SUPERFLUIDITY

a super-fluid, cloud-native, converged edge system

Research and Innovation Action GA 671566

Deliverable I3.2:

Initial security framework specification

DELIVERABLE TYPE:	REPORT
DISSEMINATION LEVEL:	PU
CONTRACTUAL DATE OF DELIVERY TO THE EU:	01/04/2016
ACTUAL DATE OF DELIVERY TO THE EU:	01/04/2016
WORKPACKAGE CONTRIBUTING TO THE DELIVERABLE:	3

EDITOR(S):	OMER GUREWITZ
AUTHOR(S):	MARK SHIFRIN (BGU), ASAF COHEN (BGU), OMER GUREWITZ (BGU), COSTIN RAICIU (UPB), Erez Biton (NOKIA-IL),
INTERNAL REVIEWER(S)	PEDRO ANDRES ARANDA GUTIERREZ (TID), JUAN MANUEL SÁNCHEZ (TELCARIA), RAÚL ÁLVAREZ PINILLA (TELCARIA)

ABSTRACT: This internal report summarizes security risks and the resulting security-related challenges in the SUPERFLUIDITY architecture. We show that besides

the ordinary security requirements of similar architectures, the migration of previously conventionally secluded application families to the virtual public environment opens new security threats and hence imposes new security requirements and challenges. We show that even though some security mechanisms can be addressed by the deployment of dedicated security VNFs providing security services to other VNFs or side-guarding the network, additional security mechanisms need to be addressed in symbiosis with the non-secured VNFs and as part of their design.

KEYWORD LIST:

INDEX

List of Figures

List of Tables

Glossary

(TO BE FILLED OUT IN THE FINAL VERSION OF THE DELIVERABLE)

SUPERFLUIDITY DICTIONNARY	
TERM	DEFINITION

Table 1: SUPERFLUIDITY Dictionary.

1 Introduction

1.1 Deliverable Rationale

1.2 Quality Review

Review Team member responsible of the deliverable: _____

VERSION CONTROL TABLE			
VERSION N.	PURPOSE/CHANGES	AUTHOR	DATE

1.3 Executive summary

Any architecture which relies on a real-time Cloud infrastructure is susceptible to security hazards. This applies more to those relying on virtualization of network functions. Accordingly, any such architecture is facing great challenges identifying and taking into account these security hazards. In this report we describe the security challenges and hazards that the SUPERFLUIDITY architecture will need to address.

1.3.1 Deliverable description

In Section 2 we provide a high level description of the main security challenges that the SUPERFLUIDITY architecture will have to address. In particular, it starts with a short discussion on security issues concerning general cloud computing and virtualization. Then we combine these two aspects and formulate the potential security challenges in the context of the SUPERFLUIDITY architecture.

In Section 3 we elaborate on each of the topics given in Section 2.3 including some examples and possible directions for which the architecture should further be explored throughout the project.

Section 4 further explores the security hazards that SUPERFLUIDITY architecture is facing by analysing several use-cases described in WP2. We provide a brief general description of each use-case and

describe possible security issues related to this particular use-case in respect to the security problems and concerns raised in Sections 2 and 3.

Section 5 concludes the report and provide some insight and conclusions.

1.3.2 Summary of results

In this report we review possible threats and corresponding security related challenges in the SUPERFLUIDITY architecture. We show that due to the migration of previously conventionally secluded application families to the virtual public environment, new security challenges and hence new security threats arise. We further show that pure security services can be deployed as isolated VNFs, providing services to other VNFs. Nonetheless, in some cases, the functional architecture should take into account and should be implemented in a symbiosis with accompanied security functions. Furthermore, we show that in some cases some security functionality should be integrated within the VNF design.

2 Security challenges

In this section we describe the security challenges that the SUPERFLUIDITY architecture will need to address. First we briefly review general cloud computing security hazards. Then we review security issues related with virtualization. And finally, we review issues with network processing on commodity hardware (i.e. NFV). Then, we combine these three aspects and formulate the possible security challenges in the SUPERFLUIDITY architecture.

2.1 Typical security issues in cloud computing

In the most simplistic view, there are two sides of security as far as virtualization on cloud is concerned. The provider must ensure security and privacy on provided infrastructure, while the user must take measures to secure their application(s) by strong passwords, encryption and authentication. The four phases of control which are normally applied to the cloud are logically depicted in Figure 1. Deterrence can be obtained by bringing to the users' attention the potential risk of not securing their data and punishment in case of attempted violations. Prevention control reduces the vulnerabilities e.g., by reducing the number of possible targets and/or hiding some of them. Detection takes action on incidents already in progress. Finally, the correction phase supplies a toolbox for fast recovering from the attacks. Clearly, the provider's objective is to try to stop the potential attacker in the earliest possible phase. All these phases should be adopted by NFV control, by adjusting to the new vulnerabilities as described in the sequel.

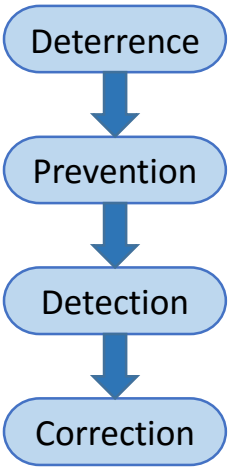

Figure 1: Cloud security control phases

Measures on the provider's side include having trusted employees and tracking and monitoring actions, especially any unusual activities. Moreover, the remotely stored data should be isolated from any potential undesired access, even if multiple copies are stored and/or cached in CDNs.

