IOT on campus

Infrastructure and platform choices for the SURF community

1. SURF explores the feasibility of an Internet of Things applications ecosystem for the Dutch research & education community.

In SURF, the Dutch education and research institutions work together on ICT facilities and innovation in order to make full use of the opportunities offered by digitisation. The evolution of 'Internet of Things' applications for the SURF community in the Netherlands is one of the topics that SURF is asked to address. SURF has a strong track record using a 'federative approach', trying to provide standard interfaces for cooperation while maintaining local autonomy and flexibility of the member institutions.

With the success of eduroam and data management solutions in mind, SURF investigates the possibility to create an 'open common IoT platform' and infrastructure

With the success of initiatives such as the network, storage, SURFconext, SURFwireless and eduroam in mind SURF is now considering to enable and facilitate a 'common platform' for IoT applications where this can offer common advantages for institutions. This way the use of IoT applications can be encouraged and facilitated and a country wide IoT platform is created for research, knowledge transfer, pilots and efficient application development.

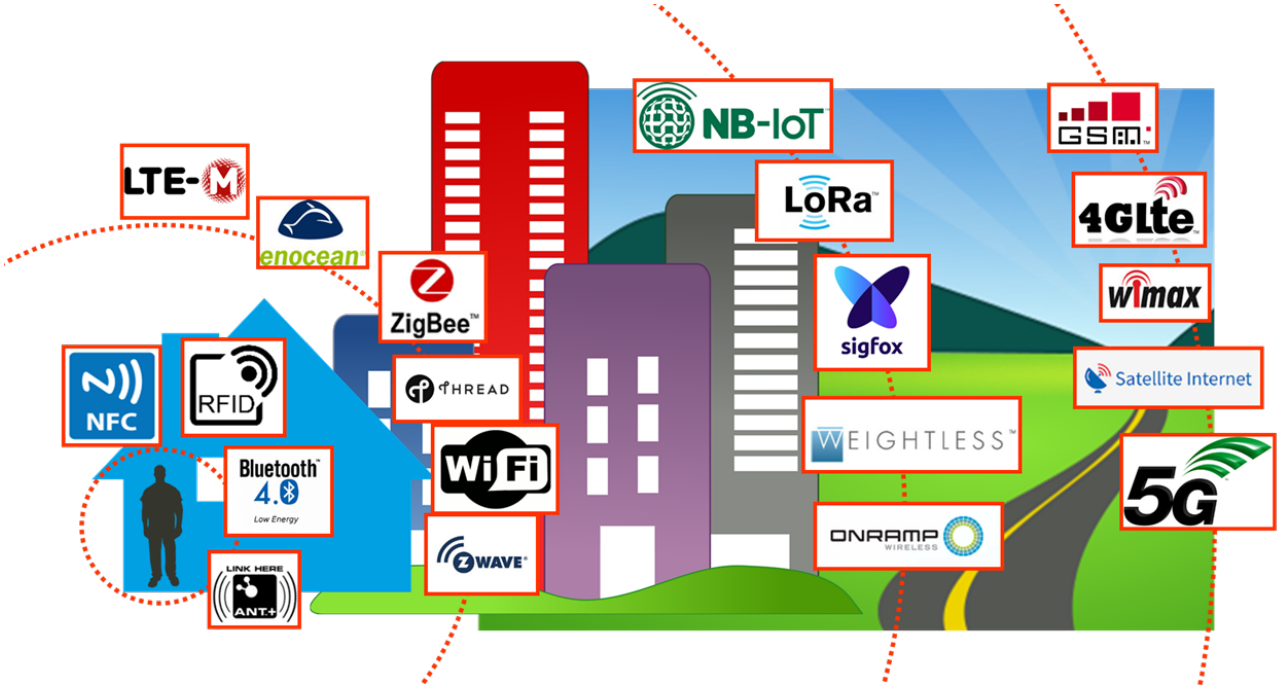

Figure 1: Variety of network protocols that can be used for IoT applications and that are each more or less tailored for different working areas such as personal, residential, campus- or city wide and national or worldwide.

A major challenge is that both for IoT network protocols and for IoT platforms the current market offers a variety of options, as standards are still evolving and maturing. The situation has slightly im-

proved from the wide variety of options that were emerging in 2015 when the Dutch Ministry of Economic Affairs published the Stratix report "Internet of Things in the Netherlands: Applications, trends and potential impact on radio spectrum"[1]. But there are still a lot of non-interoperable IoT technologies and 'Internet of Things' and 'IoT platforms' are still container concepts that keep their position high on the 'peak of inflated expectations' in the Gartner Hype Cycles.

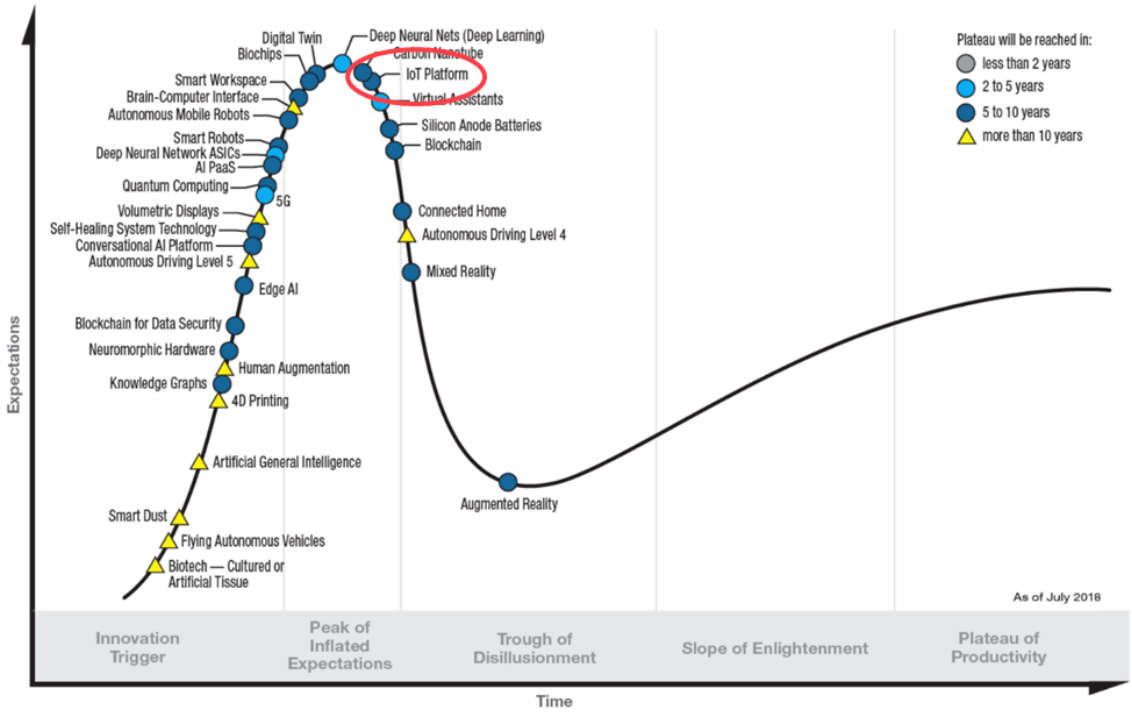

Figure 2: The globally operating market research and advisory firm Gartner regularly publishes the position of trends and buzzwords on the 'hype cycle', the phases between an innovative idea that gets traction and large scale productivity of the innovation. For many years 'IoT' – and in the last few years 'IoT data platform' has been positioned on the 'peak of inflated expectations', as is shown here on the Gartner Hype Cycle in 2018. (source: Gartner)

Ideally an IoT platform - which is in practice an IoT data platform in combination with an IoT network platform - should provide an extendable middleware layer to communicate that simplifies the harvesting of data from sensors. However there are many options to choose from: the IoT suffers from platform fragmentation and lack of technical standards[2]. This not only applies to the many variations in IoT network technologies, but even more to the wide variety of IoT data platforms that emerge. All this makes it difficult to develop applications that work consistently between different inconsistent technology ecosystems.

➔ This paper gives some examples of IoT pilots carried out in the Netherlands and elsewhere, and provides an overview of the options for network protocols and IoT data platforms. Also some options and opportunities for SURF and the SURF member institutes are discussed.

