



Università
Ca' Foscari
Venezia

Second Cycle Degree Programme in
Computer Science (Class LM-18)

MASTER THESIS

Automatic Verification of Grsecurity RBAC Policies

SUPERVISOR

Prof. Riccardo FOCARDI

AUTHOR

Marco SQUARCINA

Matriculation Number 814359

CA' FOSCARI
DORSODURO 3246
30123 VENICE

ACADEMIC YEAR
2013/2014

CA' FOSCARI UNIVERSITY OF VENICE

Abstract

Department of Environmental Sciences, Informatics and Statistics

Master Degree

Automatic Verification of Grsecurity RBAC Policies

by Marco SQUARCINA

Role-based Access Control (RBAC) is one of the most widespread security mechanisms in use today. Given the growing complexity of policy specifications arising from system administrators' needs, formally verifying that access control systems enforce some defined security invariants is a problem of crucial importance. In a paper [1] which has been presented at the *25th IEEE Computer Security Foundations Symposium*, we developed a framework for the formal verification of `grsecurity`, an access control system developed on top of the Linux kernel. In this thesis we present this work and improve the framework by considering the interaction with the underlying operating system. This refinement allows for a reduction in the number of transitions within the labelled transition systems resulting from policies. Additionally, we deal with the problem of automatic verification of `grsecurity` policies by defining a set of security invariants. Based on our abstract semantics, we implement `Granalyze`, a model checker that accounts for the verification of real policies. We report on the results of the experimental analysis conducted on existing public servers running `grsecurity`.

Acknowledgements

First of all, I would like to express my most sincere gratitude to my supervisor Professor Riccardo Focardi for the trust placed in me over the years and for his friendly guidance. I wish to thank Professor Flaminia Luccio for her continuous support and for her patience. Thanks also to all the Professors met during the Bachelor and Master degrees for their inspiring influence.

I want to thank Dr. Stefano Calzavara without whom this work would not have been possible. I am also thankful to my friends who supported me in writing, and motivated me to strive towards my goal: Claudio Bozzato, Matteo Centenaro, Gian-Luca Dei Rossi, Emanuele Giaquinta, Francesco Palmarini, Marco Signori, Mauro Tempesta and all the tremendous people who participated in the Hacking Group of our department.

A special thank to my family: my father Renzo, my mother Luciana, my sister Margherita, my aunt Gabriella and my uncle Sandro, my grandmother Natalia and my grandfather Nerino who left us during the writing of this thesis. At the end I would like to thank my beloved Paola who spent sleepless nights with me and was always my support. Words cannot describe how much I love you.

Contents

Abstract	i
Acknowledgements	ii
Contents	iii
List of Figures	v
List of Tables	vi
1 Introduction	1
1.1 Contributions	3
1.2 Structure of the Thesis	4
2 Role Based Access Control Models	5
2.1 A Common Terminology	5
2.2 Early Formal Models	6
2.3 RBAC	8
2.3.1 RBAC 92	8
2.3.2 RBAC 96	11
2.3.3 Standard RBAC	13
3 Background on grsecurity	15
3.1 Grsecurity RBAC	15
3.2 RBAC Policies	16
3.3 User and Group Identifiers	17
4 A Formal Semantics for grsecurity	19
4.1 Policies	19
4.2 Pathnames and Matching	21

4.3	Role Assignment	22
4.4	Semantics	24
5	Verification of grsecurity Policies	26
5.1	An Abstract Semantics for Grsecurity	26
5.2	Correlating the Two Semantics	30
5.3	Security Analysis	32
6	Gran: a Tool for Policy Verification	37
6.1	Implementation Overview	37
7	Granalyze: an Improved Security Analyser	40
7.1	An Updated Framework	41
7.1.1	Concrete Semantics	41
7.1.2	Abstract Semantics	41
7.2	Granalyze	42
8	Case Studies	46
8.1	Verification of Existing Policies	46
8.2	Exploits Through “setuid” Binaries	48
8.3	Information Leakage Analysis	50
9	Conclusion	52
A	Grsecurity Policy Modalities	54
	Bibliography	57

List of Figures

2.1	RBAC relationships	8
2.2	Example of a role hierarchy	10
2.3	Family of models in RBAC 96	12
2.4	RBAC 96 structure	12
6.1	<code>gran</code> help message	39
7.1	<code>granalyze</code> help message	45

List of Tables

2.1	A simple access matrix	6
2.2	NIST RBAC levels	14
3.1	A snippet from a <code>grsecurity</code> policy	16
4.1	Semantics of <code>grsecurity</code>	23
5.1	Abstract semantics of <code>grsecurity</code>	29
7.1	Updated semantics of <code>grsecurity</code>	42
7.2	Updated abstract semantics of <code>grsecurity</code>	43
8.1	A snippet of a flawed <code>grsecurity</code> policy	49
A.1	Role Modes	54
A.2	Subject Modes	55
A.3	Object Modes	56

Dedicated to Paola...

Chapter 1

Introduction

Role-based Access Control (RBAC) is one of the most widespread security mechanisms in use today and has been the subject of extensive research for more than a decade now. The central idea in the RBAC model is to factor the assignment of access rights into two steps, separating the distribution of permissions to system-specific roles, from the assignment of users to roles, so as to simplify the overall access control management task [2].

Most of the research work on RBAC (see, e.g., [3–5]) has focused on policy verification, a problem of critical importance for system administrators, and a challenging one due to the complexity of the policies to be verified and to the state changes that arise in their management. State-change is not specific to RBAC: traditional access control frameworks such as those studied in the seminal work of [6] include rules to affect the structure of the access control matrix, defining the permissions granted to subjects on objects. Modern administrative RBAC (ARBAC) systems present similar features by providing system administrators with expressive languages for manipulating RBAC policies by re-assigning users to roles and/or modifying the assignment of permissions to roles [2].

Policy verification in access control systems has traditionally been stated as a safety question, answered by means of a reachability analysis: for instance, user-role reachability in ARBAC systems formalizes the problem of determining whether, given an initial policy state, a target user and a role, there exists

a sequence of state changes leading to a state in which the target user is impersonating that role. As it turns out, this kind of analysis is challenging, as the procedural nature of state change languages often creates subtle, undesired effects that are hard to anticipate without the aid of a tool for analysis.

Model checking [7] has emerged as a promising technique for automated policy verification [5, 8, 9]. The idea is exactly as in program verification, with the set of state-change rules playing the role of the program to be tested, and the reachability question as the property of interest: to counter the state explosion that often affect the analysis, making it unscalable to the point of making the problem intractable [9], researchers have advocated the usage of abstraction techniques [10].

In the present thesis we continue along this line of research, focusing our attention on the formal verification of `grsecurity` [11], an access control system developed on top of the Linux kernel. `grsecurity` is deployed as a patch to the OS kernel that installs a reference monitor to mediate any access to the underlying OS resources; it supports the definition and dynamic enforcement of fine-grained access control policies to let users operate on objects (resources) via the subjects (executable files) provided by the underlying file system. Users are organized in roles, which are in turn structured as to identify a subset of privileged roles with higher capabilities on the system resources, and administrative control of the access control policies.

The verification problem in `grsecurity` presents much of the complexity of Administrative RBAC systems due to the presence of policy state changes: these may arise either from explicit administrative actions for manipulating users and roles, as well as from the interaction between `grsecurity`'s access control and the facilities provided by the underlying operating system for setting user ids, hence dynamically changing users and associated roles by executing binaries operating in *setuid* mode [12]. This dependency from state changes on the executable binaries of the underlying file system further complicates the model checking problem, as it causes the size of the search space to grow unbounded in the number of states and transitions.

We tackle the problem by resorting to an abstraction technique, by which the behavior resulting from the unbounded set of subjects available in the underlying file system is captured by the finite number of subjects that are listed in the security policy, which represents the input of the model checker. We prove the abstraction sound and complete, and employ it to carry out a reachability analysis on RBAC policies targeted at unveiling (potential) security leaks, leading to unintended accesses to sensitive resources.

1.1 Contributions

The contribution of this thesis may be summarized as follows:

- we develop a formal semantics for **grsecurity**'s RBAC system, based on a labelled transition system; besides providing the fundamental building block for our analysis, the LTS semantics has proved interesting in itself, as it made possible to understand the subtleties of **grsecurity**'s RBAC rules, and to unveil a flaw arising from the interplay between the access control systems supported by Linux and **grsecurity**. As we discuss in Section 4.4, this flaw makes it possible to unexpectedly bypass the imposed **grsecurity** capability restrictions when executing a **setuid/setgid** binary [13];
- we introduce an abstract semantics which provides a bounded, yet sound and complete, representation of the dynamic evolution of the **grsecurity** policy states arising in the formal semantics; based on that, we develop a framework for reachability analysis aimed at detecting the presence of access leaks in any given policy;
- we implement our framework in **gran**, a tool for the automatic analysis of **grsecurity** policies: the tool takes as input an RBAC policy, a user u , a set of initial states for u (associated with the possible subjects that may impersonate u) and a target file / object o , and checks whether there is a path of state changes leading to a state that grants u access to o ;
- we introduce a refined version of the framework which matches with the latest release of **grsecurity** and allows to reduce the number of transitions

within the labelled transition by considering the interaction with the underlying operating system; we implement this framework in **granalyze**, a tool for policy analysis that aids to identify the minimal Trusted Computing Base (TCB) which should be considered in order to satisfy a given set of security invariants;

- we provide a report of experiments we conducted with the analysis of policies in use on existing, commercial servers running **grsecurity** to implement their RBAC systems.

1.2 Structure of the Thesis

Chapter 2 recalls the main role-based access control models; Chapter 3 reviews the basic concepts and notions behind **grsecurity**; Chapter 4 presents our formal semantics of the **grsecurity** RBAC system; Chapter 5 describes the abstraction for the verification of **grsecurity** policies, and shows its formal correspondence to the previous semantics; Chapter 6 describes **gran** (**grsecurity** analyzer), a tool that automatically looks for security leaks in real **grsecurity** policies; Chapter 7 presents a refined version of the framework and introduces **granalyze**; Chapter 8 illustrates **gran** at work on some case studies; Chapter 9 concludes with final remarks and a discussion of related and future works.

Chapter 2

Role Based Access Control Models

The field of computer security started to grow rapidly with the introduction of the early time-sharing computers in the 1970s. Access control became soon a topic of crucial importance in government and military applications where strong security was required. At the same time, researchers started to analyze the problem providing a mathematical and formal description of access control models [14].

In this chapter the main access control models are briefly introduced along with the motivations that led to the development of RBAC as an alternative to the well established MAC and DAC approaches.

2.1 A Common Terminology

Access control models can be formally specified by means of entities and the relations between them. The main notions are *users*, *sessions*, *subjects*, *objects*, *operations* (or *actions*) and *permissions* [14].

The term *user* refers to people interacting with the computer system. An instance of a user's interaction with the system is called *session*. A single user is

often allowed to perform multiple logins under different IDs. *Subjects* are the active entities, usually processes, acting on behalf of an user. An *object* can be viewed as a computer resource like a file or a device. Objects are traditionally passive entities that contain or receive information. Subjects interact with objects through *operations* depending on the *permissions* (or *privileges*) given to the subject. *Permissions* are often represented as a combination of objects and operations [14].

2.2 Early Formal Models

One of the first formal definitions of an access control model is due to Lampson [15], who introduced an *access matrix* to mediate the access of subjects to objects. In this representation, each row defines a subject, whilst objects are arranged over the columns. Thus, the element in position (i, j) identifies the permissions assigned to the subject i on the object j . An example is depicted in Table 2.1: subject A is assigned read and write permissions on o_1 , read access on o_2 and write access on o_3 . B can only read o_1 and o_3 .

	o_1	o_2	o_3
A	r,w	r	w
B	r		r

TABLE 2.1: A simple access matrix

During the 1970s, Bell and LaPadula [16] provided a mathematical model of the military multilevel security policies focused on *clearance levels*. Two basic properties are defined: the *simple security property* and the **-property*, commonly known, respectively, as *no read up* and *no write down*. The first property states that a subject cleared at a given level should not read objects above that level, e.g., an user with **confidential** clearance should not read **top-secret** information. The **-property* defines that subjects can only write objects at their own level or above, e.g., an user cleared at the **secret** level cannot write **confidential** objects, but is allowed to gain write access on **top-secret** information. Additionally, in the BLP model every user must also be entitled to all the categories associated with an object. For instance, in order to read a document classified as

[**secret**, **nuclear**, **NATO**], a subject must be cleared at a level equal or above **secret** and be explicitly authorized to the **nuclear** and **NATO** categories. This kind of policy ensures that information cannot be downgraded through unintentional or malicious actions.

