Making a Business Out of (Predictive Application Management in) the Fog*

Giuseppe Astuti, Antonio Brogi, Stefano Forti
Department of Computer Science
University of Pisa, Pisa, Italy

***Abstract*—Managing large, highly distributed IoT applications over heterogeneous Fog infrastructures so to meet all their stringent QoS (as well as hardware and software) requirements is intrinsically difficult. Different simulation and predictive methodologies have been proposed to estimate key performance indicators of eligible application deployments and managements so to identify the best candidates. In this paper, we describe the current business model environment and discuss two possible business models for creating value from a company provisioning predictive Fog application management services.**

***Index Terms*—predictive application management, business model environment, business models, Fog computing, IoT**

I. Introduction

Modern applications are usually organised as large, highly distributed software systems [1]. With the continuous growth of the Internet of Things (IoT), such applications are made of interacting services (or components) that should be deployed to pervasively distributed infrastructures spanning the continuum from the IoT to the Cloud (e.g., Fogs, Cloudlets, Cloud-Edges [2]). Also, many of these applications being life-, business- or mission-critical, they need the deployment infrastructures to guarantee stringent Quality of Service (QoS) requirements, as well as all software and hardware requirements.

Whilst reasoning on software systems is currently exploited in industrial settings to optimise the results achievable by Continuous Integration (CI) tools [3], the Fog creates room for applying akin strategies to the subsequent step, performed by Continuous Deployment (CD) tools. Indeed, the problem of deciding how to deploy and manage complex distributed software architectures on top of equally complex distributed infrastructures in a QoS-aware manner is provably difficult [4]. Much literature has tackled such problem in the last years by proposing possible approaches to improve application placement and management via simulation and prediction of Key Performance Indicators (KPIs) [5]. In our previous work, we also proposed and prototyped methodologies to determine QoS- [6], cost- [7], and security-aware ([8], [9]) deployments of Fog applications (as in FogTorchΠ [10], [11]) as well as to assess correctness and effectiveness of their subsequent management operations (as in FogDirMime [12]).

In this paper, we focus on describing the market opportunities represented by offering predictive application management services in Fog computing scenarios. As new Fog services are expected to undergo a compound annual growth rate of around 119% in the years 2018–2022 (with a consistent market share increase from 15.7% to 20.4%)[13], this seems an interesting line to investigate. New business models will be key to pursue this expected growth. The original contribution of this paper is two-fold. Firstly, it details the current business model environment in which a predictive Fog application management services company can try to go to market. Secondly, it discusses two prototype business models for such company.

To perform our analysis, we follow the Business Model Canvas (BMC) methodology by Osterwalder and Pigneur [14], which focuses on an integrated suite of tools that can be used to design new business models as well as to analyse and improve existing ones. The BMC methodology consists of (iterating) two main steps:

1) defining a *Business Model Environment* that can be used to understand the context (i.e., the *design space*) in which a business is created,
2) defining one or more *Business Model Canvases* on how such business can create, deliver and capture value within the described environment.

Throughout the paper, we will focus on a potential new business that provides an ecosystem of services to perform predictive application deployment and management and offers a tool capable of monitoring both current and historical infrastructure conditions (i.e., end-to-end links QoS, available nodes and hardware resources, available IoT devices) to collect input data for its predictive engines. We assume that such predictive engines can then recommend the best deployment(s) and management(s) for input Fog applications, by relying on analyses and simulations.

The rest of this paper is organised as follows. Sect. II describes the Business Model Environment in which predictive Fog application management companies could currently design their structure. Sect. III discusses two potential prototype Business Model Canvases for such company. Finally, Sect. IV draws some concluding remarks.

II. Business Model Environment

The description of the *Business Model Environment* (BME) permits to set the *design space* that one can consider when defining possible business models for a new or an existing business [14]. The BME describes the environmental elements

* Work partly supported by project "*DECLWARE: Declarative methodologies of application design and deployment*" (PRA_2018_66) funded by University of Pisa.

Corresponding author: Stefano Forti (stefano.forti@di.unipi.it).

upon which the business has no control over, such as competitors, changing customers needs and desires, technological and legal trends, and the overall state of the global and local economy. Therefore, it permits to identify threats, opportunities and constraints to design more robust business models.

In this section, after identifying three different *actors* involved in the process of defining business models for predictive Fog application management (Sect. II-A), we detail

- the *Market Forces* (Sect. II-B), which comprise market issues, market segments, needs and demands, switching costs and revenue attractiveness, that are currently shaping the market for predictive application management in Fog computing,
- the *Industry Forces* (Sect. II-C), which cover stakeholders, competitors, new entrants, suppliers and other value chain actors, as well as substitute products and services that can affect the business strategies for the considered product,
- the *Key Trends* (Sect. II-D), which include regulatory, technology, socio-economic, societal and cultural trends that currently characterise the business environment, and
- the *Macro-Economic Forces* (Sect. II-E), which comprehend global market conditions, capital markets, economic infrastructure and commodities, that can influence the adoption of predictive management of Fog applications.

A. Actors

In what follows, we will always refer to three main actors which characterise the BME for predictive Fog application management. Namely:

- *Asset Manufacturers* (AMs) which are producing hardware components and systems (e.g., access points, routers, switches) to build Fog infrastructures (e.g., CISCO [15] and Dell [16]), and often take care of deploying and managing their equipment for other service providers (e.g., Telco operators),
- *Infrastructure Providers* (IPs) that *(i)* already own and manage part of a potential Fog infrastructure (e.g., Telco or Cloud providers), *(ii)* are planning to invest in building a new one (e.g., new Fog providers [17]), or *(iii)* are willing to voluntarily share their computing capabilities in opportunistic or crowd-based Fog networks (e.g., private users/businesses [18]),
- *Application Operators* (AOs) that deploy their application services to Cloud and Fog nodes, e.g., private users, business employees or whoever manages the life cycle of one or more applications, being responsible for guaranteeing the QoS effectiveness of the application management.

