A reference architecture for the IoE

Clive Longbottom, Service Director, Quocirca

The internet of things (IoT) is a hot topic. Quocirca is seeing many problems in the way organisations are approaching the concepts surrounding the IoT and architecting a platform to fully support the influx of a large number of devices and data, onto that platform.

The difference between the IoT and IoE

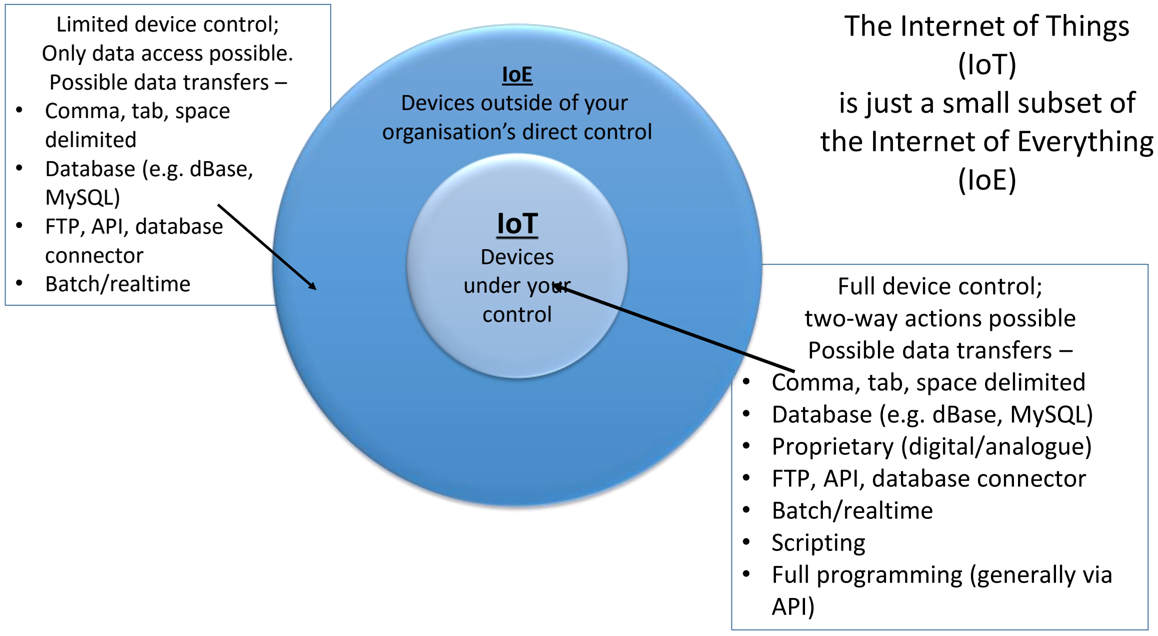

IoE
Devices outside of your organisation's direct control
IoT
Devices under your control

Limited device control;
Only data access possible.
Possible data transfers –

- Comma, tab, space delimited
- Database (e.g. dBase, MySQL)
- FTP, API, database connector
- Batch/realtime

The Internet of Things (IoT) is just a small subset of the Internet of Everything (IoE)

Full device control;
two-way actions possible
Possible data transfers –

- Comma, tab, space delimited
- Database (e.g. dBase, MySQL)
- Proprietary (digital/analogue)
- FTP, API, database connector
- Batch/realtime
- Scripting
- Full programming (generally via API)

Figure 1: The difference between the IoT and IoE

Quocirca proposes a basic architecture to help organisations avoid the many pitfalls of embracing the IoT across their own networks and beyond.

What is the IoT?

The idea behind the generally accepted term of the IoT, is to intelligently enable devices, so that they can provide data that can then be analysed and actions taken as necessary. In many cases, this will require bi-directional communication between the device and any analytics system; as the data is analysed, an event may need to be kicked off that causes the device to do something to remediate a situation. For example, a flow valve on a production line reports that the flow of liquid is at a certain rate. The analysis of the data shows that the flow rate should be increased by an amount; the valve is told to open up to make this so. In other cases, it may be a different device that takes the action – for example, a thermistor shows that the temperature in an environment is at a specific level; analysis shows that this is too high, and so an action is sent to a cooling device to bring the temperature back within limits.