A milestone towards the standardization of access control models is constituted by the *Trusted Computer System Evaluation Criteria* (TCSEC) [17] (also known as the *Orange Book*) published by the U.S. Department of Defense (DoD) in which two important models for military access control systems were introduced in detail: *Mandatory Access Control* (MAC) and *Discretionary Access Control* (DAC). In the DAC model, the access of a subject to an object is permitted depending on the identity of the subject or on the group to which the subject belongs to. The access control is *discretionary* due to the fact that a subject with a given set of permissions is allowed to pass those permissions to another subject, e.g., the owner of a file can make that resource readable by anyone. On the contrary, the MAC model recalls the principles of *multilevel security policy* pioneered by the BLP. In systems implementing MAC policies, access to resources is always mediated by the access control system using rules imposed externally. Users are thus limited in the actions they can take, avoiding leakages of permissions to other subjects or users.

Despite the efforts to promote the adoption of TCSEC-compliant systems as security solutions for commercial organisations, most commercial firms recognized that DAC and MAC were not sufficient for their needs. As Clark and Wilson [18] pointed out, TCSEC-oriented systems are focused on information flow and secrecy, but the main concern of commercial organisations involves the integrity of information. Therefore, they introduced a new model for business security practices hinged on the concepts of well-formed transactions and separation of duty (SoD) to account for data integrity and consistency of changes of critical data. Implementing these rules in a computer system resulted impractical, since it turned out to be as challenging as implementing information flow policies [14].

2.3 RBAC

A NIST¹ study conducted in the early 1990s showed that access control needs of commercial and government organizations were not met by products available at that time [19]. For many enterprises the actual owner of system resources is the organization, thus a discretionary policy on the users' part was not appropriate. On the other hand, traditional MAC systems implemented via multilevel policies, as mentioned before, were focused on confidentiality resulting inadequate to fulfill the organizations' needs. There existed a general demand for subject-based security policies to restrict access depending on the function or role of an user within the enterprise [14].

2.3.1 RBAC 92

In 1992, Ferraiolo e Kuhn [20] addressed the problem by introducing a novel access control model based on roles which extended and formalized some of the functionalities found in applications already available at the time, such as databases and ATM systems. In this seminal paper, roles are defined as set of permissions. As depicted in Figure 2.1, users are assigned to roles, thus users' permissions are only those of the roles to which users are identified to.

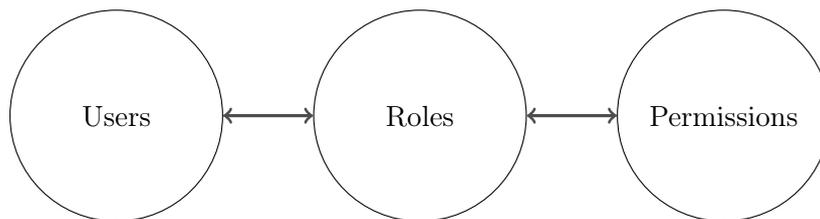


FIGURE 2.1: RBAC relationships

¹U.S. National Institute of Standards and Technology

Three basic rules are required [20]:

1. *Role assignment*

A subject can execute a transaction only if the subject has selected, or been assigned to, a role. The identification and authentication process (e.g., login) is not considered a transaction. All other user activities on the system are conducted through transactions. Thus, all active users are required to have some active role.

2. *Role authorization*

A subject's active role must be authorized for the subject. Together with rule 1, this rule ensures that users can take on only roles for which they are authorized.

3. *Transaction authorization*

A subject can execute a transaction only if the transaction is authorized for the subject's active role. With rules 1,2, this rule ensures that users can execute only transactions for which they are authorized.

A simple formal description of the aforementioned concepts in terms of sets and relations is depicted below.

Definition 2.1. RBAC 92 formal model

- For each subject, the active role is the one that the subject is currently using

$$AR(s : subject) = \{\text{the active role for subject } s\}$$

- Each subject may be authorized to perform one or more roles

$$RA(s : subject) = \{\text{authorized roles for subject } s\}$$

- Each role may be authorized to perform one or more transactions

$$TA(r : role) = \{\text{transactions authorized for role } r\}$$

- Subjects may execute transactions. The predicate $exec(s, t)$ is *True* if the subject s can execute the transaction t at the current time, otherwise it is *False*:

$$exec(s : subject, t : tran) = True \Leftrightarrow s \text{ can execute transaction } t$$

- Role assignment rule:

$$\forall s : subject, t : tran \cdot exec(s, t) \Rightarrow AR(s) \neq \emptyset$$

- Role authorization rule:

$$\forall s : subject \cdot AR(s) \subseteq RA(s)$$

- Transaction authorization rule:

$$\forall s : subject, t : tran \cdot exec(s, t) \Rightarrow t \in TA(AR(s))$$

A key concept in the RBAC 92 model is that roles are hierarchical, i.e., roles can inherit permissions from other ones. This feature enable the definition of security policies closely related to organizations' actual structure. Figure 2.2 presents an example of an hierarchical relationship within an hospital [20].

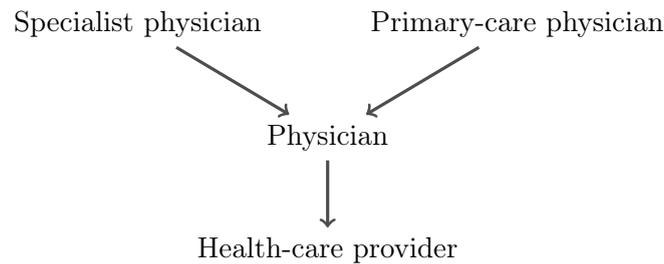


FIGURE 2.2: Example of a role hierarchy

One of the most common method of implementing access control in computer systems is through *access control lists* (ACLs). ACL is a mechanism that implements access control for an object by associating it with list of user that are permitted to access the resource and stating the access modes granted to each

user [21]. Whilst this approach allows to easily identify the users that retain access to a given object, it is impractical to list the objects that a given user is allowed to access. Additionally, in most computer systems, ACL are specified also using groups (i.e., collections of users), thus access to resources is determined by the IDs of both users and groups. It follows that ACLs make the operation of adding a permission to an object easy, but revoking all the permissions is hard due to the users/groups interplay.

The RBAC 92 model, on the contrary, requires that access to resources is performed exclusively through roles. Roles, indeed, are not collections of users, but set of permissions. In the enterprises, roles are often statically defined, whilst users and permissions are liable to change over the time. Having identified the static and dynamic entities of the system, the administration of an access control policy is greatly simplified due to the fact that, most of the time, a change in the policy simply involves moving an user to a different role or redefining the permissions of a given role.

2.3.2 RBAC 96

Four years after the introduction of RBAC 92, Sandhu et al. presented a modular framework for role based access control systems (RBAC 96) [22]. This framework breaks down RBAC in four conceptual models depending on the features included (see Figure 2.3). $RBAC_0$ is the base model that defines the minimal set of features of a system implementing RBAC. $RBAC_1$ and $RBAC_2$ include $RBAC_0$ plus support for, respectively, hierarchies and constraints such as SoD. $RBAC_3$ includes all aspects of the lower-level models.

Figure 2.4 depicts the main aspects of RBAC 96. Relations UA and PA are pivotal to the model since they define, respectively, the assignment of users and permissions to roles: an user could be a member of many roles and a role could have many users associated; similarly, a role can enlist a set of permissions and a permission can be granted to many roles. The term *session*, in this model, identifies the set of active roles for a given user. User's permissions are then the permissions defined for all his active roles within a session. Each session is tied

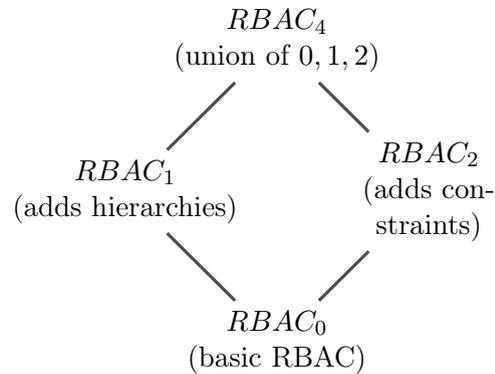


FIGURE 2.3: Family of models in RBAC 96

to a single user, but an user is allowed to start multiple concurrent sessions, each associated to a different combination of active roles.

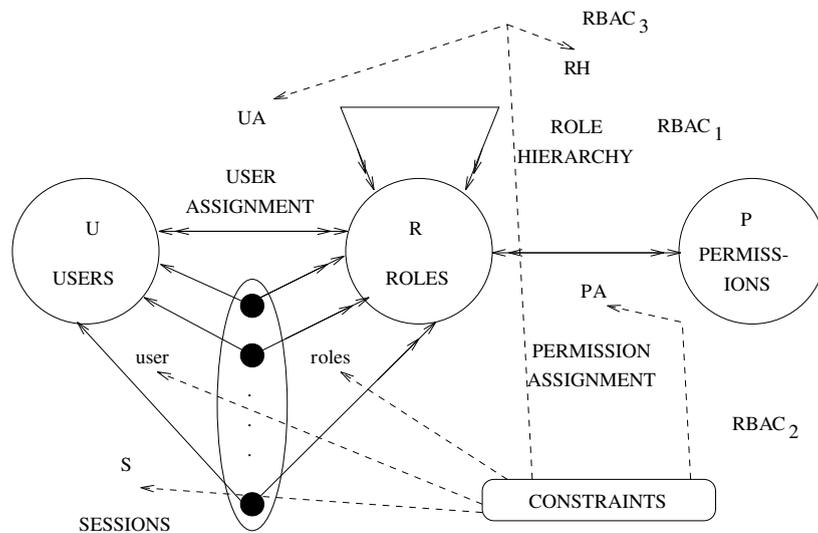


FIGURE 2.4: RBAC 96 structure

The following is a formal description in terms of sets and relations of the base model $RBAC_0$ as originally formulated [22].

Definition 2.2. $RBAC_0$ components

- U, R, P, S
respectively the sets of users, roles, permissions and sessions,

- $PA \subseteq P \times R$
a many-to-many permission to role assignment relation,
- $UA \subseteq U \times R$
a many-to-many user to role assignment relation,
- $user : S \rightarrow U$
a function mapping each session s_i to the single user $user(s_i)$ which is constant for the session's lifetime,
- $roles : S \rightarrow 2^R$
a function mapping each session s_i to a set of roles $roles(s_i) \subseteq \{r \mid (user(s_i), r) \in UA\}$ which can change with time, and session s_i which has the permissions $\bigcup_{r \in roles(s_i)} \{p \mid (p, r) \in PA\}$.

2.3.3 Standard RBAC

In 2000, NIST initiated a standardization process of the RBAC model publishing a proposal [23] developed on top of RBAC 96. The resulting model, known as NIST RBAC, is made up of levels with increasing functional capabilities. The four levels are Flat RBAC, Hierarchical RBAC, Constrained RBAC and Symmetric RBAC; their main features are presented in Table 2.2. This framework has been adopted as an ANSI/INCITS standard in 2004 [24].

Level	Name	Functional Capabilities
1	Flat RBAC	<ul style="list-style-type: none"> • users acquire permissions through roles • must support many-to-many user-role assignment • must support many-to-many permission-role assignment • must support user-role assignment review • users can use permissions of multiple roles simultaneously
2	Hierarchical RBAC	<ul style="list-style-type: none"> • Flat RBAC • must support role hierarchy (partial order) • level 2a requires support for arbitrary hierarchies • level 2b denotes support for limited hierarchies
3	Constrained RBAC	<ul style="list-style-type: none"> • Hierarchical RBAC • must enforce separation of duties (SoD) • level 3a requires support for arbitrary hierarchies • level 3b denotes support for limited hierarchies
4	Symmetric RBAC	<ul style="list-style-type: none"> • Constrained RBAC • must support permission-role review with performance effectively comparable to user-role review • level 4a requires support for arbitrary hierarchies • level 4b denotes support for limited hierarchies

TABLE 2.2: NIST RBAC levels

Chapter 3

Background on grsecurity

`grsecurity` is a patch for the Linux kernel focused on security at the operating system level. It provides many different features on latest stable kernels, implementing a “detection, prevention, and containment” model [11]. In addition to the role-based access control (RBAC) system, which is the focus of this thesis, `grsecurity` offers protection mechanisms against privilege escalation, malicious code execution and memory corruption; it also implements an advanced auditing system. `grsecurity` is typically adopted by hosting companies to harden web servers and systems providing services to locally logged users [25].

3.1 Grsecurity RBAC

`grsecurity` complements the standard discretionary access control (DAC) mechanism provided by Linux with a form of mandatory RBAC, providing an additional layer of protection. In the rest of the thesis we identify `grsecurity` with its RBAC system.

The specification of the access control requirements is provided by a *policy*, whose structure is described in Section 3.2. The policy defines the available roles, which can be of four different types. *User* roles are an abstraction of standard users in Linux systems, i.e., they provide a hook to extend the traditional DAC permission system with more sophisticated mechanisms, available only in `grsecurity`.

Group roles provide a similar device for actual groups of the system. *Special* roles, instead, are not directly associated to traditional users and groups, and they are intended to provide extra privileges to normal accounts. A *default* role applies when no user, group, or special role can be granted. The mechanism of role assignment is discussed in Section 4.3.

