Edge Cloud to Cloud Integration for IoT

Version 1.0

Guenter I. Klas

Feb 04, 2016

Abstract

Telecoms and Internet industries work on concepts and architectures that will support new enterprise products and solutions for the Internet of Things (IoT), in addition to new products and services for the consumer sector. The concepts known for example as mobile edge computing, edge cloud computing and fog computing aim at bridging a gap between devices of the IoT and more and more adopted cloud computing infrastructure of enterprises. The gap emerges as the classic, centralised (and often hyper-scale) data center paradigm becomes insufficient for the needs of new IoT applications that are characterised by being real-time, requiring very low latency and getting deployed in the millions or even billions, thereby multiplying the ingress data volume to central data centers. This paper argues that the business case for edge cloud infrastructure might in the end depend on successful adoption by enterprises that make heavy use of M2M and IoT. If that is indeed the case, integration of edge computing infrastructure with NFV/SDN concepts might be necessary, but not sufficient. Missing is the integration of edge cloud platforms with data centers and private clouds of IoT enterprises, as well as the integration with public hyper-scale clouds.

Tags: Mobile edge computing, edge cloud computing, edge cloud, edge computing, fog computing, fog node, 5G, cloud computing, cloudlet, micro data center, open edge computing, intelligent cloud, OEC, MEC, ETSI MEC, NFV, SDN, low latency, tactile internet, Internet of Things, IoT, distributed cloud applications, real-time analytics, programmable networks, network slicing, IaaS, PaaS, SaaS, Industrial Internet Consortium, IIC, Cisco, Jasper Wireless, CloudFoundry.

One question to start the discussion from is this: How would enterprises in the M2M and Internet of Things (IoT) sector benefit from IT infrastructure and cloud computing at the edge of telecoms networks?

First, we have a look at the state of computing at the edge of networks.

1. Where are we with edge cloud computing today?

The telecoms industry is undergoing a major technical revolution, whether in the US, Europe or Asia. Some of the keywords that characterise this revolution are:

- *Programmable networks* realised through virtualization (in particular network function virtualization NFV) and software defined networking (SDN).
- *Network slicing*, i.e. the ability to provision a bespoke telecoms network with particular characteristics for specific purposes or tenants out of existing telecoms resources. A network slice is essentially a chunk of virtual (and maybe physical) network resources, reserved, configured and provisioned for a particular business use end-to-end (e.g. a network slice for messaging, or a network slice for a low power M2M wide area network or a network slice for TV/video distribution).
- *Edge computing or edge cloud computing*, i.e. the capability to provide compute, storage and network resources at the edge of telecoms networks, whereby the edge means in closer proximity to end users (whether human users, machines or things of the Internet of Things).
- *5G radio interface technology*, supposed to be highly flexible and enabling a vast variety of new use cases, some of them requiring super-high bandwidth or ultra-low latency.

Thus, *edge cloud computing* is one element in this ongoing stream of technical innovation. It comes in various flavours and we have already discussed *mobile edge computing* (MEC) and *fog computing* in other papers (see [1] and [2]).

For the sake of the discussion here, let's adopt the notion of ***edge cloud computing***. What we mean with that is:

- *Virtualised mini-data center infrastructure,* possibly in a variety of locations owned or controlled by telecoms companies (or other enterprises). Sometimes they are called micro data centers or cloudlets or fog nodes or MEC servers (and of course, there are a number of differences). As there is a choice of terminology, for the purpose here let's call such a small footprint IaaS environment a cloudlet (indicating a mini-version of a cloud computing environment).
- *Distributed*: The idea is that telecoms companies (and maybe other enterprises as well) provide such cloudlets in a spectrum from low-end compute and storage capability to high-end, *in diverse locations* across the country. Example candidate locations are base station sites (typically tens of thousands of locations in individual countries), network aggregation hubs (hundreds to thousands), switching sites, telecoms data centers (maybe in single digit to double digit numbers) etc. Actual locations will very much depend on the economics, network latency, latency requirements of applications, more generally customer needs and customer preferences.
- *IaaS/PaaS*: Such cloudlets act as infrastructure as a service (sometimes also as platform as a service). This means, that application software can be deployed to such micro data centers in a rather agile way, on demand, e.g. in form of virtual machine images. In addition, the IaaS may cater for on-demand migration of application code between cloudlets during runtime to cope with mobility of devices (e.g. in order to keep application code in proximity of a device).
- *Special services*: In addition, some of these cloudlets may offer special services, which are only available through them as they have privileged access to network information. For example, a cloudlet may host a service that provides access to real-time information regarding the radio network conditions which a human user or a thing of the IoT may face, just now, in exactly this particular radio micro or macro cell. Such real-time context information can be used e.g. to improve video streaming to end users.

