Evolving the service provider architecture to unleash the potential of IoT

etisalat

Introduction

The Internet of Things (IoT) has the potential to transform our everyday lives. Sensors and devices will capture and transmit data from billions of previously unconnected objects, generating greater insight and creating new ways for devices to autonomously interact with each other.

These insights, automated decision-making and device interactions have many potential applications across different industry verticals. For example, smart city applications (e.g. lighting, parking and waste management) will lead to the more-efficient use of public infrastructure, new healthcare devices and applications will provide additional patient information to doctors enabling better diagnosis and care, and sensors on industrial equipment will create a safer field-force environment.

Etisalat recognises the transformational impact IoT will have on society and is leading the way in turning this vision of a future connected world into a reality. Etisalat, as a communications service provider (service provider), will play a key role in building and managing the underlying infrastructure to support the Internet of Things. New enabling technologies and capabilities (e.g. network functions virtualisation, software defined-networking, edge or fog computing, 5G) will allow service providers to support the wide-ranging requirements of IoT services/applications, helping to deliver the full potential of IoT.

This whitepaper outlines Etisalat's vision of how the service provider network will support IoT. The discussion will take into consideration the impact of other key enabling technologies on the future of the network and IoT.

The report first sets out the landscape of IoT services, defining IoT and discussing the potential roles for service providers. An IoT reference model is then used to highlight the building blocks of an IoT solution. Next, the network architecture, designed to support the wide-ranging requirements of IoT services, is outlined and described. The impact of other key enabling technologies is presented in detail. Finally, an assessment of how IoT will impact service provider operations is discussed.

Mr Hatem Bamatraf CTO, Etisalat International

"Etisalat has been at the forefront in creating new technology breakthroughs in the UAE and the IoT is in our DNA. We believe the IoT is vital for regional development and have made a commitment to build best-in-class IoT capabilities to offer to the market. The '*Evolving the service provider architecture to unleash the potential of IoT*' whitepaper is an initiative that is aligned with the Etisalat 2020 technology vision."

Mr Khaled Ismaeel AlBelooshi VP, Fixed Networks, Etisalat International

"New technologies will be pivotal for IoT, enabling a much greater range of future IoT services. Using the capabilities from NFV / SDN, edge (fog) computing, low-power technologies and 5G radio access networks, service providers will be able to support the wide-ranging requirements of IoT. Etisalat Group is committed to developing the underlying infrastructure and technologies to support the future of IoT."

Executive Summary

The Internet of Things (IoT) presents an opportunity for service providers to expand their portfolio of solutions and offer new services to the existing customer base, as well as to new customers. Etisalat has shortlisted IoT Services as one of three key future service categories driving growth, in addition to Communication Services and Content+ Services[1]. Etisalat will continue to be at the forefront of the industry when it comes to IoT, pioneering the adoption of new technologies and embarking on new initiatives.

Throughout this paper, Etisalat illustrates the network architecture and technologies which will be necessary to support IoT and realise its own objectives for IoT:

Figure 1: Etisalat Internet of Things focus

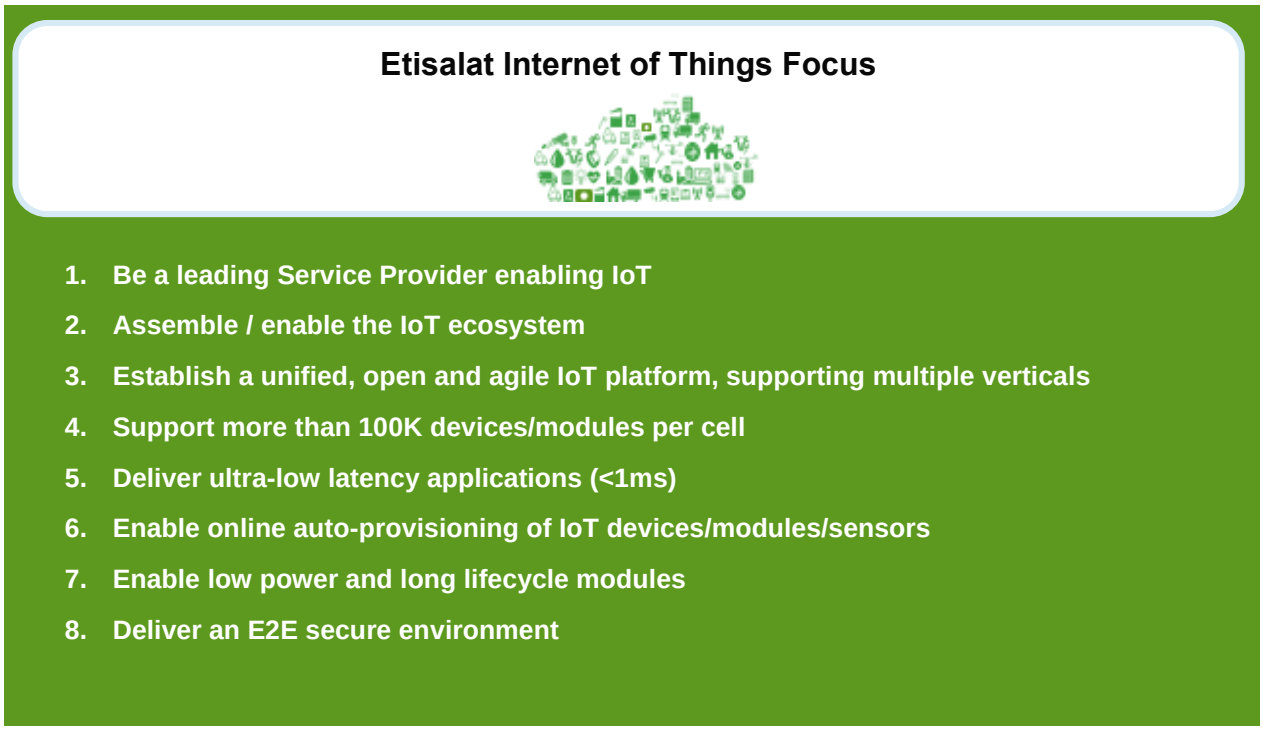

Source: Etisalat

In this paper, Etisalat recognises that the service provider's role within the IoT ecosystem will evolve to meet the customers' needs. Throughout this evolution, service providers will continue to provide the underlying network infrastructure enabling the communication of IoT devices. However, the role a service provider plays may extend beyond the network infrastructure, depending on the strategy of different service providers.

Regardless of the role of the individual service provider, the network infrastructure will still need to be designed in a way that will support a multitude of IoT connections, with a variety of network requirements. For example, smart city surveillance cameras will have reliability requirements in order to stream large amounts of video data without any risk of delay. Alternatively, simple sensors for a smart grid will have very low throughput levels, but may be located in remote areas, which have coverage and energy usage implications. Etisalat believes that the most challenging requirements to address in the future will be: high data rates; low power, low cost connections; massive numbers of connections; deep coverage; ultra-low latency; advanced data analytics;

[1] Etisalat – 2020 Landscape

security and resilience, or local autonomy. Satisfying all these requirements at once is not achievable with today's networks.

Given these requirements, the future service provider will need to manage diverse network infrastructure; for example, access technologies will continue to develop to meet the needs of new IoT use cases. Access networks will involve a mix of fixed & wireless technologies, long-range & short-range networks, capillary networks, as well as the potential use of both licensed & unlicensed spectrum.

Figure 2: Service provider network architecture to support IoT

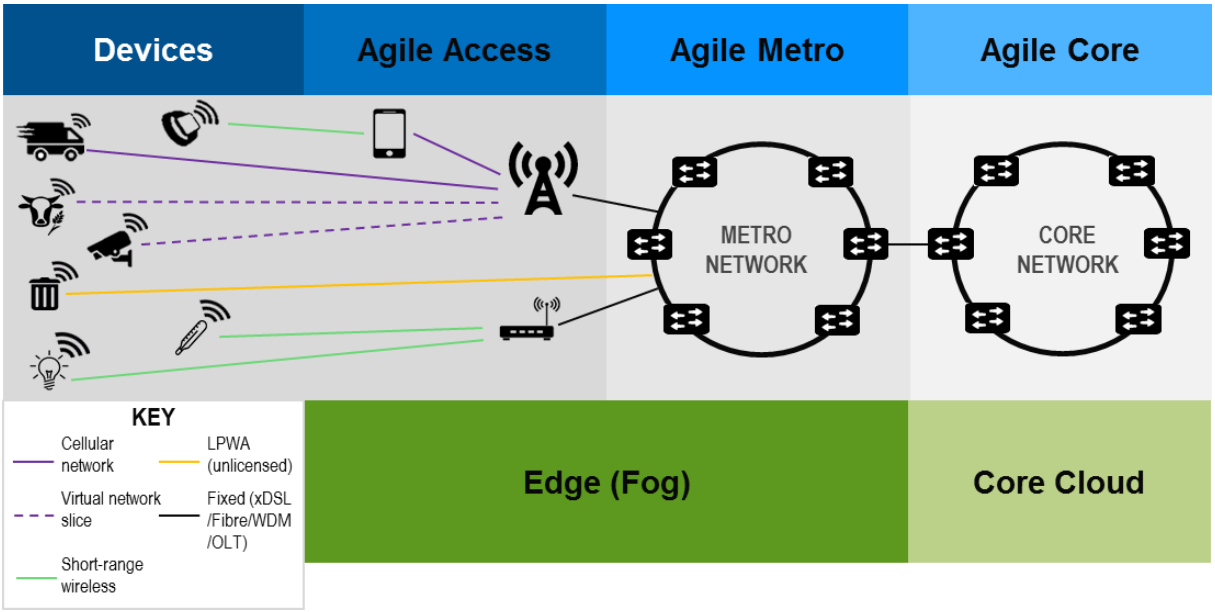

Source: Etisalat

In order to manage this infrastructure, service providers will need new capabilities, brought about by investing in new technologies and developing their networks. In this paper, Etisalat outlines NFV & SDN, edge (fog) computing, new access technologies for low-power IoT applications and 5G as the key technologies necessary to ensure the network is ready for the future of IoT.

NFV & SDN

Service providers will continue to transform their networks through the implementation of NFV & SDN. NFV & SDN provide the capabilities to manage the full range of IoT use cases over one network infrastructure. Network slicing, running multiple logical network instances over the same network infrastructure, will allow network resources to be allocated flexibly and in real-time. This will allow the network to efficiently accommodate IoT applications with radically different network requirements, creating different network 'slices' to simultaneously manage applications with high and low bandwidth requirements.

Edge (fog) computing

Cloud computing will continue to play a key role in IoT, whether it be for storing and analysing data from multiple sources, for cloud-based platforms (application, device management, connectivity management, etc,), to ensure greater security or to provide remote processing power. However, for some IoT use cases, where there are stringent latency or reliability/resiliency requirements, distant centralised cloud computing is not a suitable solution; compute needs to happen closer to the edge of the network, nearer to the device itself. Following on from the transformation to NFV & SDN, Etisalat believes that the traditional separation of network and applications will begin to blur. Service providers are expected to address this by making use of edge (fog)

computing, which will provide new, differentiated capabilities, allowing service providers to offer cloud computing and IT capabilities throughout their network, closer to the end-user or device.

Etisalat has identified three main reasons to implement edge (fog) computing for IoT:

1. **Network Efficiency:** The amount of data the network will be required to transmit will increase significantly over the next 5-10 years (some of which will be due to the rise of IoT), creating potential network capacity problems. Edge (fog) computing will be able to reduce overall network traffic and ensure better network performance for all users, as service providers can more efficiently aggregate and filter information that does not need to be sent back to the core.

2. **Creating Resilient Infrastructure**: Edge (fog) computing can increase the resiliency of local networks and operations. Localised computing can allow a system to operate even if outages have brought down other parts of the network or operations.

3. **Enabling Low-Latency Applications:** Bringing computing closer to the edge of the network will support applications that have stringent latency requirements. Etisalat expect the applications that will utilise this low-latency capability will be more futuristic (as the infrastructure is not currently in place today).

Access technologies for low-power IoT applications

New access technologies will be an important enabler for the growth of IoT, as a large proportion of massive machine-type communications use cases (those with a huge number of connected devices) are not fully addressed by current cellular networks. Today's networks do not provide a cost-effective option for connecting devices that transmit minimal amounts of data, whereas low-power technologies have the potential to connect devices over a long range, without placing a strain on energy resource. The landscape for low-power technologies is mixed between proprietary low-power wide area (LPWA) technologies using unlicensed spectrum (e.g. Sigfox, LoRa, Ingenu) and 3GPP licensed technologies (e.g. NB-IoT/LTE Cat-M2, LTE Cat-M1, EC-GSM.) The decision regarding which technologies to support will largely depend on the individual market characteristics and the future network landscape.

5G

Lastly, the evolution to 5G will gradually ensure significant improvements over the 4G/LTE standards currently in use. It is being designed with a focus on networks that enable machine-type communication (specifically M2M and IoT), in addition to improved 'traditional' mobile telephony. 5G will be fundamental in achieving certain network requirements for IoT: low latency, high throughput and better coverage; all in a more agile way, at a much lower cost.

As well as these enabling technologies, service provider's support systems and operations will need to adapt to support the Internet of Things. IoT will bring with it new customers, services and business models, creating potential challenges for service providers' OSS/BSS systems. This will be seen in the management and on-boarding of IoT devices, the diversity of business models and evolving billing arrangements. In order to mitigate and manage these challenges, service providers will need to adapt their OSS/BSS systems. This involves creating an open and flexible infrastructure using cloud-based OSS/BSS systems to ensure adequate elasticity and scalability, in addition to a comprehensive orchestration architecture, which oversees both virtual and physical aspects of the network.

Contents

Table of Exhibits

Defining the Internet of Things

The Internet of Things (IoT) refers to a growing network of physical objects ('things') and the ability of these objects to communicate and interact with/without the need for human engagement. Figure 3 explores this definition in further detail:

Figure 3: Defining the Internet of Things

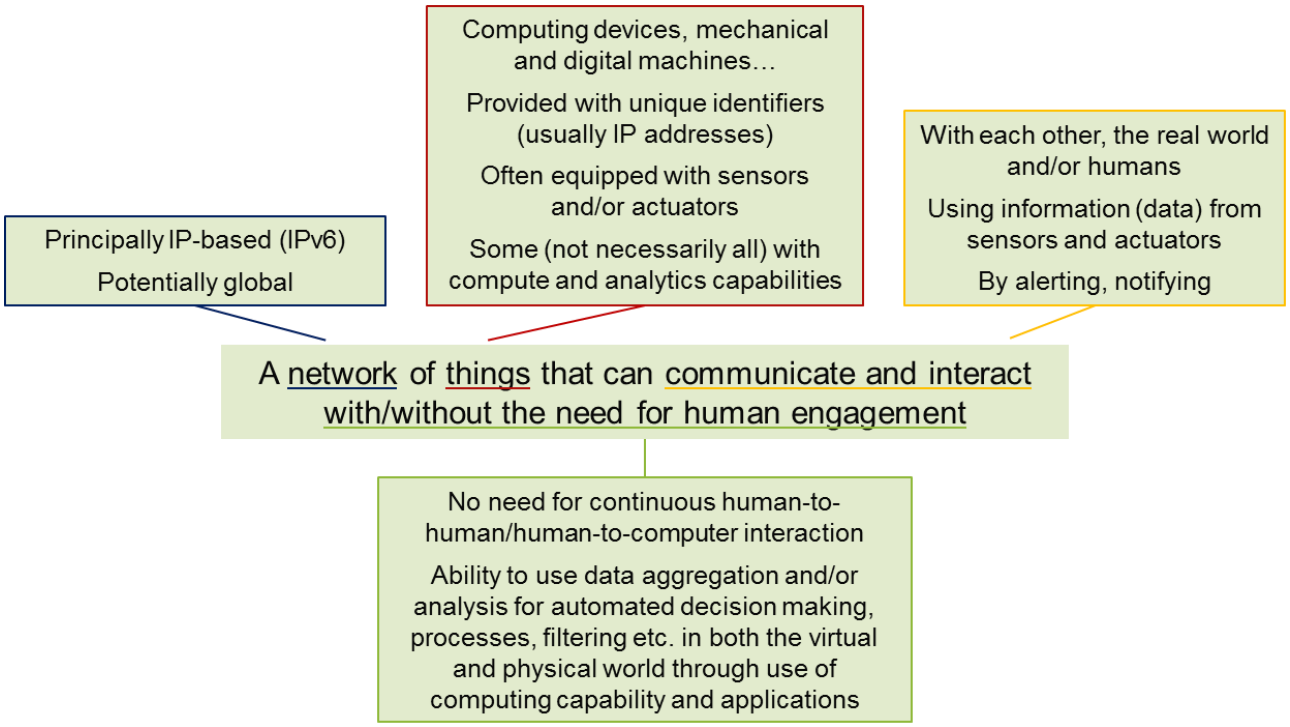

Source: Etisalat

IoT technology creates opportunities to monitor, sense and remotely control objects in the physical world. Potential advantages of this include increased productivity, reliability and quality of service, as well as reduced costs brought about by process automation and improved decision-making. This has the potential to create a number of new consumer applications as well as driving business process transformation, increasing efficiency and ultimately generating new revenue streams for organisations across the globe. One of the key characteristics of the Internet of Things is the role of IT (compute and storage) in delivering greater automation: empowering humans to focus time and energy on more productive tasks. However, humans can also use information transmitted or derived from IoT 'things' to enable smarter decision-making.

