

Aggregating and Deploying Network Access Control Policies

Joaquín G. Alfaro ^{†,‡}

Frédéric Cuppens [‡]

Nora Cuppens-Boulahia [‡]

[†] Universitat Oberta de Catalunya
Rambla Poble Nou 156,
08018 Barcelona - Spain
joaquin.garcia-alfaro@acm.org

[‡] GET/ENST-Bretagne,
2, rue de la Châtaigneraie,
35576 Cesson Sévigné - France
{frederic.cuppens,nora.cuppens}@enst-bretagne.fr

Abstract. The existence of errors or inconsistencies in the configuration of security components, such as filtering routers and/or firewalls, may lead to weak access control policies — potentially easy to be evaded by unauthorized parties. We present in this paper a proposal to create, manage, and deploy consistent policies in those components in an efficient way. To do so, we combine two main approaches. The first approach is the use of an aggregation mechanism that yields consistent configurations or signals inconsistencies. Through this mechanism we can fold existing policies of a given system and create a consistent and global set of access control rules — easy to maintain and manage by using a single syntax. The second approach is the use of a refinement mechanism that guarantees the proper deployment of such a global set of rules into the system, yet free of inconsistencies.

1 Introduction

In order to defend the resources of an information system against unauthorized actions, a *security policy* must be defined by the administrator of an information system, i.e. a set of rules stating what is permitted and what is prohibited in a system during normal operations. Once specified the complete set of prohibitions and permissions, the administrator must decide which *security mechanisms* to use in order to enforce the security policy. This enforcement consists in distributing the security rules expressed in this policy over different security components, such as filtering routers and firewalls. This implies cohesion of the security functions supplied by these components. Indeed, security rules deployed over different components must be consistent, addressing the same decisions under equivalent conditions, and not repeating the same actions more than once.

A first solution to ensure these requirements is by applying a formal security model to express network security policies. In [12], for example, an access control language

based on XML syntax and supported by the access control model Or-BAC [1] is proposed for specifying access control meta-rules and, then, refined into different firewall configuration rules through XSLT transformations. In [15], another top-down proposal based on the RBAC model [18] is also suggested for such a purpose. However, and although the use of formal models ensures cohesion, completeness and optimization as built-in properties, in most cases, administrators are usually reluctant to define a whole security policy from scratch, and they expect to recycle existing configurations previously deployed over a given system.

A second solution to guarantee consistent and non-redundant firewall configurations consists in analyzing and fixing rules already deployed. In [14], for example, a taxonomy of conflicts in security policies is presented, and two main categories are proposed: (1) intra-firewall anomalies, which refer to those conflicts that might exist within the local set of rules of a given firewall; (2) inter-firewall anomalies, which refer to those conflicts that might exist between the configuration rules of different firewalls that match the same traffic. The authors in [14] propose, moreover, an audit mechanism in order to discover and warn about these anomalies. In [4, 3], we pointed out to some existing limitations in [14], and presented an alternative set of anomalies and audit algorithms that detect, report, and eliminate those intra- and inter-component inconsistencies existing on distributed security setups — where both firewalls and NIDSs are in charge of the network security policy.

The main drawback of the first solution, i.e., refinement processes such as [12, 15], relies on the necessity of formally writing a global security policy from scratch, as well as a deep knowledge of a given formal model. This reason might explain why this solution is not yet widely used, and most of the times policies are simply deployed based on the expertise and flair of security administrators. The main drawback of the second solution, i.e., audit processes such as [14, 4] for analyzing local and distributed security setups, relies on the lack of knowledge about the deployed

policy from a global point of view — which is very helpful for maintenance and troubleshooting tasks.

In this paper we propose to combine both solutions to better guarantee the requirements specified for a given network access control policy. Our procedure consists of two main steps. In the first step, the complete set of local policies — already deployed over each firewall of a given system — are aggregated, and a global security policy is derived. It is then possible to update, analyze, and redeploy such a global security policy into several local policies — yet free of anomalies — in a further second step. We need, moreover, a previous step for retrieving all those details of the system’s topology which might be necessary during the aggregation and deployment processes (cf. Section 2). The use of automatic network tools, such as [19], may allow us to automatically generate this information, and properly manage any change within the system.

The rest of this paper has been organized as follows. We first present in Section 2 the formalism we use to specify filtering rules, and the network model we use to represent the topology of the system. We describe in Section 3 our mechanisms to aggregate and deploy firewall configuration rules, and prove the correctness of such mechanisms. We present some related work in Section 4, and close the paper in Section 5 with some conclusions and future work.

2 Rules, Topology and Anomalies

We recall in this section some of the definitions previously introduced in [4, 3]. We first define a filtering rule in the form $R_i : \{cnd_i\} \rightarrow decision_i$, where i is the relative position of the rule within the set of rules, $decision_i$ is a boolean expression in $\{accept, deny\}$, and $\{cnd_i\}$ is a conjunctive set of condition attributes (*protocol, source, destination*, and so on), such that $\{cnd_i\}$ equals $A_1 \wedge A_2 \wedge \dots \wedge A_p$, and p is the number of condition attributes of a given filtering rule.

