

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
21 November 2002 (21.11.2002)

PCT

(10) International Publication Number
WO 02/093314 A2

(51) International Patent Classification⁷: **G06F**

(21) International Application Number: PCT/US02/15421

(22) International Filing Date: 14 May 2002 (14.05.2002)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/292,088 17 May 2001 (17.05.2001) US

(71) Applicant (for all designated States except US): **DECRU, INC.** [US/US]; 275 Shoreline Drive, Suite 450, Redwood City, CA 94065 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **AVIDA, Dan** [IL/US]; 2312 Casa Bona Avenue, Belmont, CA 94002 (US). **PLOTKIN, Serge** [IL/US]; 573 Villa Real, Palo Alto, CA 94306 (US).

(74) Agents: **GLENN, Michael, A.** et al.; Glenn Patent Group, 3475 Edison Way, Suite L, Menlo Park, CA 94025 (US).

(81) Designated States (national): AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, DE, DK, DM, EE, 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, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW.

(84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

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.



WO 02/093314 A2

(54) Title: ENCRYPTION BASED SECURITY SYSTEM FOR NETWORK STORAGE

(57) Abstract: The presently preferred embodiment of the invention provides an encryption based security system for network storage that separates the ability to access storage from the ability to access the stored data. This is achieved by keeping all the data encrypted on the storage devices. Logically, the invention comprises a device that has two network interfaces: one is a clear text network interface that connects to one or more clients, and the other is a secure network interface that is connected to one or more persistent storage servers. Functionally, each network interface supports multiple network nodes. That is, the clear text network interface supports multiple client machines, and the secure network interface supports one or more storage servers.

Encryption Based Security System for Network Storage

5

BACKGROUND OF THE INVENTION

TECHNICAL FIELD

The invention relates to data security. More particularly, the invention relates
10 to an encryption based security system for network storage.

DESCRIPTION OF THE PRIOR ART

Computers are connected to storage devices such as disks and disk arrays by
15 network connections such as Ethernet, Fibre Channel, SCSI, iSCSI, and
Infiniband. Such connections use packet-based protocols to send data,
commands, and status information between computers and storage devices.
The data stored on such storage devices is often of a proprietary nature, and
the owner of such data wishes to prevent unauthorized users from reading or
20 modifying the data.

In the case of networked computer storage, unauthorized users can in many
cases gain access to the data stored in such devices. It would be
advantageous to provide a system that prevents unauthorized users from
25 understanding the data. It would also be advantageous if such system

enables both detection of data modification by unauthorized users. At the same time, to simplify integration of such a device with a plurality of computers and storage devices, it would be desirable if such system operated in a completely transparent fashion, whereby no modification is required to
5 either the computers or the storage devices.

SUMMARY OF THE INVENTION

The presently preferred embodiment of the invention provides an encryption
10 based security system for network storage that separates the ability to access storage from the ability to access the stored data. This is achieved by keeping all (or part of) the data encrypted on the storage devices. Logically, the invention comprises a device that has two network interfaces: one is a clear text network interface that is used to communicate to one or more clients, and
15 the other is a secure network interface that is used to communicate with one or more persistent storage servers. Functionally, each network interface supports multiple network nodes. That is, the clear text network interface supports multiple client machines, and the secure network interface supports one or more storage servers.

20

The invention is preferably a device, placed in the network, on the data path between the computer and the storage device, which intercepts all the packets that flow between the computer and storage device. The device distinguishes between data, command, and status. The device encrypts,
25 using a key, the data sent from the computer and decrypts the data sent from

the storage device, while passing through without modification command and status information.

In an alternative configuration, the device resides on the logical data path. In this case the client computers communicate to it as if it was a storage device (or several storage devices) while the device itself communicates with the storage devices as if it is a client computer.

The device can use a plurality of keys to encrypt and decrypt data, and a methodology to select the key based on user identification, data location on the storage device, file name, permission structure and other factors. Multiple such devices can share some or all keys and key use methodology between them, such that some or all of the data encrypted by one such device can be decrypted by another such device.

15

Such devices are able to operate in a transparent fashion to both the computers and storage systems that are exchanging data through the devices, such that no modification is required to either computers or storage devices to enable the storage of encrypted data on the storage device, and the subsequent retrieval and decryption of such data.

20

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a block schematic diagram of a typical networked storage architecture;

25

Fig. 2 is a block schematic diagram of an inline deployment architecture according to the invention;

Fig. 3 is a block schematic diagram showing a deployment using a switch
5 according to the invention;

Fig. 4 is a block schematic diagram of a Type-1 configuration according to the invention;

10 Fig. 5 is a block schematic diagram of a Type-2 deployment according to the invention;

Fig. 6 is a block schematic diagram of a Type 3 deployment according to the invention;

15

Fig. 7 is a block schematic diagram showing remote data sharing in NAS according to the invention

Fig. 8 is a block schematic diagram of an device and access controller
20 configuration according to the invention;

Fig. 9 is a block schematic diagram of a direct-connect configuration according to the invention;

25 Fig. 10 is a block schematic diagram of a single switch solution according to the invention;

Fig. 11 is a block schematic diagram of a dual switch or zoned single switch solution according to the invention;

- 5 Fig. 12 is a block schematic diagram showing remote mirroring for private and shared disaster recovery according to the invention;

Fig. 13 is a block schematic diagram of an device according to the invention;

- 10 Fig. 14 is a block schematic diagram of a hardware configuration of an device according to the invention;

Fig. 15 is a block schematic diagram of a system according to the invention;

- 15 Fig. 16 is a block schematic diagram of a checksum processing data flow according to the invention;

Fig. 17 is a block schematic diagram of a checksum processing data flow for no block aligned ops according to the invention;

20

Fig. 18 is a block schematic diagram of a checksum processing data flow showing the fixing of edge effects according to the invention; and

Fig. 19 is a system diagram according to the invention.

25

DETAILED DESCRIPTION OF THE INVENTION

The presently preferred embodiment of the invention comprises a network device that resides in a network and that secures data residing in network storage. A typical network storage architecture as is known in the art is shown
5 in Fig. 1. The servers 11 shown on the left access the storage devices 12 on the right through a communication fabric 13, *i.e.* a network.

There are two main types of storage architectures to consider in connection
10 with the invention disclosed herein:

- In a block-based architecture, the communication between a server and a storage device occurs in terms of sectors or blocks. For example, the server might request "block number 17" from a storage device. In this
15 case, the storage device is usually not aware of the file system used to organize the data on the device. Examples include SCSI over Fibre Channel (usually referred to as SAN), SCSI over IP (iSCSI).
- In a Network Attached Storage (NAS) architecture, the access is in terms of files. For example, a server might request to read a specific file from a
20 specific directory. In this case, the storage device is usually responsible for creating, maintaining, and updating the file system. Examples of NAS include NFS and CIFS.

Increasingly, data storage is outsourced. Even when it is not, it is often undesirable to let an information technology group that maintains the storage
25 devices have unrestricted access to all of the data stored on these devices. A

key goal of the herein disclosed invention is to separate the ability to access storage (for purposes such as maintenance, backup, sizing, partitioning, etc) from the ability to access the stored data. This is achieved by keeping all (or the sensitive part of) the data encrypted on the storage devices.

5

Possible integration of a presently preferred embodiment of the herein disclosed invention in the network is shown in Fig. 2. In this (inline) configuration, all traffic between the servers 11 and the storage devices 12 is routed through a network element 20 that comprises the herein disclosed invention. The invention analyzes the traffic, identifies the payload (data),
10 encrypts/decrypts it (depending on the direction of traffic flow), and forwards the modified traffic.

An alternative (switched) architecture is shown in Fig. 3. In Fig. 3, the invention virtualizes some of the storage devices, and the clients 11 can
15 access the devices 12 either directly or through the herein disclosed device 20. Again, all (or part of) the data written to the storage devices through the invention are encrypted.

20 Logically, the invention comprises a device that has two network interfaces: one is a clear text network interface that connects to one or more clients, and the other is a secure network interface that is connected to one or more persistent storage servers. Functionally, each network interface supports multiple network nodes. That is, the clear text network interface supports
25 multiple client machines, and the secure network interface supports one or more storage servers.

Both in the context of NAS and in the context of SAN, the invention supports active failover and clustering. In particular, in NAS environment, the failover is completely transparent to the client computers and to the storage servers.

- 5 In cases where multiple invention devices are deployed, all can be controlled globally, from a single point.

The invention supports initial encryption and re-encryption with a different key on-the-fly, without taking the system off-line.

10

Since the invention decodes all the storage protocols and separates control from payload, it can be set to translate between different storage protocols. For example, it can be setup to mount an NFS share and present it to a client computer as a CIFS share. Another example of translation is between iSCSI
15 and Fibre Channel: the invention can connect to a Fibre Channel array and present it as an iSCSI array.

The invention can be configured to enforce traffic-shaping policies on specific subsets of storage traffic. In particular, it can give higher priority to storage
20 traffic between certain client hosts and storage devices or, alternatively, it can limit such storage traffic to not more than a given percentage of the available storage bandwidth. In NAS environment, traffic shaping policies can also take into account specific user names associated with the traffic streams.

- 25 The invention seamlessly integrates with existing networking infrastructure. It is possible to attach the invention to a network (NAS or SAN), initialize it, and

start working. There is no need to reconfigure the servers. In the “switched” deployment case, the clients might need to be reconfigured to point them to access the storage through the invention instead of accessing it directly.

The invention supports industry standard network management interfaces such as SNMP, Web, RMON. I also supports standard physical interfaces
5 such as Gigabit Ethernet, RJ-45 and fibre; Fibre Channel copper and fibre.

NAS deployment and configurationsThe following discussion describes several configurations for NAS (NFS and CIFS) deployments.

10

Type-1 configuration

This is the simplest configuration where the invention is connected to the enterprise network at an arbitrary point. The configuration is shown in Fig. 4.
15 Data from clients 11a, 11b flows through the enterprise network 13 to the inventive device 20 (shown on Fig. 4 by the numeric designator (1)), is encrypted by the device, and sent further to the disks 12a, 12b (shown on Fig. 4 by the numeric designator (2)). A read operation proceeds in the reverse direction. Observe that in this case no assumptions are made as to
20 where the inventive device is placed in terms of the enterprise network. In particular, it is not assumed that the device is on the same subnet as NAS clients or servers.

To access a file through the inventive device, the client has to have one of the allowed IP/MAC address combinations (the list of allowed combinations is
25 determined by the administrator). In addition, if desired, the client can first be

required to authenticate to the inventive device. Authentication possibilities include user_name/password, NTLM, NTLMv2, Kerberos V5, etc.

Type-2 configuration

5

In this configuration (see Fig. 5), the clients 11a, 11b are connected to the clear text port 51 of the inventive device 20. The other port 52 is connected to the rest of the enterprise network 13. All of the traffic to and from the enterprise network (both storage related and other) passes through the
10 inventive device. The packet filtering setup of the inventive device prevents a rogue client connected to the enterprise network, *e.g.* Net 79 in the example, to pretend to be connected to the client network (Net 78 in the example). Because the device does not serve decrypted data on the storage port, the rogue client does not have access to clear text data. Fig. 5 shows only a
15 single subnet on the client side. It is straightforward to extend this setup to allow more than a single subnet on the client side.

Type-3 deployment

This deployment is similar to the Type-1 deployment. The main difference is
20 that the clients 11a, 11b are on the same subnet as the inventive device 20 (Net 78 in the example in Fig. 6). The router 60 is assumed setup in a way that prevents a rogue client from pretending to be on Net 78 when it is not physically connected to this subnet.

More precisely, in this example, the two client subnets (Net78 having clients
25 11a, 11b and Net 90 having clients 41a, 41b) are connected through a switch 61 that supports VLANs (802.3p protocol). Net-78 is mapped to VLAN-10 and

Net-90 to VLAN-20. The trunk carrying both VLANs connects the switch to the router.

Consider a client that sends a packet, *e.g.* read request, to the client port of the inventive device. Packet filtering in the inventive device drops all packets
5 with source addresses other than Net 78. If a rogue client spoofs the source IP address, the reply packet does not reach him unless he is physically connected to Net 78 or somehow succeeds in sitting on VLAN 20. Thus, as long as the switch and the router are secured, the reply never reaches the rogue client.

10

Remote file sharing

A basic setup for remote file sharing is shown in Fig. 7. The example shows one main location 70 ("Headquarters") and one remote location 71. A similar setup is possible for multiple remote locations. The clients 11c, 11d at the
15 remote location see the inventive device 20a as a file server. They access the files presented by the inventive device in the same way as they accesses any other NFS or CIFS server.

When a client in the remote location is trying to read a file from the Headquarters server, the following sequence of events occurs:

- 20 1. First, the client is required to authenticate to the (remote) inventive device 20a. The authentication is performed in exactly the same way as when the NAS servers are directly accessed by the inventive device.
2. If the authentication is successful, the remote device packages the request and forwards it through an IPSEC tunnel to the headquarters device 20b. In
25 this example, it is assumed that the devices 20a and 20b have already authenticated and have created an IPSEC tunnel between them.

3. Upon receiving the request, the headquarters device 20b checks whether the remote device has the authority to execute the request.
4. If the authority exists, headquarters device 20b reads the data from the NAS server, packages it appropriately, and forwards it through the IPSEC tunnel (shown by the numeric designator (1) on Fig. 7). Note that it does not decrypt the data.

Other operations such as write, create file, are executed in a similar way. When headquarters clients try to access data from the remote location, the roles of headquarters/remote are reversed.

10

From the administrative perspective, the permissions are set in the following way:

1. First, Cryptainers that are accessible by the specific remote device are identified.
 2. Then, the administrators of the remote devices define users that can access the Cryptainers.
- If in the above example the headquarters does not have any clients, the associated inventive device does not need to have any Cryptainer keys. In this case it serves only as an access controller. This setup is shown in Fig. 8, in which the inventive device 20 operates in connection with access controllers 80.

Secure remote mirroring

The invention further comprises a device for NAS that supports private and shared remote secure mirroring. The architecture is described in Fig. 12, which shows two companies connecting through a network to storage. Company A has two clients 11a and 11b, while Company B has a single client 11c. Company A has two inventive devices 20a and 20b through which it connects to the shared storage, while Company B has a single device 20c. Inventive device 20d serves as access controller in this setup, serving data only to devices 20a, 20b, and 20c; it does not decrypt the data. Remote mirror storage device 120 is connected to a special port on the inventive device 20c and serves as a private (non shared) mirror of Company B. Storage device 121 is connected to a similar port of inventive device 20d and can be setup to serve as a shared mirror for both companies. Because the data stored on this mirror are encrypted with keys available only to the inventive devices residing in the corresponding companies, one company can not read other company's data even though they share the mirror.

SAN Deployment and Configurations

Direct-connect configuration

Single-switch solution

A single-switch solution is shown in Fig 10. Assuming the switch 90 (or hub) is not zoned, the client sees both raw (original) disks and disks through the inventive device.

Multiple-switch solution

25

In this configuration (Fig. 11), the servers do not see disks directly. The only way they can access the disk is through the inventive device. For all practical purposes, the same configuration can be achieved by zoning a single switch, instead of using two switches 90a, 90b. In this case, ports in one zone are used to connect the servers and the server side port of inventive device, while the other zone is used to connect the disk arrays and the storage side port of the inventive device.

Secure remote mirroring

In all of the SAN configurations above, the data on disk arrays are encrypted. Any mirroring solution can now be used to mirror the data to an offsite location. Because the data are encrypted, security requirements on the offsite location are reduced. An alternative is to use a special version of the inventive device that supports mirroring directly, as it is shown in Fig. 12.

Fig. 12 shows two companies connected to a shared storage site through a fibre channel network. Both companies use the shared mirror facility 121. Because they do not share the keys, neither company can access the data of the other one. In addition, company B is shown to have a private mirror 120 as well.

Hardware configuration

A block schematic diagram of the inventive device 20 is shown in Fig. 13. Storage traffic packets are disassembled by a disassembly modules 130, 135 the payload is encrypted or decrypted by respective modules 131, 134 based on the command, *i.e.* read/write, and in view of key management module 136, the result

is assembled back into a legal packet by an assembly module 132, 133 and sent forward.

A hardware configuration of the inventive device is shown in Fig. 14. The presently preferred embodiment of the invention is a device that is based on a standard dual-CPU 142, 143 architecture with dual PCI buses 140, 141. All of the encryption/decryption operations are done inside the Packet Processor, DPP 144. The only two places where secret data is stored in the system is inside DPP and inside the system Smart Card 145.

10 *Cryptainers*

The ACL management and permissioning is done in terms of “cryptainers”. In the context of SAN, cryptainer can be, for example, an actual disk, a region on a disk, or several regions on one or more disks. In context of NAS, cryptainer is a collection of files and directories. The inventive device offers users and administrators a mechanism for creating secure Cryptainers. The following discussion describes details of a presently preferred implementation and method for creating Cryptainers in the context of NFS and CIFS.

20 *NFS Cryptainers*

Introduction

In general, from the user perspective, a Cryptainer behaves as a physical device in a sense that it is forbidden to move files (mv) from one device to another and hard-link files across devices. By “move is forbidden”, we mean

that the move will involve actual copy of the data followed by deletion of the original.

Assumptions regarding setup

5

In the following discussion, several assumptions are made:

1. The inventive NFS device is installed with NFSBOX as the DNS name for the client-side interface.

10

2. The clients are able to reach this interface through the networking infrastructure.

3. The NFS server(s) are reachable from the server interface.

15

4. A management workstation is connected to the client side of the device and the appropriate software is installed.

20

5. Users, clients, and servers were defined by the administrator through a management interface.

6. Servers are allowed to export their shares to the server-side interface of the inventive device (NFSBOX).

25

7. There are two servers: Srv1 exports Srv1:/srv1/home/a, Srv1:/srv1/home/b; Srv2 exports Srv2:/export1.

8. The inventive device is set to export all of the above three shares.

9. Every share exported by the inventive device has a name defined by the
5 administrator.

During the setup, the user is encouraged to use canonical names for the
exported shares:

10 NFSBOX:/Srv1/home/a,
NFSBOX:/Srv1/home/b,
NFSBOX:/Srv2/export1.

The administrator is able to use arbitrary names. To make the description as
15 general as possible, the use of non-canonical names is assumed, even
though this may not be a preferred practice.

In particular, in the following, it is assumed that the names are:

20 NFSBOX:/Srv1/home/a,
NFSBOX:/home/b,
NFSBOX:/export.

Note that, in this example, the names of the exported shares are independent
25 of their original names, which may not be a good IT practice.

It is assumed that the client has mounted the shares exported by device on /decru/a, /decru/b, /decru/c.

Working with Clear text Data

5

The mounts /decru/a,b,c behave exactly as regular NFS mounts. The user can create directories, create files, and move files from one directory to another as long as both are under the same mount point.

10 Creating Cryptainers

A Cryptainer can be created by the user through CLI or a user-GUI. The process of the creation is described below.

15 There are two main choices:

- An existing directory is converted to become a Cryptainer. For example, the user has a directory in /decru/a/examples/dir1. After the conversion, the dir1 directory becomes the new root of a Cryptainer. All existing data in
20 dir1 is encrypted.

- A new Cryptainer is created in an existing directory. For example, assuming /decru/a/examples exists, the user can create /decru/a/examples/dir2. A new directory, dir2, is created and ready for use
25 as a Cryptainer. Any new file created in dir2 or copied there is immediately

encrypted. During creation of a new Cryptainer, the user-GUI suggests to the user that a distinctive name be given to the Cryptainer.

Comments

5

- Although it is possible to create nested cryptainers, in the following discussion we assume that nesting is not allowed.

10

- Initially, a Cryptainer is constructed with access rights given only to the user creating it – “owner”. It is possible to pass ownership to another user.

- The process of changing the access rights and further securing the Cryptainer with a password or a smart card is described below.

15

- During creation of the Cryptainer the user can choose whether the encryption includes file names and directories.

- In the preferred embodiment, the Cryptainer name is always clear text.

20 Working with Cryptainers

The user sees a Cryptainer as a regular directory. For performance reasons, there are several restrictions on operations involving Cryptainers. In general, the restrictions are the same as if a Cryptainer is a separate mount point of a

25 share:

- It is impossible to “mv” Cryptainers from one share to another.
- It is impossible to “mv” a file or a directory structure across a Cryptainer boundary.

5

- It is impossible to hard-link across Cryptainer boundary.

Implementation

10 General assumptions and requirements:

1. Unix users are assumed to use NFS file servers.
2. Restrictions on number of Cryptainers, users, shares, are discussed
15 below.
3. The user can apply all the regular file manipulation operations to
directories and files within a Cryptainer. These operations include: mv, cp,
rm, rmdir.
20
4. In UNIX there is no inheritance of ACLs; each directory/file has its own
ACL settings. Accessing a file or a directory requires access rights to all
the directories above it. Same for Cryptainers.
- 25 5. The solution supports common backup utilities and configurations.

6. Both initial encryption of a non-empty directory and re-encryption of an existing Cryptainer is supported.
 7. Power failure of the device does not result in loss of data even during re-
5 encryption.
 8. Users should be able to recover backup and recover Cryptainers.
 9. A Cryptainer can be recovered to a destination different from the original
10 destination.
- Re-encryption - it is not always the case that the initiating user has the appropriate permissions to read and write all the files. The system checks that the user is the owner of all files in a Cryptainer to initiate re-
15 encryption. Same restriction for initial creation of a Cryptainer from non-empty directory.
 - Re-encryption is done in-place.
 - 20 • Byte range locking during re-encryption – while a specific file is being re-encrypted, accesses by clients to this file are allowed.

