General

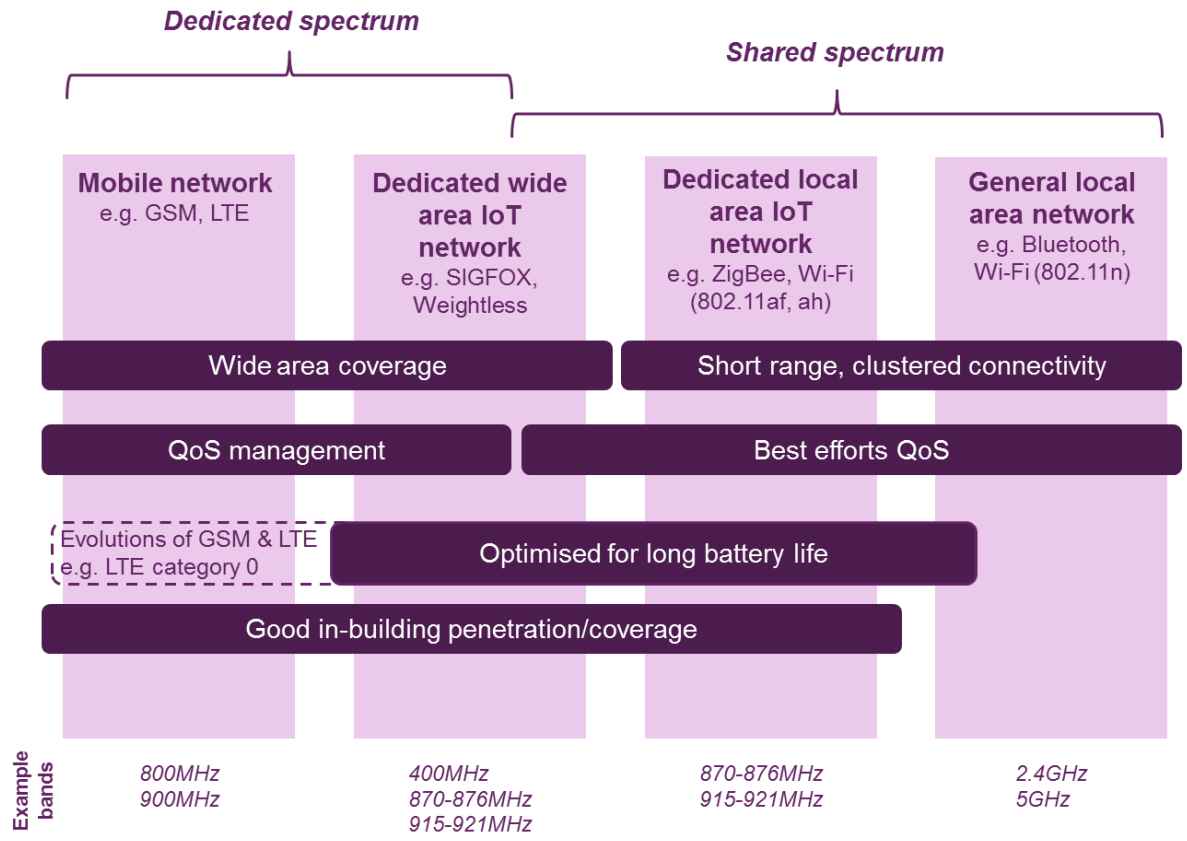

The figure (Figure 1 in the Consultation Document) is misleading in two respects. QoS is an application requirement that may be best effort or managed according to application needs and is the same regardless of the spectrum being dedicated or shared. Battery life, similarly, is unrelated to spectrum needs although it is clearly a major factor for many applications using devices with no power source of their own. We do not think these elements should be included in a framework for assessing spectrum requirements.

Clustering is not exclusive to short-range systems, although the intent is clear in the figure. Clusters of devices would also be served by wide-area systems. For mobile devices, clusters may form and disperse according to the dynamics of the applications and their users.

In-building penetration in the bands allocated to IMT is variable and connectivity could be compromised, affecting IoT applications requiring reliable connectivity, although this could be mitigated by the anticipated dense deployment of small cells enabled in future 3GPP Releases.

It is arguable that the machine-type communications (MTC) capabilities being standardised for wireless mobile cellular networks in 3GPP could meet many needs for wide-area IoT systems, complemented by short-range technologies in existing ISM bands operated as dedicated private networks or ad-hoc consumer installations. However the economic viability for LTE has yet to be proved, e.g. would resources to support IoT applications for a very large number of devices generating limited traffic (one scenario) be made available by LTE operators who could make more return from high-value mobile broadband services.

