Access Control Issues in Utilizing Fog Computing for Transport Infrastructure

Stavros Salonikias[1], Ioannis Mavridis[1], Dimitris Gritzalis[2]
[1] Dept. of Applied Informatics, University of Macedonia
156 Egnatia st., 546 36, Greece
{salonikias, mavridis}@uom.gr
[2] Information Security and Critical Infrastructure Protection (INFOSEC) Laboratory,
Dept. of Informatics, Athens University of Economics & Business
76 Patission ave., Athens, 104 34, Greece
dgrit@aueb.gr

Abstract. The integration of the information and communication technologies of cloud computing, Software Defined Networking (SDN) and Internet of Things (IoT) into traditional transportation infrastructures enables the evolution of Intelligent Transportation Systems (ITS). Moreover, the specific requirements for real-time applications and service provision near to consumers introduce the utilization of fog computing as an extension of cloud. However, such a movement affects security aspects and poses new access control challenges. In this paper, we study the operational characteristics of a proposed ITS paradigm utilizing fog computing and identify corresponding access control issues. To address these issues in such a versatile and highly distributed environment, we present the key pointers of an attribute-based access control scheme suitable for fog computing. This paper aims to set a basis for further work in refining, verifying and validating the proposed solution.

Keywords: Critical Infrastructures, Transport Infrastructure, Intelligent Transportation System, Internet of Things, Fog Computing, Software Defined Networks, Access Control, ABAC, Context-awareness.

1 Introduction

The integration of information and communication technologies (ICT) into traditional transportation infrastructures enables the evolution of Intelligent Transportation Systems (ITS) [5] to support safe, efficient, sustainable and environmentally friendly transportation facilities. In ITS, the concept of device ubiquity [21] that refers to the ability to be present but not intrusive and yet interact with physical environment in a calm way, is extensively realized.

Actually, fabricating ICT in transport infrastructures using different kinds of devices (from sensors and actuators to computing systems embedded into vehicles) as well as various types of services to create a large transport ecosystem, depicts the concept of

Internet of Things (IoT) [16]. A "thing" is a stakeholder (e.g. a sign mark or a milestone) which has been assigned an identity (e.g. RFID metal id, IPv6 address) and can pervasively be present and interact with other things within the scope of a network [3].

The ICT infrastructure that supports and utilizes the IoT contains various types of network connections in order to supply services that are provided from datacentres residing in the cloud. However, cloud-based infrastructure realizes a centralized model for service consolidation and offering through wired backbone connections. The fact that in IoT environments there may be an enormous number of stakeholders requesting services, producing and consuming data, likely leads to congestion and arbitrary delays in servicing all requests. Additionally, there are cases where time available for data transfer and decision making process is so critical that must be counted near to zero (i.e. vehicle collision avoidance).

An architecture that extends cloud to facilitate service deployment near the edge of the network, meet latency requirements, achieve context-awareness and accommodate large numbers of nodes was introduced in 2012 by Cisco Systems under the term "fog computing". Fog realizes a highly distributed virtualized platform that provides processing, storing and networking services between the cloud and the edge of the network [4]. Thus, fog can be considered as an extension of cloud, stretching it to utilize resources from devices near the edge.

Network devices are usually autonomous, configured via a web or command line interface, whereas all high level organizational policies should be transformed into low-level multiple configurations [8]. The fact that many vendors provide management controllers does not drastically improve the situation since in real world infrastructures there exist multi-vendor and even end-of-support devices. Software Defined Networking (SDN) decouples the control plane from the forwarding plane [6]. Thus, network traffic paths are no longer configured in every device but are centrally managed by the SDN controller.

Due to importance in economy, transport infrastructures face multiple threats [18]. Moreover, a single incident can cause cascading effects [9] on other critical infrastructures due to dependencies among them [10]. More specifically, ITS mainly suffer from ICT threats. To mitigate relevant risks, effective countermeasures should be taken in order for confidentiality, integrity and availability to be ensured. Thus access to system resources must be controlled in an authenticated and authorized manner.

The main contributions of this paper is to study the particular functional and non-functional requirements posed on an ITS paradigm utilizing fog computing, spot the corresponding access control issues and propose an approach to address them via an effective access control system based on proper mechanisms.

