

Pushing forward Security in Network Slicing by leveraging Continuous Usage Control

Barbara Martini¹, Paolo Mori², Francesco Marino³, Andrea Saracino², Alessio Lunardelli², Antonio La Marra², Fabio Martinelli², Piero Castoldi³

Abstract—Fifth generation (5G) softwarized network systems will allow to flexibly partition the network infrastructure into logically-independent network slices, hosting end-to-end Network Services able to dynamically meet the diverse requirements of vertical industries. However, the high-dynamicity of NFV-related operations and the interdependence of multiple slices running on top of a shared underlying infrastructure pose peculiar security challenges. In this paper we investigate how such challenges can be addressed in the context of the MANagement and Orchestration (MANO) security functions within the ETSI NFV Architectural framework. In particular, we target access control and authorization functions, and we discuss how advancing them for network slicing deployments with continuous and closed-loop Usage Control (UCON) mechanisms. We also present a Proof of Concept of a MANO framework extended with UCON capabilities able to regulate the access and the use of network slices according to customizable security policies. Preliminary performance evaluation proved the effectiveness of the proposed approach with minor impact on the user experience and prompt reaction time to security policy violations.

Index Terms—network slicing, security orchestration, access control, usage control, Open Source MANO.

I. INTRODUCTION

Software-Defined Networking (SDN) and Network Function Virtualization (NFV) are enabling a new way to deploy networks as softwarized infrastructures, i.e., *network slices*, that are provisioned as end-to-end Network Services (NSs) composed of Virtualized Network Functions (VNFs) to provide customized network, compute and storage capabilities and flexibly address the requirements of emerging 5G use cases and vertical applications, e.g., e-health, cloud robotics. The adoption of some degree of automation and orchestration is key for convergent slice resource management and service delivery operations, as addressed by the MANagement and Orchestration (MANO) functions within the ETSI NFV Architectural framework. However, despite the well-recognized advantages, the wide adoption of orchestrated NFV/SDN deployments and network slicing could be prevented by security threats and vulnerabilities that may pose additional challenges not emerged before [1].

Firstly, peculiar security challenges may arise from the high-dynamicity of NFV-related settings and operations [1]. Indeed, frequent migrations of virtual machines running Virtualized Network Functions (VNFs), dynamic formation of

virtual network zones, and on-demand orchestration eventually enlarge the attack surface of network services. Secondly, NFV products that are not yet mature make the landscape of security offences and risks even more wide and unclear. Finally, having the underlying resources shared between slices and services may further generate serious security threats. For instance, a compromised VNF could access other victim VNFs through shared resources, leading to the disruption of the entire cloud infrastructure and services.

The above challenges can be even more exacerbated if we consider the highly composite and complex NFV ecosystem and scenarios featured by:

- multiple layers of the NFV architectures including the (i) NFV Infrastructure (NFVI), i.e., the pool of physical/virtual computing, storage and network capabilities; (ii) VNFs, i.e., the pool of network functions running in virtual machines deployed in the NFVI; (iii) NFV MANO, i.e., orchestration and lifecycle management functions of physical/virtual resources, VNFs, and NSs. Moreover, the interworking is also relevant with the Operating/Business Support System (OSS/BSS) which provides service delivery and business operations at the level of network slices [2];
- different types of VNFs (e.g., virtual routers/switches, intrusion detection systems, firewall) interworking at different levels (i.e., data, control, application plane) with diverse service components (e.g., orchestrators, service delivery controllers, vertical applications) to address a composite scenario of network slicing use cases, ranging from the provisioning of virtual network assets (e.g., virtual Radio Access/Core Networks) to virtual service infrastructures (e.g., virtual Content Delivery Networks) or a combination of thereof [1];
- diversity of stakeholders involved in 5G use cases (i.e., network operators, verticals, service providers, end users) also including third parties (e.g., VNF developers). In high-dynamic service scenarios enabled by 5G use cases, ensuring trust between software components (either VNFs or service components) and consistent security policies between verticals and providers points of view is still a challenge [3].

