Trust, Security and Privacy through Remote Attestation in 5G and 6G Systems

Ian Oliver
Nokia Bell Labs
Cybersecurity Research
Espoo, Finland
Email: ian.oliver@nokia-bell-labs.com

***Abstract*—Digitalisation of domains such as medical and railway utilising cloud and networking technologies such as 5G and forthcoming 6G systems presents additional security challenges. The establishment of the identity, integrity and provenance of devices, services and other functional components removed a number of attack vectors and addresses a number of so called zero-trust security requirements. The addition of trusted hardware, such as TPM, and related remote attestation integrated with the networking and cloud infrastructure will be necessary requirement.**

I. Introduction

The increasing digitalisation and advantages there of, of systems through the use of cloud, IoT and mobile technologies is well known and documented and needs no further explanation. Similarly the increasing security concerns are well documented [1], [2].

As the ubiquity of digitalisation increases we can no longer sole rely upon perimeter security of systems through the protection of networking traffic, firewalls, PKI etc. Indeed the necessity of security extends both throughout a system and across its lifecycle, especially during the supply-chain phases where are increasingly an integral part of the security story [3], [4], [5], [6].

Technologies such as trusted computing where explicitly designed to address the points of device identity and integrity and have found a particular niche in the protection of the boot and supply-chain of server class machines through the use of the Trusted Platform Module [7] and remote attestation . This relatively inexpensive technology also finds use in smaller devices such as those employed in IoT and mobile infrastructure.

Further we can extend the concept of trust to other elements in a system, for example, container and virtual machine technologies, as long as we can preserve the chain of trust for some core root(s) of trust. This will at least allow us to establish the provenance, identity and integrity of such components across the supply-chain and at run-time with remote attestation.

Once this has been established we can further extend this technology to address other aspects of the system, such as data provenance and privacy, even enabling find grained homomorphic encryption, data notarisation etc.

This paper describes possible implementation and use of trusted computing and remote attestation in cloud and mobile infrastructure, particularly targetting current 5G implementations but also addressing future 6th generation approaches. We present a brief summary of our proof of concept in the medical domain where the integration of 5G slicing, trust properties and remote attestation will bring additional security and are currently being developed to explore these concepts.

The rest of the paper is organised to present a background on the technologies, how they link together and then the vertical use cases, namely, medical and railways, with a discussion on the kinds of learnings found in these domains so far.

II. Background

We present here the three primary technologies that play a role in the development of a trusted, mobile infrastructure, namely: trusted computing, the 5G (and by extension 6G) infrastructure and the cloud computing environments either used to implement or be supported by these.

A. Trusted Computing

Trusted computing defines a suite of mechanisms by which an element - typically a hardware device - has its identity and integrity established by some provable means. Trusted computing can therefore be utilised in the establishment of a verifiable supply-chain as well as the the more traditional role in a device's boot-process.

The canonical example of a trusted (or more strictly a trustable) element is one that contains a hardware root-of-trust such as a Trusted Platform Module (TPM)

Attestation is the process by which this identity and integrity information is verified and validated against some known good values. An element which successfully passes the verification and validation is described as being trusted.

Attestation comes in two forms: local and remote. The former form is performed on-board the device providing the trust as was the case with the TPM 1.2 and its LCP mechanism. The latter form - remote attestation - allows the checking and storage of known good values to be federated making is more suitable for larger and distributed environments as well as having the potential to admit more types of trusted elements, eg: containers, virtual elements and so on.

The definition of trusted is restricted to those whose state is in some sense good, though it does not preclude that a good element contains flaws such as security or other issues. However the important point is that this information is known and can be acted upon. More formally we can described this process as in the diagram below based upon the concepts described in [8]:

$$\begin{array}{l} Element \xrightarrow{measure} Measurement \\ \qquad\qquad\qquad\quad \downarrow attest \\ \qquad\qquad\qquad\quad Claim \\ \qquad\qquad\qquad\quad \downarrow verify \\ \qquad\qquad\qquad\quad Result \xrightarrow{decide} Decision \end{array}$$

In the case of a TPM 2.0 based system - an element in this terminology - the measure process results in measurements stored in the TPM's platform configuration registers (PCRs), a claim is generated through the TPM's quote functionality) and passed to a remote attestation service where the claim's cryptographic signature, device identity (through an attestation key) and contents - attested value, clock and other meta-data - are checked. Based upon the known good values a positive or negative decision will be made. Each of these structures may be composite, for example, an element may be attested in a number of different ways each contributing to the overall result.