The structure of this paper is as follows. In section 2 we present the ITS deployment that is utilising fog computing. Based on this paradigm, section 3 presents identified access control issues. In section 4, the stepping stones in defining a suitable attribute-based access control approach in fog computing are discussed. The paper concludes in section 5.

2 ITS paradigm

In this paradigm, we assume a transport environment supported by a proper ICT infrastructure to implement the ITS under consideration. In such an ITS, consumers (vehicles and humans as passengers or pedestrians) are provided with safety services (collision avoidance, emergency response), infotainment services (informatory services, advertising, multimedia entertainment content) and routing services (collision or congestion avoidance). We distinguish four main areas (fig. 1): Core ICT (CI), Road Side (RS), Vehicles and Humans (VH), and Sensors and Actuators (SA).

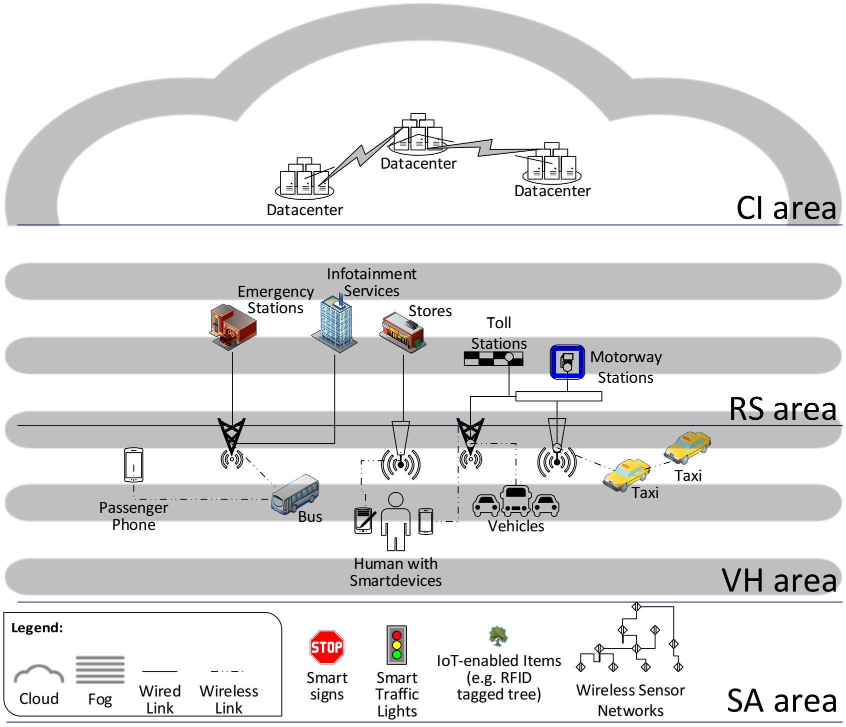

Fig. 1. Main areas of ITS paradigm

CI provides a solid platform for applications that are used to provide services to consumers, according to the Software as a Service (SaaS) model of cloud computing. Content and infrastructure management along with resources for data processing and data warehousing are provided. CI services are mainly implemented in the cloud, which is built using a distributed architecture over high-bandwidth wired links between datacenters to provide high availability of operations.

Service consumers are inhibiting the VH area. More specifically, consumers are considered to be vehicles appropriate equipped or humans (pedestrians or passengers) carrying smart devices (e.g. smartphones) connected to the ITS network. From the network point of view, consumers are mobile nodes, connected wirelessly to access stations and are arbitrarily entering and leaving the system. For the consumers to obtain authenticated access to services provided by the CI, connections from VH area to CI area are established. In a heavily inhibited road though, vehicle or human requests to CI for services can readily escalate in a large number.

