

Experiences in Trusted Cloud Computing

Ian Oliver^(✉), Silke Holtmanns, Yoan Miche, Shankar Lal,
Leo Hippeläinen, Apao Kalliola, and Sowmya Ravidas

Security Research Group, Nokia Bell Labs, Espoo, Finland
{ian.oliver, silke.holtmanns, yoan.miche,
shankar.lal, leo.hippelainen, aapo.kalliola,
sowmya.ravidas}@nokia-bell-labs.com

Abstract. While trusted computing is a well-known technology, its role has been relatively limited in scope and typically limited to single machines. The advent of cloud computing, its role as critical infrastructure and the requirement for trust between the users of computing resources combines to form a perfect environment for trusted and high-integrity computing. Indeed, the use of trusted computing is an enabling technology over nearly all ‘cyber’ areas: secure supply chain management, privacy and critical data protection, data sovereignty, cyber defense, legal etc. To achieve this, we must fundamentally redefine what we mean by trusted and high-integrity computing. We are required to go beyond boot-time trust and rethink notions of run-time trust, partial trust, how systems are constructed, the trust between management and operations, compute and storage infrastructure and the dynamic provisioning of services by external parties. While attestation technologies, so-called run-time trust and virtualized TPM are being brought to the fore, adopting these does not solve any of the fundamental problems of trust in the cloud.

Keywords: NFV · Trusted computing · Security · Telecommunications · Cloud

1 Introduction

The telecommunication cloud, colloquially known as Telco Cloud, is a fast-growing area of development for telecommunication infrastructure companies. Many telecommunication functions are and will be deployed in virtualized forms. Such functions known as VNFs - virtualized network functions – range from firewalls and routers to more esoteric systems such as the HLR and VLR etc. (Home Location and Visitor Location Registers) supporting the mobile networks and even base stations and software components of antenna systems and the radio network.

This shift in deployment is primarily due to the additional flexibility in terms of functionality, scalability and cost that the cloud provides [1]. This is especially true in terms of provisioning new equipment which would have been hardware based in the past and now is little more than spinning up a number of new VNF instances for that network function as AT&T pointed out in [5].

Given this flexibility and the fact that Telco Cloud systems are effectively mission critical systems the security of such systems is paramount. However, security is a broad

term and encompasses many areas. One area that has not been addressed is how the integrity of the system is ensured overall. That is, how can we be sure that the VNFs being loaded and launched have not been tampered with; similarly this extends to the actual Telco Cloud itself and its Management And Operations (MANO).

ETSI have defined a reference architecture as shown in Fig. 1 with the system being split into 3 major parts: the MANO, NFVI and VNF layers (OSS/BSS - operating/business systems support is not considered here). These layers are conceptual descriptions of elements and should not be confused with physical architecture. It is quite likely that MANO components be provided within the NFVI as well as VNFs in their own right. Similarly some network functions may extend over the NFVI and VNF layers, e.g.: the physical-software combination on antennas for example.

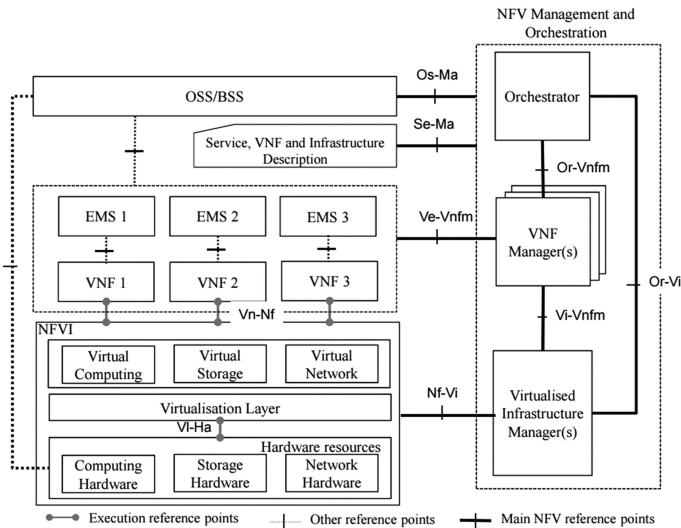


Fig. 1. ETSI reference architecture framework [3]

Unlike traditional cloud based systems, trusting the type, status and integrity of the hardware and platform provisioning, as well as the systems being built upon that is critical.

