R. Rios, R. Roman, J. A. Onieva, and J. Lopez, "From Smog to Fog: A Security Perspective", *2nd IEEE International Conference on Fog and Edge Mobile Computing (FMEC 2017)*, pp. 56-61, 2017.
http://doi.org/10.1109/FMEC.2017.7946408
NICS Lab. Publications: https://www.nics.uma.es/publications

From Smog to Fog: A Security Perspective

Ruben Rios, Rodrigo Roman, Jose A. Onieva and Javier Lopez
Network, Information, and Computer Security (NICS) Lab
Universidad de Málaga
Campus de Teatinos s/n, 29071. Málaga, Spain
{ruben,roman,onieva,jlm}@lcc.uma.es

***Abstract*— Cloud computing has some major limitations that hinder its application to some specific scenarios (e.g., Industrial IoT and remote surgery) where there are particularly stringent requirements, such as extremely low latency. Fog computing is a specialization of the Cloud that promises to overcome the aforementioned limitations by bringing the Cloud closer to end-users. Despite its potential benefits, Fog Computing is still a developing paradigm which demands further research, especially on security and privacy aspects. This is precisely the focus of this paper: to make evident the urgent need for security mechanisms in Fog computing, as well as to present a research strategy that is being undertaken within the SMOG project, in order to enable a trustworthy and resilient Fog ecosystem.**

I. Introduction

Cloud computing has long been regarded as one of the cornerstones of future Internet systems, mainly due to its ability to accommodate the computational and storage power that mobile and other computing devices, such as sensor nodes, lack [1]. However, the emergence of novel application scenarios (e.g., Cyber-Physical Systems [2]) with new requirements has called into question the ability of Cloud computing to be the one-size-fits-all solution.

The main drawbacks of Cloud computing come from its centralized nature, which prevents it from providing reasonable response times, mobility support or being aware of the context of the users [3]. This has motivated the emergence of Edge Computing paradigms, similar to the Cloud in the sense that they provide computation and storage capabilities but are placed closer to end-users, where data are being generated. This does not necessarily mean getting rid of the Cloud but instead, having a multi-tier architecture where some operations can be performed in a timely fashion by handling them in nearby edge data centers rather than relying on a distant cloud server at the backbone. Put simply, not every piece of data needs to go to the Cloud.

Fog Computing [4] is one example of the Edge Computing ecosystem[1]. As shown in Fig. 1, fog devices coexist with Cloud servers and end-user devices in a three-tier architecture. Here, fog nodes are heterogeneous devices geographically distributed, which offer services to the devices in their local environment. For example, in the figure we can observe that a roadside fog node can be used by autonomous vehicles

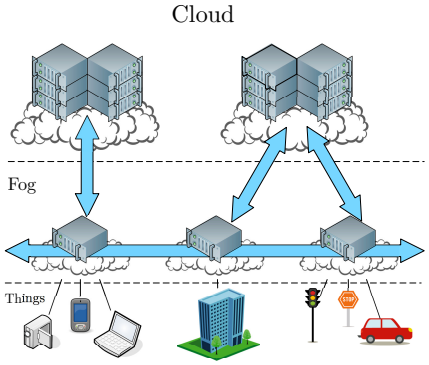

Fig. 1. Fog Computing Vision

(i.e., driverless cars) to make real-time decisions based on the traffic conditions in their vicinity. Precisely, fog nodes can communicate with one another in order to exchange information and services. Similarly, fog nodes have interfaces with centralized data centers in the Cloud to enable global coordination and other highly resource-intensive operations. Nevertheless, fog nodes can also operate autonomously.

It is important to stress that the Fog infrastructure is not monolithic, but rather it is an ecosystem where multiple infrastructure providers coexist and cooperate with each other [6]. In such an environment, fog nodes and the services deployed in the cloud can belong to multiple actors: from private users to mobile network operators. This cooperative and federated ecosystem allows users and service providers to benefit from the advantages of the Fog (mobility, knowledge of context, etc.) regardless of their location. In the example scenario described above, the driverless cars could make use of any suitable roadside fog nodes to gain local awareness of their vicinity. Besides, these roadside fog nodes could be used to provide support for other services, such as providing computationally intensive operations to constrained devices, or managing and processing data generated by security cameras, smart meters, and others.