This is fine as far as it goes – but there are two main ways of regarding devices that will be providing data to your network – those that are under your control, and those that are under someone else's control. Making a stark differentiation between these two environments is a key to ensuring that the right architecture is put in place.

To this end, Quocirca makes a split in definition between two different environments (see Figure 1):

- **The IoT** – this consists of all the devices that your organisation has total control over. The incoming data belongs to the organisation; events can be triggered based on that data, to command devices within the IoT to take specific actions.
- **The IoE** – the internet of everything is a superset of the IoT. It includes all those devices that are outside the organisation's control, but that are either presenting data of use to the organisation, or where it is difficult or impossible to stop them from passing data across the organisation's network. The latter includes consumer IoT devices (those that are owned and managed within that employee's environment) – for example, an employee's smartwatch attempts to synchronise data to other devices the employee owns, or they attempt to control their home heating system remotely from a device they are using on a corporate network. To gain the most out of the employee's own systems, they will want access via the corporate network. It may be possible to try and prevent this – but this may be counterproductive.

Putting in place an architecture that enables a full mixed IoT/IoE strategy to be implemented will be key to how well organisations perform in the future.

For the remainder of this paper, the term 'IoE' will be used when talking about the wider landscape of all devices that may impact an organisation's network, with 'IoT' being used when talking only about those devices that are under the full control of the organisation.

The issues of the IoE

The standard approach to understanding an IoE framework is to picture every device being attached to the global network, enabling everyone to have access to every device, or securing them from access, by building in security and access controls at the device level. This would not be a good idea. Not only does it threaten security through a huge increase in the attack surface, it also needs a full, global move over to IPv6 from IPv4 to ensure everything can have its own IP address. It also means that the short, small-packet data traffic created by such a global network of devices would bring private and public networks to their knees.

> *"The standard approach to understanding an IoE framework is to picture every device being attached to the global network, enabling everyone to have access to every device, or securing them from access, by building in security and access controls at the device level. This would be disastrous. Not only does it threaten security through a huge increase in the attack surface, it also needs a full, global move over to IPv6 from IPv4 to ensure everything can have its own IP address. It also means that the short, small-packet data traffic created by such a global network of devices would bring private and public networks to their knees."*

Current estimates for the number of IoE devices globally hover around the 20-50 billion level by 2020. Quocirca believes that these figures are far too low. By 2020, there will be around 8 billion people on the planet. Therefore, the prediction suggests there will be around 3-6 devices per person. Sure, many individuals will have no personal IoE devices by that stage, but consider where we are now. An average person, in a developed region, will have at least one smartphone, along with a tablet and/or laptop. Inside their home, they may have a smart, internet connected TV, a streaming radio service, one or more computers, along with games consoles

and other internet connected devices. They may be slightly more technically advanced and be using internet-connected security systems – IP cameras, intrusion detection systems and so on. They may have a Hive or Nest controller for their home heating and possibly lighting. They may have a new car, that reports data back to the manufacturer so that it can monitor the health of the vehicle, calling it in for a service as and when makes sense.

As the person is driving that vehicle, they will be picked up by multiple different CCTV systems. They may have their speed clocked by a speed camera. Consider the not-so-far-off future. As we move to driverless cars, the vehicle becomes full of devices that need to talk to each other, to central data controllers and to all the other vehicles around it on the road. Street furniture, such as traffic lights, roadworks and other items that could have an effect on the driver's journey, will need to become increasingly IoE enabled. As they park up, the car park may be using a number plate recognition system that ties back to any driving licensing authority if they then do not pay for the parking.

As the person boards the train, they may use a near field contact (NFC) card to pay for their journey. The train itself has a multitude of connected systems feeding data back to central control, as well as from there to individual's smartphones to keep them up to date with the train journey's progress. The train has WiFi, enabling the person to connect through to their office and home networks, alongside the global internet to synchronise data and send back information from, say, wearable health monitors and fitness devices.

Arriving at the office, the individual uses their smart badge to get through security, and is logged as they move throughout the building. They buy a cup of coffee using a smart card – and so on.

The reach of the IoT already is massive – there is a high probability that the current estimates of 50M for the number of IoT devices by 2020 is at least an order of magnitude too small.

