

Management of Exceptions on Access Control Policies*

J. G. Alfaro^{1,2}, F. Cuppens¹, and N. Cuppens-Boulahia¹

¹ GET/ENST-Bretagne, 35576 Cesson Sévigné - France,
{frederic.cuppens,nora.cuppens}@enst-bretagne.fr

² Universitat Oberta de Catalunya, 08018 Barcelona - Spain,
joaquin.garcia-alfaro@uoc.edu

Abstract. The use of languages based on positive or negative expressiveness is very common for the deployment of security policies (i.e., deployment of permissions and prohibitions on firewalls through single-handed positive or negative condition attributes). Although these languages may allow us to specify any policy, the single use of positive or negative statements alone leads to complex configurations when excluding some specific cases of general rules that should always apply. In this paper we survey such a management and study existing solutions, such as ordering of rules and segmentation of condition attributes, in order to settle this lack of expressiveness. We then point out to the necessity of full expressiveness for combining both negative and positive conditions on firewall languages in order to improve this management of exceptions on access control policies. This strategy offers us a more efficient deployment of policies, even using fewer rules.

1 Introduction

Current firewalls are still being configured by security officers in a manual fashion. Each firewall usually provides, moreover, its own configuration language that, most of the times, present a lack of expressiveness and semantics. For instance, most firewall languages are based on rules in the form $R_i : \{condition_i\} \rightarrow decision_i$, where i is the relative position of the rule within the set of rules, $decision_i$ is a boolean expression in $\{accept, deny\}$, and $\{condition_i\}$ is a conjunctive set of condition attributes, such as *protocol* (p), *source* (s), *destination* (d), *source port* (sport), *destination port* (dport), and so on. This conjunctive set of conditions attributes, i.e., $\{condition_i\}$, is mainly composed of either positive (e.g., A) or negative (e.g., $\neg A$) statements for each attribute, but does not allow us to combine both positive and negative statements (e.g., $A \wedge \neg B$) for a single attribute, as many other languages with

* This work was supported by funding from the French ministry of research, under the *ACI DESIRS* project; and the Spanish Government (CICYT) project *SEG2004-04352-C04-04*.

full expressive power, such as SQL-like languages [8], do. The use of more general access control languages, such as the eXtensible Access Control Markup Language (XACML) [10], also present such a lack of expressiveness. This fact leads to complex administration tasks when dealing with exclusion issues on access control scenarios, i.e., when some cases must be excluded of general rules that should always apply.

Let us suppose, for instance, the policy of a hospital where, in general, all doctors are allowed to consult patient’s medical records. Later, the policy changes and doctors going on strike are not allowed to consult medical records; but, as an exception to the previous one, and for emergencies purposes, doctors going on strike are still allowed to consult the records. Regarding the use of a language with expressiveness enough to combine both positive and negative statements, one may deploy the previous example as follows. We first assume the following definitions: (A) “Doctors”; (B) “Doctors going on strike”; (C) “Doctors working on emergencies”. We then deploy the hospital’s policy goals, i.e., (1) “In Hospital, doctors can access patient’s medical records.”; (2) “In Hospital, and only for emergency purposes, doctors going on strike can access patient’s medical records.”; through the following statement: “In Hospital, $(A \wedge (\neg B \vee C))$ can access patient’s medical records”.

The use of languages based on partial expressiveness may lead us to very complicated situations when managing this kind of configurations on firewalls and filtering routers. In this paper, we focus on this problem and survey current solutions, such as first and last matching strategies, segmentation of condition attributes, and partial ordering of rules. We then discuss how the combination of both negative and positive expressiveness on configuration languages may help us to improve those solutions. This strategy allows to perform a more efficient deployment of network access control policies, even using fewer rules, and properly manage exceptions and exclusion of attributes on firewall and filtering router configurations.

The rest of this paper is organized as follows. Section 2 recalls our motivation problem, by showing some representative examples, surveying related solutions, and overviewing their advantages and drawbacks. Section 3 then discusses our approach. Section 4 overviews some related work, and, finally, Section 5 closes the paper.