B. Market Forces

Market Issues

CISCO envisions 500 billion devices to be connected to the Internet by 2030 [19]. In the context of continuous growth of the IoT and of the data it produces, there is a clear need for processing and filtering information before it reaches the Cloud so to enforce prompter decision making and reactions [10], [20]. In addition to Fog computing, various other paradigms have been proposed (e.g., Mist computing [21], micro-clouds [22], Edge Computing [23], Osmotic Computing [24]) that will give rise to large and complex infrastructures deployed over the continuum from the IoT to the Cloud and characterised by high geo-distribution, node heterogeneity, seamless connectivity and mobility of nodes and devices [25]. Such new infrastructures will extend *utility computing* towards the Edge of the network (and therefore towards less capable, possibly user-owned devices), following the trend inaugurated by the introduction of the Cloud in the last decade.

On the other hand, the architectural evolution of software systems (from multi-component [26] to Software Oriented Architectures [27] to micro-services [28]) has made more and more complex the structure of the applications that will have to be deployed and managed on top of the pervasively distributed infrastructures mentioned before. Furthermore, most of the future IoT applications, relating to different verticals such as smart buildings [8], Industry 4.0 [29], e-health [30] and self-driving cars [31], are considered business-, mission- or even life-critical. This translates into stringent QoS requirements (e.g., latency, bandwidth, uptime, security) that they have to meet to run properly or to accommodate the deployer's needs (e.g., operational costs, energy consumption), in addition to classical software and hardware constraints.

Overall, simultaneously and optimally managing a complex system (the application) on top of another (the infrastructure) is intrinsically challenging and requires suitable tooling [4]. Furthermore, bad or poor management choices can lead to unsatisfactory service QoS, waste of electrical power or money, and – worst of all – life losses due to bad service placement, unbalanced loading of Fog devices and service downtime, respectively. Considering as potential customers for predictive application management the actors we introduced at the beginning of this section, the key issues they are facing relate to taming the scale and complexity of Fog architectures and of modern distributed software systems, whilst guaranteeing all of their QoS requirements.

Market Segments

As aforementioned, throughout this analysis of the BME, we consider *Asset Manufacturers*, *Infrastructure Providers* and *Application Operators* as the main Customer Segments which can be interested in purchasing predictive management services. In this regard, it is interesting to try to spot new, emerging segments within them.

First, besides established *Asset Manufacturers* (e.g., CISCO [15], Dell [16]), smaller start-up companies (e.g., Nebbiolo Technologies [32]) are pioneering the manufacturing market to produce their Fog-tailored hardware and platforms, and to provide QoS-aware resource management in support of real-time IoT applications. They represent a valuable niche because they aim at quickly differentiating their offer from those of established companies by relying on more agile business structures and by representing a (first) better answer to the market issues described before.

Second, traditional Cloud *Infrastructure Providers* (e.g., Amazon, Google) will be joined by new Fog computing entrants as well as by telecommunication providers (e.g., Vodafone) that already own and manage pervasive Edge infrastructures as well as local datacentres. It is worth noting that telecommunication providers will likely host both added-value services for their clientele, and their customers' application deployments. They might also plan to federate with others so to further improve the QoS they can support, and fight traditional *Infrastructure Providers* on market.

Last, *Application Operators* foresee the possibility to improve their revenues by improving the QoS or the services they offer to their customers. Among them, those which are also producing customised hardware to run part of their application (e.g., Philips Hue, Ikea Tradfri), represent a particularly interesting segment. They might see in the new paradigms an opportunity to crunch Big Data at a finer grain by relying on their customers' power consumption and owned assets. Eventually, this leads to increase their revenues by selling the derived insights, whilst reducing the costs related to elaborating all collected data.

Needs and Demands

Asset Manufacturers face the issue of producing a competitive (hardware and software) platform to permit applications and devices management within an infrastructure composed of their hardware and, possibly, third-party assets. CISCO, among others, is already providing some tooling to manage applications over their infrastructures. For instance, CISCO Fog Director [12] enables the programmatic management of large-scale production deployments of Fog applications, by providing centralised management services that span the entire application lifecycle, and that can be accessed via RESTful APIs. Naturally, to keep pace with the described market shift and to maintain their shares, *Asset Manufacturers* will consider enriching their platforms by embedding predictive and QoS-aware features in their products so to allow their customers to evaluate *a priori* their deployment and management choices.

Similarly, *Infrastructure Providers* have to face (optimal) resource management and allocation [33] and (optimal) workload allocation [34], to properly operate their (heterogeneous and pervasive) Fog systems and evaluate the service level guarantees they can offer to their customers. Not only the adoption of predictive models that enable assessing their internal application management would lead to improved QoS (e.g., uptime, availability) but also to savings in terms of operational costs (e.g., power consumption). Last, but not least, the adoption of such techniques would lead to the definition of new and suitable Service Level Objectives (SLOs) for their Service Level Agreements (SLAs) in Fog landscapes.

