4

The FORTIKA Accelerated Edge Solution for Automating SMEs Security

Evangelos K. Markakis[1], Yannis Nikoloudakis[1], Evangelos Pallis[1], Ales Černivec[2], Panayotis Fouliras[3], Ioannis Mavridis[3], Georgios Sakellariou[3], Stavros Salonikias[3], Nikolaos Tsinganos[3], Anargyros Sideris[4], Nikolaos Zotos[4], Anastasios Drosou[5], Konstantinos M. Giannoutakis[5] and Dimitrios Tzovaras[5]

[1]Department of Informatics Engineering, Technological Educational Institute of Crete, Greece
[2]XLAB d.o.o., Slovenia
[3]Department of Applied Informatics, University of Macedonia, Greece
[4]Future Intelligence LTD, United Kingdom
[5]Information Technologies Institute, Centre for Research & Technology Hellas, Greece
E-mail: Markakis@pasiphae.teicrete.gr; Nikoloudakis@pasiphae.teicrete.gr; Pallis@pasiphae.teicrete.gr; ales.cernivec@xlab.si; pfoul@uom.edu.gr; mavridis@uom.edu.gr; geosakel@uom.edu.gr; salonikias@uom.edu.gr; tsinik@uom.edu.gr; Sideris@f-in.co.uk; Zotos@f-in.co.uk; drosou@iti.gr; kgiannou@iti.gr; tzovaras@iti.gr

4.1 Introduction

Although the recent trend for the term "cyber-attack" is restricted for incidents causing physical damage, it has been traditionally used to describe a broader range of attempts to make unauthorized use of an asset related to computer information systems, computer networks, or even personal computing devices. As such, a cyber-attack aims to steal, alter a targets' system/data, or even destroy targets by gaining access into a targeted system. In this respect, a whole new industry has been shaped around the need for protection against cyber-attacks, i.e. the "cyber-security" domain,

which primarily deals with the protection of systems (incl. HW/SW & data) connected to the internet against cyber-attacks and should not be necessarily mixed with the domain of Information Technology (IT) Security (see Figure 4.1) that mainly refers to the protection of information. Cyber-security, on the other hand, is the ability to protect or defend the use of cyberspace from cyber-attacks by securing "things", vulnerable through ICT.

The first cyber-attack was recorded in 1989, in the form of a computer worm (i.e. malware), while their number has significantly grown in the following years (see Figure 4.2). Equal growth has been noted in the level

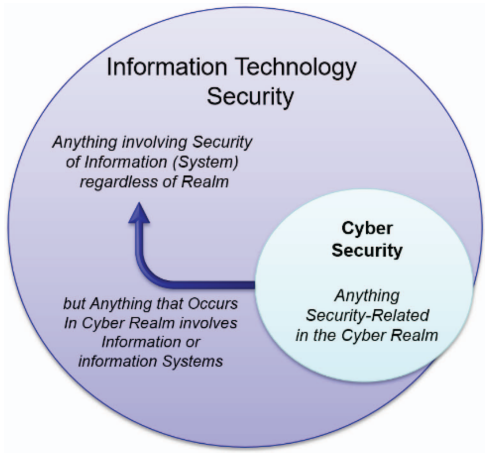

Figure 4.1 Information technology security vs cyber-security.

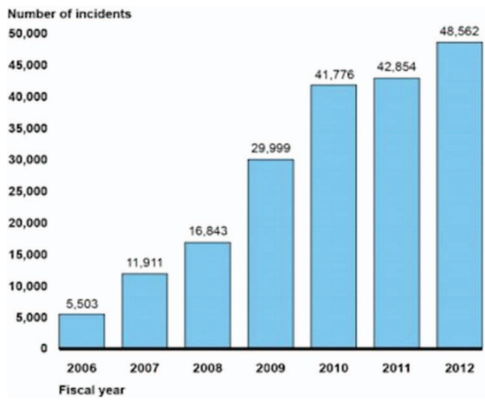

Figure 4.2 Incidents reported to US-CERT, Fiscal Years 2006–2014.

(*Source: GAO Analysis data of US-CERT*).

of both the threat they pose and the sophisticated manner with which they are launched and/or acting. Specifically, cyber-security threats have evolved from standalone threats that could affect single targets, to more complicated scenarios, where threats could be self-replicated, mutated and expanded to other devices and/or networks via the internet. Finally, the evolution of the exploitation manner of modern cyber-attacks is also extremely interesting. For instance, traditional ways for (i) harming infrastructures through DDoS attacks, (ii) misusing them through malwares, and (iii) mitigating them through identity spoofing are nowadays considered outdated and new emerging threats and attack scenarios are emerging, which aims at disusing sensitive soft assets through ransomware that directly lead their endangerment and their potential loss.

Unavoidably, cyber-security becomes, of great importance due to its increasing reliance on computer systems. Recently, in the era of the Internet of Things (IoT)[1], a large number of connected devices, located at the edge of the Internet hierarchy, generate massive volumes of data at high velocities. This turns centralized data models non-sustainable, since it is infeasible to collect all the data to remote data centres and expect the results to be sent back to the edge, with low latency.

Based on the constantly increasing dependency of the global economy on inter-connected digitization (i.e. world-web-web, smart-grids, IoT nets, direct communication links between platforms, etc.), it is the integrity and the availability of the prompt & uninterrupted interconnectivity that attracts great focus and investment from major players in the market. Similarly, the trend towards IoT and digital innovation, forms a flourishing business landscape for SMEs. However, this is put at stake due to the uncertain, cumbersome and most importantly costly nature of holistic cyber-security solutions. Specifically, although tailored solutions capable of providing the appropriate cyber-security levels for big companies appear, they can hardly be adapted to other environments and thus, lack in scalability, which makes them unsuitable for smaller enterprises.

The implementation of a complete and reliable edge computing security framework seems to be a promising alternative to protect an IoT environment and the overall network of an SME. In order to fulfil the IoT requirements, modern trends dictate that more resources (incl. computation, storage &

[1]IoT: The network of physical devices with connectivity (i.e. connect, collect & exchange data). The term was first introduced, when the amount of connected devices outnumbered the humans connected to the internet.

networking) must be located closer to users and the IoT devices, at the edge of the networks, where data is generated (i.e. "*edge computing*"), so as to (i) reduce data traffic especially in Internet backbone, (ii) provide in-situ data intelligence, (iii) reduce latency[2] and (iv) improve the response speed.

This way, cyber-security solutions will become more, in terms of both applicability and adaptability per use-case. Toward this direction, monolithic approaches are not enough; in-situ analysis based on usage Behaviour Analytics and Security Information & Event Management (SIEM) systems, customized at the for certain edge, seem able to offer a plausible and affordable solution, if offered as a modular product of adequate granularity in terms of offered services, so as to form an attractive product, easily customizable to the needs of the each customer.

Toward this direction, edge solutions introduce 5 major challenges[3] that require attention, namely (i) the massive numbers of vulnerable IoT devices, (ii) the NFV-SDN integrated edge cloud platform, (iii) the privacy & security[4] of the data, (iv) the interaction between edge & IoT devices and (v) the Trust & Trustworthiness.

