Micro-Segmenting 5G

Olli Mämmelä, Jani Suomalainen, Kimmo Ahola, Pekka Ruuska, Mikko Majanen and Mikko Uitto
VTT Technical Research Centre of Finland Ltd, Finland

Keywords: 5G, Mobile Network, Security, Micro-segmentation, Virtualization, Isolation.

Abstract: The forthcoming 5G mobile networks shall be heterogeneous in nature and embody a large number and variety of devices. Moreover, Internet of Things applications – like surveillance and maintenance – will use 5G extensively due to its high availability and quality of connectivity. However, the heterogeneous services, applications, users, devices, and the large amount of network traffic will bring challenges for the security of the mobile network. It will be important to provide isolated segments from the network for applications that require a high level of security. This paper presents the potential of micro-segmenting 5G networks. Micro-segmentation is a concept that has been considered in data center networking to enforce the security of a data center by monitoring the flows inside the data center. In this paper we describe how the micro-segmentation concept could fit into the 5G security architecture and provide scenarios of how software mobile networks can facilitate securing IoT.

1 INTRODUCTION

Mobile networks have substantially evolved throughout the years. The first generation (1G) of mobile networks embodied analog telecommunications standards, such as Nordic Mobile Telephone (NMT). 1G systems evolved into the second generation of mobile networks (2G), in which digital telecommunication standards, such as the Global System for Mobile Communications (GSM) were included. 2G systems also introduced the Short Message Service (SMS), plain text-based messages. The information transfer rates were relatively low in 1G and 2G systems and these systems were able to provide only elementary services. The information transfer rate was improved in the third (3G) and fourth (4G) generation of mobile networks. This enabled the introduction of new use cases and services, such as mobile video and Voice over IP (VoIP). The upcoming fifth generation of mobile networks (5G) introduces even further improvements to the information transfer rate, latency, and availability as well as reducing management and operating costs.

The completely new use cases, scenarios, and services included in 5G bring large concerns as to the security of the mobile network. For example, an Industrial Internet of Things (IIoT) company may use the 5G network for its operations, such as video monitoring and factory control. If an adversary gets access to the service and the network, the consequences could be crucial. The more operations that are placed in the network, the more opportunities there are for an adversary to do significant damage. Consequently, the network and the used service should be highly secure and isolated from the rest of the network. The threat landscape should be minimized.

The topics of virtualization, Network Function Virtualization (NFV) and Software-Defined Networking (SDN), and big data analysis are said (5G PPP Architecture WG, 2017; 5G PPP Security WG, 2017) to be important aspects of 5G mobile network security. Moreover, it will be important to provide customized network security depending on the demands of the used service. For example, if a high data rate or a small delay is desired for a service, the security level of the service should be designed so that these two Key Performance Indicators (KPIs) are kept at the desired level. The security level of the service may depend on multiple factors, such as access control, security monitoring, privacy, trustworthiness metrics, and isolation, to name a few. For example, the security level of the service could be customized by providing the possibility of using different authentication methods with different privacy characteristics. 5G networks will use virtualization heavily, because there is a need for efficient and dynamic techniques

Mämmelä, O., Suomalainen, J., Ahola, K., Ruuska, P., Majanen, M. and Uitto, M.
Micro-Segmenting 5G.
DOI: 10.5220/0006662700170028
In *Proceedings of the 3rd International Conference on Internet of Things, Big Data and Security (IoTBDS 2018)*, pages 17-28
ISBN: 978-989-758-296-7

to deploy virtual network functions (VNFs). Therefore, efficient and effective access control and isolation mechanisms are mandatory for virtualization platforms, such as Moby containers (Moby project, 2017; Docker Inc., 2017b).

This paper continues our previous work (Mämmelä et al., 2016) on the micro-segmentation concept. Our previous work focused on introducing the concept, originating from data centers, into 5G mobile network security. This paper provides further insight into the potential of the concept and provides use cases for leveraging software mobile networks to secure Internet of Things (IoT).

We survey related work in Section 2, describe the micro-segmentation approach for 5G mobile networks in Section 3, and provide a description of the testbed and use cases in Section 4. The conclusion and possibilities for future work are presented in Section 5.

2 RELATED WORK

2.1 Security in 5G Mobile Networks

5G mobile networks (Agiwal et al., 2016) are currently being designed and the vision is that 5G will embody completely new services, applications, and services compared to the previous mobile networking scenarios. This brings closer attention to the security of the network. With the advent of the IoT, more and more devices are becoming connected to the Internet and companies will be using 5G for many of their operations. Examples include remote video surveillance, remote control, autonomous vehicles, and remote surgery, etc. In a case where an attacker obtains access to the service and the network, the consequences could be dramatic. Moreover, the novel and diverse business requirements of vertical sectors have rendered current network security approaches inadequate. Consequently, network security must be carefully designed in 5G.

5G security relates to various requirements (5G PPP Security WG, 2017; 3GPP, 2017a). 5G should provide a higher, or at least equal, security and privacy level compared to that of 4G. This means that 5G should be able to deliver a Service Level Agreement (SLA) for availability, security, resilience, latency, bandwidth, and access control from an end-to-end perspective to verticals. As 5G infrastructures are heterogeneous and complex, they require security to be dealt with at multiple levels and across domains. This means that automation of 5G security will be important. 5G networks should include security monitoring that can be used for detecting advanced cyber security threats. Also, coordinated monitoring between different stakeholders and systems should be supported. We describe an approach for an adaptive management of security and also describe an interface for sharing real-time security monitoring information between different 5G stakeholders.

Mobile networks deploy different access control mechanisms. Subscribers' access control has been based on Authentication and Key Agreement (AKA) protocol (3GPP, 2017b), which decouples access control policies from enforcement. For instance, in 4G the service operator enforces access control in base stations (eNodeB) and uses a Mobile Management Entity (MME) to query authorization from the Home Subscriber Service (HSS) in a subscriber's home network. Operators also deploy firewalls and IPsec tunneling approaches to secure core networks, service providers, and subscribers. Firewall mechanisms – like Network Address Translation (NAT) and screening in the network edges – hide addresses of vulnerable services and ensure that all incoming data flows originate from the trusted IP address spaces.

