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

D2.3 Tracking Scientific, Technology and Business Trends (Version 3)

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:	M33
Submission Date:	30.09.2019
Version:	2.0
Status	Final
Author(s):	Michael J. McGrath, John Kennedy (INTEL), Jens Jensen, Shirley Crompton (STFC), Anna Queralt, Daniele Lezzi (BSC), Jasenka Dizdarevic (TUBS), Sašo Stanovnik, Aleš Černivec, Manca Bizjak, Jolanda Modic (XLAB), Roi Sucasas Font, Lara Lopez Muniz (ATOS), Glauco Mancini, Antonio Salis (Engineering), Eva Marin Tordera, Xavier Masip (UPC), Denis Guilhot (WSL), Cristóvão Cordeiro (SIXSQ)
Reviewer(s)	Ana Juan Ferrer (ATOS) Xavi Masip (UPC)

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	1/07/2019	Initial ToC and doc structure	Michael McGrath/John Kennedy (Intel)
0.2	26/07/2019	Integrated contributions from Intel, BSC, SixSq, UPC, ATOS, WOS, ENG, XLAB.	Michael McGrath (Intel)
0.3	2/09/2019	STFC contributions and overall content review.	Jens Jensen (STFC)
0.4	3/09/2019	Initial draft version	Michael McGrath (Intel)
0.5	6/09/2019	Document Review	Xavi Masip (UPC)
0.6	9/09/2019	Completion of review edits	Michael McGrath (Intel)
0.7	10/09/2019	References Fixes, Additional content in sections 2.4.1, 2.4.2 and 2.4.3	Sašo Stanovnik (XLAB), Michael McGrath (Intel)
0.8	11/09/2019	Reference fix and minor edits	Michael McGrath (Intel)
0.9	13/09/2019	Document Review	Ana Juan Ferrer (Atos)
1.0	16/09/2019	Updates to sections 2.2.2, 2.7, 3.6, 4.6	Daniele Lezzi (BSC), Xavier Masip (UPC), Denis Guilhot (WSL), Antonio Salis (ENG)
1.1	16/09/2019	Reference updates, section 5.0 updates and general content edits	Michael McGrath (Intel)
1.2	17/09/2019	Updates to sections 2.7, 4.6, 5.0	Lara Lopez Muñiz (Atos), Jens Jensen (STFC), Sašo Stanovnik (XLAB), Michael McGrath (Intel)
1.3	18/09/2019	Minor text edits	Michael McGrath (Intel)
2.0	30/09/2019	Quality check performed. Document ready for submission	María Teresa García (ATOS), Michael McGrath (Intel)

Table of Contents

List of Figures

List of Tables

Executive Summary

The objective of this deliverable is to provide a final update on the 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 on the trends previously documented in deliverable D2.2 [1] which was submitted in M21. As per the previous deliverable, each chapter concludes with a “key takeaways” section summarising the key points of interest, enabling readers to understand project’s main priorities. The key takeaways also outline the project’s perspective on potential further research areas in relation to Fog to Cloud technologies and capabilities. This is the third and final version of the deliverable and is aligned to iteration 2 (IT-2) of the project.

1. Introduction

1.1. Purpose

The term Fog-to-Cloud (F2C) [1] continues to refer to the resources created by merging cloud and fog computing, which in turn creates the need for new, open and coordinated management ecosystems. This is where mF2C provides value with its management framework which is designed to be open, secure, decentralised and provides multi-stakeholder support. Other key features include novel programming models, privacy and security, data storage techniques, service creation, brokerage solutions, service level agreement (SLA) policies, and resource orchestration methods. The mF2C project is building a proof of concept (PoC) system and platform based on an innovative distributed system architecture validated through real world use cases.

Workpackage 2 (WP2) is focused on studying the current state-of-the-art in fog, cloud, network and IT infrastructure technologies, with the goal of identifying technologies relevant for the deployment of the mF2C management framework, i.e. sensors, smart end-devices, connectivity, and advanced cloud services. Previous versions of this deliverable D2.1 [2] and D2.2 [3] provided an overview of the scientific, technology and business trends in fog computing which were relevant to the project at the time of publishing. The first version of the deliverable (D2.1) was aligned to iteration 1 (IT-1) of the project, the second version (D2.2) was aligned to the preliminary stages of iteration 2 (IT-2), while this version (D2.3) is the final project report in IT-2. Given the short length of time between the previous version of the deliverable (12 months), the technology trends described in D2.2 remain both valid and relevant for this version of the deliverable. The drive towards Fog/Edge computing continues on its current trajectory with an on-going focus on the massive amounts of IoT generated data, an ability to aggregate and reduce data at source in order to ensure real-time decision making, data anonymization, privacy protection and increased autonomy.

1.2. Document Overview

Section 2 reviews the scientific trends of Fog and Cloud computing, beginning with an assessment of the contributions relating to service and resource management emphasising the need for decentralised and hierarchical solutions in order to meet the challenges generated by fog computing. While the convergence between High Performance Computing (HPC) and cloud and big data technologies remain valid, attention has progressed to joint management of communications and computation resources through the use of Artificial Intelligence (AI), Deep Reinforcement Learning (DRL) and Blockchain for edge resource management.

With the emergence of fog/edge computing, there is increasing attention on the convergence of communications and computing, and solutions to provide joint management of both resource types. Recent work has focused on approaches to provide both improved system performance and efficient resource utilisation which are challenges relevant in the context of mF2C. In a similar vein both AI and DRL approaches have been applied to fog radio access networks in order to deal with the dynamics of edge compute environments. While blockchain has traditionally been associated with recording cryptocurrency transactions in a secure manner, recent research has introduced the concept of a mobile blockchain where Multi-access Edge Computing (MEC) is leveraged to execute computation-heavy proof-of-work (PoW) tasks offloaded by mobile users. These approaches have relevance to mF2C due to issues related to the high device mobility, their limited energy budgets and the impact of the network on the performance of the entire framework.

A key trend in the data domain is the development of software architecture for extreme-scale data analytics by taking advantage of the resources provided by the edge-to-cloud continuum. A variety of data management solutions for different use case domains such as the management of visual data sets in media applications have emerged. These alternative approaches are compared and contrasted against the one used in mF2C. Industry best practise is moving towards converged and hyper-

converged data centre infrastructures. The parallels to the technology challenges relating to this trend is reflected upon in the context of the challenges faced by mF2C and solution requirement similarities.

Problems related to task scheduling and offloading mechanisms due to issues related to device mobility, limitations in device energy budgets and the impact of the network (latency, monetary cost, and bandwidth) on the performance of the entire framework are examined. Capabilities such as application fragmentation in order to offload units of computation to available resources, the scheduling model and the management of parallelism are considered.

The use of machine learning based approaches for security applications has attracted a lot of interest in the research community over recent years. However, there are significant challenges remaining in order to realise fully automated operations without human supervision. Research into approaches which can successful bridge this gap remains an area of active investigation. Also in the security domain data protection remains an on-going battle. Unfortunately approaches to date have not worked well. In order to address these shortcomings, privacy engineering has emerged as a new field of research and practice. This new field aims to bridge the legal, technical and engineering approaches to support engineers in systematically identifying and addressing privacy and data protection concerns during the development lifecycle. Finally, research initiatives dealing with the resource management in the fog and cloud, considering QoS targets, are updated. Section 2 finishes by providing an update on the current state of the art in AI and Machine Learning (ML) applied to fog and cloud computing in order to improve performance.

Section 3 reviews the key technology trends of relevance to mF2C, evaluating different tools and platforms currently available that enable the management of features such as storage, compute, machine instances, and containers. In the previous versions of this deliverable [D2.1 and D2.2], we reviewed the technological trends with respect to the management of Cloud, Fog and IoT devices. While these technical areas of focus still remain valid, this section has been updated to reflect the latest trends in Cloud and IoT management. Specifically, the key open-source cloud and IoT management platforms and tools, which are starting to gain traction, are reviewed. In the Cloud domain this includes tools such as Apache CloudStack and Cloudify, while in the IoT management domains the tools include KAA, DeviceHive and DSA.

In addition to the data management technologies for edge to cloud environments, various database vendors have started to address the needs of the edge-to-cloud scenario by expanding upon their previous focus on the high performance capabilities. In edge-to-cloud solution offerings the focus is on providing enhanced analytics support through extreme scalability, high performance, and support for specific data models related to edge data sources, such as sensors. In some cases, the database system is specifically designed for a particular data model, such as time series databases, e.g. InfluxDB. Other approaches are based around a flexible data model that can also store data in other structures e.g. MongoDB. Recent approaches which address the challenge of data persistence on computationally limited edge devices while providing high performance capabilities are reviewed. In programming models, the significant trend centres on how to implement AI applications which can run on computation-capable edge devices. The key challenge now being addressed is the provisioning of approaches which abstract application partitioning and AI model inference across the lifecycle of deployment, communications management, and failure recovery for software developers.

The dominance of Kubernetes and related solutions for container orchestration continues to grow. There is increasing traction amongst the major cloud solution providers with tools such as Azure AKS, IBM IKS, Google GKE, Red Hat OpenShift and Amazon EKS now available to customers. Growing adoption is also increasing the market value for container based solutions which is expected to more than double over the next five years to reach USD 4.98 billion by 2023[1].

[1] MarketsandMarkets - https://www.marketsandmarkets.com/Market-Reports/application-container-market-182079587.html

A key development in the standards domain was the amalgamation of the OpenFog Consortium and the Industrial Internet Consortium in January 2019, and their subsequent publication of the IoT reference architecture in June 2019. In the previous deliverable versions the activities of the ISO/IEC subcommittees with respect to IoT and Cloud were reported. This work is now being supported by activities from other organisations in the broader context of the technology building blocks of Fog-to-Cloud solution architectures. Activities of note include NIST's recent recommendations on a Fog Compute Model, the open connectivity foundation's reference connectivity models and ETSI's work in relation to Multi-access Edge Computing (MEC). Finally, section 3 concludes by reviewing recent technology trends in edge computing. There is growing focus by edge computing providers and users on solutions that can provide AI, high reliability and fast decision making at the edge, whilst maintaining data anonymization and high portability characteristics that are often required in edge computing.

Section 4 reflects on the latest cloud, fog and IoT trends in the context of evolving 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. The section reviews the continued expansion in cloud computing market with a Compound Annual Growth Rate (CAGR) of 18% and a predicted market size of USD623.3 billion by 2023[2]. Increasing adoption of IoT technologies, the growth of smart cities and the penetration of e-commerce are driving market growth. Growth in IoT market is acting as a key source of big data with global IoT devices expected to generate 90 zettabytes[3] of data by 2025. Section 4 also examines how big data is impacting business, the economy and the key technology challenges associated with effectively leveraging big data.

Business continue to worry about an ever increase range of security threads with a corresponding increase in overhead. In order to address this continuously growing overhead, efforts are focused on improved "solutions" for automated threat detection and remediation. However, there is no universal or no "one size fits all" solution for the horizon for businesses. Humans will continue to play a critical role in both mitigation of security threats by improved training and awareness programs. From a response perspective, humans will be necessary for the foreseeable future to ultimately assess an incident. Section 4 concludes by examining digital businesses and their impact on more traditional businesses which are having to rapidly evolve in response to market disruption with a significant focus on innovation rather than traditional success measures such as improved productivity and customer experience.

[2] MarketsandMarkets - https://www.marketsandmarkets.com/Market-Reports/cloud-computing-market-234.html

[3] A zettabyte is 1,000 exabytes, or one million petabytes. 90 zettabytes is $9x10^{22}$ bytes.

1.3. Glossary of Acronyms

Acronym	Definition
AI	Artificial Intelligence
API	Application Program Interface
B2B	Business-to-Business
B2C	Business-to-Consumer
CapEx	Capital Expenditure
CAGR	Compound Annual Growth Rate
CRDT	Conflict-Free Replicated Data Types
CT	Communication Technology
DL	Deep Learning
DSA	Distributed Services Architecture
DRL	Deep Reinforcement Learning
DRLA	Deep Reinforcement Learning based Allocation
ECN	Edge Computer Networks
ETSI	European Telecommunications Standards Institute
F2C	Fog-to-Cloud
FaaS	Function as a Service
FogMNW	Fog computing enabled mobile communication network
F-RANs	Fog radio access networks
GDPR	General Data Protection Regulation
HPC	High Performance Computing
IoE	Internet of Everything
IIC	Industrial Internet Consortium
IP	Internet Protocol
IoT	Internet of Things
JTC	Joint Technical Committee
LPWAN	Low-power wide-area network
M2H	Machine-to-Humans
M2M	Machine-to-Machine
MAC	Media Access Control
MEC	Multi-access Edge Computing
MFA	Multifactor authentication
ML	Machine Learning
MQTT	Message Queuing Telemetry Transport
NFV	Network Function Virtualisation
NIST	National Institute of Standards and Technology
OCF	Open Connectivity Foundation
OpEx	Operational Expenditure
OT	Operational Technology
P2P	Peer-to-Peer
PbD	Privacy by Design
PET	Privacy-Enhancing Technologies
PoW	Proof-of-work
PSO	Particle Swarm Optimization
QoS	Quality of Service
QUIC	Quick UDP Internet Connections
RAS	Reliability, Availability, and Serviceability
SEDS	Self-Evolving Detection Systems

SLA	Service Level Agreement
SDN	Software Defined Networks
SSL	Secure Socket Layer
TET	Transparency-Enhancing Technologies
TLS	Transport Layer Security
TOSCA	Topology and Orchestration Specification for Cloud Applications
VDMS	Visual Data Management System
VM	Virtual Machine
VNC	Virtual Network Computing
VPN	virtual private network
UDP	User Datagram Protocol

Table 1-1. Acronyms

2. Scientific Trends

In this section, an overview of the scientific trends relevant in the context of the mF2C project is provided. This section has been updated from the previous deliverable version [D2.2] which was submitted in M21 to reflect changes of significance in the ensuing 12 month interval. However, it should be noted that many trends that were relevant at the time of the previous deliverable version submission remain the same. The section includes the scientific trends in the areas of service management, resource management, end-devices, HPC programming models trends, science applications, security, the convergence of AI and computing. These are areas where significant advances have observed and therefore form a key part of the updates presented in this deliverable.

2.1. Service Management, Resource Management, End-devices

In D2.2, we reported that a significant number of recent works have promoted the use of decentralised and hierarchical approaches to deal with service/resource management in fog and edge computing environments. Specifically, the trend of using game theoretical concepts remains relevant. In addition, the following trends have recently emerged in the literature:

2.1.1. Joint Management of Communication and Computation resources

With the advent of fog/edge computing, there is increasing attention on the *Convergence of communications and computing,* and as a consequence the joint management of both resource types. Zhou et al. [4] present a fog computing enabled mobile communication network (FogMNW) architecture. They demonstrate that the systematic management of communication and computing resources improves system performance and increases the efficiency of compute resources. A similar joint management approach is also taken in [5], where the authors simultaneously optimise the task assignment and the radio / computation resource allocation in order to meet the low latency requirements of MEC applications. Jošilo and Gyorgy [6] consider a scenario where the MEC operator jointly manages wireless and computation resources following two policies; the first one aims to minimise costs, while the second one ensures a time-fair resource allocation. Numerical results show that using the first approach yields lower task completion times in comparison to the second approach.

2.1.2. AI/DRL Based Resource Management

A number of studies have recognised the potential of AI and in particular deep reinforcement learning (DRL) in highly-dynamic fog/edge environments. For instance, in [7], the authors designed a DRL-based edge resource management scheme to address the requirements of Internet of Vehicles applications. Sun et al. [8] employed a DRL approach that combines both communication mode selection and resource management in fog radio access networks (F-RANs). Zeng et al. [9] propose a DRL-based service migration scheme to minimise system operational cost for highly dynamic edge computing environments.

2.1.3. Blockchain and Edge Resource Management

A series of recent studies have combined the use of blockchain with edge computing in order to propose solutions for resource management in the resulting architecture. For example, Xiong et al. [10], introduce the concept of mobile blockchain where MEC is leveraged to execute computation-heavy proof-of-work (PoW) tasks offloaded by mobile users. A Stackelberg game-theoretical model is then proposed to reach an optimal edge resource management policy in the mobile blockchain. In a similar context, Xiong et al [11] propose a two-stage Stackelberg game for mobile blockchain-driven edge resource management. Specifically, the authors consider both the maximisation of the edge service providers' profit together with the individual utilities of the miners. Luong et al. [12] present an auction scheme based on deep learning to address the edge resource allocation problem in a mobile blockchain.

2.2. Scientific Trends Emerging from the HPC Domain

This section highlights the convergence between High Performance Computing (HPC) and cloud and big data technologies with a specific focus on the data management and programming models trends.

2.2.1. Data Management Trends

The collaborative use of edge resources to improve resource management and bandwidth utilisation, as previously reported in D2.2 [3], continues to attract attention. In addition, new research projects led by various HPC institutions have been initiated during the course of 2019 including the ELASTIC project [13], which has the goal of developing a software architecture for extreme-scale data analytics by taking advantage of the resources provided by the edge-to-cloud continuum. Here the data management functionality plays an important role in providing applications with seamless and efficient access to data regardless of its location in the edge-to-cloud hierarchy.

New research results have been published over the last 12 months, including those in emerging workshops in flagship HPC conferences such as USENIX HotEdge, or in distributed systems conferences and journals. Research has focused on particular domains, such as visual data, or on more general purpose areas such as replica placement and synchronisation. The following summarises the most relevant research outputs and highlights the key differences with the mF2C data management solution.

Altarawneh et al. [14], propose a software architecture with two main components: a streaming analytics framework, and a Visual Data Management System (VDMS) [15], which supports machine learning and analytics workloads utilising visual data. While the system has been instantiated for a retail analytics specific use case, it has been designed to provide a general visual fog framework for ad-hoc video analytics, which could be applied to other application domains such as smart cities or agriculture, assuming visual datasets supported by VDMS.

With a different focus, but still within the media domain, Elgazar et al. [16] focuses on solving the problem of limited storage on small-scale edge devices. The authors propose a smart media compression solution for distributed storage systems that dynamically adjusts compression parameters, in order to reduce the amount of unnecessary compression, thus reducing energy consumption while providing smaller user file access delays. This paper proposes an optimisation that can be applied to media content stored on systems such as EdgeStore [17], a distributed storage solution that offloads files to household devices such as desktops, tablets, or mobile phones, thus taking advantage of their available storage capacity. This kind of system is focused on file storage, and is therefore not be suitable to support the data management functionalities in mF2C, which require a fine-grained data management to support storage and querying of data items in a database-like fashion.

A database approach is described by Meiklejohn et al. in [18] to solve consistency and synchronisation problems between replicas in edge environments. Data structures provided by the data store are Conflict-Free Replicated Data Types (CRDT) [19], which ensure convergence in the event of concurrent modifications at multiple locations. CRDTs provide interesting properties in those cases where eventual consistency suffices, but cannot guarantee the consistency level required by some applications, and in particular to manage the mF2C platform.

PathStore [20] also deals with consistency problems at the edge. Here replicas are created on demand by applications, and the solution guarantees eventual consistency. As in the previous case, it does not provide strong consistency guarantees between replicas, but according to the authors the system could be extended in the future to provide stronger consistency requirements.

Other papers such as [21] deal with replica placement in edge systems. In this case, an algorithm for the creation and removal of replicas guided by continuous monitoring of data requests is proposed. This algorithm optimises data access, but the required monitoring increases the use of resources in

edge devices. Also, as it is based on estimations, unnecessary replicas can be created in some cases. The data management solution for mF2C leverages knowledge about the hierarchy and requests from policies and execution runtime in order to place the data only where it is needed. Additionally, this proposal deals with data in the form of files, which is not appropriate for mF2C.

Also regarding file replica placement in edge environments, the solution outlined in [22] is built on top of a previous proposal dealing only with placement and not replication [23]. In this previous work, data was placed as close as possible to IoT devices, and now has been extended with creation of replicas to ensure availability. As with the previous approach, the disadvantages for mF2C is that data is managed at file level, and that unnecessary replicas may be created.

A different approach is followed in ElfStore [24] which proposes an edge store for streams of data blocks. It allows users to search on the metadata, since the content of the blocks is opaque as in the case of files. The replica placement is decided by an overlay sitting in reliable fog devices that monitor the edge resources. In contrast to mF2C, where the data management functionality supports the operation of the platform and execution of applications, the goal here is to balance capacity usage and ensure data durability for storage.