This chapter presents an analysis on the cyber-threats landscape within generic ICT environments and its impact on SMEs, it also covers the different standardization and certification schemas that would help SMEs to support a cyber-security strategy and takes into consideration, standardization and best practices for the FORTIKA ecosystem and deployment. Additionally, the modular, edge-based cyber-security solution of the FORTIKA concept[5] is promoted within the current article. The resources required from a potential SME customer are efficiently managed, while a dedicated marketplace is a repository that can extend the basic version product with affordable functionalities tailored to the needs of each SME. On top of the latter, one can

[2]Given the complexity of cyber-security tasks and the latency imposed by the network distance between the client and the cloud infrastructure, one can deduce that cloud computing architecture, is by-design unsuitable for time-sensitive applications. The advancements in Edge Computing [1–3] allow for the efficient deployment and delivery of minimum-latency services.

[3]J. Pan, Z. Yang, "Cybersecurity Challenges and Opportunities in the New "Edge Computing + IoT" World", Association for Computing Machinery, 2018, doi: 10.1145/3180465.3180470

[4]Business data can be either sensitive or non-sensitive, depending on the type of business and the type of transaction. In any case, the sensitive and classified data must be stored and managed in a "regulated zone". With sophisticated encryption and key management, cloud storage platforms can qualify as a legitimate solution for storing and maintaining such data.

[5]https://cordis.europa.eu/project/rcn/210222/factsheet/en

selectively build the appropriate cyber-security solution that matches their needs, through combination of the correct bundles.

4.2 Related Work and Background

The increasingly connected world of people, organizations, and things is driven by the vast proliferation of digital technologies. This fact guarantees a promising future for cyber-security companies but poses a great threat for SMEs. According to Symantec [4], 60% of targeted attacks in 2015 aimed at small businesses, while "more than 430 million new unique pieces of malware were discovered". According to FireEye [5], 77% of all cybercrimes target SMEs. Simple endpoint protection through antivirus has become by far inadequate, due to the complexity and variety of cyber-threats, as well as the integration of multiple digital technologies in business processes, even in small enterprises. Modern cyber-security solutions for businesses, which are designed to provide multilayer proactive protection, use heuristics and threat-intelligence technologies to detect unknown threats, protecting a wide range of devices (e.g., PCs, servers, mobile devices, etc.) and business practices (e.g., BYOD, remote access, use of cloud-based apps and services, etc.). Due to this complexity, no single security solution can effectively address the whole threat landscape. Threats may range from relatively harmless, abusive content (such as spam messages) and other low-impact opportunistic attacks, to very harmful (malicious code), while they can escalate to targeted attacks (e.g., spyware, denial of service, etc.), with major operational and economic consequences for the enterprise.

According to ENISA [6] the top-5 threats in 2016 are mainly network-based. Consequently, a cost-effective solution for such threats could prove decisive for the future of SMEs and cannot be provided by one of the traditional methods.

Social engineering is another typical form of threat. This can be manifested either by a deceptive e-mail, installation instructions for a "free" or even "trial" piece of software, bogus sites, etc.

Moreover, Internet of Things (IoT) applications, such as healthcare and assistive technologies promise a higher level of quality of life for citizens around the world; on the other hand, however, they increase the attack surface, considerably. Legacy systems, implantable devices, and wireless networks are also eligible attack domains. Embedded systems are used more and more, e.g. in modern cars. Controlling and manipulating such entities can provide attackers with enormous power. The same holds for critical infrastructures

and drones. Therefore, the cyber security research community, needs to address those issues.

SMEs consist of diverse businesses that usually operate in the service, manufacturing, engineering, agroindustry, and trade sectors. SMEs can be innovative and entrepreneurial, and usually aspire to grow. Nevertheless, some stagnate and remain family owned. There is no single, uniformly-accepted definition of SMEs. Many definitions exist whereby SMEs are classified by different characteristics, including, but not limited to profitability, turnover, sales revenue, or the number of people employed.

The European Union defines an SME combining the number of employees, along with revenue and assets. A medium-sized enterprise [7], is defined as an enterprise which employs fewer than 250 persons and whose annual turnover, does not exceed €50 million or whose annual balance-sheet total does not exceed €43 million.

SMEs represent the "middle class" of entities using computers, with single or home users at the bottom of the hierarchy and large companies or organizations at the top. As such, SMEs lack the resources typically available in the case of large organizations, while, at the same time, they need continuous and secure operation of their systems in order to function. Security can be quite expensive and since low-investment consequences on it are not evident until a significant incident takes place, it is often very tempting to allocate the minimum of resources for it.

However, a significant security incident can prove fatal for an SME, either directly (e.g., cessation of business transactions) or indirectly (e.g., bad reputation causing most of the customers to walk away or litigation). Most SMEs do not consider themselves as having data that is of interest to cyber-criminals and quite often dismiss the need for adequately addressing vulnerabilities in their infrastructure. In reality, the opposite is true; every enterprise today collects data on employees, clients, and vendors that are of interest to cyber criminals. Consequently, it is crucial to develop cyber-security products that would focus on the needs of SMEs. Challenges for mitigating cyberthreats must be addressed and highlighted, and the need to mitigate the identified risks must be addressed as well.

FORTIKA aims to establish a reliable and secure business environment for SMEs, that will provide and ensure business continuity. The FORTIKA solution is composed of modules that are designed to provide a cohesive and cost-effective set of services that address those issues. These modules are described below.

4.3 Technical Approach

This section presents the high-level deployment diagram (see Figure 4.3) of the FORTIKA modules in the two main FORTIKA systems, namely the Cloud and the SME. In the Cloud, the Marketplace, its Dashboard, and Cloud platform related modules (i.e. Orchestrator, Cloud security, Cloud Storage) are deployed; further to that, several constituent components of FORTIKA cyber security appliances (i.e. ABAC, SEARS, Encrypted data search engine, redBorder Manager) are also deployed there. At the SME level, there are two distinct cases of deployment. In the first case, the deployment is performed at the FORTIKA-GW where the GW's operational modules (depicted in red colour) and the FORTIKA security modules (lightweight modules in the

