Global Mobile Radar

January 2018

About the GSMA

The GSMA represents the interests of mobile operators worldwide, uniting nearly 800 operators with more than 300 companies in the broader mobile ecosystem, including handset and device makers, software companies, equipment providers and internet companies, as well as organisations in adjacent industry sectors. The GSMA also produces industry-leading events such as Mobile World Congress, Mobile World Congress Shanghai, Mobile World Congress Americas and the Mobile 360 Series of conferences.

The Global Mobile Radar series focuses on potential drivers of innovation and disruption across the digital economy. These reports highlight potential scenarios and examine the implications of these disruptions for a range of industry players, including the mobile operators. The reports are intended to be the basis for discussion and do not represent official GSMA positions on these future developments.

Contents

1

New disruptions on the Global Mobile Radar

This edition of the Global Mobile Radar comes as we look forward to Mobile World Congress 2018. Industry leaders from across the mobile ecosystem will gather in Barcelona to network, showcase and exchange ideas for the successful digital transformation of the industry.

As ever, Mobile World Congress looks set to provide a whole host of examples of the rapid rate of innovation in our industry and how it is becoming increasingly dynamic and disruptive. Identifying and acting on new sources of growth and opportunity, as well as the challenges ahead, remains top of the agenda for all in our ecosystem.

The Global Mobile Radar continues to provide a guiding light on the changing mix of disruptions and innovations to help you make forward-looking strategic decisions. Certainly many of the key themes running across this edition of the Global Mobile Radar will be a source of discussion and debate at Mobile World Congress. We look at the following:

- How **blockchain** has some growing up to do. We argue that the key challenge over the next 18 months is to see the development of more practical use cases and the growth of scale developer communities.
- How the growth in adoption of **edge computing** will play out. We examine how this new wave brings the two major infrastructure-based industries of the ICT world – cloud computing and telecoms – into competition for enterprise customers seeking to embrace IoT.
- How renewed attention on **augmented reality** is bringing it into the mainstream for consumers. We outline the current landscape and use cases, and examine the three key building blocks being assembled that should enable the full potential of AR to be realised.
- How **HD mapping** is proving a critical point of competition in the ecosystem. We argue that the race to win in this area is still open. Through a series of charts we also outline some of the major players, start-ups and venture-capital activity.

I hope the topics featured in this edition of the Global Mobile Radar help inspire fresh perspectives within your organisations and help you gain a view of the road ahead for our industry.

Laxmi Akkaraju

Chief Strategy Officer
GSMA

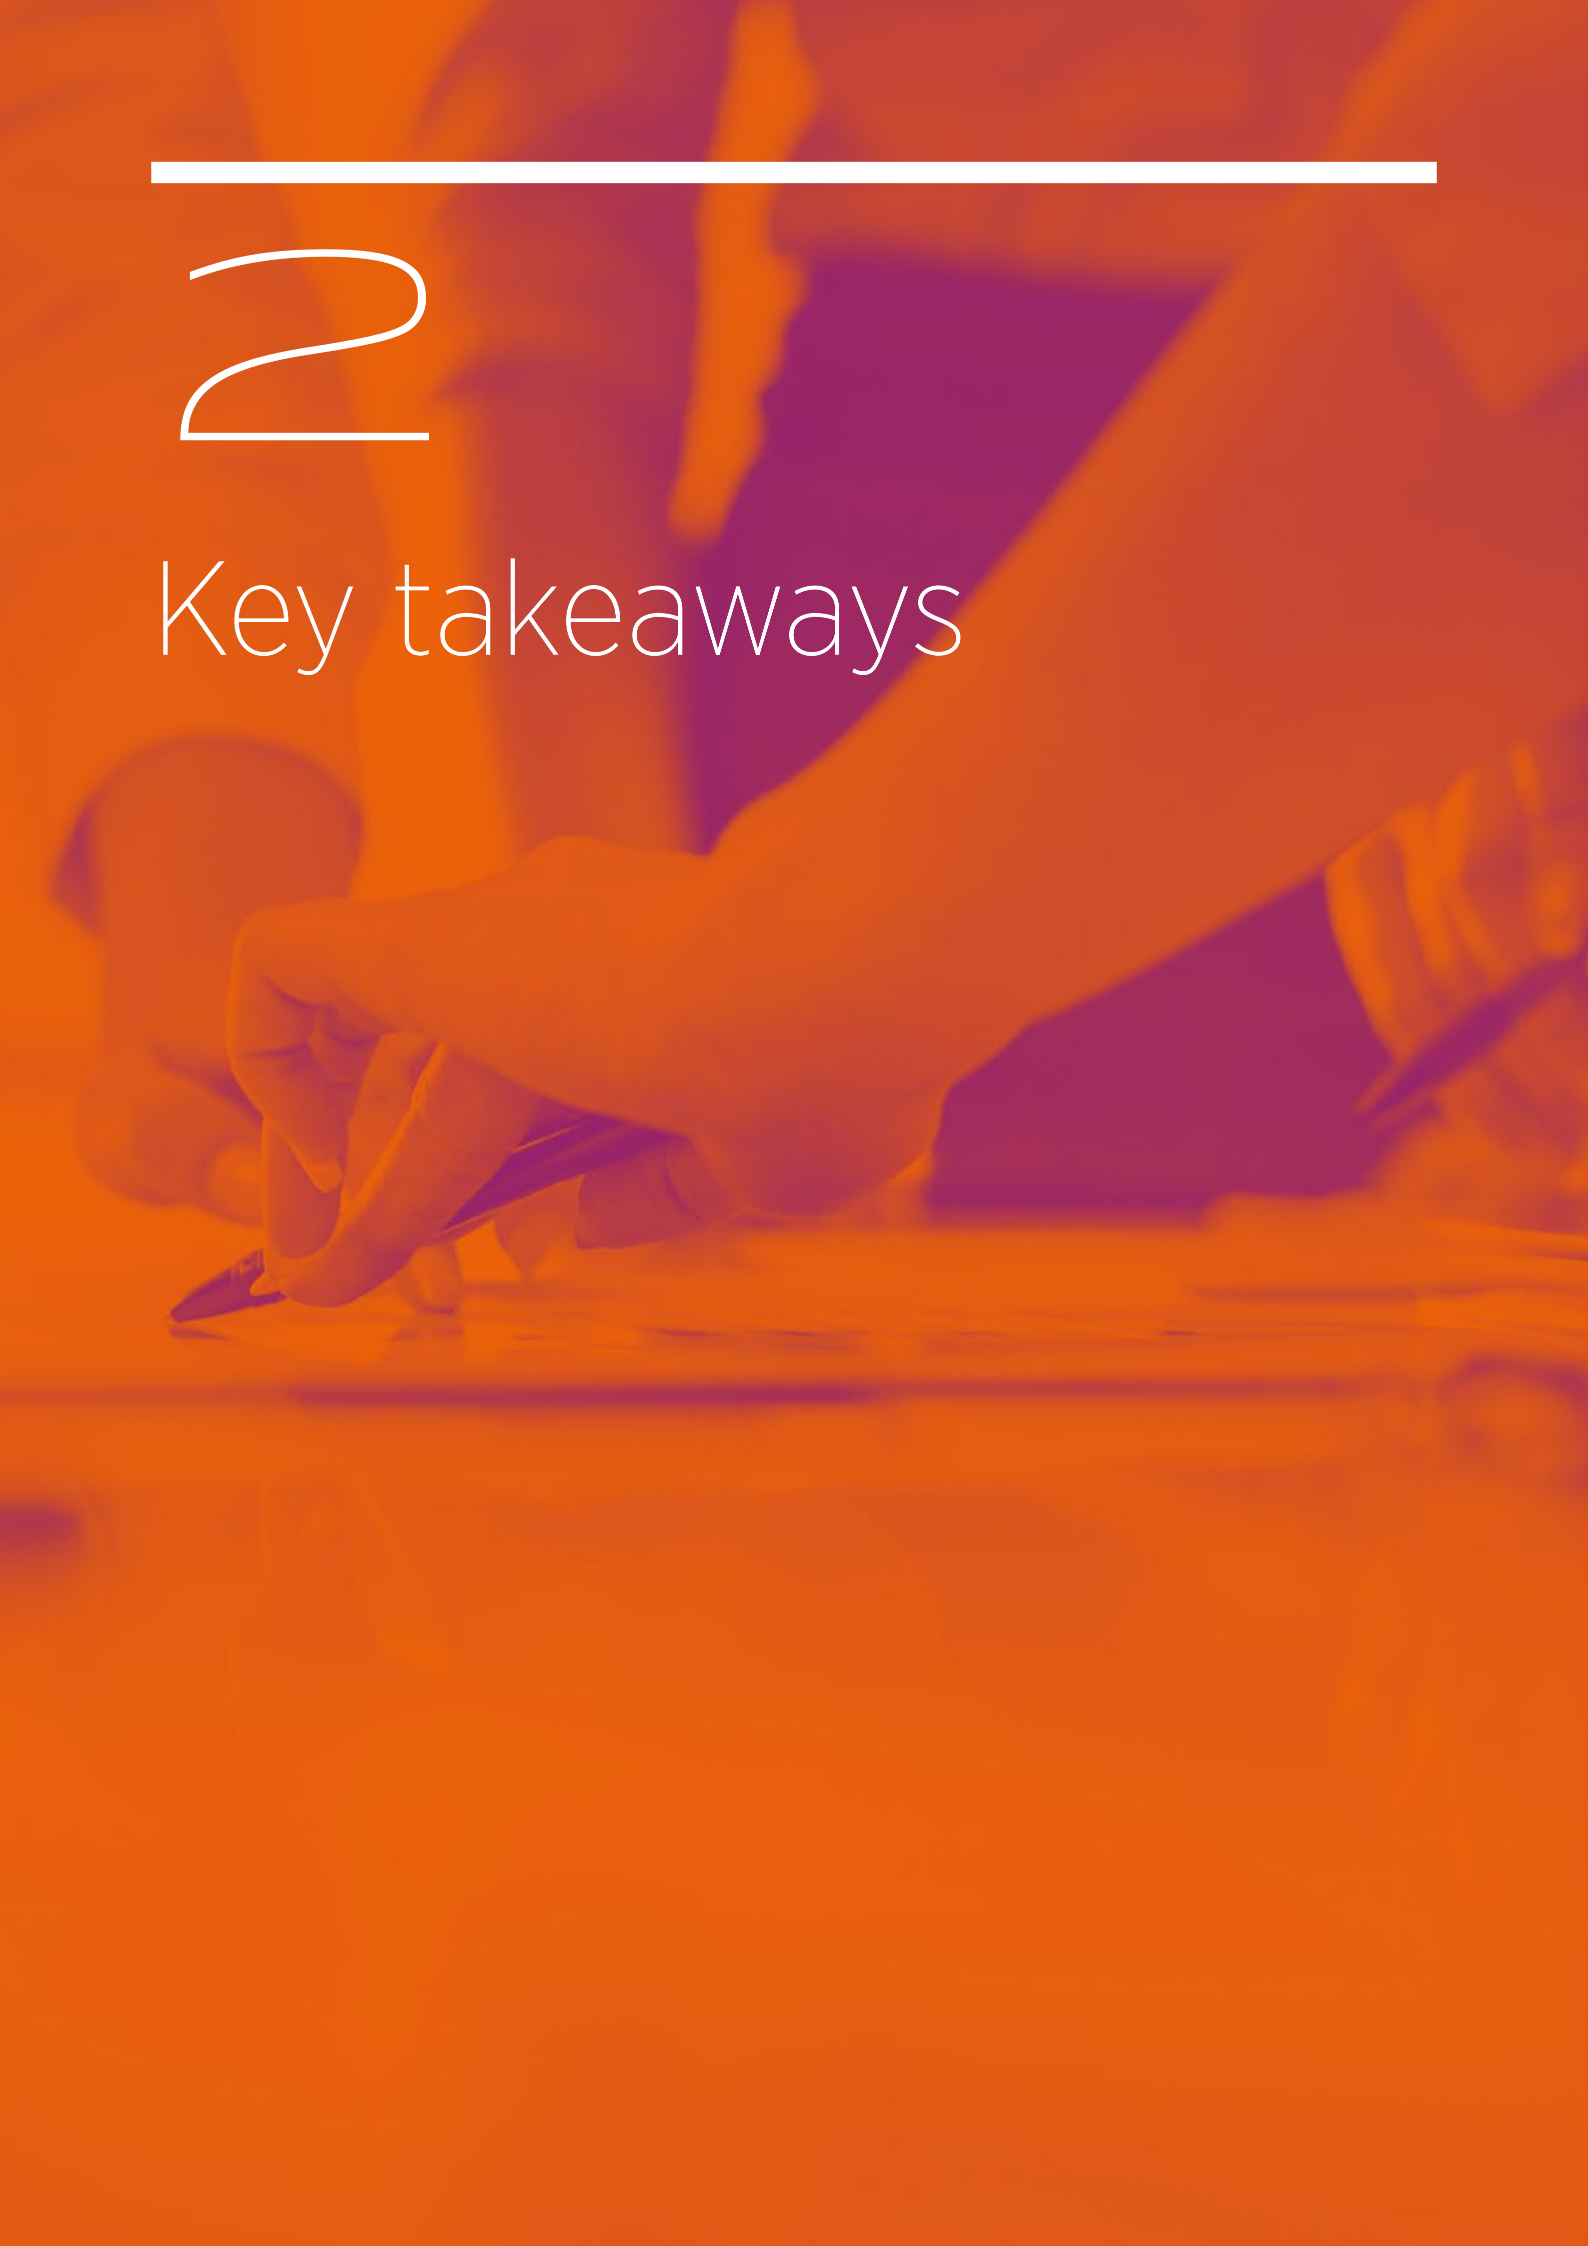

2

Key takeaways

Blockchain: growing up is hard to do

- The digital economy is generating exponential growth in the number of person-to-person transactions. Meanwhile, growth in the Internet of Things will lead to billions of devices interacting, transacting and sharing information. These developments are raising challenges around trust, identity and authorisation that existing software and organisational structures are struggling to address. Blockchain offers the promise of addressing these and similar challenges, and a future where almost every task and payment has a digital record and signature that can be validated, stored and shared.
- The initial focus of blockchain was on financial payments and transactions, with much debate around the role of cryptocurrencies and the growing number of initial coin offerings. Attention is now shifting to broader uses that focus on the principles of trust and ownership. These include identity management and executing smart contracts, with IoT a potential key focus.
- Several different types of blockchain (private versus public) and different underlying protocols exist, with different characteristics and use cases. As challenges around scaling and transaction processing time are addressed, enterprise adoption is likely to shift from permissioned to permissionless blockchains, driven by the rate of innovation in the latter. Indeed, the distinction may increasingly blur as blockchain-as-a-service providers look to offer the best attributes of both worlds.
- The key challenge over the next 18 months is to see both the development of more practical use cases and the growth of scale developer communities. Ethereum appears to be winning on the latter front, offering far broader and more flexible use cases than bitcoin. It is time for corporates to engage with the technology and explore trials and potential use cases.

Edge computing: when telco met cloud

- Edge computing describes the shifting of processing and networking power closer to the end user – whether a consumer smartphone connection, enterprise hosting environment, government database or inanimate thing.
- The model is gaining traction, driven by the need to provide connectivity for low-latency applications in IoT; immersive content consumption (augmented, virtual and mixed reality); and use of intelligent analytics.
- It brings the two major infrastructure-based industries of the ICT world – cloud computing and telecoms – into competition for enterprise customers seeking to embrace IoT.
- Early advantage lies with Amazon and Microsoft. Each has highly scaled in-built cloud infrastructure that can be parlayed to provide edge services for existing enterprise customers; high profitability; and significant cash to invest in incremental capacity to cement first-mover advantage.
- Operators and other telcos come at edge computing from a different angle, with cost efficiencies and latency reductions the principal advantages in the immediate term. In theory, new ways of selling network access such as network slicing open up but these remain untested. In any case, caution must be applied where competition overlaps exist against cloud incumbents.
- The left-field option is reimagining the network as a platform to spawn an ecosystem of developers making services based on the capabilities enabled by ultra-low latency. Common standards have been promulgated by ETSI, and the first release of MEC APIs occurred in July 2017 – a promising development to realise true global scale. On business model, open must really be open.
- Our expectation is the growth in the adoption of edge will play out in phases as economics improve, upgrades are made to increase efficiency (such as nano-processing) and acceptance grows. The ultimate end point is where the distinction between a centralised cloud and edge computing blurs or even disappears, although we would not expect this before 2030.

Realising the full potential of AR

- The recent announcements from both Apple and Google concerning new developer kits have refocused attention on the opportunities in augmented reality (AR), opening the door to a raft of new AR apps that run on existing high-end smartphones. This will bring AR increasingly into the consumer mainstream.
- However, the smartphone's 'magic window' is clearly a sub-optimal form factor if the true potential of AR is to be realised and mass-market adoption reached. An alternative form factor will likely revolve around some form of headset or glasses. There have been significant advances in both hardware and software since the launch of Google Glass, including improvements in processing power, miniaturisation and the supporting artificial intelligence needed to power a true AR experience. The longer term solution may be to offload some of the heavy processing load to edge computing assets, allowing a standalone wearable device but a vast array of supporting infrastructure (both hardware and software).
- The building blocks for a genuine AR consumer experience are now being assembled. There are three key components to realising the full potential of AR: the right hardware form factor with a natural user interface; a pervasive 3D digital map of the world; and access to an advanced intelligence and huge variety of data to add the 'right' information onto this digital canvas.
- In the long term, the AR platform will be transformative for businesses and consumers alike. As AR and VR increasingly merge, the new technology offers seamless integration with people's daily lives and a fusing of the physical and digital worlds. Significant improvements in artificial intelligence will deliver hyper-aware applications that are able to provide timely and contextually relevant information and guidance.

3

Blockchain: growing up is hard to do

3.1 Executive summary

Blockchain is an innovative approach to governance for networks and machines. Blockchains are a new form of database, which by design are distributed and decentralised. The openness of blockchain and the ability of all participants to view and validate transactions on the chain allow blockchain to act as a verifiable record of truth. This has led to it being referred to as a 'trust engine'.

The digital economy is generating exponential growth in the number of person-to-person (P2P) transactions, raising challenges around trust and identity that existing software and organisational structures are struggling to address. Similarly, the Internet of Things (IoT) will lead to billions of devices interacting, transacting and sharing information – a major challenge in areas such as authorisation for existing systems. A number of industries such as energy are now moving to decentralised models, and over time many more industries will embed networks and communications functions deeply into their structures. Blockchain offers the promise of addressing these and similar challenges, pointing to a future where almost every task and payment has a digital record and signature that could be validated, stored and shared.

The initial focus of blockchain was on financial payments and transactions, with much debate around the role of cryptocurrencies and the growing number of initial coin offerings (ICOs). Attention is now shifting to broader uses that focus on the principles of trust and ownership. Use cases include identity management and executing smart contracts, with IoT a potential key focus. In telecoms, a number of specific use cases are already being investigated, from data integrity and fraud prevention, to roaming and billing solutions.

Several different types of blockchain (private versus public) and different underlying protocols (bitcoin, Ethereum, Ripple) exist, with different characteristics and use cases. Bitcoin and public blockchains are both open source and democratic, with no central authority – features that are often championed by the developer community. But these strengths can also become weaknesses, with challenges around scale and decision making. As these challenges are addressed, enterprise adoption is likely to shift from permissioned to permissionless blockchains, driven by the rate of innovation in the latter. Indeed, the distinction may increasingly blur as blockchain-as-a-service providers look to offer the best attributes of both worlds.

Many technical and wider organisational and societal issues need to be addressed before blockchain can be widely accepted. While blockchain is an accurate record of itself, there is the classic "garbage-in" risk if the data it is recording is not accurate, with the danger of creating a permanent record of faulty data. In addition, the irreversibility of transactions and records could make it difficult for users to agree compromises in the event of a dispute.

The current focus around cryptocurrency valuations is, though, a distraction from the real potential utility of blockchain technology. The key challenge over the next 18 months is to see both the development of more practical use cases and the growth of scale developer communities. Ethereum appears to be winning on the latter front, offering far broader and more flexible use cases than Bitcoin. It is certainly time for corporates to engage with the technology and explore trials and potential use cases.

3.2 Context: ledgers, trust and the digital economy

Databases are not new: software versions have been around for decades, and the original versions can be traced back to wax tablets. A core feature of the existing database model is one of centralised trust, with siloed information pools managed by a number of 'trusted' providers (banks, governments or companies). As the digital economy has grown to shape and influence many features of both personal and corporate life, and the ledgers themselves have become digitised, this centralised model has remained the underlying paradigm of trust.

The key question is whether the centralised ledger system and its underling trust paradigm are truly scalable in the digital age. The digital economy is generating exponential growth in many areas, such as with the number of P2P transactions, raising the challenge of how to maintain the trust and identity that existing software and organisational challenges are struggling to address. Similarly, IoT will lead to billions of devices interacting, transacting and sharing information. These communications and transactions need to be secured and managed in a cost-efficient manner.

Trust relies on each individual's reputation and digital identity, such as seller and buyer feedback and ratings on sites such as eBay and Airbnb. One suggestion is that elements of this identity and profile score could be stored and managed in a blockchain (a shared, trusted public ledger). For individuals, these elements may include a variety of factors including financial or professional histories, medical information and consumer preferences. A company could also maintain its identity and reputational feedback in a manner that can establish its trustworthiness as a business partner or vendor.

