



Usage control policy enforcement in SDN-based clouds: a dynamic availability service use case

Khalifa Toumi, Muhammad Sabir Idrees, Fabien Charmet, Reda Yaich,
Gregory Blanc

► To cite this version:

Khalifa Toumi, Muhammad Sabir Idrees, Fabien Charmet, Reda Yaich, Gregory Blanc. Usage control policy enforcement in SDN-based clouds: a dynamic availability service use case. HPCC 2016: 18th IEEE International Conference on High Performance Computing and Communications, Dec 2016, Sydney, Australia. pp.578 - 585, 10.1109/HPCC-SmartCity-DSS.2016.0087 . hal-01449235

HAL Id: hal-01449235

<https://hal.science/hal-01449235>

Submitted on 23 Aug 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Usage Control Policy Enforcement in SDN-based Clouds: A Dynamic Availability Service Use Case

Khalifa Toumi*, Fabien Charmet*, Muhammad Sabir Idrees†, Reda Yaich† and Gregory Blanc*

**Institut Mines-Télécom, Télécom SudParis, Université Paris-Saclay*

9, rue Charles Fourier, 91011 Evry Cedex, France

Email: name.surname@telecom-sudparis.eu

†Institut Mines-Télécom, Télécom Bretagne

2, rue de la Châtaigneraie, 35512 Cesson-Sévigné, France

Email: muhammad.idrees, reda.yaich@telecom-bretagne.eu

Abstract—With the growing interest in Software Defined Networking (SDN) and thanks to the programmability provided by SDN protocols like OpenFlow, network application developers have started implementing solutions to fit corporate needs, like firewalls, load balancers and security services. In this paper, we present a novel solution to answer those needs with usage control policies. We design a policy based management framework offering SDN network security policies. This approach is used to enforce performance requirements (e.g., to ensure a certain level of network connectivity). A top-down approach is proposed, in order to refine the policies into the appropriate network rules, via the OpenFlow protocol. Finally, we implement the solution with an availability service use case and we provide a set of experiments to evaluate its efficiency.

Keywords—Usage Control; SDN; OpenFlow; Network Policies; Service Availability

I. INTRODUCTION

Everything-as-a-Service is emerging as a business transformation approach, far exceeding the usual cloud service provider stack (IaaS, PaaS, SaaS) and attracting many players to the market. This often implies a relaxed security in an environment where the concepts of external threats and internal assets is now blurred [1]. In such clouds, ensuring the availability of this wealth of diverse services may require as much flexibility as it is needed to orchestrate them. In this regard, the Software Defined Network (SDN) paradigm is a realistic contender to enhance the security of a cloud computing environment thanks to its network-wide visibility and programmability features. It introduces an additional abstraction layer between the control plane and data, which are now decoupled. Decision process is now centralized in a single component, namely the SDN controller, which offers a global view of the network to the operators. Thus, they are able to anticipate network resources misuse and plan network bandwidth allocation more efficiently for the hosted services.

Recent research works have already highlighted the promises of SDN by proposing new security services [2]–[5]. In particular, countering massive threats such as distributed denial of service (DDoS) attacks has been extensively studied [4] as SDN key features are expected to minimize the delay from detection to mitigation. Other aspects of the SDN paradigm have also been leveraged: FRESCO [2] provides a modular framework to develop security applications over SDN while SIMPLE [3] builds

upon the decoupling of the control plane from the data plane to ease the process from policy expression to middlebox traffic steering. Chaining services is made possible not only for physically-hosted services, but also to compose network virtual functions as in Slick [6]. [5] proposes a new energy efficient routing algorithm with SDN for data center networks. This approach can schedule flows to the queues of each link to minimize the energy for the given volume of data center traffic.

Although the importance of these solutions, it still has a critical need on a whole framework that provides orchestration and automatic management of network policies [7]. Indeed, we need a flexible and generic mechanism that permits to respond automatically to the detected alerts and problems. Moreover, the solution should provide a high abstract description of the security services to be used with different cloud providers without dependence from their used technology.

However, SDN and the OpenFlow protocol, a de-facto standard for the southbound interface, do not allow to trivially acquire services requirements and translate them into forwarding policies. Indeed, among the previous examples, writing applications that could be directly enforced in the network usually relies on a domain-specific language, slowing its adoption. On the other hand, when high-level languages are provided to express policies, they can only be translated to the extent of a service chain, which is not directly deployable to the forwarding plane. Therefore, there is a need for a policy-based management framework able to collect high-level inputs from operators and provide appropriate reaction scripts that can configure SDN equipments. Additionally, the programmability of the network is rarely exploited to introduce autonomicity: a policy management approach should be context-aware, so that policies are dynamically enforced depending on the network and security states.

In this paper, we propose a policy-based management approach that collects and processes the output of monitoring tools (e.g., incident alerts and network metrics) and provides the necessary OpenFlow rules to reconfigure the data plane. Our framework offers a set of formal security policies and SDN templates which will be used to automate the network configuration and reactions by managing and deploying rules, as well as, detecting potential conflicts. For this purpose, we use a high level formalism, the OrBAC [8] model, which

main benefits are: a) the introduction of a distinctive entity, the *organization*, permitting the representation of global entities in which a policy is applied, e.g., an SDN infrastructure, b) the usability of administration, c) the diversity of access types (permission, prohibition and obligation), and d) the definition of a high abstraction level which simplifies the management of policy rules and the detection of conflicts.