Establishing trust in the Telco Cloud is complex but necessary and comes with a series of additional challenges which do not occur in either traditional data centers or cloud systems. The challenges of incorporating trust in this environment has not been dealt with in detail in the current literature nor in any current product offerings.

In this paper we present a number of critical definitions when working with trusted cloud and VNFs, how attestation and signing can be utilized in a dynamic service delivery scenario and a range of outstanding problems that need to be addressed before anyone can claim that they are running a trusted environment.

2 Background

In [13] cloud security is considered as one of the main topic areas of research and trusting the infrastructure as one of the important challenges faced by the cloud users: Integrity, confidentiality and auditability proofs to the service provider for ensuring secure data transfer and state of the system are critical. Explicitly stated in the above paper are the requirements for trusted hardware and a trusted virtualization layer. It will become obvious here that this is just one part of the overall system that requires such trust.

In [7] a set of possible attacks is listed specifically pertaining to the tampering of a cloud environment and its virtualized workload. The attacks presented mostly deal with attacking the Virtual Machines (VMs) such as capturing VM snapshots, analyzing memory dumps of VM and attacks performed on VM migration. The authors also list the possibility of circumventing the current protections in the cloud environment; however, they do not propose any solution or mitigation for the specified attacks. Here, though, it can be seen that integrity protection is the required mechanism.

In [10] are presented the challenges and the requirements that emerging technologies need to satisfy, in order to establish trust in cloud, specifically platform integrity. They additionally present certification of the cloud and require mechanisms for establishing trust in cloud.

In [12] the authors identify the important key problem of the lack of trust architecture for Network Functional Virtualization (NFV - Telco Cloud); specifically:

- Ensuring NFVI security against intrusion attacks and possible countermeasures.
- Providing security services/functions in an efficient and economical way.
- Provide VNFs based on NFVI in a trustworthy way; especially in scenarios where multiple vendors uses same underlying infrastructure.
- Establishing trust on VNF-VNF communication.

The European Telecommunications Standards Institute (ETSI) has provided a white paper [2, 8] on NFV which further presents challenges associated with NFV, such as security and resilience:

- Establishing trust in the platform or NFVI: The goal is to verify that the platform is in an expected state.
- Establishing trust in software, policies and processes, including VNF, MANO and other NFV components
- Supplying guidance for operational environment such as MANO and Element Management System (EMS)
- Defining trust relationships between virtualization resources for trust life cycle management.

In summary, all of the works presented emphasize providing trust in the platform components - the layer known as the Network Function Virtualization Infrastructure (NFVI) and then specifically only on the hardware, operating system and hypervisor components.

No work known to the authors at this time addresses specifically the problems relating to the integrity of the VNF and MANO components, which encompasses both integrity and confidentiality of these.

Furthermore introducing trust is not just a security issue but also one of component identity and of system resources. This latter case then implying that the Telco Cloud needs to additionally manage itself the workload according to safety-critical and fault-tolerant principles.

3 Establishing NFVI and VNF Integrity

The NFVI consists of the hardware, operating system and the virtualization layer. Providing a trusted NFVI, at least in the single physical machine case is a relatively simple (and solved) task.

3.1 Single Machine Trust

A single machine trust is provided with a trusted platform module (TPM) chip that stores keys, certificates and other confidential data as well as the cryptographic hashes of selected system components.