Given that we have a rough idea of what edge cloud platforms (cloudlets, fog nodes, MEC servers etc.) are, the next question is:

What sort of application software might get deployed to such mini-clouds?

The industry expectation today appears to be threefold:

1. First, it might make economic sense to deploy certain *telecoms network functions* to such distributed edge computing locations. The reason is that 100% centralisation of network functions may not be practical in real life.
   Let's stop for a moment here and think about what's happening these days. The telecoms carriers are just in the process of saying good-bye to an era which was characterised by lots of bespoke telecoms hardware/software boxes (switches, routers, service nodes, gateways etc.) that were widely spread across a country in lots of different locations. In the new world, networks are programmable and network software functions get decoupled from the underlying commodity compute, storage and networking hardware. Though super-fast transport networks must still criss-cross the geography of a whole country to reach the thousands of radio base stations for macro and small cells, and fixed access locations, the telecoms software that controls those physical networks and supports services like voice/video and data communication can be centralised and deployed in cloud computing data centers in a way similar to enterprise software.
   At the extreme end, one may concentrate all such software into a single telecoms data center. That's indeed extreme and unlikely to happen for bigger countries. More likely is that the

concentration happens into a few major data centers with a number of much smaller satellite locations. The latter ones are more spread out, geographically distributed. They may host *network functions* (in particular functions standardised by 3GPP) where it's better to get the job done in a distributed way compared to backhauling every bit of traffic into a main data center. This relieves the backhaul networks to the main data centers. Note that such distributed telecoms network functions will be operated by the carriers (assuming it proves to be economic).

2. Second, the edge locations may host another type of software that is also controlled and operated by the telecoms carriers but not specifically standardised across the industry regarding its main functionality. This is software that improves the workings of the 'smart data pipe'. An example is a video optimisation engine or a piece of software that extracts information about the current available bandwidth for a user and signals this information to a content server (like YouTube). This type of edge software enables better service delivery e.g. to end consumers (a better quality of experience for watching YouTube videos when on the move). Various application scenarios are described in the ETSI ISG MEC white paper [3].

3. Third, the IaaS that carriers can make available in a centralised way or in geographically distributed locations, can host 3rd party application code (3rd party as seen from the carrier's perspective). The incentives for enterprises to deploy application code to the edge of telecoms networks are as follows:
   (i) offloading computation from devices to the network to decrease device weight, increase battery lifetime, and possibly benefit from more computing power and storage.
   (ii) deploying application code to the proximity of devices in order to achieve very low latency in the communication between devices and such application code.
   Third parties are e.g. enterprises from various industry sectors (starting from automotive to healthcare to energy etc.). This is now where two rather different worlds meet: the telecoms world with its future ability to offer distributed cloud computing (as IaaS, PaaS) and the M2M/IoT enterprise world.

Point 1) is something the carriers need to consider themselves to judge economic merit. The case for point 2) above is well made by the ETSI ISG Mobile Edge Computing group which is defining an overall architecture for mobile edge computing, some middleware services and APIs to expose such services to application code hosted on a MEC platform. Point 3) for now remains a proposition made by the telecoms industry to the rest of the world, in particular to vertical industry sectors. It requires bringing together the telecoms world with various other industry sectors in order to crystallise the opportunities. This requires sustained effort and the ETSI MEC Industry Enablement Group is part of this ongoing effort [4].