Moreover, our framework includes a new policy instantiation engine to map OrBAC policies and monitoring notifications for context awareness, and an SDN policy deployment engine to translate the inferred rules into OpenFlow configurations. This last feature allows an administrator to write rules for an SDN infrastructure without deep knowledge of the OpenFlow standard. Different security services may be implemented easily with our framework such as availability, attack mitigation, etc. In this paper, we will demonstrate the availability service as a case study.

The rest of the paper is structured as follows. In Section II related work is discussed. In Sections III, IV and V, the proposed framework, architecture, components and workflow, are detailed respectively. In Section IV, the experimentation is presented and discussed. Finally, in section V, the conclusion with some lines of future work are presented.

II. RELATED WORK

While SDN is expected to fulfill many promises with respects to alleviating the load of network operators [9], [10], it often fails to provide a context-aware framework for policy-based network management (PBNM), as it does not fulfill its requirements per se [11], in particular high-level policy definition is not natively supported. Additionally, a recent survey by Kreutz et al. [12] points a few challenges inherent to SDN low-level rule management, among others, dynamic policy enforcement and conflict detection are quoted.

One of the salient features of the SDN paradigm is that it makes programming the network possible, but not necessarily easy. Indeed, many controller platforms require that high-level policies be managed at the individual switch level, in terms of rules installed in each switch [13]. However, the community generally agrees that SDN applications should specify high-level policies in two distinct parts: an endpoint policy, that specifies actions to packets at the controller domain level; a routing policy, that specifies paths taken by the traffic according to traffic engineering considerations [13]. Earlier, Foster et al. [14] had noted that leveraging SDN control languages to program networks according to some policy was not straightforward. Their early state-of-the-art analysis of the OpenFlow controller languages has exposed some substantial difficulties: (1) network functions written in these languages do not compose; (2) low-level hardware-dependent semantics complexify how rules are written and distributed; (3) network programs' two-tiered architecture (spanning both the controller and the switch) incurs concurrency issues. To address these issues, they proposed Frenetic [14], a high-level network

programming language, and in particular, a network policy management library that includes a set of source-level operators for constructing and manipulating streams of network traffic. They later extend their work in the Pyretic language [15] by introducing an abstract packet model, a sequential composition operator and a topology abstraction using network objects. Concurrently, Voellmy et al. have proposed Procera [16], a functional reactive programming language like Frenetic, with which it shares similar constructs. But unlike Frenetic which applies only to packets, Procera handles any event stream allowing for both sequential and parallel composition, and uses function values to represent high-level policy. While these languages were all declarative, Voellmy et al. took another direction when proposing the Maple language [17] which introduces *algorithmic policies*, high-level policies that decide the behaviors of an entire network. Kang et al. have considered an abstraction similar to Maple, i.e., the whole network is *one big switch* [13], but are more concerned than other proposals on how to satisfy user-defined policies and adhere to the underlying hardware resources constraints at the same time. Extending such work into the domain of PBNM, Machado et al. [11] propose a policy authoring framework that is able to capture business-level requirements in the shape of Service-Level Agreements (SLAs) and translate them into low-level system configurations automatically. It leverages network infrastructure information obtained through SDN features to perform logic reasoning on the policy objectives. In contrast, FRESCO offers a development environment that is specialized to serve the needs of security applications [2]. It facilitates the rapid design, and modular composition of OF-enabled detection and mitigation modules. A similar effort has been introduced with *EnforSDN* [18] where middleboxes are described as two distinct layers: a centralized policy resolution layer (usually performed at the controller) and a policy enforcement layer, implemented in each individual switch.

Most of these frameworks usually assume specific knowledge on the user-side, either through the expertise of a programming language or a security tool. The policies define through such frameworks may often blend several sources, which may lead to conflicting rulesets, even more in a context-aware environment. This requires a comprehensive reasoning environment able to handle contexts, and as well as detecting and resolving conflicts. Tripathy et al. [19] have recently addressed such issue through a policy management framework while ensuring the consistency of deployed policies with the centralized network-wide policy. In this paper, we are leveraging a common usage control model language to specify the endpoint policy of the network. As OrBAC is an all-purpose policy model language, it is expected to be more powerful in dealing with contexts, conflicts and usages, which can be wider than the scope of the existing DSLs. Recently, Li et al. [20] have also demonstrated the possibility to translate SLAs into OrBAC rules, potentially extending its usage to cloud service providers and customers exchanges. On the other hand, we have provided a number of generic OpenFlow rule constructs

that allow the translation of OrBAC rules into individual switch configurations. The particular use case we present is a proof that programming the network can be addressed as designing an organization policy. For the moment, the results are similar to legacy mechanisms used to provide resilience to cloud services (replication, load-balancing, etc.). However, leveraging the SDN paradigm for network virtualization allows for increased flexibility in managing resources, which can be costly in replication strategies.