Management of access control policies in firewalls is a challenge. Networks are complex (with a large amount of actors and functions) and communication is heterogeneous (protocols ranging from IPsec to Radius and Diameter, IPv4 and IPv6 as well as LTE protocols like Stream Control Transmission Protocol, GPRS Tunneling Protocol, and Session Initiation Protocol) as well as dynamic (new protocols and communication paths emerge as the network evolves). Hence, controlling which nodes can communicate using what protocols yields a massive amount of access control polices. In 5G, new radio access interfaces and cooperation between operators may further increase the paths that must be protected with access control approaches. To ease operations and costs, typical firewall solutions already support centralized management. However, in practice the security policies are not so fine-grained as they could be in theory. Access control is typically enforced only in critical paths in the network. In 4G, networks firewalls have been used, e.g. between:

- Mobile infrastructure and the Internet – to protect infrastructure and subscribers against remote attacks.
- Core networks and service provider's infrastructure – to protect infrastructures from attacks originating from mobile subscribers.
- Core networks between partners – to protect important centralized functions, such as HSS, from signaling attacks.

Other challenges for firewalling approaches include speed and reliability requirements. The network must be working 99.999% of the time. The costs of firewalling, both investment and operation, are critical factors when operators define the scope and granularity of their firewall deployments. Our work provides an alternative approach for access control. We propose an – easy to understand and easy to manage – concept for fine-grained access control.

Our approach can be flexibly used to protect different parts and individual services and assets of the mobile network. The proposed approach targets easing the management and reducing configuration errors (as a micro-segment must be configured to address only those threats that are relevant for the segment user) as well as lowering costs (as firewalls can be replaced with software networks).

2.2 Software Networks and Slicing

The key elements of 5G network design and security will be network virtualization (Khan et al., 2012; Wang et al., 2013; Liang and Yu, 2015), network slicing (Ericsson, 2014; Ericsson, 2015b; Ericsson, 2015a; NGMN Alliance, 2016), and network programmability (5G PPP Architecture WG, 2017). In network virtualization, the logical network components are decoupled from hardware. This means that it is possible to create isolated parts of the network. Network slicing will be an essential concept of 5G mobile networks. In this approach, nodes and communication related to particular applications are isolated from each other. A single network slice is thus a logical instantiation of a physical network with all the functionalities needed to run a particular service. Finally, the goal of network programmability is to control the behavior and communication of network devices and flows with software while operating independently from network hardware.

All these elements can be implemented with SDN and NFV technologies. The former may be used for monitoring and controlling specific network parts, while the latter can be used for the virtualization of mobile network entities and functions. In a 4G network, these functions are Packet Data Network Gateway (PGW), Serving Gateway (SGW), Mobility Management Entity (MME), load balancing, traffic monitoring, and QoS, etc.

The vision of 5G is as a flexible and dynamic system in terms of use cases, user equipment, radio access network technologies and core network services. The complexity of the system will become even more evident if new deployment models, such as third party and multi-operator deployments, and new third party APIs, enabling, e.g. full configuration control of network functions, are supported. This means that it is important to reduce the complexity of operating security aware 5G services, provided by a Mobile Network Operator (MNO) or Virtual Mobile Network Operator (VMNO), specifically in terms of trust. Users and operators would therefore need up-to-date information on the trustworthiness of the 5G network.

2.3 Micro-Segmentation

The concept of micro-segmentation was originally introduced by VMware to data center networking (VMware, 2014; Miller and Soto, 2015; Fulton III, 2015; Huang, 2015). The main idea is that in addition to network security at the perimeter, data center security should focus on the attacks and threats coming from the internal network. Data center network security has generally considered securing the network perimeter by having firewalls that filter incoming network traffic to the data center.

However, once adversaries get past the perimeter by bypassing the firewall, they are free to move laterally inside the data center and carry out their attacks. Micro-segmentation in effect enables security monitoring inside the data center in addition to the traditional network security at the perimeter, i.e. between internal components in addition to between the external and internal network. In general, the traffic inside the data center is differentiated into small isolated parts, i.e. micro-segments depending on the traffic type. With micro-segmentation, a strict micro-granular security model can be adopted that ties security to individual workloads. Also, it provides the agility to provision policies automatically. VMware has developed an NSX Network Virtualization and Security Platform (VMware, 2017), which enables micro-segmentation. Basically, through NSX it is possible to create whole networks in software and embed them into the hypervisor layer, which is abstracted from the underlying physical hardware. Software-defined policies can enable more flexible network security than a manual configuration work would.

Security advantages of micro-segments include, e.g.:

- Writing policies in a centralized manner for small (application specific) networks is simpler and less error prone than defining low level device configurations in a distributed manner or when defining policies for large heterogeneous networks.
- Micro-segmentation provides an easy to understand concept for adjusting and controlling security. For each micro-segment, the amount of secu-

rity policies is small and policies are easy to keep consistent.

- Control can be automated to adjust situations and threats in the data plane.
- Centralized controllers have knowledge of the state of the whole network. Consequently, it is easier to develop sophisticated control functions. This awareness can be extended if controllers can cooperate across domains.
- Each micro-segment can enforce different access control policies. Such fine-grained controls can be used to prevent some threats from spreading. In case an adversary successfully compromises one micro-segment, the threats do not spread to micro-segments that have been isolated from the compromised one.
- Traffic flows within a slice or micro-segment are more homogeneous. Hence, security monitoring can be more easily customized and focused compared to that of heterogeneous networks with traffic flows from various applications.
- Security enforcement (blocking of malicious nodes) is done in low-cost software switches. Blocking can be executed at any switch in the micro-segment that is the closest to the attack source. Hence, effects of attacks to network can be minimized.

Thus, the concept increases the effectiveness and trustworthiness of network security controls but may also lower the operating costs.

