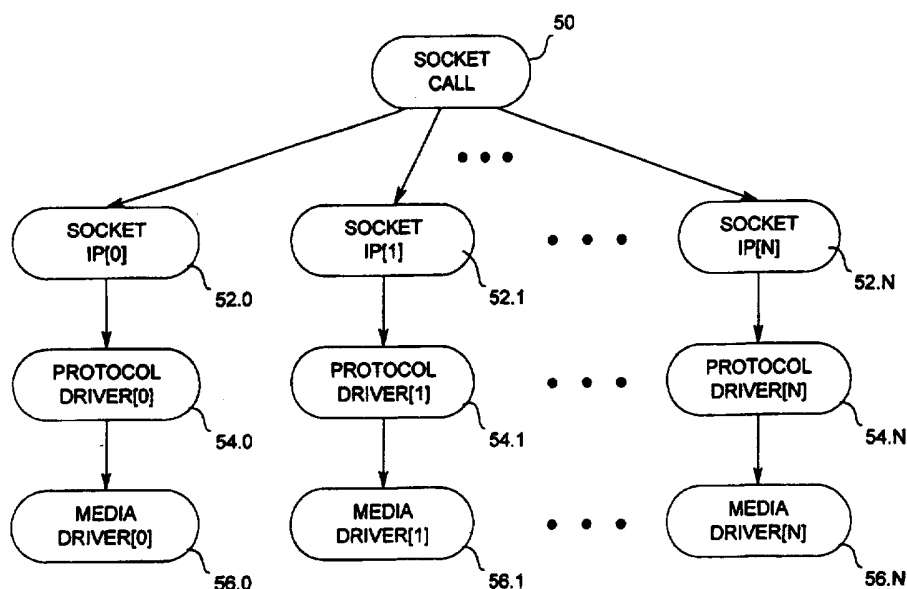




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F	A2	(11) International Publication Number: WO 97/29413 (43) International Publication Date: 14 August 1997 (14.08.97)
<p>(21) International Application Number: PCT/US96/05228</p> <p>(22) International Filing Date: 15 April 1996 (15.04.96)</p> <p>(30) Priority Data: 08/599,232 9 February 1996 (09.02.96) US</p> <p>(71) Applicant: SECURE COMPUTING CORPORATION [US/US]; 2675 Long Lake Road, Roseville, MN 55113 (US).</p> <p>(72) Inventors: GOODERUM, Mark, P.; 1115 Paul Parkway, No. 212, Blaine, MN 55434 (US). VU, Trinh, Q.; 6045 Lyndale Avenue South, No. 223, Minneapolis, MN 55419 (US). ANDREAS, Glenn; 171 - 83rd Avenue N.E., No. 304, Fridley, MN 55432 (US).</p> <p>(74) Agent: VIKSNINS, Ann, S.; Schwegman, Lundberg, Woessner & Kluth, P.O. Box 2938, Minneapolis, MN 55402 (US).</p>		<p>(81) Designated States: AU, CA, JP, European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).</p> <p>Published <i>Without international search report and to be republished upon receipt of that report.</i></p>

(54) Title: SYSTEM AND METHOD FOR ACHIEVING NETWORK SEPARATION



(57) Abstract

A system and method of achieving network separation within a computing system having a plurality of network interfaces. A plurality of burbs or regions is defined, wherein the plurality of burbs includes a first and a second burb and wherein each burb includes a protocol stack. Each of the plurality of network interfaces is assigned to one of the plurality of burbs and more than one network interface can be assigned to a particular burb. Processes are bound to specific burbs when they try to access that burb's protocol stack and communication between processes assigned to different burbs is restricted so that a communication between a process bound to one burb must pass through a proxy before being sent to a different burb.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AM	Armenia	GB	United Kingdom	MW	Malawi
AT	Austria	GE	Georgia	MX	Mexico
AU	Australia	GN	Guinea	NE	Niger
BB	Barbados	GR	Greece	NL	Netherlands
BE	Belgium	HU	Hungary	NO	Norway
BF	Burkina Faso	IE	Ireland	NZ	New Zealand
BG	Bulgaria	IT	Italy	PL	Poland
BJ	Benin	JP	Japan	PT	Portugal
BR	Brazil	KE	Kenya	RO	Romania
BY	Belarus	KG	Kyrgyzstan	RU	Russian Federation
CA	Canada	KP	Democratic People's Republic of Korea	SD	Sudan
CF	Central African Republic	KR	Republic of Korea	SE	Sweden
CG	Congo	KZ	Kazakhstan	SG	Singapore
CH	Switzerland	LI	Liechtenstein	SI	Slovenia
CI	Côte d'Ivoire	LK	Sri Lanka	SK	Slovakia
CM	Cameroon	LR	Liberia	SN	Senegal
CN	China	LT	Lithuania	SZ	Swaziland
CS	Czechoslovakia	LU	Luxembourg	TD	Chad
CZ	Czech Republic	LV	Latvia	TG	Togo
DE	Germany	MC	Monaco	TJ	Tajikistan
DK	Denmark	MD	Republic of Moldova	TT	Trinidad and Tobago
EE	Estonia	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	UG	Uganda
FI	Finland	MN	Mongolia	US	United States of America
FR	France	MR	Mauritania	UZ	Uzbekistan
GA	Gabon			VN	Viet Nam

SYSTEM AND METHOD FOR ACHIEVING NETWORK SEPARATION

5

Background of the Invention

Field of the Invention

The present invention relates to computer security, and more particularly, to an apparatus and method for providing increased computer security by
10 assigning processes to regions and restricting communication between regions.

Background Information

There has been an explosion in the growth of computer networks as organizations realize the benefits of networking their personal computers and workstations. Increasingly, these networks are falling prey to malicious
15 outsiders who hack into the network, reading and sometimes destroying sensitive information. Exposure to such attacks has increased as companies connect to outside systems such as the Internet.

To protect themselves from attacks by malicious outsiders, organizations are turning to mechanisms for increasing network security. One such
20 mechanism is described in "SYSTEM AND METHOD FOR PROVIDING SECURE INTERNETWORK SERVICES", U.S. Patent Application No. 08/322078 filed October 12, 1994 by Boebert et al. Boebert teaches that modifications can be made to the kernel of the operating system in order to add type enforcement protections to the operating system kernel. This protection
25 mechanism can be added to any other program by modifications to the program code made prior to compiling. It cannot, however, be used to add type enforcement protection to program code after that program code has been compiled.

Type enforcement is a file level mechanism for which provides con
30 What is needed is a way of adding computer security protection to compiled program code to increase secure operation of such code.

Summary of the Invention

The present invention is a system and method of achieving network separation within a computing system having a plurality of network

interfaces. A plurality of burbs or regions is defined, wherein the plurality of burbs includes a first and a second burb and wherein each burb includes a protocol stack. Each of the plurality of network interfaces is assigned to one of the plurality of burbs and more than one network interface can be assigned to a particular burb. Processes are bound to specific burbs when they try to access that burb's protocol stack and communication between processes assigned to different burbs is restricted.

According to another aspect of the present invention, a secure server is described which includes an operating system kernel, a plurality of network interfaces which communicate with the operating system kernel and a plurality of burbs, including a first and a second burb. Each network interface is assigned to one burb from the plurality of burbs and each burb includes its own protocol stack for handling communication across the network interfaces assigned to that burb. Each protocol stack is associated with only one burb. Finally, communication between a process bound to one burb must pass through a proxy before being sent to a different burb.

Brief Description of the Drawings

Fig. 1 is a representation of a system having an internal and external interface connected via two separate protocol stacks;

Figs. 2a-d are representations of communication protocols;

Fig. 3 is a representation of one form of interprocessor communication which can be used in a system having a plurality of separate protocol stacks;

Fig.4 is a more detailed representation of one embodiment of the form of interprocessor communication shown in Fig. 3;

Fig. 5 is an alternate embodiment of the system of Fig. 1, in which all communications between regions (burbs) pass through system space before being passed to another burb;

Fig. 6 is a flowchart illustrating the steps taken in securing compiled program code according to the present invention;

Fig. 7 is a flowchart illustrating the steps taken in applying network separation to compiled program code according to the present invention; and

Fig. 8 is a representation of a system built using the steps shown in Figs. 6 and 7.

Detailed Description of the Preferred Embodiments

5 In the following Detailed Description of the Preferred Embodiments, reference is made to the accompanying Drawings which form a part hereof, and in which are shown by way of illustration specific embodiments in which the invention may be practiced. It is to be understood that other embodiments may be utilized and structural changes may be made without
10 departing from the scope of the present invention.

Computer systems which use a single communications protocol stack to handle communication between an internal and an external network are widely in use. This is the communication model used, for instance, in BSD 4.4. The problem with such a system is that once a process receives privileges within
15 the system, it can use those privileges to access other network files. This can lead to a dangerous breach of network security. Two approaches can be used to beef up the security of such a system: type enforcement and network separation. Type enforcement adds an additional level of protection to the process of accessing files. Network separation divides a system into a set of independent
20 regions. Through network separation a malicious attacker who gains control of one of the regions is prevented from being able to compromise processes executing in other regions. Type enforcement and network separation will be described next.

Type Enforcement

25 In "SYSTEM AND METHOD FOR PROVIDING SECURE INTERNETWORK SERVICES", U.S. Patent Application No. 08/322078 filed October 12, 1994, Boebert et al. describe a way of extending type enforcement protection to a computer system having both an internal private network and an external public network. In one embodiment of such a system, a secure
30 computer is used to connect a private network having a plurality of workstations to a public network. A protocol package (such as TCP/IP) running on the secure computer implements a communications protocol used to communicate between

each workstation and the secure computer. A Local Cryptography function can be integrated into the protocol package in order to protect and authenticate traffic on the private network.

Program code running on the secure computer is used to
5 communicate through the private network to the workstation's protocol package. In one embodiment, the secure computer is an Intel Pentium-based machine running a hardened form of BSD386 Unix. A system based on a 90 MHz Pentium microprocessor with 32 megabytes of memory, 2 gigabytes of hard disk space, a DAT tape for backup and a CD-ROM for software loads has been found
10 to be adequate.

Likewise, program code running on the secure computer is used to communicate through a public network interface to a public network such as the Internet. In an Internet embodiment, the program code used to communicate with the Internet is part of a set of Internet protocols which communicate with
15 computers on the Internet through an Internet connection. In one embodiment, different protocols and cryptographic methods may be used when communicating with different entities on the Internet. In one embodiment, a tcp wrapper package operating in the Internet protocols is used to sit on the external, public network so that information about external probes can be logged. It is
20 most likely that the open nature of the public network will favor the use of public-key cryptography in this module.

As noted above, in one embodiment the secure computer is an Intel Pentium-based machine running a hardened form of Berkeley's BSD386 Unix. In that embodiment, BSD386 is hardened by adding a type enforcement
25 mechanism which restricts the access of processes to data. Type enforcement operates in conjunction with page access control bits in the virtual page translator of the Pentium to control access to objects stored in the memory of the secure computer. To accomplish this, system calls in the basic BSD386 kernel were modified so that type enforcement checks cannot be avoided. Certain other
30 system calls were either disabled or had certain options disabled.

The type enforcement controls are enforced by the kernel and cannot be circumvented by applications. Type enforcement is used to implement data flow structures called *Assured Pipelines*. Assured pipelines are made possible by the so-called "small process" model of computation used by Unix.

5 In this model, a computational task is divided up into small virtual units that run in parallel to each other. Unix provides a crude and loosely-controlled way of sharing data between processes. Type enforcement supplants this with the rigorously controlled, configurable structure of assured pipelines. It should be noted that Type Enforcement works best as a supplement to the normal Unix
10 permissions. That is, the Unix permissions are the first line of defense; when processes get past the Unix permissions, however, they run into the type enforcement checks.

In one embodiment, the secure computer has been configured under BSD386 to run in one of two states: administrative and operational. In the
15 administrative state all network connections are disabled and the Server will only accept commands from a properly authenticated System Administrator accessing the system from the hard-wired administrative terminal. This feature prevents anyone other than the System Administrator from altering the security databases in the secure computer.

20 In the operational state the network connections are enabled and the Server will execute only software which has been compiled and installed as executable by an assured party.

The two states are reflected in two separate kernels. The administrative kernel is not subject to type enforcement. Instead, it is network
25 isolated and accessible only to authorized personnel. This means that in administrative kernel mode, the secure computer cannot be seeded with malicious software by any but the people charged with system administration.

On the other hand, the operational kernel is subject to type enforcement. This means, for instance, that executable files stored in the
30 memory of the secure computer cannot be executed without explicit execution privileges. In one such embodiment, executable files cannot be give execution

privileges from within the operational kernel. Instead, the secure computer must enter administrative kernel to grant execution privileges. This prevents execution of malicious software posted to memory of the secure computer. Instead, only executables approved by operational administrators while in

5 administrative kernel mode ever become executable within operational kernel mode of the secure computer. In one such embodiment, administrative kernel can be entered only from either a manual interrupt of the boot process to boot the administrative kernel or by booting the secure computer from a floppy that has a pointer to the administrative kernel.