Another proposal dealing with the availability and distribution of data at the edge is outlined by Vasconcelos et al. [25]. Here the authors propose a stability function that takes into account how long devices remain in the network in order to decide where to place the data. In mF2C, the suitability of each node to hold replicas at any point in time is dynamically determined by the policies, which select appropriate leader and backup nodes. The data management component reacts to the policies decisions by placing the replicas in selected nodes, and automatically removing them when the policy component revises its decisions.

2.2.2. Programming Models Trends

In deliverable D2.2 [1] the scientific trends on programming models utilised in the development of big data applications in distributed environments were described. In this deliverable, the focus is extended to include problems related to task scheduling and offloading mechanisms. These focus areas are important due to issues relating to device mobility, limitations in device energy budgets and the impact of the network (latency, monetary cost, and bandwidth) on the performance of the entire framework. Here capabilities such as application fragmentation in order to offload units of computation to the appropriate available resources, the scheduling model and the management of parallelism are considered.

CloneCloud [26] offers the developer a thread level granularity mechanism. The most salient point of CloneCloud is its partitioning mechanism that combines static code analysis with dynamic application profiling to select the optimal migration and re-integration points. When a thread reaches a migration point, it suspends, and its state (including virtual state, program counter, registers, and stack) is sent to a synchronised clone. When the migrated thread reaches a re-integration point, it is similarly suspended and sent back to the mobile device. The drawback of this system is that it still requires the developer to manage both threads and application parallelism.

Cuckoo [27] hides the partitioning problem by exploiting a service component of the Android operating system. During the build process, stubs generated to access service components are replaced by invocations to the Cuckoo framework that decides, at run-time, whether to run the service on the local device or a remote implementation. Since the framework only replaces calls, all parallelism must be explicitly managed by the programmer of the service invocations.

ThinkAir [28] provides a mechanism to automatically parallelise the execution of an offloaded method based on ranges of input variables. The main drawback of ThinkAir is that the offloading mechanism works synchronously: an executing thread is suspended until method invocation is performed and its result collected. Thus, any subsequent method invocation is not executed until previous ones are executed even when they could run concurrently.

Mobile Fog [29] is a high level programming model for future Internet applications that are geospatially distributed, large-scale, and latency-sensitive. The goal is to allow applications to dynamically scale based on their workload, using on demand allocated resources in the fog and in the cloud. In Mobile Fog, an application consists of distributed Mobile Fog processes which are mapped onto distributed computing instances in the fog and cloud, as well as various edge devices. This Mobile Fog API does not hide the distribution of the infrastructure to the application and as a result requires a significant programming effort by the application developer.

Recently, more generic framework for machine learning have gained popularity such as Tensorflow [30], Theano [31], Caffe [32], PyTorch [33] and Keras [34]. Most of these environments provide a Python interface, which simplifies their use and supports different types of parallelism, such as multi-GPUs. A Python interface is also provided by the MLlib [35], Spark's machine learning library, which is comprised of various algorithms for classification, regression, collaborative filtering, clustering and decomposition. Another popular machine library in Python is Scikit-learn [36], which provides simple and efficient tools for data mining and data analysis. The framework proposed by mF2C is based on COMPSs which provides more flexibility than the programming frameworks outlined above due to their primary focus on providing big data related solutions. Moreover, they don't address the challenges related to the composition of distributed applications and their deployment and execution on the edge with transparent offloading to cloud platforms depending on load and constraints requirements. In addition, COMPSs applications can run seamlessly, without changes, on traditional batch systems or on clouds using a Function as a Service (FaaS) model. In recent proposals, IoT devices can pre-process generated information themselves and trigger successive computation according to the (FaaS) paradigm [37] [38]. Commercial cloud providers have started to offer fog-like services which also focus on providing ML capabilities at the edge, however they limit the type of functions that can run as FaaS. As public clouds do not fulfil intensive computing use cases, a myriad of open-source serverless frameworks supporting building FaaS environments for on-premise offerings are available. However, they require a set of additional components in order to reach a similar level of functionality to that offered by AWS Lambda, Google Functions or Azure for example. In this sense, O-SCAR [39] is a platform aiming to provide automatic encapsulation of functions in Docker containers, combining elastic Kubernetes back-ends and S3-like data storage APIs.

2.3. Applications in Different Science Areas, Data Centres, Big Data Processing

With industry best practice moving the direction of converged and hyper-converged data centre infrastructure, there are potentially more data sources in the data centre, and closer integration between them. Data centres are of course equipped with sensors – temperature, smoke, fire, voltages, airflow/water flow – and increasingly hardware support for Redfish/Swordfish is improving to the point where devices can actually provide useful data. Last year's (2018) Supercomputing conference hosted a workshop on data centre automation and control (Data-center Automation, Analytics, and Control (DAAC), https://daac-general.nsfcac.org/); and for 2019 and 2020 it is expected the workshops will merge with green computing. While there are already various industry "solutions" for data centre automation, taken together, there should now be enough data to build DAAC solutions on open source offerings, with the aim of providing the features to coordinate the deployment and integration of diverse (virtualised) data centre resources, such as compute, storage, and networking. It would require SLA management (for example, to manage green computing resources or resources for confidential data), resource management and scheduling, landscaping, as well as deployment, monitoring and telemetry acquisition for runtime tasks and data flows. If this sounds familiar, the suggestion is not necessarily to build a hyper-converged data centre on the mF2C platform, but at least to recognise that the problems identified are very similar in nature.

More generally, the concept of "Industry 4.0" is comprised of a large number of diverse devices, from IoT devices at the edge to "big data analytics" capabilities that include cloud and fog computing. Already, users provide feedback and reviews of products, and can even engage with the development

of new products through crowdfunding platforms or social media. Like the data centre example described above, Industry 4.0 can also include sensors in the production environment which can save time and money for the producer, or can add value to the product by asserting securely that the product was produced in a sustainable and ethical environment for example.

2.4. Security Trends

Despite the growing relevance of privacy concerns and the profusion of research results, the development of privacy-friendly ICT systems remains more a matter of craftsmanship rather than one of engineering. The European General Data Protection Regulation (GDPR), mandatory since May 2018, in Article 25 explicitly mentions the concepts of data protection by design and by default. This approach, more often known as Privacy by Design (PbD), requires privacy and data protection to be considered from the start of a project and throughout its lifecycle in order to effectively protect data subjects and their personal data.

Economic operators have more data at their disposal than ever before, so they progressively formulate their strategies, base their business decisions, and create revenues from the data they collect through their services. However, while (societal and economic) opportunities to benefit from data grow, legal frameworks continuously reinforce requirements for responsible data management. Service Providers are therefore in ever increasing need of technologies that enable them to collect, analyse, and share data of their users in order to innovate, optimise, and grow their businesses, while at the same time ensuring them of adherence to ethical business practices and legal compliance of their services with data protection and data trading regulations.

Technological advances in the area of (Industrial) IoT/edge/fog/cloud computing, Big Data, and HPC have enabled the development of significant innovations. However, these technologies enable collection, processing, and sharing growing amounts of personal data, thus posing a major risk to the privacy of people. These advances have also significantly increased the number of connected devices, which rapidly and significantly expands the attack surface of any organisation.

2.4.1.Artificial Intelligence and Machine Learning

Currently ML is applied to three major cybersecurity problems [40]: (1) intrusion detection, (2) malware analysis, (3) spam and phishing detection. Papers [40] [41] list the current solutions and maturity levels of machine learning approaches for different cyber security applications. Apruzzese et al. [40] describes the original taxonomy of ML cyber security approaches and also analyse the main limitations of the various approaches. The remainder of this overview presents the latest advancements in the application of ML-based approaches for tackling challenges associated with intrusion detection in computer systems.

2.4.1.Analysis of Encrypted Traffic

As a means of tackling some aspects of the ever-growing concerns for privacy, we have witnessed increased utilisation of encrypted internet communication over the past decade. Moreover, due to the larger uptake of Secure Socket Layer (SSL)/Transport Layer Security (TLS) and the introduction of new protocols (TLS 1.3, Quick UDP Internet Connections (QUIC)) it is now harder to detect such attacks. Despite the use of secure communication, network intrusion detection (IDS) solutions can still tackle detection of botnets (recognising patterns of bots) and network instances using domain generation algorithms (DGA). Apruzzese et al. [40] shows that methods used to tackle DGA problems are still not sufficiently mature to be used without human supervision.

A framework for deep-learning-based traffic classification is presented in [42]. Classification of encrypted traffic is a hard task due to the lack of representative features within the traffic while in secure mode (using TLS). Few studies have shown successful classification of TLS 1.2 and virtual private network (VPN) traffic in User Datagram Protocol (UDP) mode, but since TLS 1.3 and QUIC provide only

few unencrypted fields within their packets, these methods are not useful in real world environments. Most of ML methods are still not useful in operational networks which utilise BBon protocols [42].

2.4.2. Better Detection and Prevention of (Known and Unknown) Cyber-attacks

In contrast to signature-based network IDS software which dominates the commercial IDS landscape, anomaly-based intrusion detection remains largely an ongoing research topic, promising to tackle the most significant shortcoming of signature-based intrusion detection systems (IDSes), i.e., their inability to detect previously unseen attacks, including 0-day exploits. ML approaches are natural candidates for detecting 0-day vulnerabilities, as samples of 0-day exploits can be associated with new classes of traffic which did not exist in the training set. Several successful demonstrations of anomaly detection, either on hosts or within networks, have been presented in the literature, many of them based on shallow ML or deep learning (DL) approaches.

The abundance of IoT devices and the fact that their security is lagging behind offers an opportunity to hackers, making them desirable targets for inclusion in botnets used for launching large-scale distributed denial-of-service (DDoS) attacks. As the botnets' ability to evade discovery evolves in sophistication, detecting them becomes more challenging requiring new mechanisms for discovery and mitigation. Li et al. [43] analyses the Rustock botnet and propose key characteristics for botnet detection based on domain names analysis for botnets adopting fast-flux.

Despite these successes, research by Abeshu and Chilamkurti [44] demonstrates the superiority of DL-based approaches for detecting distributed attacks with previously unseen or slight mutations of known attack patterns (as is largely the case), compared to traditional ML-based solutions. The work makes its case in a fog-to-things setting, arguing that fog-to-things computing may be the ultimate beneficiary of DL approaches for attack/anomaly detection due to the massive amount of data produced by IoT devices which enable deep models to have superior learn capabilities in comparison to shallow ML algorithms. Priyadarshini and Barik [45] offer another demonstration of a DL model for detecting DDoS attacks based on a Long Short Term Memory algorithm, and its successful application to a fog network. To summarise, when it comes to anomaly detection, the literature has shown (see [44] [45] [46]) that carefully selected DL-based models offer better detection accuracy than conventional ML approaches, in various deployments.

Despite the enormous potential and research efforts over the last number of years, there are few implementations of anomaly-based IDSes by commercial vendors [47] employed in real-world settings. However, one example is a suite of solutions for detecting 0-day exploits and insider threats is provided by Darktrace [48]. Their approach is based on a proprietary unsupervised ML approach and as a result very little information is available in public domain on how the technology works apart from the advertising literature. Another effort is Hogzilla [49], an open-source network IDS employing ML methods such as k-means, random forests and superbags, to offer detection capabilities.

2.4.3. Holistic Solutions

Despite the abundance of research papers focusing on the intersection of ML/AI and cybersecurity (many of them covering anomaly/intrusion detection), most of them focus either on a specific ML algorithm, or only consider a particular type of data. The lack of holistic approaches of incorporating ML methods to anomaly detection is highlighted in [50] where Hariharan et al. present a prototype implementation of a real-time automated framework for classifying anomalies, achieving high efficiency and a 95% accuracy.

The need for a robust solution incorporating ML approaches not only into one particular area (for instance file inspection), but on various levels (for endpoint, application and network security), is also evident in the commercial products offered by various cybersecurity vendors. The most comprehensive solutions available on the market today are Self-Evolving Detection Systems (SEDS). SEDS systems are typically backed by a combination of different ML and/or DL algorithms. Inevitably, the self-evolving aspect of such solutions involves continuous learning that require large quantities of

data, which might be impractical (or too expensive) for deployments in small- and middle-sized operational environments

In terms of commercial efforts, Fortinet, as an early adopter of AI, launched their self-evolving detection system FortiGuard AI, which uses a DL model built around an artificial neural network. The latter is used in production and is available as a part of Fortinet's threat intelligence back-end [51] [52]. The company claims its success and effectiveness even with zero-day malware. Another product advertising heavy utilisation of ML is eSentire's Managed Detection and Response (MDR), a cybersecurity AI software that uses anomaly detection to identify security threats in enterprise networks [53] [54]. Enterprise Immune System [55] and Darktrace Antigena [56] offered by the aforementioned Darktrace also fall into the category of more advanced, holistic solutions.

However, in spite of the evident effect that the popularization of ML approaches has had on the development of commercial antivirus, IDS and threat intelligence solutions, information concerning their true role within such products is typically scarce and (likely intentional) somewhat vague.

2.4.4.Supporting Cyber-Situational Awareness and Incident Response

To the best of our knowledge, despite ongoing efforts in the research community, none of the present AI methods for intrusion detection and prevention removes the need for human intervention. Even with a high degree of automation and integration with existing security tooling, a skilled security analyst or network administrator is still required to assess alarms issued as a result of false positive detections by an AI system. However, with the increasing performance and detection accuracy of such systems, their usage promises to alleviate the workload associated with the analysis of (possibly vast amounts of) false alarms [57] thus positively affecting productivity of CERT and CSIRT teams. In addition, as clearly stated in [40], the autonomous capabilities of ML algorithms in general must not be overestimated, since the absence of human supervision could further facilitate skilled attackers - or malicious insiders - to infiltrate, steal data, and even sabotage an enterprise. Based on current trends, we expect an increase of AI-based tools for cyber-threat intelligence and establishing cyber-situational awareness, yet human operators will continue to play a crucial role in incident response. Indeed, the existing products on the markets such as the aforementioned eSentire MBR solution recognise the limitations of using ML alone, along with the need to complement it with human intervention [53], which is evident from their service offerings.

A related topic is the user of the edge device(s), for whom cybersecurity awareness and training is equally relevant. Phishing attacks remain a persistent threat, now augmented by advanced variations such as "smishing" (SMS-assisted phishing). In any security aware environment, it is necessary not only to educate end users, but also to detect when they have misbehaved, and react accordingly. The best practice is to provide security training modules to users that deploy in the same way as attacks – through email, chat – which have no malicious payload but instead direct exploited users to learning pages on security best practices. Of course, this does not stop the malicious insider threat, which remains a concern to the information security industry – but monitoring that detects an external attack can often be extended to include suspicious and unusual behaviour arising from a comprised account or an insider starting to behave maliciously.

Increases in popularity and utilisation of sophisticated ML and AI approaches in cybersecurity promises improved and robust defence mechanisms. However, it also introduces new attack vectors and tools for cybercriminals. Calderon [57] while focusing on the benefits of AI in cybersecurity, also warns against the additional risks brought by their usage, emphasising the importance of understanding both their benefits and risks, and finding a balance between them.

An interesting fact to consider when deploying ML-backed cybersecurity solutions is that existing ML systems themselves are, by design, prone to some vulnerabilities that a motivated attacker may exploit to cover up their malicious activity and avoid detection. These are being actively researched in the scope of adversarial ML, the study of ML vulnerabilities. In general, the idea of adversarial ML is

to feed the input layer of the model with data that diverts the system into making a prediction favourable for them. This can be achieved by (1) data poisoning attacks where the training data is tampered with, resulting in the model consistently making mistakes, and (2) adversarial examples, where inputs to the existing model are designed in a manner that causes misclassification [58]. While the research community presents various examples of different attacks against ML approaches, demonstrated for spam detection, malware analysis and intrusion detection [59] [40], the true extent to which exploiting these vulnerabilities is feasible in practice remains unclear. Nevertheless, according to Goodfellow et al. [60] adversarial examples represent a concrete and difficult safety issue that requires a serious research effort.

What's more, not only can the hackers trick cybersecurity controls backed by ML or even exploit additional vulnerabilities introduced by them, they can also leverage the same technology for their own malicious purposes. According to Brundage et al. [58] the same advanced technology that state-of-the-art ML/AI approaches are using to detect 0-day vulnerabilities may be (ab)used to distribute sophisticated malware. Thus, the dual-use concerns for AI are becoming just as relevant in the field of cybersecurity as they are in several other areas where it can thrive.

2.4.5.Data Protection

Despite the interest sparked by Privacy by Design in the regulatory and policy domains, industry does not seem to have kept the pace. According to recent surveys [61] industry focus is on privacy management (e.g. data inventories, consent management, breach response) and protecting access to information through extensive legal policies and corresponding IT configurations (e.g. firewalls and access control). Unfortunately, these approaches do not seem to have worked well, as demonstrated by the endless flow of reports of privacy violations at scale and technology companies' failure to fulfil basic data protection requirements. Clearly, there is an urgent need also on the industry side to develop products and services that embed privacy principles by design, and not only as an afterthought.

From a technical perspective, computer scientists have carried out extensive research to 1) identify different threats that new technologies pose to privacy, and 2) propose innovative technologies that address some of them — e.g. by developing the so-called Privacy-Enhancing Technologies (PETs) and Transparency-Enhancing Technologies (TETs). However, these researchers have rarely addressed the systematisation or generalisation of their approaches so that engineers could adopt and integrate them into their solutions [62]. Therefore, engineers still need well-elaborated, feasible, and appropriate methods and supporting techniques and tools to adequately understand and mitigate privacy risks for existing as well as to-be-designed ICT systems.

To address these needs, privacy engineering has emerged as a new field of research and practice focused on bridging the legal, technical and engineering approaches to support engineers in systematically identifying and addressing privacy and data protection concerns during a development lifecycle. Privacy engineering is still in its infancy [63] with a lack of proper methodologies, techniques and tools that support privacy engineers' work. The state of the art for privacy engineering methodologies [64] [65] either assumes a waterfall development process imposing complex activities disconnected from mainstream software engineering practices or remains at an abstract level, thus hindering their adoption. Currently, only some parts of privacy engineering are appropriately addressed by research work, mostly in the areas of threats analysis and requirements engineering [66]. These are valuable contributions that can help identify what a system must or must not do, however, they provide insufficient support on how to achieve an appropriate solution (as they are disconnected from the myriad of PETs and TETs available) and evaluate whether the engineer did a good job.

In addition, privacy engineering research is challenging as it demands the collaboration of a multi-disciplinary and inter-sectoral network for two primary reasons. Firstly, privacy engineering constitutes a research problem involving multiple disciplines, and therefore requires joint, novel inter-

disciplinary research efforts focused on analysing, synthesising and harmonising the approaches of the diverse disciplines into a coordinated and coherent whole, in order to efficiently address the challenges associated with the topic. Secondly, privacy engineering research requires the participation of various stakeholders to succeed. Academic research could give rise to a set of methods, techniques and tools which support engineers in building privacy-friendly systems. But putting this vision into practice requires the flow of information and transfer of knowledge from research to practice and vice versa.

2.5. Resource Management and QoS

In this section, we investigate current research proposals dealing with the resource management in fog and cloud, considering QoS parameters. One proposal is presented by Gill et al. [67], where they propose a new resource management framework in order to optimise multiple QoS parameters such as, response time, network bandwidth, energy consumption and latency, simultaneously by means of a Particle Swarm Optimization algorithm (PSO). The results show that in comparison with two similar techniques, the new proposal demonstrates better performance in terms of latency, energy consumption, and network bandwidth consumption.

Another proposal by Cardellini et al. [68], addresses the elasticity of resources for geo-distributed systems running over multiple edge/fog and cloud infrastructures. Specifically, the authors argue that a decentralised system to control the elasticity and then to meet certain QoS attributes is required.

Mahmud et al. [69] present a latency-aware approach for placing application modules on distributed fog nodes, considering applications with different latency requirements. Considering these latency requirements, the policy identifies which applications should be placed on lower fog levels and which should be placed on higher levels. The latency QoS parameter is optimised together with minimisation of energy consumption by reducing the number of active (executing) fog nodes.

Peng et al. [70] present iCloudFog, a scalable and agile integrated Cloud–Fog architecture, which provides Fog and/or Cloud resources in response to IoT data processing requests. They identify the main challenges in this architecture as being the IoT job scheduling with QoS requirements. These job or tasks requirements are the maximum allowed latency, security level, and requirements related to the nodes where the task is executed, such as computing, storage and bandwidth requirements.

Akintoye and Bagula in [71] presents the task allocation and the virtual machine placement problems in a single cloud/fog computing environment; and propose two algorithm based resource allocation solutions: one task allocation algorithm, binding-based, and the second one based on a genetic algorithm for the virtual placement problem, both guaranteeing the QoS requirements in terms of resource allocation cost.