In this paper we elaborate how to enhance security in network slicing and orchestration following an end-to-end approach, as required by the complexity and dynamicity of NFV deployments and stakeholder ecosystem. Among the different security areas, we target in particular the access

³ TeCIP Institute, Scuola Superiore Sant’Anna, Pisa, Italy <name.surname>@santannapisa.it

²IIT-CNR, Pisa, Italy <name.surname>@iit.cnr.it

¹ PTN Lab, CNIT, Pisa, Italy <name.surname>@cnit.it

control and the authorization area which is considered one of the most relevant and critical due to the high dynamicity of 5G service scenarios [1]. Indeed, the highly-changing context of users, resources and services enabled by 5G use cases pose specific access control challenges that cannot be addressed by traditional authorization systems. In 5G scenarios, the access context featured by a set of security-relevant attributes (e.g., security state of a VNF, user reputation) could in fact evolve in such a way that the policy which initially granted the access rights might not be satisfied anymore at a later time as a consequence of actions performed by the user (e.g., inexperienced user setting up HTTP ports instead of HTTPS ports) or due to external events (e.g., a VNF gets infected by a virus). Hence, it is advisable that attributes that could change value over time are continually re-assessed throughout the period of time the accessed resources are used, to detect improper or malicious actions or situations.

In this direction, we propose to enhance network slicing and orchestration security by integrating continuous and closed-loop access control mechanisms, in order to address the previously described security challenges. We leverage the Usage Control (UCON) mechanisms to regulate the access and the use of network slices throughout their entire lifecycle on the basis of security-relevant mutable attributes that are continuously assessed against usage control policies to promptly react and prevent any improper or malicious usage of NFV resources and network slices from users (e.g., Verticals, VNFs, end users).

II. SECURING THE NETWORK SLICING AND ORCHESTRATION

NFV MANO functions provide the orchestration and lifecycle management of physical/virtual resources and VNFs to automatically establish NSs and to monitor their integrity and performance during the entire lifecycle to assure the proper operation of network slices.

The security of NFV resources and network slices that are managed and orchestrated by MANO platforms can be addressed through the automation of security management functions, in order to allow network operators to keep up with the extreme dynamicity and complexity of NFV deployments [4]. In this direction, security functions are being more and more implemented and deployed as Virtualized Security Functions (VSFs), like virtual Identity and Access Management Systems [5], virtual Firewalls [6], and virtual Network Intrusion Detection Systems (NIDS) [7]. Moreover, the unprecedented network flexibility and the use of software-based solutions allow to benefit of great opportunities to deliver the relevant level of security by exploiting dynamic deployment, orchestration and composition of VSFs [3]. For instance, the possibility to monitor and orchestrate VNFs and NSs allows to timely detect anomalous behaviours and promptly take adequate countermeasures, e.g., (re)configuration to increase component isolation, late-chaining of virtual firewall in the slice.

Within ETSI, an initiative is undergoing to extend the NFV MANO architecture with a security lifecycle management

module, namely NFV Security Manager (NSM), supposed to interwork with the MANO functions and in charge of reacting to the internal and external events which might prevent a secure network services operation in NFV environments [8]. Leveraging this component, the network security administrators are not required anymore to interact with the network management platform every time an intervention is needed. On the contrary, they can encapsulate their knowledge in properly designed security policies, provide them to the NSM and let the network management platform take care of their enforcement.

Following [8], [9] identifies the extensions to the NFV MANO architecture and corresponding reference points related to security management and monitoring operations, while [10] specifies the interfaces supported over such reference points as well as the information elements exchanged over these interfaces.

The work described in [11] can be considered a preliminary implementation of the NSM for a security use case related to the access control to VNFs. Using a security-as-a-service approach, an access control component for a cloud infrastructure is launched by the NSM aiming at examining user authentication and permission against policy rules. The access control is performed at the time the NFV resources are accessed in the cloud, before initiating network slicing or allocating virtual resources. Hence, no continuous control is performed while the VNFs are used after network slice set-up to prevent improper or malicious usage of VNFs or slices.

Among the MANO platforms, the Open Network Automation Platform (ONAP) offers some interesting components toward supporting NSM functions. ONAP features generalized closed-loop process enabling real-time reactions to user-specified actionable events. Recently, ONAP has been extended with a specific support for access control based on XACML (eXtensible Access Control Markup Language). However, at the time of writing the usage control is not natively supported.