The process of digitisation and decentralisation is also affecting a number of industries. A leading example is the power industry, where electricity generation is changing from centralised power stations and widespread hub & spoke distribution grids, to a distributed and localised generation model. New technologies such as solar and wind power generate the power, which is consumed (or stored) locally but increasingly also fed back into local networks. This requires new forms of (secure) network and communication to control.

In the future we could see almost every task and payment having a digital record and signature that could be validated, stored and shared. As well as reducing the friction in many transactions, the creation of a single trusted record means existing intermediaries such as lawyers, accountants and bankers might no longer be necessary. The scale deployment of blockchain technology could involve significant disruption for these industries and potential job losses. Individuals, organisations, machines and algorithms would freely transact and interact with one another with little friction or external supervision.

3.3 Defining blockchain

Blockchain can be broadly defined as a mechanism for reaching consensus around the state of a shared database between multiple third parties where trust is an issue.

Beyond this high-level definition, there are a number of important features that distinguish blockchain from other data products. One is that blockchain is a decentralised, distributed ledger; there is no single record keeper as with more centralised systems. This allows blockchain to address issues of trust among parties where trust is an issue, and can be used to protect identity where users wish to remain anonymous.

Source: GSMA Intelligence

1 Blockchain: decentralising transactions

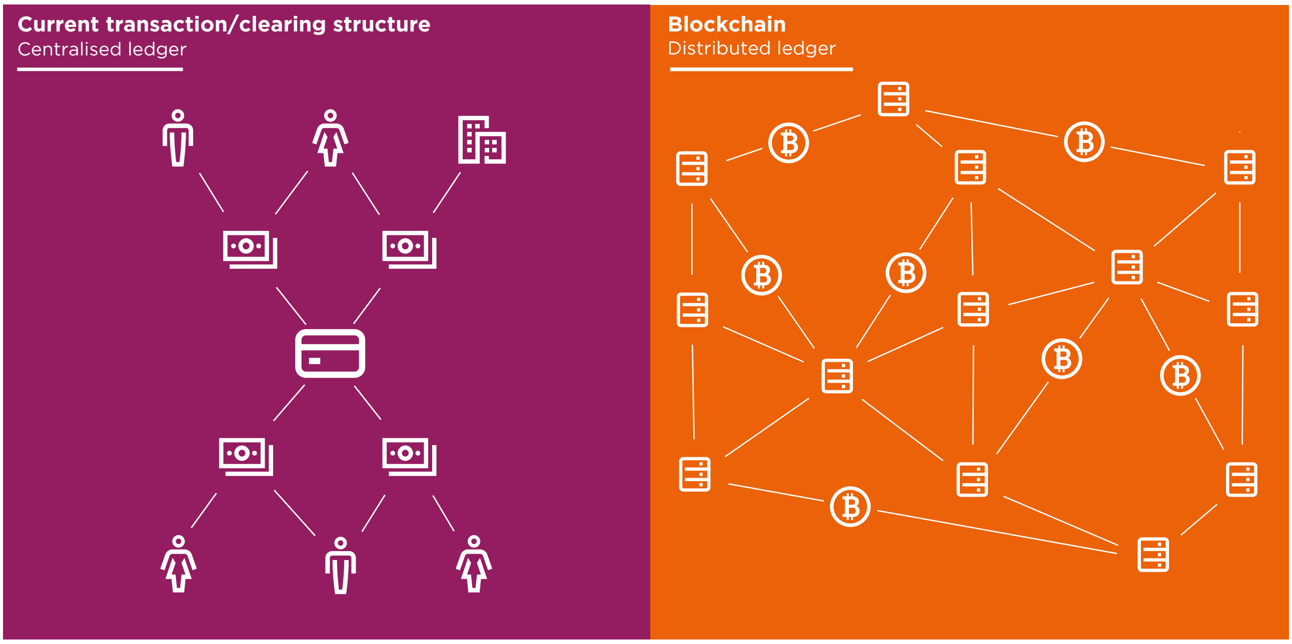

Blockchain is therefore both open and secure: in theory, every participant in the network can verify the correctness of transactions. Transactions are validated by the use of network consensus methods and cryptographic technology. The data cannot be changed or removed once it is in the blockchain – a key design characteristic that safeguards the data and creates a verifiable record of transactions. Trust is not established externally by a central authority or a third-party auditor, but continuously within the network.

3.4 Open source and technology adoption: developers strike back

Blockchain has been described by Harvard Business Review as a 'foundational' technology, with parallels to the growth of distributed networking technology (the development of TCP/IP protocols that laid the foundations for the creation and growth of the internet).[1] TCP/IP first appeared in 1972 and was used for email among a limited group of researchers. However, it has gone on to revolutionise communications, fundamentally disrupting the then prevailing circuit-switched standard for managing connectivity.

While TCP/IP represented a paradigm shift in networking technology, which significantly reduced the cost of connections, blockchain could offer a paradigm shift and significant cost reduction for transactions.

There are additional elements to the blockchain story that help explain its success to date and support among the developer community. One key element is its open source and decentralised nature, providing an opportunity to recast the digital economy in a more democratic manner. Collective action is now moving well beyond simple code sharing to create new capabilities and mechanisms for managing the digital economy, with the goal of offsetting the power of the dominant digital hubs. The early stages of the internet saw many open source solutions, but over time the internet has become dominated by a limited number of tech players who in many respects now control the internet. In theory, the lack of any central authority makes it impossible for any one entity to control blockchain (the distinction between permissioned and permissionless blockchains is discussed subsequently).

At the same time, the lingering aftermath of the financial crisis has further undermined faith in both central institutions more generally and national currencies more specifically. The rise of bitcoin and a whole range of other cryptocurrencies should in part at least be seen as driven by those seeking an alternative to traditional financial institutions and to escape the scrutiny of existing governments and regulators. When discussing the outlook for blockchain, it is important to remember the political and indeed quasi-philosophical drivers of some of the more vocal proponents of the technology.

1 "The Truth About Blockchain", Harvard Business Review, January-February 2017 Issue

3.5 The role of coins

A lot of the press coverage and controversy surrounding blockchain concerns the specific role of tokens or 'coins', with a raft of new coin offerings in recent months (so-called initial coin offerings or ICOs). A key distinction is that these coins or cryptocurrencies are in fact applications that sit on top of blockchain, not the other way around. ICOs have undoubtedly been successful, with total ICO fund raising in the first nine months of 2017 running in excess of $1.7 billion).[2] However, broader discussions around the issue of ICOs, and for example their potential to disrupt other sources of financing such as venture capital, are beyond the scope of this analysis.

Source: Coinmarketcap.com

2 Growth of cryptocurrencies

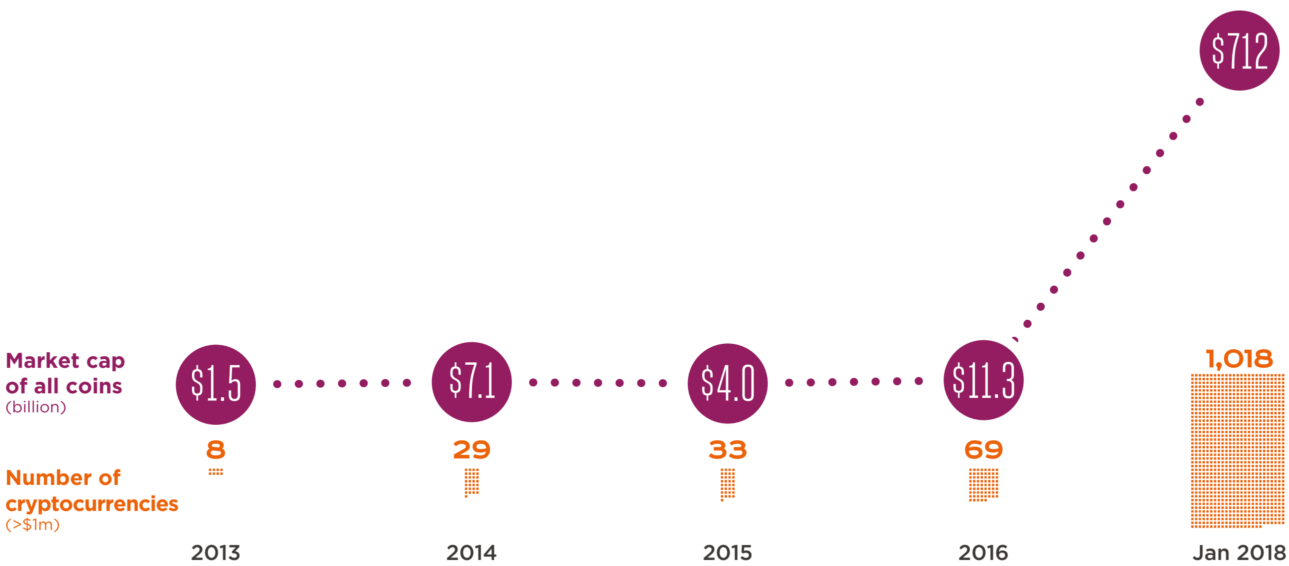

There is a misconception as to the role of these coins. When bitcoin was initially launched, it was widely seen as a breakthrough in financial technology and a new form of payment. However, bitcoin was essentially two innovations wrapped into one:

- a store of value for people looking for an alternative to the existing financial system
- a new way to develop and fund open networks (without recourse, for example, to venture-capital funding).

It is the former use case that has perhaps attracted some of the greatest criticism, with for example JP Morgan Chase CEO Jamie Dimon saying of bitcoin: *"It's worse than tulip bulbs. It won't end well. Someone is going to get killed."*

The real value of tokens lies in the second innovation. Open systems have never had a way to fund the development of a network, with the result that much of the internet today is effectively privately owned (think of the so-called 'FANG' companies: Facebook, Amazon, Netflix and Google). Tokens provide a way

2 "ICO funding hit a record $800 million in Q2 2017", TechCrunch, September 2017

not only to define a protocol but to fund the operating expenses required to host it as a service. Bitcoin and Ethereum are the two highest profile and most successful of the new tokens, with each relying on thousands of servers around the world (referred to as “miners”) that run their networks, updating the underlying blockchain.

This process of maintaining the integrity of the blockchain involves significant hosting costs, with token rewards distributed automatically to computers on the network that undertake this work (“mining rewards”). Tokens provide a model for creating shared computing resources (including databases, processing capacity and file storage) while keeping the control of those resources decentralised (and without requiring an organisation to maintain them).

However, it is also possible to have tokenless blockchains. Depending on the nature of the blockchain, there may be no need for a token reward system. In public or open blockchains (generally described as ‘permissionless’), there is a need for some sort of incentive scheme for block validators for contributing their computing power to update and maintain the integrity of the chain. However, in a more private (or ‘permissioned’) blockchain, such rewards may be unnecessary. One or several founding organisations may undertake the task as they rely on the integrity of the blockchain for their own business purposes, or they may create contractual arrangements with third parties to do so.

Much of the recent press reporting of blockchain has focused on cryptocurrency valuations and the likelihood of a speculative bubble. However, blockchain as a store of value is only one and perhaps the least interesting potential use case of blockchain. It is likely that as new use cases emerge and applications are developed, attention will focus more on the role of tokens to develop and fund these services and less on their speculative nature.

3.6 Types of blockchain

Public blockchains, including bitcoin, Ethereum and Hyperledger, are built to be accessible by anyone with adequate technology, which has so far meant a computer with sufficient computing power and access to the internet. However, from an enterprise perspective, most firms – including those in the financial services industry – are looking at options around private and 'permissioned' blockchains. Table 1 highlights some of the key differences between the different types of blockchain.

Source: coindesk.com, GSMA Intelligence

1 Types of blockchain and their characteristics

	Public (permissionless)	Private (permissioned)	Hybrid/consortium
Access	Open read/write access to database	Permissioned access to database	Contains elements of both public and private depending on use case
Speed	Slower	Faster	
Security	Proof-of-work/ state	Pre-approved participants	
Identity	Anonymous	Known identities	
Asset	Native assets	Any asset	
Network	Public internet	Private networks	

In a fully private blockchain, permissions to write to the blockchain are restricted to one or a limited number of organisations, in a similar manner to many existing database products. Other organisations may be able to view the data but not alter it. The ability to assign different levels of permission to various network participants is particularly suited to use in commercial contexts, such as financial services and beyond, where certain actions and information are not intended to be public. This means participants retain the benefit of a shared infrastructure while maintaining a level of security and privacy.

A third category, variously referred to as hybrid or consortium blockchains, shares characteristics of the first two types and is effectively partially decentralised. The consensus process can then be shared across a number of organisations, each of which could operate a node on the network. A minimum number of participating organisations could be required to sign each 'block' of data and so establish the consensus. The right to view the data could be open to the public or a larger number of organisations, or the use of some form of access point (such as via an API) that could allow third parties to query parts of the chain and receive cryptographic proofs of some parts of the blockchain state.

Source: GSMA Intelligence

2 Public blockchains

	Hyperledger	Ethereum	Ripple	Bitcoin
Description	General purpose blockchain	General purpose blockchain	Payments blockchain	Payments blockchain
Governance	Linux Foundation	Ethereum developers	Ripple Labs	Bitcoin developers
Currency	None	Ether	XRP	BTC
Mining reward	NA	Yes	No	Yes
Consensus network	Pluggable PBFT	Mining	Ripple Protocol	Mining
Network	Private or public	Private or public	Public	Public
Privacy	Open to private	Open	Open	Open
Smart contracts	Multiple programming language	'Solidity' programming language	None	Possible but not obvious

3.7 Potential use cases

The most advanced use cases for blockchain today centre on its use for financial transactions and the management of financial assets. A number of leading banks have already taken steps to use blockchain as part of their existing operations, while a growing number of start-ups are targeting use cases such as cross-border payments or low liquidity foreign exchange markets.

Source: GSMA Intelligence

3 Blockchain use cases

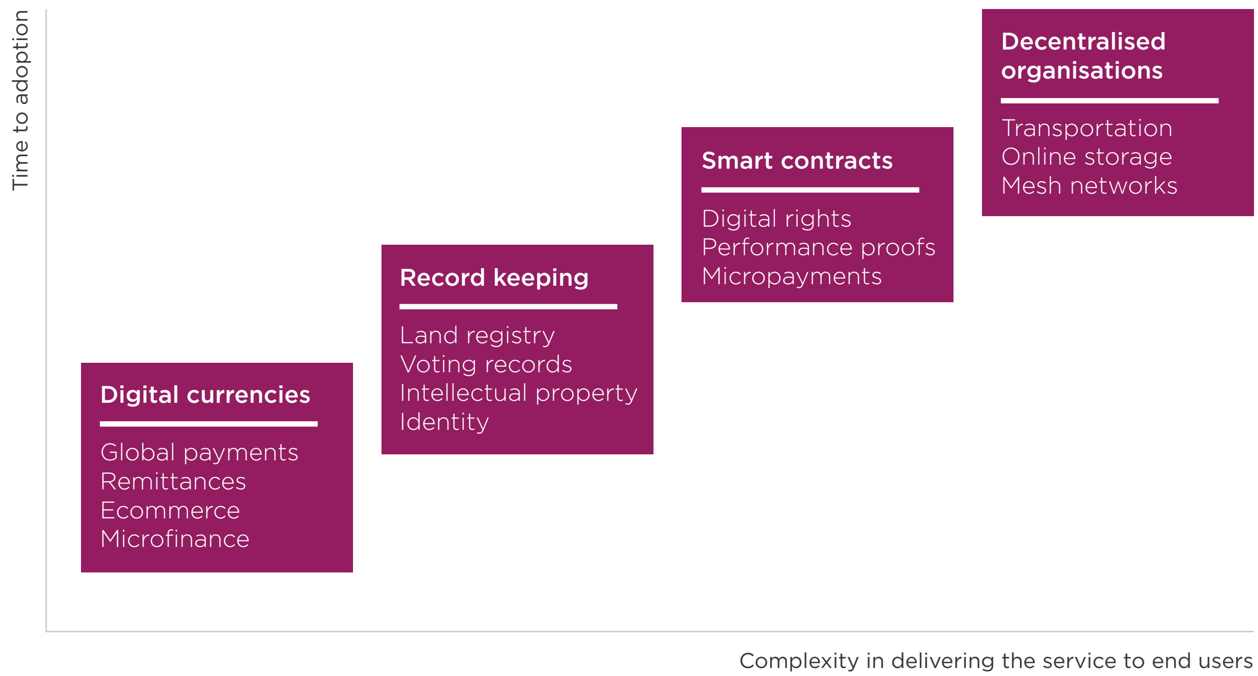

Beyond the exchange and management of financial assets, potential use cases in the broader economy centre on storing digital records and executing smart contracts:

- **Digital records:** blockchain can establish secure and unchangeable digital records, covering financial transactions but also the digital records of physical assets. This could include anything from payments to property transfers or voting records. Use cases could then extend to keep a digital record of an individual, with the ability to control which parts of that record are revealed and to whom.
- **Smart contracts:** blockchain can be used to establish the terms and conditions for a contract, and then the shared network resources automatically execute the contract when the terms have been met. The blockchain can also then monitor compliance with the contract terms, such as any payment required (which could also be automatically implemented through the blockchain). Such smart contracts would significantly reduce the friction involved in any transaction, including both the cost and speed to execute.

Governments around the world are looking at the potential of blockchain in areas such as managing identity, voting records and recording assets such as land registry. These are all areas that relate closely to the key principles of ownership and trust. Estonia is one of the leading examples of a digital government that is using blockchain to deliver real-world solutions.

Estonia and blockchain

The Estonian government has developed ‘e-Estonia’ into one of the most advanced digital societies in the world, with a vision of automatic e-services available 24 hours a day.

The Estonian government has been testing blockchain since 2008, and since 2012 it has been in operational use in registries across a number of areas, including health, judicial, legislative, security and commercial code systems. There are also plans to extend its use to other spheres such as personal medicine, cyber security and ‘data embassies’.

Estonia uses Keyless Signature Infrastructure (KSI) blockchain technology, which is also now being trialled in other European countries and the US, focusing on smart cities, healthcare and digital government.

Blockchain and identity management

Blockchain could be used to create an identity management system, which would benefit individuals that need to prove their identity, and organisations that have a requirement to establish a customer's identity (for example, those subject to KYC regulations).

The user's identity starts its journey into the distributed ledger as a self-asserted block/record containing the user's identity claims (hashed) and the user's public key, all signed with the user's private key, providing indelible proof of the claim's existence. The actual identity information is encrypted and stored separately in a container, either on the user's device or in external secure systems/cloud. An alternative approach would be for the initial user record to be established by an issuer that can provide verifiable claims based on a previous KYC process (such as a bank or mobile operator).

Other entities, such as a bank or electricity provider, with which the user has a relationship are also represented within the ledger with their own sets of hashed attributes and public keys. These entities can establish relationships with the user by signing the particular hashed attributes of the user that are relevant to that relationship or contributing claims that can be added to the user's container.

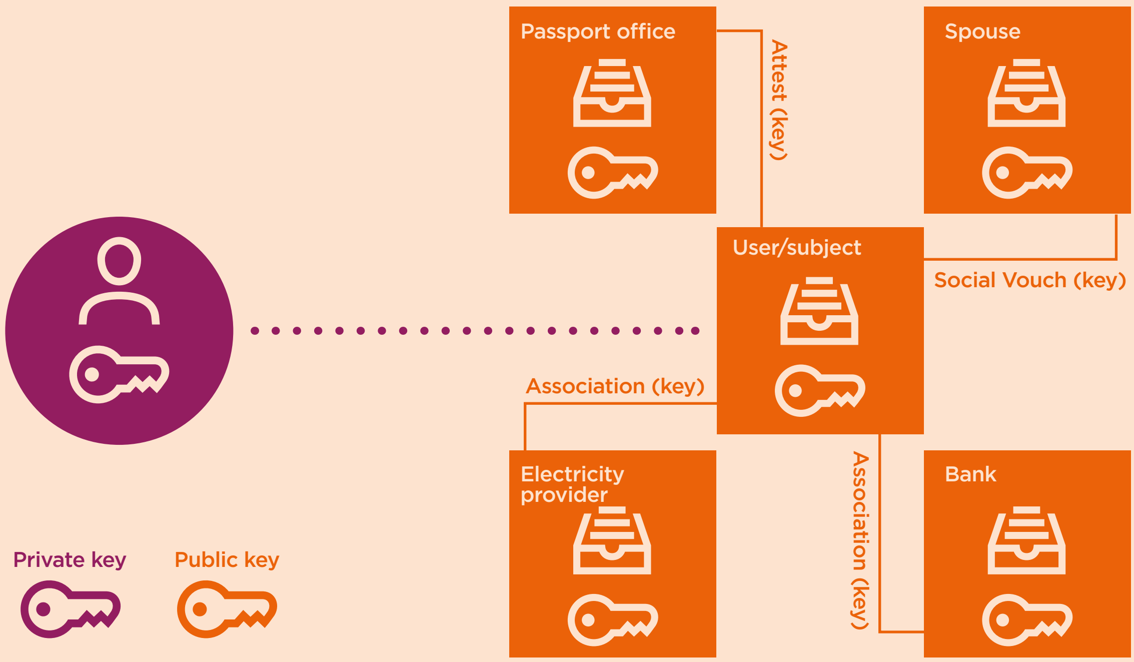

Each organisation can decide whether to trust credentials in the container based on which organisation verified or attested to them. As more and more relationships are established for the user within the ledger, confidence in the accuracy of the attributes – and hence the identity itself – grows organically. In addition, as more transactions take place involving the user (with other users or entities verifying or trusting the hashed attributes of the user), the 'reputation capital' of the identity also grows.