III. SDN POLICY ENFORCEMENT AND MANAGEMENT FRAMEWORK

In this paper, we design a new framework that will offer a high description of network security policies. It will permit to react and re-configure dynamically an SDN network against incidents and attacks. Different components are defined:

- **Context and Monitoring Module - CMM:** A monitoring service must be defined in order to provide reports on the status of VMs, the network and the detected attacks or malicious behaviors. These status will define conditions or alerts that monitor the changes in the system behavior. Several intrusion detection systems and network monitoring tools may be used in a multi-cloud environment. In this paper, we do not focus on the choice or the update of a monitoring solution. However, we will work on the management of the received notifications in order to define contexts that will be used to activate or deactivate the adequate reaction policies.
- **Security Policy Module - SPM:** SPM is composed of two sub components 1), Security Policy Modeling and 2), Security Policy Enforcement Engine. The SPM specifies and models the security requirements to be taken into account within the applications. The security policy to be applied is a set of rules defining the objectives that need to be respected by the system. It includes different rule types: permissions, prohibitions and obligations. Section IV presents how to specify and model the SDN security policy using the Or-BAC model and it gives different elements and components used to build such a policy. Whereas, policy decision engine is the architecture engine that will receive the requests and the notifications. Based on them, it will activate and deactivate the adequate policies and rules. This engine has to manipulate a set of pre-defined network security policies based on the activated context. More details about this module is presented in Section IV-B.
- **SDN Security Module - SSM:** SSM is responsible for the implementation of the current policy in the Policy Enforcement Points (PEPs). PEPs comprise the OpenFlow switches and the SDN controller. The PEP will enforce automatically the selected policy by




Figure 1. Workflow of the policy enforcement approach

the SPM using the OpenFlow scripts. More details about this module could be found in Section V.

Fig. 1 illustrates the workflow of our approach. A monitoring engine continuously supervises the data plane (2) to detect any failures, that may be system- or network-based. This component updates accordingly the virtual machines' status, as well as the network status based on alerts received from local or external sources (1). When resources are critically damaged (failure, error, etc.), the OrBAC-PDP (Policy Decision Point) is notified (3). This policy engine activates contexts based on the status of the networks or virtual machines (4). Appropriate remediation policies are then selected and pushed into the data plane in order to preserve the availability of the targeted services (5, 6, 7). The deployment of these policies is ensured by the SSM module based on predefined SDN templates.

A. Availability service use case

With our framework, several network security policies may be proposed as the availability and attack mitigation services. The first one aims to enhance the application performance against system and network failures. It will offer a technique to distribute traffic between different servers when it is needed. It will apply some techniques as the load balancing and the replication strategies based on the SDN technologies and only when some conditions are detected. Indeed, our solution may adapt the dynamic changes of the cloud (adding and deleting VMs, modification of the allocated resources, migration of a VM, etc). The second service is responsible to react against the detected attacks. It offers the cloud providers to complement their practical defense systems. In the rest of this paper, we will focus more about the availability service that will be used as a case study of our proposal. All the presented examples will be related to this service.

In this section, we present two possible reactions that are managed by our availability service.

- The first one is based on the use of a replica server with degraded Quality of Service (QoS). This scenario presents the case where the provider pledges to preserve the availability of the service without necessarily maintaining the quality of the network. The replica is deployed to receive live copies of user requests, but does not honor them as long as the

main server is running. However, if connectivity to the main server is lost or get degraded, the replica takes over the main server and see its QoS increased to act as the main service:

During the whole lifecycle of the service, two servers are kept alive, one, the main server, with the requested QoS and the second, the backup or replica, with degraded QoS. The role of the availability policy is to manage the performance of the networks and to redirect the flows to the replica with the modification of its QoS in case of a connection problem with the main server. In other words, even if the main server could not provide the required QoS, we provide new paths with the needed performances for the replica to keep the requested QoS.

- The second service proposes a load balancing on demand. At the beginning, only one server provides a service to the customers. Next, based on the network status, a second server is provided in order to achieve the required QoS. This policy provides network resilience for the customer's service by categorizing the network status into three different states. The *NORMAL* state, where the service can answer requests properly and does not suffer from bandwidth nor computation power shortage. The *MIDDLE* state corresponds to a service with several simultaneous connections or less than 50% of its available bandwidth remaining. Finally, the *CRITICAL* state characterizes an offline system. In this scenario, the second machine is solicited depending on the state of the main machine.

IV. SECURITY POLICY MODULE - SPM

A. Modeling Security Policies

Based on the different requirements and challenges highlighted in the Section I, we need to select a security policy model which is capable of modeling various types of security requirements (access control, usage control, and information flow, etc.), provide a way to enforce security requirement dynamically, and most importantly security policy model should be capable of handling conflicts in security policies. In this section, we will present our solution to model SDN network which is adequate for handling most of these requirements.