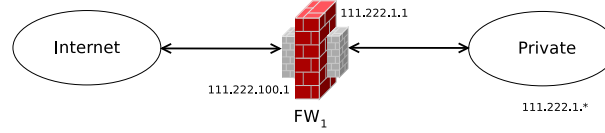
2 Management of Exceptions via Partial Expressiveness

Before going further in this section, let us start with an example to illustrate our motivation problem. We first consider the network setup shown in Figure 1(a), together with the following general premise: “In Private, all hosts can access web resources on the Internet”. We assume, moreover, that firewall FW_1 implements a closed default policy, specified in its set of rules at the last entry, in the form $R_n : deny$. Then, we deploy the premise over firewall FW_1 with the following rule:

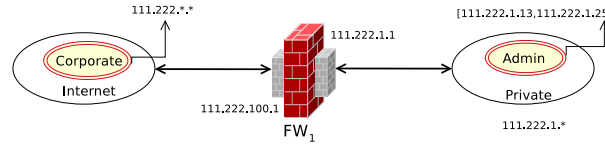
$$R_1 : (s \in Private \wedge d \in any \wedge p = tcp \wedge dport = 80) \rightarrow accept$$

Regarding the exclusion issues pointed out above, and according to the extended setup shown in Figure 1(b), let us assume that we must now apply the following three exceptions over the general security policy:

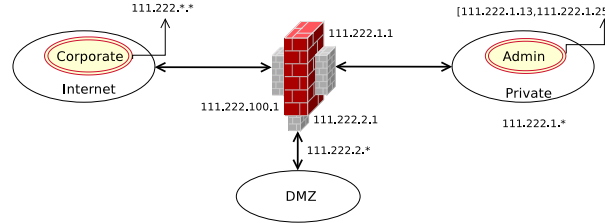
1. The interfaces of firewall FW_1 (i.e., $Interf\text{-}fw = \{111.222.1.1, 111.222.100.1\}$) are not allowed to access web resources on the Internet.
2. The hosts in Admin are not allowed to access web resources.
3. The hosts in Corporate do not belong to the zone Internet.



(a) Simple access control policy.



(b) Same policy with some excluded zones.



(c) Extended access control policy.

Fig. 1. Sample access control policy setups.

According to the first exception, we should exclude the IP address 111.222.1.1 from the hosts of *Private*. Similarly, we must exclude the whole set of hosts in zone *Admin* from the zone *Private*, and the whole set of hosts in zone *Corporate*, i.e., the range 111.222.*.*, from *Internet*. The use of a language with expressiveness enough to combine both positive and negative statements may allow us to deploy the previous policy goal, i.e., “All the hosts in $(Private \wedge \neg Admin \wedge \neg Interf\text{-}fw)$ are allowed to access web resources on $(Internet \wedge \neg Corporate)$ ”, as the following single rule:

$$R_1 : (s \in (Private \wedge \neg Admin \wedge \neg Interf\text{-}fw) \wedge d \in (any \wedge \neg Corporate) \wedge p = tcp \wedge dport = 80) \rightarrow accept$$

However, the lack of semantics and expressiveness of current firewall configuration languages (specially the impossibility for combining both positive and negative statements on single condition attributes) forces us to use different strategies to make up for this lack of expressiveness. We overview in the following sections some possible solutions for applying the previous example by means of such languages.

2.1 First Matching Strategy

Most firewalls solve the managing of exceptions by an ordering of rules. For instance, the configuration language for IPTables, the administration software used to configure GNU/Linux-based firewalls through the Netfilter framework, is based on a *first matching* strategy, i.e., the firewall is parsing the rules until a rule applies. When no rule applies, the decision depends on the default policy: in the case of an open policy, the packet is accepted whereas if the policy is closed, the packet is rejected. Other languages, like the configuration language of IPFilter, the administration software for configuring FreeBSD-, NetBSD- and Solaris 10-based firewalls, apply the opposite strategy, called *last matching*. Similar approaches have also been proposed in other security domains, such as the formal access control proposed in [9] to specify protection policies on XML databases. Through a first matching strategy, one may specify the handling of exceptions in the form $R_1 : (s \in (A \wedge \neg B)) \rightarrow \text{accept}$ by means of the following ordering of rules:

$$\begin{array}{l} R_1 : (s \in B) \rightarrow \text{deny} \\ R_2 : (s \in A) \rightarrow \text{accept} \end{array}$$

