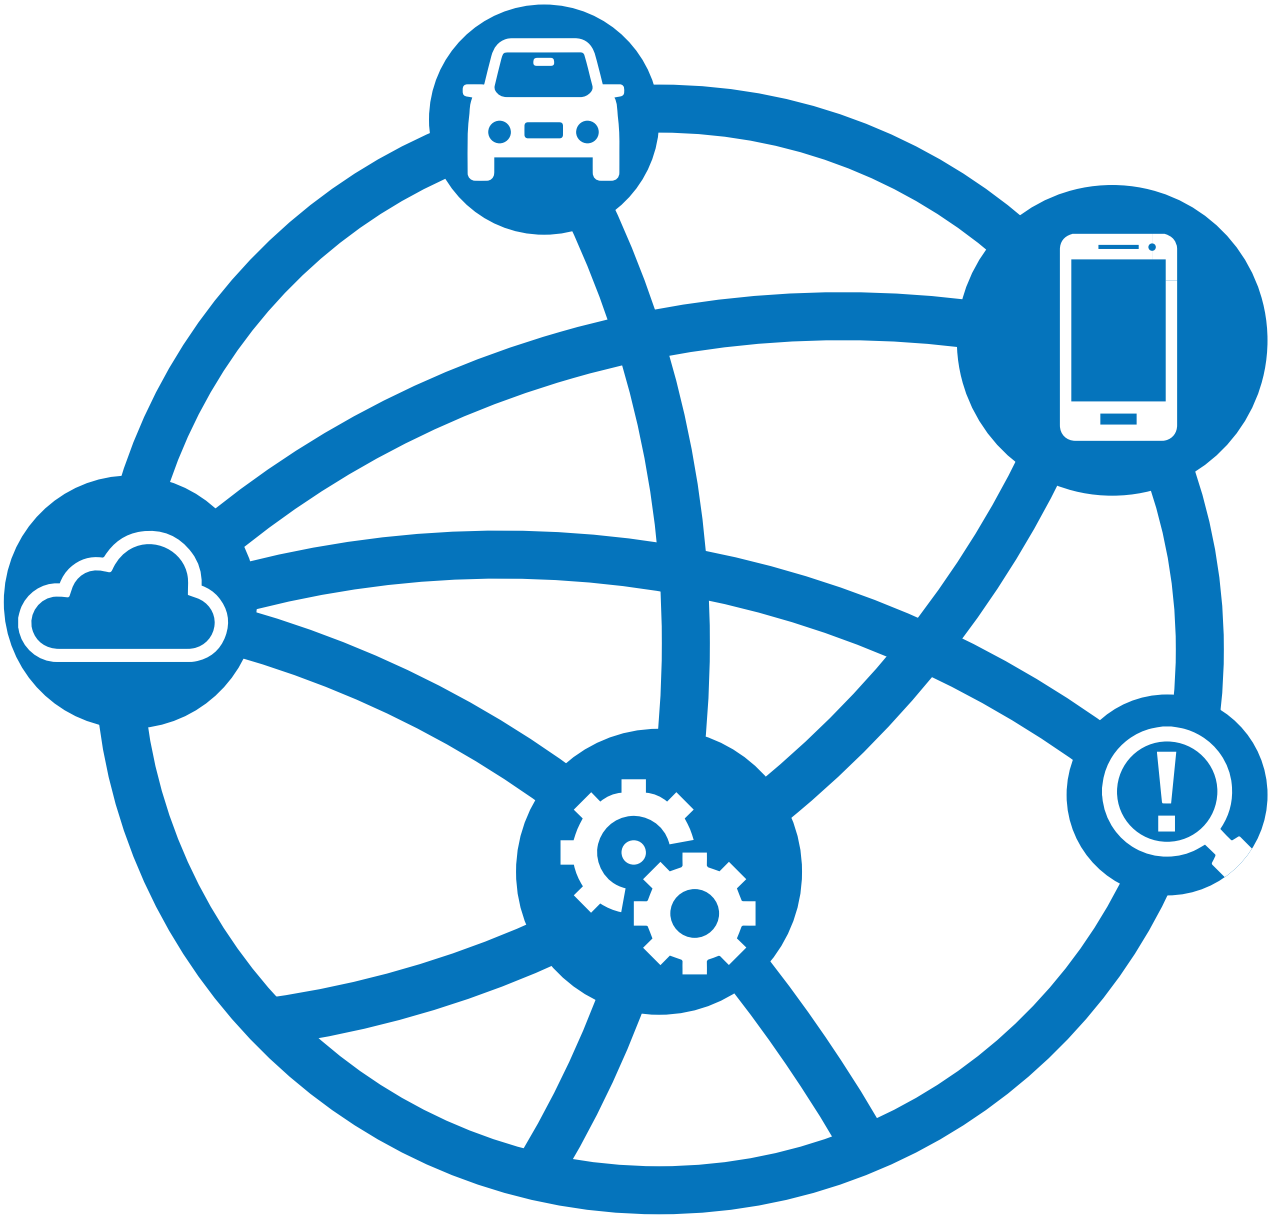

HOW TO ACHIEVE NETWORK OPTIMISATION FOR IoT

SPONSORED BY AERIS

aeris

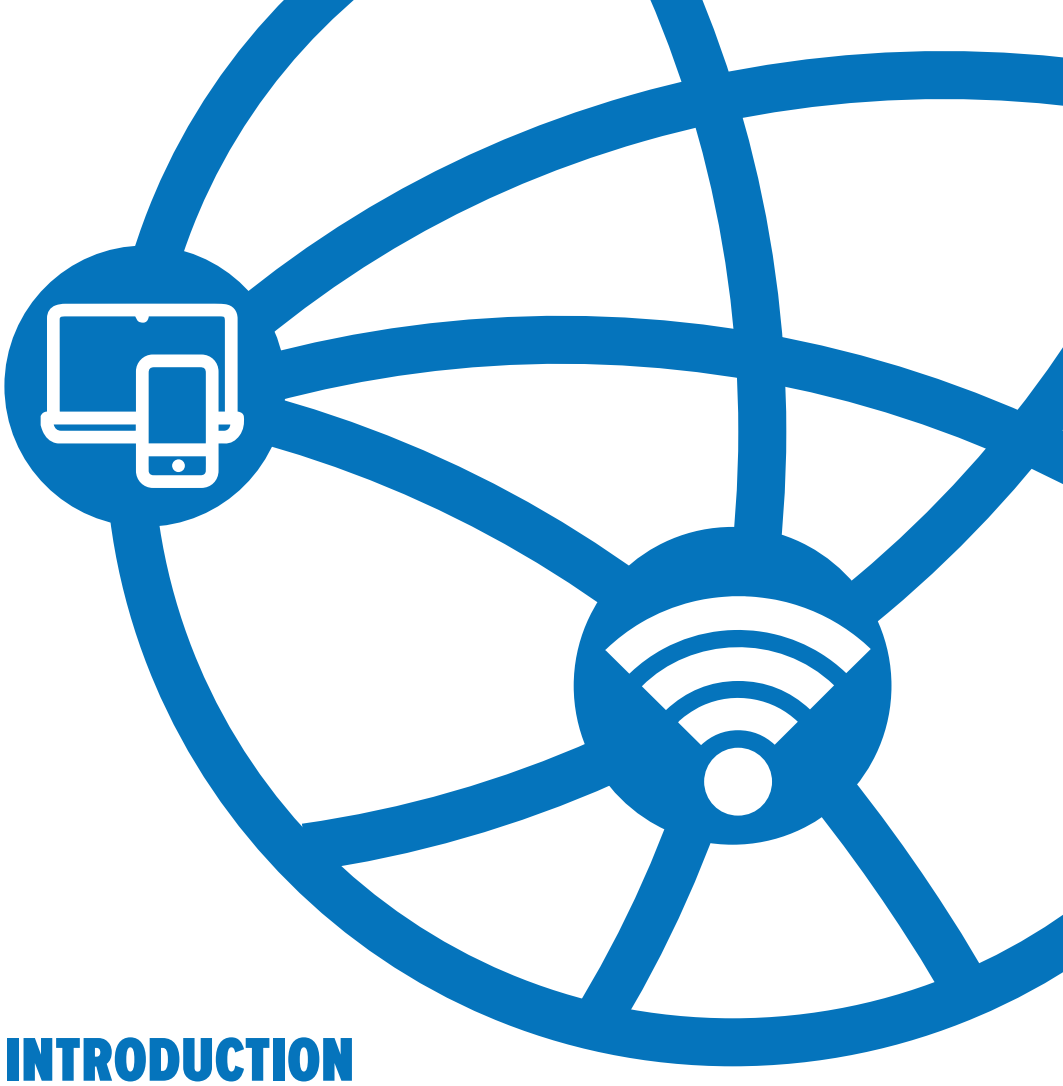

INTRODUCTION

With analyst firms forecasting exabytes of IoT-related network traffic by the mid-2020s and billions of IoT devices projected to be connected, it's clear that IoT networks need to be optimised to minimise the impact of IoT traffic on other services using cellular networks and to ensure maximised utilisation of network resources on other types of networks. The application traffic is only part of the challenge, the number of devices involved in IoT will mean a radical increase in the volume of non-app related traffic as devices send messages in the control plane of networks.