Traditionally security policies define a set of rules which are intended to control the access to resources and their usage during this access. RBAC [21] is the most well known model defining security policies as a set of permissions. Permissions are rules granting the right to access a specific resource. The approach used by RBAC model to specify security policies is static. Later, the OrBAC model [8] proposed to introduce another type of rules named prohibitions which are denying rights to access specific resources. Obligations are the third type of rules specified to enforce the usage control. They are actions that the system or users are required to take. These actions are essential for the expression and enforcement of a large number of requirements. In SDN network, we consider these three types of security rules. We model these rules using the OrBAC model. Instead of modelling the policy by using the concrete implementation-related

concepts of subject, action and object, the OrBAC model suggests reasoning with the roles that subjects, actions or objects are assigned within an organization. The role of a subject is simply called a role as in the RBAC model. The role of an action is called activity and the role of an object is called view. The security rules specify that some roles are permitted or prohibited or obliged to carry out some activities on some specific views under certain conditions. The main advantage of OrBAC model in the SDN network is that the security rules do not apply statically but their activation may depend on contextual conditions [8]. For this purpose, the concept of context is explicitly included in OrBAC. Many types of contexts have been defined within the OrBAC model in order to take into account different types of conditions. The activation of some contexts activates the rules defined under that specific context.

1) Specification of Access and Usage Control Rules:

Using the concepts introduced by the OrBAC model and based on a first order logic formalism, the specification of the abstract permissions or prohibitions is defined using a 5-places predicate as follows:

Definition 1: an OrBAC security rule is defined as:

$\text{Security_Rule}(\text{Org}, \text{Role}, \text{Activity}, \text{View}, \text{Context})$ where $\text{Security_Rule} \in \{\text{Permission}, \text{Prohibition}\}$.

The obligations differ from permissions/prohibitions by controlling the behavior of the system based on specific events that may occur. For the definition of obligations, we consider two different contexts: the obligation's activation context and the obligation's violation context.

- *Obligation Activation*: An obligation has an activation event after which it becomes effective. This event may be a temporal or an action-based event [22].
- *Obligation Violation*: An obligation has a violation event which specifies when it is violated. This event may be an action-based, temporal or a relative temporal deadline [22].

From its specification until its fulfillment, an obligation can have different status, apart from *Obligation Activation* and *Obligation Violation*:

- *Obligation Deactivation*: An active obligation is deactivated when its context ceases to hold. The deactivation of an obligation depends on the type of the obligation. Some obligations may remain required forever after their activation until they are fulfilled [22].
- *Obligation Fulfillment*: An obligation is fulfilled when its action is taken. Thus, the obligation ceases to be required [22].

Thus, the specification of the obligations can be expressed as follows:

Definition 2: an OrBAC obligation is defined as:

$\text{Obligation}(\text{Org}, \text{Role}, \text{Activity}, \text{View}, \text{Activation_Context}, \text{Violation_Context})$.

Considering our OrBAC security policy model, we have designed a set of entities and components that will permit to define an SDN security policy. In the following

we present some examples of privileges and conditions for our SDN network availability use case (cf. section III):

- P_1 : A subscribed system administrator in Cloud Service provide (CSP) facility is authorized to access various SDN controller services in a default context:
 - P_1 : `Permission(CSP, system-admin, access, sdn-controller-services, default-context)`
- P_2 : The SDN controller is allowed to activate loadBalancing-on-demand which is running on the main server by using the best path in a normal context:
 - P_2 : `Permission(CSP, SDN-Controller, activate, loadBalancing-on-demand, normal-context)`
- O_1 : The SDN controller is obliged to self-activate loadBalancing-on-demand which is running on the main-server, when the network status is partially critical:
 - O_1 : `Obligation(CSP, SDN-Controller, self-activate, loadBalancing-on-demand, partial_critical_network_status, system-timeout)`
- O_2 : The SDN controller is obliged to redirect loadBalancing-on-demand to the second-server, when the network status is critical:
 - O_2 : `Obligation(CSP, SDN-Controller, redirect, loadBalancing-on-demand, critical_network_status, system-timeout)`

Based on the OrBAC derivation mechanism [8], we first define the set of concrete subject [admin-user-1, controller-1], action [access-cmd, activate-cmd, self-activate-cmd, redirect-cmd] and objects [sdn-service-pack, s1]. In the next step, we start by empowering, considering, and using these subjects, actions, and objects into roles, activities and views, respectively. For instance, we empower admin-user-1 subject into system-admin Role, access-cmd action is considered for Access activity, and we use sdn-service-pack object into sdn-controller-services view. Thus, the concrete rules look like as follows:

- P_1' : `Is_Permitted(admin-user-1, access-cmd, sdn-service-pack)`
- P_2' : `Is_Permitted(controller-1, activate-cmd, s1)`
- O_1' : `Is_Obligated(controller-1, self-activate-cmd, s1)`
- O_2' : `Is_Obligated(controller-1, redirect-cmd, s1)`

As mentioned before, the approach that we are considering is dynamic since it takes into account the changes that may happen during the application execution. These changes are modeled in the security policies (as shown

in security rules: P_1, P_2, O_1, O_2) as a specific context. Thus, in order to invoke the permissions or obligation, the activation context should. More specifically for obligations, when *Activation Context* is evaluated as true, the obligation is considered active and it should be fulfilled before the *Violation Context* holds. Whereas, if *Violation Context* holds, then the obligation is considered violated. In section IV-B, we will explain in a detail how the context is managed and evaluated dynamically to enforce security policies.

