Transport Network Slices with Security Service Level Agreements

P.Alemany*, D.Ayed†, R.Vilalta*, R.Muñoz*, P.Bisson†, R. Casellas*, R.Martínez*
*Centre Tecnològic de Telecomunicacions de Catalunya (CTTC/CERCA), Castelldefels, Spain,
Email:palemany@cttc.es
†Thales SIX GTS (TSG), France, Email: dhouha.ayed@thalesgroup.com

Abstract

This paper presents an initial architecture to manage End-to-End Network Slices which, once deployed, are associated with Security Service Level Agreement(s) to increase the security on the virtual deployed resources and create End-to-End Secure Network Slices. Moreover, the workflows regarding the Network Slicing provisioning and the whole SSLA Lifecycle management is detailed.

I. Introduction

Network Slicing is becoming a 5G networks backbone as it is being investigated and discussed by both research and industry actors. Multiple articles have presented challenges at different layers [1] or more specifically on wireless networks [2]. Since then, multiple investigations have shown results related to some of the challenges: multiple virtualisation technologies used in a Network Slice (referred as Slice) [3], Slices deployments in different network domains (edge/cloud) [4], security aspects like using isolation to make reliable Slices [5] or how Slices help to mitigate Distributed Denial of Service (DDoS) attacks [6].

A commonly investigated topic is Service Level Agreement (SLA). Typically its use is focused on ensuring a proper Quality of Service (QoS) based on the network resources -e.g. bandwidth, latency, etc.- assigned to a service. Yet, there is another way to keep the desired QoS: the service security. Instead of requesting network resources to increase the bandwidth, it might be possible to maintain the QoS by defining a set of parameters related to information security instead of the information traffic. A Security SLA (SSLA) looks towards to define a set of requirements in order to ensure the safeness of a service -e.g. information integrity, encryption, etc.- in front of a possible problem affecting the information and accessing the service.

The main goal of the paper is to define and develop a framework that allows a Slice provider -i.e. the owner of the Slice descriptors with the verticals high-level requirements- who acts as a broker relying on several Service Providers (SPs) providing various network services to deliver Slices controlled by Security SLAs to the verticals/end-users. Each provided Slice has to be covered by a Security SLA that specifies the security grants offered.

The article is organized as follows: section II introduces the SSLA model used, sections III shows the architecture's framework and its internal modules, section IV describes a use case and the Slice provisioning and monitoring workflows using the selected use case and, finally, section V presents the future work and conclusions.

II. Security Service Level Agreement

The Security Manager (Mngr) takes care of the SSLA lifecycle in a Slice: a) it gathers the verticals/end-users security requirements; b) deploys the necessary security controls to enforce the agreed SSLA by enriching or configuring the SPs' services; c) monitors in real-time the SSLAs fulfillment d) detects violations in security provisioning level based on an analytic engine and notifies both end-users and SPs; e) reacts in real-time to adapt the provided level of security or to apply proper countermeasures.

To automate the SSLA lifecycle in a Slice, an SSLAs machine readable format is adopted based on the SPECS [7] SSLA model extended to support slicing. This model is based on a WS-Agreement XML schema that is extended with security-related information allowing to specify the following sections in a Slice term description:

- Slice resource providers: The available resource providers (appliances, networks, etc.) infrastructure.
- Slice required security capabilities: A capability is a set of security controls, the NIST's Control Framework [8] is used to specify these security controls.
- Security Metrics: They are referenced in the Slice service properties and used to define Security Service Level Objectives (SLOs) in the guarantee terms section. A metric specification includes information about it and also how to process the SLOs, such as the metric name, definition, its measurement scale and the expression used to compute its value.

III. Security SLA Manager for Securing Network Slicing

Fig. 1 presents the architecture designed to create and manage secured Slices. This architecture is divided in two main blocks: the Slice Mngr (Service Orchestrator) and the Security Mngr.

The Slice Mngr is in charge to orchestrate the Slice lifecycle. Its architecture keeps the structure presented in [9]:

1) Slice LCM: It manages the Slices lifecycle through the Network Slice Templates (NSTs) and Network Slice Instances (NSIs) data objects.

2) Slice2NS mapper: It maps any Slice-related action from an NSI object to the individual NS components and backwards -i.e. gathering individual NS information into the Slice-.
3) NST and NSI DBs: These DBs contain the NSTs -i.e. the data schema defining the components within a with the Slice- and the NSIs -i.e. the records with the deployed virtual resources information-.