Between CI and VH areas, in close vicinity to the latter, RS area is deployed to mainly facilitate network connectivity between service providers and consumers. Servers and network nodes are usually connected in a wired manner, sometimes using redundant paths. Applications deployed to this area provide services to near consumers. By doing so, services can easily be localized and directly access context parameters. Moreover, by shortening network distance and lessening connected users, network latency is lowered. In order to provide services of local interest, instead of distant datacenters, servers can be installed either in various service provider premises (in the RS area) or in the embedded computing systems of vehicles and the smart devices carried by humans (in the VH area). An example scenario of providing such services to consumers is advertising offers in gas from stations in a particular city. This information is important to vehicle drivers in the same city but rather useless to a driver located many miles away.

In the SA area IoT enabled devices with temporal or permanent presence on the road are assumed. These devices are wirelessly connected using short-range connections (e.g. Bluetooth, ZigBee) and can be considered stationary nodes in relevance to the carrier (i.e. the speed meter of a bus). A sensor network attached to a car is stationary regarding the car, just as an actuator attached on the side of the road (e.g. a smart traffic light).

In a routine scenario, a cloud service deployed in CI area is used to provide alternative routes in case of traffic blockage. For the service to be aware of traffic conditions (contextual information), traffic data should be provided from sensors near the problematic area. When a vehicle needs to use the service, it has also to provide its current location and destination. Then, the cloud service processes and compares the data with similar datasets from other vehicles. Finally, the computed routes are transmitted back to vehicles. However, the whole process required for transferring and processing may take a significant amount of time making the returned information possibly obsolete. Regarding a first case when the vehicle is a tourist bus, if alternative routes are no longer possible (vehicle has already been blocked in traffic before receiving alternative routes), tourists may have to bear a delay in their vacation. In a second case, where the vehicle is an ambulance (or any other emergency vehicle) urging to a destination in an urban environment, alternative routes must be received as soon as possible, before the vehicle become stacked and rerouting is no longer possible. If additional contextual information like weather or road condition should also be sent to the service, computation time and data transfer time may become longer.

The above scenario indicates that direct accessing of services from CI suffers from the following limitations:

- CI area is physically located away from consumers,
- network congestions and arbitrary delays are likely to be happen when a service is simultaneously offered to a vast number of consumers,

- quality of service can be negatively affected due to increased latency from network delays while a real-time service is being requested,
- increased volume of transmitted data after context-aware applications request for contextual information to be transmitted to the cloud,
- limited scope when a service is provided in a restricted area with local interest only.

To overcome the above limitations, fog computing can be utilized to extend cloud functionality in RS and VH areas, using resources from both. By extending cloud with fog computing, services can take advantage of both RS and VH area characteristics to achieve (1) low latency, (2) application localization and decentralization, (3) direct interference between applications and consumers (vehicles or humans). It is worth to make clear that by contributing their ICT resources to fog, vehicles and humans can act either as service consumers or providers.

More specifically, using fog computing to deploy services providing alternative routes, data would be uploaded to a nearby fog server, less congested than a cloud server. At the same time, a fog server located close to the scene could already be aware of contextual information like traffic conditions. Accordingly, total time for uploading, processing and downloading rerouting instructions to vehicle would be less.

Another fog advantage is that it can be used to create federations [4] between different administrative domains. Luan et al [12] present a situation where a store using its fog server distributes a digital flyer. Vehicles near roadside access points get the flyer and push it further to other vehicles. Using federations, the same store can simultaneously provide additional services from another federated domain.

Fog computing is usually implemented over heterogeneous networks. This happens because in the RS area networks usually preexist. To substitute all devices to a single-vendor infrastructure is usually not an option mainly due to cost. Device heterogeneity though leads to high administrative costs and may pose a negative trend on fog adoption. To create a vendor-independent environment that permits unified management, SDN technology is an interesting option to be implemented [8]. Using SDNs in RS area allows for a centralized orchestration [19] of the fog underlying network via SDN controllers. The SDN Controller (control plane) implements two interfaces. The Southbound Interface (SBI) which logically connects controller with SDN-enabled hardware devices (data plane) and the Northbound Interface (NBI) which exposes controller API (application plane) to make network flows programmable by external applications.

In Table 1 we summarize the main characteristics that describe each area behavior in the ITS environment.

