




# InnovAction

#3 | 2018



The best way to predict  
the future is to create it.

*Alan Kay*





# InnovAction

#3 | 2018







# The (digital) transformation

In this digital era, populated by the digital natives, telecommunication and media industries have been facing major disruptive and unpredictable forces, having to rapidly adapt and rethink their strategies while still delivering top quality service to their clients and value to all stakeholders. The real market battle today is not with the traditional players but with the global players, like the so-called GAFA.

To win this battle, to be the fastest in this race, CSPs need to transform themselves into DSPs - digital service providers - fuelled by data, by offering true digital services, with a unique, outstanding and differentiated customer experience that promotes the same feelings, the same pleasure despite being bought online or offline.

The needed transformation must address several dimensions along the journey to become an agile company, to compete in the digital arena and to lead in customer experience and digital services. This edition of the InnovAction addresses these dimensions and proposes some strategies for those seeking to succeed in this aggressive but quite exciting times and full of opportunities.

Once again, Altice Labs publishes a technological publication that intends to be relevant for all TELCO and IT service providers, as well as act as a vehicle to share the most relevant technologic developments and knowledge throughout the Altice Group's operations. Our commitment to technological innovation and the future is real, and this magazine helps to reinforce it.

I hope you enjoy reading it, as much as we enjoyed writing it.

**Alcino Lavrador**  
*General Manager of Altice Labs*

# Editorial note

All the telecommunications players have now realised the importance of metamorphosing into a digital service provider if they ever aspire to excel in this new digital era, where not only the customers but also the experience provided to them should be the focus. Having this in mind, and under the theme of the data-driven digital service provider, Altice Labs publishes the third edition of InnovAction, with a total of fifteen articles focusing on the following themes:

- **Entering the digital services industry: streamlining the CSP's metamorphosis:** summarises the way a service provider that wants to prevail in the digital services industry must be aware of the lifecycle of the digital services it intends to provide as well as to adapt them and adapt itself to the speed of the digital space.
- **Enterprise product platform enabling the digital services industry:** provides a vision on how to excel in the digital industry by leveraging the full power of an enterprise's vast product catalogue.
- **Autonomous cognitive management at the core of the DSP agile operations:** introduces Altice Labs' strategy and pragmatic approach to implement a new generation of cognitive operations support systems, a new generation of operations platform.
- **What strategy for information security:** exposes what should be the main concerns of DSP regarding the EU GDPR, namely, security by design, security in operations, secure third parties and security culture.
- **An approach to data monetization:** presents an approach that may improve the value of information and knowledge retrieved from data while, at the same time, opening new horizons for decision makers and promoters to take the most out of the data they have access to.
- **Virtual assistants, the future digital interface:** highlights the strategy that Altice Labs is following to address the virtual assistant trend (and reality): the BOTSchool platform. It will allow offering a new, simplified and more natural channel over the existing services.
- **Mining the TV consumption:** a basis for a TV personalisation system: points out a possible solution to mitigate a problem felt by TV viewers regarding the continuous growth of channels and content available in a typical interactive TV service.
- **Technology radar, for the win:** synthesises how Altice Labs addresses technological surveillance in order to better shape its product roadmaps and give a perspective about which technologies will have a higher relevance within the organisation in the near future.



- **Web Components, a compatibility solution in a cross framework environment:** addresses the advantages of web components as a technology based on standards defined by the W3C that allows for the creation of components with high potential of reuse and more immune to constant changes in web applications.
  - **Personalisation of smart cities powered by UX methodologies:** focuses the importance of understanding the journey experience of a citizen around the city as the only sustainable strategy to design a city that is smart for itself and for the ones that live in it.
  - **UX process for product co-design empowering personal engagement:** describes the implementation of a UX process in a specific Altice Labs' product, where techniques were applied with and to stakeholders, consolidating a strategy to bring value to a product supported on a matured user experience.
  - **NB-IoT@Altice Portugal:** analyses the Narrow Band IoT technology and discusses its characteristics and use cases of application from Altice Portugal's perspective.
  - **An overview of the benefits of super fast broadband networks to rural areas:** reviews the role and importance of superfast broadband networks in the development of the economy of the rural areas.
  - **Use cases: the 5G fuel:** clarifies some 5G use cases and its challenges, offering some suggestions to overcome them while creating new sources of revenues.
  - **The digital transformation of the central office:** gives a brief overview over the next generation central office, its current state of the art, as well as offers a proposal from Altice Labs to this ecosystem, particularly regarding the PON access.
- Lastly, it is essential to leave a note of gratitude and acknowledgement to all those who helped, once again, to build this publication. To all authors, for their contributions and for sharing the knowledge resulting from their research. To all technical and editorial reviewers that, even without being named in the magazine, did a relevant and meticulous job on improving the quality and excellence of all articles and, therefore, of the entire magazine. To our designer that, in a short period, did an excellent job on granting an appealing and coherent design across the entire magazine.
- Once again, InnovAction shows that, all together, it is possible to build a prominent and influential publication from Altice Labs to the Altice Group and all their customers and partners.

**Ana Patrícia Monteiro**  
**Pedro Carvalho**  
**Arnaldo Santos**  
*Altice Labs*

# Contents

## 01

Entering the digital services industry: streamlining the CSP's metamorphosis

(p. 8)

## 02

Enterprise product platform enabling the digital services industry

(p. 22)

## 03

Autonomous cognitive management at the core of the DSP agile operations

(p. 34)

## 04

What strategy for information security?

(p. 46)

## 05

An approach to data monetization

(p. 56)

## 06

Virtual assistants, the future digital interface

(p. 64)



## 07

Mining the TV consumption: a basis for a TV personalisation system

*(p. 74)*

## 08

Technology radar, for the win

*(p. 86)*

## 09

Web Components, a compatibility solution in a cross framework environment

*(p. 100)*

## 10

Personalisation of smart cities powered by UX methodologies

*(p. 108)*

## 11

UX process for product co-design empowering personal engagement

*(p. 116)*

## 12

NB-IoT@Altice Portugal

*(p. 126)*

## 13

An overview of the benefits of superfast broadband networks to rural areas

*(p. 138)*

## 14

Use cases: the 5G fuel

*(p. 144)*

## 15

The digital transformation of the central office

*(p. 158)*

## 8 Entering the digital services industry: streamlining the CSP's metamorphosis



# 01

## Entering the digital services industry: streamlining the CSP's metamorphosis

**Mário Rui Costa, Altice Labs**

[MCosta@alticelabs.com](mailto:MCosta@alticelabs.com)

**Paulo Pereira, Altice Labs**

[paulo-s-pereira@alticelabs.com](mailto:paulo-s-pereira@alticelabs.com)

The world is now digital, a challenge where the experience associated with the services is almost at the same level of relevance as the outcome of it, requiring service providers to reinvent themselves. The industry has realised that, and strong steps are on the way to address this challenge.

Starting from a new network evolution, moving into new levels of experience, up to addressing the ultimate challenge: how to make trade on this new paradigm as part of a much needed cultural shift?

---

### Keywords

Digital services; Digital experience; Experience-driven; Multi-sided business models

## Introduction

Millennials, generation Z, those are already the customers of today and tomorrow. They influence, directly and indirectly, the way purchases are made, and hence the way businesses position themselves. This is true for every other industry but is even so much more relevant for tech-based industries.

The communications industry is being urged into transforming itself under the risk of implosion if not doing so. As everyone in the industry knows, the footprint is moving away, both on the financial aspects but, even more important, on the way the relationship with customers is evolving into the perception of the value they have from the services they get in exchange of their money.

The world is now digital, a challenge where the experience associated with the services is almost at the same level of relevance as the outcome of it, requiring service providers to reinvent themselves.

The industry has realised that, and strong steps are on the way to address this challenge. Starting from a new network evolution, where speed and capacity go hand-in-hand with native support of digital services requirements; moving into new levels of experience, from a customer standpoint leveraged by new paradigms of business and operational management focused on bringing the success of the Internet age organisations into the communications industry, up to addressing the ultimate challenge: how to make trade on this new paradigm as part of a much needed cultural shift?

## The digital services industry challenge

### The digital space

Today, Communication Service Providers (CSP) are players of a tough game that goes way beyond communications. The previous challenge of enabling an always-connected lifestyle gave place to the challenge of allowing the digital engagement of persons and things, which take connectivity as granted.

The rise of the Internet as a global network, and the incremental digitalisation of everything created a new generation of digital engaged persons, also known as digital natives that aim to thrive in the digital space, exposed by digital channels and where, among the traditional and commoditised tasks, one can:

- Communicate using a mesh of voice, text and video, with persons, communities and non-human digital entities powered by artificial intelligence (AI);
- Manage digitally its body dynamics and its home place;
- Educate, be educated and work remotely engaging with teammates spread across the planet.

The multitude of digital services created to support the digital space encapsulates all the life dimensions of a person nowadays, enabling the extension of an actual organic existence into a new intangible dimension of reality, the digital space (see **Figure 1**).

One essential characteristic of the digital space is the lightning speed dynamics of digital services lifecycle. Some digital services are long-term companions (exist for a long time) of a digital native: this is the case of services used to support communication one-2-one or one-2-many; and some digital services (the majority of them)

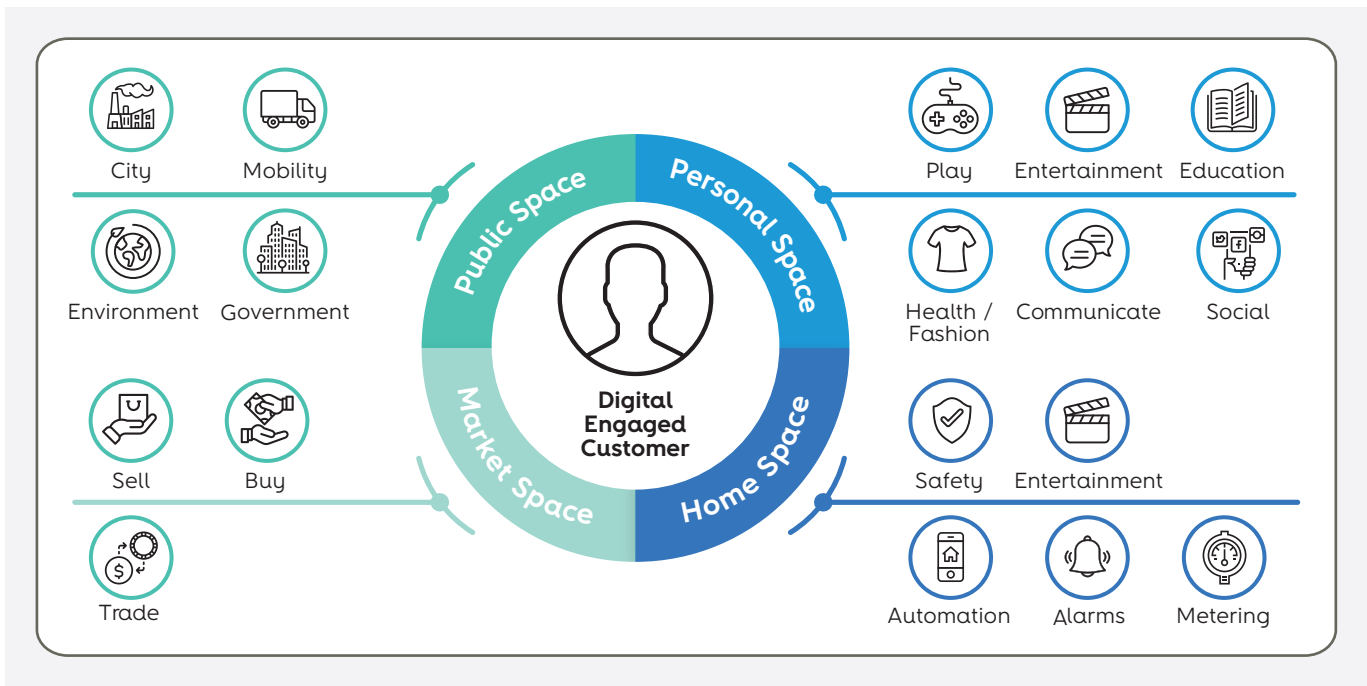


FIGURE 1 – The digital space service dimensions

are short-term companions (exist for short time frames) and are seen by a digital engaged person as a consumable/disposable asset (this is the typical case of some game applications or entertainment contents).

A service provider that wants to prevail in this industry must be aware of the lifecycle nature of the digital services it intends to provide and be prepared to work at the speed of the digital space, bringing and dawning services and updates at a pace never seen before for these players.

## The digital experience

Digital natives have key characteristics that will drive the choice of their Digital Services Providers (DSP). For a digital native, as important as having the right set of services is to have the right style of relationship with the service provider, fully supported on the top of omni digital channels. Therefore, the relationship they expect encapsulates a specific set of values that define the experience profile in the digital space – the digital experience, as seen in **Figure 2**.

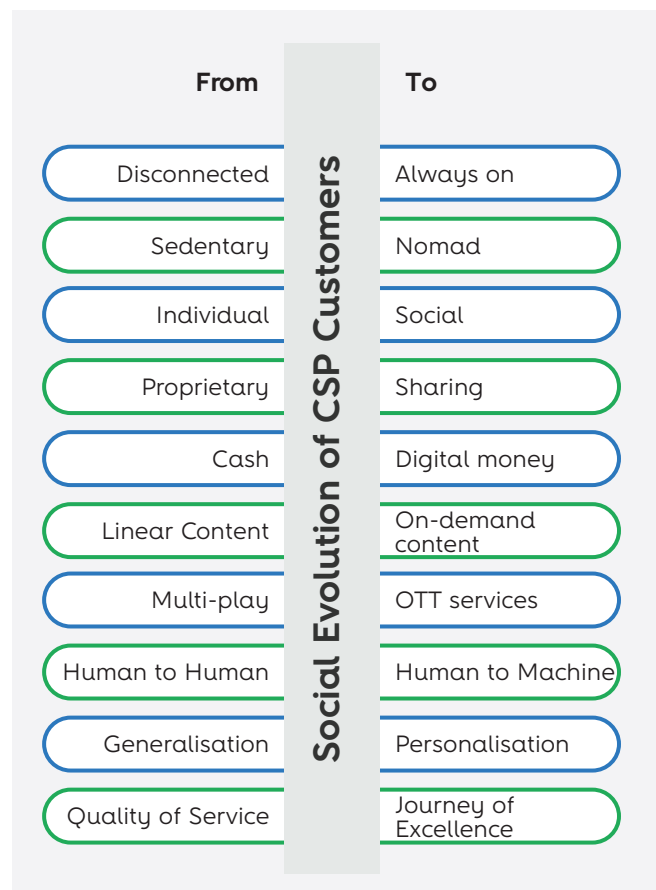


FIGURE 2 – The digital experience

# Evolving into the digital services industry: drivers, strategies and approaches

## Evolution drivers

With the communication industry blending into the digital services industry, facing the digital space and digital experience challenges, the service provider success will be highly dependent on how complete, integrated, rich and diverse will be its digital services portfolio.

Additionally, service providers face unprecedented growth in connectivity demand, either wireline or mobile, which leads to massive network evolution investments. The return on this investment will be in a very long-term and compromises the near and medium term business case, forcing service providers to run away (or be less dependent on) from a connectivity-driven business model into a business model where they may obtain more revenue from customers using their communications infrastructure and especially from (future) customers that do not use their connectivity but their services.

Simultaneously a driver and a challenge, corporate culture is also a key topic to address. Once the change is accepted, and new behaviours are adopted, corporate culture will become, just by itself a key driver to engaging into this new Digital wave.

## Business strategies

As mentioned above, evolving to the digital services industry requires changes in the existing services portfolio, creating the opportunity to reshape the business positioning and strategy of a service provider, namely through:

- **Expanding the digital relevance on the business:** A service provider may opt to leverage existing business segments and focus on expanding its digital relevance. This may be achieved through

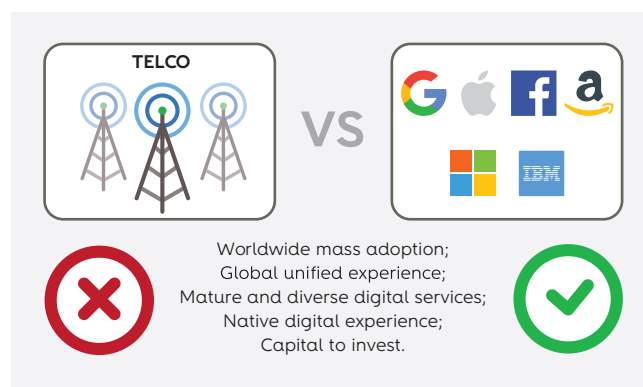
the expansion of current services portfolio with new sets of digital services and through the implementation of the expected digital experience for both types of services.

- **Expanding the support to new digital ecosystems:** governments, transports, cities, health, manufacturing, among others, are ecosystems being profoundly transformed by the rise of Internet of things (IoT) and digitalisation. Providing them unique digital platforms, with unique digital services, allowing to leverage the communications infrastructure of a service provider, thus entering a market that may reach USD 580 Billion by 2026 [1].

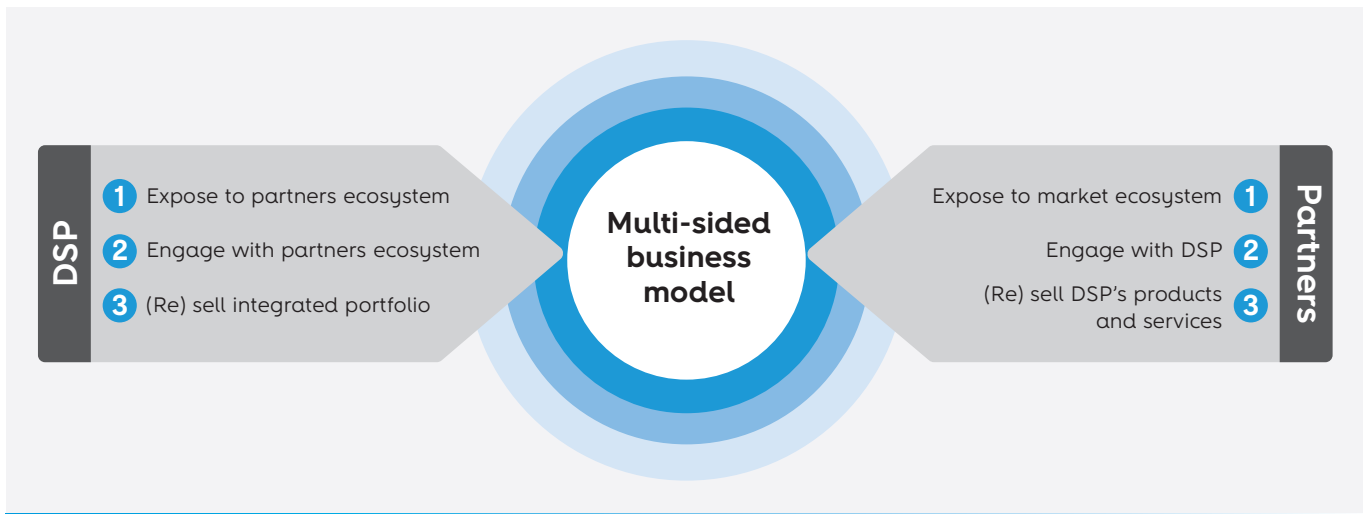
The above approaches are not mutually exclusive, i.e., a service provider may opt to adopt both business transformation strategies. Utilising a mixed approach will be the most common case, as the level of challenges and investment required for addressing new ecosystems will not allow easy wins on the short-run.

## Business approaches

An additional challenge that traditional CSP face when evolving to the digital services industry is the fierce competition of the already established "Internet age" companies, also known as Hyperscalers or over-the-top (OTT) organisations, like Google, Amazon, Facebook and Apple (GAFA), as well as big ICT companies like Microsoft and IBM – see **Figure 3**.



**FIGURE 3** – Digital space competition: TELCO vs GAFA-MI



**FIGURE 4** – Multi-sided business approaches

Service providers must quickly understand that trying to compete with these giant players alone, without even considering them in their evolution strategy, would for sure lead to significant business shrink in the medium to long-term. Thus, and as illustrated in **Figure 4**, a smarter approach is to engage in multi-sided business approaches, establishing partnerships not only with these giants but also with other ICT international companies, as a mean to a) exponentially grow the digital services portfolio, b) integrating digital services from partner companies, c) reduce capital expenditures required to implement new digital services from scratch and d) promote own services with reselling partners, augmenting through them the existing market presence.

This approach brings, as a drawback, operational margins that are entirely different than the ones the industry is used to. However, this is the lower price required to pay to progressively adapt to this changing world, to become relevant in the current Internet/digital services scene and conquer some additional time to adjust and create new offers for the digital space.

## Entering the digital services industry: key transformations

### Network evolution as a foundation for digital services

#### WiFi, the new primary need

Wi-Fi is probably one of the most silent transformations we have been seeing in the communication industry over the time. It started as a timid approach to mobility and has been winning space in everyday life in a way that it has now become one of the most required features wherever someone goes. Actually, it is now a primary need especially when focus on the current segment of digital natives and also on the above-mentioned need to develop and improve digital services (portfolios).

This means Customer Premises Equipments (CPE), present in the access layer, will need to dedicate increasing attention to Wi-Fi, supporting extension technologies, such as WiFi Mesh, and being able to integrate them into the cloud-based virtualised service execution chain to enable an

evolved and high-performance service experience. Why? Simply because digital natives have usage patterns that require additional monitoring, extreme personalisation, additional flexibility to share services among them, to add and drop new services quickly and in a self-sufficient way.

### Virtualised cloud-based service approaches, a must have reality

Alongside with the business approaches of big Internet age organisations, two other breakthrough paradigms are paramount for those organisations success: flexibility and agility. By using the underlying networks as an asset, they have created OTT or overlay solutions that excel on the capacity to adapt to the customer expectations and easily and quickly deliver what customer expects as service value, scaling as needed, seamlessly. Service providers competing in the digital space need to bring this level of flexibility and agility to its main asset: the network, since the existing ones are not up to the challenge, inasmuch as there is no easy way to isolate a new service, to scale its resources if it reaches a sudden success, to quickly introduce new service dependencies or even new services from partners.

Therefore, Software Defined Networks (SDN) are seen as the keystone component that will bring the capacity to introduce flexibility at scale for service providers, by focusing on isolating functional layers on a tiered approach and by pushing more sophisticated functions upwards into what we can see as a virtualised services cloud environment. In fact, SDN provides the foundation for future network concepts such as 5G, cloud central office, multi-access edge computing and network slicing. Relying strongly on concepts that are used throughout the Internet (cloud infrastructure, micro-services, orchestration, automation, API-driven exposure and programmable hardware) it brings along new flexibility levels that allow to achieve the required agility in the network and to improve the delivered service experience (e.g., multi-access edge computing with deployment of services closer to the user in edge data centres).

The network evolution challenges above mentioned extending naturally into what nowadays is known as B/OSS. A new generation of digital enabling solutions is required to enable, at the business and operations layers, at least the same levels of agility and flexibility made available by the service infrastructure. It's common sense that there is no point on having extreme agility in new service creation if there isn't the capability to expose those services at the market speed, with adequate costs and while providing the experience expected by the target customer segment. This has never been truer as when entering the digital space.

## Building a digital services architecture

### Digital services platforms

The already mentioned imperative to have digital enabling solutions with the levels of agility and experience required by the digital space and digital natives drives the need to define a new systems organisation (herein named platforms) and processes, i.e., a new architecture for digital enablement.

This architecture, which will coexist alongside with existing legacy B/OSS, is composed by six core platforms, each one focused on: digital experience, customer, product, service operations, monetization, data management and API management, illustrated on **Figure 5**.

Being the customer the most important business entity in any services industry and being the digital experience probably the most determinant factor for success in the digital services industry, the prime concern when defining this architecture was the customer relationship functionalities and the experience they provide. This principle drove, when setting the architecture, to start with the **digital experience platform** and to position it in the core of the architecture.

Remaining platforms of the architecture were designed to provide the required set of operations to manage the business entities used in the



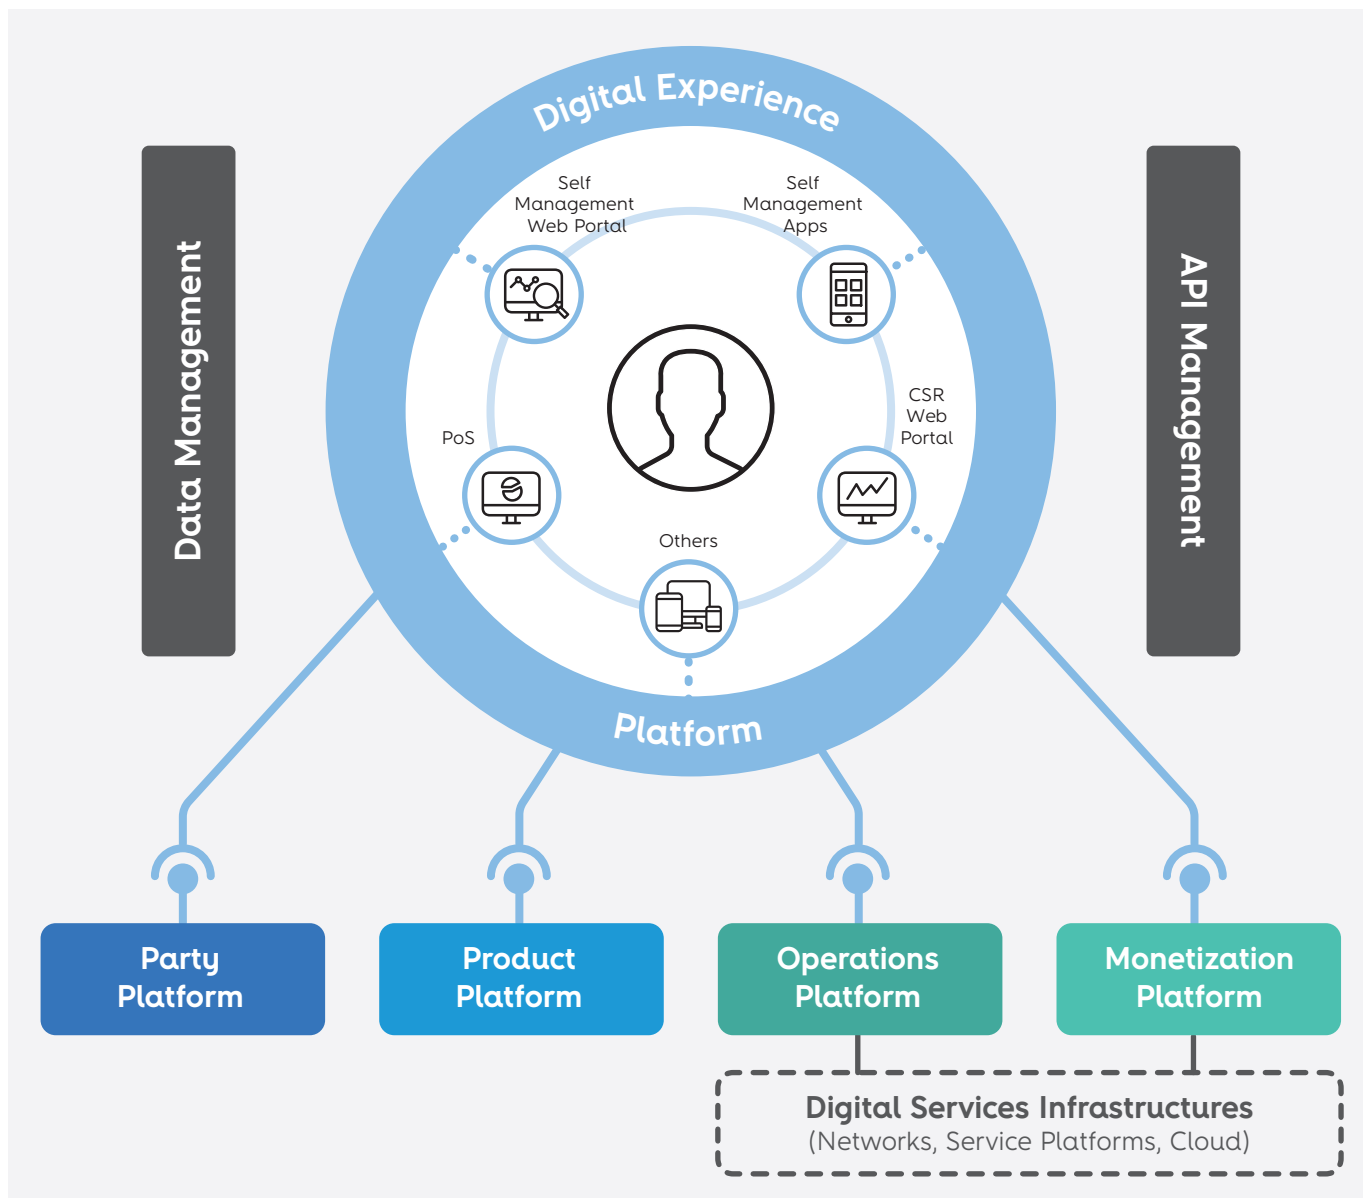


FIGURE 5 – The digital services platforms architecture

implementation of relationship processes in the digital experience platform while, at the same time, to support some critical requirements of the digital space and digital experience:

- **Party platform:** to manage customers and partners, promoting the agile engagement with partners in multi-sided business approaches;
- **Product platform:** to manage the lifecycle of the service provider’s commercial offers, supporting its agile blending with the partner’s commercial

offers. It also promotes the agility needed for the newly shortened lifecycle of blended offers;

- **Operations platform:** to manage the infrastructure required to deliver the services included in the service provider’s commercial offers, supporting unparalleled efficiency required to compete with the extremely demanding digital space competition;
- **Monetization platform:** to manage the full revenue chain, supporting the management of one-sided business agreements and multi-

sided business agreements revenues. It also provides the levels of flexibility to implement the diverse palette of monetization models required by the digital space.

Additionally to the platforms described above, two instrumental platforms were identified as required: the **data management platform**, aiming to be the data hub owning transversal information produced and used by remaining platforms, and that can then expose via APIs for further data monetization opportunities; and the **API management platform**, fundamental for assuring governance of API exposed to engage with partners, and to guarantee the levels of security and robustness in business interactions that cross organisation boundaries.

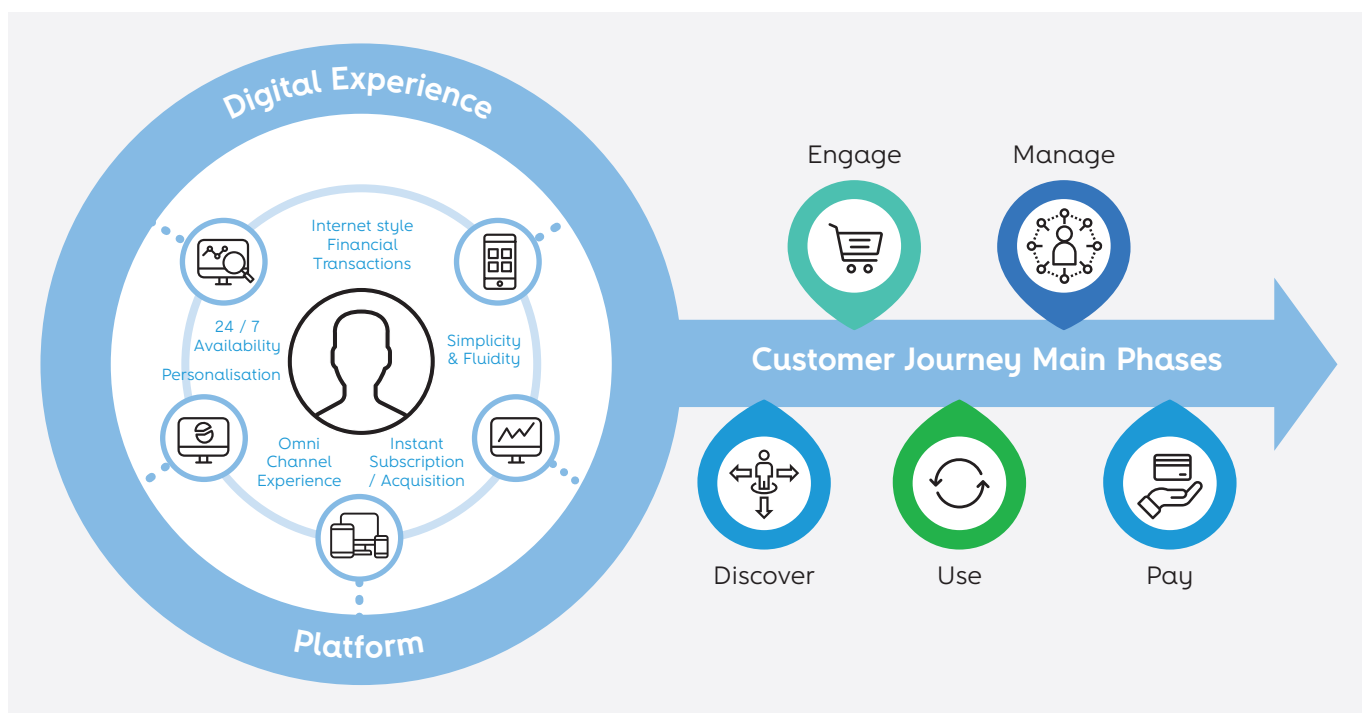
In fact, the API identified for all platforms already listed and below described correspond directly to a TM Forum Open API [2], and in this sense, this initiative was the primary inspiration to create the digital services platform architecture herein presented. TM Forum's suite of more than fifty REST based open API defines a standard specification and reference implementation for the

set of services required to create an ecosystem of digital services, where any organisation may participate as a provider, consumer or integrator of digital services, and where all business relationships are fully supported by programmatic means and automated. In this sense TM Forum's Open API were engineered to leverage a digital service API economy, maximising rapid integration, collaboration and innovation.

### Defining an experience-centric architecture – the digital experience platform

**Figure 6** details the **digital experience platform** mentioned in the previous chapter, the core platform of a DSP.

In a nutshell, the overall satisfaction of a digital native can be derived from two main dimensions: a) how I relate with my services provider; b) how the digital services I use fulfil my digital space needs. The first dimension encapsulates the discover, engage, manage and pay customer journey phases and is fully supported by the **digital experience platform**.



**FIGURE 6** – Digital experience platform

Being experience-centric, through its northbound interfaces this platform is able to:

- Support digital and omni channel requirements;
- Enable new interaction paradigms such as virtual assistants and augmented reality/virtual reality (AR/VR);
- Support interactions with high levels of availability, responsiveness, simplicity, convenience and personalisation, as demanded by digital natives.

Some of the above-outlined characteristics, namely availability, simplicity and personalisation cannot be guaranteed exclusively by the **experience platform**, insofar as having those characteristics imposes requirements to the enabling **business and operations platforms** (party, product, operations, monetization) and in concrete in the API exposed by those platforms.

## Enabling business and operations platforms

### Party Platform

The **party platform** is responsible for managing the actors of the digital space, exposing API to support the full lifecycle of the following business entities: party, accounts, agreements and loyalty programs.

The party business entity, managed by this platform, may represent the following business roles:

- **Customer** - an entity that establishes an agreement with the service provider with the intent to acquire/subscribe product or services offerings;
- **Partner** - an entity that establishes an agreement with the service provider with the intent to resell/supply product, services or resources offerings.

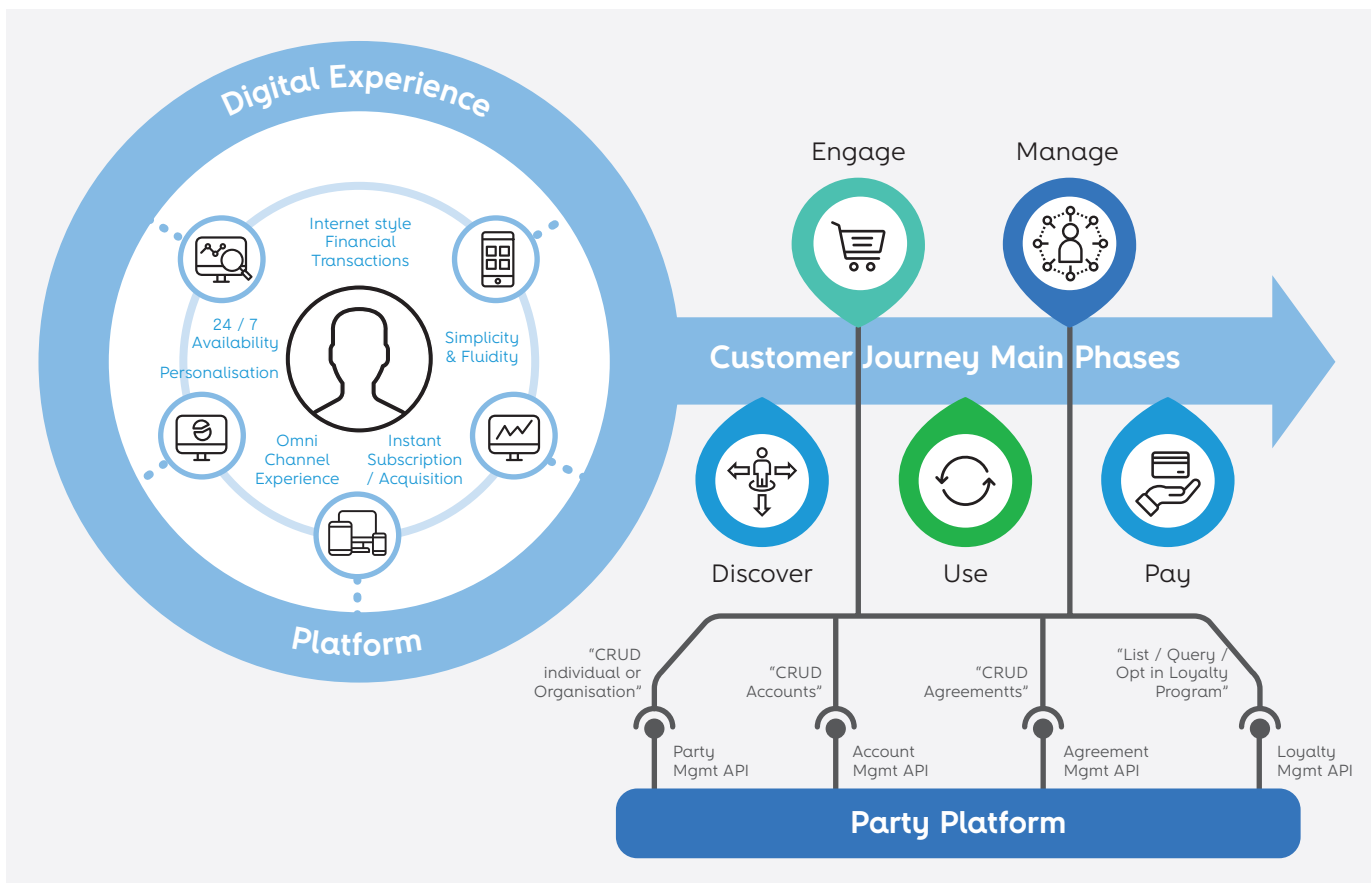


FIGURE 7 - Party platform supporting the engage and manage phases

**Figure 7** specifies the API exposed by the **party platform** and how they support the engage and the manage phase of a customer journey.

In addition to the requirements imposed by the **experience platform** (already described in the previous chapter), some essential conditions are paramount to be endorsed by the **party platform**, more specifically, the support to agile and automated engagement of partners since it make multi-sided business approaches a reality.

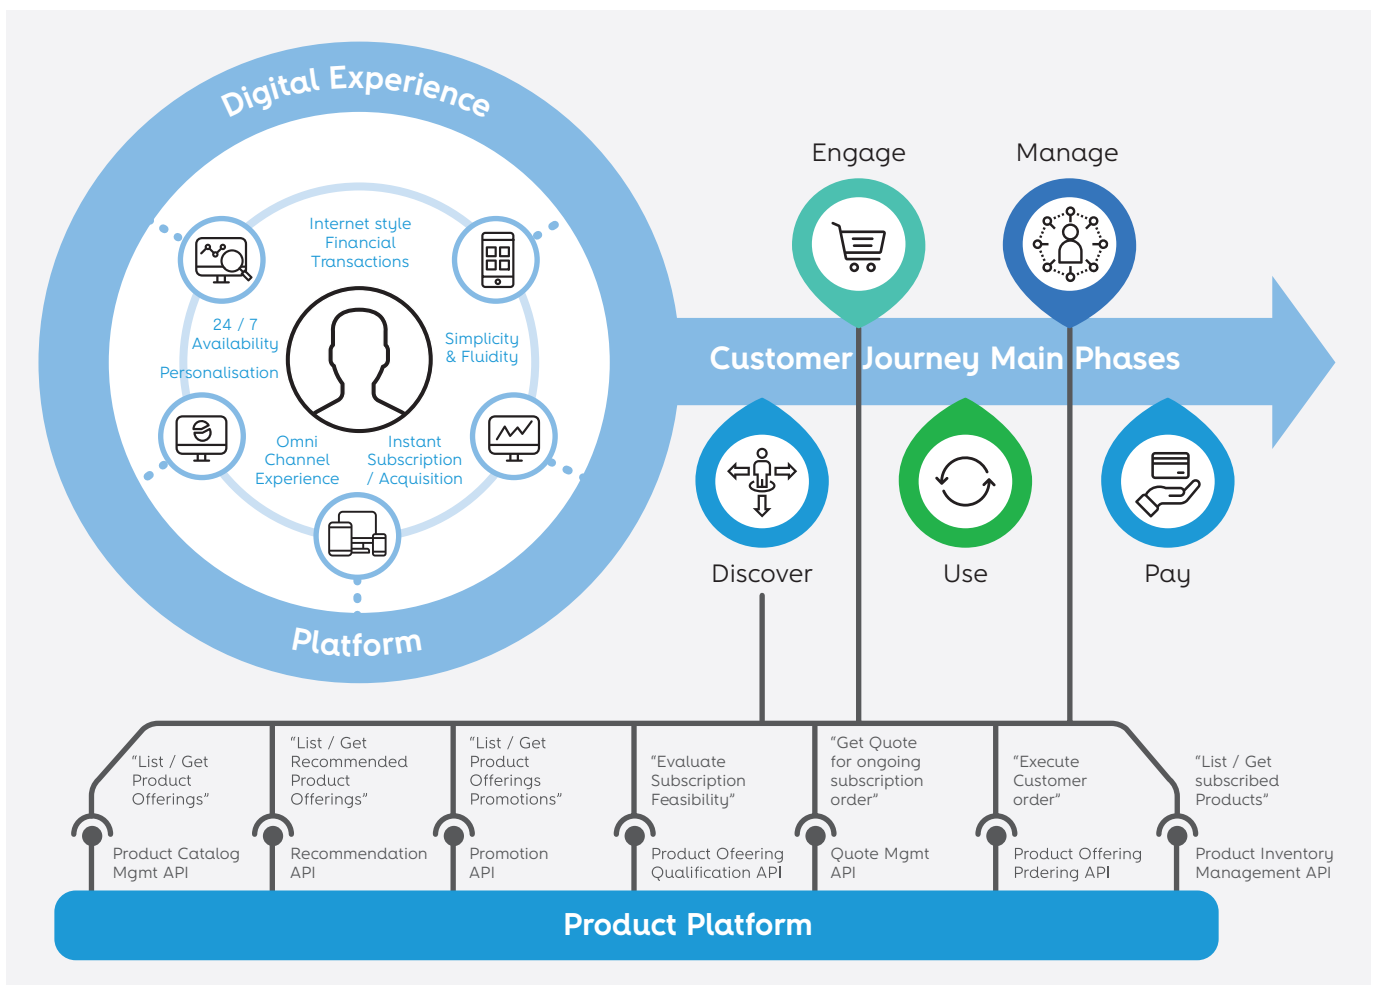
### Product platform

The **product platform** is responsible for mastering product offerings, the assets exchanged in the digital economy encapsulating digital services. For this, it exposes API to support the full lifecycle of the following business entities: product

offerings, products, recommendations, promotions and quotes.

**Figure 8** illustrates the API exposed by the **product platform** and how they support the discover, engage and manage phases of a customer journey.

In addition to the requirements imposed by the **experience platform**, some key requirements are paramount to be supported by the **product platform**, namely, supporting extreme agility to launch new product offerings and fully automated integration of partner's assets. These two requirements will ensure the business entities are up to date with the digital space market dynamics in a multi-sided economy space.



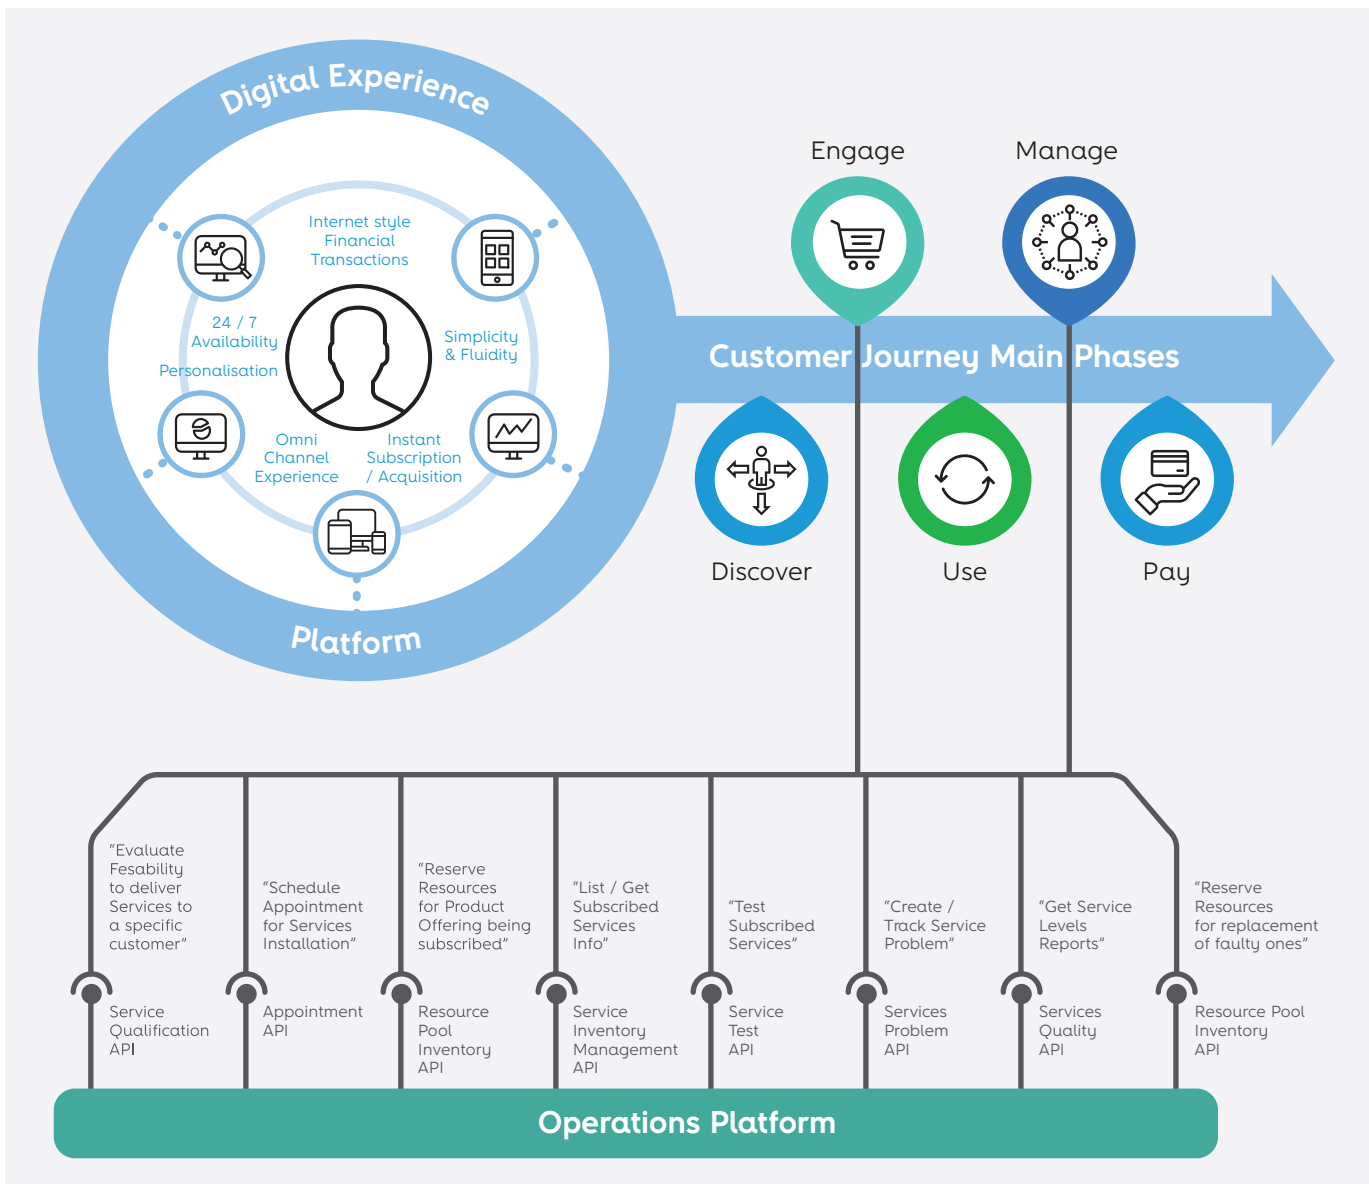
**FIGURE 8** – Product platform supporting the discover, engage and manage phases

### Operations platform

The **operations platform** is responsible for controlling the infrastructure required to deliver services to customers and partners, providing API to manage the next business entities: service, resource, service order, service qualification, appointment, resource reservation, service test, service problem and service level objective.

**Figure 9** lists the API exposed by the **operations platform** and how they support the engage and the manage phases of a customer journey.

In addition to the requirements imposed by the **experience platforms**, some key requirements are paramount to be supported by the **operations platform**, specifically, by supporting not only extreme agility to create new types of services (in line with the quickness expected on the product platform), using own resources or integrating partners assets, but also extreme efficiency to operate the services and underlay infrastructure which, in return, will create relevant operating income in the digital space.



**FIGURE 9** – Operations platform supporting the engage and manage phases

### Monetization platform

The **monetization platform** is responsible for mastering the balances and bills business entities, as well as other directly associated business entities (payment, topUp).

Finally, **Figure 10** illustrates the API exposed by the **monetization platform** and how they support the manage and pay phases of a customer journey.

In addition to the requirements imposed by the **experience platforms**, some key requirements are paramount to be supported by the **monetization platform**, in concrete, the flexibility to integrate the diverse palette of digital payment methods, payment institutions and monetization models, plus the ability to manage complex revenue schemes involving partners.

### Conclusion

Digitalisation and experience-driven approaches are the boosters of the service provider willing to adapt and thrive into the new world brought by Internet players and digital natives. However, with limited growth expectations, disruptions and convergence on the core business it means there is a need to move into unfamiliar areas. Business innovation is required.

Different approaches can be adopted, namely, acquisition of industry players thriving in adjacent market opportunities or genuine partnership at a go-to-market level with both reference Internet and content players or disruptive newcomers. Whichever combination

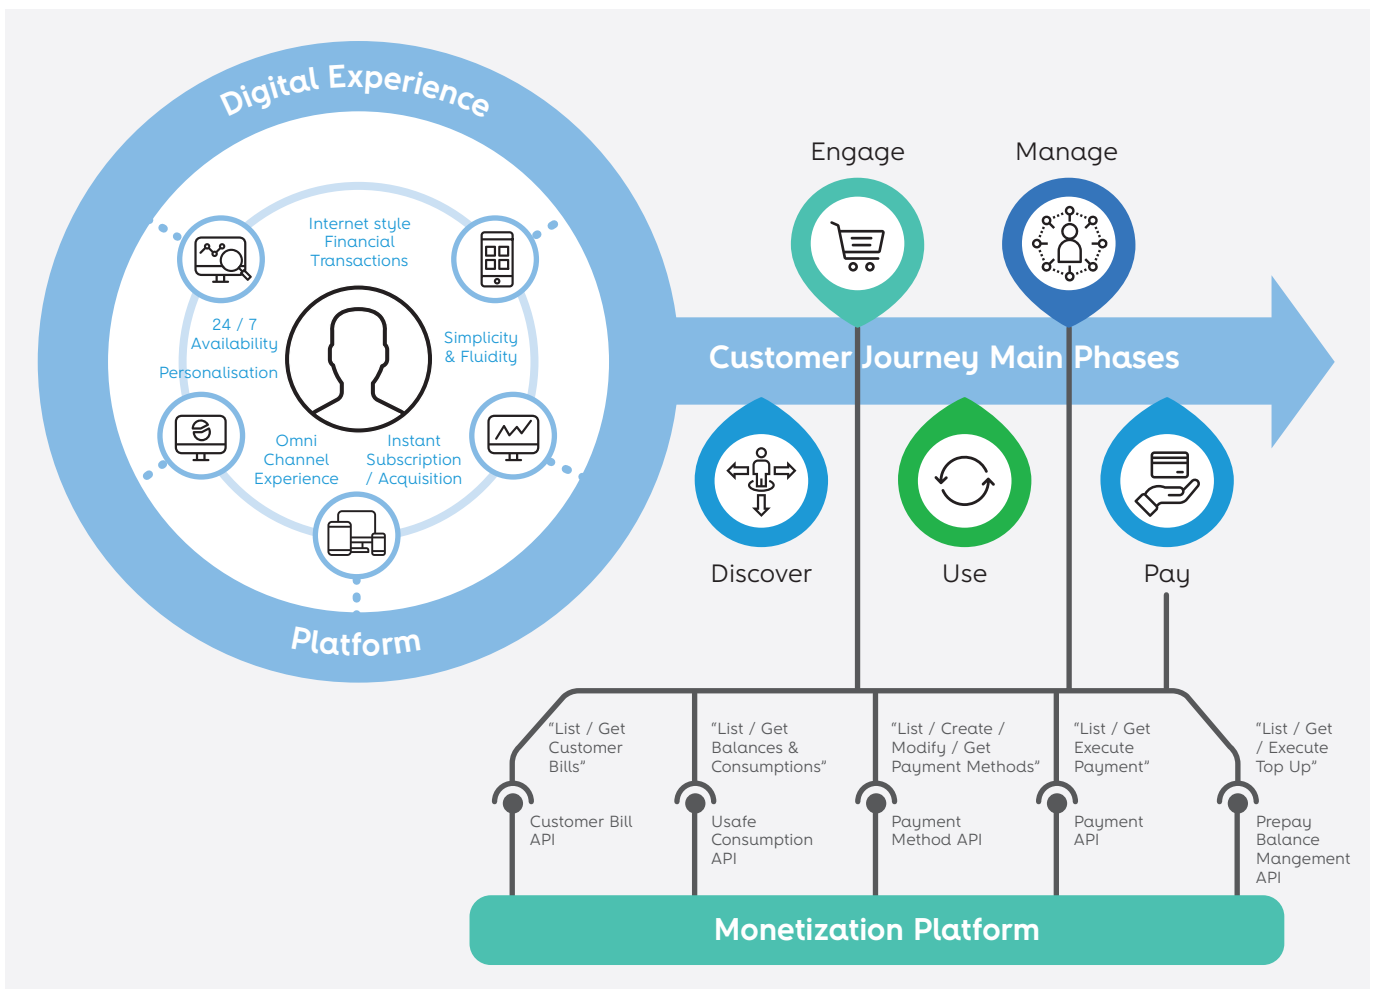


FIGURE 10 – Monetization platform supporting the manage and pay phases

of options is selected, a digital services platforms architecture is mandatory as it will bring along a toolkit of business and technology agnostic API and services that will enable the ultimate transformative point on this approach: new, multi-sided business approaches, promoted in an agile and Internet-like experience.

