

The Implementation of Lisex, a MLS Linux Prototype

Maximiliano Cristiá Gisela Giusti Felipe Manzano
{mcristia, ggiusti, fmanzano}@fceia.unr.edu.ar
GIDISS*
Facultad de Ciencias Exactas, Ingeniería y Agrimensura
Universidad Nacional de Rosario
Argentina

Abstract

In this article we describe the design and implementation of a Linux multi-level secure (MLS) file system containing access control lists (ACL). The resulting prototype is called Lisex. We implemented Lisex from model formally written and verified in Coq. We used abstract data types (ADT) to implement some data structures. Hence, we show the methodology that we have applied to program from formal specifications using ADTs.

Keywords: MLS, Linux, Formal Methods, Confidentiality, Trojan Horses, Abstract Data Types.

1 Introduction

Lisex 0.0 is a Linux prototype planned to assess the usability and compatibility level of a UNIX-like operating system when a high secure access control mechanism is included. The stronger access control mechanism is meant to increase the resistance of the system to attacks against confidentiality performed by means of Trojan horses in installations where information security is, at least, very important. Resistance to this kind of attacks is gained by including a multi-level secure (MLS) model inside the operating system [1, 9]. Thus, the most important goal of the project was to include a MLS model in the Linux kernel and to assess the compatibility and usability of the resulting system. A second goal was to evaluate and fine tune a development process based on formal methods.

The MLS model included in Lisex 0.0 [5] is an adaptation to the Linux file system of the widely known Bell and LaPadula security model [3, 4]. We decided to incorporate the MLS model just in the file system because in UNIX-like operating systems most of the entities are managed through this subsystem. In this way, by modifying the file system interface we readily get a prototype that imposes MLS controls in most of the system.

The implementation of Lisex 0.0 started from a system and security formal model described and verified in Coq [2]. Although specifications contain enough information for programmers, it was necessary to describe a design before the coding phase because some abstractions do not directly fit in the Linux design. In this intermediate phase we decided to include abstract data types (ADT) in order to be able to change the implementation of new data structures with minimum impact on the rest of the system.

This paper is organized as follows. The first two sections are introductory: section 2 is a brief introduction to multi-level security, and section 3 describes the architecture of the Linux file system.

*Grupo de Investigación y Desarrollo en Ingeniería de Software y Seguridad

Sections 4 and 5 constitute the core of this article. The former describes the design and implementation of Lisex, the latter describes how we used the formal model as implementation guide. Finally, section 6 compares this article with similar approaches, and section 7 contains our conclusions.

2 Multi-level Security Concepts

As usual, we define Computer Security as the combination of confidentiality, integrity, and availability. Confidentiality requires that only authorized users read protected information; integrity requires that only authorized users modify protected information and only by the authorized means; and finally, availability requires that users can access information every time they need it [1, 9]. Our present work deals only with confidentiality.

In a computer system, the most lethal attacks against confidentiality are those conducted using Trojan horses. A Trojan horse is a program that will most often appear to provide some desired or usual function to serve as a lure to attract the program into use by an unsuspecting user, and a covert function to attack the system [1]. Today's mainstream operating systems, including the most popular UNIX flavors, lack security mechanisms to avoid these kind of attacks or at least reduce its frequency or damage [13]. Moreover, it is impossible to reach an acceptable level of resistance to these attacks without deeply changing the philosophy of protection [9]. That is, if an unsuspecting user executes a Trojan horse in these operating systems, the program then becomes a process like any other one, authorized to request the same services as a benign program. In this case, the Trojan horse will use the set of permissions granted to the user firstly to obtain the target information and secondly, to copy or deliver it to the attacker (for example, it could either copy the information in `/tmp` or e-mail it). This is possible, at least in part, because ordinary operating systems enforce a discretionary access control (DAC); that is, an access control policy where ordinary users are authorized to change security attributes.

In the early '70s, researchers at The Mitre Corporation devised a security model known as BLP [3, 4] that features a high level of resistance to attacks against confidentiality performed by means of Trojan horses. Furthermore, this model represents an abstraction, generalization and formalization of the USA Department of Defense (DoD) security policy for handling sensitive information. From that time on, security models that in one way or another represent a generalization of the DoD's security policy are known as multi-level security (MLS) models.