From the above consideration, we conclude that the NSM scope of operation should be broadened in terms of both involved entities and time span, in order to address an end-to-end security vision in highly-changing scenarios of users, resources and services enabled by 5G hyper-connected scenarios and use cases. In this paper, how to advance security support in network slicing in such a direction is elaborated. In particular, in line with [8], [9], [10], we discuss security challenges for network slices considered in end-to-end view as virtual infrastructure and service assets for verticals to run 5G use cases. This approach extends the security discussion for NFV network services in literature [1], where security is mainly addressed at each layer of NFV environments without considering the slice service level. To address the envisioned challenges, we present a flexible security solution that, although in principle also applicable to e2e security management for any network service, we demonstrate to be specifically effective to address also some pending security problems in NFV environments as stated by ETSI and as elaborated in the use case developed in Section V [12].

III. ADVANCING END-TO-END SECURITY IN NETWORK SLICING

To give an end-to-end vision to securing network slices, in this paper we identify three different security contexts, depicted in Figure 1, which an evolved NSM should be aware of to be able to ensure a consistent slice security in 5G virtual infrastructures. In the following, the terms vertical and tenant are used interchangeably.

The three security contexts are classified based on their operational scope in which the NSM acts specified by (i) the semantics of information considered to take security decisions (e.g., service-related information), (ii) the entities affected by countermeasures (e.g., slices, VNFs or the service as a whole), and (iii) involved stakeholder(s) (i.e., the network operator or the tenant). More specifically:

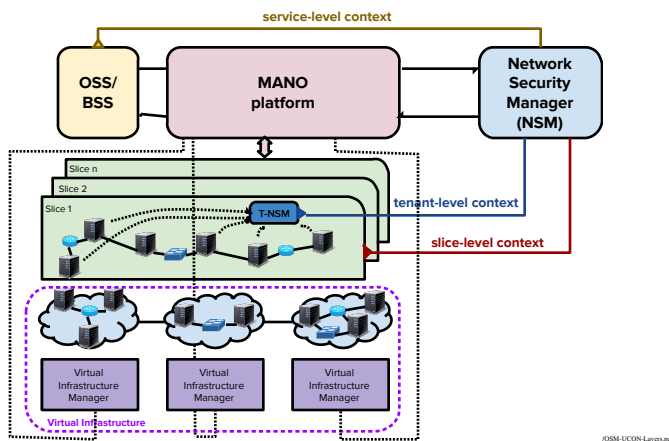


Fig. 1. NSM security contexts in NFV deployments

- **service-level context**, encompassing everything that regards slices meant as virtual infrastructure service assets delivered to Verticals but which can nonetheless provide the network operator with useful information to keep the (virtual) infrastructure service secure as a whole. The NSM probes to retrieve security-relevant information are placed within the network operator OSS/BSS, thereby collecting attributes of slice services delivered to tenants, e.g., number of slices per Vertical, Vertical reputation rank. At this level, the NSM is also able to bind slices with tenants, so that everything happens at the slices can be brought back to the behaviour of the Vertical (e.g., SLA/SLS violations or slice misuse from tenants, tenant reputation degrades), and the NSM can enforce the related countermeasures, e.g., the slice suspension, not only in the originating slice but also in all the slices owned by that tenant. An example of security threat as result of service misuse could be that slices may become points of attack. To mitigate this threat, the NSM could grant slices to trusted Verticals only (i.e., with high reputation) or could impose a limit on the maximum number of slices granted to a Vertical.
- **slice-level context**, encompassing everything concerning slices meant as instances of a virtual infrastructure underpinned by an end-to-end Network Services, i.e., a

composite set of running interconnected VNFs, that the network operator can directly control and measure. At this level, the NSM probes to retrieve security-relevant information are managed by the MANO platform and placed in the cloud and network resource domains to collect VNF monitoring data. For example, anomalous network traffic patterns may reveal infected VNFs which can be promptly isolated or suspended. Similarly, the opening up HTTP ports at VNFs instead of HTTPS ones may expose the slice to attacks. The countermeasures are taken at the level of constituent VNFs (e.g., activating a VSF like a packet inspector, isolating the suspicious VNFs) thereby recovering a secure set-up or a situation under control for the network operator. At this level, countermeasures are also meant to protect tenants whose slices have been compromised and to ensure the continuity of the corresponding network services. Countermeasures could be also propagated, if necessary, to other slices through the service context, e.g., suspending the slice instances with VNFs similar to the infected one.

- **tenant-level context**, encompassing everything that regards the software components (i.e., tenant applications) running into the slices delivered to tenants. Since this context is within the tenant domain, the tenant is the only responsible for the management and orchestration of software components (i.e., network appliances, application service control and orchestration functions, application software modules). Hence, at this level the tenant is the main actor, which may have stricter security policies than the network operator. In this case, the tenant enforces its own security policies through a tenant-managed security component (T-NSM) running into the slice [8].