Gill et al. [72] propose ROUTER, a resource management manager for fog computing considering QoS parameters such as system response time, network bandwidth, energy consumption and latency, not on individual basis but rather as a simultaneous multi-parameter optimisation for decision making. The problem is solved based on a multi-objective Particle Swarm Optimisation algorithm (PSO), and the performance is compared with other two similar management techniques, showing tangible benefits in terms of reduced network bandwidth, latency and energy consumption.

In the dynamic context of mF2C, the QoS providing component relies on Deep Reinforcement Learning to select devices to execute a service, excluding those that can potentially cause SLA violations. Other research efforts take a similar approach to allocating resources in dynamic environments. For example, Wang et al. [73] propose a Deep Reinforcement Learning based Allocation (DRLRA) that allocates computing and network resources in an adaptive manner and apply it to Multi-Access Edge Computing (MEC) environments. Their results show a reduction in average service time and a better balance in the use of resources.

2.6. Convergence of AI and Computing

In the previous version of this deliverable the main trends with respect to the convergence of AI and computing were identified. Two key trends were highlighted. Firstly, works that focus on exploiting the inherent parallelism in fog to cloud systems to execute complex AI algorithms. Secondly, approaches which apply AI to the control and management of fog-to-cloud systems, for example, to improve the match between requested services and available resources. In this deliverable the second trend is further investigated by reviewing works relating to the use of AI and ML to improve the performance of fog to cloud systems published in the last 12 months.

Aryal and Altmann, [74], propose an optimal allocation algorithm (genetic algorithm) for virtual machines (VMs) in a cloud federation environment. The allocation is a multi-objective optimisation, considering parameters such as the computing resource capability of VM instances, the application footprint, and the distance of VM instances. The main novelties of the proposal are twofold: the eligible resources are part of a federated cloud, and the multi-objective nature of the optimisation.

Focusing on cloud access radio networks, Chen and Leung in [75] introduce cognition-based communications, consisting of two layers namely cognition and communications, where AI is not only applied to cloud resource allocations, but also to the network management optimisation. The main contribution of this paper is the global optimisation of resource allocation, considering as resources both network and computing resources, using techniques such as data mining, deep learning, ML and AI; furthermore, the network is configured and managed by means of Network Function Virtualisation (NFV), Software Defined Networks (SDN) and network slicing.

Carnevale et al. in [76] propose a smart orchestrator based on AI for Osmotic computing. The orchestrator optimises the deployment and the migration of microservices based on the requirements of both: resources (i.e., load balancing, reliability, availability) and the applications to be executed (i.e., detection, implementation, awareness of the context, proximity, QoS). The AI module learns through monitoring of the Osmotic resources deployed on Cloud, Edge and/or IoT, and its target is a prediction model to deploy/migrate the microservices based on previous experiences.

La et al. [77] propose optimisation of system operations and improving network performance in fog computing, in terms of delay and energy consumption. Their research is focussed on device-driven and human-driven intelligence. IoT and edge devices at the edge of the network (fog) are endowed with smarter capacities, which allow fine-granularity information extraction in order to make local decision about resource management. On the other hand, human behavioural patterns can also be exploited to train the network to be smarter. To verify these two approaches, device-driven and human-driven, the authors propose two algorithms, one of them optimises the tasks offloading in the presence of different fog nodes, and the second considers user behaviour to perform adaptive low-latency Medium Access Control (MAC)-layer scheduling among sensor devices.

2.7. Key Takeaways

The key areas of focus in relation to scientific trends of significance for the project are as follows:

- Increasing focus on the joint management of communications and computing resources in fog-to-cloud environments. Management approaches are increasingly leveraging AI and Deep Learning approaches to improve system performance and resource utilisation efficiency. Novel approaches to resource management based on blockchain have also emerged. Additional research is necessary to efficiently allocate potential AI deployments on distributed edge infrastructure, while considering the specific constraints edge devices. Blockchain approaches should be also tuned to meet the particularities of the fog and edge, due to the limited computational and storage capacity of these devices in order to store and compute the block.

- A number of new data management solutions have been reported in the literature over the last 12 months. Currently, these approaches are out of scope for mF2C due to their specific focus on file storage which does not enable a fine-grained control on the data to be replicated in order to avoid unnecessary copies and transfers. Beyond this, the challenge will be to provide a holistic data management solution, from the edge to the cloud that facilitates the efficient management and analysis of the large amounts of diverse data acquired from different sources. Future research should also address of how new functionalities, such as storage or network can be offered within a system.
- Task scheduling and offloading mechanisms due to issues relating to device mobility, limitations in devices energy budgets and the impact of the network on the performance of fog-to-cloud management framework remains an area of active focus. Optimisation of task scheduling on different nodes within the infrastructure, based on QoS constraints and the seamless execution of the tasks following the FaaS model are relevant beyond mF2C. This approach will enable the interoperability with commercial solutions (such as Amazon Lambda) still avoiding lock-in to a single provider. Finally, additional research on mobility patterns and mobility forecast strategies is necessary in order to efficiently manage resources for an optimal offloading deployment.
- The literature indicates that carefully selected DL-based models offer better detection accuracy than conventional ML approaches, in various deployments.
- The gap is the application of AI/ML based approaches for cyber security is the continued need for human intervention. However, fully autonomous operations are unlikely for the foreseeable future. Human operators will continue to play a crucial role in incident response as reflected by commercial solutions (e.g. eSentire MBR) which recognise the limitations of using ML without complementary human intervention. However, additional research is necessary to identify scenarios where human intervention may be minimised, in order to negate mandatory roles in semi-autonomous systems.
- Privacy engineering has emerged as a new field of research and practice that aims to bridge the legal, technical and engineering approaches to support engineers in systematically identifying and addressing privacy and data protection concerns during a development lifecycle.
- From a security awareness perspective best practice is to provide security training modules to users that deploy in the same manner as real attacks – through email, chat – which have no malicious payload but instead direct users to learning pages on security best practices. Many commercial providers already provide basic training for corporate customers (e.g. with workers in the field, using mobile phones and/or VPN), but there is little available for IoT, due to the heterogeneity of devices. As a future extension, there is an opportunity to develop IoT security-awareness training for end users, which can then be sold to companies that provide IoT “solutions” so they can tailor the training and provide it to their own end users. Another approach to security is to harden the devices; Microsoft’s Azure Sphere is an example of how devices can receive field updates. As regards the outcome of the mF2C project specifically, it should be clear from previous deliverables that the basic security components are already there, but will need increasingly higher TRLs based on experience gained through field tests of use cases. From the consumer perspective, the most obvious benefit would arise from integration with the user’s choice of social media and/or cloud platforms; specifically, Facebook, LinkedIn (for business travellers), MicrosoftLive, and Google’s id. Apart from letting users authenticate with pre-existing ids, mF2C could be extended with features from these platforms – clouds for application integration, and social media for sharing.
- To further increase the level of security in mF2C through Artificial Intelligence or Machine Learning approaches, research into SIEM (Security Information and Event Management) will be necessary. This will enable comprehensive monitoring of the complete infrastructure and,

by feeding the data into an AI processing pipeline, supporting better responses to known and unknown attacks.

- New approaches to resource management in order to maintain QoS in terms of latency, bandwidth and energy consumption etc. continue to emerge based on the techniques such as Particle Swarm Optimisation, genetic algorithms and Deep Reinforcement Learning to provide adaptive allocation of resources and to identify service on-boarding locations. Future research may need to consider quality as perceived by the user (QoE) in a resource management and allocation strategy which is adaptable to the requirements of human experience. An obvious benefit would be the ability to target execution on green computing platforms, which might deliver results slower than more costly HPC platforms, but are still sufficiently performant.
- Smart orchestrator solutions based on AI for Osmotic computing have emerged. Reported solutions focus on requirements (resources and application) fulfilment for microservice deployment and migration. Features include an AI module capable of learning through monitoring of Osmotic resources deployed across the Cloud, Edge and/or IoT, continuum. The target is a prediction model to deploy/migrate the microservices based on previous experiences which has relevance in the future evolution of mF2C.

3. Technology Trends

Previously the key technological trends regarding the management of Cloud, Fog and IoT devices have been reviewed. Although those revised trends are still valid, this section provides an update on the latest trends in Cloud and IoT management.

3.1. Tools, Platforms, IoT

This section focuses on some of key open-source cloud and IoT management platforms and tools, which are starting to gain traction in the fog-to-cloud domain.

3.1.1. Cloud Management Tools

Apache CloudStack[4] provides an open-source and multi-hypervisor, multi-tenant based Infrastructure-as-a-Service (IaaS) platform. It automates the deployment of virtual machines and provides a comprehensive management suite to manage deployed VMs. It also provides public and private cloud services, as well as a full and open native API for users to interact with the different components of the platform. Interestingly, it provides support for most of the popular hypervisors (i.e. - VMware, KVM, Citrix XenServer, Xen Cloud etc.). This cloud management tool offers a virtual network computing (VNC) AJAX client to the users for accessing available computational resources using any of the latest internet browsers without adding any add-on or plug-in.

Cloudify[5] is another open-source cloud management tool motivated by the goal of providing improved flexibility and scalability functionality for existing technological solutions. Adhering to the OASIS Topology and Orchestration Specification for Cloud Applications (TOSCA) standard, helps to describe any application or network services in a generic, human-readable and intuitive modelling language. It provides pre-installed network functions which has helped the platform to gain popularity. It also provides a complete UI and various dashboards which allow users to fully control applications or network services. An important feature of the tool is the functionality to orchestrate various distributed network resources, edge devices and applications, therefore ensuring improved scalability among the various computational resources.

3.1.2. IoT Management Tools

KAA[6] is an open-source IoT middleware solution focused on providing improved scalability to enterprise IoT development. It develops an end-to-end IoT solution by connecting applications and smart devices. The platform allows developers to manage IoT devices within a system and also helps to orchestrate end-to-end data processing among IoT resources. It ensures improved communications, control and interoperability among the system resources. The platform consists of flexible microservices and has the functionality to adapt the behaviours of these microservices. It also provides a vast range of networking technology options (i.e., Wi-Fi, Ethernet, ZigBee, MQTT, CoAP, XMPP, TCP, HTTP, etc.) in order to facilitate the most appropriate solution for establishing communications between servers and endpoints.

Similar to KAA platform, DeviceHive[7] is also an open-source platform for IoT development distributed under and Apache 2.0 licence, which facilitates quicker connectivity between the edge devices and cloud resources. The platform uses a container based approach to provision the components of the platform. It provides support for most of the “Big data” management tools such as Apache Spark, Cassandra, Elasticsearch etc. The platform supports batch processing and machine learning functionalities for application collected data, which has helped the platform to gain traction with both

[4] Apache CloudStack - https://cloudstack.apache.org/index.html
[5] Cloudify - https://cloudify.co
[6] KAA - https://www.kaaproject.org
[7] DeviceHive - https://devicehive.com

developers and users. By providing support for MQTT, REST API and Websockets, DeviceHive enables direct connectivity to edge devices.

As the name suggests the Device Services Architecture (DSA)[8] is a distributed open-source IoT platform, providing inter-communication, logical facilities and applications within each layer of its infrastructure. The developer's vision for DSA is the unification of information among disparate devices, services and applications enabling a structured and adaptable real-time data model. In a decentralised and distributed manner, it helps purpose-built products and services (i.e. DSLinks) to interact with each other. The platform also features support multi-tenant applications.

3.2. Technology Trends Emerging from HPC

3.2.1. Data Management Trends

In addition to data management technologies for edge to cloud environments, various database vendors which previously focused on offering high performance capabilities, have recently started to focus on the edge-to-cloud scenario. Their aim is to support analytics by providing extreme scalability, high performance, and support for specific data models relating to sensor data. In some cases, the database system is specifically designed for a particular data model, such as the time series database InfluxDB[9]. Other vendors provide support for time series data by means of a flexible data model that can also store data in other structures, such as MongoDB[10], Cassandra[11], or CrateDB[12].

However, the challenge is how to provide high performance data persistence capabilities on resource limited edge devices. Also, due to the intrinsic characteristics of edge-to-cloud scenarios, devices which can arbitrarily join and leave the infrastructure at any time, need to be supported. Commercial products such as eXtremeDB[13], IBM Informix[14], Redis Enterprise[15] and HarperDB[16], and the open-source YottaDB[17] and dataClay[18] offer these functionalities.

Similar to dataClay, the technology at the core of the mF2C data management functionality, eXtremeDB supports in different flavours, from HPC to edge, so the same system can be used from the edge to the cloud in order to facilitate data management operations. Regarding the sharing of data between devices, eXtremeDB provides per-type and bi-directional replication from edge devices to the cloud, and vice versa. However, the level of each node and the direction of the data transfer is fixed, which prevents the flexibility of dynamically changing the role of a node according to the needs, as required in mF2C.

3.2.2. Programming Model Trends

Previous deliverables have addressed the requirements of programming frameworks to design applications which can be executed across the edge/cloud continuum. The key trends in commercial offerings which are extending traditional cloud services to IoT devices and enabling applications to

[8] Device Services Architecture (DSA) - (http://www.iot-dsa.org/)
[9] https://www.influxdata.com/
[10] https://www.mongodb.com/use-cases/internet-of-things
[11] https://academy.datastax.com/use-cases/internet-of-things-time-series
[12] https://crate.io/
[13] https://www.mcobject.com/extremedbfamily/
[14] Philip Howard. IBM Informix and the Internet of Things. White paper, IBM 2016. Available: https://www.ibm.com/downloads/cas/GKJ3QNOL
[15] https://redislabs.com/blog/ideal-iot-edge-database-redis-enterprise/
[16] https://www.harperdb.io/
[17] https://yottadb.com/use-cases/internet-of-things/
[18] 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)

seamless execute code on sensors producing data and moving it to the cloud when needed have been analysed.

Here we move the focus from execution management to the application level, where the trend is on the implementation of AI applications capable of running on computational edge devices. The main challenge to address is that existing approaches for developing and deploying cloud and edge software are not well integrated and leave to the software developers with the responsibility of partitioning the application and the AI models inference across the full ecosystem, of explicitly managing deployment, communications, and failure recovery. Programming frameworks such as COMPSs provides an integrated solution to define high-level annotations for constraints and code dependencies including performance parameters to inform the allocation of tasks on edge or cloud resources. As outlined in D2.2, tasks are defined following a Function as a Service (FaaS) model. A number of ML algorithms implemented in COMPSs are already available as a Distributed Computing Library (dislib)[19] inspired by scikit-learn, that simplifies the task of developing applications by providing a common interface to all algorithms. COMPSs also provides the Tiramisu[20] framework, a data analytics tool for processing, transforming and exploiting embedded data obtained through deep learning models. Furthermore, as a parallel framework, COMPSs exploits inter-node and intra-node parallelism and executes TensorFlow tasks as external processes.

3.3. Cloud Orchestration Platforms, Virtualisation, Containers

Previously we have described how Kubernetes had established itself as the de facto container orchestration tool. According to recent statistics, such as those presented by Ritesh Patel [78], Portworx21 and Aqua22 [79], this trend has continued over the last twelve months, where the adoption of Kubernetes solutions by both small and large enterprises has grown significantly [80]. Statistics also indicate that Kubernetes based solutions are the most widely utilised container orchestration tools by companies and large organisations. These tools, usually offered as part of global cloud solutions, include Azure AKS23, IBM IKS24, Google GKE25, Red Hat OpenShift26 and Amazon EKS27, where Azure AKS (from Microsoft) is the most dominant.

Some of these solutions are based on custom / commercial distributions of Kubernetes, where new functionalities and features have been added. One of these solutions is Openshift, which provides a variety of additional features such as an improved security layer and better management of container images. Other cloud Kubernetes solutions simplify the creation and management of Kubernetes clusters and application execution for users, by offering them as a set of layers that abstract the underlying complexity.

Finally, according to MarketsandMarkets research analysis [81], the projected market value for application container technologies, which include the containers orchestration and platforms such as Docker and Kubernetes, will double in value and is expected to reach over USD 4.98 billion by 2023. Therefore it is expected that the growth of these technologies will continue for the foreseeable future.

[19] https://github.com/bsc-wdc/dislib
[20] https://www.bsc.es/research-and-development/software-and-apps/software-list/tiramisu
[21] Portworx - https://portworx.com/
[22] Aqua Security Software Inc. - https://www.aquasec.com/
[23] Azure Kubernetes Service (AKS) - https://azure.microsoft.com/es-es/services/kubernetes-service/
[24] IBM Cloud Kubernetes Service - https://www.ibm.com/cloud/container-service
[25] Google Kubernetes Engine - https://cloud.google.com/kubernetes-engine/
[26] Red Hat Openshift - https://www.openshift.com/ https://www.openshift.com/products/container-platform
[27] Amazon Elastic Kubernetes Service - https://aws.amazon.com/eks

3.4. Role of Standards in Technologies

IoT, Fog, Cloud Computing are under pinned by the ability of the constituent components to communicate in a seamless manner. In order to achieve this level of connectedness which is built on existing standards it is also necessary to develop new standards with the intent to support new concepts and capabilities [82].

Previous deliverables D2.1 and D2.2 introduced a number of the key standards organisations and initiatives potentially relevant to mF2C. This deliverable details the developments of importance during the course of the project. Of particular importance have been the various ISO/IEC JTC1 sub-committees, the OpenFog Consortium, and its collaboration with IEEE in terms of impact to standards relevant in the context of mF2C project.

3.4.1.ISO/IEC JTC1 SC38 Cloud Computing and Distributed Platforms

ISO/IEC JTC1 continues to be the body responsible for defining technical standards at a formal international level. It continues to refine its set of working groups and sub-committees, the most relevant sub-committees including SC38 and SC41. Insights and perspectives are fed into the ISO/IEC JTC1 standardisation processes via the national standards organisations of participating countries. It should be noted that the standards authored by ISO / IEC JTC1 are often high-level and descriptive in nature. Technical specifications of APIs are often developed within Industry Groups. If appropriate measures are taken they may then be ratified by the relevant ISO/IEC JTC1 sub-committee.

Although focusing on centralised cloud systems in the past, SC38 now pursues a broader work programme with explicit references to Edge Computing in particular. Work Group three which is 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 outlined in Table 3-1.

Standard	Description
ISO / IEC 19086-3:2017	Service Level Agreement Framework - Part 3: Core conformance requirements or service level agreements (SLAs) for cloud services based on ISO/IEC 19086‑1 and guidance on the core conformance requirements.
ISO / IEC 19941:2017	Specifies cloud computing interoperability and portability types, the relationship and interactions between two the aspects. The standard also specified the common terminology and concepts used to discuss interoperability and portability, particularly relating to cloud services
ISO / IEC 19944:2017	Extends the existing cloud computing vocabulary and reference architecture in ISO/IEC 17788 and ISO/IEC 17789 to describe an ecosystem involving devices using cloud services. The standard also describes the various types of data flowing within the devices and cloud computing ecosystem. In addition the data categories and data use.

Table 3-1 ISO / IEC Cloud Computing Related Standards

At present ISO/IEC JTC 1/SC 38 is actively developing a number of standards and technical reports which are of relevance to mF2C. Standards and technical reports being actively developed include:

- ISO/IEC DIS22624 – Cloud Computing – Taxonomy based data handling for cloud services
- ISO/IEC CD22123 – Cloud Computing – Concepts and terminology
- ISO/IEC PDTS 23167 – Cloud Computing – Common Technologies and Techniques

- ISO/IEC PDTR 23188 – Cloud Computing – Edge computing landscape

3.4.1.1. ISO/IEC JTC1 SC41 Internet of Things and Related Technologies

ISO/IEC JTC1 SC41 - Internet of Things and related technologies was inaugurated during 2017 with a mandate to serve as the focus and proponent for the joint technical committee 1's (JTC 1) standardisation program on IoT. In addition, it provides guidance to JTC 1, IEC, ISO and other entities developing IoT-related applications. To date, SC41 has published two documents as outlined in Table 3-2 of particular relevance to mF2C.

Publication ID	Title	Description
ISO/IEC 20924:2018	Internet of Things (IoT) – Vocabulary	Provides a definition of Internet of Things along with a set of terms and definitions forming a terminology foundation for IoT
ISO/IEC 30141:2018	Reference Architecture	Specifies a general IoT Reference Architecture in terms of defining system characteristics, a Conceptual Model, a Reference Model and architecture views for IoT

Table 3-2 ISO/IEC JTC1 SC41 IOT Publications

In addition, SC41 is working on a suite of relevant standards including:

- ISO/IEC NP 30149 – Trustworthiness framework
- ISO/IEC NP 30161 – Requirements of IoT data exchange platform for various IoT services
- ISO/IEC NP 30162 – Compatibility requirements and model for devices within industrial IoT systems
- ISO/IEC NP TR 30164 – Edge Computing

An agreement has also been reached between SC41 and SC38 on standards related to Edge computing which will be developed in cooperation with each other.