BLP classifies system entities into objects and subjects. Objects are those system resources or data repositories that must be protected (such as files, directories, terminals, and printers). Subjects are the active entities of the system, that is those entities capable of requesting system services (typically they are processors, process, users, etc.). Since our work deals only with a file system, from now on we will talk of processes, users, individuals and subjects as synonymous; and of files, directories or objects in the same sense. Then, BLP associates an *access* or *security class* to each object and subject. Access classes can be modified just through a highly controlled procedure and performed by the most trusted personnel –i.e. BLP is a mandatory access control (MAC) model. The structure of an access class is made up of two parts [9]:

Security level consisting of one of a few names such as *TOPSECRET*, *SECRET*, *CONFIDENTIAL*, *UNCLASSIFIED*

Category set consisting of zero or more names like *NATO*, *CIA*, *NUCLEAR*, etc.

Access classes are represented with ordered pairs; for example $(SECRET, \{NATO, NUCLEAR\})$ has *SECRET* level and category set $\{NATO, NUCLEAR\}$. The set of security levels must be linearly ordered, for instance:

$$UNCLASSIFIED < CONFIDENTIAL < SECRET < TOPSECRET$$

On the other hand, categories are independent of each other and they are not ordered. Access classes classify information and users by their confidentiality or responsibility degree. This is accomplished by defining a partial order relation over the set of access classes. Access class (n_1, C_1) *dominates*, written \succeq , access class (n_2, C_2) if and only if $n_1 \geq n_2$ and $C_1 \supseteq C_2$.

The idea behind this scheme is to set the confidentiality level of information and the trust level of users managing it: the “higher” the access class the more confidential is the information, and more trust is required and given to the individuals working with it.

In terms of DoD’s security policy, an individual can read a document if and only if the access class of the former dominates that of the last. When this access control rule is formalized in a security model it is known as *simple security* [9].

Bell and LaPadula noted that preserving simple security in a computer system is not easy. Processes have complete control of their memory spaces and the operating system kernel has no way to monitor what processes do with their data. Because of this fact, Bell and LaPadula required another property if multi-level security is to be preserved at all. This extra property today is known as *confinement property* but originally it was named **-property*. An informal statement of confinement is as follows:

- A process may open a file in read mode if and only if the access class of every other resource opened in write mode dominates the access class of the requested file
- A process may open a file in write mode if and only if the access class of every other resource opened in read mode is dominated by the access class of the requested file

The intention behind this rule is to avoid that once a Trojan horse has opened a file for reading it would be able to open a lower file for writing and thus it has the chance to copy the information from the first into the second. It is obvious that confinement is too restrictive because it forbids many legal flows.

3 The Architecture of the Linux File System

One of the most important features of Linux is that it can use several file system formats (such as EXT2, FAT, MINIX, iso9660, reiserfs, etc.). Each of them is separated from the operating system by a software layer known as Virtual File System (VFS). The VFS provides an abstraction inside the kernel which allows Linux to manage many different file system implementations. In turn, these file systems show a common software interface to the VFS. Details of each file system are translated by the VFS so that the rest of the kernel and the application software can use any file no matter where it is stored. From now on, if a file system can be managed by the VFS we will refer to it as a *physical file system* (PFS).

3.1 The Virtual File System (VFS)

The VFS keeps in memory information about each PFS that is being used, and updates these data when they are changed due to creation, modification, or deletion of files or directories.

The VFS uses a data structure called *inode* to represent files, directories, devices, and other entities. This structure stores enough information to read and modify these entities, and data such as permissions, access dates, etc.

Furthermore, each inode includes functions that allows the VFS to transparently interact with the particular PFS to which it belongs to. Some of these functions are grouped in a data structure called `inode_operations`, which is made up of a number of pointers to functions provided by the module that implements the corresponding PFS.

The VFS implements system calls, such as `stat`, `open`, `creat`, `chmod`, `chown`, etc. Usually, these system calls receive a file name, and the VFS translates it into an inode, which may or may not be in memory. If this inode is not in memory, then the VFS, using an inode operation, requests to the corresponding PFS to map its internal representation into an inode.

3.2 The Second Extended File System (EXT2)

The EXT2 file system is the most widely used PFS in the Linux community. It defines the file system topology by describing each file with a data structure called, rather confusingly, `ext2_inode` which is similar to the one used by the VFS. This likeness allows an efficient interaction between EXT2 and VFS. EXT2 inodes have a fixed size and contain data such as file owner and group, and pointers to the information stored in the disk.

4 Lisex's Design and Implementation

In this section we will show how we mapped some of the features of the security model [5] onto the Linux file system, and the design decisions that we took in order to implement the model. As we have said in section 2 our security model is an adaptation of BLP. This means that we have to implement access classes for files, directories and users, as well as simple security and confinement properties.