In this paper we identify the Usage Control framework as an effective reference for a NSM featuring continuous and closed-loop access control functionalities in the MANO framework, at all the above security contexts.

IV. USAGE CONTROL: REFERENCE MODEL AND IMPLEMENTATION

Several access control models have been defined in the literature, e.g., Role-based Access Control (RBAC) and Attribute Based Access Control (ABAC). The RBAC model makes the access control decisions based on the user role, while the ABAC model advances RBAC in considering a wider set of attributes in the policy to describe security-relevant features of the environment, users, and resources, e.g., user reputation or VNF security state (clean/infected). However, both the above models assess the policy at the time of access request only. This means that, as the access is granted and the user starts using the resource, the security policy is not re-evaluated anymore. Since the access context could change, this may cause failures to detect security violations.

To extend the temporal span in which security policies are assessed, the Usage Control (UCON) model can be adopted. As shown in Figure 2, beyond the traditional *access control phase* with the policy (called *pre-policy* in the UCON model) evaluated and enforced at the time the access is requested,

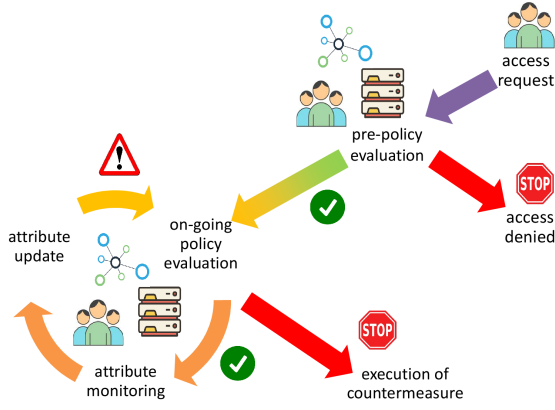


Fig. 2. Usage Control workflow

the UCON model considers additional policies (i.e., *ongoing-policies*) that are continuously evaluated while the resources are in use (*usage control phase*), thereby enabling the prompt detection of improper or malicious usage of resources from users. In the event of an *ongoing-policy* violation, proper countermeasures (specified in the *ongoing-policy*) are enforced aiming at avoiding or minimizing the duration of the improper or malicious use and, hence, mitigating their negative effects.

The Usage Control Systems (UCS) is a software platform implementing a UCON-based authorization system as an extension of the XACML standard (a well-known standard specifying a language to express ABAC policies, an architecture, and a workflow for their enforcement). The UCS is highly configurable in terms of policy specifications, attribute sources and countermeasures specification [13].

V. USAGE CONTROL IN THE MANO FRAMEWORK: USE CASE AND WORKFLOW

In line with ETSI specifications, the UCS provides NSM functions specified in [8] for securing network slices in the area of Access and Usage Control. In particular, the UCS implements policy-based security monitoring and management functions in Fully-Active mode of operation using the the *Sc-Or* reference point to interwork with the MANO platform [9]. Moreover, the UCS is compliant with the interfaces specified for the *Sc-Or* reference point to collect status information (e.g., monitoring data from the virtualized infrastructure), to use lifecycle management reports (e.g., NS set-up request), and to enforce security policies (e.g., NS or VNF termination) [10].

In relation to the service-level security context, it is worth pointing out that [9] [10] do not specify any reference point between the OSS/BSS and the NSM (i.e., UCS) since they assume the MANO platform acts as a proxy in the event the NSM needs to collect information from the OSS/BSS. The foundation behind this position is that ETSI assumes a tight coupling of the slice lifecycle operations to NSs [2]. A closer look to [10] shows that this proxy functionality cannot be actually leveraged to retrieve OSS/BSS-specific data needed by the UCS. Moreover, it is not specified in [10] if the same approach is to be applied to the security notifications and policy enforcement commands. These aspects can be regarded as shortcomings and a direct interaction between the OSS/BSS

and the NSM should be actually supported in real-world deployments. Indeed, there is a consensus on the need for a kind of network slicing-aware NFV orchestration where truly slice management and orchestration functionalities extended at the slice service level (i.e., managed at the OSS/BSS) are considered relevant to effectively address vertical requirements in 5G deployments, e.g., arbitration of virtual resources across slices to address concurrent SLAs [14][2]. We think that a similar service-level scope should be also considered in support of slice security management and monitoring functions with the OSS/BSS coming into play to prevent security threats in cooperation with the NSM (i.e., UCS). However, at this time a direct interaction between the OSS/BSS and the UCS is challenged by the lack of a unique standard on the OSS/BSS side to expose service/business-level data.