The block/record representing a digital identity in the ledger is identified using the public key associated with the user, and the corresponding private key is the credential that the user needs to keep protected. In a sense, therefore, the public key can be considered equivalent to a user ID and the private key equivalent to a password or biometric.

Initiatives such as Sovrin are tackling many of these identity issues with the aim of setting up a public, but permissioned, distributed ledger as the basis for an identity trust framework. Other companies are building solutions targeted at specific identity use cases, such as airport travel (ShoCard) and e-government (Procivis). A new consortium comprising Cisco, Bosch and others has started work on using blockchain technologies to address the identity management issues associated with IoT.

In telecoms, there are a number of potential touch points beyond identity. Ongoing discussions have included the use of private or permissioned blockchains to address issues such as improving processes and security, either within operators or between a group of operators and/or with suppliers. More general use cases include the following:

- **IoT:** the use of encrypted digital identity and smart contracts could have a range of applications to help connected objects and machines communicate to each other and authorise processes and transactions. The use of localised mesh networks and smart contracts could allow devices and machines to operate in areas of limited connectivity such as remote locations.
- **Fraud prevention:** this remains a significant challenge for operators, whether in areas such as roaming or subscriber identity fraud. In the first case, a smart contract could establish a data transfer and payment between the host and originating networks. In the latter case, a blockchain-based eSIM solution could use both public and private keys. The former could be used to identify a device on the network and the latter used by the individual subscriber to access specific services.
- **Micropayments:** these have been suggested as a solution in a number of situations, including as an alternative to advertising for funding internet or other forms of content; or paying for connectivity.
- **Business process efficiency:** areas include number portability and eSIM provisioning.
- **Roaming/interconnect management:** a permissioned blockchain with nodes for each operator in a bilateral roaming arrangement could simplify processes related to subscriber authentication and the settlement of roaming fees (via smart contracts).

Operator blockchain initiatives

SoftBank, Sprint and Far EasTone have created a new consortium aimed at helping operators adopt blockchain technology. The consortium is using a platform developed by TBCASoft – a startup developing blockchain services for the telecoms industry. The group aims to attract other operators to join the consortium and is exploring a range of use cases. These include examining how blockchain can reduce friction in the flow payments between various players in the telecoms ecosystem, in terms of operators paying each other (such as roaming), authenticating users and using cryptocurrencies to reload prepaid plans.

Blockchain challenges

It's also important to clarify in all these cases what blockchain is not – namely, an instant panacea for any of the existing problems that corporates and individuals face in the increasingly complex and interconnected digital economy. Issues around security and the need to integrate with existing systems are seen as major obstacles to IoT uptake, and challenges that blockchain could potentially address. However, to do so, it would need to be adopted across all ecosystem players, with the danger of replicating the challenge of platform fragmentation that currently bedevils IoT (unless the ecosystem adopts a public rather than private blockchain model).

Despite the hype, there remain a number of significant challenges before blockchain can enter the mainstream:

- **Energy consumption**: the distributed ledger structure has advantages but also some significant costs in terms of energy consumption. The larger a blockchain becomes, the more computing power it requires, with suggestions that the computing power required to process bitcoin already exceeds that of the world's fastest 500 supercomputers combined.
- **Processing speeds:** on bitcoin, each new 'block' of transactions that is added to the blockchain is only 1 MB in size. Bitcoin can process a maximum of seven transactions per second at most, leading to significant delays when it comes to updating each new transaction. Ethereum uses a different validation model, as well as a shorter block time compared to bitcoin, which means that transactions can be confirmed in less than 20 seconds.
- **Privacy:** any public blockchain is by its nature open to anyone, with everyone on the network able to view every transaction that has been recorded. In addition, though one of the advantages of blockchain is that it is immutable, this could become a problem if incorrect or even illegal data is recorded.
- **Interoperability:** challenges around avoiding competing standards and incompatible platforms are already significant issues in the digital economy. To date, the fragmented landscape of competing blockchains has failed to produce international standards for the technology.
- **Security:** anyone with the encryption key can read the encrypted data if the key is made public; but, at the same time, if the key to unlock the blockchain is lost it cannot be recovered. Security also remains vulnerable to technological advance or deliberate backdoors.
- **Limited ability to store data:** although blockchains are effective records of transactions, they have limited ability to store data. For example, bitcoin blockchain records the inputs and outputs of every coin on the network, as well as the content of an additional field that allows for up to just 40 bytes of metadata per transaction.
- **Interaction with external systems:** blockchains are not designed to query websites or other databases. This presents a challenge in the case for example of a smart contract: how does it know that the terms of the contract have been met?

There are solutions underway to address some of these challenges, most of which are technical in nature. The Lightning Network offers one potential solution to the scalability issue, with the concept that transactions are undertaken in separate ledgers off the blockchain, but with the potential for these to be validated on the blockchain if the need arose.

Hashgraph is a relatively new distributed ledger that is significantly faster than existing blockchains, with consensus achieved through a 'gossip protocol'. Efforts to address the speed and capacity issues have also seen bitcoin and Ethereum fork (or split). These forks can be contentious and sometimes arise due to challenges in agreeing the way forward.

When it comes to smart contracts or the need to interact with external data sources, one solution could be to develop trusted sources that allow the introduction of external data into the blockchain. These have been referred to as 'oracles'[3] – services that live off the blockchain (although these could just as easily be referred to as 'middleware').

However, because the external sources can drive behaviour and provide input that in turn is used to calculate a price or fee, or determine if an obligation is met, strong trust in these sources needs to be established between counterparties. In order to obtain that trust, the source needs to be well known with a strong identity, provide proofs as to the integrity of the data it is providing and have well defined rules as to when the interaction can take place. In addition, there are a range of policy, organisational and social barriers that would need to be addressed for blockchain to see mainstream adoption. For example, policymakers and regulators will have a common interest in ensuring that the adoption of technological innovation is subject to a consistent set of protections across the industry.

Many of the challenges listed above will be addressed (and indeed are already being tackled). The use of off-chain processing of transactions is being promoted by both bitcoin and Ethereum as a way to increase the transaction capacity and processing speeds, avoiding the resource-intensive processing requirements of the proof-of-work consensus approach. Only the opening and closing transaction of these side-chains are then recorded in the main blockchain. However, such solutions lose one of the key principles of blockchain – namely, the creation of an immutable record of all transactions that can be viewed by the entire network.

3 "How Smart Contracts Work", spectrum.ieee.org, September 2017

3.8 Future outlook

The sharing and digital economies seem natural areas for blockchain-based peer-to-peer networks to attempt to disrupt incumbents. There is the potential for a 'blockchain Uber' or 'blockchain Amazon', run on open source blockchain software and over decentralised networks where value could be more equally distributed across network participants. For example, there are already companies such as OpenBazaar (a decentralised marketplace similar to Amazon or eBay); while Akasha and Steem are distributed social networks (not unlike Facebook).

There are a growing number of blockchain platforms and many new services being trialled. However, no single platform has demonstrated the scale or network effects that are key success factors in the digital economy. In theory, achieving scale in the blockchain environment should be easier than in the existing internet economy. In the latter case, friction between proprietary APIs and platforms means that there is, for example, no common standard for a cloud storage application, but rather a number of proprietary offerings. The decentralised and open nature of blockchain and common standards reduce these frictions and act as a bulwark against the centralising tendencies of the internet economy. The dominance of many of the existing digital hubs comes in large part from the scale effect and positive feedback loops. It remains to be seen whether blockchain can truly disrupt these digital hubs, or whether some of its use cases and capabilities will simply be adopted by the existing dominant players to reduce transactional friction and perhaps drive even greater scale.

Many corporates are already struggling with how to adopt a range of promising technologies, including big data and analytics, the umbrella term of artificial intelligence and the potential of IoT. While blockchain may fit into some of these areas, particularly IoT, it is likely to take some time for the technology to reach a level of maturity to give companies the comfort to move from trial to real-world applications. The financial sector is already seeing the most blockchain activity, both from incumbents and potentially disruptive start-ups. Beyond this vertical, the most promising use cases are likely to be those that centre on the core principles of trust and ownership. These include government services, individual identity management and smart contracts. It is worth emphasising though that blockchain remains a relatively early-stage technology that has seen limited scale deployments to date.

Among the many questions that still need to be answered around the future development of blockchain, the following stand out:

- What is the 'killer' application or use case for blockchain? Beyond financial transactions, this remains unclear, as does the question of when (or if) blockchain will become a mainstream consumer play.
- Can blockchain solve the challenges of digital identity? Will a single digital identity evolve that can be used in multiple scenarios, or will the number of digital identities simply mirror the number of physical identities that each individual carries (from driver's licence to bank cards)?
- What legal changes will be needed to recognise smart contracts? While regulators are beginning to address the questions of cryptocurrencies (in some instances by simply banning them), smart contracts today have no legal status. Should existing laws be modified or new ones be created? And what limits (if any) should be applied to the fields in which smart contracts could be used?
- How will an open source technology move to common standards, and address the issue of interoperability between competing platforms and protocols? While the rapid pace of innovation and decentralised nature of blockchain argue against premature efforts to regulate and standardise, interoperability is key for blockchain and associated applications to scale.

However, as these scalability challenges are addressed, enterprise adoption is likely to shift from permissioned to permissionless blockchains. This will be driven by concentration of developer attention and consequent rate of innovation in the latter. Indeed, the distinction may increasingly blur as blockchain-as-a-service providers (such as IBM and Microsoft) look to offer the best attributes of both worlds.

Despite the uncertainties, corporates should certainly continue to explore the use of blockchain, whether individually or through industry consortia, with an initial focus on internal rather than client-facing processes. Sensible starting points would be non-core or non-critical operations, with the goal of expanding to more sensitive areas and then client-facing functions once confidence in the technology has been established.

4

Edge computing: when telco met cloud

4.1 Executive summary

Over the last 40 years, computing power and processing have alternated between cycles of centralised and decentralised architectures. The current model of cloud, built on data centres, is centralised but we believe a fourth wave – edge computing – is in gestation that would mark a shift back towards decentralisation. This incipient shift is being driven by the need to provide connectivity for low-latency applications in IoT, immersive content consumption (augmented, virtual and mixed reality), use of intelligent analytics, and the drive to reduce mobile network operating costs set against inexorable rises in data consumption.

A basic definition for edge computing is the move of processing and networking power closer to the end user – whether a consumer smartphone connection, enterprise hosting environment, government database or inanimate thing. The key factors determining whether something needs edge are latency (usually less than 10 milliseconds) and real-time analytics. High potential use cases include industrial IoT verticals (e.g. wind turbines, electricity grids, transportation, logistics, advanced robotics), AR and VR, safety-critical smart city applications and possibly autonomous vehicles.

From a technology point of view, edge computing is still in its early stages, with the majority of enterprise workloads handled by centralised data centres. The decentralised nature of edge clouds and overhead costs of operating a large number of micro data centres mean that the cost of running applications will be higher in the early adoption phase when there are fewer tenants to spread computing loads across. Redundancy will need to be provisioned in a different way to conventional cloud given the distance between data centres. Pricing structures are also fair game: one of the chief advantages of edge is the intelligent division of tasks handled by central cloud (heavy lifting processing) and edge cloud (latency-sensitive, geolocation). If edge processing can be 'chunked' into individual tasks, the next logical step is that customers should only pay for what they use as opposed to flat fees.

Our expectation is therefore that the growth in adoption of edge will play out in phases as economics improve, upgrades are made to increase efficiency (such as nano-processing) and acceptance grows.

This wave brings the two major infrastructure-based industries of the ICT world – cloud computing and telecoms – into competition for enterprise customers seeking to embrace IoT. From a revenue perspective, we forecast more than 50% of incremental revenue in IoT will come in enterprise and industrial verticals. Early advantage lies with Amazon and Microsoft (and to a lesser extent Google as the number three cloud company). Amazon and Microsoft launched commercial edge products in 2017, AWS Greengrass and Azure Edge. Each has highly scaled in-built cloud infrastructure that can be parlayed to provide edge services for existing enterprise customers; high profitability; and significant cash to invest in incremental capacity to cement first-mover advantage.

Operators come at edge computing from a different angle. One advantage is to increase network cost efficiency by reducing the volume of transport sent to central cloud servers. Edge is also an infrastructure prerequisite for many 5G use cases if sub-1 millisecond latency promises are to be met. Low-latency IoT is most interesting. If quality of service (QoS) could be guaranteed, which operators are very good at, the concept of network slicing could provide a monetisation model for edge IoT. The challenge is scale: Amazon and Microsoft have global scale and high pricing power. Operators have global scale in aggregate but on an individual company basis coverage footprints for edge computing will be limited in early-stage deployments while commercial returns are tested.

The left-field option is reimagining the network as a platform to spawn an open ecosystem of developers making services based on the capabilities enabled by ultra-low latency. Common standards have been promulgated by ETSI, and the first release of MEC APIs occurred in July 2017 – a promising development to realise true global scale. On business model, open must really be open: a set of common standards everyone is working from and openly welcoming in innovation from start-ups to established companies. Charging developers for API access is a non-starter.

Apple and Google also have stakes in edge, if indirectly. With the iPhone X and ARKit, Apple has signalled its intention to push AR as the next 'killer app'. Google has pushed VR more, although it will only be a matter of time before AR features on Android to the same extent. Both could use the same ecosystem pull model to entice developers and vertical sectors to develop AR services; if anything, this is going to be bigger than the app economy that helped fuel smartphone growth given that AR is a transformative technology in play for a plethora of industries from media to fashion to transportation.

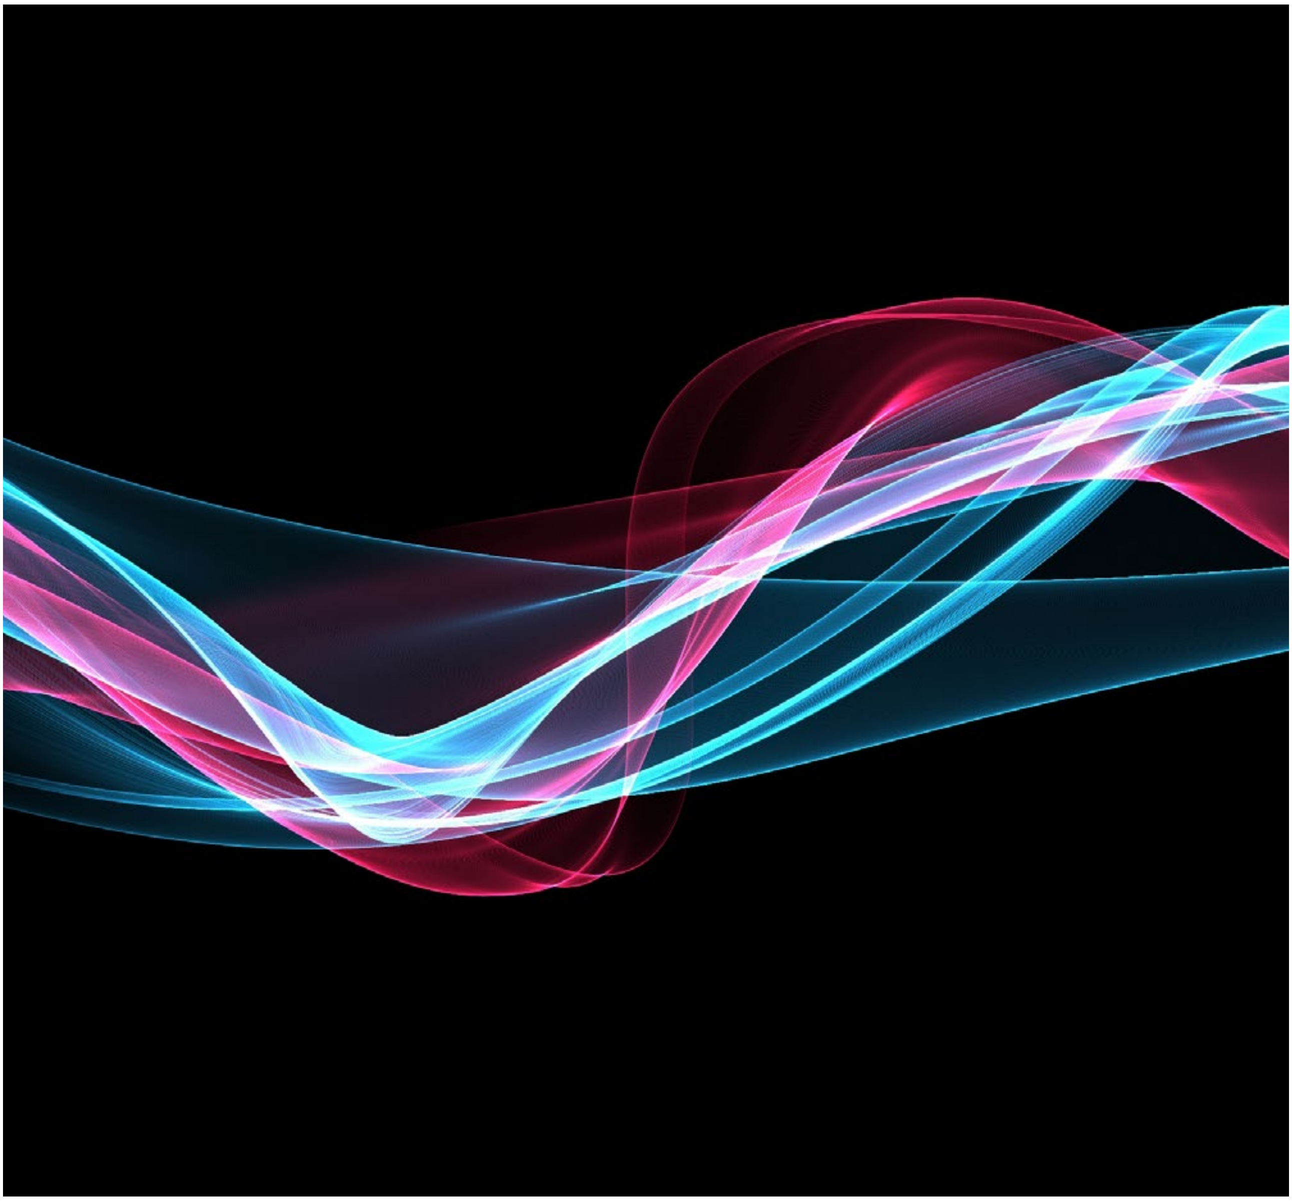

4.2 Moving (back) to a world of decentralised computing

Throughout the history of computing over the last 40 years, power and processing have alternated between cycles of centralised and decentralised architectures. Mainframes were the first real computing 'era' in the late 1960s and 1970s, operating under a highly centralised and controlled regime, mostly because they were the preserve of the military and others charged with national security responsibilities.

The PC era and the Wintel duopoly in the 1990s is best characterised through the client-server model, shifting the pendulum towards decentralisation. Mass-market adoption of the internet, enterprise IT and smartphones in the 2000s catalysed the wave of cloud computing anchored in massive server farms. While it is difficult to define hard boundaries, we believe a fourth wave – edge computing – is in gestation and would mark a shift back towards decentralisation.