So, with somewhere between 20 and 500 billion IoT devices worldwide – what can be done to ensure that the organisation's IT as well as the global IT platform can deal with it?

There is a need for a more intelligent architecture to be put in place – now, rather than later.

A reference architecture for the IoE

Clive Longbottom, Service Director, Quocirca

The wrong architecture

The ‘one ring to rule them all’ IoE

- Bung all device data into one database for direct analysis:

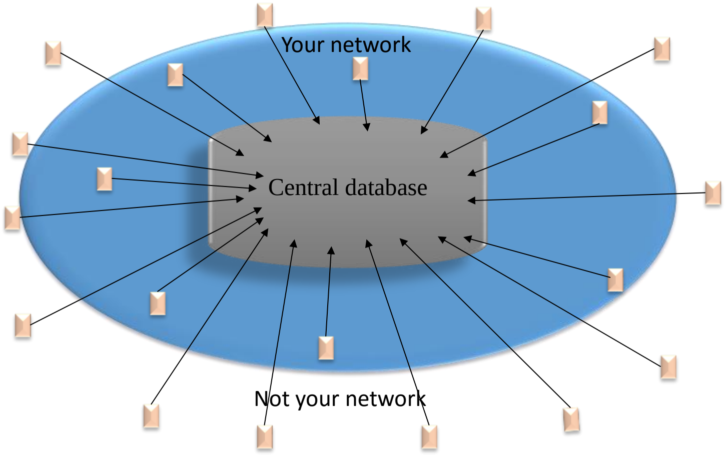

Figure 2: The wrong architecture for an IoE architecture

Let’s consider the approach of openly connecting every device to the network. For this to work, every device has to have direct access to a central database. This requires every device – both inside and outside the organisation’s network – to have a unique identifier. For this to be the case, staying with IPv4 addressing will be impossible. IPv6 is a necessity – not just on your organisation’s network, but on every other organisation that you will be dealing with, and across every network that will be used (see Figure 2).

As of May 2016, Google measured IPv6 native traffic to its websites at 11.48% - double the figure of a year ago, but still far too low to for an IoE architecture dependent on the use of IPv6. Even if an organisation implemented IPv6 and all of its partners did too, there would still be a requirement for the networks connecting them to also be IPv6 – possible across expensive dedicated networks; not so possible when looking at using the public internet.

However, IPv4 addresses have run out. Except they haven’t, really. Through the use of Network address translation (NAT), there can be an infinite number of devices with the same IP address – and they can still operate as part of an IoT/IoE architecture.

How? By using IoT aggregators.

IoT Aggregators

An IoT aggregator is a system that puts network intelligence close to a group of devices. Through this means, the aggregator can use NAT to address each IoT device separately, while still maintaining a unique identifier for each device.

