Platforms: Infrastructures for a New Economy

An IEEE Digital Reality White Paper

January 2020

Roberto Saracco, Co-chair, IEEE Digital Reality Initiative

digitalreality.ieee.org

Contents

The Digital Transformation is made possible by technology advances and is steered by economic and societal factors. It is a whole system transformation that requires the availability of a tremendously complex infrastructure, similar to how the atoms economy has evolved over centuries by creating in parallel tremendously complex and intertwined infrastructures. Logistics value chains have become extremely effective and extremely complex. Computers have boosted their efficiencies and made possible the creation and management of even more complex infrastructures. These computers supporting infrastructures are slowly morphing into infrastructures of their own. The large data centers that are supporting shipping of parcels and containers are now managing the shipping of bits in a structured way (e.g. by adopting blockchain).

These infrastructures, or platforms, have been around for quite some time and exist independently of the Digital Transformation. However, this latter has created an economic environment where platforms become a fundamental component. This whitepaper presents a synthesis of discussions with key industry players regarding platforms in this age of Digital Transformation.[1]

Digital Transformation leverages (as much as possible) bits (data). This, on the one hand, decreases (in most cases, actually slashes) the cost of developing services and products but on the other hand, requires the availability of an infrastructure to manage those bits. This infrastructure should be reliable, should support huge storage and intensive processing, should be easy to reach, and should ensure basic services like authentication, ownership and so on. In other terms, this required infrastructure turns out to be complex and expensive.

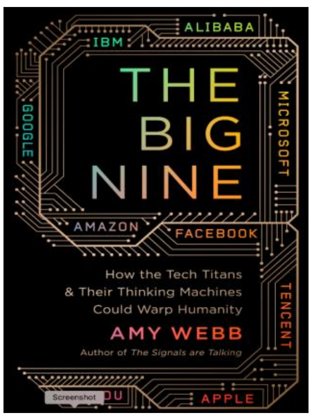

Figure 1 The cover book of "The Big Nine: How the Tech Titans and Their Thinking Machines Could Warp Humanity" which describes how nine global companies are powering, and shaping, the digital world. Image credit: Amy Webb

This complexity and expense causes the production, delivery and operation of these platforms to be provided by only a few companies. These platforms then beget ecosystems populated by a myriad of service providers who take advantage of the low cost of "using" the platform to create value and go to market. Today there are very few companies that can sustain global horizontal platforms, like the nine addressed in the Amy Webb book "The Big Nine"[2], usually known as the G-MAFIA (Google-Microsoft, Amazon, Facebook, IBM, Apple) and BAT (Baidu, Alibaba, Tencent) plus a number of others addressing important verticals, like industry, IoT, cities and so on.

To support storage processing and massive distributed access, these infrastructures use distributed clouds, thus insuring the required reliability and availability. These clouds are huge, expensive to create and operate, and can only survive economically by addressing a very large market that generates an almost statistically uniform use of the resources. This statistical uniformity is achieved through the diverse, scattered demand of thousands of users and results in very low cost (and price) per user. This creates the oligopolies we have today and makes it difficult for new entrants to grab a share of the market.

The exact number of servers in these giants' nine clouds is not known, but for each of them the number is in the millions (In October 2018, WikiLeaks released an Amazon Atlas listing the location of Amazon servers taken from an internal document of 2015, mentioned here to indicate how understandably secretive these companies are and how difficult it is to have up to date information).

[1] This topic of platforms is addressed in the Digital Transformation Course developed by EIT Digital (available now) and IEEE Digital Reality Initiative available 1st Quarter 2020.

[2] Webb, Amy, *The Big Nine: How the Tech Titans and Their Thinking Machines Could Warp Humanity*, PublicAffairs, 2019.

The need for accessibility to these platforms was the driver for many telecommunications operators to step into the horizontal platform market, but they cannot compete with the big nine since the access infrastructure has become a commodity and the scale factor is completely unbalanced towards the big nine. Telecommunications operators basically have country wide markets; the big nine have worldwide (or Chinese) markets; and in this business volume wins. Yes, there are plenty of small data centers (some as small as a few clustered servers), but they address niches and cannot compete on price with the big nine.

1. There are platforms and then there are platforms

The previous section considered the "big nine" mostly from the point of view of their capability to offer storage, processing and access (this latter mediated by the Internet). However, these nine platforms are quite different from one another:

- Amazon likely provides the biggest platform in terms of capabilities. Actually, one might say that it is multiple platforms in one. It is offering pure storage and computation services (Amazon Web Services - AWS), but it is also a marketplace including managing payments, a media distributor (music, video, books…), and a logistic delivery platform;

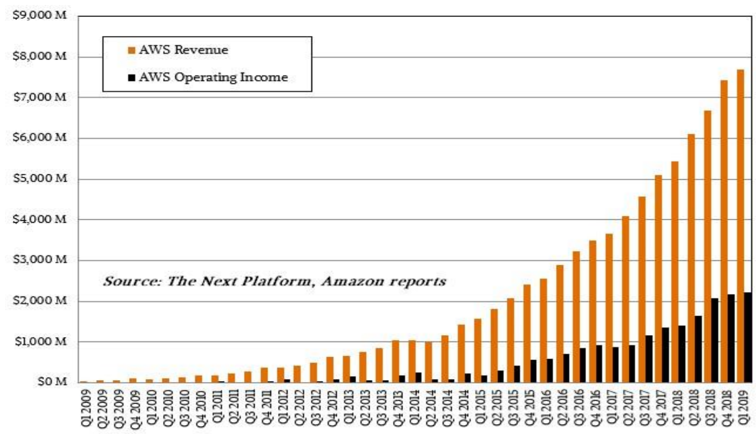

Figure 2 The Amazon AWS platform has become the largest systems business in the world, as of April 2019. Image credit: The next platform, Amazon Report

- Google is a mail platform, a service platform (e.g. Google Maps), an advertisement platform as well as a storage and processing platform. (Google previously offered free storage; recently their policy changed. More than 15GB requires a subscription to their storage service, starting at 1.99$ a month for 100GB of storage). Its strength is clearly offering services (new ones are being added continuously, such as real time translation and voice interfaces), and they are sustained by indirect revenues through advertisement;
- Apple is a marketplace platform selling billions of apps, but it is also a development platform for apps, a media distributor (music and more recently video and television), a provider of backup and storage services, and a payment and credit card platform. Apple is also a platform of devices; all of their devices (in the billions) share a common operating system (OS) which creates a huge marketplace fully orchestrated by Apple. Google is also offering a de-facto standard operating system, Android, but its grip on the market is not as strong as Apple.

As one can see, in the case of the big nine, there is a common denominator formed by the huge capacity and global access of the data centers each of them own. Their strength is that each of the variations of their platform space (services) reinforces the other.

An alternative way of looking at platforms is from the business viewpoint, as shown in Figure 3.

Eight types of platforms are identified:

- **Technology**: these are the platforms offering pure storage and processing resources. Azure (Microsoft), AWS (Amazon) are the main ones, but also Bluemix (IBM) and many platforms offered by telecommunications companies;f
- **Computing**: these are the very well-known platforms that were once called operating systems, such as Android, iOS, MacOS and Windows, and the less successful Tizen (Samsung) or HarmonyOS (Huawei). Linux would obviously fit in this category, but it is seldom addressed as a platform since it lacks the commercial halo that characterizes the other platforms in this area;
- **Utility**: these are platforms used for specific services, like Google Search, Google Maps or Kayak;

- **Interactions**: these are the platforms creating social networks and supporting interaction among their members, like Facebook, LinkedIn, Snapchat, as well as those supporting specialized focus groups like DPReview or IEEE Collaboratec;

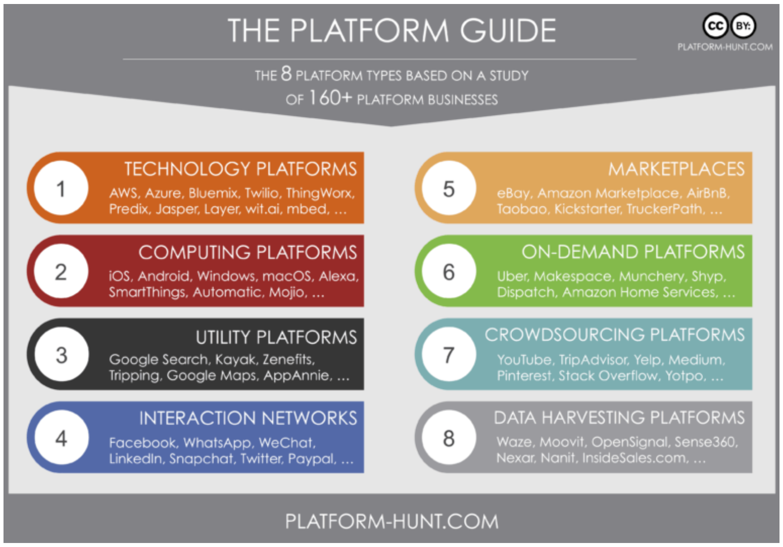

Figure 3 A possible classification of platforms based on business models. Source: Platform Hunt

- **Marketplaces**: these are platforms providing companies and individuals a way to offer goods with Amazon leading the pack, eBay, AirBnB, Kickstarter, and others;
- **On-demand**: these are the platforms supporting the Gig economy, like Foodies, Uber, Makespace;
- **Crowdsourcing**: these are the platforms harvesting people's contributions to create value, like TripAdvisor whose strong point is the reviews provided by people who ate in a restaurant or visited a location. These reviews are used by other people to decide what to do and what to buy. Other crowdsourcing platforms are the ones where people provide content, like YouTube, Pinterest or Yelp;
- **Data harvesting**: these are also platforms that capitalize on people contribution but in terms of data, like Waze that let people send info about traffic situation or Moovit (transportation availability).

There is also another classification of platforms which is based on the capability of attracting vertical sectors, i.e., creating ecosystems:

- IoT platforms
- data platforms
- industry platforms
- smart cities platforms
- healthcare platforms
- device/product platforms
- knowledge platforms

The list is clearly long and growing. The latter classification will be used throughout this whitepaper as it can be used to address several aspects of the roles played by platforms in the digital transformation.

2. IoT Platforms

IoT now number in the hundreds of billions so it should not be a surprise that the IoT platform market is huge and is soaring. It is difficult to agree on a number since it depends on the definition of an IoT and how they are counted: a connected sensor is an IoT. A smartphone has some 14 different sensors in it, each generating data that can be retrieved through the network. There are 3.3 billion smartphones today in the world; hence there are close to 50 billion IoT embedded in smartphones. However, some would disagree by not counting them as such but as one phone is one IoT. According to the *Zinnov Zone IoT Platforms 2019 the* market has reached $1.9B in 2019 and it is expected to pass the $9B by 2024, a CAGR of 35%.

