
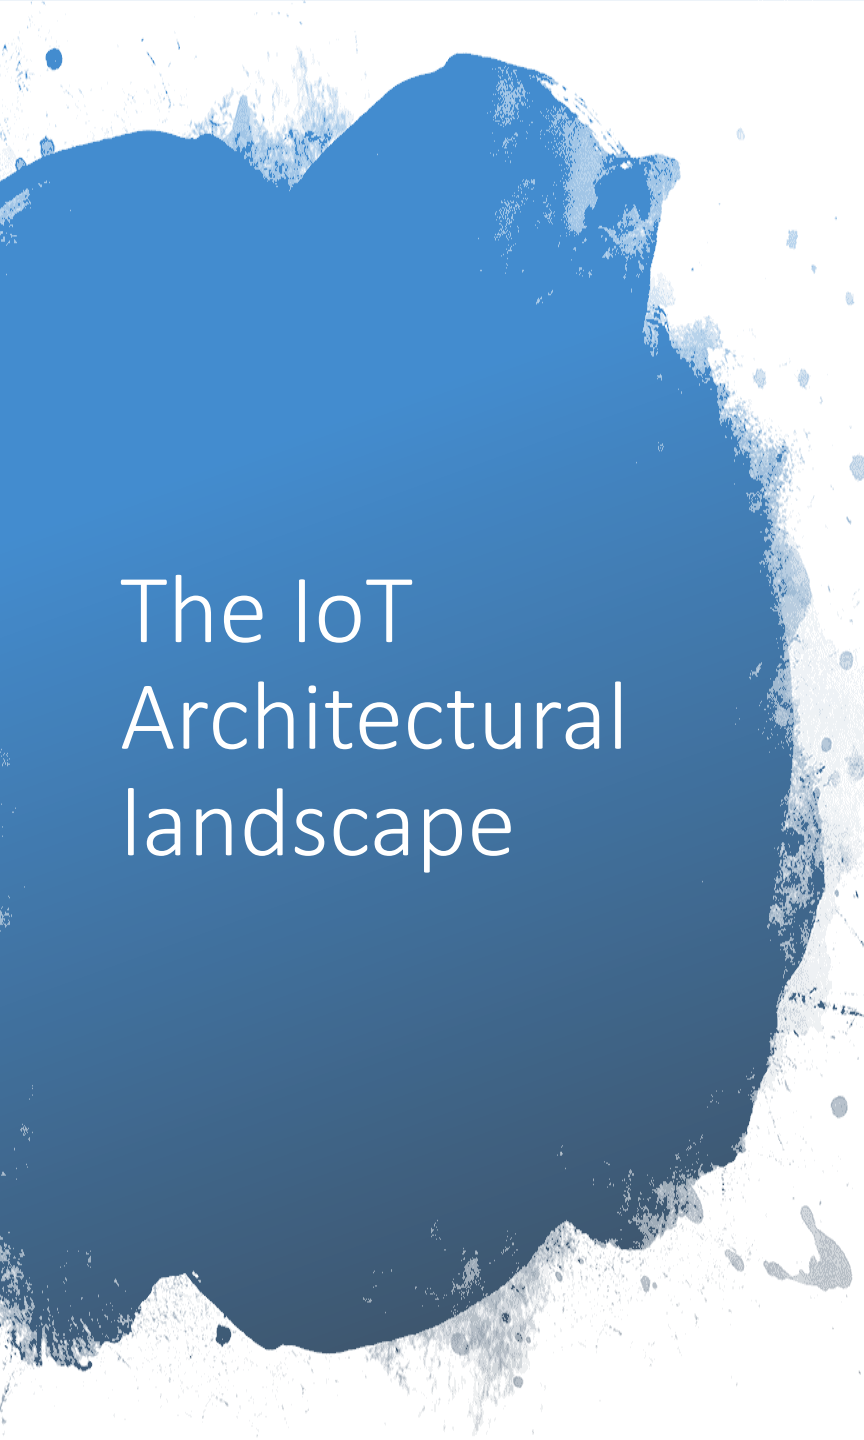


# Internet of Things System Architectures

Paul Patras





# The IoT Architectural landscape

- Thousands of new applications spanning numerous domains.
- Each comes with its own requirements; combining these leads to complex (difficult to manage) and often proprietary systems.
- Defining a unified architecture is challenging and interoperability problematic if too many standards to chose from.
- Documentation scattered and often difficult to navigate.
- Efforts to define common frameworks