Finally, many wide area IoT systems will be supported by fixed and mobile, narrowband and broadband, satellite services. Satellite communications must be included in the above overview.

1.46 IoT definition, applications and demand

The range of IoT devices, applications and supporting services that is likely to emerge across different industry sectors, along with views on potential market size. We are particularly interested in stakeholders' definitions of the IoT and views on which applications are likely to dominate and the characteristics of these applications (in terms of their range, quality of service, connection speed and data throughput, radio cost, battery life etc.).

Section 1.1 of the Consultation Document mentions certain broad domains with a focus on the possible benefits that might accrue to individuals. We agree with this limited characterisation as stated but draw attention to the significant enterprise (private and public) involvement needed to deliver those benefits, and probable regulatory issues within and beyond Ofcom's remit (please see response to 1.52 below).

Thus, in our view the IoT is a natural organic evolution towards a common platform, (the Internet) of many disparate technologies supporting applications and services that are useful to consumers, businesses and society, hitherto implemented in a haphazard ad-hoc way without the need for, or opportunity to, achieve any integration. Some systems are widely used, for example SCADA in industrial applications. In fact most deployed systems can be found in professional industrial situations, not in consumer domains.

The nature of these historical applications and services, considered alongside other Internet applications, in particular social networking and access to entertainment, suggest that there is no dominant application. There will be certain trends from time to time and one may be more evident than others at any instant.

The IoT is the enabler for integration and, with significant additional effort, interoperability. The extent of the need for integration depends on the benefits to be gained from it, e.g. the sharing of information across domains (as noted in 1.4.2). One particular benefit of a common platform and a high level of interoperability is that devices (sensors, actuators, user terminals) can be reused by several applications.

Some estimates of the number of IoT devices – sensors and effectors but not user terminals – are alarming (e.g. 50 billion) but the assumptions that they are based on are quite crude and take little account of the realities of system design and deployment. However, it is clear that a large number of devices will be deployed and that they will require a scalable and extensible naming/numbering system.

We agree in general with the analysis of IoT application, or system, characteristics given in section 1.13. We recommend adding mobility as this could significantly affect the dynamics on IoT applications and require more harmonisation and uniformity of service provision. For example, a person who owns a collection of IoT-connected devices connected via Bluetooth, say for giving indications relating to a medical condition, must be able to lead as normal a life as possible, including travelling. At present many such systems only work in the home. Enabling them to work when the person is on the move, including travel abroad, presents difficult challenges and may require regulatory oversight for this class of application..

1.47 Spectrum requirements

The need for additional spectrum to meet the expected demand for wireless connections between IoT devices. In particular, we would welcome views on which specific frequency bands are desirable, the need for internationally harmonised bands, whether additional spectrum should be made available on a licensed or licence exempt basis, and whether shared or dedicated spectrum bands will be needed.

Additional spectrum, harmonisation

The need for additional spectrum, dedicated or shared, depends on several factors, such as:

- The nature of the applications – do they require long-term reliable connectivity at high data-rates or are they short-lived, generating small messages at low but bursty rates, and tolerant to disruption (or any point between these extremes);
- The growth in the number of IoT devices and applications;

- The additional overheads required to maintain security or to assure protection of safety–critical services;
- The clustering of devices implementing application instances, and the distribution of systems, especially those that require dedicated spectrum

Concerning specific bands, we note the WRC-15 Agenda Item 1.1, concerning with possible allocation of various higher-frequency UHF bands for use by IMT. If agreed then additional spectrum would become accessible to IoT applications via LTE MTC, subject to the commercial considerations noted above (in General comments). We support Ofcom's position for this topic in the WRC, including the need for harmonisation, provided aviation services are protected.

We note also the continuing evolution of the 3GPP Releases and the new capabilities being developed for LTE and LTE-Advanced with the ultimate target of a unified air-interface in "5G" suitable for IoT devices that could span all the dimensions indicated in Figure 1. The use of mm-wave spectrum up to 300 GHz would have a profound impact on the way IoT devices communicate with each other. Existing spectrum allocated to fixed-link back-haul below 90 GHz could also be reallocated for use by IoT devices for LTE MTC at future WRCs.