Regarding the access control setup shown in Figure 1(b), together with the set of policy goals and exceptions defined above, i.e., “*All the hosts in (Private \wedge \neg Admin \wedge \neg Interf-fw) are allowed to access web resources on (Internet \wedge \neg Corporate)*”, a possible solution for such a motivation example through a first matching strategy shall be the following set of rules:

$$\begin{array}{l} R_1 : (s \in 111.222.1.0/24 \wedge d \in 111.222.0.0/16 \wedge p = \text{tcp} \wedge \text{dport} = 80) \rightarrow \text{deny} \\ R_2 : (s \in [111.222.1.13, 111.222.1.25] \wedge d \in \text{any} \wedge p = \text{tcp} \wedge \text{dport} = 80) \rightarrow \text{deny} \\ R_3 : (s \in 111.222.1.1 \wedge d \in \text{any} \wedge p = \text{tcp} \wedge \text{dport} = 80) \rightarrow \text{deny} \\ R_4 : (s \in 111.222.1.0/24 \wedge d \in \text{any} \wedge p = \text{tcp} \wedge \text{dport} = 80) \rightarrow \text{accept} \\ R_5 : \text{deny} \end{array}$$

Although this strategy offers a proper solution for the handling of exceptions, it is well known that it may introduce many other configuration errors, such as *shadowing* of rules and *redundancy* [1, 2], as well as important drawbacks when managing rule updates, specially when adding or removing new general rules and/or exceptions. For example, if we consider now the extended access control policy shown in Figure 1(c), together with the insertion of the following general rule: “*In Private, all hosts can access web resources on the zone DMZ*”; and the insertion of the following exception to the previous rule: “*The interfaces of firewall FW₁ (i.e., Interf-fw = {111.222.1.1, 111.222.2.1, 111.222.100.1}) are not allowed to access web resources on the zone DMZ*”; we shall agree that the resulting rules according with these two new premises are the following ones: $R_i : (s \in 111.222.1.1 \wedge d \in 111.222.2.0/24 \wedge p = \text{tcp} \wedge \text{dport} = 80) \rightarrow \text{deny}$; $R_j : (s \in 111.222.1.0/24 \wedge d \in 111.222.2.0/24 \wedge p = \text{tcp} \wedge \text{dport} = 80) \rightarrow \text{accept}$. Such new rules must be inserted in the previous set of rules as shown in Figure 2.

Notice that, in the previous example, the only possible ordering of rules that guarantees the defined assumptions forces us to place the new general rule in the second

$R_1 : (s \in 111.222.1.1 \wedge d \in any \wedge p = tcp \wedge dport = 80) \rightarrow deny$ $R_2 : (s \in 111.222.1.0/24 \wedge d \in 111.222.2.0/24 \wedge p = tcp \wedge dport = 80) \rightarrow accept$ $R_3 : (s \in [111.222.1.13, 111.222.1.25] \wedge d \in any \wedge p = tcp \wedge dport = 80) \rightarrow deny$ $R_4 : (s \in 111.222.1.0/24 \wedge d \in 111.222.0.0/16 \wedge p = tcp \wedge dport = 80) \rightarrow deny$ $R_5 : (s \in 111.222.1.0/24 \wedge d \in any \wedge p = tcp \wedge dport = 80) \rightarrow accept$ $R_6 : deny$
--

Fig. 2. Set of rules for our second motivation example.

position of the set of rules as $R_2 : (s \in 111.222.1.0/24 \wedge d \in 111.222.2.0/24 \wedge p = tcp \wedge dport = 80) \rightarrow accept$. Let us also notice that the related rule to the local exception “The interfaces of firewall FW_1 are not allowed to access web resources on the Internet”, i.e., the former rule $R_3 : (s \in 111.222.1.1 \wedge d \in any \wedge p = tcp \wedge dport = 80) \rightarrow deny$, is now a global exception, and it must be placed in the first position of the set, i.e., it must be placed as $R_1 : (s \in 111.222.1.1 \wedge d \in any \wedge p = tcp \wedge dport = 80) \rightarrow deny$.