As explanatory use case, suppose that a Vertical (e.g., Video Content Provider) asks a Network Service Provider (NSP) to set-up a slice service in order to run a business (e.g., live video delivery for sport events). In the network slicing orchestration context, the application of a slice service request operation is typically mediated by the NSP OSS/BSS to which the Vertical firstly needs to register. At the registration time, the Vertical is assigned with a role (e.g., *gold*, *bronze*), e.g., based on the negotiated service level. The NSP OSS/BSS is used later by the Vertical to apply the slice service request specifying the, e.g., list of VNFs (e.g., virtual streaming servers, caches, and transcoders), VNF connectivity properties¹, list of logical network appliances (e.g., firewall, Network Intrusion Detection System), level of management and control capabilities (e.g., *basic* with only VNF monitoring capabilities, *advanced* with VNF monitoring and configuration capabilities). In fact, the Vertical may want to take care of configuration tasks on VNFs (e.g., VNF functional configuration, software upgrades and connectivity configurations) to enforce desired VNFs settings according to its service requirements (e.g., set-up and reconfiguration of service chains, configuring video streaming sessions).

If the Vertical uses the slice or configure the VNFs in an improper or insecure way, this could result in security threats (e.g., loss of VNFs integrity and control, loss of slice availability, traffic modification) which could damage, besides the Vertical itself, the NSP as well. Hence, the NSP may want to enforce policies to impose best practices or to prevent security threats [12]. For instance, in the *access control phase*, the NSP may want to (i) prevent that untrusted Verticals (i.e., with low reputation) are granted with slice services, (ii) limit the number of slices granted to a Vertical, (iii) impose that a firewall is specified by the Vertical as constituent logical network function of the slice, (iv) regulate the level of VNF control and management capabilities assigned to the Vertical based on its role. Accordingly, an example of *pre-policy* is the following:

Pre-policy: A Vertical V can be granted a slice *if* (the number of slices already running for that Vertical is less than S) *AND* (the Vertical reputation rank is

¹The capacity requirements have been omitted because not relevant for this use case

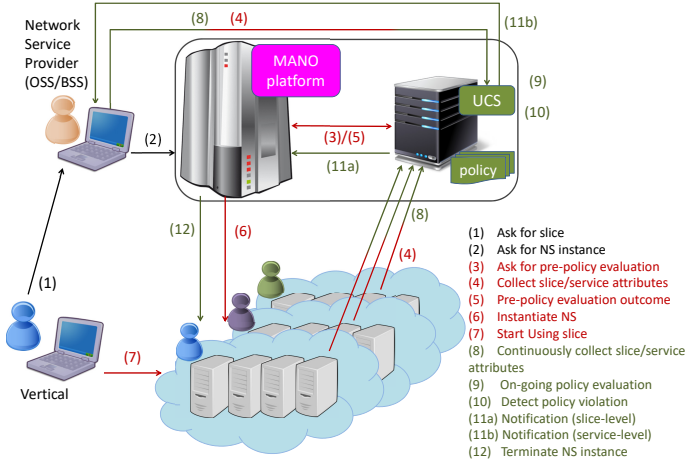


Fig. 3. Workflow of Usage Control in the NFV MANO framework

above a given threshold R) AND (the slice includes a VNF implementing a firewall) AND [(the Vertical role is equal to “bronze” AND the requested level of management and control capabilities is *basic*) OR (the Vertical role is equal to “gold” AND the requested level of management and control capabilities is “advanced”)].