  - ITU-T, 3GPP, ETSI, IETF
  - EC (via collaborative projects),
  - Industry consortium (e.g. Open Connectivity Foundation),
  - Big players (e.g. Cisco).
- We will not cover everything, but focus on the key principles these architectural patterns share and examine some examples.

# Key considerations



What application domains should be covered?



Where to place the “intelligence”?



What networking structure should be employed?



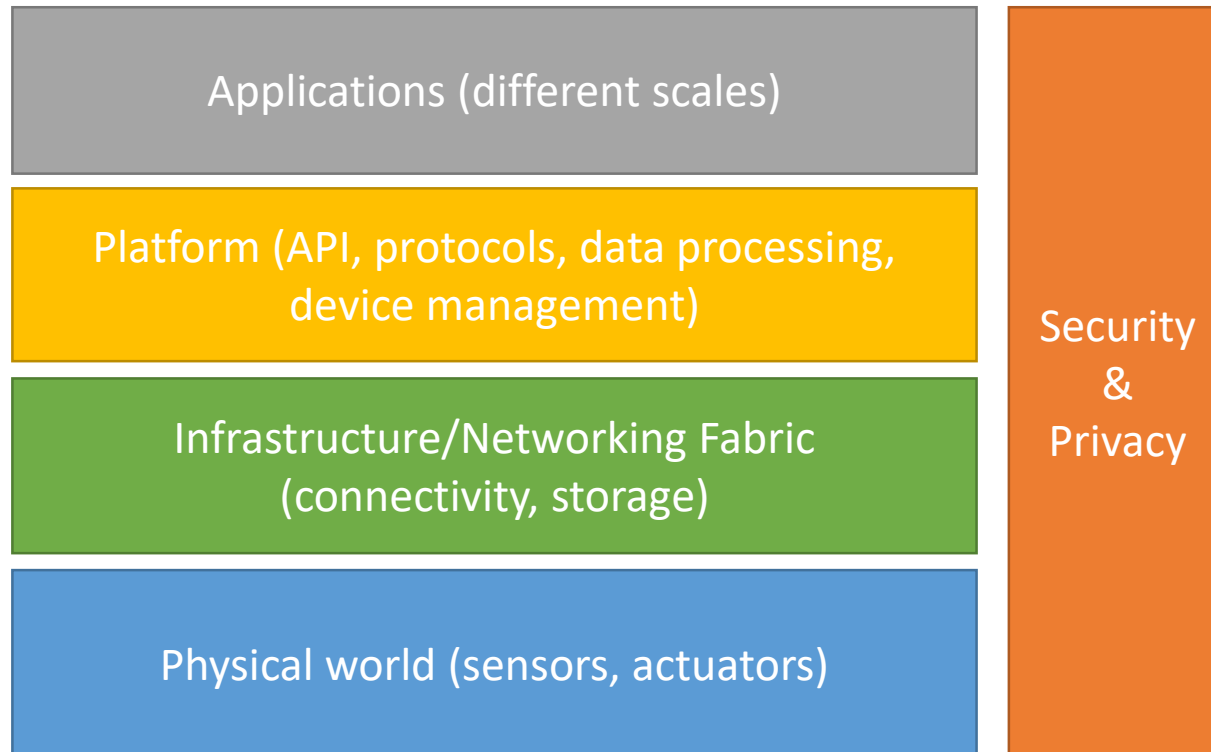
How to modularize systems, so as to manage complexity and enable programmability?