Every time we identified a possible modification we tried to isolate the entities likely to be changed behind abstract interfaces. It is worth noting that this technique is seldom used in the Linux kernel implementation. This approach allow us to select simple, less error prone, although inefficient implementations that can later be optimized without a sensible impact on the rest of the system.

Since the C programming language does not provide sophisticated mechanisms to prevent programmers from accessing modules' secrets, we had to instruct them to be disciplined in using just the modules' documented interfaces.

This section is organized as follows: first we show the changes at the VFS level by describing some of the new data structures and some modifications to existing system calls as well as the new ones; at the end of the section we give a brief account of the changes introduced at the EXT2 level.

4.1 Changes at the VFS Level

The VFS is the software module in charge of mediating the access to information requested by processes. Thus, it is imperative to modify its implementation in order to make it enforce the new security model. The most important design decision was to consider the content of files and directories as the set of objects to be protected. Clearly, this decision lefts unprotected many other entities (such as

access dates and DAC attributes) managed by the file system that can be used as high bandwidth covert channels. However, we persisted in our decision because Lisex 0.0 is no more than an experimental prototype. To protect files and directories implies to modify the inode data structure in order to add new access control attributes.

We choose to enforce security with a rather traditional approach, that is, we codified new access rules in system calls. Hence, part of the development involved the modification of several system calls. The existence of new security attributes, and the fact that applications need to consult and modify some of these attributes, made it necessary to include new system calls with these responsibilities.

4.1.1 Modifications to the inode Data Structure

We added two fields to the `inode` data structure in order to manage its ACL and security class (SC):

```
struct inode {
    ...
    struct vfs_acl  vfs_acl;           [a]
    sc              *sc;
}
```

The data type of each field is implemented as an ADT. Next, we describe with some detail the design and implementation of `sc` (see more in [10]).

The `sc` ADT. The data structure used to implement `sc` is shown below:

```
struct sc {
    unsigned int level;
    int categories[MAX_CAT_LEN];     [b]
    int size;
}
```

where

level: Access level.

categories: An array implementing the category set.

size: The amount of valid entries in `categories`.

`MAX_CAT_LEN` is a static definition used to limit the amount of elements in a category set. We have defined the ADT interface so that a change to this or other secrets has no impact on `sc`'s clients—for instance, it would be possible to use a linked list instead of an array. Table 1 summarizes `sc`'s interface.

Functions `sc_cat_first`, `sc_cat_next` and `sc_getcat` use the `sc_cat_iterator` type. Since eventually it could be necessary to implement the category set with a more efficient data structure, and that always will be necessary to iterate over it, we defined an iterator or selector independent of the particular implementation of this set. The value of `it` passed to `sc_cat_next` and `sc_getcat` must be obtained by calling `sc_cat_first` or `sc_cat_next`. For example, if the set is implemented with a linked list, `sc_cat_first` could return a pointer to the first node, `sc_get_cat` could get the value of that node, and `sc_cat_next` could return a pointer to the next node. One more advantage

sc ADT interface	
Function	Comment
<code>sc_init()</code>	Returns an initialized security class (SC)
<code>sc_destroy(sc)</code>	Releases the resources used by the SC <code>sc</code>
<code>sc_setlevel(sc, l)</code>	Sets <code>sc</code> level to <code>l</code>
<code>sc_getlevel(sc)</code>	Returns the level of <code>sc</code>
<code>sc_getsize(sc)</code>	Returns the amount of entries used in <code>sc</code>
<code>sc_cat_first(sc)</code>	Returns an iterator or selector pointing to the first valid category stored in <code>sc</code>
<code>sc_cat_next(sc, it)</code>	Returns the selector that points to the next valid category with respect to <code>it</code> stored in <code>sc</code>
<code>sc_getcat(sc, it)</code>	Returns the category pointed by <code>it</code>
<code>sc_clearcat(sc)</code>	Deletes all the categories of <code>sc</code>
<code>sc_addcat(sc, cat)</code>	Adds a category to the category set of <code>sc</code>

Table 1: Parameter `sc` is always a variable of type `sc`. This interface is defined in `sc.h`.

of using this iterator is that buffer overflows are confined to the iterator functions, making it easier to verify that they do not occur.

4.1.2 Changes to Existent System Calls

Changes to system calls were introduced due to the new access control model which, by incorporating ACLs and SCs to files, directories, and user accounts, requires to modify their semantic. Some system calls that have been modified are listed below¹:

```
long sys_chmod(const char * filename, mode_t mode)
long sys_open(const char * filename, int flags, int mode)
long sys_creat(const char * pathname, int mode)
long sys_stat(char * filename, struct __old_kernel_stat * statbuf)
```

It is worth noting that the signature of each system call is exactly the same as the original. Then, it would be possible to run any Linux application on Lisex.