Use of NAT as a device hierarchy management tool

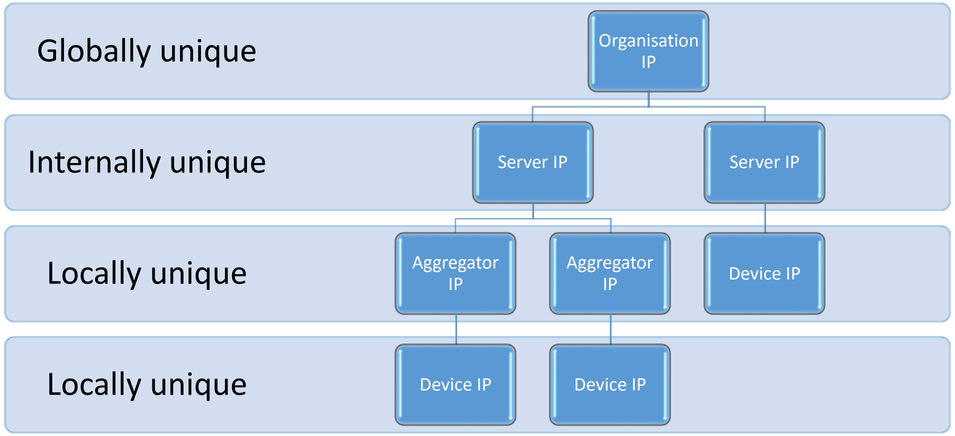

Figure 3: Use of NAT as a device hierarchy management tool

To fully understand this, a look into how NAT works is required. A corporate network will have an IP address for each device, in the form of xxx.xxx.xxx.xxx for an IPv4 address. The organisation itself must be able to represent itself to the outside world via a unique IP address – for example, the BBC.co.uk website has an IP address of 212.58.244.70; Microsoft.com has 191.239.213.197. These IP addresses are 'hidden' behind the organisation's domain name. Domain name service (DNS) servers maintain lists of domain names against IP addresses – in this case, www.bbc.co.uk will point to 212.58.244.70. The BBC could change its IP address – all it then needs to do is to change the DNS entry to reflect this.

When it comes to dealing with everything behind that public address, the BBC has two options:

- Behind this corporate IP address, the organisation can choose to have every item within its network have a globally unique address – in the case of the BBC, every device would have to have IP addresses starting with 212.58., but it could use any numbers at the end.
- It can choose to use NAT.

The most common NAT addresses used behind a corporate IP address are 198.162.xxx.xxx, 172.16.xxx.xxx and 10.xxx.xxx.xxx. If a full discovery of every device on the planet was carried out, there would be thousands of devices with the same IP address – and yet the overall global network continues to work. This is because the devices behind a NAT gateway are viewed as being unique – the gateway manages how that device interacts with the rest of the world.

As an example, let us consider that BBC network. It represents itself to the rest of the world as 212.58.244.70. Behind that address, it can point to a range of other devices through the use of port redirection – if it wants a different server to deal with web traffic, it could use a NAT table to take all traffic coming in over port 80 and redirect it to another server. This server can be within the private IP block – so a web server that is internally designated as 192.168.1.1 can be seen as the outside world as 212.58.244.70:80 (see Figure 3).

With the IoT, the same approach can be taken. For example, assume that a simple aggregator is responsible for 5 devices – 4 thermistors and a control valve for letting cooling air flow through an area (see Figure 4). Within the aggregator, these can be seen as 192.168.1.1, 192.168.1.2, 192.168.1.3, 192.168.1.4 and 192.168.1.5. Potentially, other aggregators around the network will have similar devices with the same IP addresses.

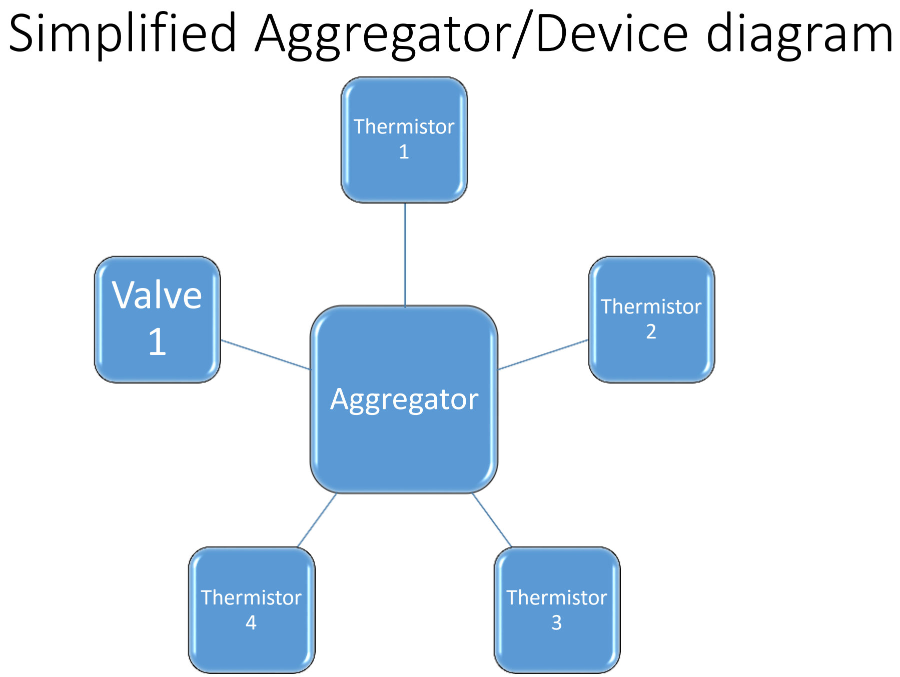

Figure 4: Simplified Aggregator/Device Diagram

However, if the aggregators have their own IP addresses (say, 192.168.10.1, 192.168.10.2, etc.), then NAT can be easily used. With the aggregator we are looking at (say, 192.168.10.1), it can set up different ports that address different devices. So, our first thermistor can be seen by the rest of the network as 192.168.10.1:90, the second thermistor as 192.158.10.1:91 and so on. An internal DNS server can then be set up that maps these to more common names – for example, thermistor1.bbc.co.uk would be matched to 192.168.10.1:90.