Service providers need to design in advance the experience they want to provide to customers, assuring an experience-centric relationship able to establish a strong brand culture where customers will become the first promoters and the first source of publicity. Experience is the primordial

requirement to customers: providing a good experience over digital services will allow to win in the digital space; providing inadequate experience is a business killer regardless of how excellent services are. The digital services platforms architecture is engineered to be experience-centric, aiming to assure a first-class experience.

If a service provider succeeds in combining the digital services platforms architecture with a well-performed network evolution, there is a pretty good chance that it will surpass the digital services industry challenges and thrive in this new amazing digital world, the digital space.

---

## References

- [1] TM Forum, "Open Digital Architecture," 2018.
- [2] TM Forum, "TM Forum Open APIs," 2018. [Online]. Available: <https://www.tmforum.org/open-apis/>.





# 02

## Enterprise product platform enabling the digital services industry

**Vitor Santos, Altice Labs**

[vsantos@alticelabs.com](mailto:vsantos@alticelabs.com)

**José António Almeida, Altice Labs**

[jose-s-almeida@alticelabs.com](mailto:jose-s-almeida@alticelabs.com)

**Pedro Nuno Pinto, Altice Labs**

[pedro-n-pinto@alticelabs.com](mailto:pedro-n-pinto@alticelabs.com)

**João Pias, Altice Labs**

[jpias@alticelabs.com](mailto:jpias@alticelabs.com)

Being a digital service provider means, among other things, to have unparalleled flexibility and agility to set up and provide amazing customer journeys, products and services, as required by the digital services industry. To excel in this new scenario, it is paramount that digital services providers can leverage the full power of an enterprise's vast product catalogue. The proposed article exposes this vision and elaborates on how service provider processes interact with a product platform to enable the digital services ecosystem.

---

### Keywords

Product platform; Catalogue-driven; Multi-sided business models; Customer engagement

## Product management challenges in the digital services industry

The telecommunications sector was built around the concept of delivering products, which provide services closely bound to the communications network infrastructure. However, nowadays, *“telecom companies face increasingly tough times as digitisation reshapes the industry landscape”*. Nevertheless, *“digitisation is not just a threat; it also offers telecom companies an opportunity to rebuild their market positions, re-imagine their business systems, and create innovative offerings for customers”* [1] that need to charm them. This is only possible by adding innovative products and services that a service provider, by itself and considering the market’s speed, is unable to deliver alone, leading to a search for partnerships. Dynamic product portfolio leveraged by partnerships.

One of the key aspects on how digital transformation can bring new revenues to telecommunications operators (TELCO), also seen as service providers, is the pursuit of new revenue streams in adjacencies, *“most companies are in the early stages of developing their digital-services portfolios”* and some of those portfolios are being built on top of partnerships, which bring external expertise, in a typical vertical integration of different solutions. Success, however, will only be possible with *“platform-based solutions”* that *“minimise costs in production and delivery by using automated and low touch processes”* [1] since *“becoming a platform for [a] TELCO isn’t easy, given the extensive legacy of systems, processes, and culture.”* [2]

As an example, *“let’s say a driver with a SIM card in the dashboard declares to the car ‘I need to buy coffee’”* [2]. In what way can one build a value-based solution that is able to fulfil a customer’s need? One would need a

car manufacturer, a communications service provider (CSP), a payment processor and a coffee shop to build a partnership. The entire solution would need to deliver a voice-activated service integrated into the car, with integrated navigation, which would use a SIM card, with an associated communications plan, and a coffee shop network that is able to process coffee requests with drive-through access. In this described scenario, the complexity of partner relationships can overgrow and, as such, a TELCO that attempts to implement this aggregated offer in legacy B/OSS platforms would find that it would need a large IT project to integrate everything and deliver a simple use case like this.

Each partner in the digital services ecosystem brings specific capabilities. The real value, however, is in the ability to create a relevant product offering that encapsulates pertinent digital capabilities from these partners. Additional value may come from the orchestration of the lifecycle of the product offering and associated services. Excelling in this orchestration and in the ability to automate and facilitate targeted commerce offers an opportunity to traditional telecommunications service providers.

## Products personalisation

One of the consequences of an active ecosystem where many partnerships may flourish, leading to a dynamic product portfolio with a multitude of value propositions, is the question of how customers can find the products they want. For example, an Amazon customer has access to as many as 562M different products. Even if divided into categories, the *“Clothing, Shoes, and Jewelry has the highest number of products: 166.28M products!”* and the lowest in the top 10 of categories has 16M products [3]. Several other Internet platforms have to handle with the same complexity, like for example Netflix, where Chris Jaffe, the vice president of product innovation, characterised the problem the company was facing: *“[we need to] solve the ‘rabbit hole problem’ - having so much content that users get lost”* [4].

So, the questions are: how to build customer-personalised experiences? How to automatically match a customer with the right product offerings to increase revenue? Service providers have access to a plethora of information regarding their customers, which allows profiling activities to match customer profile with the right product offers that fulfil customer needs and preferences. The solution may be delivered from the creation of products with a high degree of options, to provide customers with the ability to customise that product for their needs.

In a nutshell, the ability to use transversal customer information to target product recommendations and fine-tune product promotions, combined with the flexibility to create highly configurable product offerings, are key aspects to be supported by a product platform in the digital services industry.

## Products with immediate availability

Amazon's 1-Click Order was a game changer in the digital era. The idea that a customer *"could enter in their billing, shipping and payment information just once and then simply click a button to buy something going forward was unheard of when Amazon secured the patent in 1999, and it represented a breakthrough for the idea of hassle-free online shopping."* [5]

In addition, Salesforce conducted a survey to consumers and business buyers globally, where one of the options chosen by 69% of the respondents was: *"I expect Amazon-like buying experiences"*. [6]

This hassle-free online shopping idea has become so much a part of our lives that the World Wide Web Consortium initiated work on a standard to streamline the online "check-out" process in 2015 [7]. The main goal is to simplify the check-out or ordering process. Moreover, the same way a customer might expect the ordering process to be simple and hassle-free, the expectation regarding fulfilment will also be the same. In fact, PwC's Global Consumer Insights Survey found that over

40% of consumers are willing to pay extra to have same-day delivery, with 10% of respondents expecting that as the "normal" level of service. [8]

Also, if the product offerings are of digital nature, the expectation regarding fulfilment of such an order is that it will be instantaneous, following the Buy and Use Right Now (BURN) approach, in an automated way. However, since a particular product offering could be the combination of services and physical resources from different partners, and the product platform owner must drive their activation, the implementation of the BURN approach may be a true challenge. To overcome it, the platform architecture behind the (digital) product must be designed with seamless integration principles in mind.

## Creating a product platform for the digital services industry

### The product platform role

The traditional product platform used in the TELCO industry must evolve to respond to the digital services industry challenges outlined in the "Products personalisation" section, going way beyond product modelling.

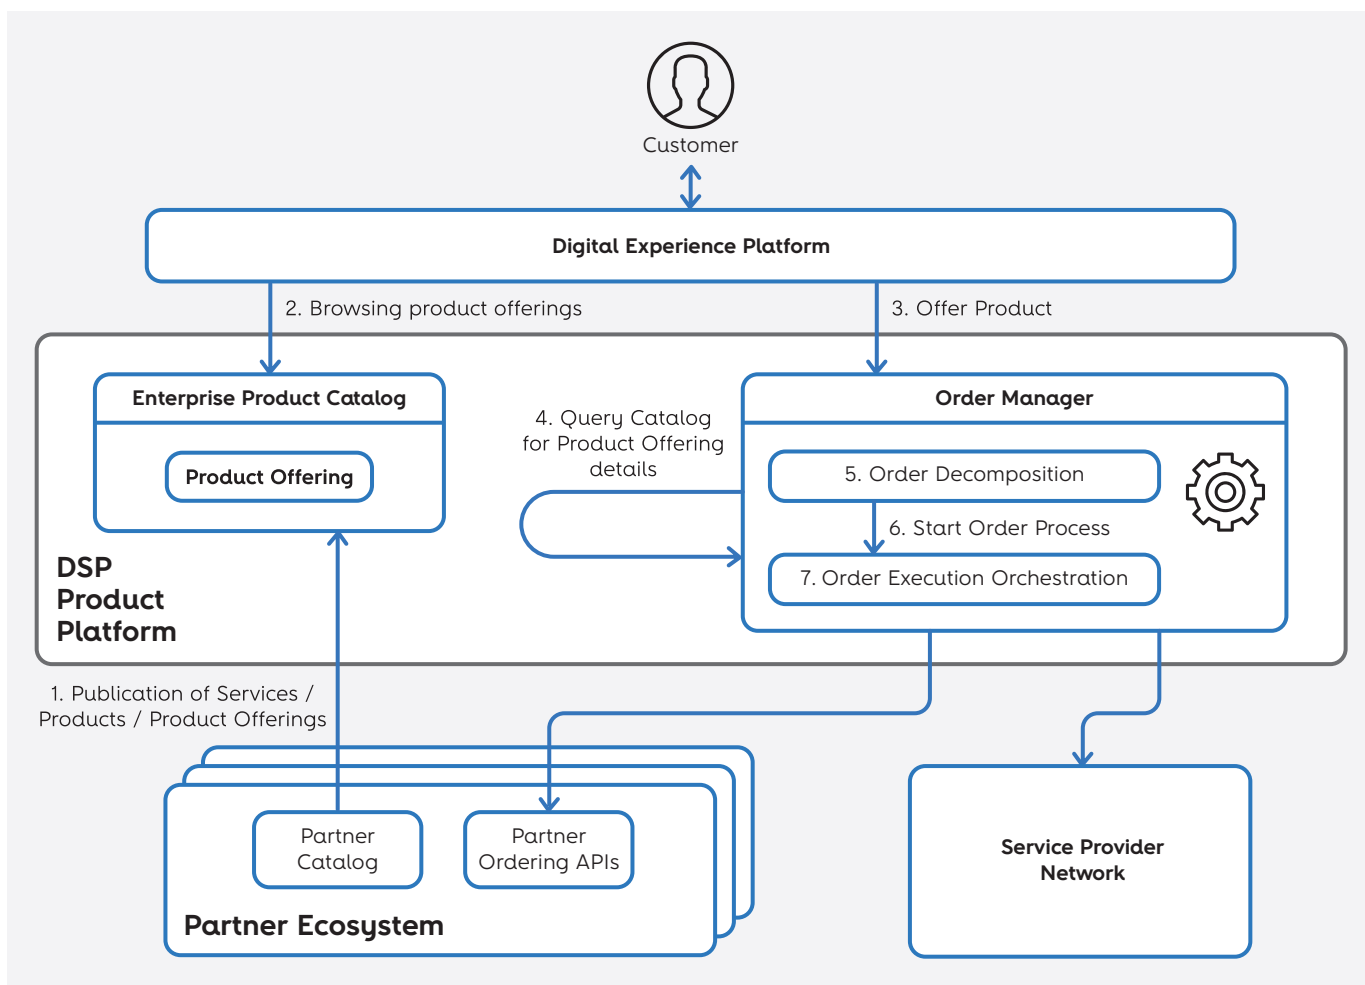
Most state-of-the-art product catalogues already support highly configurable product offerings where the end customer is able to choose its flavour among a diverse set of possibilities for each service being purchased, such as characteristics of voice, internet and content services that meet their personal preferences. However, the short cycles of service innovation in the digital services industry require the ability to create "ready to launch" highly configurable product offerings in hours or days, or at most within one or two weeks. As so, product platform

must evolve to a) support collaborative and agile work tools for marketing teams; b) ensure that when a new product offering is created, it is ready to be presented in the service provider channels, to be purchased and delivered to customers.

To assure this, the traditional product catalogue must evolve its functional scope and be able to instruct the orchestration of customer orders to purchase and manage product offerings inside product platforms. A catalogue-driven orchestration process, like the one described in **Figure 1**, is a generic process able to support the orchestration of provisioning tasks for any type of product offering without requiring coding of new workflows from scratch, and making possible to create a brand new service and pricing assembly without the need to mess up with provisioning workflows.

Thriving in the digital services industry requires from a service provider the ability to create a very relevant bouquet of integrated digital services and support, as shown in **Figure 1**, the integration of partner services when modelling their own product offerings while maintaining the levels of agility and personalisation outlined above. Besides this, the access to the right set of information and operations in key points of the interaction with a customer that is discovering product offerings, purchasing them or managing previously purchased ones, is vital in creating the digital experience expected by customers living in the digital space.

In short, the product platform is the nuclear system for the business management of a service provider. Such a platform is responsible for the support of the aforementioned key aspects of product management in the digital services industry.



**FIGURE 1** – Product platform with an effective catalogue-driven approach

## Enablement of new business models

The creation and growth of digital ecosystems have led to the emergence of multi-sided business models (Figure 2), instead of a traditional single party supply chain. This approach has significant implications for its participants, and for customers, in particular, the multi-sided business models offer a greater choice of services, richer experiences which may lead to higher levels of satisfaction.

Accustomed to be in the middle of many service chains, the CSP are naturally well placed to become the curators between producers and consumers, benefiting from their ownership of rich customer knowledge. A positive feedback loop effect occurs thanks to a network effect, in which the large community of a business platform increases the value produced for each producer. Nevertheless, for service providers' partners some challenges remain:

1. Effective collaboration between all parties requires transparency;
2. Tracking and performance control of each partner's contribution is necessary;
3. High degree of agility regarding partnering onboarding and processes is essential;

4. Interoperability of processes between all parties systems needs a standard set of API.

Thus, the product platform will be essential for the multi-sided business models since it will provide the following functionalities:

- Enterprise product catalogue - the commercial catalogue where the CSP will design product offerings by incorporating products and/or services provided by partners;
- Federation mechanisms for partner catalogues - exposition of catalogue entity management API that allows the publication of services and/or products made available by partners;
- Catalogue-driven product offerings provisioning - dynamic and automatic provisioning process guided by product offering information, able to orchestrate the provision of products and services in the service provider and partners domains.

Besides all the benefits already described, another advantage of a truly flexible platform is the possibility of creating new models with multiple financial flows, for example:

- Partners act as re-sellers of the assets produced by the service provider that may be integrated (or not) with their own assets;

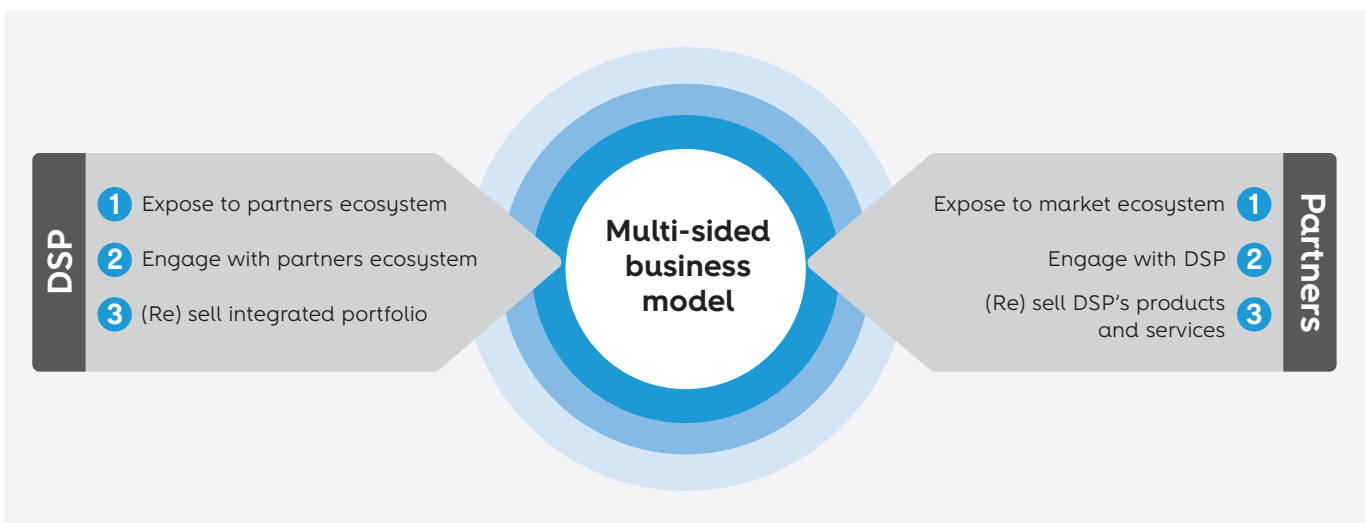


FIGURE 2 – Multi-sided business approaches

- Revenue sharing, where the revenue generated from sales is distributed to all partners that contributed with assets for the product offering;
- Service provider as a distributor of assets produced by partners.

## Enablement of the digital experience (platform)

### Building relevant offer recommendations and marketing actions

A business to thrive must fulfil the needs of its customers, which is only possible by raising product awareness and the creation of a sales channel. In the digital space, channels are guided by a digital experience platform that enables an omnichannel approach. The current challenge is how to optimise the customer-product-profit relationship in a way that will lead to customer loyalty by improving the customer (digital) experience.

From a technological viewpoint this challenge can be described in the following way [9]:

- Find the customers for the available offers: promotions and advertisements;
- Find products for a particular customer: search and recommendation;
- Determine optimal product properties: pricing and assortment.

The above steps should focus on customer information and product offerings information owned by the product platform to generate the following insights:

- Inbound interactions:
  - What offers must be proactively advertised and recommend in a channel?
  - When a customer browses on offers, what offers should be presented and how?

- Outbound interactions:
  - What promotions and recommendations should be notified to each customer?
  - What new offers should be created to fulfil customers' needs and interests?
  - What promotions should be created to maximise offer acceptance and revenue?

Traditionally, these questions (and the steps above listed) have used human insight and knowledge, but the reality is that with Machine Learning (ML) technologies it is now possible to optimise and ensure that the end results are more relevant. Or, as McKinsey stated in an article regarding modern marketing, *"technology has finally advanced to the point where marketers can use real-time data in a way that is both meaningful to customers and profitable for companies"* [10].

### Dynamic, online and contextual customer engagement

Whenever an existent customer, or even a potential new customer, engages through any digital channel, there is the expectation of a pleasant customer experience. In fact, *"three quarters expect 'now' service within five minutes of making contact online."* [11]. Failure to live up to that expectation will lead the customer to search for other competitors' products and solutions. 25% of customers who have a single bad experience will switch provider [12].

So, what defines a pleasant customer experience? Seamless omnichannel interactions, simple interfaces, fluid and responsive interactions and the ability of presenting relevant information for that specific customer and interaction context are the key ingredients. These characteristics should be expressed through the main customer journeys interaction phases, as described in **Figure 3**.

#### Discovery phase

In the discovery phase (see **Figure 3**), the customer has identified a need and is looking for information on products that might fulfil it. The customer shall use a channel that is controlled

by an enterprise product platform. This platform must ensure that the information presented in all channels is consistent, up-to-date and aligned with the available product offerings and applicable business rules, as well as it supports guidance and content information associated with products.

The activity “browse available products” presents a sales opportunity for custom-tailored product offerings using the following methods:

- **Recommendations** - the presentation of recommendations that specifically targets the customer and attempts to infer product suitability for him;
- **Promotions** - the presentation of promotions, on related browsed or previously acquired product offerings, with potential incentives/benefits may lead to a purchase/subscription.

### Engagement phase

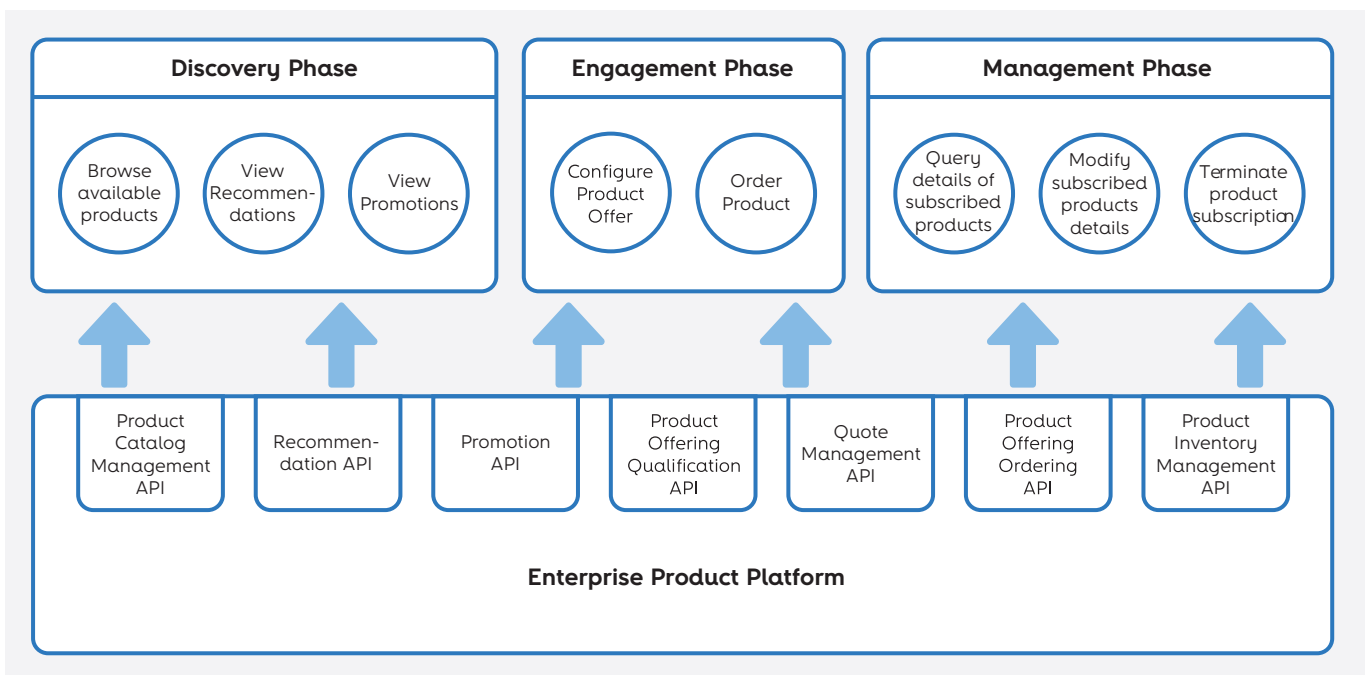
After the customer has discovered the right offer, an interaction may occur for the purchase, and this will be the beginning of the engagement phase. In the engagement phase the customer

should perform a sequence of activities (see **Figure 3**), which must also be supported by the product platform:

1. **Product offering configuration** - provide the right interaction processes with the required information to best characterise the product (optional components included) and allow its configuration. This will be used to generate a quote dynamically;
2. **Product order** - when the customer checks out its purchase, the interaction process will issue a product creation order to the product platform containing the shopping cart items that need to be provisioned.

### Management phase

After a customer has acquired a product, interactions to query, change characteristics or cancel subscribed/bought products occur in the management phase. The product platform is the cornerstone for this phase providing all required functionalities (see **Figure 3**) to support the distinct types of interactions that may occur at this phase:



**FIGURE 3** – API set exposed by the product platform to support the discovery, engagement and management phases of the customer journey

- **Product query** - the product platform controls the product inventory;
- **Product modification** - the product platform provides to interaction processes the relevant product offering information and options the customer may change;
- **Product cancellation** - the product platform also provides to interaction processes the relevant product offering information to support product cancellation orders.

### Automated fulfilment in multi-sided digital ecosystems

In multi-sided digital ecosystems with a mesh of interconnected partners, the complexity of the fulfilment process increases. At the same time, customer expectations remain the same, with the customer being oblivious and uninterested in this complexity. The usage of a standardised set of API, as proposed by TM Forum OPEN API's Initiative [13] is a toolset to solve this problem, as it defines a standard set of API to enable automated fulfilment of multi-partnered product offering.

The product platform must thus orchestrate product orders in the service provider systems, as well as partner's systems, ensuring they are also effectively configured, which means it must have the capability of controlling fulfilment in two independent domains.

### Real-time price quotes

A service provider playing in the digital space needs to optimise its sales process to excel in a continuously changing market where products may have significantly lower lifetimes than in the past. A digital marketplace should provide automated tools to increase sales efficiency and speed, thereby eliminating manual quotation processes.

The product platform provides the means to accelerate sales through a set of features, including the automatic creation of dynamic real-time quotes to inform customers during the product offering configuration stage (engagement phase),

eliminating errors and ambiguity, with real-time information on the pricing impacts of customer choices.

The product platform provides a quote management API (see **Figure 3**) that supports quote requests for the entire contents of a shopping cart, including the configuration of each included product offering, and applicable cross offer business rules. This capability is paramount to maintain the end customer always and accurately informed on the charges due, independently of offer complexity.

### Seamless integration

For the product platform, the ability to establish and manage new business relationships depends on the existence of standard shared contracts (that acts like templates) used by all partners in the digital ecosystem. Through those contracts, it is possible to establish dynamically new business agreements and to effectively integrate assets between partners in the digital services chain automatically, without human intervention. This defines the "seamless integration" paradigm in the digital services value chain, which is the key ingredient to achieve minimum time-to-market in such complex ecosystems.

Once again, TM Forum Open APIs [13] are the toolset to implement this seamless integration, providing the set of standard operations and shared information model required to put a digital ecosystem to work. The set of APIs that needs to be exposed by a product platform, to support the full lifecycle of product offerings while respecting the digital experience values, is represented in **Figure 3**.

## Conclusion

A product platform provides a set of information and functions that help to shape how a service provider competes and collaborates with both traditional CSP and also with the already well established over-the-top organisations.



This article laid the most important features a product platform is meant to have to be a catalyst for customer satisfaction, partner engagement and digital relations in general. Well implemented, the product platform is a strategic asset to seize the opportunities of the digital services competitive landscape, enabling the business models and customer experience expected by digital natives.

A product platform supported on the top of standard Open APIs is the key to cope with the ever growing business dynamics without becoming a bottleneck and a source of frustration. Altice Labs is strategically following this path, evolving its legacy product catalogue and customer order management solutions to become a product platform as a key enabler for the digital services industry.

## References

- [1] P. Caylar and A. Ménard, "How telecom companies can win in the digital revolution," *Digital McKinsey*, 2016. [Online]. Available: <http://www.mckinsey.com/business-functions/digital-mckinsey/our-insights/how-telecom-companies-can-win-in-the-digital-revolution>.
- [2] A. Behan, "Platform Economics and The Platformification of Telco," *IBM Telecom, Media and Entertainment Blog*, 2014. [Online]. Available: <https://www.ibm.com/blogs/insights-on-business/telecom-media-entertainment/platform-economics-and-the-platformification-of-telco/>.
- [3] "How Many Products Does Amazon Sell? - January 2018," *ScrapeHero*, 2018. [Online]. Available: <https://www.scrapehero.com/many-products-amazon-sell-january-2018/>.
- [4] L. O'Reilly, "Netflix lifted the lid on how the algorithm that recommends you titles to watch actually works," *Business Insider*, 2016. [Online]. Available: <https://www.businessinsider.com/how-the-netflix-recommendation-algorithm-works-2016-2>.
- [5] R. P. Wagner and T. Jeitschko, "Why Amazon's '1-Click' Ordering Was a Game Changer," *Knowledge@Wharton*, 2017. [Online]. Available: <http://knowledge.wharton.upenn.edu/article/amazons-1-click-goes-off-patent/>.
- [6] "What are customer expectations, and how have they changed?," *Salesforce*, 2016. [Online]. Available: <https://www.salesforce.com/research/customer-expectations/>.
- [7] "W3C Starts Web Payments Standards Work to Streamline the Online 'Check-out' Process," *W3C*, 2015. [Online]. Available: <https://www.w3.org/2015/09/webpaymentswg.html.en>.
- [8] "Signed, sealed, delivered (and regularly returned) - Global Consumer Insights Survey 2018," *PwC*, 2018. [Online]. Available: <https://www.pwc.com/gx/en/industries/consumer-markets/consumer-insights-survey/delivery-expectation.html>.
- [9] I. Katsov, *Introduction to Algorithmic Marketing*. Grid Dynamics, 2017.
- [10] J. Boudet, B. Gregg, J. Heller, and C. Tufft, "The heartbeat of modern marketing: Data activation and personalization," *McKinsey & Company*, 2017. [Online]. Available: <https://www.mckinsey.com/business-functions/marketing-and-sales/our-insights/the-heartbeat-of-modern-marketing>.

- [11] "Customer experience: New capabilities, new audiences, new opportunities." McKinsey & Company's, p. 104, 2017.
- [12] Thunderhead, "The Cost of Crappy Customer Experiences Infographic," 2015. [Online]. Available: <https://www.thunderhead.com/blog/customer-engagement/the-cost-of-crappy-customer-experiences-infographic/>.
- [13] "Leading Service Providers Drive Adoption of TM Forum Open APIs to Create New Digital Ecosystem Growth, Enabling Rapid Creation & Delivery of Innovative Digital Services," *TM Forum*, 2016. [Online]. Available: <https://www.tmforum.org/press-and-news/open-apis/>.





# 03

## Autonomous cognitive management at the core of the DSP agile operations

**David Murargi, Altice Portugal**

[david-l-murargi@telecom.pt](mailto:david-l-murargi@telecom.pt)

**Luís Castro, Altice Labs**

[luis-a-castro@alticelabs.com](mailto:luis-a-castro@alticelabs.com)

**Pedro Miguel Neves, Altice Labs**

[pedro-m-neves@alticelabs.com](mailto:pedro-m-neves@alticelabs.com)

Legacy OSS are becoming operations platforms able to achieve the levels of configuration and automation required to support the new business dynamics of the digital space in a context of software-defined networks. At the same time, OSS are becoming cognitive-enabled, able to shortcut assurance and fulfilment processes with artificial intelligence, contributing to providing unparalleled levels of experience expected by digital natives. This article introduces Altice Labs' strategy and pragmatic approach to implement a new generation of cognitive operations support systems, a new generation of operations platform.

---

### Keywords

SDN; NFV; Autonomous cognitive OSS; Operational agility; NetOps

## Introduction

Current human, labour-intensive, time and resource-consuming network management tasks, carried out by operations teams, should be conducted autonomously by following a much more efficient, intelligent and automated machine-oriented approach. Network transformation and the evolution from the traditional communications industry to the digital services industry demands a paradigm shift in the way networks are planned, deployed and operated.

The ongoing network transformation towards network functions virtualisation (NFV) and software-defined networks (SDN) is becoming a reality. This transformation relies on automated and scalable network infrastructure to meet the exponentially-increasing demands on mobile and fixed broadband access, concerning the number of connected users, required bandwidth and coverage, and bringing along service agility and operational costs saving opportunities. However, it also brings much more complexity to existing operations processes: higher data volume, lower latencies, more connected devices, frequent changes of service topologies, among others.

The evolution to the digital services industry imposes additional challenges on the agility, flexibility and robustness needed to manage the lifecycle of services and underlying resources. In the digital services industry, service innovation cycles are becoming much shorter, service activation is expected to be instantaneous, and services are supposed to be always available and highly responsive. Due to these requirements, the digital services industry is incompatible with human dependencies on service activation, optimisation and recovery situations.

The primary goal of this article is to describe Altice Labs' vision towards autonomous cognitive-driven operations and how Altice Labs is creating the operations platform architecture able to cope with the above challenges.

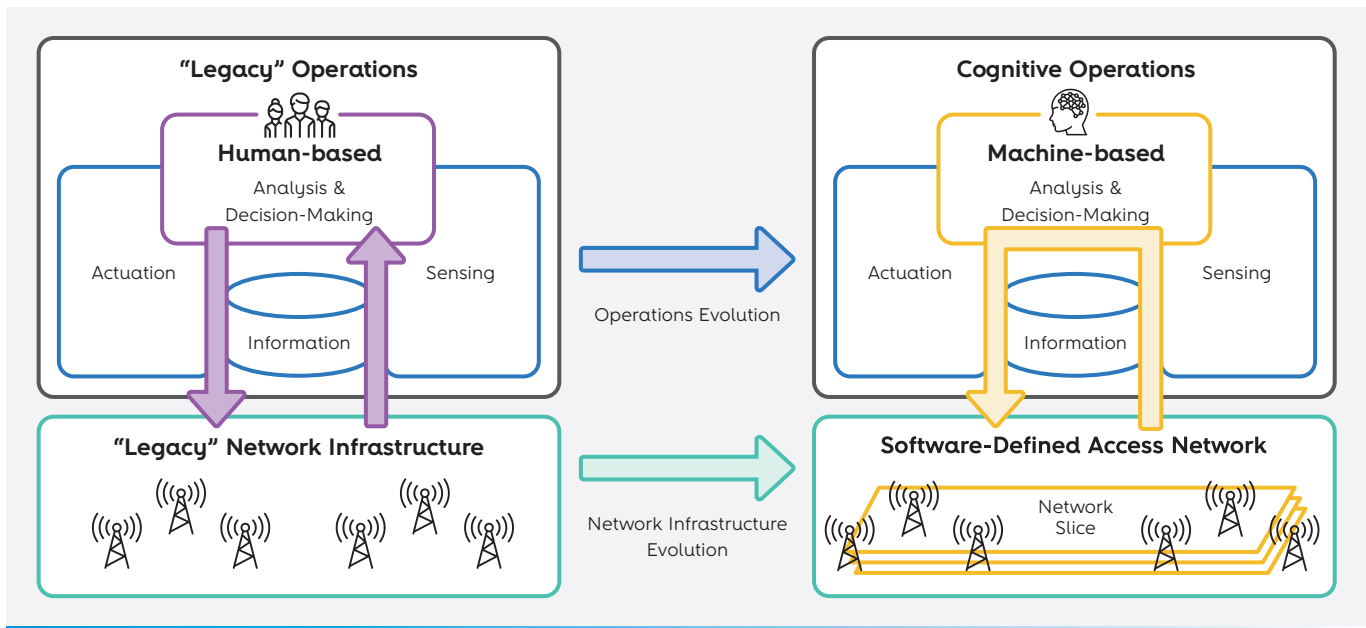
## Challenges towards autonomous cognitive operations

### Implementing process intelligence and automation

The current management paradigm, illustrated in the left-hand side of **Figure 1**, poses a number of challenges to network operators. Nowadays, operators have to do their best to detect and mitigate all sorts of problems in the networks, such as link failures, performance bottlenecks, cyber-attacks, quality of service (QoS) degradation, software bugs, and hardware faults, among others. Due to the human-based process, teams are organised by technology silos that are prepared to handle all sorts of requests and analysis for that domain. Existing solutions typically require to manual equipment reconfiguration or even the deployment of new equipment and functionalities. These tasks cannot be performed without affecting, even for a limited time, the normal operation of the network, which causes service disruptions and violations in service level agreements (SLA), thereby incurring in increased operational and capital costs, and compromising end users' quality of experience (QoE).

Recent research in SDN and NFV resulted in the emergence of new capabilities that significantly improve agility, flexibility and cost efficiency to manage network functions, which are the foundations to trigger a paradigm shift in the way network operations are planned, deployed and managed, called cognitive network management. This approach consists of implementing machine-based intelligence to support the creation of autonomous processes to manage complex networking scenarios.

One of the main impacts of this paradigm is the significant reduction of operational costs. Proactive and reactive actions are automated to resolve or mitigate networking problems, thereby



**FIGURE 1** – Towards a cognitive operations paradigm

minimising the human effort in maintenance and troubleshooting tasks, and leading to significant opex reduction. This paradigm is illustrated in the right-hand side of **Figure 1**.

Three key capabilities must be provided to create a solution able to support the implementation of autonomous, cognitive and highly intelligent processes:

- **Automated network monitoring (sensing):** the key challenge is to build transversal and automated monitoring capabilities crossing all network domains. These can be achieved through the automated deployment of virtualised probes in the network infrastructure to facilitate system-wide distributed monitoring. Virtualised probes are activated across the access and backbone network infrastructures to enable end-to-end user, service and network awareness through the collection of metrics from all network elements. Collected information must feed data analysis algorithms, like data analytics, data mining or machine learning (ML), in order to create key indicators that may translate to i) service affecting conditions (e.g., network failures, performance bottlenecks, security breaches,

intrusions); ii) conditions that may evolve to service affecting issues in the future; iii) non-optimal service delivery to specific users, i.e., detection of cases where the service topology being used to deliver a service to end users can be optimised in order to minimise allocated resources or the service QoS.

- **Cognitive framework:** the ability to define high-level tactical corrective and preventive measures to respond to the diagnosed conditions. Tactical measures may correspond to reactive actions to fix or mitigate existing network issues or may correspond to proactive actions to prevent the evolution of the diagnosed condition to an effective service-affecting anomaly. These actions typically correspond to services and/or network functions lifecycle management requests (e.g. automated instantiation, configuration, scalability, reconfiguration of connectivity logical topology, etc.).
- **Automated & dynamic service provisioning:** automated and intelligent processes to manage the lifecycle of services and network functions. These comprise the dynamic selection of the best locations and resources for services deployment (or migration) considering the

requirements associated with the specific service instance being provisioned (for instance the contracted QoS). This process also includes the provisioning of the key performance indicators to be produced for the service instance.

## Future mode of operations – NetOps culture transformation

By itself, the process automation and intelligence mentioned above will not suffice to achieve the levels of operational efficiency and customer experience required by the digital services industry. A transformation in legacy network operations work structure and work processes is required [1]. Legacy network operations usually work on the top of “silo-oriented monitoring”, i.e., typically, they operate as a sum of discrete units focused on meeting operational level targets for each technology silo (like networks, service platforms, IT infrastructure, ...). This is mostly a consequence of how, generally, legacy monitoring tools are engineered to manage the network and IT elements in separate domains.

Traditional silo-oriented work processes impede the creation of a clear, useful model, one that translates how combinations of monitored information over multiple dimensions (time, space, etc.) can relate with customer services and how they impact on his quality and perceived experience. The absence of these models brings two business threatening impacts:

1. the service provider is not aware of, and does not control how its operation is impacting quality and experience;
2. the service provider does not know what information to value from the immenseness of collected events and metrics.

The above situation established a culture in network operations centres (NOC) based on firefighting all threats individually without precise knowledge of priorities.

A transformation in operations work processes must be accomplished to implement customer

centric-operations processes. Such operations processes are defined so that each customer's end-to-end service is managed by dedicated teams as a whole, making the operations work structure aligned with the business and the services provided to end customers.

## Autonomous cognitive operations use-cases

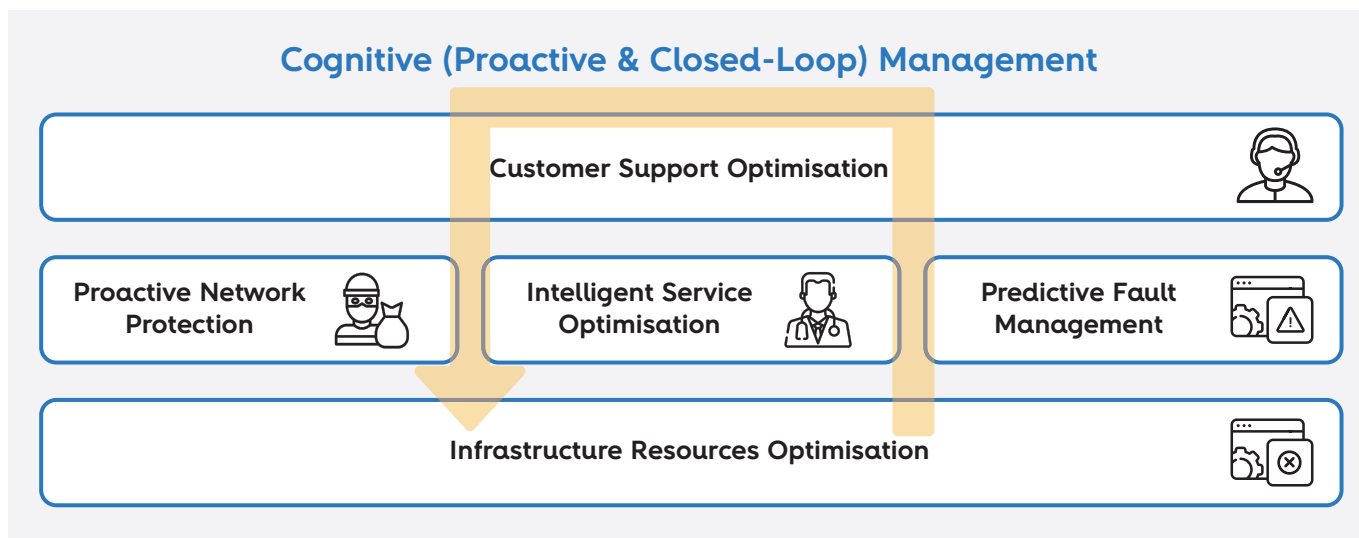
Following the major challenges towards a cognitive-enabled management system presented in the previous section, herein we present a set of use-cases that highlight the autonomous management capabilities of a digital service provider (DSP) enhanced through artificial intelligence (AI) techniques [2]. These use cases, illustrated in **Figure 2**, aim to achieve proactive/preventive measures against potential or forthcoming network problems in terms of fault/failure risks (predictive fault management), cyber-attacks (proactive network protection), QoS/QoE maintenance and enhancement (Intelligent Service Optimisation), infrastructure issues (infrastructure resources optimisation) and improve the customer support (customer support optimisation).

The successful design and implementation of the use cases will significantly improve the overall performance of future networks in reliability, availability, security, user experience, etc., all of which will jointly contribute to reduce service creation time and opex/capex, to minimise service downtime, and to optimise users QoE.

### Predictive fault management

As operators move to SDN/NFV deployments, the complexity with carrying out their operations across both physical and virtual elements increases. The reactive, time consuming and knowledge-based processes used today will soon be inadequate, and one of the major impacts will occur on service assurance domain.





**FIGURE 2** – Cognitive operations use-cases

Traditionally, the network operation processes use events, alarms and performance KPI that sense the network and drive the day-to-day activities to identify and troubleshoot problems, reach their root causes and mitigate the impact of resource failures on the services they operate. More advanced processes and tools that aid network operations teams use complex correlation rules to reduce the number of alarms and increase the accuracy in the identification of network issues/constraints. Many of these issues automatically generate trouble tickets (TTK) that are delivered to the appropriate team to solve them, thus minimising service impact. This process is based on known patterns and current knowledge across the teams for the technologies, services and vendors that defines the network ecosystem of each operator. This reactive response approach leads to inefficient team management and aggressive support SLAs (activated to minimise the impact of a problem). Today, the major effort in network operations processes is spent between the problem detection and resolution, where most situations have a real impact on service availability and user experience.

The DSP needs to evolve to be more proactive, focussing on early problem resolution and using advanced analytics to identify hidden patterns quickly. In this transformation, operations

processes evolve from a manual definition of rules for event correlation, based on known patterns, to predictive analytics that can identify the rules and minimise human intervention automatically. Based on long historical data and ML algorithms it is possible to anticipate failure events, for example, when a particular mobile site goes down or when a network resource is reaching the limit capacity causing congestion and affecting subscribers’ QoE.

ML algorithms will also augment the capability to identify problems. As so, it is critical to integrate the maximum number of information sources that can feed the analytics engine in real-time with data from events, performance management, network interventions, user and control plane data, weather conditions among other relevant data sources that are available in each DSP B/OSS ecosystem. Adding more information sources to enhance root cause identification in today reactive-response processes does not scale due to the necessary human intervention in the definition of rules and pattern knowledge.

With a proactive response approach, DSP will operate their networks differently as, for example predictive maintenance activities can be scheduled to anticipate resource failure. Fixed-scheduled activities may be planned more efficiently, avoiding unnecessary truck rolls, minimising resources downtime and service availability.

## Proactive network protection

Future networks are expected to support new business domains (known as verticals), and many of these new application areas will result in new and demanding requirements on safety and security. Cybersecurity became a top priority for governments, organisations and private industries across the world, aiming to provide high levels of resilience against cyber threats. Therefore, novel solutions are required for early detection and mitigation of potential cyber-attacks, especially in the upcoming 5G environments, which will bring a massive number of connected devices. Distributed deny of service (DDoS) and virus spreading pattern attacks on cognitive radio networks are amongst the potential threats for mobile networks and are, therefore, the target for this type of use case. Both of these potential attacks are likely to take on a new critical relevance in 5G networks: mobile subscribers' devices may come under the control of the attackers using bots, or zombies, who establish a botnet to perform DDoS attacks, or even trigger self-propagating viruses to manipulate radio elements as a means of gaining operational advantages in 5G networks.

In detail, this type of use-case, and in particular the botnet detection procedure, can be decoupled into two complementary phases, namely:

1. High-level detection to proactively identify suspected attacking channels, through monitoring and analysing network traffic flows;
2. Low-level fine-granularity detection through deep packet inspection (DPI) to confirm the attacking channels detected in the previous phase.

Monitoring network flows during the first high-level detection phase allows analysing big volumes of data quickly and in near real-time. The massive amount of network traffic makes an in-depth analysis of network packets unfeasible at the first step, and for this reason, the DPI analysis is conducted in a second step only between the peers identified as suspects during the first step.

Once the second detection phase confirms the actual existence of the botnet, the reaction consists of deploying and enforcing a virtualised and personalised honeynet to clone the botnet zombies and emulate their behaviours. As a result, the real attacker (i.e., the botnet's owner) will not be aware that the honeynet has disabled part of the attack actions. In parallel, besides confirming the botnet through DPI analysis, the second phase also triggers AI procedures based on ML techniques, that will fetch the historical data for the confirmed botnet flows and identify traffic patterns. As an outcome, a new and complex diagnosis rule, based on the analysed traffic flows, is produced allowing for the early identification of future cyber attacks without requiring the placement of intrusive DPIs on the data path.

## Intelligent service optimisation

Planning the network capacity and coverage requires a profound knowledge of network parameters for each vendor, network releases and activated features. This is a complex activity for mobile networks that are more vulnerable to changes in their behaviour due to external factors, congestion and mobility issues. The introduction of 5G networks will amplify the challenges in this domain since the number of network parameters will increase, contributing to a more complex and time-consuming activity for DSP engineers that need to implement optimal parameter configuration.

Although in the future this complexity will increase, the current networks technology state already represents a great challenge for operators. The undergoing evolution in networks introduces every year at least one or two network upgrades that require specialised knowledge on how to set up and optimise the network configurations. In the past years, the self-organising networks (SON) solutions promised to speed up the optimisation activities running periodically and identifying the best configuration for network parameters. Based on performance information and network parameters, centralised SON features are able to (re)tune parameters to optimise coverage,

handover and capacity of the mobile services. However, actual centralised SON solutions actuate based on fixed rules, and some of the results are not as good as DSP were expecting.

The future of service optimisation will be an evolution of today's SON solutions leveraged by ML algorithms capable of handling the specific aspects of each cell, based on location, usage type, neighbours, time of the day and other variables. The programmable capabilities of the underlying networks will further accelerate the SON interactions, in real-time and with more visibility. This can be achieved through a centralised SON framework that ties data from the network, OSS and external sources, analyses performance bottlenecks and proactively identifies the most appropriated tactics and actions (e.g. network functions reconfiguration, deployment, replacement ...) over the network infrastructure.

## Customer support optimisation

Customer support is one of the service providers' major problems when it comes to operational expenditures and customers dissatisfaction. Although advanced, highly-customisable and very intuitive/user-friendly diagnosis platforms are already available for performing tests over the network and for enforcing selective corrective actions to mitigate existing problems, there is always a considerable human effort involved in solving customer reported issues. As a consequence, the customer experience is severely impacted, and the operational costs increase significantly.

Nowadays, when a customer reports a symptom, the call-centre operator initiates a diagnosis process to determine the symptom root-cause and mitigate the problem. The operator uses a diagnosis platform and can run a set of tests over the customer network infrastructure to identify the cause of the problem and to enforce a set of corrective actions to mitigate the issue. Nevertheless, the analysis and decision task is a human-based process, relying entirely on the

operator to rapidly select the most appropriate batch of tests to identify the problem, as well as the set of actions required to mitigate or correct it. If the operator is a highly-experienced one, and the symptom reported is simple/well-known, the time required to diagnose and solve the problem is usually low, and the customer impact is also low. However, if the operator is dealing with a complex, non-typical situation, requiring multiple tests over different network segments or equipment for root-cause analysis and for deciding on the appropriated actuation procedures, the time required to solve the customer issue increases exponentially, resulting in higher operational costs, service downtime and ultimately customer dissatisfaction.

The goal of this use-case is to minimise the impact of the customer-reported network symptoms by rapidly and efficiently deciding and recommending the best diagnostic tests and corrective actions to mitigate the reported problems. In detail, three phases are foreseen to achieve a highly dynamic and automated scenario:

1. In the first phase, the most conservative and less intrusive one, the cognitive-based diagnosis platform automatically decides and recommends the set of required tests and corrective actions to mitigate a reported customer problem. Although the decisions and analysis are automated, the tests and actions enforcement on the network is the responsibility of the operator. The main benefit of this scenario is the improvement of the customer experience by minimising the problem resolution time.
2. In a second, intermediate phase, in addition to automating the analysis and decision process, the actuation (tests and actions) on the customer or the access network are also automated. At this phase, the service provider can leverage the full potential of batch routine tests as automated diagnosis and correction procedures. This phase also creates the opportunity for DSP to expose test functionalities to end customers, allowing

them to self-diagnose and self-correct anomalies, thus increasing the customer experience. In this scenario, the call centre operator has a passive role.

3. Finally, in a third phase, the most advanced and desirable one, through continuous monitoring and analysis over the customer premises equipments (CPE) and access network historical data, the cognitive diagnosis platform will be able to predict and mitigate customer related issues before they happen. This phase presents a highly dynamic and automated scenario considerably improving the customer experience and simultaneously diminishing the operational costs.

In summary, the main benefit of this use-case is to improve the customer experience by minimising problems resolution time and, in a more advanced scenario, even predicting and proactively solving the customer problems.

## Autonomous cognitive operations architecture

Network management in 5G networks must handle the unprecedented complexity inherent to the flexibility, pluggability, and hierarchical composition capabilities that 5G enables. It is no longer possible to prescribe what needs to be done per every conceivable system state, as there are too many options and heterogeneous scenarios. Thus, the system's evolution encompasses autonomous network management that utilises ML to understand and predict the network behaviour and proactively steer it towards its desired state. Rather than using an imperative network management approach, a declarative one is being designed in which cognition is used to learn the best tactics to achieve declared goals.

