

(19) World Intellectual Property Organization
International Bureau



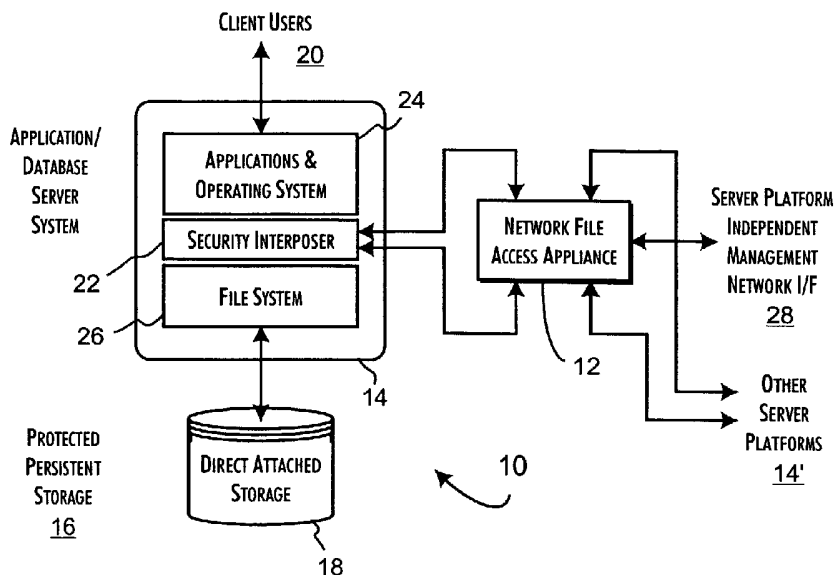
(43) International Publication Date
29 April 2004 (29.04.2004)

PCT

(10) International Publication Number
WO 2004/036350 A2

- (51) International Patent Classification⁷: **G06F**
- (21) International Application Number: PCT/US2003/031261
- (22) International Filing Date: 1 October 2003 (01.10.2003)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
10/271,050 16 October 2002 (16.10.2002) US
- (71) Applicant: **VORMETRIC, INC.** [US/US]; 3131 Jay Street, Santa Clara, CA 95054-3308 (US).
- (72) Inventors: **LO, Mingchen**; 275 Ondina Drive, Fremont, CA 94539 (US). **NGUYEN, Tien**; 10105 Stern Ave, Cupertino, CA 95014 (US). **PHAM, Duc**; 10412 Menhart Lane, Cupertino, CA 95014 (US). **ZHANG, Pu**; 6404 Mojave Drive, San Jose, CA 95120 (US).
- (74) Agent: **ROSENBERG, Gerald**; NewTechLaw, 285 Hamilton Avenue, Suite 520, Palo Alto, CA 94301 (US).
- (81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VC, VN, YU, ZA, ZM, ZW.
- (84) Designated States (*regional*): European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PT, RO, SE, SI, SK, TR).
- Published:**
— without international search report and to be republished upon receipt of that report
- For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*

(54) Title: SECURE FILE SYSTEM SERVER ARCHITECTURE AND METHODS



(57) Abstract: A data server platform (14) includes a security file system layer (22) interposed between the platform operating system kernel (24) and file system (26). The secure file system layer is structured to implement a file access control function that selectively constrains data transfer operations initiated through the operating system kernel by an application program to transfer file data through the file system with respect to a persistent data store (18). A file access controller (12), implemented independent of the operating system kernel, is coupled to the security file system layer and supports the file access control function by defining permitted file data transfers through the file system. Management of the file access controller separate from the data server platform ensures that any security breach of the platform operating system kernel cannot compromise the function of the security file system layer.

WO 2004/036350 A2

- 1 -

[0001] SECURE FILE SYSTEM SERVER
 ARCHITECTURE AND METHODS

[0002] Inventors:
 Duc Pham
 Tien Le Nguyen
 Pu Paul Zhang
 Mingchen Lo

[0003] Background of the Invention

[0004] Field of the Invention:

[0005] The present invention is generally related to systems and methods of protecting persistently stored data from unauthorized access and modification and, in particular, to a system and method of reliably securing persistent data managed through a file system server operating in the role of an application or database service provider platform.

[0006] Description of the Related Art:

[0007] Maintaining robust and verifiable security over persistently stored data has and continues to be a primary requirement in the operation of commercial, governmental, and essentially all other computing environments. Achieving a desired high level of security, however, is not commonly realized for a variety of reasons. With the fast-paced development particularly of the Internet infrastructure, the diversity of network architectures, infrastructure devices, and computer protocols has rapidly increased. Thus, the number, variety, and complexity of computer systems and network

- 2 -

components that must be cooperatively managed to establish at least a minimum uniform level of security within some defined domain has correspondingly increased. The practical difficulties of coordinated management of the different systems and components, as well as systemic failures to protect against both known and previously unrecognized security attack approaches, also contribute to the vulnerabilities of systems and networks to security breaches.

[0008] Conventionally, security protections for a given domain are layered and specialized depending on the operational nature of corresponding individual computer systems and network components. Typically, these security protections are implemented variously as password challenges and data and connection filters layered over the core functionality of a computer system or network component. The conventional implementation of security functions in software in computer systems and network components implicitly recognizes the inherent complexity of establishing robust security mechanisms and the very practical need to frequently apply patches to close both previously unrecognized and newly emergent vulnerabilities.

[0009] Different architectural approaches have been explored to minimize the vulnerabilities of different security mechanisms to attack and, ultimately, loss of security over persistently stored data. US Patent No. 5,007,082, issued on April 9, 1991 to Cummins, describes an early data security system applicable to file data transfers. Balancing the need for security, transparency of use and compatibility, Cummins describes a hardware specific, software-based encryption control system that interoperates with the platform operating system at the basic I/O system (BIOS) level. File data transfer operations provided by the BIOS are selectively re-vectorred to compatible routines implementing encryption and decryption functions against file level data. The described functions perform file-level encryption and

- 3 -

decryption. Consequently, full file reads and writes are required to support application program read/write support. Compatibility further requires that the relevant file data memory buffers must be maintained in an unencrypted state to support dynamic read/write operations directed by executing applications.

[0010] US Patent No. 5,584,023, issued December 10, 1996 to Hsu describes a similar, but more advanced software-based file-data encryption system. An operating system kernel mode driver is used to re-vector selected file data related operating system calls at the system call interface level. The underlying operating system provided file-oriented system calls are wrapped to support block-level encryption and decryption services, where the block size is determined by the nominal operation of the operating system. File blocks retrieved from encrypted files are generally maintained in an encrypted state while managed by the operating system within the kernel space buffer cache. The encryption and decryption algorithms utilize password key to select encryption code tables defined against user related processes and initialized with the invocation of an initial user process. Separate file attribute tables, populated from data appended to the individual disk files, define the encryption attributes of individual actively accessed files. The various code and attribute tables are dynamically allocated and limit pointer references to increase the difficulty in tracing the data structures and the corresponding operation of the kernel mode driver.

[0011] A more involved, but similar encryption system is described in US Patent No. 6,249,866, issued June 19, 2001 to Brundrett et al. An extensive modification of the operating system is described to integrate both file and file system encryption functions into a logically unified view of the underlying file system space. Operating system calls to read and write data are evaluated on interception to determine if the calls are directed to an encrypted file system

- 4 -

directory or file. System call file data transferred relative to an encryption flagged directory or file is encrypted or decrypted utilizing a key associated with the target directory or file utilizing a kernel mode driver layered above a conventional file system, such as the NT file system (NTFS). The encryption keys and encryption driver management systems are implemented as a combination of operating system kernel functions and user-mode key management applications.

[0012] While these conventional security systems provide a significant degree of security over persistent stored data, each fails to establish a comprehensive security system. Each of these systems remains particularly vulnerable to basic Trojan attacks for obtaining passwords and encryption keys, thereby permitting complete conversion of the security systems to support inappropriate access to and modification of the persistent stored data. Furthermore, these systems provide no protection against the execution of user-mode programs that may exploit vulnerabilities in the operating system to gain unlimited root or administrator control over the operating system. An intruder can then either directly circumvent the kernel password and encryption mechanisms or breach the security of the password and encryption key management systems to obtain the passwords and keys. In either case, the intruder again obtains unencumbered access to the ostensibly secured persistently stored data over the heightened encryption-based security capabilities with little greater difficulty than exploiting the typically limited security protections afforded by the operating system itself.

[0013] Ultimately, the security systems described by Cummins, Hsu, and Brundrett et al. rely on the basic security subsystems of the local operating system to prevent attack on the underlying encryption mechanisms. Where utilized on systems that are part of an extended security domain, these and similar systems also inherently rely on whatever cooperative management

- 5 -

policies are enforced for the coordinated configuration of the required password and encryption key management systems. Unintended errors and perhaps more typically lapses in the consistent and comprehensive management of the security mechanisms protecting the security domain only increase the availability of operating system vulnerabilities that may be exploited to penetrate the security domain and inappropriately permit access to persistently stored data.

[0014] Consequently, there is a clear need for mechanisms to secure persistent data that are ultimately reliable and cooperatively manageable.

[0015] Summary of the Invention

[0016] Thus, a general purpose of the present invention is to provide an efficient and effective mechanism for reliably securing persistent data in a manner eminently subject to cooperative management and control within a security domain.

[0017] This is achieved in the present invention by providing, on a data server platform, a security file system layer interposed between the platform operating system kernel and file system. The secure file system layer is structured to implement a file access control function that selectively constrains data transfer operations initiated through the operating system kernel by an application program to transfer file data through the file system with respect to a persistent data store. A file access controller, implemented independent of the operating system kernel, is coupled to the security file system layer and supports the file access control function by defining permitted file data transfers through the file system.

[0018] An advantage of the present invention is that the management of the file access controller is separate from the data server platform. This

- 6 -

ensures that any security breach of the platform, including specifically the operating system kernel, does not and cannot compromise the essential function of the security file system layer as supported by the independently managed file access controller.

[0019] Another advantage of the present invention is that the security file system layer permits a broad range of security information to be collected and processed through the file access controller as a basis for determining permissions for file accesses. Collection of authentication, authorization and authenticity information can be efficiently performed by or through the security file system layer. The independently secure evaluation of the collected information ensures the secure function of the security file system layer in selectively enabling file data access and modification.

[0020] A further advantage of the present invention is that minimal platform modifications are required to ensure the security of persistent platform accessible data. Implementation of the security file system layer is consistent with the architectural model of conventional operating systems supporting virtual file system support.

[0021] Still another advantage of the present invention is that the secure file system layer, in combination with the file access controller, provides a comprehensive security envelope relative to persistent stored data accessible through a protected computer platform. The security envelope includes identity authentication, access authorization and application authenticity to ensure that only well-identified users, subject to defined policy permissions can access persistent stored data through certified applications. The integrity of the persistent stored data is further preserved through data encryption subject to dynamic tamper detection.

- 7 -

[0022] Brief Description of the Drawings

[0023] These and other advantages and features of the present invention will become better understood upon consideration of the following detailed description of the invention when considered in connection with the accompanying drawings, in which like reference numerals designate like parts throughout the figures thereof, and wherein:

[0024] Figure 1 is a top-level block diagram illustrating the intended operating environment of a preferred embodiment of the present invention;

[0025] Figure 2 is a detailed architectural block diagram of a preferred implementation of the present invention;

[0026] Figure 3 is an architectural block diagram of a preferred, fixed scale appliance embodiment of the present invention;

[0027] Figure 4 is a process flow diagram illustrating the deep packet analysis processing provided in accordance with the present invention to support authentication and access qualification of client file oriented network requests directed to network storage resources;

[0028] Figure 5A provides a process interaction diagram showing the interoperation of platform processes with a secure file system layer executed by a platform computer system;

[0029] Figure 5B illustrates the process flow integration of a preferred embodiment of the present invention into the login and execution process operations in accordance with a preferred embodiment of the present invention.