Further checks such as matching the TPM's Endorsement Certificate against the TPM manufacturer's certificate are supported as well the ability to store other certificates and

information inside the TPM's NVRAM as necessary. This mechanism is used in the assurance of provenance of the device in the supply-chain. Other routes to verification are also available, for example, the UEFI event log structure or Intel's TXT logs generated by a suitable boot loaded such as tboot can also provide points for cross-reference as well as other pertinent information about the trust of a device or system.

The attestation mechanism is used to generate a claim, an example of which is shown in figure 1. The form of a claim may vary depending upon the element being attested though gernerally this is a TPM 2.0 and the payload of the claim is the aforementioned quote. The interaction between the RAS and the attestee is via a trust agent (TA) which exposed necessary functionality to the RAS. We have implemented an intent based system which attempts to communicate with the attestee and make a request according to any given policy. This allows us freedom to both select the underlying transport protocol but also to add different kinds of attestable element as required.

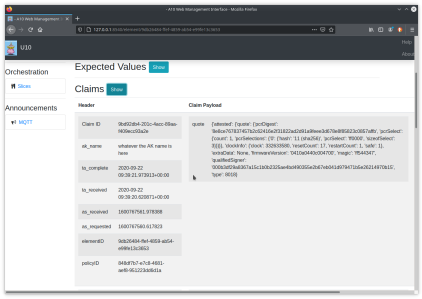

Fig. 1. Example of a Claim

The verification process is handled separately and is made by applying a set of rules to the claim to match it against expected values. An example of a result is shown in figure 2.

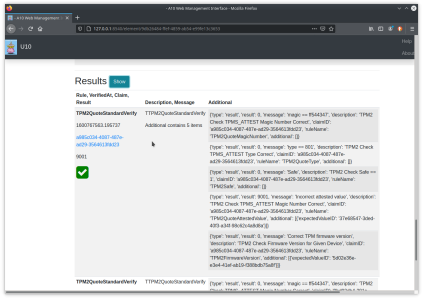

Fig. 2. Example of a Result

Rules for a TPM 2.0 quote are based around verification of the claim signature against the attestation key, the expected values and other provided meta-data. The rules can also be configured to request historical information over past claims and perform more detailed analysis as required.

The RAS is designed NOT to provide a trust decision but rather to provide the results for the individual rules. This effectively presents a report of what has been verified and validated. Trust decisions are typically not binary, especially in the cases where we have to deal with network failure and possible other failure modes of the system.

B. 5G and 6G

The 5th generation telecommunications systems is an evolving set of technologies which extends from the radio technologies through to the infrastructure providing a distributed application platform for operators and businesses. We concentrate here on the the 5G Core and related supporting functionality such as that provided by edge or NFV computing environments [9] or other kinds of orchestration [10], [11].

Specifically the slicing functionality can be utilised not just to provide bandwidth and latency guarantees through a lightweight mechanism for partitioning network and UE combinations but also for the assertion of security properties such as device or element trust. While the current 5G implementations regarding slicing are limited, extensions to this in the 6G specifications which will evolve into the 5G systems, will allow potentially unlimited slices and rapid construction of these slices.

C. Cloud Computing

Cloud computing is now a familiar paradigm for implementation and deployment of applications and services. Edge cloud and its various forms provide a localised distribution model for cloud which takes advantage of locality to provide improved service level agreements such as latency, real-time properties, management, localised processing etc.

The primary security concern from the trusted computing perspective in the cloud has been the trustworthiness of the servers running the system [12]. This is reflected in the typical measurements made as part of the remote attestation process, ie: hardware, firmware, operating system and hypervisor, plus selected run-time files or executables. Establishing these measurements and attestation provides assurance that the cloud infrastructure has not been tampered with.

III. Trust in 5G and 6G

In this section we explore how trusted computing finds a place within the 5G (and always by implication 6G) networks. We will explicitly make a distinction between the cloud and the network infrastructure levels, then explore how these layers then interact and finally explore some implications of this in the actual running of the systems.