Finally, *Application Operators* are naturally interested in the QoS aspects of the services they offer to their customers. Particularly, in addition to the classical uptime and availability, more and more they consider network QoS and security as leading parameters for their application management choices. Evidently, another main concern of *Application Operators* relates to the possibility of assessing beforehand their management choices and improving them by *a priori* determining trade-offs among inversely correlated metrics (e.g., cost vs. QoS-assurance, security vs. response time), which lead to considerable savings whenever they have to deploy to a mixture of (multiple third-party) infrastructures.

Switching Costs

Despite much research is investigating the field of application management methodologies in Fog scenarios [5], Fog computing overall is still an emerging area and predictive application management seems to have no competitors on the market yet. It is therefore important for a company which aims at selling predictive Fog application management to start a consistent branding process as soon as possible, whilst providing top of the range quality for their services. This, along with suitable lock-in strategies, would create costs for customers evaluating a switch to a competitor.

Usually, being among the firsts on the market enables a company to establish strong brand recognition and customer loyalty before other competitors enter the arena [35]. It also permits to improve the service and to set the market price for the new offer before others do. However, being the current Cloud scenarios dominated by large *Infrastructure Providers*, new entrants in the market should identify possible solutions to avoid being copied and overwhelmed by bigger players. We currently envision two main possibilities: (*i*) forming partnerships with existing *Infrastructure Providers* and taking care of tailoring the predictive methodologies to their needs, whilst selling them either independently or through the *Infrastructure Provider* platforms, or (*ii*) selling the predictive solutions to the highest bidder, after completing a strong branding process and product tuning.

Revenue Attractiveness

Asset Manufacturers could include predictive application management within their management platforms relying, for instance, on *licence-based* formulae. Whilst there can be some profit margins in this approach, *Asset Manufacturers* can consider it as a short-term trial to evaluate the decision of including predictions in their systems. This can lead to the commitment towards (*i*) extending their own management system with predictions, (*ii*) acquiring the company providing the new features, or, more unlikely, (*iii*) abort the predictive management line. Overall, (*ii*) seems to represent the most lucrative possibility for them.

Analogously, *Infrastructure Providers* will be willing to pay for a predictive application management service depending on the increased profits and visibility this can give to them. Differently from *Application Operators*, they can pay for the service they use also in terms of data collected during actual application management in order for the company selling predictive solutions to increase the quality of their results. Similarly to *Asset Manufacturers*, they might consider adopting these solutions as an exploratory trial. However, they could also see outsourcing the predictive services as a source of time and money savings, actually easing federation with other providers that rely on the same external solution.

Finally, *Application Operators* will be attracted by the possibility of improving deployment QoS at reduced costs. They will be likely to pay for achieving automated management of their applications and a larger amount of predictable performance indicators.

C. Industry Forces

Competitors

Whilst (apart from researchers in the field) we were not able to identify direct business competitors in the predictive application management field, many dominant players exist among major established *Asset Manufacturers* and *Infrastructure Providers*. They could easily exploit the trust of and the commercial relationship with their customers to start integrating predictive application management in their tooling. Also, they can rely on the fact that their assets (e.g., access points, smart-speakers) are already installed at the customers' premises and can act as Fog nodes to deploy services. However, they could mainly apply such advantage within the borders of their product lines or infrastructures, respectively. On the contrary, a third-party provider for predictive application management can base its analyses on data collected from different hardware or infrastructures and produce more accurate (and satisfying) results. Large established companies might trust more such new provider as an intermediary, with respect to disclosing data to their direct competitors.

New Entrants

As the threat of new entrants is one of the forces that shape the competitive structure of an industry [36], we have to consider them in our analysis. In addition to established *Asset Manufacturers* and *Infrastructure Providers* interested in Fog computing, there can be new smaller ones (e.g., the niche market segments we identified before) that decide to invest part of their effort in designing and implementing their own predictive systems. Overall, we envision they could fear the costs involved in following such line and suffer from the same issues of the competitors described in the previous paragraph.

Substitute Services

The current market does not exhibit substitute products for a predictive Fog application management service. *Application Operators* mainly proceed by manually tuning their deployments and management policies, whilst coping with many uncertainties (e.g., infrastructure variations, workload peeks) that lead to a time-consuming and costly *trial & error* procedure.

Suppliers and Other Chain Value Actors

All three actors we identified can represent fruitful value chain actors. *Asset Manufacturers* are very likely to provide support for integrating predictive methodologies in their ecosystems, and also attempt to acquire and incorporate the predictive management business. *Infrastructure Providers* will be strongly involved as suppliers in the process of producing high-quality predictions by collecting the data needed to execute predictive experiments. As it is crucial to obtain such data, we foresee the possibility of inaugurating partnerships or, as for *Asset Manufacturers*, of considering to acquire the *know-how* of the provider of the predictive application management service.

Stakeholders

Stakeholders depend on the structure of the business, that is, different organisations of the company lead to different entities involved [37]. We will, therefore, focus on general candidates to be stakeholders, and classify them between internal and external to the business. Internal stakeholders are represented by employees, managers and owners. Since the predictive methodologies are still in their infancy, we foresee the crucial need for the business to rely on strong R&D human resources. Among the external stakeholders, besides *Asset Manufacturers* and *Infrastructure Providers*, we identified:

- *Technology Influencers*, such as the OpenFog Consortium [20], that might accelerate the adoption of Fog computing,
- *National or local governments* that could be interested in regulating Fog computing landscapes as well as in creating new infrastructures and incentives for citizens and businesses to use or develop new IoT services (e.g., smart-cities, Industry 4.0),
- *Academic and industrial researchers* that contribute to the growth of IoT, Fog computing, and naturally, predictive application management methodologies and prototypes.