Figure 4 A graphical overview of IoT platforms. Notable are the number of players and the subdivision in various categories. Image credit: Vinay Solanski, IIM Ahmedabad based on data from Zinnov Zone

There are plenty of IoT platforms; the Zinnov Zone 2019 report put the number to over 400, and again the number depends the definition of a platform. It is interesting that the IoT platforms can be clustered based on what they focus on, as shown in Figure 4. Note that there are several overlaps, as expected:

- Hardware IoT/device management platforms
- Connectivity management platforms
- IoT applications enablement platforms

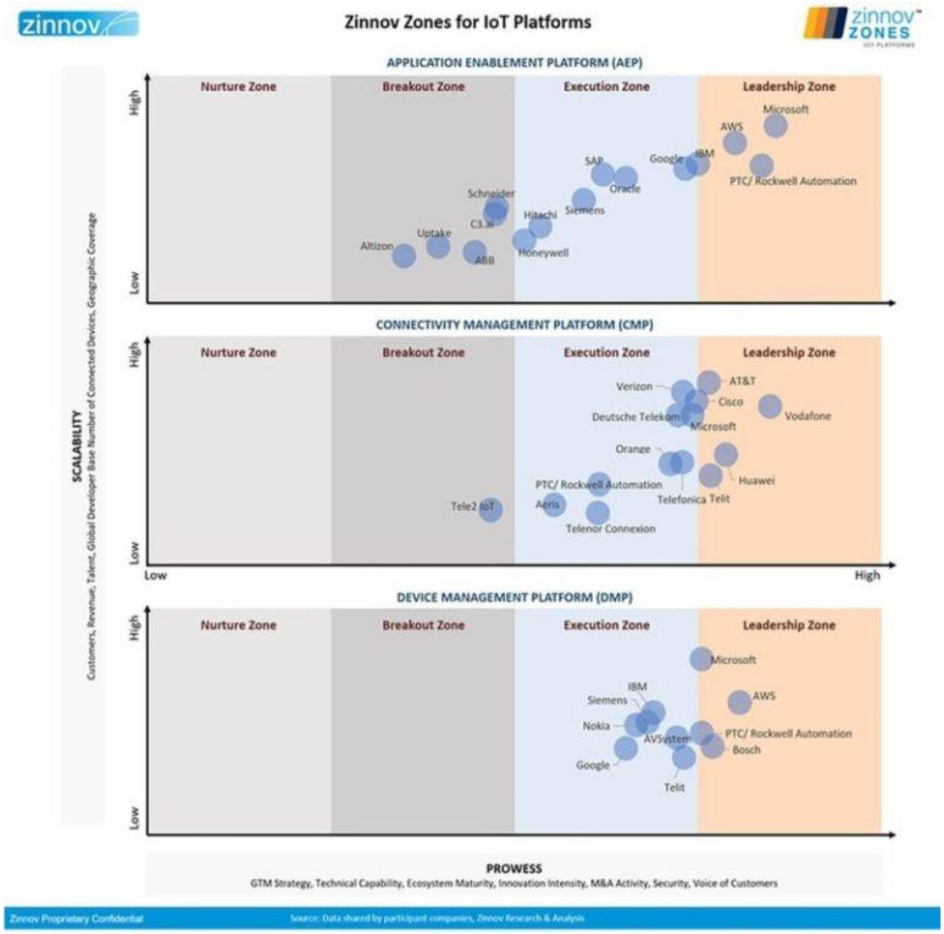

Figure 5 The main players in the IoT platform space .Telcos are providing the connectivity management platforms as expected. Image credit: Zinnov Zones

In addition, there are companies producing IoT (sensors and devices), companies providing storage and computation infrastructures for IoT, and companies providing security tools and services for IoT - a quite extensive landscape. Notice that in Figure 5 the Chinese companies are missing in the area of sensors and devices, and yet, they have the lion's share of the production of IoT (most are based in Wuxi).

IoTs are the "physical" starting point of the Digital Transformation. They are the "tools" transforming atoms into bits (sensors). More and more products are designed to be connected and to generate data for remote monitoring as well as for building services. This latter is a crucial aspect of the digital transformation business model, the one that fosters the creation of value, for both the product producer as well as for the ecosystems that can grow around a product.

The platforms become the enabler of value generation. The IoT/device management platforms space is populated by industries that have adopted the Industry 4.0 paradigm, like Siemens, and by software companies that leverage the storage and computation infrastructure to host the virtual image of the IoT and of the data they generate (i.e., a Digital Twin and its shadow). It is interesting to note that some of these companies are seeing IoT and IoT platforms as a way to change their business model.

The connectivity management space is dominated by telecommunications companies, as expected. They have struggled in this arena because of their usual definition of clients. There was much heated discussion in Telecom Italia at the turn of the century, for example, when it was clear that IoT will become "network users," between those who spearheaded supporting the IoT and those who saw them as very low income generators and as such to be turned down as independent "clients". That was a mistake, and today it is becoming clear that the value of the traditional client is decreasing. The value of IoT is not increasing but their volume is, making up for the tiny revenues generated by a single instance. Besides, controlling the IoT would have allowed the companies to play the role of intermediator towards the application and services. Telecommunication companies are now scrambling to offer connectivity services tailored to IoT, but they may have lost the bigger opportunity of becoming the infrastructure to manage them and leverage them towards becoming service providers.

The biggest business slice is in the applications platforms which offer horizontal (like MS, AWS) or vertical service support (like Siemens for Industry 4.0, SAP for enterprise management). There are also companies that support both segments (like IBM).

3. Data Platforms

IoT convert atoms (and more) into data. The data are the raw material in the Digital Transformation and need to be harvested, stored, correlated and analyzed in order to extract value. Data platforms provide this functionality. Most of the time the owner of the platform is not the owner of the data, and quite often there may be several owners of data on the same platform. Also, note that a data platform may not exist in a specific physical location. It can be hosted on a technology platform such as AWS or Azure.

Figure 6 The top 16 data analytics platforms according to InformationWeek.

A data platform should provide, and it usually does, functionalities for neutralization of the data (i.e., its decoupling from the originator and/or owner), for authentication (making sure that the data is trustworthy by authenticating the source) and for monitoring (tracing the use of the data). Additionally, it can provide services, like generation of alert signals when a data value changes or the correlation of several data creating a metadata that reflects the status of the correlation.

Neutralization of data is tricky when dealing with multiple streams. As an example, the position of a phone can be decoupled from its owner. However, if the system supports the correlation of several data streams, like the subsequent position of the phone and its monitoring over a period of time (like a few weeks) it may become straightforward to identify the owner of that phone by realizing that such a phone is moving every morning during working days from A to B. That would suggest the owner lives in A and works in B. More analysis may identify the owner. The approach to neutralization should be structured to make external correlation feasible, i.e., correlation should be a service provided by the platform resulting in metadata. It is this metadata that is released as answers to queries. Rather than releasing specific data values the platform releases aggregated values and metadata to safeguard neutralization.

The authentication of the source (in general not disclosed by the platform), is an important component to build trust on the platform. Watermarking of the data disclosed is a critical function to ensure that the value generated by the data can be tracked and part of it returned to the data owner. This is clearly a grey area, and there are many ways to approach it, none of which are 100% fool-proof.

There are quite a number of "big" companies offering data platforms as shown in Figure 6, and to a certain extent any of the vertical platforms further discussed have data support facilities. In addition, there are plenty of service oriented data platforms, like the ones hosting photos or the ones created by municipalities.

Data platforms are relatively young, just a bit more than 15 years old, with Apache Hadoop being one of the first open source platforms to become widely adopted (e.g., by Yahoo and Facebook), Some data platforms are specifically focusing on some application segment, like Cloudera introducing enterprise support in 2008. Data platforms are expected to keep growing and evolving as more and more data are being produced, harvested and leveraged. In 2003 (note the Berkeley study "How much Information") there was an estimated production of 5 exabytes worldwide; by

2011 the amount had reached 1.8 zettabytes (360 times more, an almost exact doubling every year which is faster than Moore's law); and it will reach 3.8 yottabytes in 2020 which corresponds to 1.5GB per person per day. Most of these data will not be stored or leveraged through data analytics but even a fraction of them is a mind-boggling amount that is fueling the Digital Transformation.
The 31 most current data platforms are presented here.

4. Industry Platforms

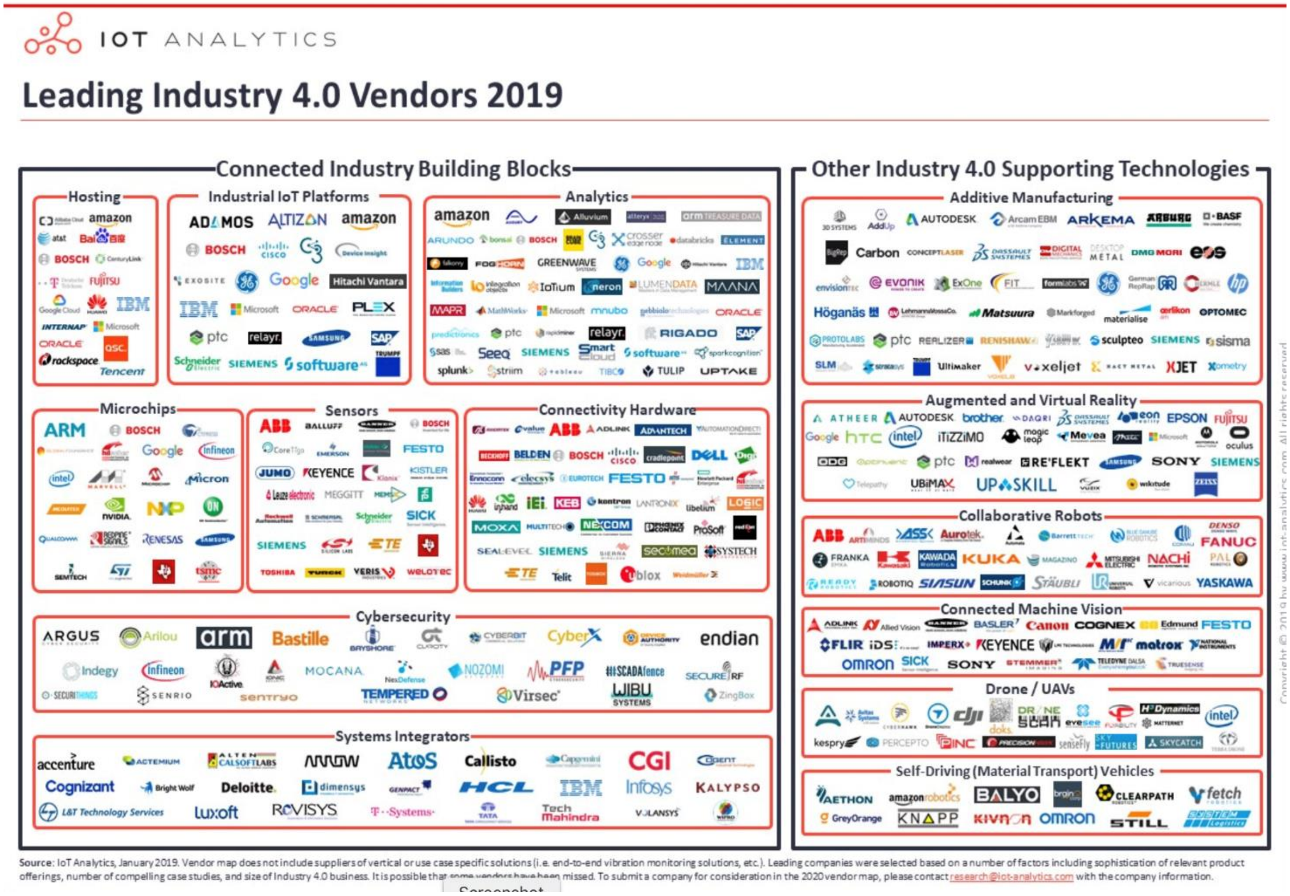

Figure 7 Leading companies in the Industry 4.0 space. Image credit: IOT Analytics