3.4.2. Alliance for Internet of Things Innovation (AIOTI)

AIOTI was launched in March 2015 by the European Commission and Internet of Things stakeholders. The overall goal of the initiative is the creation of a dynamic European IoT ecosystem to unleash the potential of IoT. Working groups within AIOTI 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 continues to provide updates to the project on direction and opportunities to influence.

3.4.3. OpenFog Consortium

The OpenFog Consortium was established in November 2015 to address technical challenges in Fog computing. In January 2017 they released details of their OpenFog Reference Architecture [83]. In August 2018 the consortium's collaborative efforts with the IEEE led to the publication of the IEEE 1934 standard which adopted the OpenFog Reference Architecture for Fog Computing [84]. The standard is a structural and functional prescription of an open, interoperable, horizontal system architecture for distributing computing, storage, control and networking functions closer to the users along a cloud-to-thing continuum of communicating, computing, sensing and actuating entities.

The standard is based of eight core principles, known as pillars encompassing (i) security; (ii) scalability; (iii) openness; (iv) autonomy; (v) reliability, availability, and serviceability (RAS); (vi) agility; (vii) hierarchy; and (viii) programmability [85].

It encompasses various approaches to disperse Information Technology (IT), Communication Technology (CT) and Operational Technology (OT) Services through information messaging infrastructures as well as legacy and emerging multi-access networking technologies. The "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".

In January 2019, the OpenFog Consortium amalgamated with the Industrial Internet Consortium (IIC) [86]. Subsequently, all OpenFog Consortium activities have been integrated into the IIC Working Group activities. The IIC published their Industrial Internet of Things Volume G1: Reference Architecture v1.9 in June 2019. This forms a foundation upon which various technical publications have been organised. These documents describe how heterogeneous devices may interoperate, from sensors and actuators at the industrial edge through various levels of gateways and hubs out to the Wide Area Network and beyond to the cloud.

3.4.4. National Institute for Standards and Technology (NIST)

In the US, NIST published recommendations (500-325) for a Fog Computing Model [87]. While not a standard is does provide pointers towards NIST's thinking on how to approach standardisation in relation to fog and mist computing and their relationship to cloud based IoT computing models. The document also characterises properties and aspects of fog computing, including service models, deployment strategies together with a baseline of what fog computing is, and how it may be used.

3.4.5. Open Connectivity Foundation (OCF)

The Open Connectivity Foundation has created an extensive and growing reference set of models to enable the discovery and control of arbitrary devices. These models include definitions of the interfaces to these devices, and are published at http://oneiota.org/. The OCF continues to sponsor IoTivity, an open-source reference implementation, published under an Apache 2.0 license, currently at version v2.0.1.

3.4.6. Linux Foundation

On a related note, the Linux Foundation has recently evolved its EdgeX Foundry into the broader LF Edge Foundation. This open-source project was launched in January 2019 with significant industry backing. It is building an open source framework for the edge and its revised charter has recently been adopted. Numerous open-source projects are included within its framework such as the Akraino Edge Stack (first release June 2019), and the Edge Virtualization Engine (EVE).

3.4.7. European Telecommunications Standards Institute (ETSI)

ETSI continues to be very active in numerous areas relevant to IOT/Edge. As well as a suite of standards that enable 5G, ETSI also has a Multi-Access Edge Computing (MEC) Industry Specification Group which is developing a suite of standards to enable cloud-computing and an IT service network at the network edge.

3.4.8.Open Grid Forum (OGF)

The Open Grid Forum continues its work to promote interoperation in the area of grids and clouds. Recent updates to the GLUE schema have aligned with work on data centre automation to suggest a GLUEfish, a means to combine data centre data with service information. OGF's network architecture (NSI) also promotes a secure peer to peer communications, relying on rerouting messages when connectivity is lost which has similarities with use case 2 in mF2C. OGF noted the relation to IETF's RFC 8453 (abstraction and control of traffic engineered networks). Finally, OGF was asked to kindly remark on mF2C's use of WS-Agreement, even if not in its original SOAP profile, in OGF's liaison statement to the September 2019 ISO/IEC JTC1 SC38 meeting in Stockholm.

3.5. Technology Trends in Edge Computing

The relevant edge computing trends which were initially identified and evaluated in D2.1 [3] and updated in D2.2 [1] to reflect key developments of significance in the interim time window between the two versions. This version repeats that process by providing a final update on the key trends over the last twelve months since the release of D2.2.

The search for faster and more powerful computing in proximity to sensor data source is driving the computing community to advance both their software and hardware offerings, to provide solutions that are smaller and more efficient, and capable of processing large amounts of data without leaving the edge. More and more, edge computing providers and users are seeking solutions that can provide artificial intelligence, high reliability and fast decision making at the edge, whilst maintaining data anonymization and high portability characteristics that define edge computing. Building on top of the topics addressed in section 3.5 of D2.2 [1], this section provides updates and highlights new trends in edge computing that have arisen in the interim.

3.5.1.Reference Solutions

The solutions referenced in past deliverables remain valid. The key features of these solutions are as follows.

NuvlaBox[28]

The NuvlaBox has been referenced in previous deliverables, however given that it is an integral part of mF2C and considering its recent re-designed, it is worth mentioning again. In order to provide better support to container technologies, SixSq has decided to open-source[29] and re-engineer the NuvlaBox in such a way that its architecture is now fully container compliant. The new NuvlaBox structure is completely decoupled, and can be distributed in two different ways:

[28] https://sixsq.com/products-and-services/nuvlabox/overview
[29] https://github.com/nuvlabox

- via the NuvlaBox Engine (see Figure 3-1) – the core of the NuvlaBox, which can be deployed on any device through Docker Compose or Docker Swarm, turning that device into a functional NuvlaBox, ready for edge applications.

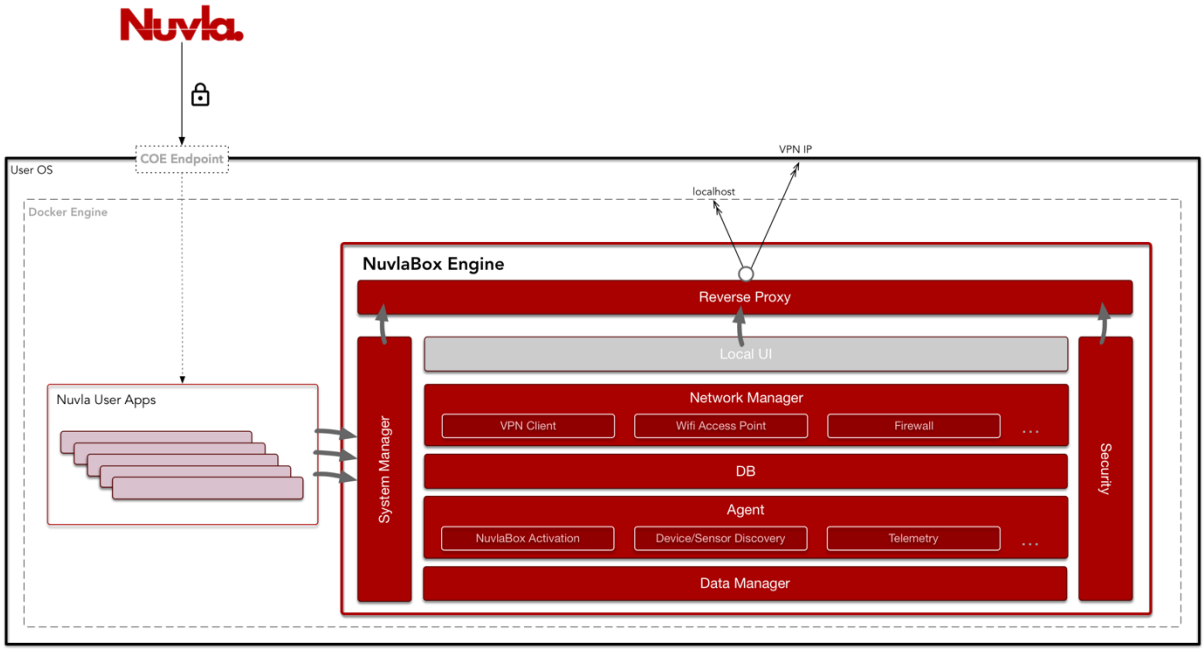

Figure 3-1 - NuvlaBox Engine architecture

- via the NuvlaBox OS (see Figure 3-2) – a specialised Linux-based operating system, containing the NuvlaBox Engine plus additional configurations and optimisations to increase security and improve performance of the device at the edge.

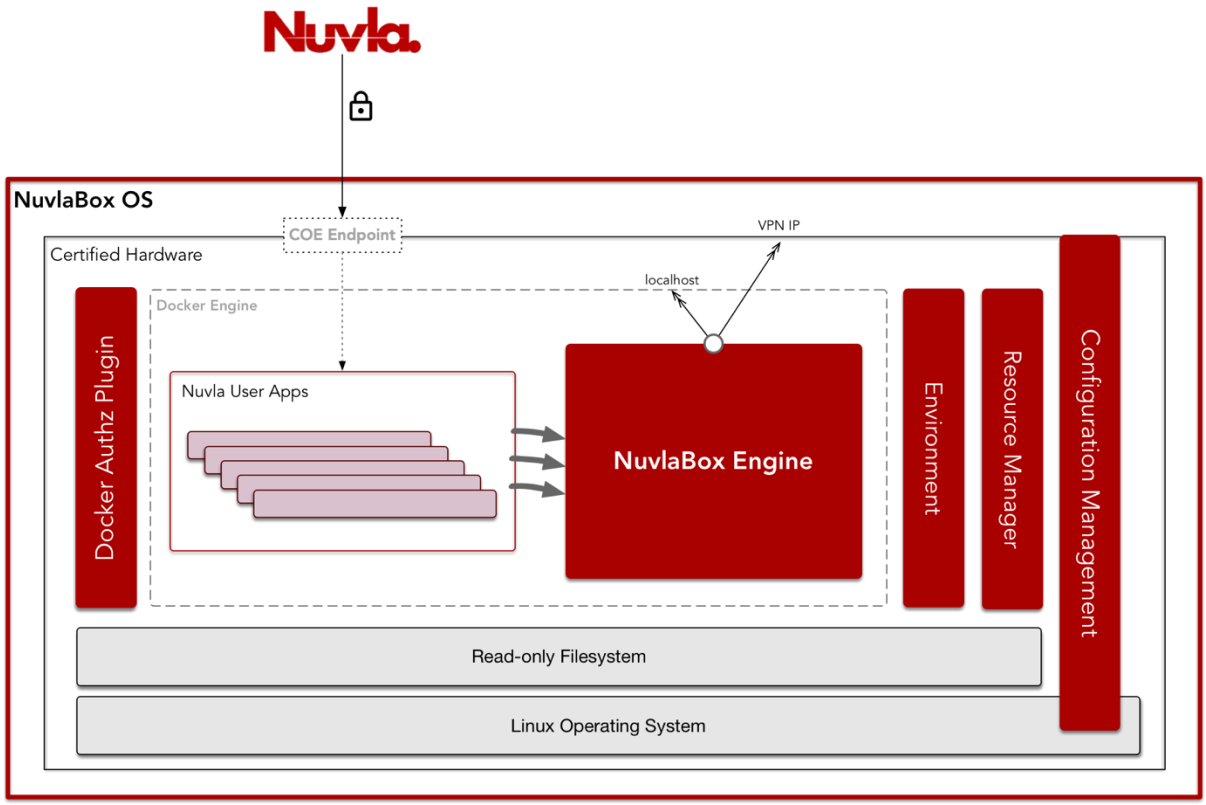

Figure 3-2 - NuvlaBox OS architecture

Eclipse ioFog™ [30]

This is an Eclipse Foundation[31] project provided by Edgeworx[32]. It is an edge computing platform for deploying, running, and networking distributed microservices. It enables users to utilise any device (provided it fulfils the minimum requirements) and to turn it into a distributed Edge Compute Network (ECN).

An Edge Compute Network (ECN) running ioFog consists of one or more devices, referred to as nodes. Each node runs a daemon service called Agent which is responsible for one or more microservices running on that particular node. Another software element is the Controller, which is used for the orchestration and tracking of different Agents. Finally, if inter-node communication is required by the microservices, ioFog includes an optional daemon called the Connector which assists in providing automatic discovery and NAT traversal, brokering direct peer-to-peer (P2P) communication when possible. This is a free and open-source project.

EdgeX Foundry™ [33]

This is a "vendor-neutral open source project hosted by The Linux Foundation building a common open framework for IoT edge computing". Comprising of a loosely-couple microservice architecture, this software platform enables an ecosystem of plug-and-play components that unifies the marketplace and accelerates the deployment of IoT solutions.

Customers can deploy a mix of plug-and-play microservices on compute nodes at the edge, where they sit in the solution stack, and the use case, according to the platform architecture shown in Figure 3-3.

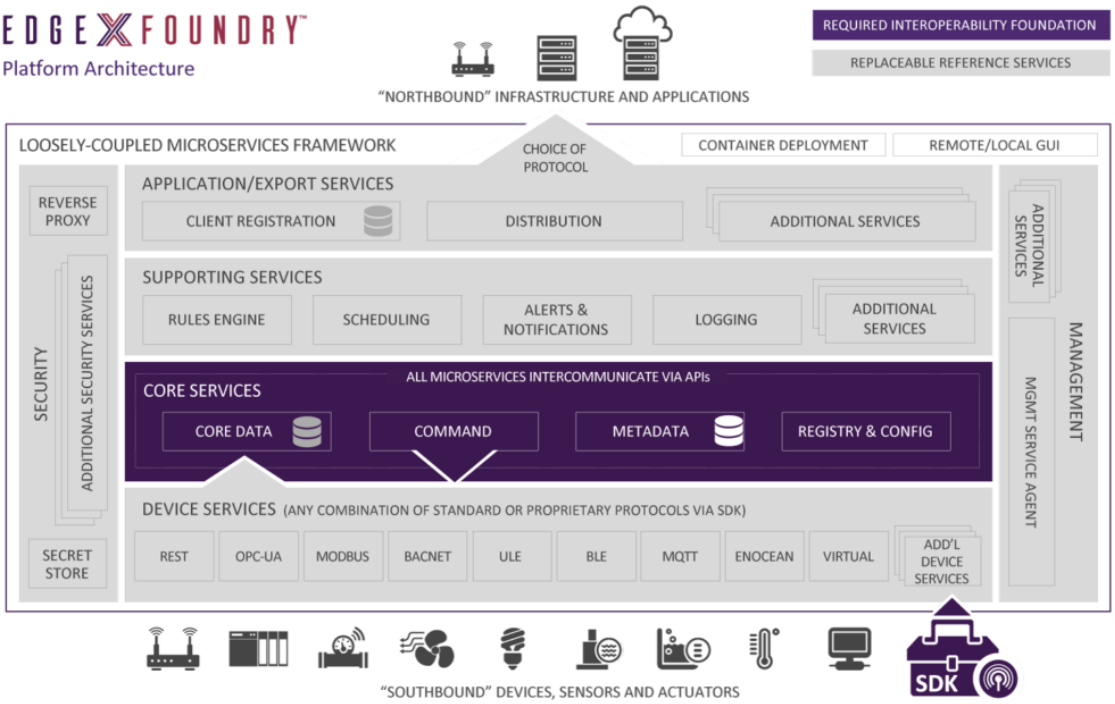

Figure 3-3 - EdgeX Foundry platform architecture[34]

[30] https://iofog.org/
[31] https://projects.eclipse.org/projects/iot.iofog
[32] http://edgeworx.io/
[33] https://www.edgexfoundry.org
[34] https://www.edgexfoundry.org/about/

3.6. Key Takeaways

The following are the key technology trends of significance for the project:

- A number of open source Cloud Management tools (i.e., CloudStack, Cloudify) are available which provide solutions for resource management and scheduling out of the box.
- IoT Management solutions (i.e., KAA, DeviceHive, DSA) orchestrate and interconnect sensing devices, data, and applications over the Web, and in doing so address some of the key challenges of interoperability and scalability associated with large IoT deployments.
- Database vendors are evolving their offerings to provide solutions which are more suitable for computation on resource limited edge devices. Other improvements include better support for analytics at the edge and accommodating dynamic behaviours at the edge due to devices joining and leaving.
- Implementation of AI applications capable of running on computational edge devices is a growing trend. However, in order to address abstracting the complexity of AI application lifecycles from developers a number of programming frameworks have emerged such as COMPSs.
- Kubernetes and related solutions such as Openshift have emerged as the dominant container orchestration approaches for companies and cloud solution providers such as Amazon and Microsoft. The market for container related technologies is expected to double over the five years to USD 4.9 billion by 2023.
- The amalgamation of the OpenFog Consortium with the IIC in January 2019 and their subsequent publication of the IoT Reference architecture in June 2019 was a key development as it increases the momentum around a common reference architecture.
- The activities of standards organisation such ISO/IEC in IoT and Cloud are now being complimented by tangential and parallel activities by other organisation in the broader context of IoT/Cloud related technologies. Increased levels of standardisation and convergence across standard complimentary initiatives will be key to growing the IOT/Edge/Cloud compute market and reducing adoption costs. A key future challenge will be to prevent fragmentation in standards due to competing commercial and national interests
- Edge computing solutions both open source and commercial, have evolved their architectures to provide better support to container technologies such as NuvlaBox. A number of edge open source edge computing platform initiatives for deploying, running, and networking distributed services at the edge based on a microservices architecture are gaining increased visibility. Solutions of note include Eclipse ioFog and EdgeX Foundry.

4. Business Trends

This section covers the topics of cloud/fog/edge computing, IoT, Big Data and security from a business perspective, as well as the evolution of digital business. The key market trends relevant to the project are identified, in order to inform the sustainability strategy for the project being developed within WP6.

4.1. Cloud, Fog and Edge Computing

The Global cloud computing market is growing at a compound annual growth rate (CAGR) of 18% according to MarketsandMarkets [88]. As a result, the market is expected to reach USD 623.3 billion by 2023, an increase of USD 351.3 billion from the USD 272 billion achieved in 2018. Market growth can be attributed to the growing adoption of cloud services globally, and increased adoption of hybrid cloud solutions.

As shown in Figure 4-1Figure 4-1 Cloud Computing Market by Region, 2016-2023 , the same report highlights the growth in the European cloud market mainly supported by increased adoption by SMEs'. Benefits such as customised offerings, reduced operational costs, scalability and flexibility are driving rapid adoption of cloud services by SMEs.

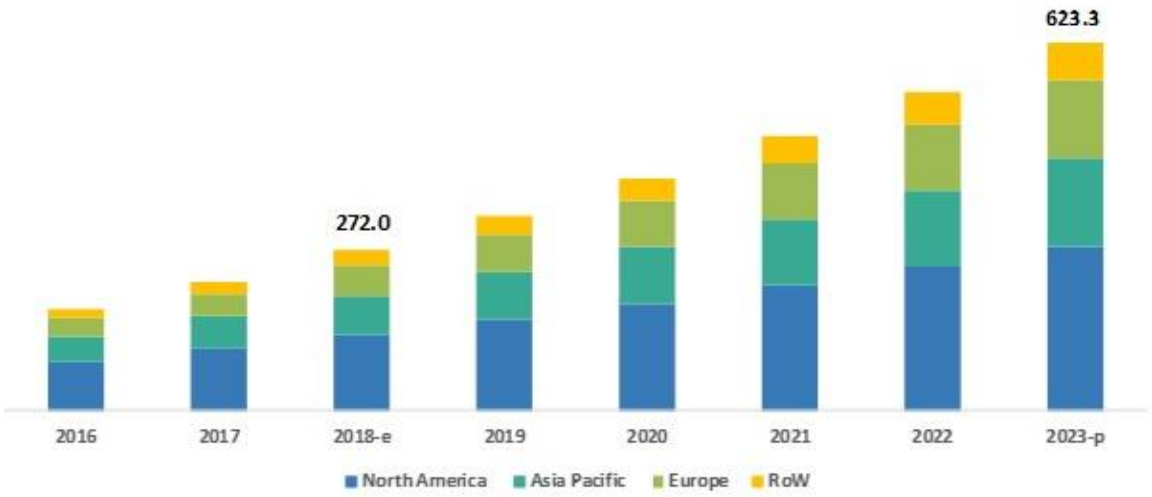

Figure 4-1 Cloud Computing Market by Region, 2016-2023 [88]

A recent study conducted by Adroit Market Research [89] shows the driving factors for the increased adoption of cloud solutions are faster services, lower costs and better security capabilities. The same study reflects the critical role played by cloud computing in digital transformation. For this reason, critical sectors, such as banking, financial services and insurance (BFSI sector) are increasingly demanding cloud services while holding a 7.5% of the global market share.