During the *usage control phase*, the NSP may want to (i) prevent that untrusted Verticals (i.e., with degraded reputation) keep using the slice, (ii) impose that the Vertical uses encrypted communication channels to/among VNFs featuring at least a given level of encryption strength in order to have communications protected at all levels, (iii) impose limitations on the kind of ports that VNFs expose to the Internet, e.g., HTTPS only, to reduce the attack surface, (iv) maintain the integrity of the VNF software in order to avoid intrusions in the slice and, hence, set-up a periodic check of the software integrity status of VNFs in order to detect any kind of corruption or presence of a virus, and (v) promptly detect intrusions, attempts of traffic or data modifications or attacks with signaling flooding and replaying through a NIDS embedded at slice level. Accordingly, an example of *ongoing-policy* is the following:

Ongoing-policy: A Vertical V can run a slice as long as (the Vertical reputation rank is above a given threshold R) AND (the encryption level on all the slice connections is higher or equal to EL) AND (the integrity level of the VNFs belonging to the slice is higher than a threshold IL) AND (the time elapsed from the last integrity check on the slice components is less than AT) AND (a network intrusion attempt has not been observed) AND (the VNFs in the slice expose services on HTTPS ports only).

Figure 3 shows the workflow to provide Usage Control in the MANO framework. Steps 1-2 describe the normal workflow in current MANO platforms. In step 1, the Vertical asks the Network Service Provider (mediated by an OSS/BSS) to set-up a slice and in step 2 the Network Service Provider submits the request for NS instantiation to the MANO plat-

form.

Steps 3-5 describe the extension to the workflow in order to include the UCS for the *access control phase*. Following the slice set-up request, the UCS evaluates (a set of) attributes against the *pre-policy* to make the access control decision. The values of attributes are collected at step 4 from both the virtual infrastructure (i.e., slice-level attributes directly or mediated by the MANO platform) and the OSS/BSS (i.e., service-level attributes). If the UCS decides the Vertical is entitled to get the new slice, the NS supporting the required slice is instantiated and properly configured by the MANO platform and then the Vertical starts using it (steps 6-7).

Steps 8-12 describe the further extension to the workflow involving UCS for the *usage control phase*. To this purpose, as soon as the Vertical starts using the slice, a background process begins in order to continuously collect (a set of) attributes. Again, the values of attributes are collected from both the virtual infrastructure (i.e., slice-level attributes) and the OSS/BSS (i.e., service-level attributes). As soon as any of those attribute values undergoes to any change while the slice is running, the UCS evaluates the updated attributes against the *ongoing-policy* to make a decision on a possible on-going security threat. If the UCS observes a policy violation, it notifies the MANO platform (step 11a) in case the security violation is at the slice-level (e.g., an intrusion has been detected) or the OSS/BSS (step 11b) in case of policy violation at the service-level (e.g., the reputation of the Vertical goes under the threshold). In turn, the notified entity enforces the appropriate countermeasure. If the MANO platform is notified, then it may enforce the termination of the NS instance and of all virtualized resources supporting the slice (step 12). If the OSS/BSS is notified, then it may take countermeasures such as downgrading the level of control and management capabilities assigned to the Vertical.

VI. PROTOTYPE VALIDATION AND PRELIMINARY RESULTS

To validate the proposed approach, we implemented a prototype integrating the UCS with OSM release SEVEN, as shown in Figure 4. In the prototype, some slice management-related OSS/BSS functions are provided by OSM, while the others we needed have been implemented in ad-hoc software module [15].

Firstly, we instrumented the Information Model - North-Bound Interface (IM-NBI) and the OSS/BSS components with a Policy Enforcement Point (PEP) to intercept security-relevant operations from the Northbound APIs (e.g., slice set-up requests) and to trigger the UCS to perform the *pre-policy* and *ongoing-policy* evaluations. In case the *ongoing-policy* is violated, the two PEPs handle (asynchronous) messages from the UCS to enforce the countermeasures specified in the policy. For instance, to terminate a slice, the PEP in the IM-NBI invokes the slice termination interface on the NorthBound API of the IM-NBI.

Secondly, we developed a set of plugins, called Policy Information Points (PIPs), which are embedded in the UCS for collecting updated attribute values from the technology-specific software components in charge of handling such