There is a good case to allocate a modest amount of dedicated spectrum for use by IoT applications operated over a wide area. This was historically quite successful, e.g. for smart-metering in the 1990s in 434 MHz and 868 MHz so there is a precedent for such a need. Although these were ISM bands, their usage and the distribution and range of deployed systems seems to have been at a level that tended not to lead to problems of coexistence. The lack of harmonisation of these bands limited deployment of these systems however, so we recommend that such an allocation if made should also be harmonised. We support the current approach proposed for M2M in the 700 MHz band systems currently under discussion in preparation for WRC'15.

Finally, we draw attention to the services operated as MSS and FSS. There is a large amount of spectrum that could be assigned for IoT applications by satellite operators, particularly for the FSS. The satellite terminal configurations, especially antenna sizes, at higher frequencies could be economic for some IoT applications. Using satellite services would extend the reach of IoT applications and require no additional spectrum allocations. Given the trend to release many frequencies for use by LTE, Ofcom must ensure that sufficient spectrum continues to be available for satellite communications.

Licensed or exempt

The above examples include both licensed and license-exempt spectrum. We recommend that such a decision be taken case-by-case according to the technical, legal and commercial requirements.

Shared or dedicated

The need to dedicate spectrum or allow sharing depends on the willingness of operators of IoT devices, and the ability of their applications, to coexist. We believe that sharing should be encouraged.

1.48 Network-related issues

We are interested in views on a number of IoT network and infrastructure related issues, including:

1.48.1 Approaches to delivering IoT services

Broadly, services could either be delvered using conventional mobile networks, in general licence exempt bands or via bespoke networks that are optimised for the IoT. Other approaches may exist between this range of options. We are interested in opinions on the approaches to delivering IoT services that will likely emerge, citing advantages, disadvantages and views on which applications might be better suited to a particular approach.

Concerning network-related issues, distinct from spectrum issues (see 1.47 above), we do not believe that there is any specific IoT optimisation that would be within Ofcom's remit alone. However we do consider that, for example, having due regard for the integrity, availability and confidentiality requirements of IoT applications operated for personal medical purposes or for energy management (power retailer or consumer), authorities responsible for oversight of those services must engage with Ofcom on issues concerning communications quality whether fixed or wireless.

1.48.2 Degree of openness

IoT services could be deployed over entirely open networks, i.e. any manufacturer's device conforming to a particular technical standard can be connected; or over a closed network, in which the operator controls which devices can access the network. We are interested in views on which of these (or similar) approaches might develop, whether particular services are suited to an approach and what the implications might be for the development of the IoT. We are also interested in views on the role of open versus proprietary standards.

It is highly likely that providers of derived IoT services will emerge that use only open networks while having commercial relationships with regulated operators of healthcare, transportation, energy management and other applications that are constrained by their respective regulators concerning what can, or not, be connected. We consider that this trend must be encouraged to promote the IoT. Inevitably other compliant devices will become connected to these derived service providers, possibly because individual consumers elect to do so in ignorance of possible consequences to personal privacy.

Proprietary standards implemented in closed networks should not be ruled out provided coexistence with other standards can be assured. Most proprietary specifications have mechanisms for connecting to web-based service platforms that allow sensor and actuator devices to be accessed indirectly.

There are many specifications for M2M systems that have been standardised by national and international standards bodies, e.g. ISO/IEC/SC25/WG1. They include systems for domestic use and other for professional industrial applications (SCADA). They are highly diverse in their use of spectrum and protocol specification. Systems that use them may be large and numerous but are effectively closed with respect to other systems using different standards.1.49 Security and resilience

Across the range of IoT services there are likely to be a variety of security and resilience requirements. At one extreme there may be applications that can be supported on a best efforts basis, whereas other applications may need to be highly available and resistant to malicious attack. We are interested in views on the steps required to enable the IoT to support high levels of security and resilience.

Security is achieved through protection mechanisms at various layers and management functions that support them. For example, encryption using private keys can be applied at individual link layers or end-to-end between application functions with support from a key manager. This is not sufficient: information must be authenticated to verify its source and be authorised for use at its destination. Producers at the source and consumers at the destination must be trusted and accredited. Overall this constitutes cybersecurity.

Provision of an accredited information assurance management framework and associated services that ensure the required availability and protection from attack is the responsibility of the IoT application provider. In a complementary way, compliance with the procedures declared by the IoT application provider to assure confidentiality, integrity and availability is a shared responsibility of the consumer.

The steps required to protect IoT applications could be onerous depending on the sensitivity of the information to be protected. This sensitivity is unlikely to be a concern for national security, although extreme scenarios could be devised to substantiate this: it is more likely that personal data, e.g. somebody's location, could be attacked and misused.

It will be essential to establish and maintain trust in any cybersecurity measures that are put in place.