D. Key Trends

Technology Trends

We focus on technology trends that can affect Fog landscapes, by speeding up the adoption of new technologies. On one hand, the diffusion of Internet access is improving. Indeed, as of June 2017, 51% of the world's population has Internet access, and in 2015, the International Telecommunication Union estimated about 3.2 billion people can get connected [38]. For instance, the Internet connection speed in the United States, on average, raised from 3.65 to 18.75 Mbps in the years 2007–2017 [39]. On the other hand, computational devices are becoming much cheaper against an increase in their performance and computing power. For instance, Docker [40] has called a challenge to maximise the number of containers that can be deployed to a Raspberry Pi 2 [41] that currently costs about 40$. Currently, the record corresponds to running 2334 containerised web-servers on the same piece of hardware [42]. In addition, IoT applications are shaped in highly distributed architectures and often subject to QoS and real-time constraints, as mentioned in the previous sections. Also, industries are investing money and effort in improving and automating the management of distributed applications. For instance, TOSCA is a standard that facilitates the platform-independent description of Cloud applications [43] and its runtime is expected to enable fully automated plan-based deployment and management [44]. Last but not least, in a world where the Everything-as-a-Service (XaaS) paradigm [45] is established, we foresee the possibility for new, highly specialised companies to provide Management-as-a-Service (MaaS) to *Asset Manufacturers*, *Infrastructure Providers* and

Application Operators in Fog scenarios. All these possible trends clearly constitute enabling factors for the widespread adoption of Fog computing, which is expected to potentially reach a market of more than $18 billion dollars worldwide by 2022 [13].

In this context, threats to predictive management can be mainly identified in the presence of devices produced or managed by *Asset Manufacturers* and *Infrastructure Providers* in our homes (e.g., Google Home, Amazon Dot) and cars (e.g., Android Car). They can actually translate into a consistent competitive advantage for big players that, however, are expected to be able to perform application management only on their own infrastructure (i.e., without considering third-party assets to better guarantee QoS requirements).

Regulatory Trends
The recent rise of the Fog has not yet been regulated by *ad hoc* laws and regulations. Of some interest for the European market, can be the recent EU-GDPR, designed to harmonise data privacy laws across Europe and reshape the way organisations approach data privacy [46]. On the other hand, despite IoT devices can be easily hacked due to their ease of access ([47], [48]), no organic, well-structured regulations are available yet. Naturally, we expect the Business Environment to evolve as legislators produce new matter to regulate these subjects.

Societal and Cultural Trends
In the last decades, we are moving from the mass production era to the *mass customisation* era [49]. New IoT applications constitute a clear example of this phenomenon, being capable to adapt offered services to actual customers needs, by monitoring and learning from their habits or interests. New societal and cultural trends are characterised by a shift in which experiences are more important than products [50]. Presumably, such a trend towards customisation will impact new technologies, such as Fog computing. Indeed, the new family of applications that will be deployed to the Fog will require QoS specifications to be fulfilled. Whilst some of these requirements will be application-specific (e.g., latency, bandwidth, hardware) or driven by the *Application Operators*' desiderata (e.g., operational costs, power consumption), other (e.g., data privacy) might involve end-users preferences.

Finally, the decreasing price of hardware is accompanied by a decentralising shift from the Cloud to the Edge of the network and by a greater users awareness on trust, privacy and security issues. This will probably require to realise new human-centred designs for computer systems so to put users in the control loop [51]. Also, it could represent an occasion to include end-users in the application management process both as *Infrastructure Providers* (with their personal devices) and as *Application Operators* (offering their services to others).

Socioeconomic Trends
The World Economic Forum [52], a non-profit Swiss organisation, reports that by 2025 more than half of all current job tasks will be performed by machines as opposed to 29 per cent today [53]. They found out that the rapid change of machines and algorithms, or computer processes that are designed to solve problems, could create 133 million new job opportunities in place of 75 million that will disappear before 2022 [53].

Additionally, large private and public investments in smart energy grids, in transportation and in healthcare, along with increased automation in manufacturing will create new potential market space for Fog computing. Finally, the growth of the middle class – in both developed and developing economies – is a phenomenon that derives from the advent of the Digital Era [54], and it makes more and more people capable of purchasing on-demand, highly customisable services (e.g., Netflix, Uber).

E. Macro-Economic Forces

Global Market Conditions
As highlighted by J.P. Morgan [55], the global market outlook for 2019 is characterised by an overall growth of 2.9%, and "*after a volatile end to 2018, tentative stability has returned to risky markets at the start of the new year, with investors seeing some reversal of the losses experienced in December. Growth momentum has slowed, but the deceleration phase should end before midyear with supportive and flexible policy actions–notably China easing and the Federal Reserve pausing. Recession risks, in the meantime, remain modest for the year ahead*". Current unemployment rates go from the 4% in the United States up to the 6.8% in the European Union. The effect The World Economic Forum in their 2017 survey reports that investments in the tech sectors are always a safe choice, ensuring high Return of Investment (ROI) [56]. They also claim that seven technologies are transforming modern industries. Among those, Artificial Intelligence, Autonomous Vehicles, IoT and connected devices, robots and drones, and custom manufacturing machines stand out in the context of our analysis. Indeed, those relate to Fog computing and to the establishment of new enabling architectures spanning the Cloud-IoT continuum, what will likely require exploiting new predictive deployment and management methodologies.

Capital Markets
In 2018, venture capital got bigger and continued to expand globally [57]. In Europe, many incentives exist for people starting new businesses. For instance, Italian governments are pushing towards founding new businesses with incentives such as national funds for start-ups [58] and Industry 4.0 [59].

