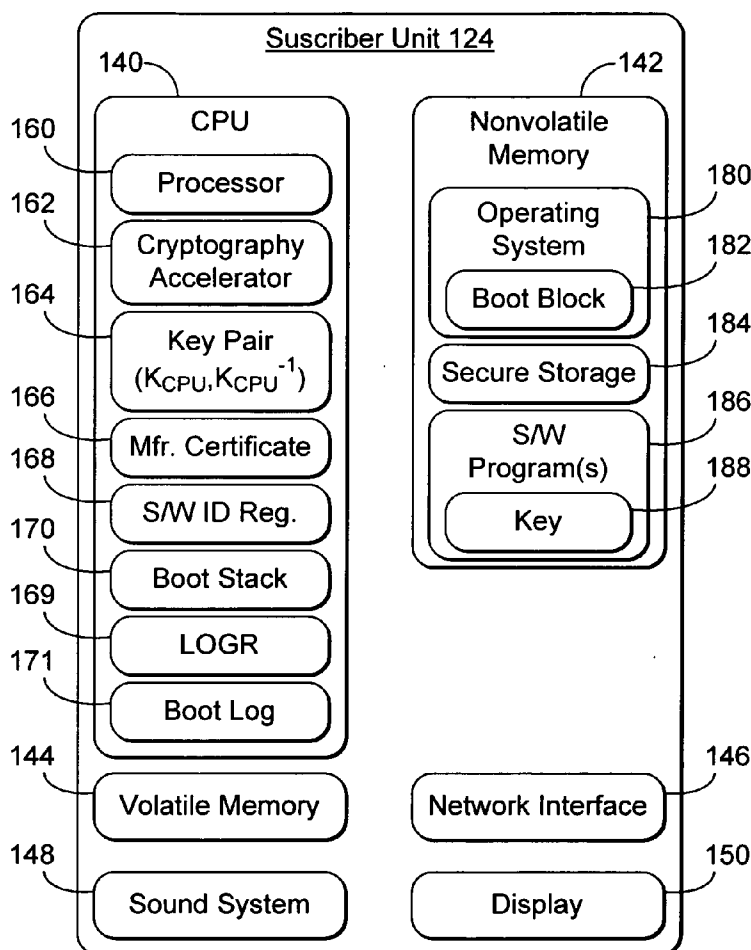
US 20050289067A1

(19) **United States**(12) **Patent Application Publication**  
**Lampson et al.**(10) **Pub. No.: US 2005/0289067 A1**(43) **Pub. Date: Dec. 29, 2005**(54) **SYSTEM AND METHOD FOR SECURE  
STORAGE OF DATA USING A KEY**(60) Provisional application No. 60/105,891, filed on Oct.  
26, 1998.(75) Inventors: **Butler W. Lampson**, Cambridge, MA  
(US); **John D. De Treville**, Seattle, WA  
(US); **Paul England**, Bellevue, WA  
(US)**Publication Classification**(51) **Int. Cl.<sup>7</sup> ..... H04L 9/00**(52) **U.S. Cl. .... 705/51**

Correspondence Address:

**LEE & HAYES PLLC****421 W RIVERSIDE AVENUE SUITE 500  
SPOKANE, WA 99201**(57) **ABSTRACT**(73) Assignee: **Microsoft Corporation**, Redmond, WA(21) Appl. No.: **11/207,917**(22) Filed: **Aug. 19, 2005****Related U.S. Application Data**(63) Continuation of application No. 10/430,994, filed on  
May 7, 2003, which is a continuation of application  
No. 09/227,568, filed on Jan. 8, 1999.

In one aspect, a data structure to be encrypted is received in a device, the data structure including content along with a statement of conditions under which the content may be decrypted. The data structure is encrypted using a symmetric key of a processor of the device. In another aspect, a data structure is decrypted using a processor symmetric key. A statement of conditions under which content in the data structure can be decrypted is obtained, and testing is performed as to whether the conditions are satisfied. The decrypted content is returned only if the conditions are satisfied.