Table 1. Main characteristics for each ITS area

	CI	**RS**	**VH**	**SA**
Network area	Core	Field	Edge	Field - Edge
Architecture/Topology	Centralized	Distributed	Mobile	Wireless Sensor Networks (WSN) [2]
Bandwidth	Very High	High	Low	Very low
Latency	Low	Low	High	Low

Network Availability	High	Medium	Low	Low
Number of connections	Few	Many	Limited	Limited
Context awareness	No	Yes	Yes	No
Device heterogeneity	No	Yes	Yes	No
Geographic distribution	Moderate	High	Very High	Low
Access Medium	Wired	Wired	Wireless	Wireless
Mobility	No	No	Yes	No
Proximity to consumers	Far	Near	Near	Near
Implementation location	Datacenters	Roadside (Buildings or/and outdoors)	Outdoors	Outdoors
Connections	Controlled	Controlled and arbitrary	Arbitrary	Managed within WSN

3 Access Control Issues

The utilization of fog computing in developing ITS, as presented in the paradigm of the previous section, results in arising a number of security and privacy considerations that are mainly regarding controlling permissions on the use of resources and services. Access control seeks to prevent activity that could lead to breach of security by ensuring in general that accesses to system's resources and services occur according to the rules defined in corresponding security policies [17]. Towards implementing a suitable access control solution various functional and non-functional system features must be taken under consideration.

In the dispersed fog environment presented in the previous section, vehicles and humans are entering and leaving domains arbitrarily. Service consumers and providers can act as access control subjects and objects interchangeably. In a gas station scenario, stores (located in RS area) are able to push information to vehicles (in VH area). Then, vehicles can pass this information to other vehicles etc. This indicates that in the ITS fog entities can dynamically operate either as subjects or as objects. The crossing of administrative domains is of special interest, as a vehicle or human moves, and appropriate administrative functions are needed. An administrative domain is a set of network entities on which a single security policy is employed by a single administrative authority [20].

Services are deployed by local providers from RS area, wherein a number of different administrative domains are interoperating and usually forming federations. These administrative domains provide, besides their own, services from other domains of the same federation. An important factor is that services can be deployed by local providers in RS area without the intervention of a particular administrative authority. Service owners should be able to set appropriate access control policies to limit services consumption. The propagation of access control policies can be delegated to administrative authorities, in order to enable interoperability between the rest of service providers and the whole ITS environment

Ambient environment plays a key role in an ITS. For example, in a circumstance with a heavy rain, it is not desirable to distract the vehicle driver with non-critical information, such as the nearest museum (unless he specifically requests for such information). Hence, context awareness becomes a critical feature of access control. Abowd et al. [1] define context as any information that can be used to characterize the situation of an entity or environment. An entity is a person, place, or object that is considered relevant to the interaction between a user and an application, including the user and application themselves. Accordingly, a system is context-aware if it uses context to provide relevant information and/or services to the user, where relevancy depends on the current task of the user [1].

In order for an ITS to support context-aware services, information related to personal data (i.e. current location of vehicle or human) is transferred, processed and stored in the fog and the cloud. Thus, the transport infrastructure should ensure protection of privacy. Marx [13] perceives privacy violation when one of the four following borders is crossed:

- Natural Borders that protect a person from direct observation.
- Social Borders that protect a person's social life.
- Spatial Borders that protect periods of a person's life.
- Transitory/Ephemeral Borders that protect fleeting moment actions.

These borders indicate that privacy can be extended in many directions. Ephemeral borders can be crossed as access control logs expose a person's activity at a specific moment. Moreover, natural borders can be crossed when access control request sequence can indicate a set or actions. To protect privacy, access control should be implemented following a method where person's identity is not used or the contextual information cannot reveal it. Another important consideration is that privacy concerns should not prevent accountability.

The high mobility of sensors and actuators in the SA area, in conjunction with their limited power and computation resources, should also be taken under consideration during access control system development.

Table 2 summarizes access control issues in fog computing.