Technically speaking, autonomous management is achieved by enhancing an autonomous

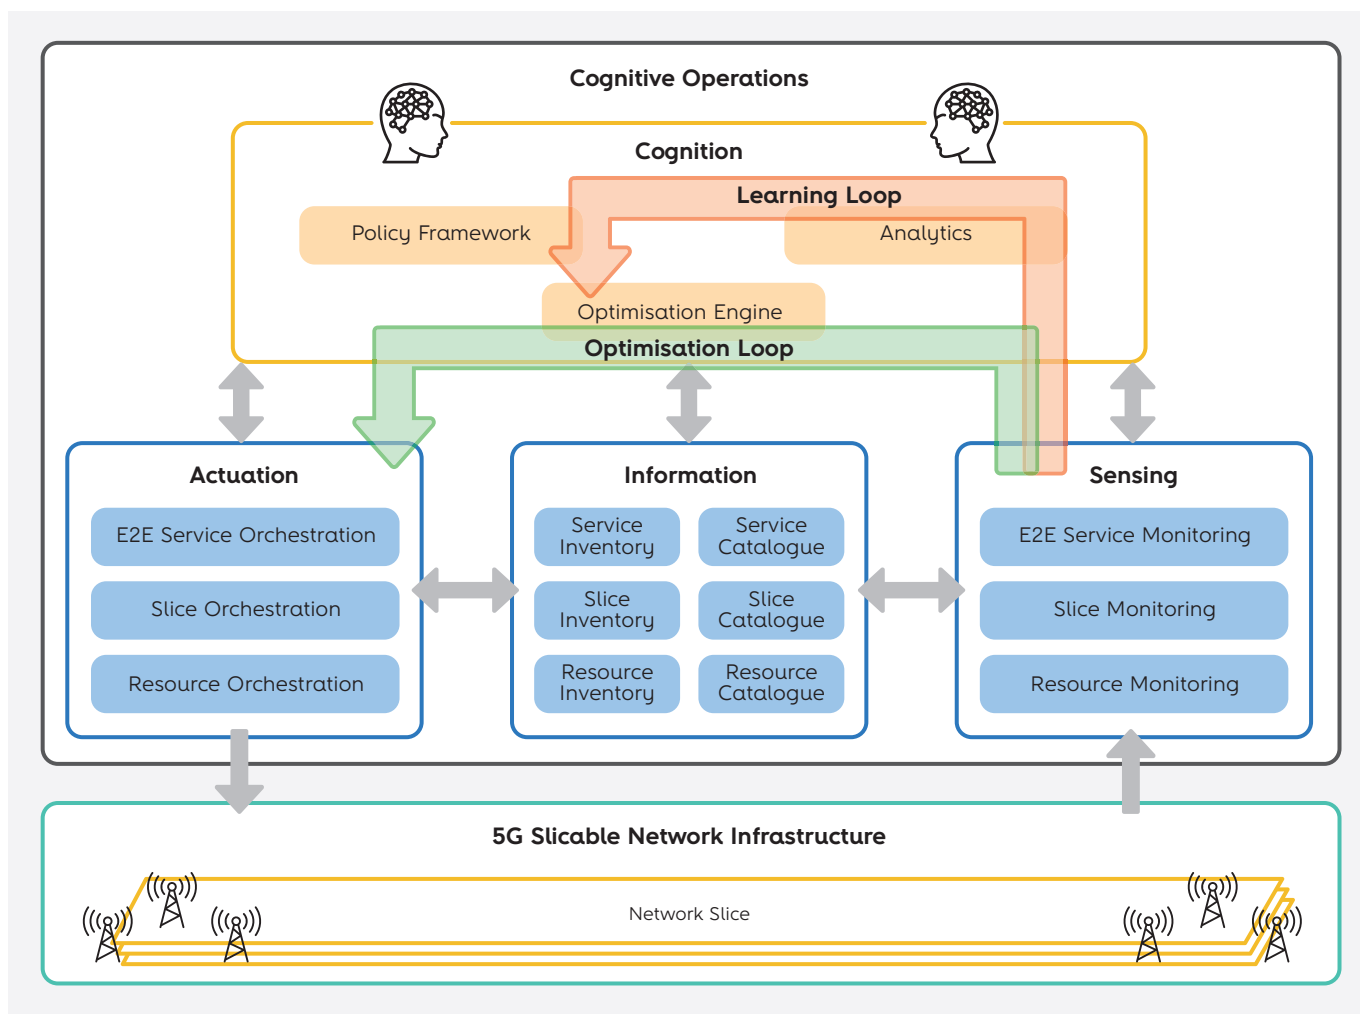
computing control loop with ML capabilities. Monitor-analyse-plan-execute (MAPE) is a common model to describe an autonomous computing loop, where the monitoring component obtains data from the controlled system through sensors, data is then transformed into information through analytics and optimisation, autonomous decisions are converted into concrete execution plans, and finally, the execution component orchestrates the plan via actuators. The MAPE loop is governed by management policies and knowledge, targeting a catalogue-driven management system in which information cannot be statically configured to cover all possible states, but it can be learned and continuously refine the MAPE loop to achieve end-to-end QoE.

In this section, we present a very high-level perspective of an autonomous cognitive network management architecture (illustrated in **Figure 3**).

Within the autonomous cognitive operations architecture, four key frameworks are identified: **information, actuation, sensing and cognition**. For each framework, three "invisible" horizontal layers are defined to abstract the resources, slice and service details, thereby guaranteeing a clear separation of concerns between these logical architecture components.

The **information framework** is responsible for maintaining all the required knowledge for the remaining architecture components. As a logically centralised entity, it avoids replication and incoherent information to be spread in several architecture components. In more detail, it provides a heterogeneous set of catalogues and inventories, and while the catalogues keep the information related to the architecture capabilities and offerings concerning network functions, services and network slices, the inventories keep track of all provisioned instances of end-to-end services, mapping and correlating all the components regarding running functions, network slice instances and service instances.

The **actuation framework** provides a set of coordination functions required to onboard, deploy, configure and maintain services and network slices



**FIGURE 3** – Autonomous cognitive operations architecture & key loops

as the combination of different network functions. The role of the service orchestrator is to coordinate the actuation management processes to fulfil and guarantee continuous delivery of services to individual customers in a multi-tenant environment. Services are implemented through network slices which consist of an isolated set of heterogeneous resources deployed and managed at runtime with the use of the slice orchestration functionality. This functionality coordinates one or more technology-independent domains by interacting with their respective management systems. The coordinated integration of distinct domains enables building isolated and unified slices that can be managed during their complete lifecycle by the slice orchestrator. To achieve this, the slice orchestrator component will use aggregation capabilities to

combine information and operations at slice-level. Finally, the resource orchestration is the lowest orchestration level within the actuation framework and provides the building blocks to interact with physical equipment and functions, virtual functions and SDN components for lifecycle management.

The **sensing framework** encompasses the collection of both structured and unstructured data from various sources and at multiple layers in a flexible way, either to support cognitive-oriented management procedures or to expose desired KPI towards vertical customers. More specifically, the resource monitoring is responsible for collecting information from the resources composing the data plane, thus enabling the monitoring functionalities of the upper levels.

The data collected, along with topological and aggregated traffic information, is persisted in raw-data storage and serves as a foundation for higher-level monitoring capabilities (slice and service) as well as main input for the cognitive/ML loops found at the cognition framework. The slice monitoring is responsible for collecting and generating data related to slice level performance indicators, for instance, end-to-end bandwidth and latency between deployed VNF instances among other metrics. To do so, it leverages on the information provided by the lower level monitoring entity. These performance indicators will then feed the cognition framework for further deriving new predictions and policies, as well as to trigger slice actuation (through the optimisation engine) to maintain desired SLA levels. Finally, the service monitoring adds high-level information regarding specific service instance performance to the already collected slice level and resource level information and exposes it towards cognition loops or vertical users for KPI reporting.

Finally, the **cognition framework** aims to ensure smart management operations, such as fault, performance and security management. It encompasses the ML operations (prediction, classification, clustering, ...) to identify faulty network components, detect anomalies, such as network cyber-attack and malware, predict SLA violations, etc. [3]

The **analytics** component provides for the AI techniques of the architecture. It is composed of a set of AI processes that are capable of making autonomous decisions that impact the network, slice and service status, from planning to deployment and maintenance. The analytics functionality should also be able to provide standard methodologies to apply several analysis processes over the gathered data. This approach allows for a generic module that is able to build predictive models, apply ML techniques or perform any other statistical treatment.

Policies represent rules and restrictions that govern the behaviour of a system. The **policy framework** component guides the decision-making process regarding what the system

should or should not do. Most decisions that are taken across the whole system are determined by rules that implement policies at various levels, ranging from decisions at the resource level, taken in real-time from locally generated information to business decisions based on information aggregated in time and space from various sources. Different policy paradigms may be considered, e.g. imperative and/or declarative, to be mapped in distinct management layers, i.e. service, slices and resources. Although related, they are distinct in nature, and a change on a policy will have effects that ripple across the whole system and eventually lead to changes in policies that depend on it.

Finally, the **optimisation engine** is intended to provide optimisation capabilities at the resource, slice and service levels to maintain the desired SLA guarantees. To achieve this, the policy framework defines the overall optimisation goals regarding declarative policies, i.e., a set of operative guidelines that lead the optimisation procedure in deciding the best actions for SLA guarantees. Under the policies, optimisation engine consumes monitoring events from the sensing framework and delivers actuation actions (e.g. deploy VNF, reconfigure VNF) towards the actuation framework to enforce the optimisation decisions.

One of the most important concepts of this type of system architecture is the ability to apply cognitive management by enhancing an autonomous, policy-based, closed loop with ML capabilities. Two main policy-based autonomous closed loops are foreseen within the system architecture:

- 1. Cognitive learning loop:** this loop includes the intelligent, e.g. ML procedures which, as a result, can update or create new policies for the optimisation closed loops at different levels (described next). It involves components and interactions from the **sensing** and the **cognitive frameworks**;
- 2. Cognitive optimisation loop:** this policy-based closed loop involves the **sensing, cognition** and **actuation frameworks** and is used mostly for real-time sensing and actuation on the

network. It is based on policies and models fed by the cognition framework.

Altice Labs is evolving their OSS frameworks towards this reference architecture. This transformation is being initiated with the developments of two major modules of NOSSIS, named Assurance and Fulfillment, where a new architecture for both sensing and actuation layers is currently being developed.

## Final remarks

The ongoing network transformation and the evolution to digital services industry bring new significant challenges to DSP. To address these challenges, DSP need to embrace a transformation in their operations so they can reach unparalleled levels of efficiency and agility required by the digital services industry while addressing the requirements to manage a new software-defined network architecture that is also a tool for enhanced agility and efficiency.

Transforming existing operations means the adoption of a new operations platform able to provide new levels of automation and intelligence. The use of operational data in real-time, with capabilities to correlate, aggregate and close the loop with actuation processes orchestrated







by a cognitive layer, able to learn and handle more complex aspects of the new virtualised and programmable network infrastructure, is being defined as the evolution in OSS that will handle this future challenges. In the past years, data processing technologies have registered a significant evolution in the capacity to process and store data in real-time. Now and for the near future, the larger investments are in AI technologies that will lead to the cognitive features to provide these superior capabilities. AI must be associated with ML algorithms to learn and build the rules that need to be deployed in underlying OSS processes, closing the sensing-cognitive-actuation control/closed loops.

Altice Labs intends to be a technology partner to support the DSP transformation towards autonomous cognitive enabled operations, providing the market with a new generation of operations platforms. Several projects are being conducted internally in Altice group and externally with the participation in European projects such as H2020 SelfNet [4] and SliceNet [5], in which cognitive service operations are being prototyped. Moreover, the undergoing transformation towards Altice Labs' NOSSIS One architecture will foster this transition, which will make carrying out service and network design, service fulfilment and service assurance much easier across the whole hybrid networks.

## References

- [1] TM Forum, "OSS of the Future," 2017.
- [2] A. Okeleke, "Using AI to Improve CSP Network Operations: Use Cases." Ovum, 2018.
- [3] ETSI, "Experiential Networked Intelligence," ETSI, 2018. [Online]. Available: <https://www.etsi.org/technologies-clusters/technologies/experiential-networked-intelligence>.
- [4] 5GPPP, "SelfNet," 5GPPP, 2018. [Online]. Available: <https://selfnet-5g.eu/>.
- [5] 5GPPP, "SliceNet," 5GPPP, 2017. [Online]. Available: <https://slicenet.eu/>.

46 What strategy for information security?





# 04

## What strategy for information security?

**Paula Cravo, Altice Labs**

[pcravo@alticelabs.com](mailto:pcravo@alticelabs.com)

**Nuno Alexandre Seixas, Altice Labs**

[nuno-a-seixas@alticelabs.com](mailto:nuno-a-seixas@alticelabs.com)

In the aftermath of the EU GDPR, technology companies are finally looking at how this new regulation can be made effective in their own ways of working. Data-driven DSP are among the most concerned ones for information security since their business models are, largely, mostly on data processing. With this article, we intend to show four strategic themes that those companies need to address, namely, Security by Design, Security in Operations, Secure Third Parties and Security Culture.

---

### Keywords

Security by design; Security in operations; Secure third parties; Security culture

## Introduction

According to Symantec's 2018 Internet Security Threat Report, 2017 witnessed an increase both in the volume and in the threat landscape, with attackers using new avenues of attack and covering their tracks in the process. From WannaCry and Petya/NotPetya to coin miners, threats came from new and unexpected sources [1].

The year of 2018 also brought new and more demanding security and privacy regulations, namely, Europe's General Data Protection Regulation (GDPR), with fines that do impose for companies and organisations the need to ensure compliance.

With this article, we intend to show four strategic themes that those companies and organisations need to address to comply with more demanding security standards and regulations, namely, Security by Design, Security in Operations, Secure Third Parties and Security Culture. More specifically, we shall discuss the possible impact from this context change into Digital Service Providers (DSP).

## Security context update

The impact of a cyberattack on companies and enterprises range from economical to reputational damages. Economic costs are simple to account for, and include significant financial loss arising from:

- theft of corporate information;
- theft of financial information;
- theft of money or disruption to trading;
- loss of business or contract;
- costs associated with repairing affected systems, networks and devices.

On the other end, reputational damages are more difficult to account for but may cause more lasting damage as trust is an essential element of the customer relationship, and cyberattacks may erode this relation. This deterioration can potentially lead to loss of customers and sales, as well as a reduction in profits.

In 2017 some cyberattacks hit the headlines all over the world: WannaCry, Petya/NotPetya, Bad Rabbit and so on. More recently, rampant data breaches and personal data exposure lead to the shutting down of Google Plus services and a major investigation on Facebook, among other worrying events around the world.

As so, it's no wonder that authorities all over the world are concerned about the way enterprises are to deal with security issues and promulgate regulations and laws, like EU GDPR, EU ePrivacy, the US Privacy Act, or the Health Insurance Portability and Accountability Act (HIPAA), initiatives that aim to protect individuals and enterprises in a consistent manner.

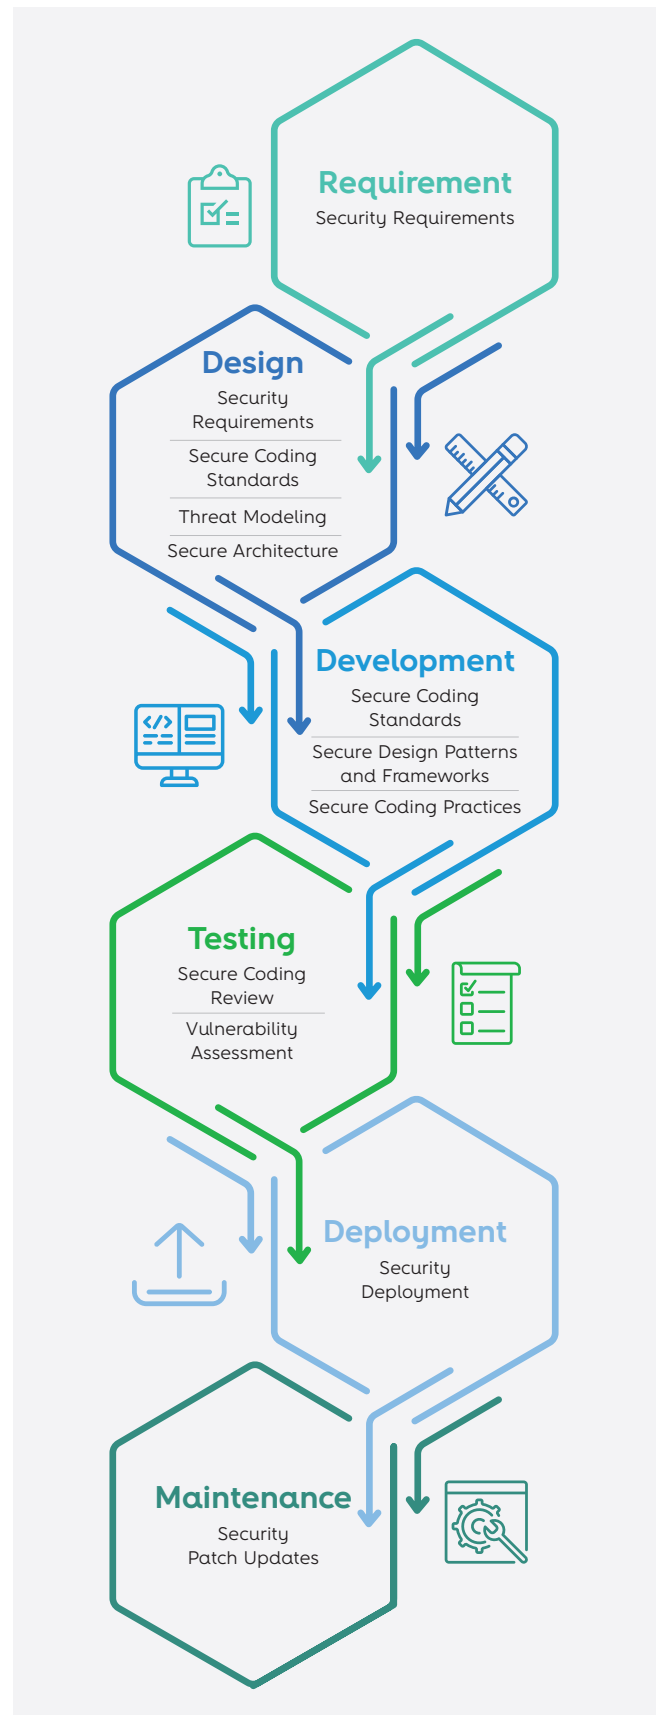
On the organisations' side, these digital security requirements demand a clear definition of an enterprise security strategy and a re-tailoring of existing policies. This demand applies even more to DSP due to the type of environment they operate in, developing digital products and services that are expected to attend a great number of customers from a wide range of distinct locations. Besides, DSP are the main drive to the acceptance of products and services that are supported by new or disruptive technologies: mobile application, services in the cloud, Internet of everything, artificial intelligence (AI) and machine learning (ML) supported services, network virtualisation, and so on, conducting to the customer's behaviour change, from a traditional client towards an always-on, self-serviced, tailored-product one.

In the next sections, we will discuss briefly some aspects that DSP must take into consideration when devising their security strategy.

# Security by design

When dealing with security-related issues, DSP need to take into account the way they build and deploy those services, therefore, incorporating security by design principles. Currently, while building a digital service, the architecture team needs to consider the functionality, the operational aspects, but also, the security ones. Specifically, from a security perspective, it is now important that this domain must be brought into the development lifecycle right from the start. The digital services need to build in the capacity of sustaining strong security events without disrupting the service or the integrity of the data being processed. Following the Open Web Application Security Project (OWASP), digital services and their applications need to be able to maintain three basic security principles: confidentiality of information; integrity of data and finally, service availability [2]. To assure these principles, software development teams need to evaluate their own making not just on a functional basis but also on a security one. Therefore, for each change made into the service, the team needs to follow a distrust approach, rather than a trust everyone approach that used to be taken.

This distrust approach means that the software development team needs to analyse their software to find out what kind of mechanisms are in place that protect data confidentiality; what kind of mechanisms can be used to attack the system and finally, if there is a need for this function to be available anytime. Such approach is different from the traditional one, that focus solely on building new functionalities and then adding some security concerns at the end of the development cycle. With this approach, security is considered in every development phase. In fact, modern development methodologies, like Agile [3] and DevOps [4], are not compatible with the traditional approach and demand security to be considered at every stage, as shown in **Figure 1**.



**FIGURE 1** – Secure software development process [5]



FIGURE 2 – OWASP security by design principles

To implement this new approach, security by design must also be a concern for software architects and developers and not just for security specialists. These professionals need to follow some common principles, as presented by OWASP and displayed in **Figure 2**.

These principles altogether with the best development practices can help minimise security risks within systems. Approaches and techniques such as threat analysis, static and dynamic source code scanning and penetration testing can efficiently help reduce the number of possible vulnerabilities.

## Security in operations

Following the idea of implementing security through distrust explained in the previous section, in operations this can also be a starting principle for security. When we talk about security in operations, this usually means to address behaviours, habits and data that can pose a threat to the security of any given system.

To analyse and implement security in operations, we should look at five different aspects:

- **Processed data identification:** it's essential that the organisation explicitly knows the type and quality of data that it has under its guard. Only after identifying all the data being processed and produced can the organisation start the definition and implementation of security mechanisms. In this data can be included everything, but to make it manageable, the organisation should focus on the ones that are more sensitive, like client's personal data, intellectual property, financial information, client's data and employee data.
- **Possible threats identification:** based on the data identified before, the organisation needs to analyse the major threats for each one, to then prepare itself for avoiding and addressing them. While analysing these threats, the organisation needs to look for internal threats, in their systems and teams,

but also for threats that might operate from outside, both illegally and legally.

- **Vulnerabilities and security holes:** it's also important to identify the current status for existent safeguards and determine what the weaknesses that might allow for intended use are. For each one of these vulnerabilities, there should be a risk analysis on the impact and probability of using this vulnerability. This analysis will help determine the vulnerabilities that are more important to be addressed and therefore, define priorities for an action plan.
- **Define and implement a plan to eliminate threats and mitigate risks.** Once the vulnerabilities and risks are identified, it's time to take action. A set of countermeasures should be defined and implemented so that the organisation can increase the confidence in their own security. The countermeasures can and should address all the different people in the organisations, making the action much more effective.
- **Cross organisation security awareness:** for all these aspects to be put to work, it's important that security is not just a matter of security teams, but a matter of concern for the whole organisation, from practitioners to senior management. This involvement is of paramount importance since they can grant responsibilities and resources to make security happen. Without this awareness and support, these initiatives will lose impact, and the vulnerabilities and risks will not be adequately addressed.

## Secure third parties

In the previous sections, we talked about how an organisation can implement mechanisms that are meant to address security at various levels. Nonetheless, once it integrates different third-party components or companies, its mechanisms need to take that fact into account or might risk

being completely void. That's why it's important to address some of the major concerns when dealing with third-parties: data transfer, workflow management and dependencies management.

The first aspect, data transfer, is related to how third-parties deal with all the data they receive from the original organisation, in the context of a project. According to EU GDPR, when, in the context of a service, an organisation transfers personal data to another organisation, it does not stop being responsible for processing that data. Also, if the organisation intends to use the results from that data processing in its services, it needs to account for all the protection mechanisms that are part of GDPR [6]. Because of these legal aspects, it's important that all third-parties have explicit mechanisms for data processing and protection.

The second aspect, workflow management comprises the way the external entity is executing its work and if that workflow is secure and does not compromise the security of the overall system. In this workflow, all the controls mentioned above need to be executed. If they are not executed, there is a real danger that the overall security strategy can be at stakes.

The third aspect, dependencies management, relates to how the project team deals with dependencies from third party suppliers. First of all, those dependencies should be explicit, because if they are implicit, they cannot be managed or accounted for security. Second, every dependency needs to be evaluated regarding security, and specifically, how it relates to the security strategy and security controls presented before.

In a similar way to the internal controls, the outside third-parties need to have several different risk mitigating mechanisms. The security risk analysis begins when selecting the supplier. All the prospective suppliers should be analysed concerning security controls and security assurance. This analysis is especially critical when these third-parties are going to have access to customer data, and therefore, they should agree

to protect that data in a correspondent way to how the organisation is doing. Also, they need to agree to make explicit evidence of those controls to assure at any given time they are addressing those matters.

The way these mechanisms are implemented depends mainly on the type of relationship with the third party but also on the kind of data and information being processed and accessed.

## Security culture

Wrapping up the security aspects we discussed in the previous sections, a security culture makes the security strategy real and effective. It exists in an organisation when employees carry out a default behaviour when facing cybersecurity-related cases where no concrete guidelines exist. Because good security practices are crucial to the development of consumer trust and confidence, the organisation's security culture is foundational for the success of its information security program. As companies face an accretion on the number of continually evolving cyber threats, greater effectiveness will be achieved if cybersecurity values are embedded in the organisation's culture, and the board as well the operational staff are committed to this goal.

Security culture is primarily for humans, not for computers, and one of the features of a sustainable security culture is that it must be deliberate and disruptive to foster change and better security. However, it must also be engaging, challenging and fun, driving people to participate in it, and also rewarding, i.e., for people to invest time and effort they need to understand what they will get in return [7]. A culture of security must involve everyone, starting at the top, with CEO's and boards engaged in support of the security program.

Awareness campaigns help to improve employees security competencies and consequently of all the organisation, minimising security risks

and incidents. Furthermore, this care must be extensible to every employee, contractor and third-party, ensuring that everyone understands their responsibilities to protect the organisation's data and assets, and effectively manage security risks. General awareness campaigns must be complemented with application security knowledge for developers and testers within the organisation, to result in the delivery of secure products and services.

Even the company's customers must be targeted on the organisation's security culture, and be part of active awareness programs to be more conscious of the security risks and know how to protect themselves.

The ultimate goal of a good security culture is to change behaviours, which will be a major determinant as to whether a company will survive in today's threat environment. Apathy, silos and self-interest are enemies of creating a security culture, and providing constant reminders of the importance of security reinforces the significance security has for the entire organisation. Of all technological measures that could be taken to ensure a minimum level of security compliance to the organisation and the company's products and services, none will be more effective and present a higher return of investment than training and cultural awareness [8]. When a security culture is sustainable, it transforms security from a one-time event into a lifecycle that generates security returns forever.

## Conclusions

New technologies like the Internet of things, cloud, software-defined networking (SDN), network function virtualisation (NFV), and others will have a tremendous impact on DSP's core business model and their digital transformation. For DSP, this means they need to think how to proactively secure their assets, from network architecture to data and its processing.

As DSP realise the importance of a security strategy for their business, they are assigning an increased budget to security initiatives. Gartner forecasts that worldwide spending on information security products and services will reach more than 114 billion USD in 2018, and in 2019, the market is forecast to grow to 124 billion USD [9].

Some of the top drivers for DSP security investment are security risks, business needs and industry changes. However, privacy concerns are also becoming a key factor in this domain.

Cyber threats landscape for 2019 is not reassuring. Besides ransomware, cryptojacking and phishing, we should expect AI-powered attacks and even some probable nation-state attacks that will affect the organisational status quo and even the individual perception of security and privacy [10] [11].

DSP, delivering digital services and products, are set to be one of the most desired targets for any or all of these threats. Therefore, by its own nature, DSP need to be at the frontline of engaging and addressing security concerns.





One of the ways DSP may address it is by following the strategy that we described previously. Ranging from building the right amount of security into their products and services (security by design) to the overall operational deployment (secure operations) and without forgetting the importance of correctly addressing and involve third parties into this strategy (secure third parties). Finally, and probably the most important one, DSP need to promote a cyber-security culture where everyone in the organisation needs to be aware and willing to act in order to maintain a secure environment. By having a secure culture, the organisation needs to ingrain a behavioural change into everyone involved (internal teams, third parties and clients), so that the previous measures are effective and not betrayed by some individual misbehaviour.

## References

- [1] Symantec, "Internet Security Threat Report - Volume 23," 2018.
- [2] "Security by Design Principles," *The OWASP Foundation*, 2016. [Online]. Available: [https://www.owasp.org/index.php/Security\\_by\\_Design\\_Principles](https://www.owasp.org/index.php/Security_by_Design_Principles).
- [3] J. Rasmusson, "Agile in a Nutshell," 2018. [Online]. Available: <http://www.agilenutshell.com/>.
- [4] G. Kim, P. Debois, J. Willis, and J. Humble, *The DevOps Handbook: How to Create World-Class Agility, Reliability, and Security in Technology Organizations*. IT Revolution Press, 2016.
- [5] "Certified Application Security Engineer (CASE)," *EC-Council*, 2018. [Online]. Available: <https://www.eccouncil.org/programs/certified-application-security-engineer-case/>.
- [6] "General Data Protection Regulation," *The European Parliament and the Council of the European Union*, 2016. [Online]. Available: <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32016R0679&qid=1542977568879&from=EN>.
- [7] C. Romeo, "6 ways to develop a security culture from top to bottom," *TechBeacon*, 2016. [Online]. Available: <https://techbeacon.com/6-ways-develop-security-culture-top-bottom>.
- [8] G. Fredriksen, "Fostering A Culture Of Security," *Forbes*, 2018. [Online]. Available: <https://www.forbes.com/sites/forbestechcouncil/2018/06/05/fostering-a-culture-of-security/#4e6652e5461a>.
- [9] S. Moore and E. Keen, "Gartner Forecasts Worldwide Information Security Spending to Exceed \$124 Billion in 2019," *Gartner*, 2018. [Online]. Available: <https://www.gartner.com/en/newsroom/press-releases/2018-08-15-gartner-forecasts-worldwide-information-security-spending-to-exceed-124-billion-in-2019>.
- [10] "9 cyber security predictions for 2019," *CSO from IDG Communications, Inc*, 2018. [Online]. Available: <https://www.csoonline.com/article/3322221/security/9-cyber-security-predictions-for-2019.html>.
- [11] P. Kantaria, "Cybersecurity trends for 2019: Every company worldwide 'will be hit by a cyberattack,'" *Power Technology*, 2018. [Online]. Available: <https://www.power-technology.com/digital-disruption/cybersecurity-trends-for-2019/>.







# 05

## An approach to data monetization

**Ricardo Jorge Ferreira, Altice Labs**

[ricardo-j-ferreira@alticelabs.com](mailto:ricardo-j-ferreira@alticelabs.com)

**Filipe Cabral Pinto, Altice Labs**

[filipe-c-pinto@alticelabs.com](mailto:filipe-c-pinto@alticelabs.com)

**Mário Moreira, Altice Labs**

[moreira@alticelabs.com](mailto:moreira@alticelabs.com)

**Penousal Machado, Faculty of Sciences and  
Technology - University of Coimbra**

[machado@dei.uc.pt](mailto:machado@dei.uc.pt)

Organisations' data, when enriched and/or correlated with third party specific data, and presented (and visualised) in a way humans naturally understand and rapidly grasp, can significantly improve the degree of information and knowledge retrieved from data, thus opening new horizons for decision makers and promoters, placing their outlook novel and innovative value into added business opportunities that will drive the market forward and will result in a market advantage.

---

### Keywords

Data; Data monetization; Data visualisation; DataPlaxe

## Data nowadays

The act of good management and decision-making requires skills, knowledge and information built upon up-to-date data. This is true for people (self-management), organisations and machines alike.

Until recently, collecting, storing and processing large amounts of data to support improved decision-making required precious resources. Furthermore, rarely the data stored had a direct financial return thus burdening data ownership. Currently, with the advent of big data systems, this paradigm has changed, and data possession has become a precious asset with direct monetization capability. The availability of cloud platforms, the proliferation of real-time Internet of things (IoT) devices and always-on data connections, combined with the emergence of advanced data mining and data analytics algorithms and automated machine learning methods, has leveraged enough burdens for big data to take-off. This is to say, the capability of cost-effectively collect, store, process and correlate data from several sources to, in the first place, produce real-time information and KPI that are of the foremost interest (for the market, as well as for data producers, data owners and data scientists) and, in the second place, allow to monetize their specific data and generate added value for all actors involved, has become a reality.

Organisations' data, when enriched and/or correlated with third party specific data, and presented in a way humans naturally understand and rapidly grasp, can significantly improve the degree of information and knowledge retrieved from data, thus opening new horizons for decision makers and promoters, placing their outlook novel and innovative value into added business opportunities that will drive the market forward and will result in a market advantage. However, and although the business space unlocked by this data usage paradigm and big data technologies seems to exhibit a high potential, *a priori*, some challenges are identified to candidate service providers in this space. One of them is to effectively

manage the overall supporting computational environment (that makes available online one-stop-shop processes, gives access to online data and analytics features capable of promoting data monetization), a complex task that involves:

- Setting up, operating and maintaining an integrated framework with the capability of effectively and efficiently perform data ingestion;
- Support an IoT hub to mediate between IoT devices and the application components;
- Support big data storage and processing capabilities;
- Exhibit a rich analytics with advanced visualisation and reporting tools.

## Data challenges in the era of IoT and big data

Since its inception, the Internet has enabled the exchange of data between systems and devices connected in a global network, enhancing the offer of the most various services based on online information (with unique characteristics and specific service needs). Indeed, the single transmission of a device is not a cause for concern, but when faced with millions of sensors and actuators exchanging information, some questions are already raised.

On one hand, it is necessary to make the ingestion of the data in a data management platform giving support to different protocols and architectures. Regarding the specificities of the IoT world, the above mentioned data management platforms must support different modes of communication, both synchronous and asynchronous, and be prepared to deal with the IoT service requirements (such as registering devices, announcing and discovering new

connected objects, or the subscription of IoT events by different entities).

On the other hand, due to its adoption by different areas of activity, IoT is becoming a strong content provider of big data systems since it allows to collect a tremendous amount of data that, when crossed with other information sources, may produce high-value knowledge that can be used in smarter, goal-oriented services. For this to happen, it is necessary to have the above-mentioned data management platforms capable of processing complex events and analyse batch and streaming data coming from heterogeneous systems. The insights achieved after data treating can be made available for free (in an open data model) or in a paid information model, promoting the creation of new businesses in the data arena.

## The importance of visualisation tools

In every industry, there is a need to get full value from the massive amounts of information gathered within each organisation. In big companies, data is collected faster than analysed, making many organisations look for better ways to obtain value from their data and compete in the marketplace [1]: Statistics, database technologies and machine learning techniques are essential to collect, process and organise large data sets, but visualisation tools make it possible to interpret the data further using the most powerful of human senses: vision [2].

Visualisation – the graphic representation of data – aims to provide visual means to extract useful information from data, overcoming human and computational limits. Card et al. give the most common definition of information visualisation – “*The use of computer-supported, interactive, visual representations of abstract data to amplify cognition*” [3]. Therefore, visualisation is perceived as a means of cognitive augmentation, making visible the information and relations in data. On the other hand, visual analytics combine advanced data analysis algorithms with interactive visualisations to provide means for an

effective understanding, reasoning and decision-making from large and complex data sets [4]. Ultimately, information visualisation, in particular, visual analytics, aim at “*detecting the expected and discover unexpected*” [4], leading to the discovery of new knowledge, supporting decision-making, engaging the user, and raising awareness.

Information visualisation techniques and methods are a powerful tool when it comes to making sense of data and decision-making [5]. Nowadays, visualisation goes beyond traditional static forms of representation (presented in animated, interactive or real-time forms) [6]. Modern visual analytics systems should effectively and efficiently handle the following aspects: (i) *real-time interaction*, i.e. support the interaction with big datasets, maintaining the system response in the range of a few milliseconds; (ii) *on-the-fly processing*, i.e. support of on-the-fly visualisations over large and dynamic sets of volatile raw (i.e., not pre-processed) data; (iii) *visual scalability*, i.e. provide effective data generalisation mechanisms to address problems related to visual information overplotting; and (iv) *user assistance and personalisation*, i.e. enhance user comprehension by allowing the creation of different user-defined exploration scenarios [7].

Nowadays, a small percentage of companies have storage systems which can adequately accommodate the enormous amount of data generated by IoT. As such, the current challenges of data management research communities and large web enterprises involve the development of consistent and scalable systems, capable of managing enormous amounts of data for a large number of users prioritising data based on needs and value [8, 9]. This exponential growth of data is fuelling the business intelligence and analytics to unveil new opportunities for innovation and enhancement of all techniques from data acquisition, knowledge extraction and dissemination of information [10]. Data analysts involved in the development of these techniques usually interact through a shared interface, such as cloud-based ones [11]. Cloud-based business intelligence and analytics platforms allow analysts

to share data, generally for a price or to monetize it by exchanging it for other benefits. Therefore, data monetization can be defined as the conversion of intangible value data into real value [11].

To provide value, and thereby improve its monetization, data must be visualised, communicated, and shared with stakeholders efficiently. For that matter, the new scientific discipline of data visualisation addresses an effective knowledge representation [12]. According to Chen and Zhang, both aesthetic form and functionality are necessary to convey the knowledge to the broadest possible audience and therefore increase the data monetization [13].

That said, visualisation tools provide means to explore, analyse and communicate data in a dynamic environment, transforming large amounts of data into useful and valuable information [14]. Its real purpose is to induce insights, which can be in the form of the discovery of unforeseen patterns, support decision-making, or communicate knowledge visually. However, the problem of outdated visualisation techniques is also propagated to the majority of existing visualisation tools. In particular, the prefabricated visualisation tools, which are bound to a limited set of functionalities and graphical options, thus limiting the possibilities to take the full advantage of domain-specific data reducing its real value. This

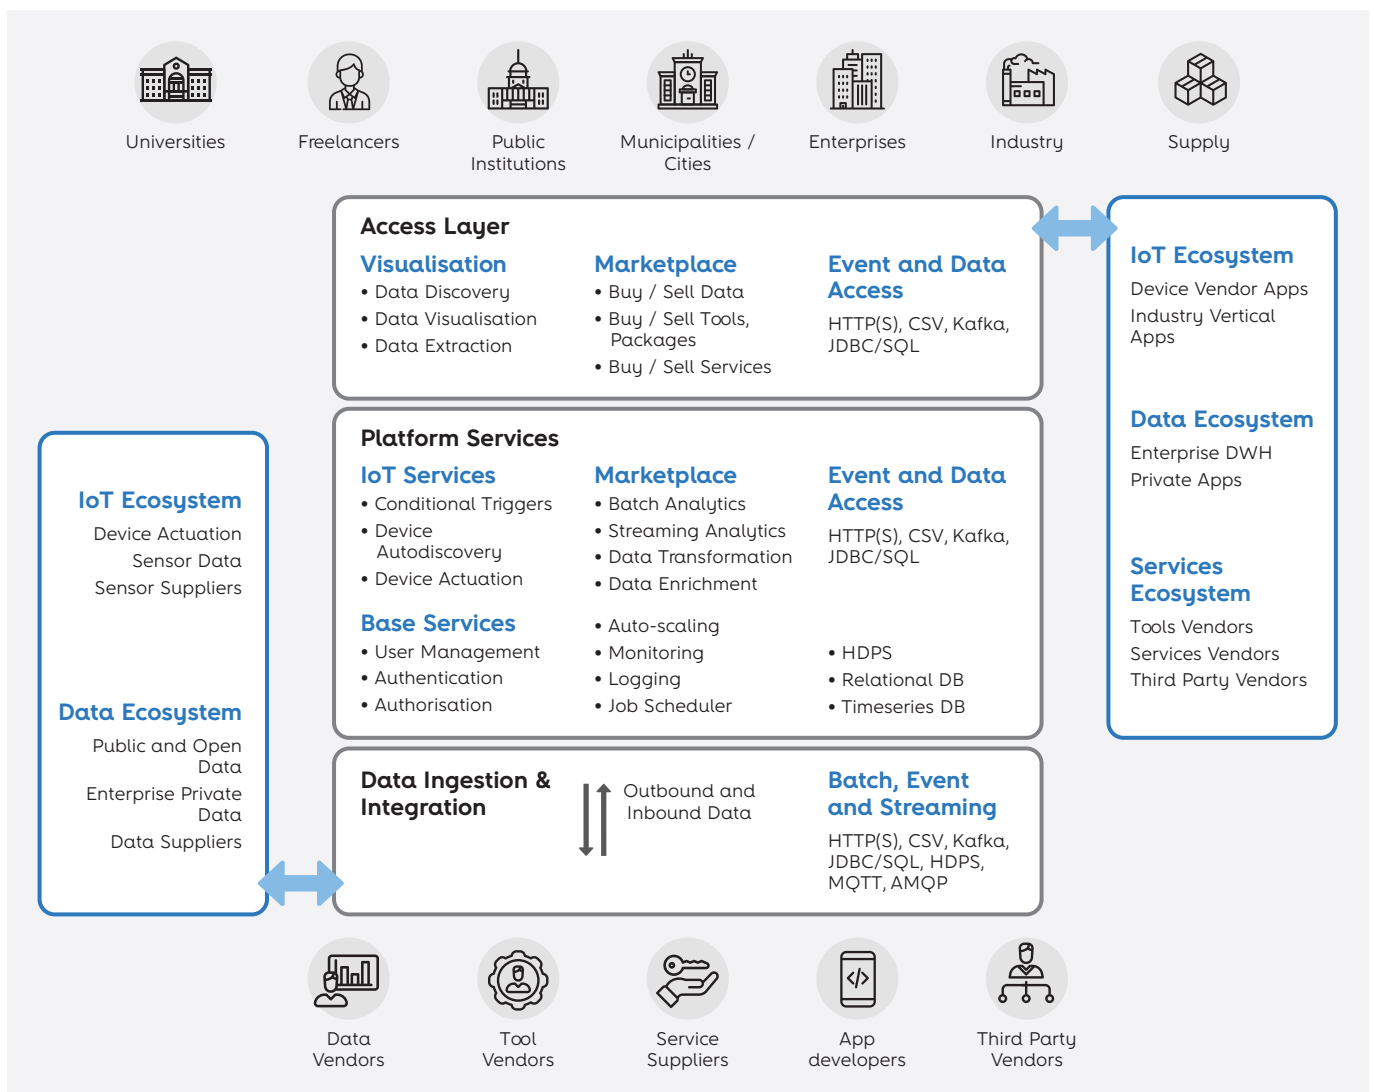


FIGURE 1 - DataPlaxe functional architecture

increases the need for the development of new visualisation tools based on interaction techniques and visualisation models appropriated to different datasets and tasks, improving information retrieved from data and its monetization.

## Altice Labs' approach – DataPlaxe

With all of the above in mind, Altice Labs is developing a platform with the purpose to not only collect, orchestrate, process, transact and perform analytical calculations over large dataset volumes, but also to offer suitable and user-friendly visualisation and reporting tools – The DataPlaxe platform.

DataPlaxe, as illustrated in **Figure 1** (see previous page), is expected to enable the ingestion, formatting, correlation, enrichment and storage of high dataset volumes from several heterogeneous sources, permitting access to both raw and processed data and computed KPI, either in open or restricted access mode, thus supporting diverse business models and use cases. Furthermore, DataPlaxe should exhibit the capabilities of an IoT data hub, supporting the direct connection to and from IoT devices and gateways, in line with the smart cities and industry 4.0 requirements.

The DataPlaxe capabilities are thought to be made available as services in a cloud SaaS

subscription model. Following the one-stop-shop concept, the services offered to the potential customers may range from a data lake for large dataset volumes storage, to data collection, ingestion and processing to advanced analytical and machine learning execution algorithms.

Summing up, the DataPlaxe services ecosystem will allow to create and leverage a dynamic business framework where individuals, companies and institutions may not only store and process their data but may also monetize it (since this platform intends to help unlock the potential value from the data generated in our ever-growing digital society).

## Conclusion

In this article we looked at the unquestionable value of data, when associated with big data and analytics tools, and presented some of the requirements and challenges of building and operating an effective big data ecosystem. It is now clear that a multitude of factors must be taken into consideration when deciding upon how to monetize your data with the resources you have at your disposal.

Furthermore, considering the challenging context of the big data business space, the article points out the DataPlaxe platform, an Altice Labs' approach, as an alternative to building a big data ecosystem, thus presenting an additional path to be considered in the decision process.

## References

- [1] LaValle, S., Lesser, E., Shockley, R., Hopkins, M. S., & Kruschwitz, N. (2011). Big data, analytics and the path from insights to value. *MIT sloan management review*, 52(2), 22.
- [2] McAfee, A., Brynjolfsson, E., Davenport, T. H., Patil, D. J., & Barton, D. (2012). Big data: the management revolution. *Harvard business review*, 90(10), 60-68.

- [3] Card, Stuart K., Jock D. Mackinlay, and Ben Shneiderman (1999). *Readings in Information Visualization: Using Vision to Think*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.
- [4] Keim, D., Andrienko, G., Fekete, J. D., Görg, C., Kohlhammer, J., & Melançon, G. (2008). Visual analytics: Definition, process, and challenges. In *Information visualization* (pp. 154-175). Springer, Berlin, Heidelberg.
- [5] Fekete, J. D., Van Wijk, J. J., Stasko, J. T., & North, C. (2008). The value of information visualization. In *Information visualization* (pp. 1-18). Springer, Berlin, Heidelberg.
- [6] Cottam, J. A., Lumsdaine, A., & Weaver, C. (2012, October). Watch this: A taxonomy for dynamic data visualization. In *Visual Analytics Science and Technology (VAST), 2012 IEEE Conference on* (pp. 193-202). IEEE.
- [7] Bikakis, N. (2018). *Big Data Visualization Tools*. *arXiv preprint arXiv:1801.08336*.
- [8] Lee, I., & Lee, K. (2015). The Internet of Things (IoT): Applications, investments, and challenges for enterprises. *Business Horizons*, 58(4), 431-440.
- [9] Agrawal, D., El Abbadi, A., Antony, S., & Das, S. (2010, March). Data management challenges in cloud computing infrastructures. In *International Workshop on Databases in Networked Information Systems* (pp. 1-10). Springer, Berlin, Heidelberg.
- [10] Chen, H., Chiang, R. H., & Storey, V. C. (2012). Business intelligence and analytics: from big data to big impact. *MIS quarterly*, 1165-1188.
- [11] Najjar, M. S., & Kettinger, W. J. (2013). Data Monetization: Lessons from a Retailer's Journey. *MIS Quarterly Executive*, 12(4).
- [12] Nasser, T. and R.S. Tariq. 2015. Big Data Challenges. *Journal of Computer Engineering & Information Technology* 4(3):1-10.
- [13] Chen, C. P., & Zhang, C. Y. (2014). Data-intensive applications, challenges, techniques and technologies: A survey on Big Data. *Information Sciences*, 275, 314-347.
- [14] Shneiderman, B. (2003). The eyes have it: A task by data type taxonomy for information visualizations. In *The Craft of Information Visualization* (pp. 364-371).







# 06

## Virtual Assistants, the future digital interface

**Jorge Miguel Sousa, Altice Labs**

[jorge-m-sousa@alticelabs.com](mailto:jorge-m-sousa@alticelabs.com)

**Jorge Miguel Monteiro, Altice Labs**

[jorge-m-monteiro@alticelabs.com](mailto:jorge-m-monteiro@alticelabs.com)

**Albino Manuel Martinho, Altice Labs**

[albino-m-martinho@alticelabs.com](mailto:albino-m-martinho@alticelabs.com)

**Sara da Silva Furão, Altice Labs**

[sara-s-furao@alticelabs.com](mailto:sara-s-furao@alticelabs.com)

Studies predict that by 2020 consumers will manage 85% of their relationship with any enterprise without interacting with a human. Virtual Assistants are becoming the new apps and in essence, companies using them will be the ones making a difference staying ahead of the curve.

This article describes the BOTSchool platform from Altice Labs, allowing to offer a new, simplified and more natural channel over the existing services.

---

### Keywords

BOTs; Virtual Assistants; NLU; NLP; NLG; Omni channel

## Introduction

According to Gartner, end-user spending for the worldwide virtual assistant (VA)-enabled wireless speaker market will reach 2.1 billion USD by 2020, up from 360 million USD in 2015. By 2020, 3.3% of global households are forecast to have adopted a VA-enabled wireless speaker [1]. The VA are an evolution of the existent chatbot systems, being smarter and closer to human behaviour. A bot (short for “robot”) is an automated program that runs over the Internet. Some bots run automatically, while others only execute commands when they receive specific input. A chatbot is, therefore, an automated program that tries to simulate a conversation with human users.

Studies predict that by 2020 consumers will manage 85% of their relationship with any enterprise without interacting with a human. VA are becoming the new apps and in essence, companies using them will be the ones making a difference staying ahead of the curve [2]. Examples of a traditionally “low-tech” business adopting VA to grow their business are appearing everywhere. With this new turn towards a more fluid communication and a technology that better understands the human intent, the opportunities for brands to develop impactful marketing campaigns are abundant. In order to keep ahead and make sure they are reaping all the benefits of the virtual assistant revolution, marketers are adopting a multichannel strategy that incorporates digital assistants into their existing marketing ecosystem [3]. Things like conversational language, optimising voice search with the terms **how**, **why**, **when**, **where**, **what** and **who** will be part of any business, where these new digital entities will be the lead actor on the online exchange with consumers in the future.

Having these opportunities in mind, this article describes the BOTSchool platform from Altice Labs, whose objective is to help on the creation of personal virtual assistants, tailored for the needs of individuals, Small and Medium-sized

Enterprises (SME) and Large Enterprises, allowing to offer a new, simplified and more natural channel over the existing services.

## A new digital transformation

Whether integrated into a device or piece of software, nearly all assistants respond to the same objective: creating personalised conversations and interactions with consumers. Virtual assistants do not just find your answers, they will also know you, being able to help on a daily basis: appointments from your calendar, promises to get back to someone via e-mail, alerts to the deadline for a task on your to-do list, when it is time to take medication, help monitor your health and much more. It is predicted that by 2021 almost 1.8 billion people will be using some sort of virtual assistant [4]. There is no doubt that consumers see the value in using digital assistants to make lives easier and more productive.

In our assessment of this technology, we see digital virtual assistants as a revolution, and not just as “fad”.

Typically the newer generation is more receptive in the adoption of recent technology. However, given the reach and applicability of virtual assistants, adoption has traversed the traditional demographic and generational divide, with research showing that the adoption and usage of virtual assistants of any kind rising across all groups [5].

What makes this technology so special and receptive is the ability to use it and understand the way humans communicate, such as voice or writing. The dialogues, context perception or even cognitive recognition as understanding, speech, or even the recognition of emotions, which will change the way people use and access technology, encouraging more proactive engagement on the part of the users.

According to Mary Meeker's 2018 Internet Trends Report [6], by 2020, 50% of all searches will be made via voice. When we ask for a question, we use expressions such as **how, why, when, where, what** and **who**, eventually sharing to reveal more information about the true intention, or even show the true emotions about the subject in question. This deeper knowledge of the consumer mentality and/or subconscious creates a set of data and insights that allow a better definition of the strategy to address the true motive in question.

These sources of information, as well as real knowledge about the consumer, allow brands, companies and their marketing teams to better identify new opportunities or new strategies, improving the overall interactions' efficiency with users, allowing in the end, to know/recognise the best campaign/offer for a person or and more important what truly motivates/excites them. This digital channel is a new window, always ready to use and available to the user, representing the beginning of a new era, where the technology will become available to all, regardless of the limitations that may exist on the user's side.

When a company virtual assistant is part of a conversation, your brand is effectively engaging one-to-one in real-time with a consumer, helping to form brand loyalty and increase basket conversion rates.

## Market trends and business overview

In order to better understand the market dynamic and opportunity, it is presented a brief business market overview of the ecosystem. The first big wave started on the chatbots, with messaging apps like Facebook Messenger (since 2016), WeChat (since 2013), WhatsApp, Kik, Slack, Line, Telegram, or simply via SMS. They are used for business-to-consumer (B2C) customer service, sales and marketing.

Other companies explore ways to use chatbots internally, for example for Customer Support, Human Resources, or even in Internet of things (IoT) projects. Overstock, for one, has reportedly launched a chatbot named Mila to automate certain simple, yet time-consuming, processes for requesting a sick leave. SAP partnered with Kore Inc, a US-based chatbot platform vendor, to build enterprise-oriented chatbots for certain SAP products like SAP Hana Cloud Platform, SAP Cloud for Customer (C4C), SAP Success Factors and Concur. Other large companies such as Lloyds Banking Group, Royal Bank of Scotland, Renault and Citroën are now using automated online assistants instead of call centres with humans to provide the first point of contact.

In 2016, Facebook Messenger allowed developers to place chatbots on their platform. There were 30 000 bots created for Messenger in the first six months, rising to 100 000 by September 2017.

Since September 2017, this has also been as part of a pilot program on WhatsApp, the world's most widely used chat app - airlines KLM and Aeroméxico both announced their participation in its testing; both airlines had previously launched customer services on the Facebook Messenger platform.

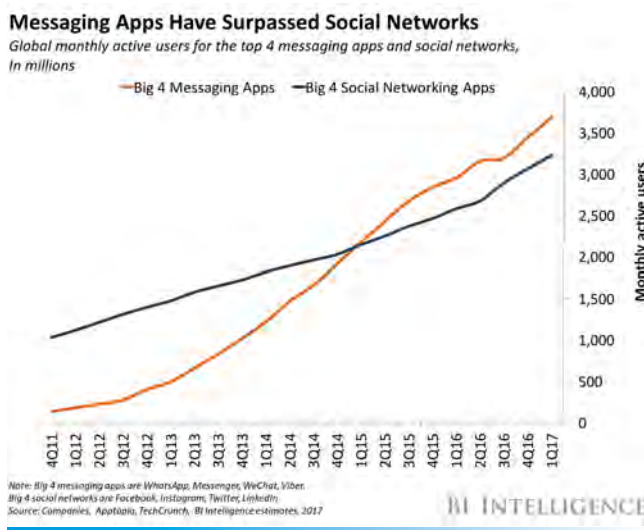
Services that companies provide via messaging apps may be extensive, encompassing chatbots, but also including live chat, and push notifications with purchase confirmations, trip information and promotions. In the case of airlines, there may be interactive functions such as flight search, check-in, and obtaining boarding passes, while e-commerce retailers might push, for example, order tracking notifications.

The bots usually appear as one of the user's contacts, but can sometimes act as participants in a group chat.

