



<http://www.diva-portal.org>

Postprint

This is the accepted version of a paper presented at *Workshop on Decentralized IoT Systems and Security (DISS) 2019*.

Citation for the original published paper:

Ahmadi Mehri, V., Ilie, D., Tutschku, K. (2019)

Designing a Secure IoT System Architecture from a Virtual Premise for a Collaborative AI Lab

In: USA

N.B. When citing this work, cite the original published paper.

Permanent link to this version:

<http://urn.kb.se/resolve?urn=urn:nbn:se:bth-17550>

Designing a Secure IoT System Architecture from a Virtual Premise for a Collaborative AI Lab

Vida Ahmadi Mehri
Blekinge Institute of Technology
vida.ahmadi.mehri@bth.se

Dragos Ilie
Blekinge Institute of Technology
dragos.ilie@bth.se

Kurt Tutschku
Blekinge Institute of Technology
kurt.tutschku@bth.se

Abstract—IoT systems are increasingly composed out of flexible, programmable, virtualised, and arbitrarily chained IoT elements and services using portable code. Moreover, they might be sliced, i.e. allowing multiple logical IoT systems (network + application) to run on top of a shared physical network and compute infrastructure. However, implementing and designing particularly security mechanisms for such IoT systems is challenging since a) promising technologies are still maturing, and b) the relationships among the many requirements, technologies and components are difficult to model a-priori.

The aim of the paper is to define design cues for the security architecture and mechanisms of future, virtualised, arbitrarily chained, and eventually sliced IoT systems. Hereby, our focus is laid on the authorisation and authentication of user, host, and code integrity in these virtualised systems. The design cues are derived from the design and implementation of a secure virtual environment for distributed and collaborative AI system engineering using so called AI pipelines. The pipelines apply chained virtual elements and services and facilitate the slicing of the system. The virtual environment is denoted for short as the *virtual premise (VP)*. The use-case of the *VP* for AI design provides insight into the complex interactions in the architecture, leading us to believe that the *VP* concept can be generalised to the IoT systems mentioned above. In addition, the use-case permits to derive, implement, and test solutions. This paper describes the flexible architecture of the *VP* and the design and implementation of access and execution control in virtual and containerised environments.

I. INTRODUCTION

The Internet of Things (IoT) is evolving. Originally, it aimed at wirelessly interconnecting cyber-physical devices (sensors, actuators, wearables, vehicles, home appliances) with control applications [1]. However, complex enhancements of the IoT have been recently developed. IoT is increasingly viewed as a service- and Cloud-based system, e.g. [2]. Such systems leverage, for example, the intermittent nature of wireless connectivity or provide for stateless information access using RESTful services. The concepts of Fog computing [3] and Multi-access Edge Computing (MEC) in public cellular networks [4] also complement the original IoT idea. Here, edge devices are located close to sensors and perform substantial amounts of computation, storage, communication (e.g. high bandwidth or multi-homing for reliability). Furthermore, slicing concepts for IoT systems have been proposed [5]. They typically focus only on slicing techniques for the network segment and use SDN (Software Defined Networking) concepts to control the data forwarding.

Our hypothesis is that IoT systems accelerate even stronger towards flexible, virtualised, programmable, sliceable, and arbitrarily chained IoT elements and services using portable code. We substantiate this hypothesis by highlighting three technology trends: a) small but powerful single-board computers became popular, such as Raspberry PI [6] with 14Mio. units sold by 2017 [7]; b) various hypervisors and container environments became available for this class of devices [8] and c) component-based application development [9] and orchestration of containers [10] have gained traction in application and service provisioning.

While it becomes clear that future IoT systems pick up concepts and mechanisms from virtualisation, slicing and code portability, limited research has been conducted on the impact of those technologies on IoT security architectures and mechanisms. A recent analysis [11] shows that network security, data encryption, API security and authentication and PKI infrastructures are among the most required technologies.

Our approach is to focus on authentication and authorisation mechanisms in IoT systems that allow for virtualisation, slicing and code portability. However, many virtualisation and slicing technologies are still evolving. For example, Docker container security is increasingly addressed [12], [13], however not yet solved. This status impacts the engineering of ICT systems in general and of IoT systems in particular.