Also note that for points 1) and 2) above, the application software deployed to the edge is not directly revenue generating for network operators, whilst for point 3) it would most likely be, as the edge cloud computing platform is essentially offered to other enterprises as IaaS (or even PaaS and one day SaaS).

Regarding point 3) above, the questions to ask are:

- When the edge computing platforms are there, who are the enterprise customers using the platforms?
- Why would enterprises that adopt M2M and IoT technologies be interested in distributed edge clouds or distributed edge IT?

Let's have a look at what's going on in the world of M2M and IoT.

2. Where is the M2M and IoT business going?

The future is bright for M2M and IoT. There is no lack of highly exciting forecasts. Louis Columbus e.g. provides a good summary of recent forecasts for 2016 and beyond in [5]. The following catches our attention:

Growth is specifically forecast for the industrial Internet of Things and IoT in areas like manufacturing, smart grid, connected vehicles, connected cities, connected homes and connected wearable systems.

So this gives us an idea of the areas where new innovation will call for new solutions. All of above areas will be flooded with emerging new use cases.

New use cases are one thing. Monetization and commercialization is another thing. And then of course: *What types of technology are needed* to enable those uses cases in the years to come? That's the crystal ball reading exercise. Let's look at some further fresh forecasts from Jan 2016 for the near future of the Internet of Things (so we are not even crystal-ball-gazing into 2020):

1) Intelligent connectivity, remote management, control and coordination

This is our own bold forecast. Things of the IoT need to be connected over the most suitable connectivity mechanisms to the back office, but that's not where the evolution ends. Those *things* need to be managed, configured, provisioned, sometimes controlled in real-time, even co-ordinated as groups (e.g. wind farm turbines, or fleets of drones or farming vehicles). In addition, different things of the Internet of Things need to be able to talk to each other (e.g. intelligent traffic light to an autonomous car, or car A to cars B and C for coordinated overtaking and lane changing in a world of autonomous driving).

2) Intelligent linkage of IoT devices with cloud platforms

Andrew Morawski forecasts in [6], that business will *link "M2M with cloud platforms* and big data analytics". Analytics "has now gone from a 'nice to have' to an essential part of a holistic M2M solution". M2M clearly moves beyond connected objects (connected fridges, connected smart gas and electricity meters etc.). Solutions become more holistic and that implies making full use of the power of cloud computing as well (whether it is for real-time analytics, use of machine learning or other reasons).

This also fits to what we heard on 27 Jan 2016 from Wind River [7]: "Wind River®, a global leader in delivering software for the Internet of Things (IoT), announced it is collaborating with IBM to advance IoT deployments for industrial customers with new *"edge-to-cloud" recipes* designed to simplify and accelerate the development of smart connected devices". And the announcement adds:

"The new recipes guide customers on how to integrate services from IBM Watson IoT Cloud Platform with products from the Wind River Helix™ portfolio. Industrial customers using the recipes will now be able to connect industrial devices running Wind River software to the IBM Watson IoT Cloud Platform and access IBM Bluemix cloud services and analytics, allowing IoT developers to more quickly and easily develop smart connected devices."

And as Wind River explains: "Device Cloud securely manages the movement of data from devices, through gateways, and into the enterprise." [8]

Mind that Wind River / IBM are not the only pair driving in this direction. Another one is ARM / Microsoft Azure.

Note 1: IBM Watson IoT is about extending the power of cognitive computing to the things in the IoT [9]. Cognitive computing, as Harriet Green from IBM says in [10] is about "systems that learn at scale [from interaction with humans and their environment], reason with purpose, and interact with humans naturally." For recent news, see also [11].

Note 2: Maybe there is a link between above and also ETSI Mobile Edge Computing, given that Intel is an active member in ETSI ISG MEC and Wind River is a wholly owned subsidiary of Intel Corporation.

3) Real-time analytics for the industry

Further underpinning above, Mary Catherine O'Connor suggests the following in [12]: "In 2015, there was a big shift toward using intelligent devices at the edge of a network in order to enable filtering that reduces the volume of incoming data. This year [i.e. 2016], the focus will shift to *adding analytics at the edge*, which requires highly specialized software. We expect to see more end users in the enterprise space deploying advanced real-time data streaming analytics software platforms." Well, our question is: Where will they deploy such software? Where is *the edge*?