3.2 RBAC Policies

A policy defines the permissions given to each role for the different *objects* stored in the file system. A further level of granularity is introduced through the standard notion of *subject*, i.e., an abstraction of a process. Namely, permissions are not directly assigned to roles, since this would lead to a very coarse form of access control; rather, permissions are defined for pairs of the form role-subject. For instance, user `alice` could be granted read access to the object `/var` only through the subject `/bin/ls`.

```

role alice u {
  role_transitions professor
  subject / {
    /
    /bin                x
    /boot               h
    /dev                h
    /dev/null           w
    /dev/pts            rw
    /dev/tty            rw
    /etc                r
  }

  subject /bin/su {
    user_transition_allow root
    group_transition_allow root
    /                  h
    /bin               h
    /bin/su            x
    /dev/log           rw
  }

```

TABLE 3.1: A snippet from a grsecurity policy

Table 3.1 presents a snippet from a `grsecurity` policy. Even if it does not show all the features provided by `grsecurity` RBAC, it allows us to introduce the most important elements considered in our formalization. The policy defines a user role (flag 'u') `alice`, which is permitted to impersonate the special role `professor`. Transitions to specific users and groups of the underlying Linux system can be allowed or forbidden at the subject level, e.g., by the `user_transition_allow` attribute.

Permissions are specified in terms of access modalities for the objects in the policy. In this case, any process executed by `alice` is assigned the permissions defined for subject `/`, except for process `/bin/su` which specifies its own set of modalities. In general, any process is accredited a set of access rights for any object in the system, according to a hierarchical matching mechanism. For instance, subject `/bin/su` inherits the rights on `/etc` by the less specific subject `/`, while it overrides the permissions for `/bin` with its own. Similarly, accesses to object `/dev/log` by subject `/` are resolved in terms of the modalities listed for the less specific object `/dev`. Complete details on this mechanism are provided in Section 4.2.

The modalities we consider mirror standard Linux permissions for reading 'r', writing 'w' and executing 'x', plus a hiding mode 'h'. Subjects are completely unaware of the presence of any hidden object, e.g., the process `/bin/ls` does not even list the directory `/boot` when it is launched by `alice`. We ignore other available modalities, which are either irrelevant for our setting (e.g., 'p' for ptrace rejection) or identifiable with one of the previous modalities (e.g., 'a' for appending).

Appendix A lists the available modalities for roles, subjects and objects supported by `grsecurity`.

3.3 User and Group Identifiers

Before digging into the internals of `grsecurity`, we need to briefly review how users and groups are identified in Linux systems. At the kernel level, users and

groups are not distinguished by names, but by numbers. We refer to these numbers as user identifiers (UIDs) and group identifiers (GIDs) respectively. When a process is started, Linux assigns it a pair of identifiers, set to the UID of the invoking user. These identifiers are called the *effective* UID and the *real* UID of the process, respectively. The effective UID determines the privileges granted to the process and is employed, e.g., for standard DAC enforcement; the real UID, instead, affects the permissions for sending signals.

This apparently simple mechanism is complicated by an important subtlety related to the execution of particular binaries in the file system. Namely, any file f may be granted the “setuid” permission, with the following effect: when f is executed, the effective UID of the process is set to the UID of the *owner* of f , irrespective of the UID of the invoking user; the real UID, instead, is set to the UID of the caller. This allows for temporary acquisition of additional privileges to perform specific tasks.

We conclude by pointing out that changing to a particular UID is considered a sensitive operation in Linux systems and requires the process to possess the *capability* `CAP_SETUID`. Capabilities provide finer-grained distribution of privileges among processes since Linux 2.2. Remarkably, capabilities are bypassed when a “setuid” binary is executed, i.e., a process spawned by a “setuid” binary is *always* allowed to set its effective UID to the UID of the owner.

All the previous discussion applies similarly to GIDs.

Chapter 4

A Formal Semantics for grsecurity

We propose a formal semantics for `grsecurity` in terms of a labelled transition system. We write $f : A \rightarrow B$ when f is a total function from A to B , while we use $f : A \mapsto B$ when f is partial. We let $f(a) \downarrow$ denote that f is defined on a . Let $f : A_1 \times \dots \times A_n \mapsto B$ and let a_i range over A_i , for any $k \leq n$ we stipulate $f(a_1, \dots, a_k) \downarrow$ if and only if $\exists a_{k+1}, \dots, \exists a_n : f(a_1, \dots, a_n) \downarrow$. Finally, we let $\mathcal{P}(A)$ stand for the power set of A .

4.1 Policies

We presuppose denumerable sets U of users and G of groups, ranged over by u and g respectively. We also let T denote the set of role types $\{\mathbf{u}, \mathbf{g}, \mathbf{s}\}$ ranged over by t ; C denote the set of capabilities $\{\mathbf{set_uid}, \mathbf{set_gid}\}$ and M the set of access modalities $\{\mathbf{r}, \mathbf{w}, \mathbf{x}, \mathbf{h}\}$. A policy P is a 8-tuple:

$$P = (R, S, O, perms, caps, role_trans, usr_trans, grp_trans),$$

where:

- R is a set of roles, ranged over by r . We let R_t denote the set of roles of type t and we assume that R_t and $R_{t'}$ are disjoint whenever $t \neq t'$;
- S is a set of subjects, ranged over by s , and O is a set of objects, ranged over by o . Both subjects and objects are pathnames, as we discuss below;
- $perms : R \times S \times O \mapsto \mathcal{P}(M)$ defines the permissions granted by the policy. Namely, if $m \in perms(r, s, o)$, then subject s running on behalf of role r has permission m on object o ;
- $caps : R \times S \mapsto \mathcal{P}(C)$ determines the capabilities allowed by the policy, i.e., if $c \in caps(r, s)$, then subject s running on behalf of role r can acquire capability c ;
- $role_trans : R \rightarrow \mathcal{P}(R_s)$ defines which special roles can be impersonated by a given role;
- $usr_trans : R \times S \mapsto \mathcal{P}(U)$ defines which user identities can be assumed by a subject running on behalf of a given role;
- $grp_trans : R \times S \mapsto \mathcal{P}(G)$ defines which group identities can be assumed by a subject running on behalf of a given role.

We require a number of well-formedness constraints on policies which formalize a corresponding set of syntactic checks performed by `grsecurity`. Recall that we write $perms(r, s) \downarrow$ to denote $\exists o \in O : perms(r, s, o) \downarrow$.

1. $\forall r : perms(r, /) \downarrow$, i.e., all roles define at least the subject “/”;
2. $\forall r, \forall s : (perms(r, s) \downarrow \Rightarrow perms(r, s, /) \downarrow)$, i.e., every subject in every role defines at least the object “/”;
3. there exists a default role “-” such that $\forall t : - \notin R_t$.

Throughout this thesis, most definitions (notably, the semantic rules in Tables 4.1 and 5.1) and notation are to be understood as parametric with respect to a given policy. To ease readability, we do not make such dependency explicit, and just assume P as the underlying policy instead.

4.2 Pathnames and Matching

Subjects and objects are collectively represented within `grsecurity` policies as pathnames, and these, in turn, are defined as sequences of “/”-separated names (or wildcards) as customary in Unix systems. For ease of presentation, we henceforth disregard wildcards and assume the following simplified structure of pathnames (that always presupposes a trailing “/”). Let n be a non-empty string non including “/”, and let “.” note string concatenation. Pathnames are defined by the following productions:

$$p ::= / \mid / \cdot n \cdot p$$

Pathnames are ordered according to the standard prefix order, so that p is smaller (more specific) than, or equal to, p' whenever p' is a prefix of p . Formally, the ordering relation, noted \sqsubseteq , is the smallest relation closed under the following rules:

$$\begin{array}{c} \text{(P-TOP)} \\ p \sqsubseteq / \end{array} \qquad \begin{array}{c} \text{(P-PATH)} \\ \frac{p \sqsubseteq p'}{/ \cdot n \cdot p \sqsubseteq / \cdot n \cdot p'} \end{array}$$

Clearly, \sqsubseteq is a partial order: this ordering is paramount in `grsecurity`, as it constitutes the basic building block underlying the mechanisms for associating subjects to processes, and for checking access rights on objects. Specifically, when a process spawned by the execution of a file f running on behalf of a role r tries to access a file f' , `grsecurity` matches f against the most specific subject s defined in role r such that $f \sqsubseteq s$. Similarly, f' is matched against the most specific object o , defined in subject s of role r , such that $f' \sqsubseteq o$. The permissions of o are then retrieved to evaluate whether the process can be granted access to f' . For instance, according to the policy in Table 3.1, process `/bin/cat` is granted read access to `/etc/fstab`, since `/bin/cat` matches the subject “/” and `/etc/fstab` matches the object `/etc` defined there.

We formalize the matching relation as follows. For any set A of path names, we let $\min(A)$ denote the minimum element of A according to the ordering \sqsubseteq , whenever such an element exists. Given a pathname p , we define the *matching*

subject for p in role r as

$$\text{match_subj}(p, r) = \min(\{s \mid p \sqsubseteq s \wedge \text{perms}(r, s) \downarrow\}).$$

Analogously, we define the *matching object* for p in role r under subject s as

$$\text{match_obj}(p, r, s) = \min(\{o \mid p \sqsubseteq o \wedge \text{perms}(r, s, o) \downarrow\}).$$

Proposition 4.1 below and the assumption of well-formedness of the policy imply that $\text{match_subj}(p, r)$ is always defined; instead, $\text{match_obj}(p, r, s)$ is defined only if $\text{perms}(r, s) \downarrow$.

Proposition 4.1 (Chain Property).

$$\text{If } p \sqsubseteq p' \text{ and } p \sqsubseteq p'', \text{ then } p' \sqsubseteq p'' \text{ or } p'' \sqsubseteq p'.$$

Proof. By induction on the sum of the depths of the derivations of $p \sqsubseteq p'$. Base case is $p \sqsubseteq / = p'$ which by (P-TOP) implies $p'' \sqsubseteq p'$. Inductive case is when $p \sqsubseteq p'$ since $p = / \cdot n \cdot \hat{p}$ and $p' = / \cdot n \cdot \hat{p}'$ with $\hat{p} \sqsubseteq \hat{p}'$. Now, if p'' is $/$ we trivially have $p' \sqsubseteq p''$. Otherwise, since $p \sqsubseteq p''$ by (P-PATH) it must be $p'' = / \cdot n \cdot \hat{p}''$ with $\hat{p} \sqsubseteq \hat{p}''$. By induction we have $\hat{p}' \sqsubseteq \hat{p}''$ or $\hat{p}'' \sqsubseteq \hat{p}'$ that, by applying (P-PATH), gives the thesis. \square

4.3 Role Assignment

Each process in `grsecurity` has a role and a subject attached to it. The assignment of the subject to the process is performed by matching the name of the running file against the list of subjects of the current role, as discussed in Section 4.2. Roles, instead, are assigned according to the hierarchy “special - user - group - default”. Special roles are granted through authentication to the `gradm` utility and are intended to provide extra privileges to normal user accounts: as such, they have the highest priority. User roles, instead, are applied when a process either is executed by a user with a particular UID or changes to that UID. This is possible, since the name of every user role must match up with the name

$$\begin{array}{c}
\text{(SETR)} \\
\frac{\hat{r} = \text{role}(r_s, u, g) \quad r'_s \in \text{role_trans}(\hat{r}) \cup \{-\}}{\langle r_s, u, g, f \rangle \xrightarrow{\text{set_role}(r'_s)} \langle r'_s, u, g, f \rangle} \\
\\
\text{(SETU)} \\
\frac{\hat{r} = \text{role}(r_s, u, g) \quad s = \text{match_subj}(f, \hat{r}) \\
u' \in \text{usr_trans}(\hat{r}, s) \quad \text{set_uid} \in \text{caps}(\hat{r}, s)}{\langle r_s, u, g, f \rangle \xrightarrow{\text{set_UID}(u')} \langle r_s, u', g, f \rangle} \\
\\
\text{(SETG)} \\
\frac{\hat{r} = \text{role}(r_s, u, g) \quad s = \text{match_subj}(f, \hat{r}) \\
g' \in \text{grp_trans}(\hat{r}, s) \quad \text{set_gid} \in \text{caps}(\hat{r}, s)}{\langle r_s, u, g, f \rangle \xrightarrow{\text{set_GID}(g')} \langle r_s, u, g', f \rangle} \\
\\
\text{(EXEC)} \\
\frac{\hat{r} = \text{role}(r_s, u, g) \\
s = \text{match_subj}(f, \hat{r}) \quad o = \text{match_obj}(f', \hat{r}, s) \\
\mathbf{x} \in \text{perms}(\hat{r}, s, o) \quad \mathbf{h} \notin \text{perms}(\hat{r}, s, o) \\
u' \in \text{usr_trans}(\hat{r}, s) \cup \{u\} \quad g' \in \text{grp_trans}(\hat{r}, s) \cup \{g\}}{\langle r_s, u, g, f \rangle \xrightarrow{\text{exec}(f')} \langle r_s, u', g', f' \rangle}
\end{array}$$

TABLE 4.1: Semantics of grsecurity

of an actual user in the system, i.e., there exists a bijective partial mapping from UIDs to user roles. It is worth noticing that only the *real* UID of the process is considered for role assignment. Group roles behave similarly to user roles, but they are applied to a given process only if no user role is associated to the process UID. The default role is chosen when no other role can be given.