2.4 Security for the IoT on Mobile Networks

In general, the security mechanisms designed to protect communications must provide assurances for confidentiality, integrity, authentication, and non-repudiation of the information flows (Granjal et al., 2015). For IoT based communication availability and resilience are also two important requirements. For example, IoT embedded devices could be infected by a botnet (Traynor et al., 2009) that is capable of executing Distributed Denial of Service (DDoS) attacks. Privacy is also one of the major concerns in IoT systems, as there can be applications that handle critical information, such as medical records or factory controls. Moreover, IoT devices do not have methods for protecting the privacy of location information. IoT devices also have limited processing and battery capabilities and they may not be able to support all available authentication methods. IoT devices have different behavior patterns compared to user terminals. In IoT networks, large amount of devices may become active at the same time and send short signaling bursts to the network. Hence, a large number of simultaneously IoT authenticating devices may cause an overload situation for authentication services (Jover, 2013).

3 MICRO-SEGMENTING 5G

This section describes how the concept of micro-segments can be utilized in 5G mobile networks. First, we provide definitions and discuss the role of micro-segments in a 5G ecosystem. Then, we propose a framework to illustrate the technical building blocks that are needed to extract the potential from micro-segments.

3.1 Micro-Segments in a 5G Ecosystem

In a 5G network, a micro-segment can be defined as a logical network portion that is decoupled from the physical 5G hardware and dedicated to a particular application or user group and that hosts functions with the same or similar security requirements. Micro-segments can be considered as a sub-slice in which the emphasis is more on network security. While an end-to-end slice provides all the necessary 5G functions that are needed to organize an end-to-end service, a micro-segment may contain only a subset of functions (sometimes only one).

Several micro-segments can be chained together to create end-to-end connectivity that maintains application isolation. Each micro-segment can have its own fine-grained access controls as well as tuned security policies and mechanisms supporting unique application specific trust models. Micro-segmentation enables the creation of smaller and less heterogeneous parts in the network. This means that better efficiency and accuracy can be achieved for security.

For instance, requirements for an end-to-end slice could be to include VNFs for radio access technology (RAT) and traffic acceleration from the access network; for packet routing and mobility management from the serving network; for subscriber authentication and accounting functions from the home network; as well as functions for Internet connectivity from the transport network. For a micro-segment, the amount of functions can be smaller. RAT and acceleration functions in base stations or mobile edge data centers at different locations can be in different micro-segments. Mobility management, routing and authentication functions reserved for different quality of service levels can be kept in isolated micro-segments.

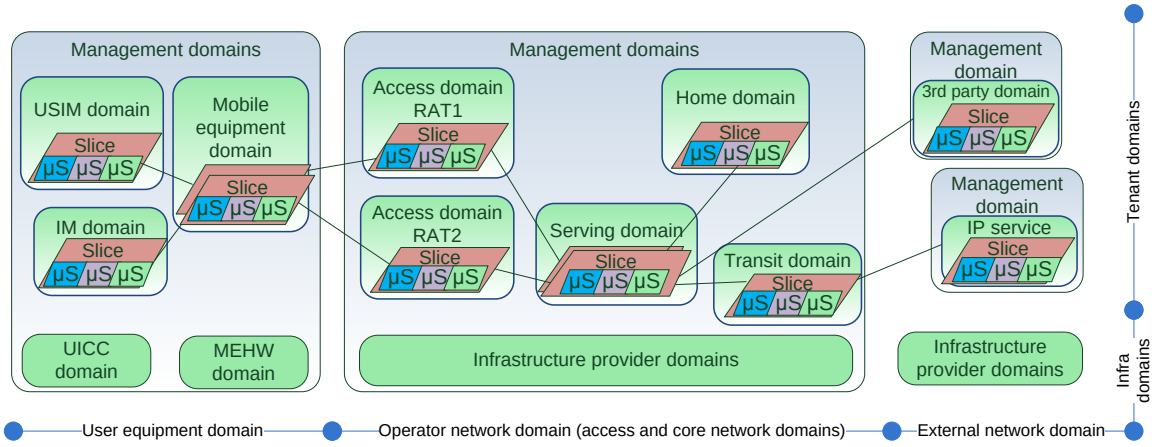

Figure 1: Micro-segmentation in 5G security architecture (5G-ENSURE, 2017; 5G-ENSURE, 2016) - a micro-segment is a fine granule security block that controls flows within administrative domains and application-specific slices.

Solutions that are known to be more vulnerable, e.g. some authentication mechanisms may be vulnerable to DoS attacks, can be kept in a micro-segment to which only trusted nodes are given access.

Micro-segments can be deployed to different parts of the 5G architecture. Figure 1 shows how the micro-segments map to the 5G security architecture (5G-ENSURE, 2017; 5G-ENSURE, 2016). The illustrated domain model is an evolution of the 3GPP's security architecture: it emphasizes the emerging infrastructure sharing by separating domains horizontally into infrastructure (hardware) and tenants (software). The management domains highlight the focus of 5G towards more flexible and cost-efficient management. Furthermore, the emerging concept of slicing is illustrated with domains distributed across different tenant domains. As slices, micro-segments can be located in all domains. In user equipment domains, the micro-segments are visible in device specific application isolation methods. In access, serving, transport, home networks, and potentially also in 3rd party and IP service providers networks, micro-segments are deployed as logical network portions.

Micro-segmentation can be used to divide 5G networks – horizontally and vertically – into portions in which different security controls can be applied. Micro-segmentation enables application-specific customization of security services. A micro-segment could, for instance, provide its own AAA and security monitoring functionalities. Consequently, different segments have different security or trust levels, i.e. can be trusted to address different security risks.

Generally, there are four entities involved in the micro-segmentation concept: a micro-segment provider, a micro-segment controller, a micro-segment subscriber, and a security function provider. A micro-segment provider is the infrastructure provider that provides micro-segments for micro-segment controllers on top of virtualized hardware. A micro-segment controller is an entity that controls the software switches and SDN controllers in micro-segments and collects event information, e.g. network statistics. This actor can be the MNO. A security function provider is an entity providing micro-segment specific security controls. For instance, there may be security monitoring and inferencing functions acquiring security information from micro-segments through the APIs that the micro-segment controller keeps open. The actor can be a micro-segment provider, a third-party, or the micro-segment subscriber. A micro-segment subscriber is an organization, company, or application service provider requiring isolated 5G network services for a particular application.