One can clearly see from the aforementioned scenarios that a number of challenges need to be met before Fog ecosystems become a reality. Among these challenges, security is of utmost importance [7]. Without proper security mechanisms, the potential benefits of this paradigm will be tarnished by the disastrous consequences and damage that attackers may bring about. For example, in the roadside

[1]There are other Edge Computing paradigms, such as Multi-Access (formerly Mobile) Edge Computing (MEC) and Mobile Cloud Computing (MCC), with virtually no difference between them except for the underlying infrastructure and the entity managing them [5].

fog scenario outlined above, an attacker may launch denial of service attacks or even physically destroy part of the infrastructure to prevent the vehicles from making timely decisions in the case of emergencies. However, although Fog computing security is still in its infancy, most of the security solutions developed for this context do not truly take into consideration the complexity of an interconnected environment, as they mainly focus mostly on protecting an isolated set of fog nodes belonging to a single administrative (trust) domain (i.e. managed by one single owner).

Therefore, it is extremely necessary to explore what mechanisms could enable a resilient and secure Fog ecosystem. This is precisely the main goal of this paper, to present and describe the necessary steps for protecting Fog infrastructures that take into account their inherent complexity. This research strategy has been developed and is currently being followed within the SMOG (Security Mechanisms for Fog Computing) project. This project consists of two complementary subprojects: SMOG-CORE, devoted to the development of security services for the infrastructure, and SMOG-DEV[2], responsible for the development of solutions to enable the secure interaction between end-users and the infrastructure. Although, end-user devices are a relevant part of the ecosystem, they are not described in this paper.

The rest of this paper is organized as follows: in Section II, we introduce the main threats that affect Fog ecosystems. Section III provides an overview of the state of the art on Fog computing security mechanisms. Section IV analyzes the specific security needs of this ecosystem, and then describes the fundamental and advanced security services that are needed to protect Fog infrastructures. Then, Section V discusses the suitability of the proposed security services. Finally, Section VI details the conclusions of the paper.

II. Main Threats

Security in Fog computing is a challenging task for various reasons. First, the Fog is a semi-distributed environment consisting of heterogeneous devices belonging to different trust domains and managed by entities with different technological backgrounds. Second, it conjugates a number of supporting technologies and scenarios with their own security issues, including various communication and connectivity protocols (e.g., 5G, Wifi, Zigbee), virtualization technologies (e.g., virtual machines, containers), and so on. These enabling technologies must not only be secure in isolation but also when integrated with other technologies within the Fog. Third, the security mechanisms must be in harmony with the intrinsic requirements of the Fog paradigm, thus precluding the use, in certain cases, of computationally intensive protocols and algorithms: there is no use in providing perfect security if usability or latency are not improved compared to Cloud environments. Fourth, security design needs to integrate business relationships among multiple stakeholders across different domains.

[2]This part of the project is being conducted by Carlos III University of Madrid (UC3M) in Spain.

Security threats and attacks can appear at any level of the infrastructure with varying severity. Attackers targeting end-user devices or attacks launched from these devices have a limited impact, usually only partially affecting the local environment. On the other hand, by attacking the infrastructure at the edge (i.e., fog nodes), the impact of the attacker is broader as he gains full control over a local environment. Finally, some attacks can target the core infrastructures and threats exist in the interactions between fog nodes and centralized cloud services.

In the following, we describe the major threats that might affect Fog environments at different levels of the infrastructure. Not surprisingly, most of these threats also affect some traditional interconnected systems and data centers, but they present new nuances and their impact will differ:

a) ***Denial of Service*** *(DoS):* these attacks are intended to prevent legitimate users from accessing the services provided by Fog infrastructures. Therefore, these threats to availability can occur at different levels. At the end-user level, an attacker can jam the wireless communication channel or exhaust their resources (e.g., with malware) to prevent certain users from communicating with the infrastructure. At a higher level, the attacker can physically destroy a fog node, thus disrupting all the services being offered at the local level. The attacker can also launch attacks to deplete the resources of fog nodes as a means to prevent it from allocating new resources for other users or delaying its responses. Successfully taking down the services at the network core is far more difficult given the resources available to cloud service providers. Notwithstanding, the attacker can target the infrastructure supporting these services (such as in the DDoS to Dyn ISP [8]) or rely on a larger pool of attackers to launch a distributed DoS.