Many banks and insurers, media and e-commerce companies, airlines and hotel chains, retailers, health care providers, government entities and restaurant chains have launched their own chatbots to answer simple questions,

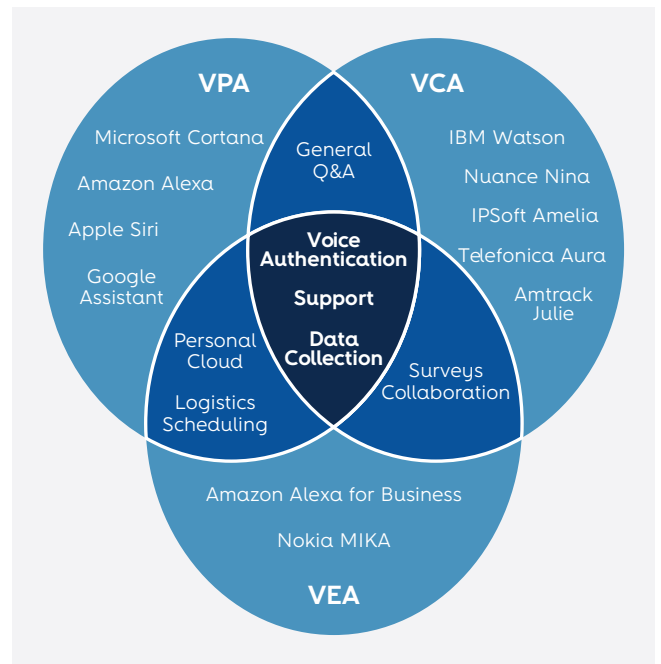
increase end customer engagement, promote their products and services, and give their customers a more convenient and easier way to order from them.

To have an idea of the trend and impact on people of chats, the four top messaging apps surpassed the top four social networks in terms of global monthly active users in 2015, and they continue to grow. The more activity messaging apps receive from users, the greater the opportunity for chatbots to facilitate usage of messaging apps (see **Figure 1**) [7].



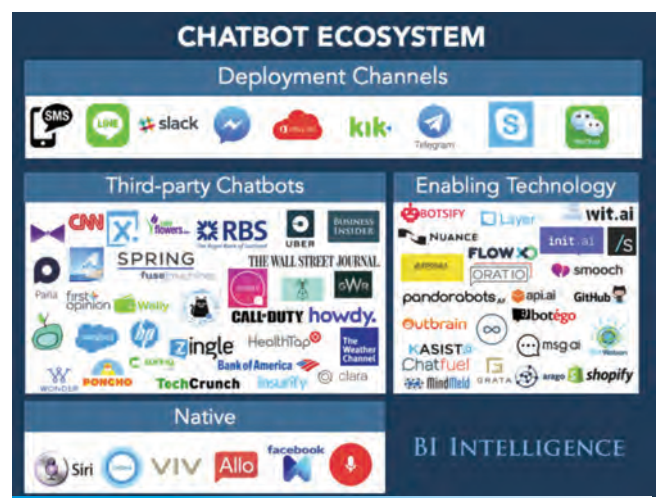
**FIGURE 1** – Messaging Apps vs social network evolution [7]

Previous generations of chatbots were present on company websites, e.g. Ask Jenn from Alaska Airlines which debuted in 2008 or Expedia’s virtual customer service agent which launched in 2011. The newer generation of chatbots, a virtual assistant, includes IBM Watson-powered “Rocky”, introduced in February 2017 by the New York City-based startup and e-commerce platform Rare Carat to assist novice diamond buyers through the daunting process of purchasing a diamond. Today, there are several different types of VA: virtual personal assistants (VPA), virtual customer assistants (VCA) and virtual employee assistants (VEA), shown in **Figure 2** [3].



**FIGURE 2** – Conversational platform: typical use cases [3]

Indeed, we verify that a virtual assistant SaaS business ecosystem has been steadily growing. While a 2017 study showed only 4% of companies had launched a chatbot, a late 2016 study showed that 80% of businesses intended to have one by 2020. According to Business Insider [8], **Figure 3** represents a big picture of the Chatbot Ecosystem.



**FIGURE 3** – Chatbot ecosystem [8]

# BOTSchool

As any new emerging technology, the greatest challenge for chatbot's mainstream adoption by companies is in its move from the domain of tech-savvy experts to the domain of regular users.

Many companies currently struggle with digital transformation and the need to support all different channels such as web or even social networks. But even worse, they are faced with issues related to Artificial Intelligence [9], which require a great deal of knowledge and effort to be addressed. Issues like: 1) need to have a large training data set, correctly categorised in order to ensure a correct response/accuracy; 2) inexistent of collaboration on the self-training, which make the learning process harder; 3) hard to take advantage of external data sources, like API or even information that can exist in webpages; 4) content/knowledge is not built to be reused

among Virtual Assistants, which difficult the creation of other assistants; 5) difficult to control and guide the channel, adapting the content to the channel behaviour or capacities; 6) need to have deep know-how in Artificial Intelligence (AI) and Machine Learning (ML) algorithms and 7) conflict between dialogues and intents during the training or dialogues.

When digital transformation projects are coupled with the technologies complexity, the failure probability, dragged-out projects or unsatisfactory results are increased.

The BOTSchool solution from Altice Labs was created by based on existing problems and in the necessity of delivering value both to the chatbot's end-users and the companies that employ them, in an easy and fast to implement way. **Figure 4** shows the BOTSchool solution architecture. The BOTSchool was therefore designed with a simple mission/purpose: **"Simplify the human interactions with the digital world"**.

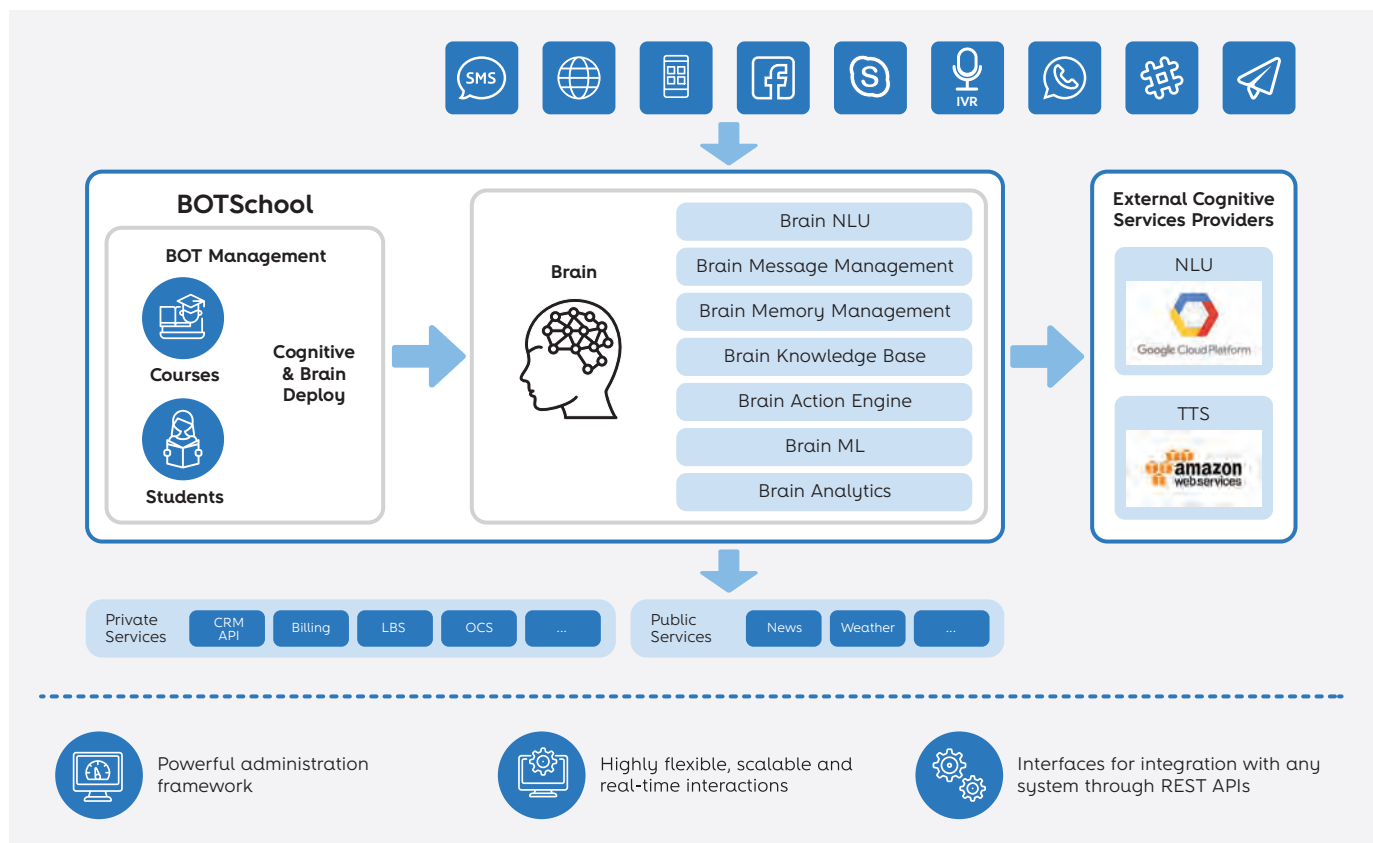


FIGURE 4 - BOTSchool solution architecture

In this way, the BOTSchool's objective is the supply of the following service/functionalities:

- Simple self-care, where the Virtual Assistant can be managed
- Easy on-boarding dialogues training
- Integration with external API
- Ability to supply the Virtual Assistant in the social networks e.g: Facebook, Skype
- Interactions via chat or voice
- Easy integration in any webpage
- Interactions overview/analytics.

## BOTSchool NLP

Natural Language Processing (NLP) is the umbrella term that describes the ability to break down what the human says, understand what it means, decide on an appropriate action and form a response delivered on comprehensible

language back to the human user. As described on **Figure 5**, both Natural-Language Generation (NLG) and Natural Language Understanding (NLU) are closely related to NLP, which provides the processing ability to the two other technologies [9].

The effectiveness of an NLP system depends on its ability to understand what the human has said and then comprehend it correctly. As the system learns and grows, its understanding of human language and thought processes becomes more and more natural and life-like.

## BOTSchool NLU

NLU stands for Nature Language Understanding, having the objective to recognise the intention of an interaction. When an intention is received, the NLU engine vectorises the text, confronting the data matrix with the Machine Learning (ML) model. As a result, one or a set of intentions is returned with the score/accuracy associated with each of them.

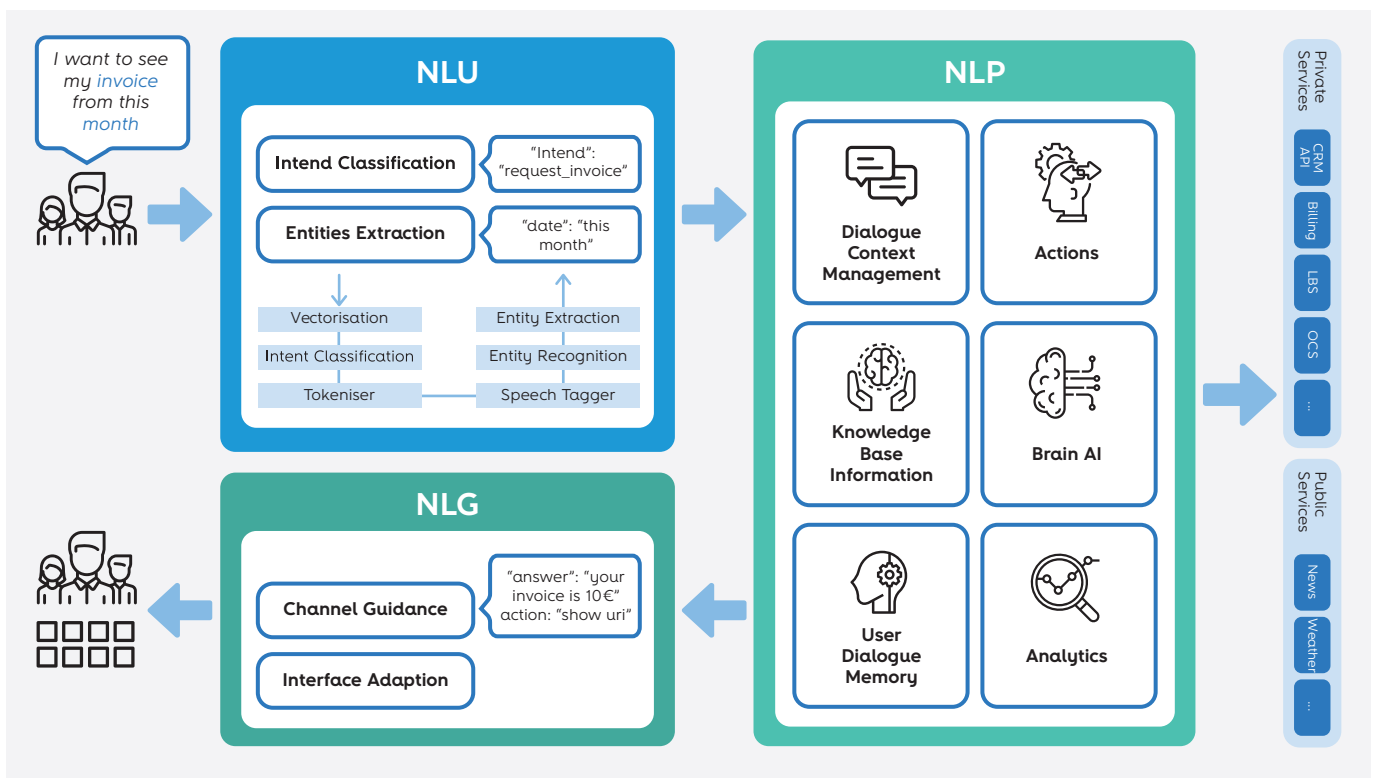


FIGURE 5 – Processing flow of a phrase/dialogue



The learning process uses the ML model based on supervised learning, achieved by using a training data set. We use this machine learning approach, because from the training dataset we know the correct answers, ensuring “as a teacher” that the learning process goes as intended. The algorithm iteratively makes predictions on the training data that can be corrected/made more accurate by the teacher.

There are three main components in the training data set: text, intent and entities, being denominated as “training phrase” to the search query (text). Each training phrase is associated with at least one intent, and similar training phrases can have the same intent. A training phrase has at least an intent and an intent can have more than one training phrase. The entities are specific parts of the training phrase which need to be identified. Entities can span multiple words and in fact, the value field does not have to correspond exactly to the substring in the training phrase. That way is possible to map synonyms, or misspellings, to the same value. Entities that have the same value will be treated as synonyms.

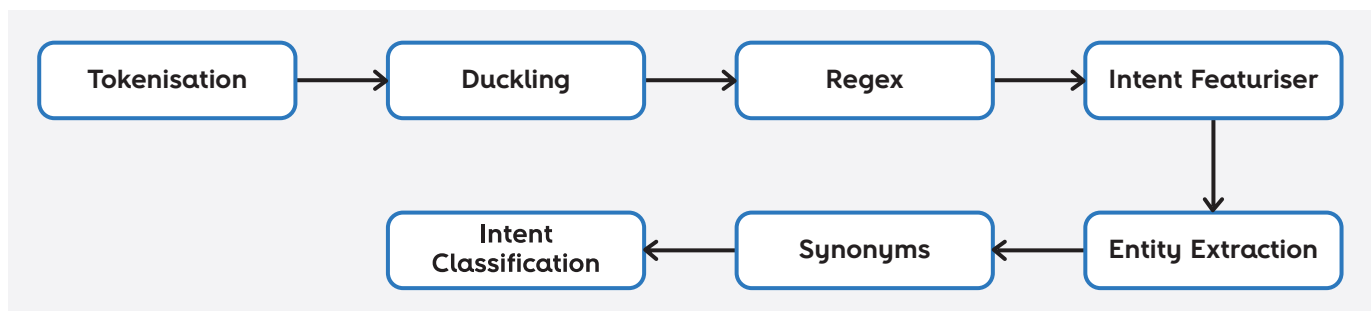
The training process can be tuned by parameterisation when a train data set is executed. For instance, information about VA language or accuracy pretended by the ML are things that can be used to control the ML behaviour.

The training itself is based on a pipeline, where each step has specific functions. These components are executed one after another in a so-called processing pipeline. Each component process has his input and creates an output.

The output can be used by any component that comes after this component in the pipeline.

The training data set is fed into the NLP/NLU pipeline and gets converted into an ML model that performs intent classification and entity extraction. There are several classifiers, with the confidence score being calculated based on the used classifier. The confidence score is not a true prediction probability, it is just a metric defined by the model that describes how similar the input was from the training data. The pipeline process used by the BOTSchool can be found below (see **Figure 6**).

- **Tokenisation:** During processing, first the text is tokenised, i.e. segmented it into words, punctuation and so on. This is done by applying rules specific to each language.
- **Duckling:** Add duckling support to the pipeline to unify entity types (e.g. to retrieve common date/number formats). For example, turn expressions like “next Saturday at 9 am” into actual date/time objects that can be used. Can also handle durations like “one hour”, amounts of money, distances, and ordinals.
- **Regex:** supports intent and entity classification. During training, the regex intent creates a list of regular expressions defined in the training data format. If an expression is found in the input, a feature will be set, that will later be fed into intent classifier/entity extractor to simplify classification (assuming the classifier has learned during the training phase where an intent can be retrieved).



**FIGURE 6** – Pipeline process

- **Intent Featuriser:** In order to apply machine learning algorithms to conversational AI, we need to build up vector representations of conversations. Each story corresponds to a tracker which consists of the states of the conversation just before each action was taken.
- **Entity Extraction:** Implements conditional random fields to do named entity recognition. Conditional Random Field algorithm (CRF) can be thought of as an undirected Markov chain where the time steps are words and the states are entity classes. Features of the words (capitalisation, POS tagging, etc.) give probabilities to certain entity classes, as are transitions between neighbouring entity tags: the most likely set of tags is then calculated and returned.
- **Synonyms:** If the training data contains defined synonyms this component will make sure that detected entity values will be mapped to the same value.
- **Intent Classification:** The intent classifier trains a linear Support Vector Machine (SVM) which gets optimised using a grid search. In addition to other classifiers, it also provides rankings of the labels that did not “win”. The BOTSchool ML intent classifier needs to be preceded by a featuriser in the pipeline. This creates the features used for the classification.

After the pipeline execution, we got a training model. When the systems receive a text phrase, it returns an intent rank with a score confidence along with the predicted intent. In addition, we created a “Context Manager API” that returns all contexts of a session. With this, we are able to verify if the input contexts of all intents presented on the ranking match with the contexts of the session. The intent chosen to use at the answer generation is the one that matches all input contexts with the session contexts. In case, we have more than one, the one with the best score is chosen.

## Conclusion

Digital transformation is one of the biggest challenges companies face nowadays. Customers are increasingly demanding contextualised, quicker, and tailored responses. The existing information is growing, faster and faster. Virtual assistants, with cognitive abilities, are actually acting as catalysts for this new wave/tsunami, and being expected to change the way we relate to the digital world. It is within this framework and given the difficulties that exist for a normal company to adopt this new AI technology, that the Altice Labs’ BOTSchool product was created. A solution designed by companies for companies, where simplicity, flexibility and agility are central parts of a successful strategy and implementation.

---

## References

- [1] Gartner, “Press Release - Gartner Says Worldwide Spending on VPA-Enabled Wireless Speakers Will Top \$2 Billion by 2020,” 3 October 2016. [Online]. Available: <https://www.gartner.com/newsroom/id/3464317>.
- [2] Gartner, “Gartner Customer 360 Summit 2011,” 30 March 2011. [Online]. Available: [https://www.gartner.com/imagesrv/summits/docs/na/customer-360/C360\\_2011\\_brochure\\_FINAL.pdf](https://www.gartner.com/imagesrv/summits/docs/na/customer-360/C360_2011_brochure_FINAL.pdf).

- [3] Samantha Searle, Annette Jump, Jessica Ekholm, Anthony Mullen, "How Virtual Assistants, Immersive Experiences and Robots Will Impact Your Organization," Gartner, 26 March 2018. [Online]. Available: <https://www.gartner.com/doc/3869766/virtual-assistants-immersive-experiences-robots>.
- [4] Eden Zoller, "The Evolution of Consumer AI Assistants," Ovum, 11 September 2018. [Online]. Available: <https://ovum.informa.com/resources/product-content/the-evolution-of-consumer-ai-assistants>.
- [5] Julie A. Ask, Michael Facemire and Andrew Hogan, "The State Of Chatbots," Forrester, 20 October 2016. [Online]. Available: <https://info.247.ai/rs/074-HBW-141/images/Forrester-The-State-Of-Chatbots.pdf>.
- [6] Mary Meeker, "Internet Trends 2018 Report," Kleiner Perkins, 30 May 2018. [Online]. Available: <https://www.slideshare.net/kleinerperkins/internet-trends-report-2018-99574140/3>.
- [7] Business Insider, "Messaging apps are now bigger than social networks," Business Insider, 20 September 2016. [Online]. Available: <https://www.businessinsider.com/the-messaging-app-report-2015-11>.
- [8] L. Beaver, "CHATBOTS EXPLAINED: Why businesses should be paying attention to the chatbot revolution," Business Insider, 4 March 2016. [Online]. Available: <https://www.businessinsider.com/chatbots-explained-why-businesses-should-be-paying-attention-to-the-chatbot-revolution-2016-7>.
- [9] Svetlana Sicular, Kenneth F. Brant, "Hype Cycle for Artificial Intelligence, 2018," Gartner, 24 July 2018. [Online]. Available: <https://www.gartner.com/doc/3883863/hype-cycle-artificial-intelligence>.



## 07

# Mining the TV consumption: a basis for a TV personalisation system

**Bernardo Cardoso, Altice Labs**

[bernardo@alticelabs.com](mailto:bernardo@alticelabs.com)

**Jorge Ferraz de Abreu, University of Aveiro**

[jfa@ua.pt](mailto:jfa@ua.pt)

With the continuous growth of channels and content available in a typical interactive TV service, viewers can become increasingly frustrated, struggling to select which programs to watch. Personalisation have been pointed out as a way to mitigate this problem, being an active area of research in Altice Labs and University of Aveiro, especially now with the recent Sofia TV UI launch.

---

## Keywords

MEO Sofia UI; TV recommendations; Catch-up TV; Linear TV; EPG; STB

## Introduction

MEO launched the new Set-Top Box (STB) 4K Wi-Fi “Sofia” in March, with access to a refreshed user interface, anchored in a logic of content discovery, customisation and speed. This unveiling was also the moment used to inform that the service now materialises a more personalised and engaging attitude of MEO, incorporated into the new MEO brand signature “humaniza-te” (humanise yourself), referring to the new contexts of artificial intelligence.

This need for personalisation comes from the perception that the amount of content available on pay-TV platforms has increased significantly. Not only has the number of available channels grown far larger, but additional features and products have emerged to increase available content and improve management and access to the already existing. These functionalities include, for example, VOD, catch-up TV, DVR, SVOD and other sources of video content available through the Internet [1].

All this increment in content and functionalities results in a wide offer of possibilities for the users, and this extends even to the devices in which the consumption can be enjoyed. Today, watching an hour-long TV show does not necessarily involve using a regular television set - it can be done, anytime, anywhere, using increasingly accessible and omnipresent devices such as computers, tablets and smartphones. However, there are still some factors that make the big screen in the living room the preferred device to visualise audio-visual content. It's not by chance that, despite more and more types of screens being available, television sets continue to sell well and are becoming bigger and more advanced [2, 3].

## Personalisation in linear TV

With the increasing number of channels, zapping and TV magazines are not effective means for

content selection anymore. The EPG (Electronic Program Guide) with its traditional grid aspect is a tool to help linear TV consumption on interactive platforms but has lost its appeal long ago. On VOD platforms, recommendation systems come as a proposal to improve the process of discovering new films to watch, with relative success. These systems end up having an even more effective impact on SVOD platforms, especially Netflix. However, the various attempts to apply the same recommendation algorithms to linear contents didn't have a great success [4].

Linear content has some particularities that make the creation of an effective recommendation systems quite complex. The main difficulty is that the contents are available only for a short time and are constantly changing [5]. Even systems that have access to a catch-up TV platform need to cope with the fact that new “recordings” are continuously replacing old items, since the temporal window of catch-up is often limited to only 7 days [1].

However, the consumption of television content often follows a regular pattern with a recurring type of viewing. When watching linear TV, the viewer tends to have consumption habits associated with certain moments of the day and diligently follows certain programs across a small number of channels. In contrast, in a VOD system, the viewer aims to discover new movies to watch, and usually doesn't watch the same movie several times [5]. This regularity in television consumption was also identified by other authors [6-8], who use it as a basis for proposing new approaches to linear television recommendations, taking advantage of elements such as the day of the week, the time of day, the channel, the type of program, as well as the very people who watch the content, that is, the viewer and eventual companions.

## Sofia TV UI personalisation

For the development of the linear TV personalisation available in Sofia TV UI, Altice Labs relied on some of the aforementioned

concepts, more precisely the day of the week, the time of day and the channel. The current algorithm provides a personalised list of channels, taking into account the time each channel is watched by each STB. This personalisation, however, is not a simple list of the top channels consumed by that STB. As the aggregates are done considering the day of the week, there are different mini-tops for each Monday, Tuesday, Wednesday, etc. Additionally, each day is broken in several time periods, namely: morning, noon, afternoon, primetime and night. In this way the recommendations presented in a Saturday morning are completely independent from the primetime in Monday. The reasoning is that in a typical family household, Saturday mornings are used by children to watch cartoons, whereas Mondays TV primetime is a family affair, and night time watching is only available to grown-ups. This is consistent with experiments conducted by both the academia and the industry [8, 9].

There are also some additional biases introduced in the algorithm, to increase prevalence of channels watched in the most recent week. This is done by using a moving average during the calculation, which gives more weight to more recent data, than historical one. The reason for this is that TV shows debut and retire in different channels all the times, and channel preferences also vary with time. For programs like news casts, that are watched with a higher regularity, this has no impact, but for a user that used to watch a soap opera on channel A and then started to follow a new soap opera on channel B, it could be very important.

The simplified overall workflow for the linear TV personalisation on Sofia UI is like this:

1. In each STB there is a process assessing, once a minute, if the user is watching a live channel, and if so, records the channel in temporary memory buckets;
2. These buckets are reported to the server side every couple of hours or when they are full;
3. On the server side this new data is merged with the previous information, and a new

personalised recommendation is created for the corresponding day of the week and time slot.

Recommendations are then presented to the user as a series of entries on the first item of MEO's main menu, like the one seen in **Figure 1**. Although the accounting is done on a channel-based approach, the recommendations are presented as a series of programs from those channels. The focus in presenting the information in a more program oriented, or a more content oriented way is pervasive in the Sofia UI refresh. For instance, on the new browse bar this concept has been implemented as a continuous strip of programs



**FIGURE 1** – Main menu with linear TV recommendations



**FIGURE 2** – New browser bar

as can be seen in **Figure 2**. In this screen we can see that the current program appears selected on the centre. The left side show programs from the past that can be restarted. The right side presents upcoming programs, enabling the user to get more information or schedule a recording.

But the transition from a channel to program is even more evident in the new catch-up TV application, for which the content has now a much more noticeable position, and channels, while still being there, have a less prominent placement.

In the catch-up TV application there are also many new personalisation features. The first one is to present to the user the channels that have been watched more recently, **Figure 3**. This goes



**FIGURE 3** – Catch-up TV application with “Recent Channels” and “Continue Watching” lists

in line with the idea that most users only watch content from a limited group of channels, and to present these channels on top will function as a welcoming shortcut. Next, the UI shows the programs that the user began watching and didn’t finish. This is also a very personalised and very simple way to enable users to continue viewing something they started before. It also works as an easy way to access other episodes from a recurrent show, and likewise an additional shortcut for the programs the user usually watches, by using the Transversal Program Info (TPI, depicted in **Figure 4**).



**FIGURE 4** – Transversal Program Info with access to all the episodes of a series

Finally, in the last program list of the catch-up TV application, the user can find a set of content recommended for his household, based on previous content consumptions. This approach attempts to propose content the user could be interested in, but didn’t watch yet, **Figure 5**, thus potentially interesting content.

The solution is a hybrid approach that uses Azure Cognitive Services in the cloud to produce recommendations for catch-up TV programs. These recommendations are then cached locally on-premises to ensure the best performance.

Although these personalisation features had a very good acceptance in user surveys conducted





**FIGURE 5** – Catch-up TV application with “Recommend for You” list

by the marketing area, there is a set of limitations already identified in this initial setup. To start with, personalisation is split in two worlds, one for live and another for catch-up TV. For instance, if a user started watching a program live and didn't finish it, this program will not appear in the “Continue Watching” List (**Figure 3**), since only content watched on-demand goes into that list. Similarly, the “Recent channels” list refers only to recently used channels on the catch-up TV application and not to recently channels watched live.

While the recommendations for live channels (**Figure 1**) consider the time of day and day of the week, none of the personalised lists from the catch-up TV app does it. In households with children this usually generates a list full of kids programming, independently of the time of day the user access it. This could be fine if it is 7:00 p.m. and the app is been used by a child but is not so nice when used at 11:00 p.m. by the father or the mother.

Another identified issue is that, in the Live suggestions (**Figure 1**), data collection and recommendation calculation are done considering only the channel and not actually the program. Normally this is not a concern, because channels have a regular programming schedule, thus the same program usually broadcast in the same time slot. However, when trying to create

a unified recommendation app, having only information about the channels and not the programs watched live doesn't merge easily with information about individual programs watched on-demand from catch-up TV.

Lastly, the personalisation and recommendations are scattered throughout the user interface. This is intentional, because it enables the user to have additional programming suggestions present in many areas of the UI, but there is no single place where the user can get a set of recommendations that propose simultaneously live and on-demand programs. This, inadvertently, also means that the user must explore the user interface in order to get some recommendations, contrasting with what happens in Netflix, for example, where recommendations are the first thing the user sees upon turning on the Netflix application.

## Going deeper in personalisation

To address the identified issues and simultaneously introduce a set of innovative features, Altice Labs and University of Aveiro are currently researching and developing an interactive TV application conveying the following objectives:

- Have a recommendation algorithm, based on the consumption data of each individual STB, considering not only the programs watched live, but also the ones from catch-up TV and from the DVR. This algorithm will use intrinsic characteristics of the programs (for example: football games are different from a news telecast, which are different from movies). However, the algorithm will continue to consider as very relevant the time of day and day of the week where consumption normally takes place.
- Be proactive, that is, the system will automatically propose contents at certain moments, without an explicit request from the

user. This can happen, for example, when the STB is turned on or when it is resumed from standby, when a zapping process is detected or when the end of a program is detected.

- Have a retrospective consumption view, that is, the system will have a feature to allow the user to view their past consumption history, in order to be able to continue previous consumptions and give positive or negative feedback on these previous views. For the system it will not matter whether the programs were watched live or on-demand, both will be presented in a consistent timeline. This will apply the concept of the web browser history to a TV consumption setup.

Aligned with these objectives, the project intends to use a content-based filtering algorithm (CBF) and augment it with the temporal context [5]. The reason to choose a CBF algorithm instead of the much widely used collaborative filtering (CF) is that on mixed TV content different programs have completely different liveness or “shelf time”. For example, a football game has much more value when watched live, than when watched in a time shifted mode. In that respect, football, as most other live sports, is kind of a “hot meal”, that should be enjoyed as soon as possible, and that is not good re-heated in a microwave oven. A different kind of content is news. In this case it can be consumed live or a little bit later without a problem, but gets stalled shortly after, like bread. Then there are serialised contents, like soap operas and series. This kind of content has a long “shelf life”, but usually needs to be watched in a sequential fashion. In the other hand, movies are perhaps the kind of content with the longest “shelf life”, and although they have more value when first exhibited, they keep being valuable over time.

The intrinsic differences among the different kinds of content are highly relevant, most notably for the ordering and timeliness of presenting the recommendations. A concrete example will be the scenario of a STB in which there is a high consumption of football matches. In this case, in a proactive mode, when this STB is switched on and a football game is being broadcast, even

if it has already started a few minutes ago, the personalisation system will automatically switch to that channel. The rationale is that sport has such a short “shelf life” that it makes little sense to tune to the start of the game. Indeed, some consumers of sports content hate the IPTV and OTT systems because they have a few seconds delay over other transmission methods [10].

In contrast, in a news scenario, a STB where the newscast is normally watched at 8:00 p.m., when it is turned on at 8:45 p.m., it will proactive start the playback of the newscast from the beginning. However, the system will no longer have the same behaviour if the box is turned on at 10:00 p.m. The idea is that news have a short life expectancy, that tolerates some delay but beyond a certain threshold they become old news.

For different contents, which have longer shelf life, i.e., where deferred consumption is not a problem, the system will check if there is a strong temporal affinity between its consumption and a particular time of day/week. For example, children’s content that has a high consumption rate in the early morning, will only be proposed if it is within its normal time window. This way the system will never suggest children’s content at 10:00 p.m., if it is not normally watched at this time of the day.

Finally, other content which is not highly perishable and does not have high temporal affinities with certain hours or days of the week, will be proposed based on a classification that will use the content popularity within this STB, i.e. the system will propose first the content most watched on that STB. However, it will also consider their shelf life in terms of the catch-up TV window and, in that way, it will propose first the oldest non-watched content available in catch-up. Let’s take the example of a STB where the series A, B and C are normally viewed, but not always on the same day of the week or at the same time of day: in this case the system will start playing the episodes that are closest to being removed from the catch-up system.

The overall idea behind this approach is that the personalisation system will work a bit like a secretary or an intelligent assistant, who manages

her boss's agenda by scheduling recurring meetings for the usual time – for example: newscast every day at around 8:00 p.m., cartoons at 7:00 a.m. in workweek days and all morning at weekends, etc. Other things are scheduled according to the availability of the boss's time, in the most urgent order. However, important football games broadcasts may have priority over all things and implicate rescheduling other activities. This idea also embodies one of the general objectives of any personalisation system, which is to reduce and automate the content selection process, in this case restricting it to the contents that are typically consumed in a given STB, thus reducing the need for the user to go hunt content throughout different apps.

## Implementing it on the field

The currently developed application was inspired by existing areas of the Sofia UI, for instance the main screen looks like the new browse bar (Figure 2). As previously explained the browser bar on the standard Sofia UI allows the users to know what program is airing, what was aired recently (on the left side) and upcoming programs from a



FIGURE 6 – Initial screen for the prototype interactive TV application

FIGURE 7 – Prototype interactive TV application showing upcoming suggestions



single channel (on the right side). The prototype implementation for the proactive personalisation interactive TV application uses the same approach, the start screen shows, in the middle, the most relevant recommendation for that moment in time (**Figure 6**). If the user navigates to the right, it gets additional recommendations that make sense at that time of day, firstly the 8 more relevant programs inline, and then, if he wishes, a full grid for all the upcoming suggestions (**Figure 7**). If the user navigates to the left side, he can review the programs that have been watched in that STB, again inline he gets the first 8 programs, and then he can have a grid that can go back in time a full week (**Figure 8**).



The system being developed incorporates also a proactive mode (that can be turned off) where content is automatically proposed at certain moments, instead of the more traditional method of waiting for an explicit request from the user. There are three moments where the system will spontaneously engage. The first one is when the STB is turned on or when resuming from standby. This appears to be a good opportunity to present content suggestions to users, since today the STB just starts with the last channel used. In this situation the system will not only propose the content, but will automatically begin the content reproduction, following the rules and examples provided before, sports are tuned live, other contents are reproduced from the beginning. Another interesting moment to make content suggestions is when a program ends, and in this way the system can bring a sense of binge watching to a traditional pay-TV system. These two moments can make the platform a bit more aligned with current UX trends in the OTT video domain. Finally, we are researching ways to detect what we call mindless zapping or mindless channel surfing, that is, we are trying to identify when the user is hunting for something to watch, just changing channels continuously. Since this is a moment where the user is not actually watching something, it appears an excellent occasion to nudge in a suggestion system.



**FIGURE 8** – Prototype interactive TV application showing previously watched programs

## Outcomes from the initial prototypes

A set of prototypes has been developed to validate these ideas. Initially the focus was on the recommendation engine and algorithms and the main idea behind the algorithm is to infer a usage pattern from a STB. For this to actually work, the algorithm needs to recognise that, although each day has different programs, the separate episodes of a soap opera are really instances of that same soap opera, and should be recognised as such in the identified pattern. This is easy for recurring programs like soap operas, series, reality shows, etc., which have season and episode information on the EPG. It is also not difficult for things like regular newscasts because they all have the same name and air at more or less precise hours. However, against initial expectations, the system still found a lot of single events, which were not recognised as instances of any series. In a deeper analysis, these programs were found to be mostly sports broadcasts. This represents a bigger challenge than other kinds of programs. There are no references in the MEO EPG to allow to group this kind of events, not even any kind of information about the sport modality involved. To tackle this issue, the system will need to have a much-enriched EPG data that annotates each event with the sports modality, the competition and the teams playing. For the next prototype, this information will be obtained from a third-party provider and it will be integrated in a project specific database.

Another insight, provided by the initial prototypes, is that the history screen, which shows all the programs the STB has shown in the past, could be improved in line with recent trends, mostly in the mobile area, to show where the user has been spending his time. Features like Screen Time in iOS 12 or Dashboard in Android, help the user to understand where he spends his time on his mobile phone, and similarly the interactive TV app could also inform the user how much time is spent on a particular STB, and in which kind

of programs. The main issue currently is that the system doesn't know if someone is actually watching the TV, or even if the TV is turned on - this will influence a lot this analysis and it has been already previously identified as an issue with the current data recollection. More research is still needed to try to minimise this problem, in order to implement this potentially relevant feature.

Other topic that the prototype doesn't address is the existence of multiple profiles in a shared STB. In most households the TV is a device used by multiple persons, whether by different people at different times of the day, or by more than one person at a time. In addition, there are times when the TV is switched on and no one is watching, or the person in front of the TV is not paying attention, using it as background noise. These realities make an effective multi-profile implementation challenging. In which profile should a program count as seen when more than one person is watching? How to perceive that no one is watching, so that this TV watching time does not count? Besides that, the platform where the prototype is implemented doesn't have any notion of profiles and implementing it throughout is beyond the scope of this research.

Although a multi-profile scenario is not built into the system, a great emphasis will be placed on the time of day and the day of the week in which consumption occurred, and it is expected that this will address some of the shortcomings of not having multi-profiles. We hope that the close relationship between the moments of consumption and the contents consumed in those moments is sufficiently strong so that the personalisation makes sense [8]. The goal is that if at 7:00 a.m. an STB sees "Zig Zag" on RTP2 (a children's content) the system will propose this program only at that time and do not propose it at 10:00 p.m., when children should already be asleep. However, if, normally, at 9:00 p.m., a soap opera is watched in an STB, the system will suggest it, even though on that day only the husband is watching TV and he has no interest in soap operas. The purpose of this personalisation system is to function as

a system that automates the routine, not the exceptions. Moreover, this personalisation system is in fact applied at the STB level, and not to the person level. If the time context will be sufficient to convey a perception of personalisation to each person, is an issue to be validated during the user evaluation phase.

Although not implemented yet, the final interactive TV application prototype that will be evaluated will have privacy preserving features, like incognito mode and history clearing. It will also have the ability to fully disable data collection, so the user can control when the system can gather information and when it shouldn't.

## Conclusions

Personalisation is rightfully considered one of the keys for the future in user interaction on pay TV platforms. In line with that idea, this article

presented the current state of personalisation in the new MEO Sofia UI, from a technical point of view and also presenting where in the user interface the personalisation is materialised. Next, some of the issues identified in the current implementation were also highlighted.


As usual in Altice Labs, research for possible improvements is already ongoing and prototypes are being developed. This article reveals some of the most promising ideas and the current state of implementation, uncovering some of additional insights and areas where more research and future work is still needed.

The proposed system is being designed and developed in order to automatically recognise a STB consumption pattern and uses this inferred pattern to automate the day-to-day operations of TV content selection, trying to simplify the user's life. The next step is to actually try this system in the field, first with beta testers and then, if successful, launch it in production.

## References

- [1] Abreu J, Nogueira J, Becker V, Cardoso B (2017) Survey of Catch-up TV and other time-shift services: a comprehensive analysis and taxonomy of linear and nonlinear television. *Telecommun Syst* 64:57-74.
- [2] statista (2018) Global TV unit sales from 2012 to 2017. [Online]. Available: <https://www.statista.com/statistics/461316/global-tv-unit-sales/>.
- [3] statista (2018) Forecast of worldwide smart TV unit shipments from 2011 to 2017. [Online]. Available: <https://www.statista.com/statistics/314616/smart-tv-unit-shipment-worldwide-forecast/>.
- [4] Churchill EF (2013) Putting the Person Back into Personalization. *interactions* 20:12-15.
- [5] Turrin R, Condorelli A, Cremonesi P, Pagano R (2014) Time-based TV programs prediction. 1st Work. Recomm. Syst. Telev. Online Video ACM RecSys.
- [6] Aharon M, Hillel E, Kagian A, et al (2015) Watch-It-Next: A Contextual TV Recommendation System. In: *Proceedings, Part III, Eur. Conf. Mach. Learn. Knowl. Discov. Databases - Vol. 9286*. Springer-Verlag New York, Inc., New York, NY, USA, pp 180-195.

- [7] Gonçalves D, Costa M, Couto FM (2016) A Flexible Recommendation System for Cable TV. 3rd Work. Recomm. Syst. Telev. Online Video (RecSysTV 2016).
- [8] Vanattenhoven J, Geerts D (2015) Contextual aspects of typical viewing situations: a new perspective for recommending television and video content. *Pers Ubiquitous Comput* 19:761–779.
- [9] Gonçalves D, Costa M, Couto FM (2016) A Large-Scale Characterization of User Behaviour in Cable TV.
- [10] Mekuria R, Cesar P, Bulterman D (2012) Digital TV: the effect of delay when watching football. In: *Proc. 10th Eur. Conf. Interact. tv video - EuroiTV '12*. ACM Press, New York, New York, USA, p 71.





# 08

## Technology Radar, for the win

**Fausto de Carvalho, Altice Labs**

[CFausto@alticelabs.com](mailto:CFausto@alticelabs.com)

**Mário Moreira, Altice Labs**

[Moreira@alticelabs.com](mailto:Moreira@alticelabs.com)

The technology radar elaborated by Altice Labs is about change. It's not a "reference stack", a repository of approved technologies. It offers a significant help for technological risk prevention as well as trend anticipation. It also provides information that can be used to better shape the product roadmaps and, on a global perspective, it gives to every technical team a perspective about which technologies will have a higher relevance and/or demand within the organisation in the near future.

---

### Keywords

Technology; Tools; Language & frameworks;  
Methods & practices; Big data technologies

## Introduction

Altice Labs is responsible for monitoring the tech community, aiming to foresee technological and architectural trends with short and medium term impact on the market. Among other initiatives, and learning from experiences within other organisations, this activity is carried out methodically by a Technology Coordination team who maintains a wiki work space for open discussion and publishes a main document, named the Technology Radar [1].

The Technology Radar has been updated yearly since 2013. Its major objective is to disseminate a set of recommendations and conclusions to technical managers, system architects, researchers and all those who are dealing with software technology, so that this knowledge may be taken into account in the systems' design process, enabling a more conscious planned evolution of those systems. Furthermore, it also serves as a guide for exploratory innovation projects, since it provides insight on newcomer technologies that can be experienced in those projects.

This document offers a significant help for risk prevention as well as trend anticipation. It provides information that can be used to better shape the product roadmaps. On a global perspective, it gives to every technical personnel

a perspective about which technologies will have a higher relevance and/or demand within the organisation in the near future.

The Technology Radar is about **change**. It's not the "reference stack", a repository of approved technologies. Instead it focuses only in technologies that are in use at Altice Labs and for a variety of reasons should stop being used, or in technologies that are not being used and should be. It won't say anything about technologies that aren't used and shouldn't be used or about technologies that are used and should continue to be used, that is, when there is no need to change.

For easier understanding and usage, the Radar is organised in four categories: "Platforms", "Tools", "Languages and Frameworks" and "Methodologies and Practices". Sometimes a technology may fit into more than one category, and in those cases a best fit is used, avoiding duplications.

For each category, there are four possible recommendations: Explore, Experiment, Adopt and Avoid. These are shown in **Table 1**, reflecting the set of conclusions at a given moment. Over time, recommendations may change.

It's important to keep in mind that this work has a clear focus on the context of systems development at Altice Labs. In other contexts those same technologies can have different recommendations.

 <p><b>EXPLORE</b></p> <p>A new technology that may have high potential. Not safe yet to use in mission critical systems. May be used in exploratory projects and other low risk situations.</p>	 <p><b>EXPERIMENT</b></p> <p>A mature technology that can be experimented in production systems. Must be used in project with moderated risk.</p>	 <p><b>ADOPT</b></p> <p>The best technologies that we aren't using yet. It provides good benefits in comparison with alternative technologies currently in use. Can and should be used in mission critical systems in a generic way.</p>	 <p><b>AVOID</b></p> <p>Obsolete technology. There are better alternatives. It should be planning abandoning these technologies and replacing them for better options.</p>
---	--	---	---

**TABLE 1** – Recommendations Levels and Criteria

Attempting to anticipate technology and market trends is far from being an exact science. Our evaluation takes into account multiple factors/ points of view, such as the following:

- **Benefits:** Which benefits the usage of this new technology will bring to Altice Labs' products and solutions (ex: improved efficiency, reduced development effort, better performance, reduced costs);
- **Cost of change:** Costs incurred by adopting the new technology, easiness of adoption, major entrance barriers (ex: disruption with existing procedures, methodologies and technologies, learning curve, tools availability);
- **Maturity and support:** Asserts the technology maturity and robustness in order to be used in critical systems that are required to be available 24x7 and how easily a problem can

be fixed (ex: how easily can we have support from third parties);

- **Popularity and roadmap:** What's the current community size and growth trend, real world usage references, planned roadmap features, interfaces quality and stability, existing libraries, etc. Will the technology be available in the years to come?
- **Technology internal usage:** Is the technology already being used within Altice Labs? If so, is it being used properly, misused, underused or overused?

In the next chapters we will look briefly into several examples of technologies addressed in the current Technology Radar edition, shown in **Table 2**, which sets the course for this year's activity and onwards.

	EXPLORE	EXPERIMENT	ADOPT	AVOID
PLATFORMS	Node.js Vert.x	Cassandra HAproxy (+)	Jetty Nginx PostgreSQL Redis Kafka Hadoop Elasticsearch (+) Graylog (+)	
TOOLS	Mesos	Docker (1) Kubernetes (1)	Ansible Cucumber Gradle etcd (1) Prometheus (+)	
LANGUAGES & FRAMEWORKS		SASS Dropwizard Lua Storm Go Flink Kotlin (+)	JRuby Play Framework Rails Foundation Framework (+) Angular (+)	GWT XML Struts2 SpringMVC
METHODS & PRACTICES		DevOps Cloud Native Microservices	REST Continuous Integration Scrum Database Migrations	WS- Manual Testing Monolithic Systems

TABLE 2 - Technology Radar 2017 summary

This edition introduced new technologies like Kotlin and the move to “Experiment” for Docker and Kubernetes. But the biggest change was motivated by the adoption of the Altice Labs’ Foundation Framework, a framework for lowering the implementation costs of microservices architectures, based on Cloud Native Foundation components, like Elasticsearch, Graylog, Prometheus and Etcd.

There’s also a very important and separated special section focused on big data technologies, being an outcome of the experience and work done over recent years in Altice Labs.

## Experiment

As described above, “Experiment” recommendation applies to a mature technology that can be experimented in production systems, but should be used only in projects with low or moderated risk. The following are some practical examples addressed in Technology Radar current edition.

### Docker

Docker [2], usually referred as the “new virtualisation”, is having a major impact in cloud and/or hybrid environments, as well as in distributed systems. No cloud platform will be seriously considered unless it supports Docker containers [3].

In a nutshell, Docker provides something similar to virtualisation, but in a more lightweight way, allowing multiple “virtual machines” running on the same physical machine without incurring in the virtualisation costs (hypervisor function).

Docker is seen as a critical component for the success of microservices-based architectures, and also as a possible alternative for more efficient use of bare-metal, supporting multiple systems instead of using virtual machines.

### Kubernetes

Kubernetes [4] is an open source container cluster manager/orchestrator, originally designed at Google and now maintained by Cloud Native Computing Foundation. It allows to deploy, scale and operate applications that are deployed in containers (Docker) across a cluster of hosts. We believe that Kubernetes will be the standard approach to interact with clusters of containers, even when they run on top of Mesos.

### Go

Go [5] is a programming language developed at Google in 2007. Despite the fact of being a relatively recent programming language, it has gained a lot of momentum and for good reasons: it’s a statically typed language, compiles to native code, and has a C-like syntax; it’s a very efficient language (close to C performance) providing abstractions that simplify the process of developing concurrent computing routines through “goroutines” and channels.

Nowadays there is a lot of production-ready software written in Go, such as Docker, InfluxDB, NSQ, Etcd, Consul, etc. Currently, Altice Labs still has few expertise and experience on this language, so it’s not advised to use it in critical systems. However, we foresee that its usage will increase, and we anticipate Go as an alternative to be used in systems where we currently use C/ C++ or Java.

### Kotlin

Kotlin [6] is a statically typed programming language that is 100% interoperable with Java and Android and can target the JVM, Android, JavaScript and native.

Kotlin has both object-oriented and functional constructs. You can use it in either object oriented and functional programming styles, or mix elements of the two. With first-class support for features such as higher-order functions, function types and lambdas, Kotlin is a great

choice if you're doing or exploring functional programming.

If you are coming from Java, Kotlin has some great advantages. Kotlin is more concise: rough estimates indicate approximately a 40% reduction in the number of lines of code. It's also more type-safe, e.g. gives support for non-nullable types, making applications less prone to NullPointerException (NPE) kind of exceptions. Other features including smart casting, higher-order functions, extension functions and lambdas with receivers, provide the ability to write expressive code as well as facilitate the creation of domain specific language (DSL). There's also an automated Java-to-Kotlin converter built into the integrated development environment (IDE) that simplifies migration of existing code.

## Flink

Flink [7] is a streaming dataflow engine that provides data distribution, communication and fault tolerance for distributed computing over data streams. Flink, like Storm, is tailored for stream processing, but also handles batch processing. Additionally, it includes several libraries for domain-specific use cases, like complex event processing (CEP), machine learning (ML), and graph processing.

It's the richest of all the big data streaming/batch frameworks we analysed so far, supporting features like delivery guarantees, fault tolerance, out-of-order processing, windowing time and count based, back pressure, check pointing and low latency, just to name a few, and with very good performance.

## Microservices

Microservices is a way of breaking large software projects into smaller, independent and loosely coupled modules. Individual modules are responsible for highly defined and discrete tasks and communicate with other modules through simple and universally accessible APIs [8].

There are many advantages and some drawbacks on using this approach, but we think that advantages largely outnumber drawbacks. Some of the advantages are: microservices are smaller and easier to understand, can use a different technology stack and are easier to deploy and replace; a system implemented by a set of microservices is easier to scale and more resilient, because if one of the microservices fails, in most cases it will affect only a small part of the system. However, system architects must learn how to design a system using microservices and how to deal with eventual consistency issues, distributed systems and more complex operations.

## Adopt

The best technologies that we aren't using yet get the "Adopt" recommendation in the Technology Radar. They provide good benefits in comparison with alternative technologies currently in use. Therefore, all references with this classification can and should be used in all sort of systems, even mission critical, in a generic way. Next we present some examples.

## PostgreSQL

PostgreSQL [9] is currently the most advanced open-source database. It's also the one that shares more similarities with Oracle database in terms of features and usage (SQL dialect). The use of this database can significantly reduce the total cost of acquisition (TCA) and ownership (TCO), without compromising functionalities or performance.

PostgreSQL has excellent references in terms of reliability and robustness, without constraints that other open source RDBMS (ex: MySQL) present. When Microsoft acquired Skype, PostgreSQL was supporting all Skype clients (around 1000 millions) [10]. It's also the type of database that supports name servers from Internet top domain ".org".

## Kafka

Kafka is a distributed messaging system developed at LinkedIn, being now an Apache project. Horizontal scalability, durability, replication and ability to handle multiple consumers without sacrificing performance, are the key design drivers, and they are delivered in a natural out-of-the-box way. Data ingestion is a common Apache Kafka use case, acting as a data-hub and buffer for other applications, such as big data platforms.

Despite being recent, it is a mature project currently used in production systems at LinkedIn for some time. LinkedIn announced in September 2015 that it was able to process about 1.1 trillion messages per day [11].

## Ansible

Ansible [12] is a tool that allows IT tasks automation, similarly to Puppet or Chef. However, what distinguishes Ansible from their counterparts is its simplicity and needlessness of installing agents in every managed host (i.e., it implements a push approach, unlike Puppet or Chef, which support a pull approach), as it uses secure shell (SSH) remote execution facility to perform tasks on managed hosts.

A comparative analysis performed by the Technology Coordination team, considered Ansible as the best choice, and we believe that its usage in systems' installation, updating or configuration processes, will yield to a return as high as the system's complexity.

By using Ansible, it's possible to manage dependencies in a distributed way, ensuring system consistency. In other words, Ansible can do for your system what RPM can do for a package. In some systems at Altice Labs, it was possible to reduce the installation and configuration time from days to just few minutes, eliminating errors introduced by manual execution and creating repeatable procedures that can be used whenever required.

## Cucumber

Test automation, when properly done and in a systematic manner, is probably the most effective way to ensure the quality of a system. However tests can also be a source of problems, due for instance to mismatches between what's supposed to be tested and what is really tested. One of the main reasons for mismatches is the completely informal and ad-hoc way how most tests are described.