[1] https://www.stratix.nl/internet-of-things-in-the-netherlands/
[2] As is mentioned on https://en.wikipedia.org/wiki/Internet_of_things

2. Devices, networks, data platforms and applications servers are all needed to make the Internet of Things work

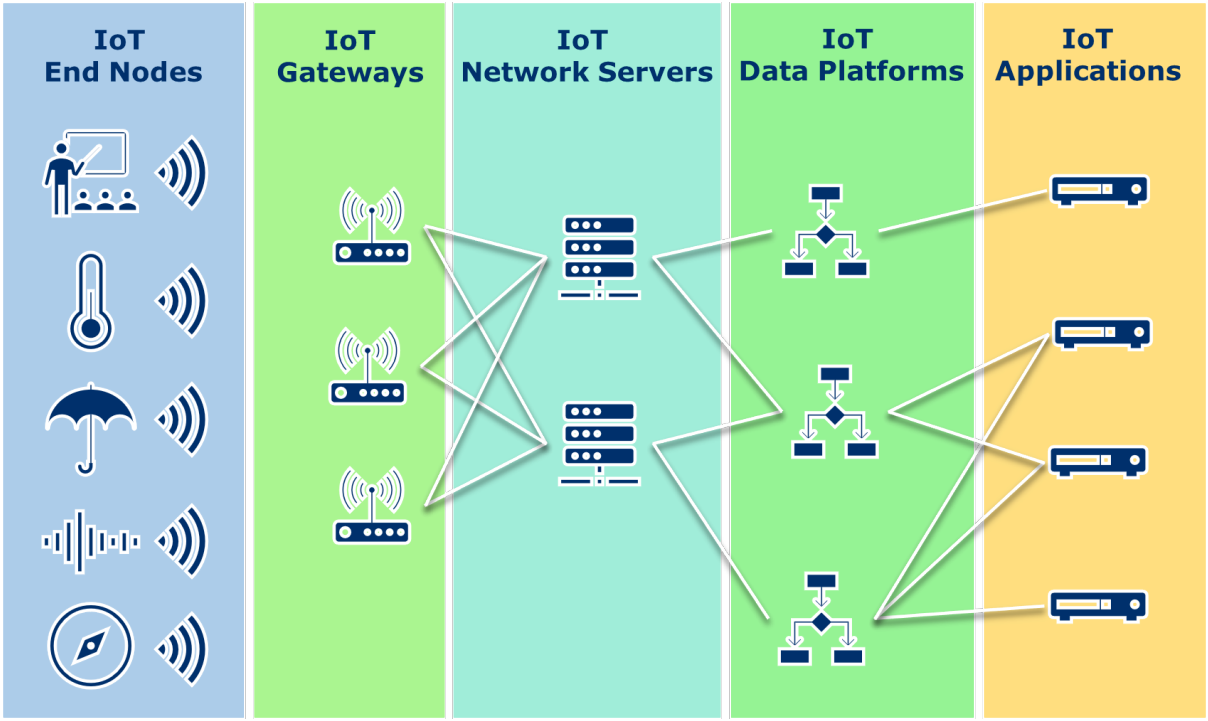

Figure 3: General overview of IoT components

Several elements are needed for IoT applications to work and have to interact with each other. The general concept is illustrated in Figure 3.

In general, IoT devices are cheap, low power devices that send small amounts of data from sensors or receive small amounts of data to perform simple tasks such as opening or closing a lock or give some kind of indication. These devices are connected via a wireless network with concentrators or gateways. An IoT network server is used for routing and addressing information from and to the End Nodes via the gateways. An IoT data platform is generally used to categorize and store data from the end nodes and as a middleware layer for IoT applications.

For each element of the Internet of Things many variants exist

One of the promoted advantages of IoT is the reuse of data and use of 'big data' applications that can combine the data of different sources. In practice many applications nowadays still form isolated islands. But combining these islands also brings challenges: reuse of data and platforms also comes with security and privacy issues. Not only are there several different network technologies that can be used, the choice of which has impact on the devices, gateways and network servers, but there are also several different implementation variants for data platforms and applications.

➔ Combinations of network servers, data platforms and applications, and collaboration and data exchange between them vary greatly between IoT solutions.

3. Successful pilots show the potential of IoT networks and applications for SURF member organisations

Several research projects and pilots at universities and research institutes in the Netherlands aim to find ways to reuse and share network resources and data platforms for IoT applications. A number of examples:

The Smart Campus in Groningen

SURF, together with Hanzehogeschool in Groningen, is running a pilot with IoT in one of Hanze's campus buildings. About 1.000 sensors measure occupancy, temperature, CO2, humidity and other variables such as sound levels and movement intensity.

The main goal of this pilot is to find out which sensor and data is useful to improve quality and flexibility of education and research. But the pilot also aims to experience the impact of operating a more scaled up IoT network, managing many sensors inside campuses and develop a data platform.

Figure 4 Sensor Smart Campus Groningen

More information can be found on: https://www.surf.nl/smart-campus

Data Value Center- Smart Industry (DVC-SI)

SURF is participating in DVC-SI, an initiative in the Province of Noord-Brabant, the Netherlands to stimulate innovation via data driven projects, workshops and training. The initiative is also part of the Dutch smart industry program. This program is currently being put into practice within so-called field labs. Field labs are practical environments within which companies and knowledge institutions are developing, testing and implementing Smart Industry solutions. SURF is delivering Cloud-, IOT -and 5G knowledge for fieldlabs on the Brainport Industry Campus in Eindhoven.

For more information see: https://www.smartindustry.nl/datavaluecenter/

The Green Village in Delft

The Green Village's goal is to accelerate the development and implementation of radical innovations by providing an environment or 'field lab' where technologies and partners can cooperate that otherwise not likely would have been combined. The innovative use of IoT technologies and 5G networks is one of the research areas.

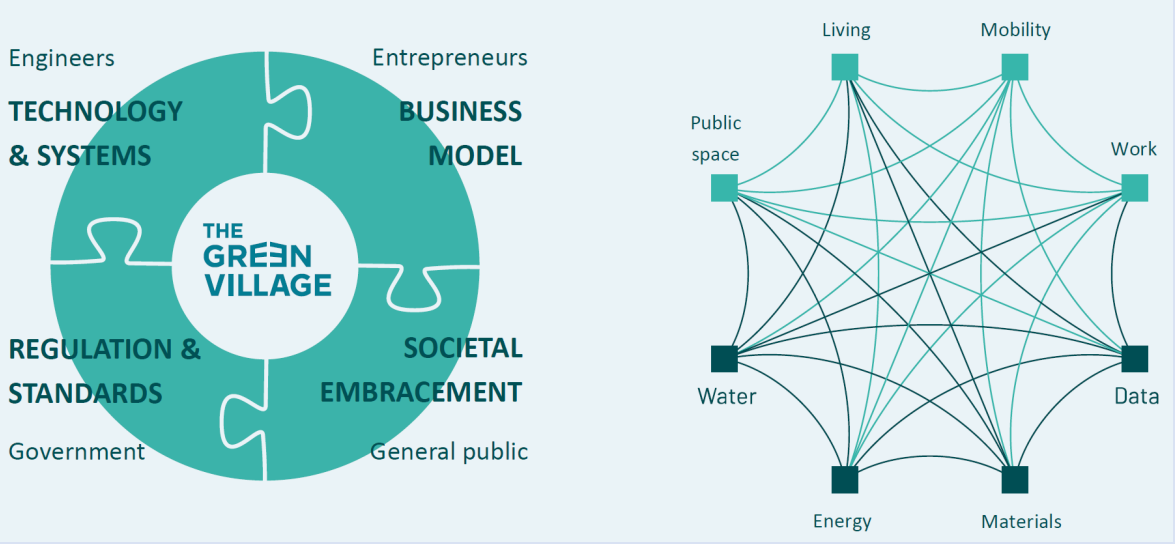

Figure 5: The Green Village research, stakeholders and application types

See https://www.thegreenvillage.org/ for more information.

Research and demonstration projects show the need for open and extendable network and platform components to enable reuse and innovation

Twente47 IoT projects

In the Twente region University of Twente and Saxion Hogeschool cooperate with province, municipalities and companies in a number of IoT pilots, varying from adaptive rostering and measuring occupation of class and conference rooms, to water management using sensors that communicate using a LoraWan network.

See https://www.twente47.online/projecten for more information

SPIN: Security and Privacy for Inhome Networks, from SIDN labs

The SPIN open source platform from SIDN labs was designed to protect users and the Internet against unsafe IoT devices with respect to security, privacy and safety issues.

See https://www.sidnlabs.nl/nieuws-en-blogs/spin-a-user-centric-security-extension-for-in-home-networks for more information.

Wireless Leiden IoT in Leiden

In Leiden the accelerator project from Economie071 drove a city-wide IoT project, together with several partners including Wireless Leiden, SURF, Leiden University and LUMC, Hogeschool Leiden, ROC/ID College together with the Municipalities of Leiden, Voorschoten, Leiderdorp, Katwijk, and Oegstgeest.