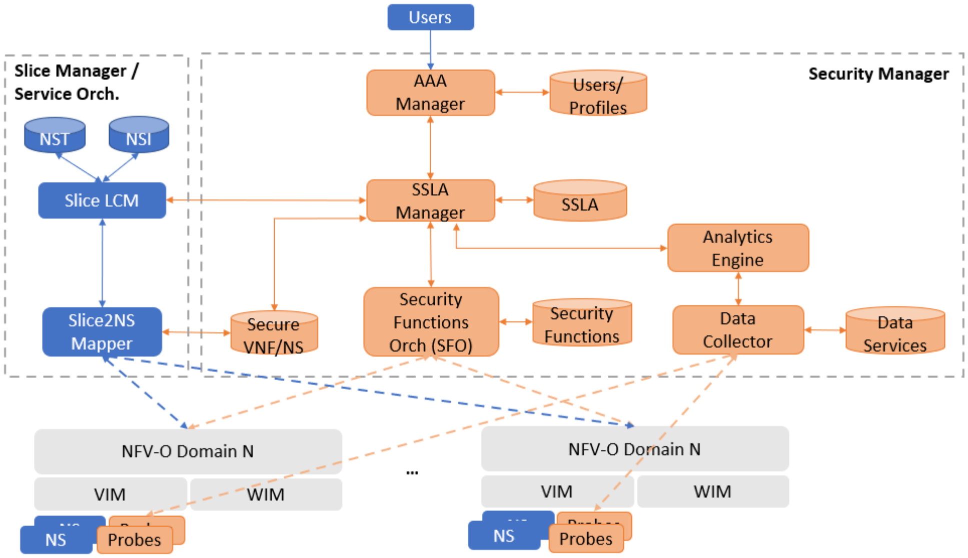

Fig. 1: Secured Network Slicing Architecture

The second enabler and the main focus of this paper is the Security Mngr. Its main objective is the management of SSLAs associated to Slices and if an SSLA violation occurs, to inform the Slice Mngr about the proper actions to solve it. Its internal components are:

1) AAA Mngr: Module in charge of allowing the access into the Secured Slice component and its internal features. Its main functionality is to control the users access and reject possible non authorised access attacks.
2) SSLA Mngr: It manages the lifecycle of the SSLA assigned to a Slice. Its functionalities are: a) to get the SSLA data object from the SSLA Data Base (DB) and pass it to the Slice LCM, b) to command the Security Functions Orchestrator (SFO) to launch the Security Functions (SFs) -i.e. probes- and so, to apply the Security Service Function Chaining (SSFC) in order to start monitoring and, c) the management of the policy action to solve a SSLA violation.
3) SFO: Module in charge to manage the SFs lifecycle for each Slice instance and monitor its security performance to warn the SSLA Mngr about a SSLA violation.
4) Analytics Engine (AE): Evaluates the data forwarded by the Data Collector and determines whether a SSLA associated to a Network Service has been violated or not based on parameters defined by the SSLA Mngr.
5) Data Collector (DC) : It gathers and organises all the information coming from all the SFs in order to forward it to the AE.
6) Users/Profiles DB: It keeps a list of who can access the framework.
7) SSLA DB: It saves the generated SSLA data objects with the trio: Security Controls/Metrics/SLO.
8) Security Functions DB: Contains the set of SFs to be launched when a Slice is instantiated and get the related data to monitor and ensure the SSLA is fulfilled. Together with probes, another possible SF could be the NIST security controls to address different security purposes (e.g. privacy).
9) Data Services DB: It keeps the incoming data/events collected by the DC module, in order to keep track and log the performance.
10) Secure VNF/NS DB: This last DB, has the list of Validated/Certified VNFS and NSs by a trust authority that a Slice can use. If a Slice needs a VNF/NS which is not in that DB, it will not be deployed as the internal element is not considered a secured element.

IV. Use Case Description

In this section a simplistic use case is presented in order to describe how the internal modules within the previously described architecture interact among them during the service provisioning and monitoring phases.

A. Scenario

Fig. 2 shows the use case selected to describe the architecture workflows involving the previously presented modules.

The use case is composed by three elements: a Slice composed by a single Firewall Network Service (FW) and attached to the Slice, two SFs: a Traffic Generator (TG) and a Traffic Sink (TS). The TG generates traffic that is sent to the FW, which must filter which traffic to forward or not to the TS using its internal rules. While the data reaches the TS, this sends up the IP addresses which were transmitted so the AE can validate whether the FW has done its functionality or not.