3.2 Micro-Segmentation Framework

The micro-segmentation concept is based on SDN and NFV technologies. A micro-segment is essentially a software network that is located on top of virtualized hardware infrastructure. The realization and orchestration (i.e. autonomous provisioning) of micro-segments requires different management, control and security functions. The proposed framework, illustrated in Figure 2, identifies these building blocks of micro-segments.

A virtualization platform separates 5G infrastructure from 5G software functionality. The platform provides a common interface that hides the details of hardware components (starting from base stations to switches and cloud computing centers) from SDN (software switches and controllers) and 5G VNFs (software components providing mobile network functionality). The virtualization is controlled

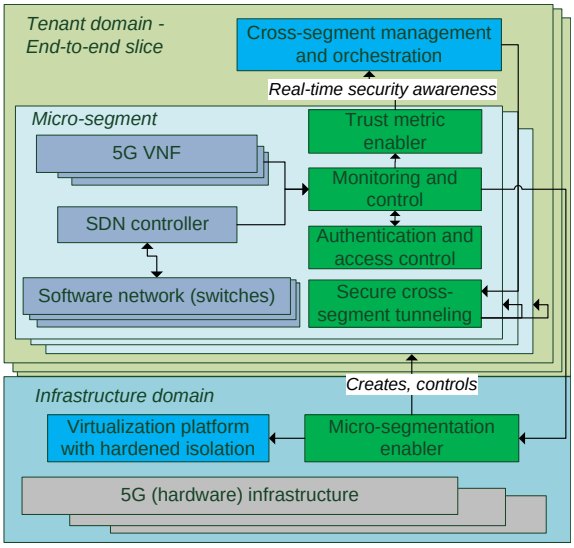

Figure 2: Micro-segmentation framework.

by *the micro-segmentation enabler*, which uses the virtualization platform to create coarse granule slices and fine-granule micro-segments. The virtualization platform is responsible for isolating the micro-segment from other applications and users using the same infrastructure. Consequently, the platform must apply strong – hardened – access control and trust verification mechanisms, at operating system and hardware level, to prevent intrusions from other micro-segments.

5G subscribers may be able to handle different risks and thus they may have different requirements for security levels. Micro-segments host application or subscriber specific security functions, such as *authentication and access control*, as well as a monitoring and inference-based security controller. *Authentication and access control* function can support different mechanisms – with different security levels and costs – which are suitable for the micro-segment subscriber. *The monitoring and inference-based security controller* collects security event information from SDN and VNFs within the micro-segment and – using application specific inferencing algorithms – deduces risk-level information, and also adapts a micro-segment's behavior and control functions.

End-to-end connections through 5G networks are created dynamically by chaining micro-segments and their services. Multi-segment and multi-domain cooperation is coordinated by cross-segment management and orchestration functions. Micro-segments support the creation of end-to-end connections over multiple segments by providing tunneling functions for generating secure Virtual Private Networks (VPNs) between segments.

To enable managers to orchestrate connectivity that fulfill subscribers' trust requirements, the managers must know how trustworthy each available segment is. Real-time awareness on the trust-level of each micro-segment must be provided dynamically for cross-segment managers to enable them to select and change connected segments, e.g. in case of a detected threat or attack situation.

A trust metric enabler provides an interface to acquire knowledge on the security state of a micro-segment. It provides real-time metrics that help both cross-segment managers and micro-segment subscribers to track segment's current trust-level. On the other hand, a trust metric enabler can be used to deliver inferred knowledge with coarse granularity so that the information remains simple and usable as well as hiding privacy and operator critical details.

4 USE CASES FOR THE IoT IN 5G

The building blocks of the micro-segmentation framework can be based on alternative technologies and implementations. We developed a testbed that was used to demonstrate three IoT use cases related to monitoring, authentication, and trust. The testbed and use cases are described in the following subsections.

4.1 Testbed Implementation

The testbed consists of software security enablers ; open source components for big data analytics, network virtualization and softwarization; IoT devices acting as service provider and adversaries; as well as a software implementation of mobile networks.

The technology selections for the testbed are illustrated in Figure 3. We developed three enablers to realize the micro-segmentation framework: a Trust Metric Enabler, a Security Monitoring Enabler, and a Micro-Segmentation Enabler - cf. (5G-ENSURE, 2017) for specifications and implementation details. Further, for the use cases, the testbed was integrated with a Privacy Enhanced Identity Protection Enabler (Baltatu et al., 2017, p. 89–104) and a Compliance Checker Enabler (Klaedtke and Sforzin, 2017, p. 252–260).

The micro-segmentation enabler facilitates the creation, deletion, and control of micro-segments. The enabler uses a modified OpenVirteX software (OpenVirteX Project, 2017) to create micro-segments and uses Ryu for an SDN controller (Ryu SDN Framework Community, 2017). The development and testing was done using the Mininet environment (Mininet Team, 2017), which emulates an

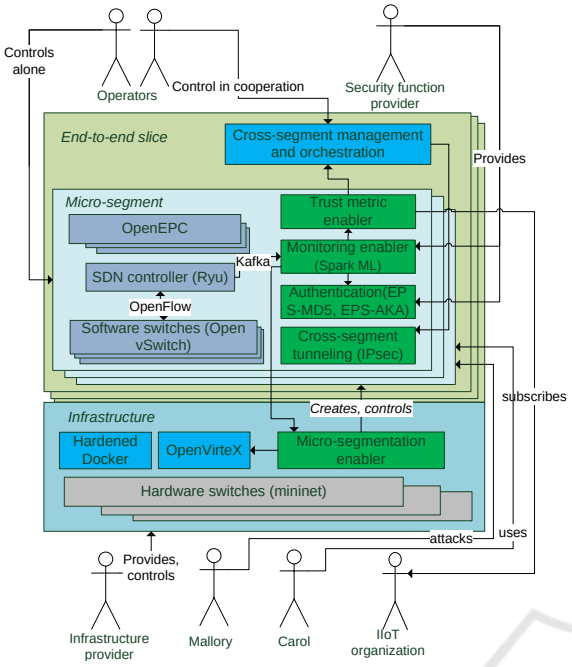

Figure 3: Micro-segmentation technology selections for prototyping and roles in the IoT use case.