2) *Conflict Management*: In SDN networks, security rules can be hierarchically structured so that they are inherited in the organization, role, activity and view hierarchies. Since a security policy can be inconsistent because of conflicting security rules (for example a permission can be in conflict with a prohibition). Our solution should be able of managing the conflicts at abstract levels. In our approach, we provide two possible reactions regarding this problem:

- **Separation Constraints**: When assigning a concrete entity to an abstract one, the separation constraints are verified to prevent the user from violating one and potentially generate concrete conflicts. The separation constraints are not exclusively used to prevent concrete conflicts but can be used to specify separation of duty.
- **Rule Priorities**: Define a priority order between the two conflicting rules. We associate the authorization rules with priorities in order to evaluate their significance in conflicting situations. Priorities between access control rules may be sometimes derived from the rules syntactical format.

B. OrBAC-PDP: Security Policy Enforcement Engine

In the previous sections, we introduced our concepts of SDN security policies. In this section, we present how the policy engine module analyzes and interprets these security policies and dynamically enforce them. This module acts as a Policy Decision Point (OrBAC-PDP) with a public SDN API in order to manage security policies. It allows to take a decision (allow/deny) according to specific actions to be applied within the application and also trigger obligations for example in the case of usage control rule enforcement or in the case of any security flaw detection. More specifically, once security policies are loaded in the OrBAC-PDP, the next step is to dynamically control the enforcement of these security policies based on the changed occurring in the environment. Thus, the set of active/inactive policies needs to be synchronized with the set of active/inactive rules which are already enforced by the SDN controller. For this purpose, we have defined an enforcement approach based on the concept of push and pull modes, depending on the type of the security rules that we are considering.

- **Pull Enforcement**: is a type of mode in which SDN controller invokes OrBAC-PDP to retrieve all the active rules related to permissions and prohibitions. The SDN controller calls, using the variables <subject, action, object>, the API method `Is_permitted` in order to evaluate the access as shown in code snippet in the Listing 1.

```

1 resultPermission = policy.is_Permitted(
  _subject, _action, _object);
2 resultProhibition = policy.is_Prohibited(
  _subject, _action, _object);
3 if ((resultPermission == true) && (
  resultProhibition == false)){
4   notificationList.prepareActiveRuleList(
    policy);
5 } else{
6   notifyNegativeEvaluationResult(_subject
    + "is not allowed to perform the
    specific" + _action + "on the" +
    _object);}

```

Listing 1. Security policy pull enforcement

While evaluating the security policies, the OrBAC-PDP consults the context and monitoring module in order to retrieve an actual status of the context as shown in Figure 1. Based on the context status and the elements in the query, the OrBAC-PDP checks its concrete rules (see rules P_1, P_2, O_1, O_2) to decide about the access. In our approach as a response the OrBAC-PDP send a list of active and inactive rules to the SDN controller. Based on this list, as shown in Listing 3, the SDN controller activate or deactivate different network configurations.

- **Push Enforcement:** *is a type of mode in which, instead of SDN controller, the OrBAC-PDP pushes the changes, that occur in the security policy, to the SDN controller in the form of notifications to activate or deactivate different network configurations.* Our core purpose to define push mode is to control obligation enforcement during the application execution. As discussed in the previous section IV-A, these obligations can have different states (Activation, Deactivation, Violation, Fulfillment), during the whole lifecycle of the application execution. In our approach, we enforce these different changes of obligation through considering and evaluating different states of the context. For this purpose, the context and monitoring modules are actively monitoring the changes happening in the environment and regularly notify the status of context to the OrBAC-PDP. Once the OrBAC-PDP has been notified that a context is in a new state, an analysis phase takes place in the OrBAC-PDP to check if there is any obligation that is activate, deactivate, fulfilled or obligation is violated.

```

1 while(_context.getContextState())
2 {
3   resultObligation= policy.Is_Obligated(_subject
  , _object, _action);
4   if (resultobligation == true)
5   {
6     for (Iterator<CConcreteObligation> it =
      policy.GetConcreteObligations().
      iterator(); it.hasNext();) {
7       CConcreteObligation
        _obligation = it.next();
8       if (_obligation.IsActive()) {
9         notificationList.
          prepareActiveRuleList
            (policy);
10        ... } } }