Now we have a more scalable and manageable means of bringing devices into the IoT platform. IPv6 is no longer needed, and each device is addressable via the network.

Using Aggregators as Intelligent Systems

Another problem with having a global IoE architecture where every device is a peer on the network, is in dealing with how much intelligence such an approach requires. Many devices are low cost – for example, a thermistor will be a few pounds at most. Building in sufficient intelligence for such a device to be secure, to be able to deal with its own data and to make this available to a central database in a usable manner, will require many times that amount spending on the device. For the majority of organisations, this becomes a show stopper. Every device needing to be replaced with a far more expensive one, just doesn't make sense.

A reference architecture for the IoE

Clive Longbottom, Service Director, Quocirca

Using aggregators enables a different approach. The money that would have to be spent on making each device intelligent can be combined, and a small proportion of it spent on making the aggregator more intelligent. In this case, existing devices can still be part and parcel of the IoE; they continue to feed data in whatever format they have done in the past. The aggregator takes this data and carries out any extraction, transformation and load (ETL) actions that are required to normalise the data and make it available for analysis.

> *"Quocirca recommends that each aggregator looks after around 10 – 100 IoT devices or IoE feeds. This provides the economy of scale that is required to make building the intelligence into the aggregator cost effective. For example, if it would cost £10 per device to build basic intelligence into an IoT device, at 10 devices, that is an aggregate cost of £200 that can be put into the intelligence in the aggregator itself. At 100 devices, it becomes £2,000. Above 100 devices, there are the risks that a security breach of a single device behind the aggregator could lead to all devices in that area being compromised."*

A further advantage of this approach is that each group of devices is in an 'air lock', removed from the rest of the network due to the aggregator. Therefore, even if the security of a device itself can be compromised, additional intelligence can be built in to the aggregator to deal with this – pattern matching, anti-virus capabilities, deep packet inspection and so on. An aggregator-based architecture builds security into the IoT and IoE.

Quocirca recommends that each aggregator looks after around 10 – 100 IoT devices or IoE feeds. This provides the economy of scale that is required to make building the intelligence into the aggregator cost effective. For example, if it would cost £10 per device to build basic intelligence into an IoT device, at 10 devices, that is an aggregate cost of £200 that can be put into the intelligence in the aggregator itself. At 100 devices, it becomes £2,000. Above 100 devices, there are the risks that a security breach of a single device behind the aggregator, could lead to all devices in that area being compromised.

What else can an Aggregator do?

An aggregator can add still more value to an IoE network. For example, in the simplified case above, of an aggregator with just five devices, there will be rules set as to what is and is not to be expected. In this case, the temperature as measured by the thermistors should be within the limits of 17-20°C.

If all four thermistors are continually letting the aggregator know that the temperature is 17°C, there is no need for that information to be passed beyond the aggregator itself. Everything is working within limits, why swamp the wider network with unwanted data?

However, let's assume that one thermistor posts an 18°C measurement. This is still within limits – but it is a change. The aggregator can store that data to see what happens next – if the thermistor comes back with a 17°C reading, then all is well – the aggregator can then drop that data and wait for more data to come through.

However, if that same thermistor next reads 19°C, the aggregator may have enough intelligence built into it, to take autonomous control and tell the valve to open slightly to bring the temperature back down again. It can poll the other three thermistors to see if they have noticed any discrepancies. It could then send the data back to a more intelligent central system, where enhanced analytics can be carried out and more complex actions decided upon.

These central environments then need access to the data that the aggregators are seeing. It is fine having aggregators that push data down when they see something happening, but this is not enough for an intelligent IoE architecture. The centre also needs the capability to request the aggregators to send the information to it.