Industry has adopted a number of standards and platforms in the various phases of the value chain, platforms for the supply chain, for manufacturing, and for delivery. In these last two decades the new paradigm of Industry 4.0 has been defined and is now making traction in several industries.

Industry 4.0 is the result of technology evolution (sensors/IoT, robotization, softwarization, Artificial Intelligence and autonomous systems), of a global value chain with a multiplicity of suppliers on one side and a multiplicity of markets on the other, and the possibility to transfer and execute a significant part of the processes in cyberspace (CAD, data analytics, simulation, additive manufacturing, Digital Twins).

It is important to recognize that Industry 4.0 (or the evolution of Industry 3.0) is not about technology but rather about a changing the world market where of course technology is a crucial enabler. The softwarization of products is accelerating their evolution (the life cycle gets shorter), making it possible to go for scale and niche/customization at the same time, making it possible to leverage an ecosystem (if you look from the point of view of the big players) and also enabling small players to take advantage on products created by the big companies. The acceleration is in the market, and industry must respond. The response is to further shorten the product's life cycle and increase the spectrum of the features possible by leveraging cyberspace, i.e., the Digital Transformation.

Data are becoming more and more central to industry, as a tool to increase industry effectiveness (in supply manufacturing and delivery), as a way to increase features, and as a way to decrease operation/maintenance cost turning these phases into opportunities to deliver services. Several

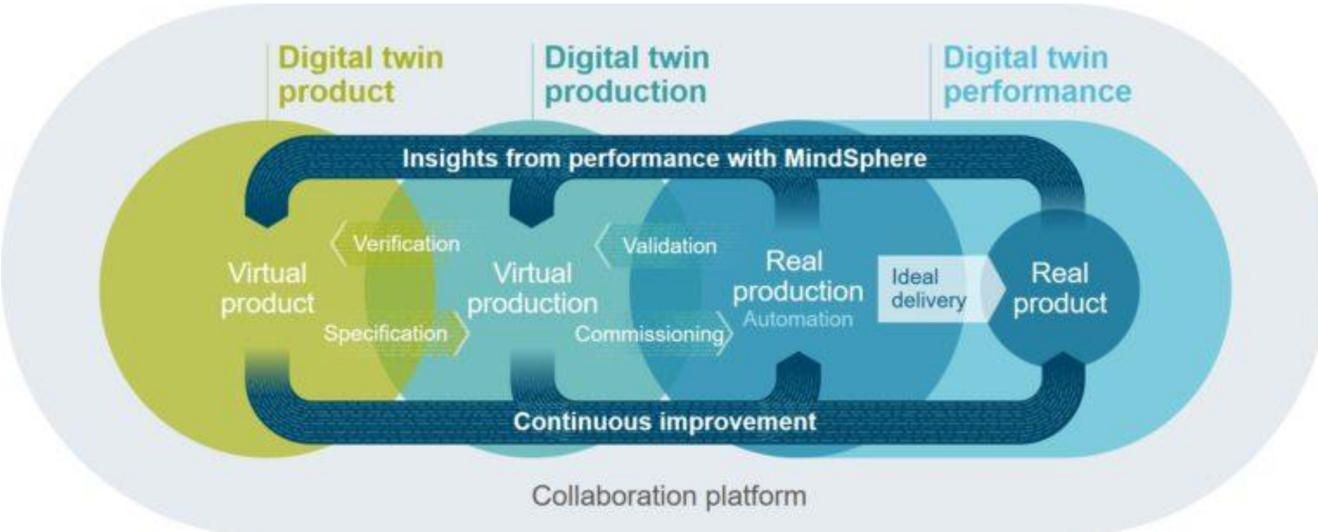

Figure 8 The Collaborative Platform supporting Digital Twins as the glue across the various phases of the manufacturing value chain. Image credit: Siemens

companies are leveraging this transformation each one targeting a specific aspect (see Figure 8).

Digital Twins are becoming a crucial technology in Industry 4.0. Siemens identifies three roles for Digital Twins in industry:

- **Product Digital Twins**: provide a virtual representation of a product. They can be used as the starting point for designing a product, simulating its characteristics and how it would fit in the market and cooperate with other entities;
- **Production Digital Twins**: used to validate production processes before anything goes into production on the shop floor. All manufacturing equipment have their digital twins which can be used for simulating the whole production process;
- **Performance Digital Twins**: data are generated throughout the production process and by the product during its operation. These data are captured by the performance digital twin for actionable insight and decision making.

Digital twins are supported by the Siemens Product Lifecycle Management (PLM) platform and by the other software packages forming Xcelerator, a complete suite applicable to several industry verticals, including support to Electronic Design Automation (EDA), Application Lifecycle Management (ALM), Manufacturing Operations Management (MOM), Embedded Software and Internet of Things (IoT).

GE, one of the first to apply Digital Twin technology in manufacturing and in the digital transformation of industry, spun off a slate of software packages that have now become the Predix Platform for industrial machines. The software is made available in the form of Platform as a Service (PaaS).

As one can imagine there are a variety of platforms used by the industry, each one specialized for a given vertical. It is important to note the shift from management of data (Industry 3.0) to the leveraging of data (Industry 4.0) and the use of data similar to autonomous systems taking a life of their own.

5. Automotive Industry

In the manufacturing industry the car platform is a quite different sort of platform. It has been used by the car manufacturing industry for over 50 years now as a way to decrease the cost of production. The idea is to have a common base: floor pan, drivetrain, suspensions and axles, and to build different cars on that common base. The customer will see, choose and pay on the basis of what is built on the platform. Indeed, most customers and drivers will not even realize that their Cadillac has exactly the same "platform" as a much cheaper Chevrolet. That is the beauty for a car manufacturer: decreasing cost and decoupling it from the price the market is willing to pay.

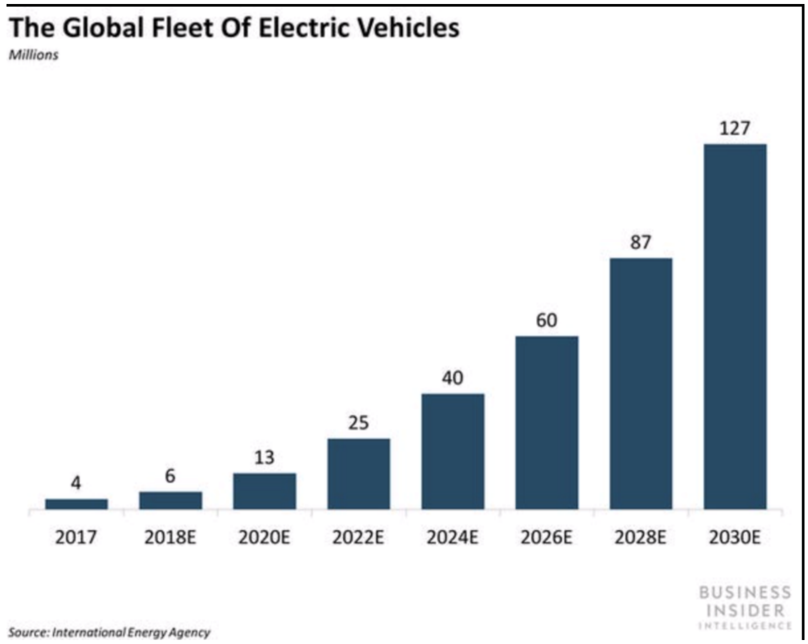

Figure 9 The expected growth in production and sales of Electric Vehicles. Image credit: Business Insider Intelligence

In the beginning the platform was designed to be used by a specific car manufacturer but since the decoupling was so successful several manufacturers agreed to share the same platform, to further decrease cost.

It may seem inappropriate to describe a platform that is rooted in the classical "old-time" industry in relation to the present industry evolution. As a matter of fact, however, in the last 60 years, the car has evolved dramatically, more in terms of production than in terms of end result (although it is a perceived fact that today's cars are much more performant, energy savvy, and safer than their relatives from 60 years ago).