To overcome this dilemma, we applied a prototype-based approach for identifying the design cues to for security architectures and mechanisms in virtualised IoT systems. We start our investigation by analysing and comparing today's IoT architectures and draw parallels from the use-case of AI system engineering using a virtualised AI lab, denoted as secure *VP*. The lab permits the usage of a marketplace for AI artefacts and pipelines for component-based development. The use case is of particular challenge since it includes code mobility, multi-party collaboration including security, DRM, and privacy requirements for the aforementioned parts and edge devices. The *VP* aims at simplifying and orchestrating the security and security policy management. That means, the *VP* relieves the system designer from addressing the security mechanisms and management, incl. security policy enforcement (e.g. for DRM and data privacy). The system designer can focus on the AI application.

The paper is organised as follows: Sec. II describes our design concept, analyses and compares the various IoT architectures. Sec. III describes a component-based AI development concept, its requirements and the suggested architecture for the *VP*. Sec. IV focuses on the authentication and authorisation

of users, hosts, and code in virtual labs using code mobility. Sec. V discusses the lesson learned and generalisation of the mechanism, incl. their vulnerabilities. Sec. VI summarises our finding and provides a brief outlook.

II. TOWARDS FUTURE SECURE, DISTRIBUTED AND FLEXIBLE IOT SYSTEM ARCHITECTURE

A. Design Concept

Engineering security is a multidisciplinary endeavour [14]. Designing ICT architectures, and IoT systems, in particular, are typically a rigorous and disciplined system engineering task [15]. However, it is also acknowledged in [15] that there is an increasing disagreement between theory and practice in system design. This gap is mainly accounted to the huge push for fast system integration, e.g. by market needs. We additionally attribute it to the lack of means for modelling and validating techniques for large-scale systems using pre-maturing technologies. To overcome this dilemma, [15] establishes the vision of autonomous IoT systems that can cope with the complexity through tighter integration of computing elements and compensating the reduced human intervention by enabling self-adaptation. While we believe that self-management mechanisms are of future importance, fully autonomous security control is hard to establish currently. Therefore, we focus on adaptivity in security, cf. Sec. II-B for a definition, as a first approach.

Requirements versus Architecture: Our scientific approach is not to design a fully flexible IoT architecture from scratch. We derive our design cues for security architecture from the interaction of security requirements and system architecture. In order to have a valid body of requirements and possible solutions, we consider the challenging use-case of collaborative AI system design.

From a distance, requirements specify what a system must be capable and the architecture describes how a system will be organised and behave such that it will fulfil these requirements. Hence, requirement engineering precede typically architecture design. However, when using immature technologies, architectural design choices easily influence these techniques. Therefore, we start our approach to security design by considering the IoT architectures first, cf. Sec. II-B and then define requirements for security mechanisms, in particular, authentication, described in Sec. III.

B. Comparison of IoT Architectures

Figure 1 outlines the three major IoT architectures introduced in Chapter I. In general, IoT architectures consists of three different layers namely: a) the *Things Layer*: sensors and actuators, b) the *Network Layer*: forwarding data or service routing (incl. access points and gateways), and c) the *Application Layer*: implementing data processing and control.

Table I compares the architecture of IoT concepts. It uses the features: compute and storage location, service chaining, slicing, enabling collaboration, flexibility and adaptivity towards application workflows, scalability, and security and trust.

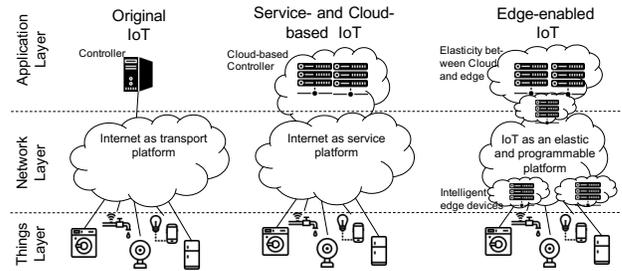


Fig. 1. Types of IoT Architectures

The original IoT concept demonstrated impressively the capabilities of the applications and the scalability of the Internet as a transport platform. Service- and Cloud-based IoT enabled a larger range of applications by augmenting the original IoT with the elasticity and flexibility of Cloud and service computing. Edge-enabled IoT permits fine-grained resource elasticity between Cloud and edge and also the slicing of the network segment. All in all, over the years the IoT concept became more flexible in terms of resource usage. However, Table I also reveals that these IoT concepts still lack in terms of flexibility of security. Flexible security, of course, aims at high security but also at specialising and adapting the security to needs of the application needs, e.g. specific data privacy is enforced at participating Cloud and edge nodes. In addition to security, trust is needed for collaboration. Hence, it also must be adaptively established for those entities who would like to cooperate on data or application.