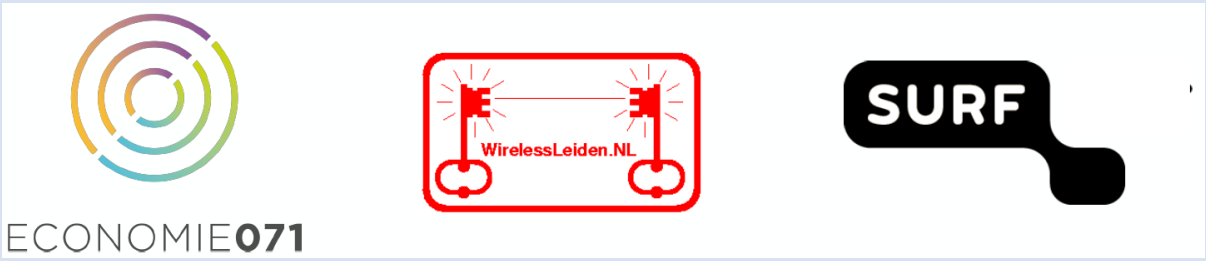

Figure 6: Stakeholders in Wireless Leiden Internet of Things

The presence of Wireless Leiden in the region allowed for a quick deployment on various sites to host LoRaWAN gateways, installed together with SURF in the Economie071 project. This allowed for both research on the LoRaWAN network (coverage) inside an urban area, and for several projects using the infrastructure.

The IoT network in Leiden and surroundings formed a basis for a number of pilots and research projects around subjects such as smart streetlights, indoor classroom environment variable measurements, smart agriculture, smart retail and smart cities. In many of these projects students of universities and other education and research institutes were involved.

See: https://www.economie071.nl/projecten/internet-of-things/ for more information.

5Groningen

5Groningen is an initiative of Economic Board Groningen (EBG) and has been made possible by 10 partners, including SURF. Within the 5Groningen initiative, a number of (regional) 5G use cases are being developed and tested with pre-5G technology, including IoT-technology. One of the use cases that stood out, was about the “smart potato”, where the external conditions of the potato were being monitored.

Figure 7: Logo 5Groningen

Fore more information, see: https://www.5groningen.nl

Regional (data) hubs in the Netherlands

SURF and Commit2Data are connecting regional data initiatives (data hubs) on a national level. A regional (data) hub exists, in most cases, of universities, local government and local business. Currently there are six datahubs forming the "Dutch Coalition of Data Innovation Hubs". Wageningen Data Competence Center (WDDC); Big Data Center Almere (BDVC); Twente 47; Data Value Center- Smart Industry (DVC-SI); Big Data Innovation Hub Zoetermeer; The Data Federation Hub (Groningen).

Every datahub has an area of (data) expertise or is organized around certain themes like energy, water, high tech, urban, agriculture, logistics, etc. SURF facilitates these by organizing meetings, sharing lessons learned and best practices and by enabling IOT/5G solutions, technical infrastructures and data expertise.

For more information see: https://commit2data.nl/coalitie-nederlandse-data-innovatie-hubs

➔ In the Netherlands many research projects use LoRaWAN. Some are experimenting with 5G technologies which are still under development[3]. A wide range of IoT applications are being developed and tested. Applications that use indoor climate monitoring or room occupancy seem to be popular as such applications are the subject of several pilots.

[3] In the Netherlands LoRaWAN and 5G seem the most popular IoT research network technologies. In other countries such as Denmark other network technologies such as Sigfox are also used, see for example http://www.nordic-iot.org/events-page/.

4. Which IoT network technology to use?

In some ways it is quite amazing that even after a decade or maybe even two decades of discussion on networking technologies for IoT there is still no easy consensus, such as there appears to be in computing. In computing it's "simply" Wi-Fi of some flavour to connect most terminals, Ethernet to connect servers and access points, Bluetooth or USB to connect peripherals to terminals, and GSM/LTE to connect while mobile.

It's so easy and standardized that we expect it to work everywhere. Of course a deep dive will show there may be issues with versions, flavours, supported standards etc. But at the end of the day most consumers don't have to think about the networking technology, because it will just work with every-thing else they have at home.

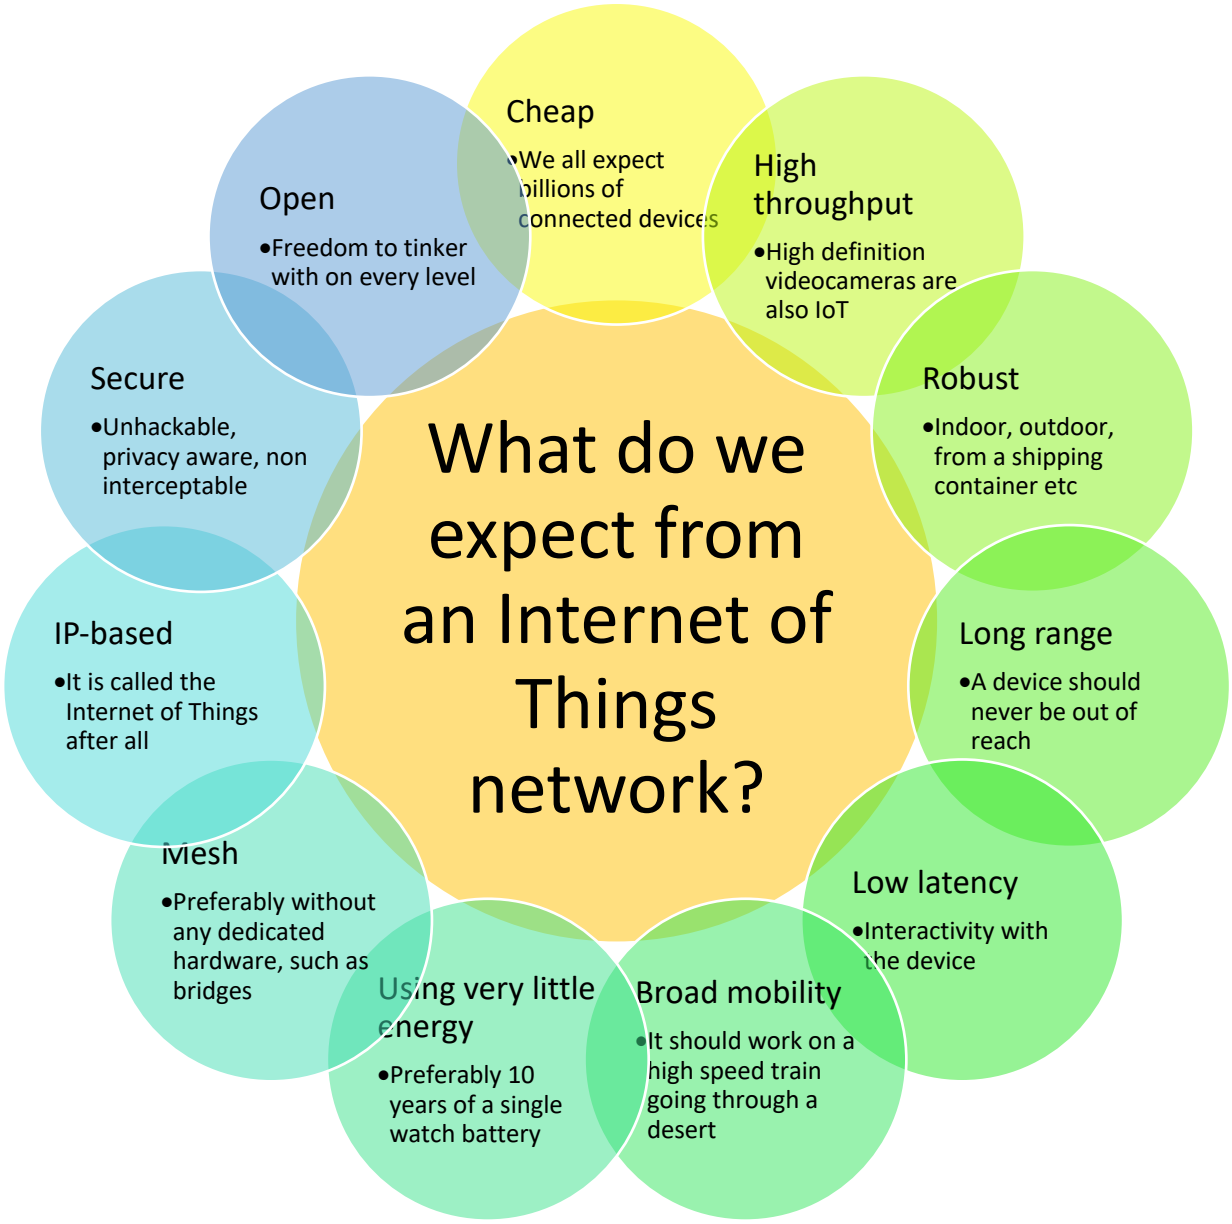

Figure 8: Expectations of an IoT network

The main cause (next to politics and competition) is the varying demands IoT users have. The ideal IoT networking technology would be the one that supports a variety of requirements, as was illustrated by Figure 8.