[0030] Figure 6 provides a process interaction diagram illustrating the preferred exposure of network storage resources provided in a preferred embodiment of the present invention to provide multiple qualified views of the underlying file data;

- 8 -

[0031] Figure 7 is a software block diagram illustrating the preferred components implementing network packet protocol processing in accordance with a preferred embodiment of the present invention;

[0032] Figures 8A-D illustrates the preferred decomposition of file data through the network packet protocol processing implemented in accordance with a preferred embodiment of the present invention;

[0033] Figure 9 is a software block diagram illustrating an extended network packet protocol processing including firewall processing in accordance with a preferred embodiment of the present invention;

[0034] Figures 10A-B illustrate the process flow of a file system read request and response performed in accordance with a preferred embodiment of the present invention;

[0035] Figures 11A-B illustrate the process flow of a file system file create request performed in accordance with a preferred embodiment of the present invention; and

[0036] Figures 12A-B illustrate the process flow of a file system write request and response performed in accordance with a preferred embodiment of the present invention.

[0037] Detailed Description of the Invention

[0038] The present invention utilizes and extends the operation of a secure network file system appliance to establish a security envelope protecting persistent stored data accessible through various computer system platforms. An exemplary server platform protected environment 10 is shown in Figure 1. A secure network file access appliance 12 is preferably implemented in the environment 10 to support the secure operation of one or more computer server platforms 14, 14' relative to protected persistent storage resources 16,

- 9 -

such as direct attached storage 18. For purposes of the preferred embodiments of the present invention, the computer system platforms 14, 14' are database and application server platforms supporting local and remote client systems and users 20. The secure network file system appliance is integrated with the computer system platform 14 through a security interposer layer 22 established between the operating system kernel 24 and a file system 26 through which data is transferred relative to direct attached storage 18.

[0039] The security interposer layer 22 selectively routes file oriented data transfers between the operating system kernel 24 and file system 26 through the secure network file access appliance 12 to encrypt and decrypt the file data stored to the protected persistent storage resources 16 subject to access policies implemented within the secure network file access appliance 12. In accordance with the present invention, the file data encryption maintains the logical file-oriented structure of the data and is thus transparent to the persistent storage resources 16. Additionally, the secure network file access appliance 12 can implement IP firewall functions, limiting potential attacks on the security of the secure network file system appliance from the computer server platforms 14, 14'.

[0040] Clients and users 20 are unaffected by the security support and encryption functions of the secure network file access appliance 12, yet are secured against unauthorized access of the encrypted content. Actively used file data encryption keys are preferably held and managed within the secure network file access appliance 12 essentially independent of the computer server platform 14. A logically and preferably physically separate private network interface 28 is supported by the secure network file access appliance 12 to control effectively single-point secure management over the secure network file access appliance 12 by a centralized security management system. Thus, the encryption keys utilized by the secure network file access appliance

- 10 -

12 are not accessible in usable form as a consequence of a security breach of the computer system platform 14, direct attached storage 18, or client users 20.

[0041] Preferably, the secure network file access appliance 12 processes file data read and write requests in aggregate at wire-speed and with minimal latency in qualifying the access privileges of each read, write, and related file access request, to selectively encrypt and decrypt file data transferred, and further selectively compress and decompress the transferred file data. The round-trip encryption of file data ensures that all protected data stored by persistent storage resources 16, including the direct attached storage 18, are secure both as transferred and while statically stored. Round-trip compression can substantially reduce the needed file data transfer bandwidth, particularly where the transfers are for repeated mass archival backups.

[0042] The preferred structure 40 of a computer system platform 14 is shown in Figure 2. The platform 14 conventionally includes an operating system kernel 42 supporting execution of applications, such as a database management system (DBMS) 44 and other server applications 46, in a user mode execution space. The operating system kernel 42 also preferably supports execution of an authentication agent program 48 substantially, if not completely, within a kernel mode execution space. A virtual file system switch (VFS) 50 provides a conventional interface to any number of different conventional file systems (xFS) 52 as necessary to access conventional direct attached storage 18 and, for example, storage devices 54 accessible through a storage area network 56.

[0043] In a first preferred embodiment of the present invention, the security interposer layer 22 is implemented utilizing a secure network file system (SNFS) 58 and a conventional network file system (NFS) client file system daemon (NFSd) 62. The secure network file system 58 is preferably

- 11 -

based on a conventional network file system (NFS) implementation used to route network file transfer requests and data through the secure network file access appliance 12. The secure network file system 58 includes modifications to enable collection of user, process and session information through an interface 60 to the operating system kernel 42, in regard to specific network file transfer requests, and to provide this information to the secure network file access appliance 12 as a basis for determining whether to permit the corresponding network file transfer to proceed.

[0044] Network file data transfers permitted by the secure network file access appliance 12, including any associated transparently encrypted and compressed data, are processed through the NFS daemon 62 and an appropriate file system 52 as necessary to transfer data relative to the persistent storage resources 16. Alternately, the secure network file access appliance 12 may determine to route the network file request and data through a network infrastructure, as generally represented by router 64, to conventional network attached storage 66 operating as a network file system client. Utilization of the conventional NFS protocols between the secure network file access appliance 12 and at least the NFS client daemon 62 enables substantial flexibility in using direct attached storage 18, storage area networks 56, and network attached storage 66.

[0045] For a second preferred embodiment of the present invention, the necessary complexity of maintaining NFS protocol compliance and the overhead of executing NFS client daemons 62 may be avoided by implementing an RPC-based communications protocol between the security interposer layer 22 and secure network file access appliance 12. While using RPC-based messages functionally similar to the NFS protocol, the requirements for formal NFS compatibility are obviated and both server and client communications functions can be merged into a single, concisely defined

- 12 -

secure virtual file system layer (SVFS). In this embodiment, the security interposer layer 22 again implements the interface 60 to the operating system kernel 42 and supports any combination of an overlay interface to a conventional file system 52, a captive file system, generally equivalent to the file system 52 though implemented as an internal component of the security interposer layer 22, and a direct interface to the direct attached storage 18 in support of a file system externally implemented as part of the secure network file access appliance 12.

[0046] Implementation of comprehensive access policy controls in the secure network file access appliance 12, essentially independent though additive to those of the computer system platform 14, enables centralized secure file data access management. The access permissions and other controls implemented by multiple computer system platforms 14, 14' and potentially other network components implementing the persistent storage resources 16 are difficult to globally maintain through additions and reconfigurations of the various storage devices 18, 54, 64. The access policy controls provided by the secure network file access appliance 12 are significantly more comprehensive, flexible, and administratively uniform than conventional access permissions supported by computer system platforms 14, 14'.

[0047] Authentication controls are also supported by the secure network file access appliance 12 as a complement to the access policy controls. For the preferred embodiments of the present invention, the security interposer layer 22 interoperates with the authentication agent program 48 installed and executed on the computer system platforms 14, 14' to enable user and client authentication, including authentication over user sessions and processes. Specifically, the security interposer layer 22 supports a modified file system interface, compatible with the virtual file system switch 50 as implemented by

- 13 -

the conventional operating system kernel 24, to provide selective authentication processing of file system requests directed to the protected persistent storage resources 16. For the preferred embodiments of the present invention, the file system switch interface security interposer layer 22 is mounted through the file system switch 50 against the directory nodes representing protected persistent storage resources 16. Authentication logic provided in the agent program 48, preferably executing in kernel space, is called through the operating system kernel 42 in response to file system operations directed against the security interposer layer 22. Through the operating system kernel 24, the agent program 48 has access to user, client process, application, and session information. Where attended user authentication is required, the agent program 48 preferably interoperates through the operating system kernel 42 to assert an authentication dialog for a user 20. User responsive information can then be authenticated using standard authentication controls, such as LDAP and other network available authentication servers (not shown). Alternately, or in combination, the user authentication response information can be transmitted to the secure network file access appliance 12 for security qualification.

[0048] Authentication of user applications 44, 46 executed within the application execution space supported by the operating system kernel 42 is performed autonomously through the agent program 48. Preferably in response to a first file system operation by a user 20 application, as received by the security interposer layer 22, or on notice from the operating system 42 of the invocation of the user application, the agent program 48 generates a secure hash identification of the loaded binary image of the user application. This hash identifier and the application file attributes are then transmitted to the secure network file access appliance 12 for verification. An authentication response is returned to the agent program providing verification status. A

- 14 -

verification failure or other exception indicated by the secure network file access appliance 12 preferably results in a disallowance of the requested file system operation.

[0049] Unattended execution of applications by the computer system platform 14, such as on booting of the platform 14, can also be supported through the application authentication mechanism. Preferably, an application launcher utility is scripted to execute on boot. Through application authentication of the utility, the absence of attended user authentication derived information is not treated as an exception by the secure network file access appliance 12. The application launcher utility is then enabled to launch a designated application.

[0050] The state of user and application authentication, in combination with user session and associated process identifiers, is preferably maintained by the agent program 48 or within the security interposer layer 22. In the preferred embodiments of the present invention, this authentication information and the digital signature of the agent program 48 are combined and sent encrypted to the secure network file access appliance 12 with each file system request received by to the interposer layer 22. Where an NFS communications protocol is utilized between the security interposer layer 22 and secure network file access appliance 12, the NFS packets are modified to include the user and agent authentication information with the file system requests. In the preferred embodiment, an NFS packet header field is extended, preferably by redefinition of an existing field, to store and transfer the user and agent authentication information. Additionally, periodic or heartbeat status remote procedure call (RPC) packets are sent by the agent program 48 to the secure network file access appliance 12 reflecting the current state of the user and agent authentication information. Computer system platform 14 changes relevant to authentication, including specifically

- 15 -

terminations of processes and user sessions, are thereby rapidly noticed to the secure network file access appliance 12.

[0051] A preferred, fixed scale, hardware platform 70 for the present invention is shown in Figure 3. The platform 70 is preferably implemented on a motherboard supporting the Intel® E7500 chipset 72, dual 2.2 GHz Intel® Xeon™ processors 74 (Intel Corporation, Santa Clara, California; www.intel.com), and a 1-Gbyte 200-MHz Double Data Rate (DDR) main memory array 76. The chipset 72 supports six PCI-X buses 78, individually capable of over 8-Gbps throughput and an aggregate throughput of at least 24-Gbps. A basic configuration of two 1-Gbps network interface controllers 80, supporting ingress and egress network connections, and one 10/100 Mbps network interface controller 80, supporting the management network connection 28, are connected to the PCI-X bus 78. A base configuration of three HiFn™ 7851 security processors 82 (HiFn, Inc., Los Gatos, California; www.hifn.com) provides hardware accelerated encryption and compression support for the generic data processing and control function of the processors 74. The security processors support symmetric programmable length block encryption algorithms, including 3-DES, at throughputs in excess of 400-Mbps per chip and programmable length block compression algorithms, including LZS, at throughputs in excess of 80MBps.

[0052] Other peripherals 84, including a BIOS program and boot hard disk drive, are supported through the chipset 72 to enable basic operation of the hardware platform 70. Preferably, the hardware platform 70 boots and runs a Linux™ based operating system, based on a commercial distribution of Red Hat™ Linux (Red Hat, Inc., Raleigh, North Carolina; www.redhat.com). The software-based authentication and access functions of the secure network file access appliance 12 preferably load and execute in the Linux kernel space.

- 16 -

Administrative and support utilities are preferably implemented as user-mode applications and daemons.

[0053] The logical control and protocol processing functions implemented in the control programs executed on a hardware platform 70 for a preferred embodiment of the present invention are shown in Figure 4. Inbound file requests are received as network data packets containing the various network file system messages implemented by a network distributed file system, such as the network file system (NFS), common internet file system (CIFS), or whatever secure virtual file system (SVFS) protocol is implemented by the security interposer layer 22. These network data packets are processed to expose the control information 114 contained in the protocol layers of each received data packet and the packet payload data 116 for examination and processing.

