OFCOM Response
To be completed at
https://stakeholders.ofcom.org.uk/consultations/iot/howtorespond/form

1. IoT definition, applications and demand
The market size is almost limitless.
I have talked to organisations considering rolling out 300 million or more devices in 5 years.
I work as CTO for an agri-informatics firm that can see a demand for a sensor per 10 metres on farms if the price comes down.
Battery-life and radio is a challenge, but low-tech solutions already exist.
Indeed, many deployments use near-field comms to gather data 'haphazardly'
(eg a farmer as he walks across his fields) with local base stations (smartphone, tablet, etc)
acting as the IoT enabled router.
Such approaches then use heuristics to identify devices that have low data-counts
or "haven't checked in recently" for explicit data collection.

Connection speed and data throughput is not a challenge provided one does not assume
that IoT devices are like web browsers!
Using sensible message-queue protocols,
rather than REST (see 4 security and resilience) and custom data compression
(eg only sending sensor points that have changed by significant value, followed by the occasional key frame, a bit like digital TV)
enable quite sophisticated data gathering.
For example, Smith Electric Vehicles were effectively recording over 32 data points
every few seconds per truck, and transmitting them over 2G, back in 2010.
Quality of service and resilience is also straightfoward to address, too (see 4 again).

AB>>>>>>>>>>>>>>

Market segment	Quality of delivery High means in order exactly once	Bandwith Low means prompt delivery.	Latency required
Healthcare			
Hospital automation			
Drug delivery	High	Low	Low
Patient biometric monitoring	High	High	Low
Safety monitoring	Medium	Medium	Low
Care policy monitoring	Low	Low	High
Outpatient			
Remote monitor	Medium	Medium	High
Remote diagnosis	Medium	High	Low
Recreational			
Fitness monitoring	Low	Low	High

Smart Cities			
Traffic			
Monitoring	Low	Medium	Medium
Active Management	High	Low	Low
Toll charges	High	Medium	High
Severe weather / flooding / wind			
Monitor	Low	Medium	High
Alert delivery	High	Low	Low
Civil defense	Medium	High	Low
Policing			
Face/biometric recognition	High	High	Medium
Vehicle recognition	High	High	Medium
Mobile device recognition	High	High	Medium
Offender tagging	High	Low	Medium
Agricluture			
Crop health monitor	Low	Low	High
Machinery monitor and control	High	Medium	Low
Retail			
RFID sensing			
Theft prevention	Low	Medium	Low
Auto checkout	High	High	Low
Stock level monitoring	High	Medium	Medium
Logistics			
Location tracking	Medium	Medium	High
Theft prevention	Low	Medium	Low
Domestic			
Heating			
Control	High	Low	Low
Monitor	Low	Low	Medium
Fire detection	High	Low	Low
Intrudor detection	High	Low	Low
At risk individual monitor	High	High	Low
Vehicle			

Fault diagnosis	High	Low	High
Location monitoring	Low	Medium	High
Remote assist			
Vehicle unlock	Medium	Low	Low
Emergency callout	High	Low	Low
Insurance use based charging	Medium	Medium	High
Inter vehicle			
Convoy management	High	High	Low
Collision avoidance	High	High	Low
Utilities			
Electricity			
Remote Metering	Medium	Low	High
Supply termination/restore	High	Low	Medium
Problem Determination	High	Low	Low
Quality monitoring	Low	Low	High
Consumption feedback	Low	Low	High
Demand control	High	Low	Low
Gas			
Remote Metering	Medium	Low	High
Consumption feedaback	Low	Low	High
Water/Sewerage			
Remote Metering	Medium	Low	High
Consumption feedback	Low	Low	High

MQTT is particulary appropriate for High Quality of service and/or low latency applications
where data must not be lost or must be delivered promptly.
<<<<<<<<<<<<<<AB

2. Spectrum requirements
Internationally harmonised,
which are licence exempt (or freely licensed) would cause an explosion in IoT usage
and be to our country's advantage.
If companies can research,
develop and deploy new technology in a simple framework with minor costs,
they will do so here.
And in doing so, we enable that great British strength - the inventor.
Most IoT work today comes from start-ups and small companies.
Evidence from the MQTT mailing list shows that most innovation in this sector starts from the grass roots,

not large corporates. Growing this sector should be a priority, and spectrum is a way to drive it.

3. Network-related issues
a. Approaches to delivery IoT services.

b. Degree of openness
MQTT is a candidate Oasis standard, governed by Oasis IP policies.