A further remark is in order for role assignment: even though user roles are assigned by just looking at the real UID of the process, the presence of “setuid” binaries must be considered with care. We recall that a process spawned by a “setuid” binary sets its effective UID to the UID of the owner; however, *even unprivileged* (i.e., without the capability CAP_SETUID) processes can always set their real UID to their effective UID [12]. Binaries with the “setuid” permission set may then come into play during the role assignment process. As usual, similar considerations apply for “setgid” files.

4.4 Semantics

We assume an underlying file system, i.e., a subset of a denumerable set of pathnames F , ranged over by f . Let r_t range over $R_t \cup \{-\}$, a *state* is a 4-tuple $\sigma = \langle r_s, u, g, f \rangle$ describing a process spawned by the execution of file f . The process may be impersonating a special role (when $r_s \neq -$) and is running with real UID set to u and real GID set to g . We identify UIDs and GIDs with elements from a subset of U and from a subset of G , respectively. The role associated to σ is determined by the first three components of the tuple, according to the following function *role*:

$$\text{role}(r_s, u, g) = \begin{cases} r_s & \text{if } r_s \in R_s \\ u & \text{if } r_s \notin R_s, u \in R_u \\ g & \text{if } r_s \notin R_s, u \notin R_u, g \in R_g \\ - & \text{otherwise} \end{cases}$$

The function formalizes the role assignment process, according to the hierarchy discussed before.

Attacker Model Our semantics tracks all role transitions and subject changes allowed to a given process. The semantics depends on an underlying Linux system hosting `grsecurity`, characterized by a set of users, a set of groups and a file system, as it is apparent by the format of the states. However, we do not explicitly model any change to the previous sets and we just assume them to be denumerable; we can imagine to pick different sets after each transition, to account for the evolution of the system as a result of background operations. Intuitively, we consider a worst-case scenario, in which any possible action not conflicting with the RBAC policy is eventually performed by the process. Of course, the resulting LTS has an infinite number of states and transitions: this problem will be tackled in Chapter 5, where we will propose an abstract, finite-state semantics, specifically designed for automated security analysis.

Transitions The transition rules are reported in Table 4.1. Rule (SETR) accounts for login operations to special roles: such transitions must be allowed by the *role_trans* function. When $r'_s = -$, the rule models a logout from a special role, which is always permitted. Rule (SETU) describes a change of the process UID, which must be allowed by the *usr_trans* function; moreover, the process must possess the capability `set_uid`, as we discussed in Section 3.3. Notice that s is the matching subject for file f in the role \hat{r} associated to the current state. Rule (SETG) details a similar behavior for changing the process GID. Finally, rule (EXEC) accounts for the execution of files and is the most interesting rule. The invoked file must indeed be executable and it must not be hidden, since hidden files are not visible to unauthorized processes. The execution of the file may lead to a role change, as we explained in Section 4.3. Since we do not model which “setuid” and “setgid” binaries are actually present in the file system and we do not explicitly keep track of changes to file permissions, we simply assume that the execution of the file may trigger any user or group transition allowed by the policy for the current state. Of course, we also consider the possibility that the execution does not alter the identifiers of the process.

This subtle behavior when executing setuid/setgid programs was unknown before we started our formalization. In Chapter 8, we will illustrate that it is potentially harmful for security. This has also been reported to the main developer of `grsecurity`, who confirmed our findings. A fix has already been implemented in the latest stable release of `grsecurity` [13]. The solution consists in requiring the capabilities `CAP_SETUID/CAP_SETGID` to perform role transitions, even upon execution of setuid/setgid binaries. A formal semantics of `grsecurity` that accounts for the aforementioned fix is presented in Chapter 7.

Chapter 5

Verification of grsecurity Policies

While suitable for describing the operational behavior of `grsecurity`, the semantics presented in Chapter 4 is not amenable for security verification, as we discuss below. We thus propose a different semantics, designed for security analysis, which is an abstraction of the previous one, while being suitable to be model-checked. We also outline some properties of `grsecurity` policies which we consider interesting to verify and we formalize them in our framework.

5.1 An Abstract Semantics for Grsecurity

The main problem with the presented semantics is that it hinges on many elements specific to the underlying Linux system hosting `grsecurity`, i.e., users, groups and files. Remarkably, all these elements are inherently dynamic, so any changes to them must be accounted for in the semantics to get a sound tool for security analysis. As a result, the corresponding LTS has infinite states and transitions, making security verification difficult to perform, inaccurate, or even infeasible. We thus design a simple *abstract* semantics for `grsecurity`, depending only on the content of the policy, which can be reasonably assumed to be static. If the policy happens to change during the lifetime of the hosting system, we simply consider a different LTS and we perform again any relevant analysis.

We start from some simple observations. First, we note that users and groups are immaterial to **grsecurity**, as only the role assigned to a process is relevant for access control. Second, we observe that also the actual content of the file system is somewhat disposable, since all granted permissions are determined by finding out a matching subject or object. We thus define an abstract state as a 4-tuple $\sigma_a = \langle r_s, r_u, r_g, s \rangle$ describing a process spawned by the execution of some file $f \sqsubseteq s$. The role assigned to the process is again determined by the first three components of the tuple and can be retrieved by overloading the type of the function *role* defined previously.

We first abstract from impersonation of user identities. The intuition here is that, in general, only a subset of the users has an associated user role, according to the definition of the policy, and all other users can be identified by **grsecurity** to the special identity “-”. We thus define the abstraction of a user u , denoted by $\llbracket u \rrbracket$, as follows:

$$\llbracket u \rrbracket = \begin{cases} u & \text{if } u \in R_u \\ - & \text{otherwise} \end{cases}$$

We define the abstract version of the *usr_trans* function, noted $\llbracket \text{usr_trans} \rrbracket$, as the partial function with the same domain of *usr_trans* such that for, every r and s , we have:

$$\llbracket \text{usr_trans} \rrbracket(r, s) = \{ \llbracket u \rrbracket \mid u \in \text{usr_trans}(r, s) \}$$

In other words, transitions to users with no associated user role are collapsed to transitions to the special identity “-”. We introduce analogous definitions also for groups and group transitions.

$$\begin{aligned} \llbracket g \rrbracket &= \begin{cases} g & \text{if } g \in R_g \\ - & \text{otherwise} \end{cases} \\ \llbracket \text{grp_trans} \rrbracket(r, s) &= \{ \llbracket g \rrbracket \mid g \in \text{grp_trans}(r, s) \} \end{aligned}$$

We still need to address the most challenging task for the definition of the new semantics, i.e., the approximation of the behaviour of **grsecurity** upon file executions. The idea is to identify the executed file f with its matching object o : as a consequence of this abstraction, we can only find out an approximation for the

subject to assign to the new process. This is done in terms of a set of possible matches, elaborating on the following observations:

- since o is the matching object for f , then f must be at least as specific as o ($f \sqsubseteq o$). Thus, we can take as an *upper bound* for the new subject the most specific subject which is no more specific than o , i.e., the subject $\min(\{s' \mid o \sqsubseteq s'\})$. For instance, the execution of the file `/bin/ls`, matching the object `/bin/ls`, may lead to the impersonation of the subject `/bin` only if the more specific subject `/bin/ls` does not exist;
- since we do not know how much specific is f , every subject s' no more generic than o ($s' \sqsubseteq o$) may be a possible match. However, we can filter out all the subjects which would be associated to the execution of a more specific object o' which overrides o , i.e., we consider the set $\{s' \mid \text{match_obj}(s', r, s) = o\}$, where r and s identify the current role and subject. For instance, the execution of a file in `/bin`, matching the object `/bin`, may lead to the impersonation of subject `/bin/ls` only if there does not exist the object `/bin/ls`. Indeed, when object `/bin/ls` exists, the execution of the file `/bin/ls` matches `/bin/ls` and not the less specific object `/bin`. Note that the file `/bin/ls` may even be non-executable, according to the policy specification for the object `/bin/ls`.

This reasoning leads to the following definition of *image* of an object o , given a role r and a subject s :

$$\begin{aligned} \text{img}(o, r, s) = & \{s' \mid \text{match_obj}(s', r, s) = o\} \\ & \cup \{\min(\{s' \mid o \sqsubseteq s'\})\} \end{aligned}$$

Again Proposition 4.1 and the well-formation of the policy imply that such a notion is always well-defined.

We finally present in Table 5.1 the reduction rules for the abstract semantics.

Rule (A-SETR) is identical to rule (SETR), while rule (A-SETU) is the counterpart of (SETU), abstracting from the users of the system. When $r'_u = -$, the rule matches a transition to a user with no associated user role. Clearly,

$$\begin{array}{c}
\text{(A-SETR)} \\
\frac{\hat{r} = \text{role}(r_s, r_u, r_g) \quad r'_s \in \text{role_trans}(\hat{r}) \cup \{-\}}{\langle r_s, r_u, r_g, s \rangle \xrightarrow{\text{set_spec}(r'_s)}_a \langle r'_s, r_u, r_g, s \rangle} \\
\\
\text{(A-SETU)} \\
\frac{\hat{r} = \text{role}(r_s, r_u, r_g) \quad \hat{s} = \text{match_subj}(s, \hat{r}) \quad r'_u \in \llbracket \text{usr_trans} \rrbracket(\hat{r}, \hat{s}) \quad \text{set_uid} \in \text{caps}(\hat{r}, \hat{s})}{\langle r_s, r_u, r_g, s \rangle \xrightarrow{\text{set_user}(r'_u)}_a \langle r_s, r'_u, r_g, s \rangle} \\
\\
\text{(A-SETG)} \\
\frac{\hat{r} = \text{role}(r_s, r_u, r_g) \quad \hat{s} = \text{match_subj}(s, \hat{r}) \quad r'_g \in \llbracket \text{grp_trans} \rrbracket(\hat{r}, \hat{s}) \quad \text{set_gid} \in \text{caps}(\hat{r}, \hat{s})}{\langle r_s, r_u, r_g, s \rangle \xrightarrow{\text{set_group}(r'_g)}_a \langle r_s, r_u, r'_g, s \rangle} \\
\\
\text{(A-EXEC)} \\
\frac{\hat{r} = \text{role}(r_s, r_u, r_g) \quad \hat{s} = \text{match_subj}(s, \hat{r}) \quad \mathbf{x} \in \text{perms}(\hat{r}, \hat{s}, o) \quad \mathbf{h} \notin \text{perms}(\hat{r}, \hat{s}, o) \quad r'_u \in \llbracket \text{usr_trans} \rrbracket(\hat{r}, \hat{s}) \cup \{r_u\} \quad r'_g \in \llbracket \text{grp_trans} \rrbracket(\hat{r}, \hat{s}) \cup \{r_g\} \quad s' \in \text{img}(o, \hat{r}, \hat{s})}{\langle r_s, r_u, r_g, s \rangle \xrightarrow{\text{exec}(s')}_a \langle r_s, r'_u, r'_g, s' \rangle}
\end{array}$$

TABLE 5.1: Abstract semantics of grsecurity

rule (A-SETG) behaves in the same way for group roles. Finally, rule (A-EXEC) accounts for the execution of processes. Again, the choice of the new user and group role assumes a worst case scenario, in that every user and group transition which is allowed by the policy is taken into account by the rule. The new subject is drawn from the image of an executable object, according to the described approximation.

We conclude this section with two observations on the abstract semantics. First we note that, for any finite policy, the resulting LTS has a finite number of states and any state has a finite number of outgoing transitions, since both states and labels are built over finite sets. The LTS can then be effectively explored using standard techniques. We also underline our design choice to include states whose subject is not defined in the current role. Indeed, in our semantics we enforce an explicit match of the current subject against the subjects defined for the role.

This choice leads to an increment of the size of the LTS, since we introduce a number of somewhat equivalent states; however, such a decision allows for a much more accurate security analysis, as we discuss in Section 5.3.

5.2 Correlating the Two Semantics

We now prove that the abstract semantics in Table 5.1 is a sound approximation of the concrete semantics in Table 4.1, in that every transition in the concrete semantics has a corresponding transition in the abstract semantics.

Formally, we abstract a file f in terms of the most specific subject which is no more specific than f itself, i.e., we let $\llbracket f \rrbracket = \min(\{s \mid f \sqsubseteq s\})$. We can now define the abstraction of a concrete state $\sigma = \langle r_s, u, g, f \rangle$ as the abstract state $\llbracket \sigma \rrbracket = \langle r_s, \llbracket u \rrbracket, \llbracket g \rrbracket, \llbracket f \rrbracket \rangle$.