Source: GSMA Intelligence

1 Edge computing marks a shift back towards decentralised computing

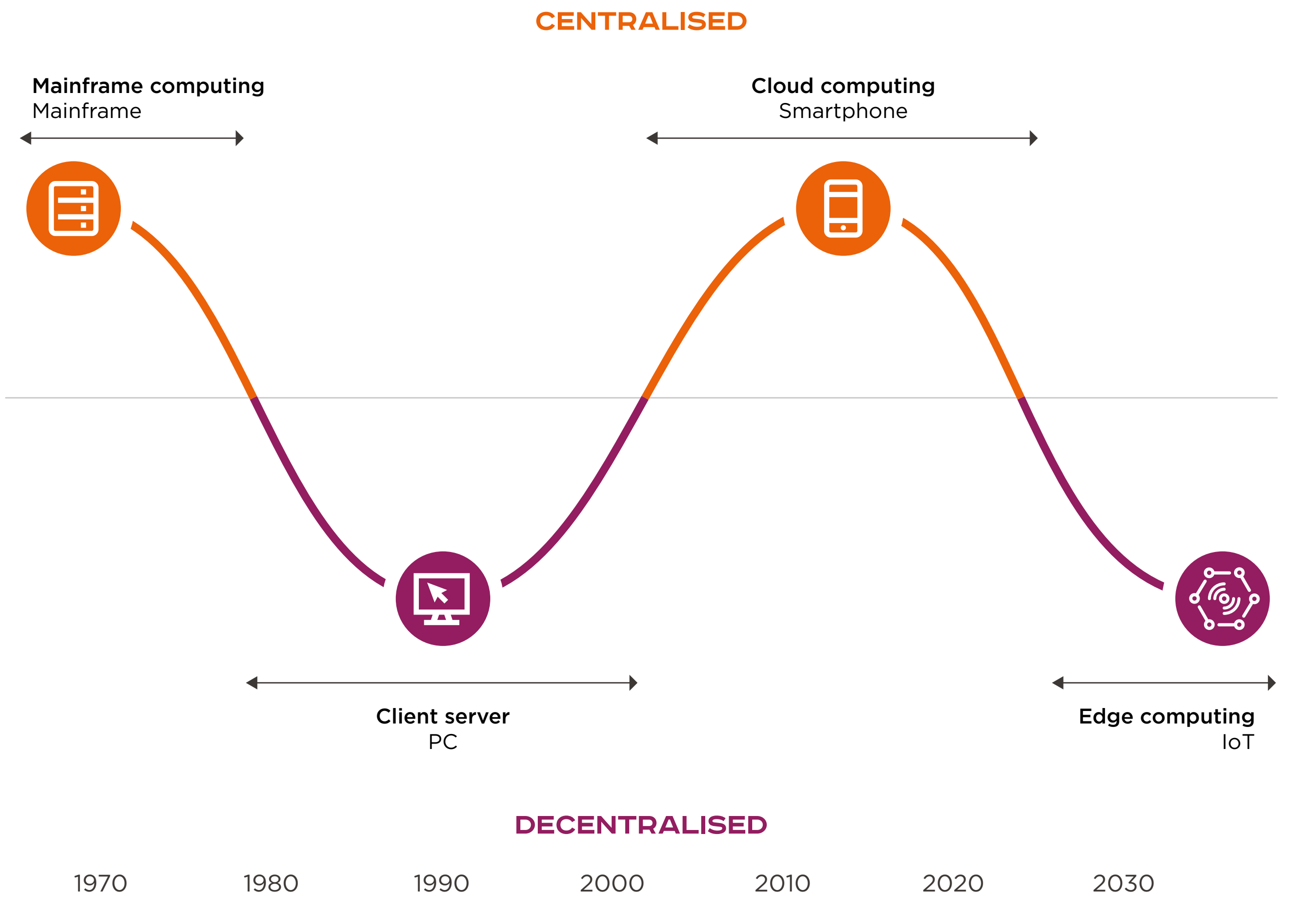

Note: Predominant device category shown within each era

A basic definition for edge computing is the move of processing and networking power closer to the end user – whether that's a consumer smartphone connection, enterprise hosting environment, government database or inanimate thing. This incipient shift is being driven by several factors: the need to provide connectivity for low-latency applications in IoT; the nascent but touted medium of immersive content consumption (augmented, virtual and mixed reality); the use of intelligent analytics; and the drive to reduce mobile network operating costs set against inexorable rises in data consumption.

Edge computing is effectively an umbrella term for the coming together of the two major infrastructure-based industries of the ICT world: cloud computing and telecoms (see Figure 2). Cloud computing is built on a vertical stack of computing capacity mostly sold to enterprises and government, with each ascending level offering an increased degree of customisation. Infrastructure, platform and software have all come to be offered 'as a service', with an increasing level of consulting support from cloud providers keen to help clients navigate the challenge of digitisation in whatever shape it takes. Telecoms is built on mobile and fixed line connectivity offered to consumer and enterprise segments. Until now, these industries have largely competed in parallel as opposed to against each other. Edge changes that because it represents uncolonised ground of mutual interest, especially for enterprise customers that make up the so-called Industry 4.0.

This fourth wave will represent the largest and most complex shift to date. Much of the cloud computing and telecoms network infrastructure is already in place, but it will take time and tactical nous for adoption on the business side. In the remainder of this analysis we explore the applications opened up by edge computing before discussing the competitive implications between cloud and telco companies.

Source: GSMA Intelligence

2 Edge computing is at the nexus of the cloud and telco worlds

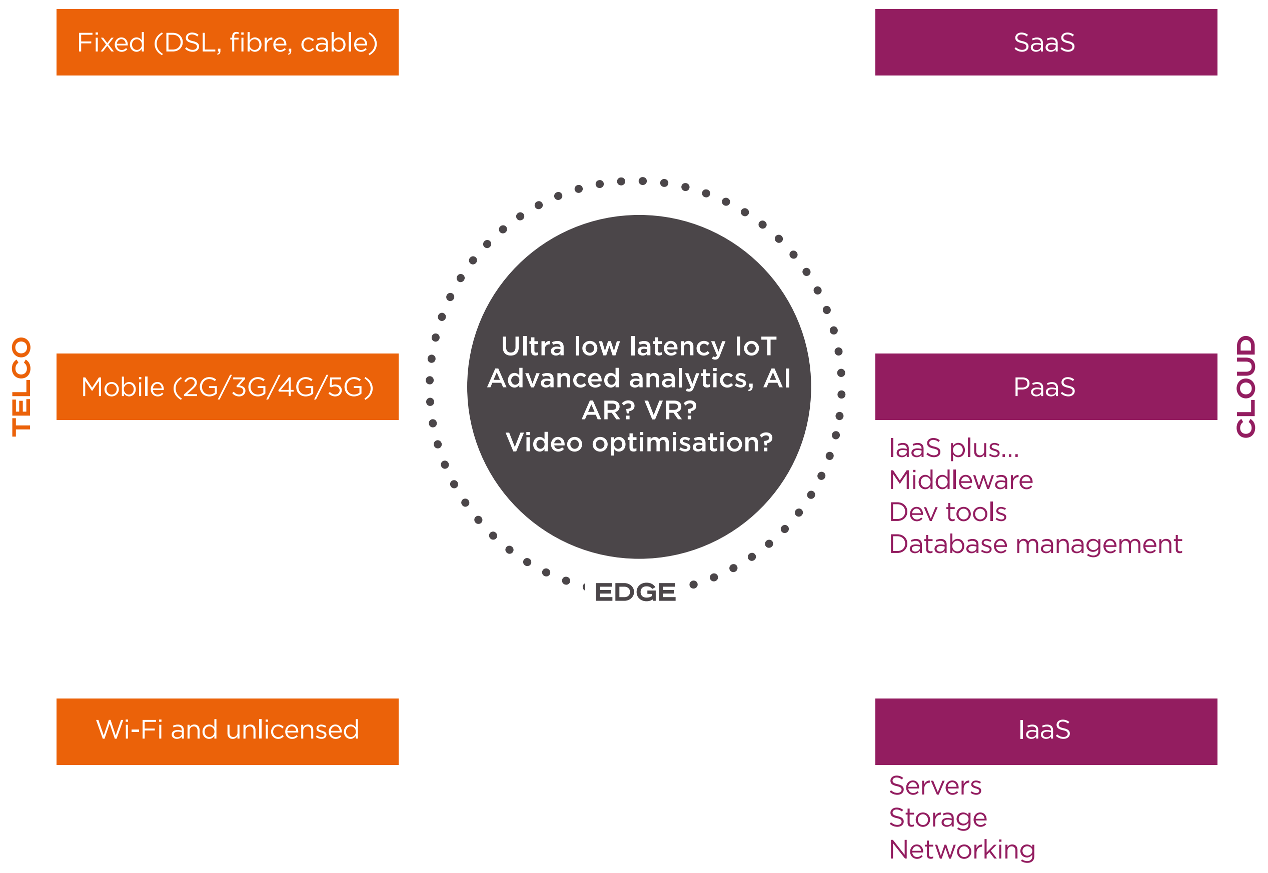

4.3 Latency and analytics define the sweet spot

The development of edge computing has proceeded somewhat like two sides – cloud and telco – seeking to locate the centre of a labyrinth from opposite directions. Both can see the promise of reaching a common end point but have different ways and strategies to get there.

The concept goes back to 2009 from a seminal research paper by Professor Mahadev Satyanarayanan at Carnegie Mellon University who proposed the development of virtual machine (VM) ‘cloudlets’. At the time, cloud computing and the data centres that support it were only beginning to gain scale and IoT was in its infancy so the concept never really gained commercial traction. The paper has, however, proved prescient: Amazon and Microsoft – the two dominant players in enterprise cloud – have launched edge products in 2017. Fog computing, a related concept of interconnected distributed devices capable of local processing to better service IoT, was also promulgated in 2015 through a consortium led by Cisco, Intel, ARM, Microsoft and Princeton University, although not Amazon.

The best way to think about the applications opened up by edge is on the vectors of latency and analytics. The most suitable applications, shown in Figure 3, demand very low latency (often less than 10 milliseconds), with functionality driven off real-time analytics capability:

- Ultra low latency IoT, particularly for industrial verticals – a non-exhaustive list would include monitoring wind turbines or electricity grids, transportation and logistics, remote installations such as oil refineries and mining rigs, aircraft and drones (for surveillance, delivery or other purposes) and robotics in high-tech manufacturing facilities. Smart city infrastructure does not generally require low latency but there are some exceptions such as traffic management, automated stop lights and real-time analysis of CCTV footage (to identify a criminal suspect).
- AR and VR – these are regularly posited as 5G use cases and as such require 1 Gbps throughput *and* sub-1 ms latency to render content and potentially tailor the experience for each person in real time in response to analytics feedback. Gaming use cases present a particular strain.
- Autonomous vehicles – these require edge functionality, though it remains to be seen what that looks like to satisfy rigorous safety threshold requirements.

Source: GSMA Intelligence

3 Use cases for edge computing

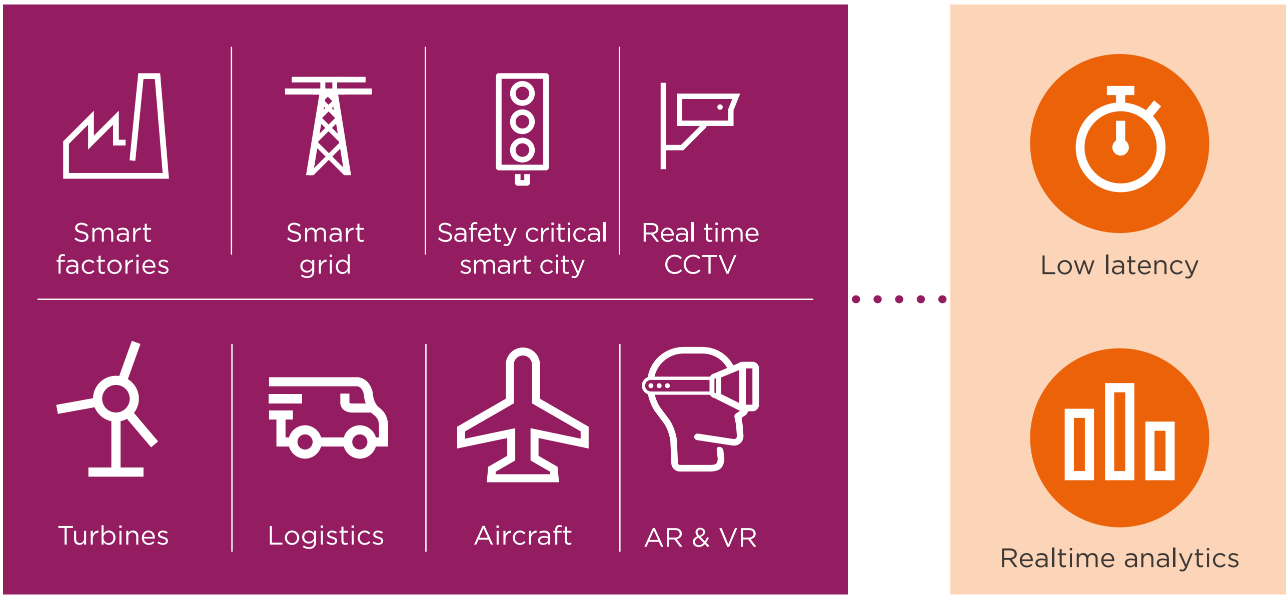

Meanwhile, *mobile* edge computing (MEC) has evolved in parallel as an offshoot that applies cloud principles to mobile telecoms networks. Cloud servers are situated within the radio access network as a means to provide processing power and analytics at the edge, rather than having to shunt traffic entirely to the core (see Figure 4). For example, a user's location can be triangulated within a group of cell sites, which raises the possibility of location-based services. Edge servers could also absorb a degree of content caching, providing a complement or potential replacement for servers currently used by content distribution networks (CDNs) such as CenturyLink/Level 3.

Source: GSMA Intelligence

4 MEC brings cloud to the mobile network edge

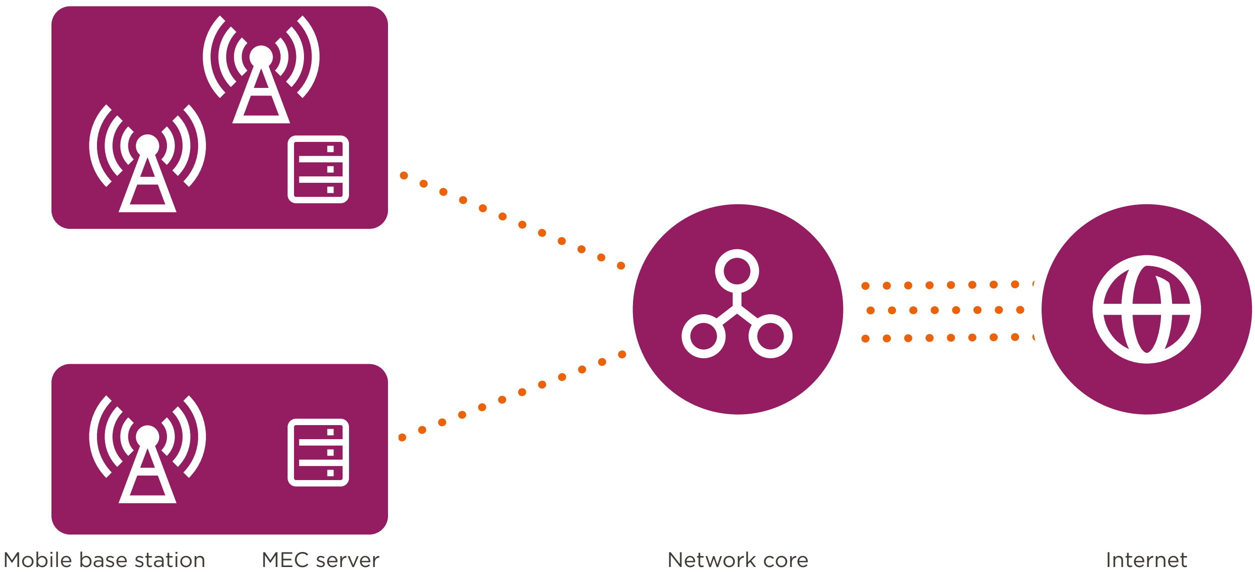

4.4 Competitor landscape: advantage Amazon and Microsoft

Our competitive analysis focuses on the two principal sectors with long-term strategic interests in edge computing: cloud and telecoms. Apple, through its iPhone X upgrade and newfound focus on AR, could also dictate part of the edge trajectory, even though it is not a direct investor in infrastructure.

Cloud

Previous computing cycles have lasted 15–20 years; if we take 2017 as the base year with the launch of AWS Greengrass and Azure Edge, the era is in the early adoption phase. Cloud, however, is not new, with a recent BDO survey[1] reporting that around 75% of US tech company CFOs see cloud as the technology with the highest expected impact on their business in 2017. This favours Amazon and Microsoft because each has highly scaled cloud infrastructure that can be used and parlayed to provide edge services for existing enterprise customers.

The economics of a scaled cloud footprint are highly profitable. Amazon, Microsoft and other cloud providers such as Google and Oracle report their footprints in terms of geographic availability zones, which is a region with one or more data centres (for example, US East or US West). Translated into square footage, we estimate Amazon controls around 8 million square feet of owned and leased data centre capacity, compared to 5 million for Microsoft and 4 million for Oracle. Because of multi-tenancy within a single data centre, scale has a multiplier effect on revenue. Amazon controls 59% more data centre capacity than Microsoft but its cloud revenues are 144% higher (based on the 12 months to June 2017 – see Table 1).

Profitability is high for both companies. AWS carries an operating income margin of 22%. Viewed in isolation this is not ground breaking but it matters hugely in the context of Amazon's low-margin businesses in e-commerce and content. Without AWS, Amazon would be loss making overall (see Figure 6). Microsoft does not report an operating income margin for Azure, but it does split out gross margin, which has increased to 57% as of September 2017 compared to 49% a year earlier.

So high are these margins and so dominant a position have they forged that other enterprise service providers, even giants such as GE and Salesforce, have conceded that they cannot compete solely by building their own data centres and now also rent capacity from the incumbents. The economics translate to the edge. Amazon Greengrass has a pricing structure based on flat-rate billing per device, providing in-built positive operating leverage as companies with large fleets or device portfolios sign on. Additional charges apply for messages sent to the central cloud.

1 2017 Technology Outlook Survey, BDO USA

Source: company reports, GSMA Intelligence

1 Cloud economics driven by scale – Amazon a clear leader

	Amazon (AWS)	Microsoft (Azure)	Oracle
Cloud revenue ($ million)*	14,528	5,960	4,572
Geographic availability zones	44	36	24
Datacentre capacity (million square feet)**	8.1	5.1	3.9
Cloud revenue per square foot ($)	1,789	1,169	1,168

*12 months to June 2017

**Reported figure as of December 2016 for Amazon; estimates made for Microsoft and Oracle based on published availability zones, each of which houses one or more data centres

Source: company results

6 AWS drives Amazon's profits

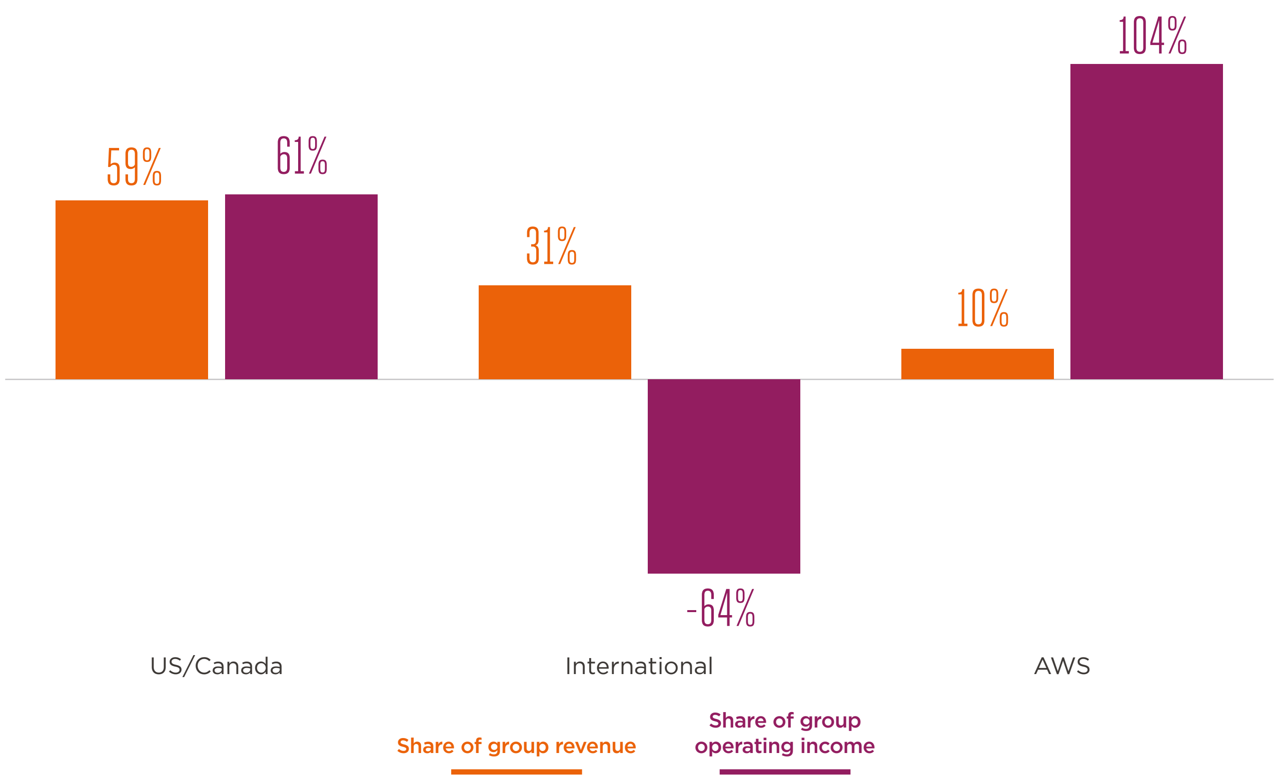

Note: segment figures calculated by dividing income at the segment level into total group income Figures are for the 12-month period to June 2017

Speed to market is also key. Amazon plans to expand its central cloud capacity into a further 17 availability zones, Microsoft by 8. It is no coincidence that their footprints largely overlap and are concentrated around cities and other commercial centres given the enterprise customer base. There is more pie to go around than either could eat on their own, so both can continue to build out capacity at rapid pace to cement their first-mover advantage. Amazon currently offers Greengrass in four availability zones (two in the US, Tokyo and Sydney) – although, as 2017 is year zero, we expect this to expand rapidly from 2018 along with feature updates such as analytics capability on the device.

Telecoms operators

Operators come at edge computing from a different angle. One obvious rationale is to increase network cost efficiency. As traffic loads increase, the more processing that can be handled at the edge, the more money is saved on transport to central cloud servers. Related to this is the repositioning of CDN servers. Edge is also an infrastructure prerequisite for 5G if sub-1 ms latency promises are to be met. Low-latency IoT is most interesting, though it is also where most of the competition overlap exists with cloud incumbents. AR, VR, industrial robotics and selected smart city infrastructure are examples. Operators have made significant strides in deploying NB-IoT networks to service low-power IoT devices that do not require low latencies. Given that low-power devices and sensors make up the bulk of the addressable IoT market size in volume – perhaps 80–90% – this is a positive step. Low latency is the other side of the IoT coin.