We define now a set of functions to determine which firewalls of the system are crossed by a given packet knowing its source and destination. Let F be a set of firewalls and let Z be a set of zones. We assume that each pair of zones in Z are mutually disjoint, i.e., if $z_i \in Z$ and $z_j \in Z$ then $z_i \cap z_j = \emptyset$. We define the predicates $connected(f_1, f_2)$ (which becomes *true* whether there exists, at least, one interface connecting firewall f_1 to firewall f_2) and $adjacent(f, z)$ (which becomes *true* whether the zone z is interfaced to firewall f). We then define a set of paths, P , as follows. If $f \in F$ then $[f] \in P$ is an atomic path. Similarly, if $[p.f_1] \in P$ (be “.” a concatenation functor) and $f_2 \in F$, such that $f_2 \notin p$ and $connected(f_1, f_2)$, then $[p.f_1.f_2] \in P$. Let us now define functions $first$, $last$, and $tail$ from P in F such that if p is a path, then $first(p)$ corresponds to the first firewall in the path, $last(p)$ corre-

sponds to the last firewall in the path, and $tail(f, p)$ corresponds to rest of firewalls in the path after firewall f . We also define the order functor between paths as $p_1 \leq p_2$, such that path p_1 is shorter than p_2 , and where all the firewalls within p_1 are also within p_2 . We define functions $route$ such that $p \in route(z_1, z_2)$ iff path p connects zone z_1 to zone z_2 , i.e., $p \in route(z_1, z_2)$ iff $adjacent(first(p), z_1)$ and $adjacent(last(p), z_2)$; and $minimal_route$ (or MR for short), such that $p \in MR(z_1, z_2)$ iff the following conditions hold: (1) $p \in route(z_1, z_2)$; (2) there does not exist $p' \in route(z_1, z_2)$ such that $p' < p$.

Let us finally close this section by overviewing the complete set of anomalies defined in our previous work [4, 3]:

Intra-firewall anomalies

- **Shadowing** – A configuration rule R_i is shadowed in a set of configuration rules R whether such a rule never applies because all the packets that R_i may match, are previously matched by another rule, or combination of rules, with higher priority.
- **Redundancy** – A configuration rule R_i is redundant in a set of configuration rules R whether the following conditions hold: (1) R_i is not shadowed by any other rule or set of rules; (2) when removing R_i from R , the security policy does not change.

Inter-firewall anomalies

- **Irrelevance** – A configuration rule R_i is irrelevant in a set of configuration rules R if one of the following conditions holds: (1) Both source and destination address are within the same zone; (2) The firewall is not within the minimal route that connects the source zone to the destination zone.
- **Full/Partial-redundancy** – A redundancy anomaly¹ occurs between two firewalls whether the firewall closest to the destination zone blocks (completely or partially) traffic that is already blocked by the first firewall.
- **Full/Partial-shadowing** – A shadowing anomaly occurs between two firewalls whether the one closest to the destination zone does not block traffic that is already blocked by the first firewall.
- **Full/Partial-misconnection** – A misconnection anomaly occurs between two firewalls whether the closest firewall to the source zone allows all the traffic — or just a part of it — that is denied by the second one.

¹Although this kind of redundancy is sometimes expressly introduced by network administrators (e.g., to guarantee the forbidden traffic will not reach the destination), it is important to report it since, if such a rule is applied, we may conclude that at least one of the redundant components is wrongly working.

3 Proposed Mechanisms

The objective of our proposal is the following. From a set F of firewalls initially deployed over a set Z of zones, and if neither intra- nor inter-firewall anomalies apply over such a setup, we aim to derive a single global security police setup R , also free of anomalies. Then, this set of rules R can be maintained and updated² as a whole, as well as redeployed over the system through a further refinement process. We present in the following the main processes of our proposal.

3.1 Aggregation of Policies

Our aggregation mechanism works as follows. During an initial step (not covered in this paper) it gathers all those details of the system's topology which might be necessary during the rest of stages. The use of network tools, such as [19], allows us to properly manage this information, like the set F of firewalls, the set of configurations rules $f[rules]$ of each firewall $f \in F$, the set Z of zones of the system, and some other topological details defined in Section 2. An analysis of intra-firewall anomalies is then performed within the first stage of the aggregation process, in order to discover and fix any possible anomaly within the local configuration of each firewall $f \in F$. In the next step, an analysis of inter-firewall anomalies is performed at the same time that the aggregation of polices into R also does. If an anomaly within the initial setup is discovered, then an aggregation error warns the officer and the process quits. Conversely, if no inter-firewall anomalies are found, then a global set of rules R is generated and so returned as a result of the whole aggregation process.