Key verticals, applications and use cases

Objects or 'things' which could fall into the scope of the Internet of Things include electronic appliances, lightbulbs, thermostats, heating elements, security sensors, cars, and many more. The ability to connect low-cost devices with limited computing power and energy requirements means that applications can be found across nearly every field.

Figure 4: Key verticals for the Internet of Things

Source: Etisalat

While the Internet of Things is not constrained to these fourteen verticals, Etisalat believes that they represent the most prominent application areas in terms of potential and current activity.

Many use cases involve collection and analysis of sensor data in order to augment and improve existing processes. For example, in a vehicle manufacturing plant, IoT sensor data is already used to measure production line efficiency. Environmental sensors (e.g. moisture sensors, thermometers, air flow sensors) can determine whether atmospheric conditions are optimal for spray painting of vehicle body parts, and automatically make changes (e.g. open/close vents, adjust air conditioning systems) as appropriate.

The Internet of Things will also enable applications that have never been possible before. This includes 'new' verticals such as smart home, smart building and smart cities, where sensor data has the potential to automate every part of our daily lives. Systems as diverse as domestic lighting installations, elevators in office buildings and city-wide traffic management systems can safely share data in real-time. Tasks that previously required human engagement, such as monitoring safety cameras, switching devices on and off, unlocking, opening and closing doors and adjusting temperature gauges can be efficiently automated, based on broad datasets including information such as user location or weather conditions.

The future of the market

Industry experts worldwide agree that the Internet of Things is going to grow massively over the coming decade. Analysts forecast that the worldwide IoT installed base will exceed 28 billion by 2020, while total market

revenue could exceed US$7 trillion[2]. Much of this growth will depend on the introduction of new technologies that enable the unique requirements of IoT networks.

Figure 5 shows how new service provider technologies will coincide with, and help to support, the rise in IoT. New technologies and capabilities, such as 5G and NB-IoT, will allow service providers to address a wider range of IoT use cases, including supporting applications utilising long-life IoT devices in remote locations as well as applications that have stringent latency requirements.

Figure 5: IoT growth (billions of connected things) and technology/capability timeline

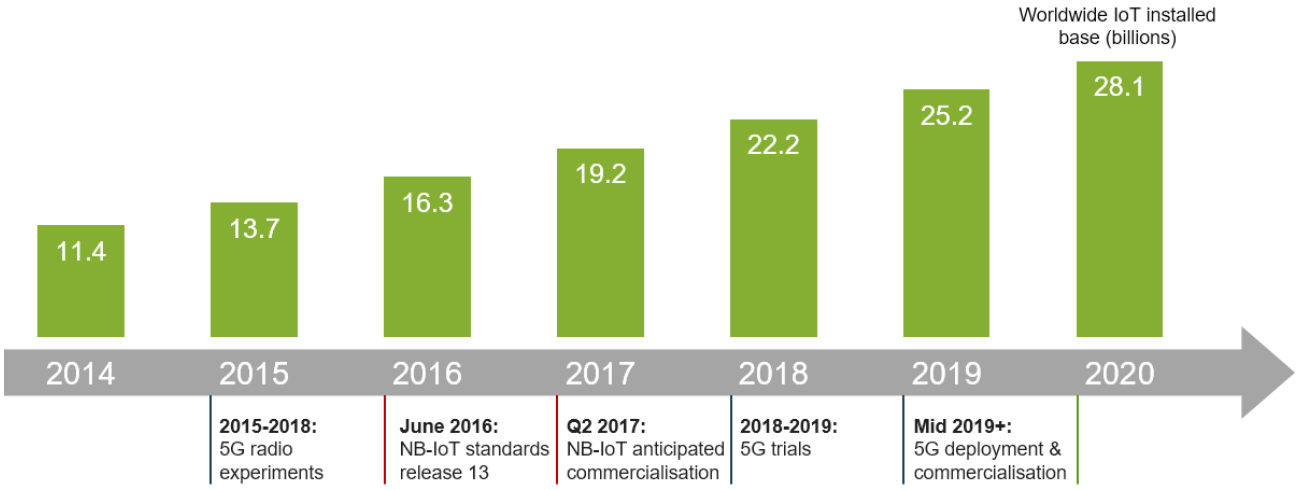

Source: IDC, STL Partners

As IoT technology matures, the Internet of Things will become an integral part of our world, transforming the way we do business and go about our daily lives.

The role of service providers

Communications service providers (service providers), including Etisalat, have unique capabilities and expertise that means they are well-placed to take the lead in enabling the Internet of Things. For service providers, networking and connectivity is a core competency. Service providers will be able to build on existing expertise, infrastructure and spectrum to deploy and manage IoT networks. In addition, service providers have strong backgrounds in functions such as billing, customer support and security, to which IoT services will need access reliably at scale.

[2] IDC: Worldwide and Regional Internet of Things (IoT) 2014–2020 Forecast

Figure 6: The IoT value chain

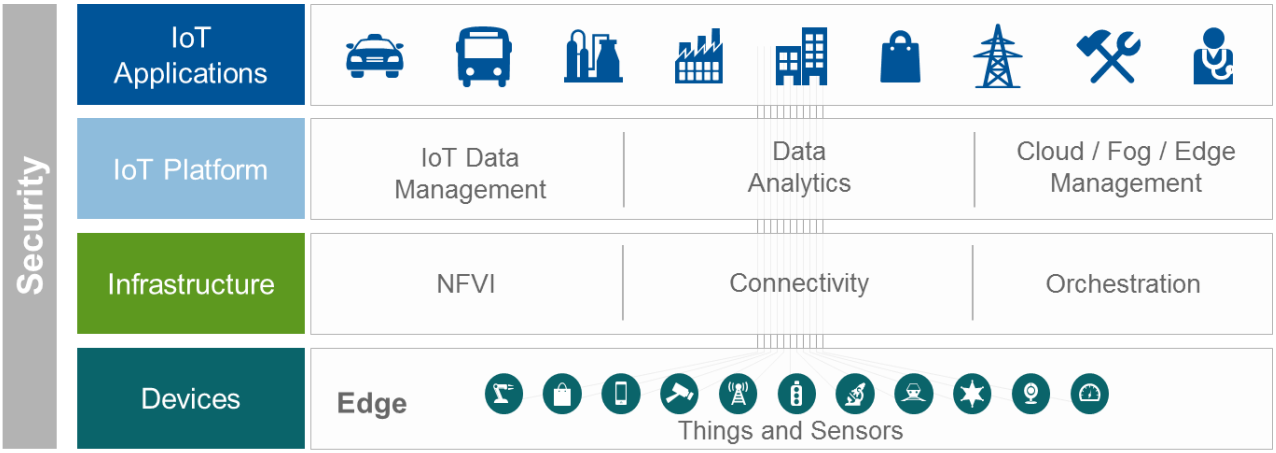

Source: IoTWF; Etisalat

Service providers may seek to play various roles across the IoT value chain. Roles could range from simply providing and managing the underlying connectivity infrastructure, to creating horizontal platforms, providing functionality and capabilities for a wide-range of IoT applications or providing end-to-end IoT solutions, including application development, for specific industry verticals. The role a service provider elects to play is a commercial decision depending on its strategy and market characteristics.

If service providers aspire to play a greater role in enabling the Internet of Things, for example through the provision of horizontal IoT platforms, they must embrace partnerships and seek to build ecosystems. Service providers can provide the underlying capabilities, including networking, compute, storage, billing, security and support. However, to enable the true potential of IoT they should create ecosystems, enabling the development of a wide-range of applications that will serve multiple industries.

Whilst the potential roles may vary, service providers will need to ensure that the underlying infrastructure is able to support the wide-range of IoT use cases. To do this, service providers may need to develop new networking and compute capabilities. For example, some low-latency IoT applications will require compute at the ‘edge’ of the network, closer to the end-device. Service providers may need to develop edge or fog computing capabilities to support these services.

The architecture presented in Figure 7 maps how service providers can support the range of IoT applications, including applications that require fog or edge computing. They can provide the underlying connectivity, compute and storage capabilities (in both the cloud and the fog), ensure security across the network, cloud and applications and provide analytical capabilities.

Figure 7: IoTWF Service Provider WG: Service Provider IoT Reference Architecture

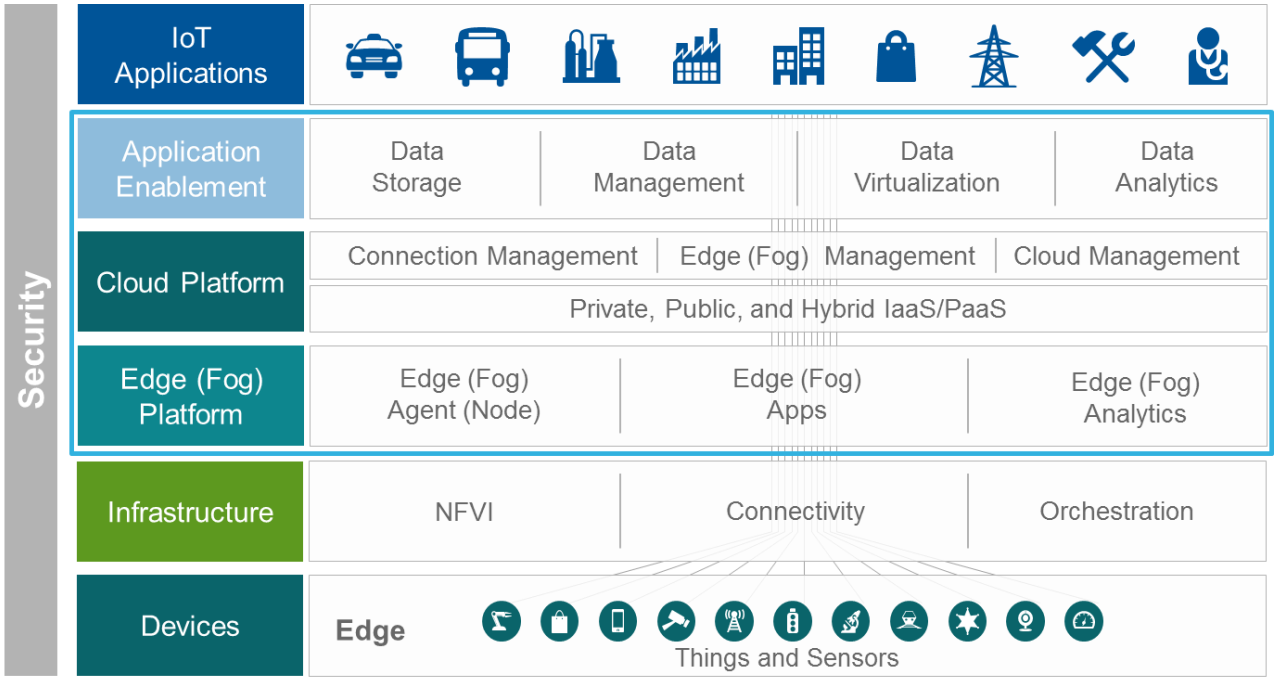

Source: IoTWF; Etisalat

Horizontal capabilities will allow service providers to address the cross-section of capabilities required for IoT applications. This includes:

- **Edge (fog) Platform:** IoT applications will require a combination of different types of connectivity (mobile, Wi-Fi, low-power etc.) along with the capability to carry out computing both at the edge, closer to the device, and in the cloud. Service providers will need to plan the deployment of edge (fog) computing capabilities in a scalable manner based on the range of use cases they look to support. Edge or fog computing will enable a range of IoT use cases, particularly in areas where large data volumes, remote locations or stringent latency requirements make local processing more optimal than shipping data to the cloud. Examples of these include remote site management and Oil and Gas exploration. Edge (fog) computing is discussed in more detail later in this report.

- **Cloud Platform:** Cloud is an integral component of any IoT solution. Service providers need to be able to assemble and integrate a critical set of services which can enable them to provide a scalable managed services platform for vertical applications. These include:
  - Connectivity, compute and storage management
  - Security services across network, cloud, and applications
  - Analytical services across applications

- **Application Enablement:** This will support the storage and transfer of data from devices or things to applications; this includes management of data, analytics and APIs as well as data visualisation. Service providers will also need to ensure security is managed effectively across the various layers of the IoT stack.

Beyond these capabilities, many service providers will be able to take advantage of their (often regional) geographic footprint, as well as existing brand recognition, to establish leadership in this growing field.

IoT Reference Model

Network, compute, application, and data management architectures that are IoT-ready require a different communication and processing model. Today, there is no standard way of understanding or describing these models for IoT. As a result, the lines are blurred between IoT devices/systems and non-IoT devices/systems.

Etisalat recognises that the IoT World Forum's Internet of Things (IoT) Reference Model is a useful step towards standardizing the concept and terminology surrounding the IoT. This section sets out the key levels of functionality, defined by the IoT World Forum's Reference Model, to provide readers with a baseline for understanding IoT requirements. The following section will discuss how the service providers' underlying infrastructure can support IoT.

In an IoT system, data is generated by multiple kinds of devices, processed in different ways, transmitted to different locations, and acted upon by applications. The proposed IoT reference model is comprised of seven levels. Each level is defined with terminology that can be standardized to create a globally accepted frame of reference. The IoT Reference Model does not restrict the scope or locality of its components. For example, from a physical perspective, every element could reside in a single rack of equipment or it could be distributed across the world. The IoT Reference Model also allows the processing occurring at each level to range from trivial to complex, depending on the situation. The model describes how tasks at each level should be handled to maintain simplicity, allow high scalability, and ensure supportability. Finally, the model defines the functions required for an IoT system to be complete.

Figure 8 below illustrates the IoT Reference model and its levels. It is important to note that with IoT, data flows in both directions. In a control pattern, control information flows from the top of the model (level 7) to the bottom (level 1). In a monitoring pattern, the flow of information is the reverse. In most systems, the flow will be bidirectional.

