

The Extended Cloud: Review and Analysis of Mobile Edge Computing and Fog from a Security and Resilience Perspective

Syed Noorulhassan Shirazi, Antonios Gouglidis, Arsham Farshad and David Hutchison

(Invited Paper)

Abstract—Mobile Edge Computing (MEC) and Fog are emerging computing models that extend the cloud and its services to the edge of the network. The emergence of both MEC and fog introduce new requirements, which mean their supported deployment models must be investigated. In this paper, we point out the influence and strong impact of the *extended cloud* (i.e., the MEC and fog) on existing communication and networking service models of the cloud. Although the relation between them is fairly evident, there are important properties, notably those of security and resilience, that we study in relation to the newly posed requirements from the MEC and fog. Although security and resilience have been already investigated in the context of the cloud - to a certain extent - existing solutions may not be applicable in the context of the extended cloud. Our approach includes the examination of models and architectures that underpin the extended cloud, and we provide a contemporary discussion on the most evident characteristics associated with them. We examine the technologies that implement these models and architectures, and analyse them with respect to security and resilience requirements. Furthermore, approaches to security and resilience-related mechanisms are examined in the cloud (specifically, anomaly detection and policy based resilience management), and we argue that these can also be applied in order to improve security and achieve resilience in the extended cloud environment.

Index Terms—Cloud computing, cloud resilience, edge computing, mobile edge computing, fog computing, cloud security.

1 INTRODUCTION

CLOUD computing technologies are becoming increasingly important since they provide a wide range of beneficial properties, such as on-demand self-services, resource pooling, rapid elasticity, measured services, etc. Although several technologies are used for the conceptualisation of the cloud model, the main enabling technology is virtualisation [1]. Cloud technologies have managed to mature over the years, to the extent that major companies, such as Amazon, Google, Ebay, etc., invested a lot of their resources in developing and using cloud infrastructures for the provision of their services. Many organisations now face the question as to when and how to cloudify their existing IT infrastructures. According to a survey, conducted by the Cloud Industry Forum (CIF)¹, 63% of UK businesses are planning to migrate their entire IT infrastructure to the cloud in the near future. CIF also recorded the opinion of 250 senior IT and business decision makers at the end of 2015, and found that 78% of UKs organisations are already using the cloud. This indicates an increase of 53% between 2011 and 2015. Moreover, it is predicted that by 2018, approximately 85% of businesses will use the cloud [2].

The main reasons for the cloud's increased popularity lie in its supported business models, which eventually result in reducing costs and offering greater scalability and on-

demand resource provisioning services. Since most cloud providers specialise in both hardware and software technologies, cloud users can be relieved of the need to have in-house teams to conduct maintenance operations on the infrastructure. Motivated primarily by economies of scale, cloud environments are also being used by sectors operating in the area of critical infrastructures [3], [4], [5], [6], [7]. Further characteristics of the cloud (e.g., support of ubiquitous connectivity, elasticity, scalable resources and ease of deployment) render these computing environments applicable in domains such as the Internet of Things (IoT) – sensors, mobile devices etc., [8]. This emerging trend of IoT deployments is introducing new requirements that existing cloud settings cannot satisfy adequately [9], [10], [11]. These requirements include, but are not limited to, geo-distribution, low latency, location awareness, and mobility support [12], [13].

In order to fulfil the above mentioned requirements, the research community has proposed new technologies, namely the *edge* [14] and *fog* [15]. These technologies, which we collectively label the *extended cloud*, allow computing needs to be performed closer to the source of data. Eventually this will improve the quality of provided services since it would result in reducing the delay in conveying data between end nodes and the cloud. A typical conceptual architecture of the extended cloud is shown in Fig. 1. These technologies enable the support of new applications and services, e.g., Google Now² and Foursquare³ that

• S. Shirazi, A. Gouglidis, A. Farshad and D. Hutchison are with InfoLab21, School of Computing and Communications, Lancaster University, LA1 4WA, UK (e-mail: n.shirazi@lancaster.ac.uk; a.gouglidis@lancaster.ac.uk; a.farshad@lancaster.ac.uk; d.hutchison@lancaster.ac.uk)

1. <https://www.cloudindustryforum.org/>

2. <https://www.google.co.uk/landing/now/>

3. <https://foursquare.com>

are both location-aware applications for mobile platforms. Further types of supported application include, amongst others, traffic control management for autonomous vehicles, robotics, public safety, augmented reality, etc. [16].




Fig. 1. The extended cloud

The edge or Mobile Edge Computing (MEC) and fog models offer similar, yet different methods for monitoring, processing, and conveying data. Moreover, each model has its own advantages and disadvantages, which make them individually preferable in different scenarios. Edge computing refers to data processing at the edge of a network close to the data source. For example, sensors in industrial IoT applications can capture streaming data towards optimising the production. This is generally achieved by connecting sensors to Programmable Automation Controllers (PAC) that handle processing and communication. In the case of edge computing, IT and cloud computing capabilities are provided within the Radio Access Network (RAN), which is close to mobile subscribers. This results in offering context aware application and services with ultra-low latency and high-bandwidth requirements [17]. By contrast, fog computing necessitates data processing from the edge to the cloud. It allows data collection, processing and storage at the local area network end of the device using a *fog node*; hence it offers less latency when compared to the cloud. Compared to edge computing, fog is more scalable as it offers a centralised model for processing data. However, the edge model has an advantage over the fog, as there are fewer points of failure. Specifically, devices are more independent in terms of decision making, i.e., making decisions on whether information should be stored locally or sent to the cloud. Table 1 provides comparative information among the cloud, edge and fog models – compiled by Cisco⁴ – and highlights some of the limitations in the cloud with regard to real-time application requirements. These new computing models provide a new ecosystem and an opportunity for providers to collaborate and develop new business models. This is feasible through the deployment of new services

4. <http://blogs.cisco.com/perspectives/iot-from-cloud-to-fog-computing>

for users and tenants based on their service requirements. Nevertheless, the edge and fog do not replace the cloud, but rather we consider them as a non-trivial extension of cloud. Crucially, security and resilience issues in the context

TABLE 1
Cloud computing vs. Fog and Edge computing

Requirements	Cloud Computing	Fog/Edge Computing
Latency	High	Low
Location of Service	Within the Internet	At the edge of the local network
Geo-distribution	Centralized	Distributed
Mobility	Limited support	High support
Location awareness	No	Yes
Type of last mile connectivity	Leased Line	Wireless
Distance between client and server	Multiple hops	One hop

of cloud computing can be foreseen to impact both edge and fog. We now elaborate on these vital security and resilience concerns in Sections 1.1 and 1.2, respectively.