We present in Algorithm 1 our proposed aggregation process. The input data is a set F of firewalls whose configurations we want to fold into a global set of rules R . For reasons of clarity, we assume in our algorithm that one can access the elements of a set as a linked-list through the operator $element_i$. We also assume one can add new values to the list as any other normal variable does ($element_i \leftarrow value$), as well as to both remove and initialize elements through the addition of an empty set ($element_i \leftarrow \emptyset$). The internal order of elements from the linked-list, moreover, keeps with the relative ordering of elements.

The aggregation process consists of two main phases. During the first phase (cf. lines 2 and 3 of Algorithm 1), and through an iterative call to the auxiliary function *policy-rewriting* (cf. Algorithm 4), it analyzes the complete set F of firewalls, in order to discover and remove any possible intra-firewall anomaly. Thus, after this first stage, no useless rules in the local configuration of any firewall $f \in F$ might exist. We refer to Section 3.2 for a more detailed description of this function.

²Due to space limitations, these operations are not covered in the paper.

Algorithm 1: aggregation(F)

```

1 /*Phase 1*/
2 foreach  $f_1 \in F$  do
3   policy-rewriting( $f_1[rules]$ );
4 /*Phase 2*/
5  $R \leftarrow \emptyset$ ;
6  $i \leftarrow 0$ ;
7 foreach  $f_1 \in F$  do
8   foreach  $r_1 \in f_1[rules]$  do
9      $Z_s \leftarrow \{z \in Z \mid z \cap source(r_1) \neq \emptyset\}$ ;
10     $Z_d \leftarrow \{z \in Z \mid z \cap destination(r_1) \neq \emptyset\}$ ;
11    foreach  $z_1 \in Z_s$  do
12      foreach  $z_2 \in Z_d$  do
13        if ( $z_1 = z_2$ ) or ( $f_1 \notin MR(z_1, z_2)$ ) then
14          aggregationError();
15          return  $\emptyset$ ;
16        else if ( $r_1[decision] = "accept"$ ) then
17          foreach  $f_2 \in MR(z_1, z_2)$  do
18             $f_2rd \leftarrow \emptyset$ ;
19             $f_2rd \leftarrow \{r_2 \in f_2 \mid r_1 \vee r_2 \wedge$ 
20               $r_2[decision] = "deny"\}$ ;
21            if ( $\neg empty(f_2rd)$ ) then
22              aggregationError();
23              return  $\emptyset$ ;
24            else
25               $f_2ra \leftarrow \emptyset$ ;
26               $f_2ra \leftarrow \{r_2 \in f_2 \mid r_1 \vee r_2 \wedge$ 
27                 $r_2[decision] = "accept"\}$ ;
28              foreach  $r_2 \in f_2ra$  do
29                 $R_i \leftarrow R_i \cup r_2$ ;
30                 $R_i[source] \leftarrow z_1$ ;
31                 $R_i[destination] \leftarrow z_2$ ;
32                 $i \leftarrow (i + 1)$ ;
33                 $r_2 \leftarrow \emptyset$ ;
34          else if ( $f_1 = first(MR(z_1, z_2))$ ) then
35             $f_3r \leftarrow \emptyset$ ;
36            foreach  $f_3 \in tail(f_1, MR(z_1, z_2))$  do
37               $f_3r \leftarrow \{r_3 \in f_3 \mid r_1 \vee r_3\} \cup f_3r$ ;
38            if ( $\neg empty(f_3r)$ ) then
39              aggregationError();
40              return  $\emptyset$ ;
41            else
42               $R_i \leftarrow R_i \cup r_1$ ;
43               $R_i[source] \leftarrow z_1$ ;
44               $R_i[destination] \leftarrow z_2$ ;
45               $i \leftarrow (i + 1)$ ;
46               $r_1 \leftarrow \emptyset$ ;
47          else
48            aggregationError();
49            return  $\emptyset$ ;
50 policy-rewriting( $R$ );
51 return  $R$ ;
```

During the second phase (cf. lines 5–51 of Algorithm 1), the aggregation of firewall configurations is performed as follows. For each permission configured in a firewall $f \in F$, the process folds the whole chain³ of permissions within the components on the minimal route from the source zone to the destination zone; and for each prohibition, it directly keeps such a rule, assuming it becomes to the closest firewall to the source, and no more prohibitions should be placed on the minimal route from the source zone to the destination zone. Moreover, and while the aggregation of policies is being performed, an analysis of inter-firewall anomalies is also applied in parallel. Then, if any inter-firewall anomaly is detected during the aggregation of rules $R \leftarrow aggregation(F)$, a message of error is raised and the process quits.