Therefore, the aggregators need to be able to store some data – but this is not every last bit of data that has been gathered over the last day or so. Rules need to be in place to define the life cycle of data, so minimising the storage needs of the aggregators.

As an example here, consider an aggregator at one end of a building that has noticed a temperature rise from the thermistors it has responsibility for. Is this purely localised, or is it more widespread? By the centre being able to request data from all aggregators, it can find that those aggregators close to the one that has raised an alarm, are seeing small temperature rises too – not enough to trigger an alarm, but enough when analysed to match a pattern with the original alarm. Although the aggregators themselves do not have enough intelligence built into them to be able to notice this trend and take actions, the central controller can. The aggregator can figure out from such data, for example, that a fire has started somewhere and is spreading in a specific direction.

What does this mean to the overall IoE architecture?

So far, the discussion has mainly been around the IoT – those devices that belong to the organisation, and where those devices capable of taking an action can be tasked to do so via an intelligent gateway.

However, the IoT spreads well beyond the boundaries of a given organisation.

Take a simple example, a retailer will have suppliers and customers. To move things along from one part of this value chain to another, there will generally be the need for an intermediary, such as a logistics company (or, for goods delivered online, a network services provider). As a retailer, it makes sense to know whether the supplier already has an item in stock and how soon it can deliver it. Once goods have been despatched, knowing where they are in the logistics chain helps to ensure that customers are kept informed, and remain happy.

The retailer therefore, has to have access to data from along that value chain – a chain that can get complex if the supplier is taking components from other suppliers, and the retailer's customers are middle men, who are then selling on the goods to other outlets and/or consumers.

"The retailer has no control over the devices that are creating the data and no control over the security of these devices. All data coming in from outside the direct control of the retailer has to be regarded as hostile; as possibly carrying harmful payloads. The devices responsible for these data streams are in the IoE – they need a different approach, but one that has commonality with the IoT approach."

The retailer has no control over the devices that are creating the data and no control over the security of these devices. All data coming in from outside the direct control of the retailer has to be regarded as hostile; as possibly carrying harmful payloads. The devices responsible for these data streams are in the IoE – they need a different approach, but one that has commonality with the IoT approach.

A reference architecture for the IoE

Clive Longbottom, Service Director, Quocirca

In a full IoE model architecture, there would be little that can be done to mitigate such security issues. The device connects directly to the central database so any impact on the data is direct and full – all the data in that store has now to be regarded as possibly compromised and/or stolen.

By using aggregators, outside data can be held at the network edge until it has been examined and analysed to make sure that it has not been compromised. Payloads can be searched using anti-virus and advanced threat detection software; unusual patterns of data can be identified and either put into a separate area, until manual investigation can be carried out, or dumped, as being too out of pattern.

> *"As an example of an aggregator, consider something that is very similar to a hardened PC – a main CPU with motherboard, memory, storage, an operating system and network connectivity. The aggregator can therefore, be programmed as a normal system, with updates being pushed from a central control point to each aggregator as required. Using software defined constructs, different functions can be defined for different use cases – deep packet inspection to analyse external data streams; filters and action event triggers for internal IoT activities. Events can also be triggered through external IoE data analysis - but these generally be to alert the organisation owning the device of an issue, rather than kicking off an automated remediation action."*

The aggregator can also carry out necessary actions to ensure that the organisation does not fall outside of its legal obligations. For example, if a downstream supplier in the value chain, sends through data that contains personally identifiable information (PII) or credit card details, these could cause a problem if they are accepted, and stored by the organisation. The aggregator can examine the data and fully redact such data, either securely erasing it from the stream, or keeping only such data as is necessary for fidelity's sake (for example, the last four digits of a credit card number).

As the aggregators are intelligent devices, they should also be more flexible and have a longer life. By using aggregators that are designed to make use of a software defined model, changes in how data needs to be gathered, managed, analysed and reported against, can be managed through the writing of new functions that can be pushed out to the aggregators. Should such a need occur in an environment where every device has to have the intelligence built into it, in many cases, this would require the complete replacement of the devices involved.