Greater connectivity between devices and applications will inevitably lead to new vulnerabilities that can be threatened with attack. The cybersecurity of IoT systems is likely to become an issue and greater awareness of this topic and associated regulatory measures will be required of IoT device and application developers.

Expanding on our comment on QoS as collection of requirements that are independent of the technology platform that implements the IoT application, this is not the case for the ability of the

technology platform to fulfil those requirements. Many different wired and wireless technologies will be combined to provide the common IP network layer. There will often be several routing choices. These may be visible to the IoT application and allow it to make choices based on link quality, e.g radio-aware routing.

Devices that are mobile are likely to change their point of attachment, also making choices based on link quality, so performance may vary during the lifetime of the IoT application. All these factors affect the compliance of the IoT application with its desired targets.

Resilience is a property of a system defined by its ability to survive disruptions. Disruptions can take many forms, planned or accidental, that may be caused by physical processes (weather, damage, or RF interference), network behaviour (congestion, loss of routing, mis-configuration), or by the application itself. They may be caused, or revealed, by malicious attack or human error (user or administrator). Historical distributed ICT applications make provision for these to the extent that most of their functionality is concerned with recovery from disruptions. IoT applications will be no different. Resilience of the communications service may not be sufficient for these reasons. However in some cases it may be necessary. For example, devices of low size, weight and power (SWAP) may not be able to maintain buffered data to retransmit after a network failure. Their data may have priority, criticality and urgency that influences intervening network path and routing priorities. The network service provider plays a role in delivering such data intact and in time. Ofcom and the concerned application authorities may have to establish performance requirements for end-to-end communications based on standard message types in addition to maintaining availability and the required service quality.

These considerations will affect the relationship between consumers, application suppliers, application service providers and network service providers profoundly.

1.50 Data privacy

We are interested in the nature of privacy and data protection issues that may arise through the development of the IoT, including views on approaches to appropriately manage personal or commercially-sensitive data.

Apart from the issues noted in our other responses, we do not believe that there are specific IoT issues to address concerning data privacy. Data protection regimes that apply to other data collected by electronic means should be applied to the IoT.

1.51 Numbering and addressing

We are interested in views on the likely nature of demand for device addresses and to what extent this demand might be for electronic addresses and/or telephone numbers. We are also interested in the extent to which demand for device addresses, in the form of telephone numbers, IP addresses or other identifiers, could be a barrier to the deployment of IoT services.

Sections 1.32 to 1.38 of the Consultation Document summarise relevant concerns about numbering and addressing of IoT connected devices. We agree with the statements in general but note the following relating to those statements:

- Few existing IoT applications require public IP addresses or end-to-end connectivity. Where it is needed, web-based platforms and IP namespace extension tools (NAT) are used as enablers;
- The MS-ISDN numbering system used in mobile networks has sufficient flexibility and extensibility that would not be challenged by very large numbers of IoT devices requiring such addresses;
- IPv4 public addresses have been in short supply for over 10 years. There is sufficient churn that availability has not been compromised as badly as expected so far and a viable market in the asset has worked adequately, again: so far. We do not believe that the IoT will change this significantly or suffer in new ways different from those endured by other Internet users;
- NAT is sustainable but only because the need for individual public IP addresses has been supplanted by indirection mechanisms. However, for other reasons (see below) we do not expect this situation to change;
- Mechanisms unique to IPv6 originally have been implemented in the IPv4 space. These facilitate local private addressing for ad-hoc local network connection, or device discovery for example;

- IPv6 is universally enabled in most products. It is not used because the bulk of web services are installed in the IPv4 space. It could be used in an entirely private closed system.

However, network layer addressing is not the preferred means of accessing IoT devices (sensors and actuators). Protocols such as CoAP and MMCP use Universal Resource Identifiers (URIs) that select a capability, not a specific device – although they can be stated as IP addresses when these are known.

Therefore, we do not believe that availability of network layer addresses is a barrier to deployment of IoT services.

1.52 Devices

We would welcome stakeholders' views on technical and commercial developments that could affect the cost and capability of IoT devices, in particular in relation enabling the manufacture of low cost devices with low energy consumption and long battery life. We are also interested in views on the role that existing or emerging device operating systems will play.

The price and capability of IoT devices will reflect the extent to which they solve problems for consumers – i.e. do they fill a niche or are they a mass market high volume product. Low cost, low energy consumption, long battery life are desirable but the need will vary according to application.