Next, we will show with some detail how we have incorporated the new access controls to the `open` system call. We choose this system call given its relevance in security models rooted in BLP. Any process must invoke `open` before it could be able to access a file. Hence, this call implements one of the most important functions of the reference monitor concept [9].

In Linux, `open` calls an auxiliary function named `permission`, with the responsibility to make DAC checks. In fact, within this function the kernel decides to call an inode operation to perform a file system-dependent DAC check; or to call a VFS auxiliary function, called `vfs_permission`, to perform a (possible) different, system-wide verification.

In our implementation, `open` invokes a new auxiliary function named `may_open`. This function calls `permission` and `mac_permission` in that order. Lisex's implementation of `permission` removes the alternative of calling a file system-dependent function. In this way, our implementation mandates the kernel to use a system-wide ACL-based access control, which is implemented with a modified version of `vfs_permission`.

¹For a complete list, see [10].

`mac_permission` implements the necessary preconditions so that `open` preserves simple security and confinement (section 2). The first precondition is implemented rather straightforwardly with a `sc`'s interface function, called `sc_compare`. This function tells whether the access class of the calling user (`current->fsuid`) dominates the access class of the file to be opened (`inode->sc`).

```
int mac_permission (struct inode *inode, int mode)
{
    if(sc_compare(inode->sc, subjectsaget(current->fsuid)))           [c]
        return -EACCES;
    ...
}
```

The implementation of the precondition to enforce confinement was in itself the most complex task we performed. Hence, we opted to write code that is easy to verify but rather inefficient, with the intention of reducing the causes of errors (see section 5). Confinement requires to implement different controls depending on the mode in which the requested file is to be opened; but it is always necessary to compare the access class of this file against the access class of files already opened by the calling process. This fact led us, quite naturally, to write two clearly different code fragments, one for each mode. However, given that Linux holds two different lists of files that can be accessed by a given process, we divided these two fragments in two sections each. In consequence, the implementation of this precondition is made of four code fragments two of which, [d] and [e], are shown below.

Let us consider a process requesting a file to be opened in `READ` mode. Thus, in order to preserve confinement, it is necessary to iterate over all the files opened in `WRITE` mode (`file->f_mode`), verifying in each case that the access class of each of them dominates the access class of the file to be opened (i.e. `o_inode->sc`). As we have mentioned above, this verification must be done over two different lists. The first list holds the open files (`current->files->fd`):

```
...
if(mode & FMODE_READ)
{
    files = current->files;
    max_fds=files->max_fds;
    for(fd=0; fd<max_fds; fd++)
        if((file=files->fd[fd]))
            if (file->f_mode & FMODE_WRITE)           [d]
            {
                o_inode = file->f_dentry->d_inode;
                if(sc_compare(inode->sc, o_inode->sc))
                    return -EACCES;
            }
    ...
}
```

While the second list records the memory mapped files (represented by the nodes of the list `current->mm->mmmap` that point to files in disk, i.e. `mmmap->vm_file` not zero):

```

...
if(current->mm)
    for(mmap = current->mm->mmap;mmap;mmap = mmap->vm_next)
        if((file=mmap->vm_file))
            if (file->f_mode & FMODE_WRITE)
            {
                o_inode = file->f_dentry->d_inode;
                if(sc_compare(inode->sc, o_inode->sc))
                    return -EACCES;
            }
...

```

[e]

4.1.3 New System Calls

We added new system calls to change or consult ACLs and SCs of files and users. Table 2 contains a brief description of each of them.

System Call	Comment
<code>aclstat(filename, acl_statbuf, len)</code>	Returns the ACL of file <code>filename</code> in <code>acl_statbuf</code>
<code>acladd(filename, id, mode, type)</code>	Adds permissions stored in <code>mode</code> to the user or group <code>id</code> in the ACL of <code>filename</code>
<code>acldel(filename, id, mode, type)</code>	Removes <code>mode</code> from the set of permissions of user or group <code>id</code> from the ACL of <code>filename</code>
<code>oscstat(filename, level, cats, len)</code>	Returns the level and the category set (in <code>level</code> and <code>cats</code> , respectively) of the access class of <code>filename</code>
<code>chobjsc(filename, level, cats, size)</code>	Changes the access class of <code>filename</code> for the access class represented by <code>level</code> and <code>cats</code>
<code>ownerclose(pid, filename)</code>	If process <code>pid</code> has opened <code>filename</code> then removes it from its list of open files

Table 2: New system calls.