Cucumber [13] introduces a simple language (Gherkin) that follows some formalities which can be used by non technical people to describe system features in almost plain english sentences. Cucumber parses those sentences in order to execute tests as described, significantly reducing the mismatch between test descriptions and what's really tested.

Cucumber adoption has grown significantly at Altice Labs and it's the recommended way for test automation usage.

## Prometheus

Prometheus [14] is a new generation open source monitoring and alerting solution, originally built at SoundCloud. It's now a Cloud Native Computing Foundation project, capable of dealing with dynamic and elastic cloud infrastructure. Prometheus has a built-in internal time series database, supporting a multi-dimensional data model and a flexible query language, providing great visualisation interface and integration with Grafana. It's very easy to operate and supports alarms based on its query language and a specific module to send notifications or even silencing. Out of the box, it offers many integrations with most common systems, and it's very easy to integrate with, as it supports many client libraries in several programming languages.

## Foundation Framework (by Altice Labs)

Moving Altice Labs' systems into microservices architectures supported by containers in cloud or bare-metal infrastructure, introduces new challenges that come from the dynamic nature of the underlying infrastructure. Applications must now deal with dynamic configuration, service discovery, log centralisation and distributed tracing, and new forms of monitoring and alerting.

The Foundation Framework is a selection of technologies that address these issues, allowing the development teams to focus on product development and still have a similar approach dealing with the same issues across Altice Labs' products.

## Database migrations (by Altice Labs)

During the product lifecycle there will be probably several product updates. Almost certainly your system will have a database with important information that you want to make sure it's safely kept whenever the system is upgraded. A backup may not be enough as the new version may change the database schema (for example splitting a table, changing few columns, etc). Thus it's important that the database schema evolves with the software version in a synchronised and controlled way, ensuring that data inside the database is not lost but is correctly "migrated" to the new schema.

Doing data migration by hand is tedious and error prone and you risk losing your most valuable asset: your data! That's why it's very important to use a tool to handle database migrations in a deterministic and proven way, which you can test and validate before doing it in production.

Altice Labs has an in-house tool for database migrations called "Migrate" with unique features not available in similar tools.

## Big data technologies

"Data is the new oil", it is said, and operators are sitting on top of terabytes of data stored in silos and scattered across the organisations. In order to exploit the full potential of this stored data, they must have the right solution to harness the volume and variety of data, as well as the speed of data change, to help them to correlate and process in order to transform that data into actionable insights. This is not possible without big data and advanced analytics.

In the forthcoming future, these platforms will be even more crucial to the business, as they will be the supporting foundations for more advanced approaches like machine learning and artificial intelligence. Failing to keep up with these technologies will have major negative consequences on business, as the competition moves on. However, implementing a successful solution is not an easy task. These technologies are recent, and there is a shortage of technicians with these skills. Traditional vendors, especially those with business intelligence background, rushed to launch products in this area trying to catch most of the initial investments and take advantage of the lack of knowledge of the customers. However some of these products fall short and are not up to the challenge.

Altice Labs has a team working on big data technologies since 2012. This team has been experimenting and validating the technologies to be used by Altice Labs' products and is also managing the big data platform from Altice Portugal since 2016, and has been leading, in partnership with Altice USA, the definition of the standard technology stack for the Altice Group Common Data Model initiative.

The standard technology stack definition is an undergoing task at the time of this writing. The next sections provide a glimpse of the work being conducted, which is aligned with the Technology Radar methodology. In **Table 3** you can see the corresponding evaluation table.

	Explore (1)	Experiment (2)	Adopt (A)
Data ingestion - batch	Apache NiFi, StreamSets	Sqoop, Flume	Python, Shell
Data ingestion - streaming			Kafka, REST API
Data ingestion - shared filesystem			GlusterFS
Hadoop base distribution			Cloudera
Data layer - incoming stage			HDFS, Kudu
Data layer - gold and other stages			HDFS, Impala
Data processing - batch			PIG, Python
Data processing - streaming	Spark	Storm, Flink	
Working in-memory storage			Redis
Data warehouse		Apache Drill, CitusDB	PostgreSQL
Data science playground - storage		Apache Drill, CitusDB	HDFS, PostgreSQL, Impala
Data access APIs		Apache Drill, REST API	
Cluster management and monitoring			Cloudera Manager
Process orchestration		Airflow	
Data governance, lineage and audit	Cloudera Navigator, Apache Atlas		

TABLE 3 – Technology Radar: big data technologies

## Data ingestion

Data ingestion in batch mode is probably the most common way of ingesting data into big data platforms, by periodic imports from external databases or from files containing data records, network events, logs, etc. Flume and Sqoop are standard tools already included in the Hadoop distributions for this purpose. Flume is more oriented to import files and logs, and Sqoop for moving data between relational databases and Hadoop. Both are mature, reliable and easy to use, so you should experiment them before trying something else.

However, in some cases where we must comply with strict high performance requirements, or specific workflows, we may need to implement

specific data import scripts and Python [15] and Shell can be your tools of choice because both play well in the Hadoop environment, and they allow a much faster implementation than using programming languages like Java.

Another way of ingesting data is doing it event by event in streaming mode (ex: web pages clicks, network alarms, etc). The tool we recommend for handling these events is Kafka. It's high performant, reliable, secure and distributed, and has been proving in very high demanding environments. For situations where using the Kafka API may not be so practical, a REST interface on top of Kafka is the way to go.

Data ingestion should not go directly into Hadoop. The first stage of data ingestion is



to bring data from the external system into a staging area where it can be prepared to move into Hadoop. A shared filesystem is required to guarantee that no data is lost in case of failure of the server processing data ingestion, so that the job can continue on another server. We found that GlusterFS [16] is a good solution for this purpose, as it removes the dependency from external storage (ex: NFS).

## Data storage

At the core of the big data platform is the Hadoop distribution, containing most of the components needed. The two major Hadoop distributions are Cloudera and Hortonworks and they have more in common than differences. We believe that Cloudera [17] has an edge over Hortonworks because it includes Apache Impala, which is presently the fastest SQL available on Hadoop, and Apache Kudu that solves a big problem on data ingestion when dealing with late information (data that comes late and forces an update on already written files).

The incoming stage is where the data is stored after the ingestion. For specific data sources, some data may arrive out of order or later, with delays of hours or even days. As Hadoop distributed file system (HDFS) doesn't allow for updates, it is more efficient to store the incoming area for these data sources into Kudu as it allows updates. Kudu is an open source complement to HDFS and Apache HBase. It is designed to complete Hadoop's storage layer, enabling fast analytics on fast data.

Normal bounded data should be stored in HDFS. The HDFS storage is the most cost effective way of storing big volumes of data for long time and should be the default option. However, you may need to speed up access to some data that is used more frequently or to provide an SQL on Hadoop interface to that data. For that purpose, the most efficient option is to use Impala. Impala is a SQL analytic engine and fully integrated analytic database architected specifically to leverage the flexibility and scalability of Hadoop. Impala brings scalable parallel database

technology to Hadoop, enabling users to issue low-latency SQL queries to data stored in HDFS and HBase, without requiring data movement or transformation.

## Data processing

The default processing model for big data in Hadoop is still MapReduce and the easiest way of processing large data sets using a high-level language fully integrated with Hadoop and taking advantage of the resources in the cluster is by using Apache Pig. The leading characteristic of Pig programs is that their structure is amenable to substantial parallelisation, which enables the handling of very large data sets. For small data sets Python is a very good choice. It is a powerful easy to use language, with a rich set of libraries for data manipulation.

Processing data in stream is gaining popularity as it produces results much faster than by using traditional MapReduce, at the cost of big CPU and memory usage. It's more complicated to write code for stream processing data than for traditional MapReduce by using Pig, but if you really need faster results, there are frameworks that may help you. Apache Storm is a generic distributed real-time computation system that is mature, robust and fast, capable of processing over a million tuples per second per node. It's scalable, fault-tolerant, guarantees that your data will be processed, and it's easy to set up and operate. Apache Flink is also an open-source stream processing framework for distributed, high-performing, always-available and accurate data streaming applications. It's not as generic as Storm so both have different optimal use cases. A special mention to Apache Spark as it is probably the most hyped and well known streaming framework. However, in our tests we've found it to be less performant and less robust than Storm or Flink. But the real deal breaker is the dependency on Scala and the problems caused by the incompatibility of binary libraries across Scala versions. At some point you'll want to use some module that requires a different and incompatible set of libraries, all because of the Scala version

used. However, as it's very popular and many libraries and frameworks depend on Spark, its use may be unavoidable, so plan carefully and use it with extreme caution.

A common pattern when processing data is that you get some information from a record that will be needed to correlate or enrich other record that will be found later. In order to take advantage of Hadoop, you should build a distributed and easily parallelised process so you could use many cores and many servers on your cluster. Thus, you need a common place to store temporary data that must be very efficient and performant, and the in-memory Redis [18] database may just be what you need. The fact that it supports auto-expire keys after some time also simplifies the management of data stored and the programming model.

## Data warehousing, science playground and data access API

Data warehousing is still the best solution for some use cases, namely when there's simply the need to store aggregated information and KPI, to allow fast access using more traditional technologies which are quite efficient for interactive use by typical analytics solutions. It complements big data nicely. However, if the lion's share of data is in Hadoop, data warehousing can be much simpler than it used to be. As stated above, PostgreSQL is the most advanced open source database, very robust and performant, and may be a very good option that can save you a lot of money. CitusDB [19] can be installed on top of PostgreSQL to scale out horizontally into a massively parallel processing (MPP) solution. Apache Drill is an excellent choice to provide a common way to access your data, which can be used by tools like Tableau, Qlick View and others, as they see your data through a JDBC interface even if data is in HDFS or NoSQL databases, unstructured or semi-structured. You can even join data from different data sources as if they were tables in the same database.

The Data Science Playground is an environment containing a subset of the data, be it aggregated or raw data, where data scientists or advanced business users can safely explore and experiment new models in order to discover new information. This playground environment may be implemented on relational databases using PostgreSQL and CitusDB and also on HDFS and Impala for holding the raw data.

There are many situations where data is processed in Hadoop in order to be used by external platforms. The results of a recommendation engine, such as client segmentations, predictions, etc, need to be accessible by external applications that will use that information to accomplish their purposes. To give a simple enough access to that data in a controlled and secure way, a REST interface may be the best approach. Drill can also be used to provide a SQL like interface into that data or as a possible REST interface.

## Cluster management, monitoring and process orchestration

The Cloudera Manager included in the Cloudera distribution is a very good and complete tool providing a view of the cluster services installed and how healthy they are. It's also used to manage the installed components and their respective configurations.

The number of jobs and processes running in a cluster is expected to grow over time and it's important to understand the relations and dependencies between all those jobs. If an ingestion job fails, it's useless to run other jobs that will need the data that would come from the failed ingestion job. This could lead to a domino effect of failed jobs. Therefore, it's important to have a tool to do all the orchestration. Airflow [20] is such a tool, originally from AirBnB, which allows programmatically authoring, schedule and monitoring workflows. You can use Airflow to author workflows as directed acyclic graphs (DAG)

of tasks. The Airflow scheduler executes your tasks on an array of workers while following the specified dependencies. Rich command line utilities make performing complex surgeries on DAG a snap. The rich user interface turns visualising pipelines running in production, monitoring progress and troubleshooting issues when needed into easy tasks. Workflows are defined as code and as such they become more maintainable, versionable, testable and allowing for collaboration.

## Data governance, lineage and audit

It's very important to have governance of the data inside the big data platform and a tool to assist that governance, supporting features like data classification, lineage, auditing, searching and to define access policies on that data. We haven't yet concluded a detailed analysis on what should be the best tool for this job, so we can only recommend you to experiment with Apache Atlas or Cloudera Navigator, the later if you have Cloudera Enterprise version.

## Wrap-up

This article describes how Altice Labs' Technology Radar is organised and how it is used to disseminate key insights on software technology to technical managers, system architects,

researchers and all those who are dealing with software technology, thus positively influencing exploratory innovation projects, system design processes, planned evolution of systems, product roadmaps and the overall quality of the organisation product portfolio.

This approach, methodically fulfilled by a Technology Coordination team, doesn't aim to maintain a reference stack, but rather to periodically issue a set of recommendations about changes on the way we, at Altice Labs, use "Platforms", "Tools", "Languages and Frameworks" and "Methodologies and Practices" in our system development process and product portfolio management. For each technology we propose "Explore", "Experiment", "Adopt" or "Avoid" advice, along with comments and guidance, with some examples being presented.

Big data is a special area for which Altice Labs has been successfully applying the Technology Radar approach, experimenting and validating technologies since 2012. We detailed a bit of the ongoing work on establishing the Altice standard technological stack.


Based on the achievements so far, this approach is proving to be a valuable tool for monitoring the tech community and to anticipate technological and architectural trends with stronger impact on Altice Labs' activity. Technology Coordination will therefore continue to embrace the needed tasks to regularly update and improve the Technology Radar, for the win!

## References

- [1] N. Ford, "Build Your Own Technology Radar," ThoughtWorks, 2015. [Online]. Available: <https://medium.com/featured-insights/build-your-own-technology-radar-27ef3c2e0bc2>.
- [2] Docker, "Docker," Docker Inc., 2018. [Online]. Available: <https://www.docker.com/>.
- [3] DataDog, "8 Surprising Facts About Real Docker Adoption," 2018. [Online]. Available: <https://www.datadoghq.com/docker-adoption/>.

- [4] Kubernetes Authors, "Kubernetes," 2018. [Online]. Available: <https://kubernetes.io/>.
- [5] The Go Project, "The Go Programming Language," 2018. [Online]. Available: <https://golang.org/>.
- [6] The Kotlin Community, "Kotlin," 2018. [Online]. Available: <https://kotlinlang.org/>.
- [7] The Apache Software Foundation, "Apache Flink - Stateful Computations over Data Streams," 2018. [Online]. Available: <https://flink.apache.org/>.
- [8] C. Richardson, "Microservices Architecture," Kong, 2018. [Online]. Available: <https://microservices.io/>.
- [9] The PostgreSQL Global Development Group, "PostgreSQL: The World's Most Advanced Open Source Relational Database," 2018. [Online]. Available: <https://www.postgresql.org/>.
- [10] T. Hoff, "Skype Plans For PostgreSQL To Scale To 1 Billion Users," 2008. [Online]. Available: <http://highscalability.com/skype-plans-postgresql-scale-1-billion-users>.
- [11] N. Narkhede, "Apache Kafka Hits 1.1 Trillion Messages Per Day - Joins the 4 Comma Club," Confluent, 2015. [Online]. Available: <https://www.confluent.io/blog/apache-kafka-hits-1-1-trillion-messages-per-day-joins-the-4-comma-club/>.
- [12] Red Hat Ansible, "Ansible," Red Hat, Inc. [Online]. Available: <https://www.ansible.com/>.
- [13] The Cucumber Community, "Cucumber," 2018. [Online]. Available: <https://cucumber.io/>.
- [14] Prometheus Authors, "Prometheus," 2018. [Online]. Available: <https://prometheus.io/>.
- [15] Python Software Foundation, "Python," 2018. [Online]. Available: <https://www.python.org/>.
- [16] The Gluster Community, "Gluster," Red Hat, Inc, 2018. [Online]. Available: <https://www.gluster.org/>.
- [17] I. Cloudera, "Cloudera," 2018. [Online]. Available: <https://www.cloudera.com/>.
- [18] Redis Labs, "Redis," 2018. [Online]. Available: <https://redis.io/>.
- [19] I. Citus Data, "Citus Data," 2018. [Online]. Available: <https://www.citusdata.com/>.
- [20] The Apache Software Foundation, "Apache Airflow," 2018. [Online]. Available: <https://airflow.apache.org/>.





## 09

# Web Components, a compatibility solution in a cross framework environment

**Rui Filipe Martins, Altice Labs**

[rui-f-martins@alticelabs.com](mailto:rui-f-martins@alticelabs.com)

**Carlos Manuel Figueiredo, Altice Labs**

[carlos-m-figueiredo@alticelabs.com](mailto:carlos-m-figueiredo@alticelabs.com)

**Nuno Miguel Simões dos Santos, Altice Labs**

[n-simoes-santos@alticelabs.com](mailto:n-simoes-santos@alticelabs.com)

The degree of complexity associated with web applications development has been growing in the last years. At the same time, frameworks focused on simplifying and speeding up the development process are being released at a fast pace, making it difficult to develop future-proof applications. Web Components is a technology based on standards defined by the W3C that allows for the creation of components with high potential of reuse and more immune to constant changes in web applications.

---

## Keywords

Web Components; Web development; Interoperability; Cross-framework

## Introduction

The range of web development frameworks focusing on the front-end has been growing rapidly in the last years. Nowadays this range is quite extended and it is expected that it keeps on growing, with tools ever more evolved who serve the needs of developers. Those aim to ease and speed up the process of development as well as allowing the creation of products with professional rigour.

However, in a cross-framework context, compatibility issues on the development and implementation of visual components arise. Each framework has a different approach to solve the same problem, so there is not a pattern to develop components that can be reused. The ideal scenario would be the development of components that could be reused independently of the selected framework, despite the internal processes imposed by each, in terms of maintenance or product integration.

At the corporate level, where the technological environment demands having to manage products based on different technologies, this issue matters the most. An interoperable solution that unifies the cross-framework workflow would reduce the costs and at the same time quicken the process of development, resulting in an organised and efficient development environment, improving both the developer and the customer experience.

## Corporate innovation

In a business environment, the constant replacement of frameworks is not feasible at all, since product lifecycles are usually long and the technology used in a product tends to be maintained for some years in order to justify the initial investment. The overhead of technological change in a product is always great, and it is not so strange to find products in this environment still with outdated technologies in the face of the latest trends. The very culture of the developers

and some reluctance to change are factors that lead one to prefer to use a technology that is already well known and with which they are productive rather than risking change.

In a context where products developed by different teams use a variety of technologies but should nevertheless provide the end user with the same experience, it is essential to promote such sharing of experience with the use of a common framework. The framework created by Altice Labs' Usability and User Experience team is framework user experience and interfaces (FUXI). It allows maintaining the several components that serve multiple technological universes, as many as the products that the company has. It is, therefore, a complex task because it requires the creation of the same component in different technologies and at the same time entails high maintenance costs.

The ideal scenario would be the creation of components in a single technology so that each of the frameworks could use them, independently of their internal processes, both for maintenance and for integration in products.

Web Components represent a technology where it is possible to create components compatible with development frameworks such as Angular or React and are increasingly supported by browsers.

## Web components technology

Current web development standards require, among other things, that an application be able to run the same way in different browsers, that it adapts to different screen sizes, be prepared for mobile phones, watches and even televisions [1]. However, the base development technologies (HTML, CSS and JS) are not up to the challenge, so the community has come over time to develop solutions that ease this process. The result is the creation of an ecosystem in constant



expansion, with diverse development frameworks being released. These are aimed at building applications with professional rigour, with a wide range of features, while facilitating the process for the programmer.

Due to the nature of the web technologies on which these frameworks are supported (HTML, CSS and JS) each has its own proprietary development environment with a set of tools available. However, this strategy leads to a blockage regarding interoperability in a cross-framework development process, so choosing a framework is the first major technological decision when starting a project. This is due to the temporal costs of replacing a framework being high because it would mean start from scratch. Another important factor is the high volatility of the ecosystem, which makes the choice of a framework a high-risk investment because currently, the pace at which technologies become outdated is alarming. This is a matter of major importance in companies where decisions of this nature imply longer investments since the associated costs do not allow frequent technological changes. Finally, the last obstacle is the lack of interoperability, in this case, between frameworks. As mentioned above, each framework requires exclusivity on the part of the programmer, that is, the development process is unique. This is also true for other important aspects such as the reuse of visual elements.

In the current context of web development, it is expected to be possible the creation of a versatile component that is capable of being reused and manipulated according to the project needs. This is possible when it is only a framework because it will have an architecture prepared to recycle its components, but not when it comes to a cross-framework process. Therefore, in addition to the exclusivity required by these, creating universally reusable components through these technologies is practically impossible [1]. However, in parallel with this ecosystem, W3C, an international community responsible for developing standards to ensure long-term web growth, has since 2012 published standards that introduce Web Components.

Web Components are a set of technologies that allow the creation of reusable components, encapsulated from the remaining code where they are injected [2]. They follow the same modular concept of the proprietary components of the frameworks as Angular and React, but unlike these, they are capable of being integrated into the various development frameworks due to their interoperability. The various advantages over these components will be discussed below.

## Interoperability

The biggest selling point of this technology is that it is interoperable between frameworks so it can be used in different contexts. In a development context where multiple teams work on different projects using multiple technologies simultaneously, it is of the utmost importance to create components that are reusable, minimising costs by reducing development time.

## Easy maintenance

The objective of these components is to be simple, focused on only one task and optimised for it. In addition, Web Components are based on pure web technologies. Adding the two points yields components with reduced maintenance needs.

## Encapsulation

One of the technologies associated with Web Components is the Shadow DOM. This makes encapsulation of components possible when integrated into another environment. This way it is possible to maintain the structure of it intact, as well as the associated styles and all the behaviour, avoiding conflicts.

## Reliability

Since it is only needed to create the components once, they will be optimised and will meet the requirements of the user. Reusing the same components previously tested in different environments, the level of confidence in them will

be high, as opposed to a development method that implies their constant development.

Because Web Components standards are relatively recent, native support by browsers is not totally covered, with Edge still in development and Firefox only adding full support in late 2018. However, it is possible to get the functionality of Web Components through polyfills. Polyfills are excerpts of code that provide the technology needed to simulate native browser functionality but are not implemented, thus addressing the lack of native support [3]. The current support state in the main browsers is shown in **Figure 1**.

## Component fx-tabs

Given the wide range of possibilities for implementing this technology, an in-depth analysis was carried out to expand the knowledge on the different approaches. In this way, it was possible to test the ones that showed to be more relevant to the case study. The study of Web Components was pertinent in that they are based on four independent standards, the use

of them all is not mandatory, and that they are still in development. In this way, the application context of the Web Components in Altice Labs, as well as the limitations presented by them, would determine its implementation. It was intended to explore the advantages and disadvantages of using each standard as well as its limitations, such as browser support and their combination. Thus, a detailed analysis of each standard was conducted in order to collect information relevant to the development of the prototype.

Having said this, and taking into account the scope of the investigation and the purpose of Web Components, a prototype was developed with the purpose of testing the ability of this technology to serve the needs of the company.

To better test the technology, the prototype was based on a custom component created by the Usability and User Experience team called “fx-tabs” that is in use across the various products developed by Altice Labs.

Browser support	CHROME	OPERA	SAFARI	FIREFOX	EDGE
HTML TEMPLATES	✓ STABLE	✓ STABLE	✓ STABLE	✓ STABLE	✓ STABLE
CUSTOM ELEMENTS	✓ STABLE	✓ STABLE	✓ STABLE	✓ STABLE	✓ POLYFILL ✗ DEVELOPING
SHADOW DOM	✓ STABLE	✓ STABLE	✓ STABLE	✓ STABLE	✓ POLYFILL ✗ DEVELOPING
ES MODULES	✓ STABLE	✓ STABLE	✓ STABLE	✓ STABLE	✓ STABLE

**FIGURE 1** – Browser support for Web Components standards [3]

## In-depth

Web Components technology, while thought of as a whole, was designed to be flexible. This is achieved through the independence of the standards on which it is based, which have the capability to fulfil their purposes without the intervention of other standards. Thus they allow the conjugation between them in the most pertinent way. Therefore, the analysis of the technology only makes sense if the standards are treated independently.

One of these standards is **Custom Elements** that allows the creation of tags in order to simplify HTML structures by encapsulating more complex structures into simple, readable and easier to handle components. So the biggest advantage that this standard presents us is to be interpreted as an HTML tag by the browser. This achieves a high level of interoperability that promotes productivity and optimisation of resources.

Because this standard has not yet been able to gather a strong consensus on the part of the whole community, this lack of definition is impacting the adoption by some browsers with Edge still not supporting it. However, there are polyfills that make this standard compatible in all browsers.

The **HTML Imports** standard allows the import and use of HTML files inside other files of the same type, thus enabling the development of Web Components based on HTML. Of the four developed standards, this is the one that has less adoption by the browsers, because it represents an approach that does not meet the current trends of web development, that currently focus on sending to the JavaScript actions like the import of modules or the declaration of HTML.

One of the advantages that this standard brings is the control of imports, managing to solve the duplication of equal resources even if they are imported several times or from different locations. However, today's modern browsers are also able to manage imports and optimise them in a similar,

automatic way. Another relevant advantage is the readability of the code in HTML files. The structure of an HTML-based Web Component makes the readability of the code improved. Unlike when it is developed using JS, the HTML syntax in a JS file must be contained in a string that makes the highlighting of the element tags indistinguishable, affecting readability. Because it has not been adopted by the major browsers and the community in general, the HTML Imports standard has been deprecated in favour of **ES Modules**, which does the same function but over JS files instead of HTML and is established in the web.

Contrary to the previous standard, **HTML Templates** has high support from browsers. This is also due to the consensus and stability of the same. Despite this, Internet Explorer 11 stands out as the only modern browser to not support this standard. In fact, Internet Explorer 11 does not support any of the four pillars of Web Components, limiting the development of components based on them to this browser. It is a standard that allows the declaration of content that must be rendered only when it is cloned or instantiated via JS, rather than rendered "naturally" when it is loaded. We can point out as the main advantage of this standard the important role it plays, as it allows the declaration of portions of code for future use. It encourages the development of components with clear and simple purposes so that they can be reused. In this way, it is in line with the new practices of modularisation of the web.

Finally, the last standard of the Web Components is the **Shadow DOM**. This standard allows the component to be encapsulated, so it does not let CSS or JS from the main page affect the component and vice versa. Thus it is possible to guarantee that the behaviour of the component and its visual appearance is the same in different environments where it is incorporated. As has been seen in previous standards, to fix this limitation of browser support, there are polyfills that somehow mimic their behaviour. However, in this case, the polyfills are not able to fully cover the gaps.

The most important advantage is the massive reduction of the risk of collision of styles between the component and the environment where it is integrated. In this way, it is not necessary to resort to elaborate nomenclature techniques or to apply styles with an extra rigour in order not to affect external elements, as this is not possible through this standard. Moreover, by treating the component in isolation, the integrity of the component is guaranteed.

Another feature that raises the value of this standard is the ability for the programmer to control what may or may not be styled externally in the component. This requires that the properties of the component that can be externally styled through CSS variables must be made explicit. Thus, in a business environment where control over the manipulation of resources is of utmost importance, this functionality is valued. In addition, in a company, keeping the same look & feel throughout their products is extremely important for users to feel comfortable and confident when interacting with them, so the components that make them must ensure that their identity is maintained.

The current state of the Shadow DOM is not consolidated, so it is expected that this standard changes over time as it has done before. Some of the aspects of the Shadow DOM that are unresolved are the lack of support for some rules and the lack of consensus as to the styling of components. In addition, the available polyfills are currently under development, without support for key features such as encapsulation of externally inserted style documents, or dynamically inserted styles using JS, so they are not yet a viable alternative.

## How to

To develop web components, it is necessary to follow a set of rules and good practices that contribute to a better standardisation of these components. It is necessary to take into account a set of characteristics that they must have in

order to have equivalent behaviours to the native components of the HTML and as such can be used in a transparent way by web frameworks (e.g. Angular).

These rules and good practices will dictate, for example, how a component should be instantiated or how any interaction should be performed in or out of the component. It is with this premise that this section arises, to clearly describe a set of rules and good practices that drive the development and lifecycle of this nature of components.

The decision to create a web component should be based essentially on its potential of reuse, which may be at the organisation level or simply at the application level.

Since Web Components technology is a set of independent standards, each with a different level of development and adoption by the community, it is relevant to make clear the best development practices. The advantages of using Custom Elements as an interoperability solution between development frameworks are obvious. Support from browsers is not the best. However, polyfills cover this flaw without significant costs associated, allowing a component to be created once and reused in a variety of contexts.

The use of Shadow DOM as an option to encapsulate JS and CSS is only feasible if it is possible to control the browsers that users will use because their support is low and in this case the polyfills currently require a restructuring of both the component and the document where it would be integrated, in browsers that do not support it natively. In addition, the polyfills do not provide the key features of the above mentioned Shadow DOM.

Regarding the communication of the components with the framework where they are inserted, a set of good practices should be followed that avoid the occurrence of synchronisation failures between them. To assign states to a Web Component, the best practice is to use properties rather than attributes because unlike properties, attributes only reflect their initial state even if they are changed, and can mislead the external environment. The

use of attributes is convenient when the initial state of a component is important to retain. As for how to transmit information from the Web Component to the outside, this must be accomplished through the propagation of events. In this way, it is possible to prepare the environment to perform something when listening to the desired event [4].

Finally, the use of an external library such as Polymer aims to minimise the time spent in creating the boilerplate of the component, as well as in its preparation for production. It natively incorporates the Shadow DOM into the component, makes the structure of the DOM more organised/easy to interpret, and the command-line tool becomes very useful in the development phase. On the other hand, the extra load on the browser and the significant changes between versions are important aspects to take into account. Therefore, the use of the library depends on the personal preference and the nature of the component itself, being recommended a native implementation for basic components.

## Wrapping up and next steps

This article is based on a master's thesis in Communication and Multimedia whose objective

was to investigate the state of the art of using Web Components in the cross-framework problem, to develop a functional prototype that would serve as proof of concept as well as a test to evaluate the capabilities of a Web Components based solution.

The diversity of web development frameworks is increasing, as well as the need to optimise software development processes to be more a reuse and share of transversal artefacts when inserted in an enterprise ecosystem that already shares a common look and feel across its products. The Web Components come to give mechanisms that allow focussing the development of web component in market standards, compatible with the different existing frameworks and whose level of acceptance by the community has been growing gradually and that is expected in the future to be a reality with the full support of all existing browsers.

We believe in the potential of Web Components to mitigate the difficulties experienced by developers in multi-technology scenarios and using different development frameworks, even if to this became a reality there should be a greater evolution of the Web Components technologies and consequent changes on the development process within the company teams.

---

## References

- [1] T. Savage, "Componentizing the Web," Commun. ACM, pp. 1–20, 2015.
- [2] A. L. Martínez-Ortiz, D. Lizcano, M. Ortega, L. Ruiz, and G. López, "A quality model for web components," ACM Int. Conf. Proceeding Ser., vol. Part F1263, pp. 430–432, 2016.
- [3] "webcomponents.org," 2018. [Online]. Available: <https://www.webcomponents.org/>.
- [4] R. Dodson, "Custom Elements That Work Anywhere," 2016. [Online]. Available: <http://robdodson.me/interoperable-custom-elements/>.



# 10

## Personalisation of smart cities powered by UX methodologies

**Cátia Raquel Oliveira, Altice Labs**

[catia-r-oliveira@alticelabs.com](mailto:catia-r-oliveira@alticelabs.com)

**Inês Freitas Oliveira, Altice Labs**

[ines-f-oliveira@alticelabs.com](mailto:ines-f-oliveira@alticelabs.com)

**Filipe Cabral Pinto, Altice Labs**

[filipe-c-pinto@alticelabs.com](mailto:filipe-c-pinto@alticelabs.com)

**Isabel Borges, Altice Labs**

[Isabel@alticelabs.com](mailto:Isabel@alticelabs.com)

If one thinks the city as a product or service, the entire experience of using it and living in it has to be measured. As a product or service, also the city is composed of a set of different layers which evolve independently but interconnected. Evaluating customer/citizen experience and engagement is the only sustainable way to design a city that is smart for itself and for the ones that live in it. Therefore, this article will focus on the importance of understanding the journey experience of a citizen around the city and the best way of engagement.

---

### Keywords

User experience; Citizen; Citizen-oriented; User-centric

## The city and the citizens

A city does not exist without people, whether residents or visitors. People represent, first and last, the entity that the city must preserve and serve [1].

The relocation of citizens within big cities raises some questions: are the cities ready for receiving and keeping the citizen? Do the cities provide the right conditions to create desire on the citizen as part of the community and as an individual element? How can the city give the best experience for its users? Having these questions in mind, two types of cities may be identified:

- Those who inspire and are an object of desire, characterised by high affluence. These cities require a fast effort to adapt to the behaviours and needs of citizens and should adopt a proactive adaptation strategy;
- Those who attract less and that need to be rediscovered and rethought to capture more interest from citizens, offering better and differentiating means and conditions to solve their needs and pains through an active listening strategy.

Cities have shown to make increasing efforts to identify and implement mechanisms, much supported by technology and digitalisation, to become more attractive and to offer better experiences towards a best quality of life. The technology factor is worth little by itself, although, at first analysis, creates a feeling of modernisation, which may be considered more attractive. Consider, for example, the Internet of things (IoT). This trend allows, through a network of sensors, to capture indicators from a closed space, a building, a condominium or a city. The information obtained will enable to make informed decisions, increasing objectivity in the definition of priorities regarding the actions to be carried out and anticipating the impact in case of success or failure.

However, a city should consider more than simple data collection to make strategic decisions; it should also consider, as the central focus, the user experience (UX) of the citizen (resident or non-resident). Thereby, besides data interpretation, it is crucial to understand the context where the data is collected to cross-reference this information with the real citizens' needs and to be creative when designing solutions that intend to attract and retain not only the citizens but also third parties.

## The best experience for the resident

Cities follow the motto to be "the best city to live in", which raises the need to rethink the way city strategies are planned, managed and built in today's reality.

Several changes are taking place, as it is the case with the more pronounced dissemination of an increasing number of events that allow living and promoting the city. Digital communication channels, namely social networks, appear as the first choice for the dissemination of such activities, as well as distinct actions for city maintenance, all of which bring life and more significant potential for the resident and non-resident user. These channels promote communication between the city and the citizen but are far from solving all the communication issues. Cities still struggle today with the need to empower the citizen with mechanisms that increase involvement and awareness, making the user a promoter of change and decision [2]. Actually, some studies (like observation and surveys) reveal that citizens either do not know which mechanisms exist or those that exist (online or offline) do not work in an appropriate way regarding their needs, being sometimes considered outdated. Furthermore, despite citizens being able to use available communication mechanisms and artefacts, they feel there is no feedback over the treatment that the city provides to the information that is shared.

So, the best city to live in is the one that can identify all the different areas necessary for the daily life of the citizen, to strategically direct efforts



for such areas; is the one that is not seen as a combination of buildings and infrastructures but the one considered a living organism that promotes a better quality of life and work by involving their citizens through mechanisms that let them express their needs and priorities, ideas and feedback.

## The best experience for the non-resident citizen

Tourism, as a business area of great importance to cities, is an exciting opportunity to show what they have to offer regarding heritage, culture, leisure, gastronomy, among other areas. It can also be a factor for attracting people to live, or just for marketing, fostering the evolution of the city's economy and the creation of services that respond to tourists, promoting job creation to answer these needs and potentiate new business models.

The digital landscape supports a more modern knowledge of the city, for example, through the use of mobile applications (apps). However, the investment to be made in specifying and implementing strategies for these apps still require some effort that may not yet have the proper return since, being used for short periods of time, tourists feel little need to install such apps. The alternatives are:

- A more objective rationale for a differentiating experience, which may include different channels, as electronic billboards, to communicate and integrate tourists, regardless of the duration of their stay. These channels should not only make part of the city as a public medium but also part of hotels and public transportation, among others.
- Designing mobile apps that follow disruptive approaches that, even for short periods of time, potentiate the installation of the app. For example, a non-intrusive gathering of tourist information (using, for instance, user-interface onboarding strategies) will not only provide better targeted and location-based information to the tourist, but also improve the information provided to the city's electronic billboards, promoting its sustained growth.

## The influence of digitalisation in the (smart) city

### The past

It is not necessary to think about a distant past to associate cities to a cadence of administrative routines, shrouded in paperwork, following process flows thematically distributed by the different areas. The objective of improving urban space has always existed, but the logic that prevailed was centred on the service provided, as understood by cities' decision-makers, and not so much in the interaction with its residents or visitors. It was a period with distinct characteristics, with established structures, where much of the basic needs, with particular emphasis on infrastructural issues, were to be met. The cadence of events followed its logic, and the vast majority of citizens were not involved and were far from decisions.

### The digital transformation

The digital transformation that has been affecting society as a whole has not left out the city, its governance and the way in which it is lived, since the citizen acquired new characteristics, gain access to new tools and to more information. The evolution of information and communication technologies (ICT), as well as process reorganisation and digitalisation (both internally and externally, most of the services were virtually exposed, allowing through web access to meet part of the city daily management needs and life in the city) and access to information, promoted changes in the daily city management, on urban planning and in the form of interaction between the citizen and the municipality.

Under the digital transformation trend and the motto of the smart cities, cities have been embracing various initiatives to combine sensing technologies with different areas of action, intending to create cities that are more intelligent,

sustainable and citizen-oriented. There are also moves to share data openly to promote transparency in governance and to foster the creation of innovative initiatives open to all citizens. The exposure of city data involves the citizens and engages them in city management activities. All of these initiatives, proof of concept and implementation on the ground are particularly interesting, but the digitisation of cities, and their evolution to smart cities, has been happening in a disorderly way, creating islands between functional areas. The adoption of heterogeneous, closed systems, from different manufacturers, with proprietary interfaces, has been a common mistake that, in the long run, will require a complete overhaul of the city's operating systems.

## Digital integration

Cities now seek to standardise their information systems by following a particular strategy, with objectives and priorities depending on their specificity, and are reluctant to accept new equipments, platforms or even solutions that do not integrate through a transparent way with existing platforms. The dispersion by different systems becomes unmanageable and proves to be of doubtful utility. It is necessary to undo built digital verticals and to base them on a common foundation, which allows reading, managing, analysing and acting on the different systems through open interfaces, well known by the various stakeholders. Assuming the integration or renewal of the demanding systems will be a critical but fundamental part of the evolutionary process so that there is no disruption of services or loss of investments already made. The creation of urban governance centres, based on urban platforms that provide useful information services to the different actors of the ecosystem, allowing them to see not only the big picture but also the detail of data, appears to be the right path to more effective municipal management and for more informed decision making, as it will be based on knowledge and not on personal feelings.

However, and although the new information systems have been providing more accessible services in online platforms, through several touch points, significantly reducing the on-site visits to municipal services, the experience for the citizen, as a user, is not yet agile, being far from a satisfactory journey. Life has become more digital, and the city will have to offer seamless digital journeys to its inhabitants and visitors to improve the quality of life in urban space, to improving citizens' perception of their relationship and involvement with it.

## The influence of industry 4.0 in the (smart) cities

The first industrial revolution happened supported by the steam machine resulting in the improvement of productivity and transports, enabling faster product commercialisation between cities. The second revolution supported by the electricity brought illumination to the city, and with it, the ability to develop new businesses and mass manufacturing. The third revolution took place hand in hand with the ICT and impacted the city tremendously as we knew today, enabling the manufacture of smart finished products and, through Internet of everything, its global interconnection, opening the way to new paradigms and applications that will make cities smarter, more inclusive and sustainable, with self-organised capabilities able to enhance the quality of urban living.

## Industry and digital transformation

Digital transformation in the industry focuses first on the change of the factory's shop floor, increasing the automation and autonomy inside the factory, where the final product itself integrates various sensors connected to the Internet. In a second phase, the industry thinks of

the dematerialisation of the processes, spanning the entire value chain, from raw material suppliers to retail distributors. This phase has a tremendous impact on cities and strategic planning, integrated policies and final deployment. Addressing the challenges in a city context through an industry 4.0 (I4.0) approach will help realise the inclusive and sustainable industrial development through its main steps as efficiency, production capability, infrastructure, logistics and innovation.

## Data correlation to leverage new services

The massive volume of data generated by the modern factory created by I4.0 initiatives, as well as the possibility of correlating all the underlying data across the entire value chain, will allow the creation of new services and new opportunities that will hatch in the services sector, usually located in the city, promoting a substantial increase of new dynamics and business models. The integration of I4.0 and smart cities available data will shape our daily lives, allowing new business models, applications and solutions that can improve the ever-growing urban population substantially. In the end, everything will be interconnected, with people residing in the cities supported by an ever-more robotised and sophisticated sectors (a popular trend [3]), evolving towards the suppression of citizens' needs at any and every stage.

## Thinking and designing a city to capture a better life experience for all

### The true smart city

Design is increasingly gaining importance when an architect and planner think the city, as it

allows determining how residents and visitors live their environment, seeking to follow standards for issues such as access. Nowadays, the city's digital services and interfaces (as an integral part of how people and companies interact with the offer made by the city) are a concern, together with the dynamic of the city, determined for example through density, mix and access (density of a city's buildings; the way people and activities are mixed; and the access or transport networks that are use to navigate through them [4]).

The UX, although still incipient in urban planning, is essential within this framework, making it possible to innovate and reducing the gap between cities and true smart cities, with a real citizen-centric approach. For example, prototyping can be very useful for capturing user feedback early in a project lifecycle when errors or modifications are inexpensive and easy to fix (in the context of a smart city, prototypes should be considered in a private or small/medium-sized public place, such as airports, universities, concert halls and retail centres).

When applied to the context of a (smart) city, UX scales the knowledge about the services that people use and need the city to provide, allowing to increase effectiveness and to identify priorities to ensure that the dynamics of the user with the city flows accordingly to what he feels about the city and the services themselves [5].

## Measuring initiatives to capture an experience tailored to the citizen needs

Results are being gathered from many initiatives around the world, which follow the axes already considered within smart cities' concept definition, such as citizens' well-being, civic awareness and experience, environmental improvement and personalisation of both citizens and visitors' experience.

Before identifying the initiatives that make sense in the context of a smart city, it is crucial to discover the underlying need. Part of this

discovery involves talking to people (citizens and visitors) to identify strategies for collecting data about their needs and values in the context of different domains of their life, such as work, leisure, health and family. It is then necessary to measure the impact of the defined strategies for the different touch points (from the public and private sectors) of the user's journey during each initiative, whenever there is an exchange of information between different entities, stakeholders, interfaces and systems, to ensure not only the improvement of his experience but also the prosperity and sustainability of the city as a smart entity (one that best combines people, processes and technologies).

In this context, UX promotes a mindset that puts the user first and promotes the right balance between business, technology and real users. Actually, to gain a deeper understanding of the user, it is possible to use some methodologies (such as card sorting, contextual research and interviews) that allow to better understanding the user's journey, to identify the touch points that provide differentiating data and insights and relate this information to generate results.

## User journey map as a tool for creating a 360-degree vision

To map the user's journey, one should take into account a 360-degree perspective of the city-wide activity through all touch points. This mapping should serve as a strategic tool in the course of the process to create a product, service or ecosystem from the city to its user. The structured and unstructured data that fits here can feed forecasting models, allow the customisation of systems and advanced analysis models (such as citizens' feelings or opinions mining) to promote a better understanding of citizens.

Concerning the services provided, it also involves a holistic perspective of the employees, who feed the experiences provided to citizens. Detailed mapping will identify ways to improve employees' productivity, operational performance, and decision-making capacity which, in turn, will affect directly

and indirectly the service and the experience of the citizen in his interaction with the city.

## Practical case: experience maps as foundations for UX alignment and vision for city hall worker profile

As a practical case, it was applied a technique called experience maps in the context of a research project supported on a platform designed for a city manager (someone who needs to have the big picture of the city, like problems that need to be solved, water consumption irregularities, and have information that helps making decisions to improve the city). This map allowed creating an alignment of what would be the expected experience for this manager.

This methodology started with the observation of a real context to represent the whole experience of the manager in the execution of his tasks. This observation enabled to realise how some of these tasks could be improved or simplified by identifying critical and time-consuming points through the process, thus creating new opportunities for improving the workflow. Besides this, it was also possible to verify that a set of functionalities had the potential for improving this profile experience, only by organising the interface.

The next step was an experience map, focused on the manager. This map, in a first analysis, took into account the person characterisation. After characterising the profile, identifying objectives and possible frustrations within the workflow process, a map of experiences was created (independently of the actual description of the journey) to describe the desired experience for the manager. The advantage of this step being done right after observation is that the desired experience is much richer, as it was focused on real problems.

Specifically applied to the interface of the platform above mentioned, the experience map allowed creating and aligning a perspective for the interface as a whole, in a structured and organised way:

- Reorient some critical points of the interface even without complete knowledge of the information that could be provided to the manager (not the ideal situation);
- Expand the short and long-term range of requirements, extra features and opportunities to the web platform, as well as reinforcement to the mobile platform (focused in a more technical profile, someone who is on the street supporting the manager in checking and maintaining public places);
- Realign strategies between the web and mobile platforms.

assume a central place to trigger digitalisation, but the citizen gets to lead the way.

Using a set of technologies for modernising the city and turning it into a more attractive place to live and visit, city's managers are now considering the citizen as a user that should be heard and engaged to more efficiently and well-succeed develop their cities.

UX design processes, methodologies and techniques urge as strategies' foundations for smart cities, as it supports its essence on the users, their needs, pains, thoughts and perspectives. The city provides the context. The citizen is the city's user.

It is believed that a perfect balance between technology, design and industry will sustain smart cities growth in the future, considering that smart cities must rapidly evolve from closed initiatives to transversal journeys, for which the citizen assumes a central role.

## Conclusions

Nowadays trends are changing cities and the way citizens live in it, promoting citizen-centric strategies to become smarter. Technologies

## References

- [1] A.-M. Slaughter, "3 responsibilities every government has towards its citizens," *World Economic Forum*, 2017. [Online]. Available: <https://www.weforum.org/agenda/2017/02/government-responsibility-to-citizens-anne-marie-slaughter/>.
- [2] L. ADLER, "How Citizens See It - Kansas City's Citizen Survey Adds Citizens' Perceptions to the Equation," 2015. [Online]. Available: <https://datasmart.ash.harvard.edu/news/article/how-citizens-see-it-677>.
- [3] R. K. and S. S. James Manyika, Susan Lund, Michael Chui, Jacques Bughin, Jonathan Woetzel, Parul Batra, "Jobs lost, jobs gained: What the future of work will mean for jobs, skills, and wages," *McKinsey&Company*, 2017. [Online]. Available: <https://www.mckinsey.com/featured-insights/future-of-work/jobs-lost-jobs-gained-what-the-future-of-work-will-mean-for-jobs-skills-and-wages>. [Accessed: 20-Aug-2010].
- [4] Kim Dovey and Elek Pafka, "What makes a city tick? Designing the 'urban DMA,'" *PHYS ORG*, 2016. [Online]. Available: <https://phys.org/news/2016-11-city-urbandma.html#jCp>.
- [5] E. Union, "About Citizen Focus," *EIP-SCC*, 2018. [Online]. Available: <https://eu-smartcities.eu/clusters/3/description>.



# 11

## UX process for product co-design empowering personal engagement

**Maria Inês Almeida, Altice Labs**

[maria-a-almeida@alticelabs.com](mailto:maria-a-almeida@alticelabs.com)

**Cátia Raquel Oliveira, Altice Labs**

[catia-r-oliveira@alticelabs.com](mailto:catia-r-oliveira@alticelabs.com)

**Gonçalo Nuno Gaspar, Altice Labs**

[goncalo-n-gaspar@alticelabs.com](mailto:goncalo-n-gaspar@alticelabs.com)

**Inês Freitas Oliveira, Altice Labs**

[ines-f-oliveira@alticelabs.com](mailto:ines-f-oliveira@alticelabs.com)

With the digital transformation and the proliferation of digital devices, it is extremely important to create simple, intuitive, usable and attractive products and services that provide users a positive and enriching experience, building the foundations for a successful customer experience strategy that will foster growth and loyalty. This article aims to describe the implementation of a UX process in a specific Altice Labs' product, where techniques were applied with and to stakeholders, consolidating a strategy to bring value to a product supported on a matured user experience.

---

### Keywords

Co-design; UX process; Design thinking; Product profile; Proto-personas

## Introduction

Nowadays, with the growing diversity of products and services available in the market, customers are becoming more rigorous and demanding. It is crucial to meet their expectations, guaranteeing the delivery of enriching experience, and building strategies that promote user loyalty. It is also essential to create simple, intuitive and visually appealing products and services that provide users with a positive experience [1], [2].

For this reason, to think about the usability of products and how they are constructed is, at the same time, to promote ease of use and improve the quality of products [3]. Human-computer interaction (HCI) supports the idea that the interaction between a computer and a human should be simple, similar to open dialogue between humans, and therefore the construction of systems must take into account human needs, without neglecting the technological ones [1] [2].

In this way, it is essential to include users and stakeholders in the process of developing a product or service. This approach is called cooperative or participatory design (Co-Design) and aims to understand user's difficulties, desires and expectations and to align them permanently with the product specification, to create a more efficient and satisfying interaction experience, with a suitable time-to-market [4]–[6].

At the same time, user research or user experience research focuses on understanding the user's behaviours, as well as their motivations and needs. We use some methods and techniques (if possible, in the early stages of the product lifecycle), allowing usability and product teams to collect rich user information, and analyse that data to be considered in the various phases of a project [7]–[9].

User experience (UX) is an area that covers the user and their context of interaction with a particular product, worrying about their emotions and the quality of their experience in general.

Therefore, the importance of interaction design is related to the fact that the main purpose is to create products that can help people in their daily and professional lives, creating experiences that improve the way they work, interact and communicate [5].

This article aims to describe the organisation and implementation of a UX process for the construction of the Altice Labs' NOSSIS Fulfillment product vision, where techniques were applied with and to stakeholders, consolidating a strategy to bring value to a product supported on a matured user experience. Involving stakeholders at the designed process is the first way of operating product thinking design. The article reports the contributions of the specified UX process when applied to NOSSIS Fulfillment product, part of OSS the portfolio.

## NOSSIS Fulfillment

Altice Labs' NOSSIS Fulfillment implements service provisioning and addresses the set of functionalities provided by two already existing Altice Labs' products: the Order Manager (OM) and the Network Activator (NA). OM is responsible for verifying and controlling the execution of orders, and for coordinating the workflow of other intervening systems, such as inventory, activation and management systems, while NA is a service activation platform that serves as a mediator between the various systems supporting the operations and the network. NOSSIS Fulfillment replaces these two products, providing the functions mentioned above.

NOSSIS Fulfillment is the centrepiece of the provisioning processes on OSS architecture, responsible for guaranteeing that all provision processes are executed in a coordinated and coherent way. Doing this means to orchestrate all the necessary activities, which may comprise resource allocation, triggering workforce management systems, network configuration, among others. Being so, NOSSIS Fulfillment is



at the epicentre of the OSS change movement towards digital services, network function virtualisation and the evolution of traditional TELCO services. It represents an opportunity for improvement in the relationship between the digital service provider (DSP), the customer service provider (CSP) and the customer, through processes automation, and allowing for an increase in operational gains and customer satisfaction.

## Drawing the process

A process, considering its different stages, can incorporate different methods and techniques to accomplish different objectives. In the case of a UX process such methods and techniques can be of investigation and development (I&D), user research, usability and user-centred design (UCD), applicable in evaluation UX.

The UX process under discussion was specified to work with the idealisation and design of the NOSSIS Fulfillment, as well as its development and testing, providing a partnership work called Co-Design.

This process seeks a change of rational that is based mainly on a methodology called Service Design Thinking. This methodology, at its first stage, identifies the needs of the user/customer and, besides, it also determines goals for the product or service to be developed, or recognises problems if the product or service already exists. As such, it is essential to think of products as services, focussing on enhancing the best experience for users and, consequently, to increase customer loyalty and, finally, market value. This approach is essential for a knowledge-based economy, where services incorporate new business models that serve the users' needs and create socio-economic value in society [10].

## Thinking the product

Considering the experience that is to be created for the user, thinking or designing a product is more than just thinking about its features. The users' tasks must be identified, as well as the finality of those tasks, what the user is trying to achieve and the associated difficulties, taking into account the time to market and the associated revenue. The reasoning of the ones who draw and the ones who develop, assuming a synergy between the two, is to understand what the product is, what it is going to be used for, what problems the product solves and what needs are satisfied by its usage.

The user assumes a prominent role in this process. Much of the information gathered to be worked in the context of a UX process comes from observing the users and talking to them in their working context, as well as to related stakeholders, to perceive and capture information that allows guiding the product to solve the identified problems quickly. This work sustains the basis for thinking a product/service starting from the identification of the real needs translated to first-to-be developed requirements, the identification of the target user (who deals directly with the problem), designing the procedure to create a better solution and, finally, to identify new functionalities for the product or to simplify the existing ones. A product or service that guides the user, simply and easily, regarding the accomplishment of specific tasks, guaranteeing a good experience, is a product or service with more value for the market. In this sense, the visual design and the interaction design for creating an easy-to-use and appealing product, while important for visual perception and differentiating experience, do not guarantee the generation of value for the market. Some questions assume a predominant part in the UX process:

- What problem does the product/service intend to solve?
- What needs does it satisfy?

- What is the guiding vision?
- What is the target audience?
- What strategy should be adopted to achieve the proposed objectives?

The identification of functionalities appears in the process as a later step.

So, as mentioned, we must think about the product from the beginning by taking into account the entire experience of the user [11].

## Working in Co-Design

### Product profile

The starting point of the journey for NOSSIS Fulfillment Co-Design is the perception of the product as well as of others that are directly or indirectly related, namely NA and OM. In addition to this technological and functional framework, we started by analysing the competition and parallel products in the market, allowing us to build a positioning map, a map of functionalities and benchmarking. For this initial analysis, competitive market research was elaborated, as it is important to create a comparative scheme for several similar products of other companies. This step accomplished in Co-Design enabled the framing of technological functionalities within users' roles and correspondent tasks. The target audience and user profiles were identified. For this process, it's essential to gather and relate all this information from the beginning of the project.