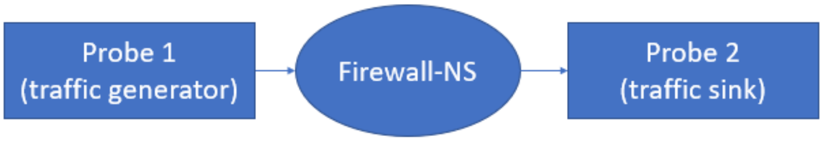

Fig. 2: Firewall use case architecture

B. Secure Network Slicing Provisioning

The deployment of a Slice with the SSLA associated and its related SFs is presented in Figs. 3 and 4a, for a better understanding and image understanding, this process is divided in three parts:

The first part is divided in three steps: the user access and authorisation to deploy a Slice, the SSLA selection and the Slice deployment request as Fig. 3a shows. Everything starts with the user accessing to request a Slice instance (1). The access is verified (2) by the AAA Mngr and the Users Profiles DB and, if allowed, the request reaches the SSLA Mngr (3). Then, the Slice associated SSLA descriptor is selected from the SSLA DB (4) and forwarded to the Slice LCM (5). This gets the reference NST (6) to create the NSI (7) descriptor and save it (8).

The second part focuses on the Slice deployment, so as presented in Fig.3b, it follows from where the first part finished with the NSI ready to deploy: the Slice LCM sends the NSI to the Slice2NS Mapper (9) to verify that all NS/VNFs composing the NSI are in the Sec VNF/NS DB and considered secured to be deployed with respect to properties that need to be guaranteed (10). If one of the elements is not in the DB (10a), the Slice LCM warns the SSLA Mngr (10b) and this to the user (10c) so the Slice is not deployed. Otherwise, each NS is requested to be deployed by the NFVO (11). When all are deployed, the Slice LCM is warned (12), the NSI descriptor is updated (13) and the SSLA Mngr is informed (14).

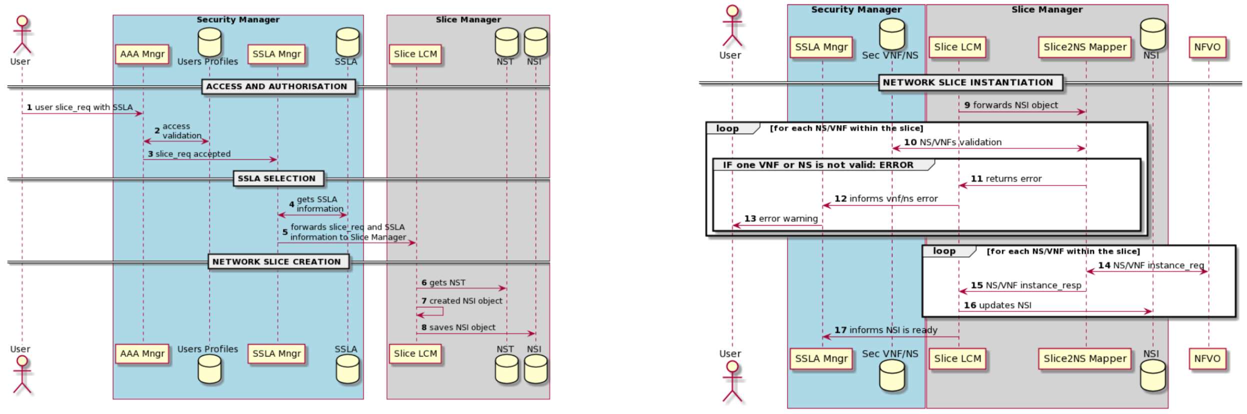

(a) Access, Authorisation, SSLA selection and Slice creation

(b) Network Slice Instantiation step

Fig. 3: Network Slice Instantiation workflow (parts 1 and 2)

The third part, as Fig. 4a shows, has three steps: the monitoring parameters selection and the SFs deployment to get statistics from the Slice and the Slice instantiation completeness. So, the SSLA Mngr sends to the AE the required information to be aware and recognise an SLA violation (15). This requests to the DC (16/17) which parameters must gather from the SFs (yet to be deployed) and confirms the correct configuration to the SSLA Mngr (18). Then, the SSLA Mngr requests (19) to deploy a set of SFs to the SFO (which are kept in the Security Functions DB (20)) and then it requests (21) to the NFVO to deploy them with the NSI (22). When all the SFs are ready, the SFO warns the SSLA Mngr (23) and this informs the user that the Slice is ready to use and be monitored (24).