NFS Implementation

25 Framework

Cryptainer keys (CK) are used to encrypt data, filenames and directory names within a given Cryptainer.

The first block of each file contains:

5

- Cryptainer Handle (CryptainerID) in plain-text
- Two random values R1, R2 encrypted with the CK

10 • Additional administrative information

Creation of a new Cryptainer results in creation of a directory with the Cryptainer name and a .decru file in this directory with meta-data about the Cryptainer. For example, creating a new Cryptainer named
15 /decru/a/examples/dir2 results in creation of a directory named
NFSBOX:/Srv1/home/a/examples/dir2.

Conversion of an existing directory into a Cryptainer leaves the files where they are on the server in terms of the directory path, encrypts them, and
20 creates the .decru file under the root name. The files are encrypted in-place and hence their NFS handles are not changed in this process.

To improve performance, one might add .decru files with appropriate metadata in any one of the directories.

25 The client can neither read nor delete the decru files.

Use Cases

Adding shares to the device export list

- 5 This step must be completed before Cryptainers can be created
1. The administrator provides the server name
 2. The Device Manager adds the server to the list of virtualized servers
 3. The Device Manager displays a list of shares for the server
 - 10 4. Administrator selects shares to be exported through the inventive device
 5. The Device Manager notifies the NFS Mount proxy of the new share information

Creating a Cryptainer

- 15
1. Using the management tool (either Web or CLI-based), the user selects an existing share
 2. User enters a path and the name of the new Cryptainer
 3. The Device Manager creates a new CryptainerID for the new Cryptainer
 - 20 4. The Device Manager gives the NFS Proxy (NP) the share, path, Cryptainer name and the CryptainerID and asks the NP to create the new Cryptainer directory
 5. The NP creates the Cryptainer directory (un-encrypted).
 6. The NP informs the Device Manager upon completion and asks the DPP
 - 25 for a new Cryptainer Key. (It is returned encrypted)
 7. The CryptainerID is stored in the .decru file

Assigning Permissions on a Cryptainer

1. The owner of the Cryptainer using the admin tools asks the Device
5 Manager to edit permissions on the Cryptainer.
2. The updated information is stored in the Configuration DB (CDB).

Client Procedures

- 10 Client procedures can be grouped into three categories for the purpose of determining a CryptainerID:

1. Procedures that take a file handle to an existing object:

15	GETATTR	Get file attributes
	SETATTR	Set file attributes
	ACCESS	Check access permissions
	READLINK	Read from symbolic link
	READ	Read from file
20	WRITE	Write to file
	COMMIT	Commit cached data

2. Procedures that take a file handle to a parent directory and a name:

25	LOOKUP	Lookup filename
	CREATE	Create file

	MKDIR	Create Directory
	SYMLINK	Create symbolic link
	MKNOD	Create a special device
	REMOVE	Remove File
5	RMDIR	Remove Directory
	REaddir	Read from Directory
	REaddirPLUS	

3. Procedures that take multiple file handles:

10

	RENAME	Rename File or Directory
	LINK	Create Link to an Object

15 The preferred procedure for determining the CryptainerID given a file handle to an existing object is:

1. The user provides a file handle and data
2. The NP looks in the file handle cache to determine the CryptainerID
3. If the file handle is not in the cache, the NP reads the file meta-data (first
20 block of the file) and gets the CryptainerID
4. The NP passes the user's credentials + CryptainerID to the Device Manager
5. The Device Manager checks the ACL and sends back the Cryptainer Key (CK) and a permission bit mask
- 25 6. The CK, the random values, and the data are sent to the DPP for encryption if the operation requires encryption or decryption

The procedure for determining the CryptainerID for an object given a parent directory and name is:

- 5 1. The client provides a file handle to parent directory and a filename
2. The NP looks in the file handle cache to determine the CryptainerID
3. If not found, the NP looks for the .decru file in parent directory (handle to
10 this directory was given) and extracts the CryptainerID if file found. If the
.decru file is missing the NP walks up the directory tree until it reaches the
root directory, each time looking for .decru file and checking the directory
handle against cache. If neither .decru found nor cache hit, the original file
is assumed to be in clear text. Otherwise the CryptainerID is extracted
either from the cache or from the .decru file.
- 15 4. To improve performance, the system can maintain a cache data structure
that holds all of the paths to Cryptainers and all of the file handles for the
directories on these paths.

Creating a new file

- 20 1. CryptainerID of the directory is obtained using the algorithm above
2. The NP passes the user's credentials + CryptainerID to the Device
Manager
3. The Device Manager checks the ACL and sends back the Cryptainer Key
(CK, encrypted with Master Key) and a permission bit mask
- 25 4. The NP asks the DPP for 2 random values (file keys)

5. The NP sends the file name to DPP to be encrypted using the CK and creates the file
6. The NP writes the first record that contains Cryptainer ID and the encrypted R1/R2 values and some additional administrative information.
- 5 7. A handle to the file is returned to the user and saved in temporary cache

The procedure for RENAME - Moving a File (mv) or directory

1. The user provides a file handle to the source directory and source name
10 and a file handle to the destination directory and a destination name
2. The NP uses the procedure above to find CryptainerIDs for both source and the destination
3. The CryptainerID of the destination directory is compared to the CryptainerID of the source directory
- 15 4. If same CryptainerID and if NFS allows a move, then a regular move is done
5. If NFS doesn't allow the move, the move is not allowed
6. If Cryptainer handle is different or if a move is attempted clear->Cryptainer or Cryptainer->clear the operation fails

20

The procedure for LINK – Create hard link

1. CryptainerIDs of source and destination are found.
2. Link created only if both are clear text or both are in the same Cryptainer.

25

CIFS Cryptainers

The following discussion describes a preferred implementation method for the CIFS Proxy (CP) behavior. To make the description more specific, several
5 assumptions are made:

- CIFS device is installed with CIFSBOX as the DNS name for the client-side interface. (This name can be arbitrary legal DNS name, CIFSBOX is used just as an example)
 - 10 • The clients are able to reach this interface through the networking infrastructure.
 - The CIFS server(s) are reachable from the server interface.
 - A management workstation is connected to the client side of the device and the appropriate software is installed.
 - 15 • Users, clients, and servers were defined by the administrator through a management interface.
 - Servers are allowed to export their shares to the server-side interface of the inventive device.
 - There are two servers: Srv1 with shares \\Srv1\share-a, \\Srv1\share-b;
20 Srv2 with shares \\Srv2\share.
 - The inventive device is set to export all of the above three shares.
 - Every share exported by the inventive device has a name defined by the administrator.
- 25 During the setup, the user is encouraged to use canonical names for the exported shares:

\\CIFSBOX\Srv1_share-a

\\CIFSBOX\Srv1_share-b

\\CIFSBOX\Srv2_share.

5

The administrator can use arbitrary names. To make the description as general as possible, the use of non-canonical names is assumed, even though this may not be a good practice.

10 In particular, in the following, it is assumed that the names are:

\\CIFSBOX\a

\\CIFSBOX\b

\\CIFSBOX\c

15

Note that in this example the names of the exported shares are independent of their original names, which is may not be a good IT practice.

It is assumed that the client has mounted the shares exported by device on

20 F:, G:, H:.

Working with Clear text Data

The mounts \\CIFSBOX\a,b,c behave exactly as regular CIFS mounts. The
25 user can create directories, create files, and move files from one directory to another as long as both are under the same mount point. Moves from one

share to another share on device are converted by the client into copies. In particular, this is true when using drag-and-drop in Windows Explorer. The command line “move” is converted into a copy+delete.

5 Creating Cryptainers

A Cryptainer can be created by the user through CLI or a user-GUI. The process of the creation is described above. There are two main choices:

- 10 1. An existing directory is converted to become a Cryptainer. For example, the user might have a directory \\CIFSBOX\dir1. After the conversion, the dir1 directory becomes the new root of a Cryptainer. All existing data in dir1 are encrypted.
- 15 2. A new Cryptainer is created in an existing directory. For example, assuming \\CIFSBOX\examples exists, the user can create \\CIFSBOX\examples\dir2. A new directory, dir2, is created and ready for use as a Cryptainer. Any new file created in dir2 or copied there is immediately encrypted. (The file name is encrypted if user sets “encrypt
20 file names” property during creation of the Cryptainer.) During creation of a new Cryptainer, Decru user-GUI suggests to the user that a distinctive name be given to the Cryptainer.

Comments

- Although it is possible to support nested Cryptainers, the assumption in the following discussion is that they are not supported.
- Initially, a Cryptainer is constructed with access rights given only to the user creating it, the “owner”. Ownership can be passed to a different user using CLI or user-Gui.
- The process of changing the access rights and further securing the Cryptainer with a password or a smart card is described in subsequent sections.
- During creation of the Cryptainer, the user can choose whether the encryption includes file and/or directory names.
- The Cryptainer name is always clear text.

Working with Cryptainers

The user sees a Cryptainer as a regular directory. There are several restrictions on operations involving Cryptainers.

- It is impossible to move (as opposed to “copy”) Cryptainers from one share to another.
- It is impossible to “mv” a file or a directory structure across a Cryptainer boundary.

Implementation

General assumptions and requirements

5

- Windows users are assumed to use CIFS file servers.
- Restrictions on number of Cryptainers, users, shares, are discussed elsewhere.
- The user is able to apply all the regular file manipulation operations to
10 directories and files within a Cryptainer.

- The solution supports common backup utilities and configurations.
- Both initial encryption of a non-empty directory and re-encryption of an existing Cryptainer is supported.
- Power failure of the device doesnot result in loss of data even during re-
15 encryption

- Users are able to recover backup and recover Cryptainers
- A Cryptainer can be recovered to a destination different from the original destination.
- When creating a Cryptainer, user can decide on whether to encrypt file
20 names and/or directory names.

- Re-encryption - it is not always the case that the initiating user has the appropriate permissions to read and write all the files. One must check that the user is the owner of all files in a Cryptainer to initiate re-
25 encryption. Same restriction for initial creation of a Cryptainer from non-empty directory.

- Re-encryption is done in-place.
- Byte range locking during re-encryption, same as for NFS described earlier.

5

CIFS Implementation

Framework

- 10
- Cryptainer keys (CK) are used to encrypt data, filenames and directory names within a given Cryptainer.
 - The first block of each file contains:
 - Cryptainer Handle (CryptainerID) in plain-text
 - Two random values R1, R2 encrypted with the CK

15

Some other information

- Creation of a new Cryptainer results in creation of a directory with the Cryptainer name and a .decru file in this directory with some meta-data
20 about the Cryptainer, including its handle/ID.
- Conversion of an existing directory into a Cryptainer leaves the files where they are on the server (in terms of the directory path), encrypts them, and creates the .decru file right under the root name.
- To improve performance, one might add .decru files with appropriate
25 metadata in any one or all of the directories.
- The client can neither read nor delete the .decru files.

Use Cases

Adding a shares to the export list:

5

This step must be completed before Cryptainers can be created:

1. The administrator provides the server name
2. The Device Manager adds the server to the list of virtualized servers
- 10 3. The Device Manager displays a list of shares for the server
4. Administrator selects shares to be exported through the inventive device
5. The Device Manager notifies the NFS Mount proxy of the new share information

15 Creating a Cryptainer

1. Using the management tool, the user selects an existing share
2. User enters a path and the name of the new Cryptainer
3. The Device Manager creates a new CryptainerID for the new Cryptainer
- 20 4. The Device Manager gives the CP the share, path, Cryptainer name and the CryptainerID and asks the CP to create the new Cryptainer directory
5. The CP creates the Cryptainer directory (un-encrypted).
6. The CP informs the Device Manager upon completion and asks the DPP for a new Cryptainer Key. (It is returned encrypted)
- 25 7. The CryptainerID is stored in the .decru file

Assigning Permissions on a Cryptainer

1. The owner of the Cryptainer using the admin tools asks the Device Manager to edit permissions of the Cryptainer
- 5 2. The updated information is stored in the Configuration DB (CDB).

Client Procedures

In CIFS, a user must first connect to a specific \\server\share via a
10 TREE_CONNECT. The server reply to a TREE_CONNECT includes a tID that
is used in all future commands sent to the server. All commands after the tree
connect reference objects on the server in one of three possible ways:

1. Using the tID returned by TREE_CONNECT and an absolute path relative
15 to this tree
2. Using an fid returned by call using a tID and path
3. Using a source tID, path destination tID, path

There is one exception to this rule. NT_CREATE_ANDX has an optional
20 parameter (RootDirectoryFid) that can be used in combination with a relative
path.

The procedure for determining a CryptainerID given a tID and a path is:

- 25 1. The user provides a tID and a path

2. Starting at the top of the path the proxy looks up the path in the CryptainerID cache.
3. If the path is not found the proxy starts the CryptainerID recovery procedure. An example of a client side path of \a\b\c is used where b is the
5 Cryptainer.
4. Proxy attempts to read \a\.decru from the server. If the file is found the Cryptainer handle is read from .decru and the path\CryptainerID are put in the CryptainerID cache.
5. If the .decru file is not found under \a, the proxy attempts to read
10 \a\b\.decru.
6. The CP passes the user's credentials + CryptainerID to the Device Manager
7. The Device Manager checks the ACL and sends back the Cryptainer Key (CK) and a permission bit mask
- 15 8. The proxy calls the DPP to encrypt the partial path "c" and creates \a\b\Enc[c]

After the CK has been determined, it is necessary to handle creation of files separately from the other procedures.

20

Creating files

1. The CP asks the DPP for two random values (file keys)
2. The CP encrypts the file name using the CK and creates the file
- 25 3. The CP write the first record that contains
4. CryptainerID

5. Two random values encrypted with the CK (returned by DPP)
6. Some additional administrative information.

The FileID in the server response is returned to the user and is saved in
5 temporary cache.

Other procedures

The command is passed through with the file/directory name modified to
10 \a\b\Enc[c]

References via Tid and Path

CREATE_DIRECTORY: Create Directory

15

The create directory message is sent to create a new directory. The appropriate TID and additional pathname are passed. The directory must not exist for it to be created.

20 **DELETE_DIRECTORY:** Delete Directory

The delete directory message is sent to delete an empty directory. The appropriate TID and additional pathname are passed. The directory must be empty for it to be deleted.

25

CHECK_DIRECTORY: Check Directory

This SMB protocol is used to verify that a path exists and is a directory. No error is returned if the given path exists and the client has read access to it.

5 **NT_CREATE_ANDX:** Create File

This command is used to create or open a file or a directory.

OPEN, OPEN_ANDX, TRANS2_OPEN2: Open File

10

This message is sent to obtain a file handle for a data file. This returned FID is used in subsequent client requests such as read, write, close,

CREATE: Create File

15

This message is sent to create a new data file or truncate an existing data file to length zero, and open the file. The handle returned can be used in subsequent read, write, lock, unlock and close messages.

20 **DELETE:** Delete File

The delete file message is sent to delete a data file. The appropriate TID and additional pathname are passed. Read only files may not be deleted, the read-only attribute must be reset prior to file deletion.

25

QUERY_INFORMATION: Get File Attributes

This request is sent to obtain information about a file.

SET_INFORMATION: Set File Attributes

5

This message is sent to change the information about a file.

CREATE_TEMPORARY: Create Temporary File

- 10 The server creates a data file in DIRECTORY relative to TID in the SMB header and assigns a unique name to it.

CREATE_NEW: Create File

- 15 This message is sent to create a new data file or truncate an existing data file to length zero, and open the file.

TRANS2_QUERY_PATH_INFORMATION

- 20 This request is used to get information about a specific file or subdirectory.

TRANS2_SET_PATH_INFORMATION

This request is used to set information about a specific file or subdirectory

25

TRANS2_CREATE_DIRECTORY

This requests the server to create a directory relative to TID in the SMB header, optionally assigning extended attributes to it.

5 *References via fid*

CLOSE: Close File

The close message is sent to invalidate a file handle for the requesting
10 process. All locks or other resources held by the requesting process on the file should be released by the server. The requesting process can no longer use FID for further file access requests.

FLUSH: Flush File

15

The flush command is sent to ensure all data and allocation information for the corresponding file has been written to stable storage. When the FID has a value -1 (hex FFFF) the server performs a flush for all file handles associated with the client and PID. The response is not sent until the writes are complete.

20

READ, READ_ANDX: Read File

The read message is sent to read bytes of a resource indicated by FID in the SMB protocol header.

25

WRITE, WRITE_ANDX, WRITE_AND_CLOSE: Write Bytes

The write message is sent to write bytes into the resource indicated by FID in the SMB protocol header.

5 **LOCK_BYTE_RANGE:** Lock Bytes

The lock record message is sent to lock the given byte range. More than one non-overlapping byte range may be locked in a given file. Locks prevent attempts to lock, read or write the locked portion of the file by other clients or
10 PIDs. Overlapping locks are not allowed. Offsets beyond the current end of file may be locked. Such locks do not cause allocation of file space.

UNLOCK_BYTE_RANGE: Unlock Bytes

15 This message is sent to unlock the given byte range. OFFSET, COUNT, and PID must be identical to that specified in a prior successful lock. If an unlock references an address range that is not locked, no error is generated.

SEEK: Seek in File

20

The seek message is sent to set the current file pointer for FID.

LOCK_AND_READ: Lock and Read Bytes

25 This request is used to lock and read ahead the specified bytes of the file indicated by FID in the SMB header

READ_RAW: Read Raw

The SMB_COM_READ_RAW protocol is used to maximize the performance
5 of reading a large block of data from the server to the client. This request can
be applied to files and named pipes.

WRITE_RAW: Write Raw Bytes

10 The Write Block Raw protocol is used to maximize the performance of writing
a large block of data from the client to the server. The Write Block Raw
command's scope includes files, Named Pipes, and spooled output (can be
used in place COM_WRITE_PRINT_FILE).

15 SET_INFORMATION2: Set File Information

SMB_COM_SET_INFORMATION2 sets information about the file
represented by FID. The target file is updated from the values specified. A
date or time value or zero indicates to leave that specific date and time
20 unchanged.

QUERY_INFORMATION2: Get File Information

This SMB gets information about the file represented by FID.
25

LOCKING_ANDX: Lock or UnLock Bytes

SMB_COM_LOCKING_ANDX allows both locking and/or unlocking of file range(s).

5 TRANS2_QUERY_FILE_INFORMATION

This request is used to get information about a specific file or subdirectory given a handle to it.

10 *References via [Tid1, Tid2] and [OLDFILENAME, NEWFILENAME]*

MOVE: Rename File

The source file is copied to the destination and the source is subsequently
15 deleted. OLDFILENAME is copied to NEWFILENAME, then OLDFILENAME is deleted. Both OLDFILENAME and NEWFILENAME must refer to paths on the same server. NEWFILENAME can refer to either a file or a directory. All file components except the last must exist; directories will not be created.

20 NEWFILENAME can be required to be a file or a directory by the Flags field. The TID in the header is associated with the source while TID2 is associated with the destination. These fields may contain the same or differing valid values. TID2 can be set to -1 indicating that this is to be the same TID as in the SMB header. This allows use of the move protocol with
25 SMB_TREE_CONNECT_ANDX.

COPY: Copy File

The file at SOURCENAME is copied to TARGETFILENAME, both of which must refer to paths on the same server.

5

The TID in the header is associated with the source while TID2 is associated with the destination. These fields may contain the same or differing valid values. TID2 can be set to -1 indicating that this is to be the same TID as in the SMB header. This allows use of the move protocol with

10 SMB_TREE_CONNECT_ANDX.

*References via Tid***SMB_QUERY_INFORMATION_DISK: Get Disk Attributes**

15

This command is used to determine the capacity and remaining free space on the drive hosting the directory structure indicated by TID in the SMB header.

Crypto Framework

20

Encryption of FilesOverview

Each file in the NFS/CIFS environment and each block in the FC environment is encrypted with a “strong” length key of 128, 192, or 256 bits. The files are arranged into Cryptainers. The Cryptainers represent encryption granularity.

- 5 In the NAS product, a Cryptainer can be a file, a set of files, a directory, a volume, or a set of volumes. In the SAN product a Cryptainer is identified by initiator id, target id, lun, and LBA range (block addresses). It is also possible to address Cryptainers by WWN. It is also possible to create a cryptainer that corresponds to several regions on one or more disks.

10

Properties of the encryption

Strength: Encryption is as strong as or stronger than ECB using 3DES or AES.

15

Block comparison: It is impossible to check whether two different blocks have the same clear text data without decrypting them.

- 20 *Bit flip:* It is cryptographically difficult to change ciphertext in a way that result in a flip of a specific bit in the output.

Additional properties

25

- Files can be copied or renamed without re-encryption, as long as they are not moved from one Cryptainer to another.
 - Support symbolic and hard links in NFS.
 - The encryption scheme is very similar for all the projects (CIFS, NFS, FC).
- 5 • Increase in the total number of head movements of a disk due to the encryption are minimized.

Preferred Solution

10

Preferred framework

- One key is assigned to each Cryptainer – CK
 - Each data block B is associated with two BlockUniqueIDs – ID1, ID2.
- 15 • Each file in NFS/CIFS and each Cryptainer in FC is associated with two randomly created values: R1 and R2.
- ID1 and ID2 are created by XORing BlockID with R1 and R2, respectively.
- A different function can be used instead of XOR.