Let us for example assume that during the aggregation process, a filtering rule $r_i \in f_i[rules]$ presents an inter-firewall irrelevance, i.e., r_i is a rule that applies to a source zone z_1 and a destination zone z_2 (such that $s = z_1 \cap source(r_i) \neq \emptyset$, $d = z_2 \cap destination(r_i) \neq \emptyset$) and either z_1 and z_2 are the same zone, or firewall f_i is not in the path $[f_1, f_2, \dots, f_k] \in MR(z_1, z_2)$. In this case, we can observe that during the folding process specified by Algorithm 1, the statement of line 13, i.e., $(z_1 = z_2)$ or $(f_i \notin MR(z_1, z_2))$, becomes *true* and, then, the aggregation process finishes with an error and returns an empty set of rules (cf. statements of lines 14 and 15). Similarly, let us assume that $r_i \in f_i[rules]$ presents an inter-firewall redundancy, i.e., r_i is a prohibition that applies to a source zone z_1 and a destination zone z_2 (such that $s = z_1 \cap source(r_i) \neq \emptyset$, $d = z_2 \cap destination(r_i) \neq \emptyset$, and $[f_1, f_2, \dots, f_k] \in MR(z_1, z_2)$) and firewall f_i is not the first component in $MR(z_1, z_2)$. In this case, we can observe that during the folding process specified by Algorithm 1, the statement of line 34, i.e., $f_i = first(MR(z_1, z_2))$, becomes *false* and, then, the aggregating process finishes with an error and returns an empty set of rules.

Let us now assume that $r_i \in f_i[rules]$ presents an inter-firewall shadowing, i.e., r_i is a permission that applies to a source zone z_1 and a destination zone z_2 such that there exists an equivalent prohibition r_j that belongs to a firewall f_j which, in turn, is closer to the source zone z_1 in $MR(z_1, z_2)$. In this case, we can observe that during the folding process specified by Algorithm 1, the statement of line 38 detects that, after a prohibition in the first firewall of $MR(z_1, z_2)$, i.e., $f_j = first(MR(z_1, z_2))$, there is, at least, a permission r_i that correlates the same attributes. Then, the aggregating process finishes with an error and returns an empty set of rules. Let us finally assume that $r_i \in f_i[rules]$ presents an inter-firewall misconnection, i.e.,

³The operator “ \sim ” is used within Algorithm 1 to denote that two rules r_i and r_j are correlated if every attribute in r_i has a non empty intersection with the corresponding attribute in r_j .

r_i is a prohibition that applies to a source zone z_1 and a destination zone z_2 such that there exists, at least, a permission r_j that belongs to a firewall f_j closer to the source zone z_1 in $MR(z_1, z_2)$. In this case, we can observe that during the folding process specified by Algorithm 1, the statement of line 21 detects this anomaly and, then, the process finishes with an error and returns an empty set of rules.

It is straightforward then to conclude that whether no inter-firewall anomalies apply to any firewall $f \in F$, our aggregation process returns a global set of filtering rules R with the union of all the filtering rules previously deployed over F . It is yet necessary to perform a post-process of R , in order to avoid the redundancy of all permissions, i.e., *accept* rules, gathered during the aggregating process. In order to do so, the aggregation process calls at the end of the second phase (cf. line 50 of Algorithm 1) to the auxiliary function *policy-rewriting* (cf. Algorithm 4). We offer in the following a more detailed description of this function.

3.2 Policy Rewriting

We recall in this section our audit process to discover and remove rules that never apply or are redundant in local firewall policies [10, 11]. The process is based on the analysis of relationships between the set of configuration rules of a local policy. Through a rewriting of rules, it derives from an initial set R to an equivalent one $Tr(R)$ completely free of dependencies between attributes, i.e., without either redundant or shadowed rules. The whole process is split in three main functions (cf. algorithms 2, 3 and 4).

The first function, *exclusion* (cf. Algorithm 2), is an auxiliary process which performs the exclusion of attributes between two rules. It receives as input two rules, A and B , and returns a third rule, C , whose set of condition attributes is the exclusion of the set of conditions from A over B . We represent the attributes of each rule in the form of *Rule[cnd]*⁴ as a boolean expression over p possible attributes (such as source, destination, protocol, ports, and so on). Similarly, we represent the decision of the rule in the form *Rule[decision]* as a boolean variable whose values are in $\{accept, deny\}$. Moreover, we use two extra elements for each rule, in the form *Rule[shadowing]* and *Rule[redundancy]*, as two boolean variables in $\{true, false\}$ to store the reason for why a rule may disappear during the process.

The second function, *testRedundancy* (cf. Algorithm 3), is a boolean function in $\{true, false\}$ which, in turn, applies the transformation *exclusion* (cf. Algorithm 2) over a set of configuration rules to check whether the first rule is redundant, i.e., applies the same policy, regarding the rest of rules.

⁴Due to space limitation, we use the notation A_i and B_i as an abbreviation of both $A[cnd][i]$ and $B[cnd][i]$ during the statements of lines 6–12.

Finally, the third function, *policy-rewriting* (cf. Algorithm 4), performs the whole process of detecting and removing the complete set of intra-firewall anomalies. It receives as input a set R of rules, and performs the audit process in two different phases.