Proposition 5.1 (Identity Preservation). *The following equalities hold:*

- (i) $role(r_s, u, g) = role(r_s, \llbracket u \rrbracket, \llbracket g \rrbracket)$;
- (ii) $match_subj(f, r) = match_subj(\llbracket f \rrbracket, r)$.

Proof. For (i) we observe that $u = \llbracket u \rrbracket$ if $u \in R_u$ and $g = \llbracket g \rrbracket$ if $g \in R_g$. When, instead, $u \notin R_u$ and $g \notin R_g$ we have $role(r_s, u, g) = role(r_s, -, g) = role(r_s, \llbracket u \rrbracket, g)$ and $role(r_s, u, g) = role(r_s, u, -) = role(r_s, u, \llbracket g \rrbracket)$, giving the thesis. Item (ii) holds since $\{s \mid f \sqsubseteq s\} = \{s \mid \llbracket f \rrbracket \sqsubseteq s\}$, by definition of $\llbracket f \rrbracket$. \square

Lemma 5.2 (Abstract Execution). *If $match_obj(f, r, s) = o$, then $\llbracket f \rrbracket \in img(o, r, s)$.*

Proof. We first observe few, auxiliary properties. Let $p \sqsubseteq p'$, then one has:

- (a) $\llbracket p \rrbracket \sqsubseteq p'$ or $p' \sqsubseteq \llbracket p \rrbracket$;
- (b) $\llbracket p \rrbracket \sqsubseteq \llbracket p' \rrbracket$;
- (c) $match_obj(p, r, s) \sqsubseteq match_obj(p', r, s)$.

(a) follows directly by Proposition 4.1 from the observation that $p \sqsubseteq \llbracket p \rrbracket$; (b) and (c) follow immediately by noting that $\{\hat{p} \mid p' \sqsubseteq \hat{p}\} \subseteq \{\hat{p} \mid p \sqsubseteq \hat{p}\}$, by transitivity of \sqsubseteq .

We are now ready to prove the Lemma. We must show that either $match_obj(\llbracket f \rrbracket, r, s) = o$ or $\llbracket f \rrbracket = \min\{s' \mid o \sqsubseteq s'\} = \llbracket o \rrbracket$. Since $match_obj(f, r, s) = o$, we have $f \sqsubseteq o$. Then, by (a) we can distinguish two cases, namely $\llbracket f \rrbracket \sqsubseteq o$ or $o \sqsubseteq \llbracket f \rrbracket$:

- let $\llbracket f \rrbracket \sqsubseteq o$ and let us assume by contradiction that $match_obj(\llbracket f \rrbracket, r, s) \neq o$. Since $match_obj(f, r, s) = o$, we have $match_obj(o, r, s) = o$, which implies $match_obj(\llbracket f \rrbracket, r, s) \sqsubset o$ by (c) and assumption $match_obj(\llbracket f \rrbracket, r, s) \neq o$. Given that $f \sqsubseteq \llbracket f \rrbracket$, we then have $match_obj(f, r, s) \sqsubseteq match_obj(\llbracket f \rrbracket, r, s)$ by (c), which implies $match_obj(f, r, s) \sqsubset o$ by transitivity, giving a contradiction;
- let $o \sqsubseteq \llbracket f \rrbracket$ and let us assume by contradiction that $\llbracket o \rrbracket \sqsubset \llbracket f \rrbracket$, i.e., $\llbracket o \rrbracket \sqsubseteq \llbracket f \rrbracket$ and $\llbracket o \rrbracket \neq \llbracket f \rrbracket$. Since $f \sqsubseteq o$, we have $\llbracket f \rrbracket \sqsubseteq \llbracket o \rrbracket$ by (b), thus we have $\llbracket f \rrbracket = \llbracket o \rrbracket$ by antisymmetry, giving a contradiction.

□

Theorem 5.3 (Soundness). *If $\sigma \xrightarrow{\alpha} \sigma'$, then there exists a label β such that $\llbracket \sigma \rrbracket \xrightarrow{\beta}_a \llbracket \sigma' \rrbracket$.*

Proof. By a case analysis on the rule applied to derive $\sigma \xrightarrow{\alpha} \sigma'$. If the rule is (SETR), the conclusion follows by the first item of Proposition 5.1. If the rule is (SETU) or (SETG), the conclusion relies on both items of Proposition 5.1 that imply that \hat{r} and \hat{s} in the abstract semantics are the same as \hat{r} and s in the concrete semantics and consequently, $\llbracket u' \rrbracket \in \llbracket usr_trans \rrbracket(\hat{r}, \hat{s})$ and $\llbracket g \rrbracket \in \llbracket grp_trans \rrbracket(\hat{r}, \hat{s})$. If the rule is (EXEC), we conclude again by Proposition 5.1, in combination with Lemma 5.2 which additionally implies $\llbracket f' \rrbracket \in \text{img}(o, \hat{r}, \hat{s})$. □

Interestingly, our formalization enjoys also a completeness result, which states that every transition in the abstract semantics has a corresponding transition in the concrete semantics for some Linux system hosting grsecurity, as far as

there exist at least one user and one group that do not have a corresponding role defined in the policy.

Lemma 5.4 (Concrete Execution). *If $s' \in \text{img}(o, r, s)$ and $\text{perms}(r, s, o) \downarrow$, then there exists f such that $\llbracket f \rrbracket = s'$ and $\text{match_obj}(f, r, s) = o$.*

Proof. Since $s' \in \text{img}(o, r, s)$, we can distinguish two cases. If $s' = \llbracket o \rrbracket$, we let $f = o$. Otherwise, if $s' \sqsubseteq o$ and $\text{match_obj}(s', r, s) = o$, we let $f = s'$. \square

Theorem 5.5 (Completeness). *Consider a policy such that $\exists u, g : u \notin R_u, g \notin R_g$. If $\sigma \xrightarrow{\beta}_a \sigma'$, then there exist a label α and two concrete states $\hat{\sigma}, \hat{\sigma}'$ such that $\llbracket \hat{\sigma} \rrbracket = \sigma$, $\llbracket \hat{\sigma}' \rrbracket = \sigma'$ and $\hat{\sigma} \xrightarrow{\alpha} \hat{\sigma}'$.*

Proof. By a case analysis on the rule applied to derive $\sigma \xrightarrow{\beta}_a \sigma'$. For rules (A-SETR) (A-SETU) and (A-SETG) the concrete states are the same as the abstract ones apart from the special identity “-” that is mapped to the u or the g that we have assumed not to belong to R_u and R_g . We rely on Lemma 5.4 for finding a f in concrete rule (EXEC) such that $\llbracket f \rrbracket$ is the same as s' in the abstract rule (A-EXEC). \square

5.3 Security Analysis

Policies in `grsecurity` are much more concise and readable than policies for other access control systems as, e.g., `SELinux` [26]. However, the plain syntactic structure of the policy does not expose a number of unintended harmful behaviors which can arise at runtime. Just to mention the simplest possible issue, the system administrator may want to prevent user `alice` from reading the files in `bob`’s home directory, but any permission set for role `alice` may be overlooked, whenever `alice` was somehow able to impersonate `bob` through a number of role transitions.

We now devise a simple formalism for verifying through our semantics if a policy is “secure”. The usage of the inverted commas is intended to denote the intrinsic difficulty in answering such a question, due to the lack of an underlying *system* policy, stating the desiderata of the system administrator. General approaches

to RBAC policies verification consider a declarative notion of error in terms of satisfiability of an arbitrary query [9]; more practical works, instead, are tailored around specific definitions of error, like the impersonation of undesired roles [10]. Here, we adopt the latter approach and we validate the policy with respect to some simple requirements on information access, which we consider desirable goals for realistic policies. This is a precise choice, since our research targets the development of a tool which should be effectively usable by system administrators. Of course, our semantics can easily fit different kind of analyses, possibly extending or generalizing those presented here.

The basic ingredient for verification consists in defining which permissions are effectively granted to a given state. Namely, we introduce two judgements $\sigma \vdash \text{Read}(f)$ and $\sigma \vdash \text{Write}(f)$ to denote that file f is readable (writeable) in state σ . The definition of such judgements arises as expected.

(L-READ)

$$\frac{\begin{array}{l} \hat{r} = \text{role}(r_s, u, g) \\ s = \text{match_subj}(f, \hat{r}) \quad o = \text{match_obj}(f', \hat{r}, s) \\ \mathbf{r} \in \text{perms}(\hat{r}, s, o) \quad \mathbf{h} \notin \text{perms}(\hat{r}, s, o) \end{array}}{\langle r_s, u, g, f \rangle \vdash \text{Read}(f')}$$

(L-WRITE)

$$\frac{\begin{array}{l} \hat{r} = \text{role}(r_s, u, g) \\ s = \text{match_subj}(f, \hat{r}) \quad o = \text{match_obj}(f', \hat{r}, s) \\ \mathbf{w} \in \text{perms}(\hat{r}, s, o) \quad \mathbf{h} \notin \text{perms}(\hat{r}, s, o) \end{array}}{\langle r_s, u, g, f \rangle \vdash \text{Write}(f')}$$

We assume to extend such rules to abstract states, in terms of the judgements $\sigma \vdash_a \text{Read}(f)$ and $\sigma \vdash_a \text{Write}(f)$.

Lemma 5.6 (Safety). *Let \mathcal{J} be either $\text{Read}(f)$ or $\text{Write}(f)$:*

- (i) *if $\sigma \vdash \mathcal{J}$, then $\llbracket \sigma \rrbracket \vdash_a \mathcal{J}$.*
- (ii) *if $\sigma \vdash_a \mathcal{J}$, then there exists a concrete state σ' such that $\llbracket \sigma' \rrbracket = \sigma$ and $\sigma' \vdash \mathcal{J}$.*

Proof. This immediately follows by Proposition 5.1. \square

All the security analyses we propose below are based on the reachability of a state with given permissions. Lemma 5.6, in combination with Theorem 5.3, guarantees that the properties can be soundly validated on the abstract semantics; in combination with Theorem 5.5, instead, it ensures that any security violation found in the abstract semantics has a counterpart in some Linux system. Thus, verification turns out to be decidable and can be effectively performed, as we discuss in Section 6. For readability, we state the analyses only for the concrete semantics.

Specification of the analyses The first analysis we propose focuses on direct accesses to files, both for reading and for writing. In particular, we want to verify if a user u can eventually get read (write) access to a given file f . While easy to specify, we believe that such property fits the needs of many system administrators, since the operational behaviour of `grsecurity` is subtler than expected. The formal description of the property we consider is reminiscent of similar specifications through temporal logics for verification like CTL and LTL [27, 28]. Namely, we define two judgements $\sigma \vdash \text{ERead}(f, \sigma')$ and $\sigma \vdash \text{EWrite}(f, \sigma')$ to denote that file f can *eventually* be read (written) in state σ' , starting from state σ .

$$\frac{\text{(L-ERead)} \quad \sigma \xrightarrow{\alpha_1} \dots \xrightarrow{\alpha_n} \sigma' \quad \sigma' \vdash \text{Read}(f)}{\sigma \vdash \text{ERead}(f, \sigma')}$$

The rule for $\sigma \vdash \text{EWrite}(f, \sigma')$ arises as expected. We just write $\sigma \vdash \text{ERead}(f)$ and $\sigma \vdash \text{EWrite}(f)$ when σ' is unimportant.

Given a user u , we denote with $\mathcal{S}(u)$ the set of the *initial states* of u . Any initial state for u has the form $\langle -, u, g, f \rangle$, where g is the primary group assigned to u by the underlying Linux system and f is a possible entry point for u . For instance, `/bin/bash` may be the standard entry point for users interfacing to the system through a “bash” shell. Here, we just assume to be given a set of initial entry

points for any user and we defer the discussion on the definition of such sets to Section 6. Note also that for initial states we are assuming that the user is not acting under any special role, since impersonation of such roles may happen only through authentication to the `gradm` utility, after a standard login operation to the Linux system.

Definition 5.7 (Eventual Read Access). A user u can *eventually read* file f if and only if there exists $\sigma \in \mathcal{S}(u)$ such that $\sigma \vdash \text{ERead}(f)$.

Eventual write access is defined accordingly.

We now build on our first analysis to specify a stronger property, inspired by the literature on information flow control [29]. We note, however, that in our setting we do not have any explicit notion of security label, so we focus on flows among different roles. Namely, if a user u_1 can read the content of file f and then write on an object o readable by u_2 , then there exists a possible flow of information from u_1 to u_2 through o . This is an adaptation to our framework of the well-known “star-property” [30].

Definition 5.8 (Reading Flow). There exists a *reading flow* on file f from user u_1 to user u_2 if and only if:

- (i) there exists $\sigma \in \mathcal{S}(u_1)$ such that $\sigma \vdash \text{ERead}(f, \sigma')$ and $\sigma' \vdash \text{EWrite}(o)$ for some o ;
- (ii) there exists $\sigma'' \in \mathcal{S}(u_2)$ such that $\sigma'' \vdash \text{ERead}(o)$.

Writing flows can be dually defined to address integrity issues. Again, this is just a reformulation into our setting of a standard property [31].