b) ***Data leakage****:* these are attacks on confidentiality and user privacy. An attacker located in the vicinity of end-users can eavesdrop on the communication channel in order to extract meaningful data from the packet contents but also from the packet headers, such as with whom they are interacting. By exploiting the features of the wireless channel, the attacker can also triangulate the location of a particular user, thus posing a threat to location privacy. The attacks mentioned so far are usually performed by external adversaries but internal adversaries can also do harm in this respect. Consider, for example, the case of honest-but-curious services providers. At the edge, service providers have access to all the information being stored and processed in the fog nodes. Fortunately, fog nodes have only a partial view of the network, that is, data being generated at a particular location. Moreover, cloud data centers will have a globally partial view of the data. In other words, Cloud providers have access to the data generated by all collaborator fog nodes, although the data will presumably be processed at each local environment prior to its transmission. This is indeed good from a privacy point of view. Beyond honest-but-curious adversaries, some attackers might be able to compromise some of these nodes, including features provided by the virtualization environment, and thus be in the same position

as the service providers.

*c) **Manipulation**:* these are attacks against the integrity of the communications and the services deployed by the Fog. First, an attacker can manipulate the traffic traversing the network by, for example, modifying or replaying some packets. Moreover, malicious or compromised end-user devices can report incorrect values or fake data in order to disrupt services or calculations. For example, compromised industrial sensors can report false data about the status of the system to cause some harm, like in the Stuxnet case [9], where compromised PLCs (programmable logic controllers) were capable of changing the spinning speed of centrifuges in a nuclear plant while continuing to report normal values to prevent the attack from being detected and stopped. Also, an adversary can manipulate the services deployed by the infrastructure once he has gained privileged access (e.g., by exploiting a vulnerability in the software) to the fog nodes or the cloud. The severity of the attack very much depends on the level at which services are manipulated. The services provided by fog nodes are usually more time-critical but have only a local impact, while manipulation of services at the core may affect the whole system, mainly in the control plane. Nevertheless, in principle it should be more difficult to manipulate services at higher levels than at lower levels because they are managed by experienced cloud providers.

*d) **Impersonation**:* these are also attacks on integrity, but in this case they target the identity of different elements (actors, software, and hardware) of the system. Attackers may try to impersonate other users in order to gain access to services for which they are not authorized. Besides impersonating users of the system, an attacker can also try to deploy rogue infrastructure elements. A sufficiently powerful adversary can pose as a fog node in order to gain access to and possibly manipulate any information being transmitted by the devices. In the case that the adversary is not powerful enough, he can still try to perform a man-in-the-middle attack by placing himself between the end-user and the fog node, obtaining similar results. The main limitation of this type of attacks is that the response times of the communications may vary considerably depending on the capabilities of the adversary. Similar to other attacks, impersonating infrastructure elements at the backbone may be more difficult because they usually deploy better protection mechanisms. However, it is important to take them into consideration since a successful attack at this level may have a severe impact on the whole system.

Given the threats mentioned above, the need to develop specific security services for the Fog is indisputable. In the following section, we inspect the literature in order to determine the level of maturity of research in Fog security.

III. Existing Research in Fog Security

Despite the existing body of research in security-related topics in analogous paradigms, especially in (Mobile) Cloud Computing [5], little work has been done to specifically protect Fog infrastructures. So far, most of the research in the area has concentrated on authentication and access control, while other services remain mostly unexplored.

Stojmenovic et al. [10] propose two preliminary techniques that can be used for authentication and authorization in Fog environments without an online cloud server. In the first approach, the Cloud service provider sends its own credentials to the Fog, so that it can later authenticate the user by itself. Due to the shortcomings of this scheme, the authors introduce another scheme based on a special type of Attribute-Based Encryption (ABE) called ciphertext-police ABE, where ciphertexts are associated with access policies and keys with attributes. Zhang et al. [11] also elaborate on the idea of using CP-ABE to enable flexible data access control. The proposed scheme reduces the latency and overhead introduced by encryption and decryption by securely outsourcing costly operations to fog nodes. A clear limitation of these solutions is that they consider only Fog environments that depend on a single infrastructure provider.