10 The flow of data between processes is limited to transfers through assured pipelines controlled, in one embodiment, by the access enforcement mechanism of the Intel Pentium processor. Virtual memory translation circuitry within the Pentium processor includes a mechanism for assigning access privileges to pages of virtual memory. This ensures that control is imposed on

15 every fetch from, or store to, the machine memory. In this way, the protection is made continuous. The Pentium access control mechanism enforces the following modes of access:

- 20 ● Read Only (R): Data values may be fetched from memory and used as inputs to operations, but may not be modified or used as program text.
- Read Execute (RE): Data values may be fetched from memory and used as inputs to operations, and may also be used as program text, but may not be modified.
- 25 ● Read Write (RW): Data values can be fetched from memory and used as inputs to operations, and may also be stored back in modified form.
- No Access: The data cannot be fetched from memory for any purpose, and it may not be modified.

30 These hardware-enforced accesses can be used to force data flowing from the internal private network to the Internet to go through a filter

process, without any possibility that the filter is bypassed or that filtered data is tampered with by possibly vulnerable software on the Internet side of the filter.

The access a process has to a data object via type enforcement is defined by an entry in a central, protected data structure called the *Domain*

- 5 *Definition Table* (DDT). A Domain name denotes an equivalence class of processes. Every process in execution has associated with it two Domain names which are used to control its interaction with object and with other Domains. The real Domain of a process is used to control Domain to Domain interactions and to grant or deny special, object-independent privileges. The effective
- 10 Domain of a process is used to control its access to objects. The real and effective Domains of a process will generally be identical; the circumstances in which they differ are described below.

- A Type name denotes an equivalence class of objects. Objects are, in general, the "base types" of BSD/386 Unix: files, directories, etc. There
- 15 are eight default subtypes: *file*, *directory*, *socket*, *fifo*, *device*, *port*, *executable*, and *gate*. The implied default subtype *pipe* is, in effect, untyped because no check is made on access to pipes. Type names consist of two parts and, in the preferred embodiment, are written in documentation and comments as creator:subtype. The creator field is the four-character name of the Domain
- 20 which created the object. The subtype field denotes the "class" of the object within that Domain. Subtype names are also four characters long and may contain any printable character except '*' or whitespace.

- Subtypes will not be shared; thus Mail: file means, in effect, "the files private to the Mail Domain." When objects are created they are
- 25 automatically assigned the appropriate default subtype. Objects which are to be shared between Domains must have their subtype changed from the default to an explicit subtype.

Subtypes can be assigned one of three ways:

- By having a default subtype assigned when the object is created
- 30 by the operational kernel.

- By having an explicit subtype assigned by the privileged *chtype* or *fchtype* syscalls. Thus a file which was to be shared between the Mail Domain and some other Domain would first be created as Mail:file and then changed to, e.g., Type Mail:Publ. If an
5 subtype is changed to a default subtype, then the object becomes private.
- By having a default or explicit subtype assigned administratively by the administrative kernel.

The default subtypes exec and gate are "static." The operational
10 kernel will not create any objects of those subtypes, change those subtypes into any other subtype, or change any other subtypes into a gate or exec.

The Domain/Type relationship is used to define the modes and consequences of accesses by processes to objects. The modes and consequences of accesses are defined by access attributes which are store in the DDT database.

15 The DDT database is "indexed" by three values:

- The effective Domain of the process requesting the access or action.
- The creator field of the object Type.
- The subtype field of the object Type.

20 The result of "indexing" is the retrieval of a set of access attributes. The term "attribute" is used instead of "mode" because some of the attributes define immediate side effects. The selection of attributes was governed by the following considerations.

- To constrain the modes of access which processes may exercise
25 on objects.
- To prevent the execution of any application software other than that which has been installed through the controlled administrative environment.
- To enable the spoofing of attackers so that the attack response
30 facilities can be used to trace them at the physical packet level.

This required a more sophisticated response to illegal accesses than just shutting down the offending process.

Gating permits a process to temporarily become a member of another Domain. The "home" or permanent Domain of the process is called its real Domain and the temporary or assumed Domain is called the effective Domain. Implicit gating is used when it is necessary to strictly control the manner in which the effective Domain's accesses are used. Implicit gating "ties" the temporary Domain change to a specific executable which has been subjected to extra scrutiny to insure that the effective Domain's accesses are used safely. The "tying" of the Domain change is done because the Domain change is a side effect of *execve*'ing a special executable: one whose subtype is *gate*. Implicit gating also allows Domain changes to be defined by changing the Type of an executable instead of inserting explicit calls into the source code.

Explicit gating is used when a looser control on the temporary Domain transition is appropriate, or when the "tying" of the gating to a specific executable would require excessive restructuring of existing software.

Domain changes are controlled by the DIT. The logical structure of the DIT is a table with an entry for each Domain. The logical structure of each entry is that of two pointers, one to a list of allowed real Domains and the other to a list of allowed effective Domains. Thus, if a process executed a *makedomain* or *changedomain*, the real Domain of the process selects the entry and the Domain given by the *domainname* argument must be on the list of allowed real Domains for the Domain change to happen. Likewise, if a process executes a *gate*, the Domain given in the *domainname* argument must be on the list of allowed effective Domains. Finally, if a process executes an *execve* of an executable whose subtype is *gate*, the *creator* Domain of that executable must appear on the list of allowed effective Domains.

Certain kernel syscalls are restricted to processes executing out of privileged Domains. In one embodiment two levels of checks are made. First, the normal BSD UNIX permissions are checked; if these permissions cause the

operation to fail, the system call returns the normal error code. If the UNIX permissions are adequate, the type enforcement (TE) privileges are checked next, (and thus in addition to the UNIX permissions).

The following BSD system calls have been modified to properly
5 implement type enforcement. The modified calls have been grouped into four groups for ease of explanation.

The first group of system calls that require modification are those that set or affect the identity and/or state of the computer. Two of these system calls affect the computer's internal time: *settimeofday* and *adjtime*. Both of these
10 system calls have been modified to require the <can_set_clock> privilege before the request will be honored. In the event of a privilege violation, the system call will raise an Alarm, will not honor the request, but will return success.

Other system calls which affect the computer's notion of self identity are *sethostname* and *sethostid*. Both of these system calls have been
15 modified to require the <is-startup> privilege before the request will be honored. In the event of a privilege violation, the system call will raise an Alarm, will not honor the request, and will return the EPERM error flag. The last system call affects the computers runtime status, *reboot*. The *reboot* system call has been modified to require the <admin-reboot> privilege before the request will be
20 honored. If the request is honored, the computer will boot to the admin kernel (single-user mode only with networking disabled). In the event of a privilege violation, the system call will raise an Alarm, will not honor the request, and will return the EPERM error flag.

The second group of system calls that require modification are
25 those that allow interaction with the computer's filesystem. The *open* system call has been modified to become the primary TE check. After performing the normal BSD UNIX permission checks, the TE check is performed. An Alarm is raised if the TE check returns null (no permissions), or if the caller asks for read but the <ddt_read> privilege is not set, or if the caller asks for write but the
30 <ddt_write> privilege is not set. The *creat* system call has been modified to set the new file's Type to <creator:file>. Additionally, the creation of a new file

implies a write operation on the directory, which in turn implies that the TE-modified *open* system call will be used to open the directory file, which in turn implies that TE can be used to control the success or failure of the *creat* system call. The *unlink* and *rename* system calls are modified in like manner. The

5 *unlink* system call requires the <ddt_destroy> privilege. The *rename* system call requires the <ddt_rename> privilege on the "from" file, and if the "to" file exists, it further requires the <ddt_destroy> privilege on the "to" file. In the event of a privilege violation, both the *unlink* and *rename* system calls will raise an Alarm, will not honor the request, but will return success. The *access* system call is

10 modified to require the <mode> privilege on the file pointed to by the path. In the event of a privilege violation, the *access* system call will raise an Alarm, will not honor the request, but will return success. The *chflags*, *fchflags* and *quotactl* system calls are modified in alike manners. All are modified to perform no functions. Attempts to call them will raise an Alarm, will not honor the request,

15 and will return EPERM. The *mknod* system call is modified to perform no function. Attempts to call it will raise an Alarm, will not honor the request, and will return EPERM.

The third group of system calls that require modification are those concerning process creation, maintenance and tracing. The *fork* system call has

20 been modified so that the child process inherits both the real and effective Domains of the parent process. The *execve* system call is modified to require the <ddt_exec> privilege on the file pointed to by the path before the request will be honored. The real and effective Domain of the process remain unchanged. In the event of a privilege violation, the system call will raise an Alarm, will not

25 honor the request, but will return success. The *ktrace*, *ptrace* and *profil* system calls are modified in alike manners. All are modified to perform no function. Attempts to call them will raise an Alarm, will not honor the request. The *ktrace* and *ptrace* system calls will return EPERM, whereas the *profil* system call will return EFAULT.

The *mprotect* system call is modified to perform no function. Attempts to call it will raise an Alarm, will not honor the request, and will return EPERM.

The fourth group of system calls that require modification are those that relate processes to user ids. The *setuid* and *seteuid* and *old.setreuid* system calls are modified in alike manners. All are modified to require the <suppress_su_alarm> privilege before the request will be honored. In the event of a privilege violation, the system call will raise an Alarm, will not honor the request, and will return success. The *acct* system call is modified to perform no function. Attempts to call it will raise an Alarm, will not honor the request, and will return EPERM. The *setlogin* system call is modified to require the <can_setlogin> privilege. In the event of a privilege violation, the access system call will raise an Alarm, will not honor the request, but will return success.

A final set of system calls consists of those that are removed entirely from the BSD UNIX kernel. This set of system calls includes: *obs_vtrace*, *nfssvc*, *asynch_daemon*, *getfh*, *shmsys*, *sfork*, *getdescriptor*, and *setdescriptor*.

Network Separation

The goal of network separation is to provide an operating system kernel with support for multiple networking protocol stacks. A system having such an operating system kernel is illustrated in Fig. 1. System 10 is split into separate regions, with a domain 16 and a protocol stack 12 assigned to each region. Each protocol stack (12.0, 12.1) has a fixed set of interfaces bound to it. For example, two or more ethernet drivers can be connected to, for instance, protocol stack 12.0.

A given socket will be bound to a single protocol stack 12 at creation time. Each protocol stack 12 will have its own independent set of data structures including routing information and protocol information. No data will pass between protocol stacks without going through proxy space and being sent back down another protocol stack by a proxy program 14. Proxy 14 acts as a go-

between, therefore, for transfers between domains 16.0 and 16.1. No user applications will have direct access to either network.

The embodiments discussed will only cover common networking support, such as the media layer drivers like PPP, ethernet, and SLIP and the TCP/IP suite of protocols. The BSD kernel supports other protocols such as CCITT/X.25 and XNS. These are not covered here, although it should be apparent that network separation can be extended to other protocols by dividing the protocol stacks associated with the protocol layers into multiple protocol stacks with the number of (and the name of) stacks being the same across all protocol suites.

The following terminology will be used throughout the document:

BSD The BSD/OS 2.0 Unix operating system as based upon the BSD 4.4 Lite distribution. In the case of the networking functionality the code is virtually identical across all 4.4 derived Unixes (NetBSD1.0, FreeBSD 2.0, and BSD/OS 2.0). The only significant differences are the hardware device drivers in terms of how they autoconfig during boot, which ones are present, and their internal code. Their interfaces to the rest of the networking code is effectively identical.

BSD kernel The BSD/OS 2.0 kernel.

Burb The aggregation of a protocol stack with all the processes that can access that stack. Processes that can access a particular protocol stack are said to be bound to that protocol stack.

NBURBS The number of protocol stacks or network interfaces to which all networking processes and related entities are bound to.

Kernel space The address and code space of the running kernel. This includes kernel data structures and functions internal to the kernel. Technically the kernel can access the memory of the current process as well, but "kernel space" generally means the memory (including code and data) that is private to the kernel and not accessible by any user process.

User space The address and code space of processes running that isn't kernel space. A running process will often execute inside of the kernel during system calls, but that's kernel space since the kernel *always* is running with the

context of the current process except during the few moments of the actual context switch - in a kernel without SMP support and kernel threads. In general user space refers to code and data allocated by a code that is loaded from an executable and is available to that process, not counting the kernel private code data structures.

Protocol stack The set of data structures and logical entities associated with the networking interfaces. This includes sockets, protocol drivers, and the media device drivers.

Link level & hardware drivers These are the (almost in some cases) bottom layer drivers that talk a particular physical and/or link level protocol (ethernet, PPP, SLIP). They may be a hardware driver, in the case of most ethernet drivers, or they may sit on top of yet other drivers, such as PPP on a TTY on the COM port driver or even a parallel port Ethernet driver on top of the LPT port driver. Most of these drivers have two layers, a generalized layer that handles the common parts of the link level protocol (ethernet, ppp, etc.) and the hardware specific driver.