As an example of an aggregator, consider something that is very similar to a hardened PC – a main CPU with motherboard, memory, storage, an operating system and network connectivity. The aggregator can therefore, be programmed as a normal system, with updates being pushed from a central control point to each aggregator as required. Using software defined constructs, different functions can be defined for different use cases – deep packet inspection to analyse external data streams; filters and action event triggers for internal IoT activities. Events can also be triggered through external IoE data analysis - but these will generally be to alert the organisation owning the device of an issue, rather than kicking off an automated remediation action.

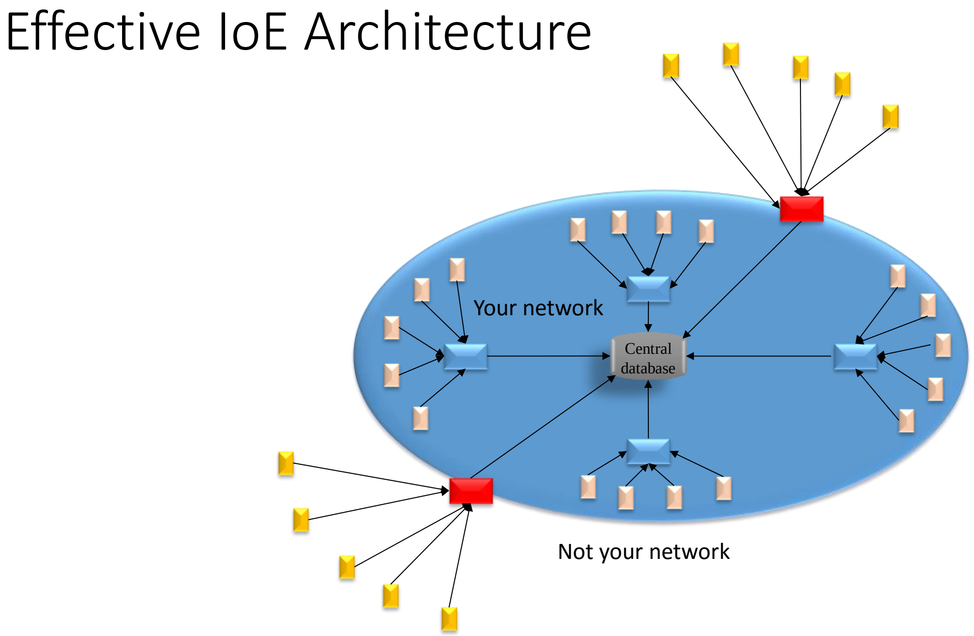

Figure 5: IoE Reference Architecture

The IoT/IoE Reference Architecture

This now brings us to how Quocirca believes a reference architecture should be constructed to deal with the emergence and evolution of the IoT/IoE.

The use of aggregators is a core part of constructing such an architecture. By using aggregators, the IoT and IoE can be divided up into manageable 'chunks', either grouping devices together by type, location or overall function. The aggregators improve security, allow for older existing devices to be kept in place, and negate the need for a move to IPv6. They minimise network traffic, so maintaining the network bandwidth for more important data traffic. They spread the load, meaning the central data acquisition, analysis and event management centre can be more easily provisioned against a lower cost platform.

All IoT devices should operate behind an aggregator. The aggregator applies simple local intelligence and shields the main network from any security breaches on any of the devices. All IoE devices should operate behind more intelligent edge of network aggregators, that carry out deeper inspection and patter matching of all external data to ensure that it is non-malicious and is clean of any information that would compromise the organisation, should it store it on its own systems.

Such an architecture enables an organisation, to more easily deal with the evolution and maturation of the IoE, to deal with events and trends across the IoE more directly and effectively, and also to embrace the different types of IoE that will occur. Back to the example of the employee as a consumer – by using the right aggregation devices,

all traffic to do with an employee wanting to control their home via Nest, can be routed through a specific aggregator that air locks all the activity from the rest of the corporate network.

The IoE is already here – it is big, and it is growing. Getting the architecture to deal with it will be very costly – in many cases, the costs of replacing bad architectures could break a company. Get it right from the start – embrace change, while ensuring that data traffic is minimised and security is maximised.

For further information, please look at Quocirca's web site at www.Quocirca.com, or get in touch with Clive Longbottom at Clive.Longbottom@Quocirca.com

quocirca