1.1 Security Aspects

In spite of the merits imposed by moving to the edge and fog, a shift to the extended cloud brings its own challenges. In cloud environments users have less control over the hardware, software and data. The loss of control over data and lack of transparency give rise to many security concerns, which cause uncertainty for organisations that want to ‘cloudify’ their IT infrastructure. This is highlighted in recent reports published by Vision, which show an increased reluctance by companies to migrate their infrastructure to the cloud due to security concerns [18]. Similarly, another recent study by Alert Logic states that application attacks aimed at cloud deployments grew by 45% over the period of a year [19]. The complexity of the underlying infrastructures also introduce a number of challenges, including misconfiguration, and malware. Furthermore, failures in the cloud may incur significant costs. For instance, on the 29th of June, 2010, amazon.com experienced problems in placing orders using its platform. Based on their 2010 quarterly revenues, such downtime resulted in a loss of approximately \$1.75 million per hour to Amazon. More recently in another incident, on October 21st, 2016, dyn.com suffered from a high speed Distributed Denial-of-Service (DDoS) attack that affected numerous websites. In this case, enterprises that relied on Software-as-a-Service (SaaS) for running critical business operations experienced failures. Thus, we argue that due to the widespread use of cloud infrastructures for hosting critical services, any potential disruption in the cloud would have a great impact on the reliant services, e.g., in health, safety, security or economic well-being of citizens or the effective functioning of governments [20], [21]. While some issues can be addressed using existing mechanisms (see Section 4), there are additional threats to the edge and fog, which pose risks to the cloud. Sections 2

and 3 highlight some of these threats in the mobile edge and fog, respectively.

1.2 Resilience as a Cloud Requirement

Failures in cloud services may result in having a great impact on the security and safety of the virtualised infrastructures, and therefore render the requirement of supporting resilience in the cloud more important than even before. Resilience is defined as the ‘ability of a system to provide an acceptable level of service in the presence of challenges’ [22].

The *acceptable level of service* in the above definition depends on user expectation. Nowadays, users require instantaneous access to information that is available at all times: the *always-on, always-available* service. Due to the presence of various challenges, resilience is sometimes evaluated through two key metrics, i.e., recovery time objective (RTO) and recovery point objective (RPO). RTO refers to the duration of time within which a system must be restored after a disruption to avoid a break in business continuity, and RPO refers to the nearest point to where a system may be recovered after a disruption [23]. Furthermore, resilience is concerned not only with the *Availability* of services, but also with maintaining the *Confidentiality* and *Integrity* of the information in the face of challenges [9].

Resilience needs to become a fundamental property of the cloud service provisioning platforms. However, the innate and often desirable properties offered by cloud-based models – such as elasticity, virtualisation, scalability and geo-distribution of devices – render the implementation of standard resilience solutions more problematic.

- Elasticity in the cloud computing refers to the dynamic adaptation of capacity to meet workload demands. It is defined as ‘the degree to which a system is able to adapt to workload changes by provisioning and de-provisioning resources in an autonomic manner, such that at each point in time the available resources match the current demand as closely as possible’ [24], [25]. Providing resilience and service guarantees in the presence of various challenges requires resources in terms of redundant storage for backup-recovery and additional network capacity, etc. However, due to varying customer workloads and the requirement of elasticity in the cloud, resource allocation and availability is constantly changing making it difficult to provide resilience and service guarantees. Service providers are faced with the challenge of determining how an application on the cloud (e.g. a fog node) should be configured, and how much of the resources should be protected.
- Virtualisation in cloud computing also introduces new sources of threats and failures. Complex interactions between the applications, and workloads sharing the same physical infrastructure make it difficult to predict demands. When it comes to the provision of virtual resources, cloud virtualisation must extend beyond the data center to reach the edge and fog nodes. This extension further reduces control over the underlying physical hardware. The difficulty of predicting interactions and adapting the system accordingly makes it hard to provide dependability guarantees in terms of availability and responsiveness, as well as resilience to external perturbations such as security attacks [26].

- The extended cloud environments achieve higher utilisation of physical resources through the consolidation of workloads. However, this makes them more vulnerable to threats resulting from unpredictable resource demands as well as operational failures, such as hardware and software failures, unforeseen load fluctuations, network attacks, etc.
- The geographical distribution in edge/fog based service architectures introduces various dependencies that impede the satisfaction of SLAs. This distribution may be between clients and services, as well as the physical infrastructure behind services. Cloud providers are coping with the need for assuring the availability and reliability (i.e. resilience) of services, which are distributed over the extended cloud architecture. The geographical distribution may also impose legal and privacy implications when data is hosted on outsourced and shared infrastructures that are under a different legal jurisdiction than that of the owner of the data. Privacy clauses in legislation – such as the Health Insurance Portability and Accountability Act (HIPAA)⁵ and the Telecommunications Act of 1996⁶, as well as financial regulations such as the Sarbanes-Oxley Act⁷ – obstruct the applicability of extended cloud solutions in their respective industries, and may adversely affect resilience.

The extended cloud model is multifaceted since services may be provided by an assortment of heterogeneous platforms, and resources may be shared between entities. The applications running on these platforms may employ best practices to varying degrees, and thus result in increasing the risk of unforeseen threats. These risks are exacerbated due to multiple administrative domains between the application and infrastructure operators, and reduce the end-to-end system visibility. To best of our knowledge, there is limited understanding on how to provide resilience mechanisms for the extended cloud that can collectively address the various challenges in a coordinated manner.

1.2.1 Resilience Strategy

The ResumeNet⁸ project evaluated a framework whereby a number of resilience principles and a strategy (i.e. D^2R^2+DR for Defend, Detect, Remediate, Recover, Diagnose and Refine) are defined, which is outlined in Fig. 2. At its core, a control loop offers a number of conceptual processes that are used to realise the real-time aspect of the D^2R^2+DR strategy and consequently implement resilience; this can be exploited in introducing resilience for the cloud. Based on the resilience control loop, other necessary elements for the detection of challenges and remediation can be derived, such as anomaly detection and policy engines respectively, that aim to build situational awareness and control mechanisms.