4. Security and resilience
The cultural view must be inculcated that lack of consideration for security is not acceptable.
As a diabetic with high-blood pressure,
I don't want someone hacking the actuator on my remote blood pressure monitor and crushing my arm.
More insidiously, I don't want my readings subtly altered,
or simply not delivered because a juvenile cracker's getting off on using all my bandwidth for streaming cat videos.

a. Security
One of the great challenges of IoT in adoption is acknowledging the need to secure data,
even when it apparently is not 'valuable',
identifiable, or private. This is because implementors:

- Firstly, do not recognise the potential for 'data leaks' and accidents
  ("whoops, we shouldn't have transmitted that to over there")
- Secondly, underestimate the potential for others to make such data valuable or identifiable
- Thirdly, view privacy differently to some of their potential customers
- Fourthly, don't understand the technology and its vulnerabilities.

To mitigate this, it should be basic, common practice that:-
- All data is transmitted over a secure channel (eg TLS tunnel, VPN, or IPsec)
- All messages (payloads) are separately encrypted,
  and remain so when stored on a server or intermediary (and to which the intermediary has no secret key)
- All devices, where functionally capable, store all data at rest encrypted (ie encrypted file systems)
- All devices authenticate with servers and vice versa
  (this is mostly easily and openly done using TLS with X.509 client certificates)

In doing these things, we should adopt open,
well-understand and freely implemented standards that can be used with the widest possible range
of devices and operating systems.
My recommendation would be to make use of TLS 1.2 (or better) with X.509 client certificates.
However, one should note other standards, such as ISO 29192.
We should not implement our own new security standards.
The early mistakes of WiFi in trying to invent their own security standard (WEP) is salutory indeed.

The OASIS MQTT committee, of which I am Co-Chair, has published on this subject two pieces of relevance:

- The MQTT 3.1.1 standard, Section 5 Security (http://docs.oasis-open.org/mqtt/mqtt/v3.1.1/mqtt-v3.1.1.html)
- Open guidance on how to adopt the NIST Cybersecurity Framework (http://docs.oasis-open.org/mqtt/mqtt-nist-cybersecurity/v1.0/mqtt-nist-cybersecurity-v1.0.html), which includes an example implementation with field devices. It also provides acknowledgment and links to other industry-wide standards that are commonly-agreed to the best practice examples of securing data and services.

b. Resilience

Resilience is an interesting concept.
When one is deploying device counts in the millions,
failure becomes the new normal.
And not just for the devices - for bandwidth, for controlling servers, and everything else.
It's a bit like that analogy you hear about Google and disk drives.
So what becomes interesting is how one designs for degraded service.

Firstly, you throw out all you ever new about the web, and client-server design - the idea that a browser connects to a website. Using technologies such as REST is just a rookie mistake. Then you go looking for how people with large data volumes and lots of problems have gone about solving the problem. And you find two interesting sources. The pipeline, etc folks using SCADA, and the banks doing trading or the academics solving big numerical problems. In other words, you look at two opposite ends of the spectrum*1:

- how do embedded people do it at scale?
- how do traders do it at scale?

And the answer is message queuing (MQ) and heuristics. Let me explain.

When you've got so many devices, a poll model (go and talk to every device) doesn't work.
Nor does synchronous communication - which is how websites work over HTTP (or using REST).
Instead, what you want is a loose-coupling (a pussh model),
where different devices talk to each other when they can, and buffer up data (messages) for when they can't. They don't expect an immediate response, either - if at all. Such an approach can also accommodate different qualities of service, too, giving different levels of resilience.

Inevitably, some messages will get lost when scale is that large (eg a mobile lost signal).
Perhaps it doesn't matter - we call that At-Most-Once messaging.
A use for that might a thermal probe tick.
Perhaps we want to make sure our message got their,
so we ask the receiver to sign for it, like in the mail.
We call that At-Least-Once messaging - if we don't get the signature back, we send the message again.
Lastly, we might really care they get the message once and once only.
We call that Exactly-Once messaging. That requires more bandwidth.

So where do heuristics fit in? Surely,
if I want to check the health of all my devices,
I talk to them? Well, yes, you can, and on a small network you could, too...
But in practice,
it's more instructive to look at patterns of behaviour.
For a group of devices, what where the message transmission patterns last week?

Do they get heavy at 1pm on Fridays? So why has device 1Z45GHJ not got the same profile?
(Perhaps it connects via a home hub, and the consumer turns it off when she's at work to save electricity).
By studying message flows,
and types of messages,
one can learn far more about scaling and making resilient one's infrastructure*2.
As the Army says,
no plan survives contact with the enemy
- and it's the same with large nets of devices.
One can predict behaviour before rollout, but then...