A. Trusted Cloud

The simplest form of trusted cloud is establishing trust in the underlying hardware. This is typically made utilising the Trusted Platform Modules and any measurements made during the boot sequence of those machines. The placement of remote attestation services (RAS) is made in the cloud's management and operations infrastructure (MANO). This now affords the opportunity for other MANO elements to interface with the RAS. The links between the RAS and the structures outside of the MANO depend upon the kinds of elements that are able to be attested - if, as is usually the case, these are TPM 2.0 based elements then the attestation endpoint, called a trust agent, is placed on that physical machine with access to the TPM [13].

The main interactions with the RAS come from the VIM and lesser so the NFVI. There are a number of interactions between the VIM and the devices in the NFVI based on booting, rebooting etc, where the RAS would be invoked. Access from the RAS to those individual devices can be made either over a direct interface (Ras-Ta) or proxied by the VIM over the Nf-Vi interface. Similar cases exist for the NFVI-RAS interactions, though with containers especially and with virtual machines, these interactions are managed by the NFVI due to the lack of trusted (or even physical hardware) and the use of virtual TPMs on a per-virtual machine basis for example. One thing to consider here is the maintenance of the chain-of-trust between the Ve and NFVI layers [14], [15]

We have explored the following aspects of trusted cloud:

- Device (Server) Trust
- VNF/Container Workload placement

- VNF and Container Trust

Workload placement is described in [16] and can be realised in a number of ways. The simplest is just to ensure that all devices in the cloud are trusted (at least since previous reboot) and those elements which are not are excluded until their trustworthiness can be established or restored. The more complex mechanisms involve policies regarding the kind of trusted environment required and the level of trust required, for example, specific operating system and firmware environments as established by their TPM measurements.

VNF and container trust are more complex as the they involve both the image and then instances of those elements. We simplify this to just image and instance: specifics on VNF trust can be found here[17], [16] and container trust. While we are familiar with downloading images from some repository and performing signature and simple integrity checks to establish validity, it is rarely enforced across the whole supply-chain. While signatures and measurements in the form of cryptographic hashes are used, these do not guarantee the provenance of the element as is extensively described in [18]

B. Trusted 5G/6G Core

There are two mechanisms for establishing a trusted 5G core, the first is to run the core itself upon trusted hardware and infrastructure - this is the mechanism effectively described in III-A. In this section we describe how remote attestation can be added as a core service and the possible interactions between it and other components. The simplest case is that involving the access and mobility management function and the user equipment. Figure 3 shows the main components we have considered currently with the inclusion of the Remote Attestation Services (RAS) as a 1st class component inside the core.

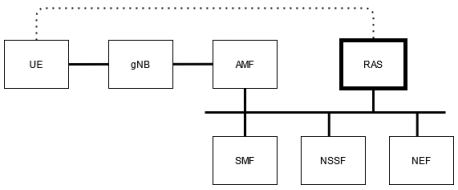

Fig. 3. 5G Core Main Services with Remote Attestation

1) AMF-RAS: As the Access Management Function - AMF - is responsible for the identification of a piece of user equipment it makes sense that via the N1 and N2 interfaces it reports to the RAS to create or update entries for attestable elements. Whether an element is attestable would be reported during initial communications with that UE. As multiple AMFs may be present the use of the GUAMI to further identify the provenance of the requests to the RAS would be included - primarily for auditing purposes at this time.

The AMF, on request or independently of the RAS, would require a claim (or claims) from the UE about its identity and integrity. In the case of a wholly new UE this would cause a new element record to be generated with the relevant identity and claim information from the UE. There are cases where the RAS may be pre-provisioned with element data, in which case an attestation would be caused to happen and a check for differences in the identity and any changes to the integrity according to the attestation rules.

Where a change has occurred a decision is required by the RAS whether that device is still trusted. This would require the RAS requesting information from the TA on the UE - denoted by the dotted line in fig.3. This communication may occur over N1 as we have prototyped. If the device is still trusted then this is communicated to the AMF and on to other 5G services as required. Where a device loses its trust status then a number of actions can take place from sandboxing to full removal of that UE from any more interaction with that 5G core. We discuss the interaction with the slicing management NSSF later.

2) NSSF-RAS: We consider slicing to be a key function for trusted elements in a 5G network. While 5G concentrates on network slicing, 6G systems will allow much finer grained and much more dynamic slicing opportunities. We extend this with the notion of trust slicing [19] as an additional property associated with a slice. In the proof of concept systems these slices are independent of the network slicing concept and work is underway on mechanisms for possible unification of these through some new or existing policy mechanism.