Another authentication scheme for Fog computing is proposed in [12]. The scheme enables authentication between end-users and fog nodes. To that end, users are required to hold a master key that they use to mutually authenticate with any fog node associated with a particular cloud server. More precisely, the protocol consists of three phases. At the initialization phase, the Cloud sends each fog node a unique identifier, signed with its private key. During registration, end-users must contact the cloud server to get their master key. In parallel, the Cloud sends each fog server the identity of the registered user together with an associated key, which is derived from the user's master key. The authentication phase is a typical challenge response protocol based on very few symmetric key operations and one hash function. Again, this solution is useful only for fog nodes belonging to a single administrative domain.

A policy-based security framework is presented in [13] to support secure sharing, collaboration and data reuse in Fog environments. Similar to the approach by Stojmenovic et al., the proposed policy management framework adopts attribute-based authentication for identifying their users and verify whether the attributes of a user entitle him to access a particular resource or service. The framework consists of a number of modules for defining rules, storing them, and making decisions on user's service requests and data migration among different fog nodes. According to the authors, the modules can be plugged, in real-time and the policy enforcer module can reside either in a fog node, a cloud server or within an end-user device. However, as stated by the authors this is a preliminary framework, which does not consider all the nuances of federated Fog ecosystems.

The authors in [14] concentrate on secure threat information exchange. They emphasize on the need to use standards to define cyber-threat information and to securely exchange these data within the context of the Internet of Things and Fog environments. To that end, they basically adapt an existing intrusion detection system of their own that consists of a number of sensors in various local environments and a central server for correlating data. The sensors use standard STIX

expressions for representing threat and attack information and send it to the central server, which can later be shared with other servers using the TAXII standard. Despite the best efforts of the authors to use standard languages and protocols, this solution has not been devised with full-fledged Fog ecosystems in mind.

In [15], Wang et al. analyze Fog security issues from the perspective of digital forensics. In their analysis, they state that Fog forensics is partially related to Cloud forensics, although it has certain specific challenges due to the nature of the Fog: (i) the preservation of the chain of custody and integrity of the evidences, (ii) the dependence on the service provider(s) for the acquisition of the evidences, (iii) the preservation of privacy of other users in multi-tenant environments. These challenges are more acute in Fog environments due to the geographical distribution of the fog nodes which are possibly managed by stakeholders belonging to different administrative domains with varying technological background. Nonetheless, there is also an advantage to the distribution of fog nodes with respect to the need for less computational resources for managing digital evidences since they are limited to their local context. Unfortunately, no solutions are provided in this respect.

IV. SMOG Research Strategy

It has been shown in Section III that Fog security is an almost unexplored area with very limited results from the research community. Moreover, these isolated solutions have oversimplified the actual vision of the Fog ecosystem, by considering the Fog as a hierarchical infrastructure with a single Cloud service provider on top, rather than a federated system with numerous service providers belonging to different trust domains. If this partial view of the problem continues, it will doubtlessly lead to incomplete security solutions and eventually to attacks. It is therefore necessary to explore which security services are needed in this context – and how to tackle them.

In actual Fog environments with multiple interacting service providers, infrastructures and services, it is paramount to provide a service for the identification of all these entities in order to enable authentication and access control. These services are, in turn, necessary for the establishment of secure communication channels between elements of the Fog ecosystem even if they belong to different security domains. With these basic services, not only is it possible to constitute the secure and federated environment that the Fog promises, but also to provide the fundamental security mechanisms that are needed by the virtualization infrastructure: virtualized services must be identified, they must be able to securely migrate, it is necessary to control which resources they can access, and so on. Moreover, as the Fog is an heterogeneous and semi-distributed environment, it is crucial to deploy situational awareness mechanisms to monitor the status of the infrastructure.

The aforementioned mechanisms provide the tools required to interconnect and deploy services in the Fog. However, these basic services do not consider all the security