20 Each data block B (64b for 3DES; 128b for AES) is encrypted as follows:

$$B' = \text{ENC}_K[B \text{ xor } ID1] \text{ xor } ID2$$

ENC is the chosen encryption function

25

Decryption is done by:

$$B = \text{DEC}_k[B' \text{ xor ID2 }] \text{ xor ID1}$$

DEC is the reverse function of ENC

5

Comments

R1 and R2 are key material and are stored encrypted with CK

10

Non-Block Aligned Disk Access

Overview

15

Because the crypto algorithm is block oriented, there is a need to process complete blocks (usually 16 or 8 bytes long) to encrypt/decrypt the text successfully.

20 Issues exist in the following scenarios:

- A read or write request starts at byte sb where $[sb \bmod \text{block-size}] \neq 0$
- A read or write request ends at byte eb where $[eb \bmod \text{block-size}] \neq 0$
- A special case exists in the last few bytes of a file if *padding* was done

25 while writing it. In FC there is no padding.

Preferred Solution – General

Preferred framework for NAS

- 5
- All files are padded with 120 random bits
 - The non-aligned bytes at the end of the file are not encrypted (they contain random bits generated by the inventive device)
 - When processing the file we always ignore that last 120 bits

10 Proposed framework for FC

- Always read and write aligned blocks – no padding is needed.

All FC and SCSI block sizes must be a multiple of 512 bytes

15 **Preferred Solution – NAS Read**

- Read requests should be manipulated by the system to be block aligned. The extra bytes should be removed before returning to the user.
- When seeking to end of file (EOF), we ignore the last 120 bits.

20

Preferred Solution – NAS Write

- 25
1. The system reads the first and last blocks that contain the *sb* and *eb* of the write request (assuming that the write request is misaligned at both ends)
 2. The two blocks are decrypted

3. The required plain-text bytes are concatenated at each end of the write block
4. The aligned blocks are encrypted and written to disk
5. If it is the last block of the file we pad it with 120 random bits and encrypt only the aligned portion of it (for example, if the last block of the files contain only three bytes, append 120 random bits to it. If AES is used, then block size is 128 bit. Therefore, encrypt the first 128 bits and write the last 16 bits to disk with no encryption).

10

Encryption Algorithms

Overview

15 The preferred system uses two encryption algorithms: 3DES and AES (Rijndael).

- 3DES uses a key of 168 bits and a block size of 64 bits.
- The AES implementation uses a key of 128 bits and a block size of 128 bits.

20

Key Management

Framework

25

Terminology

- Hardware Token – a smart-card (or a pc-card, or an ibutton)
- ITKN – Internal token, inserted in the inventive device
- MTKN – Management token, required to do normal administration
- 5 • DRTKN – Data Recovery token, required for certain sensitive operations
- PTKN – Personal token, used by end-users to secure their Cryptainer keys
- DPP – Device PCI card that implements the crypto logic
- Management Station – a PC running a Java applet and a web browser used to remotely manage the devices
- 10 • $E_{Key}(\text{Message})$ = Encryption of Message by Key
- $[\text{Message}]_{AuthKey}$ = Private AuthKey signature of Message, verifiable with Public AuthKey

General

- 15 • The presently preferred embodiment comprises level 3 FIPS certification for DPP board and level 2 for the whole device.
- The inventive device contains a mechanism to zeroize clear-text keys on demand.
- At least three Hardware Tokens (smart-card, pc-card, or an ibutton) are
- 20 provided with each device, one of each type.
- The Hardware Token must support the following: PKI, authentication, enough memory to store the information described in section 4 below + some temporary space for the crypto operations.
- Key material never leaves/enters DPP in clear text.

25

- The device has two security levels: initially it uses device generated keys (Initialization Keys). After initialization by the client, all operations are based on locally generated keys (Operation Keys) (there is no back door to the device).

5

- A Cryptainer can be accessible through one or more devices. Each Cryptainer has a *master* device. Access control modifications are managed by the master. The Cryptainer becomes accessible through other devices only with the master's approval.

- 10 • The user can disable remote configuration

Key Generation

Master Key – generated at initialization time by request of user, by DPP.

15

Cryptainer Keys - In the general case, keys are injected into the device.

Special cases of the above are:

- 20 1. Keys are generated inside the device (and “self injected”)
- 2. Keys are generated in one device and transmitted to another (clustering, replication, hot standby,)
- 3. Keys are restored from backup

Initial Authentication Key Set(AKS) - 3des static keys used to create session keys (see Visa Spec), created during device manufacturing. Session keys are used to secure DPP-ITKN and MTKN-ITKN communication.

- 5 *Local AKS* - Replaces initial AKS during device initialization.

Local Key Pair - ITKN and each DRITKN have their own key pairs, generated during initialization.

- 10 *Device Public Key* – generated once and is identical for all devices

Device Authenticity Key Pair – generated at manufacturing time, unique for each ITKN and DRITKN, the public key is signed by device Private Key, also the signature includes the type of key, *i.e.* ITKN, MTKN, or DRITKN.

15

Key Storage at Runtime

- The device Public Key is stored in the Hardware Token or the main board memory. The Master Key (MK) is kept inside the DPP in clear-text. To enable
20 the “Change MK” operation, the DPP stores two MKs: Incoming Data MK and Outgoing Data MK. The Incoming MK is used to decrypt incoming key material; the Outgoing MK is used to encrypt outgoing key material. In normal operation both are the same. When there is a need to change the MK a new key is generated and placed as the Outgoing MK, and all the Cryptainer Keys
25 are inserted using the old key and extracted using the new key. Cryptainer

Keys (CK) are stored in the main board memory encrypted using the MK.
DPP contains its AKS entry.

ITKN Keys:

5

Device Public Key

ITKN Authenticity Key Pair, signed by the device

Local AKS

Initial AKS – discarded after device initialization.

10

ITKN Local Key Pair, used to encrypt MK and sign DR Public Keys.

DRTKN Keys:

Device Public Key

15

DRTKN Authenticity Key Pair, signed by the device

DRTKN Local Key Pair: for secret sharing – sw based MK recovery.

MTKN Keys:

Local MTKN <-> ITKN AKS entry: enables secure channel.

Initial MTKN <-> ITKN AKS entry: enables secure channel.

20

Persistent Key Storage

1. MK is stored in the file system encrypted using the ITKN local Public Key,
and additionally saved via secret sharing. Secret shares are stored
25 encrypted with public keys of combinations of DRTKN cards, allowing
recovery of MK using “m out of n” DRTKNs.

2. CK are persisted into the crypto config db, which is stored in the file system. The data is encrypted using the MK.
3. Before encryption, each key is modified by inserting zeros in several places in the key. After each decryption, existence of these zeroes is
5 verified, providing error-detection capability
4. The “crypto config db” associates each CryptainerID with a key.
5. The entire AKS is stored in the ITKN. The DPP and MTKN each have their respective entries in the AKS. The AKS in the DPP is stored in battery backed RAM or inside the microcontroller residing under physically secure
10 cover on the DPP board. The AKS is backed up to CryptoDB using Secret Sharing, similar to MK.
6. The initial and local key pairs of the ITKN and DRTKN are stored on their respective smart cards.

15 *Authentication*

1. DPP to ITKN – authenticate using the Authentication Key Set
2. Admin User to MTKN/DRTKN – authenticated based on username and password
- 20 3. MTKN to ITKN – authenticated using AKS entry.
4. DRTKN to ITKN – authenticated using device signed public keys.
5. DPP implements identity-based authentication, to comply with FIPS level 3

25 *Key Distribution*

Communication with the HW Token

1. Some functions of the HW Token require authentication and a creation of a secure channel others do not.
- 5 2. The HW Token enforces the defined security policy (it executes functions only if the right security settings exist).

DPP, ITKN, and CPU Communication Channels

- 10 1. The communication channel between the DPP and the ITKN is secured using a session key that is only known to them. MTKN to ITKN session keys are derived from a different static AKS entry.
2. The communication channel between the DPP and the CPU is not secured – cryptographically sensitive information is not passed on this channel in
15 clear-text.
3. The communication channel between the CPU and the Hardware Token is not secured – cryptographically sensitive information is not passed on this channel in clear-text.
- 20 4. DRTKNs and ITKNs authenticate by trading signed public keys. They are capable of using these keys to create a secure channel, but generally do not.

Basic Procedures

25

See Fig. 15 in connection with the following discussion.

1. DPP to Internal HW Token 155, 158 (ITKN) Authentication
2. The CPU 156, 159 asks the DPP to create an Authentication Message (the logic of the packet assembly can be handled by the CPU). The
5 Authentication Message is based on the AKS (or the IAKS during initialization).
3. The CPU sends the Message to the ITKN
4. The ITKN sends its reply and the CPU transfers that to the DPP
- 10 5. The CPU transfers the DPP's response to the ITKN
6. At this point the ITKN and the DPP have a virtually secure channel of communication that the CPU cannot access. They use this channel to exchange key material. The CPU does all the data line management. This authentication is done using an implementation of the authentication
15 mechanism.

Mgmt Station to MTKN/DRTKN Authentication

1. The Mgmt Station 152 in the management PC 150 sends the right
20 user/password to the crypto card 157, 160.
2. This authentication does not result in a secure channel.

Mgmt Station to inventive device Authentication

- 25 1. The Mgmt Station authenticates to the MTKN (see above)
2. The Mgmt Station asks the MTKN to create an Authentication Message

3. The Mgmt Station sends the Message to the ITKN through the CPU
4. The ITKN sends its reply and the CPU transfers it to the Mgmt Station that transfers it to the MTKN
5. The MTKN response is sent back to the ITKN
- 5 6. At this point the ITKN and the MTKN have a secure way to communicate that the CPUs cannot access. This authentication is done using our implementation of the authentication protocol.

DRTKN to ITKN

10

1. DRTKN sends ITKN DRAuth PubKey and $[DRLocal\ PubKey]_{DRAuth}$
2. ITKN verifies DRAuth signed by the device, and DRLocal signed by DRAuth.
3. ITKN sends IAuth PubKey and $[ILocal\ PubKey\ and\ hash(R1)\ and$
 15 $E_{DRLocal}(R1)]_{IAuth}$ to DRTKN. R1 is random number generated each time for this operation.
4. DRTKN verifies IAuth Pub Key with its copy of the device Public Key and decrypts R1.
5. DRTKN sends $hash(R2\ ||\ hash(R1\ ||\ R2))$ and $E_{ILocal}(R2\ ||\ hash(R1\ ||\ R2))$.
 20 R2 is random number generated each time for this operation.
6. DRTKN and ITKN have now traded local Public Keys. They trust them because of a chain Public Key -> Authentication Public Key -> Local Public Key, R1 and R2 prevent replay and establish a 3des session key (R1 || R2).
- 25 7. Because above is not dependent on AKS, DRTKNs can be shared among different devices. (ITKN and MTKN are device specific)

Secret Sharing

1. MM asks ITKN to secret share Data among M DRTKNs, requiring N for
5 recovery.
2. ITKN signs secret with Local Key Pair, the data saved to disk is secret concatenated with the signature.
3. ITKN creates N choose M secret parts (kept on card).
4. MM reads [ID and DRLocal PubKey]_{ILocal} off disk and sends to ITKN.
- 10 5. ITKN validates its own signature on the public key, and encrypts that part of the secret with DRLocal PubKey, also signs secret with own Local Key.
6. MM saves secret to disk. Repeat for all M DRTKN public keys. (Examples in this document use n=3).

15 Recover Secret

1. User-1 authenticates to DRTKN-1
2. DRTKN-1 creates two Recover Secret requests, which are saved locally:
3. DRTKN-1 Auth PubKey and [DRTKN-1 local PubKey and nonce and E_{DR2-}
20 $_{Local}(\text{Secret-2})_{DRAuth}$
4. DRTKN-1 Auth PubKey and [DRTKN-1 local PubKey and nonce and E_{DR3-}
 $_{Local}(\text{Secret-3})_{DRAuth}$
5. User-2 authenticates to DRTKN-2
6. DRTKN-2 decodes Request 1 and returns $E_{DR1 local}(\text{Secret-2 and nonce})$
- 25 7. User-3 authenticates to DRTKN-3
8. DRTKN-3 decodes Request 2 and returns $E_{DR1 local}(\text{Secret-3 and nonce})$

9. User-1 authenticates to DRTKN-1.
10. DRTKN-1 decrypts responses from 2and3, verifies nonces.
11. DRTKN-1 decrypts its own secret.
12. DRTKN-1 returns $\text{Secret-1 xor Secret-2 xor Secret-3} = \text{Original Secret}$

5

Scenarios

Device Manufacturing and HW Token Manufacturing

- 10 1. For each device 153, 154, create one internal HW Token 155, 158.
2. Seed it as described in the Key Generation section.
3. In addition, provide one or more identical Data Recovery HW Tokens.
4. Seed them as described above. Their exposed functionality is different from the regular HW Tokens. Each has its own key pair.
- 15 5. In addition, provide one or more identical Management HW Tokens 151.
6. Seed them as described above. Their exposed functionality is different from the regular HW Tokens.
7. The client receives (at minimum) a device, three hardware tokens, and
20 one hardware token reader that can be attached to a PC 150, which serves as a management station.
8. The Manufacturing process is irreversible and cannot be repeated.

25 Device Initialization (by client)

The initialization process is very sensitive and should only be done in a safe environment by trusted people.

1. Assign an IP address to the device and connect to the network. The device can be identified using a value derived from the IAKS.
2. Attach hardware token reader to some pc (Mgmt Station) and insert the MTKN into it.
3. Connect using some interface (Web, or CLI) to the device. The connection is secured using SSL or a similar technology.
4. The user authenticates to the MTKN using a user/login.
5. The Mgmt Station authenticates to the device as described above. From now on all communication between the two is done using the secured channel.
6. The Mgmt Station asks the MTKN to start the Personalization Process.
7. The MTKN generates an encrypted Personalization Request that is sent to the device.
8. The CPU on the device initiates an authentication between the DPP and the ITKN as described above.
9. The DPP generates a random number using its TRNG. This number is passed to the ITKN and seeds the PRNG.
10. The ITKN generates a new AKS and key pair, updates its own registry, and sends it to the DPP.
11. The ITKN and the DPP re-authenticate using the new AKS and establish a new session key.
12. The DPP generates the MK and passes it to the ITKN using the secure channel.

13. The ITKN encrypts it using its Public Key and saves it.
14. The CPU asks the ITKN for the $ENC_{L_{PubK}}[MK]$ and persists it to the file system.
15. The ITKN is now in token data sync mode.
- 5 16. The user defines the username and password of the HW Token administrator.
17. All the new information (new AKS, user info) is sent from the ITKN to the initiating MTKN using the secure channel between them. The MTKN uploads the new information.
- 10 18. The user now chooses the parameters for secret sharing, and must initialize the number of DRTKN's indicated.
19. The DRTKN and ITKN authenticate using their device Authenticated Public Keys. Each DRTKN generates a local key pair and sends the public key to the ITKN during the synchronization process. That key is signed by
15 the ITKN and then saved to disk. (See the Secret Sharing scenario for more information)
20. The ITKN considers itself initialized when it has updated at least one MTKN and however many DRTKNs are selected by the user.
21. At this point, Management Module asks ITKN to secret share MK, AKS.
20 Secrets are saved to disk, along with $E_{MK}(IAuthPubKey \text{ and } [ILocalPubKey]_{IAuth})$.
22. The Initialization process is irreversible and cannot be repeated.
23. The device needs to be rebooted to start normal operation.

25 Device Boot

1. The ITKN and the DPP authenticate as described above.
2. The CPU reads the $E_{L_{pubK}}(MK)$ from the file system. The encrypted MK is sent to the ITKN.
3. The CPU tells the DPP to ask the ITKN for the MK. The ITKN sends the
5 decrypted MK to the DPP over the secure channel. The loaded MK is used for both incoming and outgoing key material traffic.
4. The CPU reads the $E_{MK}(CK)$, $E_{MK}(\text{crypto config db})$, and the $E_{MK}(\text{Cryptainer certificates})$ and stores them in memory in an encrypted form
5. Per Cryptainer the device checks with the Cryptainer masters if anything
10 has changed via MK.

Device Recovery after AKS Zeroization

Requires return to device manufacturer.

15

1. Add Uninitialized MTKN
2. Existing MTKN authenticates
3. Existing management user authorizes a new MTKN by sending MTKN
Certificate = Signed MTKN Public Key
- 20 4. ITKN returns $E_{MTKN_{Pub}}(MTKN-AKS)$
5. Uninitialized MTKN decrypts and loads MTKN-AKS.

Lost MTKN

- 25 1. User authenticates to MTKN-A.
2. MTKN-A establishes secure channel w/ ITKN.

3. MTKN-A requests AKS change.
4. ITKN changes MTKN entry in AKS.
5. MTKN-A updates AKS.
6. Other MTKN's update w/ ITKN. Uninitialized MTKN's are willing to load an
5 AKS entry.
7. MTKN-A requests finalize of AKS Change.
8. ITKN no longer accepts MTKN AKS update requests. (Lost MTKN has old
AKS, so no longer works)

10 Lost DRTKN

1. User authenticates to MTKN
2. MTKN establishes secure channel w/ ITKN
3. MTKN sends ITKN update request, using DR1, DR2, DR3.
- 15 4. Management module sends secrets to ITKN, which creates secret
recovery requests (this is the secret recovery, but initiated by ITKN):
5. ITKN Auth PubKey, [ITKN local PubKey and nonce and $E_{DR1-Local}(\text{Secret-1})]_{IAuth}$
- 20 6. ITKN Auth PubKey, [ITKN local PubKey and nonce and $E_{DR2-Local}(\text{Secret-2})]_{IAuth}$
7. ITKN Auth PubKey, [ITKN local PubKey and nonce and $E_{DR3-Local}(\text{Secret-3})]_{IAuth}$
8. DRTKN1 user authenticates
9. DRTKN1 checks ITKN Auth PubKey w/ Decru Public Key, and checks
25 signature on ITKN local public key and secret.
10. DRTKN1 decrypts Secret-1, and returns $E_{ITKN local}(\text{nonce and Secret1})$

11. Repeat process with DRTKN2 and DRTKN 3.
12. ITKN decrypts responses from DRTKNs, verifies nonce.
13. ITKN recovers MK by Secret1 xor Secret2 xor Secret3.
14. ITKN enters DRTKN replacement mode.
- 5 15. DRTKN authenticate via public keys w/ ITKN. (see Basic Procedure 4)
16. ITKN signs and asks Management module to save DR Local PubKey.
17. Management Module and ITKN redo secret sharing (Basic Procedure 5)

Lost ITKN

10

1. New ITKN initializes new local key pair.
2. DRTKN-A and new ITKN handshake with public keys.
3. Recover AKS and MK from secret sharing with DRTKNs.
4. DRTKN-A sends AKS and MK to new ITKN, encrypted with ITKN Local
- 15 Key.
5. DPP and ITKN authenticate.
6. ITKN saves new $E_{L_{PubK}}(MK)$.
7. Management Module asks CryptoCard to send $E_{SK}(IAuthPubKey$ and $[ILocal PubKey]_{IAuth})$ to ITKN.
- 20 8. ITKN verifies old ITKN certificate.
9. Management Module loads old signed DR Local Public Keys into ITKN
10. ITKN verifies and signs each with new ITKN Local Key Pair and MM saves to disk.
11. MK, AKS secret sharing recomputed. New ITKN certificate saved to disk:
- 25 $IAuthPubKey$ and $[ILocal PubKey]_{IAuth}$
12. Device boot to return to normal operation.

Data Recovery without the device

1. CryptoDB entries recovered from backup tape.
- 5 2. Run Secret Recovery Process on MK, using DRTKNs.
3. Cryptainer Keys can now be decrypted manually.

Add a New Cryptainer

- 10 1. The user attaches the device's MTKN to a Mgmt Station
2. MTKN <-> ITKN authenticates.
3. The user defines the new Cryptainer
4. If in master/slave relationship, $E_{MK}(CK)$ sent to slave.
5. Optionally, new Cryptainer can be created without the need for MTKN.

15

Change Keys*Token Local Key Pair*

20 ITKN -

The MK needs to be decrypted using the old one and encrypted using the new one.

Secret Shares must be regenerated, and DR Local PubKeys must be resigned.

25

DRTKN -

Process similar to Lost DRTKN.

Master Key

- 5 1. The ITKN and the DPP authenticate as described above.
2. The CPU reads the $ENC_{L_{PubK}}[MK]$ from the file system. The encrypted MK is sent to the ITKN.
3. The CPU tells the DPP to ask the ITKN for the MK. The ITKN sends the decrypted MK to the DPP over the secure channel.
- 10 4. The Administrator changes the DPP mode of operation to maintenance mode".
5. The MTKN tells the DPP to generate a new MK. The DPP stores this new key as the Outgoing Master Key.
6. The device loads all the $ENC_{Old-MK}[CK]$ into the DPP and reads the $ENC_{New-MK}[CK]$.
- 15 $_{MK}[CK]$.
7. The new MK is persisted using the same procedure as described in "Device Initialization" (secret share also updated).
8. Replacing the old Key Stores and MK with the new Key Stores and MK should be done in a transactional way.
- 20 9. Old MK and Key Stores are deleted.

Cryptainer Key

1. This operation is only allowed to the Cryptainer Master
2. The Master notifies the slaves and access to the Cryptainer is blocked
3. The Master changes the Cryptainer Key and re-encrypts the Cryptainer
- 25 using the new key

Create Master/Slave Devices

1. Master fully initialized. Master DRTKNs are introduced to Slave ITKN.
2. Master DRTKNs recover Master MK and send to Slave ITKN.
- 5 3. Slave set to slave mode. Master configured with Slave address.
4. MK can now be used for Master to send information to Slave. Enc(CK)
look the same in both dbs.
5. Break a Master/Slave Pair
6. User authenticates to slave MTKN.
- 10 7. Slave MTKN and slave ITKN authenticate.
8. User gives cmd, Slave CryptoDB forgets MK, forgets Cryptainers,
9. Master MTKN shuts off DB duplication to slave.