4) Augmented reality and cognitive assistance for industrial use

As Mary Catherine O'Connor says further in [12]: "Manufacturers are seeing benefits from equipping employees with wearable devices that serve as digital assistants to boost the speed and accuracy of everything from parts assembly to maintenance and repair." There is indeed some evidence that the industry is making good progress on this front.

Regarding augmented reality and cognitive assistance, see below videos that give an idea of what's to come:

First, the videos from BMW on AR in [13] and [14], second, a video from Daqri about smart helmets in [15] with some more interesting information in [16] and third, a video from Carnegie Mellon University in [17] about wearable cognitive assistance (supported by edge cloud computing).

Our question is: Where is the software typically running that enables such AR and cognitive assistance? Is there simply sufficient computing power on a below-1-kg AR helmet or on a fashionable smart glass?

5) Trend towards IoT as a Service (IoTaaS)

Brad Shimmin highlights this in [18]. He argues: "While IoT will still struggle at the outer, hardware-centric edge with incompatibility and performance/security concerns, enterprise buyers will see some relief and opportunity as technology providers such as Google, SAP, Microsoft, Oracle, IBM, and Cisco begin hosting IoT platform services[1]. Platforms such as Microsoft's Azure IoT Suite will evolve to help solve some of the serious issues surrounding data capture and integration." In addition, we should add that companies from vertical industries not only connect their "things" to the Internet e.g. for good old-fashioned telematics purposes, but evolve their solutions to full-fledge cloud-based management systems. See the article in [19] titled "From Farming To Big Data: The Amazing Story of John Deere".

Our hypothesis: If the giants of cloud hosting move to bigger and bigger hyper-scale data centres (say with a pair of data centers per country) and the companies from vertical industries either work with those IoT cloud platform services or bet on their own (few) corporate data centers or private clouds, then the space between those few cloud computing centers and the billions of devices in the IoT becomes rather big. Proponents of fog computing would say: there is lots of space for fog, that fills the void between a handful of super-scale cloud data centers and many interconnected, real-time control seeking, sensor data spitting things of the IoT.

Our question: Who fills this void?

[1] We note the acquisition of Jasper Wireless by Cisco as announced in Feb 2016.

6) Edge IT infrastructure for IoT

David Schatsky and Avinav Trigunait forecast in [20] in Jan 2016: "Business interest in the Internet of Things continues to grow, with two trends driving an expansion of IoT applications and possibilities: the emergence of new wireless communications networks and the *arrival of 'edge computing' IT infrastructure*." They summarise the benefit of edge computing IT infrastructure as follows: "The arrival of 'edge computing' IT infrastructure, which facilitates analyzing and acting on IoT sensor data close to the source, making applications more responsive to rapidly changing local conditions while avoiding communications bottlenecks." At least for now, they see edge analytics as a prime user of edge computing IT infrastructure (referring also to Cisco acquiring Parstream in 2015 and HP Enterprise, Dell, Intel having launched edge analytics-focused products including gateways, platforms, and chipsets).

Andy Thurai and Mac Devine from IBM forecast in [21] in Jan 2016: "We started seeing companies moving from the central cloud platforms toward the edge, or toward decentralizing the cloud. This is partly because, with the proliferation of IoTs, operations technologies (OT) and decision intelligence need to be closer to the field than to the central platform." And regarding technological shifts, they say: "Because of the invention of micro services, containers, and APIs, it is easy to run these smaller, self-contained, purpose-driven services that specifically target only certain functions that are needed at the edge. The ability to use containers for mobility and the massive adoption of Linux will enable much thicker, monolithic services previously running centralized to be "re-shaped" into a collection of smaller, purpose-driven micro services. Each of these can be deployed and run on the edge as needed and on-demand." The term they use for computing at the edge is "edge IT and edge processing".