Algorithm 2: $\text{exclusion}(B,A)$

```

1  $C[\text{cnd}] \leftarrow \emptyset;$ 
2  $C[\text{decision}] \leftarrow B[\text{decision}];$ 
3  $C[\text{shadowing}] \leftarrow \text{false};$ 
4  $C[\text{redundancy}] \leftarrow \text{false};$ 
5 forall the elements of  $A[\text{cnd}]$  and  $B[\text{cnd}]$  do
6   if  $((A_1 \cap B_1) \neq \emptyset$  and  $(A_2 \cap B_2) \neq \emptyset$  and ... and
    $(A_p \cap B_p) \neq \emptyset)$  then
7      $C[\text{cnd}] \leftarrow C[\text{cnd}] \cup$ 
8      $\{(B_1 - A_1) \wedge B_2 \wedge \dots \wedge B_p,$ 
9      $(A_1 \cap B_1) \wedge (B_2 - A_2) \wedge \dots \wedge B_p,$ 
10     $(A_1 \cap B_1) \wedge (A_2 \cap B_2) \wedge (B_3 - A_3) \wedge \dots \wedge B_p,$ 
11     $\dots$ 
12     $(A_1 \cap B_1) \wedge \dots \wedge (A_{p-1} \cap B_{p-1}) \wedge (B_p - A_p)\};$ 
13   else
14      $C[\text{cnd}] \leftarrow (C[\text{cnd}] \cup B[\text{cnd}]);$ 
15 return  $C;$ 

```

Algorithm 3: $\text{testRedundancy}(R,r)$

```

1  $i \leftarrow 1;$ 
2  $\text{temp} \leftarrow r;$ 
3 while  $\neg \text{test}$  and  $(i \leq \text{count}(R))$  do
4    $\text{temp} \leftarrow \text{exclusion}(\text{temp}, R_i);$ 
5   if  $\text{temp}[\text{cnd}] = \emptyset$  then
6     return  $\text{true};$ 
7    $i \leftarrow (i + 1);$ 
8 return  $\text{false};$ 

```

During the first phase, any possible shadowing between rules with different decision values is marked and removed by iteratively applying function *exclusion* (cf. Algorithm 2). The resulting set of rules obtained after the execution of the first phase is again analyzed when applying the second phase.

Each rule is first analyzed, through a call to function *testRedundancy* (cf. Algorithm 3), to those rules written after the checked rule but that can apply the same decision to the same traffic. If such a test of redundancy becomes *true*, the rule is marked as redundant and then removed. Otherwise, its attributes are then excluded from the rest of equivalent rules but with less priority in the order. In this way, if any shadowing between rules with the same decision remained undetected during the first phase, it is then marked and removed.

Algorithm 4: $\text{policy-rewriting}(R)$

```

1  $n \leftarrow \text{count}(R);$ 
2 /*Phase 1*/
3 for  $i \leftarrow 1$  to  $(n - 1)$  do
4   for  $j \leftarrow (i + 1)$  to  $n$  do
5     if  $R_i[\text{decision}] \neq R_j[\text{decision}]$  then
6        $R_j \leftarrow \text{exclusion}(R_j, R_i);$ 
7       if  $R_j[\text{cnd}] = \emptyset$  then
8          $R_j[\text{shadowing}] \leftarrow \text{true};$ 
9 /*Phase 2*/
10 for  $i \leftarrow 1$  to  $(n - 1)$  do
11    $R_a \leftarrow \{r_k \in R \mid n \geq k > i$  and
12    $r_k[\text{decision}] = r_i[\text{decision}]\};$ 
13   if  $\text{testRedundancy}(R_a, R_i)$  then
14      $R_i[\text{cnd}] \leftarrow \emptyset;$ 
15      $R_i[\text{redundancy}] \leftarrow \text{true};$ 
16   else
17     for  $j \leftarrow (i + 1)$  to  $n$  do
18       if  $R_i[\text{decision}] = R_j[\text{decision}]$  then
19          $R_j \leftarrow \text{exclusion}(R_j, R_i);$ 
20         if  $(\neg R_j[\text{redundancy}]$  and
21          $R_j[\text{cnd}] = \emptyset)$  then
22            $R_j[\text{shadowing}] \leftarrow \text{true};$ 

```

Based on the processes defined in algorithms 2, 3, and 4, we can prove⁵ the following theorem:

Theorem 1 *Let R be a set of filtering rules and let $\text{Tr}(R)$ be the resulting filtering rules obtained by applying Algorithm 4 to R . Then the following statements hold: (1) R and $\text{Tr}(R)$ are equivalent; (2) Ordering the rules in $\text{Tr}(R)$ is no longer relevant; (3) $\text{Tr}(R)$ is free from both shadowing and redundancy.*