[0054] Additionally, application and status information is gathered by an agent monitoring process 118 listening on a dedicated network port connected to the computer system platforms 14, 14'. Client status information, obtained from heartbeat network packets, is relayed to an authentication and access control process 120. Continuity of a client heartbeat is used to maintain a client authorization session. User authentication session information, minimally reflecting that a user authentication sequence mediated by the agent program 48 has completed successfully, can also be provided to the authentication and access control process 120 within the heartbeat data packets. Transmission of user authentication session information at checkpoint intervals serves to protect against conversion of any client process for the execution of unauthorized applications. Where the authentication and access control process 120 operates directly as an authentication server, user and client identifiers and the user password acquired by the agent program 48 are relayed through the agent monitor process 118. Authorization responses are

- 17 -

generated and returned by the authentication and access control process 120 based on the user and client authentication policy information maintained by the authentication and access control process 120.

[0055] Figure 5A shows a configuration 130 of the computer system platform 14 employing the second preferred embodiment of the security interposer layer 22 where the secure virtual file system 132 effectively combines the function of the secure network file system 58 and network file system daemon 62. While the secure virtual file system 132 may also use the NFS protocol for transferring file transfer requests and data with the network file access appliance 12, the secure virtual file system 132 preferably implements a non-standards compliant RPC-based message transfer protocol to obscure the information transferred between the computer system platform 14 and network file access appliance 12. A conventional interface to the virtual file system switch 50 is supported so that the secure virtual file system 132 appears to the switch 50 as an ordinary file system. The secure virtual file system 132 implements the extended operating system kernel 42 interface 60 to support operation of the authentication agent program 48. The secure virtual file system 132 also implements a conventional file system overlay interface 134, permitting functional capture and utilization of conventional file systems 52. Dedicated or proprietary file systems 136 may also be closely coupled to the secure virtual file system 132.

[0056] For both computer system platform 14 configurations 40, 130, user authentication enforcement is enabled by requiring a call to the agent program 48 in connection with the initialization of a new user process 138. As shown in Figure 5B, following from the initialization of the new user process 151, the network file access appliance 12, through the secure virtual file system 132, causes the agent program 48 to present a session login 152 to the user 20 associated with the client process 138. User authentication is

- 18 -

performed directly by a user mode component of the agent program 48 through a conventional authentication service, such as LDAP, against a user login and password. Alternately, user authentication can be direct through a pluggable authentication module generally consistent with DCE/OSF-RFC 86.0 (Unified Login with Pluggable Authentication Modules (PAM); www.opengroup.org/tech/rfc/rfc86.0.html). Upon authentication of the user 20, the agent program 48 initiates an authenticated user session 153 by submitting an authenticated user identification along with the login process identifier (LPID), user identifier (UID), and group identifier (GID) established by and collected from the operating system kernel 42. This session identification information is submitted through 154 the secure virtual file system 132 to the network file access appliance 12. A secure session key is then returned 155 to the agent program 48 for use in connection with subsequent file system requests initiated from the authenticated client process 138.

[0057] An authenticated user process 138 issues file requests to the operating system kernel 42, which are then passed to the secure virtual file system 132. A kernel mode portion of the agent program 48, operating in conjunction with the secure virtual file system 132, determines the source process identifier 142 for each file request 134, as received by the operating system kernel 42, by accessing conventional operating system kernel 42 structures. The authenticated user session information maintained by the agent program 48, located by the determined process identifier, is then provided to the secure virtual file system 132 for inclusion in the file transfer requests processed through the network file access appliance 12.

[0058] Client processes 144 spawned from an authenticated process 138 remain part of the parent authenticated user session. The chain of parent process identifiers is traced 142 by the agent program 48 to associate file requests 146 from child processes 144 with corresponding authenticated user

- 19 -

sessions. Preferably, to support access management at the level of individual processes, both the authenticated user login parent process identifier (LPID) and the current process identifier (PID) are provided to the modified network layer for inclusion in the session and process corresponding file requests forwarded to the secure network file access appliance 12.

[0059] In a preferred embodiment of the present invention, the authenticated user session information, including a session identifier generated by the agent program 48, is encrypted using the session key as obtained through a secure key exchange with the agent monitoring process 118. The resulting extended NFS requests thus securely transports the session control information, including at least a session identifier, user identifier, group identifier, and process identifiers to the secure network file access appliance 12.

[0060] Preferably, the agent program 48 further supports authentication of user applications 44, 46, particularly including login shell applications, as loaded for execution in the authenticated user session processes 138, 144. Digitally signed applications loaded for execution can be verified conventionally by the agent program 48 against digital certificates obtained from a trusted PKI-based, LDAP or other authentication server. Application authentication information, such as the identity of the authentication server and certificate, can be potentially included, by the secure virtual file system 132, with the session information provided with corresponding file requests to support auditing of independently verified applications.

[0061] Autonomous application authentication by the agent program 48 is also supported through the secure network file access appliance 12. On the loading of an application for execution in a process 138, 144, the agent program 48 is called and executes, through the operating system kernel 42, to locate 148 the application binary image and retrieve the application file

- 20 -

attributes, including the application filename, path, permissions, and file size. A secure hash signature is generated for the application binary. In a preferred embodiment of the present invention, a 20-byte hash signature is generated using the SHA-1 algorithm. An application authentication request, containing the hash signature, file attributes and a secure application token, is then passed to the secure network file access appliance 12 in an RPC directed to the agent monitoring process 118. The secure application token preferably includes a public key, of a public/private key pair stored by the secure network file access appliance 12 or trusted third-party authentication server, an application name, and a structure containing a secure hash signature of the application binary image and the application file attributes encrypted with the public key. The token is prior administratively generated through the secure network file access appliance 12 or other trusted application authenticator against an administratively determined authentic application. The tokens for authenticated applications are stored on or otherwise made accessible to the computer system platforms 14, 14'. The application file name located for the loaded binary image is used to further locate a corresponding token by the agent program 48.

[0062] On presentation of an application authentication request, the secure network file access appliance 12 compares the public key provided within the token against known valid public keys prior administratively registered with the secure network file access appliance 12. The decrypted token hash signature and file attributes are verified against the hash signature and file attributes separately provided in the request by the agent program 48 and a return RPC communicates the verification status to the agent program 48. Where the loaded application fails authentication, the corresponding application process 138, 144 can be terminated. Alternately, subsequently

- 21 -

received network file system requests 140, 146 from an unauthorized application can be ignored or refused by the secure virtual file system 132.

[0063] Periodically, the agent program 48 may regenerate and resubmit 156 user session and application authentication requests to the secure network file access appliance 12. Alternately, or in addition, the secure network file access appliance 12 may direct the agent program 48 to regenerate and resubmit 157 application authentication information for specific user sessions. Thus, within an otherwise authenticated user session, the application authentication provisions of the present invention can enforce explicit and functional limitations on user process execution to a well defined set of authenticated applications.

[0064] Referring again to Figure 4, packet control information 114 and application information 122, exposed by packet processing 112 and as received from the agent monitoring process 118, is provided to the authentication and access control process 120 for each network file data packet received by the secure network file access appliance 12. Preferably, the authentication and access control process 120 includes a policy store representing the administratively determined, functionally supported operations of the secure network file access appliance 12. The policies are preferably stored in a high-performance hash table permitting a policy lookup against the information 114, 122 as presented to the authentication and access control process 120. Audit logs of the file requests, as well as error logs and logs of refused operations are produced by the authentication and access control process 120.

[0065] Policy sets applicable to a received network file packet can be progressively discriminated based on any of the data provided in the packet control information 114. In particular, IP layer data provides source and destination IPs, permitting specific access constrains to be defined against

- 22 -

defined clients, individually or by subnets. The standard NFS/CIFS layer data provides the requesting user UID and GID, as well as the fully qualified file or directory reference, including generally a mount point, file system path, and applicable file name. The application information 122 layer identifies the user session and provides the execution and parent process identifiers. Where utilized, the application information 122 layer also provides the application name and signature. Successful discrimination of the policy sets against the provided information 114, 122 enables and qualifies the processing of network file packets transported relative to the persistent storage resources 16.

[0066] Preferably, the handling of the various possible types of policy set discrimination failures is defined by the policy sets. Discrimination failures will typically include user authorization failures and unauthorized application execution attempts, unauthorized source IP addresses, and improper file references due to unavailability of the referenced file or lack of adequate user, group or file permissions. Depending on the nature of the failure, the discrimination failure handling defined by the policy sets will direct the production of detailed audit and error log entries and immediate issuance of administrative alarms, including potentially the automated generation of email and voice messages. The policy set discrimination failure handling preferably further defines the type and content of any NFS/CIFS network file error data packets generated by of the NFS/CIFS state machine 124 and returned to a computer system platform 14, 14'.

[0067] In accordance with the present invention, the progressive discrimination of the policy sets also determines the active application of encryption and compression to the packet payload data 116. For inbound network file data packets from computer system platforms 14, 14', any combination of data provided in the control information 114, 122 can be

- 23 -

utilized as a signature identifying whether the packet payload data is to be encrypted against a particular encryption key and compressed using a particular compression algorithm. A preferred basic policy set essentially defines the combinations of source IPs, user identifiers, and group identifiers permitted access through the mount point and, further, a default encryption key to be used, particularly for file creation. Multiple policy sets can be applicable to the same mount point, differing in the specification of source IPs, user identifiers, and group identifiers or by specification of additional control information, such as the path specification and file-type extension for the network file identified in the request. The policy sets are administratively managed to ensure that unique combinations of the provided control information resolve to distinct policy sets. Where path specification information is utilized to establish the scope of other wise matching policy sets, a best match of the path specification, file name, and file extension is preferably used to discriminate the default applicability of data encryption and compression.

[0068] Network file packets returned from persistent storage resources 16 are similarly processed 112 to expose the packet control information 114 and permit a combination of data to be considered in determining whether accompanying packet payload data requires decompression and decryption. While, in accordance with the present invention, encrypted network data packets returned from the persistent storage resources 16 can be presumed secure, examination of the control information 114 through authentication and access processing 120 enables an appropriate authentication of the source and sequence of the returned network file packets.

[0069] Preferably, packet payload data presented to the secure network file access appliance 12 and determined to be encrypted or compressed is processed into a sequence of logical access blocks (LABs) through an

- 24 -

encryption and compression process 126. As part of the encryption and compression process 126, each logical access block is, in accordance with one preferred embodiment of the present invention, marked with at least an indirect identifier of the applicable encryption key and compression algorithm. Thus, while the decompression and decryption status of outbound network data packets may be suggested by a source directory specification, the applicable encryption key and compression algorithm is determined based on the encryption and compression identifiers associated with the logical access blocks. Decryption and decompression of the logical access blocks are, therefore, not essentially dependent on the directory specification or other independently alterable aspects of the network file.

[0070] Discrimination of applicable policy sets is, in accordance with the preferred embodiments of the present invention, expanded through the support by the secure network file access appliance 12 of multiple, inbound virtual mount points for various persistent storage resources 16. As shown in Figure 6, multiple virtualized mount points /dev/hd_a, /dev/hd_b, /dev/hd_c, and /dev/td_d may be predefined administratively in the configuration of the secure network file access appliance 12 for the benefit of applications 162, 164 executed by the computer system platforms 14, 14'. These virtual mount points are independently associated through a defined mapping with the same, as by alias, or separate real mount points supported by various real storage resources 166, 168. Computer system platform 14 requests to mount any of the virtual mount point represented storage resources 166, 168 can be qualified and constrained by policy sets that, at a minimum, serve to validate the existence of the virtual mount point.

[0071] In accordance with the present invention, the virtual mount points further expand the ability to discriminate applicable access policy sets for the different applications 162, 164 executed by the computer system platform 14.