Under the D^2R^2+DR framework, there must exist components that are capable of reconfiguring devices, in response to challenges, using policies. Reconfiguration need

5. <http://www.dhcs.ca.gov/formsandpubs/laws/hipaa/Pages/1.10HIPAATitleInformation.aspx>

6. <https://www.fcc.gov/general/telecommunications-act-1996>

7. <http://www.soxlaw.com/>

8. <http://www.comp.lancs.ac.uk/resilience/>




Fig. 2. The D^2R^2+DR resilience strategy [22]

not apply to the same components on which the detection was based. A policy engine is responsible for mapping detection events to reconfigurations, accepting a resilience strategy expressed as a collection of policies. The generic nature of D^2R^2+DR makes it ideal for application to the extended cloud. For illustration purposes, Fig. 3 provides a high-level mapping of the D^2R^2 part of the strategy to the reference architecture for Mobile Edge Cloud architecture from ETSI⁹. The mobile edge platform operator has access to physical devices, and some virtual components hosting the applications, which can be monitored to inform the detection process. The operator can also reconfigure these devices and applications, in response to detected challenges using policies. At the network level, D^2R^2 may exist as monitoring and reconfiguration points. Within the inner D^2R^2 loop, some interaction between architecture layers may exist in the form of events and reconfigurability exposed by the lower layers of the framework [27].

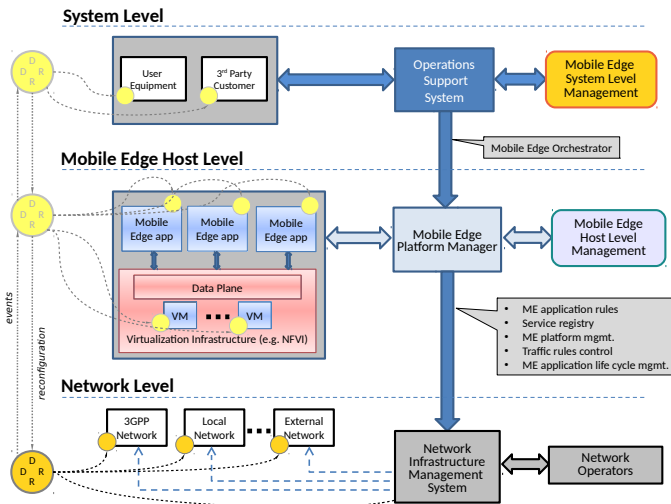


Fig. 3. A resilience-oriented view of the MEC reference architecture [27]

The remainder of this paper is organised as follows. Section 2 and Section 3 focus on the mobile edge architecture and the fog respectively, from a unique perspective of security and resilience. Section 4 discusses security and resilience requirements in the cloud and considers the implementation of resilience mechanisms towards an overall cloud architecture. Finally, we provide concluding remarks in Section 5.

9. <http://www.etsi.org/>

2 MOBILE EDGE COMPUTING

Mobile wireless communication has received a lot of attention and has become hugely popular during the past decade. This is mostly due to the emergence of mobile devices, such as smart-phones and tablets. These devices are usually equipped with multi-core processors, various sensors (e.g., high-quality camera, GPS, barometer, etc.), as well as running a plethora of applications that managed to improve the productivity of mobile users.

Meanwhile, there is an increased demand for ever higher bandwidth rates when connecting to the Internet, and thus a need for new designs when it comes to next generation mobile networks (notably 5G). Some of its characteristics include the support of $1000\times$ the number of connected devices, $100\times$ the user data rate, $1/5\times$ the end-to-end delay and $1/1000\times$ the service deployment time [28].

Some 5G goals include the enhancement of traditional mobile services (e.g. telephony services) and provision of new services/applications (e.g. IoT, vehicular communication, augmented reality, live video streaming, mobile gaming, mobile crowd sensing, etc.). Common requirements of these examples are the demand for higher bandwidth rates, computing resources and very low delays when accessing services. The European 5G PPP (Infrastructure Public Private Partnership) research body recognises Mobile Edge Computing (MEC) to be one of the key technologies for 5G networks [29].

MEC was initially defined by IBM and Nokia Siemens to describe the computing facilities within a mobile base station. Toward its standardisation, the European Telecommunications Standards Institute (ETSI) launched an Industry Specification Group (ISG) in December 2014 (i.e. ISG MEC) that resulted in producing documents providing a series of specifications for the MEC framework [27] and an architecture [30]. MEC is defined by ETSI to be capable of providing ‘... an IT service environment and cloud computing capabilities at the edge of the mobile network to reduce latency, ensure highly efficient network operation and service delivery, and offer an improved user experience’. The MEC definition has recently been extended to apply to multiple communication technologies and referred to as Multiple-access Edge Computing. However, in this paper, we choose to use the definition provided by the ETSI MEC ISG (we also call it *edge* for short). Next, we introduce, initially, mobile services/applications that can serve as use-cases for the MEC model. For the deployment of services, we review the ETSI ISG reference architecture and framework. Finally, we elaborate on the security and resilience challenges, and provide our thoughts for the future of the MEC.

2.1 MEC Reference Architecture

The ETSI GS MEC 003 V1.1.1 (2016-03) document specifies a framework and reference architecture for the MEC [30]. The framework describes the structure of the MEC environment, and the reference architecture illustrates the main functional blocks and defined reference points between them.

Fig. 4 depicts the proposed MEC reference architecture, where the general entities and their functional blocks are illustrated. Entities are grouped in three levels, namely host level, network level and system level.

The Mobile Edge (ME) host and network: The host level is represented by the middle level of MECs reference architecture. It includes the ME host and ME host management entities. The ME host is comprised of the virtualisation infrastructure (VI) and the ME platform and applications. The VI provides the computing, storage and networking infrastructure for running ME applications. The ME platform facilitates the running of ME applications on a particular virtualised infrastructure and provides ME applications as services. Specifically, the platform provides the ability to advertise, discover and consume ME applications through a service registry. The ME platform also applies traffic rules and provides DNS services. The network level provides the connectivity between internal and external entities of MEC.

The ME management: The ME management task is separated across different levels. The ME platform management includes a platform manager and a VI Manager (VIM). The ME system level management operates as an orchestrator. Thus, its main responsibilities include functionalities such as performing the optimal service provisioning (considering the ME VI available resources) and preparation of the ME virtual infrastructure (for deploying applications). Moreover, it maintains the ME resources (e.g., topology, available ME host resources, etc.) by triggering the instantiation, relocation and termination of applications. It also carries out integrity and authenticity checks for the ME application packages and does policy enforcement. The Operations Support System (OSS) receives the request for the ME APP deployment or termination from the Customer Facing Service (CFS) or directly from user end (UE) applications and passing the authorised requests to the orchestrator. To facilitate the communication of UE APPs to the orchestrator for requesting instantiation and possibly relocation of ME application, the user application life-cycle management (LCM) proxy is provided.