Commodities and Other Resources
Since the technology sector is growing, the medium cost of an employee is higher than it used to be in the past, e.g., Forbes reports that in the last ten years there has been a major increase in junior and super-talented developers worldwide [60]. In its survey, Forbes claims that the top students from the top U.S. colleges have seen the highest gains, e.g., entry salaries for the top people go from $80k to $140k in just a few years. Whereas some of the next tier students may have seen their salaries increase from $60-$80k. In Germany and in the United Kingdom, a developer salary amounts to $60k per year while in other countries such as France and Italy is relatively lower ($40k per year).

Economic Infrastructure
As this part is very specific to the region in which a company operates, we leave its analysis for future studies.

III. Business Models for Predictive Fog Application Management

In the following, we describe two prototype business models derived from the Business Model Environment described in the previous section. They both aim at creating value from a SaaS service that provides automated predictive management of Fog applications so to relief *Application Operators* from manually tuning their choices via a *Trial & Error* process.

A. License-based Business Model

The business model canvas of Figure 1 foresees the possibility that *Asset Manufacturers* outsource the management of the infrastructures they sell (and sometimes maintain) for third-party *Infrastructure Providers* to a predictive management service. *Asset Manufacturers* represent the main Customer Segment of this scenario. They are expected to consider extending their management platforms with predictive, QoS-aware management capabilities, and they might consider it beneficial to save the effort and the costs involved in designing and implementing those new features.

Asset Manufacturers should pay yearly license fees (Revenue Streams) to access the predictive application management service and to include it in their ecosystems (Value Proposition). Indeed, this can bring new customers (attracted by the guarantees derived from predictive management) to them and additional profits (if they decide to sell predictive functionalities as an add-on service to their customers).

The predictive management service provider should then contact them through a B2B sales division and through an appealing Web site and targeted marketing campaigns (Channels). Marketing is crucial to establish and maintain the relationship with the *Asset Manufacturers*, along with a usable API to access the service autonomously after paying license fees, and the possibility to get on-demand support in case of problems with the running service (Customer Relationships).

A set of reliable state-of-art predictive methodologies, continuously improved and extended, is key to successfully offer and deliver a competitive service (Key resources). The business we foresee should rely on skilled and expert personnel taking care of research and development of new prototype and production functionalities, marketing and sales, and customer support service (Key activities).

Finally, the license fee should be carefully designed so to cover personnel and Web marketing costs, along with the utility computing resources to be purchased to actually run the service and obtain output predictions (Cost Structure).

B. Freemium Business Model

The business model canvas of Figure 2 instead foresees the possibility of directly targeting *Application Operators* as Customer Segment, aiming at facilitating their QoS-aware management of Fog applications. Whilst the canvas remains unaltered for what concerns Channels, Customer Relationships, Key Activities, Key Resources and Cost Structure, it substantially changes in all other parts.

Indeed, freemium models require the costs of non-paying customers to be covered by another part of the business model or by another Customer Segment. We, therefore, split *Application Operators* into two segments, i.e. free and subscribing customers. Paying customers will have access to the complete version of the considered predictive application management service and to premium customer support. On the contrary, free customers will have limitations to their accounts, e.g. bounded number of monthly predictions, lower quality results, no access to the support service.

Another difference with respect to the model sketched in Figure 1, is in that *Infrastructure Providers* are listed among the Key Partners. Obviously, their partnership represents a fundamental strategic alliance with non-competitors as they have to run a monitoring service on top of their infrastructures to feed data to the prediction engine of the predictive management service. From *Infrastructure Providers* being key partners, it is possible to envision that they will be willing to pay and exploit the predictive management services portal, to also get customised advertisements. This further increases the Revenue Streams, thus making the business model more sustainable in the long run.

IV. Concluding Remarks

In this article we have analysed, exploiting the Business Model Canvas methodology of [14], two possible business models for marketing the provisioning of predictive application management services to different customer segments.

Obviously, the effectiveness of the discussed business models relies on the assumptions we have made on how market forces, industry forces, key trends, and macroeconomic forces will unfold. For instance, the analysis of Sect. II does not consider any competitor for predictive management services as – to the best of our knowledge – none is currently on market. The appearance of such competitors would require revising both the Business Model Environment and the business models.

In perspective, while the freemium business model targeting *Application Operators* (discussed in Sect. III-B) can be run by an independent company, the license-based business model targeting *Asset Manufacturers* (discussed in Sect. III-A) may ultimately lead to considering the option of selling the business to the highest bidder within the Customer Segment.

Needless to say, such future possible scenarios, as well as the very success or failure of the discussed business models will obviously depend on their actual implementation (and refinements) and on the market reactions.

References

[1] J. Soldani, D. A. Tamburri, and W.-J. Van Den Heuvel, "The pains and gains of microservices: A systematic grey literature review," *Journal of Systems and Software*, vol. 146, pp. 215–232, 2018.
[2] S. Yi, C. Li, and Q. Li, "A survey of fog computing: Concepts, applications and issues," in *Proceedings of the 2015 Workshop on Mobile Big Data*, Mobidata '15, pp. 37–42, ACM, 2015.