2.2 Virtualization security issues

Virtualization brings an additional layer of connectivity and the underlying control, with new security concerns. In particular, local hypervisors can be targeted by potential attackers. These attacks aim at disrupting the virtualization activity and correctness, in order to possibly damage or exploit the providers, for example, by diverting their resources. The orchestrator is a global target (e.g., by breaching the administrator's workstation) which, if successfully attacked, can cause downtime of the entire CDN.

As for security issues, as far as the virtualization technology in general is concerned, virtual platforms are subject of specialised attacks which include compromising VMs, hosting servers, malware use among others. The attacks on hypervisors might be especially severe for the overall functioning. These attacks can be both targeted on hosted and bare-metal hypervisors.

We divide the security issues concerning virtualization into the following categories:

- Snapshot and logging - Enterprises periodically store snapshots of their VMs for the possible backup. These logs include anti-malware software which should constantly be kept updated. Hence once the system is rolled back to the older version (e.g., during recovery event) some security patches could remain old. This provides a vulnerability which can be exploited.
- Virtual networking - is only one example of virtualization. To demonstrate the possible issues of virtual networking, consider a setup where multiple virtual machines are connected over a virtual switch thus providing a virtual network. Hence, since the traffic does not traverse the

actual physical network interface, in the case where one VM attacks another VM, the event will not be detected by legacy intrusion detection system (IDS) or data loss prevention (DLP) agent. This raises a question of additional overhead expressed by supplemental virtual firewalls. These risks are described in more detail in 2.3.

- Compliance - Co-hosting of several VM for different enterprises on the same domain (either physical or virtual) can come in contradiction with SLAs, regulation policies and different security demands. In particular, mixing of confidential and non-confidential information on the same storage can become an issue.
- Dynamical considerations - This issue will be in particular of interest as far as NFV is concerned. The variety of security issues with VM movements and migrations is represented but not limited to, unsecured migrations due to the load balancing, migrations due to physical failures, introduction of migration-related security policies, secured changes of physical address while attaining same virtual address, logging of migrated VMs.
- Deployment configuration and installation - This category considers deployments and erasure of VMs, rebooting of VMs after software installations. In the case where the attacker may know the reboot timings, this can be exploited for timed attack. Concurrent reboot of multiple machines can also create performance issues.
- Management and control - This category contains all issues mentioned above in the sense that a novel security-aware orchestration might be needed in order to make massive network virtualization feasible.

2.3 Security challenges in Network Function Virtualization.

Network function virtualisation promises a world where network functions are software running on commodity hardware as virtualised entities. The selling points of NFV are quick instantiation times, ease of scaling and updating of functionality. However, network configuration has traditionally been a very difficult task: ensuring a network complies with the operator policies (e.g. ACL) is difficult even with hardware appliances. Finally, running a large software base from third party vendors opens networks to attacks.

The SF vision of fast orchestration of a mix of components implemented by different vendors leaves it exposed to two possible security risks:

- Policy compliance violations – by wrongly configuring an RFB, or by deploying the wrong RFB, the resulting network configuration violates the policy of the network operators. Examples here include lack of isolation between infrastructure-critical and client traffic, allowing client traffic to reach operator services that it should not reach, etc.
- Low-level exploits – software exploits of RFB implementations fall in this category. The results of the exploits may be disabling the function of the RFB and leading to policy compliance

violations, or even worse: control of the underlying infrastructure and affecting other tenant's or the operator's own traffic.

2.4 Security challenges in SUPERFLUIDITY architecture

We now combine both aspects reviewed in previous subsections and categorize NFV specifics and SUPERFLUIDITY architecture in particular, regarding security.

We divide security issues related to NFV according to the following categories.

- Authentication and identification - This category covers issues involving password management, identification management, remote identification, virtualization of authentication servers, location of virtual machines and storage units which hold identity data, securing of identification queries and similar topics.
- Tracing and monitoring - On-the-fly control of data packets belongs in this category. The objective of these functionalities is two-fold. The first objective is to prevent potential attacks. The second is to create logs in order to facilitate the recovery process once a security incident has taken place. The possible attacks include Denial of service attacks, malware-initiated attacks on virtual resources, attempts to hamper hypervisors, etc. The monitoring tools would be represented but not limited to virtualization of firewalls, deep packet inspection (DPI).
- Policy compliance violations – verifying that the deployed network data-plane behaves according to the high level operator policy. There are two approaches in this category: testing and static analysis. Testing implies injecting test packets into the network and observing whether the outputs match the policy. Static analysis implies taking a snapshot of the data-plane configuration and testing it offline to see if it matches the policy for any given packet.
- Memory safety of deployed code – ensuring that NFV code is safe is critical to the correctness and appeal of NFV. There are two ways to do this: either write the code in a memory-safe language, such as Java (but give up performance), or use C and apply various techniques to reduce exploitation vectors (e.g. use ASLR, DEP, stack canary values, etc).
- Availability and feasibility - This category can be subdivided into two different topics. First one deals with HW targeted attacks, that may cause HW failure and loss of confidential or any other data. The loss of packets also belongs in this subcategory. The second subcategory deals with technical and architectural additions which should be implemented in order to support the security of virtualized functioning in general. This refers to virtualized network interface (vNIC), new tunnelling solution, etc.
- Performance scalability - Once hardware based network services are virtualized, a question of performance cost is raised. In particular, one should assure that the implementation of new security add-ons does not inflict a severe performance penalty as well as does not impair the performance of corresponding services as compared before the virtualization. Clearly, firewall processing should not become a traffic bottleneck. Any movement of the traffic congestion

from one point of network to another as a result of virtualization is undesirable. Hence, even if not inherently security related issue, performance scalability should be simultaneously addressed and understood.