Resilience has always been a great challenge for cloud computing, as well as for the MEC model. With regard to the edge model, it is even more challenging as there are interactions with various access technologies, such as WiFi, Bluetooth and LTE; different MEC applications and multi-tenant infrastructures all render the deployment of resilience mechanisms a technically challenging problem. The described reference architecture for the MEC can be improved by introducing resilience mechanisms across different layers. Resilience can be improved by considering relevant requirements when provisioning a MEC application and by deploying monitoring mechanisms to become aware of the current resilience level of the system/network. To improve our understanding of the security and resilience challenges for MEC, we consider threats and challenges that may target the MEC model. In [31], there is an attempt to summarise these threats across the various levels in MEC – the main issues are listed in the following.

Infrastructure Threats: The edge network infrastructure is part of the ‘last-mile network’, and thus exploits different technologies to build the network. This makes the edge infrastructure prone to several attacks. For example, Denial-of-Service (DoS) attacks and wireless jamming can easily consume the bandwidth, frequency band and computing resources at the edge. Man-in-the-middle attack is another example; it is used to inject or eavesdrop traffic from the edge. In addition, wireless communication technologies – being the dominant edge networking technology – make data accessing over the air a more probable source of attack.

Virtualisation: The MEC faces also various threats and challenges related to the virtualisation technology, which is usually deployed to share the resources of the MEC. Such attacks include, but are not limited to, DoS attacks, where a single or multiple malicious virtual machine(s) can deplete the host computing/networking resources shared between all VMs. Rogue VMs can run malicious applications for hacking other VMs or used for eavesdropping data.

Privacy leakage: Adversaries can gain access to the information store in the edge cloud. The extent of the leakage is limited as the MEC stores information only about local users. However, the type of information in these stores (users’ personal information) makes this one of the great challenges in deploying edge cloud. Encryption can be a solution to this challenge.

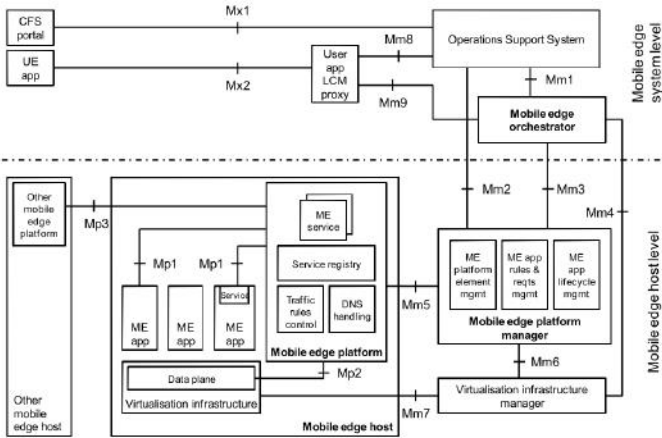


Fig. 4. The MEC Reference Architecture [30]

2.2 Security and Resilience Related Challenges

Security and resilience are key issues when considering MEC services such as messaging, navigators, etc., as well as privacy of data when this is shared through a MEC infrastructure (e.g., video, photos, various sensor data).

2.3 Future Directions

5G is the next generation of mobile communication and MEC is considered likely to help make it a success. The standardisation research bodies such as ETSI and open source communities (e.g., ONS and M-CORD project) already contributed in defining the MEC reference architecture. However, the concepts of security and resilience have not been addressed adequately when designing the MEC framework. Thus, better integration of security, resilience and privacy preserving mechanisms in the MEC reference architecture should be amongst the future directions. Another potential future direction could be the merging of activities in NFV and MEC. The deployment of MEC using an NFV environment is under investigation by a newly established working group within ETSI. Specifically, the goal is the integration of

the MEC and NFV [32]. Toward that direction Sciancalepore et al. proposed MANO+ in [33] as the result of comparing two reference architectures for the management of MEC and NFV.

3 FOG COMPUTING

The *Fog Computing* model is conceived as an extension of the cloud. The term ‘fog’ was originally defined by Cisco in an attempt to describe the need for an enabling platform that will be able to cope with the requirements posed by critical Internet of Things (IoT) services. In [15], fog computing is characterised as ‘a highly virtualised platform that provides, compute, storage, and networking services between end devices and traditional cloud computing data centres, typically, but not exclusively located at the edge of network.’

Fog computing is usually considered in application scenarios where data needs to be collected close to the edge devices; the number of devices is large; devices are physically distributed; and, there is a requirement for low latency services [34]. Fig. 1 depicts the conceptual relation among the cloud, the fog, and the IoT devices. Briefly, IoT devices consist of sensors that send data to the *fog nodes*. The latter can receive feeds from IoT devices in real time, can run applications for performing operations to the collected data (e.g. analytic), and provide short term storage of data. Moreover, fog nodes can send aggregated information to the cloud, usually in predefined intervals. With regard to the cloud – it aggregates and collects data from the fog nodes, and processes it, e.g., performs an analysis of the aggregated data, which can be used later for decision making purposes. The cloud can also send new directives to the fog nodes based on the results of the data analysis to fulfil specific operational requirements.

3.1 State of the Art

A first implementation of the fog computing concept has been realised by Cisco. Specifically, IOx is Cisco’s solution for developing and deploying fog applications. Cisco provides a framework for hosting applications and services developed not only by it, but by its partners and third-party developers. This has been achieved through the development of a concrete architecture¹⁰ for deploying fog applications, which consists of four main components, i.e., Cisco IOx, fog director, an SDK and development tools, and fog applications.

The OpenFog Consortium¹¹ was formed in late 2015 by a variety of founding members, viz. ARM, Cisco, Dell, Intel, Microsoft and Princeton University, and yet more contributing ones (e.g., AT&T, Foxconn, Hitachi, etc.), which strive jointly towards the definition of a reference architecture for the fog computing model. The consortium has published (in early 2017) in [35] a working document that provides a description of a high-level system architecture for fog nodes and networks. The main pillars that are examined in the context of an architecture include ‘... security, scalability, openness, autonomy, RAS (reliability, availability and serviceability),

agility, hierarchy, and programmability’. The investigation of these main pillars will eventually result in coping with the main challenges of fog computing.

3.2 Security and Resilience Related Requirements

Although several security and resilience requirements in the fog may be the same as in the cloud – as a result of both being virtualised platforms – we attempt to elicit, in the following, requirements that are related more to the fog model.

Security is defined by the OpenFog Consortium as one of the main pillars for investigation in fog computing. Specifically, end-to-end security is required for all applications in the fog model, and thus requiring several attributes of security to be examined, viz. ‘... privacy, anonymity, integrity, trust, attestation ...’ [35].

Privacy of data in fog computing is a requirement that has to be fulfilled when considering the analysis of sensitive information. In support of this requirement, optimal selection of fog nodes to conduct potential data analysis processes will be considered.