III. PIPELINE- AND COMPONENT-BASED DEVELOPMENT OF AI SYSTEMS

The concept of a pipeline- and component-based system development was recently successfully applied for engineering data-driven AI systems [17]. Its main objective is to reduce design costs and time by re-using knowledge manifested in so-called *AI artefacts*. These artefacts are objects, code and functions used in AI development and can be chained to form an *AI engineering pipeline*. Their use permits AI specialists to focus on AI functions and to handover the result to another specialist. This approach may improve the quality of the AI solution by use of specialisation. Google has recently published a concept for a Cloud-based pipeline [18] and the H2020 project Bonseyes suggests a distributed one [19].

AI Pipeline

An excerpt of a Bonseyes AI training pipeline for Keyword Spotting (KWS) is depicted in Fig. 2, cf. [20]. The KWS applies a Deep Neural Network (DNN) model for identifying keywords in spoken sentences. This pipeline includes four steps: Data Sourcing, Data Preparation, Training, Target Benchmarking. The steps are needed to obtain and pre-process data and to train and benchmark of the DNN model. Each step is embedded in an artefact and instantiated as a Docker container. The output of an artefact is handed over as input to the succeeding one. Hence, the pipeline forms an *AI service function chain*. In Bonseyes, the pipeline is distributed, i.e.

Feature	Original IoT	Service- and Cloud-based IoT	Edge-enabled IoT	Virtual Premise
Compute location	at controller	in Cloud	elastically at Cloud / edge	elastically and programmable at Cloud and edge
Data location	at controller	in Cloud and service proxies	elastically at Cloud / edge	elastically at Cloud / edge
Service chaining	not indented	in Cloud	in Cloud and edge	per pipeline in Cloud and edge
Slicing	not indented	in network (when using SDN)	in Cloud and network (when using SDN)	application overlay per pipeline
Collaboration	single application	in Cloud	in Cloud (not across edge)	per pipeline
Flex. / adaptive to app. workflow	static by design	in Cloud	elastic and programmable	through programmability and orchestration
Scalability	by transport and processing capacity	by Cloud and service elasticity	by Cloud and edge elasticity	per pipeline through Cloud / edge elasticity
Security and trust	static by secured transmission, nodes and PKI	static by secured transmission, nodes, service API and PKI	predefined by architecture [16]	flexible per pipeline through programmability, policies and orchestration

TABLE I
COMPARISON OF THE DIFFERENT TYPES OF IoT IMPLEMENTATION

Docker containers are executed at arbitrary hosts (incl. data-center or local developer workstations). The outcome of the training is an AI model that can be converted into code, which even can be deployed on IoT devices (if capable enough).

AI Marketplace

A core element of the Bonseyes project is the *AI Marketplace (MP)*, cf. Fig. 3., which enables the exchange of AI artefacts and the component-based AI design using artefacts. Developers can develop artefacts, provide and upload them into the MP as containers. The containers are stored by the MP in the *Bonseyes' Private Docker repository*. In turn, AI system designer can search the MP for artefacts, download from the BPDR, and place them into pipelines. In this way, the marketplace enables *collaborative development*, i.e. multiparty cooperation in AI design.

Challenges in Component-based Development Using an MP

Bonseyes use of an MP for AI artefact impose these specific *high-level challenges on security related functions*:

- 1) Protection of DRM (digital rights management): code and data are of value and developers might want to protect and manage their digital rights.
- 2) Protection of privacy and data ownership: individuals might agree on using their data for AI training but might not want to waive data privacy and data ownership.