Next we describe the implementation of `chobjsc`. This system call is used to change security classes of files and directories. `chobjsc` is secure if the following two preconditions are met:

1. The user who issues the call must be a MAC administrator.

```

asmlinkage long
sys_chobjsc (char *filename, int level, int *categories, int size)
{
    ...
    if(!is_secadm())
        return -EACCES;
    ...
}

```

[f]

`is_secadm()` checks whether category `SCADMIN` belongs to the category set of the calling user.

2. The file whose access class is about to be changed must not be open.

```
...
error = -EBUSY;
if (is_open (nd.dentry->d_inode))
    return error;
...
```

[g]

`is_open(i)` checks whether the file pointed to by inode `i` is being used by some process.

Once all preconditions are met, the level and category set received as input are set as the new access class of the object using the interface of the `sc` ADT:

```
...
sc_clearcat (inode->sc);
sc_setlevel (inode->sc, level);
error = sc_usergetncat (inode->sc, categories, size);
...
}
```

[h]

where `sc_usergetncat` copies the category set `categories` of size `size`, from user space into the access class `inode->sc`, in kernel space.

4.2 Changes at the EXT2 Level

With the present design the VFS mandates PFSs to make new security attributes persistent. Thus, a PFS had to be modified if a MLS model will be implemented. We choose to modify EXT2 because it is the most widely used PFS in the Linux community. Hence, it was necessary to modify the structure of the EXT2 inode, and the inode operations responsible of reading from and writing to disk. By taking advantage of some advanced EXT2 features, it was possible to include this modifications without any trouble.

More precisely, we changed the `ext2_inode` data structure as follows:

```
struct ext2_inode {
    ...
    struct ext2_acl_entry ext2_acl[EXT2_ACL_LEN];
    __u32 sc_level;
    __u32 categories[EXT2_CAT_LEN];
    __u32 cat_len;
}
```

[i]

where

`ext2_acl`: The ACL of the file. `ext2_acl_entry` is defined as follows:

```
struct ext2_acl_entry
{
    __u32    id;
    __u8     mode;
}
```

[j]

`sc_level`: Level of the access class of the file.

`categories`: Category set of the access class of the file.

`cat_len`: The amount of valid entries in `categories`.

The reader may note that in the `ext2_inode` representation we do not use the `sc` ADT, because `ext2_inode` must define the physical structure of the inode, hence there is no chance to abstract it.

The inode operations responsible for translating an inode from the EXT2 level to the VFS level (and vice-versa) are `ext2_read_inode` and `ext2_update_inode`. `ext2_read_inode` is responsible for building, from the new fields of `ext2_inode`, the ACL and the access class (`sc`) of the VFS inode. Next we show the translation of `sc`:

```
void ext2_read_inode (struct inode * inode)
{
    struct ext2_inode * raw_inode;
    ...
    cat_len = le32_to_cpu(raw_inode->cat_len);
    sc_setlevel(inode->sc, le32_to_cpu(raw_inode->sc_level));           [k]
    for(i=0; i<cat_len; i++)
        sc_addcat(inode->sc, le32_to_cpu(raw_inode->categories[i]));
    ...
}
```

Some fields of `inode` are set by the VFS before calling `ext2_read_inode`. These fields are used to find the corresponding EXT2 inode (`raw_inode`). With this information and using the appropriate interface, the ACL and the SC are stored in `inode`.

5 Programming from Formal Specifications

One of the project's goals is to apply formal methods during the software life cycle. We have applied formal methods to describe and verify the system and security formal models [5], then we used the system model as test case source [7], and finally we used it as an implementation guide. In the first two activities The Coq Proof Assistant was used². In this paper we will concentrate on how we used the system model as an implementation guide.

The system formal model is an abstract state machine representing the VFS interface. The state of this machine is an abstraction of the file system state, and state transitions are system calls. Thus, we want to show how the model guided us in order to implement a state variable and a state transition operation (i.e. a system call). One of the most important state variables is the function that maps objects onto access classes, called *objectSC*. Next we will show the process we followed to implement *objectSC*, and the operation *Open*.

5.1 Implementing a State Variable

From now on, we will describe some portions of the formal model using first order logic instead of the original language; otherwise it would be necessary to explain Coq's syntax and semantics and that

²Coq is a proof assistant based on the Calculus of Inductive Constructions, which in turn is a form of Type Theory.

is out of the scope of this article. In the formal model *objectSC*, the function that maps files and directories onto access classes, has the following type:

$$objectSC : OBJECT \mapsto SecClass$$

where

- *OBJECT* is the set or type denoting all the possible system's objects (i.e. files and directories)
- \mapsto denotes a finite partial function
- *SecClass* is the set or type of access classes, defined as follows:

$$SecClass \cong Record\{level : SECLEV; categs : \mathbb{P} CATEGORY\}$$

where *SECLEV* is a finite ordered set denoting all the possible security levels, and *CATEGORY* is a finite set that designates all the categories.