Personal Key Store Creation

15

1. MTKN authorizes Cryptainer set to Personal SmartCard key store.
2. An uninitialized Personal SmartCard generates a key pair.
3. Personal SmartCard sends certificate to Device. Certificate: [Local Public
Key]_{AuthKey} [Auth Public Key]_{Decru}
- 20 4. Device verifies signature, and authentication signature of local public key.
5. Device replaces $E_{MK}(CK)$ with $E_{LocalPublicKey}(E_{MK}(CK))$.

Personal SmartCard Login

- 25 1. User loads User Authentication applet on pc.

2. User authenticates to Personal SmartCard 157, 160 (Fig. 15) (pin or login/password).
3. Device sends: nonce, $E_{\text{LocalPublicKey}}(E_{\text{MK}}(\text{CK}))$.
4. Personal SmartCard decrypts and replies with $[\text{nonce}, E_{\text{MK}}(\text{CK})]_{\text{LocalPrivKey}}$
5. Device now has $E_{\text{MK}}(\text{CK})$ and user can access Cryptainer.

The Crypto Database

Overview

10

The Crypto Database (CDB) is used to store the relationship between users, Cryptainers, and keys.

The implementation should support the following features:

15

1. Given a Cryptainer and a User the CDB should either provide a key or deny access
2. The data in the CDB should be secure – the CPU on the motherboard should not be able to examine keys, modify keys,
- 20 3. Changes in Access Control should be done in a secure way
4. Backup should be supported
5. Cryptainer Keys can be changed by an authorized user. Administrative tasks and disk access should be handled correctly during the key change process (decrypt, encrypt, update the crypto db,)
- 25 6. Cryptainer Keys can be canceled by an authorized user.

7. A Key cannot be lost (lost Key == lost data). Therefore, new keys can become active only after some backup was done. The db should track the keys' state.

8. Keys must be stored with a checksum.

5

An advanced version of the Crypto DB should also support:

1. End-users should have the option to keep their keys off the device, *e.g.* in a personal smart card.

10 2. Administrators should be allowed to flag a Cryptainer as one that requires off- device storage of keys.

3. In case off- device storage of keys is used there needs to be support for a recovery process in case the keys are lost.

4. Individual users as well as Groups can be associated with a Cryptainer.

15 The group membership can change over time.

Comment: implementing the advanced features gives a solution that is much more secure. If the keys are not stored on the device there is no way an attacker can get clear-text data.

20

Basic Features Framework

The CDB (Crypto Database) is stored in the general purpose ACL structure (the ACL structure has three DAG's: Principals, Objects, and Permissions).

25

1. Cryptainers are stored as objects, and Users are associated with the Cryptainer and are given the appropriate Permissions.
2. One of the Cryptainer's properties is the CK – specifically, store $ENC_{MK}(CK)$
- 5 3. Each Cryptainer has an Owner. Only the Owner controls access rights to the Cryptainer.
4. Administrators are allowed to create Cryptainers, but are not allowed to modify them. After creation only the Owner can modify them.

10 Scenarios – Basic Crypto DB

Creating a Cryptainer

1. The Cryptainer is added to the Objects DAG
- 15 2. An owner is associated with the Cryptainer
3. A new key is created and associated with the Cryptainer
4. Note: no Recovery Agent schema is needed in this case. A backup of the Cryptainer key is available through the system's regular backup and the fail-over configuration.

20

Adding a user to a Cryptainer

1. The Owner adds the user
2. The user is assigned the required permissions

25

Removing a user from a Cryptainer

1. The permission is taken from the user

Adding a group to a Cryptainer

5

1. The Owner adds the group
2. The group is assigned the required permissions

Removing a group from a Cryptainer

10

1. The permission is taken from the group

Accessing data inside a Cryptainer

- 15 1. The user's ACL is checked to see if the user has permission
2. If it does then the $ENC_{MK}(CK)$ is retrieved and given to the DPP

Adding a user to a group

- 20 1. No special action is required

Removing a user from a group

1. No special action is required

25

Changing a Cryptainer key

1. Notify the STPs to block access to the Cryptainer and flush its cache
2. Create new key
3. Decrypt using old key and encrypt using new one
- 5 4. Notify the STPs to unblock access to the Cryptainer

Adding or Removing of a Cryptainer's owner

1. Can be done by one of the owners
- 10 2. Can be done by the Administrator using the DRTKN (or some other privileged token)

Advanced Features Framework #1

- 15 The CDB (Crypto Database) is stored in the general purpose ACL structure (the ACL structure has three DAG's: Principals, Objects, and Permissions).
 1. Cryptainers are stored as objects, and Users are associated with the Cryptainer and are given the appropriate Permissions.
 - 20 2. Each User has a "Key Store" – the store associates Cryptainers available to the user with keys.
 3. The keys in the Key Store are stored as $ENC_{UserPubK}(ENC_{MK}(CK))$
 4. Each Cryptainer has an Owner. Only the Owner controls access rights to the Cryptainer.
 - 25 5. Administrators are allowed to create Cryptainers, but are not allowed to modify them. After creation only the Owner can modify them.

Scenarios – Advanced Features

Creating a Cryptainer

5

1. The Cryptainer is added to the Objects DAG
2. An owner is associated with the Cryptainer
3. A new key is created
4. The following is added to the Owner's Key Store: (Cryptainer ID, encrypted
10 Cryptainer key)
5. The owner may select one or more Recovery Agents, and a Secret Sharing Schema. The $ENC_{MK}(CK)$ secret is split according to the Secret Sharing Schema and added to the Key Stores of the Recovery Agents.

15 Adding a user to a Cryptainer

1. The Owner provides the CK encrypted with the MK
2. The following is added to new user's Key Store: (Cryptainer ID, encrypted
20 Cryptainer key)

Removing a user from a Cryptainer

1. The following is removed from the user's Key Store: (Cryptainer ID,
25 encrypted Cryptainer key)

Adding a group to a Cryptainer

1. The Owner provides the CK encrypted with the MK
2. The following is added to every user's Key Store: (Cryptainer ID, encrypted Cryptainer key). Note that the information about the key is propagated down the tree to every user that belongs to this group.

Removing a group from a Cryptainer

1. For each user in the group the encrypted Cryptainer key is removed from the Key Store
2. If a user belongs to some other group that has access to the Cryptainer then the key should not be removed from the Key Store

Accessing data inside a Cryptainer

1. The $ENC_{UserPubK}(ENC_{MK}(CK))$ for the Cryptainer is sent to the user.
2. The user decrypts it using its private key and sends it back
3. The $ENC_{MK}(CK)$ is given to the DPP

Adding a user to a group

1. The Owner provides the CK encrypted with the MK. Note: a User can be added to a group only by an Owner of the associated Cryptainers.
2. The following is added to the user's Key Store: (Cryptainer ID, encrypted Cryptainer key).

Removing a user from a group

1. The encrypted Cryptainer keys associated with this group are removed from the Key Store
- 5 2. If the user belongs other groups that have access to one or more of the above Cryptainers, then these keys should not be removed from the Key Store

Changing Cryptainer key

10

1. Notify the STPs to block access to the Cryptainer and flush its cache
2. Create new key
3. Decrypt using old key and encrypt using new one
4. Modify the Key Store of all the Users that have the old key
- 15 5. Notify the STPs to unblock access to the Cryptainer

Adding or Removing of a Cryptainer's owner

1. Can be done by one of the owners
- 20 2. Can be done by one of the current users using the DRTKN (or some other privileged token)

Recovering a lost Cryptainer key

- 25 1. The necessary number of Recovery Agents (as defined in the Cryptainer's creation process) should provide their part of the shared secret.

2. The recovered $ENC_{MK}(CK)$ is added to the (new) owner's key store.

Advanced Features Framework #2

- 5 This is a variation of the solution above.
 1. Each group has a key pair
 2. When a user is added to a group, the user is given the private key of the group
 - 10 3. Each group has its Key Store. (Therefore, there is no need to propagate key information down the tree to users).

Discussion

- 15 When creating a Cryptainer the user selects the key storage mechanism for the Cryptainer:
 - (a) off-device key storage; or
 - 20 (b) simple key storage.

A Cryptainer's Key should either be kept in the device or outside of the device. It does not make sense that some of the Cryptainer's Users keep the key on the device and some keep it on a smart card.

25

Changing CK in any of the advanced solutions is difficult

In advanced solution #2, after a user leaves a group the user still has the group's private key.

5 *System (Security) Configuration*

Overview

The following discussion describes a few possible configurations of the system. The system can operate in a range of security levels. Usually, there is a trade-off between administration effort and security and between cost and security.

Separating Data and Keys

15

The highest level of security requires separation of the keys from the data. This can be done by handing the keys to end-users as described in the "Advanced Crypto DB" framework.

20 Users have several ways to store their keys. All the mechanisms described below can be combined. For example, most users go with the plain vanilla, some are interested in the additional passphrase protection, and a few ask for the PKI smart card.

25 PKI Smart card – this smart card stores its owner's authentication information and a Key Pair, and is capable of running RSA internally. The public key is

used to encrypt the user's key on the device. Without the user's smart card reading the data is as difficult as breaking the block cipher. This solution requires a smart card reader and a smart card for all the users who are interested in this solution. Software must be installed on the users' PC's.

5 Instead of a physical Smart Card, one can use virtual smart card implementation, e.g. RSA SecureID+Web Passport.

Smart card - this smart card stores its owner's authentication information and a Key Pair. It is not capable of running RSA. The RSA functions run on user's PC CPU. It is less secure than the first one because the user's private
10 key is given in clear text to the user's CPU. This smart card is much cheaper (~\$3 as opposed to ~\$18). As before, this solution requires a smart card reader and a smart card for all the users who are interested in this solution. Software needs to be installed on the users' PC's.

15

Software Only – the user's PC stores a Key Pair encrypted with a passphrase. The user enables the software by typing the passphrase, after that the software can decrypt the user's key store. Less secure than above because there is no hardware token and because the private key is given to the PC's
20 CPU.

On the other hand, there is no need to have a smart card reader, a smart card, and the software is very simple and is not expected to have a complex installation procedure.

25 Key Servers – in this case the secret key is shared between the crypto device and the key server. The key server releases its part of the secret only after

explicit permission from the user. This solution is easier to manage and it provides higher level of security compared to the solution without it.

5 Simple Solution – this is the plain vanilla solution where keys are stored in the ACL associated with the Cryptainer.

Random Bits

Overview

10

Random bits are generated by a TRNG located on the DPP. The bits are encrypted using the DPP's AKS (see below). Bits from the TRNG are used to seed any PRNG in the system.

15

Preferred framework

1. Integrate a TRNG into our hardware.
2. Use the PRNG that exists on the smart card
- 20 3. Seed the PRNG using the TRNG
4. Information from the TRNG to the crypto chip is encrypted using some internal unique key (in the crypto chip) that cannot be accessed from the outside.

25 *Checksum Processing*

Overview

The goal of this document is to describe the process of checksum computation for both NFS and CIFS projects. To improve performance, it is preferred to offload checksum computation to hardware. Some of the computation can be offloaded to the Ethernet NIC. Other computation (in particular, for UDP packets) are offloaded to the DPP hardware.

Interaction between the different layers of the networking protocols is presented in Fig. 16. TCP and UDP packets are processed differently.

10

TCP: Used for CIFS 161 and NFS 162 over TCP 163. Here, rely on the checksum computed by the NIC. It is advantageous to check the checksum again, just on the payload. This requires an additional stream of information from the TCP layer upwards.

15

UDP: Used in NFS.

Outgoing: The checksum over the payload is computed by the DPP and saved inside MBufs that hold the UDP packet. After the UDP layer 164 has added the UDP headers, use the precomputed payload checksum to compute the complete UDP packet checksum, store the result in the packet, and forward it to the IP layer 165.

Incoming: Checksum over the UDP headers is computed between the IP and the UDP layers and stored in Mbufs. Add to it the UDP checksum from the UDP header. The result should be a checksum of the payload only. If

necessary, also store the offset, *i.e.* whether the UDP header length is odd or even. Before processing the packet, NFS calls the DPP to decrypt the payload and compute the checksum over the payload. It then compares the result with the checksum passed from the UDP layer.

5

UDP checksum for Block Aligned Ops

Outgoing packets

10 For outgoing UDP packets the data path is as follows:

1. Application layer: (NFS running on UDP)

Depending on whether one is reading from or writing to disk, the NFS proxy decrypts and encrypts data respectively. In either case, the NFS proxy makes the crypto hardware compute the post-crypto checksum (the checksum computed after the crypto operation). This checksum, along with the data, is sent down to the socket layer using a `so_send` (actually a `ds_sosend`) call . The `so_send` forwards the mbuf to the UDP layer

15
20

2. UDP layer:

The UDP output function extracts the checksum from the mbuf that it received. Once it has the checksum for the udp payload, it fixes the checksum value taking into account the header fields including a pseudo-IP header. The checksum is computed including the src IP

25

address and destination IP address. The headers for UDP and the 2 IP addresses are of even length. This correct checksum value is given by:

$$\text{NewCksum} = \sim(\text{Old_Cksum} + \text{sum of 16 bit words in the headers})$$

where + is done with end around carry.

The new value is written to the right location in the UDP header and the packet is sent out just like any other normal packet.

10

Incoming packets

For incoming UDP packets the data path is changed from the normal path as follows:

15

1. UDP layer:

The UDP input function, which receives the packet, fixes the checksum to remove the effects of the UDP header by subtracting out the header sum

20

as follows:

$$\text{NewCksum} = \sim(\text{Old_Cksum}) - \text{sum of 16 bit words in the headers}$$

where the – takes into account end around carry.

25

Then, the UDP headers are stripped from the mbuf and the new checksum value is written into the mbuf. Then the mbuf is sent off to the so_recv function.

5 2. Application layer: (NFS running on UDP)

The nfs proxy extracts the checksum value from the mbuf. Now irrespective of whether the system is reading from or writing to disk, data are encrypted/decrypted only on the outgoing side.

10 There are two options of where to verify the checksum:

• Perform checksum verification on the input side itself by issuing an explicit checksum command without involving crypto operations. This is a rather expensive proposition because it goes over the same data
15 twice once for check summing and again for encryption/decryption.

- Perform checksum verification along with the crypto operations. This is achieved by using the hardware to compute a pre-crypto checksum along with the post-crypto checksum (needed for the outgoing packet). This solution avoids the double traversal over the data.

20 The only problems with this are:

- a. The hardware has to compute both the checksums simultaneously (which is not much of a problem because things can be done in parallel).

- b. Any error in an incoming packet is determined only after the data has been cryptographically operated upon. This means that there is a long gap between receiving the UDP packet and finding the error. This may be a problem.
- 5 The preferred embodiment employs b. above, *i.e.* compute both the pre and post crypto checksums simultaneously.

Checksum for Non Block Aligned Ops

10 **Misaligned Reads from Disk**

A non-aligned read request from the client is handled as follows (see Fig. 17):

- 15 1. Fetch all the needed aligned blocks from disk 170, *i.e.* read excess data from disk in the first and last blocks.
2. Decrypt all the fetched data 171.
3. Send out only the required data to the client in a UDP packet 172.

This requires:

- 20 1. Verifying the checksum in the UDP packet 173 that is fetching the data from the disk. This is achieved by comparing the received checksum 174 with the pre crypto checksum 175 returned by the hardware.
2. Computing the checksum for the data to be sent to the client by fixing the edge effects 176 in the post crypto checksum 177 returned by the hardware.

Thus, the hardware must compute the pre and post crypto checksum on all the data that it is handling. Verification is a 16-bit comparison between the received checksum and the pre-crypto checksum.

5 Fixing Edge Effects

Fixing the post crypto checksum to remove the effect of the extra data handled is done as follows:

- 10 1. Compute the sum of the overhanging bytes at the left end by looking at them as 16 bit words. Use a padding byte of 0 if there are an odd number of overhanging bytes. On the average, the system adds $B/2$ bytes = $B/4$ 16 bit words. (B=8 for DES/3DES, 16 for AES)
2. Subtract out this sum from the total sum found from the post checksum.
- 15 3. If the number of bytes overhanging was odd, do a swap of the bytes of the above answer.

The same process is repeated for the other overhanging right end. Thus, on the average, there is one compare for verify and $B/2$ short adds, two short subtracts and one byte swap in the fixing and verification process.

20

Misaligned Writes to Disk

A nonaligned write request to disk involves the following steps

- 25 1. Fetch the two end blocks from disk in encrypted form.

2. Decrypt the fetched blocks and write into the ends of a temp buffer.
3. Write into the temp buffer at the correct offset the data that came from the client.
4. Encrypt the temp buffer.
5. Send out the whole temp buffer to the disk to write.

As far as checksums are concerned, the system:

1. Verifies the checksums in the UDP packets that fetching the two end blocks from the disk in Step 1 above. Because the block size is small this is done in software. This adds a cost of 'B' short adds.
 2. Verifies the checksum in the UDP packet that came from the client with the data to be written.
 3. Computes the checksum for the UDP packet to be sent to the disk with the encrypted data to be written.
- The first goal is done in software and the last two are done with hardware support as shown in Fig. 18, where data from the client 180 are converted to plain text in a temporary buffer 181 and where encrypted data are written to disk 182. A checksum process verifies the received data 183, fixing edge effects 184 as discussed above, using the pre crypto checksum 185 and post crypto checksum 186, which is the sent value 187.

As in the misaligned reads case, the hardware must compute the pre and post checksum of all the data it is handling. It does not have to worry about starting or ending the check-summing process at specific locations.

- 5 The removal of edge effects is done for the pre-crypto checksum in this case unlike the read case. This fixing process is identical to the fixing done in the read case, except that the software has to remember the plain text versions of the two end blocks found in Step 2 above. Once fixing is done, it can be verified that the fixed version and the checksum received with the client data
10 match.

The checksum for the encrypted data being sent to the disk is the post crypto checksum given by hardware.

15 **Hardware's Checksum interface**

In all the above cases, it is not necessary to tell the hardware when to start or stop computing the checksum. Thus, it is fine if the hardware always computes on all data that it encrypts or decrypts both before and after the
20 crypto operation. This avoids having to set up a separate 64-bit descriptor in hardware to control checksumming. Setting up the encryption or decryption descriptor is good enough for checksumming too.

But the absence of a separate checksum descriptor poses the question as to
25 how the hardware returns the two 16-bit checksums that it computes to software. A couple of solutions are listed below:

1. The hardware could overwrite the first 32 bits of the descriptor for which the checksum was computed with the checksums. This is probably not a good solution because the software must watch out for overwriting the results inadvertently and the software has to be given some mechanism of looking at the descriptor table.
2. Another solution is to pre-allocate a fixed area in memory where the hardware writes the 32-bit value into a location indexed by the descriptor number. The software should examine this location only after it is sure that the corresponding crypto operation has been completed.

Sending Checksum across network layers

There are presently three possible ways of sending the checksum between the application and UDP layers:

15

- a. Checksum as part of the UDP payload

From Application to UDP:

20

The application layer computes the payload checksum using crypto hardware and writes it down as the first two bytes of the data field of the mbuf. It also writes the next two bytes of the mbuf data with a special status word that indicates that this is a packet whose checksum must be fixed.

The socket layer forwards the mbuf to the UDP layer. Within the UDP layer we extract the checksum data from the mbuf as the first four bytes of the mbuf's data (two byte checksum and a two byte signal indicating that this indeed a "Device" situation where the checksum must be fixed), prepend the UDP header to the mbuf in such a way that it overwrites the first four bytes of the original mbuf and, finally, fix the checksum value taking into account the 'pseudo_ip header' and the UDP header values, before sending the mbuf to the IP layer.

10

To Application from UDP:

The UDP layer computes the checksum for the payload only by subtracting out the effect of the UDP header. Then, while stripping the UDP headers, it leaves the last four bytes of the UDP header to remain in the mbuf. These four bytes are overwritten with the 2-byte payload checksum and a 2-byte status word. This mbuf is sent to the layer.

15
20

The socket layer forwards the mbuf to the nfs layer, which strips the first four bytes of the mbuf data to retrieve the checksum and status.

b. Checksum as a field in the mbuf

25

The Mbuf structure has a 16-bit checksum field, which is used if checksum is computed by a hardware accelerator. Further,

there is a flag bit in the Mbuf that indicates whether the checksum data in the mbuf is valid or not.

From Application to UDP:

5

The application layer computes the payload checksum using crypto hardware and writes it down as the 2-byte `csum_data` field of the mbuf. It also sets the `csum_valid` bit of the mbuf data.

10

The socket layer forwards the mbuf to the UDP layer. Within the UDP layer, fix the checksum value in the `csum_data` field to account for headers and then write the correct value in to the UDP header.

15

To Application from UDP:

20

The UDP layer computes the checksum for the payload only, by subtracting out the effect of the UDP header. Then, after stripping the UDP headers, it writes the `csum_data` field of the mbuf and sets the valid bit. This mbuf is sent to the layer.

The socket layer forwards the mbuf to the nfs layer, which retrieves the checksum from the mbuf field.

25

- c. Use a control Mbuf separate from the data Mbuf

From Application to UDP:

5 Both the socket layer and the UDP layer below it accept a control mbuf parameter as part of their 'so_send' and 'udp_output' and interface. Write the checksum in this control mbuf along with status. The socket layer hands off the control mbuf to the layer below without examining it. The UDP layer can examine this control mbuf to get the checksum for the data, fix it, and then write the correct checksum value in the UDP header.