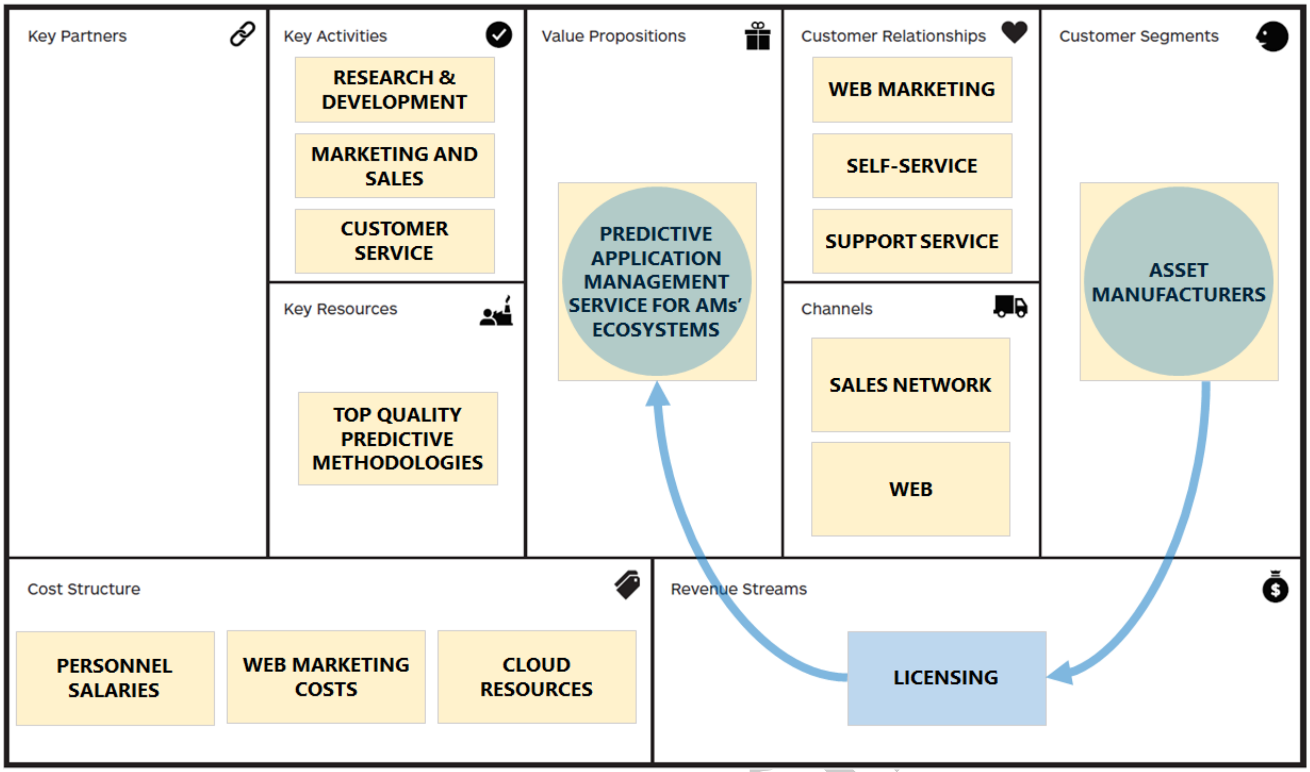

Fig. 1. License-based BMC.

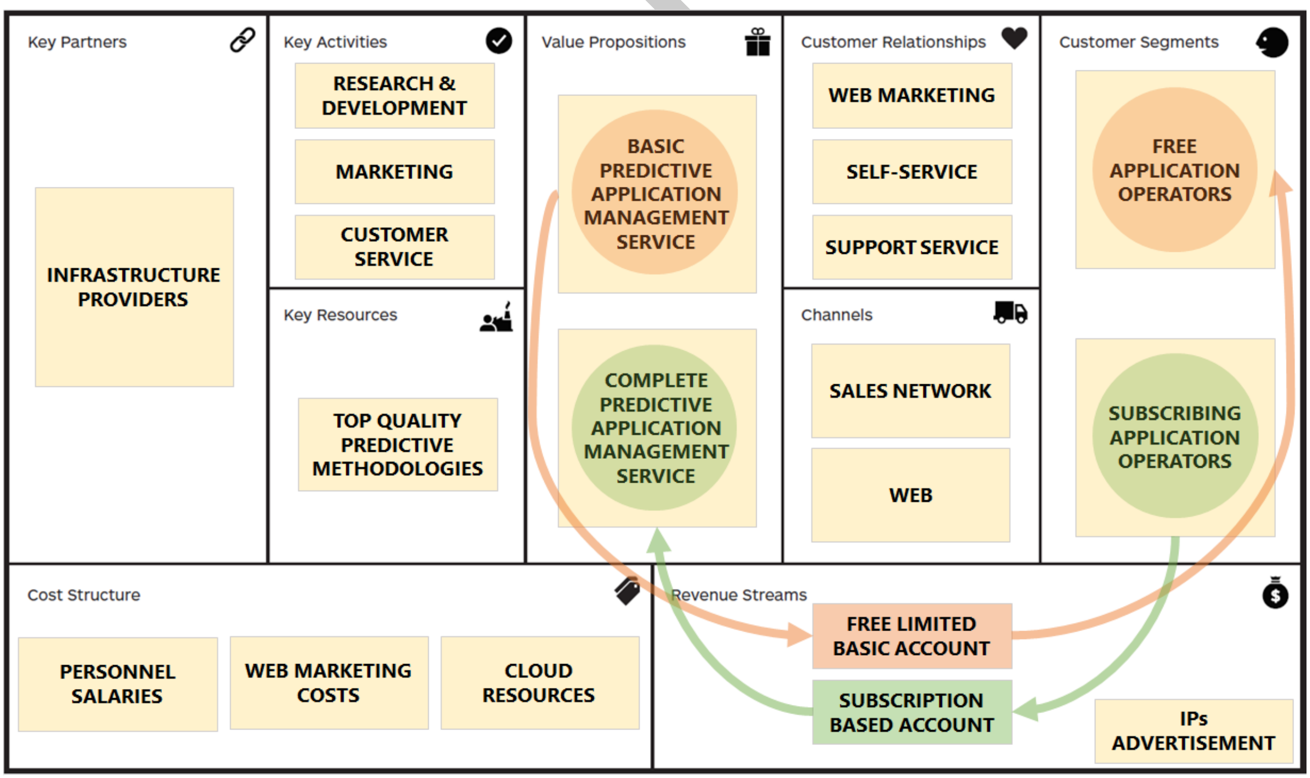

Fig. 2. Freemium BMC.

