Towards an Open, Secure, Decentralized and Coordinated Fog-to-Cloud Management Ecosystem

D2.2 Tracking Scientific, Technology and Business Trends (Version 2)

Project Number	**730929**
Start Date	**01/01/2017**
Duration	**36 months**
Topic	**ICT-06-2016 - Cloud Computing**

Work Package	**WP2, Technology survey, business models and architectural definition**
Due Date:	*M21*
Submission Date:	*Date of actual submission –02/10/2018*
Version:	*1.0*
Status	*Final*
Author(s):	Alec Leckey, John Kennedy (INTEL), Jens Jensen, Shirley Crompton (STFC), Anna Queralt (BSC), Jasenka Dizdarevic (TUBS), Matija Cankar (XLAB), Roi Sucasas Font, Lara Lopez Muniz (ATOS) Glauco Mancini (Engineering) Eva Marin Tordera (UPC), Denis Guilhot (WOS), *Cristóvão Cordeiro (SIXSQ)*
Reviewer(s)	Ana Juan Ferrer (ATOS) Xavi Masip (UPC) *Admela Jukan (TUBS)*

Keywords
Fog, Cloud, Edge

Project co-funded by the European Commission within the Seventh Framework Programme		
Dissemination Level		
PU	Public	**X**
PP	Restricted to other programme participants (including the Commission)	
RE	Restricted to a group specified by the consortium (including the Commission)	
CO	Confidential, only for members of the consortium (including the Commission)	

This document is issues within the frame and for the purpose of the mF2C project. This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 730929. The opinions expressed and arguments employed herein do not necessarily reflect the official views of the European Commission.

Version History

Version	Date	Comments, Changes, Status	Authors, contributors, reviewers
0.1	2/08/2018	Initial ToC and doc structure	Alec Leckey (INTEL)
0.2	18/9/2018	First draft with all content submitted	All Partners
0.3	26/9/2018	Added missing business section, ch4	Lara Lopez Muniz (ATOS)
0.4	27/9/2018	Conclusion, comments addressed and formatting changes	Alec Leckey (INTEL)
0.5	28/9/2018	Additional reviewer comments addressed	Lara Lopez Muniz (ATOS), Alec Leckey (Intel)
0.6	01/10/2018	Quality checks	María Teresa García González (ATOS)
0.7	1/10/2018	Minor edits to formatting	Alec Leckey (Intel)
0.8	2/10/2018	Quality checks	María Teresa García González (ATOS)

Table of Contents

Version History 3
Table of Contents 4
List of figures 5
List of tables 5
Executive Summary 6
1. Introduction 7
1.1. Purpose 7
1.2. Structure of the document 7
1.3. Glossary of Acronyms 8
2. Scientific Trends 9
2.1. Service management, resource management, end-devices 9
2.2. Scientific trends coming from the HPC area 10
2.2.1. Data management trends 10
2.3. Applications in different science areas, data centres, big data processing 12
2.4. Security trends 13
2.4.1. Artificial intelligence and machine learning 13
2.4.2. Data protection 14
2.5. Resource management and QoS 16
2.6. Convergence of AI and computing 17
2.7. Key Takeaways 18
3. Technology Trends 20
3.1. Tools, platforms, IoT 20
3.1.1. Cloud management tools 20
3.1.2. IoT Management tools 20
3.2. Technology trends coming from HPC 21
3.2.1. Data management trends 21
3.2.2. Programming model's trends 22
3.3. Cloud Orchestration Platforms, Virtualization, Containers 22
3.4. Role of standards in technologies 23
3.5. Technology trends in edge computing 24
3.5.1. Reference solutions 24
3.5.2. Containers in the edge 25
3.5.3. Edge and IoT 25
3.6. Key Takeaways 26
4. Business trends 27
4.1. Cloud, fog and edge computing 27
4.2. Internet of Things 31

4.3. Big data and IoT 34
4.4. Security Trends 37
4.5. Digital Business 38
4.6. Key Takeaways 39
5. Conclusions 41
References 43

List of figures

Figure 1 US Cloud market forecast 2014-2025 [GVR1] 27
Figure 2 Global Cloud storage market 2017-2025 [INK1] 28
Figure 3 Emerging Market Frameworks and Definitions [WBR1] 28
Figure 4 Edge computing market size forecast in the US from 2017 to 2025, by segment (in million $) [STA1] 29
Figure 5 Size of Fog computing market opportunity by vertical market, 2019 and 2022 [451R1] 30
Figure 6 Artificial Intelligence – Direct & Enabled Revenue, 2014-2025 (USD Million) [GVR4] 30
Figure 7 Global, publicly announced IoT projects [IoT analytics:2018] 32
Figure 8 Sectorial map of IoT patents [Relecura:2017] 33
Figure 9 Market pulse report IoT, 2017 [GRWTH17] 34
Figure 10 Source: recap-project.eu [RECAP] 35
Figure 11 McKinsey, Fuelling growth through data monetization, December 2017 [Gartner2017] 36
Figure 12 Digital Business Lenses for Digital Business Opportunity [GAR5] 39

List of tables

Table 1. Acronyms 9

Executive Summary

The objective of this deliverable is to track scientific, technology and business trends in the area of Fog and Cloud computing that are relevant to the mF2C project. This deliverable provides an update to trends necessary to be aware of in the project that were documented in deliverable D2.1 submitted in M03. As per last time, each chapter ends with a *"key takeaways"* section summarizing the main points of focus, helping readers to understand priorities in the project. This is the second version of the deliverable and is aligned to iteration 2 (IT-2) of the project. A final report will be submitted in M34 which will include global reporting on technology, business models and scientific trends.

1. Introduction

1.1. Purpose

We continue to use the term Fog-to-Cloud (F2C) to refer to the resources created through the merging of cloud and fog computing, which in turn creates the need for new, open and coordinated management ecosystems. This is where mF2C will provide value with its management framework that has been designed to be open, secure, decentralized, multi-stakeholder, including novel programming models, privacy and security, data storage techniques, service creation, brokerage solutions, SLA policies, and resource orchestration methods. It is the intention of this project to create a proof of concept system and platform based on an innovative distributed system architecture validated with real world use cases.

WP2 is tasked with studying state-of-the-art fog, cloud, network and IT infrastructure technologies, with a goal of identifying technologies relevant for the deployment of the mF2C management framework, i.e. sensors, smart end-devices, connectivity, and advanced cloud services. D2.1 [D21] provided an overview of all the scientific, technology and business trends in fog computing relevant to the project and which awareness was required. This was an initial version of the deliverable (v1) and was aligned to iteration 1 (IT-1) of the project. This is the second version of that deliverable and is aligned to iteration 2 (IT-2). Given the short length of time between reports (18 months), the technology trends described in the initial report remain both valid and relevant today. The drive towards Fog/Edge computing continues to be unchanged and relates to the massive amount of IoT generated data, an ability to aggregate and reduce data at its source, to ensure real-time decision making, to ensure data anonymization and privacy protection, and finally an increased autonomy.

1.2. Structure of the document

Chapter 2 reviews the scientific trends of Fog and Cloud computing, beginning with an assessment of contributions relating to service management, and resource management emphasizing the need for decentralized and hierarchical solutions in order to meet the challenges brought up by fog computing. While the convergence between High Performance Computing (HPC) and cloud and big data technologies remain valid, attention has progressed to solving latency issues and related challenges that arise in fog/edge infrastructures. These relate to task scheduling and offloading mechanisms due to issues related to the high mobility of the device, the limited availability of energy of the devices and the impact of the network on the performance of the entire framework. The exploitation of Fog computing in new domains provides exciting opportunities to solve new real-world use cases, e.g., flash flooding or forestry monitoring systems. Any technology that enables the collection, processing, and sharing of personal data, will pose privacy risks. This increasing number of connected devices therefore expands the attack surface of any organisation, so proposed solutions are documented.

AI algorithms can benefit from the parallelism of mF2C architecture, allowing the distributed implementation of Machine Learning (ML) algorithms. The project reviewed some works that propose the use of ML, fuzzy neural networks, Markov models or genetics algorithms to collect and process data coming from sensors in fog nodes and/or cloud.

Chapter 3 looks at technology trends evaluating different tools and platforms currently available that enable the management of features such as storage, compute, machine instances, and containers. Previously, we reviewed the technological trends regarding the management of Cloud, Fog and IoT devices, and while still valid, we update the latest trends in Cloud and IoT management. Again with HPC, we focus on data management solutions for the edge to cloud continuum as this new research area – the combining of edge computing and HPC - is now emerging, with many proposals and proof-of-concept evaluations, but not many implemented solutions. We put special emphasis on containers and their orchestration as this project heavily relies on this technology, e.g., all mF2C services running in IT-1 were based on Docker containers. So the confirmation of Kubernetes as the leading orchestration engine for containers is noticeable. The OpenFog Consortium reached a milestone the IEEE published IEEE 1394 standard for Adoption of their OpenFog Reference Architecture for Fog Computing. The new standard addresses the need for an end-to-end, interoperable solution positioned along the things-to-cloud continuum. Finally, we discuss reference architecture solutions in edge computing.

Chapter 4 reviews cloud, fog and IoT trends with respect to business needs. The benefits provided by these technologies will be translated into interesting business opportunities for utilities and manufacturers in order to reduce Operating Expenditure (OpEx) and Capital Expenditure (CapEx) investments. We review the latest business surveys showing IoT only starting to gain momentum in the enterprise in last 2 years, the evolution towards edge and fog computing, and the IoT technology trends in 2018-2019 appear to be blockchain and Artificial Intelligence (AI). We see market estimates on the value IoT will deliver differing across multiple sources, with expectations close to US$6.5 trillion. This is based on the number of connected devices reaching between 20 and 30 billion by 2020. We review IoT related patent filings, and finish with a review of blockchain in edge/fog scenarios.

Finally, we conclude with the key takeaways this document provides for the project. These include the areas of focus and where existing solutions can help accelerate development of the key components of the project architecture. A final version of this report is due in M34 which will include global reporting on technology, business models and scientific trends.

1.3. Glossary of Acronyms

Acronym	Definition
ACO	Ant Colony Optimization
AI	Artificial Intelligence
ANN	Artificial Neural Network
CAGR	Compound Annual Growth Rate
CapEx	Capital Expenditure
DC	Data Centre
DDoS	Distributed Denial of Service
DQL	Deep Q-learning
F2C	Fog-to-Cloud
FE	Functional Encryption
GA	Genetic Algorithms
GDPR	General Data Protection Regulation
HE	Homomorphic Encryption
HPC	High Performance Computing
IIoT	Industrial Internet of Things
IoT	Internet of Things
M2M	Machine to Machine
ML	Machine Learning
MSCOG	Multi-Slot Computational Offloading Game
NVM	Non-Volatile Memory
OGC	Open Geospatial Consortium
OpEx	Operational Expenditure
OTA	Over the Air
PET	Privacy Enhancing Technology
RoI	Return on Investment
SIEM	Security Incident Event Management
SMC	Secure multi-party computation
SOC	Security Operations Centres
SPA	Student Project Allocation
VM	Virtual Machine
VNF	Virtual Network Functions

Table 1. Acronyms

2. Scientific Trends

In this section, we give an overview of the scientific trends significant for the mF2C. We extend the section from the previous version of deliverable D2.1 submitted in M03 with recent trends or changes, but it should be noted that many trends that were relevant at the time of the submission remain the same. The section includes scientific trends in the areas of service management, resource management, end-devices, HPC (including data management and programming models trends), science applications, security, convergence of AI and computing.

2.1. Service management, resource management, end-devices

In this subsection, we update the recent trends observed in the literature dealing with the resource management in fog and cloud environments.

Aligned with what has been reported in D2.1 [D21], recent works draw attention to the fact that most of the challenges faced while designing resource management solutions for fog computing systems originate from the inherent fog characteristics, including heterogeneity, dynamicity, geo-distribution, and the existence of multiple entities owning the devices taking part in the fog system. In addition, most works try to optimize a common set of metrics, consisting in energy efficiency, and quality of service (mostly in terms of latency) as well as load balancing.

It is worth highlighting that there is a general emphasis on the need for decentralized and hierarchical solutions in order to meet the challenges brought up by fog computing, which is aligned with the key mF2C project principles. In fact, authors in [JD18] propose a decentralized algorithm allowing devices to coordinate their periodic offloading decisions to achieve an efficient management of communication and computing resources in a fog computing system. A different view is presented in [WZC17] where a coordinator server is added to ensure cooperative operation between different Local Fog Servers (LFSs) in a fog network. The coordinator provides an inter-fog resource management with the aim of reducing the dropping rate in congested LFSs, whereas the LFSs themselves deal with the intra-fog resource management with the aim of improving the energy efficiency.

Common topics that are jointly addressed with resource management in a fog computing context are matching, mapping and allocation such as the works in [WSM18], [TD17] and [FMG18]. More specifically, Fogernetes [WSM18] is a platform based on an existing container management tool i.e. Kubernetes and it allows matching the requirements of fog application components with device capabilities by using a labelling system. Within such a system, labels can be either assigned automatically, based on hardware capabilities or based on specific user-defined behaviour and they can include geographical location, possible device extensions, expected performance as well as connectivity options. Authors in [TD17] present an algorithm for mapping IoT application modules within a Fog-Cloud infrastructure. The main idea behind this algorithm is to first place application modules on eligible nodes in the Fog layer. Once those nodes are overloaded or if no eligible node is found in the fog, the placement is done in the Cloud layer. Simulation results highlight the obtained gains in terms of application latency, network usage and energy consumption, when compared with traditional cloud-only placement strategies. A similar work that rather considers more layers in the hierarchy can be found in [FMG18] where the concept of “communities” is used to represent a group of fog nodes. It is based on a hierarchical structure where the smallest (i.e. lowest-layer) communities are the ones where the fog nodes are directly connected to nearby access points whereas higher-level communities represent wider geographical areas. Therefore, in order to optimally allocate a VM to a given service to minimize the user-fog node latency, it is necessary to find the smallest community that contains the AP that the user is connected to and that is capable accommodating the user request. Then, within the chosen community, the VM allocation could be based either on traditional load balancing or on consolidation techniques.

When it comes to formal modelling of resource management problems in fog computing systems, a trend that is increasingly gaining in popularity consists in using game theoretical approaches. For instance, the work in [KSM18] proposes a potential game-based approach for optimal resource-demand management in the specific case of a vehicular fog computing environment. Within this game, decisions are made based on the demands' priorities, the required CPU capacity as well as the vehicle's current energy, thus allowing to minimize the latency and to optimize the computing and energy resource utilization. Authors in [ZZN17] present a three-layer game framework, dealing with the different interactions between the stakeholders in a fog-cloud computing system. More specifically, the problem of finding an optimal matching between the physical resources offered by the fog nodes and the virtualized resources requested by the clients is modelled based on the "student project allocation" problem and the SPA-(S,P) algorithm is used to obtain a stable matching result. In addition, the work presented in [JD18] proposes a so-called Multi-Slot Computational Offloading Game (MSCOG) to obtain decentralized offloading decisions, while minimizing the response times and the corresponding energy consumption.

Finally, it is worth to mention that regarding the management of IoT end devices, what was reviewed in the previous version of this deliverable, [D21], is still relevant. And the main conclusion was that although there are many contributions in the area of IoT standards, there is not a widely accepted consensus on the solution to be adopted. In general, we can say that the IoT management reviewed proposals rarely address the IoT management from a whole entire perspective, including data, resources, service, network, etc. The same conclusion we achieved when reviewing the main contribution regarding addressing and naming of IoT end devices, there is no well adopted and final solution for addressing this high heterogeneity of devices.

2.2. Scientific trends coming from the HPC area

2.2.1. Data management trends

While the trends in the HPC area reported in D2.1 [D21] regarding new storage devices such as non-volatile memories, or the convergence between HPC and cloud and big data technologies are still valid, a lot of attention has been paid in the last months in solving the latency issues and challenges that arise in fog/edge infrastructures.

A preliminary work in this direction is [CLP17], where a set of storage requirements in a fog/edge computing infrastructure are identified, namely low access time, network containment between sites, availability of data in case of partitioning, and support for user's mobility. In this context, this work evaluates three popular distributed storage systems used in datacentres with the goal of determining whether they would be suitable in a fog/edge setting. The conclusion is that these systems encounter difficulties to scale and are not well suited for the local activity expected in a fog context.

While edge and cloud systems share similar needs such as scalability and fault tolerance, significant differences exist between their requirements [RG18]. First, cloud systems are hosted in datacentres which guarantee a certain degree of connectivity and latency between nodes, while in the edge they use much more unstable wireless communication links. Also, in the cloud data is distributed by design to accommodate large data sizes, while in the edge data is distributed inherently due to the distributed nature of the data sources. Finally, nodes at the edge are much more heterogeneous and diverse than in a cloud datacentre.

Thus, solutions for data management in a fog-to-cloud environment are being investigated in projects led by HPC institutions, such as the recently started CLASS project [CLASS], which aims to develop a novel software architecture to help big data developers to combine data-in-motion and data-at-rest analysis by efficiently distributing data and process mining along the compute continuum (from edge to cloud) in a transparent way, while providing sound real-time guarantees.

Also, new preliminary solutions towards a holistic view of data management from the edge to the datacentre are being proposed in emerging workshops in flagship HPC conferences related to these topics, such as USENIX HotCloud and HotEdge. In the following we summarize these proposals.

A data-centric communication approach is designed in [PA+18] by enhancing network connectivity with local storage services at the edge of the network, in order to buffer data generated at the edge prior to synchronization with the cloud. These local edge data repositories enable users to upload data to local storage resources as they move along. This vision paper focuses on the network infrastructure level, and thus it is complementary to other approaches that deal with data closer to the application level.

At a higher level of abstraction some application-specific proposals exist, such as [RG18], which assumes a specific kind of application running on the infrastructure, in this case machine vision applications. Assuming this prior knowledge, an architecture optimized for this kind of applications is proposed. However, this does not necessarily fit other kinds of applications, and implies underutilizing the powerful resources at the edge, which are required for computation-intensive applications such as this one. Other examples are [BH+18, DG+17].

There exist more general approaches such as CloudPath [MS+17], which enables different kinds of applications, including workloads that aggregate data, such as IoT applications, or services that cache data and process information at different layers. This proposal is based on Cassandra [LM10], a key-value store that provides eventual consistency between replicas, and this is the only kind of consistency that this solution offers. However, many applications require strong consistency guarantees, and coping with eventual consistency in the application layer requires significant development time [SV+13].

Although also designed on top of Cassandra, FogStore [GR18] is able to provide different consistency guarantees to clients, based on their context. In particular, the criterion is that close replicas need to be strongly consistent because they are supposed to be relevant, while eventual consistency suffices for the ones located at a certain distance. This is a finer-grained approach, but it still manages all the data in the same way, independent of the application or the kind of data. Also based on Cassandra is DataFog [GXR18], but it is focused on providing the best replica placement for performance at the edge and not on consistency.

Given that mF2C is a general-purpose platform, it requires a more customizable and flexible data management solution that provides easy and customizable control both on the data placement and the consistency for different kinds of data. This flexibility enables to use the same data management platform both for the application data, as well as for the data needed to manage the fog-to-cloud platform (devices, users, service catalogue, service execution, ...) in such a way that the required dynamicity, fault-tolerance and efficiency are provided.