And of course you can't have all of these at the same time and so there are many trade-offs and those trade-offs turn out to be different for user, manufacturer, protocol and application.

The table below describes the main differences of the most common IoT network technologies currently available. Operating range and speed or net throughput turn out to be the two most defining elements of most technologies. It is for this reason that in this section we're grouping technologies based on range and speed. Also the power consumption is important as many IoT devices need to be able to operate on batteries for considerable lengths of time.

Table 1: Comparison of network technologies with regard to Power consumption, transmission speed, network type (Personal Area, Local Area, Wide Area or Low Power Wide Area), Operating range, Frequency band and whether the radio frequency band can be used license free or is licensed. (Source: Stratix based on information from PredictableDesigns, http://www.iea-4e.org/document/384/energy-efficiency-of-the-internet-of-things and other sources)

	Power	Speed	Type	Range	Frequency	License
Bluetooth	Low	2-3 Mbps	PAN	50m	2.4 GHz	no
Bluetooth LE	Very low	1 Mbps	PAN	50m	2.4 GHz	no
ZigBee	Very low	250 kbps	PAN	100m	915 MHz / 2.4 GHz	no
Z-Wave	Very low	100 kps	PAN	150m	868 MHz / 908 MHz	no
6LowPAN / Thread	Very low	Low	PAN	100m	2.4 GHz	no
Wi-Fi	High	100-250 Mbps	LAN	100m+	2.4 GHz / 5 GHz	no
GSM/GPRS	Very high	Moderate	WAN	35 km	900 MHz / 1.9 GHz	yes
LTE / 5G	Very high	High	WAN	Long	Various	yes
LoRa / LoRaWAN[4]	Low	27 kbps	LPWAN	10km+	868 MHz	no
Sigfox	Very low	100 bps	LPWAN	20km+	868 MHz	no
NB-IOT	Moderate	250 kbps	LPWAN	20km+	Various	yes
LTE-M	Moderate	1 Mbps	LPWAN	Long	Various	yes

Figure 9 shows a visual comparison of the typical operating ranges and speeds of the technologies mentioned. The graph provides a comparison in orders of magnitude using circles for each technology. The real operating ranges depend on aspects such as geographical features, line of sight, indoor or outdoor devices etc, and the actual achievable operating speeds depend on several aspects such as the number of devices and gateways that are being deployed, and may also vary per generation of technology standards.

[4] LoRa is the wireless technology, LoRaWAN represents the datalink layer needed for (larger) networks. The LoRa alliance supports LoRaWAN and interoperability of LoRaWAN products.

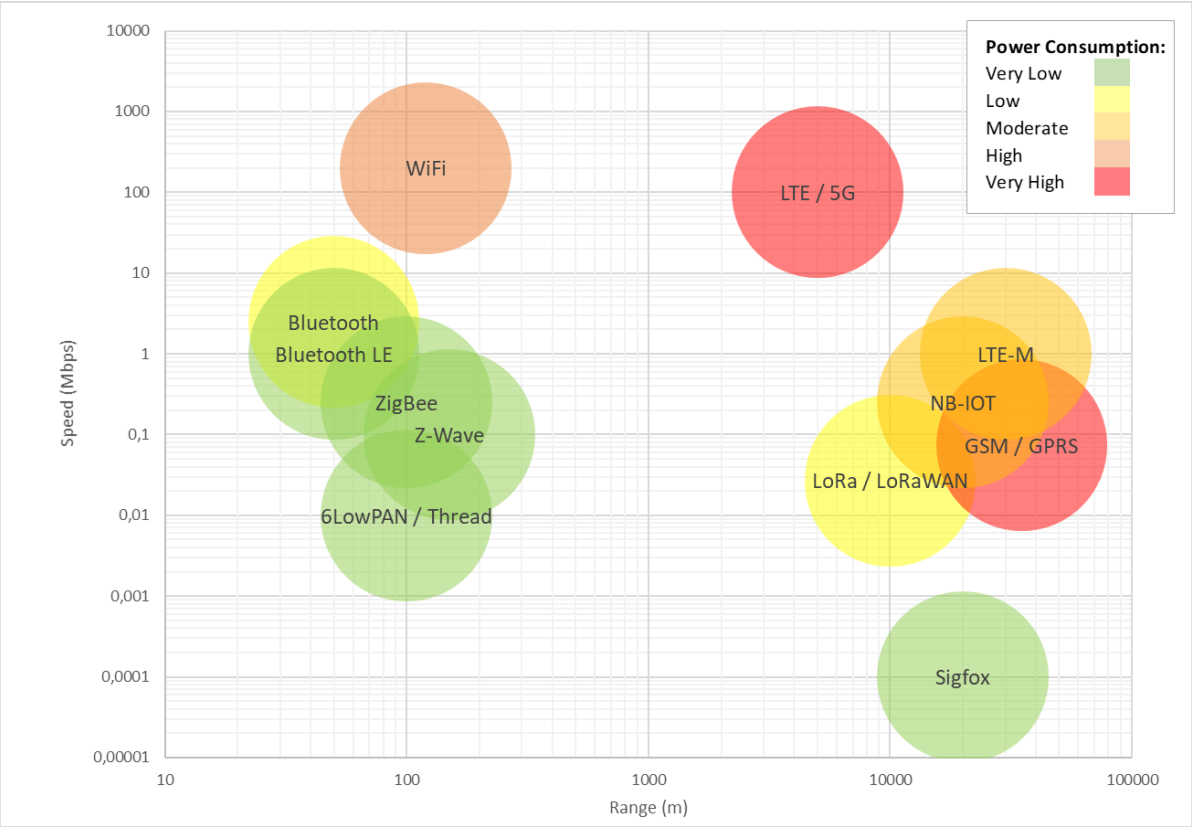

Figure 9: Comparison of network technologies
(sources: http://www.iea-4e.org/document/384/energy-efficiency-of-the-internet-of-things, PredictableDesigns, Stratix)

When we consider the typical use cases and data use of applications and we consider the typical operating environments and mobility requirements of the elements of the applications, the network technologies can be grouped in most likely combinations. The matrix in Table 1 shows these combinations.

Table 2: Use cases and network technologies

Use case	Some readings a day	interactivity	video
Home/small office		Z-wave, Zigbee, Bluetooth Mesh, Thread, 6LowPan	Wi-Fi
Large office/campus	LoraWan, Sigfox	LoRaWan, Z-wave, private LTE	Wi-Fi
Nationwide (populated)	LTE-M, NB-IoT	LTE-M	LTE
Global[5]	LTE-M, NB-IOT	LTE-M	LTE

The demands of SURF and its customers put on wireless technologies for the IoT are different than those of consumers or large companies. The result is, unfortunately, that it will be hard to use a "standard" such as Bluetooth, Wi-Fi, Ethernet, 3G/4G. SURF customers generally operate a number of larger buildings, sometimes in a campus setting, sometimes distributed over one or more cities.

[5] For global operations several satellite technologies are also an option, but in general these solutions are costly and less scalable although they can provide a feasible solution in areas where no other networks are present. Several low orbit satellite systems are in development that may potentially change the economic constraints economy for IoT solutions in the long term.

These are semi-open locations and the organizations generally cooperate with local and global communities to develop and refine new uses and applications. For SURF these constraints mean that it wants a cheap technology which costs only a few Euros per device to deploy, requires little certification and where it doesn't matter that students or staff experiment with and break a few units, but the technology does offer range and doesn't require a new network node every few meters. In the next paragraphs we will show the capabilities of these technologies, because they will often be mentioned as viable alternatives to LoRaWAN. On the other hand it will also show why sometimes other protocols are needed and used.

In and around the home / small office

There are quite a lot of wireless technologies for in and around the home and small office, most of them make use of the ISM bands in either the 433MHz, 868MHz or the 2.4GHz and 5GHz. Wi-Fi is present everywhere, however it is too power hungry and technically complex for many battery operated IoT applications. Even for many applications where power is available in the home, there is a focus on using lower-powered technologies. For example although there are Wi-Fi-lightbulbs, the market leaders, Philips Hue and Ikea focus on Zigbee.