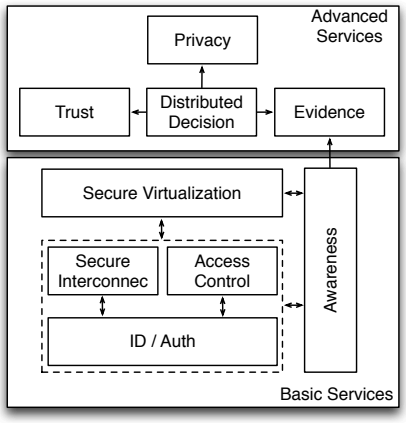

Fig. 2. SMOG Services

considerations that the interaction between Fog and Cloud infrastructures from different service providers might entail. Thus, some advanced security services must be offered to promote the collaboration between the elements of this federated ecosystem. In this respect, the provision of trust management services is essential as some of the interactions in the Fog may occur between previously unknown entities. Another relevant mechanism when data and services are constantly moving to and from different domains is privacy. The Fog must provide privacy services not only to their users but also to the service providers as they are involved in interactions with one another that they might not wish to disclose. Finally, as the Fog will be exposed to threats and attacks it must provide a service for handling digital evidences based on the situational awareness mechanisms described above. This service can also benefit from the provision of distributed consensus mechanisms, such as blockchain, to enable a transparent and verifiable evidence service. This sort of distributed mechanisms can also be used to enhance and support trust and data sharing decisions.

The particular security requirements necessary to protect the Fog arise naturally when the full vision of the Fog is considered. Moreover, the interdependencies between these services (see Fig. 2) require that they are tackled in order, as some mechanisms are required by others. First, it is necessary to develop the fundamental services for the protection of the infrastructure and then build advanced security services for the cooperation and interaction between entities. These are precisely the steps we are taking in the SMOG project and which we show below.

A. Fundamental Security Services

The services described in this section provide the basic mechanisms for enabling a federated and secure Fog infrastructure.

1) Secure interconnection of Fog elements: First of all it is vital to protect the communications. The risk of external manipulation and eavesdropping of the information flows between Fog elements is very high due to the cooperative nature of Fog nodes, their geographical distribution, and the

use of an amalgam of network technologies. Therefore, it is vital to develop services that allow Fog elements to negotiate security parameters as well as establishing credentials in heterogeneous environments, where not only do devices have different capabilities but also multiple service providers exist. To that end, it will be necessary to explore a number of strategies, such as entities federation, for the seamless and secure integration of current protocols and standards.

2) Authentication and Authorization for the Fog: One of the main challenges in Fog environments is to consistently identify and authenticate all the elements within the Fog ecosystem. This includes not only the communicating entities and infrastructures but also the virtual applications and services that are executed in and migrate to and from Fog nodes. Besides authentication, its is paramount to determine which are the privileges of these entities, that is, what actions they are authorized to perform. This problem is more acute when, as stated above, applications can be virtualized, replicated and migrated across domains. Furthermore, this fundamental service needs to factor in users' mobility and need for fast registration with easy-to-use technologies like haptic approaches.

3) Protection of Virtualized Environments: Virtualization is one of the most prominent services that a Fog environment has to provide to unleash the full potential of this paradigm. Thus, it is of evident importance to provide a set of security mechanisms to validate the correct deployment, operation and migration of virtualized services and applications across the Fog. These mechanisms should be able to verify that virtualized services can access only the data they are entitled to, since malicious code may be launched to the Fog nodes, to exploit vulnerabilities or to obtain information from other running services. Similarly, the mechanisms devised should prevent Fog nodes run by honest-but-curious service providers from gaining access to or modifying the results of the virtualized services.

4) Situation Awareness Mechanisms: The Fog must provide mechanisms to monitor the status of fog devices as well as the services deployed in them. This is imperative not only as a means to identify the presence of intruders in the system but also to detect anomalous behaviors caused by flawed devices or software bugs. This will aid not only in the detection of problems in a local context but will also help to detect more complex situations, such as advanced persistent threats. To that end, and given the distributed and heterogeneous nature of the Fog ecosystem, devising a normalized language for securely exchanging information across the Fog is imperative.

B. Advanced Security Services

The services described in this section give the entities of the Fog support for securely interacting and cooperating with one another.

1) Trust Services among Fog entities: The Fog ecosystem is composed of multiple stakeholders from geographically distant locations and presumably with no prior knowledge about each other. Also, certain parts of the infrastructure might be compromised. Therefore, interacting with entities in the Fog may be a risky business especially when these entities are unknown. Therefore, the Fog must offer support for interactions in the presence of uncertainty. This can be achieved by means of trust and reputation services. These services will rely on information about previous interactions and the context surrounding the entities or interacting with it. Again, this information must be exchanged in a normalized way to ensure the interoperability between different trust and reputation models.