For industrial and enterprise segments, whether for a factory that manufactures high-tech equipment, an energy company monitoring the output of a wind farm, or a municipal government operating a smart city grid, the issue comes down to scale and quality of service. Telcos have global scale in aggregate but on an individual company basis coverage footprints for edge computing will be limited in early-stage deployments while commercial returns are tested. In this sense, the investment in edge servers should not be judged in isolation but rather as part of 5G benefits post-2020 (particularly cost savings). This is a sensible approach; taking on Amazon, Microsoft or Google at their own game is not.

The left-field option is reimagining the network as a platform to spawn an open ecosystem of developers making services based on the capabilities enabled by ultra-low latency. ETSI established an industry standards group in 2014 to define common protocols for MEC. The accompanying whitepaper articulated the laudable goal of aiming to “unite the telco and IT-cloud worlds” and “benefit a number of entities within the value chain, including mobile operators, application developers, over-the-top (OTT) players, independent software vendors, telecoms equipment vendors, IT platform vendors, systems integrators and technology providers”[2] – in other words, everyone. Not surprisingly, this idealistic view of value creation has not yet come to pass. But diagrams buried away in technical sections illustrate a path for third-party developers to access operator networks through common APIs.

Deutsche Telekom has a live accelerator (hub:raum) in Poland with a testing environment to try this out. On the research front, Deutsche Telekom and Vodafone are both involved with the Open Edge Computing collaboration with Carnegie Mellon. The first release of common APIs for MEC from ETSI came out in July 2017. These are positive signs, with APIs particularly important; in the absence of common standards, developers would have to design apps separately for individual operators – a major disincentive. In terms of business model, open must really be open: a set of common standards everyone is working from and openly welcoming in innovation from start-ups to established companies. Charging developers for API access is a non-starter. Finally, there is speed. In a platform model, the owner (operators) acts as gatekeeper for the apps that are deployed. The turnaround time for, say, a VR app would need to be in hours or days, not weeks.

2 http://www.etsi.org/technologies-clusters/technologies/multi-access-edge-computing

Apple and Google – spoilers?

For operators, opening up the network to exploit edge computing has the most potential in enterprise IoT. In consumer, the picture is more sobering, in large part because of Apple. With the iPhone X, Apple has signalled its intention to push AR as the next 'killer app'. The iPhone X features a new A11 bionic chip that uses higher grade silicon to enable neural learning and more data processing on the device. Like Amazon and Microsoft in enterprise, Apple has a running start in the form of iPhone and iPad owners that number in the hundreds of millions. The other side of its dual-sided model – developers – now have ARKit. Android has yet to feature AR to the same extent but it will only be a matter of time (as is the case for VR support). AR is likely to be a 5–10-year progression because a number of technical hurdles still need to be overcome. However, the scenario is developing where OS winners from the smartphone era – Apple and Google – use the same ecosystem pull model to entice developers and vertical sectors to develop AR services. AR is a transformative technology, meaning it is likely to attract disparate sectors from media to fashion to transportation. The risk for operators is missing out on service-level innovation and being relegated to bit pipes.

4.5 Future outlook

From a technology point of view, edge computing is still in its early stages, with the majority of enterprise workloads handled by centralised data centres. The decentralised nature of edge clouds and overhead costs of operating a large number of micro data centres mean the cost of running applications will be higher in the early adoption phase when there are fewer tenants to spread computing loads across.

Redundancy will need to be provisioned in a different way to conventional cloud given the distance between data centres. Pricing structures are also fair game: one of the chief advantages of edge is the intelligent division of tasks handled by central cloud (heavy lifting processing) and edge cloud (latency-sensitive, geolocation). If edge processing can be 'chunked' into individual tasks, the logical next step is that customers should only pay for what they use, as opposed to flat fees.

Our expectation is therefore that the growth in the adoption of edge will play out in phases as economics improve, upgrades are made to increase efficiency (such as nano-processing), and acceptance grows. We have outlined these phases in the figure below with the help of Macrometa Corp, a Silicon Valley venture-backed edge cloud platform provider currently in stealth. The ultimate end point is where the distinction between a centralised cloud and edge computing blurs or even disappears, although we would not expect this before 2030.

Source: Macrometa Corp, GSMA intelligence

7 Expected progression of the edge computing paradigm

1
CENTRAL CLOUD DOMINATES
2008–2020

- Cloud computing anchored in large, centralised data centres
- Mix of public, private and hybrid set ups
- Edge mostly a concept, not a commercial proposition

2
HYBRID CORE/EDGE
2017–2022

- Edge computing in early stages; AWS, Microsoft kickstart commercial launches
- Economics still less favourable than core cloud
- Intelligent sharing of tasks: core handles heavy lifting (e.g. database updates), edge handles latency-sensitive applications, geolocation, analytics

3
NATIVE EDGE
2022–2030

- Edge economics improved as workloads from larger number of enterprises spread across micro data centre footprint
- Other tech upgrades such as nano-processing
- Acceptance and trust of the technology grows with proof points
- Services designed *for* the edge. Such "native" apps include AR, VR, MR, autonomous vehicles, industrial robotics

4
PARITY?
2030–

- Economics and trust of edge computing reaches parity with core cloud
- Distinction between distributed and central cloud goes away

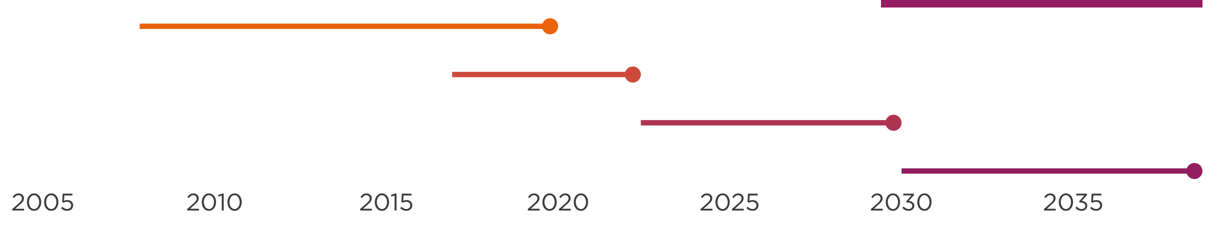

From a competitive standpoint, the fault line between cloud and telco sectors is likely to play out over this same 10–15-year period in mutual pursuit of connecting the 25–30 billion objects in IoT. From a revenue perspective, we forecast more than 50% of incremental revenue in IoT will come in enterprise and industrial verticals (Industry 4.0). For this reason, early advantage lies with cloud players who have in-built data centre scale, specifically Amazon and Microsoft with the two biggest footprints, live edge products in the market, and large enterprise client bases to sell into. Revenue from edge services is likely to be incremental to core cloud server operations because of the specific latency requirements from edge applications. Greengrass and Azure Edge are therefore products underpinning new revenue streams. The substantive difference in approaches is whether cloud platforms should be mainly closed (Amazon) or open (Microsoft).

The outlook is more nuanced for telecoms operators, who face the risk of being boxed in by Amazon/Microsoft/Google in enterprise and industrial, and by Apple and Google in consumer. In enterprise and industrial IoT (such as auto manufacturers and energy utilities), telcos can potentially win on latency and overall network reliability, but coverage footprints *per operator* will be low at first given the cost in deploying edge infrastructure. In the near term, the more likely return from edge is in cost savings reaped through transmission over 5G networks post-2020. Longer term, opening up the network to third-party developers offers potential for a range of IoT verticals (consumer less so given the threat from Apple), and if truly open could represent a sea-change in how business is done.

5

Realising the full potential of AR

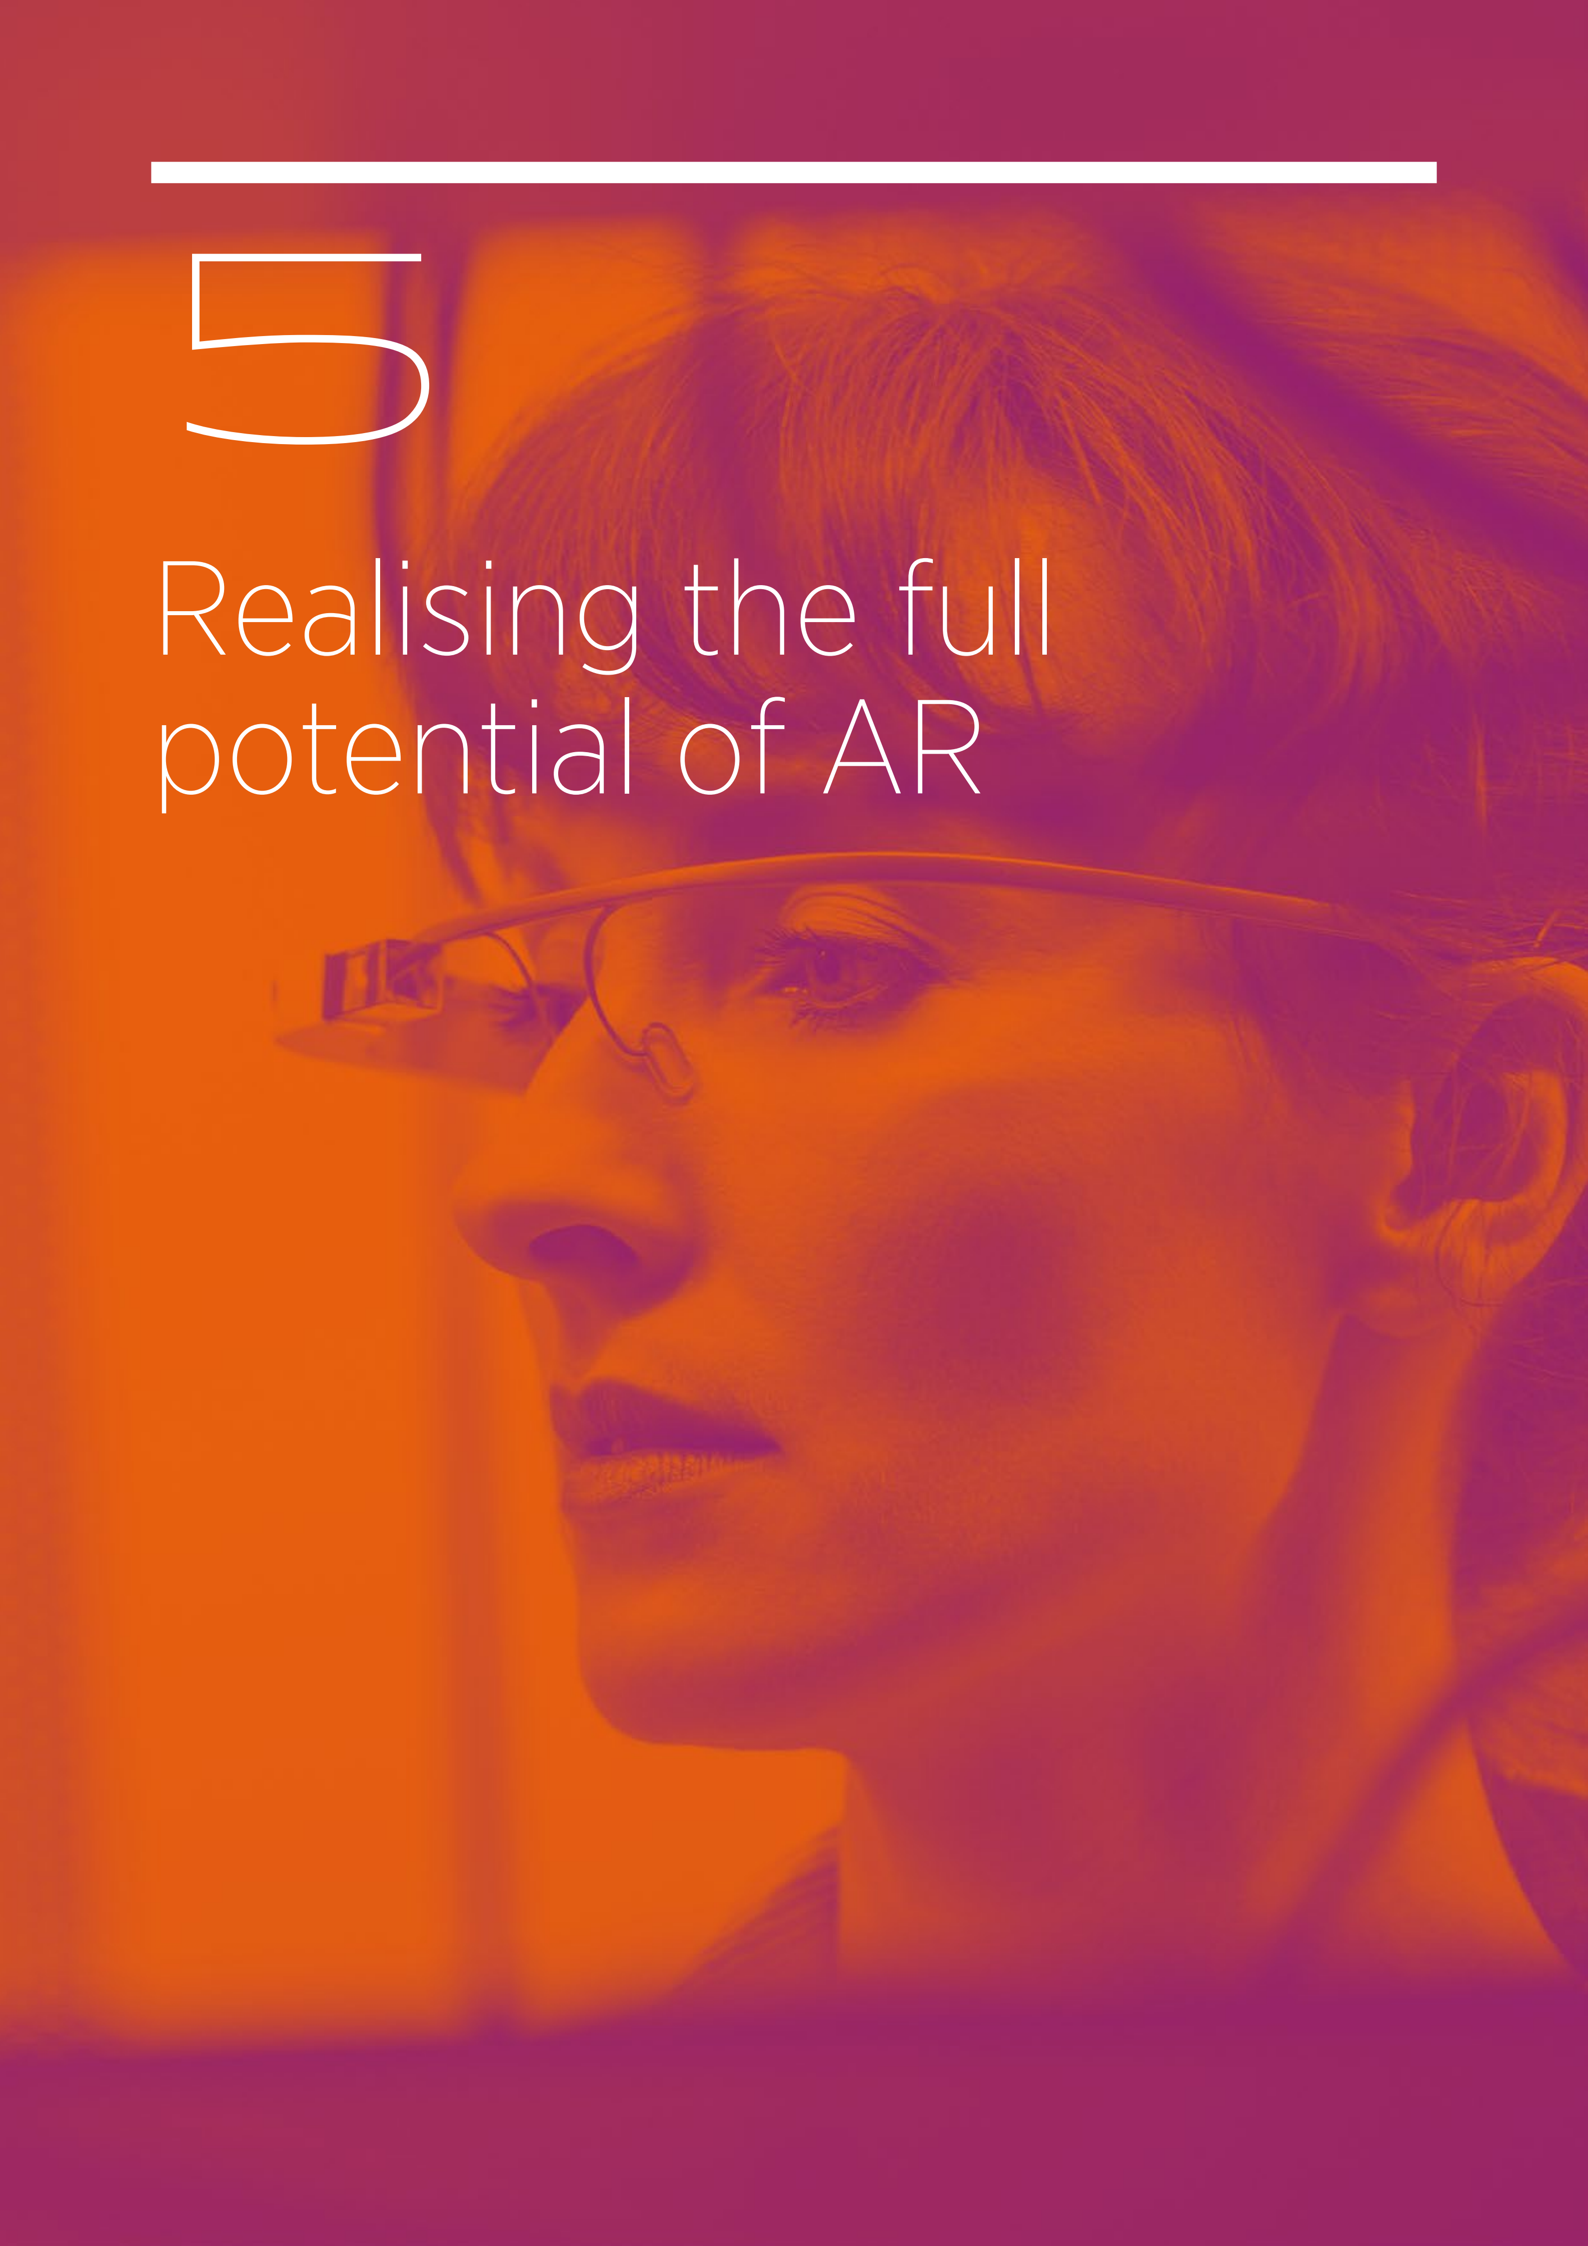

5.1 Executive summary

The recent announcements from both Apple and Google concerning new developer kits (ARKit and ARCore respectively) have refocused attention on the opportunities in augmented reality (AR), opening the door to a raft of new AR apps that run on existing high-end smartphones. This will bring AR increasingly into the consumer mainstream, with a growing range of compatible devices attracting apps and developers, including those disillusioned with the slower pace of development in virtual reality (VR).

However, the smartphone's 'magic window' is clearly a sub-optimal form factor if the true potential of AR is to be realised and mass-market adoption reached. An alternative form factor will likely revolve around some form of headset or glasses; the failure of Google Glass is perhaps over-applied as an analogy for wearables in general. There have been significant advances in both hardware and software since the launch of Glass, including improvements in processing power, miniaturisation and the supporting artificial intelligence (AI) needed to power a true AR experience.

A number of AR glasses are already coming to market, but with relatively limited functionality. A mass-market wearable form factor is at least two years away. In the short to medium term glasses are likely to remain tethered to a smartphone, reflecting the trade-off between an acceptable form factor and limitations in heat and power. The longer term solution may be to offload some of the heavy processing load to edge computing assets, allowing a standalone wearable device but a vast array of supporting infrastructure (both hardware and software).

Ultimately hardware may commoditise. The real value lies in delivering the combination of big data analytics and pervasive AI capabilities that allows the delivery of a seamless user experience and provides contextually aware information and services at the right time.

Previous platform shifts have seen new winners emerge, and paradigm shifts are by their nature typically disruptive. While the scale and reach of the current internet giants is unprecedented in the history of modern capitalism, the established incumbents typically look invincible right up to the moment that they no longer are.

There are broader implications for both individuals and society. Smartphones have fundamentally altered how individuals interact with each other and with digital content. The shift with AR could be even more profound, as the boundaries between the digital and real worlds blur, allowing the manipulation of digital images as if they are real. This in turn moves the spotlight to issues regarding the control of data and trust & privacy in an AR world.

In the long term, the AR platform will be transformative for businesses and consumers alike. As AR and VR increasingly merge, the new technology offers seamless integration with people's daily lives and a fusing of the physical and digital worlds. Significant improvements in artificial intelligence will deliver hyper-aware applications that are able to provide timely and contextually relevant information and guidance.