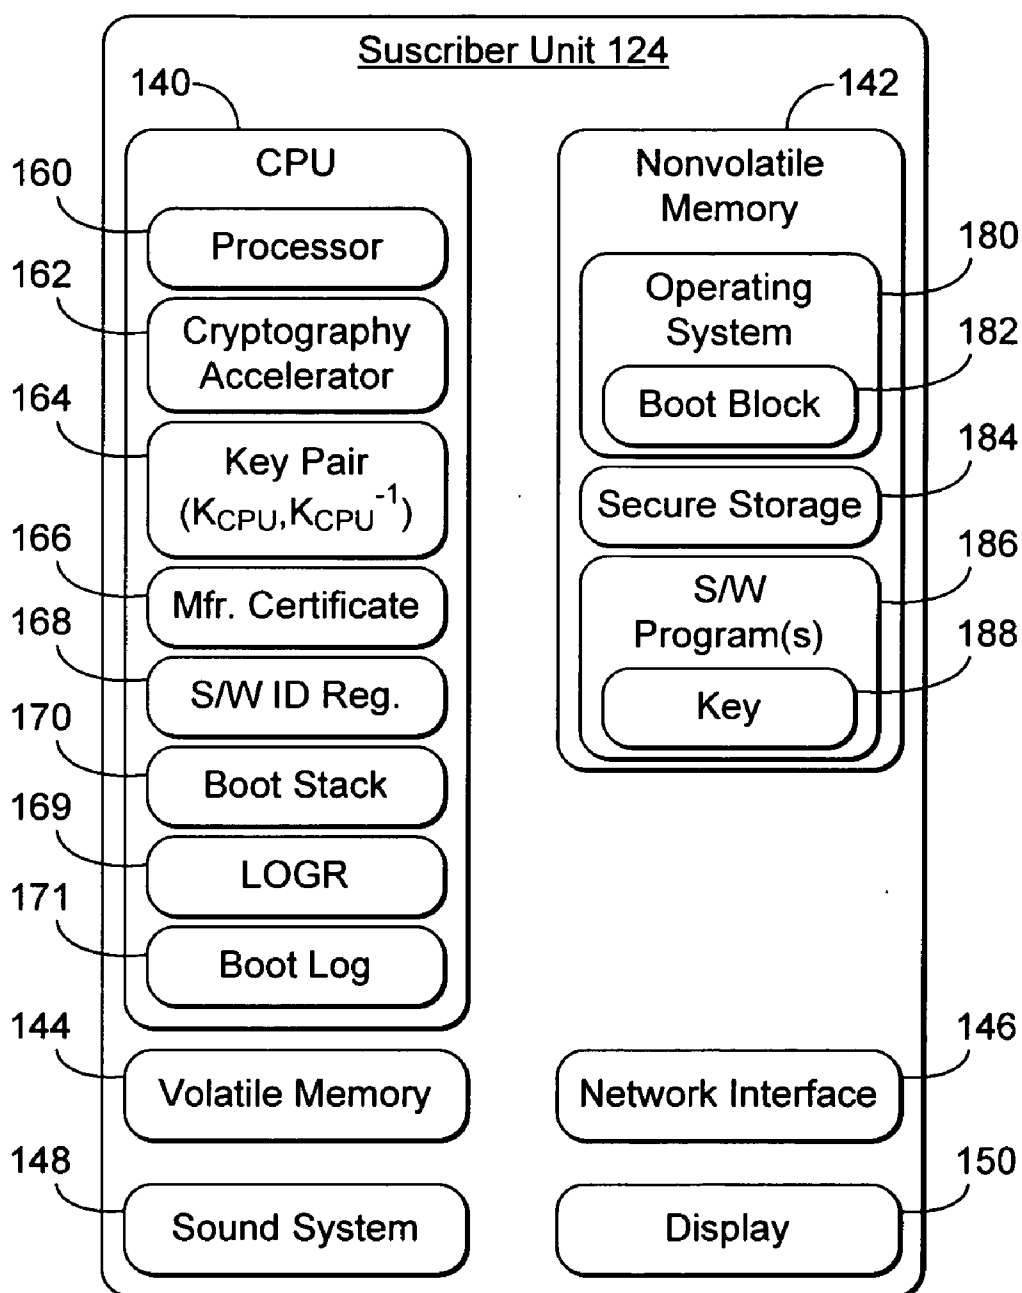



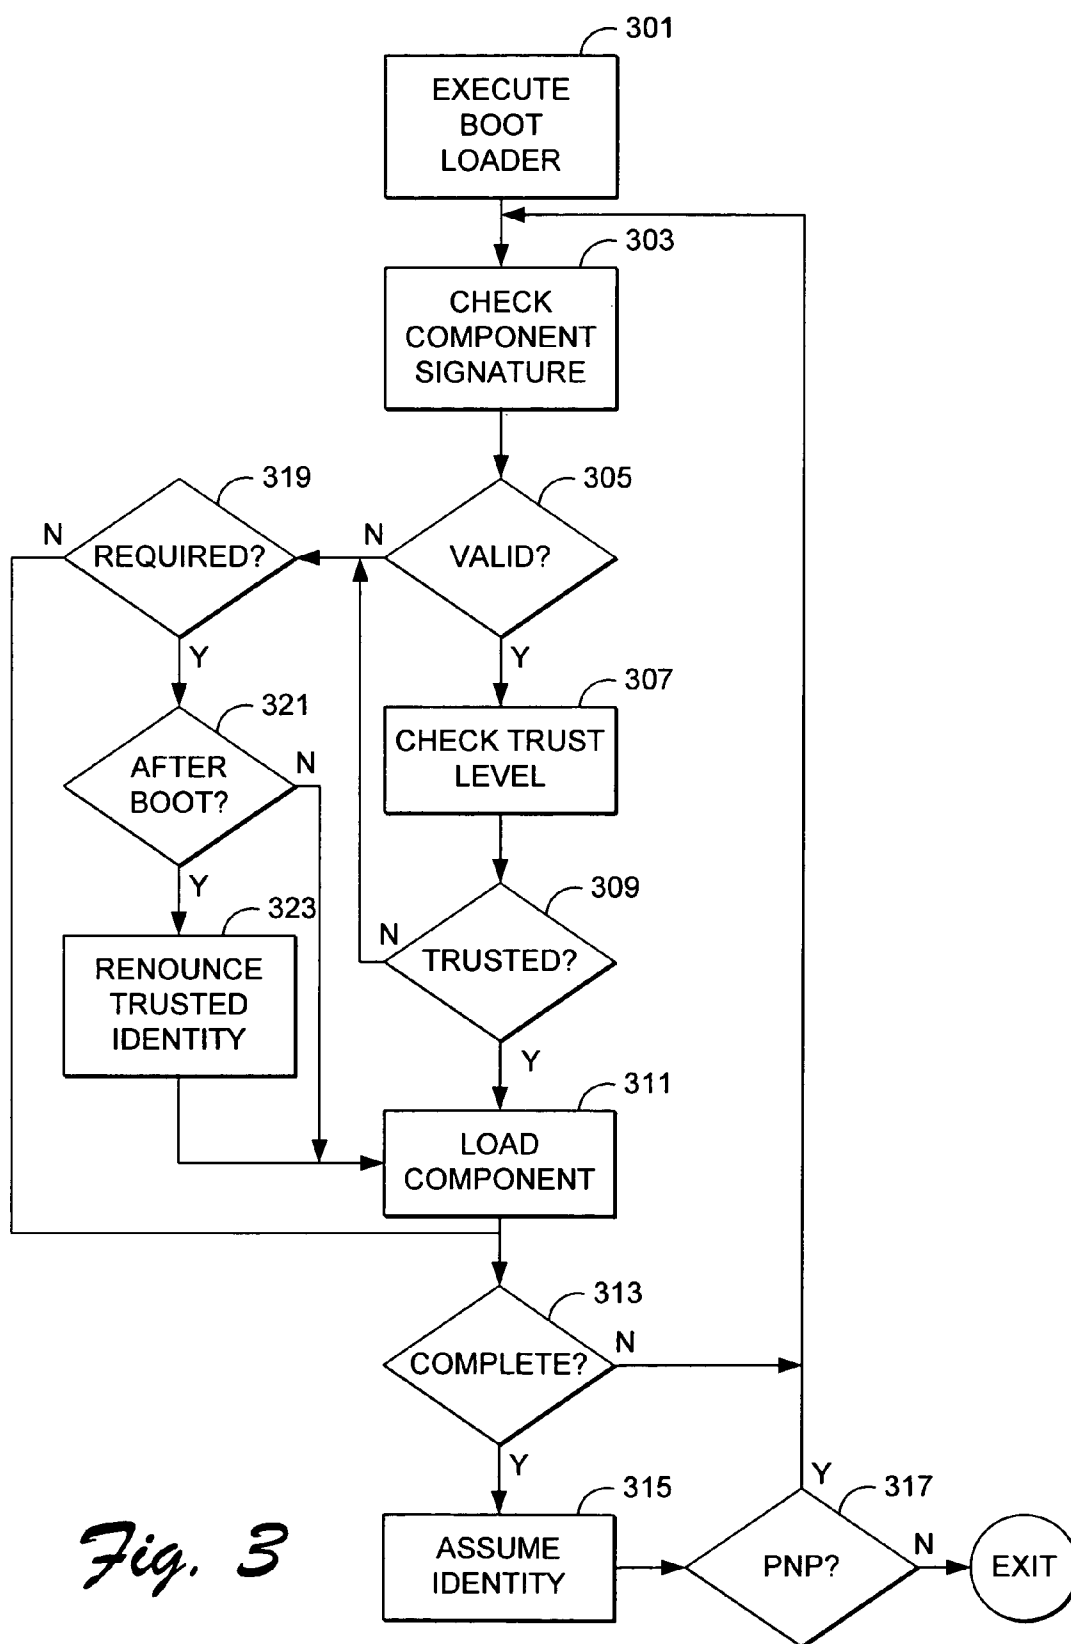
Fig. 1A



*Fig. 18*



*Fig. 2*



*Fig. 3*

400 →

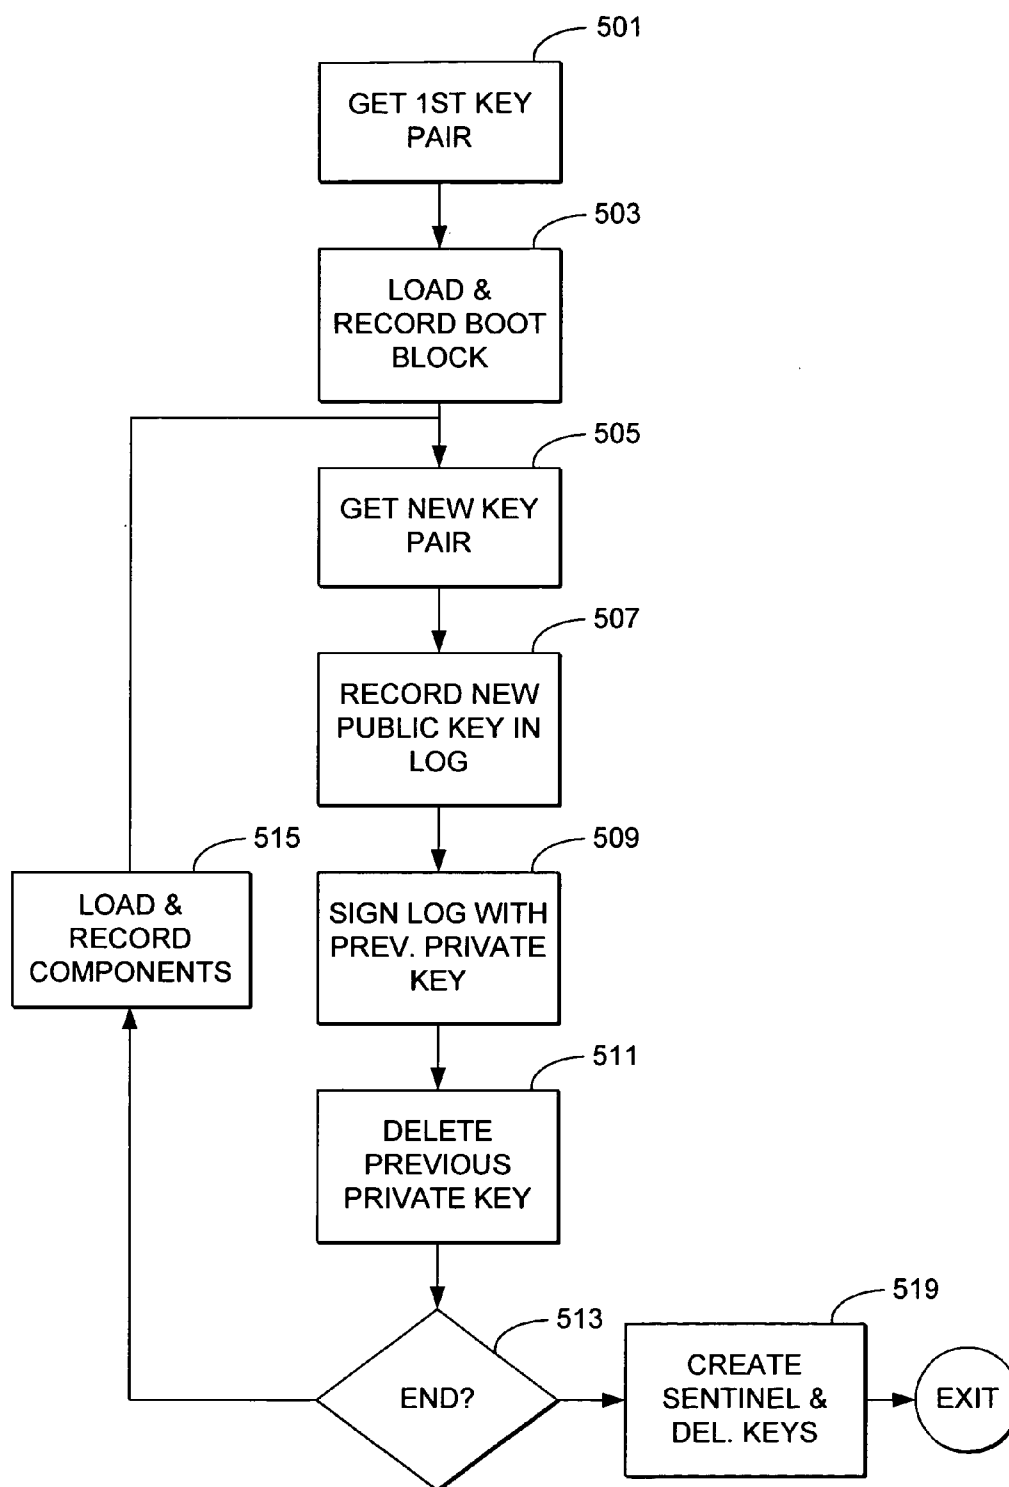
| 401 | 403<br>NAME | 405<br>VERSION | 407<br>SIGNER |
|-----|-------------|----------------|---------------|
|     |             |                |               |
|     |             |                |               |
|     |             |                |               |
|     |             |                |               |
|     |             |                |               |

*Fig. 4*


600 →

| 601 | 605<br>[BOOT BLOCK] | 606<br>$k_1$ | 607<br>$k_0^{-1}$ |
|-----|---------------------|--------------|-------------------|
| 601 | [COMPONENTS]        | $k_2$        | $k_1^{-1}$        |
| 601 | [COMPONENTS]        | $k_3$        | $k_2^{-1}$        |
|     |                     |              |                   |
|     |                     |              |                   |
|     |                     |              |                   |
| 601 | [COMPONENTS]        | $k_N$        | $k_{N-1}^{-1}$    |
| 603 | [Ø]                 | $k_N^{-1}$   |                   |



*Fig. 6*




*Fig. 5*



*Fig. 7A*



*Fig. 7B*



*Fig. 7C*









Fig. 8



*Fig. 9*




*Fig. 10*


1100 

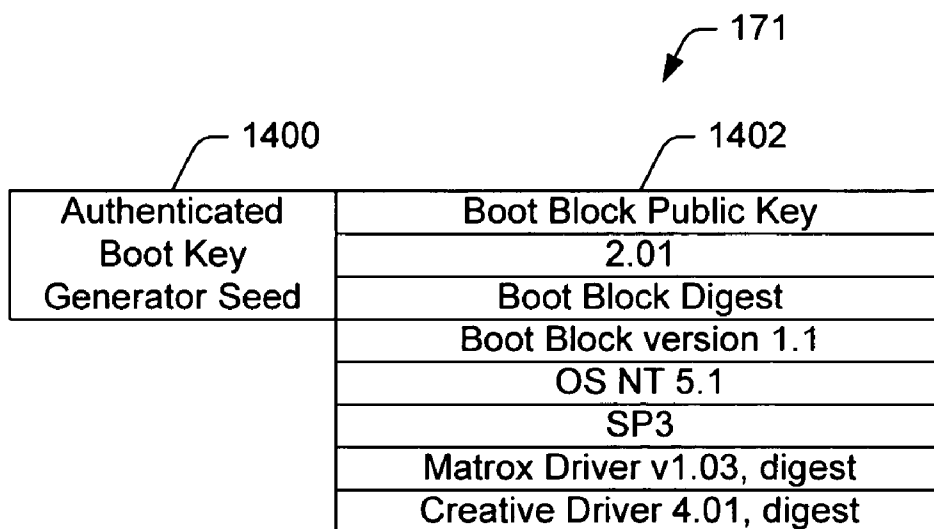
|      |                    |
|------|--------------------|
| 1101 | USAGE COUNTER      |
| 1103 | DERIVATION RIGHTS  |
| 1105 | EXPIRATION COUNTER |
| 1107 | SUBLICENSE RIGHTS  |
|      | ⋮                  |

*Fig. 11*



*Fig. 12*

*Fig. 13*



*Fig. 14*

## SYSTEM AND METHOD FOR SECURE STORAGE OF DATA USING A KEY

### RELATED APPLICATIONS

[0001] This application is a continuation of U.S. patent application Ser. No. 10/430,994, filed May 7, 2003, entitled "System and method for secure storage data using a key", which is hereby incorporated by reference herein, and which is a continuation of U.S. patent application Ser. No. 09/227,568, filed Jan. 8, 1999, entitled "Key-Based Secure Storage". U.S. patent application Ser. No. 09/227,568 is a continuation-in-part of U.S. provisional patent application Ser. No. 60/105,891 filed on Oct. 26, 1998, which is herein incorporated by reference, and is related to co-pending and co-filed U.S. patent application Ser. No. 09/266,207 titled "System and Method for Authenticating an Operating System to a Central Processing Unit, Providing the CPU/OS with Secure Storage, and Authenticating the CPU/OS to a Third Party", Ser. No. 09/227,611 titled "Loading and Identifying a Digital Rights Management Operating System", Ser. No. 09/227,559 titled "Digital Rights Management", and Ser. No. 09/227,561 titled "Digital Rights Management Operating System".