The 433MHz band is popular for cheap wireless in and around the home, such as weather stations. It's also popular with homebrew IoT applications. The range in this band is quite good and transmitters and receivers tend to be cheap. However, this comes at a cost: most technologies in this band are unsophisticated in terms of security (no authentication or encryption), spectrum use (no bi-directional communication or listen before send) and proprietary. In practice this means anyone can sniff the data and read it. As long as you know the right identifier for device, it will accept any command send and it is often unable to communicate state back nor able to set up a mesh. As a result even enthusiasts consider this band and the technology that currently uses it, unfit for serious projects.[6]

The 868-band ranges from 863MHz to 870Mhz, however the license exempt applications that are of particular interest to this report are in the high end of the band. Other applications such as RFID can be used in the lower part of the band. Figure 10 shows how differentiated the use of the band is with microphones in the bottom part, RFID in the middle and other (alarms and home automation) in the higher end. The fragmentation of the higher end of the band in terms of power and duty cycle and makes it even harder to implement.

Efforts to introduce "one device fits all" smart home hubs, such as the Smartthings orHomey 'home hubs' that would take away much of the complexity from the user have seen some adoption, but these hubs don't work with every device or protocol. And with Google's recently termination of the "Works with Nest" programme, the fate of some of these hubs is in doubt as well. Other manufacturers might follow suit and lock their devices from third party direct access.[7]

[6] Doe-het-zelf-domotica
Het nieuwe tweaken https://tweakers.net/reviews/3911/4/doe-het-zelf-domotica-het-nieuwe-tweaken-433mhz.html

[7] https://tweakers.net/nieuws/152554/google-stopt-op-31-augustus-met-works-with-nest-apis.html

Most technologies for in and around the home are based on IEEE 802.15.4. It enables a wireless mesh between devices. The specification operates on one of three possible unlicensed frequency bands:

- 868.0 – 868.6 MHz: Europe, allows one communication channel (2003, 2006, 2011[5])
- 902 – 928 MHz: North America, up to ten channels (2003), extended to thirty (2006)
- 2400 – 2483.5 MHz: worldwide use, up to sixteen channels (2003, 2006)

Unfortunately, the specification is rather limited and omits many of the vital elements to create a thriving ecosystem. As a result implementers had to build a number of higher layers to actually make the technology usable for end-users.

The clear market leader for short range, interactive IoT is Zigbee. It is a low-power, mesh capable IEEE 802.15.4 based technology that uses the 2.4 GHz and 868 MHz bands. Though there is a Zigbee alliance behind the protocol, this alliance hasn't been able to create a uniform well-supported protocol. There are a number of different flavours of the Zigbee protocol, which are not necessarily interoperable. The protocol also isn't transparent to the application. Where a Wi-Fi-device will connect to almost any Wi-Fi-access point, there is no guarantee that a Zigbee device will be able to connect to a Zigbee bridge. Zigbee being in the same band as Wi-Fi means that it may suffer from interference and range can be limited by the same factors that limit Wi-Fi-range.

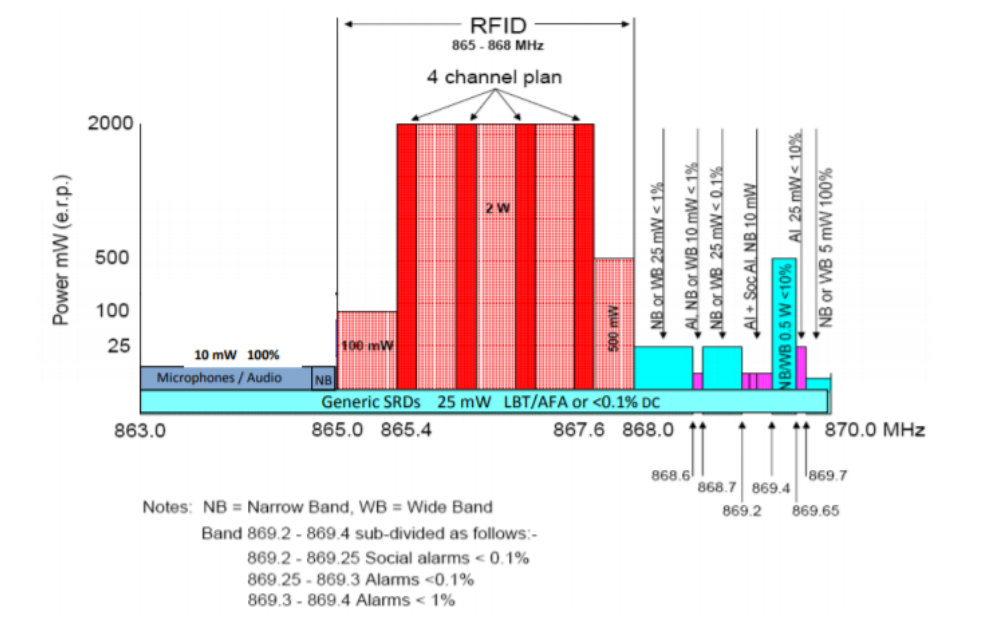

Figure 10: Overall band plan for 863-870 MHz (source: https://www.ofcom.org.uk/__data/assets/pdf_file/0025/38095/final_report.pdf)

In summary there is no standardization yet in the sensor networks for home applications. Wi-Fi dominates the high-end, but there is no standard that rules the low-end. Purely from a technical view, 6LoWPAN with its support of Ipv6 looks good, Zigbee has the most brands, Bluetooth Mesh has the

brand recognition, but none have yet won the market. New brands are even introducing new proprietary protocols and gaining market share.[8] It is very uncertain if there will be any drive towards standardization.

Z-Wave[9] is competing protocol to Zigbee, that is also mesh, but makes use of the 868MHz band. It used to be a more closed specification than Zigbee, with one company promoting the technology. As a result its interoperability between manufacturers and applications is wider. Unfortunately it works in different bands in different countries and regions. The protocol is quite popular in the United States and is used by some lighting and security brands. However, it has received less support from major brands in the EU, likely because of the reduced availability of bandwidth there. As a result it is less widely used than Zigbee.

Bluetooth was developed to connect peripherals to computers, but has developed into the default protocol to connect smartphones and computers to a wide group of personal devices. Earlier versions were deemed too power hungry, but this changed with Bluetooth Low Energy that is now part of every smartphone, smartwatch etc. As of 2017 there is a wireless mesh capable Bluetooth specification, that makes use of Bluetooth Low Energy, that is more aimed at IoT devices, instead of at peripherals. So far the standard hasn't seen large deployment, however the list of certifications shows that Xiaomi and Samsung are deploying the technology.[10] Applications appear to be speaker/home-hubs and connected lighting. There is some integration of Bluetooth LE in Wifi Access Points, for example for location services, which might increase the use of Bluetooth LE in the future.

A mention should also be given to 6LoWPAN and Thread. Based on the same 802.15.4 standard that Zigbee uses, 6LowPAN is an Ipv6 capable alternative to Zigbee. In theory Zigbee devices could even have 6LoWPAN added through a software update. Thread is a further specification over the open 6LowPAN stack which is promoted by Google Nest. Though 6LoWPAN and Thread are in use, they haven't broken through on a massive scale. Some implementations, such as the tado smart thermostat, don't expose the IP-addresses to the outside world, but instead use Ipv6 only for local addressing. It is unclear however what will really happen with Thread and 6LoWPAN in the market. There are few news stories/twitter chatter about the protocols, and it wasn't supported in Google's Home line of devices[11], until the Nest Hub Max in May 2019. However, Apple did join the Thread group in 2018, so something might be forthcoming.[12]

Large office/campus

To deploy sensors and other smaller devices across a large office building or campus is hard, especially in terms of wiring. Even inside buildings in many places it is preferred to use devices with long lasting batteries instead of devices that have to be plugged in a 230V socket. The home networking protocols generally work in the 2.4 GHz and therefore don't have the required range to serve an entire building or campus. It requires a network of IoT-access points of some sort. The market for these

[8] https://support.wyzecam.com/hc/en-us/articles/360026028731&?section=zigbee-zwave
[9] The current owner Silicon Labs has decided to open Z-Wave up
[10] https://www.bluetooth.com/bluetooth-technology/topology-options/le-mesh/mesh-qualified/
[11] https://www.cnet.com/news/nest-hub-max-google-debuts-10-inch-assistant-smart-display-for-230/
[12] https://www.the-ambient.com/guides/thread-group-explained-908

types of devices is however quite limited. In the United States Zigbee and Z-Wave have found some application, particularly in the 900 MHz band, but much of the automation is still wired. This might be less flexible, but its reliability in the operation of doors, thermostats, lights etc. is solid. Automation generally takes place only when there is a pressing business need or when there is renovation. Supermarkets are an example of an environment that has gone a different route. For smart prices tags and electronic shelf labels in many supermarkets make use of optical communication. Such technologies have currently limited applicability outside specific logistical applications. There are however developments to bring LiFi further in the market with higher speeds, so this could be an outside contender[13].