Table 2. Access control issues in fog computing

Privacy violation	Fog decentralization may require data interchange between administrative domains. Private data should be protected.
Coherency	A unique and coherent access control system should be end-to-end applied through multiple administrative domains.
Context-awareness	Capturing, transferring, processing and storing of rich and quickly changing contextual information (e.g. weather conditions, temperature, time, etc.) should be effectively managed to support access control decisions.
Resource restriction	Access control system should not drain power or resources from devices with limited capacity

Network availability	Access control should be able to provide a level of functionality even in case of network unavailability
Decision latency	The system should minimize time required to grant or deny a request. Delays in the process may lead to undesirable effects for human safety and infrastructure integrity.
Management	Policy management, which includes the ability to create, update and delete policies and to notify stakeholders for changes, should be supported.
Accountability	Tracks concerning malicious activity should not be lost across administrative domains.

4 Proposed Approach for Access Control in Fog Computing

Based on the operational and security characteristics of utilizing fog computing in transport infrastructure, as presented and discussed in the previous sections, we propose the main features of a proper access control system.

4.1 Attribute-Based Authentication and Authorization

In traditional ICT systems, subjects are assigned with an identity. During user identification this identity is presented to the system and then is verified during authentication, usually by using a password. After a successful authentication process, access control involves the authorization process which is to decide whether an access request should be permitted or denied.

For the arbitrarily changing number of subjects and objects in VH area, assigning and verifying identities for every entity is not possible. On the one hand not all services might be publicly available. On the other hand some services might require consumer identification. Even for publicly available services, there are circumstances when access should be denied based on the contextual information. This urges the need for an access control system that can consider information describing subjects, objects and the context of operation.

NIST has released the Special Publication (SP) 800-162 [7] to provide a definition of Attribute-Based Access Control (ABAC), as an access control approach based on attributes, and a guidance in implementing it within an organization. Attributes are characteristics, defined as name-value pairs, which can contain information about subjects, objects and context. Context attributes, or environmental conditions, allow ABAC implementations to be context-aware, thus making it an ideal candidate for ITS applications, where context is a factor that affects the entire system behavior.

Identity-based authentication is not a prerequisite for ABAC. Nevertheless, when required, identities can be used provided they have been assigned to subjects as attributes. To ensure certified exchanging of attributes, the utilization of proper attribute certificates has already been proposed in previous work [14].

In an attribute-based access control system, authorization is performed according to a security policy that is mainly defined on the basis of subject and object attributes instead of any identities. Security policy refers to the set of rules, laws and practices that regulate how an organization manages, protects, and distributes sensitive information [15]. ABAC utilizes two sets of policies. Digital Policies (DP) and Meta Policies (MP). DPs dictate access control rules that access control system enforces. MPs are used for managing DPs. An example of MP is the definition of priorities that should be assigned for the case of conflicting DPs.

4.2 Reference Monitor Implementation

Reference monitor (RM) is consisted by the Policy Decision Point (PDP) and the Policy Enforcement Point (PEP). Access decisions are made in PDP and are then fed into the PEP that enforces them, thus granting or denying subject access requests against objects. These decisions are based on DPs and corresponding MPs.

In traditional computing environments, the reference monitor is usually integrated with the objects to be protected or implemented in a central location to process all requests for distributed objects [11]. In a dynamic and highly distributed ITS environment, where access control decisions should be taken instantly, these approaches are both not viable. The former requires processing power that is not available in all entities and the latter would impose a significant latency for all access requests and access decisions to be sent over the network. Moreover, the number of requests can escalate to as high as to impose further delays in processing time. To overcome this issues, we propose a distributed reference monitor, based on ABAC, as presented in Figure 2.