### FIELD OF THE INVENTION

[0002] This invention relates generally to computer operating systems, and more particularly to systems and methods for secure storage of data using a key.

### COPYRIGHT NOTICE/PERMISSION

[0003] A portion of the disclosure of this patent document contains material which is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure as it appears in the Patent and Trademark Office patent file or records, but otherwise reserves all copyright rights whatsoever. The following notice applies to the software and data as described below and in the drawings hereto: Copyright© 1998, Microsoft Corporation, All Rights Reserved.

### BACKGROUND OF THE INVENTION

[0004] More and more content is being delivered in digital form, and more and more digital content is being delivered online over private and public networks, such as Intranets, the Internet and cable TV networks. For a client, digital form allows more sophisticated content, while online delivery improves timeliness and convenience. For a publisher, digital content also reduces delivery costs. Unfortunately, these worthwhile attributes are often outweighed in the minds of publishers by the corresponding disadvantage that online information delivery makes it relatively easy to obtain pristine digital content and to pirate the content at the expense and harm of the publisher.

[0005] Piracy of digital content, especially online digital content, is not yet a great problem. Most premium content that is available on the Web is of low value, and therefore casual and organized pirates do not yet see an attractive business stealing and reselling content. Increasingly, though, higher-value content is becoming available. Books and audio recordings are available now, and as bandwidths increase, video content will start to appear. With the increase

in value of online digital content, the attractiveness of organized and casual theft increases.

[0006] The unusual property of digital content is that the publisher (or reseller) gives or sells the content to a client, but continues to restrict rights to use the content even after the content is under the sole physical control of the client. For instance, a publisher will typically retain copyright to a work so that the client cannot reproduce or publish the work without permission. A publisher could also adjust pricing according to whether the client is allowed to make a persistent copy, or is just allowed to view the content online as it is delivered. These scenarios reveal a peculiar arrangement. The user that possesses the digital bits often does not have full rights to their use; instead, the provider retains at least some of the rights.

[0007] "Digital rights management" is therefore fast becoming a central requirement if online commerce is to continue its rapid growth. Content providers and the computer industry must quickly provide technologies and protocols for ensuring that digital content is properly handled in accordance with the rights granted by the publisher. If measures are not taken, traditional content providers may be put out of business by widespread theft, or, more likely, will refuse altogether to deliver content online.

[0008] Traditional security systems ill serve this problem. There are highly secure schemes for encrypting data on networks, authenticating users, revoking certificates, and storing data securely. Unfortunately, none of these systems address the assurance of content security after it has been delivered to a client's machine. Traditional uses of smart cards offer little help. Smart cards merely provide authentication, storage, and encryption capabilities. Ultimately, useful content must be assembled within the host machine for display, and again, at this point the bits are subject to theft. Cryptographic coprocessors provide higher-performance cryptographic operations, and are usually programmable but again, fundamentally, any operating system or sufficiently privileged application, trusted or not, can use the services of the cryptographic processor.

[0009] There appear to be three solutions to this problem. One solution is to do away with general-purpose computing devices and use special-purpose tamper-resistant boxes for delivery, storage, and display of secure content. This is the approach adopted by the cable industry and their set-top boxes, and looks set to be the model for DVD-video presentation. The second solution is to use secret, proprietary data formats and applications software, or to use tamper-resistant software containers, in the hope that the resulting complexity will substantially impede piracy. The third solution is to modify the general-purpose computer to support a general model of client-side content security and digital rights management.

[0010] This invention is directed to a system and methodology that falls generally into the third category of solutions.

[0011] A fundamental building block for client-side content security is a secure operating system. If a computer can be booted only into an operating system that itself honors content rights, and allows only compliant applications to access rights-restricted data, then data integrity within the machine can be assured. This stepping-stone to a secure

operating system is sometimes called "Secure Boot." If secure boot cannot be assured, then whatever rights management system the secure OS provides, the computer can always be booted into an insecure operating system as a step to compromise it.

**[0012]** Secure boot of an operating system is usually a multi-stage process. A securely booted computer runs a trusted program at startup. The trusted program loads an initial layer of the operating system and checks its integrity (by using a code signature or by other means) before allowing it to run. This layer will in turn load and check the succeeding layers. This proceeds all the way to loading trusted (signed) device drivers, and finally the trusted application(s).

**[0013]** An article by B. Lampson, M. Abadi, and M. Burrows, entitled "Authentication in Distributed Systems: Theory and Practice," ACM Transactions on Computer Systems v10, 265, 1992, describes in general terms the requirements for securely booting an operating system. The only hardware assist is a register that holds a machine secret. When boot begins this register becomes readable, and there's a hardware operation to make this secret unreadable. Once it's unreadable, it stays unreadable until the next boot. The boot code mints a public-key pair and a certificate that the operating system can use to authenticate itself to other parties in order to establish trust. We note that in this scheme, a malicious user can easily subvert security by replacing the boot code.

**[0014]** Clark and Hoffman's BITS system is designed to support secure boot from a smart card. P. C. Clark and L. J. Hoffman, "BITS: A Smartcard Operating System," Comm. ACM. 37, 66, 1994. In their design, the smart card holds the boot sector, and PCs are designed to boot from the smart card. The smart card continues to be involved in the boot process (for example, the smart card holds the signatures or keys of other parts of the OS).

**[0015]** Bennet Yee describes a scheme in which a secure processor first gets control of the booting machine. B. Yee, "Using Secure Coprocessors", Ph.D. Thesis, Carnegie Mellon University, 1994. The secure processor can check code integrity before loading other systems. One of the nice features of this scheme is that there is a tamper-resistant device that can later be queried for the details of the running operating system.

**[0016]** Another secure boot model, known as AEGIS, is disclosed by W. Arbaugh, D. G. Farber, and J. M. Smith in a paper entitled "A Secure and Reliable Bootstrap Architecture", Univ. of Penn. Dept. of CIS Technical Report, IEEE Symposium on Security and Privacy, page 65, 1997. This AEGIS model requires a tamper-resistant BIOS that has hard-wired into it the signature of the following stage. This scheme has the very considerable advantage that it works well with current microprocessors and the current PC architecture, but has three drawbacks. First, the set of trusted operating systems or trusted publishers must be wired into the BIOS. Second, if the content is valuable enough (for instance, e-cash or Hollywood videos), users will find a way of replacing the BIOS with one that permits an insecure boot. Third, when obtaining data from a network server, the client has no way of proving to the remote server that it is indeed running a trusted system.

**[0017]** On the more general subject of client-side rights management, several systems exist or have been proposed to

encapsulate data and rights in a tamper-resistant software package. An early example is IBM's Cryptolope. Another existent commercial implementation of a rights management system has been developed by Intertrust. In the audio domain, AT&T Research have proposed their "A2b" audio rights management system based on the PolicyMaker rights management system.

**[0018]** Therefore, there is a need in the art for a digital rights management operating system that protects content downloaded from a provider while operating on a general purpose personal computer without the need of specialized or additional hardware.

## SUMMARY OF THE INVENTION

**[0019]** Systems and methods for secure storage of data using a key are described herein.

**[0020]** In accordance with one aspect, a data structure to be encrypted is received in a device, the data structure including content along with a statement of conditions under which the content may be decrypted. The content is encrypted using a key.

**[0021]** In accordance with another aspect, a data structure is decrypted using a key. A statement of conditions under which content in the data structure can be decrypted is obtained, and testing is performed as to whether the conditions are satisfied. The decrypted content is returned only if the conditions are satisfied.

**[0022]** In accordance with another aspect, a data structure to be encrypted in a device is obtained. The data structure includes content along with a statement of conditions under which the content may be decrypted. A seal operation is invoked to have the data structure encrypted using a symmetric key of a processor of the device.

**[0023]** In accordance with another aspect, an unseal operation is invoked in order to have a data block decrypted using a key. In response to invoking the unseal operation, the decrypted data block is received only if conditions under which content in the data block can be decrypted are satisfied.

**[0024]** In accordance with another aspect, the processor may contain a fixed per-processor symmetric key  $K_S$  which can be used to encrypt a data structure containing content along with a statement of the conditions under which it may be decrypted; key  $K_S$  is also used to decrypt the data structure, test the conditions, and either return the content or fail. Key  $K_S$  is to be used only for this pair of operations, which are referred to as "Seal" and "Unseal".

## BRIEF DESCRIPTION OF THE DRAWINGS

**[0025]** FIG. 1A is a diagram of the hardware and operating environment in conjunction with which exemplary embodiments of the invention may be practiced;

**[0026]** FIG. 1B is a diagram of a client computer for use with exemplary embodiments of the invention;

**[0027]** FIG. 2 is a diagram illustrating a system-level overview of an exemplary embodiment of the invention;

**[0028]** FIG. 3 is a flowchart of a method to be performed by a client when booting or loading system components according to an exemplary embodiment of the invention;

[0029] FIG. 4 is a diagram of a certificate revocation list data structure for use in an exemplary implementation of the invention;

[0030] FIG. 5 is a flowchart of a method to be performed by a client to create a boot log according to an exemplary embodiment of the invention;

[0031] FIG. 6 is a block diagram of an exemplary boot log created using the method of FIG. 5;

[0032] FIGS. 7A, 7B and 7C are block diagrams of boot blocks for use in an exemplary embodiment of the invention;

[0033] FIG. 8 is a block diagram of key generation functions according to an exemplary embodiment of the invention;

[0034] FIG. 9 is a diagram of a rights manager certificate data structure for use in an exemplary implementation of the invention;

[0035] FIG. 10 is a diagram of a required properties access control list data structure for use in an exemplary implementation of the invention; and

[0036] FIG. 11 is a diagram of a license data structure for use in an exemplary implementation of the invention.

[0037] FIG. 12 shows an example of a signed boot block created by signing a block of code.

[0038] FIG. 13 shows steps in a method for performing an authenticated boot operation on an operating system.

[0039] FIG. 14 shows an exemplary structure of a boot log.

## DETAILED DESCRIPTION OF THE INVENTION

[0040] In the following detailed description of exemplary embodiments of the invention, reference is made to the accompanying drawings, which form a part hereof, and in which is shown by way of illustration specific exemplary embodiments in which the invention may be practiced. These embodiments are described in sufficient detail to enable those skilled in the art to practice the invention, and it is to be understood that other embodiments may be utilized and that logical, mechanical, electrical and other changes may be made without departing from the spirit or scope of the present invention. The following detailed description is, therefore, not to be taken in a limiting sense, and the scope of the present invention is defined only by the appended claims.

### Hardware and Operating Environment

[0041] FIG. 1A is a diagram of the hardware and operating environment in conjunction with which embodiments of the invention may be practiced. The description of FIG. 1A is intended to provide a brief, general description of suitable computer hardware and a suitable computing environment in conjunction with which the invention may be implemented. Although not required, the invention is described in the general context of computer-executable instructions, such as program modules, being executed by a computer, such as a personal computer. Generally, program modules include

routines, programs, objects, components, data structures, etc. that perform particular tasks or implement particular abstract data types.