One of the goals of a secure operating system kernel is to allow dividing the network interfaces into distinct regions so that there is assurance that packets are never quietly passed through the kernel between those regions. In one embodiment, the following rules must be met to ensure secure control of packet transfer between regions:

- A single user process can only send and receive information (packets) from one region at a time.
- Once a process is bound to a region, it can only access that region.
- Incoming packets can only go to processes that are in the region associated with the interface the packet arrived on.
- Any data passing through the computing system to a different region must come into user process and be handed off to a different process that has access to the other region, a proxy program, to be sent out again.

The reasons for these design requirements are fairly simple. The goal is that any information passing through the computer between different regions has to be passed through a predefined, or assured, pipeline.

There are three ways of achieving this goal:

- 5 **1) Packet Filtering** Packet filtering would be done using conventional approaches similar to current firewalls and screening routers. This could be enhanced with additional filters based on interface rather than address to prevent certain types of spoofing.
- 10 **2) Packet Separation** A given message, incoming or outgoing, would be assigned a type based on the interface it arrived on or the socket it was sent from. The packet would then be thrown out at the top or bottom level if that type didn't match the type of the interface it was being sent on.
- 15 **3) Separate Protocol Stacks** The protocol stacks would be separated into multiple instances with a given interface or socket existing on only one.

15 In arriving at our current approach, the packet filtering solution was quickly discarded. Although you could eliminate several kinds of spoofing, many of them would still remain. For filtering to work at all, IP packet forwarding must be enabled, but then you are vulnerable to all sorts of nasty things sneaking through the computer system using the normal sorts of attacks
20 with various forms of spoofing, piggybacking, and just plain misconfiguration of complex filtering expressions.

 The packet typing approach had merit, but required extensive code changes as all of the various layers would have to be rewritten to handle additional fields in all of the various structures (mbufs, sockets, etc.). The other
25 problem with the packet typing approach was how to deal with kernel internally generated messages like ICMP messages.

 The separate protocol stack approach was selected because it gives a very clear split of the regions making assurance easier. It also requires surprisingly few code changes. Rather than change the existing data structures,
30 they are simply replicated as needed for the various regions. The initial reference has to be fixed, but most references are done by cached pointers within

the various structures, so many of the functions and layers of the protocol stack require no changes at all. For instance, the hardware drivers require absolutely no changes. The socket code requires only a minor change during the socket instantiation, subsequent references done using the already fixed pointer to the particular protocol driver. This also gives a performance win because there are basically no lookups done dynamically, only an extra level of indirection during initialization, but not subsequently.

The general structure of the design chosen involves duplicating all of the protocol stacks where each stack is independent of the others. The routing to a given stack is done at the very top or bottom of the dataflow so that a given packet, piece of data, control message, etc. is bound to a particular stack at creation.

A name was needed for each of these protocol stacks, and the processes and related entities bound to them, the name that was chosen was *burb*.

Second, the decision whether to replicate and subdivide into separate burbs is made on a protocol family by protocol family basis. Generally all external supported networking protocol families will be divided. The internal Unix domain socket family won't be replicated. This is important because the Unix Domain sockets (AF_UNIX) will continue to exist in one common region for the entire system. This feature becomes important for the pipeline structure given later.

To identify the network/socket access a process is given a burb id that says which burb it is in. In one embodiment, the burb ID is an ordinal number.

The general picture of the protocol stacks in a normal BSD kernel is shown in Fig. 2a. An expanded view of this picture to include the various IP protocol drivers is shown in Fig. 2b.

When a socket is actually created, it is explicitly, during the initial *socket()* call, bound to one of the protocol drivers via a pair of pointers in the socket structure, the *so_proto* member (a pointer to a *struct protosw*), and the

so_pcb member. The *protosw* structure is unique per protocol driver, but there is only one per protocol driver. The *so_pcb* pointer is an opaque pointer to a socket specific structure that has the protocol specific information for a socket (such as the TCP state information for a TCP socket).

5 At the bottom level, the generic media drivers 24 (ethernet, ppp, slip) put incoming packets on a per protocol family (IP, XNS, CCITT, etc.) queue, *ipintrq* in the case of IP.

The picture then for a created socket becomes something like the system of Fig. 2c.

10 So, the fundamental concept of the kernel portion of the design is that the protocol stacks are replicated N times (with N being a fixed small constant, typically 2-10), with all protocol level and other shared data structures being replicated. In practice, each burb/protocol stack has a name, that name maps onto a number, starting at 0, that is used as an index. Each static structure
15 is converted into an array. This also makes it easy to change the number of burbs supported. Fig. 2d illustrates a system having two burbs and limited to just IP. A generic system having N burbs is shown in Fig. 3. A specific TCP/IP system having N burbs is shown in Fig. 4.

20 Once the socket is created then the picture becomes much simpler again, because the individual data structures are then bound explicitly and no lookup or search stage occurs again. All activity then occurs within one of the burbs (e.g., *TCP[0]*).

The advantage of duplicating the protocol stacks in their entirety is that the number of data structures to be duplicated, and the points that those
25 are referenced, is manageable, much more so that making a given datum have a type and/or domain and having to "route" within the kernel based on that at every layer. As mentioned before, the basic approach is that all shared data structures are replicated N times, 1 for each burg. These data structures are a reasonably finite list. They include:

30 **Protocol Domain List and Tables (domains, protosw)** - When a socket is created, a family and a type is specified, such as socket (*AF_INET*,

SOCK_STREAM). The socket creation routine first looks up the top level protocol in the family list, ala: AF_INET, which returns a 2nd level table of protosw structures that has the individual protocol drivers. The 2nd level table is in the form of XXX domain, such as inetdomain. The first level list is called

5 "domains".

The design replicates only the tables (when that family supports separate burbs, such as IP, route). In the case of things like the AF_UNIX domain, each list will point at the same table. The lookup will then return a protosw structure unique to the instance of the protocol driver, such as

10 inetsw[proto][burb], instead of inetsw[proto]. The routing structure will be replicated into routesw [NBURBS].

Protocol Interrupt Queues (ipintrq) - External networking protocols each have an input queue for incoming packets. Each queue is a simple mbuf list. So queues for protocols will be replicated, becoming things like ipintrq [NBURBS]

15 instead of just ipintrq.

IP Fragmentation Queues (ipq) - Similar to the IP interrupt queue, there is a reassembly queue, ipq, to reassemble fragmented IP packets. All network interfaces have a maximum transmission unit (MTU). Outgoing packets larger than the MTU are fragmented by the IP layer into smaller packets for

20 transmission. When the fragmented packets arrive at the destination, they are placed on a reassembly queue waiting for reassembly. This queue is also replicated into ipq [NBURBS].

Interfaces (ifnet) - Each interface, including the loopback will be bound at boot time to a given burb by a one time system call. A burb field will be added to the

25 ifnet structure, ifnet.if_burb, which is unique per interface so that it knows which burb it's in and which interrupt queues to send data to. The binding will also be one shot, so that once an interface is bound, it can't be rebound to a different burb.

Other data structures are already unique per interface. The other

30 item to be replicated is the list of interfaces. There will be one list per interface.

The interface binding call will set the burb of the interface and tell it which interface list to put itself on during initialization.

Loopback Interface (loif) - The loopback interface IP address will have to be unique per interface. In other words, the recommended loopback address

- 5 127.0.0.1 cannot be shared between the interfaces. The loopback interface data structure will be replicated, loif [NBURBS]. Loopback addresses will be subnetted, with netmask 255.255.0.0, to give each interface a unique IP address, 127.burb.0.1.

Protocol Control Blocks (tcb, udb, rawcb) - Each protocol maintains its own

- 10 list of protocol control blocks, which hold various pieces of information required for the socket operations. These lists will be replicated and maintained on a burb basis: tcb [NBURBS] for TCP, udb [NBURBS] for UDP, and rawcb [NBURBS] for routing sockets.

Routing Tables (rt_tables) - There is a single master routing table, rt_tables

- 15 [family], that is a linked list of routes. This table will be separated into one table per burb, rt_tables [family][NBURBS]. There will be no master routing table.

These routing tables will be used as lookup tables so that the incoming end of a proxy knows which outgoing proxy to connect to. For outgoing packets, the search always starts with the local table. If a route cannot

20 be found, then the system searches other tables. If a route is found on a non-local table, then the networking code will hand the packets off to the proxy processes.

ARP Tables and Queues (linfo_arp, arpintrq) - There are no data structures at the common link level (ethernet, ppp, slip) except the ARP table and the ARP

- 25 interrupt queue, linfo_arp and arpintrq. This again is a list of interfaces with just the information needed for ARP and a queue for incoming ARP packets. Both of these data structures will be replicated, linfo_arp [NBURBS] and arpintrq [NBURBS].

Sockets (struct socket) - The data structures for a socket are created

- 30 dynamically for each socket, so no fixed replication needs to be done. The alternation of the protocol driver lookup code means that a socket will

automatically get the pointer to the protocol driver in the correct burb without any changes to the socket code except the domain list search routines. A burb field will be added to the socket structure, socket.so_burb, so that it knows which burb it's in.

- 5 There are some other issues. A message doesn't normally have an intrinsic burb, but it will be added as a debugging aid during development. Each major data structure that is burb dependent (the input queues, the protocols structures, etc.) will always have a field giving its burb. So at various points, these fields can be compared against during development as a debugging aid.
- 10 Critical hand off stages (like socket initialization) will have sanity checks that are done all the time because the overhead is negligible.

 The kernel is implemented to respond to connection requests that are not intended for the computer system itself. For instance, when an internal client requests a connection to an external Internet host, the kernel will answer

15 the request as if it were the external host. Meanwhile, it will attempt to set up the connection with the external host. This process occurs transparently to the internal client. Such support allows the secure computer to answer in a secure manner connections intended for outside boxes.

 Each burb will have its own routing tables. The general algorithm

20 for each incoming packet will be, see if it's to one of the local addresses, if so accept it (as this is standard behavior). If not, route it within the burb. If that fails, then see if there is a route in another burb (using other burbs' routing tables). If not, then toss the packet. If so, accept the packet and pass it up the stack, which eventually should arrive at the proxies.

25 The existing TCP and UDP drivers will happily accept the packet and pass it onto a socket if it exists. If it's to an already established TCP connection or a UDP socket, it will get there normally. If it's a new incoming TCP connection that is accepted, it will create the socket, and fill in the local address from the incoming packet. The process doing the accept () can call

30 getsockname () to get that address. In other words, the UDP and TCP drivers trust the lower layer IP driver that the packet is for the machine, so no change to

those layers need to be done at all to support this, only the secondary routing lookup in the IP layer.

This functionality means a couple of things. For this to work, there still have to be unique IP addresses all around, so no hacks are done to support hidden IP addresses that conflict with external ones. Also, there can only be one default route on the machine still. If there were one per burb, the burb routing lookup wouldn't know where to send it. This isn't hard. Generally the default route will be for one of the outside burbs. Other burbs will simply have to have complete routing for the internal network. This isn't a problem as it's really already the case. A host with a complex internal network already has to have proper routing for all of those networks if it uses a default route.

This also means there can be no duplicate routes in the replicated routing tables. If there were duplicate routes, the routing lookup would always return the first route found.

15 Network Separation Using Type Enforcement

The degree to which network separation protects a network from malicious attack is enhanced through the use of type enforcement.

As noted above, in a type enforcement scheme processes and files are assigned to a domain and subtype, domain:subtype. Process-to-file, or domain-to-domain:subtype, access is controlled by the Domain Definition Table (DDT), that specifies the permissions:

25 create - ddt_create
 write - ddt_write
 read - ddt_read
 destroy - ddt_destroy
 execute - ddt_execute
 rename - ddt_rename

Subtypes allow a finer division of file types (file = file, dir = directory, sock = socket, exec = executable, etc.).

Interactions between the processes, or domain-to-domain, such as signals, are enforced by the Domain Interaction Table (DIT), that specifies the allowed signals (sHUP, sABRT, sJob, etc.)

An instantiated domain is one that has several instances with
 5 identical access rights to itself and other domains. Take two instantiated domains, www0 and www1. They both have identical access to all other domains and subtypes. They also have equivalent access to their instances and no access to other instances. In the domain specification language, these domains are specified with a 'X' as the last character, where X represents the
 10 burb id. So if you say:

wwwX:

has read access to www#:conf

has write access to Slog:sock

the DDT/DIT is created like:

15 www0:

has ddt_read access to www0:conf

has ddt_write access to Slog:sock

www1:

has ddt_read access to www1:conf

20 has ddt_write access to Slog:sock

If another domain is given access to a general instantiated type, then it gets equal access to all instances, for instance:

Admn:

has write, read, create access to www\#:conf

25 the DDT/DIT is created like:

Admn:

has ddt_write, ddt_read, ddt_create access to www0:conf

has ddt_write, ddt_read, ddt_create access to www1:conf

30 There are a set of special domains that are used to actually grant access to the networks themselves. These domains are in the form of

Protocol:port. Protocol is a domain name corresponding to the protocol type.