Data collected, conveyed, and processed in the fog needs to be appropriately anonymised. Nevertheless, information from data must be carefully removed in order not to affect any inter-connected processes using that information (e.g., intrusion detection system, anomaly detection algorithms, etc.), preserving at the same time the required level of anonymity.

The integrity and availability of the fog computing infrastructure and of data must not be compromised [34]. This stems from the need for the fog to operate reliably. Specifically, IoT devices are required to cope with an abundance of data that may affect decision-making processes, and thus may have a great impact on people – especially in the case of using the fog to host critical services.

Trust models and reputation systems are also considered to be an importance aspect in the fog. Yi et al., in [36] refer to these requirements and highlight the need for a solution applicable in fog computing. Among other solutions, technologies such as Trusted Execution Technology (TXT) along with Trusted Platform Module (TPM) have been discussed in the literature for ensuring such requirements. In addition to the investigation of trust models, attestation must be ensured. Attestation is considered as allowing a remote device to formally prove (e.g. through the use of a cryptographic scheme) its trustworthiness to a remote verifier [37].

In the presence of network challenges, resilience must be ensured in the fog nodes [38]. This requires the monitoring of the fog infrastructure, prompt detection of challenges that disrupt the normal operation of the network/system, and attempting to mitigate and ultimately recover the network/system to its original, normal, operational state.

3.3 Security and Resilience Related Challenges

Although the range of challenges can be quite extensive, we choose to list in the following a set of security and resilience related challenges, that need to be overcome in order to further promote the concept of fog computing.

10. <https://developer.cisco.com/site/iox/docs/#introduction-to-iox>

11. <https://www.openfogconsortium.org/>

Authentication – i.e. the process of verifying the identity of a user or process – is identified in [39] as a security issue in application scenarios of using fog computing in infrastructures such as the Smart Grid. The problem lies in the way authenticating of smart metering devices on the edge (i.e. at the consumer) is done, which involves the threats of tampering the device, reporting false data, etc. In order to overcome this issue, proposed approaches embrace cryptographic schemes, e.g. key exchange algorithms and use of a public key infrastructure. However, these solutions do not scale well in these environments [40], and emerging technologies are also considered in that context (e.g. biometrics).

The level of heterogeneity and the distributed nature of the devices in fog computing may raise issues when it comes to access control. Access control offers mechanisms to control and limit the actions or operations that are performed by a user or process on a set of objects, i.e. system resources. Inter-domain collaboration among systems may need to ensure additional security concepts such as that of secure inter-operation among the resources in the fog [41], [42], [43]. Although several solutions cope with secure inter-operation, their application in the context of fog computing may need to consider the emergent requirements set by the fog model.

Resilience covers several aspects in a network or system, namely, defence and detect mechanisms, as well as remediation and recovery strategies. The fog computing model imposes a great diversity in devices, network communication links, applications, etc., and thus makes its security a cumbersome process [44]. Sometimes resilience is examined and ensured on a single layer of an architecture (e.g., physical layer, network layer, etc.). However, we argue that resilience in the fog – as well as in other models and architectures – needs to be examined on multiple layers, including the investigation of potential interconnections between the layers. Examining resilience merely on a single layer would eventually fail to deliver the required level of resilience for the fog.

3.4 Future Directions

In this section, we briefly refer to future directions with regards to fog computing

Standardisation of the fog framework can be a future direction for fog computing. This could be conducted by standardisation bodies and/or major firms, forums, and consortia related with the fog model. Currently, the Open-Fog Consortium appears to provide information with regard to the system architecture for fog nodes and networks. This information describes a high-level view of them, but future steps are directed to the provision of low-level and detailed information for such architectures, including a list of formal requirements that will set the basis for test-beds.

We also strongly believe in the need for considering and elevating the importance of resilience in the fog – several pillars are considered in existing documented frameworks, but resilience is not covered explicitly. We advocate this requirement because several security approaches may fail, and thus resilience should be supported intrinsically to cope with such failures. Security controls and mechanisms offer a

considerable level of protection, but in practise these may fail (e.g. poor implementations). Resilience embraces the idea that networks/systems must have the right mechanisms in place to promptly identify and avoid challenges to the normal operation of a network/system. Furthermore, in the case of a challenge, the network/system needs to be in a position to mitigate and recover to its original, normal operation. Thus, we advocate that the definition of resilience requirements, and resilience-aware frameworks (e.g., resilience-by-design, real-time resilience), should be topics for future investigation.

4 SECURITY MECHANISMS FOR THE CLOUD

The introduction of the extended cloud platform must not affect network availability. Hence, a cloud provider should offer the level of resilience and address the high-availability requirements demanded by its users, tenants and their applications. In case of an adverse event, a resilience mechanism is required to promptly detect it and limit its effects on the normal operation of the system, and thus, for the cloud to be robust and resilient. Specifically, offering protection against any performance related anomalies such as mis-configuration will require a cloud platform to have the appropriate detection and remediation mechanisms to ensure that the provided level of service is within acceptable limits. Thus, in the presence of a challenge or a threat to the system, it has to be ensured that remediation and recovery measures are in place to prevent any disruption to the normal operation of the service. In the following, we discuss such detection and remediation mechanisms, which will eventually help to engineer resilience in the cloud.

4.1 Anomaly Detection

Anomaly detection systems identify events that appear to be anomalous with respect to the normal behaviour of the system. They have an understanding of normal system behaviour and issue alerts whenever the behaviour changes from the norm. The underlying assumption in anomaly detection is that such changes are normally caused by malicious or disrupting events. Anomaly detection has been studied within diverse research areas and application domains. However, for cloud environments, anomaly detection techniques are still evolving due to the fact that they present several challenging problems. For instance, in Infrastructure as a Service (IaaS) the customer is responsible for the correct operation of its own software, while the cloud provider is responsible only for the underlying infrastructure resources. This increases the importance of anomaly detection and remediation mechanisms. However, the detection system is effective only if the alert it generates is timely, accurate and provide actionable information to the administrators to respond to potential threats.

In the extended cloud (i.e., edge, fog and core) the network traffic comes from multiple heterogeneous domains. Moreover, it changes rapidly with respect to its behaviour patterns due to heterogeneity of the tenants using the cloud and the elasticity of the exposed services. Under such circumstances there are many challenges faced by underlying anomaly detection techniques such as mis-configurations, or

simply by high volumes of legitimate traffic. The importance of anomaly detection in such environments is due to the fact that anomalies in data translate to important actionable information. For example, in the case of fog, the anomalous traffic pattern could mean the resource requests from IoT devices suddenly increase rapidly, therefore causing a form of denial-of-service attack which could eventually hamper service availability to authorized users.