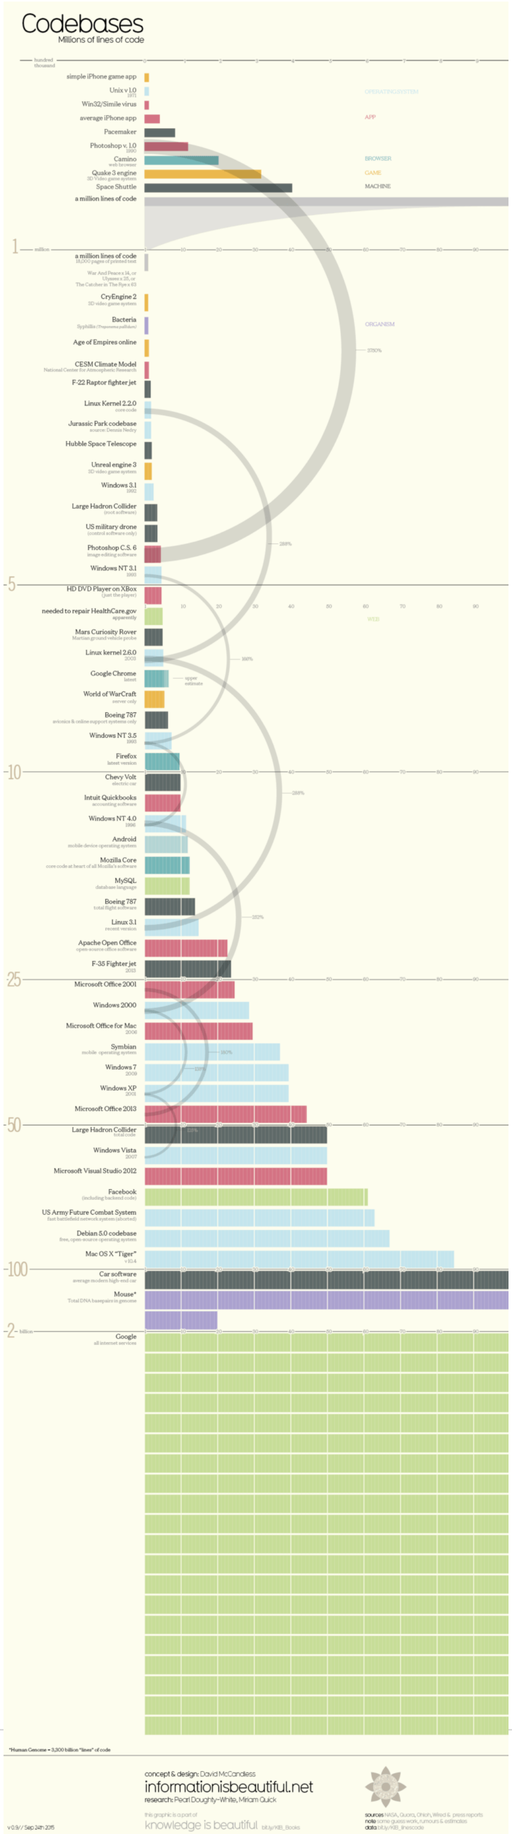

Figure 10 An interesting comparison of lines of code used in different "products". Cars have more than a Boeing airliner. Image credit: David McCandless [*https://informationisbeautiful.net/visualizations/million-lines-of-code/*](https://informationisbeautiful.net/visualizations/million-lines-of-code/)

Robotization in the assembly line is now an accomplished fact as is the computer design of cars and computer simulation of their performance. Even more important, the whole supply and delivery chain is now controlled by computers, and many of the car's functionalities are software based.

Over these last decades cars have become computers with four wheels. Today the cost of a car is unbalanced towards the software and electronics, both in terms of design as well as manufacturing. The shift towards the transformation of cars from product to services has started and is part of the digital transformation of the industry.

However, much remains to be done. The software in our cars is a collection, mostly uncoordinated, of software designed and produced by different companies with plenty of duplication. An average car purchased today may have over 100 million lines of code, and these are expected to double or triple in the next decade. These lines of code are distributed over 30 to 50 ECUs (Electronic Control Units) communicating across CANs (Controlled Area Networks, note the plural for networks). Now if 100 million lines of code seem excessive, consider that a fighter plane like the F-22 has less than 2 million lines of code, and the newest Boeing 787 has "just" 14 million lines of code. In comparison the number of lines of code in a car becomes even more staggering.

Is it because cars are more complex than a Boeing 787? No, it is because cars are not an integrated system; they are an assembly of many systems. Rather than having a single software system, a variety of software systems are used, and the number of lines of code skyrocket.

The car industry has evolved linearly over these decades, but now it is facing a revolution, actually a double revolution. On the one hand electric motors will be replacing internal combustion engines, and on the other hand, autonomous driving systems will be replacing the driver. Both revolutions will occur over the next two decades.

As shown in Figure 9, in 2029-2030 the car industry expects to sell 40 million electric cars; that is, 1 out of 3 cars will be an electric car (assuming 60 million total cars sold per year in that period. A peak was reached in 2016 with 72 million cars sold excluding commercial vehicles. Over the last 3 years there has been a reduction in sales with a small but growing erosion created by electric cars sale). In the following decade the impact of electric cars will be significant on today's car manufacturers and on their manufacturing processes. An electric car requires different kinds of labor skills and is even more amenable to automated production and assembly than a conventional car. Looming job reduction is already generating concern in unions in the US. Electrical vehicles will rely even more on software for their management.

The revolution in driving (autonomous driving) is flanking the electric cars revolution and will lead to an even more radical change in the architecture of the vehicle. More software will be involved but, most likely, future cars will have fewer lines of code than today's cars because their software will be designed top down as an integrated system.

Both revolutions will give rise to new industry platforms, this time oriented to support software creation and integration. The digital transformation will be fully realized in these platforms, and also in the business environment leading to the creation of an ecosystem of providers of features implemented through software.

There are already several companies offering software for connected cars; there is an initiative for a car operating system (Automotive Grade Linux); and a few platforms are emerging in the area of self driving cars. Additionally four companies, Apple, BlackBerry, Ford and Google, are creating and delivering platforms to create and host smart car applications.

6. Smart Cities Platforms

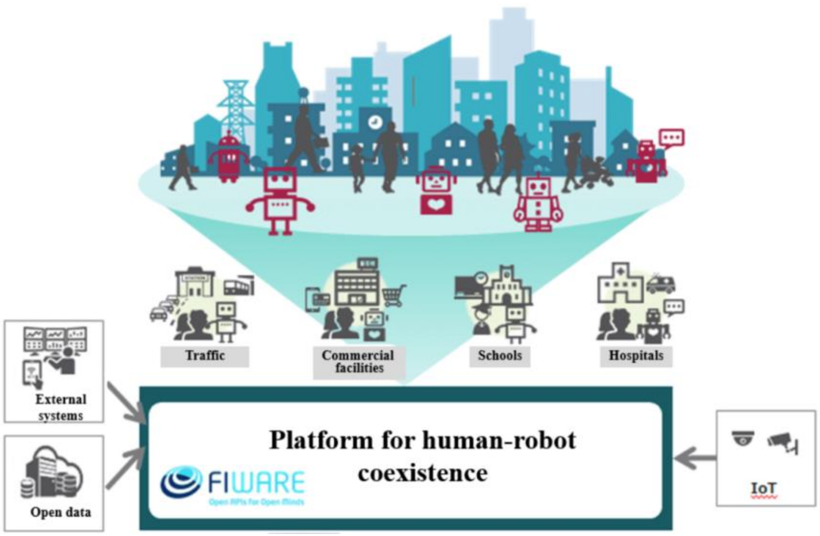

Figure 11 The FIWARE platform has been specifically designed to enable smart cities development by aggregating data generated by sensors disseminated in the city and hosting a variety of services. Image Credit: FIWARE Foundation

Industry platforms were addressed in previous sections noting the multitude of data that are created by industry processes and in perspective by the products as they will be used. Digital Twins play a crucial role in that context as Digital Transformation takes over and permeates industry.

Now, scale all that by at least two orders of magnitude and enter the smart city arena. Here the complexity is much higher. In an industry, and even in the manufacturing value chain involving several players, data are structured and the overall processes are well defined. Not so in a city.

A city is a living organism in a dynamic equilibrium, challenged every day by new situations. Besides, a fundamental component of cities is the citizens who are by nature unstructured.

A city generates an amazing quantity of data streams, most of them created independently of one another (Amsterdam was found to have 12,000 data sets and corresponding streams, basically unrelated), and yet the real value is derived by applying correlation or data analytics to these originally independent data streams. It is this variety and the ever changing dynamics that make cities different from industry. Hence the need for specific support.

There are a number of platforms available whose main goal is to collect data from the variety of IoT (sensors) disseminated in the city, store it in a way that data analytics can be effective and open it up, in a controlled way, to third party use. A crucial aspect in this data management is to ensure both privacy and openness, so that those data can become valuable. This seems to be a contradiction in terms, yet it is exactly the goal of a city platform. Actually, the better a platform is at preserving data privacy and opening up data, the more valuable the platform.

Neutralization of data is a first step but once many data streams are opened up, correlation may circumvent neutralization. More sophisticated approaches to privacy preservation need to be put in place. An approach is to create metadata in the platform and release only these metadata, not allowing the direct access to the raw data provided by sensors. In this way there is a hard decoupling that is more robust in preserving privacy. Another, complementary approach, is through some sort of watermarking or blockchaining the data values released by the platform so that it is possible to track their use and take action in case of improper usage.

Platforms operating in an industrial framework are tasked with ensuring the quality of the data, which implies making sure that the sensors providing the data are working in a proper way and within a pre-defined tolerance range. On the contrary, platforms operating in a smart city framework are tasked with ensuring the quality of the metadata, in spite of the possibly low quality of the raw data provided by sensors. This is because the sensors are (usually) not owned by the entity (municipality) managing the platform. The accuracy of the metadata can be ensured by analysis of the redundant data provided by the large number of sensors.

Another crucial difference between an industrial platform and a smart city platform is the support of reusability of services. Applications operating on an industrial platform are generally customized to the specific industrial environment they support. A different industry will develop its own specific application. A smart city platform should manage data and open data in such a way that applications developed for a certain city can be reused with minimal effort for other cities.

This is what characterizes, as an example, FIWARE, a platform resulting from cooperative work promoted by the European Future Internet Program. It is important to note that FIWARE is a platform designed to foster the Digital Transformation of cities, leveraging technology but having the goal to execute the Digital Transformation, which implies supporting the variety of processes in a city in an economically affordable way. Part of this affordability is rooted in the possibility of reusing applications and services.

7. Healthcare Platforms

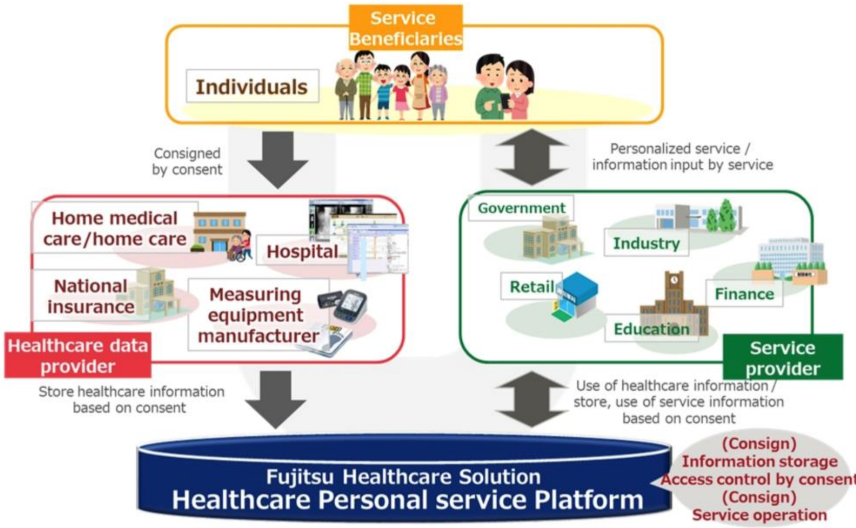

Figure 12 Personal Healthcare platforms can be a significant part of the future in Healthcare. Image credit: Fujitsu

The previous section presented smart cities platforms noting the importance of privacy since personal data is prevalent in smart cities. These concerns are increased by a few orders of magnitude in the domain of healthcare platforms.