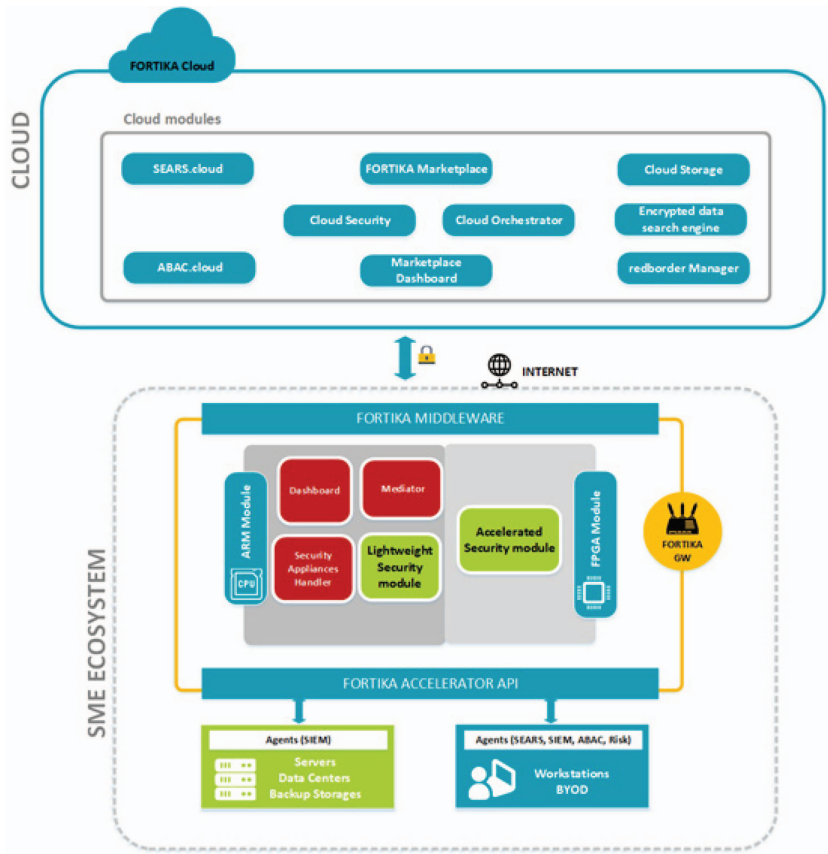

Figure 4.3 FORTIKA deployment diagram (High level).

ARM, heavyweight modules in the FPGA) can be found. In the second case, the Agents (software units collecting information and forwarding it to the GW's cyber-security appliances for processing/analysis) are deployed in the workstations and servers of the SME.

4.3.1 FORTIKA Accelerator

FORTIKA Accelerator: The FORTIKA security accelerator (FORTIKA gateway) is connected and offers unlimited expandability (by simply connecting as many accelerators in series) in terms of processing power and storage capacity, and scalability through a modular connection of two or more accelerators. Its user interface guides the enterprise administrator to appropriately define and configure the company's security & privacy policy, along with the level of encryption (information classification) and the corresponding data availability (privacy) within the enterprise and 3rd parties (e.g. suppliers, partners/ collaborators, customers, other parties), thus covering a wide range of use case scenarios. The system users/admins are kept informed at any time via comprehensive visual analytics while being able to interfere with the functionality of the presented solution in an effortless and user-friendly way.

FORTIKA Accelerator Architecture: Acceleration has been a hot topic in computing for the past few years, with Moore's law and the associated performance bumps slowly crawling to a halt. Currently, most industrial leaders accept that one form of acceleration will be used to provide the compute capacity required to cope with the large flows of data being created in the modern, widely interconnected world. FORTIKA, leverages acceleration in the form of programmable logic devices (FPGA-enabled gateway), to deliver high-performance security applications to SMEs. FPGAs offer an efficient solution in terms of performance, flexibility and power consumption. To achieve this, the FPGA must be made accessible as a resource over the network while allowing users to remotely deploy resource-demanding compute tasks on the device. This requires a middleware, either in software or in hardware to allow for the discovery of the programmable logic resources and the exchange of information between the marketplace, which is responsible for determining the appropriate infrastructure for the deployment of a task, and the accelerator module in order to determine where tasks should be deployed.

The FORTIKA accelerator module (Figure 4.4) utilizes an FPGA SoC embedded device which combines ARM processors with programmable logic in one integrated circuit. This device allows an optimal division of labour

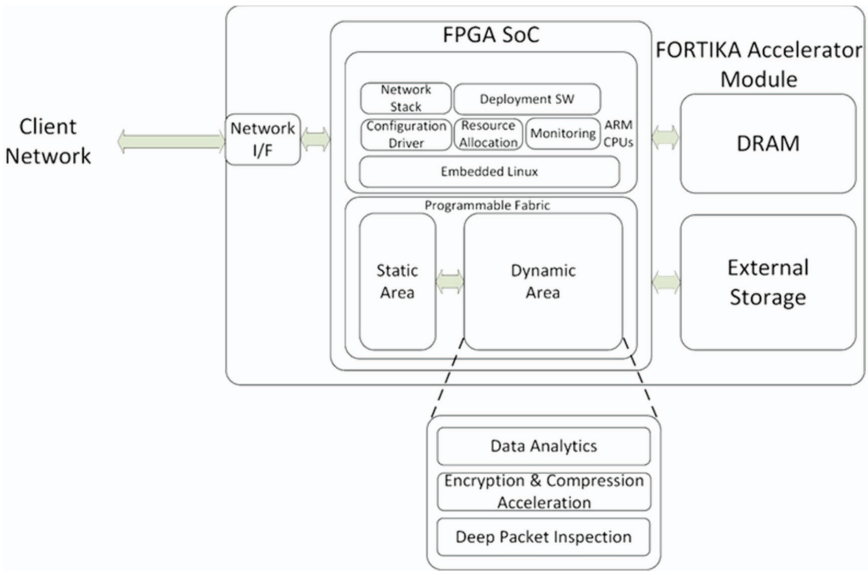

Figure 4.4 FORTIKA accelerator architecture.

between software and hardware and allows system designers, to offload computationally intensive tasks to the hardware while using the software for any light-weight, non-critical issue. FORTIKA has inherited several features from the T-NOVA FPGA-powered cloud platform, which uses OpenStack running on the CPUs to deploy tasks on the programmable logic but extended and adapted the platform to meet FORTIKA's edge demands.

The FORTIKA Middleware (MDW) (Figure 4.5) aims to facilitate a) the interactions between the FORTIKA GW and the FORTIKA marketplace; b) the loading of the security bundles to the FORTIKA accelerator; c) the exchange of data between the ARM deployed security bundles and their FPGA deployed counterparts; and d) the SW developers in producing accelerated security bundles that can be deployed in the FORTIKA accelerator. To put things in context, the following picture shows which (sub)systems, the MDW (pink Note boxes) aims to "glue" and what activities to facilitate, inside the FORTIKA architecture.

To achieve these objectives the MDW consists of several components namely the Security Bundle Handler (SBH), the LwM2M client, and the Synthesis engine. The first one provides the deployment and management of the bundles in the FORTIKA GW (both in the ARM and the FPGA parts). The second one provides the communication engine/channel which is used