- 25 -

The control information 114 provided with each file request directed to the secure network file access appliance 12 identifies a file request corresponding target mount point. In accordance with the preferred embodiments of the present invention, the authentication and access control process 120 logically selects an applicable policy set based on the identified virtual mount point. The further constraints represented by the selected policy set are concurrently used to determine how the network file data packet is to be processed. For example, otherwise authorized applications 162, 164 accessing the storage resource 166 through the /dev/hd_a virtual mount point may be constrained to read-only file data transactions. The separate policy set associated with the /dev/hd_b virtual mount point may support read-write access by only a well defined set of UIDs, further constrained to file data requests originating from a defined subnetwork.

[0072] As another example, read-write access of the storage resources 166 by the application 164, administratively limited to providing backup services, may be specially supported through the virtual mount point /dev/hd_c. Preferably, the policy set associated with the mount point /dev/hd_c preferably enables read-write access to the storage resources 166 though specifically disallowing decryption of existing encrypted files. The policy set for the virtual mount point /dev/td_d, in complementary fashion, provides for the encryption and compression of unencrypted files read from the storage resources 166 for writing to the archival storage resources 168. The /dev/td_d policy set symmetrically limits the decryption of files read from the archival storage resources 168 to those encrypted by the backup application 164. Consequently, a user with limited backup access rights can fully administer the backup and restore of files without breach of the secure storage of previously encrypted files. Thus, distinguishing policy sets based on virtualized mount points provides an extensive degree of flexibility in managing

- 26 -

the access rights of the applications 162, 164 executed on behalf of a community of clients and users 20.

[0073] Network file packets permitted or refused by operation of the authentication and access control process 120 are signaled to an NFS/CIFS state machine 124, as shown in Figure 4. The sequences of network file packets representing select file data transactions, including specifically NFS/CIFS transactions, are tracked by the NFS/CIFS state machine 124, in accordance with the present invention, to support the selective encryption and compression of NFS/CIFS network packet transferred file data and manage the attendant changes in the size and structure of network files as stored by the persistent storage resources 16. Mount and unmount request RPCs are essentially atomic operations between the computer system platforms 14, 14' and the secure network file access appliance 12. On receipt of a mount request, access is optionally determined by the authentication and access control process 120 based on the applicable policy set and a determination that the underlying network storage resource 16 identified with the corresponding real mount point is available. An RPC response acknowledging the success or failure of the mount or unmount request is then returned.

[0074] The NFS/CIFS state machine 124 tracks the state of each NFS/CIFS transaction processed through the secure network file access appliance 12. The principle NFS/CIFS transactions tracked include Read, Write, and Create. All other NFS/CIFS defined transactions (generically Requests) are also tracked by the NFS/CIFS state machine 124. The Read transaction, following from an inbound read request for file data defined by an offset and range, involves building a corresponding read request with the read offset adjusted back to an encryption and compression block boundary and the range adjusted to allow for the encryption and compression of the file data through to the end of a block boundary. The next states include issuing

- 27 -

the read request to the persistent storage resources 16, receiving a responsive series of network read file data packets, and processing, as needed, to decrypt and decompress the received packet payload data. The final read transaction states include extracting the read file data for the originally requested offset and range and building and returning one or more network file data packets with the read file data.

[0075] An NFS/CIFS Write transaction requires a read/modify/write operation where existing stored file data is encrypted or compressed. A write transaction includes receiving a write request, building a lock request with a write lock offset adjusted back to an encryption and compression block boundary and the range adjusted to allow for the encryption and compression of the file data through to the end of a block boundary. The next transaction states include issuing a read request for any initial and final partial file-data page including the adjusted write offset and range terminus, decrypting, decompressing and modifying the read data page to include the corresponding parts of the file write data as received from the client, encrypting and, as appropriate, compressing the file write data, and building and issuing corresponding write requests to the storage resources 166. The final write states include building and sending an unlock request to the storage resources 166 and building and sending a write request reply to the client.

[0076] NFS/CIFS Requests, such as get and set attributes, get access permissions, and make directory, are generally atomic transactions managed by the secure network file access appliance 12 to support infrastructure compatibility with the storage resources 166. Request transactions involve receiving a client request and building and sending a corresponding request to the storage resources 166. Upon receipt of a request response from the storage resources 166, adjustments are made for the reported file size and other attributes of the network file as stored on the storage resources 166

- 28 -

depending on the particular request involved in the transaction. A corresponding request response is then constructed and sent to the client.

[0077] An NFS/CIFS Create transaction involves receiving a file create request, constructing a file management header for the new file, and building and sending a corresponding request to the storage resources 166. Upon receipt of a request response from the storage resources 166, a corresponding request response is again constructed and sent to the client.

[0078] Figure 7 provides a block diagram and flow representation of the software architecture 170 utilized in a preferred embodiment of the present invention. While the preferred embodiment utilizes separate ingress and egress network interfaces to the computer system platforms 14, 14', the two interfaces can share a single network interface adapter, provided the total bidirectional file transfer bandwidth is within the capabilities of the adapter. For either case, inbound network communications are processed through a first network interface 172. Network file data packets representing file transfer requests originating from applications 44, 46 are processed 174 to expose and deliver the network control information 114 for authentication processing 176. Application control information 122 collected from corresponding agent applications 48 are provided through an agent interface 178 in support of the authentication processing 176.

[0079] Based on interactions with a policy parser 180, selected elements of the network and application control information 114, 122 are compared with authentication parameters maintained in a policy data store 182. The policy parser 180 preferably implements decision tree logic to determine the level of authentication required for processing the network file request represented by the network file data packet received and whether that level of authentication has been met.

- 29 -

[0080] The network and application control information 114, 122 is also processed 184 to determine whether the authorized user is permitted access to the corresponding persistent storage resources 16. The policy processor 180 and policy data store 182 operate to determine whether the access attributes provided with the network file request are appropriate to enable access to the specific persistent storage resources 16 identified by the network file request.

[0081] While logically separate operations, the authentication and access processing 176, 184 are preferably performed concurrently. In a preferred embodiment of the present invention, a basic decision tree logic sequence considers the logical combination of network file operation requested, virtual mount point, target directory and file specification, user UID and GID, and the client/user session and process identifiers. Also considered is application authentication data provided with the network file request and as prior provided by the agent program 48 and the continuity state of the client session as periodically reported by the agent interface 178. Additional state data accumulated in relation to the nature, timing, and frequency of network file access requests is considered. This state data is accumulated by the secure network file access appliance 12 to support static time scheduling and quota controls over file access requests as well as dynamic traffic shaping of the file access operations processed through the secure network file access appliance 12. The accumulated state data also permits dynamic detection of patterns in file access requests that threshold qualify as intrusion attempts or other circumstances warranting issuance of an administrative alarm. The decision tree evaluation considers prior sequences of file access requests and thereby qualifies the permitted support of a current network file access request.

[0082] Policy data is administratively established to define the set of virtual mount points and the mapping of virtual mount points to real mount

- 30 -

points. The policy data can also variously define whether application authentication is to be enforced as a prerequisite for session execution or operative response by the secure network file access appliance 12, a limited, permitted set of authenticated digital signatures of execution or response enabled applications, whether user session authentication extends to spawned processes or processes with a different UID or GID, and other data that can be used to match or otherwise discriminate, in operation of the policy parser 180, against the control information 114, 122. This administratively established policy data is logically accessed from the policy store 182 by the policy parser 180 in the evaluation of the network and application control information 114, 122. For the preferred embodiments of the present invention, the decision tree logic and policy data are stored in a hash table permitting rapid evaluation of the network and application control information 114, 122.

[0083] The network and application control information 114, 122, as well as the determined results of the authorization and access processing 176, 184 are control inputs to an NFS/CIFS state machine process 186. Non-file data messages, including various NFS/CIFS request and reply messages involved in the read, write, and create NFS/CIFS transaction sequences, are prepared and forwarded 188, 190 directly from the state machine process 186 to the inbound network interface 172 and an outbound network interface 192. Policy data needed to support the generation of network file request and reply data packets, such as virtual to real mount point mapping data, is accessed from the policy data store 182 as needed.

[0084] Where file data is included in a network file data packet inbound from an application 44, 46, the packet payload data 116 is processed 194 into a sequence of logical access blocks (LABs), provided the network file data packet is qualified through access processing 184 for encryption or

- 31 -

compression. The packet payload data 116 of unqualified network file data packets are processed 194 unchanged into network data packets and provided to the network interface 192 for transmission through the security interposer 22 to the persistent storage resources 16.

[0085] As represented in Figure 8A, the packet payload data of network file data packets corresponds to read and written portions of a file 220 recognized by a file system 26. Individual packet payload data 222, generally as shown in Figure 8B, is preferably processed 194 into a sequence of logical access blocks 224_{1-N}, as shown in Figure 8C with each logical access block containing a corresponding portion of the packet payload data 222. In an initial embodiment of the present invention, the file management header 226 is virtualized for all files associated with a real mount point and locally stored by the hardware platform 70 effectively as part of the policy data held by the policy store 182. The applicable file management header is retrieved as part of the policy set applicable to the requested virtual mount point. The preferred embodiments of the present invention provide for the creation of a file management header 226 in connection with each Create file NFS/CIFS transaction. In one embodiment, the file management header 226 is created and written to the persistent storage resources 16 effectively as the first file data block as part of the creation of the file 220 on the persistent storage resources 16. One or more logical access blocks 224 can thereafter be appended to the file as created on the persistent storage resources 16 and, subsequently, read and written in random order. Alternately, to optimize the storage and retrieval of data with respect to the persistent storage resources 16, individual or subsets of logical access blocks 224 and the file management header 226 can be written to separate I/O pages within the same or different file spaces and storage devices. In either case, in accordance with the present invention, qualified file data reads and writes directed to the persistent storage

- 32 -

resources 16 are performed as discrete, logical access block-aligned transfers encompassing the offset and range of a client network file data request.

[0086] The file management header 226 and logical access blocks 224 are repackaged in network file data packets as otherwise ordinary blocks of file data for transport to the persistent storage resources 16. The encryption and/or compression of network file data by secure network file access appliance 12 is thus entirely transparent to the reading and writing of relative to the persistent storage resources 16 by operation of the present invention.

[0087] A preferred structure of the file management header 226 is shown in Figure 8D and further detailed in Table I below. Preferably, the file management header 226 includes a unique file GUID 228, security parameter index (SPI) 230, and a security signature 232. The file GUID 228 is preferably a SHA-1-based secure hash of data related to the file, such as the user UID, and file creation time to provide a 160-bit unique random identifier for the file. The security parameter index 230 is preferably a composite of security information including an encryption key identifier (Key) 234, a security options array (Idx) 236, and file related information (Info) 238.

[0088] The encryption key identifier 234 is preferably an encrypted representation of the encryption key name utilized to encrypt the file data contained in the logical access blocks of the file 220. Encryption key name/key value pairs are utilized by the secure network file access appliance 12 are administratively defined and stored in the policy data store 182. When, as a product of access processing 184, an encryption key is associated with a new file, the corresponding encryption key name is securely digested, again preferably using the SHA-1 algorithm, and stored in the key identifier field 234 of the file management header 226.

[0089] The security parameter index 230 may optionally also include a linked list storing, in encrypted form, the encryption key value for the file 220.

Each entry in the linked list includes a public key, encrypted key value tuple. The public key corresponds to a trusted encryption key agent server and the encrypted key value is encrypted with the public key of the agent. On retrieval of the network file data by a different secure network file access appliance 12', the public key identified agent server can be used to recover the encrypted key value. Providing support for multiple independent agent servers ensures that the encrypted key value can always be recovered.