C. Secure Network Slicing Monitoring

Fig. 4b presents the monitoring workflow once the Slice is deployed and monitored to ensure the SSLA. It starts with each Slice associated SF sending the monitored data to the DC (1). This gathers the data based on the IP addresses and sends the different data sets to the AE (2). Then, this module can validate if all IP addresses are allowed to pass or not (3). If no SSLA is violated (B), the AE saves the logs and awaits for new incoming information (4b). If an SSLA is violated (A), the AE reports to the SSLA Mngr the non allowed IP addresses that passed through the firewall (4) and the SSLA Mngr looks for

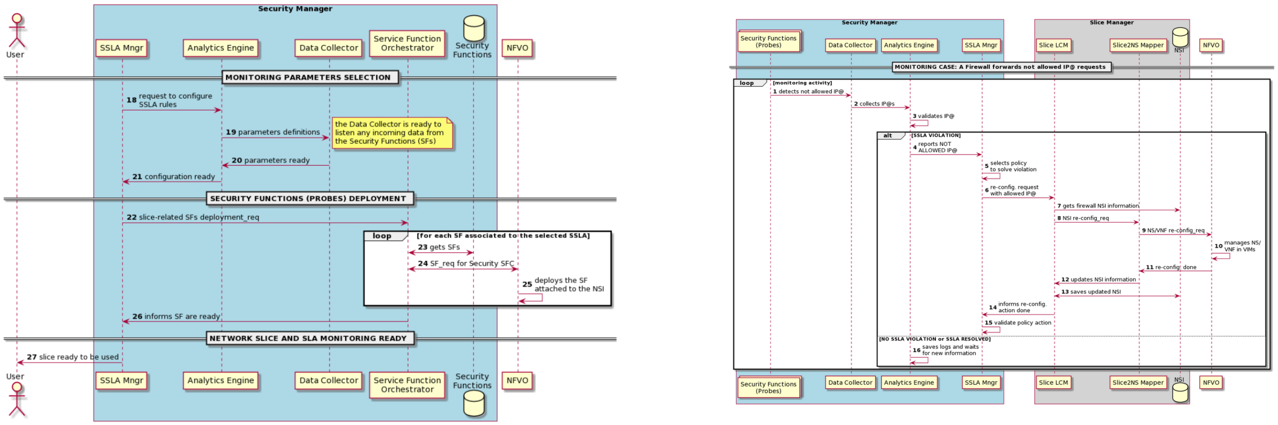
Security Manager
User
SSLA Mngr
Analytics Engine
Data Collector
Service Function Orchestrator
Security Functions
NFVO
MONITORING PARAMETERS SELECTION
18 request to configure SSLA rules
19 parameters definitions
the Data Collector is ready to listen any incoming data from the Security Functions (SFs)
20 parameters ready
21 configuration ready
SECURITY FUNCTIONS (PROBES) DEPLOYMENT
22 slice-related SFs deployment_req
loop
[for each SF associated to the selected SSLA]
23 gets SFs
24 SF_req for Security SFC
25 deploys the SF attached to the NSI
26 informs SF are ready
NETWORK SLICE AND SLA MONITORING READY
27 slice ready to be used
Security Manager
Slice Manager
Security Functions (Probes)
Data Collector
Analytics Engine
SSLA Mngr
Slice LCM
Slice2NS Mapper
NSI
NFVO
MONITORING CASE: A Firewall forwards not allowed IP@ requests
loop
[monitoring activity]
1 detects not allowed IP@
2 collects IP@s
3 validates IP@
alt
[SSLA VIOLATION]
4 reports NOT ALLOWED IP@
5 selects policy to solve violation
6 re-config. request with allowed IP@
7 gets firewall NSI information
8 NSI re-config_req
9 NS/VNF re-config_req
10 manages NS/ VNF in VIMs
11 re-config. done
12 updates NSI information
13 saves updated NSI
14 informs re-config. action done
15 validate policy action
[NO SSLA VIOLATION or SSLA RESOLVED]
16 saves logs and waits for new information