Ok, so there is some mention of edge processing, edge IT, edge-to-cloud, and decentralisation of cloud here and there. One may indeed conclude that some of the fancy computation that is required for above areas may benefit from some IT computing power and storage in close proximity to where it is actually needed (whether coordinating multiple autonomous vehicles on a road, performing fast object recognition and machine learning in real-time for wearable cognitive assistance, or running some real-time analytics on a high data rate video stream).

Let's check with the **Industrial Internet Consortium (IIC)** in the assumption they have a clear vision of the Industrial Internet's future shape. Their future-proven reference architecture [22] splits the end-to-end architecture into three tiers: Edge tier, platform tier, enterprise tier. The edge tier is simply defined as:

"The edge tier collects data from the edge nodes, using the proximity network. The architectural characteristics of this tier, breadth of distribution, location, governance scope and the nature of the proximity network, vary depending on the specific use cases." And here we go with what these tiers are supposed to do: "From the tier and domain perspective, the edge tier implements most of the control domain; the platform tier most of the information and operations domains; the Industrial enterprise tier most of the application and business domains."

The edge tier typically hosts sensors, actuators, controllers and application gateways. Where is the cloud computing at the edge? Well, though not explicitly hardcoding "edge cloud computing" into their reference architecture, the report [22] admits: "The actual functional mapping of IIS tiers [IIS = industrial internet systems] is usually not as simplistic and highly depends on the specific of the system use cases and requirements. For example, some functions of the information domain [includes e.g. data ingestion, data transformation, streaming and batch analytics] may be implemented in or close to the edge tier, along with some application logic and rules to enable intelligent edge computing". So then, here we finally go with a more tangible hint towards the potential use of *edge computing IT*.

3. What requirements will IoT enterprises have on edge cloud computing?

Let's assume edge (cloud) computing will gain a degree of relevance in the next 5 years. What might be the requirements of enterprises on any such cloud computing?

A German report [23] provides a few hints. The following percentage of interviewed cloud computing customers in Germany considers the listed requirements as Must-have:

- Support of hybrid cloud concepts: 35% of interviewed enterprises.
- Flexible adaptability of solutions: 55%.
- Ability to integrate solutions: 74%.
- Data hosted exclusively in-country: 83%.

In addition we believe the following points will be part of the requirements set:

- Security, in particular as an edge cloud platform operated by telecoms service provider can be considered a variant of a public cloud, seen from an enterprise that considers to deploy an application to the edge cloud.
- Integration with DevOps environments of enterprises.
- Single point of control for operation and management of the application during runtime incl. visualisation, reporting etc.
- Single point of control for multi-cloud application deployment (across e.g. private cloud and an operator's edge cloud) and for life cycle management.
- Ability to integrate own data center and private cloud with different edge cloud platforms from different service providers.
- Ability to easily use APIs as exposed by the edge cloud platforms.

Above cross-cloud integration aspects are particularly important as enterprises and application developers opting into hybrid clouds seek integrated service models and move more and more towards agile development, continuous integration and continuous delivery.

As a result of above, we believe, that edge (cloud) computing platforms (MEC servers, cloudlets, and some fog nodes) will need to be well integrated with the data centers and private clouds of IoT enterprises and possibly also with main public clouds as from Amazon, Google, Microsoft, HP, IBM, Salesforce.com and others. What's needed is *edge cloud to cloud integration*.

4. Consequences for the telecoms industry and conclusions

From above discussion we observe the following:

1. The telecoms industry prepares for edge computing IT platforms (e.g. in the form of ETSI's mobile edge computing, or in the form of fog computing).
2. The M2M/IoT industry (here meant the enterprises adopting M2M and IoT) is keen to link IoT-specific functions at the edge to the cloud, in order to leverage cloud computing. It also works on adding more powerful software to the edge (like real-time analytics).

Industries do not always perfectly work together, nor have the members of one industry perfect insight into what other industries are doing (they may not even seek it).