```

Listing 2. Security policy push enforcement

Thus, based on this analysis, the OrBAC-PDP retrieves all the active obligations, see code snippet in Listing 2, and push these active obligations to the SDN controller to dynamically enforce security policies. In our approach, we use the same message format defined in Listing 3 to notify active obligations. The SDN controller receives notification about a new security policy to be deployed, or about changes on the current security policies deployed and generates a new security adaptation plan (a list of actions – activate/deactivate/reconfigure – to be performed in the network) conformed to the security policy.

```

1 <xml version ="1.0" encoding="UTF-8" standalone
  ="Yes"?>
2 <Notification>
3   <RuleStatus>
4     <Rules>
5       <Rule Id="P1", Type ="Permission",
        Value ="active">
6         <Status>"active"</Status>
7         <Subject>"admin-user-1"</Subject>
8         <Action>"access-cmd"</Action>
9         <Object>"sdn-service-pack"</Object>
10        </Rule>
11        <Rule Id="P2", Type ="Permission",
        Value ="inactive">
12          <Status>"active"</Status>
13          <Subject>"controller-1"</Subject>
14          <Action>"activate-cmd"</Action>
15          <Object>"s1"</Object>
16        </Rule>
17        <Rule Id="O1", Type ="obligation",
        Value ="active">
18          <Status>"active"</Status>
19          <Subject>"controller-1"</Subject>
20          <Action>"self-activate-cmd"</Action>
21          <Object>"s1"</Object>
22        </Rule>
23        <Rule> ... </Rule>
24      </Rules>
25    </RuleStatus>
26  </Notification>

```

Listing 3. Security policy: Active/Inactive rules

V. SDN SECURITY MODULE - SSM

This module aims to deploy the selected rules by the SPM into the SDN switches. Indeed, any OrBAC security rule should be translated to a set of three possible actions: adding, deleting and/or modifying a rule from the flow table of an SDN switch. In this module, we have designed three SDN templates for each action. Indeed, each template is a predesigned form of an openflow script related to the three possible actions (adding, deleting and updating). They provide the consistent code that will permit to create the final OpenFlow scripts. The creation of these scripts will be based also on a set of inputs that should be prepared by the SSM module.

With our approach, the deployment of an OrBAC security policy is based on the Algorithm 1. It will take two inputs:

- The security policy that contains the list of rules to apply and to deactivate and
- The network topology that contains the list of switches, servers, the links between them with the different configuration details as (the IP address, ports, etc). It will be used to configure the OpenFlow scripts.

Based on these inputs, our algorithm will compose each rule on the security policy to a set of adding, deleting and/or updating rule. The network topology will permit to enrich these rules with the required information as the switch ID, the path to deploy and the servers IP address. Finally, when all the OrBAC security rules are analyzed, our algorithm will deploy the needed strategy in the controlled switches based on the OpenFlow templates and the rest API of the controller.

Algorithm 1 SSM: SDN policy deployment

Input: Security Policy SP; Network Topology NT

Output: OpenFlow scripts deployed in the SDN switches

```

flow To_add[],To_delete[],To_update[];
  ▷ These arrays will contain respectively the list of
  ▷ OpenFlow rule to add,to delete and to update
for  $R_i$  in SP do
   $R_i$ .extract_Add_Rule(To_add);
   $R_i$ .extract_Delete_Rule(To_delete);
   $R_i$ .extract_update_Rule(To_update);
end for
Concertize(To_add,To_delete,To_update,NT);
Deploy_Policy();

```

VI. IMPLEMENTATION AND EVALUATION

In this section, we evaluate the our framework with a use case on network availability.

A. Implementation

In order to implement the use case, we have set up a testbed (see Figure 2) featuring several hardware servers and OpenFlow (OF) switches:

- 4 Alcatel OS6860 switches, OF-capable with 24 ports running at 1 Gbps.
- Two web servers running Apache v2.4.7 providing a set of video files. (Named main server and second server in the following)
- Two clients streaming video from the server using Tapas [23], an adaptive video streaming client. It simulate a video player without adding the computational overhead of the decoding process.
- An SDN controller managing the switches. In these experiments, we have installed Floodlight¹ on a virtual machine. The OrBAC PDP, the security policies and the SDN security module are deployed in the same virtual machine for eased communication. This machine will be responsible for the reaction strategies and the automatic deployment.

¹Floodlight is an open source controller, available at <http://www.projectfloodlight.org/floodlight/>




Figure 2. Exprimental testbed.

The use case represents a client watching videos over the internet. In our setup, the clients will be virtual machines and video will be played using Tapas. The testbed will serve as the provider's infrastructure. In this paper, we do not cover the network monitoring and notification processes.

B. Evaluation

Network Failure Scenario: In this experiment, we want to prove the feasibility of our solution when several non-simultaneous network failures occur. The situation is presented in the second availability service (load balancing on demand).

Initial situation: During the initial deployment, the context is set as normal-context. Based on this context the OrBAC-PDP will activate a set of rule. Among them, we find the following rule:

- P_2 : `Permission(CSP, SDN-Controller, activate, loadBalancing-on-demand, normal-context)`

These latter will be mapped to openflow scripts that will deploy the best path for the client. As shown in Figure 5, flow rules are pushed toward the different switches during the first period. During this phase, only one client is using the web service. Figure 3 shows an average latency of around $500 \mu s$ between client 1 and the server. As shown in Figure 4, the client's requests increases the packet-per-second rate received by the main server.




Figure 3. Network failure experiment: The latency of the client 1 and 2.

Incident: A partial network failure occurs, which means a network problem is detected in the best path. We simulate this by unplugging a cable between switch 1 and switch 3. We observe the loss of connectivity at the 9th period in Figure 4 (no data received by the Apache server).




Figure 4. Network failure experiment: Number of packets received by the deployed servers.