2.3. Applications in different science areas, data centres, big data processing

Fog computing is being exploited in numerous domains where efficient processing of location sensitive sensor data and near real-time decision making is required. The development is to a certain extend driven by the continuous expansion of IoT devices, both in terms of volume and new functionalities. In recent months, we see application in new areas such as monitoring flash floods and valuable trees in remote forests, etc.

Rain in Saudi Arabia is infrequent but could be intense; irregular, torrential flash floods have caused significant damages to properties and disruption to society. [Guesmi:2018] proposed an early flash flood warning system based on analysing real-time hydrological data from wireless smart sensors. These sensors would be deployed in strategic locations within the monitored areas and measures attributes such as rainfall, water flow/level as well as capturing images. The sensor data are aggregated via distributed 'computation' nodes which run real-time prediction algorithms to evaluate the risk of flooding and to trigger flood alerts if the predefined threshold is reached.

The forest monitoring system [Pooja:2018] works on similar principle using smoke and different types of motion sensors to detect forest fires or illegal logging/illegitimate activities that threaten valuable trees such as teak and sandalwood in large swathe of remote forest reserves. The system streams data over ZigBee to distributed local micro-controllers for processing, the results are then forwarded to a receiver which aggregates and evaluates if an alert should be generated.

These two use cases share the common feature of localised processing of data close to source, a typical feature of the fog computing paradigm. Processing data close to source enables low latency and minimises network traffic as well as the complexity of the computation task. As IoT applications expand, the need for localised processing is driving the growth of edge data centres (also see Section 2.2.1 for a discussion of data management trends). [EDC:2018] reports that currently around 10% of enterprise-generated data is created and processed outside a traditional data centre or in the cloud, but by 2022, this figure is forecast to reach around 50%. The demand is giving rise to the growth of micro-modular and regional or local data centres, e.g. [SCHR:2018], [ECX:2018]. Unlike full scale data centres in the cloud, edge data centres will have few or even no IT staff on site to maintain and secure the infrastructure. It is all the more important for applications to incorporate a coherent policy towards securing data in transit and at rest to mitigate risks inherent in the F2C environment. mF2C - which incorporates secure by design - addresses this critical requirement through its unified framework for coordinating and managing distributed F2C components.

With regard to low-latency applications, Health IoT is the archetypal F2C use case that involves real or near real time processing of heterogeneous data streaming from different types of sensors monitoring patients' vital signs. The problem with connected healthcare devices is illustrated by infusion pumps [WIKI1]. These are devices connected to a patient's bloodstream into which they inject medicines, hormones, or nutrients. In the past, these machines were connected only to the patient, but these days they connect wirelessly to other systems, in order to receive updates on doses, alert medical personnel about deviations from the threshold state, or just send a message that they are functioning correctly. There are multiple points in such a system that can be compromised, e.g. by exploiting the sensors and their underlying communication network or any system and application vulnerabilities. Obviously, the implications of something going wrong with these devices can be quite serious: too much, or too little, medicine could cause harm to the patient. In addition, patient health records contain sensitive data that need to be protected in line with GDPR [EU2018] principles. In Sections 2.4 and 4.4, we will look at security trends in more details.

2.4. Security trends

The significant technological advances in the area of IoT/edge/fog/cloud computing, Big Data, and HPC have enabled the development of great innovations. However, these technologies enable collection, processing, and sharing growing amounts of personal data, thus posing a major risk to the privacy of people, and they highly increase the number of connected devices, which rapidly and vastly expands the attack surface of any organisation.

The evolution of intelligent things and the continuous adaptive security approach are creating a new cybersecurity trend [Pan17]. The adoption and integration of artificial intelligence and machine learning techniques into ICT infrastructures is bringing about the next generation of cybersecurity solutions. As cyber-attacks become more sophisticated, frequent, and complex, the self-learning tools and platforms that can be trained instead of just programmed, can accelerate and automate the counter fight. However, these technologies depend on vast amounts of continuously available and, most importantly, reliable data [Dua17].

Data brings many and huge benefits to the digital innovations, but it comes with great security risks to information systems, organisations, and individuals. Unauthorized manipulations of data can cause severe safety issues [Che17], data breaches can cause identity thefts and financial frauds [War17], and increasingly more popular ransomware attacks can easily put critical infrastructures on lockdown [BH16]. In today's data-driven world, data protection, individuals' privacy, and infrastructure security are more intertwined than ever before, and their integrated consideration for minimizing security incidents and data breaches is one of today's major cybersecurity trends [Fim17] [Pet18].

2.4.1. Artificial intelligence and machine learning

ML and AI approaches are becoming an emerging trend in the field of cybersecurity, offering alternative solutions to many interesting research challenges and practical use cases.

Analysis of encrypted traffic

As a means of tackling some aspects of ever-growing concerns for privacy, we have witnessed increased utilization of encrypted internet communication over the past decade. Protocols like IPsec and SSL/TLS allow entities to communicate privately, however providing no guarantees that a private communication is also a safe one, since malicious content may as well be carried in an encrypted form. Yet, analysis of encrypted traffic is beyond the capabilities of modern firewalls, resulting in malicious payload conveniently bypassing existing security controls [Yoon17]. Thus, there is a growing practical need for performing such analysis even when communication is encrypted. Fast and technically simple approaches for classification of encrypted traffic (such as port-based classification, termination proxies, and special implementations of SSL/TLS protocol), were either proven to be easily bypassed, raise significant privacy concerns, or just don't scale well. On the other hand, ML approaches cannot be bypassed, yet they retain the privacy of communication as they require no prior decryption.

ML-based approaches rely on characterization of unique statistical behaviour, which can be used to isolate individual encrypted flows. Various supervised ML algorithms like support vector machine and naive Bayes were applied for identification of application flows [Oka11]. Further, Bayesian network classifier was used in [Mai18-1] to identify anomalies within the isolated flows, which successfully demonstrated identification of DDoS and request/response poisoning attacks [Mai18-2] when the traffic was encrypted with SSL/TLS. At present, the research community is putting significant effort towards expanding the set of attacks, protocols and contexts in which ML-based approaches can be applied to detect malicious behaviour hidden in encrypted communication, in order to make the technology more useful for real-world deployments. On the other hand, Cisco recently released a family of switches that integrate their innovative product, Encrypted Traffic Analytics (backed by multi-layer supervised ML) [Cisco18], suggesting that the technology has already reached some level of maturity. Cisco's technology is said to accurately classify application flows and even promises threat detection in encrypted payloads, giving a good prognosis for further development of ML-based analyses of encrypted traffic for both research and commercial sector, and their usability.

Better detection and prevention of (known and unknown) cyber-attacks

What's more, ML/AI-based methods for detection of cybersecurity threats surpass the traditional signature-based methods that have come to the end of their useful application due to the sheer volume of potentially harmful security events, and increasing complexity of cyber-attacks enabled by computing paradigms like IoT. As a consequence, existing IDS/IPS solutions relying on these human-generated signatures of malicious network traffic and harmful software cannot keep up with such vast scale of attacks. In addition, they require significant system resources in order to be effective [Yu17]. ELIDe (extremely lightweight intrusion detection) which is a ML approach utilizing hash kernels and supervised linear classifier, emerged as an alternative for constrained appliances that yielded promising results on tactical systems as well as mobile devices [Yu15, Yu17]. Still, even in less resource-constrained environments, the need for techniques supporting real-time detection and prevention of cyber-attacks that go beyond signature-based approaches persists. Deep learning based on a convolutional neural network (CNN) has already been successfully applied to detection of new malicious source code variants in [Cui18], where features were extracted from grayscale images of the code. It even proved to achieve much better accuracy and speed than conventional malware detection mechanisms. Other AI approaches for malware detection and classification based on deep learning also report great accuracy improvements over traditional mechanisms in experimental setups [Meng17, Iera18]. At present, ML/AI techniques seem be one of the few alternatives promising detection of 0-day exploits. Yet, despite several cybersecurity vendors already incorporating ML-backed malware/attack detection suites into their solutions, some specific challenges persist. Kaspersky Lab [Kas18], themselves dealing with integration of ML into their antivirus products, emphasize the importance of training the model with the representative data (which is often very hard to achieve in practice, resulting in significant accuracy degradation over time if the model is not continuously retrained), and keeping multiple (ML and non-ML-based) detection methods in a multi-layered synergy [Kas17].

Supporting cyber-situational awareness and incident response

Another interesting research direction is application of ML and AI techniques to cyber-situational awareness and (potentially autonomous) incident response. For cyber-analysts in security operations centres (SOCs), establishing cyber situational awareness is becoming an increasingly challenging task. AI/ML-based solutions hold much promise as support mechanisms that could replace human input for initial cyber incident analysis tasks (triage process) due to their ability to classify cyber incident reports, find related ones, eliminate irrelevant information, etc. A tool based on deep auto encoder neural network, that demonstrates exactly these features, was developed in [Graf18].

2.4.2. Data protection

The massive collection and continuous availability of data constitute an enormous chance for the data-driven world. However, data is very often not handled in a secure, privacy-friendly, or transparent way, which causes distrust in data collection (i.e. privacy concerns) [EDE18] [OECD17] and insecure data processing (e.g., data breaches, identity thefts) [VER18] [Coo18].

Individuals are getting very cautious about sharing their personal data and very sensitive about how their personal data is used. Additionally, recent studies [BDVA17] on privacy-protection mechanisms show that simple approaches like masking or removal of unique identifiers in datasets (e.g., names, social security numbers, etc.), are insufficient to protect privacy. Increasingly more privacy-aware individuals, the lack of efficient privacy-protection mechanisms, and the new data protection and privacy regulations (General Data Protection Regulation [EU16] and ePrivacy Regulation [EC17], which, if adopted, will replace the ePrivacy Directive [EU02]) are shifting the scientific focus towards strong cryptographic privacy-enhancing and anonymity-enabling approaches. These approaches enable individuals to stay anonymous in front of their digital service providers and thus remain in full control over their privacy without having to place trust in their service providers or third parties to not misuse or share their personal information with others without their knowledge and consent. One of the popular examples of such technology are blockchain-based cryptocurrencies, such as Bitcoin [BIT], Mixcoin [BN+14], Zerocash [BC+14] or Monero [MON], which provide means for (pseudo-) anonymity in financial transactions. Here, financial transactions are made, with cryptography [Nak08], in a peer-to-peer fashion, without any centralised entity involved, and recorded in a public ledger to ensure integrity. Since funds are not tied to individuals but rather associated with wallet addresses, such technology provides (pseudo-) anonymity. Another example of privacy enhancing technologies (PETs) are attribute-based credentials (ABCs) [CK+12] and group signatures [CL02] (or related approaches such as direct anonymous attestation [BCC04]), which provide powerful means for anonymous or pseudonymous authentication. Such schemes involve users, issuers, and verifiers. A typical flow is that users obtain credentials from an issuer and can then anonymously demonstrate the possession of a valid credential to verifiers. Either such a credential simply proves membership to some user group (e.g., can prove that the user has the right to access a building) or in case of ABCs allows to selectively reveal information about attributes encoded into the credential (e.g., that a user is at least 18 years in age) without revealing anything else. Conceptually, these technologies date back to the 1980's [Cha85, CH91] and a lot of theoretical progress has been made since. While early approaches relied on the factoring setting [CL01, CL02, BCC04], the most efficient approaches today rely on the use of elliptic-curves and bilinear-maps [CL04, HS14, RVH17]. There is also strong interest in primitives that are secure against powerful quantum computers (e.g., schemes based on lattices [CNR12, BCN17, BK17]).

Even though a large amount of theoretical work has been done on cryptographic PETs, there is a lack of implementations available for the industry to fully adopt the technology. There are several EU research projects such as PRIME [PRI], PrimeLife [PL], ABC4Trust [ABC], MATTHEW [MAT], and CREDENTIAL [CRE], which offer prototypes for privacy enhancing technologies. Additionally, some prototypes have also been developed by the industry, such as Idemix from IBM [IBM02] and U-Prove from Microsoft [Mic12]. However, none of the implementations is complete, actively maintained, easy to use or published in open source, which puts a lot of interest in a complete, efficient, and scalable software library for cryptographic PETs.

The pervasive nature of today's technology implies high demand not only for solutions that protect the data but also for solutions that enable computations over protected data. To this end, the cryptographic community has proposed novel cryptographic concepts like secure multi-party computation, functional encryption, and homomorphic encryption (from partially, over somewhat, to fully homomorphic encryption). Secure multi-party computation (SMC) [GL05, BC+09, Kre17] ensures confidentiality in settings where two or more parties jointly and securely evaluate arbitrary functions over their inputs while keeping those inputs private. Functional encryption (FE) [BSW11] is an emerging paradigm that enables fine-grained access to encrypted data, overcoming the all-or-nothing law of the traditional public key encryption, where the encrypted data is either fully available or fully hidden and thus all entities in a system have to be in either of the two states as well. In the context of the FE approach, decryption keys are associated with a function *f* and when applied to encryptions of a message *m*, the decryption yields *f(m).* Moreover, in an FE system, we can issue decryption keys for different users, so that the users obtain a different view of the message plaintext. This means that we can cryptographically regulate who has access to what data or to what function of the data, enabling partial or full anonymity of data subjects. When/if security breaches occur, they are isolated and cannot affect the security or privacy of a whole system. In the last two decades, different flavours of FE emerged; from predicate encryption [KSW08], attribute-based encryption [GP+06], and identity-based encryption [BF01], all providing a powerful tool for minimizing information leakage and maximising security and privacy. Finally, homomorphic encryption (HE) [Gen09] refers to the encryption technique that allows for computations to be done over encrypted data. The results of such computations are encrypted and when decrypted, they match the result of the operations as if they had been performed on the plaintext. Depending on the type of operations (e.g., addition, multiplication), their combination (e.g., addition and multiplication, multiplication and division), and the number of repetitions (e.g., one operation, limited number of operations) allowed on cipher text, HE schemes are categorised into three groups. Partially homomorphic encryption [LCM16] allows only one type of operation unlimited amount of times. Somewhat homomorphic encryption [DPZ12] schemes enable some types of operations being performed unlimited amount of times. Fully homomorphic encryption [Gen09] allows different combinations of different operations unlimited amount of times. Although these schemes provide a powerful solution for ensuring data protection, they are (still) too costly in terms of computation to be practical in real-world. Therefore, the scientific community is still working on follow-up improvements.

2.5. Resource management and QoS

The incremental development from Fog and Edge computing from our last version [D2.1] of this deliverable, has brought a more mature research context to the area of Edge and Fog QoS and Resource management. Examples of these advances are available in diverse areas and approaches.

Mahmud in [MAHMUD] provides a very complete overview of most common approaches for Resource and Service provisioning metrics and Service Level objectives. It recognises that main criteria used nowadays for resource management in Fog computing are time, data, cost and energy management metrics. For Service Level Objectives, it gathers on-going interest in Latency, Cost, computation, data and application management, as well as, power handling. For future directions in this area, it recognises Context-aware Resource and Service Provisioning, including latency optimisation, as it is the approach in mF2C, as a promising area for further development.

Brogi in [BROGI] analyses existing challenges in the deployment of composite applications in heterogeneous large Fog environments. It presents FogTorch which models Cloud+Fog+IoT scenario in order to select the best deployment in this composite application taking into account QoS offered by Fog environments and composite application execution requirements. While approach and starting point for this work and mF2C existing developments can be comparable; a significant difference is that mF2C considers runtime deployment decision while FogTorch considers this decision to be taken at application design time.

Also, Interestingly for mF2C developments and taking a similar approach to its multi-layered Fog to Cloud architecture Tocze in [TOCZE] which acknowledges the need of further research in such environments in areas of resource estimation, discovery, sharing and allocation and related migration, as areas that can have significant influence in timeless and QoS Fog to Cloud computing. Of this list mF2C is in the position of providing relevant advances which specifically focus on resource discovery, sharing and allocation.

2.6. Convergence of AI and computing

In the previous version of this deliverable, [D2.1], we identified the main trends when considering the convergence of AI and computing. On the one hand, we found that many companies like Google, Amazon, Microsoft and IBM had incorporated AI in their platform-as-a-service or software-as-a-service solutions. On the other hand, and also in the scope of the mF2C project, we found a convergence or collaboration in the opposite way. Parallel computing, which is one of the pillars of the mF2C architecture, can help the development of more performance efficient ML or genetic algorithms. In this sense, AI algorithms can benefit from the inherent parallelism in the proposed mF2C architecture, allowing the distributed implementation of ML algorithms, and then improving the performance when computing data. Specifically, we reviewed some works in the literature that proposed the use of ML, fuzzy neural networks, Markov models or genetics algorithms to collect and process data coming from sensors in fog nodes and/or cloud. The conclusion was that in the mF2C system we need the best possible combination of AI techniques that can effectively be adopted in different parts of the system to take an intelligent decision based on the processed data near the end users, providing low latency as well as enhanced security, as required by critical medical and many commercial applications.

However, the revision made in the previous deliverable only considered the use of AI when processing the users' application data, but, why not using also AI in the control and management processes as part of the mF2C system? Før example, when matching the services to be executed on the available resources, AI becomes an extraordinary opportunity to make the best possible selection. We find in the recent literature proposals dealing with the use of different AI techniques in the management of Cloud computing. Examples of these proposals can be found in [GJK18], [DMD17], [SV17] and [GSR17]. In [GJK18] authors propose a system to securely monitoring Cloud computing to prevent unauthorized tasks injection and modification, as well as to optimize process scheduling and maximization of resource usage; all this based on intelligent agents equipped with Artificial Neural Network (ANN). On the other hand, in [DMD17] an Ant Colony Optimization (ACO) algorithm is proposed to optimize the load balancing between Virtual Machines (VMs) in the entire system. Also, in [SV17] authors propose the use of Fuzzy techniques, specifically the kernel fuzzy c-means clustering algorithm, for clustering resources before the allocation of them by means of optimization techniques. Finally, in [GSR17], a learning automata is used to achieve a trade-off between power consumption reduction from one side, and SLA violation percentage from the other side when optimizing the VMs placement.

When extending the scope to fog and edge computing, even less contributions can be found that use AI for optimizing the allocation of resources or other management processes. They are mainly cloud computing management techniques, extended to consider some devices at the edge. The RECAP project (https://www.recap-h2020.eu/) and its paper [OBC17], in a preliminary approach, proposes as one of its pillars the intelligent automation of management processes; for example, when placing applications in the different resources, considering both fog and edge; as well as when managing and placing the Virtual Network Functions (VNFs). Also considering cloud and edge resources, the work in

[AA18] proposes a genetic algorithm for optimal VM placement for a cloud federation architecture, also including edge resources. The project is aware of current work being done in AWS GreenGrass or Azure IoT Edge which also focus on providing Machine Learning capabilities at the edge. Other interesting developments worth mentioning here is Google Federated Learning for Edge / Cloud synchronized model training [GOO17].