The TPM can be used during the platform boot time to achieve platform trust. The TPM contains the platform configuration registers (PCRs) which stores the cryptographic hash measurements of software components such as the BIOS, boot loader, OS and hypervisor etc. The Core Root of Trust Measurement (CRTM) is achieved by enabling preceding boot components to measure the following boot component to form the chain of trust [6, 8, 9, 11]. During a trusted boot, these components can be measured and verified against the good known values specified in TPM's Launch Control Policy (LCP). If the trust chain is broken, then the system can be halted or can be started according to LCP specified by the platform admin. Launch control policies are the list of policies that verifies if the system meets the required criteria and further decides if the platform has to be launched or not.

In a cloud environment any failure during the boot sequence can result in a number of situations that need to be handled by the MANO:

- failure of the machine to start at all
- machine entering a safe-mode (and possibly reporting this to the MANO)
- machine continuing boot regardless of the integrity measurements (not recommended!)

3.2 Multiple Machine Trust

If multiple machines are provided, those with TPMs start as if each were a single machine. The use of an attestation service, commonly called remote attestation to highlight this service's independence, is required to monitor the integrity state across all machines.

The attestation service in this case simply queries each machine’s TPM (at least for those that have TPMs and those that have booted), fetches the TPM’s PCR (Platform Configuration Register) values and compares them against good known values stored in its database. The Attestation server provides the result back to the verifier, on whether the host is trusted or not as shown in Fig. 2.

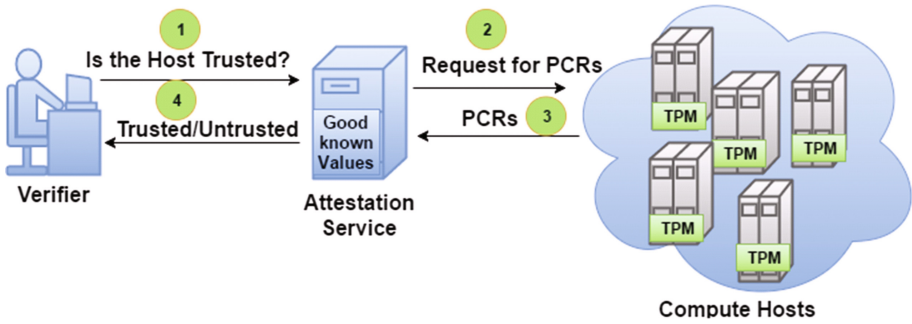


Fig. 2. Workings of a remote attestation server

Using the attestation server in this mode we can also see the lifetime of the integrity checking process in Fig. 3. Here we see that once the machine is running successfully, or is at least in a state where it can respond to the attestation server it can answer any request made [14].

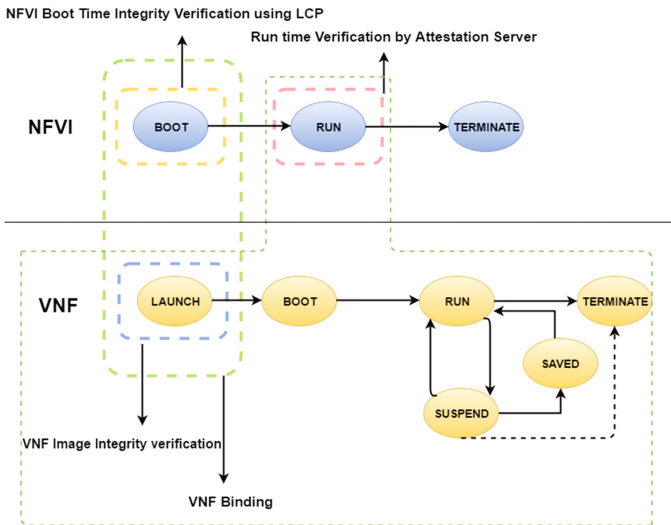


Fig. 3. “Run-Time” attestation of an NFVI element timeline

The problem here however is that while this is known as “run-time” attestation, it is little more than occasional polling and re-measurement of already known and measured components. While this method has weaknesses, it suffices as a rudimentary mechanism for ensuring trust during the running time of those machines [4].