There are currently several healthcare platforms in operation supporting hospital processes and healthcare processes in general, but the future is about healthcare platforms supporting people's health.

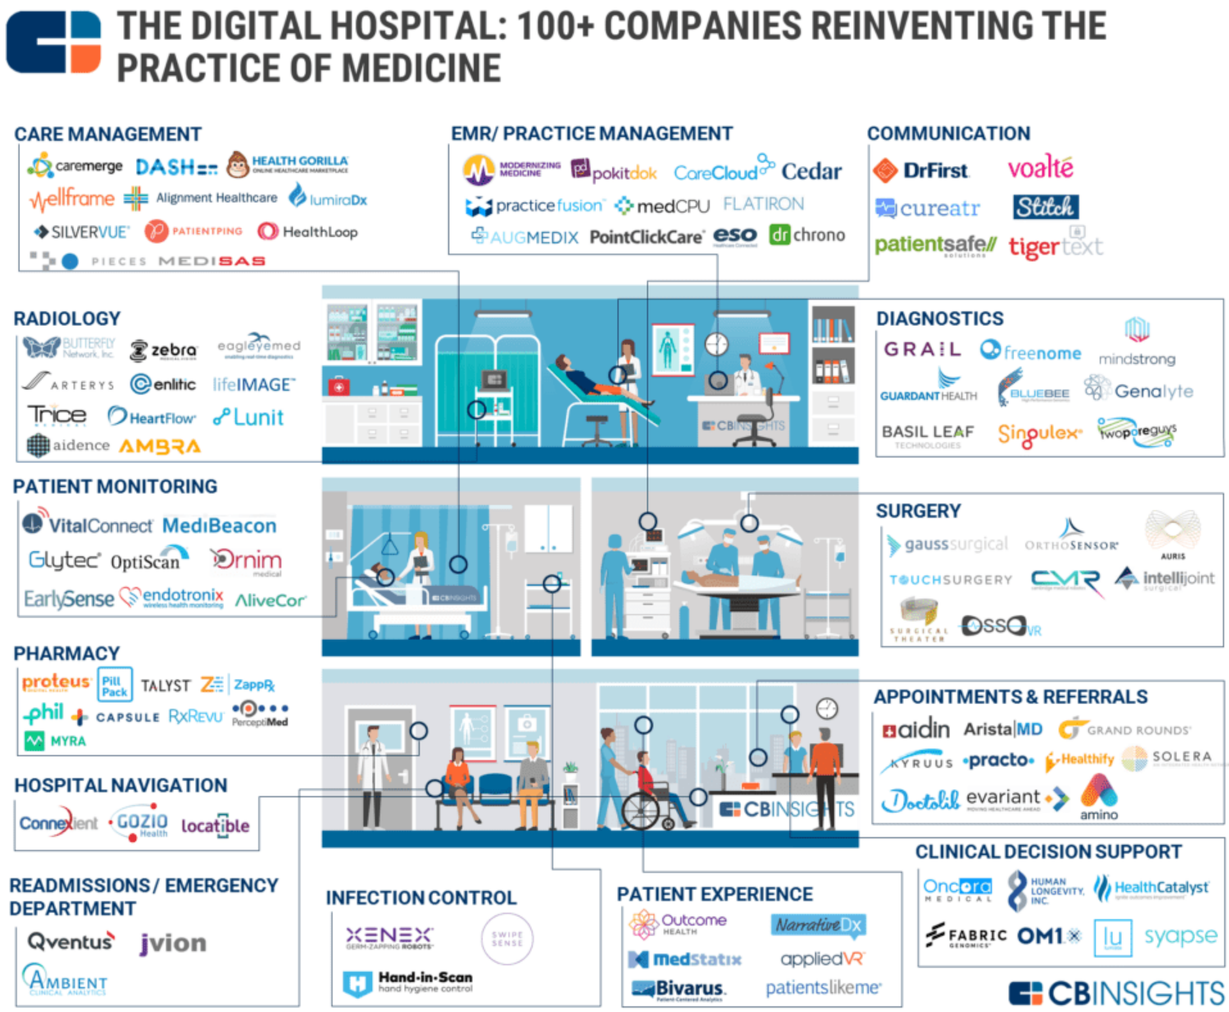

Figure 13 The ecosystem of digital health with companies clustered in various areas. Image credit: CBINSIGHTS

The healthcare industry is undergoing a digital transformation but so far any real transformation has been slow. Gartner estimated a penetration of Digital Transformation in Healthcare to be around 10%. (Compare this to the music sector that is almost at 100% or entertainment, 50%. Even industry has progressed faster, with an estimated penetration of 20%, low but still double of healthcare). This does not mean that computers, data, and automation have not spread into the healthcare space. Figure 13 depicts over one hundred companies supporting digital healthcare in various areas: care management, practice management, clinical decision support, appointments and referrals, patient monitoring, clinical decision support, pharmacy, and more. And, of course, there are many computers to assist in surgery, infection control, radiology, and so on.

The digital transformation in healthcare requires more than all of this. It needs to leverage data to support the business processes and the business intelligence on one side (see the SAP HANA Platform in Figure 14) and the creation and leveraging of data for personal healthcare.

General Electric is betting on digital twins as a way to link patients to healthcare services and is creating an ecosystem of development partners. Fujitsu has launched a Solution Healthcare Personal

Service Platform that seems to be the direction for a future made possible by the Digital Transformation.

Figure 14 Data generated by IoT (wearable, diagnostic and monitoring devices) are leveraged to support healthcare business intelligence and processes, each one feeding the other in a never ending loop. Image credit: SAP

One of the ways to address the issue of strong control of data, privacy and at the same time maintaining open data so that community healthcare can benefit from personal data of millions of people is to create personal platforms, each one containing and managing the data of a single person (its Digital Twin) and sharing those data on specific conditions, controllable by the owner (the patient). The physical location of these personal healthcare platforms should be immaterial, although some may claim that as an intermediate step the smartphone may serve as our personal platform. The reasoning is that the smartphone is already becoming a hub for wearables which are providing health related data. Of course, the smartphone has its twin in a (trusted) cloud but the perception is that patients will be able to hold their healthcare data in their hands. The smartphone can also be the gateway to access healthcare services provided by chatbots (The book, *Cell* by Robin Cook, presents a scary fictional account of this possible evolution).

8. Device/Product Platforms

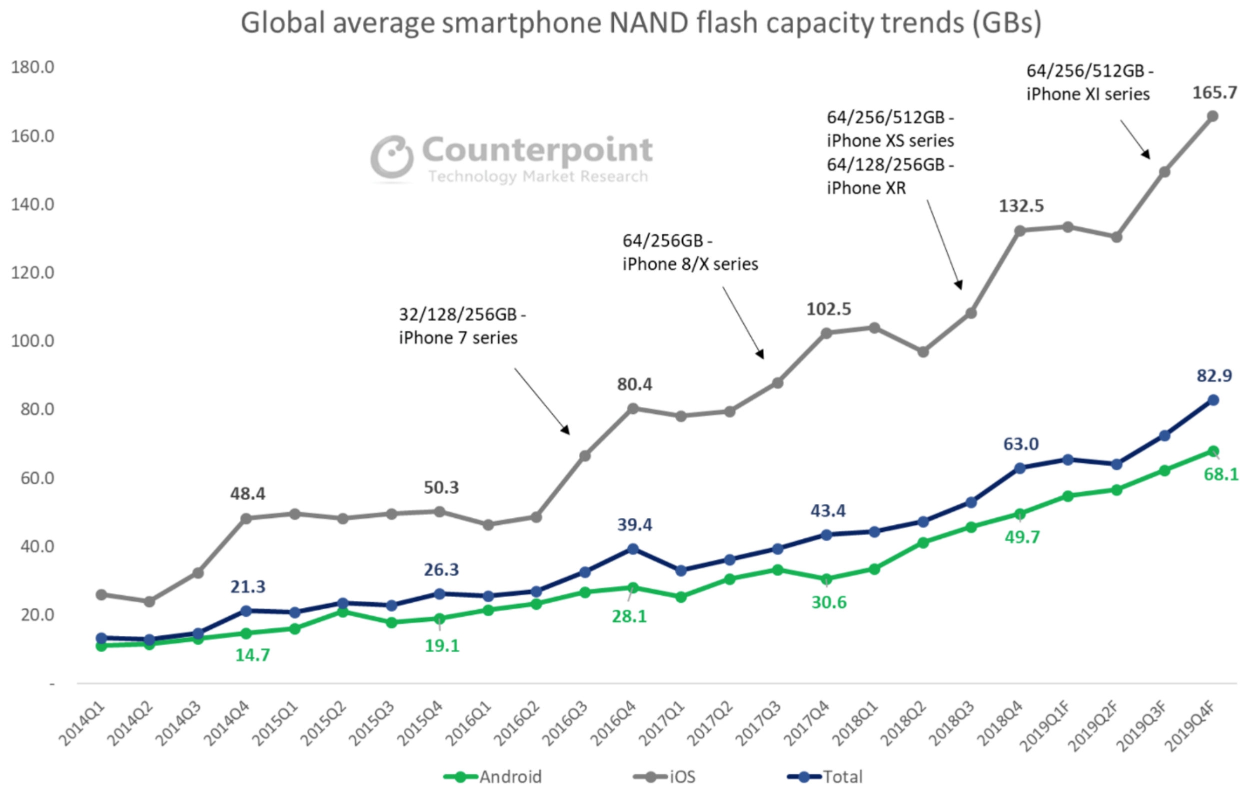

Figure 15 Average storage capacity in smartphones keeps growing. The top of the line smartphones have reached 512GB, and microSD cards to extend storage capacity have decreased in cost (1TB for $30 as of October 2019). Average storage is likely to hit 1 TB in the second half of the next decade, allowing the use of smartphones as massively decentralised data centers and as massively distributed platforms. Image credit: Counterpoint

The beginning of this whitepaper stated that a technical platform supports processing, storage and communications. Consider a smartphone. It supports processing, storage and communications and does much more, like sensing. The power of the smartphone is often underestimated. Its power equals the power of a supercomputer a few decades ago even though it is a far cry from today's data centers. However, if all smartphones (there are 3 billion of them and growing) are considered, their collective power exceeds the power of the largest data center.

Smartphones represent a massively distributed platform. However, looking at smartphones, and more generally at devices, in terms of platforms means considering them not as a technical platform but more in the sense of a market platform.