As we can observe, the use of this strategy will continuously increase the complexity of the firewall’s configuration as the combination of rules will also do. Furthermore, we can even propose combinations of rules that will not be possible to implement by simply ordering the rules. For instance, let us consider the following two condition attributes A and B , such that $A \cap B \neq \emptyset$, and the following two rules: $R_1 : (s \in (A \wedge \neg B)) \rightarrow accept$; $R_2 : (s \in (B \wedge \neg A)) \rightarrow accept$. As we have seen in this section, the use of a first matching strategy should easily allow us to separately implement these two rules as follows:

$R_{1,1} : (s \in B) \rightarrow deny$ $R_{1,2} : (s \in A) \rightarrow accept$	$R_{2,1} : (s \in A) \rightarrow deny$ $R_{2,2} : (s \in B) \rightarrow accept$
--	--

However, the simple ordering of rules for such an example will not allow us to find out any appropriate combination of rules R_1 and R_2 . Instead, we should first compute $A \cap B$ and then transform the previous rules as follows:

$R_{1,1} : (s \in (A \cap B)) \rightarrow deny$ $R_{1,2} : (s \in A) \rightarrow accept$	$R_{2,1} : (s \in (A \cap B)) \rightarrow deny$ $R_{2,2} : (s \in B) \rightarrow accept$
---	---

and finally deploy the following set of rules:

$R_1 : (s \in (A \cap B)) \rightarrow deny$ $R_2 : (s \in A) \rightarrow accept$ $R_3 : (s \in B) \rightarrow accept$

We can thus conclude that through this strategy the handling of exceptions can lead to very complex configurations and even require additional computations and transformations processes. The administration of the final setup becomes, moreover, an error prone difficult task. Other strategies, like the segmentation of condition attributes or the use of a partial order of rules, will allow us to perform similar managements with better results. We see these other two strategies in the following section.

2.2 Segmentation of Condition Attributes

A second solution when managing exceptions on access control policies is to directly exclude the conditions from the set of rules. In [6, 5], for example, we presented a rewriting mechanism for such a purpose. Through this rewriting mechanism, one may specify the handling of exceptions in the form $R_1 : (s \in (A \wedge \neg B)) \rightarrow \text{accept}$ by simply transforming it into the following rule:

$$R_1 : (s \in (A - B)) \rightarrow \text{accept}$$

The deployment of our motivation example, i.e., “All the hosts in (*Private* \wedge \neg *Admin* \wedge \neg *Interf-fw*) are allowed to access web resources on (*Internet* \wedge \neg *Corporate*)”, through this new strategy, will be managed as follows. We first obtain the set of exclusions, i.e., (*Private* $-$ *Admin* $-$ *Interf-fw*) and (*Internet* $-$ *Corporate*):

<pre> Private = 111.222.1.* Admin = [111.222.1.13, 111.222.1.25] Interf-fw = {111.222.1.1, 111.222.100.1} Private - Admin - Interf-fw \rightarrow [111.222.1.2, 111.222.1.12] \cup [111.222.1.26, 111.222.1.254] Internet = *.*.*.* Corporate = 111.222.*.* Internet - Corporate \rightarrow [0.0.0.1, 111.222.255.254] \cup [111.223.1.1, 255.255.255.254] </pre>
--

Then, we must deploy the following rules:

<pre> R₁ : (s \in [111.222.1.2, 111.222.1.12] \wedge d \in [0.0.0.1, 111.222.255.254] \setminus \wedge p = tcp \wedge dport = 80) \rightarrow accept R₂ : (s \in [111.222.1.26, 111.222.1.255] \wedge d \in [0.0.0.1, 111.222.255.254] \setminus \wedge p = tcp \wedge dport = 80) \rightarrow accept R₃ : (s \in [111.222.1.2, 111.222.1.12] \wedge d \in [111.223.1.1, 255.255.255.254] \setminus \wedge p = tcp \wedge dport = 80) \rightarrow accept R₄ : (s \in [111.222.1.26, 111.222.1.255] \wedge d \in [111.223.1.1, 255.255.255.254] \setminus \wedge p = tcp \wedge dport = 80) \rightarrow accept R₅ : deny </pre>