Therefore the concept of a “trusted cloud” is better defined as one where there exists at least one trusted NFVI element capable of running trusted workload, ostensibly as a trusted virtual machine/virtualized network function.

3.3 Trust Failure

One thing that should now be noted is that under the circumstances where trust cannot be achieved, or, where the provision of trusted resources “fails”, i.e.: becomes unavailable, needs to be handled.

We can divide workload into three categories:

- Those that do not require trusted resources
- Those that should have trusted resources
- Those that must have trusted resources

The latter two categories we define as soft and hard-trusted respectively.

Under NFV element failure, it is typical that VMs are migrated to other working machines – resources permitting. If a trusted workload requires migration then a suitably trusted NFV element must be found.

If such a resource is not found then in the case of a hard-trusted VM the VM is simply terminated (in a safe and secure manner). A soft-trusted VM can be migrated if suitable mitigations can be put in place to ensure the integrity of the workload and the platform on which it is running. This might entail usage of network slicing or VNF wrapping to protect and isolate that workload.

Of course such a situation will inevitably introduce more load to the system and additional latency. However we would argue that under such a situation preserving the service might be more important than preserving all of the service level agreements.

4 Establishing VNF Integrity

We propose using an attestation server to remotely verify the platform trust state along with a security orchestration component (known as TSecO) nominally implemented within MANO to perform VNF related security operations. Figure 4 shows the placement of this security orchestration (and attestation) within the MANO context. These operations include:

- verifying VNF image integrity before launch
- binding some VNF to a certain NFVI

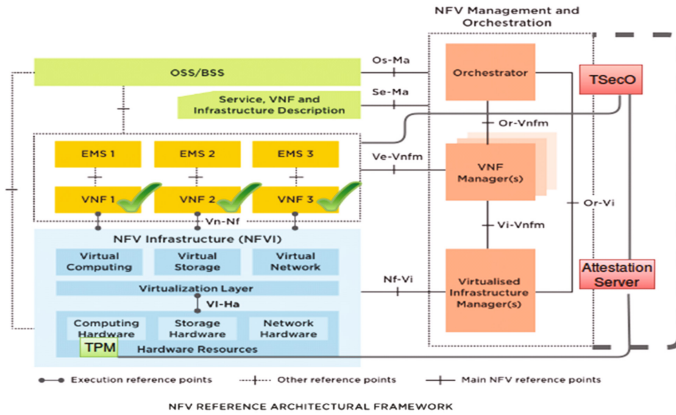


Fig. 4. NFV architecture with security orchestration and attestation server explicitly Shown [3]

VNF binding can be useful for cases such as binding VNFs to platforms which reside in certain geographical location to comply with data sovereignty or with lawful intercept regulations.

Verifying the integrity of VNF images during their launch time is crucial to enhance trust in NFV. Consider an attack scenario where the VNF image database in NFVI is hacked by an attacker (internal or external), or that a VNF in transit over public (or even private networks) is tampered with.

In such situation, it is essential to detect these attacks and provide the guarantee of VNF integrity to the service providers. In our approach, we propose a method to assure the VNF image integrity. In this method, hash digest of VNF image should be calculated first and signed by a signing authority.

The signing authority will generate the signature file which can be stored in the TSecO database. The Virtual Infrastructure Manager (VIM) in the MANO stack, such as OpenStack, can be modified to send the VNF launch requests to TSecO before launching the VNF. The launch request should contain the fresh hash digest of the VNF image and VNF identifier. TSecO should verify the validity of this signature using the hash digest, and sends the response back to the VIM whether it should proceed to launch the VNF or not. Such a mechanism can avoid launching of tampered VNFs and also detects if there has been any unauthorized modifications made in the VNF image. We propose that signature verification should be performed externally and preferably by involving a trusted third party as it prevents from malicious administrators intentionally tampering the VNF images [15].