Creating an app for iOS or Android means creating an app that can address a market of billions of potential customers. This represents the huge value of these platforms.

As a matter of fact, the enabler is the operating system (more monolithic in the case of iOS, with a few variations in the case of Android) but the operating system is just a basic component. What makes the platform effective is the support infrastructure to offer, deploy and sell (which is the purpose of the Apps stores). It is this latter that has stimulated the growth of apps, making it easier for developers to reach the market. Additionally, the infrastructure offers the support tools for developing apps and a brand that captivates the market.

These devices' platforms are going to evolve in the coming years. The smartphones and devices will become even more powerful, with more storage, processing power and connectivity. People are already engaging with devices for a considerable part of our day, and as devices become wearable (and a few embedded), more engagement will occur. In addition, the infrastructure (like the Apps stores and developer kits) will provide even more support for the go-to-market.

As powering of these devices becomes less problematic (wireless charging that will keep your devices always charged as you move from one ambient to the other), it will also become possible to leverage the technical platform created by clusters of devices. This will increase the power of each device, since whenever needs arise the device will be able to harvest the spare capacity of nearby devices.

In the second part of the next decade, vehicles will become platforms in terms of apps, as are smartphones today. Soon it is expected that vehicles in an urban environment will become a communications platform, i.e., to manage connectivity, each vehicle will become a network node. 5G will be providing the first basic functionality although 6G will be required to see vehicles and other devices become real network nodes.

As a closing point, each smartphone is expected to become a technical platform hosting much more sophisticated personalization functionality. Each person's smartphone will store personal data and the software to leverage these personal data using a digital twin. The data harvested from the web will be displayed, when needed, on the smartphone screen, or on any screen in the surroundings that would better fit the task, rendered in a way that makes those data significant. So, if two people make the exact same query on the web, different renderings of the query result will be presented, since the smartphones will create different visualizations to fit the specific interest and knowledge level of each person. Notice that this rendering process will take into account not just preferences (use table rather than graphics, one font or another), otherwise called a syntactical rendering, but also be based on knowledge, skill, and experience, called semantic rendering.

9. Knowledge Platforms

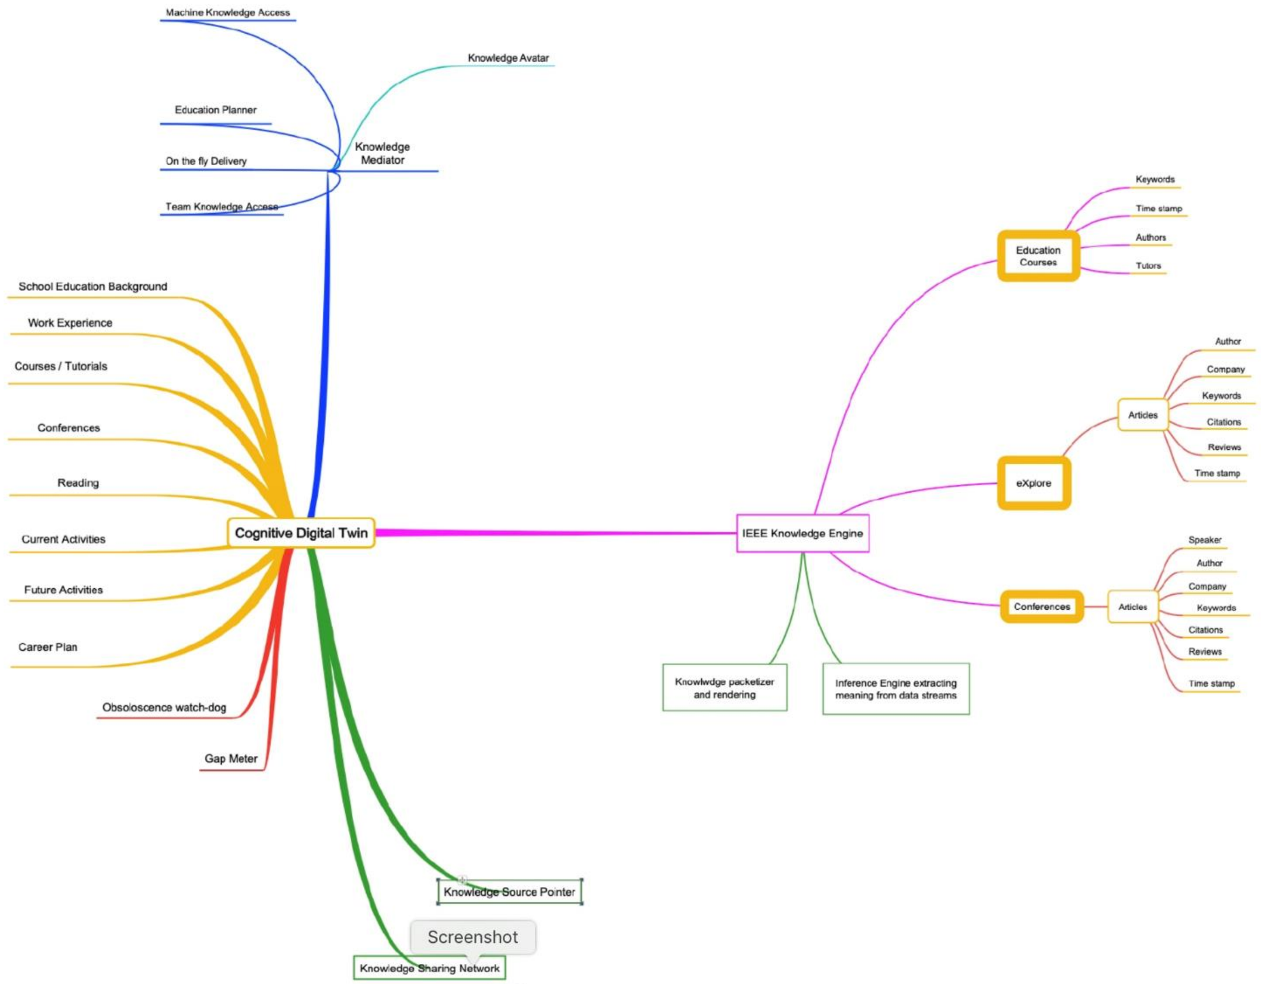

Figure 16 Schematics of a cognitive digital twin and its use to detect knowledge gaps. Image credit: IEEE Future Directions

The last section concluded with a discussion of "knowledge" and pointing out that smartphones, which can become or embed digital twins, might become knowledge intermediators in our access to the knowledge on the web. This ties in nicely with the last category of platforms: knowledge platforms.

A Google search of knowledge platforms will result in a list of companies and services delivering various forms of knowledge, each one specialized in a given sector, like agriculture, high school math, do-it-yourself advice and so on.

The knowledge platform defined here is something different: a platform that supports the creation and sharing of knowledge at a semantic level, taking into account the needs and capabilities of the person or entity requiring such knowledge.

Today there are plenty of platform dealing at a "syntactic" level, easing the exchange of documents in various forms. The various education platforms are an example, like Coursera and edX. IEEE Xplore® is another example of a platform for sharing documents. These services are platforms rather than databases because they are providing a sort of rubber stamp on the documents being shared (in the case of IEEE Xplore all documents have been peer-reviewed, in the case of Coursera and edX the courses have been checked for quality and so on).

Current education platforms are not that different from general content platforms, like Netflix or YouTube. These latter also have (sometimes minimally) curated content, and sometimes that content is also of educational nature, designed to share knowledge. However, there is little, if any, consideration to the user capability to acquire the knowledge in the way it is presented.

On YouTube I can watch a clip in Chinese, and the transfer of knowledge, given my understanding of Chinese, will be zero. Yet, YouTube will duly show me that Chinese clip. Similarly, I can ask to follow a course on Coursera without the pre-required knowledge that would allow me to understand it, hence never acquire the knowledge being presented. I can log on IEEE Xplore and read a paper that although written in English looks Arabic to me. Again: net transfer of knowledge equals zero.

There are now studies on Cognitive Digital Twins that may result in the creation of true knowledge platforms.

IEEE Future Directions is initiating an ambitious project to develop a knowledge platform that can transform the huge content base of IEEE Xplore into Knowledge-as-a-Service (KaaS). This will be done by creating cognitive digital twins of the users of IEEE Xplore (using an opt-in approach) and having those digital twins interact with the KaaS to transfer knowledge on demand as well as to be notified on the availability of knowledge that may be needed by that person, given her current activities or the ones expected in the near future. As an example, a cognitive digital twin of a person who took a Wireless Certification (WCET) this year will be notified of the evolution in wireless technology, systems, applications that suggests the need for updating the knowledge of that person. The cognitive digital twin, being aware of the current and short term needs of the person will interact with IEEE Xplore and the IEEE education courses portfolio to design a roadmap of knowledge transfer that may fit that person (also taking into account her time availability).

The project is staged in three phases, one per each of the coming years:

- 2020 – delivery of a web interface to create the person’s cognitive digital twin
- 2021 – delivery of tools to assess needs and extract/render knowledge
- 2022 – delivery of a knowledge map, dynamically tuned based on the changing user needs and on the evolving knowledge space that can be used to identify knowledge gaps and actions to fill them

The program is first targeting the needs of personal knowledge and then the one of enterprise, institution knowledge with the ambition of providing knowledge services of various types, from course to consultancy and resources (people, companies, AI tools).

Something that gets close to this concept of knowledge platform is the one of UnanimousAI. This is a company and a service, that shares knowledge in the medical domain. It assumes that those connecting to it are medical doctors, hence it presents knowledge in a way that is customized for them. The knowledge base is created in two ways:

- by roaming the web to retrieve the latest information made available in the medical field (accessing trusted magazines, journals and data banks in the medical area) and extracting knowledge using artificial intelligence software;

- by growing the knowledge base through interaction with medical professionals that present cases in their practice every day, asking for references to other doctors who may have faced similar cases.

The access to the knowledge base is mediated by natural language based interfaces, and the more interactions take place the more UnanimousAI grow its knowledge through Machine Learning.