Current domains and what type of socket they map to include:

	<u>Domain Name</u>	<u>Network Domain</u>	<u>Socket Family/Protocol</u>
5	tcpX	TCP	AF_INET/SOCK_STREAM
	udpX	UDP	AF_INET/SOCK_DGRAM
	ipoX	Raw IP	AF_INET/SOCK_RAW
	rteX	Route	PF_ROUTE/SOCK_RAW

10

For sockets that have port numbers, such as tcp and udp, the subtype is an asciiish number that is the prot that binding is allowed to. For instance, the ftp0 domain would have access to ftp 0:0020 and ftp 0:0021 to give it access to the ftp control and data ports (20 and 21).

15

The access given corresponds to the directionality of communication. There must be ddt_create privilege to create the socket at all. For datagram sockets, ddt_read gives the ability to read () or recvmsg (), and ddt_write gives access to write () and sendmsg ().

20 With stream sockets, ddt_read gives access to bind (), listen (), and accept () so that the process can receive an incoming connection. ddt_write gives access to bind () and connect () so that the process can initiate an outgoing connection.

The following subtypes give access to random and reserved ports:

- 25 rall - bind to all reserved ports explicitly
- rany - bind to any arbitrary reserved ports
- nall - bind to all non-reserved ports explicitly
- nany - bind to any arbitrary non-reserved ports

30 For portless types, the subtype grants access to a specific type of socket. The types currently include 'icmp', only valid under the IP domain (ipoX:icmp), which gives access to ICMP sockets. Some more network subtypes:

sock - flag access to a network type

conf - access to network configuration capability

Subtype 'sock' is used during the socket creation so the kernel can verify access to a given network domain. The port isn't known at this point, and an exhaustive search of the DIT would be expensive, so any domain that has
5 access to any networking domain has to be given access to the 'sock' subtype as well.

Network configuration programs such as ifconfig, nss, and sysctl must have access to "network_domain:conf". For example: tcpX:conf allows the Unix commands ifconfig to configure network interfaces and sysctl to set
10 networking configurables.

A network interface is bound to a burb at boot time. A given socket is bound to a particular burb at the socket creation time. The rules for this are pretty simple.

An example of a type enforced, network separated system 70 is shown in Fig. 5. In Fig. 5, ftp0 and ftp 1 are two instantiated domains (72.0 and 72.1, respectively). Each instantiated domain has access to its own protocol stack (74.0 and 74.1). Transfers from one domain to another are made through a ftp proxy 75 via calls to kernel 76.

Processes running in an "instantiated domain" such as ftp0,
20 intrinsically have access to one and only one burb, the burb that matches the numeric part of the domain name. In these cases, the system just binds that socket to the appropriate burb. It should be noted that a process is not actually bound to a burb until it tries to make a direct or indirect socket call to that burb.

Programs running within other domains either have no access or
25 access to all burbs. In these cases, the caller must specify which burb they want. If they don't, it's an error. If they do specify a burb and they don't have access privileges to that burb, it's also an error.

This specification is implemented via a new system call, socketburb (). Basically, socketburb () replaces the old socket () system call. It
30 performs the same function and takes one additional argument, burbid. The old socket () still exists, but its gut has been replaced by a call to socketburb ().

```
int socketburg (int domain, int type, int protocol, unsigned long burbid);
```

Ported applications and new programs will use the new `socketburg ()`. Existing programs, in instantiated domains, using `socket ()` wouldn't have to be changed and still work. Burbification of `socket ()` calls to `socketburg ()` should use the library call `domain_to_burb ()` to convert and pass the burbid argument:

```
sock1 = socket (AF_INET, SOCK_STREAM, 0);
```

```
sock2 = socket (AF_INET, SOCK_RAW, IPPROTO_ICMP);
```

would become:

10

```
int domain_to_burb (long domain);
```

```
sock1 = socketburg (AF_INET, SOCK_STREAM, 0, domain_to_burb  
(domain));
```

```
sock2 = socketburg (AF_INET, SOCK_RAW, IPPROTO_ICMP,  
15 domain_to_burb (domain));
```

Interfaces are bound to a particular burb. This is implemented by a few new `socket ioctl(s)`s which set and get the burb id of an interface.

Access to a socket for controlling interfaces and other shared parameters is controlled by access to the various conf subtypes. Parameters specific to a protocol, like TCP tunables, are controlled by access to the protocol specific type, such as 'tcp0:conf'. IP layer things like the IP address and burbness of an interface are controlled by the types like 'ip0:conf'.

The following interfaces have been modified or added to support type enforcement, burbification, and transparent proxies. The objectives of these calls have not changed, although access to these calls has been made more restricted. Some new error codes, `errno`, have been added in the case of type enforcement failures. Only the modifications are described below.

socket () - is used by a process to create a socket that is bound to a particular burb. If `socket ()` is called, the socket is either bound implicitly to the burb id that matches the number of the instantiated domain, or fails.

30

socketburb () - a new system call and it is used to create a socket in a given burb. The burb filed is passed as an additional argument to the socketburb () call:

```

int socket (int family, int type, int protocol);
5  int socketburb (int family, int type, int protocol, unsigned long burb);
    int domain_to_burb (long domain);
    sock1 = socket (AF_INET, SOCK_STREAM, 0);
    sock2 = socketburb (AF_INET, SOCK_RAW, IPPROTO_ICMP,
                        domain_to_burb (domain));

```

10 These new error codes, errno, are returned by socket () and socketburb ():

 [EDBOM] returned and an audit generated if the calling process is not in a burb-bound domain.

 [EPERM] returned and an audit generated if the calling process
15 doesn't have ddt_create privilege.

bind () - assigns a local address and port number to a socket. The bind () call has been extended to allow a process to set the local address arbitrarily, it needs ddt_rename access to "ipoX:spuf". Currently, no domains have this spoofing capability, i.e., no domains have ddt_rename access to "ipoX:spuf". The calling

20 process also must have the following privileges to successfully bind:

```

ddt_read      - tcpX:port
ddt_read      - udpX:port
ddt_read      - ipoX:port, for raw IP
has_rootness  - reserved ports (0-1023)

```

25 [EADDRNOTAVAIL] returned and an audit generated if the calling process doesn't have ddt_rename privilege.

 [EPERM] returned and an audit generated if the calling process doesn't have ddt_read or has_rootness privilege.

connect () - initiates a connection from a socket to a specified address and port
30 number. The calling process must have these privileges to the listed domains: subtypes to successfully connect.

ddt_write - tcpX:port

ddt_write - udpX:port

[EPERM] returned and an audit generated if the calling process doesn't have ddt_write privilege.

- 5 The ioctl () call is used to manipulate the underlying device parameters. Several new socket ioctl () commands have been added. Most will work with a "struct ifreq" argument and be handled similar to the other set/get ioctls of an interface. If an ordinal value is being set or retrieved such as the burb id, then it will be passed in the ifr_metric field.

- 10 The following ioctl () commands now require is_startup or is_admin privilege to execute: SIOCSIFFLAGS, SIOCSIFMETRIC, SIOCADDMULTI, and SIOCDELMULTI. EPERM is returned and an audit is generated if the calling process doesn't have the correct privilege. These new ioctl () commands are added to support burbness in the secure computer:

- 15 **SIOCGIFBURB/SIOCSIFBURB** - get and set the burb id of a network interface. The burb id is passed in the ifr_metric field. To set the interface's burb id, the calling process must have is_startup or is_admin privilege. A sample call would look like:

```

20      int ioctl (int socket_descriptor, unsigned long command, char *argp);
      struct ifreq ifr;
      int burb, so;
      ifr.ifr_metric = burb;
      if (ioctl (so, SIOCSIFBURB, (caddr_t) & ifr) < 0) perror ("ioctl
          (SIOCSIFBURB)");
25      if (ioctl (so, SIOCGIFBURB, (caddr_t) & ifr) < 0) perror ("ioctl
          (SIOCGIFBURB)");

```

[EPERM] returned and an audit generated if the calling process doesn't have is_startup or is_admin privilege.

- 30 [EINVAL] returned and an audit generated if the calling process doesn't have access to subtype ":conf".

SIOCGSOCKBURB - get the burb that a socket is bound to. This is mainly a debugging aid. The burb id is passed in the `ifr_metric` field.

SIOCGMSGBURB - get the burb of the pending message. It should always be the same as the burb of the socket. This is mainly a debugging aid. The burb id

5 is passed in the `ifr_metric` field. A sample call would look like:

```

    struct ifreq ifr;
    int burb, so;
    if (ioctl (so, SIOCMSGBURB, (caddr_t) & ifr) < 0 perror ("ioctl
        (SIOCMSGBURB)");
10     burb = ifr.ifr_metric;
        [ENOENT] returned and an audit generated if the socket cannot
        be read.
```

SIOCGMSGIFNAME - get the name of the interface that the pending message arrived on. It is used as a debugging aid or to filter on the basis of the specific

15 interface of the message. The interface name is returned in the `ifr_name` field. A sample call would look like:

```

    struct ifreq ifr;
    int so;
    if (ioctl (so, SIOCGMSGIFNAME, (caddr_t) & ifr) < 0 perror ("ioctl
20     (SIOCGMSGIFNAME)");
    printf ("interface : %s", ifr.ifr_name);
        [ENOENT] returned and an audit generated if the socket cannot
        be read.
```

SIOGSUBURBRT - returns a burb id for a given IP addresses. The kernel

25 searches each routing table and returns the first route found. The IP address is passed in the `ifr_addr` field. The burb id is returned in the `ifr_metric` field on success. Since there can be more than one burb, this command is used by an incoming proxy to find which outgoing proxy to connect to. A sample call would look like:

```

30     struct ifreq ifr;
    int so;
```

```

    ifr.ifr_addr = tempaddr;
    if (ioctl (so, SIOCGBURBRT, (caddr_t) & ifr) < 0) perror ("ioctl
        (SIOCGBURBRT)");
    printf ("route %d", ifr.ifr_metric);

```

5 [EHOSTUNREACH] returned and an audit generated if no route can be found for the given address.

The getsockopt () and setsockopt () calls manipulate the options associated with a socket. The following socket options have been added:

```

    #include <sys/types.h>
10    #include <sys/socket.h>
    int getsockopt (int so, int level, int optname, void *optval, int *optlen);
    int setsockopt (int so, int level, int optname, const void *optval, int
        optlen);

```

SO_STATE - returns the state of the socket, the so_state field of the struct socket. This is used by applications that wish to be more intelligent in their handling of a socket depending on its state.

SO_SOCKBURB - returns the burb id of the socket, similar to ioctl (SIOCGSOCKBURB). A sample call would look like:

```

    int value, length, error;
20    length = sizeof (value);
    if ((error = getsockopt (sockfd, SOL_SOCKET, SO_SOCKBURB, &
        value, & length) != 0) return error;
    return value;

```

SO_MSGBURB - returns the burb id of the pending message, similar to ioctl (SIOCGMSGBURB) via an unsigned long.

[ENOENT] returned and an audit generated if the socket cannot be read.

SO_MSGIFNAME - returns the burb id of the pending message, similar to ioctl (SIOCGMSGIFNAME) via an unsigned long.

30 [ENOENT] returned and an audit generated if the socket cannot be read.

Most of the kernel auditings will be generated by type enforcement checks such as `check_ddt()`, `check_dit()`, and `check_ddt_net()` and user privilege check such as `psuser()`. These audits are logged to `/var/log/audit.asc` and `/var/log/audit.te.asc`.

- 5 Network audit is another type of kernel audit. These are logged to `/var/log/audit.asc` and `/var/log/audit.attack.asc`. Currently, network audits are generated for ICMP, TCP, UDP, and IP packets.

ICMP - ICMP messages are used to communicate error and administrative messages between systems. ICMP audits are the only network audits that can be
10 configurable. This is implemented via the Unix command `sysctl` and the field `burb.X.net.inet.ip.icmpaudit`, where X is the burb id. The default state of the secure computer is level 1, audit only icmp-redirects. ICMP audits can be configured at 3 levels:

- 0 = no icmp audits
- 15 1 = icmp redirects only (default)
- 2 = all icmp messages, except echo and reply

IP - if the source-routed packets option is disabled (0), the kernel will generate a "source-routed packets dropped" audit for each source-routed packet. The source-routed packet option is configurable via the `sysctl` command, by disabling
20 or enabling the fields `burb.X.net.inet.ip.forwarding` and `burb.X.net.inet.ip.forwsrcrt`. As default, both of these fields are disabled so the secure computer will not forward any source-routed packets.

TCP - the kernel will generate an audit for any TCP attempts to synchronize sequence numbers without a valid protocol control block. These types of
25 packets are viewed as hostile probing attempts.