OpenFlow-supported network and hosts, using Open vSwitch virtual switches (Open vSwitch, 2017) and Linux namespaces. The enabler also provides a web-based GUI for illustrating the topology of a virtualized network (see Figure 4). Secure cross-segment connections are enabled with IPsec. Furthermore, the micro-segmentation enabler provides an interface that can be used to collect traffic statistics as well as information topology change and authentication events.

Authenticator / access controller – This entity is included in the micro-segmentation enabler and it controls which node can access which micro-segment. The access control is accomplished using the Extensible Authentication Protocol authentication with the support of two alternative methods: EAP-MD5 (Extensible Authentication Protocol Method for MD5 hash) and EAP-AKA (Extensible Authentication Protocol Method for 3rd Generation Authentication and Key Agreement). The authentication in wired Ethernet environments is done using the Extensible Authentication Protocol over LAN (EAPoL). Actual implementation uses wpa_supplicant supplicant (Malinen, 2013b), hostapd authenticator (Malinen, 2013a), and FreeRADIUS (FreeRADIUS, 2017) server. When using the Privacy Enhanced Identity Protection Enabler as an additional access control service, there is a need to use modified wpa_supplicant and hostapd versions Privacy Enhanced Identity Protection Enabler (Baltatu et al., 2017, p. 89–104).

The security monitoring enabler, which tracks and autonomously controls a micro-segment's security, was build on top of the Kafka (Kafka Project, 2017) and Spark (Spark Project, 2017b) frameworks. Kafka provides a message brokering solution based on the publish and subscribe paradigm. Spark provides libraries for near-real time analysis and processing of streaming data. Both frameworks are very scalable as they inherently support cluster-based computing. The monitoring enabler is a Spark/Python application that subscribes and processes monitoring data streams (Kafka/JSON) from the micro-segmentation enabler. The enabler captures authentication and topology events, analyses network traffic statistics for anomalies, performs risk-level analysis on detected anomalies, and triggers control actions on the micro-segment.

The anomaly detection feature was implemented using streaming-k-means (Freeman, 2015) Spark's machine learning library (Spark Project, 2017a). The algorithm builds a model of the normal behavior by clustering traffic event information. Anomalies are detected by calculating distances between modeled cluster center points and new event information and checking whether anomaly thresholds are exceeded. The streaming-k-means learns normal behavior continuously and is therefore suitable for dynamic 5G environments where new nodes may roam to the network and new connections may be initiated at any time. When the algorithm detects an anomaly, which sufficiently exceeds thresholds levels, it quarantines suspected nodes from the micro-segment and publishes warning notifications.

The trust metric enabler provides an interface for cross domain exchange of security monitoring information. The implementation is a python application integrated to trust metric clients and a security monitoring enabler with Kafka. The clients request trust metrics by specifying trust models with three primitives: a trust model may require a particular function to be available as well as to set upper and lower limits for specific security related KPIs. For instance, a trust model may require use of particular security algorithms and follow the trustworthiness of devices in a micro-segment. The trust metric enabler tracks that these requirements are fulfilled and provides near real-time status information for the micro-segment subscribers. The implementation also provides a web-based GUI - trust indicator - for visualizing different trust metrics as 'traffic lights'.

The use case application is comprised of *an IoT real-time monitoring video*, which holds a video server running in Raspberry Pi 3 as well as a laptop video client. Dynamic Adaptive Streaming over HTTP (DASH) protocol was utilized in order to pro-

vide a possibility for adaptation against network attacks by dropping the video quality. For this, the actual attacks could be visualized in the video player.

Adversaries – *a malicious IoT botnet* trying to disable use case applications – were modeled using Raspberry Pi 3 devices with hping packet generators, capable of generating a large amount of crafted packages towards victim interfaces.

OpenEPC (OpenEPC Project, 2017) implements LTE core network functions as well as an emulated base station (eNodeB) and user terminal. It was utilized to generate traffic flows for the use case analysis, though the component was not integrated to the operative micro-segmentation testbed.

4.2 Use Cases

This subsection illustrates the potential of micro-segmentation with three use cases[1]. Actors and their roles in relation to technology are illustrated in Figure 3. The general storyboard is as follows. A company has a video surveillance equipment in its factory where cameras and other equipment are connected using a 5G network. Carol, a company employee, uses her 5G mobile device to view the video surveilance service. This service should be highly secure and isolated from the rest of the network so that malicious nodes are not able to access or disturb it. Carol is receiving data to her smart phone from a factory surveillance camera that is connected to an IoT network. An attacker that wants to invade the factory – Mallory – has two options: either blackout the camera by using an IoT botnet (Subsection 4.2.1) or track the movement of Carol (Subsection 4.2.2). The company purchases hardened security from the operator and receives near real-time information on the trustworthiness of the network (Subsection 4.2.3).

Physical topology and micro-segment topology (logical) used in the use cases are shown in Figure 4. The physical topology consists of eight OpenFlow enabled switches, two servers, one authenticated client, and two devices belonging to the malicious IoT botnet. The servers, client, and IoT botnet are shown in white. A micro-segment (shown in green) resides on top of the physical topology. OpenVirteX virtualizes the switches and the virtual network sees only its part of the network topology. The micro-segment is isolated and communication between other micro-segments is not supported. The micro-segment topology consists of four virtual switches (shown in green), a server, a client, and two devices belonging to the malicious IoT botnet.

[1]A video animating the use cases is available in https://youtu.be/xfuBEpt4l8Y.

4.2.1 IoT Botnet Detection

The storyboard of the first use case is the following. An attacker tries to blackout the camera by launching denial of service attack from an IoT botnet. In the case of such an attack or intrusion, security monitoring should be able to detect malicious nodes and remove them.

In this scenario, the micro-segmentation enabler is used for creating and deleting micro-segments, adding and deleting nodes from micro-segments, and providing strong access control to the micro-segment. The security monitoring enabler monitors traffic flows inside the micro-segment, uses machine learning to model 'normal' behavior and to detect any anomalous behavior.

In the scenario, we also have several connected IoT devices. As these IoT devices are harder to keep up-to-date and may not have the capabilities for heavy security mechanisms, they have been compromised by Mallory and turned into a botnet. Mallory can then instruct this botnet to initiate a denial of service (DoS) attack against the monitoring camera. When an anomaly revealing such an attack is detected, the micro-segmentation enabler will be automatically be contacted to quarantine suspected flows. The advantages provided by micro-segmentation are evident from the quick recovery of service quality.