- **Survivability and damage isolation** - This category refers to the recovery. Certain measures should be ready to invoke in order to isolate the unit that was infected by malware or compromised by any other means, and assure fast recovery to the full working status and minimization of damage. This involves deployment of special backup and logging packages by cloud providers.
- **Data plane forwarding** - Consider the virtualized network interface (vNIC). The traffic forwarding should be transparent. Practically, the implementation of virtualization comprises virtual switching. The question raised is the performance of encapsulation and opening of the traversing packets, including tunnelling solutions, and deep packet inspection. All of the above should be created on the virtualized level.

3 SUPERFLUIDITY Security challenges

In this section we elaborate on each of the topics presented in Section 2.3 and provide some examples and possible directions that should further explored throughout the project.

3.1 Authenticity and identification

The subject of identification mostly refers to the interaction with humans. Hence, due to the human intervention the flow intensity in this case is expected to be comparatively sparse unlike in the case of application oriented data flows. However, the packets contain most sensitive information, hence will always be encrypted.

3.1.1 Mobility problem

SUPERFLUIDITY will support movement of virtual servers, because of the abstraction of the virtual network from physical devices. Hence, it is critical that the dynamics of the VNFs will be followed by the dynamics of the security controls. We describe this security challenge by an illustrative use-case example.

Example (Virtualization of a server which asks for authentication upon access)

A web server, storage server, FTP server and others which are accessed by authorized users, need their authorization functioning to be virtualized concurrently with their main activity. The implication of such a virtualization includes performing authentication by VM. Hence, either storage at remote VMs of the database of user's passwords, or opening specialized secured links which serve the authentication process should be assured. In addition, the movements, i.e., migration of VM should

comply with the security needs. Hence, any kind of migration should be enhanced by specialized security and encryption protocol.

The problem is depicted in Figure 2. The left case shows the problem where the secure database (of users' passwords) should be migrated from one NFV node of presence (NoP) to another. The right case shows the problem where the database is not virtualized, which raises the need for specialized secure flow. Note, that this flow may also change its physical characteristics. Hence, the example demonstrates that constant motion and dynamic nature of the virtualized services should be supported by all adjacent security tasks.

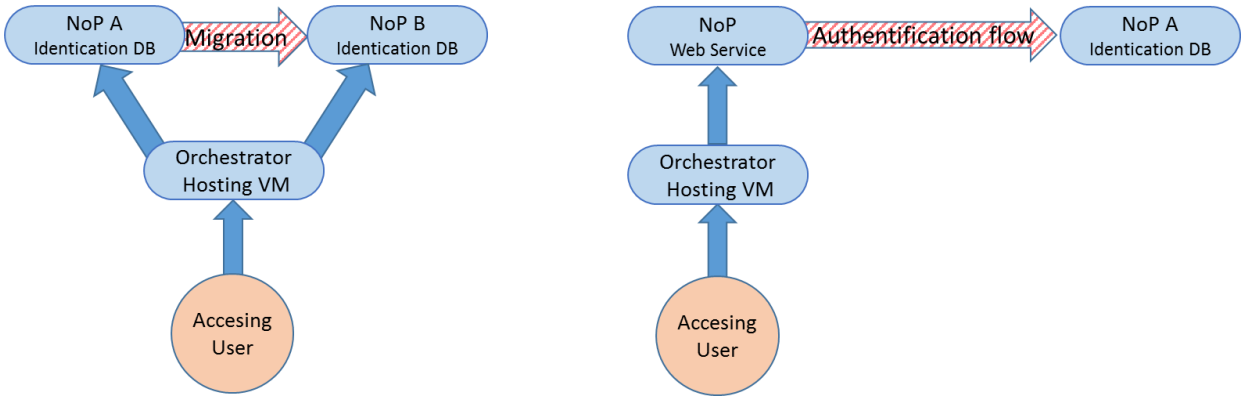

Figure 2: Possible cases of authentication by virtualized server

The encryption is always activated when personal credentials are being accessed, queried or changed.

We next apply two known identification and authentication methods and describe how they will be integrated into NFV-extensive environment. In particular, we describe the solutions that assure these methods conform with NFV security demands. Namely, we propose usage of Biometrics-based identification and Federated Identity paradigm.

We suggest a solution, CloudId, which applies to both methods.

3.1.2 Biometrics

Storage of confidential identification data (ID) in the forms of passwords does not answer the security and privacy guarantees which are normally demanded by users. Usage of biometric encrypted identification provides more confidentiality. This is because biometric data is extremely difficult to decipher. The identification process of each query is done by calculating a distance metric between the stored pattern and the pattern sent in the initiator's query, see e.g., (E. Pagnin 2014).