The main advantage of this approach, apart from offering a solution for the management of exceptions, is that the ordering of rules is no longer relevant. Hence, one can perform a second transformation in a positive or negative manner: positive, when generating only permissions; and negative, when generating only prohibitions. Positive rewriting can be used in a closed policy whereas negative rewriting can be used in case of an open policy. After this second rewriting, the security officer will have a clear view of the accepted traffic (in the case of positive rewriting) or the rejected traffic (in the case of negative rewriting). However, it also presents some drawbacks. First, it may lead to very complex configuration setups

that may even require a post-process of the different segments. Second, it may involve an important increase of the initial number of rules². Nevertheless, such an increase may only degrade the performance of the firewall whether the associated parsing algorithm of the firewall depends on the number of rules. Third, the managing of rule updates through this strategy may also be very complex, since the addition or elimination of new exceptions may require a further segmentation processing of the rules. Some firewall implementations, moreover, are not able to directly manage ranges (e.g., they can require to transform the range $[111.222.1.2, 111.222.1.12]$ into $\{111.222.1.2/31 \cup 111.222.1.4/29 \cup 111.222.1.12/32\}$), and should require the use of third party tools.

2.3 Partial Ordering of Rules

To our knowledge, the most efficient solution to manage the problem of exceptions on access control policies would be by means of a strategy based on partial ordering of rules. Notice that in both first and last matching approaches (cf. Section 2.1), the interpretation of the rules depends on the total order in which the rules are specified, i.e., a total order describes the sequence of rules from a global point of view. However, this ordering of rules can also be implemented in a partial manner, where a set of local sequences of rules are defined for a given specific context.

In the case of NetFilter-based firewalls, for instance, a partial ordering of rules may be achieved through the chain mechanism of IPTables. In this way, we can group sets of rules into different chains, corresponding each one to a given exception. These rules are, moreover, executed in the same order they were included into the chain, i.e., by means of a first match strategy. When a specific traffic matches a rule in the chain, and the decision field of this rule is pointing out to the action *return*, the matching of rules within the given chain stops and the analysis of rules returns to the initial chain. Otherwise, the rest of rules in the chain are considered until a proper match is found. If no rule applies, the default policy of the chain does. Thus, through this new strategy, one may specify the handling of exceptions in the form $R_1 : (s \in (A \wedge \neg B)) \rightarrow \textit{accept}$ as follows:

$$\begin{array}{l} R_1 : (s \in A) \rightarrow \textit{jump_to chain}_A \\ R_2^{\textit{chain}_A} : (s \in B) \rightarrow \textit{return} \\ R_1^{\textit{chain}_A} : \textit{accept} \end{array}$$

Regarding the scenario shown in Figure 2, i.e., “(1) All the hosts in (*Private* \wedge \neg *Admin* \wedge \neg *Interf-fw*) are allowed to access web resources on (*Internet* \wedge \neg *Corporate*); (2) All the hosts in (*Private* \wedge \neg *Interf-fw*) are allowed to access web resources on *DMZ*”, we can now implement such premises via two chains, *private-to-internet* (or *p2i* for short) and *private-to-dmz* (or *p2d* for short), as follows:

² This increase is not always a real drawback since the use of a parsing algorithm independent of the number of rules is the best solution for the deployment of firewall technologies [14].

```

R1 : (s ∈ 111.222.1.0/24 ∧ d ∈ any ∧ p = tcp ∧ dport = 80) → jump_to p2i
R2 : (s ∈ 111.222.1.0/24 ∧ d ∈ 111.222.2.0/24 ∧ p = tcp ∧ dport = 80) → jump_to p2d
R3 : deny

R1p2i : (s ∈ 111.222.1.1) → return
R2p2i : (s ∈ [111.222.1.13, 111.222.1.25]) → return
R3p2i : (d ∈ 111.222.0.0/16) → return
R4p2i : accept

R1p2d : (s ∈ 111.222.1.1) → return
R2p2d : accept

```

Let us now consider the same rules specified in the syntax of NetFilter. The first two rules create a chain called “private-to-internet” (or *p2i* for short) and a chain called “private-to-dmz” (or *p2d* for short). The third rule corresponds to the positive inclusion condition for the first general case (this way, when a given packet will match this rule, the decision is to jump to the chain *p2i* and check the negative exclusion conditions). Similarly, the fourth rule corresponds to the positive inclusion condition for the second general case. We shall observe that in order to deploy this example over a firewall based on Netfilter we should first verify whether its version of IPTables has been patched to properly manage ranges. We must also correctly define in the final IPTables script those variables such as \$PRIVATE, \$DMZ, etc.

```

iptables -N p2i
iptables -N p2d

iptables -A FORWARD -s $PRIVATE -p tcp -dport 80 -j p2i
iptables -A FORWARD -s $PRIVATE -d $DMZ -p tcp -dport 80 -j p2d
iptables -A FORWARD -j DROP

iptables -A p2i -s $INTERF_FIREWALL -j RETURN
iptables -A p2i -s $ADMIN -j RETURN
iptables -A p2i -d $CORPORATE -j RETURN
iptables -A p2i -j ACCEPT

iptables -A p2d -s $INTERF_FIREWALL -j DROP
iptables -A p2d -j ACCEPT