For a while these longer distance applications were thought to be the area where mobile operators would excel in. These firms were talking quite a lot about the Internet of Things, but customers found that their business needs weren't addressed. Particularly the need for low-cost, technologically simple and robust, energy efficient and secure devices wasn't met by the mobile network industry. GSM was used for all kinds of applications, however these needed to have a power source and implementing them wasn't cheap. Building indoor/campus networks with GSM also required the cooperation of an operator and as a result was in many countries impossible. Until recently the Netherlands was one of the few countries that allowed private cellular networks for indoor and campus type applications. These Private GSM/LTE networks were built using license exempt (or a light license) spectrum in the former DECT-Guardband (1800MHz) and in the 3.5GHz band. In the last two years other countries have followed the example of the Netherlands and have cleared spectrum for Private LTE/5G in the 3.5GHz. In general between 50 and 100MHz will become available.

Several competing IoT network technologies and standards are still evolving and competing. At this moment, LoRaWAN and potentially Sigfox and LTE-M seem to provide the closest match to the requirements of the SURF community

The result of the mobile sector's apathy was that several firms developed low-cost alternatives aimed at IoT applications. These so-called Low-Power Wide Area Networks (LPWAN) found quick attention from enthusiasts around the world. One of the first and most influential to propose a technology was Neul with its Weightless specification. Neul was acquired by Huawei and many of its ideas live on in the NB-IOT standard of the 3GPP. Most other LPWAN technologies, such as Lora and Sigfox focus on the ISM bands, which do not require licenses. An issue with this spectrum and this type of application is that there are many competing uses of ISM bands, while at the same time it carries quite far, so noise, interference and collisions have all the to be taken into account. [14]

[13] See https://www.stratix.nl/optical-wireless-communication-internet-via-licht/ for the Stratix report on trends in Optical Wireless Communication

[14] https://www.thethingsnetwork.org/forum/t/universal-lora-wan-gateway-limitations-because-physics/1749

LoRa is a proprietary LPWAN technology of Semtech. It makes use of the ISM bands, such as the 868MHz in Europe. Though the technology is proprietary there are some different vendors for chips. The deployment of LoRa is not limited to one network operator, there can be multiple and it is even possible to set-up a private network. The result has been that there are proprietary commercial networks, such as KPN in the Netherlands, as well as open non-for profit networks such as 'The Things Network' operating in parallel in the same spectrum. On The Things Network a duty cycle of 30 seconds is imposed, which means devices only send once per 30 seconds. Particularly the ease of use with which users can experiment has led to a wealth of initiatives and innovation.

A significant benefit of the implementation of LoRa is that there can be multiple competing gateways and networks, with different business models. It's possible to roll-out on a campus using a private node and not interconnecting with other networks. It's also possible to buy LoRa service from a mobile operator. For SURFnet this means that LoRa can be deployed in situations, where there are just a few students hacking away at an idea, to supporting the operation of an institute across multiple campuses, all the way to aiding a researcher who wants to deploy sensors across the whole country. The less than guaranteed way of operating however means that LoRa is less useful in situations that require interactiveness or high reliability.

Sigfox was promoted as a more open alternative to some proprietary LPWAN technologies. This was however primarily true for the silicon manufacturers, as several other aspects are limited to the Sigfox firm, that acts as a global network provider. In the Netherlands AEREA is the network provider. Sigfox is limited in its application by allowing only 140 upstream messages per day and 2 downstream messages. As a result it isn't capable of supporting interactive applications the way Sigfox broadcasts is, by sending the same message three times, thereby reducing how much can be sent, but being less susceptible to interference.

Nationwide and Global

Once nationwide or global coverage is necessary for a project, mobile networks in licensed bands generally have the upper hand. However, LoRaWan can operate in this category as well. Its range is exceptional and particularly in the Netherlands, there are few obstructions to the network. It is therefore a reasonable assumption many locations in the country can be reached from the site of a SURF customer. Therefore it might well be possible to have a "near"-nationwide LoRaWan network that is fully SURFcontrolled. If necessary the networks of KPN or "The Things Network" can serve as backup. This way there is quite some room for experimentation.

Other networks for nationwide IoT coverage are based on LTE, either LTE-M or NB-IoT. LTE-M is supported by all three mobile networks in the country. NB-IoT only by Vodafone and T-Mobile. LTE-M can function as a normal IP-based network with some tricks, such as long sleep times, to extend batter life to several years per device. NB-IoT promises even better battery life and range then LTE-M but does so at the cost of a reduction in interactiveness. The implementation of T-Mobile for example expects users to send messages at most once per 10 minutes. LTE-M is therefore a better fit for applications where service engineers or researchers occasionally have to interact with the device, have to update it, control something on location etc. We haven't mentioned GSM here as the technology is said to be decommissioned by 2025.

On a global scale most mobile technology can work everywhere people normally live and work. LTE-M and NB-IoT do suffer from reduced roaming options as the technology hasn't been rolled out everywhere yet and not all roaming functions are globally supported. Mobile technologies are hampered by high roaming costs in some countries.

This can sometimes be dealt with by using a global SIM, with multiple operator profiles or an eSIM. Unfortunately both solutions aren't supported everywhere by all operators. If it really has to work everywhere, satellite is the only option. There are some promising technologies for low-power global satellite technology, however those are out of scope for this report.

➔ Many IoT network technologies exist and are evolving. At this moment for low power applications with low data rates in a campus environment, the most suitable network technology that is standard, cheap and uses license free spectrum is LoRaWAN.

5. Which IoT data platform to use?

An IoT data platform should provide an extendable middleware layer that simplifies the harvesting of data from sensors. The purpose of the IoT data platforms in general is to provide a system for data management, storage and sharing that is flexible and scalable, and provides basic and extendable functionality and features for analytics and visualisation of data, and that is able to interwork with IoT networks.

Figure 11: Expectations of an IoT data platform

An IoT data platform is a collection of individual services enabling the functionality mentioned above. These services are well tied together to provide an added value as a full end-to-end solution. A high-level overview an IoT data platform is shown in Figure 12. Important components are a component for data storage and data management consisting of 'distributed log' or 'message broker' technologies to harvest data from different sources and distribute the data to applications and databases, and a protocol set for sending and receiving messages between the datbases, IoT applications and IoT devices.

However, there are many IoT data platforms and message protocols to choose from; the IoT suffers from platform fragmentation and the use of different standards, and for IoT data platforms this is even

more the case than for the network technologies, as a wide variety of IoT data platforms have emerged. There are several distributed log technologies that can be used for the data storage and data management such as Kafka, Amazon Kinesis, Microsoft Azure Event Hubs en Google pub/sub[15]. There are several (realtime) message protocols that can be used, such as http, MQTT, XMPP, STOMP, AMQP, WAMP[16], that are relatively comparable in capabilities but do not interoperate. All this makes it difficult to develop applications that work consistently between different inconsistent technology ecosystems.

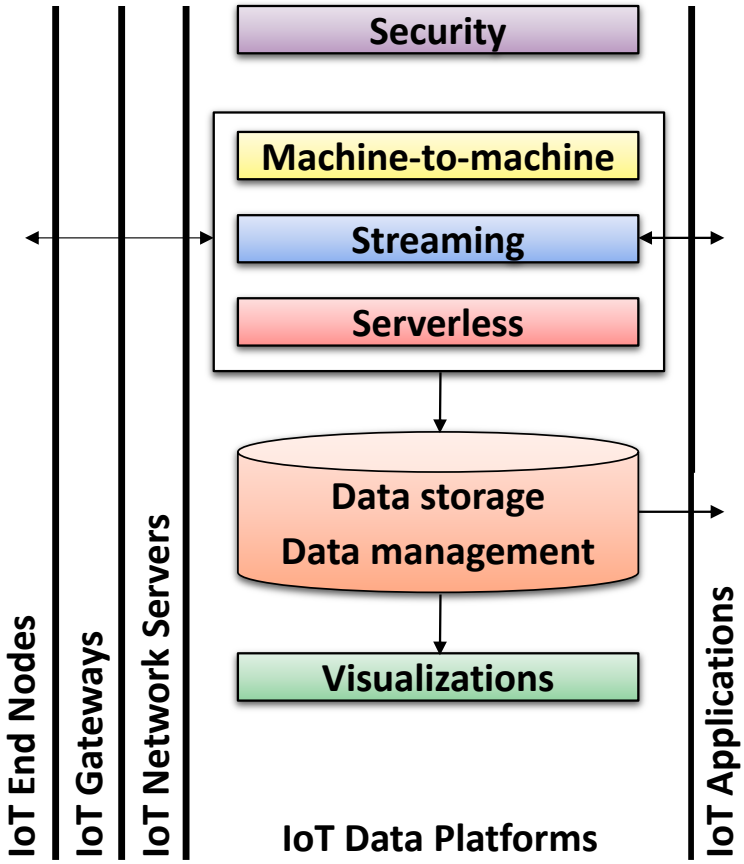

Figure 12: Schematic overview of an IoT data platform and its position in the IoT ecosystem

There are numerous IoT platforms developed to serve a specific purpose in a specific domain, e.g. home automation. Such platforms are not flexible enough for most research and development projects. On the other hand, public cloud providers, such as Amazon or Microsoft, offer generic IoT solutions and multitude of other services that can suit almost any project. One can either directly use these dedicated IoT services, or choose to develop a custom solution from independent components. The former leads to larger costs on the longer go, the latter requires skilled resources and initial development investments. The IoT data platform by SURF is positioned in the middle by providing core IoT services with a competitive price tag based on modular components. In addition, there is a possibility to store data in the SURF datacenter.