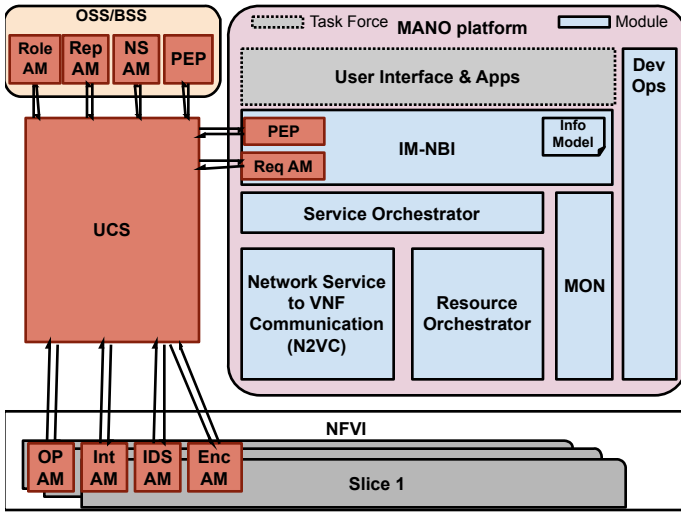


Fig. 4. Integration of UCS in the OSM platform: software architecture

attributes, called Attribute Managers (AMs). More specifically, we implemented the PIPs to interact with the following AMs: *Rep AM* (reputation of the tenant), *Role AM* (role of the tenant), *NS AM* (number of slice running for the tenant), *Req AM* (level of management and control capabilities requested by the tenant and presence of a firewall VNF in the slice request), *OP AM* (list of open ports on the VNFs of the slice), *Int AM* (integrity level of the running VNFs, and time elapsed from the last integrity check), *IDS AM* (network intrusions), *Enc AM* (presence of a non-encrypted connections to/among VNFs).

To validate the approach, we performed some experiments with OSM running on a Linux server equipped with 4 Intel Xeon gold 6140m CPU @ 2.295Ghz and 8GB of RAM memory, and the UCS running on another Linux server equipped with one Intel Xeon W3565 CPU @ 3.193 Ghz and 8GB of RAM memory. We configured OSM to use OpenStack for the provision of the virtual resources to set-up slices. We used a slice descriptor composed of two Ubuntu-based VNFs connected through a virtual link, and we referred to the policy examples in Section V. As for AMs deployment, typically standard tools are used which feature highly different delays in generating attribute values. On the other hand, from UCS design point of view, we are interested in evaluating the impact of the overhead introduced by the UCS by itself. Hence, we emulated the AM presence with custom web services in order to decouple the delay performance of the UCS from the one of AMs. This is equivalent as to assume AMs computing attribute values independently on the UCS operation, and providing the latest computed value upon a request from the UCS.

Firstly, we performed an experiment (1000 repetitions) to compute the delay introduced by the UCS in the overall slice creation process. We find out that the average time required by the UCS to evaluate the *pre-policy* (including communications with the PEP) is 893.7 ms (Standard Deviation: $\rho = 7.22$), while the average time required to set-up the slice is 48.72s. Consequently, the overhead introduced by the UCS is about 1.8% of the slice set-up time, thus it can be definitely con-

sidered not significant. Considering that higher set-up time is expected in a more realistic slice scenario, such an impact can be even lower.

Secondly, we performed another experiment (1000 repetitions) to compute the time required by the UCS to detect an *ongoing-policy* violation and trigger a countermeasure. To this purpose, we artificially communicated to the UCS that the reputation of the tenant was degraded to cause an *ongoing-policy* violation. Starting from this time we measured the time elapsed before the PEP receives the violation notification from the UCS. We measured, on average, 1315.2 ms (Standard Deviation: $\rho = 13.55$), which states a prompt reaction from the UCS.

VII. CONCLUSION

We investigated the adoption of the UCON model to provide advanced secure slice orchestration functionalities in the MANO framework. We validated our proposal through a Proof-of-Concept implementation which demonstrated minor impact on the user experience and prompt reaction time to security policy violations. It is worth pointing out that the choice to use Virtual Machine-based network functions in the Proof-of-Concept does not affect the generality of the proposed solution because the operation of UCS is independent on the technology used for the virtualization of network functions. Hence, our approach can be similarly adopted in Container-based 5G deployments.

As future work, we plan a more comprehensive evaluation of the proposed approach in real-world environments addressing realistic slice and vertical scenarios such as Content Delivery Network and Enhanced Packet Core. Also Containers will be used as an alternative virtualization technology to deploy network functions. Moreover, we plan to assess the impact of the overhead due to continuous collection of attribute values, e.g., in terms of higher load and delay introduced by AMs. Finally, we plan to promote the Proof-of-Concept to the OSM community toward a possible integration of the prototype in some future OSM releases. As long-term plan we envision to further extend the proposed security support for slice orchestration to consider the past behaviour of the tenant to decide the actions it is allowed to perform in the future (i.e., history-based access control).