Ideally, of course, the protocol messaging queuing (MQ) uses such be open so that anyone can take part.
And it should be independent of transport, too.
Today, there are only two open, standards-backed MQ protocols: AMQP and MQTT.
AMQP is for the data centre, for bank-to-bank transactions.
It's complex but works with superb guarantess, but heavy bandwidth usage.
MQTT is lightweight and simple to use - it even works on an Arduino,
and their potential in-chip hardware implementatiosn coming.
It's also de-facto what's being used.
Ford Automotive is a large early adoptor,
along with Duke Energy.
Why? Because its openness makes it easy to adopt and free-to-use.

*1 Interesting, the middle - desktop PCs, enterprise data centres, website hosting - is not where you go.
*2 We see the same lesson from companies selling tools to analyse web logs or security access logs.
  The volumes of data are too much (I once registered 10,000 hacks an hour on a simple web server,
  and that's the tip of iceberg),
  so heuristics wins the day.

5. Data privacy
It is essential that the internet of things makes data privacy and security an overriding concern that trumps protocols,
implementation needs and time-to-market; it needs to be baked in.
IoT devices are likely to be small and very infrequently changed; some may end up baked into concrete, or put down wells or other inaccesible (or hazardous) locations. As such, it is essential that data privacy solutions address device lifetime. So what analogous sector could we look at to determine what risks the future holds?

Given that IoT devices are likely to be rapidly commoditised to reach a very low price point,
yet continue to increase in technical 'grunt', it seems reasonable to choose low-priced,
electronic FMCG. Devices such as TVs or home routers. And there in lie our warnings:-

- Just about every make of home router on the internet has been either broken it,
successfully hijacked or otherwised abused in the last 24 months;
- Home TVs have been found to violate settings explicitly made by consumers to not transmit viewing data 'back-to-base';
- Devices come pre-enabled with insecure configurations, such as 'phone home',

and the like;
- Car manufacturers have actively denied, and prosecuted,
security researchers who have found gaping holes in electronic keys and diagnostics;

Even when these issues are exposed, fixes from manufacturers have been slow to appear,
if at all - and only if the product is very current. And even then, there exists no way for
consumers not subscribed to deeply technical news sources to become aware of such problems.
And to top it all, even the professionals aren't safe.
A highly experienced co-worker of mine chanced earlier this year on a botnet controller being run off his home router.

So how do we solve these intertwined issues that _will occur_ for the internet of things?

In two ways:-

- Firstly, it must a mandatory requirement that devices are field upgradeable,
remotely. This is not particularly onerous, and protocols such as MQTT make this a snap to do.
- Secondly, and quite controversially,
manufacturers should be obliged to release the firmware for their devices,
and make such information as is necessary available for purchasers to modify or replace firmware on devices.

The second point seems onerous,
but it isn't really; most devices are going to be quite simple with the majority of logic in them already have been derived from third-party open source software and freely available. With the rapid growth and conversion of embedded devices towards Linux, and less so, the BSDs, manufacturers stand to lose very little; the only losers are the long declining market segment of proprietary operating system manufacturers. Indeed, today, many companies that want mass adoption of their SoCs (system on a chip), already do this (eg Allwinner and Cubiebox), and others freely license to achieve it (eg Ardunio, OpenCores).

Such an approach will mean that there is far less likely to be 'sunk cost' in devices; indeed,
it opens the way to novel uses unthought of by the makers
(Who remembers the internet enabled coke machine? Did Coca-Cola design it?).
And it plays to a great British strength - the tinkerer,
the inventor, the grandad in the shed.
Without their ilk, televisions and LCDs would never have come about.
The MQTT protocol's wide adoption has come from the ground up - not the top down like so many others,
and because it is open.
Open software, together with open standards, is the key to the future,
and the value (ie profit) to be made in IoT is from scaling and managing devices,
not manufacturing them. Scalers and managers would greatly

6. Numbering and addressing
We expect the growth of devices to be very great indeed; in the order of 100s millions,
if not billions, over the next few years. As such, this is likely to place significant demand for two key things:-

- Device numbers
- Device identifiers, in the form of keys or certificates

Such volumes require not only a way of issuing unique values that can meet demand,
but routing, registries (authorities) and revocation mechanisms that can scale with such demand.

Use of a publish subscribe paradigm can be particularly useful here. MQTT defines a simple and well specified
topic matching scheme that enables the following.
- Light touch adminstration of the producer and consumer name spaces.
- Independence of the producer and consumer, i.e. each can be created and started without reference to the other.
- Character based names with few practical limits, the major burden that remains is for the application creator to architect
a flexible naming structure for their application.