Figure 8: IoT World Forum – Internet of Things Reference Model

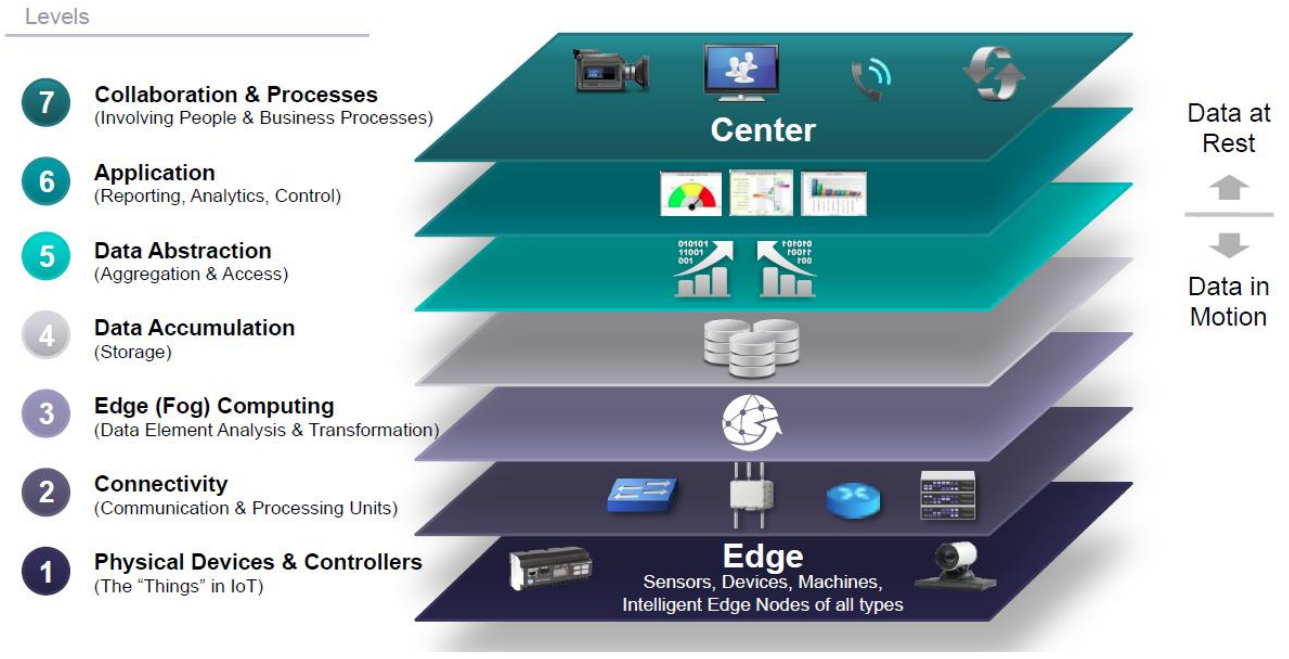

Source: IoTWF

Level 1: physical devices and controllers

The IoT Reference Model starts with Level 1: physical devices and controllers that might control multiple devices. These are the "things" in IoT, and they include a wide range of endpoint devices that send and receive information. Today, the list of devices is already extensive; devices are diverse, and there are no rules about size, location, form factor, or origin. Some devices will be the size of a silicon chip. Some will be as large as vehicles. The IoT must support the entire range.

Level 2: connectivity

Communications and connectivity are concentrated in Level 2. The most important function of Level 2 is reliable, timely information transmission. This includes transmissions:

- Between devices (Level 1) and the network
- Across networks (east-west)
- Between the network (Level 2) and low-level information processing occurring at Level 3

Traditional data communication networks have multiple functions, as evidenced by the International Organization for Standardization (ISO) 7-layer reference model. However, a complete IoT system contains many levels in addition to the communications network.

One objective of the IoT Reference Model is for communications and processing to be executed by existing networks. However, some legacy devices are not IP-enabled, which will require introducing communication gateways. Other devices will require proprietary controllers to serve the communication function. However, over time, standardization will increase. As Level 1 devices proliferate, the ways in which they interact with Level 2 connectivity equipment may change. Regardless of the details, Level 1 devices communicate through the IoT system by interacting with Level 2 connectivity equipment.

Level 3: edge (fog) computing

The functions of Level 3 are driven by the need to convert network data flows into information that is suitable for storage and higher level processing at Level 4 (data accumulation). This means that Level 3 activities focus on high-throughput data analysis and transformation, which may be "light": computationally undemanding and simple instructions or routines.

Given that data is usually submitted to the connectivity level (Level 2) networking equipment by devices in small units, Level 3 processing is performed on a packet-by-packet basis. This processing is limited, because there is only awareness of data units – not "sessions" or "transactions." Level 3 processing can encompass different functionality, including:

- Evaluation: Evaluating data for criteria as to whether it should be processed at a higher level
- Formatting: Reformatting data for consistent higher-level processing
- Expanding/decoding: Handling cryptic data with additional context (such as the origin)
- Distillation/reduction: Reducing and/or summarizing data to minimize the impact of data and traffic on the network and higher-level processing systems
- Assessment: Determining whether data represents a threshold or alert; this could include redirecting data to additional destinations

Level 4: data accumulation

Networking systems are built to reliably move data. The data is "in motion." Prior to Level 4, data is moving through the network at the rate and organization determined by the devices generating the data. The model is event driven. As defined earlier, Level 1 devices do not typically possess computing capabilities themselves. However, some computational activities could occur at Level 2, such as protocol translation or application of

network security policy. Additional compute tasks can be performed at Level 3, such as packet inspection or triggering an alert. Driving computational tasks as close to the edge of the IoT as possible, with heterogeneous systems distributed across multiple management domains represents an example of edge or fog computing. Edge or fog computing will be a distinguishing characteristic of the IoT.

However, most applications cannot, or do not need to, process data at network wire speed. Applications typically assume that data is “at rest” –or unchanging–in memory or on disk. At Level 4, Data Accumulation, data in motion is converted to data at rest.

As Level 4 captures data and puts it at rest, it is now usable by applications on a non-real-time basis. Applications access the data when necessary. In short, Level 4 converts event-based data to query-based processing. This is a crucial step in bridging the differences between the real-time networking world and the non-real-time application world.

Level 5: data abstraction

IoT systems will need to scale to a corporate – or even global – level and will require multiple storage systems to accommodate IoT device data and data from traditional enterprise ERP, HRMS, CRM, and other systems. The data abstraction functions of Level 5 are focused on rendering data and its storage in ways that enable the development of simpler, performance-enhanced applications. The data abstraction level must process many different things. These include:

- Reconciling multiple data formats from different sources
- Assuring consistent semantics of data across sources
- Confirming that data is complete for the higher-level application
- Consolidating data into one place (with ETL, ELT, or data replication) or providing access to multiple data stores through data virtualization
- Protecting data with appropriate authentication and authorization
- Normalizing or denormalizing and indexing data to provide fast application access

Level 6: application

Level 6 is the application level, where information interpretation occurs. Software at this level interacts with Level 5 and data at rest, so it does not have to operate at network speeds. The IoT Reference Model does not strictly define an application. Applications vary based on vertical markets, the nature of device data, and business needs. For example, some applications will focus on monitoring device data. Some will focus on controlling devices. Some will combine device and non-device data.

Level 7: collaboration and processes

The IoT system, and the information it creates, is of little value unless it yields action; this action can be automated (through actuators) or can occur due to data informing people and processes. Applications (Level 6) can therefore give people the right data, at the right time, so they can do the right thing. Often, the action needed requires more than one person. People must be able to communicate and collaborate, sometimes using the traditional Internet, to make the IoT useful. Communication and collaboration often requires multiple steps and it usually transcends multiple applications. This is why Level 7, represents a higher level than a single application.

Service Provider Architecture to Support IoT

To create the infrastructure to support IoT, we need to first assess and understand the wide range of IoT use cases and their associated requirements. Below, Etisalat discusses the characteristics of IoT use cases, presents the key requirements to support IoT and outlines a service provider network architecture to enable IoT.

IoT use cases

IoT will encompass many use cases, spanning many verticals, with the number and complexity of use cases steadily increasing as technology advances. Use cases range from those relying on large numbers of sensors transmitting relatively small amounts of data periodically (massive machine-type communications), such as smart city deployments, to use cases that require almost real-time response to high volumes of transmitted data in order to function (mission-critical communications), such as self-driving vehicles. The vast majority of use cases will sit somewhere along the spectrum between these extremes.

Figure 9: Different categories of use cases mapped to their most important communications requirements

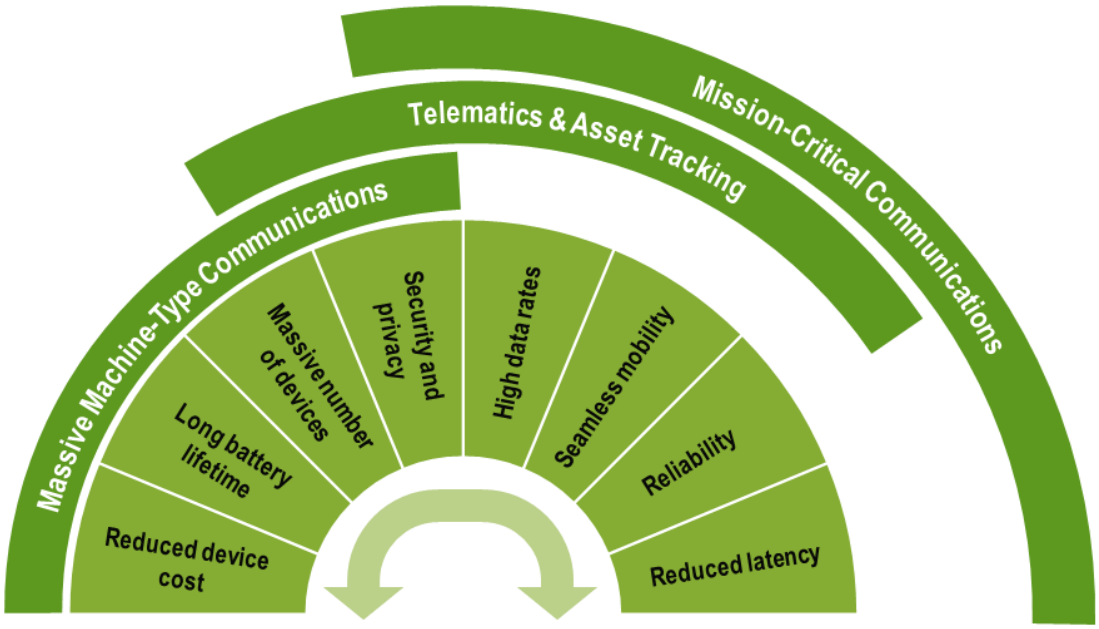

Source: Etisalat

IoT communications requirements

Etisalat recognises that the future infrastructure enabling the Internet of Things will need to support wide-ranging requirements due to the variety of potential use cases. Today's networks already address some of the criteria, particularly around security, capacity and coverage, however, some specific requirements from more forward-looking IoT use cases will present extra demands on the network that, to date, have not been relevant for mobile broadband and previous M2M applications.

The most significant requirements to support IoT are summarised below:

1. **Many more connections**
   As IoT continues to grow and the number of connected 'things' increases, the network will need to handle a surge in connected devices. Forecasts estimate the number of connected devices will exceed 28 billion by 2020[3], although not all of these devices will be connected directly to a wireless network. This increase in number of devices will, in some locations, translate into a dramatic rise in the density of devices within cell-sites.

2. **Low power usage**
   Traditional M2M SIM-based connections, using the current mobile network infrastructure, support IoT connections that typically have access to a reliable power source or limited lifetimes. They are unable to support connections that consume very low levels of power and can operate independently for long periods. In order to support simple, cheap devices, such as basic sensors, that can operate for years (indeed decades) on standard batteries, new technologies need to be deployed, and/or current networks upgraded. For many applications, changing device batteries in the field will not be practical or viable, so the battery will determine the effective life of the device; this could be ten to twenty years on a standard 2.5Ah battery. See the section, Access Technologies for Low-Power IoT Applications, for more discussion on this topic.

3. **High data rates**
   Some IoT applications will need to ensure data is transmitted at high rates in order to deliver full functionality. This is particularly important for IoT applications using video or video analytics. For example, within a smart city deployment, thousands of cameras may be installed across the city; when these cameras recognise an incident, they will stream the footage and large amounts of data will be stored and analysed in the cloud. The future network will need to be able to be handle these swathes of data going through the network.

4. **Deep coverage/penetration**
   Many IoT use cases need to cover "hard-to-reach" areas. Examples include: a smart parking sensor situated in a carpark deep underground; a medical sensor attached to a patient in a rural area with limited coverage; or a sensor on an off-shore wind-turbine. Ubiquitous coverage will also be important for moving devices, such as those used for fleet management or asset tracking, as the sensor will need to transmit data regularly as the object moves.

5. **Ultra-low latency**
   In the future, IoT will include mission-critical applications, whereby real-time action occurs based on in-bound network data. For example:

   - Semi-autonomous unmanned vehicles (drones, driverless cars or mobile robots) need to receive in-bound information in real time to make decisions on their trajectory. This will be complimentary to the on-board sensors and navigation systems and will improve performance as well as safety.
   - Augmented reality devices for field-force workers may need to interact in real-time with machinery to ensure worker safety. For these applications to work effectively, there should be close to zero latency.

6. **Advanced data analytics**
   For many IoT applications, collecting and reporting data is the pre-cursor to analysing the data to automate decision-making. This will include aggregating data from a variety of sources to trigger an actuator. For example, a smart city traffic management solution will collect data from cars, public transport, road sensors

[3] IDC: Worldwide and Regional Internet of Things (IoT) 2014–2020 Forecast

and CCTV video footage. Real-time analytics can then control traffic signals and notify drivers of any changes to their route in order to reduce congestion.

7. **Security & data management**
   As the physical world becomes more-closely linked to the virtual world, the potential for a security breach increases. Security and privacy issues arise around control of things/devices and access to confidential information, whether it be a medical devices and data, an industrial machine or a car. The network has a key role in providing defence measures against attacks, particularly if the device is connected end-to-end by the network operator.

8. **Resilience & local autonomy**
   As the world becomes more dependent on IoT applications, the network will be relied upon to provide very high-reliability connectivity, which is resilient to outages. However, some customers will have added requirements for localised autonomy (for example in the event of a major disaster).

Network architecture to support IoT

In order to support these various, and in some cases seemingly contradictory sets of requirements, Etisalat recognises that the network architecture will need to evolve.

Building on a virtualised infrastructure

The future network architecture to support IoT builds on key aspects outlined in a previous whitepaper, Etisalat – 2020 Landscape[4]. This whitepaper sets out the vision of the virtualised network in 2020, embracing Network Functions Virtualisation (NFV) and Software-Defined Networking (SDN) technologies. Many access and core network functions will be virtualised, for example load balancing, firewalls, routing and session border controls. These virtual network functions will be able to use standardised commercial off-the-shelf (COTS) hardware to run network functions that are no longer tied to proprietary, dedicated hardware. This will create a more flexible and programmable network.

IoT applications can also leverage the physical infrastructure running virtual network functions (VNFs). This infrastructure can potentially provide computation for applications closer to the edge of the network, reducing latency, demands on backhaul, providing localised autonomy and added security (for example, against Distributed-Denial-of-Service attacks by thousands of compromised IoT devices). Etisalat believes that the traditional separation of network and applications will begin to blur. Etisalat anticipates that service providers will address this through edge (fog) computing.

Network architecture for IoT

The diagram below presents the service provider network architecture to support IoT. Essentially, this consists of two main elements:

1. ***Agile Infrastructure*** (Agile Access, Agile Metro, Agile Core). These are essentially “transport” components of the network.

2. ***Cloud Core and Cloud Edge (Fog)***: These include the “virtualised” components of the network across the data plane (VNFs) and the control plane (SDN), as well as services. Since IoT applications will also benefit from cloud & edge (fog) capabilities (for application computation), service providers will seek to pool resources for both network and IoT functionality. In the short term, and at the very minimum, these

[4] Etisalat – 2020 Landscape

common resources will be physical (i.e. power & space), over the longer term, they will converge to a single virtualised fabric.

Figure 10: Service provider network architecture to support IoT

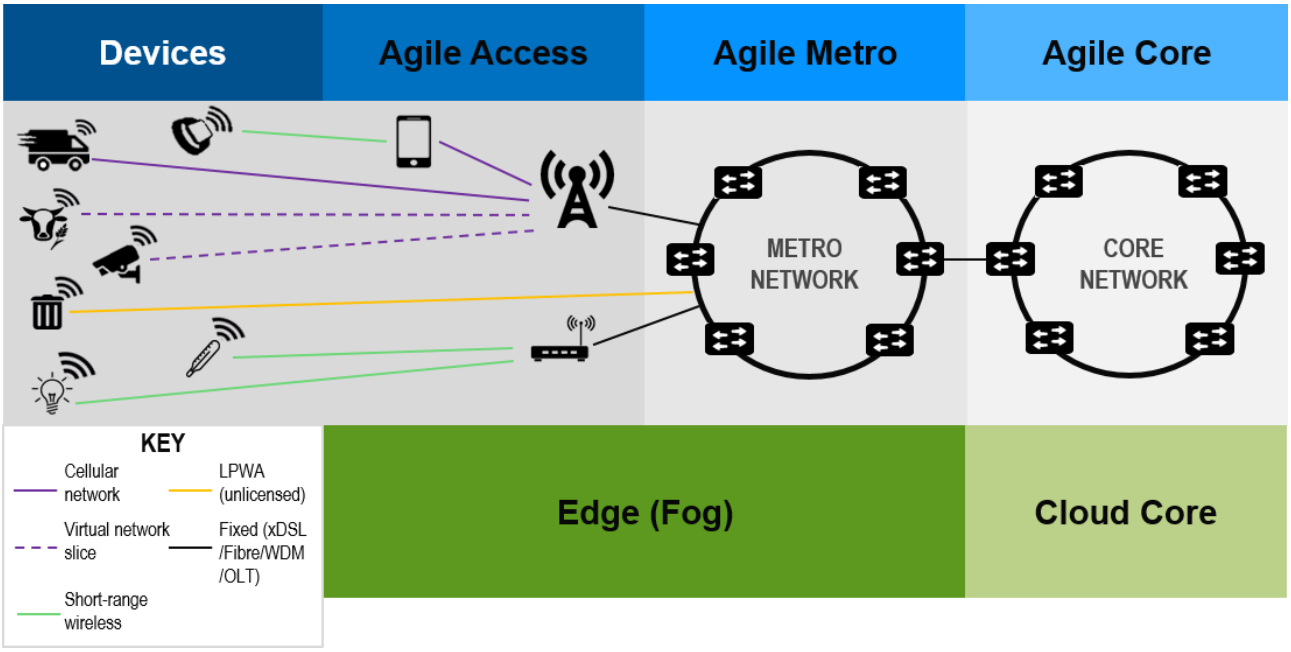

Source: Etisalat

Devices

This layer includes the various IoT devices (sensors, machines, actuators) connected to the network. Some IoT devices will connect to the network directly via fixed or wireless connections, whereas others may connect to the telecoms network through capillary networks. Capillary networks include local area networks (for example in the home or offices) connecting devices in a confined area via short-range links to a gateway or hub, which then connects to the access network via a fixed or mobile broadband connection.