Response: Once the failure is detected, the context `partial_critical_network_status` is activated.

This context triggers the following new rule:

- O_1 : Obligation(CSP, SDN-Controller, self-activate, loadBalancing-on-demand, partial_critical_network_status, system-timeout)

Let's analyse the reaction of our system after the deployment of this rule: Figure 5 from period 11 to 12 shows the OrBAC-PDP deploying the new path, allowing the traffic to reach server 1 with a lower QoS. In Figure 5, between the periods 9 and 12, we notice that flow rules are pushed on the switches. Once new flow rules are effective, the connectivity is re-established for client 1, with a much higher latency though (> 2 ms). In the meantime, a new client attempts to access the service, but is redirected to the backup server in accordance to the availability policy. The QoS for client 2 is the same as for client 1 in a normal situation. Latency in periods [1-7] for client 1 is

similar as the client 2 latency during periods [16-29] (see Figure 3).




Figure 5. Network failure experiment: Number of packets sent by the controller.

Incident: Now a complete failure occurs in the network, implying that there is no network path leading to server 1. In Figure 3 from the 19th period to the 25th, the server 1 is not reachable anymore. Therefore, the communication between client 1 and server 1 must be redirected toward the second server.

Response: The context is now set as `critical_network_status`. The OrBAC engine activates and deactivates different rules. Let's take the example of the following rule:

- O_2 : Obligation(CSP, SDN-Controller, redirect, loadBalancing-on-demand, critical_network_status, system-timeout)

The OrBAC-PDP will ask the controller to deploy new openflow rules to the switches, thus enabling a new path for client 1 with a new server S2.

We observe, in Figure 4, the second server taking over the first one in processing the requests from client 1. The latency of the two clients are now similar, at around 500 μ s.

VII. CONCLUSION

In this paper, we have designed a flexible and generic framework that will manage and define dynamically the network security policies and push/pull them to the SDN equipment. This solution permits to respond automatically to the detected alerts and problems. It offers also to an administrator or a developer to specify modify rules or to configure a service without a deep knowledge of the openflow standard.

We illustrated our work with a real SDN testbed and ran different experiments with an availability security policy as a case study. The results assure the feasibility and the importance of our solution that may enhance the performance of the application.

Future work aims at completing the implementation with the appropriate monitoring tools to assist network administrators in tasks such as alert data extraction and verification of the VM status. We also plan to evaluate our solution in more complex network scenarios, including different network operators.

ACKNOWLEDGMENT

This research has been supported in part by the European Union's Horizon 2020 Research and Innovation

Programme under grant agreement No. 643964 (SUPER-CLOUD), and by the French national Fund for Digital Society (FSN). The opinions expressed in this paper are those of the authors and do not necessarily reflect the views of the European Commission, or the French Government.