What about costs and scalability?

# Stakeholders slowly converging

At a high level, the shared view looks like this



# Immediate advantages

- Enabling software/app developers to build applications without having to understand the specifics of a device – Platform as a Service (PaaS)
- Better sharing and strict partitioning of network and computing resources (slicing); taking away from service providers the burden of building and managing a network– Infrastructure/Network as a Service (IaaS/Naas)
- Allowing IoT device manufacturers to focus strictly on improving their performance, power consumption, etc. – expose only well-defined interfaces to software platforms.

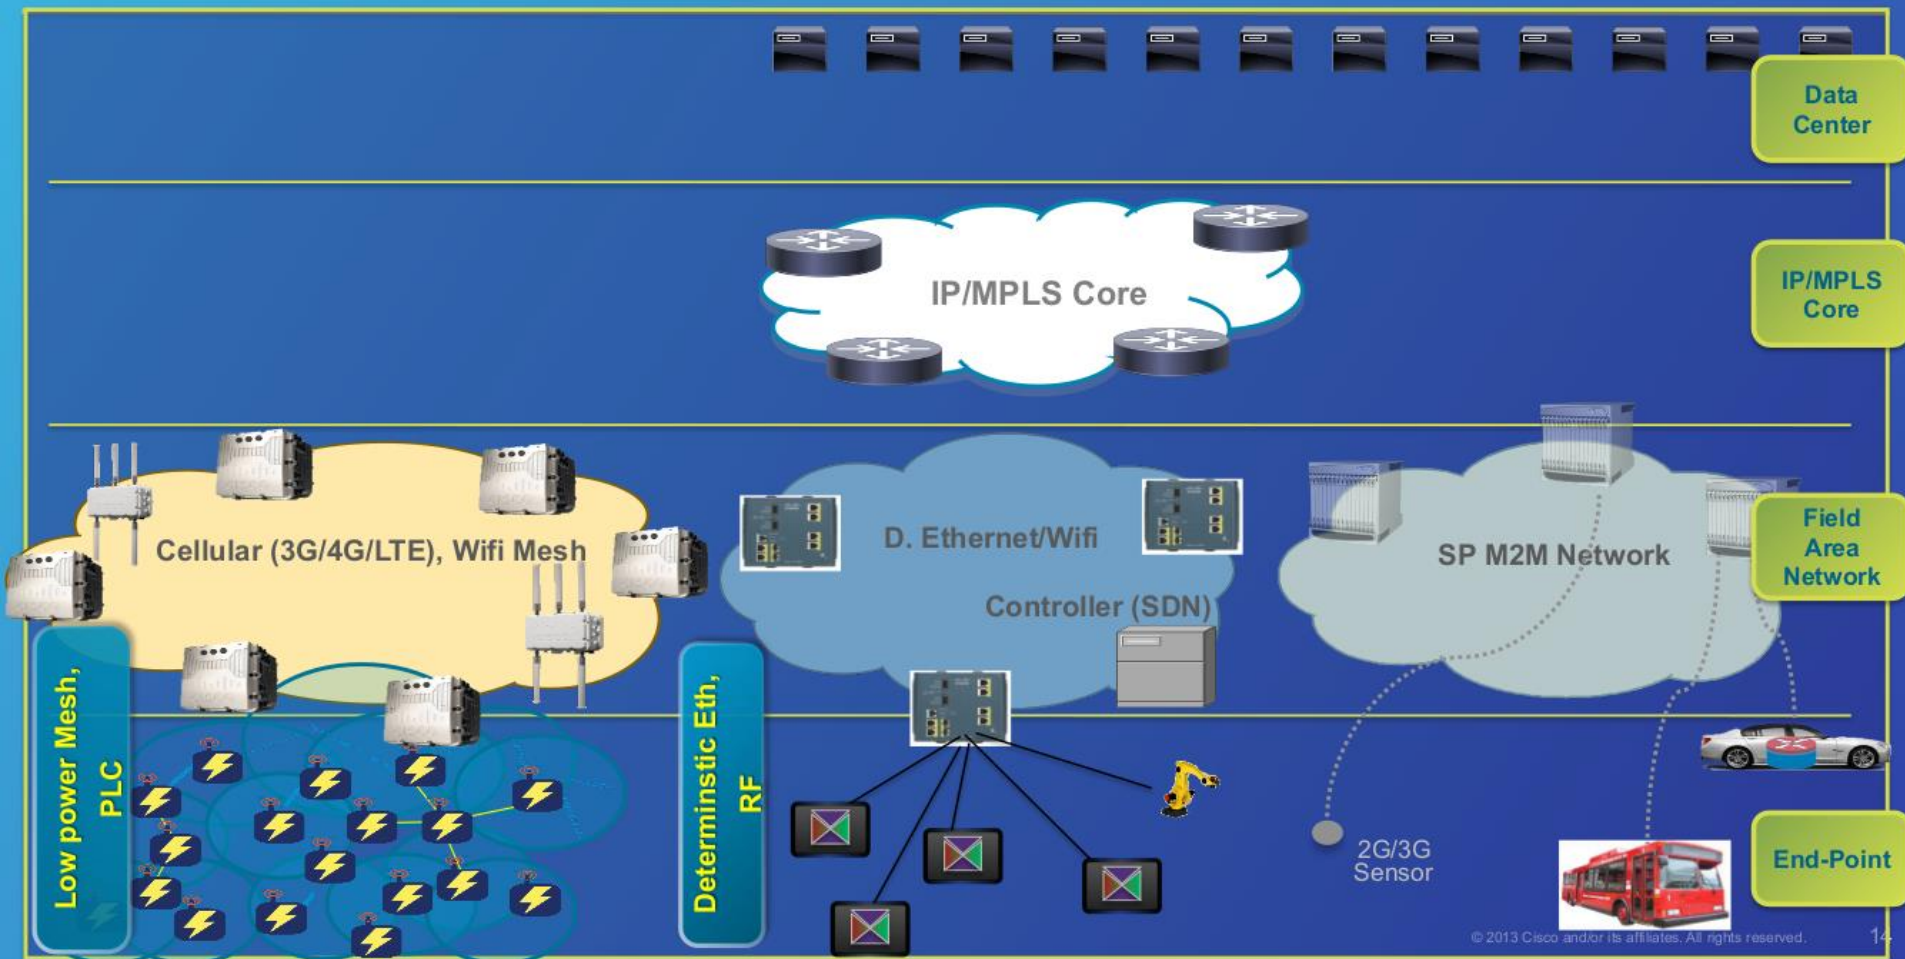
# The challenging part: Security

- Hardware isolation (eFuse, ARM TrustZone)
- Middleware (Intel MPX- Memory Protection Extensions)
- Network isolation (VPN, SDN)
- Software isolation (Sandboxing)

End-to-end security not straightforward – encryption remains the only option, but sometimes expensive (computing power, communication overheads)