Despite the usefulness of anomaly detection, applying it operationally in the context of the extended cloud involves many other challenges with regard to performance and scalability. The on-demand provisioning nature of these environments necessitates that anomaly detection in the cloud be based on real time monitoring. Further, real time monitoring demands can also change significantly over time. Hence, the detection should not only achieve high scalability, but also embrace changes in monitoring demands. The detection must also provide good multi-tenancy support to ensure that multiple tenants/providers can use anomaly detection at the same time. Previous research has created scalable methods for real time data collection [45], [46], to support online detection based on data mining and machine learning approaches [47], [48], [49], [50]. However, while monitoring has been feasible at scale, detection is typically performed after a volume of monitoring data has been stored on disk, which impedes the scalability of real time detection.

The authors in [51] carried out an experimental evaluation of state-of-the-art anomaly detection techniques to assess their monitoring and detection capabilities in multi-tenant cloud infrastructures under elastic behaviour. It is shown that elasticity of the cloud makes the traffic distribution observed by anomaly detection techniques unstable and makes it difficult to predict normal behaviour. Dapper [52] uses a static sampling strategy which is homogeneous across all nodes in the network and thus makes it inflexible for multi-tier applications running at the core. Wang et al. [53] proposed the EbAT system to allow the on-line analysis of multiple metrics obtained from system-level components (e.g., CPU utilisation, memory utilisation, read/write counts of the OS, etc.). The system showed potential in detection accuracy and monitoring scalability, but it was not evaluated in the context of adequately pragmatic cloud scenarios. Guan et al. [54] and Garfinkel et al. [55] proposed multi-level anomaly detection techniques to detect intrusions at different levels of a cloud system. The techniques appear to be rather inflexible and the application of those techniques in an operational context requires better clarification. Similarly, Dastjerdi et al. [56] proposed an approach based on mobile agents for an intrusion detection system for cloud systems. However, scalability appears to be a problem due to the high number of virtual machines that are required to be attached to the agent. The authors in [57] present a lightweight real-time detection technique that is suitable for edge and fog computing environments and can also address application-related issues which are manifest in performance anomalies. However, an adequate evaluation of a cloud-specific scenario is missing.

4.2 Policy based Resilience Management

Management and resilience of the cloud environments are closely linked. The extended cloud will make it difficult to

plan effective management due to varying user demands, co-hosted VMs and the arbitrary deployment of multiple applications. Generally, management policies are used to govern the behaviour of a system. These management policies can be generally seen as: *'the constraints and preferences on the state or the state transition, of a system and is a guide on the way to achieving the overall objective which itself is also represented by a desire system state'* [58].

Challenges to the operation of a cloud infrastructure can occur rapidly, requiring a quick response in order to maintain acceptable service levels. In order to mitigate a challenge, complex multi-phase strategies are required, which combine various monitoring and detection mechanisms that in turn influence the behaviour of remediation mechanisms. Policy based management has proven to be very effective for complex system management as evident from previous literature. In [59] the authors presented a cloud resilience management framework that uses policy-based management techniques for the configuration of resilience strategies. These techniques allow descriptions of real-time adaptation strategies, which are separate from the implementation of the mechanisms that realise the strategy. This separation allows changes to be made to strategies without the need to take resilience mechanisms off-line. At its core is a control loop comprising of a number of conceptual processes that realise the real time aspect of the existing resilience strategy, and consequently implement resilience for cloud. The *ResumeNet*¹² project provides blueprints and design guidelines for a cloud resilience management framework. The proposed resilience strategy can be applied to deal with challenges in fog and edge platforms to explore the implications of multi-staged and collaborative detection. The *TClouds*¹³ project targets cloud computing security and minimization of the widespread concerns about the security of personal data by putting its focus on privacy protection in cross-border infrastructures and on ensuring resilience against failures and attacks. *Cloud Controls Matrix (CCM)* [60] is specifically designed to provide fundamental security principles to guide cloud vendors and to assist prospective cloud customers in assessing the overall security risk of a cloud provider. The CSA (Cloud Security Alliance) CCM (Cloud Control Matrix) provides a control framework that gives a detailed understanding of security concepts and principles that are aligned to the Cloud Security Alliance guidance in 13 domains. The foundations of the Cloud Security Alliance Control Matrix rest on its customized relationship to other industry accepted security standards, regulations, and controls framework such as the *ISO 27001/27002*, *ISACA CoBIT*, *PCI* and *NIST* and will evolve to provide internal control directions for *SAS 70* attestations provided by cloud providers.

5 CONCLUSION

In this new era of software-driven communication technologies (esp. for edge and fog computing), key properties of security and resilience remain open for further investigation. Some issues with regard to security and resilience

12. www.resumenet.eu

13. www.tclouds-project.eu

have already been solved in the context of the cloud, and thus are expected to be resolved in emerging – yet related – technologies too. This is mostly due to existing similarities between the cloud and the new software-driven communication technologies; the latter rely heavily on the main concepts of the cloud, i.e. elasticity, on-demand service provision, etc. However, the requirements posed by emerging services (e.g., low latency, support of location awareness and mobility, etc.) strongly suggest the need to re-address security and resilience, and to investigate them in their new application contexts. Identifying the similarities among the computing models may also provide useful directions to their future development. Thus, we anticipate the work presented in this paper will serve as a precursor to the further investigation of both security and resilience with regard to virtualised and software-driven communication technologies, notably mobile edge computing and the fog.

ACKNOWLEDGMENT

The research presented in this paper is sponsored by UK Engineering and Physical Sciences Research Council (EPSRC) via grant agreement no. EP/L020009/1: Towards Ultimate Convergence of All Networks (TOUCAN).