Through the use of the enablers, we are able to highlight the following points: 1) threats by the IoT towards 5G infrastructure/services can be mitigated by isolating the traffic flows (with SDN based virtualization) and 2) access controlled and monitored micro-segments enable quick detection of threats and autonomous reaction. In non-segmented networks such a reaction would not be possible as some services may occasionally experience heavy loads. But in application specific micro-segments large amounts of messaging towards a service that has been previously only transmitting, can be interpreted as an attack that can be blocked without human intervention.

4.2.2 Alternative Authentication Methods

The micro-segmentation concept in a 5G network can support different authentication methods depending on the required security level of the service. For example, SIMless authentication can be provided for cheap embedded devices. Authentication mechanisms for battery restricted devices may also be more lightweight when the application can accept lower security or privacy levels. For better security and privacy, EAP-AKA could be provided.

In the second use case, the storyboard is the following. An attacker will try to track the movement of

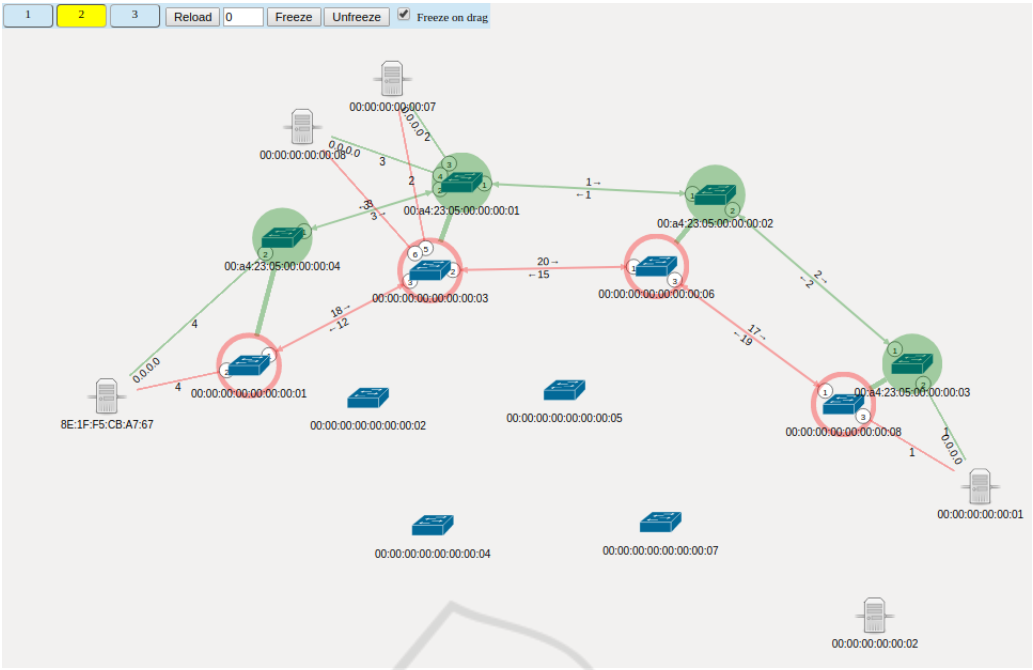

Figure 4: Micro-segment topology.

Carol. A possible adversary could do a considerable amount of damage to the factory when Carol is not at the factory. The tracking attack (Herzberg et al., 1994) is possible as Carol is using the network with her smartphone and her phone is associated with an (unprotected, unencrypted) international mobile subscriber identity (IMSI), which identifies Carol to the network.

For achieving a secure and private service to operate the video service, the micro-segmentation enabler will create an isolated micro-segment for it. For access control, the EAP-AKA implementation of the Privacy Enhanced Identity Protection enabler will be used. The enabler will protect the long term identifier (IMSI) with attribute-based encryption. In this way, the IMSI is hidden and it is not visible for tracking by possible adversaries.

In this scenario, it is possible to show that 1) we can provide services that are isolated and provide authentication that fulfills subscribers' requirements in respect to privacy-level or capabilities of the devices and 2) micro-segments can support different authentication mechanisms. Authentication mechanisms can be easily customized according to the micro-segment subscribers' requirements.

4.2.3 Trust Measurements

In subscribing to a micro-segment and using it in surveillance, Carol's company need to be aware that the network can be trusted and if something happens they must be notified in at once. The trustworthiness of the micro-segment depends on several factors that are specified by the micro-segment subscriber. In this use case the following metric is used to define whether the segment is trusted or not:

- Functions and services hosted in the micro-segment must be deployed to trustworthy platforms. Particularly, hosting containers must be up-to-date and isolated from other applications.
- The amount of devices that use lightweight and potentially vulnerable security mechanisms is limited.
- The micro-segment must have security functions for authentication and authorization as well as for security monitoring. Security functions mechanisms must be verifiable, running, and working.

In this use case, the services that a micro-segment hosts are brought in a Moby container (Moby project, 2017) that is an open source version of Docker (Docker Inc., 2017b). The network could be attacked by an adversary that utilizes a vulnerable container with outdated software. For instance, the docker engine/hypervisor may also host a malicious container able to also compromise other applications, which are running in the same engine, and thus affect the security and behavior of the micro-segment. Hence, the security level of each container is measured and isolation is enforced. Essentially, in this

use case the micro-segment subscriber has specified a policy that the Docker engines should be dedicated for services. When starting containers, with docker-compose (Docker Inc., 2017a), we check that this policy is enforced.

The security monitoring enabler must be running in the micro-segment in order to detect and react to potential DoS situations. We use the compliance checker to verify that the micro-segmentation enabler really quarantines the flows detected by the monitoring enabler. The compliance checker gets notifications from both enablers and verifies that they occur within a given time window.

Micro-segmentation and security monitoring enablers also collect information on what authentication mechanisms devices are using to access the micro-segment. In this case, the trust level decreases if there are MD5 authenticated IoT devices in the network. Large amounts of IoT devices that authenticate to the micro-segment using lightweight authentication mechanisms (and that are more likely to fall into a botnet) pose a threat to the availability of services.