A weakness in MQTT is thats its names (Topic names an Topic filters) are not time limited, they do not expire.
This needs to be addressed. Equally any other IOT provider needs to consider how the name and address spaces
will be decluttered.

Device Numbers
As such, we believe that IPv6 based addressing best serves the needs of device numbering; it meets all the requirements for volume:-
- A sufficiently large range of numbers, unlike telephone numbers or IPv4;
- A mature and efficiently designed routing protocol suitable for vast numbers of devices;
- A long-established, well understood and fair method for issuing unique values;
- Likewise, in conjunction with sister protocols such as Mobile IPv6, DNS, DHCPv6,
  Nearest Neighbour and self-discovery, a way to revoke and re-use addresses

Device Identifiers
It is essential that the internet of things makes data privacy and security an overriding concern that trumps protocols, implementation needs and time-to-market; it needs to be baked in.
As such, devices will need to be identified in a way that transcends their IP address or telephone number.
The right technology for this today is probably X.509; it meets most of the requirements for volume:-

- X.509 certificates are unique
- Registries and Authorities exist
- Practice and Usage is commonplace and well-understood
- The security model is tried and known

However, revocation is fraught with difficulties.
Large scale rollouts (eg NHS Spine) have shown that X.509's revocation protocols,
such as OCSP and CRL, do not work well at scale when their is frequent churn
(in the NHS case, contract and bank staff turnover); for large volumes of frequently re-used devices,
one can imagine that revocation will be the new normal.
X.509 also places artificial constraints on certificate expiry and usage which have been exploited

in the past by vendors to 'price-out' competitors and non-preferred business sectors.

7. Devices

As devices become more capable for a lower price,
I fully expect Linux-based operating systems (including Android) to become much more prevalent.
Economies of scale will push devices towards the most commoditised hardware (eg ARM SoCs).
As the market grows, new entrants will want to 'get up and going' and prototype devices without paying for a production run,
specially licensed OSs or SDKs. Many entrants will not be traditional embedded system manufacturers,
but kickstarter and crowdfunded efforts who'll want to use commodity tools they're familiar with.
We should do our best to enable this community,
as it is where most innovation in IoT ideas are coming from to date.

Consequently, we should look to enable common,
easily understood and readily-available building blocks that can be combined 'lego-like' to create devices:
we should off-the-shelf protocols and software libraries.
Such solutions may not be as pure or singly focused as a bespoke effort,
but would reach the market sooner and so firm up demand for aspiring firms quicker;
turn around for 'succeed or fail' will be shorter and investment costs lower.
I've seen this already - IoT children's dinosaurs changing colours,
or specialised drone IoT uses, prototyped in weeks not years.

8. Digital literacy
Consumers need to be educated to buy devices that
  (a) use open security protocols,
  (b) can be freely upgraded and
  (c) by default, don't share their data in unexpected ways (eg TV viewing preferences).
Such education is almost certainly best backed by the equivalent of a CE, Kite Mark or Egg stamp,
with advertising and awareness being simple ("don't buy it if doesn't have the stamp").
Device testing and approval can then be contracted out, allowing British firms to develop leadership in this area,
or self-certifying, with a simple one-off 'type approval' check list (eg Please provide the URL for firmware updates, etc).

9. Data analysis and exploitation
Demand? Definitely. Many companies are already scaling up.

For sensors, capturing and anylising the huge volume of data is an even bigger challengs than gathering and delivering it.
Over burdensome regulation (for example to store and provide access to large volumes of data) might inhibit innovation.

10. International Developments
The OASIS MQTT committee has developed an open IoT protocol (http://docs.oasis-open.org/mqtt/mqtt/v3.1.1/mqtt-v3.1.1.html) and security framework (http://docs.oasis-open.org/mqtt/mqtt-nist-cybersecurity/v1.0/mqtt-nist-cybersecurity-v1.0.html). It is backed by

leading software and hardware firms, including Cisco, IBM, RedHat, Software AG, Solace, VMware, TIBCO, stormmq, LogMeIn, M2MI, Blackberry and others. (Membership: https://www.oasis-open.org/committees/membership.php?wg_abbrev=mqtt). Tt is in widespread use today, including PoCs by Ford and Duke Energy (https://gigaom.com/2013/11/03/plugging-interoperability-into-the-nations-electric-grid/).

The MQTT standard is proceeding to become an ISO standard via the ISO/IEC JTC 1 process.