3.3 Deployment of Rules

We finally present in Algorithm 5 our proposed refinement mechanism for the deployment of an updated global set of rules. The deployment strategy defined in the algorithm is the following. Let F be the set of firewalls that partitions the system into the set Z of zones. Let R be the set of configuration rules resulting from the maintenance of a given global set of rules obtained from the aggregation process presented in Section 3.1 (cf. Algorithm 1). Let $r \in R$ be a configuration rule that applies to a source zone z_1 and a

⁵A set of proofs to validate Theorem 1, as well as a complexity analysis of function *policy-rewriting* (cf. Algorithm 4) and its performance in a research prototype, is provided in [10].

destination zone z_2 , such that $s = z_1 \cap \text{source}(r) \neq \emptyset$ and $d = z_2 \cap \text{destination}(r) \neq \emptyset$. Let r' be a rule identical to r except that $\text{source}(r') = s$ and $\text{destination}(r') = d$. Let us finally assume that $[f_1, f_2, \dots, f_k] \in MR(z_1, z_2)$. Then, any rule $r \in R$ is deployed over the system as follows:

- If $r[\text{decision}] = \text{accept}$ then deploy a permission r' on every firewall on the minimal route from source s to destination d .
- If $r[\text{decision}] = \text{deny}$ then deploy a single⁶ prohibition r' on the most-upstream firewall (i.e., the closest firewall to the source) of the minimal route from source s to destination d . If such a firewall does not exist, then generate a deployment error message.

Algorithm 5: `deployment(R,Z)`

```

1 policy-rewriting (R);
2 foreach  $r_1 \in R$  do
3    $Z_s \leftarrow \{z \in Z \mid z \cap \text{source}(r_1) \neq \emptyset\}$ ;
4    $Z_d \leftarrow \{z \in Z \mid z \cap \text{destination}(r_1) \neq \emptyset\}$ ;
5   foreach  $z_1 \in Z_s$  do
6     foreach  $z_2 \in Z_d$  do
7       if  $r_1[\text{decision}] = \text{"accept"}$  then
8         foreach  $f_1 \in MR(z_1, z_2)$  do
9            $r'_1 \leftarrow r$ ;
10           $r'_1[\text{source}] \leftarrow Z_1$ ;
11           $r'_1[\text{destination}] \leftarrow Z_2$ ;
12           $f_1[\text{rules}] \leftarrow f_1[\text{rules}] \cup r'$ ;
13        else if  $r_1[\text{decision}] = \text{"deny"}$  then
14           $f_1 \leftarrow \text{first}(MR(z_1, z_2))$ ;
15          if ( $\neg \text{empty}(f_1)$ ) then
16             $r'_1 \leftarrow r$ ;
17             $r'_1[\text{source}] \leftarrow Z_1$ ;
18             $r'_1[\text{destination}] \leftarrow Z_2$ ;
19             $f_1[\text{rules}] \leftarrow f_1[\text{rules}] \cup r'$ ;
20          else
21            deploymentError ();
22            exit ();

```

It is straightforward now to prove that the deployment of a given set of rules R through Algorithm 5 is free of either intra- and/or inter-firewall anomalies (cf. Section 2). On the one hand, during the earliest stage of Algorithm 5, the complete set of rules in R is analyzed and, if necessary, fixed with our *policy-rewriting* process (cf. Section 3.2, Algorithm 4). Then, by Theorem 1, we can guarantee that neither shadowed nor redundant rules might exist in R . Moreover, it also allows us to guarantee that the order between

⁶This decision is a choice for avoiding inter-firewall redundancy in the resulting setup.

rules in R is not relevant. On the other hand, the use of the deployment strategy defined above allows us to guarantee that the resulting setup is free of inter-firewall anomalies. First, since each permission r_a in R opens a flow of permissions over all the firewalls within the minimal routes from the source to the destination pointed by r_a , and since any other rule r' in R cannot match the same traffic that r_a matches, we can guarantee that neither inter-firewall shadowing nor inter-firewall misconnection can appear in the resulting setup. Second, since each prohibition r_d in R is deployed just once in the closest firewall to the source pointed by r_d , and since any other rule r' in R cannot match the same traffic that r_d matches, we can guarantee that any inter-firewall redundancy can appear in the resulting setup.

4 Related Work

A first solution to deploy access control policies free of errors is by applying a refinement mechanism. Hence, following such a top-down mechanism, one can deploy a global security policy into several component's configurations [12, 7, 15].

In [12], for example, a formal approach based on the OrBAC model [1] is presented for this purpose. There, a set of filtering rules, whose syntax is specific to a given firewall, may be generated using a transformation process. The authors in [7], on the other hand, use the concept of roles to define network capabilities and refinement of policies. Indeed, they propose the use of an inheritance mechanism through a hierarchy of entities to automatically generate permissions.