Managing networks for IoT is a different discipline to that of smartphone networks because of the different profile of IoT traffic which is subject to pronounced, yet predictable peaks and troughs. Sensor networks, for example, may communicate on the hour, every hour, but other apps such as connected car services might put more stress on the network if traffic becomes bunched up in a jam. Optimisation is therefore a critical goal and organisations are looking to ensure they can optimise as IoT services mature and the mass market of billions of devices emerges.

WHY NETWORK OPTIMISATION MATTERS FOR IoT

The reasons for optimising networks are not just technical, there are business reasons for network optimisation, too. If you consider IoT apps you might have low throughput and data volume from the app itself but the presence of a device could generate more control plane messaging than an organisation might like to support. An app, for example, could only need to send a few megabytes of data per device per month but the control plane messaging could have a dramatically higher volume than app traffic itself. This non-app focused traffic places a significant burden on the network. A way to address the issue is to specify equipment that has been designed to limit control plane traffic for IoT devices. It's important to recognize IoT devices have different requirements to smartphones and therefore do not need to

The author, **Syed Hosain**, is the chief technical officer of Aeris where he is responsible for the architecture and future direction of Aeris' networks, development programmes and technology strategy

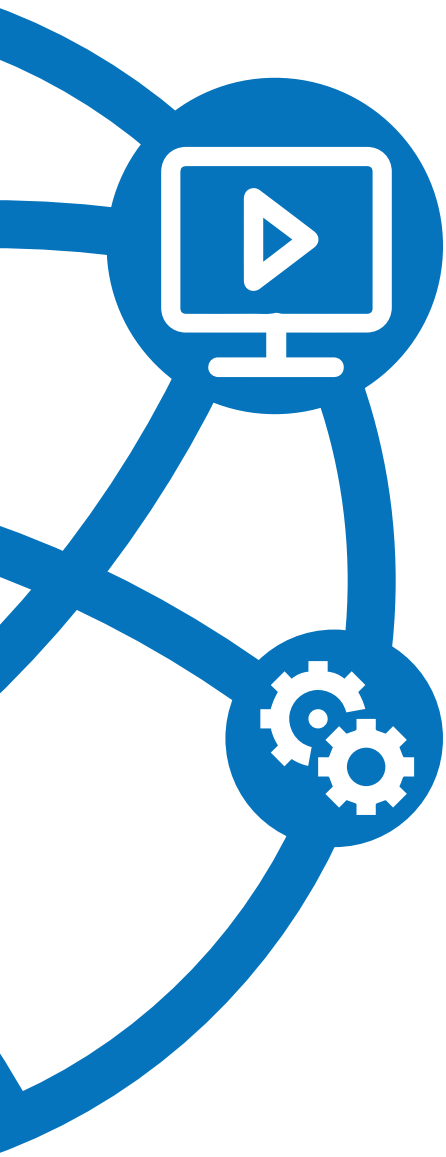

communicate with the network via the control plane in the same way or regularity. Another key point regarding why network optimisation matters is the sheer volume of endpoint devices involved in IoT. Assuming you believe the predictions that there will be tens of billions of devices there is going to be tremendous scale to contend with. Even small throughput at this scale means the impact will be enormous.

WHY IoT NETWORK OPTIMISATION IS DIFFERENT TO TELEPHONY

The demands of IoT apps are and will continue to be quite different to the traditional smartphone environment which is why systems that are capable of handling a huge volume of relatively small data communications are required. Efforts such as NB-IoT and other non-cellular bearers such as low power wide area (LPWA) radio networks, are seeking to address part of this challenge but that does not enable the industry to ignore the fundamental differences between IoT traffic and smartphone traffic.

IoT devices behave very differently to smartphone or telephony usage. Most people will be connected via their smartphones for relatively long periods with hard to predict bursts for activities such as downloads. IoT devices, in contrast, tend to communicate with much greater regularity. In addition, IoT devices have been designed with for extended battery life so they power up to transmit and then power down which means there are important data management issues to deal with.