Hence, *objectSC* is a finite partial function from objects onto access classes. It is partial because in a given state of the system not every object is present in the file system. *objectSC* is implemented by including a variable of type `sc` in the VFS `inode` data structure, which is the internal representation of a file. Since the kernel preserves a functional relation between inodes and files, we can guarantee that the implementation verifies the properties of \mapsto .

Let us inspect *SecClass* in more detail. Its definition suggests that the implementation should have two variables. The first one should vary over a type denoting a finite ordered set, thus it is natural to select `unsigned int` for it. The second variable should implement a finite set. However, given that sets are not a basic C type, we needed to program functions preserving the essential mathematical properties of finite sets. In this way, we decided to implement *categs* as an array of type `int`, plus the functions listed in Table 1. For example, `sc_addcat` adds a category to the array avoiding duplicates, and `sc_compare` compares two access classes by taking account of the elements present in the arrays but not their positions. In consequence, *SecClass* formalization guided us toward an implementation based on an ADT.

5.2 Implementing a State Transition Operation

We will show the relation between the operation *Open* and its implementation as a system call. Again, we choose this operation due to the importance it has in BLP based models. In Lisex 0.0, *Open* implements both DAC and MLS models, but here we will focus just in the latter. The operations of the system formal model are described by their pre and postconditions. Preconditions are used to guarantee that operations execute only when their execution do not put the system at risk, and postconditions establish the state change that must be performed. Thus, the formal specification of *Open* is as follows³:

$$Open(s, u?, o?, m?, t) \cong OpenRead(s, u?, o?, m?, t) \vee OpenWrite(s, u?, o?, m?, t)$$

where *s* and *t* are the start and next state respectively, user *u?* tries to open object *o?* in mode *m?*, and

³Here we will show a simplified version of the operation, but it captures its essence anyway.

$$\begin{aligned}
OpenRead(s, u?, o?, m?, t) &\hat{=} \\
& m? = READ \\
& \wedge preSimpleSecurity(s, u?, o?) \\
& \wedge preStarPropertyRead(s, u?, o?) \\
& \wedge postOpenRead(s, u?, o?, t)
\end{aligned}$$

$$\begin{aligned}
OpenWrite(s, u?, o?, m?, t) &\hat{=} \\
& m? = WRITE \\
& \wedge preSimpleSecurity(s, u?, o?) \\
& \wedge preStarPropertyWrite(s, u?, o?) \\
& \wedge postOpenWrite(s, u?, o?, t)
\end{aligned}$$

Clearly, the formal specification suggests to implement two code fragments, one for each term of the disjunction. Although this strategy (possibly) generates inefficient code, it is less error prone or, at least, the program is easier to verify given its similarity with the specification. We focus on *OpenRead*, thus we expand *preSimpleSecurity*:

$$\begin{aligned}
preSimpleSecurity(s, u?, o?) &\hat{=} \\
& s.objectSC(o?).level \leq s.subjectSC(u?).level \\
& \wedge s.objectSC(o?).categs \subseteq s.subjectSC(u?).categs
\end{aligned}$$

Now, the specification asserts that the access class of the calling user must be compared with the access class of the object to be opened in the current state of the system. This comparison is performed with `sc_compare` (see Table 1). Program fragment [c] implements this part of the specification.

Next we expand *preStarPropertyRead*:

$$\begin{aligned}
preStarPropertyRead(s, u?, o?) &\hat{=} \\
& \forall b \in \text{dom } s.openFiles \bullet \\
& \quad u? \in s.openFiles(o?).ActWriters \\
& \quad \Rightarrow s.objectSC(o?).level \leq s.objectSC(b).level \\
& \quad \wedge objectSC(o?).categs \subseteq objectSC(b).categs
\end{aligned}$$

To see the relation between *preStarPropertyRead* and the implementation pay attention to code fragments [d] and [e]. The universal quantification in *preStarPropertyRead* states that any implementation must iterate over all the open files (of the requesting process). However, the antecedent of the implication says that the iteration must be done just over the files opened in *WRITE* mode. This is implemented with a decision inside the loop in [d] and [e] where its condition determines whether *u?* has opened the file *b* in that mode (also in [d] and [e]). Finally, the inner sentences of the `if` structure implement the comparison between the access classes of *o?* and *b*. When this comparison is not right, `mac_permission` returns with a convenient error.