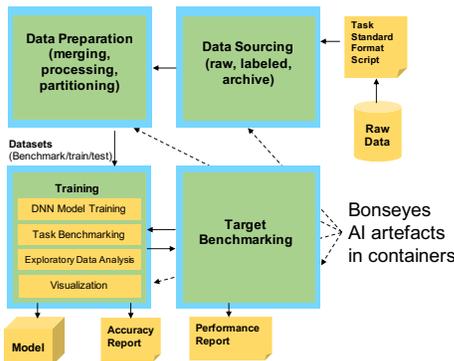


Fig. 2. An AI data collection and training pipeline, after [20]

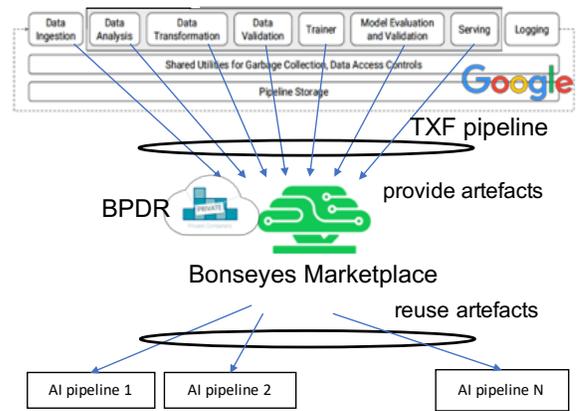


Fig. 3. Bonseyes AI marketplace, using a pipeline from [18]

- 3) Compliance in collaboration: digital rights, privacy and data ownership need to be obeyed across various physical location and legal entities when collaborating.
- 4) Fair enumeration: individuals, developers or companies need to be reimbursed when their digital intellectual property or data is used.

These high-level security challenges need to be translated into *technical security solutions*: An AI marketplace architecture is required, which enables authorised access, integrity and compliance verification for artefacts in a distributed environment, i.e. across various physical locations. Moreover, the use of virtualisation technology for artefacts requires the a) compliance verification of the virtual environment at a host, as well as b) the validation of the integrity of the artefact, i.e. the Docker image of the artefact use not changed. Hence, authentication, authorisation, and policy enforcement are the core elements of our solution.

The Virtual Premise

Our solution for solving the technical challenges is to enhance the MP by a *secure VP*, cf. Fig. 4. The core functions of the VP are authentication, authorisation and policy enforcement. VP places the AI pipeline in the centre and builds a virtual border around it. The pipeline is instantiate on

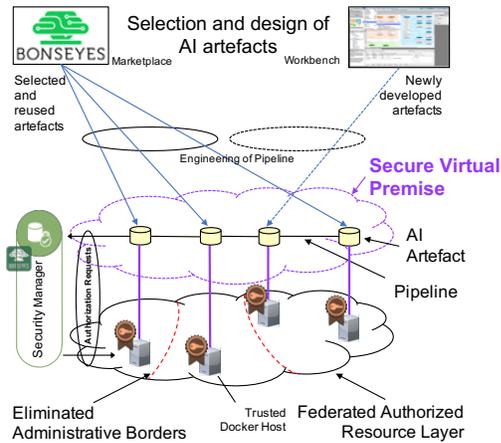


Fig. 4. Architecture of the Virtual Premise

federated authorised resources. The pipeline elements/artefacts can be downloaded from the Bonseyes MP (or created and registered in MP using a developer’s workbench). The VP uses network slicing and virtualisation techniques to provide isolation of pipelines, prevents unauthorised pipeline communication with outside and unifies the security and access policies per pipeline. It allows workflows between authorised artefacts through secure programmable and polymorphic interfaces. The VP has similarities to an application overlay and builds a *virtual AI laboratory* per pipeline. The inside of the VP is a trusted area where artefacts comply with each other on DRM, privacy and collaboration policies, while the outside is untrusted. The artefacts, hosts and users of a pipeline must be authenticated by the *Security Manager (SM)* in order to join the VP. The policies of the VP are enforced locally on the Docker hosts participating in the VP. The split between security policy decision entities and enforcement points is discussed in cf. Sec. IV.

In terms of IoT architectures, the VP can be considered as an enhancement of the Edge-enabled IoT concept which focus on slicing with respect to adaptive and flexible security policies, cf. Fig. 5. The IoT features of the VP are summarised and compared with other IoT architecture in the last column of Tab. I. It is obvious that the VP enables flexible security and trust per AI pipeline or IoT service chain.