In general, IoT devices involve a lot of regularity in transmission which is not the case with humans using their smartphones. This regularity makes it easier for IoT service providers to monitor patterns and establish if a device isn't working or something is happening that hasn't happened before. This is essential for IoT because, unlike the customer to service provider relationship with smartphones, projects involve the management of 10,000 devices in contrast to the single user of a cellphone. Therefore, if problems can be discovered automatedly and, ideally, predicted before they happen, a truly valuable service will be provided.

WILL THE MARKET BE READY FOR MEGA VOLUMES OF DEVICES?

The likely emergence of billions of devices will not be a surprise to the IoT industry. In fact, it's the indicator of success for IoT and most IoT specialists have developed their systems and products, as well as their business plans, with huge scale in mind. To that end, Aeris, in common with some other providers, has completed simulations to determine how many cellular devices can be attached to our network before systems start to fall down.

In addition, Aeris has found with its LTE networks, which rely entirely on cloud, that it can scale up easily. Aeris is able to do this because it can spin up how we support the network in a variety of ways and this is not tough on a server because Aeris is able to distribute the data across multiple server locations very quickly. This is quite different from a traditional cellphone network provider which is very localised. Nevertheless there are going to be a large number of devices and there will be customer challenges there.

THE NEED FOR FLEXIBLE SCALABILITY

The need to scale up flexibility is an issue that affects the entire market. Nobody knows at exactly which point an individual IoT application will take off and nobody can predict which types of IoT app will lead the market. Therefore, service providers have to have the capability to scale up on demand from a relatively low base. Aeris, for example, has approaching ten million devices in its network and we has simulated at 100 million devices to ensure it can support 10x growth. Organisations have to be able to handle not only the average load but also the peaks. For example, many IoT devices are set to communicate precisely on the hour and that creates a spike in traffic.

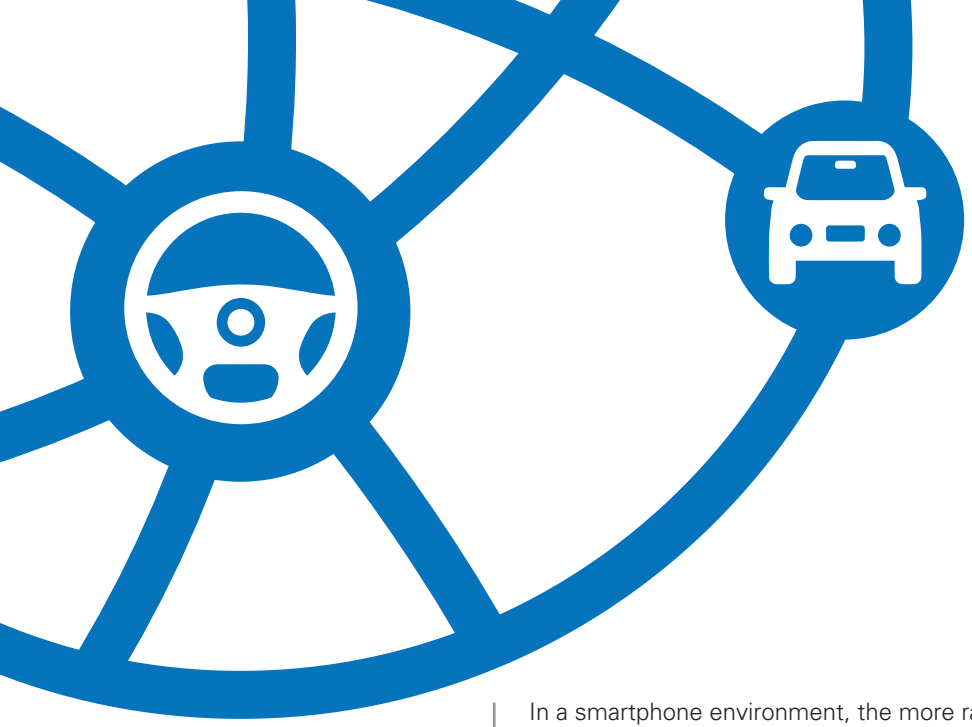

In a smartphone environment, the more random nature of individual usage smoothes out the peaks and troughs but IoT traffic is less spread out so you have to ensure you have the flexibility to accommodate the peaks. The variance from peak to trough can be 10x but this is on a very regular lifecycle so there is greater predictability.