10. Platforms Supporting the Gig Economy

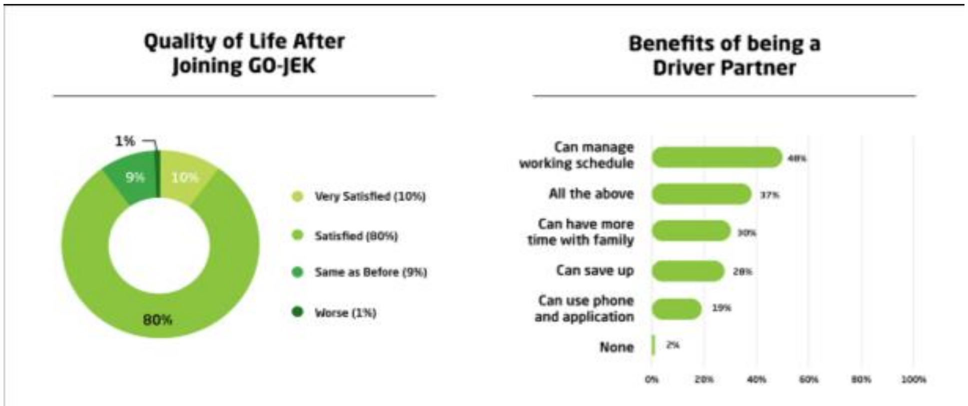

Figure 17 The majority of GO-JEK drivers are satisfied with the increased income since they joined GO-JEK and 90% of them are reporting an increased quality of life. Image credit: GO-JEK

Platforms are slashing the transaction cost in connecting providers of products and services to the marketplace. As such they are ideal tools to support the Gig Economy.

A clear example linking offering with demand can be found in platforms like Gojek and YEGOMOTO, one supporting a variety of services in Indonesia starting with Uber-like transportation, and the other also Uber-like but using motorcycles in Rwanda. These two countries are emerging economies and have little business infrastructures compared to Western countries, and thus creating business is much more difficult. Not so with the Digital Transformation, where power is shifting to small players (or, at least, small players get a chance). Platforms, like the ones just mentioned, enable a sparse but numerous "workforce" to reach the marketplace.

Platforms are lowering the market entry levels while at the same time providing a digital infrastructure that flanks (upgrading) physical infrastructures where they exist or provides a viable alternative to them (as might be the case in emerging economies). Of course a minimum of physical infrastructures is required to access the digital ones, i.e., a pervasive, affordable, wireless network which is now (almost) ubiquitously available, even in emerging economies.

It is interesting to note that while in established economies platforms are decreasing the market value in the specific segment addressed by the platform by decreasing the transaction cost (hence a competitive market results in a decreased price to the end user), in emerging countries, platforms are enabling a new market, hence creating jobs and value (that market goes from not existing to adding some value). It is also important to notice that platforms, supporting transactions in the digital space, have very little friction, hence the market growth can be very quick (no scaling issues, something that must be addressed in the physical world). This can be seen in platforms like Gojek mentioned before that have seen a tremendous growth: from 100,000 orders a day in 2015 to over 100 million orders a day in 2018 (a thousand-fold increase). It is also important to notice that once a platform has been established to support a service (like car hailing), it is straightforward to have the same platform support other services. Gojek has grown from one service in 2015 to 18 services in 2018. In other words, platforms become neutral infrastructures to support massively distributed business. This is clearly supporting and growing the gig economy. Again, while established economies are seeing a growing concern of job depreciation (companies are turning to the gig economy to save money, e.g., in food distribution undercutting pre-existing salary levels), in emerging economies there is not a competition with existing jobs but rather job creation. As mentioned earlier, in Indonesia and Rwanda, the platforms have created jobs, resulting in an increased wealth and in a satisfaction of those gig workers that is reaching 90%.

11. Platforms as Enablers of Business (and Services)

In general, key people (CEOs, CTOs, Directors) in several industries covering a broad spectrum of verticals, from mechanical to textile, from food to transportation, have indicated that industry is not interested in technology (and specifically in platform technology) but in what a platform can do to the bottom line of a company. A nice citation mentioned Theodore Levitt, economist and professor at Harvard, saying:

People don't want a quarter inch drill; they want a quarter inch hole.

Industries are seeing a transformation of the supply and delivery chain and a different relationship with the market that is bringing platforms and ecosystems to the fore. Platforms are seen as a way to reduce cost, increase efficiency and, even more importantly, as a way to connect the platform owner to the market. This connection generates data which can be leveraged to fine tune products and to offer services. Additionally, platforms are seen as a tool supporting new business models and in particular the "**Outcome Economy**."

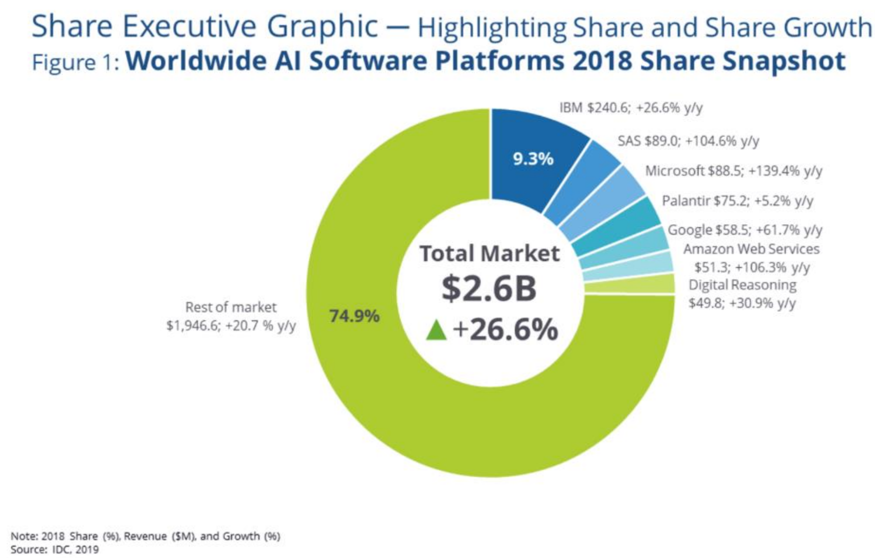

Figure 18 The worldwide AI platforms market in 2018 is $2.6 billion. AI is seen by several industries as a "must" to make sense of the huge amount of data produced along the value chain. Graphic source: IDC

The Outcome Economy is a concept that has been socialized by the Economic World Forum noting that:

> *Companies will shift from competing through selling products and services to competing on delivering measurable results important to the customer.*

This means that rather than charging a specific price for a product or a service, a company will charge a fee based on the end result that the product or service will provide for the end customer. This approach is becoming possible because of connectivity and IoT. IoT makes this business model possible by providing product performance and use metrics to the provider. McKinsey expects the Industrial IoT (IIoT) market to reach $7.7 trillion by 2025 with SAS foreseeing 55 billion IIoT in operation at that time.

Notice that this is good for the user since costs are based on the value accrued, and it is good for the provider because it allows a continuous relationship with the user fostering further business opportunities. Besides, this model avoids the risk of commoditization since the platforms connect the producer to the customer/user, and this connection that delivers the specific value. Platforms, controlled by the producer, are the tools that support the tracking and the connectivity. The platform accrues usage data that in turn can be monetized.

Figure 19 Textile manufacturing involves very sophisticated and extensive plants having multiple hundred thousand components (320,000 bearings, 1,250 motors, 3,500 transmission belts, 50,000 spindles to produce 43,000 kg of fabric a day). Downtime is extremely costly and can kill margin. Image credit: Indian textile journal

An example is the growing interest of car insurance companies to offer more advantageous pricing based on actual use by having the customer install a tracking device in his car. The device tracks such metrics as how far and where the car is driven, how fast the car is driven, and so on.

Another example is the sellers of complex systems, like textile manufacturing equipment (converting fiber into yarn, yarn into fabric, fabric into dyed or printed fabric and then transformed into clothes) that is packaged with maintenance functionality, transforming maintenance into preventive maintenance by using IoT and predictive software. In this area, as in several others, the margins are razor thin: the spinning of yarn to create fabric produces a revenue of $3 per kilogram with a margin of a few cents per kilogram. The plant works 8,500 hours per year (there are only 8,760 hours per year so essentially continuous operation), hence preventive maintenance is crucial to avoid downtime. These plants are operating in countries far away from the producer of the equipment, and remote monitoring to provide proactive and preventive maintenance is essential. The digital transformation is changing the value chain and making these new business models possible.

12. Business-oriented Application Programming Interfaces

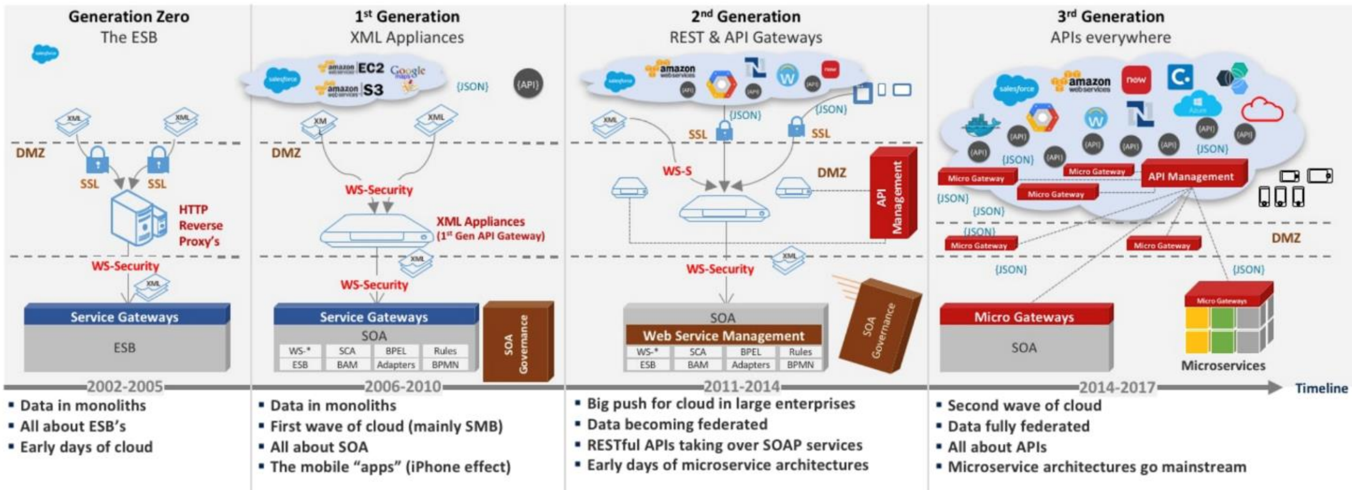

Figure 20 The evolution of APIs, from API to access monolithic data to API over distributed clouds accessing federated data. Image credit: Oracle

The previous section described the relevance platforms and ecosystems have to industry. Actually, the two are connected, or better yet, need to be connected to match industry interest. A platform provides its services through Application Programming Interfaces (APIs). For industry these APIs are crucial and need to be designed having business in mind, while in many cases platforms that result from research endeavors have APIs designed from a technical point of view. This is a strong concern of many industry leaders. The creation of an ecosystem is important, and in many cases it is what generates revenues. However, industry needs to orchestrate this ecosystem, and this cannot be done using contracts, since there is no direct relationship between the industry owning and opening the platform and the variety of players that develop platform features (data and applications). APIs become the contracts regulating the access.