Definition 5.9 (Writing Flow). There exists a *writing flow* on file f from user u_1 to user u_2 if and only if:

- (i) there exists $\sigma \in \mathcal{S}(u_1)$ such that $\sigma \vdash \text{EWrite}(o)$ for some o ;
- (ii) there exists $\sigma' \in \mathcal{S}(u_2)$ such that $\sigma' \vdash \text{ERead}(o, \sigma'')$ and $\sigma'' \vdash \text{EWrite}(f)$.

Note that both previous definitions ignore flows generated by multiple interacting users through a set of intermediate objects. While there is no technical difficulty in generalizing the definitions to such cases, we note that the current formulation already describes very strong properties.

The last analysis we consider accounts for a dangerous combination of permissions over the same object. Namely, if a user can acquire both permissions 'w' and 'x' on o , then o can be exploited for malicious code injection. *grsecurity* identifies this as an important problem, so it prevents the administrator from granting both said permissions for the same object; however, such a situation can arise at runtime, so we consider interesting to monitor it. We omit the formal specification of the analysis, much along the same lines of the previous proposals. We refer to Section 8 for details on our experiments.

Chapter 6

Gran: a Tool for Policy Verification

We present `gran`, a security analyser for `grsecurity` policies. The tool is written in Python and comprises around 1000 lines of code. At the time of writing, `gran` development stopped in favour of the new tool, `granalyze`, introduced in Chapter 7; the source code for a beta release of `gran` can be downloaded at <http://github.com/secgroup/gran>.

6.1 Implementation Overview

Given a `grsecurity` policy, `gran` performs a pre-processing, which involves the expansion of the `include` and `replace` directives. These are just syntactic sugar, used to import fragment of other policies and to define macros, respectively. The tool then generates a model of the policy based on our formalization, i.e., it constructs a tuple $(R, S, O, perms, caps, role_trans, usr_trans, grp_trans)$.

Roles, subjects and objects are retrieved simply by parsing the policy specification. The generation of *perms* involves an unfolding of the pre-processed policy, to cope with the inheritance mechanism of `grsecurity`. We recall that, if a subject *s* does not specify any permission for object *o*, but a less specific subject defines an entry for it, then *s* inherits the same permissions for *o*. The only

exception to this rule is when the subject specifies the “override” mode ‘o’, which prevents this behaviour. Thus, the permissions stored in *perms* correspond to a properly unfolded version of those specified in the original policy.

Every capability is allowed by default, so for every role r and subject s we initially let $caps(r, s) = C$, then we remove any forbidden capability. Addition and revocation of capabilities is performed through the rules `+CAP_NAME` and `-CAP_NAME`, respectively. The overall result is order-sensitive, i.e., specifying first `+CAP_SETUID` and then `-CAP_SETUID` forbids the capability, while swapping the rules allows it. We also account for inheritance of capabilities among subjects defined in the same role.

Transitions to special roles are forbidden by default, so for every role r we initially let $role_trans(r) = \emptyset$ and then we introduce in the set only the transitions explicitly allowed by the attribute `role_transitions`. Conversely, transitions to user roles are allowed by default, so we let $usr_trans(r, s) = R_u \cup \{-\}$ for any role r and subject s not providing any further specification. We recall that we abstract users with no associated user role by the distinguished identity “-”. The attribute `user_transition_allow` can be used to restrict allowed user transitions to the ones specified. Conversely, the attribute `user_transition_deny` can be used to permit all users transitions except those listed. The two attributes cannot coexist. If subject s in role r specifies a set U of allowed user transitions, we let $usr_trans(r, s) = \{\llbracket u \rrbracket \mid u \in U\}$. Conversely, if U is a set of denied user transitions, we let $usr_trans(r, s) = (R_u \setminus U) \cup \{-\}$. We apply a similar processing to construct the function grp_trans .

The tool disregards features that are not modeled, such as resource restrictions and socket policies. Domains, i.e., sets of user or group roles sharing a common set of permissions, are handled through unfolding as a set of user or group roles. Nested subjects are not supported, since the learning system of `grsecurity` does not account for them. In fact, `grsecurity` features the possibility to automatically generate a policy by inferring the right permissions from the standard usage of the system, to avoid burdening the user with the necessity of specifying all the details about access control. Since most users perform a full system learning

and then tweak the generated policy around their own needs, we think nested subjects can be safely disregarded by our analysis.

After the parsing of the policy, `gran` generates all the possible states of the model and computes the set of the transitions. The tool implements all the analyses described in Section 5.3: the initial states and the sensible objects to consider for verification can be specified through command-line parameters. As a default choice, `gran` generates an initial state for each non-special role in the policy, assuming “/” as the subject entry point. If no target is specified for the analysis, `gran` infers a set of sensible resources by the specification provided in the configuration files of the learning system.

The help message of `gran` is depicted in Figure 6.1. Further information regarding the usage of the tool can be found on the project webpage [32].

```
$ gran -h
usage: gran [-h] [-a] [-b] [-e ENTRYPOINTS] [-t TARGETS]
           [-l LEARNCONFIG] [-d] [-v]
           policy

a security analyser for Grsecurity RBAC policies.

positional arguments:
  policy                policy file to be analyzed

optional arguments:
  -h, --help            show this help message and exit
  -a, --admin           include administrative special roles in the
                        analysis
  -b, --bestcase        assume there are not setuid/setgid files in the
                        system
  -e ENTRYPOINTS, --entrypoints ENTRYPOINTS
                        specify the entrypoints file
  -t TARGETS, --targets TARGETS
                        specify a file with sensitive targets
  -l LEARNCONFIG, --learnconfig LEARNCONFIG
                        specify the learn_config file to get sensitive
                        targets
  -d, --debug           enable debugging mode
  -v, --version         show program's version number and exit
```

FIGURE 6.1: `gran` help message

Chapter 7

Granalyze: an Improved Security Analyser

As stated in Chapter 4, we reported a potentially harmful behavior when executing `setuid/setgid` programs to the main developer of `grsecurity`. A fix has been implemented in the latest stable release of `grsecurity` [13]. The solution consists in requiring the capabilities `CAP_SETUID/CAP_SETGID` to perform role transitions, even upon execution of `setuid/setgid` binaries. As follows, after the fix it is not possible to infer the active non-special role of a process from its real UID (GID).

For instance, let's assume two user roles to be defined in the policy, namely `alice` and `bob`. We also assume that the policy comprises the subject `/bin/test` for the user role `alice` and that the capability `CAP_SETUID` is not granted to `/bin/test`. When `alice` executes the `setuid` program `/bin/test` owned by `bob`, the effective UID of the process is set to the UID of `bob`. Since Linux allows unprivileged processes to set the real UID to the value of the effective UID [12], the process `/bin/test` is permitted to set its real UID to `bob`. Whilst this operation would have involved the process to change its effective role to `bob`, the latest releases of `grsecurity` require the `CAP_SETUID` capability to allow the role transition.

7.1 An Updated Framework

It follows that the semantics presented in Chapters 4,5 are not adequate to model the new behaviour of `grsecurity`. Therefore, we introduce a new semantics developed on top of the formal framework previously reported to match with the latest releases of `grsecurity`.

7.1.1 Concrete Semantics

We assume an underlying file system, i.e., a subset of a denumerable set of path-names F , ranged over by f . Let r range over $R \cup \{-\}$, a *state* is a 4-tuple $\sigma' = \langle r, u, g, f \rangle$ describing a process spawned by the execution of file f . The process impersonates the role r and is running with effective UID set to $u \in U$ and effective GID set to $g \in G$.

The transition rules are reported in Table 7.1. The main difference with respect to the previous semantics involves the rule (EXEC') since the execution of a file does not lead to a role change. As before, rule (SETR') accounts for login operations to special roles. Rule (SETU') describes a change of the process UID if the process possesses the capability `set_uid`. Rule (SETG') details a similar behavior for changing the process GID.

The function $role'$ formalizes the role assignment process and provides a generalisation of the function $role$:

$$role'(r, u, g) = \begin{cases} r & \text{if } r \in R_s \\ u & \text{if } r \notin R_s, u \in R_u \\ g & \text{if } r \notin R_s, r \notin R_u, u \notin R_u, g \in R_g \\ - & \text{otherwise} \end{cases}$$

7.1.2 Abstract Semantics

We define an abstract state as a 4-tuple $\sigma'_a = \langle r, r_u, r_g, s \rangle$ describing a process spawned by the execution of some file $f \sqsubseteq s$. The role assigned to the process is

$$\begin{array}{c}
\text{(SETR')} \\
\frac{r'_s \in \text{role_trans}(r) \cup \{-\} \quad \hat{r} = \text{role}'(r'_s, u, g)}{\langle r, u, g, f \rangle \xrightarrow{\text{set_role}(r'_s)} \langle \hat{r}, u, g, f \rangle} \\
\\
\text{(SETU')} \\
\frac{s = \text{match_subj}(f, r) \quad \text{set_uid} \in \text{caps}(r, s) \quad u' \in \text{usr_trans}(r, s) \quad \hat{r} = \text{role}'(r, u', g)}{\langle r, u, g, f \rangle \xrightarrow{\text{set_UID}(u')} \langle \hat{r}, u', g, f \rangle} \\
\\
\text{(SETG')} \\
\frac{s = \text{match_subj}(f, r) \quad \text{set_gid} \in \text{caps}(r, s) \quad g' \in \text{grp_trans}(r, s) \quad \hat{r} = \text{role}'(r, u, g')}{\langle r, u, g, f \rangle \xrightarrow{\text{set_GID}(g')} \langle \hat{r}, u, g', f \rangle} \\
\\
\text{(EXEC')} \\
\frac{s = \text{match_subj}(f, r) \quad o = \text{match_obj}(f', r, s) \quad \mathbf{x} \in \text{perms}(r, s, o) \quad \mathbf{h} \notin \text{perms}(r, s, o) \quad u' \in \text{usr_trans}(r, s) \cup \{u\} \quad g' \in \text{grp_trans}(r, s) \cup \{g\}}{\langle r, u, g, f \rangle \xrightarrow{\text{exec}(f')} \langle r, u', g', f' \rangle}
\end{array}$$

TABLE 7.1: Updated semantics of grsecurity

again determined by the first component, role r . Table 7.2 presents the reduction rules for the new abstract semantics which heavily rely on the functions and relations presented in Chapter 5.

7.2 Granalyze

We introduce **granalyze**, a security analyser for **grsecurity** policies inspired by **gran**. The tool is written in Python and consists of around 1500 lines of code. At the time of writing, **granalyze** is still under active development; the source code will be released for download at <http://github.com/secgroup/granalyze>. In contrast to **gran**, the new tool is heavily object-oriented and well documented to ease the understanding of its code base. Despite of implementing most of the features of **gran**, **granalyze** is written entirely from scratch.

$$\begin{array}{c}
\text{(A-SETR')} \\
\frac{r'_s \in \text{role_trans}(r) \cup \{-\} \quad \hat{r} = \text{role}'(r'_s, r_u, r_g)}{\langle r, r_u, r_g, s \rangle \xrightarrow{\text{set_spec}(r'_s)}_a \langle \hat{r}, r_u, r_g, s \rangle} \\
\\
\text{(A-SETU')} \\
\frac{\hat{s} = \text{match_subj}(s, r) \quad \text{set_uid} \in \text{caps}(r, \hat{s}) \quad r'_u \in \llbracket \text{usr_trans} \rrbracket(r, \hat{s}) \quad \hat{r} = \text{role}'(r, r'_u, r_g)}{\langle r, r_u, r_g, s \rangle \xrightarrow{\text{set_user}(r'_u)}_a \langle \hat{r}, r'_u, r_g, s \rangle} \\
\\
\text{(A-SETG')} \\
\frac{\hat{s} = \text{match_subj}(s, r) \quad \text{set_gid} \in \text{caps}(r, \hat{s}) \quad r'_g \in \llbracket \text{grp_trans} \rrbracket(r, \hat{s}) \quad \hat{r} = \text{role}'(r, r_u, r'_g)}{\langle r, r_u, r_g, s \rangle \xrightarrow{\text{set_user}(r'_g)}_a \langle \hat{r}, r_u, r'_g, s \rangle} \\
\\
\text{(A-EXEC')} \\
\frac{\hat{s} = \text{match_subj}(s, r) \quad \mathbf{x} \in \text{perms}(r, \hat{s}, o) \quad \mathbf{h} \notin \text{perms}(r, \hat{s}, o) \quad r'_u \in \llbracket \text{usr_trans} \rrbracket(r, \hat{s}) \cup \{r_u\} \quad r'_g \in \llbracket \text{grp_trans} \rrbracket(r, \hat{s}) \cup \{r_g\} \quad s' \in \text{img}(o, r, \hat{s})}{\langle r, r_u, r_g, s \rangle \xrightarrow{\text{exec}(s')}_a \langle r, r'_u, r'_g, s' \rangle}
\end{array}$$

TABLE 7.2: Updated abstract semantics of grsecurity

Given a `grsecurity` policy, `granalyze` performs a preprocessing similarly to `gran`. The tool then produces the abstract syntax tree of the parsed policy, unfolds permissions and capabilities to cope with the inheritance mechanism of `grsecurity` and generates a model based on our formalization.