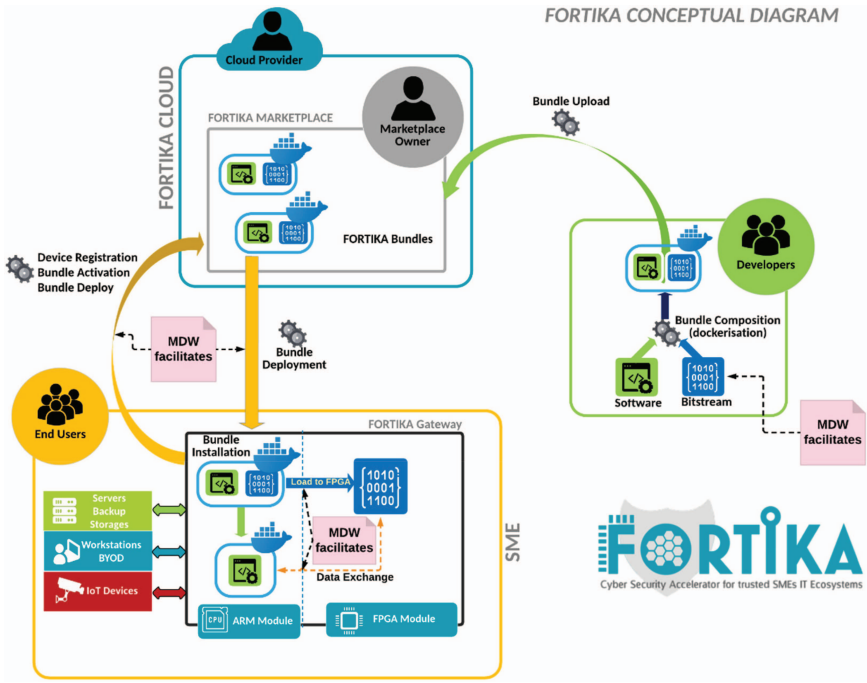

Figure 4.5 Middleware use in FORTIKA.

to interact with the FORTIKA marketplace, whereas the last alleviates the development of accelerated security bundles by hiding the complexity of HW design and configuration from the FORTIKA SW developers. As Figure 4.6 indicates, the first two components are deployed in the FORTIKA Accelerator (GW), whereas the last one is currently deployed in a Virtual Machine located at FINT's cloud infrastructure. So far, the Synthesis engine and the GW's MDW components (SBH and LwM2M client) do not have any interaction as their activities are under different scopes.

Developing applications for the FPGA requires knowledge of the HW platform and its specifics, something that can discourage SW developers from building applications for the FORTIKA accelerator. In the project's context, we tackle this issue by exploiting the fact that the FPGA application development is divided in two phases, namely the Front-End design and the Back-End design [8]. For the Front-End design phase, the FORTIKA developers are using High Level Synthesis tools (i.e. Vivado HLS suite [3]) (Figure 4.7) which allows them to write their FPGA applications in high level languages,

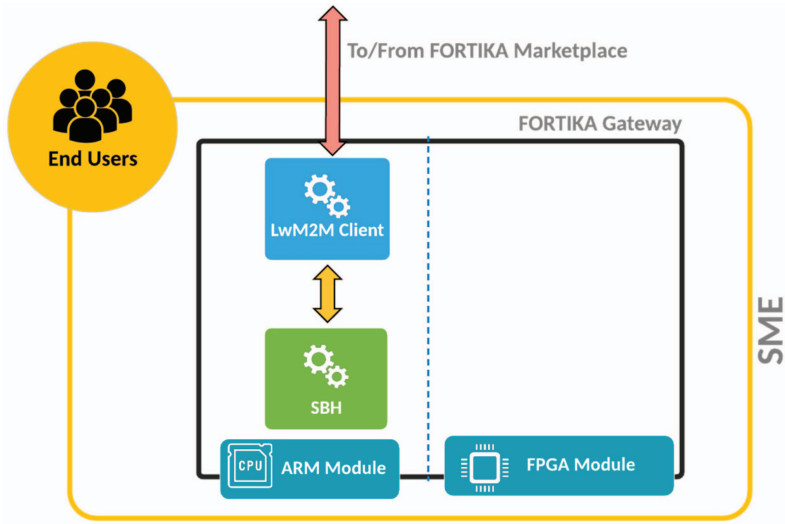

Figure 4.6 SBH and LwM2M client components of the middleware.

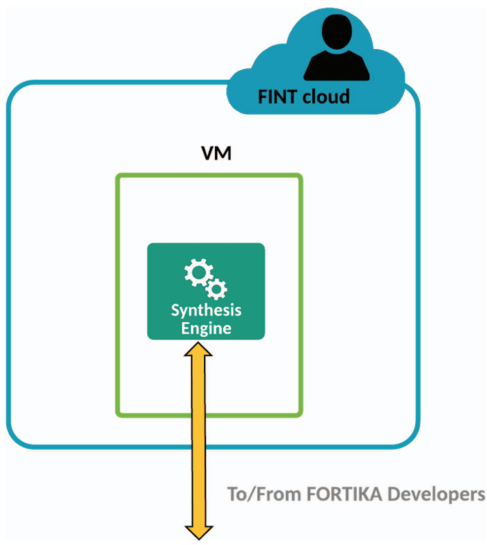

Figure 4.7 Synthesis engine component of the middleware.

such as C/C++, thus avoiding to use low level hardware specific languages (e.g. VHDL) that require knowledge of the HW specifics. After writing their code, the developers can use Vivado HLS (Figure 4.8) to produce artefacts

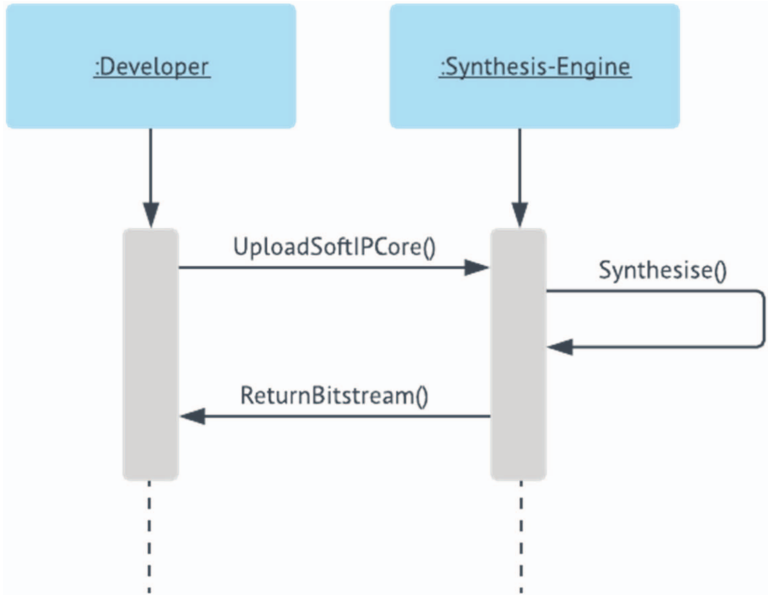

Figure 4.8 Synthesis sequence steps.

that are known as Soft IP (Intellectual Property) cores. These IP cores are used in the Back-End design phase for producing the final bitstreams, that can run on the actual FPGA; however, the Back-End design phase requires the knowledge of specific parameters of the used HW design, thus making it a hard task for the standard SW engineers; therefore, it is this design phase that the FORTIKA MDW aims to facilitate by providing a service that takes as input the produced soft IP core, runs the low-level synthesis (process of the Back-End design phase), and then returns to the developers the final bitstream. In this context, the following diagram depicts, the sequence of steps that are followed from the Synthesis Engine for implementing this task.

The *UploadSoftIPcore()* represents the function that allows developers to upload the produced soft IP cores to the Synthesis Engine. Currently, these IP cores are received via email, however at the next versions of the MDW the cores will be uploaded via a web form; this web form is planned to be provided from the Marketplace dashboard. The *Synthesise()* function, performs the low-level synthesis that produces the final bitstream. The *ReturnBitStream()* function, represents the push of the synthesised bitstream to the developer.