[0042] Moreover, those skilled in the art will appreciate that the invention may be practiced with other computer system configurations, including hand-held devices, multi-processor systems, microprocessor-based or programmable consumer electronics, network PCs, minicomputers, mainframe computers, and the like. The invention may also be practiced in distributed computing environments where tasks are performed by remote processing devices that are linked through a communications network. In a distributed computing environment, program modules may be located in both local and remote memory storage devices.

[0043] The exemplary hardware and operating environment of FIG. 1A for implementing the invention includes a general purpose computing device in the form of a computer 20, including a processing unit 21, a system memory 22, and a system bus 23 that operatively couples various system components, including the system memory 22, to the processing unit 21. There may be only one or there may be more than one processing unit 21, such that the processor of computer 20 comprises a single central-processing unit (CPU), or a plurality of processing units, commonly referred to as a parallel processing environment. The computer 20 may be a conventional computer, a distributed computer, or any other type of computer; the invention is not so limited.

[0044] The system bus 23 may be any of several types of bus structures including a memory bus or memory controller, a peripheral bus, and a local bus using any of a variety of bus architectures. The system memory may also be referred to as simply the memory, and includes read only memory (ROM) 24 and random access memory (RAM) 25. A basic input/output system (BIOS) 26, containing the basic routines that help to transfer information between elements within the computer 20, such as during start-up, is stored in ROM 24. The computer 20 further includes a hard disk drive 27 for reading from and writing to a hard disk, not shown, a magnetic disk drive 28 for reading from or writing to a removable magnetic disk 29, and an optical disk drive 30 for reading from or writing to a removable optical disk 31 such as a CD ROM or other optical media.

[0045] The hard disk drive 27, magnetic disk drive 28, and optical disk drive 30 are connected to the system bus 23 by a hard disk drive interface 32, a magnetic disk drive interface 33, and an optical disk drive interface 34, respectively. The drives and their associated computer-readable media provide nonvolatile storage of computer-readable instructions, data structures, program modules and other data for the computer 20. It should be appreciated by those skilled in the art that any type of computer-readable media that can store data that is accessible by a computer, such as magnetic cassettes, flash memory cards, digital video disks, Bernoulli cartridges, random access memories (RAMs), read only memories (ROMs), and the like, may be used in the exemplary operating environment.

[0046] A number of program modules may be stored on the hard disk, magnetic disk 29, optical disk 31, ROM 24, or RAM 25, including an operating system 35, one or more application programs 36, other program modules 37, and program data 38. A user may enter commands and information into the personal computer 20 through input devices



such as a keyboard 40 and pointing device 42. 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 21 through a serial port interface 46 that is coupled to the system bus, but may be connected by other interfaces, such as a parallel port, game port, or a universal serial bus (USB). A monitor 47 or other type of display device is also connected to the system bus 23 via an interface, such as a video adapter 48. In addition to the monitor, computers typically include other peripheral output devices (not shown), such as speakers and printers.

[0047] The computer 20 may operate in a networked environment using logical connections to one or more remote computers, such as remote computer 49. These logical connections are achieved by a communication device coupled to or a part of the computer 20; the invention is not limited to a particular type of communications device. The remote computer 49 may be another computer, a server, a router, a network PC, a client, a peer device or other common network node, and typically includes many or all of the elements described above relative to the computer 20, although only a memory storage device 50 has been illustrated in FIG. 1. The logical connections depicted in FIG. 1 include a local-area network (LAN) 51 and a wide-area network (WAN) 52. Such networking environments are commonplace in offices, enterprise-wide computer networks, intranets and the Internet.

[0048] When used in a LAN-networking environment, the computer 20 is connected to the local network 51 through a network interface or adapter 53, which is one type of communications device. When used in a WAN-networking environment, the computer 20 typically includes a modem 54, a type of communications device, or any other type of communications device for establishing communications over the wide area network 52, such as the Internet. The modem 54, which may be internal or external, is connected to the system bus 23 via the serial port interface 46. In a networked environment, program modules depicted relative to the personal computer 20, or portions thereof, may be stored in the remote memory storage device. It is appreciated that the network connections shown are exemplary and other means of and communications devices for establishing a communications link between the computers may be used.

[0049] The hardware and operating environment in conjunction with which embodiments of the invention may be practiced has been described. The computer in conjunction with which embodiments of the invention may be practiced may be a conventional computer, a distributed computer, or any other type of computer; the invention is not so limited. Such a computer typically includes one or more processing units as its processor, and a computer-readable medium such as a memory. The computer may also include a communications device such as a network adapter or a modem, so that it is able to communicatively couple to other computers.

[0050] One exemplary embodiment of a suitable client computer is described in the related application titled "System and Method for Authenticating an Operating System to a Central Processing Unit, Providing the CPU/OS with Secure Storage, and Authenticating the CPU/OS to a Third Party," and illustrated in FIG. 1B as subscriber unit 124. The CPU 140 in the subscriber unit 124 is able to authenticate the

identity of the boot block and OS components that have been loaded into the computer, and to provide quoting and secure storage operations based on this identity as briefly described next. Full descriptions of various embodiments for the subscriber unit 124 are provided in the related application.

[0051] FIG. 1B shows general components in the subscriber unit 124. They include a central processing unit (CPU) 140, nonvolatile memory 142 (e.g., ROM, disk drive, CD ROM, etc.), volatile memory 144 (e.g., RAM), and a network interface 146 (e.g., modem, network port, wireless transceiver, etc.). The subscriber unit 124 may also include a sound system 148 and/or a display 150. These components are interconnected via conventional busing architectures, including parallel and serial schemes (not shown).

[0052] The CPU 140 has a processor 160 and also can have a cryptographic accelerator 162. The CPU 140 is capable of performing cryptographic functions, such as signing, encrypting, decrypting, and authenticating, with or without the accelerator 162 assisting in intensive mathematical computations commonly involved in cryptographic functions.

[0053] The CPU manufacturer equips the CPU 140 with a pair of public and private keys 164 that is unique to the CPU. For discussion purpose, the CPU's public key is referred to as " $K_{CPU}$ " and the corresponding private key is referred to as " $K_{CPU}^{-1}$ ". Other physical implementations may include storing the key on an external device to which the main CPU has privileged access (where the stored secrets are inaccessible to arbitrary application or operating systems code). The private key is never revealed and is used only for the specific purpose of signing stylized statements, such as when responding to challenges from a content provider, as is discussed below.

[0054] The CPU manufacturer may further embed a second secret key  $K_2$  in the CPU 140 or other secure hardware. The second key is distinct from the first key pair, and is used to generate a secure storage key, as is described below. Alternatively, as described below, a symmetric key  $K_S$  may be used with "Seal" and "Unseal" operations to encrypt a data structure along with a statement of the conditions under which the data structure may be decrypted.

[0055] The manufacturer also issues a signed certificate 166 testifying that it produced the CPU according to a known specification. Generally, the certificate testifies that the manufacturer created the key pair 164, placed the key pair onto the CPU 140, and then destroyed its own knowledge of the private key " $K_{CPU}^{-1}$ ". In this way, only the CPU knows the CPU private key  $K_{CPU}^{-1}$ ; the same key is not issued to other CPUs and the manufacturer keeps no record of it. The certificate can in principle be stored on a separate physical device associated with the processor but still logically belongs to the processor with the corresponding key.

[0056] The manufacturer has a pair of public and private signing keys,  $K_{MFR}$  and  $K_{MFR}^{-1}$ . The private key  $K_{MFR}^{-1}$  is known only to the manufacturer, while the public key  $K_{MFR}$  is made available to the public. The manufacturer certificate 166 contains the manufacturer's public key  $K_{MFR}$ , the CPU's public key  $K_{CPU}$ , and the above testimony. The manufacturer signs the certificate using its private signing key,  $K_{MFR}^{-1}$ , as follows:

Mfr. Certificate = ( $K_{MFR}$ , Certifies-for-Boot,  $K_{CPU}$ ),  
signed by  $K_{MFR}^{-1}$

[0057] The predicate “certifies-for-boot” is a pledge by the manufacturer that it created the CPU and the CPU key pair according to a known specification. The pledge further states that the CPU can correctly perform authenticated boot procedures, as are described below in more detail. The manufacturer certificate **66** is publicly accessible, yet it cannot be forged without knowledge of the manufacturer’s private key  $K_{MFR}^{-1}$ .

[0058] The CPU **140** has an internal software identity register (SIR) **168**, which contains the identity of an authenticated operating system **180** or a predetermined false value (e.g., zero) if the CPU determines that the operating system **180** cannot be authenticated. The operating system (OS) **180** is stored in the memory **142** and executed on the CPU **140**. The operating system **180** has a block of code **182** that is used to authenticate the operating system to the CPU during the boot operation. The boot block **182** uniquely determines the operating system, or class of operating systems (e.g. those signed by the same manufacturer). The boot block **182** can also be signed by the OS manufacturer.

[0059] Another implementation in which a ‘chain of certificates’ leading back to a root certificate held by the processor manufacturer is also acceptable.

[0060] The CPU **140** has an internal software identity register (SIR) **168**, which is cleared at the beginning of every boot. The CPU executes an opcode “BeginAuthenticated-Boot” or “BAB” to set an identity of a corresponding piece of software, such as operating system **180**, and stores this identity in the SIR; the boot block of the operating system (described below) is atomically executed as part of the BAB instruction. If execution of the BAB opcode and the boot block fails (e.g., if the execution was not atomic), the SIR **168** is set to a predetermined false value (e.g., zero). This process is described below in more detail under the heading “Authenticated Boot”.

[0061] The CPU **140** also utilizes a second internal register (LOGR) **169**, which holds contents produced as a result of running a LOG operation. This operation, as well as the register, is described below in more detail.

[0062] The CPU **140** also maintains a “boot log” **171** to track software modules and programs that are loaded. In one implementation, the boot log **171** is a log in an append-only memory of the CPU that is cleared at the beginning of every boot. Since it consumes only about a few hundred bytes, the boot log **171** can be comfortably included in the main CPU. Alternatively, the CPU **140** can store the boot log **171** in volatile memory **144** in a cryptographic tamper-resistant container.

[0063] A further implementation is by means of a software module that allows each section of the booting operating system to write entries into the boot log that cannot be removed by later components without leaving evidence of tampering. Yet alternatively, the SIR can hold a cryptographic digest of a data structure comprising the initial boot block and the subsequent contents of the boot log. The operation of appending to the boot log (call this operation “Extend”) replaces the SIR with the hash of the concatenation of the SIR and the entry being appended to the boot log. A straightforward implementation of this operation may be seen to modify the SIR, potentially disallowing future “Unseal” operations that depend on the value of the SIR.

Note, however, that the operating system, when booting, can choose to add elements to the boot log without loading the corresponding components, and so a more privileged combination of software components can impersonate a less privileged one. This allows the controlled transfer of secrets across privilege levels. In this approach, software will keep its own plaintext copy of the boot log entries, along with the initial value of the SIR following boot, and this plaintext copy is validated by knowledge of the current composite SIR.

[0064] As an optimization, regardless of the implementation of the boot log, the OS may choose not to extend the boot log with the identities of certain software components, if these components are judged to be as trustworthy as the OS itself, or if they will execute only in a protected environment from which they will be unable to subvert operation.

[0065] FIG. 12 shows an example of a signed boot block **190** created by signing the block of code **182**. It contains the BeginAuthenticatedBoot opcode **192**, a length **194** specifying the number of byte in the block of code, the code **182**, a signature **196**, and a public key **198** used to verify the signature **196**. The boot block will also contain as a constant or set of constants, keys, or other information **199** that is used to validate the subsequent operating system components (for instance a public key or keys). In this implementation, the CPU will set the SIR to the public key of the boot block, but only if the boot block code signature is correct for the stated boot block public key.