Finally, and framed in the development of this mF2C project, in a previous deliverable D5.1 we proposed the use of reinforcement learning in one of the blocks of the mF2C architecture, the QoS provider. The decision whether a certain agent (resource) can or cannot be used for a certain service instance is based on the number of SLA violations occurred in previous executions of that specific service. In order to determine if the suggested agents by the service instance should be used, the QoS provider uses the number of service executions and the number of SLA violations, to calculate a ratio that is used as the input for the Deep Q-learning (DQL) algorithm. Based on this approach, and from a research-oriented approach, the paper [DCB18] proposes a whole service management system for a Fog-to-Cloud architecture, where the application of ML algorithms in different components of this service management (Service Classifier, Resource Provider, Quality of Service, and Analytics) unit is analysed. Also, in the context of the mF2C project, the work in [SGM18] proposes the use of ML techniques to build an adaptive cost model for managing the resources. With this adaptive cost model, the system will be able to allocate more sophisticated and optimized solutions for satisfying the requested service.

2.7. Key Takeaways

We summarize the key areas of focus in scientific trends significant for the project:

- The challenges faced designing resource management solutions for fog computing systems continue to be heterogeneity, dynamicity, geo-distribution, and multiple owners of the devices taking part in that fog system
- A number of decentralized algorithms allowing devices to coordinate their periodic offloading decisions to achieve efficient management of communication and computing resources in a fog computing system have appeared
- The new platform "Fogernetes" (based on the existing container management tool Kubernetes) allows matching the requirements of fog application components with device capabilities through a labelling system
- Core layers in the hierarchy can be found in [FMG18] where the concept of "communities" is used to represent a group of fog nodes.
- Game theory approaches have gained popularity in the formal modelling of resource management problems in fog computing systems, e.g., vehicular fog computing environment, e.g., Multi-Slot Computational Offloading Game (MSCOG) obtains decentralized offloading decisions, minimizing the response times and the corresponding energy consumption.
- Trends in the HPC area currently focused on solving latency issues and challenges that arise in fog/edge infrastructures when accessing storage. Systems encounter difficulties to scale and are not well suited for the local activity expected in a fog context.
- Preliminary solutions towards holistic views of data management from the edge to the datacentre are being proposed in in flagship HPC conferences, e.g., USENIX HotCloud and HotEdge
- A number of projects are researching how to fragment applications in order to offload the parts of the computation to the resources, the scheduling model and the management of parallelism: CloneCloud, Cuckoo, ThinkAir, MobileFog
- ML and AI approaches are becoming the emerging trend in the field of cybersecurity, offering alternative solutions to research challenges and practical use cases, helping to identify

anomalies within the isolated flows, which successfully demonstrated identification of DDoS and request/response poisoning attacks

- Software cannot keep up with the scale of attacks - existing IDS/IPS solutions rely on human-generated signatures of malicious network traffic). ML/AI techniques seem be one of the few alternatives promising detection of 0-day exploits
- There are several EU research projects offering prototypes for privacy enhancing technologies: PRIME, PrimeLife, ABC4Trust, MATTHEW, and CREDENTIAL, which (PET)
- AI algorithms can benefit from the inherent parallelism in the proposed mF2C architecture, allowing the distributed implementation of ML algorithms, and then improving the performance when computing data
- Building adaptive cost model for managing the resources using ML techniques will support allocating more sophisticated and optimized solutions for satisfying SLA’s

3. Technology Trends

Very much like the former version of this deliverable [D21], in this section we address several technology trends in different areas which are relevant to mF2C, and we try to highlight any new ones that might have been getting more attention lately. It is important to note that since the submission of D2.1 [D21], which was in March 2017, many of the technology trends remain the same, thus for further details about a recurrent trend one should refer to the previous deliverable [D21].

3.1. Tools, platforms, IoT

In the previous version of this deliverable [D21], we reviewed the technological trends regarding the management of Cloud, Fog and IoT devices. Although those revised trends are still valid, in this section we are going to update the latest trends in Cloud and IoT management.

3.1.1.Cloud management tools

Apache Mesos

Apache Mesos (http://mesos.apache.org/) is an open-source project to manage computer clusters. Mesos uses Linux Cgroups to provide isolation for CPU, memory, I/O and file system. Mesos is comparable to Google's Borg scheduler, which is a highly secretive platform used internally to manage and distribute Google's services. Interestingly, Mesos deals with resource management and scheduling and is often described as a “distributed systems kernel” that allows manage thousands of servers using containers to host applications. It also provides a set of daemons, which expose resources to a centralized scheduler. So that, allows tasks to be distributed across nodes, and thus load balancing takes place on different resources in the cloud or traditional systems. Mesos can be easily fitted with large, distributed databases such as Hadoop and Cassandra.

Cloudability - (https://www.cloudability.com/platform/)

It is one of the popular cloud management platforms. Cloudability helps the organizations to properly monitor the cloud resources consumption. By keep tracking the resource usage related information, it provides various reports to analyse the cost related issues and most important, it helps the organizations to optimize the expenses. It supports multiple public, private, and hybrid cloud service providers.

3.1.2. IoT Management tools

SensorThings API - (https://github.com/opengeospatial/sensorthings)

It is an open and unified framework, which allows to interconnect the sensing devices, data, and applications over the Web by addressing the syntactic interoperability and semantic interoperability of the Internet of Things. It can easily work with the existing IoT networking protocols such CoAP, MQTT, HTTP, 6LowPAN. Interestingly, it is non-proprietary, platform-independent, and perpetual royalty-free API, which helps to significantly simplify and accelerate the development of IoT applications. By using this API, application developers can connect to various IoT devices and create innovative applications without worrying the daunting heterogeneous protocols of the different IoT devices. It can also be embedded within various IoT hardware and software platforms, so that the various IoT devices can effortlessly connect with the OGC standard-compliant servers around the world. The Open Geospatial Consortium (OGC) SensorThings API defines standardized interfaces for sensors in the Web of Things (WoT) and Internet of Things (IoT), two terms that are frequently used interchangeably. Most importantly, these standardized interfaces will permit the proliferation of new high value services with lower overhead of development and wider reach and will also lower the cost for sensor and gateway providers.

ThingSpeak - (https://thingspeak.com/)

It is an open IoT platform which comes with MATLAB analytics facilities. This IoT analytics platform allows to aggregate, visualize and analyse live data streams in the cloud. ThingSpeak provides instant visualizations of data posted by your devices to ThingSpeak. By executing the MATLAB code in ThingSpeak, developers can be able to perform online analysis and processing of the data as it comes in. ThingSpeak is often used for prototyping and proof of concept IoT systems that require analytics. Most importantly, it is easy to configure and also it can easily work with popular IoT protocol.

Thingsboard.io - (https://thingsboard.io/)

ThingsBoard is an open-source IoT platform for data collection, processing, visualization, and device management. It enables device connectivity via industry standard IoT protocols - MQTT, CoAP and HTTP and supports both cloud and on-premises deployments. ThingsBoard assures scalability, fault-tolerance and performance of the system, so that the captured data can’t be loose. It controls the IoT entities in secure way using rich server-side APIs. It is also defining the relations between devices, assets, customers or any other entities. Also, it is collecting and storing telemetry data in scalable and fault-tolerant way. It also defines data processing rule chains by transforming and normalizing the device data. Also, it integrates devices by connecting to legacy and third-party systems with the help of existing protocols.

3.2. Technology trends coming from HPC

3.2.1. Data management trends

The technology trends regarding non-volatile memories, active storage, parallel file systems and NoSQL databases were already reported in D2.1 [D21] and are still valid. In this second version of the deliverable, we will shift the focus to data management solutions for the edge to cloud continuum.

As explained in section 2.2.1, this research area that combines edge computing and HPC is now emerging, with many proposals and proof-of-concept evaluations, but not many implemented solutions.

An exception is the Cisco Kinetic Edge & Fog Processing Module [Kinetic], which facilitates the deployment of data processes to edge and fog. This is a commercial solution that enables IoT applications for advanced monitoring and diagnostics, focused on industrial environments. Its Data Control Module filters, aggregates, and compresses time series data at the edge or in the fog, or in the cloud. It provides an enterprise IoT solution for operations and decision making. A more general commercial datastore suitable for IoT is Redis Enterprise [RedisEnt], which supports different data types and consistency levels.

A non-commercial datastore is dataClay [dataClay, MQ+17], which is also applicable to more general settings, as required in this project. It is not limited to a particular kind of data and provides the flexibility to manage replicas and the synchronization between them according to the needs of each specific application. This allows to decide, for each kind of data, how many replicas need to exist, where they need to be placed to ensure low latency and fault tolerance, and which is the required synchronization so that the requirements of the application are satisfied. This high degree of flexibility enables an efficient use of resources, avoiding unnecessary communications and delays, paying the price of replication and (degree of) synchronization only when required.

3.2.2.Programming model's trends

In D2.1 we have addressed one of the requirements for fog computing, the need of a programming framework that allow users to program applications for heterogeneous devices. In this deliverable we extend the focus to programming solutions that cover issues related to the fog-to-cloud paradigm including real time processing, latency and transparent management of a decentralized, heterogeneous and dynamic set of resources. The aim of the project is to provide a programming framework to develop applications that involve the use of traditional cloud systems, smart end-user devices, and IoT sensors.

Looking at the commercial offerings, the big companies in cloud are extending their services to include fog tools. Azure [Azure:2018] provides a platform to support the automatic scaling of Azure Functions on IoT devices. Amazon Greengrass [AWS:2018] allows to run AWS Lambda functions on connected devices, using the cloud for management, analytics, and durable storage. Both solutions provide a serverless architecture where users are only required to write the code without caring of the resources needed for the computation. Google has launched Android Things [AndroidThings:2018] to develop apps for IoT devices with existing Android development tools, APIs, and resources along with new APIs that provide low level I/O and libraries for common components like temperature sensors, display controllers, and more, available on certified hardware.

The streaming feature is one of the most common requirements, taking into account that some application areas may require the possibility of accepting streamed input data (from sensors or other sources of dynamic data) and streamed output data (visualization, monitoring, etc). Apache Spark [APA15] is one of the most widely adopted programming framework also due to the availability of libraries that are suitable for fog, as for example Spark Streaming. COMPSs is being extended to support streams as data inputs and outputs. The previously referenced CLASS project software architecture will integrate the Map/Reduce and the task-based programming models (implemented with COMPSs) into a unified language to enhance portability.

3.3. Cloud Orchestration Platforms, Virtualization, Containers

In D2.1, section 3.3, we described the technology trends related to the orchestration of containers in the cloud and fog area (Docker swarm, kubernetes, etc.), as well as the trends related to the cloud management and cloud orchestration tools (OpenStack, Microsoft Azure, Slipstream, etc.). We put special emphasis on containers and their orchestration as this is one of the technologies this project relies on. In fact, at the end of IT-1 all the services running on mF2C are based on Docker containers.

This section briefly updates all this information. Although most of the trends and technologies described in this previous deliverable are still valid, the confirmation of kubernetes as the leading engine for containers orchestration has changed the overall picture. This was partly due to the use of Docker and the popularity of this technology.

During the last year Kubernetes has established himself as the de facto container orchestration and management engine, with companies like VMware (VMware Kubernetes Engine [VMW18]), Pivotal, Openshift, Docker [DOC1], Microsoft (AKS) [MIC17] and Amazon (Amazon Elastic Container Service for Kubernetes) [AMA17], among many others, including it in their offerings. Although Apache Mesos (Marathon) and Docker Swarm are used by many companies, they are being relegated to a second plane. Surveys [CNC1] like the ones from CNCF [CNC2], reflect this leading in the containers

management world. In fact, according to reports like the one from sysdig.com [SIS2017], kubernetes is the container orchestrator used by most of the people that work with Docker containers.
Finally, Kubernetes has recently become the first CNCF project to graduate [CNC2018], hence being officially recognized as fully mature open source project.

3.4. Role of standards in technologies

Deliverable D2.1 introduced a number of the key standards organisations and initiatives potentially relevant to mF2C. Subsequent developments of particularly relevant standards have centred on the various ISO/IEC JTC1 sub-committees, the OpenFog Consortium, and its collaboration with IEEE.

ISO/IEC JTC1 SC38 Cloud Computing and Distributed Platforms

Although focusing on centralised cloud systems in the past, SC38 is now pursuing a broader work programme with explicit references to Edge Computing in particular. Work Group three dedicated to Cloud Computing Fundamentals has completed development of several relevant standards and is working on a number of relevant technical reports.
Recently published standards include the following:

- ISO / IEC 19086-3:2017 - Service Level Agreement Framework - Part 3: Core conformance requirements
- ISO / IEC 19941:2017 - Interoperability and Portability
- SO / IEC 19944:2017 - Cloud services and devices: Data flow, Data categories and data use

The following standards and technical reports are in development:

- ISO / IEC AWI TS 23167 - Common Technologies and Techniques
- ISO / IEC PDTR 23186 - Framework of trust for processing multi-sourced data
- ISO / IEC NP TR 23187 - Interacting with Cloud Service Partners (CSNs)
- ISO / IEC NP TR 23188 - Edge Computing Landscape

It should be noted that standards authored by ISO / IEC JTC1 SC38 are generally high-level and descriptive in nature. Technical specifications of APIs are typically developed by Industry Groups. Various insights and perspectives from mF2C are being fed into the discussions and national body contributions via the Irish mirror committee to SC38.

ISO/IEC JTC1 WG10 Internet of Things

This working group has been disbanded and efforts transferred into the new SC41.

ISO/IEC JTC1 SC41 Internet of Things and Related Technologies

This sub-committee has recently been inaugurated and to date has published two documents of relevance to mF2C:

- ISO/IEC TR 22417:2017 - Internet of Things Use Cases
- ISO/IEC 30141:2018 - Reference Architecture

It is working on a suite of relevant standards including:

- ISO/IEC CD 20924 - Definition and Vocabulary
- ISO/IEC AWI 21823-1 Interoperability for Internet of Things - Part 1: Framework
- ISO/IEC NP 30147 - Methodology for trustworthiness of IoT system/service

Agreement has also been reached between SC41 and SC38 for standards related to Edge computing to be developed in cooperation with each other.

Alliance for Internet of Things Innovation (AIOTI)

AIOTI working groups are focused on research and innovation, policy issues and proposed standards, as well as horizontal, cross-disciplinary activities focused on hot topics in the field. They have published 12 reports covering IoT policy and standards issues. The organisation has provided detailed recommendations for future collaborations in the *Internet of Things Focus Area* of the 2016-2017 Horizon 2020 programme. Engineering is currently a member of this organisation and will continue to provide updates to the project on direction and opportunities to influence.

OpenFog Consortium

The OpenFog Consortium is continuing its pursuit of providing a cross-industry perspective on Fog computing. It has recently formally documented Security Approaches and Requirements, to which mF2C has specifically responded. It has also updated its Glossary of Fog Computing Terms.

In August 2018, the OpenFog Consortium reached a milestone with their collaboration with the IEEE. The IEEE published IEEE 1394, the IEEE standard for Adoption of OpenFog Reference Architecture for Fog Computing. The new standard “is intended to address the need for an end-to-end, interoperable solution that is positioned along the things-to-cloud continuum. The new standard supports multiple industry verticals and application domains and is designed to enable services and applications to be distributed closer to the data-producing sources and/or the information-consuming users”.

3.5. Technology trends in edge computing

In deliverable 2.1 [D21], we have reviewed the (at the time) existing trends in edge computing. We’ve also introduced the reader to “What is Edge Computing” and provided a few promising reference solutions within the area.

Edge computing continues to evolve by pushing computing power closer to IoT sensors. This reduces the latency created by sending data to the cloud, which by itself ends up increasing the IoT data processing efficiency. Besides the reduction of the reliance on networks, edge computing has been growing stronger when it comes to address the IoT data deluge, and these are some of the motivation factors:

- Ability to aggregate (reduce) data at its source;
- Real-time decision making;
- Data anonymization and privacy protection;
- increased autonomy.

Building on top of the topics addressed in section 3.5 of D2.1, we’ll use this document to provide updates and highlight new trends in edge computing that have arisen in the meantime.

3.5.1. Reference solutions

In D2.1, several reference solutions were presented and are still valid to this day. In this document, we will extend those references by adding a few more examples of promising solutions in the market.

NuvlaBox Nano

The use of the NuvlaBox [NuvlaBox] in smart cities and other edge scenarios has already been described in past deliverables. The novelty now being implemented by SixSq is an extension of this generic edge device, the NuvlaBox Nano, which is a single board computer equipped with Docker. This new cloud-in-a-box form factor uses Docker instead of the traditional virtualisation layer and includes the capability-based leader election functionality. This allows clusters of devices to self-form Docker swarms over Wi-Fi. SixSq's original SlipStream-based software ecosystem will be upgraded, adding device management capabilities beyond Docker, like the discovery and bridging of IoT sensors to multi-container applications being deployed remotely by Nuvla, via the Docker Engine API. The portability, modularity and homogeneity offered by Docker will allow SixSq to perform dynamic and horizontal scalability at the edge, by merging cloud-based devices with Docker-powered single-board computers, while keeping a homogeneous software stack, from the cloud to the edge.

Akraino Edge Stack

The Akraino Edge Stack [Akraino] is a Linux Foundation project which will develop a fully integrated edge infrastructure solution, through an open source software stack that improves the state of edge cloud infrastructure for carrier, provider, and IoT networks. Supported by multiple companies (including the mF2C partner, INTEL), this open source software stack will make use of the work already done by AT&T and INTEL to address critical infrastructure requirements. It will enable high performance, reduce latency, improve availability, lower operational overhead, provide scalability, address security needs, and improve fault management.

The Akraino Edge Stack community aimed for a first release of its code during the second quarter of 2018. More details can be found at Akraino's Wiki page [AkrainoWiki].

3.5.2.Containers in the edge

The adoption of containers at the edge is not new and as referred in the previous version of this report, it is already being explored by several mF2C partners. Within the project itself, containers have been chosen as the backbone for the deployment of the mF2C agents and also for the execution of services in the fog, for IT-1.

As mF2C now moves towards the second phase of the project, IT-2, the use of containers seems even more natural, as these have been gradually a subject of discussion within the container-based communities, when it comes to their usability and usefulness at the edge. During 2018, the most prestigious container industry conference, DockerCon, has been requesting and strongly addressing the topic of the use of containers at the edge. During the DockerCon San Francisco 2018, back in June, several edge use cases where presented, where Docker was being used to package software for distribution at the edge, easing the management of applications and reducing the gap between cloud and edge.

More on this topic can be found online at Docker's new solution for edge, Docker Edge [DockerEdge].

3.5.3.Edge and IoT

As generally predicted back in 2016 by a study from IDC [IDCFS], IoT platforms are still on the rise, as recently supported by Gartner's 2018 Hype Curve [GHC18], where it is predicted that IoT will reach its plateau of productivity within the next 5 to 10 years.

As IoT grows, so does edge (and consequently fog) computing. With fog computing, instead of having dedicated devices at the edge which are tightly coupled with the IoT sensors, the idea now is to make use of any capable generic devices (like the NuvlaBox [NuvlaBox]), including IoT gateways and even sensors, to offload the intelligence down to the local area network level of network architecture.

Other trends in this field include the use of the extended sensor capabilities in IoT to bypass the edge by having the sensors themselves communicating directly with the cloud when necessary.