# What happens in practice?



# Cloud vs Fog vs Edge

Cloud computing dominated the networked systems landscape until recently

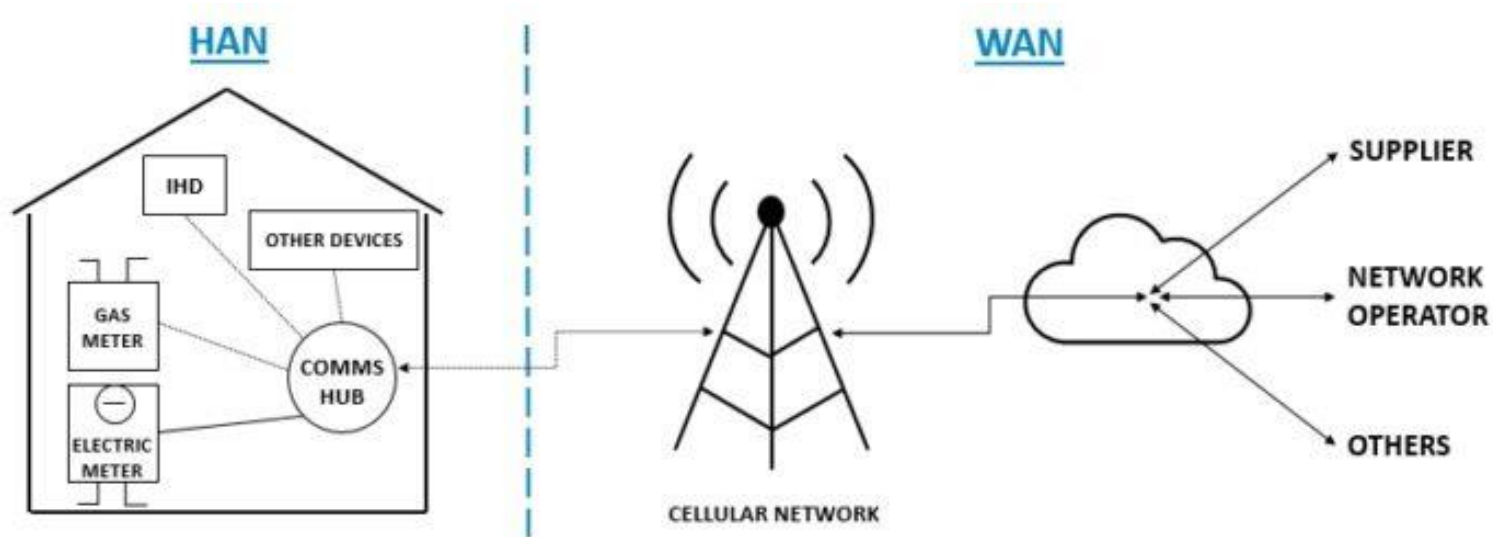
- End-devices merely information gatherers
- All intelligence in the cloud (relational DBs, analytics, web interfaces, control functions)

As the number of devices grows, applications diversify and generate more data, this will not scale

- Routing and storage costs
- Signalling overheads
- Latency inappropriate for real-time apps



# Example: Smart metering



Source: iwireless-solutions.com

- Simple sensing devices; data relayed over cellular; processing by multiple entities in the cloud