### Users' profile

User research associated with the context of business-to-business (B2B) products is hard, due to the number of users available and the criticality of their work. Often, we consider a small group of users with crucial functions that, in the case of miss-perception, interaction or failure,

may result in critical consequences and complex repercussions, hard to analyse and solve.

In this context, we only considered internal stakeholders, although ideally, we should have also considered the final users. The stakeholders considered were: the development team, the product architecture team, and the product manager. Because of the genesis of the product, we also included in the process a set of NA and OM internal users. These stakeholders were selected because they, directly or indirectly, make part of the process, are related to the product, and also are closer to the final users

The creation of a stakeholders map allowed us to organise and guide the contact with the users, guaranteeing our knowledge about their context and relation with the product. At this stage, the product manager, with his strategic vision of the product, played a prominent role in guiding and focus, given the divergent possibilities that this analysis could take. This focus guaranteed the timing delineation of the work to accomplish.

### Sensing the product over time

Then, we started interviewing the stakeholders. At this stage, we wanted to identify the roadmap, how OM and NA related to NOSSIS Fulfillment, what were the expectations for the product and the alignment with external competitors, as well as the internal positioning and strategy. A mind map was used to present the results of the interviews since it allowed us to visually explore concepts and ideas, converting a list of extensive information into an organised and more appealing diagram that translates a common language.

### Characterising the product

By now, we have fulfilled the conditions to characterise the product and define the most suitable language and communication to present the product. At this stage, we worked in Co-Design to specify different affinity diagrams, aiming to identify the most and least relevant characteristics

and five concepts that define the product. Why five? Because we needed a short and concise list of topics that we could use to rapidly characterise the product as if it was to be in the market. These five concepts have to be the ones that allow us to define the product and to be the distinguishing characteristics seen by the market.

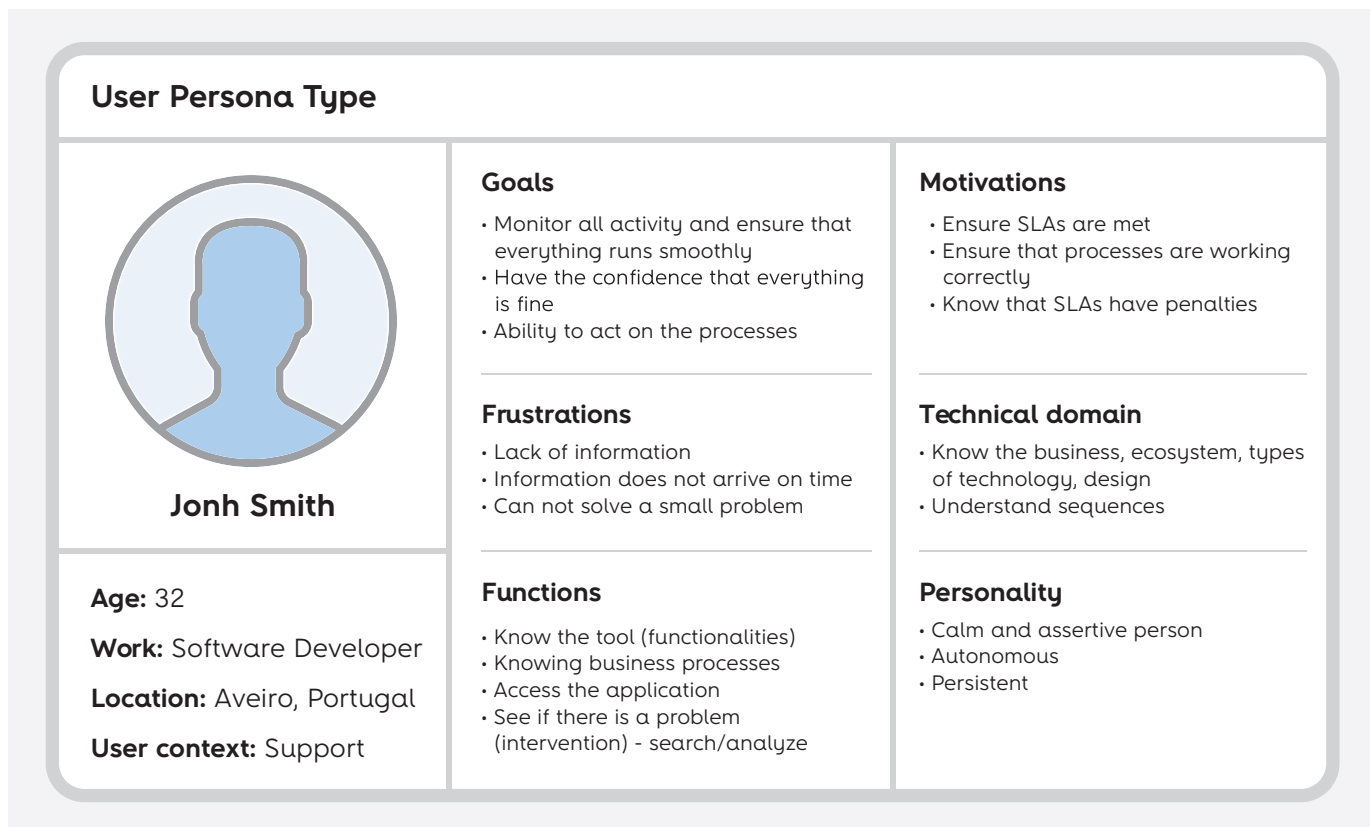
The information gathered in the interviews proved to be fundamental to the success of this stage, supporting all future analysis. Also concerning the work developed for characterising the product, its personality must be defined in Co-Design, as well as use cases to identify the relevant aspects, the end-user gains and example usage scenarios. These topics will be further discussed in the following sections.

### 5.5. Proto-personas

To identify the different user profiles that interact with the product, the development of proto-

personas in Co-Design with the stakeholders is fundamental. Proto-personas are fictional characters drawn from user research, to serve as example cases for user experience development. The goal is to create profiles of real users, but with fictional information, representing their needs. For example, one must identify the name, age, image, work or activity, objectives, technical domain, the context of use, relevant demographic information, functions and motivations. Its credibility is fundamental [12].

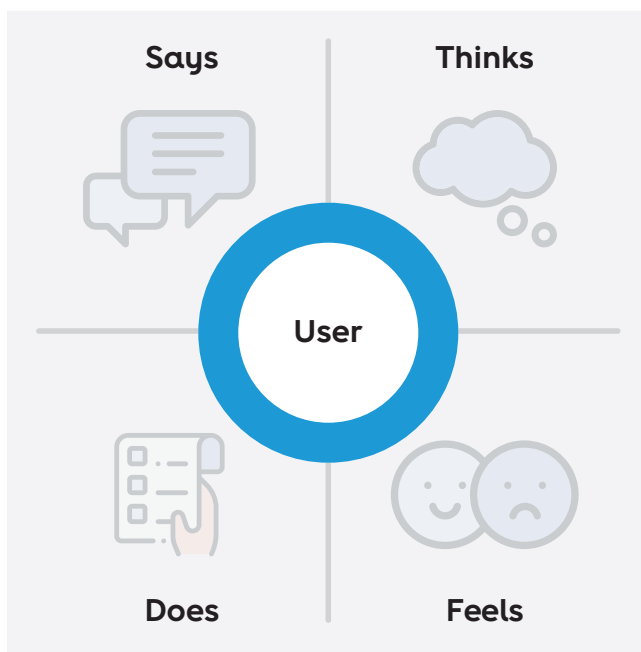
Here, the stakeholders are essential for the rapid identification of some of the user’s needs/complaints, their workflows and goals in the accomplishment of tasks, due to their past interactions with the final users. The main difficulty in using this technique with the stakeholders is that the characterised proto-persona must be described in maximum detail according to the real user and not according to a perception one has. **Figure 1** shows an example of a persona.



**FIGURE 1** – Proto-personas (in Co-Design)

## Empathy

Then, we proceeded to create an empathy map for each proto-persona (i.e., each user profile identified), also in Co-Design with stakeholders, as shown in **Figure 2**. With this artefact, we wanted to understand the mentality and thoughts of each type of user, to know what they say, think, make and feel whenever they have to perform actions during their day to day work. The goal is to work with real examples to translate a similar experience from the one of a real user whenever he/she interacts with the product. Although this technique is used to gather much information from the users, its application proves to be challenging, because it requires a non-technological, and more emotional, analysis.



**FIGURE 2** - Empathy map (in Co-Design)

## Vision: the experience offered

The next step was to deepen the Co-Design team's knowledge about the work process of each user profile for a better understanding of the different touchpoints. As such, both proto-personas and empathy maps served as the basis for creating experience maps for each type of

user. This type of map is divided into four parts: phases, actions, thoughts and emotions. We can increase value to the experience map by adding the context of use and needs or improvements/opportunities. Also, presenting the user actions in chronological order helps to create a holistic view of a user's experience with the product. After its implementation, the map was presented and analysed by the stakeholders.

With the information collected from the proto-personas and the empathy map we constructed experience maps that allowed to visualise the end-to-end user experience, to accomplish:


- daily work phases;
- what the user does at each stage;
- touchpoints with other teams/devices;
- the thoughts/feelings;
- some needs and opportunities for the product.

Based on the information gathered with the experience maps, we built an expectation map where, for each user's profile, we identified possible working process improvements and tools.

## Product positioning

After understanding the different types of NOSSIS Fulfillment users, we started working on positioning the product to validate in co-design with the stakeholders the benchmarking analysis that was done previously and to define an alignment between the all the stakeholders involved in the internal product positioning. A focus group session was carried out, involving the interviewed stakeholders. For this session, we used card sorting, a technique that allowed us to understand how the information presented is structured and organised. With this, we asked the stakeholders to group cards with terms describing the product by similarity/concept and to assign a designation for each one. It served for creating the first level of positioning concerning the business model intended to adopt. Here, we achieved two different goals: 1) we analysed the market

positioning for NOSSIS Fulfillment by presenting the initial benchmarking research together with competitors market analysis; 2) we combined the basic and differentiating concept of NOSSIS Fulfillment, defining an internal positioning for the product and contributing to the alignment of all stakeholders for the product roadmap. **Figure 3** shows an aspect of this session.



**FIGURE 3** – Product Positioning (focus group/card sorting – in Co-Design)

## Vision: strategy

As a result of the various Co-Design sessions with the stakeholders, we prepared and presented a strategy document. In addition to the artefacts constructed in Co-Design, we also produced strengths, weaknesses, opportunities, and threats (SWOT) and risk/impact analysis charts. This document intended to serve as basis and justification for the product manager's options regarding the defined roadmap.

## Product personality

Finally, supported by the initial sessions with the stakeholders, we addressed the personality of the product as the first step to devise the identity of NOSSIS Fulfillment.

To design the identity of a product is essential to define the positioning and the way the product

presents itself to the market. The key to attract the client is to delineate an objective and clear product personality and to value characteristics and attributes of the product. On the other hand, thinking about the identity towards guaranteeing a differentiated offer strengthens the bond a client establishes with a company [13].

To encourage creativity, we asked stakeholders to think of the product as if it was a person, a superhero or a cartoon character that best represents and characterises NOSSIS Fulfillment. We then defined a design persona that described and characterised the product as an imaginary person and created personality maps for each user profile. Whenever possible, we also addressed the voice tone per user and copywriting, typographies and representative colour pallet according to the personality established. These steps allowed creating a unified product identity and visual layout.

## Results

The document for the vision of NOSSIS Fulfillment and its presentation to the stakeholders added up different perspectives and knowledge. Having different parties involved guaranteed a mutual alignment and promotion of discussion among the participants, encouraging the creation of a shared vision.

All the artefacts produced along this process had an impact and added value to the product because they were obtained in Co-Design with the stakeholders. More specifically for the NOSSIS Fulfillment, it meant a roadmap revision to accommodate new requirements that weren't yet under consideration. This process also allowed to identify guiding lines to the product evolution with the goal to offer its users a better user experience and increased customer satisfaction.

Documentation of the process allowed its reanalysis and re-use in different product and service contexts, as well as the analysis of the best

methodologies and techniques for the defined objectives and information to collect. The specific process was a drive for the goal of creating a holistic vision for guaranteeing a better end-user experience and digital differentiation, but one must be aware that this is influenced by external factors, such as time-to-market, technological evolution, clients who already use the solution in hands, among others.

## Conclusion

In the end, this process and the techniques applied in Co-Design with the stakeholders of the product guarantee the joint construction of a vision for NOSSIS Fulfillment.

Two fundamental questions can qualitatively measure the importance of applying a UX process in Co-Design to NOSSIS Fulfillment: How effective and useful is the UX process? What is the value of applying a UX process in Co-Design to the product?

We can conclude that a UX Co-Design process allows promoting stakeholders under different perspectives, not usually considered in a regular process. This process enabled us to approach the theme differently and produce new ideas for the product. The stakeholders involved considered this process highly useful to provide a guideline for product evolution to increasingly offer its users the best user experience and increase their satisfaction.

From the product managers' perspective, it is important to highlight the joint creation of a visual language (the artefacts resulting from the application of the different methodologies and techniques), which helped to systematise and clearly state what is often known only empirically, inducing alignment and focus on the product, but also taking into account the perspective of the user. If systematically applied, the UX process in Co-Design supports a continuous improvement for product evolution, allows a rationale for the

product roadmap, encourages the aligning of ideas, promotes skills adjustment (management, product, UX, development) and organises a sustained rationale for the different teams involved.

From the UX designers' perspective, the process is useful to align different stakeholders and the different skills required for the creation of the product/service. It supports focussing on the end user and allows a UX rationale for the product, service or working process. In the case of NOSSIS Fulfillment, as well as for different products, the benefits are collected in the medium/long term. Despite this timeframe to sense the results on the product side, for UX designers the result is immediate and allows the creation of a consistent, sustained vision, with the right balance between the objectivity and subjectivity inherent to UX/visual design. It enables creating solid bases for a product personalisation strategy, consequently making the product closer to the user.

Another aspect to highlight is the importance of adopting interactive and dynamic techniques in the teams' working processes, which enable problems solving for existing products and thinking about new solutions for products to be or being developed.

*\* This article was based on the dissertation for the Master in Multimedia Communication of Maria Inês Almeida, developed at Altice Labs and sponsored by Inova-Ria Research Scholarship Program - GENIUS. The objective of her work was to design a methodological heuristic (a set of principles and guidelines based on empirical experience), supported by UX methods and techniques applicable to Altice Labs' business ecosystem.*

## References

- [1] H. I. T. Dourado, "Análise comparativa entre Lean UX e métodos tradicionais de UX Design," Universidade do Porto, 2014.
- [2] F. Karray, M. Alemzadeh, J. A. Saleh, and M. N. Arab, "Human-Computer Interaction: Overview on State of the Art," *Int. J. Smart Sens. Intell. Syst.*, vol. 1, pp. 137-159, 2008.
- [3] J. Nielsen, "Usability 101: introduction to usability," *Nielsen Norman Gr.*, 2012.
- [4] J. Preece et al., *Human-computer interaction*. Addison-Wesley, 1994.
- [5] J. Preece, H. Sharp, and Y. Rogers, *Interaction Design - Beyond Human-computer Interaction*. Wiley India, 2013.
- [6] J. Rubin and D. Chisnell, *Handbook of usability testing : how to plan, design, and conduct effective tests*. Wiley, 2008.
- [7] M. Kuniavsky, *Observing the User Experience: A Practitioner's Guide to User Research*. Morgan Kaufmann Publishers, 2003.
- [8] usability.gov, "User Research Basics," *U.S. Department of Health & Human Services*, Oct-2018. [Online]. Available: <https://www.usability.gov/what-and-why/user-research.html>.
- [9] The Usability Body of Knowledge, "User Research," *Usability Body of Knowledge*, 2012. [Online]. Available: <http://www.usabilitybok.org/user-research-methods>.
- [10] M. Stickdorn and J. Schneider, *This is Service Design Thinking: Basics-Tools-Cases*. BIS Publishers, 2012.
- [11] N. Blaase, "Why Product Thinking is the next big thing in UX Design," *Medium*, 2015. [Online]. Available: [https://medium.com/@jaf\\_designer/why-product-thinking-is-the-next-big-thing-in-ux-design-ee7de959f3fe?hsenc=p2ANqzt-v333zfuj8-xipHrBosO7C2ZHHmICS Mk8atFVp\\_KkEA4BA4Quv29ookhG3PjJXwV3sKAwECU3esJmfkVP4LzU1\\_6Cfg&hsmi=20492920](https://medium.com/@jaf_designer/why-product-thinking-is-the-next-big-thing-in-ux-design-ee7de959f3fe?hsenc=p2ANqzt-v333zfuj8-xipHrBosO7C2ZHHmICS Mk8atFVp_KkEA4BA4Quv29ookhG3PjJXwV3sKAwECU3esJmfkVP4LzU1_6Cfg&hsmi=20492920).
- [12] J. J. Garrett, *The Elements of User Experience : User-Centered Design for the Web and Beyond, Second Edition*, Second Edi. 2010.
- [13] D. da P. Oliveira, "Personalidade da Marca: uma forma de fidelizar o cliente. - Artigos - Marketing - Administradores.com," 2011. [Online]. Available: <http://www.administradores.com.br/artigos/marketing/personalidade-da-marca-uma-forma-de-fidelizar-o-cliente/58868/>.





# 12

## NB-IoT @Altice Portugal

**Luís Lamela, Altice Portugal**

[luis-m-lamela@telecom.com](mailto:luis-m-lamela@telecom.com)

**Álvaro Gomes, Altice Labs**

[AGomes@alticelabs.com](mailto:AGomes@alticelabs.com)

**Isabel Borges, Altice Labs**

[Isabel@alticelabs.com](mailto:Isabel@alticelabs.com)

The connectivity between physical devices and the Internet is key to convey sensor information to powerful platforms and applications. One of the most promising LPWAN technologies used for physical device data communication is the Narrow Band IoT, a standardised technology that fulfils the above requirements, while keeping the security, scalability and reliability of 3GPP mobile networks. This article presents the Narrow Band IoT and discusses its characteristics and use cases of application from Altice Portugal's perspective.

---

### Keywords

M2M; LPWAN; Technology comparison; NB-IoT applications; golabs.IoT

## Introduction

The first years of the 21st century witnessed major developments in electronics technology, information and communications technologies, ubiquitous connectivity and manufacturing processes, leading to a significant drop in costs in almost all components, equipment and even services.

Currently, the Internet of things (IoT) has the leading role in the convergence of the technologies mentioned above, and it is estimated that it will continue this way, according to the figures presented in some recent forecasts for its growth [1] [2] [3].

All available data, coming from a multitude of IoT sources and contexts, must be gathered, prepared, processed and visualised to deliver insight (or knowledge) on the best decisions to take.

The connectivity between physical devices and the Internet is key to convey sensor information to powerful platforms and applications. As a matter of fact, new paradigms to support machine-to-machine (M2M) communications are being created since the communication behaviour of machines is different from that used by humans. On the other hand, the volume of traffic is greatly reduced for the most typical IoT applications. Finally, the physical devices are typically deprived of mains electrical energy, since they can be used in inhospitable environments and simple, battery operated low-cost devices could be deployed in massive scales.

Energy consumption, however, was an issue, due to the high energy demands of existing telecommunications protocols and physical interfaces, leading to limited autonomy. This obstacle led to the development of a new wave of wireless telecommunication technologies, usually called low power wide area networks (LPWAN).

The LPWA networks needed to minimise energy consumption on physical devices when transmitting or receiving information. Simultaneously, as the goal was to support massive numbers of devices, the networks themselves would have to be very cost-efficient, minimising the number of active elements

and other resources required to provide service to large areas with thousands or millions of devices.

One of the most promising LPWAN technologies used for physical device data communication is the narrow band IoT (NB-IoT or LTE Cat NB1). This technology was standardised by the 3GPP for LPWAN, fulfilling the above requirements, while keeping the security, scalability and reliability of 3GPP mobile networks (the standardisation body that defined UMTS, LTE, 5G, etc.).

This article presents the NB-IoT technology, and discuss its characteristics and use cases of application.

## NB-IoT technology overview

Almost all the current technologies supporting the LPWAN market are non-standardised. Therefore there are weaknesses like poor reliability, inadequate security and high operational and maintenance costs. Furthermore, new and specific overlay network deployment is complex. These reasons led to the standardisation, by 3GPP, of a technology that could be reliable, scalable and taking advantage of the mobile network already installed: the NB-IoT. In June 2016, 3GPP Release 13 contemplated the “LTE Evolution: NB-IoT: Cat NB1” [4].

The main goals for NB-IoT protocol design are presented on **Table 1**.

To achieve these goals, NB-IoT uses specific techniques:

- **Spectral efficiency** - Use of “narrowband” concentrates energy for higher powered, but shorter transmissions, using less energy than low powered but longer transmissions. Orthogonal frequency-division multiplexing (OFDM) encoding ensures high information density.
- **Power saving** - Power saving mode (PSM) and

High device density	Low-cost devices	Superior battery autonomy	Enhanced coverage	Easy integration	3GPP security, roaming and QoS
Up to 50 000 devices per cell and data rates of tens of kilobits per second, with a low frequency of communications.	Modules should cost less than 5 USD.	Up to 10 years autonomy with a low-cost battery.	20dB more than GPRS coverage, reaching remote or difficult access places.	Simple upgrade to a 3GPP mobile network.	Use of LTE mechanisms for security and roaming.

TABLE 1 - NB-IoT goals

extended discontinuous reception (eDRX) allow the device to agree with the network on periods where its radio can be switched off without losing its network registration, therefore saving energy. When leaving PSM or eDRX states, the device does not need to re-register to the network.

- **Simple module design** - NB-IoT modules only need to support one antenna for half-duplex communication and have reduced requirements for processing power and memory, due to the lower bandwidth requirements (tens of kilobits per second).
- **Enhanced coverage** - NB-IoT uses specific

encoding and modulation schemes to support a 164dB link budget.

- **Mobile network integration** - NB-IoT is an upgrade to the operator’s existing LTE network, taking advantage of the already existing antennas and nationwide coverage.

A simplified functional architecture of NB-IoT is depicted in **Figure 1**. At the lower level, there are NB-IoT devices, like sensors, actuators and gateways that are able to connect to the network through an NB-IoT base station (BS). The core network routes the NB-IoT devices’ data to an IoT platform. The IoT platform will ingest the data, process it and expose it to the application layer.

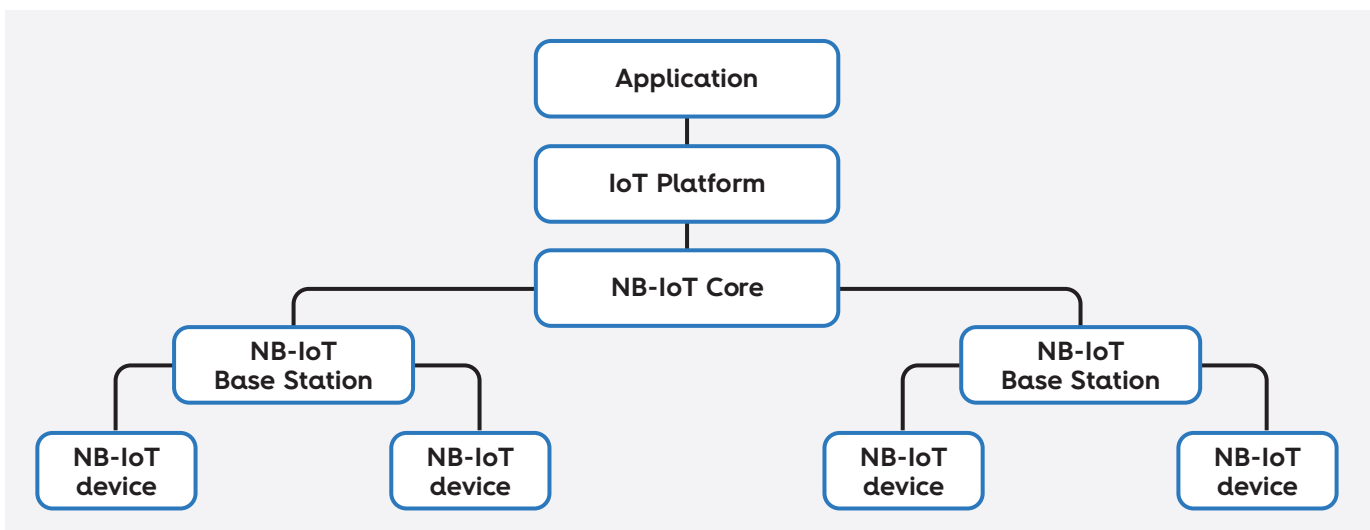
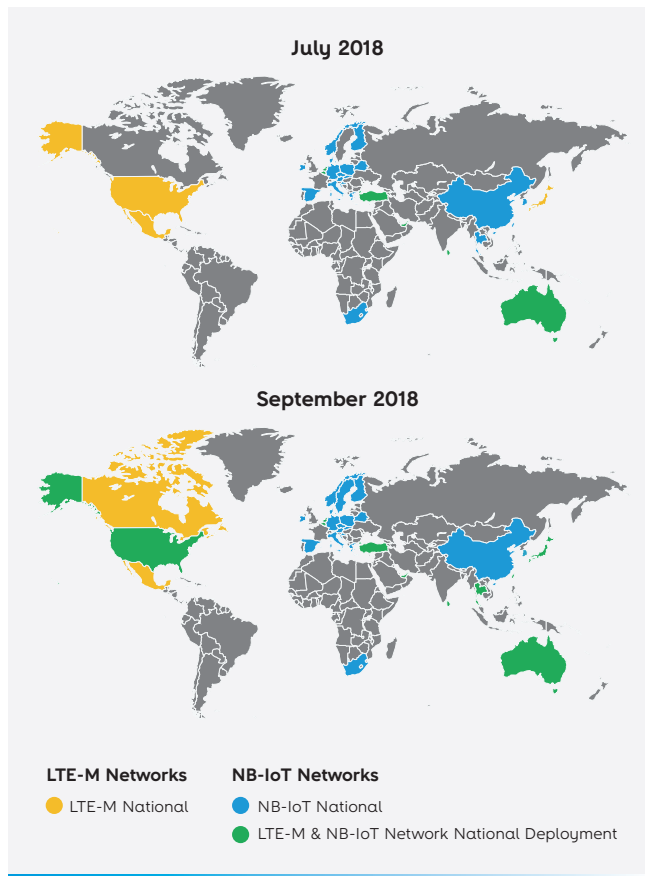


FIGURE 1 - NB-IoT functional architecture

**Figure 2** shows the IoT mobile deployment, comparing maps of July and September 2018. It is visible the adherence to mobile IoT. In [5] is the list of all mobile IoT commercial networks (total 66 launches as of September 2018) and the providers.



**FIGURE 2** – Mobile IoT deployments [6]

NB-IoT specifications were further enhanced in 3GPP Release 14 (June 2017) with some crucial features. Release 13 NB-IoT device category (Cat) is referred to as Cat NB1 and Release 14 devices are referred to as Cat NB2. Cat NB2 devices have the following features which are not available in Cat NB1:

- Positioning using observed time difference of arrival (OTDOA);
- Mobility enhancement from seamless cell re-selection;
- Support for higher data rates;
- New frequency bands;

- New power class (14dBm).

NB-IoT Cat NB2 is not yet currently commercially available.

## LPWAN technology comparison

The LPWAN technologies provide cost effective devices, enabling a long battery life, a large coverage and serving an extremely high density of users. This business requires massive deployments and user acceptance in order to become economically viable. There are presently several technologies available to address these requirements. The most relevant today are:

- **LoRa** – Proprietary, LoRaWAN open standard of LoRa Alliance, operate over industrial, scientific and medical (ISM) bands.
- **SIGFOX** – Proprietary, developed by SigFox company, operate over ISM bands.
- **3GPP** – Open standards, operate over LTE and GSM licensed bands:
  - **LTE-M** (or LTE-MTC) – Formally known as enhanced machine type communication (eMTC) is the first LTE evolution for IoT communications enabling a wide range of services.
  - **NB-IoT** – New LTE solution to support ultra-low bit rate applications.
  - **EC-GSM-IoT** – GSM enhanced technology to support low-power wide-area needs. EC-GSM can be deployed on existing 2G networks, being a suitable option when 4G is not available.

Some other technologies have also been developed, such as DASH7, Weightless, RPMA, but for the moment don't have enough commercial traction to become relevant.

Parameter	LoRa	SIGFOX	LTE-M	NB-IoT Cat NB1
Range [7]	< 15Km	< 50Km	< 11Km	< 15Km
Maximum coupling loss [8]	157dB	153dB	160dB	164dB
Max uplink peak data rate [8]	50Kbps	100bps	1Mbps	250Kbps
Spectrum	Unlicensed EU 868MHz US 915MHz [7]	Unlicensed EU 868MHz US 915MHz [7]	LTE bands	LTE in-band, guard bands, stand-alone
Bandwidth	< 500KHz [7]	100KHz [7]	1.08MHz [9]	180KHz [9]
Radio technology	Spread spectrum [7]	Ultra narrow band [7]	OFDM	OFDM
Bidirectional modes [8]	Yes	Yes	Yes	Yes
Uplink maximum output power	14, 20dBm [10]	14, 22dBm [11]	20, 23dBm [9]	23dBm [9]
Re-use existing cellular networks	No	No	Yes	Yes
Autonomy [7]	> 10 years	> 10 years	> 10 years	> 10 years

TABLE 2 – LPWAN technology comparison

Table 2 lists the main technical characteristics of the most relevant LPWAN solutions presently available.

Proprietary solutions (SIGFOX, LoRa, ...) have lower prices on devices and network equipments and can achieve higher levels of penetration with simpler deployments, while 3GPP cellular technologies offer the quality of mobile networks, enable higher throughputs and take advantage of existing operational and business systems. The LPWAN proprietary technologies have been in the field for some time, while standard solutions (3GPP) are already available but still given the first steps.

The unlicensed spectrum used by proprietary technologies could represent a difficulty concerning reliability and service level assurance due to the high number of competing technologies sharing the same spectrum, while licensed spectrum used by 3GPP standardised systems guarantee a better control of the quality of service.

The cost of communication modules continues to decrease, as shown in Figure 3, with SIGFOX presenting the most competitive prices (around 1 USD per module vs 12 USD per NB-IoT module by the end of 2017) [12].




FIGURE 3 – Declining cost of LPWA communication modules [15]

Deployment scale plays an important role in costs, with China's current nation-wide bet on NB-IoT greatly contributing to the current 5 USD cost for a single band NB-IoT module [13].

NB-IoT can support applications requiring only limited data connections at low cost. NB-IoT is the most flexible technology regarding spectrum usage and can be deployed on 2G, 3G or 4G bands. LTE-M targets more advanced services, allowing higher bandwidths, mobility and voice calls, with a more limited coverage and autonomy.

## NB-IoT applications

NB-IoT is particularly suited to applications involving a large number of battery-operated devices with relatively small and infrequent data transmissions. A typical case is smart metering. With 6+ million electricity meters and 4+ million water meters in Portugal, there is a need to support an overall high number of devices, but also high-density scenarios (city centres) is clear. Additionally, water and gas meters cannot be mains powered, and the utilities' business cases usually require unattended operation for several years, which leads to a requirement to support efficient, low-power operation. Data transmission is usually limited, ranging from a few bytes to a few kilobytes per day. Finally, the communication network needs to be able to reach meters' installation places like remote sites or basements. NB-IoT provides adequate answers to these requirements, adding the benefits of 3GPP mobile network security, licensed spectrum to ensure reliability while avoiding the vendor lock-in.

NB-IoT application examples include:

- Smart metering;
- Environmental/agriculture monitoring;
- Smart parking;
- Street lighting;

- Asset tracking;
- Fire, flood, gas alarms;
- Smart locks;
- Urban shared bicycles;
- Waste management;
- Infrastructure/facilities monitoring;
- Emergency buttons/keyfobs.

As seen in the above examples, NB-IoT strong network coverage, low cost, energy efficiency and reliability enable connectivity for many different devices, enhancing capabilities and functions. Connecting these devices also enables the collection of usage information, as a tool for continuous service improvement and other data monetization oriented applications.

Consider a shared bicycles service. Besides enabling smartphone-based lock/unlock, an NB-IoT bicycle is also able to periodically gather location data which can be aggregated and used to determine typical patterns for demand essential information to improve bicycle dock location and capacity.


## NB-IoT and the DSP

Digital transformation encompasses the deep transformation of society and businesses through the use of digital technologies, always with a strong focus on the improvement of people's quality of life, and will cause a real disruption over the business landscape. Much of this transformation comes with an exponential data growth, together with an increasing eagerness of better customer experience throughout a recentring on customer strategy. IoT, no doubt, is the technology that most contributes to this new world and connectivity is key to supporting this new way of life.

Operators are moving from communication service providers to full digital service providers

(DSP), and while they continue to ensure the best connectivity ever, they are also stepping up to higher layers in the value chains, building new applications and services. Having the vision and the possibility of an end-to-end approach positions them quite well, and that's why many operators are investing in NB-IoT and performing trials and commercial launches.

According [14], the 5 areas of IoT disruption are those presented in **Figure 4**. The same report states that the two main barriers to full IoT adoption are security and cost, with security at the top.



**FIGURE 4** – The 5 areas of IoT disruption [14]

IoT security-related technology, concerning asset discovery, profiling and tracking, authentication, network-based protection, visibility through monitoring, detection and response are all technologies that can be handled by using NB-IoT.

Managed connectivity powered by a secure mobile network, able to interface with network activation modules from the network side and with the IoT platform for the other side, will be easier since it is already done on mobile networks

and applications, relying on the subscriber identity module (SIM).

The advent of 5G will enable mission-critical IoT, broadband and time sensitive use cases that will integrate with other less stringent IoT use cases, consolidating the digital transformation.

## NB-IoT deployment at Altice Portugal

Altice Portugal is continuously evolving its offer of IoT solutions. Besides Wi-Fi and GPON, 2G (GPRS), 3G (UMTS) and 4G (LTE) connectivity are widely available through its national network, in commercial packages tailored to the specific needs of the Internet of things services. A self-care portal (IoT Connect) allows the user full control over the connectivity services, with easy to use interfaces to manage status, price plans, diagnostics, etc.

NB-IoT national connectivity was recently launched, benefitting from the same wide availability and ease of use.

Altice Portugal connectivity options are complemented by the *IoT Place* platform, which enables quick application prototyping, development and deployment, reducing the development cycle and ensuring security and scalability.

Together, connectivity and platform are the perfect foundation for competitive solution developers to bring new and innovative IoT products to the Portuguese market. Altice Portugal strongly supports several IoT projects, frequently partnering with other companies to create an interesting and diverse portfolio of vertical IoT solutions. Furthermore, through initiatives like the startup accelerator ENTER, the Altice Innovation Award prize and the IoT Challenge Hackathon, Altice Portugal actively promotes innovation and entrepreneurship centred in the Internet of things.

## Altice golabs.ioT – A GSMA Open IoT lab

The Internet of things is undergoing an “explosive” phase where, besides NB-IoT, several key technologies and services are reaching operational maturity, while the overall pace of innovations in technologies or business approaches is still accelerating.

Portugal is a good example, affirming itself as a European startup hotbed, including many IoT focused projects. A proof of this is the yearly IoT Challenge promoted by Altice, which assembles more than ten new companies working in the IoT field for a much disputed technical demonstration hackathon.

The complexity and multidisciplinary of many IoT opportunities require an ecosystem-oriented approach, where multiple partners with distinct specialities need to work together for effective solutions to reach the market.

Altice Portugal, being the leading telecom operator in Portugal, is in a privileged position to

foster this ecosystem and is assisting Portuguese companies and startups in their IoT projects through several initiatives, *golabs.ioT* being the most recent example. *golabs.ioT* integrates the GSMA global network of Open IoT Labs, a coordinated industry-wide effort to foster standards-based IoT technologies.

Launched in February 2018, *golabs.ioT* is an open IoT Lab, which is accelerating IoT projects by:

- Making a full IoT testbed available to companies and startups, including NB-IoT;
- Assisting in the validation of IoT devices and applications;
- Accelerating IoT solutions development with the technical expertise of its specialists;
- Promoting technical workshops and training sessions;
- Being a focal point for partnerships with relevant industry players, academy and others.

*golabs.ioT* complements other Altice Portugal initiatives fostering innovation and entrepreneurship, as shown in **Figure 5**.

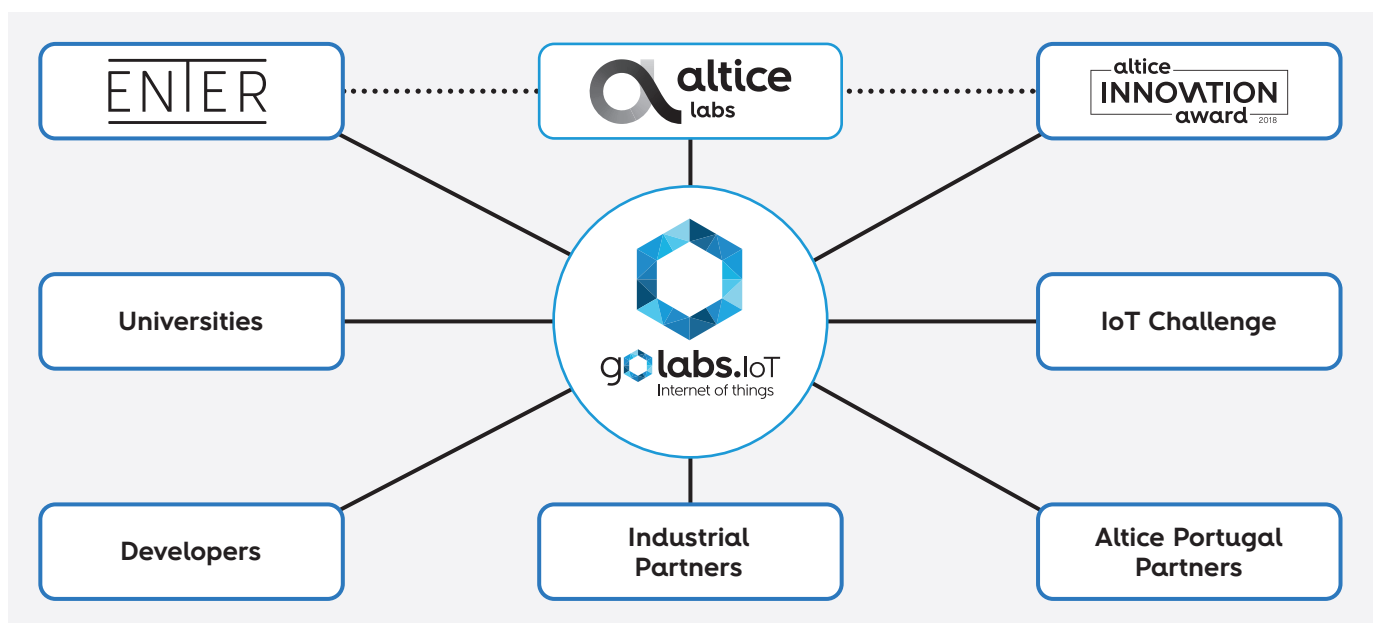


FIGURE 5 – *golabs.ioT*



Since its launch, *golabs.ioT* has been involved in several groundbreaking IoT projects, helping companies to define robust, secure and scalable service architectures, choose the most appropriate technologies and protocols, identify the best partners and suppliers and conduct a thorough test plan, thus ensuring these innovative solutions reach the market in a timely manner and adhering to a high-quality standard.

## Conclusions and recommendations

NB-IoT solutions can reuse the existing mobile infrastructure, being easily deployed over the existing 2, 3 and 4G mobile networks, and taking advantage of its licensed spectrum to guarantee a better quality of service when compared with the unlicensed, shared, and sometimes overcrowded spectrum used by proprietary solutions.

NB-IoT can also rely on 4G mobile network OSS and analytical and business platforms to deploy and manage IoT services. The use of common systems and platforms to support both IoT and the traditional mobile services can eventually bring gains concerning capex and opex savings.

3GPP is a very dynamic organisation continuously working to enhance its standard architectures and providing backward compatibility with the previous releases. NB-IoT deployments are expected to rise significantly in the forthcoming years as the technology matures. Adopting 3GPP systems is a guarantee that the evolution of new solutions is done smoothly, when needed, and respecting past investments.

Machina Recherche [15] forecasts that the total number of IoT connections will grow from 6 billion in 2015 to 27 billion in 2025, a compound annual growth rate (CAGR) of 16%. 11% of the connections in 2025 will use LPWAN connections, such as SIGFOX, LoRa or NB-IoT. IoT is growing at an accelerated pace as humans are becoming surrounded by all types of sensor and actuators, from wearables to home automation, from smart cities to industry 4.0, from entertainment to healthcare, among many others. This is the foundation for a set of new areas that harness the potential of IoT to create disruption and innovation. One of those areas is the communication area.

Among LPWAN technologies, the most promising one is NB-IoT that fits the best in offering a rich set of functionalities that could be used together with other technologies to give the best user experience, either at the personal or the professional space.

## References

- [1] L. Columbus, "2017 Roundup Of Internet Of Things Forecasts," *Forbes*, 2017. [Online]. Available: <https://www.forbes.com/sites/louiscolombus/2017/12/10/2017-roundup-of-internet-of-things-forecasts/#5853581d1480>.
- [2] Gartner Inc., "Leading the IoT," 2017.
- [3] M. Shirer and M. Torchia, "IDC Forecasts Worldwide Spending on the Internet of Things to Reach \$772 Billion in 2018," *IDC*, 2017. [Online]. Available: <https://www.idc.com/getdoc.jsp?containerId=prUS43295217>.

- [4] ETSI TS 136 141 V13.6.0 (2017-01), "LTE; Evolved Universal Terrestrial Radio Access (E-ULTRA); Base Station (BS) conformance Testing (3GPP TS 36.141 version 13.6.0 Release 13)," 2017.
- [5] GSM Association, "Commercial Launches," GSMA, 2018. [Online]. Available: <https://www.gsma.com/iot/mobile-iot-commercial-launches/>.
- [6] GSM Association, "Mobile IoT Deployments," GSMA, 2018. [Online]. Available: <https://www.gsma.com/iot/deployment-map/#deployments>.
- [7] S. Tabbane, "IoT Network Planning," *ITU Asia-Pacific CoE Progr. "Developing ICT Ecosyst. to harness Internet Things,"* 2016.
- [8] Franklin Heath Ltd, "LPWA Technology Security Comparison," p. 20, 2017.
- [9] 3GPP, "Standards for the IoT," 2016. [Online]. Available: [http://www.3gpp.org/news-events/3gpp-news/1805-iot\\_r14](http://www.3gpp.org/news-events/3gpp-news/1805-iot_r14).
- [10] LoRa Alliance Technical Marketing Workgroup, "LoRaWAN What is it?," 2015.
- [11] SIGFOX, "Sigfox Technical Overview," 2017.
- [12] B. Ray, "Market Trends: LoRa Offers Low-Risk, High-Reward LPWA Opportunity." Gartner, Inc., p. 14, 2017.
- [13] GSM Association, "China Telecom Offers NB-IoT Nationwide," GSMA, 2018. [Online]. Available: <https://www.gsma.com/iot/rollout-china-telecom/>.
- [14] S. Ganguli and T. Friedman, "IoT Technology Disruptions: A Gartner Trend Insight Report," 2017.
- [15] Machina Research, "Press Release: Global Internet of Things Market to Grow to 27 Billion Devices, Generating USD 3 Trillion Revenue in 2025," 2016.





# 13

## An overview of the benefits of superfast broadband networks to rural areas

**Cláudio Rodrigues, Altice Labs**

[claudio-e-rodrigues@alticelabs.com](mailto:claudio-e-rodrigues@alticelabs.com)

**Frederico Carvalho, Altice Labs - Viseu**

[frederico-a-carvalho@alticelabs.com](mailto:frederico-a-carvalho@alticelabs.com)

Broadband and high capillarity, digital networks may improve the life in remote, underpopulated and rural areas in several different ways. With the increase of superfast broadband, better connectivity allied to IoT is the key to boost rural economies and may also be the key to creating new revenue sources for DSP.

---

### Keywords

Superfast broadband networks; Rural areas; IoT

## The problem to solve

The world's population growth forecasts are frightening from all viewpoints, especially when focusing on the expected growth for cities. The World Bank foresees [1] that by 2050 and considering the whole world:

- About 70% of the world's population will be urban;
- Cities and urban areas will generate more than 80% of the GDP;
- Cities and large metropolitan areas will consume 2/3 of all the available energy;
- Urban areas will be accountable for the production of 70% of the greenhouse gas emissions.

Therefore, in one hand, cities will become extremely hard to live in and to manage and, on the other hand, rural areas will be highly underpopulated, with severe economic problems. Having in mind the second consequence just mentioned, this article will review the importance of superfast broadband networks and how they can help to mitigate isolation and economic problems in rural areas.

## A possible solution: superfast broadband

As an infrastructure, broadband networks can be seen in a similar way to the development of roads, railways, power grids and the way that these networks have transformed economic activities for individuals, companies and governments. The broadband networks being built today have been likened to *"the roads and bridges and railways that were built in previous times"* [2]. These new infrastructures have made economies more efficient, enabled new economic

activities which did not exist beforehand, and brought additional new economic growth [2]. Broadband networks have also profoundly changed the social environment in which we live, work and interact, allowing:

- Governments to deliver services more effectively;
- Businesses to operate more efficiently;
- Local businesses or households from smaller or isolated places to connect with larger markets, benefiting largely from the Internet. In fact, the Internet often serves as a pre-requisite for development opportunities, for example supporting short supply chains, accessing new markets, provision of services, education and training, and better quality of life.

Today, ultra-high-speed broadband infrastructure became a determining factor in ensuring the economic fortune of cities and regions. Fiber-to-the-home (FTTH) in rural areas, for example, can contribute to equality of opportunity for all citizens, promoting info-inclusion and the development of human capital, as well as it allows governments to:

- Increase the economic attractiveness of their local region, retain citizens and businesses, and attract new businesses;
- Reduce the urban-rural divide by enabling e-health, e-education, e-government, thus widening access to services typically located in cities;
- Enable people from more isolated regions to have more direct communication with family, friends and colleagues;
- Preserve local culture by promoting local crafts, foods, and performing arts.

Also, the Communication on Connectivity for a European Gigabit Society [3] sets a vision for Europe, where availability and take-up of very high capacity networks will enable

the widespread use of products, services and applications in the Digital Single Market. This vision relies on the three main strategic objectives for 2025: gigabit connectivity, 5G coverage for all urban areas and all primary terrestrial transport paths and access for all European households to connectivity offering at least 100Mbps. In fact, in an increasingly information-based economy, broadband networks are vital arteries of economic activity through which the supply and consumption of information content, services and applications flow [2]. Summing up, better connectivity is the key to boost rural economies, accomplishing the objectives of the European Union to boost rural economies.

### An add-on: IoT

Today, there are more than 7 billion people on the planet, a figure that's expected to reach 9.4 billion by 2050 [4]. To face this scenario, natural resources must be meticulously and carefully managed to make sure everyone survives and has a good quality of life. IoT can

help the world to face this scenario, potentiating a resource management system. The goal of this type of management system is to serve more efficiently and effectively the population, ensuring all necessities are coverage, integrating and making the relationship between demand and supply closer. For example, the European Commission wants to motivate younger people to return to rural areas where the economy is mostly agriculture and livestock-based, and as a consequence, the dependence of natural resources such as water is critical, since farmers depend on this to produce [5]. Thus, implementing the smart management of water is crucial to the local economy.

The IoT exploits the digitisation of everything that humans interact with, processing and correlating information that people cannot easily see and manage [6]. In rural areas, where new digital infrastructures are now available, it is possible to implement the same principles of smart cities to provide better management services, which impacts in people's lives [6] – see **Figure 1**).

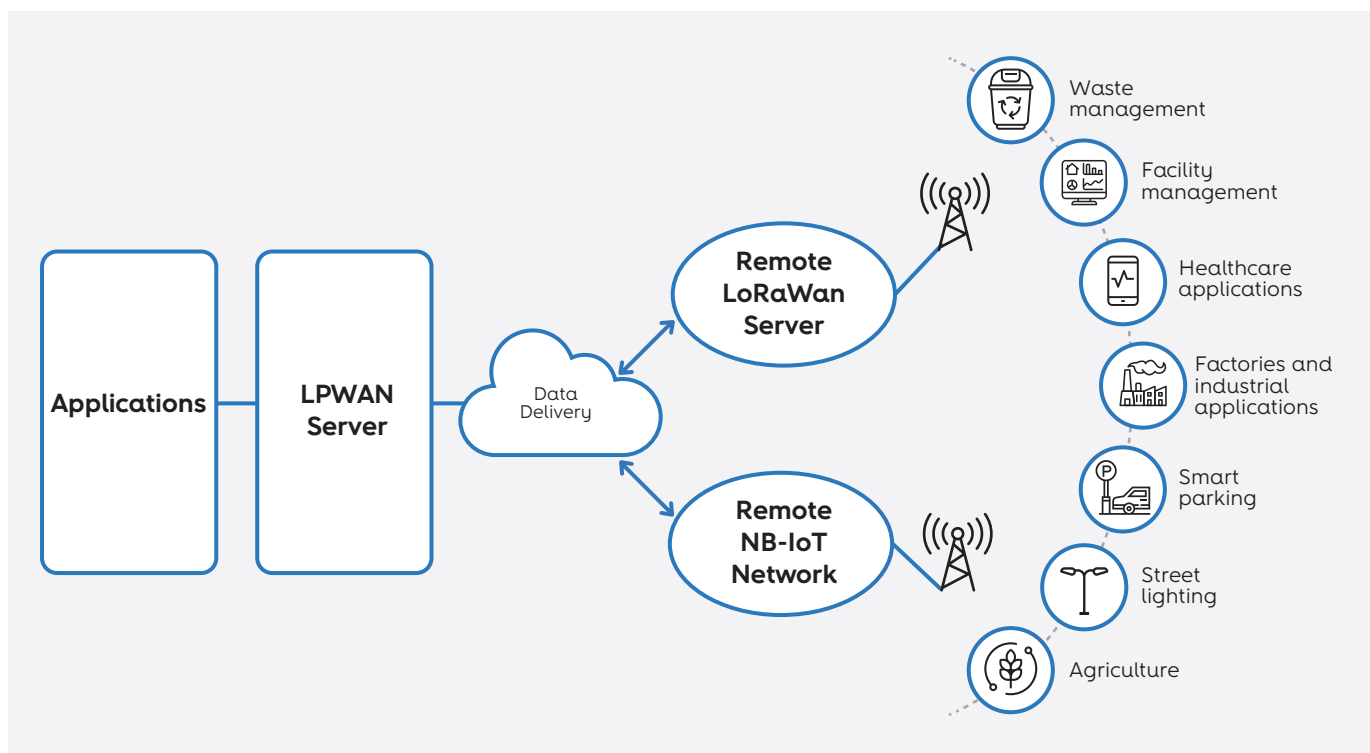


FIGURE 1 – IoT applications

The idea of a hyper-connected society as a consequence of everything being connected and correlated may also be extended to rural zones, helping to attract qualified and highly-skilled people invigorating their economies while, at the same time, potentiating future businesses. In fact, the implementation of smart sensor and actuator networks based on network technologies like 4G-LTE, Wi-Fi, LoRaWAN or NB-IoT opens chances to monitor all aspects of the environment and enables local administrations, based on the diversified information collected by sensors (for example: temperature, wind, ultra-violet index, air quality, security video, river levels, forest fire), analysed in real-time by M2M communications, correlated using several algorithms and visual displayed (see **Figure 2**) to take better decisions, adopting safer and smarter behaviors [7].

## Wrap-up

Superfast broadband networks are helping to increase the attractiveness of rural areas while expanding their economies. The traditional role of regions is changing in all possible scenarios, and that may contribute to attracting young professionals with more diverse interests in technology, business and the environment. In this work is presented, reviewed and analysed the role of superfast broadband and IoT networks to the recovery and development of the economy of the rural areas, helping to improve the life of those who live there in several different ways. Therefore, digital service providers should pay more attention to (the future) smart rural regions as possible markets for their growth.