3.6. Key Takeaways

We summarize the key areas of focus in technology trends significant for the project:

- A number of Cloud Management tools (i.e., Mesos, Cloudability) provide solutions for resource management and scheduling out of the box
- IOT Management solutions (i.e., SensorThings, ThingSpeak, and ThingsBoard.io) interconnect sensing devices, data, and applications over the Web, therefore addressing interoperability
- A new research area combining edge computing and HPC appears to be emerging. While several proposals and proof-of-concept evaluations, there appears to be not many implemented solutions.
- There appears to be a need for a programming framework allowing users program applications for heterogeneous devices, due to issues related to the fog-to-cloud paradigm. These include real time processing, latency and transparent management of a decentralized, heterogeneous and dynamic set of resources
- The OpenFog Consortium reached a milestone with the IEEE publishing their standard IEEE1394, the Adoption of the OpenFog Reference Architecture for Fog.
- Kubernetes still remains the leading engine for container orchestration. This is due to the popularity of Docker and supporting tools, i.e., VMware Kubernetes Engine, Pivotal, Openshift, Docker, Microsoft (AKS) and Amazon's AECSKT
- A number of new reference architectures have appeared on the market that deliver edge computing solutions, including both open source and proprietary products and services. These include Nuvla Box Nano and Akraino Edge Stack
- Several edge use cases were demonstrated at DockerCon San Francisco 2018, where Docker was used to package software for distribution at the edge, easing the management of applications and reducing the gap between cloud and edge.

4. Business trends

This section covers the latest business trends related to the topics covered by mF2C, identifying potential market gaps as well as expectations and needs. The content presented here will be later used to update the business goals of the project that the mF2C architecture must fulfil. On the other hand, the analysis of latest business trends will be used to identify the possible ways for developing the business models for the project.

4.1. Cloud, fog and edge computing

According to Research and Markets [RES1], the global cloud market is expected to reach $1.250 billion by 2025, growing at a CAGR of 27.5%. This is a significant growth in a well-established market.

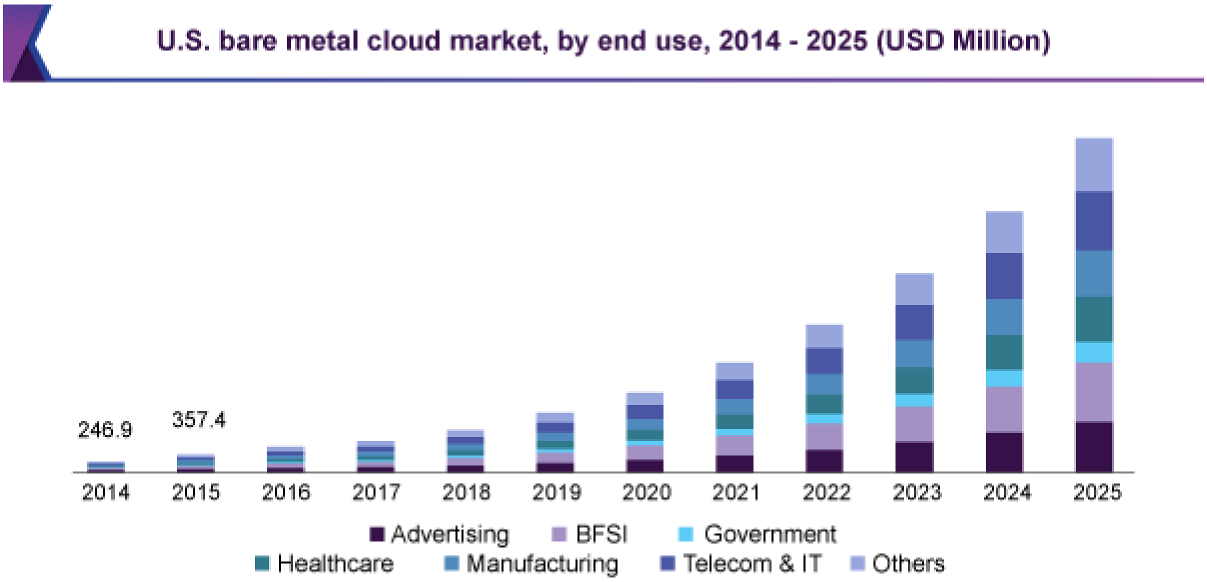

Figure 1 US Cloud market forecast 2014-2025 [GVR1]

The figure above [GVR1] shows the evolution of cloud in traditional market segments defeating the reluctance to its further adoption. At the same time, its adoption has had a significant growth in sectors like Manufacturing, traditionally more related to operational technologies rather than information technologies, what is also supporting the fog market growth. The emergence of big data and the increased need for data analytics, storage, decentralization and easy-to-install features are the key drivers that are fostering the cloud adoption in these monolithic sectors, while data loss and the unstoppable need for high computational power is pushing companies to the cloud [HTF1].

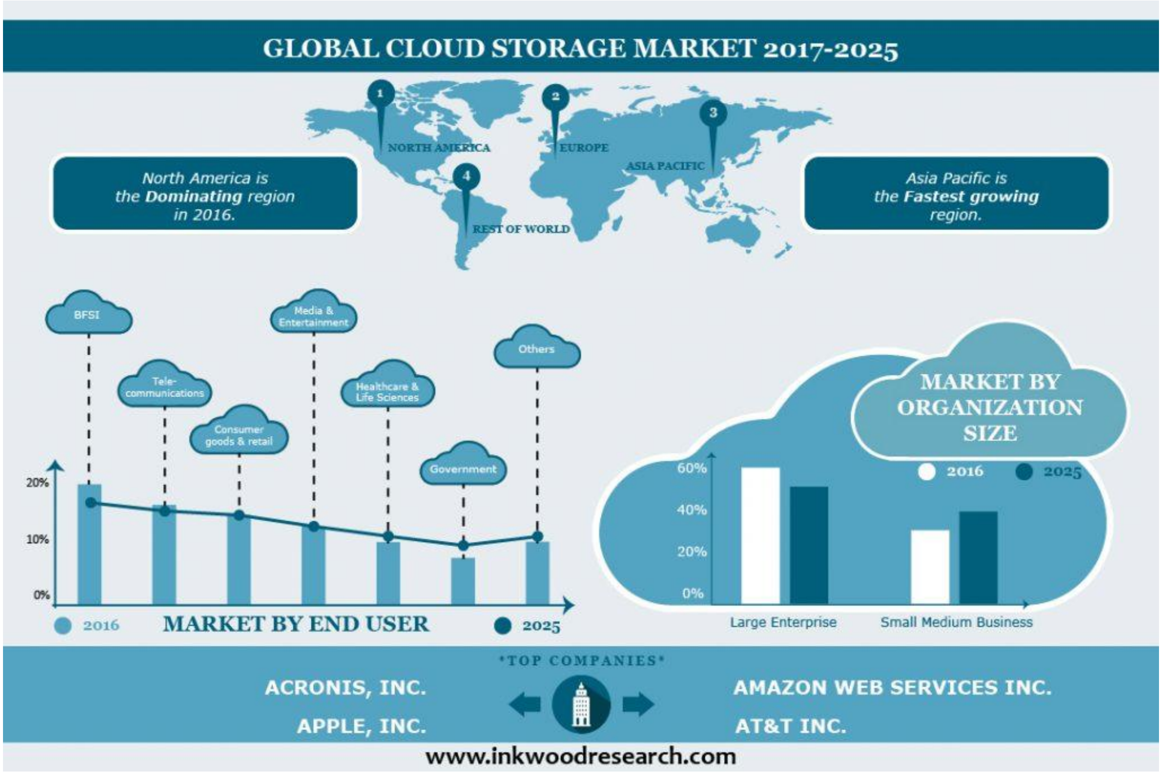

Figure 2 Global Cloud storage market 2017-2025 [INK1]

According to an Inkwood Research study [INK1] (see Figure 2), 90% of the existent businesses is adopting cloud-based solutions independently of their application domain. The same study shows that while large enterprises have been doing it in the last years, SMEs are expected to significantly do it in the coming ones. This is driving cloud to a new completely fragmented market, creating new niches for a wide range of solutions (see Figure 3).

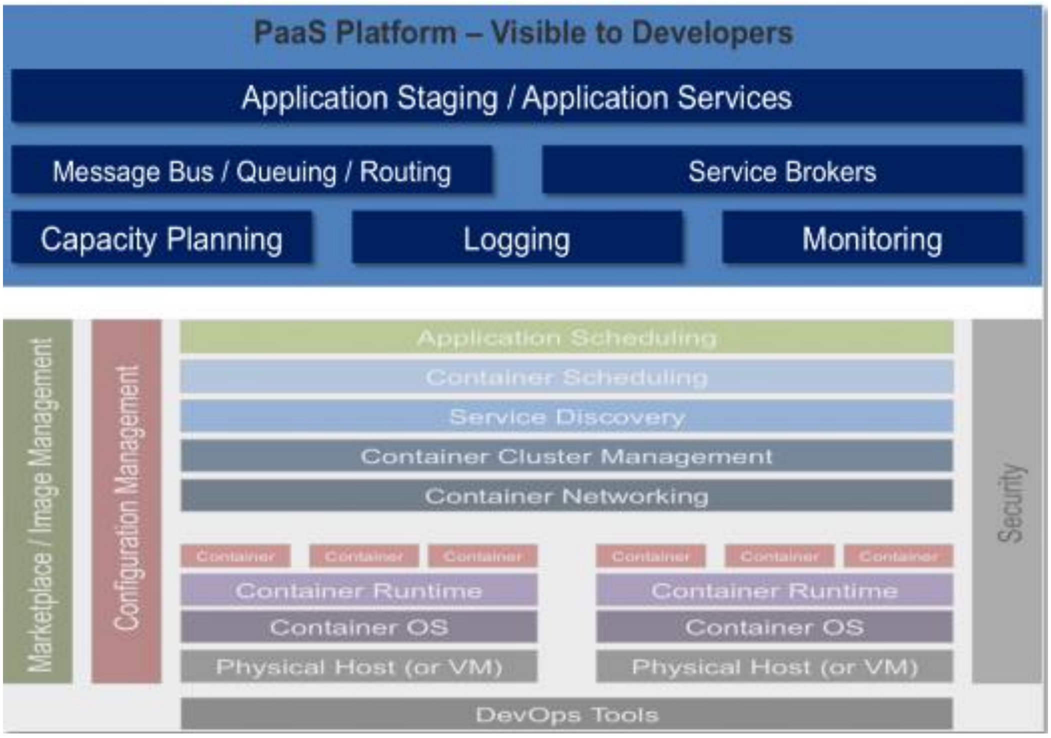

Figure 3 Emerging Market Frameworks and Definitions [WBR1]

This is reinforced by the analysis performed by Gartner [GR1] which considers that, although IaaS market is experiencing the fastest growth, SaaS market remains as the largest segment. However, as said before, the need of data analytics is making the database platform as a service (dbPaaS), as the fastest growing segment, being expected to reach $10 billion by 2021.

However, in this continuously growing market, 'cloud' will disappear [INFR1]. This doesn't mean that cloud solutions will disappear, it means in fact that cloud will be of so commonly use by 2025 that the term will be substituted by specific offering terms.

On the other side of the chain, global edge computing market is also experiencing a significant growth, as shown in Figure 4, and it is expected to reach $3.24 billion by 2025 at a CAGR of 41% [GVR2]. As it happens within the cloud market, the increased usage of IoT devices is pushing for the adoption of new and advanced technologies that can deal with the big amount of vast data generated every day avoiding delays in its processing. And again, as it is happening within the cloud market, edge analytics market is the segment experimenting the most significant growth at a CAGR of 22.2% reaching $13.44 billion by 2025 according to Research and Markets [RMR1].

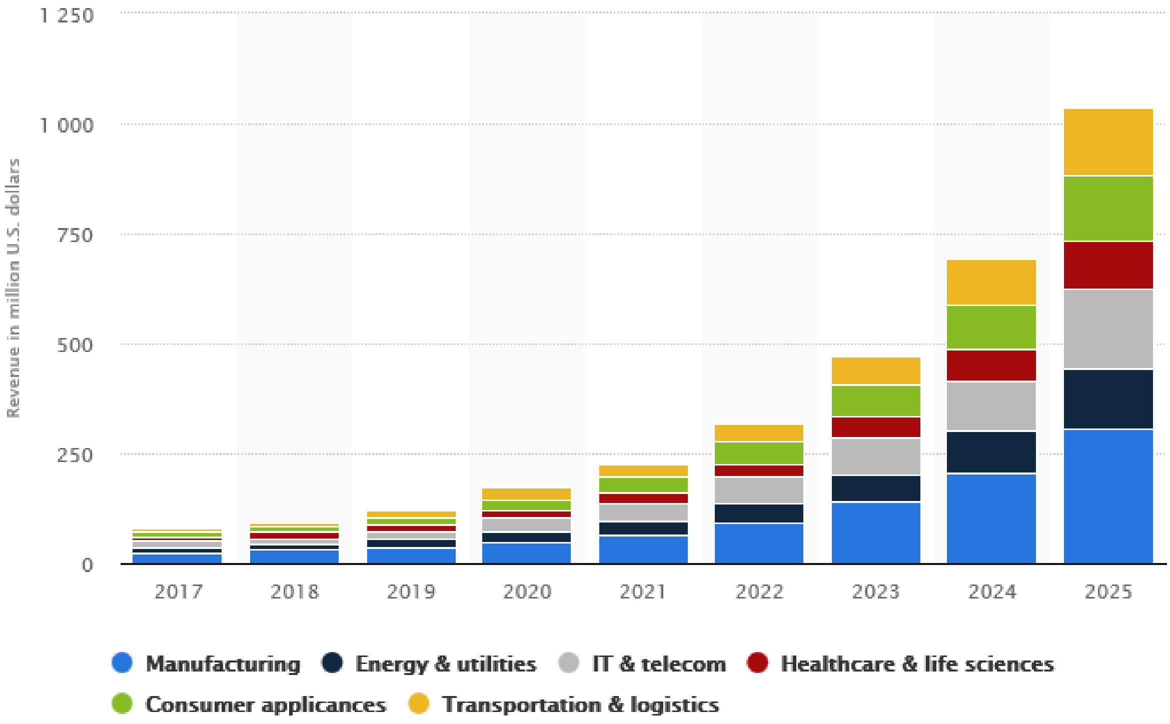

Figure 4 Edge computing market size forecast in the US from 2017 to 2025, by segment (in million $) [STA1]

However, some key aspects, such as the existent jungle of standards and the new GDPR [EUG1] recently adopted, are limiting the market growth. At the same time, while edge analytics reduces the latency of cloud analytics, there is still a gap in computation power to leverage the potential of IoT. Fog appears as the driver for solving these problems as reflected in its exponential growth (61.3% CAGR, $617.3 million by 2025) [GVR3], as shown in Figure 5. The support of resource heterogeneity, machine-to-machine communication, reduction of operational costs and possibility of distributed data analytics are the key pillars of the fog computing growth. However, as it happened with the edge computing, the lack of standards is limiting its possibilities and creating a need for homogenized and interoperable developments.

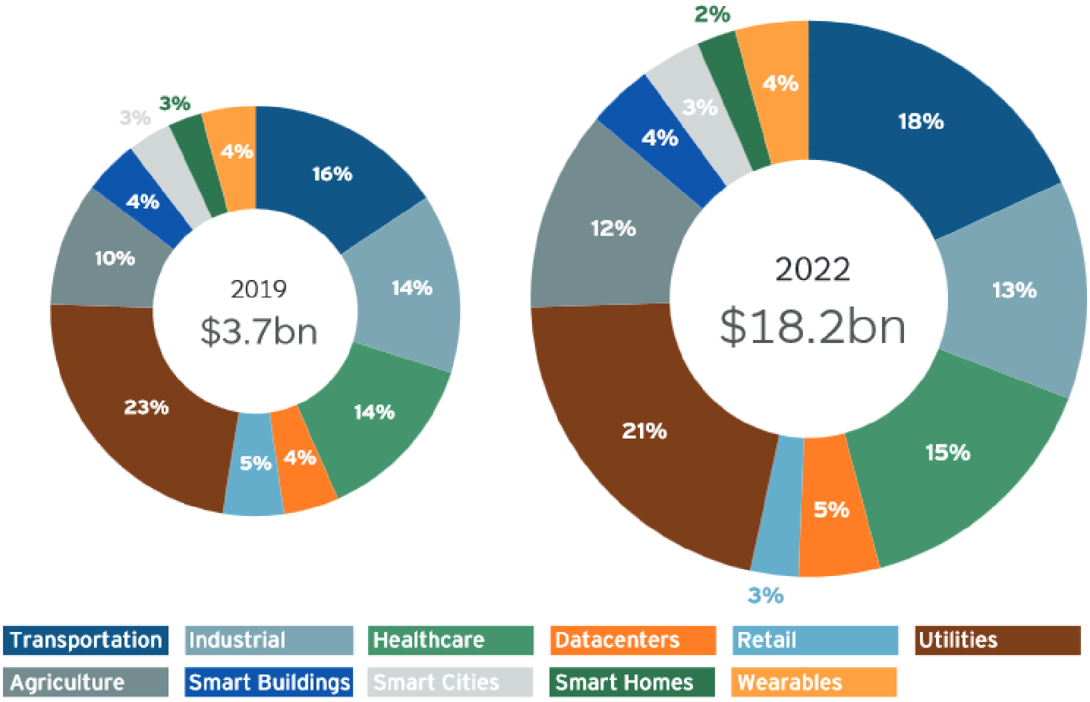

Source: 451 Research OpenFog project analysis

Figure 5 Size of Fog computing market opportunity by vertical market, 2019 and 2022 [451R1]

In order to increment the benefits of cloud, fog or edge, new technologies are arising, such as artificial intelligence (AI). According to Forrester [FORR1], which defines artificial intelligence as '*Cognitive computing*', the big data phenomenon is pushing for new technologies based on AI not only to maximise the types of data that can be analysed, but also to get improved results to remove silos of knowledge giving access to business insights that drive action. Markets and Markets [MMR1] sizes this blooming market in $190.61 billion by 2025, being the manufacturing industry the biggest niche. Thus, AI is expected to be the next boom in the technology market opening the door to a new era of AI-based applications.