# Fog computing

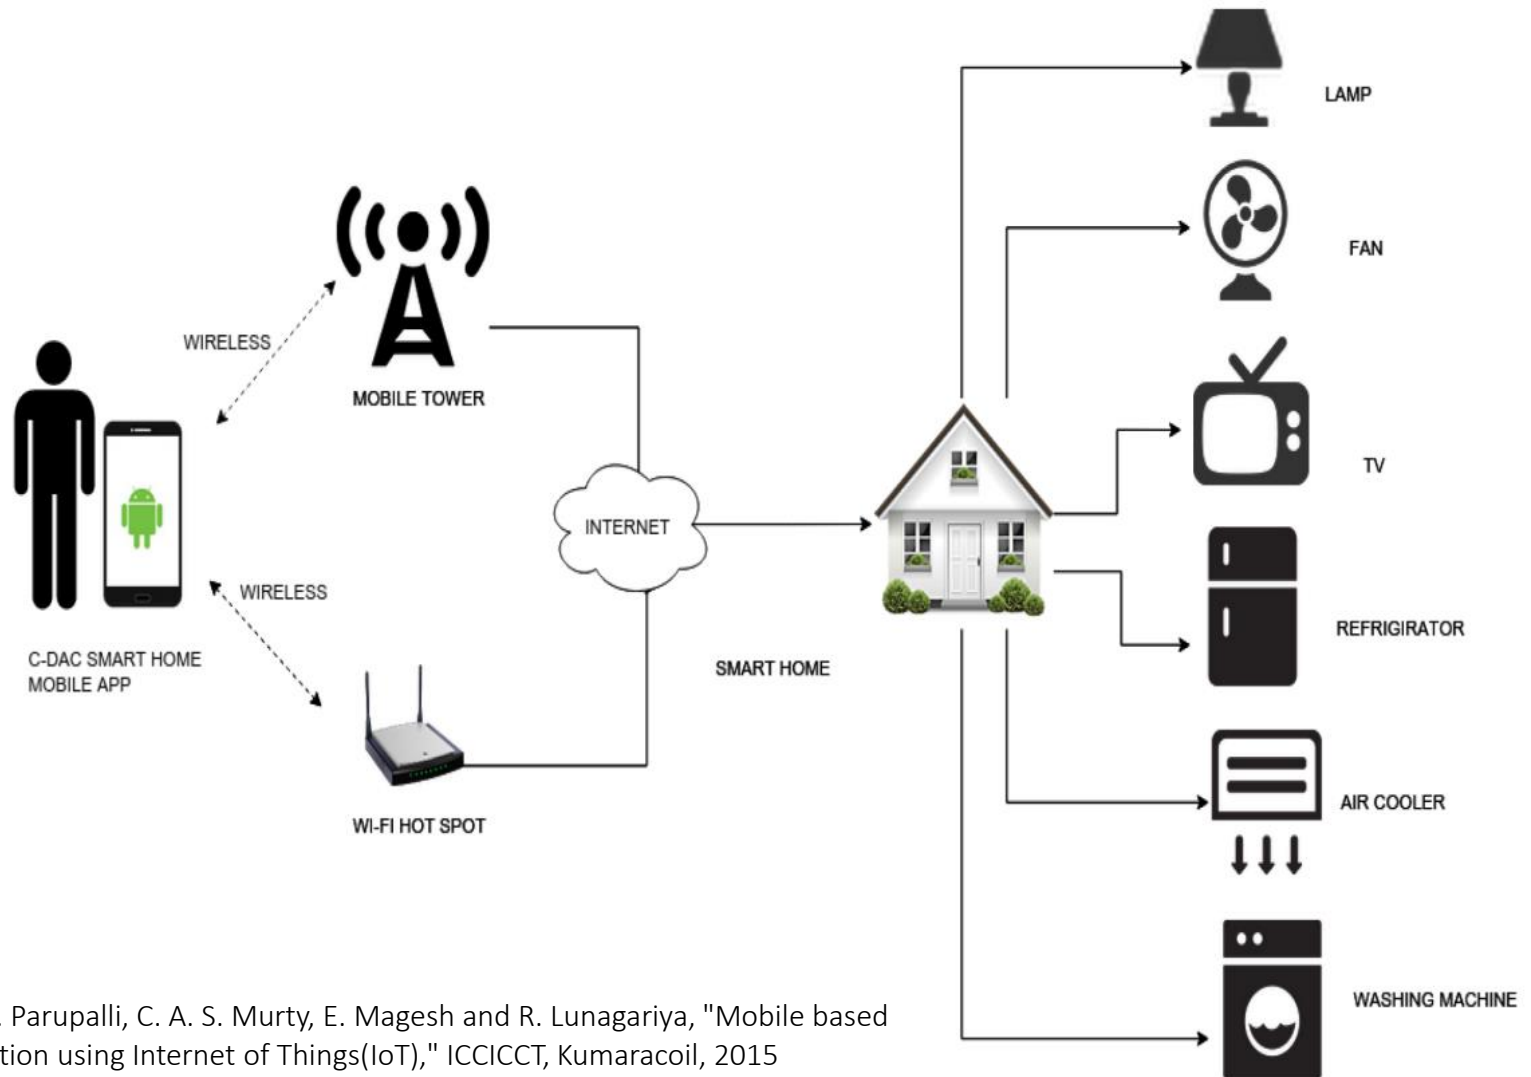
- Pushing some of the effort closer to the device, i.e. to access networks/gateways.
- This includes data aggregation, compression, (partial) processing; making localised decisions.
- IoT device
  - Not required to be extremely smart, i.e. unique address and ability to communication directly with the cloud
  - No substantial storage
  - Battery powered (lifetime)