[15] see for example https://blog.scottlogic.com/2018/04/17/comparing-big-data-messaging.html
[16] see for example https://deepstreamhub.com/blog/an-overview-of-realtime-protocols/

By providing an IoT data platform, SURF can enable and facilitate the use of connected devices in research and education, and support IoT innovation and standardization

There is demand for a data platform from a non-commercial party that aims to be as transparent as possible such that data and software can be re-used by implementing open data exchange protocols and using open source software. Furthermore, by choosing the protocols and components SURF can stimulate technical standardization in the IoT field.

The SURF IoT data platform consists of several core services that enable the sensor data traffic and sharing in a scalable and a secure way. The core components are illustrated in Figure 13:

Figure 13: Core components of SURF iOT data platform[17]

The core functionality of the platform can be extended by additional services based on the needs of particular projects. This can be demonstrated by the pilot IoT engagements by SURF. In the case of the Green Village data sharing platform, the main focus is given to connecting a wide range of IoT devices and offering a convenient way of sharing data from various sources. The former is achieved by supporting multiple communication protocols, the latter is enabled by a data management layer, including data format standardization and metadata extraction (see Figure 14).

[17] Serverless technology makes it possible to implement event-driven actions in the cloud/computer cluster without the need to worry about the resource planning and maintenance on the developer side.

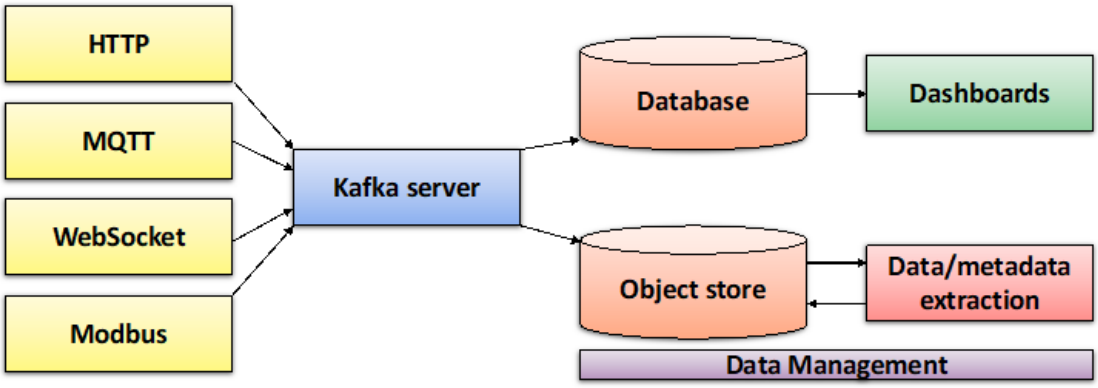

Figure 14: Components of the Green Village data sharing platform.

Kafka[18] serves as the backbone of the platform, providing a secure, highly available and fault tolerant streaming service. Another example is a SURF IoT data platform delivered for a citizen science project, where the requirements are to provide a lightweight platform enabling insights and analytics. Popular low-threshold technologies from the user perspective are chosen in this case (see Figure 15). An MQTT[19] broker and serverless technology provide the core functionality of the platform and Jupyter notebooks give a convenient way for analyzing data and sharing results.

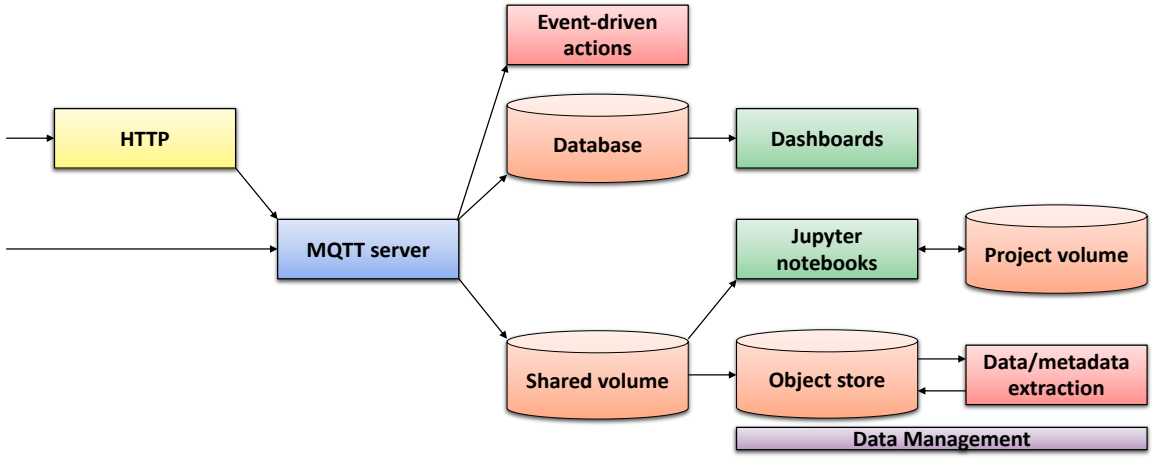

Figure 15: Components of a SURF IoT data platform delivered for a citizen science project.

➔ The world of IoT data platforms is even more diverse as IoT network technologies are. Institutes ideally should be able to make their own choice for an IoT data platform, but a basic open and extendible alternative should be available. Where possible such a data platform should also offer basic functionality while taking important issues such as governance, security, privacy and legal issues into account.

[18] Kafka is a popular open-source stream-processing software platform for handling high-throughput real-time data feeds with low latency. https://kafka.apache.org/

[19] MQTT is a lightweight messaging protocol suitable for IoT devices, optimized for high-latency or unreliable networks. http://mqtt.org/

6. What security and privacy considerations are important for IoT?

Security and privacy are the Achilles heel of many IoT projects and they are in many ways intertwined. The goal of many IoT projects is to generate actionable (sensor) data and then to base automated or user controlled actions on that data. The issue with both security and privacy of these data and actions is that it is really hard to define up front, what the privacy and security impact is of any of these data. Most projects do make some assessment of the data they collect and the security issues and threats, however their thinking is generally goal oriented and not adversarial. The thoughts are therefore towards: What data do I need to reach my goal? What systems do I need to control? Based on these analysis, procedures for authentication, software updates, processes and procedures for day to day use etc are defined. The thought process isn't adversarial; What could I do with this data/system/process/procedures if I wanted to use it against someone or something? Data that appears to be benign in the context for which it is used, can be used adversarial in another context

An example to illustrate: Data that is collected on air quality and use of class rooms can be combined to see roughly what percentage of students attend a lecture by a teacher. These data would normally be collected to improve ventilation, control lighting, save energy etc. However combining the data might be used to rank teachers at an institution in how well they perform in their task of teaching. If attendance decreases over time, that might indicate students aren't interested. An unscrupulous administrator might combine these data and use them in ways that aren't intended and the data might not even indicate how well teachers perform. Bad intentions with good data still leads to problems.

If adversarial thinking is applied then every part of an IoT system and any combination of these parts can become vulnerable to attack. In recent years it has become clear that even errors in CPUs[20], baseband chips etc can be exploited to be abused higher up in the stack. There hardly is any part of IoT system that isn't software controlled and therefore isn't vulnerable to injection of some code or manipulation in some way or another. Many of these elements are completely out of reach to even the manufacturer of the device that's using them. Even when a vulnerability is found in parts of the system that a manufacturer can control, it is often hard to distribute updates to IoT devices as these are often not capable or designed for large updates to core files and systems in the device. Sending a patch over LoRaWan for example might be rather complex and saturate the link for a significant amount of time.