Bonseyes AI Artefact Architecture

The second major component of the VP concept is the architecture of AI artefact, which is shown in Fig. 6. The artefact is implemented as a Docker container which is enhanced by a specific *Bonseyes Layer (BL)*. The BL serves three main purposes: a) to intercept all communication into the artefact; b) to provides secure polymorphic APIs to access AI functions and content; c) to authenticate the artefact with the SM and the Docker host executing the artefact; and d) verify the functional integrity of the underlying host platform, for maintaining the confidentiality and integrity of data processing and data exchange (e.g. enforcing access control and DRM services). The BL is called as soon as the container is executed, thus ensuring that all communication into the

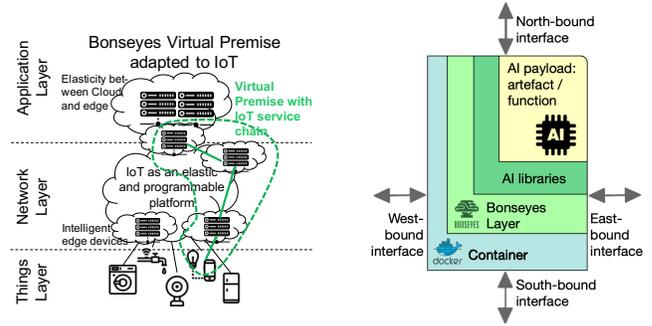


Fig. 5. Virtual Premise as an IoT Architecture

Fig. 6. A Bonseyes’ Docker container comprising an AI artefact.

container is intercepted and secured. The secure APIs for the workflow between artefacts are referred to as the *east-bound interface* and the *west-bound interface*, respectively. The API for the execution of the AI function and towards the workbench is denoted as the *north-bound interface*, while the one towards the host (incl. virtualisation environment and filesystem access) is denoted *south-bound interface*. The AI function together with associated executable code is stored in component called *AI payload* and the required libraries for this function might implemented as a layer in the container.

IV. AUTHENTICATION AND ACCESS MANAGEMENT

The MP offers an interface through which artefacts can be uploaded into or downloaded from the BPDR. In addition, the MP supports payment services for commercial artefacts. AI artefact usage can be conditioned by licensing terms, privacy requirements or local legislation. Furthermore, hosts available for artefact execution may belong to various administrative domains each free to enforce different policies in terms of which artefacts and users are allowed to execute on their systems as well as how resource allocation should be performed. Therefore access to hosts and artefacts is coordinated through the SM to enable compliance with this type of constraints through a distributed license management system.

In addition, the SM is the fundamental component that enables mutual authentication and functional integrity checks for the components belonging to a VP. This involves both the AI artefacts (*i.e.*, the containers and their contents) as well as the supporting infrastructure (*i.e.*, the server hosts executing the containers). The main concern in this case is that hosts or containers are tampered in order to circumvent the license management system and bypass the usage constraints associated with an artefact (*e.g.*, using the artefact without a valid license). However, as past experience has shown, being able to counter *any* attack on the license management system, including those where malicious code is injected into the system, is a very difficult issue. Therefore, we make a set of reasonable simplifications to relax the difficulty of the problem. In particular, we assume that the SM is not compromised and that administrative domains do not collude with attackers and instrument the hosts to bypass the license system. Also, we assume that Docker images uploaded to BPDR are searched for malicious software before becoming

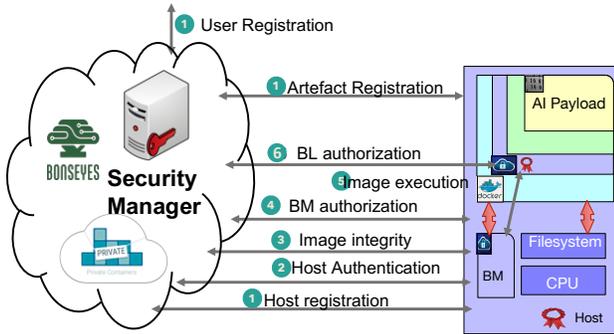


Fig. 7. The communication between Host and SM

available for download. Furthermore, we assume that the SM is equipped with a Certificate Authority(CA) that enables it to issue certificates for other entities in the MP. Essentially, the SM is the trust anchor for the system. Finally, the SM is equipped with a HTTPS server that can do certificate-based authentication.