A slice can be constructed to hold a set of elements - UEs - all of which must adhere to a given trust policy. Upon successful attestation the AMF can instruct the NSSF to include that element in a given slice. We consider three trust slices to always be present: trusted, untrusted and untrustable. The latter slice is utilised for those UEs who do not support attestation. The untrusted slice would be utilised for those devices which have failed their verification. This slice can then be sandboxed by any mechanism necessary - we might take advantage of network slicing here to partition these devices and monitor their traffic for anomalies [20], [21].

C. UE-RAS

For a UE to be trustable it must response to messages from the RAS using some form of trust agent. A protocol for this has been described earlier and this in the 5G context may be part of the N1 protocol. The UE is required to return a claim - a TPMS_ATTEST structure in the case of a TPM 2.0 root of trust. It is assumed that the identities of TPM and the UE are linked in the long-term though there is still discussion about exactly how this takes place. At a minimum if this link is broken then it can be tracked over time.

The next question is what information is measured and stored on the TPM in the PCRs. At present no TCG profile exists for such systems as it does for x86 or automotive systems [22], [23] - the latter provides a basis for the kinds of hardware used to implement the UE functionality in 5G. As with any trusted device an immutable core root of trust must be provided. Analogously to the x86 specification we propose a similar structure of CRTM, firmware, configuration information, additional ROMs and their configurations (if present) and then finally a PCR for any user defined information. Further information may be provided in the TPM's NVRAM, for example, certificates from manufacturers - again analogous to the x86 situation which specifically reserves NVRAM areas for manufacturer certificates.

D. Interactions Between Core and Cloud

For the purposes here we assume that the 5G core is providing networking services and management to a cloud environment - this is the typical case - and provides us with the following architecture as shown in figure 4.

E. Device/UE Architecture

In figures 5 and 6 we present two architectural patterns for the construction of devices [24], [25], for example, an IoT device that communicates over 5G. We misuse the term UE to denote the 5G radio hardware and supporting circuitry and antenna, while device is used to denote the circuitry (and software stack) providing the overall functionality.

In figure 5 the device and and UE are separate and independent only communicating across some bus - this pattern is typically seen in many internet enabled devices. The bus here might be anything from on board SPI to USB/Serial etc. To establish a link between these components an interface must

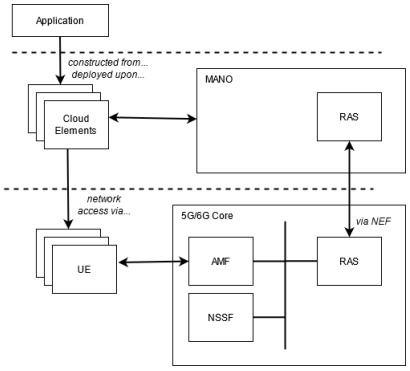

Fig. 4. Core Cloud Interaction

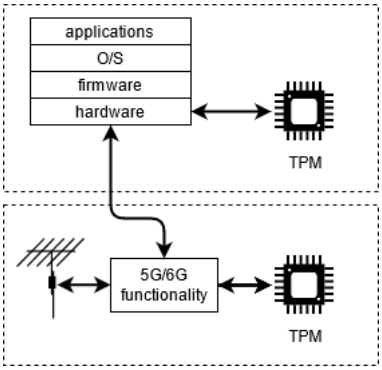

Fig. 5. Device with Two TPM Modules

be provided for one side to request measurements and TPM related information from the other, eg: endorsement certs, quotes etc.

In experiments with x86 hardware, we have used the PCR 2 and PCR 3 registers to store information from both sides as part of the optional component measurements.

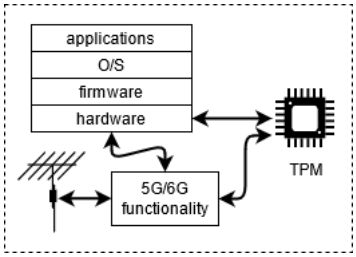

Fig. 6. Device with a Single TPM Module

In figure 6 both components share a TPM over some suitable bus. In this case the measurement sequence is more complex and needs to be customised for that specific device though this suits much smaller devices and FPGA/ASCI deployments better. The advantage in this situation is that we have a single identify that can be communicated to the AMF or cloud MANO (and thus the RAS) as required. It may be that the 5G Core RAS requests different quote information than the MANO but the identities can be easily matched.