[0066] In an alternative implementation, the SIR is set to the cryptographic hash or digest of the code and constants that make up the boot block. The signature **196** and public key **198** are then not needed.

[0067] A key observation of both of these implementations is that no one can boot an untrusted operating system in which the SIR is set to the value of a trusted operating system.

[0068] Once booted the operating system **180** and the applications named in the license or ACL by the content provider can set aside space **184** in memory or disk **142** to hold the digital content from the content provider in a secure manner, without fear of other operating systems or rogue applications reading the data in the space. The persistent content is protected by encryption using a key that is generated based in part upon a seed supplied by an authenticated and trusted OS, in part by a secret key stored in the CPU, and in part by the software identity register (SIR). (Alternatively, the persistent content is stored using the “Seal” and “Unseal” operations, described below in more detail, or using the processor’s public key pair for encryption.) The persistent content is stored with a license or ACL naming the applications that can use the content and the terms under which they can use it.

[0069] Software programs **186** (the applications) are also shown stored in memory **142**. These programs may be used to render or otherwise play the content. Each program **186** has an associated key or digest **188** for unique identification.

#### Authenticated Boot

[0070] Traditional approaches to secure boot attempt to secure the BIOS or other loader, and have the BIOS check

later components before allowing them to execute. In contrast to this traditional approach, the authenticated boot process allows any software at any point in the boot sequence to initiate an authenticated boot.

[0071] FIG. 13 shows steps in a method for performing an authenticated boot operation on the operating system 180. These steps are performed by the CPU 140 and OS 180 resident in the subscriber unit 124. At step 1300, the CPU executes the BeginAuthenticatedBoot opcode 192 in the signed boot block 190 to set an identity for the operating system 180. The identity can be a digest of the boot block's opcodes and data, or the public key 198 corresponding to a signature on the boot block of the operating system.

[0072] The BeginAuthenticatedBoot opcode 192 and the boot block 190 execute as one atomic operation, with the implication that if they execute completely and correctly, the resulting operating system can be trusted. Measures are taken to ensure that the CPU is not interrupted and that the boot code that has just been validated cannot be modified. This can involve locking the memory bus and switching off interrupts. It could also involve having the CPU watch for interrupts or for writes by other bus agents and invalidate the authenticated boot sequence if they occur. The BAB opcode 192 can be executed at any time, with one exemplary time being at the start of the OS loader, right after the OS-selector executes. An alternative implementation is to provide both a BeginAuthenticatedBoot (BAB) and an EndAuthenticatedBoot (EAB) instruction. The BAB instruction computes the secure hash of the boot block and the EAB instruction sets the SIR if the execution of the boot block was not interrupted or potentially modified by memory writes from another processor or another bus master.

[0073] Execution of the BeginAuthenticatedBoot opcode 192 sets the internal software identity register 168 to either (1) the OS's identity (i.e., boot block digest or OS public key 198) if the operation is successful, or (2) zero if some event or circumstance has potentially subverted operation. Assuming the operation is successful (i.e., the "yes" branch from step 1302), the SIR 168 is now a unique number or other value that represents the identity of the operating system 180 (step 1304). Any two processors running the same operating system will produce the same SIR. If the BAB opcode operation is unsuccessful (i.e., the "no" branch from step 1302), the SIR is set to zero (step 1306).

[0074] It is noted that different operating systems may be serially booted on the subscriber unit 124. Executing the BAB opcode 192 for different signed OS boot blocks results in different SIR values. However, it is possible for multiple boot blocks to result in the same SIR, when desired.

[0075] At step 1310, the CPU 140 fills the first entry on the boot log 171 with the public key (or digest) of the boot block 182. From now on, any running code can append data to the boot log 171, and it is generally used by code in the boot chain to identify code versions as they are loaded and executed. As noted earlier, appending data to the boot log can be simulated by modifying the SIR via the "Extend" operation.

[0076] The boot block 182 is free to load the next set of blocks in the boot-chain (step 1312). At step 1314, the boot block 182 checks the validity of the modules (by signature or other means) and loads them so that they can be executed.

An identity for each module is appended to the boot log 171. The OS will also retain additional information on components that it loads (e.g., version numbers, device driver IDs, etc.). Loading and executing the code may result in loading more code, validating it, and executing it, etc. This process continues through to the loading of device drivers. When the boot sequence is complete, the OS is operational and the software identity register and the boot log store non-modifiable data captured during the boot sequence. We can recommence loading new device drivers at any point, possibly causing the operating system to become less privileged, with the possible termination of access to protected content.

[0077] The CPU can generate a signed certificate containing the boot log data to attest to the particular operating system (including drivers) that is running. It could also generate a signed statement containing just the SIR. FIG. 14 shows an exemplary structure of a boot log 171. It contains a seed field 1400 and a block ID field 1402. The block ID field 1402 holds identities of the blocks of code that are loaded and verified on the subscriber unit. The block ID field 1402 can hold text or binary data.

[0078] The SIR or the seed field 1400 holds an authenticated boot key generator seed. The CPU uses the seed in field 1400 to generate keys unique to the OS and processor. Since the first entry of the boot log 171 can only be generated by the execution of a particular boot block or the holder of the boot block private key, the keys can only be re-generated by the same OS, or another OS from the same publisher under control of the publisher. OS-specific key generation provides a building block for secure persistent storage of data and the continued enforcement of digital usage rights even if the computer is physically compromised, or the computer is booted into another operating system. Use of OS-specific storage keys for secure storage is described below in more detail.

[0079] Alternatively, the processor may use the "Seal" and "Unseal" instructions to store persistent protected content, or when possible may encrypt it with the processor's public key and decrypt it with the "Unseal" instruction, which is called "Reveal" when used with public keys. These operations are described below in more detail under the heading "Secure Storage".

#### System Level Overview

[0080] A system level overview of the operation of an exemplary embodiment of the invention is described by reference to FIG. 2. A subscriber computer 200, such as client computer 20 in FIG. 1, is connected to a content provider server computer 220, such as remote computer 49, through a wide-area network, such as WAN 52. Processes performed by the components of the subscriber computer 200 and the content provider 200 are illustrated by arrows in FIG. 2. Many of these processes incorporate either public/private key pairs, digital signatures, digital certificates, and/or encryption algorithms, or a combination of these standard cryptographic functions. Such functions are assumed to be provided by the CPU of the subscriber computer in the descriptions that follow, but can be provided by other well-known cryptographic mechanisms as will be immediately understood by one skilled in the art.

[0081] The content may be essentially any type of content that can be expressed as digital data, including video, still

pictures, audio, graphical images, and textual data or executable content (computer programs). Examples of possible content include feature-length movies, TV shows, games, software programs, news, stock information, weather reports, art, photographs, and so on.

[0082] To prevent their content from being stolen or misused, content providers will download content only to known software, and therefore only to subscriber computers that can prove that their operating systems will enforce the limitations the provider places on the content. Such a digital rights management operating system (DRMOS) must load and execute only OS components that are authenticated as respecting digital rights ("trusted"), and must allow access to the downloaded content by only similarly trusted applications.

[0083] The first requirement is met in the exemplary embodiment of the invention by having all trusted operating system-level components digitally signed by their developers or a trusted third-party, with the signature acting as a guarantee that the components respect digital rights. The signature is validated before the component is loaded. The resulting DRMOS is assigned a unique trusted identity, as explained in detail below, which is recorded in an internal register in the CPU, such as SIR 168 in FIG. 1B. FIG. 2 illustrates a DRMOS 205, with its identity 206, after it has been loaded into the CPU 201 of a subscriber computer 200 through such a loading process 1.

[0084] The second requirement has two aspects. First, trusted applications must be identified in some fashion, and, second, the DRMOS must prevent non-trusted applications from gaining access to the content when it is stored, either permanently or temporarily, on the subscriber computer.

[0085] In the exemplary embodiment shown in FIG. 2, a trusted application 209 has agreed to operate in accordance with the limitations placed on content by a provider. The trusted application 209 is identified through a "rights manager" certificate 210. In one embodiment, the rights manager certificate 210 extends a standard digital certificate, which includes such items as date of publication and name of the application, by adding a list of services, or properties, provided by the application, i.e., content type handled, version of the application, whether it saves content to disk, etc. For purposes of the exemplary embodiment shown in FIG. 2, the certificate 210 also identifies the trusted application; alternate mechanisms for identifying a trusted application are described later in the methods section.

[0086] The DRMOS 205 provides key-secured storage for permanently stored content to prevent unauthorized access to the content. For temporarily stored content, the DRMOS 205 prevents an untrusted process from reading the memory holding the content. These and other safeguards are also described in detail below. The permanent and temporary storage within subscriber computer 200 are collectively represented by device 203, which is illustrated in FIG. 2 as a disk drive. Such illustration is not intended to limit the range of devices that can serve as secured storage for a DRMOS.

[0087] Turning now to the remainder of the processes depicted in FIG. 2, application 209 requests 2 the download of content 221 from provider 220. The DRMOS 205 sends a message 3 to the provider 220 requesting the content 221.

The content provider 220 transmits a challenge message 4 to the DRMOS 205 asking for the identity of the CPU 201, the DRMOS 205, and the application 209. The DRMOS 205 transmits a response message 5 containing a certificate 202 for the CPU 201, its own identity 206, and the rights manager certificate 210 for the application 209.

[0088] The challenge-response process follows the common protocol for such interchanges, the difference being only in the data exchanged between the subscriber computer and the content provider. In one exemplary embodiment of a suitable challenge-response process described in the related application titled "System and Method for Authenticating an Operating System to a Central Processing Unit, Providing the CPU/OS with Secure Storage, and Authenticating the CPU/OS to a Third Party," the certificate 202 contains the challenge message 3, the identity of the DRMOS 206, the public key of the CPU 201, and data representing all software components that are currently loaded and executing on the subscriber computer 200. The certificate 202 is signed using the private key of the CPU 201. The content provider 220 examines the CPU certificate 202, the DRMOS identity 206, and the properties specified in the rights manager certificate 210 to determine whether it should establish a trust relationship with the DRMOS 205 on the subscriber computer 200.

[0089] In an alternate exemplary embodiment, the challenge-response protocol runs over a secure connection such as SSL (Secure Socket Layer) or TLS (Transport Level Security), which relies on a session key to encrypt the data transferred between the subscriber computer 200 and the content provider 220. This stops an attacker (such as the legitimate owner of the machine) from rebooting the PC into a different operating system after the DRMOS has authenticated itself, or using a different computer on the network for snooping on the data destined for the DRMOS.

[0090] If the trust relationship is established, the provider downloads 6 the content 221, an access predicate 222, and a "license" 223 to the DRMOS 205 on the subscriber computer 200. The access predicate 222 specifies the properties that an application must have in order to process the content 221, such as read-only or minimum/maximum video resolution. The access predicate 222 may also specify specific applications or families of applications allowed to process the content 221. The license 223 places restrictions on the use of the content 221 by an approved application, such as the number of times the content can be accessed or what derivative use can be made of the content. A media server of the content provider may be configured to download the entire content as a file, or to stream the content continuously over the network. As an example, the content provider may implement a server computer system comprising one or clustered server computers that handle requests from subscribers, manage the digital files locally, and facilitate delivery of requested digital files over a network to the subscriber 200.