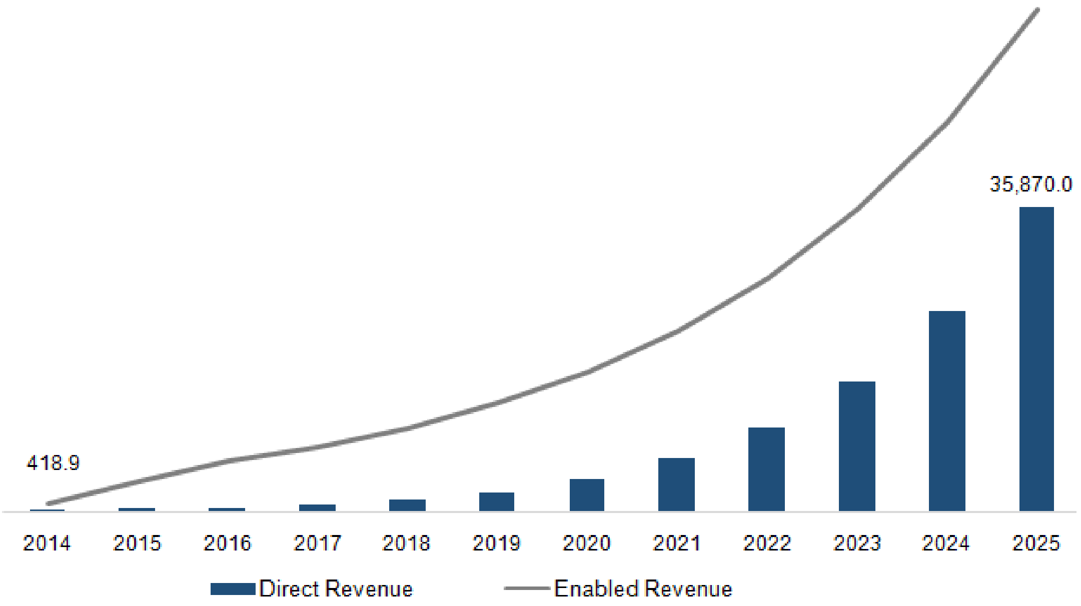

Figure 6 Artificial Intelligence – Direct & Enabled Revenue, 2014-2025 (USD Million) [GVR4]

However, although the expected forecasts for the AI market, see Figure 6, there are still some challenges for this market to be sorted out, mainly related to security and reliability, and more specifically to deep learning, as it is considered a bottleneck for the market growth. This is expected to be bypassed by the increased prominence of GPU-based applications that will increase the storage capacity, computation power and parallelization of deep learning technologies. Taking this into account, the global deep learning market is expected to reach $10.2 billion by 2025, at a CAGR of 52.1% [GVR5]. This growth will directly influence the Software-as-a-Service market (SaaS), as most of the applications will be powered by machine and deep learning to enhance the current data analytics offerings, as well as the Mobile Cloud market, with a new set of innovative and improved customer experience applications based on deep learning capabilities. Other markets such as cybersecurity, fraud detection, or data analytics are also expected to benefit from deep learning mainly by the increased usage of data mining. At the same time, the global chip market is experiencing a significant growth at a CAGR of 39.9% [ALR1] mainly due to the emergence of quantum computing and the usage of deep learning chips in robotics.

All markets analysed in this section present a high interdependency mainly based on the rise of Internet of Things and Big Data needs, as depicted in the following subsections.

4.2. Internet of Things

Internet of Things (IoT) is the base technology of digital business, essentially a network of dedicated physical objects (things) outfitted with data-collecting technologies and capable of communicating and interacting with their internal states or the external environment. The data generated has a wide range of uses but is commonly seen as a way to determine the health and status of things, from which a company can learn behaviour and usage, react with preventive action or augment or transform business processes. Most companies use the IoT to shift from corrective/reactive activities and maintenance plans to a system that reflects the real-time status of components.

Initially, the IoT was seen as a technology with the potential to solve all the IT and business problems the organizations faced. But without a concrete use case and a thought-through architecture and implementation, the IoT was essentially a solution looking for a problem [Gartner:2017]. 2015 was a year of education for IoT, in which evangelisation was performed to the enterprise sector, so it understood what the technology could offer to their organization. According to Verizon [Verizon:2017], 2016 was the year IoT gained significant momentum in the enterprise. In 2017, the companies started the implementation of proof-of-concepts and exploratory projects [IOTSWC:2017]. In 2018, IoT has entered into the next stage of business, bringing value instead of merely demonstrating its potential, and the actual integration of IoT in the companies has been performed, instead of only pilot projects [i-scoop:2018].

Marc Hung, Gartner Research vice president, reminds in a public document [Gartner:2017] that IoT will not always rely on the cloud, but also on the edge. For instance, when a fitness wristband is connected, most of the data treatment lies in the smartphone app, although some operations can be performed in the cloud so that the user can share fitness metrics results with friends or a healthcare provider. Considering mF2C project, novel technologies as Fog devices and Cloud management layers are cornerstone of automatic applications, and they are of great interest for this project since they provide a mandatory combination of programmable connectivity, rapid service provisioning and service chaining.

Apart from the growing evolution towards edge and fog computing, the IoT technological trends in 2018-2019 will mainly be blockchain, Artificial Intelligence and LPWAN. Finally, the increase of connected devices will put pressure on the services they offer to comply with strict bandwidth, delay, jitter, packet loss, redundancy, quality of service and reliability requirements. Similar requirements will be expected from the business applications. The benefits provided by IoT may then be translated to interesting business opportunities for utilities and manufacturers in order to reduce OPEX via CAPEX investments.

By 2020, more than 65% of companies will adopt IoT products, in contrast to the approximately 30% that have already done so. In 2010, Ericsson and Cisco announced predictions of 50 billion connected devices by 2020. The figure has dramatically been reduced since then and by 2020, Gartner now estimates that there will be four times more internet-connected things than humans, modifying the service sector dynamics for marketing, sales and customer service [Gartner:2017]. The estimation for the number of connected devices by 2020 ranges from 20 to 30.7 billion, depending on the source ([IoT_ecosystem:2015] [Gartner:2017], [Forbes:2018], [Ericsson:2015], [IHS market:2018], [Stringify:2018]). The following image, Figure 7, illustrates the number of IoT projects identified by segment and location, based on 1600 public projects. Most of them are located in the US, smart city being the leading segment, followed by Connected Industry and Connected Buildings. Connected buildings has experienced the largest segment increase since 2016. In that same year, the leading sector was Connected Industry. Although the US leads in deploying IoT in operations at full scale (44%), followed by the UK (41%) and Germany (31%) [CAPGEMINI:2018], the increase in smart city projects is led by Europe (164 projects), through cities such as Singapore and Barcelona. [IoT analytics:2018].

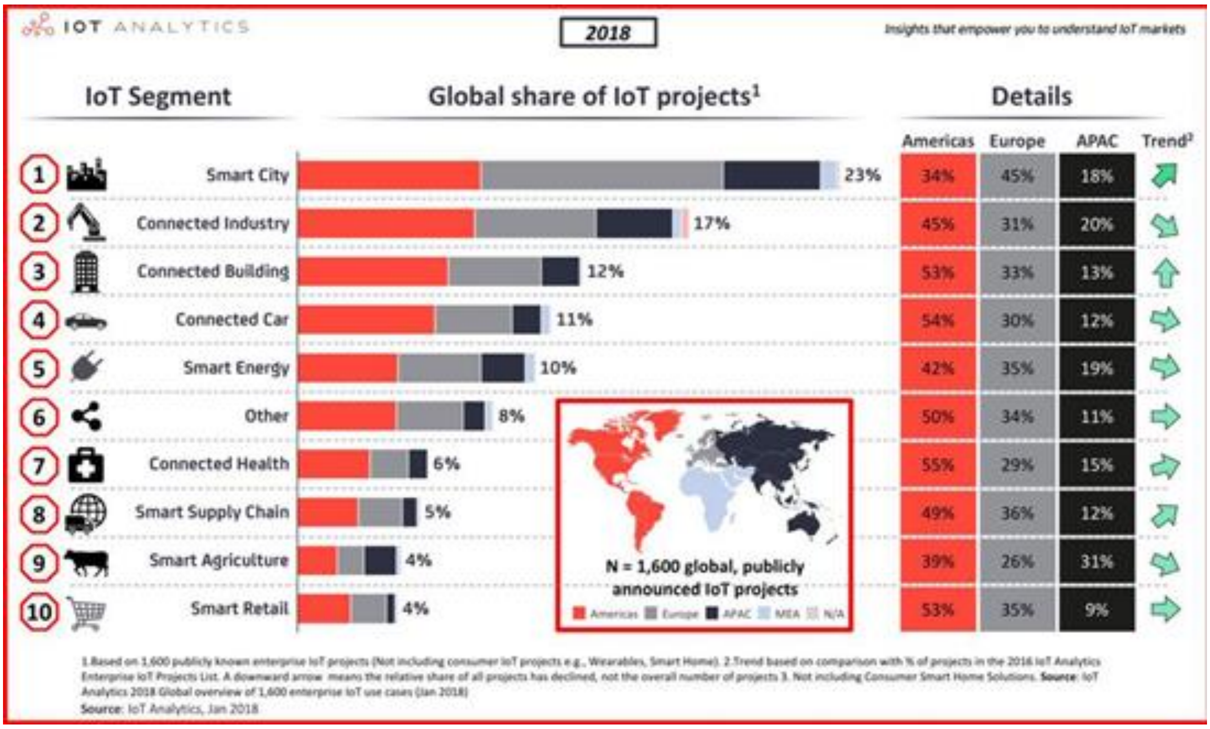

1.Based on 1,600 publicly known enterprise IoT projects (Not including consumer IoT projects e.g., Wearables, Smart Home). 2.Trend based on comparison with % of projects in the 2016 IoT Analytics Enterprise IoT Projects List. A downward arrow means the relative share of all projects has declined, not the overall number of projects 3. Not including Consumer Smart Home Solutions. **Source**: IoT Analytics 2018 Global overview of 1,600 enterprise IoT use cases (Jan 2018)

Source: IoT Analytics, Jan 2018

Figure 7 Global, publicly announced IoT projects [IoT analytics:2018]

It is worth noting that a survey was performed across executives of big companies and the result was that the absence of industry-wide IoT standards, together with security, interoperability and cost considerations, represented over 50% of their concerns about IoT [Verizon:2017]. This topic will be looked into in section 4.5.

The market estimates differ greatly depending on the source for instance the value IoT is expected to reach by 2024 is close to 6.5 trillion US dollars, compared to 1.2 in 2017, according to Energis Market research [Energias:2018]. In that same year, M2M devices and services are predicted to be worth $2.5 trillion. More conservative values state that the $1 trillion IoT spending mark will not be surpassed until 2020 [i-scoop:2018], while Visiongain assesses that the IoT market will generate revenues of $1,352 billion in 2018. As for the investment, according to IDC [i-scoop:2018], total spending on IoT in 2018 will reach $772.5 billion, with an expected $13 trillion return on investment (ROI) by 2025 [IoT ecosystem:2015]. McKinsey predicts the IoT market will be worth $581B for ICT alone by 2020, growing at a Compound Annual Growth Rate (CAGR) between 7 and 15%, and the Industrial Internet of Things (IIoT) market to reach $123B in 2021, attaining a CAGR of 7.3% through 2020. On the long term, Accenture forecasts IIoT can add as much as $14.2T to the global economy by 2030 [i-scoop:2018].

The number of IoT related patents has also exploded in the past few years, a significant turning point being 2012, although they are spread over several players and the bulk of the patents is not owned by a single company or a small group of companies. Samsung for instance, owns more than 4500 patents, Qualcomm more than 2800, LG and Huawei both over 2000, with several companies from diverse sectors like consumer electronics, telecom and software owning over 1000 patents (Sony, Ericsson, Nokia, Siemens, NEC, INTEL…).

Most patents concern EU and China, Europe being far behind, although most big players (except Sony) use the PCT (Patent Cooperation Treaty) route. Qualcomm for instance is the leader in PCT filings, which denotes its intention to license its technologies worldwide.

The network layer is where most of the current patents lie, the number of patents filed in this sector has been the highest of all sectors each year over the past ten years. This leaves big opportunities in the other areas of the IoT ecosystem, nevertheless every IoT-related technology except processors have shown a clear increase from 2011. The following figure [Relecura:2017], Figure 7, represents the players and the number of patents they own, on a sectorial map.

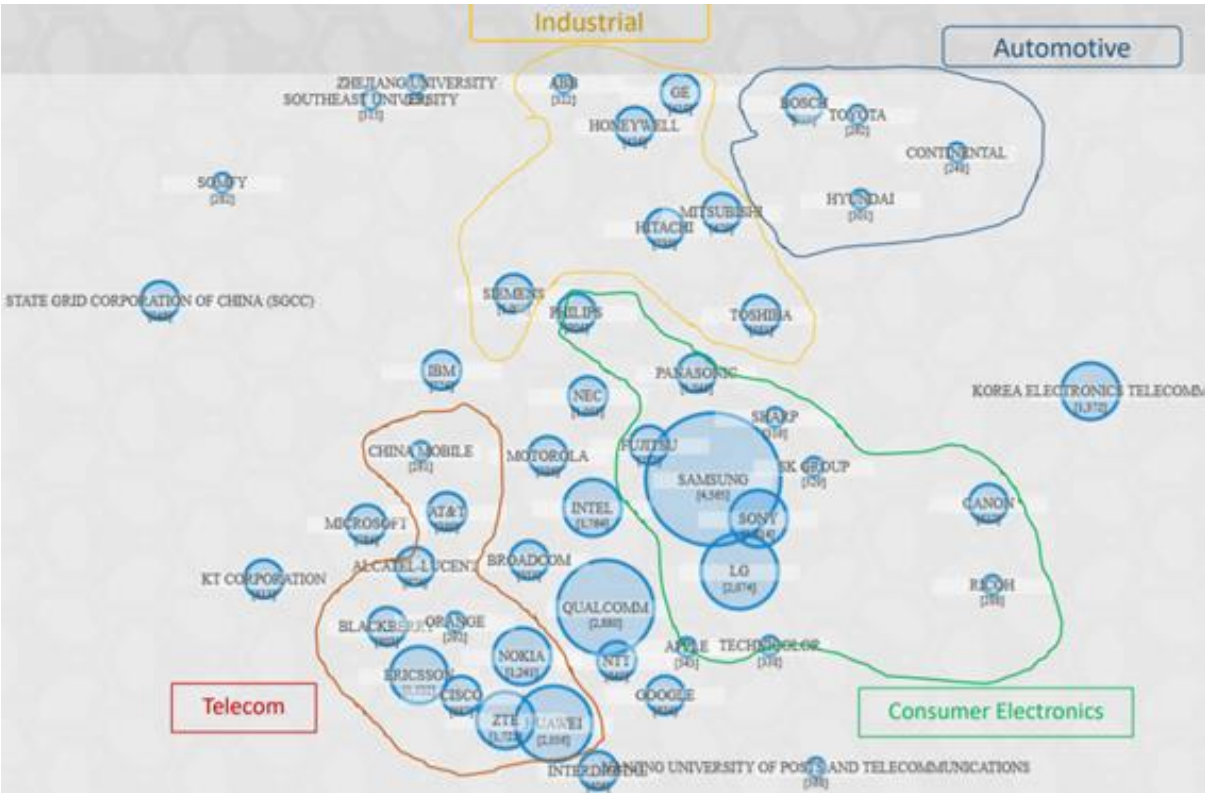

Figure 8 Sectorial map of IoT patents [Relecura:2017]

4.3. Big data and IoT

New technologies that allow to revolutionize, speed up and optimize production processes are coming. From steam we are moving to the digital and therefore we are entering the 4th industrial revolution. And coal that fed steam boilers today is called data. Data that in huge quantities are generated, arrive, transit and come out in all the modern companies. Data floods given the increasing quantity that is daily produced and exchanged, and data of a structured and unstructured types must therefore be carefully managed to gain possible advantages in terms of business. And, above all, data come from "things" that spreads on the enterprise systems, on the environment that surrounds us and also, thanks to wearables, on the people.

The evaluation of IoT applications in the environments where these systems will be deployed (e.g. cities, offices, shopping centres, hospitals, industries, etc.) shows a broader view of potential benefits and challenges, highlighting how various IoT systems can maximize value, in particular where they interact [McKinsey2015]. Interoperability between IoT systems is critically important: when IoT systems communicate each other their value is multiplied, so interoperability is an important enabler for maximizing benefits. But interoperability in the IoT field is a big challenge.

A solution for the IoT interoperability could be to have a single, unified communication and software framework but due to the diversity and "run fast" nature of the IoT this is not an option. Diversity in the IoT is an aspect that must be accepted and managed. Three are the key elements of IoT interoperability. The first key element is multimode radios to allow diverse IoT devices to talk to each other, e.g. Home hubs such as routers and gateways supporting Wi-Fi, Bluetooth and 802.15.4 technologies can also support interoperability by acting as bridges across multiple ecosystems. The second key element for the interoperability of IoT is software flexibility that enables support for connectivity frameworks, cloud services and multiple protocols. Of course, memory and processing power are critical considerations to enable software flexibility in IoT devices. The third key element is the need for hardware-based security and strong security at each node [Qualcomm2017].

The current use of IoT data is mainly to address anomaly detection, real-time control, smart metering and asset tracking so additional value remains to be captured, by using more data, as well as deploying more sophisticated applications such as using performance data for predictive maintenance to predict and prevent breakdowns, or to analyse workflows to optimize operating efficiency [Forrester2016]. But a lot of IoT data are also coming from innovative IoT projects in the smart cities, logistics and agriculture sectors that focus on human-centred and outcome-based approach [Deloitte2018]. IoT can be a key source of big data to be analysed to capture value [Forrester2016].

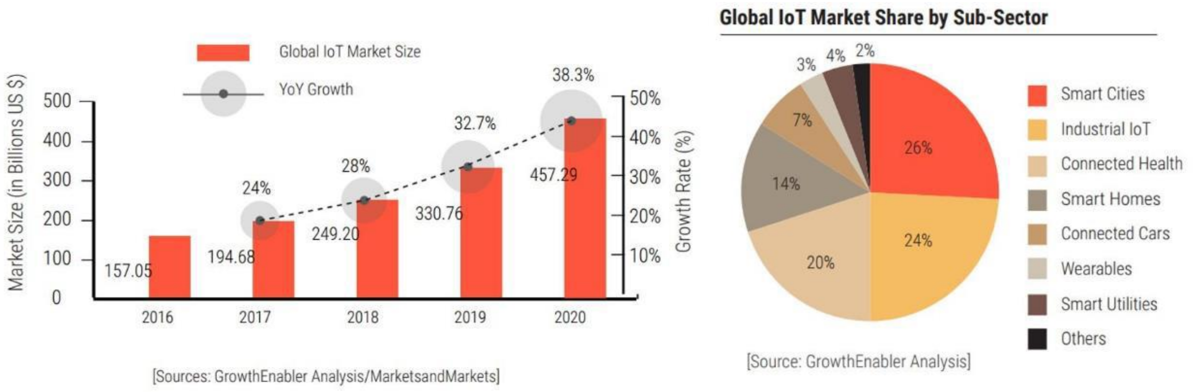

[Sources: GrowthEnabler Analysis/MarketsandMarkets]

[Source: GrowthEnabler Analysis]

Figure 9 Market pulse report IoT, 2017 [GRWTH17]

Business-to-Business (B2B) applications can create more value than pure consumer applications. While consumer applications such as fitness and e-Health monitors, home automation, intelligent speakers and self-driving cars attract the most attention and have tremendous potential for creating significant value, there is even greater potential value from IoT use in B2B applications. B2B market is expected to generate nearly 70 percent of potential value enabled by IoT [McKinsey2015].

IoT Value chain integration for organizations in B2B industries enables also direct interaction with end consumers, by extending cooperation to a business-to-business-to-consumer (B2B2C) model (see Figure 10). IoT devices embedded within or integrated with physical products enable organizations to extend their value chain beyond direct buyers or suppliers in the chain [EY2018].