REFERENCES

- [1] P. Mell and T. Grance, "The NIST definition of cloud computing," 2011.
- [2] "Why cloud computing should be at the heart of your business," <http://www.telegraph.co.uk/business/ready-and-enabled/cloud-computing/>.
- [3] G. G. Research, <http://www.geglobalresearch.com/innovation/works-bring-air-traffic-management-cloud>, sep 2011.
- [4] PCWord, "Spanish bank to move 100,000+ employees to google apps," http://www.pcworld.com/article/247851/spanish_bank_to_move_100000_employees_to_google_apps.html, jan 2012.
- [5] R. Blaisdell, "Cloud benefits in the energy and utility industry," <https://www.rickscloud.com/cloud-benefits-in-the-energy-and-utility-industry/>, Feb. 2012.
- [6] L. Reading, "Orange/sita cloud prepares for take-off," <http://www.lightreading.com/services-apps/cloud-services/orange-sita-cloud-prepares-for-takeoff/d/d-id/693612>, Feb. 2012.
- [7] Y. Sverdlik, "Nasdaq and amazon launch cloud platform for financial services," <http://www.datacenterdynamics.com/content-tracks/colo-cloud/nasdaq-and-amazon-launch-cloud-platform-for-financial-services/69672.fullarticle>, Sep. 2012.
- [8] E. Mingozzi, G. Tanganelli, C. Vallati, and V. Di Gregorio, "An open framework for accessing things as a service," in *Wireless Personal Multimedia Communications (WPMC), 2013 16th International Symposium on*. IEEE, 2013, pp. 1–5.
- [9] J. Moura and D. Hutchison, "Review and analysis of networking challenges in cloud computing," *Journal of Network and Computer Applications*, vol. 60, pp. 113–129, 2016.
- [10] S. Agarwal, M. Philipose, and P. Bahl, "Vision: the case for cellular small cells for cloudlets," in *Proceedings of the fifth international workshop on Mobile cloud computing & services*. ACM, 2014, pp. 1–5.
- [11] S. Choy, B. Wong, G. Simon, and C. Rosenberg, "The brewing storm in cloud gaming: A measurement study on cloud to end-user latency," in *Proceedings of the 11th annual workshop on network and systems support for games*. IEEE Press, 2012, p. 2.
- [12] R. P. Padhy and M. R. Patra, "Managing it operations in a cloud-driven enterprise: Case studies," *American Journal of Cloud Computing*, vol. 1, no. 1, pp. 1–18, 2013.
- [13] M. M. Islam, S. Morshed, P. Goswami, and B. Dhaka, "Cloud computing: A survey on its limitations and potential solutions," *International Journal of Computer Science Issues*, vol. 10, no. 4, pp. 159–163, 2013.
- [14] P. Garcia Lopez, A. Montresor, D. Epema, A. Datta, T. Higashino, A. Iamnitchi, M. Barcellos, P. Felber, and E. Riviere, "Edge-centric computing: Vision and challenges," *ACM SIGCOMM Computer Communication Review*, vol. 45, no. 5, pp. 37–42, 2015.
- [15] 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*. ACM, 2012, pp. 13–16.
- [16] H. T. Dinh, C. Lee, D. Niyato, and P. Wang, "A survey of mobile cloud computing: architecture, applications, and approaches," *Wireless Communications and Mobile Computing*, vol. 13, no. 18, pp. 1587–1611, 2013. [Online]. Available: <http://dx.doi.org/10.1002/wcm.1203>
- [17] M. Patel, B. Naughton, C. Chan, N. Sprecher, S. Abeta, A. Neal et al., "Mobile-edge computing introductory technical white paper," *White Paper, Mobile-edge Computing (MEC) industry initiative*, 2014.
- [18] V. Solutions, "The 2016 state of resilience: Keep your data moving forward," Tech. Rep., 2016.
- [19] A. Logic, "The changing state of cloud security," Tech. Rep., 2015.
- [20] M.A.C.Dekker, "Critical cloud computing: A CIIP perspective on cloud computing services," European Network and Information Security Agency, Tech. Rep., 2012.
- [21] S. Berman, L. Kesterson-Townes, A. Marshall, and R. Srivathsa, "The power of cloud. driving business model innovation," IBM Institute for Business Value, Tech. Rep., 2012.
- [22] J. P. Sterbenz, D. Hutchison, E. K. Çetinkaya, A. Jabbar, J. P. Rohrer, M. Schöller, and P. Smith, "Resilience and survivability in communication networks: Strategies, principles, and survey of disciplines," *Computer Networks*, vol. 54, no. 8, pp. 1245–1265, 2010.
- [23] M.A.C.Dekker, "Resilience in the era of enterprise cloud computing," IBM Global Technology Services, Tech. Rep., 2014.
- [24] N. R. Herbst, S. Kounev, and R. H. Reussner, "Elasticity in cloud computing: What it is, and what it is not." in *ICAC*, 2013, pp. 23–27.
- [25] N. Herbst, R. Krebs, G. Oikonomou, G. Kousiouris, A. Evangelinou, A. Iosup, and S. Kounev, "Ready for rain? a view from spec research on the future of cloud metrics," *arXiv preprint arXiv:1604.03470*, 2016.
- [26] S. Kounev, P. Reinecke, F. Brosig, J. T. Bradley, K. Joshi, V. Babka, A. Stefanek, and S. Gilmore, "Providing dependability and resilience in the cloud: Challenges and opportunities," in *Resilience Assessment and Evaluation of Computing Systems*. Springer, 2012, pp. 65–81.
- [27] E. M. E. Computing, "Framework and reference architecture," *ETSI GS MEC*, vol. 3, p. V1.
- [28] "Understanding 5G: Perspectives on future technological advancements in mobile," 2014. [Online]. Available: <https://www.gsmaintelligence.com/research/?file=141208-5g.pdf{&}download>
- [29] Members of the 5G Infrastructure Association, "5G Vision," p. 16, 2015. [Online]. Available: <https://5g-ppp.eu/wp-content/uploads/2015/02/5G-Vision-Brochure-v1.pdfwww.5g-ppp.eu>
- [30] E. GS, "GS MEC 003 - V1.1.1 - Mobile Edge Computing (MEC); Framework and Reference Architecture."
- [31] 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*, p. 28, 2016.
- [32] DGR/MEC-0017MECinNFV, "Mobile Edge Computing (MEC) Deployment of Mobile Edge Computing in an NFV environment," 2016. [Online]. Available: https://portal.etsi.org/workapp/WorkProgram/Report_WorkItem.asp?wki_id=49447
- [33] V. Sciancalepore, F. Giust, K. Samdanis, and Z. Yousaf, "A double-tier MEC-NFV architecture: Design and optimisation," in *2016 IEEE Conference on Standards for Communications and Networking (CSCN)*. IEEE, oct 2016, pp. 1–6.
- [34] Cisco, "Fog computing and the internet of things: Extend the cloud to where the things are." Cisco, 2015, accessed: 2017-02-15.
- [35] OpenFogConsortium, "Openfog reference architecture for fog computing," 2017.
- [36] 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*. ACM, 2015, pp. 37–42.
- [37] 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, 2016.