2) Distributed Decision Making: Fog systems must implement advanced cryptographic mechanisms to support the execution of distributed processes without depending on an online central authority and, possibly, in the presence of untrustworthy entities. It is therefore necessary to study the applicability of existing mechanisms (e.g., secure multiparty computation, distributed ledger technologies) to the particular requirements of Fog environments. Moreover, the Fog should enable the processing of obfuscated or encrypted data from the local environment and from other Fogs in order to enable secure and distributed data mining.

3) Privacy Support: The proximity of the Fog to the users makes the data handling process extremely sensitive. The data managed by the Fog can be directly associated with the users in the local context and this locality makes it difficult to apply some typical protection mechanisms. Furthermore, IoT and mobile devices will acquire much more sensitive data than in traditional scenarios. This, in conjunction with the aggregation and analytics capacity of the fog, challenges users' privacy.

Still, the infrastructure must provide mechanisms to allow end-users to specify their own privacy requirements as well as mechanisms to enforce them. Thus, the Fog should give users support to determine which mechanisms (e.g., differential privacy [16]) are most suitable to protect their data, and help them to make decisions on data sharing (or data partition) with other Fog elements. It should also offer contextual information to the end-devices, such as the number of entities in their vicinity, in order to facilitate the application of some privacy techniques, like k-anonymity.

4) Digital Evidence Management: Regardless of the implementation of multiple security mechanisms, the complexity and ubiquity of Fog environments make them an attractive target for attackers. Therefore, it is reasonable to retrieve a set of evidences from the services deployed in the Fog, which can provide information about its operation and incidents. There should exist a service capable of managing the evidence. This includes determining potential types of evidence, and building mechanisms for exchanging such information with trustworthy sources. A key aspect is the introduction of mechanisms to ensure the integrity of the system, possibly based on the distributed decision making mechanisms, such as distributed ledgers (i.e., blockchain [17]), which, in turn, enables transparency, verifiability and traceability services to be defined.

TABLE I
THREATS LANDSCAPE

SoTA	Security Services	Threats			
		D	L	M	I
[12], [10]	ID / Auth				✓
[11], [13]	Access Control	✓	✓	✓	
	Sec. Interconnection	✓	✓	✓	✓
	Sec. Virtualization	✓	✓	✓	
[15], [14]	Awareness / Evidence	✓	✓	✓	✓
	Trust Mngmnt.			✓	✓
	Distrib. Decision			✓	
	Privacy		✓		

V. DISCUSSION

In this section we show that the security services identified during the development of the SMOG project provide us with sufficient protection mechanisms against the main threats identified in Section II. The threats considered are denial of service (D), data leak (L), manipulation (M) and impersonation (I), as depicted in TABLE I. In addition, this table illustrates the need for more research on security, based on the current state of the art (SoTA) in the field. Note that some topics are still totally unexplored in the literature while others are only partially addressed by the very few existing solutions.

TABLE I clearly illustrates that most of the threats can be covered by the basic security services. For example, identification and authentication services principally cover the threat of impersonation, while access control can provide protection against malicious entities trying to access or manipulate data and services, and even prevent some sort of denial of service – particularly those trying to deplete the resources of fog nodes. On the other hand, the advanced security services provide enhanced protection in some specific situations where cooperation is necessary. As an example, trust management services can aid in situations of uncertainty when interacting with unknown entities. Trust services can also can help in the detection of compromised or manipulated elements of the infrastructure based on, for example, their reputation.

Note that the services of situational awareness and evidence management have been placed together in the table. The reason is that awareness on its own cannot protect against any of the threats but is fundamental for their detection and prevention.

VI. CONCLUSION

The main goal of this paper is to stress the need for going beyond the existing research on Fog security, considering the needs of an ecosystem where multiple trust domains coexist and interact with each other. To this end, we have provided i) an overview on the major security services that are needed in order to reduce the impact of the security threats of Fog computing, ii) an analysis of their interdependencies and how to tackle them, and iii) a summary of the current state of the art related to these services. We will continue working on these issues under the umbrella of the SMOG project.

ACKNOWLEDGMENTS

This paper has been partially funded by the Spanish Ministry of Economy and Competitiveness through the SMOG project (TIN2016-79095-C2-1-R).