10 From UDP to Application:

15 The so_recv interface also has a control mbuf parameter. So is for sending the checksum up the stack. The udp layer creates this control mbuf and writes the payload checksum in it along with status. The socket layer forwards this control mbuf to the application. The application layer reads the checksum and frees the control mbuf.

Comments

20

- Methods a. and c. above allow a status word along with the checksum to be communicated across the layers. This is useful to turn on the checksum fix selectively only if the sender is the device proxy. It is possible to leave the checksum fix feature always on for all udp packets, if the device proxy is the only application using udp. Method b. above also allows a limited form of status in the mbuf flags by setting or unsetting the csum_valid bit.
- 25

It is possible to avoid the status bytes in a. and c. by examining the udp payload to check if it is from the device proxy.

- In all three methods on the input side, if checksum fix is to be done selectively only for packets destined for the device proxy, the udp layer must examine the udp payload to determine which application is receiving the mbuf finally. It is not enough to just know to which socket the mbuf is being sent, the system needs info across two layers. This could be expensive. The system can avoid this if it is assumed that only the device proxy uses udp.
- 10
- All three methods involve changes in:
 - The nfs proxy to take into account hardware check summing and for implementing one of the above three methods to sending and receiving checksum values.
 - The udp layer, more specifically the `udp_usrreq.c` file containing the udp input and output functions
 - Some changes may be required in the socket layer to make sure that the control mbuf is sent as is for method c.
- 15
- It is more efficient to have the hardware compute the pre and post crypto checksum simultaneously.
- 20
- Method b. is presently preferred because it is the easiest to implement.

Checksum definition

The UDP checksum on a sequence of bytes is computed by doing the following:

1...Sum up the sequence looking at it as 16 bit words. This summing is done
 5 with an end around carry. Thus, if a carry ever occurs it is added to the LSB of the sum.

2...If a last byte exists which has not been added, *i.e.* the total number of bytes is 0, conceptually pad it with a trailing 0 byte and add to the sum.

3...Invert the 16-bit sum, *i.e.* take its one's complement.

10

Justification for byte swap while fixing odd overhangs

Consider the case when the left end block is:

15 A B C D E F G H

Where each of A to H is 8 bits.

The hardware is going to compute the checksum as:

20

A B +

C D +

E F +

G H +

25 =X Y say -----(1)

Now if the overhang is A B C, i.e. only checksum for DEFGH is wanted, compute:

$$\begin{aligned}
 & D E + \\
 5 \quad & F G + \\
 & H 0 \\
 & = X'Y' \text{ say}
 \end{aligned}$$

This can be obtained from the checksum for ABCDEFGH as follows:

10

$$XY = A B + C D + E F + G H$$

$$= A B + C 0 + 0 D +$$

$$E F +$$

$$G H +$$

15

$$0 0$$

$$= A B + C 0 + 0 0 +$$

$$E D +$$

$$G F +$$

20

0 H (because by moving down the last byte of each number, the total should not change)

$$= \text{Sum of overhang} + Y'X'$$

Hence $X'Y' = \text{byte swap}(XY - \text{sum of overhang})$.

DPP Micro Controller

25

Introduction

One embodiment of the invention provides a PCI DPP for FIPS-140 level-3 certification. Among other things, level-3 certification requires identity-based authentication before execution of any crypto-related operation. In the preferred embodiment, this means authentication is required for MK and AKS operations, CK operations including context setting, and encryption and decryption of data.

10 The following discussion describes a preferred implementation that satisfies FIPS requirements.

The Preferred System

15 The preferred system (see Fig. 19) includes an x86 board 190 with a PCI DPP 191. The PCI card includes an FPGA 192 that contains the PCI core 193, a command FIFO 194, RAM 195, crypto engine 199, and interfaces to SRAM 196, Flash 197, and a Micro Controller 198. The Micro Controller is a typical smart card controller. It contains a processing unit, PKI module, 3DES
20 module, RAM, and EEPROM. A (slow) serial communication line connects the FPGA to the Micro Controller.

Solution Framework

25 1. Operations on the PCI card requires authentication

2. Authentication is identity based.
3. There is only one active session at a time. A user's session terminates when the user explicitly logs out or when another user logs in.
4. A user can assume one of the following roles:
 - 5 i. Crypto Officer – this user is allowed to load MK, and modify users.
 - ii. Crypto User – this user is allowed to do all operations except for MK load, and users modification.
 - iii. Maintenance Role – this user is allowed to do maintenance work.
5. The DPP has several modes of operation:
 - 10 i. Un-initialized Mode – the module stays in this mode until a MK is loaded by a Crypto Officer
 - ii. Initialized Mode – the module is in this mode after a valid MK was loaded
 - 15 iii. Maintenance Mode – when a user assumes a maintenance role. All key material is zeroized upon entering or exiting this mode.
 - iv. In-session Mode – there is an active session (a user is logged in)
 - v. No-session Mode – no active session (no user is logged in)
6. All the DPP commands after a valid login are associated with the most recent user session.
- 20 7. The present system is limited to a maximum of ten users
8. There are two types of operations the client can invoke:
 - i. FPGA Operations – these operations are verified and processed by the FPGA
 - ii. Micro Controller Operations – these operation are given as is to the
25 Micro Controller (“pass through”)
9. Key Material Storage:

i. Key Material that is needed across power cycles is stored in the
Micro Controller RAM

ii. This RAM is zeroized when necessary

5 10. Device Public Key is loaded into the Micro Controller and cannot be
modified after the card left the factory.

11. There is a way to field upgrade the Micro Controller code.

12. The FPGA doesn't initiate communication with the Micro Controller. The
initiation is controlled by the DPP's user (client).

10 13. The Micro Controller output (see Table "A" below) is prefixed in such a
way that the FPGA can direct the output to the right location

14. The DPP exposes a register that signals if the Micro Controller is busy or
not. It is the client's responsibility not to send requests to the Micro
Controller when it is busy.

15 Although the invention has been described herein with reference to certain
preferred embodiments, one skilled in the art will readily appreciate that other
applications may be substituted for those set forth herein without departing
from the spirit and scope of the present invention. Accordingly, the invention
should only be limited by the Claims included below.

20

CLAIMS

1. An encryption based security apparatus for network storage, comprising:
- 5 one or more storage devices;
- one or more client devices; and
- an encryption device for separating access to said one or more storage devices via said one or more client devices from access to client data stored on said one or more storage devices;
- 10 wherein said encryption device encrypts all client data that is stored on said one or more storage devices.
2. The apparatus of Claim 1, said encryption device comprising:
- at least two network interfaces, comprising:
- 15 a clear text network interface that is connected to said one or more clients; and
- a secure network interface that is connected to said one or more storage devices;
- wherein each network interface supports multiple network
- 20 nodes.
3. The apparatus of Claim 1, wherein said encryption device is located in a network, on a path between said one or more client devices and said one or more storage devices.

25

4. The apparatus of Claim 1, wherein said encryption device intercepts all packets that flow between said one or more client devices and said one or more storage devices.

5 5. The apparatus of Claim 1, wherein said encryption device distinguishes between data, command, and status information.

6. The apparatus of Claim 1, wherein said encryption device encrypts said client data sent from said one or more client devices to said one or more storage devices using an encryption method; wherein said encryption device
10 decrypts said client data sent from said one or more client storage to said one or more client devices using said encryption method; and wherein said encryption device passes command and status information through without modification.

15

7. The apparatus of Claim 1, said encryption device comprising:

a plurality of keys for encrypting and decrypting said client data; and

a selection module for selecting a particular one of said plurality of keys based upon at least one of client identification, data location on said one or
20 more storage devices, file name, permission structure, and other factors.

8. The apparatus of Claim 7, wherein a plurality of encryption devices share keys and selection modules between them, wherein client data encrypted by one encryption device can be decrypted by another encryption device.

25

9. The apparatus of Claim 1, wherein said encryption device operates in a transparent fashion to both of said one or more client devices and said one or more storage systems;

5 wherein no modification is required to either of said one or more client devices and said one or more storage devices to enable storage of encrypted client data on said one or more storage devices, and to enable subsequent retrieval and decryption of said client data.

10 10. The apparatus of Claim 1, wherein said encryption device performs any of decoding device storage related traffic, extracting of a payload, and transparent encryption of said payload.

15 11. The apparatus of Claim 1, wherein said encryption device operates in any of SAN, NAS (NFS/CIFS), and HTML environments.

20 12. The apparatus of Claim 7, wherein said encryption keys are generated inside said encryption device and are not transmitted therefrom in cleartext form.

25 13. The apparatus of Claim 7, wherein said encryption keys are encrypted with a master key that is generated inside said encryption device.

14. The apparatus of Claim 13, wherein said master key is never transmitted from said encryption device.

25

15. The apparatus of Claim 13, wherein said master key is transmitted from said encryption device using a shared key, wherein said key is broken into a plurality of key parts, wherein possession of some subset of said plurality of key parts is sufficient to reconstruct said master key.
- 5
16. The apparatus of Claim 7, wherein a user is allowed to have an additional key, wherein said user key, a Cryptainer key, and a master key are required to encrypt said client data.
- 10 17. The apparatus of Claim 7, wherein keys are assigned to any of a region and a directory.
18. The apparatus of Claim 1, further comprising:
- a client application; and
 - 15 a server comprising a server application;
- wherein user passwords are propagated to said encryption device; and wherein said server can not access client data except as permitted by said user.
- 20 19. The apparatus of Claim 1, wherein said encryption device comprises:
- a module for initially encrypting and re-encrypting said client data without interrupting access to said one or more storage devices by said one or more client devices.
- 25 20. The apparatus of Claim 1, wherein said encryption device comprises:
- a module for administrator management of said client data;

wherein said client data are not readable by said administrator.

21. The apparatus of Claim 1, wherein said encryption device comprises:

5 a module for allowing a first user to control which other users may
access said first user's client data without explicitly sharing a first user's key
with said other users.

22. The apparatus of Claim 1, wherein said encryption device combines
encryption with mirroring.

10

23. The apparatus of Claim 1, wherein said encryption device permits any of
private and shared mirror configurations.

24. The apparatus of Claim 1, wherein said encryption device permits
15 selective remote client data access.

25. The apparatus of Claim 1, wherein said encryption device permits a
failover configuration in connection with encryption and decryption of storage
device related traffic.

20

26. The apparatus of Claim 1, wherein said encryption device transparently
forwards traffic that is not storage related.

27. An encryption based security apparatus for network storage, comprising:
25 an encryption device, located in a network, on a path between one or
more client devices and one or more storage devices, for separating access

to said one or more storage devices via said one or more client devices from access to client data stored on said one or more storage devices;

said encryption device comprising at least two network interfaces, comprising a clear text network interface that is connected to said one or more clients; and a secure network interface that is connected to said one or more storage devices; wherein each network interface supports multiple network nodes;

wherein said encryption device encrypts all client data that is stored on said one or more storage devices.

10

28. An encryption based security method for network storage, comprising the steps of:

providing one or more storage devices;

providing one or more client devices; and

separating access to said one or more storage devices via said one or more client devices from access to client data stored on said one or more storage devices with an encryption device;

wherein said encryption device encrypts all client data that is stored on said one or more storage devices.

20

29. The method of Claim 28, said encryption device comprising:

at least two network interfaces, comprising:

a clear text network interface that is connected to said one or more clients; and

a secure network interface that is connected to said one or more storage devices;

wherein each network interface supports multiple network nodes.

30. The method of Claim 28, wherein said encryption device is located in a network, on a path between said one or more client devices and said one or more storage devices.

31. The method of Claim 28, wherein said encryption device intercepts all packets that flow between said one or more client devices and said one or more storage devices.

32. The method of Claim 28, wherein said encryption device distinguishes between data, command, and status information.

33. The method of Claim 28, wherein said encryption device encrypts said client data sent from said one or more client devices to said one or more storage devices using an encryption method; wherein said encryption device decrypts said client data sent from said one or more client storage to said one or more client devices using said encryption method; and wherein said encryption device passes command and status information through without modification.

34. The method of Claim 28, further comprising the steps of:

providing a plurality of keys for encrypting and decrypting said client data; and

providing a selection module for selecting a particular one of said plurality of keys based upon at least one of client identification, data location on said one or more storage devices, file name, permission structure, and other factors.

5

35. The method of Claim 34, wherein a plurality of encryption devices share keys and selection modules between them, wherein client data encrypted by one encryption device can be decrypted by another encryption device.

10 36. The method of Claim 28, wherein said encryption device operates in a transparent fashion to both of said one or more client devices and said one or more storage systems;

wherein no modification is required to either of said one or more client devices and said one or more storage devices to enable storage of encrypted
15 client data on said one or more storage devices, and to enable subsequent retrieval and decryption of said client data.

37. The method of Claim 28, wherein said encryption device performs any of decoding device storage related traffic, extracting of a payload, and
20 transparent encryption of said payload.

38. The method of Claim 28, wherein said encryption device operates in any of SAN, NAS (NFS/CIFS), and HTML environments.

39. The method of Claim 34, wherein said encryption keys are generated inside said encryption device and are not transmitted therefrom in cleartext form.

5 40. The method of Claim 34, wherein said encryption keys are encrypted with a master key that is generated inside said encryption device.

41. The method of Claim 40, wherein said master key is never transmitted from said encryption device.

10

42. The method of Claim 40, wherein said master key is transmitted from said encryption device using a shared key, wherein said key is broken into a plurality of key parts, wherein possession of some subset of said plurality of key parts is sufficient to reconstruct said master key.

15

43. The method of Claim 34, wherein a user is allowed to have an additional key, wherein said user key, a Cryptainer key, and a master key are required to encrypt said client data.

20 44. The method of Claim 34, wherein keys are assigned to any of a region and a directory.

45. The method of Claim 1, further comprising the steps of:

providing a client application; and

25 providing a server comprising a server application;

wherein user passwords are propagated to said encryption device; and

wherein said server can not access client data except as permitted by said user.

46. The method of Claim 28, further comprising the step of:

5 initially encrypting and re-encrypting said client data without interrupting access to said one or more storage devices by said one or more client devices.

47. The method of Claim 28, further comprising the step of:

10 providing a module for administrator management of said client data; wherein said client data are not readable by said administrator.

48. The method of Claim 28, further comprising the step of:

15 allowing a first user to control which other users may access said first user's client data without explicitly sharing a first user's key with said other users.

49. The method of Claim 28, wherein said encryption device combines encryption with mirroring.

20

50. The method of Claim 28, wherein said encryption device permits any of private and shared mirror configurations.

51. The method of Claim 28, wherein said encryption device permits selective

25 remote client data access.

52. The method of Claim 28, wherein said encryption device permits a failover configuration in connection with encryption and decryption of storage device related traffic.

5 53. The method of Claim 28, wherein said encryption device transparently forwards traffic that is not storage related.

54. An encryption based security method for network storage, comprising the steps of:

10 providing an encryption device, located in a network, on a path between one or more client devices and one or more storage devices;

separating access to said one or more storage devices via said one or more client devices from access to client data stored on said one or more storage devices;

15 said encryption device comprising at least two network interfaces, comprising a clear text network interface that is connected to said one or more clients; and a secure network interface that is connected to said one or more storage devices; wherein each network interface supports multiple network nodes;

20 wherein said encryption device encrypts all client data that is stored on said one or more storage devices.