4.3.2 Fortika Marketplace

To facilitate competition and support different value chain configurations, a novel Marketplace Platform is introduced, allowing FORTIKA users to interact with Service Providers and multiple third-party Security Function Developers, for selecting the best service bundle that suits their needs. For this reason, the Marketplace incorporates a prototype that aims to introduce and promote a novel market field for security services, introducing new business-cases and considerably expanding market opportunities by attracting new entrants to the cyber-security market. SMEs and academia can leverage the FORTIKA architecture by developing innovative cutting-edge Security Functions, that can be included in the Function Store, and rapidly introduced to the market, thus avoiding the delay and risk of hardware integration and prototyping. By utilizing a common web-based graphical user interface, the Marketplace constitutes the environment where customers can:

- Place their requests for FORTIKA services and declare their requirements for the corresponding security functions
- Receive offerings and make the appropriate selections, considering the offered Service Level Agreements (SLAs)
- Monitor the status of the established security services and associated security functions, as well as perform, according to their rights, management operations on them (Service monitoring and management will be enabled via a graphical Service Dashboard to be implemented)

The overall concept for security functions trading, deployment and management within the Marketplace is depicted in Figure 4.9, where third-party Security Function developers (1) advertise their available virtual security appliances and users may acquire them for customized service creation/utilization. More specifically, users' requests (2) are received via the Brokerage Module as part of the Marketplace Platform, which is responsible for a) analysing their requirements, b) matching the analysis results with the available resources, maintained by the "Management & Orchestration" module along with the Security Functions aggregated at the Store (4), and c) initiating an auction process for all valid solutions under various merchandise policies and the available SLA models. Upon successful SLA establishment and Functions trading, the Orchestration module deploys the Security Function onto the underlying infrastructure (5), maintaining its control, customization and administration.

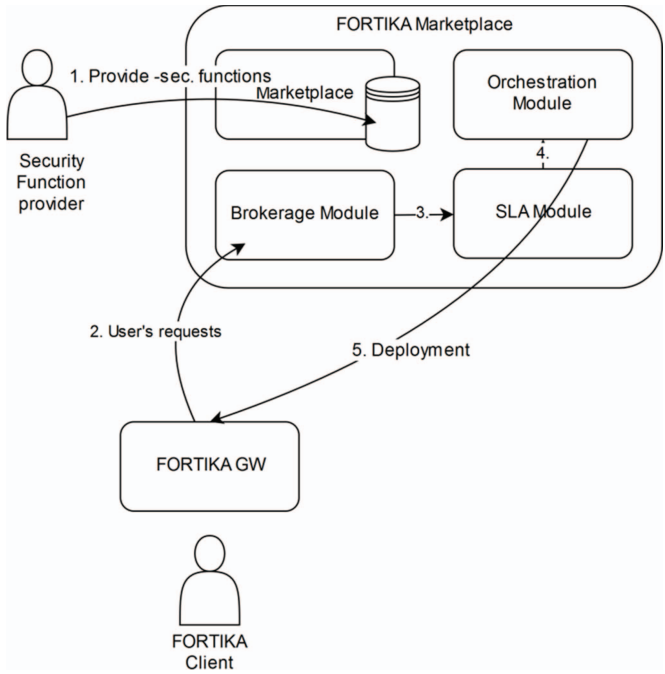

Figure 4.9 Process of deployment and management within FORTIKA marketplace.

To carry out Security Functions discovery provided by third-party developers, similarity-based algorithms such as Nearest Neighbour will be exploited by the Brokerage module to perform service matching. To speed up this process, FORTIKA will study and identify the most appropriate data structures for establishing a competent resource and service description schema for Security Functions matching the brokerage. A principal target is to identify mandatory and optional fields within the schema so as to allow a configurable degree of exposure of resources and services, associated Security Functions and SLAs to all involved actors, according to the confidentiality requirements of each. The integration of the FORTIKA Middleware appliance in existing networks requires seamless connectivity, according to usability and automation standards and guidelines. The appliance will integrate an OpenFlow Ethernet switch with physical Ethernet ports routing and security capabilities (firewall, IPS, IPSec). The appliance will also provide the required processing and storage to enable applications available through

FORTIKA Marketplace to be locally deployed but orchestrated according to rules computed in the cloud. The FORTIKA Marketplace will enable service providers to deploy and promote integrated security services through a web-based user-friendly interface with personalization features. Depending on the service design requirements, the FORTIKA Marketplace will be deployed in the cloud. Deployment of the Marketplace is not limited to public or private cloud. Due to the dynamical deployment mechanisms leveraging tools like Ansible and Docker, and the use of standards (TOSCA) for the services definition, FORTIKA consortium is not limited to any type of cloud resources.

FORTIKA Appliances (Virtual or Physical) will be managed through a FORTIKA-specific management network, using a personalized cloud service. For this reason, an integrated management platform will be deployed which will offer a consistent and unique administrator front end, for both the Middleware appliance configuration as well as installed modules configuration and management. The administrator front end, will allow management of the Security Functions' lifecycle.

Finally, the connection of FORTIKA Middleware appliances with the orchestrator in the cloud, is a critical point since protecting the integrity and confidentiality of data traveling in the fog area is crucial for middleware adoption and end-user trust to FORTIKA. For this reason, FORTIKA Middleware and FORTIKA cloud services communicate over secure channel leveraging LWM2M protocol. This is the back-channel used for management of the FORTIKA Appliance with the running Middleware.

4.4 Indicative FORTIKA Bundles

4.4.1 Attribute-based Access Control (ABAC)

Access control can be defined as a security service, co-existing with others, that aims to limit actions or operations of legitimate entities against requested resources [9]. Over the years, many access-control models have been proposed with the prevalent ones being MAC, DAC and RBAC [9]. In the recent years, information systems are able to interact with the environment, the context, thus a need for a novel approach in controlling access on context-aware information systems arose. As a result, Attribute-Based Access Control (ABAC) was proposed. ABAC policies are able to include attributes of the

subject (requestor), the object (requested resource) and the context (environment). So, in contrary to legacy models, based on identities, a higher level of versatility and control can be achieved.

FORTIKA implements ABAC by providing a cloud-based access control solution which will be highly benefited from the FORTIKA Gateway appliance, to control access to SME ecosystem resources, based on policies that the SME will be able to create and manage.

A system that implements ABAC, consists of the following components [10] (Figure 4.10):

- Policy Administration Point (PAP), that is used to create, store, test and retrieve access control policies. Since the PAP component will be hosted in FORTIKA cloud, a multi-tenant environment will be deployed so that SME administrator users will have access to own organization policies only.
- Policy Information Point (PIP), that retrieves all necessary attributes and authorization data required by PDP in order to reach an access control decision. PIP in FORTIKA is implemented twofold both in the cloud and in the fog, since attribute values are collected from both the cloud and from SME premises.
- Policy Decision Point (PDP) that evaluates access requests against policies so that access control decision is computed.