However, their work does not fix, from our point of view, clear semantics, and their concept of role becomes ambiguous as we pointed out in [12]. Another work based on policy refinement is the RBNS model [15]. However, and although the authors claim that their work is based on the RBAC model [18], it seems that they only keep from this model only the concept of role. Indeed, the specification of network entities and role and permission assignments are not rigorous and does not fit any reality [12].

The use of these refinement proposals [12, 7, 15] ensures cohesion, completeness and optimization as built-in properties. However, it is not always enough to ensure that the firewall configuration is completely free of errors and, often, administrators are reluctant to follow such a proposal. For this reason, we extended in this paper the approach presented in [12], offering to administrators the possibility of aggregating existing configurations before moving to such a refinement approach.

Support tools, on the other hand, are intended to directly assist administrators in their task of configuring from scratch firewall configurations. Firewall Builder [16], for example, provides a user interface to be used to specify a network access control policy and then this policy is auto-

matically translated into various firewall configuration languages such as NetFilter [20], IpFilter [17] or Cisco PIX [9]. It also provides higher portability. For instance, if in a given network infrastructure, IpFilter is replaced by NetFilter, it will not be necessary to completely reconfigure NetFilter. Firewall Builder will automatically generate the rules necessary to configure this firewall.

However, we observed some problems when using Firewall Builder. First, we noticed that it might generate incorrect rules. In the case of NetFilter, for example, we experienced the generation of rules associated to FORWARD when they should be associated to OUTPUT and INPUT chains. Second, we noticed the generation of redundant rules, although such redundancy was not specified within the policy. Third, it includes a mechanism called *shadowing* to detect redundancy in the policy. However, this shadowing mechanism only detects simple redundancy that corresponds to trivial equality or inclusion between zones. More complex redundancies (as the anomalies defined in Section 2) are unfortunately not detected.

Some other proposals, such as [14, 21, 4, 3], provide means to directly manage the discovery of anomalies from a bottom-up approach. For instance, the authors in [14] propose a set of algorithms to detect policy anomalies in both single- and multi-firewall configuration setups. In addition to the discovery process, their approach also attempts an optimal insertion of arbitrary rules into an existing configuration, through a tree based representation of the filtering criteria. Nonetheless, we consider their approach as incomplete. Their discovery approach is not complete since, given a single- or multiple-component security policy, their detection algorithms are based on the analysis of relationships between rules two by two. This way, errors due to the union of rules are not explicitly considered (as our approach presented in [4, 3] does).

Although in [5] the authors pointed out to this problematic, claiming that they break down the initial set of rules into an equivalent set of rules free of overlaps between rules, no specific algorithms have been provided for solving it. From our point of view, the proposal presented in [21] best addresses such a problem, although it also presents some limitations. For instance, we can easily find situations where the proposal presented in [21] reports partial redundancies instead of a single full redundancy. Moreover, neither [14] nor [21] address, as we do in this paper by extending the approach presented in [4, 12], a folding process for combining both analysis and refinement strategies.

5 Conclusions

The existence of errors or anomalies in the configuration of network security components, such as filtering routers or firewalls, is very likely to degrade the security policy of a

system [13]. This is a serious problem which must be solved since, if not handled correctly, it can lead to unauthorized parties to get the control of such a system.

We introduced in Section 1 two main strategies to set firewall configurations free of errors. The first approach is to apply a formal security model — such as the formal model we presented in [12] — to express the security policy of the access control for the network, and to generate the specific syntax for each given firewall from this formal policy — for instance, by using XSLT transformations from the formal policy to generate specific Netfilter configuration rules [20]. A second approach is to apply an analysis process of existing configurations, in order to detect configuration errors and to properly eliminate them. In [4, 3], for instance, we presented an audit process based on this second strategy to set a distributed security scenario free of misconfiguration.

We presented in Section 3 how to combine both approaches in order to better guarantee the requirements specified for a given network access control policy. Thus, from an initial bottom-up approach, we can analyze existing configurations already deployed into a given system, in order to detect and correct potential anomalies or configuration errors. Once verified those setups, we offer to the administrator a folding mechanism to aggregate the different configurations into a global security policy to, finally, express by using a sole formal model, the security policy as a whole. The security officer can then perform maintenance tasks over such a single point, and then, unfold the changes into the existing security components of the system.

As work in progress, we are actually evaluating the implementation of the strategy presented in this paper by combining both the refinement process presented in [12] and the audit mechanism presented in [4, 3] (both of them implemented through a scripting language as a web service [8]). Although this first research prototype demonstrates the effectiveness of our approach, more evaluations should be done to study the real impact of our proposal for the maintenance and deployment of complex production scenarios. We plan to address these evaluations and discuss the results in a forthcoming paper.

On the other hand, and as future work, we are currently studying how to extend our approach in the case where the security architecture includes not only firewalls but also IDS/IPS, and IPSec devices. Though there is a real similarity between the parameters of those devices' rules (as we partially show in [4, 3] for the analysis of anomalies), more investigation has to be done in order to extend the approach presented in this paper. In parallel to this work, we are also considering to extend our approach to the managing of stateful policies.