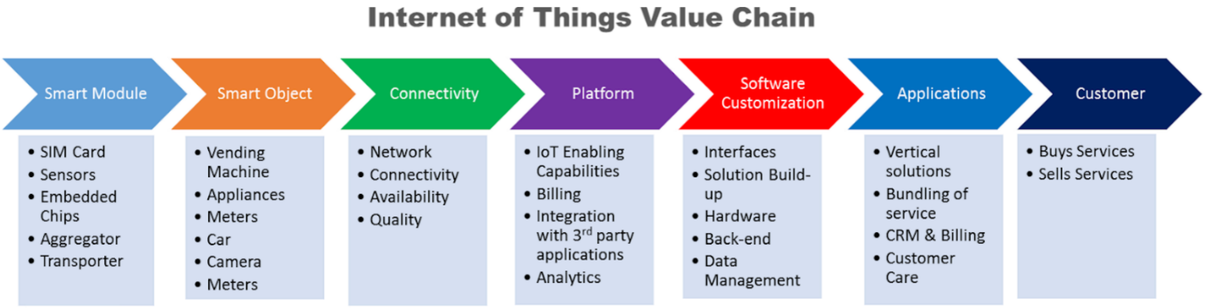

Figure 10 Source: recap-project.eu [RECAP]

Industry 4.0 is the digitization of the manufacturing sector, with embedded sensors in virtually all product components and manufacturing equipment, ubiquitous cyber physical systems and analysis of all relevant data [McKinsey2015bis]. The industry is evolving around IoT technology and IoT will change the bases of competition and drive new business models for user and supplier companies. The Internet of Things will enable—and in some cases force—new business models. For example, with the ability to monitor machines that are in use at customer sites, makers of industrial equipment can shift from selling capital goods to selling their products as services. This “as-a-service” approach can give the supplier a more intimate tie with customers that competitors would find difficult to disrupt [McKinsey2015].

To realize the full potential from IoT applications, technology will need to continue to evolve, providing more sophisticated and robust data analytics. In particular services and IoT analytics and applications are expected to capture 60% of the growth from IoT in the next 5 years [BCG2017]. As an example of the support of the latest end-to-end technology to this trend there are several public cloud services like AWS IoT Analytics and Google Insight that allows you to run sophisticated analytics on a large quantity of IoT data.

In almost all deployed or foreseen settings, IoT systems raise questions about data security and privacy, solution providers and enterprises need to work together to protect break points, as well as enables rapid detection and mitigation of security breaches [Deloitte2016].

Most organizations, taking advantage of the IoT opportunities, will require leaders to fully adopt data-driven decision making [McKinsey2015]. This point is rising among directors that see the Big Data framework, as an investment priority for the coming years [PoliMI2016]. A data-driven decision-making path from Descriptive Analytics, what’s happened, to Predictive Analytics, what will happen, to Prescriptive Analytics, what can be done, and to Automated Analytics, with automatic actions without human control when fast response is mandatory, e.g. in finance or health scopes is here to stay [Gartner2016].

A scalable infrastructure able to process large amount of data in real time is also needed. The analysis phase requires an evaluation of complex architectures that combine capabilities of real time and batch processing. The IoT will speed up this evolution path because IoT produces huge quantities of a type of asset that can be sold or exchanged: the data, coal but also diamond of the 4th industrial revolution. The ability to identify facts, hidden relations in the data available to organizations, not only allows to optimize processes and increase competitiveness, but also can open new opportunities for value creation.

Data monetization, in a data-based economy, can generate new revenues through the sale or exchange of data in the possession of the organization and the exploitation of these data can be a driver for the generation of new products and services. Many companies are launching data-focused businesses, but data monetization is still in its early days. Getting data monetization right requires significant effort, but it's becoming critical for staying ahead of traditional competitors and new disruptors [McKinsey2017].

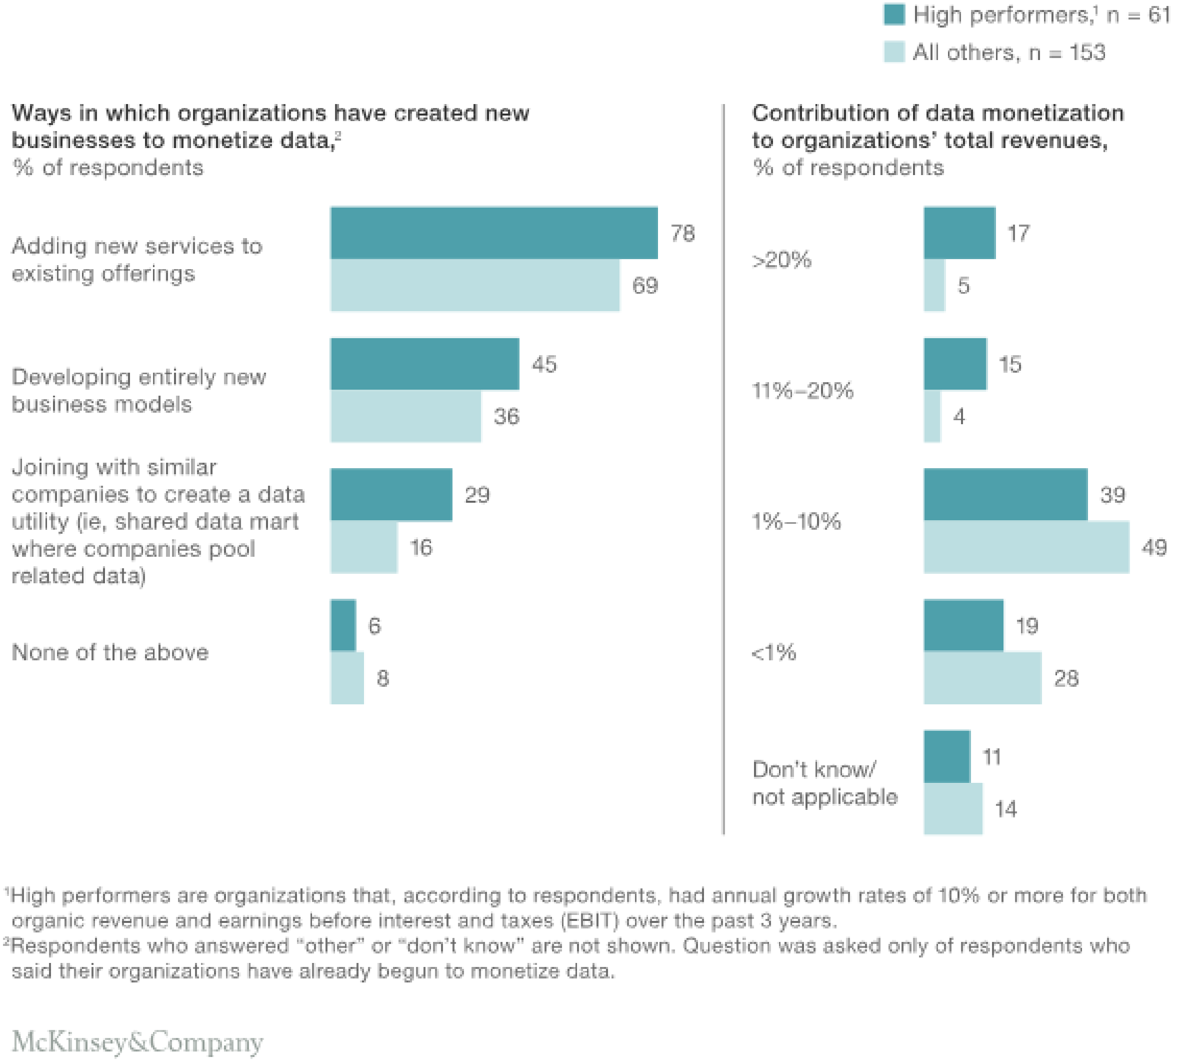

Figure 11 McKinsey, Fuelling growth through data monetization, December 2017 [Gartner2017]

The evolution of the data monetization, see Figure 11, phenomena can be slower than expected due to several factors like the lack of professional expertise, e.g. one of the possible limits can be a lack of data science specialists that will inhibit 75% of organizations from achieving the full potential of IoT in the Big Data environment until 2020 [Gartner2017].

4.4. Security Trends

The big message in the IT industry in 2017 was ransomware. Ransomware is still widespread in 2018 - in fact, some industry insiders say there is more of it this year.

Dashboards would enable monitoring of remote systems and suspected attacks and would provide controls to enable an incident handler to react, including, perhaps, getting an edge device to wipe its own memory (in case it contains sensitive information), or at least quarantining it.

The principal threats highlighted by the industry, other than ransomware, include crypto jacking, container security, mobile devices and IoT devices. Another threat theme flagged by a few companies is, perhaps surprisingly, open source software. The “threat” seems to be that commercial applications may have included open source software (e.g. libraries), either inadvertently or to speed up development, or through a supply chain. At the simplest level, the software has then been “contaminated” by the licence (such as GPL), but they are also looking for bugs and vulnerabilities (through an agile “security devops” approach), and providing additional security management through containerisation.

Less mature, but of broad interest across the industry is using machine learning and AI to assist and automate reactions to security incidents. Keywords in AI are to ensure the right amount of data is available: as the domain is sufficiently narrow, AI is expected to do reasonably well, aiding and advising an incident handler (“augmented intelligence”), and, as is increasingly expected from AI, explaining the reasons for choices.

With reference to IoT and F2C security, the focus is currently on using blockchain for building distributed trust, enforcing access control and privacy. Blockchain is a specialised, decentralised data storage technology, perhaps better known as the platform that powers cryptocurrencies such as bitcoins. As the technology matures, it has outgrown its original purpose and found many different applications in IoT security such as self-sovereign identity, e.g. uPort [uPort:2018], trusted M2M marketplace for sensor data streams, e.g. StreamR [Streamr:2018], IOTA [IOTA:2018] and security protocol and infrastructure for IoT device interactions. E.g. Atonomi [ATML:2018], etc.

Blockchain’s popularity is rooted in its ability to enhance trust among participants across a peer network. The F2C environment is inherently untrustworthy as unlike an enterprise system, it has no boundary; an F2C system is built on autonomous nodes that opted to work together but these nodes may not trust each other [Cisco:2018]. Thus blockchain is viewed as the panacea for instilling trust in F2C systems. Indeed, the OpenFog Consortium is committed to develop a composable blockchain architecture to secure the fog environment [Irwan:2018].

Blockchain enhances trust through two main mechanisms:

- Its use of consensus algorithms to ensure that the participants agree on the state of the stored data; thereby assuring the trustworthiness of the transactions. The consensus, in turn provides a yardstick by which participants can self-police each other’s behaviours, and
- Its immutable fabric. Every blockchain transaction is built on every previous transactions. Tampering with the stored transactions will break the consistency of the transaction chain and the current consensus.

Blockchain also offers other useful features [Gupta:2017] for enhancing security and privacy in F2C systems such as:

- Redundancy and availability: The chain data are replicated selectively across the network nodes in near real time, thereby avoiding single points of failure

- Self-healing: the nodes run automated, self-audit process regularly to synchronise their data. In an environment of volatile network connection, this feature allows a node to self-heal and achieve consistency
- Secure, private and indelible: access control (for permissioned blockchain) and cryptography prevent unauthorised access to help maintain privacy of the transactions. Details of the transactions and the identities of the participants can be masked for added protection. Blockchain transactions are indelible, changes to a transaction must be corrected using a new counter transaction. As transactions cannot be deleted, this contravenes the GDPR “Right to be forgotten” principle. The best practice is to store encrypted private data off-chain and the transaction stores as evidence a cryptographic hash of the data
- Transparent and auditable: transactions are time-stamped and replicated in near real time. Participants in a transaction have access to the same information and can validate the transactions independently. The indelible and immutable features enable audit trails and facilitate compliance to the GDPR principles of accountability and security (confidentiality and integrity)
- Process orchestration: most blockchain platforms supports embedding code to automate business rules and smart contracts. For example, [Zhang:2018] prototyped a decentralised access control framework for IoT using smart contract
- Sustainable: a blockchain network is decentralised; it is not owned or controlled by a single organisation. Nodes can join and leave but its continuous existence does not depend on any individual entity. This model fits in with the autonomous nature of the fog.

4.5. Digital Business

Digital transformation is supporting the appearance of new digital businesses for commercializing non-tangible services. In such a competitive business arena, organizational change is inevitable and every day more and more companies from traditional sectors are embracing cloud services as a first step for a full digital transformation. Virtual enterprises are less now sci-fi and more a blooming trend worldwide, where everything is connected through Internet, fostering the rise of new businesses.

A good example of this transformation is the rising of novel businesses based on the principle of sharing economy for tangible goods and services. These fair revenue schemes now appear to be based on the share of digital resources in i.e. cloud or IoT platform federations.

On the other hand, opposite to what it is commonly understood, the rise of innovative and trustable open source software solutions is fostering a wider adoption of cloud-based solutions, as enterprises can find high quality offerings at lower prices. Open source communities’ products such as Docker, OpenStack or MySQL are paving the road for digital transformation across businesses.

Thus, it is expected the raise of new business that can maintain and leverage the European competitiveness worldwide.

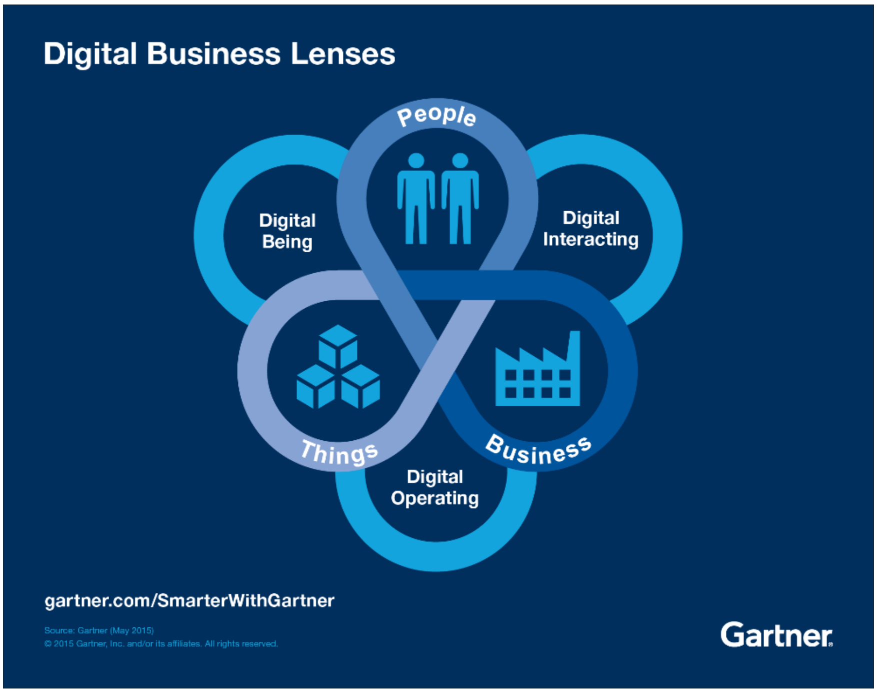

Source: Gartner (May 2015)
© 2015 Gartner, Inc. and/or its affiliates. All rights reserved.

Figure 12 Digital Business Lenses for Digital Business Opportunity [GAR5]

Gartner considers digital business from two different perspectives: people and business interactions, and things and business interactions, in order to determine the best path for each service vendor, see Figure 12.

Based on the trends already identified along this chapter, mF2C will develop a set of business models based on digital business to business (B2B) interactions for its different stakeholders, adapted to their expectations and needs after the appropriate lens, as suggested by Gartner, as well as some value proposition models for fostering the adoption of mF2C, based on the added value that such a solution can bring to the final customers (business-to-customer (B2C) interactions), reinforced with the information extracted from the project use cases.

4.6. Key Takeaways

We summarize the key areas of focus in business trends significant for the project:

- Sectors like Manufacturing, big data, data analytics, and storage are still being seen as the main drivers of cloud adoption driving a global cloud market to reach an expected $1.25 billion
- SaaS market remains the largest segment, with IaaS experiencing fastest growth. Data analytics is driving the database platform as a service (dbPaaS) expected to reach $10 billion by 2021.
- Global deep learning market expected to reach $10.2 billion by 2025, influencing SaaS market - most apps will be powered by machine and deep learning enhancing the current data analytics offerings
- Reports are showing IoT is only starting to gain momentum in the enterprise in the last 2 years, the focus being unclear before on how it would solve IT and business problems
- The growing evolution towards edge and fog computing, the IoT technological trends in 2018-2019 will mainly be blockchain, Artificial Intelligence and LPWAN
- By 2020, more than 65% of companies will adopt IoT products, in contrast to the approximately 30% that have already done so, connected devices by 2020 ranges from 20 to 30.7 billion

- absence of industry-wide IoT standards, together with security, interoperability and cost considerations, represented over 50% of their concerns about IoT
- market estimates differ greatly depending on the source for instance the value IoT is expected to reach by 2024 is close to US$6.5 trillion
- Total spending on IoT in 2018 will reach $772.5 billion, with an expected $13 trillion return on investment (ROI) by 2025
- IoT related patents increased: Samsung 4500 patents, Qualcomm 2800, LG and Huawei both over 2000, with several companies from diverse sectors like consumer electronics, telecom and software owning over 1000 patents (Sony, Ericsson, Nokia, Siemens, NEC, INTEL…).
- security trends in the IT and IoT security industry now focus more on integrated “solutions” including dashboards for security monitoring and SIEM (Security Incident Event Management).
- Blockchain outgrown its original purpose and found many different applications in IoT security such as self-sovereign identity. Focus is currently on using blockchain for building distributed trust, enforcing access control and privacy
- The rise of innovative and trusted open source software solutions is fostering a wider adoption of cloud-based solutions, as enterprises can find high quality offerings at lower prices, eg, docker, OpenStack or MySQL

5. Conclusions

In summary, all the technology trends described in D2.1 are still valid and progressing with the evolution of IoT, Edge and Fog computing. The mF2C project is already profiting from some of these trends, eg, some of the existing implementations already show great potential. The adoption of Docker during IT-1 is one such example of how mF2C is trying to be at the frontline of the most promising technologies for edge computing, and how these are being applied, seamlessly and throughout all computing layers from the cloud to the fog. Some work is still required though, as the project moves into IT-2, mostly on how to ensure a secure management and execution of services within mF2C and how to address the bridge to IoT sensors (source of data). This is where mF2C is hoping to learn and further develop the existing scenarios provided by reference solutions like the NuvlaBox Nano and the Akraino Edge Stack.

The scientific trends reviewed in chapter 2 show Cloud and Fog computing as being conceptually similar, but the challenges faced designing resource management solutions for fog computing systems continue to be heterogeneity, dynamicity, geo-distribution, and multiple owners of the devices taking part in that fog system. Partial solutions are appearing that attempt to address these problems, eg, the labelling system within Fogernetes, decentralized algorithms to offload decisions and tasks, or game theory approaches in MSCOG. The project needs to ensure that it incorporates these solutions within its modules but remain aware that it was shown that some of these systems encounter difficulties when scaling and are possibly not ideal for activities expected in a fog context. Solutions for offloading application parts of compute to local resources remains a key research area and a key functionality offered by the mF2C framework, so application fragmentation solutions should be investigated further (eg, CloneCloud, Cuckoo, ThinkAir, MobileFog). As security remains a key architectural component, the alternative solutions of Machine Learning and AI to identify anomalies (eg, DDoS and request/response poisoning attacks) should continue to remain a priority due to the trend that software is struggling to keep up with the scale of attacks so these ML/AI techniques provide promising detection of 0-day exploits. AI algorithms can benefit from the inherent parallelism in the proposed mF2C architecture, allowing the distributed implementation of ML algorithms, and then improving the performance when computing data.