UDP - similar to TCP, the kernel will also generate a probing attempt audit for any UDP attempts to connect without a valid protocol control block.

There are a number of system utilities that configure or retrieve the network states. These utilities are also modified to support burbness.

arp - displays and modifies the Internet-to-Ethernet address resolution. The art display and command line interface will remain the same. Only the internal implementation will be modified to handle the replicated arp tables.

ifconfig - is used to configure network interface parameters. It is modified to

- 5 accept a "burb id" entry so that it can be used to set the burb of an interface. It will also print the burb of an interface as part of its normal output. Example: if config ef0 inet 172.17.128.79 netmask 255.255.0.0 burb 0.

inetd - is the internet-service daemon that runs at boot time and listens for connections on certain internet sockets. This daemon will be replaced by the

- 10 Network Services Sentry, NSS. NSS will be "burb aware" and able to start network services in the appropriate domain. Resident servers will be launched from a similar master utility, such as /etc/rc, so that they run in the correct burb. For more information, see the proxy design specification.

netstat - formats and displays various network-related status and data structures.

- 15 The formatting will be modified so that when appropriate, the -I or -i flags and when printing IP sockets, it will include a burb id column in its output.

sysctl - is used to retrieve kernel state and allows processes with appropriate privilege to set kernel state. Network states are replicated appropriately into NBURBS. For the curious, execute "sysctl-a" and you'll be glad you did it.

- 20 All services supported on the secure computer, such as telnet and ftp, pass their data through a set of proxies (see. e.g. ftp proxy 75 in Fig. 5). There is one proxy for each type of service, http proxy for web, telnet proxy for telnet, and so on. The proxy will accept () the connection or rcvmsg ()/rcvmsg () the packet and look at the incoming address, either in the mesg structure for
25 UDP or via getsockname () for TCP. It will then find out which burb has a route to that destination, via ioctl (SIOCGSUBURBRT), and establish a socket connection in the outgoing burb.

That connection will be via a Unix domain socket. Each service will have a proxy directory in /var/run, such as /var/run/telnetp/. In that

- 30 directory each outgoing proxy will have already opened a domain socket. These sockets exist in the file namespace, so they already have access control via the

domain and type of the socket node. Thus, the burb functionality guarantees that a given process can only access one burb, and the type enforcement on the socket guarantees that process can't cheat and access another service's proxies. The types on the sockets can give a finer control within a service if needed.

- 5 So the following example shows a telnet proxy servicing a telnet from the external network Net1, into the internal network, Net0. /etc/rc is the boot up script, telnetd is the telnet daemon, and nss is the Network Services Sentry daemon.

```

10       -----
          launches an nss for each burb
                  (nns0, nns1)
          launches resident servers for each burb
          launches service proxies (telnetp)

                                  telnetd-out [Net1]
                                  -----
                                  psocket = socket (AF_UNIX,
                                                      SOCK_STREAM)
                                  bind (psocket, "/var/run/proxies
                                                      /telnet/out-net1")
                                  listen (psocket)
                                  accept (psocket)

                                  nss1 [Net1]
                                  -----
                                  receives connect () and accept()s
                                  launches telnetd-in.

15       telnetd-in [Net0]
          -----
          does getsockname () to get dest
          does SIOCGSUBURBRT on dest
          does a connect
20               (/var/run/proxies/telnet/out-Net1)

                                  accept () returns

          writes destination to proxy

                                  read()s destination
                                  does a connect () to the real destination
                                  passes back success

          select () loop passing data
25       cleanup on close

                                  select () loop passing data
                                  cleanup on close

```

Packets can be filtered as part of the process. Filtering could go before the proxy, it could go after the proxy, or it could go inside the proxy.

Some things, like ftp, are more complex. The proxies have to watch for PORT and PASV requests and set up the data channel and proxy it.

- 5 Additional channels can be either setup in advance with a standard protocol, or created on the fly by having proxies set up a control channel and bind additional separate domain sockets.

- UDP proxies are also more complex. The proxies may have to maintain state information about requests, so that they know which request goes
- 10 where. A UDP (or other datagram) proxy can work two ways. It can either have a single proxy per burb that just forwards and filters packets one at a time (in which case the connecting Domain socket is a datagram socket as well), or it can create a proxy for each "pseudo" connection and track state to optimize things (like sending acks, etc.). In this case there could be a pair of proxies on the fly
- 15 that created a stream connection between themselves. They could be routed by port (for protocols that have a unique port and each end negotiated on the fly, like gopher or DNS), otherwise a single proxy still has to handle incoming packets and dispatch them.

- Which approach used depends on a few things. If address hiding
- 20 and masquerading is done, state must be maintained. If not, then simple packet forwarding can be done. Some protocols, like NFS (if ever done) will require extensive state.

Type Enforcement Protection of Compiled Program Code

- Type enforcement can be used to extend the access protections
- 25 inherent to a particular program, without access to the source code for that program. An overview of the process used to install any pre-compiled application binary (i.e. the program) onto the secure computer system is shown in Fig. 6. The first step (100) is to install the binary to a location from which it can be executed. Some (but not all) examples of how this can be accomplished
- 30 are: (a) copying from the installation floppy disk(s) to the internal hard drive, (b) copying from the installation cdrom to the internal hard drive, (c) mounting

external media containing the application binary, (d) copying the application binary via the network to the internal hard drive.

Once the application binary is accessible for execution, the secure computer system is placed in a different state than the fully secure operational state. This state is called the development state, since in best practice (i.e. for highest security) this state is only available to developers. The development state has all of the type enforcement checks that the operational state has; however, the consequences of a violation are different. In operational state, a type enforcement violation results in denying the requested access. In development state, a type enforcement violation results in allowing the requested access. In both the operational and development state, a log entry is produced which records the type enforcement violation. This logging capability is needed to determine which type enforcement violations the program is generating.

It should be noted that the system does not have to be placed in development state to perform this procedure. All that is absolutely required is the logging capability. However, by placing the system in development state, the time and effort required to perform this procedure is greatly reduced. In practice, the procedure detailed in Figure 6 is followed under in the operational state only when an error was made when performing the procedure under the development state.

At step 102 the program is executed. Because this program was developed using non-type enforcement techniques, this program will make assumptions about its execution environment that will not be true under the type enforcement execution environment. These faulty assumptions will manifest themselves as type enforcement violations. Some sample faulty assumptions that the program can make include (but is not limited to): (a) the program can create files in the system-wide temporary file area, (b) the program can read files that have Unix "global-read" permissions set to true, etc. These faulty assumptions are a direct consequence of the program being non-type enforcement aware: the program is unaware that, in addition to the Unix security restrictions that it is expecting, there are type enforcement security restrictions.

Every type enforcement violation must be successfully dealt with in order for the program to function properly while under the operational state.

There are three general ways to successfully deal with a type enforcement violation. They are: (1) ignore the violation, (2) remove the request that generated the violation, (3) grant additional permission(s) to remove violation. The criteria for choosing one branch over another branch are functionality and security. The less critical the functionality, the more likely that the violation can be ignored or the request removed. The more critical the functionality, the more likely that additional permission(s) will be granted.

10 Some violations can be safely ignored. For example, some programs attempt to determine how busy the system is, in order to behave differently under different conditions. If, however, the program is unable to determine how busy the system is (because of a type enforcement violation), the program will simply assume that the system is not busy and continue. For some
15 programs this is an acceptable response, and requires no further action; in these cases (108), type enforcement violations can be ignored.

 Some violations can be removed by changing the environment that the program is running under. For example, some programs attempt to change their effective user identification after some resource has been acquired.
20 The most common situation of this kind is when the program runs as "root" long enough to open a file, but then changes to "wwwdaemon" immediately thereafter. Under type enforcement, this ability to change to another user will result in a violation under the default operational state. In cases such as these, it is more security effective to remove the need for the program to change to
25 another user by following these steps:

 (1) start the process as user "wwwdaemon" directly (thus, the program will no longer require the ability to change to another user);

 (2) set the Unix permissions on the file such that it can be read by user "wwwdaemon"; and

30 (3) set the type enforcement permissions on the file such that it can be read by the domain in which the process started in step (1) is executing.

So, in this example, the original type enforcement violation is removed at 110 by changing the environment external to the program. Note that in a non-type enforcement system, step (2) results in a security vulnerability. It is only through step (3) that the security vulnerability is closed.

5 In the above example, one can think of the original type enforcement violation as being transformed into another type enforcement violation. The ability to "change user" was transformed into the ability to "read a file".

 Some violations can be removed by granting the program
10 appropriate type enforcement permission(s), so that what was a violation now is a granted permission. Once it is determined that the violation can neither be ignored nor transformed, this action is performed (112). There are three general kinds of type enforcement violations. They are: (1) File access, (2) Process-to-Process communication, and (3) System Environment Access.
15 Therefore, a check must be made at 114 to determine the variety of type enforcement violation.

 If the violation concerns general file access, a check is made at 116 to see if the file is a shared file. Programs typically need access to many different sorts of files: configuration files, temporary or scratch files, permanent
20 data files, etc. The only factor that distinguishes these sorts of files from another in the type enforcement sense is whether or not the file is a shared file or not. An example of a non-shared file could be a temporary file. Here, only the program needs access, and no other. An example of a non-shared file could be a database file containing a list of every user of the system.

25 If the file is shared, then a shared type must be used. Sometimes, this shared type will not yet exist, and must be created. This can occur, for instance, when two new programs both access the same new file (for example, a configuration file that one program writes and the other program reads). Sometimes, the shared file will already exist. This can occur, for instance, when
30 the new program accesses a system-wide file (for example, a database file

containing a list of every user of the system.) If the shared file already exists, it will already have a type. This type is used at 118.

If the file is not shared, a new type must be created, and the file is set to this new type.

5 In either case (both shared and non-shared), the final step is to add at 120 a direct permission to the type enforcement database. This type enforcement permission will grant the domain (in which the program is running) permission on the type (and the file is then set to that type). The fact that the file is not shared is manifested in that there will not be another domain with
10 permission on the type. The fact that the file is shared is manifested in that there will be another domain with permission on the type.

The second kind of violation concerns process-to-process communication. Some programs require the ability to communicate with another program via a process-level mechanism. To deal with this kind of
15 violation, one determines the domain of the second program (the program that the first program is trying to communicate with). Then, at 122 a direct permission is added to the type enforcement database. This permission will grant the sending domain (in which the first program is running) permission on the receiving domain (in which the second program is running). Permissions of
20 this kind are called domain-to-domain interactions, and effectively control the permitted lines of communication between processes.

The third kind of violation concerns system environment access. This kind captures all the kinds of violations that do not concern file access or process-to-process communication. An example would be the ability to
25 determine how busy the system is. Previously, this functionality was used to demonstrate a case where the ability could be removed. However, if the ability was required, then dealing with the type enforcement violation would be handled under this kind of violation. To deal with this kind of violation, one determines the specific functionality required. Then, at 124 a direct permission is added to
30 the type enforcement database. This permission will grant the program's domain permission to invoke the required system function.

The Use of Both Type Enforcement and Network Separation to Protect
Compiled Program Code

The process described in the previous section is sufficient to handle those instances where programs do not have to be compartmentalized to achieve sufficient security (e.g. utilities such as screen savers or logging utilities). In that case, all programs can reside in and be bound to one region or burb. For systems which can be partitioned into distinct burbs, however, an additional degree of protection can be added by incorporating network separation.

10 In the case, for instance, where the program(s) to be installed can reside in multiple burbs, yet must maintain the ability to communicate with each other via a networking interface, network separation is used to increase the overall security of the system. One example of such a case is where the system consists of two programs: a production-side program and an administration-side
15 program. In these instances, it is required that the production-side program reside in an unprotected (e.g. Internet) burb. But one does not want the administration-side program residing in an unprotected (e.g. Internet) burb. So instead, the administration-side program is placed in a protected (e.g. internal) burb.

20 However, by splitting the two programs (production and administration) into two burbs, one has again changed the actual execution environment away from the expected execution environment. The two programs expect (and rely on) the execution environment to allow process to process communication (e.g. signals) between the two programs. On a non-Type
25 Enforced, non-Network Separated system, this assumption is valid. However, on a Network Separated system, the act of placing the administration-side program in a different burb than the production-side program will have the Network Separation side effect of disallowing process to process communication between the administration-side program and the production-side program. From a
30 functionality standpoint, this is not acceptable.

To restore the required program to program communication functionality, one is required to perform steps outside of those explained in the previous section. The problem to be solved concerns the conflicting restrictions placed by the type enforcement system and the network separation system.

5 In Network Separation, if one program has been "bound" to one burb (e.g. Internet) and another program has been "bound" to another burb (e.g. internal), the Network Separation mechanism will preclude the ability for the first program to establish a connection to the second program via a network communication connection; thus it precludes the ability for program to program
10 communication using the network. One way of viewing Type Enforced Network Separation is to acknowledge this lack of direct network to network communication: since the networks cannot be shared, they are separated.