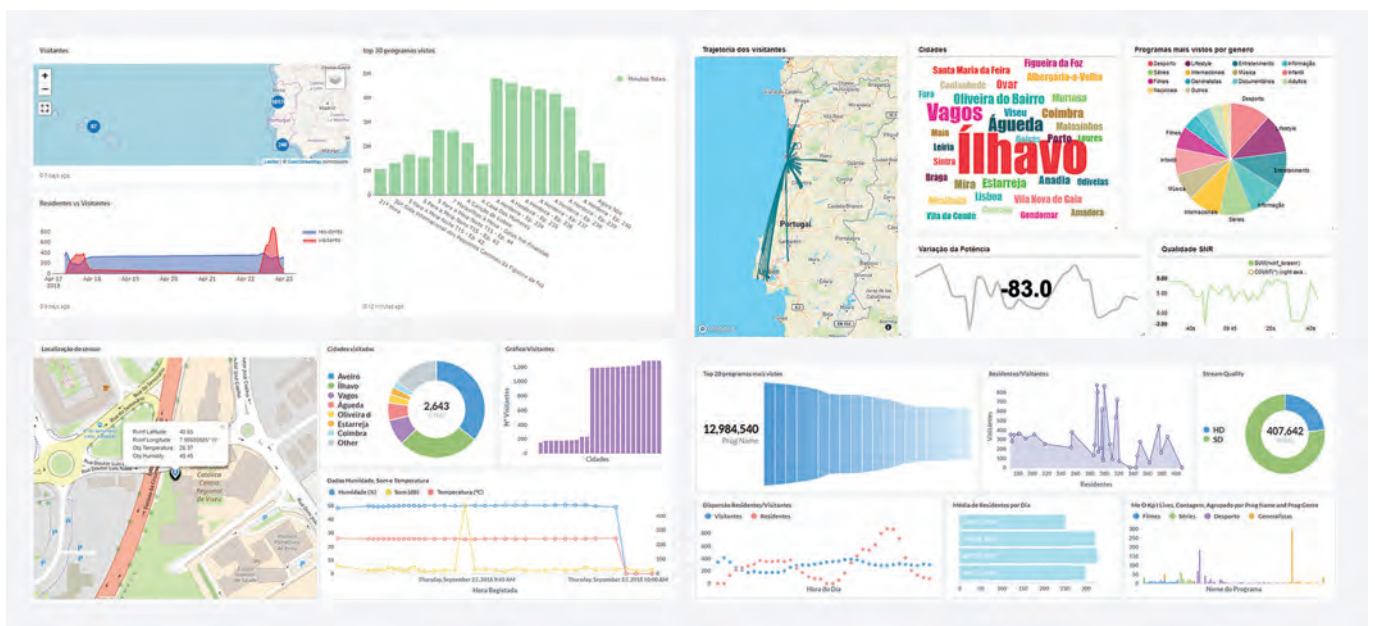



FIGURE 2 – Example of IoT dashboards



---

## References

- [1] The World Bank, "Global Platform for Sustainable Cities," 2016. [Online]. Available: <http://www.worldbank.org/en/topic/urbandevelopment/brief/global-platform-for-sustainable-cities>.
- [2] Commonwealth Telecommunications Organisation, "The Socio-Economic Impact of Broadband in sub-Saharan Africa: The Satellite Advantage," 2015.
- [3] European Commission, "Connectivity for a European Gigabit Society," 2018. [Online]. Available: <https://ec.europa.eu/digital-single-market/en/policies/improving-connectivity-and-access>.
- [4] United Nations, "World Population Prospects: The 2017 Revision," 2017. [Online]. Available: <https://www.un.org/development/desa/publications/world-population-prospects-the-2017-revision.html>.
- [5] European Parliament, "Precision Agriculture and the Future of Farming in Europe - Scientific Foresight Study," 2016. [Online]. Available: <https://publications.europa.eu/en/publication-detail/-/publication/40fe549e-cb49-11e7-a5d5-01aa75ed71a1/language-en>.
- [6] European Commission, "Cross-Cutting Business Models for Internet of Things (IoT) - Final Report," 2018. [Online]. Available: <https://ec.europa.eu/digital-single-market/en/news/cross-cutting-business-models-internet-things-iot>.
- [7] European Commission, "Smart Cities," 2018. [Online]. Available: [https://ec.europa.eu/info/eu-regional-and-urban-development/topics/cities-and-urban-development/city-initiatives/smart-cities\\_en](https://ec.europa.eu/info/eu-regional-and-urban-development/topics/cities-and-urban-development/city-initiatives/smart-cities_en).



# 14

## Use cases: the 5G fuel

**Francisco Fontes, Altice Labs**  
[Fontes@alticelabs.com](mailto:Fontes@alticelabs.com)

5G establishes a new milestone in mobile communications, representing a major disruption when compared with the smooth evolutionary steps previously observed with other generations. Notwithstanding the industry hype, the fact is that 5G, as it is being defined, has the potential to become a general purpose technology, which opens opportunities for providers to find new sources of revenue, compensating a stagnated ARPU, in a saturated and competitive market, where, to avoid churn, more bandwidth and data volumes are continuously being offered at reduced prices. Thus, 5G is expected to prop up an unlimited variety of use cases, touching and impacting all sectors of society and economy.

This article identifies the most promising use cases that will leverage on 5G and the technical characteristics required to address the respective requirements.

---

### Keywords

5G; Use cases; 3GPP; ITU; eMBB; mMTC; URLLC; 5GAA; 5G-ACIA

# Introduction

## Context

Mobile broadband (MBB) has long been a major telecommunications evolution catalyst, driven by the ever-growing demand for bandwidth on the go. 5G will be no exception in this regard. However, considering the past experience, no significant ARPU increase shall be expected in this segment, unless new use cases, addressing new markets, are found.

In this quest for new sources of revenue, 4G has already been enhanced with capabilities to connect “things” to the Internet. Entering the Internet of things territory (IoT), this is the case for cellular IoT (C-IoT), supported by enhanced machine-type communications (eMTC) and narrowband IoT (NB-IoT). This is a new, emerging market, with new actors and ecosystems, with business models still to be settled and balances to be found.

5G is being specified and marketed in this demanding context, with milestones being set by the International Telecommunication Union (ITU) standard, IMT-2020 [1]. 5G emerges at a time of MBB enhancements, C-IoT monetization and massification, and extension to the new critical IoT domain, open by the much lower latency, associated with effectively reliable communications.

## 5G definition

Approximately every 10 years a new generation of wireless, cellular, communications emerges, improving the previous generation in terms of performance. 5G, the next generation, is required to provide considerable performance improvements, supporting a much more diversified spectrum of use cases.

A 5G system builds on a new set of functional entities. A 5G new radio (5G-NR) interface, connects the user equipment (UE) to access nodes (next generation NodeB - gNB), supported by a new core network (5G CN). These components were defined in a way to integrate and exploit

emerging technologies, like software-defined networks (SDN), network function virtualisation (NFV) and multi-access edge computing (MEC), contributing to the required performance, operational efficiency, flexibility and availability.

5G-NR aims at being a universal wireless interface, able to operate in a large range of frequencies (<1GHz and >24GHz) and connecting “smartphones, cars, utility meters, wearables, and much more” [2]. In the same way, the 5G CN aims at being a common core, able to interconnect “at least the new RAT(s), the Evolved E-UTRA, non-3GPP accesses and minimise access dependencies ” [3]. This new core presents a modularised architecture, based on elementary network functions, enabling flexible and efficient network slicing, allowing direct interaction among those functions, via the exposed services, following service-based architecture (SBA) principles [4].

5G will leverage on slicing [5] to create virtual networks over a common, shared infrastructure, each being able to present differentiated behaviour and capabilities, thus contributing to address a wide range of use case requirements. Slicing will be end-to-end, including both the radio interface and the core network.

## Targeted performance

The 5G-NR is one of the most innovative aspects of the technology contributing significantly to the technology performance gains. It incorporates many new features (e.g. massive-MIMO, mmWave, sub-carriers spacing, diversity and high-order modulations), being the enabler for many application scenarios. Those are mandatory to answer the performance enhancements set by ITU-R, including lower latency, more bandwidth, higher moving speeds, enhanced reliability and increased capacity to connect low cost and low power devices.

**Table 1** summarises the objective performance for the technology to be considered IMT-2020 certified [6]. This table includes enhanced mobile broadband (eMBB), massive machine type communications

	Downlink	Uplink	Comments
Peak data rate (Gbit/s)	<20	10	Defined for a single mobile station
User experienced data rate (Mbit/s)	100	50	
Peak spectral efficiency (bit/s/Hz)	30	15	
Average spectral efficiency (bit/s/Hz/TRxP)			TRxP: transmission reception points
• Indoor Hotspot – eMBB	9	6.75	
• Dense Urban – eMBB	7.8	5.4	
• Rural – eMBB	3.3	1.6	
Area traffic capacity (Gbps/m <sup>2</sup> )	10		Indoor Hotspot – eMBB
Latency (ms)			Control plane latency: transition time from a most “battery efficient” state (e.g. Idle state) to the start of continuous data transfer
• User plane	4 (eMBB) / 1 (URLCC)		
• Control plane	Max 20 / Recomm. 10		
Connection density (devices per km <sup>2</sup> )	1 000 000		mMTC scenarios
Reliability (success probability)	Min 1-10 <sup>-5</sup>		URLCC scenarios
Mobility (km/h)	500		High speed vehicular (rural eMBB)
Bandwidth	Up to 1GHz, in high-frequency bands (>6GHz); min 100MHz		Maximum aggregated system bandwidth

**TABLE 1** – Minimum requirements related to technical performance for IMT-2020 radio interface(s) [6]

(mMTC) and ultra reliable, low latency communications (URLLC) requirements, scenarios that will be further detailed in the next sections.

## 5G calendar

5G will come in two phases, aligned with 3GPP specifications calendar: phase 1 (Release 15), finished in June 2018 and phase 2 (Release 16), expected to be ready by July 2019 [7].

Release 15 provides the foundational system architecture, radio interface, protocols, procedures

and basic features for a first deployment and exploitation of the technology, being “forward compatible”, easing the addition of future features in Release 16 and beyond. Even though eMBB is the focus of Release 15, new innovative massive and critical C-IoT use cases may start to be exploited, for instance, based on the provided low latency. The full range of requirements and enabled use cases will only be addressed by Release 16, the one to be used for IMT-2020 standardisation submission, by the end of 2019. Release 16 will incorporate several enhancements, like multi-connectivity and data duplications (reliability) enhancements.

Thus, while the first phase focuses on enabling most exciting mobile broadband experiences and partially enhancing existing C-IoT functionalities, the second phase will take cellular connectivity to a variety of industries, transforming many, and potentially creating others.

As an intermediate milestone, 3GPP approved in December 2017 the Non-Standalone Architecture (NSA) specifications, with 5G-NR being used to provide additional downlink connectivity, but supported by an evolved packet core (EPC) / long term evolution (LTE) control and data plane. This allows, for instance, 5G based fixed-wireless accesses (FWA) to be deployed, competing with fibre based ones. NSA has several options being option 3 the one mentioned here. More information can be found in [8].

Aligned with the standardisation calendar, Release 15 based products will start emerging during 2019. Meanwhile, several trials and experimentations are taking place all around the globe. Europe is very active, with several

initiatives taking place. The overall pan-European 5G trials roadmap and relevant standardisation, regulatory and ecosystems time plan are summarised in **Figure 1** [9].

## 5G use cases

5G aims at fuelling all use cases that need cellular connectivity, via 5G-NR. However, it goes even further than that, when defining a 5G core network (5GC) able to connect other wireless (including non-3GPP, e.g. wireless local area network - WLAN) and wired access technologies. Thus, 5G targets a wide range of use cases, even those suited to be served by fixed technologies (e.g. gigabit passive optical networks - GPON), like remote surgery. The next sub-section describes some of the use cases expected to be supported by 5G and how they have been organised by ITU's radiocommunication sector (ITU-R).

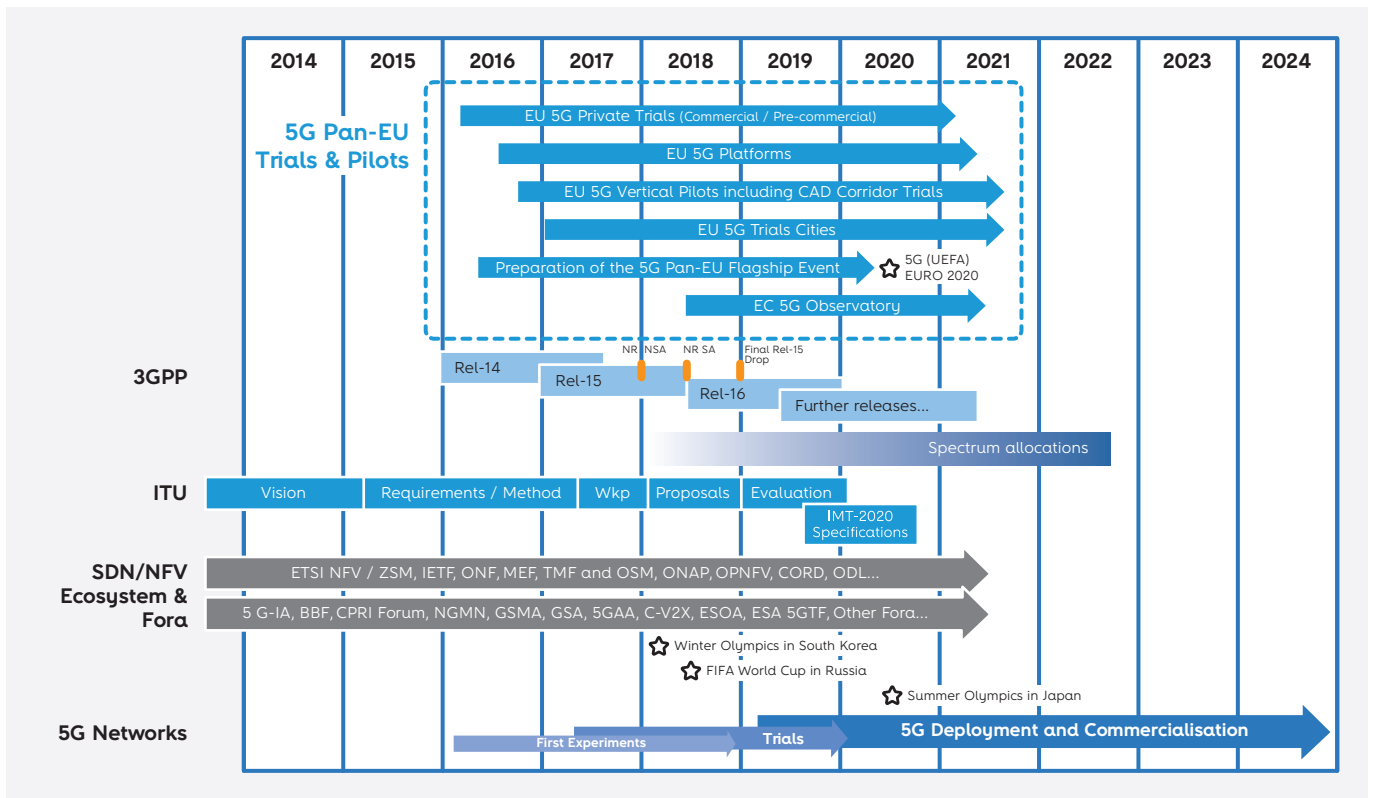


FIGURE 1 – 5G pan-European trials roadmap – time plan

## Use cases as defined by ITU-R

5G standardisation work at ITU started with the identification of the use cases to be supported and the definition of required performance (as explained in the targeted performance section), resulting in a clear need to go far beyond what 4G technology (IMT-Advanced) offers as of today. Those usage scenarios, following ITU’s telecommunication standardisation sector (ITU-T) terminology, were organised according to three main axes, depicted in **Figure 2** [10]:

- 1. Enhanced mobile broadband (eMBB):** use cases exploiting throughput and higher mobility speeds, targeting human-centric communications; lower latency may also apply.
- 2. Massive machine type communications (mMTC):** use cases characterised by a large

number of connected devices, low bandwidth and data volumes requirements, targeting low-cost devices communications; may also be referred to as massive IoT.

- 3. Ultra reliable, low latency communications (URLLC):** use cases for device communications that require very low latency and highly reliable communications, for remote actions execution and feedback; may also be referred to as critical IoT.

This use cases taxonomy became the reference for the industry, including the work being conducted in 3GPP. Accordingly, three service/slice-type (SST) values have been specified [11], with values 1, 2 and 3, respectively.

The identified usage scenarios present very distinct performance requirements. For instance, requirements for HD video streaming to a smartphone while commuting are very different

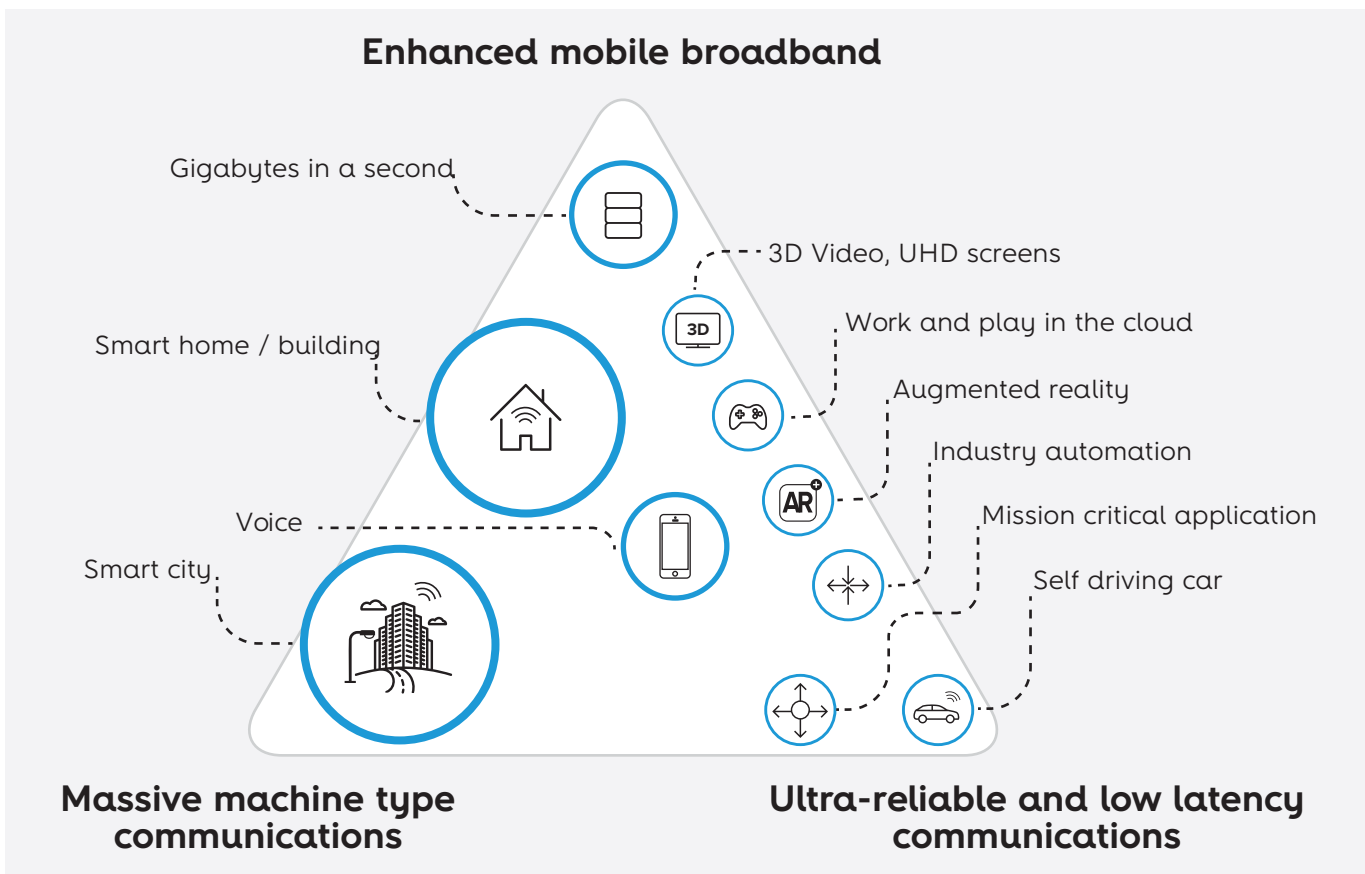


FIGURE 2 – 5G identified usage scenarios organisation [10]

from the ones related to weekly water meters readings located some floors underground. One of the main challenges for 5G is to concurrently support such a diverse range of use cases, reliably. This is where diversified radio spectrum, spanning from low to high frequencies, and 5G end-to-end network slicing capabilities, are required.

Next sections go further in the several use cases analysis, using ITU-R taxonomy.

## Enhanced mobile broadband

MBB has been the focus of previous wireless generations, providing fast, mobile, access to services. Their major beneficiaries have been over-the-top (OTT) service providers, exploiting the provided best-effort connectivity facet. Even if 3G and 4G are enabled with QoS mechanisms, no significant services use it, except operators' voice service on IP [12].

Due to its large massification, MBB will continue being key to the wireless industry businesses. The provided 5G new features and performance improvements will open opportunities for other MBB services and applications. Relevant aspects are increased bandwidth, higher mobility, lower latency, higher users' density and better location accuracy. eMBB use cases have the opportunity to go beyond basic mobile Internet, exploiting interactive services, rich media and entertainment, and tactile Internet. Video-based services are expected to continue being the main drivers for bandwidth consumption.

## Massive machine type communications

Traditional cellular data communications, such as LTE, require complex devices and drain significant energy from their batteries due to their operation mode, not optimised for the transmission of infrequent small amounts of data and low or no mobility at all (e.g. meters for reading water levels or gas and electricity consumption). MTC, the 3GPP name for C-IoT communications, answer the needs, summarised in four Cs, to appropriately connect "things": coverage (radio network), current

(battery drain), capacity (network as an all), cost (devices and services).

In that scope, 3GPP added specific features to existing "generations", via extended coverage GSM IoT (EC-GSM-IoT), eMTC and NB-IoT. Besides 3GPP work, other non-3GPP solutions exist, like LoRa [13] and SigFox [14], all belonging to low power wide area (LPWA) technologies. Before, this market was served only by personal area network (PAN) technologies, such as Bluetooth, and WLAN technologies. Each has their own use cases and market [15].

3GPP will propose to ITU-T both eMTC and NB-IoT as 5G standards, guaranteeing that 5G-NR provides mechanisms to support the coexistence between eMTC/NB-IoT and NR on the same carrier. Thus, similarly to MBB, 5G will enhance existing MTC features, increasing the capacity for higher devices and connections density, with even better power efficiency.


Expected use cases in this area are based on current 4G use cases, with a likely widespread adoption of IoT in the houses, cities, agriculture and industry in general. Besides the devices and connectivity, the full IoT business exploitation requires the integration with a platform able to collect all the data, process it, provide indicators and execute actions.

## Ultra reliable, low latency communications

In addition to eMBB and mMTC, 5G is expected to support mission-critical services, enabled by low latency and increased reliability. This opens a new area, not addressed before by cellular wireless communications. Thus, the most innovative 5G application areas are expected to emerge here.

Industry 4.0 is characterised by the incorporation of ICT technologies in their processes. Wireless communications, in the form of 5G, is just one of those. A suitable wireless technology has the potential to enhance the industry operations and businesses, fostered by features like the





**FIGURE 3** – Exemplary application areas of 5G in the factory of the future [16]

nonexistence of cables in the factory floor and unlimited indoor and outdoor mobility (see examples in **Figure 3**).

5G may have a significant role as the elected wireless communications technology, both outdoor and indoor. For instance, 3GPP Release 16 shall incorporate time sensitive networking (TSN) features [17].

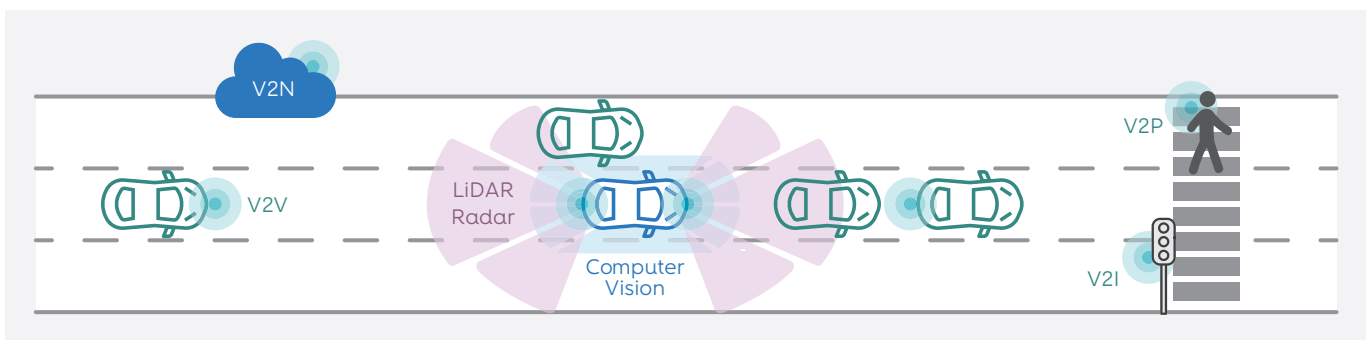
Two of the major economic sectors, automotive and factories, have already organised to exploit the new characteristics of 5G. In the next sections, they will be used to exemplify what is expected to get from URLLC 5G use cases.

### 5G role for the automotive industry

While artificial intelligence (AI) is basilar for driverless, autonomous cars, 5G will enable the creation of real advanced and disrupting

experiences, only possible with large coverage, low latency, reliable and fast wireless connectivity. “All of the smart cars on the market today have one big weakness in common; they rely solely on their own sensors and cameras for manoeuvring and navigation. The new 5G mobile networks, however, have been designed with the automotive industry in mind” [18].

The 5G Automotive Association (5GAA) [19] focuses on cellular vehicle-to-everything (C-V2X) and is a reflection of the 5G emerging ecosystems. Their use cases fit in the vehicle-to-everything (V2X) paradigm: vehicle-to-vehicle (V2V), vehicle-to-pedestrians (V2P) and vehicle-to-infrastructure (V2I) communications. Besides these shorter-range device-to-device (D2D) scenarios, vehicles communicating with the network (V2N), using traditional cellular links (C-V2X), are also considered in **Figure 4**.



**FIGURE 4** – V2X environment [20]

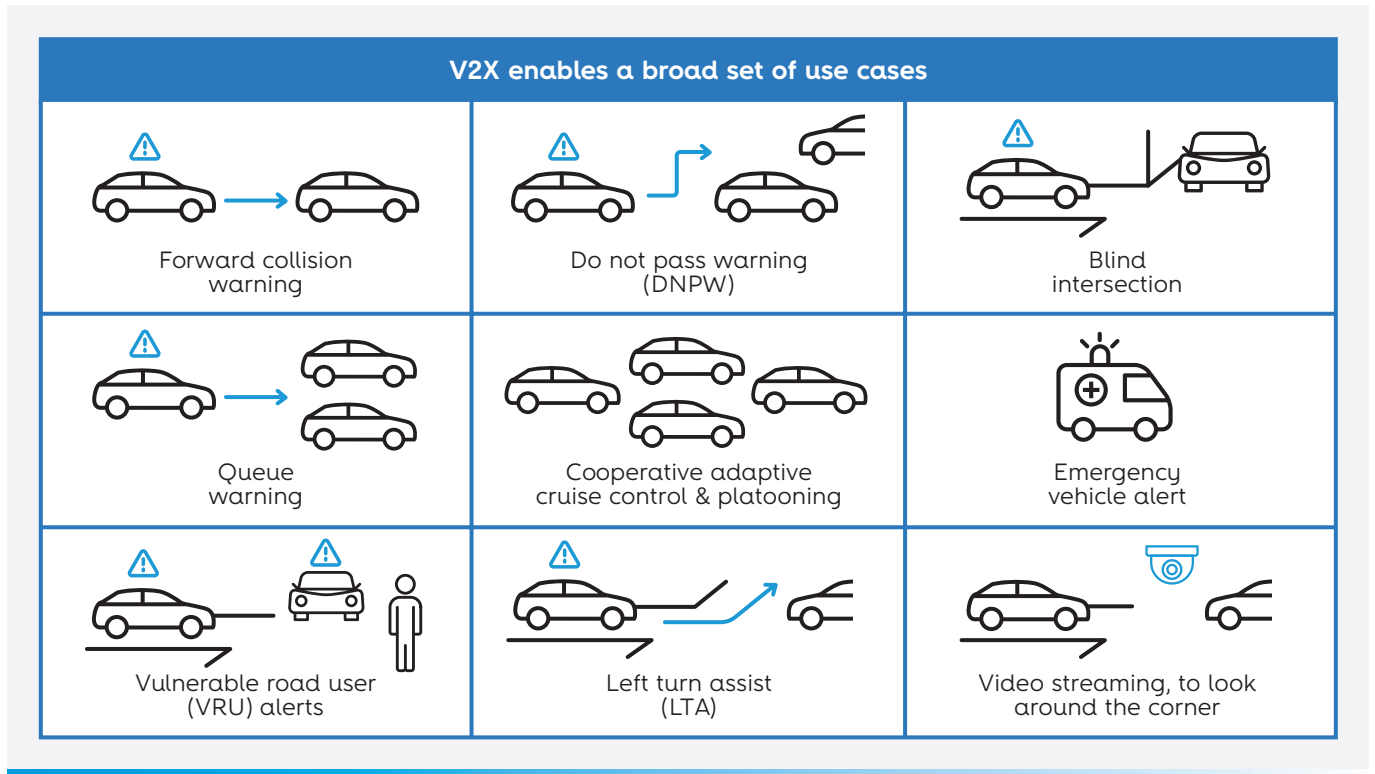


FIGURE 5 – V2X example use cases

Besides infotainment, which fits in eMBB scope, a wide range of use cases emerge, as shown in **Figure 5**.

3GPP in [21] and [22] identifies V2X requirements, to be answered by 4G and 5G, which are summarised in **Table 2**. Additional requirements for other use cases can be found in [23].

### 5G role in the industry

Wired infrastructures have been so far dominant in industrial environments. Furthermore, for some specific control applications, Ethernet-

based solutions, like EtherCAT, with an ongoing trend toward TSN, are adopted. However, there are advantages associated with the use of wireless technologies in factories and industrial environments, especially 5G.

The usage of ICT technologies in industrial environments, supported by 5G low latency and high reliability, will be one of the most innovative and demanding aspects of 5G. That is the focus of the 5G Alliance for Connected Industries and Automation (5G-ACIA) [24], stating that “promising application areas range from logistics for supply and inventory management, through robot and

	Max end-to-end latency (ms)	Reliability (%)	Min required communication range (meters)
Vehicles platooning	10 – 500	90 – 99.99	80 – 350
Advanced driving	3 – 100	90 – 99.999	360 – 700
Extended sensors	3 – 50	90 – 99.999	50 – 1000

TABLE 2 – Requirements for some V2X use cases

Industry vertical	Application
Smart factory/industrial automation	Industrial control, robot control, machine to machine, process control
Healthcare industry	Remote diagnosis, emergency response, remote surgery
Entertainment industry	Immersive entertainment, online gaming
Transport industry	Driver assistance applications, enhanced safety, autonomous driving, traffic management
Manufacturing industry	Motion control, remote control, AR and VR applications
Energy sector	Smart energy, smart grid

**TABLE 3** – Industries use cases

motion control applications, to operations control and the localisation of devices and items” [25].

Some of the foreseen industry use cases are shown in **Table 3** [26].

For some of the identified smart factory/industrial automation use cases, the following **Table 4** summarises the main requirements and also characterises foreseen deployments [16].

Latency is indirectly referred in “cycle time”, which measures the time required from issuing a command until an answer is received about the completeness of its execution. Thus, latency, being unidirectional, must be less than half of the indicated cycle time value. Even if 3GPP Release 15 is expected have the capability to achieve a latency of 1ms with a reliability of 99.999%, Release 16 is needed to provide the required higher reliability and efficiency.

Use case (high level)		Availability	Cycle time	Typical payload size	# of devices	Typical service area
<b>Motion control</b>	Printing machine	> 99.9999%	< 2ms	20 bytes	>100	100m x 100m x 30m
	Machine tool	> 99.9999%	< 0.5ms	50 bytes	~20	15m x 15m x 3m
	Packaging machine	> 99.9999%	< 1ms	40 bytes	~50	10m x 5m x 3m
<b>Mobile robots</b>	Cooperative motion control	> 99.9999%	1ms	40-250 bytes	100	< 1km <sup>2</sup>
	Video-operated remote control	> 99.9999%	10 - 100ms	15-150 kbytes	100	< 1km <sup>2</sup>
<b>Mobile control panels with safety functions</b>	Assembly robots or miling machines	> 99.9999%	4 - 8ms	40-250 bytes	4	10m x 10m
	Mobile cranes	> 99.9999%	12ms	40-250 bytes	2	40m x 60m
<b>Process automation (process monitoring)</b>		> 99.99%	> 50ms	Varies	10 000 devices per km <sup>2</sup>	

**TABLE 4** – Industrial use case requirements

3GPP, in TR 22.804 [27], also identifies eight vertical domains of interest, as follows: rail-bound mass transit, building automation, factory of the future, e-health, smart city, electrical power distribution, central power generation, programme-making and special events (PMSE) and smart farming.

LTE- and WLAN-based alternatives exist and are already being used in less demanding situations, but these do not provide the 5G lower layers with the performance required by demanding industrial environment, where low latency (cycle time), jitter, and high reliability/availability are mandatory. In such environments, availability and reliability may have a strong impact on production.

## 5G business opportunities

5G being, per se, just another connectivity technology, requires significant investments (e.g. in a heterogeneous spectrum, dense cells deployment, high capacity transport network, and distributed data centres, from core to edge) to support the full range of expected use cases. To exploit those investments, operators must take an active approach, guaranteeing their involvement and share on the emerging new market. A shift is expected from the B2C market (where OTT service providers are the main beneficiaries of the operators' investment), to the B2B market, where service value is based in service characteristics (e.g. low latency), rather than on duration and volume.

As shown in previous sections, 5G is expected to enable use cases and services in new areas, supported by the new technology capabilities. Building new eco-systems and business models, the expected economic impact of 5G is significant, according to several studies conducted by such entities as IHS, Ericsson and ADL. According to IHS, by 2035, the 5G value chain will be the direct responsible for 3.5 trillion USD and enable a total of 12.3 trillion USD for global economic output; by that time it will support 22 million jobs [28].

Similarly, a report from Ericsson and ADL estimates that by 2026 (first years of full 5G phase 2 exploitation), from a total estimation of 619 billion USD, manufacturing presents the best revenue opportunities for telecom operators (113 billion USD), closely followed by energy and utilities (101 billion USD). On the lower positions are retail, with 29 billion USD, and agriculture, with 9 billion USD [29].

However, while waiting for industries to understand, join and exploit 5G, the initial 5G deployment period (until 2023) will be mostly for common citizens (eMBB) to just upgrade from 4G to 5G speeds. Thus, no significant growth in the number of subscriptions is expected, and ARPU may well continue to drop. mMTC and URLLC industry use cases will have different associated business models, and operators must be part of them. IoT, edge and other service platforms, content and applications, must be part of operators' portfolios, to capitalise 5G investment. Currently, operators own one asset, which OTT service providers and other cloud providers are not able to offer: geographical proximity to customers. But for this to happen, the investment strategy in gigantic data centres, providing big storage and processing power, needs to be now complemented with smaller, distributed edge data centres, enabling the exploitation of the low latency services.

## Preparing for 5G

New business ecosystems are expected to emerge, especially as a result of the B2B market expansion. New value chains, integrating consumers, service and network providers, will be created. Markets will be self-organised in verticals, as the mentioned 5GAA and 5G-ACIA already denote. Representing a large number of users and entities, these will have a significant negotiation power facing mobile operators, infrastructure providers and standardisation. Tailored connectivity offers are

expected, complemented by appropriate service platforms, where the value will be in other service characteristics, rather than just in speed or data volumes.

The creation of 5G related associations (e.g. 5GAA and 5G-ACIA), emerges as a demonstration of 5G ecosystems being built beyond the traditional relationships between Internet service providers (ISP)/telecom operators and OTT companies. 5G and the massive/critical IoT world are different and require different approaches. Even if these new organisations have the power to assemble almost all the required components, answering the needs of the target activity, they need tailored connectivity with the outside world. Operators must integrate these ecosystems, bringing other arguments they only have like support for edge computing.

In fact, some of the expected performance enhancements have implications beyond 5G. While the increased data rates require operating at higher frequencies, which implies an ultra-dense network (UDN) via the deployment of small cells, low latency requires the deployment of data centres close to cell sites. Thus, 5G will require significant changes in supporting infrastructure. Much more data will be exchanged with cell sites and more capillarity to reach dense and distant antennas need to be established at low costs. This will be especially challenging for indoor, where most of the MBB services (e.g. HD video) are consumed, requiring operation at higher frequencies (mmWave).

Second, generic computational resources are required to run services and network functions as software modules, at different locations. The need for edge computing, flexibility and agility, will push virtualisation and computational power to be distributed along the path from core data centres, to very close to users, at central offices and even cell sites.

Operators still deploying 4G, and struggling to reduce capex every year, will put pressure on the reduction of 5G total cost of ownership (TCO). A fully end-to-end programmable network, leveraging on NVF, SDN, edge computing,

cloud radio access network (RAN) and network slicing, will be key for operators to define a 5G deployment strategy and make money out of 5G.

B/OSS must be very efficient, flexible and operate in much faster, closed, automated loops, powered by artificial intelligence (AI) and machine learning (ML). TM Forum appoints its Open Digital Architecture (ODA) as providing the design principles fundamental for an IT infrastructure, able to efficiently monetize, with no silos, on the emerging 5G era. TM Forum indicates that *“up to 72% of 5G revenue growth is dependent on transformation of operational and business support systems (OSS/BSS)”* [30]. Besides the contribution to optimisation and self-healing mechanisms, some new businesses, for instance, the new ones to be exploited around 5G slicing capabilities.

The embedded universal integrated circuit card (eUICC), also known as embedded subscriber identification module (eSIM), presents large advantages, one of them being the capacity to easily be used to equip terminals which traditionally do not have a SIM card, as is the case of IoT. Additionally, eSIM-equipped devices may explore the capacity to manage several operator profiles simultaneously. Besides laptops, many devices are already being shipped with eSIM, like Apple’s SmartWatch, the Google Pixel 2 phone and the Microsoft Surface Pro tablet.

## Final remarks

According to the Gartner Hype Cycle from August 2018 [31], 5G is entering the peak of inflated expectations phase. This is reflected in 5G wide span of targeted use cases, involving wireless and also wired communications. However, early 5G deployments will be an evolution of what we know today: even more bandwidth on the go. First use cases, rely on 5G-NR just for the data plane, resorting to 4G for the control plane and core network. Fixed wireless accesses are the main driver for this.

According to [32], 5G rollout models are crystallising in the following trends:

- **Gigabit broadband to the home:** high-speed, fixed, broadband service to homes;
- **NextGen mobile user experience:** enhanced mobile broadband (e.g. AR/VR) services for individuals;
- **Future corporate networks:** high reliable, low latency, high speed, services to individual enterprises;
- **Digital industrial ecosystems:** bringing together services from different providers and industries;

- **NextGen infra-as-a-service:** provide global services to other operators.


5G eye-catching use cases, like tactile Internet, factory automation, remote machines operation and remote surgeries, still need significant work on specification and experimentation, until the technology is cost-effective and mature enough to be trustable for such applications. Digital transformation, being generally observed, constitutes an important driver for 5G adoption.

5G deployments do not depend solely on technology availability but also its support. Thus, NFV maturity, transport capacity and data centres distribution are needed to exploit the full potential of 5G.

## References

- [1] International Telecommunications Union, "ITU towards 'IMT for 2020 and beyond,'" *ITU-R*, 2018. [Online]. Available: <https://www.itu.int/en/ITU-R/study-groups/rsg5/rwp5d/imt-2020>.
- [2] M. Branda, "Qualcomm takes a significant step toward making 5G NR a reality," *Qualcomm Technologies, Inc.*, 2016. [Online]. Available: <https://www.qualcomm.com/news/onq/2016/06/29/qualcomm-takes-significant-step-toward-making-5g-nr-reality>.
- [3] 3GPP, "TR 23.799 - Study on Architecture for Next Generation System (Release 14)," 2016.
- [4] 3GPP, "TR 29.891 - 5G System - Phase 1; CT WG4 Aspects (Release 15)," 2017.
- [5] T. Toving, "Management, Orchestration and Charging for 5G networks," *3GPP*, 2018. [Online]. Available: [http://www.3gpp.org/NEWS-EVENTS/3GPP-NEWS/1951-SA5\\_5G](http://www.3gpp.org/NEWS-EVENTS/3GPP-NEWS/1951-SA5_5G).
- [6] ITU-R, "M.2410-0 - Minimum requirements related to technical performance for IMT-2020 radio interface(s)," 2017.
- [7] A. Weissberger, "Busting a Myth: 3GPP Roadmap to true 5G (IMT 2020) vs AT&T 'standards-based 5G' in Austin, TX," *IEEE Communications Society*, 2018. [Online]. Available: <http://techblog.comsoc.org/2018/07/08/busting-a-myth-3gpp-roadmap-to-true5g-att-standards-based-5g-inn-austin-tx/>.
- [8] GSM Association, "Road to 5G: Introduction and Migration," 2018.
- [9] 5G Infrastructure Association, "5G Pan-European Trials Roadmap Version 3.0," 2018.
- [10] ITU-R, "M.2083-0 - IMT Vision – Framework and overall objectives of the future development of IMT for 2020 and beyond," 2015.

- [11] 3GPP, "TS 23.501 - System Architecture for the 5G System; Stage 2 (Release 15)," 2018.
- [12] GSM Association, "IR.92 - IMS Profile for Voice and SMS, Version 12.0," 2018.
- [13] "LoRa Alliance," 2018. [Online]. Available: <https://lora-alliance.org/>.
- [14] "Sigfox," 2018. [Online]. Available: <https://www.sigfox.com>.
- [15] Altice Labs, "IoT Cellular Networks," 2017.
- [16] 5G-ACIA, "5G for Connected Industries and Automation (White Paper - Second Edition)," 2018.
- [17] "Time-Sensitive Networking Task Group," 2018. [Online]. Available: <http://www.ieee802.org/1/pages/tsn.html>.
- [18] A. Tøgard, "The impact of 5G: How will 5G affect the automotive industry and mobility?," *Wallenius Wilhelmsen ASA*, 2017. [Online]. Available: <https://www.2wglobal.com/news-and-insights/articles/features/the-impact-of-5g-how-will-5g-affect-the-automotive-industry-and-mobility/>.
- [19] "5G Automotive Association (5GAA)," 2018. [Online]. Available: <http://5gaa.org/>.
- [20] 5G Automotive Association, "The Case for Cellular V2X for Safety and Cooperative Driving," 2016.
- [21] 3GPP, "TS 22.185 Service requirements for V2X services; Stage 1 (Release 15)," 2018.
- [22] 3GPP, "TS 22.186 - Enhancement of 3GPP support for V2X scenarios; Stage 1 (Release 16)," 2018.
- [23] M. Boban, A. Kousaridas, K. Manolakis, J. Eichinger, and W. Xu, "Use Cases, Requirements, and Design Considerations for 5G V2X," *IEEE Vehicular Technology Magazine*, p. 10, 2018.
- [24] 5G-ACIA, "5G Alliance for Connected Industries and Automation (5G-ACIA)," 2018. [Online]. Available: <https://www.5g-acia.org/>.
- [25] 5G-ACIA, "How can it be used?," 2018. [Online]. Available: <https://www.5g-acia.org/index.php?id=5801>.
- [26] 5G Americas, "New Services & Applications With 5G Ultra-Reliable Low Latency Communications," 2018.
- [27] 3GPP, "TR 22.804 - Study on Communication for Automation in Vertical Domains (Release 16)," 2018.
- [28] IHS Economics & IHS Technology, "The 5G economy: How 5G technology will contribute to the global economy," 2017.
- [29] Ericsson, "The 5G Business Potential, Second Edition," 2017.
- [30] C. Haslam, "5G monetization: Operational imperatives," *TM Forum*, 2018. [Online]. Available: <https://inform.tmforum.org/research-reports/5g-monetization-operational-imperatives/>.
- [31] K. Panetta, "5 Trends Emerge in the Gartner Hype Cycle for Emerging Technologies, 2018," *Gartner*, 2018. [Online]. Available: <https://www.gartner.com/smarterwithgartner/5-trends-emerge-in-gartner-hype-cycle-for-emerging-technologies-2018/>.
- [32] K. Taga, R. Swinford, and G. Peres, "5G deployment models are crystallizing." *Arthur D. Little*, p. 26, 2017.





# 15

## The digital transformation of the central office

**Rui Calé Martins, Altice Labs**

[cale@alticelabs.com](mailto:cale@alticelabs.com)

**André Brízido, Altice Labs**

[andre-d-brizido@alticelabs.com](mailto:andre-d-brizido@alticelabs.com)

**Miguel Santos, Altice Labs**

[miguel-c-santos@alticelabs.com](mailto:miguel-c-santos@alticelabs.com)

In today's passive optical networks and over all the broadband access network structure, performance and reliability have reached paramount levels, but the new digital economy is also striving for flexibility and agility. This and fierce competition are pushing hard for a new approach, which is being fuelled by various aspects of network softwarization, with a strong impact on a sensitive cornerstone for network operators - the central office. This article provides an overview of Altice Labs' insight into this ecosystem and its proposals for the new cloud central office, particularly for the PON access.

---

### Keywords

Central office; CloudCO; SDAN; BAA; Virtual fiber box; PON

## Introduction

Today, central office (CO) is a major bottleneck preventing communications services providers (CSP) to be more agile. Several access networks are usually terminated there, each one with its own set of equipment serving their own silo, operated and managed by different teams (even when they are performing the same functions), all this in a cluttered environment.

Nevertheless, the very existence of an entity such as the CO represents a major asset for CSP in a digital services world and its placement at the edge of the access networks may even become a new business opportunity by monetizing proximity to end-users. For example, next-generation applications crave for low latency and high bandwidth cloud access as close as possible to the end user, on the edge of the networks. Who else will be able to host a function that is close enough to network terminals as to provide ultra-low latency response? To be able to explore new business opportunities, the CO needs a major transformation to make it more agile and flexible in the termination of access networks and first aggregation stages. To address such transformation, softwarization of the CO is a key factor that can be achieved through software-defined networks (SDN) and network function virtualisation (NFV).

This article presents the next generation CO, the current state of the art and what are Altice Labs' proposals in this area together with the planned efforts to make them a reality.

## A common view for access networks

### A new network structure

Softwarization [1] is already starting to change networks. As many network functions gradually

stop being boxed in network elements and emerge as disaggregated software running on IT infrastructure, functions are recombined into a new physiognomy for networks, where IT-centered CO work as a common termination for access networks and connect to a number of networked data centres, forming a cloud that extends from the edge to the core of networks.

## The central office

The CO is the main communications aggregation point, and it has to evolve to support the network of the future and the new wave of services requiring very-high access speeds and ultra-low latency. Factors like next generation access networks, the convergence of wireless and wireline networks, virtualisation through NFV and SDN, and the need for flexibility in the operation are guiding the evolution of the CO.

Thus, the CO is the space where the advantages of softwarization become evident for access networks: common network functions running in software on standard IT infrastructure may serve passive optical networks (PON), cable, mobile or other access networks, regardless of their different natures.

To materialise this vision, an SDN/NFV based approach is proposed (**Figure 1**):

- Specific physical aspects of access networks are addressed by the software-defined access nodes (**SDAN**), like the SDAN-OLT or the SDAN-BBU. This physical element converts the access media to Ethernet, interfacing to the central office switching fabric and performing data plane and the most basic control plane functions;
- Specific adaptation and control logic that can be run outside of the SDAN is delivered by an **hardware abstraction** module which is able to expose the controlled resources to an SDN controller. In the case of xPON networks it allows to control an optical line termination (OLT) and its optical network units (ONU) as a single L2/L3 node;

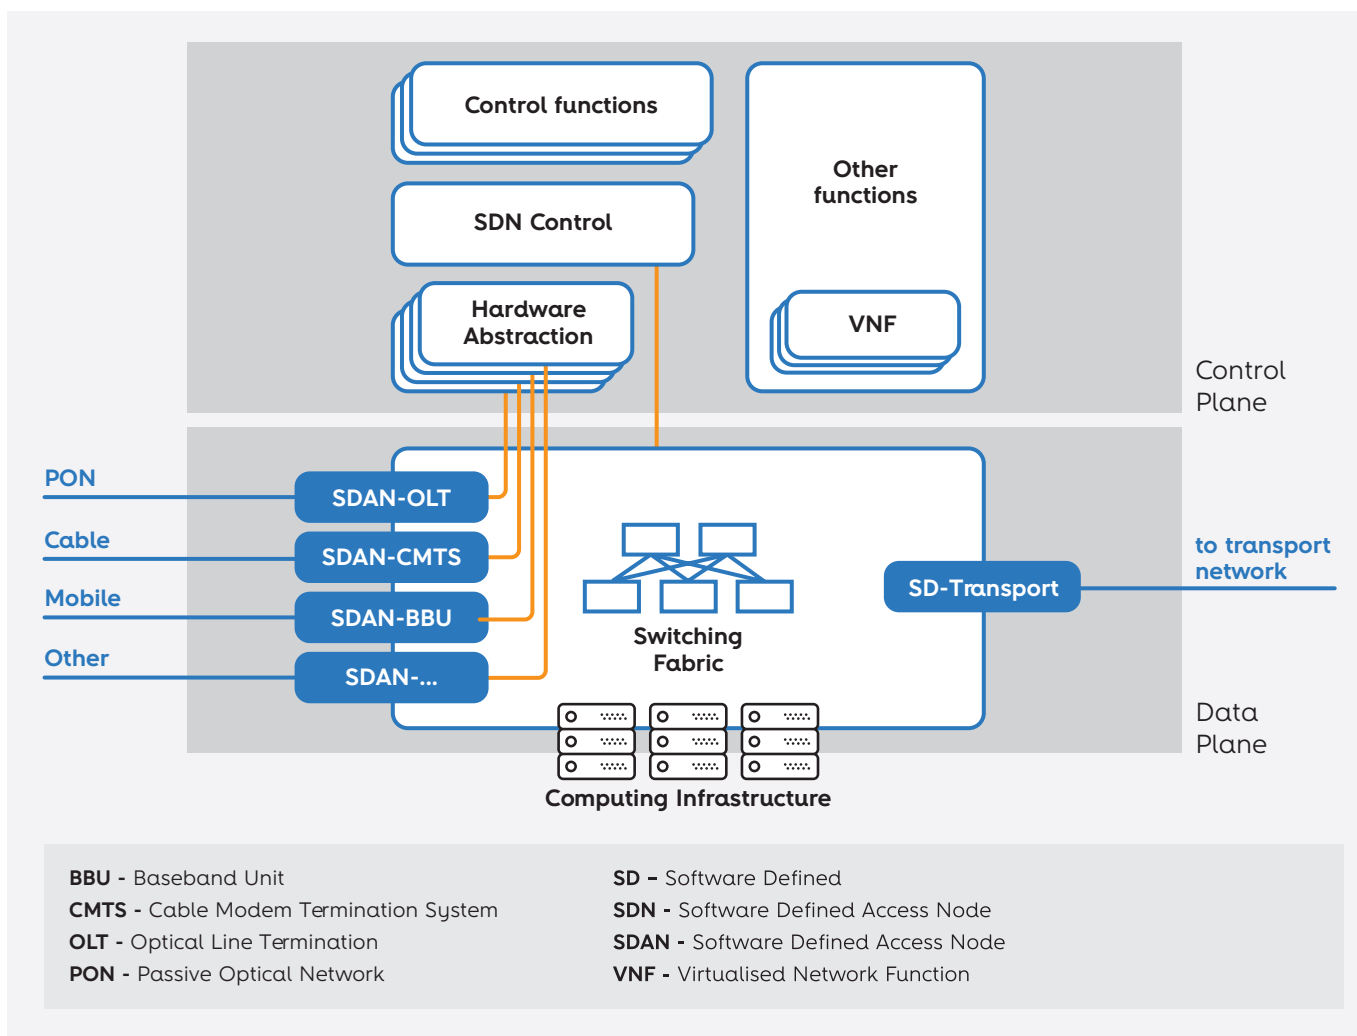


FIGURE 1 – Central office for multiple access networks

- The “heart” of the control plane is the **SDN controller** (or controllers). It exposes the access network, as well as the data plane capabilities of the CO to SDN applications that implement service control. These applications may be at different integration levels within the SDN controller, but in general, one can say that the SDN controller presents an API to the network;
- Most control functions will be virtualised and disaggregated, running on the computing infrastructure. Some data plane functions may also be virtualised;
- Other functions, usually virtualised network functions (**VNF**), may be orchestrated to be instantiated both in the data plane and in the

control plane, performing various roles that span from service-specific network functions to 3rd party functions like customer VNF, e.g. those that have to be run on the network edge to enforce ultra-low latency requirements. The softwarization of many functions is making it possible to distribute them more rationally, dynamically running them where (and when) they are more needed, more efficient, or less expensive, be it at the CO, customer premises equipments, or core data centres;

- The transport network may follow the same architecture: a software-defined data-plane element under the control of an SDN controller. The transport network is outside the scope of this article.

Not represented here, but still extremely important are the orchestration functions that are needed to manage the lifecycle of virtualised functions and services. The orchestration functions are outside the scope of this article.

The next section sums up the approach that the Broadband Forum (BBF) [2] is standardising, as well as other relevant initiatives for the central office evolution into what is called the cloud central office (CloudCO).