1) Identities: AS described earlier the interaction between various elements such as UE's or devices, containers etc in a cloud environment with the management and operations - AMF and MANO respectively is known, though in the former case as we have discussed work is still required here.

A device containing some UE element joining a 5G network is required to be provisioned a priori with either or both of the remote attestation services, or, can be performed in-situ. This latter case has issues regarding the supply-chain security and the device's provenance.

Once the AMF or MANO record the element's properties (or ascertain them) in the RAS it is required to establish some kind of link between these records. This may depend upon the construction of the device in question. If a single TPM is present, and, that TPM is being used for both the identity and integrity of the UE and the device then this can be established simply through the comparison of the TPM's own endorsement keys and key identifiers.

If two separate TPMs are present then the link must be established on the device through some reasonable cryptographic means, eg: sharing of public keys an be admissible to some cryptographic challenge.

2) Trust Slicing: While 5G supports network slicing, albeit in a limited manner currently, 6G is expected to support 100s or 1000s of slices in a dynamic manner. The notion of a slice can be extended [19] to admit trust as a property.

Cloud systems also support a partitioning mechanism, for example, a container deployment may specify an internal network, or, a cloud environment may be partitioned at a more abstract level using namespacing - a technique used in Kubernetes and K3S.

Given that the two layers have slicing schemes we can link these together such that devices that have passed attestation can be admitted to decided trusted slices. Similarly devices that can not be attestated can be similarly partitioned. Each of these slices then can have additional functionality and networking properties as required - a slice containing trusted elements may then be allowed additional access to data or better networking quality etc.

Further to this we can also create slices with additional trust policies. A device in a slice with these may be required to undergo much more rigorous verification to be admitted or to remain a member of that slice.

3) Other Considerations: The systems we have constructed to demonstrate this are by their very nature designed to demonstrate these features and not address aspects such as scalability etc. At least in the traditional IT domains the issues of scalability have been addressed and the inclusion of attestation in the system integrity verification process have not proved to be significant. The amount of bandwidth required for such attestation is very low in terms of 100s of bytes of information for a TPM 2.0 quote. The frequency of attestation is similarly low only occurring at significant system changes, eg: rebooting or readmitting of a device to an AMF.

One area where latency is an issue is in real-time systems where the processing window is small. The typical time to generate a TPM quote for example is around 200-500ms with attestation taking an additional 100-200ms - though the latter figure is more due to the inefficiencies in our prototypical code. For systems, such as railway signalling this might be too much of a time penalty. Furthermore the location of the RAS or multiple RAS and their own internal database consistency needs to be taken into consideration; work is progressing in this area.

The issue of UE privacy [26] needs to be addressed in that it is possible, even in the presence of temporary identifiers such as the 5G C-RNTI, temporary C-RNTI etc, a record in a remote attestation service for a given UE would necessarily contain that TPM's identity, usually as an endorsement public key or some derived key therefrom. Though, in the obscure case where a TPM is moved between devices or some form of spoofing or replacement of the UE is made then this might

Domain	Properties
IT/Server	server trust and scalability of remote attestation
Industry 4.0	IoT trust, data flow, information provenance, privacy
Medical	network slicing: trust, device identity
Railway	latency, real-time
Aerospace	latency, real-time, distribution, bandwidth, resiliency

TABLE I
PROPERTIES OF VARIOUS DOMAINS

prove to be a useful mechanism for tracking. In the cases we have explored at this time we have not directly addressed this problem.

IV. CASE STUDIES

To explore these concepts we developed a proof of concept for a number of domains each with differing properties and requirements of trust [27]. We will describe the two main domains: medical and railways which we have used to drive the exploration of trust concepts.

The Industry 4.0 domain has provided a conceptual basis for much of this work in that we have a plurality of small (IoT) devices and different architectures. The main aspects here have been what structures are required by those devices - in terms of hardware and firmware - for trusted operation, and, how trust supports data collection and by implication data provenance, integrity and security. From this we have been able to develop the basic attestation environment and explore how it fits in with the 5G core. To this, we have developed in addition to the remote attestation, integration with the AMF and UE, integration between Edge MANO/attestation and 5G remote attestation and ideas about how data provenance and fault tolerance might be achieved while adhering to trusted operations. Further to this the use of roots of trust with identities also allows us to take advantage of 'hyper-local' homomorphic encryption, ie: individual sensors themselves being able to rapidly encrypt data and release that for processing in the edge environment [28].