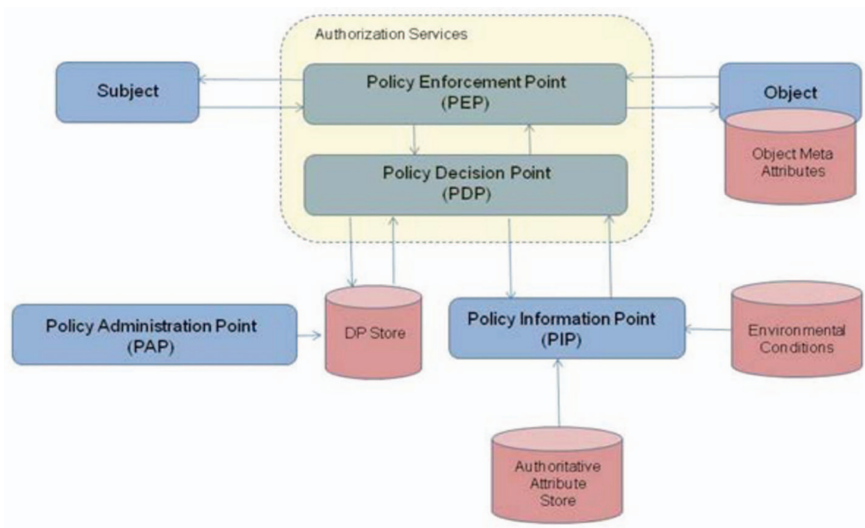

Figure 4.10 ABAC components [10].

- Policy Enforcement Point (PEP) which is the component where an access control request is generated and access decision is enforced.

The Fortika ABAC service is designed as a three-layered approach (Figure 4.11). In terms of component placing and communication architecture, the PIP and PAP components, as well as the related Policy Repository, will be deployed in the cloud (ABAC.Cloud). This will allow for rapid policy replication in case of multi-site SMEs and, additionally, will permit for replacing an on premise FORTIKA appliance without any prior consideration for existing attributes and policies. Moreover, cloud can provide adequate processing and storage resources to create a user-friendly administration environment.

On the other hand, to avoid any issues with network latency or network unavailability [11], the PDP component will be held in the fog area (ABAC.fog). More specifically, PDP will be held in FORTIKA's physical or virtual appliance hosted in SME premises, thus accelerating decision making. Additionally, to better support contextual attributes, a local PIP along with a local attribute repository (currently labelled NA-PIP) will accompany PDP and communicate with cloud PIP to exchange attribute information.

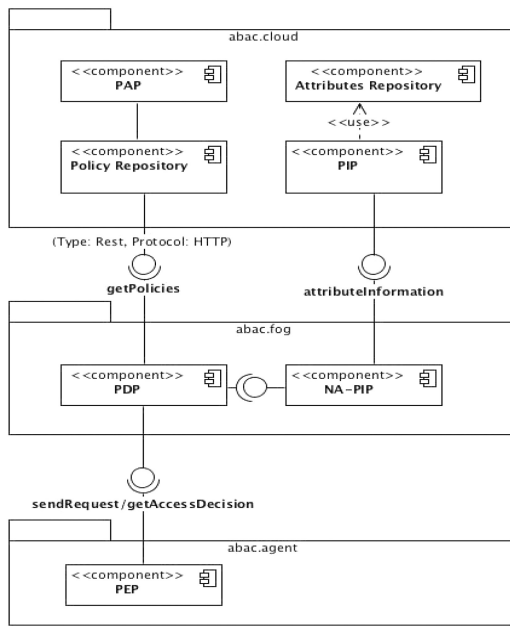

Figure 4.11 ABAC layered approach.

Finally, the PEP component will be initially integrated into a prototype agent for client devices. Nevertheless, ABAC solution will provide the appropriate API, for other compatible PEP components to be able to utilize FORTIKA's ABAC service.

FORTIKA ABAC implements the XACML framework [12] and is based exclusively on open-source technologies, developed with Java and Java EE using Maven. ABAC.Cloud is based on WSO2 Identity Server which is licensed under Apache 2.0 license, whereas ABAC.fog is based on Balana XACML and has been developed to provide a RESTful API to PEPs. The API exposes services according to OASIS REST Profile for XACML 3.0 version 1.0 [13]. This enables potentially any vendor or integrator to utilize FORTIKA ABAC.fog and consume authorization services, constituting FORTIKA ABAC an Authorization as a Service (AuthZaaS) offering.

4.5 Social Engineering Attack Recognition Service (SEARS)

Social engineering attacks are usually an important step in the planning and execution of many other types of cyber-attacks. The term 'social engineering' refers to physiological, emotional and intellectual manipulation of people into performing actions or revealing confidential information. As defined in [14], social engineering is: "*a deceptive process whereby crackers 'engineer' or design a social situation to trick others into allowing them access an otherwise closed network, or into believing a reality that does not exist.*"

The increased usage of electronic communication tools (email, instant messaging, etc.) in enterprise environments results in the creation of new attack vectors for social engineers. However, a successful social engineering attack could result in a compromised SME's information system. Thus, several attempts have been made in the research field to provide technical means for detecting such attacks in early stages. Works that are near to a prototyping level are SEDA [15] and SEADM [16]. Furthermore, interesting efforts that are still under development in the research laboratory are [17] and [18].

Social Engineering Attack Recognition System (SEARS) will operate in the application layer and will be able to compute communication risk and therefore prevent personal or corporate data leakage by raising alerts to the employees when the chat conversation reaches a specific risk threshold [19]. SEARS is a collection of autonomous services that collaborate with

each other through technology-agnostic messaging protocols, either point-to-point or asynchronously. The development of SEARS components follows the microservices design approach. Namely, each component is consisted of a number of independent microservices that serve distinct functionalities of the whole system.

SEARS components will be placed in the three layers of FORTIKA's architecture, as follows:

Client layer:
The SEARS Agent (SEARS.agent) is a service that monitors, captures and pre-processes an employee's social media communications. It is also capable of receiving the total risk value and alerting the user for possible social engineering attack attempts. SEARS.Agent is deployed on end-user's device in a form of a docker container or as local service and continuously monitors and captures an employee's social media communications. SEARS users are registered SME employees as interlocutors (e.g. working on live chat service) or corporate IT administrators (Figure 4.12).

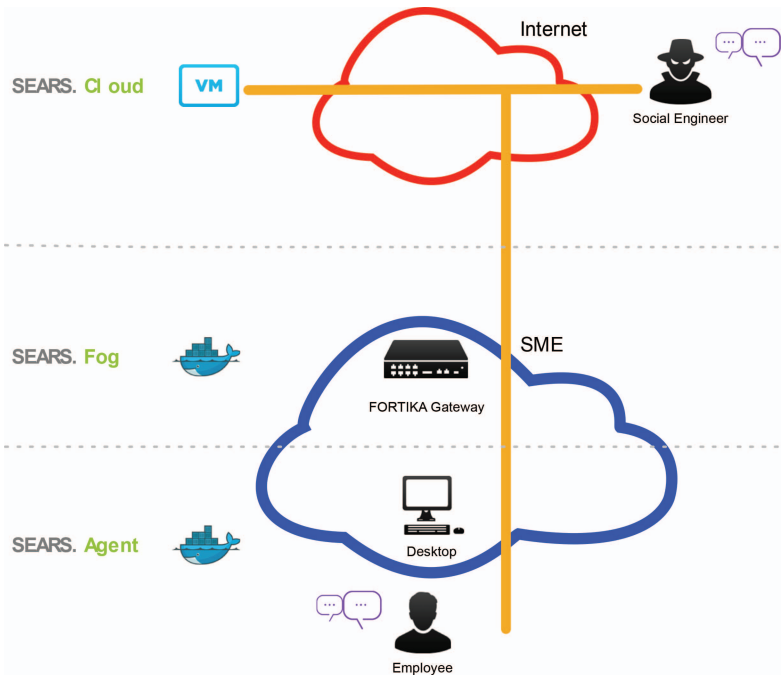

Figure 4.12 SEARS architecture.