The current situation in early 2016 bears the following risks:

- Risk that the telecoms and Internet industries define system architectures, interfaces and APIs for edge computing and fog computing without sufficient awareness of what the ultimate IoT enterprise customers (who are active in a variety of industries from health to automotive) actually need. In this regard, efforts of ETSI ISG MEC's Industry Enablement Group to reach out to companies in many other 'vertical' industries are very important.
- Risk that telecoms solutions for edge computing turn out to be weak with regards to edge cloud to private/public cloud integration, though they might be strong with regards to edge cloud to network function virtualisation integration. The latter focuses on leveraging synergy

with NFV and SDN technologies that get adopted in telecoms core and access networks. Those synergies are expected to turn into economic benefits including cost savings for the network operators. However, these synergies are not necessarily translating into economic benefits for 3rd party tenants of an edge cloud.
- Risk that the network operators' business case for investment in edge (cloud) computing infrastructure is weakened should above 2 risks sufficiently materialise. As outlined in section 1, we only expect 3rd party enterprise applications deployed to edge (cloud) computing platforms to become *directly* revenue generating (e.g. in B2B, B2B2B or B2B2C business models).

Our conclusion is that apart from seeking harmonisation with NFV/SDN concepts (e.g. for the Mobile Edge Computing architecture in ETSI ISG MEC) for the justified sake of leveraging synergies with the ongoing network transformation of telecoms operators, it is important to consider edge cloud to cloud integration as long as M2M/IoT enterprises are expected to be 3rd party application developers and tenants in edge computing infrastructure run by network operators.

Today, we see different candidate solutions that should be further considered and investigated:

- Integration using APIs exposed by edge (cloud) computing platforms like MEC servers, cloudlets and fog nodes to management and orchestration software used in private data centers, private clouds and other main clouds.
- Deployment of common PaaS layers across edge (cloud) platform and private data centers, private clouds and public clouds. This is important when the technology at the edge differs substantially from the technology used in the non-edge environments, e.g. with regards to hypervisors and cloud management systems and when the goal includes to simplify the burden for 3rd party enterprises. Related key words are multi-cloud deployment and cloud brokerage. An example open source candidate for such cross-cloud integration could be CloudFoundry.

Above all, gaining the perspective of the users of edge (cloud) computing infrastructure, cloudlets, MEC servers and fog nodes is highly important. Some of the candidate users may well be found in industry sectors as different as agriculture (e.g. precision farming), healthcare (e.g. robot-assisted), logistics, smart cities or automotive.

Disclaimer: Views stated, opinions expressed or analyses provided are solely those of the individual authors as private persons and shall not be considered or interpreted as views, opinions, or analyses of any of the companies they may be liaised with, employed by or do work for.

References

[1] Klas, Guenter. "About Mobile Edge Cloud Computing – An Introduction". 8 Nov 2015. http://yucianga.info/wp-content/uploads/2015/11/15-11-08-About-Mobile-Edge-Cloud-Computing-%E2%80%93-An-Introduction-v1.pdf

[2] Klas, Guenter. "Fog Computing and Mobile Edge Cloud Gain Momentum - Open Fog Consortium, ETSI MEC and Cloudlets", 22 Nov 2015. http://yucianga.info/wp-content/uploads/2015/11/15_11_22-_Fog_computing_and_mobile_edge_cloud_gain_momentum_Open_Fog_Consortium-ETSI_MEC-Cloudlets_v1_1.pdf

[3] ETSI ISG MEC White paper. "Mobile-Edge Computing". https://portal.etsi.org/Portals/0/TBpages/MEC/Docs/Mobile-edge_Computing_-_Introductory_Technical_White_Paper_V1%2018-09-14.pdf

[4] ETSI MEC Industry Enablement Group ToR.
https://portal.etsi.org/Portals/0/TBpages/MEC/Docs/ToR_MEC_IEG.pdf

[5] Columbus, Louis. “Roundup Of Internet of Things Forecasts And Market Estimates, 2015”. January 3, 2016
http://softwarestrategiesblog.com/2016/01/03/roundup-of-internet-of-things-forecasts-and-market-estimates-2015/