Given these assumption, we have defined a process that verifies if a host is allowed to execute an artefact, and similarly if the artefact is allowed to execute on the given host. The process consists of a set of discrete phases listed below and shown in Fig. 7:

- 1) Artefact, host, and user registration
- 2) Host authentication
- 3) Image integrity check
- 4) Host authorisation for execution of artefact
- 5) Image execution (BL start)
- 6) Artefact authorisation for execution on host

In general, phases 1 and 2 can be performed at any time. However, given a specific artefact-host-user triplet, the other phases cannot execute correctly before they are completed successfully. Artefact registration entails that an image is created and uploaded to the BPDR. Then, the SM adds the BL layer to the image. It also inserts an artefact identifier (a human readable string) which is then signed with the SM private key yielding the *artefact ID*. This ID is shared by all instances of the same image and is used when the artefact states its type in phase 6. Finally, the SM's root certificate is added to the trust database in the image and a message digest (*e.g.*, SHA-256) is computed over the final image. The artefact ID and the digest are stored as a pair in the SM database.

For hosts undergoing the registration process the identity of the host owner is tied to the host's hardware ID (the hardware ID can be computed by, *e.g.*, reading specific entries in the SMBios.). If the host platform contains a trusted platform module (TPM), the hardware ID can be replaced with the Endorsement Key (EK) certificate for the TPM. The owner (typically the system administrator or the AI developer) is responsible for verifying that hosts are free from malware before taken in use. As part of the registration, the owner can also specify if there are restrictions in the use of the host (*e.g.*, if it is reserved for a specific developer, or restrained from executing artefacts from a specific vendor). A software component called Bonseyes Module (BM) that is tied to the hardware ID/certificate is installed during the final stage of

the registration process. The BM plays the same role for the host as the BL does for the artefact. Developers need also to register and obtain a user ID. This is a one time process, and the ID is carried over to any pipeline or VP that the developer is working on.

The hosts are configured to start the BM after the operating systems has booted up. This is where phase 2 begins. The BM will open a HTTPS connection to the SM and use the host ID (and EK certificate, if available) to authenticate itself with the SM in order to be added to the pool of available hosts. The SM will refuse to authenticate hosts that are missing from registration database or for whom the authentication fails for any other reason. If the authentication succeeds, the SM will securely transfer cryptographic material to the BM, allowing the BM to self-authenticate to other entities. This includes a time-limited, but renewable, asymmetric key pair, a certificate for the public key and the SM root certificate. The confidentiality and integrity of the cryptographic material is protected by HTTPS.

Phase 3 begins when a developer instructs the SM to assign the pipeline to a VP for execution. Each artefact can be pinned to a specific host in the VP or the SM can automatically assign them to hosts based on a predefined policy. The SM coordinates with the BMs in copying the artefacts to the hosts and in checking the integrity of the received Docker images using strong message digests (*e.g.*, SHA-256).

During phase 4, hosts request from the SM policies surrounding the execution of the artefact under the controller of the developer. The request contains the message digest of the Docker image, the host ID, the user ID and a nonce. The message digest allows the SM to determine the identity of the artefact being executed, the host ID provides the identity of the host, user ID denotes the developer executing the artefact, and the purpose of the nonce is to prevent replay attacks. The request is signed with the host's private key to protect the integrity of the message and to prove to the SM that the request is really coming from this host. The signed request is sent over HTTPS to the SM.

Upon retrieving the contents of the message and verifying its integrity the SM constructs an entry consisting of five items: image digest, user ID, host ID, nonce and host IP address. This combination, *runtime ID*, uniquely identifies the artefact instance about to be started. The reason for storing the IP address is to enable SM to open connections towards the host, if necessary. Then, the SM looks up an existing *host policy specification*, based on the contents of entry. If one is found, it is associated with the entry and stored in an internal database. The policy specifies the constraints of the host in running the artefact under the control of user ID.

If the received license specifies that the host is allowed to execute the artefact, phase 5 can begin and the BL is started. Otherwise an error message is sent to the developer and the artefact execution is aborted. When the BL is started it expects to received a copy of host license from the BM. This begins phase 6. The BL uses the SM public key to verify the signature from the license and determine if it is being executed on a