Fog layer:
SEARS components in the fog area (SEARS.fog) will be deployed in FORTIKA physical or virtual appliance, hosted in SME premises. SERS.fog receives the captured data and stores it (Detection Storage component) locally for further pre-processing (Pre-processing component), using Natural Language Processing techniques. The pre-processed data is then anonymized and sent to the cloud (SEARS.cloud). The Detection Engine receives the particular risk values from the SEARS.cloud and then calculates the total Social Engineering Risk value, stores it in the Detection Repository and sends it to the SEARS.client. SEARS.Fog is deployed on FORTIKA Gateway in the form of a docker container.

Cloud layer:
The pre-processed data received from the SEARS.fog is stored in the SEARS Storage component of SEARS.cloud, in order to be used by the Risk Estimation component to calculate values of particular risks. These values are then sent to the SEARS.fog. The following components are part of SEARS.Cloud core functionality and implemented by several microservices.

- **Document Classification (DC)**:
  Text dialogue, in the form of an anonymized TF-IDF matrix, is processed and classified as dangerous or not. The real text dialogue is processed at the SEARS.Agent, where an anonymized frequency vector is delivered to SEARS.Cloud, where the classification takes place.
- **Personality Recognition (PR)**:
  Each of the interlocutors is being classified based on his/her writings. The processing/classification takes place at the SEARS.Cloud using the previous anonymized frequency vector.
- **User History (UH)**:
  Each previous text chat between the two specific interlocutors is represented as probability (decimal number) and it is stored at SEARS.Cloud.
- **Exposure Time (ET)**:
  The duration of an employee's online presence is being depicted as a decimal number stored at SEARS.Cloud

SEARS offers the ability to communicate the estimated risk values to other modules of FORTIKA. The outgoing information is provided using a standard HTTP POST method. All data is encoded using the JavaScript Object Notation (JSON) format and follow the structure of SEARS Output JSON Schema. Moreover, all data transfers are being carried out using REST

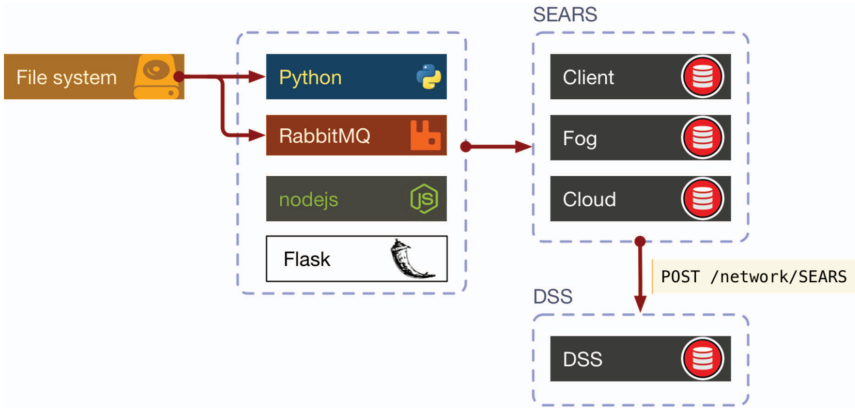

Figure 4.13 SEARS conceptual design.

APIs through HTTPS protocol thus the communication channel cannot be compromised. The SEARS conceptual design as a whole is presented in Figure 4.13.

SIEM

The Security Information and Event Management System (SIEM), is a solution able to analyse information and events collected at different levels of the monitored system in order to discover possible ongoing attacks, or anomalous situations. FORTIKA includes a customized SIEM solution, able to deal with specificities of its different technologies and components.

The network provides real-time traffic data to the SIEM system. The system in turn, forwards the data for processing to both the Anomaly detection and the behavioural analysis components. The Anomaly detection component analyses the data in order to detect anomalies, utilising both automatic anomaly detection algorithms, such as Local Outlier Factor and Bayesian Robust Principal Component Analysis, as well as visual analytics methods, such as k-partite graphs and multi-objective visualizations. The Behavioural analysis component processes the network data in order to identify abnormal traffic patterns that may indicate that a malicious event such as a DDoS attack is in progress. The output from both components is then passed to the Visualization component for presentation to the user, or to the Hypothesis Formulation component. The Hypothesis Formulation component performs a statistical analysis of the output data of the Anomaly detection and the

Behavioural analysis component, through a series of hypotheses in order to determine whether these data express a normal or a usual traffic pattern or behaviour. The analysis data can be subsequently fed back to the Anomaly detection and the behavioural analysis components for further analysis.

4.6 Conclusion

FORTIKA architecture proposes a hybrid (hardware software) cybersecurity solution suitable for micro, small and medium-sized enterprises allowing them to continuously integrate novel cyber-security technologies and thus reinforce their position and overall reputation in the European market. Concluding, this paper introduced a novel architecture that aims at reshaping the cyber-security landscape in order to provide an end-user-friendly solution targeting towards moving security near the network edge. This architecture is based upon two pillars: A near-the-edge security-accelerator, which is able to "accelerate" security in the place where the problem is formulated, and a Cloud Marketplace which provides a unified portal for enabling security for FORTIKA end-users. The preliminary evaluation of the presented work illustrated that users (SMEs) can identify which cyber-security solutions are suitable for their enterprises and seamlessly deploy them on their infrastructures (FORTIKA gateway). Additionally, security-solutions' developers/providers can easily offer their services through the FORTIKA marketplace, which also allows them to interact with users and offer custom-tailored cyber-security solutions (brokerage), thus extending their marketing opportunities. The presented work is an ongoing EU-funded Horizon 2020 project, and currently runs the second year of development. Several complex and intuitive features are to be developed in the near future and thus more detailed and elaborate reporting of the work will be presented through publications and public workshops, as well as from the project's social media accounts (Facebook, Twitter, YouTube, etc.).

Acknowledgment

This work has received funding from the European Union's Horizon 2020 Framework Programme for Research and Innovation, with Title H2020-FORTIKA "cyber-security Accelerator for trusted SMEs IT Ecosystem" under grant agreement no. 740690.