Since biological data can vary, this metric does not necessarily provide a clear binary answer. As a result, intelligent machine learning algorithms tailored for usage with the biometric information are


stored in the system. On the other hand, the content stored in the public domain is much more sensitive. Hence, the following demands of remote storage of biometric data are introduced:

1. The biometric data should be encrypted
2. The query should be encrypted
3. The initiator of the query cannot see the stored encrypted biometric data
4. The storage owner cannot see the result of the test of the query; only the binary answer correct/incorrect query is seen to the (cloud) storage owner.

In order to address all the above demands, specialized encryption methods should be employed.

The authors in (E. Pagnin 2014) are aware of hill-climb attacks aimed at the core of the biometric content. Note that these threats should be effectively eliminated. Also note that the mere fact of visibility of queries traversing the virtualized network, which is concurrently accessed by various tenants can provide an opportunity for information leakage.

3.1.3 Federated identity management

Due to the virtualization of certain network services, e.g., deployment of home gateways as VNF, the identification processes will be performed on virtualized servers. We now discuss federated identity.

The federation of user identity (FI) grants users which already have secured access to one domain seamless access to another domain. The objective is to minimize the amount of effort put by a user to administration.

Note that in several use-cases the access is performed *on behalf* of a user. The most widespread example is usage of social plug-ins.

Generally, the overall security level is assumed to increase by FI. This is supposed to be fulfilled by one-time, possibly periodically updated, authentication. The corresponding identity information is then exploited across multiple systems and websites, including third-party websites. Hence, the paradigm of FI is naturally generic, and should suit all technologies and standards.

A representative list of standards and technologies being employed by the FI application includes SAML (Security Assertion Markup Language), OAuth, OpenID, Security Tokens, Web Service Specifications, Microsoft Azure Cloud Services, and Windows Identity Foundation.

3.1.4 CloudID-based solution

The solution can be based on the CloudID platform (M. Haghighat 2015).

The solution objective is to assure that both the cloud provider and any potential third-party attackers are restricted. In both cases, the restriction spreads on two subjects:

1. The sensitive data which is virtually statically stored at the cloud (i.e., the cloud provider itself)
2. The content of any of the potential individual queries.

The solution in (M. Haghighat 2015) suggests creating encrypted search queries, by means of a k-dimensional tree structure (*k-d tree structure)* of the search process in the encrypted data (details can be found in the cited paper).

3.2 Tracing and monitoring

This section deals with monitoring and securing of NFV function on separate VM-level, and on the inter VM level. Most of the VM security was already addressed, however special care should be taken in the NFV context. We mention the following vulnerabilities

- Autostart hooks. Various autostart options implanted by malware, side application, which could be running on the same machine. Such side application can refer to any, seemingly light, network service running in the background. See Example 3.2.1
- Sleeping loops and deliberate delays. Sleeping loops and delays can be imposed on VMs thus deliberately slowing down connections (See Example 3.2.2.

Consider the following examples:

Example 3.2.1

Consider malware application X being autostarted each time VM a.b.c.d is rebooted, and/or new VNF is deployed. The application takes the shares of the machine cycle. Hence, even if it cannot harm any VNFs concurrently deployed on the same VM, it degrades its performance.

Example 3.2.2 DoS attack in the context of NFV - *Consider web server deployed as VNF. Now consider a dummy connection which opens TCP port and connection with minimal MSS and minimal CWND. The connection will be particularly slow because of these MSS and CWND constraints. Now, in order to slow it down further, it will delay packets by activating exponential back-off on the other side, and will send a single MSS to avoid closing the connection. Engaging multiple connections of such a type can potentially harm networking functionality of the other side.*

Designated Monitoring Centres (DMC)

The directions to the solutions are as follows:

Virtualized and specialized malware protection should be added to each VM. Logging should be intensively practiced. This includes both frequent snapshot and also logging of the movements. That is, in case VM moves from HW A to HW B and then to HW C, these movements should be logged.

Specific monitoring centres should be deployed on cloud owned HW.

Additional solution: Providing full packages of VNF as a Service (VNFaaS) which handle/incorporate security.

3.3 Policy compliance violations

Each network operator has a set of policies that its network configuration must obey. Example policies include isolation client traffic from operator traffic, isolating tenants renting processing from the operator, ensuring all traffic from clients is NATed and firewalled, etc. In a dynamic 5G network, where network configuration changes constantly due to the instantiation and removal of network functions, how can the operator make sure its policies always hold?

Checking policy compliance is a well-studied field, with many solutions that fall in two categories: active testing or static analysis.

Active testing means injecting packets at many vantage points and verifying that the outcome obeys the policy. This solution is simple to implement, however it cannot guarantee policy compliance: random test packet generation only covers a small part of the possible packets, and correctness must be proven given all possible packets. To increase coverage, tools like ATPG [KazemianConext2012] or BUZZ [SekarNSDI2016] choose test packets after they run static analysis on the network configuration and figure out "classes" of packets.

All static analysis tools (including HSA, SymNet, Veriflow, NOD, etc) use a snapshot of the data plane forwarding state and a model of the processing performed by each box. In this model, a generic packet is injected and different properties are checked such as reachability and loop freeness. The big advantage of static analysis over testing is that it can give much better coverage of the possible set of packets and thus better guarantees of policy compliance. One downside, however, is that generating models to use with static analysis is a hard thing to do.