So, the problem here can be restated as follows: given a program bound to one burb and another program bound to a different burb, how can we
15 allow the two programs to communicate with each other without having access to the source code for either program.

The procedure to solve this problem is shown in Fig. 7:

First, at 140 decide to which burb each of the separate programs is to be bound. In the case of a production-side program and an
20 administration-side program, it is clear that the production-side program must be bound to an Internet burb, and it is clear from a security standpoint that the administration-side program should be bound to an internal burb. In the case of more complicated sets of programs, tradeoffs must be made between availability and security. Binding a program to an Internet burb increases availability but
25 also enormously increases security risks. Binding a program to an internal burb decreases security risks, but results in a decrease in availability.

Second, at 142 determine how the two programs interact. Again, it is important to note that source code is not available for this step. If source code is available for inspection and modification, then this procedure does not
30 apply. Without source code, this step is accomplished by examining the type enforcement violations generated by either or both programs, as required.

Third, at 144 determine whether the type enforcement and network separation violations discovered in the second step can be eliminated. Techniques for removing TE violations (such as those shown in Fig. 6) can be used. In addition, one may be able to eliminate a type enforcement violation by
5 placing one of the programs into a situation where it is "partially bound" to a particular burb. A "partially bound" program is a program which can interact with a specific burb under a reduced set of Network Separation rules. That is, a partially bound program can perform process to process communication with another burb but not network to network communication. Specifically, it is
10 possible for the partially bound program in a first burb to signal a program that has been completely bound to another, separate burb, while simultaneously maintaining other connections to the first burb. Therefore, communication between burbs assigned to different burbs is permitted in a limited way and the problem is solved.

15 In one embodiment, one creates this "partially bound" program by using a type enforcement database trick. Normally, a program "P" is completely bound to a burb when the following syntax is used:

"P" runs in domain www#

which means that program "P" can access any of the burbs 0 through 9 (as
20 signified by the character '#' above). When program "P" first starts running, it has its choice of burbs (e.g. www0, www1, ... www9) to establish a connection. But after the first connection, the program "P" is fully bound to that burb and the domain that is associated with that burb, and cannot establish a connection to any network in another burb.

25 A program "P" is "partially bound" to a burb when the following syntax is used:

"P" runs in domain adm0

Here, the program "P" loses some of its flexibility (it can now connect only to burb adm0) but it gains the "partially bound" status. Furthermore, since such an
30 approach is treated as if there are no other instantiated domains (i.e. no adm1,

adm2, etc.), program P running in domain adm0 can now connect to all other
"#" -defined domains.

For the production-side program "PP" and administration-side
program "AA" example above, a full example might be:

5 "AA" runs in domain adm0

"PP" runs in domain prd#

Thus, since the program "AA" is "partially bound" to 'adm0', there can be no
domains adm1 through adm9. Program "AA" has, however, gained the ability to
interact with prd0 through prd9. Thus, the required signalling ability is
10 re-established.

If it is not possible to place all but one of the programs in a
"partially bound" burb, then this procedure must stop without succeeding. The
only alternative at this point is to place all programs in the same burb. Then,
since the additional Network Separation checks are no longer an issue, this
15 procedure degenerates into the type enforcement only approach described
previously.

Secure Commerce Server

One embodiment of a binary program made secure through a
combination of type enforcement and network separation is described next. The
20 Netscape Commerce Server, version 1.1, available from Netscape Computer,
Inc., is a server for use in conducting commercial transactions over the Internet.
This server provides user authentication, SSL protection to http connections, and
a forms interface for server administration. As such, it is critical that care be
taken to protect a malicious attacker from gaining access to and subverting the
25 program.

A type enforced, network separated embodiment 160 of the
Netscape Commerce server system is shown in Fig. 8. The Netscape Commerce
server system includes two distinct servers: the Commerce server 162 and the
administration server 164. In the preferred embodiment, one administration
30 server 164 should be able to administer and configure a plurality of separate
Commerce servers. In the system of Fig. 8, the dashed lines indicate a user level

permission check such as the Unix file permissions. In addition, components of the system are placed within domains 166, 168 and 172 for a higher level of security.

- Commerce server 162 is installed via an HTML forms driven
- 5 interface. Administration and configuration after installation is done in the same manner. Installation operates as follows:
1. The user runs the script ns-setup in the source tree (the directory structure from which files are copied to install the Netscape Commerce Server).
 2. The user is prompted for the hostname of the host machine.
 - 10 3. A temporary configuration is written in /tmp and a special WWW server is started using this configuration (the installation server).
 4. The installation server binds to an arbitrary, non-reserved port.
 5. The user is prompted for the name of a Web browser to start up. It then starts the browser on the port the installation server is listening on.
 - 15 6. The rest of the installation is HTML forms-driven through the browser. Various items such as port number for the Commerce server, UID to run the server under, install directory, logging, administration password, and other server configuration are entered via three forms.
 7. The configuration is verified and the actual installation takes place:
 - 20 1. create destination directory and subdirectories
 2. create document root directory
 3. create server startup and shutdown scripts
 4. create server configuration files
 5. copy all binaries, administration forms, and other files from
 - 25 source tree to destination tree
 6. change owner of directories
 7. start administration server
 8. start Commerce Server
 9. remove files created in /tmp
 - 30 8. Installation complete.

Many Commerce servers 162 can be run on the same machine, binding to different ports (or even the same port, using different IP addresses) and having their own configuration. Each server preforks 16 processes (number is configurable) to serve requests. This can load the system down if many
5 servers are running.

After installation, the Commerce server is administered via forms using any forms-capable browser and connecting to a separate administration server running on its own port. This server is started and stopped manually, except during installation when it is started automatically. (It is recommended
10 by Netscape that it be stopped when not being used.) The administration server is also configurable via forms (it configures itself). It can be configured to require a password to access it, and to allow only certain hostnames or IP addresses to connect to it.

The Commerce server, once running, serves Web pages from a
15 directory tree whose root is configurable. CGI script location, URL mappings, user authentication, access control by host, logging, and other items are also configurable via the administration server. Configuring some items causes creation of files, such as generating a key, installing a certificate, and creating a user database.

20 **High Level Design**

This section describes what must be done to port the Netscape Commerce Server to the secure computer. Since we have no source code from Netscape, we cannot modify how the server or installation process operates. Thus, design will focus on what needs to be done outside the Netscape source
25 code to handle type enforcement.

The installation process will be done completely in the administrative kernel (no type enforcement checks), so setting file types in the source tree is not necessary. All file types must be set after the install process is complete since Netscape is not type-enforcement (TE) aware. To accomplish
30 this, a script is used to set all types appropriately. The script must be able to find the appropriate directories, because paths are configurable, and the server files

are in a directory named for the port it binds to. This info is passed via arguments or entered by the installer via prompts.

Commerce Server 162 runs in domain 166. In one embodiment, this is the same domain in which the CERN server (not shown) runs. Domain
5 166 is a bound (burbed) domain which is extended to handle the Commerce Server. That is, the Commerce Server is extended to be able to bind to port 443 (https), 80 (http), 8000, 8001, and 8080. In addition, to allow site flexibility, Commerce Server 162 must be able to bind to selected reserved port ranges.

Administration Server 164 runs in a new Netscape Admin
10 domain, domain 168. Domain 168 must be a burbed domain, since it does a non-burb-aware socket () call. Transfers between domains must be done through proxy programs acting in concert with kernel 170.

CGI scripts run in a CGI processor 171 in burbed CGI domain
172. Since, however, the Netscape server cannot be modified to do a
15 makedomain () to run the CGI script in the proper domain, there will need to be a :tran type added to CGI domain 172. Then , all CGI scripts will be of type tran. In an alternate embodiment, a line could be added to the CGI script itself which causes the script to automatically transition into CGI domain 172. To run
20 a CGI script, Commerce Server 162 initiates a new process which executes in CGI processor 171 within CGI domain 172. The results are then transferred back to Commerce Server 162. In one embodiment, the CERN server also executes CGI scripts within CGI domain 172.

For installing two servers on the same port, but with different IP addresses:

- 25 1. Follow the install directions twice, specifying everything the same except a different server name (which will need to resolve to the correct IP address), bind address, and document root.
2. Make an IP alias so that both IP addresses used for the bind
addresses in the previous step access the secure computer. Type:
30 ifconfig efl alias 172.17.128.199 to alias the IP address
172.17.128.199 to the external interface.

3. Access both servers using the different IP addresses.

Files Manifest

The following are files associated with the Commerce Server. In this section, /server-root/ will denote the path where the server files are installed.

- 5 Also, the subdirectory https-443 will denote the directory containing all files specific to the server running on port 443. The actual name will contain the port number the server binds to in place of '443'.

● **Executables**

/server-root/bin/https/ns-httpd

- 10 The Commerce server.

/server-root/admserv/ns-admin

The administration server.

/server-root/admserv/servlist

Part of the administration server.

- 15 /server-root/bin/https/admin/bin/*

CGI scripts for the admin server that generate forms, interpret responses, change config files, administer servers, generate keys, etc.

/server-root/extras/database/batchddb

Utility to convert NCSA-style user database to DBM database.

- 20 /server-root/extras/database/changepw.cgi

CGI script to allow users to change their own password via forms.

/server-root/extras/log_anly/analyze

Analyze access logs.

/server-root/extras/log_anly/a_form.cgi

- 25 CGI script to analyze access logs via forms.

/server-root/https-443/restart

Restart the Commerce server.

/server-root/https-443/rotate

Rotate the log files.

- 30 /server-root/https-443/start

Start up the Commerce server.

/server-root/https-443/stop

Shut down the Commerce server.

/server-root/start-admin

Start up the administration server.

5 /server-root/stop-admin

Shut down the administration server.

- **Configuration files**

/server-root/admserv/admpw

User and password file for access to administration server.

10 /server-root/admserv/ns-admin.conf

Configuration file for the administration server.

/server-root/bin/https/admin/html/*

Forms templates for the administration server.

/server-root/bin/https/admin/icons/*

15 Icons for the administration server.

/server-root/mc-icons/*

Icons for the Commerce server, used for gopher and ftp listings.

/server-root/userdb/*

User databases. (Initially empty)

20 /server-root/https-443/config/admin.conf

Commerce server configuration file.

/server-root/https-443/config/magnus.conf

Commerce server configuration file.

/server-root/https-443/config/obj.conf

25 Commerce server configuration file.

/server-root/https-443/config/mime.types

Commerce server configuration file.

/server-root/https-443/config/ServerKey.der

Key pair file. Initially non-existent, generated by admin server when requested by the administrator. Path is configurable.

30

/server-root/https-443/config/ServerCert.der

Certificate file. Initially non-existent, installed by admin server from the Certificate Authority's certificate response when requested by the administrator. Path is configurable.

- **Log files**

5 /server-root/admserv/errors

Log file for administration server.

/server-root/https-443/logs/access

Commerce server access log. Path is configurable.

/server-root/https-443/logs/errors

10 Commerce server error log.

/server-root/https-443/logs/secure

Commerce server secure log. (All https accesses are logged here with the keysize used).

- **Temporary files**

15 /server-root/admserv/pid

Contains the process ID of the administration server when running.

/server-root/https-443/logs/pid

Contains the process ID of the Commerce server when running.

- **System files**

20 /etc/spwd.db

System password file containing encrypted passwords.

/etc/pwd.db

A shadow password file.

- **Other files**

25 /server-root/extras/database/changepw.htm

HTML form to allow users to change their own password.

/server-root/extras/log_anly/a_form.html

HTML form to analyze access logs.

/server-root/docs/index.html

30 Default home page for the Commerce server. Path is configurable.

CGI scripts

Path is configurable.

HTML pages

Path is configurable.

As noted previously, Netscape Commerce Web server 162 runs in the Web server domain (domain 166), the same domain that the CERN Web server runs in. Therefore, the Web server domain needs execute permission to the following:

/server-root/bin/https/ns-httpd

In addition, the Web server domain needs read permission to these files:

/server-root/mc-icons/*

/server-root/userdb/*

/server-root/https-443/config/admin.conf

/server-root/https-443/config/magnus.conf

15 /server-root/https-443/config/obj.conf

/server-root/https-443/config/mime.types

/server-root/https-443/config/ServerKey.der

/server-root/https-443/config/ServerCert.der

/server-root/extras/database/changepw.htm

20 /server-root/extras/log_anly/a_form.html

/server-root/docs/index.html

The Web server domain needs write permission to these files:

/server-root/https-443/logs/access

/server-root/https-443/logs/errors

25 /server-root/https-443/logs/secure

/server-root/https-443/logs pid

Since the server normally runs as root, if it were able to execute arbitrary code, this would be bad. We can address this by not having it run as root, taking advantage of the type enforcement protection of the sockets. This allows users other than root to connect to low numbered sockets. Other than
30 needing to bind to port 80 or 443, there is no other reason the server needs to run

as root. Therefore, the server will run as user "www" and group "www" with no need to run as root.