References

[1] H. Madsen, G. Albeanu, B. Burtschy, and F. Popentiu-Vladicescu, "Reliability in the utility computing era: Towards reliable fog computing," in International Conference on Systems, Signals, and Image Processing. IEEE, jul 2013, pp. 43–46. [Online]. Available: http://ieeexplore.ieee.org/document/6623445/

[2] Y. Nikoloudakis, S. Panagiotakis, E. Markakis, E. Pallis, G. Mastorakis, C. X. Mavromoustakis, and C. Dobre, "A Fog-Based Emergency System for Smart Enhanced Living Environments," IEEE Cloud Computing, vol. 3, no. 6, pp. 54–62, nov 2016. [Online]. Available: http://ieeexplore.ieee.org/document/7802535/

[3] C. Dobre, C. X. Mavromoustakis, N. M. Garcia, G. Mastorakis, and R. I. Goleva, "Introduction to the AAL and ELE Systems," Ambient Assisted Living and Enhanced Living Environments: Principles, Technologies and Control, pp. 1–16, jan 2016. [Online]. Available: https://www.sciencedirect.com/science/article/pii/B9780128051955000016 Books (IDEA/IGI, Springer and Elsevier).

[4] B. McKenna, "Symantec's Thompson pronounces old style IT security dead," Network Security, vol. 2005, no. 2, pp. 1–3, 2005. [Online]. Available: http://linkinghub.elsevier.com/retrieve/pii/S135348580500, 1947."

[5] Cyber Security Experts & Solution Providers |FireEye." [Online]. Available: https://www.fireeye.com/

[6] "ENISA Threat Landscape Report 2016 — ENISA." [Online]. Available: https://www.enisa.europa.eu/publications/enisa-threat-landscape-report-2016. [Accessed: 20-Nov-2017].

[7] E. O. Yeboah-Boateng, Cyber-Security Challenges with SMEs in Developing Economies: Issues of Confidentiality, Integrity & Availability (CIA). Institut for Elektroniske Systemer, Aalborg Universitet, 2013.

[8] E. K. Markakis, K. Karras, A. Sideris, G. Alexiou, and E. Pallis, "Computing, Caching, and Communication at the Edge: The Cornerstone for Building a Versatile 5G Ecosystem," IEEE Communications Magazine, vol. 55, no. 11, pp. 152–157, nov 2017. [Online]. Available: http://ieeexplore.ieee.org/document/8114566/

[9] R. S. Sandhu and P. Samarati, "Access control: principle and practice," IEEE Communications Magazine, vol. 32, no. 9, pp. 40–48, Sep. 1994.

[10] V. C. Hu et al., Guide to Attribute Based Access Control (ABAC) Definition and Considerations (Draft). 2013.
[11] S. Salonikias, I. Mavridis, and D. Gritzalis, "Access Control Issues in Utilizing Fog Computing for Transport Infrastructure," in Critical Information Infrastructures Security, 2016, pp. 15–26.
[12] "eXtensible Access Control Markup Language (XACML) Version 3.0." [Online]. Available: http://docs.oasis-open.org/xacml/3.0/xacml-3.0-core-spec-os-en.html. [Accessed: 18-Jan-2019].
[13] "REST Profile of XACML v3.0 Version 1.0." [Online]. Available: http://docs.oasis-open.org/xacml/xacml-rest/v1.0/csprd03/xacml-rest-v1.0-csprd03.html. [Accessed: 18-Jan-2019].
[14] B. H. Schell, B. Schell, and C. Martin, Webster's New World Hacker Dictionary. John Wiley & Sons, 2006.
[15] M. D. Hoeschele and M. K. Rogers, "CERIAS Tech Report 2005–19 Detecting Social Engineering," 2004.
[16] F. Mouton, L. Leenen, and H. S. Venter, "Social Engineering Attack Detection Model: SEADMv2," 2015, pp. 216–223.
[17] R. Bhakta and I. G. Harris, "Semantic analysis of dialogs to detect social engineering attacks," in Semantic Computing (ICSC), 2015 IEEE International Conference on, 2015, pp. 424–427.
[18] S. Uebelacker and S. Quiel, "The Social Engineering Personality Framework," 2014, pp. 24–30.
[19] N. Tsinganos, G. Sakellariou, P. Fouliras, and I. Mavridis, "Towards an Automated Recognition System for Chat-based Social Engineering Attacks in Enterprise Environments," in Proceedings of the 13th International Conference on Availability, Reliability and Security – ARES 2018, Hamburg, Germany, 2018, pp. 1–10.
[20] C. Liu, Y. Mao, J. E. Van der Merwe, and M. F. Fernández, "Cloud Resource Orchestration: A Data-Centric Approach," in Proceedings of the biennial Conference on Innovative Data Systems Research (CIDR), 2011, pp. 241–248. [Online]. Available: http://www2.research.att.com/maoy/pub/cidr11.pdf
[21] A. Dubey and D. Wagle, "Delivering software as a service," The McKinsey Quarterly, vol. 6, no. May, pp. 1–12, 2007. [Online]. Available: http://www.pocsolutions.net/Delivering_software_as_a_service.pdf
[22] K. Lane, "Overview Of The Backend as a Service (BaaS) Space," 2013. [Online]. Available: http://www.integrove.com/wp-content/uploads/2014/11/api-evangelist-baas-whitepaper.pdf

[23] S. A. Fahmy, K. Vipin, and S. Shreejith, "Virtualized FPGA accelerators for efficient cloud computing," in Proceedings IEEE 7th International Conference on Cloud Computing Technology and Science, CloudCom 2015. IEEE, nov 2016, pp. 430–435. [Online]. Available: http://ieeexplore.ieee.org/document/7396187/

[24] J. A. Williams, A. S. Dawood, and S. J. Visser, "FPGA-based cloud detection for real-time onboard remote sensing," in Proceedings 2002 IEEE International Conference on FieId-Programmable Technology, FPT 2002. IEEE, 2002, pp. 110–116. [Online]. Available: http://ieeexplore.ieee.org/document/1188671/

[25] S. Byma, J. G. Steffan, H. Bannazadeh, A. Leon-Garcia, and P. Chow, "FPGAs in the cloud: Booting virtualized hardware accelerators with OpenStack," in Proceedings 2014 IEEE 22nd International Symposium on Field-Programmable Custom Computing Machines, FCCM 2014. IEEE, may 2014, pp. 109–116. [Online]. Available: http://ieeexplore.ieee.org/document/6861604/

[26] L. Xu, W. Shi, and T. Suh, "PFC: Privacy preserving FPGA cloud A case study of MapReduce," in IEEE International Conference on Cloud Computing, CLOUD. IEEE, jun 2014, pp. 280–287. [Online]. Available: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6973752

[27] K. Karras, O. Kipouridis, N. Zotos, E. Markakis, and G. Bogdos, "A Cloud Acceleration Platform for Edge and Cloud," in EnESCE: Workshop on Energy-efficient Servers for Cloud and Edge Computing, 2017. [Online]. Available: https://www.researchgate.net/publication/313236609

[28] Nikoloudakis, Y, Pallis, E, Mastorakis, G, Mavromoustakis, CX, Skianis, C & Markakis, EK 2019, 'Vulnerability assessment as a service for fog-centric ICT ecosystems: A healthcare use case' Peer-to-Peer Networking and Applications. https://doi.org/10.1007/s12083-019-0716-y