It should be noted that OpenStack already provides for a hashing mechanism for any VNFs stored in its system. However this mechanism is internal and would be one area that is a target for a compromise in the first place.

The solution we present above relies upon addressing the virtual machine launch mechanism within the hypervisor. Nominally we should be able to state that if a machine is trusted then this mechanism itself should not have been tampered with. In order to ensure this and also to address hashing calculations, we can further use

processor extensions such as Intel's SGX or ARM's TrustZone to protect this processing and utilize keys from the TPM to encrypt and decrypt critical code as necessary. The overhead of this lies mainly in the complexity of code and not in the temporal overhead which is dominated by transfer times of multi-gigabyte VM images.

Further, the already known mechanism of information hiding - decryption of specific files using the keys present in the TPM - can be used in this case and indeed it is recommended that VNFs are shipped in an encrypted form. The disadvantage here is that this now complicates the management of keys in any key management infrastructure.

5 Details to Establish VNF Integrity

5.1 TPM Binding

Consider the case where the service providers want their VNFs to run only on systems with certain platform configurations such as preferred operating system (OS) and Hypervisor type etc. This is particularly true in Telco Cloud where hardware optimizations are necessary depending on the cloud workload, for example, lawful intercept requires certain provisions to be made, including geographical trust.

Binding or pinning VNFs to certain NFV platforms require policies that should be satisfied for the binding to be successful. The policies would contain the platform configurations that the NFV platform must possess in order to launch the VNF. In order to solve the challenge of VNF binding, we need to address the following:

1. Determine if a VNF requires binding.
2. How to retrieve the platform configuration state of NFV platform.
3. Implementing the binding mechanism with associated policies.
4. A mechanism to verify the binding rules before launching any VNF.

We propose another approach to solve above challenges. In this approach, each trusted NFV platform should register its PCR hash measurement values to an attestation server. NFVI VIM should be modified to send the VNF launch requests to TSecO, which include the VNF identifier and binding policy regarding the destination host. TSecO should fetch the PCR measurements of NFV platform from the attestation server and verify the binding policy to find the destination host. The response should be sent back to VIM containing the destination host ID where VNF should be launched.

5.2 VNF Snapshotting

VNF instances running in one Telco Cloud might need to be migrated to another cloud infrastructure due to the various reasons such as disaster recovery, high availability and fault tolerance etc. Therefore, it is necessary to create the snapshot image of the running VNF instance which would contain all the running software code loaded into the VNF memory. Now the original VNF image from which the VNF instance was launched, is modified hence the signature verification would fail for this new VNF snapshot. In order to trust this VNF snapshot, it needs to be re-signed and verified in the similar manner as the original VNF image.

5.3 Intra and Inter-NFV Trusted Communication

In current OpenStack implementation, traffic exchanged among NFVIs and also among VNFs, is in plain text. Applications running in VNFs are responsible to encrypt their own traffic. Also, the management traffic exchanged among the NFVI nodes is not secured. The adversary can launch the man in the middle attack by first capturing the traffic, inserting the malicious code and replaying it. Therefore, it is mandatory to imply the mechanisms to encrypt all the inter and intra VNF traffic without VNFs and NFVIs needing to worry about it. This could be performed by modifying the existing networking layer of the OpenStack cloud to incorporate these changes. These mechanisms also introduce new challenges such as identity management and cryptographic key management etc.

6 Trust as an Identity Management Problem

Trust is very much framed as a resource problem, though in many interactions trust is often decided on the identity of the two parties. Protocols such as TLS and SSL are effectively based on the sharing of secrets only known to given pairs.

6.1 Element Identification

Trust can also be expressed as identity problem and this is especially true when looking at trust relationships between elements within the NFV reference architecture. For example, as clouds become more distributed it is not unconceivable that the MANO for one cloud might be running on a totally separate cloud. In this, and even self-contained cases, it will become a necessity to ensure point-to-point and group communications - network, API or other - are trusted in some form.