```

The main advantages of this strategy (i.e., partial ordering of rules) are threefold. First, it allows a complete separation between exceptions and general rules; second, the ordering of general rules is no longer relevant; and third, the insertion and elimination of both general rules and exception is very simple. We consider, moreover, that a proper reorganization of rules from a total order strategy to a partial order one may also help us to improve not only the handling of exception, but also the firewall’s performance on high-speed networks [15, 11]. In [15], on the one hand, the authors propose a refinement process of rules which generates a decision-like tree

implemented through the chain mechanism of IPTables. Their approach basically reorganizes the set of configuration rules into an improved setup, in order to obtain a much flatter design, i.e., a new set of configuration rules, where the number of rules not only decreases but also leads to a more efficient packet matching process. In [11], on the other hand, the authors also propose a reorganization of rules in order to better deploy the final configuration. Nevertheless, both authors in [15] and [11] do not seem to address the handling of exceptions, neither expressiveness aspects of their configuration language – that seems to rely upon partial expressiveness languages.

3 Use of Full Expressiveness

Notice that the solutions above overviewed are always based on partial expressiveness, i.e., they implement security policies by means of security rules whose condition attributes are mainly composed of either positive (e.g., A) or negative (e.g., $\neg A$) statements, but they do not allow us to combine both positive and negative statements (e.g., $A \wedge \neg B$) for a single attribute at the same time. Although we have seen in the previous section that these languages may allow us to specify any possible security policy, they can lead to very complex configurations when dealing with the management of exceptions. However, the use of both negative and positive statements for each condition attribute may allow us to specify filtering rules in a more efficient way. The use of a structured SQL-like language [8], for example, will allow us to manage the handling of exceptions in the form $R_1 : (s \in (A \wedge \neg B)) \rightarrow \text{accept}$ through the use of queries like the following ones:

<pre>select decision from firewall where (s ∈ A) ∧ (s ∉ B)</pre>	<pre>select decision from firewall where (s ∈ A) minus select decision from firewall where (s ∈ B)</pre>
--	--

However, these kind of languages are not currently being used for the configuration of firewalls or similar devices – at least not for managing exceptions on access control policies, as defined in this paper. We consider that they will allow security officers to deploy the security policies in a more efficient manner, as well as to properly manage the handling of exceptions on access control policies. Let us for example assume that the configuration language we have been using along the examples of this paper allows us the combination of either positive (e.g., A) and negative (e.g., $\neg A$) statements for each attribute of a single filtering rule. For the sake of simplicity, let us just assume the use of a 2-tuple for specifying both positive and negative values of each attribute (e.g., $R_i : (s \in (A \wedge \neg B)) \rightarrow \text{accept}$ becomes $R_i : (s[+] \in A \wedge s[-] \in B) \rightarrow \text{accept}$). Let us also assume that both positive and negative values are initialized to \emptyset by default. Let us finally assume that we rewrite the matching algorithm implemented in our hypothetical firewall FW_1 into Algorithm 1. In this case, we can easily deploy the first motivation example based on Figure 1(b)'s setup, i.e., “All the hosts in ($Private \wedge \neg Admin \wedge \neg Interf-fw$) are allowed to access web resources on ($Internet \wedge \neg Corporate$)”, as follows:

Algorithm 1: MatchingAlgorithm

```

input : (1) firewall's filtering rules:  $r_1 \dots r_n$ ;
          (2) firewall's default policy: policy;
          (3) packet:  $p$ 
output: decision
1 decision  $\leftarrow$  policy;
2  $H \leftarrow$  GetPacketHeaders ( $p$ );
   /* Let  $r_i = (A_1^i[+] \in V_1^+) \wedge (A_1^i[-] \in V_1^-) \dots (A_p^i[+] \in V_p^+) \wedge (A_p^i[-] \in V_p^-) \rightarrow d_i$ , */
   /* where  $A_{1..p}^i[+]$  and  $A_{1..p}^i[-]$  are, respectively, the set of positive and negative */
   /* attribute conditions of rule  $r_i$ ; and  $V_{1..p}^+$  and  $V_{1..p}^-$  are, respectively, the set */
   /* of positive and negative attribute values of rule  $r_i$ ; */
3 for  $i \leftarrow 1$  to  $n$  do
4   if  $(H_1 \cap V_1^+ \neq \emptyset) \wedge (H_1 \cap V_1^- = \emptyset) \dots (H_p \cap V_p^+ \neq \emptyset) \wedge (H_p \cap V_p^- = \emptyset)$  then
5     decision  $\leftarrow d_i$ ;
6     break; /* Leave the loop */
7 return decision;

```

```

 $R_1 : (s[+] \in 111.222.1.0/24 \wedge s[-] \in \{[111.222.1.13, 111.222.1.25] \setminus$ 
       $\cup 111.222.1.1\} \wedge d[+] \in any \wedge d[-] \in 111.222.0.0/16 \wedge p[+] = tcp \setminus$ 
       $\wedge dport[+] = 80) \rightarrow accept$ 
 $R_2 : deny;$ 

```

Regarding the second motivation example, i.e., “(1) All the hosts in (Private \wedge \neg Admin \wedge \neg Interf-fw) are allowed to access web resources on (Internet \wedge \neg Corporate); (2) All the hosts in (Private \wedge \neg Interf-fw) are allowed to access web resources on the zone DMZ”, we can now properly specify the resulting set of rules as follows:

```

 $R_1 : (s[+] \in 111.222.1.0/24 \wedge s[-] \in \{[111.222.1.13, 111.222.1.25] \setminus$ 
       $\cup 111.222.1.1\} \wedge d[+] \in any \wedge d[-] \in 111.222.3.0/24 \wedge p[+] = tcp \setminus$ 
       $\wedge dport[+] \in 80) \rightarrow accept$ 
 $R_2 : (s[+] \in 111.222.1.0/24 \wedge s[-] \in 111.222.1.1\} \wedge d[+] \in 111.222.2.0/24 \setminus$ 
       $\wedge p[+] = tcp \wedge dport[+] \in 80) \rightarrow accept$ 
 $R_3 : deny;$ 

```

As we can observe, the use of a language based on both positive and negative statements, when specifying the condition attributes of the security rules of a firewall, allows us a more efficient deployment of policies, even using fewer rules. We therefore consider that the little modification we must perform to improve the expressiveness of current firewall configuration languages may allow us to better afford the managing of exceptions on network access control policies. To verify such an assumption, we implemented a proof-of-concept by extending the matching algorithm of IPTables through a Netfilter extension. Due to space limitation, we do not cover in the paper this first proof-of-concept. However, a report regarding its implementation and performance is provided at the following address <http://www.crimplatinum.org/fex/report.pdf>.

4 Related Work

To our knowledge, very little research has been done on the use of full expressiveness languages for the management of firewall configuration as we address in this paper. In [12], for instance, a SQL-like query language for firewalls, called Structured Firewall Query Language is proposed. The authors do not seem to address, however, whether such a language can be used for examining incoming and outgoing traffic, neither to accept nor discard such traffic. The language seems to only be used for the understanding and analysis of firewall's functionality and behavior, rather than be used to perform packet matching or for expressiveness improvement purposes. Similarly, the authors in [13] propose a firewall analysis tool for the management and testing of global firewall policies through a query-like language. However, the expressiveness power of such a language is very limited (just four condition attributes are allowed), and we doubt it may be useful to address our motivation problem.

Some other approaches for the use of formal languages to address the design and creation of firewall rules have been proposed in [4, 7, 3]. However, those approaches aim at specifying and deploying a global security policy through a refinement process that automatically generates the configuration rules of a firewall from a high level language. Thus, the problem of managing exceptions is handled in those works at a high level, rather than a concrete level, and so, the proper configuration once solved the managing issues shall be implemented through one of the strategies already discussed in Section 2. Finally, some proposals for the reorganization of filtering rules have been presented in [15, 11]. However, and as we already pointed out in Section 2, those approaches do not seem to address the handling of exceptions, neither expressiveness aspects of their configuration languages. Their reordering process aim at simply improve the firewall's performance on high speed networks, rather than to offer an easier way to manage the exclusion of condition attributes.