[0090]

Table I
Management Header Structure

Struct MGT_BLOCK {	
U32 File_GUID[5];	// 160-bit unique random GUID for File
U32 Mgt_Hdr_Ver;	// 32-bit version identifier for this structure
U32 Size_Mgt_Blk;	// Size of the management block structure
U32 Options[];	// Option include
	// --IntegrityMode: to compare digital
signatures	
	//--OutOfBand: out-of-band meta-data
used	
	// --CypherName: encryption algorithm ID
	// --ComprName: compression algorithm ID
	// --UserEncryption: Key_GUID is a user key
	// --GroupEncryption: Key_GUID is a group
key	
	// --HaveKeys: has list of agent encrypted
keys	
U32 Key_GUID[5];	// 160-bit GUID for Key, generated by // SHA-1(KeyName)
U32 Creator_GUID[5];	// 160-bit GUID identifying the file creator
BYTE Init_Vector[8];	// Initial seed value for LAB encryption; // encryption seeds are a function of // Init_Vector + LAB Offset
U32 Padding[];	
U32 CRC;	// To verify management header block integrity
BYTE Signature[128];	// Signature, signed with PrivKey for

Table I
Management Header Structure

```

// PublicKey_Verify Pre-computed.
// Signs only static part of the structure to
// avoid overhead on each file under the
same
// volume/policy. CRC is signed as the last
part
// so that changing to any part of the whole
// block is detected.
*Key_Table // Linked list of Public Key, agent encrypted
// LAB Symmetric Key tuples
}

```

[0091] The security options array 236 provides an indexed list of the security functions applied to the logical access blocks 224 associated with file management header 226. These options preferably include identifiers of the whether encryption is used and the applicable encryption algorithm, whether compression is used and the applicable compression algorithm, whether the encryption key name lookup should be user or group based, whether an agent encrypted key list is present, and whether tamper detection through digital signature checking is to be enforced. The file related information 238 fields provide storage for various other information, such as a GUID corresponding to the file creator.

[0092] Finally, the security signature 232 provides storage for a cyclic redundancy check (CRC) value and digital signature. The CRC value is preferably computed over the binary value of the preceding portions of the file management header 226 to permit block integrity checking. The digital signature is computed for the preceding portions of the file management header 226 including the CRC field to enable detection of tampering with any portion of the file management header 226.

- 35 -

[0093] A preferred in-band structure of logical access blocks 224 is also shown in Figure 8D. The primary fields of a logical access block 224 include a LAB data field 240, a LAB signature field 242, and an optional LAB compression header 244. The LAB data field 240 contains an encrypted and/or compressed portion of the packet payload data 222. The size of the LAB data field 240 is nominally set as a multiple of a natural or convenient block size recognized by the file system 26 and further chosen for block encryption algorithm efficiency.

[0094] In accordance with the present invention, segmentation of the packet payload data 222 into the logical access blocks 224 enables reasonably sized blocks of file data to be encrypted and compressed as atomic units. Smaller segments sizes are preferred for obtaining relatively efficient random read/write operations directed to the file 220 as stored by random access devices within the persistent storage resources 16. Larger segment sizes are preferred for lower processing overhead, greater encryption and compression efficiency, and where the target device within the network storage resources 16 is a streaming access device, such as a conventional tape drive. Preferably, the packet payload data 222 segment size has a block modulo of eight bytes with a minimum size of 512 bytes and a nominally preferred size of 1024 bytes for random access devices. For streaming access devices, larger block sizes on the order of 8096 bytes may be preferred.

[0095] Where the last segment of the packet payload data 222 is less than the nominally preferred segment size, a smaller block size is used. This smaller block size is chosen to be the largest modulo eight byte block size that is the same or smaller than the size of the last segment. All but at most seven bytes of the last segment are then block encrypted. Any remaining segment bytes are then XORed with a mask value generated by the encryption of an

- 36 -

eight-byte length, zero-value string and then appended to the block encrypted portion of the last segment.

[0096] The LAB compression header 242, preferably included only where the packet payload segment held by the logical access block 224 is compressed, includes fields specifying the offset and range of the file data contained within the LAB data field 240. Dependent on the underlying data values and the stream compression algorithm applied, the segment length or range of the packet payload data 222 stored in the LAB data field 240 is variable. The segment length is manipulated to obtain compressed data that closely approaches the preferred LAB data field size. Padding is provided to reach a modulo eight-byte encryption block compatible size. At a minimum, the range value identifies the actual compressed data carried in a completed logical access block 224.

[0097] The LAB signature 244 is preferably computed as a secure digest of the LAB data field 240 and, where present, the LAB compression header 242. In the preferred embodiments of the present invention, an SHA-1 algorithm is used to create the LAB signature 244. The security of each logical access block 244, when retrieved to the secure network file access appliance 12, can be assured against tampering by recomputing the secure digest of the LAB data field 240, including any LAB compression header 242, and comparing against the LAB signature 244. For a preferred variant of the present invention, network file data is stored as logical access blocks 224 containing only unencrypted, uncompressed LAB data 240 and LAB signatures 244. While the efficiency of random access over network file data is maintained, modifications potentially due to improper tampering with the contents of the network file are nonetheless detectable on an individual logical access block 224 level. The conventional necessity of reading the entire network file to compute a secure digest to detect tampering is not required.

- 37 -

[0098] In an alternate embodiment of the present invention, an error correction trailer 246 is provided to store an ECC value computed over the LAB data field 240, any LAB compression header 242 and the LAB signature 244. ECC values are computed on creation of the logical access blocks 244. Upon retrieval of logical access blocks 244, the ECC value is used to correct bit errors that may occur as a consequence of extended network infrastructure transport of the logical access blocks 244. In particular, bit errors may be introduced by network routers operating at the TCP layer and above. Such infrastructure induced bit errors are otherwise detected from the LAB signature 244, but are then indistinguishable from data tampering. Use of the error correction field 246 serves to independently protect the integrity of the logical access blocks 244.

[0099] The file management header 226 and the headers 244 and trailers 242, 246 of the logical access blocks 244 may be included in-band, or in-file, as generally represented in Figure 8D, as part of the file 220 as ultimately stored by the persistent storage resources 16. Different in-band layouts can also be used to optimize access to the logical access block data 240. The file management header 226, digital signatures 242, and compression headers 244 can be collected into one or more in-band super blocks. The size of these super blocks and the remaining logical access block data 240 can be sized to optimize I/O performance of the persistent storage resources 16.

[0100] Alternately, and potentially preferred, only the logical access block data 240 is stored by the persistent storage resources 16 in-band as the network file 220. The file meta-data, including the management header 226 and the headers 244 and trailers 242, 246, corresponding to a network file 220 are stored in a separate, meta-data or shadow file. Any parallel storage structure that maintains the relationship between the shadow file and the in-

- 38 -

band network file 220 may be used. The shadow files can be created and stored on the network resources 16 within the same storage space as the network files 220, within a different storage space potentially physically remote from the network files 220, or on the platform 70 provided the parallel association of the shadow files with the network files 220 is maintained. For example, shadow files can be stored in the same directory with the counterpart network files 220 and identified by file names that are a defined permutation of the network file 220 file names. The shadow files can alternately be stored in a parallel directory structure diverging from a defined root or relative root node of the persistent storage resources 16. In either case, the defined relationship between the shadow files and the corresponding network files 220 is determined and known to the secure network file access appliance 12, which can ensure the parallel reading and writing of the shadow files with corresponding reading and writing of the network files 220.

[0101] Referring again to Figure 7, the packet to LAB processing 194 preferably utilizes, as required, the hardware accelerators 62 to perform encryption 196 and compression 198 over the segments of packet payload data 222. The logical access blocks 224_{1-N} , together containing the packet payload data 222 of a network file data packet, are then collected into a new network file data packet and passed to the network interface 192 for transport to the networks storage resources 16.

[0102] Network file data packets received through the network interface 192 are similarly processed 200 to expose and deliver the network control information 114 for authentication and access processing 176, 184 and logical access blocks 224_{1-N} contained in the packet payload data to a logical access block to packet data process 202. The provision for authentication and access processing 176, 184 permits even distributed, potentially client-based network storage devices to be equally secured and made accessible as other

- 39 -

persistent storage resources 16. In the preferred embodiments of the present invention, minimal authentication and access processing 176, 184 is performed for network file data packets received from dedicated persistent storage resources 16.

[0103] The logical access blocks 224_{1-N} received in the packet payload data are processed 202 to apply error correction, where the error correction field 246 is present, and validate the integrity of the LAB data fields 240, including the LAB compression headers 244 if present, against the digital signature 242 values. The file management header 226 is read, typically in advance, by the NFS/CIFS state machine process 186 to obtain the encryption key identifier from the field 234 and compression algorithm identity, if applicable from the options index field 236. The LAB data fields 240 are then decompressed 204, if applicable, and decrypted 206. The NFS/CIFS state machine process 186, based on the pending inbound file data read request transaction, identifies an offset and range-selected portion of the combined logical access block 224_{1-N} data representing client read requested data. The selected data is then incorporated into a network file data packet and provided to the network interface 172 for transport to the transaction identified application 44, 46.

[0104] For the preferred embodiments of the present invention, an administration interface 208 provides access to and configuration of the policy parser 180 and policy data store 182. A network communications interface 210 provides access to the administration interface 208 independent of the inbound and outbound network interfaces 172, 192.

[0105] The software architecture 170 is preferably extended, as shown in Figure 9, to provide additional security appliance-oriented features. The extended architecture 250 includes IP filter layers 252, 254 implementing firewall-type filtering for network connections made through the network

- 40 -

interfaces 172, 192. A filter rules store 256 preferably maintains iptables-type specifications that define the IP addresses, network protocols, and internet ports permitted to pass network packets through the IP filter layers 252, 254. Preferably, the IP filter layers 252, 254, and particularly the inbound IP filter layer 252, is set to reject all connections except those pertaining to network file access operations, including the NFS, CIFS, RPC, and mount protocols. These network file data packets passed by the IP filter layers 252, 254 are directed for packet/LAB processing 258 as performed by the software architecture 170. Unauthorized connection attempts and access requests lacking adequate policy-based permissions are therefore preferentially received, detected, and audited by the software architecture 170.

[0106] The flexible analysis capabilities of the authentication and access controls 176, 184 and policy parser 180, particularly based on access to the full set of control information 114, 122, allows a more refined identification of potential abuse patterns and a wider variety of remedial actions, including dynamically blocking specific source IPs, logging detailed information, and issuing real-time administrative alerts. The security and reporting strength of the firewall filters 252, 254 is appropriate for handling connection attempts unrelated to the primary functions of the secure network file access appliance 12. The firewall filters 252, 254 may also be utilized to proxy selected network data packets, including potentially network file data packets, through the secure network file access appliance 12, utilizing a bypass route 260. In the case of VPN 42 and network file access appliance 12' designated source IP addresses and protocols can be identified and appropriately bypassed 260.

[0107] For the fixed scale, hardware platform 70, the firewall filters 252, 254 are preferably implemented through the kernel execution of the operating system iptables module by the main processors 54. On the scalable hardware platform 80, the firewall filter layers 252, 254 are preferably implemented on

- 41 -

the ingress and egress processors 86, 88, with the bypass routed network packets being passed directly between the ingress and egress processors 86, 88. The filter rules maintained in the filter rules store 256 are administered through the administration interface 208.

[0108] An NFS/CIFS read transaction 270, structured in accordance with a preferred embodiment of the present invention, is shown graphically in Figure 10A. A read target file, consisting of a file management header 226 and a sequence of logical access blocks 224_{1-N} , exists on the persistent storage resources 16. In general, an inbound read request identifies an offset and range of data to read 272. Outbound read requests are issued to read 274, 276 the file management header 226 and an encompassing, block-aligned sequence of logical access blocks 224_{A-X} . The read request 276 retrieves the requested logical access blocks 224_{A-X} in a series of one or more network file data packets, which are then processed to complete the inbound read request by returning one or more network file data packets containing the read request data 272.

[0109] The specific processing 280 associated with an NFS/CIFS read transaction 270 is shown in Figure 10B. The secure network file access appliance 12, on receiving a firewall-filtered file data read request, exposes 282 and parses 284 the network control information 114 against the policy rules and data 182, 184. A policy compliance failure is reported 286 by return issuance of an NFS/CIFS appropriate reply network data packet.