Like `gran`, the tool disregards features that are not modeled, such as resource restrictions, socket policies, nested subjects and most of the modalities reported in Appendix A which are not relevant to the analysis. Unlike `gran`, `granalyze` handles domains as normal user/group roles. This is accomplished by extending the abstraction function for users and groups to deal with domains. For instance, given a user `alice` member of the user domain `customers`, we have that $\llbracket \text{alice} \rrbracket = \text{customers}$. This elegant solution leads to a substantial reduction in the number of roles and, thus, in the total number of states. Moreover, it is more consistent with the actual behavior of `grsecurity` as compared to the previous approach.

Most notably, **granalyze** allows to inspect the filesystem for `setuid/setgid` files. Moreover, it identifies objects which may be modified to be `setuid/setgid`. By computing the interplay between the policy and the underlying filesystem, **granalyze** produces an accurate model that accounts for a reduction in the number of transitions within the labelled transition systems arising from policy analysis.

The tool implements most of the analyses described in Section 5.3 except for *reading/writing flow* detection¹. On top of these features, **granalyze** deals with the problem of automatic verification of **grsecurity** policies by defining a set of sensitive pathnames which should be protected from untrusted access (e.g., `/etc/shadow`) and a set of trusted applications which are allowed to access sensible objects (e.g., `/bin/passwd`). In our formulation, the set of trusted applications constitutes the TCB of the system with respect to the security invariants expressed by means of the access mode on the sensitive pathnames. The tool computes and point out the TCB to help **grsecurity** administrators to identify possible vulnerabilities.

The help message of **granalyze** is reported in Figure 7.1. Further information can be found on the project webpage [33].

¹such properties, indeed, resulted too strict on real systems and difficult to grasp by the **grsecurity** community

```

$ granalyze -h
usage: granalyze [-h] [-A] [-a] [-fs {none,all,scan}] [-lp LP] [-lf LF]
                [-dp DP] [-df DF] [-r [RULES [RULES ...]]] [-t]
                [--protected-paths PROTECTED_PATHS]
                [--trusted-apps TRUSTED_APPS] [--learnfile LEARNFILE]
                [--fullstart] [-v] [-V]
                [policy]

a security analyser for Grsecurity RBAC policies.

positional arguments:
  policy                policy file to be analyzed

optional arguments:
  -h, --help            show this help message and exit
  -A, --admin           include administrative special roles in the analysis
                        (default: false)
  -a, --auth           include transitions to special roles that require
                        explicite authentication (default: false)
  -fs {none,all,scan}  set the desired type of interaction with the
                        underlying filesystem."none": assume there are no
                        setuid/setgid files; "all": all files could be
                        setuid/setgid and owned by every possible user/group
                        combination; "scan": perform a dynamic analysis of the
                        filesystem to detect the actual set of setuid/setgid
                        files (default: "none")
  -lp LP              path of a precomputed Policy object to load
  -lf LF              path of a precomputed Filesystem object to load
  -dp DP              path to which the Policy object will be saved
  -df DF              path to which the Filesystem object will be saved
  -r [RULES [RULES ...]], --rules [RULES [RULES ...]]
                        set of security invariants the policy should satisfy
  -t, --traces        show execution traces
  --protected-paths PROTECTED_PATHS
                        list of highly sensitive paths that should be
                        protected from untrusted access. The syntax used to
                        specify paths is reminiscent of gentoo USE flags
                        declaration. Remove a path from the set of hardcoded
                        protected paths by prefixing the path with the minus-
                        sign or just add extra paths, e.g. "-/etc/ppp
                        /etc/nginx". To clear the set of stored paths use "-*"
  --trusted-apps TRUSTED_APPS
                        specify the list of trusted apps using the same syntax
                        as described above
  --learnfile LEARNFILE
                        parse the learn_config file in order to populate the
                        list of protected paths with the resources specified
                        as high-protected-path
  --fullstart         assume the full set of starting states as entrypoints
                        when verifying the provided rules. This is usually not
                        needed besides of special roles
  -v, --verbosity     print a detailed output for diagnostic purposes
  -V, --version       show program's version number and exit

```

FIGURE 7.1: granalyze help message

Chapter 8

Case Studies

We illustrate the outcome of practical experiments with `gran` and we give general considerations about possible vulnerabilities detected by our tool. We omit to report the results of experiments performed using `granalyze` since it is still in an early stage of development and the output of the tool is subject to be changed. We point the readers to the webpage of the program [33] which will be constantly updated.

8.1 Verification of Existing Policies

We asked the `grsecurity` community for policies to be verified using `gran`. Unfortunately, most system administrators are unwilling to provide their policies, since they can reveal a number of potentially harmful information about the system. However, we managed to gather a small set of real policies and we analyzed them with our tools. Due to privacy reasons, we cannot reveal any detail of such policies, so we report a properly sanitized outcome of the verification process. Our results were favorably welcome by the lead developer of `grsecurity`, who proposed us to integrate the functionalities of our tools in the `gradm` utility for policy management [34]. We consider this an important opportunity to continue our investigation on a larger scale, since users are for sure more comfortable to provide us the results of the validation rather than to disclose their policy.

We performed the verification of five different policies: the first and the second one from small web servers, the third one from a server running at our department, the fourth one generated by the learning system of `grsecurity`, and the fifth one from a large web server. In all cases, `gran` performed very effectively, providing the results of the analysis in less than one minute on a standard commercial machine. The output of the analysis was manually reviewed, looking for possible vulnerabilities: the process took from 10 to 30 minutes for each policy.

We start by reporting on direct accesses to sensitive information. In some cases, we noticed that critical files like `/etc/shadow` were readable by untrusted users. Even if this is not a vulnerability by itself, since the underlying DAC enforced by Linux does prevent this behavior, we believe that this is a poor specification at the very least. Indeed, system critical files on a hardened server are better protected also by a MAC policy. Interestingly, a similar warning sometimes applies also for resources which are publicly readable, according to the default settings of Linux DAC, but are considered highly sensitive by the standard configuration files of the learning system of `grsecurity`. Examples of such resources include files as `/proc/slabinfo` and `/proc/kallsyms`, whose content may be potentially exploited by an attacker. We also noticed a dangerous specification in one of the analyzed policies: subject `/etc/cron.monthly` was provided almighty access to the system. This can have a tremendous impact on security, since cronjobs are usually executed with `root` privileges, thus mostly bypassing standard DAC. We argue that such a dangerous specification was provided for convenience, since scheduled jobs may need many different access rights and a careful assignment of permissions should feature very high granularity. Finally, we noticed that at least one of the users was not fully aware of the workings of the inheritance mechanism and, by manually tweaking the policy after the learning process, had created some unwanted cascade propagation of permissions.

We also performed some tests based on the other kinds of analyses described in Section 5.3. In particular, we noticed that unwanted writing accesses are much less frequent than undesired reading accesses: this is comforting and it was somehow expected, since the learning system of `grsecurity` tends to grant really few write permissions. The analysis also highlighted that usually only “physical” users, i.e., users with shell access to the system, have the opportunity to get

both write and execution permissions over the same object, thus compromised services are unlikely to execute arbitrary code. Users, instead, probably need such permissions to effectively work on the system.

We think that the overall security of the analyzed `grsecurity` policies was fairly satisfying. We argue that much of the robustness, especially against undesired write accesses, comes from the sophisticated learning system of `grsecurity`, which tries to grant minimal privileges to each user. Indeed, the analyzed auto-generated policy turned out to be quite resilient to vulnerabilities; unfortunately, most administrators need to manually tweak the policy to get an usable system for their users and the overall impact of local changes may be easily overlooked. We think that our tool helps in getting the big picture on the security of the system.

8.2 Exploits Through “setuid” Binaries

The analysis presented in the previous section was performed using the “-b” option of `gran`, which discharges potential attacks due to the “setuid” flaw pointed out in Section 4. We decided to make this choice, since a fix is already merged in `grsecurity`, and in our worst-case scenario almost every object of the policy turns out to be potentially vulnerable. Precisely, the amended abstract semantics assumes that no role transition can be performed upon execution. Here, we discuss the impact of the flaw we found out, by describing a realistic scenario where it can be harmfully exploited by an attacker.

One of the goal of `grsecurity` is to try to drop many of the privileges normally granted to `root`, thus limiting the impact of many known vulnerabilities; however, during the learning process, some background operations may be overlooked, leading to undesired assignment of permissions. For instance, let us assume that an administrator starts full system learning to generate his own policy, when a scheduled cronjob performs an access to a sensitive resource: in this case the learning system could provide `root` with liberal access rights on the resource, since it would consider it as a normal system behaviour. If the administrator does not take care in manually strengthening the policy after the learning process, “setuid”

```

role root uG
role_transitions admin
...
subject /usr/sbin/cron o {
user_transition_allow alice
group_transition_allow users
    /                h
    /usr             h
    /usr/sbin/cron  rx
}

role alice u
...
subject / {
    /usr/bin
}
subject /usr/sbin/cron {
    /usr/bin          rx
}
subject /usr/bin/python2.7 o {
    /                h
    /tmp             rw
    /home            r
    /home/alice/bin  r
    -CAP_ALL
}

role bob u
subject / o {
    /                h
    /bin             x
    -CAP_ALL
}
subject /bin/bash {
    /tmp             rw
    /home/bob        rw
}

```

TABLE 8.1: A snippet of a flawed `grsecurity` policy

binaries can lead to unintended impersonation of a powerful `root` role, bypassing the capability system. Indeed, even if the learning process tries to forbid as many capabilities as possible to user roles, such a practice does not offer the expected level of security, due to the subtle interplay between `grsecurity` and Linux.

8.3 Information Leakage Analysis

We conclude our experiments by performing an information leakage analysis on a policy we generated for testing. We agree on the common statement that compartmentalization between users is a too strict property for many realistic systems; still, it can be interesting in some highly sensitive settings [30, 31]. Our sample policy is shipped with the `gran` package and a subset of it is depicted in Table 8.1.

When we process the policy with `gran`, we find out that user `alice` is able to share some confidential information in her home directory with her accomplice `bob` through a leakage on `/tmp`. The attack is mounted on top of `cron`, which our experiments seem to identify as a subtle subject. The output of `gran` looks as follows:

```
[!!] Indirect flow found for target
/home/alice on object /tmp
Traces for writing:
[1] root:U:/usr/sbin/cron
    -set_UID(alice)->
    alice:U:/usr/sbin/cron
    -exec(/usr/bin)->
    alice:U:/usr/bin/python2.7
Traces for reading:
[1] bob:U:/
    -exec(/bin)->
    bob:U:/bin/bash
```

We assume that `alice` can schedule her tasks through `cron`. The daemon initially runs as `root`, changes its identity to `alice` and selects for execution a Python script in `/home/alice/bin`. The subject `/usr/bin/python2.7` defined for `alice` can read `/home/alice/bin` and write on `/tmp`. `bob` cannot directly read `/tmp`, but he can execute `/bin/bash` and get read access on `/tmp`. It is worth noticing that `alice` cannot directly execute Python to get write access on `/tmp`, since her default subject `"/` does not allow execution of files under `/usr/bin`.

Our tool is then able to identify a subtle and unintended flow, which is unlikely to be noticed by just looking at the policy. We think that **gran** can help in strengthening the system against leakage of some particularly sensitive targets.

Chapter 9

Conclusion

We have presented a framework for a formal, automated analysis of `grsecurity`'s RBAC system, to help system administrators validate and maintain their policies. As we have illustrated, our endeavor has proved useful, as `gran` has unveiled a series of ambiguities and unexpected behaviors that have been reported to the main developer, confirmed and fixed. Our results have been well-received by the lead developer of `grsecurity`, who have requested to integrate the tool in the official `gradm` utility for policy administration.

To the best of our knowledge, the present work is the first research focusing on the verification of `grsecurity` RBAC policies. However, there exists a huge literature on the analysis of (A)RBAC policies in general, mainly targeted to the isolation of restricted classes of policies whose verification is tractable.

Sasturkar et al. [3] show that role reachability is PSPACE-complete for ARBAC and identify restrictions on the policy language to partially tame this complexity; similar results are presented by Jha et al. in later work [9]. Li and Tripunitara [35] perform a security analysis on restricted ARBAC fragments and identify a specific class of queries which can be answered efficiently. Stoller et al. [36] isolate subsets of policies of practical interest and develop algorithms to analyze them; their techniques are implemented in the RBAC-PAT tool [5], which supports also information flow analysis much in the spirit of the one provided by `gran`. Jayaraman et al. [10] propose Mohawk, a model-checker implementing

an abstraction-refinement technique aimed to error finding in complex ARBAC policies.

Contrary to all these works, our framework is not targeted to the analysis of generic ARBAC policies, but of real, full-fledged `grsecurity` RBAC policies, which turn out to be amenable for efficient static verification. A formal comparison with previous work, however, might be useful to understand how possible extensions of `grsecurity` would impact on the complexity of the analysis. We leave this as future work.