REFERENCES

- [1] M. Pattaranantakul *et al.*, “NFV security survey: From use case driven threat analysis to state-of-the-art countermeasures,” *IEEE Communications Surveys Tutorials*, vol. 20, no. 4, pp. 3330–3368, Fourthquarter 2018.
- [2] H. Khalili *et al.*, “Network slicing-aware nfv orchestration for 5g service platforms,” in *2019 European Conference on Networks and Communications (EuCNC)*, 2019, pp. 25–30.
- [3] E. Dotaro, “5G network slicing and security,” <https://sdn.ieee.org/newsletter/january-2018/5g-network-slicing-and-security>, accessed on 28 May 2020.
- [4] G. Gardikis *et al.*, “Shield: A novel nfv-based cybersecurity framework,” in *2017 IEEE Conference on Network Softwarization (NetSoft)*, July 2017, pp. 1–6.
- [5] Z. Zhang *et al.*, *Identity and Access Management in NFV*, 2017, pp. 135–155.
- [6] J. Deng *et al.*, “On the safety and efficiency of virtual firewall elasticity control,” in *2017 Network and Distributed System Security Symposium (NDSS)*, 2017.

- [7] H. Li *et al.*, “vNIDS: Towards elastic security with safe and efficient virtualization of network intrusion detection systems,” in *2018 ACM Conference on Computer and Communications Security (CCS)*, 2018, pp. 17–34.
- [8] ETSI Industry Specification Group, “Network Functions Virtualisation (NFV) Release 3; Security; Security Management and Monitoring specification,” *ETSI GS NFV-SEC 013 V3.1.1*, 2017.
- [9] —, “Network Functions Virtualisation (NFV) Release 3; Management and Orchestration; Architecture enhancement for Security Management Specification,” *ETSI GS NFV-IFA 026 V3.2.1*, 2019.
- [10] —, “Network Functions Virtualisation (NFV) Release 3; Management and Orchestration; Sc-Or, Sc-Vnfm, Sc-Vi reference point - Interface and Information Model Specification,” *ETSI GS NFV-IFA 033 V0.9.0*, 2020.
- [11] M. Pattaranantakul *et al.*, “Secmano: Towards network functions virtualization (nfv) based security management and orchestration,” in *2016 IEEE Trustcom/BigDataSE/ISPA*, 2016, pp. 598–605.
- [12] ETSI Industry Specification Group, “Network Functions Virtualisation (NFV); NFV Security; Security and Trust Guidance,” *ETSI GS NFV-SEC 003 V1.1.1*, 2014.
- [13] P. Mori *et al.*, “Usage control on cloud systems,” *Future Generation Comp. Syst.*, vol. 63, pp. 37–55, 2016.
- [14] F. Malandrino *et al.*, “Service shifting: A paradigm for service resilience in 5g,” *IEEE Communications Magazine*, vol. 57, no. 9, pp. 120–125, 2019.
- [15] A. L. Marra *et al.*, “Demonstration of secure slicing using etsi mano enhanced with usage control capability,” in *2019 IEEE Conference on Network Softwarization (NetSoft)*, 2019, pp. 254–256.

Antonio La Marra is CEO at Security-Forge SME. Previously, he was Research Assistant at IIT CNR, working on usage control and mobile/IoT security.

Barbara Martini is Head of Research at CNIT, Italy. Her research interests include service orchestration in SDN/NFV/5G environments, network control/management architectures, network security.

Piero Castoldi is Full Professor at Scuola Superiore Sant’Anna, Pisa, Italy. His research interests cover reliability, switching and control for optical networks, application-network cooperation and cloud networking.

Alessio Lunardelli is Research Assistant at IIT CNR, Italy, working on data privacy and security, focusing on access/usage control.

Francesco Marino is Research Engineer at École Polytechnique Fédérale de Lausanne. His research interests include security and privacy in softwarized 5G networks, IoT, smart environments.

Fabio Martinelli is Senior Researcher at IIT CNR, Italy, leading the Cyber Security Project. His research interests include security and privacy in distributed and mobile systems, and foundations of security and trust.

Paolo Mori is Researcher at IIT CNR, Italy. His research interests include trust, security and privacy in distributed systems, focusing on access/usage control, and Blockchain technology.

Andrea Saracino is Researcher at IIT CNR, Italy. His research interests include usage control, IoT, distributed system security and mobile malware analysis.