Research and Markets [90] present a thorough market analysis of Cloud Computing which shows growth at a CAGR of 27.5%, reaching USD 1,250 billion by 2025. Their analysis show the booming economies of emerging regions is supporting the entry of new players into the market while supporting the consolidation of existing ones, through an increasing demand for cloud applications.

All these reports illustrate that cloud computing is not only a well-established market, but there is still room for further expansion. Furthermore, the increasing adoption of IoT technologies, the growth of smart cities and the penetration of e-commerce are driving market growth, according to Wise Guy Reports [91]. At the same time, reductions in free trade, interest rates and data localisation could hinder in growth fog-to-cloud as proposed by mF2C.

Machine-to-machine (M2M) communications, decision-making capabilities, reduced operating costs, distributed data analytics and real-time tracking are the key factors for adopting fog computing solutions by end-use industries. Smart manufacturing sector has the highest expected growth rate (a CAGR of 60%), while the smart grids segment is expected to have one of the highest market share (over 20%) by 2025. Overall, this can be translated into a market worth USD 617.3 million and a CAGR of 61.3% [92] from a market previously valued at USD 22.19 million in 2017 [93].

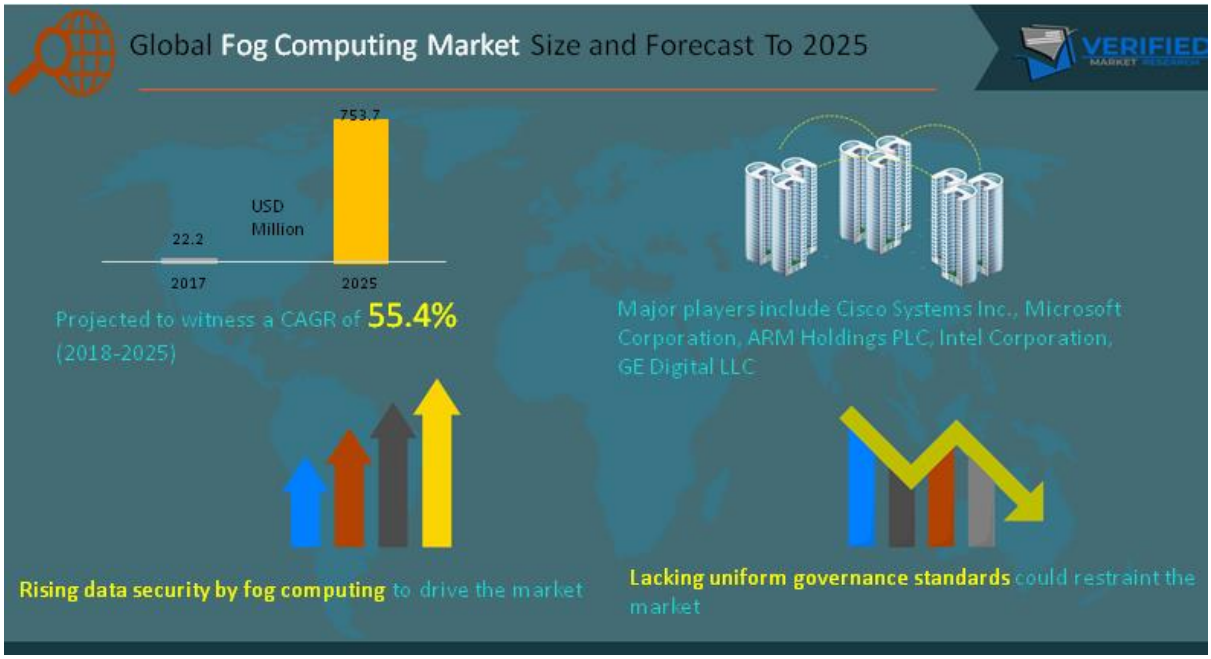

Figure 4-2 Global Fog Computing Market Analysis [94]

Verified Market Research [94] defines fog computing as "*a decentralized computing infrastructure that is used to enhance the efficiency and data computing capabilities for cloud computing. It is capable of decentralizing the computing resources required for processing in the most efficient manner. […] Rising adoption of smart devices has fuelled the growth of fog computing market*".

This decentralised architecture provides high-speed data analytics and shorter response time capabilities, pushing large volumes of data to the cloud. However, the lack of standards is clearly hampering the adoption of fog solutions by companies. Different bodies, such as the OpenFog Consortium (see section 3.4 for more details), are pushing for comprehensive standards to be adopted which can support acceleration in the fog computing market.

The fog approach supports Internet of Everything (IoE) where all devices are connected, allowing more interactions between them. However, it is also expected to dynamically adapt to user needs leading digital transformation by connecting cloud and edge computing.

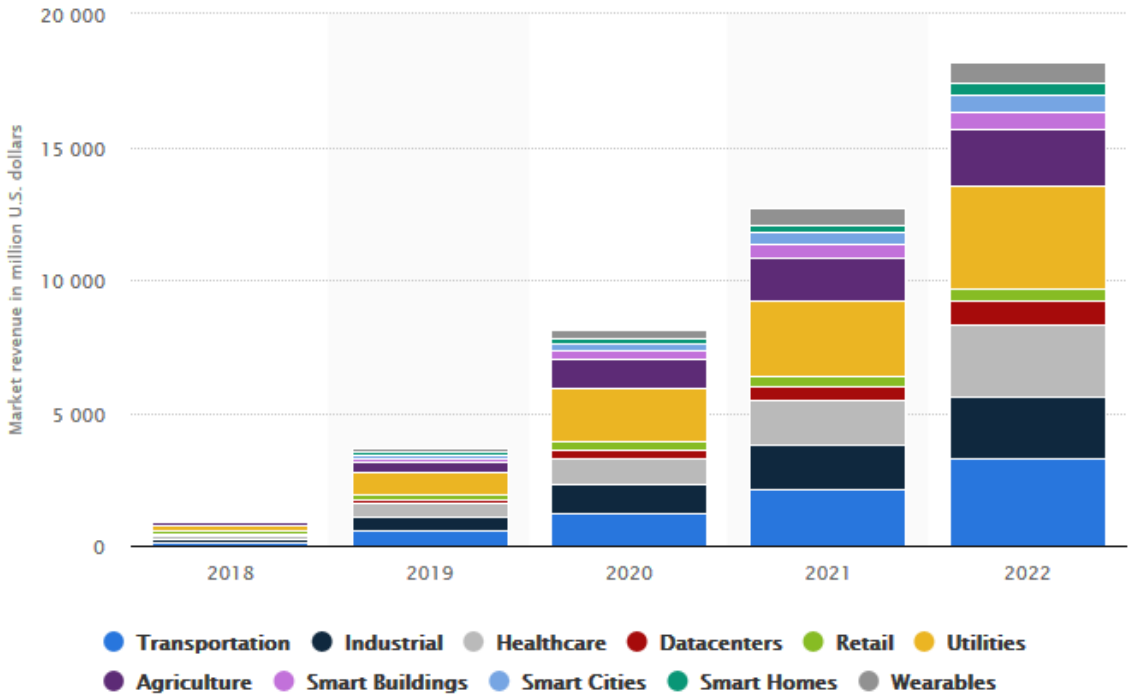

Figure 4-3 Global fog computing market size forecast (2018 to 2022), by vertical [95]

In order to improve daily operations, more and more organisations are moving some of their operations to the edge. This represents a significant growth in the edge computing market, which represented USD 1.7 billion in 2017 and it is expected to reach USD 16.5 billion by 2025, at a CAGR of 32.8% [96]. Only in Europe, edge computing market is expected to grow at a CAGR of 29.3% reaching a value of USD 1.94 billion by 2023 [97]. Digitisation of factory plants is a key factor in this growth, currently 91% of European corporations are investing in this first step towards digital transformation. Based on this premise of transformation, Grand View Research predicts even higher growth where the edge computing market reaches USD 28.84 billion by 2025 at a CAGR of 54% [98].

With the inclusion of edge computing as a major trend for cloud by Gartner [99] in 2018 (see Figure 4-4), the relationship between cloud, fog and edge has become a closed loop circle.

Figure 4-4 Hype Cycle for Cloud Computing, 2018 [99]

4.2. Internet of Things

Definitions of IoT tend to revolve around McKinsey's notion of 'sensors and actuators embedded in physical objects and devices that are linked through wired or wireless networks, often using the same Internet Protocol (IP) that connects the Internet'. To make this simpler, IoT simultaneously refers to the ever-widening network of physical objects that use an IP address for internet connectivity, and the communications that happens between these objects and other Internet-enabled devices and systems, allowing them to exchange data [100]. Nevertheless, IoT is a very wide definition for technologies that encompasses everything from consumer devices to industrial applications. The future of the technology will depend on how these technologies are made scalable. The key to its adoption will be the readiness of the public and/or the enterprise [101].

As IoT grows in importance — Gartner predicts the number of connected things in use will hit 14.2 billion in 2019, and grow to 25 billion by 2021, or 34 billion by 2020, up from 10 billion in 2015 according to Forrester — increasing numbers of previously manual processes will become automated. However, the technology used is not always necessarily easily understandable by their users [101]. The IoT ecosystem includes many technologies that enable consumers, businesses and governments to connect, control and extract value from their connected objects in diverse environments, including manufacturing, agriculture, transportation, smart cities, construction or oil and gas. IoT devices will account for 24 billion of these, while traditional computing devices (e.g. smartphones, tablets, smartwatches, etc.) will comprise a further 10 billion. It is estimated that nearly $6 trillion will be spent on IoT solutions over the next five years.

Interestingly, the more the IoT market matures, the more it fragments; IoT solutions then tend to be absorbed into other markets. As a result of this fragmentation, vendors and service providers alike will coalesce around large enterprise platforms that offer the necessary standardised architecture and stimulate smaller vendors to develop applications. In order to succeed in such a highly fragmented market, consolidation will be imperative. Businesses will endeavour integrate new IoT technologies into their existing software solutions, usually provided by big vendors such as Microsoft, Intel, IBM, Microsoft, SAP and Oracle, among others. System integrators and smaller organisations building IoT applications can benefit from this; integration will also be influenced by standards consolidation [100].

The IoT market remains fragmented, but there is a clear split between providers offering general IoT platforms and insight-centric providers offering applications or end-to-end solutions. The former tend to only gather and visualise data, while the latter tend to be more pain-point focused and can offer greater actionable insights. Buyers looking to invest in an IoT solution, need take sufficient time to make sure that they understand exactly what a vendor is offering and if their product can best serve their business needs. Companies may find that a general IoT platform is a worthwhile investment for their long-term growth strategy, and that they can integrate applications into this platform to solve specific needs later, preferably aided via the original provider's ecosystem. Alternatively, depending on the company's budget, it may be better to only invest money into solutions that specifically address an organisation's needs and has proven outcomes [100].

IDC estimates that worldwide spending on IoT in 2019 will reach $745 billion, with discrete manufacturing having the highest level of investment ($119 billion) whilst process manufacturing will reach $78 billion [102]. This represents an increase from the $646 billion spent in 2016, with annualised growth rate of 15.4%. Fortune Business Insights values the global market at $190.0B in the year 2018 and anticipates that it will reach $1111.3B by 2026. In their opinion, the CAGR will be 24.7% [103].

It is worth noting that the IoT market grew slower than previous forecasts of $772.5 billion by 2018. This is due to a variety of factors, however IDC still predicts that worldwide IoT spending will maintain its double-digit CAGR throughout the 2017-2022 period and will surpass the $1 trillion mark by 2022, two years later than previously expected. In June 2019, the forecasts were updated and IoT spending for 2019 is now expected to reach $726 billion (instead of $745 billion) [102]. Mordor Intelligence

predict that the IoT market will have a CAGR of 21% over the 2019 to 2024 period. The main reasons for this growth rate will be the development of wireless networking technologies, the emergence of advanced data analytics, a reduction in the cost of connected devices and an increase in cloud platform adoption. The retail segment will witness the highest growth whilst the region with highest growth will be Asia [104].

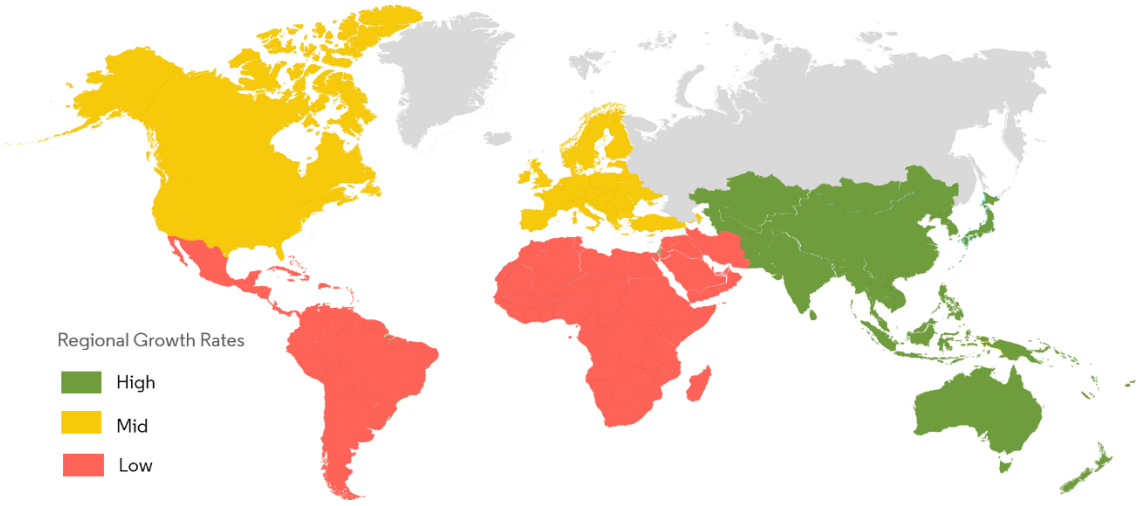

Figure 4-5 IoT Market growth rate by region - 2019 to 2024 (Mordor intelligence [103])

IoT is growing considerably in popularity in commercial companies amongst business decision makers, IT decision makers, and developers, who are driving incorporation of IoT into their businesses. It is worth noting, most of them are satisfied with the results being achieved. As an outcome, the enthusiasm for IoT adoption is growing globally and across industry sectors. Microsoft recently published a report which outlined that 85% of the enterprise IoT decision makers they surveyed indicated they have at least one IoT project in-flight, whether in the learning, proof of concept, purchase, or use phase and many of them reported having one or more projects currently deployed. They report similar results across various countries surveyed including the US, UK, Germany, France, China and Japan and in different sectors (manufacturing, retail/wholesale, transportation, government and healthcare). Adoption rates are projected to increase by 9 points over the next two years, meaning that 94% of businesses will be using IoT by the end of 2021 according to Microsoft [105].

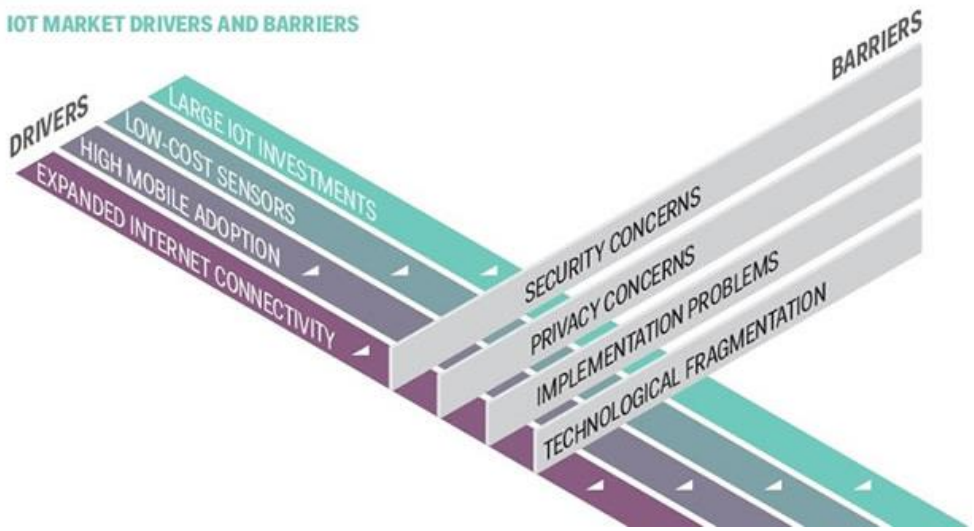

Figure 4-6 IoT Market drivers and barriers (Mordor intelligence [106]).

I-Scoop presents four drivers for the adoption of IoT by the market, namely the expansion of Internet connectivity, the high adoption of mobile and technologies such as low-power wide-area network (LPWAN), the persistent decline in the cost of sensors and the growth of large IoT investments. As IoT technology adaptation increases, security is emerging as a significant potential barrier to further adoption. Security should not be an afterthought and solution providers have started integrating improved security capabilities into their products. IoT actually presents more weaknesses than would be normally be found in standard IT systems. Security by design must take centre stage in an integrated way across technology stack. The second IoT barrier relates to privacy concerns, which also needs a 'by design' approach intrinsically connected to security. The final two IoT market barriers are implementation problems and technological fragmentation [106]. The first three items represent the main challenges and opportunities for improving the mF2C framework. Improvements in the security and privacy aspects of developments would make mF2C even more competitive in the future, reducing possible concerns of both users and clients. Ease of implementation is also key in the adoption of the framework by the industry, and will make or break any implementation based on the outputs of the project.

The landscape of IoT vendors and companies has grown increasingly complex over the last number of years. A comprehensive picture of vendor landscape is presented in Figure 4-7.

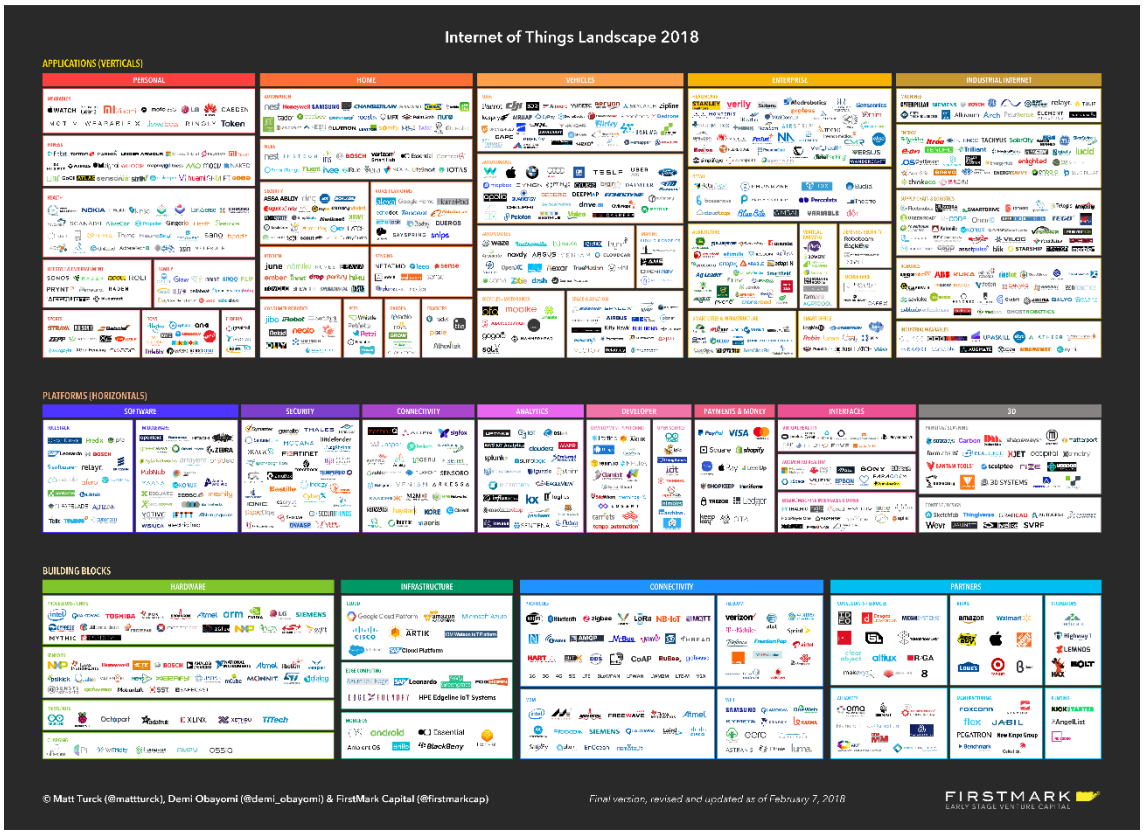

Figure 4-7 IoT landscape 2018 [McKinsey].