There are a number of further issues that are part of our plans for future work, some of which we discuss below. First, it would be desirable to extend `granalyze` with support for the networking components. For this purpose, it would be worthwhile to consider a recent publication where the Authors provide a formal model of `Netfilter`, a firewall system integrated in the Linux kernel [37]. Secondly, it would be interesting to integrate our tool with existing model checkers. This could be done by implementing a back-end module for the generation of the model checker input after the policy parsing. As a result, `granalyze` would be amenable for additional formal reasoning (e.g., by providing queries expressed in some temporal logic form), while, at the same time, being still able to perform a concrete analysis of real systems.

While considering huge policies or runtime policy state changes, as in the case of ARBAC models [2][38], our intuition is that the model checking approach would be infeasible due to a state explosion problem. To counter this issue, we plan to translate the LTS reachability problem into a program verification task suitable for information leakage analysis [39]. The idea of analysing RBAC policies through program verification using abstract interpretation techniques is not new [38]; however, to the best of our knowledge, this would be the first attempt at performing an information leakage analysis on such models. The goal is to explore the feasibility of this approach, as it could potentially lead to new research directions. Additionally, the results may be general enough to fit a range of (A)RBAC models.

Appendix A

Grsecurity Policy Modalities

Mode	Meaning
u	This role is a user role. That is, the role name must be an existing user on the system
g	This role is a group role. That is, the role name must be an existing group on the system
s	This role is a special role, meaning it does not belong to a user or group and does not require an enforced secure policy base to be included in the ruleset
l	This role has learning enabled
A	This role is an administrative role, thus it has special privileges that normal roles do not have. In particular, this role bypasses the additional ptrace and library loading restrictions
G	This role can use gradm to authenticate to the kernel. A policy for <code>gradm</code> will automatically be added to the role
N	This role does not require authentication. To access this role, use <code>gradm -n <rolename></code>
P	This role uses Pluggable Authentication Modules (PAM) for authentication
T	This role has Trusted Path Execution (TPE) enabled
R	The role is persistence. When shell/session in which authorization was done is terminated, spawned processes won't be dropped to non-special role

TABLE A.1: Role Modes

Mode	Meaning
a	Allow this process to talk to the <code>/dev/grsec</code> device
d	Protect the <code>/proc/<pid>/fd</code> , <code>/proc/<pid>/mem</code> , <code>/proc/<pid>/cmdline</code> , and <code>/proc/<pid>/environ</code> entries for processes in this subject
h	This process is hidden and only viewable by processes with the <code>v</code> mode
i	Enable inheritance-based learning, causing all accesses of this subject and anything it executes to be logged as originating from this subject. The policy generated from this learning will have the inheritance flag added to every file executed from this subject
k	This process can kill protected processes
l	Enables learning mode for this process
o	Override ACL inheritance for this process
p	This process is protected; it can only be killed by processes with the <code>k</code> mode, or by processes within the same subject
r	Relax <code>ptrace</code> restrictions (allows ptracing of processes other than one's own children)
s	Enable <code>AT_SECURE</code> when entering this subject. This enables the same environment sanitization that occurs in <code>glibc</code> upon execution of a <code>suid</code> binary
t	Allow ptracing of any process (do not use unless necessary, allows <code>ptrace</code> to cross subject boundaries). This flag also allows a process to use <code>CLONE_FS</code> and execute a binary that causes a subject change
v	This process can view hidden processes
x	Allows executable anonymous shared memory for this subject
A	Protect the shared memory of this subject. No other processes but processes contained within this subject may access the shared memory of this subject
C	Auto-kill all processes belonging to the attacker's IP address upon violation of security policy
K	When processes belonging to this subject generate an alert, kill the process
O	Allow loading of writable libraries
T	Deny execution of binaries or scripts that are writable by any other subject in the policy. This flag is evaluated at policy enable time

TABLE A.2: Subject Modes

Mode	Meaning
none	Lack of any of the below modes implies “find” access to the object. The object can be listed and have its ownership, size, etc. information obtained, but cannot be read or modified
a	This object can be opened for appending
c	Allow creation of the file/directory
d	Allow deletion of the file/directory
f	Needed to mark the pipe used for communication with init to transfer the privilege of the persistent role; only valid within a persistent role. Transfer only occurs when the file is opened for writing
h	This object is hidden
i	This mode only applies to binaries. When the object is executed, it inherits the ACL of the subject in which it was contained
l	Allow a hardlink at this path. Hardlinking requires a minimum of c and l modes, and the target link cannot have any greater permission than the source file
m	Allow creation of setuid/setgid files/directories and modification of files/directories to be setuid/setgid
p	Reject all ptraces to this object
r	This object can be opened for reading
t	This object can be ptraced, but cannot modify the running task. This is referred to as a “read-only ptrace ”
w	This object can be opened for writing or appending
x	This object can be executed (or mmap 'd with PROT_EXEC into a task).

TABLE A.3: Object Modes

Bibliography

- [1] M. Bugliesi, S. Calzavara, R. Focardi, and M. Squarcina, “Gran: Model checking grsecurity rbac policies,” in *Proceedings of the 2012 IEEE 25th Computer Security Foundations Symposium, CSF '12*, (Washington, DC, USA), pp. 126–138, IEEE Computer Society, 2012.
- [2] R. S. Sandhu, V. Bhamidipati, and Q. Munawer, “The arbac97 model for role-based administration of roles,” *ACM Trans. Inf. Syst. Secur.*, vol. 2, no. 1, pp. 105–135, 1999.
- [3] A. Sasturkar, P. Yang, S. D. Stoller, and C. R. Ramakrishnan, “Policy analysis for administrative role based access control,” in *CSFW*, pp. 124–138, IEEE Computer Society, 2006.
- [4] A. Armando and S. Ranise, “Automated symbolic analysis of arbac-policies,” in *STM* (J. Cuéllar, J. Lopez, G. Barthe, and A. Pretschner, eds.), vol. 6710 of *Lecture Notes in Computer Science*, pp. 17–34, Springer, 2010.
- [5] M. I. Gofman, R. Luo, A. C. Solomon, Y. Zhang, P. Yang, and S. D. Stoller, “Rbac-pat: A policy analysis tool for role based access control,” in *TACAS* (S. Kowalewski and A. Philippou, eds.), vol. 5505 of *Lecture Notes in Computer Science*, pp. 46–49, Springer, 2009.
- [6] M. A. Harrison, W. L. Ruzzo, and J. D. Ullman, “Protection in operating systems,” *Commun. ACM*, vol. 19, no. 8, pp. 461–471, 1976.
- [7] E. M. Clarke, O. Grumberg, and D. A. Peled, *Model Checking*. The MIT Press, 1999.

-
- [8] N. Zhang, M. Ryan, and D. P. Guelev, “Synthesising verified access control systems through model checking,” *Journal of Computer Security*, vol. 16, no. 1, pp. 1–61, 2008.
- [9] S. Jha, N. Li, M. V. Tripunitara, Q. Wang, and W. H. Winsborough, “Towards formal verification of role-based access control policies,” *IEEE Trans. Dependable Sec. Comput.*, vol. 5, no. 4, pp. 242–255, 2008.
- [10] K. Jayaraman, V. Ganesh, M. V. Tripunitara, M. C. Rinard, and S. J. Chapin, “Automatic error finding in access-control policies,” in *ACM Conference on Computer and Communications Security* (Y. Chen, G. Danezis, and V. Shmatikov, eds.), pp. 163–174, ACM, 2011.
- [11] B. Spengler, “Increasing performance and granularity in role-based access control systems.” <http://grsecurity.net/researchpaper.pdf>, 2004.
- [12] “man page for function `setreuid`.” <http://linux.die.net/man/2/setreuid>.
- [13] B. Spengler, “Changelog of grsecurity.” <http://grsecurity.net/changelog-stable2.txt>, February 2012. commit 3981059c35e8463002517935c28f3d74b8e3703c.
- [14] D. F. Ferraiolo, D. R. Kuhn, and R. Chandramouli, *Role-Based Access Control*. Norwood, MA, USA: Artech House, Inc., 2003.
- [15] B. W. Lampson, “Dynamic protection structures,” in *Proceedings of the November 18-20, 1969, fall joint computer conference*, AFIPS ’69 (Fall), (New York, NY, USA), pp. 27–38, ACM, 1969.
- [16] D. Bell and L. LaPadula., “Secure computer systems: Unified exposition and multics interpretation,” in *Technical Report ESD-TR-75-306, MTR-2997, MITRE, Bedford, Mass*, 1975.
- [17] D. of Defense, *Trusted Computer System Evaluation Criteria*. DoD 5200.28-STD, National Computer Security Center, Dec 1985.
- [18] D. D. Clark and D. R. Wilson, “A comparison of commercial and military computer security policies,” *IEEE Symposium of Security and Privacy*, vol. 0, pp. 184–194, 1987.

-
- [19] D. G. Ferraiolo, D. and N. Lynch, "An examination of federal and commercial access control policy needs," in *Proceedings of NISTNCSC National Computer Security Conference*, (Baltimore, MD, USA), pp. 20–23, 1993.
- [20] D. Ferraiolo and R. Kuhn, "Role-based access control," in *In 15th NIST-NCSC National Computer Security Conference*, pp. 554–563, 1992.
- [21] R. Shirey, "Internet Security Glossary, Version 2." RFC 4949 (Informational), Aug. 2007.
- [22] R. S. Sandhu, E. J. Coyne, H. L. Feinstein, and C. E. Youman, "Role-based access control models," *IEEE Computer*, vol. 29, pp. 38–47, 1996.
- [23] R. Sandhu, D. Ferraiolo, and R. Kuhn, "The nist model for role-based access control: towards a unified standard," in *Proceedings of the fifth ACM workshop on Role-based access control*, RBAC '00, (New York, NY, USA), pp. 47–63, ACM, 2000.
- [24] D. Ferraiolo, R. Kuhn, and R. Sandhu, "Rbac standard rationale: Comments on "a critique of the ansi standard on role-based access control"," *IEEE Security and Privacy*, vol. 5, pp. 51–53, Nov. 2007.
- [25] "Sponsor page of grsecurity." <http://grsecurity.net/sponsors.php>.
- [26] M. Fox, J. Giordano, L. Stotler, and A. Thomas, "SELinux and grsecurity: A case study comparing linux security kernel enhancements." University of Virginia.
- [27] E. M. Clarke, E. A. Emerson, and A. P. Sistla, "Automatic verification of finite-state concurrent systems using temporal logic specifications," *ACM Trans. Program. Lang. Syst.*, vol. 8, no. 2, pp. 244–263, 1986.
- [28] A. Pnueli, "The temporal logic of programs," in *FOCS*, pp. 46–57, 1977.
- [29] A. Sabelfeld and A. C. Myers, "Language-based information-flow security," *IEEE Journal on Selected Areas in Communications*, vol. 21, 2003.
- [30] D. E. Bell and L. J. LaPadula, "Secure computer systems: Mathematical foundations," tech. rep., MITRE Corporation, 1973.

-
- [31] K. J. Biba, “Integrity Considerations for Secure Computer Systems,” tech. rep., USAF Electronic Systems Division, 1977.
- [32] M. Bugliesi, S. Calzavara, R. Focardi, and M. Squarcina, “Homepage of gran.” <https://github.com/secgroup/gran>.
- [33] M. Squarcina, “Homepage of granalyze.” <https://github.com/secgroup/granalyze>.
- [34] B. Spengler, “Private communication,” February 2012.
- [35] N. Li and M. V. Tripunitara, “Security analysis in role-based access control,” *ACM Trans. Inf. Syst. Secur.*, vol. 9, no. 4, pp. 391–420, 2006.
- [36] S. D. Stoller, P. Yang, C. R. Ramakrishnan, and M. I. Gofman, “Efficient policy analysis for administrative role based access control,” in *ACM Conference on Computer and Communications Security* (P. Ning, S. D. C. di Vimercati, and P. F. Syverson, eds.), pp. 445–455, ACM, 2007.
- [37] P. Adão, C. Bozzato, G. Dei Rossi, R. Focardi, and F. Luccio, “Mignis: A semantic based tool for firewall configuration,” in *Proceedings of the 2014 IEEE 27th Computer Security Foundations Symposium, CSF ’14*, IEEE Computer Society, to appear.
- [38] A. L. Ferrara, P. Madhusudan, and G. Parlato, “Security analysis of role-based access control through program verification,” in *Proceedings of the 2012 IEEE 25th Computer Security Foundations Symposium, CSF ’12*, (Washington, DC, USA), pp. 113–125, IEEE Computer Society, 2012.
- [39] M. Zanioli and A. Cortesi, “Information leakage analysis by abstract interpretation,” in *Proceedings of the 37th international conference on Current trends in theory and practice of computer science, SOFSEM’11*, (Berlin, Heidelberg), pp. 545–557, Springer-Verlag, 2011.