[0110] Where the read request complies with the defined policy requirements, the file related access control information is optionally read 288 from the persistent storage resources 16 to confirm existence of the file and evaluate applicable read data permissions. Where the permissions check is performed and fails, nonexistence of the file or inadequate permissions are reported 286 without issuing the read file request to the persistent storage

- 42 -

resources 16. The file meta-data, including the file management header 226 for the request target file, is also read 288 from the network storage resource 16. A block-aligned logical access block offset 290 and range 292 are determined and used to create and issue an outbound read request directed to the persistent storage resources 16. The read data offset is adjusted to account for the size of the file management header 226 as stored at the beginning of the file. Where the logical access blocks 224_{A-X} contain compressed data, file data reads of the LAB compression headers 244 may be required to determine adjustments to both the read data offset and an encompassing read request range.

[0111] As the requested logical access blocks 224_{A-X} are received 294, error correction is applied 296, depending on whether the LAB ECC field 246 is present, decrypted 298 utilizing the key associated with the key name determined from the key identifier field 234 of the file management header 226, and decompressed 300, depending on whether the file management header 226 includes the compression option and identifies a corresponding algorithm. The LAB digital signatures 242 are used to check the integrity of the retrieved file data. A failure of the integrity check for any of the logical access blocks 224_{A-X} may result in a re-reading of some or all of the logical access blocks 224_{A-X} , to protect against soft-errors, with persistent errors being ultimately reported by the return issuance of an NFS/CIFS appropriate error network data packet. Preferably, both soft and persistent errors are logged by the secure network file access appliance 12. Persistent errors, recognized through the operation of the NFS/CIFS state machine processing 186 of the inbound read request, are further preferably asserted against the policy parser 180 for evaluation and subsequently issued 302 as a tampering alert message through the administrative interface 208. Finally, as file data is received and processed in response to the outbound read request, the file data identified in

- 43 -

the inbound read request is assembled 304 into one or more reply network file dat packets and returned.

[0112] An NFS/CIFS create file transaction 310, as shown graphically in Figure 11A, preferably operates to create a new file containing a new file management header 226. As further detailed in Figure 11B, a create file request process 320 initially exposes 322 and parses 324 the network control information 114, with any policy compliance failures resulting in the return issuance of an NFS/CIFS appropriate reply network data packet. Provided the file create request complies with the defined policy requirements, directory information is optionally read 328 from the persistent storage resources 16 to obtain the target file creation permissions. Where the permissions check is performed and fails, non-existence of the target directory and inadequate permissions are reported 326 without asserting a create file request to the persistent storage resources 16.

[0113] A file management header 226 is then created 330. Through operation of the NFS/CIFS state machine processing 186, the policy parser 180, based on the stored values provided from the policy data store 182, generates and provides the necessary values for the security parameter index 230. In particular, the policy parser 180 preferably associates encryption keys and compression choices against directory specifications, including mount points. Thus, the target location of the file to be created is utilized to determine whether encryption and compression are to be applied and the applicable key and algorithms for implementation. A secure identifier based on the key name and compression and compression algorithm identifiers are computed and stored in the new file management header 226 along with computed CRC and signature values.

[0114] The NFS/CIFS state machine 186 next provides for the creation and issuance 332 of an NFS/CIFS create file request to the persistent storage

- 44 -

resources 16 utilizing the directory specification provided by the inbound create file request. For in-band storage of the file management header 226, an NFS/CIFS file write request, containing the file management header 226, is then created and issued 334 to the persistent storage resources 16. Where a shadow meta-data file is designated for use, an NFS/CIFS file create and write requests, the latter containing the file management header 226, are created and issued 334 to the persistent storage resources 16 to create the shadow file. Finally, an NFS/CIFS appropriate create file reply network data packet is returned to the client.

[0115] An NFS/CIFS write transaction 340, structured in accordance with a preferred embodiment of the present invention, is shown graphically in Figure 12A. The write of file data to an existing file in the persistent storage resources 16 uses a read, modify, write procedure. An inbound write data request specifies an offset and range of write data 342 that is provided in a transaction sequence of one or more network file data packets. In most instances, the write request data will be unaligned to the logical access blocks 224_{1-N} existing in the stored file. The file management header 226 and any partially overlapped logical access blocks 224_A , 224_X are preemptively read 344, 346, 348, permitting the overlapped logical access blocks 224_A , 224_X to be decrypted and decompressed as required. An overlay of the inbound write data 342 with the block-aligned read data is then performed. The resulting block-aligned write data is then processed into logical access blocks 224_{A-X} and written 350 in a write transaction sequence of one or more network file data packets to the persistent storage resources 16.

[0116] The preferred process 360 of performing an NFS/CIFS write request transaction is shown in Figure 12B. The received write file data request is received and processed 362 to expose the network control information 114. This information is then parsed 364 against the established policies 180, 182,

- 45 -

with any compliance failures being reported 386. The network control information 114 is then further processed 368 to identify the target file stored by the persistent storage resources 16, create and issue read requests to obtain the file meta-data, including the file management header 226. The logical access block offset and range are then determined 370, 372, adjusting as needed for the presence of the file management header 226 and compression of the logical access block 224 contained data. A file lock is asserted against the range logical access blocks 224_{A-X} . The initial and terminal logical access blocks 224_A , 224_X are read 374 from the persistent storage resources 16, corrected 376 if the LAB ECC field 246 is present, decrypted 378, and decompressed 380, as needed. Integrity failure errors are reported 382. Data from the terminal logical access blocks 224_A , 224_X are merged 384 with the write data 342 and the combined data is resegmented 386, compressed 388 as appropriate, and encrypted 390. As applicable, LAB ECC values are computed and added 392 to the assembled 394 series of logical access blocks 224_{A-X} . As the logical access blocks 224_{A-X} are assembled, one or more write network file data packets are constructed and sent to the persistent storage resources 16. Once the writing the logical access blocks 224_{A-X} has completed, the file lock is released.

[0117] Thus, a secure file system server and methods of efficiently and reliably securing persistently stored data accessible through the file system sever has been described. In view of the above description of the preferred embodiments of the present invention, many modifications and variations of the disclosed embodiments will be readily appreciated by those of skill in the art. It is therefore to be understood that, within the scope of the appended claims, the invention may be practiced otherwise than as specifically described above.

- 46 -

Claims

- 1 1. A secure server platform providing protected access to persistent data,
2 said secure server platform comprising:
3 a) a persistent data store supporting storage of predetermined
4 files;
5 b) a computer server system, including an operating system
6 kernel supporting execution of an application program with respect to a user,
7 a first file system layer, responsive to predetermined data transfer operations
8 managed by said operating system kernel relative to said application program
9 to transfer file data with respect to said persistent data store, and a second file
10 system layer interposed between said operating system kernel and said first file
11 system layer, said second file system layer establishing a file access control
12 function selectively constraining said predetermined data transfer operations;
13 and
14 c) a file access controller coupled to said second file system layer
15 to support said file access control function, wherein said file access controller
16 implements an access policy list establishing a predetermined correlation
17 between said user and said predetermined files determinative of permitted
18 ones of said predetermined data transfer operations, and wherein said access
19 policy list is maintained within said file access controller independent of said
20 computer server system.

- 1 2. The secure server platform of Claim 1 wherein said second file system
2 layer includes an operating system interface coupled to said operating system
3 kernel to enable collection of authorization information with respect to said
4 user and session information with respect to said application program and
5 wherein said file access control function, in response to said predetermined

- 47 -

6 data transfer operations, selectively provides said authorization and session
7 information to said file access controller.

1 3. The secure server platform of Claim 2 wherein said authorization
2 information includes a secure identification of said user and said session
3 information includes a session key identifying the session, established by said
4 operating system kernel, within which said application program executes.

1 4. A secure server system, supporting application access to file data
2 persistently maintained in a file data store, said secure server system
3 comprising:

4 a) a cryptographic data processor, responsive to session data, operative
5 to encrypt and decrypt file data streams dependent on encryption keys, said
6 cryptographic data processor including an encryption key store; and

7 b) a computer system, coupled to said cryptographic data processor,
8 including an operating system kernel and a file system responsive to file data
9 requests provided by said operating system kernel with respect to a
10 predetermined file data store,

11 said file system being operative to route, with respect to a
12 predetermined file data request, a file data stream through said cryptographic
13 data processor, said file system identifying predetermined session data with
14 said predetermined file data request to selectively enable the encryption and
15 decryption of said file data stream, and wherein said encryption key store is
16 secure against access through said file system.

1 5. The secure server system of Claim 4 wherein said cryptographic data
2 processor includes first and second access interfaces, said first access interface
3 being coupled to said file system and said second access interface providing

- 48 -

4 management access to said encryption key store independent of said first
5 access interface.

1 6. The secure server system of Claim 5 wherein said first access interface
2 provides for the bidirectional transfer of file data streams.

1 7. A method of securing the transfer of persistently stored data between
2 a computer system and a persistent data store, wherein said computer system
3 includes a processor supporting the execution of an operating system kernel
4 and a file system to transfer data with respect to said persistent data store, said
5 method comprising the steps of:

6 a) associating session information with data transfer requests provided
7 from said operating system kernel;

8 b) providing said session information and said data transfer requests
9 to an independent security processor system to determine permitted data
10 transfer operations; and

11 c) routing, between said operating system and said file system, the
12 persistently stored data corresponding to said permitted data transfer
13 operations exclusively through said independent security processor system.

1 8. The method of Claim 7 further comprising the step of acquiring said
2 session information through said operating system kernel.

1 9. The method of Claim 8 further comprising the step of decrypting, by
2 said independent security processor system, persistently stored data retrieved
3 through said file system, said step of decrypting utilizing decryption keys
4 identified only by reference by said data transfer requests.

- 49 -

1 10. The method of Claim 9 wherein said step of decrypting includes
2 retrieving a predetermined decryption key in correspondence with a
3 predetermined data transfer request, wherein said predetermined decryption
4 key is retrieved independent of said computer system.

1 11. The method of Claim 10 wherein, in performance of said step of
2 decrypting, said predetermined encryption key is held by said independent
3 security processor system inaccessible from said computer system

1 12. A method of securing persistently stored file data stored by a persistent
2 storage device and accessible through a computer server system including a
3 processor, an operating system executable by said processor including an
4 operating system kernel providing a first file system interface and a second file
5 system interface to said persistent storage device, said method comprising the
6 steps of:

- 7 a) coupling an independently operating encryption processor
8 supporting the transparent encryption and decryption of persistent file data to
9 said computer server system through a defined communications channel;
- 10 b) associating session information obtained through said operating
11 system kernel with a predetermined persistent file data transfer request
12 identifying predetermined persistent file data;
- 13 c) providing said predetermined persistent file data transfer request and
14 said session information to said independently operating encryption processor
15 through said defined communications channel; and
- 16 d) routing said predetermined persistent file data as transferred between
17 said first and second file system interfaces through said independently
18 operating encryption processor.

- 50 -

1 13. The method of Claim 12 further comprising the step of maintaining
2 inaccessible, by said independently operating encryption processor from said
3 computer server system through said defined communications channel, an
4 encryption key associated with said predetermined persistent file data transfer
5 request.

1 14. The method of Claim 13 further comprising the step of determining, by
2 said independently operating encryption processor, said encryption key from
3 said predetermined persistent file data transfer request.

1 15. The method of Claim 14 wherein said step of determining
2 predetermines the access authorization of said predetermined persistent file
3 data transfer request with respect to said session information and an access
4 policy store maintained by said independently operating encryption processor.

1 16. The method of Claim 15 wherein said step of maintaining further
2 maintains inaccessible, by said independently operating encryption processor
3 from said computer server system through said defined communications
4 channel, said access policy store.