The technology trends reviewed in chapter 3 included both Cloud and IoT Management tools which provide solutions for resource management and service scheduling, and for interconnecting sensing devices, data, and applications over the Web to address interoperability. There appear to be an opportunity to extend the state of the art lies in the merging of high performance computing at the edge, given the lack of many solutions in this field so the project will continue to monitor this area. The second key area identified here is the lack of any programming framework to allow developers create software for heterogeneous devices, due to issues related to the fog-to-cloud paradigm. So the mF2C framework will solve real word problems if they can provide solutions for real time processing, latency and transparent management of a decentralized, heterogeneous and dynamic set of resources. Containers is appearing to be the key technology for service placement at the edge due to the lack of virtualization support on compute restricted devices. Therefore, the project should remain focused on container technologies and supporting tools. Finally, new reference architectures that deliver edge computing solutions have appeared on the market which the project should continue to monitor.

Finally, the business trends reviewed in chapter 4 evaluated research from different business analysts groups to help feed the business value that the mF2C framework could generate. These include reports showing adoption of IoT products is rising with estimates of connected devices in the range

20-30billion by 2020. These reports identify Blockchain, AI and LPWAN as key to achieving this which were documented earlier in the report. Regarding Blockchain, we're starting to see different applications of the technology beyond its original purpose in the fields of IoT security, self-sovereign identity, distributed trust, enforcing access control and privacy. So the project should look to include this technology in these areas. Unfortunately, these include differing market estimates depending on the source suggesting the value IoT reaching anywhere up to US$6.5 trillion by 2024. This large number assumes a total spending on IoT in 2018 reaching $772.5 billion, with an expected $13 trillion return on investment (ROI) by 2025. These numbers will help feed our value proposition statement captured in the annual report on exploitation due at the end of year 2. Finally, there is opportunity for the project to contribute to standards due to an absence of IoT standards, including security, interoperability and cost considerations, which represent greater than 50% of concerns regarding IoT coming from industry sources.

With this deliverable, the project extends its understanding of the current scientific, technical and business trends in Fog, Edge and Cloud computing. This will feed the design of the 2nd version of the project architecture due M25. Awareness of these trends helps focus the project on the problem areas requiring new solutions and steer it away from problem areas that solutions already exist. We will continue to avail of existing codebases to accelerate the prototyping of the Platform Manager and Agent Controller. Based on the identified trends in this report, the project will develop business models for its potential stakeholders. This will also include value proposition models for the adoption of mF2C, based on the value the framework could bring to end users.

References

[D21] D2.1 Tracking Scientific, Technology and Business Trends (Version 1) [Online]. Available: http://www.mf2c-project.eu/wp-content/uploads/2017/04/mF2C-D2.1-Tracking-Scientific-Technology-and-Business-Trends-Version-1.pdf

[DockerEdge] Docker Edge. https://www.docker.com/solutions/docker-edge

[CLASS] CLASS Project. Edge and Cloud Computation: a Highly Distributed Software for Big Data Analytics
[Online]. Available: https://class-project.eu/

[dataClay] dataClay storage platform [Online]. Available: https://www.bsc.es/dataClay

[Kinetic] Cisco Kinetic [Online]. Available: https://www.cisco.com/c/en/us/solutions/internet-of-things/iot-kinetic.html

[RedisEnt] Redis Enterprise [Online]. Available: https://redislabs.com/blog/ideal-iot-edge-database-redis-enterprise/

[BH+18] Neha Belapurkar, Jacob Harbour, Sagar Shelke, Baris Aksanli. Building Data-Aware and Energy-Efficient Smart Spaces. IEEE Internet of Things (2018). DOI: 10.1109/JIOT.2018.2834907

[CLP17] Bastien Confais, Adrien Lebre, Benoît Parrein. Performance Analysis of Object Store Systems in a Fog and Edge Computing Infrastructure. T. Large-Scale Data- and Knowledge-Centered Systems 33: 40-79 (2017)

[EU2018] EU GDPR https://www.eugdpr.org/

[DG+17] Utsav Drolia, Katherine Guo, Jiaqi Tan, Rajeev Gandhi, Priya Narasimhan. Cachier: Edge-Caching for Recognition Applications. ICDCS 2017: 276-286

[GR18] Harshit Gupta, Umakishore Ramachandran. FogStore: A Geo-Distributed Key-Value Store Guaranteeing Low Latency for Strongly Consistent Access. DEBS 2018: 148-159

[GXR18] Harshit Gupta, Zhuangdi Xu, Umakishore Ramachandran. DataFog: Towards a Holistic Data Management Platform for the IoT Age at the Network Edge. HotEdge 2018

[LM10] Avinash Lakshman, Prashant Malik. Cassandra: a decentralized structured storage system. Operating Systems Review 44(2): 35-40 (2010)

[MQ+17] Jonathan Martí, Anna Queralt, Daniel Gasull, Alex Barceló, Juan José Costa, Toni Cortes. Dataclay: A distributed data store for effective inter-player data sharing. Journal of Systems and Software 131: 129-145 (2017)

[APA15] Apache Spark https://spark.apache.org/

[GOO17] Federated Learning Collaborative https://ai.googleblog.com/2017/04/federated-learning-collaborative.html

[MS+17] Seyed Hossein Mortazavi, Mohammad Salehe, Carolina Simoes Gomes, Caleb Phillips, Eyal de Lara: Cloudpath: a multi-tier cloud computing framework. SEC 2017: 20:1-20:13

[PA+18] Ioannis Psaras, Onur Ascigil, Sergi Rene, George Pavlou, Alexander Afanasyev, Lixia Zhang. Mobile Data Repositories at the Edge. USENIX HotEdge 2018

[RG18] Arun Ravindran, Anjus George. An Edge Datastore Architecture For Latency-Critical Distributed Machine Vision Applications. USENIX HotEdge 2018

[SV+13] Jeff Shute, Radek Vingralek, Bart Samwel, Ben Handy, Chad Whipkey, Eric Rollins, Mircea Oancea, Kyle Littlefield, David Menestrina, Stephan Ellner, et al.F1: A distributed sql database that scales. Proceedings of the VLDB Endowment, 6(11):1068–1079, 2013.

[IDCFS] Carrie MacGillivray, Vernon Turner, Ruthbea Yesner, Jill Feblowitz, Kimberly Knickle, Lionel Lamy, Milly Xiang, Andrea Siviero, Mike Cansfield. IDC FutureScape: Worldwide Internet of Things 2016 Predictions. Nov 2015 - IDC FutureScape

[Chun:2011] Chun, B.G., Ihm, S., Maniatis, P., Naik, M., Patti, A.: Clonecloud: Elastic execution between mobile device and cloud. In: Proceedings of the Sixth Conference on Computer Systems. pp. 301–314. EuroSys ’11, ACM, New York, NY, USA (2011), http://doi.acm.org/10.1145/1966445.1966473 .

[Kemp:2012] Kemp, R., Palmer, N., Kielmann, T., Bal, H.: Cuckoo: A Computation Offloading Framework for Smartphones, pp. 59–79. Springer Berlin Heidelberg, Berlin, Heidel- berg (2012), http://dx.doi.org/10.1007/978-3-642-29336-8 4

[Kosta:2012] Kosta, S., Aucinas, A., Hui, P., Mortier, R., Zhang, X.: Thinkair: Dynamic resource allocation and parallel execution in the cloud for mobile code offloading. In: 2012 Proceedings IEEE INFOCOM. pp. 945–953 (March 2012)

[Hong:2013] Hong, K., Lillethun, D., Ramachandran, U., Ottenwa¨lder, B., Koldehofe, B.: Mobile fog: A programming model for large-scale applications on the internet of things. In: Proceedings of the Second ACM SIGCOMM Workshop on Mobile Cloud Computing. pp. 15–20. MCC ’13, ACM, New York, NY, USA (2013), http://doi.acm.org/10.1145/2491266.2491270

[Gartner:2017] Leading the IoT – Gartner insights on how to lead in a connected world report, Mark Hung

[Verizon:2017] State of the Market: Internet of things 2017 – Making way for the entreprise

[IOTSWC:2017] IT and OT debate at IOTSWC17, October 2017

[IoT ecosystem:2015] https://internet-of-things-innovation.com/iot-ecosystem/

[Ericsson:2015] https://www.ericsson.com/assets/local/news/2016/03/ericsson-mobility-report-nov-2015.pdf

[IHS market:2018] https://ihsmarkit.com/index.html

[Stringify:2018] https://www.stringify.com/

[CAPGEMINI:2018] internet of things study, march 2018

[IoT analytics:2018] https://iot-analytics.com/top-10-iot-segments-2018-real-iot-projects/] .

[McKinsey:2018] https://www.mckinsey.com/business-functions/digital-mckinsey/our-insights/the-iot-as-a-growth-driver

[WIKI1] Infusion Pumps https://en.wikipedia.org/wiki/Infusion_pump

[Relecura:2017] Relecure IP Intelligence report, IoT – Internet of things – Technology landscape and IP commercialisation trends (May 2017). https://relecura.com/reports/IoT_IP_Landscape_Commercialization_May2017.pdf

[Forbes:2018] https://www.forbes.com/sites/louiscolumbus/2018/06/06/10-charts-that-will-challenge-your-perspective-of-iots-growth/#4b6dc4913ecc

[i-scoop:2018] https://www.i-scoop.eu/iot-2018-1/

[Energias:2018] https://globenewswire.com/news-release/2018/04/17/1479964/0/en/Global-Internet-of-Things-IoT-Market-to-witness-a-CAGR-of-26-6-during-2018-to-2024-Energias-Market-Research-Pvt-Ltd.html

[Visiongain:2018] https://www.visiongain.com/report/internet-of-things-iot-market-report-2018-2028/

[Azure:2018] https://azure.microsoft.com/en-us/services/functions/

[AWS:2018] https://aws.amazon.com/greengrass/

[AndroidThings:2018] https://developer.android.com/things/

[Cisco:2018] Blockchain and Fog: Made for Each Other, June 2018. https://blogs.cisco.com/innovation/blockchain-and-fog-made-for-each-other

[uPort:2018] Open Identity System for the Decentralized Web. https://www.uport.me/ (accessed on 29 August, 2018)

[Streamr:2018] Streamr https://www.streamr.com/ (accessed on 29 August, 2018)

[IOTA:2018] IOTA https://www.iota.org/ (accessed on 29 August, 2018)

[ATML:2018] Atonomi https://atonomi.io/ (accessed on 29 August, 2018)

[Irwan:2018] Irwan, Susanto, Redesigning Security for Fog Computing with Blockchain. https://www.openfogconsortium.org/redesigning-security-for-fog-computing-with-blockchain/ (accessed on 29 August, 2018)

[Gupta:2017] Gupta, Manav, Blockchain for Dummies. https://www-01.ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=XIM12354USEN (accessed on 29 August, 2018)

[Zhang:2018] Zhang, Yuanyu, Kasahara, Shoji, Shen, Yulong and Jiang Xiaohong, Smart Contract-Based Access Control for the Internet of Things. https://arxiv.org/abs/1802.04410 (accessed on 29 August, 2018)

[Guesmi:2018] Guesmi, Hattab, Wireless Smart Sensor Networks for Real-time Warning System of Flash Floods and Torrents in KSA. International Journal of Computer Applications, Vol. 165:6, May 2017.

[Pooja:2018] Pooja, H.K., Pallavi, R., Vedhitha, N.S., Roa, A. and M.S. Chandini, A Forest Monitoring System based on GPRS & Powered by IoT. International Journal for Scientific Research and Development Vol. 6, Issue 03, June 2018. http://www.ijsrd.com/articles/IJSRDV6I30778.pdf (accessed on 29 August, 2018)

[EDC:2018] IoT Growth Pushing the Data Centre to the Edge. https://www.comparethecloud.net/articles/iot-growth-data-center-edge/ (accessed on 30 August, 2018)

[SCHR:2018] AST Modular/Schnedier Electric. https://www.schneider-electric.com/en/brands/ast-modular/ast-modular-index.jsp (accessed on 30 August 2018)

[ECX:2018] Edgeconnex. http://www.edgeconnex.com/ (accessed on 30 August, 2018)

[Che17] A. Chesla, "When an attack means murder: The IoT healthcare security vulnerability," 2017. https://medium.com/@Avi.Chesla1/when-an-attack-means-murder-the-iot-healthcare-security-vulnerability-b29ff83a2a0f

[War17] C. Warren, “How Common is Identity Theft? (Updated 2017) The Latest Stats,” 2017. https://www.lifelock.com/education/how-common-is-identity-theft/

[BH16] Becker’s Healthcare, “Hospitals are hit with 88% of all ransomware attacks,” 2016. https://www.beckershospitalreview.com/healthcare-information-technology/hospitals-are-hit-with-88-of-all-ransomware-attacks.html

[Fim17] M. Fimin, “Five biggest security technology trends for 2018,” 2017. https://www.itproportal.com/features/five-biggest-security-technology-trends-for-2018/

[Pet18] C. Pettey, “Gartner Top 6 Security and Risk Management Trends For 2018,” 2018. https://www.gartner.com/smarterwithgartner/gartner-top-5-security-and-risk- management-trends/

[Pan17] K. Panetta, “Gartner Top 10 Strategic Technology Trends for 2018,” 2017. https://www.gartner.com/smarterwithgartner/gartner-top-10-strategic-technology-trends-for-2018/

[VER18] Verizon, “Verizon’s 2018 Data Breach Investigations Report,” 2018. https://www.verizonenterprise.com/verizon-insights-lab/dbir/

[SYM18] Symantec, “2018 Internet Security Threat Report,” 2018. https://www.symantec.com/security-center/threat-report

[BDVA17] Big Data Value Association (BDVA), “European Big Data Value Strategic Research and Innovation Agenda, Version 4.0,” 2017. http://www.big-data-value.eu/bdva-announces-the-official-sria-version-4/

[Dua17] J. Duarte, “The real danger of Artificial Intelligence – it’s not what you think,” 2017. https://hackernoon.com/the-real-danger-of-artificial-intelligence-its-not-what-you-think-f7fdc7059cf8

[Yoon17] J. Yoon, K. Shin, Y. Won, “Encrypted Network Traffic Analysis Method via Secure Socket Layer Handshake Control,” in Advanced Multimedia and Ubiquitous Engineering. FutureTech 2017, pp. 61-66, 2017. https://doi.org/10.1007%2F978-981-10-5041-1_11

[VMW18] “Introducing VMware Kubernetes Engine”, June 2018: https://cloud.vmware.com/community/2018/06/26/introducing-vmware-kubernetes-engine-vke/

[DOC1] "Run Swarm and Kubernetes Interchangeably" https://www.docker.com/products/orchestration

[MIC17] “Introducing AKS (managed Kubernetes) and Azure Container Registry improvements”, October 2017: https://azure.microsoft.com/en-us/blog/introducing-azure-container-service-aks-managed-kubernetes-and-azure-container-registry-geo-replication/

[AMA17] “Amazon Elastic Container Service for Kubernetes”, November 2017: https://aws.amazon.com/es/blogs/aws/amazon-elastic-container-service-for-kubernetes/

[CNC1] Survey Shows Kubernetes Leading as Orchestration Platform: https://www.cncf.io/blog/2017/06/28/survey-shows-kubernetes-leading-orchestration-platform/

[CNC2] The Cloud Native Computing Foundation (CNCF) was founded in 2015 by companies like Google, IBM, Docker as one of the Linux Foundation Collaborative Project, and it’s dedicated to make “cloud native computing universal and sustainable”: https://www.cncf.io/

[SIS2017] The 2017 Docker Usage Report: https://sysdig.com/blog/sysdig-docker-usage-report-2017/

[CNC2018] “Kubernetes Is First CNCF Project To Graduate”, March 2018: https://www.cncf.io/blog/2018/03/06/kubernetes-first-cncf-project-graduate/

[Oka11] Y. Okada et al., “Comparisons of Machine Learning Algorithms for Application Identification of Encrypted Traffic,” ICMLA 2011, pp. 358-361. https://doi.org/10.1109/CISDA.2011.5945941

[Mai18-1] H. L. Mai et al., “Towards a security monitoring plane for named data networking and its application against content poisoning attack,” NOMS 2018, pp. 1-9, 2018. https://doi.org/10.1109/NOMS.2018.8406246

[Mai18-2] H. L. Mai et al., “Implementation of content poisoning attack detection and reaction in virtualized NDN networks,” ICIN 2018, pp. 1-3, 2018. https://doi.org/10.1109/ICIN.2018.8401591

[Cisco18] Cisco, “Encrypted Traffic Analysis”, whitepaper, 2018. https://www.cisco.com/c/dam/en/us/solutions/collateral/enterprise-networks/enterprise-network-security/nb-09-encrytd-traf-anlytcs-wp-cte-en.pdf

[Yu15] K. F. Yu, G. S Payer, “Porting Extremely Lightweight Intrusion Detection (ELIDe) to Android”, ARL-TN-0681, US Army Research Laboratory, 2015. http://www.dtic.mil/docs/citations/ADA622509

[Yu17] K. F. Yu, R. E. Harang, “Machine learning in malware traffic classifications,” MILCOM 2017 https://doi.org/10.1109/MILCOM.2017.8170769

[Cui18] Z. Cui et al., “Detection of Malicious Code Variants Based on Deep Learning,” in IEEE Transactions on Industrial Informatics, vol. 14, no. 7, pp. 3187-3196, 2018. https://doi.org/10.1109/TII.2018.2822680

[Meng17] X. Meng et al., “MCSMGS: Malware Classification Model Based on Deep Learning,” International Conference on Cyber-Enabled Distributed Computing and Knowledge Discovery, pp. 272-275, 2017. https://doi.org/10.1109/CyberC.2017.21

[Iera18] C. Ieracitano et al. “Statistical Analysis Driven Optimized Deep Learning System for Intrusion Detection. ” arXiv preprint, 2018. https://arxiv.org/pdf/1808.05633.pdf

[Kas18] Kaspersky Lab, https://www.kaspersky.com/.