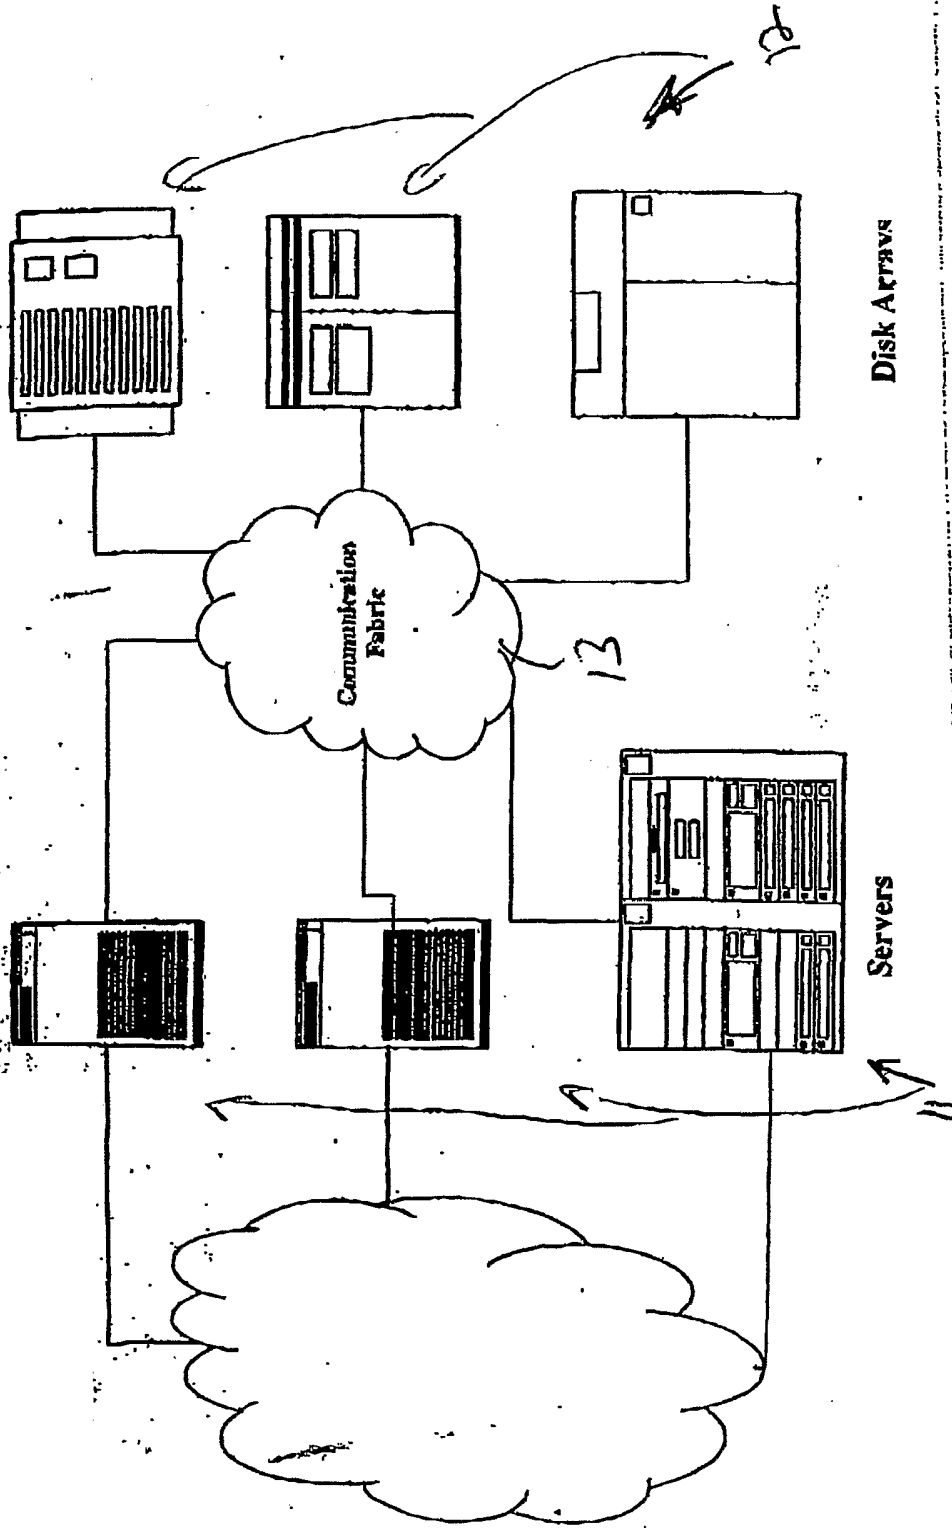
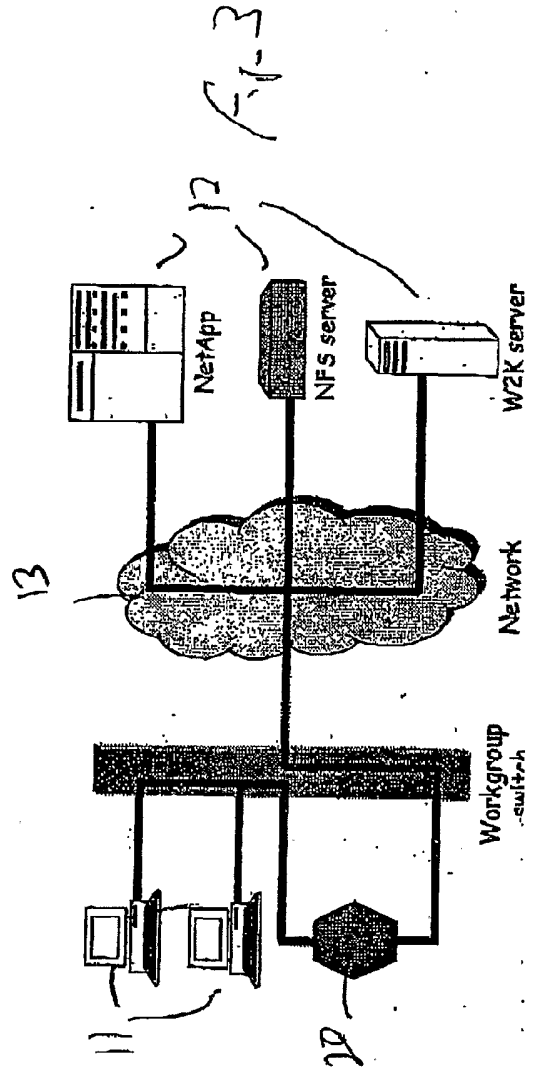
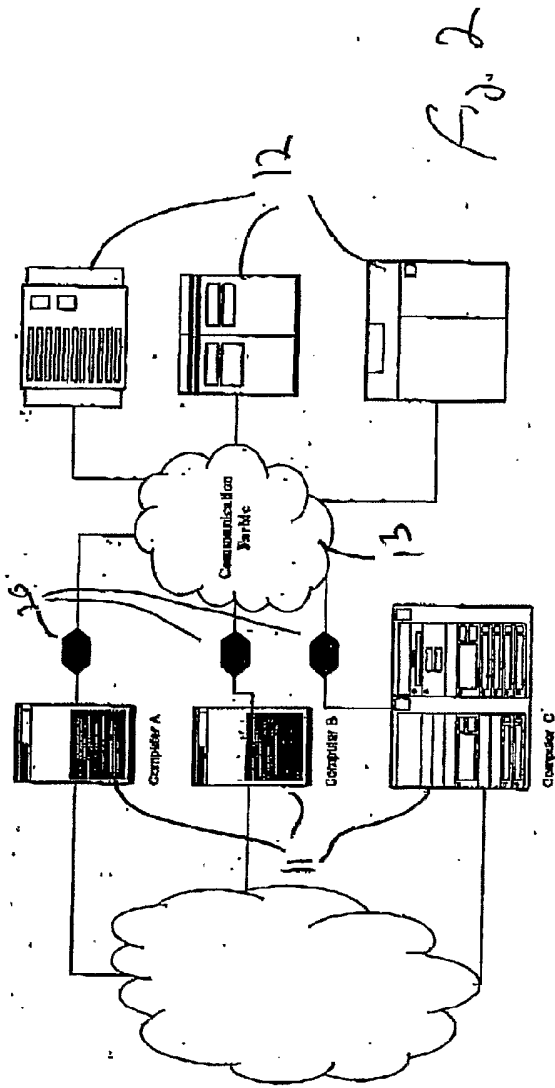
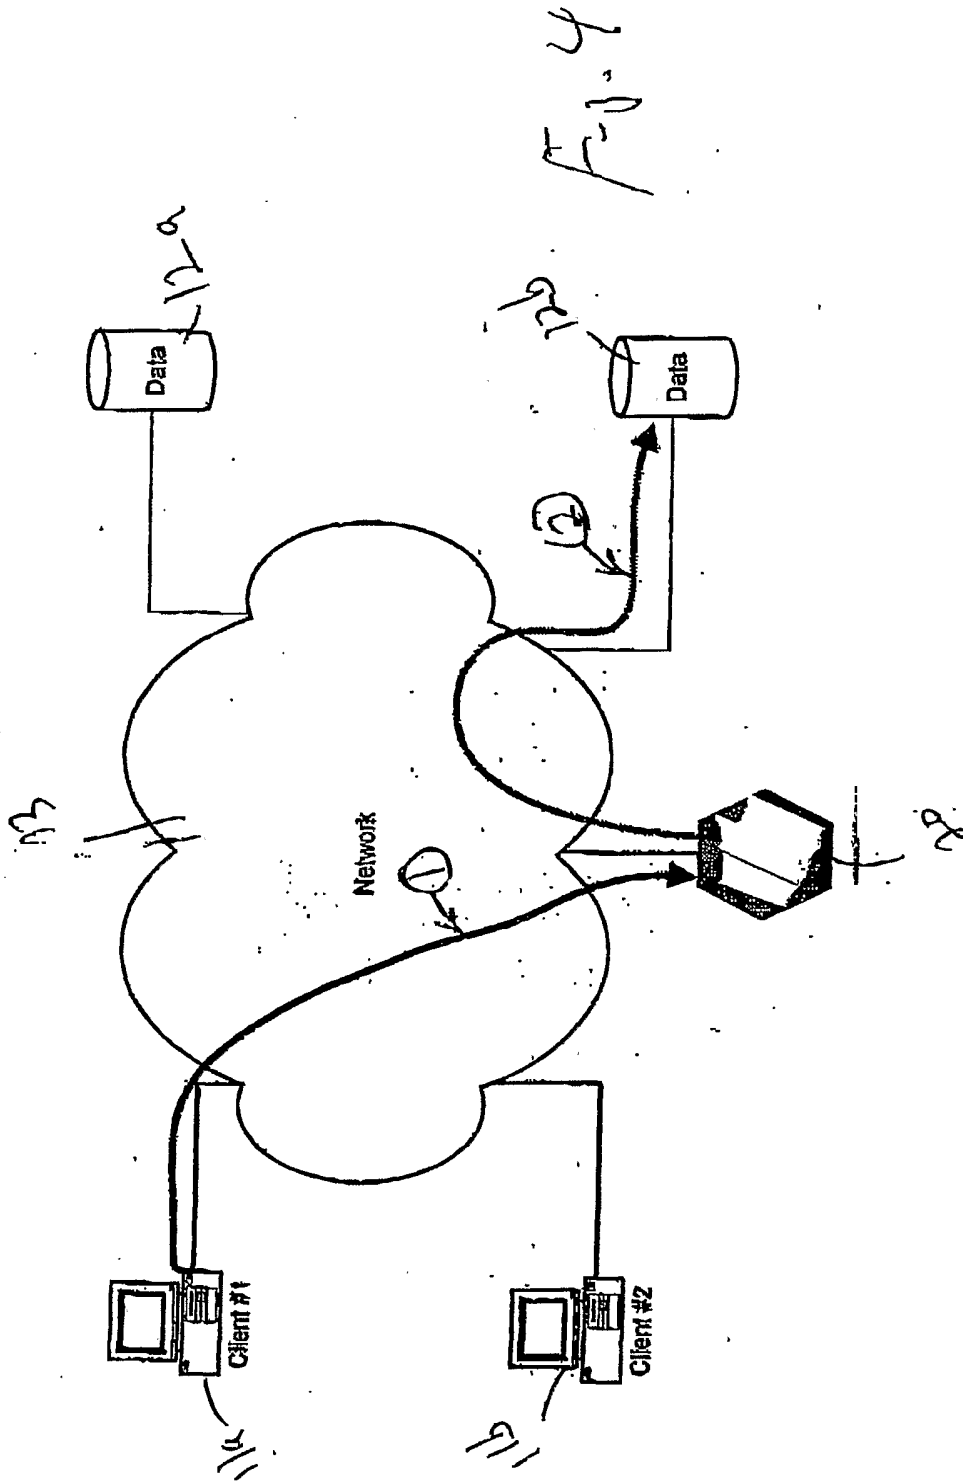
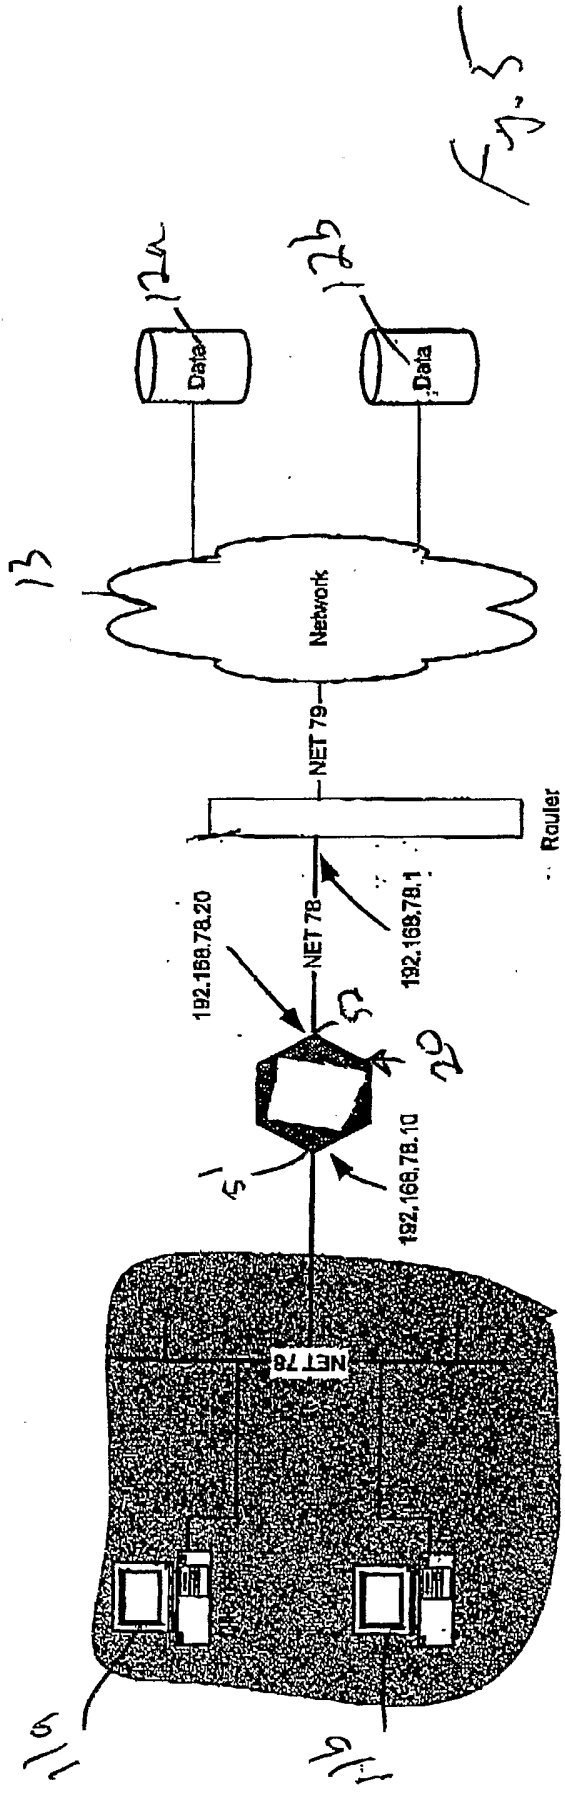


Fig. 1 (Prior Art)







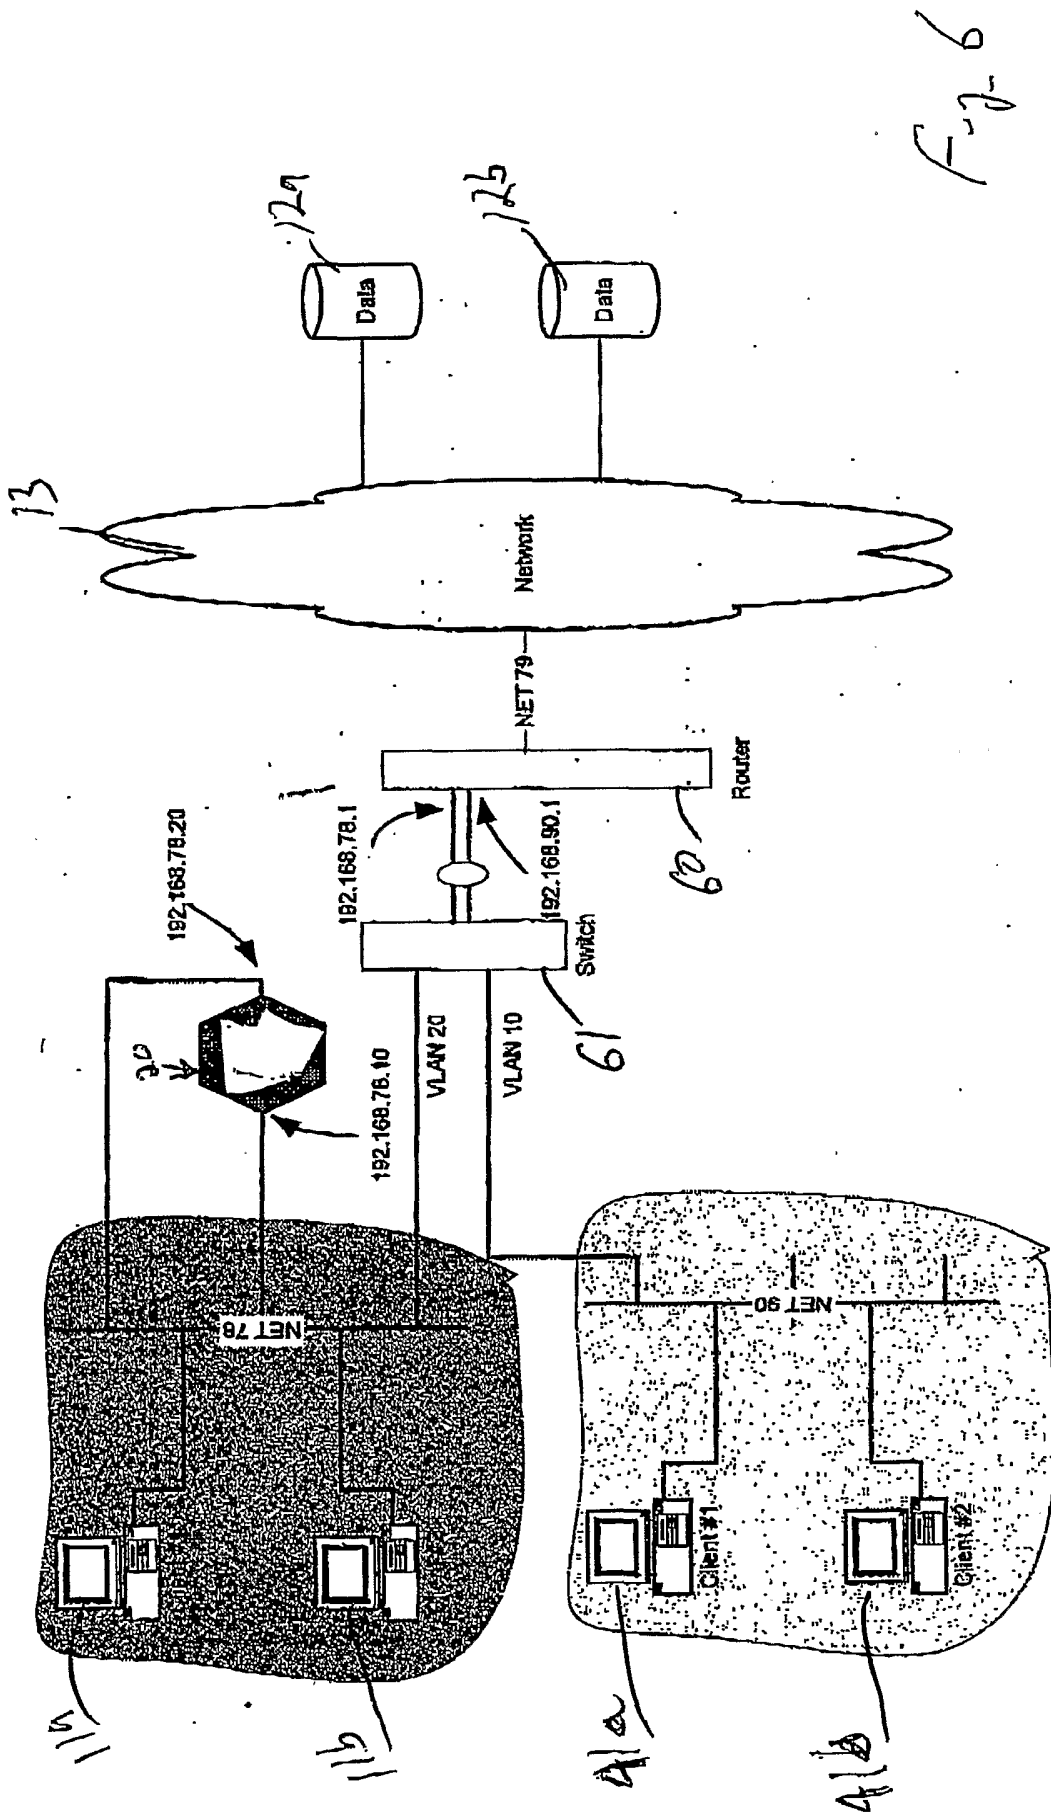
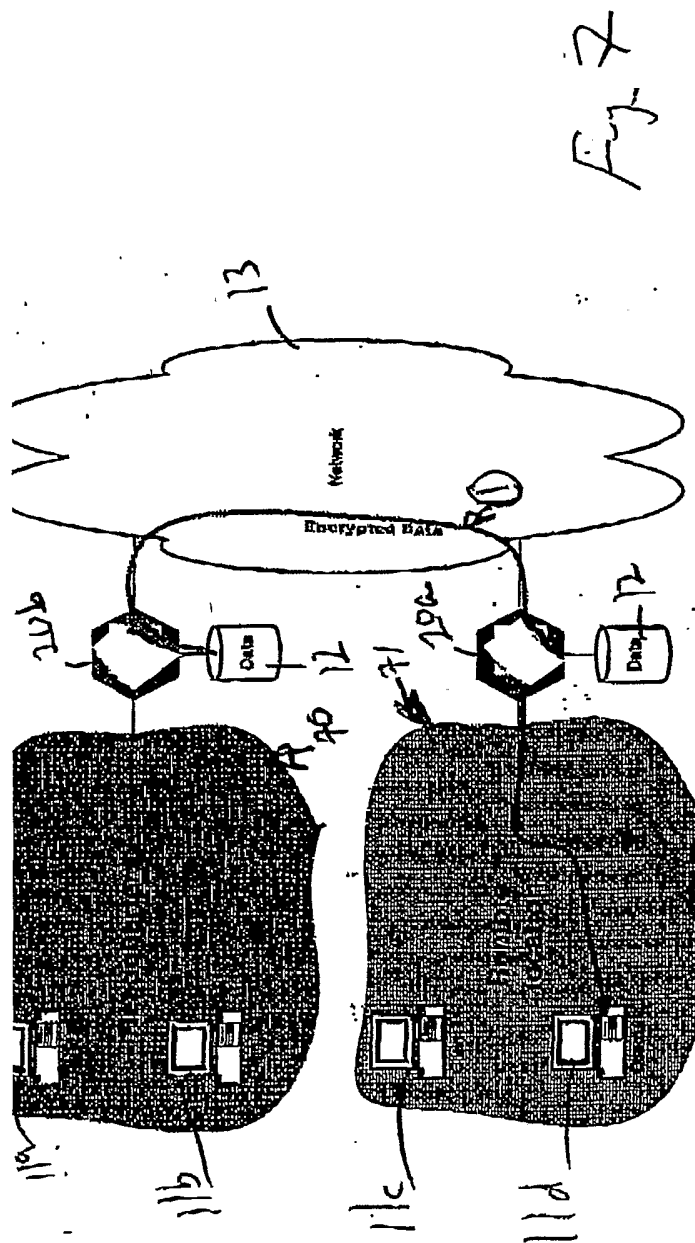
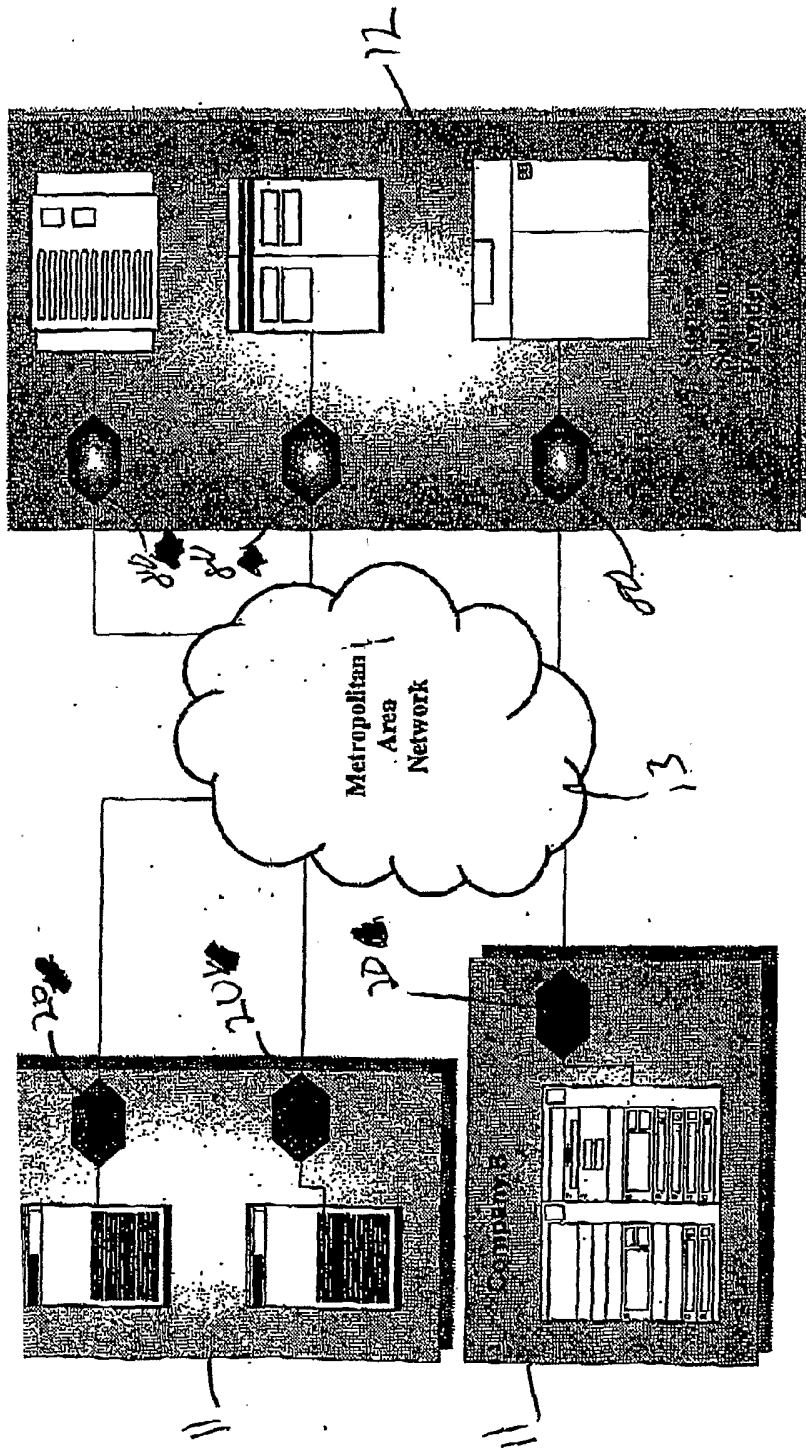