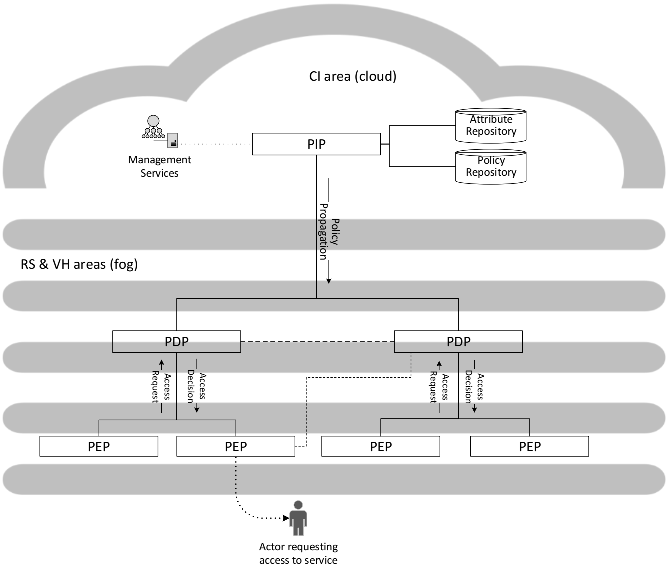

Fig. 2. Proposed ABAC implementation

In the proposed ABAC implementation, security administrators set policies that are stored in a logical Policy Information Point (PIP). Policies can be set using User Interface (UI) or another suitable method (e.g. via a web service). We characterize PIP as logical since it is presented as a single entity although it is more likely to be implemented in a distributed manner within the CI area. PIP contains all policy and attribute information for one or multiple domains and propagates policies to PDPs deployed mainly in RS area. A PEP enforces access decisions and can be implemented on every single device. For example, PEP can reside on a network switch, where it disables a port if attributes of a requestor subject lead to a deny decision by a PDP located in a fog server. PEP does not require significant computing resources to consume, since most of the required computation takes place in the PDP.

4.3 Policy Propagation

For system continuance and redundancy, PEPs can be bound to one or more PDPs. This can only be functional when all PDPs are simultaneously informed of any policy change. PDP policy synchronization requires a proper policy propagation control method. As mentioned above, MPs contain information for managing DPs. All policies are created, modified or deleted by policy managers in the scope of an administrative domain.

In case of a single PDP, all changes in policies stored in a PIP can directly be propagated to the PDP and then immediately be effective for the whole system. However, in case of distributed PDPs, where communication networks may impose delays due to outages or breakdowns, a change in the policy set (PIP) should be propagated based on specific rules that take under consideration, among others, the current network conditions. For example, it may be preferable not to send policy updates to any PDP unless all required PDPs are reachable.

To cope with such issues, we propose the introduction of propagation rules described by a new policy set, called Propagation Policies (PP). PPs are policies that define how PIP policies are updated, propagated to PDPs and exchanged between PDPs. Setting the propagation policies alone is not enough. DPs and MPs should be securely transmitted and verified. For this purpose, we aim at extending the definition and use of access rule certificates [14].

4.4 Offline Operation

A major concern in high-available systems is ensuring network availability. This is usually done by providing redundant links. Nevertheless, a case of network failure is still a possibility, especially in a mobile environment. Failures can affect either PIP-to-PDP or PDP-to-PEP communication.

The intended attribute-based access control system should support redundant links between PIP and PDPs, as well as between PDP and PEPs. When a connection between PDP and PIP fails, the PDP that cannot reach PIP should be able to use a neighbor PDP to fetch policy information required to answer requests. For the case of a complete communication failure, a small set of default policies could be locally stored. On the other hand when PEP cannot reach any PDP, a default action should be enforced.

5 Conclusion

Transport infrastructure, apart from the significant economic impact that justifies the characterization critical, accommodates a vast number of consumers as vehicles and humans. Integrating ICT to transport infrastructure can improve the overall experience and advance safety.

In this paper, we described an ITS paradigm that results in dividing the ICT environment into four areas (CI, RS, VH and SA). These areas facilitated our study for the adoption and utilization of fog computing in ITS scenarios that introduce specific requirements. Summarizing our study, fog enables ITS to extend services from the cloud, near to consumers, overcoming problems that a centralized approach induces. Based on the ITS paradigm, a number of specific access control issues were located and discussed. Our proposed approach towards addressing these issues is to develop a distributed access control system that instead of identities only, will utilize attributes for authentication and authorization purposes.

Our future research effort aims at refining the access control deployment for single or multiple ITS domains, along with providing a solid theoretical basis, using appropriate formal methods, on which the proposed solution can be validated and verified. This will include effective solutions for reference monitor implementation, secure policy propagation and support for offline operation, with an utmost goal of a real-world implementation.