NETWORK OPTIMISATION ACROSS MULTIPLE TECHNOLOGIES

This is an issue that's going to affect most of the IoT market. However, IP technologies can be treated similarly and normalization between network technologies is a capability that can be provided. With Wi-Fi, cellular or NB-IoT IP data is already provided but on Sigfox or other low power bearers messages will need to be treated differently. Aeris, for example, achieves normalization of this sort of network traffic by delivering it as IP encapsulated messages.

In addition to the challenge of handling different network technologies, IoT deployments, especially global ones, rely on the networks of many different carriers. Aeris encounters this issue frequently. It can provide services using 500 different operators which creates some challenges in device normalization. For smartphones this isn't an issue for the network because the manufacturer will address localization issues but for IoT devices the challenge is greater. Aeris often has customer deployments in multiple countries where it has to normalize the behaviour of the device.

The company has identified 20 global zones in which it treats several countries in the same way. For instance, Aeris has zones such as Europe and South America so when a device is shipped to one zone it has a good understanding of what the billing will be. It's very important that a device in London will operate in the same way as a device in Madrid. There will, of course, be some differences in the radio frequencies used but the back-end piece is essentially the same.

EXAMPLE 1: NETWORK OPTIMISATION FOR AUTOMOTIVE APPS

Aeris encountered a situation several years ago with an automotive industry customer. The problem was that a vehicle would send a crash notification SMS and the service would be followed up with a call from the car user. The data centre people would correlate the information from the SMS with the incoming call but about 35-30% of the time the caller ID wouldn't show up because the location could require utilization of a small local carrier.

The data centre operators then had to DTMF tone sequence analysis to identify who the call was from. The automotive company wanted us to find out who calls were from and recognized that we could extract a notification of a call delivery attempt so out systems would see a call being attempted. We could automatically engage an app to put the call on momentary hold and query it. Knowing the text had just been received from a specific location meant we could then route the call to an operator that has the SMS on screen. This process went from taking 30 seconds on quickly resolved calls to less than 15 seconds.

This functionality along drove the car manufacturer to select us and we went from identifying the problem to having a commercially deployed solution in eight weeks.