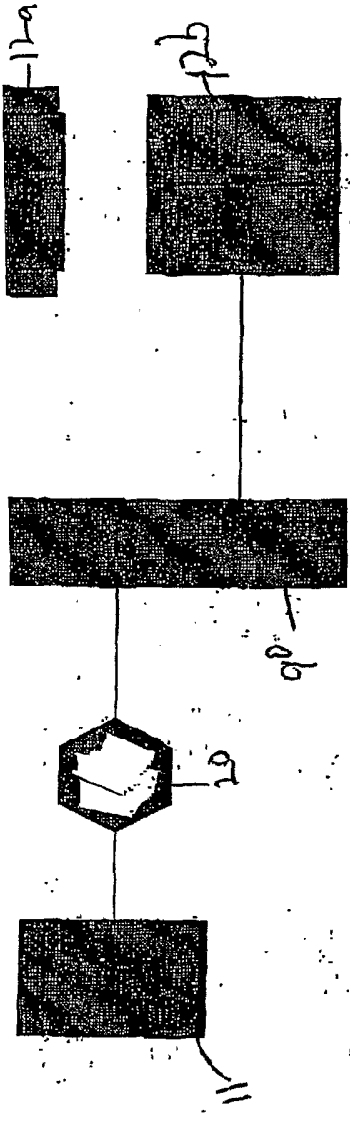
Fig. 6





— [hexagon] — appliance — [hexagon] — access controller

Fig. 8



Disk Arrays

Server

Fig. 9

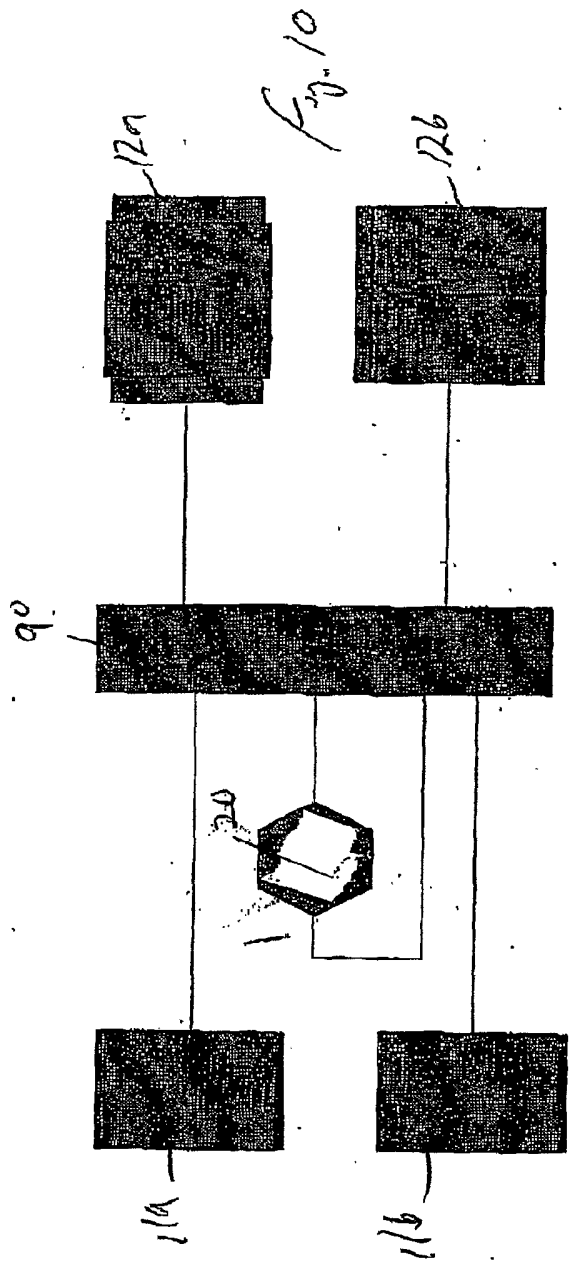
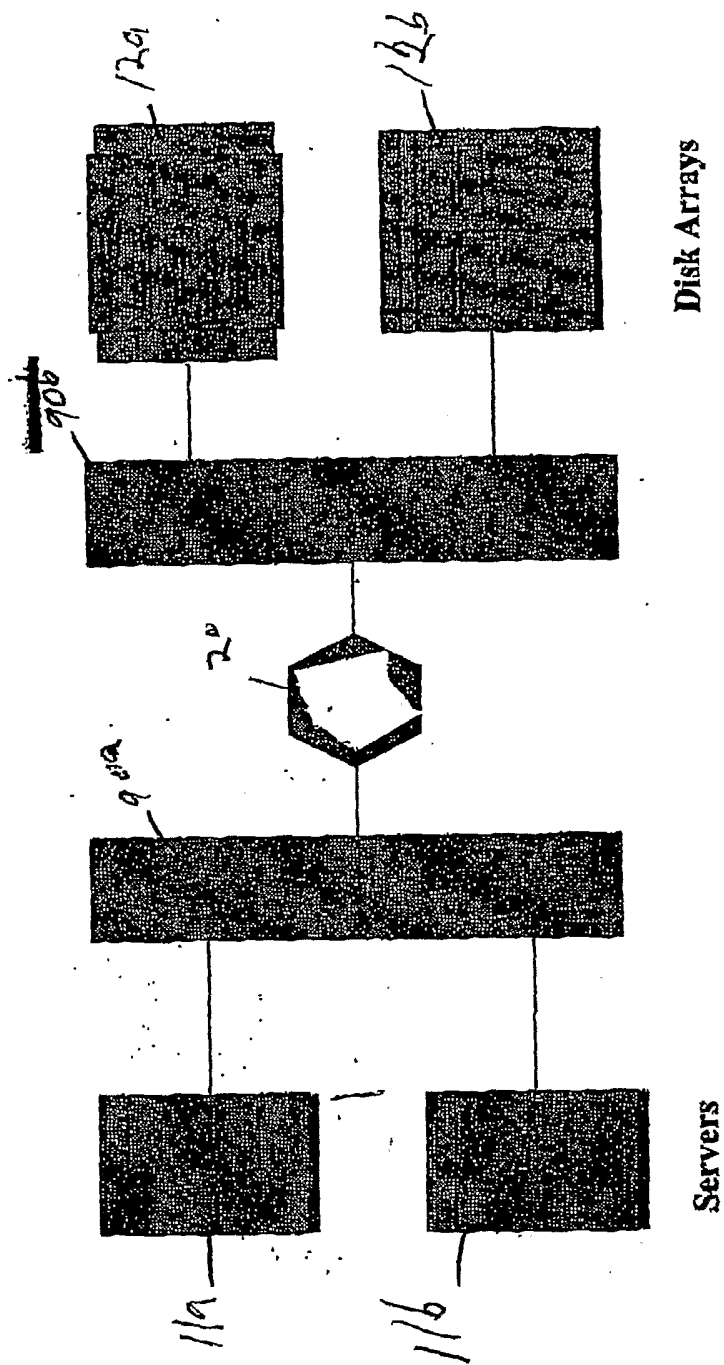