REFERENCES

- [1] J. Cropper, J. Ullrich, P. Frühwirth, and E. Weippl, "The role and security of firewalls in IaaS cloud computing," in *Proceedings - 10th International Conference on Availability, Reliability and Security, ARES 2015*, 2015, pp. 70–79.
- [2] S. Shin, P. Porras, V. Yegneswaran, M. Fong, and G. Gu, "FRESCO: Modular Composable Security Services for Software-Defined Networks," in *Network and Distributed System Security Symposium (NDSS 2013)*, 2013, pp. 1–16. [Online]. Available: <http://www.csl.sri.com/users/vinod/papers/fresco.pdf>{\%}5Cnpapers3://publication/uuid/5A7E2F2B-FBAC-48CF-BDCE-0C4A1A4604FB
- [3] Z. A. Qazi, C.-C. Tu, L. Chiang, R. Miao, V. Sekar, and M. Yu, "SIMPLE-fying Middlebox Policy Enforcement Using SDN," *SIGCOMM Comput. Commun. Rev.*, vol. 43, no. 4, pp. 27–38, 2013. [Online]. Available: <http://doi.acm.org/10.1145/2534169.2486022>
- [4] Q. Yan, F. R. Yu, Q. Gong, and J. Li, "Software-defined networking (SDN) and distributed denial of service (DDoS) attacks in cloud computing environments: A survey, some research issues, and challenges," pp. 602–622, 2016.
- [5] G. Xu, B. Dai, B. Huang, and J. Yang, "Bandwidth-Aware Energy Efficient Routing with SDN in Data Center Networks," in *High Performance Computing and Communications (HPCC), 2015 IEEE 7th International Symposium on Cyberspace Safety and Security (CSS), 2015 IEEE 12th International Conferen on Embedded Software and Systems (ICSS), 2015 IEEE 17th International Conference on*. IEEE, 2015, pp. 766–771.
- [6] B. Anwer, T. Benson, N. Feamster, D. Levin, and J. Rexford, "A Slick Control Plane for Network Middleboxes," in *Proceedings of the Second ACM SIGCOMM Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN '13. New York, NY, USA: ACM, 2013, pp. 147–148. [Online]. Available: <http://doi.acm.org/10.1145/2491185.2491223>
- [7] S. Richard, "Effective Network Orchestration Starts by Automating Provisioning," Gartner, Tech. Rep., 2015. [Online]. Available: <https://www.gartner.com/doc/3122020/effective-network-orchestration-starts-automating>
- [8] A. A. El Kalam, R. El Baida, P. Balbiani, S. Benferhat, F. Cuppens, Y. Deswarte, A. Mieke, C. Saurel, and G. Trouessin, "Organization Based Access Control," in *Proceedings of the 4th IEEE International Workshop on Policies for Distributed Systems and Networks POLICY 2003*, 2003, pp. 120–131. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1206966>
- [9] H. Kim and N. Feamster, "Improving network management with software defined networking," *IEEE Communications Magazine*, vol. 51, no. 2, pp. 114–119, 2013.
- [10] P. Sharma, S. Banerjee, S. Tandel, R. Aguiar, R. Amorim, and D. Pinheiro, "Enhancing network management frameworks with SDN-like control," ... *Management (IM ...*, pp. 688–691, 2013. [Online]. Available: http://ieeexplore.ieee.org/ielx7/6560458/6572961/06573054.pdf?tp=\\&arnumber=6573054\\&isnumber=6572961\\&5Cnhttp://ieeexplore.ieee.org/xpls/abs{_}all.jsp?arnumber=6573054
- [11] C. C. Machado, J. A. Wickboldt, L. Z. Granville, and A. Schaeffer-Filho, "Policy authoring for software-defined networking management," in *Proceedings of the 2015 IFIP/IEEE International Symposium on Integrated Network Management, IM 2015*, 2015, pp. 216–224.
- [12] D. Kreutz, F. M. V. Ramos, P. E. Verissimo, C. E. Rothenberg, S. Azodolmolky, and S. Uhlig, "Software-defined networking: A comprehensive survey," *Proceedings of the IEEE*, vol. 103, no. 1, pp. 14–76, 2015.
- [13] N. Kang, Z. Liu, J. Rexford, and D. Walker, "Optimizing the "One Big Switch" Abstraction in Software-Defined Networks," *Context'13*, p. 17, 2013. [Online]. Available: <http://www.cs.princeton.edu/{~}jrex/papers/rule-place13.pdf>
- [14] N. Foster, R. Harrison, M. J. Freedman, C. Monsanto, J. Rexford, A. Story, and D. Walker, "Frenetic: a network programming language," *Proceeding of the 16th ACM SIGPLAN international conference on Functional programming - ICFP '11*, p. 279, 2011. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2034812\\&5Cnhttp://dl.acm.org/citation.cfm?doid=2034773.2034812>
- [15] C. Monsanto, J. Reich, N. Foster, J. Rexford, and D. Walker, "Composing software-defined networks," *Proceedings of the 10th USENIX conference on Networked Systems Design and Implementation*, pp. 1–14, 2013. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2482626.2482629{\\%}5Cnhttp://www.frenetic-lang.org/pyretic/>
- [16] A. Voellmy, H. Kim, and N. Feamster, "Procera: a language for high-level reactive network control," in *Proceedings of the first workshop on Hot topics in software defined networks*, 2012, pp. 43–48.
- [17] A. Voellmy, J. Wang, Y. R. Yang, B. Ford, and P. Hudak, "Maple: simplifying SDN programming using algorithmic policies," *Proceedings of the ACM SIGCOMM 2013 conference on SIGCOMM*, pp. 87–98, 2013. [Online]. Available: <http://doi.acm.org/10.1145/2486001.2486030>
- [18] Y. Ben-Itzhak, K. Barabash, R. Cohen, A. Levin, and E. Raichstein, "EnforSDN: Network policies enforcement with SDN," in *Proceedings of the 2015 IFIP/IEEE International Symposium on Integrated Network Management, IM 2015*, 2015, pp. 80–88.
- [19] B. K. Tripathy, A. G. Sethy, P. Bera, and M. A. Rahman, "A Novel Secure and Efficient Policy Management Framework for Software Defined Network," in *Computer Software and Applications Conference (COMPSAC), 2016 IEEE 40th Annual*, vol. 2. IEEE, 2016, pp. 423–430.
- [20] Y. Li, N. Cuppens-Boulahia, J. M. Crom, F. Cuppens, and V. Frey, "Expression and enforcement of security policy for virtual resource allocation in IaaS cloud," in *IFIP Advances in Information and Communication Technology*, vol. 471, 2016, pp. 105–118.

- [21] R. S. Sandhu, E. J. Coyne, H. L. Feinstein, and C. E. Youman, "Role-Based Access Control Models," pp. 38–47, 1996. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs/_all.jsp?arnumber=485845
- [22] Y. Elrakaiby, F. Cuppens, and N. Cuppens-Boulahia, "Formal enforcement and management of obligation policies," *Data and Knowledge Engineering*, vol. 71, no. 1, pp. 127–147, 2012.
- [23] L. De Cicco, V. Caldaralo, V. Palmisano, and S. Mascolo, "TAPAS: a Tool for rApid Prototyping of Adaptive Streaming algorithms," in *Proceedings of the 2014 Workshop on Design, Quality and Deployment of Adaptive Video Streaming*. ACM, 2014, pp. 1–6.