genuine host. If that is the case, the BL constructs a request for SM consisting of the license from BM, its artefact ID, and an nonce. The BL is not aware of the image digest. Therefore, the it includes the artefact ID to let SM detect if the license presented was generated for a different artefact type. The request is sent to the SM over HTTPS. The SM uses the artefact ID to lookup the image digest stored during phase 1 and compares it with the one stored in the license. If there is no match, an error is returned to the artefact which stops execution. Otherwise, the SM looks up an existing *artefact policy specification*. If none is found, a default one is used. The looked up artefact policy is associated with the runtime ID maintained by SM. The SM prepares then a reply consisting of the policy and cryptographic material generated for the BL (similar to items sent to BM in phase 2). This enables the BL to self-authenticate and communicate securely with other entities of the east-/west-/north-/south-bound interfaces defined earlier. The reply is signed by the SM and the resulting *artefact license* is transported over the established HTTPS connection to the BL. The BL learns through the license if it is allowed to execute the artefact code of it supposed to cease execution. Also, the BL is expected to forward the license to the BM before a timeout occurs. If that happens, the BM will stop BL's execution. At this point, a mutual trust has been created and the artefacts and a pipeline can begin execution.

V. GENERALISATIONS FROM THE VP AS DESIGN CUES FOR IOT SYSTEMS WITH FLEXIBLE SECURITY

The specifications of the high-level security and trust requirements and implementation of mechanisms for collaborative AI design using pipelines and artefacts gives detailed insights into the security architectures and mechanisms for IoT system using similar concepts. We will generalise our findings:

- 1) *Collaborative and distributed workflows in IoT service chains require security and compliance verification and enforcement along the workflow, i.e. in an "horizontal way" (authentication for cooperation)*. The policies, e.g. for DRM or privacy, need to be enforced on the involved IoT element, e.g. edge devices, along the east/west-bound interfaces of the artefacts.
- 2) *Code mobility in IoT, i.e. use of portable container, requires security and compliance checking on the IoT host, i.e. in a "vertical way" (authentication for resources)*.
- 3) *The use of a virtualisation concepts and of containers in IoT need the verification of integrity of containers and hosts*. Therefore, the container and the host need to implement a module/layer to verify each others integrity.
- 4) *Docker's container technology is a great tool for component-based system design for AI system and probably also for IoT service chains*.
- 5) *Docker's mechanisms to ensure that code in the container can not be by-passed are still limited*.
- 6) *A centralised Security Manager is typically a single-point-of-failure/attack and needs to be distributed in future*.

An initial analysis of the security threads against the VP is given in [21]. It lists 13 simple threats by unexperienced

malicious users, ranging from by passing license constraints to artefact execution on wrong VPs, and techniques to protect against them. While these techniques are not yet fully implemented in the Bonseyes MP and VP demonstrator, their availability suggests that VP-enabled architectures can be secured against them. More difficult to solve are advanced threats from skilled malicious users. A major approach is to raise the overall trust level by using TPMs.

VI. CONCLUSION

The investigation of the security, DRM, and trust requirements for pipeline-based collaborative AI system design and the specification of the VP enabled the investigation of similar concepts for future IoT architectures. These IoT systems will apply service chains where flexible and adaptive security per service is needed.

The suggested solution builds a *VP* which forms a virtual boarder around the pipeline or service chain. This chain facilitates also the slicing of the full infrastructure, i.e. of the application infrastructure and of the network. The use of containers in such IoT systems requires the authentication of both, containers and execution environments. Container and virtualisation technologies are still rapidly evolving in their capabilities. Hence, their future capabilities might impact the features of the VP, e.g. the VP might need to enforces that a container starts always the Bonseyes Layer either by Docker mechanisms or by other verification mechanisms. Another future feature for the VP might be the automatic policy compliance matching, enforcement and eventually policy orchestration, e.g. using distributed ledgers [22]. Another future enhancement is the use of distributed TPMs [23].

ACKNOWLEDGMENTS

The authors would like to thank the Bonseyes team for extensive discussions, in particular our gratitude goes to Tim Llewellyn, Samuel Fricker, Lorenzo Keller and Yuliy Maksimov. This project has received funding from the European Unions Horizon 2020 research and innovation programme under grant agreement No 732204 (Bonseyes). This work is supported by the Swiss State Secretariat for Education Research and Innovation (SERI) under contract number 16.0159. The opinions expressed and arguments employed herein do not necessarily reflect the official views of these funding bodies.