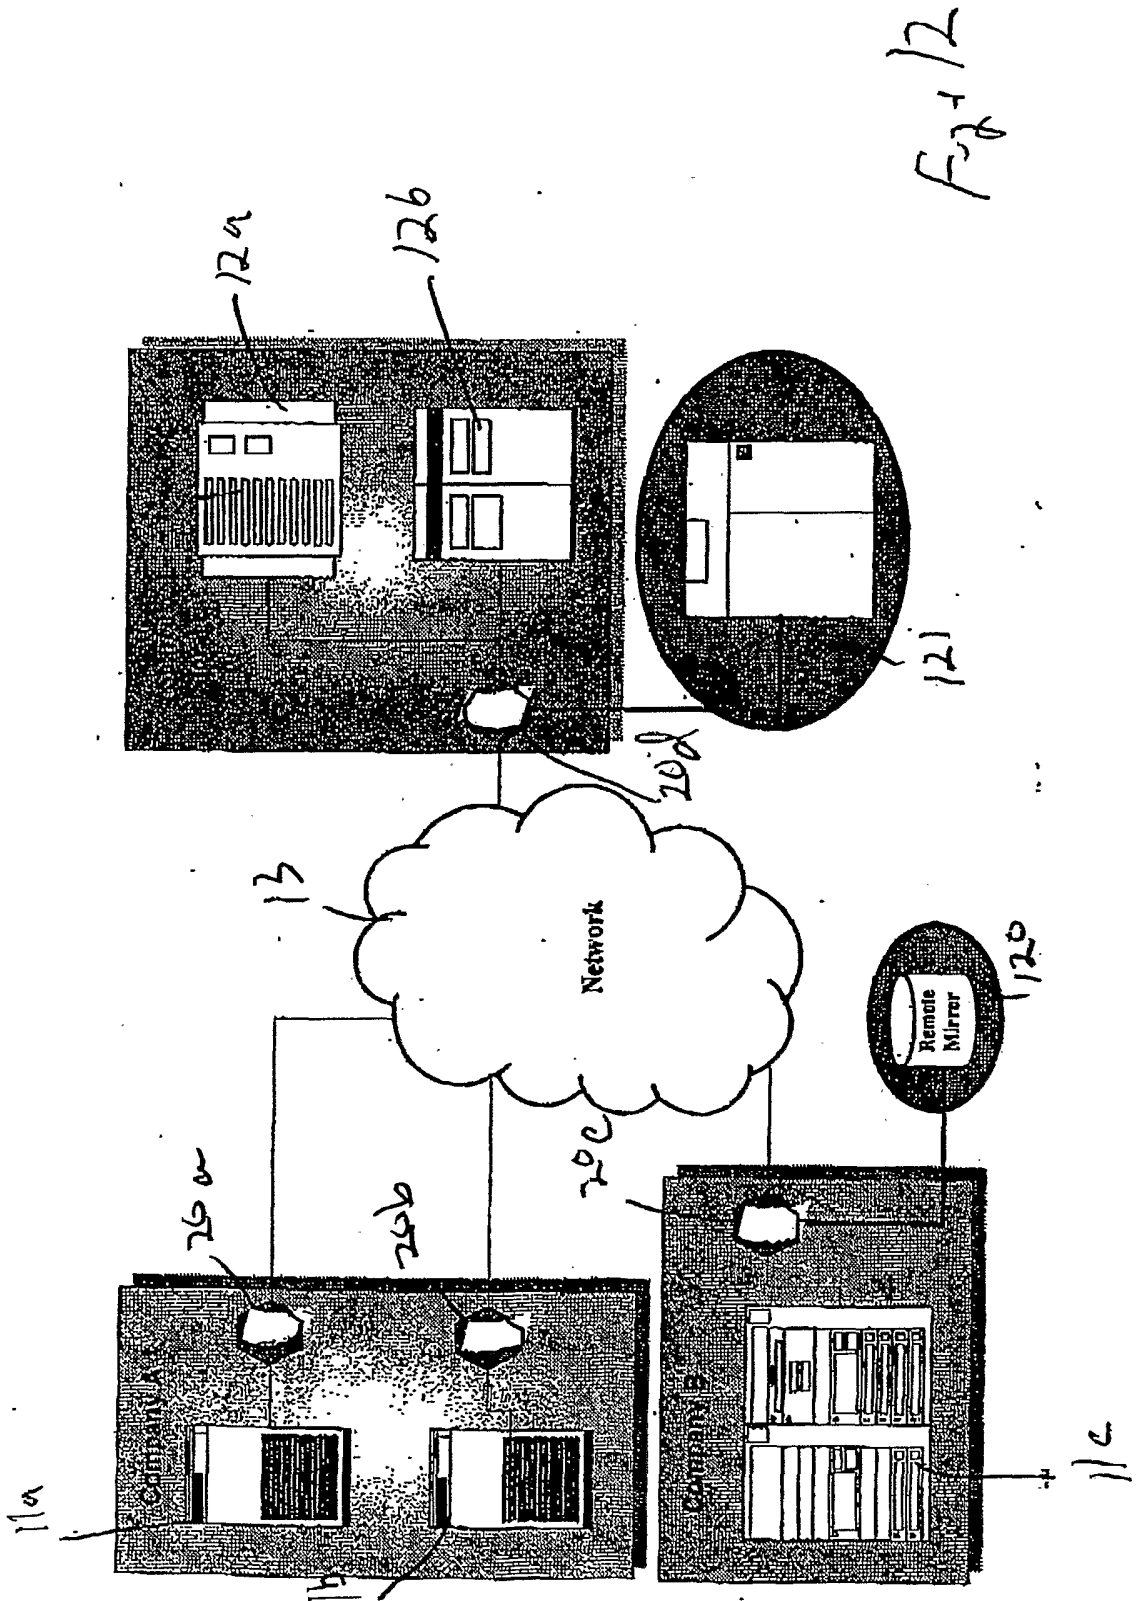
Fig. 10

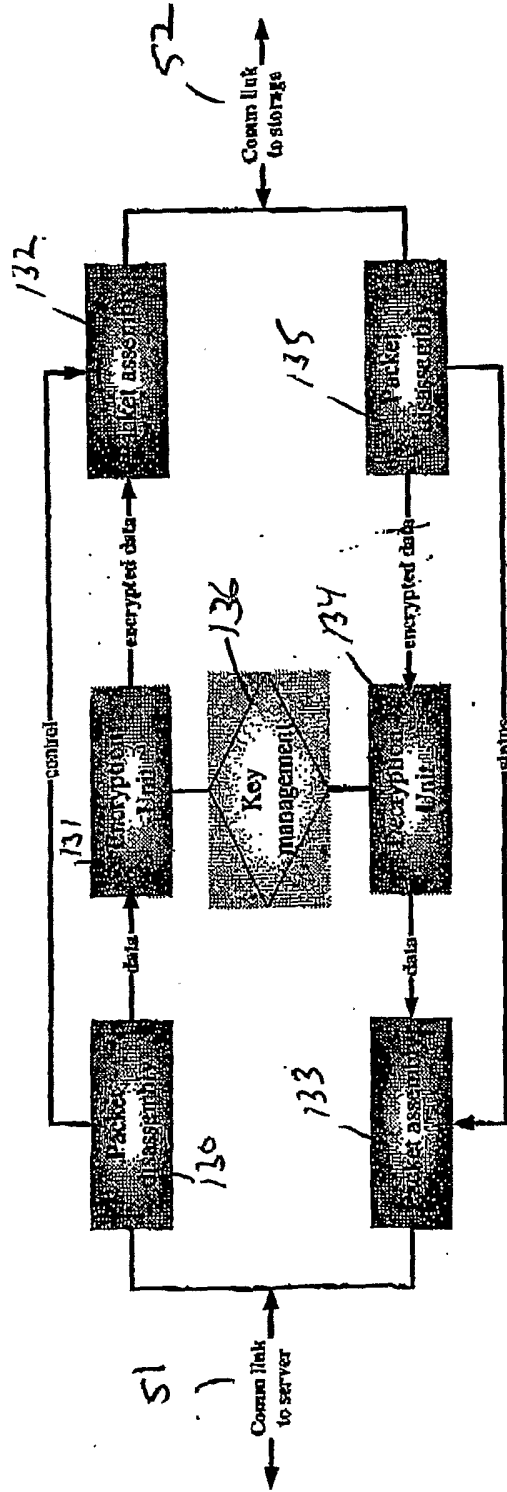


Disk Arrays

Servers

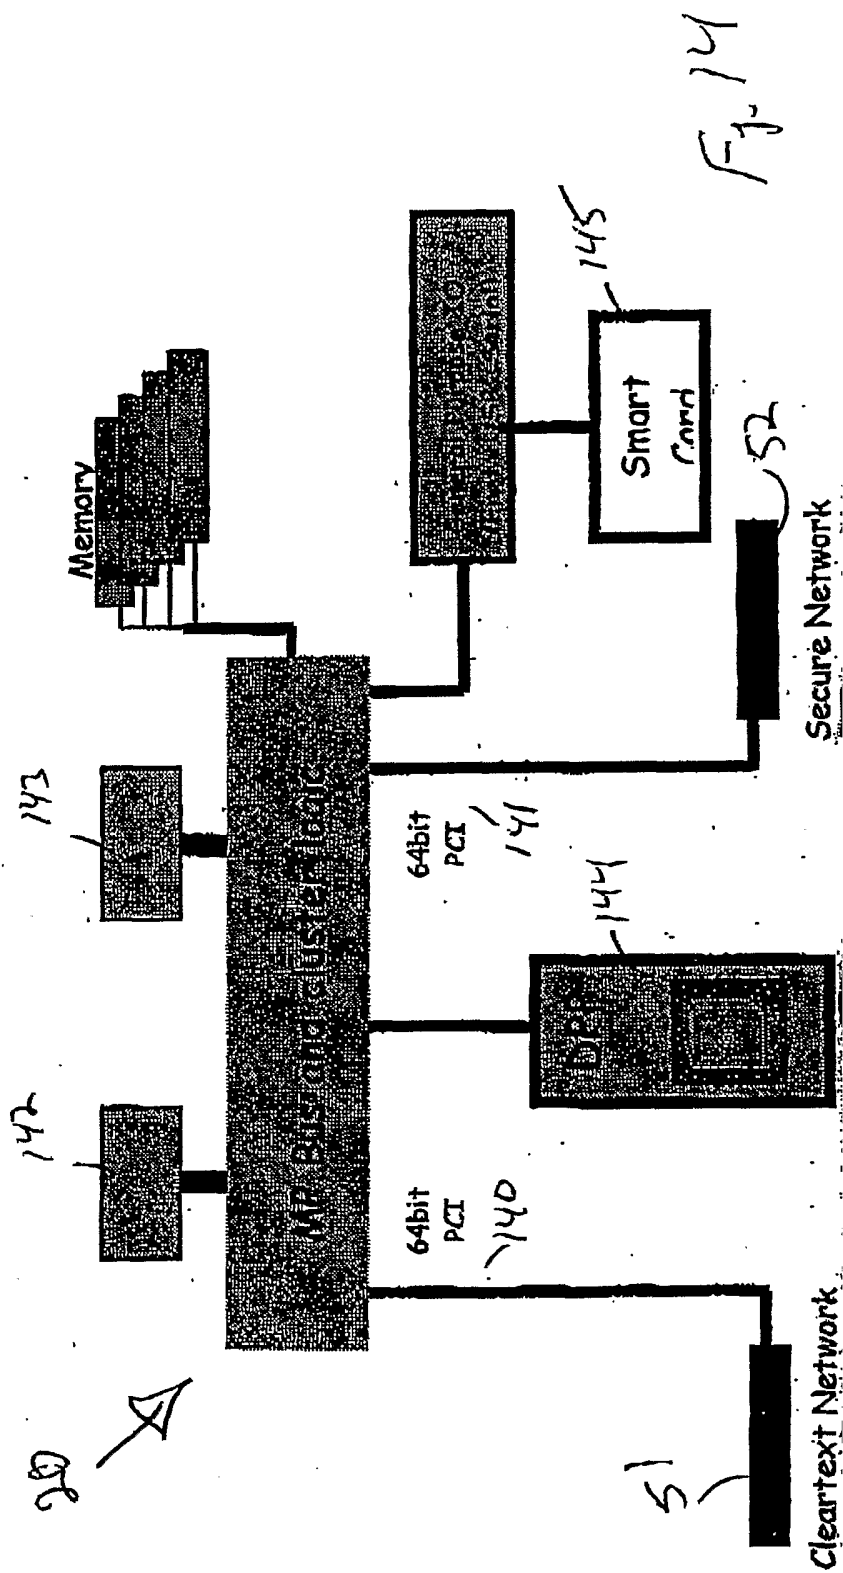
Fig. 11



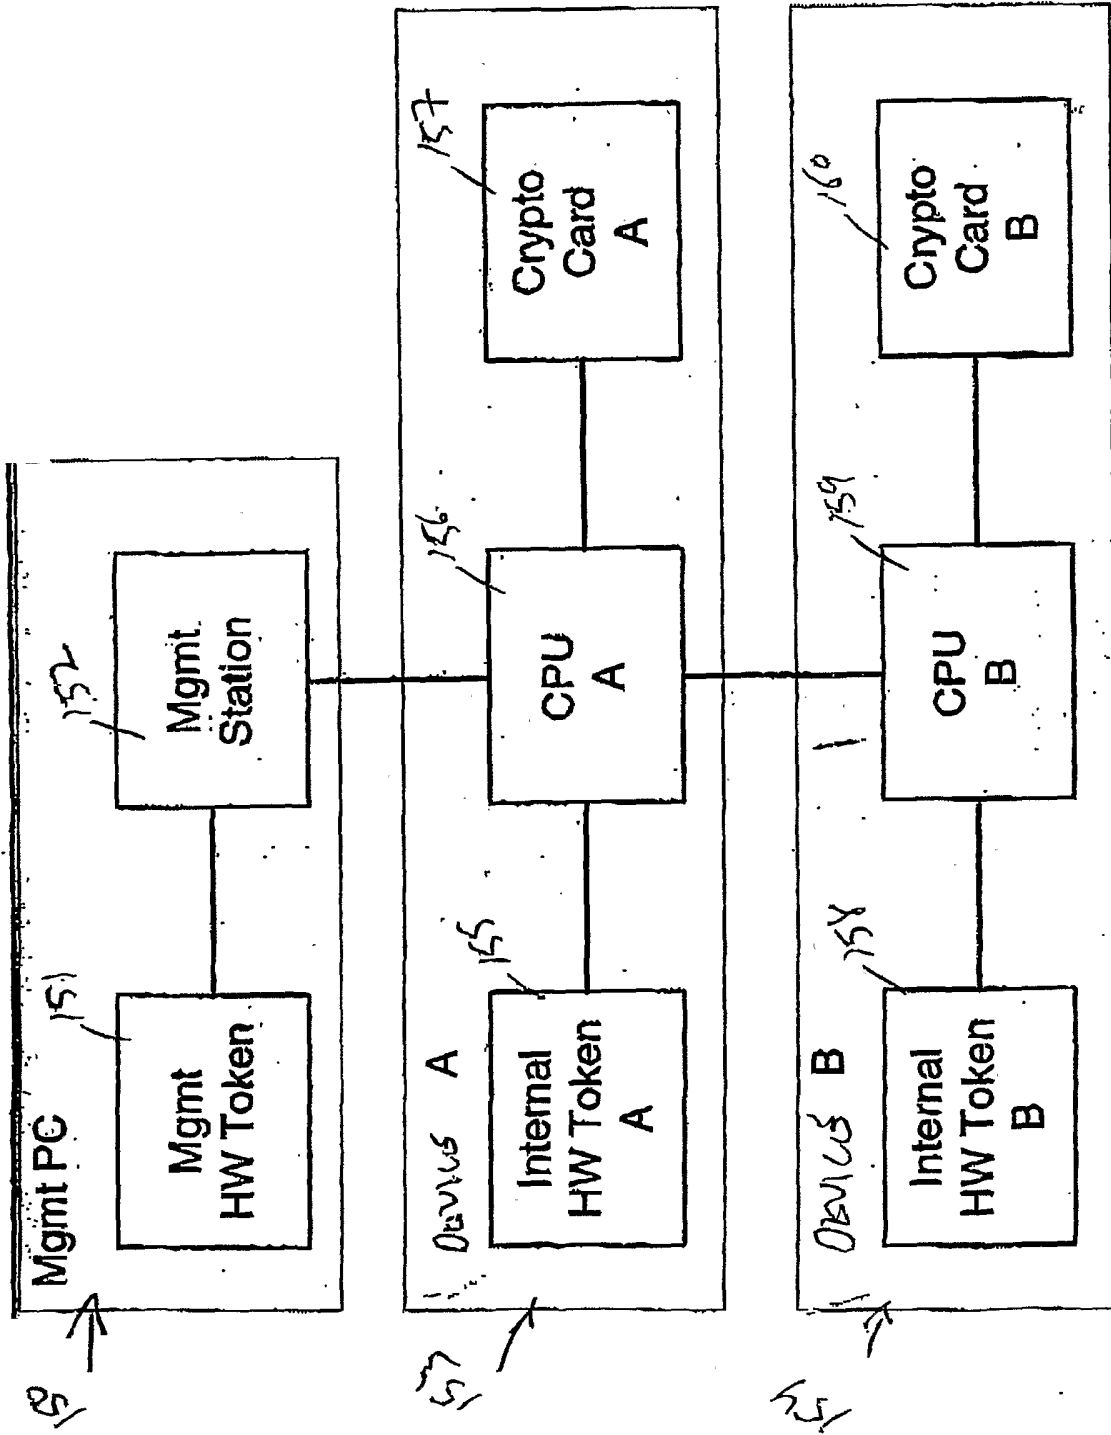


F-0-13

A 20



Figs 13



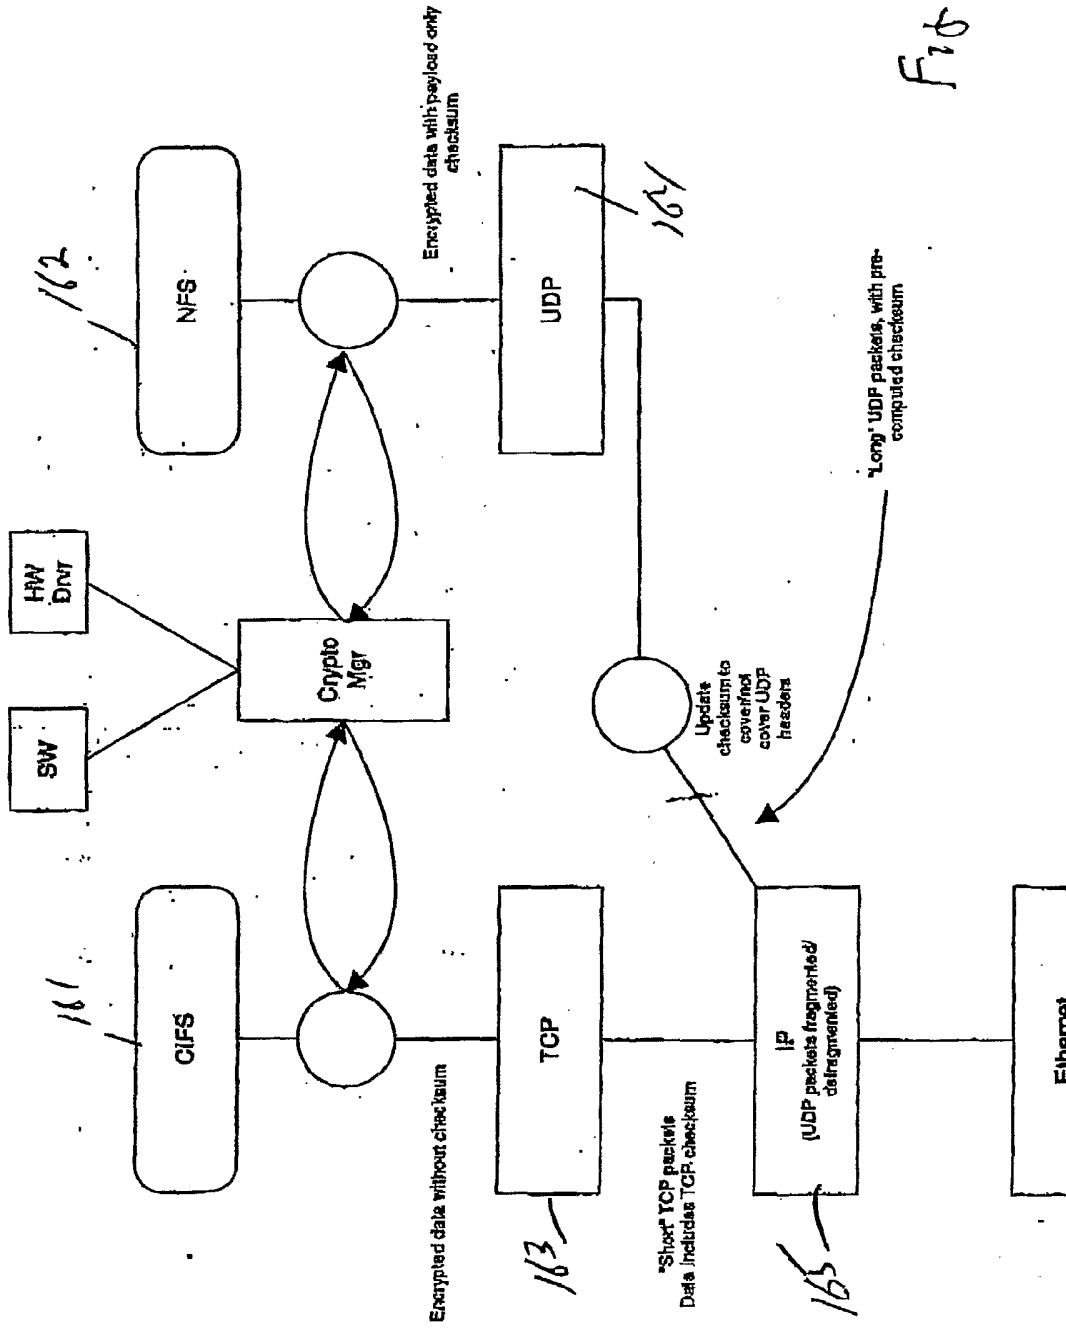
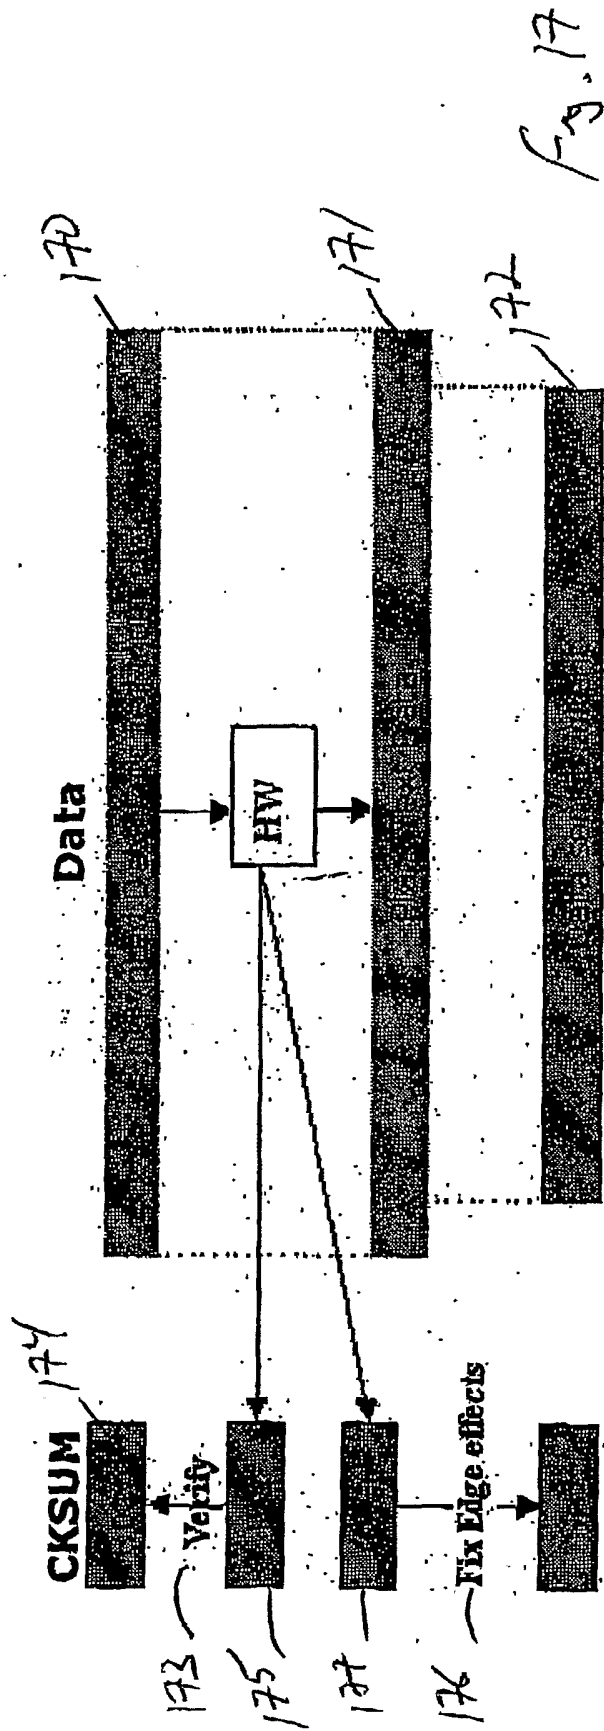


Fig. 16



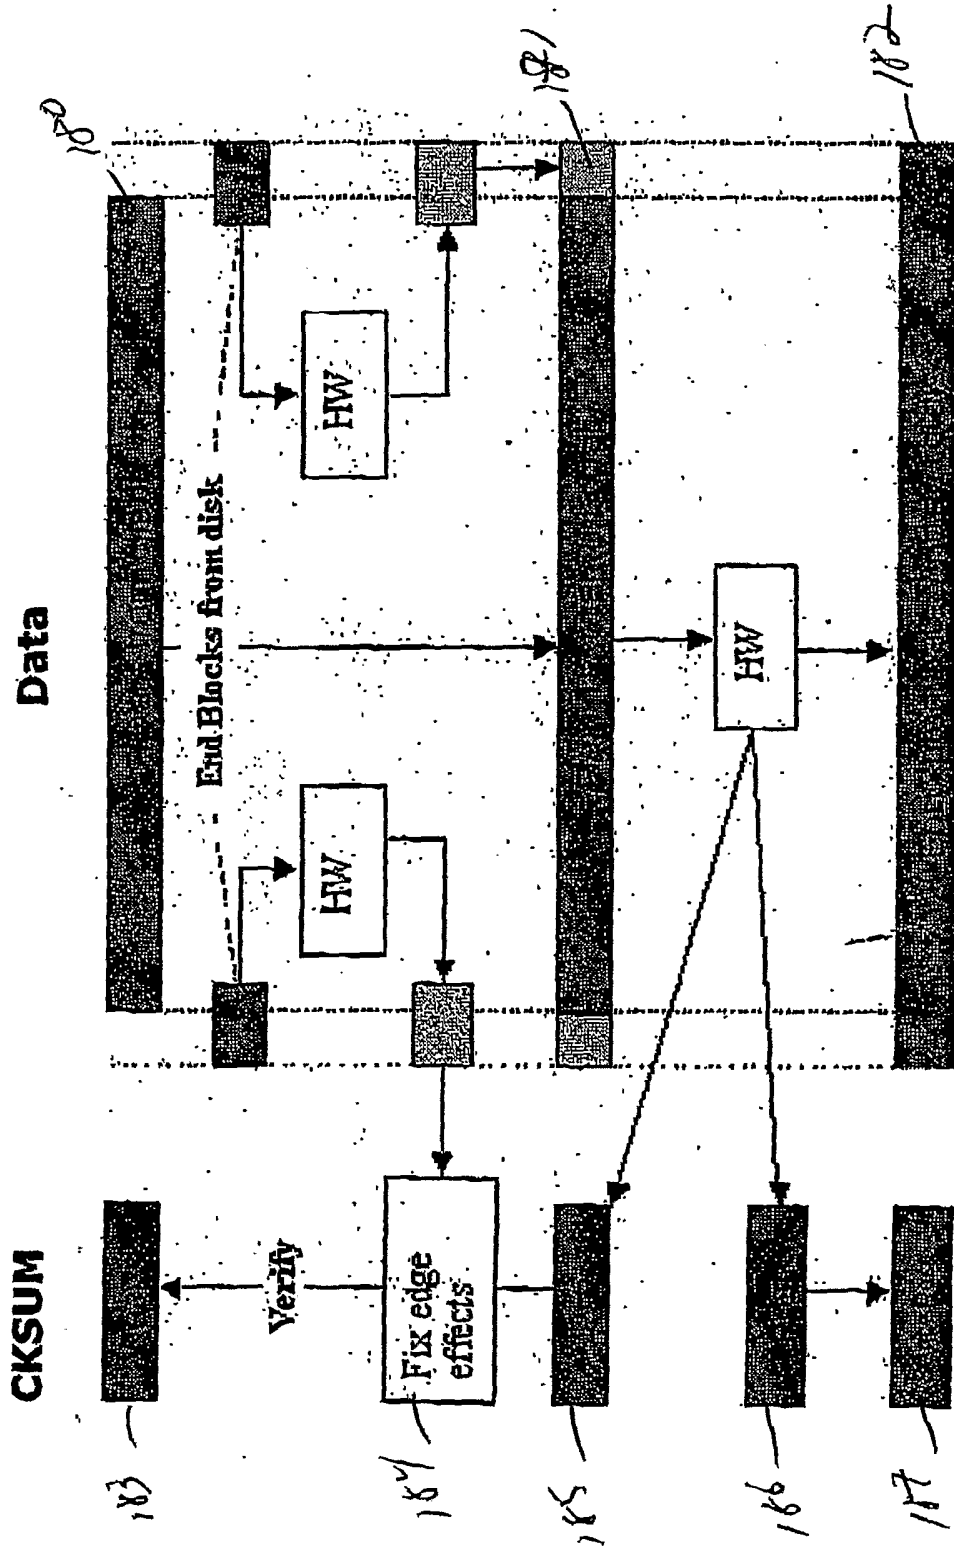


Fig. 18

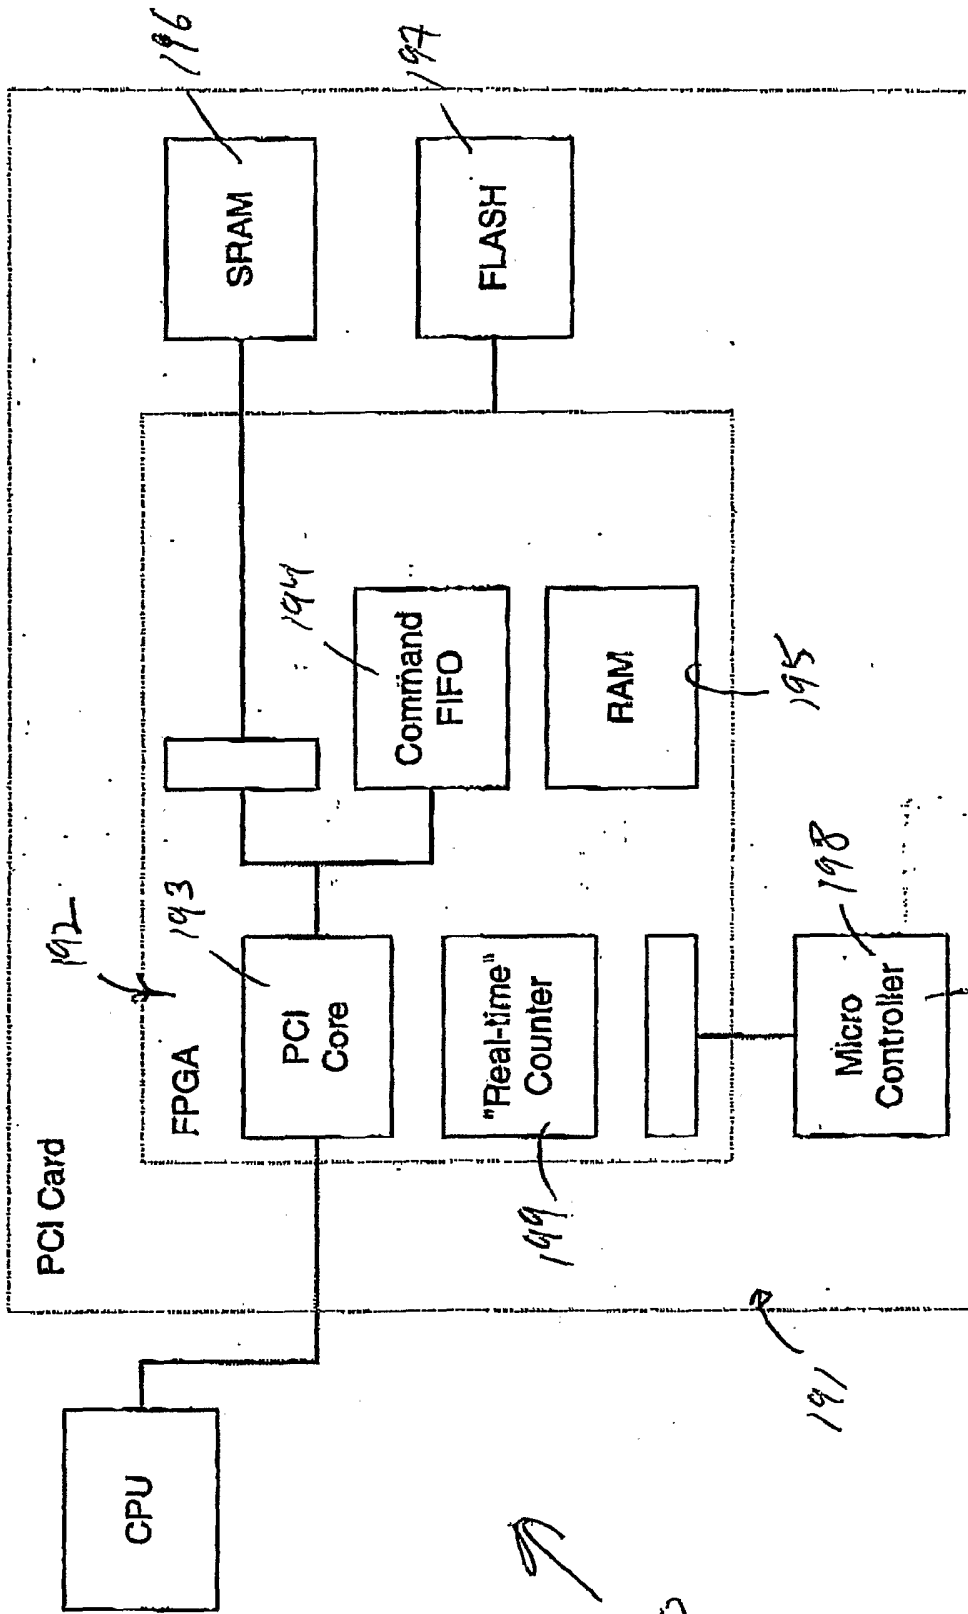


Fig. 19