5 Conclusions

In this paper we have studied current strategies in order to manage and deploy policy exceptions when configuring network security components, such as firewalls and filtering routers. As we have discussed, those components are still being configured by security officers in a manual fashion through partial expressiveness based languages. We have also discussed how the use of these languages can lead to very complex configurations when dealing with exclusions of general rules that should always apply. We finally pointed out to the necessity of full expressiveness for combining both negative and positive conditions on firewall languages in order to improve this management of exceptions on access control policies. As we have seen, the simple modification of a general packet matching algorithm can allow us to perform a more efficient deployment of policies by using almost always fewer rules.

As work in progress, we are actually evaluating the implementation of the strategy presented in this paper over NetFilter-based firewalls. For the moment, we have slightly modified its matching process according to the algorithm shown in Section 3,

through the rewriting of a new matching process for IPTables. This first proof-of-concept demonstrates the practicability of our approach. However, we must conduct more experiments to study the real impact on the performance of Netfilter through real scenarios when using our proposal. We plan to address these evaluations and report the results in a forthcoming paper.

References

1. Alfaro, J. G., Cuppens, F., and Cuppens-Boulahia, N. Analysis of Policy Anomalies on Distributed Network Security Setups. In *11th European Symposium On Research In Computer Security (Esorics 2006)*, pp. 496–511, Hamburg, Germany, 2006.
2. Alfaro, J. G., Cuppens, F., and Cuppens-Boulahia, N. Towards Filtering and Alerting Rule Rewriting on Single-Component Policies. In *Intl. Conference on Computer Safety, Reliability, and Security (Safecomp 2006)*, pp. 182–194, Gdansk, Poland, 2006.
3. Alfaro, J. G., Cuppens, F., and Cuppens-Boulahia, N. Aggregating and Deploying Network Access Control Policies. In *Symposium on Frontiers in Availability, Reliability and Security (FARES), 2nd International Conference on Availability, Reliability and Security (ARES 2007)*, Vienna, Austria, 2007.
4. Bartal, Y., Mayer, A., Nissim, K., Wool, A. Firmato: A novel firewall management toolkit ACM Transactions on Computer Systems (TOCS), 22(4):381–420, 2004.
5. Cuppens, F., Cuppens-Boulahia, N., and Alfaro, J. G. Detection and Removal of Firewall Misconfiguration. In *Intl. Conference on Communication, Network and Information Security (CNIS05)*, pp. 154–162, 2005.
6. Cuppens, F., Cuppens-Boulahia, N., and Alfaro, J. G. Misconfiguration Management of Network Security Components. In *7th Intl. Symposium on System and Information Security*, Sao Paulo, Brazil, 2005.
7. Cuppens, F., Cuppens-Boulahia, 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*, pp. 203–218, 2004.
8. Date, C. J. A guide to the SQL standard. Addison-Wesley Longman Publishing Co., Inc. Boston, MA, USA, 1989.
9. Gabillon, A. A formal access control model for XML databases. Lecture notes in computer science, 3674, pp. 86–103, February 2005.
10. Godik, S., Moses, T., and et al. eXtensible Access Control Markup Language (XACML) Version 2. Standard, OASIS. February 2005.
11. Hamed, H. and Al-Shaer, E. On autonomic optimization of firewall policy organization, *Journal of High Speed Networks*, 15(3):209–227, 2006.
12. Liu, A. X., Gouda, M. G., Ma, H. H., and Ngu, A. H. Firewall Queries. In *Proceedings of the 8th International Conference on Principles of Distributed Systems (OPODIS-04)*, pp. 197–212, 2004.
13. Mayer, A., Wool, A., Ziskind, E. Fang: A firewall analysis engine. *Security and Privacy Proceedings*, pp. 177–187, 2000.
14. Paul, O., Laurent, M., and Gombault, S. A full bandwidth ATM Firewall. In *Proceedings of the 6th European Symposium on Research in Computer Security (ESORICS 2000)*, pp. 206–221, 2000.
15. Podey, B., Kessler, T., and Melzer, H.D. Network Packet Filter Design and Performance. *Information Networking*, Lecture notes in computer science, 2662, pp. 803–816, 2003.