We did not make any modification to the current postcondition because is it the same as the one we specified.

We believe that the specification structure is so similar to the code structure that it is possible to put more confidence in the correctness of the implementation. However, since our implementation is mixed with code written without precise specifications, the correctness of Lisex 0.0 depends on the correctness of Linux.

6 Related Work

The enhancement of UNIX-like operating system with MLS controls dates back to 1975 when the UCLA Data Secure UNIX project was carried on. At least the articles we were able to check [16, 18], differ from the present work in (a) the UCLA project addressed the construction of a complete security kernel, when, to date, we only deal with the file system; (b) they report the specification and verification of models but they do not focus on the implementation source code; (c) they do not directly implement a BLP-like security model (despite their model could be adapted to verify Bell-LaPadula). Further, their formal, abstract models are described and verified with languages and tools quite different from The Coq Proof Assistant. In [18] it is shown how the UCLA team constructed part of a proof of correctness of the implementation source code with respect to a low-level specification. Our approach (shown in section 5) is less rigorous than the one applied by Walker and others, but many times less expensive. On the other hand, they addressed, like us, the broader problem of producing a secure system with formal methods.

In 1984 Kramer [12] proposes to enhance the UNIX operating system with MLS controls and some other security features. In that paper the author accounts for a number of desired security features to incorporate into the UNIX kernel but he never describes a low level design, internal data structures or a formal model. Moreover, the work described in that article addresses a complete system, thus audit, personnel, and authentication issues are considered. Linus IV and Lisex are both based on the Bell-LaPadula security model, and both systems implement ACLs. Also, both systems attempts to separate the responsibilities held by `root` in different user accounts (not shown in our present paper, see [6]). Kramer's work tries to preserve the UNIX semantics and compatibility as much as possible, goals that we too tried to achieve.

Two years after the publication of Kramer's paper Gligor and others presented the design and implementation of Secure Xenix Workstations [11]. Despite that Gligor's paper shows a more detailed design than Kramer's, it does not show source code, kernel internal data structures or any other details at implementation level. Furthermore, they do not mention a formal model on which the implementation is based -although they mention the Bell and LaPadula security model. Once more, this paper refers to a complete system where many fetarures not present in Lisex were taken into account.

There are other similar projects that had been conducted by some UNIX vendors but we were unable to find detailed information about the design or implementation of their systems.

None of the aforementioned projects seems to be related with open source software, nor with a BLP-like system model of UNIX (or at least with a model of a UNIX-like file system) formally verified with a proof asistant.

More recently, the NSA⁴ has started a project called Security-Enhanced Linux (SELinux) [14, 15, see also <http://www.nsa.gov/selinux>] which has many things in common with Lisex but also some differences. Obviously both try to enhance Linux with mandatory security controls, but NSA's work differs from ours in that (a) they enhance security controls by adding a Linux kernel module and not by a direct kernel re-implementation, and (b) they incorporate a general security server rather than a direct implementation of some MLS model. This server can be replaced by others that implement different security models. The example security server delivered with the current version of SELinux does not work with a MLS model, instead it implements a form of the Clark-Wilson integrity model, Type Enforcement, and IBAC. On the other hand, despite the security arquitecture of SELinux is

⁴National Security Agency, U.S.A.

rooted in a formal model developed for the DTOS operating system [17] we found no evidence that this model is being used in the current development.

7 Conclusions

For the past couple of years it seems that the security community has a renewed interest in operating systems capable of preserving the confidentiality of information but without losing compatibility nor usability. For this reason, we decided to implement a MLS model in the kernel of an open source operating system. Eventually we opted for Linux because it is widely used. We are strongly convinced that by choosing to modify an existing operating system rather than to build a new one from scratch, we saved a lot of time and effort, and we gained a lot of experience. Moreover, this strategy allowed us to count in a few months with a prototype on which to prove our ideas, validate our methodology, and analyze the level of usability and compatibility of the system.

Our first conclusion is that the usability level of a direct implementation of a BLP-like security model is too low for most users of modern computing environments. Moreover, our experiments suggest that by moving to an information flow model [8] we can increase the level of usability without losing security. We hope to count in a few months with Lisex's successor, called GIDISS Trusted Linux, which will implement such a model.