REFERENCES

- [1] F. Mattern and C. Floerkemeier. From the Internet of Computers to the Internet of Things. In *From active data management to event-based systems and more*. Springer, 2010.
- [2] ITU-T. Overview of the internet of things. ITU-T Recommendation ITU-T Y.4000/Y.2060 (06/2012), 6 2012.
- [3] F. Bonomi, R. Milito, J. Zhu, and S. Addepalli. Fog Computing and its Role in the Internet of Things. In *Proceedings of the First MCC workshop on Mobile cloud computing*, 2012.
- [4] T. Taleb et al. On multi-access edge computing: A survey of the emerging 5G network edge cloud architecture and orchestration. *IEEE Communications Surveys & Tutorials*, 19(3):1657–1681, 2017.
- [5] V. Sciancalepore, F. Cirillo, and X. Costa-Perez. Slice as a Service (SlaaS) Optimal IoT Slice Resources Orchestration. In *Proc. of GLOBE-COM 2017*, 2017.
- [6] Raspberry Pi Foundation. Raspberry Pi. www.raspberrypi.org.

- [7] L. Tung. Raspberry Pi: 14 million sold, 10 million made in the UK. ZDnet Information available at:<https://www.zdnet.com/article/14-million-raspberry-pis-sold-10-million-made-in-the-uk/>, 2017.
- [8] P. Bellavista and A. Zanni. Feasibility of Fog Computing Deployment Based on Docker Containerization over RaspberryPi. In *Proc. of the 18th Int. Conf. on Distributed Computing and Networking*, 2017.
- [9] P. Mohagheghi and R. Conradi. An empirical investigation of software reuse benefits in a large telecom product. *ACM Transactions on Software Engineering and Methodology (TOSEM)*, 17(3):13, 2008.
- [10] B. Burns, B. Grant, D. Oppenheimer, E. Brewer, and J. Wilkes. Borg, Omega, and Kubernetes. *Queue*, 14(1):10:70–10:93, January 2016.
- [11] G. Press. TechRadar: 6 Hot Internet of Things (IoT) Security Technologies. Forrester Research Information available at: <https://www.forbes.com/sites/gilpress/2017/03/20/6-hot-internet-of-things-iot-security-technologies>, 2017.
- [12] H. Gantikow, C. Reich, M. Knahl, and N. Clarke. Providing security in container-based hpc runtime environments. In *International Conference on High Performance Computing*, pages 685–695. Springer, 2016.
- [13] R. Yasrab. Mitigating docker security issues. *arXiv preprint arXiv:1804.05039*, 2018.
- [14] C. C. Wood. Why information security is now multi-disciplinary, multi-departmental, and multi-organizational in nature. *Computer Fraud & Security*, 2004(1):16–17, 2004.
- [15] J. Sifakis. System Design in the Era of IoT? Meeting the Autonomy Challenge. In *Proc. of the 1st Int. Workshop on Methods and Tools for Rigorous System Design (MeTRiD 2018)*, 2018.
- [16] R. Roman, J. Lopez, and M. Mambo. Mobile edge computing, fog et al.: A survey and analysis of security threats and challenges. *Future Generation Computer Systems*, 78:680–698, 2018.
- [17] S. Yegulalp. Data in, intelligence out: Machine learning pipelines demystified. Information available at: www.InfoWorld.com, August 2018.
- [18] D. Baylor et al. TFX: A tensorflow-based production-scale machine learning platform. In *Proc. of 23rd ACM SIGKDD*. ACM, 2017.
- [19] T. Llewellynn and other. Bonseyes: platform for open development of systems of artificial intelligence. In *Proc. of the Computing Frontiers Conference*. ACM, 2017.
- [20] M. de Prado, J. Su, R. Dahyot, R. Saeed, L. Keller, and N. Vallez. AI Pipeline – Bringing AI to You. Submitted; available on request from: <https://www.bonseyes.com/>, 2018.
- [21] V. A. Mehri, D. Ilie, and K. Tutschku. Privacy and DRM requirements for collaborative development of ai application. In *Proceedings of ARES*. ACM, Hamburg, Germany, August 2018.
- [22] Tim Swanson. Consensus-as-a-service: a brief report on the emergence of permissioned, distributed ledger systems. Report available at: <http://www.ofnumbers.com/wp-content/uploads/2015/04/Permissioned-distributed-ledgers.pdf>, 2015.
- [23] P. Wagner et al. Distributed usage control enforcement through trusted platform modules and sgx enclaves. In *Proc. of 23rd ACM on Symp. on Access Control Models and Tech*. ACM, 2018.