In Superfluidity we will use the Symnet symbolic execution framework to test policy compliance. Symnet has been developed in the previous Trilogy2 FP7 project and Superfluidity, and we will continue developing it and applying it to real networks. Symnet works on models written in SEFL, a specialized language that has been optimized to ensure fast symbolic execution. We have created SEFL models for routers, switches and many other boxes (see [StoenescuSigcomm16]). We are currently in the process of using Symnet to ensure policy compliance of Openstack tenant configurations ([StoenescuLanman2016]).

3.4 Defending against memory exploits

Defending against memory exploits for network functions written in C can be done by using traditional tools including Address Space Randomization (ASLR), Data Execution Protection (DEP) and write integrity protection (stack canary values, WIT [CadarOakland2008]). On top of this basic protection, fine grained virtualization of network functions can further limit the damage of faulty NF can incur to the rest of the network.

In Superfluidity, we take a different approach to achieving memory safety. We will develop an approach to automatically generate C implementations from SEFL models that are optimized for symbolic execution. This work has two distinct parts:

- a framework that enables equivalence-preserving optimizations on SEFL models to allow successive optimizations that are faster to run
- a translation from SEFL to C that provably preserves the memory safety properties.

This work is just starting.

3.5 Performance and scalability of secured NFV

Virtualization of multiple services will open new security problems, which should be effectively treated. However, the implementation of new secure privacy-related and authentication functionalities should be aware of performance demand of the corresponding VNFs. In particular, deployment of VNFs, jointly with their additional security assuring components, should be still cost effective for the enterprise.

We explain both performance and scalability issues in firewall Example 3.3.1.

Example 3.3.1 (Scalability of virtualization: HW based Firewall)

Consider a legacy firewall, the throughput of the designated device is high; still it is not virtually unlimited. Another advantage of the device is low delay. This is not the case if the firewall is deployed as NFV. The virtualized firewall will be able to support high load amounts and will not be vulnerable to intermittent load peaks.

However, this poses new challenges pertaining to the concurrent secure and effective functioning.

First, the delay will increase. The reason for this can be also attributed to the additional networking which is added to the previous functioning.

Second, the new networking linking the enterprise with the virtualized firewall should be implemented in a secured manner, hence the packets which traverse the (virtual) link should be additionally secured. This may add additional overhead which can affect the performance.

Finally, the newly virtualized firewall should run on isolated and secured VM.

3.6 Availability and feasibility

In the case of VNFaaS, availability of a sufficient number of deployable instances should be assured.

Shortage of VNFs, or lack of scalability can be abused by a potential attacker. An additional vulnerability is the inability to effectively access the secured storage, or the storage that contains sensitive data.

Denial of service attacks can also hamper the availability, as explained next.

3.6.1 Denial of service attack and tracking attack in NFV environment

Consider the identification server (IS) which stores identity details. The potential attacker could deploy an attacking system, which comprises virtualized switches and routers, virtualized DHCP server, that would produce differentiated queries to be sent to the IS. In particular, all queries will be seemingly different and will be hard to see if coming from the single identity by the IS. The virtualized structure of the IS and adjacent network functions could be deliberately misused by potential attacker, see e.g., (Wueest 2014). The objective of such an attack could be two-fold:

- Sending multiple queries in order to flood the IS, thus causing a denial-of-service (DOS) attack.
- Tracking of queries.

In contrast to the case where the IS resides in the user's static private premises behind a firewall, the virtualized IS is in the neighbourhood of the public resources. In order for a DoS attack to succeed, special planning from the side of the attacker should be applied, yet it is theoretically possible. Hence, IS should be employed together with defending envelope, including virtualized firewall, with specially tailored routing rules.

The tracking attack is designated to follow different types of queries in order to extract any kind of side information about both the initiator(s) of the queries and about the confidential identification data stored at the IS. It is not clear how this threat can be alleviated.

3.7 Survivability and damage isolation

Once network services are virtualized using the cloud's infrastructure, the responsibility for disaster treatment is taken out from the domain of the enterprise. Typically, the recovery issues should be a part of the SLA. Hence, in case the enterprise has well-defined strict disaster recovery demands, the deployment of the VNF should be performed on the suitable NFVI nodes, which can address these demands. In this case, it is an enterprise's responsibility to assure this availability. Otherwise, in case NFVSaaS is provided, special recovery procedures should be ready and presented to the NFV user.

Normally, the disaster is attributed to one or more than one of the following events (which can be triggered by natural causes, e.g., fire, natural disaster, intended physical attack, etc.) have occurred:

- CPU fallout - Disaster of local dimensions, which lead to stoppage of several deployed services.
- Electricity down - Usually one specific geographical point is affected. Causes all deployed VNF to be stopped. Causes stoppage of all services incoming /outgoing to the points located on the site. Hence VNFs deployed on other sites can no longer access.
- Disk failure - This issue become especially dangerous then the disk serves for authentication of virtually deployed resources. In addition to the damage caused by the stoppage of the


service and obscurity of how to handle the existing user's authenticated sessions, the event can be maliciously exploited for the additional attacks, thus providing the vulnerability in the face of the "second-wave" disaster.

- Disrupted link - In the case where a physical line is damaged all associated data transfers with this link are stopped.

Note that we excluded from the list above "disasters" which happen due to extreme system overload. This is because we assume that the load, load balancing and QoS provisioning is taken care of at other management levels. Hence, load peaks should not lead to the events associated with disasters.