Agile Access (access layer)

The Agile Access layer of the network architecture represents the access network, where devices connect to the service provider's network. To address the variety of IoT applications, this will include a combination of wired (e.g. Ethernet) and wireless (radio) connections. Etisalat anticipates that future radio access networks will include a range of licensed spectrum connections; legacy 2G / 3G / 4G networks, including low-power variations and future 5G networks as operator networks evolve towards 5G.

The access layer may also include other networks outside the traditional telecoms network infrastructure. IoT applications may use unlicensed spectrum network technologies, including Low-Power Wide Area (LPWA) technologies. The diversity of use cases (and requirements) will mean that service providers will need to consider various deployment approaches as "legacy" cellular may not offer the most cost-effective, or energy-efficient connections for a low-cost battery-powered device. The different access technologies addressing this particular IoT segment is discussed in more detail in the Access Technologies for Low-Power IoT Applications section.

Agile Metro (aggregation layer)

The Agile Metro (aggregation layer) connects the access layer with the core layer and seeks to aggregate the various access networks in order to transport data to the core. Agile Metro is envisaged to provide a metropolitan area network that will replace the need for multiple separate fixed links (for enterprise customers and backhaul connectivity for mobile networks, etc.) with a single flexible network that can aggregate varied traffic requirements. Etisalat predicts that the future Agile Metro will need to support much higher bandwidth requirements; bandwidth at this layer will likely rise from 1Gbps-10Gbps to the range of 100Gbps-400Gbps, some of this is due to the growth of IoT.

Agile Core (core layer)

Agile Core refers to the optimised core network layer. Etisalat views NFV and SDN as important technologies that will enable a more flexible core network that can efficiently support IoT use cases. A combination of virtualised functions and SDN will enable the 'slicing' of multiple mobile core instances for IoT use cases rather than a single 'one-size-fits-all' core network.

The growth of IoT will contribute to the increase in the risk of potential congestion at the core layer. Whereas traditional broadband largely consists of IP streaming and downlink data, data in the Internet of Things will use uplink communication as data is captured by devices and transmitted into the network. The future bandwidth requirement at this layer is forecast to reach in the range of 400Gbps-1Tbps, compared to that of traditional core networks which ranged between 10Gbps-multiples of 10s of Gbps.

Cloud Core

The Cloud Core relates to cloud computing resources for core networking (e.g. SDN controllers) and other centralised networking, storage and compute (e.g. orchestration, OSS/BSS). Service Providers' Cloud Core will, over time, extend to support IoT applications' centralised compute and storage needs.

Edge (fog)

Edge (fog) computing extends cloud computing all the way to the edge of the network. The Edge (fog) Cloud resides anywhere between the Agile Metro and the device. The edge (fog) computing architecture optimises the distribution of network, compute & storage to maximise network efficiency and improve application performance, where appropriate. This is discussed in more detail in the Edge (fog) computing section in this report.

Illustrating the network architecture through use cases

Etisalat's vision of the network architecture to support IoT takes into consideration the many use cases IoT presents. Using three distinct IoT use cases, the features of the architecture are brought to life and applied in practice.

This mapping to the network architecture illustrates one potential approach to deliver the use cases. However, there will often be more than one way to support the use cases; for example, some use cases could use capillary networks or connect devices directly to the cellular network.

1. Smart Farms

Smart Farms are able to monitor and analyse the soil / environmental conditions as well as third party data (e.g. weather information) to enable automated / smarter decision making around engaging certain farm systems (e.g. irrigation system, crop treatment) or other intervention (harvest).

Figure 11: Implementing Smart Farms over service provider network

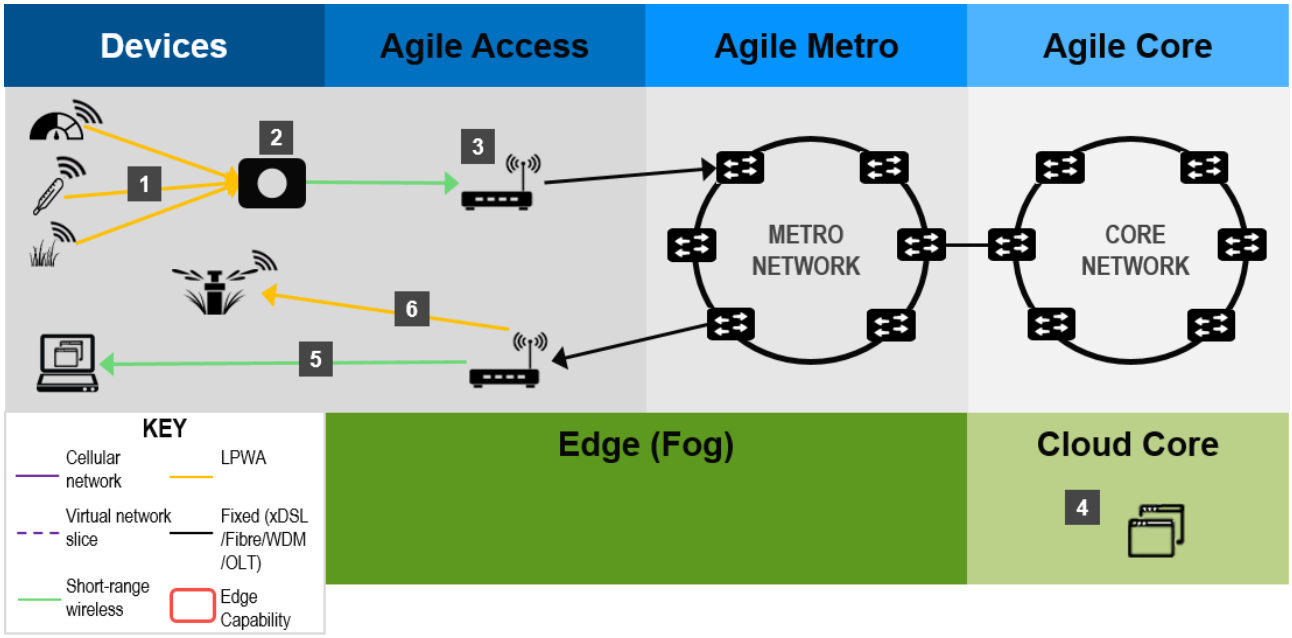

Source: Etisalat

[1] Sensors on the farm send data (such as soil and environmental conditions) to an IoT gateway.

[2] The IoT gateway aggregates the data from these various sources to be sent to the router.

[3] The aggregated data is sent periodically (e.g. once per hour) to the smart farm application platform in the cloud.

[4] The application platform provides further aggregation and analysis, combining information from sensors with 3rd party data (e.g. weather forecasts).

[5] This information is then displayed to the farmer on a web platform.

[6] Additionally, the application platform uses data analysis to automate decision-making to engage certain farm systems, for example automatically turning on the farm's irrigation system.

2. In-Transport CCTV

Increased security and response to incidents on public transportation through automated analysis of CCTV footage. Edge (fog) computing can process CCTV footage from public transport close to the edge of the network to identify instances of disturbance, risk, crime etc. If an incident has occurred, automated decision-making can rapidly alert the relevant emergency services, with the footage of the incident streaming centrally. If no incident has occurred, the video footage will not be transmitted across the network, reducing backhaul.

Figure 12: Implementing In-Transport CCTV over service provider network

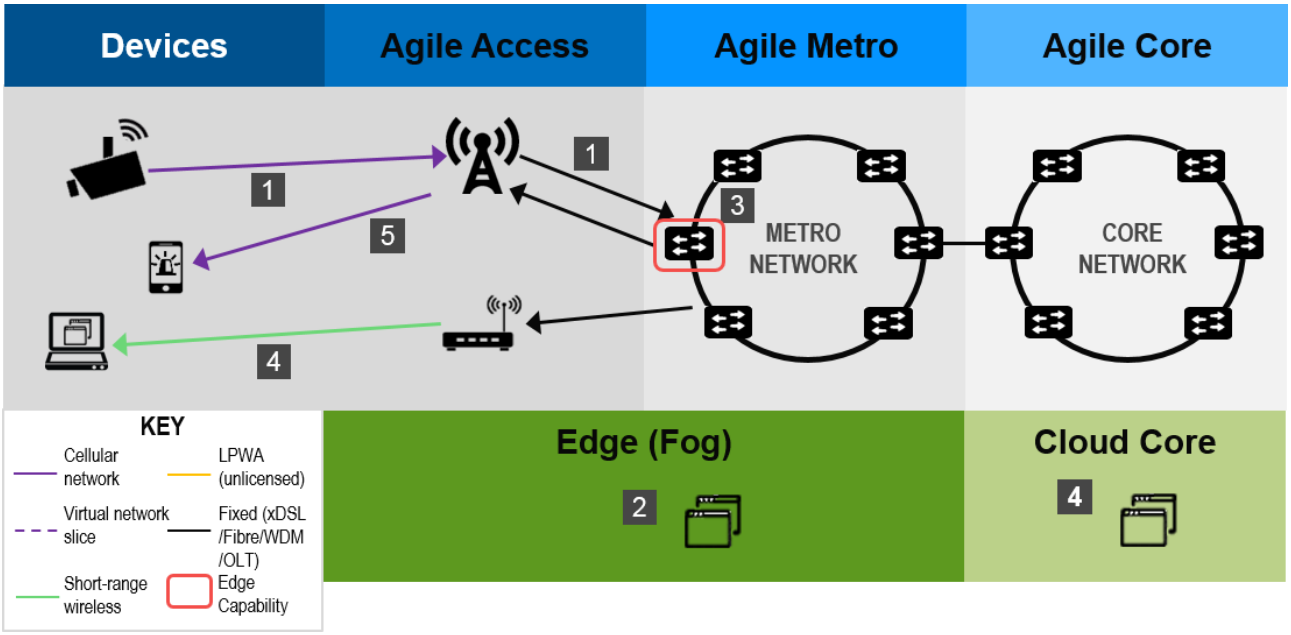

Source: Etisalat

[1] Video camera located on public transport sends live stream to network.

[2] Nearest available Edge Node runs event-based video analytics in real time, and can identify risk, disturbance, crime etc. Video data is not continuously transmitted across network, but Edge Node caches last ten minutes of video.

[3] When event (road accident, disturbance, crime etc) identified, Edge node begins live stream to cloud, and sends video of event from cache.

[4] Video can be accessed by local authorities via a web platform.

[5] Depending on outcome of event-based video analytics (e.g. serious road accident), Edge Node can send alert directly to local authorities.

3. Augmented Reality for Field-Force Safety

Data about field-force workers, captured by Augmented Reality (AR) devices and data on the environment (e.g. environment conditions, machine/equipment data) can be processed locally to ensure rapid response/action if a dangerous situation is occurring, ensuring greater field-force worker safety. For example, if the sensors detect high levels of heat discharge in the area of the field worker, machine process speeds are reduced or non-critical devices shut down.

Figure 13: Implementing AR for Field-Force Safety over service provider network

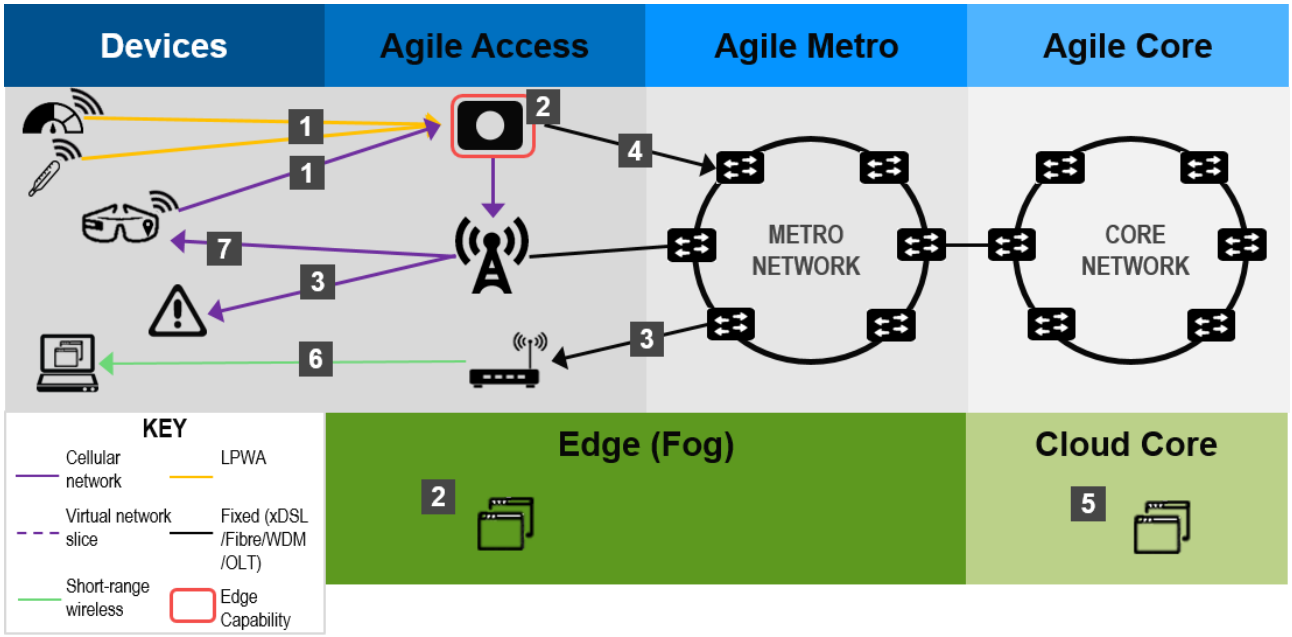

Source: Etisalat

[1] In-field sensors transmit data on equipment status and the field environment and the augmented reality glasses send data on the field worker, for example on the field worker's location.

[2] Edge computing at the IoT gateway will be used for analysis and processing closer to the edge to detect any changes in the environment and/or equipment which may pose a danger to the worker's safety.

[3] In this case, this will then trigger any alerts or actions. For example, if the sensors detect the environment is unsafe (e.g. environment is too hot due to machine over-use) an automatic action will be taken (e.g. machine will shut-down or reduce output) to protect worker safety.

[4] The edge node will also aggregate data which will be sent to the application platform in the cloud.

[5] The application platform performs analysis and further aggregation of the data.

[6] This aggregated data will then be displayed on the central office web platform (e.g. for field-force monitoring).

[7] Some data will also be displayed on the augmented reality glasses to keep the worker informed of the status of the field environment.

Note: Augmented reality may require additional network capabilities (for example edge processing of video) to ensure low latency rendering.

New Technologies Enabling IoT

The IoT ecosystem will be complex; it will not depend solely on one type of technology, rather many technologies will interact to enable a wide variety of IoT applications. New innovations on the horizon, across multiple technologies, will provide additional capabilities that will make the potential of IoT possible.