The availability, reaction time, cyberphysical integrity (i.e. accurate measurement and assertion of physical quantities), will be determined by the regulatory domain (if any) in which the devices operate. A temperature sensor that is supplied as a part of a heating system may require no regulation in that context. However, if used to maintain a life-critical temperature for a baby or critically ill person, it becomes mission critical and subject to a different regulatory regime.

The device operating system and the applications that it supports must have a design assurance level that reflect the safety critical requirements of its use, which may change from time to time. The design and validation of complex hardware and software for safety-critical systems is a well understood process, e.g. the Eurocae ED79A process used in the aviation industry.

1.53 Digital literacy

We welcome views on the role of digital literacy in underpinning the growth in take-up of IoT devices. What steps, if any, will be required to enable citizens and consumers to understand the potential benefits and risks of the data created by their devices being shared? What steps is industry taking to address this challenge?

Citizens and consumers appear to have little regard for usage of the data that they create through use of IT web services and the like at present. We do not believe that we can influence this complacency significantly but will make best endeavours to do so.

However, consumers are sensitive to intrusions by government and business that are perceived to be invasions of their persons or home for the purposes of collecting their data. For example smart-metering studies show that unless there is a regulatory requirement, many citizens reject such monitoring even though they accept the potential benefit in reducing carbon footprint or delivering other environmental benefits.

IoT applications will succeed if they solve problems and offer benefits. Industry should also ensure that the risks are given as much attention as the benefits in training and instructional materials.

1.54 Data analysis and exploitation

The capture, analysis and exploitation of "big data" from multiple devices and applications to provide new, innovative services. We are interested in views on whether there will likely be demand for such services, on the nature of the services and whether there are any barriers to their development.

We believe that the demand for such services will be very high.

1.55 International developments

In the longer term, IoT equipment is likely to be developed for a regional or global market; this will be necessary to drive down device costs and achieve economies of scale. We welcome views on

relevant international activities, such as the development of common technical standards, trials and commercial deployments.

The development of international standardisation of IoT-like specifications has been underway for more than 30 years, e.g. in CENELEC or ISO/IC/JTC1/SC25/WG1. Many of these standards are used only in a few countries, often only in one country, and have been taken up on a modest to large scale.

ETSI has developed its own set of specifications, supported by the oneM2M initiative. A wide-area low-power specification, LTN, has also been published recently.

The SCADA domain also has some useful lessons. While technical standards for SCADA messaging exist, regulatory requirements based on the application scope (criticality, security, safety, spectrum availability etc.) that address the needs of operating the services are yet to be clarified.]

It is likely that the most credible solutions at network level and above will emerge from the IETF. In fact, most new developments provide for IP connectivity and it has been integrated into several existing specifications. Many of the issues concerning routing for low capability devices and link quality awareness are already in IETF drafts. IETF standards also cover messaging protocols that can be used in IoT applications, and protocols for security.

Google and Facebook also develop their own protocols and APIs based up IETF protocols. These could be a strong driver in the development of the IoT.

We see no need for IoT trials but do not discourage them, for example the experimentation done in the European Commission's IoT projects under FP7 and possible large scale pilots (which are not trials) funded under the FIRE, FIPPP or 5GPPP initiatives.

1.56 Ofcom's role

We recognise that the IoT is a fast-moving area in which industry is well-placed to create a range of innovative technologies and services. To enable us to best support these efforts, we welcome stakeholders' views on our role across the range of policy issues raised in this document, including spectrum management, network resilience and security.

More broadly, and in light of our general responsibility to encourage investment and innovation, we welcome stakeholders' views on the role we should take in driving the development of the IoT.

IoT applications bring specific application domains under separate regulation together with communications functions that are within Ofcom's remit. Ofcom must engage with other regulators to ensure that the communications services have the required communications performance and security to ensure compliance with the accompanying application-specific regulations. Other regulators must ensure that Ofcom is aware of their required communications performance for operators of IoT services.

The Consultation Document seeks views on barriers to exploitation of "big-data" generated by the IoT and Ofcom's role in removing those barriers. IoT data may be more intimate at a personal level, e.g. reference to vital signs that could be used by third parties to advocate inappropriate medical treatment. Clearly mechanisms should be in place to protect consumers from this type of intrusion and possibly existing means available to Ofcom are reusable. Conversely, this kind of data if anonymised could be valuable to emergency services or it could be used in managing an epidemic or predicting needs for medical supplies. The barrier in this case could be that there is as yet no precedent for discriminating such cases, and little knowledge of how to apply current rules a burst of infringements arising from misuse of very large amounts of data.