The basic requirements for successful disaster recovery are as follows:

- Understand the dimensions of the disaster- The cloud provider should realize what are the damaged/infected/harmed nodes, what are the locations of the affected sites . It should be made clear what are the geographical coordinates of the CPUs and disks which are involved. In the special case of the networking layer collapse, the location of the damaged links, lines, optical fibres should be understood.
- Isolation of the damaged location - It is crucial to prevent any leakage of the damage to other sites.
- User acknowledgment - The enterprise should immediately understand the magnitude of the disaster and to receive the recovery prognosis. The NFV service provider can offer alternative solutions, but this information will allow the customer enterprise to make effective decisions on its own.
- Log and snapshot - This phase should be performed concurrently with the temporal (substitute) deployment. The logged data should be used in order to restore the functioning. Special care and security measures should be taken, because any potential attacker currently *knows* that logs will be taken out of the log servers.
- Temporary deployment/solution - NFVs should be offered backup infrastructure in order to revive the functioning as soon as possible.
- Repair - The actual elimination of the damage.
- Report and analysis - Conclusions should be drawn for future prevention.

4 Security analysis of SUPERFLUIDITY use-cases

We now analyse several use-cases for security issues. For each use case we bring first the general description and the methods of the virtualization, next we mention possible security issues which arise when the service is deployed as VNF, and finally we provide possible solution directions.

4.1 5G RAN - Network slicing

General description

5G run will support split of micro-services, will be dynamically scalable in both ways.

Slices should provide networking allocations for various services. (Each slice for different service type.) The services should be dynamically allocated on the cloud. Hence, this is a heavy networking resources exploiting service. The virtualization would provide the elasticity and scalability of the slicing resolution. The service is large and complex, hence can be logically partitioned into (not necessarily disjoint) sub-services.

Security challenges

As a networking service - the virtualization of such paradigm will involve profound security control which will involve treatment of all risks described above.

Possible solution direction

We mention here general security issues concerning the data flows. The sub-modules of this service block (control, load balancing) are discussed separately below. Secure slicing service should be provided. The challenge is to provide security without affecting the QoS resolution and scalability, to supply secure control plane for the various networking slices and to concurrently provide the basic radio resource management (MAC services).

4.1.1 Wireless Software Defined fronthauling (WSDF)

General description-WSDF

Provides the control plan for the 5G - RAN, and actually acts as a main managing unit of the 5G RAN. It controls the capacity allocation, opening and closing of new resources, reallocation of provided network slices. As such, it operates with resources such as Ethernet ports, SDN agents, Load balancing functionalities, Interface and synchronizations with and among micro services. In addition, once deployed as a complex set of VNFs, this component will be responsible for the effective virtualization. This will involve virtual resources allocation and further orchestration

Security challenges - WSDF

All attacks mentioned in this document can be a potential threat for this complex control system. Hence, the design of the control plane should be based on the security needs.

Possible solutions - WSDF

In order to cover multiple security issues associated with the control plane the best solution is conceived in the design of security-aware service. This will include integration of the fronthaul network in a control plane separately deployed (see (A. de la Oliva 2015) proposing such an idea) on the private cloud. Namely, this is suggested to be implemented in the following layers.

- Module-based security. The system will be deployed as a hierarchical structure of services controlled by a main orchestrator. This orchestrator will be located in secure cloud and connected to secondary orchestrators by secured links. See Figure 3 for the possible deployment scheme.
- Secondary orchestrators are also advised to be located in the secure cloud. The main reason for this secure structure is to hide the most sensitive parts from all potential threats. Hence, once the entire cloud RAN fronthauling centre is virtualized, it is proposed that secured HW will be used.

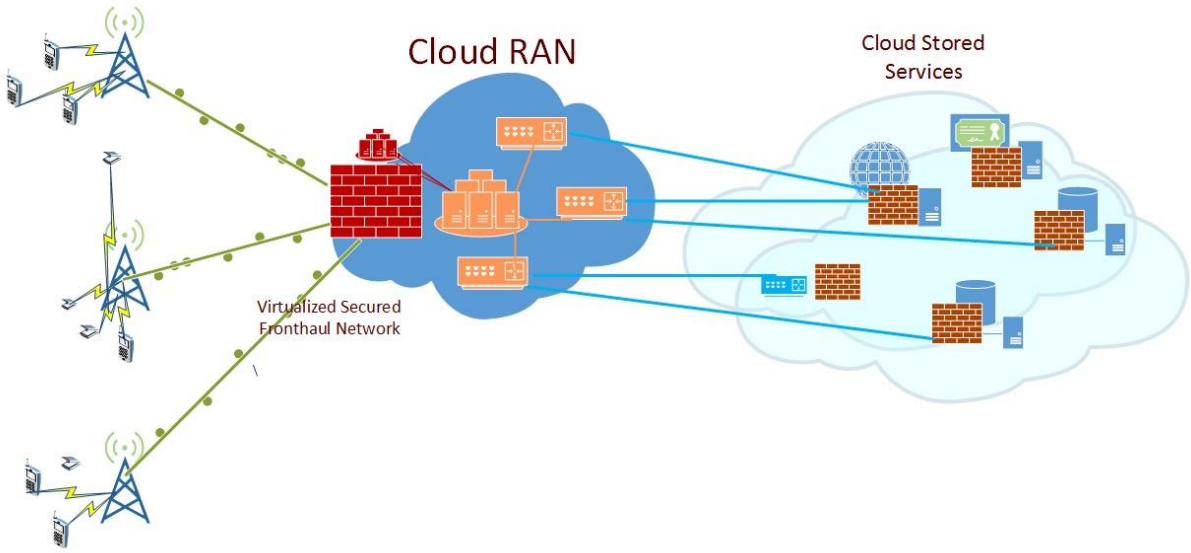

Figure 3: Deployment of virtualized cloud RAN. The main fronthaul layer is mainly attributed to the control interface between the cloud-RAN and users. It is proposed to be deployed at separate secured cloud domain. Hence, the cloud RAN will be deployed in the private cloud which will comprise dedicated security measures specifically applied to this designated domain.