In an earlier edition of the Global Mobile Radar,[1] we looked at the evolving landscape around AR and VR, reviewing the key challenges these emerging technologies need to address before they can gain more widespread adoption. The last six months have seen a raft of announcements in the field of AR, with new developer kits from both Apple and Google. Here we examine the competitive landscape and likely future evolution of this new technology.

5.2 Current landscape

While the range of AR apps and services is still relatively limited, some early offerings have had considerable success and give some indication of the potential for mass-market adoption. Pokémon Go became the fastest mobile game in history to surpass $1 billion in worldwide gross revenues. Meanwhile, Snapchat has been experimenting for some time with AR filters and more recently with 'world lenses', which allow users to place 3D objects into scenes and manipulate them as if they are real-life objects.

Advancements in AR, particularly for consumer services, took a significant step forward in 2017 with the launch of two developer kits: ARKit from Apple and ARCore from Google. ARKit was first announced in June 2017 and demonstrated in September as part of Apple's iOS release and new iPhone model launches. ARKit makes it easy for a software designer or app maker to position a digital object in a real-world 3D space.

Google announced ARCore in August 2017. It follows the company's previous AR tool, Project Tango. This was released in 2014 but failed to gain traction; it relied on specific hardware that was only available in a limited number of devices. In contrast, ARCore requires only an Android-enabled device, but at present is limited to a handful of high-end devices.

Both ARCore and ARKit are software development kits (SDKs) that allow developers to access data from smartphones' cameras and sensors, recording data such as planes, spacial dimensions and speed of movement. This allows developers to model virtual 3D interfaces, which can then be used for a variety of applications, including shopping, gaming and education. Both kits will initially rely on smartphones, particularly newer smartphones (such as the Google Pixel 2, Samsung Galaxy S8 and iPhones 8 & X) with more advanced cameras with 3D perception and location/motion sensors.

These new SDKs and the growing number of devices that are compatible with them will create an attractive marketplace for developers, which should see a surge in AR content and applications. Certainly some of the developers that have been supplying content to the more niche VR market are likely to move across to the rapidly growing opportunity in consumer AR.

1 Global Mobile Radar, GSMA, May 2017

Source: GSMA Intelligence

1 Examples of AR consumer use cases

Type	Description	Examples
Interior design	Ability to superimpose furniture onto real rooms, change its orientation, etc. Can help designers' clients visualise options.	IKEA Place; Target
Instructional	For example, point your phone's camera at an appliance and it can give you operational info.	Hyundai Virtual Guide for auto repairs
Mapping and location information	Gives contextual information about sights and landmarks through your camera.	The UK Ordnance Survey Maps app
Spatial awareness/ measurements	With ARKit and ARCore, you can measure the distance between two points or dimension of objects.	AR Tape Measure
Retail – fashion	AR-enabled "virtual changing room" for e-commerce sites.	Several
Retail – restaurants	Access interior visuals, menus, reviews and table reservations by pointing camera at the exterior of a restaurant.	Vino Levantino wine bar in New York City – iPads and an AR application show desserts in 3D
Educational	Learning through visualisation, which can make grasping vocabulary and spelling easier and more engaging.	Very Hungry Caterpillar AR
Kitchen	Decorate cakes virtually; access recipes and nutritional information.	Food Network AR
Art	Artistic objects projected onto real-world space.	Jeff Koons with Snapchat Lens
Gaming	Largely based on the core AR principle of objects/characters projected into the real world.	The Walking Dead: Our World

Source: GSMA Intelligence

1 Consumer AR: expected timeline of development

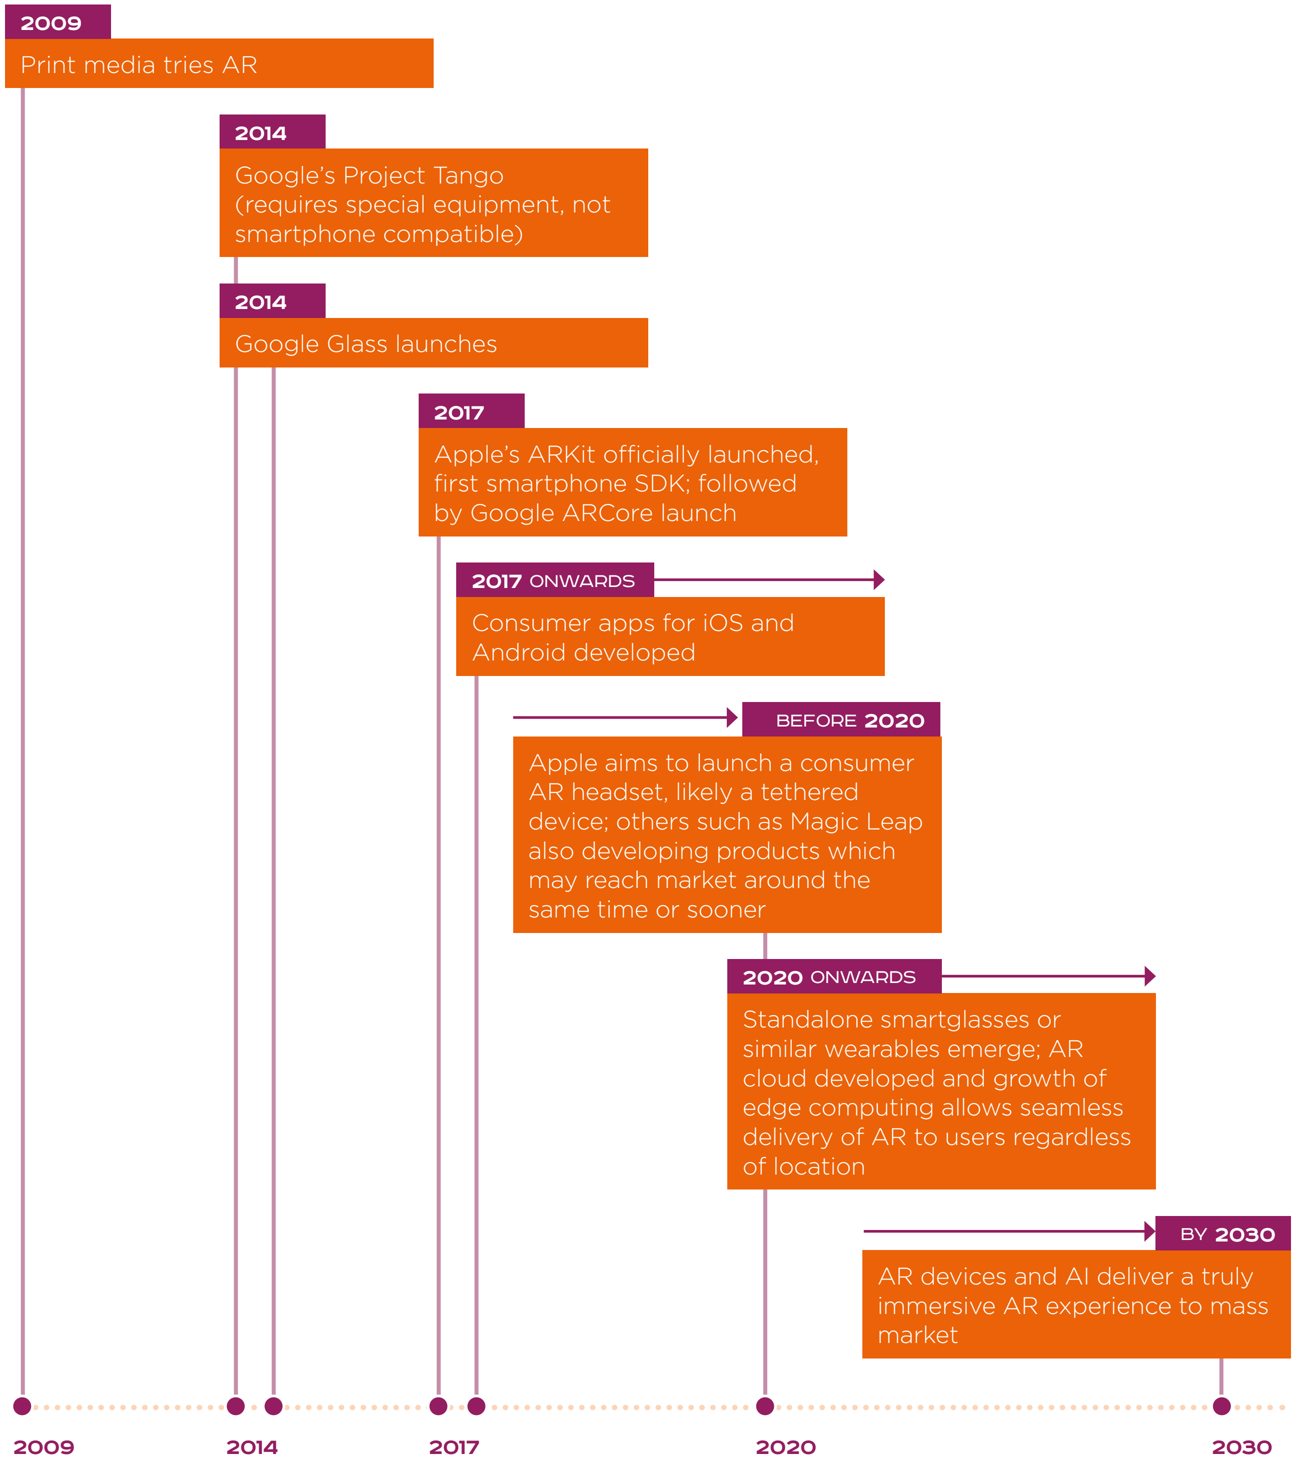

5.3 Key components to realising the full potential of AR

The building blocks for a genuine AR consumer experience are now being assembled. There are three key components to realising the full potential of AR: the right hardware form factor with a natural user interface; a pervasive 3D digital map of the world; and access to an advanced intelligence and huge variety of data to add the 'right' information onto this digital canvas.

Source: GSMA Intelligence

2 Key components required to realising AR's full potential

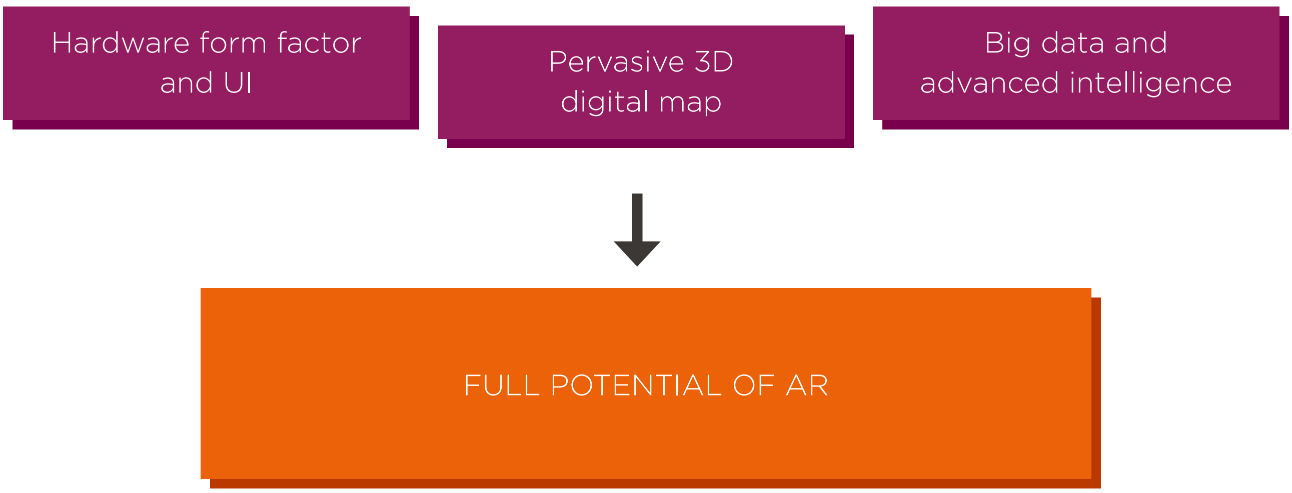

Hardware: wearables versus the 'magic window'

Recent moves by both Google and Apple focus on maximising the utility of smartphones as the main hardware platform for AR. However, the smartphone screen (or so-called 'magic window') suffers from many challenges from a consumer utility perspective and appears unlikely to be the winning form factor that will allow AR to become mainstream and integral to daily life.

Despite the likely advances and the large number new apps and products that will be launched over the coming year, it seems unlikely that smartphones can provide the required user experience.

> "Holding your arm out to look through a tiny screen has got to be one of the worst form factors ever accidentally invented by man."[2]

What seems more likely is that the form factor will focus on some form of headset or glasses. However, alternatives to smartphones are still at an early stage of development, with the industry still in discovery mode. The first generation of headsets have tended to be large and uncomfortable.

2 "A Few Good Apps Won't Save Apple AR - The Future Will", Forbes, October 2017

> “We know what we really want: AR glasses. They aren’t here yet, but when they arrive they’re going to be the great transformational technologies of the next 50 years.”[3]

Google Glass is a salient reminder of the dangers of releasing an immature technology to an expectant public without a clear purpose or use case. However, the relevance of Glass as an indicator for consumer willingness to adopt an AR wearable can perhaps be overstated, particularly given the lack of clear use cases or killer apps for Google Glass. It did not offer a true AR experience; rather, it offered a limited internet feed on a small screen and the ability to take pictures and videos. It also scored poorly on the aesthetic front.

Start-up Magic Leap recently announced a 2018 launch for its first headset, although with an initial high-end and enterprise focus. Press reports have suggested that Apple has plans for its own glass product to succeed the iPhone, with the goal of reaching the market in 2020.[4]

Even before the launches of a Google Glass 2 or Apple Glass, a number of new AR glasses are already available or about to be launched. While these may lack the true immersive capability of the full AR experience, they should help prove both the utility of AR glasses and help build consumer acceptance of AR glasses as a mainstream product:

- **Everysight Raptor AR cycling glasses:** this headset is due for release in the US in early February, with a base model price of $579. The glasses are designed to replace on-bike computers, are tethered to a smartphone and can be connected to other sensors. As well as accessing calls and messages, the glasses will project performance and mapping information onto the glass lens. The glasses will also take pictures and video based on either touch or voice commands.
- **NUVIZ heads-up display for motorcyclists:** a relaunch of an earlier product that failed to make it to commercial launch. NUVIZ is a smartphone-sized device that attaches to a helmet and tethers to a smartphone. NUVIZ uses a GPS signal to display a speedometer and can also display a route map, with a camera for photos and video. A Bluetooth connection to the phone allows the device to play music or make phone calls through a headset installed in the helmet.
- **Epson Moverio BT-300FPV Drone Edition:** a specialist headset designed for drone users, offering a transparent heads-up display of the drone's live video feed or flight statistics, while also allowing the user to watch the drone itself. It is currently priced at $699.
- **Google Glass enterprise edition:** Google is currently focusing on the enterprise space with its AR eyewear. The enterprise edition has a number of features that differ from the original consumer version, including upgraded hardware, a red light indicator when recording video, and a transparent display (Glass Pod) that is detachable so can be used in safety glasses.
- **Vuzix Blade 3000:** designed as smart sunglasses rather than AR goggles, this offers a full colour display that mirrors what is being viewed on the smartphone screen. The glasses are compatible with Android and iOS smartphones, and can be controlled via head motions, a built-in touchpad or speech recognition. Full release dates are still to be confirmed, but the company has indicated a price of around $1,000 and demonstrated production build units at the recent CES event.

3 Oculus’s chief scientist Michael Abrash at Facebook’s F8 developer conference, April 2017
4 "Apple Is Ramping Up Work on AR Headset to Succeed iPhone", Bloomberg Technology, November 2017

These devices are typically either fairly niche in application (such as cycling or drone use) or designed as a smartphone complement. Certainly, they fall short of the immersive AR experience, which may mean waiting for Google Glass 2.0 or Apple Glass before consumers get a truer sense of the capabilities of AR. There are still a number of challenges that need to be addressed if a suitable form factor is to emerge and see everyday usage, particularly around aesthetic appearance and comfort:

- **Tethered or standalone:** the development of eSIMs provides the opportunity for glasses and other wearables to connect to cellular networks independently of devices such as smartphones or tablets. However, considerations around size and weight could mean displacing some of the processing load and wide area connectivity to a smartphone.
- **Cost:** hardware cost will be a key determiner of take-up. Apple has already shown how it is possible to continue to push the boundaries of what consumers will pay for a smartphone. But if AR glasses are to become mainstream, then high-end smartphone prices appear likely to indicate the top end of what consumers are willing to pay.
- **Processing power:** AR devices will likely use mobile chipsets, which are seeing increasing advances in both CPU and GPU performance. However, as John Carmack from Oculus has noted, the power of the PC will never get to a mobile platform.[5] It is possible to get up to one order of magnitude faster, but there are challenges in dealing with heat and power that will remain limiting factors for mobile devices.
- **Battery life and efficiency:** these are also key challenges for a portable AR device, though there is significant investment from VC firms into start-ups in this area, which is driving ongoing progress from a technological perspective.

Beyond these hardware challenges, there is the question of how the user interface will evolve. The ideal AR interface does not require a controller or the need to touch any form of physical device. One option would allow users to manipulate AR digital objects using their hands, as if they existed in the physical world. The user interface would then be based on touch and gestures that individuals are familiar with in the real world. Eye tracking and smaller gestures are alternatives, while voice interfaces have seen a renaissance with the development of smart speakers. It is likely that AR devices would allow a variety of user interfaces, depending on the situation and use case.

AR cloud and instant localisation

To be useful for individuals and offer a seamless experience, AR needs to understand exactly where a person is, a process known as localisation. This currently relies on the use of a SLAM (simultaneous localisation and mapping) system, which involves the creation of a graph of 3D points, providing a map of the local environment against which a device can orientate itself.

At present, when an AR app launches, it checks if there is an existing map of the current location (either stored in the device or potentially in the cloud). As few such maps exist today, the device camera creates a new 3D map based on its field of view. The more detailed the map, the greater the processing power and memory needed to generate it. There are also memory challenges in storing lots of mapping data points (or point clouds) – a challenge currently solved by deleting older data points as the device field of vision moves onto new locations.

An alternative to constantly creating new maps would be the ability to draw on existing maps, stored either locally or in the cloud. As well as being open to all users, the development of what has been termed the 'AR cloud'[6] could transform AR from a private to a shared experience. This would allow one individual to leave AR data and content that would then be viewable to other users (a closed user group or everyone).

This is far more detailed than Google Maps, with accurate data on an individual's surrounding environment and physical context, which a device can rapidly access and use to establish an exact position. The goal is an accurate 3D map available of almost every physical location, both public and private. This would effectively create a canvas onto which AR apps and services can be overlaid. It transforms AR from what is today a fairly private and individual experience to one that becomes more public and social.

5 Oculus Connect 4, Day 2 Keynote, October 2017
6 "ARKit and ARCore will not usher massive adoption of mobile AR", Super Ventures Blog, September 2017

There are several challenges to developing this AR cloud: will it be owned by one company or a collaborative initiative? How will offices and shopping centres be incentivised to map their spaces? Will individuals want to share their own private spaces? For consumers, it may well develop as part of the entertainment experience, with data shared with apps and services to unlock their full potential. For businesses and retailers, commercial opportunities around advertising and boosting retail sales are likely to be the drivers.

Big data, AI and the software ecosystem

Edge computing will also play a role in the development of AR. It is likely to be a key building block of the AR cloud. New solutions could also offload the rendering of real-time graphics from devices (smartphones today but even smaller form factors like wearables tomorrow) to edge computing resources. This reduces the need for processing power on the device, and – by drawing on edge rather than centralised cloud resources – addresses the critical latency issue that can be a major problem for the user experience.

The move to edge computing is in part driven by the realisation that for a growing number of functions it is not economic to move large amounts of data across the network and the cloud. The commercial deployment of 5G networks over the next decade will also help address issues around latency, combining with edge computing assets to provide the real-time tracking and digital overlay that AR requires.

The other fundamental requirement of AR will be 'intelligent' data – data that follows an individual around and can be used to provide appropriate information and content exactly when it is required. Though a shopping trip would benefit from information on prices and promotions, being bombarded with offers while on the way to a meeting is not going to encourage mainstream adoption. Timely route information, on the other hand, could be useful.

Advanced AI capabilities will be crucial to harnessing the vast amounts of consumer data generated and assembling the appropriate information at the right time. This will require a more generalised and pervasive AI than currently available, although advances in key areas such as neural networks continue apace.

AR will also require a new software ecosystem. As long as AR remains a largely smartphone-based technology, the existing app ecosystem will be how users experience AR and how developers are paid. However, a new wearable form factor would inevitably move beyond apps, relying instead on advanced data analytics and AI to solve the current discovery problem that undermines the app store model.

This in turn raises questions around business models and routes to monetisation for content and service providers. While the boundaries between AR and VR will increasingly merge, it may be that there are different routes to monetisation depending on application. For entertainment and gaming, existing subscription and purchase models will continue to apply, with VR opening the door to premium options that offer a more immersive experience.

5.4 Competitive implications

AR has been grandly announced as the 'fourth computing platform', following the growth of PCs, the internet and smartphones. Regardless of whether this grand promise will be realised, there are some fundamental questions as to the nature of the AR platform: will it be hardware or software based, or somewhere between the two?