As an example, the Trentino Open Data Platform is clustering and opening data enabling service creation. The APIs have been designed to ensure the access to metadata hiding the raw data (to preserve ownership and privacy), and the interactions are monitored for possible revenue sharing.

As mentioned previously, there are the big "platforms", probably better identified as ***data infrastructures*** (the ones of the Google, Microsoft, Amazon, Facebook, IBM and Apple, - "G-MAFIA", Baido, Alibaba, Tencent - "BAT"), and then there are many smaller platforms serving the digital transformation of industry which includes the ones provided by several Telecom Operators. These smaller platform providers cannot compete in terms of price and performance with the bigger ones but serve specific industry needs supporting industry driven ecosystems, the transformation of products into services, or the flanking of services to products. These latter are the enabler for both the outcome economy and the further step to the pull economy where the revenues and pricing is no longer based on the sale of a product/service but by the revenues that product/service generates. This requires much better knowledge of the client and sharing of data among the industry and its clients.

Another example is an agreement with the municipality of Turin for the deployment of 5G. The offer to the municipality includes the creation of clouds at the edges (some 10-20 clouds for the city of Turin, 1 million people) that can be used by industries as competitive tools. 5G is indeed a good tool for supporting cloud at the edges with its flexible access architecture potentially orchestrated by the end user. This evolution is supported, and requires, a third generation of API. API are interfaces which have become more and more sophisticated providing a more effective and "abstract" way of interacting including a shift from needing to know syntax to needing to know semantics:

- **Interface basics**: Different data bases, each one monolithic, can be located in a cloud and the access is via an Enterprise Service Bus. Accessing and exchanging data requires knowledge of how the data is identified and where it is stored.
- **First generation API** were part of a Service Oriented Architecture (SOA) where sharing of files were supported by Server Message Blocks. Data remained monolithic, and the overall framework was the one of XML (eXtended Mark-up Language) and object oriented programming (SOAP – Simple Object Access Protocol) used by Apps. The service is the unit interacting (in basic interfaces the basic interacting unit is the instruction code), and from a perceptual point of view users started to look for information. For example, rather than looking for a repository of the weather forecast and then looking for the specific relevant data, the user started with a location of interest and the relevant weather information was displayed.
- **Second generation API**, adopted by enterprises, particularly large enterprises making use of Cloud, needing to pull together a variety of data/data bases (federated data), shifted to micro-service architecture to increase flexibility. Here there is a shift from SOAP to RESTful API (HTTP REquest for Get, Put, Post, and Delete data). The access is to metadata because these are providing the information. The interfacing occurs at a semantic level (for example, I need to go from Milan to Turin), and the code looked at the various possibilities searching different data bases returning the information resulting from the aggregation and rendering of various sources.
- **Third generation API** is the growing wave of API supporting data and applications distributed on clouds at the edges, with fully federated data and micro-service architecture, particularly useful in the ecosystem context. The focus is on the relationships among data distributed in clouds. Different from the second generation, processing occurs in each single cloud; there is not necessarily a master aggregating data. This ensures privacy and ownership protection. The final computation or rendering takes place at the user side customizing the rendering to the specific contextualized need. This is particularly relevant in creating value through Augmented Reality where information flow needs to be contextualized to the local ambient. It is also a crucial component/tool in the Digital Transformation.

This third wave is much more semantic oriented in terms of interactions and comes with a set of tools allowing the development of APIs and their sharing plus the availability of services like the one to support security in the interactions (OWASP).

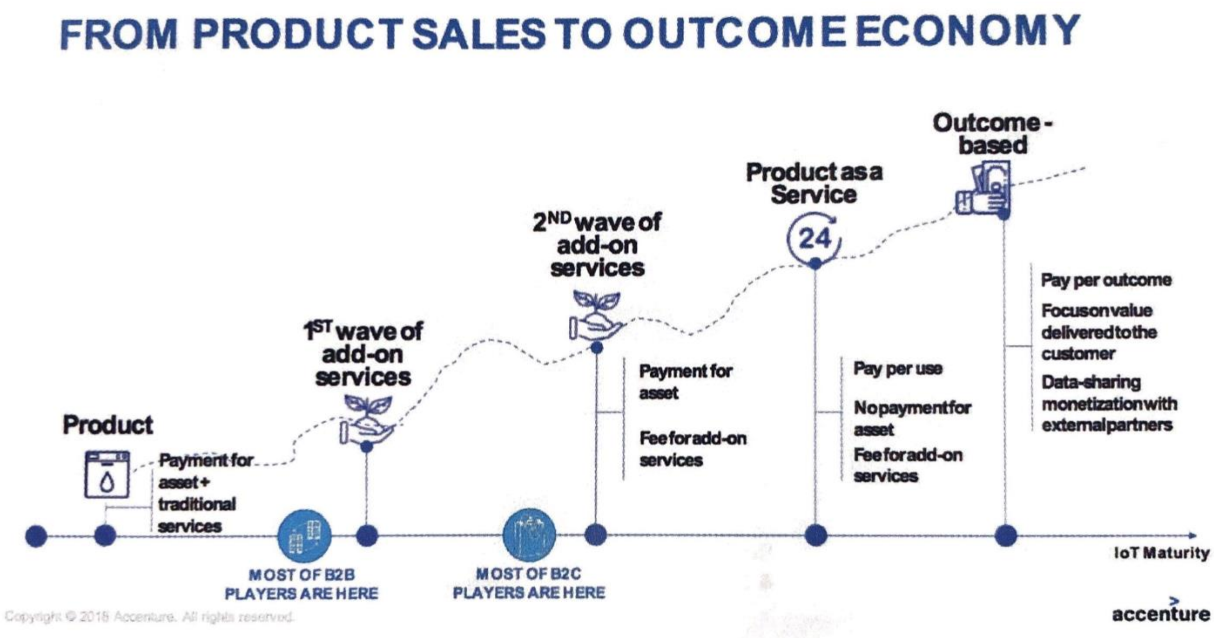

Figure 21 The steps in the shift from products to outcome economy. Image credit: Accenture – Industry Big Event, Milan 2019

Now that the evolution of APIs can handle contextual processing and rendering, the use of platforms will enable of the shift from products to services (and ecosystems). The evolution of products is:

- first towards flanking them with add-on services (see Figure 21),
- then to sell Products as a Service (PaaS) or Platforms as a Service (like the Michelin Effitires, Effifuel, and Effitrailer products that are converting the sales of tires into the sale of services, or Rolls Royce and GE offering turbine power by the hour)
- to finally reach the point of getting a fee in return from revenues generated by the client using a product (outcome economy).

While the intermediate steps are generating recurrent revenues, the latter is leveraging data (IoT data monetization) and is often leading to a provider/client partnership enabled by IoT.

Of course, this is easier said than done. According to Accenture, some 60% of European manufacturing companies are considering the shift but just a few, so far, have actually managed the transition from Proof of Concept (POC) to actual large scale implementation.

Another example of platforms supporting the shift towards manufacturing services is the one of Siemens Additive Manufacturing Network (AMN) designed to support the aggregation of ecosystems. Siemens is providing Product Lifecycle Management (PLM) Software as part of the platform which can work on federated data allowing effective cooperation among different companies, each one using CAD for the design of various components including warehouse management and IoT data aggregation. Elesa, as an example, is a manufacturer of plastic components which needs to manage an inventory of 20 million parts since everything is custom-made. They are using Siemens AMN to produce and collaborate with their partners and with their clients.

Another interesting characteristic of the Siemens platform is the management of Digital Twins. As products are manufactured Digital Twins are created and these contain the start of the Digital Thread[3]. The platform will seamlessly aggregate data from the Digital Shadow and support the use of the Digital Twin to deliver services (like predictive maintenance). The Digital Twins may in the future become active components in the exploration of data lakes produced by IoTs in a given ambient, something that is particularly important in manufacturing plants as well as in the management of clients of specific products/services since this can create knowledge across the usage pool that can feed back into manufacturing and new releases or new products (Industry 4.0).

Actually, Industry 4.0, so called by Western countries, is the new economic model made possible by IoT, data, AI and advanced infrastructures (5G, Edge clouds and Digital Twins and is called Society 5.0 in Japan to convey the deep implication that this evolution is bound to have on society and the fact that the evolution should start from society.

Figure 22 Ever considered buying a 570S McLaren? It costs around $200,000 but if you could get it for $20,000, why not? Image credit: McLaren

Notice how this approach of digital twins, threads and shadows could be mirrored in other sectors and industries, including the IEEE document process, where any article being peer-reviewed could generate a knowledge digital twin embedding the reviews as a thread and supporting knowledge management applications. And, of course, IEEE could start this "revolutionary" approach from the perspective of those, members or not, accessing its knowledge services.

To conclude this white paper, note an aspect that may not be obvious in the transformation of products into services. Services are made of bits and as such can be easily duplicated at basically zero cost. If a manufacturing company manages to transform a product into a service it can decrease cost and can offer a lower price to the customer expanding its market.

Now, for some products this would seem impossible. For products with both hardware and software components, the overall cost can be decreased when considering software, with software being shared among many clients and therefore dramatically enjoying economy of scale. But what about the "hardware part"? As an example, consider a self-driving car. The software can be duplicated over thousands of cars decreasing the cost "per-car" but a car still needs to be manufactured for each customer. This describes the present value chain and the present ownership and usage model of cars.

Consider this: a smartphone costs between $2,000 and $8,000 per kilogram. A car costs between $40 and $70 per kilogram - much cheaper, and one could correctly assume there is less margin in a car than in a smartphone.

A car is so cheap, per kg, because it is used in a marginal way, on average just 5% of the car is used, time-wise. Now consider a self-driving car that can be shared by several users to increase its

[3] A Digital Twin consists of a Digital Model (often created in the design phase) of a Digital Thread aggregating all data gathered during the manufacturing and in the operation of the product and a Digital Shadow mirroring the status of the product at any specific time.

usage time to some 50-75% of the time. You would be able to push the price per kg to $1,000, still cheaper than a smartphone but closer. Actually, a car manufacturer would be able to increase the price of the car to $2,000 per kg, and the end user will still pay less than what he is paying today (some 200$ per kg) but he will buy only 1/10 of the car so his total cost will be half of what he is paying today. A $200,000 car all of a sudden becomes affordable.

This is the magic of sharing, of course, a way to increase the usage of resources where everybody benefits for the increased system-wide efficiency. This is what the digital transformation and platforms (IoT, AI, and the like) make possible.

As platforms support the transformation of products into services they are increasing the efficiency of the system, enabling different business models and supporting the shift to the outcome economy where people pay based on the revenues generated (or increased value perceived). Probably a funny way but nevertheless an intriguing way to look at things.

Acknowledgement
The editing and polishing of the text of this paper by Theresa Cavrak is gratefully acknowledged.