References

1. Abowd, G.D., Dey, A.K., Brown, P.J., Davies, N., Smith, M., Steggles, P.: Towards a better understanding of context and context-awareness. In: Handheld and Ubiquitous Computing. pp. 304–307. Springer (1999).
2. Akyildiz, I.F., Su, W., Sankarasubramaniam, Y., Cayirci, E.: Wireless sensor networks: a survey. Computer Networks. 38, 393–422 (2002).
3. Atzori, L., Iera, A., Morabito, G.: The Internet of Things: A survey. Computer Networks. 54, 2787–2805 (2010).
4. Bonomi, F., Milito, R., Zhu, J., Addepalli, S.: Fog Computing and Its Role in the Internet of Things. In: Proc. of the MCC Workshop on Mobile Cloud Computing. pp. 13–16. ACM, USA (2012).
5. ETSI - Intelligent Transport, http://www.etsi.org/technologies-clusters/technologies/intelligent-transport. (last accessed 20/05/2015)
6. Hakiri, A., Gokhale, A., Berthou, P., Schmidt, D.C., Gayraud, T.: Software-Defined Networking: Challenges and research opportunities for Future Internet. Computer Networks. 75, Part A, 453–471 (2014).
7. Hu, V.C., Ferraiolo, D., Kuhn, R., Schnitzer, A., Sandlin, K., Miller, R., Scarfone, K.: Guide to Attribute Based Access Control (ABAC) Definition and Considerations. National Institute of Standards and Technology (2014).

8. Kim, H., Feamster, N.: Improving network management with software defined networking. IEEE Communications Magazine. 51, 114–119 (2013).
9. Kotzanikolaou, P., Theoharidou, M., Gritzalis, D.: Assessing n-order dependencies between critical infrastructures. International Journal of Critical Infrastructures. 9, 93–110 (2013).
10. Kotzanikolaou, P., Theoharidou, M., Gritzalis, D.: Interdependencies between Critical Infrastructures: Analyzing the Risk of Cascading Effects. In: Critical Information Infrastructure Security. pp. 104–115. Springer (2013).
11. Lampson, B., Abadi, M., Burrows, M., Wobber, E.: Authentication in Distributed Systems: Theory and Practice. In: Proceedings of the 13th ACM Symposium on Operating Systems Principles. pp. 165–182. ACM, USA (1991).
12. Luan, T.H., Gao, L., Li, Z., Xiang, Y., Sun, L.: Fog Computing: Focusing on Mobile Users at the Edge. arXiv:1502.01815 [cs]. (2015).
13. Marx, G.T.: Murky conceptual waters: The public and the private. Ethics and Information Technology. 3, 157–169 (2001).
14. Mavridis, I., Georgiadis, C., Pangalos, G.: Access-rule certificates for secure distributed healthcare applications over the Internet. Health Informatics Journal. 8, 127–137 (2002).
15. Mavridis,I., Pangalos,G.: Security Issues in a Mobile Computing Paradigm. Communications and Multimedia Security. Vol.3, 60-76 (1997)
16. Miorandi, D., Sicari, S., De Pellegrini, F., Chlamtac, I.: Internet of things: Vision, applications and research challenges. Ad Hoc Networks. 10, 1497–1516 (2012).
17. Sandhu, R.S., Samarati, P.: Access control: Principle and practice. IEEE Communications Magazine. 32, 40–48 (1994).
18. Theoharidou, M., Kandias, M., Gritzalis, D.: Securing Transportation-Critical Infrastructures: Trends and Perspectives. In: Georgiadis, C.K., Jahankhani, H., Pimenidis, E., Bashroush, R., and Nemrat, A. Al (Eds.) Global Security, Safety and Sustainability & e-Democracy. pp. 171–178. Springer (2012).
19. Truong, N.B., Lee, G.M., Ghamri-Doudane, Y.: Software defined networking-based vehicular Adhoc Network with Fog Computing. In: 2015 IFIP/IEEE International Symposium on Integrated Network Management (IM). pp. 1202–1207 (2015).
20. Vázquez-Gómez, J.: Multidomain security. Computers & Security. 13, 161–184 (1994).
21. Weiser, M.: The Computer for the 21st Century. Scientific American. 265, 94–104 (1991).