[0091] When the DRMOS 205 receives the content 221, the access predicate 222 and the license 223, it determines whether the content should be permanently stored in a key-secured storage. If so, it requests an application storage key from the CPU 201. In the present example, the application storage key is specific to the application 209 that requested the content 221. The content 221 and the license

**223** are encrypted using the application storage key and the access predicate **222** is attached to the encrypted information. If the content **221** is to be stored only temporarily, the **DRMOS 205** places various safeguards around the memory area holding the content so that the content cannot be accessed by an untrusted application. The generation of application storage keys and the memory safeguards are described in detail below.

[**0092**] Each time application **209** wants to access the stored content **221**, it passes its rights manager certificate **210** and the appropriate application storage key (action **8**) to the **DRMOS 205**. The **DRMOS 205** validates the key and compares the rights manager certificate **210** against the access predicate **222**. Assuming the storage key is authenticated and the rights manager certificate **210** satisfies the access predicate **222**, the content **221** and the license **223** are decrypted. The **DRMOS** determines if the application's use of the content is permitted under the license **223** and allows access **9** if it is.

[**0093**] The system level overview of the operation of an exemplary embodiment of the invention has been described in this section of the detailed description. A series of processes and data structures on a subscriber computer control the loading of a digital rights management operating system, identify the **DRMOS** and trusted applications to a content provider, and secure content downloaded by the provider to the subscriber computer. While the invention is not limited to any particular hardware and software, for sake of clarity only a minimal hardware and software configuration necessary to process multimedia has been assumed for the subscriber computer.

#### Methods of Exemplary Embodiments of the Invention

[**0094**] In the previous section, a system level overview of the operation of exemplary embodiments of the invention was described. In this section, the particular methods performed by a subscriber computer, or client, of such exemplary embodiments are described by reference to a series of flowcharts and operational diagrams. The methods to be performed by the client constitute computer programs made up of computer-executable instructions. Describing the methods by reference to flowcharts and operational diagrams enables one skilled in the art to develop such programs including such instructions to carry out the methods on suitable computerized clients (e.g., on the processor of a client executing the instructions from computer-readable media). Data structures necessary to perform the methods are also described in this section. The methods of the content provider server computer are described to complete the understanding of the methods performed by the client.

[**0095**] Although many of the methods are interrelated, they have been divided into four groups to facilitate understanding. The boot/load process and various mechanisms for creating identities for different versions of a digital rights management operating system (**DRMOS**) are first described. The functions that must be provided by the **DRMOS** to ensure the enforcement of the content providers' rights are described next. The third group consists of methods directed toward providing permanent storage of the content on the subscriber computer once downloaded, and protecting that content from unauthorized access. Finally, the identification of trusted applications and the rights management functions are described.

#### [**0096**] Booting/Loading and Identifying the **DRMOS**

[**0097**] Referring first to **FIG. 3**, a flowchart of a method to be performed by a subscriber computer according to an exemplary embodiment of the invention is shown. This method is inclusive of the acts required to be taken by the computer to boot a **DRMOS** or to load additional components after the boot process is complete. Exemplary embodiments of boot block data structures are described below in conjunction with **FIGS. 7A-C**.

[**0098**] Shortly after a computer is turned on or is reset, a small program called a boot loader is executed by the CPU (block **301**). The boot loader loads a boot block for a particular operating system. Code in the boot block then loads various drivers and other software components necessary for the operating system to function on the computer. The totality of the boot block and the loaded components make up the identity of the operating system.

[**0099**] For a **DRMOS**, that identity can be trusted only if the boot block and the loaded components are trusted. In the embodiments described herein, all components are signed by a trusted source and provided with a rights manager certificate. An exemplary embodiment of the rights manager certificate is described below in conjunction with **FIG. 9**.

[**0100**] The operating system checks the signature of a component before loading it (block **303**). If the signature is valid (block **305**), the component has not been compromised by someone attempting to circumvent the boot process and the process proceeds to check the level of trust assigned to the component (block **307**). If the signature is not valid (or is there is no signature) but the component must be loaded (block **309**), the operating system will not assume the identity of a **DRMOS** upon completion of the boot process as explained further below.

[**0101**] A plug-and-play operating system provides an environment in which devices and their supporting software components can be added to the computer during normal operation rather than requiring all components be loaded during the boot process. If the device requires the loading of an untrusted component after the boot process completes, a plug-and-play **DRMOS** must then "renounce" its trusted identity and terminate any executing trusted applications (block **323**) before loading the component. The determination that an untrusted component must be loaded can be based on a system configuration parameter or on instructions from the user of the computer.

[**0102**] Assuming the signature is valid (block **305**) and the component is trusted (block **309**), it is loaded (block **311**). The trustworthiness of a component can be decided using various criteria. In one embodiment, only components provided by the operating system developer are trusted. At the other end of the scale, in another embodiment, all components are assumed trustworthy by the **DRMOS**, leaving the final decision to the content provider as described in more detail below. Still a third alternate embodiment provides that components signed by any of a select number of entities can be considered as equivalent to components provided by the **DRMOS** developer. In this embodiment, the identity of the resulting operating system is considered equivalent to the "pure" **DRMOS** provided by the **DRMOS** developer. The content provider decides whether it trusts the equivalent operating system.

[0103] Furthermore, not all versions of a component may be trusted. Because the rights manager certificate contains the version number of the component, it can be used to verify the trust level of a particular version. One embodiment of the loading process checks a component certification revocation list (CRL) to determine whether a component signature has been revoked. The CRL can be provided by the content provider or the DRMOS developer. An exemplary embodiment of a CRL is illustrated in FIG. 4. Each entry 401 contains the name of the component 403, the version 405, and the signer 407 whose signature is revoked. The particular CRL used becomes part of the operating system identity using a standard hashing function described further below.

[0104] Alternatively, if the rights manager certificates on the components are short-lived and must be renewed periodically, then a version that is found to be untrustworthy will not have its certificate renewed. This alternate embodiment requires a secure time source to be available on the subscriber computer so the user cannot simply turn back the system clock on the subscriber computer. A monotonic counter in the CPU can serve as this secure time source since it only counts up and cannot be reset "back in time." For example, a monotonic counter that is periodically incremented while the CPU is active, and that cannot be reset, can be used in conjunction with a secure time service, such as a secure Internet time service, to provide a lower bound on the current time in a trusted manner. Such exemplary use of a monotonic counter is described in detail below as part of the functions of the DRMOS.

[0105] Once all components are loaded, the operating system assumes its identity (block 315). In one embodiment, a one-way hashing function provided by the CPU is used to create a cryptographic "digest" of all the loaded components. The digest becomes the identity for the operating system and is recorded in an internal register in the CPU. Alternate methodologies of assigning an identity to the loaded components are equally applicable as long as a non-trusted configuration cannot have the same identity as a DRMOS. Signing the operating system identity with a private key particular to the type of CPU serves to identify both the operating system and the processor on which it is executing.

[0106] If all computers were identically configured, a single, signed operating system identity would suffice to authenticate a particular operating system executing on a particular type of CPU. However, computers contain a myriad different hardware components, and the corresponding supporting software components are frequently updated to add enhancements and fix problems, resulting in a virtually unlimited number of operating system identities. Therefore, the content provider would have to maintain a registry of each subscriber's DRMOS identity or delegate that function to a trusted third party.

[0107] The problems attendant on having a vast number of DRMOS identities can be alleviated in at least three ways. First, an identity is generated or assigned for the basic configuration of each operating system. Such a basic configuration includes only components supplied by the operating system vendor. The identity is generated (or assigned) and stored when the basic components have been loaded. Different versions of the basic operating system will generate (or be assigned) different identities.

[0108] Once the basic configuration of a DRMOS is loaded and its trusted identity is stored, subsequent components required to support the particular hardware configuration must be verified and loaded as explained in conjunction with FIG. 3. Such additional software components can also include updates to the basic components provided by vendors other than the operating system developer. Each additional loaded component has an individual identity (such as a cryptographic digest) generated/assigned and stored. All the identities are uploaded to the content provider when the DRMOS identity is requested. Because the basic DRMOS and additional components always have the same identities when executing on a specific type of processor, the content provider has only to maintain a list of the identities for the combinations of the basic DRMOS and additional components that the provider trusts. Each identity uploaded is then checked against the list.

[0109] In a second alternate embodiment, the operating system maintains a "boot log," containing the identity of the basic DRMOS and the identities of the additional OS components that have been loaded. The identity is a cryptographic digest of the code for the component, or a well-known name, or any other string that is uniquely associated with the component. The CPU also maintains a composite identity register that holds a one-way cryptographic function of the boot log. Whenever a component is loaded, its identity is appended to the boot log and folded into the composite identity register, such as by setting this register to a secure hash of its old value appended with the new component's identity. Whenever the CPU certifies the current value of its composite identity register, it also verifies that the operating system's boot log has not been tampered with. Because the log is indelible, the loaded component cannot erase the record that shows it was loaded.

[0110] An alternate exemplary embodiment of the boot log holds the entire boot log in the CPU. The DRMOS uses an instruction provided by the CPU that appends the identity of each loaded component to the log. The CPU then signs the boot log to attest to its validity and delivers the signed boot log to the content provider as the identity for the DRMOS.

[0111] In another alternate embodiment, DRMOS uses a chain of public and private key pairs newly generated by the CPU to create an indelible boot log. The method is shown in FIG. 5 and an exemplary embodiment of the completed boot log is illustrated in FIG. 6. The boot loader generates or obtains a first key pair ( $K_0, K_0^{-1}$ ) and records the first key pair in memory (block 501). The first public key is also saved to secure storage in the CPU. The boot loader loads the boot block into memory and records the identity of the boot block in the boot log (block 503). Before turning control over to the boot block code, the boot loader obtains a second key pair ( $K_1, K_1^{-1}$ ) (block 505), writes the second public key ( $K_1$ ) to the boot log (block 507), and then signs the boot log with the first private key ( $K_0^{-1}$ ) (block 509). The boot loader deletes the first private key ( $K_0^{-1}$ ) from its memory (block 511) and relinquishes control to the boot block.

[0112] The boot block code loads additional components into memory, records the identities of those components to the boot log (block 515), obtains a third key pair ( $K_2, K_2^{-1}$ ) (block 505), appends the boot log with the third public key ( $K_2$ ) (block 507), and signs its portion of the boot log with

the second private key  $K_1^{-1}$  (block 509). The boot block erases the second private key ( $K_1^{-1}$ ) (block 511) from memory and turns control of the boot process over to the first loaded component. Each loaded component that will load additional components obtains a new key pair ( $K_n, K_n^{-1}$ ) and uses the private key of the previous key pair ( $K_{n-1}^{-1}$ ) to sign its portion of the boot log. The boot process continues in this iterative manner through until all components are loaded or, in the case of a plug-and-and play DRMOS, a content provider challenge is received (block 513).

[0113] When a non-plug-and-play DRMOS resumes control after the final component is loaded, it places a "sentinel" on the boot log (block 519) to indicate that the log is complete and to prevent a loaded component from deleting entire lines of the log. The characteristics of the sentinel are that is a known, unique value that is signed using the last private key ( $K_n^{-1}$ ). In the present embodiment, the sentinel is a signed zero entry. The DRMOS deletes the last private key and all public keys from memory after creating the sentinel.