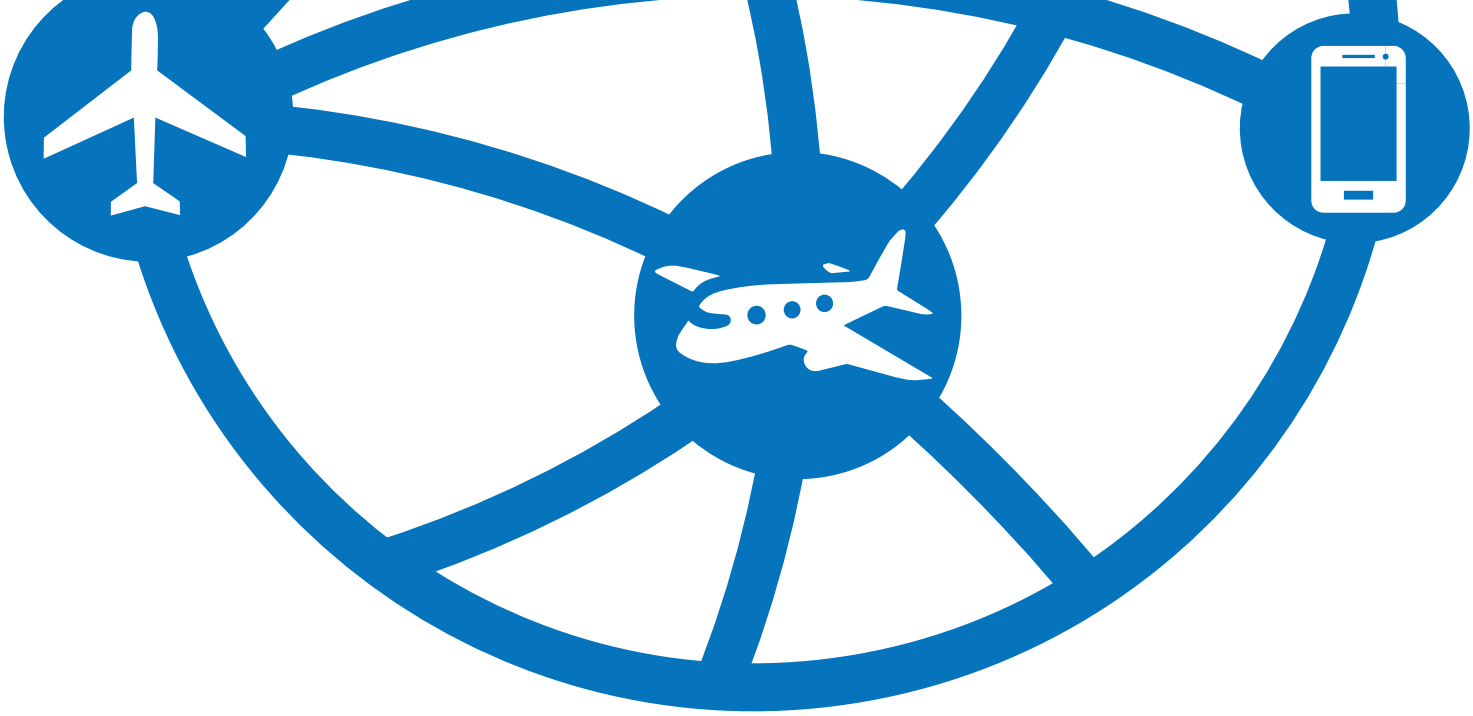

THE AERIS APPROACH: AERCLOUD AND AERVOYANCE

In the past, Aeris was purely a connectivity service provider competing with mobile operators but also in partnership with them for connectivity provision. The data would flow through Aeris systems and we would deliver it to data centres to handle it.

Aeris decided to create systems that would look at the content in the data instead of just delivering it. Aercloud was created so Aeris could store data and look at the triggers within it so the customer organisation could decide what to do. Aercloud stores, manages and enables Aeris to look at the data in a streaming fashion so it can provide appropriate alerts to the customer.

Aercloud is added to with the Aervoyance offering which provides analysis that looks not only at the streaming data but also enables customer organisations to look at the past to analyse issues.

EXAMPLE 2: NETWORK OPTIMISATION FOR A PLANE MANUFACTURER

Aeris has deployed Aervoyance for a plane manufacturer and, while partial insights can be gathered from communications when the plane is in the air, the bulk of the data is transmitted when the plane arrives at its destination gate. Aervoyance is able to look at the data and also perform predictive analytics.

Aeris does not currently provide customer organisations with choices or decision points but the company plans to start using Aervoyance to look for patterns that humans can't recognise using specific algorithms it has developed. The company's expectation is that customers will be surprised at what it can uncover utilising this new functionality.

NETWORK OPTIMISATION BENEFITS

Network optimisation brings a series of benefits to organisations which include: the ability to save money, improved operating efficiency and faster time to scale up or deploy.

Network optimisation saves money by enabling throughput to be optimised. Network optimisation improves efficiency because you manage the transport of data from the low throughput devices and IoT devices tend to be application-specific so we have customers that like to be notified and start running analysis as soon as the data arrives.

There are also things that can be done to save further money. From a financial perspective, when there is an overage scenario, such as if a device goes haywire and starts to transmit continuously, the organisation will want to shut it off immediately. Aeris can set thresholds and limits and look out for situations that are abnormal. In addition, network optimisation certainly helps capacity planning. IoT apps tend to be very uniform and send similar amounts of data each time they communicate so you can manage them effectively. That predictability means upgrades can be made in a timely way so the service doesn't get caught out by an increase in devices and usage. The equation is simple, multiple by the average traffic by the number of devices deployed for the app and you get the likely network requirement.

CONCLUSION

If you look at the monthly revenue from smartphones and compare that to IoT devices, there's a tremendous difference. In the US, for example, there are a large number of IoT cellular devices yet the total revenue being generated by the carrier from them is less than 1% of their total revenue. However, all of those devices could have a significant impact because of all their control plane messaging.

The industry needs to start looking at this in a very critical way and the issues will probably be addressed better by specialists because carriers focus on smartphone requirements. We won't see changes being made by carriers to better optimise IoT devices because these are costly changes and, because of the low percentage of IoT revenue they generate today, such activity isn't a priority for them. Specialists like Aeris will put the effort in where a carrier won't to optimise the network and eliminate unnecessary control plane traffic from IoT devices.

MARKET INSIGHTS

Infonetics Research forecasts that **70% of total network traffic by 2020 will be IoT or M2M-based**.

Machina Research projects that not only will cellular connections increase from 250 million to 2.3 billion in the next decade but that **traffic will grow even more quickly from 200 petabytes in 2014 to 3.2 exabytes in 2024**.

Even with billions of IoT connections live, by 2023 the **revenues from all LPWA connections in the world will only be around half the current size of the French mobile market**, reports analyst firm **Analysys Mason**.

aeris

www.aeris.com