The trust metric enabler is able to receive trust measurements from the micro-segment through the micro-segment monitor. The trust metric enabler will combine this information and notify the company of changes in micro-segments trust. The company can visualize the trust information for Carol, e.g. by displaying green, yellow or red icons or showing textual information stating that the company's trust requirements are either met or are not satisfactory. Through the use of the trust metric enabler, service providers and end-users are made more aware of a network's trustworthiness.

4.3 Benefits and Costs of Micro-Segmentation

Micro-segmentation aims to divide the network into smaller segmented security zones. This has the potential to improve the speed of detecting intrusions and deleting malicious nodes. Also, with micro-segmentation, the attacks can be tied to a specific location and the spread of attacks can be minimized. Micro-segments can also be customized based on the needs of the application. For example, some micro-segments may require a strong authentication protocol and encrypted traffic and some may only allow particular types of communication or protocols.

4.3.1 Managing Diversity in Mobile Networks

To evaluate the potential of micro-segments, we measured the diversity of traffic flows in our mobile network testbed. The aim was to study how heterogeneous and diverse the communication is in different parts of the network and thus enable us to assess whether micro-segmentation could be applied to isolate more homogeneous parts of network.

We used IoT and video streaming applications to generate realistic traffic flows and to trigger signaling from our OpenEPC LTE functions. Then we collected traffic statistics flowing between different LTE interfaces: from user terminal to eNodeB (net_c), from eNodeB to serving gateway (net_d), from serving gateway to packet gateway (net_b), and from packet gateway to internet gateway (net_a).

The traffic statistics crossing different mobile network interfaces, and caused by two applications (plain IoT, as well as combined IoT, and video) are illustrated in Figure 5. It illustrates how the amount of packets and the amount of involved protocols increases when the application amount increases. The application layer (IP) traffic flows through the whole network but layer-2 signaling depends on the interface. The numbers over 100% in the figure are due to the GPRS tunneling between eNodeB and serving gateway (the packets in the tunnel can have two IP and UDP headers depend-ing on the inner packet). CoAP packets were sent only once per 2 seconds, while video packets were trans-mitted with a much smaller packet interval. This explains the difference in CoAP packet percentage numbers, as well as other protocol numbers, between the two scenarios.

The results illustrate that micro-segmentation can therefore be used to isolate different application cases to get more homogeneous traffic flows that are easier to monitor. As the traffic volumes of IoT traffic are much more moderate compared to those of combined video and IoT data, the monitoring functions in a micro-segment dedicated for just IoT traffic needs significantly less resources and can also perform more complex analyses.

Micro-segmentation can also be utilized when defining authorization policies for different domains. For instance, as stream control transmission protocol is present and normal only in an access network (eNodeB to serving gateway) it can be filtered from micro-segments where a packet gateway or an internet gateway is hosted. Hence, packet and internet gateways can be isolated from vulnerabilities related to these protocols.

Scenario:	**IoT + video**				**IoT only**			
Capture point:	net_a	net_b	net_c	net_d	net_a	net_b	net_c	net_d
Protocol	%Packets	%Packets	%Packets	%Packets	%Packets	%Packets	%Packets	%Packets
Ethernet	100,00	100,00	100,00	100,00	100,00	100,00	100,00	100,00
Logical-Link Control	0,73	0,46	0,73	0,66	29,13	26,19	25,56	15,88
Spanning Tree Protocol	0,73	0,46	0,73	0,66	29,13	26,19	25,56	15,88
Internet Protocol Version 4 or IPv6*	99,10	99,54*	99,18	195,67	64,08	73,81*	70,61	93,13
User Datagram Protocol	64,04	163,34	64,08	160,86	56,31	119,05	57,83	81,12
Session Initiation Protocol	0,06	0,06	0,06	0,06				
Real-Time Transport Protocol	62,54	62,95	62,69	61,92				
MP4V-ES	62,54	62,95	62,69	61,92				
GPRS Tunneling Protocol		99,23		97,74		64,29		44,42
GPRS Tunneling Protocol V2		0,03						
Domain Name System	0,09	0,09	0,09	0,09	1,94	1,59	1,28	0,86
Constrained Application Protocol	1,35	0,97	1,24	1,06	54,37	53,17	56,55	35,84
Transmission Control Protocol	0,05		0,05	0,05				
Stream Control Transmission Protocol				0,20				4,29
S1 Application Protocol				0,05				
Internet Control Message Protocol (v6*)	0,35	0,28*	0,31	0,24	7,77	9,52*	12,78	7,73
Data	34,66		34,75	34,32				
Address Resolution Protocol	0,17		0,09	1,41	6,80		3,83	35,41

Figure 5: Traffic statistics in mobile network interfaces with IoT and video applications.

4.3.2 Costs

Micro-segmentation can also lead to additional costs compared to normal SDN-based slicing. These additional costs include, e.g. the overhead needed for the control messages that security monitoring uses. That overhead can be changed with the polling interval of flow information from the micro-segmentation enabler. A longer polling time decreases the overhead, but the acting time on anomalous events increases. Other additional costs are caused by the accuracy of rules inside switches: higher granularity of monitored traffic needs more specific rules and therefore the amount of rules increases, as do the monitoring costs: the processing cost of running anomaly detection algorithms based on the large amount of data gathered from the micro-segment switches and the storage cost of that information.

5 CONCLUSIONS AND FUTURE WORK

This paper presented the potential of micro-segmenting 5G mobile networks. We described the role of micro-segments in the 5G ecosystem and provided use cases that considered securing a video monitoring service in an IoT network. We also described the benefits and costs of the approach.

Future research is needed to ease deployment and to minimize management costs of micro-segmentation. Automatizing dynamic creation of micro-segments and instantiation of security functions, particularly in end-to-end multi-domain scenarios, is complex and solutions for orchestration are needed. Micro-segmentation-as-a-service (with e.g. container based packaging) might be one development idea to make the approach easily deployable.