[0114] Because a plug-and-play DRMOS cannot arbitrarily declare that all components are loaded at the end of the boot process, the DRMOS cannot add a sentinel to the end of the boot log at that time. Instead, the DRMOS attests to its most recent public key  $K_n$  as well as its first public key  $K_0$  to certify the contents of the boot log when challenged.

[0115] Using a chain of key pairs 606, 607, as shown in FIG. 6, guarantees the boot log reflects the loaded components. Each public key in a log section is used to authenticate the signature on the next section. The first public key remains in memory to authenticate the signature on the boot block section of the log. While each set of components is free to load more components, a component cannot change the recording of its identity in a previous portion of the log because doing so would cause the validity check on the corresponding signature to fail. Similarly, a section in the middle of the log cannot be deleted because that would break the chain of keys. Deleting multiple sections of the log through to the end also breaks the chain. In this case, attempting to insert a new sentinel in an effort to make the log appear unaltered will fail because the private key necessary to add the sentinel is not longer available. Finally, the entire boot log cannot be replaced since the signature on the boot block section of the log would not be validated by the first public key.

[0116] Turning now to the boot block, one exemplary embodiment suitable for use with a digital rights management operating system is shown in FIG. 7A. The boot code 701 is signed (signature 703) by the developer of the DRMOS using its private key. The corresponding public key 705 of the developer is attached to the boot block 700. In an alternate embodiment, the public key 705 is not attached to the boot block 700, but instead is persistently stored in an internal register in the CPU. The public key 705 is used to validate the signature 703.

[0117] If the DRMOS developer's private key used to sign the boot block is compromised, the key pair must be changed and thus all boot blocks must be reissued to subscriber computers. FIG. 7B illustrates an alternate embodiment of a boot block that ameliorates this problem. Boot block 710 comprises a basic boot section 711 and an intermediate boot section 713. The basic boot section 711

contains boot code 715 that validates and loads the intermediate boot section 713 and components not provided by the DRMOS developer. The intermediate boot section 713 contains boot code 717 that validates and loads components from the DRMOS developer. The intermediate boot section 713 is signed with a special boot block private key. The basic boot code 715 uses a corresponding boot block public key 719 stored in the basic boot section 711 to validate the intermediate boot section 713. Components 727 from the DRMOS developer are signed 729 with the developer's standard private key and the intermediate boot section 713 uses the DRMOS developer's standard public key 721 to validate those components.

[0118] If the standard private key used to sign components is compromised, the developer creates a new standard key pair and provides a replacement intermediate boot block 713 containing the new standard public key. Replacement components signed with the new standard private key are also issued. Because the special boot block private key is used for few, if any, other purposes than signing boot blocks, it is less likely to be compromised and replacement of the basic boot section 711 will rarely be necessary.

[0119] In FIG. 7C, an alternate embodiment of the single section boot block 730 also uses a special boot block key pair. The boot block 730 contains the special boot block, or master, public key 733. The master private key is used to certify ephemeral keys that are valid for a short period of time. Certificates signed 737 by the master private key attest to the validity of the ephemeral keys. A component is signed with one of the ephemeral private keys and the corresponding certificate 739 is attached. The boot block determines that the certificate on the component is valid using the master public key. When the ephemeral key expires, the DRMOS developer issues replacement components. As with the two-section boot block shown in FIG. 7B, the master private key is only used to sign the certificates for the ephemeral keys so it is less likely to be compromised. Because the ephemeral keys are valid for only a short duration, public release of a private ephemeral key has limited impact.

[0120] Functions of a DRMOS

[0121] As described above, components may be valid only until a specified date and time, and content may also be licensed only until a certain date and time. The monotonic counter described earlier can also be used to ensure that the computer's clock cannot be set backwards to allow the replacement of a trusted component by an earlier, now untrusted version. The DRMOS connects on a regular basis to a trusted time server and presents the value of its monotonic counter, whereupon the trusted time server returns a certificate binding that value to the current time. If the monotonic counter is updated periodically, such as every hour that the DRMOS is running, then the monotonic counter in conjunction with the most recent time certificate can serve as a useful approximation to a trusted clock.

[0122] A DRMOS must also protect the content once it is loaded into the client computer's memory by a trusted application. In particular, the DRMOS must prohibit the use of certain types of programs and refrain from performing certain common operating system procedures when content is in memory.

[0123] An example of one kind of procedure that must be prohibited is loading a kernel debugger because it would

allow the user to make a copy of the content loaded in memory. If the user of the subscriber computer attempts to load a kernel debugger into memory, the DRMOS can either 1) refuse to load the debugger, or 2) renounce its trusted identity and terminate the trusted application that was accessing the content before loading the debugger. In the latter case, the memory must also be purged of the content before the debugger is loaded. The choice of action can be pre-determined or chosen by the user when the user attempts to load the kernel debugger. One of skill in the art will immediately identify other types of programs that will need to be treated in the same manner as a kernel debugger.

[0124] Virtual memory operating systems maintain a page file that holds sections of program code and data that are not currently active. Under normal circumstances, the contents of the page file are accessible by the user of the computer, either by running a privileged program or by booting another operating system that allows inspection of the disk. Therefore, a DRMOS must either protect content stored on the page file or must not page content and similar protected information at all.

[0125] Protecting content on the page file can be accomplished in at least three ways. First, the DRMOS can prohibit all "raw" access to page file device when a trusted application is running. Second, the DRMOS can terminate all trusted applications and erase the page file before allowing such access. Third, the DRMOS can encrypt the content and similar protected information before writing it to the page file.

[0126] Often, a DRMOS must allow the user to perform certain standard functions but prohibit other, related functions. The DRMOS can assign the user permissions based on the granularity of the normally permitted function. For example, the DRMOS can allow the user to delete an entire content file but not a portion of it. Another example is that the DRMOS can allow the user to terminate all the threads of execution for a trusted application but not just a single thread.

[0127] Finally, a DRMOS must protect the trusted application itself from tampering. The DRMOS must not allow other processes to attach to the process executing the trusted application. When the trusted application is loaded into memory, the DRMOS must prevent any other process from reading from, or writing to, the sections of memory allocated to the trusted application without the explicit permission or cooperation of the trusted application.

#### [0128] Key-based Secure Storage

[0129] In order to protect content permanently stored on the subscriber computer, the DRMOS must provide a secure storage space. In essence, the DRMOS must securely store private keys or session keys for use with encrypted content, or provide some other mechanism for keeping these keys secret from other OSs or system level software. These keys can be used for the secure storage and retrieval of protected information. In the exemplary embodiments described in this section, the information to be stored in a protected format is encrypted using one of a set of keys that may be generated by a function 800 provided by the CPU. The storage key generation process is tightly coupled to the DRMOS so that the same key cannot be generated by the CPU for an unrelated operating system, or by any software

on another computer. Three types of storage keys are envisioned as illustrated in FIG. 8: an OS storage key 801, an application storage key 811, and a user storage key 821. Each key is specific to the entity that requests it.

[0130] Beginning with the OS storage key 801, the DRMOS passes a "seed" 803 as an operand of a key-generation instruction ("GenerateKey") 805 to the CPU and receives an OS storage key based on the seed 803 and the identity of the DRMOS. The CPU will always return the same OS storage key 801 when the same seed 803 is provided by the same DRMOS but will return a different OS storage key if the same seed 803 is provided by an unrelated operating system. Because an unrelated operating system cannot get the same key 801, it cannot read any data encrypted by the DRMOS.

[0131] In an alternate embodiment, only a single operating system storage key is used by the DRMOS as described below. Therefore, in this embodiment only the identity of the DRMOS is factored into the key generation function 800 and the seed 803 is not necessary.

[0132] An application storage key 811 is generated when an application calls an operating system instruction ("GenerateAppKey") 815 using a seed 813. The DRMOS passes the seed 813 through an application-specific one-way hash function 817 to produce a hashed seed 819. The hashed seed 819 is then passed to the CPU through the GenerateKey instruction described above. The resulting application storage key 811 is returned to the application for use. Because the GenerateKey function uses the operating system's identity, the same application executing under an unrelated operating system cannot get the same key, and therefore cannot access the encrypted data, even if it requests the key using the same seed 813. Similarly, an unrelated application using the same seed 813 gets a different key because the DRMOS passes the seed 813 through a different hash algorithm for each application.

[0133] In an alternate embodiment, the operating system stores decryption keys for applications using its own identity; the applications call the operating system to retrieve application keys. This also provides a way for an application to allow other applications access to its key and therefore to the content encrypted by the key. Instead of creating a secret using a seed 813, the application passes in the access predicate for the content. The access predicate designates values that must be present in the rights manager certificate for an application wishing access to the content. An exemplary embodiment for an access predicate is shown in FIG. 9 and described in detail in the following section. The DRMOS supplies the seed 813 that is required to generate the application specific key and passes the seed 813 through a generic one-way hash. The DRMOS encrypts the seed 813 and the access predicate using an OS storage key and associates the encrypted access predicate with the encrypted seed. When any application requests access to a key protected by an access predicate, the DRMOS compares the criteria in the access predicate against the rights manager certificate of the requesting application. An application that meets the criteria is given access to the seed 813 and therefore to the application storage key. Because the seed 813 is encrypted using an OS storage key, an application that is running under an unrelated operating system will be unable to gain access to the encrypted data because the unrelated operating system cannot decrypt the seed 813.



[0134] Finally, a particular user can request a key that is based on a user identity assigned by the DRMOS or another facility that guarantees a unique identity for each user. The user supplies a seed **823** in a "GenerateUserKey" call **825**. The operating system passes the seed **823** through a one-way hash **828**, and then passes the resulting first hashed seed **827** through a keyed hash routine **829** to generate a second hashed seed **833**. The operating system factors the user identity **831** into the keyed hash routine **829** so that the second hashed seed **833** is unique to the user. The second hashed seed **833** is passed to the CPU, which returns the user storage key **821**. As described above, only the same user will be able to access data encrypted with the storage key **821** when the DRMOS that generated the key is executing. Analogously, the keyed hash routine **829** guarantees that the user storage key will not duplicate either an OS storage key or an application storage key based on the same seed. Such a facility is used when downloaded content can be accessed only by a particular user. Moreover, if downloaded content is to be accessed only by a particular user and by a particular application, the secret to be stored may be divided into parts, with one part protected by an application-specific key and the other part protected by a user-specific key.

[0135] Once the data is encrypted using the storage keys, there must be a way to recover the keys when the DRMOS identity changes (as when the operating system is upgraded to an incompatible version or an unrelated operating system is installed) or the computer hardware fails. In the exemplary embodiments described here, the keys are stored off-site in a "key vault" provided by a trusted third party. In one embodiment, the DRMOS contains the IP addresses of the key vault providers and the user decides which to use. In another embodiment, the content provider designates a specific key vault and the DRMOS enforces the designation. In either embodiment, when the user requests the restoration of the storage keys, the key vault provider must perform a certain amount of validation before performing the download. The validation process can include such actions as recording the identity of the original operating system (or computer) in a revocation list, checking the frequency of the requests, and requiring a credit card number before downloading the storage keys.

#### [0136] Rights Management

[0137] Most operating systems do not directly process media content, such as video or audio. That function is usually available through special application programs. Therefore, a content provider must not only trust the operating system but must also trust the application that will process the content. Content also can be accompanied by a predicate stating which applications are to be trusted to access that content, and this statement can include a list of generic properties that implicitly define a set of applications. Further associating a rights manager certificate with the application provides identification of the application and certification of its properties. This allows the content provider to determine if the application fulfills the requirements of the content provider before downloading content, and also allows the operating system to restrict future access to only the appropriate applications.