Acknowledgements

This work was supported by funding from the French ministry of research, under the *ACI DESIRS* project; the Spanish Government (CICYT) projects *TIC2003-02041* and *SEG2004-04352-C04-04*; and the Catalan Government (DURSI) grants *2006FIC00229* and *2006BE00569*.

References

- [1] Abou el Kalam, A., Baida, R. E., Balbiani, P., Benferhat, S., Cuppens, F., Deswarte, Y., Miège, A., Saurel, C., and Trouessin, G. Organization Based Access Control. In *IEEE 4th Intl. Workshop on Policies for Distributed Systems and Networks*, pages 120–131, Lake Como, Italy, 2003.
- [2] Alfaro, J. G., Cuppens, F., and Cuppens-Boualahia, N. Analysis of Policy Anomalies on Distributed Network Security Setups. In *11th European Symposium On Research In Computer Security (Esorics 2006)*, pages 496–511, Hamburg, Germany, 2006.
- [3] Alfaro, J. G., Cuppens, F., and Cuppens-Boualahia, N. Towards Filtering and Alerting Rule Rewriting on Single-Component Policies. In *Intl. Conference on Computer Safety, Reliability, and Security (Safecom 2006)*, pages 182–194, Gdansk, Poland, 2006.
- [4] Alfaro, J. G., Cuppens-Boualahia, N., and Cuppens, F. Complete Analysis of Configuration Rules to Guarantee Reliable Network Security Policies. In *International Journal of Information Security*, Springer, 7(2):103-122, April 2008.
- [5] Al-Shaer, E. S., Hamed, H. H., and Masum, H. Conflict Classification and Analysis of Distributed Firewall Policies. In *IEEE Journal on Selected Areas in Communications*, 23(10):2069–2084, 2005.
- [6] Al-Shaer, E. S., Hamed, H. H., and Masum, H. Modeling and Verification of IPsec and VPN Security Policies. In *13TH IEEE International Conference on Network Protocols (ICNP'05)*, pages 259–278, 2005.
- [7] Bartal, Y., Mayer, A., Nissim, K., and Wool, A. Firmato: A novel firewall management toolkit. In *IEEE Symposium on Security and Privacy*, pages 17–31, Oakland, California, 1999.
- [8] Castagnetto, J. et al. (1999). *Professional PHP Programming*. Wrox Press Inc, ISBN 1-86100-296-3.
- [9] Chapman, D. and Fox, A. (2001). *Cisco Secure PIX Firewalls*. Cisco Press.
- [10] Cuppens, F., Cuppens-Boualahia, N., and Alfaro, J. G. Detection and Removal of Firewall Misconfiguration. In *Intl. Conference on Communication, Network and Information Security (CNIS05)*, pages 154–162, 2005.
- [11] Cuppens, F., Cuppens-Boualahia, N., and Alfaro, J. G. Misconfiguration Management of Network Security Components. In *7th Intl. Symposium on System and Information Security*, Sao Paulo, Brazil, 2005.
- [12] Cuppens, F., Cuppens-Boualahia, N., Sans, T. and Mieke, A. A formal approach to specify and deploy a network security policy. In *2nd Workshop on Formal Aspects in Security and Trust*, pages 203–218, 2004.
- [13] D. Geer. Just How Secure Are Security Products? *IEEE Computer*, 37(6):14–16, June 2004.
- [14] Hamed, H. H. and Al-Shaer, E. S. Taxonomy of conflicts in network security policies. In *IEEE Communications Magazine*, 44(3):134-141, 2006.
- [15] Hassan, A. and Hudec, L. Role Based Network Security Model: A Forward Step towards Firewall Management. In *Workshop On Security of Information Technologies*, Algiers, 2003.
- [16] Kurland, V. (2004). Firewall Builder. In *11th DFN-CERT Workshop*, Hamburg, Germany, 2004.
- [17] Reed, D. IP Filter. [Online]. Available from: <http://www.ja.net/CERT/Software/ip-filter/ip-filter.html>
- [18] Sandhu, R., Coyne, E. J., Feinstein, H. L., and Youman, C. E. Role-Based Access Control Models. *IEEE Computer*, 29(2):38–47, 1996.
- [19] Skybox Security, Inc. Security Risk Management and Network Change Management Solution from Skybox Security.
- [20] Welte, H., Kadlecik, J., Josefsson, M., McHardy, P., and et al. The netfilter project: firewalling, nat and packet mangling for linux 2.4x and 2.6.x. [Online]. Available from: <http://www.netfilter.org/>
- [21] Yuan, L., Mai, J., Su, Z., Chen, H., Chuah, C., and Mohapatra, P. FIREMAN: a toolkit for FIREwall Modeling and ANalysis. In *IEEE Symposium on Security and Privacy*, pages 199–213, 2006.