Work has been made in the railway domain on the use of 5G for security, real-time communication for signalling coupled with the use of edge cloud as the mechanism by which the operations on the locomotive are deployed [29]. This work demonstrated while the use of 5g slicing to provide the bandwidth and latency guarantees for railway signalling traffic coupled with the establishment of the trust of the devices within that slice. One result from this was the necessity to distribute remote attestation to the edge elements in order to eliminate any latency caused by the request for attestation into the latency of the railway signalling commands [30].

We now describe the medical domain case study in a little more deatil addressing possible use cases for attestation and slicing with respect to the 5G/6G core functionality.

A. Medical Domain

Most medical devices are based around the microcontroller or one or more single board computers. For example a home CPAP machine is little more than a set of pressure/flow sensors, display and motor for blowing air. Additional systems such as mobile connectivity, display, data capture etc may be provided. A hospital patient monitor may be constructed from a number of SBCs, for reliability and resiliancy purposes, to provide monitoring, data capture and analuysis, interfacing (eg: to anaesethetic or veltilator systems), alarms and display etc.

In both cases these systems rely upon a degree of unfamiliarity and obscurity of the system to provide security. Some systems now allow remote operation and integration with cloud services for data capture and even over-the-air software updates. Extraction, modification and loading of firmware on these devices, while not trivial, is relatively easy through accessing the reflashing pins for example. Debugging tools suchas Ghdira can then be utilised to analyse the contents of flash memory and the executables for interesting functionality and objects, eg: public and private keys.

Given the known tampering vectors in the supply-chain and run-time there are ample opportunities to affect the behaviour of these devices. At a device and software level, the ability to measure and attest the state of the hardware, firmware, software etc as well as the device identity has obvious advantages.

1) Description of the POC: The proof of concept system we constructed consists of a number of trusted devices communicating over a 5G network. A doctor may issue devices to a patient and collect data from these devices.

- Patent Specific devices
- Edge cloud and Network slices
- Device trust
- Data provenance

To support the required data collection, processes and services for a patient an Edge cloud is set up specifically for that patient-doctor combination. Automated deployment of Edge clouds is a known functionality. In this case we utilise the remote attestation to establish both the trust of the underlying hardware allocated and of the services and supporting infrastructure being loaded. For example, if heart rate data is being collected then suitable container images are loaded to collect and process this data - these images are required to be trusted.

The devices being issued to the patient are likely to be pre-provisioned and thus known to some central, health care database. This information can be utilised to preload the RAS for the above edge cloud. As these devices are also mobile they would contain 5G/6G hardware - again trusted or at least trustable as described earlier.

A slice (or slices!) in the 5G/6G core would be constructed to support communication inside and externally with the edge cloud and these devices. A key concept here is that there would exist a slice in which all devices, services etc are trusted. Communication in and out of this slice can then be more tightly controlled and internally additional network level protection may be put in place, as well as latency, real-time and bandwidth guarantees or requirements. This network slice is then linked to the edge slice or namespace.

When a device connects to the mobile network the AMF and other core functionality can then attest that device and allocate it to the necessarily network slice.

Once the device is connected to the network it can then communicate with the OSS/BSS layer providing the medical services. This layer can further call the MANO and RAS to establish the identity and integrity of this device. These calls to the Edge RAS will necessarily call the Core RAS to ensure that both layers are trusted.

B. Data Provenance

If the identity of a device can be established then this identity can be combined with the data to establish its provenance. The usual scheme of signing data with a digital signature, for example, a suitable singing key linked to the TPM provides one mechanism for establishing this.

We can enhance this by also including with the signature a quote of the integrity of the device at the time of data collection/transmission etc.

This would allow any received of data to both verify the device from which the data was collected and also to interact with the RAS to establish whether that data was collected at a point in time when that device/service was in a trusted state. The attestation environment we have constructed allows this kind of historical information to be ascertained additionally.

Failure to establish the identity of the sending party is usually indicative of some kind of fraud. Failure to establish

the integrity of the sending device however may indicate a wider range of problems but not necessarily invalidate the data. Knowledge that there is a problem can effect a much more detailed and intelligent forensics and decision over the validity of that data.