1 17. The method of Claim 16 wherein said step of maintaining inaccessible
2 prevents access to said encryption key and said access policy store through
3 said defined communications channel independent of a security breach of said
4 computer server system.

1 18. A computer system implementing a server architecture enforcing secure
2 authentication and access control to file data, said computer system
3 comprising:

- 51 -

4 a) a first processor system, coupled to said persistent file data store,
5 including a server memory area and a server processor providing for the
6 execution of an operating system within said server memory, wherein said
7 operating system includes an operating system kernel, having a kernel data
8 access interface, a filesystem switch, supporting routing of file data requests
9 provided from said operating system kernel, and a security interposer layer
10 coupled to said kernel data access interface and said filesystem switch, said
11 security interposer layer being responsive to received file data requests to
12 associate session information obtained from said operating system kernel with
13 said file data requests; and

14 b) a second processor system coupled to said security interposer layer
15 and responsive to said session information as associated with said file data
16 requests to selectively enable said file data requests relative to the transfer of
17 file data in response to said file data requests.

1 19. The computer system of Claim 18 wherein said second processor
2 system includes a security processor providing for the execution of a control
3 program that implements an access policy list enabling said security processor
4 to identify, based on said session information, permitted ones of said file data
5 requests, wherein said second processor system interoperates with said security
6 interposer layer to enable said permitted ones of said file data requests.

1 20. The computer system of Claim 19 wherein said control program further
2 implements an authorization control list and wherein said second processor
3 system interoperates with said security interposer layer to enable the execution
4 of a predetermined application program within said server memory area by
5 said server processor.

- 52 -

1 21. The computer system of Claim 20 wherein said security interposer layer
2 presents a file system interface to said filesystem switch and wherein said
3 security interposer layer is coupled to said filesystem switch as the exclusive
4 route path for predetermined file data requests.

1 22. The computer system of Claim 21 wherein said security interposer layer
2 is coupled to said filesystem switch as the exclusive route path for said file data
3 requests received by said filesystem switch.

1 23. The computer system of Claim 20 wherein said security interposer layer
2 utilizes an RPC-protocol to communicate with said second processor system.

1 24. The computer system of Claim 23 wherein said security interposer layer
2 utilizes an NFS protocol to communicate with said second processor system.

1 25. The computer system of Claim 18 wherein said security interposer layer
2 provides for the transfer routing of file data through said second processor
3 system and wherein said second processor system includes a security processor
4 operative to decrypt file data transferred through said second processor
5 system.

1 26. The computer system of Claim 25 wherein file data is bidirectionally
2 transferred between said security interposer layer and said second processor
3 system.

1 27. The computer system of Claim 26 wherein said second processor
2 system stores an encryption key for use by said security processor in decrypting
3 file data, wherein said security interposer layer provides a reference to said

- 53 -

4 encryption key determined from said session information to said second
5 processor system, and wherein second processor system maintains said
6 encryption key independent of said security interposer layer.

1 28. The computer system of Claim 27 wherein said security processor
2 provides for the execution of a control program that implements an access
3 policy list enabling said security processor to identify, based on said session
4 information, permitted ones of said file data requests, wherein said second
5 processor system interoperates with said security interposer layer to enable
6 said permitted ones of said file data requests.

1 29. The computer system of Claim 28 wherein said control program further
2 implements an authorization control list and wherein said second processor
3 system interoperates with said security interposer layer to enable the execution
4 of a predetermined application program within said server memory area by
5 said server processor.



1 30. The computer system of Claim 29 wherein said security interposer layer
2 presents a file system interface to said filesystem switch and wherein said
3 security interposer layer is coupled to said filesystem switch as the exclusive
4 route path for predetermined file data requests.

1 31. The computer system of Claim 30 wherein said security interposer layer
2 is coupled to said filesystem switch as the exclusive route path for said file data
3 requests received by said filesystem switch.


1 32. The computer system of Claim 31 wherein said security interposer layer
2 utilizes an RPC-protocol to communicate with said second processor system.

- 54 -

- 1 33. The computer system of Claim 32 wherein said security interposer layer
- 2 utilizes an NFS protocol to communicate with said second processor system.