The answer to this will have direct implications as to which companies will emerge as winners in the AR ecosystem. A hardware platform will inevitably favour Apple and potentially some of the existing scale smartphone vendors, while a software platform would play more to social media and messaging giants.

The emergence of a new technology platform runs the risk of disrupting the current competitive dynamics, as we saw with previous iterations of the dominant computing platform.

As we have seen with Apple in the smartphone ecosystem, the winner may be the company that ties the hardware and software together. Some of the existing AR applications need expensive hardware to function, though the danger of this approach is the limited uptake that prevents mass-market adoption.

The geographical dimension should also not be overlooked, with China in particular likely to see its internet players emerge as frontrunners in the race to take AR into the mainstream.

Data will be a key differentiating factor in consumer AR. Ownership of (or access to) the richest and most diverse data will become a key competitive differentiator. Google is among the best placed today in terms of the variety and amount of data it can access. However, as consumers move away from traditional search, the question is whether the company can maintain its lead as a new platform emerges.

Whether the next generation of AR devices are standalone devices or still tethered to a smartphone, the data-intensive requirements of AR provide an interesting revenue source for the operator community. Indeed, the promises of 5G around low latency, massive bandwidth and ubiquitous coverage are important elements for a mass-market AR play.

New AR devices may also find their way into the stable of operator supported (and even subsidised) devices, alongside smartphones and tablets today. Whether operators can play higher up the value chain in part is likely to depend on the outcome of some of the competitive dynamics described previously. If operators can deliver the promise of 5G alongside the wealth of subscriber location data, they could establish themselves as a key enabler of the AR ecosystem.

5.5 Future outlook

The building blocks for a genuine AR consumer experience are now being assembled. There are three key components to realising the full potential of AR: the right hardware form factor with a natural user interface; a pervasive 3D digital map of the world; and access to an advanced intelligence and huge variety of data to add the 'right' information onto this digital canvas. All three of these need to come together to deliver the full immersive augmented reality experience, something that at the current trajectory of technological advancement could take up to a decade.

Once these factors converge, it appears likely that new AR wearables can approach the level of consumer adoption that smartphones enjoy today, at least in the developed world. In the shorter term, the likely proliferation of AR apps and services will drive consumer interest in AR as well as generating increasing attention from the developer community.

The early winners from AR will be familiar names. Apple's strong consumer positioning and integrated proposition will likely establish it as an early frontrunner. Google will have perhaps a final shot at a mass-market consumer hardware proposition with Google Glass 2.0. However, as we have seen in the smartphone market, there is a tendency towards commoditisation even at the top end of the market, suggesting hardware alone will not be a key differentiating factor.

The growth of AR content will also provide a new wave of engagement for the existing social media and messaging platforms, including Facebook in western markets and Tencent in China. However, while AR may add to the consumer experience of social media, this does not represent the true realisation of the potential for consumer AR.

The greatest beneficiaries will be the companies that supply the advanced AI and software infrastructure that powers the AR experience. Previous platform shifts have seen new winners emerge, and paradigm shifts by their nature are typically disruptive. While the scale and reach of the current internet giants is unprecedented in the history of modern capitalism, the established incumbents typically look invincible right up to the moment they no longer are.

There are broader implications for both individuals and society. Smartphones have fundamentally altered how individuals interact both with each other and with digital content. The shift with AR could be even more profound, as the boundaries between the digital and real worlds blur, along with the manipulation of digital images as if they are real. While the use of Google Glass in public settings created some controversy, the AR experience could allow individuals to exchange a wealth of information about themselves, filtered according to context and location.

This in turn shifts the spotlight to issues around the control of data and trust & privacy in an AR world: in particular, under what circumstances consumers share their data, and how they want companies and advertisers to access their own AR world. The position of data privacy will be under even greater scrutiny: will individuals be able to exert more proactive control over the use of their data than is the case today? These issues will shape future business models for AR and may challenge the advertising-funded model that underpins much of the digital experience today.

6

Mapping's fourth wave

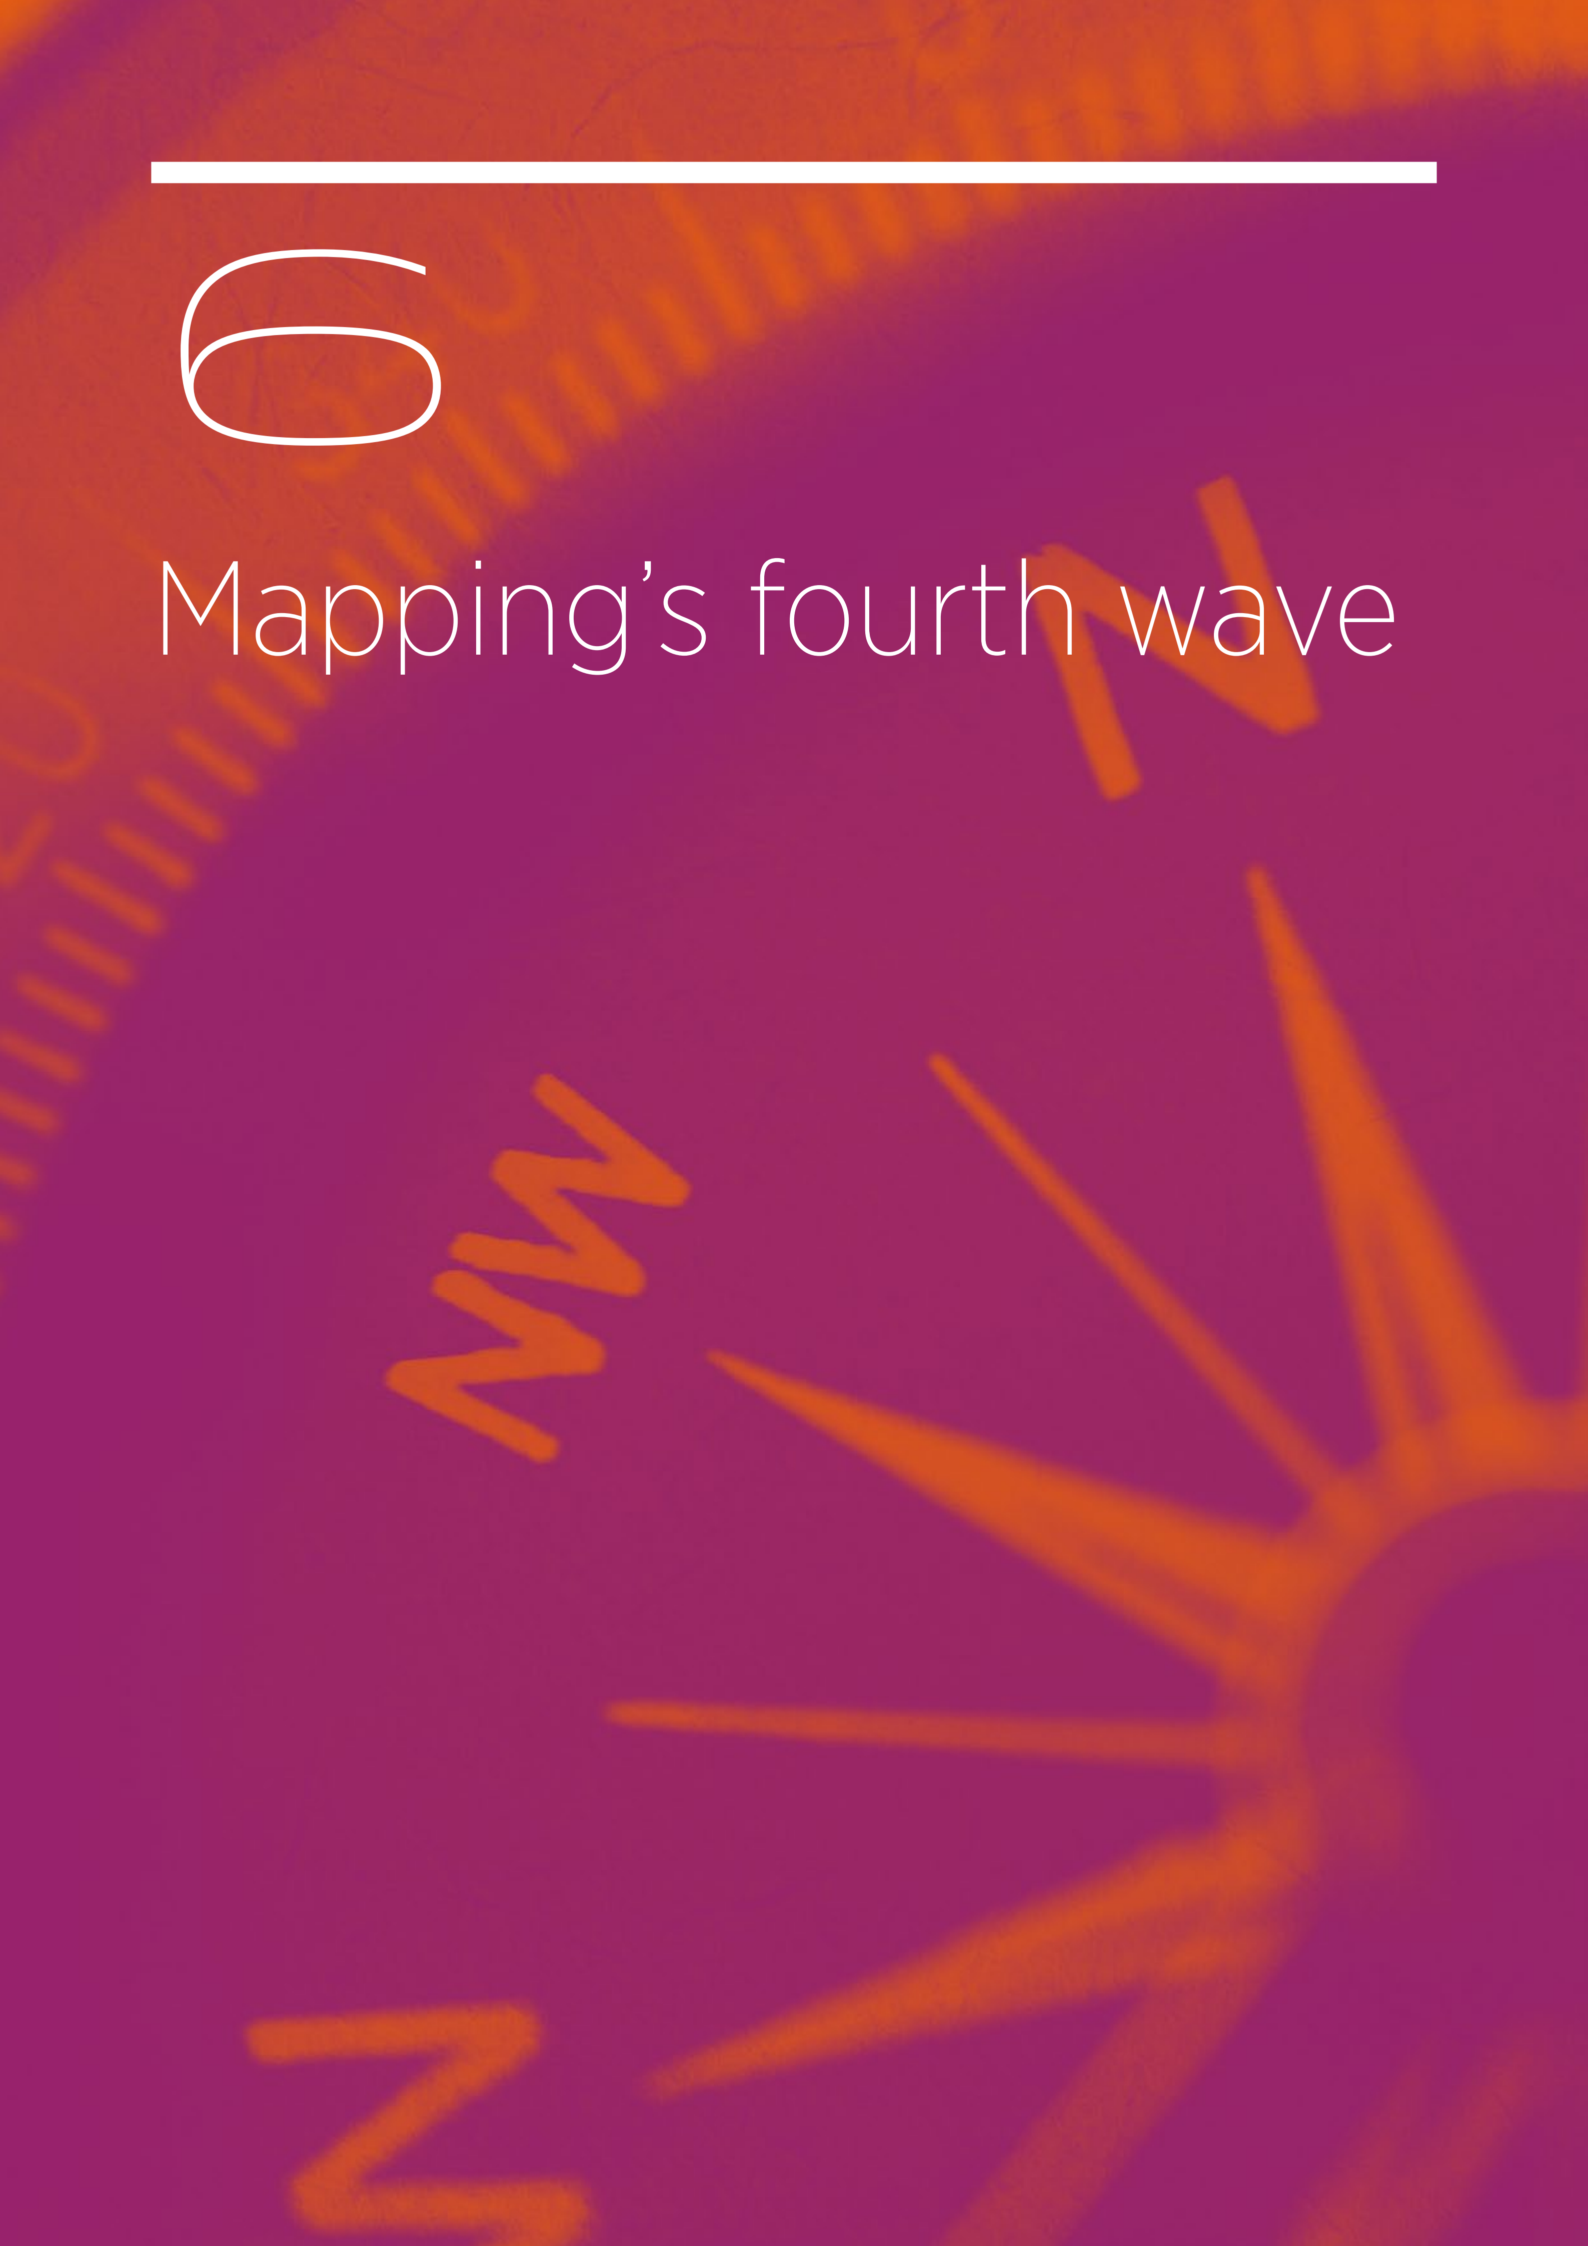

6.1 The 2,000 year-old platform powering the autonomous vehicle era

The earliest maps date back to the 14th century BC: cave drawings that recorded identifiable stars and constellations. Ancient Babylons, Greeks and Asians all created maps with geological features, roads, towns, city borders and directions.

By the 6th century BC, pre-Socratic ancient Greek philosopher Anaximander was the first to publish a map of the known world and is considered by many to be the first mapmaker. Anaximander depicted a flat, circular Earth centred on modern Greece.

In the early to mid-2000s OpenStreetMap, Google Maps, Mapquest and others began to digitise the physical, paper map. More than 2,000 years later, maps have become central to the digital ecosystem – from simple navigation around cities, to the underlying engine for e-commerce, transport and many billion-dollar 'unicorns'.

Mapping is now entering a fourth age – a race to build high-definition mapping fit for the robotics systems of an autonomous vehicle future.

Source: GSMA Intelligence

1 Anaximander's map of the world

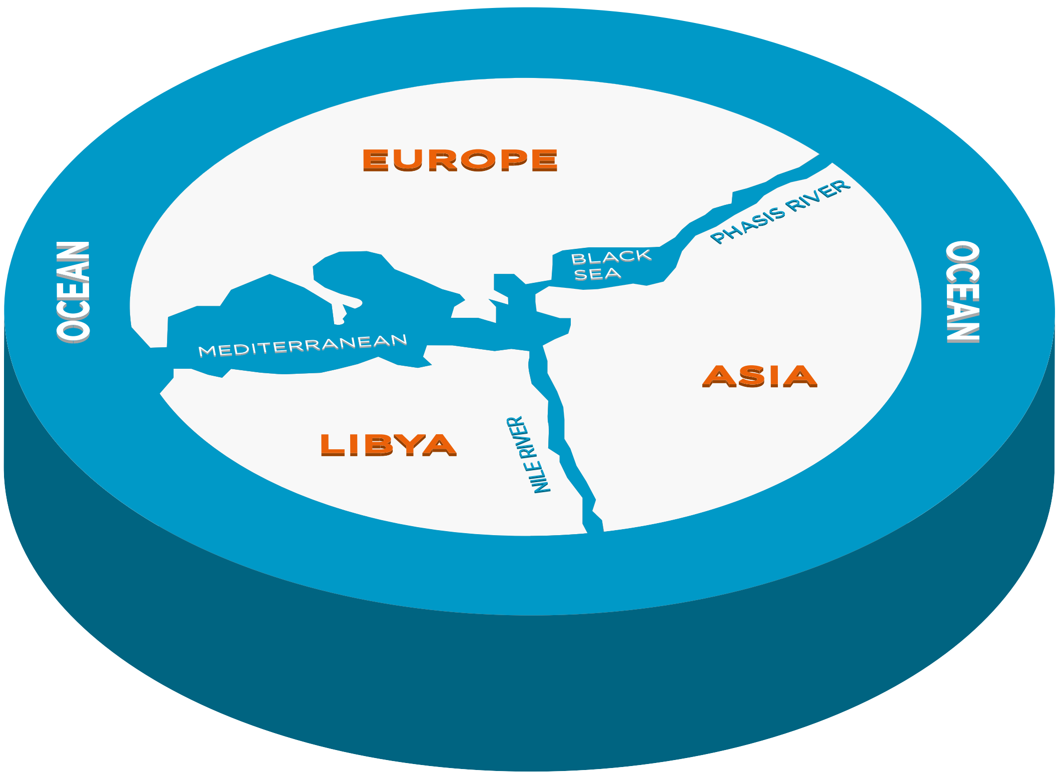

Source: GSMA Intelligence

2 Four ages of mapping

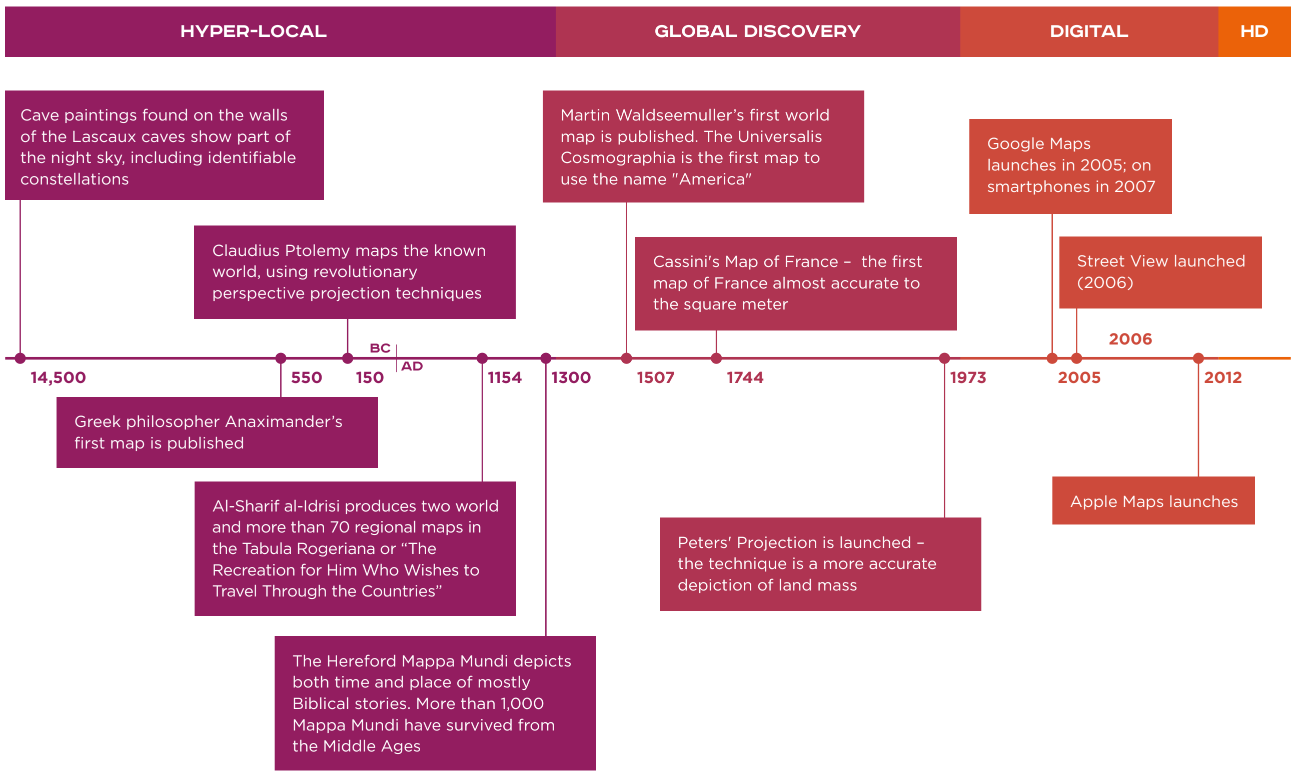

6.2 Google won the smartphone mapping platform war...

Google Maps has grown to 1 billion active users across desktop and mobile applications. It is now the dominant digital mapping product.