As TPM already provides mechanisms for unique keys - each TPM has its own unique public/private key pair at manufacturing time - this mechanism can be utilized in the process of establishing the identity of VNFs, NFVI, MANO etc. - each of which have their own groups of identities. We can therefore see this as an extension to existing PKI systems in some form, but also one that includes the hardware root of trust.

Further to this there are a number of areas of potential exploration such as the use of distributed/decentralized and auditable storage of identities and their management - ostensibly blockchain based technology - incorporated with TPM for establishing trust. As an aside here this may also provide a solution to a billing/charging problem at the same time.

6.2 Multiple Roots of Trust

As cloud systems increase in size there becomes the necessity to support more decentralized and failure tolerance mechanisms, especially in MANO which is often framed as being monolithic in nature.

This implies the idea of multiple attestation mechanism which together can give a much better overall trust by reducing the byzantine failure possibilities. The authors are not aware of any work in this area at present.

7 Challenges

The two solutions presented above address trust as a resource management problem in that they attempt to force VNFs as virtual machines to load and start on given systems with certain pre-established properties encoded as cryptographic hashes. Thinking of trust in this manner leaves three major areas to address - these are presented briefly below.

7.1 Load Management

Given that a Telco Cloud will conceivably consist of both machines that have successfully started in a trusted mode and machines that are not required to start in a trusted mode (as opposed to machines that have failed trust), as well as VNFs that require signature verification as well as binding, we are presented with the problem of where to start VNFs.

The default case is that VNFs will be started on the next, most powerful (in terms of CPU and memory resources typically) machine. This can lead to the situation where potentially all trusted NFVI resources are taken by VNFs that do not require trust thus meaning that VNFs that do require trust can not start despite available, although untrusted, resources.

From a traditional point of view this means that resource allocation and balancing now needs to address the likely requirement of running trusted workload. In many cases this mean that workload will not be optimally placed to ensure trusted resources are available.

One solution is always to move or migrate workload not requiring trust away from trusted machines when trusted workload is required, however migration will invariably imply loss of service while migration is taking place.

7.2 Service Resilience

As noted in the earlier section migration of virtual machines (meaning VNFs) is an expensive process and should be avoided. However in cases where a trusted machine fails it becomes necessary to reallocate workload to other machines.

In the presence of trust and additionally in the presence of TPM binding this becomes more difficult in that suitably trusted resources may not be available, even though these resources are trusted in some sense.

Provision of any Service Level Agreement (SLA) is usually paramount, especially in telecommunication systems. In this respect we need to differentiate between levels of trust and decide on a per-VNF basis whether that VNF requires hard, soft or no trust.

Hard trust basically states that if no suitably trusted resource is available then that VNF either does not start, or if already running, is terminated without any change of migration.

Soft trust states that if no suitably trusted resource are available then mitigations can take place until the MANO can reconfigure the NFVI to provide such resources. This admits the SLAs but with some risk - both from a security perspective and an overall system SLA. Mitigations here would include wrapping the VNFs - in terms of SDN reconfiguration and network slicing and isolation and utilizing additional anomaly detection mechanisms. Overall provision of some level of trust becomes an expensive process.

Unknown or non-existent trust can only be accepted when the VNF does not require any trusted resources. A system that freely migrates and allocates VNFs and other workload without respect for trust would be highly dubious.

7.3 Insecurity Through Trust

Given a mixed NFVI environment as described earlier we will see patterns in resource allocation. Knowing which machines are trusted therefore makes these machines more of a target to attackers. For example, if lawful intercept would only occur on trusted machines, then this is knowledge that reduces significantly the attacker's "search space" for suitable targets. Current research suggests this is theoretical, the authors are under no impression that such information would not be used by an attacker for gain.

8 Summary

This paper provides an overview of the challenges in incorporating trust in Telco Cloud/NFV and discusses some of the approaches to address them. We have explained the stages of constructing a trusted Telco Cloud and discussed the challenge of platform trust. Further, we have devised methods such as VNF integrity verification and VNF binding to an NFV platform. We have also looked into the aspects of resource management and meeting the SLA requirements.