[0138] One exemplary embodiment of a right manager certification is shown in FIG. 9. A list of application properties **903** is appended to the digital certificate fields

**1001** standard in some digital certificate format such as X.509. The certificate names the application. Each entry **905** in the list **903** defines a property **906** of the application, along with optional arguments **907**. For example, one property might be that the application cannot be used to copy content. Another example of a property is one that specifies that the application can be used to copy content, but only in analog form at 480P resolution. Yet another example of a property is one that specifies that the application can be used to copy content, but only if explicitly allowed by an accompanying license. Additional examples include the right to store an encrypted copy of the content and to restrict such storage to only certain, acceptable peripheral devices. The property **906** can also be used to specify acceptable helper applications, such as third-party multimedia processing stacks or other libraries, to be used in conjunction with the application named in the certificate. The certificate is signed by an operating system vendor, content provider, or third party, certifying the properties of the application.

[0139] Because the content provider must trust the DRMOS and application to protect the content from misuse once downloaded, the content provider attaches an access predicate to the content. This access predicate can also include a license to the content. The basic functions of both the access predicate and the license, which were described in the system overview, are explained in detail next.

[0140] In one embodiment, the access predicate takes the form of a required properties access control list (ACL) as shown in FIG. 10. The required properties ACL **1000** contains a basic trust level field **1001**, which specifies the minimum rights management functions that must be provided by any application wishing to process the content. These minimum functions can be established by a trade association, such as the MPAA (Motion Picture Association of America), or by the DRMOS vendor. A unique identifier is used to reference a list of the minimum functions. The minimum functions list can include CPU, DRMOS, and application specific requirements.

[0141] The required properties ACL **1000** can also contain one or more extended trust level fields **1003**. The extended trust level fields **1003** contains identifiers that specify additional rights management function that must be provided by the subscriber computer. For example, a required properties ACL can require that only a certain version of a particular application be allowed access to the content. The required properties ACL **1000** is compared against the certificates for the CPU, the DRMOS, and the application starting at the hardware level, i.e., CPU, DRMOS, application name, version, and specific properties for the application. One of skill in the art will readily recognize that the required properties ACL **1000** can require that all properties must be present, or at least one of the properties, or some specified subset.

[0142] The content license (FIG. 11) imposes additional restrictions on what kind of processing can be performed on the content once an application has access to the content. As described briefly above, the license data structure **1100** can limit the number of times the content can be accessed (usage counter **1101**), determine what use can be made of the content (derivation rights **1103**), such as extracting still shots from a video, or building an endless loop recording from an audio file, or an time-based expiration counter **1105**.

[0143] The license can also specify whether or not a trusted application is permitted to validate other client

computers and share the content with them (sublicense rights **1107**), in effect having the subscriber computer act as a secondary content provider. The sublicense rights **1107** can impose more restrictive rights on re-distributed content than those specified in a license for content downloaded directly from the original content provider. For example, the license **1100** on a song purchased directly from the music publisher can permit a song to be freely re-played while the sublicense rights **1107** require a limit on the number of times the same song can be re-played when re-distributed. To enforce the sublicense rights **1107**, in one embodiment, the trusted application modifies the original license **1100** to specify the additional restrictions and downloads the modified license with the re-distributed content. In an alternate embodiment, the original content provider downloads a sublicense along with the content and that sublicense is re-distributed by the trusted application when it re-distributes the content. The sublicense is structurally identical to the license data structure **1100** although the content of the fields differs.

[0144] Additional licensing restrictions will be readily apparent to one skilled in the art and are contemplated as within the scope of the invention.

[0145] The license **1100** is stored with the content on secured storage. In one embodiment, the required properties ACL **1000** is also stored with the license **1100** and the content. In an alternate embodiment, the ACL **1000** is secured separately and controls access to the storage key for the content as described above.

[0146] In the embodiments described above, the DRMOS is responsible for checking the required properties ACL and for enforcing the licensing restrictions. By providing the validation functions in the DRMOS, the functions are centralized and can be utilized by any process. In an alternate embodiment, the validation functions concerning the application are coded directly into the trusted applications programs. A similar effect is achieved in yet another alternate embodiment that places the application validation functions in a library that is incorporated into the trusted applications.

[0147] One of skill in the art will immediately perceive that certain rights are more easily enforced at the DRMOS level, such as the right for a certain application to access a key or other content, or the ability to open a file a limited number of times, while other types of rights are best enforced by the application itself. Since the DRMOS enforces the restriction that only explicitly stated applications can access restricted content, the application can be trusted to enforce the additional restrictions. Alternate embodiments in which the enforcement of certain rights is allocated to the DRMOS and the enforcement of others to the application is therefore contemplated as within the scope of the invention.

[0148] As described above in conjunction with FIG. 2, the content provider **220** delivers content to the subscriber computer **200** after trust is established by transmitting the appropriate certificates/identities for the CPU, the DRMOS, and the application to the provider. The content can be explicitly encrypted by the content provider for this combination of CPU, DRMOS, and application, as described above, or, if the content is sent over a secured link (with, for example, Secure Socket Layer services), the content provider can encrypt the content using the session key for the secure link. In the latter embodiment, the DRMOS writes the

encrypted content to permanent storage and uses one of the storage keys generated by the CPU to securely store the session key for later use. Alternately, the content provider can choose not to encrypt the content if it is transmitted to the application in a secure fashion, in which case the application performs the encryption if it stores a persistent copy of the content.

[0149] The particular methods performed by a subscriber computer of an exemplary embodiment of the invention have been described. The methods performed by the subscriber computer have been shown by reference to flowcharts, operational diagrams, and data structures. Methods performed by the content provider have also been described.

#### Secure Storage

[0150] As an alternative to the GenerateKey operation, two new operations referred to as “Seal” and “Unseal” may be introduced, which provide the ability to seal secrets only for subsequent use on the same machine.

[0151] The “Seal” instruction takes as inputs an arbitrary block of data, the current OS identity (the SIR), and a target OS identity (a specified SIR value that must be current at the point of future decryption). The processor encrypts this data structure using a symmetric key,  $K_S$ .

[0152] The data block can now only be decrypted via an “Unseal” operation on the same processor, using the same symmetric key. This symmetric key is only used by the “Seal” and “Unseal” operations, and will only decrypt the secret if the target OS identity is equal to the current value of the SIR. If this check succeeds, the processor decrypts and returns the secret, otherwise it returns an error.

[0153] In this way, a processor can store encrypted information that can be decrypted only by the same processor running a specified operating system.

[0154] As a special case, the operating system can choose to seal information for a different operating system whose identity it knows and trusts. An example of this occurs when the operating system is about to be upgraded and has a signed certificate from the operating system vendor confirming the identity of the new operating system. In this case the operating system will seal its secrets for the new operating system that is about to run.

[0155] Alternatively, another approach is to employ encryption with the processor’s public key and decryption using the “Reveal” operation, as described earlier. Instead of using  $K_S$  for encryption and decryption, the processor’s public key pair is used. This allows the “Seal” operation to be performed in software, even on another processor.

#### Exemplary Chipset Implementation

[0156] The fundamental requirements of atomicity and privileged access to keys for the microcode that implements authenticated boot can be met in a variety of alternative implementations. In one implementation, components in the chipset may examine the bus to infer operation and permit or deny access to keys depending on the code executing. Components on the chipset can also examine the bus for unauthorized agents writing to protected code, or reading unauthorized secrets.

[0157] An agent on the bus can also check for unauthorized interrupts during the execution of the authenticated operations or execution of the boot block.

[0158] Similarly, there is no fundamental requirement for the microcode that implements the authenticated boot operations to be physically resident on the microprocessor chip. It could also be stored in ROM, EPROM, or protected flash memory in a physically separate device on the bus.

#### BIOS Implementation

[0159] The authenticated boot technique can be implemented by existing CPU operating modes using code in the computer's BIOS code. The System Management Mode (SMM), supported by Intel microprocessors, provides for a region of memory that is inaccessible to normal operating system operation, but can provide subroutines that operating systems or applications can use. Such SMM protected memory could be used for the storage of keys and the code that manages those keys.

#### CONCLUSION

[0160] A digital rights management system has been described whereby certain cryptographic secrets are reliably associated with a particular digital rights management operating system or family of operating systems running on a particular general-purpose computer. The operating system uses these secrets to authenticate itself to a third party, to receive encrypted content or information, to store this content or information securely, and to retrieve it later. No unrelated operating system or other software running on the same computer can obtain these secrets and perform these actions, nor can any operating system or other software running on any other computer. By using these cryptographic secrets, the digital rights management operating system can recursively provide derived cryptographic secrets for the same uses by applications running on the same operating system on the same computer.

[0161] Although specific embodiments have been illustrated and described herein, it will be appreciated by those of ordinary skill in the art that any arrangement which is calculated to achieve the same purpose may be substituted for the specific embodiments shown. This application is intended to cover any adaptations or variations of the present invention.

[0162] For example, those of ordinary skill in the art will appreciate that various combinations of the exemplary embodiments are applicable to solve the digital rights management problem depending on the exact computing environment. Furthermore, those of ordinary skill in the art will recognize that the invention can be practiced on a large scale although illustrated herein with only a single subscriber and content provider.

[0163] The terminology used in this application with respect to is meant to include all hardware and software configuration and all networked environments. Therefore, it is manifestly intended that this invention be limited only by the following claims and equivalents thereof.

We claim:

1. A method comprising:

receiving a data structure to be encrypted, wherein the data structure includes content along with a statement of conditions under which the content may be decrypted, wherein the statement of conditions comprises an operating system identity that an operating system executing on the device must have in order for the content to be decrypted; and

encrypting the content using a key.

2. A method as recited in claim 1, wherein the key comprises a symmetric key.

3. A method as recited in claim 1, wherein the key comprises a symmetric key of a processor of the device.

4. A method as recited in claim 1, wherein the operating system identity is maintained in a software identity register (SIR).

5. A method as recited in claim 1, wherein the operating system identity is identified in a signed certificate from an operating system vendor.

6. A method as recited in claim 1, wherein the operating system identity is for an operating system that is different than an operating system executing when the content to be encrypted is received.

7. A method as recited in claim 1, wherein encrypting the content comprises encrypting the data structure.

8. A system comprising:

means for obtaining a block of data to be encrypted, a current operating system identity, and a target operating system identity; and

means for invoking a seal operation to have the block of data encrypted by a processor of the system using a symmetric key of the processor.

9. A system as recited in claim 8, further comprising means for invoking the seal operation so that only an operating system having the target operating system identity can decrypt the encrypted block of data.

10. A system comprising:

means for invoking an unseal operation in order to have a data block decrypted using a key; and

means for receiving, in response to invoking the unseal operation, the decrypted data block only if conditions under which content in the data block can be decrypted are satisfied.

11. A system as recited in claim 10, wherein the key comprises a symmetric key.

12. A system as recited in claim 10, wherein the key comprises a symmetric key of a processor of the device.

13. A system as recited in claim 10, wherein the conditions comprise an operating system identity that an operating system invoking the unseal operation must have in order for the content to be decrypted.

14. A system as recited in claim 13, wherein the operating system identity of the operating system invoking the unseal operation is different than an operating system identity of an operating system that previously had the data block encrypted.

15. A system as recited in claim 10, further comprising means for receiving, in response to invoking the unseal operation, an error indication if the conditions are not satisfied.