The CGI bin executables present a potential security concern.

Since the Commerce server is allowed to execute these files, if someone were
5 able to put their own executable on the secure computer, the server may be compromised. TE only allows the Web server domain to transition to the CGI domains and no others. It does not have permission to create or write to files of any executable type. This helps prevent the possibility of an external user subverting the server and uploading an executable file.

10 The entire directory tree containing the html documents which the Commerce server accesses will be write protected from the Web server domain. A separate writable directory will be set aside for all the data created by external users. This is where the CGI executables will put their results. The CGI domain will be allowed to create files in this directory, but not to read or destroy them.
15 These files may later be read and moved internally by the webmaster using mail or FTP.

The password administration program will run in the Netscape Admin domain (domain 168). This domain will have all the accesses required to maintain the Admin server and all Commerce servers on the secure computer.

20 The Admin domain needs execute access to the following:

/server-root/admserv/ns-admin
/server-root/admserv/servlist
/server-root/bin/https/admin/bin/*
/server-root/extras/database/batchdb
25 /server-root/extras/log_anly/a_form.cgi
/server-root/extras/log_anly/analyze
/server-root/https-443/restart
/server-root/https-443/rotate
/server-root/https-443/start
30 /server-root/https-443/stop
/server-root/start-admin

/server-root/stop-admin

The Admin domain needs read access for:

/server-root/bin/https/admin/html/*

/server-root/bin/https/admin/icons/*

5 /server-root/extras/database/changepw.htm

/server-root/extras/log_anly/a_form.html

HTML pages

The Admin domain needs both read and write access for:

/server-root/admserv/admpw

10 /server-root/admserv/ns-admin.conf

/server-root/admserv/pid

/server-root/admserv/errors

/server-root/userdb/*

/server-root/https-443/config/admin.conf

15 /server-root/https-443/config/magnus.conf

/server-root/https-443/config/obj.conf

/server-root/https-443/config/mime.types

/server-root/https-443/config/ServerKey.der

/server-root/https-443/config/ServerCert.der

20 /server-root/https-443/logs/access

/server-root/https-443/logs/errors

/server-root/https-443/logs/secure

/server-root/https-443/logs/pid

CGI scripts run in CGI domain 172. The following files will be

25 of cgix:tran type and will run in the CGI domain:

/server-root/extras/database/changepw.cgi

/server-root/extras/log_anly/a_form.cgi

CGI-scripts

Commerce Server 162 is controlled by the various configuration

30 files maintained by administration server 164. In one embodiment, it may be

advantageous to add the ability to start and stop administrative server 164 from another system administration program running in another domain.

Since the secure computer uses type enforcement on sockets, we can allow users other than root to bind to reserved ports. As a result, we run the Commerce server as user id "www" and group id "www". The "www" user is similar to "nobody" and has no log in capabilities. The administration server will still run as root. In addition, scripts that start and stop the servers need to execute /bin/sh (type \$Sys:shel). These scripts are also not type-aware, so they need to be modified to execute in the correct domain.

In one embodiment, the following domains are used to define the commerce server system. It should be apparent that other combinations of domains, privileges or access rights could be used. In the list, each domain can have the following privileges: "is_admin" indicates the domain is an administrative domain, "has_rootness" means that the domain can violate Unix permissions if the process UID is root, "can_setlogin" indicates that the domain can set the login name of a process. In addition, each domain's DIT could have the following permissions: "dt" indicates that transition to the named domain is allowed, "sABRT" indicates that sending an ABRT signal to the named domain is allowed, "sJob" indicates that sending a job control signal to the named domain is allowed, "sHUP" indicates that sending a HUP signal to the named domain is allowed, and "sUser" indicates that sending a USR1 or USR2 signal to the named domain is allowed.

Finally, each DDT uses the following permissions: "w" is permission to write, "r" is permission to read, "d" is permission to destroy, "n" is permission to rename, "c" is permission to create and "e" is permission to execute.

Domain: CGI 'cgix'

{ CGI scripts run here - limited access to things }

Privs:

DIT:

CGI sABRT { }

Subtype: Script 'scrip'

{ Used for CGI shell scripts }

DDT:

52

```

        CGI    e r    { }
        Webmaster wr d n c    { }
        www    e r    { }
    Subtype: Transistion 'tran'
5      { Used for CGI compiled executables, not shell scripts. }
      DDT:
        Webmaster r e w n { }
        www    e    { }
    Subtype: File 'file'
10     { }
      DDT:
        CGI    wr d n c    { }
        Webmaster wr d n c    { }
    Subtype: Directory 'diry'
15     { }
      DDT:
        CGI    wr    { }
        Webmaster wr d n c    { }
        www    { }
20     Subtype: Save 'save'
      { }
      DDT:
        CGI    c    { }
        Webmaster wr d n c    { }
25
    Domain: NetscapeAdmServ 'nas0'
      { This domain is for the Netscape Admin Server. }
    Privs:
      can_setlogin
30     DIT:
        CGI    dt    { }
        NetscapeAdmServ sJob sABRT { To control itself. }
        www    dt sJob    sHUP { }
    Subtype: Executable 'exec'
35     { }
      DDT:
        NetscapeAdmServ r e    { }
        Webmaster    r e    { }
    Subtype: Config 'conf'
40     { The Netscape Admin Server configuration file }
      DDT:
        Webmaster    r w n c d    { }
        NetscapeAdmServ r w d n c    { }
    Subtype: File 'file'
45     { Default file type. This is the type created whenever the Admin
      Server creates a file (Key/Certificate files, User Databases,
      temporary files, etc.). }

```

```

      DDT:
        Webmaster   w r d n c      { }
        NetscapeAdmServ   w r d n c      { }
        www   r      { To read Key and Certificate files, and
5                      User databases created by the Netscape
                      admin server. }

      Subtype: Directory 'diry'
        { }
      DDT:
10      NetscapeAdmServ   w r d n c      { }
        Webmaster   w r d n c      { }

      Domain: Webmaster 'Webr'
        { This domain is for web admin work. }
15      Privs:
        is_admin
        has_rootness

      DIT:
        CGI   dt sJob      { }
20      NetscapeAdmServ   dt sJob      sABRT      { }
        Webmaster   sJob sHUP sABRT      { }
        www   dt sJob sHUP sABRT      { }

      Subtype: File 'file'
        { Default file type. }
25      DDT:
        Webmaster   e w r d n c      { }

      Subtype: Executable 'exec'
        { }
      DDT:
30      Webmaster   e      { }

      Subtype: Directory 'diry'
        { }
      DDT:
        Webmaster   w r d n c      { }
35

      Domain: www 'www#'
        { The www domain is where any Web servers will be run, such as the
        Commerce Server. }
      Privs:
40      DIT:
        CGI   dt      { }
        www   sUser sJob sHUP sABRT      { Needs to control itself. }

      Subtype: Config 'conf'
        { The httpd configuration file(s). }
45      DDT:
        NetscapeAdmServ   w r      { }
        Webmaster   r w n c d      { }

```

```

        www r { }
Subtype: File 'file'
    { Default file type. This is the type created whenever the
      Commerce Server creates a file (log files, temporary files, etc.). }
5   DDT:
        NetscapeAdmServ r n { }
        Webmaster w r d n c { }
        www w r d n c { }
Subtype: Page 'page'
10  { web pages }
    DDT:
        NetscapeAdmServ r { }
        Webmaster r w n c d { }
        www r { }
15  Subtype: Directory 'diry'
    { }
    DDT:
        CGI r { }
        NetscapeAdmServ w r { }
20  Webmaster w r d n c { }
        www w r d c { Web stores cached files in dirs, needs to
                        make and destroy as needed}

Domain: wwwCommon 'wwwc'
25  { The common Web server domain (for items shared across burbs). }
    Privs:
    DIT:
    Subtype: Executable 'exec'
        { This is the Web Server executable type. }
30  DDT:
        Webmaster e { }
        www e { Need to be able to run the server. }
Subtype: Config 'conf'
    { }
35  DDT:
        NetscapeAdmServ w r d n c { }
        Webmaster w r d n c { }
        www r { }
Subtype: Directory 'diry'
40  { }
    DDT:
        CGI r { }
        NetscapeAdmServ w r { }
        Webmaster w r d n c { }
45  www r { }
Subtype: Page 'page'
    { }

```


DDT:

5 CGI r { }
 NetscapeAdmServ r { }
 Webmaster w r d n c { }
 www r { }

 Although the present invention has been described with reference
 to the preferred embodiments, those skilled in the art will recognize that changes
 may be made in form and detail without departing from the spirit and scope of
10 the invention.

What is claimed is:

1. A method of achieving network separation within a computing system having a plurality of network interfaces, including a first and a second network interface, the method comprising the steps of:
 - 5 defining a plurality of burbs, wherein the plurality of burbs includes a first and a second burb and wherein each burb includes a protocol stack;
assigning each of the plurality of network interfaces to one of the plurality of burbs, wherein more than one network interface can be assigned to a particular burb, wherein the step of assigning includes the steps of assigning the
10 first network interface to the first burb and the second network interface to the second burb;
binding processes to burbs; and
restricting communication between processes bound to different burbs.
- 15 2. The method according to claim 1, wherein the step of restricting communication between processes includes the steps of:
 - limiting each single user process to sending and receiving information from only one burb during the lifetime of that process;
preventing a process from accessing burbs other than the burb to which
20 the process is bound;
limiting transfer of incoming packets such that incoming packets can only go to processes bound to the burb associated with the network interface the packet arrived on;
defining a proxy; and
25 ensuring that data passing from the first network interface to the second network interface must pass through the proxy before moving from the first burb to the second burb.
3. The method according to claim 1, wherein each burb has its own routing
30 table and wherein the step of restricting communication includes the step of

examining an incoming packet to determine if its destination is an address in the burb's routing table.

4. The method according to claim 1, wherein the step of restricting
5 communication includes the steps of:
 defining a plurality of types;
 assigning each data object to one of the plurality of types, wherein the
 step of assigning data objects includes defining type enforcement rules for
 accesses by processes to said data objects; and
10 applying a type enforcement check of accesses by processes to data
 objects.
5. A secure server, comprising:
 an operating system kernel;
15 a plurality of network interfaces which communicate with the operating
 system kernel; and
 a plurality of burbs, including a first and a second burb;
 wherein each network interface is assigned to one burb from the plurality
 of burbs, wherein each burb includes a protocol stack for handling
20 communication across the network interface, wherein each protocol stack is
 associated with only one burb and wherein communication from a process bound
 to the first one burb must pass through a proxy before being sent to the second
 burb.
- 25 6. The secure server according to claim 5, wherein each burb has its own
 routing table and wherein each burb further includes means for examining an
 incoming packet to determine if its destination is an address in the burb's routing
 table.
- 30 7. The secure server according to claim 5, wherein the operating system
 kernel includes type enforcing means for applying a type enforcement check of

accesses by processes to data objects, wherein the type enforcing means includes type defining means for defining the modes and consequences of accesses by processes to data objects.

- 5 8. The secure server according to claim 5, wherein the operating system further includes a domain interaction table which enforces interactions between processes.