Our work opens up new research directions for enhancing the trust in Telco Cloud but also highlights the current naivety and dangers of adopting trusted and high-integrity computing technologies without a full understanding of the implications of said technologies.

Trusted computing address a major and critical area of system security and privacy and it is therefore paramount that this technology be properly conceptualized and implemented in order the gain the advantages it can bestow within a cloud environment.

Acknowledgments. This work was made under the DIMECC Cyber Trust Program (Finland).

References

1. Yang, W., Fung, C.: A survey on security in network function virtualisation. In: IEEE NetSoft Conference and Workshops, pp. 15–19 (2016)
2. Operators, N.: Network functions virtualization, an introduction, benefits, enablers, challenges and call for action. In: SDN and OpenFlow SDN and OpenFlow World Congress (2012)
3. ETSI, G.: Network functions virtualisation (NFV): architectural framework. ETSI GS NFV 2 (2), V1.ETSI (2013)
4. Haldar, V., Chandra, D., Franz, M.: Semantic remote attestation: a virtual machine directed approach to trusted computing. In: USENIX Virtual Machine Research and Technology Symposium (2004)
5. Han, B., Gopalakrishnan, V., Ji, L., Lee, S.: Network function virtualization: Challenges and opportunities for innovations. *IEEE Commun. Mag.* **53**(2), 90–97 (2015)
6. Krautheim, F.J., Phatak, D.S., Sherman, A.T.: Introducing the trusted virtual environment module: a new mechanism for rooting trust in cloud computing. In: Acquisti, A., Smith, S. W., Sadeghi, A.-R. (eds.) *Trust 2010. LNCS*, vol. 6101, pp. 211–227. Springer, Heidelberg (2010). doi:[10.1007/978-3-642-13869-0_14](https://doi.org/10.1007/978-3-642-13869-0_14)
7. Rocha, F., Correia, M.: Lucy in the sky without diamonds: stealing confidential data in the cloud. In: 2011 IEEE/IFIP 41st International Conference Dependable Systems and Networks Workshops (DSN-W), pp. 129–134 (2011)
8. ETSI, G.: NFV Security and Trust Guidance, ETSI GS NFV-SEC 003 V1.1.1 (2014)
9. ETSI, G.: Network Function Virtualization: Trust; Report on Attestation Technologies and Practices for Secure Deployments, ETSI GS NFV SEC 007 V0.0.3 (2015)
10. Khan, K.M., Malluhi, Q.: Establishing trust in cloud computing. *IT Prof.* **12**(5), 20–27 (2010)
11. Stumpf, F., Benz, M., Hermanowski, M., Eckert, C.: An approach to a trustworthy system architecture using virtualization. In: Xiao, B., Yang, Laurence T., Ma, J., Muller-Schloer, C., Hua, Yu. (eds.) *ATC 2007. LNCS*, vol. 4610, pp. 191–202. Springer, Heidelberg (2007). doi:[10.1007/978-3-540-73547-2_21](https://doi.org/10.1007/978-3-540-73547-2_21)
12. Yan, Z., Zhang, P., Vasilakos, A.V.: A security and trust framework for virtualized networks and software-defined networking. *Secur. Commun. Netw.* **9**, 3059–3069 (2015)
13. Zhang, Q., Cheng, L., Boutaba, R.: Cloud computing: state-of-the-art and research challenges. *J. Int. Serv. Appl.* **1**(1), 7–18 (2010)
14. Sowmya, R., Lal, S., Oliver, I., Hippelainen, L.: Incorporating Trust in NFV: Addressing the Challenges, *IEEE ICIN* (2017, to appear)
15. Lal, S., Sowmya, R., Oliver, I., Taleb, T.: Assuring VNF image integrity and host sealing in Telco Cloud, *IEEE ICC* (2017, to appear)