Key service provider technological innovations include:

- NFV & SDN
- Edge (Fog) Computing
- Network, Compute & Storage Orchestration
- Access Technologies for Low-Power IoT Applications
- 5G

These technologies will create more flexible network infrastructure, better able to meet the wide-ranging requirements of IoT. Etisalat anticipates that the combination of these new technologies will allow service providers to enable IoT.

NFV & SDN

Etisalat believes 'NFV & SDN driven network transformation' is one of the most disruptive technologies for service providers, eventually leading to the transformation of all key aspects of networks and operations. This transformation towards virtualised, programmable networks will also deliver many of the capabilities required to support IoT services, including enabling network slicing.

NFV involves the virtualisation of key network functions such as load balancing, firewalls, routing and session border controls; the creation of these virtual network functions (VNFs) enables the use of more standardised hardware as network functions are no-longer tied to the physical infrastructure. Moving from proprietary, dedicated hardware to more-flexible, multi-purpose hardware will create a more cost-effective and programmable network.

Related to NFV, **SDN** will also play a key role in the realisation of IoT. SDN refers to the physical separation of the network control plane, the system determining where traffic is sent, from the network forwarding or data plane, the system that forwards traffic to its destination. This separation enables much greater network control, as the network can be directly programmable, with the underlying infrastructure abstracted from applications and network services.

Together these new network technologies will provide key capabilities that will enable IoT.

As discussed earlier, IoT will have wide-ranging network requirements due to the variety of potential use cases, with some use cases only requiring kilobits of sensor data and others relying on real-time streaming video. NFV & SDN provide the capabilities to manage the full range of IoT use cases over one network infrastructure. Network resource can be allocated flexibly and in real-time, accommodating use cases that have high bandwidth requirements whilst simultaneously providing resource for low-bandwidth use cases.

NFV & SDN also create the potential for service providers to prioritise certain IoT applications over others; for example, IoT services that are 'mission-critical', (e.g. autonomous vehicles, advanced robotics) where there are major implications from a failure or latency in the network, could be prioritised over use cases where data is not required in real-time.

Additionally, Etisalat anticipates that IoT will lead to many more connected devices than exist today. NFV & SDN will allow service providers to readily provision and serve these new devices; rather than deploying physical hardware into the network to provision and manage these devices, service providers will be able to automatically implement VNFs using software, making it faster and more cost-efficient to deal with the large number of connections.

Furthermore, as opposed to today's networks, not all network functions will be required for all IoT services. For example, most IoT use cases will not require many of the networking functions which are necessary to support voice connections. Also, since many IoT connections will be static and will only transmit limited amounts of information periodically, they will not require complex functions that support mobility (e.g. for seamless handover when travelling at speed between cells).

NFV & SDN begin to change the network cost structure for a service provider, reducing capex and opex as hardware is more standardised, virtualised and shareable through many re-programmable functions. The speed and cost of innovation will also decline dramatically. This will allow networks to support use cases that would previously be unviable to become economically possible. For example, use cases that would previously have required the deployment of physical hardware whilst only offering the potential for marginal or un-proven revenues would not be considered. This change in the cost model and the automation in the network will allow service providers to address significantly more use cases. Critically, it will also allow operators to quickly and inexpensively "retire" functions that are no longer required.

Network slicing

A key capability of NFV & SDN is the ability for the network to cater to the different use case requirements via network slicing, running multiple logical network instances over the same network infrastructure. By using the same physical infrastructure, slicing enables extensive resource re-use, including radio resources, to optimise the use of the network environment.

Figure 14: Network slicing uses the capabilities of SDN and NFV for IoT use cases

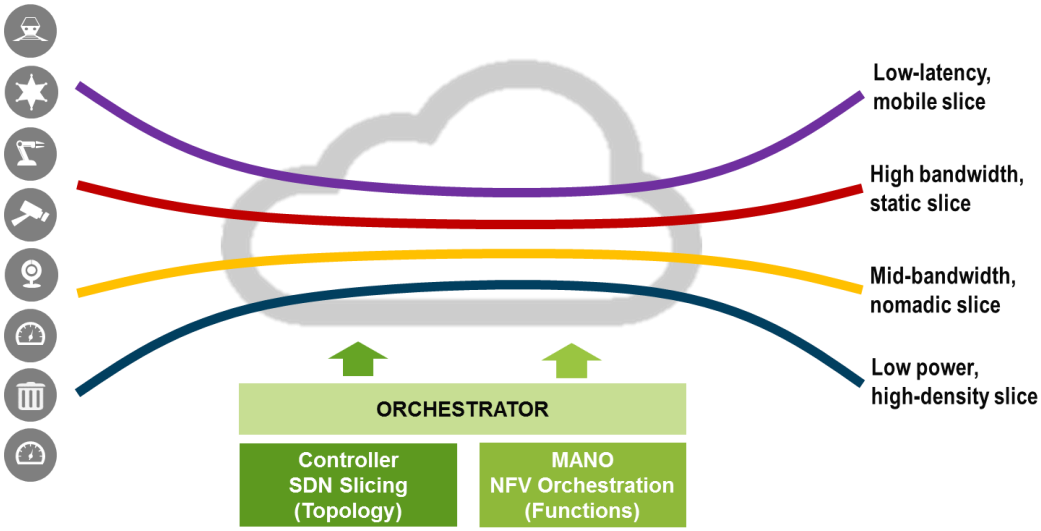

Source: Etisalat

The technology behind network slicing leverages the capabilities of NFV and SDN. Each logical network slice will be associated with a particular connection type, including a specific collection of network functions and specific radio access technology settings that are combined together for the specific application type. SDN is used to make the network more programmable by configuring slicing and network service chains using the SDN controller. Using NFV, different core network functions can be allocated to each slice to avoid the use of unnecessary functionality.

Whilst network slicing will be possible on 4G networks that have implemented NFV & SDN, the deployment of 5G (underpinned by NFV & SDN) will take full advantage of these new capabilities. Applying this concept to IoT use cases, we can examine a number of examples of network slices:

- Network slicing to support massive machine-type communications sensors/devices
- Network slicing to support mission-critical IoT services
- Network slicing to support IoT video sensor gateway applications

Each of these use cases have different network requirements and therefore require different network functionality and resource.

For example, for massive machine-type communications sensor networks, the key connectivity requirements include dense coverage and low-power, low-frequency data transmission. This would imply the need for a light duty core without mobility management features in the core network functions. Alternatively, the mission-critical IoT services slice would need the core functions to run at the edge to minimise transition delay and ensure high-reliability and low-latency. Finally, IoT video sensor gateway applications require analytics at the edge to filter the video stream, therefore only a light duty core is necessary but edge computing functionality is required.

Figure 15 below highlights how 5G network slices will support different IoT use cases:

Figure 15: 5G network slices implemented on the same infrastructure

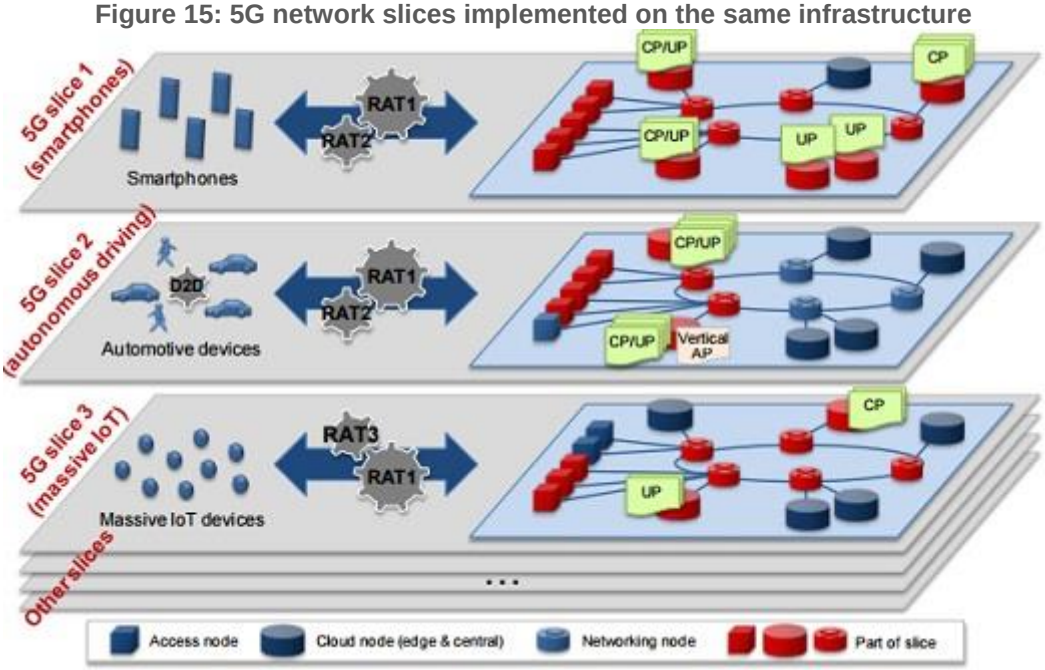

Source: NGMN

Edge (fog) computing

Edge (fog) computing will provide new, differentiated capabilities, allowing service providers to offer cloud computing and IT capabilities throughout their network, closer to the end-user or device. This has the potential to enable a host of new IoT use cases, some that require ultra-low latency and others that avail of real-time network information.

Edge (fog) computing is a natural extension of cloud computing. As more objects or 'things' become connected, more compute capability will be required. Edge (fog) will provide compute closer to these 'things', improving application performance as well as ensuring efficient management of the network.

Figure 16: Edge (fog) computing occurs at aggregation and local level

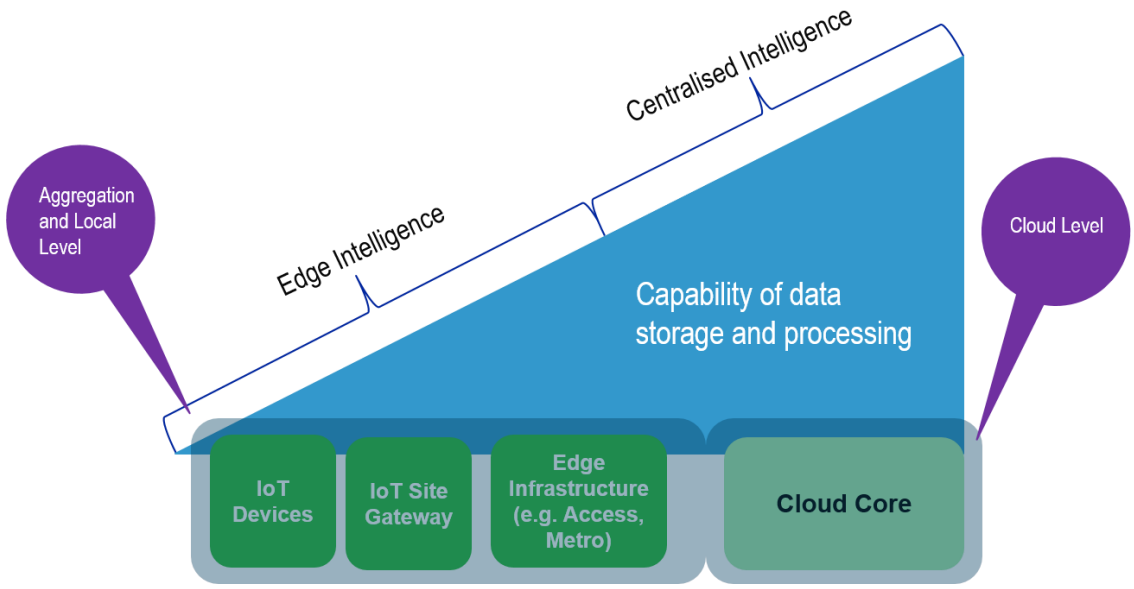

Source: Etisalat

Edge (fog) computing capabilities will reside much closer to the end-user device than when they occur in the Cloud Core (or data centre). Potential locations for edge computing include an IoT gateway or a point on the service provider's network (e.g. at a base station). This will enable information to be processed faster, reducing latency for certain applications and providing extra reliability, allowing applications to run in a standalone manner (i.e. without connectivity to the core).

Whilst this edge (fog) computing capability will be available to support IoT, Etisalat anticipates that the majority of data processing and storage will still occur in the Cloud Core. Edge (fog) computing capabilities will only be leveraged when it makes sense from a commercial or network efficiency perspective. This means that any given IoT application may be partly running some (light) routines in the edge (fog), as well as in the Cloud Core. Managing the edge (fog) infrastructure/capabilities is discussed in more detail in the next section.

Service providers are well-positioned to develop these capabilities and provide or enable IoT services that leverage edge (fog) computing due to their inherently distributed infrastructure. Etisalat anticipates that edge (fog) computing will be a key future competitive advantage for service providers.

Etisalat recognises three main benefits from Edge (fog) computing with regard to IoT:

- **Network Efficiency:** Edge (fog) computing capabilities will make the transmission of data across the network smarter and more efficient. The amount of data the network will be required to transmit will increase significantly over the next 5-10 years, creating potential network capacity problems. This increase in data will be more pronounced with the rise of millions more IoT sensors and devices continually capturing and transmitting data.

  Edge (fog) computing will bring analytical capabilities closer to the edge of the network; service providers will be able to more efficiently aggregate information as well as filter information that does not need to be sent back to the core. This will reduce overall network traffic and ensure better network performance for all users.

- **Creating Resilient Infrastructure:** Edge (fog) computing can increase the resiliency of local networks and operations. Localised computing can allow a system to operate even if outages have brought down other parts of the network or operations.

- **Enabling Low-Latency Applications:** Bringing computing closer to the edge of the network will support applications that have stringent latency requirements. The current model of transmitting information to the core and back may result in lag times that will impact mission critical IoT applications. Edge (fog) computing will significantly reduce latency, enabling potentially new IoT use cases that depend upon the rapid computation of information.

  Etisalat expects applications that will utilise this low-latency capability will be more futuristic (as the infrastructure is not currently in place today).

Edge (fog) computing will also provide additional benefits outside of IoT. For example, next generation immersive content (i.e. virtual reality, augmented reality) will also require lower-latency compute capabilities as the interactive environment will need to change in real-time based on the user's movements.

Etisalat will deploy edge (fog) computing infrastructure where it makes commercial sense. This could be at the customer site, or close to the edge at a mobile base station, at a cell aggregation site or at the Radio Network controller site. The deployment of edge computing infrastructure will depend on commercial viability on an opco-by-opco basis.