- [38] S. Yi, Z. Hao, Z. Qin, and Q. Li, "Fog computing: Platform and applications," in *Hot Topics in Web Systems and Technologies (HotWeb), 2015 Third IEEE Workshop on*. IEEE, 2015, pp. 73–78.
- [39] I. Stojmenovic and S. Wen, "The fog computing paradigm: Scenarios and security issues," in *Computer Science and Information Systems (FedCSIS), 2014 Federated Conference on*. IEEE, 2014, pp. 1–8.
- [40] S. Yi, Z. Qin, and Q. Li, "Security and privacy issues of fog computing: A survey," in *International Conference on Wireless Algorithms, Systems, and Applications*. Springer, 2015, pp. 685–695.
- [41] A. Gouglidis, I. Mavridis, and V. C. Hu, "Security policy verification for multi-domains in cloud systems," *International Journal of Information Security*, vol. 13, no. 2, pp. 97–111, 2014. [Online]. Available: <http://dx.doi.org/10.1007/s10207-013-0205-x>
- [42] A. Gouglidis and I. Mavridis, "domrbac: An access control model for modern collaborative systems," *Computers & Security*, vol. 31, no. 4, pp. 540 – 556, 2012. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0167404812000144>
- [43] C. Dsouza, G.-J. Ahn, and M. Taguinod, "Policy-driven security management for fog computing: Preliminary framework and a case study," in *Information Reuse and Integration (IRI), 2014 IEEE 15th International Conference on*. IEEE, 2014, pp. 16–23.
- [44] 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*, 2016.
- [45] M. L. Massie, B. N. Chun, and D. E. Culler, "The ganglia distributed monitoring system: design, implementation, and experience," *Parallel Computing*, vol. 30, no. 7, pp. 817–840, 2004.
- [46] P. Yalagandula and M. Dahlin, *A scalable distributed information management system*. ACM, 2004, vol. 34, no. 4.
- [47] P. Bahl, R. Chandra, A. Greenberg, S. Kandula, D. A. Maltz, and M. Zhang, "Towards highly reliable enterprise network services via inference of multi-level dependencies," in *ACM SIGCOMM Computer Communication Review*, vol. 37, no. 4. ACM, 2007, pp. 13–24.
- [48] E. Kiciman and A. Fox, "Detecting application-level failures in component-based internet services," *Neural Networks, IEEE Transactions on*, vol. 16, no. 5, pp. 1027–1041, 2005.
- [49] C. Wang, K. Viswanathan, L. Choudur, V. Talwar, W. Satterfield, and K. Schwan, "Statistical techniques for online anomaly detection in data centers," in *Integrated Network Management (IM), 2011 IFIP/IEEE International Symposium on*. IEEE, 2011, pp. 385–392.
- [50] Q. Guan, S. Fu, N. DeBardleben, and S. Blanchard, "Exploring time and frequency domains for accurate and automated anomaly detection in cloud computing systems," in *Dependable Computing (PRDC), 2013 IEEE 19th Pacific Rim International Symposium on*. IEEE, 2013, pp. 196–205.
- [51] S. N. Shirazi, S. Simpson, A. Marnierides, M. Watson, A. Mauthe, and D. Hutchison, "Assessing the impact of intra-cloud live migration on anomaly detection," in *Cloud Networking (CloudNet), 2014 IEEE 3rd International Conference on*, Oct 2014, pp. 52–57.
- [52] B. H. Sigelman, L. A. Barroso, M. Burrows, P. Stephenson, M. Plakal, D. Beaver, S. Jaspan, and C. Shanbhag, "Dapper, a large-scale distributed systems tracing infrastructure," Technical report, Google, Tech. Rep., 2010.
- [53] C. Wang, "Ebat: online methods for detecting utility cloud anomalies," in *Proceedings of the 6th Middleware Doctoral Symposium*. ACM, 2009, p. 4.
- [54] Y. Guan and J. Bao, "A cp intrusion detection strategy on cloud computing," in *International Symposium on Web Information Systems and Applications (WISA), 2009*, pp. 84–87.
- [55] T. Garfinkel, M. Rosenblum *et al.*, "A virtual machine introspection based architecture for intrusion detection." in *NDSS*, vol. 3, 2003, pp. 191–206.
- [56] A. V. Dastjerdi, K. A. Bakar, and S. G. H. Tabatabaei, "Distributed intrusion detection in clouds using mobile agents," in *Advanced Engineering Computing and Applications in Sciences, 2009. ADV-COMP'09. Third International Conference on*. IEEE, 2009, pp. 175–180.
- [57] S. N. Shirazi, S. Simpson, A. Gouglidis, A. U. Mauthe, and D. Hutchison, "Anomaly detection in the cloud using data density," in *IEEE International Conference on Cloud Computing*, 2016.
- [58] C. Goh, *A Generic Approach to Policy Description in System Management*. Hewlett Packard Laboratories, 1997.
- [59] S. N. Shirazi, S. Simpson, S. Oechsner, A. Mauthe, and D. Hutchison, "A framework for resilience management in the cloud," *e & i*

Elektrotechnik und Informationstechnik, vol. 132, no. 2, pp. 122–132, 2015.

- [60] C. C. L. Team, "Cloud security alliance cloud controls matrix v1.1," Tech. Rep., 2010.



Syed Noorulhassan Shirazi is a Research Associate in the School of Computing and Communications at Lancaster University, UK. His current research focuses on anomaly based challenge detection techniques in elastic cloud deployment scenarios. His research interests include security and resilience of computer networks and networked systems.



Antonios Gouglidis is a Senior Research Associate at Lancaster University, and currently involved in the EU FP7 funded project HyRiM. Previously, he has worked in academia as a Security Researcher; in industry as a Software Engineer; and in the public sector as an IT Training Consultant. He received his PhD in Applied Informatics from University of Macedonia, Greece; MSc in Mathematics from the Aristotle University, Greece; MSc in Computer Science from Lancaster University, UK; and BSc in IT

Engineering from the Alexander Technological Educational Institute of Thessaloniki, Greece. His research interests include security, resilience, access control, and formal methods.



Arsham Farshad is a Research Associate within the School of Computing and Communications at Lancaster University. His main research interests lie in monitoring, measurement and management in next generation of communication networks including both wireless and data networks. His recent research focuses on network services deployment and management for heterogeneous networks. He holds a PhD in Informatics from the University of Edinburgh.



David Hutchison is Professor of Computing at Lancaster University and founding Director of InfoLab21. He has served on the TPC of top conferences such as ACM SIGCOMM, IEEE Infocom, and served on editorial boards of Springers Lecture Notes in Computer Science, Computer Networks Journal and IEEE TNSM, as well being editor of the Wiley book series in Computer Networks and Distributed Systems. He has helped build a strong research group in computer networks, which is well known internationally for

contributions in a range of areas including Quality of Service architecture and mechanisms, multimedia caching and filtering, multi-cast engineering, active and programmable networking, content distribution networks, mobile IPv6 systems and applications, communications infrastructures for Grid based systems, testbed activities, and Internet Science. He now focuses largely on resilient and secure networking, with interests in Future Internet and also the protection of critical infrastructures including industrial control systems.