REFERENCES

[1] A. Botta, W. de Donato, V. Persico, and A. Pescap, "On the integration of cloud computing and internet of things," in *2014 International Conference on Future Internet of Things and Cloud*, Aug 2014, pp. 23–30.

[2] K. D. Kim and P. R. Kumar, "Cyber-physical systems: A perspective at the centennial," *Proceedings of the IEEE*, vol. 100, no. Special Centennial Issue, pp. 1287–1308, May 2012.

[3] OpenFog Consortium Architecture Working Group, "Openfog Architecture Overview," The OpenFog Consortium, White Paper, February 2016. [Online]. Available: https://www.openfogconsortium.org/wp-content/uploads/OpenFog-Architecture-Overview-WP-2-2016.pdf

[4] F. Bonomi, R. Milito, J. Zhu, and S. Addepalli, "Fog computing and its role in the internet of things," in *Proceedings of the First Edition of the MCC Workshop on Mobile Cloud Computing*, ser. MCC '12. New York, NY, USA: ACM, 2012, pp. 13–16. [Online]. Available: http://doi.acm.org/10.1145/2342509.2342513

[5] 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*, pp. –, 2016. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0167739X16305635

[6] M. Chiang and T. Zhang, "Fog and iot: An overview of research opportunities," *IEEE Internet of Things Journal*, vol. 3, no. 6, pp. 854–864, Dec 2016.

[7] S. Yi, C. Li, and Q. Li, "A survey of fog computing: Concepts, applications and issues," in *Proceedings of the 2015 Workshop on Mobile Big Data*, ser. Mobidata '15. New York, NY, USA: ACM, 2015, pp. 37–42. [Online]. Available: http://doi.acm.org/10.1145/2757384.2757397

[8] F. Y. Rashid, "Dyn DDoS attack exposes soft underbelly of the cloud," online, Oct 2016. [Online]. Available: http://www.infoworld.com/article/3134023/security/dyn-ddos-attack-exposes-soft-underbelly-of-the-cloud.html

[9] D. Kushner, "The Real Story of Stuxnet," *IEEE Spec*, 2013. [Online]. Available: http://spectrum.ieee.org/telecom/security/the-real-story-of-stuxnet/

[10] I. Stojmenovic, S. Wen, X. Huang, and H. Luan, "An overview of fog computing and its security issues," *Concurrency and Computation: Practice and Experience*, vol. 28, no. 10, pp. 2991–3005, 2016, cpe.3485. [Online]. Available: http://dx.doi.org/10.1002/cpe.3485

[11] P. Zhang, Z. Chen, J. K. Liu, K. Liang, and H. Liu, "An efficient access control scheme with outsourcing capability and attribute update for fog computing," *Future Generation Computer Systems*, pp. –, 2016.

[12] M. H. Ibrahim, "Octopus: An edge-fog mutual authentication scheme," *I. J. Network Security*, vol. 18, no. 6, pp. 1089–1101, 2016. [Online]. Available: http://ijns.femto.com.tw/contents/ijns-v18-n6/ijns-2016-v18-n6-p1089-1101.pdf

[13] C. Dsouza, G. J. Ahn, and M. Taguinod, "Policy-driven security management for fog computing: Preliminary framework and a case study," in *Proceedings of the 2014 IEEE 15th International Conference on Information Reuse and Integration (IEEE IRI 2014)*, Aug 2014, pp. 16–23.

[14] M.-G. Ionita and V.-V. Patriciu, "Secure threat information exchange across the internet of things for cyber defense in a fog computing environment," *Informatica Economica*, vol. 20, no. 3, pp. 16–27, Sept. 2016.

[15] Y. Wang, T. Uehara, and R. Sasaki, "Fog computing: Issues and challenges in security and forensics," in *2015 IEEE 39th Annual Computer Software and Applications Conference*, vol. 3, July 2015, pp. 53–59.

[16] C. Dwork, "Differential privacy," in *33rd International Colloquium on Automata, Languages and Programming, part II (ICALP 2006)*, vol. 4052. Venice, Italy: Springer Verlag, July 2006, pp. 1–12. [Online]. Available: https://www.microsoft.com/en-us/research/publication/differential-privacy/

[17] C. Cachin, "Architecture of the hyperledger blockchain fabric," in *Workshop on Distributed Cryptocurrencies and Consensus Ledgers*, 2016.