Figure 17: A comparison of potential edge (fog) computing deployment scenarios to cloud computing

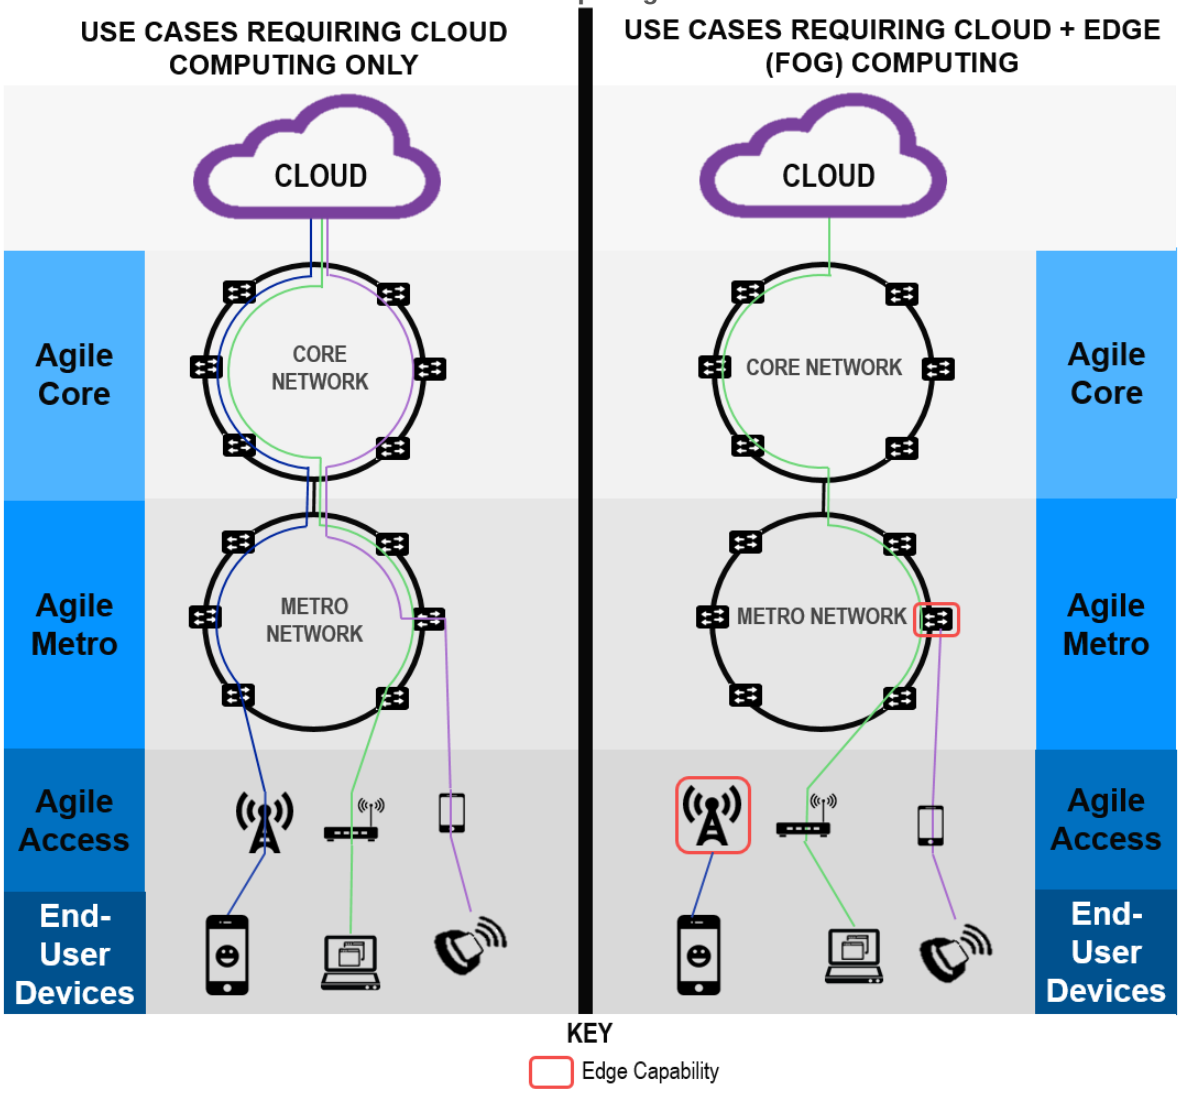

Source: Etisalat

Figure 17 outlines potential deployment scenarios for a service provider. The 'cloud computing only' use cases, demonstrate how all the data from the devices or things are transmitted all the way through the network to the cloud. Processing occurs in the cloud and the corresponding information is sent back to the devices or things.

The Cloud + Edge use cases illustrate the potentially shorter cycle, leveraging edge (fog) computing capabilities. In one case, the mobile base station has some edge (fog) computing capabilities and so the information is not transmitted all the way across the network. Similarly, it also presents a use case where the edge (fog) computing capability is further within the network, at the Agile Metro layer. However, these are simplified examples. In practice, the same application may run some computation in the edge (fog) and some in the cloud, with the mix determined by performance, resource availability and cost optimisation.

Figure 18: Potential deployment locations for edge (fog) computing

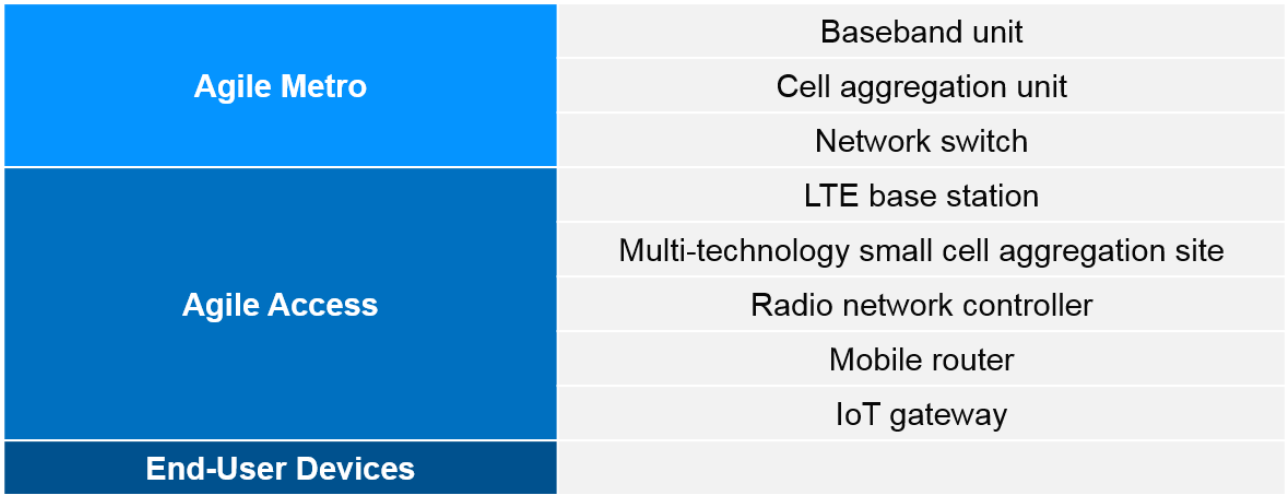

Source: Etisalat

For Etisalat, the use of edge (fog) computing capability will depend on when it makes sense practically. This means selecting to distribute computing functionality in cases when it is needed and to centralise computing when there is no requirement for edge computing. Etisalat will also open-up its edge (fog) computing capabilities to third parties, allowing others to utilise this infrastructure and create low-latency applications.

Network, compute & storage orchestration

Orchestration is central to a service provider's role supporting IoT; service providers will have to manage networks that support different types of devices with differing requirements. NFV & SDN will create the flexibility and automation required to manage the network dynamically.

The below architecture in Figure 19 illustrates how the additional functionality from NFV & SDN will support IoT, through common orchestration across networking, cloud and edge (fog) capabilities.

Figure 19: Adapted IoTWF Service Provider IoT Reference Architecture

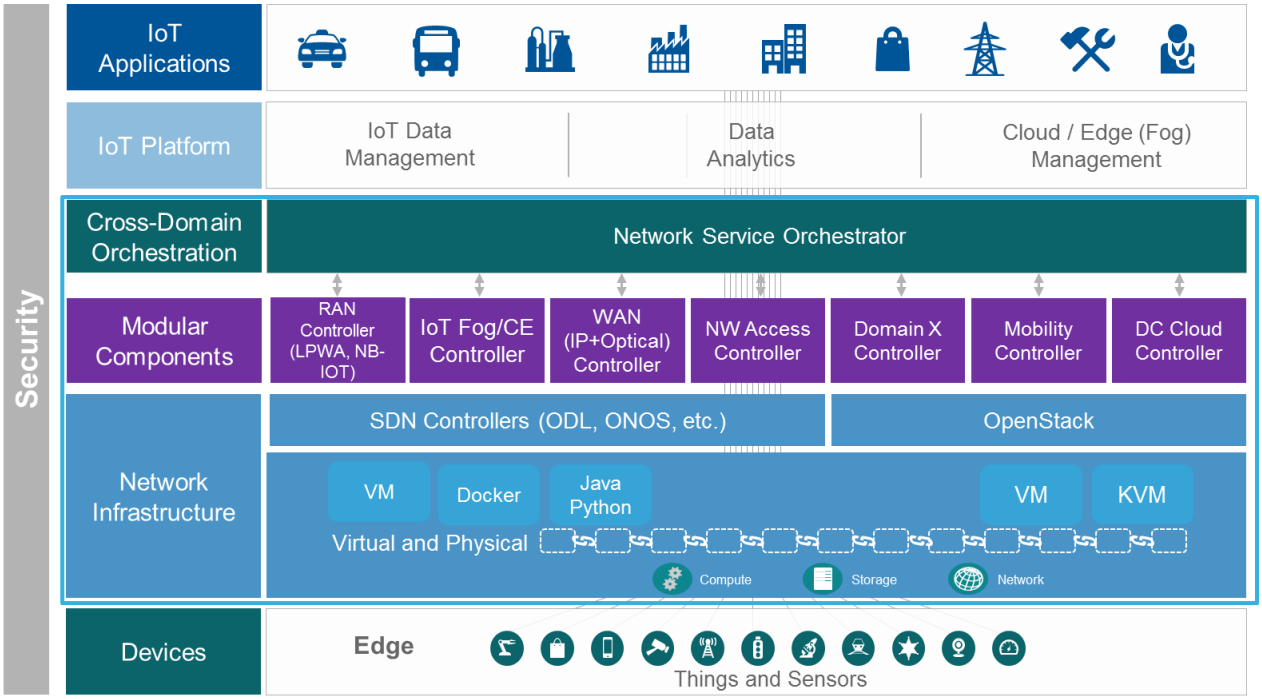

Source: IoTWF; Etisalat

Figure 19 shows the logical network architecture to support IoT. The Network Service Orchestrator provides end-to-end orchestration spanning multiple domains across the network, enabling:

- Service instantiation: The creation of a virtual instance of the service
- Service chaining: A chain of connected virtual network services
- Service monitoring: Ensuring the service/services are performing effectively
- Service scaling: The ability to create resource to support multiple services to scale

Network Service Orchestration will enable service providers to fulfil every aspect of a service, end to end, across the entire services stack. This includes managing NFV & SDN as well as the traditional physical network.

The European Telecommunications Standards Institute (ETSI) have published a reference architecture for NFV management and orchestration (MANO). Etisalat endorses ETSI's framework for NFV deployment. The architecture presented in Figure 20 below maps IoT to the ETSI MANO reference architecture.

Figure 20: Mapping IoT to ETSI MANO Reference Architecture

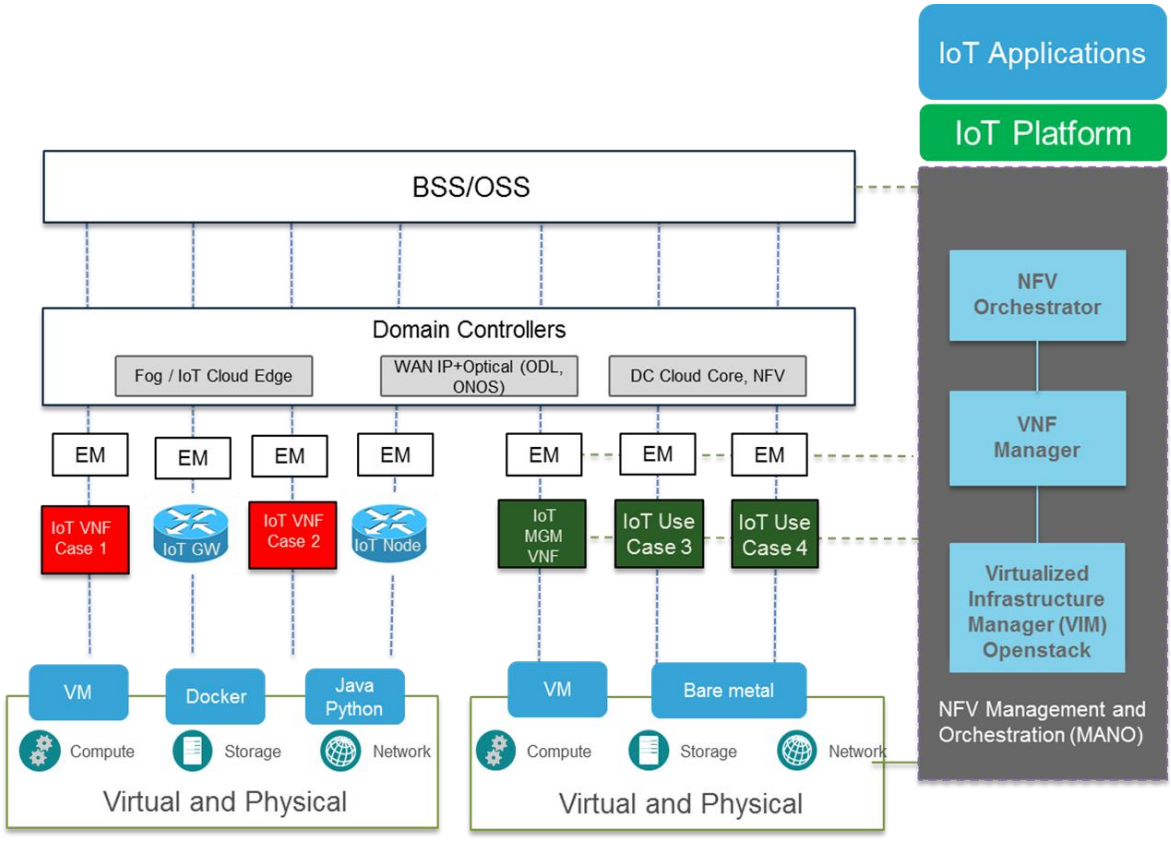

Source: ETSI; Etisalat

The NFV MANO (NFV Orchestrator, VNF Manager and Virtualized Infrastructure Manager) can control and manage IoT VNFs and services, in a similar way to other network services.

- Virtualised Infrastructure Manager (VIM): Manages the underlying NFVI resource i.e. network, compute, storage). The VIM will create and tear-down virtual machines to support network and IoT services.
- VNF Manager: Responsible for VNF lifecycle management (i.e. creating, managing and terminating VNF instances).
- The NFV Orchestrator is responsible for on-boarding and managing new network services and VNFs, including potential IoT VNFs. This ensures that network and IoT services can be delivered end-to-end.

Managing edge (fog) infrastructure

Managing the edge (fog) infrastructure will prove to be one of the most essential components in order to fully reap the benefits of edge (fog) computing. Figure 19 highlights the role of the Cloud / Edge (Fog) Management platform to ensure edge (fog) computing works in practice by setting the policy and configuration rules for forwarding traffic to/from edge (fog) applications.

The platform would be located at the edge (fog) nodes, at the edge of the network, where data is ingested from the IoT devices. It then directs different types of data to the optimal place for analysis, based on three important

factors: performance, resource availability and cost optimisation. As resources are constrained and the costs of edge (fog) computing are higher than cloud computing, only the most time-sensitive data (with high performance requirements) will be analysed at the edge of the network; data that is less time sensitive is either sent to aggregation nodes or to the cloud for historical analysis, big data analytics and/or long-term storage.

Orchestration of the edge (fog) will be administered by the Fog/CE controller (see Figure 19). This will enforce the policy, set by the Cloud / Edge (Fog) Management platform, to determine the split between centralised and distributed workloads, which will allow the operator to optimise resources against performance dynamically. Edge (fog) computing will primarily be done at the packet level, not the session level, therefore the overall application runs both on the edge and the cloud; the routines which will be distributed will be determined by the Fog/CE controller.

A key component of management and orchestration of edge (fog) computing is that it is done centrally and co-ordinated across nodes and devices. The model used for this has adapted the ETSI MANO model (see Figure 20) to ensure operational efficiency, as, over time, the edge (fog) hosting location will be able to converge with broader NFV and cloud services.

Access technologies for low-power IoT applications

New access technologies are a key enabler for the growth of the Internet of Things. A large proportion of IoT use cases, namely massive machine-type communications (applications with a huge number of devices) are not well served by current cellular networks as they are not a cost-effective option for connecting devices which transmit minimal amounts of data, often small data messages which may only be sent once a day or once every hour.

Figure 21: Low-power applications are a significant subset of total IoT applications

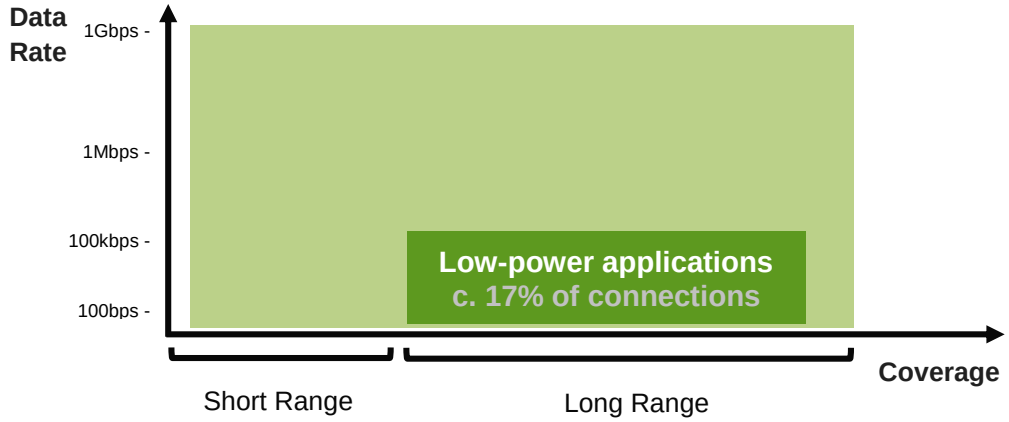

Source: Etisalat

These low-power IoT applications have specific network requirements, which differentiate them from other types of IoT use cases:

- **Low energy**
  Devices are likely to be powered in ways that will constrain the amount of energy that can be consumed for connectivity. For example, some devices may be battery-powered and placed in remote areas. Replacing or charging batteries is therefore not practical nor cost effective, so batteries need to last for long periods of time, sometimes years or even decades to be cost-effective. Alternatively, some devices will not even use batteries, but use energy harvesting capabilities close to the device, which are only able to provide a small amount of power. An example of this would include using solar energy to power smart farm sensors which measure light, humidity, pressure, etc.