Note that RAN constitutes a set of services with complex structure of various immediately responsive subservices. Hence, in order to be efficiently operated and concurrently secured, it should be *designed* taking into account the aforementioned security demands.

4.1.2 Dynamic MAC services allocation in Cloud RAN

General description - Dynamic MAC services

This service will cope with dynamicity of the wireless networks by addressing NFV resource management in the face of load changes, handovers, rapid movements, and by providing virtualized load balancing.

Security challenges - Dynamic MAC services

Load-balancing operations are vulnerable to attacks both on transport and on message level. The transport level, if load balancer (LB) is virtualised, presumes deployment of the *attached firewall*. See Microsoft implementation guidelines (Microsoft 2011)

Hence, *virtualised* firewall will be attached to each instance of radio LB.

In the case of multiple instances and in the case of certain users, messaging between the LB and user or among various LBs should be enabled. Hence, these messages should be delivered by means of a secured channel and properly encrypted.

4.2 Virtualization of home environment- Virtual home gateways (VHG)

General description - VHG

Aims to move traditional functions (e.g., firewall, parental control, NAT, etc.) residing at the customers' home to a virtual home gateway in the cloud. Allows deploying new services, or upgrading them without changing anything on customers' equipment.

Security challenges - VHG

The main concern in this VNF is privacy and vulnerabilities associated with identity, federated identity and authentication. Additional concern is related to intrusions of all kinds. For example, users which experience malware effect can potentially harm other users by spreading the malware via virtualized centre.

4.3 Local breakout - 3GPP optimization

General description - 3GPP optimization

This service's purpose is to apply to situations where communicating sides are attached to the same edge of the network, i.e., geographically close. In this case, it is desirable that the traffic can flow between them directly, not going to the mobile core, which is today the default behaviour.

Security challenges - 3GPP optimization

The main security concern here is disrupting the correct functioning of this service. In this case, the opposite effect can be achieved, once closed users will communicate via remote site.

Therefore, the implementation of this VNF should be performed with security-aware insight.

The servers or orchestrators which will be responsible for the connectivity should be properly firewalled and the messages between them should be provisioned by means of secured links and enhanced by encryption. This service may also suffer from the effect of a malicious user that intentionally reports false location.

4.4 Virtualized Internet of things- VIoT

General description - VIoT

Virtualization of Machine 2 Machine communication, control center of "things", IoT Virtual Network.

In this category, we also include Virtualized wireless sensors network (VWSN), and hierarchical WSN which in general is described as a lightweight application service, demanding fast control, with multiple low bandwidth communication links.

Security challenges - VIoT

The effective management of IoT devices will be provided by specialized services deployed on the cloud as dedicated IoT managing VNFs. Hence the actual control and orchestration of IoT devices and units will be implemented on virtualized manner as a dedicated IoT VNF.

The implementation of control of IoT is vulnerable to various potential attacks, on the cloud.

Each IoT instance, whether virtualized or not, incorporates inherently security-related hazards mainly related to the privacy of the corresponding enterprises. Therefore, the elimination of all possible threats should be treated within the framework of secured identification and authentication mechanisms, specifically designed for the NFV demands.

4.5 Virtual CDN for TV content distribution

General description - CDN

Assume content cache deployment, close to network edge based on behaviour analysis and forecasts.

Such a deployment is ordinarily associated with virtualised CDN allowing several players to deploy their own CDN, according to their rules and needs.

Note that current Internet contents distribution CDNs are based on traditional content caching algorithms, based on observed content popularity. Hence, these CDNs apply similar distribution rules for all contents.

However, various types of services require different constraints to cache contents. In particular, geographical binding has a crucial impact on latency of content delivery, hence, the content is presumed to be cached in accordance to users' locations, which is especially important for applications like TV.

Usually, the contents of interest are stored in the *edge repository* and become available after their scheduled transmission time. However, after that, they stay available in the edge for only a defined period of time (that is, the interest in certain TV content becomes quite rapidly outdated) after that, they become available only from a central data centre, freeing storage space for other contents.

Security challenges

In the described context, the security will depend on the time-stamp of the content. Hence, we propose to logically divide the content according to one of two possible stages.

1. Is of high interest. This content will be in high demand in the next two days.

2. Archived content. After approximately, e.g., 48 hours, the content loses wide public interest, hence may be moved to the remote, occasionally accessed storage.

Clearly in the first phase, the content will be classified as highly sensitive, thus vulnerable to the various types of attacks, especially DoS attacks. Therefore, CDN for TV content which is deployed as a set of VNFs will need special attention of security enhanced services during the first stage.

After that stage, the security effort can be reduced to the normal standards.

4.6 Pure security service and their virtualization

In this subsection we mention virtualized services which will augment other services by providing additional security and trust application.