FIG. 1

2/10




FIG. 2A




FIG. 2B

SUBSTITUTE SHEET (RULE 26)

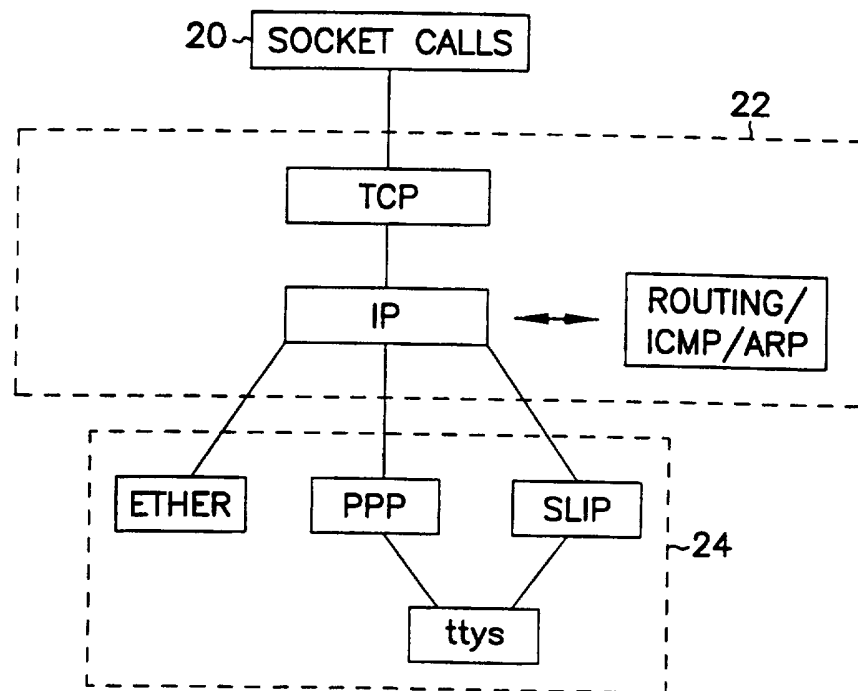


FIG. 2C




FIG. 2D




FIG. 3

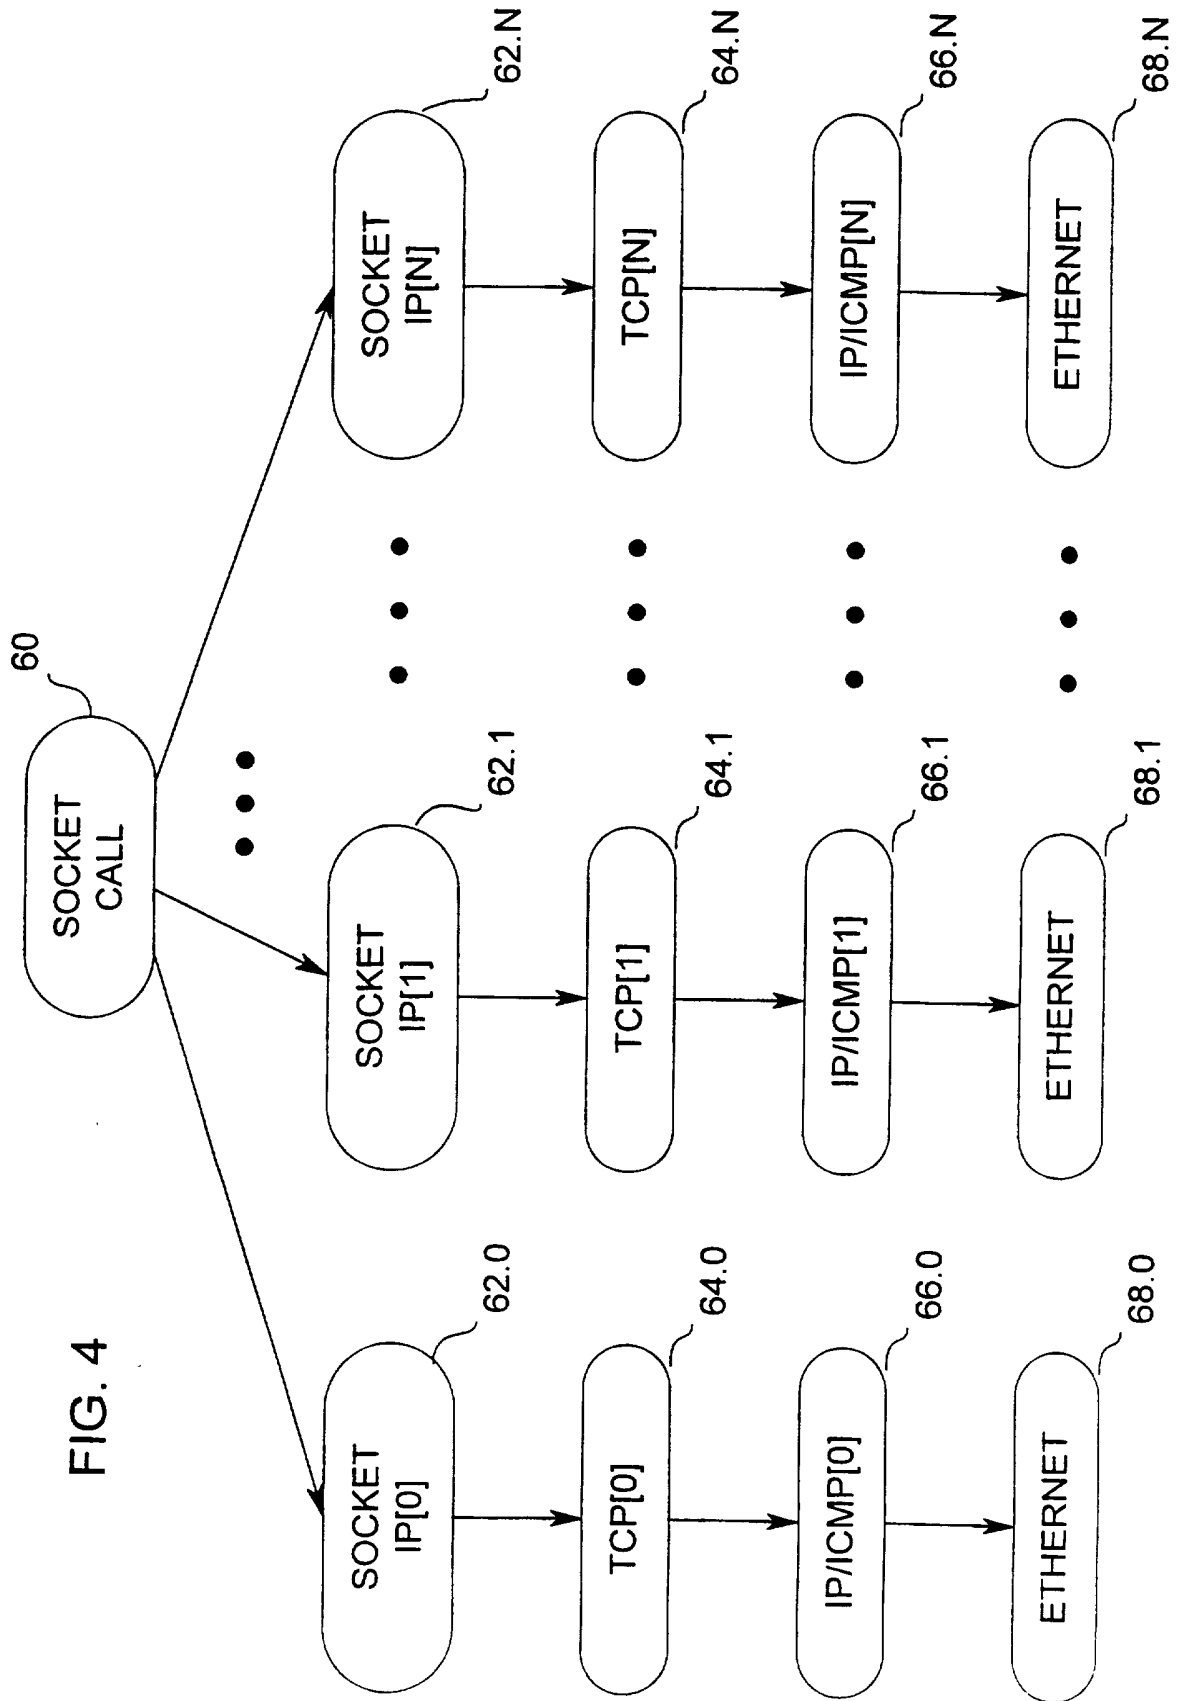


FIG. 4


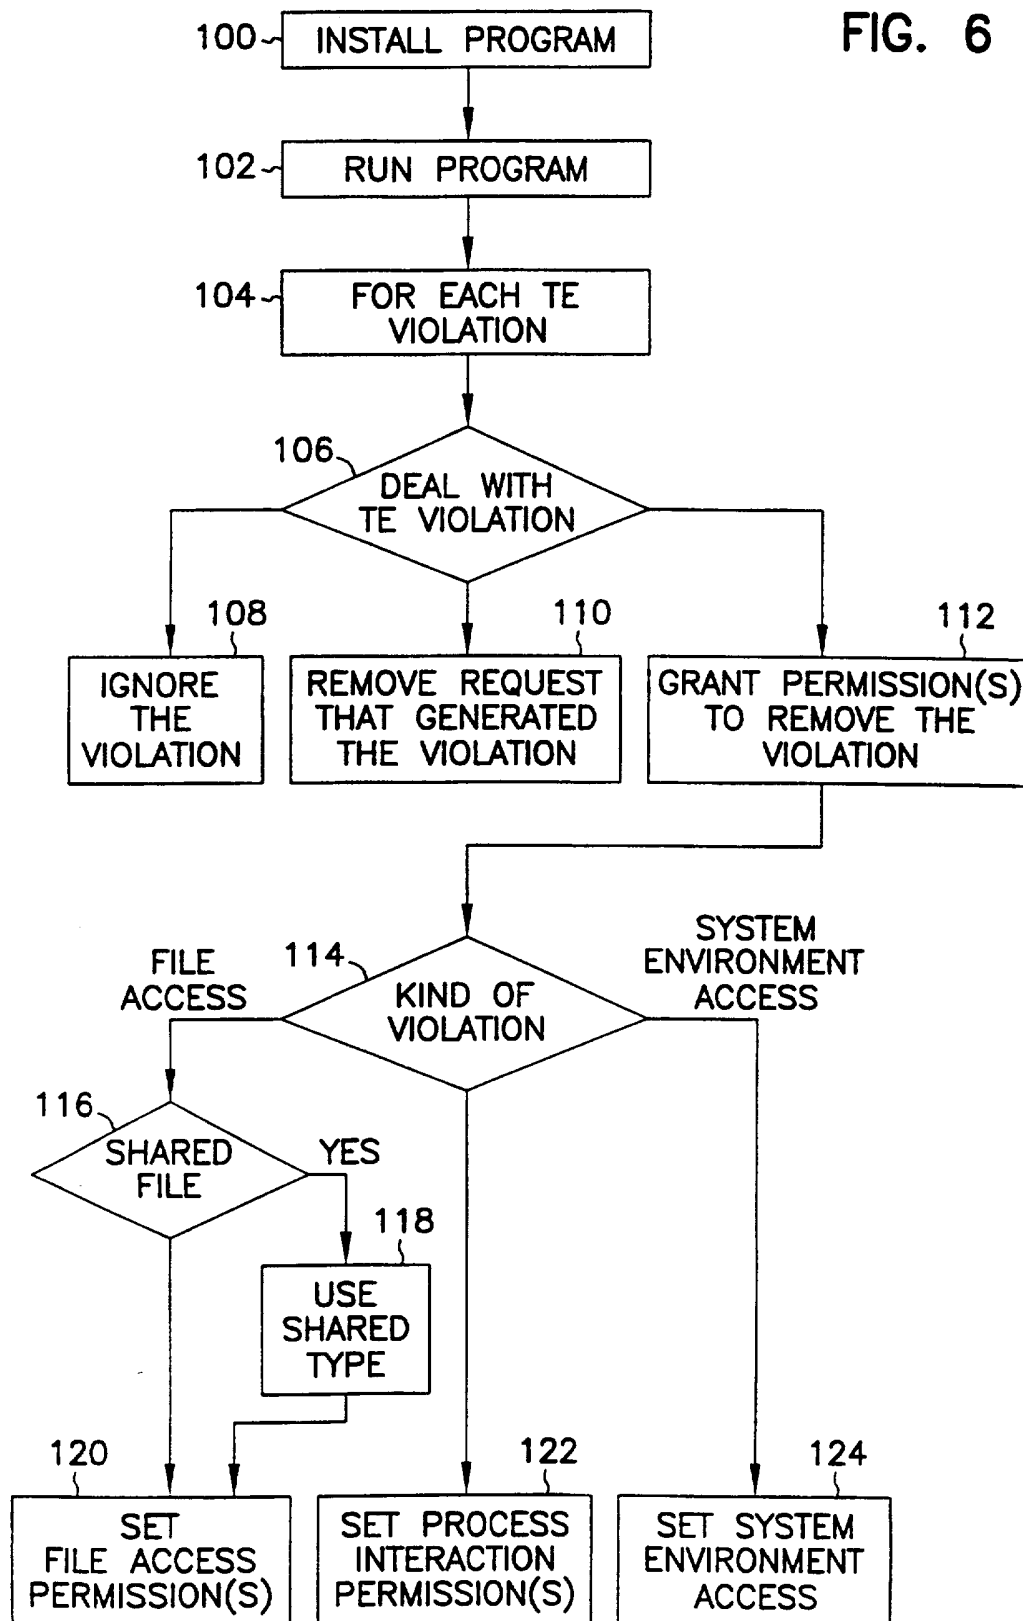


FIG. 5

FIG. 6



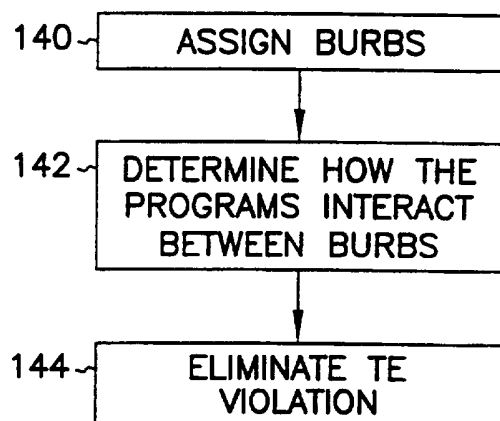


FIG. 7