2/11






FIG. 4




FIG. 5A

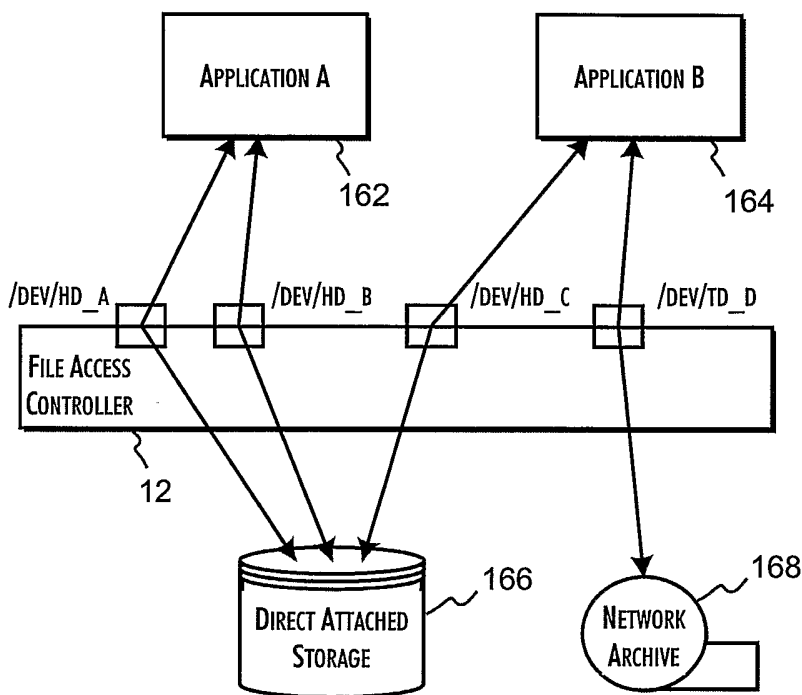


FIG. 6

5/11




FIG. 7

FIG. 8A




FIG. 8B




FIG. 8C


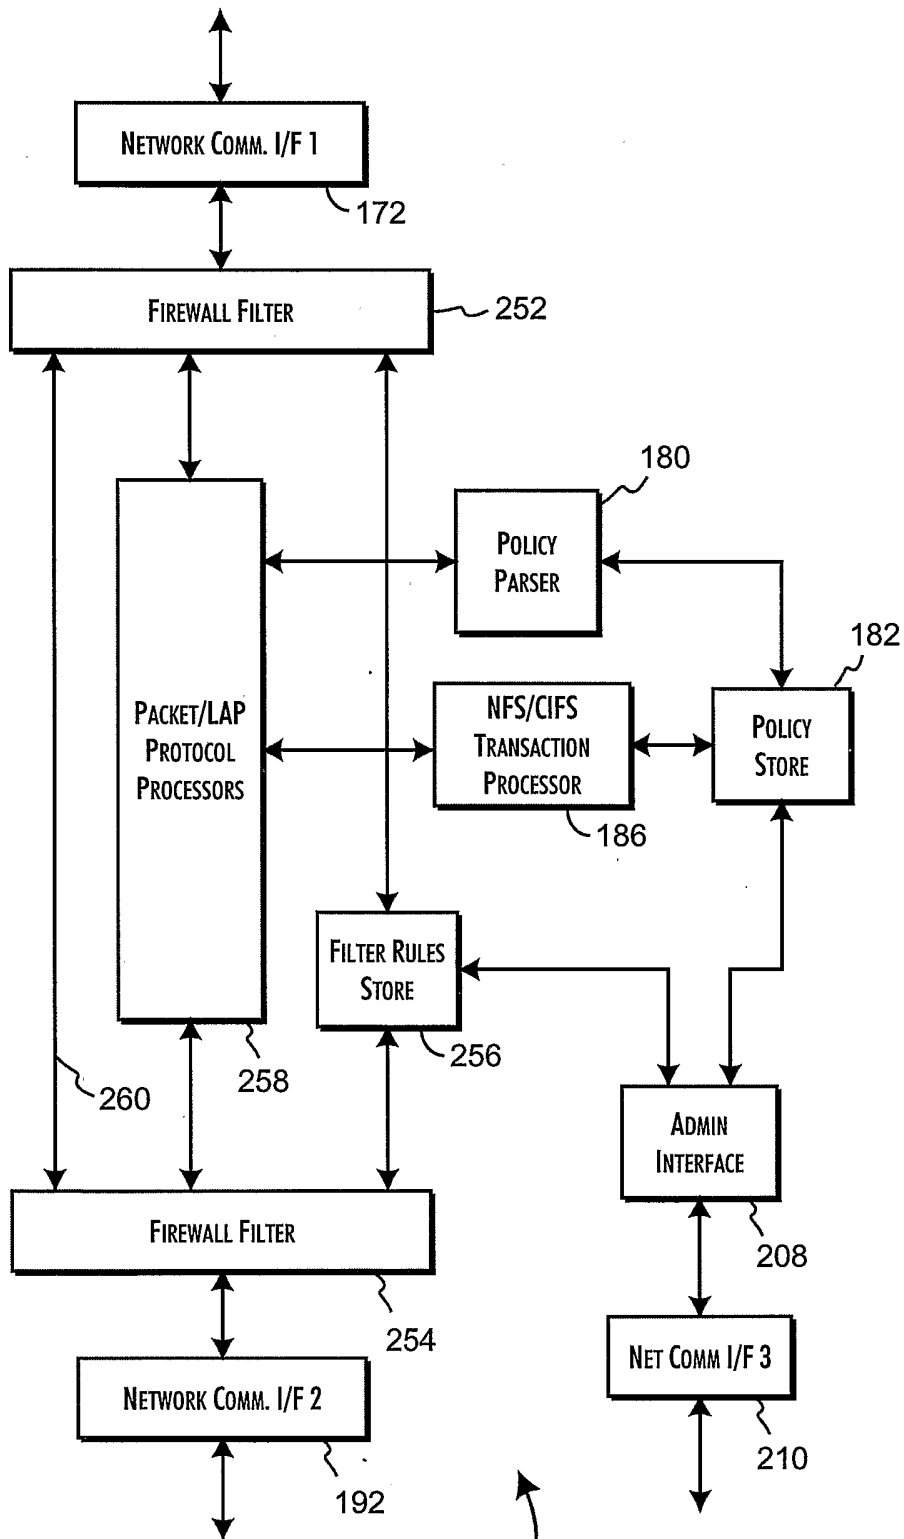


FIG. 8D

7/11



250 FIG. 9




FIG. 10A

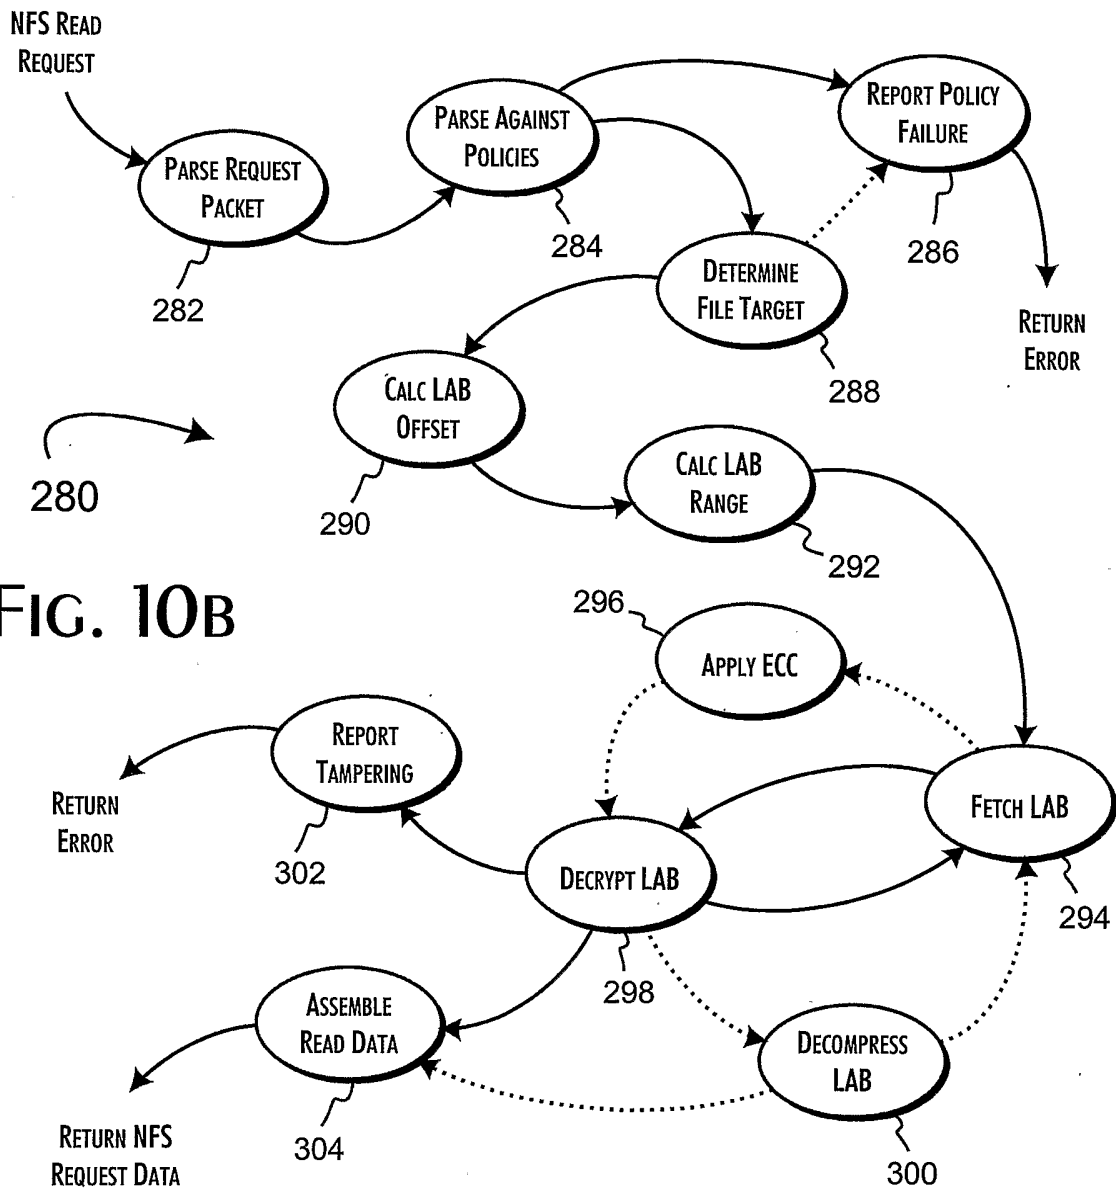


FIG. 10B

9/11

FIG. 11A

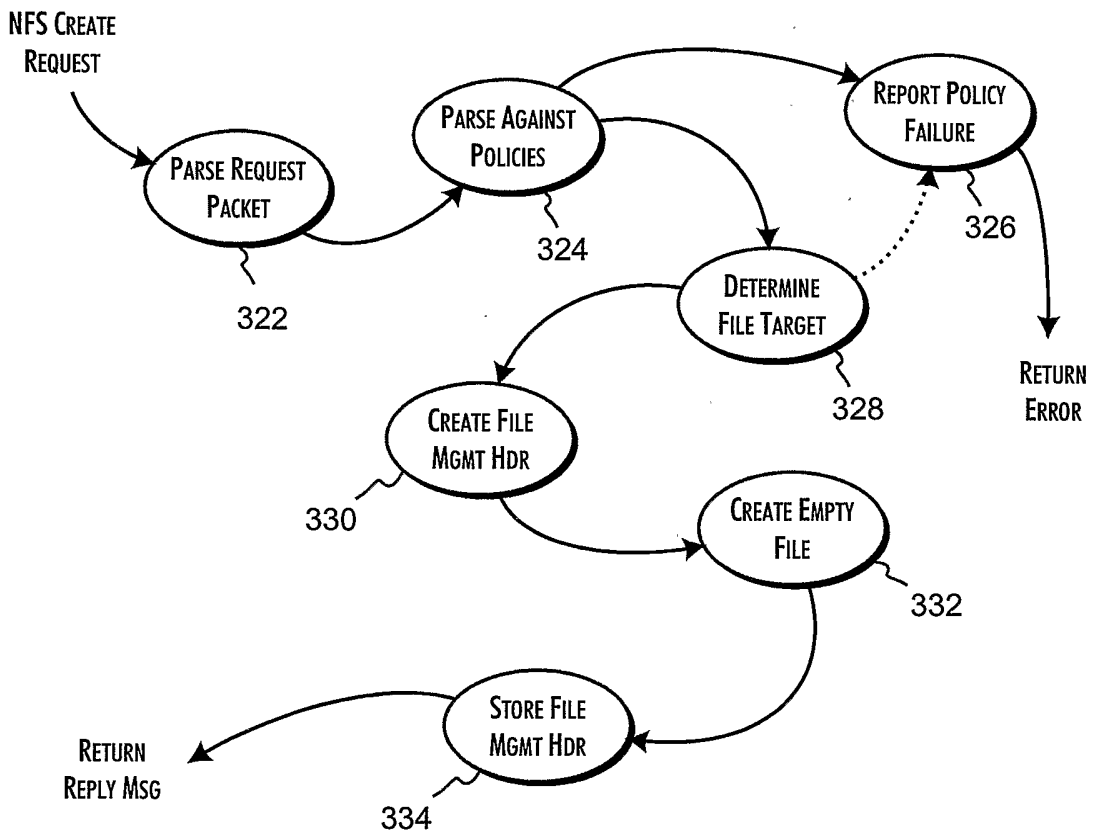






FIG. 11B





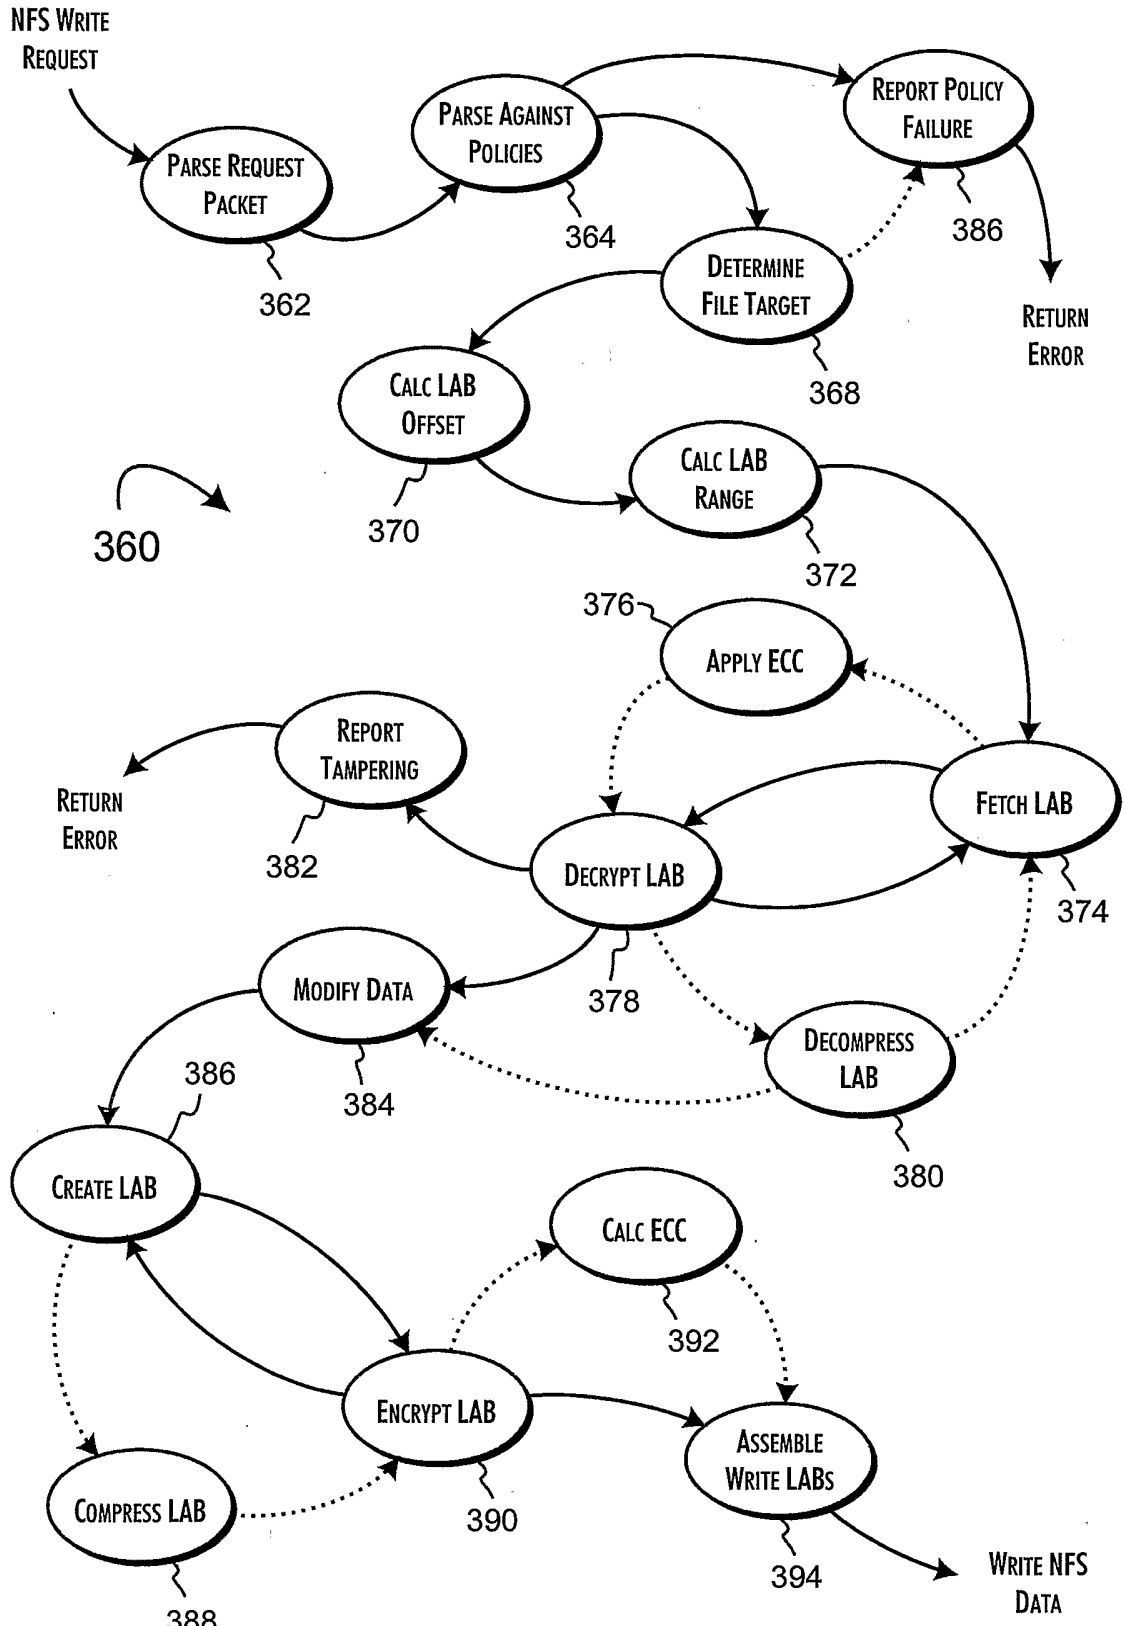


FIG. 12B