McKinsey has identified the market opportunities in the IoT sector depending on the potential, the market growth and the technology maturity, as represented in Figure 4-8. It is worth noting that devices in general are at sufficient technology readiness level to be commercialised and the market size is relatively large, however the largest market (i.e. business applications) is still in need of mature technologies and therefore represents a huge commercial opportunity.

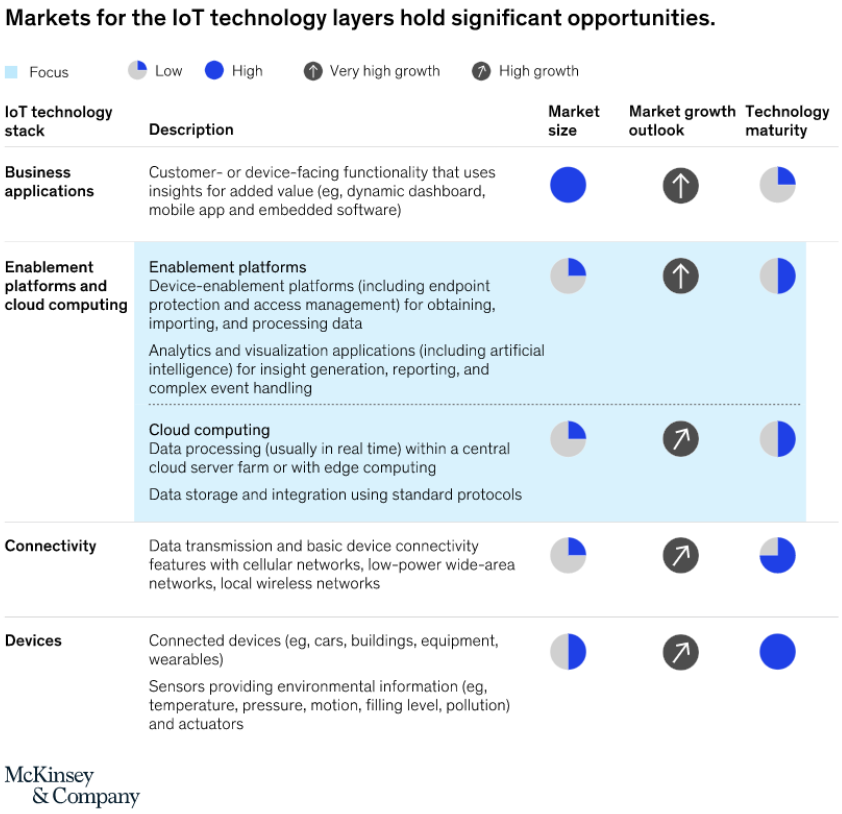

Figure 4-8 Market opportunities for the IoT sector [107].

Regarding intellectual property in IoT, Figure 4-9 provides a representation of the patent landscape in 2018 based on research by Relecura [108]. It is interesting to note that Avago technologies leads in the number of patent applications, instead of one of the traditional multinationals such as Samsung.

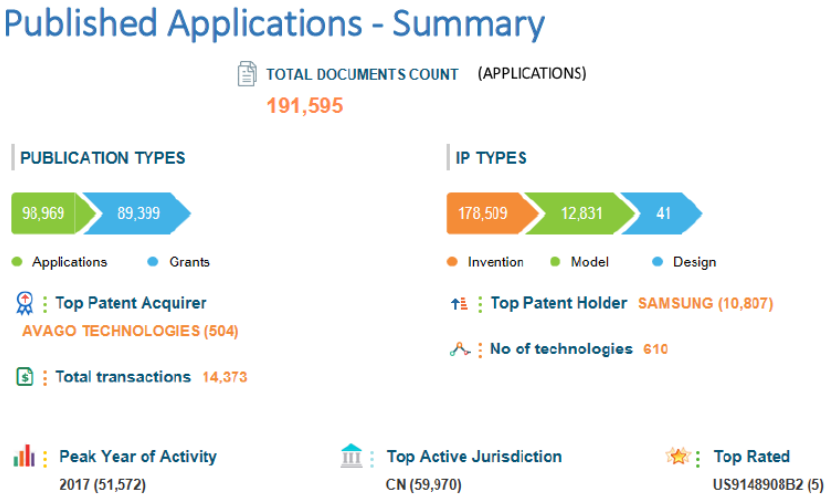

Figure 4-9 IoT Patent landscape [108].

Finally, one growing and specific niche of the IoT market is Industrial IoT, which represents a significant growing opportunity. The market was valued at $176B for 2018, growing at a CAGR of 27% with a potential market value of $934B by 2025 according to a KeyBanc Capital markets study. Worldsensing

has focused on this market for several years and is growing as a key partner in the construction, mining and utilities sectors.

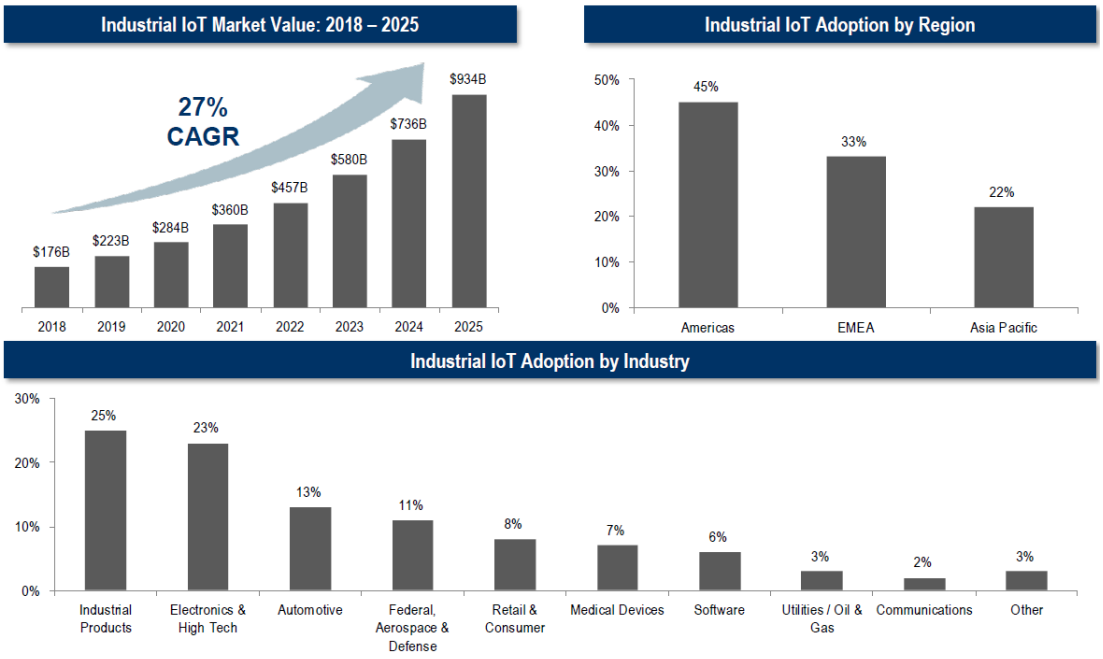

Figure 4-10 Industrial IoT market overview [109].

4.3. Big Data and IoT

Technologies such as Big Data and the Internet of Things (IoT) are shaping our lives and disrupting traditional businesses. The IoT market is exhibiting continuous growth as billions of devices, services and systems become connected, mainly driven by ubiquitous and cheaper sensors that convert physical measurements to digital data. It is predicted that IoT devices worldwide will generate 90 zettabytes (90.000.000.000 Terabytes) of data by 2025 [110].

Data is sent by sensors to centralised Big Data platforms that aggregate, process, store, analyse and visualise this data in order to create insights and improve the operational efficiencies of processes. However centralised architectures increase latency due to data transport overheads such as a saturation of the network bandwidth.

These shortcomings are leading to the evolution of computing platforms from centralised architectures to distributed or decentralised architectures with a focus on fog computing and AI capabilities closer to sources of data, such as edge centric computing. In fact some related technologies are already available such as ML inference for AWS IoT Greengrass, AWS DeepLens or Google's Edge TPU.

The edge computing model is well suited to IoT applications because of several key benefits, including near real-time analysis of data, lower costs related to operations and data management, reduced data transmission to a cloud backend reducing network overheads.

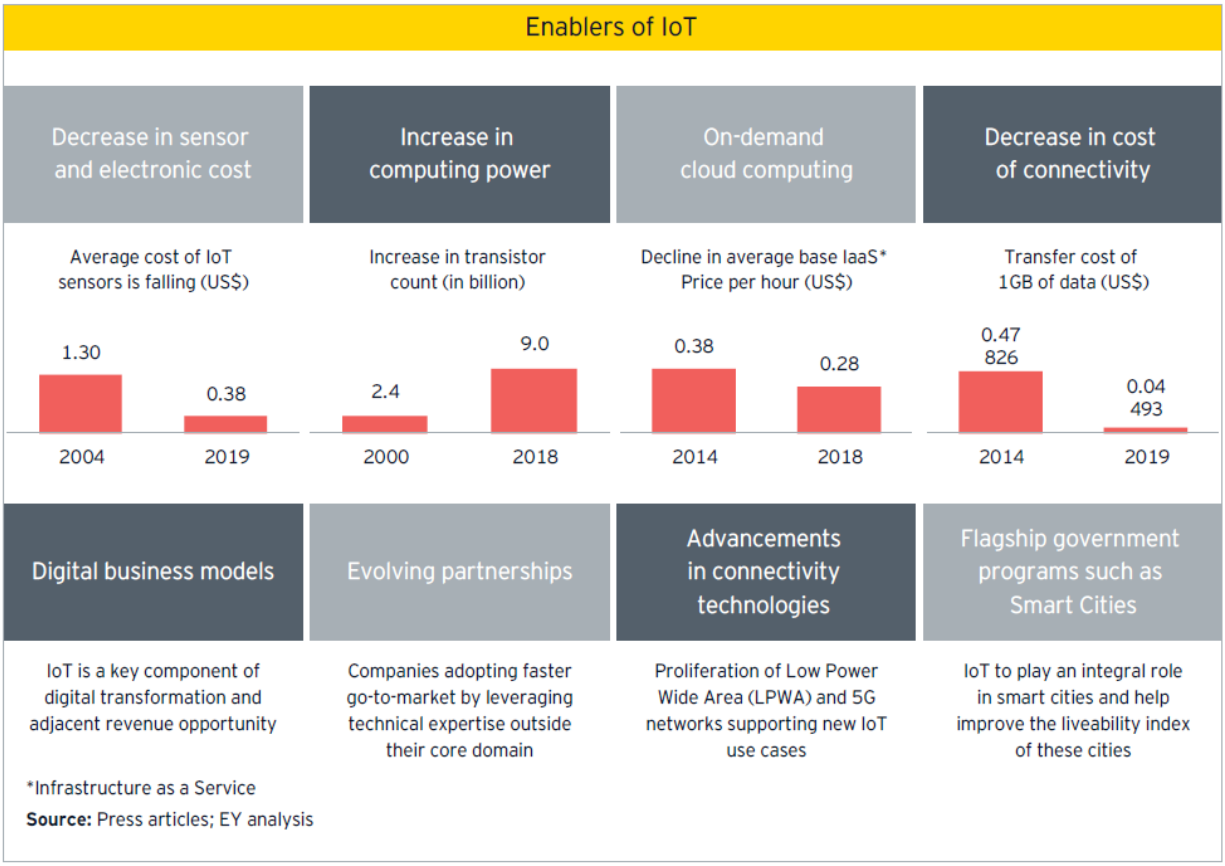

Figure 4-11 Future of IoT, 2019 [110]

The use of data is now essential in the decision-making processes of companies and institutions. Current technologies allow ever greater diffusion of the "datafication" processes, converting anything (films, books, vocal messages, body movements, etc.) into digital formats. Sources can be found in any device, sensor, operating system, search engine or social network.

Big data represents the key productive factor in a data-driven economy; there are several areas, both private and public, where the use of analysis techniques of big data has allowed the creation of new services, improvements to existing ones, innovative production and distribution processes, enabling products and services (even non-digital) to provide a better respond to the needs of consumers and citizens [111].

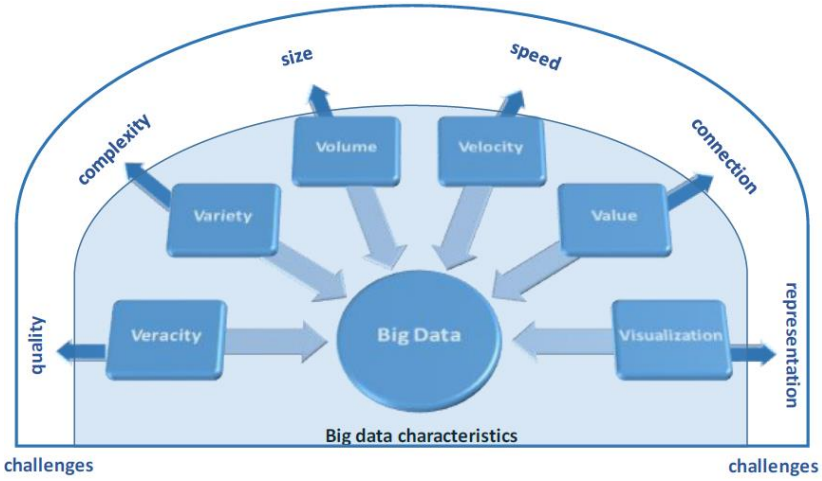

Figure 4-12 Big Data interim report in the context of joint inquiry on "Big Data", 2018 [111]

In addition to the undisputed economic and social benefits derived from the advent of the data-driven economy supported by Big Data systems, there are some risk factors which could potentially result in market failure such as information asymmetries and market power positions. The emergence of new and possibly discriminatory practices, among which those linked to price are the most widespread. Price discrimination coupled with modern online profiling techniques could be detrimental to specific categories of users (consumers, workers, publishers, etc.) with possible discrimination, even involuntarily, to differences in the population based on ethnicity, race, sexual orientation, and health condition, thus violating the privacy of data owners.

In terms of the technological challenges created by the IoT evolution, current technologies present some bottlenecks such as poor scalability, security issues and difficulties with installation, fault tolerance, maintenance and low performance. Hence, we need to adapt these technologies to provide solutions to other problems. IoT and Big Data interplay with each other where IoT plays the role of a data source unit.

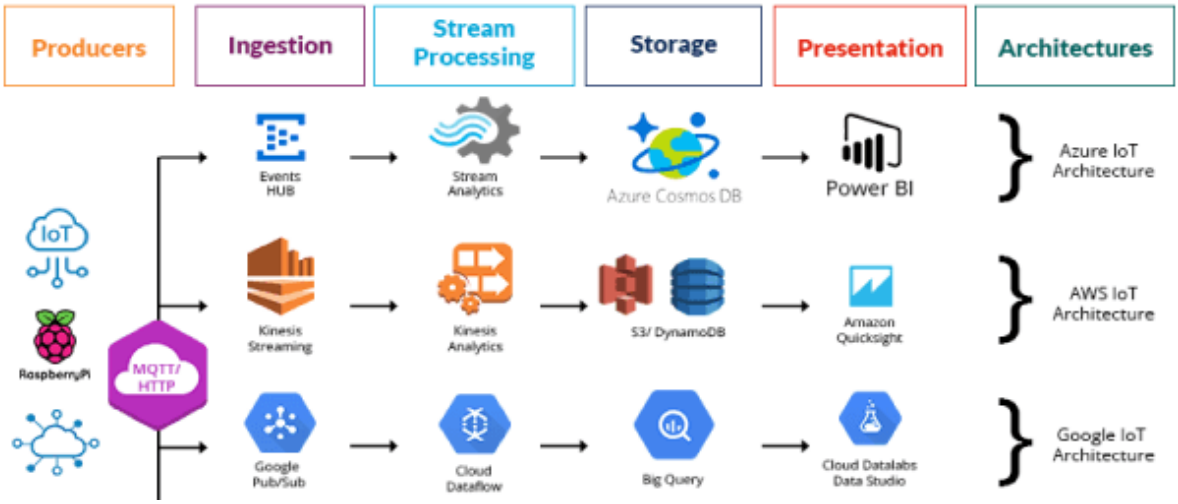

Figure 4-13 Relationship between IoT, Big Data and Cloud Computing, 2019 [112]

There is a strict relationship between IoT and Big Data. IoT is an opportunity to streamline operations in many sectors by enabling interactions between machines and humans (M2H) and between machines (M2M). In most cases sensor-generated data are fed to the big data system for analysis and reporting which is the main point of interaction between the two technologies as show in Figure 4-13. The intersection of the IoT and Big Data has created new IT challenges regarding data storage integration, and analytics. However, it has created significantly more opportunities than challenges [112].

We often quote the power of data, comparing them to the wealth of oil owners. But unlike oil, data are unlimited and are created exponentially by the action of machines and people. A datum always has a history, it is not interchangeable: each one is different from the other. Data are immaterial and as a result intangible assets, therefore in order to obtain value from their personalisation, data must process in order to extract information and obtain knowledge. Data must be used as soon as possible and as a consequence primarily by those who produce them [113].

IoT's big data promises to help companies understand customer needs, market dynamics, and strategic issues with unmatched precision. But in pursuing of this goal, organisations will amass previously unimaginable quantities of information. Data marketplaces offer them an innovative way to turn some of that data into cash and reap the benefits that will accrue from building a self-reinforcing ecosystem, enabling crowdsourcing, supporting interoperability, satisfying customer data needs, and improving data quality [114].

4.4. Security Trends

Industry security solutions continue to focus on intrusions, malicious insiders, ransomware, as described in the earlier deliverables. Machine learning methods are becoming more sophisticated and able to distinguish, for example, a compromised account or an intrusion into a resource through deviations from normal behaviour patterns. However, while some steps can (and should) be taken to automatically react to an incident, every serious incident should still be investigated by a human, and, if necessary, the automated security system should be updated.

In general, and in SaaS applications, the end user is often the weakest link, and in corporate environments users are often required to go through cybersecurity awareness training. Much of the training is generic – awareness of public Wi-Fi, social engineering, malicious email attachments, phishing attacks, although the most sophisticated attacks are difficult to detect even for security-aware users. Other training is specific to the corporate environment – the type of edge devices used, what to do if the user detects an incident, and security related policies such as requirements for multifactor authentication (MFA) or hardware assisted security (e.g. “dongles.”)

In development environments (PaaS and IaaS in clouds), industry trends continue to promote “devsecops” – where software development moves to develop practices (including containerised, continuous deployment/integration), security needs to be integrated into these processes, from Secure by Design and the Privacy by Design mentioned earlier. Moreover, industry recommendations are that developers identify any open source products used in developing their applications, the thinking being that open source products are less rigorously patched and maintained compared to commercially procured equivalents, and might “contaminate” the application if the licence on the open source components does not allow proprietary reuse. However, in an open source project, this advice is not very useful – reliance on external components, possibly with other licences, is not exactly new. Whenever mF2C has relied on an external component, the project has always assessed the maturity, sustainability, licence, and other implications of inbounding it into the projects codebase.

In areas where end users cannot be trained directly in cybersecurity awareness – such as people running an app on a mobile phone, connecting to an airport’s free Wi-Fi – the app and the supporting software infrastructure needs to promote good security practices and good usability. Server side monitoring can detect whether the user is accessing services from an unusual environment, and require additional verification. Additionally, it makes sense to plan for end user devices being compromised.

4.5. Digital Business

The appearance of digital businesses has directly impacted on traditional corporate strategies. Thus, organisations have had to adapt to this new disruptive wave in order to remain competitive. Innovation is now the basis for new businesses where improved productivity and customer experience are they factors for success.

However, there is still a long way to go until organisations fully adopt digital businesses as many of them just translate ‘digital’ into ‘technology’. But ‘value’ is however the most accurate translation. Digital transformation is about creating value through defeating barriers for entering into new markets while enhancing existing businesses at the same time.

One of the most challenging scenarios for organisations it to change their traditional mind-set into a fully digital one, including innovation at all levels and relying on technology as the key driver for transformation.