# Roles of gateways

1. Data filtering and processing (e.g., aggregation of summaries, compression, etc.)
2. Protocol translation and interfacing between different connectivity technologies
3. Security (e.g., data encryption, firewalling)
4. Data flow multiplexing, packet routing

**Scalability problem:** as the number of devices grows, so will the number of gateways that are required


# Example: Home automation



K. Mandula, R. Parupalli, C. A. S. Murty, E. Magesh and R. Lunagariya, "Mobile based home automation using Internet of Things(IoT)," ICCICCT, Kumaracoil, 2015

# Example: Fitness tracking systems

The user's mobile phone acting as gateway



# Edge computing (process as much as possible where data is collected)

- Pushing processing power, communication capabilities, intelligence down at device level
- Emerging applications range from autonomous vehicles, to VR glasses, to earbuds.
- Do as much processing as required on the device, transmit only what is relevant long term or summaries
  - Low latency and decentralised decisions
  - Less signalling and communication overhead
  - Personalised experience

# Example: Deep Learning at the Edge


**Hardware Support:** Low-power chips specialised in computationally intensive tasks (IBM TrueNorth, Movidius, Huawei Kirin)

**Software:** lightweight inference frameworks optimised for constrained devices (mobile TensorFlow, Apple CoreML, DeepSense)

**Dedicated NN architectures:** Model compression (SqueezeNet, MobileNet), point-wise group convolution (ShuffleNet), model pruning (NestDNN).




# Ultimately it is about performance






# Standards (I)

- IEEE, 3GPP, IETF, and several alliances standardising the ‘language’ devices speak among each another / with gateways or cloud services.
- IEEE (the Institute of Electrical and Electronics Engineers) mostly dealing with definition of protocols for access networks
- Targeting the ISM bands with 802.15.4 (used by Zigbee) and ‘WiFi’ extensions for low power wide area networking, i.e. 802.11ah (HaLow)




# Standards (II)

- 3GPP (3<sup>rd</sup> Generation Partnership Project) focuses on specifying cellular network architectures and protocols (GSM, 3G, 4G-LTE, etc.)
- Developing standards for cellular communications tailored to IoT applications
  - LTE-M - compatible with existing LTE networks, easy to roll out, limited to max 1Mb/s speeds
  - NB-IoT - different band, lower capacity (200kb/s), different modulation, but no need for gateways.




# Standards (III)

- IETF (Internet Engineering Task Force) focusing on specifying protocols for
  - Routing (6LoWPAN, Routing Over Low-power and Lossy networks - ROLL)
  - End to end communications (TCP/IP, HTTP, CoAP, MQTT)
  - Security (DNSSEC, Datagram Transport Layer Security - DTLS)
  - Software updating (SUIT – specific to IoT)
  - Language independent formats (JSON)



# Standards (IV)


- Industry alliances
  - Bluetooth (personal area networks, application profiles)
  - ZigBee (on top of IEEE 802.15.4, inexpensive consumer/industrial)
  - LoRaWAN (low-power long range)
- Other (NIST – encryption, ISO – IoT reference architecture, ITU – recommendations, reference models)



# The complete picture



Source: AIOTI WG3 (IoT Standardisation) – Release 1



If you know what I mean

# The complete picture



Source: AIOTI WG3 (IoT Standardisation) – Release 1