(a) Monitoring parameters, SF deployment and instantiation complete

(b) Slice monitoring workflow

Fig. 4: Network Slice Instantiation (part 3) and Monitoring workflow

the policy to be apply and solve the violation (5). For the current example, a policy to be applied could be to re-configure the firewall Slice. So, the SSLA Mngr requests to the Slice LCM to re-configure the Slice passing and updated IP addresses filter (6). The Slice LCM gets the NSI descriptor (7) and informs the Slice2N mapper about which is the specific NS inside the Slice to re-configure (8). Finally, the NFVO is requested to re-configure (9) the virtual resource -e.g. a Virtual Machine or Container-. Once the re-configuration is done (10), all the modules answer back to their requester: the NFVO to the Slice2NS Mapper (11), this to the Slice LCM (12) which will update the database (13) with the new NSI data. Finally, the SSLA Mngr is informed (14) and awaits for a new SSLA violation warning from the AE (15).

V. Future Work and Conclusions

This paper has presented an architecture with a Security Manager and a Network Slice Manager working together to join E2E Slices with Security SLAs that are necessary to deploy E2E Secured Slices according to vertical security requirements. Once deployed, the slices are monitored to fulfill the expected security requirements. Furthermore, the architecture also presented the ideas of secured Network Services and Virtual Network Functions to compose the E2E Secured Slice and the Security Functions to gather the information to monitor any Security SLA violation.

The architecture presented is the first version for a Network Slice Manager with an enforced view on the Slices security. Even though the design seems to follow the correct path, some modifications are still possible and, in fact, one of the tasks to carry out in the future is to validate and ensure that there are no weak points on the processes presented. In addition to the architecture, the workflow of a simple use case is presented to understand how the framework works. Once the architecture will be more consolidated, an experimental proof of the concept will be developed in order to demonstrate the described functionalities and if necessary, to refine the architecture in order to simplify it (if possible), make it more cognitive and not limited to transport Network Services but also others, while addressing interoperability and scalability issues we may anticipate.

Acknowledgment

Work partially funded by the EC 5GPPP INSPIRE-5GPlus (871808) and MINECO AURORAS (RTI2018-099178-B-I00).

References

[1] X. Foukas, G. Patounas, A. Elmokashfi, and M. K. Marina. Network slicing in 5g: Survey and challenges. *IEEE Communications Magazine*, 55(5):94–100, 2017.

[2] M. Richart, J. Baliosian, J. Serrat, and J. Gorricho. Resource slicing in virtual wireless networks: A survey. *IEEE Transactions on Network and Service Management*, 13(3):462–476, 2016.

[3] Pol Alemany, Ricard Vilalta, Felipe Vicens, Ignacio Dominguez, Ramon Casellas, Ricardo Martínez, Sonia Castro, Josep Martrat, and Raul Muñoz. Hybrid network slicing: Composing network slices based on vnfs, cnfs network services. In *2020 IEEE Conference on Network Softwarization (NetSoft) 29 June – 3 July 2020, Ghent (Belgium)*. IEEE, 2020.

[4] Pol Alemany, Juan Luis de la Cruz, Ricard Vilalta, Ramon Casellas, Ricardo Martínez, Raul Muñoz, Ana Pol, and Antǿn Roman. Comparison of real-time communications service kpi in edge and cloud domains through multi-domain transport networks. In *24th International Conference on Optical Network Design and Modelling (ONDM), 18-21 May 2020, Castelldefels (Barcelona, Spain)*, 2020.

[5] A. Marotta, D. Cassioli, M. Tornatore, Y. Hirota, Y. Awaji, and B. Mukherjee. Reliable slicing with isolation in optical metro-aggregation networks. In *Optical Fiber Communication Conference (OFC) 2020*, page M4D.3. Optical Society of America, 2020.

[6] Danish Sattar and Ashraf Matrawy. Towards secure slicing: Using slice isolation to mitigate ddos attacks on 5g core network slices. *CoRR*, abs/1901.01443, 2019.

[7] V. Casola, A. D. Benedictis, M. Rak, and U. Villano. Sla-based secure cloud application development: The specs framework. In *2015 17th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing (SYNASC)*, pages 337–344, 2015.

[8] National Institute of Standards and Technology (NIST). Nist special publication 800-53 revision 4: Security and privacy controls for federal information systems and organizations. Technical report, NIST, 2013.

[9] R. Vilalta, P. Alemany, R. Casellas, R. Martínez, C. Parada, J. Bonnet, F. Vicens, and R. Muñoz. Zero-touch network slicing through multi-domain transport networks. In *2018 20th International Conference on Transparent Optical Networks (ICTON)*, pages 1–4, 2018.