Taking all of this into account, mF2C has developed a set of business models (including digital (B2B) and traditional ones (B2C)) that will be further analysed and combined in order to find the most suitable ones for the sustainability of mF2C.

It is important to highlight, that traditional strategies with one single business model for one product cannot be applied nowadays to new solutions in the market. This consideration will also be taken into account in order to build successful models for mF2C results.

4.6. Key Takeaways

The key areas of focus for the project within the context of relevant business trends are as follows:

- Growth in the global cloud compute market continues and is being driven by factors such as grow in M2M communications, efforts to improve decision making capabilities in businesses, efforts to reduce IT costs, faster service delivery and increasing utilisation of distributed and real-time data analytics to support operational efficiencies.
- The global cloud computing market is growing at a CAGR of 18% with the market expected to reach USD 623.3 billion by 2023, an increase of USD 351.3 from 2018. While in Europe this growth is driven mainly by SMEs, looking for customised offerings specifically adapted to their own needs.
- M2M communications, decision-making capabilities, reduced operating costs, distributed data analytics and real-time tracking are the key factors for adopting fog computing solutions by end-use industries. The smart manufacturing sector is currently showing the highest growth rate, a CAGR of 60%, while the smart grid segment has the highest market share (over 20%).
- The landscape of the cloud compute market is also changing rapidly with increases utilisation of decentralised edge-based approaches in order to improve daily operation. This trend represents a significant growth in the edge computing market, which was USD 1.7 billion in 2017 and is expected to reach USD 16.5 billion by 2025.
- The inclusion of edge computing as a major trend for cloud by Gartner in 2018 demonstrates that the relationship between cloud, fog and edge has become a closed loop. All the major cloud platforms include edge and IoT features. Being mostly containerised, mF2C should in theory be readily deployable across these platforms, but an obvious next step would be an actual instantiation across the prominent commercial providers. It would also be important to document how to do so – in order to lower the barrier to adoption for someone to utilise mF2C on their own resources. In particular, when these platforms provide their own “marketplace,” mF2C resources could be made available through those marketplaces – as mF2C is open source, a potential business model would be the provision of paid support.
- The edge computing model exhibits strong synergies with IoT applications due to key benefits such near real-time analysis of data, lower data management costs and reduced network traffic footprints.
- As the IoT market matures and diversifies, fragmentation will become an increasing important problem. Vendors and service providers will potentially coalesce on large enterprise platforms that offer the necessary scalable and standardised architecture and will also enable smaller vendors to develop applications.
- IDC still predicts that worldwide IoT spending will maintain its double-digit CAGR throughout the 2017-2022 period and will surpass the $1 trillion mark in 2022, two years later than previously expected
- IoT continues to be key driver in the underlying data for Big Data. It is predicted that IoT devices will generate 90 zettabytes of data by 2025 on worldwide basis.
- Big data represents the key productive factor in a data-driven economy where the use of analysis techniques of big data has allowed the creation of new services, drives improvements to existing ones, innovative production/distribution processes and enabling products and services to be more responsive to customer needs.
- Data marketplaces offers an innovative way to monetise data and to reap the benefits that accrue from building a self-reinforcing ecosystem, enabling crowdsourcing, supporting interoperability, satisfying customer data needs, and improving data quality.

- Security continues to improve “solutions” for automated detection and reaction, although there is no perfect solution and no “one size fits all” solution. Humans are, and will likely always be, needed to ultimately assess the incident. End users of the applications are of course also human and should be trained in cybersecurity awareness when possible. An obvious next step for mF2C would be to adapt to one or more industry cloud-integrated edge security “platforms,” in order to lower the barrier to reusing the outcomes from mF2C. We had originally intended to demonstrate one or more of these platforms within the project, but due to time and development overhead challenges as it would require re-engineering of multiple components beyond the security components this was not possible.
- Digital businesses have directly impacted on traditional corporative strategies by driving them to adapt in order to remain competitive. Innovation is now the basis for new businesses where previously improved productivity and customer experience were the primary success factors.
- Cloud and edge computing play a critical role in digital transformation as 91% of European corporations are investing in digitisation. Booming economies are increasingly demanding cloud services, but there is still room for further expansion. mF2C’s proposed approach for connecting cloud and edge computing to dynamically adapt to user needs supports digitalisation of a variety of businesses remains valid.
- Market growth is driven by the increased adoption of IoT, the growth of smart cities and the penetration of e-commerce. These, together with reductions in free trade, interest rates and data localisation could hinder in growth fog-to-cloud as proposed by mF2C.
- The current market situation and increased demand for user-adapted services allows both the entry of new players and consolidation of existing ones, creating new market opportunities for services built on top of mF2C.
- The current jungle of standards (cloud-, fog-, edge) is hampering the adoption of solutions, such as mF2C. Different bodies such as the OpenFog Consortium are pushing for more comprehensive standards which can enable acceleration within the market, bypassing some of the current adoption pain points.
- The emergence of digital business has directly impacted on traditional corporate strategies. This, together with increasing adoption of cloud and edge services, is driving evermore challenging scenarios including innovation at all levels and relying on technology as the key driver for transformation. mF2C aims to support organisations in this transformation with a set of proposed business models which are domain-independent in order to add real value to their daily operations.

5. Conclusions

The technology trends previously described in D2.2 remain valid and progress in the interim period is reflected in this deliverable version. In addition, new areas of relevance to the mF2C project are described.

The scientific trends reviewed in chapter 2 show Cloud and Fog computing as being conceptually similar, but the challenges faced in designing resource management solutions for fog computing systems continue to be heterogeneity, dynamicity, geo-distribution, and multiple owners of the devices comprising that fog system. There is an increasing focus on the joint management of communications and computing resources in fog-to-cloud environments. These management approaches are increasing leverage AI and Deep Learning techniques to improve system performance, resource utilisation efficiency and QoS. In mF2C the QoS providing component utilises Deep Reinforcement Learning to select the appropriate devices to execute a service, excluding devices that can potentially cause SLA violations. This approach is reflective of other research to address the challenges of resource allocation in dynamic environments such as MEC. In addition, new novel approaches to resource management based on blockchain have also emerged. Task scheduling and offloading mechanisms due to issues related to device mobility, limitations in devices energy budgets and the impact of the network on the performance of fog-to-cloud management framework remains an area of active focus.

A variety of generic frameworks which provide support for machine learning have gained in popularity such as Tensorflow, Caffe and Scikit. These environments provide Python interfaces which simplifies adoption and provides support for various forms of parallelism. These frameworks are specifically designed to address big data use cases. As a result they are not suitable for adoption by mF2C as they do not address the composition of distributed applications, edge deployments and execution and transparent cloud offloading in response to load and constraints requirements. Instead COMPs has been adopted which provides greater levels of flexibility and is compatible with the Function as a Service paradigm.

From a security perspective the techniques used to detect and prevent cyber-attacks also continue to evolve with ML/AI approaches continuing to attract attention. Of particular interest in the context of IoT are anomaly based approaches which utilise deep learning for the detection of distributed attacks with previously unseen or slight mutations of known attack patterns. Some commercial solutions are now emerging such as FortiGuard AI which utilise this form of approach. As these solutions provide increasing performance and detection accuracy, their adoption and usage promises to alleviate the workload associated with the analyses of (possibly vast amounts of) false alarms. However, significant challenges remain in order to realise fully automated operations that do not require human supervision. This remains an area of active research. Finally, in the security domain data protection remains an on-going battle. Unfortunately, approaches to date have not worked well as highlighted by the regular reports of major privacy violations. In order to address these shortcomings privacy engineering has emerged as a new field of research and practice. This new field aims to bridge the legal, technical and engineering approaches to support engineers in systematically identifying and addressing privacy and data protection concerns during a development lifecycle.

The technology trends reviewed in chapter 3 included both Cloud and IoT Management tools which address resource management and service scheduling, interconnection of sensing devices, data, and applications over the Web to address interoperability. These tools are critical to addressing key challenges such as interoperability, and scalability associated with large scale IOT deployments in various vertical domains. From an IT technology perspective database vendors are starting to evolve their offerings in order to provide solutions which are more suitable for deployment on computation and resource limited edge devices. These developments will provide better supports for analytics at

the edge and to accommodate dynamic behaviours due to devices joining and leaving in an ad-hoc manner. The evolvement and adoption of container related technologies has continued apace. In parallel, Kubernetes and related solutions such as Openshift have emerged as the dominant container orchestration approaches for companies and cloud solution providers such as Amazon and Microsoft. From a business perspective container related technologies offer significant opportunity with the market expected to double over the next five years to USD 4.9 billion by 2023. The last twelve months has seen some interesting developments in the technology standards space. The amalgamation of the OpenFog Consortium with the IIC and subsequent publication of their IOT Reference architecture is reflective of momentum around a standard reference architecture. More broadly the Cloud/Fog technology ecosystem has seen a variety of initiatives which in the longer term will lead to standard technology building blocks to implement emerging standards based architectures.

Finally, an updated review of the key business trends is provided in chapter 4. This review reflects research from various business analysts groups and helps to contextualise the potential business value that the mF2C framework could generate. Growth in the global cloud compute market continues and is being driven by factors such as grow in M2M communications, efforts to improve decision making capabilities in businesses, efforts to reduce IT costs, faster service delivery and increasing utilisation of distributed and real-time data analytics to support operational efficiencies. These factors are reflective of cloud computing adoption by the smart manufacturing sector which is exhibiting the highest growth rate with a GAGR of 60% which is significantly higher than the overall cloud compute market which has a GAGR of 18%. Overall the cloud compute market still represents a significant business opportunity with a projected value of USD 623.3 billion by 2023. However, the landscape of the cloud compute market is also changing rapidly with increased utilisation of decentralised edge based approaches in order to improve daily operation. This trend represents a significant growth in the edge computing market, which was USD 1.7 billion in 2017 and is expected to reach USD 16.5 billion by 2025. In parallel to the emergence of the edge compute market, the IoT is exhibiting signs of maturity. However, with increasing levels of maturity and diversity, fragmentation is becoming a growing problem. Vendors and service providers will potentially coalesce on large enterprise platforms that offer the necessary scalable and standardised architectures. Standardisation of solution architectures and technology building blocks should enable smaller vendors to develop applications for this lucrative market. IDC predicts that worldwide IoT spending will maintain its double-digit CAGR throughout the 2017-2022 period and will surpass the $1 trillion mark in 2022. Growth in IOT will continue to drive growth in big data. It is predicted that IoT devices will generate 90 zettabytes of data by 2025 on a worldwide basis. Big data represents the key productive factor in a data-driven economy where the use of analysis techniques of big data has allowed the creation of new services based on the monetisation of data and will enable products and services to be more responsive to customer needs. The emergence of these new digital business has been highly disruptive to traditional corporate strategies driving them to adapt in order to remain competitive. Innovation is now the basis for new businesses where previously improved productivity and customer experience were the primary success factors. Increased adoption of cloud and edge services, is driving evermore challenging scenarios including innovation at all levels and relying on technology as the key driver for transformation. mF2C aims to support organisations in this transformation with a set of proposed business models, domain-independent, to add real value to their daily operations.

The growing range of security threads remains an area of significant concern for most businesses. Automation of security responses to threats is an area of considerable focus, however every serious incident still requires investigated by a human, and, if necessary, the automated security system should be updated. In development environments the industry trend is the continued promotion of “devsecops”. Here software development moves to develop practices (including containerised, continuous deployment/integration). Security is integrated into these processes in the form of Secure by Design and Privacy by Design.

This deliverable provides a comprehensive understanding of the current scientific, technical and business trends in Fog, Edge and Cloud computing and their relevance to the mF2C. This information and associated insights has influenced the final project architecture. Awareness of these trends has helped to focus the efforts of the project on the problem areas requiring new solutions and approaches. Existing codebases have been effectively utilised to accelerate the development of the Platform Manager and Agent Controller functional entities. This deliverable also highlights and confirms the significant business opportunities which continue to exist and will grow over the coming years for the mF2C Framework. Finally, the key takeaways at the end of each section highlight areas for additional research and development in order to realise future versions of mF2C which reduce the barriers to adoption and support new and emerging use cases.

6. References

[1] X. Masip-Bruin, E. Marín-Tordera, G. Tashakor, A. Jukan and G.-J. Ren, "Foggy clouds and cloudy fogs: a real need for coordinated management of fog-to-cloud computing systems," *IEEE Wireless Communications,* vol. 23, no. 5, pp. 120-128, 2016.

[2] mF2C Consortium, "D2.1 Tracking Scientific, Technology and Business Trends (Version 1)," 7th April 2017. [Online]. Available: https://www.mf2c-project.eu/d2-1-m3/.

[3] mF2C Consortium, "D2.2 Tracking Scientific, Technology and Business Trends (Version 2)," 8th October 2018. [Online]. Available: https://www.mf2c-project.eu/d2-2-tracking-scientific-technology-and-business-trends-version-2/.

[4] Y. Zhou, L. Tian, L. Liu and Y. Qi, "Fog computing enabled future mobile communication networks: a convergence of communication and computing," *IEEE Communications Magazine,* vol. 57, no. 5, pp. 20-27, 2019.

[5] Q. Lu, T. Han and N. Ansari, "Joint Radio and Computation Resource Management for Low Latency Mobile Edge Computing," in *IEEE Global Communications Conference (GLOBECOM)*, Abu Dhabi, United Arab Emirates, 2018.

[6] S. Jošilo and D. Gyorgy, "Joint Management of Wireless and Computing Resources for Computation Offloading in Mobile Edge Clouds," *IEEE Transactions on Cloud Computing,* 2019.

[7] Y. Dai, D. Xu, S. Maharjan, G. Qiao and Y. Zhang, "Artificial intelligence empowered edge computing and caching for internet of vehicles," *IEEE Wireless Communications,* vol. 26, no. 3, pp. 12-18, 2019.

[8] Y. Sun, M. Peng and S. Mao, "Deep Reinforcement Learning-Based Mode Selection and Resource Management for Green Fog Radio Access Networks," *IEEE Internet of Things Journal,* vol. 6, no. 2, pp. 1960–1971, 2018.

[9] D. Zeng, L. Gu, S. Pan, J. Cai and S. Guo, "Resource Management at the Network Edge: A Deep Reinforcement Learning Approach," *IEEE Network,* vol. 33, no. 3, pp. 26-33, 2019.

[10] Z. Xiong, Y. Zhang, D. Niyato, P. Wang and Z. Han, "When mobile blockchain meets edge computing," *IEEE Communications Magazine,* vol. 56, no. 8, pp. 33-39, 2018.

[11] Z. Xiong, S. Feng, D. Niyato, P. Wang and Z. Han, "Optimal pricing-based edge computing resource management in mobile blockchain," in *2018 IEEE International Conference on Communications (ICC)*, 2018.

[12] N. C. Luong, Z. Xiong, P. Wang and D. Niyato, "Optimal auction for edge computing resource management in mobile blockchain networks: A deep learning approach," in *2018 IEEE International Conference on Communications (ICC)*, 2018.

[13] Elastic Project, "Elastic - A Software Architecture for Extreme-Scale Big Data Analytics in Fog Computing Ecosystems," 2019. [Online]. Available: https://elastic-project.eu/. [Accessed 26th July 2019].

[14] R. Altarawneh, C. Strong, L. Remis and P. Munoz, "Navigating the Visual Fog: Analyzing and Managing Visual Data from Edge to Cloud," in *HotEdge '19*, Renton, WA, US, 2019.

[15] L. Remis, V. Gupta, C. Strong and R. Altarawneh, “VDMS: An Efficient Big-Visual-Data Access for Machine Learning Workloads,” in *Conference in Neural Information Processing Systems*, Montreal, Canada, 2018.

[16] A. E. Elgazar, M. Aazam and K. A. Harras, “SMC: Smart Media Compression for Edge Storage Offloading,” in *HotEdge '19*, Renton, WA, US, 2019.

[17] A. Elgazar, M. Aazam and K. Harras, “Edgestore: Leveraging edge devices for mobile storage offloading,” in *IEEE International Conference on Cloud Computing Technology and Science (CloudCom)*, Nicosia, Cyprus, 2018.

[18] C. S. Meiklejohn, H. Miller and Z. Lakhani, “Towards a Solution to the Red Wedding Problem,” in *HotEdge '18*, Boston, MA, US, 2018.

[19] M. Shapiro, N. Preguiça, C. Basquero and M. Zawirski, “Conflice-Free Replicated Data Types,” in *Stabilization, Safety, and Security of Distributed Systems (SSS 2011). Lecture Notes in Computer Science*, vol. 6976, X. Défago, F. Petit and V. Villian, Eds., Berlin, Heidelberg, Springer, 2011.

[20] S. H. Mortazzavi, B. Balasubramanian, E. de Lara and S. P. Narayanan, “Pathstore: A Data Storage Layer for the Edge,” in *16th Annual Conference on Mobile Systems, Applications, and Services*, Munich, Germany, 2018.

[21] A. Aral and T. Ovatman, “A Decentralized Replica Placement Algorithm for Edge Cloud,” *IEEE Transaction of Network and Service Management,* vol. 15, no. 2, pp. 516-529, 2018.

[22] C. Guerrero, I. Lera and C. Juiz, “Optimization Policy for file replica placement in fog domains,” *Concurrency and Computation Practice and Experience,* 2019.

[23] C. Mahmoudi, F. Mourlin and A. Battou, “Formal definition of edge computing: An emphasis on mobile cloud and IoT,” in *3rd International Conference on Fog and Mobile Edge Computing (FMEC)*, Barcelona, Spain, 2018.

[24] S. Kumar Monga, S. K R and Y. Simmhan, “ElfStore: A Resilient Data Storage Service for Federated Edge and Fog Resources,” in *IEEE International Conference in Web Services (ICWS)*, Milan, Italy, 2019.

[25] D. Vasconcelos, V. Severino, J. Neuman, R. Andrade and M. Maia, “Bio-Inspired Model for Data Distribution in Fog and Mist Computing,” in *42nd Annual Computer Software and Applications Conference (COMPSAC)*, Tokyo, 2018.

[26] B. G. Chun, S. Ihm, P. N. M. Maniatis and A. Patti, “Clonecloud: Elastic execution between mobile device and cloud,” in *Proceedings of the Sixth Conference on Computer Systems (EuroSys '11)*, Salzburg, Austria, 2011.

[27] R. Kemp, N. Palmer, T. Kielmann and H. Bal, “Cuckoo: A Computation Offloading Framework for Smartphones,” in *Mobile Computing, Applictions, and Services MobiCASE 2010. Lecture Notes of the Institute for Computer Scieneces*, Berlin, Heidelberg, Springer, 2010, pp. 59-79.

[28] S. Kosta, A. Aucinas, P. Hui, R. Mortier and X. Zhang, “ThinkAir: Dynamic resource allocation and parallel execution in the cloud for mobile code offloading,” in *2012 Proceedings IEEE INFOCOM*, Orlando, Florida, 2012.

[29] K. Hong, D. Lillethum., U. Ramachandran, B. Ottenwälder and B. Koldehofe, “Moble Fog: A Programming Model for Large-Scale Appliction on the Internet of Things,” in *Proceedings of the Second ACM SIGCOMM Workshops on Mobile Computing*, Hong Kong, China, 2013.

[30] M. Adadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, M. Kudlur, J. Levenberg, R. Monga, S. Moore, D. G. Murray, B. Steiner, P. Tucker, V. Vasudevan, P. Warden, M. Wicke, Y. Yu and X. Zheng, "TensorFlow: A System for Large-Scale Machine Learning," in *12th (USENIX) Symposium on Operating Systems Design and Implementation (ODSI '16)*, Savannah, GA, 2016.

[31] R. Al-Rfou, G. Alain, A. Alamahairi, C. Angermueller, D. Bahdanau, N. Ballas, F. Bastien, J. Bayer, A. Belikov, A. Belopolsky and Y. Bengio, "Theano: A Python framework for fast computation of mathematical expressions," 9th May 2016. [Online]. Available: https://arxiv.org/abs/1605.02688. [Accessed 16th September 2019].

[32] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama and T. Darrell, "Caffe: Convolutional Architecture for Fast Feature Embedding," in *Proceedings of the 22nd ACM International Conference on Multimedia*, Orlando, Florida, 2014.

[33] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga and A. Lerer, "Automatic differentation in PyTorch," in *NIPS 2017 Autodiff Workshop*, Long Beach, California, 2017.

[34] Keras, "Keras: The Python Deep Learning library," 2019. [Online]. Available: https://keras.io. [Accessed 16th September 2019].