- **Low device cost**
  Massive machine-type communications will rely on many (dumb) sensors to extract data. In order to build a viable business case for these applications, the cost-per-device needs to be at a minimum, with connectivity costs reaching a few dollars per year.

- **Small data volumes**
  Many of these devices will send relatively small data messages, for example a status indicator for temperature or the location of a device. However, the frequency of messages will vary by use case. In general, the data volume from low-power devices is expected to reach up to tens of kilobytes per day.

- **Wide area**
  The types of use cases these networks will aim to serve, for example industrial, transport and logistics, etc, may require connectivity over a wide area (e.g. tens of kilometres). Some use cases will have specific coverage requirements, such as national / regional / global coverage or deep indoor connectivity.

- **Massive numbers (of devices)**
  The network will need to be able to support a massive number of devices, due to the number of sensors involved in massive machine-type communications, although density of devices will vary from cell to cell.

There have been many recent developments, harnessing both unlicensed and licensed spectrum, to create network solutions that satisfy these wide-ranging access requirements.

Figure 22: LPWA unlicensed technologies comparison

Technology	Range	Uplink Data Rate	Downlink Data Rate
LoRaWAN	2-5km (urban) 15km (rural)	300bps – 50kbps (EU)	300bps – 50kbps (EU)
Sigfox	3-10km (urban) 30-50km (rural)	100bps (140 messages/day)	Max. 4 messages of 8 bytes/day
Weightless-N	3km (urban)	100bps	No downlink
Ingenu	>500 km Line-of-Sight	AP aggregates to 624 kbps per Sector (Assumes 8 channel Access Point)	AP aggregates to 156 kbps per Sector (Assumes 8 channel Access Point)

Source: EDN

3GPP, in Release 13 in June 2016, completed the standardisation of NB-IoT, the narrowband radio technology utilising licensed spectrum; the NB-IoT ecosystem is developing, with commercialisation expected by Q2 2017. Licensed network options offer an alternative to the unlicensed LPWA options and provide added benefits, including:

- **Bi-directional communication:** Most unlicensed LPWA options are not designed to handle downlink data transfer and are more suited for uplink, however some are in the process of developing capability for both (e.g. LoRa.)
- **Higher speeds:** This will vary based on the standards agreed, but cellular-based technologies are likely to be able to handle up to 1Mbps speeds.
- **Ensured quality of service and reliability:** Although some use cases are tolerant of network delays and glitches, other use cases require a level of quality assurance.
- **Security:** Controlling the end-to-end connectivity allows operators to ensure secure connections on their networks.
- **Scalability:** Cellular networks already cover 90% of the world's population and a cellular-based solution will provide devices with roaming capability, allowing for fast, global uptake of low-power solutions.
- **Ease of implementation:** A cellular-based solution will not require the roll-out of new network infrastructure and, in the case of NB-IoT, may be deployed as a simple software update to existing LTE infrastructure.
- **One-stop solution:** Cellular connectivity can provide for a multitude of IoT applications, not solely a sub-set of low-power use cases.

The current 3GPP low-power solutions comprise of three standards. Etisalat is an active member of the GSMA Mobile IoT Initiative, which supports the 3GPP in standardising these technologies.

- **Extended Coverage – GSM (EC-GSM):** A low power evolution of GSM, based on eGPRS. It is most suitable for markets / areas (rural) with 2G networks where LTE is not currently deployed. The optimisation of EC-GSM can be deployed as a software upgrade to existing GSM networks. However, Etisalat recognises there is the potential future problem as operators start to switch off 2G networks. Therefore, the suitability of this solution will depend on the market.

- **LTE CAT-M1 (also known as LTE-M):** An evolution of LTE, which works within the normal infrastructure of LTE networks via a software upload. It is aimed for speeds up to 1Mbit/s rather than slower, low-power alternatives. It holds a larger slice of the spectrum compared to LTE Cat-M2, but will be better at supporting mobility.

- **LTE CAT-M2 (also known as NB-IoT):** Narrowband radio technology that can be deployed "in-band" in licensed LTE spectrum using a software plugin, or in the LTE carrier's guardband or standalone in dedicated spectrum. NB-IoT uses new physical layer signals and challenges to enhance extended coverage and reduce device complexity.

These three standards satisfy different sets of requirements so will each be suited to a particular set of use cases, depending on the criteria of the use case.

Figure 23: 3GPP Low-power technologies' characteristics

Technology	Cost	Range	Power	Data Rate
EC-GSM	Low	High	Low-Medium	Low
LTE Cat-M1	Medium	High (slightly lower)	Medium	Medium
LTE Cat-M2/NB-IoT	Low-Medium	High	Low-Medium	Low-Medium

Source: STL Partners

The landscape for low-power network technologies is diverse, reflecting the need to address a variety of use cases. Etisalat maintains an open approach and understands the need to be prepared for a future access network that supports a wide range of use cases. The decision regarding which technologies to support will largely depend on the individual market characteristics and the future network landscape.

5G (RAN)

The forthcoming fifth generation of mobile telecommunications technology (5G) is expected to offer significant improvements over the original 4G/LTE standards. Whilst 5G technology is yet to be standardised, and is not expected to be commercially deployed until at least 2020, a level of consensus is beginning to form between 3GPP and other stakeholders in the industry about key network requirements.

Figure 24: 5G release schedule

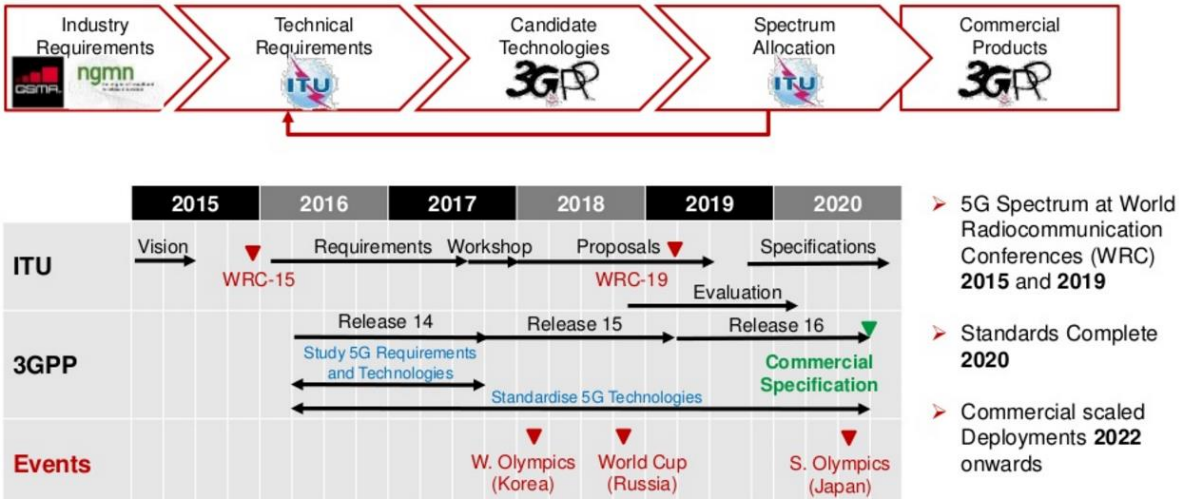

Source: GSMA

In contrast to previous standards, 5G is being designed with a significant focus on networks that enable machine-type communications (specifically M2M & IoT), in addition to improved 'traditional' mobile telephony.

Figure 25: Use cases for 5G networks

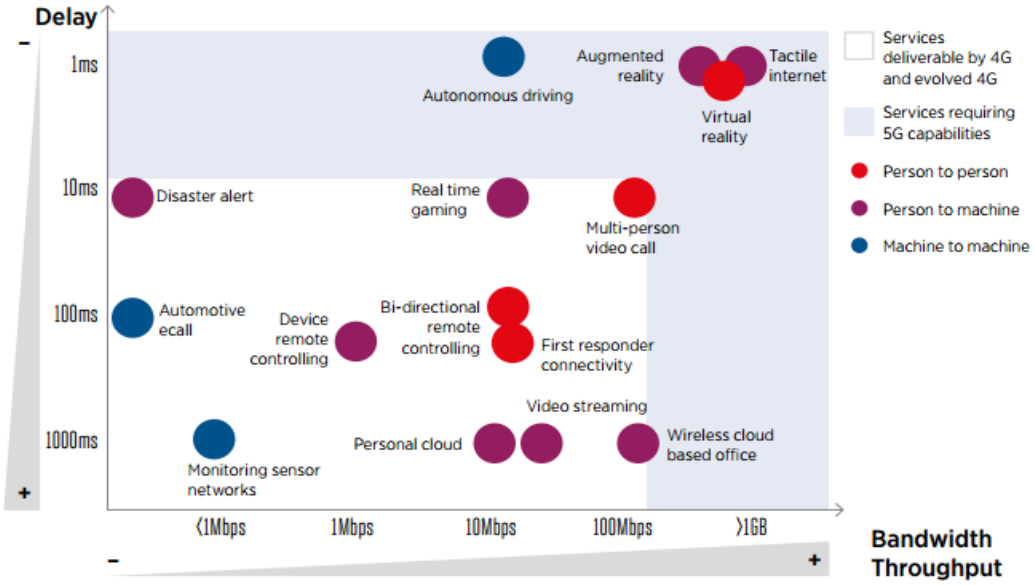

Source: GSMA – Unlocking Commercial Opportunities

As shown in Figure 255, 5G use cases are diverse, and as a result key network requirements are broad:

- **User base**
  5G networks must be able to support massive numbers of simultaneous users and connected devices: from hundreds of thousands to billions of connections, in order to enable massive sensor deployments.

- **Data rate**
  Able to support real data rates of tens of megabits for tens of thousands of users, as well as peak data rates between one and ten gigabits per second for mission-critical communications.

- **Latency**
  The GSMA has suggested that one of the unique attributes of 5G[5] will include the ability to deliver 1ms end-to-end latency, significantly lower than is possible with 4G/LTE. There are clear technical challenges with achieving this in reality, as it will change the structure of the network.

  End-to-end latency includes delays all the way from the origin of the data generated to the destination. It will not only depend on the radio interface, but on the amount of time it takes for the data to be processed on the initial device (e.g. a sensor) and at the control server in the network. Figure 26 below demonstrates an example of how the required 1ms latency is distributed for an IoT application where an actuator triggers an action based on data from a sensor. The 0.1ms latency objective at the radio interface further highlights the need for IoT applications that require these stringent latency requirements to be located at the edge. The physical distance between the device and the network will need to be minimised to under 1 kilometre, according to industry estimates, to even attempt to achieve 1ms latency.

Figure 26: Example of 1ms latency distribution

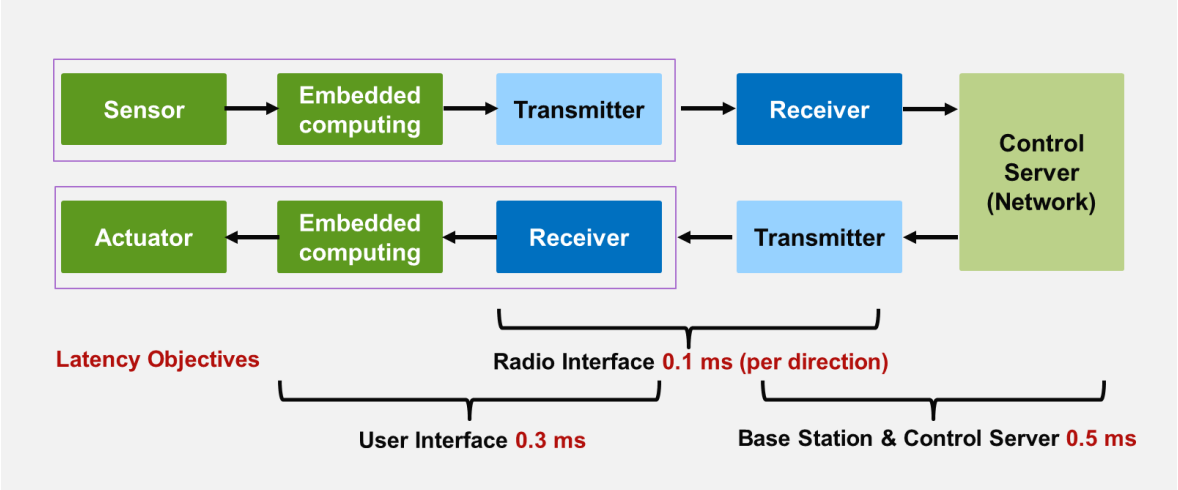

Source: ITU-T Technology Watch Report

- **Spectral efficiency**
  In order to support higher traffic and more connections, spectrum efficiency should be significantly enhanced over 4G/LTE. To be able to support the variety of spectrum bands and IoT services, new air interface technologies are being developed. These include new waveform technology, Filtered-OFDM (Filtered-Orthogonal Frequency Division Multiplexing); new multiple access technology SCMA (Sparse Code Multiple Access); a new channel code Polar Code; the full-duplex mode and massive MIMO technology.

- **Energy consumption**
  Energy-per-bit consumption should be significantly reduced to allow for low-power connected devices with battery life measured in years.

- **Coverage**
  Improved in comparison to 4G/LTE.

[5] GSMA Intelligence: Understanding 5G: Perspectives on future technological advancements in mobile

Etisalat recognises that implementation of edge computing techniques and some level of converged network access (including making use of fixed fibre backhaul) will be integral to achieving many of these requirements. However, if this can be achieved, 5G will bring many benefits to service providers supporting IoT. For example:

- 5G has the potential to act as a **one-stop networking solution**, which could supersede both existing low-power technologies (including NB-IoT) and traditional cellular networks.
- The technology will support **massive numbers of connections** out-of-the-box and be **highly scalable**, which will enable providers to implement IoT quickly and on a large scale.
- Reduced energy consumption in comparison to previous standards will enable providers to support **low-power devices** which are key to massive-scale IoT deployment.

Service Provider Operations

The sheer scale of the Internet of Things means that the IoT ecosystem will be made up of many players focused on different levels of the value-chain, some providing niche IoT products for specific verticals and others attempting to provide horizontal, cross-industry solutions. In turn, IoT applications will make use of different access technologies, protocols and network types, with management and configuration occurring over multiple IoT service platforms. Furthermore, business models will differ for different IoT services. In order to properly support these diverse IoT applications, regardless of where the service provider plays in the value chain, they must adapt and simplify their IT systems.

Impact of IoT on OSS / BSS

The IoT will bring with it new customers, services and business models, creating potential challenges for service providers' OSS/BSS systems:

- Managing and on-boarding IoT devices
  - IoT will result in many more devices and things. A service provider will need simplified and automated systems in order to support and manage these devices.
  - New IoT devices will be frequently connected to the network. Device on-boarding therefore has to be simple and automated.
  - Connectivity management platforms give service providers the ability to manage the connectivity element of the IoT solution. For example, in the case of a GSM network, service providers would be able to manage the mobile SIM. The management features vary, but generally allow service providers to turn on/off the connectivity element, set-up alerts and determine if there are problems with the connectivity element. Platforms also allow service providers to collect usage data from the connectivity element for tariffing and billing purposes.
  - Non-cellular devices/networks typically support OTA capabilities allowing devices to connect automatically.
- Diverse business models / billing arrangements creates a need for **network-aware** IT systems
  - The value of connectivity will vary depending on the IoT service. For example, for a low-power smart farm sensor, which transmits data only a few times a week/month/year on the pH of the soil, the value of the internet connection will be much lower than the connection of a health device, which can be mission-critical to its user. This will have implications on how different IoT applications are supported and monetised. This in turn will affect billing and reporting capabilities.
  - The value derived from individual network transactions will also be much smaller. This increases the need for provisioning and billing to be highly automated and flexible, to lower the cost per transaction and ensure the OSS/BSS can feasibly provide for these low-value transactions. Alternatively, it may alter the way provisioning of connectivity service is executed; low ARPU connections may be better suited to be provisioned in bulk, which will have further implications on IT systems.
  - Billing for IoT connectivity could be based on total usage, per device per year or a premium charge when connectivity is required. Managing multiple IoT services will require more sophisticated billing methods, based on real-time rating and charging, which reflects customer needs and usage.
- Customer overlap with existing services

  - IoT customers may also purchase other services/products from the service provider. To provide the best possible service to their customers, service providers must ensure they have a single common view of their customer.
  - Additionally, in order to both coordinate a unified OSS/BSS and provide the customer with a seamless experience, service providers should aim to have a single (IoT) product catalogue to allow the easy management of the various services offered.