## State of the art

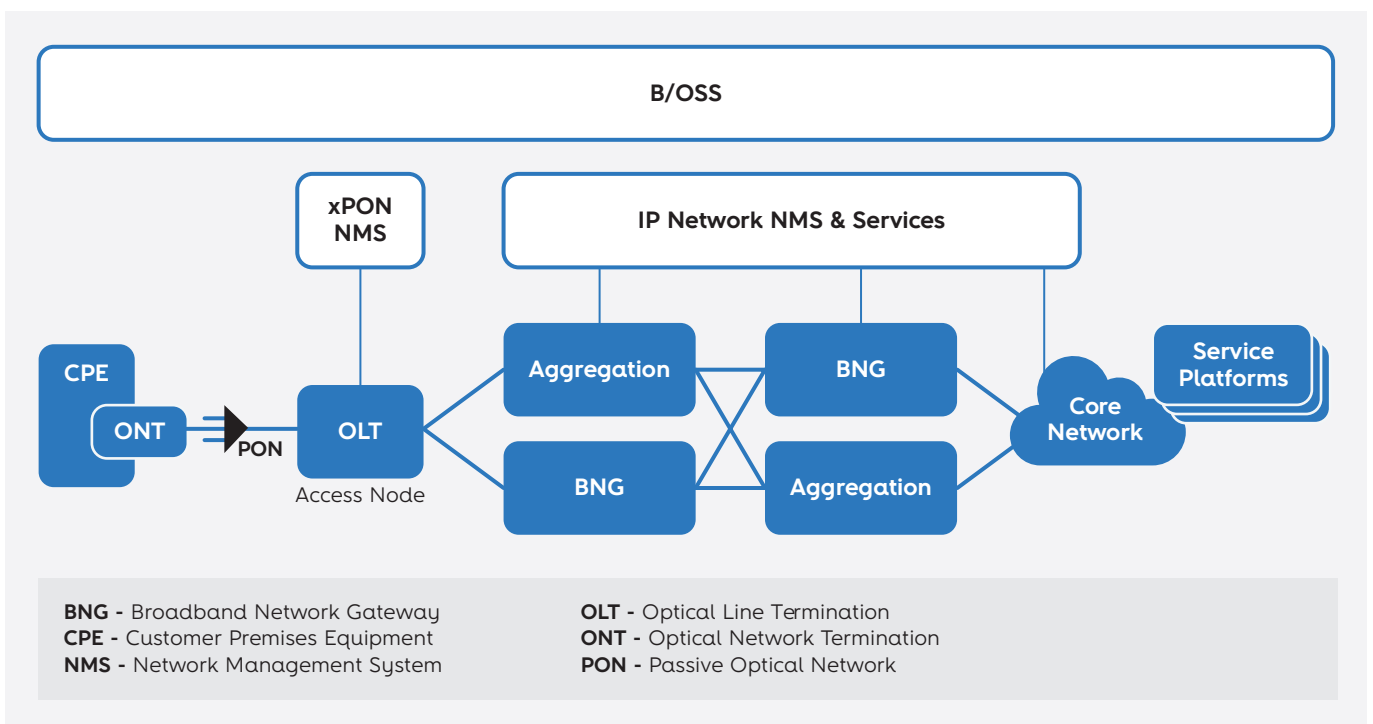
Here we describe the state of the art of CO evolution, focussing on xPON access networks, which are particularly relevant for AltiCe Labs.

### Current xPON architecture

The typical network service provider GPON access deployments follow the architecture defined in TR-101/ TR-156 [3][4]. This architecture is depicted in **Figure 2**.

The following key components can be identified:

- Customer premises equipments (CPE): the devices located at the home or the business premises of a customer (ex: router, set-top-box);
- Access node: the first aggregation point in the access network. In xPON networks this role is performed by the optical line terminal (OLT);
- Optical network termination (ONT) or optical network unit (ONU): the device that terminates the xPON network at the customer site;
- OLT: the device that terminates the xPON network at the CO side;
- Broadband network gateway (BNG): its main role is to terminate the Ethernet layer and to perform tasks such as hierarchal QoS, congestion management, access loop identification and policing on a per-user basis;
- Network management system (NMS): usually each segment of the network will have its vendor specific NMS;



**FIGURE 2** – Current xPON architecture

- B/OSS: the systems responsible for managing the operator’s network to deliver services to the end user and for monetizing it.

The Broadband Forum has extended this architecture in TR-178 [5] to add the possibility to extend IP/MPLS to the access network, to embed BNG functions in the access node and consequently to enable new types of services.

## BBF cloud central office

### Common architecture

BBF’s CloudCO initiative completely redefines the TR-101/TR-178 [3][5] architecture to enable the adoption of SDN, NFV and cloud computing. The CloudCO Reference Architectural Framework is

defined in TR-384 [6] and aligns the standards being defined for SDN and NFV with the reality of CO and the experience of “pre-standard” implementations, like central office re-architected as a data centre (CORD), described below.

Along with the architecture definition, the BBF also defines a number of CloudCO use cases and scenarios in TR-416 [7] and is defining interfaces, test cases, migration scenarios, and a number of application notes for implementable use cases.

Under this new architecture, many functions of access nodes and BNG are disaggregated and run as VNF in compute nodes. The special-purpose dedicated nodes are replaced by general purpose computing, storage and switching devices. Specific access nodes provide the I/O hardware for the data plane of each type of access technology (GPON, G.fast, etc.).

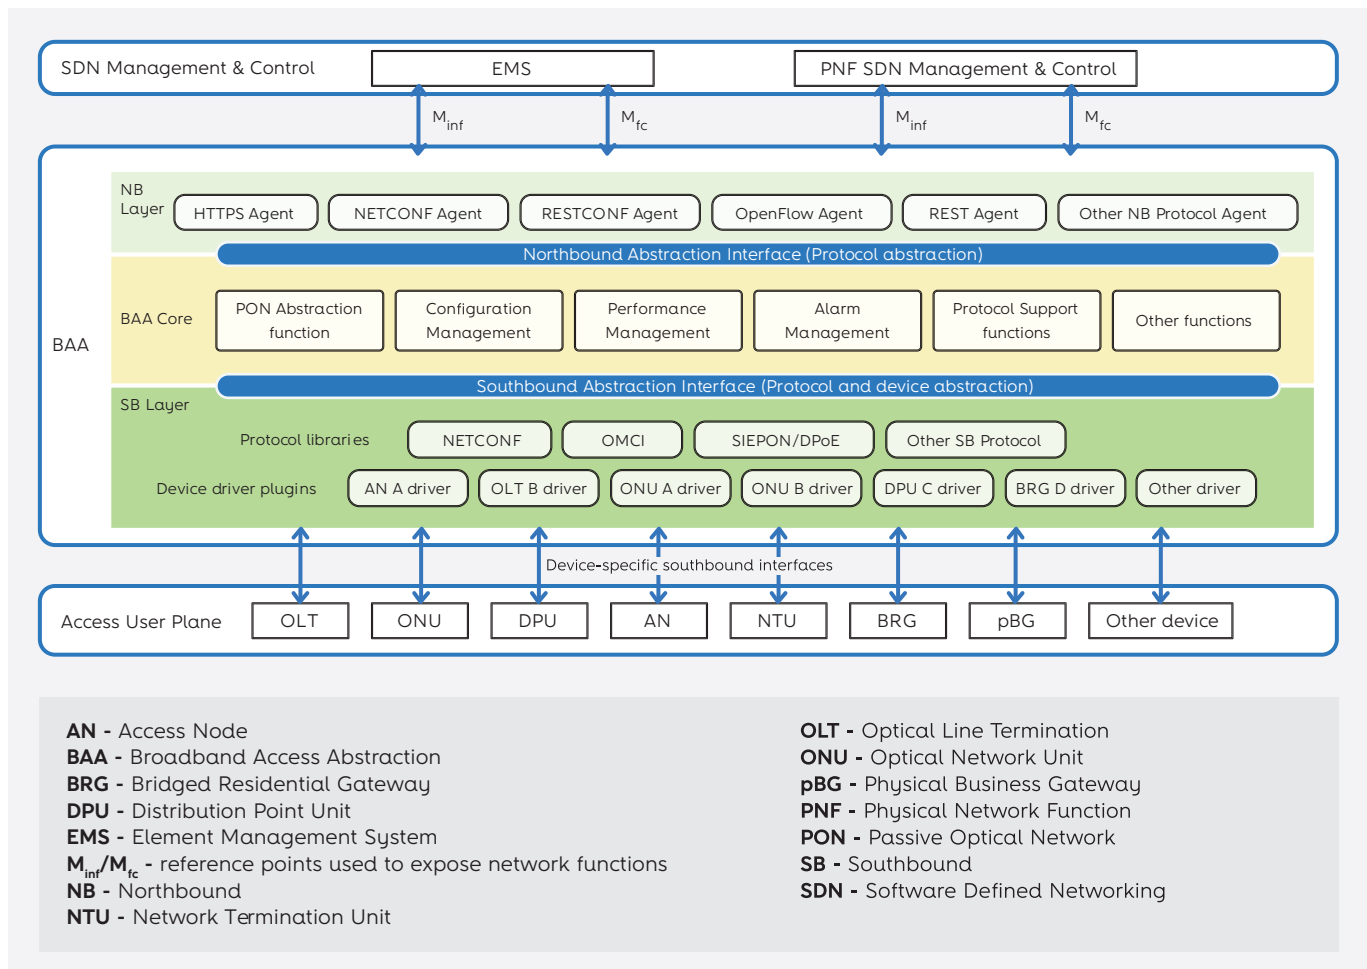


FIGURE 3 – Broadband access abstraction [6]

The B/OSS systems interact with the CloudCO domain through the northbound (NB) API of the CloudCO domain orchestrator.

## Access nodes

In a CloudCO, access devices like the OLT are disaggregated into their basic functions, keeping the data plane in a physical element (SDAN). The SDAN is “exposed” as a common switch to the SDN controller layer by the broadband access abstraction (BAA) [6], which provides an adaptation between access-specific physical functions and a common SDN control layer and replaces the specialised software that runs on the traditional access nodes (see **Figure 3**).

The BAA consists of three layers.

- NB layer: provides the northbound protocol implementation;
- BAA core: contains a set of disaggregated functions defined independently of the NB and southbound (SB) protocols;
- SB layer: converts generic commands from the northbound control and management elements device specific messages.

The NB layer provides two types of interfaces, one for management ( $M_{inf}$ ) and another for flow control ( $M_{fc}$ ) fitting the SDN paradigm of separation between the management and control planes. The BAA also provides interfaces to attach an element management system (EMS) to perform configuration, reporting and alarming functions of the access nodes.

The open broadband-broadband access abstraction (OB-BAA) [8] is an open source reference implementation of the BAA architecture currently promoted by the BBF.

## Broadband network gateway

The BNG is another key component of the CO. It is usually a very complex and expensive device that joins together a large number of

functionalities, and hence has much to benefit from disaggregation and virtualisation.

As mentioned above, TR-384 [6] also proposes the decomposition of the BNG, departing from two main contexts: control plane and data plane. In the control plane, one can find all the functions related to service and subscriber session control, while in the data plane one can find the service and routing forwarding functions. **Figure 4** shows the decomposition of the legacy BNG.

On the deployment of the disaggregated BNG, one can take benefit from virtualisation since the system is now divided into more atomic functions. The service and subscriber session control functions can be instantiated as VNF on an NFV environment, either has independent VNF or by grouping them into a single VNF. The data plane functions can be instantiated on a mix of physical network functions (PNF), on physical devices to achieve higher packet processing performance, for the forwarding functions and VNF for routing control functions.

Also, the disaggregation of the BNG functions together with the same trend for the OLT will enable the use of common software and control processes that was not possible with the current non-disaggregated architecture.

## Other pre-standard activities

- **CORD** [9] is an initiative started by AT&T, now led by the Open Networking Foundation (ONF) [10] and hosted by the Linux Foundation that aims to “*promote the building of network edge facilities from cloud hardware and software technologies*” [11]. CORD is a set of multiple open source projects, not only in the sense that the software is open source (under the Linux Foundation), but also that it is built with Open Compute Project (OCP) specifications servers and switches, as well as the access nodes [12]. All these projects comply with a common CORD architecture, each one addressing different aspects. One project that is particularly

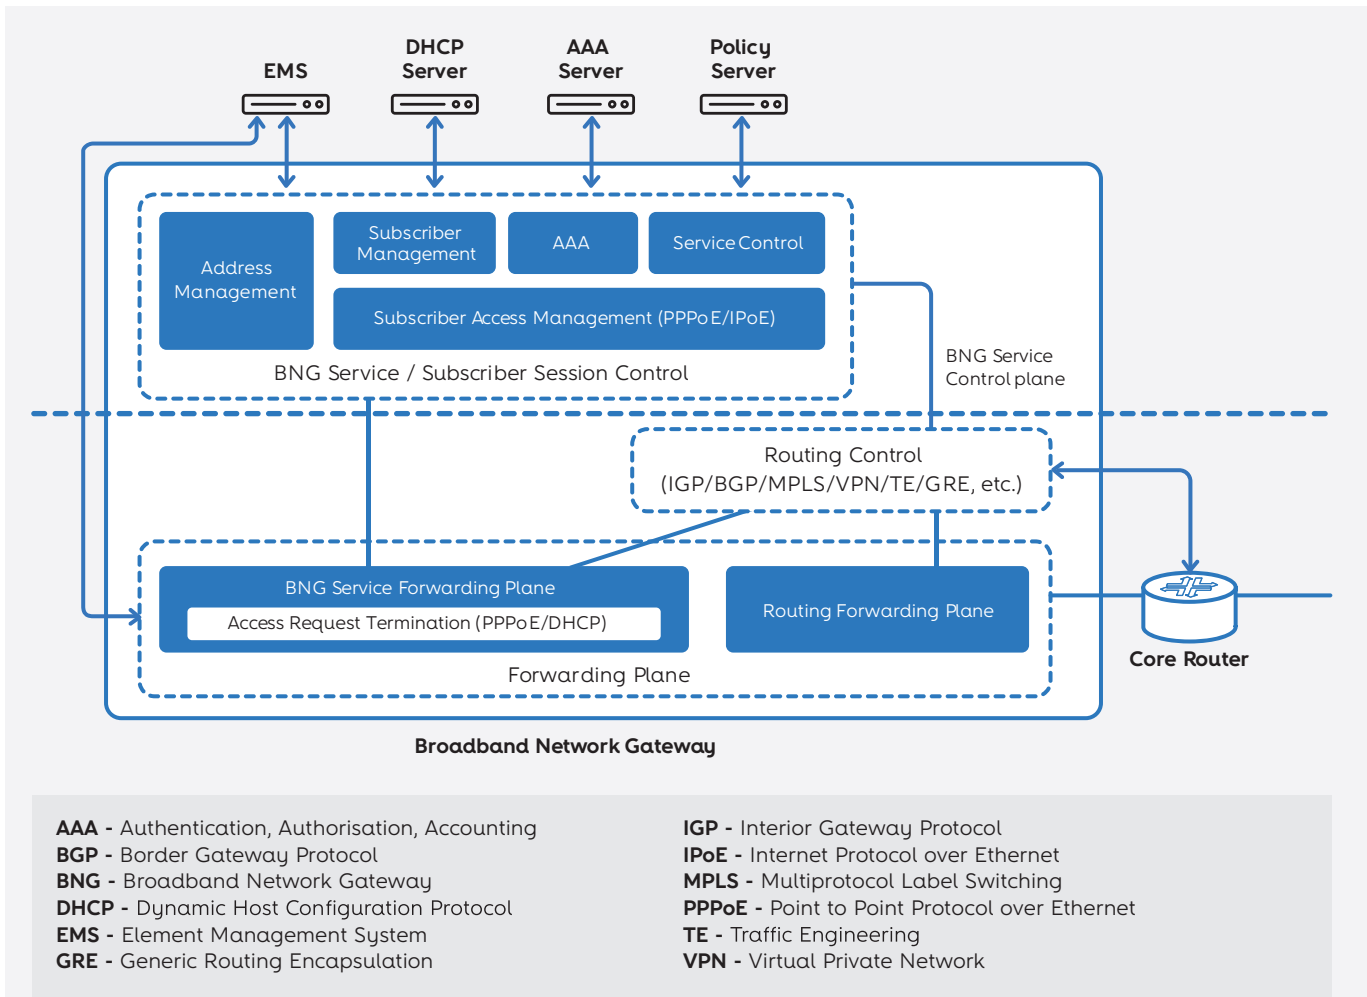


FIGURE 4 – BNG decomposition [6]

relevant for PON access nodes is the virtual OLT hardware abstraction (VOLTHA). The objective of VOLTHA is similar to the BAA layer (and its OB-BAA implementation): to provide a vendor-independent northbound interface to PON (and other) hardware devices which include protocols such as REST, OpenFlow and NETCONF.

The CORD architecture is aligned with TR-384, briefly described above. In fact, CORD has been a precursor and a major motivation for it.

- Virtual central office (**VCO**) is an open platform for NFV (OPNFV) [13] project to build demos for central office virtualisation using open source projects (OPNFV, ODL, OCP, OpenStack). The present work in this initiative, VCO 2.0,

is focused on central office virtualisation for mobile access, but the group has plans for future demos involving residential (PON, vOLT and vCMTS/vCCAP) and enterprise (SD-WAN, PBX) environments.

## MEC

Multi-access edge computing (MEC) is an ETSI study group with the goal of establishing an architecture aligned with ETSI NFV [14] for edge computing in a multi-access environment. It aims at providing application developers and content providers with cloud computing capabilities and an IT service environment at the edge of the network. MEC does not address any access-specific aspects, but still, its scope is well within the CO domain.

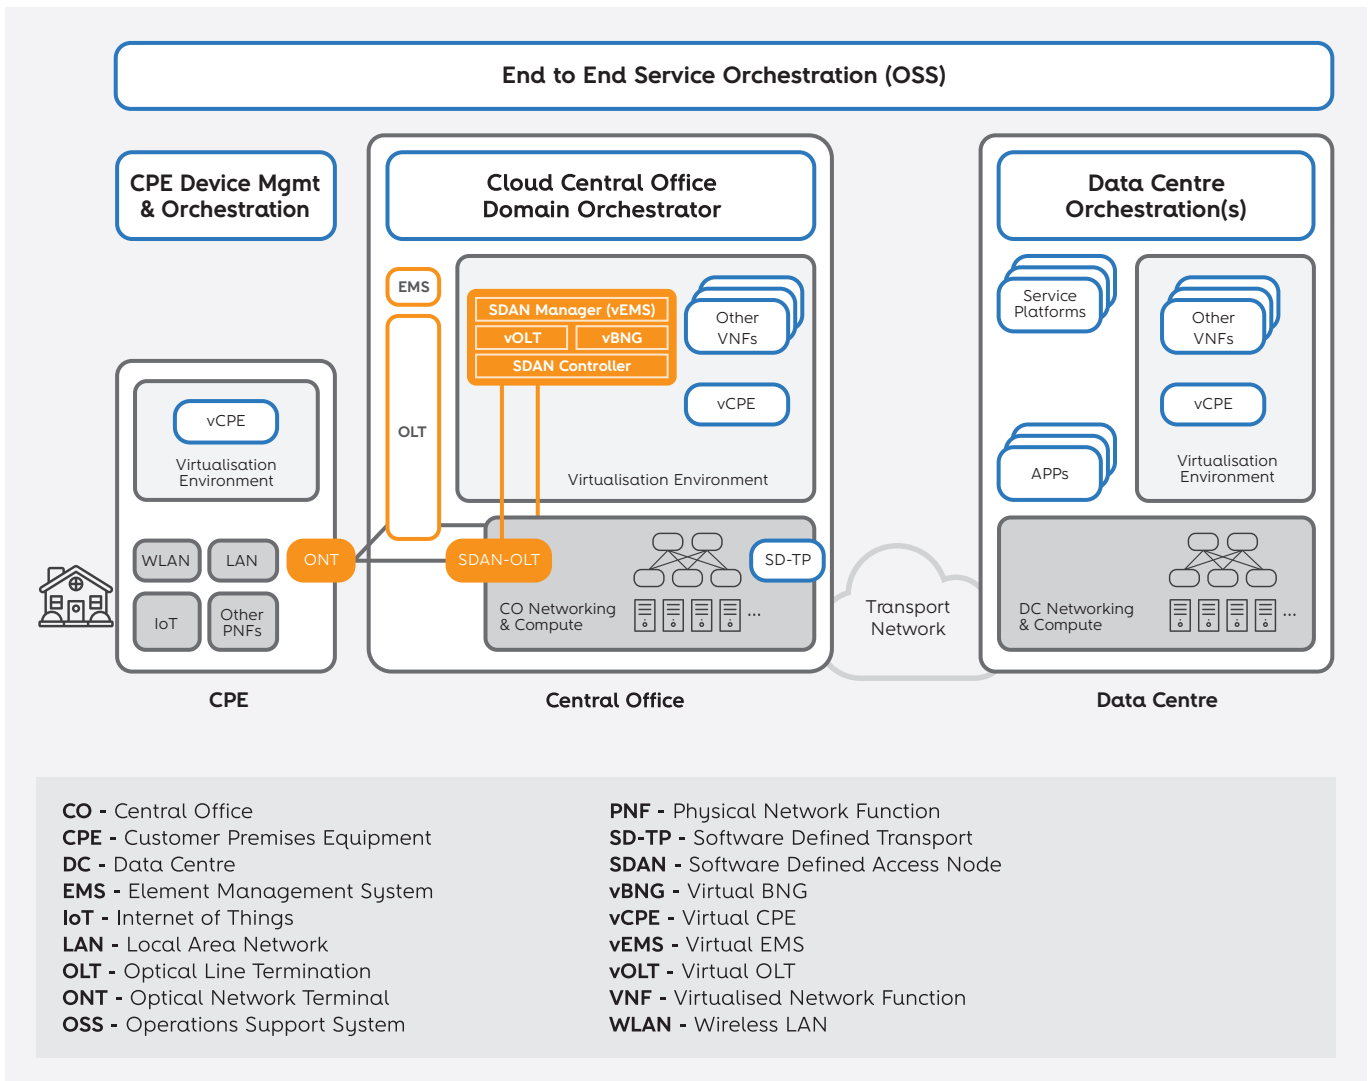
## The Altice Labs' view

The vision of Altice Labs for the CO is that it will evolve to be an IT-centered CloudCO, powered by SDN, NFV, and a network-oriented cloud approach, enhanced by highly specialised hardware for access and transport networks specific functions, as formerly described in this publication [15], and as described by the BBF in TR-384 [6].

Some aspects of the CloudCO evolution are particularly relevant in the present scenario, as they involve the roadmap of products in the Altice Labs' portfolio. This is the case of the PON

access line of products (OLT, ONT, gateways). As mentioned above, other access networks will converge on the CO, but the scenario below tries to depict a "big picture", while still trying to keep the focus on the specifics of PON access at the CO (see **Figure 5**).

At the **Central Office**, a physical **SDAN-OLT** takes care of converting data between the optical media of the PON and Ethernet packets. It differs from its precursor - the **OLT** - in that it specialises only in doing the OLT data plane functions and some very specific and time-bound control functions. The SDAN-OLT can include its own switching matrix or just connect to the switches on the CO switching fabric, e.g. as a small form-



**FIGURE 5** – General scenario with PON access



factor pluggable (SFP) switch module. In any case, the **SDAN Controller** abstracts the hardware and models **ONT** as ports of a switch. Such a model requires a software adaptation layer (such as BAA as described above in the BBF CloudCO architecture).

The SDAN Controller exposes a view of the network to a number of SDN applications that implement the control-level functions, including the control functions currently performed by the OLT and by the BNG, represented in **Figure 5** as **vOLT** and **vBNG**. In fact, these functions aren't necessarily two separate software entities. Instead, they are a set of disaggregated functions that are orchestrated to perform the duties of the OLT and BNG control planes. To name a few, these functions will include authentication, multicast control, session control, service control, address management, etc. Some of these functions are specific to the OLT or BNG, but many are common, and their disaggregation creates the opportunity for new approaches, e.g., the junction of OLT and BNG data plane functions on the same hardware under common control or the complete scattering of data plane functions across the CO switching fabric.

To guarantee connectivity between the CO and the rest of the network, software-defined transport (**SD-TP**) will tend to follow the same SDN architecture as explained for the access network, i.e., SD-TP will perform transport data plane functions, and it will be controlled by the SDN Controller that takes care of the transport network (not represented in **Figure 5**).

The control plane functions, as well as the management functions, will be typically deployed as virtualised functions, in a virtualised environment using standard off-the-shelf computing infrastructure. Along with network control functions, other VNF, as well as other applications, will be instantiated in the virtualised environment. One example of such applications is customer applications requiring properties that can only be provided at the network edge. MEC establishes an architecture to support the deployment of such functions.

The virtualisation of common data plane functions at the CO is under discussion. The increasing demand for bandwidth casts some doubt on the computing capacity needed to keep up with that growth.

The **Data Centres** will be typically connected among themselves and with the CO, with varying topology, depending on the geographic size and the client concentration of the network. These Data Centres run powerful virtualisation infrastructures capable of supporting centralised network functions, as well as service platforms.

The CPE also plays an important role with two major economic drivers:

- On one side, the large number of residential CPE pushes towards the simplification of the device, with some of its functions being virtualised at the CO;
- On the other side, the increasingly complex requirements from businesses are defining an important role for a more sophisticated CPE, capable of running virtualised functions (universal CPE - uCPE). This sets up the scenario for a cloud that extends from the customer premises to the operator core Data Centres.

The CPE example demonstrates the real impact of network softwarization: CPE functions - so far attached to the CPE box - are virtualised and can run wherever (and whenever) they are most needed or less expensive.

To coordinate all the actions necessary to organise such a complex scenario, orchestration will have to extend hierarchically from an end-to-end view of the system that is capable of understanding service as a whole to the micro-orchestration of resources to perform a certain network function.

## Altice Labs' projects

Altice Labs is currently working in different projects exploring various aspects of CO evolution which outcome will be incorporated in Altice Labs' product portfolio. Next, some of these initiatives are briefly described.

### Virtual Fiber Box

Virtual Fiber Box (VFB) is an ongoing P2020 project that started in June 2018. Altice Labs is promoting this project together with Instituto de Telecomunicações [16], to support the evolution of Altice Labs' PON portfolio in two particular areas:

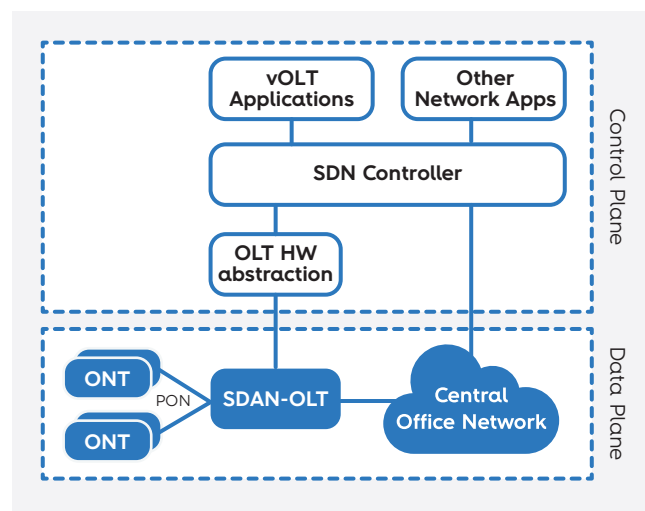
- R&D in optics, to foster the innovation that supports higher bandwidths and new technological approaches;
- The definition of an architecture for PON functions disaggregation and the development of software and hardware modules that support it.

The latter aspect is completely aligned with the proposals of the BBF and empowers Altice Labs into becoming a lead supplier, as well as an integrator, of disaggregated PON solutions.

**Figure 6** lines up the major components of the proposed architecture. In the scope of this project, Altice Labs is committed to developing the hardware of various Software Defined OLTs, the Software for the adaptation of these to an SDN Controller, and the software applications that will implement the virtualised functions of the OLT.

On the hardware component, the project will focus on the development of the SDAN concept through the following hardware platforms:

- **SDAN1** – 1 optical port OLT with inherent virtualisation characteristics in C form-factor pluggable (CFP), that will later evolve to quad small form-factor pluggable (QSFP) or SFP+. This product will be built with photonic integrated circuits (PIC);



**FIGURE 6** – Virtual Fiber Box basic architecture

- **SDAN16** – 16 port OLT with inherent virtualisation characteristics;
- **SFP1010** – 10Gbps ONU in SFP+ format with virtualisation of control functions.

The SDAN1 or the SFP1010 will enable any switch or client router to be converted into high-speed PON optical termination equipment.

On the software component, this project includes:

- The firmware for SDAN-OLT;
- OLT hardware abstraction (OLT-HA), the abstraction software that will allow the standardised control and operation of these OLT from “universal” SDN control platforms;
- Virtual OLT controller application(s), SDN application(s) that will contain a set of control functionalities disaggregated from the traditional OLT.

The OLT-HA is the software that manages the OLT and its ONT and provides northbound interfaces to the SDN controller, exposing a model of the resources it controls (in NETCONF/YANG). It corresponds to the BAA described above. The OLT-HA software can run in virtualised environments. It contains NETCONF and OpenFlow agents and the device specific drivers to the hardware (including the OMCI stack). The

drivers included in the OLT-HA interact remotely with the hardware using a proprietary interface.

The VFB project is planned to end in February 2021.

## SDAN initiative

To continue as the provider of a complete solution for the optical access network in a context of growing virtualisation and disaggregation, Altice Labs wants to embrace a CloudCO approach. To do so, Altice Labs will evolve its portfolio for access networks under a common architecture, supported on standards.

The SDAN initiative combines the efforts of the VFB project described above with other aspects of the CloudCO, namely those regarding network control usually assigned to BNG functions. It features three lines of work:

- **Virtual PON (vPON):** evolution of the ONT, OLT and the PON portfolio. This line of work is coincident with the VFB project;
- **Virtual Network Control (vNetworkControl):** development of CloudCO network control functions, capable of implementing network control functions, in particular, those performed by a BNG, within an SDAN architecture;
- **SDAN architecture:** defining and keeping a common architecture, implementing proofs-of-concept and demos, exploring new approaches for the other two lines of work.

Also, this initiative frames all experimentation and eventual participation in wider scope CloudCO

related activities. Namely, Altice Labs is actively participating in the development team of the open broadband-broadband access abstraction (OB-BAA) [8] open source project, under the auspices of BBF.

Currently, the common basic architecture is largely determined by BBF in its CloudCO initiative, described above.

## Conclusion

The evolution to a scenario like the one described above won't happen overnight, but it is definitely happening. Major operators, like AT&T, have already made trials with the described architecture [17]. Others are running trials or planning to. All of them expect to obtain important gains from it.

Nevertheless, the proposed approach assumes that important changes will occur in telecom operations. There is a major shift (which is not new) from the typical TELCO to IT-centered operations, and many operators are not prepared for this new paradigm. Also, some of the technologies involved may be not so mature as to be deployed in production, mission-critical networks.

Altice Labs is committed to evolving its portfolio, in particular, the CO PON solutions to the new paradigm, as we believe that it will bring strong benefits to our clients. We foresee that these changes are important and relevant enough to lead, not only to the re-architecture of the network but also, to important changes in the business models presented by operators and vendors alike.

## References

- [1] ITU-T, "Y.3150 : High-level technical characteristics of network softwarization for IMT-2020," 2018. [Online]. Available: <https://www.itu.int/rec/T-REC-Y.3150-201801-I/en>.
- [2] "Broadband Forum," 2018. [Online]. Available: <https://www.broadband-forum.org/>.
- [3] Broadband Forum, "TR-101 Issue 2 - Migration to Ethernet-Based Broadband Aggregation," 2011. [Online]. Available: [https://www.broadband-forum.org/technical/download/TR-101\\_Issue-2.pdf](https://www.broadband-forum.org/technical/download/TR-101_Issue-2.pdf).
- [4] Broadband Forum, "TR-156 Issue 4 - Using GPON Access in the context of TR-101," 2017. [Online]. Available: [https://www.broadband-forum.org/technical/download/TR-156\\_Issue-4.pdf](https://www.broadband-forum.org/technical/download/TR-156_Issue-4.pdf).
- [5] Broadband Forum, "TR-178 Issue 2 - Multi-service Broadband Network Architecture and Nodal Requirements," 2017. [Online]. Available: [https://www.broadband-forum.org/technical/download/TR-178\\_Issue-2.pdf](https://www.broadband-forum.org/technical/download/TR-178_Issue-2.pdf).
- [6] Broadband Forum, "TR-384 - Cloud Central Office (CloudCO) Reference Architectural Framework," 2018. [Online]. Available: <https://www.broadband-forum.org/technical/download/TR-384.pdf>.
- [7] Broadband Forum, "TR-416 - CloudCO Use Cases and Scenarios," 2018. [Online]. Available: <https://www.broadband-forum.org/technical/download/TR-416.pdf>.
- [8] Broadband Forum, "Open Broadband-Broadband Access Abstraction (OB-BAA)," 2018. [Online]. Available: <http://www.broadband-forum.org/projects/ob-projects/ob-baa>.
- [9] "CORD - Central Office Re-architected as a Datacenter," 2018. [Online]. Available: <https://opencord.org/>.
- [10] "Open Networking Foundation," 2018. [Online]. Available: <https://www.opennetworking.org/>.
- [11] Open CORD Community, "CORD," 2018. [Online]. Available: <https://wiki.opencord.org/display/CORD/Defining+CORD>.
- [12] "Open Compute Project," 2018. [Online]. Available: <https://www.opencompute.org/>.
- [13] "OPNFV," 2018. [Online]. Available: <https://www.opnfv.org/>.
- [14] ETSI, "Network Functions Virtualisation (NFV)," 2018. [Online]. Available: <https://www.etsi.org/technologies-clusters/technologies/nfv>.
- [15] R. Calé, P. Neves, M. R. Costa, C. Parada, and J. Bonnet, "NFV & SDN Innovation Strategy at Altice Labs," in *InnovAction*, no. 1, Altice Labs, 2016, pp. 34–53.
- [16] "Instituto de Telecomunicações," 2018. [Online]. Available: [www.it.pt](http://www.it.pt).
- [17] S. Hardy, "AT&T finishes open source XGS-PON field trials," *Lightwave Online Magazine*, 2018. [Online]. Available: <https://www.lightwaveonline.com/articles/2018/01/at-t-finishes-open-source-xgs-pon-field-trials.html>.



# Acronyms & Terms

<b>2</b>	<b>2G/3G/4G/5G</b>	Second/third/fourth/fifth generation mobile networks
<b>3</b>	<b>3GPP</b>	Third Generation Partnership Project, a collaboration between groups of telecommunications standards associations
<b>5</b>	<b>5G CN</b>	5G Core Network
	<b>5G IA</b>	5G Infrastructure Association, an European ICT industry joint initiative, committed to the advancement of 5G in Europe and to building global consensus on 5G
	<b>5GAA</b>	5G Automotive Association
	<b>5G-ACIA</b>	5G Alliance for Connected Industries and Automation, a global forum for collaboration between automation, engineering, and process industries, and telecom operators and suppliers
	<b>5G-NR</b>	5G New Radio
<b>A</b>	<b>AAA</b>	Authentication, Authorisation, Accounting
	<b>ADL</b>	Arthur D. Little, an international management consulting firm
	<b>Agile</b>	Iterative approach to project management and software development that helps teams deliver value to their customers faster
	<b>AI</b>	Artificial Intelligence
	<b>AN</b>	Access Network
	<b>Android</b>	Mobile operating system developed by Google
	<b>Angular</b>	TypeScript-based open-source front-end web application platform led by the Angular Team at Google and by a community of individuals and corporations
	<b>API</b>	Application Programming Interface
	<b>App</b>	Application
	<b>AR</b>	Augmented Reality
	<b>AR/VR</b>	Augmented Reality/Virtual Reality
	<b>ARPU</b>	Average Revenue Per User/Unit
	<b>AT&amp;T</b>	American multinational conglomerate holding telecommunications company

<b>B</b>	<b>B/OSS</b>	Business/Operational Support System
	<b>B2B</b>	Business-to-business
	<b>B2C</b>	Business-to-consumer
	<b>BAA</b>	Broadband Access Abstraction
	<b>BBF</b>	BroadBand Forum
	<b>BBU</b>	Baseband Unit
	<b>BGP</b>	Border Gateway Protocol
	<b>Bluetooth</b>	A wireless technology standard for exchanging data over short distances using short-wavelength ultra-high frequency radio waves
	<b>BNG</b>	Broadband Network Gateway
	<b>botnet</b>	A number of Internet-connected devices, each of which is running one or more bots
	<b>BOTSchool</b>	Virtual assistant management platform developed by Altice Labs
	<b>BRG</b>	Bridged Residential Gateway
	<b>BS</b>	Base Station
	<b>BSS</b>	Business Support System
	<b>BURN</b>	Buy and Use Right Now
<b>C</b>	<b>CAD</b>	Connected Automated Driving
	<b>CAGR</b>	Compound Annual Growth Rate
	<b>capex</b>	Capital Expenditures
	<b>Cat</b>	Category
	<b>CBF</b>	Content-Based Filtering
	<b>CEO</b>	Chief Executive Officer
	<b>CEP</b>	Complex Event Processing
	<b>CF</b>	Collaborative Filtering
	<b>CFP</b>	C Form-factor Pluggable
	<b>C-IoT</b>	Cellular-IoT
	<b>CloudCO</b>	Cloud Central Office
	<b>CMTS</b>	Cable Modem Termination System
	<b>CN</b>	Core Network
	<b>CO</b>	Central Office
	<b>Co-Design</b>	Cooperative or Participatory Design
	<b>CORD</b>	Central Office Re-architected as a Datacentre
	<b>CPE</b>	Customer Premises Equipment
	<b>CPRI Forum</b>	Common Public Radio Interface, a industry cooperation defining the publicly available specification for the key internal interface of radio base stations

<b>CPU</b>	Central Processing Unit
<b>CRF</b>	Conditional Random Field
<b>CSP</b>	Communications Service Providers
<b>CSS</b>	Cascading Style Sheets, a style sheet language used for describing the presentation of a document written in a markup language like HTML
<b>C-V2X</b>	Cellular-V2X

---

<b>D</b>	<b>D2D</b>	Device-to-Device
	<b>DAG</b>	Directed Acyclic Graph
	<b>DASH7</b>	DASH7 Alliance Protocol is an open source Wireless Sensor and Actuator Network protocol
	<b>DataPlane</b>	Big data processing platform as a service developed by Altice Labs
	<b>DC</b>	Data Centre
	<b>DDoS</b>	Distributed Deny of Service
	<b>DevOps</b>	Software development methodology that combines software development with information technology operations
	<b>DHCP</b>	Dynamic Host Configuration Protocol
	<b>DNPW</b>	Do Not Pass Warning
	<b>DOCSIS</b>	Data Over Cable System Interface Specification, an international telecommunications standard
	<b>DOM</b>	Document Object Model, a cross-platform and language-independent application programming interface
	<b>DPI</b>	Deep Packet Inspection
	<b>DPoE</b>	DOCSIS Provisioning of EPON
	<b>DPU</b>	Distribution Point Unit
	<b>DSL</b>	Domain Specific Language
	<b>DSP</b>	Digital Services Provider
	<b>DVR</b>	Digital Video Recorder

---

<b>E</b>	<b>EC</b>	European Commission
	<b>EC-GSM-IoT</b>	Extended Coverage-GSM-IoT
	<b>eDRX</b>	extended Discontinuous Reception
	<b>eMBB</b>	enhanced Mobile BroadBand
	<b>EMEA</b>	Europe, the Middle East and Africa
	<b>EMS</b>	Element Management System

<b>eMTC</b>	enhanced Machine Type Communication
<b>EPC</b>	Evolved Packet Core
<b>EPG</b>	Electronic Program Guide
<b>EPON</b>	Ethernet Passive Optical Network
<b>ESA</b>	European Space Agency, an intergovernmental organisation of European member states dedicated to the exploration of space
<b>eSIM</b>	embedded SIM
<b>ESOA</b>	EMEA Satellite Operators Association
<b>EtherCAT</b>	Ethernet for Control Automation Technology, a real-time industrial Ethernet technology originally developed by Beckhoff Automation
<b>ETSI</b>	European Telecommunications Standards Institute
<b>EU</b>	European Union
<b>eUICC</b>	embedded Universal Integrated Circuit Card
<b>E-UTRA</b>	Evolved Universal Terrestrial Radio Access
<b>Evolved E-UTRA</b>	RAT that refers to an evolution of the E-UTRA radio interface for operation in the NextGen system

---

<b>F</b>	<b>FIFA</b>	Fédération Internationale de Football Association
	<b>FTTH</b>	Fiber-to-the-Home
	<b>FUXI</b>	Framework User eXperience and Interface , developed by Altice Labs
	<b>FWA</b>	Fixed Wireless Access

---

<b>G</b>	<b>G.fast</b>	ITU-T digital subscriber line protocol standard
	<b>GAFA</b>	Google, Amazon, Facebook and Apple (group of companies made up of)
	<b>GAFA-MI</b>	Google, Amazon, Facebook and Apple - Microsoft and IBM (group of companies made up of)
	<b>GDP</b>	Gross Domestic Product
	<b>GDPR</b>	General Data Protection Regulation



<b>gNB</b>	Next Generation NodeB
<b>GPON</b>	Gigabit Passive Optical Network
<b>GPRS</b>	General Packet Radio Service
<b>GRE</b>	Generic Routing Encapsulation
<b>GSA</b>	Global mobile Suppliers Association, an organisation that aims to promote mobile phone standards worldwide
<b>GSM</b>	Global System for Mobile Communications
<b>GSMA</b>	Global System for Mobile Communications Association

---

<b>H</b>	<b>HA</b>	Hardware Abstraction
	<b>HD</b>	High Definition
	<b>HDFS</b>	Hadoop Distributed File System
	<b>HIPAA</b>	Health Insurance Portability and Accountability Act
<b>Honeynet</b>		One or more computer systems on the Internet expressly set up to detect, deflect, or counteract attempts at unauthorised use of information systems
	<b>HTML</b>	Hypertext Markup Language, the standard markup language for creating web pages and web applications
	<b>HTTPS</b>	HyperText Transfer Protocol Secure
	<b>HW</b>	Hardware

---

<b>I</b>	<b>I&amp;D</b>	Investigation and Development
	<b>I/O</b>	Input/Output
	<b>I4.0</b>	Industry 4.0
	<b>ICT</b>	Information and Communications Technology
	<b>IDE</b>	Integrated Development Environment
	<b>IETF</b>	Internet Engineering Task Force, an open standards organisation that develops and promotes voluntary Internet standards
	<b>IGP</b>	Interior Gateway Protocol
	<b>IHC</b>	Human-Computer Interaction
	<b>IHS</b>	IHS Markit, a data and information services business
	<b>IMT</b>	International Mobile Telecommunications
<b>IMT-2020</b>		An ITU-R standard for 5G

<b>iOS</b>	Mobile operating system created and developed by Apple Inc
<b>IoT</b>	Internet of Things
<b>IP</b>	Internet Protocol
<b>IPoE</b>	Internet Protocol over Ethernet
<b>IPTV</b>	Internet Protocol Television
<b>ISM</b>	Industrial, Scientific and Medical
<b>ISP</b>	Internet Service Provider
<b>IT</b>	Information Technology
<b>ITU-R/T</b>	International Telecommunication Union, Radiocommunication Sector/Telecommunication Standardization Sector

---

<b>J</b>	<b>JDBC</b>	Java Database Connectivity
	<b>JS</b>	JavaScript, a high-level, interpreted programming language
	<b>JVM</b>	Java Virtual Machine

---

<b>K</b>	<b>KPI</b>	Key Perform Indicators
----------	------------	------------------------

---

<b>L</b>	<b>L2/L3</b>	Layer 2/Layer 3 of the OSI Reference Model
	<b>LAN</b>	Local Area Network
	<b>LiDAR</b>	Light Detection And Ranging, a detection system which works on the principle of radar, but uses light from a laser
	<b>LoRa</b>	Long Range
	<b>LoRa Alliance</b>	Open, non-profit organisation dedicated to promoting the interoperability and standardisation of LPWAN technologies to drive implementation of the IoT
	<b>LoRaWAN</b>	Long Range Wide Area Network
	<b>LPA</b>	Local Profile Assistant
	<b>LPWA</b>	Low Power Wide Area
	<b>LPWAN</b>	Low Power Wide Area Networks
	<b>LTA</b>	Left Turn Assist
	<b>LTE</b>	Long Term Evolution
	<b>LTE Cat NB1</b>	Same as NB-IoT
	<b>LTE-M</b>	LTE category M1
	<b>LTE-MTC</b>	Long Term Evolution for Machine Type Communication

---

<b>M</b>	<b>M2M</b>	Machine-to-Machine
	<b>MAPE</b>	Monitor-Analyse-Plan-Execute

<b>MBB</b>	Mobile BroadBand
<b>MEC</b>	Multi-access Edge Computing
<b>MEF</b>	Metro Ethernet Forum, a global industry alliance that aim to accelerate the worldwide adoption of carrier class ethernet networks and services
<b>Mgmt</b>	Management
<b>MIMO</b>	Multiple-Input/Multiple-Output
<b>M<sub>inf</sub>/M<sub>fc</sub></b>	Reference points used to expose network functions
<b>ML</b>	Machine Learning
<b>mMTC</b>	massive Machine Type Communications
<b>mmWave</b>	Millimeter Wave
<b>MPLS</b>	Multiprotocol Label Switching
<b>MPP</b>	Massively Parallel Processing
<b>MTC</b>	Machine Type Communications

<b>N</b>	<b>NA</b>	Network Activator, developed by Altice Labs
	<b>NAT</b>	Network Address Translation
	<b>NB</b>	Northbound
	<b>NB-IoT</b>	Narrow Band IoT
	<b>NETCONF</b>	Network Configuration protocol
	<b>NFS</b>	Network File System
	<b>NFV</b>	Network Function Virtualisation
	<b>NGMN</b>	Next Generation Mobile Network, a mobile telecommunications association of mobile operators, vendors, manufacturers and research institutes
	<b>NLG</b>	Natural Language Generation
	<b>NLP</b>	Natural Language Processing
	<b>NLU</b>	Natural Language Understanding
	<b>NMS</b>	Network Management System
	<b>NOC</b>	Network Operations Centre
	<b>NoSQL</b>	Non Relational Database
	<b>NOSSIS</b>	New generation of NOSSIS with core activities for service operations, developed by Altice Labs
	<b>One</b>	Network Operation Support Systems Integrated Solution, developed by Altice Labs
	<b>NOSSIS</b>	Network Operation Support Systems Integrated Solution, developed by Altice Labs
	<b>NPE</b>	Java Null Pointer Exception
	<b>NR</b>	New Radio
	<b>NSA</b>	Non-Standalone Architecture
	<b>NTU</b>	Network Termination Unit

<b>OB-BAA</b>	Open Broadband-Broadband Access Abstraction
<b>OCP</b>	Open Compute Project, an organisation that shares designs of data centre products among companies
<b>ODA</b>	Open Digital Architecture
<b>ODL</b>	OpenDaylight, an open source SDN controller
<b>OFDM</b>	Orthogonal Frequency-Division Multiplexing
<b>OLT</b>	Optical Line Termination
<b>OLT-HA</b>	OLT Hardware Abstraction
<b>OM</b>	Order Manager, developed by Altice Labs
<b>OMCI</b>	ONT Management and Configuration Interface
<b>ONAP</b>	Open Network Automation Platform
<b>ONF</b>	Open Networking Foundation, a user-driven organisation dedicated to the promotion and adoption of SDN through open standards development
<b>ONT</b>	Optical Network Terminal
<b>ONU</b>	Optical Network Unit
<b>OpenFlow</b>	A communications protocol that gives access to the forwarding plane of a network switch or router over the network
<b>OpenStack</b>	A free and open-source software platform for cloud computing
<b>opex</b>	Operational Expenditures
<b>OPNFV</b>	Open Platform for NFV, a collaborative open source platform for network functions virtualisations
<b>OSI</b>	Open Systems Interconnection, a reference model for how applications communicate over a network
<b>OSM</b>	Open Source Mano, an ETSI-hosted initiative to develop an Open Source NFV Management and Orchestration (MANO) software stack aligned with ETSI NFV
<b>OSS</b>	Operational Support System
<b>OTDOA</b>	Observed Time Difference Of Arrival
<b>OTT</b>	Over-the-Top

	<b>OWASP</b>	Open Web Application Security Project			
<b>P</b>	<b>P2020</b>	Portugal 2020, a partnership agreement between Portugal and the European Commission			
	<b>PAN</b>	Personal Area Network			
	<b>pBG</b>	Physical Business Gateway			
	<b>PBX</b>	Private Branch Exchange			
	<b>PIC</b>	Photonic Integrated Circuits			
	<b>PMSE</b>	Programme Making and Special Events			
	<b>PNF</b>	Physical Network Function			
	<b>PON</b>	Passive Optical Network			
	<b>PPPoE</b>	Point to Point Protocol over Ethernet			
	<b>PSM</b>	Power Saving Mode			
<b>Q</b>	<b>QoE</b>	Quality of Experience			
	<b>QoS</b>	Quality of Service			
	<b>QSFP</b>	Quad Small Form-factor Pluggable			
<b>R</b>	<b>R&amp;D</b>	Research and Development			
	<b>RAN</b>	Radio Access Network			
	<b>RAT</b>	Radio Access Technology			
	<b>RDBMS</b>	Relational Database Management System			
	<b>React</b>	JavaScript library for building user interfaces			
	<b>REST</b>	Representational State Transfer, an architectural style for developing web services			
	<b>RESTCONF</b>	IETF draft that describes how to map a YANG specification to a RESTful interface			
	<b>RPM</b>	RPM Package Manager, a package management system			
	<b>RPMA</b>	Random Phase Multiple Access			
	<b>RTP2</b>	Second television channel of Rádio e Televisão de Portugal, the Portuguese public broadcasting corporation			
<b>S</b>	<b>SA</b>	Standalone Architecture			
	<b>SaaS</b>	Software as a Service			
	<b>SB</b>	Southbound			
	<b>SBA</b>	Service-Based Architecture			
	<b>SD</b>	Software Defined			
	<b>SDAN</b>	Software Defined Access Node			
	<b>SDN</b>	Software Defined Networks			
	<b>SD-TP</b>	Software Defined Transport			
	<b>SD-WAN</b>	Software-Defined networking in a Wide Area Network			
	<b>SFP</b>	Small Form-factor Pluggable			
	<b>SFP+</b>	SFP plus, an enhanced version of the SFP			
	<b>SIEPON</b>	Service Interoperability in Ethernet Passive Optical Networks			
	<b>SIGFOX</b>	French global network operator that builds wireless networks to connect low-power objects			
	<b>SIM</b>	Subscriber Identity Module			
	<b>SLA</b>	Service Level Agreement			
	<b>SME</b>	Small and Medium Enterprises			
	<b>SON</b>	Self Organising Networks			
	<b>SQL</b>	Structured Query Language			
	<b>SSH</b>	Secure Shell			
	<b>SST</b>	Service/Slice Type			
	<b>STB</b>	Set-Top Box			
	<b>SVM</b>	Support Vector Machine			
	<b>SVOD</b>	Subscription Video on Demand			
	<b>SWOT</b>	Strengths, Weaknesses, Opportunities and Threats			
<b>T</b>	<b>TCA</b>	Total Cost of Acquisition			
	<b>TCO</b>	Total Cost of Ownership			
	<b>TE</b>	Traffic Engineering			
	<b>TELCO</b>	Telecommunication Operators			
	<b>TM Forum</b>	A non-profit industry association for service providers and their suppliers in the telecommunications industry			
	<b>TPI</b>	Transversal Program Info			
	<b>TSN</b>	Time-Sensitive Networking			
	<b>TTK</b>	Trouble Tickets			
	<b>TV</b>	Television			
<b>U</b>	<b>UCD</b>	User-Centred Design			
	<b>uCPE</b>	Universal CPE			
	<b>UDN</b>	Ultra Dense Network			
	<b>UE</b>	User Equipment			
	<b>UEFA</b>	Union of European Football Associations			
	<b>UI</b>	User Interface			
	<b>UMTS</b>	Universal Mobile			

<b>URLLC</b>	Telecommunications Service Ultra Reliable and Low Latency Communications
<b>US</b>	United States
<b>USD</b>	United States Dollar
<b>UX</b>	User Experience

<b>Z</b>	<b>ZSM</b>	Zero Touch Network and Service Management
----------	------------	--

---

<b>V</b>	<b>V2I</b>	Vehicle-to-Infrastructure
	<b>V2N</b>	Vehicle-to-Network
	<b>V2P</b>	Vehicle-to-Pedestrian
	<b>V2V</b>	Vehicle-to-Vehicle
	<b>V2X</b>	Vehicle-to-everything
	<b>VA</b>	Virtual Assistant
	<b>vBNG</b>	Virtual BNG
	<b>VCA</b>	Virtual Customer Assistant
	<b>vCCAP</b>	Virtual Converged Cable Access Platform
	<b>vCMTS</b>	Virtual CMTS
	<b>VCO</b>	Virtual Central Office
	<b>vCPE</b>	Virtual CPE
	<b>VEA</b>	Virtual Employee Assistant
	<b>vEMS</b>	Virtual EMS
	<b>VFB</b>	Virtual Fiber Box
	<b>vNetwork- Control</b>	Virtual Network Control
	<b>VNF</b>	Virtual Network Functions
	<b>VOD</b>	Video On Demand
	<b>vOLT</b>	Virtual OLT
	<b>VOLTHA</b>	Virtual OLT Hardware Abstraction
	<b>VPA</b>	Virtual Personal Assistant
	<b>VPN</b>	Virtual Private Network
	<b>vPON</b>	Virtual PON
	<b>VR</b>	Virtual Reality
	<b>VRU</b>	Vulnerable Road User

---

<b>W</b>	<b>W3C</b>	World Wide Web Consortium
	<b>Wi-Fi</b>	IEEE 802.11x - Wireless Network (Wi-Fi Alliance)
	<b>WLAN</b>	Wireless LAN

---

<b>X</b>	<b>xPON</b>	Designation for several PON technologies
----------	-------------	---

---

<b>Y</b>	<b>YANG</b>	Yet Another Next Generation, a data modeling language for the definition of data sent over the NETCONF network configuration protocol
----------	-------------	---