[3] P. W. O'Hearn, "Continuous reasoning: Scaling the impact of formal methods," in *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science*, pp. 13–25, ACM, 2018.

[4] A. Brogi and S. Forti, "QoS-aware Deployment of IoT Applications through the Fog," *IEEE Internet of Things Journal*, vol. 4, no. 5, pp. 1185–1192, 2017.

[5] A. Brogi, S. Forti, C. Guerrero, and I. Lera, "How to place your apps in the fog-state of the art and open challenges," *arXiv preprint arXiv:1901.05717*, 2019.

[6] A. Brogi, S. Forti, and A. Ibrahim, "How to best deploy your fog applications, probably," in *Fog and Edge Computing (ICFEC), 2017 IEEE 1st International Conference on*, pp. 105–114, IEEE, 2017.

[7] A. Brogi, S. Forti, and A. Ibrahim, "Deploying fog applications: How much does it cost, by the way?," in *Proceedings of the 8th International Conference on Cloud Computing and Services Science*, pp. 68–77, SciTePress, 2018.

[8] A. Brogi, G.-L. Ferrari, and S. Forti, "Secure apps in the fog: Anything to declare?," in *Proc. of the 4th Int. Workshop on Cloud Migration and Architecture (CloudWays 2018)*, 2018.

[9] A. Brogi, G.-L. Ferrari, and S. Forti, "Secure cloud-edge deployments, with trust," *arXiv preprint arXiv:1901.05347*, 2019.

[10] A. Brogi, S. Forti, and A. Ibrahim, "Predictive analysis to support fog application deployment," in *Fog and Edge Computing: Principles and Paradigms, Rajkumar Buyya and Satish N. Srirama (eds.)*, Wiley, 2019.

[11] A. Brogi, S. Forti, and A. Ibrahim, "Optimising QoS-assurance, Resource Usage and Cost of Fog Application Deployments," in *Cloud Computing and Services Science Selected Papers, Communications in Computer and Information Science*, Springer. *In Press*.

[12] S. Forti, A. Ibrahim, and A. Brogi, "Mimicking FogDirector Application Management," *Software-Intensive Cyber-Physical Systems*, 2019.

[13] 451 Research, "Size and impact of fog computing market." https://www.openfogconsortium.org/wp-content/uploads/451-Research-report-on-5-year-Market-Sizing-of-Fog-Oct-2017.pdf.

[14] A. Osterwalder and Y. Pigneur, *Business model generation: a handbook for visionaries, game changers, and challengers*. Wiley, 2010.

[15] "Cisco hybrid and public solutions." https://www.cisco.com/c/en/us/solutions/cloud/hybrid-cloud-public-cloud.html.

[16] "Dell EMC - IaaS." https://www.dellemc.com/sr-me/solutions/converged-infrastructure/iaas.htm.

[17] 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.

[18] D. P. Anderson and G. Fedak, "The computational and storage potential of volunteer computing," in *6th IEEE Int. Symposium on Cluster Computing and the Grid (CCGRID'06)*, vol. 1, pp. 73–80, 2006.

[19] Cisco, "Internet of Things." https://www.cisco.com/c/dam/en/us/products/collateral/se/internet-of-things/at-a-glance-c45-731471.pdf.

[20] "Openfog consortium." https://www.openfogconsortium.org/.

[21] J. S. Preden, K. Tammeme, A. Jantsch, M. Leier, A. Riid, and E. Calis, "The benefits of self-awareness and attention in fog and mist computing," *Computer*, vol. 48, pp. 37–45, July 2015.

[22] Y. Elkhatib, B. Porter, H. B. Ribeiro, M. F. Zhani, J. Qadir, and E. Rivière, "On using micro-clouds to deliver the fog," *arXiv preprint arXiv:1703.00375*, 2017.

[23] W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, "Edge computing: Vision and challenges," *IEEE IoT Journal*, vol. 3, no. 5, pp. 637–646, 2016.

[24] M. Villari, M. Fazio, S. Dustdar, O. Rana, and R. Ranjan, "Osmotic computing: A new paradigm for edge/cloud integration," *IEEE Cloud Computing*, vol. 3, no. 6, pp. 76–83, 2016.

[25] F. Bonomi, R. Milito, J. Zhu, and S. Addepalli, "Fog computing and its role in the internet of things," in *Proceedings of the first edition of the MCC workshop on Mobile cloud computing*, pp. 13–16, ACM, 2012.

[26] G. T. Heineman and W. T. Councill, "Component-based software engineering," *Putting the pieces together, addison-westley*, p. 5, 2001.

[27] T. Erl, *Service-oriented architecture*, vol. 8. Pearson India, 2005.

[28] N. Dmitry and S.-S. Manfred, "On micro-services architecture," *International Journal of Open Information Technologies*, vol. 2, no. 9, 2014.

[29] M. Lom, O. Pribyl, and M. Svitek, "Industry 4.0 as a part of smart cities," in *Smart Cities Symposium Prague (SCSP), 2016*, pp. 1–6, IEEE, 2016.