- SIM provisioning
  - Many IoT devices will require embedded SIMs as they have long-life requirements and operate in tough-environments and hence may be sealed units.
  - The GSMA's Embedded SIM Specification provides a single, de-facto standard mechanism for the remote provisioning and management of machine-to-machine (M2M) connections, allowing the "over the air" provisioning of an initial service provider subscription, and the subsequent change of subscription from one service provider to another.
  - Companies like Jasper, now part of Cisco, can provide this capability to change SIM profiles on-demand or based on pre-defined rules or triggering events, allowing greater flexibility for swapping SIM profiles.

- Fault Management / Monitoring
  - With a multitude of connected devices, faults will pose a continual challenge; fault management / monitoring will therefore play an important role ensuring the stability of IoT services.
  - Service provider's performance/fault management platforms will need to be able to cope with a significant increase in the amount of data they handle in order to effectively monitor and react to faults, ensuring continuity of service.
  - Some devices will be able to send alerts directly if a fault has occurred, in fact this alert could be the primary function of an IoT device or service (e.g. safety-related IoT devices can send an alert if a fault has occurred). Other devices will not have the intelligence to generate an alert. These devices can be actively monitored through tools (e.g. ping) to determine if the device is working.
  - Intelligent service organisation will ensure that the service still operates even with faults occurring across of number of devices.
  - Edge (fog) computing can potentially allow faster recognition of and reaction to faults, minimising potential loss of data as well as providing additional resiliency for IoT services.

In order to mitigate and manage these challenges service providers will need to simplify their OSS/BSS systems. This involves creating an open and flexible architecture, supporting standard APIs and remaining open to DevOps.

To support IoT, the OSS/BSS must automate real-time interaction between users, partners, the network and the OSS/BSS capabilities and enable rapid self-care by users. These systems will become more data-centric, as it will be imperative to use data in order to flexibly and automatically provision and bill for services. Real-time analytics will be required to detect network and performance/experience degradation and self-heal. This will be particularly important for 'mission-critical' IoT applications.

Etisalat will build on its development of cloud-based OSS/BSS systems to be able to support its services with adequate flexibility, elasticity and scalability.

Key challenges and considerations

Ecosystem creation

Partnering will be extremely important for service providers as they look to take a leading role enabling the internet of things. Whilst some service providers may seek to develop IoT applications for specific industries, service providers will typically not have the capacity or skillsets to develop solutions.

Service providers should therefore look to build an ecosystem of partners and developers, creating applications across vertical markets, leveraging the service provider's infrastructure.

Figure 27: Service provider capabilities and the IoT ecosystem

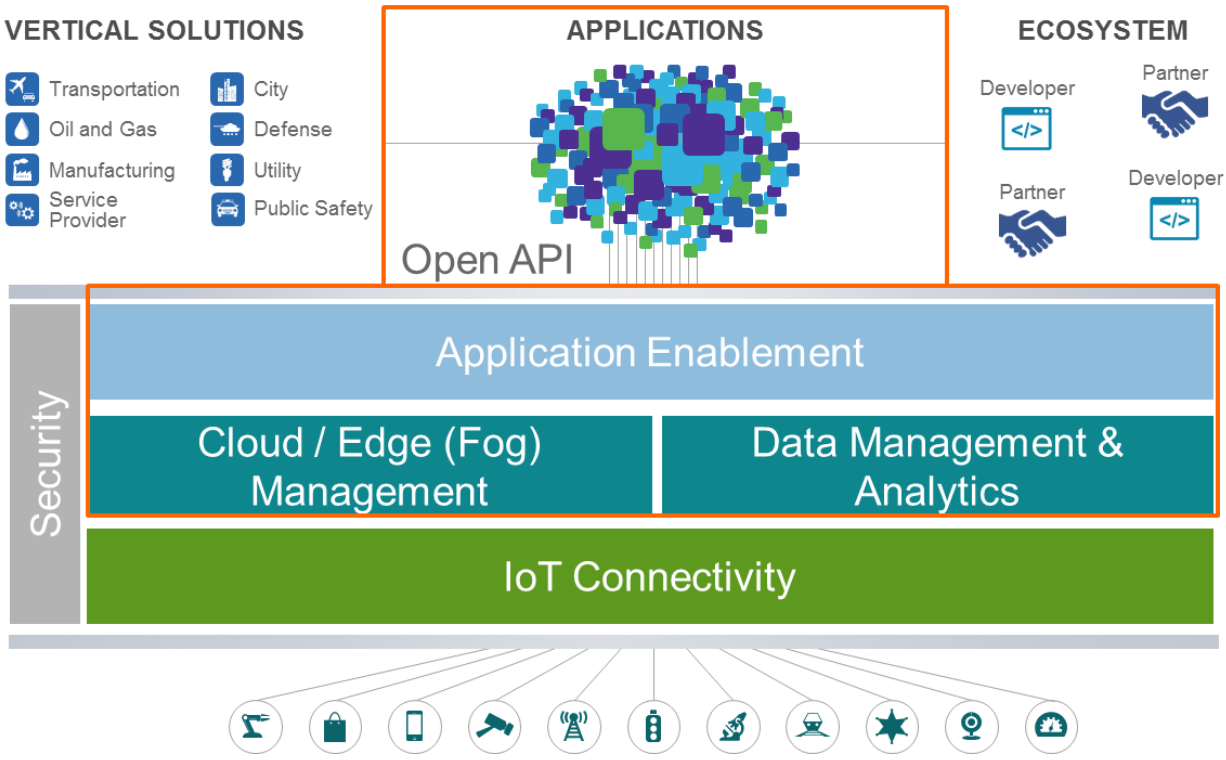

Source: IoTWF; Etisalat

Application enablement

In an IoT deployment, application enablement provides a way to store and transfer device data to applications in an efficient manner. Application enablement has a series of components including data storage, data management, data virtualization, analytics, platform rules engine and APIs.

Service providers are currently deciding how to implement application enablement to support their IoT businesses. Those that are selling a series of platform, application and integration services in some of their key target segments.

While few service providers have implemented a truly horizontal application enablement capability, there are some that have built or bought industry-specific capabilities covering sectors like automotive/transport, healthcare and connected home. These service providers have made strategic decisions, sometimes based on expertise in given industry sectors, to sell a more end-to-end solution in these industry sectors. Whether a service provider has implemented an industry-specific or a horizontal application enablement capability, almost

all service providers have partnerships and integrations to ensure the efficient storage and secure transfer of device data from the service providers' networks and edge devices to the end customers' applications.

APIs

In order to create an ecosystem for IoT application development, service providers should begin to open up APIs to third party developers and partners. Exposing APIs will allow developers to use more standardised formats and more closely integrate their IoT applications or solutions with the service provider infrastructure. This in turn will help create a more buoyant ecosystem around IoT applications that is enabled by the service provider.

Security

Network operators and IoT service providers often share similar security requirements. From an operational and cost perspective, therefore, it makes sense for providers to leverage common security solutions and avoid unnecessary duplication of security infrastructure. This applies more so if the network operator and IoT service provider are part of the same entity.

However, the unique characteristics of IoT will throw up new security challenges not encountered in traditional networking set-ups. Most important, the ever-increasing number of IP-enabled devices and services will create a larger 'attack surface', increasing the whole ecosystem's exposure to potential fraud and attack. It is therefore essential that IoT networks are properly designed such that threat and risk can be mitigated. Challenges to overcome include:

- Availability**:** ensuring constant connectivity between endpoints and their respective services
- Identity**:** secure authentication of endpoints, services and the end-user
- Privacy**:** reducing the potential for harm to individual end-users
- System integrity**:** ensuring that systems cannot be breached, and that system integrity can be verified, tracked and monitored

Potential solutions include:

- **Embedded SIM** and **UICC** (Universal Integrated Circuit Card) technology has been designed to withstand malicious attack such as glitching, side-channel analysis, passive data interception, physical tampering and identity theft.
- Secure identification can be ensured by using any, or a combination of, **IMSI** (International Mobile Subscriber Identity), **IMEI** (International Mobile Station Equipment Identity) and **ICCID** (Integrated Circuit Card Identifier).
- Providing a mechanism for keeping track of device versions and rolling out **over-the-air (OTA) updates** to maintain a desired level of security on the device. Knowing that many IoT device are designed to last a decade or more, leading IoT device and gateway vendors must make their devices updatable and upgradeable.
- IoT service providers can enable **data encryption** services to ensure integrity of communications and network resilience.

- Network privacy can be maintained either by using a combination of **L2TP** (Layer 2 Tunnelling Protocol) and **IPsec** (Internet Protocol Security), or by establishing a **dedicated instance of the core network** (with shared radio network) for IoT services.
- Use of **licensed spectrum** can reduce potential interference and maximise network availability to end-users.
- Implementation of **network operator-managed gateways** allows endpoint devices to connect securely, in a manner that best integrates into Wide Area Network security mechanisms. This also allows for monitoring and managed updates of firmware and software.
- **Storing as little data as possible at the edge**, only storing data that is necessary for accomplishing business objectives.
- Providers can leverage **data analytics**, **deep packet inspection** and **machine learning** services to identify anomalies and potential threats in IoT data which cannot be detected up by 'lightweight' endpoint devices.
- **Software-defined networking** gives the potential for automation of existing security models.

For any IoT deployment that has cloud-based components, service providers and IoT customers must ensure holistic security practices. First, multi-tenant cloud platforms must allow deployment of logically separated cloud IoT services in order to ensure that security and reliability for each tenant is never compromised by other instances running in the same cloud. Second, IoT traffic ideally should be isolated from general IT traffic to help ensure that general IT security breaches do not spread rapidly across the network and impact mission critical operations powered by IoT solutions. Finally, organizations should use network traffic monitoring solutions that rely on static rule-based approaches to setup baselines for expected traffic as well as on heuristic-based solutions that can spot unusual activity not easily detectable with predefined rules.

Analytics

Data analytics, or the discovery, interpretation and communication of meaningful patterns in data, is a key component of the Internet of Things. Growing numbers of IoT devices and sensors will generate unprecedented amounts of new data which, if leveraged properly, could generate enormous value. The GSMA, for example, predicts that an IoT big data ecosystem could bring economic benefits worth $44 billion per annum by 2025[6]. For example, data could be used:

- **To add value to existing services**
  e.g. on-the-fly systems optimisation using real-time information from sensors
- **To create new services**
  e.g. using live user data to create usage-based billing plans for consumers
- **For automated decision-making**
  e.g. using real-time environmental data to control street lighting

It is important that service providers have a way to conduct analytics in the edge (fog) platform, fairly close to the edge devices, and in the data centre or cloud. Having the ability to do analytics at both places gives a service provider the most flexibility in offering a managed service model to its customers. This becomes

[6] GSMA: Unlocking the Value of IoT Through Big Data

important to address the variety of use cases which service providers will need to support, ranging from high-bandwidth remote applications which require local high-speed processing (e.g.: Oil& Gas, Transportation, Remote sites) to high-speed cloud driven applications (e.g.: Smart cities, Retail Operations). While service providers are beginning to implement big data solutions for their own internal IT needs, a service provider's decision to implement analytics and visualization tools as a product to sell to customers is a more complex decision.

Challenges to realising the full benefits of data analytics include a lack of standardised developer-friendly data formats and APIs, difficulties in establishing trust and commercial agreements between parties, and the need to comply with laws and regulations that govern intellectual property, security, privacy, confidentiality and data sovereignty. However, in cases where customers have engaged service providers for big data services, the relationships are strengthened and both parties benefit.

Data sovereignty

The data collected by IoT devices often contains sensitive, personally-identifiable information or data sets that may have potential implications for national security. Examples include data on vital national infrastructure networks (power, communications, transport), live video from home security systems, confidential healthcare data, and personal financial information. National governments may impose strict regulations on how this kind of data can be processed, and where it can be stored. In some cases, sensitive information is not allowed to leave its country of origin. This creates legal implications for service providers planning to implement global IoT networks and leverage cloud technology.

Careful network and IT design (e.g. use of containerisation) can ensure that local data is kept local. By making use of edge computing techniques and locating data centres in-country, for example, service providers can avoid the need to send data abroad for processing and storage, thereby ensuring regulatory compliance, where relevant.

Conclusions

Etisalat recognises the transformational impact that IoT will have on our everyday lives, and is leading the way to develop the infrastructure to support and enable a wide range of IoT applications. Below we recap the key takeaways for service providers as they consider their strategy to support IoT:

- Service providers have unique capabilities and expertise that means they are well-placed to take the lead in enabling the Internet of Things.
- Service providers may choose to play various roles across the IoT value chain, including providing the connectivity infrastructure, creating horizontal platforms for wide-ranging IoT applications, or providing end-to-end IoT solutions for specific verticals.
- In order to properly cater for a wide range of IoT use cases with varying requirements, IoT infrastructure must be able to support: high data rates, low power usage, high numbers of connected things, deep coverage, ultra-low latency, advanced data analytics, security and data management capabilities, resilience and local autonomy.
- Service providers will need to make use of new technologies and capabilities in order to achieve this.
  - NFV/SDN driven network transformation will be one of the most disruptive technologies to date for service providers, eventually leading to the transformation of all key aspects of networks and operations. This transformation towards virtualised, programmable networks will also deliver many of the capabilities required to support and enable IoT services.
  - NFV/SDN capabilities will also allow service providers to adopt network slicing, which enables extensive resource re-use and optimises the use of the network environment.
  - Etisalat believes that the traditional separation of network and applications will begin to blur. Service providers are expected to address this by making use of edge (fog) computing, which will provide new, differentiated capabilities, allowing service providers to offer cloud computing and IT capabilities throughout their network, including, when commercially viable, closer to the end-user or device.
  - New access technologies will be a key enabler for the low-cost, low-power devices which are integral to massive machine-type communications.
  - The forthcoming 5G networking standards are expected to bring about improvements in speed, bandwidth and latency, allowing service providers to support mission-critical applications exchanging information in real time.
- The IoT will bring new customers, partners, services and business models, creating potential challenges for service providers' OSS/BSS systems. In addition, as service providers begin to deploy SDN/NFV, new capabilities will be required to manage and orchestrate the network infrastructure.
- To manage these challenges, service providers will need to create an open and flexible architecture, supporting standard APIs, such as that set out by ETSI's MANO project.
- Data analytics, or the discovery, interpretation and communication of meaningful patterns in data, is a key component of the Internet of Things. Growing numbers of IoT devices and sensors will generate unprecedented amounts of new data which, if leveraged properly, could generate enormous value.
- The unique characteristics of IoT will throw up new security challenges not encountered in traditional networking set-ups.
- By making use of edge computing techniques and locating data centres in-country, service providers can avoid the need to send data abroad for processing and storage, thereby ensuring regulatory compliance.

References

Etisalat Group, 2015. *2020 Landscape*

ETSI, 2014. *Network Functions Virtualisation (NFV): Management and Orchestration*

GSMA, 2015. *Unlocking the Value of IoT Through Big Data*

GSMA Intelligence, 2014. *Understanding 5G: Perspectives on Future Technological Advancements in Mobile*

IDC, 2014. Worldwide and Regional Internet of Things (IoT) 2014–2020 Forecast

IoT World Forum, 2014. *IoT Reference Model*

ITU, 2014. *The Tactile Internet: ITU-T Technology Watch Report*

NGMN Alliance, 2015. *NGMN 5G White Paper*

This whitepaper is prepared in collaboration with CISCO