The micro-segmentation approach can be utilized in different settings. A prominent micro-segmentation use case for further studies could e.g. be 5G Mobile Edge Computing, in which there is an entity, such as a server or group of servers, located between the base station (eNodeB) and the core network. This entity brings the mobile network functions closer to the edge of the network and user. In this way, it is possible to enhance the performance of the applications and reduce the delay.

ACKNOWLEDGEMENTS

This research has been performed within 5G-ENSURE project (www.5gensure.eu) and received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 671562. The work has been supported by the Challenge Finland 5G-SAFE project, partly funded by Tekes and VTT.

REFERENCES

3GPP (2017a). TR 33.899 Study on the security aspects of the next generation system, V1.3.

3GPP (2017b). TS 33.401 3GPP System Architecture Evolution (SAE); Security Architecture, V15.0.

5G-ENSURE (2016). Deliverable D2.4 - Security Architecture (draft). Technical report.

5G-ENSURE (2017). Deliverable D3.6 - 5G PPP Security Enablers Open Specifications (v2.0). Technical report.

5G PPP Architecture WG (2017). View on 5G Architecture (Version 2.0). Technical report.

5G PPP Security WG (2017). 5G PPP Phase 1 Security Landscape. Technical report.

Agiwal, M., Roy, A., and Saxena, N. (2016). Next generation 5g wireless networks: A comprehensive survey. *IEEE Communications Surveys Tutorials*, 18(3):1617–1655.

Baltatu, M., Costa, L., and Lombardo, D. (2017). Privacy Enhanced Identity Protection Enabler. 5G-ENSURE. Deliverable D3.6 - 5G PPP Security Enablers Open Specifications (v2.0). Technical report.

Docker Inc. (2017a). Docker compose - web documentation. https://docs.docker.com/compose/.

Docker Inc. (2017b). Web site. https://www.docker.com/.

Ericsson (2014). Network functions virtualization and software management. Technical report.

Ericsson (2015a). 5G Security - Scenarios and Solutions. Technical report.

Ericsson (2015b). 5G systems – Enabling Industry and Society Transformation. Technical report.

Freeman, J. (2015). Introducing streaming k-means in Apache Spark 1.2 - The Databricks Blog. https://databricks.com/blog/2015/01/28/introducing-streaming-k-means-in-spark-1-2.html.

FreeRADIUS (2017). Web site. https://freeradius.org/.

Fulton III, S. M. (2015). Microsegmentation: How VMware Addresses the Container Security Issue. The New Stack. http://thenewstack.io/microsegmentation-how-vmware-addresses-the-container-security-issue/.

Granjal, J., Monteiro, E., and Silva, J. S. (2015). Security for the internet of things: A survey of existing protocols and open research issues. *IEEE Communications Surveys Tutorials*, 17(3):1294–1312.

Herzberg, A., Krawczyk, H., and Tsudik, G. (1994). On travelling incognito. In *Mobile Computing Systems and Applications, 1994. WMCSA 1994. First Workshop on*, pages 205–211. IEEE.

Huang, G. (2015). Three Requirements For True Micro-Segmentation. Network Computing. http://www.networkcomputing.com/networking/three-requirements-true-micro-segmentation/1151379004.

Jover, R. P. (2013). Security attacks against the availability of lte mobility networks: Overview and research directions. In *Wireless Personal Multimedia Communications (WPMC), 2013 16th International Symposium on*, pages 1–9. IEEE.

Kafka Project (2017). Web site. https://kafka.apache.org/.

Khan, A., Zugenmaier, A., Jurca, D., and Kellerer, W. (2012). Network virtualization: a hypervisor for the internet? *IEEE Communications Magazine*, 50(1):136–143.

Klaedtke, F. and Sforzin, A. (2017). Component-interaction audits enabler. 5G-ENSURE. Deliverable D3.6 - 5G PPP Security Enablers Open Specifications (v2.0). Technical report.

Liang, C. and Yu, F. R. (2015). Wireless network virtualization: A survey, some research issues and challenges. *IEEE Communications Surveys Tutorials*, 17(1):358–380.

Malinen, J. (2013a). hostapd: IEEE 802.11 AP, IEEE 802.1X/WPA/WPA2/EAP/RADIUS authenticator. Web site. https://w1.fi/hostapd/.

Malinen, J. (2013b). Linux WPA/WPA2/IEEE 802.1X Supplicant. Web site. https://w1.fi/wpa_supplicant/.

Mämmelä, O., Hiltunen, J., Suomalainen, J., Ahola, K., Mannersalo, P., and Vehkaperä, J. (2016). Towards Micro-Segmentation in 5G Network Security. In *European Conference on Networks and Communications (EuCNC 2016). Workshop on Network Management, Quality of Service and Security for 5G Networks.*

Miller, L. and Soto, J. (2015). Micro-segmentation for Dummies. Technical report, VMware.

Mininet Team (2017). Web site. http://mininet.org/.

Moby project (2017). Web site. https://mobyproject.org/.

NGMN Alliance (2016). Description of Network Slicing Concept. Technical report.

Open vSwitch (2017). Web site. http://openvswitch.org/.

OpenEPC Project (2017). Web site. http://www.openepc.com/.

OpenVirteX Project (2017). Web site. http://ovx.onlab.us/.

Ryu SDN Framework Community (2017). Web site. https://osrg.github.io/ryu/.

Spark Project (2017a). Streaming K-Means Clustering - Spark Documentation. https://spark.apache.org/docs/latest/mllib-clustering.html#streaming-k-means.

Spark Project (2017b). Web site. https://spark.apache.org/.

Traynor, P., Lin, M., Ongtang, M., Rao, V., Jaeger, T., McDaniel, P., and La Porta, T. (2009). On cellular botnets: measuring the impact of malicious devices on a cellular network core. In *Proceedings of the 16th ACM conference on Computer and communications security*, pages 223–234. ACM.

VMware (2014). Data Center Micro-Segmentation: A Software Defined Data Center Approach for a "Zero Trust" Security Strategy. Technical report.

VMware (2017). NSX Network Virtualization and Security Platform - web site. https://www.vmware.com/ products/nsx.

Wang, A., Iyer, M., Dutta, R., Rouskas, G. N., and Baldine, I. (2013). Network virtualization: Technologies, perspectives, and frontiers. *Journal of Lightwave Technology*, 31(4):523–537.