[Kas17] Kaspersky Lab, “Machine Learning for Malware Detection”, whitepaper, 2017.

https://media.kaspersky.com/en/enterprise-security/Kaspersky-Lab-Whitepaper-Machine-Learning.pdf

[Graf18] R. Graf, R. King, "Neural network and blockchain based technique for cyber threat intelligence and situational awareness", CYCON 2018, pp. 409-426, 2018. https://doi.org/10.23919/CYCON.2018.8405028

[EDE18] Edelman, "2018 Edelman Trust Barometer," 2018. https://www.edelman.com/trust-barometer/

[OECD17] OECD, "Key Issues for Digital Transformation in the G20," 2017. https://www.oecd.org/g20/key-issues-for-digital-transformation-in-the-g20.pdf

[VER18] Verizon, "2018 Data Breach Investigations Report," 2018. https://www.verizonenterprise.com/resources/reports/rp_DBIR_2018_Report_en_xg.pdf

[Coo18] S. Cook, "Identity theft stats & facts: 2017 – 2018," 2018. https://www.comparitech.com/identity-theft-protection/identity-theft-statistics/

[EU16] European Parliament, "Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation)," 2016. http://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32016R0679

[EC17] European Commission, "Proposal for an ePrivacy Regulation", 2017. https://ec.europa.eu/digital-single-market/en/proposal-eprivacy-regulation

[EU02] European Parliament, "Directive 2002/58/EC of the European Parliament and of the Council of 12 July 2002 concerning the processing of personal data and the protection of privacy in the electronic communications sector (Directive on privacy and electronic communications)," 2002. http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:32002L0058:en:HTML

[BIT] Bitcoin: https://www.bitcoin.com/

[BN+14] J. Bonneau, A. Narayanan, A. Miller, J. Clark, J. A. Kroll, E.W. Felten, "Mixcoin: Anonymity for Bitcoin with Accountable Mixes", FC, 2014. https://doi.org/10.1007/978-3-662-45472-5_31

[BC+14] E. Ben-Sasson, A. Chiesa, C. Garman, M. Green, I. Miers, E. Tromer, M. Virza, "Zerocash: Decentralized Anonymous Payments from Bitcoin," IEEE SP, 2014. http://zerocash-project.org/

[MON] Monero: https://getmonero.org/

[Nak08] S. Nakamoto, "Bitcoin: A Peer-to-PeerElectronic Cash System," 2008. https://bitcoin.org/bitcoin.pdf

[CK+12] J. Camenisch, I. Krontiris, A. Lehmann, G. Neven, C. Paquin, K. Rannenberg, H. Zwingelberg, "H2.1 – ABC4Trust Architecture for Developers," 2012. https://abc4trust.eu/download/ABC4Trust-H2.1-Architecture-for-Developers.pdf

[CL02] J. Camenisch, A. Lysyanskaya, “A signature scheme with efficient protocols,” in Proceedings of SNC'02, pp. 268-289, 2002. https://doi.org/10.1007/3-540-36413-7_20

[BCC04] E. Brickell, J. Camenisch, L. Chen, “Direct Anonymous Attestation,” in Proceedings of CCS'04, pp. 132-145, 2004. https://doi.org/10.1145/1030083.1030103

[Cha85] D. Chaum, “Security Without Identification: Transaction Systems to Make Big Brother Obsolete,” Communications of the ACM 28(10), pp. 1030-1044, 1985. https://doi.org/10.1145/4372.4373

[CH91] D. Chaum, E. van Heyst, “Group Signatures,” in Proceedings of EUROCRYPT'91, pp. 257-265, 1991. https://www.chaum.com/publications/Group_Signatures.pdf

[CL01] J. Camenisch, A. Lysyanskaya, “An Efficient System for Non-transferable Anonymous Credentials with Optional Anonymity Revocation,” in Proceedings of EUROCRYPT'01, pp. 93-118, 2001. https://doi.org/10.1007/3-540-44987-6_7

[CL04] J. Camenisch, A. Lysyanskaya, “Signature Schemes and Anonymous Credentials from Bilinear Maps,” in Proceedings of CRYPTO'04, pp. 56-72, 2004. https://doi.org/10.1007/978-3-540-28628-8_4

[HS14] C. Hanser, D. Slamanig, “Structure-Preserving Signatures on Equivalence Classes and Their Application to Anonymous Credentials,” in Proceedings of ASIACRYPT'14, pp. 491-511, 2014. https://doi.org/10.1007/978-3-662-45611-8_26

[RVH17] S. Ringers, E. Verheul, J.-H. Hoepman, “An efficient self-blindable attribute-based credential scheme,” in Proceedings of FC'17, pp. 3-20, 2017. https://doi.org/10.1007%2F978-3-319-70972-7_1

[CNR12] J. Camenisch, G. Neven, M. Rueckert, “Fully Anonymous Attribute Tokens from Lattices,” in Proceedings of SCN'12, pp. 57-75, 2012. https://doi.org/10.1007/978-3-642-32928-9_4

[BCN17] C. Boschini, J. Camenisch, G. Neven, “Relaxed Lattice-Based Signatures with Short Zero-Knowledge Proofs,” ePrint Archive, 2017. https://eprint.iacr.org/2017/1123.pdf

[BK17] R. El Bansarkhani, A. El Kaafarani, “Direct Anonymous Attestation from Lattices,” ePrint Archive, 2017. https://eprint.iacr.org/2017/1022

[PRI] PRIME: http://www.prime-project.eu/

[PL] PrimeLife: http://www.primelife.eu/

[ABC] ABC4Trust: https://abc4trust.eu/

[MAT] MATTHEW: http://matthew-project.eu/

[CRE] CREDENTIAL: https://credential.eu/

[IBM02] IBM, “Identity Mixer (idemix)”, 2002. https://www.zurich.ibm.com/identity_mixer/

[Mic12] Microsoft, “U-Prove”, 2012. https://www.microsoft.com/en-us/research/project/u-prove/

[GL05] S. Goldwasser, Y. Lindell, “Secure Multi-Party Computation Without Agreement,” Journal of Cryptology 18(3), pages 247-287, 2005. https://doi.org/10.1007/s00145-005-0319-z

[BC+09] P. Bogetoft, D. Christensen, I. Damgard, M. Geisler, T. Jakobsen, M. Kroigaard, J. D. Nielsen, J. B. Nielsen, K. Nielsen, J. Pagter, M. Schwartzbach, “Secure multiparty computation goes live,” Financial Cryptography and Data Security, pages 325–343, 2009.
https://doi.org/10.1007/978-3-642-03549-4_20

[Kre17] B. Kreuter, “Secure multiparty computation at Google,” RWC'17, 2017.
https://www.youtube.com/watch?v=ee7oRsDnNNc

[BSW11] D. Boneh, A. Sahai, B. Waters, “Functional encryption: Definitions and challenges,” in Proceedings of TCC’11, pp. 253-273, 2011. https://doi.org/10.1007/978-3-642-19571-6_16

[KSW08] J. Katz, A. Sahai, B. Waters, “Predicate Encryption Supporting Disjunctions, polynomial equations, and inner products,” in Proceedings of EUROCRYPT’08, pp. 146-162, 2008.
https://doi.org/10.1007/978-3-540-78967-3_9

[GP+06] V. Goyal, O. Pandey, A. Sahai, B. Waters, “Attribute-Based Encryption for Fine-Grained Access Control of Encrypted Data,” in Proceedings of CCS’06, 2006.
https://doi.org/10.1145/1180405.1180418

[BF01] D. Boneh, M. Franklin, “Identity-Based Encryption from the Weil Pairing,” in Proceedings of CRYPTO’01, pp. 213-229, 2001.
https://doi.org/10.1007/3-540-44647-8_13

[Gen09] C. Gentry, “A fully homomorphic encryption scheme,” Ph.D. Dissertation, Stanford University, 2009.
https://crypto.stanford.edu/craig/craig-thesis.pdf

[LCM16] J. Liu, L. Chen, S. Mesnager, “Partially homomorphic encryption schemes over finite fields,” in Proceedings of SPACE'16, pp. 109-123, 2016.
https://doi.org/10.1007/978-3-319-49445-6_6

[DPZ12] I. Damgard, V. Pastro, N. P. Smart, S. Zakarias, “ Multiparty Computation from Somewhat Homomorphic Encryption,” in Proceedings of CRYPTO'12, pp. 643-662, 2012.
https://doi.org/10.1007/978-3-642-32009-5_38

[GJK18] D. Grzonka, A. Jakóbik, J. Kołodziej, S. Pllana,” Using a multi-agent system and artificial intelligence for monitoring and improving the cloud performance and security”, in Future Generation Computer Systems Volume 86, September 2018, Pages 1106-1117.
https://doi.org/10.1016/j.future.2017.05.046

[DMD17] S. Dam, G. Mandal, K. Dasgupta, P. Dutta, “An Ant-Colony-Based Meta-Heuristic Approach for Load Balancing in Cloud Computing, book chapter in Applied Computational Intelligence and Soft Computing in Engineering”, pages 29, September 2017, DOI: 10.4018/978-1-5225-3129-6.ch009.

[SV17] A. Suresh, R. Varatharajan, “Competent resource provisioning and distribution techniques for cloud computing environment” R. Cluster Comput (2017).
https://doi-org.recursos.biblioteca.upc.edu/10.1007/s10586-017-1293-6

[GSR17] M. Ghobaei-Arani, M. Shamsi, A. A. Rahmanian, "An efficient approach for improving virtual machine placement in cloud computing environment" Journal of Experimental & Theoretical Artificial Intelligence, Volume 29, 2017 - Issue 6. https://doi-org.recursos.biblioteca.upc.edu/10.1080/0952813X.2017.1310308

[OBC17] P. Östbergz, J. Byrnez, P. Casari, P. Eardley, A. Fernandez Anta, J. Forsman, J. Kennedy, T. Le Ducz, M. Noya Mariño, R. Loomba, M. A. López Peña, J. López Veiga, T. Lynnz, V. Mancuso, S. Svorobejz, A. Torneus, S. Wesnerx, P. Willis, J. Domaschka, "Reliable capacity provisioning for distributed cloud/edge/fog computing applications" in Proceedings of the 2017 European Conference on Networks and Communications (EuCNC), Oulu, Finlad June 2017. https://doi-org.recursos.biblioteca.upc.edu/10.1109/EuCNC.2017.7980667

[AA18] R. G. Aryal, J. Altmann, "Dynamic application deployment in federations of clouds and edge resources using a multiobjective optimization AI algorithm" in Proceedings of the 2018 Third International Conference on Fog and Mobile Edge Computing (FMEC), Barcelona, Spain, April 2018.

[DCB18] J. Dizdarevic, F. Carpio, M. Bensalem, A. Jukan, "Enhancing Service Management Systems with Machine Learning in Fog-to-Cloud Networks" in Proceedings of the Euro-Par 2nd F2C-DP Workshop on Fog-to-Cloud Distributed Processing, Torino, Italy, 2018.

[SGM18] S. Sengupta, J. Garcia , X. Masip-Bruin, "An Architecture for Resources Management in a Fog-to-Cloud Framework" in Proceedings of the Euro-Par 2nd F2C-DP Workshop on Fog-to-Cloud Distributed Processing, Torino, Italy, 2018.

[JD18] S. Jošilo and G. Dán, "Poster abstract: Decentralized fog computing resource management for offloading of periodic tasks," in IEEE INFOCOM 2018 - IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS), 2018, pp. 1–2.

[WZC17] W. Zhang, Z. Zhang, and H.-C. Chao, "Cooperative fog computing for dealing with big data in the internet of vehicles: Architecture and hierarchical resource management," IEEE Commun. Mag., vol. 55, no. 12, pp. 60–67, 2017.

[WSM18] C. Wöbker, A. Seitz, H. Mueller, and B. Bruegge, "Fogernetes: Deployment and management of fog computing applications," in NOMS 2018-2018 IEEE/IFIP Network Operations and Management Symposium, 2018.

[TD17] M. Taneja and A. Davy, "Resource aware placement of IoT application modules in Fog-Cloud Computing Paradigm," in Integrated Network and Service Management (IM), 2017 IFIP/IEEE Symposium on, 2017, pp. 1222–1228.

[FMG18] S. Filiposka, A. Mishev, and K. Gilly, "Community-based allocation and migration strategies for fog computing," in Wireless Communications and Networking Conference (WCNC), 2018 IEEE, 2018, pp. 1–6.

[KSM18] J. Klaimi, S.-M. Senouci, and M.-A. Messous, "Theoretical Game Approach for Mobile Users Resource Management in a Vehicular Fog Computing

Environment,” in 2018 14th International Wireless Communications & Mobile Computing Conference (IWCMC), 2018, pp. 452–457.

[ZZN17] H. Zhang, Y. Zhang, Y. Gu, D. Niyato, and Z. Han, “A hierarchical game framework for resource management in fog computing,” IEEE Commun. Mag., vol. 55, no. 8, pp. 52–57, 2017.

[VPK15] A. Verma, L. Pedrosa, M. R. Korupolu, D. Oppenheimer, E. Tune, J. Wilkes, “Large-scale cluster management at Google with Borg”, Proceedings of the European Conference on Computer Systems (EuroSys), ACM, Bordeaux, France (2015)

[GHC18] Gartner’s Hype Curve, 2018. https://www.gartner.com/smarterwithgartner/5-trends-emerge-in-gartner-hype-cycle-for-emerging-technologies-2018/

[NuvlaBox] NuvlaBox - Secure and Intelligent Edge Computing Solution. https://sixsq.com/products-and-services/nuvlabox/overview

[Akraino] Akraino Edge Stack. https://www.akraino.org/

[AkrainoWiki] Akraino Edge Stack Wiki. https://wiki.akraino.org/

[MAHMUD] Redowan Mahmud and Rajkumar Buyya. 2016. Fog Computing: A Taxonomy, Survey and Future Directions. Distrib. Parallel, Clust. Comput. , November (2016), 1–28. DOI:http://dx.doi.org/10.1002/9780470773239.ch10

[BROGI] Antonio Brogi, Stefano Forti, and Ahmad Ibrahim. 2017. How to Best Deploy Your Fog Applications, Probably. Proc. - 2017 IEEE 1st Int. Conf. Fog Edge Comput. ICFEC 2017 (2017), 105–114. DOI:http://dx.doi.org/10.1109/ICFEC.2017.8

[TOCZE] Klervie Toczé and Simin Nadjm-Tehrani. 2018. A Taxonomy for Management and Optimization of Multiple Resources in Edge Computing. (2018), 1–13.

[McKinsey2015] McKinsey Global Institute, The Internet of Things: mapping the value beyond the hype, june 2015

[Qualcomm2017] Qualcomm, How to bring true interoperability to the Internet of Things, 2017)

[Forrester2016] Forrester, Simplifying the complexity of IoT, june 2016

[EY2018] Ernst Young, Every business must monetize its IoT data to survive, 2018

[McKinsey2015bis] McKinsey & Company, Industry 4.0 How to navigate digitization of the manufacturing sector

[HPE2016] Hewlett Packard Enterprise, How to Get the Most From the Internet of Things, 2016

[BCG2017] Boston Consulting Group, Winning in IoT: It’s All About the Business Processes, 2017

[Deloitte2016] Deloitte, Internet of Things Ecosystem: Unlocking the Business Value of Connected Devices, 2016

[PoliMI2016] Politecnico di Milano, Big Data Analytics & Business Intelligence Observatory: final report 2016

[Gartner2016] Gartner Magic Quadrant for Data Warehouse and Data Management for Analytics 2016

[Deloitte2018] Deloitte, IoT Innovation Report, 2018

[Gartner2017] Gartner, Leading the IoT, 2017

[McKinsey2017] McKinsey, Fueling growth through data monetization, 2017

[GCC1] Global Cloud Computing Market Analysis & Trends – Industry Forecast to 2025, Research and Markets, https://www.researchandmarkets.com/research/54tvtd/global_cloud

[GVR1] Grand View Research, Bare Metal Cloud Market Analysis by Type (Hardware, Service), By Deployment (Hosted, On-Premise), By Enterprise, By End-use (Advertising, BFSI, Government, Healthcare), By Region, And Segment Forecast, 2018-2025, https://www.grandviewresearch.com/industry-analysis/bare-metal-cloud-market

[HTF1] HTF Market Intelligence, Global Cloud Analytics Market Size, Status and Forecast 2025, https://www.htfmarketreport.com/sample-report/1139586-global-cloud-analytics-market-3

[INK1] Inkwood Research, Global Cloud Market Forecast 2017-2025, https://www.inkwoodresearch.com/reports/global-cloud-storage-market-2017-2025/

[WBR1] Wikibon Research, Cloud Computing (2015-2025), https://wikibon.com/wp-content/uploads/Wikibon-BGracely-Cloud-Computing-Nov-20152.pdf

[GR1] Gartner Forecasts Worldwide Public Cloud Revenue to Grow 21.4 Percent in 2018, https://www.gartner.com/newsroom/id/3871416

[INFR1] Information Age, Why the term 'cloud' could be obsolete by 2025, https://www.information-age.com/term-cloud-123473665/

[GVR2] Grand View Research, Edge Computing Market $3.24 Billion By 2025 | CAGR: 41%, https://www.grandviewresearch.com/press-release/global-edge-computing-market

[RMR1] Research and Markets, Global Edge Analytics Market: (Focus on Components (Hardware, Software & Services), Country Analysis, End Users, Competitive Landscape, Market Share Analysis, and Region Specific Information) – Analysis and Forecast (2018-2025), https://www.researchandmarkets.com/research/6qr33t/global_edge?w=12

[STA1] Statista, Market size for edge computing in the U.S. 2017-2025, by segment, https://www.statista.com/statistics/909308/united-states-edge-computing-market-size-segment/

[EUG1] EU GDPR, https://eugdpr.org/

[GVR3] Grand View Research, Fog Computing Market Analysis by Solution, By Hardware (Gateways, Routers & Switches, Sensors), By Application (Connected Vehicles, Smart Grids, Smart Cities, Connected Healthcare), By Region, & Segment Forecast, 2018-2025, https://www.grandviewresearch.com/industry-analysis/fog-computing-market

[451R1] 451 Research, Size and Impact of Fog Computing Market, A report on research commissioned by OpenFog Consortium, https://www.openfogconsortium.org/wp-content/uploads/451-Research-report-on-5-year-Market-Sizing-of-Fog-Oct-2017.pdf

[FORR1] Forrester, Predictions 2017: Artificial Intelligence Will Drive The Insights Revolution,https://go.forrester.com/wp-content/uploads/Forrester_Predictions_2017_-Artificial_Intelligence_Will_Drive_The_Insights_Revolution.pdf

[MMR1] Markets and Markets, Artificial Intelligence Market by Offering (Hardware, Software, Services), Technology (Machine Learning, Natural Language Processing, Context-Aware Computing, Computer Vision), End-User Industry, and Geography – Global Forecast to 2025, https://www.marketsandmarkets.com/Market-Reports/artificial-intelligence-market-74851580.html

[GVR4] Grand View Research, Artificial Intelligence Market Analysis By Solution (Hardware, Software, Services), By Technology (Deep Learning, Machine Learning, Natural Language Processing, Machine Vision), By End-use, By Region, and Segment Forecasts, 2018-2025, https://www.grandviewresearch.com/industry-analysis/artificial-intelligence-ai-market/request

[GVR5] Grand View Research, Deep Learning Market Analysis By Solution, By Hardware (CPU, GPU, FPGA, ASIC), By Service, By Application (Image Recognition, Voice Recognition, Video Surveillance), By End-use, By Region, And Segment Forecasts, 2018-2025, https://www.grandviewresearch.com/industry-analysis/deep-learning-market

[ALR1] Allied Market Research, Deep Learning Chip Market by Chip Type (GPU, ASIC, FPGA, CPU, and Others). Technology (System-on-chip, System-in-package, Multi-chip module, and Others), and Industry Vertical (Media & Advertising, BFSI, IT & Telecom, Retail, Healthcare, Automotive & Transportation, and Others) – Global Opportunity Analysis and Industry Forecast, 2018-2025, https://www.alliedmarketresearch.com/deep-learning-chip-market

[GAR5] Gartner, Use Three Lenses to View Digital Business Opportunity, https://www.gartner.com/smarterwithgartner/use-three-lenses-to-view-digital- business-opportunity/

[GRWTH17] Market IoT Pulse Report, https://growthenabler.com/flipbook/pdf/IOT%20Report.pdf

[RECAP] RECAP Project, http://www.recap-project.eu