1) Trust Failures: If a device either at the edge or 5G level fails its attestation then the device can be removed from the network. However in the case of safety-critical systems this might not be a valid response.

Firstly however it is necessary to establish the cause of failure and whether the device really has failed, and why. Mechanisms for forensics collection and establishment of the failure have been discussed in [31], [32] - the description of a fuller root cause analysis of trust failures and forensics can be found in [33], [34]. Work is on-going on the response orchestration. Secondly, even if a device fails, it may be required that the data from that device is still received, albeit with a warning about the validity or trustworthiness of the data.

In the case of a trust failure we can instruct the NSEF to create specific sandboxed slices as necessary to hold and isolate these misbehaving devices. Similar mechanism may be replicated in the edge cloud partitioning to support the sandboxing.

V. Conclusions

This paper has outlined the integration of trusted computing with a 5G/6G supported cloud environment. Such environments contain remote attestation to assess the trustworthiness of the devices and services they contain. By doing so we can eliminate a number of attack vectors, especially those against identity and firmware tampering which will become (if not already, but largely undetected) a major attack vector. It is the author's opinion that many of these attacks have passed unnoticed already.

A number of the advantages of 5G and 6G, especially with regards to functionality such as slicing, when combined with remote attestation can be utilised to provide for a much more resiliant system under trust failure scenarios.

The use of trusted elements in safety-critical systems supported by 5G will become a necessary security property. We can not rely upon perimeter security alone to establish the security of a system. For example, railway signalling over 5G must establish both the identity and integrity of the sending and receiving devices as well as the integrity of the signalling message itself.

The addition of trust properties to any system is complex and is often mistaken for just adding a TPM or remote attestation to a system without addressing the possible interactions and additional functionality and benefits this provides. As we have seen the UE, AMF and NSEF interactions in the 5G core with the RAS is already complex. Similar complexity is seen in cloud environments, especially when container and virtual image/instance trust is also take into consideration. This paper has provided an overview and possible interactions for this.

References

[1] Y. Danidou, "Trusted computing initiative on the spectrum of eu cybersecurity legal framework," in *EU Internet Law in the Digital Era*. Springer, 2020, pp. 277–296.

[2] C. Benzaïd, T. Taleb, and M. Z. Farooqi, "Trust in 5g and beyond networks," *IEEE Network*, 2021.

[3] R. Langner, "Stuxnet: Dissecting a cyberwarfare weapon," *IEEE Security & Privacy*, vol. 9, no. 3, pp. 49–51, 2011.

[4] A. Di Pinto, Y. Dragoni, and A. Carcano, "Triton: The first ics cyber attack on safety instrument systems," in *Proc. Black Hat USA*, 2018, pp. 1–26.

[5] L. H. Newman, "Medical devices are the next security nightmare," *WIRED, Mar*, 2017.

[6] J. Rieck, "Attacks on fitness trackers revisited: a case-study of unfit firmware security," *Sicherheit*, 2016.

[7] A. Tomlinson, *Introduction to the TPM*. Cham: Springer International Publishing, 2017, pp. 173–191.

[8] H. Birkholz, D. Thaler, M. Richardson, N. Smith, and W. Pan, "Remote attestation procedures architecture," *Internet Engineering Task Force, Internet-Draft draft-ietf-rats-architecture-12*, April 2021.

[9] P. Sayyad Khodashenas, C. Ruiz, M. S. Siddiqui, A. Betzler, and J. Riera, *The role of Edge Computing in future 5G mobile networks: concept and challenges*, 04 2017.

[10] S. Wijethilaka and M. Liyanage, "Security orchestration framework for federated network slicing," 06 2021.

[11] B. Jaeger, "Security orchestrator: Introducing a security orchestrator in the context of the etsi nfv reference architecture," 08 2015, pp. 1255–1260.

[12] T. Sechkova, E. Barberis, and M. Paolino, "Cloud & edge trusted virtualized infrastructure manager (vim)-security and trust in openstack," in *2019 IEEE Wireless Communications and Networking Conference Workshop (WCNCW)*. IEEE, 2019, pp. 1–6.

[13] M. Eckel, A. Fuchs, J. Repp, and M. Springer, "Secure attestation of virtualized environments," in *ICT Systems Security and Privacy Protection*, M. Hölbl, K. Rannenberg, and T. Welzer, Eds. Springer International Publishing, 2020.

[14] A. Lioy and A. Bertorello, "Hardware-bound virtual tpm for cloud computing deep attestation," 2020.