If Linux did not exist, our work would be impossible. However, the lack of design documentation and the absence of clearly defined abstract interfaces, slowed down our progress and sometimes forced us to focus on issues not directly related with security. Because of this, we planned to use ADTs and to document every development phase. We noted that the use of ADTs notably improves the readability of code, is less error prone, and reduces the impact on future changes. At the same time, we are convinced that by choosing right implementations, abstract interfaces need not to incur in performance penalties. Furthermore, the inclusion of ADTs allowed us to attack the complexity of the problem in two stages: first, by adding part of the required functionality as modules' interfaces; and second, by choosing very simple implementations for each of them so that a prototype was readily available.

Availability of the system formal model proved to be crucial for the success of the project. It was useful as design and implementation guide, and as test oracle during the testing phase. Programmers were able to program faster avoiding many errors, as it is shown by the testing performed until now. Moreover, the model allowed us to perform a disciplined, rigorous, and structured code review, which increased our confidence in the correctness of the implementation. Future versions will be developed from formal models because we are convinced that the time and effort spent on that phase is fully returned during implementation and testing. It is important to mention that all except one of the people involved in the project, were undergraduate students. This, obviously, was not an impediment for success.

We invite the reader to check our web site (<http://www.fceia.unr.edu.ar/gidis>) for news about GIDISS Trusted Linux (Lisex's successor).

References

- [1] Marshall D. Abrams, Sushil Jajodia, and Harold J. Podell. *Information Security: an integrated collections of essays*. IEEE Computer Society press, 1995.
- [2] Bruno Barras and et. al. *The Coq Proof Assistant Reference Manual*. INRIA, 1999.

- [3] D. Elliot Bell and Leonard LaPadula. Secure computer systems: Mathematical foundations. MTR 2547, The MITRE Corporation, May 1973.
- [4] D. Elliot Bell and Leonard LaPadula. Secure computer systems: Mathematical model. ESD-TR 73-278, The MITRE Corporation, November 1973.
- [5] Maximiliano Cristiá. *Formal verification of an extension of a secure, compatible UNIX file system*. MSc thesis, Departamento de Computación, Universidad de la República, Uruguay, 2002.
- [6] Maximiliano Cristiá. *Manual del usuario y del administrador de Lisex 0.0*. Grupo de Investigación y Desarrollo en Ingeniería de Software y Seguridad, <http://www.fceia.unr.edu.ar/gidis>, 2002.
- [7] Maximiliano Cristiá, Martín Degrati, and Pablo Garralda. *Testing: casos de prueba a nivel del modelo para Lisex 0.0*. Grupo de Investigación y Desarrollo en Ingeniería de Software y Seguridad, <http://www.fceia.unr.edu.ar/gidis>, 2002.
- [8] Dorothy E. Denning. A lattice model of secure information flow. *Communications of the ACM*, 19(5):236–243, May 1976.
- [9] Morrie Gasser. *Building a Secure Computer System*. Van Nostrand Reinhold, 1988.
- [10] Gisela Giusti, Felipe Manzano, and Maximiliano Cristiá. *Manual de referencia para el programador de Lisex 0.0*. Grupo de Investigación y Desarrollo en Ingeniería de Software y Seguridad, <http://www.fceia.unr.edu.ar/gidis>, 2002.
- [11] Virgil D. Gligor, E. L. Burch, C. S. Chandrasekaran, R. S. Chapman, L. J. Dotterer, M. S. Hecht, W. D. Jiang, G. L. Luckenbaugh, and N. Vasudevan. On the design and the implementation of Secure Xenix Workstations. In *1986 IEEE Symposium on Security and Privacy*, pages 102–117, 1986.
- [12] Steven Kramer. Linus IV—an experiment in computer security. In *1984 IEEE Symposium on Security and Privacy*, pages 24–32, 1984.
- [13] P. A. Loscocco and et. al. The inevitability of failure: The flawed assumption of security in modern computing environments. In *21st National Information Systems Security Conference*, pages 303–314, October 1998.
- [14] Peter Loscocco and Stephen Smalley. Integrating flexible support for security policies into the Linux operating system. In *Proceedings of the FREENIX Track of the 2001 USENIX Annual Technical Conference*, 2001.
- [15] Peter Loscocco and Stephen Smalley. Meeting critical security objectives with Security-Enhanced Linux. In *Proceedings of the 2001 Ottawa Linux Symposium*, 2001.
- [16] G. J. Popek and D. A. Fabre. A model for verification of data security in operating systems. *Communications of the ACM*, 21(9):737–749, September 1978.
- [17] Secure Computing Corporation. DTOS generalized security policy specification. Technical report, Secure Computing Corporation, 1997.
- [18] B. J. Walker, R. A. Kemmerer, and G. J. Popek. Specification and verification of the UCLA Unix security kernel. *Communications of the ACM*, 23(2):118–131, February 1980.