4.6.1 Preventing NDP (Neighbour Discovery Protocol) spoofing

NDP is an IPv6 substitution of an ARP. It also performs IND (Inverse Neighbour Discovery) which acts similarly to RARP. Unlike ARP, it does not necessarily act on pure link layer, e.g., Ethernet. Namely, NDP acts over ICMPv6, hence it sends IP level packets. This creates the inherent vulnerability and the threat of NDP spoofing.

In order to neutralize the threat all virtualized networks should secure Neighbour Discovery (SEND) Protocol. SEND prevents a potential attacker from accessing the broadcast segment, thus preventing abusing NDPs to trick hosted VMs into sending the attacker any traffic designated to someone else. This technique is also known as ARP poisoning. SEND uses RSA encryption in order to assure security.

This practice should be specified as obligatory for secured and sensitive VNFs.

4.6.2 Protection against DDoS (Distributed Denial of Service) attacks

DDoS is a type of DOS attack where multiple infected systems, (e.g., by Trojan), are exploited to attack a single target in order to inflict the Denial of Service (DoS).

DDoS is harder to implement from the side of a potential attacker but likewise is harder to overcome.

A cleverly organized DDoS can be disguised as a common load peak situation and can thus affect hosts, hypervisors, load balancers and other virtualized objects. It can be activated from various locations, by hiding botnets within larger cloud-deployed applications, neighbouring VNFs or even enterprises. Normally, the side(s) which participated in the attack are unaware of this and also considered as victims

The most effective measure against DDoS is prevention. In particular, as long as multiple logical links would be involved in the attack, with multiple VMs, some of them possibly belonging to the same VNF, deployment should be performed with awareness of possibility of such an attack. Hence, global defending mechanism which will prevent multiple VMs from being infected by such a malware,

combined with Firewall(s) protecting from malicious flows from outside should serve as a security prototype. This prototype should be part of the VNF's *initial design*, rather than application of patches in the future. As such, it will be more effective and cost-effective in the future as well.

4.6.3 Emergency treatment and recovery VNF

Emergency recovery applications can be implemented as a instantaneously deployable VNF, which will act in the case of global disaster, down, multiple VM collapse, etc. In order to be successfully activated, it should be hosted on several hypervisors and be active in sleep mode.

For greater security, it is suggested to separately occupy a dedicated HW device. The emergency VNF (EVNF) will be implemented as self-orchestrated and self-deployable. Its tasks will include the following functionalities:

- EVNF will be automatically triggered once certain conditions will conform to the situation of system collapse.
- Detection of damaged sites by simple probing (e.g., pinging). Once probing results are ready it will send them within a designated packet to a set of manually preconfigured addresses.
- The EVNF will have access to logs and snapshots. Using this data will facilitate deployment of hypervisors and orchestrators which will serve as alternatives to those that are currently out of function.
- EVNF will be immune to all second-wave attacks by rejecting all incoming data which is not expected and from addresses which are not expected to send any packets. This will ensure any further attacking and abuse of the EVNF
- EVNF will send report messages to the enterprises with summary of the infected HW and stopped SW in the environment which is within its responsibility area.

5 Conclusion

We reviewed possible threats and corresponding security related challenges in the NFV world. NFV applications and services can be seen as combinations of cloud computing related and virtualization related objects. Hence, most of the threats are known to the technology industry from these two areas.

Nevertheless, due to the prospect of previously naturally secluded application families being deployed into the public environment, and thus being exposed to neighbouring HW and SW structures, the security challenges should be reviewed and new threats should be defined.

In some cases, VNFs should be implemented in symbiosis with accompanying security functions. In other cases, the existing services could be deployed and merely side-guarded by dedicated security functions, possibly by specially tailored separate VNFs.

We also mentioned the pure security services, which would be deployed as VNFs. In this case, these services should be appropriately revised and enhanced in order not to become a target by their own newly revealed weakness.

Finally, we conclude that in the case of large and complex services being virtualized and deployed as an entire set of various VNFs (as, for example, the 5G RAN), the deployment should be designed and implemented in a security-aware fashion. Here, inherent and pre-integrated security awareness will both provide proper defence against any potential attacks and assure performance scalability and optimization in the face of this new deployment and demands. This approach will demand initial planning investment but will prove to be much more valuable and cost-effective than later application of patches which may soon after be outdated.

6 References

A. de la Oliva, X.C. Perez, A. Azcorra, A. Di Giglio, F. Cavaliere, D. Tiegelbekkersk, J. Lessmann, T. Haustein, A. Mourad, P. Iovanna. 2015. "Xhaul: The 5G integrated fronthaul/backhaul." *IEEE Wireless Communications* 32-40.

E. Pagnin, C. Dimitrakakis, A. Abidin, and A. Mitrokotsa. 2014. "On the leakage of information in biometric authentication." *Progress in Cryptology - INDOCRYPT.* Springer. 265-280.

M. Haghighat, S. Zonouz, and M. Abdel-Mottaleb. 2015. "CloudId: Trustworthy cloud-based and cross-enterprise biometric identification." *Expert Systems with Applications* 42 (21): 7905-7916.

Microsoft. 2011. *Things to consider when implementing a load balancer with WCF.* https://msdn.microsoft.com/en-us/library/hh273122%28v=vs.100%29.aspx.

Wueest, C. 2014. "Threats to virtual environments." Symantec Research.