[15] D. Ganesan, M. Y. Sharum *et al.*, "A survey on advanced schemes applied within trusted platform modules (tpm) and iaas in cloud computing," in *2021 5th International Conference on Computing Methodologies and Communication (ICCMC)*. IEEE, 2021, pp. 317–322.

[16] B. Vigmostad, "Enhancing trust and resource allocation in telecommunications cloud," Master's thesis, Department of Computer Science, 2018.

[17] S. Ravidas, "Incorporating trust in network function virtualization," Master's thesis, Department of Computer Science, 2016.

[18] T. Victor, "Providing trusted computing services for multi-access edge cloud computing," Master's thesis, School of Science, 2021.

[19] I. Oliver, "Trusted computing and slicing in the dynamic environment," *Security of Hardware and Software Development, ETSI Security Week 2019*, 2020.

[20] Z. Kotulski, T. W. Nowak, M. Sepczuk, and M. A. Tunia, "5g networks: Types of isolation and their parameters in ran and cn slices," *Computer Networks*, vol. 171, p. 107135, 2020.

[21] S. Lal, A. Kalliola, I. Oliver, K. Ahola, and T. Taleb, "Securing vnf communication in nfvi," in *2017 IEEE Conference on Standards for Communications and Networking (CSCN)*. IEEE, 2017, pp. 187–192.

[22] *TCG PC Client Platform Firmware Profile Specification, Family 2.0, Level 00, Revision 1.05*, Trusted Computing Group, December 2019.

[23] *Protection Profile Automotive-Thin Specific TPM, Family 2.0, Level 0, Revision 1.0*, Trusted Computing Group, December 2018.

[24] K. N. McGill, "Trusted mobile devices: Requirements for a mobile trusted platform module," *Johns hopkins apl technical digest*, vol. 32, no. 2, pp. 544–554, 2013.

[25] E. Pisko, K. Rannenberg, and H. H. Robnagel, "Trusted computing in mobile platforms," *Datenschutz und Datensicherheit*, vol. 29, no. 9, pp. 526–530, 2005.

[26] I. Loutfi and A. Jøsang, "Privacy concerns of tpm 2.0," in *European Conference on Cyber Warfare and Security*. Academic Conferences International Limited, 2016, p. 205.

[27] I. Oliver, "Trusting the verticals: from a trusted 5g core to rail, automotive, medical and beyond," *5G Security for Verticals, ETSI Security Week 2020*, 2020.

[28] M. Ekblom, "Applications of Homomorphic Encryption," Master's thesis, Aalto University. School of Science, 2015.

[29] R. Bäckman, "Simulating rail traffic management with trusted computing," Master's thesis, South-Eastern Finland University of Applied Sciences (XAMK), Kotka, Finland, 2020.

[30] R. Bäckman, I. Oliver, and G. Limonta, "Integrity checking of railway interlocking firmware," in *International Conference on Computer Safety, Reliability, and Security*. Springer, 2020, pp. 161–175.

[31] I. Oliver, G. Limonta, and B. Vigmostad, "Improving system trustworthyness by combining remote attestation and root cause analysis," in *2018 11th International Conference on the Quality of Information and Communications Technology (QUATIC)*. IEEE, 2018, pp. 293–294.

[32] B. Manral, G. Somani, K.-K. R. Choo, M. Conti, and M. S. Gaur, "A systematic survey on cloud forensics challenges, solutions, and future directions," *ACM Computing Surveys (CSUR)*, vol. 52, no. 6, pp. 1–38, 2019.

[33] I. Oliver, "An approach to combining medical device fault analysis with trusted computing forensics," in *2020 IEEE International Conference on Big Data (Big Data)*. IEEE, 2020, pp. 1831–1837.

[34] A. Nieto, "An overview of proactive forensic solutions and its applicability to 5g," in *2018 IEEE 5G World Forum (5GWF)*. IEEE, 2018, pp. 191–196.

Acknowledgements

The authors wish to thank Kiti Müller (Aalto University), Victor Ţrucanu (Bell Labs), Kasper Kyllönen (Bell Labs), Ronny Bäckman (Rejlers Finland) and Marko Vatanen (JAMK: University of Applied Sciences, Jyväskylä) for their input.

This work has been partially funded by EU ECSEL Project SECREDAS (Grant Number: 783119)