[35] X. Meng, J. Bradley, B. Yavuz, E. Sparks, S. Venkataraman, D. Liu, J. Freeman, D. Tsai, M. Amde, S. Owen, D. Xin, R. Xin, M. J. Franklin, R. Zadeh, M. Zaharia and A. Talwalkar, "MLlib: Machine Learning in Apache Spark," *Journal of Machine Learning,* vol. 17, no. 1, pp. 1235-1241, 16th April 2016.

[36] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot and É. Duchesnay, "Scikit-learn: Machine Learning in Python," *Journal of Machine Learning,* vol. 12, pp. 2825-2830, 12th October 2011.

[37] N. Marz and J. Warren, Big data : principles and best practices of scalable real-time data systems, New York: Manning Publications Co., 2015.

[38] E. Jonas, J. Schleier-Smith, V. Sreekanti, C.-C. Tsai, A. Khandelwal, Q. Pu, V. Shankar, J. Carreira, K. Krauth, N. Yadwadkar, J. E. Gonzalez, R. A. Popa, I. Stoica and D. A. Patterson, "Cloud Programming Simplified: A Berkeley View on Serverless Computing," 9th February 2019. [Online]. Available: https://arxiv.org/abs/1902.03383. [Accessed 16th September 2019].

[39] A. Pérez, S. Risco, D. M. Naranjo, M. Caballer and G. Molto, "On-premises Serverless Computing for Event-Driven Data Processing Applications," in *IEEE International Conference on Cloud Computing (CLOUD 2019)*, Milan, Italy, 2019.

[40] G. Apruzzese, M. Colajanni, L. Ferretti, A. Guido and M. Marchetti, "On the effectiveness of machine and deep learning for cyber security," in *10th International Conference on Cyber Conflict (CyCon)*, Tallinn, Estonia, 2018.

[41] A. Polyakov, "Machine Learning for Cybersecurity," Medium, 4th October 2018. [Online]. Available: https://towardsdatascience.com/machine-learning-for-cybersecurity-101-7822b802790b. [Accessed 22nd August 2019].

[42] S. Rezaei and X. Liu, “Deep Learning for Encrypted Traffic Classification: An Overview,” *IEEE Communications Magazine,* pp. 76-81, 13th May 2019.

[43] W. Li, J. Jin and J.-H. Lee, “Analysis of Botnet Domain Names for IoT Cybersecurity,” *IEEE Access,* vol. 7, pp. 94658-94665, 8th July 2019.

[44] A. Abeshu and N. Chilamkurti, “Deep Learning: The Frontier for Distributed Attack Detection in Fog-to-Things Computing,” *IEEE Communications Magazine,* vol. 56, no. 2, pp. 169-175, 13th February 2018.

[45] R. Priyadarshini and R. K. Barik, “A deep learning based intelligent framework to mitigate DDoS attack in fog environment,” *Journal of King Saud University - Computer and Information Sciences,* 24th April 2019.

[46] C. Li, Y. Wu, X. Yuan, Z. Sun, W. Wang, X. Li and L. Gong, “Detection and defense of DDoS attack-based on deep learning in OpenFlow-based SDN,” *International Journal of Communications Systems,* vol. 31, no. 5, 25th March 2018.

[47] A. Shenfield, D. Day and A. Ayesh, “Intelligent intrusion detection systems using artifical neural networks,” *ICT Express,* vol. 4, no. 2, pp. 95-99, June 2018.

[48] Darktrace, “Cyber AI Platform,” September 2019. [Online]. Available: https://www.darktrace.com/en/technology/#cyber-ai. [Accessed 10th September 2019].

[49] Hogzilla, “Hogzilla IDS ...big data technologies empowering your detection capabilities,” 2019. [Online]. Available: http://ids-hogzilla.org/. [Accessed 20th August 2019].

[50] A. Hariharan, A. Gupta and T. Pal, “CAMLPAD: Cybersecurity Autonomous Machine Learning Platform for Anomaly Detection,” 29th July 2019. [Online]. Available: https://arxiv.org/abs/1907.10442. [Accessed 10th September 2019].

[51] J. Maddison, “Using Advanced AI to Stay Ahead of Cybercriminals,” Fortinet, 12th February 2019. [Online]. Available: https://www.fortinet.com/blog/industry-trends/using-advanced-ai-to-stay-ahead-of-cybercriminals.html. [Accessed 2nd September 2019].

[52] Fortinet, “Predictive Intellegence: Protect Against Tomorrow's Threats Today,” 2019. [Online]. Available: https://www.fortinet.com/solutions/enterprise-midsize-business/machine-learning.html. [Accessed 2nd September 2019].

[53] eSentire, “Managed Detection and Response,” 2019. [Online]. Available: https://www.esentire.com/mdr-better-together/. [Accessed 30th August 2019].

[54] eSentire, “A Closer Look: The sSentire Difference,” 2019. [Online]. Available: https://www.esentire.com/assets/resources/f766bc8af8/eSentire-Difference-A-Closer-Look.pdf. [Accessed 30th August 2019].

[55] Darktrace, “The Enterrprise Immune System,” September 2019. [Online]. Available: https://www.darktrace.com/en/products/enterprise/. [Accessed 10th September 2019].

[56] Darktrace, “Darktrace Antigena,” September 2019. [Online]. Available: https://www.darktrace.com/en/products/antigena/ . [Accessed 9th September 2019].

[57] R. Calderon, “The Benefits of Artificial Intelligence in Cybersecurity - Masters of Science Thesis,” 15th January 2019. [Online]. Available: https://digitalcommons.lasalle.edu/ecf_capstones/36/.

[58] M. Brundage, S. Avin, J. Clark, H. Toner, P. Eckersley, B. Garfinkle, A. Dafoe, P. Scharre, T. Zeitzoff, B. Filar, H. Anderson, H. Roff, G. C. Allen, J. Steinhardt and C. Flynn, "The Malicious Use of Artifical Intelligence: Forecasting, Prevention, and Migration," February 2018. [Online]. Available: The malicious use of artificial intelligence: Forecasting, prevention, and mitigation. [Accessed 30th August 2019].

[59] W. Xu, Y. Qi and D. Evans, "Automatically Evading Classifiers: A Case Study on PDF Malware Classifiers," in *Network and Distributed Systems Symposium*, San Diego, California, 2016.

[60] I. Goodfellow, N. Papernot, S. Huang, R. Duan, P. Abbeel and J. Clark, "Attacking Machine Learning with Adversarial Examples," 24th February 2017. [Online]. Available: https://openai.com/blog/adversarial-example-research/. [Accessed 10th September 2019].

[61] iapp, "2018 Privacy Tech Vendor Report," 2018. [Online]. Available: https://iapp.org/media/pdf/resource_center/2018-Privacy-Tech-Vendor-Report.pdf. [Accessed 10th September 2019].

[62] S. Gürses and J. M. del Alamo, "Privacy Engineering: Shaping an Emerging Field of Research and Practice," *IEEE Security & Privacy,* pp. 40-46, 6th April 2016.

[63] IEEE, "In Support of Privacy Engineering," 18th November 2018. [Online]. Available: http://globalpolicy.ieee.org/wp-content/uploads/2018/11/IEEE18021.pdf. [Accessed 10th September 2019].

[64] N. Notario, A. Crespo, A. Kung, I. Kroener, D. Le Métayer, C. Troncosco and J. M. del Álamo, "PRIPARE: A New Vision on Engineering Privacy and Security by Design," in *Cyber Security and Privacy. CSP 2014. Communications in Computer and Information Science*, vol. 470, F. Cleary and M. Felici, Eds., Springer, 2014, pp. 65-76.

[65] S. S. Shapiro, "Privacy Risk Analysis Based on System Control Structures: Adapting System-Theoretic Process Analysis for Privacy Engineering," in *IEEE Symposium on Security and Privacy Workshops (SPW)*, San Jose, CA, 2016.

[66] C. Kalloniatis, E. Kavakli and S. Gritzalis, "Methods for Designing Privacy Aware Information Systems: A Review," in *13th Panhellenic Conference on Informatics*, Corfu, Greece, 2009.

[67] S. S. Gill, I. Chana, M. Singh and R. Buyya, "CHOPPER: an intelligent QoS-aware autonomic resource management approach for cloud computing," *Cluster Computing,* vol. 21, no. 2, p. 1203–1241, June 2018.

[68] V. Cardellini, T. G. Grbac, M. Nardelli, N. Tanković and H.-L. Truong, QoS-Based Elasticity for Service Chains in Distributed Edge Cloud Environments, book chapter, Autonomous Control for a Reliable Internet of Services, vol. 2018, Springer Open- LNCS Sublibrary: SL5 – Computer Communication Networks and Telecommunications.

[69] R. Mahmud, K. Ramamohanarao and R. Buya, "Latency-Aware Application Module Management for Fog Computing Environments," *ACM Transactions on Internet Technology (TOIT) - Regular Papers, Special Issue on Service Management for IOT and Special Issue on Knowledge-Driven BPM TOIT,* vol. 19, no. 1, pp. 9:1-9:21, November 2018.

[70] L. Peng, A. R. Dhaunu and P.-H. Ho, "Toward integrated Cloud–Fog networks for efficient IoT provisioning: Key challenges and solutions," *Future Generation Computer Systems,* vol. 88, pp. 606-613, 2018.

[71] S. B. Akintoye and A. Bagula, "Improving Quality-of-Service in Cloud/Fog Computing through Efficient Resource Allocation," *Sensors,* vol. 19, no. 6, p. 1267, 2019.

[72] S. S. Gill, P. Garraghan and R. Buyya, "ROUTER: Fog enabled cloud based intelligent resource management approach for smart home IoT devices," *The Journal of Systems and Software,* vol. 154, pp. 125-138, 2019.

[73] J. Wang, L. Zhao, J. Liu and N. Kato, "Smart Resource Allocation for Mobile Edge Computing: A Deep Reinforcement Learning Approach," *IEEE Transaction of Emerging Topics in Computing,* pp. 1-1, 2019.

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

[75] M. Chen and V. C. Leung, "From cloud-based communications to cognition-based communications: A computing perspective," *Computer Communications,* vol. 128, pp. 74-79, September 2018 2018.

[76] L. Carnevale, A. Celesti, A. Galletta, S. Dustdar and M. Villari, "From the Cloud to Edge and IoT: a Smart Orchestration Architecture for Enabling Osmotic Computing," in *32nd International Conference on Advanced Information Networking and Applications Workshops*, Krakow, Poland, 2018.

[77] Q. D. La, M. V. Ngo, T. Q. Dinh, T. Q. Quek and H. Shin, "Enabling intelligence in fog computing to achieve energy and latency reduction," *Digital Communications and Networks,* vol. 5, no. 1, pp. 3-9, February 2019 2019.

[78] R. Patel, "New Survey Yields Kubernetes as Mainstream but Requires Significant Expertise to Adopt Across the Enterprise," 2019. [Online]. Available: https://www.nirmata.com/2019/01/24/new-survey-yields-kubernetes-as-mainstream/.

[79] Portworx and Aqua, "2019 Container Adoption Survey," 2019. [Online]. Available: https://portworx.com/wp-content/uploads/2019/05/2019-container-adoption-survey.pdf.

[80] L. Hecht, "Add It Up: Enterprise Adoption of Kubernetes Is Growing," December 2018. [Online]. Available: https://thenewstack.io/add-it-up-enterprise-adoption-of-kubernetes-is-growing/.

[81] MarketsandMarkets, "Application Container Market by Service (Container Monitoring, Security, Data Management, Networking, Orchestration), Platform (Docker, Kubernetes), Application Area, Deployment Mode, Organization Size, Vertical, and Region - Global Forecast to 2023," May 2018. [Online]. Available: https://www.marketsandmarkets.com/Market-Reports/application-container-market-182079587.html. [Accessed 28th August 2019].

[82] A. Sill, "Standards at the Edge of the Cloud," *IEEE Cloud Computing,* vol. 4, no. 2, pp. 63-67, 2017.

[83] OpenFog Consortium Architecture Working Group , "OpenFog Reference Architecture for Fog Computing," OpenFog Consortium, 2017.

[84] IEEE Standards Association, "IEEE 1934-2018-IEEE Standard for Adoption of OpenFog Reference Architecture for Fog Computing," IEEE, 2018.

[85] C. Puliafito, E. Mingozzi, F. Longo, A. Puliafito and O. Rana, "Fog Computing for the Internet of Things: A Survery," *ACM Transactions Internet Technologies,* vol. 19, no. 2, pp. 18:1-18:41, 2019.

[86] ICC, "The Industrial Internet Consortium and OpenFog Consortium Unite," 30th January 2019. [Online]. Available: https://www.liconsortium.org/press-room/01-31.19.htm. [Accessed 17th July 2019].

[87] NIST, "Fog Computing Conceptual Model, Recommendation of the National Institute of Standards and Technology," National Institute of Standards and Technology, 2018.

[88] MarketsandMarkets, "Cloud Computing Market by Service Model (Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS)), Deployment Model (Public, Private, and Hybrid), Organization Size, Workload, Vertical, and Region - GF to 2023," [Online]. Available: https://www.marketsandmarkets.com/pdfdownloadNew.asp?id=234.html.

[89] Adroit Market Research, "Global Cloud Computing Market Size 2018 by Deployment, By Application, By Product, By Region and Forecast 2018 to 2025," 30th May 2019. [Online]. Available: https://www.adroitmarketresearch.com/industry-reports/cloud-computing-market. [Accessed 28th August 2019].

[90] Research and Markets, "Global Cloud Computing Market Analysis & Trends - Industry Forecast to 2025," January 2017. [Online]. Available: https://www.researchandmarkets.com/research/54tvtd/global_cloud. [Accessed 28th August 2019].

[91] Wise Guy Reports, "Cloud Services Global Market Opportunities and Strategies To 2022," 18th March 2019. [Online]. Available: https://www.wiseguyreports.com/sample-request/3928608-cloud-services-global-market-opportunities-and-strategies-to. [Accessed 28th August 2019].

[92] Grand View Research, "Fog Computing Market Analysis By Solution, By Hardware, By Application, By Region, & Segment Forecasts, 2018-2025," March 2017. [Online]. Available: https://www.grandviewresearch.com/industry-analysis/fog-computing-market. [Accessed 27th August 2019].

[93] Reports Intellect , "Global Fog Computing Market Size and Forecast To 2025," January 2018. [Online]. Available: https://www.reportsintellect.com/sample-request/256872?ata. [Accessed 27th August 2019].

[94] Verified Market Research, "Global Fog Computing Market Size and Forecast To 2025," November 2018. [Online]. Available: https://www.verifiedmarketresearch.com/download-sample/?rid=5425. [Accessed 27th August 2019].

[95] Statista, "Global fog computing market size forecast from 2018 to 2022, by vertical (in million U.S. dollars)," 6th June 2018. [Online]. Available: https://www.statista.com/statistics/830451/world-fog-computing-revenue-by-vertical/. [Accessed 26th August 2019].

[96] Allied Market Research, "Edge Computing Market By Component, Application, Organization Size, and Industry Vertical: Global Opportunity Analysis And Industry Forecast, 2018-2025,"

May 2019. [Online]. Available: https://www.alliedmarketresearch.com/edge-computing-market.

[97] Research and Markets, “Global Edge Computing Market (2018-2023),” September 2018. [Online]. Available: https://www.researchandmarkets.com/reports/4667633/global-edge-computing-market-2018-2023#rela0-4667632. [Accessed 28th August 2019].

[98] Grand View Research, “Edge Computing Market Size, Share & Trends Analysis Report By Component, By End Use, And Segment Forecasts, 2019-2025,” June 2019. [Online]. Available: https://www.grandviewresearch.com/industry-analysis/edge-computing-market. [Accessed 29th August 2019].

[99] Gartner, “Hype Cycle for Cloud Computing, 2018,” 31st July 2018. [Online]. Available: https://www.gartner.com/en/documents/3884671. [Accessed 28th August 2019].

[100] I. Vilajosana - CEO Worldsensing, “The Definitive Buyer's Guide to IOT Solutions - and why you may not need an IoT Platform,” 2019. [Online]. Available: https://blog.worldsensing.com/ebook/buyers-guide-iot-solutions-not-iot-platforms. [Accessed August 2019].

[101] P. Fredric, “Gartner's Top 10 IoT trends for 2019 and beyond,” NetworkWorld, 26th November 2019. [Online]. Available: https://www.networkworld.com/article/3322517/a-critical-look-at-gartners-top-10-iot-trends.html. [Accessed 27th August 2019].

[102] M. Torchia and M. Shirer, “IDC Forecasts Worldwide Spending on the Internet of Things to Reach $745 Billion in 2019, Led by the Manufacturing, Consumer, Transportation, and Utilities Sectors,” IDC, 3rd January 2019. [Online]. Available: https://www.idc.com/getdoc.jsp?containerId=prUS44596319. [Accessed 27th August 2019].

[103] Help Net Security, “Global IoT market projected to reach $1111.3 billion by 2026,” 12th July 2019. [Online]. Available: https://www.helpnetsecurity.com/2019/07/12/global-iot-market/. [Accessed 20th August 2019].

[104] Mordor Intelligence, “Internet of Things (IoT) Market - Growth, Trends, and Forecase (2019 - 2024),” Mordor Intelligence, Hyderabad, India, 2016.

[105] Microsoft, “IoT Signals,” 2019. [Online]. Available: https://azure.microsoft.com/mediahandler/files/resourcefiles/iot-signals/IoT-Signals-Microsoft-072019.pdf. [Accessed 26th August 2019].

[106] i-Scoop, “IoT Barometer 2017/2018: number of large scale IoT projects doubled in one year,” 2019. [Online]. Available: https://www.i-scoop.eu/internet-of-things-guide/large-iot-projects-2017-2018-analysis/.

[107] F. Dahlqvust, M. Patel, A. Rajiko and J. Shulman, “Growing opportunities in the Internet of Things,” McKinsey, July 2019. [Online]. Available: https://www.mckinsey.com/industries/private-equity-and-principal-investors/our-insights/growing-opportunities-in-the-internet-of-things?cid=other-eml-alt-mip-mck&hlkid=9efc13dd456841029221011fffac9c79&hctky=10435690&hdpid=ff8df05e-dec4-40e5-a297-a0eb8988f. [Accessed 26th August 2019].

[108] Relecura, “IoT - Internet of Things,” May 2018. [Online]. Available: https://relecura.com/wp-content/uploads/2018/05/IoT_Patent_Landscape_May2018.pdf. [Accessed 26th August 2019].

[109] D. Kalez, “Industrial Tech M&A Overview,” 2018. [Online]. Available: https://www.mcrockcapital.com/uploads/1/0/9/6/10961847/2._dave_kalez_keybanc.pdf. [Accessed 26th August 2019].

[110] Ernst & Young, “Future of IoT,” 2019. [Online]. Available: https://www.ey.com/Publication/vwLUAssets/EY_-_Future_of_IoT/$FILE/EY-future-of-Iot.pdf. [Accessed 26th August 2019].

[111] AGCOM - Dept of Economics and Statistics, “ Big data - Interim report in the context of the joint inquiry on “Big data” launched by the AGCOM deliberation No. 217/17 / CONS,” June 2018. [Online]. Available: https://www.agcom.it/documents/10179/10875949/Allegato+4-9-2018/f9befcb1-4706-4daa-ad38-c0d767add5fd?version=1.0. [Accessed 26th August 2019].

[112] Whizlabs - Amit Verma, “The Relationship between IoT, Big Data, and Cloud Computing,” 28th November 2018. [Online]. Available: https://www.whizlabs.com/blog/relationship-between-iot-big-data-cloud-computing/. [Accessed 26th August 2019].

[113] S. Montegiove, “When the algorithm becomes selfish: an interview with Massimo Chiriatti,” Ingenium Magazine, 19th July 2019. [Online]. Available: http://www.ingenium-magazine.it/en/humanless-lalgoritmo-egoista-intervista-a-massimo-chiriatti/. [Accessed 26th August 2019].

[114] McKinsey & Company, “The Internet of Things, How to capture the value of IoT,” May 2018. [Online]. Available: https://www.mckinsey.com/~/media/McKinsey/Business%20Functions/McKinsey%20Digital/Our%20Insights/The%20Internet%20of%20Things%20How%20to%20capture%20the%20value%20of%20IoT/How-to-capture-the-value-of-IoT.ashx. [Accessed 26th August 2019].

[115] B. Liang, H. Li, M. Su, P. Bian, X. Li and W. Shi, “Deep Text Classification Can be Fool,” 7th January 2019. [Online]. Available: https://arxiv.org/abs/1704.08006. [Accessed 30th August 2019].