[30] G. Yang, L. Xie, M. Mäntysalo, X. Zhou, Z. Pang, L. Da Xu, S. Kao-Walter, Q. Chen, and L.-R. Zheng, "A health-iot platform based on the integration of intelligent packaging, unobtrusive bio-sensor, and intelligent medicine box," *IEEE transactions on industrial informatics*, vol. 10, no. 4, pp. 2180–2191, 2014.

[31] X. Hou, Y. Li, M. Chen, D. Wu, D. Jin, and S. Chen, "Vehicular fog computing: A viewpoint of vehicles as the infrastructures," *IEEE Trans. on Vehicular Technology*, vol. 65, no. 6, pp. 3860–3873, 2016.

[32] "Nebbiolo Technologies." https://www.nebbiolo.tech/.

[33] M. Aazam and E.-N. Huh, "Fog computing micro datacenter based dynamic resource estimation and pricing model for iot," in *Advanced Information Networking and Applications (AINA), 2015 IEEE 29th International Conference on*, pp. 687–694, IEEE, 2015.

[34] R. Deng, R. Lu, C. Lai, T. H. Luan, and H. Liang, "Optimal workload allocation in fog-cloud computing toward balanced delay and power consumption," *IEEE IoT Journal*, vol. 3, no. 6, pp. 1171–1181, 2016.

[35] M. B. Lieberman and D. B. Montgomery, "First-mover advantages," *Strategic management journal*, vol. 9, no. S1, pp. 41–58, 1988.

[36] M. E. Porter, "How competitive forces shape strategy," in *Readings in strategic management*, pp. 133–143, Springer, 1989.

[37] T. Donaldson and L. E. Preston, "The stakeholder theory of the corporation: Concepts, evidence, and implications," *Academy of management Review*, vol. 20, no. 1, pp. 65–91, 1995.

[38] "International Telecommunication Union." https://www.itu.int/.

[39] "Average internet connection speed in the United States from 2007 to 2017 (in Mbps), by quarter." https://www.statista.com/statistics/616210/average-internet-connection-speed-in-the-us/.

[40] "Docker." https://www.docker.com/.

[41] "Raspberry Pi 2 Model B - Raspberry Pi." https://www.raspberrypi.org/products/raspberry-pi-2-model-b/.

[42] "DockerCon Challenge." https://blog.docker.com/2015/09/update-raspberry-pi-dockercon-challenge/.

[43] T. Binz, U. Breitenbücher, F. Haupt, O. Kopp, F. Leymann, A. Nowak, and S. Wagner, "Opentosca–a runtime for tosca-based cloud applications," in *International Conference on Service-Oriented Computing*, pp. 692–695, Springer, 2013.

[44] "OASIS TOSCA." https://www.oasis-open.org/committees.

[45] Y. Duan, G. Fu, N. Zhou, X. Sun, N. C. Narendra, and B. Hu, "Everything as a service (xaas) on the cloud: origins, current and future trends," in *2015 IEEE 8th International Conference on Cloud Computing (CLOUD)*, pp. 621–628, IEEE, 2015.

[46] "European Union GDPR." https://eugdpr.org/.

[47] M. U. Farooq, M. Waseem, A. Khairi, and S. Mazhar, "A critical analysis on the security concerns of internet of things (iot)," *International Journal of Computer Applications*, vol. 111, no. 7, 2015.

[48] A. Brogi, G.-L. Ferrari, and S. Forti, "Secure cloud-edge deployments, with trust," *CoRR*, vol. abs/1901.05347, 2019.

[49] M. T. Fralix, "From mass production to mass customization," *JTATM*, vol. 1, no. 2, pp. 1–7, 2001.

[50] C. K. Prahalad and V. Ramaswamy, "Co-creation experiences: The next practice in value creation," *Journal of interactive marketing*, vol. 18, no. 3, pp. 5–14, 2004.

[51] P. Garcia Lopez, A. Montresor, D. Epema, A. Datta, T. Higashino, A. Iamnitchi, M. Barcellos, P. Felber, and E. Riviere, "Edge-centric computing: Vision and challenges," *ACM SIGCOMM Computer Communication Review*, vol. 45, no. 5, pp. 37–42, 2015.

[52] "The World Economic Forum." https://www.weforum.org/.

[53] W. E. Forum, "The future of jobs 2018." http://reports.weforum.org/future-of-jobs-2018/.

[54] P. Dunleavy, H. Margetts, S. Bastow, and J. Tinkler, "New public management is deadlong live digital-era governance," *Journal of public administration research and theory*, vol. 16, no. 3, pp. 467–494, 2006.

[55] J. M. Research, "Global market outlook 2019," *Technical Report*, 2019.

[56] W. E. Forum, "Digital transformation initiative." https://www.weforum.org/whitepapers/digital-transformation-initiative, 2017.

[57] J. D. Rowley, "Venture funding in iot startups continues to go up." https://news.crunchbase.com/news/q4-2018-closes-out-a-record-year-for-the-global-vc-market/, 2019.

[58] "Startup innovative (italia) – ministero dello sviluppo economico." https://www.mise.gov.it/index.php/it/impresa/competitivita-e-nuove-imprese/start-up-innovative.

[59] "Industry 4.0 – ministero dello sviluppo economico." https://www.mise.gov.it/index.php/it/industria40.

[60] "How the salaries of software engineers have evolved over the past 20 years." https://www.forbes.com/sites/quora/2016/04/14/how-the-salaries-of-software-engineers-have-evolved-over-the-past-20-years/, April 2016.