Security and privacy by design remain important for all IoT components. Security and privacy form a challenging combination with some key features of IoT applications such as low data rates, open systems and reuse of data

---

What makes IoT even more vulnerable is that it often interconnects a number of diverse and distributed systems that each have their own function, but the combination of them can lead to wanted and unwanted and even emergent behaviour. *Emergence* occurs when an entity is observed to have properties its parts do not have on their own. These properties or behaviours emerge only when the parts interact in a wider whole. It is really hard to design against unwanted emergent behaviour, that somehow causes security or privacy risks. Designers don't know what new parts will be added to the system in the future, what devices can read the sensed data and what actions can be based on them.

[20] Think Intel Rowhammer or BadUSB https://threatpost.com/new-research-same-old-problems-with-badusb/109398/

The initial thought of developers and academics is to share data widely, because that is what helps them further. However they can't control for emergent behaviour as a result.

One of the reasons that was brought forward of why Google shut down the Works with Nest programme is its fear of emergent behaviour on privacy and security. In Works with Nest users could authorize other devices in their home to work with Google's Nest systems directly. For example to read state or direct an input, for example using If This Then That (IFTTT) as a result there was no control anymore over who would have access to user data, to protect it and to protect against unwanted emergent behaviour. SURF customers will have to face this threat too

There is not one solution to the privacy and security threat. They are two faces of the same problem. If the hardware and software are insecure, data can be gathered and accessed by unauthorized persons and systems, but even when the system works as initially designed, an actor can use the system and combine different systems to compromize the safety, security and privacy of others in ways unforeseen by the designers and developers of parts of the system. There are a number of security and privacy practices out there that help mitigate basic problems, such as updating and patching, authorization controls, limiting data collection etc., etc. All of this should be done. But fundamentally only continuous monitoring for anomalous use and adversarial thinking can keep the system of systems safe. This is hard for developers who are focussed on making things work, not on how they can abuse the same systems.

However three questions are particularly important to consider when choosing the technology components for your IoT application:

- How vulnerable are the components for Denial of Service (DoS) attacks?
- How vulnerable are the components for (device) hijacking?
- How vulnerable are the components with regard to negligence with regard to installation of the necessary software of firmware updates?

If we compare NB-IOT, LoRaWAN and Wifi in this respect than with NB-IOT much of the responsibility to avoid security risks can be put with an operator, but with a cost. For wifi security is primarily the responsibility of application owner and device owner, and wifi and IP theoretically offer sufficient tools for encryption (but they have to be properly used of course). The message sizes and frequencies of LoRaWAN do not offer a huge amount of overhead for encryption so for LoRaWAN some care has to be taken by application owners and device owners with regard to security issues. But examples as SPIN from SIDN show that there are robust and efficient ways to deal with relatively unsafe devices.

➔ Security and privacy are issues that have to be taken care of at each and every level. A uniform approach can offer standard building blocks, interfaces and procedures that help to minimize risks and maximize performance.

7. SURF can improve IoT in the Netherlands by doing the right thing

SURF ant to enable and facilitate a 'common platform and infrastructure' for IoT applications where this can offer common advantages for institutions. Different options can be distinguished with regard to the role SURF can take, as the picture below illustrates.

SURF can help to build a set of open, extendible standard components that can be easily used by institutions in the Netherlands

In the Netherlands a wide variety IoT pilot projects are carried out and IoT attracts a lot of attention from institutes. Institutes want to be flexible in their approach towards IoT, but on the other hand there is a demand for some interoperability, standardization and reuse of components to avoid that every project has to reinvent the wheel with regards to connectivity, data management but also with regard to security and privacy issues.

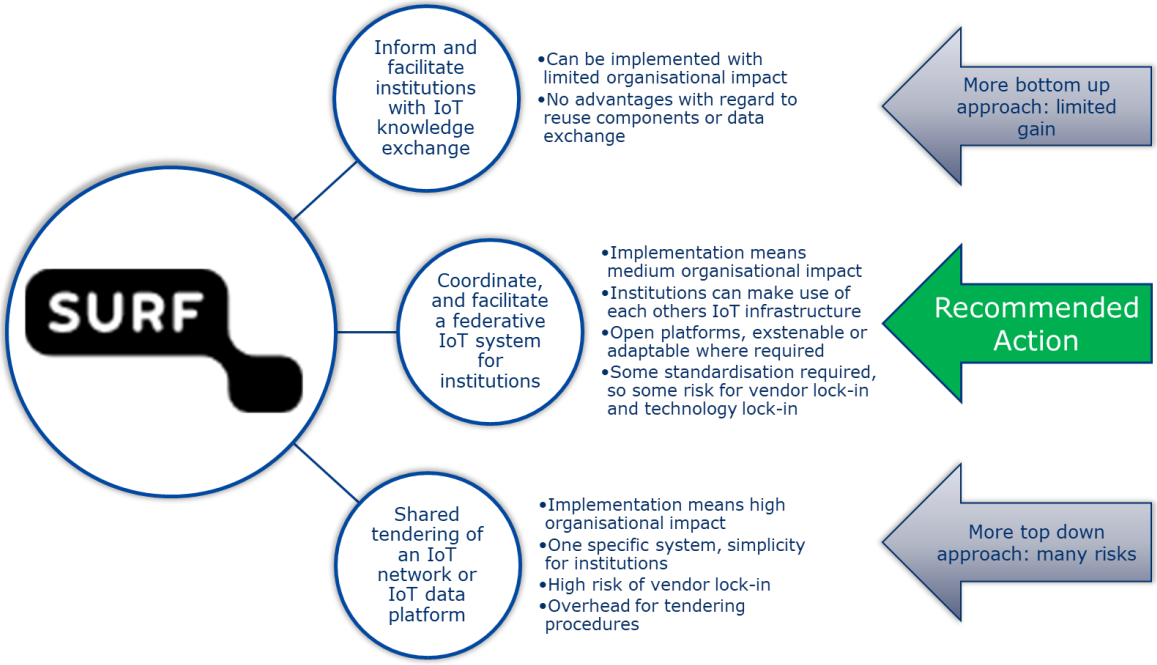

Figure 16: General options for SURF with regard to facilitating IoT

At this stage therefore SURF and the institutions it represent can work on defining an open platform together, experiment and expand and reuse elements, building on each other's successes and experiences, but maintaining flexibility.

This option is to be preferred over waiting and only informing on one hand, and trying to enforce a uniform single operator IoT approach on the other hand:

Only inform and support information exchange on IoT subjects would cost less effort, but does not really help move the development of IoT applications forward. Encouraging or enforcing more uniformity would likely cost more effort and would come with more risks with regard to technology and vendor lock-in. Also, facilitating in a federative system with standardized interfaces but flexibility towards own solutions or variations seems to best fit SURFs ambitions and expertise.

Some institutions may prefer to implement and deploy their own IoT network or data platform, while others may see the advantages of using common building blocks together with other institutions. SURF could aim to facilitate a common network technology, a common data platform or both. These choices and their advantages and disadvantages are illustrated in the diagram below.

Figure 17: What to facilitate? Network, data platform or both?

The current variety of network technologies and data platform technologies makes successful IoT introduction difficult and expensive. Considering the need for standard interfaces and network elements to make IoT applications successful, and taking into account the differences between different variants and the experiences with technologies and data platforms in pilots, it seems advisable to work towards a cheap standard ecosystem for IoT applications for the education and research community in the Netherlands. This ecosystem should contain basic elements for a IoT network components as well as basis data platform functionality that can easily be extended and adapted.

At this moment however, for SURF and the SURF institutions, the choice for a cheap, open and extendable standard is the most feasible to help IoT research forward. There are still many IoT technologies and no single standard is emerging yet. LoRaWAN is a network technology that is commonly used, offers cheap device components and uses license free spectrum. In combination with the open

IoT data platform SARA has developed this forms a standard IoT platform that is also open and extendible.

In the case other standards emerge these choices also do not provide a very high investment loss or vendor- or technology lock-in.

➔ Unlike others, SURF has a unique background in supporting research and development projects. SURF offers to become a partner or a Trusted Third Party in the area of end-to-end IoT solutions and can act as a facilitator between business, government and universities.

8. First use cases for the common SURF IoT ecosystem

The IoT ecosystem that SURF is developing should be flexible enough to support many different use cases but initially focus on particular use cases will help to start the ecosystem. In the pilots, both in the Netherlands and abroad, popular use cases are monitoring room climate for conference and class rooms and monitoring room occupancy, combined with room agendas. These are applications that have direct impact and advantages for institutions and are relatively cheap to implement. Therefore these use cases will be the basis of the first applications that will be implemented on the new SURF IoT network and data platform. Also these applications can be used as 'learning cases' for students during the implementation phases.

When the need for other additional applications arises, new initiatives can be added to the network and data platform.

➔ Devices and applications monitoring climate and occupancy in offices, lecture halls, class rooms, meeting rooms and conference rooms are only the first wave of applications on the open and extendable IoT platform.