Initially fully developed and updated in-house by Google, the platform is now editable by an army of voluntary users.

Google has increasingly added a path towards commercialisation, including APIs that allow developers to embed ads on Google Maps. Enterprises can also pay for the rights to use the Google Maps API behind protected logins and intranets. Additionally, Google uses maps within search to enrich results.

Major competitors include OpenStreetMap, a free Wikipedia-like open-source mapping database; Apple Maps, launched in 2012; Chinese search engine Baidu's maps; and Waze, an independently-operated subsidiary of Google.

Source: comScore

3 Total unique users of mapping apps on smartphones in the US (thousands, November 2017)

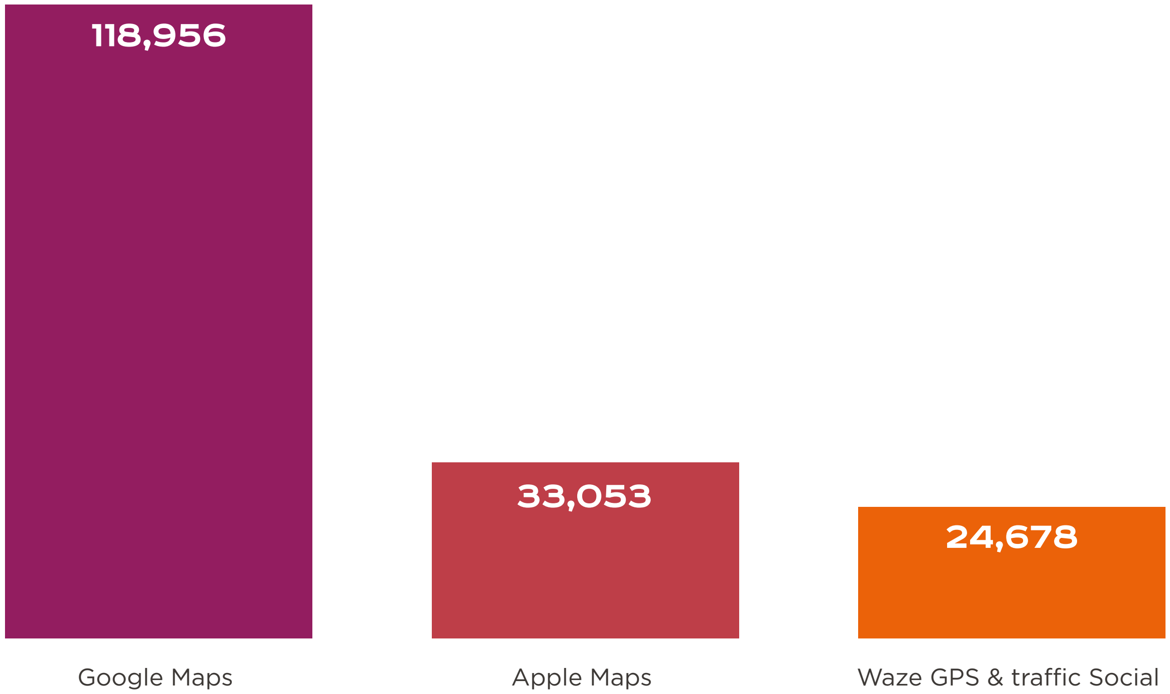

6.3 … but the race to win in high-definition mapping is still open

If you have an autonomous car, then the map is not going to be an optional feature. It's going to be a core component of the vehicle that will produce ongoing revenue.

John Ristevski, former VP, Here

High-definition maps have become a critical point of competition in the ecosystem because they will underpin the move from manual to autonomous vehicles (AVs), including the in-vehicle digital services offered in the AV era.

HD maps are essential to the functioning of AVs, particularly in scenarios where the car does not have full information on the environment through on-board sensors alone, such as in poor weather or in low visibility when objects and road markings may be obscured.

While there is some belief that neural network (NN) models will allow AVs to 'see' and interpret the environment on the fly, that vision may be some years away. In the meantime, the era of AVs, where computers make most of the decisions on roads, will require a new set of maps purposefully built for robotic not human systems. HD maps give AVs the ability to anticipate turns, junctions and objects far beyond sensors' horizons.

These HD maps offer orders of magnitude greater accuracy and detail, in a 3D environment, than the 2D systems largely used today. For example, HD maps will need to contain not only where lanes and road signs are, but where road boundaries are located, where the curbs are, and even how high the curbs are to an accuracy of centimetres.

While Google is an obvious frontrunner – due to more than a decade of mapping and imaging experience – numerous corporates and startups from the mapping, imaging and automaker sectors are competing for relevance and control of this new sector.

Source: Here

4 High-definition map for autonomous vehicles

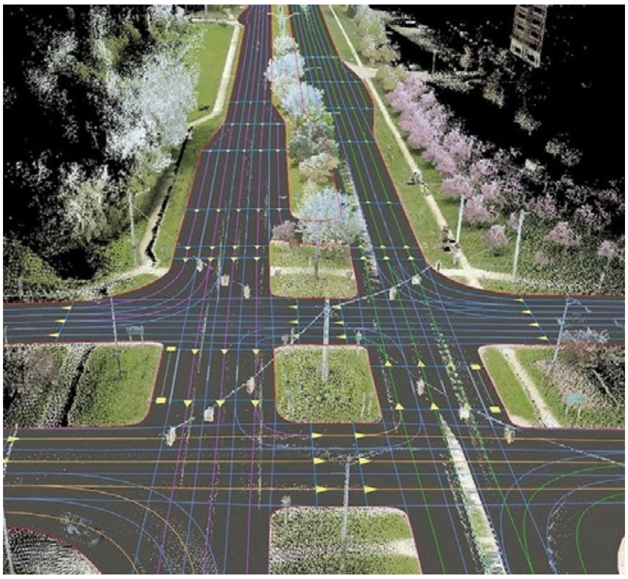

6.4 Autonomous vehicles' localisation, perception and mapping systems

The **planning and control module** decides what decisions to make and when – the central computer.

The **localisation module** tells the car where it is in 3D space, typically on a six-degree axis. For example, the system must interpret how far the car is from the next stop line or the boundaries of a road.

The **perception system** is the 'eyes' of an AV, and is based on several imaging and LiDAR sensors. The system has to see what's on the road and interpret objects and information.

The **mapping system** works in tight harmony with the other three components.

Source: GSMA Intelligence

5 Autonomous vehicle core systems

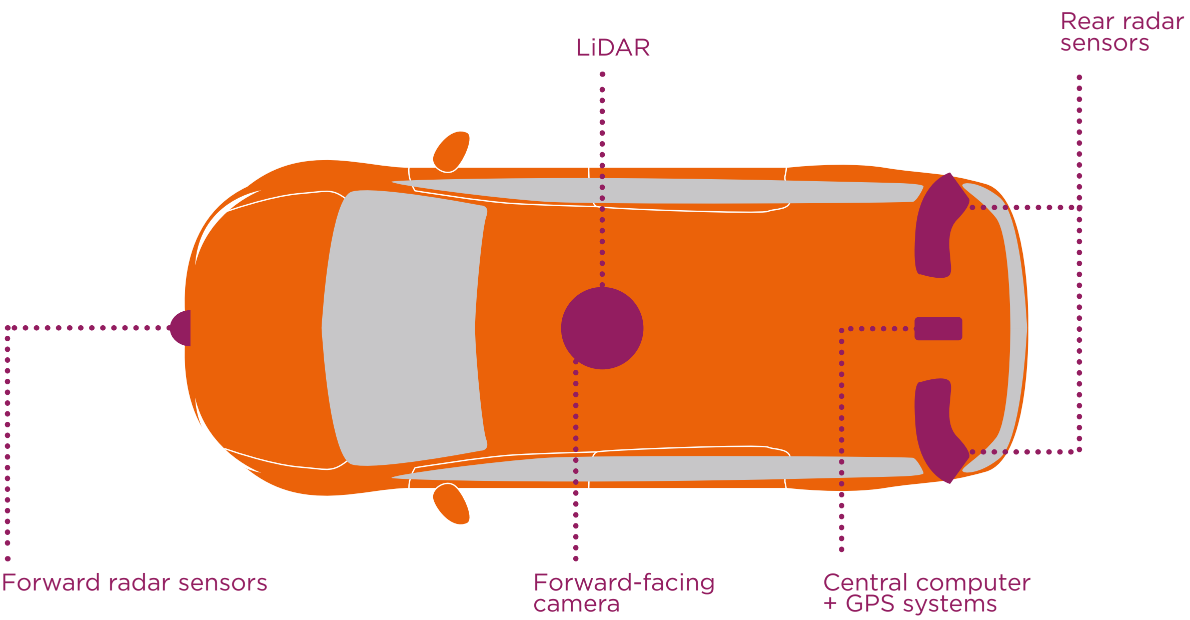

6.5 LiDAR – the sensory backbone of HD mapping

LiDAR
Light Detection And Ranging

LiDAR consists of a laser, scanner and specialised GPS receiver, and was used to map the earth from the air before adaption for AVs. It uses light in the form of a pulsed laser to measure range and variable distances to objects. These pulses are fired at up to 150,000 per second, with a sensor measuring the amount of time it takes for each pulse to bounce back. Since light is a constant of known speed, accurate distances can be measured.

This method generates precise, three-dimensional information about the shape of an environment and its characteristics. Typically both LiDAR and on-board cameras work together to collect huge amounts of data – essentially both consuming and creating maps at the same time.

The cost of an individual mechanical first generation LiDAR module was estimated at several thousands of dollars in 2015. Solid state LiDAR systems have reduced the cost of components to under $1,000 today, with a bill-of-materials for third-generation single-chip gallium nitride systems potentially falling to less than $100 according to some analysts.

Source: GSMA Intelligence

6 LiDAR

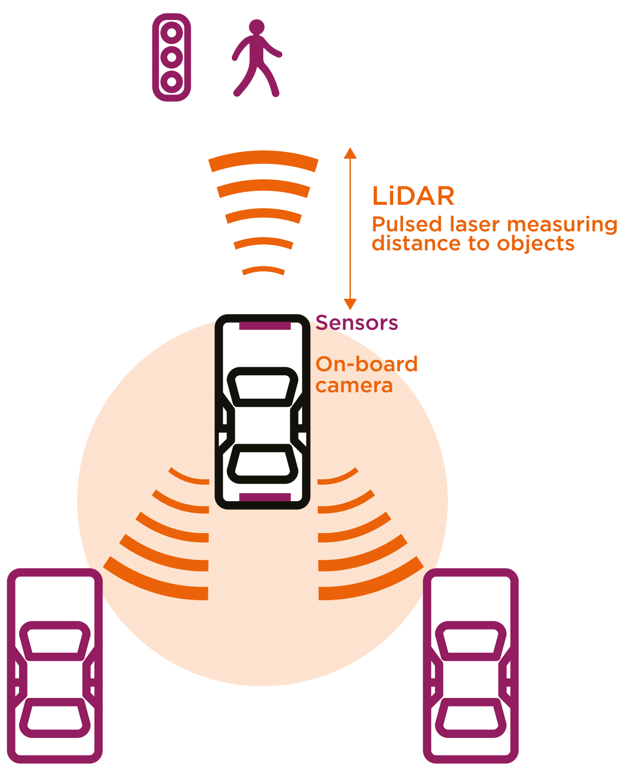

6.6 Five major players in HD mapping

Google is developing in-house high-definition maps for autonomous vehicles based on the millions of miles currently undertaken to map streets, develop Street View and test Waymo's autonomous vehicle systems. Waymo's cars have clocked more than 4 million miles of testing on public roads since the division's launch as part of Google X in 2009.

Uber uses a mix of mapping technologies, including those developed in-house together with third-party data, to provide the underlying infrastructure for its apps. The goal is to tailor systems towards service-oriented needs, such as data on traffic patterns and precise pickup and drop-off locations. Uber has been using Radar and LiDAR enhanced mapping cars in the US, Canada and Mexico.

Ford has invested in Civil Maps, a startup that provides 3D mapping technology for fully autonomous vehicles. Civil Maps uses AI and vehicle-based processing to convert on-board sensor data into HD maps. Ford is also working with DeepMap on a collaborative research project. The goal is to create high-definition maps with ultra-precise accuracy at the centimetre level.

Here Maps, once part of Nokia following the 2004 $8 billion acquisition of Navteq, provides location-driven data solutions and mapping. The division was divested in 2015, with investors now including Continental, Intel, Tencent, Audi, BMW and Daimler. Here introduced its HD Live Map product in 2016 as a system intended to provide AVs with detailed information on the environment and surroundings. It also has a crowdsourced site called Map Creator, with visitors contributing local updates.

TomTom's legacy in-car device business was disrupted by the smartphone and turn-by-turn navigation built into apps such as Google Maps, Apple Maps and Waze. Services for the AV market include HD and SD live mapping, localisation and map maintenance for autonomous driving.

6.7 Ten key startups in HD mapping

Mapbox

A platform to help users design maps and publish them across the web and mobile devices at scale. The company's tools give developers the power to make custom maps. SoftBank recently led a $164 million Series C funding round.

Total VC funding
$224.57 MILLION

GeoDigital

An analytics company for high-definition, 3D, geospatial data focused on mapping technology for autonomous vehicles.

Total VC funding
$29.12 MILLION

Civil Maps

A provider of 3D mapping technology for fully autonomous vehicles. The company uses artificial intelligence and vehicle-based processing to convert sensor data into meaningful map information.

Total VC funding
$6.7 MILLION

Carmera

Provides real-time 3D maps and navigation-critical data for autonomous vehicles, as well as 3D scene reconstruction and site analytics.

Total VC funding
$6.4 MILLION

DeepMap

Builds systems that enable self-driving cars to navigate through cities. The company will license its map-building software to automakers and other technology companies.

Total VC funding
$32 MILLION

Lvl5

Uses computer vision software to crowdsource high-accuracy maps for self-driving cars.

Total VC funding
$2.12 MILLION

Swift Navigation

Provides high-accuracy real-time kinematics, GPS and GNSS positioning technology for autonomous vehicles, as well as drones, robotics and space applications.

Total VC funding
$47.7 MILLION

Ushr

Provides HD mapping technology that combines with vehicle sensors and on-road cameras to convey "real-world" detail to AVs.

Total VC funding
$10 MILLION

Mapillary

Crowdsources maps with a service consisting of a smartphone app, website and an API. Users take street-level photos of roads, paths and buildings. Mapillary processes the images, enabling map-level access through an API for developers, app makers and mapping services.

Total VC funding
$9.55 MILLION

Mapper.ai

A stealth company developing machine readable maps.

Total VC funding
$8.35 MILLION

6.8 Venture-capital activity

Venture capital financing into high-definition mapping – a sub-group of the overall auto-tech sector – is small in global terms, but it has risen sharply in 2017.

We expect this trend to continue, given HD mapping's important role in autonomous vehicles, meaning an increase in both the number of deals and overall amount of funding through to 2018.

Today, the median deal size remains low, and 45% of deals have been at the seed stage. We expect both size of deal to grow and more funding to go into later stage companies in the coming years.

Source: CB Insights

7 HD mapping: funding and deal activity

$509.7M	41	$15.4M	$7M
Total funding	Deals	Average deal size	Median deal size

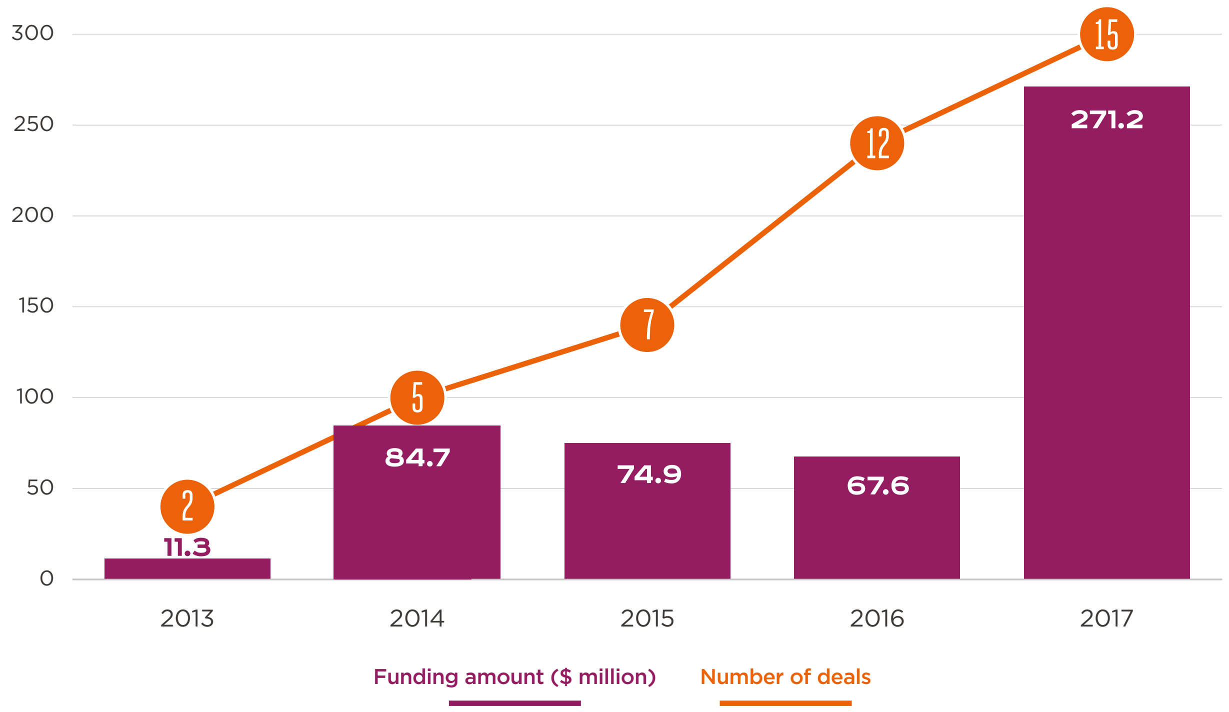

Blurred vision of the future?

Need foresighting for your business strategy?

Looking for the biggest disruptions and opportunities?

Engage with the GSMA Global Mobile Radar suite of products

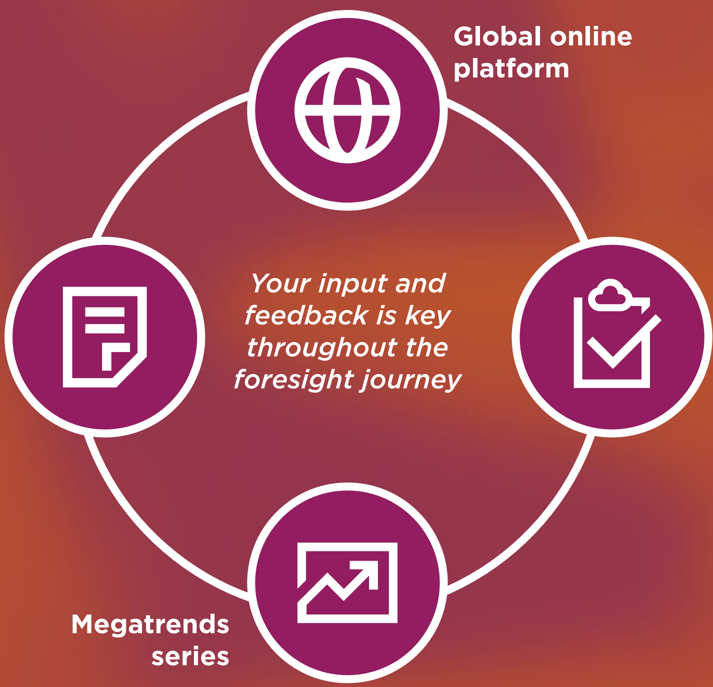

Access data, trends and insights from the GSMA through the **Global Mobile Radar** suite.

Strengthen your business strategy with our foresight intelligence into the digital ecosystem.

Participate in the global digital ecosystem with mobile operators, start-ups, think tanks, internet players and others.

Voice your opinions and requests in our webinars and events, and be part of the **Global Mobile Radar** community.

Contact

Matt Bonsall,
Director of Ecosystem Strategy

mbonsall@gsma.com

Access platform

gsmaradar.com

Access webinars

gsmaradar.com/faq

Authors

Blockchain: growing up is hard to do

David George
dgeorge@gsma.com

Edge computing: when telco met cloud

Tim Hatt
thatt@gsma.com

Realising the full potential of AR

David George
dgeorge@gsma.com

Mapping's fourth wave

Ed Barker
ebarker@gsma.com

GSMA Head Office
Floor 2
The Walbrook Building
25 Walbrook
London EC4N 8AF
United Kingdom
Tel: +44 (0)20 7356 0600
Fax: +44 (0)20 7356 0601