[6] Morawski, Andrew. “2016 Predictions: M2M and IoT to transform operations”.
http://www.rcrwireless.com/20160105/internet-of-things/2016-predictions-m2m-and-iot-to-transform-operations-tag2

[7] Wind River news room, Wind River Teams with IBM to Accelerate Edge-to-Cloud IoT Deployments for Industrial Customers. 27 Jan 2016.
http://www.windriver.com/news/press/pr.html?ID=13972

[8] Wind River Helix Device Cloud.
http://www.windriver.com/products/helix/device-cloud/

[9] IBM Watson IOT.
http://www.ibm.com/internet-of-things/

[10] Green, Harriet, IBM. “The Internet of Things in the Cognitive Era: Realizing the future and full potential of connected devices”, Dec 2015.
http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?subtype=WH&infotype=SA&htmlfid=WWW12366USEN&attachment=WWW12366USEN.PDF

[11]M2M Magazine, IBM Opens Watson IoT Global Headquarters in EU. Dec 17, 2015.
http://www.machinetomachinemagazine.com/2015/12/17/ibm-opens-watson-iot-global-headquarters-in-eu/

[12] O'Connor, Mary Catherine. “The IoT Forecast: Five Predictions for 2016”, 4 Jan 2016.
http://www.iotjournal.com/articles/view?13901

[13] BMW smart glasses. 27 Jan 2014.
https://www.youtube.com/watch?v=uoZSZe4rTyI

[14] BMW: BMW's Mini augmented reality glasses are the future of car tech.
https://www.youtube.com/watch?v=-m7B-91KBXg

[15] Daqri smart helmet.
https://www.youtube.com/watch?v=47H6ul0W1-E

[16] Ravenscroft, Tom. “Soluis and Crossrail to test smart helmets on site in Innovate UK trial”, CIOB, BIM+, 3 Dec 2015.
http://www.bimplus.co.uk/news/soluis-a3nd-cros6srail-test-smart-helm2ets-site/

[17] Carnegie Mellon University: 2D Lego Assistant.
https://www.youtube.com/watch?v=uy17Hz5xvmY&feature=youtu.be

[18] Shimmin, Brad. “Enterprise Data and Analytics Market Outlook for 2016”. Dec 7, 2015.
http://itcblogs.currentanalysis.com/2015/12/07/enterprise-data-and-analytics-market-outlook-for-2016/

[19] Marr, Bernard. “John Deere: From Farming To Big Data: The Amazing Story of John Deere”, 7 May 2015.

http://www.datasciencecentral.com/profiles/blogs/from-farming-to-big-data-the-amazing-story-of-john-deere

[20] Schatsky, David & Trigunait, Avinav. “Internet of Things, Dedicated networks and edge analytics will broaden adoption”. Jan 21, 2016.
http://dupress.com/articles/internet-of-things-iot-adoption-edge-analytics-wireless-communication-networks/

[21] Thurai, Andy, Devine, Mac. “Cloud computing moves to the edge in 2016”, IBM, 15 Jan 2016.
http://www.thoughtsoncloud.com/2016/01/cloud-computing-move-to-the-edge-in-2016/

[22] Industrial Internet Consortium.” Industrial Internet Reference Architecture”.
http://www.iiconsortium.org/IIRA-1-7-ajs.pdf

[23] Germany Trade and Invest. Opportunities in the German software market, German SMEs: Driving Digitization Forward 2016.
https://www.google.co.uk/url?sa=t&rct=j&q=&esrc=s&source=web&cd=1&cad=rja&uact=8&ved=0ahUKEwiAqfrEl9_KAhWFDA8KHQCfAd4QFggfMAA&url=http%3A%2F%2Fwww.gtai.de%2FGTAI%2FContent%2FEN%2FInvest%2F_SharedDocs%2FDownloads%2FGTAI%2FFact-sheets%2FBusiness-services-ict%2Ffact-sheet-software-industry-en.pdf%3Fv%3D2&usg=AFQjCNHNwEwlcILSJvWHlfsdpqz44XG-RQ&sig2=i7XI-mJ4ONhRYxJYyySOnA&bvm=bv.113370389,d.ZWU